

Ecological Knowledge Formation Through Ocean Robotics

Wanhong HUANG
huangwanhong@serendip.ngo

Abstract

Robotic encounters with marine environments can expand scientific observation while changing the conditions under which ecological phenomena are observed. Their epistemic value depends on instrument design, sampling histories, ecological response, and institutional authority. This exploratory paper investigates ecological knowledge formation through oceanic robotics using a provisional generative-relational (GR) vocabulary. It distinguishes questions concerning ecological reality, causally consequential relations, and situated manifestations without asserting a comprehensive ontology. The inquiry develops connected accounts of ecological generativity and material finitude, nonhuman interests, procedural justice, marine nomos, and the circulation of knowledge and heterogeneous value. Eleven topic-specific engineering reviews examine robotic platforms, biologically informed morphology, hydrodynamics, control, perception, navigation, communication, manipulation, materials, reliability, heterogeneous teams, and adaptive scientific observation. Conditional models and comparative designs examine uncertainty, ecological exposure, sampling bias, bounded robotic capability, resource constraints, and the limits of inference between contexts. The paper differentiates instrument measurements from warranted ecological interpretations; observed behavioral responses from welfare or population effects; and technical feasibility from ecological acceptability and lawful authorization. Habitat recovery, sediment transport, decommissioning, and lifecycle burdens are treated as context-dependent questions rather than consequences that can be inferred from engineering performance. Proposed tests compare GR-informed provenance and decision records with competent conventional workflows, including FAIR-, CARE-, and PROV-informed approaches, and specify results that would weaken the proposed framework. A staged research agenda addresses relational ecological epistemology, embodied robotics, organism and habitat responses, institutional practice, and integrated validation. The manuscript offers an interdisciplinary synthesis and a criticizable research programme. It reports no original marine deployment, animal experiment, completed ecological impact assessment, or site-specific legal determination.

Keywords: generative-relational inquiry; marine ecology; oceanic robotics; ecological generativity; animal welfare; procedural justice; marine nomos; adaptive control; knowledge circulation.

Discussion-paper note. This v3.6 exploratory discussion draft contains all 26 planned sections. It has undergone selected internal conceptual, mathematical, engineering, ecological, habitat, maritime-law, and cross-manuscript reviews. These do not constitute peer review, exhaustive source verification, independent proof certification, empirical validation, or jurisdiction-specific legal analysis. No original marine deployment, animal experiment, ecological impact assessment, or operational authorization is reported. Three historical preliminary sections are archived separately and are not part of the compiled manuscript.

Licence and preparation. Original material © 2026 Wanhong Huang, CC BY-NC 4.0 (<https://creativecommons.org/licenses/by-nc/4.0/>); third-party rights remain unaffected. ChatGPT assisted with research dialogue, source discovery, drafting, editing, and LaTeX; the human author assumes responsibility for claims and subsequent review.

1 Introduction

This inquiry examines how encounters with marine ecological systems can enlarge the capacities of scientific understanding and technological creation, and under what conditions those encounters can be conducted responsibly. Its subject is the formation of ecological knowledge through oceanic robotics:

the relations among living organisms, physical environments, scientific instruments, embodied actions, institutional arrangements, and the representations through which observations become intelligible. The ocean affords striking opportunities for inquiry because its organisms and processes inhabit environments that can be difficult to access, and because their dynamics may resist the categories through which human observers initially approach them. Yet the possibility of discovery must be considered together with the conditions that make discovery possible. A robotic instrument may disclose previously uncharacterized relations while disturbing the organisms involved; a scientific survey may generate valuable observations while consuming energy, occupying habitat, or changing local acoustic conditions; a new technological infrastructure may extend research capacities while redistributing ecological risks and economic benefits. These possibilities motivate an account of marine exploration in which scientific promise, ecological responsibility, and the material and institutional conditions of inquiry are investigated together.

A generative-relational perspective provides the initial philosophical orientation. The expression *generative-relational* identifies an investigative concern with how entities, distinctions, practices, capacities, and values are constituted, sustained, and transformed through relations. It does not imply that scientific observers have complete access to the processes constituting another organism's world, or that all phenomena can be explained by a single relational formalism. A provisional distinction among *ontological*, *relational*, and *manifestational* registers is useful here. The ontological register concerns questions of what exists and what forms of organization are possible, including questions that remain independent of a particular observing practice. The relational register concerns the interactions, dependencies, constraints, and historically constituted relations within which entities and capacities acquire determinate roles. The manifestational register concerns the appearances, measurements, classifications, and practical articulations through which certain aspects of a system become accessible to an observer. These registers describe different tasks of analysis; they are neither mutually exclusive compartments nor three universally established strata of reality. Their relationships require explicit assumptions in any proposed model.

Marine ecology makes these distinctions consequential. An organism can exist and act independently of a particular camera or taxonomic database. Yet its scientifically documented identity, observed behavior, and inferred ecological role depend on the instruments, sampling schemes, environmental contexts, research conventions, and inferential practices through which it is studied. A remotely operated camera may reveal an encounter between two animals without establishing whether their relation is predation, defense, competition, curiosity, or an interaction inadequately represented by an existing behavioral category. A genetic signal may support the presence of biological material without locating a viable organism at the sampling point or establishing its ecological role. The observations may be accurate within their stated measurement conditions while permitting several incompatible explanatory reconstructions. Scientific inquiry must therefore examine both the evidence and the generative operations through which a claim is made, criticized, reproduced, and transformed. The relevant object is richer than the final dataset, though the dataset remains indispensable.

The term *ecological generativity* is introduced as a provisional analytical description of the capacities through which ecological systems sustain, alter, and sometimes lose possibilities for further organization and interaction. Biological reproduction, developmental variation, symbiosis, niche construction, habitat transformation, and evolutionary change exemplify distinct mechanisms that may contribute to such capacities. Their relationships are diverse, and no claim is made that these mechanisms share a single universal measure or direction. An ecosystem may reorganize after disturbance, maintain certain relations under changing conditions, or cross thresholds at which familiar forms of persistence become difficult or impossible. Some transformations preserve important possibilities; others foreclose them. Accordingly, the existence of generative capacity should never be inferred to entail unlimited abundance, inevitable recovery, or a trajectory of continuous improvement. In this inquiry, accounts of change must specify the represented histories, partial orderings, dynamical variables, and observational supports on which particular claims depend. A privileged global "present" or universal arrow of development is not assumed merely because an operational system uses a clock.

Three research questions organize the inquiry. First, which distinctions between ecological processes,

measurement traces, and inferential claims are needed when an observing robot both samples and alters its environment? Second, how can the engineering conditions enabling observation be specified sufficiently well to test a scientific proposition without treating the robot's operational success as evidence for a biological explanation? Third, how should legal competence, organism-related concerns, and the ability to challenge evidence enter decisions about the continuation, modification, or withdrawal of a mission? These questions require different methods. Conceptual analysis can expose category errors; engineering and ecological experiments can test bounded hypotheses; institutional comparison can examine the practical availability of scrutiny and remedy. The paper develops their interfaces without claiming that one method verifies all three.

This perspective creates a difficult ethical question. An increase in *human* knowledge-producing capacity can coincide with a reduction in the ecological conditions supporting nonhuman life. A vehicle may acquire an unprecedented record of a fragile reproductive site while its light, motion, contact, or acoustic emissions alter animal behavior. A biological sample may make a novel molecular pathway available to scientific study while extraction changes a small or poorly understood population. Greater predictive accuracy may facilitate both protection and more efficient exploitation. The relevant consequences depend on the specific organisms, habitats, exposure pathways, scales of activity, and institutional decisions. An ethical analysis must therefore examine the effects of inquiry on the systems under study, rather than treating the acquisition of additional information as a self-justifying benefit. This claim calls for evidence and context-specific deliberation; it does not establish that every observation causes measurable harm.

The interests of individual animals require distinct consideration. Population-level conservation, habitat protection, and the protection of individual welfare are related but nonidentical objectives. An intervention potentially beneficial to a population may injure particular animals. A survey that leaves population abundance apparently unchanged may still cause distress, disorientation, behavioral interruption, or an alteration of relationships important to the individuals concerned. Knowledge about animal perception, sentience, physiological vulnerability, and behavioral significance varies across species. The limits of this knowledge create difficulties for both scientific interpretation and moral representation. Researchers may observe responses to a stimulus while remaining uncertain about the animal's experiential or motivational state. Such uncertainty warrants careful qualification, appropriate precaution, and the development of research methods that can detect relevant adverse responses. It cannot, on its own, determine that an organism lacks ethically significant interests. For example, the scientific and management literature on anthropogenic ocean noise documents effects that can include interference with acoustic communication, displacement, and disruption of feeding or breeding under particular exposure conditions [130]. The operational consequences for a specific vehicle or species remain questions for properly scoped assessment.

Proceduralism provides one route for organizing inquiry under such uncertainty, provided that procedural legitimacy is distinguished from substantive justice. Procedures can specify how hypotheses are assessed, how ecological risks are identified, how affected human communities participate, who may authorize an activity, what evidence must be recorded, and under which conditions an action may be suspended or challenged. Their institutional design matters because different participants may possess unequal scientific resources, interpretive authority, economic power, and practical ability to contest a decision. Nonhuman organisms ordinarily cannot articulate claims through human legal and administrative procedures; their interests are interpreted through fallible scientific and representative practices. Procedural organization should therefore make the conditions and limitations of representation visible. A formally compliant decision remains open to evaluation in light of its actual effects on organisms, habitats, and affected people. This paper investigates procedural accountability together with distributive, recognitional, ecological, and remedial aspects of justice, rather than assuming that a completed review is sufficient evidence of a just result.

The ethical question is inseparable from material finitude. Marine environments are dynamic and sometimes renewable, but the rates and conditions of biological regeneration, habitat recovery, material supply, and environmental assimilation are limited and heterogeneous. A long-lived geological structure,

a vulnerable breeding habitat, a localized population, and an offshore energy resource each require a different account of available stock, flow, disturbance, and potential renewal. Scientific and technological infrastructures themselves depend on finite energy, materials, maintenance capacity, vessels, manufacturing supply chains, and waste-management arrangements. Subsea installations may change local heat, noise, substrate conditions, or access; their ecological consequences require empirical study and comparison with nonmarine alternatives. The paper will consequently distinguish the creation of technological opportunities from the sustainability of the physical and biological arrangements through which those opportunities are realized. Possible benefits from oceanic computation, energy, construction, or biological research should be investigated alongside cumulative impacts and the feasibility of recovery, repair, removal, and restoration.

A further question concerns the circulation of what scientific encounters produce. Records, explanations, experimental methods, material samples, digital models, commercially useful discoveries, and the capacity to conduct further inquiry circulate through different channels and under different conditions. A research group may make a dataset publicly accessible while retaining exclusive access to the instruments, compute resources, trained personnel, or proprietary methods needed to reproduce its analysis. A biological discovery may create financial returns for a distant company while ecological costs remain concentrated near the site of extraction or installation. Distinct communities and institutions may also interpret the same observed phenomenon through heterogeneous scientific, cultural, legal, and economic practices. In such circumstances, circulation requires more than copying files or translating vocabulary: the receiving context may reconstruct an observation's meaning, create new uses for it, reject a proposed valuation, or require additional evidence and consent. The infrastructural requirements include traceable provenance, partial and revisable translation, clear authority boundaries, accessible means of objection, and opportunities to develop independent knowledge-producing capacities.

In considering these economic and institutional relations, the inquiry takes intellectual inspiration from Eglash's treatment of *generative justice*, particularly his attention to the generation and circulation of ecological, labor, and expressive value [35]. The generative-relational framework, marine-specific legal analysis, and proposed engineering requirements developed here remain separate investigations. The purpose of the reference is to acknowledge an important source for approaching economic activity through value creation and circulation, while examining additional questions about heterogeneous relational histories, institutional authority, and ecological conditions. A satisfactory analysis cannot presume that diverse values are completely commensurable or that financial transfers reproduce biological, cultural, or epistemic relations affected by an activity. Nor can a particular ownership form, public or private, guarantee by itself that benefits, responsibilities, and productive capacities are distributed justly. These are matters for explicit comparative inquiry.

The practice and regulation of marine exploration take place within a plurality of normative orders. Scientific conventions, professional responsibilities, environmental assessment procedures, community practices, domestic legislation, treaties, and international institutions can govern different aspects of one operation. Here *nomos* denotes the broader formation, interpretation, and transformation of such normative ordering, with positive law retaining its specific conditions of legal validity. Maritime space is not a uniform jurisdiction. Under Part XIII of the United Nations Convention on the Law of the Sea (UNCLOS), coastal states exercise defined powers concerning marine scientific research in the territorial sea, exclusive economic zone, and continental shelf; conditions differ across these maritime zones [196]. The Agreement on Marine Biological Diversity of Areas beyond National Jurisdiction, which entered into force on 17 January 2026, adds treaty obligations for its parties within its defined scope concerning marine genetic resources, conservation measures, environmental impact assessment, and capacity building [188]. These legal frameworks supply indispensable starting points, while the classification and authorization of a particular robotic mission, biological collection, or commercial installation must be determined under its actual jurisdictional and factual circumstances.

Consider a hypothetical installation constructed by participants from Countries A, B, and C that provides computational services to users in many other locations. Its technical components might include vessels, robotic installation equipment, energy links, submerged modules, communication cables, soft-

ware, and maintenance systems. It might generate research capabilities, wages, commercial revenues, knowledge, and new technological dependencies, while producing ecological effects in a particular marine environment. The authority to construct it depends partly on its precise location, purpose, and legal classification. Its economic beneficiaries need not coincide with the institutions that grant permissions or the communities potentially exposed to environmental consequences. A public benefit can arise through service provision or shared scientific capacity, but such an outcome cannot be assumed from multinational ownership or a geographically distributed customer base. The hypothetical case motivates a later investigation of heterogeneous value formation, transnational cooperation, permissions, environmental duties, and the translation of competing evaluations. The analysis will avoid inventing a general international licensing body for all oceanic production or a universal obligation to distribute commercial revenue worldwide.

Against this philosophical and normative background, present observation technologies provide concrete opportunities and limitations. The 2026 General Bathymetric Chart of the Oceans (GEBCO) reports that 28.7% of the global seafloor has been mapped to modern standards, indicating substantial remaining work in bathymetric measurement [50, 70]. This statistic concerns the seafloor and is not a measure of the proportion of oceanic life, water-column processes, or ecological relationships that has been understood. Similarly, the Global Ocean Observing System’s Essential Ocean Variables organize sustained measurement across physical, biogeochemical, biological, and ecosystem domains without supplying a complete theory of the relations among those variables [105, 54]. A pressure record, a current measurement, an image, an acoustic detection, and an environmental-DNA sample have different supports, error mechanisms, and inferential affordances. Ecological understanding depends on the way these records are combined with models, independent checks, and explicit consideration of the circumstances in which they were produced.

Oceanic robotic systems create capabilities for observing and interacting across different environments: the sea surface, water column, seafloor, and, where relevant, sediments and sub-seafloor formations. Autonomous surface vehicles, gliders, floats, freely swimming vehicles, remotely operated systems, benthic crawlers, and fixed observatories differ in their endurance, sensors, communication arrangements, ability to intervene, and degree of direct human supervision. Their bodies also enable different activities. Morphologies inspired by fish or rays may be investigated for efficient or maneuverable movement; compliant arms inspired by cephalopods may be investigated for sensitive interaction; anchored rigid systems may be preferable for tasks demanding large reaction forces, precise connections, or structural assembly. Biological analogies are hypotheses about useful mechanisms rather than evidence that an engineered machine reproduces the organism’s full capacities. Softness is not by itself a welfare guarantee, and effective swimming does not imply effective manipulation or marine construction. Heterogeneous fleets, including combinations of swimming observers, manipulators, fixed facilities, and surface support vessels, therefore deserve systematic comparative study.

The control problem follows from these differences in embodiment and environment. Aircraft, terrestrial humanoids, and marine robots all admit feedback, state estimation, planning, adaptive control, and learning-based methods at a general methodological level. Their mechanical plants, disturbances, contact relations, and observability conditions differ substantially. Oceanic vehicles experience hydrodynamic forces, buoyancy, added mass, stratification, currents, waves, and spatially varying environmental conditions. Marine dynamics also depend on astronomical tidal forcing, wind and pressure fields, biogeochemical transport, biological activity, and the actions of the observation platform itself. The ocean cannot be presumed globally reconstructible from the finite sensors of an individual mission. A useful engineering problem is therefore to act within a locally warranted account of the dynamics, revise this account when observations demand it, and recognize the boundaries of its predictive and control authority. The proposed use of *tangent-space governance* refers provisionally to locally supported sensitivity and controllability assessments; *criticality governance* refers to procedures for monitoring model breakdown, operational thresholds, and possible dynamical transitions. These terms name an exploratory synthesis requiring definitions and comparative tests, rather than a claim that new control laws have already been established.

Adaptive control and learning may improve a robot's responsiveness to changing conditions, while biologically informed control may reveal useful coordination principles across compliant bodies and distributed actuation. Neither approach independently establishes stability, reliable generalization, ecological admissibility, or legal authorization. A robot's choice of action must be evaluated at multiple levels: the physical feasibility of a maneuver, the evidential justification for a scientific measurement, and the normative conditions under which the measurement or intervention is permitted. A feasible observation can be scientifically uninformative; an informative experiment can carry unacceptable ecological risk; and a successful simulation can fail to predict a particular field environment. These distinctions motivate an architecture in which locomotion and manipulation controllers, environmental estimation, mission-level experimental planning, and independent safety and authorization review exchange information while retaining distinct functions. Potential responses to inadequate evidence include revised sampling, reduced intervention, human review, withdrawal, and termination of an operation.

The present investigation consequently distinguishes *data acquisition*, *ecological inference*, and *generative scientific inquiry*. The first concerns the production of observations under identifiable measurement procedures; the second concerns warranted claims about organisms, environments, and relations; the third concerns changes in the questions, methods, representations, and capabilities through which further inquiry becomes possible. A surprising measurement can reveal instrument error, an unmodeled environmental condition, or inadequacy in a current hypothesis family. Determining among these possibilities requires calibration, comparison, contextual evidence, and appropriately designed follow-up investigation. An AI system may assist with anomaly detection, proposed experiments, or model revision, but its suggested explanation does not establish the existence of an unfamiliar ecological mechanism. Scientific generativity is likewise not equivalent to an optimizer obtaining a larger numerical reward. The paper treats candidate new representations as proposals that must survive relevant scientific and institutional scrutiny.

The paper's principal research question is therefore: *under what epistemic, ecological, normative, and engineering conditions can oceanic robotic systems contribute to the formation and circulation of knowledge about heterogeneous marine relations while sustaining appropriate responsibility toward the organisms and environments they encounter?* This question contains several subsidiary problems. The philosophical task is to clarify the limits and uses of generative-relational accounts of ecological reality and inquiry. The ecological and ethical task is to investigate how scientific actions affect nonhuman life, relationships, and materially limited generative capacities. The institutional task is to identify the normative orders, authorities, and contestation procedures relevant to international marine research and production. The engineering task is to compare embodied forms, sensing methods, dynamics models, control strategies, energy systems, and cooperative infrastructures under explicit performance and ecological conditions. The methodological task is to design evaluations capable of challenging, rather than merely illustrating, the proposed architecture.

The intended contributions are deliberately bounded. First, the study develops a vocabulary for distinguishing ecological reality, observations, relationally situated interpretation, and the generative capacities implicated in marine inquiry. Second, it investigates how animal welfare, material finitude, proceduralism, and substantive justice constrain the formation of scientific and economic opportunities. Third, it examines marine nomos, legally relevant authority, common knowledge creation, generalized translation, and heterogeneous value circulation as practical features of cross-border ecological research. Fourth, it establishes a topic-based engineering survey spanning morphology, fluid-structure mechanics, locally adaptive and biologically informed control, sensing, navigation, manipulation, materials, power, and collective robotic operation. Fifth, it proposes falsifiable comparative research designs for studying physical performance, epistemic value, ecological disturbance, and institutional accountability. These are prospective contributions to an interdisciplinary research programme, not claims of completed field validation or general results about all marine species or technologies.

The argument proceeds from philosophy toward practice and engineering while recognizing their reciprocal dependencies. Sections 2–6 examine relational reality, ecological generativity, nonhuman life, proceduralism and justice, and cross-system knowledge formation. Sections 7–10 develop the analy-

sis of marine nomos, legal and institutional authority, common knowledge circulation, and accountable procedures. Sections 11–21 survey contemporary oceanic robotic platforms and analyze morphology, hydrodynamics, control, sensing, navigation, manipulation, actuation, material conditions, robot teams, and adaptive inquiry. Sections 22–24 address animal welfare within engineering practice and possible future marine technologies under finite ecological conditions. Sections 25–26 specify comparative evaluation, limitations, conclusions, and research continuations. This sequence is an argumentative organization, not a thesis that reality or institutional history necessarily unfolds through these stages. Detailed mathematical representations will be introduced where they support explicit scientific questions, with models of relational histories and evolving possibilities distinguished from any unexamined universal time coordinate.

This manuscript is an exploratory independent inquiry, complementary to a broader study of global observation and knowledge infrastructures. Its specific focus is the material and ecological formation of knowledge through oceanic robots and the relations enabling or constraining their operation. Some subjects, especially transnational production and the political economy of autonomous AI, invite extensive treatment in companion investigations; they are considered here to the extent that they alter marine inquiry and its responsibilities. The present manuscript develops a complete provisional argument and comparative research programme across its twenty-six sections. It reports no new marine field observations, animal experiments, robotic deployment, or installed oceanic computing facility. The viability of its concepts depends on future philosophical criticism, specialist legal and ecological review, independently reproduced engineering comparisons, and evidence gathered under appropriate authorization. The investigation begins from the possibility that encounters with unfamiliar ecological relations may change what can be known and created, while leaving open how those possibilities ought to be pursued, limited, shared, and sustained.

2 Relational Reality in Marine Ecologies

2.1 The object of a relational ecological inquiry

This section establishes the philosophical and representational premises required to investigate marine ecological systems through robotic observation. Its immediate task is to clarify what an observer may mean by an organism, an ecological relation, an environment, and an observation when the objects of inquiry exhibit interacting physical, biological, and historically contingent forms of organization. The discussion develops a generative-relational (GR) research vocabulary while distinguishing proposed analytical structures from established empirical findings. A relational formulation can direct attention toward the conditions of a phenomenon’s formation and intelligibility; it does not, by that fact alone, identify its causal mechanisms or establish the explanatory superiority of a particular model.

Conventional ecological research already examines interactions among organisms and their environments. The present inquiry therefore does not claim to introduce relationality into ecology. The processual philosophy of biology has emphasized metabolic continuity, life cycles, and ecological interdependence as reasons to investigate organisms through processes of maintenance and transformation [34]. Niche-construction research investigates how organisms alter environmental conditions that can subsequently influence other organisms and evolutionary processes [149]. Multiscale ecology, moreover, has repeatedly shown that observed patterns depend on the spatial and organizational scale at which they are described [92]. These established approaches provide interlocutors for the present analysis. The additional question posed here concerns how heterogeneous observing and acting systems generate distinct but potentially communicable representations of those processes, and how their acts of inquiry enter the relational configuration under investigation.

Consider a reef-associated organism. A photograph can record a visible body at one location, a genetic sample can support an inference about biological material, and a dissolved-oxygen probe can characterize one part of the organism’s surrounding chemical environment. The observations may be individually valid while failing to characterize feeding, reproduction, microbial association, migration,

Table 1: Analytical registers and the limits of their interpretation in a marine ecological investigation.

Register	Principal question	Marine example and inferential limitation
Ontological	What exists or can be organized independently of this particular description?	An organism and its physiological processes may persist without a camera record; their existence is not demonstrated by conceptual analysis alone.
Relational	Which interactions, dependencies, or organizing constraints matter to the research question?	Coral, microbial associates, light, nutrients, and thermal exposure participate in relations whose strengths and mechanisms require investigation.
Manifestational	Through which practices do aspects of a configuration become observable, reportable, or actionable?	Imaging, chemical sensors, and environmental-DNA assays produce different evidential records with distinct sensitivities and biases.

and environmental dependence. They may refer to the same biological individual through different operational criteria, or to overlapping ecological processes without identifying the same object. A relational account consequently requires clarity about the conditions under which an entity is individuated, a relation is proposed, and different records are said to concern a common phenomenon. These questions remain open to ecological evidence and disciplinary disagreement.

The paper treats *generative-relational system* as a provisional analytical designation for a delimited configuration of relations and capacities considered with respect to a specified inquiry. It does not imply that every system is conscious, that every relation is socially constituted, or that boundaries in nature are arbitrary. Boundaries can be imposed for scientific purposes, constrained by reproducible physical discontinuities, or associated with biological organization; multiple defensible boundaries can coexist. A marine ecosystem is neither automatically one unified subject nor merely a passive background for humans and machines. Species, individuals, communities, microbial assemblages, oceanographic processes, and technical instruments have different forms of organization and should not be assigned identical kinds of agency by terminological stipulation.

2.2 Ontological, relational, and manifestational registers

The preliminary distinction among *ontological*, *relational*, and *manifestational* registers identifies three questions that frequently become conflated in accounts of ecological observation. The ontological register concerns what entities, processes, and forms of organization exist, including possibilities not exhausted by particular observations or descriptions. The relational register concerns determinate dependencies, interactions, constraints, and modes of organization through which the identified entities persist or change. The manifestational register concerns what becomes available in observations, records, classifications, interventions, and other articulations within an identifiable investigative context. These are distinguishable analytic responsibilities, not empirically established and separately measurable layers of reality.

The ontological register permits realism about an independent marine world without presuming complete epistemic access to its organization. A scientific description may be tested against recalcitrant observations, revised through instrument calibration, or rejected because an alternative model explains more of the evidence. Realism in this restricted sense need not require that an observer's categories coincide perfectly with the individuation of marine entities. Conversely, acknowledgment of observer-dependence in the manifestational register does not imply that biological phenomena exist only when scientifically represented.

The relational register prevents the opposing simplification in which a marine organism is understood solely as a fixed bundle of intrinsic properties. The same measured physiological trait can have different ecological significance under different environmental conditions or in different assemblages. This is not

equivalent to asserting that traits lack material organization. It means that causal or functional claims about their ecological roles can require an account of the organism's relations. Coral–microorganism associations offer a useful case: contemporary reviews describe corals as hosts associated with complex microbial communities and examine how environmental pressures can alter these associations and coral condition [206]. The existence of these associations is an empirical claim, while the choice to represent a coral colony, a host–symbiont assemblage, or a reef community as the principal analytical unit is a research-design decision that must be justified by the question at hand.

The manifestational register calls attention to the difference between a response produced by a measurement apparatus and a warranted account of the ecological relation that produced it. A fluorescent signal, acoustic return, molecular read, or image segment becomes evidence through protocols of sampling, calibration, processing, and interpretation. The apparatus's registration is a real event, even when the associated ecological interpretation is mistaken. Accordingly, a manifestational account is neither a dismissal of measurement as subjective nor an assertion that data transparently reveal the whole object measured.

The three registers are coupled in both epistemic and practical directions. An assumed ontology helps determine which relations a scientist seeks; knowledge of those relations affects which measurements are designed; unexpected manifestations can support revision of the relational model and, in some cases, of the concepts used to individuate entities. The coupling does not form a guaranteed cycle of progressive improvement. Observations can be incomplete, incommensurable, or misleading; actions can destroy information-bearing conditions; and a revised representation can remain underdetermined. For this reason, the analytic scheme serves as a guide to specifying assumptions rather than a universal mechanism for converting observations into truth.

An operational passage between the registers requires more specificity than assigning a phenomenon to one of three headings. Consider the claim that a moving organism altered its behavior in response to a robotic vehicle. The *manifestational* record may consist of images, acoustic returns, navigation estimates, and instrument settings. The *relational* hypothesis concerns the dependence of a behavioral variable on vehicle approach, illumination, sound, currents, and unobserved causes. The *ontological* question concerns the organism and processes to which the hypothesis is intended to refer. Whether a machine-detected movement is a valid indicator of the attributed behavior remains a separate empirical and interpretive question. Only a controlled or otherwise credible inferential design can justify passage from the first record to the second claim; the terminology of GR cannot supply the missing evidence.

For this purpose, a claim should have an explicit *referent*, an *observation procedure*, an *inferential transformation*, and a *range of admissible alternatives*. An observation of proximity might be reproducible while its interpretation as avoidance remains weak if the vehicle approached only during a particular tidal phase. A sensor change following an intervention might be caused by altered mixing rather than animal activity. The research design should therefore record the observer's action and relevant environmental conditions, compare rival mechanisms where possible, and disclose distinctions that the instruments cannot resolve. This is an epistemic discipline of separating recorded events from interpretations about them, not a stipulation that ecological reality is created by scientific classification.

This distinction also constrains model criticism. A failure of a classifier to recognize an organism may call for revision of a sampling or recognition model; it need not establish that the biological entity has changed. Conversely, a verified ecological transition may require revision of the ontology used to individuate an assemblage rather than merely recalibration of one sensor. The appropriate target of revision is itself a question for investigation. Sections 3, 21, and 25 develop these distinctions at the levels of capacities, sampling policies, and comparative evidence.

2.3 Marine entities, dependence, and contextual relations

Marine environments expose the insufficiency of an exclusively object-centered description through their densely coupled physical and biological processes. Water movement transports nutrients, metabolites, genetic material, and heat. Stratification, currents, substrates, and habitat geometry can mediate the

encounters through which feeding, dispersal, symbiosis, and competition occur. The mechanisms and relative importance of these processes vary with ecological scale and local conditions. Rather than treating a relation between two organisms as a timeless edge with fixed meaning, an inquiry may have to investigate how the relation is conditional on environmental context, life-history stage, and the observational interval.

A concrete empirical illustration comes from long-term observations of a southern California kelp forest. Liu and Gaines [99] used empirical dynamic modeling to examine context-dependent interactions among kelp and urchin species and reported substantial variation in interaction strengths and directions with environmental conditions. The result supports the scientific importance of interaction context; it does not establish that all ecological interactions can be reconstructed from observational time series or that one inferred network uniquely identifies causality. In particular, inferential relationships estimated from time series remain subject to assumptions about sampling, unmeasured processes, and the chosen modeling approach.

We may consequently distinguish at least three ways in which an ecological relation enters research. A *material relation* concerns an interaction or dependency in the investigated system, such as nutrient transfer or predation. An *inferred relation* concerns a model-supported claim about that interaction. A *representational relation* concerns a connection within the scientist’s description, such as an edge of a causal hypothesis graph or an association between data records. These can correspond, but correspondence is something to establish, not to assume. A correlation between animal detections and temperature can describe the evidential record; identifying temperature as a direct cause may require an additional mechanism, design, or intervention. Similarly, an edge drawn between two species in a food-web model does not, without further specifications, exhaust their ecological relationship.

An individual organism may also contribute to the alteration of conditions that subsequently matter for itself and others. The niche-construction literature documents how organisms can modify physical or biotic environments, providing one scientifically developed way to study reciprocal dependence [149]. In the marine setting, habitat-forming organisms, bioturbation, microbial metabolism, and other processes invite inquiry into such reciprocity. However, describing environmental modification as a generative relation does not require assigning the organism an intention to engineer its surroundings, nor does it imply that the modification has a beneficial or stable outcome.

This distinction matters when biological analogies motivate robotic engineering. A fish’s locomotor morphology or a cephalopod’s contact behavior is situated in sensory, physiological, environmental, and developmental processes. A robotic imitation of a visible movement pattern can be useful without reproducing all of those relations. Conversely, a robotic body introduced into a habitat becomes another interacting material system. It can alter illumination, acoustic conditions, or local flow while seeking to observe an animal. The object of inquiry and the instrument are therefore distinguishable, but their realized relations may not be dynamically separable in every observation. The specific effects and magnitudes require field-appropriate measurements and are developed later in the engineering and welfare sections.

2.4 Situated observation and epistemic identifiability

A limited formal description helps make the inferential problem precise. Let $\mathcal{X}_h$ be a set of ecological configurations admitted by a stated model in investigative context h . A configuration may include physical fields, organismal states, and selected relational variables, with its boundaries and state variables explicitly declared. Let a denote an observation procedure, including a sensing modality, a platform action, and relevant calibration conditions. A deterministic idealization writes the recorded output as

$$y = O_{h,a}(x), \quad x \in \mathcal{X}_h, \quad O_{h,a} : \mathcal{X}_h \longrightarrow \mathcal{Y}_{h,a}. \quad (1)$$

The equation does not assert that the marine world is deterministic. A stochastic measurement model can replace $O_{h,a}$ with an observational distribution $P_{h,a}(dy \mid x)$; both formulations are conditional on the selected model and measurement protocol. The subscript h identifies the relevant research context rather than a universal privileged instant or globally ordered history.

For the deterministic idealization, define observational indistinguishability under procedure a by

$$x \sim_{h,a} x' \iff O_{h,a}(x) = O_{h,a}(x'). \quad (2)$$

This relation captures a limitation of a procedure's discriminatory power within the chosen hypothesis space. It does not say that x and x' are ontologically identical.

Proposition 2.1 (Task-relative identifiability under a fixed observation procedure). *Let $\mathcal{X}_h$ and $\mathcal{Y}_{h,a}$ be sets, let $O_{h,a} : \mathcal{X}_h \rightarrow \mathcal{Y}_{h,a}$, and let $q : \mathcal{X}_h \rightarrow \mathcal{Q}$ be a property of scientific interest. There exists a map $\hat{q} : O_{h,a}(\mathcal{X}_h) \rightarrow \mathcal{Q}$ satisfying $q = \hat{q} \circ O_{h,a}$ on $\mathcal{X}_h$ if and only if q is constant on every fiber $O_{h,a}^{-1}(y)$ for $y \in O_{h,a}(\mathcal{X}_h)$.*

Proof. If $q = \hat{q} \circ O_{h,a}$, any two configurations with the same recorded output have the same value of q , proving necessity. Conversely, if q is constant on each fiber, define $\hat{q}(y) = q(x)$ for any x with $O_{h,a}(x) = y$. Fiberwise constancy makes this definition independent of the chosen x , and $q = \hat{q} \circ O_{h,a}$ follows. $\square$

The proposition is a basic set-theoretic fact, not a new law of ecology. It shows precisely why a measurement can identify a task-relevant property without identifying the full ecological configuration, and why an omitted distinction can obstruct reconstruction. A camera may support a particular visible count while failing to distinguish two physiological conditions; a chemical sensor may distinguish oxygen levels without determining the composition of a microbial community. The value of the result is in forcing the investigator to specify q , the admissible ecological configurations, and the actual observational capacity rather than assuming that an increase in data volume establishes general understanding.

For stochastic measurements, the corresponding limitation is expressed through equality or closeness of the observational distributions. If two ecological configurations yield the same distribution under an allowed procedure, no method using only that procedure's output can discriminate them with certainty. Distinguishing two distributions empirically may nevertheless be difficult even when they are formally unequal, because finite samples and dependence structures restrict statistical power. Statements about statistical identifiability, practical estimability, and causal identification must therefore be separated. More detailed statistical and experimental formalization is reserved for the subsequent measurement, engineering, and evaluation sections.

Environmental-DNA observations illustrate this distinction in a particularly clear way. Aquatic genetic material can be transported and can persist after an organism has left a location; its detection consequently does not by itself establish a currently living organism at the sampling position. Goldberg and colleagues [56] discuss transport, degradation, contamination, and other interpretive constraints on ecological inference. Thus, the legitimate inferential object may be *a detected genetic signature under specified sampling conditions*, followed by a bounded ecological hypothesis. Treating that signature as an immediately observed individual would merge a manifestational record with a stronger ontological claim than the evidence warrants.

2.5 Partial histories and the problem of ecological time

The meaning of an ecological relation may depend on histories of environmental change, individual development, and previous interactions. Nevertheless, historical interpretation should not be built on an unexamined global division among a privileged present, an already determined past, and a uniformly available future. The analytical requirement is to specify the ordering relationships justified for the events and processes being represented. A monitoring series may use a clock coordinate for calibration; this is a methodological convenience with defined uncertainties and conventions, rather than a general ontological claim about the structure of ecological reality.

Let E_h denote recorded or hypothesized events relevant to an investigation and let $\preceq_h$ be a declared partial-order relation when the evidence supports one. Examples include an instrument acquisition preceding its laboratory analysis, or a water sample being collected before an associated genomic read

is produced. Different ecological processes may have interacting timescales or evidence that fails to establish a unique order across sites and instruments. In that case the representation should preserve incomparability instead of manufacturing a total chronology. Moreover, an order among records is not automatically a verified causal order among ecological events.

A provisional history object may be written

$$H_h = (E_h, \preceq_h, \lambda_h, \Pi_h), \quad (3)$$

where λ_h associates events with typed descriptions and Π_h records provenance, calibration, uncertainty, and relevant relations of dependence. This is a schema for research representation. It accommodates partial provenance, disputed event relationships, and heterogeneous clocks; the data do not thereby become a full representation of the ecological system's history. If $\preceq_h$ is used as a causal order, additional assumptions and tests are necessary. If it is only a dependency order among records, that narrower meaning should be maintained.

The distinction also matters for ecological generativity. A model may describe admissible transformations or organizational alternatives without claiming that these are unrealized events queued to occur in a universal future. Possibilities represented in a model, historical realizations reconstructed from evidence, and physically supported continuation hypotheses have different epistemic statuses. The mechanisms by which ecological generative capacities depend on material conditions will be investigated in Section 3; here the principal requirement is to avoid using an unqualified temporal sequence as the foundation for relational individuation or scientific inference.

2.6 Relational plurality, scale, and incomplete translation

An ecological phenomenon may be investigated at the scale of biochemical reactions, an organism, a population, a multispecies assemblage, a habitat, or an ocean basin. No one scale automatically answers every question. Levin's analysis of pattern and scale emphasizes how an investigator's window of observation filters ecological variability and how relevant mechanisms can operate at scales different from those at which patterns are observed [92]. A reef photograph, satellite product, water-column profile, and tissue-level genetic assay therefore cannot be expected to preserve the same relations merely because each concerns the same geographic region.

Formally, let R_A and R_B be representations developed under investigative contexts A and B . A proposed translation $T_{A \rightarrow B}$ may be defined only for certain objects or relations. For a selected task property q , translation fidelity is meaningful only after specifying compatible interpretations of that property and the conditions under which comparison is permitted. What is preserved by a map of temperature grids need not coincide with what is preserved by a translation from a community's ecological classifications into scientific taxonomic labels. The latter may additionally involve rights, consent, histories of use, and contested interpretations. A technically valid conversion of a file format cannot, without further work, establish equivalence among scientific explanations or normative judgments.

For example, a scientist studying reef microbial communities and a researcher investigating animal behavior may generate partially overlapping accounts of the same habitat. The first may individuate associations by microbial sequence and host phenotype, the second by repeated behavior and environmental context. Translation between their records can yield a productive cross-disciplinary hypothesis, yet neither description uniquely determines the other. A third participant's local ecological knowledge may introduce distinctions that neither scientific database currently encodes. Such cases call for interfaces that preserve source context, declared transformations, and unresolved alternatives, rather than a compulsory global category system. The institutional and generalized-language implications belong principally to Section 6 and the later nomos and knowledge-circulation discussion.

There is a further danger in interpreting heterogeneous systems through human-centered analogies. An organism's responsiveness to chemical gradients, sound, or light can be investigated as a physical and biological phenomenon; whether this constitutes an epistemic practice analogous to human reasoning is an additional question. The present formulation leaves room to study learning, perception, and organism–environment relationships without treating all living systems as having the same cognitive structure.

Claims about sentience, suffering, and particular welfare interests require their own species-relevant evidential arguments. They will be developed in Section 4 rather than inferred from relational vocabulary alone.

2.7 Relational geometry and a limited cellular representation

The account may be related to geometrical and topological formalisms when the scientific problem actually contains adjacency, connectivity, enclosure, cycles, or multiway interactions. For some applications, a finite cell complex or two-dimensional CW complex can organize a selected history of observation and relational hypotheses. Let

$$K_h = (K_h^{(0)}, K_h^{(1)}, K_h^{(2)}; \partial_1, \partial_2), \quad \partial_1 \partial_2 = 0, \quad (4)$$

where the sets denote cells chosen by the modeling scheme and the boundary maps satisfy the standard chain-complex identity. The construction is an optional representational tool. A vertex need not represent an organism, an edge need not represent a social relationship, and a face need not represent an ecological value. Such assignments must be specified and defended for the individual inquiry. The cellular structure can represent relational incidence even when no continuous spatial manifold is posited; conversely, fluid and biological field models may be more appropriate where quantitative spatial dynamics are indispensable.

For illustration, a cellular model of a collaborative marine observation campaign might assign vertices to measurement or sampling events, edges to typed provenance dependencies, and faces to documented multi-event comparison procedures. This particular construction describes the *investigative process*; it should not be mistaken for an exhaustive topological model of the marine habitat. A different complex might be formed from biologically justified habitat connectivity relations, but translating between these complexes would require an explicit correspondence. Homology classes or cycles could then characterize structures *within the model*, while any ecological significance would need separate physical or biological validation. In particular, a formal hole in an observational network is not itself evidence of a physical void or an unknown species in the ocean.

Nor does a discretized representation automatically preserve transformations that change the relevant units, relation types, or admissible distinctions. Cell refinements and coarse-graining may discard information; a formal isomorphism need not preserve a task-specific ecological property. This links the geometric proposal to the identifiability condition in Proposition 2.1. The topology or geometry of a representation becomes scientifically consequential only through declared measurement procedures, model semantics, and questions for which preservation or loss can be assessed. More elaborate dynamical and topological models are postponed until relevant engineering and formal problems have been specified.

2.8 Consequences for scientific responsibility

This section's distinction among reality, relations, and manifestation has practical consequences without deciding substantive ethical questions in advance. First, an observer should identify what its instrument actually registers and separate the registered event from stronger ecological interpretations. Second, scientific representations should preserve conditions of sampling, uncertainty, and revision so that disagreements about ecological identity or causal relations remain inspectable. Third, because observation can affect the investigated system, a mission must consider whether its own actions change what its instruments subsequently record. Fourth, uncertainty concerning nonhuman organisms should inform appropriate research precautions without being presented as proof of any specific animal experience. Fifth, no successful inference by itself settles who may conduct an intervention or how its consequences ought to be distributed.

These conditions can be operationalized through explicit model boundaries, alternative ecological hypotheses, provenance of transformations, tests of observational identifiability, and independent assessment of disturbance where relevant. They also produce potential failure cases: a classifier trained

on one habitat may apply its categories in a different ecological context; a detected DNA signature may be misread as an immediately present animal; an inferred species-interaction edge may be treated as a demonstrated mechanism; or a physically accurate measurement may be used to justify an intervention beyond its evidential scope. These examples represent research vulnerabilities, not findings from new experiments conducted for this paper.

The central conclusion is limited but constructive. Marine ecological reality can be investigated through many complementary representations, each generated under situated physical, methodological, and institutional conditions. A generative-relational account becomes useful when it clarifies the formation, dependence, and transformation of those representations and the ecological relations they aim to characterize. It becomes misleading when it claims to replace empirical biology, to identify a universal subject of the ocean, or to turn representational connections into established ontological facts. The subsequent section examines ecological generativity and material finitude in more detail, asking which capacities may develop, what dependencies sustain them, and how inquiry can alter those conditions. The present section supplies the epistemically modest conceptual foundations for that task.

3 Ecological Generativity and Material Finitude

3.1 Generative capacities and their conditions

This section investigates the relationships among ecological generativity, material constraints, and the persistence or transformation of marine ecological possibilities. It develops a bounded generative-relational (GR) vocabulary, examines ecological evidence concerning resource limitation and regime change, and specifies what oceanic robotics might need to observe before a claim about an ecosystem's regenerative or transformative capacities can be warranted. Its principal methodological commitment is that *the ability of a system to generate further relations or forms of organization entails neither unlimited material supply nor an intrinsic direction toward improvement*. Generativity and finitude are compatible properties of the same evolving configuration.

The expression *ecological generativity* will designate, relative to a declared scientific question, capacities through which organisms and their relations can produce, maintain, vary, reorganize, or sometimes extinguish conditions for further ecological processes. Reproduction, development, metabolism, microbial association, habitat construction, dispersal, behavioral learning, and evolutionary diversification involve different mechanisms, variables, and evidential standards. Their grouping under one investigative term does not assert a common causal law or a scalar magnitude of "generativity." Biological reproduction can increase the abundance of a population while narrowing genetic variation under some conditions. Habitat formation can create opportunities for some species and exclude others. An altered ecological configuration may sustain novel relations even as particular organisms disappear. The consequences must be described through specific histories and biological models rather than a presumption that novelty is beneficial or that persistence is always possible.

This approach intersects with established ecological inquiry. Niche-construction research examines the alteration of environmental conditions by organisms and the resulting ecological or evolutionary consequences [149]. Process-oriented accounts of biology emphasize the activities that maintain living organization [34]. Ecological studies of resilience examine the capacity of systems to absorb disturbances while maintaining or transforming recognizable patterns of organization [62]. The proposed GR contribution concerns the articulation of these mechanisms with the conditions of their investigation: which possibilities a researcher can distinguish, how observational and inferential practices modify those distinctions, and which changes in material or ecological relations make particular continuations available or inaccessible. Such an account is supplementary to empirical ecological theory and should be tested against it.

The terms *capacity*, *possibility*, and *realization* require separate treatment. A capacity concerns a relationally supported disposition or set of admissible operations; a possibility is specified relative to a model, constraints, and a particular question; and a realization is a process or history for which specified

evidence is available. None is necessarily equivalent to a future event waiting to occur. Likewise, a relation judged feasible within a formal representation may be physically infeasible because relevant dependencies were omitted. A scientific explanation should disclose the modeling conditions that connect its representations to the ecological mechanisms under study.

3.2 Biological, relational, and epistemic dimensions of generativity

Different forms of ecological generativity can be distinguished without making them mutually exclusive. *Biological generativity* concerns mechanisms such as reproduction, development, regeneration, and heritable variation. *Relational generativity* concerns changes in the interaction structures through which organisms, environments, and material processes make further interactions possible. *Environmental generativity* concerns processes such as substrate formation, nutrient recycling, and habitat modification that affect the conditions available to organisms. *Epistemic generativity*, by contrast, concerns the capacities of investigators to construct better questions, instruments, models, and interpretations. The last category belongs to the scientific inquiry and need not be a property of the ecosystem itself.

These distinctions help avoid a frequent ambiguity. When a new sensing method reveals a previously unobserved microbial association, human knowledge-producing capacity has expanded; the biological association may have existed independently of its discovery. When a restoration measure enables renewed habitat formation, the ecological possibilities may themselves have changed. When an observing robot causes animals to avoid a breeding area, the encounter can simultaneously create a scientific record and transform the relational conditions being recorded. The investigator must distinguish these situations through observations and defensible causal analysis.

A useful marine example is the ecological association among sea otters, sea urchins, and kelp. Classical field studies identified how predators can influence herbivores and thereby alter the persistence of nearshore macroalgal communities [37]. The ecological relevance of a kelp forest cannot be adequately represented by a single inventory of kelp biomass without considering species interactions, recruitment, grazing pressure, habitat conditions, and the consequences of disturbance. This example illustrates the possibility that an apparently local change in one relation alters a much larger set of conditions sustaining other relations. It does not imply that every kelp ecosystem has an identical trophic structure or that every intervention can restore a former ecological configuration.

From the perspective developed in Section 2, a scientific account must also distinguish the ontological, relational, and manifestational registers. A sensor-produced change in an ecological index is a manifestation under a defined procedure; an inferred change in trophic interactions is a relational claim; and stronger claims about what constitutes the biological system require separate ontological argument and evidence. Expanding access to manifestations may transform the investigator's categories without establishing that the underlying ecological system has become more generative. This methodological distinction will later inform both automated model revision and evaluation of robotic interventions.

A comparison of generative capacities needs a specified *capacity claim*, not a presumed common metric. One hypothesis might concern the persistence of reproductive opportunities under a known temperature and exposure regime; another might concern whether a habitat-forming species supports a wider range of settlement relations; a third might concern a research consortium's ability to develop independently testable ecological explanations. These hypotheses have different units of analysis, reference conditions, and potential counterexamples. A rise in recorded species richness at a newly surveyed site could indicate improved detection rather than newly generated biological diversity. Likewise, an increase in available genomic records could enlarge scientific possibility while leaving actual habitat viability unchanged. The model should treat such evidential possibilities separately.

To assess a proposed biological capacity, investigators should declare at least the relevant entities, enabling material relations, candidate continuation events, observational window or partial history, comparison condition, and evidence by which the claim could fail. A capacity need not be realized in the observed interval to be physically plausible, but plausibility alone is not an estimate of its prevalence. Where realization is observable, causal attribution to a specific enabling condition requires further design,

especially when environmental covariates affect both the intervention and its outcome. For ecological systems whose continuation conditions are poorly known, the correct record may be a set of rival, partly identified models, including one in which the hypothesized capacity is absent.

This method has a direct engineering consequence. A controller's ability to maintain a defined exclusion envelope is a fact about the controlled plant and the specified envelope. It becomes evidence about ecological generativity only if the envelope has been independently connected to a biologically meaningful mechanism and its uncertainty has been assessed. The distinction prevents a formal viability set from being mistaken for a demonstration of organism welfare or ecosystem recovery. It also provides a common research question linking the ecological account here to the constrained-control models in Section 14 and welfare analysis in Section 22.

3.3 Material constraints, resource flows, and heterogeneous renewability

The term *finite marine resources* covers several scientifically different phenomena. Some material stocks are nonrenewable on operational or human institutional horizons, such as particular geological formations and mineral deposits. Some stocks, including living populations, renew conditionally through biological processes but can decline under extraction, habitat loss, or unfavorable environmental regimes. Other relevant constraints arise from rates of supply or exchange: available oxygen, nutrient flux, light, usable energy, substrate, dispersal opportunities, and the assimilation of disturbances. A body of water contains extensive matter and energy, but their mere presence does not imply that a particular biological or technological process can use them at an arbitrarily chosen place, rate, or ecological cost.

Marine plankton provide a concrete illustration. Elemental requirements and environmental nutrient supply influence phytoplankton physiology and biogeochemical processes, while cellular carbon:nitrogen:phosphorus ratios vary with lineage, growth conditions, and environmental context. Consequently, one universal fixed elemental ratio cannot exhaust the relevant nutrient constraints [118]. Physical oceanic processes, including stratification, currents, mixing, and upwelling, further affect the spatial and temporal conditions of nutrient and oxygen availability. These relations are coupled: a change in mixing can modify nutrient supply, which can alter biological production, which can feed back into oxygen consumption and other chemical processes. Their interacting scales preclude a simple interpretation of the marine environment as an inexhaustible storehouse.

For a declared control volume and a particular conserved or approximately balanced constituent, a provisional bookkeeping relation is

$$R_{k+1} = R_k + I_k - O_k - H_k + D_k, \quad (5)$$

where R_k denotes a measurable stock, I_k and O_k denote influx and outflux, H_k denotes a designated removal, and D_k accounts for specified internal sources or sinks expressed in the same units. The subscript k indexes a selected sequence of measurement or accounting operations; it is not a universal ontological clock. Equation (5) is an accounting identity only when boundaries, constituents, units, and treatment of internal transformations are specified consistently. For truly conserved mass one would generally set net internal production to zero for a closed boundary, whereas chemical species can be converted into other species. A model that ignores migration, redistribution, biological storage, or measurement bias may produce a misleading estimate of available resources even while its arithmetic is correct.

Biological renewal is an additional conditional process, not an automatic credit to the stock. A population can reproduce while recruitment to the locally measured population remains insufficient to compensate for mortality or extraction. The conditions of renewal depend on age structure, genetic and physiological states, species relations, habitat, and environmental variability. Therefore, the classical question of how much material can be removed sustainably is incomplete unless the model specifies the biological and institutional meaning of "sustainably," its measurement horizon, and the responses of other components. An operational limit inferred from one population or one ecosystem state cannot automatically be transferred to a different species or region.

The material costs of knowledge production must also be counted. Oceanic research can require vessels, propulsion energy, batteries, sensor manufacturing, calibration infrastructure, sampling consumables, data centers, communications, maintenance and eventual retrieval. Deployment may impose acoustic, thermal, light, substrate, or chemical effects. The possibility of placing computing or energy infrastructure in marine environments does not imply that the environment can absorb unlimited heat or disturbance. Assessment should therefore distinguish renewable operational inputs, embodied materials, ecological exposures, and the degree to which consequences can be observed or reversed. More detailed engineering estimates belong to Sections 19 and 24; here the point is philosophical and methodological: *knowledge production itself occupies a finite material and ecological context.*

3.4 Conditional possibility spaces and represented continuation

A generative-relational formulation may help make the dependence of apparent possibilities on model choices explicit. Let h denote a represented ecological history assembled from a specified set of events, observations, and inferential relations. Let $\mathcal{C}$ designate constraints supported or stipulated by the model, and let Π designate the investigative questions and distinctions by which candidate continuations will be compared. Define

$$\text{Cont}_{\mathcal{C}}(h) = \{c : c \text{ is an admissible continuation of } h \text{ under } \mathcal{C}\}, \quad (6)$$

and the task-relative generative repertoire

$$\mathcal{G}_{\Pi}(h; \mathcal{C}) = \{\pi_{\Pi}(c) : c \in \text{Cont}_{\mathcal{C}}(h)\}. \quad (7)$$

Here π_{Π} is an explicitly defined representation of properties relevant to the inquiry. A “continuation” is a compatible extension in the selected model of partial histories; it need not inherit a universal before–after order or represent an event that has not yet been realized. Different physical and biological models can assign different sets of admissible extensions to the same incomplete observational record.

Equations (6) and (7) offer a modeling vocabulary rather than an observed numerical measure. The extension set can be unbounded in a mathematical model despite strictly limited biological resources; formal cardinality therefore cannot be interpreted directly as ecological abundance, diversity, or justice. Replacing Π by a different set of distinctions can enlarge or reduce the represented repertoire without any corresponding ecological transformation. Similarly, tightening $\mathcal{C}$ can remove modeled continuations because additional evidence has ruled them out, while the physical system remains unchanged. A claim that ecology has “lost possibilities” must consequently identify whether it concerns actual ecological change, changes to evidential constraints, or changes in the observer’s representational vocabulary.

A simple dependency relation helps specify how some possibilities are sustained. Let a modeled continuation c require several resources or relations $d \in D(c)$. Feasibility then depends on the joint satisfaction of requirements rather than a free-standing potential attributed to an organism:

$$\text{Feasible}(c \mid h, \mathcal{C}) \implies \bigwedge_{d \in D(c)} \text{Supported}(d \mid h, \mathcal{C}). \quad (8)$$

The implication deliberately runs in one direction: identified dependencies may be necessary without being sufficient. Unknown dependencies, nonlinear dynamics, historical contingencies, and incomplete observation may invalidate a candidate continuation. The model should allow dependence sets to be revised when new evidence exposes interactions previously absent from its representational language. This is one way to connect ecological generativity to a historically evolving scientific epistemology without reducing ecological phenomena to the state of the observer’s software.

3.5 Resilience, transitions, and incomplete recovery

A system’s present appearance does not determine its responses to intervention. Ecological resilience concerns the persistence or transformation of an organization’s relevant functions and relations under

disturbance; it must be defined relative to the modeled variables, perturbations, and scales. Holling's influential distinction between resilience and particular notions of stability remains important here [62]. Ecological transitions may be nonlinear, with feedbacks producing alternative persistent configurations, hysteresis, and different conditions for disturbance and recovery [160]. A stable appearance within a limited observation window is consequently weak evidence for indefinite regeneration.

Kelp forests and urchin barrens provide a marine example. Comparative evidence and manipulative experiments indicate that some temperate reef systems exhibit an asymmetry between the grazing pressure that contributes to the loss of kelp-dominated habitat and the substantially lower pressure required for recovery [97]. Such hysteresis means that simply restoring a variable to its earlier measured value need not restore the associated ecological organization. It also shows why observations restricted to biomass may fail to identify feedbacks that govern recruitment or persistence. Thresholds and mechanisms remain site-dependent; the cited literature does not establish one global transition value for all marine habitats.

We should also distinguish *persistence*, *recovery*, *restoration*, and *transformation*. Persistence can describe continuity of selected structures or functions under changing conditions. Recovery concerns a specified response of measured ecological properties after disturbance. Restoration additionally invokes an evaluative reference condition and a chosen intervention objective. Transformation can involve the emergence or disappearance of relations, possibly without return to an earlier state. These terms require different evidential and normative commitments. In particular, a habitat that contains living organisms after disturbance has not thereby recovered every previous ecological relation or every capacity important to those organisms.

This difficulty is particularly visible in slow-changing seabed environments. A field evaluation of the area affected by a 1979 deep-sea nodule-collection test found persistent physical and biological changes after approximately 44 years, alongside partial recolonization and comparatively limited residual indications in some nearby plume-affected areas [83]. These findings should neither be universalized to all subsea activity nor read as evidence that recovery is impossible. They demonstrate why effects, recovery pathways, and reference conditions must be documented separately. They also caution against the assumption that a successfully completed technological operation has only short-lived ecological consequences.

At broader scales, ocean warming, acidification, deoxygenation, and changing circulation affect marine habitats and biological processes. The IPCC's *Special Report on the Ocean and Cryosphere in a Changing Climate* synthesizes evidence of ecosystem shifts and differential vulnerabilities arising from these interacting drivers [68]. These changes can operate on scales that exceed a single robotic mission. A local measurement campaign must therefore distinguish the conditions it can identify from larger influences it cannot resolve. Ecological criticality cannot be inferred reliably from a striking sensor anomaly without testing appropriate rival explanations, including instrument faults, transient physical variability, and changes in measurement context.

3.6 Habitat dependence, functional recovery, and ecological reference conditions

The distinction between ecological persistence and generativity acquires an operational meaning where the material substrate of a community is itself one of the conditions for its continuation. In the Clarion–Clipperton Zone, standardized visual comparisons associated greater epifaunal densities with high polymetallic-nodule coverage; particular attached taxa were scarce in surveyed nodule-poor sites [203]. The inference is habitat- and survey-conditioned: nodules supply attachment surfaces for organisms that require hard substrate, whereas some sediment-associated organisms have different dependencies. This is not a claim that every nodule is biologically equivalent or that all community attributes are caused by nodule density alone. It indicates that extracting a geological material can also remove a living community's substrate, a condition that a short-run recolonization count cannot restore by itself.

Studies of mining simulations and earlier collector tracks expose why *recovery* should be reported as a vector of endpoints. A meta-analysis of 11 disturbances found substantial initial declines in many

benthic groups and evidence of partial re-establishment in some smaller and mobile faunal categories, while few monitored assemblages returned to control conditions even after decades [82]. Twenty-six years after the DISCOL experiment, photographic comparisons found continued community differences and lower suspension-feeder occurrence in disturbed areas [168]. Sampling of DISCOL sediments also documented persistent microbial and biogeochemical differences, extending the relevant endpoints beyond the organisms most visible in imagery [205]. The later assessment of the 1979 OMCO test, which documented partial re-establishment across some taxa after 44 years yet persistent direct-track changes, further complicates any undifferentiated assertion of either recovery or irreversibility [83].

These records concern particular designs for seabed mining, mechanical substrate disruption, and geographically distinct samples. They cannot be directly converted into impact factors for a hovering camera platform or a single lightweight recoverable anchor. At the same time, the measurements suggest a relevant set of *questions* for less intensive projects: Does the system disturb hard-substrate associations, redistribute fine particles, repeatedly revisit the same microhabitats, or obstruct recolonization? What fraction of the footprint experiences contact, and what comparison sites plausibly represent its alternative trajectory? A small instrument's burden should be measured at its own scale, while its scientific justification remains separate from a commercial collector's.

Seamount corals offer a second, nonidentical recovery case. Surveys of northwestern Hawaiian and Emperor seamounts reported signs of coral regrowth and greater benthic megafaunal abundances at some protected sites following decades of reduced bottom-contact fishing, despite continuing evidence of historical damage [16]. Neither a recovered colony nor a higher image-count density determines that the earlier composition, recruitment process or habitat function has returned. Conversely, documenting incomplete recovery does not justify assuming zero recovery capacity. Recovery baselines must specify what is compared: an earlier observed assemblage, a nearby reference site, a functionally comparable habitat, or a normative restoration objective. Such reference conditions are themselves affected by long-run climatic change and imperfect historical observation.

Within the paper's GR vocabulary, one can therefore distinguish continuation of *observable abundance*, *community composition*, *habitat-forming relations*, *microbial function*, and *recolonization capacity*. These dimensions can move in different directions and should not be treated as independent just because they are tabulated separately. Their empirical dependencies, measurement uncertainties and normative relevance have to be identified before any statement about an ecological generative repertoire can be warranted.

3.7 A bounded formal problem of ecological viability under partial observation

The distinction between generativity and finitude can be connected to a practical question of control and scientific responsibility. Consider a deliberately simplified ecological state space X , an action set U , environmental uncertainty W , and a discrete transition relation

$$x_{k+1} = F(x_k, u_k, w_k), \quad y_k = O(x_k). \quad (9)$$

Let $S \subseteq X$ denote a scientifically and normatively specified admissible-state region for one explicitly bounded purpose. For a state $x \in S$, define the actions capable of preserving one-step membership in S under all modeled disturbances,

$$U_S(x) = \{u \in U : F(x, u, w) \in S \text{ for every } w \in W(x, u)\}. \quad (10)$$

The set S might encode an experimentally justified ecological exposure limit or an engineering exclusion condition, but no particular S establishes animal welfare, moral legitimacy, or ecosystem integrity merely by being mathematically declared. Its bounds require independent justification.

Proposition 3.1 (One-step action ambiguity under limited ecological observation). *Suppose $x_1, x_2 \in S$ satisfy $O(x_1) = O(x_2) = y$, and that $U_S(x_1) \cap U_S(x_2) = \emptyset$. Then no deterministic observation-only policy $u = \mu(y)$ can guarantee one-step preservation of S for both states under the modeled disturbances.*

Proof. An observation-only policy assigns the same action $\mu(y)$ to x_1 and x_2 . Guaranteed preservation at x_i requires $\mu(y) \in U_S(x_i)$ by definition (10). Such a policy would therefore require an action in $U_S(x_1) \cap U_S(x_2)$, contrary to the assumed empty intersection. $\square$

The result is modest but useful. Observational indistinguishability can produce *action-level* limitations even where a controller is computationally powerful. A robot may observe an apparently similar benthic surface in two circumstances while relevant differences in hidden habitat condition, animal presence, or sediment sensitivity make different maneuvers admissible. The proposition does not claim that those ecological conditions have been empirically identified; it describes a failure mode of incomplete state information. A controller may address it by collecting discriminating measurements, reducing intervention, choosing a conservative action when one exists, or referring the decision for review. A randomized policy does not yield a uniform guarantee when each action in its support fails under at least one compatible state. The difficulty is structural rather than merely algorithmic.

An elementary numerical illustration can make the assumptions inspectable. Let $X = S = \{a, b\}$, $U = \{u_a, u_b\}$, and suppose both states produce the same observation y . Under a disturbance-free toy transition, let $F(a, u_a) = a$ and $F(a, u_b) = \dagger$, while $F(b, u_b) = b$ and $F(b, u_a) = \dagger$, where $\dagger \notin S$ denotes a failed admissibility condition. Then $U_S(a) = \{u_a\}$ and $U_S(b) = \{u_b\}$. No observation-only action is guaranteed safe across the two cases. These labels represent no particular species, geography, or operational ecological threshold; their role is to expose the inference limit cleanly.

The concept is related to viability analysis and constrained control, but its practical use requires significantly more than the finite proposition. Real ecological states are partially observable, habitat responses may be stochastic and delayed, the admissible set may itself be contested, and ecological indicators may be affected by the observer's intervention. If $U_S(x)$ is empty, the model yields no preserving action even with full state knowledge; this is a warning about the model and feasible intervention set, not a mandate to invent a safe maneuver. Nor can the proposition certify that a large marine intervention is ethically permissible. A more realistic robust or probabilistic formulation must specify credible disturbance sets, calibration, monitoring, governance procedures, and acceptable failure conditions. Sections 14, 21, 22, and 25 will investigate such requirements in relation to oceanic control, scientific inquiry, animal welfare, and comparative evaluation.

3.8 Generative scientific opportunity and ecological cost

Marine robotics can increase the resolution, frequency, and spatial reach of ecological measurement; it can also reveal processes that require new inferential categories. But the scientific value of a mission and its ecological consequences must be assessed as distinct, interacting objects. A persistent robotic observatory might record ecological variation that a short voyage could miss, while creating its own footprint through equipment, lighting, acoustic exposure, maintenance activities, and habitat occupation. A sampler might supply evidence needed to evaluate a rare population while removing individuals or material from that population. A new model of ecological recruitment might improve conservation decisions while also being usable for extraction or selective exploitation. The appropriate assessment must examine the relevant organisms, exposures, interactions, temporal scales, and institutional objectives.

The challenge is not solved by assuming that obtaining more data always expands beneficial generativity. Additional measurements may be redundant, poorly calibrated, or destructive of precisely the relations they aim to study. They may improve the research institution's predictive or productive capacities while diminishing possibilities for organisms or communities that do not participate in the decision. Even when a new technology lowers the disturbance per observation, total disturbance may rise if the technology makes much more extensive exploration attractive. The realized consequences depend on scale of deployment, behavior of operators, environmental conditions, and effective governance. These concerns motivate lifecycle assessment and independent ecological exposure analysis alongside measurements of scientific performance.

A practical evaluation can begin by keeping separately reported outcome vectors rather than aggre-

gating all considerations into one scalar objective. For a mission design u , one may report

$$\mathcal{Q}(u) = (\mathcal{I}(u), \mathcal{D}(u), \mathcal{R}(u), \mathcal{P}(u)), \quad (11)$$

where $\mathcal{I}$ records a defined scientific information or inference outcome, $\mathcal{D}$ a vector of ecological disturbance indicators, $\mathcal{R}$ material and energy requirements, and $\mathcal{P}$ the status of procedural review and applicable authorization. These quantities need not be measured in commensurable units; $\mathcal{P}$ is not a numerical reward for a machine to trade against animal harm. Relevant legal and ecological constraints should be applied when determining which designs are admissible for a given project. For unresolved uncertainties or conflicting evidential interpretations, the report should preserve the disagreement rather than assign an unjustified common score.

This approach also prevents confusion between productive potential and public benefit. A technological capacity to perform intensive marine sampling can produce valuable science, but it may also increase exclusion, ecological exposure, or dependence on privately controlled infrastructure. A public knowledge repository might circulate research outputs while providing few of the instruments or research capacities needed for others to reproduce them. The political economy of generative capacities therefore concerns control over methods, materials, interpretation, access, and the ability to contest decisions, as well as money. Sections 7–10 will examine how marine nomos and common knowledge institutions can address these questions. The discussion there will distinguish inspirations from existing generative-justice literature, including the work already acknowledged in Section 1, from the present paper’s independently developed GR models.

3.9 Consequences for ecological inquiry

The preceding analysis establishes four conditions for the remainder of this investigation. First, claims about ecological generativity should identify the processes and dependencies that warrant them, including distinctions between biological reproduction, relational transformation, and new scientific knowledge. Second, claims about resource availability should specify relevant stocks, flows, operational boundaries, regeneration mechanisms, and uncertainty; renewability is conditional and cannot be presumed from the ocean’s apparent physical extent. Third, evaluations of resilience or recovery should identify the ecological variables and relations of concern, the disturbances examined, and the reference conditions used in interpretation. Fourth, robotic mission design should account for the possibility that limitations in sensing and control create undecidable or unsafe actions even when the physical vehicle remains technically functional.

These conditions do not establish a general formula for an optimum level of ocean exploration. They establish questions that can be investigated and contested through empirical research and institutional deliberation. New technologies might make marine relations newly observable, discover unfamiliar mechanisms, or allow less intrusive forms of investigation. Their contribution to ecological life and public knowledge depends on the consequences of their implementation and on how resulting capacities and responsibilities are distributed. The next section examines nonhuman life and epistemic asymmetry, including the ethical difficulties that arise when human observers encounter organisms whose interests, capacities, and possible experiences are known only incompletely.

4 Nonhuman Life and Epistemic Asymmetry

4.1 The epistemic and ethical problem of encountering marine life

This section investigates what can be learned about nonhuman marine organisms through technologically mediated encounters, how the evidential limits of those encounters affect scientific claims, and what responsibilities arise when observation can alter the conditions of nonhuman life. The argument connects the relational account in Section 2 and the account of ecological generativity and finitude in Section 3 to questions of biological agency, possible sentience, animal welfare, and scientific restraint. Its objective

is to establish concepts and investigable constraints for oceanic research, while reserving a fuller account of proceduralism and ecological justice for Section 5 and engineering-specific protocols for Section 22.

Marine organisms present researchers with forms of organization, perception, action, and environmental dependency that need not correspond closely to human perceptual or practical categories. The encounter with a cephalopod, fish, crustacean, marine mammal, coral colony, or microbial community involves different biological mechanisms and different evidence about possible experiences and interests. A robotic platform may record movement, color change, prey capture, reproductive behavior, or response to an approaching vehicle. Such records can warrant claims about particular behavioral or physiological processes under stated observation conditions. They do not by themselves disclose the organism's complete experiential life, nor do they establish that every biologically responsive system possesses an epistemology in the human sense.

An important distinction is required between *an organism as a subject of scientific inquiry* and *an organism as a subject of experience*. The former is a role in a research relation. The latter is an empirical and philosophical claim that depends on evidence appropriate to the species and the capacity at issue. Within the present generative-relational (GR) inquiry, an organism may be studied as a participant in ecological relations without granting or denying it a complete, human-like representational language. A response to light, pressure, acoustic stimulation, or a moving robot can involve sensory transduction, physiological regulation, learning, interaction, or some combination. Assigning a specific subjective meaning requires additional justification.

The practical importance of these distinctions increases as marine research becomes more autonomous. An instrument can register a fish turning away from its path with high temporal precision while lacking reliable evidence about whether this response signifies perceived danger, curiosity, territorial defense, acoustic irritation, or an unrelated ecological event. An animal may alter its activity before it enters the visual field of the camera; it may remain outside an instrument's detection range; or its behavior may change after repeated encounters. Thus, increases in the quantity or spatial extent of observation need not produce proportional increases in understanding. The possibility of systematically interpreting another living system is an empirical research problem requiring biological knowledge, methodological comparison, and humility about residual uncertainty.

4.2 Biological agency, responsiveness, and possible sentience

The term *agency* can describe capacities for context-dependent action or regulation without assuming human deliberation. Organisms can act through sensorimotor coordination, metabolic regulation, ecological modification, developmental plasticity, or learned responses. These capacities depend on biological structures and environmental relations. For scientific purposes, agency claims should specify which choices, constraints, behavioral flexibilities, or causal mechanisms have been measured. Calling an animal an agent does not by itself establish reflective intentionality, moral personhood, or consciousness. Conversely, limited evidence of human-like planning does not imply an absence of ethically relevant capacities.

Sentience, as discussed in contemporary animal-welfare research, commonly concerns the capacity for valenced experiences, including negatively or positively experienced states. *Nociception* concerns detection and processing of potentially damaging stimuli, which may occur without demonstrating a conscious experience of pain. *Cognition* encompasses capacities such as learning, discrimination, memory, or problem-solving. These constructs overlap but are not interchangeable. For a particular marine taxon, responsible interpretation requires attention to the kind of evidence available for each claim and to plausible alternative mechanisms. The diversity of marine organisms makes a single general attribution particularly fragile.

Empirical research provides reasons to take some nonhuman marine experiences seriously. Crook's study of octopuses reported place avoidance associated with noxious treatment, preference for contexts associated with pain relief, and neurophysiological evidence consistent with affective pain processing [29]. A systematic report examining cephalopod molluscs and decapod crustaceans synthesized evidence

across several criteria and recommended their inclusion in relevant animal-welfare law [20]. These results do not authorize unrestricted extrapolation from the sampled species, experimental settings, and investigated capacities to all marine invertebrates. They nevertheless undermine a blanket presumption that complex experience can be dismissed in organisms with markedly nonhuman body plans.

The resulting epistemic position is neither an attribution of the same inner life to every organism nor a demand for direct subjective proof before any concern is warranted. Birch’s analysis of the precautionary principle in animal sentience supplies a philosophical precedent for decisions made under incomplete but relevant evidence [19]. In this paper, precaution is investigated as a bounded policy of evidence gathering, exposure limitation, and accountable choice under uncertainty. The biological content of such decisions must remain species- and context-sensitive, and normative thresholds cannot be obtained by an algorithm from observations alone.

4.3 Asymmetric access to ecological and experiential relations

The asymmetry of an encounter takes several forms. *Sensory asymmetry* arises when the observer and organism detect different modalities, ranges, or environmental cues. *Action asymmetry* arises when the robot can initiate motions, illuminate a space, collect samples, or interrupt habitat use while the organism has limited means to alter the mission’s objectives. *Interpretive asymmetry* arises because humans formulate the questions, categories, and criteria used to interpret observations. *Institutional asymmetry* arises when permissions, benefits, and remedies are established by human legal and research systems even where the consequences extend to nonhuman life. These asymmetries need not be uniform across taxa, sites, or observational practices.

The three analytical reality registers proposed in Section 2 can clarify the problem. At the ontological register, the marine organism and its environment possess properties and interactions that may extend beyond current scientific accounts. At the relational register, its capacities and interactions arise within ecological and material dependencies, including relations modified by the observing system. At the manifestational register, particular signals become available through the instruments and interpretive practices of a research programme. These distinctions do not imply that the organism’s experiential reality is transparently accessible through a human instrument or that all three registers are independently observable layers of nature. They are tools for identifying where an inference acquires evidence and where it remains dependent on assumptions.

Consider a task-specific observational model. Let x describe a configuration involving an organism, a robot, and relevant environmental relations. Let m specify a calibrated measurement practice, and let y denote its observable outcome. The measurement may be expressed schematically as

$$y \sim P_m(\cdot \mid x, u, c), \quad (12)$$

where u represents robotic operations and c the declared observational context. A welfare-relevant state or outcome w may depend on features of x that are imperfectly expressed in y . If different configurations compatible with the same observations support materially different welfare interpretations, the inference from y to w is underdetermined. Section 2 established a general identifiability condition for scientific targets; the present application concerns the unusually consequential limitations of interpreting a living organism’s interests and experiences from partial records.

It is therefore useful to keep a *plural evidential profile* rather than equating one measured signal with welfare. Such a profile may draw on behavior, physiology, exposure, ecological context, developmental and life-history knowledge, species-specific validation, and independent expert assessment. Each observation can be relevant without supplying a universally reliable decoding rule. A cessation of movement might indicate rest, concealment, freezing, reduced responsiveness, or sensor failure. An apparent attraction to a robot might be driven by light, food cues, curiosity, or a learned association, and it cannot by itself establish positive welfare. The interpretation must remain corrigible as scientific evidence changes.

This uncertainty also bears on generalized translation across relational systems. A researcher may construct a partial functional translation from an animal’s response into a model of sensorimotor coordination, resource use, or habitat preference. The translation can be useful for prediction and may support

conservation decisions. Its success at one explanatory task does not demonstrate correspondence between the model’s categories and the distinctions through which the organism regulates its own behavior. For example, an octopus’s manipulation of a shelter can be represented geometrically as contact and force, behaviorally as exploration or concealment, or ecologically as habitat use. These representations can be jointly informative without providing a lossless account of the octopus’s generative capacities or possible experience. An epistemically responsible infrastructure should preserve the task and evidential scope of each correspondence, the uncertainty of alternative interpretations, and the animal-specific data on which the correspondence depends.

A second asymmetry concerns opportunities for correction. Human participants can sometimes challenge a translation of their testimony, explain an intended meaning, or decline a research intervention. A marine organism can respond behaviorally or physically, but that response is not generally part of a shared deliberative procedure in which the researcher can verify the animal’s reasons. The absence of human linguistic testimony should not be filled by an unchallengeable human claim to have spoken on the organism’s behalf. A constructive alternative is to expose the assumptions of the interpretation, permit independent scientific challenge, observe avoidance and signs of disturbance, and provide procedures for revising or withdrawing an intervention. These procedures cannot fully eliminate asymmetry, but they can constrain the authority exercised through it.

4.4 Robotic observers as causes of the observations they record

Oceanic robotic observation is a physical encounter. For sea-surface platforms, water-column vehicles, and seabed instruments, potential exposure pathways include noise, artificial illumination, flow disturbance, contact, habitat occupation, sediment resuspension, chemical leakage, repeated approach, and the operations of deployment and retrieval. The type and magnitude of disturbance depend on the platform, organism, habitat, and conditions of exposure. A robot’s quiet propulsion or soft structure may reduce a specified mechanical exposure while leaving other relevant pathways, such as light or repeated proximity, unresolved.

Benoit-Bird and colleagues documented platform-, depth-, and illumination-dependent behavioral responses of pelagic animals during robotic surveys, showing that animal avoidance or attraction can alter the scientific sample itself [17]. The scientific implication is simultaneously epistemic and ethical. If a robot changes an animal’s position or behavior, its survey may misrepresent the baseline ecological distribution; if the response entails disruption or risk, an apparently informative survey may also be welfare-relevant. The existence of a behavioral response alone, however, does not specify its experienced valence or long-term consequence. Independent exposure and biological evidence remain necessary.

Recent studies sharpen the problem without offering a universal solution. A controlled study reported that fish avoidance of a conspecific-like robotic fish depended on approach distance, relative speed, and repeated exposure [202]. A September 2026 preprint comparing conventional and sea-turtle-inspired vehicles in a reef environment reported measurable animal responses to both platforms; between-robot differences did not remain statistically significant after correction for multiple comparisons [154]. Its preliminary status and study-specific design make it unsuitable as proof that biological resemblance guarantees less disturbance. In general, biomimetic morphology may change the animal’s response without making the encounter ecologically neutral.

A causal distinction matters here. Let $B(0)$ denote a defined behavioral outcome under a specified no-approach reference condition and $B(1)$ the corresponding outcome during a robotic approach. The contrast

$$\Delta_B = B(1) - B(0) \quad (13)$$

can describe an intervention effect for a measured response, subject to assumptions about measurement, comparability, and the relevant causal contrast. One observed approach does not reveal both potential outcomes for the same ecological unit under identical background conditions. Randomized exposure, appropriate controls, repeated independent observations, validated natural experiments, or sensitivity analysis may improve causal identification; each has practical and ethical limitations. Moreover, a measured

Table 2: Different evidential objects in the assessment of encounters with marine life. Indicators are illustrative and require taxon-specific validation.

Object of assessment	Potential observations	Principal inference boundary
Immediate behavioral response	Approach or avoidance, altered movement, habitat use	Response does not uniquely identify subjective state
Physiological condition	Injury, stress-associated measures, recovery	Indicators and baselines can be species-specific
Individual welfare	Multiple bodily, behavioral, contextual indicators	Experiential interpretation is indirectly evidenced
Population effects	Reproduction, survival, abundance trends	Population trends do not determine individual welfare
Ecological relations	Habitat change, interactions, trophic effects	Attribution can require long-horizon comparisons

Δ_B does not automatically quantify suffering, cumulative habitat disruption, or population consequences. These require additional investigation.

A research design should consequently record the robot’s trajectory and operating modes, illumination and acoustic emissions where relevant, the presence and responses of animals, environmental covariates, observation geometry, detection limitations, and the reference conditions supporting any estimate. When a mission’s scientific purpose permits, noncontact or remote sensing may reduce particular exposures; the empirical effects of each alternative remain to be tested. The later engineering sections will examine these measurements as requirements for sensing and control rather than presuming that autonomous observation is intrinsically noninvasive.

4.5 Individual welfare, ecological continuity, and the limits of aggregation

The welfare of an individual animal and the persistence of an ecosystem are related but distinct objects of evaluation. An intervention might improve a particular population-level indicator while producing adverse experiences for certain individuals. Conversely, avoiding disturbance to an individual in one encounter may leave larger habitat pressures unchanged. A research programme should specify the level of biological organization to which each claim applies: individual, group, population, community, or habitat. The ethical significance of effects across these levels requires argument and context-specific assessment. Measures of biodiversity, productivity, or scientific information gain do not constitute comprehensive measures of welfare.

The Five Domains Model offers an established structured approach for examining nutrition, physical environment, health, behavioral interactions, and mental state in animal-welfare assessment [107]. Its specific use in an oceanic research mission would require species-relevant indicators and validation; a framework developed through animal-welfare science should not be mechanically applied to unknown deep-sea organisms. Nevertheless, the distinction between measurable external conditions, plausible effects on living states, and uncertain subjective interpretation is instructive. Scientific observation can support structured concern without pretending to measure all relevant experiences directly.

The possibility of unknown or vulnerable life also raises problems of representation. An unobserved organism cannot directly appear in the evidential record supporting a proposed mission. For nonanimal marine life, including microbial communities, algae, and habitat-forming associations, the appropriate ethical object often concerns ecological integrity and relational continuity rather than a claim of individual animal sentience. A coral holobiont, for example, can be central to habitat formation and symbiotic relations without warranting a human-like attribution of its experiences. Scientific and normative consideration may be justified through these ecological dependencies even where no individual mental-state assessment is applicable. Different protective grounds should therefore remain distinguishable rather than being assembled into a single presumed category of nonhuman welfare.

Species with unusual behavior may be undersampled, while animals most able to avoid a robot may

disappear from the images on which impact judgments rely. A precautionary assessment should thus examine who or what might be systematically absent from its evidence, rather than simply concluding that an unrecorded effect is zero. Such an assessment benefits from comparative instruments, explicitly bounded confidence claims, and a record of residual ignorance.

Ecological continuity also includes material conditions that cannot be restored through an accounting transfer alone. Disturbed spawning habitat, altered acoustic environments, and physical damage to a long-lived benthic association involve different processes of regeneration and loss. The finitude analysis of Section 3 therefore constrains the interpretation of any claim that scientific, commercial, or public value gained through marine exploration compensates for ecological harm. Compensation, mitigation, restoration, and nonintervention refer to different practical and normative responses. Their adequacy depends on the relevant organism, landscape, institutional responsibility, and achievable ecological outcomes.

4.6 Precaution, refusal, and epistemically bounded admissibility

An ethically defensible research procedure need not wait for complete scientific agreement about every possible animal experience. It should, however, make the grounds for intervention inspectable. The strength of available evidence, seriousness of plausible harm, reversibility of exposure, scientific necessity, feasible less disturbing alternatives, and legal requirements can jointly inform the choice of observational method. These considerations must be separated from a claim that any one ethical rule is already universally binding. Species protections and relevant research requirements vary across jurisdictions and types of activity; Section 8 will examine their legal scope.

A useful formal distinction separates *epistemic compatibility* from *operational admissibility*. Let $\mathcal{B}(y)$ be a set of ecological configurations consistent with observation y and declared uncertainty assumptions. For a candidate action a , let $R_j(x, a)$ denote a specified adverse-outcome indicator for dimension j , and let b_j be a contextually justified limit. A proposed precautionary control condition may take the form

$$\mathcal{A}_{\text{prec}}(y) = \left\{ a \in \mathcal{A}_{\text{reviewed}} : \sup_{x \in \mathcal{B}(y)} R_j(x, a) \leq b_j \text{ for each defined } j \right\}. \quad (14)$$

The condition is meaningful only when the uncertainty set, adverse indicators, and operational limits are justified and the actions reviewed by appropriate authorities. It does not convert subjective welfare into a fully known numeric function; for some taxa, no validated R_j may be available. In such cases the absence of a quantitative indicator must not be mistaken for evidence of negligible harm. A mission may need qualitative restrictions, conservative avoidance distances, additional monitoring, independent review, or a decision not to intervene.

The proposed admissibility set can be empty, uncertain, or impossible to evaluate from available data. This possibility matters for autonomous control. A controller able to optimize locomotion within a local dynamical model might still lack grounds to initiate close contact, collect an animal, introduce sustained lighting, or approach a breeding site. The appropriate supervisory response may be to reduce exposure, gather noninvasive evidence where authorized, refer the decision to human review, or terminate the intervention. Refusal and abstention are operational outcomes to be recorded, rather than optimization failures to be hidden from the scientific account.

There is also a methodological limit to invoking precaution. If every uncertain effect were interpreted as a categorical prohibition, some noninvasive observations necessary to identify serious ecological risks could themselves become inaccessible. Conversely, treating all uncertainty as a reason to proceed would place the burden of unmeasured harm on affected organisms. A constructive procedural inquiry should compare alternative routes to knowledge, the proportionality of the proposed exposure, the feasibility of reversibility, and the public reasons supporting a decision. In Section 5, these criteria will be situated within the larger questions of procedural legitimacy and substantive ecological justice.

4.7 Scientific, legal, and institutional asymmetries

Updated comparative evidence and unresolved taxa. The 2021 expert synthesis of cephalopod and decapod sentience proposes precautionary treatment of these groups in a specified legislative context [20]. A more recent eight-criterion assessment differentiates cephalopod groups rather than transferring the evidence uniformly: it reports strong evidence in octopuses and cuttlefish, substantial evidence in squid, and particularly limited evidence for nautilus [164]. These comparative conclusions concern the strength of evidence for sentience at the taxon level; they do not disclose the experience of an individual animal during a particular robotic encounter. A specific experiment with octopuses reported conditioned place avoidance and neurophysiological responses to noxious treatment [29]. The updated review also discusses limitations in control conditions for the inference concerning the rewarding character of pain relief. This strengthens the case for describing precise endpoints and controls while preserving uncertainty, rather than treating one behavioral result as sufficient for all welfare judgments. Such evidence can inform precaution and independent review, but encounter-specific welfare inference remains dependent on the animal, exposure, observed indicators and unobserved alternatives.

The asymmetry of nonhuman encounters extends beyond laboratory epistemology. Scientists can establish research questions, control the instruments, interpret the records, own derived intellectual property, and decide which findings to circulate. Marine organisms and ecosystems usually cannot participate in such human processes through established legal channels. This does not mean their interests can be determined unilaterally by invoking scientific expertise. It motivates reasoned representation, independent monitoring, challenge procedures, and safeguards against institutional conflicts of interest.

Existing legal arrangements show both substantive developments and variations in scope. EU Directive 2010/63/EU includes living cephalopods within specified protections for animals used for scientific purposes [41]. The United Kingdom's Animal Welfare (Sentience) Act 2022 explicitly includes vertebrates, cephalopod molluscs, and decapod crustaceans in its definition of animals for the purposes of that Act [187]. Neither instrument establishes a worldwide legal standard for all oceanic robotic encounters. The particular application to field observation, animal experimentation, research institutions, and foreign marine environments depends on the relevant legal provisions, jurisdiction, and operational facts. Scientific ethics may also recommend safeguards exceeding minimum applicable legal requirements.

A further asymmetry concerns the creation and distribution of knowledge. An international research consortium might gain commercially valuable biological information through a robotic survey, while local scientific institutions bear environmental monitoring costs or lack access to the instruments needed to evaluate the claims independently. Such outcomes raise questions about epistemic participation, value circulation, and institutional authority in addition to individual welfare. Sections 7–10 will develop the marine nomos, common knowledge, and heterogeneous value-circulation questions more fully; the present section establishes why nonhuman interests and the limits of their representation must remain visible within those processes.

4.8 Research implications and transition to ecological justice

The analysis yields several requirements for subsequent inquiry. Claims about marine animal agency should state their behavioral and mechanistic scope. Claims about sentience or welfare should be supported by species-relevant evidence, preserve reasonable alternatives, and disclose uncertainty. Robotic observations should be evaluated for their capacity to alter animal behavior and the resulting scientific record. Individual welfare, population trajectories, and habitat condition should be reported as distinguishable outcomes. Finally, scientific and institutional decision processes should record the assumptions, permissible actions, independent review, and conditions for changing or withdrawing a mission.

These requirements make encounters with nonhuman life a central part of the scientific problem, rather than an external restriction added after an otherwise completed technology design. The welfare-sensitive evaluation of morphology, sensors, locomotion and autonomous control will be developed later through marine robotic engineering. At the philosophical level, a more difficult question remains: how should a normative procedure account for interests that cannot be communicated through its dominant

language, for ecological harms that may be difficult to reverse, and for unequal human capacities to define what counts as relevant evidence? Section 5 approaches these issues through proceduralism and ecological justice, preserving the distinction between legitimate procedures and ethically defensible outcomes.

5 Proceduralism and Ecological Justice

5.1 The normative problem of ecological inquiry

This section investigates how procedures for marine scientific inquiry can be justified when the activities under review affect heterogeneous human communities, nonhuman organisms, and ecological relations that are only partially understood. Its immediate objective is to distinguish procedural legitimacy, substantive ecological consequences, and the distribution of epistemic and material capacities. Building on Sections 2–4, it develops an exploratory framework for ecological justice and its implications for the formation of marine *nomos*. The applicable law of specific maritime zones and the authorization of particular marine operations will be examined in Sections 7–10; here the question concerns what a responsible normative procedure should be capable of recognizing, testing, contesting, and revising.

Suppose a research consortium proposes to deploy autonomous vehicles near a region supporting feeding, reproduction, and migration. Its engineers can estimate propulsion efficiency and nominal path-following performance. Ecologists can measure selected physical variables and record responses by organisms within the instrumental field of view. A coastal community may describe the area through fisheries, customary knowledge, or long-standing obligations; another institution may emphasize scientific discovery or conservation. The proposed operation may create useful scientific records while producing acoustic exposure, altered animal movement, material disturbance, and longer-lived obligations of monitoring. The question whether the mission is permissible cannot be determined by integrating expected measurements alone. It also concerns what is represented as an affected interest, whose evidence receives attention, who may decide, what uncertainties justify caution, and which remedies remain possible.

Proceduralism is useful because the relevant evidence and values are not assumed to be exhaustively specified in advance. However, a formally regular decision can have seriously harmful consequences. An authorized mission may be supported by inaccurate ecological baselines, omit an affected group, or induce biological changes that were not predicted. A procedure may also be nominally participatory while only well-funded participants can provide the measurements, legal expertise, or computing resources that influence a decision. Consequently, the present inquiry treats procedure as an accountable means of forming, testing, and revising normative relations, while retaining independently examinable questions of animal welfare, environmental protection, distribution, and remedy.

Established environmental-justice scholarship supplies relevant distinctions. Schlosberg’s multidimensional account emphasizes distribution, recognition, participation, and capabilities, including the possibility of extending ecological justice beyond exclusively human claims [163]. The IPBES values assessment documents the relevance of plural conceptualizations and valuations of nature to environmental decisions [69]. Neither supplies a single decision algorithm for marine robotics. Rather, they motivate an inquiry that records disagreements over the characterization of value alongside disagreements over evidence, institutional authority, and acceptable action.

5.2 Relational proceduralism and the formation of normative orders

A generative-relational account begins from the observation that the participants, claims, evidence, and rules involved in ecological decisions are situated within developing relations. A scientific description of an organism, a legal category governing protected species, and a community practice regarding a marine place may each have a distinct represented history. Their interaction can modify the meanings and practical consequences of later decisions. This does not entail that all forms of authority are equivalent

or that legal validity can be inferred from consensus. It means that identifying the applicable decision procedure requires an account of its institutional provenance, scope, participants, and relation to other normative orders.

For present purposes, *normative procedure* denotes a specified arrangement of roles, evidence standards, deliberative steps, admissible actions, decision authority, objection, monitoring, and remedy. A *procedurally justified decision* is one for which the applicable procedural requirements have been met with adequate reasons and an accessible record, subject to the limits of those requirements. *Substantive ecological justification* additionally concerns the consequences, interests, and constraints relevant to the action. *Legal authorization* denotes permission arising under an actually applicable body of law, through the competent authority and within its scope. A research ethics approval, a public consultation result, an operator's internal checklist, and a permit issued by a coastal state should therefore remain distinguishable records.

This analysis is proceduralist in a qualified sense. It places weight on how claims and decisions are generated under uncertainty, yet it does not infer ethical correctness from their having passed through a formal sequence. A procedure needs epistemic capacities sufficient for the question it addresses. It also needs ways to revise the assumptions by which relevant participants, observations, and effects are identified. A procedure that never revisits who is affected may reproduce its original omissions with increasing administrative efficiency.

The wider meaning of *nomos* in the research programme concerns the formation and transformation of normative ordering through legal institutions, scientific practices, customs, and material relations. *Nomos* is analytically broader than positive law, while the binding force of particular legal obligations remains a question of recognized sources, applicable jurisdiction, and interpretation. Sections 7 and 8 examine these legal distinctions explicitly. The concern of the present section is how a situated decision can remain open to challenge while providing the stability needed for responsible action. Adaptive governance research likewise emphasizes social-ecological learning and institutional arrangements capable of responding to changing environmental circumstances [44]. Such responsiveness must be assessed against continuity, accountability, and the risk of concentrating discretionary authority in those operating the technology.

5.3 Recognition, representation, and the limits of speaking for marine life

Recognition concerns whose existence, relations, capacities, and relevant interests become visible to a decision process. In marine scientific research, this includes human research participants, coastal communities, workers, institutions, and potentially affected people beyond an immediate study area. It also raises the more difficult question of how nonhuman organisms and ecological dependencies are to be represented without assigning them human preferences or treating an expert's interpretation as unchallengeable.

The epistemic distinctions developed in Section 4 remain essential. A measured avoidance response may support a hypothesis of disturbance without establishing the animal's complete subjective experience. Absence of a recorded response may reflect limited sensor sensitivity, inaccessible environmental context, or the organism's response outside the observation interval. A scientific assessment should therefore preserve the distinctions among behavioral observation, physiological inference, welfare interpretation, ecological consequence, and normative judgment. The ethical implication of uncertainty depends on the action and its plausible consequences rather than on an assumption that undetected interests are negligible.

Representation of nonhuman interests could involve species-relevant welfare evidence, population and habitat assessment, independent ecological review, monitoring of unintended effects, and explicit consideration of uncertain but consequential harms. Each technique has limits. An individual animal's welfare may diverge from a population-level conservation objective; maintaining a protected habitat does not automatically guarantee the welfare of every organism within it. Nor can one human representative reliably stand for every species or ecological relation. A defensible procedure should therefore identify

which claim is made on behalf of which interest, which evidence supports that claim, where competing interpretations exist, and how its assumptions can be challenged.

Representation among human participants involves comparable problems of epistemic power, although the legal and moral standing of human communities and nonhuman organisms should not be conflated. Consultation conducted exclusively in the dominant scientific language may exclude people whose knowledge is expressed through other observational practices or institutional vocabularies. Translating such knowledge into a standard data field may omit relational conditions of its interpretation. Generalized translation, considered further in Section 6, would require task-specific negotiation of meaning, provenance, and acceptable use; it cannot assume that all claims are fully convertible or transferable. Representation should accordingly preserve disagreement as information relevant to responsible decisions, rather than erase it to produce an apparent consensus.

5.4 Procedural participation, authority, and epistemic capacity

Meaningful participation requires more than a formal opportunity to submit comments. Its practical conditions may include timely notice, intelligible documentation, access to underlying evidence where lawful, scientific assistance, sufficient time, accessible forums, and procedures by which submitted information can affect outcomes. Unequal control of instruments, laboratories, computational infrastructure, modeling assumptions, and professional credentials can prevent nominally open procedures from supporting effective participation. Procedural design should therefore consider the capacity to challenge a measurement, commission an alternative interpretation, request additional monitoring, and obtain a reasoned response.

A relevant established precedent is Principle 10 of the 1992 Rio Declaration, which identifies public access to environmental information, participation in decision-making, and access to proceedings including redress and remedy [197]. The Aarhus Convention develops legally binding access rights within the scope of its parties and provisions, including environmental information, public participation, and review [198]. These instruments have distinct legal characters and territorial and institutional scopes. Their existence supports careful comparison of procedural architectures, not an inference that every marine mission is universally governed by a single consultation code.

Authority remains separate from participation. A research group may solicit public views, while a coastal state retains defined jurisdiction over the activity. Conversely, possession of legal jurisdiction does not alone answer whether affected communities have been meaningfully heard, whether evidence is adequate, or whether a scientifically significant risk has been properly considered. An independent reviewer may assess the adequacy of the environmental evidence without acquiring the authority to issue a permit. These divisions of function should be made explicit in records of scientific and institutional decision-making.

There are difficult cases where participation and ecological protection appear to pull in different directions. Disclosure of precise location data for a vulnerable species might improve independent scrutiny while facilitating disturbance, extraction, or poaching. Procedural design may therefore require restricted data access, trusted review channels, or publication of uncertainty-preserving summaries, with reasons for any restriction and possibilities for challenge. The relevant question is not simply maximal publicity. It is how to sustain reviewable decisions while protecting legitimate confidentiality, community interests, and vulnerable life.

5.5 Justice, heterogeneous value, and generative capacities

The creation and circulation of value through oceanic research present a further problem. Marine exploration may generate new scientific models, biological knowledge, software, robotic equipment, intellectual-property claims, training opportunities, and commercial products. At the same time, it consumes energy and materials, may disturb ecological relations, and depends on institutions that allocate access to scientific and productive capacity. A monetary account is important within its own

conventions, but it does not necessarily encode welfare changes, unpriced ecological dependencies, or the ability of other participants to conduct subsequent inquiry.

Consider a hypothetical installation supplied with equipment from Country A, algorithms from Country B, and operating labor from Country C, located in a maritime space subject to a further set of applicable institutions. Computational services could be sold globally. The distribution of profit, tax revenue, employment, technological expertise, data access, heat discharge, habitat impacts, and future research capacities might follow different relational paths. A legitimate financial transaction between firms does not establish that all ecological or public benefits and costs have been appropriately recognized. Similarly, a payment characterized as compensation may be incapable of reproducing an altered habitat or restoring a lost form of ecological relationship.

We therefore distinguish distributive outcomes from the distribution of *generative capacities*: access to instruments, educational resources, computation, methods, scientific networks, and the possibility of creating further knowledge or productive infrastructure. The distinctive concern here is the distribution of capacities to generate further value and knowledge across relational systems. An equitable allocation of revenue, when required or agreed, need not be sufficient to secure independence of scientific participation. Conversely, shared access to equipment cannot by itself repair animal-welfare injuries or guarantee fair conditions of labor.

At the level of evaluation, it is useful to record separate questions: who contributes to producing the value, who possesses the power to define and recognize it, who can control or restrict its circulation, who receives benefits, and who bears risks or loses capacities. These questions can be examined under different ownership and institutional arrangements. A commons-based scientific infrastructure and a privately financed consortium may each support useful research, yet either can reproduce exclusions if its decision procedures and material dependencies are inadequately examined. The paper accordingly avoids assuming that a chosen ownership category automatically determines justice. Specific institutional and legal arrangements are developed in Sections 8–10 and compared with engineering configurations in Sections 23–25.

5.6 Normative standing, representation and institutional competence

An important difficulty in developing procedural ecological justice is the coexistence of several distinct questions about standing. A marine organism may have ethically relevant interests, while a scientist has evidential standing to describe observations, an Indigenous or coastal community may possess legally protected procedural interests in a particular context, and a regulator may hold legally conferred decision-making competence. These positions cannot be exchanged. Biological vulnerability does not establish that one researcher may legally represent a community; the legal standing of a human participant does not establish the organism's precise welfare state. A constructive framework should therefore separately record an interest or concern, the evidence bearing upon it, the person or institution proposing to represent it, the source of any legal procedural entitlement, and the authority responsible for a decision.

The proposal has practical consequences for consultation. An invitation to comment does not itself demonstrate that access was timely, intelligible, adequately resourced or responsive to submitted evidence. A formal opportunity to participate also differs from a right to consent or withhold consent; the applicable source must be identified for the actual activity, right-holder and jurisdiction. Conversely, consultation requirements should not be invented from a general ethical preference and presented as positive law. The Aarhus Convention, for example, establishes specified rights of environmental information, public participation and access to justice within its parties' legal framework; it does not create a universal global permit procedure for oceanic robotics [198].

Where marine animals are affected, their interests can be represented provisionally through discipline-appropriate ecological and welfare evidence and independent scrutiny. The resulting representation remains contestable: evidence of altered behaviour may support an exposure or encounter finding without uniquely identifying pain, suffering or population decline. The ethical case for restraint or further investigation can nonetheless remain substantial under serious uncertainty. This is an argument for reason-

giving, independent scientific challenge and proportionate preventive action, not a claim that one technical indicator authoritatively settles interspecies justice.

5.7 An exploratory model of accountable ecological decisions

A modest formalization can clarify the distinct functions of scientific evidence, institutional permission, and ethical admissibility without implying that justice has become mechanically computable. Consider a particular proposed operation a , a represented evidential history h , an institutionally defined decision context j , and a set $\Omega(h)$ of ecological configurations compatible with the available evidence. These objects are task-relative representations. They do not exhaust ontological reality or establish a privileged ordering of every event.

Define a decision dossier

$$\mathcal{D}(a; h, j) = (E(a, h), N(a, j), W(a, h), P(a, h, j), R(a, h, j)), \quad (15)$$

where E records scientific evidence and uncertainty, N the relevant normative and authority claims, W species- and context-specific welfare or ecological concerns, P procedural participation and review, and R monitoring, revision, and remedy arrangements. Each component may be incomplete or disputed. The dossier is an accountable organization of reasons rather than an evaluator of moral truth.

For an operation and context, let $\mathcal{A}_{\text{legal}}(j)$ denote actions affirmatively authorized by relevant legal processes, $\mathcal{A}_{\text{technical}}(h)$ actions meeting specified engineering feasibility conditions, and $\mathcal{A}_{\text{ecological}}(\omega)$ those satisfying specified ecological constraints when the actual configuration is ω . A conservative screening set may be represented as

$$\mathcal{A}_{\text{screen}}(h, j) = \mathcal{A}_{\text{legal}}(j) \cap \mathcal{A}_{\text{technical}}(h) \cap \bigcap_{\omega \in \Omega(h)} \mathcal{A}_{\text{ecological}}(\omega). \quad (16)$$

This intersection expresses an explicitly chosen robust-screening policy, not a universal legal or ethical rule. The evidence-compatible set may be misspecified, incomplete, or so wide that it excludes useful low-disturbance observations. Legal authorization also needs to be determined from the competent authority and applicable instruments, not inferred from the notation. A different scientifically and legally justified decision rule might use graded uncertainty, bounded exposure, independent supervision, or additional observation rather than this strict intersection.

The limits of procedural evidence can be stated precisely. Let $P_j(\omega, a)$ be a *procedural record* produced in configuration ω under a specified review arrangement, and let $q(\omega, a)$ describe an ecological or welfare outcome of interest. If two possible configurations ω_1, ω_2 yield the same procedural record but satisfy $q(\omega_1, a) \neq q(\omega_2, a)$, then no function of that record alone can exactly recover q for both configurations. This is an immediate consequence of the identical-input condition. It concerns the informational limits of the selected record; it does not demonstrate that the procedure was unlawfully applied, that an unethical outcome necessarily occurred, or that scientific investigation cannot improve the record. The practical implication is that recording procedural compliance cannot replace independently specified ecological monitoring and opportunities for correction.

The two levels of assessment should therefore remain distinct. A decision can satisfy a defined procedural standard, while its projected ecological consequences remain uncertain or disputed. Later evidence may change a prediction or expose a previously unrepresented affected interest. A responsible institution needs a way to register that change, reconsider the authorized scope where legally available, and preserve the reasons for continuing, modifying, or terminating the activity. Neither the moment of authorization nor a mathematical model of local viability should be treated as an irreversible epistemic closure.

5.8 Precaution, disagreement, refusal, and remedy

Precaution becomes important when an activity could affect animals or ecological relations for which reliable welfare indicators or restoration techniques are limited. The 1992 Rio Declaration articulates

a precautionary approach in Principle 15, while its formulation concerns environmental protection under the conditions and institutional contexts in which that principle is applied [197]. A precautionary procedure needs an explicit account of potential harm, evidence quality, feasible alternatives, exposure, monitoring, and the conditions for reconsideration. It need not presume that every unknown effect warrants permanent prohibition or that every missing measurement licenses intervention.

A scientifically useful sequence may involve simulation, previously collected evidence, non-contact observation, supervised field testing, and more intrusive operations subject to separately justified review. Progression between stages is conditional on the relevant findings and permissions. A system that can recognize when it lacks evidence to act should have operational options besides confident continuation: deferral, lower-disturbance measurement, a request for human review, withdrawal, or termination. These are decisions to be specified by competent institutions and operational safety procedures, rather than assumed outputs of a universal objective function.

Disagreement should be preserved at several levels. Scientists may contest calibration, model structure, or causal inference. Participants may agree on data while valuing an ecological relation differently. A legally authorized activity may be subject to ethical criticism, and an apparently low-risk activity may remain unauthorized. Treating all of these disagreements as one scalar level of “confidence” would hide important distinctions. A decision record should indicate which propositions are contested, by whom or by what review body, under which standards, and with what possible consequences for the operation.

Remedy presents additional limits. Financial compensation, ecological restoration, medical or veterinary response, public correction of a misleading finding, and institutional reform address different losses. A harmed animal cannot necessarily be restored through a financial transfer, and restoring habitat structure does not guarantee recreation of its previous ecological relationships. Some harms may be irreversible within the observational or management horizon. The possibility of remedy is relevant to authorization, yet absence of a complete remedy does not by itself identify one uniquely appropriate policy. Explicit proportionality arguments, alternatives, and enforceable conditions are required for the particular case.

Environmental duties also appear in existing law. UNCLOS Article 192 establishes a general obligation to protect and preserve the marine environment; Articles 204–206 address monitoring, reporting, and assessment subject to their formulations and triggers [195]. The BBNJ Agreement, in force since 17 January 2026, contains additional provisions for its parties within its defined geographic and substantive scope [188]. These instruments should be interpreted in Sections 7–8 alongside relevant domestic law, specific jurisdictions, and institutional competences. They do not convert the exploratory decision dossier of Equation (15) into an internationally binding licensing test.

5.9 Procedural failure cases and the research agenda

Several hypothetical failures can test the proposed distinctions. In a *consultation without capacity* case, all participants receive a report, but only the project operator can afford access to the model and raw measurement records; the procedural opportunity exists while epistemic dependence remains severe. In a *silent-welfare* case, the robotic mission records no behavioral disturbance, but the sensor configuration cannot detect changes outside its visual field; the reported absence of response overstates the evidence. In a *compliant-externalization* case, one installation satisfies local technical requirements while a collection of similarly authorized projects creates cumulative ecological pressure that the individual reviews failed to examine. In a *translation-loss* case, a community’s account of relational obligations is reduced to a monetary valuation without retaining the qualifications attached to its use. Each scenario tests a different failure of recognition, information, aggregation, or authority.

Other failures concern the governance system itself. A precautionary review can become a means of unjustified exclusion if participation is available only to incumbents or if alternative evidence cannot enter the procedure. A continuously adaptive monitoring algorithm may reproduce its developers’ categories, falsely implying that conditions outside those categories have been assessed. A nominal remedy process may be inaccessible, too expensive, too slow, or unable to influence ongoing operations. Evalu-

ating procedures consequently requires attention to actual institutional practices, not merely the existence of published protocols.

The proposed research programme would compare procedures using separate observables: the completeness and contestability of evidence, the capacity of differently situated participants to engage, the coverage of species-relevant welfare questions, the traceability of normative authority, the sensitivity to cumulative and cross-border effects, and the accessibility and consequences of objections and remedies. These are assessable dimensions, not an exhaustive universal index of justice. Their empirical evaluation would require legal expertise, ecological field evidence, institutional study, and participation by the relevant communities. Synthetic decision cases could expose logical omissions but could not independently certify ethical legitimacy.

In summary, ecological justice in marine inquiry concerns the formation and exercise of accountable relations among those who can initiate observation, those who can interpret it, those who can authorize intervention, and the humans and nonhumans affected by it. Proceduralism supplies methods for examining how decisions are generated, while substantive welfare, ecological continuity, and the circulation of productive capacities retain independent ethical significance. The next section investigates how ecological knowledge is formed and translated across heterogeneous relational systems; Sections 7–10 then examine the nomos and institutional arrangements through which scientific inquiry and oceanic production are actually governed.

6 Ecological Knowledge Formation Across Relational Systems

6.1 The problem of heterogeneous ecological inquiry

This section investigates how ecological knowledge may be generated, examined, and circulated among systems whose observational capacities, inferential practices, conceptual distinctions, and normative commitments differ. Its objective is to connect the relational account of marine reality in Section 2 with a constructive account of cooperative inquiry. A marine robotic mission rarely produces knowledge through one instrument, one discipline, or one institution. A sensing platform may register an acoustic signal; a laboratory may infer a molecular signature from a water sample; an ecological model may estimate a population-level process; and a coastal community may describe animal behavior through a history of repeated encounters. Each contribution can be informative under its own conditions, while the relationships among them require investigation. The existence of a shared file, coordinate reference, or biological name cannot settle the meaning or justification of a joint claim.

Heterogeneity occurs at several levels. Instruments sample different physical processes and produce different error structures. Scientific communities employ distinct methods of classification, model construction, validation, and reporting. Institutions differ in the authority under which data are collected and interpreted. Communities may associate a marine habitat with livelihood, cultural continuity, ecological stewardship, or other relations that a conventional measurement inventory does not represent. These differences do not imply that every interpretation is equally well supported. They require investigators to specify which differences are empirical, conceptual, procedural, or normative, and what forms of comparison are possible. The epistemic plurality considered here concerns the conditions of inquiry, not an assertion that separate observers necessarily inhabit physically disjoint oceans.

The distinction among ontological, relational, and manifestational registers from Section 2.2 is useful in this setting. An oceanic process may continue independently of the observations available to particular participants. Its interactions and constraints may be only partially represented. Instrument outputs, models, claims, and local descriptions articulate aspects of that process in different investigative contexts. Cross-system agreement over an observable feature can offer evidence for a shared claim; agreement is not guaranteed by a common vocabulary, and disagreement need not establish the existence of different underlying physical laws. Conversely, a common sensor reading can support incompatible causal or normative interpretations. In a generative-relational (GR) inquiry, the problem is to determine which relations support particular claims and how the investigative practices themselves can change in response

to an encounter.

The practical literature already supplies important precedents. The FAIR principles address the findability, accessibility, interoperability, and reusability of research objects, including methods and workflows [214]. The Ocean Biodiversity Information System (OBIS) aggregates marine biodiversity records and applies data-quality procedures, while the Ocean Best Practices System provides a repository for documented ocean methods [67, 141]. These efforts substantially improve the conditions for comparative inquiry. Their existence also helps specify what remains at issue: interoperability of records and methods is valuable, while claims about a population, ecological mechanism, or legitimate use still require warrant within the relevant research and institutional context.

6.2 Knowledge-producing systems and represented histories

A *knowledge-producing context* will designate a delimited arrangement of observations, representational resources, inferential procedures, material instruments, and institutional practices relevant to a particular inquiry. A context might be a robotic observatory, a laboratory consortium, a fisheries research programme, or a bounded set of community-based practices. It is not presumed to be a self-contained subject or a complete account of its participants. The designation is functional and revisable; a single organization can participate in several contexts, while one investigation can cross many institutional boundaries. The term *generalized language* refers provisionally to a context's distinctions, admissible transformations, inferential relations, and practical capacities for producing further representations. The term covers more than natural-language expression and is not identified with an organism's putative subjective experience.

Consider a represented investigative history $H_i = (E_i, \preceq_i, \Pi_i)$ for context i . Here E_i is a specified set of recorded events, $\preceq_i$ a partial dependency or precedence relation supported by the representation, and Π_i a collection of provenance and interpretation links. This is an analytical record of available evidence. It need not be a complete history of the ecological process, a global temporal order, or an ontologically primitive spacetime. For example, a molecular observation may depend on collection, preservation, extraction, amplification, sequencing, and taxonomic comparison. A machine-vision identification may depend on illumination, camera geometry, training data, annotation rules, and a confidence model. Both can refer to a candidate taxon, although the relevant dependencies and failure conditions differ.

Representing the history of a finding matters when it is subsequently translated or reused. A new reference library can revise the taxonomic assignment of a stored sequence. A calibration correction can change the interpretation of previously registered oxygen measurements. A different population model can reinterpret an acoustic record without altering the original waveform. The published claim, its underlying observation, and the inferential path connecting them therefore have distinct revision dependencies. Provenance ontologies, including the W3C PROV-O Recommendation, provide established means of representing entities, activities, and agents associated with such records [220]. They offer an implementable layer for lineage and attribution; the scientific adequacy of a causal interpretation or normative permission to reuse a record is a further question.

Cross-context inquiry also requires care concerning ecological time. One system may describe processes through measurement intervals and laboratory timestamps; another through seasonal ecological recurrence, organism life history, or locally recognized environmental transitions. These descriptions can be aligned for a defined task, but alignment does not demonstrate that one clock captures every relevant order of change. The relation $\preceq_i$ records dependencies supported by the model or evidence, leaving questions about temporal comparability, intermittence, and continuation open to empirical investigation. The distinction is consequential for marine robots whose opportunistic sampling, intermittent communications, and deferred analyses generate different partial orders of data acquisition, processing, and interpretation.

6.3 Generalized translation as partial relational mediation

A generalized translation is an explicitly scoped attempt to relate selected aspects of two knowledge-producing contexts so that a specified task can be pursued. It may concern measurement units, scientific terms, inferential warrants, model structures, action constraints, ecological values, or institutional practices. The objective need not be the reproduction of a fully equivalent representation in the destination context. In some cases a useful correspondence can be constructed; in others, the appropriate outcome is a documented incompatibility, a limited analogy, or an unresolved question. Translation can also generate a new mediating procedure or experimental design in which the participating contexts jointly develop distinctions that neither previously used in the same way.

For example, a spatially localized camera observation and an environmental-DNA detection can each inform the investigation of a species. Their outputs require different interpretive procedures. The presence of a DNA marker does not by itself establish that a living organism occupied the exact sampling point when the water was collected: transport, persistence, shedding, contamination, and assay performance may intervene [56]. A translation that converts a molecular detection into an unqualified visual occurrence claim would hide precisely the differences needed for scientific assessment. Conversely, a shared occurrence record may be useful if it retains the evidence type, sampling event, uncertainty, and limits of the identification. The translation is successful for a restricted discovery task while potentially inadequate for individual-animal tracking or causal population inference.

The same logic applies to relations among scientific and situated ecological knowledge. Accounts maintained through repeated practical encounters may preserve ecological distinctions absent from an instrument’s standard variable list. Their interpretation depends on context, custodianship, language, and relevant rights and responsibilities. A scientifically productive encounter can involve negotiating what may be compared, what requires independent treatment, and which forms of circulation are authorized. The CARE Principles for Indigenous Data Governance specify collective benefit, authority to control, responsibility, and ethics, explicitly complementing the data-oriented FAIR framework [22]. They are relevant where Indigenous data and knowledge are concerned; their application must respect the specific people, institutions, and legal conditions involved. Neither the FAIR nor CARE frameworks alone resolves the entire philosophical problem of generalized translation.

A practical mediation record should consequently distinguish at least five objects: the originating claim with its provenance; the task and destination context; the proposed relation or transformation; the identified preservation and loss conditions; and the associated questions of authorization, review, and challenge. The preservation claim itself requires evidence. A mapping of taxonomic identifiers may be adequately documented while a mapping of ecological obligations remains undefined. These outcomes can coexist within the same cooperative project. Their coexistence is preferable to a single declaration that two contexts have become fully interoperable.

6.4 A task-relative formal account of translation and loss

The following formal model is offered as a limited instrument for analyzing inferential interoperability. Let $\mathcal{R}_i$ be a set of representations available in context i , and let $D_{ij} \subseteq \mathcal{R}_i$ be a domain on which a proposed mediation $T_{ij} : D_{ij} \rightarrow \mathcal{R}_j$ is defined. A translation may be partial because the receiving system lacks a corresponding category, admissible measurement, or accepted procedure. Let $q_i : \mathcal{R}_i \rightarrow Q$ be a task-relative quantity or property extracted from context i , where Q is a common assessment space selected *for that task*. A shared Q is an explicit modeling choice and cannot be presumed for every ecological value, experiential relation, or institutional authority.

When Q is equipped with a metric d_Q , an evidentially bounded translation claim can take the form

$$d_Q(q_j(T_{ij}(r)), q_i(r)) \leq \varepsilon_{ij}(r), \quad r \in D_{ij}^q, \quad (17)$$

where D_{ij}^q contains those representations on which the displayed expressions are defined. The value $\varepsilon_{ij}(r)$ describes an assumed or validated error allowance under a specified evaluation procedure. It may be unknown, and the inequality must not be presented as an empirical guarantee when no validation exists.

For nonmetric or normative properties, qualitative preservation, disagreement, and refusal records may be more appropriate than Equation (17).

A composition of two proposed translations introduces a further problem. The destination representation produced by one mediation must lie within the next mediation’s domain, and each comparison must refer to a compatible task. These conditions should be tested rather than assumed.

Proposition 6.1 (Conditional error bound for composed task-relative mediations). *Suppose T_{AB} and T_{BC} are partial mappings, $r \in D_{AB}^q$, and $T_{AB}(r) \in D_{BC}^q$. Suppose the three task representations $q_A(r)$, $q_B(T_{AB}(r))$, and $q_C(T_{BC}(T_{AB}(r)))$ belong to the same metric space (Q, d_Q) and satisfy the two respective bounds in Equation (17). Then the composed representation satisfies*

$$d_Q(q_C(T_{BC}(T_{AB}(r))), q_A(r)) \leq \varepsilon_{BC}(T_{AB}(r)) + \varepsilon_{AB}(r). \quad (18)$$

Proof. Insert $q_B(T_{AB}(r))$ as an intermediate point in the metric space and apply the triangle inequality. The two assumed task-relative bounds yield Equation (18). The domain and comparability assumptions ensure that every term is defined. $\square$

The proposition is elementary and deliberately limited. It demonstrates that separately supported bounded transformations can yield a bounded composition *under a common task and compatible domain*. It establishes neither reversibility nor semantic equivalence. If the first transformation destroys a distinction necessary for the second, if the evaluation spaces differ, or if either error bound lacks support, the result cannot be invoked. Moreover, even a numerically accurate path can be unauthorized or normatively objectionable. Its value lies in making explicit the conditions under which an engineering system may legitimately propagate a preservation claim.

The existence of a shared task variable in Proposition 6.1 is a substantive precondition of comparison. For an image classifier and an environmental-DNA assay, a carefully defined assessment space might contain taxon-specific detection statements, associated detection models, and temporal or spatial uncertainty. Even when both instruments use the same taxon identifier, a mapping to an individual organism’s contemporaneous location may be undefined. Such an absence of a justified mapping should be preserved as a missing inferential relation, rather than repaired by assigning an arbitrary numerical error. The same principle is stronger for welfare interests, community rights and legal authority: their partial mediation cannot be inferred from the existence of interoperable metadata.

A practical comparison should preserve both a semantic version of the claim and the method of its assessment. If a later taxonomic revision changes which biological entities fall under a label, the receiving dataset may remain syntactically readable while its inferential content has changed. If two observing institutions disagree about what an acoustic detection signifies, a joint investigation can compare calibration and rival hypotheses without requiring either institution to surrender the recorded distinction. This offers a concrete test of whether relational mediation improves reproducibility and explanatory discrimination beyond conventional format conversion. A successful result would be more reliable identification of disagreements and their grounds; the absence of such improvement would count against introducing an additional mediation layer.

Different translation paths can also yield different outcomes. A biological identification transferred from context A through B to C may differ from a direct A -to- C analysis because the intermediate representation records different taxonomy versions or uncertainty conventions. This constitutes a testable path-dependence problem, not an automatic demonstration that one path is scientifically superior. A mediation infrastructure should record path identity, software and vocabulary versions, unresolved discrepancies, and evidence supporting the chosen task. The existence of a direct mapping does not authorize the silent replacement of a longer inferential history.

6.5 A marine case: heterogeneous evidence for a shifting ecological assemblage

Consider a hypothetical research consortium investigating changes in the distribution of a coastal fish assemblage near an area of temperature variability. A robotic surface platform collects acoustic observa-

tions and temperature profiles; a submerged sampler produces environmental-DNA records; a laboratory develops taxonomic assignments; a modeling group estimates habitat suitability; and a locally situated research partnership contributes observations of seasonal behavior and habitat use. The example is constructed to expose representational tensions. It makes no claim that a particular observed assemblage has undergone the proposed change.

The acoustic system may detect scattering structures consistent with several organism types. It requires validation to distinguish species, schooling behavior, and environmental interference. Molecular measurements can support taxonomic inferences while remaining sensitive to transport, degradation, sample handling, and reference-database coverage [56]. Habitat models combine observations with specified environmental relationships, and their predictions depend on selected explanatory assumptions. Locally situated records may identify habitat or seasonality distinctions absent from the model, while also having their own spatial, temporal, and interpretive limitations. Each source thus contributes a partially different view of the evolving ecological assemblage.

A merely syntactic integration would place the records in one database and assign common coordinates and taxon labels. An epistemically responsible integration would retain detection modality, uncertainty, measurement conditions, versioned inferential transformations, known spatial transport processes, and the precise ecological proposition each record may support. It would also distinguish “a relevant marker was detected,” “a candidate organism was observed,” and “the modeled population is changing” as non-equivalent assertions. Data-standard infrastructures make parts of this practicable. OBIS organizes marine occurrence information with quality-control procedures and interoperable biodiversity schemas; Darwin Core provides a maintained vocabulary for biodiversity records and related observations [67, 30]. These are technical supports for defensible comparison, not guarantees of population-level causal identification.

A second problem concerns how the integrated result circulates. The partnership might seek to publish a scientifically reproducible dataset while restricting the disclosure of sensitive habitat locations or respecting community-controlled information. A modeling team may wish to release an inferred map whose uncertainty creates risks of overstating a conservation claim. Another actor may seek to use the map to authorize a physical intervention. The resulting tasks call for different reviews. Scientific support for a habitat hypothesis does not create a legal authorization to deploy robots, extract organisms, or build marine infrastructure. The distinction developed in Section 5 between epistemic evidence, institutional authority, procedural legitimacy, and substantive ecological consequences remains operative.

The case suggests a constructive experiment. Participants could evaluate three modes of cooperation using the same appropriately licensed evidence: pooling records through a common schema; supplying records together with modality-specific interpretive and provenance constraints; and developing a negotiated task-specific mediation that also records conflicts and unresolved categories. The comparison would examine inferential accuracy against independent reference observations where feasible, calibration of uncertainty, the rate of unsupported cross-modality inferences, correction traceability, practical research effort, and the ability of affected participants to challenge a representation. No result should be predicted solely from the philosophical preference for richer mediation. Its additional documentation and governance requirements may also introduce expense and complexity that require evaluation.

6.6 Learning another system’s generative representation

Generalized translation raises a further challenge when the rules by which a context produces representations are partly unknown. Conventional transfer may align existing outputs, identifiers, or probabilities. A generative inquiry may instead investigate aspects of the producing system’s distinctions, inferential operations, and experimental practices. An AI model observing scientific workflows could propose a candidate representation of the constraints governing a laboratory’s classifications. A robotic system might discover that recurrent acoustic events require a new observational category. A collaborative group could develop a new protocol linking molecular observations and repeated behavioral surveys. These possibilities concern the generation of hypotheses about another system’s practice; they do not

presuppose complete access to its internal organization.

For a specified hypothesis class $\{G_\theta : \theta \in \Theta\}$, let $\mathcal{H}$ denote the documented interaction history and $\mathcal{U}$ a set of permitted probing activities. An exploratory research procedure may compare candidate representations according to how well they account for observations under those activities. The choice of Θ , observation model, intervention conditions, and evaluation criteria determines what can be learned. Observationally equivalent hypotheses may remain distinct, and some probes may be ethically impermissible or technically inaccessible. Learning one system's predictive regularities does not establish that the system uses the same internal mechanisms as the learned model.

In the marine setting, this issue is especially salient when biological organization inspires robotic control or interpretation. A model of an animal's movements may explain certain observed trajectories without describing its lived experience or reconstructing every ecological relation shaping its behavior. A morphology-inspired controller that achieves useful locomotion similarly demonstrates an engineering capability, not a transparent translation of biological agency. Model revision can improve scientific warrant, but an increase in predictive performance alone does not license claims about sentience, welfare, or the legitimacy of invasive observation. Those claims require separate evidential and normative assessment, as argued in Section 4.

A modest implementation would maintain a register of competing generative hypotheses and associated evidence, identify conditions under which predictions diverge, and propose low-disturbance measurements capable of discriminating among them. Where available actions are insufficient to discriminate, the appropriate output may be a documented unresolved equivalence class and a research question for later investigation. Where the model's own categories fail, investigators may revise the hypothesis class or the measurement protocol itself. These changes can be recorded without asserting a universal progression from inadequate to complete understanding. The work of epistemic translation thus includes the revisability of the languages by which translations are attempted.

6.7 Knowledge circulation, recognition, and epistemic capacity

Cooperative ecological knowledge creation depends on the ability to participate in producing and evaluating evidence. A consortium can provide an open database while concentrating computational resources, instrument access, analytical expertise, or decision authority among a few institutions. Participants may consequently possess nominal access to research products while lacking practical means to develop alternative interpretations or reproduce the research process. The distinction between circulation of established knowledge and circulation of knowledge-producing capacity has ethical as well as technical significance. Ocean best-practice repositories, training networks, and collaborative observing initiatives can contribute to that capacity, but the actual distribution of facilities and authority must be investigated in each institutional arrangement [141, 186].

The conditions of participation also differ across forms of knowledge. Some ecological records involve community-held information, restricted biological locations, or practices for which general publication is contested. Respecting the rights and responsibilities of relevant knowledge holders may require scoped access, review before secondary use, and the possibility to refuse a proposed translation or dissemination. Such arrangements can coexist with substantial open-science commitments when the actual rights, duties, and public-interest considerations are made explicit. CARE and FAIR supply complementary, not interchangeable, design considerations [22, 214]. The appropriate institutional decisions depend on applicable law, affected parties, and the purpose of the inquiry.

Circulation can create additional interpretations and capacities. An acoustic observation may motivate a new molecular sampling strategy; a revised taxonomy may alter the evidential significance of an archived record; a community-based account may suggest a variable omitted from a mechanistic ecological model. These are potentially generative interactions among knowledge-producing contexts. Their value is neither exhausted by a monetary transaction nor guaranteed by maximal dissemination. Some uses may cause ecological harm, amplify asymmetries in research capacity, or undermine legitimate restrictions. The question for procedural justice is how such opportunities are created, evaluated,

distributed, revised, and, where justified, constrained.

The appropriate infrastructure is therefore an accountable federation of partial capabilities: observation and calibration, provenance and versioning, task-specific mediation, uncertain inference, independent scientific challenge, and the identification of actual institutional authority. The federation does not need to absorb every participant into one universal ontology or data model. It does need operational means of indicating when a claim is supported within a certain scope, when translation has changed the object of inference, when a record should be corrected, and when the requested use exceeds the available permission or evidence. A technical system that reports only success or failure of file conversion cannot meet these requirements.

6.8 Limits, counterexamples, and the transition to marine nomos

Several cases test the boundaries of the proposed account. In a *false-equivalence* case, two institutions use the same species label while applying different taxonomic concepts; successful identifier transfer conceals a biological distinction. In an *inference-promotion* case, an environmental-DNA signal is translated into a precise claim about the present location of a living animal. In a *path-conflict* case, two chains of model and vocabulary transformation yield materially different estimates, but only the final outputs are retained. In an *authority-confusion* case, agreement over an environmental model is interpreted as permission for a research or construction operation. In a *capacity-asymmetry* case, a purportedly open dataset cannot be independently analyzed by participants excluded from the computational or methodological infrastructure. These failures are conceptually distinguishable and call for different corrections.

A more difficult counterexample concerns projects in which insisting on exhaustive provenance and mutual interpretation would make an urgent, low-risk observation impracticable. A useful design cannot assume that maximal documentation is always proportionate. It should specify a minimum justified record for each task, the conditions under which provisional judgments may be issued, and whether a subsequently contested decision can be independently reconsidered. Similarly, retaining multiple interpretations can improve epistemic accountability while complicating coordination. The analysis must address costs, latency, scientific usefulness, and effects on participants, rather than presuming that every additional layer of mediation improves outcomes.

The formal results in this and previous sections identify restricted possibilities and limits of inference. They do not demonstrate the existence of one generative grammar shared by ecological systems, establish that every value possesses a faithful translation, or supply a metric by which the interests of different species and institutions can be aggregated. In particular, the normative significance of a marine animal's welfare cannot be recovered from an accurate species-identification pipeline alone. Real-world deployment would require task-specific scientific validation, legal analysis, appropriate participation, and ecological monitoring. Synthetic comparisons can expose failures of inference and record handling; they cannot certify the justice of actual cross-border arrangements.

Claim-level scope across relational translations. The distinction between ecological referent, model object, and observable record must remain explicit when results are translated between communities or instruments. A map from sonar return to object detection has a different evidential domain from an inference about population abundance; a map from field response to a proposed welfare concern has yet another. In each case the domain of the translation, the information discarded, the relevant validation population, and plausible alternative explanations should be inspectable. An explanatory agreement between two models does not entail that either captures every relevant relation in an ecological setting. Conversely, disagreement may reveal differences of resolution, classification or measurement before it establishes a contradiction about the referent. The comparative programme in Section 25 treats these interfaces as independently reviewable hypotheses. Their practical value must be judged against existing scientific provenance and reporting practices rather than inferred from the GR vocabulary alone.

Ecological knowledge formation across relational systems is consequently an open process of producing, testing, translating, and sometimes refusing proposed connections. The operational object is a

revisable network of claims, histories, practices, and responsibilities, not a final universal description. This completes the preliminary philosophical account of ecological inquiry: Sections 2–6 have distinguished ecological reality, generative conditions, nonhuman interests, procedural justice, and cooperative knowledge formation. Sections 7–10 investigate how these relationships are organized through marine nomos, applicable law, institutional authority, and common scientific capacity; the subsequent engineering survey examines what oceanic robotic systems can actually observe and do.

7 Marine Nomos and Normative Plurality

7.1 Normative ordering in an inhabited and technologically mediated ocean

This section investigates how normative orders govern encounters with marine ecological systems when scientific observation, material production, and technological operation cross physical, institutional, and jurisdictional boundaries. Its objective is to move from the procedural and epistemic foundations of Sections 5 and 6 to the institutions and legal distinctions that condition actual oceanic activities. The inquiry uses *marine nomos* as a provisional analytical designation for the historically formed and revisable ordering of marine space, permitted practices, recognized relations, authority, responsibility, and contestation. The term is broader than positive law. It does not imply that customary practices, scientific conventions, contractual rules, moral expectations, and binding international law possess equivalent legal status.

Oceanic research gives this distinction a practical significance. A scientific vessel may navigate a region under one legal freedom while its robotic instruments collect samples subject to another regime. An autonomous surface craft may relay signals from submerged vehicles; a seabed installation may remain in place after a research mission ends; data generated at that installation may subsequently support products and services in several states. The activities share physical and infrastructural dependencies, yet their corresponding permissions and responsibilities do not necessarily travel together. A contract among participating scientists cannot alone establish jurisdiction over a marine site. Conversely, legal competence to authorize an installation does not establish that the installation is scientifically justified, ecologically harmless, or distributively just.

The generative-relational perspective makes several distinctions especially important. Marine ecosystems contain physical and biological relations that may precede or extend beyond their recognition by human institutions. Their scientific manifestation depends partly on observing instruments, measurement models, and interpretive practices. Human normative orders are themselves relationally constituted: an authority exercises powers through recognized offices, applicable legal sources, administrative procedures, and actual capacities of enforcement and review. Analytically, these relations interact, but an environmental state cannot simply be identified with its legal classification. The existence of a protected area, for example, changes the admissible human activities within its legal scope; it does not entail that all ecological risks have been detected or that biological relations have become adequately represented.

The central practical questions are consequently disaggregated. Who may enter, measure, sample, construct, operate, extract, transport, process information, and withdraw? Which institutions can legitimately answer each question, and under what legal sources? Which affected interests and values enter the decision, including those of marine animals and people outside the immediate operational region? How are changes in technical capability, scientific understanding, and ecological condition reflected in the review of permissions? These questions require a plural description of marine nomos with explicit boundaries of authority, rather than a universal rule inferred from the simple fact that a project occurs at sea.

7.2 Maritime zones and the allocation of legal competences

The United Nations Convention on the Law of the Sea (UNCLOS) supplies a principal architecture for allocating rights, jurisdiction, and duties. For the present inquiry its significance lies partly in the differentiation of maritime zones, and partly in the fact that activities within a single physical location can

Table 3: Illustrative maritime-zone distinctions relevant to oceanic research and production. The table indicates questions for legal classification rather than sufficient permissions for any particular project.

Setting	Primary legal distinction	Project question
Territorial sea	Coastal-state sovereignty, subject to international law	Which coastal-state approvals and operational rules govern the mission or installation?
EEZ	Specified coastal-state rights and jurisdiction, including UNCLOS Articles 56 and 60	Is the structure covered; which permissions, duties, and safety requirements apply?
Continental shelf	Seabed rights distinct from the legal status of superjacent waters	Is the activity an installation, cable, drilling operation, or resource activity; what shelf regime applies?
High seas	Qualified freedoms and duties under international law	Which states and legal instruments govern the actors, craft, installation, and its effects?
The Area	Seabed and subsoil beyond national jurisdiction, with a defined mineral-resource regime	Is a proposed activity within the relevant mineral-resource mandate, or does another regime govern it?

engage different categories of legal competence. Article 56 recognizes specified coastal-state sovereign rights and jurisdiction in the exclusive economic zone (EEZ), including concerning installations, marine scientific research, and protection of the marine environment. Article 60 addresses the coastal state's exclusive right to construct and authorize and regulate covered artificial islands, installations, and structures in the EEZ, together with specified jurisdiction over them. These are defined powers within a treaty regime, not a general equivalence between an EEZ and territorial sovereignty [190].

The continental shelf requires separate analysis. Articles 77 and 78 concern coastal-state rights to resources of the shelf and the distinct legal status of the waters above it. Article 80 applies the Article 60 regime *mutatis mutandis* to installations and structures on the continental shelf, while Article 79 addresses submarine cables and pipelines. An installation positioned beyond 200 nautical miles from the coastal baselines may therefore intersect an extended-continental-shelf question while the overlying water is subject to the high-seas regime. Any legal analysis that labels the entire water column and seabed together as "international waters" risks omitting the vertical differentiation of rights [191].

On the high seas, Article 87 recognizes freedoms exercised subject to UNCLOS and other applicable international law. These include qualified freedom to construct artificial islands and other installations permitted under international law, with the continental-shelf qualifications expressly retained. Article 89 prohibits state claims of sovereignty over the high seas. Neither proposition creates a general individual entitlement to install, own, and operate any desired structure without regard to flag-state rules, other states' rights, marine-environment duties, or other applicable regimes [192]. The legal status of installations also depends on their nature, location, function, and relationship to ships, cables, pipelines, and seabed activities.

This matrix should not be applied without first establishing an actual maritime location and relevant state relationships. Archipelagic waters, internal waters, overlapping maritime claims, regional instruments, and the legal position of nonparties may require additional analysis. Determining applicable law is consequently an exercise in legal classification, source interpretation, jurisdictional inquiry, and fact verification. Engineering coordinates are necessary inputs, but alone cannot resolve the applicable normative order.

7.3 Marine scientific research, productive activity, and hybrid infrastructures

Marine scientific research is explicitly addressed in UNCLOS Part XIII. Articles 245 and 246 differentiate coastal-state powers over research in the territorial sea, the EEZ, and continental shelf, including

coastal-state consent in the latter two settings. Article 246 also specifies conditions relating to consent and circumstances in which it may be withheld. Other Part XIII provisions address information supplied to coastal states, participation, communication of results, and research installations and equipment [196]. These provisions should be interpreted by reference to the status and purposes of the actual mission. An instrument that measures ecological conditions does not, by that fact alone, determine whether an entire project is legally classified as marine scientific research.

Consider an autonomous observation platform providing publicly accessible environmental data. Its operations may involve navigation, deployment, biological sampling, remote communications, and temporary or persistent installation. Each component may demand a different legal inquiry. A shift from noninvasive observation to sample collection may implicate additional permissions and biological resource rules. A decision to add a fixed energy supply or a persistent seabed docking structure may raise installation, cable, environmental-assessment, and maritime-safety issues. An AI model used to select observations alters the scientific protocol but does not itself authorize activities that the protocol proposes.

Commercial production adds further distinctions. A subsea data center could deliver computational services around the world while depending on a coastal landing point, surface servicing operations, submarine cables, energy infrastructure, environmental cooling, and decommissioning capability. Its economic purpose cannot simply be re-described as marine scientific research because it records local temperature or supports some scientific computation. Equally, the existence of private investment or commercial customers does not prove that a given type of scientific activity is excluded from research-related provisions. The legal result depends on the actual activity and the interpretation of applicable instruments. This is particularly important for facilities combining publicly funded research, industrial experimentation, environmental monitoring, and commercial computation.

A practical institutional response is to maintain an *activity decomposition*: the physical components of a project, the operations they perform, the scientific or economic purposes asserted, the relevant jurisdictional contacts, and the sources of claimed authority. The decomposition should record uncertainty and unresolved classifications. The output is an input to competent legal and administrative review, not an algorithmically issued permission. Section 8 will consider authorization and accountability workflows in detail; the present section establishes why an undifferentiated classification of the entire installation would be analytically insufficient.

7.4 Areas beyond national jurisdiction and differentiated international institutions

The international seabed *Area*, in the UNCLOS sense, consists of the seabed and ocean floor and subsoil beyond the limits of national jurisdiction. Article 133 defines the relevant resources for Part XI as mineral resources *in situ*; Articles 136 and 137 articulate the common-heritage regime for the Area and its resources [193, 194]. The International Seabed Authority (ISA) is mandated to organize and control mineral-resource-related activities in the Area for the benefit of humankind as a whole, including related environmental protection. This specific competence does not make the ISA a universal licensing authority for all oceanic robotics, computing installations, marine biology, or non-mineral production beyond national jurisdiction [78].

An additional layer entered into force on 17 January 2026: the Agreement under UNCLOS on the Conservation and Sustainable Use of Marine Biological Diversity of Areas beyond National Jurisdiction, generally called the BBNJ Agreement. Its four principal thematic areas are marine genetic resources and benefit-sharing, area-based management tools, environmental impact assessments, and capacity-building and transfer of marine technology. The Agreement also creates institutional arrangements. Its obligations must be examined by reference to its scope, the relevant parties, and its provisions rather than extended to every form of marine economic value [188, 189, 200].

For example, the BBNJ framework can be relevant when an activity under a party's jurisdiction or control has consequences for biodiversity beyond national jurisdiction, subject to the Agreement's screening and assessment rules. It does not thereby create a universal revenue-sharing charge on ordinary

data-center computing services. Nor should its marine genetic resource provisions be treated as an automatic mechanism for allocating every benefit produced through computational analysis. The distinction between a legally defined subject matter and a more expansive philosophical account of heterogeneous value circulation must remain explicit. The former has legal consequences when its requirements apply; the latter supplies research questions about contributions, capacities, ecological burdens, and distribution.

The institutions and their mandates therefore form an overlapping, only partially integrated configuration. UNCLOS, implementing agreements, coastal-state domestic law, competent international institutions, regional environmental arrangements, relevant vessel-related rules, and enforceable contractual obligations can all contribute to governance in specified circumstances. Their interaction should be described through actual jurisdiction and norm-conflict analysis. Describing this as *normative plurality* identifies the existence of multiple potentially relevant orders; it does not excuse deciding which sources govern a disputed operation.

7.5 Ecological obligations, precaution, and differentiated legal effects

UNCLOS Part XII places protection and preservation of the marine environment within the general legal order of the sea. Article 192 states an overarching obligation, while Article 194 concerns measures to prevent, reduce, and control marine pollution. Article 206 addresses assessment, insofar as practicable, where states have reasonable grounds for believing planned activities under their jurisdiction or control may cause substantial pollution or significant and harmful changes to the marine environment [195]. These provisions establish important legal duties, although the precise standard and its application to a proposed operation require interpretation of the facts, relevant obligations, and institutional competence. A scientific simulation predicting negligible impacts does not substitute for an applicable assessment process.

The BBNJ Agreement introduces more detailed mechanisms for relevant activities within its scope, including screening, environmental assessment, public information, decision-making, and monitoring. Its procedures should not be conflated with project approval by a general supranational marine authority. Determination whether an activity may proceed remains connected to the appropriate party's jurisdiction or control under the treaty framework [189]. The binding obligations of an individual state or operator cannot be established merely by observing that the Agreement has entered into force internationally. Party status, implementing law, the specific activity, and any other relevant treaty regimes remain consequential.

Maritime operations may further involve the laws applicable to vessels. The International Maritime Organization describes MARPOL as a principal global convention addressing pollution from ships arising from operational and accidental causes [73]. Its existence should not be interpreted as a comprehensive environmental-licensing code for stationary seabed computing modules. It may instead govern relevant support vessels and other activities within its actual scope, interacting with different duties relating to a fixed installation.

These legal distinctions are important for ecological justice. A marine operator may comply with a narrowly framed noise criterion while producing effects on an organism that the criterion fails to represent. A permit may be legally effective even when subsequent measurements reveal assumptions needing revision. The appropriate research response is to identify routes for monitoring, reconsideration, challenge, or remedy rather than infer that a compliance checklist exhausts ecological responsibility. Duties owed under law, voluntary safeguards adopted by a consortium, and proposed principles of interspecies justice can reinforce one another, but should be documented as distinct normative claims. A scientifically estimated ecological threshold can inform an enforceable condition only through the relevant legal or contractual process; its transformation into an obligation is itself an institutional event.

The evidential thresholds in BBNJ environmental assessment. The distinction between screening and a full assessment warrants particular care. Under Article 30 of the BBNJ Agreement, a planned ac-

tivity within the relevant treaty scope is subject to screening when its effects may be more than minor or transitory or are unknown or poorly understood. That initial screening asks whether there are reasonable grounds to believe that the activity may cause substantial pollution or significant and harmful changes to the marine environment; the resulting determination governs whether the treaty's fuller environmental impact assessment procedure is required. Article 30 lists factors including the technology used, duration, location, site vulnerability, cumulative impacts and the extent of unresolved effects. The presence of uncertainty may therefore be legally material to the *screening question*, while a particular screening outcome should not be described as advance ecological approval [189]. Article 34 locates a decision whether an assessed activity may proceed with the Party exercising jurisdiction or control, within the treaty's conditions. These provisions must be read with the Agreement's rules on geographic and activity scope, relationship with other assessment bodies and existing instruments, party status, and domestic implementation. A mission database should accordingly record separately: a preliminary threshold determination, the actual competent decision-maker, the decision and its conditions, any public documentation, and whether subsequent monitoring warrants review. This reconstruction is an analytical aid for legal research, not a determination that a hypothetical installation requires, has obtained, or is exempt from a particular permit.

7.6 Research consent, environmental assessment and autonomous platforms

Three legal workflows that can intersect in one marine robotic mission require separate evaluation. The first concerns *marine scientific research consent*. UNCLOS Article 245 calls for express coastal-state consent in the territorial sea. Articles 246–249 address research in the EEZ and continental shelf, including coastal-state consent, advance information, coastal-state participation, research outputs and removal of equipment; Article 248 generally requires specified project information not less than six months before the expected start. Article 252 provides a carefully conditioned implied-consent procedure, rather than a universal exemption from permission, while Article 253 addresses suspension or cessation when specified conditions are not met. Articles 258–262 add research-equipment rules concerning deployment, legal status, safety zones, shipping routes and markings [196]. Their applicability depends on the project's classification, location, participating states and relevant facts. In particular, a machine-generated “consent pending” status cannot lawfully be promoted to “consent granted” merely because a local administrative timer has expired.

A second workflow concerns *assessment of environmental effects*. UNCLOS Articles 192, 194, 204 and 206 provide duties concerning protection, pollution prevention, monitoring and assessments within their specified formulations [195]. Under the BBNJ Agreement, for an activity and Party within Part IV's scope, Article 30 differentiates a screening trigger (possible more-than-minor or transitory effects, or poorly understood effects) from the later determination whether reasonable grounds indicate potential substantial pollution or significant harmful change and therefore require a fuller assessment. Articles 31–33 address the assessment process, consultation and published report; Article 34 allocates the decision on whether an activity may proceed to the Party exercising jurisdiction or control, and Articles 35–37 concern monitoring, reporting and review. Article 29 addresses interactions with other assessment regimes. None of these provisions should be read as a single worldwide certification of “oceanic sustainability” [189].

A third workflow concerns *maritime safety and autonomous surface shipping*. In May 2026 the International Maritime Organization adopted the non-mandatory Maritime Autonomous Surface Ships (MASS) Code, effective from July 2026 and focused on the covered SOLAS cargo-ship domain. The IMO specifies that other vessels, including smaller craft, are not automatically within the Code's mandatory scope, although governments are encouraged to apply its guidance where practicable. Submerged AUVs, fixed sensors and autonomous scientific equipment also require their own legal classification. The existence of a remotely controlled craft or an AI system is therefore insufficient to establish that the MASS Code authorizes the operation or replaces applicable collision-prevention, flag-state or coastal-state requirements [72].

Table 4: Distinct legal inquiries in hypothetical marine robotic projects. The classification is preliminary; permissions must be verified under applicable instruments and facts.

Workflow	Relevant evidence and legal source	What it does not establish
Scientific research	Zone and activity classification; UNCLOS Part XIII; coastal-state communications and conditions	General installation or extraction entitlement
Environmental assessment	Jurisdiction or control, effects, screening, assessment and decision; UNCLOS Part XII and applicable BBNJ provisions	Scientific certainty or an autonomous permit decision
Surface-vessel safety	Vessel type, applicable flag and coastal rules, relevant IMO instruments and guidance	Automatic MASS applicability to every robotic platform
Biological resources	Resource, sampling, location, use, party and temporal scope; applicable access/benefit-sharing rules	Universal entitlement to use data or universal profit sharing
Ethical review	Animal evidence, affected interests, proposed safeguards and participatory opportunities	Binding legal authority merely through moral endorsement

The distinction is especially useful where activities change during a mission. A sensor survey may be lawful under conditions that would not extend to collecting organisms, installing an acoustic emitter, or expanding a seabed structure. A review must relate the altered activity to the actual scope of each relevant decision instead of treating the platform's original project name as a legal category. A parallel scientific record should document which factual assumptions changed, which authority has been contacted and which legal consequences remain unresolved.

7.7 Normative plurality, situated knowledge, and asymmetric participation

Normative ordering also emerges through scientific conventions, communal practices, regional forms of cooperation, ethical commitments, and the allocation of material capacities. These forms may govern behavior in practice even when their legal status differs. For example, scientific communities may recognize standards for calibrated observations, animal exposure minimization, or data citation. Communities may maintain knowledge practices specifying when biological locations or sensitive ecological information should be shared. An operator may adopt contractual commitments concerning restoration or access to local monitoring. Some of these commitments are enforceable through applicable law, some establish professional or contractual expectations, and others remain ethical aspirations. Their relative legal force must be investigated rather than presumed.

The distinction between normative pluralism and epistemic relativism is essential. The existence of divergent scientific or cultural valuations does not mean that a predicted ecological impact is scientifically indeterminate in the same way that a rule's interpretation is legally contested. Scientific claims may be assessed through observations, methods, uncertainty, and reproducibility; legal claims must be assessed through recognized sources and procedures; ethical claims require justified consideration of affected interests and consequences. Their interaction may demand translation, but neither mutual intelligibility nor agreement necessarily follows from making a shared database or a common decision form.

This point matters for marine organisms, which ordinarily do not participate in human legal deliberation through linguistic testimony. Their interests may be investigated through behavioral, physiological, population, and habitat evidence without claiming a complete reconstruction of subjective experience. A review body might accept a particular welfare proxy yet continue to recognize the limitations of that proxy. The asymmetry is epistemic and procedural: human institutions define the categories through which affected organisms become legally or scientifically visible. The design of an accountable process should make those representational dependencies open to challenge and revise assessment protocols

when relevant scientific knowledge changes.

The cross-border dimensions are comparably significant. One participating state may possess advanced instrumentation and computation, another may provide operational labor, and a third may experience ecological burdens. Their practical capacity to challenge a scientific or regulatory assessment can diverge substantially even when procedural participation is formally available. Recognition and representation thus depend partly on the circulation of scientific methods, access to independent expertise, legal assistance, and the availability of verifiable monitoring records. This motivates a study of common knowledge-producing capacities as part of marine nomos, while the fuller political-economic argument belongs to Section 9.

7.8 Representing jurisdiction and authorization without collapsing normative orders

A limited formal device can clarify the relationships established above. Let a proposed marine activity be described by a factually supported dossier

$$d = (\ell, \sigma, \alpha, \pi, \omega, \iota, \chi), \quad (19)$$

where ℓ describes the verified spatial setting (including relevant seabed and water-column status), σ the material structures and platforms, α the operations, π their stated purposes, ω the participating actors, ι the potential ecological interactions and impacts, and χ the institutional and cross-border contacts. The dossier is a representation of known facts and uncertainties, not a substitute for an investigation of legal fact or authoritative interpretation.

Let $\mathcal{N}(d)$ denote a collection of *candidate* legal and normative instruments potentially relevant to the dossier. For each candidate source n , four separate questions arise:

$$\mathfrak{J}(n, d) = (\text{app}_n(d), \text{competence}_n(d), \text{conditions}_n(d), \text{review}_n(d)). \quad (20)$$

The first concerns applicability, the second the competence of a purported decision-maker, the third any operative conditions or limitations, and the fourth the available review or challenge mechanisms. Each response may be supported, refuted, or unresolved on the evidence available. A three-way evidential status must not be confused with the legal system's own categories of valid, invalid, or voidable decisions. In particular, uncertainty in a research database does not automatically invalidate an actual legal instrument, just as a database assertion of “supported” cannot issue a legal authorization.

Consider an investor's ownership claim $O(d)$, a positive scientific assessment $S(d)$, and authorization under the applicable law $A_n(d)$. Neither

$$O(d) \implies A_n(d) \quad \text{nor} \quad S(d) \implies A_n(d) \quad (21)$$

can be adopted as a general implication. A party may own equipment but lack permission to install it in a particular maritime zone; a rigorous environmental model may support a proposed research operation that still requires coastal-state consent. The converse also needs attention: legally authorized action may have remaining scientifically significant uncertainties, ecological costs, or ethically contested consequences. Equation (21) is a logical caution about the independence of predicates rather than a new theorem about the substantive content of international law.

One implication for practical infrastructure is that no undifferentiated “permission” flag should govern the entire lifecycle of a complex project. Each authorization record should identify a source, a competent authority, material and spatial scope, activities covered, conditions, status, provenance, and routes for verification. The system should preserve revisions and objections without replacing binding institutional decisions with algorithmic outputs. Such records can support research transparency and operational coordination, provided that privacy, sensitive ecological locations, restricted knowledge, and lawful confidentiality are handled appropriately. Their evaluable implementation is deferred to Sections 8–10.

2150 **7.9 A hypothetical transnational marine installation and open legal questions**

2151 Consider a wholly hypothetical consortium in which actors associated with States A, B, and C contribute
2152 capital, engineering expertise, and marine operations to a computational installation. The planned facility
2153 delivers scientific and commercial services globally and uses a water-column cooling interface, a seabed
2154 attachment, marine cables, and maintenance craft. Assume initially that the site lies within an EEZ
2155 claimed by a distinct coastal State D. The legal inquiry begins with the location and characteristics of the
2156 proposed structure, D's applicable coastal-state competence, the activities' environmental effects, and the
2157 relevant domestic and treaty rules. The nationalities of the investors alone do not displace jurisdiction
2158 attaching to the site. Nor does a cross-border contract determine the legal status of a scientific sampling
2159 operation performed by one of the maintenance robots.

2160 Changing only the site illustrates the relational character of jurisdictional classification. A facility
2161 located beyond national EEZs may be attached to a seabed area subject to an extended continental shelf,
2162 or it may be proposed for the Area beyond any national shelf. High-seas freedoms and seabed regimes
2163 must be separately assessed. If it uses mineral resources from the Area, the specific UNCLOS Part XI
2164 framework and relevant ISA mandate may become pertinent; the mere fact of seabed placement does
2165 not establish that the same regulator licenses the project's worldwide computing activity. Biodiversity
2166 assessment and other applicable obligations must be considered independently under their proper sources
2167 and scope. None of these examples is a finding that the hypothetical facility is legally permissible.

2168 The wider economic consequences also traverse different relations. Investors, workers, customers,
2169 researchers, and institutions may receive distinct benefits or acquire new capacities; biological commu-
2170 nities may experience thermal, acoustic, or structural effects. A payment among consortium members
2171 does not automatically compensate those ecological consequences, while a public data release does not
2172 necessarily provide affected communities with the instruments or expertise necessary to create further
2173 knowledge. The account of these heterogeneous circulations, and their contestable valuation across par-
2174 ticipating relational systems, will be developed in Section 9. The present argument is narrower: before
2175 asking how value should circulate, one must identify which activities, duties, and authority claims are
2176 actually constituted within the applicable marine normative orders.

2177 Marine nomos thus presents a difficult but productive meeting point between GR theory and insti-
2178 tutional practice. Legal regimes distribute powers by zone, type of activity, and relationship to state
2179 jurisdiction; scientific observations expose changing ecological conditions; normative practices mediate
2180 disagreements and shape what can be recognized as a relevant interest. Technological developments may
2181 generate activities for which existing categories require interpretation, coordination, or lawful revision.
2182 The task for a constructive research programme is to represent those dependencies accurately enough
2183 to support accountable decision-making, while preserving the possibility of disagreement, refusal, and
2184 legal remedy. Section 8 accordingly turns from the classification of normative orders to the concrete in-
2185 stitutional mechanisms by which marine activities can be authorized, monitored, challenged, and, where
2186 required, suspended or remedied.

2187 **8 Institutional Authority and Responsible Ecological Inquiry**

2188 **8.1 From marine nomos to accountable institutional decisions**

2189 The purpose of this section is to examine the institutional conditions under which oceanic observation,
2190 scientific intervention, and associated productive activity may be authorized, conducted, reviewed, and
2191 remedied. Section 7 differentiated normative orders and legal competences across marine spaces. That
2192 account identifies the sources and possible holders of authority; it does not, by itself, establish how
2193 a proposed activity becomes an accountable operation. An authorization is ordinarily bounded by the
2194 identity of the activity, its location, technical characteristics, purposes, duration, applicable procedures,
2195 and the competence of the decision-making institution. Its existence does not settle the accuracy of a
2196 scientific hypothesis, the acceptability of risks to marine organisms, or the distributive consequences of

economic benefits. Responsible ecological inquiry accordingly requires an analysis of institutions as ongoing configurations of decision, evidence, supervision, and possible correction.

The generative-relational perspective contributes a particular methodological emphasis. A research mission is constituted through relations among operators, instruments, living systems, locations, legal instruments, funding conditions, and the communities capable of observing and challenging consequences. The effective boundaries of these relations are not necessarily coextensive. A vessel may be flagged in one state, chartered by researchers in another, operate in a third state's exclusive economic zone, and process acquired observations in a fourth. Legal competence must be reconstructed for the particular action at issue. No single organizational name, contractual relationship, or machine-readable classification suffices to establish its full normative status.

Three evaluative questions should therefore remain separate. *Authority* asks whether a legally competent decision-maker has authorized the relevant act within its actual scope. *Epistemic justification* asks whether the underlying observations, models, and anticipated consequences are sufficiently supported for the stated scientific purpose. *Ecological and procedural justification* asks whether the interests and risks implicated by that act have received appropriately accountable treatment. Evidence may contribute to all three questions without making them equivalent. A permit does not produce ecological certainty; a credible simulation does not confer jurisdiction; and a consensus among project partners does not replace the interests of affected persons or nonhuman organisms.

For analytical clarity, this section considers four institutional activities: determining applicable authority, establishing evidence and assessment obligations, supervising changes during operations, and making accountability and remedies possible. The treatment describes general treaty rules and proposed institutional design principles. It does not render a jurisdiction-specific legal opinion or claim that any particular marine facility has been approved.

8.2 Authorization pathways for marine scientific and productive activity

The United Nations Convention on the Law of the Sea (UNCLOS) establishes differentiated regimes for marine scientific research and marine installations. Under Articles 245 and 246, coastal-state consent has particular significance for marine scientific research in the territorial sea, exclusive economic zone, and on the continental shelf, with legally specified conditions and qualifications. Research in the water column beyond the exclusive economic zone is treated under Article 257, while research in the international seabed Area is subject to Article 256 and the relevant provisions of Part XI [196, 193]. These distinctions have consequences for mobile surface vehicles, submerged sampling instruments, seabed observatories, and associated scientific work. They also caution against classifying a mixed scientific-commercial operation according to a single self-declared purpose.

Article 248 specifies information to be provided, ordinarily at least six months before a proposed marine scientific research project in the EEZ or on the continental shelf: its objectives, methods, geographical area, timing, sponsoring institutions, and opportunities for coastal-state participation. Article 249 addresses specified conditions concerning participation, access to data and samples, reporting, major changes, and removal of equipment. Articles 252 and 253 establish further procedures and consequences, including narrowly specified conditions for implied consent and coastal-state powers to require suspension or cessation of certain research activities. Such rules should be interpreted together, not simplified into a universal "permission after six months" formula. Actual notification, party status, circumstances, and the applicable legal conditions matter [196].

By comparison, Article 60 addresses artificial islands, installations, and structures of the types identified in the EEZ regime, and Article 80 applies relevant provisions on the continental shelf. A subsea computing module intended to operate for years, an anchored scientific station, and a transient sampling robot need not engage identical authorization and removal requirements. Nor should the label "research infrastructure" be allowed to obscure construction, energy provision, waste handling, navigation interference, or commercial processing associated with the facility [190, 191]. National environmental, maritime, labor, safety, and information laws may create additional duties. These domestic rules are not

Table 5: Illustrative authorization questions for distinct oceanic activities. The table is a research classification, not a permit checklist applicable in every jurisdiction.

Activity	Authority and scope questions	Evidence and oversight questions
Mobile ecological survey	Location, nature of marine scientific research, vessel and platform rules, consent conditions	Sampling methods, effects on animals, calibration, reporting, changes of mission
Permanent seabed observatory	Installation regime, site rights, cables, navigation and safety responsibilities	Baseline effects, long-duration monitoring, maintenance and eventual removal
Commercial marine computing facility	Installation and economic activity competences, operator responsibilities, connected services	Heat and noise, energy and material flows, cumulative effects, failure and decommissioning
Biological sample collection	Research and resource-related permissions, material transfer and applicable access rules	Specimen handling, welfare, ecological effects, provenance and permitted downstream use

interchangeable with the jurisdictional allocations in UNCLOS.

The table makes an analytical rather than doctrinal claim: different activities demand different jurisdictional questions, and the evidence needed to evaluate their consequences differs accordingly. In any actual project, the table must be replaced by a traceable mapping from the particular activity to authoritative sources, implementing law, and competent institutions. It is especially important to record uncertainties in legal classification. An uncertain classification is a reason for competent clarification, not evidence of the absence of regulation.

8.3 Environmental assessment, animal welfare, and incomplete ecological knowledge

UNCLOS Part XII provides a general environmental framework. Article 192 states the obligation of states to protect and preserve the marine environment. Article 194 concerns measures to prevent, reduce, and control marine pollution. Articles 204 and 205 concern monitoring and reporting; Article 206 requires an assessment, as far as practicable, where states have reasonable grounds to believe that planned activities under their jurisdiction or control may cause substantial pollution or significant and harmful changes to the marine environment [195]. This is a qualified treaty requirement. It does not imply that every sensor deployment automatically requires the same form of formal environmental impact assessment, nor that a completed document establishes the absence of ecological risk.

The Agreement on Marine Biological Diversity of Areas beyond National Jurisdiction (BBNJ Agreement), which entered into force on 17 January 2026, provides an additional treaty framework within its scope. Its environmental impact assessment provisions distinguish screening from subsequent assessment. In particular, Article 30 specifies screening when a proposed activity may have more than a minor or transitory effect, or when its effects are unknown or poorly understood; a further threshold determines when an environmental impact assessment is required. Implementation turns on party status, the activity's connection to jurisdiction or control, and the Agreement's provisions for relationships with other instruments and bodies [189, 200]. It would be inaccurate to claim that entry into force automatically established one universal licensing agency for every installation in areas beyond national jurisdiction.

For robotic ecological observation, scientific uncertainty may be especially significant when an environmental impact is both difficult to measure and capable of altering the evidence being collected. Acoustic propulsion or active sonar can change an animal's behavior; lighting can modify local activity; sampling can remove organisms or alter microhabitats; anchored equipment can create physical changes that persist beyond the observation period. The relevance and magnitude of such effects require species-, context-, and task-specific evidence. A technically gentle-looking or biologically inspired machine cannot be presumed welfare-neutral. Section 4 distinguished organism-level welfare, population relations, and ecosystem continuity. Assessment procedures should preserve these distinctions rather than averaging them into an undocumented composite value.

The institutional challenge is not merely to estimate expected harm. It is to specify which interests are considered, what constitutes sufficient evidence to proceed, how unknown impacts are represented, and who may require additional investigation. For a migratory animal or poorly characterized deep-sea habitat, the absence of observed injury may reflect limits of detection. Assessment should therefore address observational capacity and uncertainty alongside predicted impacts. Noninvasive alternatives, smaller sampling envelopes, seasonal exclusions where supported, independent observation, and mission withdrawal can be investigated as design responses. Whether any of these measures is legally obligatory depends on the applicable rules; their status as proposed safeguards must be distinguished from enacted law.

Cumulative effects warrant particular attention. Several individually limited deployments may interact through repeated exposure, shared energy and communication infrastructure, or accumulated seabed occupation. The scientifically relevant assessment unit may therefore extend beyond the footprint of a single machine and beyond the duration of one operator's permit. A consortium's internal division of projects should not be treated as proof that the ecological consequences are independent. Conversely, a claimed cumulative risk must be supported by a mechanism, affected system, and evidentiary bounds rather than inferred solely from the presence of multiple activities.

8.4 Participation, representation, and asymmetries of institutional capacity

A legally correct decision procedure may remain epistemically narrow when those who can produce relevant evidence are unequally positioned. Marine decisions frequently involve coastal authorities, flag states, scientists, operators, workers, fishers, local communities, and persons whose exposure to ecological or economic consequences differs from their formal authority. Appropriate inclusion depends on jurisdiction and the particular activity; there is no general entitlement to exercise another state's sovereign functions. Yet responsible research design can ask whose observations are missing, whether environmental information can be understood and contested, and how differences in technical capacity shape practical participation.

UNCLOS Article 244 expressly addresses dissemination of marine scientific knowledge and strengthening the research capabilities of developing states, while Article 249 provides coastal-state participation and access conditions for research falling within its scope [196]. These provisions are more specific than a general aspiration toward "open science." They also demonstrate why the circulation of a dataset should be distinguished from the capacity to reproduce, interrogate, or extend a scientific result. Providing readings from a proprietary marine instrument does not necessarily transfer calibration methods, inference software, or access to relevant scientific training.

Human representatives face additional asymmetry when decisions affect marine organisms. Organisms cannot be assumed to communicate interests through human legal procedures. Welfare studies, independent ethical review, and ecological expertise can inform representation, yet each may leave important experiential or ecological relations unrecognized. No proxy can be presumed to exhaust those interests. Section 5 therefore treated recognition as an ongoing responsibility and substantive ecological outcomes as distinct from procedural endorsement. In the present institutional context, a research proposal should record which possible interests were considered, which remain uncertain, and which alternative interpretations affected the proposed decision.

Research integrity also places limits on openness. Publishing precise locations of threatened organisms, culturally restricted knowledge, sensitive infrastructure vulnerabilities, or personally identifying data may cause harm. A useful design should record disclosure decisions and their reasons, preserve legitimate opportunities for challenge, and distinguish scientific reproducibility from unrestricted redistribution. Community-generated knowledge may also require context-specific conditions concerning acknowledgement and permitted uses. Such arrangements should be evaluated through applicable legal sources and agreements rather than reduced to automated assumptions of universal consent.

2327 8.5 Operational supervision and revisable conditions

2328 An authorization decision occurs in a changing relation between an activity and its ecological environ-
2329 ment. A robot may be deployed under one measurement plan and later encounter an organism, current
2330 regime, or physical condition that was poorly represented in the original assessment. An energy installa-
2331 tion may alter its load profile; a computing facility may expand beyond its modeled thermal envelope; a
2332 sampling programme may produce data that challenge its own ecological assumptions. Responsible su-
2333 pervision consequently requires a means to connect authorized scope, observed changes, and competent
2334 review. This need not presuppose a universal chronology: particular authorizations, observations, and
2335 revisions can be represented through documented event relations and relevant effective dates.

2336 At the engineering level, local control may respond immediately to measured disturbances. At the
2337 institutional level, however, the ability to change a control law does not entail authority to change the
2338 purpose, location, scale, or ecological effects of a mission. We should distinguish adjustments already
2339 covered by the approved operating envelope from material changes requiring notification or fresh autho-
2340 rization. The distinction is context-dependent and legally determined; Section 8.2 has noted the more
2341 specific UNCLOS research provisions concerning major changes and suspension. The institutional re-
2342 quirement is a documented interface between operational decision systems and actual legal conditions,
2343 without pretending that the software supplies the legal judgment.

2344 A provisional supervision record would identify the approved activity envelope, relevant evidence
2345 and assumptions, responsible persons, data sources, monitoring results, alert criteria, decisions taken,
2346 and any directions from competent authorities. It should support three kinds of correction. *Scientific*
2347 *correction* revises a claim or observation model. *Operational correction* modifies or halts an activity
2348 within lawful decision procedures. *Institutional correction* amends, suspends, replaces, or remedies a
2349 legally relevant decision through the proper authority. These processes can be coupled while retaining
2350 different competence and evidentiary requirements.

2351 This separation also protects against a superficially attractive but hazardous form of automated regu-
2352 lation. If an AI system identifies a likely ecological concern, it may provide a reason for human review
2353 or a precautionary operational response consistent with an established protocol. Its prediction alone can-
2354 not be treated as an administrative order with force of law. Equally, an operator's technical ability to
2355 override a robotic safety constraint does not create an ethical entitlement to do so. Escalation routes, doc-
2356 umented overrides, and post-incident review are thus substantive elements of accountable automation,
2357 even though their eventual legal design depends on the governing jurisdiction.

2358 8.6 Responsibility, environmental harm, and accessible remedies

2359 Accountability concerns what happens when scientific inquiry or productive activity causes damage, vi-
2360 olates conditions, or exposes omissions in the original decision. UNCLOS Article 235 recognizes state
2361 responsibility for fulfillment of marine-environment obligations and provides for appropriate domestic-
2362 law recourse concerning pollution damage. Article 263 separately addresses responsibility and liability
2363 regarding marine scientific research, including specified forms of pollution damage [195, 196]. These
2364 provisions should be read carefully: they do not establish an automatically available universal individ-
2365 ual claim for every ecological loss, nor resolve every private-law question arising from a multinational
2366 marine facility.

2367 A responsible research infrastructure should nonetheless make remedial questions legible. Which
2368 activity and actor are implicated? What physical and epistemic evidence links an act to a consequence?
2369 Which standards of causation and proof apply in the competent forum? Could urgent cessation or con-
2370 tainment prevent additional harm? What kinds of remediation, compensation, restoration, investigation,
2371 or institutional correction are legally or practically available? Distinguishing these issues is especially
2372 important when ecological damage is cumulative, delayed, difficult to observe, or distributed across
2373 several operators.

2374 Remedy should also include epistemic correction. If a robotic disturbance created a misleading eco-
2375 logical observation that was subsequently incorporated into shared models, the correction may need to

propagate through dependent scientific claims. The existence of a provenance link identifies candidates for review; it does not establish that every downstream inference is wrong. This distinction follows the task-relative epistemic framework developed in Section 6. A transparent record can support affected researchers and decision-makers in determining what requires re-evaluation, while preserving uncertainty about unobserved consequences.

Institutional design must consider a further asymmetry: an entity capable of constructing a complex installation may possess greater legal and technical resources than those exposed to its effects. Accessible complaint channels, independent evidence review, reasonable disclosure, and enforceable restoration or financial-security arrangements are therefore candidate mechanisms for study. Whether an instrument mandates any such measure requires source-specific examination. Its normative desirability likewise remains open to justified disagreement and competing legal obligations.

8.7 The institutional chain of reasons and reviewable decisions

The quality of an institutional record depends on what can be reconstructed from it. An authorization dossier should connect the factual submission and verified site information to the applicable instrument, competent decision-maker, reasons for applying the relevant legal classification, assessment and consultation records, an authenticated decision, conditions, time limits, monitoring duties and any subsequent amendment or suspension. Its purpose is to make existing decisions intelligible and contestable; it cannot confer powers on institutions that do not have them. The distinction between a *claim about authorization*, the primary decision record and a subsequent independent legal assessment remains essential.

A useful comparison is between *procedural compliance*, *substantive ecological evidence* and *remedial reachability*. Procedural compliance concerns whether identified steps and institutions satisfy the relevant rules. Ecological evidence concerns the reliability of observations, exposure estimates and resulting causal claims. Remedial reachability concerns whether affected actors can discover material changes, contest the reasoning and invoke an available competent process. The third question can remain unresolved even when the first two appear satisfactory; for example, an accessible decision may lack a feasible means of reviewing new data within the timeframe of an ongoing harmful exposure.

BBNJ Part IV provides a bounded illustration. For applicable activities, publication of screening decisions, opportunities for views on certain screening outcomes, consultation and publication of assessment reports, monitoring and subsequent review create differentiated accountability events [189]. They should be represented according to their own procedural triggers; they are not all equivalent to a private person's right to veto. Similarly, environmental-access rights established under particular regional instruments or domestic rules should be attributed to those sources, with their exceptions, rights-holders and enforcement routes [198]. This discipline also helps distinguish a review body's capacity to recommend corrective action from its legally conferred power to require it.

The operational research hypothesis is that a typed record which preserves these distinctions can reduce category errors and shorten the path between new evidence and competent review. A competing hypothesis is that well-maintained conventional administrative and scientific practices perform as well, or better, with less complexity. An evaluation should test both possibilities by assigning independent legal reviewers to realistic fact patterns, including incomplete party status, a disputed legal classification, altered activity, new ecological observations and contradictory decision metadata. Metrics can include false authorization assertions, omitted actors, time to identify the controlling decision and successful correction of mistaken citations. Such measures evaluate information support, not the substantive validity of a sovereign decision.

8.8 A scoped model of authorization and responsible continuation

For comparative research, let a proposed marine activity be represented as a structured description

$$a = (\ell, p, m, \sigma, \tau, \omega), \quad (22)$$

where ℓ describes the location and affected marine spaces, p the purpose, m the methods and material configuration, σ the operational scale, τ the represented operating interval or event scope, and ω the participating operators and organizational relations. These are descriptive coordinates selected for the assessment task, not a claim that an ecological activity is ontologically exhausted by six attributes. Let j denote a legally competent institution and d_j an authenticated decision record. The scope relation

$$\text{Covered}(a, d_j; L_j) \in \{\text{yes, no, undetermined}\} \quad (23)$$

asks whether a is covered by the decision under the applicable legal materials L_j . The third outcome is essential whenever classification, authenticity, territorial scope, or decision status is unresolved. In practical use, neither Covered nor L_j can be authoritatively determined by an unreviewed language model.

A second relation $\text{Evidence}(a; H)$ records task-specific scientific support and uncertainty against a represented observational history H . A third, $\text{Admissible}(a; N)$, represents independently specified ecological, welfare, and procedural constraints under an identified normative specification N . We may then define a research-oriented screening predicate,

$$\text{Screen}(a; d_j, H, N) = (\text{Covered}(a, d_j; L_j), \text{Evidence}(a; H), \text{Admissible}(a; N)). \quad (24)$$

This tuple is deliberately *not* a scalar score. A positive technical assessment cannot cancel a negative authority finding; an unknown welfare consequence cannot be converted into verified safety merely by increasing estimated information gain. Where multiple institutions possess distinct competences, the model must retain multiple scope relations rather than synthesize a fictional supreme permit.

Proposition 8.1 (Scope preservation under material change). *Suppose an authorization record d_j is interpreted under a specified legal model as permitting exactly the activities in a set P_j . If an initial activity a belongs to P_j and a materially changed activity a' is outside P_j , then authorization of a provides no authorization of a' by that record alone.*

Proof. By assumption, the authorized activities represented by d_j are exactly the elements of P_j . Since $a' \notin P_j$, d_j does not cover a' . The inference does not exclude the possibility of a separate authorization, an applicable exception, or a different legally supported interpretation of scope. $\square$

The proposition is intentionally elementary. Its value is methodological rather than doctrinal: an automated system should never assume that a permission survives an altered activity merely because the operator, machine, or general project name remains the same. Equally, the result depends on the correctness of the represented permission set, which itself requires competent legal determination. A research demonstration can test whether a system preserves this conditional distinction; it cannot prove that its representation of a real permit is legally correct.

8.9 Hypothetical transnational facility and institutional failure cases

Return to the wholly hypothetical installation introduced in Section 7.9. Investors associated with State A, engineers associated with B, and an operator associated with C propose a seabed computing facility in coastal State D's EEZ. The site includes a cooling interface, cables, vessel support, and an associated marine biological observation programme. A first institutional question concerns D's relevant jurisdiction over the covered installation and research, together with applicable national laws. Further questions concern the roles of flag states, contracting parties, employers, data controllers, and other legally relevant actors. The presence of four national connections neither creates a shared ocean sovereignty nor eliminates the need to examine each state's actual competence.

Consider a proposed increase in computing power. The operators expect additional heat dissipation and modify the cooling installation. The economic project may retain its name and ownership, but the relevant environmental assessment and permission scope may change. The legal consequence cannot be determined from the engineering fact alone: the parties must review the governing decision and applicable regulatory rules. An accountable system would expose the changed assumptions, responsible

authority, available monitoring, and unresolved classification. It would not report the expanded activity as approved merely because the initial facility was authorized.

A second hypothetical failure arises when a scientific robot discovers animals near its intended sampling track. The robot's model has insufficient evidence to estimate disturbance, but its navigation controller can continue safely in a mechanical sense. If the existing mission protocol specifies observation-only behavior under such uncertainty, an autonomous collection action would violate the research specification even if it improved a scientific score. If the applicable legal permission also excludes collection, it would raise a distinct compliance question. The example illustrates how mechanical feasibility, task design, animal welfare, and legal authority can diverge without requiring a single omniscient evaluator.

A third case concerns cumulative effects. Each consortium member proposes a technically separate survey and reports limited individual impacts. Later evidence suggests overlapping disturbance to one habitat. The institutional question is whether the applicable assessment and review frameworks require a cumulative analysis, and who can demand it. Scientifically, the hypothesis must be investigated using exposure histories, ecological baselines, and competing causal explanations. Procedurally, responsible actors should disclose the overlap and preserve the basis for review. An assertion that all activities were previously permitted does not by itself resolve the observed ecological question.

The cases point toward a precise research agenda. A prototype could be challenged with changes of operating scale, jurisdiction, research purpose, missing authority records, contradictory ecological observations, sensitive biological knowledge, and disputes about who may participate. The immediate test would be whether the system maintains the distinction between *covered*, *excluded*, and *undetermined* actions and routes consequential uncertainty to legitimate review. Subsequent work would assess the adequacy of ecological observations and institutional response. The fuller operational workflow belongs to Section 10; Section 9 first examines the common knowledge and heterogeneous value circulation generated by such activities, including the differences between lawful authorization and just distribution.

9 Common Ecological Knowledge Creation and Circulation

9.1 The production and circulation of ecological knowledge

The objective of this section is to investigate how ecological knowledge, scientific capacities, and heterogeneous forms of value are created and circulated across oceanic research networks, economic organizations, and distinct generative-relational contexts. Sections 6 and 8 developed, respectively, the epistemic problem of partial translation and the institutional problem of accountable authority. Their intersection raises a further question: when an oceanic observation or technological activity generates something valuable, which relations make that value possible, who can recognize or reproduce it, and what happens as its products and consequences circulate through different systems? The question applies both to scientific projects and to productive facilities that employ marine space, robotic technologies, energy, digital infrastructure, or biological knowledge.

An ecological observation does not emerge solely from an isolated sensor. Its credibility can depend upon instrument manufacture, calibration, navigation, an animal or habitat being observable under particular conditions, vessel operation, maintenance, model assumptions, data curation, and subsequent interpretation. Different actors contribute different capacities and may be situated in different jurisdictions. A research publication is one possible product of these relationships. Other products include datasets, tested procedures, trained personnel, computational methods, representations of ecological processes, institutional capacities, and questions that can organize further inquiry. These products do not necessarily belong to one owner, receive recognition through one institution, or circulate according to a common valuation rule.

The notion of a *knowledge commons* helps describe arrangements in which participants develop rules for shared creation, maintenance, access, and use of knowledge resources. Hess and Ostrom's treatment of knowledge as a commons emphasizes institutions and the threats posed by enclosure and inadequate preservation [61]. Contemporary open-science principles similarly concern findability, reusability, sci-

entific infrastructure, and opportunities for participation [214, 186]. The present inquiry extends this conversation to marine ecological relations and generative capacities. A repository holding observations may be accessible while the instruments, computational resources, methodological expertise, or institutional authority needed to generate comparable observations remain concentrated elsewhere. Hence the existence of public knowledge products is analytically distinct from the effective distribution of *knowledge-producing capacity*.

Scientific openness also involves justified boundaries. Location-sensitive biodiversity records, information contributed by communities under specific conditions, industrially confidential designs, and personal information may require differentiated access. The CARE principles for Indigenous data governance identify collective benefit, authority to control, responsibility, and ethics as indispensable considerations when the relevant communities' data and knowledge are involved [22]. These principles should be applied within their appropriate scope and with the participation of the relevant rights-holders. They do not establish that all ecological data have identical custodians or legal requirements. A constructive commons arrangement must be capable of documenting distinctions among publication, conditional access, refusal, stewardship, and review.

9.2 Heterogeneous value and generative capacity

The expression *value circulation* can conceal several different operations. A monetary payment transfers a legally or contractually specified entitlement in a currency. A scientific result can be copied and interpreted, yet its reproduction may require instruments and specialized knowledge. A marine organism's ecological role is embedded in habitat, population, and other biological relations that do not travel with the genetic description of the organism. A regulatory permission may authorize a specified action within a particular jurisdiction while leaving other uses unauthorized. In each case, the thing circulating, the receiving context, and the consequences of reception differ.

The distinction is especially important when economic value is generated through relationships that extend beyond a company's accounting boundary. The generative perspective on economic activity acknowledged in the Introduction motivates attention to the creation and circulation of ecological and social value. The present paper independently develops a task-specific generative-relational analysis of marine scientific production, ecological dependencies, and cross-jurisdictional institutional arrangements; its marine legal assessment and formal models require their own assumptions and evidence.

For purposes of analysis, we distinguish at least six kinds of value-bearing relationship. *Material flows* include energy, equipment, extracted samples, heat, and waste. *Monetary relations* include funding, wages, rents, fees, taxable income, and contractual payments. *Epistemic relations* include observations, model claims, inference, calibration methods, and scientific criticism. *Productive capacities* include facilities, skilled labor, software, maintenance arrangements, and abilities to undertake further inquiry. *Normative relations* include permissions, duties, restrictions, accountability, and rights to challenge decisions. Finally, *ecological relations* include organism welfare, habitat dependencies, population processes, and biophysical effects arising from the activity. These are dimensions used to interrogate a case, rather than mutually exclusive types of objects or a universal typology of reality.

A financial account may reliably report revenues and expenditures while omitting whether a participant gained the ability to conduct future research or whether an affected habitat retained the conditions required for its recovery. Conversely, a nonmonetary contribution may be scientifically or institutionally decisive without receiving a proportional financial reward. The analysis should therefore distinguish *creation, recognition, translation, access, appropriation, benefit, impairment, and remedy*. These operations can overlap, yet no logical equivalence among them is presumed.

The non-rival qualities of some digital products do not abolish material scarcity. A dataset may be replicated at relatively low marginal cost, while its collection and interpretation depend upon vessels, sensor supply chains, electricity, storage, skilled labor, ecological access, and opportunities for independent verification. Nor does the biological generativity of an ecosystem imply that its underlying resources or assimilative capacities are unbounded, as Section 3 emphasized. Any proposed expansion of oceanic

research or production must therefore represent both newly opened possibilities and the material dependencies or losses associated with them.

9.3 Cross-system circulation as partial and transformative mediation

Consider participating contexts G_A, G_B, G_C containing research institutions, firms, public authorities, and communities. Each context admits particular distinctions, interpretive practices, relevant histories, and criteria for evaluating claims. These contexts should never be equated mechanically with countries as unitary subjects: a coastal community and a national ministry in the same state may interpret a facility quite differently, while scientific groups situated in separate countries may employ closely related experimental practices.

A measurement generated in one context may circulate into another as an observation file, a derived estimate, a regulatory exhibit, or input for a new model. The circulation changes what the receiving institution is capable of investigating or doing. Yet even an apparently identical digital copy can become implicated in different epistemic and normative relations. An ecological signal interpreted as a proxy for fish abundance may be incorporated elsewhere into a commercial risk calculation, a conservation assessment, or a robot's mission plan. The validity of each interpretation depends upon its assumptions, data quality, intended task, and authorization to use relevant information.

A proposed *generalized mediation record* would therefore identify the originating evidence and practice, the task and receiving context, the transformations performed, properties provisionally preserved, known losses, contested interpretations, and permissions or constraints on further use. This extends the partial translation analysis of Section 6. Technical metadata standards such as W3C PROV-O can represent agents, activities, derivations, and provenance links; they provide useful infrastructure without determining whether an interpretation is substantively sound or legally legitimate [220].

Translation may generate new capacities rather than simply move an unchanged product. For example, researchers combining marine acoustic observations with environmental-DNA evidence could develop a new hypothesis about animal presence, yet each evidence stream has distinct sensitivity, sampling, and inferential limits. A collaborative protocol might enable experiments that none of the participating groups previously could perform. The resulting knowledge should be understood through its actual joint history, rather than imputed to a single participant because that participant owns the final database or model. Joint generativity, however, does not settle ownership, credit, or benefit-sharing automatically; those require appropriately scoped legal and institutional arrangements.

There is also a significant possibility of *failed mediation*. A terminology mapping may conceal incompatibility between observational units; a conservation model may omit a community's recognized uses; an economic contract may carry revenue while ignoring habitat effects; and a research institution may interpret silence as consent although the affected participants lacked opportunities to respond. These failures call for recorded uncertainty, contestation, and potentially refusal, rather than forced reconciliation. The system should permit translations to remain incomplete and should preserve the ability to reconstruct the evidence and decisions upon which consequential applications depend.

9.4 A typed model of circulation and its information limits

We introduce a provisional model whose purpose is to make the preceding distinctions examinable. Let I denote a finite or countable set of represented contexts and let K be a collection of specified types, such as material, monetary, epistemic, productive, normative, and ecological. A circulation record is a structured object

$$f = (i, j, k, a, h, s, r, \lambda), \quad i, j \in I, \quad k \in K, \quad (25)$$

where i and j identify represented originating and receiving contexts, a identifies a relevant activity, h its supported provenance or history, s the object or relationship being conveyed or affected, r an explicitly scoped recognition or transformation rule, and λ records relevant legal and institutional conditions. The

arrow-like origin–destination notation is a convenience for some kinds of transaction. A habitat disturbance may involve multiple affected contexts, including nonhuman ones, without a voluntary recipient. Likewise, an ecological impact need not constitute a transfer of property. These cases require an effect relation associated with the activity, rather than an artificial sale or exchange.

For an activity a , let $L(a)$ be its monetary and contractual record and let $E(a)$ be a task-specific ecological-effect description. The two records may have different evidence sources, uncertainties, and affected subjects. The following elementary result establishes a limitation relevant to claims that financial accounting could, by itself, settle all consequences of marine production.

Proposition 9.1 (Financial indistinguishability and ecological inference). *Let $\mathcal{H}$ be a set of admissible activity histories and let $L : \mathcal{H} \rightarrow \mathcal{L}$ be a financial-ledger representation. Let $E : \mathcal{H} \rightarrow \mathcal{E}$ represent an ecological property material to a declared evaluation task. If there exist $h_1, h_2 \in \mathcal{H}$ such that $L(h_1) = L(h_2)$ and $E(h_1) \neq E(h_2)$, there is no single-valued map $g : \mathcal{L} \rightarrow \mathcal{E}$ satisfying $E = g \circ L$ on all of $\mathcal{H}$.*

Proof. If such g existed, the equality $L(h_1) = L(h_2)$ would imply $g(L(h_1)) = g(L(h_2))$, and hence $E(h_1) = E(h_2)$, contrary to the assumption. $\square$

The proposition is a representation result, not a numerical estimate of environmental harm. Its antecedent is an empirically or analytically testable distinction between histories that financial records fail to separate. A detailed financial ledger may include environmental spending or liability provisions; the conclusion applies only when the particular ecological property remains unresolved by that representation. The general methodological consequence is a need for independently warranted ecological records where ecological differences matter to a decision. Conversely, detailed ecological measurement cannot alone reconstruct the allocation of legal authority or financial entitlements. An integrated infrastructure requires linked, *typed* records rather than an assumed reduction of all judgments to one account.

A circulation record is also insufficient to establish distributive justice. Determining whether a recipient obtained fair access, a contribution was recognized, or a harm received adequate remedy requires substantive criteria and competent review. Nor should the model presume that ecological interests become tradable merely because they can be represented. We can test the structure through incomplete records, contradictory valuations, revisions of reported observations, and actions whose benefits and harms reach actors absent from the original contract. These tests should reveal information loss and institutional omissions before any claim of operational adequacy is made.

9.5 Transnational marine production: a four-jurisdiction thought experiment

Consider a fictional installation developed through a consortium involving investors in Country A, engineering teams in Country B, and operators and logistics providers in Country C. The proposed facility is a marine data center supporting computational services and some environmental monitoring. Country D is the coastal state if the installation lies within a relevant part of its territorial sea, exclusive economic zone, or continental shelf. The designation of each state specifies a legal and logistical role in the hypothetical exercise; none should be regarded as a homogeneous GR subject. Communities, employees, researchers, regulators, service users, and ecological systems may have divergent interests within and across these countries.

The project could generate computing services sold internationally, profit distributions, wages, technical expertise, publicly usable environmental records, patents or trade secrets, tax obligations, and contractual permissions. It could simultaneously require electricity, materials, cables, cooling arrangements, seabed contact, maintenance vessels, and decommissioning plans. Heat release, acoustic disturbance, infrastructure placement, and interference with other legitimate marine uses would require site- and design-specific investigation; no actual impact magnitude is assumed here. Even if global customers derive substantial value from inexpensive computational services, their benefits do not establish the ecological acceptability of the site or settle duties owed under applicable law.

A site in D's territorial sea would involve a different legal configuration from a facility covered by the installation regime of D's exclusive economic zone, and both differ from a hypothetical facility beyond

Table 6: Illustrative value and responsibility relationships in a multinational marine computing project. Entries identify questions for investigation rather than legal determinations or measured outcomes.

Relationship	Possible product or consequence	Issue requiring separate assessment
Capital and enterprise	Finance, contracts, revenue and profit	Ownership, taxation, attribution and risk allocation
Engineering and re-search	Designs, operational methods, observations	Credit, effective access and replicable capability
Site and operations	Power, construction, maintenance, services	Competent authorization and lifecycle conditions
Service circulation	Computing availability, derived knowledge	Equitable access, data rules and dependency
Ecological effects	Habitat change, disturbance, uncertain recovery	Affected interests, monitoring and remedy

national jurisdiction. Sections 7 and 8 established why the status of the water column, seabed, operator, structure, vessel, cables, and processing activities must be identified separately. The distribution of corporate revenue follows still other applicable tax and contractual relationships. OECD transfer-pricing standards, for example, address related-party cross-border pricing and associated tax questions; they do not purport to calculate the total moral or ecological value generated by a marine project [150].

The thought experiment produces several possible governance designs. Under a predominantly proprietary arrangement, a consortium might retain control over computing services and operational data while contracting specified ecological monitoring and restoration obligations. Under a negotiated commons-oriented arrangement, participating institutions might establish additional shared scientific instruments, trained research teams, access to appropriately governed environmental datasets, and public reporting of ecological monitoring. Neither institutional form guarantees justice as a consequence of its label. The relevant comparison concerns actual rights and capacities, distribution of benefits and burdens, ecological conditions, and opportunities for correction. Similarly, an offer to release low-value datasets cannot be presumed an adequate substitute for restricted access to critical research methods or for ecological harm.

9.6 Existing benefit-sharing law and the limits of generalization

International legal instruments already recognize certain forms of benefit-sharing, although their scopes differ substantially. The Nagoya Protocol concerns access to genetic resources and the fair and equitable sharing of benefits arising from their utilization, together with associated traditional knowledge within its defined scope and applicable legal arrangements [167]. The Agreement on Marine Biological Diversity of Areas beyond National Jurisdiction (BBNJ Agreement), in force since 17 January 2026, addresses marine genetic resources and digital sequence information, relevant benefit-sharing arrangements, capacity-building, and the transfer of marine technology within the treaty's provisions [189, 188]. These regimes are consequential for appropriately qualifying marine biological research and utilization. They should not be misrepresented as automatic obligations to distribute the profits of every oceanic computing, transport, or energy facility worldwide.

In the hypothetical data-center case, the scientific component might acquire or analyze marine genetic resources, in which case the applicable regime would require activity-, area-, and party-specific investigation. A conventional computing-service transaction, by contrast, is not transformed into utilization of marine genetic resources merely by being located at sea. An environmental-monitoring duty imposed through a competent authorization may coexist with separately negotiated community-benefit agreements and intellectual-property arrangements. Different institutions can therefore regulate distinct relations within one physical project.

Potential mechanisms for broader participation include access to scientific infrastructure, joint training, appropriately scoped data-sharing agreements, research partnerships, independently supported envi-

ronmental monitoring, contractual benefit arrangements, decommissioning security, and remedial facilities. Their force and enforceability differ. Some may follow existing law or permit conditions; others may be freely negotiated; still others would require legal or institutional innovation. The analytical work is to specify the source of each obligation and the actor with authority to enforce or challenge it. Invoking a general ideal of global benefit supplies no substitute for these determinations.

A deeper justice problem concerns durable dependency. Where research groups receive a dataset but lack the right, equipment, or training needed to scrutinize its production, the apparent circulation of knowledge may increase reliance on the dominant producer. Conversely, collaboration can expand the range of independently testable scientific propositions and the ability of different institutions to generate further knowledge. These outcomes should be evaluated separately from the number of records transferred or the amount of financial compensation received. This is the sense in which *generative capacity* is an additional object of institutional concern.

9.7 Benefit-sharing, traditional knowledge and open-data conditions

The temporal and subject-matter boundaries of benefit-sharing are particularly relevant to a scientific data infrastructure. The BBNJ Agreement's Part II addresses activities concerning marine genetic resources (MGRs) and digital sequence information on such resources from areas beyond national jurisdiction, with Article 10 specifying temporal application and exceptions. Article 13 addresses traditional knowledge associated with these MGRs and held by Indigenous Peoples and local communities: it calls for relevant and appropriate measures aimed at access with free, prior and informed consent or approval and involvement, on mutually agreed terms. Article 14 identifies a range of non-monetary benefits, including access to samples and sequence information, FAIR scientific data and capacity-building; other provisions specify monitoring, identifiers and mechanisms for monetary benefits [189]. These differentiated arrangements must not be compressed into a single statement that "all marine data are automatically open" or that every offshore commercial activity is subject to one benefit-sharing formula.

Three distinct questions arise for an operator of a shared ecological information infrastructure. First, what is the provenance and legal classification of each physical sample, sequence record or knowledge contribution? Second, which entity holds applicable access, custodianship, confidentiality or other responsibilities, and what actual source defines them? Third, what access, benefit or research-capacity outcome has been produced for an identifiable recipient? The answer to the third question may be contested even when disclosure formally occurred. A database that lists a publicly accessible file but omits instrument dependence, necessary expertise or restrictions on reuse may overstate the effective circulation of scientific value.

Open-science and CARE-based governance considerations may be complementary while addressing different normative objects [186, 22]. Their use requires genuine attention to the rights and knowledge practices of relevant communities. A protocol allowing justified restrictions on sensitive ecological locations, protected traditional knowledge and operational security is compatible with the general goal of reproducibility when the restrictions are specified, necessary, reviewable and paired with appropriate alternative scrutiny. The research question is how to make access and accountability coexist in concrete contexts, rather than presuppose that maximum public disclosure always creates justice.

9.8 An accountable infrastructure for common production and circulation

A prospective implementation would connect a federated marine research workspace with a governance registry and a typed circulation ledger. Participating institutions could publish qualified descriptions of observing instruments, calibration histories, ecological studies, executable methods, computational dependencies, and supported knowledge claims. Provenance graphs would connect a derived product to the relevant activities, instruments, and authorized uses. A mediation dossier would record cross-context interpretations, alternative models, declared task invariants, known translation losses, and objections. Distinctly, a value-circulation register could record monetary arrangements, nonmonetary contributions,

access to productive capacities, and material ecological effects without imposing a universal unit of value.

The operational architecture must separate a proposed use from permission to perform it. An observation may be technically available while governed by limits on disclosure, sensitive location information, community protocols, intellectual property, or legal obligations. Similarly, an AI-generated inference about the condition of a marine habitat may inform a review but cannot determine the applicable law or confer an institutional authorization. A robust workflow would request specific review when the proposed action exceeds recorded permissions, invokes uncertain ecological assumptions, or materially changes the activity's expected consequences.

A further requirement is *correction reachability*. When an instrument is found to be miscalibrated or a translated ecological claim is shown to be misleading, researchers should be able to identify downstream scientific products that relied upon the affected record and communicate the revision to responsible stewards. A provenance link establishes a possible dependency, however, rather than proving the magnitude of an actual downstream effect. Reviewers must assess the affected claim, decision and model individually. The same principle applies to revised access conditions or institutional permissions: a recorded change can trigger an accountable review without pretending that software can settle the underlying legal dispute.

The governance of the infrastructure should itself be open to challenge. Public-access promises can mask asymmetric access to computational resources; centralized discovery indexes can privilege particular taxonomies; and translation services can make unfamiliar scientific contributions legible only through categories favored by the receiving institution. Participating actors therefore need meaningful routes to request correction, disclose omitted relations, challenge representations, and negotiate alternative interfaces. Marine organisms cannot use these procedures directly; animal-welfare evidence, ecological baselines, independent review, and carefully specified human responsibilities are consequently indispensable. Sections 4 and 5 developed the epistemic and normative reasons for maintaining these safeguards.

9.9 Evaluation, failure cases, and the research continuation

The framework remains a research proposal rather than a deployed system. Its adequacy should be evaluated using counterexamples that a superficial data-sharing or accounting architecture would mishandle. A first scenario would hold reported revenue constant while varying ecological disturbance, testing whether assessment procedures recognize the omitted ecological dimension identified in Proposition 9.1. A second would provide two technically interoperable observation datasets with incompatible sampling assumptions, testing whether the mediation record exposes rather than hides the limitations. A third would provide nominally open observations while restricting calibration or computation, testing whether the evaluation distinguishes data access from effective scientific capacity.

Additional scenarios should include an institution that lacks competence to authorize the proposed marine operation, a sensitive biological observation whose indiscriminate publication could cause harm, and a multinational activity whose contracts omit an affected party. Each case requires distinct outcomes for technical success, epistemic warrant, legal scope, and normative review. Abstention, unresolved translation, restricted disclosure, or a request for further evidence can be justified outcomes rather than software failures. Empirical evaluation would require relevant participant involvement, independent legal review, field-appropriate ecological expertise, and evidence of actual institutional and material effects.

This section's contribution is therefore a conditional architecture of *common knowledge creation and heterogeneous circulation*, anchored in traceable research practices and scoped normative authority. It does not imply that all interests can be reconciled, all marine resources can be made available for production, or all values can be represented in monetary terms. It also does not assert that a new international levy, global ownership regime, or universal licensing institution already exists. The next section investigates how these conceptual and institutional commitments might be translated into operating procedures for ecologically accountable marine research. Later engineering sections then examine the observational

platforms, embodied control, energy systems, and prospective marine facilities to which those procedures could apply.

10 Procedures for Ecologically Accountable Research

10.1 The problem of translating normative commitments into research practice

The objective of this section is to formulate a practical procedure for planning, authorizing, operating, reviewing, and, when appropriate, discontinuing marine scientific activities whose ecological consequences are uncertain. Sections 5–9 developed the conceptual and institutional grounds for accountable inquiry: differentiated interests, unequal capacities to participate, scoped legal competences, and heterogeneous forms of value. These grounds acquire operational meaning only when a research team can specify who reviews a proposed mission, which evidence is needed, how the activity is constrained, what constitutes a material change, and what happens if the initial justification proves inadequate. The problem is acute for robotically mediated inquiry because a vehicle may act at a distance from responsible human decision-makers, while its environment and scientific target change in ways not completely represented by the mission model.

An *ecologically accountable research procedure* is proposed here as an ordered but revisable set of institutional and technical operations linking a mission description, evidential assessment, lawful authorization, operational restrictions, environmental observation, and accessible routes for reconsideration. Its organization need not presuppose a universal direction of historical time. The terms *prior assessment*, *operation*, and *follow-up* describe dependencies within a represented mission workflow. Their relevance depends upon the actual event structure, the legal requirements governing that mission, and available evidence. A correction may reopen an earlier assessment; a discovery may require renewed consultation; and a later model may revise the interpretation of observations obtained during an earlier operation. The procedure is therefore better represented as a partially ordered and revisable decision history than as a sequence of irreversibly closed bureaucratic forms.

Four distinct judgments are required. *Technical feasibility* concerns whether the equipment and operators can perform the activity under specified conditions. *Epistemic warrant* concerns whether the measurements, research design, and interpretations are adequate to the intended scientific question. *Legal authorization* concerns whether competent authorities have lawfully permitted the described activity within their actual jurisdiction and any associated conditions. *Ecological and procedural justification* concerns the treatment of marine organisms, habitats, affected human interests, uncertainty, and possibilities of challenge. An affirmative judgment on one dimension does not determine the others. A functioning autonomous vehicle can encounter an unauthorized sampling request; a permitted survey can yield uninterpretable data; and an accurately measured impact can require a halt despite the technical feasibility of further operation.

The framework is a proposed research and governance architecture, not an assertion that one comprehensive international permit, animal-welfare protocol, or review procedure presently regulates every marine activity. Actual duties must be established for the location, technology, species, kind of activity, state obligations, and competent decision-maker. The remainder of the section examines how a general procedure can support such determinations without pretending to replace them.

10.2 Mission characterization, jurisdiction, and reviewable commitments

A practical mission dossier should begin with a bounded characterization of the proposed activity. At minimum it identifies the scientific question, observing or productive method, geographical and depth envelope, surface and submerged components, responsible organizations, platform registration and control arrangements where relevant, expected duration, sampling or manipulation actions, energy and communication arrangements, and foreseeable recovery or decommissioning measures. For ecologically sensitive activities it should also specify target and potentially exposed organisms, plausible disturbance pathways,

the ecological baselines used in assessment, material uncertainties, and the sources from which those assumptions were derived. The purpose is to make the operative scope contestable and comparable with subsequent activity, rather than to claim that a sufficiently long form eliminates unknown consequences.

Jurisdictional analysis must be applied to the particular act rather than assigned once to an entire project. A mission can combine a vessel voyage, robotic environmental imaging, biological sample acquisition, deployment of an instrument on the seabed, transboundary data transfer, and later commercial application. Different elements can engage different authorities and legal obligations. UNCLOS Part XIII addresses marine scientific research, while the regimes of installations, navigational activities, environmental protection, and the Area have their own conditions and qualifications [196, 190, 191, 193, 195]. As shown in Section 8, an institution's scientific expertise or operational ownership does not confer the power to grant a coastal-state permission. Particular research-consent mechanisms, notification duties, and amendment requirements should be established through current and competent legal review.

Each operative decision should be associated with its legal basis, issuing or responsible authority, version, spatial and functional scope, validity conditions, reporting duties, and route to review. A document described merely as *approved* carries insufficient information if it does not distinguish, for example, permission to record images from permission to collect animals, install anchored equipment, or increase acoustic output. The documentation should state what remains unresolved. Where an activity is not adequately classified, an *undetermined* state is preferable to imputing permission from the absence of a recorded prohibition.

The mission dossier should also identify substantive commitments undertaken independently of legal necessity. A project may adopt non-invasive methods, additional acoustic monitoring, or community consultation beyond its enforceable legal minimum. Such commitments are relevant to scientific and ethical review but should be marked as institutional policy, contract, research condition, or voluntary practice, as appropriate. In particular, a consortium should not present voluntary access or benefit-sharing commitments as treaty requirements unless the applicable treaty actually provides them.

10.3 Ecological assessment, uncertain baselines, and alternative designs

The assessment of ecological effects begins with a question of evidential coverage. An oceanic research site may have incomplete seasonal observations, heterogeneous records of animal presence, sparse measurements of temperature or chemistry, and little evidence about cumulative interactions with nearby activities. A baseline assembled from a short survey can be useful without representing a complete or undisturbed ecological state. Reviewers should document spatial and observational limits, instrument detection probabilities, relevant environmental variability, and the degree to which recorded conditions can be generalized to the proposed operational envelope.

UNCLOS Articles 204–206 address, respectively, monitoring the risks or effects of marine pollution, publication or communication of reports, and assessment of potential effects where the specified grounds and thresholds apply [195]. The BBNJ Agreement, in force since January 2026, establishes more elaborated screening, assessment, consultation, decision, monitoring, reporting, and review requirements for qualifying activities within its scope. Its Article 30 addresses screening in circumstances including more-than-minor or transitory potential effects or effects that are unknown or poorly understood; Articles 31–37 set out further procedures and associated responsibilities [189, 188]. Neither instrument supplies a single automatic environmental-impact-assessment threshold for every marine scientific or commercial operation. The applicable regime, party status, character of the activity, and territorial scope require specific examination.

In an accountable research design, assessment should consider meaningful alternatives as well as modifications to the preferred platform. These may include remote sensing instead of animal contact, passive acoustics instead of active transmission, opportunistic sampling instead of dedicated disturbance, a temporary vehicle mission instead of a permanent mooring, or a terrestrial or coastal facility rather than a seabed installation. Alternatives should be compared against a specified scientific objective: a less intrusive method that cannot answer the actual question is not scientifically interchangeable with

the proposed experiment, but the difference must be explained rather than asserted. A *no deployment* or *defer pending evidence* alternative should remain available where expected consequences are material and poorly characterized.

Ecological assessment should not be collapsed into one weighted utility score. Individual animal welfare, population-level outcomes, habitat connectivity, water chemistry, noise, heat, energy use, invasive-species risks, and resource extraction can require separate thresholds or review competences. Effects may be cumulative across repeated missions and across different operators; the absence of demonstrated harm in one brief pilot cannot be generalized to indefinite continuous activity. Responsible comparison should record both expected scientific benefits and the uncertainties attaching to proposed protective measures. Where decision-relevant baseline information is missing, the appropriate response may be a narrower observation, a non-intervention study, or postponement.

10.4 Marine organisms, research encounters, and operational precautions

A scientifically justified encounter with marine life can impose risks on the organism whose behavior or physiology is being investigated. Section 4 distinguished responsiveness, biological agency, possible sentience, and the limits of inferring experience from human observational categories. A practical procedure should translate these distinctions into review questions without assuming that a comprehensive account of the animal's interests is available. Relevant questions include whether contact, capture, illumination, sound, chemical exposure, sample removal, or proximity could alter behavior or welfare; whether repeated exposure can produce cumulative disturbance; whether particular life stages or ecological relationships are more vulnerable; and what observations would justify stopping or modifying the encounter.

The choice of sensing method has a welfare dimension. A soft-bodied or biologically inspired robot cannot be presumed harmless merely from its compliant morphology. Its movement, propulsion wash, materials, light, acoustic emissions, and contact geometry still require assessment. Conversely, a rigid remotely operated vehicle may sometimes achieve an observation with less intrusion than a prolonged close-following soft robot. Welfare protection therefore concerns the actual interaction among the robot, animal, habitat, and experimental protocol. Independent review is especially important when the research question concerns pain, stress, or cognition and when a proposed procedure relies upon behavioral proxies with uncertain interpretation.

Applicable national and regional frameworks differ. As one bounded illustration, NOAA Fisheries describes authorization pathways for certain marine-mammal research under United States law, including permit applications that specify methods, species, potential exposures, study areas, and review procedures [128]. Such requirements are neither universal oceanic permissions nor an exhaustive model for invertebrates or fish. A multinational mission should instead assemble the relevant species- and activity-specific rules together with institutionally justified precautions.

A proposed *encounter protocol* would specify permitted proximity or contact envelopes where supportable, sensor and emission limits, required observation capabilities, procedures for responding to detected animals, and predefined withdrawal or recovery actions. Some high-noise marine operations use designated monitoring and shutdown arrangements; their numerical criteria are highly activity- and jurisdiction-specific and should not be copied into unrelated robotic science as universal safety distances. Procedures must also identify observer detectability limits. Failing to detect an animal in poor visibility is not equivalent to establishing its absence. Where the necessary conditions cannot be observed or operationalized, authorization for an otherwise discretionary intervention should not be inferred from silence in the sensor stream.

10.5 Adaptive mission states and evidence-bounded operational authority

Operational conditions in the ocean can depart from forecast dynamics through tides, stratification, currents, sediment disturbance, shifting animal distributions, and equipment degradation. A rigid preapproved route may be infeasible or scientifically uninformative under such conditions. Yet adaptive au-

Table 7: Illustrative reviewable mission states and possible evidence-triggered responses. These are proposed operational distinctions, not universal treaty-defined categories.

State	Evidence or institutional condition	Possible response
Proposed or under review	Scope, legal competence, environmental or welfare evidence incomplete	Collect information; consider alternatives; defer authorization
Conditionally authorized	Competent decision and applicable restrictions documented	Release bounded task instructions and monitoring plan
Active within scope	Required observation and safety assumptions remain supported	Continue, log actions, re-estimate dynamics
Restricted or paused	Material uncertainty, disturbance, exceedance, or change of mission purpose	Reduce emissions or motion; withhold intervention; request review
Terminated or recovering	Recovery, withdrawal, or stopping condition is triggered	Secure equipment and site within safe operational limits
Post-operation assessment	Monitoring, incidents, samples, or derived claims require evaluation	Report, correct, remedy, or reopen decision

tonomy does not imply a general authority to redefine the mission’s legal or ethical scope. We should distinguish *control adaptation within a reviewed envelope* from *material mission alteration requiring renewed review*. The first might adjust propulsion to a measured current; the second might replace passive imaging with animal capture or move a seabed installation into an unassessed habitat.

The engineering work of later Sections 13, 14, and 21 will address local control, learning, and transition-aware supervision in detail. At this institutional level we propose a small set of reviewable operational states: *proposed*, *conditionally authorized*, *active within scope*, *restricted*, *paused for review*, *terminated or recovering*, and *post-operation assessment*. The states denote documented decisions rather than an assertion of one universally mandated international workflow. A mission can return to proposal or assessment if scientific evidence or conditions materially change. Local autonomous controllers may carry out only transitions delegated under the applicable mission instructions; the normative authority to expand those instructions remains with competent persons or institutions.

The distinction between an adaptive technical response and a new permission is particularly important for learning-enabled robots. A controller may improve local prediction by updating hydrodynamic parameters or adapting its swimming gait. Its learned policy might also propose an efficient but unreviewed maneuver near a sensitive habitat. The latter proposal requires screening against the actual authorization and ecological conditions. Learning algorithms cannot confer competence upon themselves by demonstrating improved navigation or predicted scientific information gain.

Likewise, a *criticality supervisor* must distinguish a dynamical transition or stability margin from an ecological or legal boundary. A robot may experience no impending mechanical instability while its repeated presence disrupts animal behavior. Conversely, a severe model-prediction failure can justify withdrawal even when no animal has been detected. These conditions require distinct evidence and decision rules. Pausing, requesting human review, or returning to a safer state are legitimate outcomes of autonomous research, subject to the constraints of maintaining vehicle recovery and emergency safety.

10.6 Provenance, participation, correction, and responsible circulation

Accountability continues through the production and circulation of scientific records. A mission should maintain versioned descriptions of its research question, instrument configuration, calibration, observation context, control-policy changes, environmental monitoring, approvals and amendments, significant deviations, and reviewer decisions. These records permit a future investigator to distinguish a change in ecology from a change in sampling intensity, sensor sensitivity, classification method, or vehicle behavior. The Ocean Best Practices System provides an existing infrastructure for discovering and sharing marine standard operating procedures, methodological documentation, and related practices; its existence does not, however, certify the epistemic adequacy or legal permission of every application [141,

55].

A correction procedure should connect questionable measurements to downstream claims and operational uses. If a sonar calibration changes, the team should identify the observations and derived ecological inferences potentially dependent upon the affected record. If an autonomous vehicle disturbed the animals whose abundance it was supposed to estimate, the report should acknowledge the threat to interpretation. A dependency graph is useful for locating potentially affected products, but it does not establish that all those products are substantively false. Each correction requires a scoped scientific assessment and communication to the relevant custodians. Records of correction should remain accessible to authorized reviewers, including where raw data require restrictions for privacy, cultural sensitivity, conservation, or other legitimate reasons.

Procedural legitimacy also depends on whose knowledge and interests enter the assessment. A multinational marine project may affect a coastal community that has little control over a remote scientific or commercial consortium. Relevant scientific and situated knowledge, including appropriately governed community-contributed observations, should be able to inform scoping, baseline selection, and interpretation. The design must avoid the opposite error of assuming that every community or state speaks with one voice. Participation requires attention to language, access to evidence, independent expertise, time, material capacity, and routes to contest a representation. Scientific knowledge exchange can support such capacities, but a publicly available dashboard does not guarantee effective involvement.

Animal interests and ecological relations do not enter these procedures as ordinary contractual parties. The dossier should consequently identify who has been assigned responsibility for animal-welfare evidence, ecosystem effects, independent scientific review, and possible objections to an intervention. It should allow records of uncertainty and dissent rather than forcing every participant's claims into a single harmonized valuation. Sections 6 and 9 showed why generalized translation, financial accounting, and knowledge circulation can otherwise conceal important differences across relational contexts.

Finally, an operational procedure should make remediation possible. This entails pre-identifying emergency contacts, recovery responsibilities, reporting recipients, funds or guarantees where lawfully required, and the processes by which unforeseen damage or excluded interests can be raised after a mission. Remedial measures depend on the actual law and injury; neither compensation nor publication alone should be treated as automatic restoration of ecological conditions. Procedures should permit reconsideration of future missions if records indicate persistent or cumulative effects that initial assessments failed to capture.

10.7 Normative provenance and escalation during adaptive missions

An adaptive oceanic mission requires more than a record of the latest controller output. A reviewable operational history should preserve five links: the observation and its calibration; the inferred environmental state and uncertainty; the proposed action and predicted exposure; the legal and institutionally approved operating envelope; and the human or authorized institutional response when a proposal falls outside that envelope. Each link has a different failure mode. Model uncertainty may make the ecological prediction unreliable; a missing decision record may leave legal scope unresolved; an outdated mission plan may fail to communicate a changed acoustic or material exposure. A single Boolean "safe to proceed" obscures these differences.

A prospective procedure can encode state labels such as proposed, scientifically-reviewed, legally-scope-checked, authorized-with-conditions, restricted, review-required and closed. These are design labels, not legal categories or approvals. The machine's permission to execute a specific action must refer to an independently authenticated authorization and an appropriately reviewed operational rule, including the spatial, temporal, technical and activity-specific scope. If critical metadata is missing, the appropriate software outcome is an unresolved classification and routing to a competent human process, while any immediate safety response must follow separately established operational and legal arrangements.

This design addresses an ethical tension of dynamic scientific exploration. Seeking additional in-

formation can itself increase acoustic, light, sediment or contact exposure. Conversely, prematurely terminating observation can leave serious effects undetected. A defensible mission design should therefore specify advance limits and stopping criteria, emergency branches, obligations to preserve records, and opportunities for independent review of the evidence generated by its own intervention. These requirements may derive from particular laws, permits or research-ethics procedures, or may be proposed as additional safeguards; each record should make that status explicit. Their effectiveness remains an empirical research question until tested against realistic operations and affected reviewers.

10.8 A bounded formal model of review and continuation

We can represent an operating proposal as a tuple

$$d = (a, \ell, \theta, \mathcal{E}, \mathcal{L}, \mathcal{W}, \mathcal{Q}), \quad (26)$$

where a identifies the activity and its technical specification, ℓ the spatial or jurisdictional scope, θ the relevant operational conditions, $\mathcal{E}$ the scientific and environmental evidence, $\mathcal{L}$ the documented applicable authority and restrictions, $\mathcal{W}$ the welfare and ecological assessment, and $\mathcal{Q}$ the review, participation, and remedy routes. These are typed record fields, not exhaustive representations of reality or legally authoritative determinations. Some fields can be incomplete, disputed, or subject to revision. Their usefulness depends upon the fidelity of the evidence and the competence of reviewers.

Let $\Omega(d)$ denote the set of environmental and institutional configurations still compatible with the recorded evidence for a specific decision task, and let $A_{\text{authorized}}(d)$ denote the actions whose legal and operational scope has been affirmatively established by the responsible review. For a configuration ω let $A_{\text{ecological}}(\omega; d)$ be those actions satisfying specified substantive ecological restrictions. A conservative screening set is

$$A_{\text{screen}}(d) = A_{\text{authorized}}(d) \cap \bigcap_{\omega \in \Omega(d)} A_{\text{ecological}}(\omega; d). \quad (27)$$

The intersection is a methodological *screen*, conditional upon the adequacy of Ω , the specifications of the restriction sets, and independent legal review. It is not a computational guarantee of ecological safety or a legal license. Its emptiness means that no action has been justified against all configurations included in this representation; this can motivate better observation, a narrower action, a lawful emergency response, or a review of the assumptions.

Proposition 10.1 (Expansion of uncertainty weakens a fixed conservative screen). *Fix the proposal's established authorized-action set and its ecological restrictions for all candidate configurations. If new evidence or a corrected model expands the represented compatible configurations from Ω_1 to Ω_2 , with $\Omega_1 \subseteq \Omega_2$, then the corresponding conservative action sets satisfy $A_{\text{screen}}(\Omega_2) \subseteq A_{\text{screen}}(\Omega_1)$.*

Proof. An action in $A_{\text{screen}}(\Omega_2)$ belongs to the fixed authorized set and satisfies the ecological restriction for every $\omega \in \Omega_2$. Since $\Omega_1 \subseteq \Omega_2$, it satisfies those restrictions for every $\omega \in \Omega_1$ and therefore also belongs to $A_{\text{screen}}(\Omega_1)$. $\square$

This elementary set-theoretic result provides an explicit criterion for a *screening mechanism*, not a theorem about the nature of marine ecosystems. If uncertainty expands, an unchanged conservative screen cannot become more permissive simply by ignoring additional plausible conditions. In application, however, an updated model can also change the specification of ecological restrictions or reveal that an apparent permission was narrower than recorded; the fixed-set assumptions would then fail, and a new review is required. A smaller uncertainty set need not establish truth, and the absence of an action from the screen does not by itself determine how an institution should resolve competing obligations. Emergency recovery may be governed by separate applicable rules and necessities.

The procedure also requires an *institutional abstention* state. An automated assessor should return “scope unresolved” when no competent determination is documented, “evidence insufficient” when required ecological conditions cannot be evaluated, and “review required” when an operational change

exceeds the established envelope. These outcomes should not be silently converted into permission by a downstream autonomous controller. Nor does the presence of a machine-readable authorization imply that a reviewing institution has considered the latest scientific evidence. Explicit provenance and revocation or amendment channels are therefore needed alongside the mathematical screen.

10.9 A transnational facility case, adversarial tests, and the engineering transition

Consider again the fictional consortium with investors in Country A, engineering teams in Country B, operating personnel in Country C, and a proposed coastal-state site in Country D. Its marine data center would initially be accompanied by seabed mapping, robotic inspection, and environmental monitoring. For a real project, a competent analysis would have to establish the precise maritime zone, nature of the installation and associated activities, coastal-state and flag-state competences where applicable, electrical and communications links, relevant environmental rules, and any additional applicable data or labor obligations. A research operator should not infer that a permit for environmental sampling also authorizes seabed occupation, a cooling discharge, cable work, or continuous industrial service. Equally, a license for an installation should not be presumed to authorize unlimited biological sampling.

Suppose the operator proposes to enlarge the module, add active acoustic inspection, and extend autonomous monitoring into an ecologically sensitive area. A credible accountable procedure would identify the changed actions; compare them with the actual permission and assessment boundaries; update the baseline and cumulative-impact analysis; and determine which institution must review the change. If the robot observes unusual animal behavior but its sensors cannot establish whether the reaction is caused by noise, light, current, or another source, the appropriate response is not to classify the impact as absent. It may instead limit the operation while independent observation or further modeling determines which explanations are supported.

The same case illustrates how scientific and economic circulation complicates accountability. Environmental observations could be processed by an AI system in Country B, transmitted to researchers elsewhere, and summarized for the coastal authority in D. A model update may change the estimate of habitat disturbance, while the financial contracts remain unchanged. The procedure should preserve the relation between the underlying record, its translation into a regulatory claim, and the subsequent decision, even if other project participants dispute the ecological interpretation. The question of international distribution of profits belongs to a distinct legal-economic assessment, as Section 9 explained; unresolved benefit-sharing questions should not be disguised as proof of compliance with ecological obligations.

An evaluation programme should stress-test at least five failure classes: an apparently valid permission applied outside its stated scope; a failure of animal detection mistaken for evidence of absence; a learned control policy proposing an unreviewed intervention; a changed scientific result that does not reach affected institutional users; and an economically successful operation whose ecological records reveal material unaddressed effects. Comparative tests should measure the proportion of out-of-scope actions blocked, the time and completeness of correction notifications, false restriction and false clearance rates under specified synthetic conditions, the completeness of review histories, and whether affected reviewers can understand and challenge the relevant records. These performance measures require clearly defined denominators and should not be presented as measured results in the absence of an actual evaluation.

This procedural architecture completes the move from relational and normative foundations to operational design. Its purpose is to make ecological scientific inquiry technically practicable, normatively reviewable, and receptive to correction, without implying that legality, justice, or ecological safety can be certified by one software pipeline. The next section begins a dedicated survey of existing oceanic robotic platforms and their demonstrated capabilities. Subsequent sections examine morphology, hydrodynamics, control, perception, communication, manipulation, energy, biological inspiration, and prospective marine infrastructures. Those engineering investigations will provide the empirical and technical material needed to test whether the procedures proposed here can be implemented and whether their safe-

guards perform as intended.

11 Present-Day Marine Robotic Platforms and Evidence

11.1 The evidential object of a marine robotics survey

The objective of this section is to establish which marine robotic capabilities have been demonstrated, what kinds of ecological knowledge those capabilities can support, and which conditions limit the transfer of reported achievements to other locations, missions, and organisms. The preceding sections developed a philosophical account of generative-relational inquiry and a practical account of marine nomos, authorization, review, and ecological responsibility. The engineering inquiry begins here with evidence about existing platforms. A technology that operates at sea does not thereby observe the ecological quantities relevant to a particular question; an instrument that records a quantity does not thereby establish an ecological mechanism; and a successful demonstration does not establish unrestricted readiness for persistent, autonomous or intrusive operations. These distinctions structure the survey that follows.

The marine research environment includes the sea surface and its atmosphere–ocean boundary, the water column, benthic habitats, fixed installations on the seabed, and, where a research question requires it, sub-seafloor contexts. Consequently, the term *marine robotic platform* is used for an engineered mobile or actuated system operating in relation to one or more of these environments. Robotic platforms form part of larger observing assemblages that also include moorings, cabled observatories, human-operated ships, satellite communication, laboratories, data archives, and institutions. A float that adjusts its buoyancy according to a programmed cycle and a piloted vehicle capable of collecting a specific biological sample satisfy different notions of autonomy. Treating both as “autonomous ocean robots” without describing their decision and interaction capabilities would obscure scientific differences.

We distinguish at least five kinds of evidence. *Design evidence* describes intended capabilities or specifications; *controlled experimental evidence* evaluates a mechanism under bounded test conditions; *field demonstration* documents function during an actual ocean deployment; *repeated operational evidence* documents a capability across missions and conditions; and *independently assessable scientific evidence* provides observations, uncertainties and provenance adequate for a separate research claim. These categories describe evidential functions rather than a universally ordered readiness score. An operational vehicle may have extensive records for navigation but limited independent evidence for its ecological effects; a new soft robot may have rigorous pressure-chamber tests yet a comparatively narrow set of autonomous scientific deployments. Claims about future utility should be qualified by which type of evidence supports them.

A comparative review must therefore ask what the platform can sense, where and how long it operates, which operations it can perform on the environment, how its state is estimated, what level of human oversight is available, and whether the data can be assessed independently. It must also document the logistical and ecological dependencies that make the deployment possible. These questions do not imply that greater intervention capability is intrinsically preferable to lower-impact observation. They help identify the capacities needed for a particular form of scientific encounter.

The transition from the philosophical registers to engineering evidence can be stated as a sequence of separately assessable operations. A particular platform establishes an *access condition*: a sensor can reach a spatial region and remain within a calibrated operating envelope. Its sensing and preprocessing establish an *observation condition*: a specified signal is measured with known or estimated uncertainty. A statistical or mechanistic model supplies an *inferential condition*: the signal warrants a claim about an ecological variable or process under declared assumptions. Finally, an appropriate research and institutional procedure evaluates whether generating and using that evidence is acceptable in the proposed circumstances. No earlier condition automatically validates a later one.

For example, a glider capable of repeated oxygen profiling can demonstrate depth control, endurance, and oxygen measurement performance. Interpreting a low-oxygen profile as a particular microbial mechanism additionally requires biogeochemical evidence, an appropriate model, and consideration of circu-

lation and sensor drift. Determining whether more intensive sampling is warranted requires the mission's physical footprint and legally reviewed scope. This layered comparison clarifies why an advance in controllability may be highly valuable without itself supporting a strong ecological causal claim. The same assessment logic applies to acoustic surveys, benthic visual mapping, organism encounter recordings, and remotely supervised interventions.

The engineering sections will accordingly separate *capability verification* from *ecological claim verification*. When describing a published demonstration, its platform, payload, operating regime, comparator, and evidential limits should remain visible. Where the literature establishes an instrument's technical accuracy but offers little evidence on long-run ecological impacts, the appropriate conclusion is a narrow technical result and an unresolved ecological question. Filling that gap is an independent research programme, not a transfer of certainty from the engineering trial.

11.2 Sea-surface robotic vessels and the atmosphere–ocean interface

Uncrewed surface vessels (USVs) provide a relatively mature example of oceanic robotic observation. A surface platform can use solar energy, wind or waves, or conventional propulsion, and often has more direct access to satellite communication and positioning than a submerged vehicle. Surface operations nevertheless encounter wave loading, changing winds and currents, severe weather, biofouling, collisions, navigation rules, launch and recovery constraints, and geographic limitations on reliable communications. The arrangement of energy source, propulsion system, sensors and communication links determines whether a vessel is appropriate for an extended observation mission. In particular, the absence of crew does not imply absence of supervising institutions or complete independence from shore-based operators.

NOAA documents the scientific use of Saildrone surface vehicles for air–sea measurements, marine-life observations, carbon-system measurements and seafloor mapping, including deployments in tropical cyclones [138, 140]. The scientific significance is that a mobile surface station can observe coupled atmospheric and near-surface ocean processes during conditions that can be hazardous for a staffed research vessel. NOAA's documented hurricane programme combines surface vehicles, ocean gliders and additional observations to investigate the exchange of energy and momentum between atmosphere and ocean. Such a multi-platform arrangement is more informative than treating the surface craft as an independent proxy for water-column processes [140]. Demonstrated survival in an extreme storm, however, should not be generalized into a claim that every instrument remains calibrated, that all regions can be navigated, or that one platform can resolve the complete storm–ocean dynamics.

The scientific roles of these vessels are not restricted to meteorology. Hydrographic surveying, passive acoustic observation, fishery-related sampling, and communication relay are possible where a particular platform carries the appropriate instrumentation and mission approvals. A sea-surface craft can also support a glider or profiler by relaying commands or contributing complementary measurements. These uses introduce different ecological encounters: a passive recorder, an active echo sounder, and a vessel deploying equipment may affect organisms through different mechanisms. Comparisons between robotic and crewed research should therefore account for the actual acoustic, physical, chemical and operational footprints rather than presume that a smaller or uncrewed platform is harmless.

Operational surface autonomy also has a layered meaning. A vehicle may autonomously steer toward waypoints but depend on remote scientists to select its research question, decide when to change sampling, or authorize biological intervention. Such division of tasks is often desirable, particularly when uncertain animal welfare or protected-area conditions arise. For the present paper, the important conclusion is that sea-surface robotics extends the conditions of observation while leaving scientific inference and normative authorization as separable responsibilities.

11.3 Profiling floats and buoyancy-driven gliders

Profiling floats provide one of the clearest examples of a globally distributed autonomous observation infrastructure. The Argo programme supports an international array of free-drifting robotic floats that

measure properties of the water column and transmit observations after surfacing. Conventional profiling floats obtain repeated pressure, temperature and salinity profiles; biogeochemical and deep-profiling programmes expand the kinds and depth ranges of measurements that can be collected [14, 183]. They usually are not designed to grasp objects, station-keep precisely next to an individual animal, or install instruments on the seabed. Their value arises from repeatable measurements, coverage across many sites, and common practices of instrument characterization, data handling and quality control.

Argo distinguishes rapidly released observations from subsequent delayed-mode quality-control products. Near-real-time delivery supports forecasting and operational applications, whereas expert adjustments and calibration can improve suitability for questions involving subtle long-term changes [13, 182]. This distinction offers an exemplary link to the epistemology developed in Section 6: data circulation requires not only access but also an interpretable account of correction and evidence status. A new observation may change the scientific interpretation of a previous profile; updating a dataset should not erase that dependency. Network size also cannot by itself establish spatially uniform coverage or equal observational power for all marine properties, as the programme's own status materials acknowledge geographic gaps [14].

Buoyancy-driven gliders operate through a different relation between propulsion and measurement. By changing buoyancy and using wings to convert ascent or descent into horizontal movement, Slocum and Seaglider systems can perform repeated transects with relatively low propulsion energy. The United Kingdom's National Oceanography Centre describes deployments lasting weeks to months and use of physical, biogeochemical and biological or optical sensors [125]. Their movements, sampling density, achievable station keeping and eventual routes depend on the ambient current field and available control authority. The capability to move long distances over a mission should not be equated with the ability to counter an arbitrary local current or make fine-scale observations at any selected coordinate.

Surface-wave-propelled gliders occupy an intermediate architecture: a floating surface component can carry solar-powered instrumentation and communications, while a submerged element converts wave motion into thrust [125]. This architecture makes the traditional surface-versus-submerged classification less useful than a description of energy flow, sensor placement, and mission function. Across floats and gliders, the scientific limitations include unresolved fine-scale processes, sensor drift, biofouling, variable temporal sampling and incomplete coverage of ecological relations. These platforms are central to studying physical and chemical conditions; inferring the behavior or welfare of particular marine organisms requires additional evidence and representations.

11.4 Propelled autonomous ocean vehicles and mobile scientific laboratories

Propelled autonomous underwater vehicles (AUVs) offer greater maneuverability and more directed survey capability than most freely drifting profilers. The WHOI *Sentry* vehicle, for example, is specified for scientific operations to approximately 6,000 m and supports bathymetric, acoustic, photographic and other measurements, with inertial, Doppler and acoustic navigation systems [218]. Its documented roles include investigations of hydrothermal systems and complex seafloor terrain. Such vehicles may operate close to the seafloor or in the water column, following prepared missions and, in some configurations, accepting acoustic retasking. Their autonomy usually concerns navigation and bounded mission execution, rather than unrestricted independent construction of a scientifically warranted ecological explanation.

A second engineering direction emphasizes prolonged operations and payload specialization. MBARI's long-range AUV programme reports extensive cumulative operation across a fleet and payload integrations for water sampling, plankton observation, active bioacoustics, bioluminescence and other ocean-science objectives [116]. The platform illustrates the significance of energy-aware design, onboard computation, mission software, and scientific payload interfaces. Researchers can modify the relation between mission mobility and the biological event under investigation without developing an entirely new propulsion system each time. Nevertheless, accumulated vehicle-hours are an operational evidence measure. They are not automatically evidence of the accuracy of every biological interpretation derived from those missions.

A particularly consequential development is the pairing of mobile robots with automated molecular sampling. MBARI's Environmental Sample Processor can collect and filter water, archive material for later analysis, and perform specified molecular detection procedures in situ [115]. Such instruments reduce some delays between collection and detection and may reveal transient patterns that sparse vessel-based sampling would miss. Yet the detection of DNA, genetic material or molecular markers within collected water is not equivalent to direct evidence of a living organism's current location, abundance, behavior, agency or experience. Hydrodynamic transport, degradation, assay sensitivity and contamination can alter the inferential relationship between a sample and the ecological claim drawn from it. Mobile laboratory capacity therefore increases observational possibilities while introducing additional calibration, validation and interpretation requirements.

AUV performance also depends on recovery and infrastructure. Launch conditions, mission safety envelopes, localization uncertainty, communications windows, battery reserve and surface-vessel assistance can determine which observations are feasible. In a strong or spatially variable current, a vehicle may follow a nominal track while the actual sampled water mass differs from the one intended. A recorded position can be precise in an engineering frame yet poorly matched to a biological target evolving in the flow. For this reason, mission logs should document both navigation estimates and the uncertainty or environmental assumptions relevant to the scientific interpretation. Sections 2 and 3 provide the conceptual grounds for treating those dependencies as part of a situated observation history.

11.5 Remotely operated, tethered, and hybrid intervention systems

Remotely operated vehicles (ROVs) occupy a distinct practical position because they can support real-time human observation and targeted manipulation. NOAA's *Deep Discoverer*, operated from the *Okeanos Explorer* through a tethered system, is rated to approximately 6,000 m and carries cameras, lighting, environmental sensors, sample-handling equipment and manipulator arms [137]. Operators and shore-based researchers can examine imagery and guide decisions while a dive is in progress. This arrangement enables precise interaction and rapid scientific discussion, but it also depends on ship facilities, tether management, staffing, weather, vessel availability, and the ability to interpret often ambiguous visual evidence.

Manipulation should be distinguished from free swimming. A robotic arm exerting force on a structure may destabilize the vehicle or require bracing, while contact with an animal or fragile habitat can itself alter the target of investigation. Sampling systems likewise vary in whether they suction, grip, cut, enclose or retain material, creating different evidential and welfare consequences. Established ROV operations demonstrate that marine manipulation is technically possible; they do not establish that general-purpose autonomous assembly, construction, or ecosystem-safe biological handling has been solved. The later engineering sections will examine morphology, stabilization, interaction control and construction in detail.

Hybrid vehicles investigate the practical space between tethered human supervision and untethered programmed missions. WHOI's *Nereid Under Ice* uses a lightweight fiber-optic connection for real-time communications, can maneuver far from the support ship in difficult ice-adjacent environments, and supports transitions between supervised and autonomous operation in specified configurations [215]. Its architecture reflects a task-specific negotiation among tether mass, battery power, communications, navigational risk and human oversight. An autonomous return capability is valuable in a mission that may lose a communications link; it does not imply that all uncertainties or risks have become internally diagnosable.

Stanford's *OceanOneK* further demonstrates the continued relevance of human-inspired dexterity in particular oceanic tasks. Its field programme investigated haptic teleoperation, bimanual manipulation and deep-water interaction in archaeological and exploration settings, with a demonstrated dive to approximately 852 m [172]. These achievements should be understood as evidence for a defined manipulation configuration, not a general argument that humanoid morphology is appropriate for routine marine locomotion or every engineering task. Fish-inspired propulsion, compliant tentacle-like graspers,

heavy industrial arms and remotely operated human-like hands may each serve distinct functions. The resulting research agenda concerns morphological heterogeneity rather than the selection of a universal body type.

11.6 Biologically inspired bodies and evidence of ecological encounters

Bio-inspired marine robotics is an active field in which existing demonstrations are especially important to interpret precisely. A snailfish-inspired untethered soft robot was experimentally tested at extreme hydrostatic pressure and showed fin actuation at approximately 10,900 m in the Mariana Trench; the same study reported freely swimming trials at approximately 3,224 m in the South China Sea [93]. The distinction matters. Reaching hadal depth with a functioning actuator and executing a freely swimming field trial at another depth are valuable but different outcomes. Neither result independently demonstrates repeated autonomous oceanographic sampling, environmental reconstruction, biological recognition, or large-scale marine engineering.

Soft and compliant mechanisms may distribute contact forces, tolerate deformation and exploit fluid–structure interactions, potentially enabling forms of movement and grasping unavailable to certain rigid designs. These advantages depend on actuation method, material properties, environment, control bandwidth, load, and the scientific task. Soft materials may introduce challenges in state estimation, precise manipulation, fatigue, pressure tolerance of integrated electronics, and validation after repeated deployments. “Bio-inspired” should accordingly identify a testable source of design hypotheses rather than confer a presumption of biological equivalence or ecological harmlessness.

Another approach emphasizes reducing disturbance during encounters with marine life. WHOI’s *Mesobot* is designed to observe and follow organisms in the ocean’s twilight zone, with slow propulsion, adapted lighting, extended tracking and both supervised and autonomous configurations [216]. Such a design explicitly recognizes that a robot’s own movement and illumination can change what an animal does and therefore change what a scientist observes. Reduced-disturbance design is nevertheless a hypothesis requiring context-specific validation. A quiet platform can still be conspicuous to particular animals, while an apparently unperturbed animal may experience consequences not captured by the sensors deployed.

The reciprocal significance for biology and engineering is substantial. Investigating fish locomotion, octopus coordination or benthic attachment may reveal principles of compliant propulsion and manipulation; building a robotic analogue may clarify what a proposed mechanical explanation can reproduce. However, successful robotic imitation does not demonstrate a complete account of an organism’s biological generativity. Nor does the presence of potentially useful biological mechanisms license intrusive experiments on an animal. This evidential and ethical asymmetry will inform the dedicated morphology and control analyses in Sections 12–14.

11.7 Persistent observatories, mixed fleets, and scientific circulation

Not all productive observing systems are mobile, and no single robotic platform can supply all scales of ecological evidence. Ocean Networks Canada operates cabled observatories, including NEPTUNE and VENUS, that provide power, communication and persistent observation through networks of seabed instruments, cameras, hydrophones and other sensors [143]. Such infrastructure can produce long time series and provide contextual measurements against which episodic mobile-robot observations can be interpreted. Its continuity depends on cable integrity, access to power and shore facilities, component maintenance, calibration, stable governance and continued support for data access. A permanent installation also has its own site-specific material and ecological footprint.

The combination of fixed stations, autonomous profilers, mobile vehicles, remotely operated manipulators and laboratory analysis can generate forms of evidence unavailable to each component alone. One system might detect changing oxygen concentrations; a second might image the organisms present; another might collect water for molecular analysis; and a human laboratory might refine identification.

However, these records differ in spatial support, sampling window, detection model, calibration, biological meaning and legal conditions on reuse. Combining them requires explicit metadata and task-relative translation, as developed in Section 6. A federation of platforms without adequate provenance could increase the quantity of data while obscuring the actual basis of an ecological claim.

For an engineering comparison, we propose a *platform–task–evidence dossier* rather than a single technology ranking. A dossier specifies the demonstrated activity, operational environment, sensing and intervention modes, navigation and communication dependencies, energy and support requirements, evidential status, and ecological or legal review questions. Table 8 gives a condensed comparison. Its categories describe common deployment patterns rather than capabilities guaranteed by every model.

Table 8: Representative oceanic platform families, observational strengths, and principal limitations. Performance depends on the specific vehicle, payload, deployment, and evidence.

Platform	Demonstrated or established role	Important limitation
Surface USV	Extended air–sea surveys, telemetry, surface relay	Weather and navigation restrictions; limited direct evidence about deep water
Profiling float or glider	Repeated water-column profiles and transects; distributed arrays	Coverage, calibration and drift; little ability to handle objects
Propelled AUV	Directed acoustic, optical and scientific-payload surveys	Energy, localization, support and payload-specific inference
ROV or hybrid vehicle	Human-guided inspection, sampling and manipulation	Tether or communications limits; contact effects and supervision needs
Soft or bio-inspired prototype	Selected compliant-motion and pressure-resilience demonstrations	Variable field maturity; unproven general endurance and tooling
Cabled observatory	Persistent, connected instruments and long time series	Fixed location; installation, maintenance and ecological footprint

The dossier would preserve the distinction among manufacturer specifications, research-group deployment reports, peer-reviewed experiments, released scientific datasets, and independently reproduced results. Different types of evidence may support different claims, and none should be mechanically substituted for the others. The public availability of data also varies among projects. The Argo data-distribution and quality-control arrangements and the open observing services associated with Ocean Networks Canada illustrate important practices, but neither guarantees that every ecological inference from those observations is correct [13, 143].

Resolution of platform evidence at the claim level. Comparative platform surveys become misleading when a report of *survival at depth* is silently promoted into a report of autonomous scientific operation. Li and colleagues describe a snailfish-inspired soft system whose fins were actuated during a Mariana Trench test at approximately 10,900 m and whose free-swimming test took place separately at approximately 3,224 m [93]. The former supports pressure-tolerant actuation under reported conditions; the latter supports locomotion under different conditions. Neither supplies a field evaluation of autonomous construction or organism-level welfare. Similarly, Stanford’s OceanOneK project reports bimanual haptic exploration and a particular 852 m dive; an institutional project description should be read as evidence of the reported demonstration, with independently reproducible task performance remaining a separate question [172]. These differences do not diminish either engineering achievement. They identify which claim the experiment or institutional account actually supports.

For review and future data integration, each platform assertion should be assigned a *claim-resolution pair*: the proposed property (pressure tolerance, locomotion, sensing validity, mission endurance, intervention success, or ecological response) and the strongest identified evidence type (vendor specification, research-group account, controlled laboratory test, instrumented field demonstration, field dataset, or independent replication). Evidence types are not a universal ranking across all tasks; a carefully calibrated

field dataset can be more relevant to ecological inference than a detailed propulsion bench test, whereas the latter may be decisive for a narrow actuator claim. The proposed register should include source identity, test environment, instrument calibration, comparison condition, known failure events, performance denominator, and unresolved generalization. This makes statements such as “deployed,” “autonomous” and “bioinspired” inspectable rather than treating them as transferable performance certificates.

Deployment envelopes and the status of platform specifications. A cross-platform comparison should specify the unit of evidence: a manufacturer’s rated envelope, an institution’s account of a field mission, a peer-reviewed component experiment, or a repeated autonomous scientific operation. For example, the hybrid Nereid Under Ice vehicle is described by its operator as carrying a communications-only optical microtether while drawing energy from onboard batteries; its ability to change modes does not make tethered communications, autonomous recovery, and a completed biological survey the same capability [215]. The snailfish-inspired soft robot likewise provides particularly valuable but different evidence from its approximately 10,900 m fin-actuation test and its separate approximately 3,224 m free-swimming trial [93]. A pressure rating, maximum expedition depth, instrument specification, and scientific measurement quality should therefore occupy different fields in a platform evidence register. Such a register would report the date, platform configuration, environmental conditions, human or support-vessel involvement, repeatability, calibration status, and independent availability of observations before a mission-specific ecological inference is proposed.

11.8 Technology gaps, research questions, and the transition to mechanics

Several gaps emerge from the survey. First, the ability to reach an environment is different from the capacity to make an observation that identifies a particular ecological property. Second, endurance and coverage have to be interpreted together with measurement quality and the distribution of missed phenomena. Third, locomotion and manipulation demand different mechanical capabilities; a highly efficient swimmer may be unable to apply the controlled forces required for construction or sampling, while an effective manipulator may be energetically expensive to transport. Fourth, data connectivity and autonomy remain conditioned by the geometry, currents and visibility of the oceanic environment. Fifth, biological and ecological disturbance must be measured, rather than inferred from the appealing appearance or apparent softness of a robot. Sixth, integration among platforms depends on common research practices and the ability to trace transformations of evidence across systems.

These limitations suggest several experimentally testable comparisons: energy per valid ecological observation rather than only energy per traveled distance; quality-adjusted spatial and temporal coverage rather than nominal route length; task completion under uncertain currents with and without model adaptation; manipulator accuracy and unintended contact effects under comparable tasks; correction latency and reproducibility of cross-platform conclusions; and changes in animal behavior or habitat condition relative to appropriate baselines. Measures should be reported with their denominators, selection rules, missing observations and uncertainty. The choice of such measures is itself part of scientific and normative procedure, especially when one measure can conceal costs borne by another ecological or human system.

A platform’s performance is therefore relational: it depends on the question, environment, available infrastructure, scientific representation and legitimate task. This result does not imply that engineering has no transferable physical laws or that all comparisons are subjective. Rather, it motivates defining common test conditions and the bounds within which comparison is justified. A rigorous evaluation should include failure and recovery events, unobserved disturbance, support-vessel requirements, manufacturing and maintenance demands, and ecological constraints on operations.

The next sections turn to mechanisms responsible for these differences. Section 12 investigates morphology and biological embodiment, with special attention to the heterogeneous bodies suitable for locomotion, anchoring, sensing and manipulation. Section 13 examines hydrodynamics and fluid–structure coupling; Section 14 surveys control approaches, including locally adaptive and biologically informed

strategies under only partial knowledge of the actual dynamical landscape. Subsequent sections address perception, navigation, communications, construction, materials, energy, and multi-robot scientific work. The present section establishes the empirical baseline for those analyses: oceanic engineering already comprises several capable but functionally distinct technological families, and the generation of responsible ecological knowledge requires the relations among their demonstrated capabilities, evidential limits and material consequences to remain explicit.

12 Morphology and Biological Embodiment

12.1 Morphology as an engineering and epistemic variable

The objective of this section is to investigate how the organization of robotic bodies shapes the forms of motion, sensing, contact, and scientific encounter available in marine environments. Section 11 distinguished present-day platforms according to their demonstrated tasks and evidence. The corresponding morphological problem is more fundamental: how should a body be organized when the same oceanic research programme may require prolonged surface transit, quiet movement through the water column, delicate investigation of living organisms, and stable contact with a seabed installation? An apparently versatile body can be compromised when the mechanical requirements of these activities diverge. The morphology of a robot is consequently part of the specification of a scientific operation and not merely a container for sensors and algorithms.

Here *morphology* refers to geometry, structural topology, the distribution of mass and buoyancy, stiffness and compliance, articulated or continuous deformation modes, contact surfaces, and the arrangement of actuators and sensors. It also includes the interfaces through which a robot can exchange forces and information with surrounding water, suspended objects, substrates, and human-built structures. A robot's available actions depend jointly on morphology, actuation, control, loading, and its particular environment. A fish-inspired design optimized to move through a flow field does not automatically achieve precise installation of a scientific connector; conversely, a rigid vehicle with two manipulators may accomplish a carefully stabilized intervention while being energetically inappropriate for long-duration biologically unobtrusive cruising. These are task-dependent claims, not a ranking of nature-inspired and conventional machines.

The conceptual connection to the generative-relational framework is precise but limited. A morphology influences the relations a robotic system can physically establish, the measurements it can make and the interventions it can perform. Different bodies may therefore produce different *accessible scientific encounters* even when they share the same research objective. The body does not unilaterally determine ecological reality, the interests of observed organisms, or the truth of a scientific interpretation. It conditions a practical and epistemic interface. In this sense, embodiment is relevant both to the construction of a scientific model and to the material conditions through which that model is tested.

The general robotics literature uses *morphological computation* to describe how physical properties can contribute to adaptive behavior ordinarily assigned to active control [153]. Recent surveys of marine soft robotics emphasize a related four-way coupling between locomotion, materials, distributed sensing, and control, and argue for functional mechanisms rather than literal imitation of individual animals [94, 224]. These studies motivate a systematic survey without establishing that an organism's evolved morphology, its complete regulatory system, or its ecological role can be faithfully reproduced in an engineered device.

12.2 Biological morphologies and their functional mechanisms

Marine organisms present diverse solutions to movement and interaction with fluids and substrates. Fish-like axial undulation and caudal-fin propulsion redistribute body and fluid momentum, while paired and median fins can support turning, stabilization and lower-speed maneuvering. Ray-like bodies offer another configuration, in which broad compliant pectoral fins generate traveling waves or oscillations,

Table 9: Morphological inspirations and task-conditioned limits. The cases indicate research directions rather than universal performance or ecological safety guarantees.

Morphological principle	Possible function and limitation
Fish-like body and caudal fin	Directed swimming and transit; requires tests of payload, low-speed station keeping, mission energy and disturbance.
Ray-like deformable fins	Broad-area maneuvering and distributed fluid forces; sensitive to stiffness, steering authority and scale-dependent performance.
Cephalopod-like continuum arms	Compliant reaching and contact sensing; precision loads, tool integration, repeatability and wear remain task-dependent.
Crustacean-like contact limbs	Seabed travel, bracing and tactile observation; requires habitat-sensitive traction and sediment-disturbance evaluation.
Rigid vehicles and manipulators	Tooling and substantial load transfer; reaction forces, contact damage, logistics and propulsion demand matter.
Hybrid and reconfigurable structures	Switching between movement and interaction tasks; adds mass, complexity, recovery and maintenance requirements.

with performance depending on fin geometry, deformation and environmental context. Batoid-inspired robotics has developed several such mechanisms, although a recent review emphasizes that comparisons across devices require careful matching of locomotion modes and evaluation conditions [47, 224]. The inference warranted by these biological examples is not that one type of swimming is universally optimal, but that morphology and actuation jointly select a region of possible hydrodynamic interactions.

Cephalopods introduce different functional relationships. An octopus arm can distribute contact across a deformable surface while conforming to irregular objects; suction and local sensing can support attachment and object acquisition. Liu and colleagues demonstrated an octopus-inspired soft robotic arm in which propagated bending and an approach–sense–grasp sequence reduced the need to compute an explicit full-body trajectory during particular grasping experiments [98]. Another 2026 study coupled distributed contact-sensitive suckers with peripheral control in an experimental soft arm [31]. These demonstrations support the feasibility of certain compliant interaction mechanisms; they do not establish the autonomous strength, accuracy, power budget, lifetime or functional diversity of a biological octopus. Nor is a continuum arm intrinsically appropriate for high-load fastening, sustained structural bracing or tool interfaces designed for rigid industrial wrists.

Other biological examples reveal that movement and manipulation need not be separated by the same body plan as that of a human. Crustaceans use legs and claws for substrate contact and object interaction. Crab-inspired robotic studies investigate walking, posture stabilization, and contact-dependent sensing; a field study of a legged *MR. Crab* robot documented human-supervised retrieval of inert test objects in a shallow-wave environment [57]. This evidence concerns a bounded task and location, rather than the capacity to operate indefinitely on unknown or sensitive seafloors. Some biologically inspired architectures instead use jetting, buoyancy adjustment, or flexible membranes. The 2025 *RoboNautilus* study demonstrated a cephalopod-inspired artificial siphon with integrated sensors under restricted experimental conditions, illustrating a different route to propulsion [43].

The biological source should consequently be described in terms of *mechanism, context and demonstrated transfer*. We should investigate which structures transmit forces, which sensory information affects behavior, which interactions are enabled by compliance, and what is lost when mechanisms are simplified for manufacture and control. Biological studies may also identify morphology-dependent functions that remain poorly represented in robotics. The scientific investigation of these functions should employ non-invasive observation and ethically justified intervention, rather than treat marine organisms as freely available components of an engineering design process.

12.3 Soft, rigid, and hybrid physical architectures

The mechanical distinction between rigid and soft bodies is more informative when expressed through compliance, strain limits, stiffness distribution, load-bearing capability, and response times than through a binary category. A soft fin can be mounted to a rigid sensor hull; a flexible arm may require rigid pumps or power electronics; a walking structure may combine stiff supports with compliant joints and contact feet. Materials that deform under fluid loading can store and return elastic energy, conform to irregular objects, and alter the distribution of forces. The same deformation can complicate shape estimation, introduce hysteresis or fatigue, and undermine precise positioning. Pressure, temperature, salinity, biofouling and repeated actuation may change material behavior during a mission [178, 208].

One experimentally active direction is variable stiffness. A compliant propulsor need not maintain the same mechanical response throughout a stroke or across sea states. Hu and colleagues reported an adjustable joint-stiffness mechanism for aquatic locomotion, combining a mechanics model, trajectory optimization and robotic experiments; the reported improvements were relative to specified experimental baselines and swimming conditions [65]. Such results justify investigating stiffness schedules as part of morphology–control co-design. They do not imply that all compliant robots benefit from increased flexibility or that software alone can overcome actuator, material and loading limits.

Likewise, pressure tolerance and environmental friendliness are separate properties. A highly compliant deep-sea robot may accommodate some external pressure conditions while remaining constrained by electronics packaging, connectivity, sensor payload, actuator longevity and recovery logistics. The snailfish-inspired robotic demonstration surveyed in Section 11 is important as a pressure-tolerant actuation example [93]; depth reached by a particular device is not a proxy for autonomous manipulation or for sustained marine ecological research. A rigid hull may be preferable for calibrated instruments, whereas a soft appendage may reduce a particular contact force. Neither structural classification establishes the acoustic, thermal, chemical or behavioral effects of the complete deployed system.

Hybrid architectures offer a constructive alternative to an all-purpose humanoid or all-purpose biological replica. They may contain a rigid payload core and compliant propulsion, a swimmer with detachable contact limbs, a seabed crawler with a manipulation tool, or a docking interface that allows one vehicle to work with a specialized intervention unit. The engineering issue then shifts from selecting one ideal shape to designing verified transitions between operating configurations. Such transitions may demand additional actuators, connectors, sealing surfaces, control modes and maintenance procedures, each of which can increase failure risk and ecological footprint. The capacity to reconfigure must be evaluated together with the cost of achieving and supervising that reconfiguration.

12.4 Embodied dynamics and the limits of morphological computation

Morphological computation is a useful research hypothesis when a mechanical arrangement produces favorable responses without requiring an equally complex explicit controller. Passive elastic bending, self-righting buoyancy, distributed compliance, or suction-based attachment can generate stable local behavior under bounded disturbances. Nevertheless, the term should not be used to claim that morphology somehow computes the complete marine environment or autonomously resolves a scientifically uncertain interaction. The body supplies physical dynamics that can be exploited, estimated or disrupted; their contribution is conditional on fluid loading, contact geometry, available energy and operational state.

For analytical clarity, let a candidate morphology be represented by parameters m concerning geometry, structural arrangement, material distribution, sensor placement and actuator layout. Let x denote robot state, z the locally relevant environmental state, and u an admissible control input. A schematic coupled plant can be written as

$$\dot{x} = f_m(x, u, z, w), \quad y = h_m(x, z, \nu), \quad (28)$$

where w and ν represent modeled disturbances and measurement uncertainty. Here m changes the physical dynamics f_m and the observation relation h_m ; it is not itself a control law. The local state and environmental variables require interpretation appropriate to a specified domain, and some soft-body

and fluid processes will need distributed rather than finite-dimensional representations. Section 12 uses this equation solely to identify design dependencies. Detailed marine dynamics and controller classes belong to Sections 13 and 14, respectively.

For a task q , define a *task-feasibility relation*

$$\mathcal{F}_q(m; \mathcal{Z}, \mathcal{C}) = \{(x_0, \pi) : \text{the prescribed task conditions hold under } z \in \mathcal{Z} \text{ and constraints } \mathcal{C}\}, \quad (29)$$

where π is an allowed policy and $\mathcal{Z}$ is a specified collection of environmental conditions. A nonempty feasible set establishes only task-conditioned possibilities within the model, not a guarantee for an unobserved oceanic state. *Observability* concerns whether the relevant states or task properties can be distinguished from measurements, whereas *controllability* concerns achievable state changes under a particular dynamics model; both depend on the choice of m , sensors, actuation and operating regime. Thus a body whose shape assists propulsion may impair line-of-sight imaging, mechanical isolation of a sensor, or access to an observation target. Embodiment alters what can be inferred and acted upon in different ways.

This motivates a design evaluation based on several separate quantities: task completion, propulsion energy, manipulation forces, local state-estimation error, resilience to changing loads, reliability, and measured ecological effects. Claims about a robot's embodied intelligence should specify the tested task and comparator. The same mechanical compliance that improves an approach to a delicate object may reduce the bandwidth or predictability of high-force contact. Recognizing such incompatibilities is more productive than identifying biological resemblance with overall superiority.

12.5 The locomotion–manipulation mismatch in marine construction

The distinction between locomotion and manipulation is especially consequential for proposed marine scientific stations, energy devices, sensor cables and subsea computing facilities. A swimming robot mainly manages forces needed to move its own body and payload through water. A robot installing a connector or fastening an element to a fixed structure must additionally transmit forces through external contacts, maintain alignment within tool tolerances, and resist reaction forces and currents. A free-floating manipulator applying a force or torque to an object can accelerate the vehicle body in the opposite direction unless thrusters, ballast, contacts or external supports supply an adequate counteraction. The relevant stability question is therefore a coupled vehicle–manipulator–environment problem.

Human-like two-arm arrangements may be useful when existing tools, handles and procedures were developed for human hands, as experimental oceanic humanoid-manipulation systems demonstrate [172]. They are not inherently the best marine locomotion forms. Conversely, an elongated fish-like body may traverse the water column efficiently but offer no stable work surface or reaction-force pathway for structural assembly. Octopus-inspired arms may conform to irregular biological samples or cables while struggling with accurately calibrated tightening and heavy rigid components. The design should depend on the task's required *contact wrench*, access geometry, tool compatibility, positioning accuracy and permissible disturbance, not only on the robot's swimming speed.

One possibility is to separate tasks across a heterogeneous work system. Survey swimmers can map a site, surface vehicles can relay commands or contribute positional references, specialized manipulators can handle fragile instruments, anchored robots can perform precision tool work, and ships or dedicated heavy equipment can lift and position large modules. A data-center demonstration, for instance, would more plausibly fabricate and seal major modules in controlled facilities, transport them to the site, and deploy marine robots for inspection, alignment, connection or maintenance than construct every component from raw materials on the seabed. Whether a particular installation is permissible and environmentally acceptable remains a distinct determination developed in Sections 7–10.

Task specialization also creates costs. Separate machines may require rendezvous, force-transfer fixtures, compatible docking mechanisms, recovery planning, and greater total material use. Heterogeneity is not automatically efficient or ecologically preferable. The scientific question is whether dividing tasks improves measured capability and safety after the complete logistic and environmental system is included in the comparison.

12.6 Morphological portfolios and task-relative design

To formalize the preceding argument, consider a finite family of tasks $Q = \{q_1, \dots, q_n\}$ and candidate morphologies $M = \{m_1, \dots, m_k\}$. Define a binary relation $R \subseteq Q \times M$ where $(q_i, m_j) \in R$ means that a specified body and admissible control system have demonstrated or conditionally established the requirements for task q_i in a declared environment and evaluation regime. An associated *coverage set* is

$$\text{Cov}(M; Q, R) = \{q \in Q : \exists m \in M \text{ with } (q, m) \in R\}. \quad (30)$$

This relation is deliberately typed by tasks and evidential conditions; it does not imply an overall ranking of scientific or ethical value.

Proposition 12.1 (Conditional task coverage under an expanded portfolio). *For a fixed task family Q and fixed feasibility relation R , if $M_1 \subseteq M_2$, then $\text{Cov}(M_1; Q, R) \subseteq \text{Cov}(M_2; Q, R)$.*

Proof. Every task covered by M_1 has at least one feasible morphology in M_1 . That morphology also belongs to M_2 , so the same task belongs to the coverage set of M_2 . $\square$

The proposition states only a set-theoretic property. It assumes independent availability of candidate systems, an unchanged feasibility relation and no constraints coupling deployments. It says nothing about financial cost, equipment availability, ecological disturbance, conflicting communication channels, maintenance burdens, authorization, or collective failure probabilities. If the addition of another morphology changes the logistics, collision risks, ecological exposure or legal scope of a real mission, the fixed-relation premise may fail and net mission performance may decrease. The proposition is included to make precisely this boundary explicit.

An actual portfolio decision therefore requires a richer assessment, for example

$$\text{PortfolioReview}(M, \Pi; \omega) = (\text{TaskEvidence}, \text{Energy}, \text{Reliability}, \text{EcologicalEffects}, \text{Authority}, \text{Recoverability}), \quad (31)$$

where Π specifies permitted task allocation and coordination, and ω denotes a represented operating context. No automatic sum across these differently typed assessments is assumed. A project may reject a system with superior task coverage if it entails unacceptable habitat contact or unauthorized intervention. The design problem thus connects the engineering question of heterogeneity to the procedural and ecological boundaries developed in earlier sections.

12.7 Biology-to-robotics inquiry and ecological responsibility

Biological inspiration can support a reciprocal research process. Scientists may document a marine organism's deformation, propulsion, sensing or grasping behavior; construct a constrained robotic analogue; compare the analogue and biological observations under matched conditions; and revise the explanation when the results differ. The robot can then function as an experimental instrument for discriminating among mechanical hypotheses. A successful engineered analogue is evidence that a mechanism can produce a certain outcome, but it is not proof that an organism uses exactly that mechanism or control representation. Multiple mechanisms may generate similar observable behavior. Such evidential modesty is especially important when the biological system possesses complex neural, sensory and developmental organization that the mechanical analogue omits [94, 98].

The choice of experimental organism and observation method also carries normative consequences. Close approach, artificial illumination, flow generation, sampling, attachment and repeated exposure may alter animal behavior or local habitat conditions. Softness of a robot does not establish that it is benign: compliant bodies can still entangle, obstruct, cause abrasion, release material, transmit sound, or disturb sediment. Tests of ecological compatibility should measure specific organisms' responses and physical effects under explicit baseline and control conditions; where adequate evidence is unavailable,

the responsible action may be to reduce exposure, seek additional assessment or abstain. These obligations follow the distinctions between observable response, possible sentience, procedural review and ecological justice developed in Sections 4 and 5.

The marine environment is a valuable source of biological knowledge but not a limitless repository of freely usable engineering solutions. Research into living morphologies does not justify unbounded extraction or manipulation. Existing biological knowledge, community contributions and instrument designs may also be governed by different access conditions and intellectual, scientific or collective responsibilities. This section does not attempt to establish a universal regime for biological inspiration; it identifies why its use in robot design must remain connected to the institutional record of scientific permissions, claims, risks and benefits.

What biological analogy can establish. Recent reviews organize bio-inspired systems by propulsion, morphology, distributed sensing, and control, and survey distinct ray-like propulsion and actuation architectures [94, 47]. Their comparative tables are useful for generating mechanical hypotheses, but the reported platforms differ in body length, fin geometry, excitation, tethering, actuator technology, and test medium. Even normalization by body lengths per second does not equalize energetic efficiency, sensing, payload, station-keeping, or response to living organisms. Biological analogy can motivate a controlled variation of stiffness, stroke phase, flow velocity, or contact geometry. A specific claim about energetic superiority or disturbance reduction would then require matched apparatus and measurement units, along with acoustic, flow and biological observations appropriate to the claim. This distinction also preserves the direction of inference: reproducing a swimming trajectory is not sufficient evidence for reproducing an animal's sensory experience or ecological functions.

12.8 Research agenda and transitions to dynamics and control

The preceding survey supports several experimental comparisons for future work. First, designs inspired by similar biological functions should be compared under matched environments, sensor payloads, mass accounting, power measurement boundaries and operating objectives. A fish-inspired prototype should not be declared more efficient than a conventional AUV because one reports body lengths per second and the other reports mission energy consumption with a larger payload. Second, soft and rigid manipulators should be evaluated under common object geometries, contact forces, sensing uncertainty, water motion and tool tolerances. Third, reconfigurable designs should be compared with specialized robot teams while counting the energy, logistics and maintenance required to change configurations. Fourth, welfare and ecological effects should be measured separately from task success, with attention to species-specific and cumulative responses.

A useful experimental programme would define a cross of *morphology*, *task*, and *environment*, then publish feasibility failures as well as successful demonstrations. Representative tasks might include transit in variable currents, close-range observation of a mobile organism, gentle retrieval of an inert surrogate, tool placement against a fixed frame, and recovery from a failed contact. The corresponding environments may include controlled tanks, current-generating facilities, monitored coastal sites and more demanding ocean conditions when appropriate permissions and safety preparations are in place. Evaluation should report whether an outcome is shown in simulation, controlled experimentation, field demonstration or repeated field use, preserving the evidential distinctions established in Section 11.

This also establishes an open problem in the interpretation of biological mechanisms. A robot inspired by a fish's shape may lack its sensory feedback, distributed muscle dynamics, metabolic energy supply and evolutionary context. A comparison may therefore test a single selected mechanism while leaving the whole organism's generative capacities unrepresented. For this reason, we propose *functional abstraction with explicit loss accounting*: specify the biological mechanism being hypothesized, identify which variables and couplings the engineering implementation preserves, and report which aspects are omitted or unresolved. This provides a connection to generalized translation in Section 6, where loss of task-relevant structure may invalidate an inference across systems.

The section’s principal conclusion is that a body is a conditional interface for action and inquiry. A productive marine robotics research programme will likely require several bodies, each with different relations to flow, contact, sensing, manipulation and habitat. Yet no portfolio design, biological inspiration or material compliance alone establishes control under an unobserved oceanic dynamical landscape. Section 13 therefore investigates the mechanical and hydrodynamic processes that shape these interactions, while Section 14 examines how conventional, adaptive, learned and biologically informed controllers can respond within locally justified limits. The eventual engineering objective is a system whose morphology and control improve scientific capabilities while remaining accountable to the ecological and institutional conditions of each mission.

13 Hydrodynamics, Fluid–Structure Interaction and Mechanical Models

13.1 The coupled dynamical landscape of oceanic operations

The objective of this section is to identify the physical models and empirical limits through which marine robotic systems can be described as dynamical systems. Section 12 established that morphological configurations create different possibilities for propulsion, contact, sensing, and scientific intervention. These possibilities are realized in an environment whose motion, stratification, chemistry, and biological processes may change during the operation itself. A scientifically responsible mechanical description must therefore specify both the body and the surrounding domain, together with the assumptions under which their interaction is approximated. This section addresses the models and their limits; the design and comparison of feedback controllers are reserved for Section 14.

Oceanic motion shares fundamental continuum-mechanical principles with atmospheric flow and other fluid systems. Navier–Stokes dynamics, conservation laws, rotating-frame effects, buoyancy, and turbulent transport remain relevant. Their practical expression differs because water has markedly different density and viscosity from air, marine vehicles are often nearly buoyancy-balanced, density stratification can constrain vertical exchange, surface waves and seabed boundaries alter loading, and navigation sensors may provide only intermittent or spatially limited environmental information. These contrasts depend on the platform. A sea-surface autonomous vessel, a buoyancy-driven glider, a compliant swimmer near an organism, and a seabed-contact manipulator do not encounter the same effective dynamical plant.

The ocean is also a coupled system of systems. Lunar–solar tidal forcing, wind stress, changes in heat and freshwater fluxes, Earth’s rotation, bathymetry, and the marine cryosphere contribute to circulation across different spatial and temporal scales. Temperature, salinity, and pressure shape density; physical flows transport heat, nutrients, organisms, and suspended particles; biological production and decomposition can change chemical concentrations and optical conditions [132, 124, 112, 111]. Many useful ocean models prescribe, average, or parameterize some of these processes. Their successful use does not imply that a robot has access to a complete contemporaneous state, still less that all possible continuations of the ocean have been determined.

For this inquiry it is useful to distinguish the *environmental evolution*, the *vehicle–fluid mechanical response*, the *scientific observation map*, and the *operationally admissible action set*. Each has its own evidence requirements. The environmental model may forecast a current field, while the vehicle model predicts how that current affects motion. A sensor model determines what is actually observed, while applicable ecological and institutional conditions constrain the actions that may be attempted. Agreement at one of these levels cannot be transferred automatically to the others. In particular, a numerically accurate velocity predictor is not evidence that an intervention is harmless to marine life or legally authorized.

13.2 Astronomical forcing, stratification and multiscale circulation

The gravitational configurations of the Moon and Sun make major contributions to ocean tides. Rotation, basin geometry, coastline shape, water depth, friction, and resonance influence the resulting local tidal elevations and currents. The term *astronomical forcing* in this context concerns these specific gravitational

Table 10: Representative dynamical representations for oceanic robotic inquiry. Each model family has a distinct evidence domain; listed limitations are examples rather than exhaustive conditions.

Model family	Appropriate research use	Important limitation
Tide and current products	Route planning; expected flow windows	Coastal gradients and local transients
Hydrostatic circulation models	Regional stratification and large-scale transport	Small nonhydrostatic or turbulent processes
Local CFD and turbulence models	Hull loads, vortices, wake and seabed contact	Boundary conditions and resolution dependence
Six-DOF vehicle models	Maneuvering, buoyancy and propulsion estimation	Unidentified coefficients and spatial current shear
Compliant-body FSI models	Flexible fins, arms and distributed loads	High-dimensional material and fluid uncertainty
Reactive transport models	Chemical gradients and biological transformations	Unknown sources, reactions and biological states

interactions, and should not be read as attributing operationally comparable tidal effects to the arrangement of distant stars. In coastal passages and near seabed features, tidal currents can differ markedly from a horizontally uniform periodic signal. Surface wind stress, mesoscale eddies, density gradients, and local turbulent mixing create other kinds of flow variability. NOAA distinguishes currents driven by tides, winds, and thermohaline circulation; MIT ocean-modeling research emphasizes interacting scales from internal waves to basin circulation [132, 111].

A common continuum idealization describes water motion by a velocity field $\mathbf{u}_f(\mathbf{r}, t)$, pressure $p(\mathbf{r}, t)$, and density $\rho(\mathbf{r}, t)$ in a specified domain. For an incompressible Boussinesq approximation in a rotating local frame, a schematic momentum and mass balance is

$$\partial_t \mathbf{u}_f + (\mathbf{u}_f \cdot \nabla) \mathbf{u}_f + f \mathbf{e}_z \times \mathbf{u}_f = -\rho_0^{-1} \nabla p' + b \mathbf{e}_z + \nabla \cdot \mathbf{T} + \mathbf{F}, \quad (32)$$

$$\nabla \cdot \mathbf{u}_f = 0, \quad b = -g \frac{\rho - \rho_0}{\rho_0}. \quad (33)$$

Here f is the local Coriolis parameter, $\mathbf{e}_z$ points upward, ρ_0 is reference density, $\mathbf{T}$ collects the represented dissipative stresses per unit mass, and $\mathbf{F}$ denotes additional modeled forcing. The pressure perturbation p' is defined relative to a hydrostatic reference state. Equations (32)–(33) are an illustrative reduced representation, not a claim that all conditions admit one closed Boussinesq solution. Specific hydrostatic, nonhydrostatic, compressible, free-surface, and boundary formulations require their own assumptions. MITgcm provides documented implementations and distinctions among such formulations [112].

Density is affected by temperature, salinity, and pressure. Stratification may inhibit vertical exchange while permitting internal gravity waves, localized shear, and mixing when particular conditions arise. Relevant internal-wave motions can alter an instrument's effective depth, change the local flow experienced by a vehicle, and interact with a vehicle wake. Studies of stratified vehicle wakes describe features including internal-wave radiation and horizontally spreading eddies, demonstrating why homogeneous-water tests may miss important operational effects [21, 213]. These effects are contingent on conditions and measurement methods, not universal characteristics of all submerged vehicles.

Spatial and temporal scales matter operationally. A coarse forecast may be appropriate for selecting a regional route, while station keeping next to a reef or a connector demands local estimates. A favorable mean flow can coexist with short-duration turbulence or near-boundary jets. Conversely, a local sensor reading may not reveal a regional transport process relevant to an ecological interpretation. The measurement scale and the time interval over which a model is evaluated should therefore be reported alongside its performance.

13.3 Rigid-body hydrodynamics, relative motion and model reduction

For a rigid marine robotic platform, the familiar six-degree-of-freedom formulation provides a useful starting point. Let $\boldsymbol{\eta}$ represent position and attitude, and $\boldsymbol{\nu}$ the body-fixed generalized velocity. Under an appropriate coordinate chart,

$$\dot{\boldsymbol{\eta}} = J(\boldsymbol{\eta})\boldsymbol{\nu}. \quad (34)$$

For the case in which a current-relative representation is justified, one may write the more explicit schematic dynamics

$$M_{\text{RB}}\dot{\boldsymbol{\nu}} + C_{\text{RB}}(\boldsymbol{\nu})\boldsymbol{\nu} + M_A\dot{\boldsymbol{\nu}}_r + C_A(\boldsymbol{\nu}_r)\boldsymbol{\nu}_r \quad (35)$$

$$+ D(\boldsymbol{\nu}_r)\boldsymbol{\nu}_r + g_h(\boldsymbol{\eta}) = \boldsymbol{\tau} + \boldsymbol{\tau}_{\text{ext}}, \quad (36)$$

where $\boldsymbol{\nu}_r = \boldsymbol{\nu} - \boldsymbol{\nu}_c$ is the velocity relative to the modeled current. The matrices M_{RB} and M_A represent rigid-body and added-mass contributions, C_{RB} and C_A their associated inertial terms, D hydrodynamic damping, g_h weight–buoyancy restoring terms, and $\boldsymbol{\tau}$ applied actuation. Fossen derives forms of this model and carefully specifies assumptions concerning ambient currents [45, 46]. In particular, the placement of $\dot{\boldsymbol{\nu}}_r$ and the representation of external flow cannot be exchanged casually with a zero-current model.

Added mass expresses the additional fluid inertia involved when a body accelerates in water. It depends on geometry, surrounding boundaries, and modeled flow conditions. Drag depends on relative motion, Reynolds number, surface state, and the selected empirical approximation; common quadratic laws are regime-dependent rather than complete descriptions of all marine forces. Hydrostatic restoring effects depend on buoyancy, weight distribution, center of gravity, and center of buoyancy. When a vehicle manipulates a large object, takes on a sample, or changes buoyancy, these properties may themselves change. Severe current shear, complex wakes, contact, and deforming appendages can invalidate simple coefficient models.

Different engineering decisions justify different reductions. Three-degree-of-freedom planar maneuvering models can be useful for certain sea-surface missions; a longitudinal glider model may include buoyancy and pitch dynamics; six-degree-of-freedom models help describe free-moving survey vehicles. Local computational fluid dynamics (CFD) and experiments are needed when hull appendages, wakes, and boundary effects determine performance. Reviews of CFD approaches distinguish Reynolds-averaged simulations, large-eddy methods, and hybrid methods, each with accuracy, computational, closure, and experimental-validation trade-offs [60, 64].

A model is useful insofar as it can answer a specified engineering question under documented conditions. Reproducing a tank maneuver at one speed does not automatically establish accuracy for vertical shear in a different ocean, for a robot carrying new instruments, or for operation near sensitive animals. The domain of identified parameters and possible model failure should be part of the model record. These concerns provide a physical motivation for the controller adaptation examined in the next section, without allowing control algorithms to substitute for poorly characterized mechanical dynamics.

13.4 Compliant bodies and two-way fluid–structure interaction

Rigid-body models are often inadequate for a robot whose propulsion or manipulation depends on large, distributed deformation. Flexible fins, undulating soft bodies, tentacle-inspired arms, and compliant contact surfaces exchange momentum with the surrounding fluid as they change shape. In these systems, body deformation modifies hydrodynamic loading, which in turn changes deformation: the coupling is two-way. Recent literature surveys identify unsteady vortices, flexible propulsion, added mass, and material architecture as central engineering considerations [80, 94]. The goal of such models is to explain mechanisms and predict task-relevant behavior, while recognizing the substantial cost of inferring geometry-dependent parameters.

Let $\mathbf{d}_s(\mathbf{X}, t)$ be the displacement of a deformable solid reference configuration, with density ρ_s , and let $\mathbf{u}_f$ describe the surrounding fluid. A schematic continuum solid balance and moving fluid–solid

interface conditions are

$$\rho_s \partial_{tt} \mathbf{d}_s = \text{Div}_{\mathbf{X}} \mathbf{P}_s + \mathbf{b}_s, \quad (37)$$

$$\mathbf{u}_f(\chi_s(\mathbf{X}, t), t) = \partial_t \mathbf{d}_s(\mathbf{X}, t), \quad \mathbf{X} \in \Gamma_0, \quad (38)$$

$$\boldsymbol{\sigma}_f \mathbf{n}_f + \boldsymbol{\sigma}_s \mathbf{n}_s = \mathbf{0} \quad \text{on } \Gamma(t). \quad (39)$$

Here Γ_0 is the material reference interface and $\chi_s(\mathbf{X}, t) = \mathbf{X} + \mathbf{d}_s(\mathbf{X}, t)$ maps a material point to its current position on $\Gamma(t)$. The first equation uses a reference mass density ρ_s , a first Piola–Kirchhoff stress $\mathbf{P}_s$, and body force $\mathbf{b}_s$ per reference volume. The last equation uses Cauchy stresses $\boldsymbol{\sigma}_s, \boldsymbol{\sigma}_f$ and consistently oriented normals on the *current* interface. In the absence of interfacial slip, Equation (38) matches velocities at corresponding material and spatial positions, rather than evaluating a reference displacement at an undefined current coordinate. Contact, porosity, suction, or a specified interface constitutive law may require modified conditions. The formulas provide an admissible modeling template, not a new universal constitutive law for soft marine bodies. A modern immersed Lagrangian–Eulerian formulation similarly makes the relation between Eulerian fluid variables, a Lagrangian motion map, and interface continuity explicit [86].

Two quantities help distinguish regimes. The Reynolds number $\text{Re} = UL/\nu_f$ compares inertial and viscous effects using characteristic speed U , length L , and kinematic viscosity ν_f . For periodic propulsion, the Strouhal number $\text{St} = f_a A/U$ uses actuation frequency f_a and an explicitly chosen oscillation amplitude A (here taken as peak-to-peak excursion). Their values do not independently determine performance; comparisons also depend on fin compliance, phase, geometry, loading, and the surrounding flow. When the robot interacts with a substrate or grasped object, ordinary free-swimming measures may cease to characterize the task.

This distinction is central to a heterogeneous marine fleet. A compliant swimmer can exploit body–flow interactions for travel, but an arm performing subsea assembly may need adjustable stiffness, force sensing, stable attachment, and reaction-force management. Conversely, a rigid intervention system may satisfy a precision manipulation specification while imposing an undesirable ecological footprint for quiet long-duration observation. The same engineering programme can therefore need several coupled models, each corresponding to a task and body type, instead of a single biomimetic dynamical representation.

13.5 Chemical transport, biological processes and observational feedback

Marine robotic inquiry often measures quantities that evolve through both transport and local production or consumption. Oxygen, nutrients, dissolved inorganic carbon, plankton concentrations, particles, and environmental DNA are affected by currents and mixing as well as chemical reactions, microbial activity, organism behavior, and the material conditions of sampling. A schematic scalar concentration model is

$$\partial_t c + \mathbf{u}_f \cdot \nabla c = \nabla \cdot (K_c \nabla c) + R_c(c, \boldsymbol{\xi}; \vartheta) + S_c, \quad (40)$$

where K_c is an admissible effective diffusivity tensor, R_c denotes modeled reaction or biological transformation, $\boldsymbol{\xi}$ represents other relevant states, ϑ model parameters, and S_c specified sources or sinks. The decomposition is model-dependent and may omit particulate, organism-mediated, or nonlocal transport. Equations such as (40) can be embedded in coupled ocean models; they do not mean that one local chemical sensor can infer the full biological or chemical state [112, 111].

Biological activity can create materially important feedback. Primary production changes carbon and nutrient pools; respiration and decomposition consume oxygen, and biological movement can alter where substances are concentrated or deposited. Investigations of ocean oxygen minimum zones show why transport histories, organic matter supply, and microbial transformations need joint consideration rather than exclusively mechanical interpretation [201, 119]. Some biological changes are gradual relative to robotic motion, while a localized bloom or aggregation may alter observations on a mission-relevant time scale. The adjectives *fast* and *slow* are therefore relative to the research process and the particular

variable under study; they are not inherent categories into which all physical and biological dynamics can be permanently divided.

The robot may also change what it is trying to observe. Thruster wash can resuspend particles, lights may alter some animals' behavior, and sampling may affect local chemical or biological conditions. An oxygen or particle reading recorded after vehicle arrival may consequently differ from an otherwise comparable unperturbed measurement. We need not regard every measured difference as an effect of the observer, but we must be able to formulate the alternative and, where possible, distinguish it through controls, sensor placement, staggered observations, or independent instruments.

At the institutional level, this creates a provenance requirement. A data product should link its measurement method, instrument position and motion, calibration, modeled ambient currents, possible local disturbances, and the assumptions used in subsequent ecological inference. The relationship between these fields may be represented through the paper's partial relational histories; this need not presuppose that one centralized database reconstructs all causal relations. Section 9 has already argued that publishing a number and transmitting its meaning and provenance are distinct operations.

13.6 Partial observation, local mechanical models and tangent validity

A robot acting in a particular region receives finite, noisy, and modality-specific observations. State estimation may fuse inertial, acoustic, optical, chemical, current, pressure, or external model information, but even a good estimator cannot recover every unresolved eddy, biological interaction, or chemical history. The technical problem is accordingly one of selecting models with a declared region of validity, identifying their uncertainty, and revising them when observations justify reconsideration. Our use of *tangent-space governance* here refers to local model-based reasoning in an explicitly specified state space. It is a proposed interdisciplinary synthesis, not an established general oceanographic theorem.

Let $x \in \mathbb{R}^n$ denote an appropriately chosen local state coordinate, $u \in \mathbb{R}^m$ an actuation or decision variable, and $z \in \mathbb{R}^p$ a finite representation of relevant environmental factors. A local mechanical hypothesis may be written

$$\dot{x} = F(x, u, z) + w, \quad y = H(x, z) + \epsilon, \quad (41)$$

where w represents dynamics not resolved by F , and ϵ represents observation error. The environmental representation z may be updated using marine measurements, rather than interpreted as the full state of the sea. Let $\zeta = (x, u, z)$ and fix a reference point ζ_0 . Provided F has the requisite smoothness locally, an expansion gives

$$F(\zeta_0 + \delta) = F(\zeta_0) + DF(\zeta_0)\delta + \mathcal{R}_F(\delta). \quad (42)$$

The usefulness of this statement turns on bounding its remainder and explaining what it excludes.

Proposition 13.1 (Local approximation remainder under bounded curvature). *Let $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^n$ be twice continuously differentiable on an open set containing the closed line segment between ζ_0 and $\zeta_0 + \delta$. Assume that the operator norm of its second derivative satisfies $\|D^2F(\zeta)\| \leq L$ at every point on the segment. Then the remainder in Equation (42) satisfies*

$$\|\mathcal{R}_F(\delta)\| \leq \frac{L}{2} \|\delta\|^2. \quad (43)$$

Proof. Apply the fundamental theorem of calculus to $t \mapsto F(\zeta_0 + t\delta)$ and subtract $DF(\zeta_0)\delta$. The remainder is $\mathcal{R}_F(\delta) = \int_0^1 (1-s) D^2F(\zeta_0 + s\delta)[\delta, \delta] ds$. Taking norms, using the uniform second-derivative bound, and integrating $1-s$ over $[0, 1]$ gives Equation (43). $\square$

The proposition is an elementary local Taylor estimate, included to articulate an engineering condition rather than claimed as a novel mathematical discovery. It bounds approximation error for the *specified* finite-dimensional F along the segment. It does not bound w , observation error, parameter estimation, poorly represented ecological states, or long-horizon trajectory divergence. If a Jacobian $\widehat{DF}$

and intercept $\hat{F}_0$ are estimated with respective errors ε_J and ε_0 , then a corresponding one-step vector-field bound additionally includes $\varepsilon_0 + \varepsilon_J \|\delta\|$. A viable implementation must estimate or conservatively bound these quantities, specify the chart and norms, and recognize where differentiability or the fitted model fails. Since $\zeta = (x, u, z)$ may join positions in metres, velocities in metres per second, actuator inputs, and environmental quantities in unrelated units, the norm in this estimate requires a stated coordinate scaling. A diagonal or block scaling $\tilde{\zeta} = S^{-1}(\zeta - \zeta_0)$ can express admissible deviations dimensionlessly before estimating the second-derivative bound; the resulting L and residual norm refer to that scaled coordinate system. Applying an unqualified Euclidean norm to mixed physical units would make the numerical constant meaningless.

In particular, neither a small local remainder nor a low short-term tracking error justifies extrapolation to an unobserved coastline, a new ecological regime, or an intervention in which contact changes the mechanics. A local tangent model may support short-horizon motion selection while leaving ecological admissibility unresolved. That distinction is why operational decision procedures in Section 10 were kept separate from technical performance.

13.7 Transitions, failure of local models and ecological criticality

An oceanic system may encounter changes that invalidate a controller's mechanical approximation: a change in current shear, contact with a structure, departure from a near-equilibrium buoyancy regime, deformation beyond the material calibration range, or an unexpected sensor disturbance. Some such events correspond to recognized mathematical transitions; others are operational boundary crossings without any system bifurcation. It is scientifically important to maintain this distinction. A sharp increase in the residual between predicted and measured motion may indicate model mismatch, sensor error, altered loading, or actual environmental change. Without further tests, it does not uniquely identify an ecological regime shift.

The corresponding concept of *criticality governance* is best formulated as a requirement to investigate and respond to changes in the conditions supporting a model and a mission. In physical models, one may study stability margins, eigenvalues of linearized dynamics, finite-time growth of perturbations, observable changes in operating regimes, and safe recovery regions. In ecological inquiry, proposed thresholds may concern animal behavior, habitat disturbance, dissolved oxygen, or cumulative effects. The evidence and normative role of a threshold must be identified independently. A mathematical instability indicator does not authorize an experiment; nor does a precautionary ecological limit have to be a bifurcation point in the underlying physical system.

A hypothetical reef-observation vehicle makes the distinction concrete. A regional circulation model predicts a moderate current at a planned location. On approach, onboard velocity and pressure measurements indicate intermittent flow and a substantial discrepancy from the model, while optical data become degraded by resuspension. This evidence may justify updating the local disturbance estimate, selecting a more conservative route, moving the vehicle away from the substrate, or pausing the sampling programme. It does not justify an unsupported assertion that a regional circulation transition has occurred. If independent animal-response monitoring indicates concern, that information should affect the operational admissibility review, not merely the hydrodynamic tracking objective.

A second case concerns a seabed manipulator installing a scientific connector. The force required for assembly can excite vehicle motion, affect contact forces, and disturb sediments. The effective plant switches between free motion, tentative contact, braced manipulation, and withdrawal. A single smooth global model may poorly represent all these hybrid modes; distinct mechanical models and guarded transitions may be preferable. The research problem is to establish when switching conditions can be identified and how a supervisory system can respond under incomplete evidence. The design of that supervisor belongs to Section 14; the present section establishes its mechanical justification.

Range of authority of hydrodynamic models. Fossen's marine-craft framework offers established kinematic and dynamic representations for specified craft models, including hydrodynamic dependence

on velocity relative to the surrounding water [45, 46]. The convenient representation of a locally uniform, approximately irrotational current should not silently be extended to strong spatial shear, rapidly evolving stratification, deforming fins, or disturbed sediment. Conversely, reviews of compliant bio-inspired mechanisms identify added mass, unsteady vortex shedding, and material–flow feedback as important design mechanisms, not universally calibrated force laws [80]. The relevant verification question is which measured range of Reynolds number, geometry, wave state, depth, current gradient and deformation amplitude supports a given reduced model. Whenever a simulation is used to motivate ecological sampling or close-range intervention, the model discrepancy must be reported separately from sensor error and numerical convergence. A correct numerical solution of a poorly transported flow model cannot establish an accurate local wake or exposure estimate.

13.8 Model validation and implications for subsequent control research

The principal engineering output of the preceding analysis is a hierarchy of physically justified models with explicit limits. Validation should match the form of the claim. Tank tests can estimate hydrodynamic coefficients under specified fluids, geometries and maneuvers. Flow visualization and particle-image velocimetry can expose wakes and boundary phenomena; towing tests and field comparisons can constrain maneuvering models. Flexible-body experiments can measure deformation, force, energy, and repeatability. Ocean model validation requires comparison with independent observations at relevant spatial and dynamical scales. Field tests remain essential when local shear, wave forcing, biological material, or seabed contact dominate the task [45, 60, 227, 80].

For this paper, future experiments should compare at least: a nominal simplified rigid-body plant; a current-relative model with identified coefficients; a spatially varying-flow scenario; a compliant-body fluid–structure model; and a hybrid contact case. Each comparison should report state-estimation quality, one-step prediction error, sustained prediction error, energy or effort, and model-validity violations. Those technical metrics are not substitutes for disturbance measurements, animal-welfare review, or legal authorization. A later demonstration should include unmodeled changes and adverse cases as well as favorable trials. It should also report failures of identifiability and cases where safe operation requires abstention.

This section has not established that a particular adaptive or learning-based controller can robustly control marine ecology, nor that the proposed local validity tests are sufficient for operational safety. It has instead specified why different bodies and tasks require different plant models, why oceanic forcing and biological processes can exceed simplified models, and when a local approximation may be justified. The next section can now compare classical, adaptive, learned, bio-inspired, and supervisory controllers against these physical and epistemic constraints. Its central question is how a robotic system should select an action when neither its current representation nor its accessible observations justify treating the marine dynamical landscape as globally known.

14 Control Architectures for Oceanic Robotic Systems

14.1 Control across heterogeneous oceanic dynamical regimes

The purpose of this section is to examine how oceanic robots can regulate movement and interaction while responding to an environment that is incompletely observed, dynamically heterogeneous, and ecologically inhabited. Section 13 developed the mechanical distinction between global marine flows, local vehicle–fluid models, deformable-body interactions, and operational transitions. Control now concerns the actions that can be justified under such models, the observations needed to revise those actions, and the conditions under which a controller should suspend its current objective. Its theoretical vocabulary can be shared with aviation, mobile manipulation, and terrestrial robotics. The relevant differences are principally the dynamical plant, the sensing and actuation architecture, and the constraints under which the robot operates [45, 36].

For a rigid marine vehicle, standard guidance, navigation, and control can decompose a mission reference, a state estimate, and feedback actuation. A sea-surface vessel encounters wind, waves, collision regulations, and surface-sensing opportunities. A submerged AUV encounters depth-dependent current and buoyancy forces, uncertain localization, and often intermittent communication. A compliant fish-inspired vehicle generates thrust through a body–fluid interaction in which gait and displacement are coupled. A seabed construction robot may pass between free motion, contact, anchoring, and forceful manipulation, changing the effective constraints and contact modes. These are different plants even where the same feedback mathematics applies. A controller useful for regulating course on a surface vessel cannot be presumed to stabilize a highly deformable manipulator or to apply admissible contact forces near a sensitive organism.

An additional distinction concerns the *controlled objective*. Heading and depth tracking, efficient propulsion, station keeping, contact force, habitat observation, avoidance of animals, and restoration of a damaged mission state have different measurements and success conditions. Low tracking error is insufficient evidence of successful ecological inquiry; a vehicle may follow a reference accurately while disturbing a habitat or inducing an animal’s avoidance response. Conversely, a robot that withdraws from a planned sampling path may have executed the more appropriate decision under uncertainty. This is consistent with the separation of technical feasibility, scientific warrant, institutional authority, and ecological admissibility developed in Sections 5 and 10. The present section reviews existing control methods before developing a conditional research architecture connecting local adaptation to operational criticality.

14.2 Classical feedback, robust regulation, and adaptation

Conventional feedback offers an essential baseline. Proportional–integral–derivative control, cascaded attitude and position loops, and line-of-sight guidance remain attractive because their implementation and tuning can be relatively transparent. Rigid craft can also use model-based feedback linearization, backstepping, passivity-based design, or sliding-mode methods where relevant assumptions hold. Controllers of these families have long been studied for marine guidance and tracking; they should be judged against specific tasks, actuator limits, modeling errors, and verified operating envelopes [45, 36]. The important comparison is with a competent tuned baseline, rather than with an intentionally inadequate controller.

Let x represent a locally chosen controlled state, u the allocated actuator command, r a reference, and z a partial environmental representation. A reduced plant can be written as

$$\dot{x} = f(x, u, z, \theta) + d(t), \quad y = h(x, z) + \epsilon, \quad (44)$$

where θ denotes identified or uncertain parameters and d aggregates unresolved forcing within the declared model. An adaptive design may estimate θ or a lumped disturbance from recent input–output evidence. A robust design instead specifies an uncertainty set and establishes performance or stability properties for that set. These approaches can be combined, but they answer different questions: adaptation asks what can be learned as conditions change; robust analysis asks what remains defensible across represented variations.

An observer may use vehicle velocity, attitude, pressure, Doppler measurements, current estimates, or manipulator force sensing to distinguish measured changes from predicted dynamics. Disturbance observers and robust tracking have been evaluated on real experimental vehicles, while learning-based disturbance estimates have also been integrated into constrained marine predictive controllers [59, 209]. Their reported success applies to the tested vehicles, disturbance families, and feedback rates. One cannot assume that estimation convergence is guaranteed in environments where observations fail to excite the relevant dynamics, sensors lose calibration, or environmental effects cannot be separated from body-parameter changes. Persistent excitation, structural identifiability, actuator authority, and measurement bias need independent assessment.

The same consideration arises with soft bodies. Their actuation may introduce hysteresis, configuration-dependent stiffness, delayed response, and interactions between bending and propulsion. A simplified rigid model can be useful in one operating regime but cease to be informative during large deformation or contact. Adaptive estimation then needs either a model class covering those effects or a mechanism to recognize that its current representation is inadequate. Adaptation is not a substitute for declaring what the robot measures or can safely do.

14.3 Predictive, constrained, and safety-filtered control

Model predictive control (MPC) is a useful framework for a robot that must reconcile short-horizon motion, energy and actuator limitations, and known geometric or operational constraints. At each update it estimates a current state, predicts candidate evolutions over a finite horizon using a specified model, and applies the initial portion of a selected control sequence before replanning. Nonlinear, stochastic, and learning-augmented formulations have been studied for marine vehicles under uncertain hydrodynamics and disturbances [84, 209]. Their value depends on the quality of the prediction model, available computation, feasibility of the optimization, and validity of uncertainty estimates.

A general finite-horizon problem can be represented as

$$\begin{aligned} \min_{u_{0:H-1}} \quad & \sum_{k=0}^{H-1} \ell(x_k, u_k; r_k) + V_f(x_H), \\ \text{subject to} \quad & x_{k+1} = \hat{f}_k(x_k, u_k), \\ & u_k \in \mathcal{U}_k, \quad x_k \in \hat{\mathcal{X}}_k, \end{aligned} \tag{45}$$

where $\hat{\mathcal{X}}_k$ refers to constraints represented and justified for the problem. These constraints must be specified separately from the physical conservation laws of the plant and from the legal and ethical scope of the activity. A cost on survey time or energy cannot by itself establish whether a mission may approach an animal. If ecological constraints are insufficiently observed or authorization is unresolved, the problem may instead require an alternative experiment, human review, or a withdrawal action.

Control barrier functions, viability methods, and reachability analysis offer complementary mechanisms for represented boundaries. For a control-affine hypothesis $\dot{x} = f(x) + g(x)u + w$, a differentiable function $b(x)$ can represent the nominal admissible set $\{x : b(x) \geq 0\}$. A robust barrier condition may take the form

$$\inf_{w \in \mathcal{W}} [\nabla b(x) \cdot (f(x) + g(x)u + w)] + \alpha(b(x)) \geq 0, \tag{46}$$

where $\mathcal{W}$ is a justified disturbance set and α is an appropriate comparison function. Forward-invariance arguments additionally require regularity, feasibility over the states actually reached, a justified disturbance set, correct state information, and the other conditions of the chosen barrier formulation. If only an estimated state $\hat{x}$ is available, an estimator-error set must also be propagated into the barrier test; satisfying the displayed inequality at $\hat{x}$ alone is not a guarantee about the unknown physical x . All terms in the inequality must have consistent units, with $\alpha(b)$ expressed in units of b per unit time. Measurement-robust barrier research explicitly studies this difference between true state and sensed state [28]. Research has applied barrier-based constraints to marine surface-vessel collision avoidance, including both simulation and some field trials [221, 26]. Such results concern geometrically and dynamically specified collision constraints. They do not automatically certify the protection of a species, ecosystem, or legal right whose status cannot be inferred from vehicle state alone.

A controller also needs a defined response when the nominal optimization has no feasible solution. Exiting a dynamically unsafe situation, stopping an invasive sampling action, and initiating recovery are distinct options that may require different reserve energy and maneuvering authority. A purported safety filter that can only reject commands, without providing a feasible fallback when its assumptions fail, is incomplete for an autonomous marine mission. The corresponding recovery states will be integrated into a criticality supervisor below.

14.4 Biologically inspired control and embodied coordination

Marine organisms offer diverse examples of coordinated movement under fluid forcing. Fish may generate propulsion through distributed muscle activation, body elasticity, and hydrodynamic reaction; rays rely on fin deformations; cephalopod arms combine compliant mechanical structures with distributed sensory interactions. Engineering research often draws on central pattern generators (CPGs), coupled oscillator networks, proprioceptive feedback, and mechanical compliance to reproduce selected coordination principles [58, 103, 85]. This is a functional source of hypotheses, not a claim that an artificial controller reconstructs an animal's complete neural control or experiential organization.

A schematic CPG can assign an oscillator phase ϕ_i to each actuated segment:

$$\dot{\phi}_i = \omega_i + \sum_j k_{ij} \sin(\phi_j - \phi_i - \psi_{ij}) + s_i(y), \quad (47)$$

where coupling coefficients and phase offsets determine a coordinated rhythm, while $s_i(y)$ describes measured feedback. This representation can reduce a high-dimensional gait problem to a smaller set of tunable parameters. It nevertheless requires an account of how phase dynamics are converted into forces, how the body and water react, and what stability or maneuverability properties hold under environmental changes. In a 2024 robotic-fish experiment, a CPG-inspired torque strategy with proprioceptive feedback was evaluated on an underactuated mechanism; this provides evidence for that specified morphology and experimental setting [103].

Biologically informed control also extends beyond swimming. A 2026 octopus-inspired continuum arm demonstrated simplified bending propagation, attachment sensing, and grasping of test objects, showing how sequential activation can reduce the need to solve a full-body configuration during one class of manipulation tasks [98]. A flexible arm designed for compliant exploration may still require a stiffer load path, bracing, and calibrated force sensing to operate a heavy marine connector. Different task regimes can call for different control morphologies, including hybrids with soft sensing surfaces and rigid structural support. Section 12 established the morphology problem; the control question is how to switch between the corresponding sensorimotor regimes without unsupported claims of seamless biological equivalence.

Embodied coordination presents another possible benefit: material compliance and environmental coupling can implement part of the movement pattern without computing a complete actuator-by-actuator trajectory. But a reduced command dimension does not imply reduced ecological disturbance. Fin beats, wake production, acoustic signatures, or approach behavior may remain consequential for marine animals. The appropriate test compares energy, task success, flow disturbance, acoustic exposure, and observed animal responses under specified conditions. A controller inspired by marine life may serve both engineering design and new scientific questions about biological control, provided the inference from artificial to biological mechanisms remains limited and independently investigated.

14.5 Learning-enabled policies and model revision

Learning-based control can address parts of the oceanic dynamics that are difficult to parameterize in advance. Gaussian-process regressors, neural dynamical models, residual learning, imitation learning, reinforcement learning, and physics-informed approaches may supply predictions or policies tailored to a particular vehicle. One 2024 study combines a Gaussian-process disturbance observer with MPC for a bionic vehicle, while a 2025 robotic-fish study reports an end-to-end reinforcement-learning controller evaluated with simulation-to-real transfer methods [209, 96]. These are substantive engineering developments, but their performance should not be generalized beyond the tested morphologies, environments, and evaluation protocols. The available literature also notes high training costs, differences between simulation and reality, and limited reporting of comparable validation conditions [36, 85].

A useful division is between *learning a model*, *learning a controller*, and *learning when a model is inadequate*. The first updates a representation of dynamics. The second changes how actions are selected. The third treats unusual residuals, unexplained observations, or changed contact dynamics as

evidence that an existing hypothesis requires reassessment. These activities may be jointly implemented but require distinct metrics. Improved training reward does not demonstrate that the model has become a scientifically adequate explanation of a biological process. Likewise, low simulated collision frequency cannot establish a field safety guarantee under unmodeled current regimes.

A constrained learning architecture can use model-based or adaptive control as a reference and apply an independently evaluated action screen before executing a learned proposal. In a further variant, the learned component generates candidate maneuvers or predicts flow disturbances, while a model-based supervisor controls actuation within a declared envelope. The resulting guarantee, if any, belongs to the assumptions and proof of the supervisor, including actual disturbance and sensing bounds. It cannot be inferred from the architecture diagram alone. Transfer across different bodies, seawater conditions, seasons, or biological communities must be tested independently; apparent out-of-distribution detection may itself fail when sensing degrades.

This motivates a distinction between fast changes in vehicle motion, adaptation of control parameters, and slower reassessment of the generative representation used in scientific inquiry. A robot that observes persistent unexplained biological behavior may suggest additional scientific measurements, but it should not immediately reinterpret such behavior as an environmental disturbance to be suppressed by the motion controller. The movement controller and the scientific inference system must exchange information through provenance-bearing, contestable records, as developed in Section 6. The prospective mission-science architecture is developed more fully in Section 21.

A controller can satisfy a mechanical tracking criterion while influencing which ecological phenomena are ever available for scientific inference. For example, maintaining proximity to a predicted animal aggregation may improve the volume of image data while increasing disturbance and concentrating observation on easily detected individuals. The closed-loop plant objective and the scientific sampling objective should therefore be specified as distinct, coupled design problems. A low position-tracking error cannot be treated as a representative ecological sample, just as a statistically informative sampling policy cannot authorize an otherwise impermissible maneuver. This distinction links local control to the policy-recording requirements in Section 21. The practical interface is a documented record of commanded motion, realized motion and uncertainty, sensor activation, mission purpose, and any ecological or legal condition that constrained the chosen action.

14.6 Tangent-space governance and locally justified control

The term *tangent-space governance* is used here for a proposed research synthesis: the controller should rely on a locally supported dynamical representation, track the conditions under which that representation remains credible, and refrain from extrapolating it into a purportedly known global oceanic landscape. The term does not name a newly discovered physical law. Its mathematical basis lies in familiar local linearization, uncertainty sets, robust feedback, state estimation, and finite-horizon control, with the additional methodological requirement that the validity and provenance of the model remain explicit. The local Taylor estimate in Proposition 13.1 supplies one elementary illustration of this requirement.

Suppose the feedback law, including its estimated parameters, induces a nominal closed-loop vector field $\hat{f}_c(t, x)$ and an actual field $f_c(t, x)$ for the operating regime. A set $\mathcal{D}$ and interval $[0, T]$ are declared as the validity domain. The controller should maintain estimates or conservative bounds for the model mismatch, the state-estimation error, and the proximity of the planned trajectory to represented operating constraints. Crucially, *small residuals at observed locations* do not certify that mismatch is small throughout an unobserved domain. Such a bound must follow from separate identification experiments, physics-informed constraints, validated uncertainty models, or other relevant evidence.

Proposition 14.1 (Conditional finite-horizon deviation bound). *Let $x(t)$ and $\hat{x}(t)$ solve $\dot{x} = f_c(t, x)$ and $\dot{\hat{x}} = \hat{f}_c(t, \hat{x})$ on $[0, T]$. Assume that both trajectories remain in a common domain $\mathcal{D}$, that $\hat{f}_c$ is L -Lipschitz in state there with $L \geq 0$, and that $\|f_c(t, z) - \hat{f}_c(t, z)\| \leq \delta$ for all relevant $(t, z) \in [0, T] \times \mathcal{D}$.*

Then, for $e_0 = \|x(0) - \hat{x}(0)\|$,

$$\|x(t) - \hat{x}(t)\| \leq \begin{cases} e_0 e^{Lt} + \frac{\delta}{L}(e^{Lt} - 1), & L > 0, \\ e_0 + \delta t, & L = 0, \end{cases} \quad (48)$$

for $0 \leq t \leq T$.

Proof. Writing $e = x - \hat{x}$, add and subtract $\hat{f}_c(t, x)$ in the expression for $\dot{e}$. The mismatch hypothesis and Lipschitz condition imply the upper differential inequality $D^+ \|e\| \leq L \|e\| + \delta$ wherever the upper Dini derivative is required. Grönwall's inequality yields Equation (48), with the second case obtained by taking $L = 0$. $\square$

The proposition is a standard continuous-dependence argument restated as an engineering requirement, rather than a new result in control mathematics. It bounds trajectories only when the assumed mismatch bound, local domain, and closed-loop description are actually justified. To connect this bound to a scalar constraint $b(x) \geq 0$ with b Lipschitz constant K on the region, a nominal trajectory satisfying $b(\hat{x}(t)) \geq K \bar{e}(t)$ remains inside the represented constraint set when $\bar{e}(t)$ is the valid bound from Equation (48). If the nominal clearance shrinks below that margin, the inference fails even when the nominal controller tracks its reference accurately. These are conditional statements about specified mechanical or geometric constraints. They do not guarantee marine animal welfare, independent legal approval, or the ecological validity of a model of the observed organisms.

The practical use of this result is to motivate *bounded local prediction*. A robot may choose a short maneuver when its residuals, localization, and environmental sensing support the predicted state tube, and shorten the horizon or withdraw when those conditions weaken. This can be tested against long-horizon prediction and unconstrained learning baselines. The controller's evidence should include the sensors, model version, confidence assumptions, time since calibration, and known operating regimes. In this setting, locally justified control is simultaneously a technical and epistemic discipline, while remaining separate from the normative decision about what an otherwise feasible mission should attempt.

14.7 Criticality governance and supervisory transitions

Criticality governance concerns the recognition and management of changes that threaten the conditions supporting continued operation. Mathematical criticality may involve loss of stability, a bifurcation, or transitions between attractors; operational criticality also includes actuator saturation, localization failure, insufficient reserve energy, unmodeled contact, and ecological or legal stopping conditions. The two classes should be recorded separately. An abrupt residual may arise from a sensor fault, a current change, or a contact event, without indicating a physical bifurcation. The supervisory mechanism therefore evaluates multiple hypotheses about model failure rather than assigning every warning to an environmental phase transition.

A proposed control architecture has four distinguishable interfaces. The *embodied controller* regulates propulsion, body deformation, and contact under a specified morphology. The *local adaptive estimator and regulator* maintains a bounded model of recent dynamics and updates prediction intervals. A *criticality supervisor* monitors residuals, estimated safety margins, uncertainty, energy, communication status, and authorized task scope. Finally, a *scientific and institutional mission layer* selects questions, records consequential decisions, and identifies whether continuation is within the approved research plan. These interfaces are conceptual specifications for investigation; the existence of four software components does not itself establish correct interaction or safety.

The supervisor may admit states such as *nominal*, *cautious*, *reassess*, *recover*, and *terminate*. Each transition needs a measured or documented guard, an available recovery action, and a defined authority. Hysteresis or dwell-time constraints may reduce rapid oscillation between modes, but they require tuning against both delayed response and false alarms. If an acoustic navigation fix is delayed, the feasible local maneuver set can shrink; if a robot detects possible distress in a nearby animal, it may withdraw from

the encounter pending authorized review; if a manipulator fails to maintain a bracing contact, it may release a tool only when doing so avoids additional danger. The appropriate recovery differs by setting. A sea-surface vessel may rely on navigation rules and collision sensors, whereas a deep-water glider might need to prioritize ascent conditions and energy reserves.

There is a fundamental limitation: no supervisor can guarantee a safe response when the actual plant and ecological state fall outside all justified models, or when the available actuators cannot achieve a feasible escape. The monitoring architecture must therefore retain the ability to record *undetermined* conditions, require mission redesign, and trigger external review. It should also make unresolved ecological evidence visible rather than reducing it to a confidence score that is silently traded against sampling rewards. This connects the engineering problem back to procedural review in Section 10.

Safety results in the cited control literature. The integral-control-barrier study of maritime autonomous surface ships reports a particular fully actuated vessel control problem, theoretical conditions, and numerical examples [221]. Its conclusions depend on its modeled ship dynamics, obstacle representations, observation and input assumptions; they are not experimental evidence for close-proximity ecological operation by a different class of robot. Likewise, an analytic forward-invariance condition can constrain trajectories only for the state and disturbances it actually represents. To transport a controller into a new oceanic setting, a study should separately demonstrate closed-loop tracking, robustness to estimated-state and current errors, actuator feasibility, recovery from degraded sensing, and bounded contact or exposure consequences. Collisions and welfare-related disturbances are different event classes, requiring different measurements and limits. Hence control-theoretic certification, when possible, is conditional on an explicitly identified operational domain rather than on the use of a particular algorithmic label.

14.8 Comparative evaluation, limitations, and research continuations

A credible evaluation programme needs comparisons across identical plants and disturbance realizations wherever technically possible. At minimum, experiments should include a tuned classical feedback controller; a robust or adaptive controller; an MPC implementation with declared constraints; a biologically inspired rhythmic controller where the morphology supports it; a learned or learning-augmented controller; and the proposed local-validity and criticality supervisor attached to suitable baselines. It would be misleading to force a CPG designed for fin locomotion onto an anchored industrial manipulator and then interpret its failure as evidence against biological inspiration. Controller comparisons must be grouped by compatible embodiments and tasks, with common sensing budgets, actuator constraints, and recovery requirements.

Experiments should vary the disturbance landscape: steady currents, vertical shear, wave forcing, changing density, sensor dropouts, navigation delay, actuator degradation, flexible-body deformation, and contact transitions. Evaluation should report tracking or task error, control effort, energetic cost, prediction coverage, out-of-domain detection, false-positive and false-negative alerts, intervention latency, recovery success, and rate of unrecognized constraint violations. Ecological evaluation requires additional measurements appropriate to the species and environment, such as noise, sediment disturbance, physical proximity, and observable behavioral responses. These metrics cannot be collapsed into an assumed universally valid reward. Proposed thresholds should be compared with independent ecological evidence and responsible institutional review.

The section's constructive proposal is consequently bounded. Oceanic robots may profit from combining classical, robust, adaptive, learned, and bio-inspired control, especially when their morphology and operating regime demand different methods. Tangent-space governance and criticality governance specify how those methods could be coordinated around local evidence and changes in model validity. The deviation proposition establishes only a conditional technical bound, and the examples remain hypothetical. The next engineering sections will examine sensing and localization because a control system cannot assess its dynamical envelope without sufficient observations. Subsequent sections on

Table 11: Control families and evidence requirements for oceanic robotics. These are design considerations rather than universal rankings.

Family	Principal use and current evidence	Critical comparison or limitation
Classical and robust feedback	Motion regulation, tracking, rejection of bounded disturbances; established marine practice	Operating envelope, tuning, uncertainty sets, actuator authority
Adaptive and observer-based	Online estimation of hydrodynamic changes and disturbances; field and laboratory studies	Excitation, observability, estimation bias, parameter drift
Predictive and constrained	Finite-horizon trajectories and represented bounds; numerical and selected experimental evidence	Model error, feasibility, computation, recovery policy
Bio-inspired and embodied	Rhythmic propulsion, compliant grasping, distributed coordination; platform-specific demonstrations	Morphology dependence, flow coupling, contact forces, ecological effects
Learning-enabled	Learned policies and residual dynamics; selected real-system demonstrations	Generalization, sim-to-real transfer, uncertainty calibration
Local-validity supervisor (proposed)	Evidence-bounded control and criticality-triggered revision	Validated error bounds, false alarms, safe fallback, institutional review

manipulation, distributed teams, and adaptive scientific inquiry will test whether the control interfaces remain meaningful when the task itself, the environmental interactions, and the scientific representations change.

15 Marine Perception and Biological Sensing

15.1 Perception as situated ecological measurement

The purpose of this section is to examine what marine robotic systems can presently observe, how their sensors transform ecological processes into evidence, and which uncertainties constrain the interpretation and circulation of that evidence. Section 14 considered action under partially observed dynamics. Its state estimator presupposes measurements, but the availability of a sensor output cannot establish that the intended ecological variable has been observed. The distinction between physical encounter, signal registration, inferential identification, and warranted ecological claim is fundamental to the account of relational reality developed in Section 2 and to the methodological responsibilities considered in Section 10. Here the focus is concrete: optical, acoustic, physicochemical, molecular, biological, and contact-based marine sensing, together with the conditions for evaluating their outputs.

The oceanic scope spans sea-surface platforms, the water column, seabed observatories, marine habitat interfaces, and where appropriate sub-seafloor contexts. Their sensing conditions are heterogeneous. A surface vessel may obtain atmospheric positioning and imagery of the sea surface while observing processes below through active sonar. An AUV may image a habitat at short range but have an uncertain global position. A benthic platform may support persistent acoustic and chemical measurements while seeing only a restricted volume. A mobile sampling system might encounter a marine animal whose behavior changes because of the vehicle’s light or propulsion. No universal sensor hierarchy follows from these differences. The appropriate observation system depends on the scientific question, target ecology, environmental conditions, disturbance limits, and the available calibration and validation procedures.

An abstract measurement event e can be represented by

$$Y_e \sim P_e(\cdot \mid X_{\mathcal{U}(e)}, U_e, M_e, C_e, Q_e), \quad (49)$$

where $X_{\mathcal{U}(e)}$ denotes the ecological configuration relevant to the event’s space–time support, U_e the sensing and vehicle operations, M_e the measurement apparatus and algorithmic representation, C_e envi-

ronmental context, and Q_e calibration and quality information. Equation (49) is a schema for specifying a probability model, rather than a claim that the complete ecological state is available. Event supports can overlap without sharing a globally ordered ecological history. Different observers may make non-equivalent measurements from partially related histories, as examined in Section 6. A data record should thus identify what was measured, how and where it was acquired, what was inferred, and what remains unresolved.

The Global Ocean Observing System identifies Essential Ocean Variables across physics, biogeochemistry, biology, and ecosystems, together with observation requirements and practices. These provide useful common scientific targets while allowing observation technologies to evolve; the existence of an agreed variable does not guarantee that any one platform can measure it directly, at sufficient accuracy, or in every marine setting [54, 66]. A camera’s image is neither a population census nor direct evidence of an animal’s experience; an acoustic detection is conditional on sound production, propagation, instrument response, and classification; a molecular detection records particular genetic material rather than unambiguous local organism presence. These distinctions are the organizing basis of the survey below.

15.2 Optical imaging and light–water interactions

Optical cameras and imaging systems provide spatially detailed evidence for morphology, habitat structure, benthic surveys, animal behavior, debris identification, and close-range manipulation. Conventional RGB cameras, stereo imaging, photogrammetry, video transects, line-scan systems, fluorescence instruments, laser profilers, and hyperspectral approaches serve different missions. Photogrammetric reconstruction can estimate habitat geometry where image overlap, camera calibration, lighting, pose recovery, and geometric conditions are adequate. Repeated images can support comparisons of change if viewing geometry, effort, and image processing are accounted for. Marine optical observations nevertheless depend strongly on the propagation of light through water and suspended material [6, 155].

Absorption and scattering vary with wavelength, water type, optical path, particulates, and illumination. Backscatter can obscure a target while reflections and artificial light produce apparent features absent from the subject. A common illustrative decomposition is

$$I_\lambda(p) = D_\lambda(p; d, c, a, s) + B_\lambda(p; d, c, a, s) + \epsilon_\lambda(p), \quad (50)$$

where D_λ denotes the received direct component and B_λ the backscattered contribution at image location p , both dependent on range d , environmental and illumination conditions c , absorption a and scattering s ; ϵ_λ collects residual error. The expression is deliberately schematic. The direct and backscatter coefficients are not generally identical, and direct transmission cannot be assumed to decay according to one wavelength-independent atmospheric dehazing coefficient. Akkaynak and Treibitz derived a revised marine image-formation model and validated its relevance in situ, demonstrating the practical consequence of mismatched optical assumptions [6, 7].

Algorithmic image enhancement introduces a second inferential layer. Color correction, contrast enhancement, denoising, learned restoration, and depth estimation may improve a defined downstream task, but visual plausibility does not establish faithful recovery of lost radiance or of taxonomic features. A learned image generator may create details that are unsupported by photons registered at the instrument. Scientific archives should therefore retain unaltered observations, calibration parameters, transformation versions, uncertainty estimates, and the original acquisition context. The appropriate evaluation depends on whether the task is habitat mapping, taxon recognition, geometric reconstruction, behavior measurement, or human interpretation; visually pleasing imagery is a different outcome from accurate ecological inference [155, 222].

Lighting itself is an intervention. Strong artificial illumination, close approach, vehicle shadows, laser scanning, and repeated surveys may change animal behavior or the observation environment. Specific published results make the potential sampling distortion tangible. Geoffroy and colleagues used optical and acoustic instruments to investigate pelagic responses to instrument light; the tested organisms avoided not only white and blue but also red illumination, with responses varying by irradiance,

community and context [51]. Their site- and instrument-specific findings do not supply a universal distance of avoidance, but they reject a general assurance that red survey lights are biologically invisible. In separate Monterey Bay transects, Benoit-Bird and colleagues compared robotic sampling platforms and documented attraction, avoidance and context-dependent responses to vehicle operation and illumination [17]. The contrast is instructive: the observation bias may change sign across settings and vehicle modes. A visual survey's detection function should therefore be conditioned on the vehicle's sensory footprint, with simultaneous or temporally coordinated independent measurements where feasible. A non-contact sensor can still produce ecological disturbance. Power demand and sampling duration become important for both mission engineering and welfare review. Conversely, passive daylight imaging may be inadequate at depth or in turbid water. These are conditional tradeoffs, not a reason to assume that one optical technique is universally preferable. Near-field observations should be accompanied by effort, visibility, image quality, detection limitations, and where appropriate independent evidence about behavioral disturbance.

15.3 Passive acoustics, active sonar, and interpretation of echoes

Sound can propagate through marine environments over distances at which optical imaging becomes ineffective, although propagation depends on sound-speed structure, bathymetry, scattering, ambient noise, and source–receiver geometry. *Passive* acoustic sensing records naturally occurring and human-generated signals without transmitting a probing pulse. Hydrophones on fixed seabed stations, gliders, buoys, and mobile vehicles are used to characterize marine soundscapes, detect vocalizing species, and investigate the temporal distribution of sound-producing animals. NOAA Fisheries documents applications ranging from archival bottom-mounted systems to near-real-time glider and buoy detections [131]. Such instruments cannot directly count silent animals, and a detection may not uniquely identify an individual or its ecological role. Species-specific vocal repertoires, call-production rates, classifier errors, environmental masking, and unequal listening coverage all affect what may reasonably be inferred.

Active acoustic systems use an emitted signal and measure returns. Multibeam echo sounders support bathymetry and backscatter mapping, while imaging sonars and other acoustic sensors may assist with obstacle avoidance, animal aggregation studies, seabed inspection, and contact or construction tasks. Multibeam water-column returns can reveal bubble plumes or concentrated biological layers, but a strong return alone is not a species identification. The acoustic properties of the target and water column, beam pattern, range, spatial resolution, and processing assumptions must be distinguished [135]. A fish school, suspended sediment, and a geological feature can present different signal characteristics, yet interpretation normally requires independent contextual evidence or validated classifiers.

Acoustics provides an instructive example of the distinction between physical coverage and epistemic completeness. A detector can report an animal call with high classification confidence only relative to the recordings and labels on which its calibration and validation depend. Domain shift may arise from another ocean basin, different equipment, overlapping vessel noise, or species with incompletely documented vocal behavior. Localization derived from time differences across hydrophones additionally requires propagation and clock assumptions. Passive monitoring offers important possibilities for observing marine life at distance, but absence of detected calls is not equivalent to absence of animals. Acoustic observations should record listening effort, effective detection region, classifier versions, signal quality and ecological context.

Active sensing also creates sound. Marine animals use acoustic information for communication, foraging, predator detection, and navigation; the effects of anthropogenic noise vary with source intensity, frequency, duration, exposure history, and species. NOAA notes potential acute, chronic, and cumulative ecological consequences, without implying identical effects from all sonar operations [130]. An ethically accountable system should distinguish passive recording from active emission, review possible exposure, and select measurement designs proportionate to their scientific purpose and applicable regulatory conditions. A robot cannot claim harmlessness merely because its sonar succeeds in detecting a target.

15.4 Physicochemical measurements and biological material sampling

Marine ecological states are partly characterized through temperature, pressure, salinity, oxygen, acidity, nutrients, dissolved and particulate organic matter, light, and chemical tracers. Instruments include conductivity–temperature–depth packages, oxygen optodes, pH and carbon-system sensors, ultraviolet nitrate spectrometers, fluorometers, optical backscatter sensors, microstructure probes and autonomous water samplers. These instruments have different sampling footprints, response times, drift mechanisms and power requirements. A measured chemical concentration is a local material signal conditional on instrument operation and environmental transport; it cannot on its own determine the ecological processes that produced the concentration [1, 53].

Biogeochemical Argo provides a concrete quality-control lesson. Unadjusted oxygen, nitrate, and pH measurements may contain storage or deployment drift; operators apply sensor-specific corrections and distinguish real-time from adjusted or delayed-mode records. GO-BGC explicitly cautions against scientific use of certain unadjusted nitrate and pH fields and explains the use of adjusted oxygen values. An instrument’s mere presence in the ocean is therefore insufficient for a defensible longitudinal trend [53, 15, 2]. Comparison among platforms requires more than consistent units: calibration histories, response lag, pressure effects, effective sampling depth, processing versions, and quality flags may alter whether observations are comparable.

Biological evidence can also be collected through imaging, preserved specimens, plankton nets, microfluidic sampling, environmental DNA (eDNA), RNA and protein assays, and other chemical or molecular methods. Molecular observations can disclose organisms or biological products that are difficult to detect visually. The Monterey Bay Aquarium Research Institute’s Environmental Sample Processor is an established example of autonomous collection and processing of water samples, including filtration, preservation, and selected molecular assays [115, 114]. MBARI also describes efforts to couple eDNA sampling to its long-range AUV and to develop near-real-time sequence processing. These should be distinguished from a claim that every oceanic robot already performs fully autonomous onboard sequencing and reliable ecological interpretation.

Molecular detections depend on sampling volume and location, concentration, degradation, transport, contamination control, amplification, reference libraries, and bioinformatic classification. DNA transported from another region can complicate attribution of a detected taxon to the immediate sampling site; nondetection may reflect low concentration or methodological limits. Replicate samples, field blanks, extraction controls, positive controls, and documented reference databases improve interpretability, but cannot make uncertainty vanish [101]. Sampling may also physically disturb organisms, remove biological material, or carry ecological or biosafety implications. The appropriate procedures depend on the location, sample, protected interests, and applicable permissions.

15.5 Biological inference, ecological meaning, and sensory asymmetry

The transition from a sensed signal to ecological explanation introduces additional inferential steps. Taxonomic identification, behavioral inference, abundance estimation, ecological network reconstruction, and claims about sentience have different evidential requirements. A classifier that identifies a vocalization may support one kind of claim about acoustic activity. It does not independently establish animal population size or individual welfare. An image of a coral colony records visible structure under a particular illumination and geometry, but cannot alone determine microbial associations or long-term resilience. A positive eDNA result may constrain a hypothesis about organismal occurrence without revealing the organism’s current physiological condition. Different modalities accordingly contribute partial, often non-equivalent evidence.

Observation also depends on organismal behavior. Marine animals may flee, approach, hide, change their vocal output, or fail to exhibit detectable behavior in the presence of platforms. Instrument noise and artificial illumination may alter the measured activity. In turn, animal movement and ecosystem processes can change the physical conditions of observation. Studies of observer effects are thus relevant to both epistemology and welfare. The state of the robot, the state of the animal, and the history of their

encounter cannot always be factored into independent variables; the observer is physically situated in the studied ecology, as developed in Section 4.

An important opportunity concerns *complementary observation*: combine passive acoustics, optical surveys, animal movement measurements where ethically and legally appropriate, chemical sensing, environmental DNA, and independent human ecological observations. However, agreement among modalities is only stronger evidence when their error structures and selection effects have been analyzed. For example, acoustic detections and visual encounters can be correlated because both are collected in accessible habitats or under favorable sea states. Molecular and visual observations may also share sampling-site bias. A model should retain differences in the scale, sample support, temporal window, and taxonomic resolution of each modality. Joint interpretation can proceed through explicit measurement models, uncertainty ranges, and disputable translation assumptions, rather than forcing all records into a single “true species count” variable.

This issue extends to generalized translation. A robotic acoustic representation, a molecular classification pipeline, an ecologist’s field account, and a local community’s ecological interpretation can all be relevant to a scientific question while expressing different relations and warrants. Their practical coordination may require mediation dossiers recording what was preserved, transformed, or left undetermined, as discussed in Section 6. Neither a shared database schema nor an AI-generated summary by itself establishes that the originating observation’s meaning, uncertainty, use rights, or ecological context survived the translation.

15.6 Calibration, multimodal fusion, and accountable provenance

A scientifically usable sensing system should preserve the relationship between the original signal, instrument state, environmental context, processing steps, and the claim eventually made from it. Calibration before and after deployment can test sensor response under relevant pressure, temperature, salinity, optical and chemical conditions. In situ intercomparison may reveal offsets that laboratory calibration misses. Quality control should identify missing data, drift, fouling, delayed response, saturation, clock misalignment, georeferencing error, and domain shift. The Ocean Observatories Initiative’s biogeochemical sensor guide and Argo’s documentation provide concrete precedents for method-specific uncertainty and quality-flag practices [145, 15].

For heterogeneous modalities $m \in \mathcal{M}$, let D_m denote the preserved raw or minimally processed record, π_m its calibration and transformation history, and H_m its stated measurement hypothesis. A task-relative fused inference can be schematically expressed as

$$P(X_{\text{task}} \mid \{D_m, \pi_m, H_m\}_{m \in \mathcal{M}}, C), \quad (51)$$

where C records declared contextual information. This expression alone specifies neither a correct likelihood nor independence among modalities. A fusion model must describe common sources of error, unequal detection effort, correlated disturbance, prior assumptions, and what remains outside its representation. In particular, increased model confidence after adding a sensor is not necessarily increased scientific validity if the new sensor supplies redundant or biased evidence.

A minimum provenance record should distinguish raw acquisition, time and spatial support, vehicle trajectory and actuator activity, calibration and QC versions, instrument uncertainty, software/model versions, transformations, contributors, and permissions governing subsequent use. A corrected record should remain linked to its predecessor and to claims that depend on it. This does not require a single universal observation ontology: different disciplines may maintain their own reference quantities and histories while publishing task-specific interoperable interfaces. The purpose of provenance is not to guarantee truth; it is to make inquiry reproducible, reviewable, and correctable.

15.7 Repeated surveys, nondetection, and a bounded detection proposition

A pervasive error in ecological survey design is to interpret an undetected organism as an absent organism. Let $Z = 1$ indicate that a defined target is present within a stated ecological region and observation

Table 12: Marine sensing modalities and the distinct limits on ecological inference. The examples concern representative uses, not universal performance guarantees.

Modality	Direct evidence and selected uses	Important inference or impact limits
Optical imaging	Reflected/emitted light; visible morphology, habitat geometry, behavior	Wavelength-dependent attenuation, lighting disturbance, limited range
Passive acoustics	Recorded sound; vocal activity and ambient soundscape	Silent animals, masking, uncertain attribution, sampling effort
Active sonar	Acoustic returns; bathymetry, structural echoes, aggregation signatures	Ambiguous target identity, sound exposure, propagation dependence
Physicochemical	Instrument response; oxygen, nutrients, pH, temperature, particulate signals	Drift, fouling, response lag, proxy-to-mechanism ambiguity
Molecular sampling	Collected genetic or biochemical material; selected biological detection	Transport, contamination, incomplete libraries, physical sampling cost
Contact/tactile	Force and contact response; grasping, seabed interaction, substrate inspection	Altered habitat, force uncertainty, invasiveness and permissions

interval. Let D_i indicate detection on survey occasion i using a specified protocol. Detectability is affected by habitat, organismal behavior, sensor settings, effort, and prior observations. Repetition can reduce the probability of overlooking a present target, but only under assumptions that establish nontrivial detection probability across the relevant repeated occasions. Independent repeated detections are a familiar special case; real marine surveys may contain substantial correlation and common-mode failure.

Proposition 15.1 (Conditional bound on repeated nondetection). *Suppose $Z = 1$ is fixed for the specified survey interval and that, for each $i = 1, \dots, n$ with a defined conditioning event, the probability of detecting the target on occasion i , conditional on $Z = 1$ and all previous nondetections, is at least p_* , where $0 < p_* \leq 1$. Then*

$$\Pr(D_1 = 0, \dots, D_n = 0 \mid Z = 1) \leq (1 - p_*)^n. \quad (52)$$

Proof. By the chain rule for conditional probabilities, the probability of n consecutive nondetections given $Z = 1$ is the product of the conditional probabilities of nondetection at each occasion, given presence and the preceding nondetections. The assumed lower bound p_* on each corresponding conditional detection probability bounds each factor above by $1 - p_*$. Multiplication gives Equation (52). If an earlier conditioning event has zero probability, the probability of the full nondetection sequence is already zero and the inequality holds trivially. $\square$

The proposition is a conditional statement about a defined design, not a field-estimated detection guarantee. It requires a justified lower bound on detection probability even after previous failures. A common obstruction, such as persistent turbidity, a silent animal, an unsuitable reference library, or a malfunctioning sensor, may make such a bound unavailable. Repeating the same biased measurement can then add little usable information. The proposition also says nothing about false positives, changes in organism presence, the probability of presence before sampling, or the ecological consequences of additional surveys. Those require separate models and evidence. It would be scientifically misleading to assign an arbitrary p_* and then report the resulting number as a validated absence probability.

The result offers an operational design question. When repeated surveys remain negative, should the system increase effort, choose another observation modality, adjust location and timing, request human review, or explicitly retain the absence claim as undetermined? Such decisions depend on empirical detection functions and on ecological and authorization constraints. Where further sampling risks an animal's welfare or habitat integrity, greater statistical power alone does not establish that another intervention is admissible. The proposal thus connects measurement design with the decision and refusal procedures of Section 10.

Optical enhancement and the evidence retained by a sensor. The Sea-Thru method is especially instructive for scientific interpretation: its reported correction uses a revised physical formation model and RGBD images with range information, rather than estimating an unconstrained true scene from arbitrary color images [7]. Its image comparisons support the performance of that method under the described data and water conditions. They do not establish recovery of unobserved species traits in opaque water, an unbiased abundance estimate, or the validity of every subsequently trained classifier. A recent systems review additionally emphasizes illumination, optical attenuation, scattering, and mission-specific architectures [155]. Comparative ecological studies should retain raw sensor records and depth or range estimates, validate the corrected representations on independent physical targets, report where correction is underdetermined, and then assess taxonomic or behavioral conclusions against ground-truth observations. Visual quality, radiometric fidelity, and inferential validity should be evaluated as distinct quantities.

15.8 Scientific opportunities, observer effects, and engineering continuations

The present survey identifies opportunities for biologically and environmentally sensitive observation without assuming that future sensor development will remove epistemic uncertainty. Physics-informed optical imaging can improve correction and reconstruction when its assumptions are matched to the marine setting. Passive acoustic networks may broaden observation of vocalizing animals with less acoustic emission than active sonar, while still imposing deployment and interpretation costs. Microfluidic sampling, molecular detection, and autonomous laboratory components may extend temporal coverage of intermittent biological events. Contact sensing and compliant manipulation may enable delicate handling under appropriate validated conditions. Sensor miniaturization, edge computation, and cooperative platforms may permit new combinations of these capabilities [155, 131, 114].

The most important advance may be the joint design of sensing, motion, scientific inference, and ethical constraints. A robot that detects a previously unmodeled acoustic signal might identify several competing explanations: an unfamiliar animal vocalization, a flow-induced noise source, a mechanical artifact, or an error in its recording chain. An epistemically responsible response is to preserve the raw record, revise the candidate hypotheses, consider complementary low-disturbance measurements, and document uncertainty. It is not necessarily to pursue the signal immediately or classify an unknown organism automatically. A coordinated fleet may be capable of testing more hypotheses than a single robot, but the gains must be assessed alongside duplicated bias, increased disturbance, and unequal access to the resulting knowledge.

Concrete evaluations should therefore report at least four separable outcomes: *measurement validity*, including calibration, ground-truth or reference comparisons and uncertainty; *scientific utility*, including whether the evidence changes defensible ecological claims; *operational feasibility*, including coverage, energy, storage, latency and failure recovery; and *ecological accountability*, including sound and light exposure, sampling effects, animal responses and appropriate mission restrictions. Benchmarks should compare single and multimodal systems under matched coverage and resource budgets, incorporate adversarial turbidity, acoustic masking, sensor drift, transported molecular material, unknown species, and correlated errors, and record unresolved cases without treating abstention as classification failure.

This section concludes with a conditional engineering perspective. Marine perception is a distributed process of interaction, measurement, interpretation, and correction; it cannot be reduced to a single sensor feed or an exhaustive digital image of marine reality. A platform may improve its capability to sense while remaining unable to locate itself precisely or transmit its findings, problems addressed in Section 16. The capacity to handle tools and construct installations introduces additional contact and sensing requirements, addressed in Section 17. At each stage the distinction persists between an observed signal, a warranted ecological interpretation, and an authorized intervention. These are the conditions under which new sensory capabilities can contribute to generative ecological inquiry while remaining open to scientific and ethical challenge.

16 Oceanic Navigation, Localization and Communication

16.1 Positioning as a situated scientific and operational problem

This section examines how oceanic robots establish, maintain, communicate, and subsequently assess the spatial and temporal relationships needed for scientific inquiry and responsible operations. Perception, considered in Section 15, generates records whose interpretation depends partly on where, when, and under which conditions they were obtained. Control, considered in Section 14, depends on state estimates and communication of objectives or exceptional conditions. Yet a navigation solution can be operationally useful without uniquely locating a measured organism, and a position record can appear precise while concealing an uncertain trajectory. Navigation is thus both a control capability and part of the provenance of ecological knowledge.

The oceanic operational domain includes surface vessels and floating platforms, vehicles in the water column, bottom-following robots, seabed observatories, and installations connected to shore. These systems do not share one navigation environment. A surface vehicle may receive a global navigation satellite system (GNSS) signal under suitable antenna and atmospheric conditions. A submerged vehicle usually cannot receive conventional GNSS signals directly because radio-frequency propagation through seawater is strongly attenuated. A seabed-fixed installation may have a surveyed reference position but need to establish the relative poses of its moving tools and other nodes. A float may possess accurate surface fixes at discrete times while much of its intervening subsurface path remains inferred. Real operation therefore joins partially observed motion, environmental forcing, references of differing reliability, and communication constraints [218, 219, 170].

The navigation state should be distinguished from the ecological state and from a publication-ready location claim. Let $x_t = (p_t, R_t, v_t, b_t)$ denote a chosen position, orientation, velocity, and sensor-bias state, and let z_t describe those environmental and infrastructure variables that the estimator includes. An estimate is more accurately represented by a conditional distribution or qualified feasible region than by an isolated coordinate:

$$\mathcal{B}_t = P(x_t, z_t \mid \mathcal{H}_t^{\text{nav}}, \mathcal{M}_t), \quad (53)$$

where $\mathcal{H}_t^{\text{nav}}$ is the available navigation history and $\mathcal{M}_t$ identifies models, calibrations, reference frames, and synchronization assumptions. This schematic expression presupposes a probabilistic formulation; bounded-set or interval estimators are also possible. Crucially, the confidence reported by an estimator is conditional on its measurement and disturbance models. Confidence cannot substitute for validation of those models in the actual oceanic setting.

An observation may be assigned a platform pose $\hat{p}_t$ with a covariance, yet the scientifically relevant target region can remain uncertain because sensor range, field of view, ocean current advection, propagation delays, or molecular transport have different supports. The resulting claim might concern a detected signal near an estimated track, a seabed structure relative to a map, or an organism whose location is probabilistically associated with acoustic detections. Such distinctions matter for comparisons across voyages, long-duration ecological monitoring, and decisions about approaching habitats or animals. They also explain why navigation data should be treated as revisable evidence rather than immutable metadata.

16.2 Navigation across surface, water-column, and seabed domains

Surface systems offer comparatively direct satellite position and time references. Autonomous surface vehicles can combine GNSS with inertial sensors, radar, optical perception, and marine traffic information where appropriate. The advantages are substantial for relay, mapping, and coordination, but a reliable geographic fix does not eliminate errors in heading, sensor alignment, maneuvering among other vessels, or the geographic interpretation of a remotely sensed feature. Surface conditions, wave motion, loss of communications, and the obligations associated with navigation among human-operated vessels remain independent operational issues. Persistent surface instruments can also drift, so scientific records must distinguish moving reference points from fixed survey locations [89].

For propelled vehicles operating beneath the surface, horizontal translation, attitude, and depth are usually inferred through sensor fusion. A pressure sensor constrains depth relative to a declared reference and hydrostatic assumptions, but cannot by itself determine latitude and longitude. Inertial navigation integrates accelerations and angular rates at high frequency, providing useful continuity between external measurements while accumulating error through biases and integration. A Doppler velocity log (DVL) can observe velocity relative to the seabed when bottom returns are available, or velocity relative to scattering water when configured and interpretable for water tracking. These are physically different reference relationships. Incorrectly treating water-relative velocity as ground-relative velocity can convert ocean current into a persistent navigation bias. Bottom lock can also be lost because of altitude, seabed reflectivity, beam geometry, suspended material, or instrument failure [218, 120, 170, 166].

A common kinematic reduction is

$$\dot{p}_t = R_t v_t^{\text{rel}} + u_c(p_t, t) + w_t, \quad (54)$$

where v_t^{rel} denotes vehicle velocity relative to water in body coordinates, u_c the current velocity in the chosen geographic frame, and w_t unresolved motion and approximation error. This decomposition is not a complete mechanical model: it excludes, for example, detailed fluid–structure loads and the estimator dynamics of Section 13. Nevertheless, it makes a fundamental observability problem visible. A record of relative motion can be insufficient for reconstructing an Earth-fixed trajectory when the currents encountered along the path remain unknown. An ocean glider that resurfaces with GNSS fixes can constrain its net displacement across a dive, but multiple depth-varying paths may satisfy the same endpoints. Current-informed reconstruction improves a model only insofar as current fields and motion assumptions are supported [228, 219].

Seabed-fixed or bottom-following systems introduce further alternatives. Acoustic transponder fields can establish a navigation reference within surveyed areas, while terrain-relative sonar, bathymetric matching, visual or acoustic simultaneous localization and mapping (SLAM), and local fiducials can support relative localization where environmental features are sufficiently stable and distinctive. A repeated sandy region, poorly mapped terrain, variable sediment, or biological habitat that changes between visits may defeat feature correspondence. A model that assumes an unchanging seafloor is scientifically consequential: it can mistake real ecological or geological change for a localization error, or conversely interpret an erroneous pose as environmental change. A responsible localization system should retain competing explanations when observations do not resolve them. For a fixed marine installation, high-precision docking and connector placement additionally require orientation, contact geometry, structural reference frames, and local tool calibration; global position alone does not supply assembly-grade alignment.

The same platform may traverse multiple domains. Surfacing to update GNSS can improve the global trajectory estimate, but may incur energy, schedule, collision, weather, ice, or ecological constraints. Long-range gliders often accept intermittent fixes and sparse communications to conserve energy, while intervention vehicles may invest in acoustic infrastructure and more expensive navigation instrumentation. WHOI’s *Sentry* provides an operational example: its published system description combines inertial navigation, DVL, and ultra-short or long-baseline acoustic aids for deep-sea missions, with acoustic contact providing status and retasking capability [218, 120]. These established combinations should be understood as mission-dependent architectures, not as a demonstration of universally accurate localization.

16.3 Inertial, acoustic, and environmental position references

Navigation aids differ by the physical quantity they measure and by the infrastructure required. An inertial measurement unit supplies angular rate and specific force; a DVL supplies velocity under qualified acoustic-return conditions; pressure gives a depth-related constraint; acoustic ranging uses sound travel times or bearing estimates; and mapped natural features can supply relative reference information. Combining these measurements with a nonlinear estimator, smoothing algorithm, factor graph, or

uncertainty-bounded filter is established practice. The estimator should nevertheless represent latency, outliers, cross-sensor correlation, sound-speed uncertainty, and equipment bias rather than treating all observations as independent truth.

Long-baseline (LBL) navigation uses spatially separated acoustic transponders at known or surveyed positions, allowing geometric ranging with infrastructure overhead. Ultra-short-baseline (USBL) systems estimate relative direction and range from an acoustic array, often aboard a support vessel, thereby linking the vehicle estimate to the vessel's reference state. Short-baseline arrangements occupy another geometry and installation regime. Performance depends on synchronization, propagation, anchor geometry, multipath, and calibration; there is no single accuracy figure applicable to every oceanic site. The ocean's temperature, salinity, and pressure structure changes the speed and path of sound, so a travel-time measurement should not be converted into a precise straight-line range without accounting for the applicable propagation model. Sound-speed profiles, refraction, boundary reflections, and variable currents can create both systematic errors and intermittent links [170, 81].

A minimal one-way travel-time model is

$$\tau = \int_{\Gamma} \frac{ds}{c(s)} + \Delta_{\text{clock}} + \epsilon_{\tau}, \quad (55)$$

where Γ is the unknown or modeled acoustic path, $c(s)$ its sound-speed field, Δ_{clock} relevant timing offset, and ϵ_{τ} measurement error. This expression makes clear why unknown sound-speed profiles and clock offsets are not automatically separable from geometric range. Two-way measurements remove some common timing requirements under suitable transponder protocols, but introduce their own processing-delay and motion assumptions. Bearing data can reduce geometric ambiguity, yet angular error and calibration must be assessed in the same reference frame. When a robot uses a moving ship as its acoustic reference, uncertainties in the ship's GNSS, attitude, acoustic lever arms, and instrument synchronization propagate into the inferred marine pose.

Terrain-aided navigation and SLAM reduce reliance on installed acoustic anchors but cannot manufacture distinctive landmarks where none are measurable. Repeated feature matches are only evidence of correspondence under the selected observation model. In marine habitats, features can be altered by currents, growth, sediment redistribution, and previous scientific intervention. A map suitable for obstacle avoidance may be unsuitable for detecting gradual ecological change, or vice versa. Survey plans should specify whether navigation and science reuse the same feature model; if so, errors and circular inference become a concern. A downstream ecological claim should retain enough pose and feature uncertainty to distinguish movement of the robot from change in the environmental object.

16.4 Marine communication media and network architecture

Oceanic communications are best understood as a heterogeneous network rather than a single wireless channel. At the sea surface, compatible radio and satellite systems can connect vehicles or buoys to shore services. Submerged acoustic systems can communicate over substantially greater ranges than typical optical links, but often at limited throughput and with propagation delays, multipath, Doppler effects, interference, and energy costs. Optical communication can transfer high data rates over short distances when water clarity, alignment, receiver field of view, and illumination allow. Tethered fiber links can support high-rate video and real-time ROV operation, at the expense of cable handling, dependence on surface or fixed infrastructure, and entanglement or contact risks. Conventional radio-frequency propagation is typically limited in seawater, so a terrestrial cellular or GNSS design cannot simply be submerged without adapting its physical layer [136, 179, 3].

These differences create a division of labor. Acoustic messaging can transmit bounded navigation updates, emergency signals, mission summaries, and scientific alerts while high-volume image or molecular data remain locally buffered until docking, surfacing, or a suitable short-range link becomes available. A surface relay may use GNSS and satellite connectivity while exchanging acoustic messages with mobile robots in the water column. A seabed observatory may use cables for continuous power and communication while operating acoustic or optical links to neighboring instruments. The familiar Internet

Table 13: Representative oceanic navigation and communication mechanisms, their operating roles, and major limitations. Entries describe qualitative tendencies, not guaranteed range or accuracy.

Method	Primary role or advantage	Limitation or evidence condition
GNSS / satellite	Surface location, time reference, remote upload	Requires suitable surface antenna and link; intermittent by mission design
INS / dead reckoning	High-rate local motion continuity	Drift grows between valid external constraints; bias and model dependence
DVL / pressure	Velocity relative to selected acoustic reference; depth	Bottom/water reference ambiguity, beam dropout; no horizontal fix from pressure
LBL / USBL acoustics	Geometric positioning with surveyed anchors or surface reference	Sound speed, multipath, anchor geometry, timing and infrastructure cost
Terrain / visual SLAM	Relative localization, mapping, local revisit	Feature scarcity, appearance change, correspondence and map dependence
Acoustic modem	Submerged status, range and constrained messaging	Delay, throughput, interference, sound exposure and power
Optical / tethered link	Short-range fast exchange; high-rate cable data	Alignment or tether logistics, turbidity and physical interaction

assumption of continuous, inexpensive, bidirectional connectivity is not appropriate for every marine mission. A practical network therefore needs store-and-forward queues, event priorities, time stamps, delivery acknowledgments where feasible, graceful loss of contact, and procedures for reconciling duplicates or delayed messages.

16.5 Interoperable protocols and accountable scientific data exchange

The JANUS signaling protocol illustrates one interoperability approach. NATO STANAG 4748, promulgated in 2017, establishes a common digital underwater acoustic signaling specification for discovery and initial communication. Its existence is evidence that acoustic interoperability can be standardized across selected systems; it does not supply high-bandwidth continuous connectivity or a complete scientific-data ontology [126, 179]. A robot that can exchange an initial standardized message with another platform may remain unable to translate its ecological observations or infer the other's mission permissions. The generalized translation analysis of Section 6 therefore remains relevant even when physical interoperability succeeds.

A communication architecture should record how network behavior can alter scientific claims. A delayed report of a drifting chemical plume may describe conditions no longer present at the robot's estimated position; an acoustic packet may contain a robust time stamp but an obsolete state estimate; and a compressed or filtered image may omit an anomaly noticed during later analysis of raw data. The data-management system should distinguish acquisition time from reception and processing time, sample coordinates from relay coordinates, and a message's successful delivery from verification of the underlying measurement. Correction records should propagate through locally stored data, remote repositories, and scientific products insofar as dependency information permits. These requirements connect marine communications directly with the common knowledge circulation infrastructure of Section 9.

16.6 Cooperative localization, synchronization, and distributed records

A heterogeneous scientific fleet might include a GNSS-equipped surface relay, drifting sensors, gliders, swimming robots, seabed acoustic anchors, and a specialized intervention unit. Their relative observations can jointly constrain trajectories when geometry, clock relationships, communication availability, and trust in reference states are adequate. Cooperative localization is consequently a promising research direction, particularly for lower-cost systems that cannot carry a high-grade inertial navigation suite. Field-oriented and review literature examines acoustic ranging, exchanged position constraints, and co-

operative methods for groups of ocean robots [81, 24]. However, correlated errors can spread through a network: a misplaced anchor, an incorrect sound-speed profile, or a shared clock error can bias many apparently independent vehicles. More measurements do not guarantee more independent information.

Scientific synchronization is distinct from keeping devices' wall clocks identical. A measurement record needs an interpretable acquisition time and declared uncertainty, sometimes including the duration of an exposure, the latency of a sensor's response, and the window of a biological event. Events performed by different robots may have only partially comparable time references. A seabed sensor recording an acoustic call, a surface vehicle recording a wave condition, and a sampler preserving a water specimen may refer to different physical intervals even when their data files display identical second-resolution time stamps. A provenance system should preserve these supports and the transformations through which a joint timeline is constructed. Precise clock synchronization cannot by itself establish that two observations arose from one ecological event.

Operational decisions also need freshness guarantees. The time t_a when a state was acquired, the time t_r when a message arrived, and the time t_d when it informed a decision should be distinguished. An accepted packet from an otherwise trustworthy robot might be stale at t_d because the vehicle entered a different current regime or because an animal moved. A task-relative admissibility protocol can therefore require evidence of an upper bound on message age and state uncertainty, or else direct the receiving robot to maintain separation, seek more information, or stop the affected operation. Such measures are particularly important when a remotely operated arm, moving docking unit, and nearby marine life share a restricted space. Cryptographic integrity and access controls help establish the origin and integrity of messages but do not establish that their scientific contents remain current or correct.

The communication and localization infrastructure has ecological consequences. Acoustic ranging and modems introduce sound; optical alignment beacons introduce light; installed anchors and cabled equipment occupy or disturb habitats; and surface relays or recovery operations use material and energy. These impacts vary by site, timing, exposure, and species and must be examined with qualified ecological evidence. Section 15 distinguished passive listening from active emissions, and Section 10 established the need for reviewable operation conditions. A network design that minimizes data loss may still be incompatible with an animal-welfare constraint or with an authorized exposure limit. Communication optimization should therefore be conducted within the applicable ecological and legal boundaries, not by treating all network benefits and harms as a scalar exchange rate.

16.7 Geometric ambiguity and a bounded localization proposition

Some navigation limitations are consequences of observation geometry rather than insufficient computing power. Ranging data can leave distinct positions observationally indistinguishable. Such nonuniqueness should be represented instead of allowing an estimator to report an unjustifiably unique point. A simple exact result illustrates the problem.

Proposition 16.1 (Reflection ambiguity for planar acoustic anchors). *Let P be a plane in Euclidean three-space and let all known reference anchors $a_1, \dots, a_m$ lie in P . Suppose the only measurements of an otherwise unconstrained target position p are ideal Euclidean distances $r_i = \|p - a_i\|$. If $p \notin P$, let p' be the reflection of p across P . Then $p' \neq p$ but $\|p' - a_i\| = r_i$ for every i . Thus these measurements alone do not globally identify the target's position in three-dimensional space.*

Proof. Reflection across a plane is an isometry of Euclidean space and fixes every point in the plane. Since every $a_i \in P$, the reflected distance from p' to a_i equals that from p to a_i . The two reflected positions differ because $p \notin P$. Consequently the specified range measurements coincide for two distinct candidate positions. $\square$

This proposition concerns an idealized Euclidean ranging model and a particular reference geometry. A measured depth, a verified one-sided spatial constraint, a noncoplanar anchor, a bearing observation, or an independently supported prior may remove the illustrated ambiguity. Refraction, measurement noise,

clock uncertainty, and anchor-position errors introduce further complications rather than a general guarantee of uniqueness. The point is methodological: a high-confidence numerical estimate is not evidence of identification when alternative positions remain compatible with the data and stated constraints. If two candidate positions correspond to different sensitive habitats, an operational supervisor should not silently collapse the ambiguity before evaluating approach, sampling, or construction.

The same reasoning motivates a broader observability audit for cooperative fleets. Reference-frame definitions, gauge freedoms in relative maps, anchor configuration, sensor dropout, and stale reports all determine which spatial claims can actually be warranted. Current research on energy-aware and learning-enabled oceanic localization investigates improved estimators and adaptive communication policies, yet much comparative evidence remains based on simplified simulations. Reported reductions in mean position error under a benchmark do not, by themselves, establish robust performance in an unfamiliar coastal shelf, strongly stratified region, or ecologically sensitive seabed environment [24, 170].

Protocol standardization and situated communications performance. JANUS, adopted as NATO STANAG 4748, is described in the OCEANS 2024 literature as a broadly implementable acoustic digital communications standard intended to improve interoperability among heterogeneous systems [179, 126]. This is evidence for the definition and operational relevance of a common signalling protocol, not a promise of continuous high-throughput exchange, guaranteed delivery, or mutual semantic comprehension among independent fleets. Acoustic range and data rate remain influenced by frequency, geometry, propagation, noise, and multipath, while optical links depend on water clarity and alignment. A reported link budget or successful message exchange should therefore be accompanied by actual latency, loss, synchronization and energy measurements under the mission’s environmental envelope. Interoperable packets may transmit an observation accurately while leaving uncertain its reference frame, calibration, provenance, or legitimate use; these require separate representation and review.

16.8 Research agenda and implications for responsible oceanic operations

The technological opportunity is to develop interoperable, task-aware localization and communication infrastructure in which local inference, acoustic or optical measurements, cooperative references, environmental context, and epistemic provenance can be evaluated together. Such an architecture would allow an observing robot to retain a high-resolution local map without pretending that every feature has a precisely known global coordinate; allow a surface relay to transmit compact actionable summaries without erasing the uncertainty of the original data; and allow correction of a bad anchor estimate to reach the scientific products that depended on it. It could also support cooperative scientific decisions when robots possess distinct morphologies and sensor capabilities, while preserving the autonomy to abstain when another platform’s messages cannot be interpreted or trusted for the requested task.

Research evaluations should compare at least five distinct outputs. *Localization quality* includes error against appropriately independent references, confidence calibration, trajectory closure, and sensitivity to acoustic or inertial failures. *Communication utility* includes reliable delivery, goodput, latency distributions, outage recovery, and freshness of information at decision time. *Scientific validity* includes accurate spatiotemporal support for observations and whether pose corrections alter downstream ecological inference. *Resource cost* includes energy consumption, anchor deployment, computation, logistics, and maintenance. *Ecological accountability* includes exposure to active acoustic and optical transmissions, deployment disturbance, animal-response monitoring, and compliance with documented permissions. These outcomes should be reported separately, with an explicit record of unresolved cases rather than combined into a purported universal performance score.

Adversarial evaluations should test loss of seabed DVL returns, incorrect sound-speed profiles, moved acoustic anchors, strong depth-dependent currents, uninformative visual terrain, extended communications outages, stale relay messages, common-mode synchronization faults, and disagreements between environmental maps. A plausible comparison would contrast inertial-only navigation, qualified sensor fusion, externally aided positioning, and cooperative localization under matched operational

constraints. It would also compare acoustic-only, intermittent-surface, and hybrid communication arrangements for scientific tasks requiring different evidence volumes and response times. All evaluations should clearly separate simulation, controlled-tank demonstrations, limited marine field trials, and repeated deployments across ecological settings.

The limitations of the proposed synthesis remain substantial. An estimator cannot remove an unobservability imposed by its available measurements; a learned localization policy does not certify safety under unfamiliar fluid dynamics; and a successful network message does not establish scientific truth or institutional permission. The oceanic dynamics studied in Section 13 and the local supervisory control strategies of Section 14 make these limitations operationally salient. Moreover, the physically available action may be inadmissible under the animal-welfare and procedural commitments of Sections 4 and 10. A marine robotic system should therefore regard uncertain localization, delayed communication, and inadequate references as possible reasons to revise, narrow, or suspend a mission rather than conditions that can always be compensated through more aggressive control.

This section provides the communication and spatial-inference basis for subsequent engineering analysis. The use of manipulators, contact tools, installation frames, and marine construction machinery in Section 17 makes localization more demanding by requiring reliable relative alignment and reaction-force management. Longer-lived energy and material infrastructures require persistent references and correction procedures across maintenance cycles. Throughout these developments, an oceanic navigation network is a means of supporting defensible encounters with marine environments, rather than a claim to possess a complete and globally accurate representation of the ocean's evolving relational reality.

17 Marine Manipulation, Tool Use and Construction

17.1 Intervention as a physical and epistemic encounter

This section examines how oceanic robots move from observing marine settings to making deliberately controlled physical changes within them. Manipulation includes grasping and retrieving samples, contacting or cleaning instruments, operating valves and connectors, positioning tools, and changing the configuration of an engineered structure. Construction introduces an additional requirement: multiple contact events must produce a stable assembly whose mechanical and ecological consequences persist beyond the individual maneuver. These activities connect the embodied dynamics of Section 12, the mechanical models of Section 13, the control architectures of Section 14, and the localization constraints of Section 16. They also expose the limits of treating successful robot motion as proof of successful intervention.

In the ocean, a manipulation command changes a coupled system comprising the robot, water, grasped or contacted material, supporting structures, and nearby organisms. An end effector that reaches the correct visual location may still apply an inappropriate contact force; a valve that rotates to a commanded angle may remain obstructed; a specimen that is successfully retrieved may no longer represent its undisturbed environmental condition. Scientific sampling and engineering maintenance consequently require evidence about both an action and its consequences. The physical event is not exhausted by the robot's internal report that its task planner has terminated. Tool pose, transmitted wrench, object state, collateral motion, and post-action observations should be recorded as distinct evidential quantities.

Marine intervention is also institutionally situated. Ownership of a subsea facility does not in itself establish permission to disturb surrounding seabed or marine life, and a scientific authorization need not cover unrelated productive operations. The scoped mission dossier and review procedure of Section 10 should therefore describe the objects that may be contacted, the intervention envelope, applicable jurisdiction, expected disturbances, stopping conditions, and obligations to record failures. This procedural distinction is essential for the hypothetical multinational installation developed in Sections 7–9. Its engineers might reasonably wish to replace a failed sensor or connector, while ecologists could object to a newly proposed anchoring method at the same site. Neither a demonstrated robotic capability nor a

contractual repair obligation automatically resolves the latter question.

A useful evaluation divides intervention into *approach*, *contact establishment*, *force-bearing task execution*, *release*, and *verification*. These are analytical phases rather than a necessary chronology for all robots: contacts may be maintained while a second tool approaches, a moving platform may continually re-establish alignment, and an unexpected organism may require immediate release or retreat. The phases identify which measurements support a claim of completion and where a reversible action becomes difficult to undo. Crucially, reversible robot motion and reversible ecological effects are not equivalent.

17.2 Operational systems and levels of demonstrated intervention

Remotely operated vehicles (ROVs) provide the most established examples of purposeful underwater manipulation. NOAA Ocean Exploration describes the *Deep Discoverer* ROV as equipped with two manipulator arms, sampling equipment, and storage for biological and geological material. Its human operators use live imagery and ship-supported telepresence to select and retrieve specimens; it is an operational instrument, not a demonstration that all its grasp choices are fully autonomous [137, 133]. WHOI's hybrid *Nereid Under Ice* vehicle combines onboard energy, a slender optical tether for remote interaction, and a seven-function electro-hydraulic manipulator used for sampling and in situ work. Such hybrid architectures preserve real-time human guidance under suitable communication conditions and provide a different operational trade-off from heavy work-class ROVs [217]. These systems make the maturity of teleoperated work visible, while reminding us that their supporting vessels, tethers, power systems, and experienced teams are part of the effective intervention infrastructure.

Research on intervention autonomous underwater vehicles (I-AUVs) aims to reduce continuous human control in selected tasks. Experiments with the Girona 500 platform reported autonomous docking to a specially adapted intervention panel followed by valve turning and connector insertion or removal, with tests in a tank and a harbor environment [151]. The mechanism benefited from a funnel-shaped docking structure and engineered targets intended to simplify manipulation. This is a substantive demonstration of a particular intervention workflow, rather than evidence of generalized autonomous operation on arbitrary subsea equipment. Earlier Girona research on floating-base valve turning similarly developed task-priority kinematic coordination in controlled conditions, illustrating how vehicle motion and arm motion must be managed together [25].

More recent research demonstrates cooperation between heterogeneous intervention vehicles. López-Barajas and colleagues reported a pipeline-coupler assembly experiment using Girona 500 and BlueROV2-based platforms in a controlled water tank. Their workflow combined grasping, a synchronized teleoperated approach to an assembly starting configuration, and coordinated assembly control [100]. The reported tank result should not be generalized to unrestricted autonomous subsea construction, long-duration offshore maintenance, or unstructured habitats. In particular, coordination from a prepared initial configuration is a different accomplishment from discovering, preparing, securing, joining, and subsequently certifying an unfamiliar seabed structure without assistance.

This evidence hierarchy should shape research claims. Bench tests can support actuator limits and contact models; tank demonstrations can test coordination under controlled perturbations; restricted field trials can identify effects of turbidity, currents, and natural seabed heterogeneity; repeated operational deployment can address reliability and maintenance. Each class supports different generalizations. Successful task execution at one depth does not imply pressure-tolerant operation elsewhere, and a successful grasp does not imply that the manipulated organism, sediment patch, or installation remained ecologically unchanged. The scientific question is not simply whether an intervention was completed, but which components of the coupled physical and normative task were actually tested.

17.3 Grasping, contact mechanics and tactile evidence

The contact interface is the immediate physical boundary between robot and object. Mechanical grippers, claw-like manipulators, compliant fingers, suction tools, magnetic fixtures, and specialized end effectors

Table 14: Representative evidence for marine intervention and its limits. The comparison concerns task classes, not an ordinal ranking of platforms.

System or method	Demonstrated task class	Scope qualification
NOAA <i>Deep Discoverer</i>	Human-directed biological/geological sampling with manipulator arms	Operational ROV with ship, tether and human decision support.
WHOI <i>Nereid Under Ice</i>	Hybrid survey and sampling using a manipulator and optical tether	Mission- and communications-dependent intervention.
Girona 500 I-AUV	Autonomous docking, valve actuation and connector tasks	Engineered panel; bounded tank and harbor demonstrations.
Dual I-AUV assembly study (2025)	Coordinated coupler/pipeline assembly with two vehicles	Prepared task and water-tank validation; intermediate teleoperated alignment.
Soft sample grippers	Compliant retrieval of delicate organisms	Biological sampling demonstrations; no general proof of welfare or construction capability.

differ in force transmission, spatial tolerance, object compatibility, and failure modes. A rigid parallel-jaw grasp may provide repeatable geometry for manufactured components but concentrate pressure on an irregular specimen. A compliant gripper can redistribute contact over a larger area, although excessive deformation may obscure contact location or impede precision. Suction can help retain selected delicate material but also entrain water, sediment, or neighboring organisms. A claw appropriate for stable rock retrieval may be inappropriate for living coral. Tool suitability is relational: the same mechanism changes meaning with object stiffness, geometry, surface condition, and mission purpose [48, 175].

A minimal quasi-static grasp model introduces object wrench $w_o \in \mathbb{R}^6$, contact forces $f \in \mathbb{R}^k$, and a configuration-dependent grasp map $G(q)$:

$$w_o = G(q)f + w_{\text{ext}}, \quad f \in \mathcal{C}(q, \mu, \sigma), \quad (56)$$

where w_{ext} collects modeled external loading and $\mathcal{C}$ is an admissible contact-force set depending on contact geometry, friction μ , and a chosen contact-state or material description σ . This schematic static relation omits fluid loads on individual fingers, object inertia, suction dynamics, and the possible destruction of an uncertain grasp surface. The equality can be used to reason about attainable wrenches only when the contact model remains valid. In an oceanic setting, unknown biofouling, sediment, fluid forces, or organic tissue can invalidate a friction coefficient measured in a laboratory. A contact controller should therefore distinguish the estimated object wrench from the measured or inferred interaction forces.

Tactile sensing can improve control at the point where optical or sonar interpretation becomes insufficient. Encoders indicate commanded or observed joint motion; motor current and hydraulic pressure are indirect indicators of actuator effort; wrist force/torque sensors can detect transmitted wrenches; distributed tactile sensors may identify pressure patterns, contact transitions, and slip. Underwater use introduces sealing, hydrostatic-pressure, thermal, drift, and calibration challenges, and published research treats tactile perception as an active area rather than a uniformly deployed capability [175]. No one modality directly measures an organism’s subjective experience or the long-term ecological effect of contact. Sensor fusion should not be presented as eliminating such interpretive limitations.

The distinction between contact and manipulation is particularly important when handling naturally occurring material. A robot may immobilize a specimen and alter its behavior even when the final pressure is low. A sediment core may be successfully captured while its stratigraphy is disturbed; an adhesive-covered connector may be held firmly while a nearby fouling community is detached. An accountable record should therefore preserve pre-contact observations, approach trajectories, tool configuration, contact proxies, extraction forces or limits when measurable, observed organismal response, retained sample condition, and post-contact follow-up. In some tasks, the appropriate result may be a deliberately aborted attempt or a non-contact alternative. The welfare and epistemic asymmetries considered in Section 4 remain operative at the end effector.

17.4 Coupled vehicle–manipulator mechanics and coordinated control

An arm mounted to a free-floating vehicle cannot generally be treated as a stationary industrial robot. Moving the arm redistributes momentum and alters hydrodynamic loading; applied tool forces can push the vehicle away from a target; and thruster corrections can change the geometry of contact. Conversely, a vehicle that docks, grasps a stable fixture, or rests on a suitable support obtains mechanical constraints unavailable during free flight. The appropriate model depends on which degrees of freedom are actually supported and on whether contact is maintained or intermittent. Foundational underwater vehicle–manipulator research developed coupled models incorporating nonlinear and hydrodynamic terms; subsequent work explored feedback linearization and the sensitivity of motion control to model accuracy [162, 161, 176].

For a model with base configuration η , joint variables q , and combined generalized velocity ν , a local task-velocity relation can be written

$$v_e = J_b(\eta, q)\nu_b + J_q(\eta, q)\dot{q}, \quad (57)$$

where v_e is the end-effector twist in a declared reference frame. This kinematic relation should not be confused with a force balance. For example, holding $v_e = 0$ during a connector insertion does not establish that the connector is unstressed or that the base has sufficient thrust authority against a current. A schematic coupled dynamic equation is

$$M(z)\dot{\nu} + C(z, \nu)\nu + D(z, \nu, u_c) + g(z) = B(z)\tau + J_c(z)^\top \lambda + d, \quad (58)$$

where z includes base pose and arm configuration, M is an effective inertia matrix, C and D collect inertial coupling and modeled fluid loads, g buoyancy/gravity effects, $B\tau$ actuated generalized forces, λ contact forces, and d unresolved disturbances. The sign of the contact term follows the adopted convention. The expression is a modeling template, not an identified model of a particular system. Contact transitions, hysteresis, tether forces, thruster interactions, and deformable structures may require distinct models, as discussed in Section 13.

Several control strategies can be distinguished. *Free-floating manipulation* allocates motion between vehicle and arm while estimating reaction forces. *Station-kept manipulation* uses thrusters to compensate for disturbances and arm reaction, subject to energy and actuator limits. *Docked or mechanically supported manipulation* transfers some reaction force to fixtures or the seabed, potentially improving accuracy but introducing contact hazards and dependence on fixture integrity. *Dual-arm or multi-robot manipulation* distributes grasp forces and workspace coverage, but may produce damaging internal forces if participants disagree about object pose or compliance. Task-priority kinematic methods can enforce alignment or joint-limit objectives, whereas impedance and force-aware strategies address interaction behavior under qualified contact models. None can replace a sensor and model validity argument.

Coordination also depends on observability. The global localization of a platform can be relatively coarse while the relative pose of a connector and tool is well estimated through nearby cameras and fiducials. Conversely, two cooperating vehicles can report accurate global positions yet lack the millimeter-scale relative alignment needed to mate a connector. Force-bearing tasks thus require local registration, tool calibration, pressure-compensated hardware, and a representation of stale measurements or temporary occlusion. Timing is equally consequential: a force command based on a delayed observation may be appropriate for a stationary panel but dangerous for a drifting component held by another robot. The synchronization and provenance requirements of Section 16 should be carried through to contact-level decisions.

17.5 Tools, fixtures and task-relative engineering

The design of tools and targets is often as consequential as the sophistication of the robot. A purpose-built docking receptacle, funnel guide, keyed connector, compliant alignment fixture, or remote-operable valve can turn a difficult free-space manipulation problem into a constrained sequence with measurable

end states. The Girona panel experiments demonstrate this principle: successful autonomous intervention relied on a designed interface and reliable relative positioning as well as manipulation control [151]. Engineering a structure for future robotic repair should therefore be considered early in the infrastructure design, not added only after an installation has become inaccessible or dangerous.

A modular intervention system may include grippers, sampling scoops, corers, cleaning heads, fastening or torque tools, valve keys, sensor cartridges, and electrical or fluid connectors. Each tool imposes distinct requirements for grasp stability, reaction-force absorption, orientation, cleaning, contamination control, and verification. A sediment corer requires preservation of sample stratigraphy; a valve tool requires torque limits and reliable confirmation of the resulting valve state; a connector operation requires sealing, alignment and electrical or fluid validation. In a biological study, a sample-container identity and chain of custody are part of scientific utility. In infrastructure maintenance, traceability extends to replaced components, their specifications and the observed post-repair state. Neither task is adequately summarized by an end-effector reaching a target pose.

Tool exchange creates further infrastructure dependencies. In industrial robotics, an interchangeable tool is often assumed to be available in a controlled work cell. At sea, a tool may need a pressure-tolerant coupling, positive latching, mechanical retention against currents, contamination protection, powered communication, and a recovery path if exchange fails. A subsea tool magazine or docking node can enlarge a vehicle's capabilities, but also creates an additional constructed object that may disturb the marine environment. It may be better to deploy different specialized platforms than to pursue a universally dexterous machine whose mass, seals, power, and failure modes undermine useful mission availability. This comparison follows the morphological portfolio discussion in Section 12 and will be developed at the fleet level in Section 20.

For scientific operations, the preferred tool may be one that minimizes irreversible interaction rather than one that maximizes throughput. Galloway and colleagues demonstrated compliant grippers for sampling fragile deep-reef organisms, with controlled tests and field experience showing the potential of soft end effectors [48]. Those achievements support consideration of lower-concentration contact forces; they do not prove that every contact was welfare-neutral or that a soft tool can manipulate heavy structural components. The relevant design space therefore includes tool refusal, contact-minimizing observation, and smaller or differently sequenced samples. Mission success should be evaluated against the scientific question and the permissible ecological intervention envelope.

17.6 Assembly, seabed construction and lasting environmental change

Marine construction is more demanding than a single grasp because spatial relationships and load paths must remain stable through the transition from separate components to an assembled system. Representative tasks include installing seabed observatory modules, joining pipe or cable connectors, replacing sensor frames, inspecting anchors, positioning protective housings, fastening components, and maintaining robotic docking infrastructure. Proposed autonomous assembly would require accurate mapping and site preparation, appropriate foundations or attachment points, cooperative transport, controlled insertion and joining, and verification of mechanical, electrical, and environmental outcomes. Different stages may be performed by ships, cranes, ROVs, divers where appropriate, and autonomous platforms, rather than by one robot throughout.

Mechanical tolerances must be evaluated together with disturbances. An inserted part may need a small translational error, but also a specified angular alignment, orientation of seals, allowed insertion load, absence of trapped debris, and verification of engagement. A heavy module adds seabed bearing capacity and stability requirements. Soft seabed sediments can move under distributed load; seabed fauna and benthic structures may be damaged through placement or anchoring. Turbidity generated during installation can degrade optical feedback and change the ecological conditions being observed. An assembly plan consequently requires a model of construction *process* and of the resulting long-lived material footprint. It is not sufficient to optimize a robot's trajectory while holding the environment conceptually fixed.

The 2025 two-I-AUV pipeline assembly study is informative because it tests coordinated contact and alignment with two distinct platforms, yet its prepared water-tank conditions remain far from a demonstration of generalized offshore construction [100]. In future field evaluations, differences in currents, wave-coupled support-ship motion, low-visibility sediment, nonrigid anchoring and intermittent communication should be represented as separate sources of failure. A connector might align successfully yet leak later; a fixture may become stable yet damage a habitat; a robot pair may achieve geometric assembly while generating uncontrolled internal stresses. Long-term structural and ecological outcomes therefore demand monitoring beyond the immediate contact task.

Relevant environmental impacts cannot be inferred from robot scale alone. NOAA's synthesis of submarine cable regulatory studies identifies potential effects of installation and maintenance, including physical seabed disturbance, habitat change, sediment effects, noise, electromagnetic fields and thermal effects, with severity depending on methods and ecological setting [139]. Similar impact categories can inform assessments of robotic infrastructure work, while the magnitude and legal treatment of any particular installation require site- and activity-specific analysis. Cable installation, marine scientific sampling, and robotic construction are not automatically subject to identical authorization regimes. Sections 7 and 8 supply the necessary distinction between institutional competence, environmental obligations and proposed normative commitments.

The transnational A–B–C–D marine data-center example makes these tensions concrete. Suppose a failed seabed communication coupler is held by an autonomous maintenance robot while a second robot installs a replacement. Engineers may verify electrical continuity and mechanical alignment, yet an ecological review might ask whether suspended sediment affected an adjacent filter-feeding community. The installation operator's record of maintenance expense and service recovery does not describe these consequences. A justified completion dossier would link tool actions and measurements to subsequent ecological observations, maintenance documentation, and reviewable authorization conditions. Further investigation of material durability, corrosion, seal design, and manufacturing belongs to Sections 18 and 19 rather than being treated as solved by successful assembly.

17.7 Bounded contact-force governance and a conditional proposition

A formally modest but useful supervisory problem is whether a commanded contact force remains within a specified mechanically admissible set despite a bounded error. Let $\hat{f} \in \mathbb{R}^k$ denote an estimated or commanded vector of contact forces and let the actual force be $f = \hat{f} + e$, where e lies in an uncertainty set E . Suppose an engineering review defines a contact-force region C using experimentally justified equipment, object and task constraints. Define the robustly admissible command set by the Pontryagin difference

$$C \ominus E = \{\hat{f} \in \mathbb{R}^k : \hat{f} + e \in C \text{ for every } e \in E\}. \quad (59)$$

This construction identifies commands that remain in the declared engineering force region for all modeled errors. A force set calibrated for one tool and rigid connector need not apply to a living organism, an unknown sediment support, or a different grip configuration. Moreover, mechanical admissibility is only one condition of ecological and institutional acceptability.

Proposition 17.1 (Monotonic restriction under enlarged contact uncertainty). *Let $C \subseteq \mathbb{R}^k$ be an arbitrary specified admissible contact-force set and let $E_1 \subseteq E_2 \subseteq \mathbb{R}^k$ be two uncertainty sets. Then*

$$C \ominus E_2 \subseteq C \ominus E_1. \quad (60)$$

Consequently, enlarging the assumed error set cannot enlarge the set of commands satisfying the same robust contact-force requirement.

Proof. Take any $\hat{f} \in C \ominus E_2$. By definition, $\hat{f} + e \in C$ for every $e \in E_2$. Since $E_1 \subseteq E_2$, the same inclusion holds for all $e \in E_1$. Therefore $\hat{f} \in C \ominus E_1$, establishing the claimed set inclusion. $\square$

For polyhedral engineering constraints $C = \{f : Af \leq b\}$ and a bounded error set E , Equation (59) becomes the family of conditions

$$a_i^\top \hat{f} + \sup_{e \in E} a_i^\top e \leq b_i, \quad i = 1, \dots, m, \quad (61)$$

where $a_i^\top$ is row i of A . This representation supports explicit uncertainty margins in a model-based contact supervisor. It is meaningful only if the underlying measured force, calibration, relevant contact locations and disturbance bound are credible. The proposition is set-theoretic; it establishes neither stability of a full floating-base closed loop nor animal-welfare protection. As in the local error analyses of Sections 13 and 14, the purpose is to identify a defensible local property and its validity domain.

The proposed supervision should also assess what occurs when a command becomes inadmissible. A stationary industrial manipulator may respond by freezing joints, whereas an underwater system may be pushed by currents or continue applying load after its thrusters stop. Withdrawal may itself create damaging scraping forces, and opening a gripper can release a sample into an unsuitable habitat. Therefore, *stop*, *hold*, *release*, *retreat* and *request human review* are distinct actions with physical risks. A criticality supervisor should compare feasible transitions against the present contact state, local marine dynamics, supported ecological observations and mission authorization. A decision to abstain may be an engineering accomplishment when successful completion would require assumptions that available evidence cannot justify.

Disaggregating an experimental assembly claim. The cooperative pipeline-coupler experiment reported by López-Barajas and colleagues achieved a precision task with two intervention vehicles in a controlled approximately 5 m-deep water tank [100]. The study’s described procedure includes prepared visual fiducials, magnetic assistance at the joining interface, an operator-assisted approach to the starting pose, and position-holding by one platform. These details identify an experimental contribution in coordinated grasping, alignment, and joining; they are material to reproducibility and to the interpretation of the word “autonomous.” The experiment did not expose the system to naturally changing currents, turbid sediment, uncertain external anchoring, long-duration structural loads, or habitat-sensitive field conditions. A sequence demonstrating autonomous construction offshore would additionally have to establish localization, approach, fixture preparation, joining quality, post-assembly inspection, recovery and environmental monitoring. In the meantime, the tank demonstration is better treated as a bounded enabling step than as an operationally certified construction capability.

17.8 Validation, ecological responsibility and research continuations

A defensible evaluation programme should separate mechanical, operational, scientific and ecological outcomes. *Mechanical measures* include grasp retention, contact-force and torque error, slip, tool positioning, reaction loads on the vehicle, insertion accuracy, structural integrity and failure recovery. *Operational measures* include mission duration, energy, human intervention time, communications dependence, repeatability across currents and visibility, and recovery from tool or sensor loss. *Scientific measures* include sample representativeness, contamination, spatial provenance, preservation of material condition and the validity of any claim derived from an intervention. *Ecological and institutional measures* include observed disturbances, relevant exposure estimates, animal responses, habitat alteration, cumulative interactions, permit scope and adequate post-operation review. These cannot be reduced to one universal efficiency score without obscuring value judgments about harm and acceptable evidence.

Experiments should proceed from identified tool and contact models to instrumented tanks, controlled coastal sites, representative non-sensitive installations and, only where independently authorized, ecologically consequential sites. A meaningful benchmark would compare a human teleoperator, a fixed-rule supervised controller, an impedance-aware intervention method and a locally adaptive or learning-enabled method under matched objectives and documented environmental conditions. The evidence must differentiate model-based laboratory guarantees from empirical results. Testing should deliberately introduce moving targets, inaccurate object stiffness, sensor occlusion, poor visibility, changing

currents, shifting foundations, entanglement risk, dropped communication and correlated localization errors. Adversarial trials should also include attractive but unauthorized objects, to ensure that a planner does not confuse technical reachability with permission to touch.

Biological sampling and infrastructure construction demand different evaluation designs. For a sampling robot, non-contact alternatives, organism identification, contact duration, recovery and ecological interpretation may dominate. For a repair system, sustained load, sealing, material reliability, and post-installation monitoring may dominate. The same physical tool may be useful in both domains without the evidential or normative criteria being interchangeable. Where living organisms are involved, controlled performance should not be represented as proof of absence of pain, stress, behavioral changes, or cumulative habitat effects. Welfare questions require their own methods, appropriate expertise, and ethical review; the technologies discussed here do not independently establish those results [48, 133].

The proposed GR interpretation remains deliberately limited. A robot–object–habitat relation can be represented through contact histories, material dependencies and subsequent measurements, but a single robotic representation does not exhaust the ecological relations of the encountered system. Construction creates new physical conditions for later inquiry and action; failure and repair can redirect that process in ways that cannot be inferred from an initial task score. Such observations justify treating intervention as a generative and accountable event, without claiming that a universal ontology or algorithm can encode all affected organisms, institutions and futures. The challenge is to preserve traceable evidence and opportunities for revision across partially interacting scientific and engineering practices.

This section establishes the distinction between mobility, dexterity, tool-mediated work and enduring construction. Section 18 will examine how actuators, materials and manufacturing constrain the morphology and force capabilities assumed here; Section 19 will analyze pressure, energy, thermal effects, corrosion and lifetime reliability. Section 20 will extend manipulation to heterogeneous robotic teams and infrastructure operations, including the division of labor among sensing, transportation, assembly and maintenance platforms. Throughout, the goal is an engineering research programme that can expand responsible marine scientific capacity while preserving the possibility to question, correct or refuse physically consequential interventions.

18 Actuation, Materials and Robotic Manufacturing

18.1 The material realization of ecological robotic action

Sections 12–17 examined the relation of morphology, dynamics, control, perception, navigation, and intervention. Their common material premise warrants separate examination: a command becomes a physical action only through actuators, transmissions, structural members, seals, joints, interfaces, and manufacturing processes. Oceanic deployment then modifies those components through pressure, salinity, temperature, motion, contact, and biological colonization. The objective of this section is to assess the physical means of action and the evidentiary conditions under which a design may be considered durable, maintainable, and ecologically responsible. It is an engineering literature review and a proposal for comparative experiments, rather than a report of a newly constructed machine.

Actuation and material selection should be evaluated together. A motor described by its nominal shaft torque may have insufficient available torque after transmission losses, increased seal drag, or an adverse temperature excursion. A compliant fin selected for its strain capacity may deteriorate at a bonded rigid interface. A manufactured component with attractive laboratory performance may be impractical if its failure cannot be diagnosed or repaired offshore. Moreover, low force, soft contact, and biological resemblance are properties of particular interactions rather than unconditional guarantees of negligible ecological impact. A material can remain mechanically serviceable while releasing compounds, transporting organisms, obscuring scientific samples, or changing habitat surfaces. The framework therefore separates *functional performance*, *evidence of physical integrity*, *repair capacity*, and *ecological consequence*.

The scope extends to robots operating at the sea surface, through the water column, near benthic habi-

tats, and at subsea installations. Some actuation principles can be compared across these settings, but validation conditions cannot be freely transferred. A pressure-balanced oil-filled drive, a deck-mounted electric winch, a soft arm immersed in seawater, and a mobile robot used for coastal fabrication experience distinct pressure differentials, fouling histories, chemical exposures, and maintenance opportunities. Each architecture must be understood relative to a bounded mission envelope. For a task T , material system m , manufacturing route r , and environment history h , let the experimentally estimated performance vector be

$$\mathbf{P}(T, m, r, h) = (F_{\text{available}}, \Delta x, \eta, D_{\text{observed}}, M_{\text{repair}}, E_{\text{ecol}}), \quad (62)$$

where the terms denote, respectively, task-available effort, motion range, an efficiency measure with a stated denominator, observed degradation, a maintainability measure, and a vector of ecological effects. These heterogeneous quantities cannot be combined into an invariant scalar ranking without an explicitly justified evaluation rule. A configuration favorable for broad-area surveying may be undesirable for repeated high-load assembly; a configuration suitable for delicate specimen handling may be unable to stabilize a connector against current-induced loads.

18.2 Electric, hydraulic, fluidic, and transmission-based actuation

Conventional marine robotic propulsion often relies on electric motors driving propellers or other rotating mechanisms, while intervention systems may use electric or hydraulic joints, cables, gearing, and dedicated subsea tooling. Electric drives support programmable velocity or torque regulation and may simplify modular control architectures. Their mission suitability nevertheless depends on insulation, motor and bearing design, seal integrity, corrosion mitigation, electromagnetic compatibility, and thermal limits. For a rotating output, the elementary balance $P_{\text{shaft}} = \tau\omega$ relates torque and angular speed, but electrical input, shaft output, hydrodynamic thrust, and task-delivered work are different quantities. A high nominal motor efficiency is not evidence of a high total mission efficiency.

Hydraulic actuation remains important where concentrated force, compact power density, or established manipulator interfaces are required. The idealized relation $F = pA$ between pressure difference p , piston area A , and force F is modified by friction, losses, compliance, valve dynamics, and the surrounding hydrostatic pressure. Hydraulic circuits can transmit substantial loads, while introducing potential leak pathways, fluid compatibility questions, and maintenance burdens. Appropriate descriptions distinguish a pressurized hydraulic circuit from ambient-pressure-compensated volumes: pressure compensation reduces differential-pressure loading on selected components but does not eliminate material stress, leakage, corrosion, or the need for correctly specified sealing. Similar distinctions apply to electric motors in flooded or oil-filled housings and to assemblies containing a mixture of pressure-tolerant and protected elements. A design should document which subsystem is pressure-exposed, pressure-balanced, or isolated rather than calling the whole robot pressure-proof.

Cable-driven actuation relocates heavy motors away from articulated or compliant sections, reducing local mass and simplifying some pressure-exposed mechanisms. Its transmission models must account for cable stretch, friction, backlash, routing and path-dependent curvature. Qi and colleagues' multi-section compliant arm provides a useful controlled example: they manufactured a silicone-body structure with printed rigid endcaps and identified the effects of cable indentation on the relation between cable length and bending [157]. The experimental success concerns the reported arm and test conditions. It cannot be inferred that its kinematic model applies unchanged to a salt-exposed arm enduring sediment abrasion or deep-ocean deployment.

Pneumatic and liquid-driven soft chambers generate deformation through internal-pressure differences and constrained wall geometry. Their softness enables broad contact areas and potentially gentle grasping, as demonstrated in marine biological sampling by Galloway and colleagues [48]. Yet ambient pressure changes the available differential pressure, hose compliance changes bandwidth, and leaks can contaminate a sample or damage a mission. Compressible pneumatic systems at substantial ocean depth entail additional complications that do not accompany their use in shallow laboratory tanks. Hydraulic-fluid choice, pump design, and energy consumption must therefore be evaluated jointly with actuator

structure, not treated as secondary implementation details.

For a single commanded axis, a general transmission relation can be written

$$\tau_{\text{load}}(t) = \eta_{\tau}(x, t) G(x, t) \tau_{\text{motor}}(t) - \tau_{\text{loss}}(x, \dot{x}, T, s, t), \quad (63)$$

with transmission factor G , an appropriately defined torque-transfer term η_{τ} , and a loss term depending on state, temperature, surrounding medium s , and operational history. This is a bookkeeping model, not a universal constitutive law. For nonlinear transmissions the useful formulation may instead be a full input–output map with measured hysteresis and uncertainty intervals. The critical distinction is whether load-bearing capacity is specified at the *motor*, *joint*, *end effector*, or *ecological contact*. Section 17 showed why these are not interchangeable.

18.3 Soft actuators, artificial muscles, and bio-inspired physical mechanisms

Biologically inspired aquatic robots exploit elastomeric membranes, fluid-filled chambers, flexible appendages, compliant tendons, and electroactive materials. A soft architecture can reduce local rigid-body mass, accommodate distributed deformation, and make interaction mechanics less sensitive to exact object geometry. But the relevant performance quantities include output force, stroke, actuation rate, efficiency, fatigue, material creep, electrical safety, permeability, and controllability under changing hydrostatic pressure. The engineering task is to explain which biological function a selected physical mechanism reproduces, rather than to treat external resemblance to a fish, octopus, or jellyfish as a demonstration of functional equivalence.

Dielectric elastomer actuators generate strain through electrostatic forces in deformable insulating membranes. Their behavior depends on electrical field, dielectric properties, prestretch, geometry, temperature, material fatigue, and insulation. A salient marine demonstration is Li and colleagues' snailfish-inspired, untethered soft robot with decentralised electronics embedded in silicone [93]. The article reports fin actuation during a Mariana Trench test at approximately 10,900 m and untethered swimming after release at 3,224 m in the South China Sea. It is essential to distinguish these outcomes: the trench experiment documented actuation while supported by a lander, whereas the South China Sea experiment documented free swimming. Neither by itself demonstrates long-duration autonomous exploration, biological sampling, precision manipulation, or lifecycle reliability. The high drive voltages involved also require an appropriate electrical and environmental failure assessment.

Hydraulically amplified self-healing electrostatic actuators (HASEL) combine fluidic redistribution with electrostatic forcing; Acome and colleagues reported muscle-like laboratory actuation and demonstrations of grasping and self-sensing [4]. This research establishes a physical principle and controlled-device behavior, not marine qualification. Immersion introduces insulation, seal, humidity, hydrostatic loading, dielectric breakdown, and contamination questions that cannot be removed by the phrase *self-healing*. Other artificial-muscle candidates include shape-memory alloys, ionic polymer–metal composites, conductive or magnetic composites, and stimuli-responsive hydrogels. Each trades driving voltage, bandwidth, force density, cycle life, manufacturing complexity, and environmental compatibility differently. Their specific suitability for saline water and a specified depth must be demonstrated rather than inferred from generic soft-robotics performance [5, 11].

Material compliance also changes control and sensing. A distributed soft structure stores elastic energy, deforms under hydrodynamic load, and exhibits delayed responses through viscosity, geometry, and fluid–structure coupling. An actuator's free displacement measured without load is insufficient to predict its behavior when grasping a coral-adjacent specimen or holding a tool against a fixed pipeline. For a viscoelastic element, even the elementary representation

$$\sigma(t) = E \varepsilon(t) + c \dot{\varepsilon}(t) \quad (64)$$

is only a local constitutive approximation for a specified strain and frequency range; anisotropic silicone composites, large deformations, and swelling may require nonlinear and history-dependent laws. Co-design should consequently encompass actuator geometry, physical feedback, material identification,

and the expected contact environment. The more ambitious prospect of combining pressure-tolerant soft bodies with manipulators and independent scientific instruments is reserved for the research programme in Section 17, not presumed established by a locomotion prototype.

18.4 Structural materials, joints, corrosion, and biological colonization

Marine robotics uses metals, polymers, elastomers, composites, ceramics, adhesives, and hybrid joints because its physical functions vary greatly. Metals offer machinability and load-bearing capacity but can suffer localized corrosion, galvanic coupling, or mechanical fatigue. Polymeric housings and soft components can reduce weight and enable complex geometries but introduce creep, water uptake, permeation, ultraviolet ageing at the surface, and adhesive-joint uncertainty. Composite structures can tailor stiffness and mass, while requiring scrutiny of laminate damage, delamination, impact behavior, and interface sealing. Ceramics can protect specialized sensing or pressure-related functions but may be brittle and difficult to join. Material names alone do not resolve performance: manufacturing defects, surface preparation, coatings, fastening, and site-specific chemistry are often decisive.

Salinity, dissolved oxygen, temperature, microorganisms, and the relative electrochemical potentials of joined materials shape corrosion mechanisms. A small noble-metal component coupled electrically to a larger susceptible component can change the latter's exposure; incorrect fasteners or sacrificial elements can undermine an otherwise robust assembly. Evidence should include representative coupons and full assemblies, cyclic loads, galvanic couples where relevant, crevices, connector contact resistance, and corrosion products. Results from artificial seawater provide controlled mechanistic evidence but cannot automatically substitute for field histories containing microbial communities, sediment, changing flow, and periodic drying. Research on welded nickel–aluminium bronze, for example, explicitly compares natural and simulated seawater and the relation between colonization and corrosion [33].

Marine biofouling introduces additional complications. Colonizing organisms alter surface roughness, drag, optical windows, acoustic transducer response, and access to joints. Fouling may increase energy demand, obstruct a manipulator, bias sensors, or change the material–water interface. Conversely, eliminating all colonization is not an ecologically neutral objective: coatings and cleaning processes may affect nearby organisms or disperse biological material. Reviews of marine antifouling systems document the continuing tension among performance, durability, toxicity and evaluation methods [9, 63]. The International Maritime Organization's Anti-fouling Systems Convention imposes specific restrictions for covered ships and offshore structures, including organotin and cybutryne controls; its applicability must be determined for the particular object and jurisdiction rather than mechanically extended to every small robotic instrument [74].

Material standards supply evidence only within their express scope. ISO 12944-9:2018 establishes requirements and laboratory test methods for protective paint systems on certain offshore and related *carbon-steel structures*, including specified corrosivity and immersion categories; it does not generally certify titanium, polymeric robots, electronic connectors, or the ecological benignity of a coating [76]. The principle is transferable without extending the legal or technical claim: list the actual applicability conditions before using a standard as a source of confidence. A laboratory corrosion pass for a steel coupon is a component-level observation; mission reliability additionally depends on housing penetrations, moving joints, dynamic seals, electrical contacts, structural interfaces, and replacement procedures.

18.5 Fabrication routes and the distinction between printed prototypes and qualified structures

Robotic manufacture encompasses machining, molding, casting, composite layup, bonding, welding, additive manufacturing, modular assembly, and quality assurance. Each technique constrains the geometry and material behavior available to designers. Machining supports known tolerances and inspection routines for certain rigid parts but may introduce high cost for intricate internal passages. Casting and molding support soft fluidic chambers and complex compliant geometries, yet can trap voids or create variable wall thickness. Fiber-reinforced composites offer favorable mass–stiffness relations for some

structures but complicate repair and inspection. Additive manufacturing expands accessible geometry and can integrate channels, variable stiffness regions and sensing mounts, while making print orientation, layer adhesion, porosity, residual stress, curing and material traceability central verification variables.

The multi-material print-and-cast arm described by Qi and colleagues illustrates how fabrication decisions are inseparable from actuation models [157]. Multi-material embedded printing has likewise been used to manufacture multifunctional silicone-based soft robotic components [95]. These are important research demonstrations. A printed soft gripper that succeeds on a laboratory bench has not thereby been qualified for repeated deep deployments, chemical isolation of biological samples, or safe disposal of damaged polymer components. Qualification requires physical tests directed to the relevant failure mechanisms, including accelerated ageing supplemented by representative long-duration exposures, bond strength after seawater conditioning, leak testing, abrasion resistance, and end-of-life inspection. Accelerated tests can inform mechanisms and comparative rankings, but their time-to-failure extrapolations require an identified and validated acceleration model.

An especially important distinction concerns *factory fabrication* versus *in-situ marine construction*. Printing a component on land for subsequent robotic installation is different from depositing material directly into a water-exposed structure. For submerged concrete additive manufacturing, recent reviews identify washout, fresh-mixture rheology, early-age stability, bond formation, robot nozzle control, and deposition inspection as significant unresolved engineering questions [127, 174]. These sources motivate controlled development, not an assertion that autonomous seabed printing is already industrially generalized. Likewise, submerged welding or patch repair must be evaluated against defect formation, heat-affected material, visibility, diver or robot access, and post-repair inspection. The fact that a motion planner can follow a toolpath does not establish the structural performance of the fabricated result.

For reproducible research, each manufacturing artifact should have a traceable configuration record: material batch and composition, manufacturing route, versioned geometry, processing and curing parameters, component serial identification, acceptance tests, assembly interfaces and known deviations. Such records matter for scientific reproducibility as well as engineering maintenance. If two nominally identical grippers exhibit different grasp performance, the difference may arise from casting bubbles, cable routing, post-curing or aging rather than from a revised controller. Retaining provenance allows investigators to distinguish material change from policy change, thereby extending the epistemic concerns of Sections 18 and 17 into manufacturing practice.

18.6 An evidence-graded comparison of material and actuation pathways

Table 15 aligns several families of devices with what has actually been demonstrated and what remains to be verified. The table is intentionally not a universal ranking. The same actuation family can support radically different performance depending on size, payload, pressure, contact constraints and servicing context. Evidence maturity refers to the stated class of use, not to an intrinsic maturity score for a physical principle.

Three methodological observations follow. First, evidence quality is task-specific: successful motion under pressure does not validate an actuator for precision tool use, and a successful single grasp does not establish fatigue endurance. Second, an engineering capability can be real while its durability and environmental consequences remain uncertain. Third, a comparative study must record the *unit of comparison*. Energy per unit displaced fluid, energy per completed task, work at an end-effector and battery energy per survey are not interchangeable. Nor is the mass of a printed shell a complete estimate of the environmental footprint of its production, replacement and disposal.

Comparative research should therefore use matched task definitions and report, at minimum, output load, deformation, precision, response time, power boundary, failure events, repair time, material loss and appropriate ecological measurements. A mechanical test may be combined with microscopy, water-chemistry analysis, biological sampling-control experiments, and post-test disassembly. The resulting comparison will be multidimensional. This is consistent with the general argument that heterogeneous values and forms of scientific capacity cannot be reduced to a single monetary or performance metric

Table 15: Selected actuation and fabrication pathways: evidence and bounded interpretation. The examples indicate reported capabilities, not blanket qualification for every marine mission.

Pathway	Evidence example	Engineering opportunity	Unresolved field-qualification issue
Electric/hydraulic transmission	Established marine propulsion and intervention practice; mechanical modelling literature [176]	Regulated thrust, high load or manipulability with serviceable modules	Seals, leakage, joint and power-chain efficiency, servicing
Compliant grippers and cables	Deep-reef sampling [48]; printed-endcap cable arm [157]	Distributed contact and flexible reachable geometry	Fatigue, hysteresis, bio-sample compatibility, bonded-joint ageing
Dielectric elastomer swimmer	Trench fin actuation and separate deep-water free swimming [93]	Pressure-tolerant distributed soft actuation	Endurance, insulation, high voltage, mission payload, recovery
Electrohydraulic artificial muscles	Bench experiments and soft grasping [4]	High-strain, electrically addressed compliant actuation	Submerged operation, fluid retention, material ageing
Multi-material fabrication	Embedded soft printing demonstrations [95]	Integration of channels, structures and functions	Process variance, inspection and long-term materials qualification
In-situ deposition	Emerging submerged concrete-printing research [127]	Potential localized repair and construction	Washout, layer bonding, structural test evidence, ecological effects

merely because a common dashboard is convenient.

18.7 A bounded degradation-budget model for responsible actuator use

A technically useful synthesis connects material uncertainty to permissible operational loading without equating a mathematical estimate with certified durability. Let $D(t) \geq 0$ be a scalar *validated proxy* for damage to a specified component, such as a calibrated fatigue variable or loss of a measurable integrity margin. For a bounded mission horizon $[0, T]$, suppose controlled exposure tests justify, over a stated range of temperature, salinity, pressure and mechanical loading, a differential inequality

$$\dot{D}(t) \leq a_* |u(t)|^p + b_*, \quad a_*, b_* \geq 0, \quad p \geq 1, \quad (65)$$

for almost every t , where u is a correctly normalized load measure, preferably dimensionless relative to a reported reference load. If D is dimensionless, both $a_* |u|^p$ and b_* must have units of inverse time; if a dimensional damage proxy is used, their units must instead match D per unit time. The nonnegative term b_* bounds uncommanded environmental degradation within the represented conditions. No claim is made that all degradation follows this law: corrosion, joint cracking, insulation breakdown and biofouling can require distinct, interacting state models, and a catastrophic defect can violate these assumptions.

Proposition 18.1 (Conditional finite-horizon damage budget). *Assume D is absolutely continuous on $[0, T]$, that inequality (65) holds almost everywhere, and that $D(0) = D_0$. For a validated threshold $D_{\text{lim}} > D_0$, if*

$$a_* \int_0^T |u(t)|^p dt + b_* T \leq D_{\text{lim}} - D_0, \quad (66)$$

then $D(t) \leq D_{\text{lim}}$ for every $t \in [0, T]$. This implication is restricted to the damage proxy and environmental/model domain for which the assumed upper bound is justified.

Proof. Integrating inequality (65) gives

$$D(t) - D_0 \leq a_* \int_0^t |u(s)|^p ds + b_* t.$$

All terms added to D_0 on the right are nonnegative, so the right side does not exceed its counterpart evaluated at T . Applying (66) yields $D(t) \leq D_{\text{lim}}$ for the claimed interval. $\square$

The proposition is deliberately narrow. It can support an operational monitoring policy only after the chosen proxy, exponent and bounds have been calibrated and checked against relevant observations. It does not show that $D(t)$ captures every failure mechanism, that the model assumptions persist after a seal rupture or unexpected collision, or that remaining below D_{lim} protects marine organisms. The threshold must represent a stated engineering acceptance criterion and cannot be used to certify legal authorization. For a heterogeneous multi-component robot, separate proxies and interacting failure pathways may be essential. If a load limit is repeatedly approached, the actionable inference is a need for examination, reduced duty, redesign or withdrawal within a review procedure, rather than an unrestricted extension of service life.

The model nevertheless clarifies a valuable comparison: maintaining high instantaneous actuator performance by using larger duty cycles can consume the modeled integrity budget sooner, while an architecture with lower peak capability may permit more repeated work within its validated damage envelope. Importantly, repairs, material replacement and requalification alter the state description, and should be entered into the mission history rather than erased by resetting a software counter. Coupling this monitoring approach with Sections 14 and 17 supplies a researchable interface between controller behavior, material condition and intervention permission.

Fabrication trials, reviews, and structural qualification. Recent submerged concrete additive-manufacturing publications document active development of mixture rheology, washout prevention, deposition stability, interlayer bonding and robot-process coordination [127, 174]. These reviews map design and material challenges; they do not themselves constitute a repeated, independently certified offshore construction trial. A controlled specimen test may provide evidence for fresh-state printability or a particular hardened-material property, but the inference to load-bearing service requires appropriate dimensions, cure history, fatigue, corrosion, substrate interfaces, repair strategy, and inspection. Material loss to the environment and biological exposure are additional research objects beyond structural strength. The strength of the emerging literature is that it makes the coupled material–process problem investigable. A defensible technology-readiness statement should specify whether an outcome was obtained on a bench, in a tank, on a deployed prototype, or during monitored operational service, with each step supported by its own records.

18.8 Repairability, ecological consequences, and the engineering research agenda

A robot that operates successfully once and is discarded is a different infrastructural proposition from one that can be inspected, recovered, repaired and deployed repeatedly. Mission economics often favors high performance and low mass, yet ecological accountability also requires attention to consumed materials, service vessels, spare parts, chemical leakage, discarded batteries, detached coatings and exposure introduced by repeated recovery trips. Repairability is therefore a design variable rather than a late logistical concern. Examples include modular actuator cartridges, interchangeable end effectors, accessible electrical penetrators, seal replacements, inspectable joints, mechanically separable material combinations and records of reworked components. A bonded or potted assembly may provide short-term environmental resistance while making field diagnosis and recycling more difficult. These tradeoffs should be documented rather than hidden in a nominal durability index.

The ecological implications of material selection arise both during manufacture and within the marine environment. A coating chosen to prevent drag-increasing fouling can release substances that affect organisms; cleaning may release detached biological material and particles; and lost synthetic components can persist after the intended mission. Conversely, untreated surfaces may accumulate organisms whose removal, transport or habitat effects warrant assessment. The legal treatment of antifouling substances, where applicable, and the methodological standards for environmental assessment offer relevant constraints but do not resolve the full ecological question of an individual robotic mission [74, 63]. Where evidence is incomplete, investigators should report what was measured and what remained outside the observation window, especially for chronic and cumulative effects.

A concrete research programme can compare three candidate intervention modules: a conventional joint-driven gripper, a pressure-fluidic compliant gripper, and an electrically driven soft actuator. The study would use a common set of representative manipulation tasks without assuming that all are appropriate for live-animal contact. Its initial bench phase would measure force–displacement curves, hysteresis, actuation energy and interface stresses. A separate pressure/salinity/temperature phase would examine drift, seals, stiffness changes and insulation. A subsequent long-duration immersion phase would investigate degradation, fouling and repair requirements; open-water demonstration would follow only where the relevant approvals and ecological safeguards are established. Environmental measurements would distinguish chemical leaching and particulate loss from the mechanical success of the grasp. Failure reports and rejected design configurations should remain part of the scientific record.

Fabrication experiments should test whether added complexity actually produces new generative research capacity. A one-piece printed gripper may reduce part count but increase replacement waste; an easily swapped modular joint may require more seals yet make failures diagnosable. Hybrid rigid–soft designs may distribute loads more effectively but introduce debonding and differential thermal or pressure response. A scientific comparison must retain these tradeoffs alongside acquisition costs and instrument performance. The objective is the formation of warranted, revisable knowledge about the relations among material, manufacturing process, mechanical action and marine ecology.

This section completes the material and actuation foundations needed to evaluate a physically consequential robotic platform. Section 19 will separately investigate power architecture, hydrostatic pressure, thermal management, and integrated lifecycle reliability; the present section has addressed these only insofar as they condition actuation and material behavior. Section 20 will assess heterogeneous teams and sustained infrastructure operations. Thus a progression from biological morphology to material embodiment and from local actuation to collective maintenance can be investigated without identifying technological sophistication with ecological legitimacy. The unresolved research question is which manufactured systems can reliably expand scientific and operational capacity while remaining inspectable, repairable, and subject to ecological review throughout their actual material histories.

19 Power, Pressure, Thermal Design and Lifecycle Reliability

19.1 The endurance and integrity of ecological robotic inquiry

The preceding engineering sections developed a task-oriented understanding of morphology, hydrodynamics, control, sensing, navigation, manipulation, and material realization. This section investigates the conditions under which those capacities remain available over a specified mission. A robot may be able to execute an accurate trajectory, identify a biological signal, or operate a manipulator during a laboratory trial and nevertheless be unsuitable for repeated scientific encounters. Its energy reserve may fail during recovery, a pressure enclosure may accumulate irreversible damage, sensor heat may alter calibration, or a shared communications subsystem may disable apparently redundant functions. These limitations concern the production of reliable knowledge as much as vehicle survival: incomplete recovery of samples, unrecorded operating temperatures, and intermittent instrument failure can affect what a scientific mission warrants.

The engineering object is therefore a *mission-conditioned system history*, rather than a single nameplate endurance figure. Let a designate a platform architecture, π an operational policy, $\xi_{0:T}$ a partly observed oceanic environment over a proposed interval $[0, T]$, and r a recovery and maintenance plan. A provisional mission dossier may be represented as

$$\mathcal{D}_{\text{life}} = (a, \pi, \xi_{0:T}, r; \mathcal{E}, \mathcal{P}, \mathcal{T}, \mathcal{F}, \mathcal{C}), \quad (67)$$

where $\mathcal{E}$ comprises energy records, $\mathcal{P}$ pressure qualification, $\mathcal{T}$ thermal evidence, $\mathcal{F}$ failure and maintenance histories, and $\mathcal{C}$ ecological and institutional conditions. A nominal endurance claim becomes scientifically informative only when accompanied by the duty cycle, payload, ambient conditions, reserve rule, and failure criterion to which it applies. Those qualifications distinguish an instrument drifting for

years from an actively propelled vehicle, a cable-powered intervention robot from a battery-powered hybrid, and an autonomous observatory from a remotely monitored installation.

The differences are concrete. WHOI specifies its *Nereid Under Ice* vehicle as battery powered with a communications-only optical microtether; its published specification lists a 22 kWh lithium-ion battery and an operating depth rating of 5,000 m [217, 215]. The microtether’s communication function must not be mistaken for energy transmission. MBARI’s long-range *Tethys* architecture instead emphasizes low propulsion drag, adjustable electronics power, and intermittent higher-power scientific sensing [113]. Profiling Argo floats employ another strategy: low-duty-cycle buoyancy changes and periodic satellite reporting, usually yielding multiyear observational life under a different service and recovery model [219, 181]. None of these examples establishes a universally superior energy architecture; they demonstrate the dependence of endurance on the research task and its service environment.

System endurance also has an ethical dimension. Lost equipment can become persistent debris or an unanticipated habitat structure; batteries, hydraulic media, surface coatings, and housings may create release or recovery risks; and a power-efficient mission may require stronger acoustic communications or a longer cumulative encounter with animals. Low electrical consumption is therefore an important engineering measure, but not a sufficient measure of ecological responsibility. Sections 5, 10, and 18 provide the corresponding normative, operational, and material distinctions.

19.2 Energy architectures and mission-conditioned power budgets

Marine robotic energy systems include surface-supplied electric power, rechargeable or primary batteries, buoyancy-control work, local renewable harvesters, and stationary infrastructure supplied by cables or nearby generating installations. Hybrid mission designs may exchange reduced propulsion demand for larger batteries or periodically surface to receive energy and communicate. The relevant energy budget includes propulsion, attitude and buoyancy control, sensing, payload manipulation, processing, communications, battery management, heating where needed, and housekeeping losses. These quantities vary with operational state; continuous operation cannot be inferred by adding nominal maximum loads and dividing battery capacity by that sum without specifying simultaneous use, power conversion losses, and allowable depth of discharge.

For stored usable energy $E(t)$, an accounting model is

$$\dot{E}(t) = P_{\text{in}}(t) - \sum_{j \in \mathcal{J}} P_j(t) - P_{\text{loss}}(t), \quad E_{\text{min}} \leq E(t) \leq E_{\text{max}}, \quad (68)$$

where P_{in} is validated power delivered to the storage interface, P_j enumerates state-dependent consumers, and P_{loss} incorporates conversion and parasitic losses without double counting. Capacity, permissible discharge rate, voltage limits, and state-of-charge estimation change with temperature, state of health, chemistry, pressure-exposure design, and accumulated cycles. Equation (68) is a conservation-based representation, not a fitted battery degradation model. It also distinguishes stored energy (joules) from instantaneous electrical power (watts). Reported range, days at sea, and instrument duty cycle should be compared only after converting them to an explicitly defined mission measure.

Power scheduling can have greater leverage than adding battery mass. MBARI reports that its long-range AUV architecture was designed around modest average payload demand, including selective activation of instruments whose instantaneous draw is higher than the platform’s long-term average [113, 116]. A state-indexed policy can sample densely during a transient biological encounter and conserve energy during transit. However, conserving sensor power can increase observational gaps; deferring a calibration may save energy while worsening the validity of scientific inferences. Hence sampling frequency, detection probability, scientific resolution, and battery reserve are joint decision variables. The costs of such tradeoffs should be recorded as changed epistemic opportunity and uncertainty, not hidden within a single efficiency percentage.

Energy provision differs sharply across system classes. A conventional work-class remotely operated vehicle may obtain high-power electric supply through a substantial umbilical, allowing sustained

lighting and manipulation but introducing tether handling, ship dependence, and deployment costs. A communications-only tether, such as that of Nereid Under Ice, permits more lateral range while leaving the vehicle subject to onboard battery accounting [215]. Long-lived floats minimize energy-intensive operations, while docking or inductive-charging concepts for mobile vehicles require verified coupling efficiency, docking robustness, power-transfer safety, and a defensible analysis of seabed infrastructure impacts. Solar harvesting is generally available on exposed surface platforms; a submerged platform cannot treat incident sunlight as an assumed energy source, particularly at depth. Seawater and temperature gradients create possible energy-conversion opportunities, but their net power density, thermodynamic efficiency, materials, heat exchange, and local ecological effects demand site-specific evidence [90].

The practical budget must reserve energy for navigation uncertainty, weather or currents, communication attempts, fallback control, emergency ascent where physically possible, and recovery. A nominal reserve stated as a fixed percentage may be inadequate if the required return journey is strongly direction-dependent. A more informative reserve function $R(t, \hat{\xi}, \mathcal{U})$ depends on time, estimated location and conditions, and an explicit environmental uncertainty set $\mathcal{U}$. Its values should be empirically estimated and may need immediate revision when unexpected currents or equipment faults are detected. In particular, research planners should distinguish a vehicle's capacity to reach its destination from its capacity to cease an intrusive operation and return without creating additional disturbance.

19.3 Hydrostatic loading and pressure-boundary architectures

Hydrostatic pressure is a depth-dependent environmental load with consequences for enclosures, energy storage, buoyancy elements, connectors, seals, instruments, and manufacturing choices. In a vertically referenced approximation, the absolute pressure at depth $z \geq 0$ is

$$p_{\text{abs}}(z) = p_{\text{surface}} + \int_0^z \rho(\zeta)g(\zeta) d\zeta. \quad (69)$$

The familiar rule of roughly one additional atmosphere per ten metres follows from representative seawater density and gravity, but is not an exact conversion valid without qualification at great depth [123, 79]. For an enclosure the governing load is often the *differential* pressure $p_{\text{ext}} - p_{\text{int}}$, which depends on how its interior is configured. An apparently small connector, window, cable entry, or sensor interface may be the critical pressure boundary even when the principal pressure housing passes a static calculation.

Three broad design families illustrate the tradeoffs. First, a relatively dry, pressure-isolated housing protects electronics and batteries within a mechanically qualified boundary. It needs geometrically appropriate wall design, buckling analysis, sealing, feedthrough integrity, and validation under appropriate load cycles. Second, an oil-filled or otherwise pressure-compensated compartment reduces differential pressure on selected components; its components must still tolerate the chosen fluid, permeation, pressure-related changes in dielectric properties, and thermal expansion. Third, directly pressure-tolerant electronics or materials may avoid a large dry housing but require device-specific testing, insulation, packaging, chemical compatibility, and verified fault response. Mixed architectures are common; design documentation should identify the pressure strategy of each subsystem. These families are not interchangeable merely because their depth ratings appear numerically similar.

Pressure qualification is distinct from service-life validation. A single successful descent gives evidence that specified components functioned during that history; it does not bound slow leakage, cycle-driven fatigue, seal deterioration, or corrosion over many missions. External pressure may raise buckling concerns that are not captured by simple tensile stress comparisons. Testing should distinguish hydrostatic hold, pressure cycling, thermal cycling, vibration and handling, combined loading, and post-test functional inspection. Published underwater-vehicle classification rules, including those of the American Bureau of Shipping, provide relevant engineering frameworks, but a statement that a design follows a standard is not equivalent to documentation of specific compliance or classification [10]. This paper claims no certification or approval for any proposed platform.

Pressure also changes physical and scientific measurement. A deep-water sample brought to the surface may undergo pressure-induced transformations; a pressure-conditioned sensor can respond dif-

ferently outside its calibration range. A thermally contracting enclosure alters internal clearances and connector preload, potentially coupling temperature to sealing performance. Scientific provenance should therefore include deployment depth, pressure histories, and relevant test limits, especially when results are compared across dissimilar habitats or laboratories. The biology discussed in Section 4 supplies an additional caution: organisms adapted to high ambient pressure cannot be presumed physiologically unaffected by sampling or decompression simply because a robot survives those conditions [134].

19.4 Thermal management and the ecological accounting of heat

Cold ambient seawater may appear to offer abundant cooling capacity. Yet a robot's internal temperature depends on actual heat-generation pathways, enclosure geometry, insulation, internal fluid motion, external flow, and the thermal conductivity of interfaces. A fully enclosed battery or computing compartment can develop local hot spots despite low ambient temperature. Conversely, components designed for a warmer laboratory may lose usable battery capacity or actuator performance in cold deep water. For a simplified lumped interior node with temperature T_i and environment temperature T_w , an energy balance is

$$C_i \dot{T}_i = Q_{\text{gen}}(t) - U_i A_i (T_i - T_w) + Q_{\text{coupled}}(t), \quad (70)$$

with effective heat capacity C_i in joules per kelvin, area A_i in square metres, heat-transfer coefficient U_i in watts per square metre per kelvin, and additional thermal exchanges Q_{coupled} in watts. Each term in the displayed balance is thus a power; C_i is a heat capacity, rather than a thermal conductivity. The model is appropriate only when its lumped-temperature assumption and coefficients are experimentally justified. Battery cells, electronic components, motor windings, enclosures, and external heat exchangers may require spatially resolved models or measured transient transfer functions.

Even when electrical input is known, heat production is not in general equal to input power: some energy becomes external mechanical work, stored chemical change, emitted light, or acoustics, while a fraction may dissipate at another location. A thermal test should locate losses, distinguish average from peak dissipations, and include degraded cooling conditions. Changes in salinity, sediments, biofouling, currents, and deployment orientation can change exchanger performance. Robust temperature limits must consider both overheating and excessively low temperatures, including available heating power, startup transients, and sensor calibration drift. It is inadequate to infer thermal safety from the presence of seawater alone.

For a steady-state simplification of Equation (70) with negligible Q_{coupled} and fixed T_w , $T_i - T_w = Q_{\text{gen}}/(U_i A_i)$. Thus a measured dissipation and an assumed ocean temperature cannot determine component temperature without characterizing the effective conductance. When computing, battery storage, and subsea installations are integrated, the relevant thermal field extends beyond the housing. Rejected heat is transported and mixed in the receiving marine environment; ecological consequences depend on plume geometry, duration, gradients, species, and local baseline variability. The claim that seawater is a convenient cooling medium is a statement about an engineering possibility, not evidence that heat rejection is environmentally negligible.

Microsoft's Project Natick offers a bounded example. Its North Sea experimental installation investigated subsea computing, cooling and hardware reliability; the company reported a lower observed server-failure rate than a land-based comparison group, within that project's conditions [110]. Such evidence cannot be generalized to untested data-centre locations, different enclosures, larger heat fluxes, altered energy sources, or ecological impacts without additional studies. Ocean thermal energy conversion research similarly explores thermal-resource exploitation with its own temperature and exchanger requirements, not an unlimited source of usable power or a universal way to dissipate computing heat [90]. The hypothetical A–B–C–D installation analyzed in Sections 7–9 consequently requires both a thermodynamic performance model and independent ecological monitoring of actual receiving conditions.

19.5 Integrated reliability, fault propagation, and mission recovery

Reliability cannot be inferred by multiplying component manufacturers' headline reliability values without establishing independence, operating conditions, failure definitions, and maintenance states. Oceanic platforms may experience common-cause failures: loss of a power-distribution unit may disable navigation, sensor logging, and communication; a flooded junction can invalidate apparently separate redundant paths; one navigation fault can produce an unsafe interaction that damages both the vehicle and an installation. A useful failure model begins with mission functions and traces dependencies to sensors, actuators, electrical distribution, software, connectors, enclosures, and recovery systems. Failure-mode and effects analysis, fault-tree methods, reliability block diagrams, and dynamic simulation are complementary tools, each limited by its data, state assumptions, and event coverage. A recent review of autonomous-vehicle reliability emphasizes subsystem-specific failures, incomplete observations, and the limitations of available reliability methods [23].

For a component with a well-defined failure event and an appropriate conditional hazard rate $\lambda(t | h_t)$, a survival model may be written

$$S(T | h) = \exp \left[- \int_0^T \lambda(t | h_t) dt \right]. \quad (71)$$

Here h_t represents specified stress and maintenance history. This representation does not assert a constant failure rate, an exponential service lifetime, or independent failures across components. Estimating λ requires compatible censored field histories, credible failure attribution, and an account of changes in design, mission, and maintenance. Where such records are lacking, uncertainty ranges and qualitative failure analyses are preferable to fabricated precision. Time-to-failure and probability of scientific-data validity are also distinct outcomes; an instrument may continue transmitting while its calibration is compromised.

The failure response should be mission-specific. In a free-swimming survey robot, failure detection may trigger a reduced-power return or a safe recovery trajectory; a deep vehicle may need controlled ascent and a separate surface communications strategy. In an intervention system, aborting a task could leave a suspended tool, an unlatched fixture, or a disturbed sample. A cabled installation may require isolation and logging prior to repair, and an autonomous surface craft may need collision-avoidance safeguards during degraded operations. The practical hierarchy distinguishes *detect*, *isolate*, *limit secondary harm*, *recover*, and *document*, but which actions are feasible depends on the current physical configuration and lawful operating conditions.

A robust scientific mission also maintains evidence continuity. Battery health records, pressure-leak indicators, thermal excursions, watchdog restarts, and sensor-health flags should be synchronized with observation timestamps. A measurement produced after unrecognized cooling failure may not be comparable with earlier data; a mission recovered intact may nonetheless require qualified exclusion of a subset of observations. This connection between engineering health and epistemic validity extends the provenance commitments developed in Sections 6 and 15. Long-duration observations are particularly vulnerable to calibration drift and biofouling; the Ocean Observatories Initiative's biogeochemical sensor guidance describes quality-control and correction issues that cannot be reduced to hardware uptime [145]. A reliable observation infrastructure therefore requires a maintenance plan for both the robot and the credibility of its reported knowledge.

19.6 Maintenance, repairability, and ecological lifecycle costs

A full lifecycle description follows the system from material extraction and manufacture through transport, deployment, recharge, inspection, calibration, repair, replacement, retirement, and recovery. Different activities distribute burdens differently. A highly reliable sealed device may be difficult to repair; a modular instrument may simplify component replacement but add connectors and leak paths; a remotely operated intervention fleet may reduce the need for repeated human dives while depending on energy-intensive support vessels. The choice cannot be settled by vehicle electrical power alone. An

assessment should state its functional unit, such as validated observations per spatial-temporal coverage, qualified intervention cycles, or years of maintained observatory performance, and disclose excluded support-system costs.

Maintenance is also a site-level question. Recovering a sensor every month may permit calibration, cleaning, and repair but increase vessel traffic and recurring habitat disturbance. Extending the interval may reduce surface interventions but permit substantial data drift or undetected loss of energy reserve. Sea-Bird's calibration recommendations differentiate profiling and moored sensors and emphasize field comparisons for detecting drift, while recent field investigations document the interaction of fouling and pH sensor performance [165, 102]. Neither establishes one universal service interval. Engineers must condition intervals on local fouling, operating hours, sample chemistry, equipment type, and the consequence of scientifically misleading output.

Repairability calls for documented material interfaces, connector lifetimes, modular fault isolation, replaceable seals, environmentally defensible lubricants, and recoverable components. Some architectures may benefit from robotic self-inspection or remotely conducted maintenance, but the field evidence for autonomous fault repair is less extensive than for fault detection. A review of engineering and maintenance on autonomous surface ships identifies significant unresolved practical requirements for automation, spare components, and remote intervention [8]. The distinction between diagnosing a malfunction and restoring a pressure-qualified subsystem should remain explicit. Moreover, a component repaired at sea may have to undergo renewed pressure and electrical qualification before returning to full mission authority.

Ecological lifecycle accounting records potential material losses, heat releases, construction footprints, support-ship emissions, acoustic activity, pressure-fluid escape, entanglement hazards, and persistent equipment following failed recovery. It should also identify potential contributions to scientific capacity and environmental monitoring, with neither benefit presuming permission for every intervention. The commons-based circulation framework of Section 9 supplies a complementary inquiry: who can inspect maintenance histories, challenge failure claims, access designs, and undertake repairs? These procedural capacities influence whether an installation's long-term costs and benefits are visible across different institutions and ecological relationships. No scalar lifecycle score is presumed to exhaust these heterogeneous concerns.

19.7 A robust energy-reserve proposition under bounded uncertainty

The system history motivates a limited formal result. Define a controlled mission interval $[0, T]$ and suppose stored usable energy evolves through measurable delivered supply and consumption. Let $H(t)$ represent actual received supply and $C(t)$ actual aggregate consumption including losses, and let calibrated uncertainty envelopes $\underline{H}(t)$ and $\overline{C}(t)$ be specified for this mission. The envelope functions need not be constant. Here E , $E_{\min}$ and the reserve $R(t)$ must be expressed in the same units of usable energy (e.g. joules), whereas H , C , $\underline{H}$ and $\overline{C}$ are powers (e.g. watts), integrated over a consistent time unit. For a recovery-reserve requirement $R(t) \geq 0$ and a fixed lower permissible energy $E_{\min}$, the conservative accounting trajectory is

$$\underline{E}(t) = E(0) + \int_0^t [\underline{H}(s) - \overline{C}(s)] ds. \quad (72)$$

Proposition 19.1 (Conditional energy-reserve preservation). *Assume E is absolutely continuous on $[0, T]$ and satisfies $\dot{E} = H - C$ almost everywhere; $H(t) \geq \underline{H}(t)$ and $C(t) \leq \overline{C}(t)$ almost everywhere; and $E(0)$ is known within the chosen initial accounting convention. If*

$$\underline{E}(t) \geq E_{\min} + R(t) \quad \text{for every } t \in [0, T], \quad (73)$$

then $E(t) \geq E_{\min} + R(t)$ throughout that interval.

Proof. By the differential balance and the almost-everywhere supply and consumption bounds,

$$E(t) - \underline{E}(t) = \int_0^t [H(s) - \underline{H}(s) + \overline{C}(s) - C(s)] ds \geq 0.$$

Combining this inequality with (73) gives the stated bound for each t in the modeled horizon. $\square$

The proposition is deliberately narrow. It establishes a conditional conservation-accounting implication, not a novel battery model or a certified return capability. Its premises can fail because of degraded state-of-charge estimation, latent internal faults, pressure-related capacity loss, power electronics failure, unmodeled loads, or uncertainty sets that omit a strong current. The function $R(t)$ itself must be justified by navigation, recovery, and environmental evidence, and remains only a modeled reserve rather than a guarantee that a damaged robot can use it. There is also no derived guarantee concerning maximum thermal stress, depth-rating compliance, collision, animal welfare, or lawful mission continuation. A resilient operational policy should therefore combine the reserve calculation with independent fault monitoring, thermal and pressure limits, an authorized recovery plan, and a supervisor able to revise the mission as its uncertainty grows.

The proposition permits an auditable comparison of planning regimes. One plan may use an optimistic mean current and low average instrument demand; another may use a calibrated envelope with a reserve for communication attempts and adverse transit conditions. If the second rule excludes an otherwise attractive mission, it provides evidence that the mission requires different energy, duty cycle, recovery infrastructure, or empirical characterization. It does not prove that the first plan would have failed. A useful experimental test would compare predicted lower bounds with actual energy histories across controlled perturbations, recording false exclusions, missed deficits, and the scientific measurements forgone because of conservative shutdown decisions. This establishes a pathway for falsification and improvement without turning local bounds into unrestricted claims.

Failure-mode catalogues are not mission reliability estimates. The published review by Chen and colleagues surveys reliability concepts and multiple subsystem failure categories for autonomous ocean vehicles [23]. Its synthesis of published investigations helps identify failure modes, but does not yield a transportable failure probability for an arbitrary vehicle configuration or ecological mission. A manufacturer's endurance figure, pressure rating, or nominal component life also differs from demonstrated availability under a specified mission cycle. Reliability claims require defined failures and exposure denominators, operational logs, environment and maintenance histories, dependencies among subsystems, recovery effectiveness, and a treatment of sparse or censored observations. In particular, two nominally redundant sensors may share electrical power, clock, pressure housing, or data processing and therefore fail together. A lifecycle comparison additionally needs an explicit functional unit and assumptions about servicing voyages, consumables, lost vehicles, and disposal; superior endurance does not itself imply a lower ecological burden.

19.8 Comparative qualification and the transition to heterogeneous operations

Table 16 organizes the present engineering evidence in terms of its scope. The distinction between *documented operation*, *design-rule guidance*, and *proposed generalization* is essential. A measured platform performance figure applies to a particular vehicle and configuration; a pressure code governs specified structural or operational classes; and a mathematical energy envelope applies only if its input bounds have been substantiated. These are complementary but differently warranted forms of evidence.

A staged validation programme follows. At component scale, laboratory testing should measure electrical energy conversion, discharge capability across realistic temperatures, housing and connector leak response, thermal gradients, actuation heat, and aging after pressure cycles. At subsystem scale, engineers should impose simultaneous loads rather than test electrical power, pressure integrity, and cooling in isolation. At platform scale, mission trials should record energy supply and consumption, recovery reserves, failure detection, communication loss, instrument drift, and time-to-recovery. At ecological scale, researchers should separately monitor animal encounters, heat and noise fields, material releases, habitat effects, and disturbance associated with maintenance vessels. At institutional scale, test reports and environmental conditions must remain accessible to competent reviewers with a pathway for contesting unsupported safety or reliability claims.

Table 16: Illustrative architectures, demonstrated evidence, and remaining qualifications for long-duration oceanic operations. Entries are research comparisons rather than certifications or universal performance rankings.

Architecture or source	Evidenced capability or design basis	Qualification needed for broader inference
Battery-powered hybrid NUI	Onboard electrical supply, optical communication tether; published depth and battery specifications [217]	Payload duty, thermal history, recovery reserve, intervention cycles
Long-range AUV (MBARI)	Low-drag transit, intermittent sensor operation; reported long-range platform history [113, 116]	Currents, high-load tasks, payload-specific endurance, battery aging
Argo profiling floats	Low-duty-cycle profiling and periodic satellite telemetry over multiyear service lives [219]	Different sensor loads, chemistry, calibration drift, actual recovery
Pressure-qualified components	External-pressure and underwater system guidance [10]	Scope-specific design review, cycling and connector qualification
Subsea computing experiment	Project Natick’s experimental server operation and reported reliability [110]	Larger deployments, plume ecology, location changes, full lifecycle
Long-lived observatory sensors	Quality-control guidance for drift and biofouling [145]	Independent calibration, long-term uncertainty and corrective monitoring

The hypothetical transnational marine data-center case provides a demanding integration scenario. Its operators might report excellent compute uptime while experiencing increased cooling discharge, uncertain effects on a nearby benthic habitat, or the need for frequent support-ship intervention. Investors could observe stable financial returns while scientists observe sensor degradation or missing ecological baselines. These discrepancies cannot be corrected by adding server reliability statistics to robot uptime or by treating the installation’s continued operation as an ecological endorsement. Rather, the case calls for typed, traceable records of energy sourcing, pressure-boundary maintenance, heat transport, biological observation, data validity, incident response, and decision authority, consistent with the distinctions in Sections 9 and 10.

The resulting picture of reliability is relational and conditional. Electrical endurance depends on the mission, pressure integrity on construction and exposure history, thermal viability on local transport, and scientific validity on maintained calibration and provenance. A long-lived platform can improve ecological inquiry when these dependencies are actively tested, documented, and revised. Endurance itself does not turn any platform into an epistemically reliable observer or a justifiable intervention. The following section will address how these component and vehicle limits change when multiple robotic systems, installations, and institutions share energy, tasks, navigation references, repair capabilities, and scientific records.

20 Heterogeneous Robot Teams and Marine Infrastructure Operations

20.1 Collective operations as an engineering and epistemic problem

Sections 12–19 considered the bodies, dynamics, sensing, actuation, and reliability of particular robotic systems. Oceanic scientific inquiry and infrastructure service nevertheless frequently depend on relationships among different vehicles, instruments, people, and fixed installations. A surface vessel may transport a submerged sampler’s observations to shore; a glider may profile the surrounding water while a stationary node records rapid local variability; a tethered remotely operated vehicle (ROV) may attach a connector that an autonomous vehicle cannot presently manipulate with comparable assurance. These are functional relationships, rather than interchangeable instances of an abstract robot. They require an account of which capacities are combined, how decisions are shared, what happens when communication or equipment fails, and how the resulting evidence can be understood.

The research object of this section is a *heterogeneous marine operational system*: a bounded configuration of mobile agents, stationary instruments, supporting vessels and shore institutions assembled for

scientific and practical purposes. Heterogeneity has several dimensions. Morphological heterogeneity concerns propulsion, pressure tolerance, manipulators, and environmental coupling. Epistemic heterogeneity concerns optical, acoustic, chemical and biological measurements, their sampling scales, and their calibration histories. Organizational heterogeneity concerns operators, institutional permissions, data custodians, and the responsibilities attached to activities performed in different maritime zones. Temporal heterogeneity is equally important: an ocean profile acquired over hours may be compared with high-frequency moored observations, while manipulation requiring seconds of feedback may coexist with communications delayed by minutes or longer. These distinctions connect fleet engineering to the relational and institutional inquiries of Sections 6, 7, and 9.

Let $\mathcal{R}$ denote deployable robotic assets, $\mathcal{I}$ stationary scientific and infrastructural assets, $\mathcal{H}$ human operational roles, and $\mathcal{T}$ tasks. For a mission interval $[0, T]$, define an operational dossier

$$\mathcal{D}_{\text{team}} = (\mathcal{R}, \mathcal{I}, \mathcal{H}, \mathcal{T}, \mathcal{G}_{0:T}, \Xi_{0:T}, \mathcal{A}, \mathcal{P}), \quad (74)$$

where $\mathcal{G}_{0:T}$ represents time-varying communication and service relationships, $\Xi_{0:T}$ available environmental information and uncertainty, $\mathcal{A}$ legal and organizational scopes of action, and $\mathcal{P}$ scientific provenance and ecological precautions. This is a proposed descriptive schema, not an implemented distributed operating system. A team should count as successful only relative to defined scientific, engineering, ecological, and procedural outcomes. Aggregate area covered or number of vehicle-hours is insufficient to establish accuracy, safe intervention, or warranted biological conclusions.

20.2 Demonstrated coordination and the maturity of marine operational systems

The distinction between a cooperative field experiment, a managed observational network, and an autonomous infrastructure-construction capability is critical. McCammon and colleagues reported ocean-front detection and tracking with a heterogeneous team comprising four Slocum gliders and two robotic surface samplers during a two-week Gulf of Mexico deployment. Their approach combined sequential allocation Monte Carlo tree search, environmental estimation, a drifting reference frame, and a human decision-support interface [106]. This work demonstrates an extended, field-tested form of coordinated adaptive ocean sampling. It does not imply that the same system performed general seabed assembly, continuous fully independent operation, or ecologically certified intervention.

The Ocean Observatories Initiative (OOI) presents a different form of heterogeneity. Its observatory arrays link fixed moorings, mobile gliders, scientific sensors, maintenance vessels, and data-management systems. OOI describes gliders as mobile instruments supplementing fixed observing assets, with satellite communication and common cyberinfrastructure interfaces [147, 146]. Its quality-control practices combine automated tests, operator annotations, calibration and deployment records, and human-in-the-loop review [148]. An observatory is therefore a coupled scientific and operational system even when its mobile units do not negotiate tasks directly with each other. Coordination can be organizational and infrastructural as well as a real-time multiagent algorithm.

Ocean Networks Canada (ONC) likewise operates cabled observatories that provide power and communications to seafloor instruments, while vessels and ROVs install, replace, and inspect subsea equipment [143, 144]. Continuous scientific data delivery from a fixed network and its maintenance by a human-supervised intervention system should be described as distinct functions. Their integration illustrates a mature, serviceable infrastructure pathway without establishing continuous autonomous construction. MBARI develops targeted sampling by autonomous surface and submerged vehicles, with onboard feature detection and cooperation among instruments and robots as explicit research objectives [117]. This establishes a credible active development programme, while particular claims about coordinated field capability require individual experiment reports.

The broader task-allocation literature offers classifications and optimization methods, but it also documents why allocation depends on task duration, agent capabilities, information, and whether several robots must cooperate on the same task [52, 207]. The empirical distinction matters: simulation with identical robots and idealized communications is an appropriate theoretical test but weaker evidence for

Table 17: Selected evidence for heterogeneous marine operations. The evidence categories describe documented activities, not comparative autonomy scores.

Example	Documented configuration and activity	Boundary of inference
McCammon et al. (2021)	Four gliders and two robotic surface samplers; two-week coordinated ocean-front survey	Field evidence for an adaptive observation task, not for construction or general-purpose autonomy
OOI observational arrays	Fixed platforms, gliders, ships, shore operations, data services	Operational heterogeneous observatory; direct inter-robot autonomy varies by mission
ONC cabled observatories	Powered subsea sensors serviced by ROVs and vessels	Established remote infrastructure service, typically with substantial human control
MBARI targeted sampling	AUV/ASV feature recognition, sampling, collaborative research development	Programme-level capability descriptions need task-specific validation
Multi-AUV task allocation research	Scheduling, coalition and route optimization, including simulated cases	Algorithmic performance may not survive changing currents, links, or biological constraints

a heterogeneous, geographically dispersed oceanic operation. Table 17 records the relevant maturity boundaries; subsequent engineering evaluations must report platform configuration, operator intervention, environmental conditions, sample validity, and the extent of autonomous decision-making rather than adopting a generic label of “swarm” or “multiagent” as a performance result.

20.3 Capabilities, task allocation, precedence, and team composition

Team planning begins with a task ontology grounded in observable requirements. A sampling task may require a depth-qualified sampler, sufficient battery reserve, a calibrated chemical instrument, and a viable recovery procedure. A seafloor connector task may require station keeping, an aligned manipulator, verified tooling, and a communication pathway to an authorized operator. Some scientific operations require multiple vehicles sequentially, such as identifying a front with one sensor and then characterizing it with a more selective sampler. Others require simultaneous cooperation, such as maintaining a geometric baseline for acoustic localization or coordinating a mobile manipulator with a sensing platform. Treating these arrangements as simple assignments of one task to one vehicle can obscure coalition requirements, timing, and shared bottlenecks.

Write $c_{ir}(t, \xi)$ for a validated task-relevant capability of robot i concerning requirement r , and q_{jr} for the requirement of task j . A provisional, time-indexed feasibility representation is

$$\sum_{i \in \mathcal{R}} c_{ir}(t, \xi) x_{ij}(t) \geq q_{jr} y_j(t), \quad r \in \mathcal{Q}_j, \quad (75)$$

where $x_{ij} \in \{0, 1\}$ denotes membership of i in the task’s assigned team and $y_j \in \{0, 1\}$ denotes task activation. This additive approximation can be useful for independently pooled capabilities, such as the number of independently calibrated sampling instruments. It is invalid when a requirement is intrinsically non-additive: a high-torque vehicle and a separate precision manipulator do not automatically become a single capable assembly platform. Joint feasibility must additionally include geometric, communication, temporal, and interface constraints. Thus, the capability representation is a preliminary filter rather than a certificate of coalition success.

A planning objective may combine scientific gain I_j , task value V_j , energetic cost E_{ij} , time and recovery cost K_j , and ecological exposure R_j :

$$\max_{x,y} \sum_{j \in \mathcal{T}} y_j (V_j + \lambda I_j) - \sum_{i,j} \alpha_{ij} E_{ij} x_{ij} - \sum_j (\beta_j K_j + \gamma_j R_j) y_j. \quad (76)$$

These weights are only a proposed decision-support representation. An ecological threshold or binding authorization condition cannot be made optional merely by assigning it a small penalty weight; applicable

hard constraints belong to the feasible set. Units and uncertainty in scientific gain, cost, and ecological exposure must also be disclosed. If the allocation objective combines task value, information gain, energy costs, and risk into a scalar score, its coefficients must have conversion or normalization units that make the summands commensurate within the *declared decision model*. Such weights express choices of a reviewing institution, not a discovered physical equivalence between scientific information, ecosystem effects, or rights. A vector-valued or lexicographic assessment remains appropriate where defensible trade-off weights are unavailable. The optimization may be hierarchical: institutionally authorized mission classes at the upper level; operator-reviewed task allocation in the middle; local trajectory and manipulator control below. Section 14 supplies the local-validity ideas, while Section 10 supplies the review logic.

Inter-task precedence and recovery require explicit scheduling. If a biological sampler must wait for a high-resolution survey to identify a candidate region, its arrival cannot be planned independently. A manipulator's actions may require a stabilized support vehicle, and a surface relay may need to remain within an acoustic or radio service region. Such dependencies form a temporally constrained task graph. Reallocation after equipment failure must recompute not only coverage of the abandoned task but also consequences for downstream tasks and the shared energy and communication reserves. A nominally idle robot may be essential as a relay, navigation reference, or recovery asset; treating it as a free substitute can create a new single point of failure. Technical heterogeneity must therefore be represented through concrete interfaces, qualification evidence and bounded substitution rules.

20.4 Communication, synchronization, and distributed operational authority

Distributed oceanic coordination cannot assume persistent high-bandwidth connectivity. Different links support different forms of cooperation: submerged acoustic messaging commonly trades low data rate and latency against greater reach; optical links can convey more data over relatively short, condition-dependent ranges; surface systems may use satellite or radio links; tethers and seabed cables support distinct high-throughput or power arrangements. The limitations reviewed in Section 16 become collective constraints when two vehicles need a shared view of the mission. Acoustic propagation time, multipath, changing sound speed, modem contention and sparse surfacing intervals can produce different local histories, even when all devices record their local data accurately.

Let $G_c(t) = (\mathcal{V}, \mathcal{E}_c(t))$ be a directed communication graph over robots, relays, fixed nodes and operator endpoints. A message may carry a claim about an observation or a task state rather than an authoritative world fact. A defensible envelope includes at least a sender identifier, clock basis, local timestamp, sequence number, frame and units, sensor calibration reference, estimated uncertainty, task authorization version and integrity check. Commands additionally require an issuer role, scope of permitted action and expiry or validity condition. These attributes do not guarantee authenticity or a synchronized global state; they make mismatches inspectable after a mission. Eventual synchronization of scientific datasets can be acceptable for retrospective analysis, whereas force-bearing manipulation or collision avoidance may require bounded-latency local feedback and may become inadmissible when that bound is lost.

A practical arrangement combines local autonomy with supervised mission-level planning. Vehicles should retain independently evaluated conditions for collision avoidance, abort, recovery and energy protection; the shore centre may revise a scientific plan when a link is available; two mobile robots may synchronize on a limited handshake for a short cooperative operation. Crucially, a missing acknowledgement should not be treated as agreement, and a stale authorization should not be interpreted as unlimited permission for a modified task. Proposed coordination policies should specify a *disconnection contract*: which actions can continue locally, which require fresh acknowledgment, which measurements are cached, how divergent versions are reconciled, and who is accountable for restoring or terminating the operation.

Synchronization also has a scientific dimension. Two platforms observing a dynamic front at different locations or times may produce apparently conflicting estimates because their measurements refer

to different parts of an evolving field. A common database alone cannot make their observations simultaneous or commensurable. Data assimilation must retain uncertainties in vehicle pose, clock offset, environmental advection and sensor response. When an acoustic relay also serves as a navigation reference, its relocation changes both communication and localization geometry; a planner optimizing link quality may inadvertently degrade the inferred position of observations. Team coordination therefore couples mission communications to the epistemic quality of the scientific record.

20.5 Distributed observation, information fusion, and ecological inference

Heterogeneous fleets can improve ecological inquiry by placing complementary sensors within different spatial and temporal regimes. A float profiles the water column at low energy cost; a propelled AUV follows local gradients; a moored observatory records longer histories at a fixed site; a mobile sampler obtains a specimen or biochemical assay at a strategically chosen event. However, the union of observations is informative only insofar as each measurement can be related to its environmental context and to the scientific question. The observation functions and their uncertainties remain distinct, as emphasized in Sections 6 and 15.

A conditional team information objective can be written as

$$\mathcal{J}(S \mid D_0) = I(\Theta; Y_S \mid D_0), \quad (77)$$

where Θ is an explicitly modeled ecological quantity, D_0 an existing evidence set, and Y_S observations that would result from task set S . Mutual information can guide experiment selection when a defensible probabilistic model and observation likelihood are available. It is not a direct measure of ecological value, organismal welfare, or the depth of a generative relationship. Moreover, the contributions of different platforms are generally not additive: two calibrated sensors may produce correlated evidence, and repeated sampling of a local anomaly may have diminishing informational return. In some settings complementarity can instead be synergistic, particularly when one observation resolves ambiguities left by another. An information-gain algorithm must disclose its prior, noise structure, dependence assumptions, and failures of transport from simulated to actual marine conditions [106, 117].

An operational data record should attach a *measurement lineage*: sensor and platform identity, pose and time reference, calibration and quality status, environmental estimates used for correction, data transformations, operator interventions and interpretive decisions. Existing OOI practice provides an instructive precedent: automated quality flags and human annotations accompany the data, while calibration and maintenance records contextualize apparent anomalies [148]. These practices are especially important for a fleet because a shared environmental estimate may be used by several agents and create correlated inference errors. A systematic navigation bias, an incorrectly assumed sound-speed profile, or a delayed time synchronization can propagate across the team's outputs even if each instrument individually passes nominal functional checks.

Scientific records should distinguish commands, intended tasks, attempted tasks, completed actions, measurements obtained and ecological inferences warranted. A commanded sample can fail to materialize; an obtained sample can be contaminated; an inferred biological event can reflect an instrument artifact. Conversely, an unexpected ecological signal should remain accessible for later reinterpretation even if an automated flag marks it suspicious. Such distinctions preserve the possibility of correction and challenge, rather than making a collectively managed dataset appear more certain merely because several platforms contributed to it. The design of common access and downstream value circulation remains dependent on the legal, institutional and epistemic conditions developed in Sections 8 and 9.

20.6 Collaborative intervention and persistent marine infrastructure service

Observation coordination and physical intervention impose sharply different system requirements. Mapping and adaptive sampling may tolerate delayed replanning, while coordinated manipulation near a connector or living habitat requires precise relative localization, robust force control, collision constraints,

and a clear recovery strategy. Cooperative lifting or installation introduces coupled loads and possible rigid-body constraints among vehicles. At a fixed observatory, installation may also depend on cable handling, vessel station keeping, connector protection, pressure-qualified interfaces and the availability of spare modules. Section 17 surveyed the mechanics of individual intervention systems; the additional team problem lies in interface and responsibility across those systems.

An illustrative observatory workflow comprises seabed surveying, identification of a permitted attachment site, equipment transport, remotely supervised connection, electrical and data verification, ecological inspection, and post-service monitoring. Different platforms may undertake distinct steps, but their interdependence should be treated as a validated operational sequence. ONC reports scheduled vessel expeditions in which ROVs recover, install and connect instruments to its deep-sea networks [144]. This is evidence for sustained remotely operated infrastructure maintenance. It does not demonstrate that independent underwater vehicles autonomously negotiate and execute every construction task. Similarly, promising multi-AUV manipulation and assembly algorithms should be characterized by the environment of their demonstration, the tasks performed, and the amount of human supervision before they are used as evidence for full oceanic construction capability.

A hypothetical A–B–C–D installation, introduced in Sections 7–9, clarifies the collective problem. A consortium could combine a surface ship and relay (A), an autonomous inspection vehicle (B), a manipulation-capable ROV and service crew (C), and host-jurisdiction scientific institutions and affected communities (D). Scientific survey, installation and subsequent repair would generate different evidence, rights, obligations, environmental disturbances and value flows. Data gathered by B may guide C, but a map with uncertain seabed registration cannot be silently promoted to a sufficiently precise connector-alignment instruction. D’s participation or applicable authorization may be required for particular activities without constituting a blanket delegation to the autonomous fleet. Financially shared benefits do not erase site-specific ecological consequences or responsibility for recovery and remediation.

Persistent infrastructure also changes the temporal horizon of robotic coordination. A one-time installation may enable many years of observation but create maintenance obligations, cable or battery waste, physical habitat alteration and constraints on subsequent users of the same marine space. Fleet efficiency, expressed as fewer ship days or reduced energy per completed service task, should be assessed together with material and ecological effects and with the institutional capacity to respond when equipment fails. A team that appears robust because it has numerous vehicles can be fragile when all depend on one recovery vessel, one shore-control channel, or one spare connector. Redundancy must be evaluated by failure domain, not only by vehicle count.

20.7 Communication cuts and the limits of distributed decisions

The interdependence of observation, authority and action can be given a limited formal expression. Consider two subsystems A and B that hold private binary observations $a, b \in \{0, 1\}$. Suppose a joint operation is authorized by a fixed common rule only when both predicates hold, and both subsystems must make a matching initiation decision $d_A = d_B = a \wedge b$. Assume they have no communication or common source of the other’s observation after receiving their private inputs, no additional shared information that determines those inputs, and deterministic decision rules. These assumptions are purposely narrow: actual marine operations may provide prior coordination, relays, conservative defaults or probabilistic protocols.

Proposition 20.1 (Local indistinguishability across a communication cut). *Under the preceding assumptions, there are no deterministic local rules $d_A = f(a)$ and $d_B = g(b)$ that correctly implement $d_A = d_B = a \wedge b$ for all four pairs $(a, b) \in \{0, 1\}^2$.*

Proof. Consider agent A when $a = 1$. In the joint state $(a, b) = (1, 0)$, the required decision is $d_A = 0$. In the state $(1, 1)$, the required decision is $d_A = 1$. Its local input is identical in both states, so a deterministic rule $f(1)$ cannot equal both values. Thus no such pair of local rules implements the joint rule for every state. $\square$

The proposition identifies a missing-information barrier, not an intrinsic defect of decentralized coordination. Waiting for acknowledgment, obtaining the missing predicate through a functioning relay, or assigning one subsystem the authority to decide from a subsequently reconciled record can alter the assumptions. So can restricting the feasible states through an independently validated process. Importantly, a default refusal to initiate may prevent unsafe action but cannot satisfy the specified requirement to initiate in the $(1, 1)$ state; the operator must distinguish *safety-conservative availability* from *complete task execution*. Randomization cannot guarantee exact correct initiation in every state without additional information, although it may change probabilistic performance criteria. A communication cut therefore deserves classification according to the task's information and authorization dependencies, not merely by the number of lost messages.

A practical implementation would test joint-task contracts under dropped packets, divergent clocks, faulty relays and inconsistent sensor beliefs. Each participating agent would record the rule version, evidence supporting its local predicate, messages acknowledged, and the eventual task result. For operations involving physical contact, the absence of a verified coordinated state should normally invoke a preapproved, task-specific cautious or abort mode, provided that this transition is itself safe in the physical environment. A moving manipulator cannot necessarily freeze without inducing additional load; a vehicle near a seabed habitat may require active separation to avoid impact. These transitions must be modeled and experimentally tested. The theorem alone certifies neither the safety of stopping nor the adequacy of the chosen fallback controller.

Field coordination and the limits of extrapolation. The ocean-front study of McCammon and colleagues provides a distinctive field benchmark: four gliders and two robotic surface samplers were coordinated during an approximately two-week Gulf of Mexico mission using task-specific autonomous planning and a human decision-support interface [106]. The observed system demonstrates coordination in a real oceanographic sampling context. It is not evidence of entirely unattended institutional decision-making, heterogeneous robotic construction, or reliable multi-month communication across every ocean basin. Team studies should therefore report what decisions were made onboard, what information was exchanged, what operators approved or corrected, and how performance changed when platforms or links became unavailable. Vehicle count, task duration, target dynamics, field conditions, and scientific ground truth define the proper comparison class. Multi-platform information gain remains a conditional design objective and cannot substitute for ecological inference or authorization to intervene.

20.8 System-level validation, accountability, and engineering continuations

A research programme for heterogeneous teams should proceed from component and interface verification to joint tasks and then to operationally relevant environmental trials. Controlled laboratory experiments can test common timing references, acoustic messaging, actuator handshakes, coalition scheduling and collision prevention. Hardware-in-the-loop evaluation can introduce realistic navigation drift, stochastic delays, current fields, energy degradation and lost communication. Field trials should disclose actual operating conditions, operator interventions and failures. Results from a closed-water or controlled tank trial should not be generalized to deep-ocean construction, and a successful observation-fleet deployment should not be used as evidence for autonomous infrastructure repair.

The evaluation should include at least four mutually distinguishable families of measures. *Operational performance* concerns task completion, duration, energy, successful recovery, downtime and intervention precision. *Collective robustness* concerns degraded performance under asset loss, loss of synchronization, shared-power faults and communications partition; it should compare a heterogeneous team against appropriate single-platform and homogeneous-fleet baselines. *Epistemic validity* concerns calibration, coverage biases, uncertainty propagation, reconstruction quality, provenance completeness and human review of discrepant evidence. *Ecological and institutional accountability* concerns disturbance metrics, documented avoidance and stopping responses, approved task scope, participation and the feasibility of correction or remediation. Reporting one aggregate score would obscure the fact that

high throughput can coexist with poor ecological inference or unacceptable intervention risk.

Experimental design must also consider task allocation’s distributional consequences. A fleet may concentrate high-value samples in accessible regions while systematically undersampling fragile, remote or difficult-to-observe habitats. Shared robotics may expand global access to ecological knowledge while remaining controlled by a narrow set of institutions able to fund and maintain the instruments. Decisions about where robots operate, which data are retained and whose ecological concerns lead to a pause consequently link team engineering to common knowledge production and institutional authority. The appropriate response is to make these dependencies visible and contestable through research dossiers, task histories and revision procedures, rather than presume that a more efficient algorithm resolves their normative content.

Two research continuations follow. Section 21 will investigate adaptive scientific observation and digital modeling: teams can be treated as distributed experimental instruments whose policies are revised through observation, model criticism and provenance-preserving updates. Sections 22–25 will evaluate animal welfare, ecological accountability, prospective technology and comparative validation, with special attention to physically finite environments and institutions that can review changing operational scope. The engineering claim established here is deliberately bounded: heterogeneous marine systems can produce capabilities unavailable to isolated platforms, as documented by coordinated observation and serviceable observatory infrastructure, but their collective effectiveness depends on communication, differentiated competence, material reliability, legal authority and ecological conditions that cannot be reduced to fleet size or optimization alone.

21 Adaptive Scientific Observation and Digital Modeling

21.1 Scientific observation as a changing relationship between models and interventions

The preceding engineering survey establishes that oceanic robots are not merely mobile carriers for interchangeable instruments. Their bodies, sensors, navigation estimates, communication links, maintenance dependencies, and operational permissions determine where measurements can be made and what those measurements can establish. The purpose of this section is to investigate a further capability: the capacity to revise the location, timing, and modality of observations in response to accumulating evidence. *Adaptive scientific observation* means that measurement decisions depend on a stated scientific question, observations and models available at the decision point, and bounded practical and normative conditions. It encompasses changing a profiling route to investigate a developing front, requesting a chemically specific sample after a sensor anomaly, and reallocating a heterogeneous team when an oceanic feature moves. Adaptation can improve a research campaign’s informativeness, but also introduces a dependence between where a system looks and what it subsequently reports.

This dependence is central to the generative-relational (GR) inquiry developed in Sections 2–6. An oceanic observation contributes to a situated representation of one or more evolving relational systems; the choice of observer action changes the accessible evidence and can sometimes change the material conditions being observed. An instrument’s successful detection of an oxygen minimum, for instance, establishes a calibrated observation at a place, depth, time, and resolution. It does not uniquely reconstruct microbial mechanisms, animal experiences, or the later continuation of the ecological system. Conversely, a numerical representation that reproduces temperature records may still omit organism behavior or locally important anthropogenic disturbances. Adaptation should expand the capacity to ask and test questions while retaining distinctions among measurement, inference, intervention, and warranted ecological claims.

Consider an observing campaign with physical state X_t , a robotic team configuration R_t , observation history H_t , an estimated model state M_t , and an approved set of activities $\mathcal{A}_t$. The decision process can

Table 18: Distinct interfaces in robotically mediated ecological inference. The stated checks are research-design requirements, not claims of completed validation.

Interface	Record and scientific question	Principal unsupported inference to prevent
Platform and setting	Realized action, navigation error, instrument state, relevant physical context	Requested action equals realized, ecologically neutral encounter
Signal formation	Calibration, detection model, noise, sampling support	Recorded signal uniquely identifies ecological cause
Ecological dynamics	Competing mechanism models, controls, responses under interventions	An observation model validates an organism response model
Adaptive selection	Candidate domain, access restrictions, action rationale, selection probabilities where applicable	Policy-selected data represent all eligible locations
Normative and institutional review	Approved scope, exposure assumptions, responsible authority, review triggers	Technical feasibility or information gain grants permission

be written schematically as

$$\begin{aligned}
 a_t &= \pi_t(H_t, M_t, \mathcal{A}_t), \\
 y_{t+1} &= \mathcal{O}(X_{t+1}, a_t, \eta_{t+1}), \\
 M_{t+1} &= \mathcal{U}(M_t, H_t, a_t, y_{t+1}),
 \end{aligned} \tag{78}$$

where a_t includes route, sensor, timing, or sampling choices; $\mathcal{O}$ is a modality-specific observation mechanism; η denotes measurement disturbances; and $\mathcal{U}$ is a model-revision operator. Equation (78) is a research schema, not a complete ecological dynamical model. It notably keeps the action that generated an observation within the recorded history. Removing that action would invite the mistaken interpretation of policy-selected samples as representative of locations never visited. A defensible adaptive programme accordingly records both *what was measured* and *why that location and measurement were chosen*.

One useful consistency check is to distinguish physical state evolution from the observation policy itself. In a more explicit, still schematic representation, the system can be described by

$$X_{t+1} \sim K_\theta(dx \mid X_t, a_t, c_t), \quad Y_{t+1} \sim L_\phi(dy \mid X_{t+1}, a_t, c_t), \tag{79}$$

where K_θ is a hypothesized physical and ecological transition kernel, L_ϕ is a measurement kernel, and c_t records relevant environmental context. The selected action a_t appears in both because a robot can affect the processes being measured as well as the conditions of sensing. The two kernels express different kinds of assumptions: a validated observation likelihood does not imply a validated account of how an intervention changes biological processes. In a strictly passive observational idealization, the action dependence in K_θ might be negligible for the studied quantity, but the negligibility itself is an empirical claim. Similarly, the specification of c_t cannot be assumed to cover unknown animal responses or unmeasured physical variables.

This distinction provides a precise interpretation of three common scientific failures. In an *observation failure*, a change in calibration or detectability produces a changed measurement even where the ecological process is similar. In an *intervention failure*, the platform changes a local process while the analysis treats the resulting record as an undisturbed baseline. In a *selection failure*, an otherwise reliable observation mechanism is applied at policy-preferred sites and the observed distribution is mistaken for the distribution over the declared target domain. These mechanisms can coexist. Each calls for a different study design, and none is resolved simply by increasing a robot’s collection rate.

The division is compatible with the earlier three-register discussion: a recorded signal is a manifestation, a modeled relation between action and ecological response is an explanatory hypothesis, and its referents and possible mechanisms require independent scientific scrutiny. The table supplies a more concrete bridge, not a claim that the physical dynamics, observation kernel, and legal regime form one empirically estimated global model. The statistical implications of policy-dependent selection are developed below; Section 25 considers how to design comparisons capable of separating them.

Table 19: Selected approaches to adaptive marine inquiry, with the scope of their evidence. The final column identifies unresolved questions rather than claims of demonstrated failure.

Approach	Scientific target	Documented contribution	contribution	Inference and deployment boundary
Informative route planning	Ocean field estimates	Constrained, model-based route choice [225]		Performance is conditional on model and objective; no general ecological optimum
Targeted AUV sampling	Moving fronts and episodic features	Feature-responsive coordinated sampling [117, 226]		Field scope and calibration limits
In situ bioanalysis	Water-borne organisms, molecular markers and toxins	Automated sampling and selected analyses [115]		Markers require biological interpretation
Heterogeneous front tracking	Dynamic front geometry	Six-robot field coordination [106]		Human interface; field-specific scope
Digital ocean modeling	Multiscale ocean representations	Shared modeling and scenarios [38, 40]		Task-dependent validation and uncertainty

21.2 Operational evidence for targeted and adaptive ocean sampling

The technical literature already contains credible field and laboratory foundations. Yilmaz and colleagues formulated adaptive autonomous-vehicle path planning as a constrained mixed-integer optimization problem, with uncertainty reduction in oceanic field estimates as a planning objective and explicit constraints for mobility, communication and coordination [225]. This is an instructive mathematical approach: a route can be informationally valuable under a model without being geometrically shortest, and model value must be balanced against an actual vehicle's feasible trajectory. Nevertheless, the optimality claim belongs to the specified objective and constraints, rather than to a universal ecological notion of the best observation.

At the level of operational research, MBARI describes onboard AUV and surface-vehicle intelligence for detecting and tracking oceanographic features and for responding with targeted observations. Its work includes collaborative sampling, autonomous acquisition of water samples, and adaptation to features that move or appear episodically [117, 115]. Zhang and colleagues reported coordinated and collaborative sampling by two long-range autonomous underwater vehicles, an informative demonstration of task-specific cooperation rather than proof of general marine scientific autonomy [226]. McCammon and colleagues' Gulf of Mexico campaign, considered in Section 20, likewise gives field evidence of heterogeneous teams following ocean fronts with human-supervised planning [106]. These studies should be compared by target phenomenon, real deployment conditions, onboard decision authority, required communications, and independent scientific validation.

Adaptive sampling also extends beyond vehicle trajectories. MBARI's Environmental Sample Processor (ESP) can collect and process marine water samples in situ, with some configurations providing molecular detection and others preserving samples for subsequent laboratory analysis [115]. It thereby alters the delay between biological sampling and information that can influence future sampling decisions. However, observing a genetic marker or molecular signal is not identical to proving the local abundance, welfare, or current activity of the organism associated with that marker. Similarly, an AUV that detects a plume with a Gaussian proxy model can use real-time observations to revise its route, but the resulting operational success remains contingent on model structure, sampling conditions, and the plume's dynamics [18].

A practical survey benefits from separating the scientific objects and demonstrated scope of these approaches:

Several maturity boundaries follow. A shore-updated sampling route is not equivalent to an onboard robot autonomously revising ecological hypotheses. A controlled demonstration of plume pursuit does not validate long-term biodiversity estimates. A model-derived forecast can help choose samples while

retaining substantial errors for unobserved scales or changing dynamics. These distinctions also apply in reverse: a highly capable ocean model cannot substitute for the physical ability to reach a fragile habitat, measure a relevant variable without undue disturbance, or retrieve material needed for independent verification.

21.3 Ocean state estimation, data assimilation, and digital modeling infrastructures

Adaptive field inquiry is connected to a larger infrastructure of ocean state estimation. The Global Ocean Observing System (GOOS) coordinates requirements and sustained measurement for physical, biogeochemical, and biological Essential Ocean Variables; such variables are chosen with attention to the scientific importance and practical feasibility of observation [54, 105]. Argo profiling networks, fixed observatories, ship surveys, satellite observations, robotic vehicles, and laboratory analyses contribute observations at different resolutions and levels of maturity. Their records are not automatically interchangeable. Calibration, location, temporal support, sensor response, quality control, and data-processing decisions remain part of their scientific meaning.

Numerical ocean analysis may combine observations and dynamical models through data assimilation. In simplified form, let x be a model state and y the measurement vector, with the observation operator $\mathcal{H}$ and observation-error covariance R . An ensemble or variational procedure can use discrepancies $y - \mathcal{H}(x)$ to update an estimate while accounting for background and observational uncertainty. For an illustrative linear Gaussian case,

$$x^a = x^f + K(y - Hx^f), \quad K = P^f H^\top (HP^f H^\top + R)^{-1}, \quad (80)$$

where superscripts f and a indicate forecast and analysis, respectively, and P^f is the forecast covariance. The matrix expression assumes compatible finite-dimensional quantities and an invertible innovation covariance. It illustrates estimation rather than a new algorithm. Oceanic physical processes are nonlinear, model errors are not generally Gaussian, and biological variables may lack a justified observation operator. Copernicus Marine differentiates model simulations, analyses, forecasts and reanalyses, including products that assimilate in situ and remote observations [27, 42]. A published analysis is thus a model-and-data-derived estimate, neither a complete field measurement nor an unqualified statement about ecological reality.

The term *digital twin* is useful when employed with comparable precision. The European Digital Twin Ocean initiative describes interoperable infrastructures that integrate observational data, models and interactive applications; the EDITO platform builds upon Copernicus Marine Service and EMODnet contributions [38, 40]. This is an actual public infrastructure initiative, not evidence that the ocean has already been captured by a single continuously validated model. Its stated applications range from simulations of physical processes to scenario studies involving economic and ecological decisions. A scientific digital twin must therefore declare its geographical and disciplinary scope, input datasets, temporal update rules, uncertainty representation, computational dependencies, model lineage, and appropriate decision uses. Multiple competing models may be more epistemically responsible than a single authoritative-looking rendering, especially when material and biological mechanisms are poorly observed.

A provisional model registry for an oceanic research infrastructure may store

$$\mathfrak{M}_k = (\mathcal{X}_k, \mathcal{O}_k, \mathcal{F}_k, \Theta_k, \mathcal{D}_k, \mathcal{V}_k, \mathcal{P}_k), \quad (81)$$

with represented variables $\mathcal{X}_k$, observation operators $\mathcal{O}_k$, dynamical or statistical models $\mathcal{F}_k$, parameters and uncertainty Θ_k , source-data histories $\mathcal{D}_k$, validation record $\mathcal{V}_k$, and provenance, permission, and revision information $\mathcal{P}_k$. A representation of an ecological community may have a different spatial and temporal domain from a hydrodynamic model of the surrounding water. Translating between them requires a specified interface, scale transformation, and uncertainty characterization. A visualization overlay is insufficient evidence that the models are causally or semantically integrated.

21.4 Scientific objectives, constrained experiment design, and information gain

Adaptive sampling is often described as maximizing information. That phrase conceals several choices. One can minimize forecast error for a physical field, reduce uncertainty about a biological event, detect an extreme condition early, distinguish rival dynamical hypotheses, or increase the coverage of a long-term reference network. These are different experimental objectives. Shannon entropy measures uncertainty under a specified probability distribution; realized surprisal refers to the improbability of the encountered measurement relative to a prediction; information gain measures a change in a conditional distribution or another explicitly stated inferential criterion. Neither a surprising sample nor a locally high-variance region is necessarily the most ecologically significant or the most valuable for every scientific question.

Let Z denote a scientific estimand or hypothesis, H_t the policy-aware observation history, and Y_a a potential measurement under action a . A standard expected information criterion is the conditional mutual information

$$\text{EIG}(a \mid H_t) = I(Z; Y_a \mid H_t) = \mathbb{E}_{Y_a \mid H_t} [D_{\text{KL}}(p(Z \mid H_t, Y_a) \parallel p(Z \mid H_t))] . \quad (82)$$

This quantity is nonnegative within the stated probabilistic model when the conditional distributions exist. The notation denotes conditional information given a realized history H_t , using the corresponding posterior and observation kernel; it is not an unconditional population average over every possible history. It depends on the model's prior, the measurement likelihood and the chosen estimand; it cannot certify that the hypotheses are complete or ecologically appropriate. Moreover, when the action physically changes an ecological quantity, a fixed pre-intervention Z and an action-indexed outcome Z_a are different estimands. The displayed information gain should then be interpreted for an explicitly chosen potential observation and estimand under a stated intervention model, rather than silently treating all actions as passive measurements of an unchanged Z . Gaussian-process sensor-placement research offers algorithms and conditional near-optimality guarantees for particular mutual-information objectives, but those results must not be generalized to nonlinear ocean ecosystems or non-submodular objectives [87]. Adaptive marine path planning has also used explicit predicted uncertainty fields [225]. A high information score is therefore a model-relative planning signal and an object of empirical validation.

A more realistic research design treats information alongside material and institutional constraints:

$$\begin{aligned} \max_{\pi \in \Pi_{\text{adm}}} \quad & \mathbb{E}_{\pi} \left[\sum_{t=0}^{T-1} g(Z, Y_{a_t}, H_t) \right] , \\ \text{subject to} \quad & a_t \in \mathcal{A}_t^{\text{phys}} \cap \mathcal{A}_t^{\text{authorized}} \cap \mathcal{A}_t^{\text{ecological}} , \end{aligned} \quad (83)$$

where Π_{adm} contains policies satisfying task-specific navigation, communication, energy and recovery constraints. The term g denotes a disclosed scientific gain measure, not necessarily mutual information. Ecological conditions should include applicable hard stopping criteria and review triggers rather than be represented solely as penalties exchangeable against more data. This follows the procedural framework of Section 10 and the locally bounded control discussion in Section 14.

Scientific evaluation must avoid choosing a policy solely because it improves its own model's expected score. A route designed to find a predicted feature may repeatedly rediscover that class of feature while leaving competing explanations untested. One possible approach combines targeted samples with fixed reference transects, occasional preapproved exploratory locations, and independently collected validation data. Randomized or stratified components can preserve detectability of systematic omissions. Multiple policies should be compared under equal energy, disturbance, communication and operational budgets. A counterfactual comparison also requires care: one cannot ordinarily observe the alternative sample that an unused route would have collected in the same evolving oceanic conditions.

21.5 Preferential observation, missingness, and the limits of model self-confirmation

The central inferential hazard of adaptation is *preferential sampling*: selection of measurement locations is statistically associated with the field under investigation. In geostatistics, this can invalidate analyses

that silently treat sampled locations as independent of the underlying process [32]. For example, suppose a mission prioritizes locations where its predictive model indicates high chlorophyll. The mean of its collected samples could then overrepresent bloom-associated waters relative to a specified spatial reference population. This is a selection effect even if each chlorophyll measurement is perfectly calibrated. In a dynamic sea, navigation and feature-tracking errors can further alter which local water masses were actually measured.

The selection mechanism must therefore be included in the research dossier. A record should distinguish the candidate locations available at each decision, the policy version, predicted score, estimated uncertainty, platform feasibility, applicable avoidance restrictions, actual selection, sensor observation, and any operator overrides. When sampling is randomized, the probability assigned to each feasible unit can be scientifically important; when sampling is deterministic, some units may have zero probability of inclusion. Such a dataset can remain extremely useful for local process study or event detection, even when it cannot support an unbiased mean for an entire target domain. The appropriate estimand is a substantive scientific decision and cannot be inferred after the fact from the path alone.

Three limitations deserve explicit recognition. First, adaptive policies may selectively visit *observable* features rather than features of greatest ecological significance: a highly visible physical front can dominate an objective while inconspicuous biological processes are omitted. Second, apparent improvement after an adaptive campaign can result from leakage between policy development, model fitting, and evaluation; validation data and reference models should be separated before policy comparison. Third, observation can itself change the target through disturbance from light, sound, sampling, or physical contact. If sampling actions materially modify the ecological state, the inferential target becomes action-dependent, and the simplistic picture of passively revealing a fixed hidden field fails. The welfare and experimental issues of Sections 4 and 15 therefore remain active constraints on adaptive observation.

A further problem arises when an aggregate digital representation is treated as if it were a neutral reconciliation of different institutions' evidence. Data-rich regions, well-funded sensors, accessible maritime jurisdictions, and commercially valuable activities may be overrepresented. Disagreements about habitats, lived environmental knowledge, or the impact of installations cannot always be resolved by assimilating every contribution into one numerical state vector. The relevant common infrastructure should preserve scientific disagreements, transformations, omitted evidence, and the conditions under which each interpretation may be challenged. Shared representation may facilitate inquiry without becoming a globally privileged account of every marine relationship.

21.6 Relational model revision, histories, and plural digital representations

A generative-relational approach treats the production and revision of scientific representations as events within the inquiry system. Robotic sampling changes the available evidence; analysis changes which continuations are considered plausible; institutional decisions change which subsequent observations can take place. The practical consequence is that a digital model should include a transparent revision history. A model version may be re-estimated with a corrected sensor offset, a new acoustic calibration, or a changed hydrodynamic boundary condition. These are not interchangeable operations, because they alter different claims about evidence and mechanism. A responsible record can state the relevant transformations and preserve the ability to reconstruct prior results under their original assumptions.

For models M^A and M^B produced by heterogeneous scientific communities or instruments, translation should be identified by a partially defined mapping $\tau_{A \rightarrow B}$ together with the domain in which it preserves task-relevant relations. A map from a high-resolution microbial assay to a coarser regional ecosystem indicator, for example, requires explicit aggregation assumptions and may discard species-level information. Likewise, transferring a flow-field analysis to a vehicle controller demands bounded interpolation error and validity under changed current conditions. There is no need to posit one complete "ocean model" or one ocean-wide GR subject from which every scientific object can be deduced. An interoperable infrastructure can support several partially aligned representations and keep their differences intelligible.

Model revision also concerns relations conventionally classified as historical, current, or prospective. A hindcast, a near-real-time analysis, and a forecast are distinguished here by data conditioning, operators, and intended inference, rather than by treating past and future as trivial, ontologically separate containers. A revised analysis may reclassify earlier samples because a calibration change modifies their interpretation; a prospective simulation may motivate sampling that changes the conditions of subsequent observation. This recursive relation is scientifically productive provided the representation retains explicit provenance and makes its counterfactual or predictive claims testable.

For the A–B–C–D transnational marine-facility example, consider a consortium model predicting local heat dispersion and robot access conditions. Monitoring robots measure thermal and ecological indicators, while other models describe installation operations, habitat sensitivities and regulatory commitments. A revision to the thermal model may justify changed monitoring frequency or a proposed shutdown threshold, but does not by itself revise a legal permit or establish that a disturbed organism has suffered a particular welfare impact. The institution responsible for review must be able to inspect input records, model version, validation coverage, relevant uncertainty and proposed operational consequences. In the same way, commercial ownership of computational facilities does not entail proprietary authority over every ecological observation or the interpretive rights of other affected actors. Scientific circulation, ownership, access and remedy remain differentiated, as developed in Sections 7–9.

21.7 A conditional unbiasedness result for adaptive sampling with recorded selection probabilities

A narrow finite-domain proposition makes the selection issue precise without asserting that adaptive observation is generally unbiased. Fix a finite target domain $\mathcal{S} = \{1, \dots, N\}$ with unknown but fixed values $z_1, \dots, z_N$ during the considered sampling period. The target is $\mu = N^{-1} \sum_{i=1}^N z_i$. At each time t , a policy chooses an index I_t with selection probabilities $p_t(i \mid H_{t-1})$ recorded before the draw. The history H_{t-1} may contain all previous indices and observations, and hence the policy may be fully adaptive. Assume each $p_t(i \mid H_{t-1}) > 0$ for all target units, the probability distribution is normalized and correctly recorded, the observed value equals z_{I_t} , and the variables are integrable. These strong assumptions deliberately exclude measurement error, a changing underlying field, forbidden locations and uncertain position registration. In particular, strict positivity across the entire target set is incompatible with any location assigned selection probability zero by navigation, legal, or ecological restrictions. The finite-population claim must then be narrowed to the *predeclared accessible population*, or supplemented with independent identification assumptions for inaccessible units; recording an impossible unit as if it had positive inclusion probability does not repair the inference.

Proposition 21.1 (Conditional design-unbiased estimation under adaptive selection). *Under the preceding assumptions, the estimator*

$$\hat{\mu}_T = \frac{1}{TN} \sum_{t=1}^T \frac{z_{I_t}}{p_t(I_t \mid H_{t-1})} \quad (84)$$

satisfies $\mathbb{E}[\hat{\mu}_T \mid z_1, \dots, z_N] = \mu$ for every positive integer T .

Proof. Conditional on H_{t-1} and on the fixed values $z_1, \dots, z_N$, the probability of selecting i is $p_t(i \mid H_{t-1})$. Hence

$$\mathbb{E} \left[\frac{z_{I_t}}{N p_t(I_t \mid H_{t-1})} \mid H_{t-1}, z_1, \dots, z_N \right] = \frac{1}{N} \sum_{i=1}^N z_i = \mu. \quad (85)$$

Taking expectations and averaging the T summands establishes the claim by the tower property of conditional expectation. The proof does not require successive sampling decisions to be independent. $\square$

The proof uses conditional expectation rather than independent draws. This distinction is helpful, but unbiasedness alone is insufficient for reliable inference: extremely small inclusion probabilities can produce an estimator with severe dispersion. Indeed, for fixed $z_1, \dots, z_N$, let $X_t = z_{I_t} / [N p_t(I_t \mid$

$H_{t-1})]$. The centered increments $X_t - \mu$ form a martingale-difference sequence under the proposition's assumptions. When their second moments are finite,

$$\text{Var}(\hat{\mu}_T \mid \mathbf{z}) = \frac{1}{T^2} \sum_{t=1}^T \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \frac{z_i^2}{p_t(i \mid H_{t-1})} - \mu^2 \mid \mathbf{z} \right]. \quad (86)$$

If a predeclared uniform lower bound $p_t(i \mid H_{t-1}) \geq p_{\min} > 0$ holds for all feasible target units and all realized histories, the variance is at most $(TN^2 p_{\min})^{-1} \sum_i z_i^2$. Such a bound can be very loose and says nothing about bias from changing field values, inaccurate measurement, model-based imputation, or robot-induced ecological alterations. In particular, it supplies no permission to visit locations whose inclusion is ecologically or legally prohibited.

The proposition is a design-based elementary extension of inverse-probability reasoning to history-dependent selection; it should not be construed as a novel universal estimator. It has important boundaries. Very small probabilities produce large weights and potentially severe variance. A deterministic adaptive route gives zero probability to unsampled units and violates full support. If a survey's available locations change because of safety or legal constraints, the estimand must be restricted or inference must be redesigned; it cannot be repaired by pretending forbidden units remained observable. If z_i changes as the mission proceeds, observations at different times no longer estimate the same fixed-domain mean without an explicit temporal model. Measurement error, navigation uncertainty, missed detections and destructive sampling require further assumptions and corrections. Even unbiased estimation of a variable's spatial mean does not validate an ecological explanation or an animal-welfare judgment.

The practical value of the result lies in its requirement for *policy provenance*. The policy, its candidate set, and its time-specific selection probabilities are part of the scientific record whenever design-based statements are intended. In cases where full-support randomized exploration is ecologically unacceptable, the responsible choice may be to narrow the inferential question and disclose the inaccessible domain. Institutional authorization and ecological protection take precedence over satisfying a statistical theorem's premises. The relationship between sampling capability and justified inference is thus conditional and revisable, not guaranteed by information maximization.

Experimental status of adaptive-planning claims. Yilmaz and colleagues formulate model-based adaptive sampling as a mixed-integer optimization problem maximizing a specified uncertainty-related objective under vehicle and network constraints [225]. This formal contribution establishes a tractable way to state a planning problem, but the maximized objective is not identical to ecological truth or a demonstration of increased observational validity at every site. Berget and colleagues examine a Gaussian-proxy approach for an AUV observing dispersal processes, with a particular spatial representation and proposed information-seeking routes [18]. The distinction between a validated planning formulation, a domain-specific experimental demonstration and a generally transportable ecological estimator is central. Comparative evaluation should predefine target phenomena, independent reference observations, model-mismatch diagnostics, observer effects, and the probability of sampling each eligible unit. Adaptive sampling that preferentially revisits predicted hotspots may raise informativeness for one modeled objective while reducing coverage of uncommon or poorly modeled processes. Both achievements and exclusions belong in the scientific provenance record.

Animal-mediated missingness and the adaptive-policy record. One mode of preferential sampling arises when organisms react to the scientific platform before detection. Acoustic observations ahead of robotic systems have documented both attraction and avoidance, varying with depth, platform and illumination [17]. Consequently, a robot that updates its search policy using only animals that remained in view may reinforce a systematically incomplete picture. Recorded route and light-switching probabilities are helpful, but they do not by themselves identify the biological inclusion probabilities for animals that left before the sensor could detect them. The statistical proposition in this section therefore applies only to the defined measured inclusion mechanism, for which positivity and the conditioning assumptions are

Table 20: Selected engineering sources and the conclusions their evidence can warrant. The examples are illustrative rather than an exhaustive verification of the engineering bibliography.

Study or system	Directly documented evidence	Additional evidence needed for broader claims
Soft robot at extreme depth [93]	Fin actuation near 10,900 m; separate free swimming near 3,224 m	Repeated mission autonomy, sampling validity, ecological exposure
Sea-Thru [7]	Image correction with a specific optical model and range-supported RGBD inputs	Radiometric and ecological ground truth across sites and taxa
JANUS [179]	Acoustic interoperable signalling standard and documented demonstrations	Mission-specific reliability, latency, semantic interoperability
Dual-I-AUV assembly [100]	Prepared cooperative assembly in a water tank	Current-exposed field operations, unaided installation, structural qualification
Ocean-front robotic team [106]	Two-week field campaign with four gliders, two surface samplers and human oversight	Transport to other fronts and fleets, ecological inferential accuracy
Adaptive route optimization [225]	A specified constrained planning formulation and illustrative evaluations	Independently measured estimator performance and intervention effects

justified. It does not repair unmeasured animal responses. A field protocol should specify the desired abundance or behavior estimand, distinguish policy-selection from organism-detection processes, maintain exposure and operating-mode logs, and use suitably independent instruments, matched transects, or sensitivity bounds to investigate residual nonignorable missingness. Where such calibration is unavailable, adaptive products should be reported as platform-conditional observations rather than unconditional ecological estimates.

21.8 Engineering evidence grades across the topical survey

The eleven engineering topics considered above are best compared using the actual observation or test underlying each report. Table 20 records selected source-specific examples. It does not rank scientific quality on a single scale: an analytical method and a field operation answer different questions. Neither the presence of a peer-reviewed article nor the technical sophistication of an implementation independently determines ecological validity.

For reproducibility, the proposed evidence register stores at least the bibliographic identity, engineering claim, platform or model version, setting, sample or mission size, comparison method, human participation, failure and uncertainty reports, and the exact supported inference. When data, detailed protocols, or independent replication are unavailable, that absence should be recorded rather than filled by extrapolation. This audit also provides a starting point for the comparative methods of Section 25, without treating a proposed GR extension as superior to existing scientifically competent baselines.

21.9 Sediment transport, moving exposure fields, and habitat-conditioned inference

Oceanic disturbance frequently involves a moving exposure field rather than a static circular footprint. For seabed contact operations, a monitoring system may need to distinguish the initial mechanical imprint, near-bottom suspended sediment, later redeposition, and any biological responses or altered detection conditions. Independent monitoring during a pre-prototype nodule-collector trial at approximately 4,500 m recorded a density current that moved downslope and suspended material dispersed with ambient bottom flow; elevated concentrations and redeposition varied markedly with location, altitude, and elapsed time [49]. This was a specific intensive collector trial, not a measurement of routine scientific vehicle disturbance. It nevertheless demonstrates that fixed-distance exclusion assumptions can be poor substitutes for measured local transport.

A prospective robotic field protocol would represent relevant exposure as a spatiotemporal field $C(\mathbf{r}, z, t)$ with stated units, sensor response and concentration-calibration uncertainty. It would compare observations across depth or altitude strata, gradients, current conditions, and both exposed and reference

locations. A sampled peak concentration, spatially integrated suspended mass, deposited thickness and biologically meaningful dose are different endpoints: none can be derived from another without additional transport, geometry and organism-exposure assumptions. Camera obscuration may also alter the probability of detecting animals during the same interval in which the ecological exposure changes. The resultant missingness is coupled to the intervention and cannot in general be repaired by knowing only a robot's planned station-selection probabilities.

Where a numerical plume model guides future sampling, out-of-sample concentration and deposition measurements should be used to test dispersion and settling parameters. A model that accurately reproduces a survey's optical backscatter need not reproduce toxicity, burial of filter feeders, or longer-run benthic recovery. Independently measured biological effects, rather than the output of the same model used to target the observations, must anchor those subsequent claims. No uniform conversion from the cited collector plume into a safe approach distance for unrelated devices is warranted.

21.10 Validation, institutional accountability, and transition to prospective research

Validation should separate at least four classes of success. *Technical success* concerns whether route replanning, sensors, navigation and communications function under disturbances, including delayed or lost messages. *Statistical success* concerns estimation error, uncertainty calibration, coverage, detection power and robustness to preferential sampling, benchmarked against fixed, stratified and alternative adaptive designs. *Scientific success* concerns whether the observations differentiate specific explanations or support replicable inferences about a defined ecological process. *Ecological and institutional success* concerns avoidable disturbances, traceable authorization, review of unexpected findings, data correction and practical remedy. None is interchangeable with a robot's number of samples collected or a digital twin's display resolution.

An appropriate staged research programme begins with synthetic fields for which target estimands are known, followed by controlled observations with independently calibrated instruments and then by field campaigns with explicit ecological monitoring. Synthetic benchmarks should include stationary and advecting fronts, intermittent plume events, sparsely observed taxa, sensor drift, navigation uncertainty, heteroscedastic noise, and misspecified model classes. For learning-enabled systems, training, tuning and validation datasets should be separated by mission or environmental regime as appropriate. A model-based adaptive policy should be compared with a reproducible nonadaptive policy under identical energy, field time, access and disturbance constraints. Any claim of superiority should report uncertainty and the conditions where gains disappear or reverse.

Persistent observation infrastructures raise additional design requirements. Reference transects and long-running fixed records may be scientifically valuable precisely because their measurement decisions are stable enough to support comparisons. An adaptive system that continually reallocates effort can increase local event detection while weakening long-term comparability. A robust network may therefore combine sustained baseline measurement, targeted campaigns, community-defined observation priorities and explicitly exploratory robotic sampling. The distribution of instruments and computation also has political-economic significance: actors with stronger resources may determine the questions encoded into models, which phenomena are rendered visible and who can contest conclusions. Open interfaces, audit trails and accessible validation data can improve common knowledge production, but responsible collaboration must also attend to legally protected information and the interests of ecological communities and people affected by deployments.

In this paper, adaptive scientific observation is therefore the final engineering link between embodied capability and ecological knowledge formation. It connects the physical technologies examined in Sections 11–20 to questions of evidence, representation, normativity and common productive capacity. The next section develops animal welfare and ecological accountability in the engineering design itself; Sections 23 and 24 investigate biologically informed future technologies and prospective infrastructure under material and governance constraints. Section 25 then specifies comparative research designs that could falsify the present program's assumptions. The resulting research direction is a system of inspectable

inquiries with partial, plural, and revisable representations, rather than a presumption that more sensors, more adaptive autonomy or one digital ocean model necessarily produces better ecological knowledge or justice.

22 Animal Welfare and Ecological Accountability in Engineering

22.1 The research object: engineering action as an ecological encounter

The engineering survey in Sections 11–21 examined the means by which robotic systems travel, perceive, manipulate, communicate, and revise their scientific operations. The present section investigates a constraint on the very use of these capacities: an oceanic robotic system can change the biological and material conditions under which its observations acquire meaning. The objective is to make this relationship available to engineering design, empirical evaluation, and institutional revision without equating successful observation with an ethically justified intervention. A camera that records an animal’s avoidance, a vehicle that elicits that avoidance, and a model trained on the resulting images form a coupled measurement encounter. Each element requires evaluation within its own evidential scope.

Animal welfare here concerns the condition and potentially valenced experiences of individual animals, including avoidable injury, distress, disruption, and opportunities for species-appropriate activity. *Ecological accountability* concerns the justifiability and revisability of operations in relation to organisms, habitats, ecological processes, affected people, and institutions. Neither notion can be inferred solely from the other. A campaign may produce negligible measurable population-level change yet seriously harm a small number of animals; an intervention that appears tolerable to an encountered animal may damage habitat or cause accumulative effects at a larger scale. Ecological assessment also extends to organisms for which current evidence of sentience is incomplete. The concern is therefore broader than an inventory of legally protected or charismatic taxa.

The epistemic difficulty is substantive. Fish nociception and pain have been investigated using behavioral and physiological evidence, while cephalopod and decapod sentience have generated a substantial but heterogeneous comparative literature [169, 29, 20]. Evidence across taxa is neither uniform nor reducible to a single behavioral response. An animal moving away from a vehicle might be responding to sound, illumination, water displacement, perceived predation, or some combination of these; the same outward behavior cannot uniquely establish subjective welfare. Conversely, an apparently motionless animal may be unable or unwilling to flee. The Five Domains model supplies a useful welfare-assessment vocabulary, but translating it to free-ranging, briefly encountered marine animals requires validation rather than simple transfer from managed-animal settings [107].

The generative-relational (GR) formulation developed earlier in this paper regards robotic operations as interventions within heterogeneous relational trajectories. This directs attention to what is changed by the encounter, who or what may be affected, and which inferences remain unobservable. It does not confer a privileged moral interpretation upon a robot’s internal representation. Responsible engineering must instead join mechanistic exposure models, species- and context-relevant evidence, institutional obligations, operational alternatives, and mechanisms for challenging conclusions after deployment. In this section, “responsible” describes an assessable research objective, not an assertion that any currently deployed platform has met universal ecological criteria.

Evidential levels for welfare-related inference. The Five Domains model is a structured approach to welfare assessment, with four domains concerning nutrition, physical environment, health, and behavioral interactions and a fifth concerning inferred mental state [107]. It does not transform any generic exposure indicator into an observed subjective experience. Its authors distinguish environmental risk conditions from animal-based indicators and note major knowledge gaps for some free-ranging taxa. For a robotic encounter, a measured sound field might establish a physical exposure, and an observed change of movement might support a context-specific behavioral finding. Moving from those records to claims concerning pain, distress, welfare compromise, or improvement requires additional species-appropriate

physiological, behavioral, or other justified evidence. Even then, organism-level findings and population-level ecological conclusions remain different claims. A protocol should therefore retain the evidence for each inferential transition, alternatives such as current changes or predator presence, and the uncertainty associated with missing or short-duration observations. Precautionary operating rules can be warranted under uncertainty without presenting speculative subjective-state diagnoses as scientific measurements.

22.2 Acoustic, optical, hydrodynamic, and chemical exposure pathways

Marine vehicles introduce several partly coupled exposure pathways. Propulsion and onboard machinery radiate sound; active sonar and acoustic communication add intentionally emitted signals; illumination modifies local spectral and temporal light fields; a body in motion displaces water and alters pressure gradients; sampling and manipulators can generate contact, sediment resuspension, chemical release, and biological-material disturbance. The spatial footprint and temporal persistence of these changes depend on water-column structure, bathymetry, vehicle geometry, propulsion, behavioral response, and existing background pressures. A single undifferentiated “disturbance score” obscures these mechanisms unless its empirical construction is made explicit.

Underwater sound is particularly consequential because acoustic information and communication support many marine animals. NOAA Fisheries identifies hearing effects, stress responses, habitat displacement, and disruption of feeding, reproduction, and communication among the potential consequences of human-generated ocean noise [130]. For marine mammals, Southall and colleagues propose auditory-weighted exposure criteria for temporary and permanent hearing changes, while emphasizing differences among hearing groups, source types, exposure metrics, and uncertainty [171]. These auditory-injury criteria cannot by themselves determine the onset of behavioral disturbance, much less a population consequence. Pirotta and colleagues show why links between disturbances, individual vital rates, and population responses require explicit intermediate assumptions and evidence [156]. Sound exposure modeling for a robotic mission should therefore record source spectra, temporal duty cycles, receiver contexts, propagation uncertainty, and cumulative exposure rather than merely the stated acoustic power of a transmitter.

The 2024 NOAA Fisheries hearing guidance updates acoustic criteria for the onset of auditory injury and temporary hearing-threshold shifts in specified marine mammal hearing groups [122]. Like the Southall recommendations, this is an endpoint-specific acoustical tool rather than a criterion for all behavioral responses or a generalized welfare score. Its receptor-based exposure and frequency weighting should not be confused with source power, and marine mammal auditory criteria should not be extended without validation to fishes or invertebrates. The guidance also does not independently authorize a particular mission. The International Maritime Organization’s revised ship-noise guidance addresses reduction of radiated noise from *shipping*; its revised circular and subsequent amendment form valuable engineering references but cannot be read as a robotic-animal-welfare certification standard [75, 71]. Regulatory treatment differs by activity and jurisdiction. The European Union’s Marine Strategy Framework Directive uses a descriptor concerning the introduction of energy, including underwater noise, into the marine environment. Its marine-status framework makes spatial and temporal distributions of sound an environmental policy object, without supplying a universal numerical limit for every marine robot [39]. A mission designer may use these materials to identify pressures and assessment methods while separately establishing the actual legal requirements that govern a proposed location and operation.

Optical exposure requires similar care. Artificial illumination can alter orientation, predator–prey interactions, developmental and reproductive processes, and temporal behavior in marine environments. A review of artificial light at night documents mechanisms and effects across multiple coastal and pelagic taxa and highlights context-specific uncertainties [104]. In naturally dark habitats, short-duration survey lights may be ecologically relevant despite being essential for reliable imagery. Reducing light intensity may protect some animals but decrease image quality and increase survey duration; using an alternative wavelength may alter detectability differently across taxa. In particular, experimental pelagic observations have recorded marked avoidance under white, blue *and red* instrument lights [51]. That result does

not establish equivalent responses across taxa, habitats or light levels, but it removes the basis for classifying red illumination as categorically non-disruptive. A defensible design compares integrated exposure and evidential usefulness under specified observation conditions rather than assuming a particular color is universally harmless.

Hydrodynamic effects are often neglected because the same wake and thrust used for maneuvering are treated as engineering disturbances. Close pursuit, rapid acceleration, propeller wash, near-bottom thruster operation, or manipulation around sessile organisms can alter local fluid motion and sediment exposure. Chemical and material pathways include leaks, surface coatings, corrosion products, cleaning agents, and lost fragments. The long-term biological changes reported after an experimental deep-sea mining disturbance illustrate the possibility of persistent habitat effects, although such an experiment is not an exposure-equivalent analogue of ordinary robotic survey operations [83]. Extrapolation from large-scale extraction to a small robot, or in the opposite direction, requires footprint- and process-specific justification.

The scope of acoustical criteria and dose comparisons. Underwater sound is heterogeneous in frequency, temporal structure, duration, direction, propagation and biological reception. NOAA Fisheries summarizes possible hearing, stress, displacement and behavioral pathways while explicitly conditioning their effects on source, duration and setting [130]. Southall and colleagues' updated criteria concern specified marine mammal auditory effects and hearing groups; those criteria should not be transformed into a universal bound for behavioral disturbance, fish welfare, invertebrates, or repeated multi-stressor exposure [171]. Comparisons should preserve the distinctions among source level, modeled transmission, *received* exposure at a specified receptor, and the biological endpoint for which an adopted criterion was developed. Acoustic metrics and light, hydrodynamic, thermal, and chemical exposures should remain a vector of physically defined quantities unless a justified, endpoint-specific response model supports combining them. The surrogate bound later in this section is mathematically valid under its own assumptions, but it does not validate the ecological meaning of scalarized cross-modal exposure.

22.3 Animal encounters, behavioral evidence, and observer-induced bias

A robotic observer can influence the occurrence, duration, and interpretation of animal encounters. Benoit-Bird and colleagues compared ROV and AUV surveys of midwater life in Monterey Bay using imaging and forward-looking acoustics; responses included attraction and several forms of avoidance, with effects differing by depth and vehicle type [17]. In that setting, animals reacted more strongly to the AUV than to the ROV, and illumination affected the observations. The measured response is evidence of observation bias and animal behavior within the studied campaign; it is not in itself a diagnosis of subjective distress or an estimate of demographic harm. Recent experimental work on repeated exposure to a robotic fish likewise reports avoidance dynamics in a specified experimental setting; those findings should not be projected automatically onto wild marine species, different morphologies, or long-term ecological outcomes [202]. A contemporaneous wildlife-response comparison of biomimetic and conventional vehicles is available as a preprint, but its claims should remain identified as preliminary until their methods and review status warrant stronger conclusions [154].

An animal's apparent reaction should be analyzed against explicit alternative interpretations. It may reflect response to a particular modality, unfamiliar object, prior experience, group structure, energetic state, predation risk, or environmental covariates. The absence of an observable response also has multiple interpretations: exposure below an animal's detection threshold, habituation, limited mobility, hidden stress, insufficient camera coverage, or a modality to which the relevant animal does not respond behaviorally. These possibilities make welfare inference both an experimental-design problem and an ethical problem. Passive remote measurement where practicable, independent reference observations, pre-registered encounter categories, and calibrated response windows can reduce but not eliminate inferential ambiguity.

An operational protocol should distinguish *incidental encounter*, *targeted observation*, and *physical*

interaction. In an incidental encounter, the robot can often reduce speed, avoid pursuit, and alter route without compromising a scientific question that did not require that animal. In targeted observation, the research protocol must specify the observational objective and justify its sensory footprint, repetition, and approach constraints. In an intentionally physical interaction, such as specimen collection or attachment of a device, a qualitatively different justification and authorization analysis is needed. The scientific benefit of a close image does not by itself authorize capture, touching, or experimental stimulation. In relevant United States jurisdictions, NOAA's marine life guidelines warn against close approaches and harassment of protected animals, and distinguish public viewing guidance from research permissions [129]. These distances and restrictions vary across taxa, locations, and activities; they are not universal marine-robot operating rules.

Welfare monitoring should combine, when defensible, animal trajectories, orientation, ventilation or other taxon-appropriate physiological indicators, feeding and social behavior, avoidance or attraction, and changes in habitat occupancy. However, the inferential value of each measure needs explicit baselines and calibration. A robotic sensor observing an animal only after it approached the platform yields selected evidence; animals that left before detection may be absent from the dataset. Repeated records of a seemingly undisturbed subset can therefore provide false confidence. The sampling-policy discussion in Section 21 applies equally to encounters selected by attraction, avoidance, or sensor visibility. Scientific claims should report who was observable, by which means, and under what vehicle behavior.

Separating evidence of displacement, harm, and sampling distortion. A useful distinction concerns three outcomes that can occur together but should be reported separately. *Encounter response* describes a change in observable behavior or physiology relative to an explicit comparator. *Sampling distortion* describes the difference between an ecological target quantity and the quantity recoverable through a particular detection process. *Welfare or demographic consequence* requires evidence relevant to an individual's condition, or a causal pathway connecting response to energetics, survival, reproduction, or population change. The ROV survey review documents substantial variation in illumination, speed, sampling design, and reporting practices; inconsistent recording of those features obstructs both inter-study comparability and the diagnosis of observation effects [177]. An animal can avoid a vehicle and bias an abundance estimate without demonstrated injury. Conversely, a physiological response or injury may occur without an observable departure from the camera field. Comparative protocols should treat the three outcomes as linked hypotheses and report where each remains indeterminate.

22.4 Habitat, ecological continuations, and cumulative exposure

Animal-welfare precautions must be coordinated with habitat-level considerations. Near-bottom vehicles and manipulators can contact biological structures, disturb sediments, introduce shading or heat, damage fragile benthos, or change the configuration of habitats. Moorings and persistent installations can create artificial substrate, alter flows, attract organisms, and introduce entanglement or maintenance pressures. Such effects may be beneficial for some species, adverse for others, and different over the life of the installation. Evaluating only the footprint visible during deployment may omit persistent material change and later recovery or removal activities.

Ecological continuations involve processes that exceed the temporal horizon of a mission: larval settlement, repeated disturbance, slow-growing habitat-forming organisms, seasonal movement, and coupled food webs. A disturbance that is small in one hour may recur at sensitive times or across many vehicles. For example, repeated vehicle approaches near a breeding aggregation could alter energetic expenditure without causing acute injury; a recurring inspection route could traverse the same benthic patch after each maintenance cycle. Conversely, brief changes observed during a survey cannot be presumed to have enduring consequences. The connection between short-term response and subsequent demographic effects is a research question requiring independent evidence [156].

The relevant unit of analysis is not simply the vehicle. Consider a scientific fleet supported by ships, charging and communication stations, moorings, onshore computation, and periodic material replace-

ment. Each component has a direct and indirect footprint. If the operator claims that replacing a ship survey with long-endurance robots reduces disturbance, the comparison must include the counterfactual ship operations displaced, as well as charging logistics, manufacturing, failure recovery, and extra sampling induced by lower per-mission cost. Ecological value cannot be read from energy efficiency or reduced staffing alone. Similarly, a system that accumulates high-resolution imagery could generate repeated disturbance while mainly confirming a previously identified phenomenon; additional data volume does not settle the marginal epistemic justification of the associated exposure.

A marine ecological dossier can therefore organize impacts by at least four temporal registers: pre-installation baseline; immediate operational exposure; delayed ecological response; and recovery, abandonment, or removal. These registers are practical observation windows rather than ontologically fixed divisions into past, present, and future. Their boundaries depend on ecological processes and monitoring capacities. The dossier should retain alternative plausible trajectories, changes in background conditions, the confidence attached to observed effects, and whether a project could meaningfully withdraw or alter its presence if new evidence emerged. Where restoration is difficult or uncertainty irreducible, preventing an unnecessary initial disturbance may be more defensible than treating a speculative remediation plan as sufficient compensation.

22.5 Habitat integrity and cumulative ecological response

A review of ecological encounter effects should account for habitat as well as organism-level response. Repeated mechanical contact can affect sediment structure, attachment surfaces, burrows, microbial gradients and recolonization pathways, even where a survey records no immediate avoidance by conspicuous fauna. The absence of an observed behavioral response from an individual is an inadequate surrogate for habitat integrity or subsequent population dynamics. Evidence from deep-sea disturbance experiments, including persistent differences in benthic megafauna and sediment microbial functioning after decades, supports this separation of endpoints [168, 205]. Those effects arise from intensive seafloor disturbance and require explicit exposure scaling before informing smaller research-system designs.

Cumulative disturbance must be specified across both space and time. Several visits to a station can overlap physically and biologically even when individual operations satisfy local engineering limits; independent projects can also act on one habitat through separate channels. Acoustic, optical, mechanical, chemical and thermal exposures need receptor- and habitat-specific descriptions. Simply summing incompatible quantities has no physical meaning. For example, a model may keep separate time-integrated light intensity, modeled received sound exposure, sediment deposition and contacted seabed area, and examine interactions only after a justified response model or empirically supported bounding rule has been established. An acoustic criterion for an auditory endpoint does not establish sediment tolerance or a welfare threshold.

Evidence of protected seamount coral regrowth over several decades shows that conditions for ecological continuation can persist amid visible historical damage [16]. It also cautions against equating the possibility of regrowth with rapid restoration. Habitat monitoring should therefore include recolonization, recruitment or attachment opportunity where scientifically feasible, retain control and reference sites, and allow outcome-specific monitoring horizons. Ethical and governance review should consider whether some interactions are insufficiently observable for an authorization decision at the proposed scale.

22.6 Engineering design alternatives and welfare-sensitive control

Animal-welfare concerns become actionable when compared against actual engineering alternatives. A mission may vary vehicle morphology, propulsor placement, cruise speed, acceleration profiles, acoustic communication duty, sensing modality, illumination, route, scheduling, manipulator compliance, and proximity. No single design is uniformly superior: a quiet propeller may be less efficient in a strong current, a low-light camera may require longer occupation, and an acoustic beacon may improve navigation enough to reduce the risk of unintended contact. Trade-offs should be declared as hypotheses to be

Table 21: Candidate pathways and accountable engineering responses. These are testable design considerations, not universal safe thresholds or validated species-independent effects.

Pathway	Possible exposure or effect	Design and operating alternatives	Required verification
Propulsion and sonar	Acoustic masking, injury risk under some exposures, altered behavior	Source characterization; route and duty revision; passive sensing where suitable	Received exposure and taxon-specific interpretation [130, 171]
Artificial light	Attraction, avoidance, temporal disturbance	Exposure scheduling; selective illumination; calibrated low-light imaging	Spectral irradiance, behavior and data quality [104]
Wake and close approach	Escape, resuspension, habitat contact	Clearance, low-thrust approach, speed and proximity governance	Current-relative motion, contact and sediment measures
Sampling and manipulation	Capture or tissue damage, local habitat alteration	Non-contact methods; compliant tools; task-scoped collection	Procedure-specific review and observed outcomes
Persistent hardware	Entanglement, material loss, heat, altered substrate	Recoverable fixtures, maintainability, material accounting	Lifecycle inspection and ecological follow-up
Data and autonomous decisions	Selective animal visibility, incorrect inference, unauthorized action	Encounter logs, scoped control policy, human review, model audit	Coverage, error, provenance and contestability

evaluated rather than settled by the language of “bio-inspired” or “soft” technology.

Design qualification should proceed through staged evidence. A bench assessment can characterize propulsor noise, lighting spectra, actuator force limits, materials, leakage paths, and failure modes. Controlled testing can then measure the physical field generated by a vehicle under realistic maneuvering. The next stage investigates species-relevant encounter conditions when scientifically necessary and ethically authorized; experimental exposure of animals should not be treated as a routine prerequisite when nonanimal evidence can answer the engineering question. Finally, field assessment must test whether instrumented measurements and ecological observations agree with the intended operating envelope. Each stage may falsify earlier design assumptions.

In the control architecture of Section 14, a welfare-sensitive supervisor should bound actions by the applicable authorization and a separately documented set of ecological precautions. For an encountered animal, such precautions might prioritize distance, cessation of targeted illumination, rerouting, or withdrawal if an observable behavioral condition is met. But “withdraw” is not automatically safe: rapid acceleration or reverse thrust may create an additional exposure or collision pathway. The relevant supervisory action is the one supported by a mission-specific safe-state analysis. A halt may permit a tether to drift onto fragile habitat, while a powered departure may disturb sediments. Thus, operational safety states require empirical validation for each morphology, environment, and ecological context.

22.7 A comparative protocol for observer effects and animal response

A feasible field design should establish the ecological question and its target estimand before selecting an apparently unobtrusive platform. Consider a bounded, non-contact fish or pelagic-organism survey. The research team could compare a conventional illuminated survey, an otherwise matched reduced-light survey, and, where ecologically and operationally appropriate, a fixed or stand-off acoustic reference. Assign routes or illumination blocks with a declared randomization or balancing scheme, preserving depths, observation time, habitat type and vehicle speed as recorded covariates. Where alternating illumination is used, the design must also account for possible behavioral carryover and changes in ambient conditions. Any exposure comparison should first receive the required biological, site and operational approvals; scientific randomization is not a reason to create unreviewed harm.

The protocol requires a baseline defined with respect to an explicit target, rather than the convenient but potentially distorted observations of the vehicle itself. For each block it would record calibrated spectral irradiance, source and received acoustics where feasible, hydrodynamic operating mode, distance to

Table 22: Biological evidence interfaces for prospective robotic observations. The examples are research-design requirements, not experimentally validated welfare limits.

Evidence object	Defensible observation or analysis	Remaining inferential limitation
Physical exposure	Received acoustic or light levels, mode, duration, relevant uncertainty	Receptor and biologically effective dose may remain uncertain
Encounter response	Matched trajectories, approach/avoidance, independent reference and contextual controls	Behavioral reaction does not reveal subjective state
Sampling distortion	Detection comparison by platform and operating mode; sensitivity analysis	Unobserved departures can remain unidentified
Individual welfare	Validated, taxon-relevant indicators and carefully interpreted synthesis	Species and context may preclude mental-state inference
Population or habitat	Longitudinal demographic or habitat comparison with counterfactual assumptions	Short trials rarely establish cumulative or lasting change

an animal when estimable, animal responses, environmental conditions, sensor visibility, and failures or exclusions. When independent acoustics and video disagree, that discrepancy becomes a result requiring interpretation. Outcomes should be reported in separate columns: instrument detectability; platform-specific observed counts or trajectories; standardized responses to encounter; feasible welfare-relevant indicators; and any longer-horizon ecological measures justified by the design. *No detected avoidance* is not equivalent to *no disturbance*, while a detected avoidance response cannot by itself establish lasting welfare harm.

The comparison would count against a proposed benign-design claim if a reduced-exposure platform nevertheless systematically displaces organisms, or if its apparently improved measurements are explained by attraction. An inconclusive result is also possible when effective baseline measurement, independent detection, or minimally harmful comparison cannot be established. In such cases the appropriate result is a bounded description of what can be inferred, not a default declaration of ecological neutrality. This comparative approach connects the source-specific evidence in [51, 17, 177] to the larger validity programme of Section 25.

22.8 A model of ecological encounter dossiers and reviewable autonomy

To connect engineering evidence with institutional review, let a mission encounter dossier be

$$\mathcal{D}_e = (\ell_e, \tau_e, r_e, a_e, \xi_e, y_e, \mathcal{M}_e, \mathcal{P}_e, \nu_e), \quad (87)$$

where ℓ_e and τ_e represent estimated location and time with uncertainty, r_e the robotic configuration, a_e the action sequence, ξ_e the physical exposure estimate, y_e the animal and habitat observations, $\mathcal{M}_e$ the inference model and its uncertainty, $\mathcal{P}_e$ the applicable permission and precaution dossier, and ν_e the linked review/version history. The tuple is deliberately typed: a model-predicted low acoustic dose is not an animal's reported experience, and a valid permission is not a proof of insignificant habitat impact. Some entries may be unavailable or protected. Missing entries should be identified as unknown rather than encoded as zero disturbance.

A corresponding engineering review may distinguish three questions. First, *scope*: was the activity covered by the permission and scientifically specified objective? Second, *evidence*: what exposure and response were measured, how well calibrated are they, and what competing interpretations remain? Third, *continuation*: under what justified conditions can the action proceed, be modified, be stopped, or be subjected to independent evaluation? These categories complement the legal and operational dossiers in Sections 8 and 10. They do not assign legal competence to a particular algorithm or replace regulatory and ethics review when required.

A conditional policy may assign actions to states such as ordinary operation, restricted observation, cessation of a targeted stimulus, habitat-preserving withdrawal, and incident review. Its transition criteria should be based on validated indicators, explicit uncertainty and the most defensible available action.

It should also record overrides, communication losses, operator interventions, and post hoc corrections. Machine-learning systems introduce a further issue: an animal-detection model can be uncertain precisely in rare, turbid or visually unfamiliar settings. A failure to detect an animal should not be interpreted as confirmation that none is nearby. Where the uncertainty is consequential and direct evidence inaccessible, conservative routing or escalation can be preferable to increased sensor intensity that itself produces additional exposure.

Responsible data circulation follows from these constraints. Sharing encounter records can make ecological investigations more reproducible and reveal previously unrecognized impacts. Yet public disclosure of vulnerable-species coordinates can create independent risks, while aggregate reports may erase individual encounters or local context. The open-science and common-knowledge programme in Section 9 therefore requires selective access, traceable transformations, explicit uncertainty, and procedures to correct or challenge ecological claims. A normalized “welfare index” might aid comparisons only if its constituent observations, taxon scope, weighting choices and missing data remain visible and contestable.

22.9 A conditional bound on modeled cumulative exposure

An elementary mathematical result clarifies what conservative engineering screening can establish. Consider a finite operation interval $[0, T]$ and a finite set of exposure channels $k \in \{1, \dots, K\}$, such as received acoustic energy, illumination dose, or a specified measure of sediment disturbance. Let $e_k(t) \geq 0$ denote actual channel exposure at a stated receptor or location, and let $\bar{e}_k(t) \geq 0$ be an empirically justified upper envelope that accounts for measurement and propagation uncertainty. Let $w_k(t) \geq 0$ be an adopted weighting for a clearly defined *engineering screening surrogate*, not a directly observed quantity of animal suffering. The modeled cumulative screen is

$$S_T = \sum_{k=1}^K \int_0^T w_k(t) e_k(t) dt, \quad \bar{S}_T = \sum_{k=1}^K \int_0^T w_k(t) \bar{e}_k(t) dt. \quad (88)$$

The weights and envelopes must refer to compatible units, receptor assumptions, and specified operating conditions. More precisely, for a dimensionless S_T each w_k must have dimensions reciprocal to the product of the exposure-rate units of e_k and the chosen time unit, or the separate channels must first be normalized by independently justified reference doses. Adding unnormalized sound exposure, light dose and sediment loading is dimensionally undefined. A vector of separate bounds $\int_0^T e_k(t) dt \leq B_k$ is often scientifically preferable where there is no defensible common screening scale. The formulation does not assert that distinct physical stressors can be added to produce a universal biological-welfare quantity; it merely defines a transparent surrogate for a particular screening decision.

Proposition 22.1 (Conditional containment of a cumulative exposure surrogate). *Suppose e_k and $\bar{e}_k$ are nonnegative measurable functions, $0 \leq e_k(t) \leq \bar{e}_k(t)$ for almost every $t \in [0, T]$, and w_k are nonnegative measurable functions with finite weighted integrals. Then $S_T \leq \bar{S}_T$. Consequently, if $\bar{S}_T \leq B$ for a chosen screening budget B , the actual defined surrogate satisfies $S_T \leq B$.*

Proof. For every channel, $w_k(t) \geq 0$ preserves the assumed pointwise inequality under multiplication. Integrating over the interval and summing across the finite channel set gives

$$\sum_{k=1}^K \int_0^T w_k(t) e_k(t) dt \leq \sum_{k=1}^K \int_0^T w_k(t) \bar{e}_k(t) dt. \quad (89)$$

The stated budget implication follows by transitivity. $\square$

The proposition is intentionally modest. Its conclusion can fail operationally if an exposure channel is omitted, propagation envelopes are underestimated, receptors move outside the modeled domain, biological susceptibility differs from the adopted screening assumption, or the action modifies later exposure

conditions. Moreover, $S_T \leq B$ establishes neither the absence of pain nor the absence of population effects, because a mathematical screen is not a validated biological dose–response relation. A low acoustic surrogate cannot compensate for unmodeled direct contact, and a numerically acceptable sum does not overrule binding restrictions on a protected species or habitat. The policy question of whether B is acceptable requires independent ecological, ethical, and legal justification.

The result nevertheless has a constructive role. It makes assumptions about the exposure model inspectable, encourages measurement of conservative envelopes, and clarifies what information should trigger revision. If a monitoring campaign shows the envelope is invalid, the proper continuation is not to preserve the original certificate of compliance by changing the recorded observation: the domain, model, budget, and operational authority must be reviewed. This principle connects the locally valid dynamics and supervisory governance of Sections 13 and 14 to the epistemic constraints of ecological investigation.

22.10 Validation, transnational scenarios, and a prospective research agenda

Evaluation of welfare-sensitive robotic systems should retain distinct endpoints. *Physical endpoints* include source-level and received acoustic measures, spectral irradiance, proximity, thruster wash, contact force, material release, and sediment disturbance. *Biological endpoints* include appropriately interpreted movement, feeding, communication, physiology, and longer-term habitat use. *Scientific endpoints* include detection quality, alternative-hypothesis discrimination, observer bias, and whether less intrusive measurements would suffice. *Institutional endpoints* include legal scope, procedural approval, transparent incident records, opportunity for independent assessment, effective cessation where required, and feasible remedy. Improvement on one axis cannot be presumed to compensate for deterioration on another.

The hypothetical transnational marine computing and observatory scenario developed in Sections 7–9 illustrates the stakes. Suppose operators from jurisdictions A, B and C support a marine data facility and associated robotic inspection programme near waters subject to coastal State D’s jurisdiction. The facility creates scientific data and computing capacity while inspection vehicles generate sound, light, maintenance traffic and possible near-bottom disturbance. The engineering question is whether alternative inspection configurations can achieve their specified maintenance and ecological-monitoring tasks at lower justified exposure. The legal question is which permits and duties actually apply to particular activities and locations. The distributive question concerns who receives infrastructural and epistemic value while bearing local ecological risks. These must be investigated separately and brought together through reviewable decisions; nothing in the scenario assumes a lawful installation, uniform national rules, or an established ecological impact.

Comparative validation should begin with physical characterization and nonanimal surrogates where possible, followed by ethically reviewed animal-relevant studies only where a scientifically justified question cannot be answered otherwise. Prospective field evaluation should incorporate passive reference instruments, encounter logging, blinded annotation where feasible, independent environmental sampling, control or comparison sites, and monitoring beyond the active mission where the plausible exposure requires it. Potential benefits should be stated against an explicit counterfactual; they should not be inferred from an appealing morphology, improved observation resolution, or the word “autonomous.” Failure scenarios should include false-negative animal detection, lost communications, persistent lighting after control failure, retrieval-induced habitat contact, a damaged acoustic transducer, sediment resuspension in sensitive habitat, and inability to verify an assumed low-exposure condition.

Animal welfare and ecological accountability are therefore integral to engineering rather than external commentary appended after technical development. The resulting research object comprises physical exposure, the organism’s partially observed response, changing ecological contexts, and institutions capable of revising action when evidence changes. A responsible architecture does not claim to eliminate uncertainty or determine the interests of every marine organism. It constructs narrower, falsifiable claims and connects them to alternative designs and practical restraint. Section 23 investigates which biolog-

ically informed technologies may genuinely contribute to this programme; Section 24 considers more extensive oceanic infrastructure under ecological finitude; and Section 25 develops comparative tests across scientific, engineering, and normative dimensions. No original welfare trial, field experiment, regulatory approval, or ecological-safety certification is asserted in this paper.

23 Biologically Informed Future Technologies

23.1 The research object: transferable mechanisms and changing biological contexts

The preceding engineering sections established a differentiated inventory of marine platforms, mechanical bodies, sensors, controllers, communication systems, interventions, and maintenance conditions. Section 22 then examined the ecological and welfare consequences of employing these capabilities. The purpose of the present section is to identify *experimentally investigable future technologies* that can arise from encounters with living marine systems. Its object is neither a catalogue of appealing animal shapes nor a prediction that marine robotics will converge on a preferred biological design. It is a set of candidate mechanisms, transfer conditions, evidential thresholds, and potential research continuations. Biological variation offers mechanisms worth investigating; the conditions under which they can be realized in artificial materials and useful marine operations remain separate empirical questions.

A productive distinction is made among *biological observation*, *mechanistic abstraction*, *synthetic implementation*, *system integration*, and *ecological deployment*. Observing a ray's fin motion does not establish which combinations of stiffness, unsteady lift, elastic recoil, motor timing, or sensory feedback are necessary for a particular robotic task. A controlled robotic reproduction of one fin mechanism does not establish energetic efficiency over a full autonomous mission. Integration with batteries, seals, navigation, and instruments changes the dynamics again. Finally, the deployment of a technically successful design requires environmental and institutional examination of its physical interactions. The 2026 interdisciplinary review by Li and colleagues illustrates a reciprocal scientific programme in which biology informs robot construction and robotic experiments help discriminate biological hypotheses [94]. Such reciprocity does not imply that a robotic surrogate reproduces an organism's lived ecological situation.

The generative-relational (GR) perspective gives this division a methodological role. A biological mechanism is learned within a situated ecological trajectory; a corresponding robotic mechanism is constructed in a distinct material, energetic, and institutional trajectory. The central research question is which relations can be reproduced, transformed, or deliberately discarded while retaining evidence for the desired engineering function. This formulation prevents a supposed ecological virtue from being transferred solely by visual resemblance. A compliant robot might reduce one contact stress while increasing residence time, illumination, debris risk, or propulsive disturbance. Accordingly, a promising biological analogy must specify its *target function*, *counterfactual alternative*, and *validation domain*. The five technological pathways considered below are propulsion and adaptive morphology, distributed sensing, reversible contact and manipulation, responsive surfaces, and biohybrid platforms. They are complementary possibilities, not an inevitable sequence of technological development.

23.2 Propulsion, compliant structures, and adaptive locomotor repertoires

Fishes, rays, cephalopods, and gelatinous organisms illustrate differing ways to couple deformation with surrounding flow. Their relevance to engineering lies in distributed structural stiffness, elastic energy storage, changing hydrodynamic surfaces, and sensorimotor coordination under particular Reynolds-number and environmental regimes. Section 12 distinguished biological body plans from robot architectures, while Section 13 established why successful coupling must be interpreted through fluid-structure dynamics. Here the prospective step concerns *reconfigurable repertoires*: bodies whose shape, effective stiffness, actuation timing, or operating mode can be revised in response to changing task and fluid conditions.

Manta-inspired research demonstrates several concrete fragments of this agenda. Experiments on a flexible-fin robotic manta investigated three-dimensional path following, while a soft flapping swimmer

exploited elastic snap-through and fluid momentum transfer to produce rapid experimental locomotion [108, 158]. A more recent ray-inspired swimmer using bistable actuators demonstrated configurable swimming modes and remotely controlled untethered operation under its experimental conditions [223]. These results support the feasibility of particular body–actuator mechanisms. They do not establish superior mission-level energy consumption relative to mature propeller-driven or buoyancy-driven marine systems, nor do they provide evidence of long-duration ocean deployment across pressure, fouling, or variable-current envelopes. Body-length-normalized speed is useful for comparing motion within certain classes; it does not by itself equate payload, power-system mass, scientific sensor quality, or ecological exposure.

One prospective architecture is a *mode-switching body* in which the same structural surface performs different functions, such as efficient transit, low-disturbance approach, and limited station keeping. This may reduce added moving parts, but switching can introduce fatigue, control transients, or losses in one mode that offset gains in another. The design problem is therefore not the isolated maximization of swimming efficiency. It is to identify a feasible collection of behaviors and the physical conditions under which transitions remain observable, reversible, and mechanically supported. Testing should compare equal payloads, matched missions, power including peripheral pumps and electronics, thermal behavior, and repair cycles. A proposed advantage in sensitive habitat should additionally be tested through measured acoustic, optical, and hydrodynamic exposure, not inferred from the absence of a rotating propeller.

A second direction is to use robotic reproductions to probe biological explanations. Independent variation of stiffness, fin waveform, body flexibility, and flow environment may discriminate competing hypotheses about propulsion and sensory feedback. But a robotic body cannot reproduce animal metabolism, cognition, growth, or species-specific ecological relationships merely through selected mechanical similarities. Such experiments are valuable precisely when their retained and omitted mechanisms are made explicit. They can yield both a more useful robot and a narrower, better-supported biological hypothesis.

23.3 Distributed hydrodynamic, tactile, and multimodal sensing

Biological sensory systems motivate a shift from the assumption that all useful perception must pass through a few centralized imaging or acoustic instruments. Fish lateral lines, hydrodynamic-sensitive hairs, and distributed receptors suggest that information can be extracted locally from the coupling of body motion and environmental flow. A 2026 study by Park and colleagues integrated magnetic-hair flow sensors into robotic bodies and demonstrated estimation of robot motion together with detection of imposed wake structure in the tested flows [152]. This is a mechanistic result with clear potential for near-field navigation and environmental inquiry. It is not evidence that such sensors can replace global localization, active sonar, or calibrated oceanographic current measurements in arbitrary waters. Their readings depend on placement, body-generated disturbance, changing biofouling, and the extent to which flow and self-motion can be separated.

A plausible future platform therefore integrates multiple sensing layers: local deformation and pressure for hydrodynamic contact; inertial and acoustic positioning for larger-scale motion; optical or active acoustic instruments for selected scene reconstruction; and chemical or biological sampling for ecological interpretation. These channels often observe different physical variables at different scales. The mathematical challenge is not simply to stack sensors or average measurements. It is to infer which latent physical causes remain identifiable when motion, flow, turbulence, and measurement noise are mutually coupled. In particular, a stronger correlation between a bioinspired sensor and vehicle speed during one towing-tank calibration does not guarantee that the same mapping holds when the vehicle enters a vortex, brushes a surface, or carries a different payload. The observability concerns of Section 16 thus remain operative.

Distributed sensing may also permit *local response before central interpretation*. Octopus-inspired robotic research in 2026 combined sensors embedded in suction cups with a hierarchical control sys-

tem for flexible underwater grasping [31]. The immediate engineering opportunity is that local contact measurements can modify grip or release behavior without reconstructing the whole scene. Yet such distributed control raises new interface problems: two contacts may demand conflicting actions, biological targets may be misclassified, and local sensors may disagree about forces under deformation. A coherent system should record both local trigger events and the supervisor’s authority to stop, override, or revise them. The presence of sensors in a biologically suggestive arrangement supplies neither a theory of animal perception nor a guarantee of safe physical interaction.

For scientific platforms, one further possibility is to treat the body itself as a distributed observing instrument: coordinated changes in compliant structures may provide information about local fluid mechanics. Such inferences demand calibrated mechanical and fluid models, uncertainty estimates, and independent reference instrumentation. They can be valuable when they augment rather than silently replace established observations. The aim is a multi-scale and falsifiable perception architecture, not an assertion that synthetic senses inherit the adaptive capacities of their biological analogues.

23.4 Reversible adhesion, soft grasping, and local intervention intelligence

Many marine scientific and maintenance operations require contact with irregular surfaces. Octopus suckers, remora attachment structures, and compliant appendages suggest alternatives to rigid gripping jaws, but they solve different tasks. A remora-inspired robotic adhesion disc demonstrated attachment strategies relevant to temporary transport and anchoring [210]. A comparative analysis of octopus-inspired suction cups identified substantial progress in adhesion and integration while noting unresolved deficits in sensing and fully soft system-level operation [204]. Subsequent work has demonstrated switchable attachment to challenging wet surfaces using compliant octopus-inspired interfaces [88]. These studies support candidate mechanisms for reversible contact; they cannot establish reliable anchorage on unknown marine structures contaminated by sediments, biofilms, irregular roughness, or damage.

A mechanism-sensitive design separates *approach*, *surface recognition*, *seal formation*, *holding*, *load transfer*, and *release*. A sucker that performs well during perpendicular pull-off tests may fail under shear, peel, dynamic load or repeated actuation. In an inspection mission, adhesion must also respect the substrate: a biological surface, archaeological artifact, coated cable, glass sensor housing and unfinished steel plate are not interchangeable contacts. A future intervention robot might combine suction pads, soft continuum arms, and distributed tactile feedback to localize a holding surface while keeping reaction loads away from fragile biological features. The correct validation comparison is against specified conventional tools and their measured contact distributions, not against an idealized hard gripper assumed to be harmful.

The prospective concept of *local intervention intelligence* arises when a manipulator can distinguish uncertain touch, stable contact, grip loss, and tool misalignment using peripheral signals. The 2026 sensorized octopus-like arm is relevant because it links contact sensing with behavior selection, although its experiments do not demonstrate complete autonomy for complex oceanic construction [31]. Stable attachment can enable station keeping, inspection and low-force material handling; accurate assembly additionally requires alignment, fixture tolerance, calibrated forces, vehicle stabilization, and knowledge of the structure being modified. Those requirements remain grounded in the coupled contact mechanics developed in Section 17.

The ecological question is equally specific. A soft adhesion mechanism might reduce peak impact but remove mucus, damage tissue, increase attachment dwell time, or spread biological material between sites. It may also leave fragments after failure. Future studies should quantify attach–detach cycles, local stress and strain, contamination, sensor drift, tool retrieval, and long-term substrate consequences. Demonstrating low contact pressure on an artificial surface is a useful mechanical result; declaring an unfamiliar organism safe to grasp requires evidence unavailable from that measurement alone.

23.5 Responsive skins, adaptive materials, and ecological interfaces

Cephalopod skins invite investigation of responsive optical and mechanical interfaces. Research on synthetic chromatophores has produced stretchable, stimulus-responsive hydrogel arrays with programmable color and pattern changes, drawing on selected mechano-optical principles of cephalopod skin [212]. This demonstrates a materials architecture rather than a field-ready camouflage device for oceanic robots. Light attenuation, ambient spectrum, pressure, fouling, material stability and power demands change markedly across surface, coastal and deep-ocean environments. The biological ability to change appearance also has neural, behavioral and ecological components that do not follow from reproducing one optical mechanism.

Adaptive surface technology could serve functions other than concealment. A change in surface state might provide a low-power external status indicator, reveal localized mechanical damage, vary wettability, or regulate attachment to inspection fixtures. Reversible stiffness or texture changes could help a robot move between locomotion and manipulation modes, provided these material transitions remain within validated durability limits. Each proposed use must be tested against the requirements of marine science. For example, a visible status surface is unlikely to assist an entirely dark, distant deep-sea survey; an optically responsive coating may interfere with the reflectance of calibration targets; and chemical additives or shed fragments may create ecological liabilities. Nor should concealment of a vehicle from animals be presumed to reduce disturbance: camouflage may change proximity and encounter dynamics in ways that require direct biological measurement.

Material design must account for the long-term transformations studied in Sections 18 and 19. Useful properties in fresh water or laboratory saline baths may degrade after cycles of pressure, temperature, abrasion, colonization and cleaning. A seemingly repairable soft skin may require adhesives or multi-layer fabrication that complicate separation and recycling. In a lifecycle comparison, one should record material composition, component replacement rate, detached particulate mass, cleaning chemistry and retrieval procedures alongside any immediate actuation benefit. These costs cannot be subsumed in a qualitative designation such as “biological” or “soft.”

A more defensible future direction is *ecologically informed interface engineering*: choose functions that directly address an identified operational mechanism, use synthetic materials whose aging is measurable, and design for access, repair and eventual recovery. Biological observation can help generate candidate geometries and materials, but ecological accountability depends on a distinct evidence chain linking design decisions with observed physical and biological consequences.

23.6 Biohybrid platforms, living participants, and the limits of substitution

Biohybrid systems bring a different technological and moral problem into view. Rather than implementing a mechanism in entirely synthetic materials, they incorporate living cells, tissues or organisms into a machine. Experiments by Anuszczyk and Dabiri investigated electromechanically enhanced live jellyfish with externally modulated swimming and mechanical modifications in controlled settings [12]. A 2026 study reported pressure-triggered vertical profiling by biohybrid *Aurelia aurita*, including temperature and depth measurements during limited natural-marine field trials reaching depths greater than 25 metres [159]. These are significant demonstrations of instrumented biological locomotion and field measurement within stated bounds. They do not establish the welfare of jellyfish under chronic instrumentation, suitability for arbitrary pressure or salinity regimes, long-term ecological innocuousness, or a general replacement for autonomous mechanical floats.

A living organism has metabolic needs, responses to injury, interactions with other organisms, and possible welfare interests whose scientific characterization is incomplete. A technological accounting that records only purchased electronics and electrical propulsion energy omits the conditions sustaining the animal’s activity. Consequently, low electrical expenditure cannot alone establish an ethically or ecologically inexpensive platform. Device attachment, stimulation, imposed trajectories, altered feeding, reproductive opportunities, predation exposure, retrieval, and possible release are relevant parts of the research design. Welfare evidence should be species- and procedure-sensitive, should include appro-

priate animal-relevant controls, and should be reviewed independently when the intervention warrants it. Whether a given animal falls under a specific legal animal-research regime varies by jurisdiction; a gap in positive legal coverage cannot establish the absence of ethical considerations.

Biohybrid robotics can be scientifically valuable even when wide deployment is inappropriate. It may clarify fluid–structure interactions, animal biomechanics, or the extent to which a living propulsor responds to additional loads. Some questions can be investigated through synthetic analogues, passive observation, or short and carefully reviewed trials. A staged research programme should therefore first specify why a living component is required, compare alternatives, and separate basic biological discovery from a proposed commercial monitoring service. The point is to make the technology and its justification jointly revisable, not to render all contact with living organisms categorically impermissible.

For the generative-relational framework, this distinction is fundamental: a jellyfish is an ecological participant with a trajectory of its own, not merely a replaceable actuation module. Describing its metabolism as “free energy” would erase that relationship. The prospective value of biological integration must be evaluated in relation to animal interests, ecological effects, scientific benefit, institutional responsibility, and the conditions under which an experiment can be stopped or modified. Synthetic bioinspiration and living biohybrids should consequently occupy separate columns in engineering evidence and ethics reviews.

23.7 A bounded formalism for biological-to-engineering transfer

A comparative transfer framework can clarify when a mechanism observed in one set of environments supports a claim in another. Let $\mathcal{C}$ be a metric space of operating contexts, with a chosen metric d . A context may encode pressure, salinity, temperature, relative flow, surface type, payload, and relevant action constraints; its representation is itself a modeling choice. Let $F : \mathcal{C} \rightarrow \mathbb{R}^q$ describe a measurable target response, such as a set of forces, sensor outputs, or energetic quantities, and let $G : \mathcal{C} \rightarrow \mathbb{R}^q$ denote a candidate synthetic mechanism or predictive surrogate. The framework does *not* assume that F represents a living organism’s welfare, experience, or ecological identity. These are engineering observables with defined units and protocols.

Suppose a finite calibration set $C_0 = \{c_1, \dots, c_n\}$ covers a specified operational domain $D \subseteq \mathcal{C}$ within radius ρ . Suppose measured discrepancies at the calibration contexts are bounded by ε , after any justified measurement uncertainty has been included. This gives a conditional transfer statement.

Proposition 23.1 (Coverage-bounded mechanism discrepancy). *Let D be contained in the union of the closed ρ -balls about C_0 . Suppose F and G are Lipschitz on a domain containing each line segment or admissible metric path used in the comparison, with constants $L_F, L_G \geq 0$, and $\|F(c_i) - G(c_i)\| \leq \varepsilon$ for every $c_i \in C_0$. Then every $c \in D$ satisfies*

$$\|F(c) - G(c)\| \leq \varepsilon + (L_F + L_G)\rho. \quad (90)$$

Proof. For $c \in D$, choose a calibration context c_i with $d(c, c_i) \leq \rho$. The triangle inequality gives

$$\begin{aligned} \|F(c) - G(c)\| &\leq \|F(c) - F(c_i)\| + \|F(c_i) - G(c_i)\| \\ &\quad + \|G(c_i) - G(c)\| \\ &\leq (L_F + L_G)d(c, c_i) + \varepsilon \\ &\leq \varepsilon + (L_F + L_G)\rho. \end{aligned} \quad (91)$$

□

For a practical transfer claim, ρ must be defined in a calibrated metric (including scaling across pressure, salinity, temperature and other coordinates), and L_F, L_G must be independently justified on paths to the target domain. A finite set of trials does not by itself establish those Lipschitz constants. In high-dimensional operating domains, small covering radius may require infeasibly many qualified configurations; therefore the estimate can remain mathematically correct while offering no useful numerical guarantee outside a narrow experimental region.

Table 23: Selected biologically informed mechanisms and research tests. The table separates published demonstrations from proposed cross-context capabilities.

Mechanism	Published evidential scope	Prospective test and transfer barrier
Flexible fins and mode switching	Experimental ray-inspired locomotion and configurable actuation [158, 223]	Matched payload/power comparisons; oceanic endurance, transitions, fatigue, disturbance
Hydrodynamic sensor arrays	Flow and self-motion inference on tested robotic platforms [152]	Current separation, fouling, placement robustness, sensor calibration and navigation integration
Distributed suction and grasping	Sensorized soft-arm grasping and reversible wet adhesion [31, 88]	Force and seal reliability on variable surfaces; safe release, contamination and wear
Responsive artificial skins	Stimulus-responsive hydrogel color/pattern structures [212]	Pressure, spectral environment, aging, material release, ecological purpose
Living biohybrid profiling	Short ocean field trials with instrumented jellyfish [159]	Animal-relevant outcomes, sustained duty, retrieval, ecological and legal review

The proposition is a generic approximation bound, not a new law of biomimetics. Its value is diagnostic: good agreement at a few demonstration points can support only the domain justified by calibration coverage and variability assumptions. If a robotic sucker is tested only on a clean planar plate, then organism tissue, fouled steel, rough stone, and curved polymer may lie outside D . If a soft fin undergoes snap-through or abrupt contact, smoothness assumptions and finite Lipschitz constants may fail across the relevant transition. If one observes biological behavior only within a tank, the response function itself may change during field deployment. The metric must not disguise categorical differences by assigning arbitrary small distances to incompatible environments. Measurement uncertainty and model error must be budgeted rather than omitted to make ε appear smaller.

For engineering comparison, multiple response components can be retained separately: thrust, power, contact force, sensing reliability, materials loss, and exposure proxies need not be reduced to one scalar utility. The result cannot bound unobserved welfare outcomes, omitted species interactions, or legal compliance. It does, however, suggest an actionable experiment design: explicitly state D , measure responses on a sufficiently informative calibration set, test response regularity, record exclusions, and attempt independent challenges near the proposed operating limits. No empirical values for L_F , L_G , ρ , or ε are inferred or estimated in this manuscript; those quantities would require new measurements.

23.8 A comparative research programme and transitions to oceanic infrastructure

The foregoing mechanisms should be evaluated by the evidential status of their claims rather than by their novelty or biological familiarity. Table 23 summarizes candidate technological directions, present kinds of support, and principal transfer barriers. “Demonstrated” refers to the task and environment reported in the cited research, not a complete independent qualification assessment.

Three methodological priorities follow. First, *mechanism isolation* should precede elaborate integration: compare a candidate with an otherwise comparable control that removes the hypothesized biologically derived feature. Changing shape, controller, power source, sensors, and payload simultaneously may yield a better machine while preventing attribution of the improvement. Second, *cross-context validation* should test boundaries most likely to falsify the mechanism: pressure, temperature, salinity, turbulence, fouling, changing substrates, and human intervention during maintenance. The response quantities should be specified before testing, and unexpected failures retained as scientific data rather than excluded as “nonrepresentative.” Third, *ecological and institutional review* should be designed alongside the experiment. The questions of measurable animal disturbance, retrieving a failed robot, using living organisms, chemical release, and required permissions cannot be answered by the mechanism-transfer bound of Proposition 23.1.

A practical research ladder begins with biological observation that minimizes unnecessary intrusion, proceeds through simplified synthetic models, controlled mechanical experiments and simulated dynamics, then integrated prototype tests, independently calibrated trials and limited field studies under appro-

7190 puate authorization. Advancement need not be linear: evidence of ecological disturbance may require
7191 retreat to a less intrusive design even after an impressive engineering demonstration. Research groups
7192 should retain competing hypotheses about biological mechanisms and their artificial analogues rather
7193 than allowing one successful prototype to become a generalized description of marine life. Negative
7194 results have particular epistemic value when they identify which assumed relations failed to transfer.

7195 The hypothetical A–B–C–D marine installation from earlier sections provides a prospective system-
7196 level test. A facility might use flexible-fin inspection robots, suction-based station-keeping devices, and
7197 distributed flow sensors for different maintenance and observation tasks. An improvement in one com-
7198 ponent’s local propulsion cannot establish a net decrease in fleet energy, ecological disturbance, cost,
7199 or unequal allocation of benefits. Construction, computing demand, communication relays, sensor cal-
7200 ibration, damaged parts, material retrieval, and marine habitat interactions create system-wide effects.
7201 The question is whether heterogeneous technologies can be combined without obscuring those costs and
7202 responsibilities. Accordingly, Section 24 extends the analysis to prospective oceanic infrastructure and
7203 finite ecological resources. Section 25 will articulate comparative research designs capable of falsifying
7204 proposed engineering and normative advantages. The research remains exploratory: no new marine pro-
7205 totype, animal trial, field deployment, validated transfer envelope, or ecological certification is reported
7206 here.

7207 **24 Prospective Ocean Infrastructure and Finite Ecological Resources**

7208 **24.1 Oceanic infrastructure as a lasting transformation of ecological relations**

7209 The preceding sections have examined oceanic robotic technologies from individual bodies and sensors
7210 to heterogeneous fleets, adaptive observation and biologically informed mechanisms. The present sec-
7211 tion investigates a distinct research object: *infrastructural arrangements that persist beyond individual*
7212 *missions and transform the conditions under which subsequent marine inquiry and production take place.*
7213 Such arrangements may include cabled observatories, mobile robotic service systems, marine communi-
7214 cations and energy networks, submerged computing installations, scientific sampling stations, and engi-
7215 neered seabed structures. Their significance cannot be assessed by extrapolating the performance of a
7216 successful robot. An infrastructure changes the distribution of physical access, observational opportuni-
7217 ties, material exposure, legal authority, financing and responsibility across interdependent systems.

7218 In the generative-relational (GR) vocabulary of Sections 2, 3, and 6, these transformations concern
7219 the continuation of ecological and epistemic capacities. A permanently powered instrument may reveal
7220 a seasonal biological phenomenon previously missed by intermittent cruises. The same cable corridor
7221 may fragment a sensitive seabed habitat during installation, while the associated energy supply and main-
7222 tenance vessels impose additional pressures outside the instrument’s immediate footprint. A scientific
7223 dataset might become widely accessible even as the expertise and physical infrastructure necessary to in-
7224 terpret it remain concentrated. Accordingly, an apparently beneficial knowledge infrastructure can have
7225 heterogeneous consequences rather than a single unqualified ecological value.

7226 A first analytical distinction separates *scientific observatories*, *enabling service infrastructures*, *com-*
7227 *mmercial productive installations*, and *resource-extractive facilities*. The categories are provisional, may
7228 overlap, and do not confer legal status. A shared subsea power connection can serve both open scien-
7229 tific instruments and a privately controlled processing module. An offshore renewable-energy system
7230 can supply a marine station while altering habitat conditions and maritime uses in its own right. A re-
7231 search platform supporting biological sampling has different downstream material consequences from a
7232 facility constructed to recover commercially valuable minerals. Their physical coupling does not erase
7233 differences in purpose, authorization, benefit or risk.

7234 The concept of *finite ecological resources* is correspondingly broader than a stock of commercially
7235 recoverable materials. It includes availability of habitat and substrate, tolerable acoustic and optical ex-
7236 posure, thermal and chemical transport conditions, ecological recovery capacities, energy and material
7237 inputs, and the human institutional capacities required to inspect and correct operations. These are hetero-

geneous and sometimes only partially measurable. Some are renewable under certain conditions; some may recover over timescales irrelevant to an installation decision; some possess thresholds or irreversible changes that cannot be inferred from a resource balance. Neither technical endurance nor financial profitability therefore establishes ecological viability. The aim here is to specify prospective configurations and the comparative tests needed to investigate their effects without treating future expansion as either intrinsically beneficial or intrinsically illegitimate.

24.2 Sustained observatories and maintenance-capable scientific networks

Cabled marine observatories are among the strongest existing examples of operational oceanic research infrastructure. Ocean Networks Canada describes the NEPTUNE and VENUS networks as providing continuous subsea power and fibre-optic connectivity to ocean instruments; NEPTUNE includes an approximately 850-kilometre seabed cable loop and supports observations extending to approximately 2.6 kilometres depth [143, 142]. These are institutional descriptions of operating facilities rather than independent verification of every scientific inference drawn from their sensors. Importantly, continuous communication does not abolish biological sampling limitations, instrument drift, spatial sparsity or the need for ship- and ROV-supported maintenance. The Ocean Observatories Initiative and other programmes similarly integrate fixed platforms with mobile vehicles and shore-based information services [146, 147].

The prospective scientific opportunity lies in *persistent heterogeneity*. A seabed observatory might supply calibrated background observations while mobile gliders traverse a changing front, robotic samplers collect selected biological material, and intervention vehicles retrieve fouled or damaged modules. The service network would have to reconcile different sensor accuracies, measurement times, vehicle coordinate systems, repair histories and authority boundaries. Section 21 established how adaptive targeting can create selective datasets; a sustained fixed network can aid calibration and validation but cannot provide a universal unbiased baseline merely by being stationary. Its locations were themselves selected, and the installation may alter nearby organism behavior or flow.

For future infrastructure research, *serviceability* should be treated as a design variable alongside scientific resolution. Consider two sensors that record the same variables under laboratory calibration. One is embedded in a difficult-to-recover, single-use housing; the other can be replaced by a robotic end effector without disturbing its anchoring foundation. Their lifecycle costs, calibration continuity and ecological effects may differ substantially, even before considering the accuracy of their measurements. Modular power, data and mechanical interfaces are attractive only when connector durability, manipulator alignment, tolerable contact forces, cleaning chemicals, biofouling and safe recovery have been investigated. Sections 17 and 19 provide the relevant engineering foundations. A design that avoids an annual ship visit but requires frequent replacement of unrecoverable batteries has not necessarily reduced lifecycle disturbance.

Institutional continuity is equally material. An observatory requires preservation of equipment inventories, calibration provenance, access arrangements, cybersecurity maintenance, data stewardship, environmental monitoring, and responsibilities when funding changes. Open data alone does not create an open or resilient scientific commons; research communities may also require computation, methods, technical participation and durable access to the underlying instruments. The legal and value-circulation concerns of Sections 7 and 9 therefore enter at the design stage. A credible infrastructure proposal should identify who can observe, maintain, contest, suspend and retire each component, including components financed or operated by different organizations.

24.3 Subsea computation, power corridors and thermal consequences

Subsea computing is an instructive prospective case because it couples highly visible economic value with comparatively unfamiliar ecological exposures. Microsoft Research's Project Natick deployed a full-scale module in the North Sea in June 2018 and retrieved it in July 2020 after more than two years of operation; its published project account reports lower server failure rates than a terrestrial comparison

group and envisages recycling the retrieved equipment [110, 109]. The result is relevant engineering evidence for a particular research installation. It is not a controlled demonstration that all marine data centres will be more reliable, economically preferable, carbon-neutral, ecologically harmless or generally licensable. Deployment, retrieval, system scale, data processing, power sourcing and heat dissipation must be assessed independently.

Power and thermal design illustrate this boundary. A computing module requires an electrical supply, conversion and distribution equipment, data connections, cooling interfaces and emergency responses. Offshore renewable generation can supply some electricity, but the physical generation system, intermittency management, storage, interconnection and construction lifecycle must appear in the accounting. A claim of low operational emissions does not automatically include embodied material emissions or infrastructure maintenance. Furthermore, seawater heat exchange transfers energy to an environmental sink: using ambient water as coolant does not cause the rejected thermal energy to disappear. A local temperature signal depends on discharge design, power density, ambient currents, stratification, turbulence and ecological sensitivity. Model predictions require field calibration and cannot stand in for site-specific ecological evidence.

If $P_{IT}(t)$ is computing power and $P_{aux}(t)$ accounts for supporting loads, a simplified first-law relationship is

$$P_{in}(t) = P_{IT}(t) + P_{aux}(t) + P_{loss,other}(t), \quad \dot{Q}_{env}(t) + \dot{U}(t) = P_{heat}(t). \quad (92)$$

Here $\dot{U}$ represents temporary thermal energy storage and $\dot{Q}_{env}$ denotes heat transferred to the environment across a specified boundary; the equality is a deliberately simplified control-volume balance. The actual partition must include exported energy, work, coolant transfer, geometry and boundary conditions. Even a precise power meter cannot establish the temperature experienced by an organism or the consequences for a local ecosystem. Those inferences require coupled transport models and observations of the affected ecological relationships.

Submarine power and communications cables also constitute infrastructure rather than invisible links. NOAA's review of cable studies identifies trenching, vessel operations, seabed disturbance, thermal and electromagnetic effects, and possible release of insulation or cooling fluids in certain systems as potential concerns, with impacts varying by cable type and habitat [139]. The footprint of a cable corridor can extend beyond the narrow line used to depict it on a map. Installation, burial, inspection, repair and eventual abandonment or removal have distinct effects. An honest comparison should therefore distinguish a data centre's direct substrate footprint from the connected energy and communication network and from displaced activities at other locations.

A possible research architecture is a modular marine compute-and-observation node designed to support local scientific processing, reduce selected data transfers, and retain the ability to retrieve failed modules. Such a system would have to justify why marine deployment is preferable to a shore facility or surface platform serving the same scientific questions. Reducing network latency, improving cable access or offering experimental refrigeration opportunities may be relevant engineering objectives. None supplies an automatic ecological preference. Comparative studies should use a common service demand, a matched observation schedule, lifecycle boundaries, and quantitative monitoring of thermal and construction effects. A design can be scientifically productive while remaining inappropriate in a habitat where the relevant uncertainties are too large or available safeguards insufficient.

Research demonstration, operator account, and general ecological outcome. Project Natick is an instructive case for separating evidence of feasibility from broad infrastructure comparisons. Microsoft Research's published account reports a lower server failure rate for its North Sea module than for a specified onshore comparison group [110, 109]. This is a project-reported, context-specific reliability comparison, not an independently established causal estimate for future fleets, climates or maintenance regimes. It also cannot alone establish a lower total ecological burden. A comparably scoped study would need the baseline failure denominator, component and maintenance differences, electricity mix, deployment and recovery logistics, monitoring duration, environmental heat transport and ecological

outcomes. The same evidence rule applies if a proposed installation is justified as supporting biodiversity research: beneficial scientific computation is one output and changes to ecological conditions are a different object of measurement.

24.4 Offshore energy, engineered habitats and cumulative spatial pressures

Prospective infrastructures may combine energy generation, scientific observation and ecological restoration functions. Offshore wind turbines, floating energy platforms, wave-energy converters and subsea substations offer potential power to remote instruments, but they alter ocean space as well as energy supply. A 2025 review of offshore wind and biodiversity identifies both adverse pathways, including construction noise, animal displacement, altered habitats and barriers, and locally beneficial effects such as artificial reef or refuge formation under particular conditions [211]. Neither an observed concentration of organisms around a structure nor an increase in one local biodiversity index establishes a net regional ecological benefit. Aggregation may redistribute existing organisms, change predator–prey encounters, or expose them to new hazards.

The physical placement of marine installations is therefore a multi-actor decision involving navigation, fishing, conservation, scientific research, cultural relationships, security and economic infrastructure. A seemingly vacant patch of seabed may participate in ecological flows and in human livelihoods that are not captured by a standard engineering bathymetric survey. Baselines can be incomplete because habitats change seasonally or monitoring technologies cannot reliably observe small, cryptic or migratory organisms. Cumulative assessment should consider multiple interacting developments rather than project-by-project compliance alone, while distinguishing ecological effects within and beyond national jurisdiction.

Prospective ecological infrastructure can also include purpose-built habitat monitoring and limited restorative interventions. Robots might deploy recoverable sensor frames, place experimental substrates, inspect restoration structures, or remove abandoned equipment. Such activities can generate useful knowledge about colonization, sediment dynamics or material degradation. Yet physically creating an apparent habitat does not necessarily repair the system it replaces, and the introduction of engineered surfaces can facilitate unwanted colonization or alter substrate-specific ecological relationships. Baseline and control-site comparisons are particularly important where the project itself changes the habitat whose recovery it claims to measure. Where ecological reconstruction is a stated objective, the definition of recovery must be explicit, plural when warranted, and sensitive to actual ecological histories rather than a single unexamined target state.

Energy-system comparisons should also include terrestrial alternatives, deployment logistics, and recycling. It is easy to count electricity harvested offshore while omitting fuel for installation vessels, protective coatings, converter replacements, and waste pathways. An operationally self-powered observatory can remain materially dependent on distant extraction and manufacturing. Conversely, an expensive shared cabled observatory may replace some repeated vessel expeditions and improve access to scientific data. Whether this improves environmental performance depends on measured mission substitution, system lifetime and ecological conditions. The relevant question is not whether energy and knowledge are renewable in the abstract, but which capacities are replenished at what rate, through which material processes, and with whose authority to determine acceptable losses.

24.5 Mineral extraction, juridical boundaries and differentiated production regimes

Marine mineral extraction illustrates the limits of treating infrastructure expansion as an undifferentiated oceanic future. The United Nations Secretary-General’s Scientific Advisory Board warned in 2025 that deep-sea mining could produce wide-ranging, long-lasting and potentially irreversible impacts through habitat destruction, sediment plumes and associated disruptions [199]. These concerns require empirical investigation differentiated by deposit type, mining method, depth, affected organisms and environmental transport. A development application cannot be evaluated solely by extrapolating from a successful small-scale manipulator test. Ore recovery, processing, equipment traffic, plumes, waste, monitoring

and remediation form a larger causal system. Some observed environmental changes may persist long after profitable extraction ceases.

Distinct legal regimes apply to the water column, continental shelf and international seabed Area. As discussed in Section 7, UNCLOS allocates specified coastal-State rights in the exclusive economic zone and over the continental shelf; the International Seabed Authority's competence concerning mineral resources of the Area is not universal authority over all high-seas science, subsea data processing, or communications installations [190, 191, 193, 78]. The BBNJ Agreement entered into force on 17 January 2026 and addresses, among other matters, marine genetic resources, area-based management tools, environmental impact assessment and capacity building within its defined regime [189, 188]. Its operation depends on the relevant activity, location and parties; it does not create a general international revenue-sharing levy on subsea computation or replace every sectoral instrument. A specific facility would require an activity-by-activity assessment of applicable instruments, domestic implementation and competent authorities.

This differentiated scope matters for GR political economy. Consider a hypothetical installation involving an investor in State A, robot designers in State B, operators in State C, and the affected coastal State D. Financial revenues, intellectual property, engineering capacity, ecological effects and scientific knowledge may circulate through different relations and institutional registers. The same infrastructure may create publicly useful ecological observations and commercial benefits while requiring local habitat protection, maintenance services and public or community participation. The presence of cross-border value circulation does not determine ownership, tax jurisdiction, environmental liability, benefit-sharing obligations, or authority to refuse an installation. Those consequences depend on actual legal sources and agreements. Nor should one presume that the scientific use of an extracted material compensates for loss of habitat or procedural exclusion.

The hypothetical case permits comparison of alternative production forms. A closed commercial facility, a publicly owned observatory, a shared research consortium and a commons-governed data repository may distribute access and control differently. Their ecological performance is not established by the vocabulary of ownership. A commons arrangement with inadequate maintenance funding may become a source of unrecovered debris; a carefully regulated commercial operator may supply reliable monitoring while retaining inequitable access to its data; a public facility may concentrate decision authority and exclude situated expertise. Evaluation therefore requires tracing who controls the capacity to produce further knowledge, how scientific benefits and ecological burdens are distributed, and which procedures enable contestation, correction and remedy. These questions extend the heterogeneous value analysis of Section 9 without converting it into an invented positive legal rule.

The scale and mechanism of the seabed disturbance comparator. The long-term study by Jones and colleagues revisits a 1979 experimental polymetallic-nodule mining track and reports persistent physical and biological changes in directly disturbed areas after about four decades, together with evidence of re-establishment in some groups and more limited detectable residual effects in certain plume-affected areas [83]. The relevant exposure includes nodule removal, tracked or screw-driven collector contact, and disturbed sediment. This evidence is directly relevant to proposed intensive mineral extraction. It does not provide a transferable prediction that routine camera surveys, small recoverable observatory anchors, or all robotic seabed visits will have the same magnitude or duration of effect. Conversely, the smaller footprint of a proposed device does not establish ecological safety. A defensible comparison must identify disturbed area and depth, substrate type, organism assemblage, contact mechanics, revisit frequency, plume transport, and observed recovery processes before moving from one intervention class to another. The persistence of damage and the possibility of partial biological re-establishment should both be represented, rather than compressing distinct ecological endpoints into a binary label of recovery.

Table 24: Prospective oceanic infrastructure families, evidence status and non-substitutable assessment dimensions. This is a research comparison rather than an authorization or ranking.

Infrastructure	Evidence boundary	Required comparative assessment
Cabled observatories	Sustained field operation; periodic human/ROV maintenance	Calibrated observations, spatial coverage, cable effects, continuity, repair and accessible data
Subsea computation	Retrieved research-module demonstration; no general industry validation	Matched computing service, thermal field, energy source, retrieval, end-of-life, authorization
Offshore power and cables	Established sectoral installations; habitat-dependent consequences	Full power-chain materials, acoustic/EM exposure, seabed works, interactions and cumulative effects
Robotic service hubs	Partial component- and team-level demonstrations	Docking, contact-force reliability, maintenance substitution, faults, retrieval and habitat disturbance
Mineral recovery systems	Variable pilot evidence; major long-term ecological uncertainty	Plumes, benthic loss, waste, restoration uncertainty, jurisdiction and alternatives

24.6 Comparative infrastructure options and the limits of a single benefit metric

A disciplined prospective programme should compare *service-equivalent alternatives*. A proposal for a subsea computing node should be tested against a land-based or surface-hosted system meeting the same computing and observation requirements; a cabled station against an appropriately specified fleet of autonomous instruments and scheduled research cruises; a permanent seabed repair station against recoverable modular servicing; and an energy installation against different power and logistics arrangements. These alternatives vary across spatial scale, engineering maturity, available monitoring evidence, reversibility and legal scope. A claim that one option has lower cost per observation is insufficient if it omits environmental exposure or transfers scientific costs to another institution.

The comparison table is intentionally multidimensional. Monetary expenditure, carbon emissions, research data volume, animal exposure and habitat integrity cannot be assumed commensurable. Even within ecology, one cannot automatically replace a loss of a rare habitat with an increase in the density of a more common organism. A weighted score might be useful for a defined, publicly justified decision task, but its weights, normalizations, exclusion thresholds and representation limits require examination. Some environmental or legal constraints operate as independent conditions rather than fungible entries in a benefit–cost sum. In particular, legal authorization and scientific evidence about welfare should not be converted into artificial numerical bonuses that erase deficits elsewhere.

A more informative evaluation matrix retains separate columns for operational performance, scientific inference, material circulation, ecological exposure, distributive incidence, decision authority and remedial capacity. Rows should include deployment, routine operation, maintenance, abnormal conditions, recovery and eventual decommissioning. Comparisons must be spatially explicit: a cable repair can disturb an area far from a sensor; electric-power generation and fabrication may affect different regions; an offshore facility may create effects beyond the national jurisdiction in which it is located. Where counterfactual alternatives cannot be observed directly, the design should state which models or proxies support the estimate and how those assumptions might be challenged. The comparative methodology is developed more fully in Section 25; the present section specifies the objects and stakes to be compared.

Table 25: Ecological endpoints and material decisions requiring separate evidence in prospective oceanic infrastructure assessment. The examples are assessment questions, not universal impact scores.

Assessment dimension	Candidate observations	Inferential limit or comparison requirement
Direct substrate change	Seabed contact area, hard-substrate removal, burial depth, geomorphic footprint	Area alone does not identify organism or functional loss; match substrate and habitat
Suspended and deposited material	Time–depth concentration, settling, spatial deposition, background turbidity	Concentration is not a biological dose; compare calibrated exposed and reference areas
Ecological continuation	Taxon-specific recruitment, community composition, sediment processes, recovery trajectories	Abundance recovery cannot certify restored functions, welfare or earlier assemblage
Built material circulation	Component mass, coatings, leakage, maintenance and retrieval rate	Recoverability and lifecycle burdens need material-specific inventories
Decommissioning and succession	Original habitat, colonization, removed versus retained structure, long-term monitoring	New local diversity does not establish whole-system benefit or lawful retention

24.7 Material circulation, substrate conversion, and decommissioning comparisons

A lifecycle interpretation of oceanic infrastructure must follow material through manufacture, transport, deployment, maintenance, loss, retrieval and end-of-life handling, not merely count the components visible at a deployed site. Batteries, pressure housings, polymers, anticorrosion coatings, cable sheaths, sacrificial elements, buoyancy materials and electronic modules can follow distinct recovery and leakage pathways. A proposed design should record material composition and mass, known hazardous constituents, replacement frequencies, expected losses, retrieval feasibility and where recovered components are processed. These are prospective inventory categories; the present manuscript contains no measured lifecycle inventory for a particular product. Carbon emissions, raw-material use, marine exposure and end-of-life waste must remain distinct outcomes until a justified comparative method is available.

The material context may also be ecological substrate. Survey evidence in the Clarion–Clipperton Zone connects hard polymetallic nodules to particular sessile epifauna [203]. Removing that substrate is different from collecting a detachable sensor housing. Conversely, adding an artificial structure can support attached organisms while replacing soft-sediment habitat or changing local interactions. A global synthesis of offshore artificial structures examined ecological effect sizes across reference habitats and found an evidence base focused on abundance and diversity, with far fewer records of reproduction, behavior, and survival [91]. A North Sea lifecycle-context model of offshore wind-farm biodiversity likewise found results dependent on the starting seabed and decommissioning scenario, with uncertainty about ecosystem functioning [173]. Such analyses do not establish that an artificial reef produces a net benefit across every ecological and normative dimension; they make the reference habitat and outcome selection indispensable.

The decommissioning choice illustrates the potential conflict among research, preservation and material responsibilities. Full removal may recover manufactured materials while disturbing species that settled on a foundation. Partial retention may preserve an emergent assemblage but leave a long-term material and governance obligation. Continued operation may maintain valuable ecological observations while extending cable, maintenance and energy exposures. The alternatives should be compared against a specified counterfactual, with attention to biological succession, legal requirements, recoverability and the distribution of costs over locations and actors. None of the evidence cited here constitutes authorization to leave a structure in place.

A meaningful comparative ledger preserves differences in units and evidential quality. It should also represent explicitly unknown values rather than assigning them zero. Where a proposal is described as

“restorative” or “circular,” those terms should correspond to independently testable material and ecological results: a reduced uncontrolled material burden, demonstrated retrieval, a specified ecological endpoint, and a transparent comparison with reasonable alternatives. The present paper proposes this accounting architecture but has not evaluated any candidate technology against it.

24.8 A conditional model of heterogeneous ecological budgets

To make finite-resource reasoning testable without asserting a universal ecological unit, consider a finite family K of *specified assessment dimensions*. These might represent maximum permissible sediment disturbance in a defined study region, a provisional thermal-exposure envelope, an operational energy or vessel-fuel budget, a recoverable-material obligation, or a capacity constraint on maintenance operations. Their units and thresholds differ. Each is defined only relative to an identified measurement method, spatial boundary, time horizon and normative or engineering source. Other concerns, such as animal interests, a lack of legal authorization, or possible irreversible habitat change, may require separate qualitative or threshold review and are not necessarily included in K .

Let $\mathcal{A}$ be a set of prospective lifecycle plans, where a plan a specifies construction, operation, service and end-of-life actions over a declared interval $[0, T]$. For each $k \in K$, let $q_k(a, \omega)$ be the modeled cumulative burden under an admissible uncertainty scenario $\omega \in \Omega$, with consistent nonnegative units, and let b_k be a justified available budget for that dimension. Neither q_k nor b_k is assumed to represent all ecological relations. The robust screening set is

$$\mathcal{F}(\Omega, \mathbf{b}) = \left\{ a \in \mathcal{A} : \sup_{\omega \in \Omega} q_k(a, \omega) \leq b_k \text{ for every } k \in K \right\}. \quad (93)$$

A legal-authorization predicate and a separate welfare-review requirement may further restrict the actions considered, but satisfying Equation (93) alone establishes neither.

Proposition 24.1 (Monotonicity of robust heterogeneous-budget screening). *Suppose $\Omega_1 \subseteq \Omega_2$, and the budgets satisfy $b_{2,k} \leq b_{1,k}$ for every $k \in K$. If the candidate plans $\mathcal{A}$, their modeled burdens q_k , and the meaning and units of each assessment dimension are held fixed, then*

$$\mathcal{F}(\Omega_2, \mathbf{b}_2) \subseteq \mathcal{F}(\Omega_1, \mathbf{b}_1). \quad (94)$$

Proof. Take any $a \in \mathcal{F}(\Omega_2, \mathbf{b}_2)$. For each k , the supremum over the smaller scenario set cannot exceed the supremum over the larger one. Hence

$$\begin{aligned} \sup_{\omega \in \Omega_1} q_k(a, \omega) &\leq \sup_{\omega \in \Omega_2} q_k(a, \omega) \\ &\leq b_{2,k} \leq b_{1,k}. \end{aligned} \quad (95)$$

Thus $a \in \mathcal{F}(\Omega_1, \mathbf{b}_1)$, proving the inclusion. $\square$

The result is an elementary set-theoretic property, not a new ecological law or a normative claim that more conservative screening is always the best decision. Its purpose is to prevent an incoherent inference: when validated concerns enlarge the uncertainty set or reduce a justified burden allowance, the robustly screened plan set cannot become larger under unchanged models. But the assumptions are demanding. A habitat may lack any scientifically defensible scalar disturbance budget; measurements may fail to resolve rare organisms; interactive effects can change the meaning of q_k ; and scenarios outside Ω can invalidate the model. The inclusion also says nothing about whether the budgets are legitimate, whether risks are acceptable to affected communities, or whether a proposed plan should legally proceed. If the budgets represent morally incomparable phenomena, they should remain distinct constraints rather than being collapsed into a monetized scalar.

An operational dossier would record the grounds for each threshold, the method of deriving uncertainty envelopes, the consequences of a failed condition and the authority able to revise the plan. A

rejection should preserve evidence of which constraints failed; an approval should identify the scope, expiration, monitoring and recovery responsibilities of the decision. This procedure is consistent with Sections 8 and 10, while allowing future observations to change modeled constraints. Dynamic ecological contexts may require repeated screening and independently triggered review rather than a single once-for-all certificate.

24.9 Staged demonstrations, accountable withdrawal and research continuations

The prospective programme developed here has three interdependent research pathways rather than a single prescribed development trajectory. The first concerns *sustained global observation*: interoperable observatories, autonomous sampling, verified data provenance, public scientific capacity and ecological monitoring across spatial scales. The second concerns *ecology-aware robotic intervention*: reproducible manipulation, repairs, inspections, recoverable modular infrastructure and independently observed habitat consequences. The third concerns *ecologically accountable productive systems*: experimental energy, computation and service installations whose economic and technological possibilities are examined alongside their material limits, authority relations and decommissioning obligations. These pathways can proceed at different rates, and progress in one supplies no automatic authorization for another.

A staged programme should begin with explicit alternatives, bounded trials and measurement systems capable of detecting the kinds of harm an intervention might cause. Before installing a permanent facility, investigators can compare a shore-based instrument with a surface-hosted or temporary submerged prototype, test maintenance connectors in controlled environments, and measure cooling or acoustic disturbances at progressively more realistic spatial scales. Field experiments should register baseline conditions, monitoring blind spots, thresholds for reassessment, responsibilities for animal encounters, and recovery provisions. The programme should also ask what would falsify a claimed benefit. If a subsea processor performs adequately but produces unacceptably concentrated thermal exposure, that outcome should be reported as a limitation rather than absorbed by unrelated gains in computational efficiency.

Infrastructure evaluation should also account for the capacity to *withdraw*. An installation that cannot be safely repaired, disconnected, or retrieved creates a different moral and engineering situation from one whose operations and physical presence can be reconsidered. Reversibility is often incomplete: a retrieved cable or module cannot necessarily restore the previous habitat or erase prior biological exposure. Nevertheless, planned retrieval, modular replacement and financial or institutional responsibility for aftercare make correction more practically meaningful. Proposed abandonment must be evaluated separately from decommissioning; merely ending a research project does not end the material trajectory of its equipment. Ecological monitoring after removal can be necessary even when the immediate intervention has ceased.

Knowledge governance is part of this continuity. A multinational observatory or computing installation may preserve raw data while losing calibration files, instrument histories, maintenance records or the authority needed to correct interpretations. Common repositories should therefore support provenance, versioned methods, qualified access arrangements, alternative interpretations, and explicit procedures for recognizing scientific and ecological harms. Long-term data access and collective capacity are not identical; shared governance must attend to who can construct instruments, inspect source code, allocate computational resources, challenge site selection and exercise remedial authority. These concerns carry forward the common-value framework of Section 9 without treating open publication as the totality of generative justice.

Section 25 will specify comparative research designs and validation across platform capability, inference, ecological outcomes and institutionally authorized interventions. It will distinguish what can presently be tested in tank, laboratory, controlled field trial, sustained observatory and real infrastructure settings. Section 26 will synthesize the wider theoretical stakes and their unresolved conditions. The central conclusion of this section is deliberately conditional: oceanic infrastructures may enlarge collective ecological inquiry and enable forms of productive activity that are difficult to realize on land, yet such

possibilities are situated within finite material, ecological and institutional capacities. Responsible engineering requires treating those conditions as objects of observation, contestation and revision throughout the infrastructure's trajectory.

25 Comparative Research Designs and Validation

25.1 Validation across heterogeneous claims and relational systems

The preceding sections have proposed an interconnected programme of ecological inquiry through oceanic robotic systems. Its claims, however, are different in kind. A vehicle may maintain station under a declared current regime; a sensor may reproduce a calibrated concentration measurement; an algorithm may improve the prediction of an ocean front; a biological interpretation may infer an organism's presence; a governance procedure may provide an opportunity to challenge a deployment decision. Success in one register cannot establish success in the others. The purpose of this section is to specify comparative research designs capable of separating these claims, making their evidential dependencies inspectable, and identifying observations that would weaken or overturn the proposed programme. This is a prospective methodology, not a report of completed trials.

The generative-relational (GR) framework offers a reason to conduct such comparisons without presupposing a unified evaluator. As developed in Sections 2 and 6, the capacity to observe a relational system is constrained by instruments, interpretations, histories and the positions of observers. Moreover, Section 9 distinguished access to information from the productive capacity to reinterpret and extend it. Validation therefore concerns the specific relationship between a research claim, a method of producing evidence, the relevant system and the institution that can assess or contest the finding. It is neither a substitute for ecological justice nor a purely social negotiation independent of material observations.

A *claim dossier* should consequently precede each experiment. It should record the object of inference, a precise operational claim, the physical or institutional scope, uncertainty assumptions, comparison conditions, observable consequences, alternative explanations, and a criterion that would count against the claim. For a propeller-powered autonomous vehicle, for example, the claim could concern current-relative tracking within a stated range of flow speeds, sensor errors, actuator faults and battery capacities. A separate claim would address whether its acoustic transmissions cause measurable behavioral changes in nearby fish. The first requires vehicle telemetry and reference navigation; the second requires a biologically defensible design that accounts for background variation and observation effects. Neither alone answers whether a proposed mission is legally authorized.

Not every philosophical proposition in this paper can be experimentally confirmed through a single operational indicator. Relational descriptions, ontological commitments and arguments about justice should be judged through clarity, internal consistency, explanatory fertility, comparison with alternatives and their capacity to disclose empirically consequential distinctions. In particular, the proposed distinctions among ontological, relational and manifestational registers should be tested for analytic usefulness and revised when they obscure the mechanisms under inquiry. A failure to falsify a particular robot performance hypothesis neither proves the underlying GR ontology nor establishes the truth of a moral theory. Conversely, a physical result that conflicts with a forecast may expose a specific model or measurement defect without refuting all forms of relational interpretation.

We therefore propose a *typed evidence architecture*: engineering performance, observation quality, ecological inference, welfare-relevant exposure, institutional procedure and heterogeneous value circulation each retain their own units, review criteria and competent evaluators. Cross-register comparisons are possible through explicit connecting assumptions, such as a validated relationship between acoustic source level and an exposure estimate or between an operational schedule and an environmental assessment. A connecting assumption must itself be documented and challengeable. The remainder of this section develops comparative designs for each register, then a limited mathematical account of transferring empirical performance claims between contexts.

The central cross-layer hypothesis can now be framed as a comparative research question: does

preserving the observation policy, physical-intervention history, interpretation lineage, and institutional constraints produce better-calibrated or more corrigible ecological claims than a workflow that records only processed sensor outputs? To make this test discriminating, an experimental comparison must avoid attributing gains to a new GR label when they arise from familiar, already effective practices of calibration, experimental design, or provenance. One baseline should use the best applicable conventional scientific workflow, including metadata standards, versioned code and ordinary quality control. The proposed extension should be evaluated for an *incremental* result, such as improved recovery of source errors, detection of incompatible evidential supports, warranted uncertainty coverage, or the practical ability of an independent reviewer to reconstruct a disputed inference.

The appropriate unit of comparison depends on the claim. For a measurable inference outcome, preregistered synthetic or controlled cases could assign the same data-generating process to several analysis pipelines while varying whether action records and observation-selection histories are visible to analysts. Calibration, false attribution rates, missed ecological responses, and response to intentionally injected metadata errors can then be evaluated. For an institutional outcome, matched decision vignettes could vary evidence accessibility, competent review pathways, and opportunities for correction; evaluators would document both substantive decisions and reasons for disagreement. These are prospective designs, not completed comparisons. Researchers should report implementation cost and document burden, because additional procedural complexity may sometimes impair rather than improve accountable inquiry.

A failure condition follows directly. If full provenance and cross-context mediation do not outperform a well-implemented conventional workflow on the declared tasks, the paper has no warrant to claim a practical advantage for the additional GR representation. This would not by itself refute a philosophical account of relations, but it would constrain the proposed engineering and institutional value of that account. Conversely, favorable performance on one task cannot establish a universal interpretive infrastructure: the outcome must be reported within its validated population, environment and decision context.

25.2 A staged experimental ladder and a comparison matrix

The first methodological choice concerns *where evidence is produced*. Simulation and physical testing offer complementary opportunities, but their relationships cannot be represented by one universal technological readiness level. Simulation enables controlled variation of disturbances, sensor failures, animal-response hypotheses and resource budgets; its outcomes remain conditional on model structure. Bench and pressure-chamber experiments expose particular components to measurable loads. Tank and flume tests investigate coupled motion and flow under designed boundary conditions. Coastal trials introduce environmental variation and operational constraints, while deeper-water missions add pressure, communications delay, vessel logistics and recovery risks. Sustained observatory operations may reveal seasonal drift, fouling and maintenance failures, but can be expensive and institutionally dependent. Research progression should be set by the claim and risk, not a presumption that larger deployments are always more scientifically informative.

Table 26 provides a minimum cross-domain matrix rather than a catalogue of validated achievements. Before testing, a research team should state its intervention and comparison designs, primary outcomes, sample or mission allocation rules, relevant covariates, decision thresholds, exclusions and adverse-event reporting rules. Where randomization is feasible and ethically acceptable, it helps reduce confounding; where it is impractical or prohibited, investigators should describe the assumptions needed for matched, observational or quasi-experimental comparisons. Calibration standards and negative controls differ between a dissolved-oxygen probe, a navigation stack, a manipulation actuator and a biological survey. Uniform documentation is useful; uniform measurement is generally inappropriate.

We distinguish *repeatability*, under closely matched apparatus and operators, from *reproducibility* across changed instruments, analysts or operating conditions, and from *transportability* to a distinct environmental regime. Terminology can vary across disciplines, so each protocol should define its usage.

Table 26: Illustrative comparative validation designs. Evidence levels are task-relative; no row alone establishes ecological acceptability or authorization.

Claim family	Comparator and observation	Principal validity threat
Motion and manipulation	Matched trajectories or intervention tasks; reference pose, force, energy and recovery records	Currents, fixture effects, hidden human intervention, contact mismatch
Sensor measurement	Traceable standards, co-located instruments, blanks and field calibration	Drift, biofouling, sampling disturbance and uncertain measurand
Adaptive ecological inference	Predeclared sampling domain; policy log, fixed-route and randomized reference observations	Preferential selection, correlated errors, inaccessible habitats
Animal and habitat exposure	Replicated spatial and temporal controls; acoustic, behavioral and substrate measures	Natural variability, weak detectability, cumulative effects
Infrastructure and governance	Alternative lifecycle plans; authorization, budgets, audit and contestation records	Omitted supply chains, jurisdictional scope, unequal review capacity

A tank-replicated controller may be repeatable without being transportable to stratified coastal waters. Independent data review can reproduce an inferential computation while leaving the original animal observation impossible to rerun. A precommitted protocol should record what was deliberately held constant and what variation is part of the target population of contexts.

Ocean observing already possesses partial infrastructure for such discipline. GOOS Essential Ocean Variables are selected around sustained, feasible measurements and scientific or societal needs, while the Ocean Best Practices programme promotes methodological interoperability [54, 55]. The U.S. IOOS QARTOD effort supplies quality-control practices for specified observational streams, and the Ocean Observatories Initiative combines automated checks with instrument provenance and human review [180, 148]. These approaches provide operational exemplars rather than interchangeable certification procedures for animal welfare, robot construction or legal legitimacy. A research protocol should cite the applicable standard and state explicitly when no established standard is available.

25.3 Engineering performance, physical modeling and system identification

Engineering validation should begin with a comparison against a reference that is simpler, established or otherwise sufficiently understood. Section 14 proposed tangent-space governance and criticality-aware supervision as a synthesis of established local identification, robust and adaptive control, and operational constraint review. The evaluation question is whether this synthesis improves predeclared outcomes under specified disturbances compared with matched classical feedback, receding-horizon control, a learned policy, and an appropriate human-supervised baseline. An ablation must also compare the same controller with and without local-model revision, with and without the supervisory component, and with explicit fallback behavior. Otherwise, apparent gains may result from access to additional sensors, computational resources or operator attention.

A reproducible robot study should identify the plant and sensing configuration: hull or body morphology, degrees of freedom, actuation and transmission details, inertial parameters, hydrodynamic coefficient estimates, navigation references, communications settings and software version. Trials should distinguish commanded pose from measured pose, and pose measured by the robot from an independent reference. The definitions of current-relative speed, position error, contact force, power consumed and mission completion need to be stable across competitors. Contact-rich tasks should report fixture geometry and material compliance, since successful docking against a forgiving laboratory fixture can overstate connector-handling ability in changing flows. Navigation and manipulation results should not be combined into a single success rate without explaining task composition and exposure time.

Robustness requires experiments at the limits of stated model validity. A factorial or otherwise effi-

ciently blocked design might cross several current regimes with sensor dropout, localization errors, degraded actuators, intermittent communication and different morphological configurations. High-pressure and temperature effects may require separate facilities before combined trials become appropriate. Test cases should include ordinary disturbances and deliberately identified adverse regimes, without deploying prototype systems into an ecologically sensitive area merely to generate difficult data. For learned controllers, training, tuning and evaluation environments must be separated, and performance should be reported over seeds, independent missions and uncertainty intervals rather than a selected exemplary trajectory.

Two specific failure criteria follow from earlier formal sections. The local approximation in Section 13 is credible only where the empirical residual remains compatible with its claimed validity envelope. The finite-horizon error bound in Section 14 must be evaluated using documented Lipschitz or disturbance assumptions rather than treated as a guaranteed engineering margin. If a controller consistently requires human intervention in an unreported fraction of missions, or the supposedly conservative supervisor fails to detect an identified mode change within the allowed time, the deployment-level claim should be rejected or narrowed. A technical fallback such as holding position must also be tested for feasibility in the actual flow and energy regime: holding may be less recoverable than a controlled retreat.

Measurement uncertainty is not resolved by repeated numeric output. Reference sensors and load cells can be biased, time synchronization may be inaccurate, and environmental reference measurements may be sparse. Consistent uncertainty statements should identify measurands, calibration pathways, systematic and random contributions, and covariance where material. The GUM family provides a general approach to expression of measurement uncertainty, while NIST describes the dependence of uncertainty claims on defined measurement models [77, 121]. These sources concern metrology; translating their principles to biological or normative interpretations requires further assumptions and independent evidence.

25.4 Ecological response, welfare evidence and causal comparison

Ecological validation requires a design capable of separating environmental change associated with the intervention from the natural variation of marine habitats. Robot presence, illumination, sonar, thruster wash, sediment contact, heat release and artificial structures can influence animals and substrate in different ways. No single proxy captures welfare, habitat condition and ecological generativity. A disturbed animal may alter its trajectory without demonstrable injury; an unchanged population count may conceal chronic stress or a change in recruitment; an altered acoustic recording may arise from microphone orientation or changing background vessel traffic. Accordingly, a proposal should list the ecological mechanism hypothesized, plausible competing mechanisms, and the observable evidence appropriate to each.

For site-specific ecological disturbances, before–after–control–impact (BACI) designs are one useful family. A stronger study typically considers multiple comparison locations and replicated periods rather than assuming one control site represents a valid counterfactual. Underwood’s work on *beyond-BACI* designs explains the problems posed by natural temporal variation and inadequate spatial replication [184, 185]. A subsea cable installation, for example, could be studied through seabed community measures at several matched sites before disturbance and during subsequent monitoring, with explicit limits on what habitat conditions, seasons and spatial scales the study represents. If an intervention cannot ethically or legally be randomized, a prospective observational design and sensitivity analysis may still improve inference, though it cannot eliminate every alternative explanation.

The relevant counterfactual may be an alternative operational design rather than complete absence of inquiry. Passive listening could replace a high-output acoustic survey for a narrow question; remote imaging may reduce manipulator contact; a shore-based computation system may avoid local thermal discharge while imposing other electrical and land-based burdens. Researchers should therefore evaluate candidate alternatives under a consistent defined lifecycle boundary and document where different data sources or uncertainty structures make direct comparisons weak. For animal encounters, include behavior

before approach, during operation and after withdrawal when feasible, as well as exposure conditions and detectability. The study must distinguish a measured response from an attribution of subjective pain or suffering. Such attribution may require convergent comparative biological evidence beyond the mission itself, as discussed in Sections 4 and 22.

Negative results deserve careful treatment. If repeated surveys do not observe a rare animal, Section 15 demonstrates how conditional detectability determines the evidential interpretation. For site impact, failure to detect a difference does not establish equivalence when the design lacks power to resolve a biologically meaningful effect. A credible protocol should specify a smallest consequential effect or interval of ecological concern, the expected variance, sampling power or precision, and a policy for reporting inconclusive findings. Cumulative impacts may arise over years or across installations and cannot generally be resolved through a single short field trial. Ecological baseline selection itself can favor particular states of the environment, so study design should preserve alternative ecological histories and uncertainties.

Precaution is not a mechanical consequence of a p -value. Institutional permission, animal welfare and environmental review involve independent standards and competent bodies. An experiment showing an insignificant average response cannot authorize a prohibited activity; conversely, a legal permit does not make uncertain ecological impacts scientifically understood. For studies involving living organisms, an appropriate independent ethical review must assess the necessity of the intervention and possibilities for less disturbing evidence production. These are components of the research design, not deductions from technical performance.

25.5 Adaptive sampling, model comparison and inferential stability

The adaptive observation programme of Section 21 raises a different validation challenge. When a robot selects its next measurement location using its current model, the resulting dataset depends on the model's previous beliefs and the scientific objective. An apparent improvement in finding a biological concentration or physical front may reflect effective targeting while simultaneously creating biased coverage of the broader target domain. The preferential-sampling literature formalizes such dependencies and demonstrates that ignoring the sampling process may distort spatial inference [32]. Thus research quality should be judged both by mission-directed discovery and by independently assessable properties of the inferred ecological field.

A suitable experiment compares at least three observation regimes where feasible: a predeclared fixed survey, a spatially randomized or probability-based reference design, and one or more adaptive policies. The adaptive policy's selection probabilities, sensors, map updates, mission constraints and human overrides should be logged. Comparisons must specify which objective is being optimized: detection of a localized event, regional estimation error, calibrated probabilistic forecasts, temporal coverage, data volume per joule, or the probability of disturbing organisms. A policy may dominate on target detection while performing worse on population-scale estimation. One result should not be substituted for the other. If the reference domain is physically inaccessible at many locations, the design must explicitly narrow the inferential population or state the untestable extrapolation assumptions.

Model comparisons should include independent holdout observations that were not used in sample selection, parameter fitting or post hoc tuning. Where spatial and temporal correlations are substantial, randomly withholding individual adjacent pixels or time points can create misleading independence; spatial-blocked, temporally separated or external-platform evaluation may be more informative. For probabilistic models, calibration, sharpness and appropriate scoring should accompany point accuracy. A particularly important failure case is an adaptive sampling system that converges on a seemingly stable pattern only because it repeatedly samples locations favored by its current forecast. An independently sampled disagreement is valuable evidence for model revision, not noise to discard.

Observation records need scientific provenance sufficient for reanalysis. The FAIR principles support discoverable and reusable scientific data, while the CARE principles address collective benefit, authority, responsibility and ethics in relevant Indigenous data contexts [214, 22]. The two have different

research objects and should not be conflated. A useful marine archive would connect raw and calibrated records with instrument versions, geographic and depth references, timing, quality flags, model versions, sample-selection policy, access limitations and corrections. Data governance must recognize that sensitive ecological locations, biological material, traditional knowledge and contractual datasets may warrant carefully differentiated access rather than universal public release.

The proposed GR framework also needs discriminating tests. Compare whether recording relations among observer, instrument, context and interpretation actually improves detection of model mismatch, cross-platform integration or the correction of scientific claims relative to conventional well-documented metadata and provenance systems. A claimed epistemic advantage should be narrowed if established methods deliver the same outcomes with less complexity. An open evaluation should invite independent domain specialists to identify representation losses and articulate what the framework fails to capture. Claims about generalized translation require task-specific preservation tests with stated source and target domains, not the presumption that every biological or institutional meaning admits a common numerical coding.

An independent test of ecological observation bias. The experimental ladder should distinguish calibration against static targets from comparison of sampling living organisms that respond to the platform. An instrument can be well calibrated and a vehicle can execute its route accurately while an unknown fraction of the target population leaves the observation region before detection. Empirical work with acoustic and camera-equipped robots demonstrates that responses can include attraction, avoidance and different effects of light and platform at different depths [17]. A serious ecological benchmark should therefore compare an explicitly defined target estimand with platform-conditional estimates, report background detectability, exposure conditions and plausible response mechanisms, and quantify sensitivity to pre-detection departure. A result showing no statistically detectable response under a particular sample size is narrower than evidence of negligible biological effects across missions. For welfare-related evaluation, the causal chain from exposure through response and possible physiological or demographic consequences must be examined on its own timescales and with relevant endpoints [156, 107].

25.6 Institutional procedures, heterogeneous value and comparative accountability

Institutional evaluation asks different questions from field engineering. The legal competence of an authority, the scope of a research consent, access to monitoring records, opportunities for contestation and the effectiveness of remedies are not inferred from navigation precision. Sections 7, 8 and 10 developed these as distinct institutional objects. A comparative protocol should therefore include a jurisdiction-specific legal review performed by qualified experts, the actual decision instruments, affected activities, reasons supplied, environmental conditions, routes for challenge, and provisions for correction or withdrawal. A mock procedure can test a workflow's clarity; it cannot demonstrate the enforceability of a rule in a real legal system.

The hypothetical transnational A–B–C–D oceanic facility offers a research case that can be varied without attributing misconduct to identifiable actors. One organization supplies capital, another robotics and scientific methods, a third runs the installation and a coastal authority regulates applicable activities. Compare alternative governance and value-circulation arrangements: proprietary data with shared calibration protocols; open data supported by constrained computing access; jointly governed research infrastructure; and commercial services with independent ecological monitoring. Record, separately, investment, material inputs, operating energy, ecological exposure, scientific outputs, rights to reuse data, technical participation, responsibilities for maintenance and remedial capacity. Such a typed incidence account follows Section 9. Aggregating these relations into one monetary balance would conceal distributional and normative differences.

Several criteria are independently assessable: whether a monitoring record exists, whether an environmental condition is measurable, whether a challenge was acknowledged, whether an institution issued reasons, whether correction was executed, and whether affected participants had practical access

to evidence. Other criteria require interpretation: whether participation was meaningful, whether a distribution was just, or whether an animal interest was adequately represented. Evaluators should state who supplies these judgments and how dissenting interpretations are recorded. The procedure must avoid treating an organization's self-assessment as independent verification. A comparison can be improved by independent observers, anonymized hypothetical cases, multiple decision scenarios and structured disagreement reporting, while still acknowledging that simulations lack the pressures and consequences of actual adjudication.

For international comparison, do not treat marine zones, scientific research permissions, installations and mineral-resource activities as legally interchangeable. A field protocol should preserve location, water-column versus seabed jurisdiction, flag-state and coastal-state relations where applicable, and the specific treaty and domestic instruments actually relevant. The BBNJ Agreement's mechanisms cannot be generalized into a universal license, revenue-sharing tax or enforcement authority for any subsea computing facility. Normative hypotheses about shared benefits or generative justice are appropriate objects of philosophical and institutional debate, but they require their own reasons and cannot be certified by a favorable technical experiment.

25.7 A bounded result on transferring empirical validation across contexts

The comparison programme requires an explicit limitation on the transport of laboratory and field results. A controller may be accurate in one tank and fail in a different current regime; a biological classifier may perform well near a fixed sensor but misinterpret images in turbid water. The abstract result below formalizes a minimal version of this problem without presuming that ecological systems have fixed global probability laws. It is intended as an assessment device for identified evaluation distributions.

Let $\mathcal{Z}$ be a measurable space of test instances, each including the relevant vehicle, environment, mission and sensor conditions. Let P be the evaluation distribution represented by a validated test campaign and Q a proposed target distribution, both defined for a *specified claim*. Let the predeclared loss $\ell : \mathcal{Z} \rightarrow [0, 1]$ have a fixed meaning across both contexts, and let $R_P = \mathbb{E}_P[\ell]$ and $R_Q = \mathbb{E}_Q[\ell]$. Write $d_{TV}(P, Q) = \sup_{A \subseteq \mathcal{Z}} |P(A) - Q(A)|$, where the supremum ranges over measurable sets.

Proposition 25.1 (Bounded transfer of an evaluation loss). *If P and Q are probability distributions on the same measurable test space and the same measurable loss ℓ takes values in $[0, 1]$, then*

$$|R_Q - R_P| \leq d_{TV}(P, Q). \quad (96)$$

Consequently, a justified upper bound $d_{TV}(P, Q) \leq \delta$ entails $R_Q \leq R_P + \delta$.

Proof. For every $z \in \mathcal{Z}$, the layer-cake identity gives $\ell(z) = \int_0^1 \mathbf{1}\{\ell(z) > s\} ds$. Boundedness permits integration under either probability measure. Thus

$$\begin{aligned} R_Q - R_P &= \int_0^1 [Q(\ell > s) - P(\ell > s)] ds, \\ |R_Q - R_P| &\leq \int_0^1 d_{TV}(P, Q) ds = d_{TV}(P, Q). \end{aligned} \quad (97)$$

The stated one-sided consequence follows immediately. $\square$

The proposition is a standard implication of total-variation distance, not a new theory of ocean dynamics. Its contribution here is an explicit discipline for *validation scope*. A favorable result for P cannot be transported to Q without evidence about the distributional shift, a justified shared loss and appropriate treatment of measurement uncertainty. In many oceanic settings, no defensible total-variation bound can be estimated, especially when target regimes are rare, biologically evolving or absent from a trial. The result then supplies no usable guarantee. Even when a statistical bound holds, it describes a chosen loss distribution, not the entirety of ecological relations, unknown failure classes, human responsibility or legal authorization.

A practical transfer dossier should identify which changes are being evaluated: depth, currents, turbidity, pressure, morphology, biofouling, animal communities, mission duration, institutional setting, and intervention intensity. It should distinguish covariate shifts from changes in the physical or biological mechanisms that give a loss its meaning. For an ecological claim, the target population of organisms and exposure pathways may change with location, so using the same numerical loss can itself require justification. The result also does not apply without modification to strongly history-dependent trajectories treated as independent test cases. A validation team may use blocked evaluation, hierarchical modeling, controlled stress tests or other methods suited to the actual dependence structure; the formal inequality is not a substitute for such design.

25.8 A comparative protocol for provenance and institutional mediation

The proposed GR record architecture should be evaluated against already substantial practices rather than against an artificially impoverished baseline. FAIR concerns findability, accessibility under defined conditions, interoperability, and reuse of scientific data [214]. The W3C PROV ontology offers a standard account of entities, activities, agents, and derivation relations [220]. The CARE principles concern collective benefit, authority to control, responsibility, and ethics specifically in Indigenous data governance, and their origin and normative object should remain explicit [22]. These contributions are neither interchangeable nor exhausted by a generic metadata checklist. Where marine observations involve Indigenous data or knowledge, the affected communities' relevant authority and governance processes must be evaluated on their own terms, rather than represented as a property generated by an additional software field.

A predeclared comparison could use three progressively equipped implementations over the same synthetic and, where permitted, real research dossiers. The *competent conventional baseline* would include calibrated observations, versioned code, ordinary scientific metadata and quality flags, access rules, and an auditable interpretation record. The *standards-oriented baseline* would additionally implement suitable FAIR data stewardship, PROV-derived lineage, and any contextually applicable governance processes, including CARE where relevant. The *proposed extended record* would retain those capabilities and encode explicit observation-policy transitions, intervention histories, cross-system interpretive disputes, links to reviewable authorizations, and status changes in disputed claims. These arms should be implemented by comparably trained teams, given comparable computing and annotation budgets, and tested with blinded reviewers. The last arm cannot be credited for improvements that its preceding controls already provide.

Evaluation should include tasks that can fail visibly. For example, independently trained assessors could be asked to identify the exact raw observations supporting a contested species-occurrence claim, reconstruct whether an adaptive sampling rule selectively omitted negative observations, detect a change in calibration between deployments, or locate the authority and conditions for collecting a particular biological sample. Prospective tasks could test time to reconstruct the claim, error detection, evidence coverage, justified uncertainty, reviewer agreement, correction latency, successful reanalysis, and the number of unresolved authority questions. Record-creation time, information disclosure risks, storage demand, specialist participation and the effect of unavailable historical metadata are costs that must be measured as well. Synthetic errors should be specified and hidden from evaluators before testing, and real-case findings should be adjudicated by an independent reference process where practicable.

The causal comparison would isolate at least three candidate mechanisms: additional *data captured*, improved *representation* of information already captured, and changed *review procedure* or institutional participation. Ablation of each feature can show whether a claimed advantage comes from storing intervention history, making provenance queries easier, or simply requiring more reviewer attention. A representation that increases clerical burden without reliable improvement in predeclared scientific or corrective outcomes is a meaningful negative result. If the standards-oriented baseline already achieves equivalent reconstruction and correction, the appropriate response is to narrow the GR proposal to theoretical interpretation or genuinely unserved integration cases. A favorable result in any arm remains

Table 27: Formal result families, the information needed to apply them, and possible invalidating observations. This is a verification checklist rather than independent certification.

Formal family	Operational premise needing evidence	Failure or limitation to record
Fiber factorization and information loss (§§2, 6, 9)	Declared estimand, observation map, representation metric, admissible histories	Distinct target states remain indistinguishable; translation domain is absent
Robust admissibility (§§3, 10, 17, 24)	Fixed constraint semantics, compatible uncertainty sets, known actuator/impact dimensions	Changed scope of permissions or unmodeled effect channel breaks the represented screen
Portfolio and institutional logic (§§8, 12, 20)	Fixed feasibility relation and authority scope; genuine data-access limits	Coverage expansion disappears under new costs or infeasible cooperation
Mechanical and material bounds (§§13, 14, 18, 19, 23)	Coordinate maps, scaled norms, calibrated derivatives, load and energy envelopes, domain coverage	Residuals, estimation errors or faults exceed the certified premise
Observational inference (§§15, 16, 21)	Conditional detection rates, sensor geometry, positive known sampling probabilities, stable reference estimand	Dependent miss rates exceed bounds; unreachable units or intervention bias the target
Welfare and evidence transfer (§§22, 25)	Receptor-specific exposure metrics, selected loss, same evaluation space, target-domain evidence	Screening surrogate disagrees with outcomes; distributional shift is unbounded

bounded to the evaluated records, actors, access conditions, and institutional setting; it neither generates legal permissions nor independently establishes ecological welfare.

25.9 Formal verification, dimensional consistency, and operational identification

The twenty numbered propositions distributed through this manuscript serve as conditional implications, not empirical confirmations of their premises. A formal verification pass therefore asks three successive questions: whether a conclusion follows from its written hypotheses; whether the symbols and transformations describe physically coherent quantities; and whether the premises can be measured, bounded or otherwise justified in an intended application. Correctness at the first level does not imply success at the latter two levels. The distinction is particularly consequential when a familiar inequality, such as a Taylor remainder, a total-variation bound or an inverse-probability expectation, is used alongside an exploratory ecological interpretation.

In the purely representational results of Sections 2–10, the mathematical objects are task-specific sets, maps, fibers, and uncertainty relations. Those conclusions follow by factorization, inclusions, or logical incompatibility provided the sets and mappings are declared consistently. They can establish when information is inadequate for a stated inference, or when a conservative action screen shrinks as modeled uncertainty expands. They cannot establish which biological states or legally valid permissions those sets contain. A falsification plan should therefore test both the adequacy of a selected representation and any claimed operational benefit over less elaborate conventional descriptions.

In the mechanical results, compatibility of reference frames, physical dimensions, and validity domains is indispensable. The interface velocities of Section 13 are compared at corresponding spatial and material positions, with stress measures kept in their respective configurations. Local Taylor estimates require scaled coordinates when positions, forces, control inputs, and environmental measurements are joined in one vector. Control-barrier inequalities require consistent derivative units, sound uncertainty bounds, and feasibility over the reachable states. Actuator damage and energy-reserve inequalities require a calibrated damage proxy, energy and power accounting in consistent time units, and physically defensible envelopes. Even an algebraically valid finite-horizon estimate may be useless if its constants can only be made so conservative that no relevant action remains feasible.

A worked inferential distinction is especially useful for adaptive observation. A finite-population estimator may be conditionally unbiased when every unit has a correctly logged nonzero draw probability, but a biological survey that excludes habitat under a binding no-entry condition no longer samples that entire reference population. It cannot repair the missing observations by asserting a tiny numerical probability for prohibited locations. Nor does unbiasedness imply useful precision: inclusion probabilities approaching zero inflate inverse-probability contributions, potentially producing an unstable estimate

even when its expectation is correct. The estimator should be evaluated through variance, effective coverage, actual measurement errors and the behavior of its scientific estimand under observation-induced changes. Section 21 makes these assumptions explicit.

The same logic applies to claims about organisms. A unit-consistent cumulative exposure surrogate can be bounded by calibrated envelopes. It remains a specified mathematical functional of physical variables, and its scalar value cannot replace taxon-specific welfare evidence. When acoustic, visual and sediment effects have no empirically defended common scale, the appropriate representation may consist of separate exposure limits and a record of conflicting observations. For validation transfer, total-variation distance gives a correct risk difference inequality on a shared measurable test space; it does not supply the missing bound or ensure that an ecological outcome is measurable in the same way at two sites.

Thus a formal model should enter an empirical study with an explicit *verification dossier*: typed variables and SI or dimensionless units, domains and physical frames, differentiability or probabilistic regularity conditions, calibration data, identifiability assumptions, uncertainty and failure budgets, pre-declared invalidation tests, and a responsible process for revising the model. Where assumptions are not testable because observations are unavailable, they should be reported as conjectural restrictions rather than hidden inside an apparently definitive inequality. The broader audit is recorded in the editable release alongside the source, and should remain open to independent mathematics, robotics, ecology and statistics review.

25.10 Comparative tests for habitat recovery and cumulative material burden

A habitat-sensitive research design would connect independent observations of physical exposure to organism assemblages and subsequent ecological processes. For a prospective, appropriately authorized seabed-device comparison, the primary contrasts could include a noncontact imaging trajectory, a recoverable low-contact configuration and a conventional research-platform baseline. The protocols would record pre-intervention substrate, microhabitat and visible fauna; calibrate sensor detection under disturbance and background conditions; and use multiple comparable reference sites or explicitly matched spatial blocks. The proposed design makes no assertion that these treatments have been field-tested in the present work. Where active intervention risks foreseeable harm, observational comparisons or controlled nonbiological trials should precede field manipulation, and the project may be unsuitable for intervention altogether.

The follow-up period should be chosen with the biological outcome in view. Immediate turbidity and animal movement, seasonal colonization, and slow recruitment of long-lived assemblages require distinct intervals. Evidence from mining-track revisits after 26 or 44 years supports the importance of long horizons for some intensive deep-sea disturbances [82, 83], while seamount coral recovery observations show why partial recovery should remain visible rather than being coded automatically as failure or success [16]. The durations are not transferable minimum-monitoring rules for every marine robot. Instead, they identify the need to justify a time horizon rather than assume it from a project's funding cycle.

A preregistered causal comparison should define separate estimands for physical disturbance, detection-corrected abundance, community composition, specified habitat function and lifecycle material burden. Geographic and temporal clustering, seasonal variation, habitat heterogeneity, exposure measurement error and missing observations must enter uncertainty analysis. In particular, an apparent post-deployment increase in filmed organisms may reflect attraction to new structure, a change in the camera's detection function, or a shift from one assemblage to another. Recolonization measured against a control assemblage also differs from the normative judgment that retaining an artificial structure is desirable.

A decision-focused result would present joint but noncommensurate findings, including nulls and adverse outcomes, rather than optimize a single invented "ecological generativity score." A design might outperform a comparator in successful sampling and material recovery while underperforming on mea-

Table 28: Suggested comparative tests of legal and normative information support; none reports an observed result.

Test object	Adversarial case	Observable failure or success criterion
Jurisdiction classification	High-seas water above an extended shelf, or equipment at different depths	Missed vertical distinction; unsupported rule citation
Consent and permission	Conditioned consent, changed activity, or Article 252 objections	False implication that an unreviewed operation is authorized
Assessment procedure	BBNJ screening below full-EIA threshold with subsequent concerns	Confusion of screening, public record, EIA and authorization
Source hierarchy	Regional access right applied outside its parties or MASS guidance treated as statute	Unsupported binding-force assertion
Representation and remedy	Animal-response evidence with uncertain welfare inference and inaccessible complaints	Omitted uncertainty, participation route or review authority
Value circulation	MGR-linked sequence information versus ordinary sensor telemetry	Invalid generalization of specialized benefit-sharing rules

surable disturbance or revealing insufficient power to estimate a vulnerable population's response. Reporting that conflict is an intended result, not a defect. A proposed GR provenance system offers a testable improvement only if its records change the traceability, reviewability or scientific conclusions relative to well-operated conventional baselines at justifiable overhead. External ecological and legal reviewers must be able to challenge the counterfactual, exposure measurements and meaning of any asserted recovery.

25.11 Institutional validity tests for jurisdictional and ethical claims

An evidentially disciplined evaluation of the proposed normative infrastructure should ask whether it improves reasoning about applicable rules and review pathways without claiming to automate legal judgment. Construct a benchmark of hypothetical marine activities with independently reviewed fact patterns that vary the water-column and seabed location, scientific purpose, infrastructure permanence, biological sampling, flag-state relation, party status and uncertainty about affected organisms. Ground the reference decisions in the identified legal instruments and explicit assumptions; where specialists disagree on classification, preserve the disagreement rather than impose an artificial gold-standard answer. Competing workflows should include trained specialists using ordinary document management and established scientific research records, alongside a provenance-augmented model.

Define a false-clearance rate only against cases in which the reference review has established that the proposed action is outside a specific permission; exclude genuinely disputed outcomes from that denominator and report them separately. Other metrics include the proportion of assertions linked to an accurate primary source, completeness of identified competent authorities, detection of material plan changes, quality of uncertainty disclosure, and the time needed to correct an erroneous record. Report confidence intervals and reviewer disagreement. A design is weakened if it reduces visible uncertainty by silently inventing authorization, and may be rejected if conventional workflows equal its performance without additional burdens. Specialist legal review remains necessary even if information-support metrics improve.

25.12 Manuscript-wide claim dependencies and limits of transfer

A manuscript spanning ecological philosophy, robotic engineering, statistical inference and maritime institutions presents a special risk: a result that is warranted within one field may be reused in another

Table 29: Cross-sectional claim-dependency audit. Each bridge is proposed for verification; no row reports a completed experiment or institutional approval.

Proposed inference	Additional warrant required	Invalid shortcut to exclude
Instrument signal $\rightarrow$ ecological occurrence	Detection and calibration model, defined target, independently checked observations	Treating classification confidence as ground truth or population abundance
Selected observations $\rightarrow$ target-domain estimate	Declared estimand, policy support, measurement errors, detection and temporal assumptions	Assuming adaptive selection is ignorable because its route is logged
Observed response $\rightarrow$ organism welfare	Taxon-specific exposure and behavioral evidence, rival explanations, welfare construct	Equating avoidance or attraction with a measured subjective state
Engineering feasibility $\rightarrow$ responsible deployment	Scenario-valid operational evidence, ecological comparisons, risk review and applicable permission	Using controller feasibility as environmental approval
Research records $\rightarrow$ institutional legitimacy	Correct legal sources, competent authority, affected rights and reviewable decision	Treating provenance, open access or software flags as authorization

with a stronger meaning than its evidence supports. The present audit treats a consequential claim as a scoped assertion $c = (q, D, E, U)$, where q is the proposition, D specifies the referent and conditions of intended applicability, E identifies the evidence or argument used to support it, and U records material uncertainties and exclusions. This tuple is an editorial and methodological schema, not a proposed physical ontology or an empirically established universal model. Claims with similar verbal descriptions can differ materially in D , E or U .

Suppose claim c_1 concerns the response of an imaging instrument under calibrated illumination, while c_2 concerns the presence of an organism and c_3 concerns the organism's welfare. Passage from c_1 to c_2 requires a justified detection and interpretation model, with attention to false detections, organisms that evade the platform, and transport across the sampled setting. Passage from c_2 to c_3 additionally requires taxon-sensitive evidence and a defensible account of the welfare-relevant construct. A fourth claim concerning permission to operate requires a separate legal and institutional inquiry. An accurate image, a defensible biological inference, a precautionary welfare decision and a valid authorization may coexist, but they are neither mutually substitutable nor guaranteed by a common data record.

The audit also requires a precise distinction between *validation* and *transfer*. Validation assesses a claim within the specified conditions; transfer asks whether the relevant model, instrument, organism, context or governing rule is sufficiently similar in a second setting. This latter step needs its own evidence or defensible assumptions. Identical notation in two sections does not by itself identify the same random variable, physical quantity, sampling frame or competent evaluator. For instance, a transition kernel in Section 21 has a modeled ecological-state role, whereas a stochastic controller in Section 14 concerns vehicle behavior and declared disturbances. They should be coupled only through a defined state interface and validated physical assumptions. Likewise, a control-theoretic safety margin cannot be added numerically to an ethical value, and a positive statistical estimand cannot represent a right of consent.

Established practices provide the demanding comparison baseline. FAIR data stewardship, PROV-style provenance and the CARE Principles where Indigenous data rights are relevant already address parts of what a GR-informed record would seek to preserve [214, 220, 22]. The manuscript's distinctive proposal must therefore be evaluated on a further task-specific contribution, such as better preservation of contested interpretation, ecological intervention history or correctability across institutional boundaries. Where careful conventional procedures perform equally well, that negative comparative finding should constrain claims of necessity or novelty. A useful audit record consequently identifies each claim's cited source, exact supporting result, method and population, limitations, bridging assumptions, available

counterevidence and proposed remedy for error. Source identity without source-to-claim correspondence remains insufficient.

This is a manuscript-level review of inferential architecture, not a certification that every experiment, legal interpretation or welfare hypothesis has been checked by an independent domain expert. The provenance system should make that residual boundary visible to any reader. It may also identify claims that are inadequately testable, underidentified or inconsistent with available evidence, in which case their role should be reduced to explicitly exploratory discussion or omitted.

25.13 Integrated comparative trials, falsifiers and research continuation

We conclude with an integrated but deliberately staged comparison programme. One prospective *reef-adjacent observation trial* would begin with desk-based authorization and ecological review, a measured ambient sound-and-light baseline, and a validated instrument configuration. Matched surveys could compare a fixed-route robot with a policy-aware adaptive robot and a passive or more distant alternative, provided intervention is justified and permitted. Technical outcomes would include navigation error, energy, observation completeness and calibration drift; inferential outcomes would include blinded image interpretation, independently checked biological findings and uncertainty calibration. Welfare-related observations would be independently assessed through predeclared behavior or exposure measurements, avoiding unvalidated claims about subjective animal states. Adverse responses, inadequate detection power or inability to recover a robot would be reasons to narrow or halt the research scope.

A second *observatory maintenance trial* could begin with simulated connector tasks, then move to pressure-tank and flume systems, with field introduction only after controlled reliability and recovery evidence. Compare a rigid manipulator, a compliant end effector and a human-supervised intervention baseline under matched load and visibility regimes. Test contact forces, connector condition, cumulative energy, disturbance of the adjacent substrate, intervention duration and recovery after communications loss. The criticality supervisor should be assessed against meaningful adversarial conditions, including a sudden loss of model validity. Success should require both completion and evidence that the declared operational envelope and monitoring procedures were respected; failure must include previously unreported human rescue and unobserved collateral disturbance.

A third *infrastructure alternatives study* would compare shore-based or surface-hosted instruments, a recoverable submerged prototype, and a more persistent installation across the same specified scientific objectives. Its assessment boundary would include manufacture, vessel logistics, power systems, maintenance, construction or decommissioning, habitat exposure, computing or observation benefit, and institutional control of knowledge. Where different environmental effects cannot be placed in a defensible common unit, the study should retain a vector of findings and transparent tradeoffs. Legal permissions must be reviewed for each design and site separately. The aim is a comparative account of possible continuations, not the identification of a universally optimal ocean infrastructure.

Across these studies, explicit *failure tests* are essential. The proposed GR representation should be reconsidered if its extra relational records do not improve any predeclared inference or accountability outcome over existing scientific provenance methods. A tangent-space controller's claims should be narrowed when its stated residual envelope is empirically violated. An adaptive survey should be criticized if improved event targeting entails a documented deterioration of independently evaluated field estimates without appropriate qualification. An apparent ecological safeguard should be revised when monitored organisms or habitats exhibit consequential responses that the screening model systematically misses. A participatory protocol should not be judged adequate when affected actors cannot inspect its evidence or seek practical correction. These are heterogeneous tests with different implications; none alone invalidates every part of the wider inquiry.

Finally, independent scrutiny requires a release architecture: registered questions and exclusions, versioned software and sensor configuration, calibration provenance, machine-readable quality flags, observation-policy logs, controlled access where warranted, model and analysis notebooks, null and adverse findings, and a public account of which claims remain unsupported. Scientific capacity to re-

produce a result is partly infrastructural and economic, so a research programme should report required computing resources, proprietary dependencies and the degree to which other groups can realistically perform its tests. Every published inference should identify an appropriate contact or procedure for subsequent correction, and every ecological intervention should have an accountable record of the authority and conditions under which it occurred.

These designs leave substantial uncertainties unresolved. They cannot make nonhuman experience fully legible, eliminate the dynamic character of oceanic environments, or transform every ecological concern into a scalar risk. They can nevertheless clarify where the emerging research programme has acquired trustworthy evidence and where new observations might change its direction. Section 26 will bring these methodological conclusions together with the paper's philosophical, engineering and institutional investigations, identifying continuations that remain open without conflating technical capability, knowledge, ecological responsibility and authority.

26 Conclusion and Research Continuations

26.1 The scope and result of the inquiry

This paper has investigated the conditions under which oceanic robotic systems can contribute to ecological knowledge formation while remaining answerable to the material, biological, epistemic and institutional relations in which their operations occur. Its starting point was the conjunction of a scientific opportunity and an ethical difficulty. Oceanic technologies can extend observation across the sea surface, water column, seabed and, for selected purposes, the sub-seafloor. The same vehicles, instruments, missions and supporting infrastructures also enter into the histories of the systems they seek to understand. A sampling decision affects which events are recorded; a moving observer may alter animal behavior; a physical intervention may change a habitat; a sustained installation reorganizes patterns of energy, material, access and institutional control. An adequate account of scientific success must therefore consider what was observed, how the observation came to be possible, what warranted the inference, and which further ecological and social continuations were affected.

The inquiry proceeded through four connected but analytically separable undertakings. Sections 2–6 developed a provisional generative-relational (GR) vocabulary for ecological reality, heterogeneous relations, observable manifestations, finite histories, generative capacity and partially translatable knowledge. Sections 7–10 investigated the normative plurality of marine activity, lawfully scoped authorization, animal and ecological responsibility, heterogeneous value creation and the practical conditions of review and remedy. Sections 11–21 assembled a topic-specific engineering survey spanning established marine platforms, morphology, hydrodynamics, control, perception, navigation, manipulation, actuation, materials, endurance, collective operations and adaptive observation. Sections 22–25 then examined welfare-sensitive engineering, biological mechanism transfer, prospective infrastructure under finite ecological conditions and differentiated research-validation methods. The resulting account is an architecture of questions, models and comparative tests. It establishes neither a completed robotic system nor the empirical adequacy of a universal ecological ontology.

Its principal integrative claim is conditional: the scientific value of robotic encounters can be assessed more adequately when the record of a measurement retains enough information about the observing relation, the intervention, the inferential procedure and its limitations to permit situated criticism and subsequent revision. Whether the additional representations proposed here improve particular research outcomes is an empirical question. Existing scientific disciplines already possess extensive techniques for calibration, metrology, provenance, experimental design, uncertainty analysis and governance. The proposed GR perspective contributes an organizing vocabulary for relating such practices across partially incompatible systems; it has yet to demonstrate an independent gain over well-designed alternatives. Its usefulness should be judged by those comparisons rather than by conceptual breadth alone.

The argument also yields an important limitation on integration. Mechanical feasibility, predictive accuracy, biological interpretation, welfare-relevant evidence, substantive ecological justification and

legal authority involve different predicates, evidence and institutions. A simulation can support a controller's numerical performance claim without establishing its animal-welfare consequences; an impact study can inform decision-makers without itself granting permission; a procedure can be lawfully administered while remaining open to ethical contestation. A unified project must preserve these distinctions even as it creates disciplined exchanges among them. This principle governs the research continuations below.

The strongest integrative test concerns the interfaces between conceptual and applied claims. The analytical registers do not themselves identify a microbial mechanism or authorize a sampling action. Ecological generativity describes distinct candidate capacities whose biological conditions require independent investigation. The platform survey establishes task- and regime-specific performance, while an adaptive observation scheme additionally requires evidence about action-dependent disturbance, detection and selection. The methodological integration proposed here is therefore best evaluated through explicit counterexamples and comparisons: whether recorded observation and intervention histories prevent a mistaken ecological interpretation, whether reviewers can locate an evidential or authority failure, and whether the additional architecture improves on established provenance and evaluation practices. Failure on these tests would call for revising or narrowing the programme rather than protecting its vocabulary from criticism.

26.2 Relational histories, observation and ecological knowledge

The philosophical investigation proposed three provisional analytical registers: ontological, relational and manifestational. The ontological register concerns what may exist or occur independently of a particular researcher's access; the relational register concerns interactions, dependencies and histories through which systems become coupled; and the manifestational register concerns what is made available to particular observers, instruments and interpretations. These are distinctions for inquiry, not an experimentally established decomposition of the ocean. Their epistemic importance lies in resisting an unnoticed transfer from an observed trace to a complete description of the producing ecology. Two different ecological histories may generate indistinguishable measurements under a chosen instrument and observation policy. Additional observations, independent modalities or interventions can sometimes discriminate them, although those choices may introduce new disturbances and costs.

Consequently, ecological knowledge formation was modeled as a process of revising the relation among a question, a measurement practice, a family of representations, the available evidence and the possible corrections. An observation history includes sensor characteristics, calibration, sampling effort, instrument position, environment, mission decisions, data transformations, exclusions and any recordable changes in the observing system. Such histories can support ordinary scientific provenance practices, while the proposed GR extension asks whether relations among different forms of evidence and changes in the repertoire of possible inquiry can be represented explicitly enough to improve understanding. The distinction between data acquisition, ecological inference and the production of new scientific capacities thus remains central.

An adaptive vehicle illustrates the difference. Its policy may preferentially visit areas where a model predicts a rare organism. The resulting observations may improve the study of those selected events while becoming poorly suited to an unqualified estimate of population prevalence. Section 21 showed that a design-based unbiasedness argument is available for a *fixed, finite* target under recorded positive inclusion probabilities and appropriate measurements. That result does not automatically carry over to inaccessible habitats, evolving populations or policies with unknown selection probabilities. Section 25 similarly used a standard total-variation inequality to clarify the conditions necessary to transfer a bounded evaluation loss across test distributions. Neither mathematical observation creates ecological knowledge on its own. Both make assumptions visible at points where seemingly successful observation may support an unwarranted generalization.

Cross-system translation presents a second limitation. A vehicle controller, a biological survey, an Indigenous or local knowledge practice, a marine manager and a legal decision procedure may maintain

different histories, purposes and admissible reasons. Their representations can sometimes be mediated for a specified task without being reduced to a common total description. A useful knowledge infrastructure would record the provenance and permitted uses of an interpretation, preserve unresolved conflicts, and support correction by participants with relevant expertise or standing. These are design propositions, to be evaluated against established interoperability standards and situated practices, including FAIR-oriented data stewardship, CARE considerations for Indigenous data governance where applicable, and the existing marine observing and metadata infrastructures discussed earlier [214, 22, 220]. Such standards do not themselves resolve the philosophical problem of another system's experience or the political question of who may determine its use.

More generally, the project should avoid treating a single numerical or chronological coordinate as the substance of ecological history. The temporal ordering of observations may be indispensable for scientific analysis, yet some continuations are characterized by thresholds, recurrence, path dependence, changing relations or newly possible interactions. A relational account can investigate these structures without claiming that a reconstructed history captures every dimension of an organism or ecosystem. Its scientific contribution remains proportional to its capacity to expose alternative explanations, organize measurements and allow informed disagreement.

26.3 Ecological generativity, nonhuman life and finite conditions

The proposed notion of ecological generativity concerns capacities for biological, relational and material continuation under historically situated conditions. It is intended to organize questions about reproduction, adaptive repertoires, changing ecological associations, ecosystem recovery, possibilities for further observation and the maintenance of productive scientific capacity. These possibilities are heterogeneous. An organism's physiological condition, the persistence of a habitat, a food-web relation and a research group's access to calibrated instruments cannot be assumed commensurable. Generative capacities likewise depend on energy, matter, ecological interactions, available space, disturbance histories and sometimes highly localized conditions that are poorly represented by a global efficiency metric.

This account has implications for the treatment of nonhuman life. Observed responsiveness, agency in an operational sense, indicators relevant to possible sentience and evidence about welfare are related but distinct. Evidence supporting pain or sentience in some taxa cannot be extended without qualification to every marine organism; lack of evidence under an inadequate instrument cannot be equated with absence of interests. Research on sentience, welfare and behavior supports taxon- and context-sensitive protocols rather than a universally valid measure of subjective experience [20, 107, 29]. The possibility that a research platform modifies the very behavior it records is also a methodological problem of observer interference, requiring explicit attention to vehicle characteristics, exposure pathways, baselines and independent controls [17].

The engineering sections accordingly distinguished modeled exposure from biological consequence. Acoustic output, illumination, flow and shear, contact force, heat, chemical release, retrieval failure and cumulative activity can be measured or bounded under particular assumptions; their effects on welfare or ecological viability depend on organisms, places, histories and uncertain mechanisms. Section 22 developed an explicitly conditional bound on a selected cumulative exposure surrogate. Such a bound supports a transparent control or screening exercise only within its declared sensing, aggregation and model-validity conditions. The existence of a numerical limit is not evidence of a species-independent threshold for acceptable harm. Where a relevant response is unobservable or a model fails to cover plausible mechanisms, renewed assessment, alternative design, lower exposure or withdrawal may be scientifically and normatively appropriate.

Finitude becomes more pronounced for persistent infrastructures. Cable systems, observation nodes, batteries, autonomous service operations, seabed interaction and prospective subsea computation entail material inputs, power, thermal exchanges, maintenance burdens and opportunities for cumulative ecological pressure. Some benefits may accrue to distant beneficiaries while exposure is locally concentrated. A lifecycle assessment should make its boundary, uncertainty, beneficiaries, burdens and excluded rela-

tions explicit. Where effects differ in kind, they should be retained as a vector of evidence and disputed judgments rather than concealed in a single net-benefit score. The bounded resource-screening model proposed in Section 24 is useful for checking compatibility with declared constraints; it neither establishes universal limits nor decides what a just allocation would be.

A further research problem concerns possible incompatibilities among valuable continuations. Preserving a habitat's disturbance-sensitive relations may limit access to an otherwise scientifically informative sampling site. A robot designed to avoid biological contact may consume additional power or require a longer mission. A device that improves local measurement may generate construction burdens elsewhere in its supply chain. A viable research practice acknowledges such tensions and evaluates alternative feasible arrangements rather than presuming that all scientific, ecological and economic values necessarily increase together.

26.4 The engineering programme and its evidential maturity

Oceanic robotic systems were treated throughout as differentiated configurations of bodies, sensors, actuators, energy supplies, communication links, software, operators and environmental relations. The review documented the distinctive tasks of surface craft, drifting and profiling instruments, gliders, autonomous vehicles, remotely operated vehicles, seabed platforms and fixed observatories. These systems have different observational access, endurance, controllability, recovery requirements and intervention capacities. A biologically inspired morphology may improve a constrained swimming or gripping function while having no demonstrated advantage for another task. The proposed design principle is therefore comparative: morphology, control and sensing should be evaluated jointly for a specified scientific or operational objective rather than selected from a generalized expectation of biological superiority.

The mechanical analysis assembled hydrodynamic forcing, fluid–structure interaction, pressure, thermal loads, contact and environmental transport into the study of embodied operation. It then distinguished locally validated estimation and control from extrapolation into unknown regimes. The terminology of tangent-space governance refers to using local models and sensitivity information within a supported neighborhood; criticality governance refers to detecting loss of model adequacy, approach to a declared operational boundary or possible regime transition. These are hypotheses for architecture and validation. Existing robust, adaptive, predictive, barrier-based and learning-based control methods provide important baselines. The conditional error and reserve results in Sections 13–14 and 19 require their assumptions to be checked and say little about unspecified disturbances. A responsible comparison should report failures, recovery costs, sensing uncertainty, off-nominal behavior and the controls used to establish differences in performance [45, 36, 84].

In the perception and navigation investigations, physical measurements were similarly kept separate from the scientific entities inferred from them. Acoustic returns, optical images, chemical signals, molecular samples and position estimates carry modality-specific selection and uncertainty. Sensor fusion can increase robustness when error sources and dependence are addressed, but correlated error and misleading calibration may also create apparent corroboration. A scientific observing system needs a provenance model at least as careful as its vehicle-state estimator. This requirement extends to communication-limited and heterogeneous teams: synchronization, evidence quality, delegated action limits and conflict resolution cannot be assumed from task allocation alone. The communication-cut analysis of Section 20 clarified one circumstance in which agents sharing no decisive information cannot individually distinguish different global conditions. Distributed autonomy should therefore expose its information and authority boundaries to human operators and subsequent reviewers.

Marine manipulation and infrastructure service introduce contact, reaction forces and persistence. A remotely supervised connector operation and general-purpose autonomous seabed construction occupy different evidential levels. Soft and compliant end effectors can change force distribution without establishing harmlessness; experimental bio-inspired actuators may demonstrate a mechanism without meeting prolonged pressure, corrosion, power, reliability or recoverability requirements. Across morphology, material choice and design, evidence maturity is local to the tested function and environment.

Section 23 therefore examined transfer from living systems as a context-dependent engineering investigation, with separate questions about actual biological participation, species interests, and any living component used in a biohybrid device.

A promising direction is the combination of heterogeneous instruments with explicit scientific objectives and conservative intervention constraints. Nevertheless, the systems integration problem cannot be settled by collecting independently successful modules. Sensor placement changes flow and hydrodynamic signatures; power limits change trajectories and sampling distributions; communication decisions change which evidence is available; and recovery requirements may alter the ecological footprint of a mission. The continuation of this engineering programme must test these couplings through subsystem qualification, controlled integration, stress testing, field evidence where justified, and reliable recording of deviations from assumptions. Detailed reports of negative results are particularly valuable, because a limited platform's failure may identify a crucial physical or biological boundary of the broader research ambition.

26.5 Marine nomos, knowledge circulation and institutional responsibility

The institutional investigation started from the plurality of normative settings in which marine work occurs. Jurisdiction, resource use, scientific research, environmental protection, navigation, installation activity, genetic-resource governance, animal welfare and the organization of knowledge are governed by different and sometimes overlapping instruments and authorities. Legal analysis must identify the activity, place, actors, applicable legal sources and institutional decision before inferring rights or duties. In particular, territorial sea, exclusive economic zone, continental shelf and areas beyond national jurisdiction cannot be assigned a single undifferentiated permission regime. UNCLOS separates relevant scientific-research, jurisdictional and environmental provisions, while the BBNJ Agreement establishes specified mechanisms within its own scope; neither is a universal license for autonomous scientific or productive installations [196, 190, 191, 189].

The analysis of proceduralism and ecological justice distinguished legal compliance, procedural adequacy and substantive normative justification. An institution may possess a recognized decision competence while facing serious evidential uncertainty or disagreement about affected interests. Procedural quality involves reasons, notice where applicable, transparent assessment of alternatives, appropriate participation and expert scrutiny, explicit conditions for revision, and accessible means to seek correction or remedy. Nonhuman animals cannot ordinarily participate through the same institutions as human actors; representation of their possible interests therefore requires biological evidence, ethical reasoning and scrutiny of claims made on their behalf. Formal assent by one organization cannot settle all relevant questions of harm, knowledge ownership or ecological consequence.

Knowledge circulation has an associated political economy. The production of a marine data set depends on instruments, trained labor, coastal and transnational permissions, public and private capital, environmental access and ongoing institutional maintenance. Its value may circulate through publications, infrastructure design, forecasting services, commercial applications, ecological management and the training of models. These forms of value are partly heterogeneous: scientific reproducibility, local environmental benefit, firm revenue, community control and living ecological capacity require distinct accounts. The commons literature and principles for responsible data governance illuminate some institutional alternatives [61, 214, 22], but neither open publication alone nor an asserted commons designation guarantees equitable access to productive capacities, shared decision-making or ecological protection.

A particularly important unfinished question is the distribution of capacities to *form, challenge and use* ecological knowledge. Some actors can provide vessels, computation, proprietary sensors or model infrastructure; others possess situated ecological practices and bear exposure to interventions; still others exercise licensing or supervisory authority. The resulting asymmetries affect whose questions can be pursued, which signals become recognized as evidence, and who can contest consequential conclusions. A prospective GR knowledge infrastructure should investigate the representation of evidence lineage, authority claims, restricted or consent-governed access, versioned corrections and contested interpretations.

Table 30: Research claims, appropriate evidence, and limitations on cross-domain inference.

Claim type	Principal evidence or comparison	Inference requiring separate justification
Mechanical capability	Controlled tasks, disturbance envelopes, component qualification, recovery tests	Scientific informativeness, ecological acceptability
Measurement reliability	Calibration, uncertainty budgets, quality control, instrument intercomparison	Organism-level or ecosystem-level interpretation
Ecological inference	Explicit populations, sampling policies, controls, model criticism, replication	Generality across sites, welfare or causality beyond design
Welfare and exposure	Taxon-sensitive evidence, measured exposures, alternative designs, monitored responses	Comprehensive experience, universal safety thresholds
Institutional permission	Applicable law, specified activity, competent decisions, conditions and records	Ethical justification or authorization in another jurisdiction
Knowledge and value circulation	Provenance, access patterns, participatory scrutiny, benefit and burden accounts	Equitable outcomes, universal commensurability

Such a system would support accountable circulation if it demonstrably improves the ability of appropriate participants to scrutinize and revise scientific claims. Technical traceability, however, does not settle who is entitled to obtain protected knowledge, what benefit-sharing arrangement applies, or which institution may authorize an intervention. Those require context-specific legal and normative argument.

The manuscript's hypothetical transnational installation scenarios therefore remain thought experiments about differentiated responsibilities and value circulation. They identify questions for competent legal counsel, environmental authorities, affected stakeholders and disciplinary researchers. They should not be construed as determinations concerning an actual location, enterprise, state practice or permission. Their continued value depends on producing better question inventories, more precise actor-and-activity mappings and institutionally feasible routes for disagreement and remedy.

Differentiated legal and ethical review. The completed legal-source review makes a further methodological constraint explicit. Oceanic autonomy, marine scientific research, environmental assessment, genetic-resource utilization and ecological representation are potentially intersecting but separately classified activities. UNCLOS Part XIII's differentiated research-consent provisions, the BBNJ Agreement's scope-specific environmental and benefit-sharing procedures, and the IMO's 2026 non-mandatory MASS Code cannot be fused into an undifferentiated global ocean-robotics licence [196, 189, 72]. Nor can a scientifically reasonable concern be assumed to confer a particular legal right of action. The proposed contribution is narrower and more testable: preserve the competent source, the classification assumptions, the decision and its reasons, the ecological evidence and uncertainty, and an available path to review. Whether such records improve practical accountability is a comparative research question, and whether a particular operation is lawful remains an issue for competent authorities and qualified specialist analysis.

26.6 Linking evidence, uncertainty and revision without category collapse

The validation programme in Section 25, including its cross-sectional claim-dependency audit (Section 25.12), offers a practical connection among the previous investigations. An integrated research project may share observations and records across disciplines, but should distinguish the kind of claim each result can warrant. Table 30 summarizes the main interfaces. Its categories are intended as a research and audit aid, rather than a complete taxonomy or prescribed legal checklist.

A useful operational principle is to assign every consequential claim an accountable *claim dossier*: the specific proposition under test, its domain of intended application, relevant histories, measurement and modeling procedures, independent comparisons, failure criteria, residual uncertainties, evidence-access conditions and responsible review route. Dossiers should link without silently changing one an-

Table 31: Prospective research workstreams and the evidence that would challenge them.

Workstream	Initial comparative question	Challenge or revision condition
GR representation and inquiry	Do relational histories improve predeclared cross-modal inference or correction over conventional provenance?	No repeatable gain; indistinguishable claims with higher representation cost
Oceanic robotics and control	Do task-specific morphology and bounded adaptive control improve performance under matched disturbance, energy and recovery constraints?	Model validity or recovery bounds fail; gains vanish under fair baseline
Ecological observation and welfare	Can observation designs estimate relevant response and sampling effects with independent controls?	Exposure or selection effects remain unmeasured, or evidence is inadequate for the stated organism or habitat
Knowledge and institutional procedure	Can relevant actors inspect, contest and revise consequential evidence within valid authority and access constraints?	Formal records exist without usable correction or competent authorization
Infrastructure and political economy	Do alternatives preserve specified scientific outcomes with more defensible lifecycle and distributional consequences?	Material, energy, environmental or governance burdens omitted from comparison

other's scope. A sensor classifier's positive detection may inform an ecological hypothesis; it does not automatically license close approach to the observed organism. A welfare-related response may trigger operational caution while remaining insufficient to settle a disputed theory of sentience. Legal authorization may impose obligations on data use but cannot retroactively establish scientific validity. In practice these records must remain proportionate to mission scale and institutional capacity, with protection for sensitive ecological and community information.

The argument for preserving these distinctions is epistemic as well as ethical. Overaggregation can conceal a decisive failure behind a successful average. A mission achieving high mapping completeness might impose an unanticipated disturbance on a vulnerable habitat; a low-disturbance mission might produce systematically preferential observations; a transparent data catalogue might be unusable to researchers lacking computing resources or the ability to contest interpretations. The research programme should therefore evaluate the joint feasibility of its obligations and report unresolved tradeoffs explicitly. The existence of a formally feasible control policy under a model is one form of result; evidence that a deployed system remains acceptable to the relevant human and ecological processes is another.

Revision must also be treated as substantive work. A model can be challenged by new calibration evidence, a previously inaccessible ecological indicator, a shift in local environmental conditions, a failure of an instrument, a change in legal requirements or a credible objection to the use of a data set. A responsible response may include reanalysis, retraction of a specific inference, changed operating boundaries, new participation, suspension or physical withdrawal. Some changes will be scientifically warranted without automatically dictating a particular institutional action; others will arise from legitimate authority or precaution even when the model's point estimate remains unchanged. Preserving a traceable record of these different reasons would be an important test of the proposed relational infrastructure.

26.7 A staged programme of research continuations

Further research should proceed through linked but independently criticizable workstreams. A sequence from low-intervention analysis to controlled integration and, where warranted, authorized field studies can reduce the temptation to represent architectural aspirations as demonstrated capabilities. Table 31 identifies representative near-term questions and failure conditions; the proposed stages are conditional research priorities, not a calendar or a promise of deployment.

Relational ecological epistemology and formal representation.

A first workstream should turn the provisional registers and evolving relational histories into interoperable, technically modest representations. A candidate graph or cell-complex model could record selected events, dependencies, observation actions, provenance relations, uncertainty and reversible interpretive revisions. Its purpose would be to answer declared scientific queries and represent conflicting or incomplete histories, not to claim a complete reconstruction of ecological being. Minimal test cases could compare this representation with conventional relational databases, provenance graphs, event logs and domain-specific data standards. Predeclared tasks could include explaining a classification discrepancy, reconstructing an adaptive observation policy, detecting a mismatch between sampling history and reported prevalence, or tracing a contested downstream interpretation. Evaluation should report correctness, missingness, semantic ambiguity, computational burden, maintainability and access constraints. If simpler structures perform equally well, the GR formalism should be simplified or its claims narrowed.

Formal research may separately investigate how a repertoire of possible scientific actions changes when a sensor, model, collaborator or environmental relation becomes available. One could represent histories as partially ordered observations and actions, with conditional transitions that depend on documented capacities and constraints. Such formalizations should be compared with existing methods in state estimation, active learning, causal inference, event-sourced computation and dynamical systems. Assertions concerning ecological “generativity” should identify the particular modeled capacity, feasible continuations and observational basis. A categorical, topological or geometric description acquires explanatory force only when its constructions yield meaningful invariants, distinctions, testable predictions or improved practical reasoning. A novel notation, by itself, is not evidence of a new ecological mechanism.

Embodied oceanic engineering and situated model validity.

A second workstream should build controlled, reproducible tests of differentiated marine robotic configurations. Comparisons could begin with digital hydrodynamic and sensor models, then proceed to flumes, tanks, pressure facilities and integrated testbeds before any considered field operation. Trials should include conventional and biologically inspired platforms as well as rigid, compliant and hybrid mechanisms, chosen for specific tasks rather than a ranking of whole robot families. The priority experiments are those that reveal cross-layer effects: whether a compliant manipulator changes vehicle station-keeping demands; whether a quieter propulsion mode alters endurance and biological detectability; whether an adaptive controller remains within its assumed error envelope when sensor quality deteriorates; and whether teams can preserve shared science and safety records during communication interruptions. Adverse trials should include unexpected currents, contact failure, acoustic ambiguity, thermal loads, loss of localization and practicable recovery limits.

This workstream also calls for clearly bounded definitions of tangent-space and criticality governance. A candidate implementation would specify local models, the measurable residuals that trigger reconsideration, any stability or reachability assumptions, and an independently testable action when those assumptions are breached. The evidence for success would concern defined error, energy, contact, recovery and observation objectives compared with strong baselines. It would not, without additional investigation, imply animal welfare, habitat compatibility or generalized autonomy. Independent laboratories should be able to rerun representative tests with documented equipment, configurations and software versions. Robustness across sites and ecological contexts would remain a separate question even after tank-level reproducibility.

Ecological encounters, causal assessment and nonhuman interests.

A third workstream should couple ecologically justified observation with specialist-designed assessment of organism and habitat responses. Candidate studies could compare distant or passive instruments, fixed-route vehicles, and adaptive vehicles, but only where the scientific question, species and anticipated exposures warrant the comparison. Baselines should account for diel, seasonal and environmental variability; controls should be chosen to identify the effect of the platform rather than merely the outcome of a convenient transect.

Behavioral or physiological indicators require taxon-sensitive interpretation and appropriate expertise. Multiple-control before–after designs or other valid causal strategies may be useful when their assumptions can be met [184, 185]. The failure to observe a response should be evaluated against detection power, exposure coverage and plausible unmeasured effects before it becomes a safety-related assertion.

A related ecological question concerns the persistence and distribution of consequences. Repeated low-intensity missions, a recoverable observatory, a semi-permanent cable route and an infrastructure-dependent fleet may have different footprints and potential ecological dependencies. Research should combine time- and place-bounded measurements with lifecycle inventories and independent ecological interpretation. A changing habitat may generate new research opportunities while simultaneously losing capacities valuable to its resident organisms. A framework capable of recording both outcomes is analytically preferable to one that reports scientific gain as a sufficient proxy for ecological benefit. Specified alternatives, including reduced intervention or no deployment, must remain genuine comparators.

Institutional practice, public knowledge and value circulation. A fourth workstream should examine whether institutions and evidence systems actually provide usable procedures for correction and disagreement. Its unit of analysis might be a hypothetical cross-border observing programme, an existing scientific consortium operating under its actual permissions, or a carefully designed participatory simulation. Research questions include: who can initiate or alter a research objective; which evidence is needed for each authorization; how commercial, scientific and sensitive data are classified; what costs are imposed on actors seeking review; and how objections, uncertainty or new ecological information change a mission. A cross-jurisdiction comparison should preserve differences among research consent, navigation, resource activities, installation permissions, environmental duties, intellectual-property arrangements and treaty-based mechanisms. Any empirical investigation involving affected institutions or communities would require appropriate research ethics and contextual competence.

The circulation of value also deserves separate political-economic work. One promising comparison would examine how ownership, access, governance and operational control of observation infrastructures condition the distribution of scientific capacity and economic returns. A research commons can provide meaningful collective capabilities where maintenance, computational access, methodological intelligibility and contestable governance accompany data sharing. It can also reproduce inequalities when resources or decision rights remain concentrated. Economic valuation should specify whose value is measured and which ecological consequences remain nonmonetized; legal entitlement and ethical justification require independent treatment. A future assessment of generative outcomes should consider the capacities created for further inquiry and ecological stewardship, as well as the opportunity costs and relations diminished by their production.

Integrated trials and decision points. The integrated continuation described in Section 25 offers a possible bridge among these workstreams: a reef-adjacent observation comparison, a controlled observatory-maintenance study, and a comparative assessment of surface, shore-based and submerged infrastructure alternatives. Each should advance only when its preliminary question has a justified design, reliable measurements, an appropriate environmental and institutional review pathway, and feasible suspension or recovery provisions. The transition from one stage to another should be governed by explicit criteria, including evidence quality and possible ecological burden, rather than an assumption that additional technical capacity is necessarily desirable. Where necessary evidence cannot be produced responsibly, the corresponding experiment should remain unperformed or be replaced with a non-intervention alternative.

An interdisciplinary research programme also needs negative-result preservation. A protocol that reveals uncorrectable sampling bias, a robot morphology with no task-specific advantage, a control envelope too fragile for useful application, a welfare surrogate with poor biological validity or an institutional procedure offering symbolic rather than effective review constitutes informative progress. Such results can delimit the feasible domain of a hypothesis, redirect resources, and prevent its premature application. A common release structure could include research questions, assumptions, software and hardware provenance, statistical designs, negative outcomes, review conditions, source limitations and

corrections, subject to legally and ethically justified access constraints. These elements should be designed with the relevant disciplinary and institutional communities rather than imposed as a universal reporting ontology.

Habitat succession and persistent material conditions. A further continuation concerns comparisons across ecological timescales. Long-interval seabed studies document changes that remain observable decades after some intensive disturbances, while recording selective biological re-establishment; complementary work on corals, microbial functions and artificial offshore structures demonstrates the importance of taxon, substrate and reference conditions [83, 205, 16, 91]. These findings do not imply that research robots and mining collectors have equivalent footprints. They motivate a research programme in which physical disturbance, substrate dependence, community composition, microbial processes and infrastructure material recovery are documented separately. In the proposed GR interpretation, a change in observed abundance need not represent restoration of the ecological relations that support further biological continuation. Comparative experiments and responsible decommissioning decisions require evidence of those relations, not merely an aggregate performance indicator.

26.8 Limitations, intellectual responsibilities and the prospect of continuation

Several limitations shape the interpretation of the manuscript. First, its GR vocabulary is an exploratory synthesis across philosophy, marine ecology, robotics, information systems and political economy. It cannot be validated in the abstract by citing successes of the scientific methods it seeks to organize. Its independent content must emerge through explicit definitions, cases where it discriminates alternatives, and comparisons showing what scientific or institutional work its representations enable. Some components may remain useful only as philosophical questions, and others may be absorbed by established disciplinary frameworks after specialist critique. This possibility should be considered a normal outcome of conceptual investigation.

Second, the engineering evidence reviewed is uneven in task, environment and maturity. Operational gliders, remotely operated vehicles, cabled observatories and selected adaptive sampling systems have established histories, whereas particular bio-inspired mechanisms, cooperative interventions, persistent robot-service infrastructures and generalized autonomous construction remain much more restricted in demonstrated scope. A successful report about one physical capability, or an engineering paper's preferred evaluation setting, cannot be treated as field evidence for a comprehensive ecological observing architecture. Similarly, the formal propositions developed in this manuscript are principally conditional consequences of explicitly selected assumptions, together with familiar statistical and mathematical inequalities placed in a research context. The formal-consistency pass in Section 25.9 clarifies frame compatibility, dimensional scaling, inclusion positivity, and the distinction between mathematical bounds and calibrated operational guarantees; it supplies neither automated proof certification nor validation of those premises in the sea. They are not a catalogue of newly verified oceanic physical laws.

Third, the treatment of nonhuman interests acknowledges substantial epistemic asymmetry. Animal behavior, physiology, habitat effects and population dynamics can be investigated with scientific rigor, yet interpretation of experience and welfare may remain uncertain and taxon-dependent. The paper does not establish a comprehensive theory of marine sentience, prescribe a universal moral ranking of organisms, or convert ecological plurality into a single calculable welfare currency. Precaution, meaningful alternatives, expert criticism and institutional responsibility are consequently part of the research practice itself. A technically sophisticated model may increase observability while leaving ethically relevant dimensions unrepresented.

Fourth, the normative and legal arguments are general investigations of differentiated authority and justification, rather than applied legal advice. Treaty provisions and domestic measures have specific scopes, relevant institutions and histories of interpretation. The hypothetical facilities used to explore transnational value circulation are not substitutes for competent analysis of any actual proposal. An acceptable research practice may require legal permission, environmental impact assessment, animal-

welfare scrutiny, community engagement, data-governance restrictions or additional forms of authorization according to context. The normative discussion also remains open to disagreement over representation, distribution, duties toward nonhuman life and what constitutes a defensible tradeoff among ecological and scientific goods.

Finally, the paper reports no original marine deployment, animal experimentation, completed ecological impact assessment, independent legal ruling or validated universal model of generativity. Its primary research outputs are a structured interdisciplinary question, a philosophy-to-engineering review, provisional formal and institutional distinctions, and a staged comparative research agenda. The breadth of this undertaking creates risks of uneven disciplinary depth and unrecognized conflicts among source traditions. Future revisions should prioritize primary-source verification, engagement with specialist counterarguments, transparent limits of applicability, and independent review of the empirical, legal, welfare and normative components. The intellectual precedents and existing standards cited throughout should be treated as substantive research contributions in their own right, not merely as background for a new vocabulary.

The prospect advanced by this inquiry is therefore modest in its warrant but consequential in its orientation. Robotic observation may expand the kinds of marine relations that can become scientifically legible, while also changing the histories through which those relations continue. The legitimacy of that expansion depends upon how research systems handle incomplete knowledge, living interests, material finitude, unequal capacities and revisable authority. A meaningful continuation would be one in which scientific claims become better grounded, the conditions and consequences of encounter become more visible, and affected institutions and communities acquire usable means to question decisions and pursue correction. Whether oceanic robotics can contribute to these aims, and through which combinations of engineering, science and institutional arrangements, remains a problem for critical comparison, situated experiments and continuing public and scholarly inquiry.

References

- [1] *A BGC-Argo Guide: Planning, Deployment, Data Handling and Usage*. 2019. URL: <https://www.frontiersin.org/journals/marine-science/articles/10.3389/fmars.2019.00502/full> (visited on 10/25/2026).
- [2] "A census of quality-controlled Biogeochemical-Argo float measurements". In: *Frontiers in Marine Science* (2023). DOI: 10.3389/fmars.2023.1233289.
- [3] "A Survey of Turbulence in Underwater Optical Wireless Communications". In: *Optik* 320, 172126 (2025). DOI: 10.1016/j.ijleo.2024.172126.
- [4] Eric Acome, Shane K. Mitchell, Timothy G. Morrissey, et al. "Hydraulically Amplified Self-Healing Electrostatic Actuators with Muscle-Like Performance". In: *Science* 359.6371 (2018), pp. 61–65. DOI: 10.1126/science.aao6139.
- [5] "Actuation and Locomotion of Miniature Underwater Robots: A Survey". In: *Engineering* 51 (2025), pp. 195–214. DOI: 10.1016/j.eng.2024.10.022.
- [6] Derya Akkaynak and Tali Treibitz. "A Revised Underwater Image Formation Model". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 6723–6732. DOI: 10.1109/CVPR.2018.00703.
- [7] Derya Akkaynak and Tali Treibitz. "Sea-Thru: A Method for Removing Water From Underwater Images". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 1682–1691. URL: https://openaccess.thecvf.com/content_CVPR_2019/html/Akkaynak_Sea-Thru_A_Method_for_Removing_Water_From_Underwater_Images_CVPR_2019_paper.html (visited on 10/25/2026).
- [8] Anas S. Alamoush and Aykut I. "Olçer. "Automated and Remote Engineering, Maintenance, and Repair in Maritime Autonomous Surface Ships (MASS)". In: *Journal of Shipping and Trade* 10, 20 (2025). DOI: 10.1186/s41072-025-00210-6.
- [9] Abid Ali et al. "Nature-Inspired Anti-Fouling Strategies for Combating Marine Biofouling". In: *Progress in Organic Coatings* 189, 108349 (2024). DOI: 10.1016/j.porgcoat.2024.108349.
- [10] American Bureau of Shipping. *Rules for Building and Classing Underwater Vehicles, Systems and Hyperbaric Facilities*. 2025. URL: https://ww2.eagle.org/content/dam/eagle/rules-and-guides/current/special_service/7-rules-for-building-and-classing-underwater-vehicles%2C-systems-and-hyperbaric-facilities-2025/7-uwvs-rules-jan25.pdf (visited on 10/25/2026).
- [11] "An Investigation of Innovative Strategies in Underwater Soft Robotics". In: *Engineering Science and Technology, an International Journal* 70, 102123 (2025). DOI: 10.1016/j.jestch.2025.102123.
- [12] Simon R. Anuszczyk and John O. Dabiri. "Electromechanical Enhancement of Live Jellyfish for Ocean Exploration". In: *Bioinspiration & Biomimetics* 19.2, 026018 (2024). DOI: 10.1088/1748-3190/ad277f.
- [13] Argo Data Management Team. *Argo User's Manual: Introduction and Data Modes*. URL: <https://oneargo.github.io/argo-format-user-manual/01-introduction.html> (visited on 10/25/2026).
- [14] Argo International Program. *Argo's Status*. URL: <https://argo.ucsd.edu/about/status/> (visited on 10/25/2026).
- [15] Argo International Programme. *Data FAQ*. URL: <https://argo.ucsd.edu/data/data-faq/> (visited on 10/25/2026).
- [16] Amy R. Baco, E. Brendan Roark, and Nicole B. Morgan. "Amid fields of rubble, scars, and lost gear, signs of recovery observed on seamounts on 30- to 40-year time scales". In: *Science Advances* 5.8, eaaw4513 (2019). DOI: 10.1126/sciadv.aaw4513.

- [17] Kelly J. Benoit-Bird et al. "Schrödinger's fish: Examining the robotic observer effect on pelagic animals". In: *Limnology and Oceanography: Methods* 21.9 (2023), pp. 563–580. DOI: 10.1002/lom3.10565.
- [18] Gunhild Elisabeth Berget et al. "Adaptive Sampling of Ocean Processes Using an AUV with a Gaussian Proxy Model". In: *IFAC-PapersOnLine* 51.29 (2018), pp. 238–243. DOI: 10.1016/j.ifacol.2018.09.509.
- [19] Jonathan Birch. "Animal sentience and the precautionary principle". In: *Animal Sentience* 2.16 (2017), p. 1. DOI: 10.51291/2377-7478.1200.
- [20] Jonathan Birch et al. *Review of the Evidence of Sentience in Cephalopod Molluscs and Decapod Crustaceans*. London School of Economics and Political Science, 2021. URL: <https://www.lse.ac.uk/business/consulting/reports/review-of-the-evidence-of-sentiences-in-cephalopod-molluscs-and-decapod-crustaceans> (visited on 10/25/2026).
- [21] Liushuai Cao et al. "Review on the Hydro- and Thermo-Dynamic Wakes of Underwater Vehicles in Linearly Stratified Fluid". In: *Journal of Marine Science and Engineering* 12.3 (2024), p. 490. DOI: 10.3390/jmse12030490.
- [22] Stephanie Russo Carroll et al. "The CARE Principles for Indigenous Data Governance". In: *Data Science Journal* 19, 43 (2020). DOI: 10.5334/dsj-2020-043.
- [23] Yunsai Chen et al. "Failure Modes and Reliability Analysis of Autonomous Underwater Vehicles: A Review". In: *Journal of Marine Science and Application* 24.5 (2025), pp. 900–924. DOI: 10.1007/s11804-025-00627-2.
- [24] Mainul Islam Chowdhury et al. "Next-Generation Underwater Localization: Artificial Intelligence-Based and Energy-Aware Approaches". In: *Applied Ocean Research* 165, 104842 (2025). DOI: 10.1016/j.apor.2025.104842.
- [25] Patryk Cieslak, Pere Ridao, and Mariusz Giergiel. "Autonomous Underwater Panel Operation by GIRONA500 UVMS: A Practical Approach to Autonomous Underwater Manipulation". In: *2015 IEEE International Conference on Robotics and Automation*. 2015, pp. 529–536. DOI: 10.1109/ICRA.2015.7139230.
- [26] "Collision Avoidance System for Autonomous Ships Based on Control Barrier Functions: Simulations and Experiments". In: *2026 European Control Conference*. 2026. URL: <https://ieeexplore.ieee.org/abstract/document/11625451> (visited on 10/25/2026).
- [27] Copernicus Marine Service. *Definition of the Model Data Types Available in the Copernicus Marine Service*. 2026. URL: <https://help.marine.copernicus.eu/en/articles/4575382-definition-of-the-model-data-types-available-in-the-copernicus-marine-service> (visited on 10/25/2026).
- [28] Ryan K. Cosner et al. "Measurement-Robust Control Barrier Functions: Certainty in Safety with Uncertainty in State". In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021. DOI: 10.1109/IR0551168.2021.9636584.
- [29] Robyn J. Crook. "Behavioral and neurophysiological evidence suggests affective pain experience in octopus". In: *iScience* 24.3, 102229 (2021). DOI: 10.1016/j.isci.2021.102229.
- [30] Darwin Core Maintenance Group. *Darwin Core List of Terms*. Oct. 13, 2026. URL: <https://dwc.tdwg.org/list/> (visited on 10/25/2026).
- [31] Emanuela Del Dottore, Romina Adhami, Ebrahim Shahabi, et al. "Peripheral Control Enabled by Distributed Sensing in an Octopus-Inspired Soft Robotic Arm for Autonomous Underwater Grasping". In: *Nature Machine Intelligence* 8 (2026), pp. 708–721. DOI: 10.1038/s42256-026-01230-y.
- [32] Peter J. Diggle, Raquel Menezes, and Ting-li Su. "Geostatistical Inference under Preferential Sampling". In: *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 59.2 (2010), pp. 191–232. DOI: 10.1111/j.1467-9876.2009.00701.x.
- [33] Tamsin Dobson et al. "Biofouling and Corrosion Rate of Welded Nickel Aluminium Bronze in Natural and Simulated Seawater". In: *Biofouling* 40.2 (2024), pp. 193–208. DOI: 10.1080/08927014.2024.2326067.
- [34] John Dupré and Daniel J. Nicholson. "A Manifesto for a Processual Philosophy of Biology". In: *Everything Flows: Towards a Processual Philosophy of Biology*. Ed. by Daniel J. Nicholson and John Dupré. Oxford University Press, 2018, pp. 3–46. DOI: 10.1093/oso/9780198779636.003.0001.
- [35] Ron Eglash. "An introduction to generative justice". In: *Teknokultura: Revista de Cultura Digital y Movimientos Sociales* 13.2 (2016), pp. 369–404. DOI: 10.5209/rev_TEKN.2016.v13.n2.52847.
- [36] Aly M. Eissa et al. "Trajectory tracking control of autonomous underwater vehicles: a review from classical methods to AI-based approaches". In: *International Journal of Dynamics and Control* 14, 32 (2026). DOI: 10.1007/s40435-025-01963-5.
- [37] James A. Estes and John F. Palmisano. "Sea Otters: Their Role in Structuring Nearshore Communities". In: *Science* 185.4156 (1974), pp. 1058–1060. DOI: 10.1126/science.185.4156.1058.
- [38] European Commission. *European Digital Twin Ocean*. 2026. URL: https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe/eu-missions-horizon-europe/restore-our-ocean-and-waters/european-digital-twin-ocean_en (visited on 10/25/2026).
- [39] European Commission. *Descriptors under the Marine Strategy Framework Directive: Descriptor 11, Energy, Including Underwater Noise*. URL: https://environment.ec.europa.eu/topics/marine-environment/descriptors-under-marine-strategy-framework-directive_en (visited on 10/25/2026).
- [40] European Digital Twin Ocean. *What is EDITO?* Aug. 26, 2026. URL: <https://help.edito.eu/en/articles/12569275-what-is-edito> (visited on 10/25/2026).
- [41] European Parliament and Council of the European Union. *Directive 2010/63/EU on the protection of animals used for scientific purposes*. Sept. 22, 2010. URL: <https://eur-lex.europa.eu/eli/dir/2010/63/oj/eng> (visited on 10/25/2026).
- [42] European Space Agency. *Marine Services*. 2026. URL: https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Marine_services (visited on 10/25/2026).
- [43] Dominic Flores, Sahib Sandhu, Alexander White, et al. "RoboNautilus: A Cephalopod-Inspired Soft Robotic Siphon for Underwater Propulsion". In: *npj Robotics* 3, 17 (2025). DOI: 10.1038/s44182-025-00035-2.
- [44] Carl Folke et al. "Adaptive Governance of Social-Ecological Systems". In: *Annual Review of Environment and Resources* 30 (2005), pp. 441–473. DOI: 10.1146/annurev.energy.30.050504.144511.
- [45] Thor I. Fossen. *Handbook of Marine Craft Hydrodynamics and Motion Control*. 2nd ed. Wiley, 2021. DOI: 10.1002/9781119575016.
- [46] Thor I. Fossen. *Fossen's Marine Craft Model*. URL: <https://www.fossen.biz/html/marineCraftModel.html> (visited on 10/25/2026).
- [47] Luke Freyhof, Matticus Brown, Sergio Araujo-Estrada, et al. "Ray-Inspired Robots: Recent Advances in Actuation and Control". In: *npj Robotics* 4, 1 (2026). DOI: 10.1038/s44182-025-00064-x.
- [48] Kevin C. Galloway et al. "Soft Robotic Grippers for Biological Sampling on Deep Reefs". In: *Soft Robotics* 3.1 (2016), pp. 23–33. DOI: 10.1089/soro.2015.0019.
- [49] Jason-Zois Gazis et al. "Monitoring benthic plumes, sediment redeposition and seafloor imprints caused by deep-sea polymetallic nodule mining". In: *Nature Communications* 16, 1229 (2025). DOI: 10.1038/s41467-025-56311-0.
- [50] General Bathymetric Chart of the Oceans. *Release of the GEBCO_2026 Grid*. Apr. 27, 2026. URL: <https://www.gebco.net/news/release-gebco2026-grid> (visited on 10/25/2026).
- [51] Maxime Geoffroy et al. "Pelagic organisms avoid white, blue, and red artificial light from scientific instruments". In: *Scientific Reports* 11, 14941 (2021). DOI: 10.1038/s41598-021-94355-6.

- [52] Brian P. Gerkey and Maja J. Mataric. "A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems". In: *The International Journal of Robotics Research* 23.9 (2004), pp. 939–954. DOI: 10.1177/0278364904045564.
- [53] Global Ocean Biogeochemistry Array. *Data FAQ*. URL: <https://www.go-bgc.org/data/data-faq> (visited on 10/25/2026).
- [54] Global Ocean Observing System. *Essential Ocean Variables*. 2026. URL: <https://goosocan.org/what-we-do/framework/essential-ocean-variables/> (visited on 10/25/2026).
- [55] Global Ocean Observing System. *Ocean Best Practices*. 2026. URL: <https://goosocan.org/our-work/ocean-best-practices/> (visited on 10/25/2026).
- [56] Caren S. Goldberg et al. "Critical considerations for the application of environmental DNA methods to detect aquatic species". In: *Methods in Ecology and Evolution* (2016). DOI: 10.1111/2041-210X.12595.
- [57] Yifeng Gong, Mingyu Pan, and Kathryn A. Daltorio. "Using a Small Hexapod Robot to Pick Up Large Cylinders for Munitions Response". In: *Journal of Field Robotics* 42.5 (2025), pp. 1679–1695. DOI: 10.1002/rob.22482.
- [58] S. Yu. Gordileeva et al. "Control of movement of underwater swimmers: Animals, simulated animates and swimming robots". In: (2023). URL: <https://pubmed.ncbi.nlm.nih.gov/38072505/> (visited on 10/25/2026).
- [59] Jesus Guerrero et al. "Robust super-twisting-based disturbance observer for autonomous underwater vehicles: Design, stability analysis, and real-time experiments". In: *Robotics and Autonomous Systems* 184, 104859 (2025). DOI: 10.1016/j.robot.2024.104859.
- [60] Kangjian He et al. "Overview of Research Progress on Numerical Simulation Methods for Turbulent Flows Around Underwater Vehicles". In: *Journal of Marine Science and Application* 23 (2024), pp. 1–22. DOI: 10.1007/s11804-024-00403-8.
- [61] Charlotte Hess and Elinor Ostrom, eds. *Understanding Knowledge as a Commons: From Theory to Practice*. MIT Press, 2007. ISBN: 9780262083577. DOI: 10.7551/mitpress/6980.001.0001.
- [62] Crawford Stanley Holling. "Resilience and Stability of Ecological Systems". In: *Annual Review of Ecology and Systematics* 4 (1973), pp. 1–23. DOI: 10.1146/annurev.es.04.110173.000245.
- [63] Heting Hong et al. "A Review of Experimental Assessment Processes of Material Resistance to Marine and Freshwater Biofouling". In: *Journal of Environmental Management* 357, 120766 (2024). DOI: 10.1016/j.jenvman.2024.120766.
- [64] Lin Hong, Xin Wang, and De-Sheng Zhang. "CFD-based hydrodynamic performance investigation of autonomous underwater vehicles: A survey". In: *Ocean Engineering* 305 (2024), p. 117911. DOI: 10.1016/j.oceaneng.2024.117911.
- [65] Junzhe Hu, Jeong Hun Lee, Tiancheng Wu, et al. "Programmable Online Stiffness Modulation for Optimized Aquatic Locomotion". In: *npj Robotics* 4, 44 (2026). DOI: 10.1038/s44182-026-00099-8.
- [66] Intergovernmental Oceanographic Commission of UNESCO. *The ocean observations powering safer communities and a sustainable future*. 2026. URL: <https://www.ioc.unesco.org/en/articles/essential-ocean-variables> (visited on 10/25/2026).
- [67] Intergovernmental Oceanographic Commission of UNESCO. *Ocean Biodiversity Information System*. URL: <https://www.ioc.unesco.org/en/obis> (visited on 10/25/2026).
- [68] Intergovernmental Panel on Climate Change. *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*. IPCC, 2019. URL: <https://www.ipcc.ch/srocc/> (visited on 10/25/2026).
- [69] Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Methodological Assessment Report on the Diverse Values and Valuation of Nature*. IPBES, 2022. URL: <https://www.ipbes.net/the-values-assessment> (visited on 10/25/2026).
- [70] International Hydrographic Organization. *Global seabed mapping reaches new milestone as five million square kilometres added in a year*. Apr. 28, 2026. URL: <https://iho.int/en/global-seabed-mapping-reaches-new-milestone-as-five-million-square-kilometres-added-in-a-year> (visited on 10/25/2026).
- [71] International Maritime Organization. *Marine Environment Protection Committee, Eighty-Second Session: Underwater Radiated Noise*. 2024. URL: <https://www.imo.org/en/mediacentre/meetingsummaries/pages/mepc-82nd-session.aspx> (visited on 10/25/2026).
- [72] International Maritime Organization. *IMO Adopts First Global Code for Autonomous Ships*. May 22, 2026. URL: <https://www.imo.org/en/mediacentre/pressbriefings/pages/imo-adopts-mass-code.aspx> (visited on 10/26/2026).
- [73] International Maritime Organization. *International Convention for the Prevention of Pollution from Ships (MARPOL)*. URL: <https://www.imo.org/en/about/conventions/p.s/international-convention-for-the-prevention-of-pollution-from-ships-%28marpol%29.aspx> (visited on 10/25/2026).
- [74] International Maritime Organization. *International Convention on the Control of Harmful Anti-fouling Systems on Ships*. URL: <https://www.imo.org/en/about/conventions/pages/international-convention-on-the-control-of-harmful-anti-fouling-systems-on-ships-%28afss%29.aspx> (visited on 10/25/2026).
- [75] International Maritime Organization. *Ship Noise: Revised Guidelines for the Reduction of Underwater Radiated Noise from Shipping*. URL: <https://www.imo.org/en/mediacentre/hottopics/pages/noise.aspx> (visited on 10/25/2026).
- [76] International Organization for Standardization. *ISO 12944-9:2018: Protective Paint Systems and Laboratory Performance Test Methods for Offshore and Related Structures*. 2018. URL: <https://www.iso.org/standard/64832.html> (visited on 10/25/2026).
- [77] International Organization for Standardization. *ISO/IEC Guide 98-1:2024: Guide to the Expression of Uncertainty in Measurement—Part 1: Introduction*. 2024. URL: <https://www.iso.org/standard/82708.html> (visited on 10/25/2026).
- [78] International Seabed Authority. *About the International Seabed Authority*. URL: <https://isa.org.jm/about-isa/> (visited on 10/25/2026).
- [79] Japan Agency for Marine-Earth Science and Technology. *Ocean Observation Column: Depth and Pressure*. 2021. URL: <https://www.jamstec.go.jp/jpr/topics/column-20210428/> (visited on 10/25/2026).
- [80] Hao Jiang et al. "Hydrodynamic Mechanisms, Fluid–Structure Interaction, and Material Selection in Underwater Bio-Inspired Robots: A Review". In: *Fluid Dynamics & Materials Processing* 22.6 (2026). DOI: 10.32604/fdmp.2026.082152.
- [81] Yuehan Jiang and Bernd-Christian Renner. "Low-Cost Underwater Swarm Acoustic Localization: A Review". In: *IEEE Access* 12 (2024), pp. 25779–25796. DOI: 10.1109/ACCESS.2024.3357359.
- [82] Daniel O. B. Jones et al. "Biological responses to disturbance from simulated deep-sea polymetallic nodule mining". In: *PLOS ONE* 12.2, e0171750 (2017). DOI: 10.1371/journal.pone.0171750.
- [83] Daniel O. B. Jones et al. "Long-term Impact and Biological Recovery in a Deep-Sea Mining Track". In: *Nature* 642 (2025), pp. 112–118. DOI: 10.1038/s41586-025-08921-3.
- [84] Dong-Hwi Kim et al. "Data-driven stochastic nonlinear model predictive control for 6 DOF underwater vehicle under unknown disturbance and uncertainty". In: *Ocean Engineering* 318, 120067 (2025). DOI: 10.1016/j.oceaneng.2024.120067.
- [85] Mithilesh Kumar Koiri et al. "A comprehensive review of bio-inspired swimming in robotic fishes". In: *Sensors and Actuators A: Physical* 394, 116913 (2025). DOI: 10.1016/j.sna.2025.116913.
- [86] Ebrahim M. Kolahdouz et al. "A sharp interface Lagrangian–Eulerian method for flexible-body fluid–structure interaction". In: *Journal of Computational Physics* 488 (2023), p. 112174. DOI: 10.1016/j.jcp.2023.112174.
- [87] Andreas Krause, Ajit Singh, and Carlos Guestrin. "Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies". In: *Journal of Machine Learning Research* 9 (2008), pp. 235–284. URL: <https://www.jmlr.org/papers/v9/krause08a.html>.

- [88] Chanhong Lee et al. “Octopus-Inspired Adhesives with Switchable Attachment to Challenging Underwater Surfaces”. In: *Advanced Science* 12.1, 2407588 (2025). DOI: 10.1002/advs.202407588.
- [89] Dong-Kyu Lee and Nikolai Maximenko. “Surface Drifters and Ocean Dynamics: A Review of Technological Advancements and Scientific Contributions”. In: *Ocean Science Journal* 60, 22 (2025). DOI: 10.1007/s12601-025-00217-x.
- [90] HoSaeng Lee, HyeonJu Kim, and SeungTaek Lim. “Experimental Performance Analysis of Ocean Thermal Energy Conversion Using Combined High-Temperature Waste Heat and Deep Seawater”. In: *Journal of Ocean Engineering and Marine Energy* (2026). DOI: 10.1007/s40722-026-00513-1.
- [91] Anaëlle J. Lemasson et al. “A global meta-analysis of ecological effects from offshore marine artificial structures”. In: *Nature Sustainability* 7 (2024), pp. 485–495. DOI: 10.1038/s41893-024-01311-z.
- [92] Simon A. Levin. “The Problem of Pattern and Scale in Ecology”. In: *Ecology* 73.6 (1992), pp. 1943–1967. DOI: 10.2307/1941447.
- [93] Guorui Li, Xiangping Chen, Fanghao Zhou, et al. “Self-Powered Soft Robot in the Mariana Trench”. In: *Nature* 591 (2021), pp. 66–71. DOI: 10.1038/s41586-020-03153-z.
- [94] Lei Li, Boyang Qin, Wenzhuo Gao, et al. “Bioinspired Underwater Soft Robots: From Biology to Robotics and Back”. In: *npj Robotics* 4, 25 (2026). DOI: 10.1038/s44182-026-00088-x.
- [95] Yixin Li et al. “Multi-Material Embedded 3D Printing for One-Step Manufacturing of Multifunctional Components in Soft Robotics”. In: *Additive Manufacturing* 85, 104178 (2024). DOI: 10.1016/j.addma.2024.104178.
- [96] Xiaozhu Lin, Xiaopei Liu, and Yang Wang. “Learning Agile Swimming: An End-to-End Approach Without CPGs”. In: *IEEE Robotics and Automation Letters* 10.2 (2025), pp. 1992–1999. DOI: 10.1109/LRA.2025.3527757.
- [97] S. D. Ling et al. “Global Regime Shift Dynamics of Catastrophic Sea Urchin Overgrazing”. In: *Philosophical Transactions of the Royal Society B* 370.1659, 20130269 (2015). DOI: 10.1098/rstb.2013.0269.
- [98] Jiaqi Liu, Zhichao Zhu, and Li Wen. “Underwater Soft Arm Grasping with Simplified Control Using Octopus-Inspired Bending Propagation”. In: *npj Robotics* 4, 2 (2026). DOI: 10.1038/s44182-025-00066-9.
- [99] Owen R. Liu and Steven D. Gaines. “Environmental context dependency in species interactions”. In: *Proceedings of the National Academy of Sciences* 119.36, e2118539119 (2022). DOI: 10.1073/pnas.2118539119.
- [100] Salvador López-Barajas et al. “Towards Autonomous Coordination of Two I-AUVs in Submarine Pipeline Assembly”. In: *Journal of Marine Science and Engineering* 13.8 (2025), p. 1490. DOI: 10.3390/jmse13081490.
- [101] Suxiang Lu et al. “Advances in environmental DNA monitoring: standardization, automation, and emerging technologies in aquatic ecosystems”. In: *Science China Life Sciences* 67.7 (2024), pp. 1368–1384. DOI: 10.1007/s11427-023-2493-5.
- [102] Anthony J. Lucio et al. “Obtaining Accurate, High-Frequency and Long-Term Seawater pH Data by Using Coupled Lab-on-Chip and Optode Sensing Technologies”. In: *Ocean Science* 22 (2026), pp. 1609–1623. DOI: 10.5194/os-22-1609-2026.
- [103] Gianluca Manduca et al. “A Bioinspired Control Strategy Ensures Maneuverability and Adaptability for Dynamic Environments in an Underactuated Robotic Fish”. In: *Journal of Intelligent & Robotic Systems* 110, 69 (2024). DOI: 10.1007/s10846-024-02080-9.
- [104] Laura F. B. Marangoni et al. “Impacts of Artificial Light at Night in Marine Ecosystems—A Review”. In: *Global Change Biology* 28.18 (2022), pp. 5346–5367. DOI: 10.1111/gcb.16264.
- [105] Belén Martín Míguez, Emma Heslop, Narissa Bax, et al. “GOOS Essential Ocean Variables: the backbone of a sustained and evolving global ocean observing system”. In: *Frontiers in Marine Science* 13, 1737002 (2026). DOI: 10.3389/fmars.2026.1737002.
- [106] Seth McCammon et al. “Ocean Front Detection and Tracking Using a Team of Heterogeneous Marine Vehicles”. In: *Journal of Field Robotics* 38.6 (2021), pp. 854–881. DOI: 10.1002/rob.22014.
- [107] David J. Mellor et al. “The 2020 Five Domains Model: Including Human–Animal Interactions in Assessments of Animal Welfare”. In: *Animals* 10.10, 1870 (2020). DOI: 10.3390/ani10101870.
- [108] Yan Meng et al. “Development and 3-D Path-Following Control of an Agile Robotic Manta With Flexible Pectoral Fins”. In: *IEEE Transactions on Cybernetics* 54.6 (2024), pp. 3727–3738. DOI: 10.1109/TCYB.2023.3252589.
- [109] Microsoft Research. *Project Natick: Phase 2 Timeline*. 2026. URL: <https://www.microsoft.com/en-us/research/project/natick/phase-2-timeline/> (visited on 10/25/2026).
- [110] Microsoft Research. *Project Natick: Research Overview and Phase 2 Results*. 2026. URL: <https://www.microsoft.com/en-us/research/project/natick/overview/> (visited on 10/25/2026).
- [111] MIT Department of Earth, Atmospheric and Planetary Sciences. *Ocean Modeling*. URL: <https://paocweb.mit.edu/research/ocean/ocean-modeling> (visited on 10/25/2026).
- [112] MITgcm Development Team. *MITgcm User Manual: Equations of Motion for the Ocean*. URL: https://mitgcm.readthedocs.io/en/latest/overview/eqn_motion_ocn.html (visited on 10/25/2026).
- [113] Monterey Bay Aquarium Research Institute. *Long-range AUV (LRAUV): Tethys*. 2026. URL: <https://www.mbari.org/technology/long-range-auv-lrauv/> (visited on 10/25/2026).
- [114] Monterey Bay Aquarium Research Institute. *Environmental Genomics*. URL: <https://www.mbari.org/project/environmental-genomics/> (visited on 10/25/2026).
- [115] Monterey Bay Aquarium Research Institute. *Environmental Sample Processor (ESP)*. URL: <https://www.mbari.org/technology/environmental-sample-processor-esp/> (visited on 10/25/2026).
- [116] Monterey Bay Aquarium Research Institute. *Long-range AUV (LRAUV): Research Group and Fleet*. URL: <https://www.mbari.org/team/lrauv/> (visited on 10/25/2026).
- [117] Monterey Bay Aquarium Research Institute. *Marine Robots’ Onboard Intelligence for Capturing Oceanographic Features*. URL: <https://www.mbari.org/team/marine-robots-onboard-intelligence-for-capturing-oceanographic-features/> (visited on 10/25/2026).
- [118] Allison R. Moreno and Adam C. Martiny. “Ecological Stoichiometry of Ocean Plankton”. In: *Annual Review of Marine Science* 10 (2018), pp. 43–69. DOI: 10.1146/annurev-marine-121916-063126.
- [119] NASA Earth Observatory. *What Are Phytoplankton?* URL: <https://science.nasa.gov/earth/earth-observatory/what-are-phytoplankton/> (visited on 10/25/2026).
- [120] National Deep Submergence Facility, Woods Hole Oceanographic Institution. *Sentry: Specifications*. URL: <https://ndsf.whoi.edu/sentry/specifications-sentry/> (visited on 10/25/2026).
- [121] National Institute of Standards and Technology. *Measurement Uncertainty*. 2026. URL: <https://www.nist.gov/itl/sed/topic-areas/measurement-uncertainty> (visited on 10/25/2026).
- [122] National Marine Fisheries Service. *2024 Update to: Technical Guidance for Assessing the Effects of Anthropogenic Sound on Marine Mammal Hearing (Version 3.0): Underwater and In-Air Criteria for Onset of Auditory Injury and Temporary Threshold Shifts*. NOAA Technical Memorandum NMFS-OPR-71. NOAA Fisheries, 2024. DOI: 10.25923/z6h7-5p93.
- [123] National Oceanic and Atmospheric Administration. *How Does Pressure Change with Ocean Depth?* 2026. URL: <https://oceanservice.noaa.gov/facts/pressure.html> (visited on 10/25/2026).

- 8865 [124] National Oceanic and Atmospheric Administration. *Tides*. URL: <https://www.noaa.gov/education/resource-collections/ocean-coasts/tides> (visited on 10/25/2026).
- 8866
- 8867 [125] National Oceanography Centre. *Gliders*. URL: <https://www.noc.ac.uk/our-work/marine-technology/ocean-vehicles/gliders> (visited on 10/25/2026).
- 8868
- 8869 [126] NATO. STANAG 4748: *Digital Underwater Signalling Standard for Network Node Discovery and Interoperability*. 2017. URL: https://quicksearch.dla.mil/qsDocDetails.aspx?ident_number=284318 (visited on 10/25/2026).
- 8870
- 8871 [127] Vuong Nguyen-Van and Hongjian Du. "Underwater 3D Concrete Printing: Automation Challenges, Process-Material Coupling, and Future Construction Pathways". In: *Journal of Building Engineering* 120, 115490 (2026). DOI: 10.1016/j.jobe.2026.115490.
- 8872
- 8873 [128] NOAA Fisheries. *Scientific Research and Enhancement Permits for Marine Mammals*. July 21, 2026. URL: <https://www.fisheries.noaa.gov/permit/scientific-research-and-enhancement-permits-marine-mammals> (visited on 10/25/2026).
- 8874
- 8875 [129] NOAA Fisheries. *Marine Life Viewing Guidelines*. URL: <https://www.fisheries.noaa.gov/topic/marine-life-viewing-guidelines> (visited on 10/25/2026).
- 8876
- 8877 [130] NOAA Fisheries. *Ocean Noise*. URL: <https://www.fisheries.noaa.gov/national/science-data/ocean-noise> (visited on 10/25/2026).
- 8878
- 8879 [131] NOAA Fisheries. *Technologies for Passive Acoustic Research*. URL: <https://www.fisheries.noaa.gov/new-england-mid-atlantic/science-data/passive-acoustic-technologies> (visited on 10/25/2026).
- 8880
- 8881 [132] NOAA National Ocean Service. *What's the difference between a tide and a current?* Sept. 23, 2026. URL: <https://oceanservice.noaa.gov/facts/tidescurrents.html> (visited on 10/25/2026).
- 8882
- 8883 [133] NOAA Ocean Exploration. *Sampling Procedures Manual*. 2024. URL: https://oeab.noaa.gov/wp-content/uploads/2024/10/Sampling_Manual_noaa_60804_DS1.pdf (visited on 10/25/2026).
- 8884
- 8885 [134] NOAA Ocean Exploration. *How Does Pressure Impact Animals in the Ocean?* 2026. URL: <https://oceanexplorer.noaa.gov/ocean-fact/animal-pressure/> (visited on 10/25/2026).
- 8886
- 8887 [135] NOAA Ocean Exploration. *Multibeam Sonar*. URL: <https://oceanexplorer.noaa.gov/technology/sonar-multibeam/> (visited on 10/25/2026).
- 8888
- 8889 [136] NOAA Ocean Exploration. *Ocean Exploration Technology: How Robots Are Uncovering the Mysteries of the Deep*. URL: <https://oceanexplorer.noaa.gov/explainers/technology/> (visited on 10/25/2026).
- 8890
- 8891 [137] NOAA Ocean Exploration. *Remotely Operated Vehicle Deep Discoverer*. URL: <https://oceanexplorer.noaa.gov/technology/subs-deep-discoverer/> (visited on 10/25/2026).
- 8892
- 8893 [138] NOAA Ocean Exploration. *Uncrewed Surface Vessels*. URL: <https://oceanexplorer.noaa.gov/technology/usv/> (visited on 10/25/2026).
- 8894
- 8895 [139] NOAA Office of General Counsel. *Submarine Cables: Domestic Regulation and Environmental Impact Studies*. URL: <https://www.noaa.gov/general-counsel/gc-international-section/submarine-cables-domestic-regulation> (visited on 10/25/2026).
- 8896
- 8897 [140] NOAA Pacific Marine Environmental Laboratory. *Uncrewed Surface Vehicle Hurricane Observations*. URL: <https://www.pmel.noaa.gov/usv-hurricane/> (visited on 10/25/2026).
- 8898
- 8899 [141] Ocean Best Practices System. *IOC-OBPS Repository*. URL: <https://oceanbestpractices.org/repository/> (visited on 10/25/2026).
- 8900
- 8901 [142] Ocean Networks Canada. *Observatories*. 2026. URL: <https://www.oceannetworks.ca/observatories/> (visited on 10/25/2026).
- 8902
- 8903 [143] Ocean Networks Canada. *Cabled Networks*. URL: <https://www.oceannetworks.ca/observatories/physical-infrastructure/cabled-networks/> (visited on 10/25/2026).
- 8904
- 8905 [144] Ocean Networks Canada. *Expeditions*. URL: <https://www.oceannetworks.ca/expeditions/> (visited on 10/25/2026).
- 8906
- 8907 [145] Ocean Observatories Initiative. *OOI Biogeochemical Sensors Data: Best Practices and User Guide*. 2025. URL: <https://oceanobservatories.org/wp-content/uploads/2025/07/00I-Biogeochemical-Sensor-Data-Best-Practices-and-User-Guide-Version-1.1.1-with-July-dated-GOOS-Endorsement-Certificate-on-p2-Updated-DOI.pdf> (visited on 10/25/2026).
- 8908
- 8909 [146] Ocean Observatories Initiative. *Data Explorer Documentation: Introduction*. URL: <https://dataexplorer.oceanobservatories.org/help/overview.html> (visited on 10/25/2026).
- 8910
- 8911 [147] Ocean Observatories Initiative. *Gliders*. URL: <https://oceanobservatories.org/marine-technologies/gliders/> (visited on 10/25/2026).
- 8912
- 8913 [148] Ocean Observatories Initiative. *Quality Control*. URL: <https://oceanobservatories.org/quality-control/> (visited on 10/25/2026).
- 8914
- 8915 [149] F. John Odling-Smee, Kevin N. Laland, and Marcus W. Feldman. *Niche Construction: The Neglected Process in Evolution*. Princeton University Press, 2003.
- 8916
- 8917 [150] Organisation for Economic Co-operation and Development. *Transfer Pricing*. URL: <https://www.oecd.org/en/topics/transfer-pricing.html> (visited on 10/25/2026).
- 8918
- 8919 [151] Narcis Palomeras et al. "I-AUV Docking and Panel Intervention at Sea". In: *Sensors* 16.10 (2016), p. 1673. DOI: 10.3390/s16101673.
- 8920
- 8921 [152] Myungsun Park et al. "Bioinspired Flow Sensor Enables Underwater Robots to Estimate Motion and Detect Flow Structure". In: *Science Advances* 12.23, ead2847 (2026). DOI: 10.1126/sciadv.aed2847.
- 8922
- 8923 [153] Rolf Pfeifer, Fumiya Iida, and Gabriel Gómez. "Morphological Computation for Adaptive Behavior and Cognition". In: *International Congress Series* (2006), pp. 22–29. DOI: 10.1016/j.ics.2005.12.080.
- 8924
- 8925 [154] Huy Pham et al. *Characterizing Wildlife Response to Biomimetic and Conventional Underwater Vehicles*. Sept. 18, 2026. arXiv: 2609.22594. URL: <https://arxiv.org/abs/2609.22594> (visited on 10/25/2026).
- 8926
- 8927 [155] Dimitrios Piromalis, Christos Alexandris, and Panagiotis Papageorgas. "Underwater Optical Imaging Systems: A Comprehensive Review of Technologies, Architectures and Mission-Driven Design Trade-Offs". In: *Sensors* 26.19, 6127 (2026). DOI: 10.3390/s26196127.
- 8928
- 8929 [156] Enrico Pirotta et al. "Understanding the Population Consequences of Disturbance". In: *Ecology and Evolution* 8.19 (2018), pp. 9934–9946. DOI: 10.1002/ece3.4458.
- 8930
- 8931 [157] Xinda Qi et al. "Design and Nonlinear Modeling of a Modular Cable-Driven Soft Robotic Arm". In: *IEEE/ASME Transactions on Mechatronics* 29.4 (2024), pp. 3083–3091. DOI: 10.1109/TMECH.2024.3402609.
- 8932
- 8933 [158] Haitao Qing et al. "Spontaneous Snapping-Induced Jet Flows for Fast, Maneuverable Surface and Underwater Soft Flapping Swimmer". In: *Science Advances* 10.49, eadq4222 (2024). DOI: 10.1126/sciadv.adq4222.
- 8934
- 8935 [159] Kelsi M. Rutledge et al. "Biohybrid Robotic Jellyfish for Swimming-Enhanced Vertical Ocean Profiling". In: *Biomimetics* 11.5, 325 (2026). DOI: 10.3390/biomimetics11050325.
- 8936
- 8937 [160] Marten Scheffer et al. "Catastrophic Shifts in Ecosystems". In: *Nature* 413 (2001), pp. 591–596. DOI: 10.1038/35098000.
- 8938
- 8939 [161] Ingrid Schjølberg and Olav Egeland. "Motion Control of Underwater Vehicle-Manipulator Systems Using Feedback Linearization". In: *Modeling, Identification and Control* 17.1 (1996), pp. 17–26. DOI: 10.4173/mic.1996.1.2.

- 8937 [162] Ingrid Schjølberg and Thor I. Fossen. “Modelling and Control of Underwater Vehicle-Manipulator Systems”. In: *Proceedings of the 3rd Conference on Marine Craft Maneuvering and Control*. 1994. URL: <https://www.fossen.biz/publications/1994%20Schj%C3%B8lberg%20and%20Fossen%20MCMC94.pdf> (visited on 10/25/2026).
- 8938
- 8939
- 8940 [163] David Schlosberg. *Defining Environmental Justice: Theories, Movements, and Nature*. Oxford University Press, 2007. DOI: 10.1093/acprof:oso/9780199286294.001.0001.
- 8941
- 8942 [164] Alexandra K. Schnell et al. “Sentience in cephalopod molluscs: an updated assessment”. In: *Biological Reviews* 101.3 (2026), pp. 1311–1333. DOI: 10.1002/brv.70125.
- 8943
- 8944 [165] Sea-Bird Scientific. *How Often Do I Need to Have My Instrument and Auxiliary Sensors Recalibrated?* 2026. URL: <https://www.seabird.com/faqs/1834> (visited on 10/25/2026).
- 8945
- 8946 [166] “Seamless Underwater Navigation with Limited Doppler Velocity Log Measurements”. In: *IEEE Transactions on Intelligent Vehicles* (2024). DOI: 10.1109/TIV.2024.3457772.
- 8947
- 8948 [167] Secretariat of the Convention on Biological Diversity. *Nagoya Protocol on Access to Genetic Resources and the Fair and Equitable Sharing of Benefits Arising from Their Utilization*. 2010. URL: <https://www.cbd.int/abs/doc/protocol/nagoya-protocol-en.pdf> (visited on 10/25/2026).
- 8949
- 8950
- 8951 [168] Erik Simon-Lledó et al. “Biological effects 26 years after simulated deep-sea mining”. In: *Scientific Reports* 9, 8040 (2019). DOI: 10.1038/s41598-019-44492-w.
- 8952
- 8953 [169] Lynne U. Sneddon. “Evolution of Nociception and Pain: Evidence from Fish Models”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 374.1785, 20190290 (2019). DOI: 10.1098/rstb.2019.0290.
- 8954
- 8955 [170] Fredrik Fogh Sørensen et al. “The Localization Problem for Underwater Vehicles: An Overview of Operational Solutions”. In: *Ocean Engineering* 330, 121173 (2025). DOI: 10.1016/j.oceaneng.2025.121173.
- 8956
- 8957 [171] Brandon L. Southall et al. “Marine Mammal Noise Exposure Criteria: Updated Scientific Recommendations for Residual Hearing Effects”. In: *Aquatic Mammals* 45.2 (2019), pp. 125–232. DOI: 10.1578/AM.45.2.2019.125.
- 8958
- 8959 [172] Stanford Robotics Laboratory. *OceanOneK*. URL: <https://robotics.stanford.edu/ocean-one-k.html> (visited on 10/25/2026).
- 8960
- 8961 [173] Liv Stranddorf, Francesca Verones, and Stig Irving Olsen. “Benefits and impacts of offshore wind farms on benthic marine biodiversity in a life cycle assessment context”. In: *The International Journal of Life Cycle Assessment* 30 (2025), pp. 2068–2085. DOI: 10.1007/s11367-025-02504-0.
- 8962
- 8963 [174] Yen-Fang Su et al. “Towards Scalable Underwater Concrete Additive Manufacturing for Marine and Aquatic Settings”. In: *Additive Manufacturing Letters* 18, 100397 (2026). DOI: 10.1016/j.addlet.2026.100397.
- 8964
- 8965 [175] Jia Sun et al. “A Review of Touching-Based Underwater Robotic Perception and Manipulation”. In: *Machines* 13.1 (2025), p. 41. DOI: 10.3390/machines13010041.
- 8966
- 8967 [176] Linxiang Sun et al. “Underwater Robots and Key Technologies for Operation Control”. In: *Cyborg and Bionic Systems* 5 (2024). DOI: 10.34133/cbsystems.0089.
- 8968
- 8969 [177] Darryn Sward, Jacquomo Monk, and Neville Barrett. “A Systematic Review of Remotely Operated Vehicle Surveys for Visually Assessing Fish Assemblages”. In: *Frontiers in Marine Science* 6, 134 (2019). DOI: 10.3389/fmars.2019.00134.
- 8970
- 8971 [178] Yu Jun Tan, Gianmarco Mengaldo, and Cecilia Laschi. “Artificial Muscles for Underwater Soft Robots: Materials and Their Interactions”. In: *Annual Review of Condensed Matter Physics* 15 (2024), pp. 45–61. DOI: 10.1146/annurev-conmatphys-032822-041146.
- 8972
- 8973 [179] “The JANUS Underwater Communications Standard: From Promulgation to Present”. In: *OCEANS 2024 – Singapore*. 2024. DOI: 10.1109/OCEANS51537.2024.10682419.
- 8974
- 8975 [180] U.S. Integrated Ocean Observing System. *Quality Assurance/Quality Control of Real Time Oceanographic Data (QARTOD)*. 2026. URL: <https://ioos.noaa.gov/project/qartod/> (visited on 10/25/2026).
- 8976
- 8977 [181] UK Argo. *Technology: How Do the Floats Work?* 2026. URL: <https://www.ukargo.net/technology> (visited on 10/25/2026).
- 8978
- 8979 [182] UK Argo. *Data Quality Control*. URL: https://www.ukargo.net/data/quality_control/ (visited on 10/25/2026).
- 8980
- 8981 [183] UK Argo. *Real-Time Data*. URL: <https://www.ukargo.net/data/real-time-data> (visited on 10/25/2026).
- 8982
- 8983 [184] A. J. Underwood. “Beyond BACI: The Detection of Environmental Impacts on Populations in the Real, but Variable, World”. In: *Journal of Experimental Marine Biology and Ecology* 161.2 (1992), pp. 145–178. DOI: 10.1016/0022-0981(92)90094-Q.
- 8984
- 8985 [185] A. J. Underwood. “On Beyond BACI: Sampling Designs That Might Reliably Detect Environmental Disturbances”. In: *Ecological Applications* 4.1 (1994), pp. 3–15. DOI: 10.2307/1942110.
- 8986
- 8987 [186] UNESCO. *UNESCO Recommendation on Open Science*. 2021. URL: <https://www.unesco.org/en/open-science/about> (visited on 10/25/2026).
- 8988
- 8989 [187] United Kingdom Parliament. *Animal Welfare (Sentience) Act 2022, Section 5*. 2022. URL: <https://www.legislation.gov.uk/ukpga/2022/22/section/5/enacted> (visited on 10/25/2026).
- 8990
- 8991 [188] United Nations. *Agreement on Marine Biological Diversity of Areas beyond National Jurisdiction*. 2026. URL: <https://www.un.org/bbnjagreement/en> (visited on 10/25/2026).
- 8992
- 8993 [189] United Nations. *Text of the BBNJ Agreement*. URL: <https://www.un.org/bbnjagreement/en/bbnj-agreement/text-bbnj-agreement> (visited on 10/25/2026).
- 8994
- 8995 [190] United Nations. *United Nations Convention on the Law of the Sea: Part V, Exclusive Economic Zone*. URL: https://www.un.org/Depts/los/convention_agreements/texts/unclos/part5.htm (visited on 10/25/2026).
- 8996
- 8997 [191] United Nations. *United Nations Convention on the Law of the Sea: Part VI, Continental Shelf*. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/part6.htm (visited on 10/25/2026).
- 8998
- 8999 [192] United Nations. *United Nations Convention on the Law of the Sea: Part VII, High Seas*. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/part7.htm (visited on 10/25/2026).
- 9000
- 9001 [193] United Nations. *United Nations Convention on the Law of the Sea: Part XI, Section 1, The Area*. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/part11-1.htm (visited on 10/25/2026).
- 9002
- 9003 [194] United Nations. *United Nations Convention on the Law of the Sea: Part XI, Section 2, Principles Governing the Area*. URL: https://www.un.org/Depts/los/convention_agreements/texts/unclos/part11-2.htm (visited on 10/25/2026).
- 9004
- 9005 [195] United Nations. *United Nations Convention on the Law of the Sea: Part XII, Protection and Preservation of the Marine Environment*. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/part12.htm (visited on 10/25/2026).
- 9006
- 9007 [196] United Nations. *United Nations Convention on the Law of the Sea: Part XIII, Marine Scientific Research*. URL: https://www.un.org/depts/los/convention_agreements/texts/unclos/part13.htm (visited on 10/25/2026).
- 9008
- 9009 [197] United Nations Conference on Environment and Development. *Rio Declaration on Environment and Development*. 1992. URL: <https://www.cbd.int/doc/ref/rio-declaration.shtml> (visited on 10/25/2026).
- 9010
- 9011 [198] United Nations Economic Commission for Europe. *Convention on Access to Information, Public Participation in Decision-Making and Access to Justice in Environmental Matters*. 1998. URL: <https://unece.org/environment-policy/public-participation/aarhus-convention/text> (visited on 10/25/2026).

- [199] United Nations Secretary-General's Scientific Advisory Board. *Science Brief: Deep-Sea Mining*. May 19, 2025. URL: <https://www.un.org/scientific-advisory-board/en/deep-sea-mining> (visited on 10/25/2026).
- [200] United Nations Treaty Collection. *Status of the Agreement on Marine Biological Diversity of Areas beyond National Jurisdiction*. URL: https://treaties.un.org/Pages/ViewDetails.aspx?chapter=21%5C&clang=_en%5C&tdsg_no=XXI-10%5C&src=TREATY (visited on 10/25/2026).
- [201] University of Washington, Aquatic Organic Geochemistry Laboratory. *Oxygen Minimum Zones*. URL: <https://depts.washington.edu/aog/oxygen-minimum-zones/> (visited on 10/25/2026).
- [202] Stef Van Havermaet et al. "Fish adapt and dynamically avoid an approaching robotic fish across repeated exposures". In: *Scientific Reports* 16, 14248 (2026). DOI: 10.1038/s41598-026-44115-1.
- [203] Ann Vanreusel et al. "Threatened by mining, polymetallic nodules are required to preserve abyssal epifauna". In: *Scientific Reports* 6, 26808 (2016). DOI: 10.1038/srep26808.
- [204] Stein van Veggel et al. "Classification and Evaluation of Octopus-Inspired Suction Cups for Soft Continuum Robots". In: *Advanced Science* 11.30, 2400806 (2024). DOI: 10.1002/advs.202400806.
- [205] T. R. Vonnahme et al. "Effects of a deep-sea mining experiment on seafloor microbial communities and functions after 26 years". In: *Science Advances* 6.18, eaaz5922 (2020). DOI: 10.1126/sciadv.aaz5922.
- [206] Christian R. Voolstra et al. "The coral microbiome in sickness, in health and in a changing world". In: *Nature Reviews Microbiology* 22 (2024), pp. 460–475. DOI: 10.1038/s41579-024-01015-3.
- [207] Chaoqun Wang et al. "Task Allocation for Multi-AUV System: A Review". In: *Ocean Engineering* 266, 112911 (2022). DOI: 10.1016/j.oceaneng.2022.112911.
- [208] Jie Wang et al. "From Shallow Waters to Mariana Trench: A Survey of Bio-Inspired Underwater Soft Robots". In: *Bioinspiration & Biomimetics* 21.2 (2026). DOI: 10.1088/1748-3190/ae3af2.
- [209] Kaihui Wang et al. "Bionic Underwater Vehicle: A Data-Driven Disturbance Rejection Control Framework". In: *IEEE Robotics & Automation Magazine* (2024). URL: <https://ramagazine.ieee.org/2024/04/03/bionic-underwater-vehicle-a-data-driven-disturbance-rejection-control-framework/> (visited on 10/25/2026).
- [210] Yueping Wang et al. "A Biorobotic Adhesive Disc for Underwater Hitchhiking Inspired by the Remora Suckerfish". In: *Science Robotics* 2.10, eaan8072 (2017). DOI: 10.1126/scirobotics.aan8072.
- [211] Stephen C. L. Watson et al. "Assessing, monitoring and mitigating the effects of offshore wind farms on biodiversity". In: *Nature Reviews Biodiversity* 1 (2025), pp. 581–596. DOI: 10.1038/s44358-025-00074-5.
- [212] Brennan P. Watts et al. "Synthetic Chromatophores for Color and Pattern Morphing Skins". In: *Advanced Materials* 37.35, 2505104 (2025). DOI: 10.1002/adma.202505104.
- [213] Haoxuan Wei et al. "Influence of the stratification parameters on the wake field of an underwater vehicle in a two-layer stratified flow". In: *Ocean Engineering* 312 (2024), p. 119000. DOI: 10.1016/j.oceaneng.2024.119000.
- [214] Mark D. Wilkinson, Michel Dumontier, I. Jsbrand Jan Aalbersberg, et al. "The FAIR Guiding Principles for scientific data management and stewardship". In: *Scientific Data* 3, 160018 (2016). DOI: 10.1038/sdata.2016.18.
- [215] Woods Hole Oceanographic Institution. *HROV Nereid Under Ice*. URL: <https://www.whoi.edu/what-we-do/explore/underwater-vehicles/hybrid-vehicles/nereid-under-ice/> (visited on 10/25/2026).
- [216] Woods Hole Oceanographic Institution. *Hybrid Vehicles: Nereid Under Ice and Mesobot*. URL: <https://www.whoi.edu/what-we-do/explore/underwater-vehicles/hybrid-vehicles/> (visited on 10/25/2026).
- [217] Woods Hole Oceanographic Institution. *Nereid Under Ice: Specifications*. URL: <https://www.whoi.edu/what-we-do/explore/underwater-vehicles/hybrid-vehicles/nereid-under-ice/nereid-under-ice-specifications/> (visited on 10/25/2026).
- [218] Woods Hole Oceanographic Institution. *Sentry: Autonomous Underwater Vehicle*. URL: <https://www.whoi.edu/what-we-do/explore/underwater-vehicles/aavs/sentry/> (visited on 10/25/2026).
- [219] Woods Hole Oceanographic Institution, Argo Program. *How Do the Floats Work?* URL: <https://www2.whoi.edu/site/argo/technology/how-do-the-floats-work/> (visited on 10/25/2026).
- [220] World Wide Web Consortium. *PROV-O: The PROV Ontology*. Apr. 30, 2013. URL: <https://www.w3.org/TR/prov-o/> (visited on 10/25/2026).
- [221] Jiaxue Xu et al. "Safety-Critical Parallel Trajectory Tracking Control of Maritime Autonomous Surface Ships Based on Integral Control Barrier Functions". In: *IEEE Transactions on Intelligent Vehicles* 9.5 (2024), pp. 4979–4988. DOI: 10.1109/TIV.2024.3361477.
- [222] Xiaoyi Xu et al. "Exploring underwater image quality: A review of current methodologies and emerging trends". In: *Image and Vision Computing* 154, 105389 (2025). DOI: 10.1016/j.imavis.2024.105389.
- [223] Zefeng Xu et al. "Manta Ray-Inspired Soft Swimmers Actuated by Bistable Actuators". In: *Bioinspiration & Biomimetics* 21.5 (2026). DOI: 10.1088/1748-3190/aea3e6.
- [224] Bo Yan, Hongyuan Liu, and Decai Tang. "From Biological Mechanisms to Task-Oriented Design Principles: A Critical Review of Fish-like Biomimetic Robots". In: *Biomimetics* 11.9, 630 (2026). DOI: 10.3390/biomimetics11090630.
- [225] Namik Kemal Yilmaz et al. "Path Planning of Autonomous Underwater Vehicles for Adaptive Sampling Using Mixed Integer Linear Programming". In: *IEEE Journal of Oceanic Engineering* 33.4 (2008), pp. 522–537. DOI: 10.1109/JOE.2008.2002105.
- [226] Yanwu Zhang et al. "Coordinated and Collaborative Sampling by Two Long-Range Autonomous Underwater Vehicles". In: *IEEE Journal of Oceanic Engineering* 49.4 (2024), pp. 1371–1382. DOI: 10.1109/JOE.2024.3408889.
- [227] Shuo Zhao et al. "Advances of Complex Marine Environmental Influences on Underwater Vehicles". In: *Journal of Marine Science and Engineering* 13.7 (2025), p. 1297. DOI: 10.3390/jmse13071297.
- [228] Yaojian Zhou et al. "High-Precision Surfacing Position Prediction for Underwater Gliders via Coordinate Transformation". In: *Journal of Marine Science and Engineering* 13.4, 760 (2025). DOI: 10.3390/jmse13040760.