Physics-informed representation learning for a digital-ocean twin: The AlphaOcean framework

Fengqin Yan,[1,2] Rong Fan,[1,2] Vincent Lyne,[1,3] and Fenzhen Su[1,2,*]
[1]Key Laboratory of Resources and Environmental Information System, Institute of Geographic Sciences and Natural Resources Research, CAS, Beijing 100101, China
[2]College of Resources and Environment, University of Chinese Academy of Sciences, Beijing 101408, China
[3]IMAS-Hobart, University of Tasmania, Hobart, Tasmania 7004, Australia
*Correspondence: sufz@lreis.ac.cn
Received: January 9, 2026; Accepted: January 27, 2026; Published Online: February 9, 2026; https://doi.org/10.59717/j.xinn-inform.2026.100031

Citation: Yan F., Fan R., Lyne V., et al. (2026). Physics-informed representation learning for a digital-ocean twin: The AlphaOcean framework. *The Innovation Informatics* **2**:100031.

The ocean is a critical yet under-integrated component of the Earth system, regulating climate, sustaining ecosystems, and supporting global food, transportation, and energy systems,[1] while remaining one of the least coherently integrated domains of Earth observation and modeling.[2] Although vast volumes of ocean data are available, they are highly fragmented, heterogeneous in resolution and uncertainty, and unevenly sampled in space, depth, and time. As a result, ocean science continues to rely on task-specific analyses and variable-centric models that generalize poorly across regions and scales, limiting the transformation of growing ocean data archives into transferable knowledge or persistent digital-twin representations of the ocean.

Recent advances in Earth foundation models demonstrate the potential of unified representation learning to integrate complex geospatial data through compact, self-supervised semantic embeddings.[3,4] However, directly extending land-based foundation models to the ocean is fundamentally constrained by the ocean's three-dimensional structure, sparse interior observations, strong multiscale physical–biogeochemical coupling, and scarcity of semantic labels. Existing ocean analysis frameworks further exacerbate this challenge by maintaining a separation between physically based numerical models and data-driven learning approaches, resulting in fragmented knowledge and limited reusability of representations. To address this bottleneck, we introduce AlphaOcean, a foundation framework for digital-ocean intelligence that learns a shared operational multimodal representation of the ocean state from heterogeneous observations.

FOUNDATIONAL CHALLENGES FOR DIGITAL-OCEAN INTELLIGENCE

The primary obstacle to realizing digital-ocean intelligence is not a lack of data, models, or applications, but the absence of an integrating representational layer that can unify these diverse capabilities into a coherent and persistent system. Modern ocean science already encompasses a wide range of observational platforms, numerical models, analytical tools, and application-specific workflows, each optimized for specific questions or domains. However, these components largely operate in isolation, producing outputs that are difficult to compare, reuse, or integrate across scales, variables, and contexts. Building digital-ocean intelligence therefore requires confronting a set of foundational challenges that arise from attempting to integrate heterogeneous capabilities into a shared semantic and dynamical representation of the ocean.

Fragmented observability and data incompatibility. Digital-ocean intelligence must reconcile fundamentally uneven and heterogeneous observations across the ocean system. Ocean data span satellite-derived surface measurements, sparse in situ profiles, autonomous platform trajectories, and model-generated fields, each characterized by distinct spatial coverage, temporal sampling, uncertainty structures, and observational biases. While individual applications can often accommodate these inconsistencies through task-specific preprocessing, a digital-twin ocean requires a representation that can jointly encode surface and subsurface information while explicitly accounting for missing data and uncertainty. The incompatibility of observational modalities thus poses a foundational challenge for learning unified representations that remain meaningful across regions, depths, and time scales, motivating the need to consider how physical processes couple and propagate information beyond directly observed variables.

Process coupling and physical consistency across scales. Being a watery environment, oceans are governed by tightly coupled physical, biogeochemical, and ecological processes that operate across a wide range of spatial and temporal scales. Digital-ocean intelligence cannot be achieved from surface observations alone, as surface signals are often indirect manifestations of subsurface dynamics, lateral transports, and vertical exchanges. Numerical ocean models address this coupling through explicit physical formulations and data assimilation, but their outputs are not readily reusable as semantic representations for diverse applications. Conversely, data-driven approaches that ignore physical constraints in models risk producing representations that are inconsistent, unstable, or misleading when extrapolated. Ensuring that digital-ocean representations respect essential physical relationships and multiscale coupling is therefore a central challenge.

Semantic incoherence and limited reusability of ocean knowledge. A defining feature of digital-ocean intelligence is the ability to compare, retrieve, and reason about ocean states and processes across disparate contexts. At present, ocean knowledge is encoded primarily in task-specific outputs, such as maps, indices, forecasts, or anomaly fields, which lack a shared semantic structure. This fragmentation makes it difficult to identify similarities between ocean regions, track the evolution of comparable phenomena through time, or transfer knowledge from well-studied settings to data-poor environments. The scarcity of explicit semantic labels in the ocean further compounds this issue, limiting supervised learning approaches and hindering interpretability. Overcoming semantic incoherence requires representations that capture meaningful latent structure under weak supervision and that can serve as a reusable substrate for multiple applications, setting the stage for fully spatiotemporal and system-level integration.

Spatiotemporal integration and system-level coherence. Digital-ocean intelligence demands representations that evolve coherently through space and time, rather than static snapshots or surface-only abstractions. Ocean processes are inherently dynamic, shaped by advection, diffusion, mixing, and episodic storm, atmospheric and geophysical events that connect distant regions and sustain effects over extended periods. Representations that fail to encode these spatiotemporal relationships cannot support consistent updating, forecasting, or scenario exploration within a digital twin. Moreover, the ocean's geometry imposes structural constraints that must be reflected in any system-level representation. Addressing spatiotemporal integration is therefore essential for transforming fragmented ocean capabilities into a living digital-ocean intelligence system, and it directly informs the development pathways required to construct viable Ocean Foundation Models.

DEVELOPMENT PATHWAYS TOWARD DIGITAL-OCEAN INTELLIGENCE

Rather than optimizing isolated models for individual tasks, Ocean Foundation Models must provide a persistent semantic and dynamical substrate that can accommodate heterogeneous observations, respect physical constraints, and evolve coherently through space and time. The pathways outlined below describe complementary design directions for constructing such models (Figure 1).

Multimodal representation and fusion across the ocean system. A core pathway toward digital-ocean intelligence is the construction of unified representations that can integrate heterogeneous ocean observations across modalities, depths, and scales. Satellite measurements, in situ profiles, autonomous platforms, and numerical reanalyses each provide partial and biased views of the ocean state, yet digital twins require representations that can jointly encode these disparate sources while preserving their uncertainties and complementarities. Rather than fusing data solely at the input or output level, Ocean Foundation Models emphasize learning shared latent

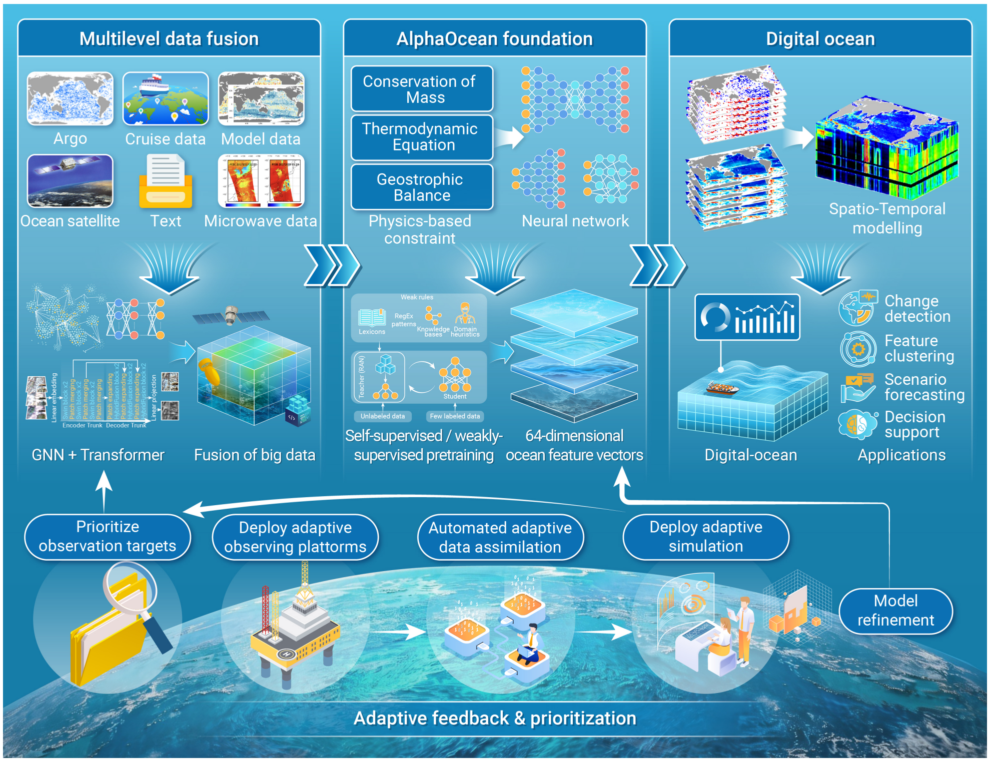

Figure 1. AlphaOcean foundation framework for digital-ocean intelligence.

modelling that encodes advection, mixing, and long-range interactions within learned representations. Rather than treating ocean observations as independent snapshots, AlphaOcean incorporates temporal context and spatial connectivity to maintain continuity and coherence as new data are assimilated. This enables persistent updating of the digital ocean state, supports forecasting and scenario exploration, and allows the digital twin to respond dynamically to perturbations. Crucially, spatiotemporal coherence also facilitates coupling with land and atmospheric digital twins, supporting integrated Earth-system intelligence built upon shared representational foundations.

SUMMARY

This study introduced AlphaOcean as a conceptual framework for Ocean Foundation Models that addresses this gap by placing representation at the center of ocean intelligence. Rather than focusing on task-specific predictions, the framework emphasizes the construction of general-purpose semantic and dynamical embeddings that can integrate heterogeneous observations, respect essential physical constraints, and evolve coherently through space and time. By identifying the foundational challenges that hinder integration and outlining development pathways that respond directly to these challenges, AlphaOcean provides a structured blueprint for advancing ocean-scale representation learning. Unlike AlphaEarth, where the primary bottleneck lies in the scarcity of labeled data rather than observations, the ocean—particularly its interior—is fundamentally constrained by sparse and uneven observations. AlphaOcean does not eliminate this physical limitation. Instead, it reframes ocean observation from isolated samples to constraints on a physically consistent, evolving state representation, thereby amplifying the informational value of each observation and reducing the effective data requirement for digital-ocean intelligence.

representations in which multimodal information is aligned within a common semantic space. Such representations enable surface observations to inform subsurface structure, sparse in situ measurements to propagate context more broadly, and model-derived fields to provide dynamical scaffolding without dominating learned semantics. Establishing this multimodal representational layer is a prerequisite for addressing physical consistency and process coupling within a digital-ocean intelligence framework.

Physics-informed representation learning and constraint integration. A second development pathway focuses on integrating physical knowledge into representation learning, not as a replacement for numerical models, but as a set of guiding constraints that shape learned embeddings. These constraints may include conservation principles, stability conditions, bounds on tracer behaviour, or simplified dynamical relationships that reflect dominant ocean processes. Embedding such information into the learning process helps ensure that representations remain coherent when extrapolated across space, time, or variable combinations, and reduces the risk of producing semantically plausible but physically inconsistent states. Physics-informed representations thus provide a critical bridge between observational integration and semantic reusability, enabling representations that are both interpretable and transferable across applications.

Semantic embedding spaces for reusability and knowledge transfer. A defining capability of digital-ocean intelligence is the ability to reuse learned knowledge across regions, phenomena, and tasks. Ocean Foundation Models pursue this capability by organizing ocean states into semantic embedding spaces that capture meaningful similarities and differences under weak or self-supervision. In such spaces, ocean regions with comparable dynamical or biogeochemical characteristics can be identified, tracked through time, and compared across contexts, even in the absence of explicit labels. This enables similarity-based exploration, clustering of ocean features, and transfer of predictive models from data-rich to data-poor environments. By shifting the focus from task-specific outputs to reusable semantic representations, this pathway directly addresses the fragmentation of ocean knowledge and establishes a foundation for scalable digital-ocean intelligence.

Spatiotemporal modelling for dynamic and coherent digital twins. Digital-twin ocean systems require representations that evolve consistently through space and time, reflecting the connectivity and dynamics of the ocean system. A final development pathway therefore emphasizes spatiotemporal

REFERENCES

1. Hoegh-Guldberg O., Northrop E. and Lubchenco J. (2019). The ocean is key to achieving climate and societal goals. *Science* 365:1372-1374. DOI:10.1126/science.aaz43
2. Révelard A., Tintoré J., Verron J., et al. (2022). Ocean integration: The needs and challenges of effective coordination within the ocean observing system. *Front. Mar. Sci.* 8:737671. DOI:10.3389/fmars.2021.737671
3. Brown C.F., Kazmierski M.R., Pasquarella V.J., et al. (2025). Alphaearth foundations: An embedding field model for accurate and efficient global mapping from sparse label data. *arXiv* preprint. DOI:10.48550/arXiv.2507.22291
4. Tollefson J. (2025). Google AI model creates maps of Earth 'at any place and time'. *Nature* 644:313. DOI: 10.1038/d41586-025-02412-1

FUNDING AND ACKNOWLEDGMENTS

This research is financially supported by the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDB0740300); the National Natural Science Foundation of China (Grant No. 42476195). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare no competing interests.