
GENERALIZABLE IMPLICIT NEURAL REPRESENTATIONS VIA PARAMETERIZED LATENT DYNAMICS FOR BAROCLINIC OCEAN FORECASTING

Guang Zhao, Xihaier Luo, Seungjun Lee, Yihui Ren, Shinjae Yoo

Brookhaven National Laboratory, USA

{gzhao, xluo, slee19, yren, sjyoo}@bnl.gov

Luke Van Roekel, Balu Nadiga

Los Alamos National Laboratory, USA

{lvanroekel, balu}@lanl.gov

Sri Hari Krishna Narayanan, Yixuan Sun

Argonne National Laboratory, USA

{snarayan, yixuan.sun}@anl.gov

Wei Xu

Brookhaven National Laboratory, USA

xuw@bnl.gov

ABSTRACT

Mesoscale ocean dynamics play a critical role in climate systems, governing heat transport, hurricane genesis, and drought patterns. However, simulating these processes at high resolution remains computationally prohibitive due to their nonlinear, multiscale nature and vast spatiotemporal domains. Implicit neural representations (INRs) reduce the computational costs as resolution-independent surrogates but fail in many-query scenarios (inverse modeling) requiring rapid evaluations across diverse parameters. We present **PINROD**, a novel framework combining dynamics-aware implicit neural representations with parametrized neural ordinary differential equations to address these limitations. By integrating parametric dependencies into latent dynamics, our method efficiently captures nonlinear oceanic behavior across varying boundary conditions and physical parameters. Experiments on ocean mesoscale activity data show superior accuracy over existing baselines and improved computational efficiency compared to standard numerical simulations.

1 INTRODUCTION

Mesoscale ocean dynamics is important for climate modeling, as they directly influence global phenomena such as heat transport, hurricane formation, and drought prediction. These dynamics are governed by nonlinear physical processes across vast spatiotemporal scales, necessitating high-fidelity simulations for accurate understanding and prediction Deser et al. (2010); Vallis (2017). However, the computational cost of high-resolution simulations of oceanic processes is prohibitive, particularly in many-query scenarios (e.g., parameter estimation) essential for handling the inherent uncertainties in initial conditions, boundary conditions, and model parameters Edwards & Marsh (2005); Almeida et al. (2010).

Surrogate models, which approximate the input-output mapping of complex models, offer a computationally efficient alternative. While traditional reduced-order models like POD Berkooz et al. (1993) and DMD Schmid (2010) are common, they often fail to capture nonlinear interactions in turbulent fluids. Recently, machine learning-based surrogates, particularly deep learning models, have shown promise in modeling spatiotemporal oceanic processes Berlinghieri et al. (2023); Johnson et al. (2023); Lu et al. (2024). Implicit neural representations (INRs) are particularly appealing due to their ability to learn continuous representations of spatial and temporal fields Xie et al. (2022); Luo et al. (2024), crucial for representing the continuous nature of ocean dynamics.

Unlike grid-based models, INRs use coordinate-based neural networks, enabling extrapolation and continuous solutions. However, existing INR methods struggle with the demands of many-query scenarios in ocean modeling. They typically require training separate networks for each set of conditions (initial, boundary, or model parameters), creating a significant computational bottleneck Lee et al. (2021); Kim et al. (2023). This hinders their practical application where rapid predictions across diverse scenarios are needed.

We introduce **Parameterized Implicit Neural Representations for Ocean Dynamics (PINROD)** to overcome this limitation. PINROD extends dynamics-aware implicit neural representations Yin et al. (2022) with parameterized neural ordinary differential equations, introducing parametric dependencies directly into the latent dynamics. This enables efficient and accurate modeling of nonlinear, high-dimensional oceanic data across diverse scenarios from a single trained model. Our experiments demonstrate that PINROD surpasses state-of-the-art methods in accuracy and computational efficiency, particularly in many-query settings.

2 METHODOLOGY

2.1 PROBLEM SETUP

Building on the need for efficient, generalizable ocean surrogate models, we aim to address the challenge of forecasting mesoscale dynamics under varying physical parameterizations. This is crucial for understanding ocean circulation sensitivity to parameter uncertainties, vital for robust climate projections. We focus on predicting the evolution of an idealized midlatitude ocean basin using the Simulating Ocean Mesoscale Activity (SOMA) test case within the Model for Prediction Across Scales-Ocean (MPAS-O) framework DeYoung et al. (2004); Ringler et al. (2013); Petersen et al. (2019). SOMA provides a computationally tractable yet realistic environment for evaluating a surrogate model. The SOMA configuration simulates an eddying ocean basin (21.5°N–48.5°N, 16.5°W–16.5°E) with curved coastlines, a continental shelf, and 32 km horizontal resolution, capturing essential mesoscale dynamics. This setup is a more realistic evolution of typical double-gyre test cases, allowing for a better assessment of a forecasting surrogate’s performance under conditions resembling real-world ocean modeling. We provide detailed data description in Appendix A.1.

Challenges. Since the dataset contains prognostic variables, including temperature, salinity, and velocity fields, each defined over 8521 hexagonal cells with 60 vertical layers, it yields more than 15 million spatial and temporal data points per variable Wolfram et al. (2015); Sun et al. (2023). Consequently, modeling this data presents challenges due to:

1. **High Dimensionality:** The fine spatial resolution and long temporal extent result in a high-dimensional dataset, posing significant computational challenges.
2. **Unstructured Grid:** Compared to Cartesian grids, the hexagonal grid complicates interpolation and generalization.
3. **Nonlinear and Parameter-Dependent Dynamics:** Ocean dynamics involve complex nonlinear interactions influenced by physical parameters like the Bottom Drag Coefficient C_d , which controls near-bottom friction, affecting deep ocean currents and energy dissipation. These parameterizations introduce further variability and complexity and is important for parameter sensitivity analysis and uncertainty quantification.

Goal. Our objective is to solve the initial value problem for parameterized prediction. Specifically, given the initial state of the ocean and a set of physical parameters, the goal is to predict the *continuous spatiotemporal evolution* of these prognostic variables at arbitrary future times. This problem setting mimics real-world scenarios where ocean models are initialized with observed data and integrated forward in time to produce forecasts.

Problem Formulation. We formulate the SOMA test case prediction task as a parameterized system of PDEs with an initial condition:

$$\partial_t u = F(t, x, u; \mu) \quad \text{with} \quad (t, x) \in [0, T] \times \mathbb{X},$$

where $u = u(t, x)$ represents the ocean state, $\mu \in \mathbb{R}^d$ denotes physical parameters, F governs the nonlinear dynamics, and $\mathbb{X} \in \mathbb{R}^3$ is the spatial domain defined by SOMA’s unstructured mesh. The

initial condition is:

$$u(0, x) = u^0(x) \quad \text{with} \quad x \in \mathbb{X},$$

where $u^0(x)$ specifies the initial ocean state. Given the initial state and physical parameter, the task involves forecasting $u(t, x)$ at arbitrary continuous spatial coordinates and

2.2 MODEL

To address the challenge of the SOMA prediction task, we introduce **Parameterized Implicit Neural Representations for Ocean Dynamics (PINROD)**, a reduced-order modeling framework integrating INRs and parametrized neural ODEs. The **PINROD** reduces the high-dimensionality of the data by encoding the spatial states into a low-dimensional latent vector, while the latent dynamics, influenced by physical parameters, are modeled using neural ODEs. This method learns the dynamic system’s flow in a continuous manner, allowing it to be trained on the unstructured mesh of the SOMA data, while enabling prediction at continuous time and spacial positions. Overall, the framework comprises the three major components:

Encoder. The encoder E_ϕ maps the high-dimensional states $u_t \in \mathbb{R}^{N_t}$ to low-dimensional latent vectors $\alpha_t \in \mathbb{R}^k$ with $k \ll N_t$ as $\alpha_t = E_\phi(u_t)$. We employ an auto-decoder Park et al. (2019) to obtain the latent vectors α_t by minimizing the decoding error. Given the decoder D_ϕ , α_t is obtained by minimize the decoding loss $\|D_\phi(\alpha_t) - u_t\|_2^2$.

PNODE. To capture the continuous temporal evolution of these latent states, we extend Neural ODEs to include explicit parametric dependencies. Specifically, we define: $\frac{d\alpha_t}{dt} = f_\theta(\alpha_t, \mu)$, where μ represents physical parameters. We embed μ via an MLP and concatenate it with α_t as the input to another MLP which outputs the predicted derivatives. This allows the learned dynamics to adapt to varying physical conditions.

Decoder. Once the latent states are evolved in time, we reconstruct the original high-dimensional fields using a coordinate-based INR decoder. In particular, we adopt a SIREN architecture, where the parameters of the SIREN network are modulated by the latent vector α_t as $u_t(x) = D_{\theta_d}(\alpha_t)(x)$, where x indexes the spatial coordinates. This allows us to produce predictions at any queried spatial location, including those not present in the training set, making the model highly flexible for unstructured meshes.

Training Procedure. We employ a two-step training strategy, in Encoder–Decoder Pretraining: We jointly optimize the encoder E_ϕ , the decoder D_ϕ , and the latent vectors $\{\alpha_t\}$ to minimize the reconstruction loss $\|D_\phi(\alpha_t) - u_t\|_2^2$. During this stage, the goal is to learn a low-dimensional manifold that faithfully represents the high-dimensional climate fields. **PNODE Training:** We fix the encoder–decoder parameters (and hence the latent vectors) and train the PNODE to capture the continuous-time dynamics. Concretely, we optimize f_θ by minimizing the temporal prediction error on the latent trajectories, with physical parameters μ included as additional inputs.

This approach allows **PINROD** to handle unstructured sampling mesh, while incorporating relevant physical knowledge into the latent-state evolution. The result is a flexible framework that can make continuous forecasts on the SOMA dataset, surpassing grid-based baselines that struggle with unstructured data.

3 EXPERIMENTS

We demonstrate the performance of **PINROD** on the SOMA dataset, comparing with two baselines FNO and VCNeF. We utilize SOMA simulations with varied bottom drag coefficient values generated from forward runs. The dataset contained 92 forward runs, each spanning 30 days of ocean states. Our prediction task involves that based on




Figure 1: The mean squared error (MSE) of temperature predictions over 29 days. FNO exhibits accumulating error due to its autoregressive strategy, while VCNeF and PINROD maintain more stable performance.

Table 1: **Quantitative results on the SOMA test set.** Metrics include mean squared error (MSE), relative L_2 error (Rel. Err.), and *average inference time* (seconds) to predict all 29 days from a single initial condition. Bolded entries indicate the best results in each category.

Method	Temperature		Salinity		LayerThickness		Infer. Time
	MSE	Rel. Err.	MSE	Rel. Err.	MSE	Rel. Err.	Time (s)
FNO	3.80e-01	5.67e-02	9.11e-02	1.09e-02	3.23e-01	1.29e-02	1.08
VCNeF	3.35e-02	1.55e-02	2.64e-03	1.57e-03	8.8e-03	3.02e-03	1.13
PINROD	2.74e-02	1.53e-02	2.38e-03	1.49e-03	7.98e-03	2.87e-03	3.43

the initial state (day 0) and the bottom drag coefficient value, forecast the ocean states for the next 29 days (days 1-29). We split 92 runs into 72 training set, 10 validation set and 10 testing set. The model performance is reported below.

Baselines. (1) Fourier Neural Operator (FNO) Li et al. (2020), FNO is a neural operator method that learns function mappings in the Fourier frequency domain. In our setup, FNO is trained to make one-day-ahead predictions and is tested to perform multi-days forecasting autoregressively. (2) Vectorized Conditional Neural Field (VCNeF) Hagnberger et al. (2024), VCNeF is a Vision Transformer (ViT) based conditional neural field. It takes physical parameters as condition and learns to solve initial value problem in continuous time. However, it requires training data on a regular grid. (3) PINROD. Our proposed method, which learns spatio-temporal dynamics on unstructured meshes.

Metrics. We evaluate performance using Mean Square Error and Relative L_2 error: $\|u_{\text{pred}} - u_{\text{true}}\|_2 / \|u_{\text{true}}\|_2$. Additionally, we report *inference time* to compare computational efficiency.

Experimental Setup The SOMA data is sampled on a hexagonal mesh within a basin domain. FNO and VCNeF require regular grid data, so we interpolate values on hexagonal mesh onto 3D grid and pad 0's for out of bound coordinates. In contrast, PINROD directly uses coordinate and value pairs for training, avoiding any interpolation error.

Results and Discussion. Table 1 provides quantitative results for three representative fields: Temperature, Layer Thickness and Salinity. From the table, we observe that PINROD consistently yields both lower MSE and lower relative L_2 error, reflecting PINROD's inherent advantage in handling unstructured dataset without interpolation. While PINROD's inference time is not the fastest among comparing methods, it remains significantly faster than numerical simulations.

Figure 1 illustrates how the temperature MSE over the 29-day forecast horizon. While FNO shows reasonable short term accuracy, its error accumulates rapidly due to autoregressive predictions. By contrast, the other two methods, VCNeF and PINROD maintain more stable errors over longer time spans. Figure 2 presents qualitative results for temperature prediction at $t = 29$ across four different depth ($z = 0, 7, 13, 19$). The figures show that PINROD preserves fine-scale structures more accurately, which is crucial in climate modeling. We show more visualizations at $t = 1, 9, 19$ in the Appendix A.2.




Figure 2: **Qualitative Comparison of Temperature Fields.** Snapshots at $t = 29$ for different depths ($z = 0, 7, 13, 19$). PINROD more accurately preserves fine-scale structures.

4 CONCLUSION

In this paper, we introduced PINROD, a novel method that combines implicit neural representations with parameterized neural ODE for baroclinic ocean forecasting, we applied it to the challenging SOMA dataset. The method addresses key challenge of SOMA dataset, including unstructured sampling, high dimensionality and parametric dependent non-linear dynamics. Our experiments demonstrate that PINROD addresses these difficulties and outperforms competing baselines. Our method advances the machine learning application in climate science, facilitate data-driven modeling of complex spatio-temporal processes in climate science problems. Future work can focus on leveraging adaptive sampling strategies and efficient training techniques to improve the model efficiency, enabling its application to more complex and large-scale oceanic datasets.

5 ACKNOWLEDGEMENT

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, and Office of Biological and Environmental Research, through the Scientific Discovery through Advanced Computing (SciDAC) program, under Award Number 9233218CNA000001 and Contract No. DE-AC02-06CH11357. Brookhaven National Laboratory is supported by the DOE’s Office of Science under Contract No. DE-SC0012704 and FWP No. CC122. This research also used resources from the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility located at LBNL.

REFERENCES

- Joao Almeida, Carlos Silvestre, and Antonio Pascoal. Cooperative control of multiple surface vessels in the presence of ocean currents and parametric model uncertainty. *International Journal of Robust and Nonlinear Control*, 20(14):1549–1565, 2010.
- Gal Berkooz, Philip Holmes, and John L Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics*, 25(1):539–575, 1993.
- Renato Berlinghieri, Brian L Trippe, David R Burt, Ryan James Giordano, Kaushik Srinivasan, Tamay Özgökmen, Junfei Xia, and Tamara Broderick. Gaussian processes at the helm (holtz): A more fluid model for ocean currents. In *International Conference on Machine Learning*, pp. 2113–2163. PMLR, 2023.
- Clara Deser, Michael A Alexander, Shang-Ping Xie, and Adam S Phillips. Sea surface temperature variability: Patterns and mechanisms. *Annual review of marine science*, 2(1):115–143, 2010.
- Brad DeYoung, Mike Heath, Francisco Werner, Fei Chai, Bernard Megrey, and Patrick Monfray. Challenges of modeling ocean basin ecosystems. *Science*, 304(5676):1463–1466, 2004.
- Neil R Edwards and Robert Marsh. Uncertainties due to transport-parameter sensitivity in an efficient 3-d ocean-climate model. *Climate dynamics*, 24:415–433, 2005.
- Jan Hagnberger, Marimuthu Kalimuthu, Daniel Musekamp, and Mathias Niepert. Vectorized conditional neural fields: A framework for solving time-dependent parametric partial differential equations. *arXiv preprint arXiv:2406.03919*, 2024.
- J Emmanuel Johnson, Quentin Febvre, Anastasiia Gorbunova, Sam Metref, Maxime Ballarotta, Julien Le Sommer, et al. Oceanbench: The sea surface height edition. *Advances in Neural Information Processing Systems*, 36:78275–78295, 2023.
- Chiheon Kim, Doyup Lee, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Generalizable implicit neural representations via instance pattern composers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11808–11817, 2023.
- Jaeho Lee, Jihoon Tack, Namhoon Lee, and Jinwoo Shin. Meta-learning sparse implicit neural representations. *Advances in Neural Information Processing Systems*, 34:11769–11780, 2021.

-
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Bin Lu, Ze Zhao, Luyu Han, Xiaoying Gan, Yuntao Zhou, Lei Zhou, Luoyi Fu, Xinbing Wang, Chenghu Zhou, and Jing Zhang. Oxygenerator: Reconstructing global ocean deoxygenation over a century with deep learning. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=0HUIInAsdoo>.
- Xihaier Luo, Wei Xu, Balu Nadiga, Yihui Ren, and Shinjae Yoo. Continuous field reconstruction from sparse observations with implicit neural networks. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=kuTZMZdCPZ>.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 165–174, 2019.
- Mark R Petersen, Xylar S Asay-Davis, Anne S Berres, Qingshan Chen, Nils Feige, Matthew J Hoffman, Douglas W Jacobsen, Philip W Jones, Mathew E Maltrud, Stephen F Price, et al. An evaluation of the ocean and sea ice climate of e3sm using mpas and interannual core-ii forcing. *Journal of Advances in Modeling Earth Systems*, 11(5):1438–1458, 2019.
- Todd Ringler, Mark Petersen, Robert L Higdon, Doug Jacobsen, Philip W Jones, and Mathew Maltrud. A multi-resolution approach to global ocean modeling. *Ocean Modelling*, 69:211–232, 2013.
- Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.
- Yixuan Sun, Elizabeth Cucuzzella, Steven Brus, Sri Hari Krishna Narayanan, Balu Nadiga, Luke Van Roekel, Jan Hüchelheim, and Sandeep Madireddy. Surrogate neural networks to estimate parametric sensitivity of ocean models. *arXiv preprint arXiv:2311.08421*, 2023.
- Geoffrey K Vallis. *Atmospheric and oceanic fluid dynamics*. Cambridge University Press, 2017.
- Phillip J Wolfram, Todd D Ringler, Mathew E Maltrud, Douglas W Jacobsen, and Mark R Petersen. Diagnosing isopycnal diffusivity in an eddying, idealized midlatitude ocean basin via lagrangian, in situ, global, high-performance particle tracking (light). *Journal of Physical Oceanography*, 45(8):2114–2133, 2015.
- Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. In *Computer Graphics Forum*, volume 41, pp. 641–676. Wiley Online Library, 2022.
- Yuan Yin, Matthieu Kirchmeyer, Jean-Yves Franceschi, Alain Rakotomamonjy, and Patrick Gallinari. Continuous pde dynamics forecasting with implicit neural representations. *arXiv preprint arXiv:2209.14855*, 2022.

A APPENDIX

In this appendix, we provide additional details about the SOMA dataset and visual examples of the predicted fields at different timesteps.

A.1 SOMA DATASET

The Simulating Ocean Mesoscale Activity (SOMA) dataset, as part of the MPAS-Ocean model, offers a sophisticated simulation of an eddying mid-latitude ocean basin designed to study mesoscale activity and its sensitivity to model parameters. The model domain spans latitudes from 21.58°N

to 48.58°N and longitudes from 16.58°W to 16.58°E, featuring a circular basin with curved coastlines, a 150km wide 100m deep continental shelf, and a longitudinally constant wind stress forcing. The SOMA dataset is implemented in MPAS-Ocean, which uses an unstructured hexagonal mesh with the 32km resolution. The dataset consists of 8521 hexagonal cells, each extending through 60 vertical layers, leading to over 15 million spatical and temporal data points per simulation.

SOMA tracks five key prognostic variables that describe ocean state evolution: layer thickness, salinity, temperature, zonal velocity, and meridional velocity. These variables are influenced by external physical parameters that introduce variability into the simulation. The four primary perturbed parameters include the bottom drag coefficient, which controls near-bottom friction; the Gent–McWilliams (GM) constant diffusivity, which governs eddy-induced advection; the Redi constant diffusivity, responsible for isopycnal mixing; and the background vertical mixing diffusivity, which regulates vertical turbulence. These parameters play a crucial role in mesoscale dynamics, making SOMA an essential dataset for parameter sensitivity analysis.

Each SOMA simulation runs for three years, with the first two years serving as a spin-up period, followed by daily data recording in the third year. The dataset includes 100 ensemble runs per perturbed parameter, allowing for extensive evaluation of ocean model responses. By enabling controlled perturbations of key oceanic parameters, SOMA provides an invaluable resource for developing machine learning surrogates and studying the impact of parameter variations on ocean circulation. It has been used to train neural network models to predict ocean state evolution and compute adjoint sensitivities for parameter optimization. The dataset’s combination of high spatial resolution, long simulation duration, and multiple perturbation parameters makes it a powerful tool for advancing climate modeling, uncertainty quantification, and data-driven ocean forecasting.

A.2 MORE QUALITATIVE COMPARISON OF TEMPERATURE FIELD

To complement the main text, Figures 3, 4, and 5 show temperature snapshots at $t = 1, 9,$ and 19 for different depths ($z = 0, 7, 13, 19$). These examples illustrate how the predictions evolve over time, and how each method captures or misses fine-scale ocean state structures.




Figure 3: Temperature field at $t = 0$.




Figure 4: Temperature field at $t = 9$.




Figure 5: Temperature field at $t = 19$.