

SSTODE: Ocean-Atmosphere Physics-Informed Neural ODEs for Sea Surface Temperature Prediction

Zheng Jiang¹, Wei Wang^{1*}, Gaowei Zhang¹, Yi Wang²

¹School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing, China

²School of Computer Science, Beijing University of Posts and Telecommunications, Beijing, China
{jiangzheng, weiwang, zhanggaowei}@bupt.edu.cn, wang@cocolabs.org

Abstract

Sea Surface Temperature (SST) is crucial for understanding upper-ocean thermal dynamics and ocean-atmosphere interactions, which have profound economic and social impacts. While data-driven models show promise in SST prediction, their black-box nature often limits interpretability and overlooks key physical processes. Recently, physics-informed neural networks have been gaining momentum but struggle with complex ocean-atmosphere dynamics due to 1) inadequate characterization of seawater movement (e.g., coastal upwelling) and 2) insufficient integration of external SST drivers (e.g., turbulent heat fluxes). To address these challenges, we propose SSTODE, a physics-informed Neural Ordinary Differential Equations (Neural ODEs) framework for SST prediction. First, we derive ODEs from fluid transport principles, incorporating both advection and diffusion to model ocean spatiotemporal dynamics. Through variational optimization, we recover a latent velocity field that explicitly governs the temporal dynamics of SST. Building upon ODE, we introduce an Energy Exchanges Integrator (EEI)—inspired by ocean heat budget equations—to account for external forcing factors. Thus, the variations in the components of these factors provide deeper insights into SST dynamics. Extensive experiments demonstrate that SSTODE achieves state-of-the-art performances in global and regional SST forecasting benchmarks. Furthermore, SSTODE visually reveals the impact of advection dynamics, thermal diffusion patterns, and diurnal heating-cooling cycles on SST evolution. These findings demonstrate the model’s interpretability and physical consistency.

Code — <https://github.com/nicezheng/SSTODE-code>

Introduction

Covering 75% of Earth’s surface, the ocean regulates climate through ocean-atmosphere interactions governing energy exchange and hydrological cycles, largely governed by Sea Surface Temperature (SST) variability (Huang and Shukla 2005). SST field constrains upper-ocean circulation and thermal structure across daily to decadal scales, modulating air-sea energy exchange (Garcia-Soto et al. 2021). As the dominant climate fluctuation of this coupling, ENSO

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

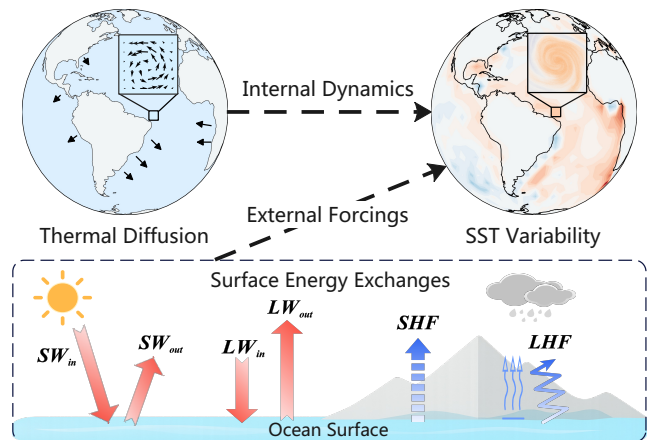


Figure 1: Two primary mechanisms govern SST variability (red: warming, blue: cooling): internal ocean dynamics (coupled advection-diffusion dynamics) and external energy exchanges. Among them, thermal diffusion is particularly crucial since it models heat spread driven by seawater movements like subgrid-scale eddies (see zoomed-in box), and coastal upwelling due to boundary effects (see black arrows). External energy exchanges mainly include shortwave radiation (SW), longwave radiation (LW), sensible heat flux (SHF), and latent heat flux (LHF).

affects the global climate and disrupts normal weather patterns (Ham, Kim, and Luo 2019). Accurate SST predictions at various time-scale are thus crucial for scientific and socioeconomic applications.

SST prediction has traditionally relied on coupled ocean-atmosphere models that numerically solve fluid dynamics equations, e.g., Navier-Stokes (Singh and Roxy 2022; Ma et al. 2021). While offering physical fidelity, they suffer from high computational cost. Deep learning models—from CNNs to Transformers—have recently advanced spatiotemporal SST forecasting by modeling complex ocean systems (Tan et al. 2023b; Shi, Zheng, and Dong 2024). For example, Wenhai (Cui et al. 2025) is a ViT-based model pretrained on 25 years of reanalysis data, enabling skillful multi-day SST forecasting (1-10 day lead times). Though predictive accuracy is essential, the omission of explicit fine-

scale thermodynamic processes in these models compromises both physical consistency and interpretability—a fundamental requirement in ocean science. Furthermore, operating at fixed intervals (e.g., daily or weekly) limits their ability to capture multiscale ocean-atmosphere interactions that evolve across varying timescales.

Recently, Physics-Informed Neural Networks (PINNs) have emerged to incorporate physical laws into model training, preventing physically implausible predictions (Li et al. 2024b,a). A typical strategy involves encoding partial differential equations into Neural ODE frameworks (Verma, Heinonen, and Garg 2024; Xiang et al. 2024). However, the existing physical-informed architectures may not be suitable for SST prediction due to the following inherent challenges.

(1) Inadequate Characterization of Seawater Movement. While the advection equations are typically the basis of existing Neural ODE models for atmospheric transport, SST evolution is governed by coupled advection-diffusion dynamics. As depicted in Fig. 1, this diffusion component is essential to model subgrid-scale heat transfer processes associated with seawater movements, such as turbulence mixing, mesoscale eddies, or coastal upwelling. In particular, due to boundary effects, coastal upwelling generates sharp temperature gradients near coastlines and leads to numerical oscillation during model training (Salois et al. 2022). **(2) Insufficient Integration of External SST Drivers.** In the upper ocean layers, due to surface energy exchanges, the external factors (see Fig. 1) further drive changes in SST tendency. Ignoring such forcings leads to biased predictions, particularly for diurnal heating-cooling cycles, regional warming trends, and seasonal anomalies.

To address the above challenges, we propose SSTODE, an ocean-atmosphere physics-informed Neural ODE framework for SST prediction. SSTODE formulates a continuous-time Neural ODE grounded in the advection-diffusion equation governing oceanic heat transport. Specifically, we incorporate a diffusion term into the Neural ODEs by applying the Laplacian of the SST field (O’dea et al. 2012), scaled by a learnable thermal diffusivity. The Laplacian operator explicitly captures the spatial curvature of SST fields induced by subgrid-scale ocean dynamics, with particular effectiveness in resolving sharp thermal gradients around coastal boundary regions. The learned diffusivity dynamically modulates diffusion strength, smoothing abrupt SST variations in regions with strong thermal gradients. The combination effectively models seawater movement and enhances the numerical stability of ODE, ultimately generating physically consistent and accurate SST predictions. Building upon the Neural ODEs, we introduce an Energy Exchanges Integrator (EEI, a time-dependent network) to further account for external SST drivers in the upper open ocean. Derived from the ocean heat budget equation (Vijith et al. 2020), EEI incorporates four physically grounded surface flux factors (i.e., shortwave radiation, longwave radiation, latent heat flux, and sensible heat flux) from ERA5 reanalysis data to quantify their collective impact on SST dynamics over the forecast horizon. Comprehensive experiments on SST forecasting benchmark and visualizations substantiate the effectiveness and interpretability of SSTODE.

Our main contributions are as follows:

- We propose SSTODE, a continuous-time Neural ODE framework that explicitly models coupled advection-diffusion processes for SST spatiotemporal prediction.
- The EEI is developed that integrates surface heat fluxes, improving ODE solutions by accounting for air-sea energy exchanges—enhancing open-ocean predictions.
- Our model achieves state-of-the-art performance on global and regional SST forecasting benchmarks, revealing three physical mechanisms behind SST evolution: 1) coherent advection flows; 2) boundary-aware diffusion patterns; and 3) external forcing of diurnal variations.

Related Work

Ocean Parameters Prediction Models

Numerical ocean models traditionally simulate ocean dynamics using physical equations but suffer from high computational cost and limited adaptability to dynamic external forcings (Veerasha 2021). Deep models for spatiotemporal forecasting, initially developed for computer vision tasks such as video and traffic prediction, have been extended to SST prediction using both recurrent-based (e.g., ConvLSTM (Lin et al. 2020), Transformers (Tang et al. 2024)) and recurrent-free architectures (e.g., SimVPv2 (Tan et al. 2025), COTERE (Shi, Zheng, and Dong 2024)). Large models have recently emerged as a promising paradigm in ocean and climate science, e.g., AI-GOMS (Xiong et al. 2023), XiHe (Wang et al. 2024), and WenHai (Cui et al. 2025), which leverage massive datasets and millions of parameters to improve prediction accuracy. While these DL models achieve impressive performance, their lack of explicit physical constraints limits their interpretability and thermodynamic consistency. Meanwhile, they are typically trained on reanalysis data sampled at fixed, discrete time intervals (e.g., daily), overlooking the fact that the seawater continuously transports heat and redistributes mass while surface fluxes add or remove energy. This discreteness violates mass conservation and introduces approximation errors. In contrast, our work focuses on ensuring physical interpretability by preserving intermediate states and enabling earlier anomaly detection through continuous-time forecasting.

Physic-Informed Neural Networks

Recent PINNs aim to combine data-driven learning with physical priors to improve prediction accuracy and physical plausibility (Xu et al. 2024; Li et al. 2024a). A straightforward way is to embed the governing partial differential equations (PDEs) in the loss function to enforce physical laws (Ghosh et al. 2023). However, these methods often require full observation of key variables to close the equations, which is challenging for advection-diffusion systems because the requisite velocity fields are typically unavailable from satellite observations (Sun and Wang 2020).

Neural ODEs extend PINNs by parameterizing time derivatives with neural networks and integrating dynamics from initial conditions in continuous time (Chen et al. 2018; Chu et al. 2024). This framework has shown success in modeling physical systems such as fluid dynamics (Rojas,

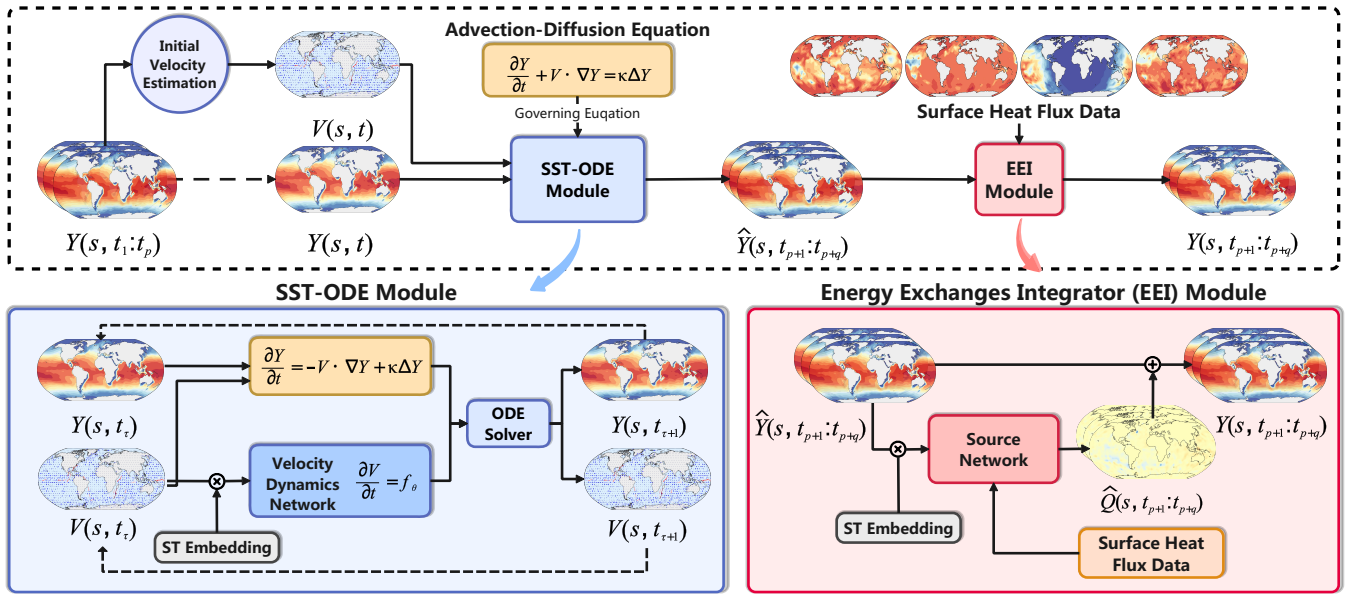


Figure 2: SSTODE framework. It comprises three modules: (1) Initial Velocity Estimation infers latent initial velocity from past SST by solving a PDE-constrained inverse problem; (2) SST-ODE integrates the advection–diffusion equation using Neural ODEs to prediction SST and latent velocity over continuous time; and (3) Energy Exchange Integrator (EEI) refines predictions using surface heat flux data via a learned source network. Spatiotemporal Embeddings (ST Embedding) encode position and time context. \otimes : concatenation. \oplus : element-wise addition.

Dengel, and Ribeiro 2021; Tian et al. 2024) and climate processes (Hwang et al. 2021; Verma, Heinonen, and Garg 2024). For example, ClimODE embeds the advection equation to model atmospheric spatiotemporal patterns, achieving promising results in weather and climate prediction (Verma, Heinonen, and Garg 2024). However, these methods often assume idealized advection in unbounded, compressible atmospheric settings, making them suboptimal for ocean modeling. Specifically, they can not adequately capture key ocean-specific dynamics: (1) intense thermal diffusion near coastlines from turbulence and eddies, (2) incompressible fluid assumptions unique to seawater, and (3) external forcings of surface energy exchange (e.g., radiative and turbulent heat fluxes). Our model addresses this limitation by explicitly integrating coupled advection-diffusion and surface energy exchange patterns into Neural ODE, enabling improved accuracy and realism of SST modeling in the upper open ocean.

Method

Overview

Problem Definition We aim to predict the spatiotemporal evolution of SST distribution over the global ocean, represented by a 2D spatial domain $\Omega = [-90, 90] \times [-180, 180]$, with latitude and longitude defining the horizontal coordinates. Let $Y(s, t) \in \mathbb{R}$ denote the SST at location $s = (x, y) \in \Omega$ and time t . Given a sequence of p past SST fields $Y(s, t_1), \dots, Y(s, t_p)$ sampled at uniform intervals Δt , the goal is to predict q future fields $Y(s, t_{p+1}), \dots, Y(s, t_{p+q})$. We frame this task as a continuous-time neural dynamical

system, initialized with the last observed field $Y(s, t_p)$ and integrated forward to generate future SST fields.

Overall framework As shown in Fig. 2, SSTODE comprises three main components: the Initial Velocity Estimation, the SST-ODE Module, and the Energy Exchange Integrator (EEI) Module. Given past SST observations $Y(s, t_1:t_p)$, the Initial Velocity Estimation module infers the initial velocity field $\mathbf{V}(t_p)$ to initialize the ODE system. The SST-ODE Module then solves a continuous advection-diffusion ODE, integrating from $Y(s, t_p)$ with velocity $\mathbf{V}(t_p)$ to produce predictions $\hat{Y}(s, t_{p+1:p+q})$ iteratively. This module captures both large-scale advection and subgrid-scale thermal diffusion in a unified, physics-driven manner. To account for external forcings, the EEI Module integrates surface heat flux data through a mixed-layer source network and adds their effects to the ODE outputs, yielding the final SST forecast $Y(s, t_{p+1:p+q})$.

Physical Dynamics Modeling

Physical Equations We formulate SST evolution using the advection–diffusion equation, which captures the core physical processes underlying spatiotemporal SST dynamics. This PDE models temporal SST changes driven by horizontal transport (advection) and subgrid-scale dynamics (diffusion), and is defined point-wise over the spatial domain as follows (Eq. 1):

$$\underbrace{\frac{\partial Y(s, t)}{\partial t}}_{\text{Temporal Derivatives}} + \underbrace{\mathbf{V}(s, t) \cdot \nabla Y(s, t)}_{\text{Advection}} = \underbrace{\kappa \Delta Y(s, t)}_{\text{Diffusion}}. \quad (1)$$

- **Advection term.** $\mathbf{V}(s, t) \cdot \nabla Y(s, t)$ represents the horizontal heat transport by ocean surface currents. Here, $\mathbf{V}(s, t) \in \mathbb{R}^2$ denotes the 2D velocity field consisting of zonal (V_x) and meridional (V_y) components, and ∇ is the spatial gradient operator.
- **Diffusion term.** $\kappa \Delta Y(s, t)$ is subgrid-scale thermal diffusion caused by unresolved processes such as turbulence and mesoscale eddies. The scalar coefficient $\kappa \in \mathbb{R}^+$ controls the strength of diffusion, and Δ denotes the spatial Laplacian operator.

ODE Dynamics To enable continuous-time modeling with neural networks, we reformulate the advection–diffusion PDEs as a system of first-order ODEs using the Method of Lines (MOL) (Iakovlev, Heinonen, and Lähdesmäki 2021). For notational simplicity, we omit the spatial coordinate s in $Y(t)$ and $\mathbf{V}(t)$; all spatial operators ∇ and Δ are applied over the discretized grid. The resulting continuous-time ODE integration is given in Eq. 2:

$$\begin{aligned} Y(t) &= Y(t_0) + \int_{t_0}^t [-\mathbf{V}(\tau) \cdot \nabla Y(\tau) + \kappa \Delta Y(\tau)] d\tau, \\ \mathbf{V}(t) &= \mathbf{V}(t_0) + \int_{t_0}^t \dot{\mathbf{V}}(\tau) d\tau, \end{aligned} \quad (2)$$

where $\tau \in \mathbb{R}$ denotes continuous integration time, allowing predictions at arbitrary time resolutions without relying on fixed discrete steps. Such continuous integration is performed across multiple time intervals during the system’s evolution. Integration begins at t_0 , the last observed timestep. The first equation describes SST evolution via the advection-diffusion equation, while the second defines the velocity field’s evolution through temporal derivative $\dot{\mathbf{V}}(\tau)$. These formulations underpin the framework presented in the subsequent sections.

Initial Velocity Estimation

The Neural ODE requires an initial condition to initiate integration, but the initial velocity field $\mathbf{V}(t_0)$ in Eq. 2 is typically unavailable in satellite data (Sun and Wang 2020). We adopt the preprocessing approach from Verma, Heinonen, and Garg (2024) and estimate $\mathbf{V}(t_0)$ by solving a PDE-constrained inverse problem. Specifically, we first approximate the temporal derivative $\frac{\partial Y(t_0)}{\partial t}$ from past p SST observations using cubic spline interpolation, yielding a smooth and differentiable proxy ground truth of local SST variation. We then optimize a learnable latent field $\hat{\mathbf{V}}_\theta(t_0) \in \mathbb{R}^{2 \times H \times W}$ by minimizing the loss of MSE between the interpolated and ODE-computed derivatives under the advection–diffusion equation, with a spatial smoothness regularization term.

$$\hat{\mathbf{V}}_\theta(t_0) = \arg \min_{\mathbf{V}_\theta(t_0)} \{ \|\mathbf{R}(t_0)\|^2 + \alpha \|\mathbf{V}_\theta(t_0)\|^2 \}, \quad (3)$$

$$\mathbf{R}(t) = \frac{\partial Y(t)}{\partial t} + \mathbf{V}_\theta(t) \cdot \nabla Y(t) - \kappa_\theta \Delta Y(t), \quad (4)$$

where κ_θ is a learnable scalar diffusion coefficient constrained via softplus, and α controls the strength of spatial

smoothing. We also place an RBF-kernel Gaussian prior on $\mathbf{V}_\theta(t)$ to encourage coherence. See Appendix A.1 for details and an initial condition robustness test.

SST-ODE Module

The SST-ODE Module jointly models the continuous-time evolution of SST and latent velocity by Neural ODEs grounded in the advection–diffusion equation. As shown in Fig. 2, the module solves the advection–diffusion equation and explicitly captures both large-scale advection and subgrid-scale thermal diffusion. Initialized with estimated velocity field $\mathbf{V}(t_0)$ and SST field $Y(t_0)$, the integration progresses in continuous time. The dynamics are split into two ODEs: one for SST evolution, and one for latent velocity evolution.

Temporal SST Evolution Given the velocity field $\mathbf{V}(t)$, the SST temporal derivative is computed using the advection–diffusion equation. At each continuous integration time τ , this yields:

$$\frac{\partial Y(\tau)}{\partial \tau} = -\mathbf{V}(\tau) \cdot \nabla Y(\tau) + \kappa_\theta \Delta Y(\tau), \quad (5)$$

where κ_θ is parameterized similarly to the initial velocity estimation. The resulting derivative is integrated from the initial SST state using a Neural ODE solver, yielding a continuous and differentiable prediction trajectory. This formulation supports flexible forecasting over arbitrary horizons while preserving a physically interpretable decomposition of SST dynamics into advection and diffusion processes.

Temporal Velocity Evolution Unlike SST, directly modeling velocity evolution is challenging due to the lack of pressure observations and the nonlinearity inherent of the Navier-Stokes equations (Hoang and Sell 2010). To address this, we parameterize the complex velocity evolution $\dot{\mathbf{V}}(\tau)$ as a neural function learned during forward ODE integration. Inspired by the influence of temperature gradients on fluid motion through buoyancy and density-driven forces (Willeit and Ganopolski 2024), we condition the velocity dynamics on both the SST field and its spatial gradients. At each integration time τ , the temporal derivative of velocity is defined as:

$$\frac{\partial \mathbf{V}(\tau)}{\partial \tau} = \dot{\mathbf{V}}(\tau) = f_v(\mathbf{V}(\tau), Y(\tau), \nabla Y(\tau), \phi(s, \tau)), \quad (6)$$

where f_v is a neural velocity dynamics network, and $\phi(s, \tau)$ denotes the spatiotemporal embedding (See Appendix A.2). To capture both local and global dependencies, f_v adopts a hybrid architecture network combining ResNet blocks and attention modules. Similar to SST dynamics, the velocity field $\mathbf{V}(t)$ is continuously updated over time by integrating $\dot{\mathbf{V}}(t)$ using a Neural ODE solver.

Energy Exchange Integrator Module

To incorporate external drivers of SST dynamics, the EEI Module estimate surface energy exchanges (the **Source Term** $Q(s, t)$), and refines the ODE-based outputs accordingly. We begin by outlining the external data sources, followed by our estimation and integration strategy.

Model	OceanVP						ERA5					
	q=5		q=7		q=12		q=5		q=7		q=12	
	MSE ↓	ACC ↑	MSE ↓	ACC ↑	MSE ↓	ACC ↑	MSE ↓	ACC ↑	MSE ↓	ACC ↑	MSE ↓	ACC ↑
ConvLSTM	0.0781	0.9985	0.1091	0.9980	0.1758	0.9967	0.1745	1.0000	0.1879	1.0000	0.2295	1.0000
MAU	0.0627	0.9988	0.0792	0.9985	0.1143	0.9978	0.0242	1.0000	0.0326	1.0000	0.0508	1.0000
PredRNNv2	0.1081	0.9980	0.1251	0.9977	0.1552	0.9972	0.1892	1.0000	0.2081	1.0000	0.2007	1.0000
TAU	0.0576	0.9989	0.0756	0.9985	0.1092	0.9979	0.0242	1.0000	0.0326	1.0000	0.0508	1.0000
SimVPv2	0.0612	0.9988	0.0815	0.9984	0.1221	0.9977	0.0227	1.0000	0.0295	1.0000	0.0450	1.0000
COTERE	0.0624	0.9988	0.0778	0.9984	0.1062	0.9979	0.0284	1.0000	0.0335	1.0000	0.0612	1.0000
ClimODE	0.0622	0.9988	0.0705	0.9986	0.1004	0.9981	0.0196	1.0000	0.0258	1.0000	0.0408	1.0000
SSTODE	0.0527	0.9990	0.0638	0.9987	0.0954	0.9981	0.0180	1.0000	0.0232	1.0000	0.0349	1.0000

Table 1: Global forecasting results on OceanVP and ERA5 at prediction steps $q = 5, 7,$ and 12 . The evaluation metrics are MSE ↓ and ACC ↑, where ↓ indicates lower is better, and ↑ indicates higher is better. Bold indicates the best result.

External Forcing Variables Inspired by the ocean mixed-layer heat budget (Shu et al. 2024), which relates surface energy exchanges to net energy fluxes Q_{net} at the air–sea interface, we adopt the following formulation:

$$Q_{\text{net}} = \frac{Q_{\text{LW}} + Q_{\text{SW}} + Q_{\text{LHF}} + Q_{\text{SHF}}}{\rho c_p h}, \quad (7)$$

where ρ , h , and c_p are the mean density, the mixed layer depth, and the specific heat capacity of sea water, respectively. The numerator includes four dominant fluxes: short-wave (Q_{SW}) and longwave radiation (Q_{LW}), latent (Q_{LHF}), and sensible heat flux (Q_{SHF}). These four components constitute the dominant forms of energy exchange at the ocean surface and are considered essential to SST evolution.

Estimation and Integration We construct the external inputs $H(s, t) = \{Q_{\text{LW}}, Q_{\text{SW}}, Q_{\text{LHF}}, Q_{\text{SHF}}\}$ by concatenating four surface heat flux variables. Since future $H(s, t)$ is unavailable during inference, we extend $H(s, t_0)$ across the prediction horizon and concatenate it with the predicted SST fields $\hat{Y}(s, t_{p+1:p+q})$ and spatiotemporal embeddings ϕ (see Appendix A.2). These inputs are fed into a time-dependent neural network f_s (Source Network in Fig. 2) to estimate the sequence of future source terms:

$$\hat{Q}(s, t_{p+1:p+q}) = f_s \left(H(s, t_0), \hat{Y}(s, t_{p+1:p+q}), \phi \right). \quad (8)$$

This temporally-aware design captures both the impact of external forcing on SST and the feedback effects from SST to surface energy exchange, implicitly modeling short-term feedback effects driven by thermodynamic coupling at the air–sea interface. The estimated source term $\hat{Q}(s, t)$ is then added as a correction to the ODE-based forecast:

$$Y(s, t_{p+1:p+q}) = \hat{Y}(s, t_{p+1:p+q}) + \hat{Q}(s, t_{p+1:p+q}). \quad (9)$$

This formulation explicitly integrates external energy inputs into the forecasting pipeline, allowing the model to adjust its forecast based on real-world environmental variability.

Experiment

Dataset

We evaluate our method on the OceanVP (Shi, Zheng, and Dong 2024) and ERA5 (Hersbach et al. 2020) benchmark.

OceanVP benchmark is built on HYCOM reanalysis (Chassignet et al. 2007), providing 3-hourly global surface temperature, salinity, and velocity fields from 1994 to 2015 at 5.625° resolution. We follow the standard split: 1994–2013 (train), 2014 (val), and 2015 (test). ERA5 offers hourly global reanalysis data from 1940 onward at 0.25° resolution. We extract SST as the target and use four heat flux components—shortwave, longwave, latent, and sensible fluxes—to drive the EEI module. To avoid potential long-term ocean memory effects (Ham, Kim, and Luo 2019), we introduce a 4-year gap by training on 2000 to 2015, validating on 2016, and testing on 2021 to 2022.

Baselines

We benchmark our method against representative spatiotemporal forecasting baselines from three categories: **Recurrent-based:** ConvLSTM (SHI et al. 2015), MAU (Chang et al. 2021), and PredRNNv2 (Wang et al. 2022) utilize recurrent structures to capture temporal dependencies. **Recurrent-free:** TAU (Tan et al. 2023a), SimVPv2 (Tan et al. 2025), and CORETE (Shi, Zheng, and Dong 2024) employ parallel architectures for improving efficiency and scalability. **Physics-Informed:** ClimODE (Verma, Heinonen, and Garg 2024) incorporates the advection equation within a Neural ODE framework without requiring velocity inputs. Neural operator (NO) models learn field-to-field mappings without enforcing explicit physical constraints and thus compared separately in Appendix A.9.

Global Ocean Forecasting Results

We evaluate SSTODE for global SST forecasting on OceanVP and ERA5, comparing against representative spatiotemporal baselines. To balance computational cost and support long-term forecasting, all data are downsampled to 5.625° resolution, 6-hour intervals, and standardized. Models observe 3 past snapshots and predict $q = 5, 7, 12$ future steps (30h, 42h, 72h). We report MSE, and Anomaly Correlation Coefficient (ACC) after de-normalization to assess both short-term variability and long-range dynamics. Table 1 shows that SSTODE consistently outperforms all baselines on both OceanVP and ERA5 across all horizons ($q = 5, 7, 12$). Compared to both recurrent-based and recurrent-free baselines, SSTODE leverages embedded physical priors for

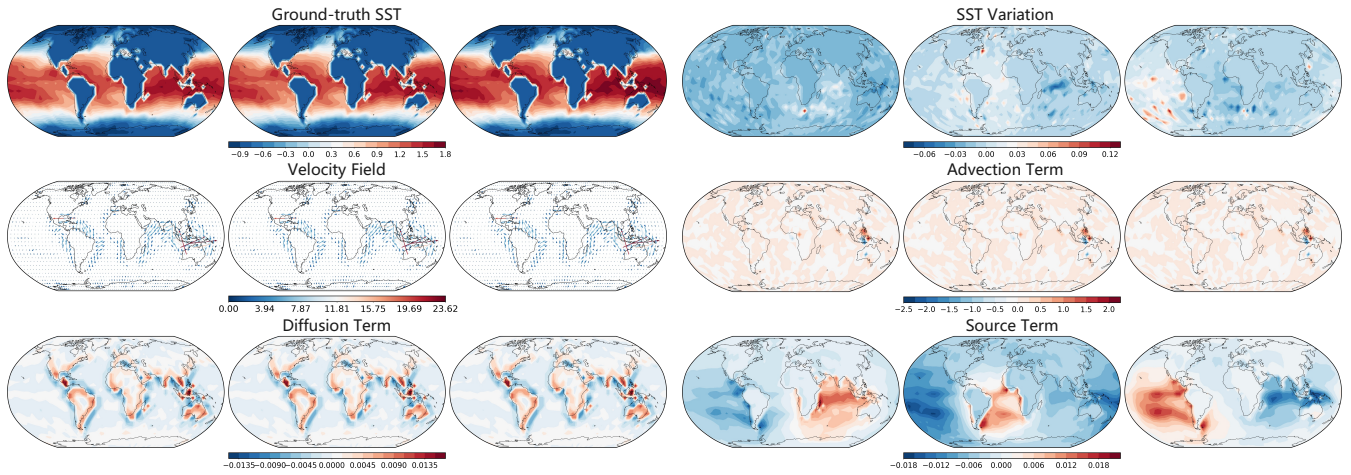


Figure 3: Visualization of intermediate dynamics across three time steps ($t = 1$ to 3). Rows show: Ground-truth SST, SST Variation, Advection, Diffusion, Velocity Field, and Source Term. SSTODE decouples these components, yielding interpretable and physically grounded results across the forecast horizon.

Model	E.P.		N.A.		S.O.	
	q=5	q=12	q=5	q=12	q=5	q=12
ConvLSTM	0.5578	1.1256	1.4052	2.9191	0.4766	1.1772
MAU	0.5590	0.8755	1.1125	1.8058	0.3504	0.7103
PredRNNv2	0.9216	1.2223	1.9601	2.6103	0.6236	0.9231
TAU	0.5736	0.9200	1.0621	1.8038	0.3014	0.6135
SimVPv2	0.5160	0.9612	1.1709	1.9543	0.3051	0.6946
COTERE	0.5736	0.8653	1.1118	1.7515	0.3259	0.6183
ClimODE	0.6064	0.8466	1.1234	1.7046	0.3098	0.5716
SSTODE	0.5157	0.7473	0.9432	1.6445	0.2641	0.5597

Table 2: Regional forecasting results (MSE \downarrow) on OceanVP across Equatorial Pacific (E.P.), Northwest Atlantic (N.A.), and Southern Ocean (S.O.) at forecast steps $q = 5, 12$. Bold indicates the best result per column. MSE is scaled $\times 1000$ for better readability.

better performance under limited observations. It also outperforms the physics-informed ClimODE by jointly modeling thermal diffusion and surface energy exchange.

Regional Ocean Forecasting Results

We evaluate our model on regional forecasting over three oceanographically diverse domains constrained by fixed spatial bounding boxes (Lee et al. 2022; Jiang and Li 2021; Kang et al. 2023): (1) the Equatorial Pacific (5°S – 5°N , 160°E – 100°W), which captures tropical SST variability linked to ENSO dynamics; (2) the Northwest Atlantic (30°N – 50°N , 80°W – 40°W), dominated by the Gulf Stream mesoscale activity; and (3) the Southern Ocean (50°S – 65°S , all longitudes), representing high-latitude regions sensitive to heat exchange and climate variability. We use the same experimental setup as in the global evaluation. Table 2 shows that SSTODE consistently achieves notable gains across all regions and both horizons ($q = 5, 12$). See Appendix A.5 for complete metrics.

Physical Plausibility Analysis

We ensure physical plausibility of SSTODE through both qualitative and quantitative analyses. Qualitatively, Fig. 3 visualizes the intermediate components of the model, which are decomposed into: (1) ground-truth SST $Y(t)$, (2) observed SST variation $\Delta Y = Y(t+1) - Y(t)$, (3) learned velocity field $\mathbf{V}(t)$, (4) advection term $-\mathbf{V}(t) \cdot \nabla Y(t)$, (5) diffusion term $\kappa \Delta Y(t)$, and (6) source term $Q(t)$. These visualizations illustrate how SSTODE decouples SST evolution into physical interpretable processes.

- **Velocity Field & Advection.** These fields reveal coherent large-scale circulations and localized high-magnitude structures. Red arrows highlight key advection zones (e.g., Western Pacific Warm Pool, Indonesian Throughflow, and Caribbean Current between North and South America), which align with known boundary current systems. This indicates that the model captures physically consistent velocity patterns essential for SST transport.
- **Diffusion term.** This component exhibits smooth and coherent spatial patterns across ocean regions, indicating that the learned diffusivity smooths temperature gradients and models gradual heat transport without overfitting to local noise. Notably, it also preserves sharp gradients near coastlines aligning with dynamic coastal processes such as coastal upwelling and boundary currents.
- **Source term.** The learned source term captures diurnal patterns aligned with radiative and turbulent fluxes, presenting strong correspondence with observed SST variations—particularly in regions undergoing day–night heating transitions. This consistency underscores the source term’s contribution to accurate forecasting and emphasizes the importance of modeling external energy exchanges for reliable SST prediction.

For the quantitative check, we assess whether the reanalysis data and model predictions satisfy the advection–diffusion assumption by computing the integrals in Eq. 4 over time and quantities. The results show that both remain constant

Model Variant	MSE ↓	MAE ↓	ACC ↑
SSTODE w/o diffusion	0.0596	0.1185	0.9989
SSTODE Kappa (Fixed)	0.0562	0.1158	0.9989
SSTODE Kappa (2D)	0.0578	0.1166	0.9989
SSTODE	0.0527	0.1107	0.9990

Table 3: Diffusion term ablation. Comparison of diffusivity κ parameterizations and removal of diffusion (w/o) cases at forecast step $q=5$.




Figure 4: Impact of diffusion modeling on prediction bias. Left to right: SST variation, learned diffusion term with strong coastal responses, prediction bias without diffusion, and reduced bias with diffusion, especially in the red box.

over time up to 10^{-6} , demonstrating the physical consistency and numerical stability of the formulation.

Robustness Analysis

Ablation Study on Diffusion Term Design We evaluate the effect of thermal diffusion in SST forecasting by removing the diffusion term and comparing three formulations of the diffusion coefficient κ in Eq. 2, ranging from fixed diffusivity (no parameterization) to spatially varying setups (high parameterization): (i) a fixed κ to a constant ($\kappa = 1$, as used in (Vijith et al. 2020)), (ii) a learnable global scalar, and (iii) a learnable spatially-varying 2D map. Experiments on OceanVP with a 5-step forecast show that removing diffusion markedly degrades performance, while the global scalar yields the best results across all metrics (Table 3). In practice, we observed that the 2D diffusivity map often fails due to over-parameterization. With only SST observed and velocity inferred latently, $\kappa\Delta T$ becomes unstable and partly interchangeable with advection term $-\mathbf{V} \cdot \nabla T$ in explaining SST tendencies. In contrast, the global scalar diffusivity acts as an effective regularizer stabilizing the ODE integration and capturing global diffusion.

Fig. 4 reveals that regions with strong SST gradients, particularly near coastlines and straits in the Northwest Pacific, are highly sensitive to diffusion modeling. Excluding the diffusion term leads to numerical artifacts during and larger errors in these areas. In contrast, the diffusion term acts as a physical regularizer, enabling the model to capture sharp gradients and coastal dynamics (e.g., coastal upwelling) more accurately and stably during training.

Ablation Study on Source Term Design We assess the contribution of each surface heat flux variable used in the EEI Module by incrementally adding SW, LW, LHF, and SHF. As shown in Table 4, adding the full source term leads to a substantial performance gain, underscoring the impor-

Model Variant	MSE ↓	MAE ↓	ACC ↑
SSTODE w/o Source	0.0595	0.1198	0.9988
SSTODE + SW	0.0546	0.1109	0.9990
SSTODE + LW	0.0560	0.1135	0.9990
SSTODE + LHF	0.0559	0.1136	0.9990
SSTODE + SHF	0.0554	0.1173	0.9990
SSTODE	0.0527	0.1107	0.9990

Table 4: Source term ablation. Performance at forecast step $q = 5$ with incremental addition of surface heat flux component (SW, LW, LHF, SHF) in the EEI module.

tance of energy forcing. Among individual terms, SW contributes the most, while all variables offer complementary benefits for accurate SST forecasting.

Longer Horizon and Sparse Sampling We evaluate temporal generalization over a longer 7-day forecast horizon with both standard 6-hourly ($q = 28$) and sparser 12-hourly ($q = 14$) intervals on OceanVP. SSTODE demonstrates robust performance against strong baselines under both settings, as reported in Appendix A.6.

Scalability to Higher Spatial Resolution We assess SSTODE’s scalability to finer spatial grids by increasing the resolution to 2.8125° . Using SST and energy forcing variables extracted from ERA5, experiments in Appendix A.7 report that SSTODE maintains strong performance against representative baselines, demonstrating its potential to scale effectively to even finer resolutions.

Model Parameters Comparison We compare model parameters under identical settings on a 2.8125° -resolution SST forecasting in Appendix A.8. SSTODE demonstrates superior parameter efficiency over both recurrent-based and recurrent-free baselines, benefiting from its autoregressive design. With a marginal increase in parameter count (4.62M vs 3.67M) over ClimODE, SSTODE achieves an $\sim 10\%$ improvement in forecasting accuracy—a favorable trade-off that remains highly acceptable for practical deployment.

Conclusion

SST prediction remains fundamentally important for understanding evolving ocean structure and its feedback with atmospheric circulation. This paper addresses the challenges of adapting Neural ODE framework for SST forecasting by incorporating explicit advection-diffusion constraints to model ocean dynamics and designing an EEI Module to account for surface heat fluxes. Extensive experiments demonstrate that our model achieves competitive performance, effectively capturing key physical processes such as advection, boundary-aware diffusion and diurnal variations. This work highlights the potential of combining physical priors with Neural ODEs for more interpretable and accurate ocean forecasting and also establishes it as a promising component for foundation models that support regional refinement and physically consistent downscaling. Future work will explore the extension to ocean subsurface thermohaline and multi-scale dynamics.

Acknowledgments

This work is partially supported by National Natural Science Foundation of China under grants 62076232 and 62172049.

References

- Chang, Z.; Zhang, X.; Wang, S.; Ma, S.; Ye, Y.; Xinguang, X.; and Gao, W. 2021. Mau: A Motion-aware Unit for Video Prediction and Beyond. *Advances in Neural Information Processing Systems*, 34: 26950–26962.
- Chassignet, E. P.; Hurlburt, H. E.; Smedstad, O. M.; Halliwell, G. R.; Hogan, P. J.; Wallcraft, A. J.; Baraille, R.; and Bleck, R. 2007. The HYCOM (hybrid coordinate ocean model) Data Assimilative System. *Journal of Marine Systems*, 65(1-4): 60–83.
- Chen, R. T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural Ordinary Differential Equations. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Chu, Z.; Ma, W.; Li, M.; and Chen, H. 2024. Adaptive Decision Spatio-temporal Neural ODE for Traffic Flow Forecasting with Multi-Kernel Temporal Dynamic Dilation Convolution. *Neural Networks*, 179: 106549.
- Cui, Y.; Wu, R.; Zhang, X.; Zhu, Z.; Liu, B.; Shi, J.; Chen, J.; Liu, H.; Zhou, S.; Su, L.; et al. 2025. Forecasting the Eddying Ocean with a Deep Neural Network. *Nature Communications*, 16(1): 2268.
- Garcia-Soto, C.; Cheng, L.; Caesar, L.; Schmidtko, S.; Jewett, E. B.; Cheripka, A.; Rigor, I.; Caballero, A.; Chiba, S.; Báez, J. C.; et al. 2021. An Overview of Ocean Climate Change Indicators: Sea Surface Temperature, Ocean Heat Content, Ocean pH, Dissolved Oxygen Concentration, Arctic Sea Ice Extent, Thickness and Volume, Sea Level and Strength of the AMOC (Atlantic Meridional Overturning Circulation). *Frontiers in Marine Science*, 8.
- Ghosh, S.; Chakraborty, A.; Brikis, G. O.; and Dey, B. 2023. RANS-PINN Based Simulation Surrogates for Predicting Turbulent Flows. *arXiv preprint arXiv:2306.06034*.
- Ham, Y.-G.; Kim, J.-H.; and Luo, J.-J. 2019. Deep Learning for Multi-year ENSO Forecasts. *Nature*, 573(7775): 568–572.
- Hersbach, H.; Bell, B.; Berrisford, P.; Hirahara, S.; Horányi, A.; Muñoz-Sabater, J.; Nicolas, J.; Peubey, C.; Radu, R.; Schepers, D.; et al. 2020. The ERA5 Global Reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.
- Hoang, L. T.; and Sell, G. R. 2010. Navier–Stokes Equations with Navier Boundary Conditions for an Oceanic Model. *Journal of Dynamics and Differential Equations*, 22: 563–616.
- Huang, B.; and Shukla, J. 2005. Ocean–Atmosphere Interactions in the Tropical and Subtropical Atlantic Ocean. *Journal of Climate*, 18(11): 1652 – 1672.
- Hwang, J.; Choi, J.; Choi, H.; Lee, K.; Lee, D.; and Park, N. 2021. Climate Modeling with Neural Diffusion Equations. In *2021 IEEE international conference on data mining (ICDM)*, 230–239. IEEE.
- Iakovlev, V.; Heinonen, M.; and Lähdesmäki, H. 2021. Learning Continuous-time PDEs from Sparse Data with Graph Neural Networks. In *International Conference on Learning Representations*. OpenReview. net.
- Jiang, L.; and Li, T. 2021. Impacts of Tropical North Atlantic and Equatorial Atlantic SST anomalies on ENSO. *Journal of Climate*, 34(14): 5635–5655.
- Kang, S. M.; Yu, Y.; Deser, C.; Zhang, X.; Kang, I.-S.; Lee, S.-S.; Rodgers, K. B.; and Ceppi, P. 2023. Global Impacts of Recent Southern Ocean Cooling. *Proceedings of the National Academy of Sciences*, 120(30): e2300881120.
- Lee, S.; L’Heureux, M.; Wittenberg, A. T.; Seager, R.; O’Gorman, P. A.; and Johnson, N. C. 2022. On the Future Zonal Contrasts of Equatorial Pacific climate: Perspectives from Observations, Simulations, and Theories. *Npj Climate and Atmospheric Science*, 5(1): 82.
- Li, W.; Liu, Z.; Chen, K.; Chen, H.; Liang, S.; Zou, Z.; and Shi, Z. 2024a. Deepphysinet: Bridging Deep Learning and Atmospheric Physics for Accurate and Continuous Weather Modeling. *arXiv preprint arXiv:2401.04125*.
- Li, Z.; Zheng, H.; Kovachki, N.; Jin, D.; Chen, H.; Liu, B.; Azizzadenesheli, K.; and Anandkumar, A. 2024b. Physics-Informed Neural Operator for Learning Partial Differential Equations. *ACM / IMS J. Data Sci.*, 1(3).
- Lin, Z.; Li, M.; Zheng, Z.; Cheng, Y.; and Yuan, C. 2020. Self-Attention ConvLSTM for Spatiotemporal Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11531–11538.
- Ma, Y.; Li, J.; Zhang, S.; and Zhao, H. 2021. A Multi-model Study of Atmosphere Predictability in Coupled Ocean–Atmosphere Systems. *Climate Dynamics*, 56: 3489–3509.
- O’dea, E.; Arnold, A.; Edwards, K.; Furner, R.; Hyder, P.; Martin, M.; Siddorn, J.; Storkey, D.; While, J.; Holt, J.; et al. 2012. An Operational Ocean Forecast System Incorporating NEMO and SST Data Assimilation for the Tidally Driven European North-West Shelf. *Journal of Operational Oceanography*, 5(1): 3–17.
- Rojas, C. J.; Dengel, A.; and Ribeiro, M. D. 2021. Reduced-order Model for Fluid Flows via Neural Ordinary Differential Equations. *arXiv preprint arXiv:2102.02248*.
- Salois, S. L.; Gouhier, T. C.; Helmuth, B.; Choi, F.; Seabra, R.; and Lima, F. P. 2022. Coastal Upwelling Generates Cryptic Temperature Refugia. *Scientific Reports*, 12(1): 19313.
- SHI, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-k.; and WOO, W.-c. 2015. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In Cortes, C.; Lawrence, N.; Lee, D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

- Shi, Z.; Zheng, H.; and Dong, J. 2024. OceanVP: A HYCOM Cased Benchmark Dataset and a Relational Spatiotemporal Predictive Network for Oceanic Variable Prediction. *Ocean Engineering*, 304: 117748.
- Shu, R.; Wu, H.; Gao, Y.; Xu, F.; Gou, R.; and Huang, X. 2024. Improved Forecasts of Global Extreme Marine Heatwaves Through a Physics-guided Data-driven Approach. *arXiv preprint arXiv:2412.15532*.
- Singh, V. K.; and Roxy, M. 2022. A Review of Ocean-Atmosphere Interactions During Tropical Cyclones in the North Indian Ocean. *Earth-Science Reviews*, 226: 103967.
- Sun, L.; and Wang, J.-X. 2020. Physics-constrained Bayesian Neural Network for Fluid Flow Reconstruction with Sparse and Noisy Data. *Theoretical and Applied Mechanics Letters*, 10(3): 161–169.
- Tan, C.; Gao, Z.; Li, S.; and Li, S. Z. 2025. SimVPv2: Towards Simple yet Powerful Spatiotemporal Predictive Learning. *IEEE Transactions on Multimedia*, 1–15.
- Tan, C.; Gao, Z.; Wu, L.; Xu, Y.; Xia, J.; Li, S.; and Li, S. Z. 2023a. Temporal Attention unit: Towards Efficient Spatiotemporal Predictive Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18770–18782.
- Tan, C.; Li, S.; Gao, Z.; Guan, W.; Wang, Z.; Liu, Z.; Wu, L.; and Li, S. Z. 2023b. Openstl: A Comprehensive Benchmark of Spatio-temporal Predictive Learning. *Advances in Neural Information Processing Systems*, 36: 69819–69831.
- Tang, Y.; Dong, P.; Tang, Z.; Chu, X.; and Liang, J. 2024. Vmrnn: Integrating Vision Mamba and Lstm for Efficient and Accurate Spatiotemporal Forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5663–5673.
- Tian, J.; Liang, Y.; Xu, R.; Chen, P.; Guo, C.; Zhou, A.; Pan, L.; Rao, Z.; and Yang, B. 2024. Air Quality Prediction with Physics-Guided Dual Neural ODEs in Open Systems. In *The Twelfth International Conference on Learning Representations*.
- Veerasha, P. 2021. A Numerical Approach to the Coupled Atmospheric Ocean Model using a Fractional Operator. *Mathematical Modelling and Numerical Simulation with Applications*, 1(1): 1–10.
- Verma, Y.; Heinonen, M.; and Garg, V. 2024. ClimODE: Climate Forecasting With Physics-informed Neural ODEs. In *The Twelfth International Conference on Learning Representations*.
- Vijith, V.; Vinayachandran, P.; Webber, B. G.; Matthews, A. J.; George, J. V.; Kannaujia, V. K.; Lotliker, A. A.; and Amol, P. 2020. Closing the Sea Surface Mixed Layer Temperature Budget from in Situ Observations Alone: Operation Advection during BoBBLE. *Scientific Reports*, 10(1): 7062.
- Wang, X.; Wang, R.; Hu, N.; Wang, P.; Huo, P.; Wang, G.; Wang, H.; Wang, S.; Zhu, J.; Xu, J.; et al. 2024. Xihe: A Data-driven Model for Global Ocean Eddy-Resolving Forecasting. *arXiv preprint arXiv:2402.02995*.
- Wang, Y.; Wu, H.; Zhang, J.; Gao, Z.; Wang, J.; Yu, P. S.; and Long, M. 2022. Predrnn: A Recurrent Neural Network for Spatiotemporal Predictive Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2208–2225.
- Willeit, M.; and Ganopolski, A. 2024. Generalized Stability Landscape of the Atlantic Meridional Overturning Circulation. *Earth System Dynamics*, 15(6): 1417–1434.
- Xiang, Y.; Li, C.; He, S.; and Chen, J. 2024. AGC-ODE: Adaptive Graph Controlled Neural ODE for Human Mobility Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 25(11): 18449–18460.
- Xiong, W.; Xiang, Y.; Wu, H.; Zhou, S.; Sun, Y.; Ma, M.; and Huang, X. 2023. Ai-goms: Large Ai-driven Global Ocean Modeling System. *arXiv preprint arXiv:2308.03152*.
- Xu, W.; Ling, F.; Han, T.; Chen, H.; Ouyang, W.; and BAI, L. 2024. Generalizing Weather Forecast to Fine-grained Temporal Scales via Physics-ai Hybrid Modeling. *Advances in Neural Information Processing Systems*, 37: 23325–23351.