

Geophysical Research Letters®



RESEARCH LETTER

10.1029/2024GL114318

Key Points:

- We develop a global, 3D, ocean autoregressive machine learning emulator for climate studies
- The emulator, based on a UNet architecture, is stable for centuries, producing accurate climatologies and variability of ocean variables
- The emulator training is robust to changes in seeds and initial conditions in the data

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

S. Dheeshjith,
sd5313@nyu.edu

Citation:

Dheeshjith, S., Subel, A., Adcroft, A., Busecke, J., Fernandez-Granda, C., Gupta, S., & Zanna, L. (2025). Samudra: An AI global ocean emulator for climate. *Geophysical Research Letters*, 52, e2024GL114318. <https://doi.org/10.1029/2024GL114318>

Received 16 DEC 2024

Accepted 15 MAY 2025

Samudra: An AI Global Ocean Emulator for Climate

Surya Dheeshjith¹ , Adam Subel¹, Alistair Adcroft² , Julius Busecke³ , Carlos Fernandez-Granda^{1,4} , Shubham Gupta¹ , and Laure Zanna^{1,4} 

¹Courant Institute of Mathematical Sciences, New York University, New York, NY, USA, ²Program in Atmospheric and Oceanic Sciences, Princeton University, Princeton, NJ, USA, ³Lamont Doherty Earth Observatory, Columbia University, Palisades, NY, USA, ⁴Center for Data Science, New York University, New York, NY, USA

Abstract AI emulators for forecasting have emerged as powerful tools that can outperform conventional numerical predictions. The next frontier is to build emulators for long climate simulations with skill across a range of spatiotemporal scales, a particularly important goal for the ocean. Our work builds a skillful global emulator of the ocean component of a state-of-the-art climate model. We emulate key ocean variables, sea surface height, horizontal velocities, temperature, and salinity, across their full depth. We use a modified ConvNeXt UNet architecture trained on multi-depth levels of ocean data. We show that the ocean emulator—*Samudra*—which exhibits no drift relative to the truth, can reproduce the depth structure of ocean variables and their interannual variability. *Samudra* is stable for centuries and 150 times faster than the original ocean model. *Samudra* struggles to capture the correct magnitude of the forcing trends and simultaneously remain stable, requiring further work.

Plain Language Summary AI tools are extremely effective in making fast and accurate predictions on weather to seasonal timescales. Capturing decadal to centennial changes, which arise from ocean dynamics, remains an outstanding challenge. We built an advanced AI model called “*Samudra*” to simulate global ocean behavior. *Samudra* is trained on simulated data from a state-of-the-art ocean climate model and predicts key ocean features such as sea surface height, currents, temperature, and salinity throughout the ocean's depth. *Samudra* can accurately recreate patterns in ocean variables, including year-to-year changes. It is stable over centuries and is 150 times faster than traditional ocean models. However, *Samudra* still faces challenges in balancing stability with accurately predicting the effects of external factors (like climate trends), and further improvements are needed to address this limitation.

1. Introduction

The recent success of emulators for components of the climate system, primarily the atmosphere, continues to produce remarkable outcomes, achieving state-of-the-art performance for weather prediction tasks (Bi et al., 2023; Kochkov et al., 2024; Price et al., 2023) and promising results reproducing climate models over decadal (Cachay et al., 2024) to multi-decadal timescales (Watt-Meyer et al., 2023).

Existing work on ocean emulation has mainly been limited to the surface and upper ocean, or to steady forcing. Several works focusing on surface ocean variables show results for timescales of years to a decade (Dheeshjith et al., 2024; Gray et al., 2024; Subel & Zanna, 2024). Emulators that include subsurface information have focused on the weekly to decadal timescales and at most the upper 1,000 m (Arcomano et al., 2023; Guo et al., 2024; Holmberg et al., 2024; Patel et al., 2024; Xiong et al., 2023). Bire et al. (2023) explored longer timescales within a simplified ocean model with idealized steady forcing. Finally, a seasonal coupled atmosphere–ocean emulator has shown promising results, considering the upper 300 m of the ocean (Wang et al., 2024). These ocean and atmosphere emulators have been used for seasonal forecasts based on reanalysis data, and to build surrogates of numerical models.

Emulators of traditional numerical climate models leverage the computational efficiency of machine learning approaches to reduce the often prohibitive computational cost of running a large number of simulations on the original (usually CPU-based) climate model. One of the main benefits of emulators is the ability to run large ensembles. Such ensembles can be used to probe the likelihood of extreme events, explore the climate response to a range of forcing scenarios (e.g., greenhouse gases), and facilitate the development of numerical models by reducing the number of perturbed parameter experiments typically used for calibration (Maher et al., 2021; Mahesh et al., 2024). Emulators can also accelerate the spin-up integration of numerical models or replace full

© 2025. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

model components in a coupled setting (Khatiwala, 2024). Finally, emulators can help with data assimilation, replacing an expensive numerical model with a fast surrogate to generate affordable ensembles or an approximate adjoint, maintaining accuracy with reduced cost (Manshausen et al., 2024).

Our goal here is to reproduce the full-depth ocean state for four 3D and one 2D prognostic variables, using a time-dependent realistic atmospheric forcing as input, extending the work of Subel and Zanna (2024); Dheeshjith et al. (2024). At rollout lengths of nearly a decade, our emulator shows considerable skill across several key diagnostics (mean and variance) when compared to the parent numerical model output, which is our ground truth. In particular, both the temperature structure as a function of depth and the El Niño–Southern Oscillation (ENSO) variability are well reproduced by the emulator.

Simultaneously capturing variables with vastly different timescales, such as velocity (which can contain fast fluctuations) and salinity (which typically fluctuates more slowly), is an outstanding issue for long integrations (already encountered by Subel and Zanna (2024)). To alleviate this problem, we introduce an additional emulator by focusing on the thermodynamic variables (i.e., potential temperature and salinity only). This additional emulator captures the slowly varying changes in potential temperature and salinity on timescales of decades to centuries.

We show that our emulator can retain skill and remain stable for centuries for experiments equivalent to both control and climate-change simulations. However, we also note that this stability is accompanied by a weak response to climate-change forcing. This work demonstrates (to our knowledge) the first ocean emulator capable of reproducing the full-depth (from the surface down to the ocean floor) ocean temperature structure and its variability, while running for multiple centuries in a realistic configuration with time-dependent forcing.

The paper is organized as follows. We discuss the data and all emulator details in Section 2. We explore the properties of the trained emulator on a test data set and report several multi-decadal experiments with a range of climate forcing in Section 3. We present our conclusions in Section 4.

2. Methods

We built an autoregressive ocean emulator from data generated by a state-of-the-art numerical ocean simulation. Below, we describe the data, the emulator, the architecture, and the training and evaluation of the emulator.

2.1. Data

The data was generated by OM4 (Adcroft et al., 2019), an ocean general circulation model that is the ocean component of the state-of-the-art coupled climate model CM4 (Held et al., 2019). The circulation model was initialized with hydrography from the World Ocean Atlas (Levitus et al., 2015) and forced with atmospheric reanalysis, following the OMIP-2 protocol, with version 1.4 of the Japanese Atmospheric Reanalysis (Tsujiro et al., 2020). The model was run for 65 years (1958–2022).

The ocean prognostic variables are potential temperature (θ_o), salinity (S), sea surface height (SSH), oceanic zonal (u), and meridional (v) velocity components. The circulation model has 75 degrees of freedom in the vertical for each 3D prognostic variable, which we conservatively remap onto 19 fixed-depth levels of variable thickness—[2.5, 10, 22.5, 40, 65, 105, 165, 250, 375, 550, 775, 1,050, 1,400, 1,850, 2,400, 3,100, 4,000, 5,000, 6,000] m to reduce the data size. We also conservatively coarsen the data in time using a 5-day simple average in geopotential coordinates, averaging over the fastest waves resolved by the circulation model (which originally used a 20-min time-step).

The native horizontal grid for the data has a nominal resolution of $1/4^\circ$, but is curvilinear and has three poles (grid singularities) inland. We further post-process by filtering with an 18×18 cell Gaussian kernel using the *gem-filters* package (Loose et al., 2022), and then conservatively interpolate onto a $1^\circ \times 1^\circ$ Gaussian grid using the *xESMF* package (Zhuang et al., 2023). Values in land are treated as missing, and missing values are imputed with zeros. Before conservative spatial interpolation, we interpolate the velocities to the center of each cell using the *xGCM* package (Abernathy et al., 2022) and rotate the velocity vectors so that u and v indicate purely zonal and meridional flow, respectively.

2.2. Ocean Emulator

The variables in the ocean emulator are:

1. The ocean state $\Phi = (\theta_O, S, SSH, u, v)$, which includes all 19 depth levels. We denote the subset of thermodynamics variables as $\Phi_{\text{thermo}} = (\theta_O, S, SSH)$, as opposed to the dynamic variables $\Phi_{\text{dynamic}} = (u, v)$.
2. Atmosphere boundary conditions $\tau = (\tau_u, \tau_v, Q, Q_{\text{anom}})$, which consist of the zonal, τ_u , and meridional, τ_v , surface ocean stress, and net heat flux downward across the ocean surface Q (below the sea-ice) and its anomalies Q_{anom} . The net heat flux is a sum of the short- and long-wave radiative fluxes, sensible and latent heating, heat content of mass transfer, and heat flux due to frazil formation (see K4 and K5 of Griffies et al. (2016) for a precise definition of the variable “*hfds*”). The heat flux anomalies are calculated by removing the climatological heat flux computed over the 65-year OM4 data set.

Our emulator, $\mathcal{F}$, is built to autoregressively produce multiple future oceanic states given multiple previous oceanic states. Specifically, we use a 2-input–2-output model configuration. Mathematically,

$$\tilde{\Phi}_{t+(n+1)\Delta t}, \tilde{\Phi}_{t+(n+2)\Delta t} = \mathcal{F}(\tilde{\Phi}_{t+(n-1)\Delta t}, \tilde{\Phi}_{t+n\Delta t}, \tau_{t+n\Delta t}) \quad (1)$$

where n is a positive integer and $\tilde{\Phi}$ represents the ocean state predicted by the emulator at time t . A depth-varying land mask is used to set land cells in the model output to zero. We use OM4 ocean states, Φ_t and $\Phi_{t-\Delta t}$, along with the corresponding atmospheric forcing, τ_t , to produce the first predictions. Subsequent ocean states are recursively produced by using previously generated ocean states as input. We illustrate the rollout process of the emulator in (Figure 1a). The use of multiple input states provides additional context to the emulator, similar to the use of model time tendencies in PDE-based numerical integrations. In all of our experiments, $\Delta t = 5$ days.

2.3. Architecture

The emulator is based on the ConvNeXt UNet architecture from (Dheeshjith et al., 2024), where the core blocks of a UNet (Ronneberger et al., 2015) are inspired by ConvNeXt blocks (Liu et al., 2022) adapted from (Karlbauer et al., 2023). The UNet implements downsampling based on average pooling and upsampling based on bilinear interpolation, which enables it to learn features at multiple scales. Each ConvNext block includes GeLU activations, increased dilation rates, and inverted channel bottlenecks. We did not use inverted channel depths and replaced the large 7×7 kernels with 3×3 kernels. We use batch normalization instead of layer normalization, as it yielded better skill. The encoder and decoder consist of four ConvNeXt blocks, each with channel widths [200, 250, 300, 400]. The dilation rates used for both the encoder and decoder are [1, 2, 4, 8]. Additionally, we include a single ConvNext block (with channel width 400 and dilation 8) in the deepest section of the UNet before upsampling. The total number of model parameters is 135 M. We apply periodic (or circular) padding in the longitudinal direction and zero padding at the poles as in (Dheeshjith et al., 2024).

The architecture is modified from Dheeshjith et al. (2024) to process multiple ocean depth levels (as opposed to surface only). In the surface ocean emulator, which contains only a single depth level, each channel is associated with a variable. In the multi-depth ocean emulator, each channel is associated with a variable and a depth level. Our main emulator $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ takes as input four 19-level oceanic variables (θ_O, S, u, v), the surface variable SSH and four atmospheric boundary conditions ($\tau_u, \tau_v, Q, Q_{\text{anom}}$). It produces five output variables (θ_O, S, SSH, u, v). As discussed above, we use a 2-input 2-output model configuration and thus, there are $(4 \times 19 + 1) \times 2 + 4 = 158$ input and $(4 \times 19 + 1) \times 2 = 154$ output channels. In addition, we build another emulator $\mathcal{F}_{\text{thermo}}$ that only uses the thermodynamic variables, $\Phi_{\text{thermo}} = (\theta_O, S, SSH)$.

2.4. Training Details

We illustrate the training of the model in (Figure 1a). We train the emulators using 2,900 data samples corresponding to the range 1975-01-03 to 2014-09-20 with the last 50 samples used for validation. Each sample is a 5-day mean of the full ocean state and atmospheric boundary conditions.

We ignore data over 1958–1975 due to the excessive model cooling, while it adjusts from the warm initial conditions. This cooling does not reflect the forcing but rather an interior ocean model adjustment (see Sane




Figure 1. (a) Schematic of the model training process, illustrating the mapping from input (ocean states and atmospheric forcing) to output (ocean states rolled out over several time steps). Initially, the ground-truth ocean states, Φ_t and $\Phi_{t-\Delta t}$, along with the atmospheric forcing, τ_t , are provided as inputs to predict $\Phi_{t+\Delta t}$ and $\Phi_{t+2\Delta t}$. Predictions, along with ground-truth atmospheric forcing, are then used as inputs for future steps in the unrolling process. (b) Time-averaged potential temperature (θ_o) depth-latitude profiles over the 8-year test set, comparing the ground truth OM4 (left) and predictions from $\mathcal{F}_{\text{thermo}}$ (middle) and $\mathcal{F}_{\text{thermo+dynamic}}$ (right). (c) root mean square error (RMSE) of 8-year test set predictions for different initial conditions of the emulators, $\mathcal{F}_{\text{thermo}}$ and $\mathcal{F}_{\text{thermo+dynamic}}$. Gray dots represent an RMSE instance of a single rollout, including runs from training on 5 unique model seeds per emulator and 2 additional rollouts initialized at states 6 months apart. Horizontal lines indicate the respective mean RMSE. RMSE is calculated over the common periods of each rollout.

et al. (2023) and Figure S3 in Supporting Information S1). Note that some regions are still cooling post-1975 in this simulation, which biased some of our testing (see results).

The loss function used for optimization is

$$\mathcal{L}_t = \sum_{n=1}^{PN} \frac{1}{C Y X} \sum_{j=1}^C \sum_{k=1}^Y \sum_{l=1}^X (\tilde{\Phi}_{t+n\Delta t}^{[j,k,l]} - \Phi_{t+n\Delta t}^{[j,k,l]})^2. \quad (2)$$

$\mathcal{L}_t$ is the total mean square error (MSE) loss function at time step t , where P corresponds to the total number of input/output states used by the model in a single step, N is the total number of recurrent passes, C , Y and X are the total number of output channels, height and width, respectively, of a single output state. Here, we set $P = 2$ to obtain a 2-input 2-output model configuration and $N = 4$ steps.

We use the Adam optimizer with a learning rate of $2e - 4$, which decays to zero using a Cosine scheduler. Our emulators are trained using 4 80GB A100 GPUs for 15 and 12 hr for the models $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ and $\mathcal{F}_{\text{thermo}}$ respectively, with a total batch size of 16.

2.5. Evaluation

To evaluate the emulators, we take our initial conditions from 2014-09-30 and produce an 8-year rollout using the corresponding atmospheric forcing. We compare the output from this rollout to held-out OM4 data to evaluate the emulator skill. In addition, we produce longer runs to assess the emulator's response, similar to control simulations, with arbitrarily long rollouts. The emulator is forced with atmospheric boundary conditions taken from 1990 to 2000, with a repeat 10-year cycle. This period is chosen specifically because it has a near-zero globally integrated heat flux forcing, which ensures minimal ocean drift. We also performed a 100-year and a 400-year control run (see Supporting Information S1).

We produce predictions using both $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ and $\mathcal{F}_{\text{thermo}}$. All evaluations use a single 40 GB A100 GPU. For each year of rollout, $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ and $\mathcal{F}_{\text{thermo}}$ take about 90.52 s and 47.2 s, respectively. Thus, for the faster emulator, a century rollout takes approximately 1.3 hr $\mathcal{F}_{\text{thermo}}$ takes approximately half the time to produce the same number of states in the rollout compared to $\mathcal{F}_{\text{thermo} + \text{dynamic}}$.

3. Results

3.1. Full-Depth Global Ocean Emulator

We begin by evaluating the emulators $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ and $\mathcal{F}_{\text{thermo}}$ against the ground truth to establish a baseline skill. Capturing the full-depth climatological profiles of potential temperature and salinity is a key target of ocean numerical climate models in general and, therefore, a key target for our ocean climate emulators. The structure of the zonal mean of potential temperature (Figure 1b) is captured by the two emulators, demonstrating significant skill at reproducing the profile from OM4 (see Figure S6 in Supporting Information S1 for salinity structure). The average mean absolute error (MAE) is $5.7 \times 10^{-3} \text{ }^\circ\text{C}$ for $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ and $4.5 \times 10^{-3} \text{ }^\circ\text{C}$ for $\mathcal{F}_{\text{thermo}}$, with a pattern correlation of roughly 0.99 for both emulators. The outputs show a robust thermocline structure, subtropical gyres, and a region of North Atlantic deep water formation. However, both emulators in the northern hemisphere show too warm and too salty high latitudes (around 55N), too cold and too fresh mid-latitudes, and Arctic signals down to 750m depth (Figures S2 and S7 in Supporting Information S1). The biases are consistent with underestimating the northward heat transport by the ocean. The potential temperature and salinity biases in the Southern Ocean for the $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ emulator are reminiscent of residual transport changes, with opposite signed biases in the Southern Ocean and in the region north of it. The $\mathcal{F}_{\text{thermo}}$ emulator is warmer than $\mathcal{F}_{\text{thermo} + \text{dynamic}}$, at most depths (Figure S2 in Supporting Information S1).

We performed several experiments to test the sensitivity of the emulators to different training choices. The emulators' skill is unchanged when using different seeds and start dates, so the trained models are statistically reproducible. We measure robustness by calculating the root mean square error of rollouts with 5 different seeds and rollouts initialized with ocean states taken 6 months apart. The RMSEs show little variance across the different trained models (Figure 1c). The standard deviation of the RMSEs across training seeds in the emulators $\mathcal{F}_{\text{thermo}}$ and $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ are 0.0033 and 0.00225, respectively.

The potential-temperature timeseries at 2.5 m and 775 m (Figure 2a) are further indicators that both emulators capture the climatological means and the upper ocean response to variable atmospheric forcing. The standard deviation of the 2.5 m potential temperature for OM4, and the emulators $\mathcal{F}_{\text{thermo}}$ and $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ are $6.8 \times 10^{-2} \text{ }^\circ\text{C}$, $4.35 \times 10^{-2} \text{ }^\circ\text{C}$ and $5.26 \times 10^{-2} \text{ }^\circ\text{C}$ respectively, while the standard deviations of the 775m potential temperature are $2.3 \times 10^{-3} \text{ }^\circ\text{C}$, $1.0 \times 10^{-3} \text{ }^\circ\text{C}$ and $2.1 \times 10^{-3} \text{ }^\circ\text{C}$, respectively. The emulators capture a large portion of the variability, but with some biases (Figure 2b). The standard deviations are calculated after removing both the trend and the climatology from the timeseries (See Figure S8 in Supporting Information S1 for




Figure 2. (a) Spatially averaged timeseries of potential temperature θ_0 at depths 2.5 m (left) and 775 m (right) over the test set comparing the ground truth OM4 (black), and predictions from F_{thermo} (red) and $F_{\text{thermo+dynamic}}$ (green). The mean prediction and its variance (indicated by shading) are plotted over five initial seeds of training for each model. (b) Nino 3.4 index timeseries over the test set for the ground truth (OM4, black) and predictions (F_{thermo} , red; $F_{\text{thermo+dynamic}}$, green). Anomalies are averaged over rolling 150-day windows. (c)–(e) Meridionally averaged depth profile of potential temperature anomalies in the tropics during the peak Nino event (marked by a black dot in the timeseries) over the test set for OM4 (c), F_{thermo} (d) and $F_{\text{thermo+dynamic}}$ (e). Anomalies in (c)–(e) are averaged over a 15-day window.

additional timeseries of potential temperature, along with salinity, zonal velocity, and meridional velocity, and Figure S10 in Supporting Information S1 for bias maps).

The emulator can skillfully emulate ENSO response in both warm and cold phases (Figure 2b) and Figure S11 in Supporting Information S1). The smallest fluctuations in the Nino 3.4 timeseries are the hardest for the emulators to capture. The emulator responses are in phase with OM4 for all years shown, but the amplitude is altered. $F_{\text{thermo+dynamic}}$ exhibits higher skill than F_{thermo} in capturing the magnitude of ENSO events. We hypothesized that providing the velocities, whose data contain shorter time-scales and larger variability, helps the emulator produce larger ENSO events. F_{thermo} still manages to detect the correct phase and structure (Figures 2b and 2d) despite producing events with smaller magnitudes, both at the surface and in the upper ocean. The emulators capture the deepening and shoaling of the equatorial thermocline from equatorial Kelvin waves for the strongest events (Figures 2d and 2e). The magnitude of subsurface anomalies for the emulators is weaker than for OM4. For the Nino 3.4 timeseries (Figure 2b), the MAE is 0.0077 °C for $F_{\text{thermo+dynamic}}$ and 0.0124 °C for F_{thermo} , with correlations of 0.905 and 0.7017, respectively. For the ENSO profiles (Figures 2c–2e), the MAE is 0.01 °C and 0.07 °C for the emulators $F_{\text{thermo+dynamic}}$ and F_{thermo} respectively, and their pattern correlations are 0.976 and 0.973, respectively.

For the ocean emulator $F_{\text{thermo+dynamic}}$ that uses all variables, we noticed that the potential temperature and salinity fields exhibit atypically high spatial variability, with scales more characteristic of velocity so we posit that this results from using velocity inputs. See Figures S16–S17 in Supporting Information S1 for maps of variability for our emulators. This result is consistent with Subel and Zanna (2024). We hypothesize that this may arise from the large separation in timescales and variability between velocity and potential temperature in the ocean.

Finally, despite capturing the mean and climatology of ocean variables, the emulators struggle to capture the magnitude of the small, but systematic potential temperature trends (Figure S1 in Supporting Information S1




Figure 3. (a) Globally averaged potential temperature (θ_0) timeseries over a 100-year control run, comparing the 10-year ground truth OM4 (black) and predictions from F_{thermo} (red) and $F_{thermo+dynamic}$ (green). (b) Time-averaged potential temperature (θ_0) depth profile over a 100-year control run, comparing the 10-year ground truth OM4 (left) and predictions from F_{thermo} (middle) and $F_{thermo+dynamic}$ (right). (c) Nino 3.4 index timeseries over a 100-year control run, comparing the 10-year repeat for the ground truth (OM4, black) and predictions (F_{thermo} , red; $F_{thermo+dynamic}$, green). (d)–(e) Meridionally averaged depth profile of potential temperature anomalies in the tropics during the peak Nino event (marked by a black dot in the timeseries) over the test set for F_{thermo} (d) and $F_{thermo+dynamic}$ (e). Anomalies are as in Figure 2.

global mean 10^{-3} °C/yr) over the same 8-year period (Figure 2a and Figures S1 and S3 in Supporting Information S1); for most depths the trained models underestimate trends by 20%–50% relative to OM4. Of the two emulators, F_{thermo} has higher skill in capturing the global heat changes (Figure S9 in Supporting Information S1). The salinity trends in OM4 are weak, due to the small forcing, and to the use of salinity restoring boundary conditions. For both emulators, the trends are 7–8 orders of magnitude less than the mean value, consistent with the numerical representation of variables within the learned models, suggesting that the models conserve properties of the OM4 data although strict conservation is not imposed (Figures S4–S5 in Supporting Information S1).

3.2. Long-Term Stability

We also evaluated the ability of the emulators to produce long control experiments, without retraining. For these experiments, we use repeat boundary conditions over 10 years (described in Section 2.5) chosen to contribute a near-zero net heat flux, allowing the emulators to run for arbitrarily long periods of time while minimizing potential temperature drift.

Both emulators converge to an equilibrium, maintaining a global mean potential temperature close to OM4 throughout a century of integration (Figure 3a). The global mean temperatures are 3.225°C/yr for $\mathcal{F}_{\text{thermo}}$ and 3.215°C/yr for $\mathcal{F}_{\text{thermo} + \text{dynamic}}$, compared to 3.219°C/yr for OM4. In addition, $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ over-predicts the variability in potential temperature, likely extrapolating some fast dynamics via the velocities variables. This issue is exacerbated in the deeper layers of the ocean, which have little variability in the original data set. The temperature structure is again well preserved for the long rollouts (Figure 3b), with different structures in potential temperature biases (Figure S12 in Supporting Information S1) than for the 8-year test data (Figure S2 in Supporting Information S1).

We examine the emulators' skill in reproducing variability over these long timescales. Since we are reusing the same 10-year cycle to drive the emulator, we expected some persistent features to appear when looking at a phenomenon such as the response to ENSO. Although both emulators can produce appropriate Nino 3.4 anomalies for the entire century rollout ((Figure 3c) and Figure S13 in Supporting Information S1), $\mathcal{F}_{\text{thermo} + \text{dynamic}}$ shows stronger peak-to-peak amplitude, but little cycle-to-cycle variability—perhaps due to the strong coupling of velocity with the wind stress forcing, whereas $\mathcal{F}_{\text{thermo}}$ shows more aperiodic variability across years.

To further test stability, we generate a 400-year rollout, with an identical forcing setup as for the century-long run. Both emulators remain stable (Figure S15 in Supporting Information S1). $\mathcal{F}_{\text{thermo}}$ has the added benefit of exhibiting long-term aperiodic variability in potential temperature and salinity, despite the repeat forcing, across the centuries. The long experiments were reproduced using a repeat forcing period from the test set that is, 2014–2022, producing similar results (Figure S19 in Supporting Information S1).

4. Discussion

We produce a computationally cheap machine-learning (ML) emulator of a state-of-the-art ocean model, namely OM4 (Adcroft et al., 2019). The ML architecture consists of a modified ConvNeXt UNet (Dheeshjith et al., 2024). The reduced order model—*Samudra*—predicts key ocean variables, sea surface height, temperature, and salinity, across the full depth of the world oceans while remaining stable for centuries. Integrating OM4 for 100 years takes approximately 8 days using 4,671 CPU cores, whereas our fastest (thermo) emulator completes the same task in about 1.3 hr on a single 40GB A100 GPU. This represents approximately a 150x increase in simulated years per day for *Samudra* compared to OM4. Some of this speed up can be attributed to *Samudra*: (a) using a 5 days time step (vs. 15 minutes in OM4); (b) operating on a spatially coarser grid. However, we note that *Samudra* makes predictions with the implicit spatial skill of the finer resolution OM4, whereas existing coarser General Circulation Model (GCMs) with eddy parameterization tend to show worse biases (e.g., see Figure 9 of Adcroft et al. (2019) for a 1/2-degree GCM).

The emulator performs well on a range of metrics related to the model climatology and its variability on the test set and long control simulations. The emulator produces accurate climatologies over the last 8 years of the OM4 simulations and is robust to changes in seeds and initial conditions. Furthermore, it can capture variability (e.g., ENSO response to forcing). Therefore, these emulators could be used to study the contemporary ocean and climate at a significant reduction in cost compared to OM4.

The emulator, however, struggles to capture trends under a range of surface heat flux forcings (see Supporting Information S1), similarly to the surface emulators in Dheeshjith et al. (2024). We performed idealized forced experiments using the same repeated atmospheric forcing generated for the control experiment and a spatially uniform linear forcing of varying magnitudes for the surface heat flux. Figure S16 in Supporting Information S1 showcases the ocean heat content trends predicted by $\mathcal{F}_{\text{thermo}}$ under linear surface heat flux increases of 1, 0.5, 0.25, and 0 W/m^2 . The patterns of ocean heat uptake are reminiscent of ocean-only and coupled forced numerical experiments (Couldrey et al., 2020; Todd et al., 2020), with dipole patterns in the Southern Ocean and North Atlantic sinking region (Figure S14 in Supporting Information S1). However, the magnitude of change is too weak compared to the forcing (Figure S16 in Supporting Information S1). Similar behavior of weak generalization under climate change is also observed in the atmosphere climate emulator, ACE (Watt-Meyer et al., 2023), but improved when a slab ocean model is added (Clark et al., 2024).

Here, we could not produce an emulator that simultaneously captures the trends in the test data and remain stable for centuries. Further work is needed to explore the reasons for the issues and would require new numerical simulations.

The lack of generalization reflected in the weak warming trends could be due to the training data. The effects of an initial drift can be alleviated by pruning years 1958–1975 from the training data, which removes the bulk of this adjustment period. Yet, different depths and regions adjust more slowly, and some of this continued adjustment may remain in the data since the time scale of equilibration of the model is hundreds of years. Another reason for the trend bias could be the forcing data sets. The atmospheric forcing imposed on the ocean implicitly results from the real ocean-atmosphere coupling. Therefore, the atmospheric forcing has felt a changing ocean circulation, particularly in the North Atlantic (Chemke et al., 2020). The resulting effect is that the “forcing” applied to the ocean emulator is not entirely decoupled from the ocean response, potentially leading to some biases in the response, as in Todd et al. (2020); Couldrey et al. (2020); Zanna et al. (2019). We alleviated these issues by adding an extra forcing input, namely the cumulative heat forcing, which led to a more skillful model capable of capturing the global warming trend. However, this model was unstable under climate-change forcing past 50 years. Alternatively, it is possible that learning to predict the model state directly may not be optimal. We explored learning tendencies, which improved performance for the warming trends but, again, was unstable over long timescales. A challenge going forward is designing faithful emulators capable of capturing trends while remaining stable in long rollouts.

Despite the limited response to future climate forcing, Samudra is skillful at emulating the contemporary ocean and is therefore an affordable emulation of expensive ocean circulation models. Without further modification, Samudra could be used in studies requiring large ensembles (e.g., uncertainty quantification, extreme events) or to enhance and accelerate operational applications (e.g., data assimilation). More opportunities emerge if we consider refining training for Samudra, for example, to revised versions of OM4 or to other models, which could greatly accelerate climate model development by allowing evaluations of long, yet affordable, rollouts. This includes coupling Samudra with atmosphere climate emulator (Watt-Meyer et al., 2023) to emulate CM4.

Data Availability Statement

The code for training the models along with generating rollouts and plots is available on GitHub at <https://github.com/m2lines/Samudra>, while the model weights and data are hosted on Hugging Face at <https://huggingface.co/M2LInES/Samudra> and <https://huggingface.co/datasets/M2LInES/Samudra-OM4>, respectively. Additionally, data from National Centers for Environmental Prediction et al. (2006) was also used in the Supporting Information S1. The code is also version tagged and archived at Dheeshjith (2025) via zenodo.

Acknowledgments

This research received support through Schmidt Sciences, LLC, under the M²LInES project. We thank all members of the M²LInES team for helpful discussions and their support throughout this project. We gratefully acknowledge Karthik Kashinath and the NVIDIA team for providing us access to NERSC resources, which were instrumental in supporting this work. This research was also supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. We also thank the reviewers for their useful comments.

References

- Abernathy, R. P., Busecke, J. J. M., Smith, T. A., Deauna, J. D., Banihirwe, A., Nicholas, T., et al. (2022). xgcm. *Zenodo*. <https://doi.org/10.5281/zenodo.7348619>
- Adcroft, A., Anderson, W., Balaji, V., Blanton, C., Bushuk, M., Dufour, C. O., et al. (2019). The GFDL Global Ocean and sea ice model OM4.0: Model description and simulation features. *Journal of Advances in Modeling Earth Systems*, 11(10), 3167–3211. <https://doi.org/10.1029/2019MS001726>
- Arcomano, T., Szunyogh, I., Wikner, A., Hunt, B. R., & Ott, E. (2023). A hybrid atmospheric model incorporating machine learning can capture dynamical processes not captured by its physics-based component. *Geophysical Research Letters*, 50(8), e2022GL102649. <https://doi.org/10.1029/2022gl102649>
- Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2023). Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970), 533–538. <https://doi.org/10.1038/s41586-023-06185-3>
- Bire, S., Lütjens, B., Azzizadenesheli, K., Anandkumar, A., & Hill, C. N. (2023). Ocean emulation with fourier neural operators: Double gyre. *Authorea Preprints*.
- Cachay, S. R., Henn, B., Watt-Meyer, O., Bretherton, C. S., & Yu, R. (2024). Probabilistic emulation of a global climate model with spherical dyffusion. *arXiv preprint arXiv:2406.14798*.
- Chemke, R., Zanna, L., & Polvani, L. M. (2020). Identifying a human signal in the north atlantic warming hole. *Nature Communications*, 11(1), 1540. <https://doi.org/10.1038/s41467-020-15285-x>
- Clark, S. K., Watt-Meyer, O., Kwa, A., McGibbon, J., Henn, B., Perkins, W. A., et al. (2024). Ace2-som: Coupling to a slab ocean and learning the sensitivity of climate to changes in CO₂. *arXiv preprint arXiv:2412.04418*.
- Couldrey, M. P., Gregory, J. M., Dias, F. B., Dobrohotoff, P., Domingues, C. M., Garuba, O., et al. (2020). *What causes the spread of model projections of ocean dynamic sea-level change in response to greenhouse gas forcing?* (pp. 1–33). Climate Dynamics.
- Dheeshjith, S. (2025). suryadheeshjith/samudra: v1.0.0 - First stable release [software]. *Zenodo*. <https://doi.org/10.5281/ZENODO.15037462>
- Dheeshjith, S., Subel, A., Gupta, S., Adcroft, A., Fernandez-Granda, C., Busecke, J., & Zanna, L. (2024). Transfer learning for emulating ocean climate variability across CO₂ forcing. *arXiv preprint arXiv:2405.18585*.
- Gray, M. A., Chattopadhyay, A., Wu, T., Lowe, A., & He, R. (2024). *Long-term prediction of the gulf stream meander using oceannet: A principled neural operator-based digital twin* (Vol. 2024, pp. 1–23). EGU sphere.
- Griffies, S. M., Danabasoglu, G., Durack, P. J., Adcroft, A. J., Balaji, V., Böning, C. W., et al. (2016). OMIP contribution to CMIP6: Experimental and diagnostic protocol for the physical component of the ocean model intercomparison project. *Geoscientific Model Development*, 9(9), 3231–3296. <https://doi.org/10.5194/gmd-9-3231-2016>

- Guo, Z., Lyu, P., Ling, F., Luo, J.-J., Boers, N., Ouyang, W., & Bai, L. (2024). Orca: A global ocean emulator for multi-year to decadal predictions. *arXiv preprint arXiv:2405.15412*.
- Held, I. M., Guo, H., Adcroft, A., Dunne, J. P., Horowitz, L. W., Krasting, J., et al. (2019). Structure and performance of GFDL's CM4.0 climate model. *Journal of Advances in Modeling Earth Systems*, 11(11), 3691–3727. <https://doi.org/10.1029/2019MS001829>
- Holmberg, D., Clementi, E., & Roos, T. (2024). Regional ocean forecasting with hierarchical graph neural networks. *arXiv preprint arXiv:2410.11807*.
- Karlbauer, M., Cresswell-Clay, N., Durran, D. R., Moreno, R. A., Kurth, T., & Butz, M. V. (2023). Advancing parsimonious deep learning weather prediction using the healpix mes. *Authorea Preprints*.
- Khatiwal, S. (2024). Efficient spin-up of earth system models using sequence acceleration. *Science Advances*, 10(18), eadn2839. <https://doi.org/10.1126/sciadv.adn2839>
- Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., et al. (2024). Neural general circulation models for weather and climate. *Nature*, 632(8027), 1–7. <https://doi.org/10.1038/s41586-024-07744-y>
- Levitov, S., Boyer, T., Garcia, H., Locarnini, R., Zweng, M., Mishonov, A., et al. (2015). World ocean atlas 2013 (NCEI accession 0114815). <https://doi.org/10.7289/v5f769gt>
- Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., & Xie, S. (2022). A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11976–11986).
- Loose, N., Abernathy, R., Grooms, I., Busecke, J., Guillaumin, A., Yankovsky, E., et al. (2022). Gcm-filters: A python package for diffusion-based spatial filtering of gridded data. [Software]. *Journal of Open Source Software*, 7(70), 3947. <https://doi.org/10.21105/joss.03947>
- Maher, N., Milinski, S., & Ludwig, R. (2021). Large ensemble climate model simulations: Introduction, overview, and future prospects for utilising multiple types of large ensemble. *Earth System Dynamics*, 12(2), 401–418. <https://doi.org/10.5194/esd-12-401-2021>
- Mahesh, A., Collins, W., Bonev, B., Brenowitz, N., Cohen, Y., Elms, J., et al. (2024). Huge ensembles part i: Design of ensemble weather forecasts using spherical fourier neural operators. *arXiv preprint arXiv:2408.03100*.
- Manshausen, P., Cohen, Y., Pathak, J., Pritchard, M., Garg, P., Mardani, M., et al. (2024). Generative data assimilation of sparse weather station observations at kilometer scales. *arXiv preprint arXiv:2406.16947*.
- National Centers for Environmental Prediction, National Weather Service, NOAA, U.S. Department of Commerce. (2006). NCEP global ocean data assimilation system (GODAS) [dataset]. *Boulder CO: Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory*. <https://doi.org/10.5065/ZAHZ-3174>
- Patel, D., Arcomano, T., Hunt, B., Szunyogh, I., & Ott, E. (2024). Exploring the potential of hybrid machine-learning/physics-based modeling for atmospheric/oceanic prediction beyond the medium range. *arXiv preprint arXiv:2405.19518*.
- Price, I., Sanchez-Gonzalez, A., Alet, F., Andersson, T. R., El-Kadi, A., Masters, D., et al. (2023). Gencast: Diffusion-based ensemble forecasting for medium-range weather. *arXiv preprint arXiv:2312.15796*.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—miccai 2015: 18th international conference, munich, Germany, october 5–9, 2015, proceedings, part iii* (Vol. 18, pp. 234–241). https://doi.org/10.1007/978-3-319-24574-4_28
- Sane, A., Reichl, B. G., Adcroft, A., & Zanna, L. (2023). Parameterizing vertical mixing coefficients in the ocean surface boundary layer using neural networks. *Journal of Advances in Modeling Earth Systems*, 15(10), e2023MS003890. <https://doi.org/10.1029/2023ms003890>
- Subel, A., & Zanna, L. (2024). Building ocean climate emulators. *arXiv preprint arXiv:2402.04342*.
- Todd, A., Zanna, L., Couldrey, M., Gregory, J., Wu, Q., Church, J. A., et al. (2020). Ocean-only fafmip: Understanding regional patterns of ocean heat content and dynamic sea level change. *Journal of Advances in Modeling Earth Systems*, 12(8), e2019MS002027. <https://doi.org/10.1029/2019ms002027>
- Tsujino, H., Urakawa, L. S., Griffies, S. M., Danabasoglu, G., Adcroft, A. J., Amaral, A. E., et al. (2020). Evaluation of global ocean–sea-ice model simulations based on the experimental protocols of the Ocean Model Intercomparison Project phase 2 (OMIP-2). *Geoscientific Model Development*, 13(8), 3643–3708. <https://doi.org/10.5194/gmd-13-3643-2020>
- Wang, C., Pritchard, M. S., Brenowitz, N., Cohen, Y., Bonev, B., Kurth, T., et al. (2024). Coupled ocean-atmosphere dynamics in a machine learning earth system model. *arXiv preprint arXiv:2406.08632*.
- Watt-Meyer, O., Dresdner, G., McGibbon, J., Clark, S. K., Henn, B., Duncan, J., et al. (2023). Ace: A fast, skillful learned global atmospheric model for climate prediction. *arXiv preprint arXiv:2310.02074*.
- Xiong, W., Xiang, Y., Wu, H., Zhou, S., Sun, Y., Ma, M., & Huang, X. (2023). Ai-goms: Large ai-driven global ocean modeling system. *arXiv preprint arXiv:2308.03152*.
- Zanna, L., Khatiwal, S., Gregory, J. M., Ison, J., & Heimbach, P. (2019). Global reconstruction of historical ocean heat storage and transport. *Proceedings of the National Academy of Sciences*, 116(4), 1126–1131. <https://doi.org/10.1073/pnas.1808838115>
- Zhuang, J., raphael, d., Huard, D., Bourgault, P., Banihirwe, A., Raynaud, S., et al. (2023). pangeo-data/xesmf: v0.8.2. *Zenodo*. <https://doi.org/10.5281/zenodo.8356796>