



RESEARCH ARTICLE

10.1029/2021MS002534

Stochastic-Deep Learning Parameterization of Ocean
Momentum ForcingArthur P. Guillaumin¹ and Laure Zanna¹ ¹Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

Key Points:

- We use data from a realistic high-resolution coupled climate model to train a neural network (NN) and perform offline tests
- We learn a stochastic parameterization of subgrid momentum forcing with the NN that generalizes well to different regimes and climates
- The stochastic parameterization results in a stable implementation in an idealized ocean model that improves the kinetic energy of the flow

Correspondence to:

A. P. Guillaumin,
ag7531@nyu.edu

Citation:

Guillaumin, A. P., & Zanna, L. (2021). Stochastic-deep learning parameterization of ocean momentum forcing. *Journal of Advances in Modeling Earth Systems*, 13, e2021MS002534. <https://doi.org/10.1029/2021MS002534>Received 2 MAR 2021
Accepted 15 JUL 2021

Abstract Coupled climate simulations that span several hundred years cannot be run at a high-enough spatial resolution to resolve mesoscale ocean dynamics. Recently, several studies have considered Deep Learning to parameterize subgrid forcing within macroscale ocean equations using data from ocean-only simulations with idealized geometry. We present a stochastic Deep Learning parameterization that is trained on data generated by CM2.6, a high-resolution state-of-the-art coupled climate model. We train a Convolutional Neural Network for the subgrid momentum forcing using macroscale surface velocities from a few selected subdomains with different dynamical regimes. At each location of the coarse grid, rather than predicting a single number for the subgrid momentum forcing, we predict both the mean and standard deviation of a Gaussian probability distribution. This approach requires training our neural network to minimize a negative log-likelihood loss function rather than the Mean Square Error, which has been the standard in applications of Deep Learning to the problem of parameterizations. Each estimate of the conditional mean subgrid forcing is thus associated with an uncertainty estimate—the standard deviation—which will form the basis for a stochastic subgrid parameterization. Offline tests show that our parameterization generalizes well to the global oceans and a climate with increased CO₂ levels without further training. We then implement our learned stochastic parameterization in an eddy-permitting idealized shallow water model. The implementation is stable and improves some statistics of the flow. Our work demonstrates the potential of combining Deep Learning tools with a probabilistic approach in parameterizing unresolved ocean dynamics.

Plain Language Summary Numerical predictions are pivotal to understanding the impact of climate change. However, those predictions are limited in accuracy by the trade-off between the models' spatio-temporal resolution and their time span, due to the large computational cost. Since small-scale dynamics impact larger-scale dynamics, a common approach is to use idealized equations, based on the practitioner's understanding of physics, to account for the impact of unresolved small-scale dynamics on the large-scale flow. However, this approach has shown its limits. Recently, several studies have considered the use of Deep Learning—a set of techniques designed to learn high-dimensional complex functions from large amounts of data—to learn the impact of small-scale dynamics on the large-scale flow. Here we apply Deep Learning methods using simulated data from a state-of-the-art climate model. Additionally, we account for the uncertainty associated with the learned representation of the impact of the small scales on the large scale. Our tests using this representation in an idealized ocean model show that some metrics are improved. Much work remains to be done to assess the success of Deep Learning in improving climate models.

1. Introduction

The climate system is governed by highly nonlinear equations, making them inherently multiscale, with small-scale processes backscattering to large scales. Fluid dynamics equations are known and valid in a continuum. However, climate models solve fluid dynamics equations on a grid, resulting in approximate solutions. Ideally, increasing the spatio-temporal resolution could improve these truncated simulations. However, even with the increasing available computational power, running high-resolution climate models over decades or centuries is not a viable approach within the near future (Balaji, 2021). Typically, the impact of unresolved small-scale processes on coarse quantities is accounted for via parameterizations. These parameterizations are commonly based on first principles (Gent & McWilliams, 1989; Germano et al., 1991;

© 2021. The Authors. Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Fox-Kemper & Menemenlis, 2008; Lilly, 1992; Marshall & Adcroft, 2010). Despite vastly improving the physics of climate models, current parameterizations continue to induce biases in simulations (IPCC, 2013).

The era of Machine Learning offers an opportunity to improve the parameterization of unresolved processes using available data from observations and limited high-resolution simulations. While some progress has been made toward *online* learning of neural networks (NN) for unresolved processes in partial differential equations (Sirignano et al., 2020), the approach is not yet ready for complex climate simulations and might not be generalizable due to model dependence. Therefore, the typical approach in atmosphere and ocean modeling consists in training Machine Learning algorithms offline, with a subgrid forcing term that is diagnosed via a filtering operation over high-resolution simulation data, used as the “truth”. Recent studies have shown the potential of Machine Learning approaches for atmospheric (Rasp et al., 2018; Yuval & O’Gorman, 2020) and ocean parameterizations (Bolton & Zanna, 2019; Zanna & Bolton, 2020) to improve simulations. So far, most studies on ocean parameterizations that use Machine Learning have been limited to the use of data from ocean-only models with idealized geometry. The viability of Deep Learning parameterizations trained using data from realistic coupled or uncoupled models and their potential to generalize to different climates both remain open questions. In addition, the stability and the physical behavior of implementations of Deep Learning parameterizations in models have also been a subject of debate (Brenowitz et al., 2020; Yuval & O’Gorman, 2020; Zanna & Bolton, 2020).

Here we address these questions by showing the high performance of a Deep Neural Network in offline predictions of subgrid momentum forcing in different climates using data from a high-resolution coupled climate model, which resolves ocean mesoscale eddies in many regions (Griffies et al., 2015; Hallberg, 2013), and show which will result in a stable implementation online. Our work focuses on parameterizing the interaction between mesoscale eddies and large-scale flow, which is key to establishing the transfer of energy between reservoirs and scales (Ferrari & Wunsch, 2009) and to establishing the large-scale ocean circulation (Waterman & Jayne, 2011). In particular, we propose a stochastic parameterization that aims to represent the inherent uncertainty of the subgrid forcing (Juricke et al., 2017), stabilize the online implementation of the parameterization (Palmer, 2012; Zanna et al., 2017) and reduce systematic biases (Berner et al., 2017; Gagne II et al., 2020). Stochastic parameterizations are in demand in particular in what has been referred to as the *gray zone* (Gerard, 2007; Jones et al., 2019; Zanna et al., 2020), where subgrid processes are partly resolved (Berner et al., 2017).

In our study, our NN model outputs the mean and standard deviation for the predicted momentum forcing, given coarse-resolution velocities. The predictions of mean and standard deviation of the momentum forcing form the basis of a stochastic parameterization that we will implement in an idealized ocean model. Our contribution therefore establishes a bridge between recent developments on Deep Learning approaches to the problem of stochastic parameterizations (Berloff, 2005; Brankart, 2013; Frederiksen & Davies, 1997; Juricke et al., 2017; Mason & Thomson, 1992; Williams et al., 2016; Zanna et al., 2017).

The manuscript is structured as follows. In Section 2, we describe the data, the NN architecture and the training procedure—which uses a probabilistic loss function. In Section 3, we conduct an offline test on a global scale, showing the ability of our NN to generalize to regions not seen during training. We also show the ability of our NN to generalize to a different climate in which CO₂ levels are higher and have affected the mesoscale variability. In Section 4, we demonstrate the potential for increased stability via an implementation of our stochastic parameterization into an idealized ocean model, without further training of the NN. Finally, in Section 5 we conclude and discuss the implications of our work and future directions.

2. Methods

In Sections 2.1 and 2.2, we describe the filtering and subsequent coarsening of the data in order to diagnose the corresponding subgrid momentum forcing necessary to force a coarse-resolution model. In Section 2.3 and 2.4, we describe a procedure that enables us to represent the uncertainty associated with the forcing using a probabilistic loss function for training. In Section 2.5, we review the structure of our proposed NN. Finally, in Section 2.6 we provide details about our training procedure.

2.1. Data for Training and Validation

Applications of Deep Learning to the parameterization of subgrid ocean momentum forcing and backscatter have been limited to ocean-only idealized models so far (Bolton & Zanna, 2019; Zanna & Bolton, 2020). In contrast, here we investigate the use of Deep Learning on data from a state-of-the-art high-resolution coupled climate model, CM2.6 (Delworth et al., 2012; Griffies et al., 2015). The nominal horizontal resolution of the ocean component of CM2.6 is $1/10^\circ$, therefore resolving mesoscale eddies in many regions of the ocean (Hallberg, 2013). The data and tools for analysis were obtained from the Pangeo platform (Abernathey et al., 2021).

The data used in the present work consists of the high-resolution simulated ocean surface velocity field \mathbf{u}^\uparrow with components u^\uparrow (zonal) and v^\uparrow (meridional), where the upper arrow is used to denote quantities that sit on the high-resolution grid. The model grid is configured according to an Arakawa *B*-grid (Griffies, 2015), with velocity points (both zonal and meridional) placed to the North-East of tracer (*T*) cells, i.e. the top-right corner of a *T*-cell. The temporal resolution of the surface velocity data is daily, and the available data span over $\sim 7,000$ days (about 20 years) for each of the two available simulations—a control simulation with pre-industrial atmospheric CO_2 levels, referred to as piControl, and a forced simulation with a 1% CO_2 increase per year, referred to as 1pt CO_2 (Griffies et al., 2015). The 1pt CO_2 simulation experiences a one percent increase of CO_2 per year from the levels of the control simulation until it reaches doubling after 70 years, at which point the CO_2 levels remain constant. The 1pt CO_2 simulation data available from Pangeo corresponds to years 60–80.

2.2. Filtering and Coarse-Graining Procedure

In this section, we describe the processing necessary to generate the training data for our NN. The procedure follows the two steps presented in Zanna and Bolton (2020): low-pass area-weighted Gaussian filtering, followed by coarse-graining. Applied on the high-resolution surface velocities \mathbf{u}^\uparrow from the CM2.6 simulations, this procedure generates coarse-resolution velocity data that mimics the simulation data from coarser models and will serve as the input to our NN. In addition, given the high-resolution and filtered velocity data, we diagnose the missing subgrid forcing of a coarse-resolution model (e.g., CM2.5) compared to its high-resolution counterpart (here, CM2.6). This missing forcing is the subgrid parameterization needed at coarse resolution to mimic the effect of unresolved scales on the large-scale flow that will be learned by the NN.

Unlike data used in previous machine learning studies (Bolton & Zanna, 2019; Yuval & O’Gorman, 2020; Zanna & Bolton, 2020), the CM2.6 grid is on a sphere. In the zonal direction, the spacing is $1/10^\circ$, and in the meridional direction, it varies from approximately $1/10^\circ$ near the equator to $1/20^\circ$ at high latitudes, so that the zonal and meridional extent of grid boxes varies with latitude. The grids of CM2.5, with $1/4^\circ$ nominal resolution, and CM2.1, with 1° nominal resolution, have a similar structure. As such, the meridional length scale used to define the subgrid eddy forcing should depend on the latitude, which would in part capture the latitudinal variations of the Rossby deformation radius. In contrast with the typical approach, rather than selecting a uniform length scale to filter the data and generate a coarse-resolution field, we select a uniform and unitless integer *scaling factor* σ , that defines the number of grid boxes from the high-resolution grid that map to a single grid box of the low-resolution grid. This is the simplest and most consistent definition of subgrid scale for the purpose of data-driven parameterizations. This unitless scaling factor applies to both the filtering and coarse-graining steps.

We now describe in details the two steps of our low-resolution data generation procedure given the fixed scaling factor σ . As a first step we apply a low-pass weighted Gaussian filter, denoted by $(\cdot)^\sigma$, to the high-resolution surface velocity data, with weights provided by grid box areas, to separate the subgrid from the resolved field (Bolton & Zanna, 2019). The standard deviation of the Gaussian kernel is set to $\sigma / 2$, such that approximately 80% of its weight falls within the interval $[-\sigma / 2, \sigma / 2]$ of length σ . Note that in using a uniform scaling factor we also allow the use of standard convolution algorithms for regularly spaced data. This would not be possible if we were using a uniform length scale as we would then have to adapt the size of the filter in terms of number of grid points as a function of latitude, incurring a high computational cost to generate the training data. The second step simply consists of a coarse-graining procedure, resulting in the

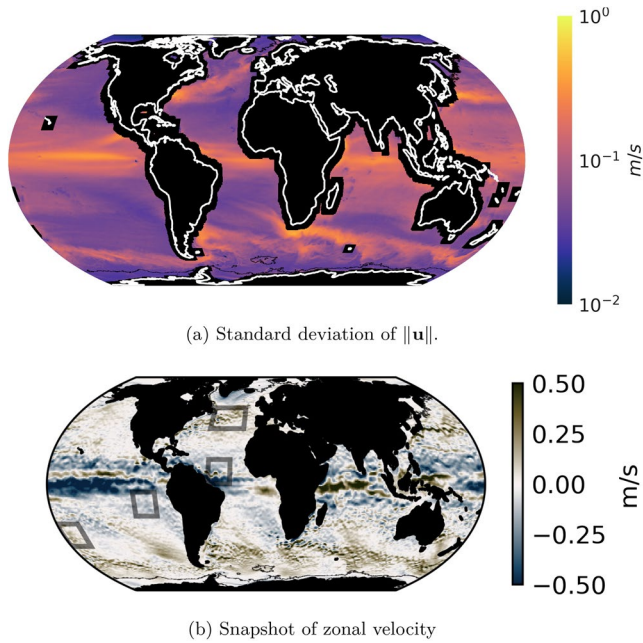


Figure 1. Filtered and coarse-grained surface velocity \mathbf{u} in $[m / s]$ from piControl used as training data: (a) standard deviation of surface velocity norm and (b) snapshot of the zonal component. The gray rectangles identify the subdomains selected for training the Neural Network (NN).

filtered and coarse-grained velocity field $\mathbf{u} = \langle \overline{\mathbf{u}^\uparrow} \rangle_\sigma$, where $\langle \cdot \rangle_\sigma$ denotes the coarse-graining operator by a factor σ . The coarse-graining operator is based on the mean function applied over squares of side length σ – equivalent to area-weighted average. After coarse-graining, the resulting grid consists of approximately σ^2 times less points than the high-resolution grid.

This filtering and coarse-graining procedure is applied to surface velocities from the CM2.6 control simulation. Figure 1 shows a snapshot of the filtered and coarse-grained surface zonal velocity. In addition, the subgrid momentum forcing on the high-resolution grid, denoted $\mathbf{S}^\uparrow = [S_x^\uparrow, S_y^\uparrow]^T$, is diagnosed via,

$$\mathbf{S}^\uparrow = \left(\overline{\mathbf{u}^\uparrow \cdot \nabla} \right) \overline{\mathbf{u}^\uparrow} - \overline{\left(\mathbf{u}^\uparrow \cdot \nabla \right) \mathbf{u}^\uparrow}, \quad (1)$$

and is then coarse-grained according to $\mathbf{S} = \langle \mathbf{S}^\uparrow \rangle_\sigma$. For exact implementation details of the procedure in the form of pseudo-code, please see Appendix B or the online code.

In our work, we primarily target parameterization for the eddy-permitting regime, or the gray zone, in which momentum parameterizations are in demand to mimic the inverse energy cascade (backscatter) processes while potentially overcoming spurious dissipation (Bachman, 2019; Charney, 1971; Jansen et al., 2019; Kjellsson & Zanna, 2017; Kraichnan, 1967; Larichev & Held, 1995; Leith, 1990; Rhines, 1977; Salmon, 1980; Treguier et al., 1997; Zanna et al., 2017, 2020). Here, we present experiments in which we set $\sigma = 4$, such that, irrespective of the subdomains of study,

the coarse-grained grid has ~ 4 times less grid boxes along each horizontal dimension. This choice of four grid boxes leads to a subgrid forcing of an ocean model at resolution 0.4° , close to the resolution of ESM4 which has a nominal resolution of 0.5° (Dunne et al., 2020).

The filtering and coarse-graining procedure is applied to both data from the piControl and 1pctCO2 simulations. The piControl data set will be used both for training and offline testing, while the 1pctCO₂ data set will be used for testing only.

2.3. Prediction: Conditional Distribution of Subgrid Scale Forcing

Our goal is to learn a parameterization, denoted by $\hat{\mathbf{S}}$, of the diagnosed *true* subgrid momentum forcing, \mathbf{S} (Equation 1), using deep learning. We propose a NN that uses maps of coarse surface velocities, \mathbf{u} , at a given time as inputs, and estimates the subgrid momentum forcing components at that same time as outputs. Here *estimates* is to be understood in a broad sense: it could be a single-value prediction or a probability distribution as discussed below.

Specifically, in this work we present a stochastic parameterization of the subgrid momentum forcing. To do so, we assume that at each grid box, the distribution of the forcing is Gaussian, conditionally on the coarse surface velocities (we do not assume that the marginal distribution of the forcing is Gaussian). We also assume that the forcing at individual grid boxes and times are conditionally independent given the coarse surface velocities.

The rationale behind stochastic approaches to the modeling of subgrid-scale forcing has been widely discussed (Berloff, 2005; Buizza et al., 1999; Crommelin & Vanden-Eijnden, 2008; Duan & Nadiga, 2007; Frederiksen & Davies, 1997; Majda et al., 1999; Palmer, 2001), in particular; they can partly account for the uncertainty in the representation for the subgrid forcing (Berner et al., 2017; Brankart, 2013; Juricke et al., 2017; Stanley et al., 2020; Williams et al., 2016; Zanna et al., 2018); and, they have proven potential in stabilizing numerical simulations (Berner et al., 2017; Palmer, 2012; Zanna et al., 2017).

One of the main sources of uncertainty in the predicted subgrid forcing comes from the fact that we only use the resolved coarse velocities to make a prediction. Given a surface velocity field over a subdomain of the oceans at a given time, we do not necessarily expect the subgrid momentum forcing to be given by a deterministic mapping. To illustrate this statement, consider Equation 1; for the problem at hand, \mathbf{u}^\uparrow is unknown, and while $\mathbf{u}^\uparrow \mapsto \mathbf{u}$ is well-defined as a mapping, it is not invertible. Thus the parameterization problem can be viewed as an inverse problem, for which probabilistic representations are a common approach (Bishop, 1991). If \mathbf{u}^\uparrow is seen as a random variable, we may want to represent the probability distribution $P(\mathbf{u}^\uparrow | \mathbf{u})$, and the same applies to the forcing. Hence we may want our NN's output to determine a parametric probability distribution rather than a single number.

Besides, a stochastic parameterization of the subgrid forcing can also account for the fact that what we call the true subgrid forcing, Equation 1, depends on our choice of filter which may not adequately represent the “missing forcing” from any given numerical model at coarse-resolution. One could train our NN with subgrid forcing generated from a variety of methods, to *partially* account for the fact that the exact subgrid forcing is not known.

The output Gaussian distribution at each location can be interpreted as an estimate of the conditional distribution of the subgrid momentum forcing given the local velocity field. Its mean represents the expectation of that conditional probability distribution, while its standard deviation represents the uncertainty around the mean. Such representation will allow deriving confidence intervals of the predicted subgrid momentum forcing (see Section 3.2). It also forms the basis of our stochastic parameterization, see Section 4 about implementation. In the next section, we will show how to learn the mean and standard deviation of the subgrid forcing from data.

2.4. Probabilistic Loss Function: From MSE to Gaussian Log-Likelihood

To train our NN, we aim to find a local minimum to a loss function $L(\mathbf{S}, \hat{\mathbf{S}}(\boldsymbol{\theta}))$ —summed over all the samples of the training data set—that represents the *mismatch* between our prediction $\hat{\mathbf{S}}$ and the true value \mathbf{S} given the current state of the parameters of the NN, represented by the vector of parameters $\boldsymbol{\theta}$. Here, \mathbf{S} , the target tensors of the NN corresponding to a single sample at a given time, has dimensions (n_C, n_x, n_y) , where $n_C = 2$ for the zonal and meridional component of the velocity field, and (n_x, n_y) is the size of the domain considered, i.e., the number of grid boxes in the zonal and meridional direction, respectively. We have ignored the number of mini-batches here for simplicity, which will be discussed in Section 2.6.

The most common loss function used in regression is the Mean Square Error (MSE), which in our case would take the form of, for a single sample,

$$L_{\text{MSE}}(\mathbf{S}, \hat{\mathbf{S}}(\boldsymbol{\theta})) = \sum_{k=1}^{n_C} \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} (\hat{\mathbf{S}}_{k,i,j} - \mathbf{S}_{k,i,j})^2, \quad (2)$$

where k denotes the index of the component of the subgrid momentum forcing (here, 1 corresponds to the zonal component and 2 to the meridional component). Despite its widespread use within the Deep Learning community for regression, the MSE loss function is not always appropriate. To justify the above claim briefly, it is common to interpret the MSE loss function from a probabilistic perspective. For simplicity, we limit the discussion to univariate random variables, but this can be easily extended to multivariate variables. Let ψ, ξ be random variables; assume that ψ is observed, and ξ is such that its conditional probability density function given ψ is a Gaussian distribution with mean μ and constant standard deviation σ ,

$$p(\xi | \psi; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(\xi - \mu(\psi))^2}{2\sigma^2}\right\}. \quad (3)$$

If we assume that $\mu(\psi)$ can be modeled by a parametric function f (in our case the NN) with parameter $\boldsymbol{\theta}$, the log-likelihood of the parameters $\boldsymbol{\theta}, \sigma$ for an independent and identically distributed (i.i.d.) sample $\{\psi_i, \xi_i\}_{i=1, \dots, n}$ will be given by,

$$l(\boldsymbol{\theta}, \sigma) = \sum_{i=1}^n \left\{ -\frac{1}{2} \log 2\pi\sigma^2 - \frac{(\xi_i - f(\psi_i, \boldsymbol{\theta}))^2}{2\sigma^2} \right\}. \quad (4)$$

Maximizing this log-likelihood (Maximizing the log-likelihood results in estimating the parameters of a probability distribution, so that under the assumed statistical model f the observed data ψ is most probable) over θ, σ can be achieved in a separable way (Davison, 2003): we first maximize over θ , which corresponds to training the NN using the MSE loss, and we then estimate σ by simply computing the standard deviation of the residuals $\{f(\psi_i, \theta) - \xi_i\}_{i=1, \dots, n}$. Hence, from a probabilistic point of view, by minimizing the MSE loss function, we are assuming a *constant* standard deviation (i.e., that does not depend on the velocity field).

In this study, we propose to relax this common assumption based on the literature and our understanding of the data. We replace the MSE loss function by a full negative Gaussian log-likelihood. Referring back to our univariate example, this would lead to replacing Equation 4 by,

$$l(\theta) = \sum_{i=1}^n \left\{ -\frac{1}{2} \log 2\pi f_2(\psi_i, \theta)^2 - \frac{(\xi_i - f_1(\psi_i, \theta))^2}{2f_2(\psi_i, \theta)} \right\}, \quad (5)$$

where our function f —which would correspond to our NN—now has two components, one for the mean, f_1 , of the Gaussian distribution, and the other one for the standard deviation, f_2 . In particular, the term corresponding to the standard deviation of the Gaussian in Equation 5, $f_2(\psi_i, \theta)$, does depend on the input ψ_i . In order to apply this to the problem of subgrid momentum forcing, we build our NN to output the two moments of a Gaussian distribution, at each location and for both (zonal and meridional) components of the subgrid momentum forcing. The output tensor, $\hat{\mathbf{S}}$, now has dimension $(2 \times n_C, n_x, n_y)$: we have four output channels ($2 \times n_C$)—the first two correspond to the means of the two components of the subgrid momentum forcing, the last 2 correspond to the associated standard deviations. Our loss function therefore takes the form of (ignoring constant terms),

$$L_G(\mathbf{S}, \hat{\mathbf{S}}(\theta)) = \sum_{k=1}^{n_C} \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \left\{ \log \hat{\mathbf{S}}_{k+2,i,j} + \frac{(\hat{\mathbf{S}}_{k,i,j} - \mathbf{S}_{k,i,j})^2}{2\hat{\mathbf{S}}_{k+2,i,j}^2} \right\}, \quad (6)$$

For ease of reading, we introduce a more natural notation, where we denote $\mathbf{S}_{C,i,j}$ the true value of the forcing, $\hat{\mathbf{S}}_{C,i,j}^{(\text{mean})}$ the mean of the predicted Gaussian distribution, and $\hat{\mathbf{S}}_{C,i,j}^{(\text{std})}$ its standard deviation, for component $C = X$ (zonal), Y (meridional) at location i, j . With this notation, Equation 6 takes the form of,

$$L_G(\mathbf{S}, \hat{\mathbf{S}}(\theta)) = \sum_{C=X,Y} \sum_{i=1}^{n_x} \sum_{j=1}^{n_y} \left\{ \log \hat{\mathbf{S}}_{C,i,j}^{(\text{std})} + \frac{(\hat{\mathbf{S}}_{C,i,j}^{(\text{mean})} - \mathbf{S}_{C,i,j})^2}{2(\hat{\mathbf{S}}_{C,i,j}^{(\text{std})})^2} \right\}. \quad (7)$$

The NN will learn to jointly optimize the two moments of the predicted Gaussian distribution, as shown schematically in Figure 2. Note that we also jointly train on both zonal and meridional components of the forcing, rather than having separate NNs for each component, as in Zanna and Bolton (2020).

2.5. Neural Network Architecture

Our NN is a Fully Convolutional Neural Network (Long et al., 2015) with a sequence of eight convolutional layers. The ReLU activation function is used for hidden layers. Given that the NN is fully convolutional, it can adapt to varying sizes of the input subdomain. We remind the reader that the input consists of two channels, one per component of the velocity field, while the output consists of four channels, two for each of the two components of the subgrid momentum forcing, see Section 2.4. We do not use any padding in the implementation of our convolutional layers. Due to the lack of padding in our NN structure, some *pixels* near the edges are lost in the application of convolutional layers. This results in the outputs predicted by our NN having spatial extent $(n_x - p, n_y - p)$, where p is a nonnegative integer that depends on the size of the kernels used in the convolutional layers.

The mean of the subgrid momentum forcing predicted by the NN can take any real value, as such we do not use any activation function in the final layer for the first two channels. However, the output predicted for the standard deviations are required to be positive. To enforce this constraint, we use a softplus function, defined by,

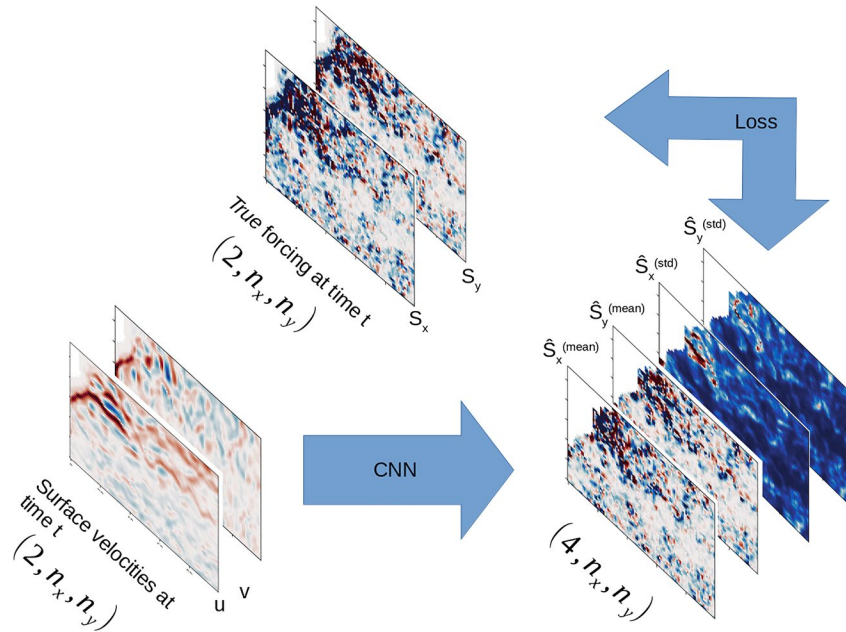


Figure 2. Our neural network (NN) outputs four maps: the two first maps are the maps of the means of the predicted forcing components, the last two maps are the standard deviation of the predicted forcing components.

$$\text{softplus}(x) = \ln(1 + \exp x) > 0, \quad (8)$$

as a final activation function for the two output channels associated with the standard deviations.

2.6. Training and Validation Procedure

We now describe our training procedure. The inputs are fed into the NN in the form of mini-batches (i.e., small batches of several samples stacked along an extra dimension), rather than individually, such that the dimensions of our input tensors are $(n_{\text{batch}}, n_C, n_x, n_y)$, where n_{batch} is the number of samples per mini-batch, set to $n_{\text{batch}} = 4$ in our experiments. A common practice in the methodology of NNs is to normalize inputs to be distributed within the interval $[-1, 1]$, to avoid vanishing and exploding gradients in the application of the back-propagation algorithm. Here, we multiply the surface velocities by a factor of 10, which approximately corresponds to normalizing by the standard deviation (≈ 0.1 m/s) of the surface velocity data from the training subdomains. This same transformation is applied in testing.

The targets used to train and evaluate our NN consist of the *true* subgrid momentum forcing \mathbf{S} for a given subdomain. We train our NN on data from the piControl simulation. We restrict the training data to a combination of four selected sub-domains of the oceans—shown as gray rectangles in Figure 1, see also Table 1—that correspond to various dynamical regimes: the Gulf Stream extension, the Equatorial Atlantic, just south of the Equatorial Pacific, and in the South Pacific gyre. Further improvements through more advanced selection and weighting of the training subdomains may improve performance. We select the first 80% of the data (approximately spanning 16 years) as training data, and use the final 15% (approximately spanning 3 years) for validation. We ignore 5% of the data (1 year) to avoid any correlation between the training data and the validation data, as it could cause validation metrics to become over-optimistic.

Table 1
Subdomains Used for Training and Validation

Subdomain	Latitude range	Longitude range
A	35.0°, 50.0°	−50.0°, −20.0°
B	−40.0°, −25.0°	−180.0°, −162.0°
C	−20.0°, −5.0°	−110.0°, −92.0°
D	0.0°, 15.0°	−48.0°, −30.0°

During the training phase samples are entirely shuffled across the time dimension, as well as across subdomains. This allows to jointly train on data from all selected subdomains simultaneously. However, this requires the input tensors obtained from all the subdomains to have the same

spatial sizes. We therefore crop the input tensors according to the smallest size across subdomains for both spatial dimensions, resulting in training samples of spatial extent $(n_x, n_y) = (38, 45)$. Given the nonMarkovian nature of turbulence (Kraichnan, 1959), further work incorporating temporal dependence and memory, rather than shuffling the data, might improve the skill of the learned parameterization (Porta Mana & Zanna, 2014).

We compute the average loss—defined in Section 2.4—over the samples of a mini-batch, across the two components of the forcing, and across both spatial dimensions. The average loss is then back-propagated to obtain the derivatives of the loss function with respect to the NN's parameters. The NN's parameters are then updated using the ADAM algorithm (Kingma & Ba, 2015). ADAM has become one of the go-to optimization algorithms in the Deep Learning community, which is in part due to its robustness to the choice of the learning rate and its quick convergence. After the NN's parameters have been updated, we repeat the same process with a new mini-batch, and so on, until all the training data has been used, which corresponds to one epoch of training. At this point, we compute the average loss over the validation data, which was not used for optimization. We track this validation loss over the whole set of training epochs and repeat this process. We implement early stopping so that training stops once the validation loss has not improved for four consecutive epochs of training. More details about our final choice of hyperparameters, such as the learning rate, hand-picked through a validation procedure, can be found in Appendix A.

3. Offline Tests on a Global Scale

We test our probabilistic deep learning parameterization on a global scale and demonstrate its generalization properties offline using test metrics introduced in Section 3.2. In Section 3.3, we first carry out a test on piControl in order to assess the ability of our NN model to generalize to regions and dynamical regimes not seen during the training phase. We then carry out a test on 1ptCO₂ in Section 3.4, where the CO₂ levels in the atmosphere reach double those of the piControl simulation. Our results show that our stochastic deep learning parameterization performs well in this new climate, without requiring further training of our NN; this is crucial if such parameterizations are to be used in models for climate projections (O'Gorman & Dwyer, 2018; Rasp et al., 2018).

3.1. Global Reconstruction for Offline Testing

We directly apply our trained NN to the global coarsened velocities for offline testing. While the training subdomains are only 38×45 points, we use the global extent of the oceans for testing which results in an input tensor of dimension $(N_x, N_y) = (900, 625)$. This is possible due to the convolutional nature of the NN. When applying the NN to global data, we extend the input velocities cyclically along the zonal dimension (since the output values of the tensor for $i = N_x$ requires information from $i = 1$), thus ensuring the output covers all longitudes. Obviously, this is not possible along the meridional dimension, thus resulting in the loss of $p = 10$ grid boxes (see Section 2.5) along the meridional dimension at high latitudes in each hemispheres.

Velocity snapshots are assembled to form small mini-batches with size 4 (equivalent to 4 days); the size is determined by the available GPU memory. Nonocean points of the input grid are stored as NaNs. In our tests, we therefore ignore locations whose receptive field intersect with a continent and show them in black in the maps shown thereafter (note, the receptive field of a neuron within the NN's output is the set of input neurons that impact its value).

3.2. Metrics and Statistics for Offline Performance

To quantify the offline accuracy of our NN's predictions of the subgrid momentum forcing, we define several metrics. We note $T = 7300$ the total number of days over which these metrics are computed.

We first define our notation for the standard MSE and correlation. To make explicit the dimension along which the data is reduced to compute these two metrics, we write $\text{MSE}_{C,i,j,-}$ for the time-mean MSE of the $C \in \{X, Y\}$ component of the forcing, where the reduction is carried out along the time axis, i.e.

$$\text{MSE}_{C,i,j,-} = \frac{1}{T} \sum_{t=1}^T \left(\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})} - \mathbf{S}_{C,i,j,t} \right)^2, \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y. \quad (9)$$

The combined MSE, that encompasses both components X and Y , can be shown on a map, and is defined as

$$\text{MSE}_{i,j,-} = \frac{1}{T} \sum_{t=1}^T \left\{ \left(\hat{\mathbf{S}}_{X,i,j,t}^{(\text{mean})} - \mathbf{S}_{X,i,j,t} \right)^2 + \left(\hat{\mathbf{S}}_{Y,i,j,t}^{(\text{mean})} - \mathbf{S}_{Y,i,j,t} \right)^2 \right\}. \quad (10)$$

We also define a scalar MSE according to,

$$\text{MSE} = \frac{1}{N_x N_y T} \sum_{t=1}^T \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \left\{ \left(\hat{\mathbf{S}}_{X,i,j,t}^{(\text{mean})} - \mathbf{S}_{X,i,j,t} \right)^2 + \left(\hat{\mathbf{S}}_{Y,i,j,t}^{(\text{mean})} - \mathbf{S}_{Y,i,j,t} \right)^2 \right\}. \quad (11)$$

In addition to the standard MSE, we define an R^2 coefficient which is normalized by the value of the true subgrid forcing such that

$$R_{C,i,j,-}^2 = 1 - \frac{\sum_{t=1}^T \left(\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})} - \mathbf{S}_{C,i,j,t} \right)^2}{\sum_{t=1}^T \mathbf{S}_{C,i,j,t}^2}, \quad C = X, Y, \quad (12)$$

$$R_{i,j,-}^2 = 1 - \frac{\sum_{t=1}^T \left\{ \left(\hat{\mathbf{S}}_{X,i,j,t}^{(\text{mean})} - \mathbf{S}_{X,i,j,t} \right)^2 + \left(\hat{\mathbf{S}}_{Y,i,j,t}^{(\text{mean})} - \mathbf{S}_{Y,i,j,t} \right)^2 \right\}}{\sum_{t=1}^T \left(\mathbf{S}_{X,i,j,t}^2 + \mathbf{S}_{Y,i,j,t}^2 \right)} \quad (13)$$

and a scalar version R^2 , according to,

$$R^2 = 1 - \frac{\sum_{t=1}^T \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{C=X,Y} \left(\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})} - \mathbf{S}_{C,i,j,t} \right)^2}{\sum_{t=1}^T \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{C=X,Y} \mathbf{S}_{C,i,j,t}^2} \quad (14)$$

We note that $R^2 \leq 1$ and if $\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})}$ is zero always, R^2 is 0. The advantage of this quantity is that it is easier to interpret when shown on a map—values close to 1 indicate that our predictions account for a large part of the average amplitude of the subgrid momentum forcing, while values close to 0 would indicate the opposite.

In order to verify that our model is not simply predicting the seasonal climatology of the subgrid momentum forcing, we define modified versions of these quantities, e.g.,

$$R_{C,i,j,-}^{2,\text{clim}} = 1 - \frac{\sum_{t=1}^T \left(\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})} - \mathbf{S}_{C,i,j,t} \right)^2}{\sum_{t=1}^T \left(\mathbf{S}_{C,i,j,t}^{\text{clim}} - \mathbf{S}_{C,i,j,t} \right)^2}, \quad C = X, Y, \quad (15)$$

where $\mathbf{S}_{C,i,j,t}^{\text{clim}}$ is the climatological C -component subgrid momentum forcing at location i, j and time t . This metric allows us to assess what percentage of the signal's variance we account for, after removing the inherent variability due to the seasonal cycle.

Another quantity of interest given our probabilistic representation of the subgrid momentum forcing parameterization is that of standardized residuals, given by,

$$\mathbf{e}_{C,i,j,t} = \frac{\hat{\mathbf{S}}_{C,i,j,t}^{(\text{mean})} - \mathbf{S}_{C,i,j,t}}{\hat{\mathbf{S}}_{C,i,j,t}^{(\text{std})}}, \quad C = X, Y. \quad (16)$$

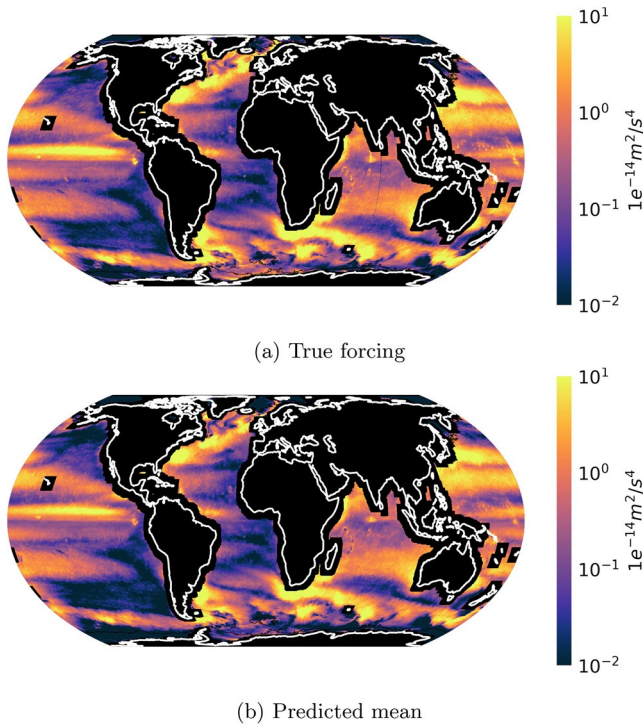


Figure 3. Time-mean variance of the norm of subgrid momentum forcing in piControl: (a) True forcing $\| S \|$; (b) predicted mean, $\| \hat{S}^{mean} \|$ in offline testing.

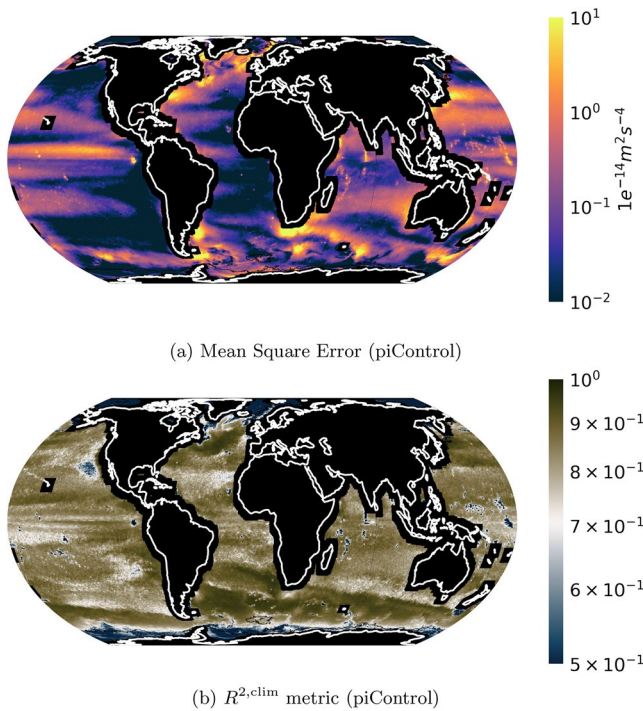


Figure 4. Time-mean (a) Mean Square Error (MSE) (10) and (b) $R^{2,clim}$, the climatological version of (13), for piControl.

Under our idealized assumption, these normalized residuals are expected to follow a standard normal distribution.

We will also use confidence intervals based on the standard deviation, to quantify the uncertainty in the predicted subgrid forcing and evaluate its performance. Under the Gaussian assumption, a 95% confidence interval corresponds to,

$$\hat{S}_{C,i,j,t}^{(mean)} \pm 1.96 \hat{S}_{C,i,j,t}^{(std)}, \quad C = X, Y. \quad (17)$$

3.3. Generalization & Dynamical Regions—Test on piControl


We carry out an offline test of our NN on global scale data from piControl. There are large variations in subgrid eddy momentum in the piControl (Figure 3a) across the oceans, with the largest amplitude occurring in eddy rich regions such as the Gulf Stream, Kuroshio, Southern Ocean and equatorial regions. There is a strong coherence between the pattern of the variance of the mean of the true subgrid forcing (Figure 3a) and that of the predicted forcing (Figure 3b). This coherence holds for the zonal and meridional component of the forcing, as shown for example in the correlation map between the true zonal forcing and the mean component of the predicted zonal forcing (Figure D1).

The time-mean MSE over both components of the forcing (Equation 11) can vary by several orders of magnitude from one region to another (Figure 4a). However, these changes are largely due to the inherent spatial variability of the subgrid forcing, evident by comparing its spatial pattern (Figure 3a) with the spatial pattern of the MSE (Figure 4a). Therefore, the $R_{i,j,-}^{2,clim}$ coefficient (Equation 15) is more informative of the NN's performance (Figure 4b).


In most regions of the oceans, our NN is able to account for more than 70% of the signal's variance, with performance nearing 90% in regions where the variance of the eddy momentum forcing is the highest, for instance in the Gulf Stream region and Southern Ocean (see the appendix for maps of the $R^{2,clim}$ computed for each component of the forcing—Figure D2—showing similar skill). These metrics indicate that the NN generalizes well to most regions, despite being trained on only four small subdomains of the oceans. However, our NN performs poorly in sea-ice covered regions, which is not surprising as the dynamics of these regions were not included in the training and widely differ from open ocean turbulence. Considering turbulence at the ocean-ice boundary will require numerical simulations that can adequately represent such processes.

The near-global ($60^\circ S, 60^\circ N$) scalar R^2 value obtained is 0.869, while for $R^{2,clim}$ we obtain 0.855; the skill demonstrates the high performance of our NN and further confirms that the NN does not merely predict large variations due to seasonal climatology. The global R^2 is higher than the average of R^2 values over the map due to the higher R^2 values in regions where the variance of the forcing is large (note that Equation 14 is not the spatial average of Equation 12).

To demonstrate some advantages of predicting the two moments of a Gaussian distribution, we focus on time series at two different locations. We compare the time series of the true and predicted zonal forcing at



(a) turbulent



(b) quiescent

Figure 5. Time series of the zonal component of the subgrid momentum forcing at (a) $30^{\circ}N, 60^{\circ}W$, a location dominated by turbulent behavior and (b) $20^{\circ}S, 104^{\circ}W$, a more quiescent location for 300 days: true forcing (solid blue), mean of the predicted forcing (orange), and 95% confidence interval (green).

$30^{\circ}N, 60^{\circ}W$, which is located within the turbulent Gulf Stream region, and at $20^{\circ}S, 104^{\circ}W$, which corresponds to a more quiescent region with less mesoscale eddy activity (Figure 5). The true zonal forcing $S_{X,i,j,t}$ is shown along with the mean prediction $\hat{S}_{X,i,j,t}^{(\text{mean})}$ and the 95% confidence interval obtained from the predicted standard deviation $\hat{S}_{X,i,j,t}^{(\text{std})}$. The forcing is generally well approximated by the predicted mean forcing, except when extremes occur. However, the true forcing is, most of the time, within the 95% confidence

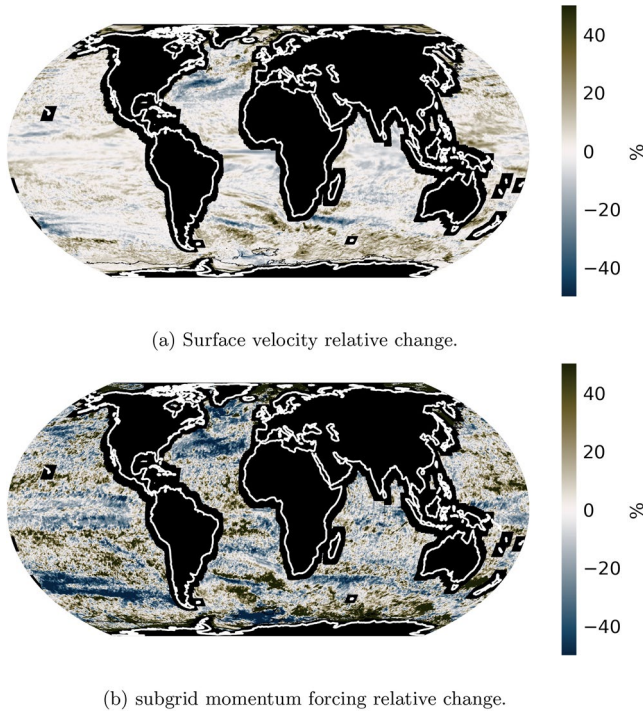


Figure 6. Relative difference between piControl and 1ptCO₂ in the standard deviation of the (a) surface velocity norm and (b) subgrid forcing norm. Positive (negative) values indicate that the variance has increased (decreasing) in the 1ptCO₂ compared to the piControl.

interval. The predicted standard deviation $\hat{S}_{X,i,j,t}^{(\text{std})}$ varies greatly across the considered time window—indicating that the uncertainty of the forcing is not constant. Our NN performs best in turbulent regions. This is in agreement with R^2 maps where higher values are observed in regions where the forcing is larger, and also with results from idealized ocean models (Bolton & Zanna, 2019; Zanna & Bolton, 2020). Finally, to investigate regions with a low R^2 score, we analyze the time series of the true and predicted meridional forcing at 29°N, 129°W (Figure D3), which corresponds to a location near the West Coast of the United States where the R^2 score is 0.532. The time series indicates that the low R^2 occurs due to a few extreme events that are not well predicted.

To further analyze our predicted forcing from the piControl data set, we study the global distribution of a stochastic simulation of subgrid momentum forcing generated using,

$$\tilde{S}_{C,i,j} = \hat{S}_{C,i,j}^{(\text{mean})} + \epsilon_{C,i,j} \times \hat{S}_{C,i,j}^{(\text{std})}, \quad C = X, Y, \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y,$$

where the $\epsilon_{C,i,j}$ are i.i.d. standard normal and the inputs to the NN are the coarsened surface velocities from piControl. The histograms of the global distribution of each component of the subgrid forcing for the true and simulated forcing show that the two distributions are very similar (Figure D4). However, the distribution of the true forcing has longer tails than that of the simulated forcing. This is partly due to our assumption that the distribution of the forcing, conditioned on the coarse surface velocity field, is Gaussian. We test this hypothesis by investigating the distribution of normalized residuals, defined by Equation 16. Figure D5a consists of the sample distribution of normalized residuals (blue), after subsampling one point out of ten along the time axis, and one point out

of five along the spatial axes, shown together with the probability density function of the standard normal distribution (red). We also present a quantile-quantile (QQ)-plot of the sampled normalized residuals in Figure D5b, using the same subsampling procedure as in Figure D5a. The normalized residuals have much heavier tails than those of a standard normal. Hence, we could improve our model by using another distribution with heavier tails, or a multimodal distribution (Bishop, 1991). This approach will likely improve the offline prediction of extreme events which we have shown is problematic in our NN.

3.4. Generalization & Climate Change—Test on 1ptCO₂

One key challenge for deep learning parameterizations in ocean and climate modeling is for them to be able to generalize to a new climate (Martínez-Moreno et al., 2021; O’Gorman & Dwyer, 2018). So far, we have only used data from the piControl simulations, both for training (Section 2.6) and testing (Section 3.3). Here, we test the trained NN from Section 2.6, without further tuning, using simulated data from 1ptCO₂. The surface velocities, associated kinetic energy, and subgrid momentum forcing, are influenced by the CO₂ forcing. The time-mean standard deviation of the surface velocity between piControl and 1ptCO₂ (Figure 6a) changes by up to 40% in some parts of the oceans. The majority of the changes are occurring in regions dominated by high kinetic energy in piControl such as the Gulf Stream region and its extension, the Kuroshio extension, or the Southern Ocean. Additionally, we identify large changes in the Indian Ocean and in the Arctic, ice-melt is likely related to changes in the latter. Similar changes in the subgrid momentum forcing between piControl and 1ptCO₂ are occurring as well (Figure 6b). The surface velocities used as inputs to the NN and the target subgrid forcing to be predicted are therefore significantly different from those of piControl.

In order to compare performance of our NN on piControl and 1ptCO₂ we use the same metrics as in Section 3.3. The MSE and $R^{2,\text{clim}}$ metrics computed over the 20 years of daily simulation data from 1ptCO₂ are

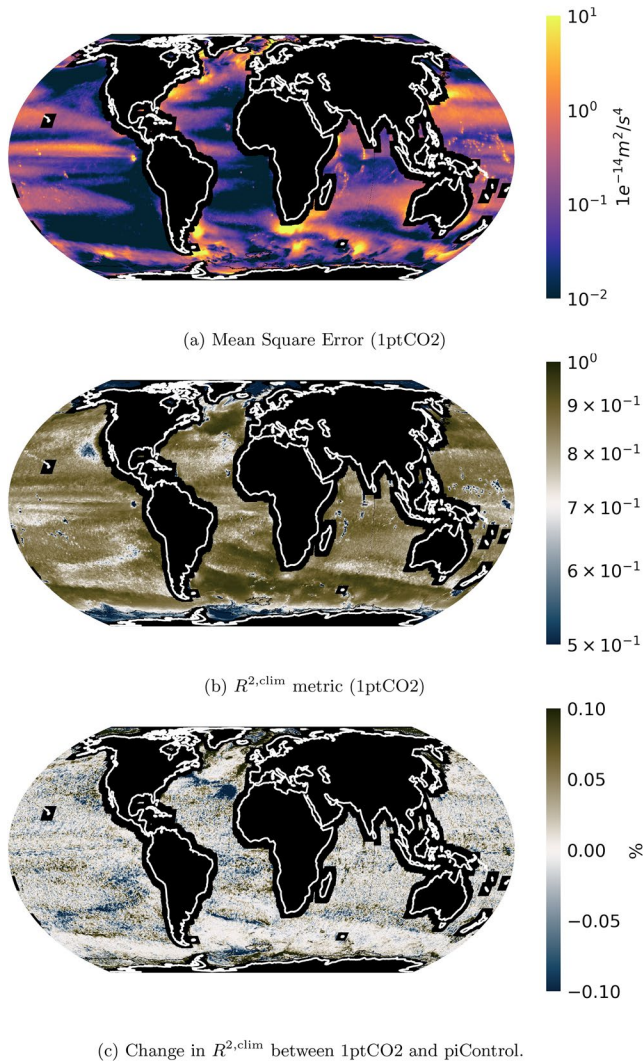


Figure 7. Performance of the trained neural network (NN) on 1ptCO₂ simulation: (1) Mean Square Error (MSE) (1ptCO₂); (b) $R^{2,\text{clim}}$ metric (1ptCO₂); (c) Change in $R^{2,\text{clim}}$ between 1ptCO₂ and piControl.

complex CM2.6 data; for further details on the model, consult these these references, or our code. The stochastic parameterization is implemented in a 40 km horizontal resolution run, and we compare the runs to a high-resolution model run at 10 km horizontal resolution, hence mimicking the change in resolution between CM2.6 simulation data and the coarse-grained data we generated to diagnose the momentum forcing. Note that we tested offline the learned CM2.6 parameterization with data from the shallow water model and shown that the learned NN model transfers well to a different data set and model before implementation, see details in Appendix C.

Unlike CM2.6 which was on a B-grid, the shallow water model is discretized on an Arakawa C-grid. Therefore, at each time step of the integration, we first interpolate the two velocity components on tracer points and then pass them through our NN. This produces, for each grid box and for each component of the forcing, a mean $\hat{S}_{C,i,j}^{(\text{mean})}$ and a standard deviation $\hat{S}_{C,i,j}^{(\text{std})}$.

The stochastic subgrid momentum forcing \tilde{S} implemented in the shallow water model is then generated (see schematics in Figure 10) according to,

shown in Figure 7. Our NN performs as well for this new climate as it did for the climate it had been trained on (e.g., compare Figure 7 with Figure 4). The time-mean $R^{2,\text{clim}}$ obtained on piControl and 1ptCO₂ show little difference (Figure 7c), except in the North-East Atlantic and in certain polar regions which were partially ice-covered in piControl, where there is a slight decrease in performance (at most 0.1) as measured by the time-mean $R^{2,\text{clim}}$. We compute scalar metrics of our NN's performance over the 1ptCO₂ simulation data, again limited to 60°S,60°N, and obtain 0.871 for R^2 and 0.858 for $R^{2,\text{clim}}$, i.e. very similar to the values obtained for piControl. The NN for subgrid momentum forcing trained on piControl data generalizes well to an unseen warmer climate as simulated by a coupled high-resolution climate model.

4. Online Implementation for an Idealized Model

Offline performance tests have not been good predictors for online performance, as shown for example in Zanna and Bolton (2020), at least not using current metrics. The coupling between the machine learning (ML) parameterization and the prognostic coarse-resolution ocean model must satisfy the same numerical stability criteria and conservation properties as any physics-derived parameterization. Therefore, good offline performance is a necessary but not sufficient condition to the success of any ML parameterization.

Zanna and Bolton (2020) implemented a convolutional NN parameterization which, while physically constrained, led to too vigorous an inverse energy cascade and energy backscatter. While the model was not numerically unstable, the behavior of the model was pushed into a different dynamical regime in which the eddy mean-flow interactions dominated over the external wind forcing. To ensure a reasonable dynamical behavior, the authors tuned down the parameterization by applying a spatially and temporally uniform multiplicative factor to reduce the magnitude of the forcing in an ad-hoc way.

The use of a stochastic parameterization has the potential to damp the eddy (and destabilizing) feedbacks seen in Zanna and Bolton (2020). Here, we use the same idealized barotropic shallow water model as in Zanna and Bolton (2020) and developed by Klöwer et al. (2018) to implement the stochastic deep learning parameterization learned from

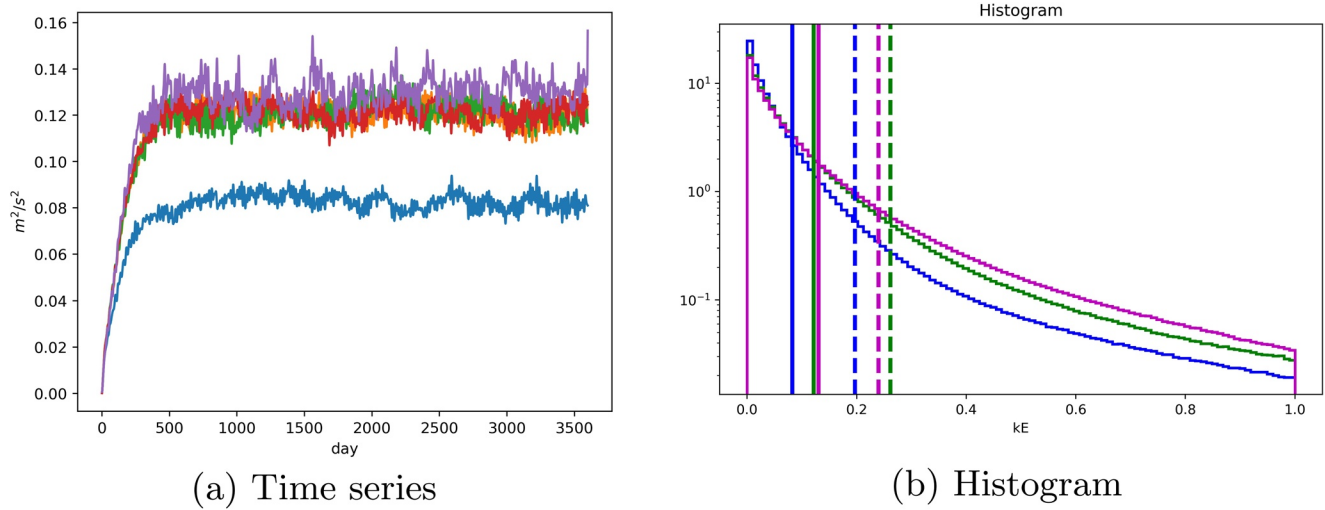


Figure 8. Kinetic energy [m^2 / s^2] (a) time series, and (b) histogram for the low resolution unparametrized simulation at 30 km (blue), low resolution parameterized ensemble member simulations (green, orange, red), and filtered + coarse-grained high-resolution simulation (purple). In panel (b) the solid vertical lines indicate the mean and the dashed vertical lines the standard deviation of the simulated kinetic energy; only one ensemble member is shown, but the other ensemble members produce similar statistics.

$$\tilde{S}_{C,i,j} = \hat{S}_{C,i,j}^{(\text{mean})} + \epsilon_{C,i,j} \times \hat{S}_{C,i,j}^{(\text{std})}, \quad C = X, Y, \quad i = 1, \dots, m_x, \quad j = 1, \dots, m_y, \quad (18)$$

where the $\epsilon_{C,i,j}$ are sampled according to i.i.d. standard normal distributions, and $m_x = m_y = 96$ is the number of points of the low-resolution grid along the x and y axes respectively. The field \tilde{S} is then interpolated back to the u and v grid for the X and Y components, respectively, and used as the value of the subgrid momentum forcing in the shallow water model.

We ran the following simulations for 10 years each: a high resolution run (10 km), a coarse-resolution run (40 km) without parameterization, and three different ensemble members of the coarse-resolution run with the stochastic parameterization. The parameterized simulations are stable and produced a physically consistent state without any tuning or scaling factor. The kinetic energy of the flow is improved: both the mean and the standard deviation are very close to the high-resolution simulation (Figure 8). Similarly to Zanna and Bolton (2020), the variance of the velocity fields (not shown) and sea surface height (Figure 9) are vastly improved by the parameterization. However, changes in the mean velocity are rather small (not shown). We believe that the simplicity of the shallow water model used in the present study is at the core of the lack of substantial improvement in the mean flow and our parameterization will need to be tested in a more complex model. Unlike Zanna and Bolton (2020), no physical constraint was imposed when learning the NN parameterization in our study; yet, we do not observe any drift in the model. Despite using zero-padding




Figure 9. Standard deviation of sea surface height [m] for (left) the low resolution simulation; (middle) one ensemble member from the parameterized versions; (right) the high resolution simulation.




Figure 10. Procedure for generating the stochastic parameterization (Equation 18) implemented in the coarse resolution idealized model, based on the trained neural network (NN).

during the implementation, the solutions near the boundaries are not strongly impacted. Overall, the coarse resolution stochastic simulations are 25% slower than the unparameterized runs but more than 40 times faster than a high-resolution simulation at 10 km on the same CPU. However, this statement is to be taken with care as the high-resolution simulation was not optimized.

5. Discussion

Current parameterizations of ocean and atmosphere processes can lead to biases in climate models and remain a large source of uncertainty in climate projections. Therefore, harnessing state-of-the-art Deep Learning and statistical methods to improve parameterizations of subgrid processes has recently raised a lot of interest (Bolton & Zanna, 2019; Rasp et al., 2018; Yuval & O’Gorman, 2020; Zanna & Bolton, 2020). Here, we have demonstrated the potential of Deep Learning approaches for the problem of ocean momentum subgrid parameterizations using data generated by a realistic coupled climate model, as opposed to data from idealized ocean-only quasi-geostrophic or primitive equation simulations (Bolton & Zanna, 2019; Zanna & Bolton, 2020).

The use of data from realistic coupled climate models to train Deep Learning is nontrivial due to the size of the data, the use of the tripolar irregular spherical grid, and the coupling between the ocean and the atmosphere. Here, we establish a filtering and coarse-graining procedure to diagnose the subgrid momentum forcing in a global model and show that using only a limited number of subdomains, we can train a NN to skillfully predict the subgrid momentum forcing over the global oceans, and in a different climate with increased CO₂ levels. However, there are several remaining challenges.

We have shown that the offline skill of the predictions is lower in regions where sea-ice is present. Therefore, to improve parameterizations of ocean mesoscale eddies in these regions, it might be necessary to acquire data that can faithfully represent these interactions with sea-ice and likely with the atmosphere in polar regions. Another outstanding challenge is to diagnose and train for subgrid momentum forcing in grid boxes near continents. Our NN accounts for some nonlocal turbulent interaction (Frederiksen, 1999; Kraichnan, 1959), however the influence of large-scale far-field on the subgrid momentum forcing could potentially be improved by increasing the size of the subdomain used for training. The current size of the subdomain was chosen for convenience, by considering the best skill with the smallest amount of data. In addition, learning the standard deviation (in addition to the mean) of the subgrid parameterization takes into account the missing nonlocal interactions to some extent (and any other missing feedbacks in the learning procedure). While larger subdomains may perform better offline, training might require larger amounts of data to avoid overfitting and may not translate into a better online implementation. Yet, other structures such as the U-net (Ronneberger et al., 2015) allow for interactions with larger scales. Our experiments with a U-net structure have not yet led to any improvements in the subgrid momentum forcing parameterization; though further investigation might be required.

To improve the generalization of data-driven parameterizations, simulations spanning the phase space of turbulence (different forcing, viscosity, etc.) could be used to train Machine Learning models and potentially offer a window into the theoretical underpinning of subgrid models (Kitsios et al., 2016; Porta Mana & Zanna, 2014). These Machine Learning models could be implemented more generally in climate model simulations. In addition, 1 / 10^o-resolution ocean models still underestimate the true mesoscale transient kinetic energy. While the use of Large Eddy Simulations (LES) and Direct Numerical Simulations (DNS)

could alleviate these issues, these models cannot be run at global scales or include all components of the climate system, all of which influence the eddy momentum forcing. Satellite observations of sea surface height could be used to fine-tune our trained NN via transfer learning. Yet, using satellite data comes with its own challenges due to sparse spatio-temporal sampling and error due to the interpolation of the gridded datasets. These challenges could be overcome in the near future as higher resolution satellite data becomes available and improved gridding procedures are being developed.

We only used information from the surface layer to train for subgrid momentum forcing (since subsurface velocity data is not publicly available). Training a NN model for each layer of a climate model would likely be limited by computational power and storage (of the NN parameters) and of limited scientific use. Instead, a single NN model of similar complexity to the one presented in this manuscript can be trained using data from all layers of a coupled model (Bolton & Zanna, 2019; Zanna & Bolton, 2020). While the current NN might generalize well to other layers, we might expect degradation in skill in regions that interact strongly with bathymetry, for example.

The NN presented here was trained to predict the parameters (mean and standard deviation) of a Gaussian probability distribution at each grid box, therefore providing a probabilistic prediction of the subgrid forcing with Deep Learning. This probabilistic approach attempts to account for both the uncertainty in the mapping between the coarse velocity field and the subgrid forcing, uncertainty in any missing interaction not included in the learning, and uncertainty in the data itself. Stochasticity has been shown to improve model bias and produce more reliable ensemble predictions (Berner et al., 2017). Besides, while most current Deep Learning approaches to the parameterization of subgrid processes have been deterministic, a stochastic approach could be key when it comes to online implementations. Several Deep Learning implementations of parameterizations trained offline have resulted in poor stability properties or unrealistic flows in online simulations. Stochasticity could provide one approach to potentially solve this issue as shown in previous work (Palmer, 2012; Zanna et al., 2017). Using an idealized shallow water model, we showed that implementing our stochastic parameterization results in stable simulations and produces a realistic flow without any tuning. However, while the stochastic parameterization vastly improved some metrics (mean and variance of the kinetic energy), the impact on other metrics were only modest (e.g., zonal velocities).

The probabilistic approach presented here to learning the subgrid forcing remains simple and could be applied to parameterizing other processes. Yet it could benefit from more advanced probabilistic modeling. While we limited ourselves to conditionally i.i.d. Gaussian distributions, our analysis of residuals shows that representing higher moments could lead to a better representation of the distribution of subgrid forcing. In addition, we do not account for the uncertainty associated with the parameters of the NN (Jospin et al., 2020). While Bayesian NNs remain computationally more expensive, recent progress on that front could be an interesting avenue of investigation, and may provide additional assurance compared to single outputs.

Finally, combining closed-form parameterizations with stochastic Deep Learning approaches could be another fruitful avenue. For instance, it would be possible to predict the mean forcing via a closed-form equation, such as done by Zanna and Bolton (2020) using equation-discovery methods, while representing higher-order moments via a probabilistic Deep Learning approach similar to that proposed in this manuscript. This approach could improve our understanding of missing processes and their representation in climate models. While the effects of Deep Learning subgrid parameterizations on climate projections remain to be ascertained, the benefits of Deep Learning could be greater if they are used to understand processes from a probabilistic perspective.

Appendix A: Neural Network Training Hyperparameters

The values of the hyperparameters used in our training procedure are provided in Table A1. The learning rate is decreased through the training procedure, hence we provide its initial value (for epoch 0) and epochs at which the rate is decreased. The provided number of epochs corresponds to the maximum number of training epochs. In practice, training usually stops earlier due to our implementation of early stopping.

Table A1
Hyperparameter Values for Training

Hyperparameter		Value
Number of epochs		100
Learning rate	Epoch 0	$5e^{-4}$
	Epoch 10	$5e^{-5}$
	Epoch 20	$5e^{-6}$
Batch size		4
Filter sizes	Layers 1 – 2	5
	Layers 3 – 8	3
Padding		No

Appendix B: Generation of Low-Resolution Data and Estimates of the Missing Mesoscale Subgrid Forcing

We provide pseudo-code for the generation of the low-resolution data based on the CM2.6 high-resolution data set. This algorithm makes use of two functions whose pseudo-code is also provided, *filter*, which applies a Gaussian filter to the high-resolution data weighted by the cell areas, and *advectons*, which computes the advection term of a discrete velocity field.

```

Data:  $\sigma, u_{i,j}^\uparrow, v_{i,j}^\uparrow, \quad i = 1, \dots, \sigma n_x, \quad j = 1, \dots, \sigma n_y$ 
Result:  $u_{i,j}, v_{i,j}, S_{X,i,j}, S_{Y,i,j}, \quad i = 1, \dots, n_x, \quad j = 1, \dots, n_y$ 
/* Pre-processing: replace NaNs (i.e. land points) with zeros */
if  $u_{i,j}^\uparrow == \text{NaN}$  then  $u_{i,j}^\uparrow \leftarrow 0$ ;
if  $v_{i,j}^\uparrow == \text{NaN}$  then  $v_{i,j}^\uparrow \leftarrow 0$ ;
/* Compute the filtered high-rez surface velocities, for each
   component */
 $\bar{u}^\uparrow \leftarrow \text{filter}(u^\uparrow);$  /* filter function defined below */
 $\bar{v}^\uparrow \leftarrow \text{filter}(v^\uparrow);$ 
/* Compute the advection term of the filtered surface velocities */
 $\psi_X^\uparrow, \psi_Y^\uparrow \leftarrow \text{advection}(\bar{u}^\uparrow, \bar{v}^\uparrow);$  /* advection function def. below */
/* Compute the filtered advection term from high-rez surface
   velocities */
 $\phi_X^\uparrow, \phi_Y^\uparrow \leftarrow \text{advection}(u^\uparrow, v^\uparrow);$ 
 $\phi_X^\uparrow \leftarrow \text{filter}(\phi_X^\uparrow);$ 
 $\phi_Y^\uparrow \leftarrow \text{filter}(\phi_Y^\uparrow);$ 
/* Compute the components of the forcing term */
 $S_X^\uparrow \leftarrow \psi_X^\uparrow - \phi_X^\uparrow;$ 
 $S_Y^\uparrow \leftarrow \psi_Y^\uparrow - \phi_Y^\uparrow;$ 
/* Apply coarse-graining by factor  $\sigma$  */
 $u \leftarrow \langle \bar{u}^\uparrow \rangle_\sigma;$ 
 $v \leftarrow \langle \bar{v}^\uparrow \rangle_\sigma;$ 
 $S_X \leftarrow \langle S_X^\uparrow \rangle_\sigma;$ 
 $S_Y \leftarrow \langle S_Y^\uparrow \rangle_\sigma;$ 

```

Function filter:

Input: $u_{i,j}^\uparrow, dx_{i,j}^\uparrow, dy_{i,j}^\uparrow, \sigma$

Output: $\bar{u}_{i,j}^\uparrow$

$A_{i,j}^\uparrow = dx_{i,j}^\uparrow \times dy_{i,j}^\uparrow;$ /* Area of the U-cell */

$$\bar{u}_{i,j}^\uparrow = \frac{\sum_{i',j'=-2\sigma}^{2\sigma} u_{i',j'}^\uparrow * A_{i',j'}^\uparrow \exp\left\{-\frac{(i'-i)^2+(j'-j)^2}{2\left(\frac{\sigma}{2}\right)^2}\right\}}{\sum_{i',j'=-2\sigma}^{2\sigma} A_{i',j'}^\uparrow \exp\left\{-\frac{(i'-i)^2+(j'-j)^2}{2\left(\frac{\sigma}{2}\right)^2}\right\}};$$

return $\bar{u}_{i,j}^\uparrow$;

Function Advection:

Input: $u_{i,j}^\uparrow, v_{i,j}^\uparrow, dx_{i,j}^\uparrow, dy_{i,j}^\uparrow$

Output: $\phi_{X,i,j}^\uparrow, \phi_{Y,i,j}^\uparrow$

$$\partial_x^\uparrow u_{i,j} \leftarrow \frac{u_{i,j}^\uparrow - u_{i-1,j}^\uparrow}{dx_{i,j}^\uparrow};$$

$$\partial_y^\uparrow u_{i,j} = \frac{u_{i,j}^\uparrow - u_{i,j-1}^\uparrow}{dy_{i,j}^\uparrow};$$

$$\partial_x^\uparrow v_{i,j} = \frac{v_{i,j}^\uparrow - v_{i-1,j}^\uparrow}{dx_{i,j}^\uparrow};$$

$$\partial_y^\uparrow v_{i,j} = \frac{v_{i,j}^\uparrow - v_{i,j-1}^\uparrow}{dy_{i,j}^\uparrow};$$

/* Note here that the 4 quantities defined above need to be interpolated back on the U-grid before the two lines below */

$$\phi_{X,i,j}^\uparrow = u_{i,j}^\uparrow \partial_x^\uparrow u_{i,j} + v_{i,j}^\uparrow \partial_y^\uparrow u_{i,j};$$

$$\phi_{Y,i,j}^\uparrow = u_{i,j}^\uparrow \partial_x^\uparrow v_{i,j} + v_{i,j}^\uparrow \partial_y^\uparrow v_{i,j};$$

return $\phi_{X,i,j}^\uparrow, \phi_{Y,i,j}^\uparrow$;

Note: the high-resolution velocities are set to zero on land point when applying the Gaussian filter.

Appendix C: Offline Test of the Learned CM2.6 ML Parameterization With Shallow Water Data

To test if the learned ML parameterization with CM2.6 data generalizes to a different model, we conduct an offline test with data from the shallow water model described in Section 4. We use daily velocities from the high resolution run, a total of 3,600 velocity snapshots. We then apply the same procedure presented in Section 2, i.e. we filter and coarse-grain the data spatially to obtain coarse velocities on a low-resolution (40 km) grid. We also diagnose a *true* forcing, which we use as a reference for comparison with outputs from our NN. The time-average R^2 is high in most parts of the domain with values above 0.9, except in the more quiescent Eastern regions and near the boundaries with values of about 0.7–0.8 (Figure C1). The lower R^2 values in quiescent regions is not surprising (due to the low forcing in that region), and consistent with offline results carried out on CM2.6 data.

The time series of the predicted mean forcing, along with 95% confidence intervals evaluated from the predicted standard deviation for both components, at locations indicated by the white and green dots in Figure C1, show good agreement with the true forcing (Figures C2 and C3). Overall the NN—trained on CM2.6 simulation data—is able to generalize to the shallow water model data without any degradation in performance, indicating that the ML model is able to capture the appropriate ocean subgrid eddy forcing simplified model.




Figure C1. Time-average R^2 coefficient maps for S_x (left) and S_y (right) for offline tests on simulation data from the shallow water model.




Figure C2. Time series of true forcing (blue), of the mean predictions (orange), and of the 95% confidence interval (dashed green) at the location indicated by the white dot in Figure C1.




Figure C3. Time series of true forcing (blue), of the mean predictions (orange), and of the 95% confidence interval (dashed green) at the location indicated by the green dot in Figure C1.

Appendix D: Complementary Figures






Figure D1. Correlation between S_X and $\hat{S}_X^{(mean)}$. Note the lower values are 0.7.



(a) $R^{2,clim}$ metric (piControl) for the zonal component



(b) $R^{2,clim}$ metric (piControl) for the meridional component

Figure D2. Map of time-mean $R^{2,clim}$ metric in piControl for (a) the zonal component (b) the meridional component. The performances for both components are similar with lowest values around 0.6.




Figure D3. Time series of the true (solid blue) zonal component of the subgrid momentum forcing, mean zonal component of our neural network (NN) (orange), and 95% confidence interval (green), at $29^{\circ}N, 129^{\circ}W$. This location was selected within the region on the West coast of the United States where the R^2 is lower; the lower R^2 value appears to be due to a few extreme events that are not accurately predicted, rather than a consistent ill-performance.



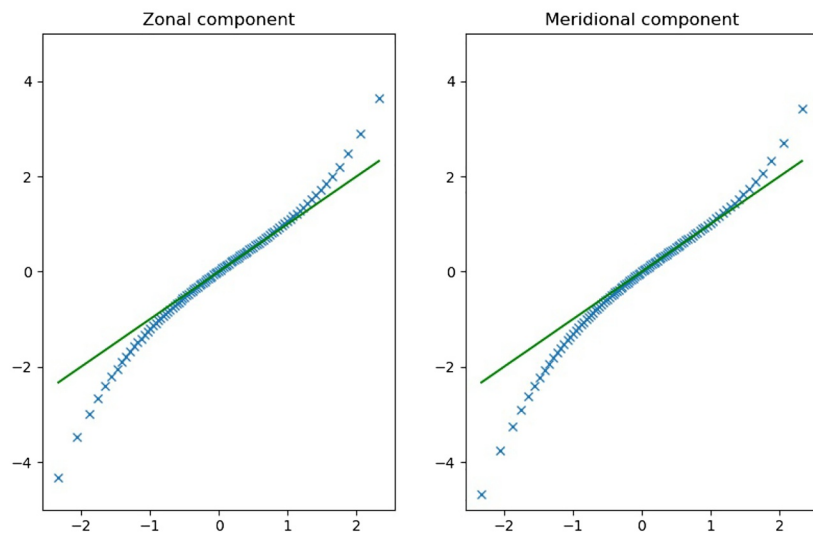


Figure D4. Sample log-probability distribution of true (purple) and stochastically simulated forcing (green) in the control simulation, for both components—zonal (left) and meridional (right). The histograms are shown on a log scale due to the hyperbolic-type distribution of the forcing.



(a) Sample distribution



(b) QQ plot

Figure D5. Distribution analysis of normalized residuals (Equation 16) of subgrid momentum forcing in the control simulation. (a) Sample distribution (blue) along with the probability density function of the standard normal distribution (red), (b) QQ plot (blue) of normalized residuals against the standard normal distribution, and line (green) defined by $y = x$.

Data Availability Statement

The authors downloaded the simulation's ocean surface velocities from a Pangeo data catalog at https://raw.githubusercontent.com/pangeo-data/pangeo-datastore/master/intake-catalogs/ocean/GFDL_CM2.6.yaml made publicly available by the Geophysical Fluid Laboratory. The code used in this study can be accessed from two repositories: <https://doi.org/10.5281/zenodo.4573438> for the data processing, neural network training and its tests; <https://doi.org/10.5281/zenodo.4573448> for its implementation in a shallow water model.

Acknowledgments

This research was made possible by Schmidt Futures, a philanthropic initiative founded by Eric and Wendy Schmidt, as part of its Virtual Earth System Research Institute (VESRI). The work was also supported in part by NSF-GEO 1912357, NOAA CVP NA19OAR4310364, and through the NYU IT High Performance Computing resources, services, and staff expertise. The authors thank Alistair Adcroft, Ryan Abernathy, Steve Griffies and V. Balaji for extremely valuable inputs at different stages of this work. The authors also thank the Associate Editor and the anonymous reviewers for their helpful comments which helped improve the manuscript.

References

- Abernathy, R., Augsburger, T., Banihirwe, A., Blackmon-Luca, C., Crone, T., Gentemann, C., et al. (2021). Cloud-native repositories for Big Scientific Data. *Computing in Science & Engineering*, 23, 26–35. <https://doi.org/10.1109/MCSE.2021.3059437>
- Bachman, S. D. (2019). The GM+E closure: A framework for coupling backscatter with the Gent and McWilliams parameterization. *Ocean Modelling*, 136, 85–106. <https://doi.org/10.1016/j.ocemod.2019.02.006>
- Balaji, V. (2021). Climbing down Charney's ladder: Machine learning and the post-Dennard era of computational climate science. *Philosophical Transactions of the Royal Society A*, 379(2194). <https://doi.org/10.1098/rsta.2020.0085>
- Berloff, P. S. (2005). Random-forcing model of the mesoscale oceanic eddies. *Journal of Fluid Mechanics*, 529, 71–95. <https://doi.org/10.1017/S0022112005003393>
- Berner, J., Achatz, U., Batté, L., Bengtsson, L., De La Cámara, A., Christensen, H. M., et al. (2017). Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565–587. <https://doi.org/10.1175/BAMS-D-15-00268.1>
- Bishop, C. (1991). *Mixture density networks*. Birmingham.
- Bolton, T., & Zanna, L. (2019). Applications of deep learning to ocean data inference and subgrid parameterization. *Journal of Advances in Modeling Earth Systems*, 11(1), 376–399. <https://doi.org/10.1029/2018MS001472>
- Brankart, J.-M. (2013). Impact of uncertainties in the horizontal density gradient upon low resolution global ocean modelling. *Ocean Modelling*, 66, 64–76. <https://doi.org/10.1016/j.ocemod.2013.02.004>
- Brenowitz, N. D., Beucler, T., Pritchard, M., & Bretherton, C. S. (2020). Interpreting and stabilizing machine-learning parametrizations of convection. *Journal of the Atmospheric Sciences*, 77(12), 4357–4375. <https://doi.org/10.1175/JAS-D-20-0082.1>
- Buizza, R., Miller, M., & Palmer, T. (1999). Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Quarterly Journal of the Royal Meteorological Society*, 125(560), 2887–2908. <https://doi.org/10.1256/smsqj.56005>
- Charney, J. G. (1971). Geostrophic turbulence. *Journal of Atmospheric Sciences*, 28(6), 1087–1095. [https://doi.org/10.1175/1520-0469\(1971\)028h1087:GTi2.0.CO;2](https://doi.org/10.1175/1520-0469(1971)028h1087:GTi2.0.CO;2)
- Crommelin, D., & Vanden-Eijnden, E. (2008). Subgrid-scale parameterization with conditional Markov chains. *Journal of the Atmospheric Sciences*, 65(8). <https://doi.org/10.1175/2008JAS2566.1>
- Davison, A. C. (2003). *Statistical models*. Cambridge University Press.
- Delworth, T. L., Rosati, A., Anderson, W., Adcroft, A. J., Balaji, V., Benson, R., et al. (2012). Simulated climate and climate change in the GFDL CM2.5 high-resolution coupled climate model. *Journal of Climate*, 25(8), 2755–2781. <https://doi.org/10.1175/JCLI-D-11-00316.1>
- Duan, J., & Nadiga, B. (2007). Stochastic parameterization for large eddy simulation of geophysical flows. *Proceedings of the American Mathematical Society*, 135(4), 1187–1196. <https://doi.org/10.1090/S0002-9939-06-08631-X>
- Dunne, J. P., Horowitz, L. W., Adcroft, A. J., Ginoux, P., Held, I. M., John, J. G., et al. (2020). The GFDL Earth System Model Version 4.1 (GFDL-ESM 4.1): Overall coupled model description and simulation characteristics. *Journal of Advances in Modeling Earth Systems*, 12(11), e2019MS002015. <https://doi.org/10.1029/2019MS002015>
- Ferrari, R., & Wunsch, C. (2009). Ocean circulation kinetic energy: Reservoirs, sources, and sinks. *Annual Review of Fluid Mechanics*, 41(1), 253–282. <https://doi.org/10.1146/annurev.fluid.40.111406.102139>
- Fox-Kemper, B., & Menemenlis, D. (2008). Can large eddy simulation techniques improve mesoscale rich ocean models? In: *Ocean modeling in an eddying regime* (pp. 319–337). American Geophysical Union (AGU). <https://doi.org/10.1029/177GM19>
- Frederiksen, J. S. (1999). Subgrid-scale parameterizations of eddy-topographic force, eddy viscosity, and stochastic backscatter for flow over topography. *Journal of the Atmospheric Sciences*, 56(11), 1481–1494. [https://doi.org/10.1175/1520-0469\(1999\)056<1481:SSPOET>2.0.CO;2](https://doi.org/10.1175/1520-0469(1999)056<1481:SSPOET>2.0.CO;2)
- Frederiksen, J. S., & Davies, A. G. (1997). Eddy viscosity and stochastic backscatter parameterizations on the sphere for atmospheric circulation models. *Journal of the Atmospheric Sciences*, 54(20), 2475–2492. [https://doi.org/10.1175/1520-0469\(1997\)054<2475:EVASBP>2.0.CO;2](https://doi.org/10.1175/1520-0469(1997)054<2475:EVASBP>2.0.CO;2)
- Gagne, II, D. J., Christensen, H. M., Subramanian, A. C., & Monahan, A. H. (2020). Machine learning for stochastic parameterization: Generative adversarial networks in the Lorenz '96 model. *Journal of Advances in Modeling Earth Systems*, 12(3), e2019MS001896. <https://doi.org/10.1029/2019MS001896>
- Gent, P. R., & McWilliams, J. C. (1989). Isopycnal mixing in ocean circulation models. *Journal of Physical Oceanography*, 20, 150–155. [https://doi.org/10.1175/1520-0485\(1990\)020<0150:IMIOCM>2.0.CO;2](https://doi.org/10.1175/1520-0485(1990)020<0150:IMIOCM>2.0.CO;2)
- Gerard, L. (2007). An integrated package for subgrid convection, clouds, and precipitation compatible with meso-gamma scales. *Quarterly Journal of the Royal Meteorological Society*, 133, 711–730. <https://doi.org/10.1002/qj.58>
- Germano, M., Piomelli, U., Moin, P., & Cabot, W. H. (1991). A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics*, 3(7), 1760–1765. <https://doi.org/10.1063/1.857955>
- Griffies, S. M. (2015). *A handbook for the GFDL CM2-O model suite*. Retrieved from https://mom-ocean.github.io/assets/pdfs/CM2_O_suite_1990_tutorial_and_analysis.pdf
- Griffies, S. M., Winton, M., Anderson, W. G., Benson, R., Delworth, T. L., Dufour, C. O., et al. (2015). Impacts on ocean heat from transient mesoscale eddies in a hierarchy of climate models. *Journal of Climate*, 28(3), 952–977. <https://doi.org/10.1175/JCLI-D-14-00353.1>
- Hallberg, R. (2013). Using a resolution function to regulate parameterizations of oceanic mesoscale eddy effects. *Ocean Modelling*, 72, 92–103. <https://doi.org/10.1016/j.ocemod.2013.08.007>
- IPCC. (2013). Climate change 2013: The physical science basis. In: *Contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change [Book]*. NY, USA. Cambridge University Press. <https://doi.org/10.1017/CBO9781107415324>

- Jansen, M. F., Adcroft, A., Khani, S., & Kong, H. (2019). Toward an energetically consistent, resolution aware parameterization of ocean mesoscale eddies. *Journal of Advances in Modeling Earth Systems*, 11(8), 2844–2860. <https://doi.org/10.1029/2019MS001750>
- Jones, T. R., Randall, D. A., & Branson, M. D. (2019). Multiple-instance superparameterization: 2. The effects of stochastic convection on the simulated climate. *Journal of Advances in Modeling Earth Systems*, 11(11), 3521–3544. <https://doi.org/10.1029/2019MS001611>
- Jospin, L. V., Buntine, W., Boussaid, F., Laga, H., & Bennamoun, M. (2020). *Hands-on bayesian neural networks – A tutorial for deep learning users*. Cornell University.
- Juricke, S., Palmer, T. N., & Zanna, L. (2017). Stochastic perturbations to sub-grid scale ocean mixing: Impacts on low frequency variability. *Journal of Climate*, 30(13), 4997–5019. <https://doi.org/10.1175/JCLI-D-16-0539.1>
- Kingma, D. P., & Ba, J. L. (2015). *Adam: A method for stochastic optimization*. arXiv preprint:1412.6980v9 (pp. 1–15). Cornell University.
- Kitsios, V., Frederiksen, J. S., & Zidikheri, M. J. (2016). Theoretical comparison of subgrid turbulence in atmospheric and oceanic quasi-geostrophic models. *Nonlinear Processes in Geophysics*, 23(2), 95–105. <https://doi.org/10.5194/npg-23-95-2016>
- Kjellsson, J., & Zanna, L. (2017). The impact of horizontal resolution on energy transfers in global ocean models. *Fluid*, 2(3), 45. <https://doi.org/10.3390/fluids2030045>
- Klöwer, M., Jansen, M. F., Claus, M., Greatbatch, R. J., & Thomsen, S. (2018). Energy budget-based backscatter in a shallow water model of a double gyre basin. *Ocean Modelling*, 132, 1–11. <https://doi.org/10.1016/j.ocemod.2018.09.006>
- Kraichnan, R. H. (1959). The structure of isotropic turbulence at very high Reynolds numbers. *Journal of Fluid Mechanics*, 5(4), 497–543. <https://doi.org/10.1017/S0022112059000362>
- Kraichnan, R. H. (1967). Inertial ranges in two-dimensional turbulence. *The Physics of Fluids*, 10(7), 1417–1423. <https://doi.org/10.1063/1.1762301>
- Larichev, V. D., & Held, I. M. (1995). Eddy amplitudes and fluxes in a homogeneous model of fully developed baroclinic instability. *Journal of physical oceanography*, 25(10), 2285–2297. [https://doi.org/10.1175/1520-0485\(1995\)025<2285:EAAFFIA>2.0.CO;2](https://doi.org/10.1175/1520-0485(1995)025<2285:EAAFFIA>2.0.CO;2)
- Leith, C. E. (1990). Stochastic backscatter in a subgrid-scale model: Plane shear mixing layer. *Physics of Fluids A: Fluid Dynamics*, 2(3), 297–299. <https://doi.org/10.1063/1.857779>
- Lilly, D. K. (1992). A proposed modification of the germano subgrid-scale closure method. *Physics of Fluids A: Fluid Dynamics*, 4(3), 633–635. <https://doi.org/10.1063/1.858280>
- Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 3431–3440).
- Majda, A. J., Timofeyev, I., & Vanden Eijnden, E. (1999). Models for stochastic climate prediction. *Proceedings of the National Academy of Sciences*, 96(26), 14687–14691. <https://doi.org/10.1073/pnas.96.26.14687>
- Marshall, D. P., & Adcroft, A. (2010). Parameterization of ocean eddies: Potential vorticity mixing, energetics and Arnold’s first stability theorem. *Ocean Modelling*, 32(3–4), 188–204. <https://doi.org/10.1016/j.ocemod.2010.02.001>
- Martínez-Moreno, J., Hogg, A. M., England, M. H., Constantinou, K. A. E., Navid, C., & Morrison, A. K. (2021). Global changes in oceanic mesoscale currents over the satellite altimetry record. *Nature Climate Change*, 66. <https://doi.org/10.1038/s41558-021-01006-9>
- Mason, P. J., & Thomson, D. J. (1992). Stochastic backscatter in large-eddy simulations of boundary layers. *Journal of Fluid Mechanics*, 242(28), 51–78. <https://doi.org/10.1017/S0022112092002271>
- O’Gorman, P. A., & Dwyer, J. G. (2018). Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events. *Journal of Advances in Modeling Earth Systems*, 10(10), 2548–2563. <https://doi.org/10.1029/2018MS001351>
- Palmer, T. N. (2001). A nonlinear dynamical perspective on model error: A proposal for non-local stochastic-dynamic parameterization in weather and climate prediction models. *Quarterly Journal of the Royal Meteorological Society*, 127(572), 279–304. <https://doi.org/10.1002/qj.49712757202>
- Palmer, T. N. (2012). Towards the probabilistic earth-system simulator: A vision for the future of climate and weather prediction. *Quarterly Journal of the Royal Meteorological Society*, 138(665), 841–861. <https://doi.org/10.1002/qj.1923>
- Porta Mana, P., & Zanna, L. (2014). Toward a stochastic parameterization of ocean mesoscale eddies. *Ocean Modelling*, 79, 1–20. <https://doi.org/10.1016/j.ocemod.2014.04.002>
- Rasp, S., Pritchard, M. S., & Gentile, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689. <https://doi.org/10.1073/pnas.1810286115>
- Rhines, P. B. (1977). The dynamics of unsteady currents. *The Sea*, 6: *Marine Modelling*, 189–318.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-net: Convolutional networks for biomedical image segmentation*. Cornell University.
- Salmon, R. (1980). Baroclinic instability and geostrophic turbulence. *Geophysical & Astrophysical Fluid Dynamics*, 15(1), 167–211. <https://doi.org/10.1080/03091928008241178>
- Sirignano, J., MacArt, J. F., & Freund, J. B. (2020). DPM: A deep learning PDE augmentation method with application to large-eddy simulation. *Journal of Computational Physics*, 423, 109811. <https://doi.org/10.1016/j.jcp.2020.109811>
- Stanley, Z., Grooms, I., Kleiber, W., Bachman, S. D., Castruccio, F., & Adcroft, A. (2020). Parameterizing the impact of unresolved temperature variability on the large-scale density field: Part 1. Theory. *Journal of Advances in Modeling Earth Systems*, 12(12), e2020MS002185. <https://doi.org/10.1029/2020MS002185>
- Treguier, A. M., Held, I. M., & Larichev, V. D. (1997). Parameterization of quasi-geostrophic eddies in primitive equation ocean models. *Journal of Physical Oceanography*, 567–580. [https://doi.org/10.1175/1520-0485\(1997\)027<0567:POQEIP>2.0.CO;2](https://doi.org/10.1175/1520-0485(1997)027<0567:POQEIP>2.0.CO;2)
- Waterman, S., & Jayne, S. R. (2011). Eddy-mean flow interactions in the along-stream development of a western boundary current jet: An idealized model study. *Journal of Physical Oceanography*, 41(4), 682–707. <https://doi.org/10.1175/2010JPO4477.1>
- Williams, P. D., Howe, N. J., Gregory, J. M., Smith, R. S., & Joshi, M. M. (2016). Improved climate simulations through a stochastic parameterization of ocean eddies. *Journal of Climate*, 29(24), 8763–8781. <https://doi.org/10.1175/JCLI-D-15-0746.1>
- Yuval, J., & O’Gorman, P. A. (2020). Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions. *Nature Communications*, 11(1), 1–10. <https://doi.org/10.1038/s41467-020-17142-3>
- Zanna, L., Bachman, S., & Jansen, M. (2020). Energizing turbulence closures in ocean models. *CLIVAR Exchanges/US CLIVAR Variations*, 18(1), 3–8. <https://doi.org/10.5065/g8w0-fy32>
- Zanna, L., & Bolton, T. (2020). Data-driven equation discovery of ocean mesoscale closures. *Geophysical Research Letters*, 47, e2020GL088376. <https://doi.org/10.1029/2020GL088376>
- Zanna, L., Brankart, J., Huber, M., Leroux, S., Penduff, T., & Williams, P. (2018). Uncertainty and scale interactions in ocean ensembles: From seasonal forecasts to multidecadal climate predictions. *Quarterly Journal of the Royal Meteorological Society*. <https://doi.org/10.1002/qj.3397>
- Zanna, L., Porta Mana, P. G. L., Anstey, J., David, T., & Bolton, T. (2017). Scale-aware deterministic and stochastic parametrizations of eddy-mean flow interaction. *Ocean Modelling*, 111, 66–80. <https://doi.org/10.1016/j.ocemod.2017.01.004>