A Deep Learning Approach to Spatiotemporal Sea Surface Height Interpolation and Estimation of Deep Currents in Geostrophic Ocean Turbulence

Georgy E. Manucharyan, Lia Siegelman, Patrice Klein

JAMES | Journal of Advances in Modeling Earth Systems

RESEARCH ARTICLE
10.1029/2019MS001965

Key Points:

- The efficacy of Deep Learning in exploiting sparse sea surface height (SSH) data is demonstrated in a quasigeostrophic model
- Residual Neural Networks are superior to linear and dynamical interpolation techniques for SSH reconstruction
- Deep networks skillfully reconstruct the unobserved deep flows from SSH data to provide a complete state estimate of geostrophic turbulence

Correspondence to:
G. E. Manucharyan,
gmanuch@uw.edu

Citation:
Manucharyan, G. E., Siegelman, L., & Klein, P. (2021). A deep learning approach to spatiotemporal sea surface height interpolation and estimation of deep currents in geostrophic ocean turbulence. *Journal of Advances in Modeling Earth Systems*, *13*, e2019MS001965. https://doi.org/10.1029/2019MS001965

Received 28 NOV 2019
Accepted 1 DEC 2020

A Deep Learning Approach to Spatiotemporal Sea Surface Height Interpolation and Estimation of Deep Currents in Geostrophic Ocean Turbulence

Georgy E. Manucharyan[1], Lia Siegelman[2], and Patrice Klein[2,3,4]

[1]School of Oceanography, University of Washington, Seattle, WA, USA, [2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, [3]Laboratoire de Météorologie Dynamique, Ecole Normale Supérieure, CNRS, Paris, France, [4]Laboratoire d'Oceanographie Physique et Spatiale, IFREMER, CNRS, Brest, France

Abstract Satellite altimeters provide global observations of sea surface height (SSH) and present a unique data set for advancing our theoretical understanding of upper-ocean dynamics and monitoring its variability. Considering that mesoscale SSH patterns can evolve on timescales comparable to or shorter than satellite return periods, it is challenging to accurately reconstruct the continuous SSH evolution as currently available altimetry observations are still spatially and temporally sparse. Here we explore the possibility of SSH interpolation via Deep Learning by using synthetic observations from an idealized quasigeostrophic model of baroclinic ocean turbulence. We demonstrate that Convolutional Neural Networks with Residual Learning are superior in SSH reconstruction to linear and recently developed dynamical interpolation techniques. Also, the deep neural networks can provide a skillful state estimate of unobserved deep ocean currents at mesoscales. These conspicuous results suggest that SSH patterns of eddies might contain substantial information about the underlying deep ocean currents that are necessary for SSH prediction. Our training data are focused on highly idealized physics and diversification of processes needs to be considered to more accurately represent the real ocean. In addition, methodological improvements such as transfer learning and implementation of dynamically aware loss functions might be necessary to consider before its ultimate use with real satellite observations. Nonetheless, by providing a proof of concept based on synthetic data, our results point to deep learning as a viable alternative to existing interpolation and, more generally, state estimation methods for satellite observations of eddying currents.

Plain Language Summary Satellite observations of sea surface height (SSH) are widely used to derive surface ocean currents on a global scale. However, due to gaps in SSH observations, it remains challenging to retrieve the dynamics of rapidly evolving upper-ocean currents. To overcome this limitation, we propose a Deep Learning framework that is based on pattern recognition extracted from SSH observations. Using synthetic data generated from a simplified model of ocean turbulence, we demonstrate that deep learning can accurately estimate both surface and subsurface ocean currents, significantly outperforming the most commonly used techniques. By providing a proof of concept, our study highlights the strong potential of deep learning for estimating ocean currents from satellite observations.

1. Introduction

Mesoscale eddies are a key component of the global ocean circulation and the Earth's climate as a whole due to their influence on mean currents, heat and salt transport, atmosphere-ocean interactions, and biological productivity (Ferrari & Wunsch, 2009; Klein et al., 2019). The strong interactions between mesoscale eddies constitutes geostrophic turbulence (Charney, 1971), characteristics of which significantly vary depending on large-scale ocean currents and the associated local instabilities as reviewed in (Tulloch et al., 2011). Eddies of O (100 km) in size are in near geostrophic balance and hence their surface streamfunction is directly proportional to their sea surface height (SSH) signature. In fact, mesoscale eddies explain most of the SSH variance in regions with strong baroclinic currents, for example, the Gulf Stream in the North Atlantic, the Kuroshio Extension in the North Pacific, and the Antarctic Circumpolar Current (Chelton et al., 2011). Thus, satellite altimeters continue providing crucial SSH observations for monitoring the spatiotemporal

evolution of the mesoscale eddy field (Chelton et al., 2011; Fu et al., 2010), although mapping the along-track altimetry observations leads to a relatively coarse resolution of gridded SSH products (Ballarotta et al., 2019; Chelton & Schlax, 2003). Hence, understanding and monitoring quantities like the oceanic kinetic energy spectrum and the associated spectral energy fluxes (Aluie et al., 2018; Scott & Arbic, 2007), understanding tracer dispersion (Abernathey & Marshall, 2013) or inferring subsurface flows (Klein et al., 2009) remain challenging because these quantities depend on higher-order SSH derivatives that are resolution-sensitive.

To increase the density of SSH observations, several altimeters have been put in orbit but their 10–20 days repeat orbits and relatively coarse along-track resolutions allow to view the ocean dynamics only down to relatively large mesoscale eddies of O (100) km wavelengths (Ballarotta et al., 2019; Chelton & Schlax, 2003; Wunsch, 2010). Thus, the SSH interpolation can be especially challenging in regions with energetic baroclinic turbulence where the evolution of small-scale SSH anomalies can be fast compared to the satellite return periods. The upcoming Surface Water Ocean Topography (SWOT) altimeter mission promises to observe ocean mesoscale eddies and submesoscale fronts (≤50 km) at unprecedented spatial resolutions, potentially resolving 15–30 km wavelengths (Fu & Ubelmann, 2014; Morrow et al., 2019). However, with its complete repeat cycle of 21 days, the temporal resolution of the altimeter is insufficient to continuously capture the evolution of submesoscale eddies. Thus, the mismatch between the high spatial resolution and the moderate temporal resolution presents a challenge for reconstructing time-continuous maps of SSH.

The existing gridded SSH products, for example, AVISO (Ducet et al., 2000), are spatially and temporally interpolated from along-track data using objective mapping and hence their accuracy and effective resolution are constrained by the density of observations and deficiencies of the interpolation technique. The temporal SSH interpolation could be conceptually viewed as reconstructing the phase-space trajectory given only partial observations of the two endpoints separated in time. A major complication arises due to the chaotic nature of ocean turbulence for which there may be multiple plausible phase-space trajectories passing within some close vicinity of any given endpoints. Thus, the task of temporal interpolation, that is, finding the true (or most likely) trajectory, becomes increasingly more difficult with an increasing time separation between observations. Most commonly used interpolation techniques, such as objective mapping or polynomial interpolation, do not attempt to make use of any potential dynamical constraints present in the data and perform well only for autocorrelated data while failing for sparse data. It is thus crucial to develop frameworks to efficiently extract information about the oceanic eddy dynamics from the spatially and temporally sparse SSH observations. Note, throughout the study, we use the term *information* qualitatively and do not attempt to quantify it using Shannon's information theory. Below we discuss how the nature of baroclinic ocean turbulence provides dynamical limitations for SSH interpolation and why deep learning might be a viable alternative to other interpolation techniques.

1.1. SSH Interpolation and the Associated Dynamical Limitations

Spatiotemporal interpolation or gridding of SSH data is inherently linked to ocean eddy physics as the success of a given technique ultimately should rely on the pertinence of its assumed model (either dynamical or statistical) that captures the essence of mesoscale eddy propagation in space and time. To illustrate this point, imagine a coherent eddy moving in a turbulent field and several altimeter tracks passing through it at different times and directions. If there is an accurate model of eddy propagation, it would allow pinpointing the observations taken over this specific eddy and combining this data to better constrain the two-dimensional eddy shape. Thus, to extract the relevant information from various altimetry tracks to the fullest extent, it is necessary to have an accurate model of eddy evolution. However, due to the stratified nature of geostrophic ocean turbulence, the unobserved deep ocean flows can affect the surface dynamics, and hence the SSH observations on their own may not be self-sufficient to infer its evolution. Given the lack of subsurface data at eddy scales, constructing a closed system of equations for SSH evolution is challenging.

Another complication for SSH interpolation arises due to the chaotic nature of geostrophic turbulence that has an increasingly higher sensitivity to initial conditions as time progresses. Upon the dynamical evolution, any ordered set of state-space trajectories will become increasingly stretched and filamented over time, eventually equidistributing its volume over the entire state space, that is, undergoing measure-theoretic and topological mixing. Thus, with increasing time separation between any two observations of a turbulent eddy field, the relation between them becomes increasingly more convoluted. For sufficiently large time

separations, the two sequential observations would become effectively independent from each other and hence interpolating between them would not be plausible. Combined with the fact that satellites only provide approximate and partial observations of the ocean, the problem of temporal SSH interpolation can be underconstrained, that is, it might not have unique solutions as not enough information is given.

Existing methods for spatiotemporal SSH interpolation can be broadly split into two distinct classes: methods that rely on a postulated dynamical model of SSH evolution and purely data-driven methods. Both methods have their advantages and disadvantages. To avoid prescribing a dynamical model, statistical models like objective interpolation (Davis, 1985; Ducet et al., 2000; Le Traon et al., 1998) rely on data only. Their premise is to incorporate spatiotemporal correlations and measurement error into a statistical model and provide the most likely estimate of the true continuous field under consideration. However, this method does not rely on any dynamical model of the eddy propagation and hence can lead to an unphysical behavior of the interpolated SSH field. Methods involving dynamical ocean models are typically based on data assimilation, a procedure that minimizes the difference between the observed and modeled fields by adjusting unknown variables like boundary and initial conditions or external forcing (see e.g., reanalysis product by Carton & Giese, 2008). While resulting in SSH fields that are dynamically constrained, this method suffers from a drawback that it requires additional observations to constrain other essential model variables like the subsurface flow and/or the density field. Also, data assimilation for complex ocean models at eddy-resolving scales is often underdetermined and is computationally demanding.

A recent study by (Ubelmann et al., 2015) demonstrated that representing SSH propagation with a single equivalent barotropic mode (that mimics the first baroclinic mode) in a quasigeostrophic (QG) model results in significant improvements in the spatiotemporal interpolation of sparse SSH observations. In particular, Ubelmann et al. (2015) considered a fundamental problem of reconstructing the SSH distribution that occurred in between two observed SSH fields separated by about 20 days, a characteristic timescale required by a set of altimeters to reconstruct a spatial SSH field. They found that integrating the earlier SSH observation forward in time (following the assumed dynamics of an equivalent barotropic mode) and averaging it with the later observed SSH anomalies that were integrated backward in time resulted in an improvement compared to conventional linear interpolation methods. In follow-up work, Ubelmann et al. (2016) generalized this temporal interpolation method to the spatiotemporal interpolation of along-track SSH observations by essentially performing data assimilation on the one-layer QG model. The advantage of the dynamical interpolation method is that it relies on the advection of potential vorticity—a nonlinear process that is inherently present in ocean dynamics and cannot be represented by linear or objective interpolation techniques.

A drawback of the dynamical interpolation is that it assumes that SSH evolves independently of deep ocean flows, considering the so-called equivalent barotropic mode dynamics (Berloff & Meacham, 1997). However, in many energetic regions of the ocean, for example, in Gulf Stream, Kuroshio or Antarctic Circumpolar Current, the currents are baroclinically unstable and hence are necessarily composed of at least two dynamically interacting vertical modes, the barotropic and baroclinic modes (see e.g., Chapter 6 in Vallis, 2017). To illustrate this point, consider the conservation of a quasigeostrophic potential vorticity q_1 in the upper-ocean layer as a model of SSH evolution at mesoscales:

$$\frac{Dq_1}{Dt} = \underbrace{\frac{D}{Dt}[\nabla^2\left(\psi_1 - R_d^{-2}\psi_1\right) + \beta y]}_{\text{Depends on partially observed } \psi_1} + \underbrace{R_d^{-2}\frac{D}{Dt}\psi_{b.t.}}_{\text{Depends on unobserved } \psi_2} \approx 0, \tag{1}$$

$$\text{where } \psi_{b.t.} = \frac{H_1\psi_1 + H_2\psi_2}{H_1 + H_2} \text{ and } R_d^{-2} = \frac{f_0^2}{g'H_1} + \frac{f_0^2}{g'H_2}, \tag{2}$$

ψ_1 and ψ_2 are the surface and deep ocean stream functions, $\psi_{b.t.}$ is the barotropic streamfunction (depth-averaged transport), R_d is the Rossby baroclinic deformation radius, f_0 and β are the Coriolis and beta-plane parameters, y is the meridional coordinate, H_1 and H_2 are the ocean layer depths, g' is the reduced gravity, and D/Dt is the material derivative accounting for advection by the surface flow (see Section 2). Note that the surface streamfunction is directly proportional to SSH: $\psi_1 = (g/f_0)$SSH, where g is the acceleration due to

gravity. On relatively short timescales, sources and dissipation of potential vorticity could be neglected and its approximate conservation provides a basic description of the eddy evolution. The terms in the Equation 1 above have been grouped into those that only depend on the partially observed ψ_1 (or equivalently SSH) and terms that depend on the unobserved subsurface flow ψ_2 (or on the barotropic flow $\psi_{b.t.}$). By considering only the equivalent barotropic mode dynamics and taking ψ_1 to be equal to the baroclinic mode, the dynamical interpolation method as described in (Ubelmann et al., 2015, 2016) discards the term in the PV-conservation equation that depends on the unobserved barotropic streamfunction, resulting in

$$\frac{D}{Dt}\left[\nabla^2\psi_1 - \left(R_d^{-2}\psi_1\right) + \beta y\right] = 0. \tag{3}$$

Since the discarded term is the only term that depends on the unknown streamfunction ψ_2, it is possible to integrate the approximate PV-conservation equation forward and backward in time given only ψ_1 observations, as was done in (Ubelmann et al., 2015). Even though in many ocean regions both deep and surface geostrophic currents are dynamically active, reconstructing SSH using the dynamical interpolation technique proved to be superior to linear interpolation methods (Ubelmann et al., 2015) because it relies, at least approximately, on the fundamental PV-conservation constraint. Nonetheless, the dynamical interpolation method can lead to significant errors (see Section 3), implying that the omitted term, while being relatively small, can substantially impact SSH evolution on timescales comparable to return periods of altimetry satellites.

1.2. The Rationale for a Deep Learning Approach.

A clear way of improving the dynamical interpolation algorithm would be to take into account the contribution of the barotropic mode to SSH evolution. However, comprehensive measurements of deep ocean currents at eddy scales are missing, posing a significant challenge of inferring them from only SSH observations. Without taking into consideration the physical processes that have led to the generation of any given SSH snapshot, there is a wide range of plausible ways in which ψ_1 could be decomposed into baroclinic and barotropic modes, each corresponding to the distinct configuration of PV anomalies in the deep and surface layers. However, considering that PV anomalies are specifically due to baroclinic instabilities obeying specific conservation laws (Equation 1), the corresponding barotropic and baroclinic modes are inherently entangled, providing at least partial constraints on how any specific SSH pattern could be partitioned into modes.

Since the QG model exhibits a highly nonlinear and chaotic behavior, an analytical approach to disentangle the modes has not been found but the evidence that data-driven approach might be relevant has been presented in the literature. In particular, the surface and subsurface flows from mooring observations are significantly correlated such that a single Empirical Orthogonal Function can explain a significant amount of variance of the overall vertical velocity profile (de La Lama et al., 2016; Wunsch, 1997). Furthermore, machine learning techniques such as self-organizing maps (Chapman & Charantonis, 2017), as well as convolutional neural networks (CNNs) (Bolton & Zanna, 2019), have been successfully used to estimate the subsurface flows from SSH data. However, the unknown term $D\psi_{bt} / Dt = \left(\partial_t + \mathbf{u_1}\cdot\nabla\right)\psi_{bt}$ in Equation 1 can only provide a substantial contribution to the PV budget if ψ_{bt} has a substantial component that is decorrelated from ψ_1 because $\mathbf{u}_1\cdot\nabla\psi_1 \equiv 0$, and $\partial_t\psi_{bt} \ll \partial_t\psi_1$ for surface-amplified flows. Thus the key for a more accurate SSH interpolation lies in estimating the component of ψ_2 that is decorrelated from ψ_1—a problem that is tightly linked to estimating eddy heat fluxes in baroclinically unstable flows. Using residual neural networks, (George et al., 2021) demonstrated that ψ_1 indeed contains substantial information about the decorrelated part of the subsurface streamfunction ψ_2, allowing to estimate about 60% of the variance in eddy heat fluxes only from SSH snapshots. Given that machine learning methods can utilize SSH patterns to estimate the component of ψ_{bt} that is uncorrelated with ψ_1 for estimation of the eddy heat fluxes, it is plausible that they could be used for SSH interpolation as well.

While ocean turbulence is chaotic and appears to be random and unpredictable, it does not prohibit characteristics that are particularly beneficial for deep learning: the emergence of underlying repeating patterns, self-similarities, and self-organization. We thus hypothesize that deep learning techniques could outperform conventional interpolation methods including linear and dynamical interpolation. In this manuscript,

we use synthetic model observations related to the QG turbulence in the Gulf Stream, typically forced by the baroclinic instability of a westerly sheared flow (Tulloch et al., 2011). These model observations only involve mesoscale eddies and no submesoscales (resulting for example from mixed-layer instabilities) and do not include other processes, such as Ekman pumping or incoherent internal tides, known to impact SSH. These model observations are used to present a first proof of concept for using deep learning to shortcut the formal process of data assimilation and reconstruct not only the interpolated SSH field but also the corresponding unobserved deep ocean currents, thus providing a complete state estimate of the baroclinic ocean turbulence.

The manuscript is organized in the following way. In Section 2, we present a range of neural network (NN) architectures, outline a set of training experiments, and describe the synthetic model of ocean turbulence that we used to evaluate the efficacy of deep learning in SSH interpolation and state estimation of both surface and deep ocean streamfunctions. In Section 3, we present examples of SSH estimates using deep NNs and compare their skills to linear and dynamical interpolation techniques. In Section 4, we discuss the broader implications of our results, outline the deficiencies and advantages of our deep learning methodology, and propose possible improvements to generalize our method for its ultimate use with real satellite observations.

2. Methods

We implement a range of deep NN architectures to address a basic question of interpolating SSH fields in baroclinic ocean turbulence. To exclude potential limitations of real-world data, our study is entirely based upon synthetic data that we generate using the QG model of baroclinic ocean turbulence. We find the QG model to be optimal for our goals as it is pertinent to many energetic regions in the ocean while being relatively simple such that a large volume of data can be generated for training and testing; furthermore, the model allows us to directly benchmark deep learning against the dynamical interpolation technique that also utilizes QG dynamics. Below we describe our NN architectures, the QG model used for the generation of training and testing data sets, and the details of the dynamical interpolation that we implemented for direct skill comparisons with deep learning and linear interpolation.

2.1. Deep Learning Framework: Residual Convolutional Neural Networks

Artificial NNs are based on the idea of approximating the "output" by taking the "input" variable and performing a large number of matrix additions and multiplications, applying nonlinearity functions, and either condensing or expanding the variable dimension as it passes from layer to layer. The resulting network contains a large number of free parameters that are later adjusted to optimize a given loss function, commonly taken as a measure of difference between the prediction and the truth. Because we are trying to extract information from the eddy patterns expressed in SSH fields, the choice of CNNs is rationalized. In passing the data from layer to layer, CNNs use a set of convolutional filters (kernel matrices with prescribed dimensions) to produce increasingly more abstract levels of data representation that are passed on to the next layer. Here we implement the ResNet architecture—a CNN with Residual Learning blocks (He et al., 2016). The Residual Learning is a process by which the data are not only transferred sequentially from one layer to another but is also transferred by skipping several layers via the so-called skip connections (blue arrows in Figure 1). The presence of skip connections can result in better performances for a wide range of computer vision problems (Targ et al., 2016).

A brief description of the ResNet architecture as shown schematically in Figure 1 follows. The input consists of two SSH snapshots represented by a (32,32,2) matrix. The very first convolutional layer takes the input and applies a set of 32 convolutional filters of size (5,5) with a stride of (1,1), followed by the batch normalization, the nonlinearity function taken to be the Leaky Rectified Linear Unit (Leaky ReLU), and the maximum 2D pooling of size (2,2) with a stride of (2,2). Next, a series of residual learning blocks follow, each consisting of two convolutional layers that take the input with M channels and apply N filters, each followed by batch normalization and Leaky ReLu, and at the very end of the residual block, its initial input matrix is added to its output (see Figure 1). The architecture has a total of 16 residual blocks containing 52 convolutional layers. The first series of residual blocks consist of three blocks that transform the input from

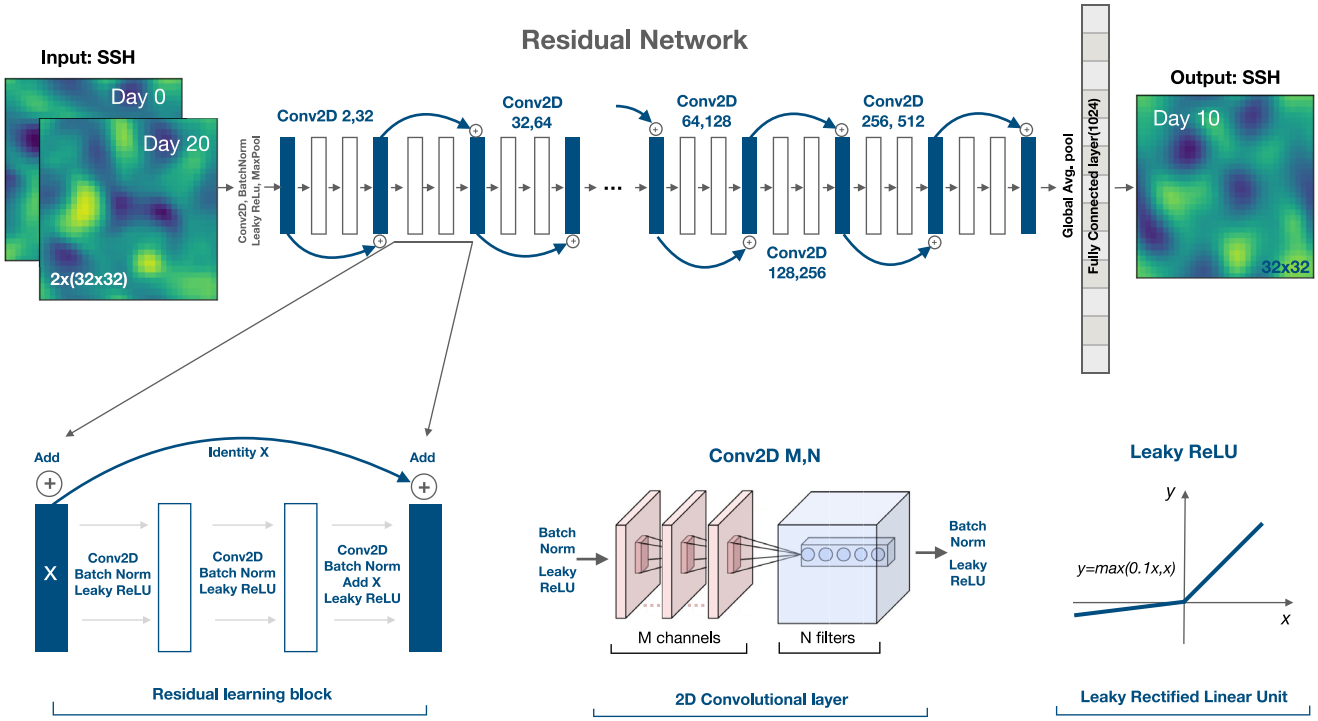

Figure 1. The ResNet architecture of a deep CNN with residual learning that was used for SSH interpolation and state estimation. The input consists of two SSH snapshots separated by 20 days. A set of convolutional layers are then applied to create an abstract representation of the input patterns in a bottleneck fashion: when image sizes decrease by a factor of two, the number of filters increases by a factor of two. Each convolutional layer is followed by the batch normalization and the application of the nonlinear function (Leaky Rectified Linear Unit). Residual learning blocks are saving the output from one layer and adding its identity to the output several layers ahead (blue arrows). The output from the final convolutional layer is subject to a global average pooling and flattening into a vector that is densely connected to the output of the appropriate dimension to represent either a single or multiple fields. CNN, convolutional neural network; SSH, sea surface height.

$M = 32$ to $M = 64$ channels while reducing the matrix rows and columns by a factor of two using the (2,2) max pooling. Next, a set of four blocks transform the input to 128 channels, a set of six blocks to 256, and a set of three blocks to 512 channels, and the matrix dimension becomes (2,2,512). Then, a global two-dimensional average pooling is applied to have a vector of length 512, which is in some experiments subjected to a dropout rate of 20%. The resulting vector is then densely connected to a vector of size 1024, which is finally reshaped to represent the output SSH snapshot of size (32,32). For our state estimation experiments with four separate fields appearing as the output matrix, the ResNet architecture remains the same except for the final dense layer being of length 4,096 and reshaped to the appropriate output size of (32,32,4).

We have explored more complex ResNets (going up to 161 convolutional layers) but also simpler CNN architectures without residual learning as well as shallow feed-forward networks (see Table 1). A brief description of the NN architectures follows. **FC**: feed-forward NN with 2 hidden layers (254 and 512 neurons correspondingly), batch normalization, and leaky ReLU as an activation function after each hidden layer. **FC_Large**: same as **FC** but with 512 and 1024 neurons in the hidden layers. **VGG**: CNN with 32 (4×4) filters in the first layer, 64 (3×3) in the second, 128 (3×3) in the third, 256 (2×2) in the fourth, with batch normalization and leaky ReLU used after each layer and the two-dimensional global average pooling before connecting to the dense layer. **VGG_Large**: same as **VGG** but using a four times larger number of filters in each convolutional layer. **VGG_Deep**: same as **VGG** but repeating each convolutional layer three times before proceeding to the next one. **ResNet_Small, ResNet, and ResNet_Large** are residual NNs with architectures as depicted in Figure 1 but with a total of 31,52, and 161 convolutional layers correspondingly; **_Dropout** denotes the use of 20% dropout rate in the last layer. We have implemented the architectures in Tensorflow/Keras and provided the Python scripts along with the training data in the Zenodo data repository (Manucharyan, 2020).

Table 1
List of NN Training Experiments Demonstrating the Achieved Prediction Skill for Temporal Interpolation of SSH Snapshots

#	Architecture	Parameters	Data samples	ΔT (days)	Skill
1	FC	1.2×10^6	2×10^5	20	0.53
2	FC_Large	6.3×10^6	2×10^5	20	0.54
3	VGG	0.5×10^6	2×10^5	20	0.63
4	VGG_Large	4.6×10^6	2×10^5	20	0.64
5	VGG_Deep	1.4×10^6	2×10^5	20	0.61
6	ResNet_Small	0.9×10^6	2×10^5	20	0.69
7	ResNet_Large	7×10^6	2×10^5	20	0.72
8	ResNet_Large_Dropout	7×10^6	2×10^5	20	0.72
9	ResNet_Dropout	4.7×10^6	2×10^5	20	0.73
10	ResNet	4.7×10^6	2×10^5	20	0.75
11	ResNet	4.7×10^6	2×10^5	40	0.44
12	ResNet	4.7×10^6	2×10^5	60	0.18
13	ResNet	4.7×10^6	1×10^5	20	0.71
14	ResNet	4.7×10^6	4×10^4	20	0.65
15	ResNet	4.7×10^6	2×10^4	20	0.58
16	ResNet	4.7×10^6	8×10^3	20	0.55
17	ResNet	4.7×10^6	4×10^3	20	0.44
18	ResNet	4.7×10^6	1×10^3	20	0.39
19	ResNet	4.7×10^6	5×10^2	20	0.33

Notes. Experiments 1–10 explore various architectures, 11–12 explore the skill deterioration with increasing time separation between the input images, and 13–19 explore skill dependence on the number of training examples. The architecture names correspond to function names in the provided NetworkArchitectures.py script that encodes their graphs using TensorFlow/Keras. The parameters column represents the number of trainable NN parameters for corresponding architectures. The Data Samples column denotes the number of input-output examples that were used in NN training. The ΔT column denotes the time separation between the two input snapshots of SSH, and the skill column denotes the maximum achieved skill on validation data.

Abbreviations: NN, neural network; SSH, sea surface height.

As a performance metric, we define the model skill that is proportional to the loss function and normalized by the standard deviation of the SSH signal in the following way:

$$Skill = 1 - \left(\frac{\overline{| SSH_{predicted} - SSH_{true} |^2}}{\overline{| SSH_{true} |^2}} \right)^{\frac{1}{2}}. \quad (4)$$

For reference, the maximum skill = 1 is achieved when the predicted and true images are identical; the skill = 0 corresponds to a prediction that makes the same error as assuming a spatially homogeneous SSH field, and negative skill implies an even worst fit. This definition of skill is more conservative than the correlation coefficient or the R-squared value; for example, ψ_2 is correlated to ψ_1 with an average correlation coefficient of 0.74 and the linear regression model has the R-squared of about 0.55 but the skill is only 0.33 if defined as in Equation 4 above. It is thus important to compare the results from different publications using consistent metrics. Here we use the skill metric that is based on the RMS-error normalized by the standard deviation (Equation 4) and, for consistency, we use the Mean Square Error (L2 norm) as the loss function for a NN to minimize during training.

Coefficients of filter matrices, along with all other weights and biases involved in the NN architecture are then iteratively optimized using the Adam optimizer (Kingma & Ba, 2015) to minimize the loss function that is the root-mean-square difference between the predicted and true SSH images (or equivalently to maximize the skill). The parameter optimization procedure requires evaluating NN predictions for a large volume of training data and hence the final optimized state of a particular NN depends only on the training data itself. To ensure that no overfitting has occurred, the NN skill is evaluated for a group of three independent data sets: training, validation, and testing. The training data are used only for the training purposes, the validation data are used to evaluate the skill of the NN and to identify a stoppage criterion for the training, while the testing data are used at the very last step to define the skill of a trained NN. All three data sets are generated from different numerical simulations to ensure that overfitting did not occur.

2.2. Synthetic Training Data: QG Model

In the absence of high-quality and/or large volumes of data, NNs are likely to overfit the training data and have poor skills when evaluated on the test data. To avoid these issues, we choose to train NNs on synthetic data generated using an idealized model of ocean turbulence—the two-layer QG model (Phillips, 1951; Vallis, 2017). The QG model is pertinent to baroclinically unstable flow and contains the propagation dynamics of large-scale ocean eddies, including advection by the mean flow, the beta drift, and the eddy interactions with the mean flow. Our choice of using the two-layer model is rationalized because (i) ocean currents are predominantly composed of the barotropic and the first baroclinic mode (Smith & Vallis, 2001; Wunsch, 1997) and (ii) it is the minimal model demonstrating the difficulty of predicting SSH evolution without direct observations of subsurface flows because both layers are necessarily dynamically active during baroclinic instabilities, and (iii) the dynamical interpolation method also relies on QG dynamics, allowing us to make a straight-forward performance comparison.

The QG model relies on the conservation of potential vorticity and simulates the mesoscale turbulence driven by baroclinic instabilities associated with the vertical shear of the mean flow, requiring a minimum

of two vertically stacked shallow layers. The conservation laws for the top and bottom layer potential vorticities, $q_{1,2}$, are written in the following way:

$$\frac{Dq_1}{Dt} = \frac{D}{Dt}[\nabla^2\psi_1 - \frac{f_0^2}{g'H_1}(\psi_1 - \psi_2) + \beta y] = 0 \quad (5)$$

$$\frac{Dq_2}{Dt} = \frac{D}{Dt}[\nabla^2\psi_2 - \frac{f_0^2}{g'H_2}(\psi_2 - \psi_1) + \beta y] = -r_{Ek}\nabla^2\psi_2, \quad (6)$$

where $\psi_{1,2}$ is the top and bottom layer streamfunctions, f_0 is the Coriolis parameter and β is its derivative in the meridional y-direction, g' is the reduced gravity, $D/Dt = \partial/\partial t + \mathbf{u}\nabla$ is the material derivative using corresponding layer' geostrophic velocity u, and r_{Ek} is the bottom drag coefficient. The relative importance of the discarded term in the PV-conservation budget in Equation 1, $D\psi_{bt}/Dt$, could be estimated by comparing its magnitude to $D\psi_1/Dt$, where both material derivatives use the velocity in the top layer. The ratio of these terms would scale roughly as the ratio of the characteristic amplitudes of the barotropic and surface streamfunctions, which we find from numerical simulations to scale as the ratio of layer depths in QG simulations of the baroclinic instabilities, that is, $[\overline{\psi}^2_{b.t.} / \overline{\psi}^2_1]^{\frac{1}{2}} \sim O(H_1 / H_2)$. Since in most ocean regions, the pycnocline is relatively shallow compared to the full depth of the ocean, the flows are surface-amplified and the discarded term is relatively small but nonnegligible and can substantially impact the SSH evolution leading to significant errors of the dynamical interpolation (see Section 3).

The QG model has been configured to represent baroclinically unstable currents that characterize the Gulf Stream mesoscale turbulence. Model parameters are as follows: the Rossby deformation radius is 40 km, the ratio of mean layer depths is 0.2, there is a steady uniform mean vertical shear of 0.2 m/s, beta plain parameter corresponds to a latitude of 40 degrees, linear Ekman friction was prescribed in the bottom layer for dissipation, and high-wavenumber motions are being filtered in Fourier space for all variables (more details could be found in [Arbic et al., 2012; Flierl, 1978]). The model domain is 1,000 km by 1,000 km and periodic boundary conditions are used. We have explored various resolutions and find that it is sufficient to use a relatively coarse grid of 32×32 to simulate baroclinic instabilities and the chaotic evolution of relatively large mesoscale eddies. While the QG model is aimed at simulating the low-dimensional evolution of the ocean, reconstructing our simulated SSH variability with a significant skill nonetheless requires over 30 principle component modes (Empirical Orthogonal Functions), implying the existence of a vast range of eddy field configurations. The QG model is integrated forward in time managing an ensemble of noisy initial conditions to produce a large volume of data: about 200,000 SSH snapshots separated by 10 days (Figure 2). Over a timescale of 20 days, the correlation between SSH fields drops to about 0.4 and it is hard to identify any persisting eddies because their shapes and intensities have dramatically changed due to interactions with other eddies (Figure 2).

We ensure that the data for training, validation, and testing come from distinct simulations to accurately access the NN generalization skill. Nonetheless, it is plausible that the use of multiple model simulations still does not lead to independent data and that analogous states might exist between the testing and training data sets. To quantify the similarity between the testing and training data, we compared each testing sample to all samples in the training data set (200K samples) and identified the training image that is closest to the testing one using *skill* (Equation 4) as a quantitative metric of similarity. For our testing and training sets that come from separate QG model simulations starting with different initial conditions, the similarity skill does not exceed about 0.2 and its average value is about 0.1 (Figure 3, gray curve), dramatically lower than the prediction skill achieved via NN interpolation (see Section 3). Given such low similarity skills, the existence of analog states in our training/testing data sets can be ruled out. The lack of analogs between the different QG simulations is expected as considering just a handful of the its principle component modes (out of about 30 modes that have significant energy) as independent degrees of freedom would already amount to astronomically low probabilities of encountering true analogs (Van den Dool, 1994).

To evaluate the efficacy of the NN approach, we consider the tasks of (i) temporal interpolation where the input consists of two SSH snapshots separated by 20 days, (ii) spatiotemporal interpolation with the same input as for the temporal interpolation but with SSH images having missing data, and (iii) the state estimation of unobserved deep ocean flows from SSH snapshots. For the temporal separation of SSH images, we

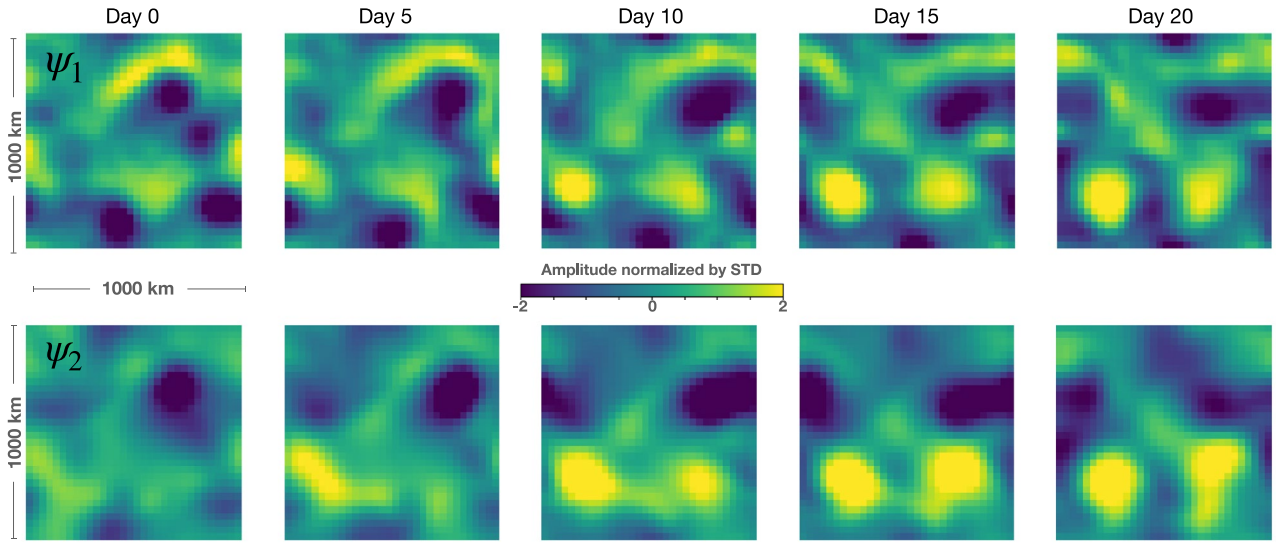

Figure 2. An example of the eddy field evolution over 20 days as generated by the QG model of a baroclinically unstable current. Top panels show surface streamfunction ψ_1 (or SSH) and bottom panels show the corresponding deep ocean streamfunction, ψ_2, both being normalized by their respective standard deviations; the domain size is 1000 × 1000 km and rows correspond to streamfunction snapshots taken 5 days apart. Note that the eddy field dramatically changes over 20 days (SSH decorrelation time scale is about 10–20 days), implying that conventional linear or optimal interpolation methods would lead to significant errors if available observations are separated by more than the decorrelation timescale. QG, quasigeostrophic; SSH, sea surface height.

choose 20 days because it is of the order of the return periods for existing altimeters and to be consistent with (Ubelmann et al., 2015), and we explore how the skill varies with increasing this timescale to 40 and 60 days (Table 1). For the spatiotemporal interpolation, we choose the area of missing data to roughly correspond to that of the SWOT observations over its return period. For a 1,000 km domain, SWOT would have about four crossings (each having a swath of 120 km) with one inclination angle and another four with an opposite angle (see e.g., Figure 1 in Gaultier et al., 2016). While SWOT would have missing data areas in the shape of a rhombus, here for simplicity we have prescribed square shapes as there is no reason to assume this would lose generality.

2.3. Dynamical Interpolation

We reproduce the dynamical interpolation methodology as outlined in (Ubelmann et al., 2015) and evaluate its skill distribution. The method consists of initializing the surface streamfunction $\psi_1=(g/f_0)$SSH and integrating a single-layer quasigeostrophic equation, that is, Equation 3. The domain size, boundary conditions, stratification parameters, and all other parameters of the single-layer model are consistent with those of the two-layer model that was used to generate the validation data. The model integration is performed for 10 days forward in time starting from the SSH snapshot on day 0 and also backward in time starting from the SSH snapshot on day 20. The backward in time integration is performed by reversing the direction of the velocity field and changing the time variable to be negative. The estimate of the SSH field on day 10 is then taken to be the arithmetic mean between the SSH fields resulting from the forward and the backward integration. The skill of the dynamical interpolation is evaluated on the testing data from the two-layer QG model and used for comparison with linear and deep learning interpolation.

3. Results

We have explored various NN architectures for the task of temporal SSH interpolation, ranging from single hidden layer networks (FC) to convolutional networks (VGG), to a more complex residual neural networks (ResNet)—all achieving skills comparable to or higher than the linear and dynamical interpolation methods (Table 1). Substantially decreasing NN complexity leads to an only slight decrease in the skill (e.g., compare experiment pairs [1, 2] or [6,10] in Table 1), while substantially increasing the complexity does not

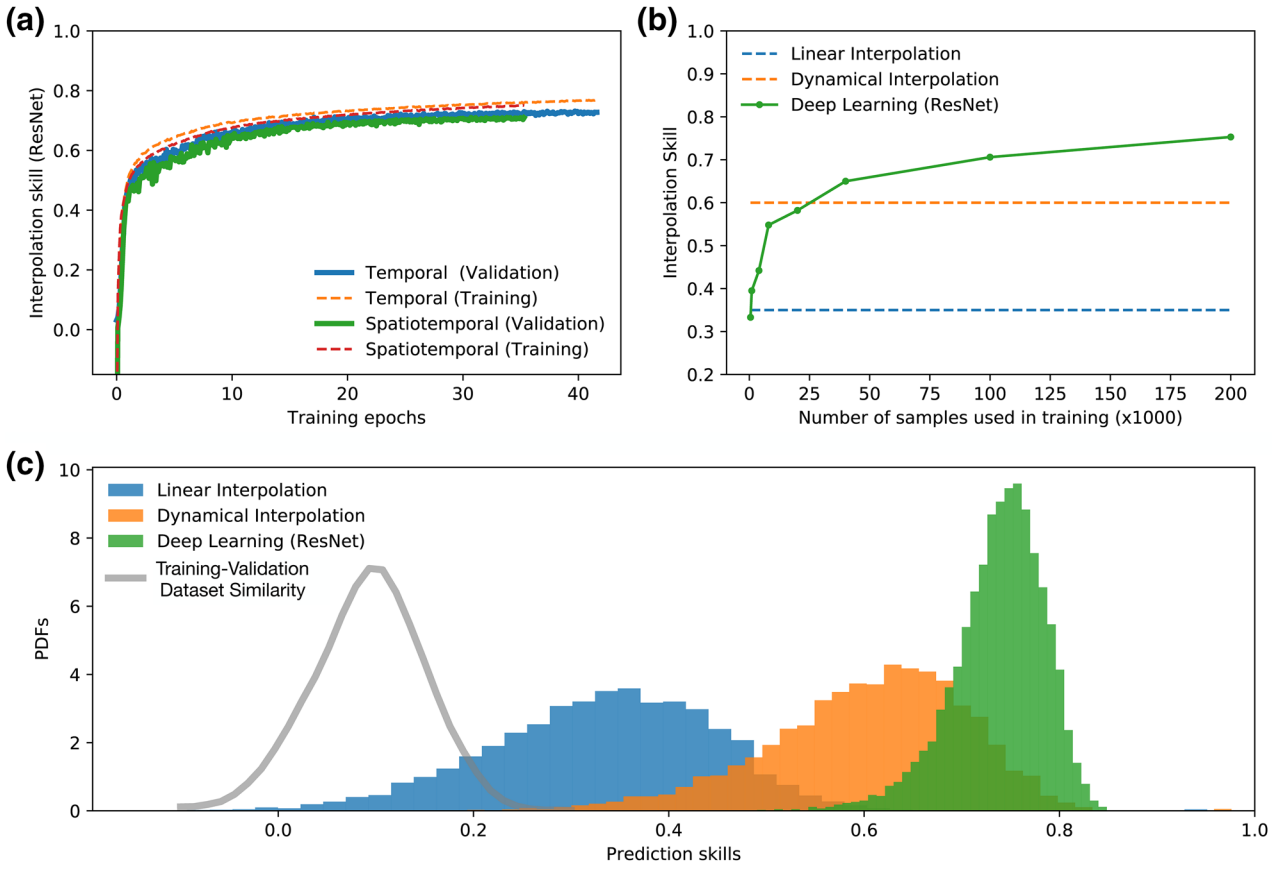

Figure 3. Performance comparison of the deep learning NN (ResNet) with linear and dynamical interpolation techniques. (a) The evolution of the ResNet model validation and training skill during its training on temporal and spatiotemporal SSH interpolation (b) The dependence of the ResNet skill on the number of data samples used in training for the temporal SSH interpolation. (c) Comparison of skill distributions of the linear interpolation (LI), dynamical interpolation (DI), and the deep learning method evaluated on the testing data set. Gray curve shows the similarity between the validation and training data sets, quantified as a probability density function of the skill for comparing a validation input (SSH pair) to the most alike input from the training data set. NN, neural network; SSH, sea surface height.

significantly improve the skill (e.g., compare experiment pairs [3,4] or [7,10] in Table 1). The highest skill of 0.75 is achieved by the ResNet architecture (Figure 1) with a total of 52 convolutional layers and about five million adjustable parameters, taking about 1 h to train on a Tesla T4 GPU on 200K data samples. We thus find the ResNet architecture to be optimal for our tasks and we use it throughout the paper to present our deep learning results, although we note that other superior architectures may exist. Below we use ResNet to demonstrate the deep learning skill in spatiotemporal SSH interpolation and state estimation.

3.1. Spatiotemporal SSH Interpolation

Upon training separate ResNets to perform temporal and spatiotemporal interpolation of SSH data, a significant performance skill is achieved with networks generating realistic SSH images with small errors (see Figure 4). The average prediction skill for both simulations plateaus at about 0.75 and it is not significantly smaller when evaluated on the test data set (Figure 3a). A few illustrative examples of eddy field evolution are shown in Figure 4a, demonstrating the nontrivial SSH evolution that occurs in a chaotic QG model of baroclinically unstable flow. In the top-raw example of Figure 4a, the strong positive SSH anomaly in the center of the domain almost completely disappears after 20 days, yet the NN is still capable to reconstruct the SSH state at day 10. For such examples when the eddy field changes dramatically with time, linear or objective interpolation techniques perform poorly as they do not rely on any dynamical model of SSH evolution and only make use of autocorrelation as a statistical model. Evaluated on a large number of testing data (10K samples), the deep learning model outperforms the linear and dynamical interpolation techniques,

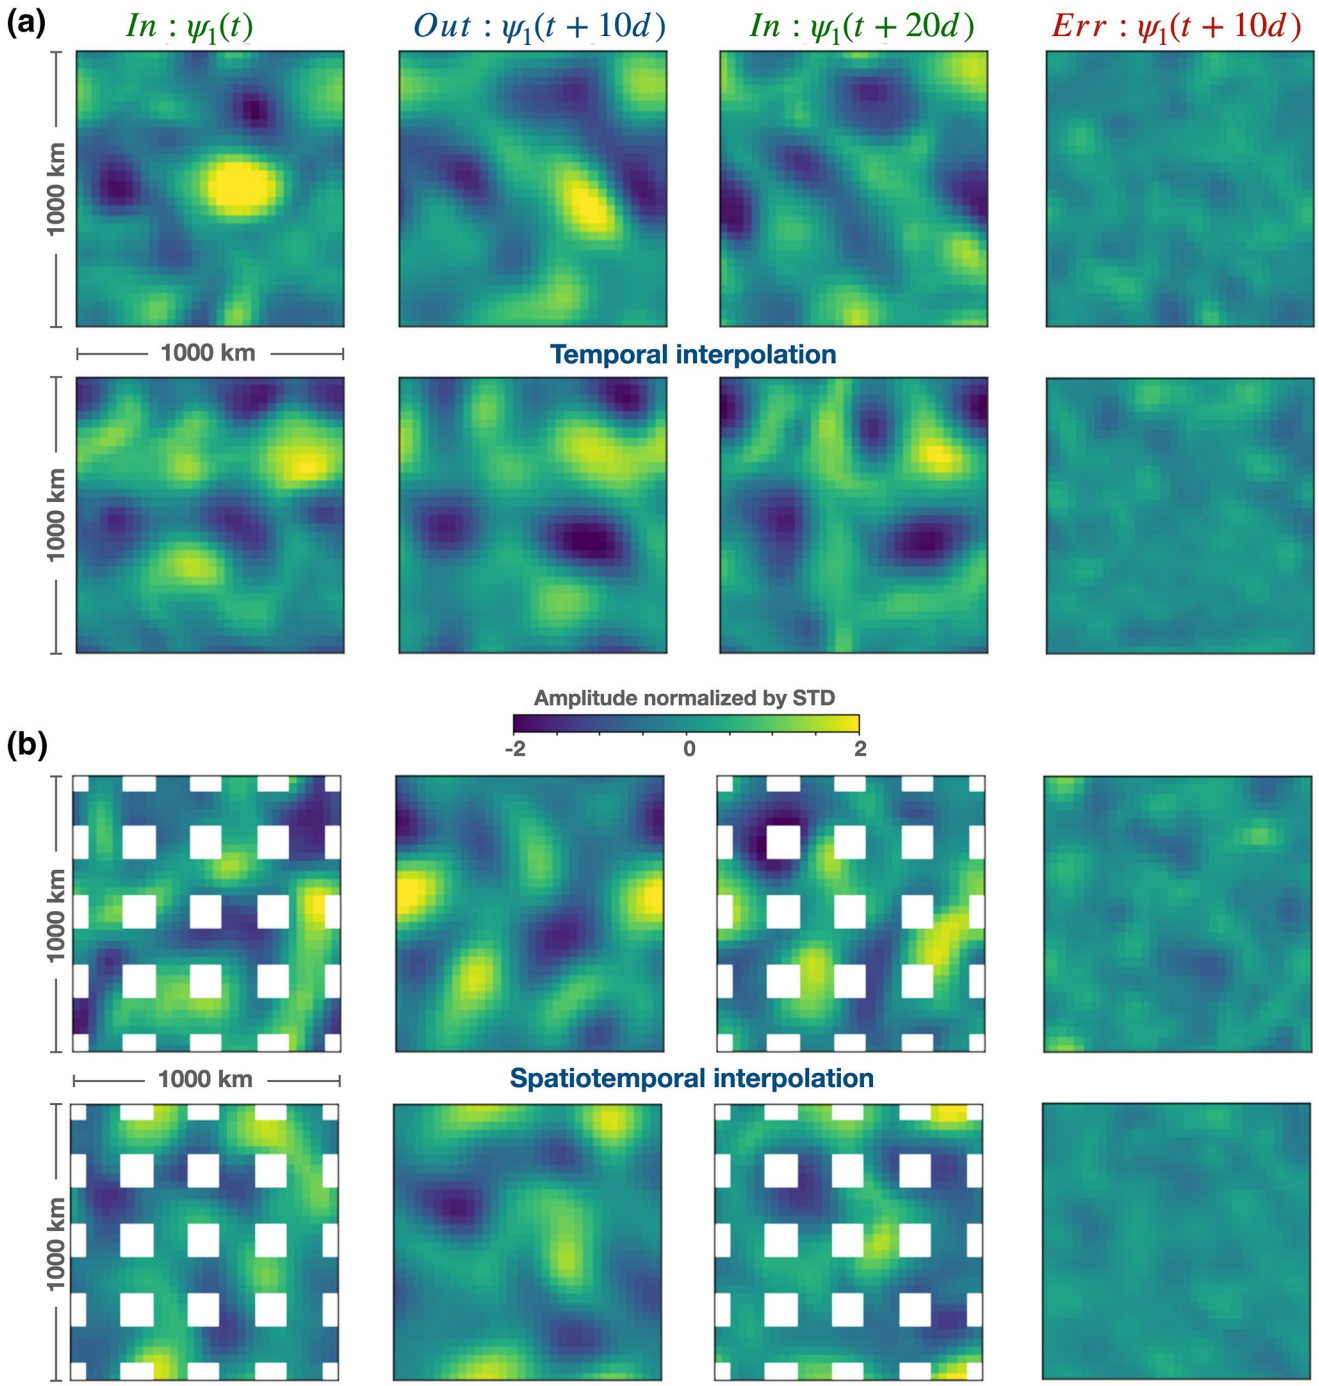

Figure 4. Examples of temporal (a) and spatiotemporal (b) interpolation of SSH data using the deep learning framework. Each row represents a randomly chosen interpolation example from the testing data set. All panels share the same color bar and display streamfunction magnitudes normalized by the standard deviation of the entire data set. The first and third column show panels with input SSH fields $\psi_1(t)$ and $\psi_1\,(t + 20d)$, second column shows the interpolated field $\psi_1\,(t + 10d)$, and the fourth column shows the interpolation error. White regions in the case of spatiotemporal interpolation denote areas of obstructed input data. SSH, sea surface height.

having not only a better average skill but also much more infrequent occurrence of low-skill interpolations, that is, much narrower skill-distribution tail in the direction of small skills (Figure 3c). Noticeably, the linear interpolation skills can be so low as to approach zero and even negative values, that is, its prediction is no better than assuming that SSH = 0 everywhere in the domain. The dynamical interpolation is much

better than the linear interpolation but still has a significant probability of poor interpolations in the skill range of about 0.4–0.6.

While the deep learning technique is superior to other methods, it is important to note that it still does not provide a perfect reconstruction and has a limit in skill bounded by about 0.85 (Figure 3c). The dynamical evolution of the ocean flow considered in our study is inherently chaotic, that is, the phase-space trajectories undergo topological mixing to the extent that the sensitivity to initial conditions increases exponentially with time. Thus, if SSH snapshots of a turbulent eddy field are separated by sufficiently large time (greater than the characteristic Lyapunov exponent timescale), there should be no physical or statistical relationship between these snapshots and hence no interpolation technique could achieve a skill significantly above zero. Indeed, given the same NN architecture and the same volume of training data, the interpolation skill deteriorates dramatically from 0.75 to 0.44 and 0.18 as the time separation between the input SSH snapshots increases from 20 to 40 and 60 days correspondingly (Table 1).

The sensitivity to the number of data samples used in training demonstrates that for the ResNet architecture, about 20–30K data samples are needed to achieve a skill comparable to the dynamical interpolation skill, and using a larger number of training samples leads to a significant skill improvement (Figure 3b). However, the skill continues to increase slowly with the number of samples (Figure 3b), with the best power-law fit for the case of 20-day SSH separation being $skill \sim N^{0.09}$, where N is the number of training samples. Extrapolating the power-law would imply that achieving the perfect $skill = 1$ would require O (10^7) training samples—a number beyond what the author's computing capabilities, though not impossible to reach on modern supercomputers. Nonetheless, estimating the necessary number of samples is only a hypothetical consideration as it is not clear if the power-law would remain the same with the increasing volume of data. In addition, it is not possible to exclude the existence of superior NN architectures that could lead to faster convergence.

3.2. State Estimation of the Unobserved Deep Ocean Flows at Mesoscales

Here we assess the efficacy of the deep learning framework in addressing the state estimation problem, that is, estimating all dynamical variables in the ocean turbulence model, which in our case of a two-layer QG model implies estimating both the surface stream function ψ_1 (or equivalently SSH) and the deep ocean streamfunction ψ_2. Conventionally, for state estimation, one needs to postulate the dynamical model and only then implement the techniques, for example, variational data assimilation or the ensemble Kalman filter techniques to estimate the unknown variables and parameters in the model at all times and everywhere within the model domain. However, we demonstrate here that the deep learning framework can provide an alternative to conventional data assimilation methods. The NN is capable of skillful reconstruction of $\tilde{\psi}_2$ based on two SSH snapshots separated by 20 days, with an average skill of 0.7 for day 0 and a skill of 0.8 for day 20 (Figure 5). While the NN provides skillful predictions for all state variables with skills ranging from 0.65 to 0.85, the best prediction skill is achieved for the deep flow at day 20 while the worst prediction is for deep flow at day 0 (compare orange and red curves in Figure 5c). This temporal asymmetry is expected in chaotic and dissipative QG dynamics, making it more difficult to estimate the past state by observing the future as opposed to estimating the future by observing the past. Thus, the two SSH snapshots must indeed be ordered in time as the PV-evolution equations allow time reversal only for sufficiently small time intervals at which the dissipation effects can be neglected.

It is important to note that only the component of ψ_2 that is uncorrelated with ψ_1 can affect the SSH evolution because the tendency due to the advection of the surface streamfunction by the surface flow is identically zero (see Equation 5). However, ψ_2 is highly correlated with ψ_1, with an average correlation coefficient is about 0.84, which is why reconstructing its full amplitude is a relatively trivial exercise. To evaluate the network ability to predict the decorrelated component, we define it as $\tilde{\psi}_2 = \psi_2 - A\psi_1$, where the constant A is the average linear regression coefficient between ψ_1 and ψ_2. Indeed, using two SSH snapshots as the input, the NN does provide a skillful estimate of $\tilde{\psi}_2$ with a relatively small error (Figure 5a). However, further exploring the limits of NNs, we identify that they are capable of reconstructing an instantaneous relation between the SSH field and deep ocean streamfunction. We train the ResNet model using a single SSH

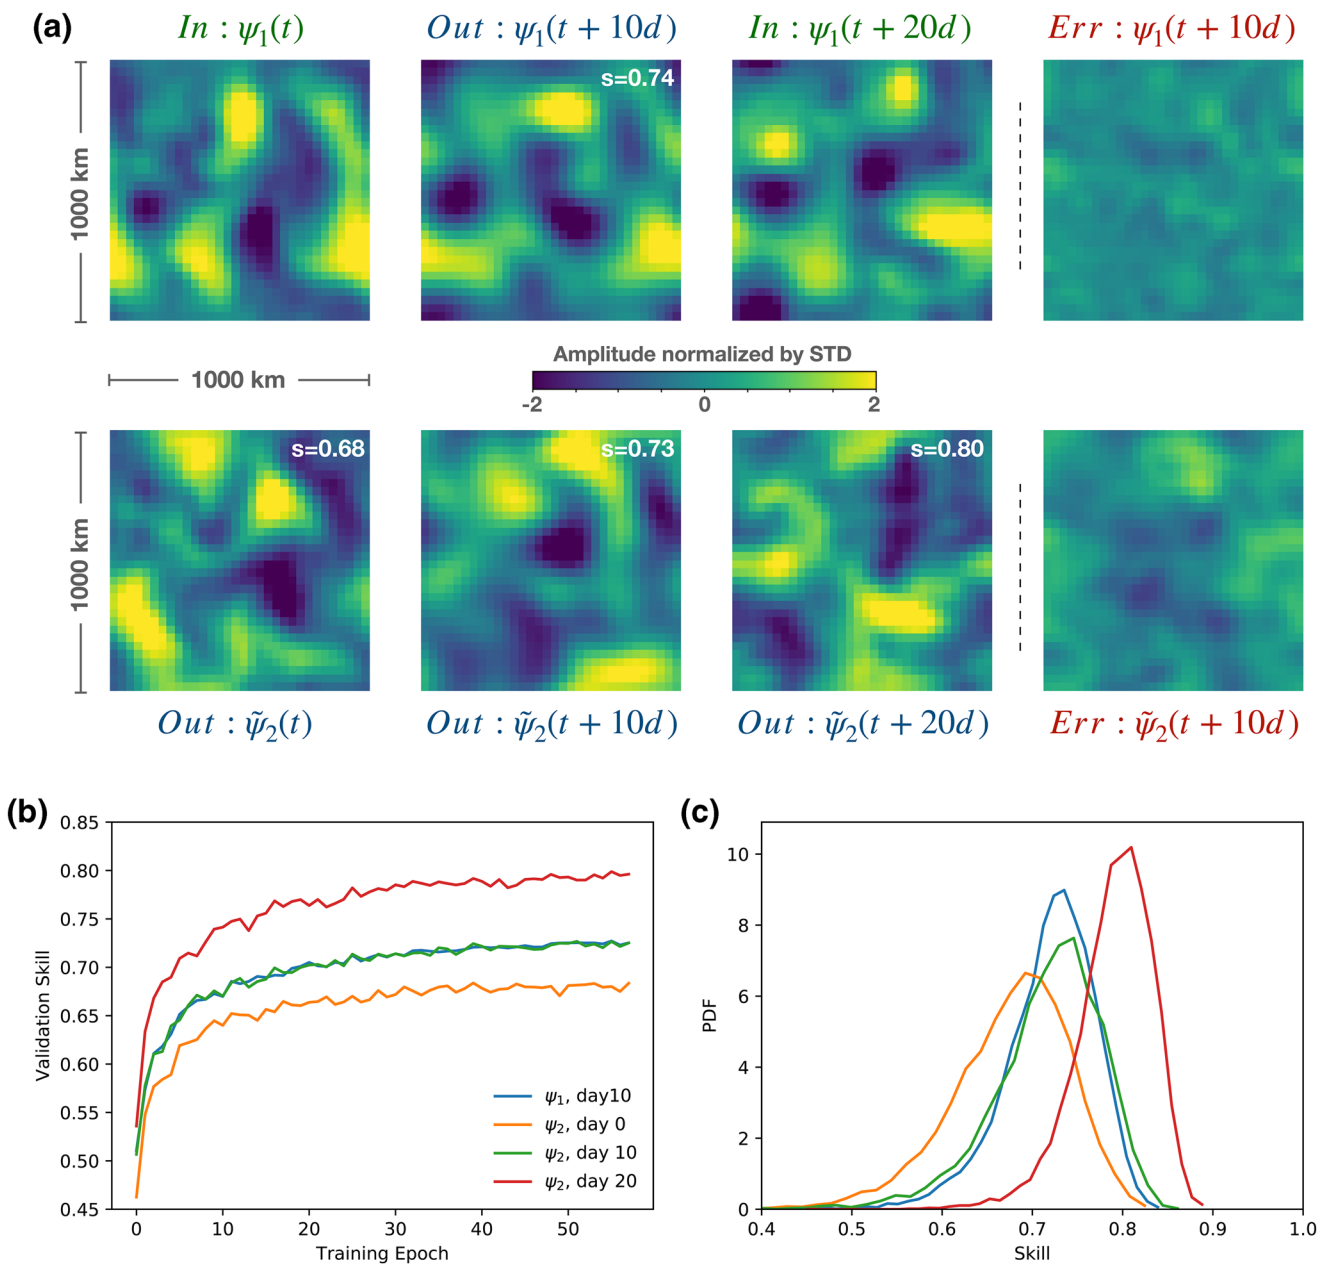

Figure 5. Examples of state estimation using deep learning (a) and its statistical skill distribution for surface and subsurface variables at different times (b). As in the case of SSH interpolation, the NN receives as input two SSH snapshots separated by 20 days, $\psi_1(t)$ and $\psi_1(t + 20d)$ (top row, first and third columns), but reconstructs not only the surface streamfunction at the intermediate time, $\psi_1(t + 10d)$ (top row, second column), but also the subsurface flow at all three times: t, $t + 10d$, and $t + 20d$. Note that ψ_1 and ψ_2 are linearly correlated with a correlation coefficient of 0.8, which is why the bottom rows in panel (a) show $\tilde{\psi}_2$, the component of the reconstructed deep flow that is not linearly correlated with the surface flow. The errors for reconstructing the day 10 surface and deep streamfunctions are shown in the last column. The probability density function of the NN skill distribution is plotted in panel (b) for all predicted variables. NN, neural network; SSH, sea surface height.

snapshot as the input and the decorrelated component $\tilde{\psi}_2$ of the corresponding deep ocean streamfunction as the output to achieve a prediction skill of 0.56, while a skill of 0.7 is achieved if using ψ_2 as the output.

4. Discussion

Our study explored the efficacy of deep learning in reconstructing the unobserved state variables of the chaotic ocean turbulence. The motivation for addressing the specific problem of SSH interpolation came from the present-day use of relatively rudimentary techniques of reconstructing continuous fields from sparse satellite data. Using synthetic data from an idealized model of baroclinic ocean turbulence, we presented the proof of concept for using deep neural networks as an efficient technique to extract nontrivial information from sparse SSH observations. Specifically, we demonstrated that residual CNNs can reconstruct SSH snapshots at the intermediate time between the 20 days separated observations with an average skill of 0.75, significantly outperforming the commonly used linear interpolation (skill = 0.35) and dynamical interpolation (skill = 0.6) techniques. We also demonstrated that the deep learning technique is flexible enough to address a more general problem of state estimation that includes reconstruction of the unobserved deep ocean streamfunction using only SSH snapshots. Nonetheless, there is an inherent lack of information in SSH-only observations that prevents any interpolation or state estimation methodology from achieving a perfect skill. After all, if SSH snapshots are separated by a sufficiently long time, there should not be any relation between them due to the chaotic nature of baroclinic ocean turbulence. Indeed, the ResNet could only achieve a maximum skill of about 0.85 for interpolation between SSH snapshots separated by 20 days, and the skill dramatically decreased to about 0.2 for the snapshots separated by 60 days. The lack of the perfect interpolation skill suggests the existence of a dynamical barrier associated with the inherent lack of information in SSH data, although it is not possible to deduce this with certainty due to potential deficiencies of the NN architecture and the limited volume of training data.

While it is challenging to interpret the SSH interpolation algorithm that was ultimately learned by the deep NN, its superiority over other methods could be associated with its ability to estimate the unobserved deep currents because they directly affect the SSH evolution (Equation 5). Taking only a surface streamfunction snapshot as the input, we demonstrated that the ResNet can estimate the underlying deep ocean streamfunction with an average skill of 0.7, which is high enough for a skillful estimate of the component of the deep streamfunction that is not linearly correlated with the surface streamfunction. Apart from deep learning, no other methods have been reported in the literature that can skillfully estimate the uncorrelated component of the deep ocean currents at mesoscales. The success of those NN architectures that rely specifically on 2D convolutions for pattern extraction implies that it may be the eddy shapes that contain the information necessary to infer deep ocean currents.

A possible physical interpretation in terms of the eddy shapes could be drawn from considering the ocean dynamics in terms of the barotropic and baroclinic modes that are nonlinearly coupled and continuously exchange energy (Larichev & Held, 1995). The surface streamfunction (or SSH) is simply the weighted sum of the barotropic and baroclinic modes while the lower layer streamfunction is their difference. The key question here is: are instantaneous observations of only surface streamfunction sufficient enough to reconstruct the corresponding barotropic and baroclinic modes? This presents an underconstrained problem as there are two unknown modes while there is only one equation connecting their sum to the SSH field and there are no analytical laws that could be inferred from the QG dynamics to provide any additional constraints on the instantaneous relationship between the modes. Nonetheless, the distinct dynamical evolution of each mode can lead to differences in their characteristic spatial patterns that could be discerned by deep learning algorithms. The baroclinic mode experiences a direct energy cascade and its spatial structures should appear more elliptical or elongated because it is stirred by the barotropic flow, especially at scales of the order of or smaller than the Rossby deformation radius. On the contrary, the barotropic mode experiences an inverse kinetic energy cascade manifested in eddy merging and a tendency toward axisymmetrization (Melander et al., 1987). While the two modes continuously interact by exchanging energy, the barotropic mode ends up strongly dominating the baroclinic mode at large scales and their amplitudes become comparable at scales of the order of the Rossby deformation radius (see Figure 4a in Larichev & Held, 1995). This implies that the barotropic mode should dominate large-scale relatively axisymmetric eddy patterns, the baroclinic mode dominates smaller-scale relatively more elliptical patterns, while both modes are present at the deformation scale. Thus, our tentative rationalization of the deep learning success is that by using convolutional filters, the NNs are effectively extracting SSH patterns at different length scales and classifying them into barotropic and baroclinic modes. After estimating the mode amplitudes based on individual SSH

snapshots and learning from many synthetic examples of SSH evolution in time, the NNs are then capable to effectively integrate the QG equations forward or backward in time for a skillful temporal interpolation between the two SSH snapshots. While the complexity of deep learning algorithms makes it impossible to interpret them, our hypothetical two-step process of the mode decomposition followed by the forward and backward integration provides a plausible dynamical rationalization for the superiority of deep learning over methods that ignore the influence of deep ocean flows on SSH evolution.

We chose to use the QG simulations of baroclinic turbulence as the synthetic training data set because it presents a hard test for the temporal SSH interpolation due to its chaotic nature and an a priori unknown impact of the dynamically active bottom layer on SSH evolution. However, for the case of submesoscale turbulence (length scales smaller than about 100 km), the question remains open as to how SWOT's 2D high-resolution swath measurements could be used to enhance the resolution of SSH data. While we expect the deep learning framework to perform well in reconstructing both large and small mesoscale eddies, its limitations still need to be understood when considering mesoscale and submesoscale turbulence as a continuum. It is thus necessary to develop more general training data sets that are representative of the SSH dynamics for any given region or process of interest. Including satellite observations from Synthetic Aperture Radars or of sea surface temperatures in addition to the SSH observations could provide additional information for improved reconstruction of SSH. The training data sets could be assembled ranging from more realistic submesoscale-resolving general circulation models to simplified stochastic models in various parameter regimes (Samelson et al., 2019). While diversifying the training data sets should increase the versatility of NN interpolation methods, the crucial constraint of their performance would likely come from the chaotic evolution of submesoscale eddies that occurs on substantially shorter timescales compared to mesoscale eddies.

While we have demonstrated the efficacy of supervised deep learning using synthetic data, its practical utility in interpolating real-world SSH observations remains to be tested. The drawback of deep learning is that it requires a large volume of training data, although there are continuously improving methods aimed at addressing this practical issue, for example, transfer learning (Pan & Yang, 2009), data augmentation (Perez & Wang, 2017), one-shot learning (Fei-Fei et al., 2006). A way toward ultimately developing the gridded SSH product using deep learning could be through training networks on a wide range of idealized and realistic models and then fine-tuning a much smaller number of NN parameters using existing satellite data. However, since the true two-dimensional SSH state is not known at any particular time, the fine-tuning of a NN cannot be achieved by defining a simple loss function as was done with synthetic data. Thus, the NN ultimately would need to use a loss function that is based purely on observations, without invoking a dynamical model to provide a true state. This issue could be addressed for example using reinforcement learning, where two-dimensional SSH fields generated by the NN would be rewarded or penalized based on the accuracy of their projection on the observed altimetry tracks that were left out from the input set of tracks. Developing deep learning SSH interpolation techniques that would steer away from solely relying on dynamical models to provide training data is a necessary next step toward practical implementation with real satellite observations. Nonetheless, our work presents an important proof of concept that SSH observations do contain dynamically relevant information about subsurface flows, and hence with deep learning it should be possible to build a skillful model of SSH evolution and as a consequence improve the existing SSH estimates.

Finally, we note another potentially important application of deep learning for state estimation at eddy-resolving scales. Since mesoscale-resolving data assimilation methods require large computations, providing an accurate initial guess would substantially reduce the number of iterations necessary for optimization. Thus, it might be possible to accelerate data assimilation methods by providing a deep learning estimate as a first guess that is already close to reality. Note that data assimilation and NNs are similar approaches in that they both use iterative procedures to find the optimal set of unknown parameters to minimize the error between the predicted and true fields. The critical difference is that data assimilation methods are based on a concrete physical model or its linearization, and hence the predicted fields conform to the desired physical constraints but the reconstruction skill relies on the accuracy of the model. Contrarily, the deep learning approach does not rely on a physical model as it is optimizing a complex nonlinear mapping function that is general enough to map the input to the output. Hence, the deep learning predictions do not have to obey

any dynamical constraints unless those have been explicitly incorporated in the loss function. Thus, we see the synergy between deep learning and conventional state estimation methods as a potential framework for constructing improved state estimates, combining the best of the two paradigms: fast data-driven state estimation via deep learning and fine-tuning by conventional data assimilation methods to ensure the strict consistency with an assumed dynamical model.

Data Availability Statement

The neural network architectures coded in Tensorflow/Keras and the training data sets are published in the following Zenodo repository: https://doi.org/10.5281/zenodo.3757524.

Acknowledgments
Georgy E. Manucharyan was funded by the ONR Grant N00014-19-1-2421 at the University of Washington and thanks Charles Trimble for providing support at the California Institute of Technology. Lia Siegelman was a NASA-JVSRP affiliate and supported by a Caltech-JPL postdoctoral fellowship. Patrice Klein was supported by the NASA NPP Senior Fellowship and the NASA SWOT project. Discussions with Dimitris Giannakis, Andrew Thompson, Clement Ubelmann, Laure Zanna, and participants of the 2019 SWOT Science Team Meeting are highly appreciated. The authors thank Glenn Flierl for providing the Matlab code for simulating the two-layer QG turbulence and Michael Dietz for the open-source implementation of ResNet, modifications of which were used in this study. The authors thank the two anonymous reviewers and the associate editor, Annalisa Bracco, for insightful comments that led to improvements of the manuscript.

References

Abernathey, R. P., & Marshall, J. (2013). Global surface eddy diffusivities derived from satellite altimetry. *Journal of Geophysical Research: Oceans*, *118*(2), 901–916. https://doi.org/10.1002/jgrc.20066

Aluie, H., Hecht, M., & Vallis, G. K. (2018). Mapping the energy cascade in the north Atlantic Ocean: The coarse-graining approach. *Journal of Physical Oceanography*, *48*(2), 225–244.

Arbic, B. K., Scott, R. B., Flierl, G. R., Morten, A. J., Richman, J. G., & Shriver, J. F. (2012). Nonlinear cascades of surface oceanic geostrophic kinetic energy in the frequency domain. *Journal of Physical Oceanography*, *42*(9), 1577–1600.

Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., et al. (2019). On the resolutions of ocean altimetry maps. *Ocean Science*, *15*(4), 1091–1109.

Berloff, P. S., & Meacham, S. P. (1997). The dynamics of an equivalent-barotropic model of the wind-driven circulation. *Journal of Marine Research*, *55*(3), 407–451.

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.

Carton, J. A., & Giese, B. S. (2008). A reanalysis of ocean climate using simple ocean data assimilation (soda). *Monthly Weather Review*, *136*(8), 2999–3017.

Chapman, C., & Charantonis, A. A. (2017). Reconstruction of subsurface velocities from satellite observations using iterative self-organizing maps. *IEEE Geoscience and Remote Sensing Letters*, *14*(5), 617–620.

Charney, J. G. (1971). Geostrophic turbulence. *Journal of the Atmospheric Sciences*, *28*(6), 1087–1095.

Chelton, D. B., & Schlax, M. G. (2003). The accuracies of smoothed sea surface height fields constructed from tandem satellite altimeter datasets. *Journal of Atmospheric and Oceanic Technology*, *20*(9), 1276–1302.

Chelton, D. B., Schlax, M. G., & Samelson, R. M. (2011). Global observations of nonlinear mesoscale eddies. *Progress in Oceanography*, *91*(2), 167–216.

Davis, R. E. (1985). Objective mapping by least squares fitting. *Journal of Geophysical Research: Oceans*, *90*(C3), 4773–4777.

de La Lama, M. S., LaCasce, J. H., & Fuhr, H. K. (2016). The vertical structure of ocean eddies. *Dynamics and Statistics of the Climate System*, *09*(1), dzw001. https://doi.org/10.1093/climsys/dzw001.1

Ducet, N., Le Traon, P.-Y., & Reverdin, G. (2000). Global high-resolution mapping of ocean circulation from topex/poseidon and ers-1 and-2. *Journal of Geophysical Research: Oceans*, *105*(C8), 19477–19498.

Fei-Fei, L., Fergus, R., & Perona, P. (2006). One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *28*(4), 594–611.

Ferrari, R., & Wunsch, C. (2009). Ocean circulation kinetic energy: Reservoirs, sources, and sinks. *Annual Review of Fluid Mechanics*, *41*, 253–282.

Flierl, G. R. (1978). Models of vertical structure and the calibration of two-layer models. *Dynamics of Atmospheres and Oceans*, *2*(4), 341–381.

Fu, L.-L., Chelton, D. B., Le Traon, P.-Y., & Morrow, R. (2010). Eddy dynamics from satellite altimetry. *Oceanography*, *23*(4), 14–25.

Fu, L.-L., & Ubelmann, C. (2014). On the transition from profile altimeter to swath altimeter for observing global ocean surface topography. *Journal of Atmospheric and Oceanic Technology*, *31*(2), 560–568.

Gaultier, L., Ubelmann, C., & Fu, L.-L. (2016). The challenge of using future swot data for oceanic field reconstruction. *Journal of Atmospheric and Oceanic Technology*, *33*(1), 119–126.

George, T. M., Manucharyan, G. E., & Thompson, A. F. (2021). Deep learning to infer eddy heat fluxes from sea surface height patterns of mesoscale turbulence. *Nat Commun*. https://doi.org/10.1038/s41467-020-20779-9

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition (770–778). *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE.

Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y., Bengio, & Y., LeCun (Eds.), *Procedding of the 3rd International Conference on Learning Representations, {ICLR} 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. Preprint arXiv:1412.6980. Retrieved from http://arxiv.org/abs/1412.6980

Klein, P., Isern-Fontanet, J., Lapeyre, G., Roullet, G., Danioux, E., Chapron, B., et al. (2009). Diagnosis of vertical velocities in the upper ocean from high resolution sea surface height. *Geophysical Research Letters*, *36*(12), https://doi.org/10.1029/2009gl038359

Klein, P., Lapeyre, G., Siegelman, L., Qiu, B., Fu, L.-L., Torres, H., et al. (2019). Ocean-scale interactions from space. *Earth and Space Science*, *6*(5), 795–817.

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.

Le Traon, P., Nadal, F., & Ducet, N. (1998). An improved mapping method of multisatellite altimeter data. *Journal of Atmospheric and Oceanic Technology*, *15*(2), 522–534.

Manucharyan, G. (2020). *State estimation of surface and deep flows from sparse SSH observations of geostrophic ocean turbulence using Deep Learning*: Zenodo. https://doi.org/10.5281/zenodo.3757524

Melander, M., McWilliams, J., & Zabusky, N. (1987). Axisymmetrization and vorticity-gradient intensification of an isolated two-dimensional vortex through filamentation. *Journal of Fluid Mechanics*, *178*, 137–159.

Morrow, R., Fu, L.-L., Ardhuin, F., Benkiran, M., Chapron, B., Cosme, E., et al. (2019). Global observations of fine-scale ocean surface topography with the surface water and ocean topography (swot) mission. *Frontiers in Marine Science*, *6*, 232.

Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, *22*(10), 1345–1359.

Perez, L., & Wang, J. (2017). The effectiveness of data augmentation in image classification using deep learning, *CoRR*, *11*. Preprint arXiv:1712.04621. Retrieved from http://arxiv.org/abs/1712.04621

Phillips, N. A. (1951). A simple three-dimensional model for the study of large-scale extratropical flow patterns. *Journal of Meteorology*, *8*(6), 381–394.

Samelson, R., Chelton, D., & Schlax, M. (2019). The ocean mesoscale regime of the reduced-gravity quasi-geostrophic model. *Journal of Physical Oceanography*. *49*(10), 2469–2498.

Scott, R. B., & Arbic, B. K. (2007). Spectral energy fluxes in geostrophic turbulence: Implications for ocean energetics. *Journal of Physical Oceanography*, *37*(3), 673–688.

Smith, K. S., & Vallis, G. K. (2001). The scales and equilibration of midocean eddies: Freely evolving flow. *Journal of Physical Oceanography*, *31*(2), 554–571.

Targ, S., Almeida, D., & Lyman, K. (2016). Resnet in resnet: Generalizing residual architectures. *CORR*. Preprint arXiv:1603.08029. Retrieved from http://arxiv.org/abs/1603.08029

Tulloch, R., Marshall, J., Hill, C., & Smith, K. S. (2011). Scales, growth rates, and spectral fluxes of baroclinic instability in the ocean. *Journal of Physical Oceanography*, *41*(6), 1057–1076.

Ubelmann, C., Cornuelle, B., & Fu, L.-L. (2016). Dynamic mapping of along-track ocean altimetry: Method and performance from observing system simulation experiments. *Journal of Atmospheric and Oceanic Technology*, *33*(8), 1691–1699.

Ubelmann, C., Klein, P., & Fu, L.-L. (2015). Dynamic interpolation of sea surface height and potential applications for future high-resolution altimetry mapping. *Journal of Atmospheric and Oceanic Technology*, *32*(1), 177–184.

Vallis, G. K. (2017). *Atmospheric and oceanic fluid dynamics*, (745). Cambridge, UK: Cambridge University Press.

Van den Dool, H. (1994). Searching for analogues, how long must we wait? *Tellus A*, *46*(3), 314–324.

Wunsch, C. (1997). The vertical partition of oceanic horizontal kinetic energy. *Journal of Physical Oceanography*, *27*(8), 1770–1794.

Wunsch, C. (2010). Toward a midlatitude ocean frequency–wavenumber spectral density and trend determination. *Journal of Physical Oceanography*, *40*(10), 2264–2281.