State estimation of surface and deep flows from sparse SSH observations of geostrophic ocean turbulence using Deep Learning

Georgy E Manucharyan [1]**, Lia Siegelman** [2]**, Patrice Klein** [2]

[1]School of Oceanography, University of Washington, Seattle, WA, USA
[2]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA

Key Points:

- Deep Learning framework is developed for SSH interpolation in a baroclinically unstable current.

- Residual Neural Networks outperform linear and dynamical SSH interpolation techniques.

- Skillful estimation of unobserved deep flows from temporally-sparse SSH observations is plausible

Corresponding author: Georgy Manucharyan, `gmanuch@uw.edu`

Abstract

Satellite altimeters provide global observations of sea surface height (SSH) and present a unique dataset for advancing our theoretical understanding of upper ocean dynamics and monitoring its variability. Considering that mesoscale and submesoscale SSH patterns can evolve on timescales comparable to or shorter than satellite return periods, currently available altimetry observations are still spatially and temporally sparse and hence it is challenging to accurately reconstruct continuous SSH evolution. Here we explore the possibility of SSH interpolation using Deep Learning — a machine learning approach that extracts information only from data. Using synthetic observations taken from an idealized quasi-geostrophic model of baroclinic ocean turbulence, we demonstrate that Convolutional Neural Networks with Residual Learning are superior in SSH reconstruction than the linear and recently developed dynamical interpolation techniques. Furthermore, the neural network can provide an accurate state estimate of unobserved deep ocean currents at mesoscales, suggesting that SSH patterns of eddies do contain substantial information about ocean interior that is necessary for SSH prediction. Our framework is highly idealized and several crucial improvements such as transfer learning and diversification of training data would be necessary to implement before its ultimate use with real satellite observations. Nonetheless, by providing a proof of concept, our results point to Deep Learning as a viable alternative to existing interpolation and more generally state estimation methods for satellite observations of baroclinic ocean turbulence.

Plain Language Summary

Satellite sea surface height (SSH) observations provide critical insights into the variability of ocean currents. However, these observations are spatially and temporally sparse, presenting a challenge for reconstructing time-continuous SSH maps particularly at resolutions containing relatively fast-evolving upper-ocean eddies. Further limitations are due to the fact that the evolution of SSH is not self-constrained as it is affected by unobserved deep ocean flows. In this study, we test a different approach to address poor temporal sampling of SSH: a machine learning framework that relies on pattern recognition in large-scale ocean turbulence. We demonstrate that deep artificial neural networks can generate a skillful state estimation of unobserved deep ocean currents and outperform conventional SSH reconstruction methods in an idealized model of ocean turbulence. In providing the proof of concept, our results strongly point at Deep Learning learning as a viable alternative to existing interpolation and state estimation methods for satellite oceanography.

1 Introduction

Satellite-derived global observations of sea surface height (SSH) has shed light on many dynamical processes including large-scale circulation, propagation of waves as well as on the evolution of the mesoscale eddy field (Chelton et al., 2011; Fu et al., 2010). Since the satellite era, an increasing amount of evidence points towards the mesoscale eddies being a key component of the global ocean circulation and significantly impact, among others, carbon sequestration, biological productivity, heat transport and thus the Earth's climate as a whole (Ferrari & Wunsch, 2009). Nonetheless, understanding and monitoring oceanic energy spectrum and associated spectral energy fluxes (Scott & Arbic, 2007; Aluie et al., 2018), understanding tracer dispersion (Abernathey & Marshall, 2013) or inferring subsurface flows (Klein et al., 2009) still remains challenging because these quantities depend on higher-order SSH derivatives and hence require high resolution and accuracy. The regularly-gridded SSH data, e.g. AVISO (Ducet et al., 2000), is spatially and temporally interpolated from along-track altimetry measurement using objective mapping and hence its accuracy is constrained by the density of observations and by the deficiencies of the interpolation technique. To provide better coverage, 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 (Wunsch, 2010).

The upcoming Surface Water Ocean Topography (SWOT) altimeter mission (Fu & Ubelmann, 2014) promises to observe ocean mesoscale eddies and submesoscale fronts ($\leq$ 50 km) at unprecedented spatial resolutions, potentially resolving 15-30km wavelengths. However, the temporal resolution of the altimeter (i.e., a complete repeat cycle of 21 days) is not sufficient to continuously capture the evolution of submesoscale eddies and fronts. The mismatch between the high spatial resolution and the moderate temporal resolution presents a challenge for reconstructing time-continuous maps of SSH. This task is especially challenging in eddy-rich regions where small-scale SSH anomalies can evolve relatively fast compare to satellite return periods, e.g. in the Antarctic Circumpolar Current, the Kuroshio Extension and the Gulf Stream, all of which are key players in the climate system. It is thus crucial to develop frameworks to efficiently extract information about oceanic eddy dynamics from the spatially and temporally sparse SSH observations.

1.1 SSH interpolation and associated dynamical limitations

Spatiotemporal interpolation or gridding of SSH data is inherently linked to ocean 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 eddy propagation in space and time. To illustrate this point, imagine a coherent moving eddy 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 to pinpoint only those tracks that have passed over this specific eddy and combine this information to constrain the two-dimensional eddy shape. Thus, without a model of eddy evolution, or more generally SSH evolution, the information from various altimetry tracks could not be used in an optimal way. However, due to the stratified nature of geostrophic ocean turbulence, the unobserved subsurface flows can affect surface dynamics and hence the knowledge of the SSH field may not be self-sufficient to infer its evolution. Given the lack of subsurface information at eddy scales, constructing a reduced self-contained model of SSH evolution is challenging.

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 having their advantages and disadvantages. To avoid prescribing a dynamical model, statistical models relying on data only, e.g. objective interpolation methods (Davis, 1985; Le Traon et al., 1998; Ducet et al., 2000). Their premise is to incorporate spatiotemporal correlations and measurement error into a statistical model that provides the most likely estimate of the true continuous state, given available observations.

However, this method does not rely on any dynamical model of the eddy propagation and hence can lead to 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 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, the drawback of this method is that it requires additional observations to constrain other essential model variables like the subsurface flow and density field. Also, data assimilation for complex ocean models at eddy-resolving scales is often under-determined and is computationally demanding.

Recent work by Ubelmann et al. (2015) demonstrated that representing SSH propagation with a single equivalent barotropic mode in a quasigeostrophic model results in significant improvements in spatiotemporal interpolation of sparse SSH observations. In particular, Ubelmann et al. (2015) considered a fundamental problem of reconstructing 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 assumed dynamical of an equivalent barotropic mode) and averaging it with the later observed SSH anomalies that were integrated backward in time, resulted in improvement compared to conventional linear interpolation methods. In a follow-up work, Ubelmann et al. (2016) generalized this temporal interpolation method to 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 process that is inherently non-linear and thus cannot be accurately represented by linear or objective interpolation techniques that do not take into account the dynamical constraints imposed on ocean flows.

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

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

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

ψ_1 is the surface streamfunction directly proportional to SSH, ψ_2 is the subsurface streamfunction, $\psi_{b.t.}$ denoting the barotropic streamfunction (depth-averaged transport), R_d is the Rossby baroclinic deformation radius, f and β are the Coriolis and beta-plane parameters, 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 Methods). On relatively short timescales, sources and dissipation of potential vorticity could be neglected and its approximate conservation provides a basic description of 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 $\psi_{b.t.}$). It is now clear that 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. Since

the discarded term was the only one that depended 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 upper-ocean currents are dynamically active, reconstructing SSH using the dynamical interpolation technique performed well, being superior to linear interpolation methods because it relied, at least approximately, on the fundamental PV-conservation constraint. Nonetheless, the dynamical interpolation method can lead to significant errors (see Results), 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 Rationale for Deep Learning approach.

A clear way of improving the dynamical interpolation algorithm would be to include 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 lead 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 two layers. However, considering that PV anomalies are specifically due to baroclinic instabilities and they obey specific conservation laws (Eq. 1), the barotropic and baroclinic modes are inherently entangled and this must provide at least partial constraints on how any specific SSH pattern could be partitioned. Since the QG model exhibits highly non-linear 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, mooring observations demonstrate that surface and subsurface flows are significantly correlated such that a single empirical orthogonal function (EOF) can explain a significant amount of variance of the overall vertical velocity profile Wunsch (1997); de La Lama et al. (2016). Furthermore, machine learning techniques such as self-organizing Chapman & Charantonis (2017), as well as convolutional neural networks (Bolton & Zanna, 2019), have been used to estimate subsurface flows from SSH data. However, the unknown term $D\psi_{bt}/Dt = (\partial_t + \mathbf{u_1} \cdot \nabla)\psi_{bt}$ in Eq. 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 $\partial_t \psi_{bt} << \partial_t \psi_1$ 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. (2019, under review) have 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 extract information from SSH patterns to estimate the component of ψ_{bt} that is uncorrelated with ψ_1 when estimating the eddy heat fluxes, here we hypothesize that machine learning techniques could outperform the dynamical interpolation methods.

2 Methods

Here we present a machine learning framework that mimics the task of dynamical interpolation, i.e. reconstructs the SSH snapshot that occurred between two given SSH snapshots separated by 20 days (Ubelmann et al., 2015). We use machine learning as a tool to shortcut the formal process of data assimilation and to establish if there are substantial possible connections between the dynamical evolution of eddies and spatiotemporal interpolation. We are interested in providing a proof-of-concept machine learning framework and understanding its dynamical limitations, i.e. those that are not subject to insufficiency or poor quality of data. We develop and test our method in an idealized framework of predicting SSH snap-
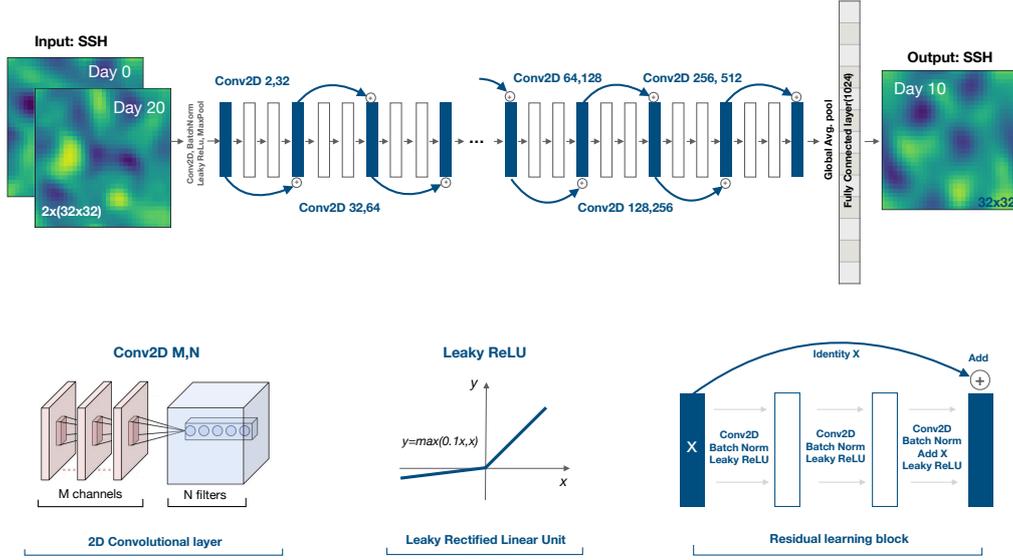
Figure 1. Architecture of the Deep Learning neural network with residual learning (slightly modified from the well-known ResNet50 architecture He et al. (2016)). The input consists of two SSH snapshots separated by 20 days (in case of SSH input with missing data, those values are set to zero). A set of convolutional layers are then applied to create abstract representation of the input patterns in a bottleneck fashion: when image sizes decrease by a factor of two, the number of filters increase 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 information from one layer and adding its identity to the output several layers ahead. The output from the convolutional layer is subject to global average pooling and flattening into a vector that is finally densely connected to the output.

shots that were generated by a quasigeostrophic (QG) model of baroclinically unstable flow. We find the QG model to be optimal for our goals as is pertinent to many energetic regions in the ocean while being relatively simple that a large volume of data can be generated for training and testing; furthermore, the model allows us to directly benchmark machine learning against the dynamical interpolation technique that also utilizes QG dynamics. Below we describe our neural network architecture for spatiotemporal SSH interpolation and the QG model used for the generation of training and testing datasets.

2.1 Deep Learning framework: Residual Convolutional Neural Networks

Artificial neural networks 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 non-linearity 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 eddy patterns expressed in SSH fields, the choice of convolutional neural networks (CNNs) is rationalized. In passing information from layer to layer, CNNs define a set of filters (kernel matrices with prescribed dimensions) and convolve images

to produce more abstract levels of information that is passed on to the next layer. Here we implement the ResNet50 architecture – a Convolutional Neural Network with Residual Learning blocks (He et al., 2016). The Residual Learning is a process by which the information is not only transferred sequentially from one layer to another but is also transferred via skip connections that add the identify of the current layer to the layer that is a few ahead (see Fig. 1); the presence of skip connections can result in better performance for a wide range of computer vision problems (Targ et al., 2016). We note that we have explored several simpler architectures like shallow neural networks with only dense connections, and simpler VGG-type architectures without residual learning but have achieved significantly poorer performance; we thus present the network architecture that have lead to a significant skill, although there is always a possibility that superior neural network architectures may exist. The graph of the architecture used in this study, outlining all hyperparameters together with the Python code of its implementation in Tensorflow/Keras as well as the training datasets can be found here `https://drive.google.com/drive/folders/1tZrpILw2m19CB1YcQABj6pnOa7_-H63n?usp=sharing`.

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:

$$Skill = 1 - \left(\frac{\overline{|SSH_{predicted} - SSH_{true}|^2}}{\overline{|SSH_{true}|^2}} \right)^{\frac{1}{2}}. \tag{3}$$

For reference, the maximum skill=1 is achieved when the predicted and true images are exactly the same; 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 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 R-squared of about 0.55 but the skill of only 0.33 if defined as in equation 3 above. It is thus important to compare results from different publications using consistent metrics. Here we stick with the skill metric that is based on the RMS-error normalized by the standard deviation as it is a natural choice for a neural network loss function to minimize during training.

Coefficients of filter matrices, along with all other weights and biases involved in the neural network architecture are then iteratively optimized using the Adam optimizer (Kingma & Ba, 2014) 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 neural network predictions for a large volume of training data and hence the final optimized state of a particular neural network depends only on the training data itself. To test if a general dependence was found the neural network skill is estimated for a group of three independent datasets: training, validation, and testing sets. Training data is used only for training purposes, validation data is used to evaluate the skill of the neural network and to identify a stoppage criterion for the training, while the test data is used at the very last step to define the skill of a trained neural network. All three datasets were generated from different numerical simulations to ensure that overfitting didn't occur and that a general law was found.

2.2 Synthetic training data: quasigeostrophic model

Deep neural networks typically require a large volume of training data to identify a general law. In the absence of high-quality or sufficiently large volume of data neural networks are likely to overfit the training data and have poor skill when evaluated on the test data. To avoid these issues we choose to train neural network on synthetic data generated using an idealized model of ocean turbulence – the two layer quasigeostrophic (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 mean flow, beta drift, and eddy interactions with mean flow. Our choice of using the
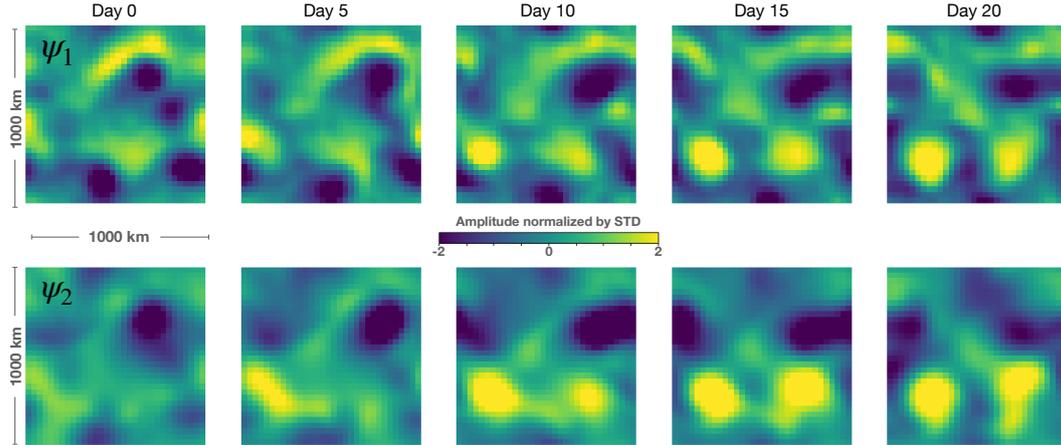
Figure 2. An example of the eddy field evolution over the course of 20 days as generated by the QG model of a baroclinically unstable current. Top panels show surface streamfunction ψ_1 (or SSH) and bottom panels show the corresponding deep ocean streamfunction ,ψ_2, both being normalized by their respective standard deviations; domain size is 1000x1000 km and rows correspond to streamfunction snapshots taken 5 days apart. Note that the eddy field dramatically changes over the course of 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.

two-layer model is rationalized because i) ocean currents are predominantly composed of the barotropic and the first baroclinic mode (Wunsch, 1997; Smith & Vallis, 2001) and ii) it is the minimal model demonstrating the difficulty of predicting SSH evolution without direct observations of subsurface flows because both layers necessarily are dynamically active during baroclinic instabilities, and iii) the dynamical interpolation method also relies on QG dynamics which allows to make a fair performance comparison.

The quasigeostrophic model relies on the conservation of potential vorticity and simulates mesoscale turbulence driven by baroclinic instabilities associated with the vertical shear of 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:

$$\frac{Dq_1}{Dt} = \frac{D}{Dt}[\nabla^2\psi_1 - \frac{f_0^2}{g'H_1}(\psi_1 - \psi_2) + \beta y] = 0 \tag{4}$$

$$\frac{Dq_2}{Dt} = \frac{D}{Dt}[\nabla^2\psi_2 - \frac{f_0^2}{g'H_2}(\psi_2 - \psi_1) + \beta y] = -r_{Ek}\nabla^2\psi_2, \tag{5}$$

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 Eq. 1, $D\psi_{bt}/Dt$, could be estimated by comparing its magnitude to $D\psi_1/Dt$, where both material derivatives use velocity in the top layer. The ratio of these terms would scale roughly as the ratio of 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 baroclinic instabilities, i.e. $[\bar{\psi}_{b.t.}^2/\bar{\psi}_1^2]^{\frac{1}{2}} \sim O(H_1/H_2)$. 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
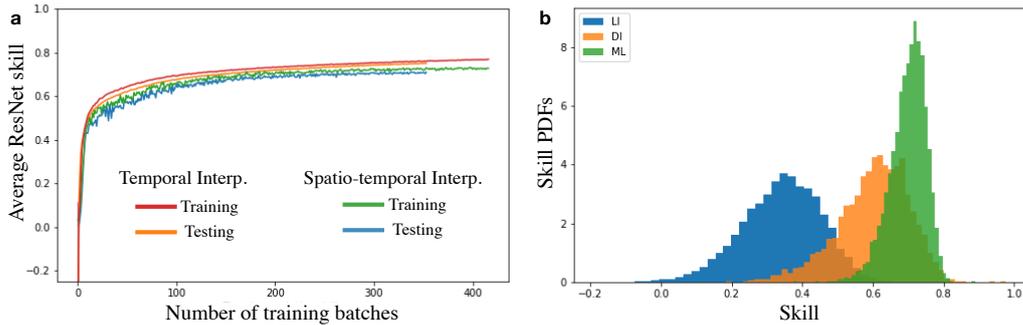
Figure 3. a) The evolution of the ResNet50 model skill during its training; 10 training mini-batches correspond to one epoch, while the actual number of batches is 128 and the total number of samples used in training is about 80,000; testing was conducted on 10,000 samples. b) comparison of skill distributions of the linear interpolation (LI), dynamical interpolation (DI), and the machine learning method evaluated on the test data set. Note that the lack of skills close to 1 hint at the existence of a dynamical barrier for SSH interpolation likely due to the chaotic nature of QG equations; i.e. the information content about the predicted SSH image is decreasing with the increased amount of time separation between the input image as the phase space trajectories are being mixed to the point where the dependence of initial conditions is being lost.

non-negligible and can substantially impact the SSH evolution leading to significant errors of the dynamical interpolation (see Results).

The QG model has been configured to represent baroclinically unstable mid-latitude currents such as the Gulf Stream or Kuroshio. 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 Flierl (1978); Arbic et al. (2012)). The QG model is integrated forward in time using an ensemble of noisy initial conditions to produce a large volume of data: about 100,000 SSH snapshots separated by 10 days (Figure 2). Over a timescale of 20 days, SSH fields become substantially decorrelated such that it is hard to identify any persisting eddies because their shapes and intensities have been dramatically changed due to interactions with other eddies (Figure 2). We ensure that the data for training/validation/testing comes from distinct simulations to accurately access the generalization skill of the neural network.

3 Results

3.1 Spatiotemporal SSH interpolation

Two separate neural networks were trained to perform two types of interpolation tasks to identify SSH field: i) temporal interpolation where the input consists of two SSH snapshots separated by 20 days, and ii) spatiotemporal interpolation with the same input as for the temporal interpolation but with SSH images having missing data. For the temporal separation of SSH images we chose 20 days because it is of the order of the return periods for existing altimeters and to be consistent with Ubelmann et al. (2015). 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 1000 km domain, SWOT would have
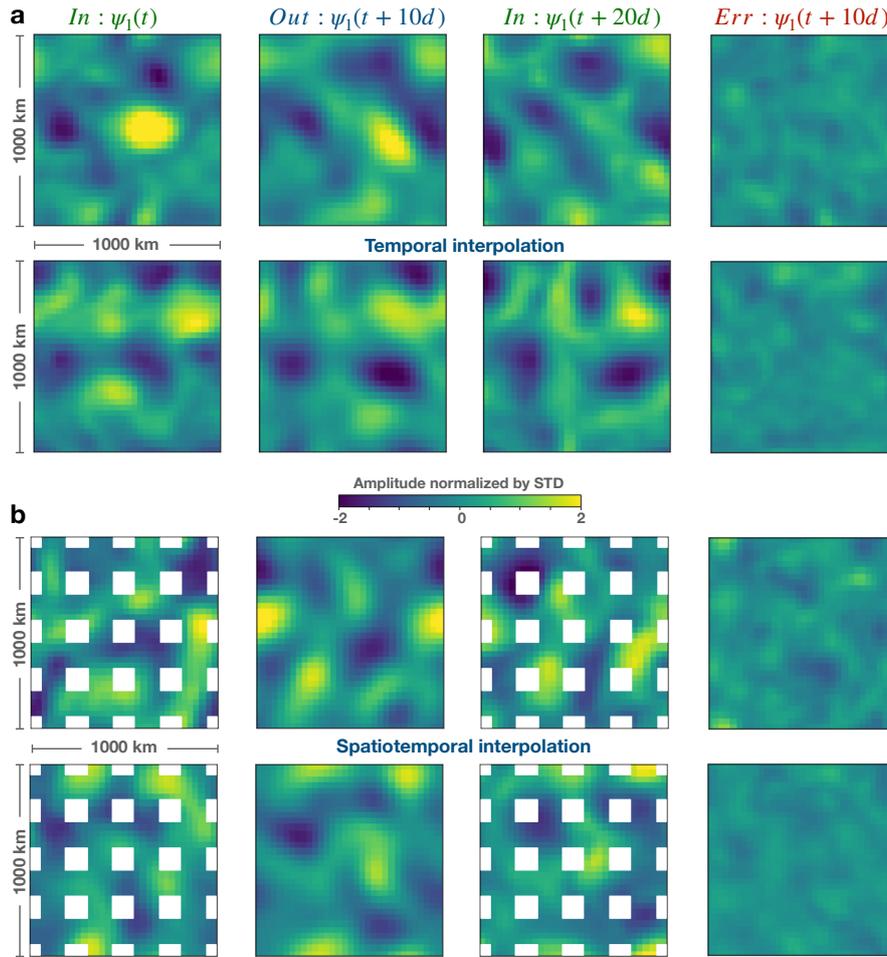
Figure 4. Examples of temporal (**a**) and spatiotemporal (**b**) interpolation of SSH data using the Deep Learning framework. Each row represents a randomly chosen interpolation example from the testing dataset (for a statistical distribution of prediction skill see Figure 3). The input SSH fields, $\psi_1(t)$ and $\psi_1(t+20d)$ are separated by 20 days and plotted in the first and the third columns correspondingly, while the predicted SSH field at day 10 ($\psi_1(t+10d)$) is plotted in the second column; the prediction error is plotted in the forth column. White regions in the case of spatiotemporal interpolation denote areas of missing data. Domain size for both input and prediction is 1000km by 1000km. SSH data for training and testing was generated using baroclinically unstable QG model of ocean turbulence with configuration pertinent to midlatitude ocean jets (see Methods).

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 any generality.

The neural networks were trained using about 100K data samples, both achieving a significant performance skill and producing realistic SSH images with small errors (see Figure 4). The average prediction skill for both simulations plateaued at about 0.7 and it wasn't significantly smaller when evaluated on the test dataset (Figure 3 a), implying that a generalized dependence has been found. A few illustrative examples of eddy field evolution are shown in Figure 4a, demonstrating the non-trivial 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 disappeared after 20 days, yet the neural network was 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 machine learning model outperformed the linear and dynamical interpolation techniques, having not only a better average skill but also much more infrequent occurrence of low-skill interpolations, i.e. much narrower skill-distribution tail in the direction of small skills (Figure 3b). Noticeably, the linear interpolation skill can be so low that it has values reaching zero, i.e. its prediction is no better than assuming that $SSH = 0$ everywhere in the domain. The dynamical interpolation is much better than that but still has a significant probability of poor interpolations in the skill range of about 0.4-0.6. While the machine 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 abounded by about 0.8. Since we are utilizing only surface observations while SSH evolution depends also on the unknown subsurface flow, it is, of course, expected that the interpolation skill would not be perfect: it is inherently a partial information problem. In addition to having partial observations, the chaotic nature of the flow also must be contributing to the skill limitation: if the SSH images are separated by a sufficiently large amount of time (greater than the characteristic Lyapunov exponent timescale), there should be no physical or statistical relationship between them and no interpolation technique could have a skill significantly above zero. Yet, the 20-day separation timescale, which is of the order of the return periods for existing altimeters, still allows one to extract sufficient information even for highly energetic baroclinically unstable flows.

3.2 State estimation of unobserved deep ocean flows at mesoscales

Here we assess the efficacy of the Deep Learning framework in addressing the state estimation problem, i.e. estimating all dynamical variables in the ocean turbulence model, which in our case of a two layer QG model implies estimating both surface and subsurface layer streamfunctions. Conventionally, for state estimation one needs to postulate the dynamical model and only then implement techniques e.g. data assimilation or the ensemble Kalman filter techniques in order to estimate unknown variables and parameters in the model at all times and everywhere within the model domain. However, we demonstrate here that the machine learning framework is capable to estimate both ψ_1 and ψ_2 with a high average skill of about 0.7 (Figure 5). It is important to note that while ψ_2 is highly correlated with ψ_1 (average correlation coefficient is about 0.8) it is the decorrelated part, $\tilde{\psi}_2 = \psi_2 - A\psi_1$, that is dynamically important for the SSH evolution. The neural network 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). It is thus clear that mesoscale eddy patterns imprinted in SSH do provide substantial information on deep ocean currents or equivalently information on the partitioning between the baroclinic and barotropic modes even for baroclinically unstable flows. Nonetheless, since both layers are dynamically active and no subsurface information is given by satellite observations, there is an inherent lack of information that is contained in sparse SSH observations and this 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 ocean mesoscale turbulence. Nonetheless, we have demonstrated here that the 20 day separation even for a strong baroclinically unstable current with a dynamically active subsurface flow does contain sufficient information for skillful interpolation and state estimation.

In providing a skillful state estimate from snapshots only, the deep neural network can essentially encode the underlying model equations and then utilize the information from subsequent SSH snapshot to even better constrain subsurface flows (see bottom panel in
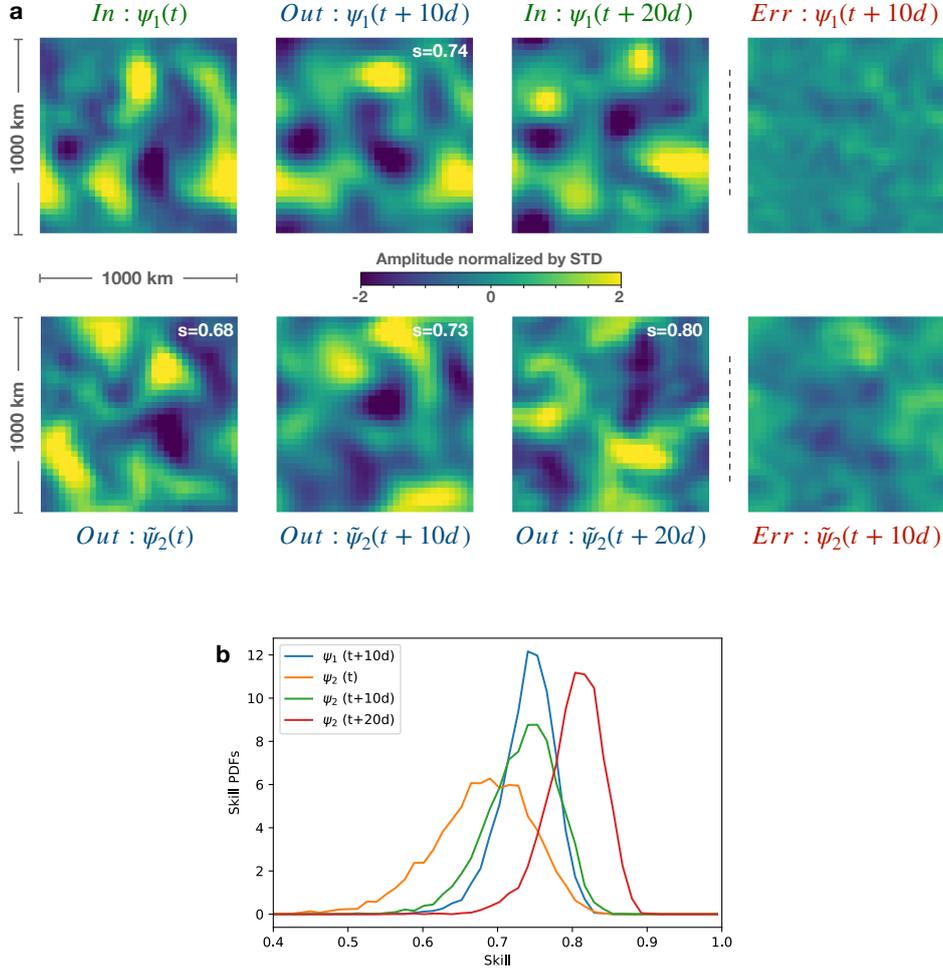
Figure 5. Examples of state estimation using Deep Learning neural network (**a**) and its statistical skill distribution for surface and subsurface variables at different times (**b**). As in the case of SSH interpolation, the neural network receives as input two SSH snapshots separated by 20 days, $\psi_1(t)$ 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, $\psi_2(t, t + 10d, t + 20d)$. 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 subsurface streamfunctions are shown in last columns. The probability density function of the neural network skill distribution is plotted in panel (**b**) for all predicted variables. Note, while the neural network provides skillful predictions for all variables (skill ranges from 0.65 to 0.85), the best prediction skill is achieved for subsurface flow at day 20 and the worst prediction is for subsurface flow at day 0.

Figure 5). It important to note the asymmetry in reconstructing subsurface velocities: at earlier times the skill is substantially worse (compare the orange and red curves in bottom panel of Figure 5). This asymmetry is expected in a chaotic and dissipative quasigeostrophic dynamics that makes it more difficult to estimate past state by observing the future as opposed to estimating future by observing the past. Thus, the two SSH snapshots must indeed be ordered in time as the PV-evolution equations only allow time reversal only for sufficiently small time intervals at which the dissipation effects can be neglected.

4 Discussion

Here we presented the proof of concept for using Deep Learning as an efficient tool to extract non-trivial information from sparse SSH observations, specifically demonstrating its utility in the spatiotemporal interpolation of SSH data and more generally in the state estimation that includes deep ocean currents apart from SSH. The Residual CNN outperformed the commonly used dynamical interpolation method. While it is challenging to precisely interpret the algorithm that was ultimately learned by the neural network, its success seems to be associated with its ability to predict subsurface flows only from individual snapshots of mesoscale eddy patterns. This separates machine learning from other methods that are incapable of dis-entangling the highly-nonlinear relation between surface and subsurface flows and hence gives machine learning the edge in constructing a more accurate model of SSH evolution, resulting in higher quality interpolation.

Here we only considered the case of mesoscale turbulence and for the case of submesoscale turbulence, 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 machine learning framework to perform well in reconstructing large and small mesoscale eddies, its limitations still need to be understood when considering mesoscale and submesoscale turbulence as a continuum. The main difficulty arises because the persistence timescale for submesoscale eddies is substantially shorter than for mesoscale eddies and no technique can bypass the inherent loss of the dependence on initial conditions for a chaotic system.

We have used a model of baroclinic turbulence as a synthetic training data because it presents a hard test for temporal SSH interpolation due to its chaotic nature and an a priori unknown impact of the dynamically active bottom layer on SSH evolution. We recognize that ocean regions where SSH variability is dominated by waves or other processes unrelated to baroclinic instabilities, a neural network that was trained to represent baroclinically unstable currents could perform poorly. Thus, it is necessary to develop more general training datasets that are more representative of the SSH dynamics for any given region of interest. Those could be ranging from more realistic mesoscale-resolving general circulation models to simplified stochastic QG-based models Samelson et al. (2019). Note, however, that it is essential for the synthetic model not to be overly simplistic to the point that it misrepresents the nature of SSH variability. The drawback of Deep Learning is that it generally requires a large number of data for training. Nonetheless, there are continuously improving methods aimed at addressing this practical issue, e.g. transfer learning (Pan & Yang, 2009) or one-shot learning (Fei-Fei et al., 2006).

A way towards ultimately developing the gridded SSH product using machine learning could be through training networks on a wide range of idealized and realistic models and then fine-tuning a much smaller number of neural network 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 neural network cannot by achieved by defining a simple loss function as was done for synthetic data. Thus, the neural network 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 neural network 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 stir away from solely relying on dynamical models to provide training data is a necessary next step towards practical implementation with real satellite observations. Nonetheless, our work presents an important proof of concept, demonstrating that SSH observations do contain dynamically-relevant information about subsurface flows and hence with deep learning it is possible to build a skillful self-contained model of SSH evolution and as a consequence improve 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 requires large computations, providing an accurate initial guess could substantially reduce the number of iterations necessary for optimization. It might be possible to accelerate data assimilation methods by providing the machine learning estimate as a first guess that is already close to reality. We thus see the synergy between machine 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 with machine learning and fine-tuning to ensure its strict consistency with a given dynamical model achieved by data assimilation.

References

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

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.

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.

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, 09). The vertical structure

of ocean eddies. *Dynamics and Statistics of the Climate System*, *1*(1). Retrieved from https://doi.org/10.1093/climsys/dzw001 (dzw001) doi: 10.1093/climsys/dzw001

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.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 770–778).

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Klein, P., Isern-Fontanet, J., Lapeyre, G., Roullet, G., Danioux, E., Chapron, B., . . . Sasaki, H. (2009). Diagnosis of vertical velocities in the upper ocean from high resolution sea surface height. *Geophysical Research Letters*, *36*(12).

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.

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

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*(2019).

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. *arXiv preprint arXiv:1603.08029*.

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*. Cambridge University Press.

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.