Deep learning to infer eddy heat fluxes from sea surface height patterns of mesoscale turbulence

Tom M George[1,2], Georgy E Manucharyan[1,3,*], and Andrew F Thompson[1]

[1]Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, 1200 E California Blvd, California, 91125, USA
[2]The Cavendish Laboratory of Physics, University of Cambridge, JJ Thompson Avenue, Cambridge, CB1 3FZ, UK
[3]School of Oceanography, University of Washington, Seattle, Washington, 98195, USA
[*]Corresponding author email: gmanuch@uw.edu

ABSTRACT

Oceans play a major role in Earth's climate by storing and transporting heat via turbulent currents called mesoscale eddies. However, direct monitoring of eddy-driven heat fluxes is currently impossible because it requires simultaneous surface and subsurface observations of velocity and heat content, while only surface properties of mesoscale eddies can be comprehensively measured by satellites in the form of sea surface height (SSH) anomalies. Nonetheless, surface and subsurface expressions of eddies are dynamically linked, suggesting that surface observations may contain at least partial information about subsurface flows and thus heat transport. Here we used deep convolutional neural networks (CNNs) to demonstrate that SSH expressions of mesoscale turbulence contain sufficient information to predict about 64% of eddy heat flux variance, with CNNs significantly outperforming other conventional data-driven techniques. Our results suggest that deep CNNs could provide an effective pathway towards an operational monitoring of eddy heat fluxes using satellite altimetry.

Introduction

Large-scale oceanic and atmospheric circulations are constrained by vertical stratification and planetary rotation. Energetic, vertically-sheared ocean flows, e.g. the oceanic Gulf Stream and Kuroshio currents, are baroclinically unstable[1–3], and generate mesoscale eddies – vortices with scales of motion of the order of 10-100 km in the open ocean[4,5]. Across most of the ocean, mesoscale eddies dominate the kinetic energy and cumulatively make the dominant contribution to the transport of tracers across frontal currents[6–8]. Due to computational limitations, even mid-latitude oceanic mesoscale eddies are not fully-resolved in climate projection models[9,10] so their impact on larger-scale circulations and tracer fields must be parameterized. Specifically, the eddy heat (or more generally, the buoyancy) fluxes directly affect the evolution of the mean ocean currents by draining their potential energy[11]. Expressing the magnitude of these cumulative eddy heat fluxes in terms of mean properties of the ocean is the essence of prevailing eddy parameterizations[12].

Comprehensive monitoring of eddy heat fluxes to test and inform physically-based parameterizations remains a major challenge for several reasons. Firstly, the eddy flux is a noisy quantity, it is not sign-definite and its domain-average can be an order of magnitude smaller than its local maximum. Secondly, the dynamics of the baroclinic instability depends on interactions between the upper and lower layers of the ocean[1] and hence direct calculations of eddy heat fluxes require *simultaneous* observations of the upper-ocean horizontal velocity field alongside the subsurface heat content anomalies, at eddy scales. However, current oceanic *in situ* observations, e.g. by ARGO floats[13], moorings[6] or ship transects, are spatially and/or temporally sparse, leaving the interior ocean poorly observed at eddy length scales such that explicit evaluation of eddy heat fluxes is not possible. Surface expressions of mesoscale eddies are monitored globally by satellite-derived estimates of dynamic sea surface height (SSH, which is proportional to the streamfunction of the surface geostrophic flow, see equations [3] & [4]) anomalies[14], which have significantly advanced our understanding of eddy propagation[15], ocean energetics[16–18], and tracer diffusivities used in parameterizations[19,20].

Since eddy heat fluxes can not be directly calculated from surface flows alone, several techniques have been proposed to use satellite SSH observations by fitting them to an appropriate model, e.g. data assimilation at eddy-permitting scales in primitive equation models[21], parameter estimation via stochastic Kalman-type filters in quasigeostrophic models[22], and statistical reconstruction of subsurface velocities using self-organising maps[23]. These different methods, as well as our own, are linked by the theme of using only surface observations to infer subsurface velocities or eddy heat fluxes despite the fact that no theoretical prediction of the surface-subsurface mutual information currently exists. It is also important to note here that subsurface flows in the ocean are highly correlated with surface flows such that the vertical distribution of currents can be

represented with just a single empirical orthogonal function (EOF or PCA) capturing over 80% of the variability[24]. However, linearly correlated components of surface and subsurface flows contribute no meaningful domain-averaged heat flux (see Methods) and eddy heat fluxes are maximised when subsurface flows are $\pi/2$-phase-shifted with respect to the surface flow. Thus, studies presenting highly-skilled reconstructions of subsurface flows from surface flows (e.g.[23,25]) may in fact only be reflections of their high degree of linear correlation (see figure 1b) from which one cannot deduce how accurate any subsequent estimation of the eddy heat flux would be. A natural question arises: how much information is contained in the SSH field with regard to heat flux estimation? On the one hand, for any given SSH pattern of eddies, many dynamically-consistent subsurface flows could be constructed to ensure that eddy heat fluxes take a range of values, implying that SSH data provides incomplete information to determine the eddy fluxes. On the other hand, the likelihood that any predefined subsurface flow emerges as a result of baroclinic instability and eddy interactions may be small, except for a narrow range of specific flows corresponding to a given SSH pattern. From a quasigeostrophic (QG) perspective[26], i.e. large-scale, slowly-evolving turbulence in rotationally-constrained fluids, circulation in the bottom layer is likely associated with potential vorticity anomalies that would also induce surface flow[27], thus affecting the distribution of the SSH field. That is to say any particular pattern of eddies in an SSH snapshot should significantly constrain the posterior distribution of subsurface flow and thus contain at least partial information about the corresponding eddy fluxes.

Given the highly non-linear and chaotic nature of mesoscale turbulence, a direct theory-based approach to reconstructing heat fluxes from SSH snapshots has not yet been developed, suggesting consideration of data-driven approaches. Here we make progress on this problem by using deep Artificial Neural Networks[28] (ANNs), which are powerful tools for extracting critical, if subtle, information from large volumes of convoluted data[29–34]. ANNs are widely used for supervised learning tasks where an approximation of an input-to-output mapping can be developed by optimising a highly nonlinear function with respect to a large number of trainable parameters. Specifically in fluid mechanics, deep neural networks have been used to address the closure problem in Reynolds-averaged Navier Stokes equations[35–37], outperforming other data-driven methods such as dimensionality-reduction via proper orthogonal decomposition[38] or dynamic mode decomposition[39,40]. In theory, deep ANNs can approximate nonlinear mappings of any complexity, provided the network contains a sufficient number of free parameters, and there exists a sufficient amount of training data.

For geostrophic turbulence[26], neural networks have been used to demonstrate a strong potential for parameterizations of eddy momentum fluxes in barotropic[36,41] and baroclinic[25] ocean gyres. The general premise of these studies is that the unresolved eddy fluxes at any particular time or location depend only on the distribution of large-scale fields that are resolved by the model. However, it is clear that external dissipation and forcing processes combined with the spatial and temporal non-locality in the evolution of eddy fluxes[42] limits the appropriate length and time scales at which the parameterization could apply. For instance, even in a basic QG model with eddies forming due to baroclinic instabilities of a spatially- and temporally-uniform mean flow[2], the domain-averaged eddy heat fluxes fluctuate on monthly timescales with an amplitude similar to their time-average value (Figure 1d). Since the mean flow does not change, it is impossible to express the eddy fluxes in terms of the mean flow at timescales of the order of months and shorter. Yet, ocean fluctuations on diurnal and seasonal timescales are critical for coupling with atmosphere and biosphere. Thus, at these relatively short timescales, mean flow observations alone are insufficient and we must search for ways to estimate the time varying components of the eddy heat flux from the information (i.e.SSH anomalies) we have available.

In this study, we strive to extract the maximum amount of information about (i.e. predict with highest confidence) the *instantaneous* domain-averaged eddy heat fluxes from SSH observations of the domain's mesoscale eddy field. We quantify the potentially existing limits of the predictive capabilities of SSH data for diagnosing eddy heat fluxes by using powerful data-driven approaches involving deep Convolutional Neural Networks (CNNs). We search for data-driven solutions to the eddy heat flux problem by constructing large volumes of training data for supervised learning using mesoscale eddy-resolving models. Since the variability of large-scale oceanic flows is predominantly contained in the barotropic and the first baroclinic modes[43–46], our research philosophy here is to use one of the most fundamental and influential models of baroclinic turbulence, the two-layer QG model[2]. Using this idealised model allows us to estimate the heat flux predictability limit free from other practical constraints such as the number of available samples, their spatial sparseness, measurement inaccuracy and external noise. We aim to reconstruct the instantaneous, domain-averaged eddy fluxes for a steady background mean flow given only the corresponding SSH snapshots. This presents a less-constrained problem compared to reconstructing long-term mean eddy fluxes, which are indeed expected to depend only on the mean flow. We will demonstrate that CNNs are powerful tools in extracting abstract information from SSH snapshots by identifying spatial patterns containing most the relevant information for flux predictions. Nonetheless, we will demonstrate that even in this idealised turbulence model, there appears to exist an upper bound on prediction capabilities of SSH data.

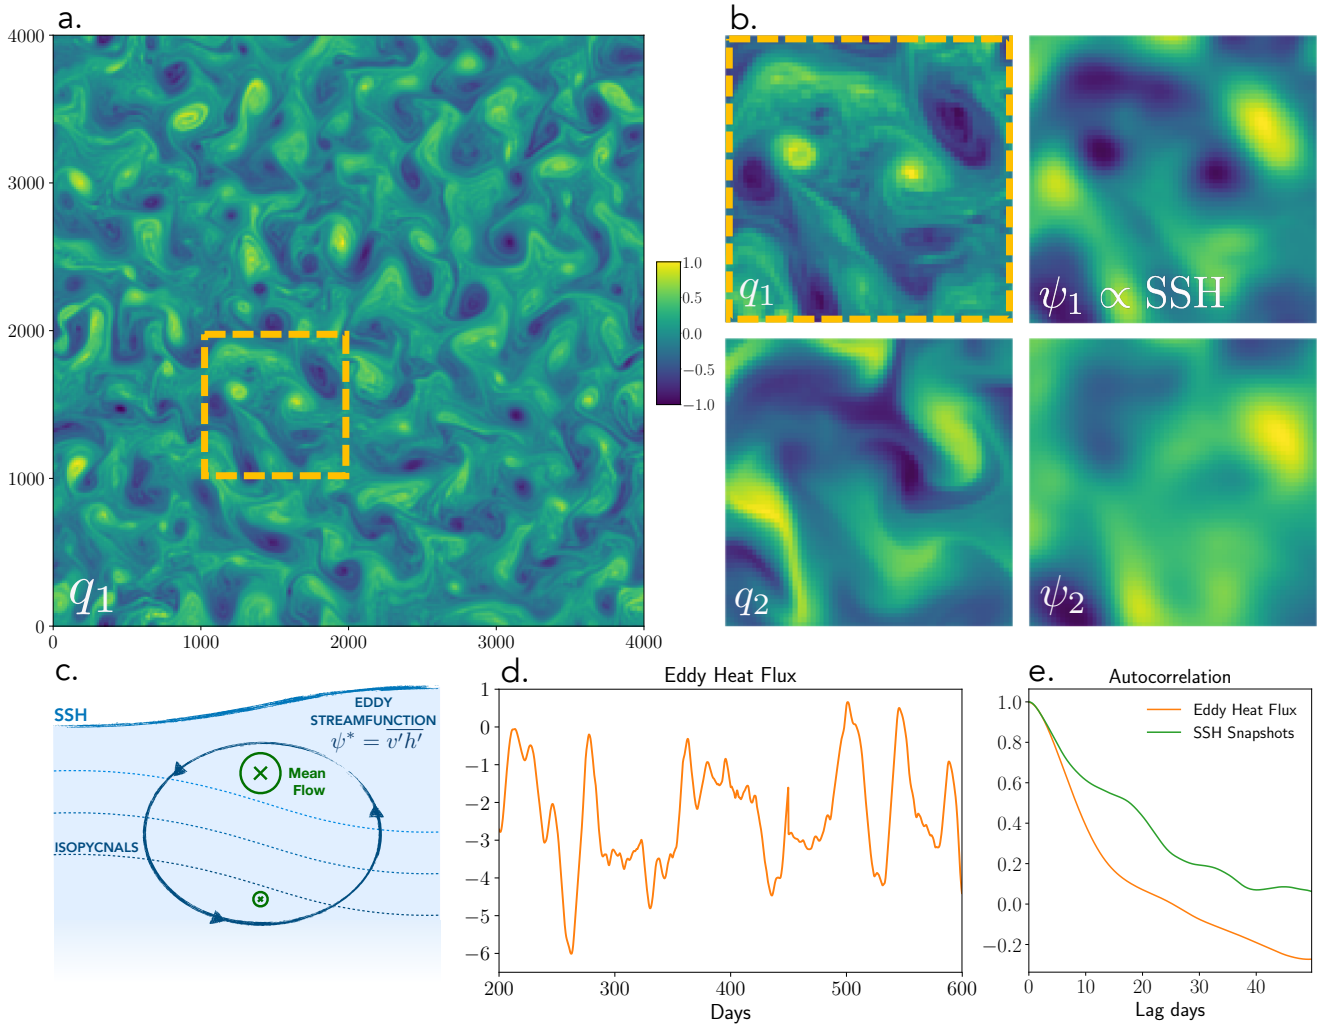

Figure 1. Simulations of baroclinic geostrophic turbulence on a β-plane (see Methods). **(a)** Snapshot of the upper-layer potential vorticity field, q_1, showing the prominence of mesoscale eddies that are centred around local extrema; orange dashed square shows a representative 1000x1000km subdomain used for CNN training. **(b)** Subdomain snapshots of potential vorticities q_i and streamfunctions ψ_i in upper ($i = 1$) and lower ($i = 2$) layers. All variables have been non-dimensionalised by scaling such that their extremal pixel has a value of modulus 1. **(c)** Cross-sectional schematic of a stratified ocean, demonstrating the thermocline tilt and corresponding vertically sheared mean flow (for the northern hemisphere); the cumulative impact of all eddies creates an overturning eddy streamfunction, ψ^*, that acts in a direction to flatten the thermocline slopes and decelerate the mean flow. **(d)** Temporal evolution of the domain-averaged eddy heat flux showing representative fluctuations within the statistically-equilibrated mesoscale turbulence. **(e)** Autocorrelation functions for the eddy thickness flux (orange) and SSH snapshots (green); note that the eddy fluxes decorrelate almost twice as fast compared to SSH snapshots.

1 Results

Eddy heat fluxes in geostrophic turbulence

We conduct an idealised numerical simulation of two-layer quasi-geostrophic turbulence with a horizontally-uniform vertically-sheared mean flow that is constant in time (see Methods for model equations and parameters). After the initial spin-up phase, the baroclinic turbulence equilibrates and mesoscale eddies approximately 200km in size are prominent throughout the domain (Figure 1a). A large, 4000x4000km doubly periodic domain (containing 100 Rossby deformation radii per side), is divided into 16 subdomains of 1000x1000km, which are considered distinct samples of mesoscale turbulence; note that the subdomain variables are no longer doubly periodic and hence individual eddies passing through its boundaries can significantly affect the

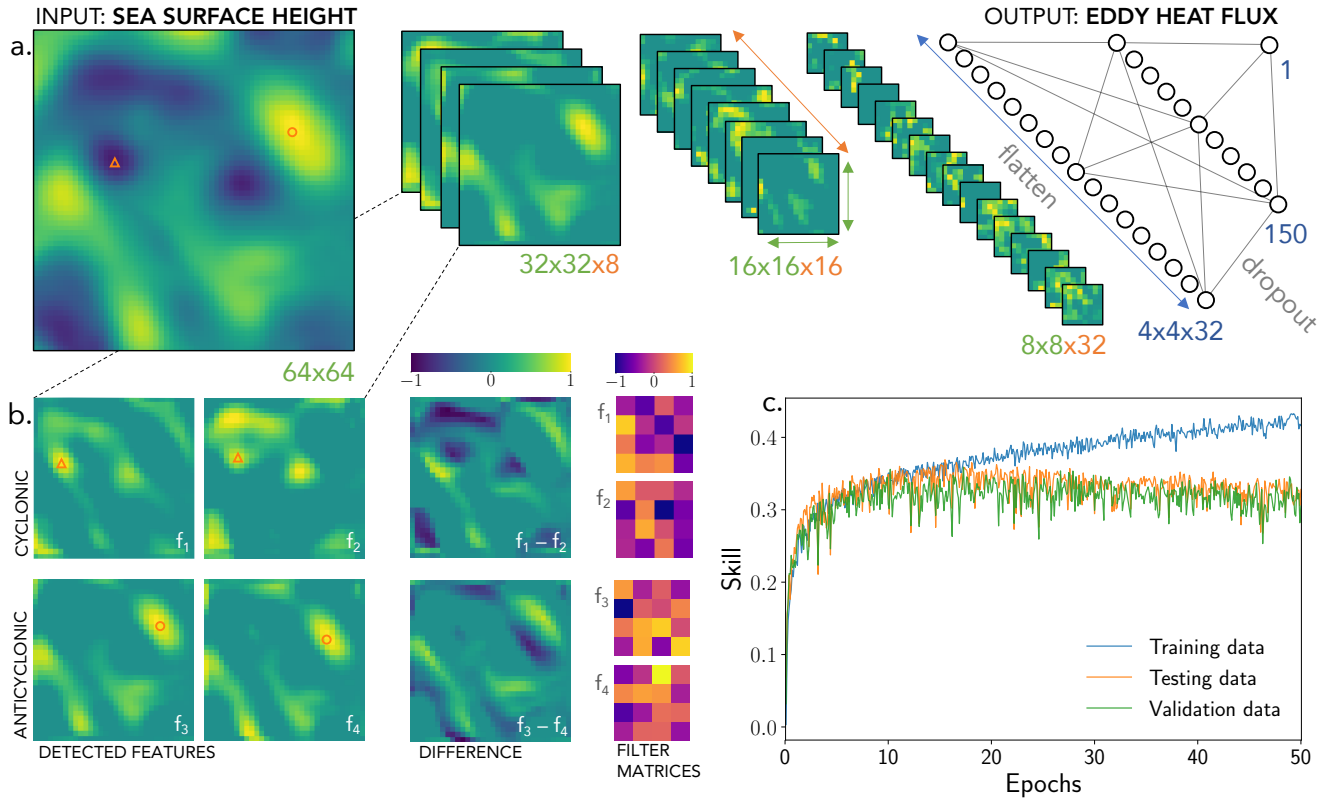

Figure 2. CNN architecture optimised to predict eddy heat fluxes from SSH snapshots. **(a)** The CNN architecture consists of multiple convolutional filters that are consecutively downsized, representing more abstract features of SSH patterns at each subsequent layer; the learned convolutional features are mapped to the eddy heat flux using the final two hidden fully-connected layers. These sub panels shows the progression of an example SSH input through each layer of the trained network. **(b)** Example of the outputs from the first layer convolutional filters after training, (left two columns) f_i, applied to the input SSH snapshot that we find are split into cyclonic and anticyclonic representative filter groups. Also shown are intra-group differences of filtered SSH snapshots (third column), and corresponding 4x4 convolutional weight matrices (fourth column). The triangle/circle markers highlight example cyclonic/anticyclonic eddy centres. **(c)** The evolution of CNN skill for training, testing and validation data (blue, orange and green curves respectively) plotted as a function of training epochs, i.e. the number of times each data point has been used in the optimisation procedure. All images are scaled such that their extremal pixel has a value of modulus 1.

eddy heat fluxes. The dynamical variables (potential vorticities and streamfunctions) in both layers have visual similarity in the form of eddies and filaments (Figure 1b). Cyclonic and anticyclonic eddies are visible in SSH snapshots despite having less pronounced filamentary structures at their edges compared to potential vorticity snapshots (Figure 1a,b). The averaged eddy heat fluxes are directed in such a way as to induce an overturning circulation that would flatten the tilted thermocline and slow down the mean flow (Figure 1c), although in our model the mean flow is prescribed and so is a perpetual source of energy.

The eddy heat fluxes fluctuate dramatically and stochastically on monthly timescales, ranging in magnitude from nearly zero to over double their mean values (Figure 1d). The heat flux decorrelation timescale, of the order of 20 days, is roughly half that of the SSH snapshots (Figure 1e), implying that a subtle change in SSH patterns could translate to a significant change in the eddy heat flux. The high sensitivity to SSH patterns reflects the fact that eddy heat fluxes are proportional to a correlation between surface velocity, which is directly related to SSH (see equations [3] & [4]), and an unknown subsurface streamfunction (equation [5]) – both of which evolve differently according to a set of highly nonlinear and coupled equations (equations [1] & [2]). It is the lack of any information about either subsurface flow or (equivalently) thermocline depth anomalies that makes the problem of eddy heat flux mapping from SSH data mathematically ill-defined. The SSH snapshots and the corresponding spatial-mean eddy heat fluxes the in subdomains are used as training data for deep CNNs.

Deep CNNs predict eddy heat fluxes from SSH snapshots

Here we discuss the CNN architecture and it's skill in predicting instantaneous domain-averaged eddy heat fluxes given only SSH snapshots. The optimal network architecture mapping the input (SSH snapshot) to the output (eddy heat flux) is conceptually shown in Figure 2a, consisting of three convolutional layers followed by two fully-connected neuron layers. Convolutional layers filter the output from the previous layer (Figure 2b, left two columns) by convolving them with 4x4 weight matrices (Figure 2b, right-most column). The information contained in the last convolutional layer is then passed to a conventional network consisting of several hidden layers of neurons and eventually mapped to the output. Each time the information is passed from one layer to another, a specific non-near activation function[47] is applied to generate abstract information that allows ANNs to construct highly non-linear mappings. The CNN has O(10^5) free parameters that are iteratively updated using a stochastic gradient descent method to maximise the prediction skill. It takes about O(10) epochs to find the optimal set of parameters (Figure 2c), with each epoch corresponding to one complete presentation of the training dataset, O(10^5) points, to the network. Note that if left for $\sim$20 epochs the network learns to approximate the *training* data even more accurately but this is a clear sign of overfitting as the skill evaluated on the *testing* data (which is not used in the optimisation) begins to fall. It is thus essential to ensure that training and testing data samples are completely independent from each other, which we do by using separate numerical simulations starting from independently seeded noisy initial conditions.

By optimising information extraction from SSH snapshots, the learned filters reflect dynamically relevant features. The set of filters in the first convolutional layer, f_i, can be split into two representative groups identifying cyclones, e.g. f_1, f_2, and anticyclones, e.g. f_3, f_4, (Figure 2b, right two columns), while slight filter differences ($f_1 - f_2$ and $f_3 - f_4$) emphasise eddy gradients and edges, particularly for dipoles 2b, middle columns). From linear stability analysis, the eddy heat fluxes should depend on the relative position and strength of eddies in both layers, with the magnitude of the flux being particularly strong in baroclinic dipoles known as hetons[48]. Thus, it is reassuring that the network has learnt to extract this type of information from SSH snapshots. In later layers the information becomes too abstract for interpretation. The average testing skill (success metric define in equation [6]) achieved by the CNN tops out at about 0.35 (Figure 2c), corresponding to a relatively high correlation of 0.8 between the predicted and true eddy fluxes. The CNN is highly efficient at extracting the required information from SSH patterns, so much so that they perform better than other data-driven methods that either disregard the two-dimensional nature of SSH data or attempt to use more simplified linear methods (see Methods and Figure 4).

Despite explaining about 60% of the eddy heat flux variance (Figure 3a), CNN predictions have several systematic biases reflecting the fundamental limitations of the information contained in SSH snapshots. Firstly, extreme values of eddy heat fluxes (over one or two standard deviation from the mean) are persistently underestimated by the CNN predictions (Figure 3b,d). Increasing the number of training examples of extreme eddy fluxes as well as testing various optimisers (Stochastic Gradient Descent, Adam, Adamax[49]) and losses (mean absolute and square error) did not significantly improve their representation in the CNN, suggesting that the limited skill and biases are not due to deficiencies of CNN architectures or the lack of data but rather caused by inherent incompleteness of the information contained in SSH snapshots. Secondly, when trained on decorrelated discrete SSH snapshots separated by 10 days, but evaluated on a continuous SSH time series, the CNN generates elevated variability at timescales shorter than about 10 days (Figure 3c). Nonetheless, the CNN still generates a relatively smooth eddy flux time series (Figure 3d) with similar statistics to the true flux. Whilst superior networks architectures that would eliminate these biases might exist, it is not evident they could achieve a significantly higher predictive skill. Indeed, increasing the network complexity or the number of training samples does not indicate any prospects for improving the skill (3e,f), implying that this may be a theoretical upper bound. Nonetheless, as Figure 3d clearly demonstrates, CNNs do provide a significantly accurate estimation of eddy heat fluxes despite using inherently incomplete information as input.

Optimal CNN complexity and required volume of data

To identify and prevent overfitting, regularisation techniques, such as splitting the data into independent training/validation/testing sets, applying random dropout to neurons and early stopping of training, have been applied during CNN optimization. For any neural network, an insufficient number of data samples would inevitably lead to overfitting and/or skill reduction thus the optimal volume of data necessary to achieve the maximum skill depends on the network architecture: higher complexity networks with larger number of free parameters generally can achieve higher prediction skills but require larger volumes of data. Specifically for the CNN architecture used in our study, it is necessary to have at least 20,000 SSH images (and their corresponding eddy heat fluxes) in order for training to achieve the maximum skill and to avoid significant overfitting (Figure 3e). When it comes to using *real* ocean data to train a neural network, even though the required number of SSH snapshots may be available from satellite altimetry databases, a severe lack of spatially- and temporally-dense interior ocean measurements means there is little chance of being able to calculate their corresponding heat fluxes. Simply put, the required number of training samples is currently too large to make practical progress and other, more efficient, network architectures must be considered to reduce this number.

We could not construct a CNN architecture that can significantly surpass the skill of 0.35, even when a total of 200,000

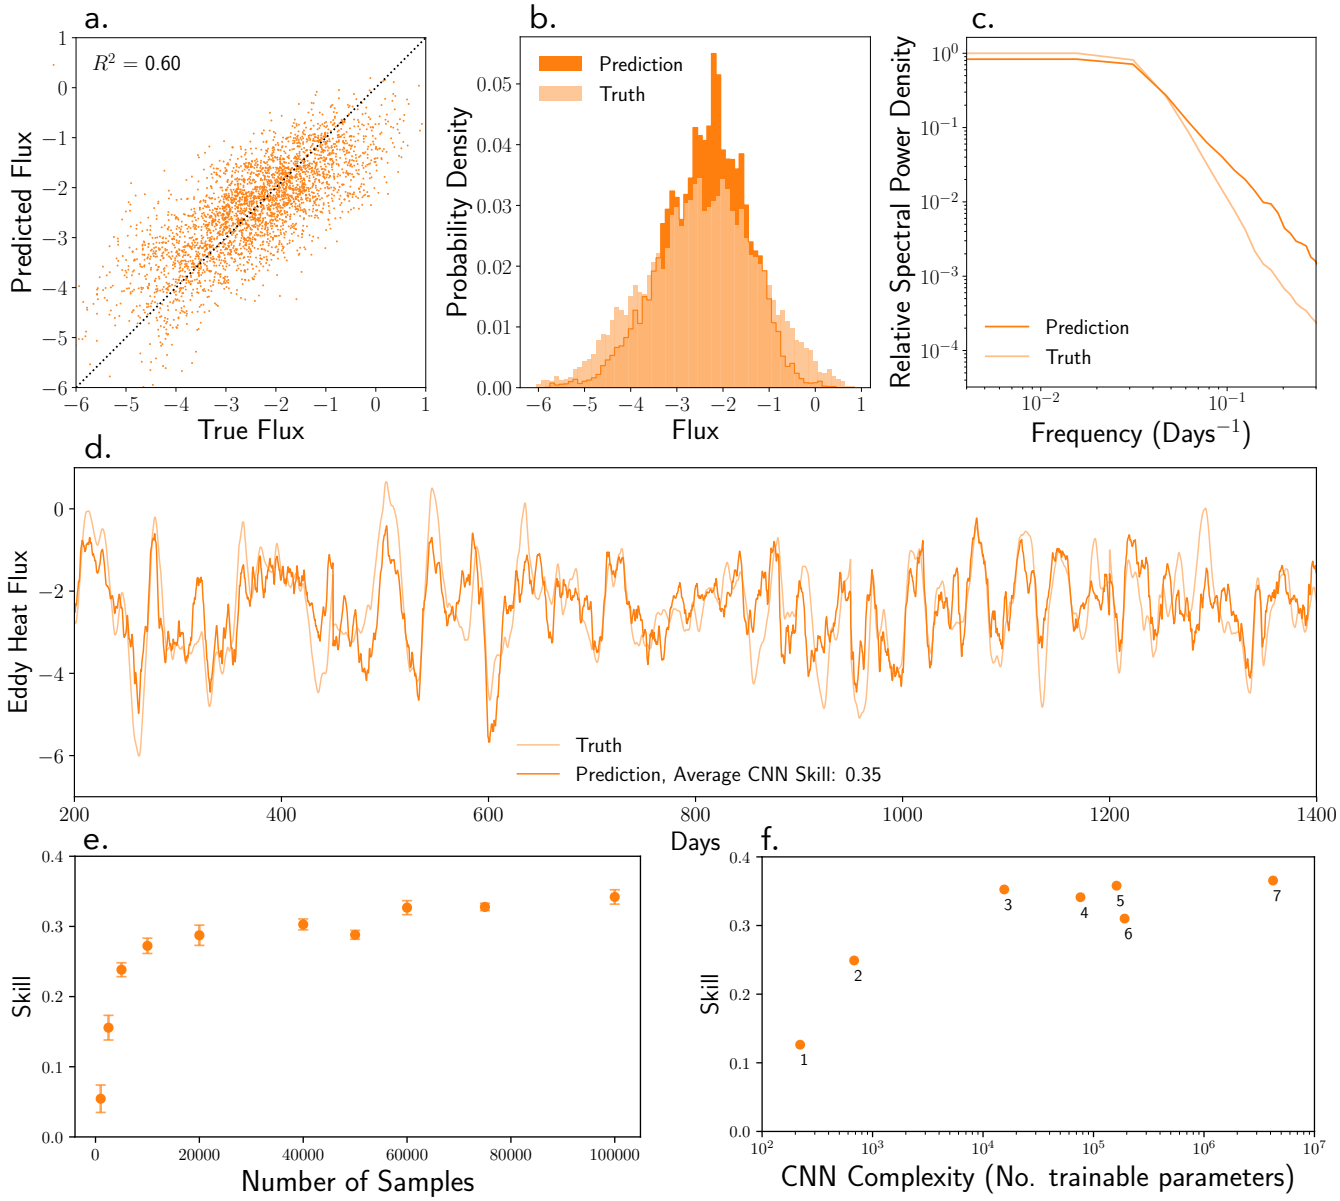

Figure 3. Performance of CNNs in estimating instantaneous domain-averaged eddy heat fluxes from SSH snapshots. Panels **a-d** show comparisons between the CNN predictions (dark orange) with true values (light orange) for various performance diagnostics. **(a)** Scatter plot of predictions for the validation data set demonstrating that CNNs explain over 60% of the flux variance (max achieved $R^2 = 0.63$). **(b)** Histograms highlighting the biases of CNNs towards underestimating extreme values of eddy fluxes. **(c)** Power spectra highlighting the CNN biases towards producing noisier time series at frequencies higher than O(0.1 days^{-1}). **(d)** Time series showing the fluctuations of the true and predicted eddy heat flux highlighting the skill and relative temporal smoothness of the prediction as well as the deficiencies of CNNs on individual events. The prediction plotted is for the CNN state with maximum validation skill, corresponding to a test data skill of 0.36 (the average skill of the NN is 0.34 as shown in Figure 4 and defined in Methods). **(e)** Sensitivity of the fully trained CNN skill on the fraction of data samples used for training (a fraction of 1 corresponds to 10^5 samples); for insufficient volume of data (below about 10^4 samples), there is a dramatic decrease in prediction skill. **(f)** Prediction skills achieved for a variety of CNN architectures ranging from simplistic to deep; the number of adjustable parameters is shown on the x-axis as a rough measure of CNN complexity while hyperparameters are referred to by numbers and described in Methods.

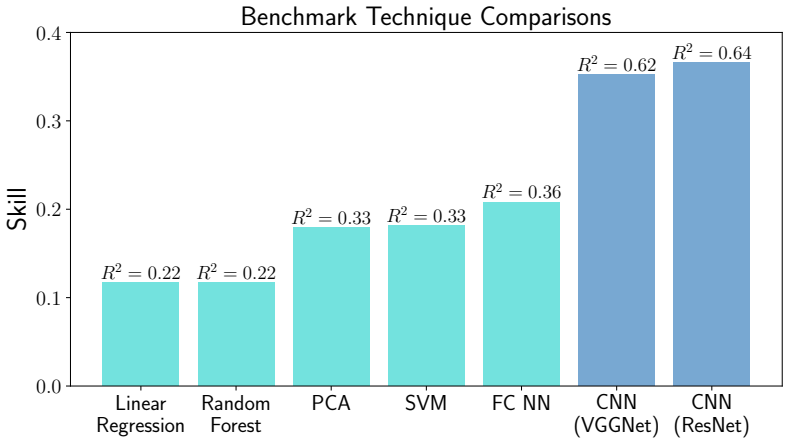

Figure 4. Performance chart for a set of various data-driven techniques (see Methods) used to estimate eddy heat fluxes from SSH snapshots. The y-axis represents their obtained skill and each column has a denoted value of R^2 that reflects the fraction of the total eddy flux variance explained by the fit. CNNs (dark blue) significantly outperform all other methods (light blue) as they are optimised to extract the most informative patterns form SSH snapshots. Note that the FC and VGGNet training time is O(1 cpu hour), whereas ResNet takes O(1000 cpu hours) to train.

training samples were used on very deep CNNs. Instead, we find that there exists an optimal CNN complexity for this problem: simpler networks cannot achieve the highest possible skill, while complex networks struggle with significant overfitting (Figure 3f). The optimal CNN architecture (Figure 2a) involves a large number of trainable parameters, $O(10^5)$, and hence should be sufficiently powerful in recovering any physically-constrained dependencies between eddy fluxes and SSH snapshots, if they exist. The existence of the upper bound of the skill implies that there are process-based limitations on the information contained in SSH snapshots, i.e. a significant fraction of subsurface eddy variability does not have a well-defined surface expression. Nonetheless, the CNNs with the achieved skill of 0.35 explain over 63% of the eddy heat flux variance, performing substantially better compared to other statistical methods including linear regression, principle component analysis, support vector machines or random forests (Figure 4).

Discussion

By training deep CNNs on data from high-resolution simulations of baroclinic turbulence, we showed that eddy patterns in SSH snapshots contain sufficient information to estimate the instantaneous domain-averaged eddy heat fluxes, accounting for over 60% of their variance. In fact, CNNs substantially outperform other traditional data-driven techniques such as principle component analysis, linear regression, random forests, statistical vector machines, or basic fully connected neural networks (Figure 4). Furthermore, tested over a broad range of network architectures, domain sizes and training data set volumes, we have identified an upper bound to the ultimate predictive power of CNNs trained on SSH data: an indication that we may have approached the theoretical dynamics-based bound on the information content in SSH snapshots. The existence of the bound would imply that analytical laws directly derived from the QG equations and linking eddy tracer fluxes to SSH snapshots do not exist (otherwise CNNs would have approximated these laws) and only data-driven regression or probabilistic models could be relevant. Nonetheless, the bound reached by CNNs is sufficiently high to potentially provide valuable information on eddy-mean flow interactions globally for the ocean, which could be used for developing and explicitly testing eddy parameterizations for climate models.

Our idealised study provides a proof-of-concept that CNNs can be used for estimating eddy heat fluxes from satellite altimetry. However, there are several remaining issues that must be addressed before one could implement this approach operationally using satellite SSH observations. Firstly, even within our controlled numerical experiments that are free of all external processes, the amount of training data necessary for the supervised learning of CNNs is substantial, $O(10^5)$ samples. Secondly, our idealised simulations of baroclinic turbulence were performed with constant mean flow and stratification parameters, thus ignoring spatially and temporally non-local eddy processes that may complicate the CNN learning and require even larger volume of training data. Thirdly, in practice, directly measuring eddy heat fluxes for supervised learning would require an eddy-resolving network of ocean instruments such as ARGO floats, gliders or moorings that are expensive to sustain on basin scales and over long periods of time. Thus, it is necessary to develop superior deep learning techniques, potentially implementing meta-learning strategies[50–52], that could reduce the volume of necessary training data by at least an order of

magnitude. Another way to approach the limitations of obtaining the training data is to explore the possibility of transfer learning[53]: here CNNs could be almost entirely trained on a large volume of model data and afterwards use a much smaller fraction of ocean observations for validation and fine-tuning of weights and biases.

Finally, while CNNs are powerful generalization tools, it is plausible that the upper bound on the predictive capability of SSH snapshots obtained from CNNs may be surpassed by other more complex and fine-tuned machine learning architectures, e.g. by variational autoencoders[54], cross-modal CNNs[55] or ultra-deep CNNs[56]. Thus, the ultimate dynamics-based upper bound remains open. To address this question, we have made the training and testing data set available to the community and welcome attempts to improve upon the bounds we found in this study, or more specifically, determine if the eddy heat fluxes can be estimated from SSH snapshots with a skill higher than 0.35 using any other methodology.

Data Availability

The necessary procedures to generate the data and reproduce the machine learning techniques have been outlined in the manuscript. We provide O(10^5) SSH snapshots of mesoscale turbulence and corresponding domain-averaged eddy heat fluxes as simulated by the two-layer QG model and split into training and validation data; the data and Python/TensorFlow scripts including neural network architectures graphs and hyperparameters that reproduce our training results can be downloaded here: https://drive.google.com/drive/u/1/folders/1-nZ4FQKkytlwjZRhBqZq-g1-TsooicKk. If additional data is needed, the QG model that was used to generate the samples is available upon request from the authors.

Methods

Numerical simulations of geostrophic turbulence.

Baroclinic turbulence was simulated using the two-layer QG equations with a vertically sheared and horizontal uniform background mean flow that is kept constant in time: an idealised view of baroclinic instability known as the Phillips model[2,27]. The model assumes conservation of potential vorticity in both layers with their anomalies from the mean state, $q_{1,2}$, defined as

$$q_1 = \nabla^2\psi_1 + \frac{f_0^2}{g' H_1}(\psi_2 - \psi_1), \qquad q_2 = \nabla^2\psi_2 + \frac{f_0^2}{g' H_2}(\psi_1 - \psi_2). \tag{1}$$

The time evolution of potential vorticity anomalies is governed by lateral advection due to eddies and mean flow:

$$\partial_t q_i + U_i \nabla q_i + \beta v_i = -r_{Ek}\delta_{i2}\nabla\psi_i, \qquad i = \{1,2\}, \tag{2}$$

where ψ_i is the perturbation streamfunction defined by

$$\begin{pmatrix} u_i \\ v_i \end{pmatrix} = \begin{pmatrix} -\partial_y\psi_i \\ \partial_x\psi_i \end{pmatrix} \tag{3}$$

where u and v are the *zonal* and *meridional* components of velocity and r_{Ek} represents the Ekman drag acting on the bottom layer. Simulations are performed for characteristic parameters of a mid-latitude baroclinic current such as the Gulf Stream: Coriolis parameter $f = f_0 + \beta y$ (with parameters f_0 and β evaluated at 40 degree latitude), the stratification parameters are chosen to result in a baroclinic Rossby deformation radius of 40 km, and the ratio of the top to bottom layer thickness is chosen to be 5. The background mean flow is uniform and constant in time, with the vertical shear of $U_1 - U_2$=0.2 m s^{-1} being sufficiently large to develop baroclinic instabilities that reinforce generation of strongly-interacting mesoscale eddies equilibrating to be O(200 km) in diameter (Figure 1a). The model dissipation is due to the bottom Ekman drag (10 day timescale), while small-scale vorticity gradients are arrested by a scale-dependent dissipation implemented as a filter in spectral space that damps high wavenumber energy in all model variables each time a Fourier transform is used to evaluate tendencies. These specific parameters were chosen so as to give statistically steady 'weak β-plane turbulence' corresponding to a midlatitude ocean. There is no reason to believe our machine learning method would not work equally well for other parameter sets, so long as the nature of the turbulence does not change.

The doubly-periodic domain was set to 4000 km in horizontal scale and the model equations are solved in spectral space using 256 Fourier modes in both directions. To ensure that the double-periodic boundary conditions do not affect the results we divide the full domain into 16 subdomains, 1000x1000km each, before using SSH snapshots for CNN training; the subdomain SSH snapshots used for training are no longer periodic and hence the corresponding eddy heat fluxes have a contribution due to eddies at the boundaries. In total, 112,000 training data images were obtained from 4 independent simulations, each initialised with independent noisy initial conditions and summing up to about 3000 years of model time (spinup data was discarded to avoid dependence of initial conditions). Since SSH snapshots decorrelate from themselves over a timescale of 20 days, the time gap between successive training snapshots was set to 10 days to avoid redundant data and to be more in line with real altimetry

data from satellites that have return periods of O(10 days). 16,000 test data images were produced from an independently seeded simulation each separated from the next by 6 hours to evaluate the smoothness of the obtained CNN mapping.

The eddy heat fluxes, HF, are defined as $HF = \overline{v_1 h_1}$, where the overline corresponds to averaging over the subdomain area, $v_1 = \partial_x \psi_1$ is the anomalous surface ocean velocity in the meridional direction perpendicular to the mean flow, and $h_1 = (f_0/g') \cdot (\psi_2 - \psi_1)$ is the thermocline depth perturbation. Note that SSH perturbations are directly related to the surface geostrophic streamfunction as

$$SSH = \frac{f_0}{g'} \psi_1, \tag{4}$$

and hence the eddy heat flux can be split into the 'trivial' component that only depends on the known SSH field and the 'coupled' component that depends on the unknown bottom layer streamfunction:

$$HF = \frac{f_0}{g'} \overline{(\psi_2 - \psi_1)\partial_x \psi_1} = \underbrace{\overline{\psi_2 \partial_x SSH}}_{\text{'Coupled flux'}} - \underbrace{\frac{g'}{f_0} \overline{SSH\, \partial_x SSH}}_{\text{'Trivial flux'}}. \tag{5}$$

The 'trivial' component of the eddy heat flux exists solely due to eddies passing through the artificially-defined subdomain boundaries and it is identically zero in a periodic domain. Not only is this component dependent only on known SSH field, and hence is trivially calculated, but it is dynamically irrelevant in the sense that it is a noisy term, highly dependent on the location of subdomain boundaries rather than on fundamental processes going on inside it. Also note that surface and deep flow streamfunction (ψ_1 and ψ_2) are significantly correlated[24,43], such that even in the baroclinically-unstable flow considered here the linear correlation explains about half of subsurface flow variance. However, given that the dynamically-relevant component of the eddy heat flux (the 'coupled flux') depends on the average of a product between ψ_1 and a horizontal derivative of ψ_2, it is clear that a component of ψ_2 that is proportional to ψ_1 only provides a dynamically-irrelevant noisy contribution to the heat flux and does not reflect the intensity of baroclinic instabilities. Thus, a dynamically-meaningful heat flux exists only due to a component of ψ_2 that is decorrelated from ψ_1. The 'coupled' flux is also affected by the boundary effects but it nonetheless contains the critical contribution from the fluxes emerging due to baroclinic instability. We thus focus on the prediction of the 'coupled flux' from SSH snapshots, noting that the 'trivial' component could be exactly calculated from SSH data and added if necessary; we choose not to include the 'trivial' component in the calculation of the prediction skill because this would artificially increase it. For any given subdomain, SSH snapshots and corresponding eddy heat fluxes are then used as training input and output for data-driven mapping methods. Importantly, we aim to predict the *instantaneous* flux given an *instantaneous* SSH snapshot and, although the CNN results are evaluated on a continuous timeseries (Figure 3d), all the training points are treated as independent and our method in no way attempts to forecast the time-evolution of SSH.

CNN architecture and performance measures

The optimal CNN architecture used in this study is schematically shown in Figure 2, consisting of 4 pooling and 3 convolutional layers followed by two fully connect layers, with the ReLu (rectified linear unit) used as a non-linear activation function as it outperformed the sigmoid and a hyperbolic tangent functions. The hyperparameters were chosen to optimise the network for the task of flux reconstruction, specifically: the convolution matrices had horizontal dimensions of 4x4 and gradient descent was achieved using Kingma and Ba's AdamOptimizer algorithm[57] with default training rate of 0.001. To reduce overfitting, the dropout with a probability of 30% was implemented between the first and second fully connected layers. The CNN was set to minimise the loss function chosen as mean squared error between the true flux $\mathbf{y}_t$ and the CNN prediction $\mathbf{y}_p$. Training proceeded in batches of 100 snapshots and was deemed complete when the performance of the network on the entire test set was not seen to improve over three epochs. The network was coded in Python using Google's machine learning package TensorFlow.

To evaluate the performance of the CNN and other data-driven methods, we use the skill, S, and the correlation coefficient, R, defined as:

$$S = 1 - \left(\frac{\frac{1}{N} \sum_{i=1}^{N} (\mathbf{y}_{p,i} - \mathbf{y}_{t,i})^2}{\sigma_{\mathbf{y}_t}^2} \right)^{\frac{1}{2}} \tag{6}$$

$$R^2 = \left(\frac{\frac{1}{N} \sum_{i=1}^{N} (\mathbf{y}_{p,i} - \overline{\mathbf{y}}_p)(\mathbf{y}_{t,i} - \overline{\mathbf{y}}_t)}{\sigma_{\mathbf{y}_p} \sigma_{\mathbf{y}_t}} \right)^2, \tag{7}$$

where $\sigma_{\mathbf{y}_t}, \sigma_{\mathbf{y}_p}$ is the standard deviation of the true and predicted eddy heat flux $\mathbf{y}_t, \mathbf{y}_p$. The skill and the regression coefficient both approach 1 for a perfect prediction, however there are important differences in the interpretation of these metrics. The

skill, a monotonically decreasing function of the loss, can be negative if the prediction is worse than the data persistence test, i.e. predicting the average of the eddy heat fluxes corresponds to a zero skill. The square of the correlation coefficient, R^2, provides a useful measure of a fraction of variance that is explained by the prediction, but it in some cases fails to be a reliable measure of accuracy as R is insensitive to shifts in the mean or multiplication by a constant multiple. Throughout the paper we specify both metrics.

Since NN training involves stochasticity in defining its initial parameters and during their optimization, we defined their performance based on an average metrics in the following way. First, we evaluate NN skill on validation data (10 times per epoch) and obtain top 100 results. Second, we evaluate the skill on the entire test data using each of the CNN model parameters corresponds to the top 100 validation skills. The average of the test skill and its standard deviation is the one we report in our study. The choice of using averages over top 100 validation skills biases our skill metrics slightly lower compared to the maximum skill, but the difference is only about 5%.

Benchmarks and additional tests.

We compare CNN performance to a number of more standard statistical techniques and summarise the results in Figure 4. The methods include:

- **Linear Regression.** First we assume $\psi_2 = \psi_1$ then perform simple linear regression on the predicted flux. To first order figure 1b shows these two fields are proportional. Given the estimated ψ_2 We calculate the eddy heat flux, which in a non-periodic domain doesn't have to be zero even if ψ_2 is proportional to ψ_1.

- **Principal Component Analysis (PCA).** By finding the PCA basis set for concatenated training-ψ_1 & ψ_2 snapshots and retaining an optimal number of modes, test-ψ_1 images can be used to find an estimate for their corresponding ψ_2 field, from which the eddy heat flux is found. PCA is also known under the names of Proper Orthogonal Decomposition or Empirical Orthogonal Functions.

- **Support Vector Machine (SVM).** Regression with a radial basis function kernel. SVMs are an early but effective form of supervised machine learning good at classification and regression.

- **Random Forest Regression.** We implement random forest with 75 trees estimators. Another commonly used machine learning algorithm for regression problems.

- **Fully Connected Neural Networks (FC NN).** We use a basic neural network with 2 hidden layers of 100 and 10 neurons respectively, ReLU activation, mean square error as the loss function, and no dropout. This is a basic form of deep supervised learning which treats the input images as flattened vectors. The results do not significantly change if a higher number of neurons is used.

- **Convolutional Neural Networks (CNNs).** These networks have the advantage of explicitly treating the input as images (spatially ordered data) by applying convolutional filters with adjustable parameters. Here we show the best results from **VGG**-type[58] and **ResNet**-type[56] architectures. Our **VGGNet** architectures are of various complexity depending on the number of convolutional filters and number of neurons used in dense layers. The hyperparameters for the CNN architectures referred by the numbers in Figure 3:

  (1) two 3x3 convolutional layers (4 filters each) and 2x2 max-pooling layers followed by a fully connected layer;

  (2) single 4x4 convolutional layer (8 filters) and a single max-pooling (4x4 strides, 4x4 poolsize) with no hidden dense layers

  (3) three 4x4 convolutional layers and 2x2 max-pooling layers followed by a hidden layer with 10 neurons;

  (4) three 4x4 convolutional layers (8,16, and 32 filters) and corresponding 2x2 max-pooling layers (4x4 poolsize) followed by a hidden layer with 128 neurons;

  (5) sixteen 4x4 convolutional layers and three 2x2 max-pooling layers followed by a hidden layer with 128 neurons (similar to the VGG16 architecture);

  (6) five 4x4 convolutional layers (8,16,32,64, and 128 filters) and corresponding 2x2 max-pooling layers (4x4 poolsize) followed by a hidden layer with 128 neurons;

  (7) corresponds to the **ResNet** architecture containing 52 layers (50 convolutional and two dense layers) with skip connections as described in[56]. Dropout of 30% was used on the last fully connected layer in all networks. The network graphs can be downloaded as png-files following the data link in the Data Availability section.

A range of additional tests were performed to confirm that a skill of 0.36 is approximately an upper bound, corresponding to a correlation with true heat flux of ≈ 0.8 with $r^2 \approx 0.6$. These test included applying filters to reduce the flux contribution from eddies crossing subdomain boundaries, increasing and reducing the QG model resolution, increasing the subdomain size to 2000km, reducing the subdomain size to 500 km, exploring the CNN architecture by changing the number of filters and fully connected layers, oversampling the eddy flux outliers to obtain a more uniform distribution among the training dataset. Without exceptions, the average test skill obtained over the final 3 epochs of training was never above 0.36, indicating that this may be a dynamically-constrained upper bound on the information contained in SSH snapshots.

References

1. Eady, E. T. Long waves and cyclone waves. *Tellus* **1**, 33–52 (1949).
2. Phillips, N. A. Energy transformations and meridional circulations associated with simple baroclinic waves in a two-level, quasi-geostrophic model. *Tellus* **6**, 274–286 (1954).
3. Charney, J. G. The dynamics of long waves in a baroclinic westerly current. *J. Meteorol.* **4**, 136–162 (1947).
4. Chelton, D. B., Deszoeke, R. A., Schlax, M. G., El Naggar, K. & Siwertz, N. Geographical variability of the first baroclinic rossby radius of deformation. *J. Phys. Oceanogr.* **28**, 433–460 (1998).
5. Nurser, A. & Bacon, S. The rossby radius in the arctic ocean. *Ocean. Sci.* **10**, 967–975 (2014).
6. Wunsch, C. Where do ocean eddy heat fluxes matter? *J. Geophys. Res. Ocean.* **104**, 13235–13249 (1999).
7. Jayne, S. R. & Marotzke, J. The oceanic eddy heat transport. *J. Phys. Oceanogr.* **32**, 3328–3345 (2002).
8. Volkov, D. L., Lee, T. & Fu, L.-L. Eddy-induced meridional heat transport in the ocean. *Geophys. Res. Lett.* **35** (2008).
9. Taylor, K. E., Stouffer, R. J. & Meehl, G. A. An overview of cmip5 and the experiment design. *Bull. Am. Meteorol. Soc.* **93**, 485–498 (2012).
10. Eyring, V. *et al.* Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization. *Geosci. Model. Dev.* **9**, 1937–1958 (2016).
11. Lorenz, E. N. Energy and numerical weather prediction. *Tellus* **12**, 364–373 (1960).
12. Gent, P. R. & Mcwilliams, J. C. Isopycnal mixing in ocean circulation models. *J. Phys. Oceanogr.* **20**, 150–155 (1990).
13. Roemmich, D. *et al.* The argo program: Observing the global ocean with profiling floats. *Oceanography* **22**, 34–43 (2009).
14. Chelton, D. B., Ries, J. C., Haines, B. J., Fu, L.-L. & Callahan, P. S. Satellite altimetry. In *International geophysics*, vol. 69, 1–ii (Elsevier, 2001).
15. Chelton, D. B., Schlax, M. G., Samelson, R. M. & de Szoeke, R. A. Global observations of large oceanic eddies. *Geophys. Res. Lett.* **34** (2007).
16. Zhai, X., Johnson, H. L. & Marshall, D. P. Significant sink of ocean-eddy energy near western boundaries. *Nat. Geosci.* **3**, 608 (2010).
17. Wunsch, C. The work done by the wind on the oceanic general circulation. *J. Phys. Oceanogr.* **28**, 2332–2340 (1998).
18. Scott, R. B. & Wang, F. Direct evidence of an oceanic inverse kinetic energy cascade from satellite altimetry. *J. Phys. Oceanogr.* **35**, 1650–1666 (2005).
19. Stammer, D. On eddy characteristics, eddy transports, and mean flow properties. *J. Phys. Oceanogr.* **28**, 727–739 (1998).
20. Abernathey, R. P. & Marshall, J. Global surface eddy diffusivities derived from satellite altimetry. *J. Geophys. Res. Ocean.* **118**, 901–916 (2013).
21. Menemenlis, D. *et al.* Ecco2: High resolution global ocean and sea ice data synthesis. *Mercat. Ocean. Q. Newsl.* **31**, 13–21 (2008).
22. Keating, S. R., Majda, A. J. & Smith, K. S. New methods for estimating ocean eddy heat transport using satellite altimetry. *Mon. Weather. Rev.* **140**, 1703–1722 (2012).
23. Chapman, C. & Charantonis, A. A. Reconstruction of subsurface velocities from satellite observations using iterative self-organizing maps. *IEEE Geosci. Remote. Sens. Lett.* **14**, 617–620 (2017).
24. de La Lama, M. S., LaCasce, J. H. & Fuhr, H. K. The vertical structure of ocean eddies. *Dyn. Stat. Clim. Syst.* **1**, DOI: 10.1093/climsys/dzw001 (2016). Dzw001, http://oup.prod.sis.lan/climatesystem/article-pdf/1/1/dzw001/9629836/dzw001.pdf.

25. Bolton, T. & Zanna, L. Applications of deep learning to ocean data inference and sub-grid parameterisation. *EarthArXiv* (2018).

26. Charney, J. G. Geostrophic turbulence. *J. Atmospheric Sci.* **28**, 1087–1095 (1971).

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

28. McCulloch, W. S. & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The bulletin mathematical biophysics* **5**, 115–133 (1943).

29. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q. (eds.) *Advances in Neural Information Processing Systems 25*, 1097–1105 (Curran Associates, Inc., 2012).

30. Russakovsky, O. *et al.* Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015).

31. Khan, J. *et al.* Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks. *Nat. medicine* **7**, 673 (2001).

32. Way, G. P. & Greene, C. S. Extracting a biologically relevant latent space from cancer transcriptomes with variational autoencoders. *bioRxiv* 174474 (2017).

33. Dieleman, S., Willett, K. W. & Dambre, J. Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Mon. notices royal astronomical society* **450**, 1441–1459 (2015).

34. Khalifa, N. E. M., Taha, M. H. N., Hassanien, A. E. & Selim, I. Deep galaxy: Classification of galaxies based on deep convolutional neural networks. *arXiv preprint arXiv:1709.02245* (2017).

35. Ling, J., Kurzawski, A. & Templeton, J. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155–166 (2016).

36. San, O. & Maulik, R. Extreme learning machine for reduced order modeling of turbulent geophysical flows. *Phys. Rev. E* **97**, 042322 (2018).

37. San, O. & Maulik, R. Neural network closures for nonlinear model order reduction. *Adv. Comput. Math.* 1–34 (2018).

38. Pearson, K. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, Dublin Philos. Mag. J. Sci.* **2**, 559–572 (1901).

39. Schmid, P. J. Dynamic mode decomposition of numerical and experimental data. *J. fluid mechanics* **656**, 5–28 (2010).

40. Kutz, J. N., Brunton, S. L., Brunton, B. W. & Proctor, J. L. *Dynamic mode decomposition: data-driven modeling of complex systems*, vol. 149 (SIAM, 2016).

41. Rahman, S., San, O. & Rasheed, A. A hybrid approach for model order reduction of barotropic quasi-geostrophic turbulence. *Fluids* **3**, 86 (2018).

42. Manucharyan, G. E., Thompson, A. F. & Spall, M. A. Eddy memory mode of multidecadal variability in residual-mean ocean circulations with application to the beaufort gyre. *J. Phys. Oceanogr.* **47**, 855–866 (2017).

43. Wunsch, C. The vertical partition of oceanic horizontal kinetic energy. *J. Phys. Oceanogr.* **27**, 1770–1794 (1997).

44. Smith, K. S. & Vallis, G. K. The scales and equilibration of midocean eddies: Freely evolving flow. *J. Phys. Oceanogr.* **31**, 554–571 (2001).

45. Fu, L.-L. & Flierl, G. R. Nonlinear energy and enstrophy transfers in a realistically stratified ocean. *Dyn. Atmospheres Ocean.* **4**, 219–246 (1980).

46. Scott, R. B. & Arbic, B. K. Spectral energy fluxes in geostrophic turbulence: Implications for ocean energetics. *J. physical oceanography* **37**, 673–688 (2007).

47. Pedamonti, D. Comparison of non-linear activation functions for deep neural networks on mnist classification task. *arXiv preprint arXiv:1804.02763* (2018).

48. Hogg, N., Stommel, H. M. *et al.* The heton, an elementary interaction between discrete baroclinic geostrophic vortices, and its implications concerning eddy heat-flow. *Proc. R. Soc. Lond. A* **397**, 1–20 (1985).

49. Chollet, F. *et al.* Keras. https://keras.io (2015).

50. Thrun, S. Lifelong learning algorithms. In *Learning to learn*, 181–209 (Springer, 1998).

51. Vilalta, R. & Drissi, Y. A perspective view and survey of meta-learning. *Artif. intelligence review* **18**, 77–95 (2002).

52. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D. & Lillicrap, T. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, 1842–1850 (2016).

53. Pan, S. J. & Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge data engineering* **22**, 1345–1359 (2010).

54. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).

55. Karazija, L., Velickovic, P. & Liò, P. Automatic inference of cross-modal connection topologies for x-cnns. *CoRR* **abs/1805.00987** (2018). 1805.00987.

56. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. *CoRR* **abs/1512.03385** (2015). 1512.03385.

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

58. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

Acknowledgements

The authors gratefully acknowledge support from Charlie Trimble as well as the David and Lucile Packard Foundation. This work was largely completed as part of Caltech's Summer Undergraduate Research Fellowship Program (SURF), and we thank the SURF staff for their assistance. This study benefited from conversations with Frederick Eberhardt and RJ Antonello. Glenn Flierl provided the QG turbulence code used in this study.

Author contributions statement

G.E.M conceived the study; T.M.G and G.E.M. performed the research; all authors analysed the results and contributed to writing of the paper; the research was supervised by A.F.T and G.E.M.

Additional information

The authors declare no competing interests.