



HAL
open science

Synthesizing Sea Surface Temperature and Satellite Altimetry Observations Using Deep Learning Improves the Accuracy and Resolution of Gridded Sea Surface Height Anomalies

Scott A. Martin, Georgy E. Manucharyan, Patrice Klein

► **To cite this version:**

Scott A. Martin, Georgy E. Manucharyan, Patrice Klein. Synthesizing Sea Surface Temperature and Satellite Altimetry Observations Using Deep Learning Improves the Accuracy and Resolution of Gridded Sea Surface Height Anomalies. *Journal of Advances in Modeling Earth Systems*, 2023, 15, 10.1029/2022MS003589 . insu-04195519

HAL Id: insu-04195519

<https://insu.hal.science/insu-04195519v1>

Submitted on 4 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License



RESEARCH ARTICLE

10.1029/2022MS003589

Key Points:

- We developed a deep learning method that significantly improves the accuracy and resolution of gridded sea surface height anomalies
- This data-driven method takes advantage of combining sea surface temperature and altimetry observations
- Inferred surface geostrophic currents are quantitatively and qualitatively more realistic than those from existing sea surface height maps

Supporting Information:

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

Correspondence to:

S. A. Martin,
smart1n@uw.edu

Citation:

Martin, S. A., Manucharyan, G. E., & Klein, P. (2023). Synthesizing sea surface temperature and satellite altimetry observations using deep learning improves the accuracy and resolution of gridded sea surface height anomalies. *Journal of Advances in Modeling Earth Systems*, 15, e2022MS003589. <https://doi.org/10.1029/2022MS003589>

Received 16 DEC 2022

Accepted 27 APR 2023

Synthesizing Sea Surface Temperature and Satellite Altimetry Observations Using Deep Learning Improves the Accuracy and Resolution of Gridded Sea Surface Height Anomalies

Scott A. Martin¹ , Georgy E. Manucharyan¹ , and Patrice Klein^{2,3}

¹School of Oceanography, University of Washington, Seattle, WA, USA, ²Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, ³LMD/IPSL, CNRS, Ecole Normale Supérieure, PSL Research University, Paris, France

Abstract Gridded sea surface height (SSH) maps estimated from satellite altimetry are widely used for estimating surface ocean geostrophic currents. Satellite altimeters observe SSH along one-dimensional tracks widely spaced in space and time, making accurately reconstructing the two-dimensional (2D) SSH field challenging. Traditionally, SSH is mapped using optimal interpolation (OI). However, OI artificially smooths the SSH field leading to high mapping errors in regions with rapidly-evolving mesoscale features such as western boundary currents. Motivated by the dynamical relation between SSH and sea surface temperature (SST) and the notion that even the chaotic evolution of mesoscale ocean turbulence may contain repeating patterns, we outline a deep learning (DL) approach where a neural network is trained to reconstruct 2D SSH by synthesizing altimetry and SST observations. In the Gulf Stream Extension region, dominated by mesoscale variability, our DL method substantially improves the SSH reconstruction compared to existing methods. Our SSH map has 17% lower root-mean-square error and resolves spatial scales 30% smaller than OI compared against independent altimeter observations. Surface geostrophic currents calculated from our map are closer to surface drifter observations and appear qualitatively more realistic, with stronger currents, a clearer separation between the Gulf Stream and neighboring eddies, and the appearance of smaller coherent eddies missed by other methods. Our map yields significant re-estimations of important dynamical quantities such as eddy kinetic energy, vorticity, and strain rate. Applying our DL method to produce a global SSH product may provide a more accurate and higher resolution product for studying mesoscale ocean turbulence.

Plain Language Summary Satellites observe small variations in the height of the sea surface but with large gaps in the observations. Having an estimate of the two-dimensional sea surface height field allows one to estimate surface ocean currents, so filling in the gaps between the observations is an important problem. The traditionally-used method for filling in the gaps between sea surface height observations struggles when there are lots of small-scale, rapidly-interacting ocean currents. We developed a deep learning (DL) model to estimate the sea surface height field more accurately. We achieved this by combining the sea surface height observations with satellite observations of sea surface temperature. The relationship between sea surface temperature and height is non-trivial, but our DL model learned to use information from the sea surface temperature observations in the places where sea surface height wasn't observed to improve the accuracy of the sea surface height estimate. We applied and tested our method in the Gulf Stream and demonstrated that our sea surface height map is more accurate than that from the traditional method. Our method provides a more accurate sea surface height map which could allow us in future to learn new lessons about small-scale surface currents in the ocean.

1. Introduction

1.1. Background

Sea surface height (SSH) maps—that is, time-varying gridded maps of the height of the ocean's surface referenced against the geoid—derived from satellite altimetry observations have been a crucial tool for physical oceanographers for many years (Abdalla et al., 2021; Fu et al., 2010). Under the assumption of geostrophy, SSH is proportional to a streamfunction for the surface currents. Hence a SSH map provides easy access to an estimated map of surface geostrophic ocean currents, which would otherwise be challenging to obtain through direct in situ observations with global coverage. SSH maps have been used to directly observe the inverse kinetic

energy cascade in the global oceans (Scott & Wang, 2005), to track and study the behavior of mesoscale eddies (Chelton et al., 2011; Fu et al., 2010), and to estimate surface eddy diffusivities (Abernathey & Marshall, 2013).

To date, the only global SSH observations have come from satellite-borne nadir radar altimeters, which measure SSH along one-dimensional tracks directly beneath the satellite's orbit resolving along-track scales of $O(20\text{--}100\text{ km})$ (Dufau et al., 2016). Reconstructing the full two-dimensional (2D) SSH field from altimeter observations requires significant interpolation in space and time. Hence, the length-scales 2D products can accurately resolve are coarser than the along-track resolution of the altimeters (Ballarotta et al., 2019).

The most widely used 2D SSH product is that generated by the “Data Unification and Altimeter Combination System” (DUACS) (Taburet et al., 2019) and distributed by the “Copernicus Marine Environment Monitoring Service” (CMEMS) (*note that this is the same product as was formerly distributed by the “Archiving, Validation and Interpretation of Satellite Oceanographic data” (AVISO) service*). This product is created using the optimal interpolation (OI) method (Bretherton et al., 1976; Le Traon et al., 1998), which is otherwise known as “objective mapping” or “objective analysis.” OI provides the best least squares linear estimator of the SSH in the gaps between observations, given an a priori model for the covariance of SSH between different locations and times, and knowledge of the instrument noise covariance. In the DUACS mapping, the SSH covariance is assumed to be a single-scale Gaussian with prescribed decorrelation length- and time-scales that are geographically varying and have been tuned empirically to account for a varying number of available satellite altimeters over the 30 years altimetry record, while the instrument noise covariance has been chosen to be suitable for each satellite altimeter mission and technology (Taburet et al., 2019).

While the DUACS SSH product continues to be valuable to the oceanography community, the OI method has been shown to introduce significant deficiencies. Amores et al. (2018) used an Observer System Simulation Experiment (OSSE) in which a DUACS-like SSH product was generated using pseudo-observations from the output of a high-resolution ocean model to study how accurately DUACS captured the model's mesoscale SSH field. They demonstrated that the eddy fields inferred from the DUACS SSH map are significantly distorted, with often multiple smaller eddies being aliased into larger eddies. The real-world accuracy of an SSH mapping method can also be studied by generating a map with one satellite altimeter's observations withheld and examining the errors of the mapped SSH compared to these independent along-track observations. In this way, Ballarotta et al. (2019) showed that the DUACS product only accurately resolves SSH signals down to time-scales of $O(30\text{ days})$ and length-scales $O(100\text{ km})$ at high latitudes, increasing to $O(800\text{ km})$ in the tropics. OI tends to smooth out small-scale, fast-evolving features and strong peaks/troughs in SSH where observations are scarce and the oceanic mesoscale is energetic.

The challenges involved in reconstructing surface currents from nadir satellite altimetry observations has led to the development of a number of methods for inferring surface ocean currents from satellite observations of other surface ocean quantities, most notably sea surface temperature (SST). Satellite-borne infrared sensors allow the observation of 2D snapshots of SST with high spatial resolution of $O(1\text{ km})$ when no clouds obscure the sensors. Since submesoscale SST features with scales $O(1\text{--}10\text{ km})$ are advected by mesoscale surface ocean currents, tracking the motion of submesoscale SST features in consecutive cloud-free SST observations allows for the estimation of the mesoscale surface currents (Bowen et al., 2002; Emery et al., 1986). Others have attempted to use a sequence of cloud-free SST observations to invert the heat conservation equation to estimate the surface currents which enter this equation through the advection term (Chen et al., 2008; Kelly, 1989; Ostrovskii & Piterbarg, 1995). Rio et al. (2016) combined the heat conservation approach with background geostrophic surface current estimates from DUACS to reduce the uncertainty in along-isotherm velocities, which cannot be estimated from SST tracer advection. The heat conservation approach requires the approximation of the main source and sink terms in the heat conservation equation. These methods allow for the retrieval of geostrophic currents smaller than the scales resolved by DUACS as well as ageostrophic currents that cannot be retrieved directly from SSH. However, the need for consecutive cloud-free SST observations severely limits their capability for consistently estimating surface currents across all regions and seasons. Another approach utilizes surface quasi-geostrophic (SQG) theory to estimate surface currents from a single SST snapshot and has been applied with promising results to microwave SST observations, which have lower spatial resolution than infrared but do not suffer gaps due to clouds (Isern-Fontanet et al., 2006, 2014). However, as will be discussed further in Section 2.1, this approach relies on particular environmental conditions that are not universally met in all regions and seasons. While the retrieval of surface ocean currents from SST observations is a dynamic and growing research field,

SSH maps from satellite altimetry along with their associated surface geostrophic currents remain the primary source of surface ocean current estimates for the ocean science community.

The surface currents derived from the DUACS SSH product are used extensively in oceanographic studies, so it is essential to ensure these inferred currents are as accurate and high resolution as possible using the existing satellite observing capabilities. There is therefore increasing focus on developing better methods for reconstructing the 2D SSH field from satellite altimetry observations (Fablet, Amar, et al., 2021; Le Guillou et al., 2021; Manucharyan et al., 2021; Ubelmann et al., 2015, 2021). Efforts to improve SSH mapping methods have been aided by the creation in recent years of community-maintained “Ocean Data Challenges” (AVISO, 2021). By establishing a common set of evaluation metrics, a common study region and time, and by sharing code and results these data challenges allow for direct comparison between different SSH mapping methods.

OI can only provide the statistical best *linear* estimator of the unobserved SSH for a given a priori covariance model. While the decorrelation length- and time-scales used in this covariance model can be tuned to best fit the data, as is done to create the DUACS product, a single-scale Gaussian covariance model and its resulting linear estimator is ultimately limited in its ability to accurately represent the non-linear dynamics of mesoscale ocean turbulence. An explicitly dynamics-based approach to SSH mapping would be to use a data assimilation framework to constrain an ocean circulation model to best match the available observations, then use this constrained model's SSH field as the estimate of 2D SSH. However, such an approach requires additional observations to accurately constrain other essential model variables such as subsurface flow and density which are not typically available at high resolution. Data assimilation using a full global ocean circulation model is very computationally expensive, thus it is only in recent years that data assimilation with mesoscale eddy-permitting model resolution ($1/12^\circ$) has become feasible (Lellouche et al., 2021). Despite the computational expense of $1/12^\circ$ ocean model data assimilation, even higher model resolutions may be required before such data assimilation approaches can accurately resolve the mesoscale eddy field. To circumvent the computational challenges of using a full primitive equation ocean model, an approach that uses data assimilation with an idealized one layer quasi-geostrophic (QG) model of ocean turbulence has been effectively applied to SSH mapping (Le Guillou et al., 2021). Such an idealized model, however, is unable to capture many of the dynamical effects observed in the real-world ocean.

The recent advances in the application of deep learning (DL) methods to the earth sciences (Sonnewald et al., 2021; Sun et al., 2022) have inspired an increasing focus on the possibility of using data-driven interpolation methods to improve SSH mapping (Beauchamp et al., 2020; Manucharyan et al., 2021). DL models can be trained to approximate highly non-linear mappings from inputs to outputs. Thus it is plausible that a DL model could be trained to recognize dynamical signatures in partial SSH observations and use these to estimate SSH in unobserved regions given sufficient training examples. By training a DL model on a large library of examples of real ocean turbulence the DL model would be able to implicitly use ocean dynamics to create a more accurate SSH mapping than would be possible with OI.

1.2. Previous Work Applying DL to SSH Mapping

A number of attempts to use DL for SSH mapping have been published recently (Barth et al., 2022; Buongiorno Nardelli et al., 2022; Fablet, Amar, et al., 2021; Fablet, Chapron, et al., 2021). These studies have mostly focused on the OSSE setting, where the SSH observations are generated from a circulation model in which the full 2D SSH field is known. Fablet, Amar, et al. (2021) and Fablet, Chapron, et al. (2021) developed a DL approach (4DVarNet) for interpolating SSH which formulates SSH mapping as a four-dimensional variational data assimilation (4DVar) problem but replaces the dynamical model and solver with trainable neural networks. They demonstrated that 4DVarNet yields a significant improvement over OI when tested in idealized OSSEs, though the improvement is more modest when applied to real-world observations.

While these results are auspicious, OSSEs can only offer a limited approximation of the real-world SSH interpolation problem since the full 2D SSH ground truth is unknown in real-world settings. It therefore cannot be used to calculate the reconstruction error to be minimized during training. A number of studies have attempted to address this issue by developing methods that can be trained exclusively on real-world observations in a manner similar to that described here in Section 4.2. Barth et al. (2020) introduced a Convolutional Neural Network (CNN) method, Data INterpolating Convolutional Auto-Encoder (DINCAE), for reconstructing sea surface temperature (SST) from partial satellite observations. In a subsequent paper, Barth et al. (2022) refined

and applied the method to reconstruct SSH in the Mediterranean from partial real-world SSH and SST observations. They found that including SST yielded a lower reconstruction error than SSH observations alone. However, an error comparison to the operational DUACS method was not made, so this method's performance is hard to compare to other SSH mapping methods. Buongiorno Nardelli et al. (2022) developed another CNN approach for reconstructing and super-resolving SSH from partial SSH and SST satellite observations in the Mediterranean with promising results. However, to split the observations into independent training and testing data sets they randomly selected days to withhold for testing, meaning adjacent days could be used for training and testing. SSH and SST fields contain significant temporal autocorrelation over time scales of days to weeks. Hence, such a sampling strategy leads to a testing data set that could be highly correlated to the training data, calling into question the independence of the performance metrics calculated.

To date, DL SSH mapping studies remain limited to theoretical OSSE studies or limited regional experiments and the gridded SSH product of choice for oceanographers remains the OI-generated DUACS product.

1.3. Problem Statement and Paper Structure

In this study, we seek to demonstrate in practice that DL can provide a more accurate SSH mapping with a more physically realistic eddy evolution than was previously possible. Further, our DL approach allows us to include co-located observations of other ocean surface variables in the input to the mapping, something that is challenging with existing SSH mapping methods. Specifically, we will demonstrate that the inclusion of SST observations as an additional source of information for the SSH mapping leads to dramatic improvements in the accuracy and resolution of the reconstructed SSH evolution. In Section 2, we provide a dynamical motivation for why SST observations are expected to improve SSH mapping and provide the rationale for our DL approach. Our method is both trained on and tested against real-world satellite observations (the datasets we use are presented in Section 3). We present our mapping method in Section 4 and how it performed during training in Section 5. In Section 6, we present the results of a regional experiment to quantitatively compare our method's accuracy to that of the conventional and several recently proposed SSH mapping methods in the Gulf Stream Extension. In Section 7, we discuss the implications of these results and outline a road-map to creating a more accurate global SSH product with DL.

2. Rationale for Our Proposed Methodology

2.1. Dynamical Motivation for Supplementing SSH Observations With SST

In addition to satellite altimetry observations, satellites routinely observe other surface ocean quantities, such as SST, sea surface salinity, and ocean color. The spatiotemporal evolution of these fields is certainly not independent from one another and from that of SSH, since they are all affected and mediated by the dynamics of the ocean (and by biological processes in the case of ocean color). Thus, if one of these other surface variables is observed in the gaps between satellite altimetry observations, and the relationship to SSH is known, then this can inform the mapping of the unobserved SSH.

In this study, we supplement satellite altimetry with SST observations to improve the SSH mapping for reasons both physical and practical. Physically, this is a natural choice since SSH and SST are known to have a close dynamical relationship. Indeed, in high eddy kinetic energy (EKE) regions (e.g., western boundary currents in the northern hemisphere and the Antarctic Circumpolar Current in the southern hemisphere), temperature explains much of the density field, especially in winter (Isern-Fontanet et al., 2006), indicating that mesoscale SST anomalies can be interpreted as a dynamical part of surface-intensified baroclinic eddies (Smith & Vallis, 2001). Hausmann and Czaja (2012) confirmed these results: SSH and SST anomalies have similar spatial scales and are characterized by a westward shift as expected from baroclinic instability, which is the main source of mesoscale eddies. In addition, a close spectral relationship between mesoscale SSH and SST can be inferred when the large-scale meridional gradient of potential vorticity at depth is linearly related to the large-scale meridional SST gradient through a function only depth-dependent (Lapeyre & Klein, 2006), which is verified in high EKE regions (Lapeyre, 2009). This relationship makes use of the SQG approximation in which the ocean interior dynamics at mesoscale is implicitly taken into account through an “effective” Brunt-Vaisala frequency, N_{eff} (Isern-Fontanet et al., 2006; Klein et al., 2019; LaCasce & Mahadevan, 2006; Lapeyre & Klein, 2006). The resulting relationship in spectral space is

$$\widehat{\text{SSH}}_{\text{SQG}}(\mathbf{k}) = -\frac{f_0}{N_{\text{eff}}|\mathbf{k}|} \alpha \widehat{\text{SST}}(\mathbf{k}), \quad (1)$$

where a “hat” symbol indicates a horizontal Fourier transform, \mathbf{k} the horizontal wavenumber vector, α the thermal expansion coefficient, and f_0 the Coriolis frequency. Isern-Fontanet et al. (2006) proposed to use this relationship to recover SSH and therefore surface ocean currents from SST observations. Practically, satellite observations of SST have higher spatial resolution and lower missing data rates than those of SSH, which emphasizes the pertinence of using SST observations because of their close dynamical relationship with SSH.

Applying Equation 1 to directly reconstruct SSH from SST observations is challenging in that the oceanic conditions described above do not always apply directly in all seasons and ocean regions. Other attempts have used a prescribed analytical relation between SST and SSH like that in Equation 1 to reconstruct the unobserved SSH field (González-Haro & Isern-Fontanet, 2014; Isern-Fontanet et al., 2014). All such methods face the challenge that the dynamical relationship between SSH and SST is non-trivial and is likely to change in space and time. Lapeyre (2009) showed that the SSH field is in some regions dominated by the SQG mode, making Equation 1 applicable, whereas in others the first baroclinic mode dominates, meaning the SSH estimated from Equation 1 would be inaccurate. Thus, methods that rely on the SQG framework (González-Haro & Isern-Fontanet, 2014; Isern-Fontanet et al., 2006, 2014) and methods that rely on a QG framework (Le Guillou et al., 2021; Ubelmann et al., 2015, 2021) will each be geographically and seasonally limited in their ability to reconstruct the SSH field by the pertinence of their respective dynamical assumptions. Another approach to using SST observations to inform the SSH mapping would be to use OI (as in DUACS), replacing the univariate SSH covariance function with a multivariate SSH-SST covariance function. However, the dynamical relationship between SST and SSH is sufficiently non-trivial that such an approach would likely only yield marginal mapping performance improvements.

2.2. Rationale for Deep Learning

DL models (specifically, neural networks) have been shown to be universal function approximators, given sufficient depth and training data (Hornik et al., 1989). By employing a DL approach for SSH mapping, we implicitly make two hypotheses. (a): Available satellite observations contain sufficient dynamical information that there exists a function mapping from the observations to an estimate of the dynamical state of the ocean, and by extension to the 2D SSH field. (b): There exists sufficient training data that we can train a DL model to approximate this function accurately enough to provide better SSH maps than OI. Both are motivated by earlier work applying DL neural networks to spatiotemporal SSH interpolation in a two-layer quasi-geostrophic (QG) model given partial observations of the surface (Manucharyan et al., 2021). George et al. (2021) also showed a DL model could learn to infer subsurface dynamics, which will, in turn, influence SSH evolution, from snapshots of the surface SSH field. The overarching conclusion from the two aforementioned studies is that the surface SSH field contains implicit signatures of the subsurface ocean dynamics and that DL models are sufficiently expressive that they can learn to recognize and exploit these signatures to more accurately reconstruct unobserved ocean variables.

The advantage of a DL approach that combines the SSH and SST fields is that a sufficiently expressive DL model, given enough training data, can, in theory, learn a dynamical relationship of arbitrary complexity. Thus, a DL model could learn to exploit the SST observations in regions and times when the dynamics makes them pertinent, while not being restricted to assuming the SQG approximation is always appropriate. The promise of a DL approach that combines SSH and SST observations to reconstruct the unobserved SSH field more accurately was also noted recently in Fablet and Chapron (2022) and Fablet et al. (2022). However, to date, there are no operationally-used SSH reconstructions that utilize SST observations or DL, and most of the studies on the topic have been theoretical in nature.

3. Data Sets

3.1. Sea Surface Height (SSH) Observations

We use Level 3 1Hz along-track satellite altimetry SSH observations from 2010 to 2020 (distributed by CMEMS, (CMEMS, 2023b)). This product consists of observations from multiple satellite altimetry missions merged,

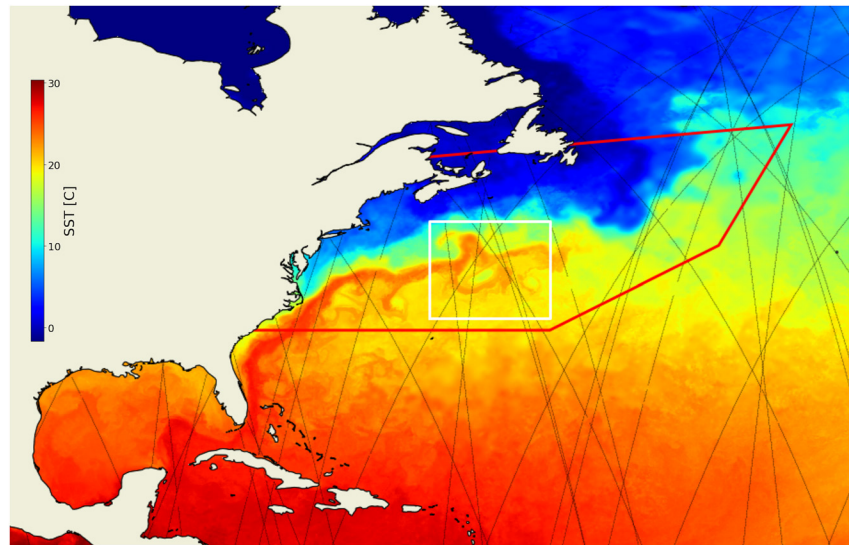


Figure 1. The typical data availability for 1 day (2017-04-11) during our testing period: the color map shows the GHRSSST Multiscale Ultrahigh-Resolution gridded SST, and the black tracks indicate the locations of satellite altimeter observations from all satellites available on that date (CryoSat-2, Jason-2, Jason-3, Haiyang-2A, Sentinel-3A, and SARAL/Altika). As discussed in Section 5, we draw training examples from within the red polygon (the Gulf Stream Extension), and test our method in the white box (where other mapping methods have been applied so we can compare to their accuracy).

calibrated, and corrected for several geophysical effects, including barotropic tides and atmospheric effects. The complete steps for producing this product are described in Taburet et al. (2019) and references therein. Note that at the Level 3 stage of the observation processing chain, the observations have not yet been mapped to a 2D grid, as shown in Figure 1. These along-track observations have been shown to accurately resolve the SSH signal down to length-scales $O(20\text{--}100\text{ km})$ with root-mean-square instrument noise of 1–4 cm depending on the location, season, and satellite altimeter in question (Dufau et al., 2016).

3.2. Sea Surface Temperature (SST) Observations

We use the Level 4 Multiscale Ultrahigh-Resolution (MUR) gridded SST product from the same years provided by the Group for High-Resolution SST (GHRSSST) (NASA PODAAC, 2023). The MUR SST product is generated by combining observations from multiple satellites and in situ sensors before interpolating to a regular grid using OI. Clouds cause gaps in the highest resolution SST satellite sensors, meaning that the OI product has high resolution in cloud-free regions but suffers artificial smoothing in cloud-occluded regions. An alternative approach would be to use the high-resolution observations directly with gaps, or to use only the lower resolution cloud-free observations provided for example, by the NASA Advanced Microwave Scanning Radiometer-EOS (AMSR-E) which would both offer a uniform spatial resolution. Until a uniformly high-resolution gridded SST product is developed, there is an inevitable trade-off between using high-resolution SST with gaps, low-resolution SST without gaps, and an OI product like MUR with non-uniform spatial resolution. Despite its deficiencies, we here use the MUR SST product since it is the highest-resolution gap-free SST product currently available.

3.3. Data Pre-Processing: Domain and Coordinate System

We pursue a patch-based approach, where SSH is reconstructed in a local square domain using observations from only its surroundings. This approach is motivated by the fact that the dynamics governing SSH evolution is predominantly local in space and time, thus estimating the SSH at a point should only require observations from some finite space-time window surrounding that point. Our method requires the input data to be defined on a regular grid, so we bin-average the available SSH and SST observations within the domain onto regular square grids as shown in Figure 2. For this study, we choose the dimensions of the square domain to have a side-length of 960 km with a grid resolution of 7.5 km (128×128 grid points). The input SSH and SST observations are first re-scaled by subtracting the mean and dividing by the standard deviation, before being bin-averaged onto

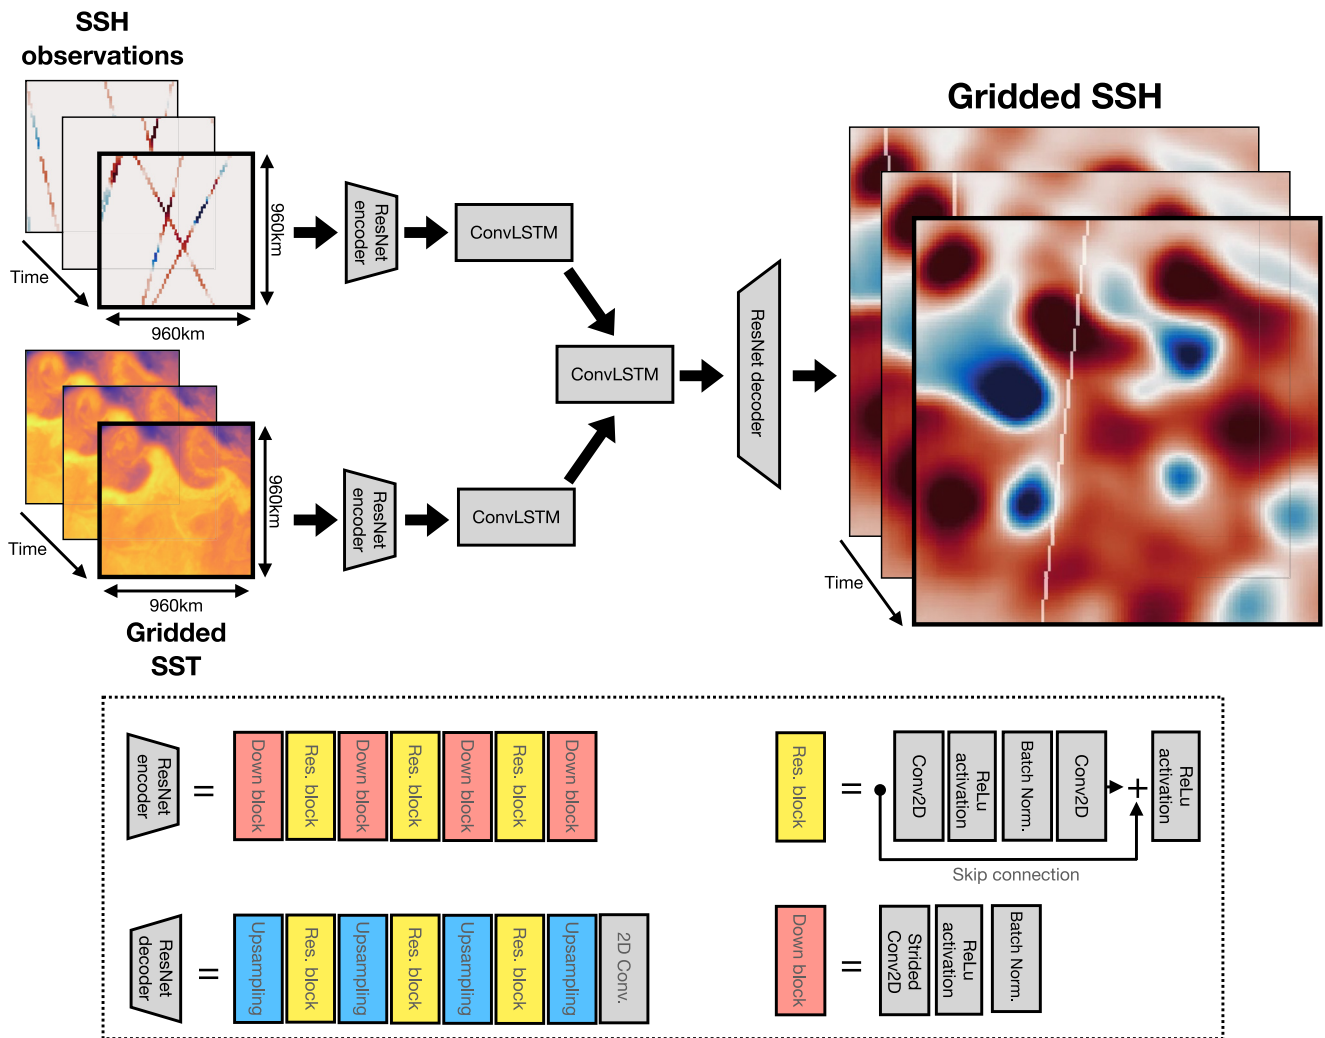


Figure 2. Schematic diagram of our deep learning method for SSH mapping. A time series of SSH and SST observations are each inputted into separate ResNet encoders (whose architecture is depicted in the lower panel) and ConvLSTM layers. Their representations are then combined through concatenation before being passed through another ConvLSTM layer and a ResNet decoder (whose architecture is in the lower panel). The loss function is then minimized along the location of the withheld altimeter observations (overlaid in white).

a 7.5 km grid with zero padding where there are no observations. Our mapping method estimates the 2D SSH field on this same grid. This choice of grid dimensions limits the computational resources required for the reconstruction while ensuring that the domain remains large enough that a significant number of interacting mesoscale eddies fit within the model's “field of view.” Defining a local orthonormal projection allows us to find the latitude and longitude for each grid point given the coordinates of the square domain's center. This projection ensures that the shape and size of eddies are not distorted at different latitudes, as would be the case if the square grids were defined in latitude-longitude space.

To avoid overfitting to the persistent SSH patterns in one such domain and with a view to creating a location-independent SSH mapping method, we generate the training data for our DL method by randomly selecting the coordinates on which to center the square domain for each training example. For most experiments in this study the length of the observation time series used to map SSH was chosen to be 30 days, and the effect of varying this length is explored in Section 5.1. The input to our mapping method for a single training example therefore consists of a time series (centered on a randomly chosen date) of bin-averaged SSH and SST observations within a square domain centered on a randomly chosen point in the ocean. The SSH and SST observations are each inputted to our neural network in the form of a 3D array (2D space plus time), with the SSH observation array containing zeroes everywhere except the places where along-track SSH was observed.

4. Deep Learning Method for SSH Mapping

4.1. Deep Learning Architecture

Here, we describe the DL neural network architecture we use to map SSH, a schematic diagram of our method is shown in Figure 2. Since the dynamical relationship between SSH and SST is non-trivial, we first use a “ResNet” CNN (He et al., 2016) to encode each variable separately in a learnable latent space representation. Conceptually, each ResNet encoder learns a mapping from the SSH/SST observations to some combined latent space in which the information from the SSH and SST observations can be combined on an equal footing. Note that trying out a more straightforward approach of combining SSH and SST as two different “channels” in the input to a single ResNet encoder yielded lower accuracy SSH mapping.

The ResNet architecture used for each encoder consists of alternating downsampling and residual learning blocks. A downsampling block consists of: a convolution with a stride of 2, followed by a rectified linear unit (ReLU) activation function, and finally, batch normalization. A residual learning block (He et al., 2016) consists of: a convolution with a stride of 1, ReLU activation, batch normalization, and another convolution of stride 1. The original input to the block is then combined with the output of the convolutions through addition (the so-called “skip connection” (He et al., 2016)), before a final ReLU activation function. Note that the same ResNet encoder is applied to each time step in the time series of input observations. Thus, the resulting latent space representation is one in which observations from each time step do not yet inform the representation at other time steps.

The resulting time series of latent space representations are then each passed through a bi-directional convolutional long short-term memory (ConvLSTM) layer (Shi et al., 2015). ConvLSTM is a type of recurrent neural network widely used for problems involving regularly-spaced spatiotemporal data which has been demonstrated to capture complex dynamical relationships, for example, to perform precipitation nowcasting (Shi et al., 2015). This layer must learn how partial observations from different times inform the state in the locations where observations are missing, that is, the dynamics governing SSH/SST evolution. Separate ConvLSTM layers are applied to the SSH and SST representations, resulting in latent space representation time series for SSH and SST where each time step is informed by observations at all other times.

The SSH and SST latent space representations are combined through concatenation. This time series is then passed through another bi-directional ConvLSTM layer which can learn relationships between the SSH and SST representations to give a combined latent space representation of the dynamical state of the ocean informed by both SSH and SST.

From this latent space representation, we finally use a ResNet to decode the representation to a gridded SSH map. The decoder is like the ResNet encoder networks, but with the downsampling blocks replaced by nearest-neighbor upsampling layers, alternating with residual learning blocks. The final layer has the dimensions of the desired final target SSH grid using a linear activation. The same ResNet decoder is applied to each step in the time series. The full neural network (ResNet encoders, then separate ConvLSTMs, then joint ConvLSTM, then ResNet decoder) provides an end-to-end trainable mapping from a time series of SSH and SST observations to a time series of 2D SSH maps.

Although our method produces a time series of SSH maps for the duration of the input observations, the reconstruction error increases away from the center of the time series. This is expected, since dynamically informative observations of past (future) states are missing for days at the beginning (end) of the time series. Thus, to generate our final SSH product, we retain only the reconstruction for the central time step, creating a time series of gridded SSH by successively shifting the input observation time series by 1 day.

4.2. Formulating a Loss Function in the Absence of Gridded SSH Data

The network is trained by minimizing the mean squared error between the reconstructed SSH and the ground truth. However, the full 2D ground-truth SSH is not available when training exclusively on real-world observations. To train using real-world observations, we implement the following strategy. We first withhold some of the SSH tracks from the input and then linearly interpolate the network's 2D reconstruction at the locations of these withheld observations to calculate the mean squared error between the reconstruction and the withheld observations. This way the error for any given training example is calculated only along a few withheld satellite tracks.

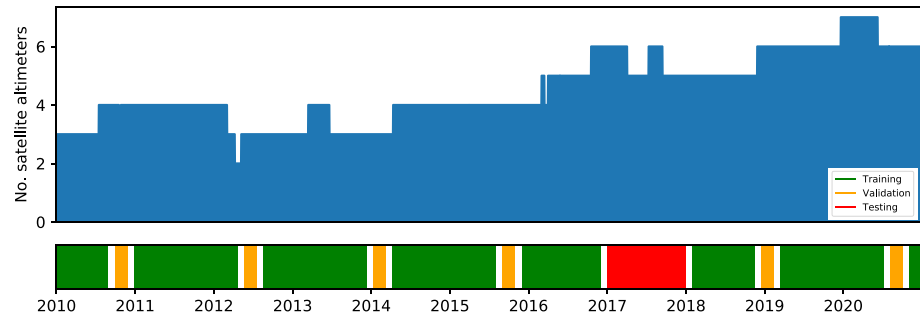


Figure 3. Dates are partitioned into non-overlapping and spaced out training (green), validation (orange), and testing (red) windows (the white bars represent the gaps between partitions when no examples are drawn). Examples are then drawn centered on days randomly drawn from the respective windows. The blue histogram shows the changing availability of satellite altimeters from which we can draw data.

Upon training, the network is forced to produce realistic 2D reconstructions throughout the domain since the withheld tracks appear at random places in the domain, so the network is unaware of where it will be evaluated.

Since there are typically several satellite altimeters operational at any time, for each training example we randomly select one of the available satellites, withhold its observations from the input, and use them as the ground truth in the loss function calculation. For each example, up to five satellite altimeters (depending on mission availability) are randomly selected to be used as the input SSH observations. The remaining satellites (or one satellite in the times when fewer than six missions are operational) are withheld for use as the ground truth when calculating the loss function.

To reduce over-smoothing and the appearance of high-frequency artifacts in the reconstruction, we include additional regularization terms in the loss function proportional to the mean squared error in the first and second along-track derivatives of SSH. Thus the cost, \mathcal{L} , we seek to minimize during training is given by:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(\frac{1}{\sigma_0^2} (\tilde{\eta}_i - \eta_i)^2 + \frac{\lambda_1}{\sigma_1^2} (\partial_x \tilde{\eta}_i - \partial_x \eta_i)^2 + \frac{\lambda_2}{\sigma_2^2} (\partial_x^2 \tilde{\eta}_i - \partial_x^2 \eta_i)^2 \right), \quad (2)$$

where N is the number of observations, η_i is the true observed SSH for the i th observation, $\tilde{\eta}_i$ is the correspondingly mapped SSH, x is a spatial coordinate following the satellite track, σ_0^2 , σ_1^2 , and σ_2^2 are the variances of η , $\partial_x \eta$, and $\partial_x^2 \eta$ respectively, and λ_1 and λ_2 are tunable parameters controlling the relative weighting for each regularization term in the loss function. In this study, we set $\lambda_1 = \lambda_2 = 0.05$. Along-track derivatives are estimated using first-order centered difference.

4.3. Choice of Training, Validation, and Testing Data Sets

The data are split into training, validation, and testing datasets to prevent overfitting. The validation data set is used to determine at what point the model is trained to its full potential. In contrast, the testing data set is withheld and used only at the end to test the model's performance on unseen data. Ideally, these datasets should be independent. To this end, we divide the available history of observations into time windows for training, validation, and testing as shown in Figure 3. The year 2017 was kept aside for testing to coincide with the testing period covered by an Ocean Data Challenge which allows us to compare our method to other proposed mapping methods (this data challenge is described in Section 6.1). The remaining times were then split into interleaved training and validation windows. We choose interleaved windows to ensure that the training and validation datasets both capture various seasons and sample interannual variability. These windows were chosen such that 80% of the dates fall within the training data set and 20% in the validation data set (specifically, we broke the available dates into 60 days chunks then chose the first 4 chunks to be training, the next 2 to be validation, then the next 4 to be training, and repeated this procedure until all dates had been assigned). Because there is significant autocorrelation in the SSH field over a time scale of $O(10\text{--}20)$ days, it is important to choose training and validation windows that are widely enough separated in time to prevent significant autocorrelation between the training and validation data. We ensure this by leaving a gap of 30 days between training and validation windows, since

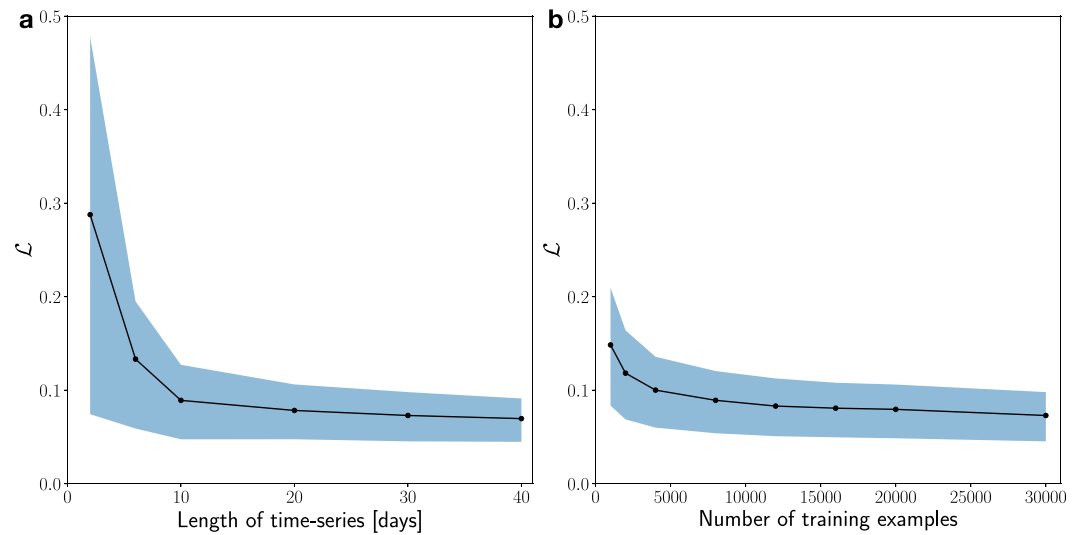


Figure 4. (a): the mean cost, \mathcal{L} , achieved on the validation data set (containing 2,000 samples) for our ConvLSTM model with SSH and SST input observations for increasing time series length. The blue shading shows the 16th and 84th percentiles, thus containing the same fraction of points as $\pm 1\sigma$ for a normal distribution. (b) The reduction in \mathcal{L} with increasing numbers of training examples for our ConvLSTM model with the length of the input observation time series set to 30 days (15 days either side of the reconstruction). Shading as for (a).

we observed the autocorrelation in the SSH field of an ocean global circulation model (GCM) (specifically, the CESM 1/10° Parallel Ocean Program model with JRA55 reanalysis forcing) to drop by ~50% in the Gulf Stream Extension over this time.

5. Training Performance

To provide a proof of concept for our method in this study we focus on mapping SSH in a region with energetic mesoscale dynamics, namely the Gulf Stream Extension. To this aim, we draw training and validation examples from the area shown in red in Figure 1. This area was chosen such that all examples feature the dynamics characteristic of the Gulf Stream Extension. In principle, we could draw training and validation examples from anywhere in the ocean, but since there are a wide range of dynamical conditions in the ocean it is challenging for a single neural network to learn to map SSH in all regions at once (this will be further discussed in Section 7.3). Training and validation examples are drawn from this region with the square domain centered at a randomly selected point and time (respecting the training-validation date partition outlined in Section 4.3). This region features an abundance of mesoscale eddies and complex jet-eddy interactions between the Gulf Stream and neighboring eddies. It is thus a challenging region for SSH mapping, where the widely-used DUACS product displays relatively high mapping errors, with the error becoming comparable to the standard deviation of the along-track SSH signals.

5.1. Determining the Optimal Length of the Input Time Series and Number of Training Examples

The SSH reconstruction on a given day is generated using a time series of observations centered on that day. Thus, a key parameter of our approach is the length of this time series. Including observations from further into the past and future increases the aggregated spatial coverage of the SSH observations and provides information about the time variation of the surface flow. To accurately reconstruct the SSH in unobserved regions, the network needs to learn an approximation of the surface dynamics, so that observations from other days can be used to inform the reconstruction.

Physically, we expect there to be some predictability time horizon beyond which observations become decreasingly helpful in informing the reconstruction. Therefore, we expect to see diminishing improvement to the reconstruction as the length of the time series is increased beyond some time period. We tested this hypothesis by varying the length of the input time series and training our ConvLSTM model, with both SSH and SST observations for input, and finding the mean value for the cost, \mathcal{L} , achieved on the validation data set (Figure 4a).

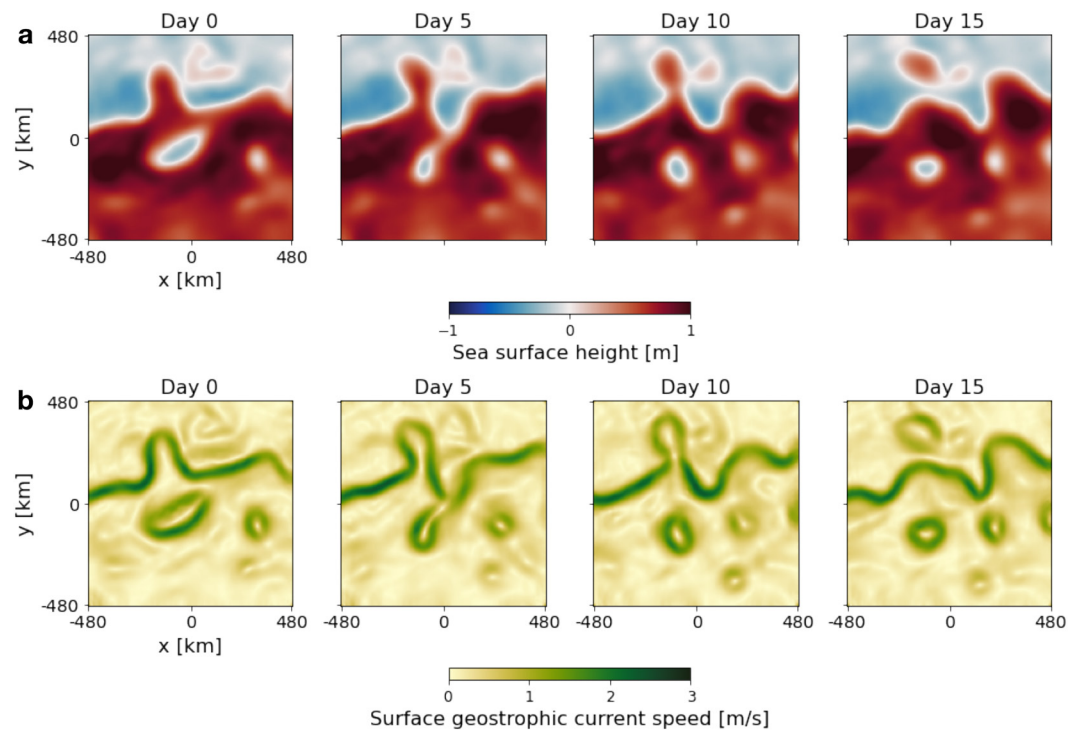


Figure 5. An example of the output from our ConvLSTM SSH + SST mapping method. Time series of (a) SSH (absolute dynamic topography) mapped using ConvLSTM SSH + SST and (b) surface geostrophic current speed calculated from the SSH field.

\mathcal{L} reduces with increasing time series length until about 20 days (i.e., ten days on either side of the reconstruction), beyond which further lengthening yields minimal improvement. As demonstrated in Figure 4b, we checked that \mathcal{L} had stopped improving significantly as a function of the number of training examples for each time series length. This time-scale could be a feature of our neural network architecture rather than a physical predictability horizon.

The results in the remainder of this study use a time series length of 30 days (i.e., fifteen days on either side of the date for reconstruction) noting that the significant computational cost of extending the time series beyond this length is unlikely to yield significant reductions in \mathcal{L} .

5.2. Quality of Mapped SSH Field

Figure 5 provides an example time series of SSH mapped using our method from the independent testing data set (i.e., during 2017). The SSH field features separated, distinct mesoscale eddies and a strong Gulf Stream, evident in the surface geostrophic currents calculated from the SSH map. The time-evolution of the reconstructed field also looks qualitatively realistic, with a Gulf Stream meander pinching off to form a new mesoscale eddy to the north of the jet. The capability of our method to produce realistic-looking SSH fields from observations, unseen during training, gives us confidence that the model has learned a robust mapping from SSH and SST observations to the 2D SSH field, rather than overfitting to the training data.

6. Inter-Comparison of Existing SSH Mapping Methods in the Gulf Stream Extension Region

6.1. Mapping Methods and Evaluation Metrics

We use the data distributed through the AVISO Ocean Data Challenge (AVISO, 2021) to provide a direct quantitative comparison of our method's accuracy to that of five established and experimental SSH mapping methods. The data consist of SSH maps created for the year 2017 (the year set aside for testing) in the region covering

(55°–65°W, 33°–43°N) using five different mapping methods, described below. Each map was created using observations from all available altimeters apart from CryoSat-2, which is withheld for use as a ground-truth for calculating the maps' errors.

We compare our approach to five other mapping methods:

- *DUACS*: the community-standard SSH product created using OI (Taburet et al., 2019). The full mapping method is not publicly available, but a map was generated excluding CryoSat-2 for the data challenge.
- *DYMOST*: the “Dynamic Interpolation” method proposed by Ubelmann et al. (2015) and evaluated by Ubelmann et al. (2016) and Ballarotta et al. (2020). This method is a variant of the OI approach. The Gaussian a priori SSH covariance model used in DUACS is replaced by a dynamically informed model based on the forward and backward in time integration of an idealized potential vorticity conservation equation. In this method the SSH evolution is assumed to be governed only by the first baroclinic mode.
- *MIOST*: the “Multiscale and Multivariate Interpolation” method proposed by Ubelmann et al. (2021). This method is another variant of the OI approach where the single-scale Gaussian covariance model for SSH used in DUACS is replaced with a wavelet basis where the amplitude of each scale component is set to match the power spectrum observed in along-track altimetry, thus assuming isotropy.
- *BFN-QG*: the “Back-and-Forth Nudging QG” method proposed by Le Guillou et al. (2021). This method uses the back-and-forth nudging method for data assimilation to interpolate the altimetry observations while respecting the dynamics of a one-layer QG model of ocean turbulence.
- *4DVarNet*: the DL method proposed by Fablet, Amar, et al. (2021) and Fablet, Chapron, et al. (2021). This approach poses the SSH interpolation problem as a four-dimensional variational data assimilation (4DVar) problem. It replaces the dynamical model and solver with trainable neural networks. We note that this method has been used in several OSSE studies, so many different versions exist in the literature. Here, we use the map provided through the data challenge which takes only SSH observations as input. Fablet and Chapron (2022) and Fablet et al. (2022) recently showed that including SST in the 4DVarNet approach could significantly improve the SSH reconstruction. However, this has only been demonstrated in the OSSE setting where the full underlying ground-truth SSH is known during training. It has not yet been demonstrated to improve SSH mapping using real-world SST observations which suffer gaps due to cloud cover if the highest resolution observations are used. Hence, we cannot directly compare the two methods with the inclusion of SST but note that the improvements seen for both methods in their respective settings are consistent with the expectation from SQG dynamics that the SST field contains relevant information about the dynamics of the SSH field.

We do not here compare our method to a full-scale ocean GCM data assimilation reanalysis (e.g., Lellouche et al., 2021), since none have been produced with CryoSat-2 observations withheld from the assimilation process, preventing a like-for-like comparison to the other methods in the above data challenge framework. The comparison of our method's SSH mapping accuracy to a GCM data assimilation scheme will be the subject of a future study.

For our analysis, we reduce the size of the data challenge's study region to (55°–65°W, 34°–42°N) to ensure the full region fits inside the “field of view” of our neural network's 960 × 960 km output. All methods are assessed in this same region over the entire year of 2017. Figure 5 shows an example of the result of applying our method to map SSH in this study region.

We use the withheld CryoSat-2 observations to calculate various metrics to quantitatively compare the performance of the SSH mapping methods. These metrics are summarized below:

- *Mean root mean square error (RMSE) (cm)*: The SSH maps are interpolated to the locations of the withheld observations and the RMSE of the map is calculated for each day where observations are available, then the mean of these daily RMSEs is taken.
- *Standard deviation of RMSE (cm)*: The standard deviation of the daily RMSE values.
- *Effective spatial resolution (km)*: Calculated using the along-track SSH spectra for each satellite pass. This is the wavelength at which the power spectral density of the misfit between the map and the observations

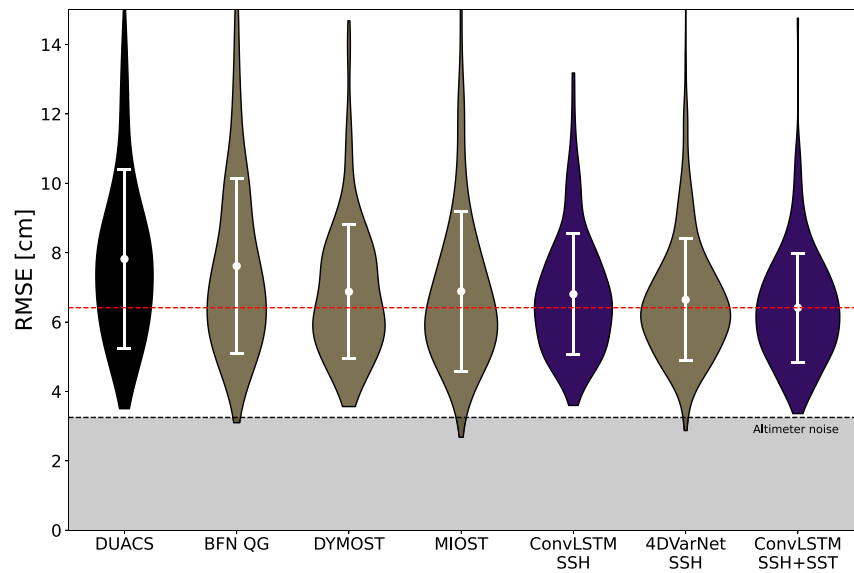


Figure 6. Violin plot showing the distribution of daily reconstruction root mean square error for each method when compared to independent satellite altimeter observations for the test period. The black violin shows the community-standard optimal interpolation product, the gold violins are experimental methods proposed by others in the literature, and the purple violins are using our new method. The width of each violin is proportional to the probability density function, the white circle indicates the mean, and the whiskers show one standard deviation. The shaded region indicates the estimated instrument noise in the observations (as reported by Copernicus Marine Environment Monitoring Service).

becomes comparable to that of the observations, quantifying the spatial scales accurately resolved by the map. See Ballarotta et al. (2019) for a complete discussion of this metric.

- *RMSE score (no units):*

$$RMSE\ score = 1 - \frac{RMSE_{map}}{RMS_{obs}}$$

A perfect map would have a score of 1, a map that predicted zeroes everywhere would have a score of 0.

- *Standard deviation of RMSE score (no units):* Standard deviation of the daily RMSE scores.

6.2. SSH Validation Against Independent Satellite Observations

6.2.1. Root Mean Square Error (RMSE)

All the experimental mapping methods were found to have lower mean RMSE values than DUACS and the spread in their daily RMSE is also lower, as can be seen in Figure 6 and Table A1. BFN-QG provides only a marginal improvement in RMSE, whereas DYMOST and MIOST significantly improve (a 12% reduction in mean RMSE). All the DL methods tested here (4DVarNet SSH and all configurations of our ConvLSTM method) yielded lower RMSE than all the other methods. 4DVarNet SSH showed a 14% reduction in RMSE while our ConvLSTM method respectively gave a 13% and 17% reduction in RMSE when run with just SSH (ConvLSTM SSH) and with combined SSH and SST observations (ConvLSTM SSH + SST).

To explore the eddy-jet configurations that present the most significant challenge to traditional altimetry mapping and further illustrate our method's improvement, in Figure 7 we present three case studies from the testing period where we contrast the DUACS reconstruction with our ConvLSTM SSH + SST method. These cases correspond to the three CryoSat-2 satellite tracks of at least 800 km in length for which the DUACS reconstruction shows its highest RMSE. In all three cases, the primary source of error is the misrepresentation of a sharp kink in the Gulf Stream in the top left corner of the domain (recall that geostrophic surface currents flow along contours in SSH). In all three cases ConvLSTM SSH + SST leads to a significantly more accurate along-track SSH profile. Physically, the eddy features in SSH are expected to evolve relatively quickly in shape and amplitude in regions where

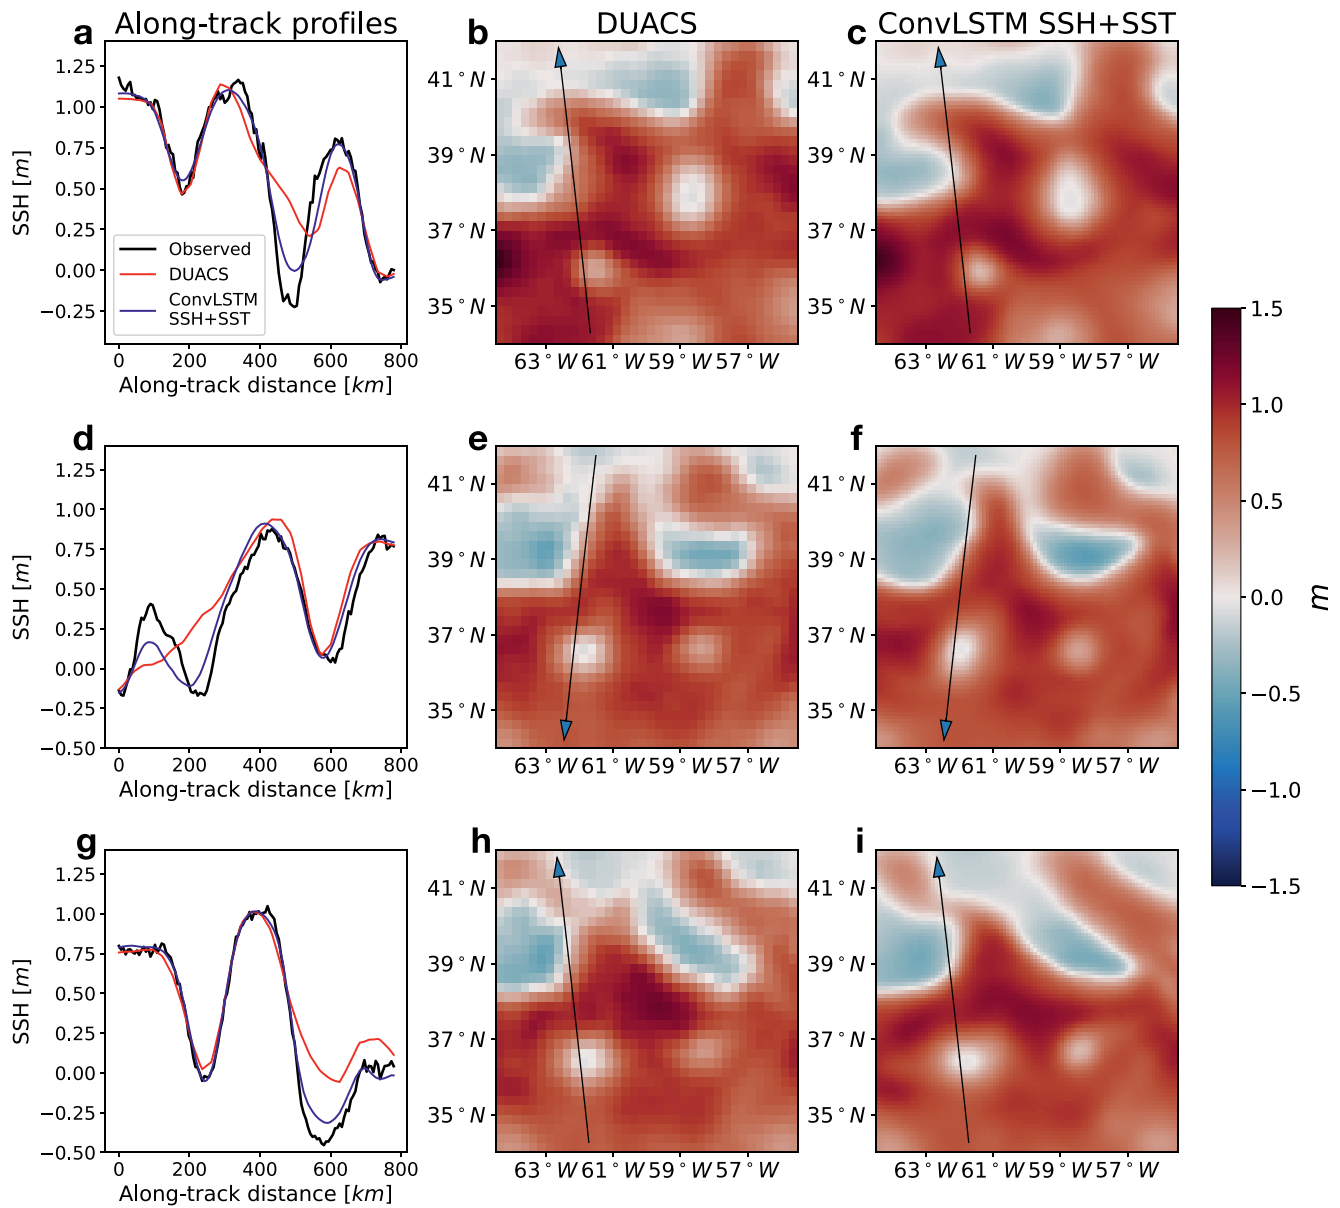


Figure 7. Example along-track SSH profiles. Shown are the three 800 km track segments (a, d, g) where Data Unification and Altimeter Combination System (DUACS) suffers the highest root mean square error during 2017 along with the track's location superimposed over the DUACS reconstruction (b, e, h) and the ConvLSTM SSH-SST reconstruction (c, f, i). In each case, DUACS over-smoothes one or more large peaks/troughs in SSH. In all of these cases ConvLSTM provides a significant improvement.

the jet (i.e., the Gulf Stream) is bent at a sharp angle, thus presenting a challenge to the simple SSH covariance model employed by DUACS. Our ConvLSTM SSH + SST method overcomes this shortcoming.

6.2.2. Effective Along-Track Spatial Resolution

All the experimental mapping methods were able to reconstruct smaller-scale SSH signals than DUACS, as can be seen in Figure 8. The spectra in Figure 8a show that the DUACS reconstruction underestimates the strength of the SSH signal at smaller scales compared to the observations. All other methods show spectra closer to the observations at these scales to varying degrees. The spectral coherence plot in Figure 8b demonstrates that the DUACS reconstruction is less accurate than all other methods at scales of 70–200 km. None of the methods accurately resolve SSH below these scales due to the limited along-track resolution of the observations and the size of the gaps between observations.

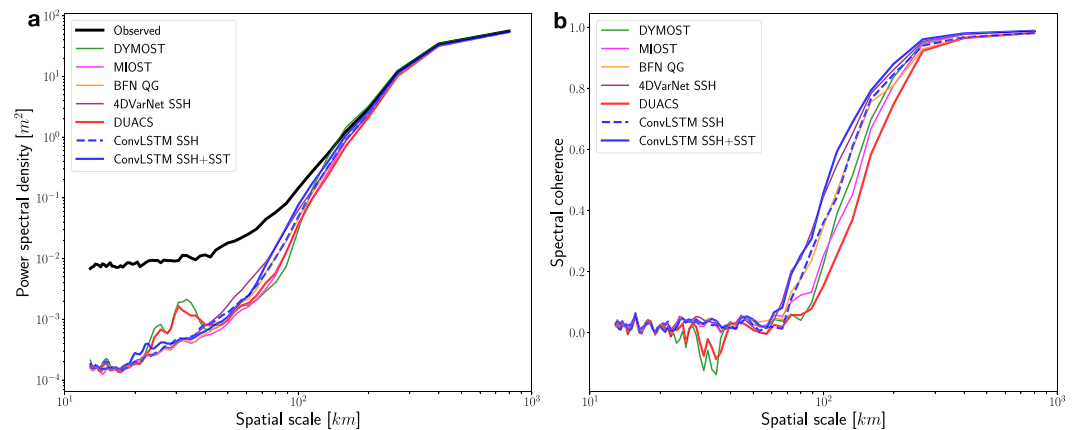


Figure 8. (a) Wavenumber spectrum of mapped SSH along the tracks of the withheld CryoSat-2 observations for each method, compared to that of the observations. The observations show a high power spectral density for small scales due to instrument noise. (b) Spectral coherence for each mapping method (approaches unity for scales at which the SSH signal is well resolved by the map).

The effective spatial resolution (the scale at which each curve in Figure 8b crosses 0.5, also listed in Table A1) for each method reveals significant differences. MIOST and DYMOST only offer marginal improvements on this metric (a 7% and 12% reduction in the smallest wavelength accurately resolved), whereas BFN-QG provides a more substantial improvement (20%). Again, all DL methods yield the largest improvements on this metric, with ConvLSTM SSH, 4DVarNet SSH, and ConvLSTM SSH + SST providing 23%, 28%, and 30% improvements respectively. The SSH map produced using our ConvLSTM SSH + SST method accurately resolved SSH signals with wavelengths as small as 104 km.

The significant improvement in all metrics by all the DL methods provides a compelling case for the use of DL in developing an improved 2D SSH product. The addition of SST observations also provides a clear improvement in our method's reconstruction accuracy and resolution.

6.2.3. Frequency Spectra of SSH Fields

Another important dynamical property of a mapped SSH field is its frequency spectrum. Unlike for the wavenumber spectrum, discussed in Section 6.2.2, estimating a ground truth from the independent altimeter observations is challenging since at any one location the temporal sampling is relatively sparse (CryoSat-2 observes the same 1/8° by 1/8° spatial bin once every 5–6 days on average in this region). Thus, we cannot calculate the frequency spectral coherence in the manner of Section 6.2.2 using just the withheld altimeter observations.

Nevertheless, we can still characterize the frequency spectrum of each SSH map and draw comparisons among the maps, as shown in Figure 9a. Compared to all other methods, DUACS has a substantially lower power spectral density at time scales shorter than 30 days, implying that DUACS may be underestimating the short-time SSH variability, as has been noted in the literature (Ballarotta et al., 2019). The other maps all show similar frequency spectra down to time scales of around 5 days, below which 4DVarNet and BFN QG show significantly higher power spectral density. Without knowing the ground truth frequency spectrum from observations, it is hard to discern which map's spectrum most accurately represents the variability at these time scales. The higher variability at short temporal scales for 4DVarNet and BFN QG could be due to the appearance of non-physical artifacts along the location of the input satellite tracks (discussed more in Section 6.3).

To provide some estimate of the SSH variability from the observations, we also calculated the second-order structure function in time, as shown in Figure 9b. To do this, we used all available satellite altimetry observations from the study region and time, selected the 100 most frequently observed 1/8° by 1/8° bins, used the observations to construct a SSH time series for each bin, and used these to estimate the second order structure function. We thus assume that the frequency spectrum of the SSH field is approximately isotropic within this region. We used the bias-corrected and accelerated bootstrap method to estimate 95% confidence intervals.

Calculating the second-order structure function for each of the SSH maps at uniformly-spaced points throughout the region allows us to compare the maps' variability to the observational estimate. We see that DUACS is indeed

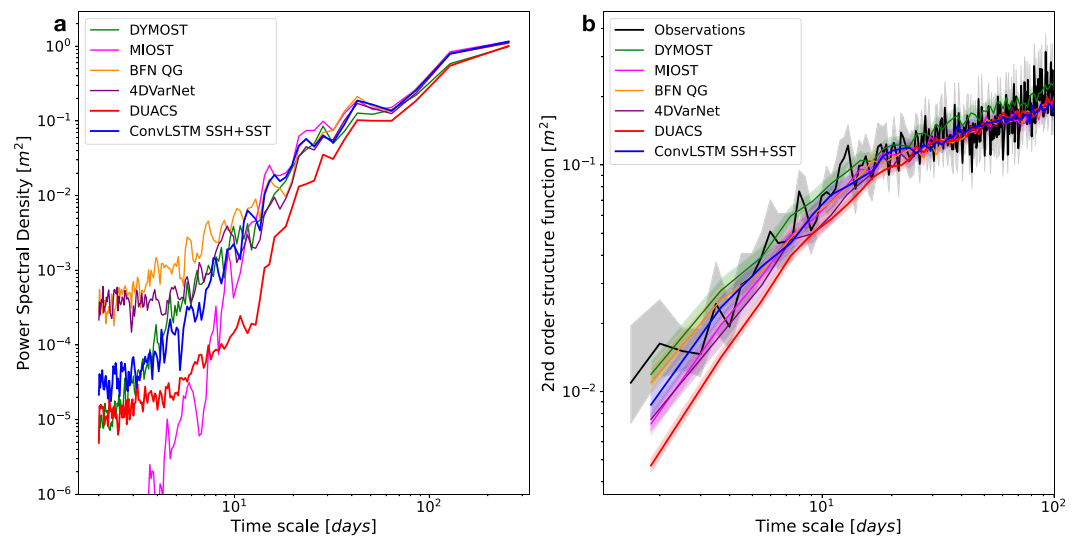


Figure 9. (a) Frequency spectrum for each SSH map averaged over points throughout our study region. (b) Second order structure function calculated from all available satellite observations (including those used to make the maps) at satellite track crossover points, compared against that of each map average over uniformly-spaced points in the domain. Error bars show the 95% confidence interval estimated using the bias-corrected and accelerated bootstrap method.

underestimating the variability of the SSH field at time scales below 20–30 days as the higher variability shown by the other maps is replicated in the observations. The differences between the other maps at short (<10 days) time scales are small and the confidence interval on the observational estimate is large since satellite altimetry provides poor sampling of the SSH field at these short time scales.

6.2.4. Seasonality of Improvement to ConvLSTM Mapping With the Inclusion of SST

The utility of SST for recovering sea surface currents in theory should depend on environmental conditions which in turn vary geographically and seasonally. Enhanced surface temperature stratification in the summer months may obscure the signature of ocean eddies, while the phase shift between mesoscale SST anomalies and surface currents is known to be smaller when the mixed layer is deep, typically in the winter months (Isern-Fontanet et al., 2008). Thus, it is anticipated that SST will be most useful in the winter months for estimating surface currents (and SSH by extension). González-Haro and Isern-Fontanet (2014) explored the geographical and seasonal variations in the accuracy of sea surface currents recovered from SST through SQG theory and for the Gulf Stream Extension region considered in this study found that the SST-derived surface currents were most accurate in the winter.

Since our DL approach does not explicitly rely on the SQG framework, it is interesting to consider the seasonality of the improvement to the ConvLSTM mapping accuracy and resolution when SST is included. We here study this seasonality for the years 2017, 2021, and 2022. These years were chosen since there were a large number of altimeters available continuously in 2021 and 2022 (6 and 5 respectively, we excluded Jason-3 in 2022 since there was a gap in coverage) and none of the three years were used during training or cross-validation. To ensure a large enough number of independent altimeter observations for robust statistics, we generated SSH maps for each possible combination of three input satellites and evaluated each map against the remaining available altimeters; this gives 4 maps for 2017, 20 for 2021, and 10 for 2022. Rolling 2-month averages of the ConvLSTM mapping RMSE (normalized by the along-track SSH RMS variations since this varies seasonally) and effective spatial resolution were calculated for all three years with and without SST.

While there is variation between the years, the improvement in both accuracy and resolution for ConvLSTM SSH + SST compared to ConvLSTM SSH was found to be slightly higher in the summer months than during the winter (Figure 10). The ConvLSTM SSH + SST normalized RMSE varied from around 8% lower than ConvLSTM SSH in the winter months to around 15% lower in the summer. While the seasonal trend was less clear in the effective resolution, the maximum percentage reduction in resolved wavelength for ConvLSTM SSH + SST was found outside the winter months for all 3 years.

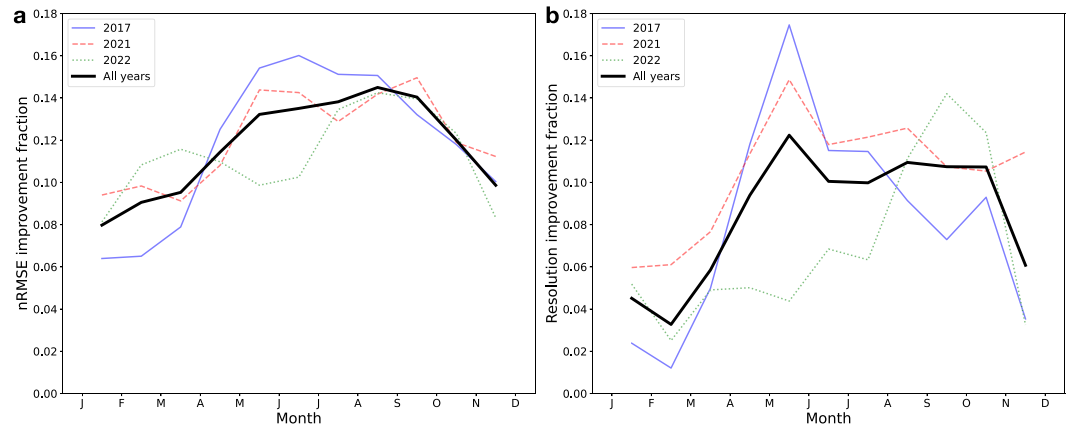


Figure 10. (a) Fraction by which the normalized root mean square error (nRMSE) was reduced with ConvLSTM SSH + SST compared to ConvLSTM SSH. Maps were created from all possible three-satellite combinations and the mean is shown for each year along with the combined mean over the 3 years (weighted to account for the different number of satellite configurations in each year). (b) Fraction by which the effective spatial resolution was reduced for ConvLSTM SSH + SST compared to ConvLSTM SSH. Again, the mean across all three-satellite configurations is shown for each year along with the combined 3-year mean.

These findings run contrary to the expectations from ocean dynamics described above. One possible explanation is the impact of the seasonally varying cloud cover on the Level 4 GHRSSST MUR SST product used in this study. Since this product combines both microwave and infrared SST observations, its spatial resolution varies dramatically in space and time. The high-resolution infrared SST sensors can only be used on cloud-free snapshots whereas the lower-resolution microwave SST sensors are insensitive to cloud cover, so the MUR SST product has a high spatial resolution on cloud-free days but becomes coarser on cloudy days when it relies on the microwave observations. Cloud cover was higher in our study region during the winter months than during the summer, so the improved performance of ConvLSTM SSH + SST during summer is likely due to the higher resolution of the SST observations rather than some ocean dynamics signature. The sensitivity of our method to the choice of SST product will be explored in a future study. We note that the Level 4 GHRSSST MUR SST product provides gridded uncertainty estimates for the interpolated SST field. While in this study we did not make use of these estimates, a possible future line of inquiry would be to include them in the input to allow the neural network to put more weight on the more reliable, cloud-free SST regions.

6.3. Surface Geostrophic Currents

The primary use of SSH maps is to infer surface ocean currents through the assumption of geostrophy. The eastward and northward geostrophic surface currents (u_g and v_g respectively) are related to the SSH, η , through

$$(u_g, v_g) = \frac{g}{f} \left(-\frac{\partial \eta}{\partial y}, \frac{\partial \eta}{\partial x} \right), \quad (3)$$

where x and y are eastward and northward spatial coordinates respectively, f is the Coriolis frequency, and g is the acceleration due to gravity. The currents are thus proportional to the first-order spatial derivatives of η .

From these currents, other physical quantities such as relative vorticity, ω , and strain rate, γ , can be calculated by taking spatial derivatives of the velocity field (corresponding to second-order derivatives of the SSH field):

$$\omega = \frac{\partial v_g}{\partial x} - \frac{\partial u_g}{\partial y} \quad (4)$$

$$\gamma = \sqrt{\left(\frac{\partial u_g}{\partial x} - \frac{\partial v_g}{\partial y} \right)^2 + \left(\frac{\partial v_g}{\partial x} + \frac{\partial u_g}{\partial y} \right)^2}. \quad (5)$$

These higher order derivatives are an important diagnostic since they give information about the dynamics (i.e., the acceleration) of the surface motions (Hua & Klein, 1998). Relative vorticity (and potential vorticity)

are crucial quantities when studying ocean turbulence because in 2D turbulence vorticity is conserved along streamlines and potential vorticity is conserved in QG and SQG motion. As such the interaction of eddies with each other and with the mean flow, as well as their equilibration, filament formation, and frontogenesis are all processes that are viewed through the lens of vorticity and strain evolution.

One such quantity that is important to consider is the Okubo-Weiss quantity, W , defined as the difference between the square of the strain rate and the square of relative vorticity (Hua & Klein, 1998; Hua et al., 1998; McWilliams, 1984; Okubo, 1970; Weiss, 1991)

$$W = \gamma^2 - \omega^2. \quad (6)$$

When W is negative, relative vorticity and therefore rotation dominates. This typically occurs in the regions inside mesoscale eddies where SSH extrema are found. However, when W is positive or close to zero, the strain rate dominates or is close to vorticity, indicating an intensification of SST fronts and associated vertical velocity. This typically occurs in regions between mesoscale eddies and at the eddy edges (Hua et al., 1998; Klein et al., 2019; Lapeyre et al., 1999).

We calculated the surface geostrophic currents, ω , and W from the DUACS and ConvLSTM SSH + SST maps. To closely match the procedure for calculating current speeds from the DUACS SSH map used to generate the surface current speeds distributed in the CMEMS Level 4 SSH product (CMEMS, 2023c), we use a 9-point stencil (as advocated in Arbic et al. (2012)) for estimating derivatives of the DUACS SSH field. Because the DUACS product is distributed on a $1/4^\circ$ (~ 30 km) grid whereas our SSH reconstruction is on a 7.5 km grid, we take all DUACS derivatives on its native grid before linearly interpolating to our higher-resolution grid.

The ConvLSTM SSH + SST reconstruction results in narrower and stronger currents than those calculated from the DUACS reconstruction, as shown in Figures 11a and 11d. The maximum current speed for the date shown in Figure 11 is 40% stronger in the ConvLSTM SSH + SST reconstruction than in the DUACS reconstruction. This result is expected since the covariance model used in DUACS leads to an overly smooth SSH field, smaller-scale features being blurred out. Hence, the magnitude of the steepest gradients of the SSH field will, in turn, be weaker.

The NOAA Atlantic Oceanographic & Meteorological Laboratory (AOML) global surface drifter program provides in situ observations of the total near-surface (15 m depth) ocean currents through satellite-tracked drifters. To quantitatively test the accuracy of the surface currents inferred from the SSH maps, we here compare the mapped currents to all available AOML drifters within our study region and time. The AOML drifter data are available through the CMEMS in situ surface and sub-surface water velocity observations product (CMEMS, 2023a). The results shown here are after we applied the wind slippage correction distributed within the product, however, removing this correction does not change the hierarchy of the methods' errors. Both ConvLSTM SSH and ConvLSTM SSH + SST show a clear improvement in current reconstruction accuracy relative to DUACS and the other methods, with a 24% and 27% decrease in surface current RMSE respectively (surface current RMSE for each method can be found in Table A1). Surface drifters don't directly measure geostrophic currents, since Ekman and ageostrophic components will be aliased onto the drifter-observed currents. However, the fact that we can more accurately estimate the surface currents through geostrophy from our improved SSH map than from the DUACS map is compelling since surface drifters are an observational data set that is entirely independent of the satellite altimetry methodology.

The apparent difference between the DUACS and ConvLSTM SSH + SST reconstructions becomes starker for the relative vorticity and Okubo-Weiss quantity, W , (shown in Figures 11b, 11c, 11e, and 11f) since deficiencies in the reconstructed SSH field are amplified as derivatives of higher order are taken. The path of the Gulf Stream is hard to identify from the DUACS relative vorticity field visually and is practically impossible to pick out in the Okubo-Weiss field. By contrast, the ConvLSTM SSH + SST map results in a relative vorticity and Okubo-Weiss fields that show a clear, physically realistic Gulf Stream trajectory. Two evident mesoscale eddies can also be seen in the Okubo-Weiss field as distinct blue regions where vorticity dominates the flow surrounded by red regions where the eddy is inducing high strain rates. One deficiency of our reconstruction is the presence of faint high frequency grid-like artifacts. Such artifacts are common when using CNNs, and their intensity decreases with increasing number of training examples. We discuss these more in Section 7.3 but note that they only become visually apparent when taking higher-order derivatives of the SSH field.

Although several of the methods in this study exhibit similar summary statistics, there are large qualitative differences in the resulting current reconstructions. In Figure 12, we show the relative vorticity field calculated from each SSH map on the same date. The DYMOST and MIOST vorticity fields are qualitatively like that of

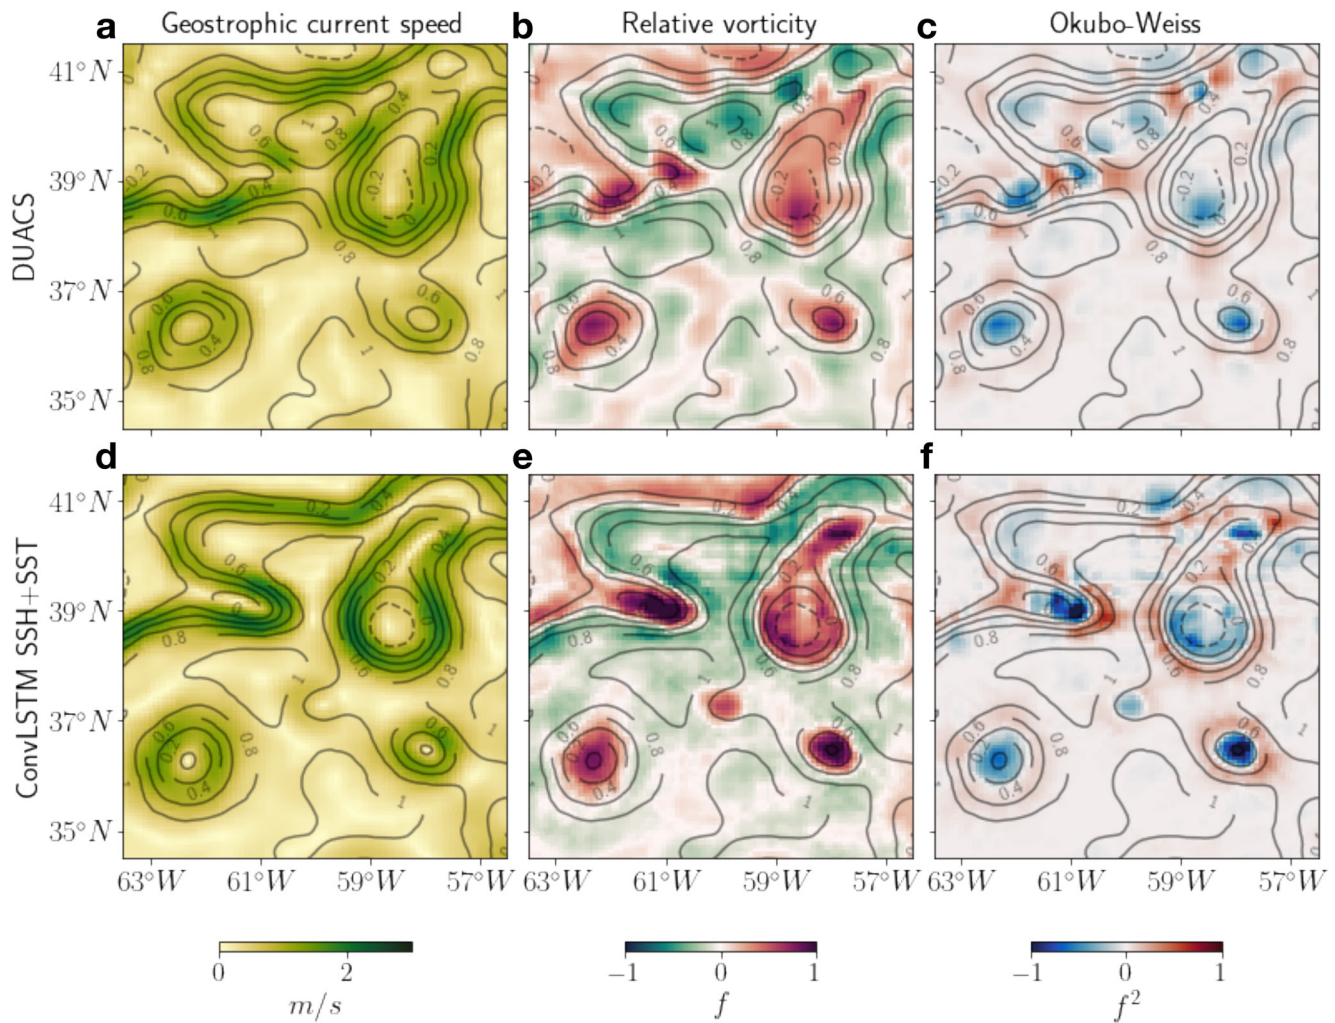


Figure 11. (a–c) Surface geostrophic current speed, relative vorticity, and the Okubo-Weiss quantity calculated from the Data Unification and Altimeter Combination System SSH map. (d–f) The same quantities calculated from the SSH map produced by our ConvLSTM SSH + SST method. Contours of the corresponding SSH maps (absolute dynamic topography) are overlaid in black. Relative vorticity and Okubo-Weiss are normalized by the Coriolis frequency, f , and its square respectively. A video of this figure is available for the full year in the (Movie S1).

DUACS so the Gulf Stream's large-scale structure is challenging to discern. By design, BFN-QG results in a physically realistic-looking vorticity field since the SSH map is the result of solving an idealized dynamical model of ocean turbulence, albeit an overly-idealized one. 4DVarNet produces a vorticity map with similar large-scale structures to other methods but with the addition of some smaller-scale features and clear artifacts. The prominent mesoscale eddies (the two distinct red regions in the lower half of the domain) exhibit high-frequency radial features in the 4DVarNet reconstruction that are likely to be CNN-induced artifacts. In this reconstruction, there are straight-line features along which the map exhibits smaller-scale features. These lines correspond to the ground tracks of the satellites providing the input SSH observations (see Movie S2 for the vorticity evolution).

These qualitative differences have impact on two key physical quantities of interest: EKE and enstrophy. Figure 13a shows each map's time- and domain-averaged EKE. DUACS has a significantly lower EKE than the other maps while DYMOST and ConvLSTM SSH + SST have the highest EKE. Larger-scale flows dominate EKE, so the fact that DYMOST and ConvLSTM SSH + SST have a substantially higher EKE than other methods implies these reconstructions result in faster large-scale currents (e.g., the Gulf Stream) than the others. Figure 13b shows the time- and domain-averaged enstrophy for each map. Enstrophy is a quantity that is typically dominated by smaller-scale flows. DUACS has a significantly lower enstrophy than the other maps, consistent

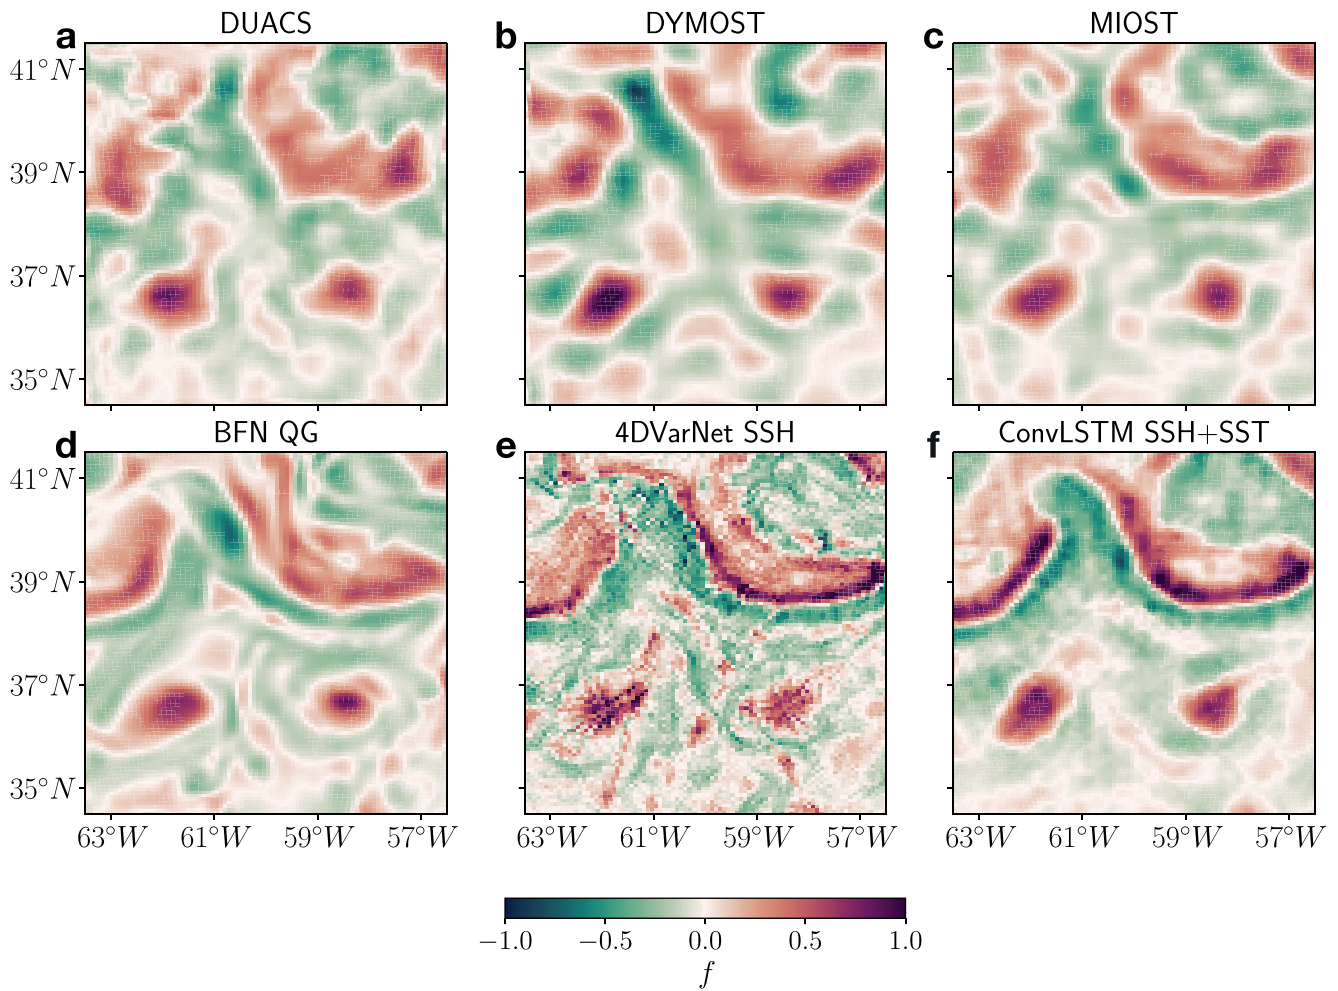


Figure 12. Relative vorticity calculated from the surface geostrophic currents from each method's SSH map, normalized by the Coriolis frequency, f . A video of this figure is available for the full year in the (Movie S2).

with the expectation that OI smooths out small-scale features. 4DVarNet has a very high enstrophy compared to

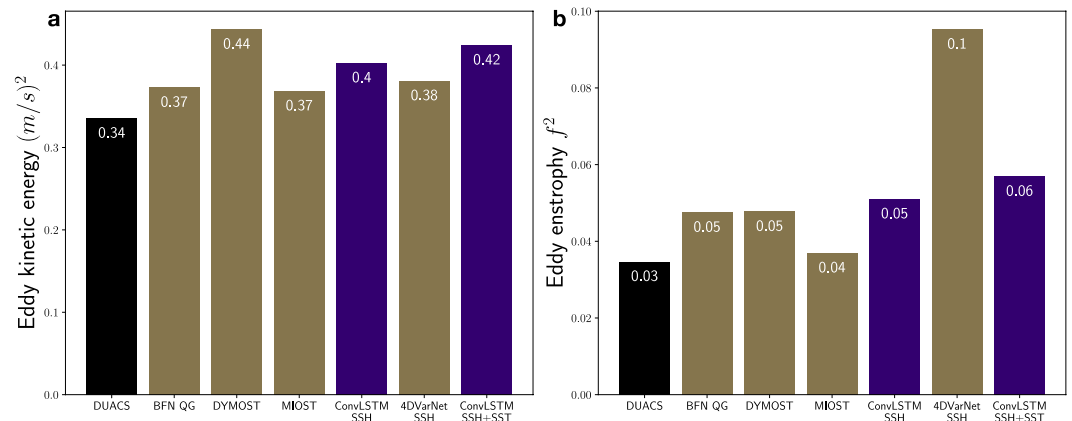


Figure 13. Domain- and time-averaged eddy kinetic energy (a) and enstrophy (b) for the geostrophic currents calculated from each SSH map. Enstrophy is normalized by the square of the Coriolis frequency, f . The black bars shows the community-standard optimal interpolation product, the gold bars are experimental methods proposed by others in the literature, and the purple bars are using our new method.

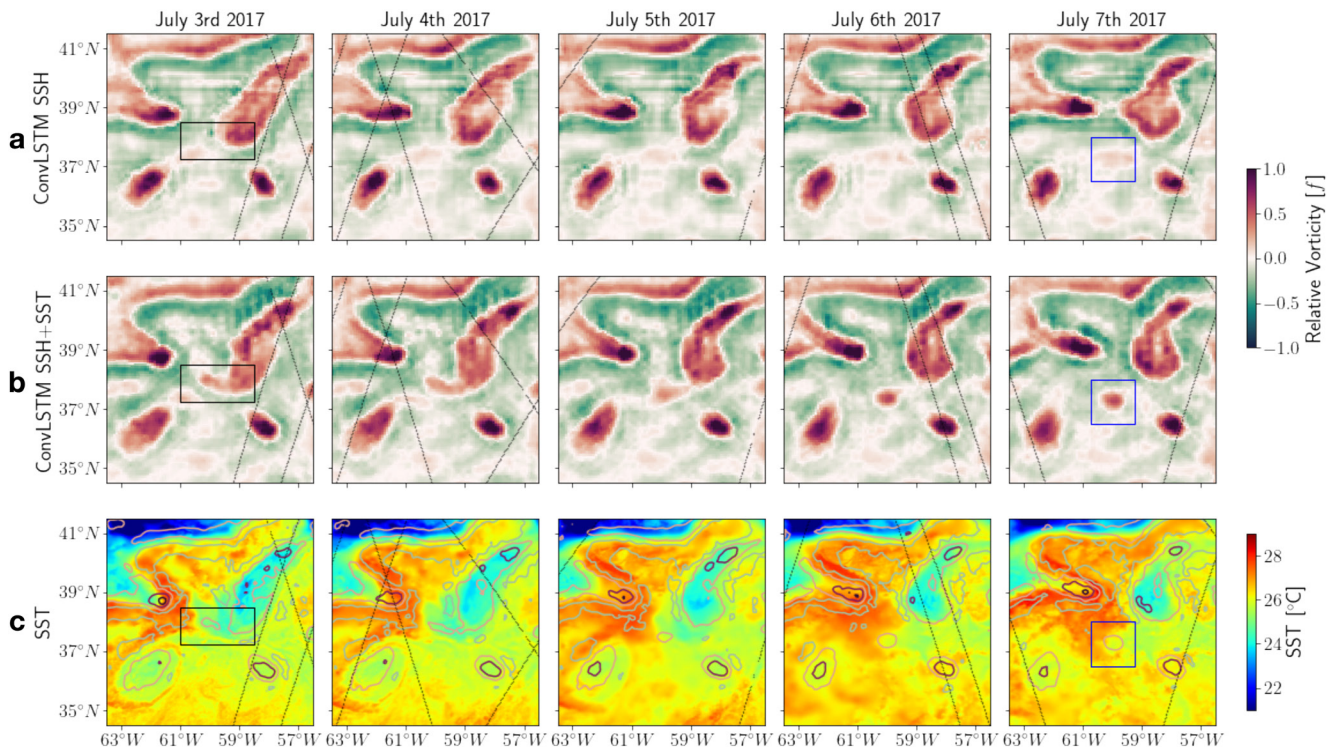


Figure 14. Relative vorticity from (a) the ConvLSTM SSH and (b) the ConvLSTM SSH + SST reconstructions for 5 consecutive days. GHRSSST Multiscale Ultra-high-Resolution SST for the same days is shown in row (c) and the locations of the satellite altimetry observations on each day are overlaid in black. ConvLSTM SSH + SST relative vorticity contours in increments of $0.5f$ from $-1.25f$ to $1.25f$ are overlaid over the SST observations. The black box on July 3rd highlights the filament from which this eddy formed.

the other methods due to the unphysical high-frequency artifacts in the relative vorticity field.

6.4. New Mesoscale Eddy Found by Synthesizing SST and SSH Observations

To further emphasize the power of using SST observations to help map SSH, we here present one case study from our testing data set where ConvLSTM SSH + SST was able to reconstruct a mesoscale eddy that was missed by all the other methods, which only used SSH observations as input. Figure 14 shows the relative vorticity fields for five consecutive days in July 2017 reconstructed using both ConvLSTM SSH and ConvLSTM SSH + SST. On July 7th in the ConvLSTM SSH + SST reconstruction there is a cyclonic mesoscale eddy in the center of the domain (boxed in blue in the figure) which is not present in the ConvLSTM SSH reconstruction or any of the other maps (not shown in the figure). None of the satellite altimeters passed over this eddy (the altimeter tracks are overlaid on the figure), but there is a clear cold signature in the SST coinciding with the reconstructed eddy. The formation process of the eddy can be seen in the SST field for the days leading up to July 7th. On July 3rd, a filament develops on the large eddy in the northeast of the domain (boxed in black in the figure). Over the following days this filament becomes unstable and sheds to form the new eddy. This formation process all happens within just 5 days and no altimeter track passes over the region during this period, making it very challenging for an SSH mapping method that only uses altimetry observations to predict this process from the available observations. Using SST observations to aid SSH mapping is thus essential if we are to capture the formation of short-lived, small-scale mesoscale eddies.

7. Discussion and Conclusions

7.1. Summary

In this study, we presented a DL framework for mapping SSH from satellite altimetry and SST observations, motivated by the close dynamical relationship between these surface ocean variables. By training and testing

our method on real-world satellite observations, we demonstrated that we could reconstruct the mesoscale SSH field with a high degree of physical realism. Including SST as additional observations caused a significant improvement in the accuracy of our reconstruction, in line with expectations from SQG theory. When compared against several other SSH mapping methods proposed in the literature, our method was able to map SSH in the Gulf Stream Extension with the highest accuracy and effective spatial resolution. Our study provides a roadmap toward replacing the decades-old OI method with state-of-the-art DL methods to produce higher-accuracy and higher-resolution global SSH maps.

7.2. Physical Implications of Improved Mesoscale SSH Field

The higher-accuracy and higher-resolution mesoscale SSH field that our DL mapping method provides is an exciting new data set for physical oceanographers. The geostrophic surface currents calculated from our SSH field resulted in a substantially (25%) higher EKE than from the DUACS SSH map. This result, if replicated in other ocean regions, suggests that current satellite altimetry-derived estimates of oceanic mesoscale eddies' contribution to the global energy budget could be a significant underestimate (Martínez-Moreno et al., 2019; Thoppil et al., 2011; Xu et al., 2014). An updated global estimate of EKE using our DL SSH mapping method is the subject of an ongoing study.

Our SSH reconstruction also resulted in relative vorticity and Okubo-Weiss quantity, W , fields that appear qualitatively to be more physically realistic, with a clearer separation between the Gulf Stream and nearby coherent mesoscale eddies. As discussed in 6.3, W is an important diagnostic for distinguishing between rotation-dominated regions (inside mesoscale eddies) and regions of high strain with SST front intensification and relatively strong vertical velocities. Thus, W is a prerequisite to get access to the 3D eddy field (Qiu et al., 2020; Siegelman et al., 2020). Another interpretation of W is that, whereas first-order SSH derivatives only give access to currents, second-order derivatives (through W) give access to accelerations and therefore to the time evolution of the eddy field (Hua & Klein, 1998; Hua et al., 1998; Lapeyre et al., 1999). One consequence is that improving the diagnosis of second-order SSH derivatives would lead us to better estimate the kinetic energy cascade, that is, the balance between the merging between eddies leading to larger eddies (inverse KE cascade) and the splitting of eddies by other eddies leading to smaller eddies (direct KE cascade) (Klein et al., 2019; Scott & Wang, 2005; Storer et al., 2022). These comments emphasize the importance of accurately diagnosing second-order SSH derivatives. Our DL SSH mapping method could allow this diagnosis to now be made from current satellite observing capabilities.

Much of what we know about coherent mesoscale eddies in the ocean comes from applying eddy tracking algorithms to SSH maps (Chelton et al., 2011; Martínez-Moreno et al., 2019, 2021). By including SST observations in the input to the SSH mapping, we demonstrated in Section 6.4 that we could reconstruct short-lived, small-scale coherent eddies that would otherwise not be captured in SSH maps. Thus, applying eddy tracking algorithms to our new SSH map could lead to changes in the conclusions drawn in these past studies, especially about the number of detected coherent eddies, their shapes and strengths.

7.3. Potential for Development of a Global DL SSH Product

While the results of this study are promising, several challenges must be overcome before a global DL SSH map can be made available to the oceanography community. The main challenge we foresee is that different regions of the ocean exhibit a diverse range of dynamical regimes. It may therefore be challenging for a single DL model to learn to map SSH accurately in all ocean regions. An approach involving an ensemble of regional DL models could be employed to overcome this challenge where bespoke models are trained for each ocean region before the resulting regional SSH maps are merged to form a global product.

The power of DL models comes from their ability to continue to improve dramatically as they are trained on larger volumes of data. While we limited ourselves here to using real-world observations for the training data set, future work could employ “transfer learning” where the DL model is trained on a large available data set drawn from a system with similar characteristics to that under consideration before being fine-tuned on the desired data set (Subel et al., 2022). In the context of SSH mapping, the abundant output from high-resolution ocean GCMs might provide a valuable data set for pre-training our DL model, though all GCMs contain imperfect representations of ocean physics and so the resulting SSH product would need to be used with caution if being used for observational validation of the GCM itself.

As highlighted in Section 6.3, CNN methods such as ours and 4DVarNet can suffer non-physical high-frequency artifacts. Without enforcing hard physical constraints on the SSH reconstruction, there is to our knowledge no way to guarantee that no such artifacts would appear in the SSH map. However, in this study we found adding terms regularizing the first and second along-track SSH derivatives significantly reduced the artifacts.

For a global DL SSH product to become widely adopted in the oceanographic community, the product must have a consistent spatial resolution that does not depend on the observational sampling. We have demonstrated here that this can be achieved by constructing a training data set where the loss function is only calculated at points not included in the input to ensure that the DL model does not learn to skillfully reproduce a high-resolution field near the input observations but with a smoother low-resolution map in the gaps between observations. The development and validation of a global SSH product using our method is the subject of an ongoing study.

The region considered in this study is particularly favorable for the use of SST to infer SSH. In regions where salinity plays a larger role in setting the surface density field and stratification, for example, at high latitudes and in the tropics, the improvement due to including SST may be less pronounced. Additionally, as discussed in Section 6.2.4, consideration will need to be given to the choice of SST product used to generate a global product since the effect of clouds on the MUR SST product used in this study could cause the accuracy and effective resolution of the SSH product to vary geographically and seasonally, posing challenges for end users seeking to interpret the dynamics of the gridded SSH field.

Other ocean observation datasets could be incorporated into a similar framework to that presented here to improve the mapping of SSH or other quantities of interest. Surface salinity and ocean color are routinely observed by satellites and may contain different dynamical signatures to the SST observations. ARGO profiles provide regular vertical profiles of temperature and salinity (among other variables) albeit with sparser horizontal sampling density than satellite observations of the surface. Since the vertical structure of mesoscale eddies plays a significant role in governing their evolution, it is possible that incorporating ARGO observations into a similar DL framework to that described here would lead to more accurate SSH mapping. Although described here in the context of satellite altimetry, our DL framework could similarly be applied to other remote sensing interpolation problems.

Finally, the recent launch of the Surface Water and Ocean Topography (SWOT) mission will soon offer, for the first time, high-resolution 2D snapshots of the SSH field through its wide swath. The development of a method to use SWOT observations to map high-resolution SSH is an area of active research (Beauchamp et al., 2020; Fablet, Amar, et al., 2021; Le Guillou et al., 2021) and our framework could also be adapted to suit the nature of SWOT observations. The launch of the SWOT mission certainly does not remove the need for a more accurate method for mapping SSH from traditional nadir altimetry. The nadir observations are available over a longer period than the likely duration of the SWOT mission, and so are an important data set for studies of inter-annual variability and climate change. SSH maps and the tracked eddies inferred from them are also widely used to aid the interpretation of in situ observations across all oceanographic disciplines, so improving the SSH maps for the pre-SWOT years would add value to many years of already-collected in situ observations. To this end, SWOT observations in future could be used as a ground truth to further validate our nadir altimeter-derived DL SSH map or calculate the loss function during training.

The effective use of a DL framework such as that outlined in this study to map the global ocean's SSH field from existing and future satellite observations would represent an exciting new paradigm for studying surface ocean dynamics and, by extension, global climate.

Appendix A: Full Performance Metric Comparison

The following Table A1 summarizes the performance of each of the mapping methods on all of the evaluation metrics considered in this study.

Table A1

Summary of the Performance Metrics Calculated for the SSH Mapping Methods Compared in This Study

Mapping methods' summary statistics						
Mapping method	RMSE [cm]	Std. of RMSE [cm]	RMSE score [no units]	Std. of RMSE score [no units]	Eff. resolution [km]	Drifter RMSE [m/s]
DUACS	7.7	2.6	0.88	0.055	149	0.349
DYMOST	6.8	2.0	0.89	0.047	131	0.342
MIOST	6.8	2.3	0.89	0.057	139	0.297
BFN-QG	7.5	2.6	0.88	0.053	119	0.290
4DVarNet SSH	6.6	1.8	0.90	0.046	107	0.271
ConvLSTM SSH	6.7	1.8	0.89	0.050	115	0.265
ConvLSTM SSH + SST	6.4	1.6	0.90	0.048	104	0.256

Note. Methods above the double horizontal line are from the literature (see the text for references) and those below are those proposed in this study, bold values indicate the highest performing method for each metric. See Section 6.1 for a description of each metric.

Data Availability Statement

The Level 3 satellite altimetry observations (CMEMS, 2023b, <https://doi.org/10.48670/moi-00146>) and AOML global drifter observations (which we took from The CMEMS in situ surface and sub-surface water velocity product: CMEMS, 2023a, <https://doi.org/10.17882/86236>) data used in this study are freely publicly available from the CMEMS data store. The GHRSSST MUR Level 4 SST product is freely publicly available from the NASA Earthdata PODAAC (NASA PODAAC, 2023, <https://doi.org/10.5067/GHGMR-4FJ04>). The Ocean Data Challenge data used in this study to compare our method to the other mapping methods (AVISO, 2021, <https://doi.org/10.24400/527896/a01-2021.005>) were developed, validated by CLS and MEOM Team from IGE (CNRS-UGA-IRD-G-INP), France and distributed by Aviso+. Code to generate training examples from these public datasets, to define and train our ConvLSTM SSH and ConvLSTM SSH + SST mapping methods, and to reproduce the results figures in this manuscript is publicly available here: <https://doi.org/10.5281/zenodo.7730739> (Martin, 2023). Also provided in this repository are the underlying data for Table A1 and .nc files containing the SSH maps we generated for the Gulf Stream study region for the year 2017.

References

- Abdalla, S., Kolahchi, A. A., Ablain, M., Adusumilli, S., Bhowmick, S. A., Alou-Font, E., et al. (2021). Altimetry for the future: Building on 25 years of progress. *Advances in Space Research*, 68(2), 319–363. <https://doi.org/10.1016/j.asr.2021.01.022>
- Abernathy, R. P., & Marshall, J. (2013). Global surface eddy diffusivities derived from satellite altimetry. *Journal of Geophysical Research: Oceans*, 118(2), 901–916. <https://doi.org/10.1002/jgrc.20066>
- Amores, A., Jordà, G., Arsouze, T., & Le Sommer, J. (2018). Up to what extent can we characterize ocean eddies using present-day gridded altimetric products? *Journal of Geophysical Research: Oceans*, 123(10), 7220–7236. <https://doi.org/10.1029/2018jc014140>
- Arbic, B. K., Scott, R. B., Chelton, D. B., Richman, J. G., & Shriver, J. F. (2012). Effects of stencil width on surface ocean geostrophic velocity and vorticity estimation from gridded satellite altimeter data. *Journal of Geophysical Research*, 117(C3), C03029. <https://doi.org/10.1029/2011jc007367>
- AVISO. (2021). Data challenge ose—2021a SSH mapping OSE [Dataset]. AVISO. <https://doi.org/10.24400/527896/a01-2021.005>
- Ballarotta, M., Ubelmann, C., Pujol, M.-I., Taburet, G., Fournier, F., Legeais, J.-F., et al. (2019). On the resolutions of ocean altimetry maps. *Ocean Science*, 15(4), 1091–1109. <https://doi.org/10.5194/os-15-1091-2019>
- Ballarotta, M., Ubelmann, C., Rogé, M., Fournier, F., Faugère, Y., Dibarboue, G., et al. (2020). Dynamic mapping of along-track ocean altimetry: Performance from real observations. *Journal of Atmospheric and Oceanic Technology*, 37(9), 1593–1601. <https://doi.org/10.1175/jtech-d-20-0030.1>
- Barth, A., Alvera-Azcárate, A., Licer, M., & Beckers, J.-M. (2020). DINCAE 1.0: A convolutional neural network with error estimates to reconstruct sea surface temperature satellite observations. *Geoscientific Model Development*, 13(3), 1609–1622. <https://doi.org/10.5194/gmd-13-1609-2020>
- Barth, A., Alvera-Azcárate, A., Troupin, C., & Beckers, J.-M. (2022). DINCAE 2.0: Multivariate convolutional neural network with error estimates to reconstruct sea surface temperature satellite and altimetry observations. *Geoscientific Model Development*, 15(5), 2183–2196. <https://doi.org/10.5194/gmd-15-2183-2022>
- Beauchamp, M., Fablet, R., Ubelmann, C., Ballarotta, M., & Chapron, B. (2020). Intercomparison of data-driven and learning-based interpolations of along-track nadir and wide-swath SWOT altimetry observations. *Remote Sensing*, 12(22), 3806. <https://doi.org/10.3390/rs12223806>
- Bowen, M. M., Emery, W. J., Wilkin, J. L., Tildesley, P. C., Barton, I. J., & Knewton, R. (2002). Extracting multiyear surface currents from sequential thermal imagery using the maximum cross-correlation technique. *Journal of Atmospheric and Oceanic Technology*, 19(10), 1665–1676. [https://doi.org/10.1175/1520-0426\(2002\)019<1665:emscfs>2.0.co;2](https://doi.org/10.1175/1520-0426(2002)019<1665:emscfs>2.0.co;2)
- Bretherton, F. P., Davis, R. E., & Fandry, C. (1976). A technique for objective analysis and design of oceanographic experiments applied to MODE-73. *Deep-Sea Research and Oceanographic Abstracts*, 23(7), 559–582. [https://doi.org/10.1016/0011-7471\(76\)90001-2](https://doi.org/10.1016/0011-7471(76)90001-2)

Acknowledgments

The research was funded by the National Aeronautics and Space Administration under Grant 80NSSC21K1187 issued through the Science Mission Directorate, Ocean Surface Topography Science Team program. S.A.M. received further financial support from the Theodore H. and Marie M. Sarchin Endowed Fellowship in Oceanography. P.K. acknowledges support from the SWOT Science Team and the QuikSCAT mission. The authors acknowledge helpful discussions with Steven Brunton, Jinbo Wang, Brian Arbic, Paige Lavin, and Maxime Ballarotta and thank all those involved in creating and maintaining the Ocean Data Challenge. S.A.M. thanks Arthur Nowell for writing support during the preparation of the first draft of this manuscript. G.E.M. thanks Charles Trimble for funding conceptual studies of DL methods for SSH mapping at the California Institute of Technology. The authors thank two anonymous reviewers for their reviews which have helped to improve the manuscript.

- Buongiorno Nardelli, B., Cavaliere, D., Charles, E., & Ciani, D. (2022). Super-resolving ocean dynamics from space with computer vision algorithms. *Remote Sensing*, *14*(5), 1159. <https://doi.org/10.3390/rs14051159>
- Chelton, D. B., Schlax, M. G., & Samelson, R. M. (2011). Global observations of nonlinear mesoscale eddies. *Progress in Oceanography*, *91*(2), 167–216. <https://doi.org/10.1016/j.pocean.2011.01.002>
- Chen, W., Mied, R. P., & Shen, C. Y. (2008). Near-surface ocean velocity from infrared images: Global optimal solution to an inverse model. *Journal of Geophysical Research*, *113*(C10), C10003. <https://doi.org/10.1029/2008jc004747>
- CMEMS. (2023a). Copernicus marine in situ—Global ocean-delayed mode in situ observations of surface (drifters, HFR) and sub-surface (vessel-mounted ADCPs) water velocity [Dataset]. CMEMS. <https://doi.org/10.17882/86236>
- CMEMS. (2023b). Global ocean along track L3 sea surface heights reprocessed 1993 ongoing tailored for data assimilation [Dataset]. CMEMS. <https://doi.org/10.48670/moi-00146>
- CMEMS. (2023c). Global ocean gridded L4 sea surface heights and derived variables reprocessed 1993 ongoing [Dataset]. CMEMS. <https://doi.org/10.48670/moi-00148>
- Dufau, C., Orszynowicz, M., Dibarboire, G., Morrow, R., & Le Traon, P.-Y. (2016). Mesoscale resolution capability of altimetry: Present and future. *Journal of Geophysical Research: Oceans*, *121*(7), 4910–4927. <https://doi.org/10.1002/2015jc010904>
- Emery, W. J., Thomas, A., Collins, M., Crawford, W. R., & Mackas, D. (1986). An objective method for computing advective surface velocities from sequential infrared satellite images. *Journal of Geophysical Research*, *91*(C11), 12865–12878. <https://doi.org/10.1029/jc091ic11p12865>
- Fablet, R., Amar, M. M., Febvre, Q., Beauchamp, M., & Chapron, B. (2021). End-to-end physics-informed representation learning for satellite ocean remote sensing data: Applications to satellite altimetry and sea surface currents. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, *V-3*(3), 295–302. <https://doi.org/10.5194/isprs-annals-v-3-2021-295-2021>
- Fablet, R., & Chapron, B. (2022). Multimodal learning-based inversion models for the space-time reconstruction of satellite-derived geophysical fields. arXiv preprint arXiv:2203.10640.
- Fablet, R., Chapron, B., Drumetz, L., Mémin, E., Pannekoucke, O., & Rousseau, F. (2021). Learning variational data assimilation models and solvers. *Journal of Advances in Modeling Earth Systems*, *13*(10), e2021MS002572. <https://doi.org/10.1029/2021ms002572>
- Fablet, R., Febvre, Q., & Chapron, B. (2022). Multimodal 4DVarNets for the reconstruction of sea surface dynamics from SST-SSH synergies. arXiv preprint arXiv:2207.01372.
- Fu, L.-L., Chelton, D. B., Le Traon, P.-Y., & Morrow, R. (2010). Eddy dynamics from satellite altimetry. *Oceanography*, *23*(4), 14–25. <https://doi.org/10.5670/oceanog.2010.02>
- George, T. M., Manucharyan, G. E., & Thompson, A. F. (2021). Deep learning to infer eddy heat fluxes from sea surface height patterns of mesoscale turbulence. *Nature Communications*, *12*(1), 1–11. <https://doi.org/10.1038/s41467-020-20779-9>
- González-Haro, C., & Isern-Fontanet, J. (2014). Global ocean current reconstruction from altimetric and microwave SST measurements. *Journal of Geophysical Research: Oceans*, *119*(6), 3378–3391. <https://doi.org/10.1002/2013jc009728>
- Hausmann, U., & Czaja, A. (2012). The observed signature of mesoscale eddies in sea surface temperature and the associated heat transport. *Deep Sea Research Part I: Oceanographic Research Papers*, *70*, 60–72. <https://doi.org/10.1016/j.dsr.2012.08.005>
- 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).
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, *2*(5), 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
- Hua, B., & Klein, P. (1998). An exact criterion for the stirring properties of nearly two-dimensional turbulence. *Physica D: Nonlinear Phenomena*, *113*(1), 98–110. [https://doi.org/10.1016/s0167-2789\(97\)00143-7](https://doi.org/10.1016/s0167-2789(97)00143-7)
- Hua, B., McWilliams, J., & Klein, P. (1998). Lagrangian accelerations in geostrophic turbulence. *Journal of Fluid Mechanics*, *366*, 87–108. <https://doi.org/10.1017/s0022112098001001>
- Isern-Fontanet, J., Chapron, B., Lapeyre, G., & Klein, P. (2006). Potential use of microwave sea surface temperatures for the estimation of ocean currents. *Geophysical Research Letters*, *33*(24), L24608. <https://doi.org/10.1029/2006gl027801>
- Isern-Fontanet, J., Lapeyre, G., Klein, P., Chapron, B., & Hecht, M. W. (2008). Three-dimensional reconstruction of oceanic mesoscale currents from surface information. *Journal of Geophysical Research*, *113*(C9), C09005. <https://doi.org/10.1029/2007jc004692>
- Isern-Fontanet, J., Shinde, M., & González-Haro, C. (2014). On the transfer function between surface fields and the geostrophic stream function in the Mediterranean Sea. *Journal of Physical Oceanography*, *44*(5), 1406–1423. <https://doi.org/10.1175/jpo-d-13-0186.1>
- Kelly, K. A. (1989). An inverse model for near-surface velocity from infrared images. *Journal of Physical Oceanography*, *19*(12), 1845–1864. [https://doi.org/10.1175/1520-0485\(1989\)019<1845:aifms>2.0.co;2](https://doi.org/10.1175/1520-0485(1989)019<1845:aifms>2.0.co;2)
- Klein, P., Lapeyre, G., Siegelman, L., Qiu, B., Fu, L.-L., Torres, H., et al. (2019). Ocean-scale interactions from space. *Earth and Space Science*, *6*(5), 795–817. <https://doi.org/10.1029/2018ea000492>
- LaCasce, J., & Mahadevan, A. (2006). Estimating subsurface horizontal and vertical velocities from sea-surface temperature. *Journal of Marine Research*, *64*(5), 695–721. <https://doi.org/10.1357/002224006779367267>
- Lapeyre, G. (2009). What mesoscale signal does the altimeter see? On the decomposition in baroclinic modes and the role of the surface boundary condition. *Journal of Physical Oceanography*, *39*(11), 2857–2874. <https://doi.org/10.1175/2009jpo3968.1>
- Lapeyre, G., & Klein, P. (2006). Dynamics of the upper oceanic layers in terms of surface quasigeostrophy theory. *Journal of Physical Oceanography*, *36*(2), 165–176. <https://doi.org/10.1175/jpo2840.1>
- Lapeyre, G., Klein, P., & Hua, B. (1999). Does the tracer gradient vector align with the strain eigenvectors in 2D turbulence? *Physics of Fluids*, *11*(12), 3729–3737. <https://doi.org/10.1063/1.870234>
- Le Guillou, F., Metref, S., Cosme, E., Ubelmann, C., Ballarotta, M., Le Sommer, J., & Verron, J. (2021). Mapping altimetry in the forthcoming SWOT era by back-and-forth nudging a one-layer quasigeostrophic model. *Journal of Atmospheric and Oceanic Technology*, *38*(4), 697–710. <https://doi.org/10.1175/jtech-d-20-0104.1>
- Lellouche, J.-M., Eric, G., Romain, B.-B., Gilles, G., Angélique, M., Marie, D., et al. (2021). The copernicus global 1/12 oceanic and sea ice GLORYS12 reanalysis. *Frontiers of Earth Science*, *9*, 698876.
- 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. [https://doi.org/10.1175/1520-0426\(1998\)015<0522:aimmom>2.0.co;2](https://doi.org/10.1175/1520-0426(1998)015<0522:aimmom>2.0.co;2)
- Manucharyan, G. E., Siegelman, L., & Klein, P. (2021). A deep learning approach to spatiotemporal sea surface height interpolation and estimation of deep currents in geostrophic ocean turbulence. *Journal of Advances in Modeling Earth Systems*, *13*(1), e2019MS001965. <https://doi.org/10.1029/2019ms001965>
- Martin, S. A. (2023). smartin98/deep-learning-ssh-mapping-JAMES-paper: v1.0.0 (Version v1.0.0) [Computer software]. Zenodo. <https://doi.org/10.5281/ZENODO.7730738>

- Martínez-Moreno, J., Hogg, A. M., England, M. H., Constantinou, N. C., Kiss, A. E., & Morrison, A. K. (2021). Global changes in oceanic mesoscale currents over the satellite altimetry record. *Nature Climate Change*, *11*(5), 397–403. <https://doi.org/10.1038/s41558-021-01006-9>
- Martínez-Moreno, J., Hogg, A. M., Kiss, A. E., Constantinou, N. C., & Morrison, A. K. (2019). Kinetic energy of eddy-like features from sea surface altimetry. *Journal of Advances in Modeling Earth Systems*, *11*(10), 3090–3105. <https://doi.org/10.1029/2019ms001769>
- McWilliams, J. C. (1984). The emergence of isolated coherent vortices in turbulent flow. *Journal of Fluid Mechanics*, *146*, 21–43. <https://doi.org/10.1017/s0022112084001750>
- NASA PODAAC. (2023). GHRSSST level 4 MUR global foundation sea surface temperature analysis (v4.1) [Dataset]. NASA. <https://doi.org/10.5067/GHGMR-4FJ04>
- Okubo, A. (1970). Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences. *Deep-Sea Research and Oceanographic Abstracts*, *17*(3), 445–454. [https://doi.org/10.1016/0011-7471\(70\)90059-8](https://doi.org/10.1016/0011-7471(70)90059-8)
- Ostrovskii, A., & Piterbarg, L. (1995). Inversion for heat anomaly transport from sea surface temperature time series in the northwest Pacific. *Journal of Geophysical Research*, *100*(C3), 4845–4865. <https://doi.org/10.1029/94jc03041>
- Qiu, B., Chen, S., Klein, P., Torres, H., Wang, J., Fu, L.-L., & Menemenlis, D. (2020). Reconstructing upper-ocean vertical velocity field from sea surface height in the presence of unbalanced motion. *Journal of Physical Oceanography*, *50*(1), 55–79. <https://doi.org/10.1175/jpo-d-19-0172.1>
- Rio, M.-H., Santoleri, R., Bourdalle-Badie, R., Griffa, A., Piterbarg, L., & Taburet, G. (2016). Improving the altimeter-derived surface currents using high-resolution sea surface temperature data: A feasibility study based on model outputs. *Journal of Atmospheric and Oceanic Technology*, *33*(12), 2769–2784. <https://doi.org/10.1175/jtech-d-16-0017.1>
- Scott, R. B., & Wang, F. (2005). Direct evidence of an oceanic inverse kinetic energy cascade from satellite altimetry. *Journal of Physical Oceanography*, *35*(9), 1650–1666. <https://doi.org/10.1175/jpo2771.1>
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., & Woo, W.-c. (2015). Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in Neural Information Processing Systems*, *28*.
- Siegelman, L., Klein, P., Thompson, A. F., Torres, H. S., & Menemenlis, D. (2020). Altimetry-based diagnosis of deep-reaching sub-mesoscale ocean fronts. *Fluid*, *5*(3), 145. <https://doi.org/10.3390/fluids5030145>
- 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. [https://doi.org/10.1175/1520-0485\(2001\)031<0554:tsaecom>2.0.co;2](https://doi.org/10.1175/1520-0485(2001)031<0554:tsaecom>2.0.co;2)
- Sonnenwald, M., Lguensat, R., Jones, D. C., Dueben, P., Brajard, J., & Balaji, V. (2021). Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, *16*(7), 073008. <https://doi.org/10.1088/1748-9326/ac0eb0>
- Storer, B. A., Buzzicotti, M., Khatri, H., Griffies, S. M., & Aluie, H. (2022). Global energy spectrum of the general oceanic circulation. *Nature Communications*, *13*(1), 1–9. <https://doi.org/10.1038/s41467-022-33031-3>
- Subel, A., Guan, Y., Chattopadhyay, A., & Hassanzadeh, P. (2022). Explaining the physics of transfer learning a data-driven subgrid-scale closure to a different turbulent flow. arXiv preprint arXiv:2206.03198.
- Sun, Z., Sandoval, L., Crystal-Ornelas, R., Mousavi, S. M., Wang, J., Lin, C., et al. (2022). A review of Earth artificial intelligence. *Computers & Geosciences*, *105*034.
- Taburet, G., Sanchez-Roman, A., Ballarotta, M., Pujol, M.-I., Legeais, J.-F., Fournier, F., et al. (2019). DUACS DT2018: 25 years of reprocessed sea level altimetry products. *Ocean Science*, *15*(5), 1207–1224. <https://doi.org/10.5194/os-15-1207-2019>
- Thoppil, P. G., Richman, J. G., & Hogan, P. J. (2011). Energetics of a global ocean circulation model compared to observations. *Geophysical Research Letters*, *38*(15), L15607. <https://doi.org/10.1029/2011gl048347>
- 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. <https://doi.org/10.1175/jtech-d-15-0163.1>
- Ubelmann, C., Dibarboure, G., Gaultier, L., Ponte, A., Arduin, F., Ballarotta, M., & Faugère, Y. (2021). Reconstructing ocean surface current combining altimetry and future spaceborne Doppler data. *Journal of Geophysical Research: Oceans*, *126*(3), e2020JC016560. <https://doi.org/10.1029/2020jc016560>
- 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. <https://doi.org/10.1175/jtech-d-14-00152.1>
- Weiss, J. (1991). The dynamics of enstrophy transfer in two-dimensional hydrodynamics. *Physica D: Nonlinear Phenomena*, *48*(2–3), 273–294. [https://doi.org/10.1016/0167-2789\(91\)90088-q](https://doi.org/10.1016/0167-2789(91)90088-q)
- Xu, C., Shang, X.-D., & Huang, R. X. (2014). Horizontal eddy energy flux in the world oceans diagnosed from altimetry data. *Scientific Reports*, *4*(1), 1–7. <https://doi.org/10.1038/srep05316>