



HAL
open science

ORCAst: Operational High-Resolution Current Forecasts

Pierre Garcia, Inès Larroche, Amélie Pesnec, Hannah Bull, Théo Archambault,
Evangelos Moschos, Alexandre Stegner, Anastase Charantonis, Dominique
Béréziat

► **To cite this version:**

Pierre Garcia, Inès Larroche, Amélie Pesnec, Hannah Bull, Théo Archambault, et al.. ORCAst: Operational High-Resolution Current Forecasts. Artificial Intelligence for the Earth Systems, In press. <hal-05052411>

HAL Id: hal-05052411

<https://hal.sorbonne-universite.fr/hal-05052411v1>

Submitted on 30 Apr 2025

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.



HAL Authorization

ORCAst: Operational High-Resolution Current Forecasts

PIERRE GARCIA,^{a,c,*} INÈS LARROCHE,^{a,d,*} AMÉLIE PESNEC,^{a,*} HANNAH BULL,^{a,*} THÉO ARCHAMBAULT,^a
EVANGELOS MOSCHOS,^a ALEXANDRE STEGNER,^a ANASTASE CHARANTONIS,^b AND DOMINIQUE BÉRÉZIAT^c

^a *Amphitrite*

^b *Inria, Sorbonne Université*

^c *Sorbonne Université, CNRS, LIP6*

^d *Univ. Brest, CNRS, Ifremer, LOPS*

ABSTRACT: We present ORCAst, a multi-stage, multi-arm network for Operational high-Resolution Current forecAstS over one week. Producing real-time nowcasts and forecasts of ocean surface currents is a challenging problem due to indirect or incomplete information from satellite remote sensing data. Entirely trained on real satellite data and in situ measurements from drifters, our model learns to forecast global ocean surface currents using various sources of ground truth observations in a multi-stage learning procedure. Our multi-arm encoder-decoder model architecture allows us to first predict sea surface height and geostrophic currents from larger quantities of nadir and SWOT altimetry data, before learning to predict ocean surface currents from much more sparse in situ measurements from drifters. Training our model on specific regions improves performance. Our model achieves stronger nowcast and forecast performance in predicting ocean surface currents than various state-of-the-art methods.

SIGNIFICANCE STATEMENT: Our study introduces a novel neural network designed to produce high-resolution, real-time 7-day forecasts of ocean surface currents. Accurate forecasting of ocean currents is important for ship routing, climate studies and tracking of pollutants. This is a challenging problem, due to sparse in situ observations, indirect or incomplete satellite observations, as well as complex ocean dynamics. We develop a multi-stage training procedure to learn ocean currents by progressively refining forecasts with data of increasing quality. Our method only uses real satellite and in situ data, including the latest generation of satellite altimetry from the recently launched SWOT satellite, and does not depend on numerical simulations.

1. Introduction

Ocean currents are an important component of Earth’s climate system, influencing weather patterns, marine ecosystems, and global heat distribution. Nevertheless, accurate reconstructions and forecasts of ocean surface current fields is a challenging task. There are two main types of observational data for estimating ocean surface currents: in situ measurements, for example from drifters, and satellite observations, for example from satellite altimetry and from multispectral optical satellites. These observational data can be assimilated in numerical models or used in data-driven methods to reconstruct and forecast ocean currents.

In situ measurements, such as those from drifter buoys transported by the currents, allow direct measurements

of ocean surface currents (Ifremer 2024; Elipot et al. 2022). However, drifters are sparse and unevenly distributed across the global ocean. Drifters are commonly used for evaluating the performance of models for estimating ocean currents (Kugusheva et al. 2024; Martin et al. 2024; Le Guillou et al. 2023b). In our work, we use drifters both as ground truth data to fine-tune our model, as well as to evaluate our method. Other sources of in situ data, for example from instruments on board ships, are generally not publically available, but limited examples of ship routes can be used to evaluate the performance of ocean current nowcasts and forecasts (Kugusheva et al. 2024).

Satellite altimetry measures sea surface height (SSH), which is related to surface currents through the geostrophic balance approximation, i.e. an equilibrium between the Coriolis and pressure forces. Far from the Equator, geostrophic currents are the main component of the total ocean surface current. Since their first launch in 1978, nadir-pointing altimeters allow low resolution estimation of SSH through sparse along-track measurements (Tapley et al. 1982). A widespread method of interpolating these sparse measures of SSH is DUACS, the Data Unification and Altimeter Combination System (Taburet et al. 2019). This method is a global optimal interpolation, a linear interpolation method with optimally chosen coefficients of covariance matrices with respect to expert knowledge and historical data. However, due to the limited by the number of along-track measurements, the global maps of SSH provided by this optimal interpolation have a coarse resolution from $1/4^\circ$ to $1/8^\circ$ in specific regional areas. In February 2022, a new satellite called Surface Water and Ocean Topography (SWOT) was launched, providing high-resolution SSH data, below the kilometer, with a 120 km

Corresponding author: Pierre Garcia, pierre.garcia@amphitrite.fr

* These authors contributed equally to this work.

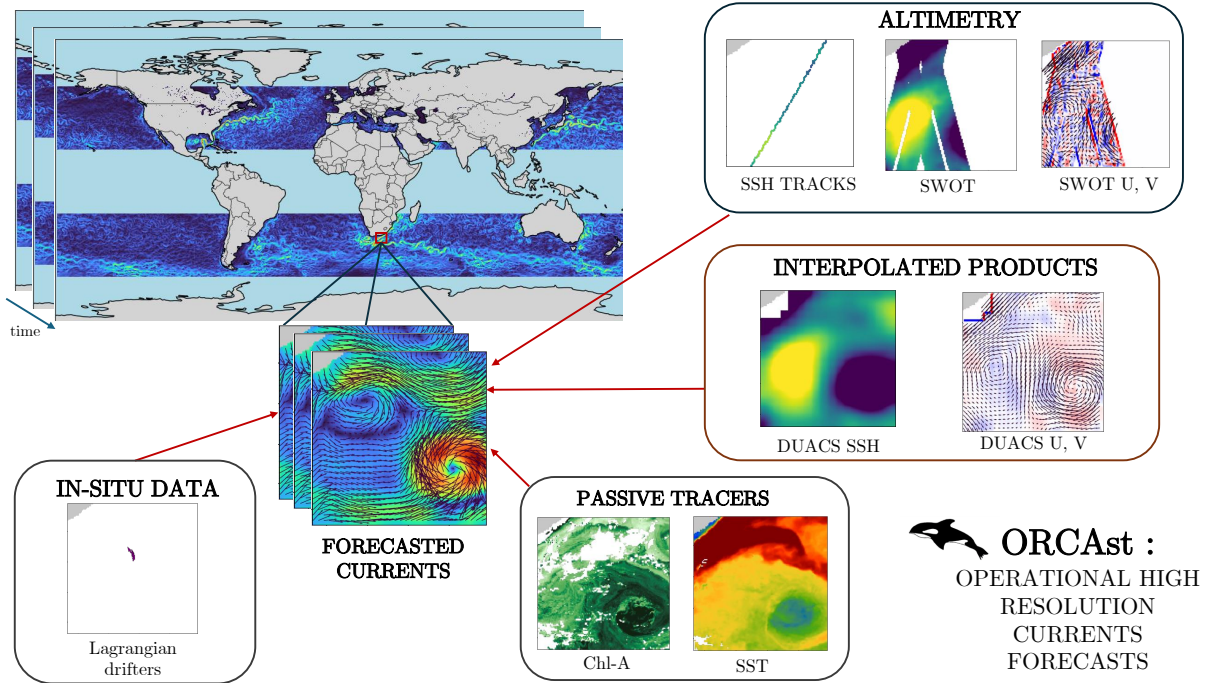


Fig. 1. ORCAst is an operational neural network model that achieves state-of-the-art forecasts of ocean surface currents in extratropical latitudes. It is trained exclusively on observational data, including drifters, SWOT and nadir altimetry, SST and CHL observations, in a multi-stage training process.

swath, significantly increasing the potential of satellite altimetry in measuring small-scale structures (Morrow et al. 2018).

In addition to SSH, ocean currents also have a physical signature on other remote sensing variables such as Sea Surface Temperature (SST) and chlorophyll-a (CHL). SST and CHL are passive tracers, or scalar quantities transported by the currents in via advection dynamics. They are measured by multispectral optical satellites, which allows for wide data coverage and high-resolution images (300m to 1km) when no clouds are present (Emery et al. 1989). Numerous studies demonstrate the importance of including multiple physical variables in reconstructing reliable surface currents (Thiria et al. 2023; Fablet et al. 2024; Martin et al. 2023; Archambault et al. 2024b; Kugusheva et al. 2024; Ciani et al. 2024).

There are two main methods of forecasting ocean surface currents using available drifter observations or observations from satellites: numerical models based on physical equations and data-driven approaches. Ocean Global Circulation Models (OGCMs) simulate ocean circulation using differential equations describing fluid dynamics (Tonani et al. 2015). Starting from an estimation of the ocean state, OGCMs can forecast key physical variables. Initial conditions of the model can be estimated by assimilating observational data from satellites or in situ sensors. For example, the Mercator global forecasting system (Drévilleon

et al. 2008; Mercator Océan International 2024) produces a global 10-day forecast of the ocean, by using the NEMO 3.6 model (Gurvan et al. 2017) and by assimilating satellite observations of SSH, SST, Sea Ice Concentration, as well as in situ measurements. Physical simulations produce physically realistic fields at the cost of complex hyperparameter tuning and computationally expensive inference, although some methods can be used to reduce the computational burden of data assimilation (Le Guillou et al. 2023a; Ubelmann et al. 2022).

Data-driven approaches, such as deep learning methods, use large quantities of data to learn the evolution of the physical system, rather than modelling the system dynamics through explicit equations. Deep neural networks have demonstrated high accuracy both in forecasting atmospheric evolution (Lam et al. 2023; Chen et al. 2023), and in forecasting ocean evolution (Wang et al. 2024; Aouni et al. 2024). However, the aforementioned approaches are trained on reanalysis data from numerical models, and not purely on observational data. More closely related to our work, various methods produce forecasts or time series of SSH and ocean surface currents using either numerical simulations or observational data for training, or a mix of both. Filoche et al. (2022) introduce a loss computed on nadir SSH observations only, and use it to perform a variational interpolation with a deep image prior. Archambault et al. (2023) extends this idea to leverage SST

information. Martin et al. (2023) present a self-supervised neural network to perform delayed time SSH interpolation from partial SSH nadir-pointing and SST satellite observations. Using a related approach, Archambault et al. (2024a) shows that multi-stage training, i.e., using different datasets successively in the learning phase, increases SSH reconstruction quality. However, all these methods produce a delayed-time series of SSH fields, and must be adapted to forecast ocean surface currents.

In this work, we introduce ORCAst, a data-driven method for Operational, high-Resolution Current forecasts, illustrated in Figure 1. The novel features of this model are the following:

- ORCAst is an operational neural network model to forecast SSH and ocean surface currents, without being limited to the geostrophic approximation, and achieving state-of-the-art performance
- ORCAst is trained purely on observational data, including drifters, SWOT and nadir altimetry, as well as SST and CHL observations,
- ORCAst learns high-resolution ocean surface currents through a three-phase training procedure, which progressively refines forecasts with ground-truth data of increasing quality.

The article is structured as follows: in Section 2, we describe the in situ and satellite data used in this study. Section 3 provides a detailed explanation of the architecture of our model ORCAst and our evaluation strategy. In Section 4, we present the experimental results demonstrating the operational benefits of ORCAst. Finally, in Section 5, we discuss the conclusions and offer perspectives for future work.

2. Data

Ocean surface currents can be estimated from direct and indirect measurements from different sources of in situ and satellite remote sensing data. In this section, we provide details on the different data sources used to train or evaluate our model: in situ measurements from drifters and from instruments on board ships (Sec. a), SSH measurements from nadir-pointing altimeters and SWOT, as well as derived geostrophic currents (Sec. b), in addition to visible and infrared satellite images of chlorophyll and sea surface temperature (Sec. c). We also describe our data preprocessing pipeline (Sec. d).

a. In situ measurements

We use two types of in situ measurements: sparse observations from drifters and estimations of ocean currents from instruments on board ships.

Lagrangian drifters are buoys equipped with a drogue and a GPS, and are advected by ocean currents. Our

dataset contains hourly measurements of global drifters (Ifremer 2024; Elipot et al. 2022, 2016). In order to smooth noise, for instance due to surface winds and waves, we compute moving averages of observations over periods of 24 hours. Drifters are sparsely and non-uniformly distributed, which can bias learning and evaluation. Figure 2 illustrates the spatial distribution of drifters in the Mediterranean, Gulf Stream and Aghulas region in 2023.

On-board instruments from ships also estimate the effect of currents, by measuring both Speed Through Water (STW) and Speed Over Ground (SOG). STW can be measured with instruments such as Wavex or Acoustic Doppler Current Profiler (ADCP), and SOG can easily be measured by regular GPS localisation. The difference between these measurements, $SOG - STW$, is an indicator of the effect of currents felt by ships along their trajectory. This data is limited to a specific examples of ship voyages and thus only used for evaluation purposes.

b. Satellite altimeter measurements

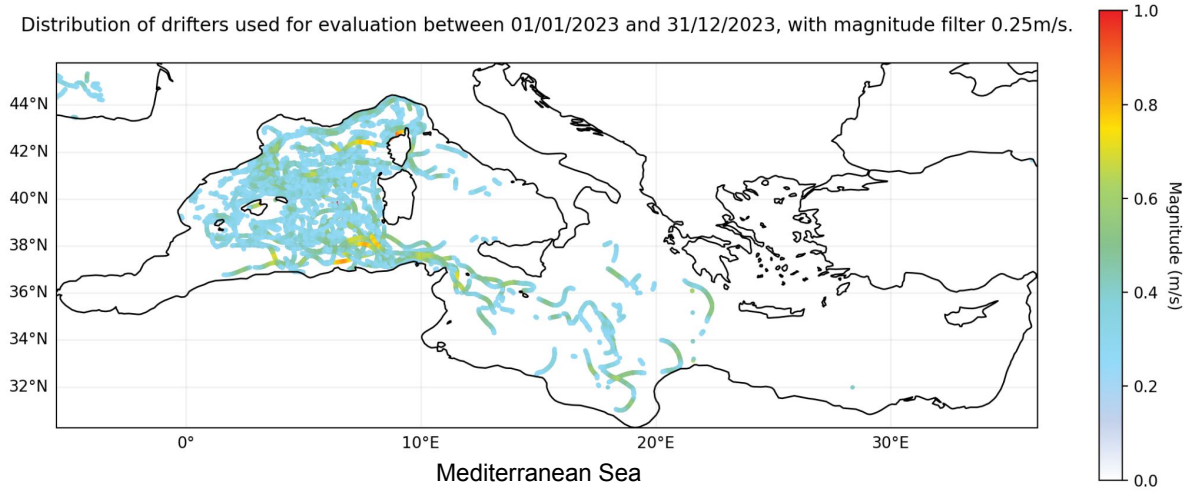
Far from the Equator, the ocean circulation is primarily driven by its geostrophic component, which can be derived from SSH. Two main types of satellite measure the SSH: nadir-pointing altimeters and the Ka-band Radar Interferometer (KaRIn) on board the SWOT satellite. Numerous interpolation methods exist to produce complete estimations of SSH and geostrophic currents such as Taburet et al. (2019) DUACS and Scott et al. (2024) NEUROST.

Nadir-pointing altimeters such as those used in various satellite missions like Sentinel-6A, Jason-3, and Cryosat-2, measure SSH along the satellite’s track with a spatial resolution of approximately 7 km. The data is processed through the DUACS protocol, which ensures consistency across multiple altimeter missions by standardizing them to a reference mission (Copernicus Marine Service 2023).

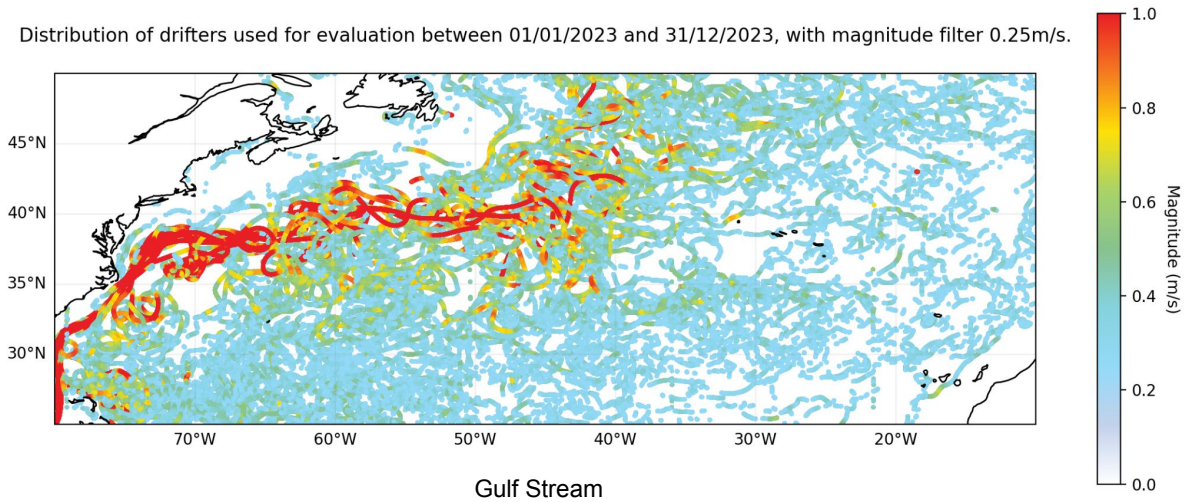
KaRIn altimetry from SWOT provides exciting opportunities for precise reconstruction of the SSH. SWOT provides 120 km-wide swaths of high-resolution observations (NASA/JPL and CNES 2024). This precision allows observation of mesoscale/submesoscale SSH patterns. However, only one satellite is presently in orbit, its revisit time period is 21 days, which leaves important areas unobserved. We have access to around one year of SWOT observations, compared to decades of nadir-pointing altimeter data. Figure 3 shows examples of fine-scale eddies visible from SWOT, but also internal gravitational waves close to the Equator. This illustrates SWOT’s potential in estimating ocean surface currents, but also its limitations, particularly in equatorial latitudes.

DUACS is a Level 4 (L4) product, providing a low-resolution (0.125°) mapping of SSH through optimal interpolation (Taburet et al. 2019) of nadir altimetry. The geostrophic currents can then be derived from this mapping, and are available at (Copernicus Marine Service

Distribution of drifters used for evaluation between 01/01/2023 and 31/12/2023, with magnitude filter 0.25m/s.



Distribution of drifters used for evaluation between 01/01/2023 and 31/12/2023, with magnitude filter 0.25m/s.



Distribution of drifters used for evaluation between 01/01/2023 and 31/12/2023, with magnitude filter 0.25m/s.

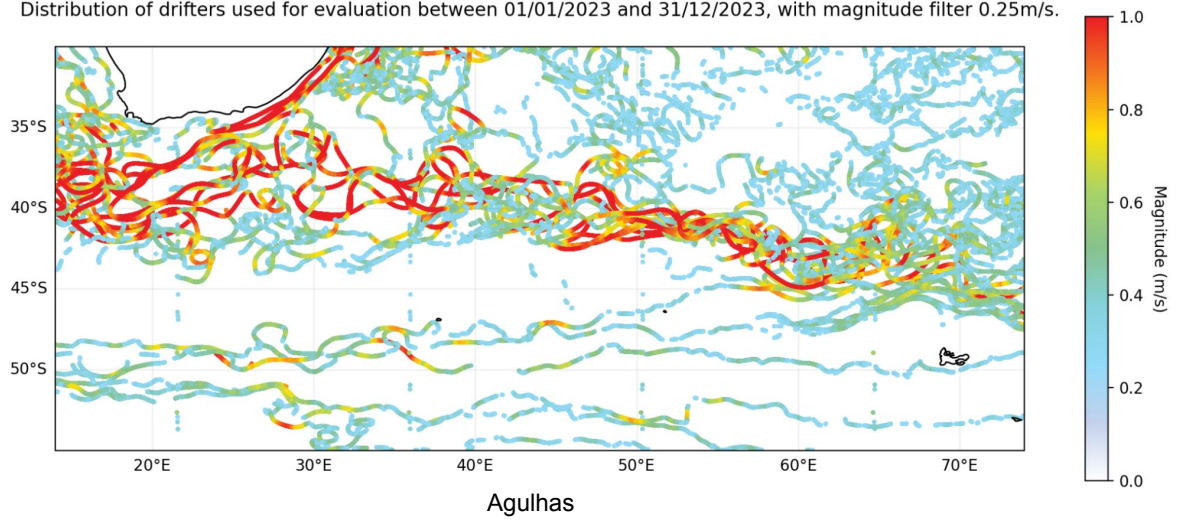


Fig. 2. In situ observations of currents from drifters in three key regions in 2023, used for evaluation.

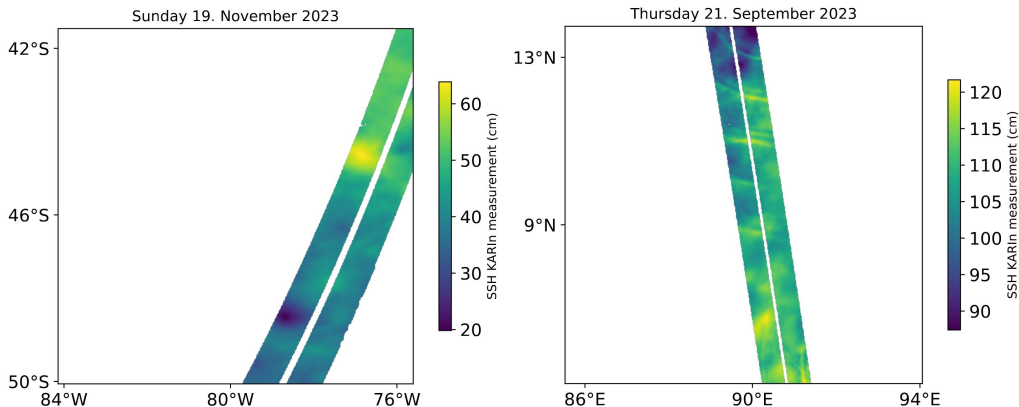


FIG. 3. SWOT SSH observation, consisting of two 50 km wide bands, separated by a 20 km gap covered by traditional nadir altimetry instruments. On the left, we note visible cyclonic and anti-cyclonic eddies at high resolution. On the right, we note visible internal gravitational waves. In particular, close to the Equator, it is more difficult to estimate ocean surface currents from SSH as a result of the reduced Coriolis effect.

2024; Copernicus Climate Service 2024). For evaluation and stage 1 training, we use a delayed-time DUACS product, which uses 6 days of future observations from nadir tracks.

NeurOST SSH is a state-of-the-art method for predicting global SSH at 0.1° resolution. Martin et al. (2024) train a neural network to predict nadir-pointing along-track SSH from SST and SSH inputs, using only observations in a self-supervised manner. This product is produced in delayed-time, integrating 14 days of future observations¹.

c. Satellite imagery

SST and CHL concentration are passive tracers of ocean currents, as water motion advects these scalar variables. Small-scale structures in ocean surface currents are visible using high-resolution satellite imagery of SST and CHL. However, both infrared and visible waves are blocked by clouds, and so information provided from these sensors is partial. Although methods exist for interpolating SST and CHL, we use L3 and L2 products and allow our model to learn to combine these different sources of information from past days.

SST is measured using infrared sensors at a high spatial resolution. We use the daily L3 products (glo 2024; MET Norway 2023), with a resolution of 0.02° .

Chlorophyll-a (CHL) is measured using data from the Ocean and Land Colour Instrument (OLCI) aboard Sentinel-3 satellites, which provides ocean color geophysical products. The chlorophyll product is provided using the OC4Me algorithm (O’Reilly et al. 1998). This algorithm calculates the concentration of algal pigments by analyzing

water-leaving reflectances across 21 spectral bands. The chlorophyll concentrations are provided in log10-scaled values. We use data available at 1.2 km resolution (EU-METSAT for Copernicus 2024).

d. Data processing pipeline

Our data processing pipeline interpolates all data sources into a common grid at resolution $1/30^\circ$. We compute means and standard deviations of the key variables - SSH, U and V components of ocean currents, SST and CHL - at a spatial resolution of 2° and a temporal resolution of 1 week. To estimate the means and standard deviations of SSH and ocean surface currents, we use the DUACS L4 product. We normalise all variables using these values to remove average seasonal and spatial effects.

3. Method

ORCAst is a deep learning model to forecast ocean currents. Our model takes as input T timesteps of past and present satellite observations of SSH and optical satellite data, and predicts τ future timesteps of SSH along with the U and V components of ocean surface currents. We learn the parameters θ of the neural network f_θ , shown in Equation 1,

$$f_\theta \left(X^{0:T} \right) = \widehat{Y}_{SSH}^{T+1:T+\tau}, \widehat{Y}_{U,V}^{T+1:T+\tau} \quad (1)$$

where X designates past and present observations and \widehat{Y} the forecasted fields. The inputs and outputs are gridded regional crops of dimension $h \times w$ and at identical spatial resolutions. In practice, we input $T = 11$ days of past and present satellite observations and output $\tau = 7$ future days of satellite observations. Inputs X and outputs \widehat{Y} are daily images of size 128×128 at $1/30^\circ$ resolution.

¹The authors of NeurOST have informed us that there were some issues in production of this data, potentially reducing quality, and that there will be a new release in the near future. We were unaware of this issue at the time of analysis, and have used the NeurOST data as provided in November 2024.

In Section a, we present our multi-arm encoder-decoder model architecture and in Section b, we describe our three-stage training procedure. Section c provides the geographic zones used for training and evaluation, and Section d outlines our evaluation metrics and our baseline comparisons.

a. Model architecture

Our model architecture, illustrated in Figure 4, is based on SimVP, a flexible encoder-decoder architecture initially designed for video prediction (Gao et al. 2022). A similar architecture is used to reconstruct a time series of SSH from nadir altimetry and SST inputs by Martin et al. (2024). The ORCAst model has the following key features:

- one 2D encoder per input variable, allowing the model to best exploit each variable and timestep with fewer parameters;
- one 2D decoder per output variable, allowing us to remove decoder arms during training (e.g. in Stage 3 of our training procedure, described in Section b), and also reducing the number of parameters;
- we replace the deep inception module from Gao et al. (2022) with a Gated Spatio Temporal attention (GSTa) based module for better efficiency;
- we add a spatio-temporal positional encoding module before the temporal translator, which allows the model to learn regional and seasonal phenomena.

2D univariate encoders. The network encodes each variable and timestep separately, using different encoders (f_{θ_i} in Figure 4) and dividing the spatial dimensions by four. The same encoder is applied to each timestep, but different encoders encode each physical variables, enabling the network to learn the different characteristics of each input data source. The encoders first apply a 4×4 convolution, followed by a GroupNorm from Wu and He (2020) and Gaussian Error Linear Unit (GELU) activation function. Then a strided convolution divides spatial dimensions by 2. This module is then repeated to produce the final latent space.

Spatial-temporal embedding. After the encoder, we apply a positional encoding, i.e. a learnable distribution shift of the latent image informing the network about the spatio-temporal coordinates of the patch (see Appendix b). The temporal dimension is then placed in the variable dimension, meaning that the model can subsequently access temporal dependencies through the variable dimension.

3D multivariate temporal translator. The Gated Spatio Temporal Attention (GSTa) module (t_φ in 4) is composed of a sequence of gated attention sub-blocks. As the the temporal dimension is merged with variables, GSTa

is able to learn the dynamics of the timeseries by multi-variables convolutions. A block of the GSTa model consists of a gated attention module, followed by a Convolutional Feed Forward Network (CFFN) from Wang et al. (2022), used with residual connections and learned rescaling, schematised by t_φ in Figure 4. This module contains dilated convolutions and an attention gating processes, allowing the model to learn long range spatial interactions. For further details about the GSTa module, we refer the reader to Yu et al. (2022).

2D univariate decoder. After the multivariate temporal translator, we apply three 2D univariate decoders in order to decode SSH, the zonal current component U , and the meridional current component V . Similarly to the encoders, the decoders g_{φ_i} apply the same operation to each timestep, but have different weights for each variable. They are composed of 4×4 convolutions followed by bilinear upsampling layers. These operations are repeated twice to obtain the original image dimensions.

b. Multi-stage training

Training a model that leverages the diverse range of oceanographic observations — such as drifters, satellite altimetry, and satellite imagery — is challenging due to the differing temporal coverage and resolutions of these data sources. One might consider developing a model to forecast ocean currents and sea surface height (SSH) by directly comparing predictions to all available measurements, using drifter velocities as targets for currents and combining observations from nadir and SWOT altimetry for SSH forecasting. However, this approach presents several significant challenges. First, observations from different sources are not always consistent; for example, inconsistencies have been identified in some SWOT and nadir crossover points, as discussed in Appendix b. Second, the temporal coverage varies greatly between data sets: SWOT provides only about a year of data, while nadir altimetry has accumulated observations over several decades (Taburet et al. 2019). Third, the density of observations differs across sources, with drifter data being far sparser than satellite measurements, making it difficult to appropriately weight each data source during model training.

Considering these difficulties, we train our model using a 3-stage training strategy, illustrated in Figure 5. The goal of this strategy is to learn a first approximation of SSH and ocean currents using nadir-pointing SSH and derived geostrophic currents from interpolated SSH fields as ground truth, and to progressively refine the learning using scarcer, but more precise or direct observations (SWOT and drifters). The inputs of the neural network are the same at each stage: nadir SSH and SST. We also include experiments on using additional CHL or SWOT observation as inputs. Our strategy contains 3 stages of training:

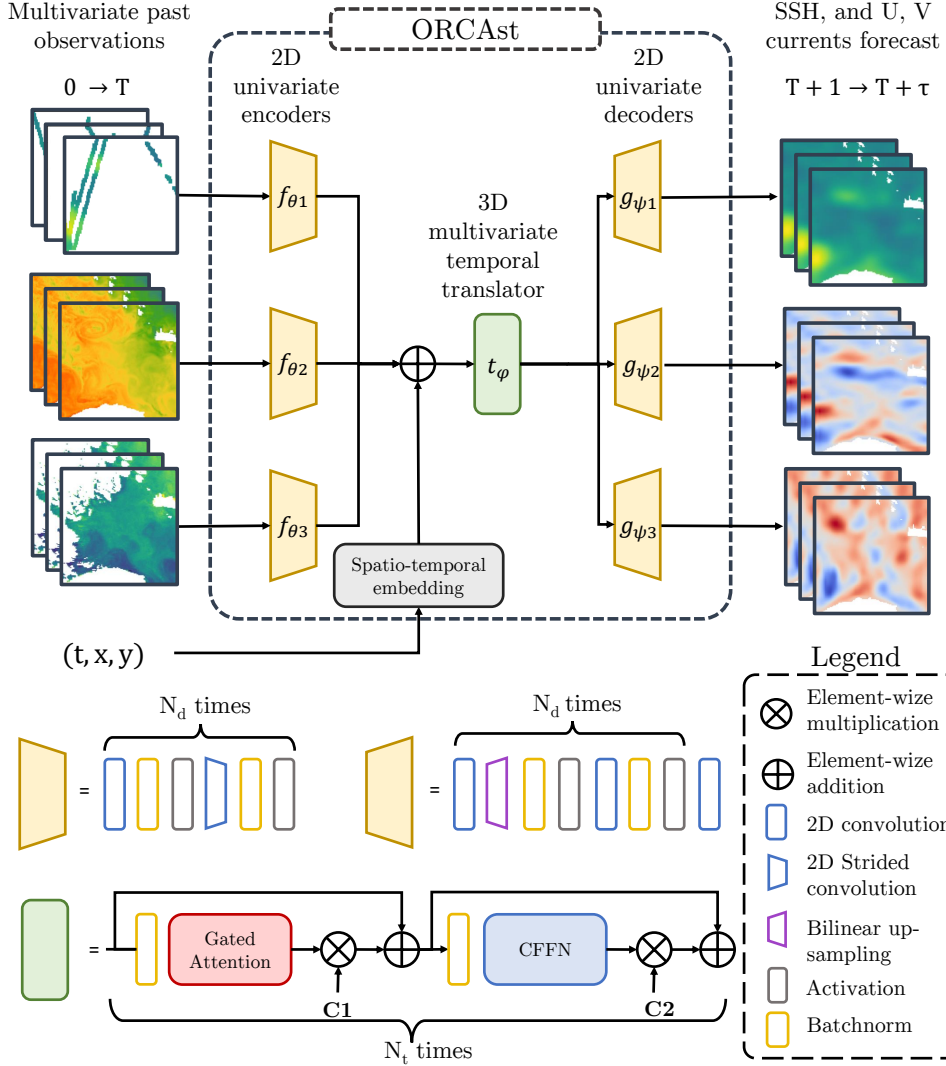


FIG. 4. Overview of ORCAst. Our models inputs temporal sequences of multivariate satellite observations (here SSH, SST, CHL) from $t = 0$ to $t = T$. Each variable is inputted to a different spatial encoder, which extracts information from each timestep separately. A spatio-temporal positional embedding is applied to inform the network of geographical coordinates and seasonal information. A spatio-temporal translator t_φ is used to produce the forecast encodings. Finally, each timestep is decoded independently with three univariate decoders g_ψ , forecasting SSH, U and V components separately.

1. Our model estimates along-track nadir SSH, and DU-ACS delayed-time geostrophic currents. As the nadir pointing observations are sparse, we first perform a masking operation before computing the loss, so that the network can be trained with incomplete fields. For further details about this methodology we refer the reader to Filoche et al. (2022); Archambault et al. (2023); Martin et al. (2023); Archambault et al. (2024b). This corresponds to Stage 1 of Figure 5.
2. Our model estimates high-resolution SWOT data of SSH and associated geostrophic currents. As SWOT produces large bands of observations, it is possible to compute spatial gradients and thus retrieve the geostrophy directly from observations without interpolation methods. However, SWOT provides partial spatial coverage, so we apply the masking operation to both SSH and surface current outputs of ORCAst before computing the loss. This corresponds to Stage 2 of Figure 5.
3. The final stage involves finetuning only the output currents, whilst freezing the SSH decoder. In this stage, the regression of ORCAst currents fits to sparse but direct drifter measurements of U and V , thereby allowing the model to learn from the total currents

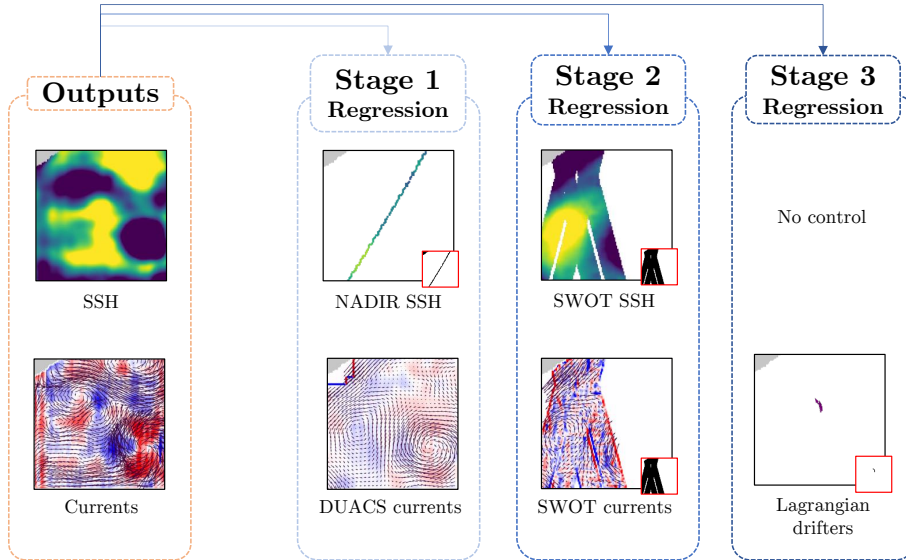


FIG. 5. ORCAst training stages. From left to right, we first show an output example of SSH and U and V components of ocean surface currents, represented as a vector field with colours denoting vorticity. Stage 1 shows nadir and DUACS regression targets, where the small red box illustrates the mask of pixels taken into account when computing the loss. Stage 2 shows SWOT regression targets, with associated masks on the loss function in the small red box. Stage 3 shows drifter regression target, with masks on the loss function in the small red box. Note that in Stage 3 we freeze the SSH decoder and train the model to learn only U and V .

rather than a geostrophical approximation. This corresponds to Stage 3 of Figure 5.

At each stage, we train the model using MSE loss. In order for the model to learn to reconstruct structures and fast-moving ocean currents, we weight the MSE loss function proportionally to the magnitude of ocean currents in the target observations, following Kugusheva et al. (2024). We use an Adam optimizer (Kingma and Ba 2014) with the learning rate set to 10^{-3} and weight decay to 10^{-3} in Stage 1, and resume training by reducing the learning rate by a factor of 10 in Stages 2 and 3. We set each epoch to 1000 randomly chosen patches. During the three stages, the model is trained using 1000, 200 and 50 epochs respectively. When training in the smaller area of the Mediterranean Sea (described in Section c), we train for 500, 200 and 25 epochs respectively, over the 3 stages. The number of epochs was chosen via a validation set. We merge patches using Gaussian kernel-weighted averaging, as detailed in Appendix A of Callaham et al. (2019).

When training the model on a single A100 GPU, running 100 epochs with 1000 samples each requires about 2 hours. The first stage of training therefore lasts roughly a day, the second stage a few hours, and the third stage about an hour. At inference time on the same GPU, generating current forecasts over the entire global area with a spatial stride of 64 pixels and a patch size of 128×128 pixels typically completes in around 5 minutes.

Stages 1 and 2 of our analysis are trained using geostrophic currents, which are the dominant components

TABLE 1. Geographical definitions of oceanic regions by latitude and longitude.

Region	Latitude Range	Longitude Range
Global (extratropical)	-60° to -20° and 20° to 60°	-180° to 180°
Mediterranean Sea	30° to 46°	-6° to 36°
Gulf Stream	20° to 45°	-99° to -34°
Agulhas	-55° to -30°	14° to 74°

of ocean surface currents far from the Equator. Tropical latitudes, where the Coriolis effect is weaker and other fast ageostrophic dynamics such as internal gravitation waves (Figure 3) or equatorial Rossby and Kelvin waves become more significant, are excluded from our analysis. See Figure A1 for a further illustration of the importance of latitude.

c. Geographic zones and time periods

The study area is divided into several subregions, defined in Table 1. We exclude equatorial regions as we train using geostrophic currents in Stages 1 and 2, which are only a reasonable approximation of the total currents far from the Equator. In order to compare the performance of regional-specific models and global models, we train our model on three regions: the Mediterranean Sea, the Gulf Stream and the Agulhas region (Table 1).

We train our model using data from 2016 to mid-December 2018 and from mid-January 2020-2022 during Stages 1 and 3. In Stage 2, we use 6 months of available

SWOT data from 2024, until July. We use 2023 to evaluate our model’s performance.

d. Evaluation metrics and baselines

We evaluate our results on sparse in situ observations collected from Lagrangian drifters in the year 2023 (see Figure 2). As we are interested in detecting eddies and other structures with fast-moving currents, we restrict our drifter evaluation set to observations above 0.25 m/s. In the following, let \hat{w} be the velocity vector predicted by the model at the position of the drifter, and $w_{drifter}$ be the velocity vector observed by the drifter. We measure prediction quality using 3 metrics, illustrated in Figure 6:

- **Percentage of correct angles** Let θ be the angle difference between vectors \hat{w} and $w_{drifter}$. Then:

$$\theta = \frac{180}{\pi} \cos^{-1} \left(\frac{\hat{w} \cdot w_{drifter}}{\|\hat{w}\| \|w_{drifter}\|} \right) \in [0; 180].$$

If $\theta \leq 45^\circ$, we consider the current **direction** as **correct**. We then compute the percentage of correctly predicted angles in our drifter evaluation set.

- **Percentage of correct magnitude**

Let ΔM be the magnitude difference between \hat{w} and $w_{drifter}$. Then:

$$\Delta M = \|\|\hat{w}\| - \|w_{drifter}\|\| \in \mathbb{R}^+.$$

If $\Delta M \leq 2.5$ cm/s, we consider the current **magnitude** as **correct**. We then compute the percentage of correctly predicted magnitudes in our drifter evaluation set.

- **Mean Error Vector Amplitude (MEVA)**

The Error Vector Amplitude is the velocity difference ΔV between \hat{w} and $w_{drifter}$, i.e.

$$\Delta V = \|\hat{w} - w_{drifter}\| \in \mathbb{R}^+.$$

We compute the mean of the Error Vector Amplitude, or the MEVA, across drifter observation in our evaluation set.

These three metrics for our predicted ocean surface current forecasts are calculated at time $T+1$ and $T+7$ days. Our model performance is compared to three alternative methods: DUACS, NeurOST and Mercator. As DUACS and NeurOST do not provide a forecasting method, we compare with persistence forecasting, i.e. compare our model’s forecasts at time $T+7$ with DUACS and NeurOST predictions at time $T+1$. Note that our method is operational, we use only data up until time T to predict forecasts of $T+1$ until $T+7$. We are thus at a disadvantage

when comparing to 6-day delayed-time ocean surface currents from DUACS or 14-day delayed-time ocean surface currents from NeurOST. Mercator (Dréville et al. 2008; Mercator Océan International 2024) provides operational near real-time forecasts of ocean surface currents, and we compare our ocean surface current forecasts to Mercator forecasts.

4. Results

In this section, we provide quantitative and qualitative results of our ocean current forecasts model, using independent in situ observations from drifters from 2023. Section a presents our results on the extratropical Global area and the importance of our proposed three-stage training procedure. Section b details model performance when trained at a global scale or trained specifically on regional areas. In Section c, we conduct ablation studies experimenting with CHL and SST inputs, as well as SWOT data as inputs. Section d illustrates the qualitative performance of our model, compared to baselines. Finally, Section e evaluates our model on real ship data, demonstrating a potential application of forecasting ocean currents in the maritime industry.

a. Global forecast

Table 2 presents the performance of our model, trained and evaluated over the extratropical Global area, using the metrics and baselines outlined in Section d. Our method consistently outperforms the baselines in both next-day ($T+1$) and 7-day ($T+7$) forecasts. For example, at $T+1$, our model achieves 85% of correct angles, compared to 78% for DUACS, 83% for NeurOST, and 70% for Mercator. Similarly, for $T+7$ forecasts, our model reaches 70% of correct angles, outperforming DUACS (62%), NeurOST (65%), and Mercator (56%). These results highlight that despite being an operational, near-real-time forecasting model, ORCAst surpasses the delayed-time methods DUACS (6-day delay) and NeurOST (14-day delay) while also significantly improving over the near-real-time numerical model Mercator. This demonstrates the ability of our deep learning-based approach to effectively capture physical relationships from observational data.

The contribution of the three training stages to the performance of ORCAst is shown in Table 3, with each stage progressively enhancing the model’s accuracy. During Stage 2, incorporating SWOT data allows the model to leverage high-resolution SSH observations, improving the percentage of correct angle predictions from 79% in Stage 1 to 83% in Stage 2. Fine-tuning in Stage 3 using in situ drifter measurements further refines the model’s ability to capture ocean surface currents beyond the geostrophic component. Stage 3 boosts the percentage of correct angles to 85%. Similarly, the 7-day ($T+7$) forecast accuracy improves steadily, with the percentage of correct angles




FIG. 6. Simplified illustration of our three metrics, from left to right: angle error, magnitude error and error vector amplitude when compared to drifter measurements.

TABLE 2. *Evaluation on extratropical Global area.* Our near-real-time model outperforms existing methods for forecasting ocean surface currents, including delayed-time (DT) methods evaluated using persistence (pers.) forecasts. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
DUACS (DT pers.)	78	62	69	68	25	37
NeurOST (DT pers.)	83	65	72	67	25	31
Mercator	70	56	68	68	25	37
ORCAst	85	70	77	69	24	30

TABLE 3. *Importance of 3-stage training.* Each training stage brings additional improvement in the considered metrics. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
ORCAst stage 1	79	64	59	57	30	50
ORCAst stage 2	83	68	73	68	27	34
ORCAst stage 3	85	70	77	69	24	30

increasing from 64% in Stage 1 to 68% in Stage 2, and reaching 70% in Stage 3. The model also experiences a substantial improvement in magnitude performance across the stages. For next-day ($T+1$) forecasts, the percentage of correct magnitudes increases significantly from 59% in Stage 1 to 73% in Stage 2, and finally to 77% in Stage 3. These results emphasize the critical role of each stage in enabling the model to effectively utilize diverse observational data for more accurate ocean current predictions.

The importance of Stage 2 in learning fine-scale structures is shown in Figure 7. During Stage 1, the model is trained to predict large-scale ocean structures using nadir altimetry data, and is not penalised for failing to detect structures below the effective resolution of nadir altimetry (~ 100 km). In Stage 2, the model is further refined to predict smaller-scale structures, visible in the SST and CHL signatures, by training on high-resolution SWOT altimetry data with a lower effective resolution (~ 15 km). We refer

to Appendix b for an ablation study on using NeurOST instead of DUACS in Stage 1.

b. Regional forecasts

This section explores the differences between training models on a large, global scale and using region-specific models. Results in Tables 4, 5, and 6 show that regional models consistently outperform the global model in the Mediterranean Sea, Gulf Stream region, and Agulhas region. Regional ORCAst models benefit from focusing on the unique physical characteristics of their specific areas, while the global ORCAst model must generalize across a wide variety of ocean dynamics. Although our global model incorporates geographical information through positional encoding, this alone is insufficient to fully resolve the nuanced differences across regions at the same capacity, as the model is still required to generalise across a much wider range of physical and environmental conditions, increasing the difficulty of the learning task.

For instance, the Mediterranean Sea is an enclosed basin with lower energy levels and distinct physical characteristics compared to the open ocean. The Mediterranean Sea also features smaller-scale structures and a reduced Rossby radius, which partly explains why the regional Mediterranean model achieves a significant improvement over the globally-trained model. As shown in Table 4, the regional model achieves 85% correct angle predictions at time $T+1$, compared to 77% for the global model, and 88% of correct magnitudes at time $T+1$, compared to 78% for the global model. For 7-day lead times, improvements are even more significant. At $T+7$, the Mediterranean model achieves 70% correct angles, compared to 58% for the global model, and 83% correct magnitudes, compared to 59% for the global model. This substantial improvement in both angle and magnitude accuracy underscores the advantage of focusing on the region’s unique dynamics.

In contrast, the Gulf Stream and Agulhas regions are highly energetic regions, presenting a unique challenge for forecasting. The regional models outperform the global model, particularly at 7-day lead times ($T+7$), showcasing the ability of regional models to capture local dynamics (Tables 5 and 6). In the Gulf Stream region, the regional model achieves 88% correct angle predictions, outperforming the global model at 85% and surpassing the perfor-

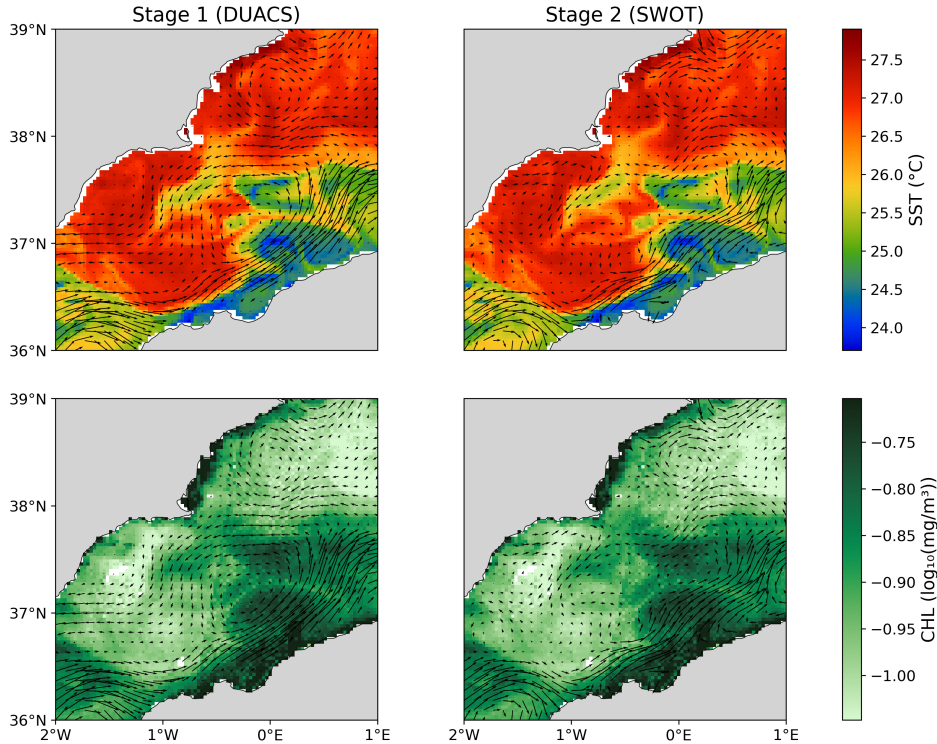


Fig. 7. Example of performance of ORCAst in the Mediterranean Sea on August 31, 2023. SST (top line) and CHL (bottom lines) are shown in the background. Black arrows denoted predicted ocean surface currents after Stage 1 (left column) and Stage 2 (right column) of the regional model trained in the Mediterranean Sea (see Table 4). In Stage 1, the model is trained to predict large-scale structures using nadir altimetry data. In Stage 2, the model is trained with fine-scale SWOT altimetry data, which allows it to be better rewarded for accurately predicting small-scale structures seen in the SST signature.

TABLE 4. *Evaluation in Mediterranean Sea.* Our model trained specifically on the Mediterranean Sea outperforms ORCAst trained on the extratropical Global area, and also achieves higher accuracy than alternative methods. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
DUACS (DT pers.)	58	55	51	51	29	31
NeurOST (DT pers.)	74	63	65	62	25	28
Mercator	51	43	64	60	32	35
ORCAst Global	77	58	78	59	26	30
ORCAst Regional	85	70	88	83	21	25

mance of delayed-time methods such as DUACS (83%) and NeurOST (86%), as well as near-real-time Mercator (77%) (Table 5). Furthermore, at $T + 7$, the regional model predicts angles with 73% accuracy, compared to 66% for the global model. Similarly, in the Agulhas region, the regional model achieves an angle accuracy of 93% at time $T + 1$, compared to 92% for the global model (Table 6). At

TABLE 5. *Evaluation in the Gulf Stream region.* Our model trained specifically on the Gulf Stream region outperforms ORCAst trained on the extratropical Global area, and also achieves higher accuracy than alternative methods. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
DUACS (DT pers.)	83	61	71	65	29	42
NeurOST (DT pers.)	86	63	72	65	25	37
Mercator	66	52	65	64	38	41
ORCAst Global	85	66	74	66	27	34
ORCAst Regional	88	73	77	70	24	31

$T + 7$, the regional model reaches 78% accuracy in angle prediction, surpassing the global model's 73% and indicating its superior ability to capture the region's ocean dynamics over extended time frames. These results emphasise the advantage of training on region-specific data for accurate, long-term forecasts in highly energetic ocean

TABLE 6. *Evaluation in the Agulhas region.* Our model trained specifically on the Agulhas region outperforms ORCAst trained on the extratropical Global area, and also achieves higher accuracy than alternative methods. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
DUACS (DT pers.)	87	67	71	67	29	35
NeurOST (DT pers.)	91	70	75	67	26	35
Mercator	77	62	64	66	37	38
ORCAst Global	92	73	77	70	25	33
ORCAst Regional	93	78	80	73	23	30

TABLE 7. Comparison of ORCAst performances on the extratropical Global area, using SST or SST + CHL as inputs, in addition to nadir SSH. Including CHL data does not bring important improvements to ORCAst currents. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
ORCAst	85	70	77	69	24	30
ORCAst +CHL	85	70	76	72	27	29

regions. We refer to Appendix b for metrics on Eddy Kinetic Energy (EKE) in these regions.

Figure 8 illustrates how ORCAst learns to accurately predict the evolution of a dynamic structure in the Agulhas region. The Mercator model incorrectly localises the structure, but predicts a westward evolution. In contrast, ORCAst not only correctly identifies the location of the structure, but also accurately forecasts its position in 7-days. This ability to both localise and predict the evolution of such features highlights the strength of ORCAst in learning complex ocean dynamics.

c. Including more input data

The flexibility of ORCAst allows us to include various inputs. In this section, we present experiments where we evaluate ORCAst using CHL or SWOT data as inputs, in addition to SST and nadir altimetry. Table 7 compares model performance using CHL as additional inputs, and Table 8 compares model performance using SWOT as additional inputs.

Table 7 suggests that there is limited potential of CHL to provide complementary high-resolution information on ocean currents, when used in combination with high-resolution SST data. Scores for the model both with and without CHL are similar.

TABLE 8. Comparison of ORCAst performances on the extratropical Global area, using nadir altimetry or SWOT and nadir altimetry as inputs, in addition to SST. Including SWOT data as inputs does not bring improvements to ORCAst currents. This model was trained as usual during Stage 1, then trained with SWOT inputs during Stage 2. It is evaluated on August-December 2023, rather than the whole of 2023, due to availability of SWOT data, thus the first line of results is not identical to the line corresponding to Stage 2 in Table 3.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
ORCAst	83	68	73	69	25	30
ORCAst +SWOT	82	67	72	68	26	31

We compare ORCAst models using nadir-pointing altimetry only or nadir and SWOT as SSH inputs, in addition to SST. As SWOT data is only available from mid-2023 onwards, we train our model using SWOT inputs during Stage 2 and evaluate our model using August-December 2023. We keep the training procedure of Stage 1 as before, using only SST and nadir altimetry as inputs. Results of this model at the end of Stage 2 are provided in Table 8. Using SWOT data as inputs seemingly does not help the model to better predict ocean surface currents. However, since there is only around one year of available SWOT data, the training and test sets are small, making it difficult to draw strong conclusions from this experiment. SWOT observations reveal unobserved mesoscale and sub-mesoscale structure, however the long revisit time and limited available data from the mission may partially explain the lack of performance gains when used as additional inputs.

d. Qualitative evaluation

An example of ORCAst predicted currents at time $T + 1$ in the Mediterranean Sea is shown in Figure 9, alongside comparisons with delayed-time models from DUACS and NeurOST, as well as ocean surface currents from Mercator. SST and CHL images are shown in the background. Although DUACS and NeurOST tend to correctly localise large-scale structures, some smaller-scale structures may be missed. In the case of DUACS, this is because the method lacks high-resolution information from SST. In the case of NeurOST, this is possibly because the model learns from nadir altimetry, with a low effective resolution, and is thus not sufficiently penalised for not learning fine-scale information visible in the SST. Figure 7 illustrates the difference between training on nadir altimetry and high-resolution SWOT altimetry.

A qualitative comparison of our model outputs to drifter observations at days $T + 1$ and $T + 7$ is shown in Figure 10. Like DUACS and NeurOST, ORCAst correctly localises the eddy at $T + 1$. However, persistence forecasts for DUACS and NeurOST do not align well with the drifter’s




Fig. 8. Predicted evolution of a structure in the Agulhas region between October 11 and October 17, 2023. SST is shown in the background, with black arrows denoting the ocean surface currents. The Mercator model misidentifies the structure’s location but predicts an eastward movement. In contrast, ORCAst both accurately locates the structure and provides a more precise 7-day forecast of its future position.

motion at time $T + 7$. In contrast, ORCAst correctly predicts the drifter’s trajectory at time $T + 7$. Mercator fails to localise the eddy at $T + 1$, as well as its evolution at time $T + 7$.

e. Evaluation using ship data

Ships generally measure their Speed Through Water (STW) and Speed Over Ground (SOG), and the difference between these measurements, $\text{SOG} - \text{STW}$, serves as an indicator of the currents experienced by ships along their trajectory.

In Figure 11, we evaluate our model using ship data from a trans-Mediterranean voyage between August 21 and August 26, 2023. We compute $\text{SOG} - \text{STW}$ and compare these observations to forecasted ocean currents along the ship’s trajectory, as provided by the ORCAst and Mercator models. Our findings indicate that ORCAst forecasts align more closely with the observed $\text{SOG} - \text{STW}$ values than Mercator forecasts, suggesting superior performance. Correct estimations of the currents felt along ship routes is important for maritime route optimisation, allowing ships to take routes following more favourable ocean currents and avoiding currents against the path of the ship.

TABLE 9. Correlation and R^2 Scores between the SOG-STW and the current impact measured by DUACS, NeurOST, Mercator and ORCAst for the trans-Mediterranean voyage illustrated in Figure 11. For DUACS and NeurOST, we consider delayed-time current estimations for August 21, 2023. For Mercator and ORCAst, we consider ocean current forecasts for August 21-26, 2023, made using data available at or before August 20, 2023.

	Pearson Correlation (r)	R^2 Score
DUACS (DT pers.)	0.73	0.50
NeurOST (DT pers.)	0.74	0.53
Mercator	0.31	-0.01
ORCAst	0.79	0.54

Moreover, for this trans-Mediterranean voyage, we compute the Pearson correlation coefficient and the R^2 score between the current impact predicted by DUACS, NeurOST, Mercator and ORCAst. We use delayed-time predictions for DUACS and NeurOST, corresponding to August 21, 2023. We use operational forecasts for Mercator and ORCAst, using only data available at or before August 20, 2023. As the Mediterranean crossing takes 5 days, we use 6 days of forecasts. Table 9 shows that ORCAst performs best, even when compared to delayed-time methods.

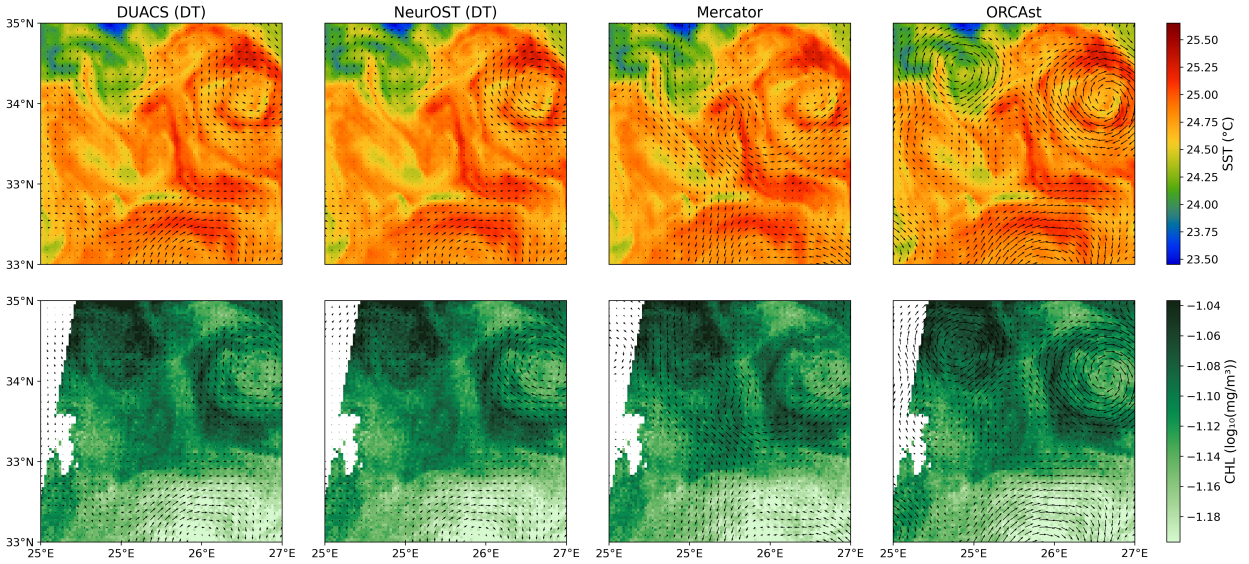


Fig. 9. Model performance in the Mediterranean Sea on June 28, 2023. The top row shows the corresponding SST images in the background, and the bottom row shows the corresponding CHL images in the background. The field of ocean surface currents for each model is represented using black arrows. ORCAst correctly localises both large and fine-scale structures, visible in the SST and CHL signatures. We note that for this example, we illustrate a higher-resolution DUACS product, only available for the Mediterranean Sea (Copernicus Marine Service 2024).

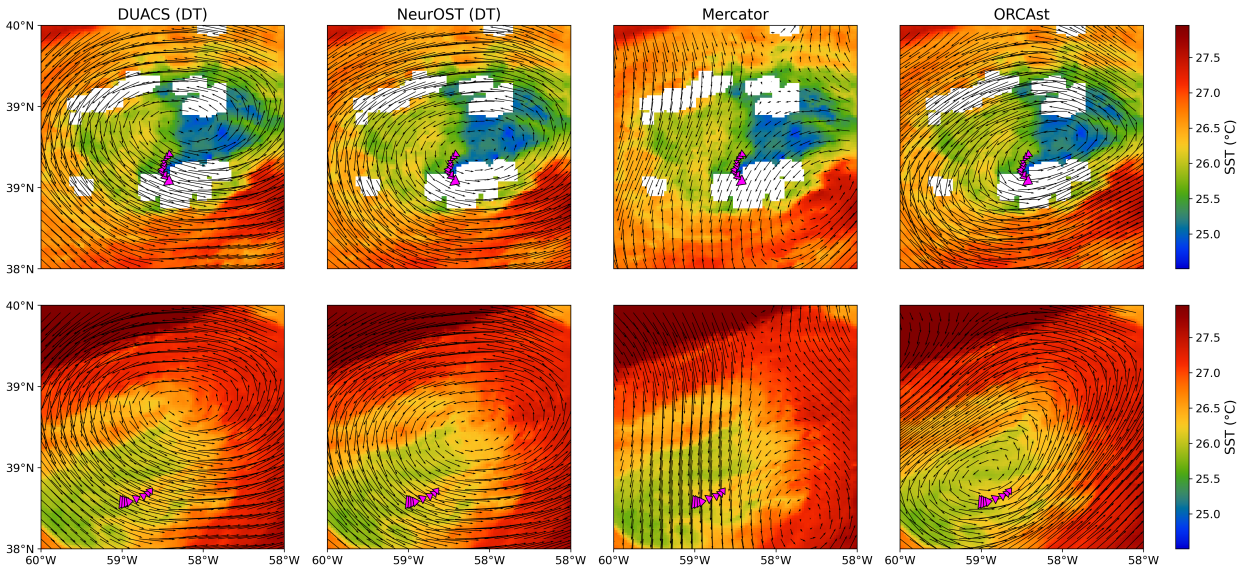


Fig. 10. Drifter (in magenta) with WMO identification number 4402878 on August 12 and August 18, 2023. The image backgrounds show the corresponding SST, with the field of predicted ocean surface currents represented using black arrows. We compare the model performance to in situ observations from drifters at different lead times. ORCAst correctly localises the eddy and predicts its evolution over 7 days.

5. Conclusion

a. Discussion

In this study, we introduce ORCAst, a deep learning framework for forecasting ocean surface currents. ORCAst produces $1/30^\circ$ high-resolution 7-day forecasts of ocean currents in Global extratropical regions. Our flexi-

ble approach allows us to fuse multivariate observations of the ocean such as in situ measurements of currents, satellite altimetry, and satellite imaging of sea surface temperature and chlorophyll-a, to improve the spatial resolution and prediction of the temporal evolution of surface currents.

For next-day predictions, we accurately predict the angle of drifters 85% of the time, surpassing delayed-time DU-

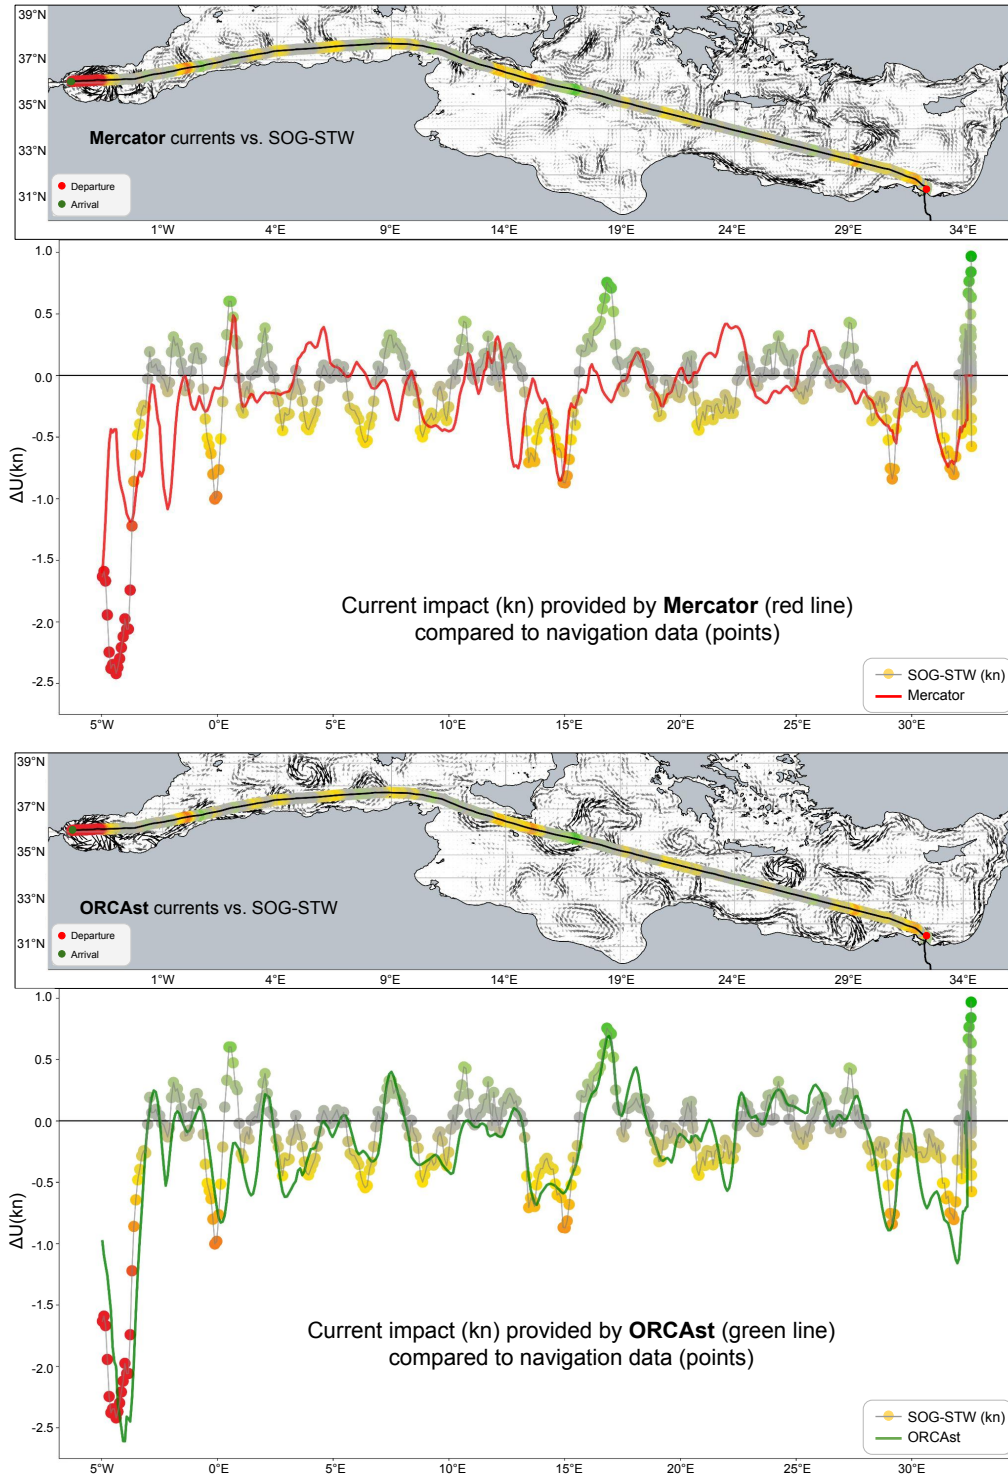


FIG. 11. Evaluation of Mercator and ORCAst on ship data for a trans-Mediterranean voyage between August 21 and August 26, 2023. Maps show ocean currents estimated by Mercator and ORCAst, with the ship route coloured proportionally to SOG-STW measurements. The graphs show SOG-STW measurements, compared with the current impact predicted by Mercator and ORCAst. The SOG-STW curve is coloured proportionally to SOG-STW, in the same way as in the map. Currents predicted by ORCAst more closely follow the SOG-STW measurements than currents predicted by Mercator. See Table 9 for a quantitative evaluation of this voyage.

ACS (78%), delayed-time NeurOST (83%) and Mercator (70%). For 7-day forecasts, we achieve 70% correct angle predictions, surpassing persistence forecasts of DUACS (62%) and NeurOST (65%), and numerical forecasts of the operational OGCM Mercator (56%). We also achieve superior performance on our metrics of correct magnitude and mean MEVA. The superior performance of ORCAst over delayed-time (DUACS and NeurOST) and near real-time (Mercator) methods highlights the potential of deep learning in leveraging large quantities of observational data to produce accurate operational ocean current forecasts.

We show that progressively training the network on sparser, but more direct, observations of currents, leads to improved performances at each training stage. We train using nadir-pointing altimetry and DUACS geostrophic currents, followed by SWOT SSH and SWOT geostrophic currents, and finally with drifter current measurements. In particular, training on drifters allows the model to learn from observations beyond geostrophic ocean surface currents.

Regional models consistently outperform the globally trained model. The Mediterranean Sea, with its unique dynamics and fine spatial structures, showed the largest improvement when using a regional model. Our regional models also outperform the Global model in energetic regions such as the Gulf Stream and Agulhas regions, better capturing localised physical dynamics over longer lead times.

The inclusion of chlorophyll-a data did not yield substantial improvements when SST data was already used, suggesting limited complementary information from chlorophyll-a for forecasting surface currents. Similarly, incorporating SWOT data as inputs did not significantly enhance model performance, likely due to the limited availability of SWOT observations (one year of data). The addition of SWOT measures over the coming years should further enhance ORCAst’s learning and performance. However, as there is little overlap between the SWOT measurements in 2023 and the drifter trajectories used for validation we cannot demonstrate yet the effectiveness of using SWOT data as inputs. We note that training our model using SWOT data as targets led to significant improvements for surface currents at meso and sub-mesoscale.

The qualitative assessments reinforce the quantitative metrics, demonstrating that ORCAst can provide detailed forecasts that align closely with observed patterns, outperforming baseline methods. Moreover, the evaluation against real ship data highlights the practical applicability of the model. Accurate current predictions are essential for optimizing ship routing and reducing emissions, making ORCAst a valuable tool for the maritime industry.

b. Perspectives

Throughout this study, we demonstrate the potential of fusing multivariate observations with deep learning. However, a number of issues still need to be addressed to improve the quality of our forecasted ocean currents.

First, ORCAst produces no forecast between latitudes -20° and 20° . This is because the geostrophic approximation is not valid near the Equator, and the link between currents and altimetry weakens. One alternative approach would be to train a deep learning model using an assimilated numerical model as a target on the Equator.

Second, to further develop the fusion of observations from various sources, we are interested in using ship data. For instance, the Automatic Identification System (AIS), provides information about the ship’s steering and speed over ground, which can help us deduce the surface water velocity. However, using AIS data to obtain indirect current measurements is difficult, and calls for additional studies. Other sources of on-board observations exist, such as Acoustic Doppler Current Profiler (ADCP) observations or Wavex sensor data. These more direct measurement of currents along the ship’s route can also be used as targets for evaluation or training, however they may also be noisy, often measuring incoherent current velocities.

Finally, from a methodological point of view, we are interested in going beyond regression models (like ORCAst). Regression models estimate the mean of a distribution, their predictions are therefore smoothed, leading to a domain gap compared to real physical samples, with smoothing getting stronger when conditions are scarce or irrelevant. On the other hand, generative deep learning allows the network to generate multiple examples from the same set of inputs, hopefully forming a physics-matching distribution. In particular, Denoising Diffusion Probabilistic Models (DDPM) and flow matching techniques have led to an impressive image generation breakthrough in the past years. Using these sampling models, we might be able to estimate the full conditional distribution of the currents, instead of only the mean forecast.

Acknowledgments. This research study was supported by a Cifre thesis grant from the ANRT 2023/1590, and partially financed by the iLab2023 grant DOS0220247/00 of the French Public Bank of Investment (BPIFrance). This work was granted access to the HPC resources of IDRIS under the allocation AD011015571 made by GENCI. The authors of this paper have affiliations with Amphitrite, a company developing maritime solutions such as optimal ship routing, using methods similar to those discussed in this study. As such, there exists a potential conflict of interest regarding the presentation and interpretation of the results. The authors declare that the study was conducted with rigor and integrity, and the interpretation of the results is based on the evidence presented in the data. Large

Language Models (LLMs) were used to assist with coding and writing, however the authors have made every effort to ensure the content meets ethical and academic standards.

Data availability statement. The original satellite data and in situ data from drifters presented in the study are openly available. We refer to the Data section of this article for citations of all of our data sources. We unfortunately are unable to diffuse the ship data used in our analysis, however this data is not necessary for the reproducing of our model and key results.

APPENDIX A

Learned Positional Embeddings

Learned positional embeddings provide insights into the spatio-temporal conditions most relevant for ORCAst to forecast ocean currents. These embeddings are represented as 32-dimensional vectors. To visualise patterns within the embeddings, we create a 3D latitude-longitude-time grid encompassing all spatio-temporal positions used during inference. To efficiently cluster points in this large grid, we apply the mini-batch k -means clustering algorithm with $k = 30$ clusters and a batch size of 8. Figure A1 plots these $k = 30$ clusters on a latitude-longitude grid for a chosen temporal dimension (first week of the year). Latitude is a strong determining factor of cluster membership, demonstrating that the positional embedding learns the strong influence of latitude in predicting ocean currents.

Furthermore, Figure A2 shows the correlation and the distance between clusters. We see a clear separation between Southern Hemisphere (clusters 0 to 18) and Northern Hemisphere (clusters 19-29) and Southern Hemisphere embeddings. The Southern Hemisphere is more represented in the clustering (19 clusters versus 11 clusters), because the ocean-to-continent ratio is higher in the South, which means that the model sees more South examples, and adapts its encoding accordingly. Generally, Southern Hemisphere clusters are correlated with each other, and negatively correlated with Northern Hemisphere clusters. Southern Hemisphere clusters are closest to other Southern Hemisphere clusters, and Northern Hemisphere clusters are closest to other Northern Hemisphere clusters.

This is expected, as the Coriolis effect varies with latitude, reaching its maximum at the poles. In the Northern Hemisphere, the westerlies drive currents eastward, while the Coriolis effect deflects them to the right, forming clockwise gyres. In the Southern Hemisphere, a similar mechanism creates counterclockwise gyres due to the Coriolis effect deflecting currents to the left. The variability in the Coriolis effect modulates the intensity and structure of ocean currents in both hemispheres.

Although we denormalise our dataset at a spatial resolution of 2 degrees and at a temporal resolution of 1 week,

our positional encoding of the day of the year theoretically allows the model to learn seasonal differences which may occur at a finer scale, such as the speed of the evolution of the eddies over the 7-day forecast, or the size of the eddies. Figure A3 plots $k = 30$ clusters at longitudes 80°W to 0°E at dates across 4 seasons. There is minor seasonal variation. In the Northern Hemisphere, some clusters corresponding to higher latitudes in July are the same as clusters corresponding to lower latitudes in January. In the Southern Hemisphere, some clusters corresponding to higher latitudes in July are the same as clusters corresponding to lower latitudes in January.

APPENDIX B

Bias of SWOT

We calculate the bias between Absolute Dynamic Topography (ADT) measured by SWOT and by nadir altimetry at the same location and time, as well as the standard deviation of their differences (Figure B1). We find that SWOT’s KaRIn instrument overestimates SSH compared to nadir altimetry by 5.26 ± 3.32 cm. This bias, which likely arises from differences in the Mean Dynamic Topography (MDT), could complicate training, as our model uses nadir SSH in Stage 1 and SWOT SSH in Stage 2. While correcting for this offset could help a model aiming to predict ADT, we do not apply such a correction here, as our focus is on predicting ocean surface currents, which should not be significantly affected by an MDT-related bias. Instead, to address potential distributional shifts between data sources, we adopt a three-stage training approach, using a different ground truth dataset at each stage. By first training on nadir and then transitioning to SWOT, the model progressively learns to account for such biases and generalise across observation types.

APPENDIX C

Ablation Study

NeurOST performs better than DUACS, as shown in Table 2, however we train using DUACS as targets in Stage 1 of our learning procedure. As shown in Table C1, replacing DUACS with NeurOST improves model performance during Stage 1 of training. However, the difference in performance difference diminishes in Stages 2 and 3.

APPENDIX D

Eddy Kinetic Energy

Eddy Kinetic Energy (EKE) is commonly used to quantify the intensity of mesoscale variability, such as ocean

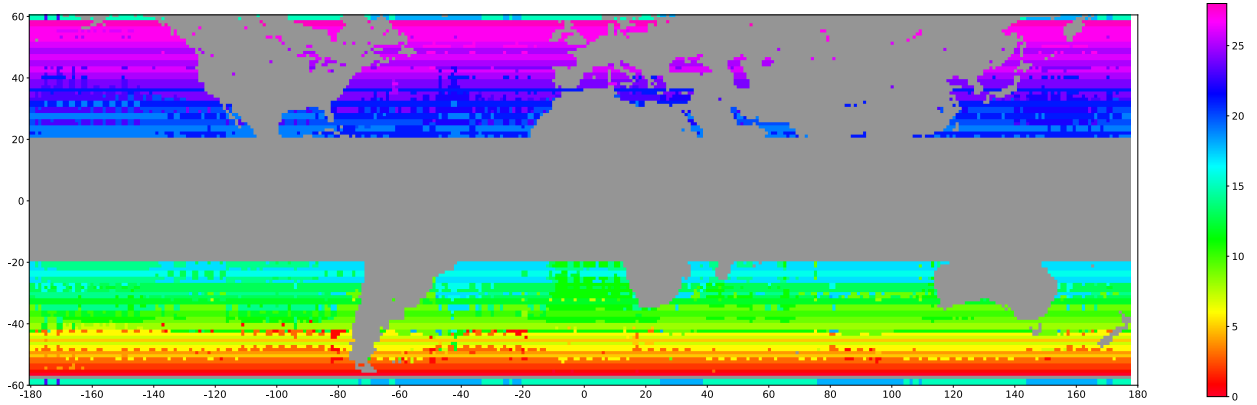


FIG. A1. Clustering of positional embeddings using k -means, where $k = 30$. Each colour represents a cluster learned by our model. We reorder the numbering and colours of the groups for clearer visualisation purposes. Land areas and the equator are in grey.




FIG. A2. The correlation (left) and the distance (right) between clusters of our positional embedding, chosen using the k -means method with $k = 30$. Geographic zones corresponding to cluster numbers are shown in Figure A1. Southern Hemisphere clusters are numbered 0-18 and Northern Hemisphere clusters are numbered 19-29.

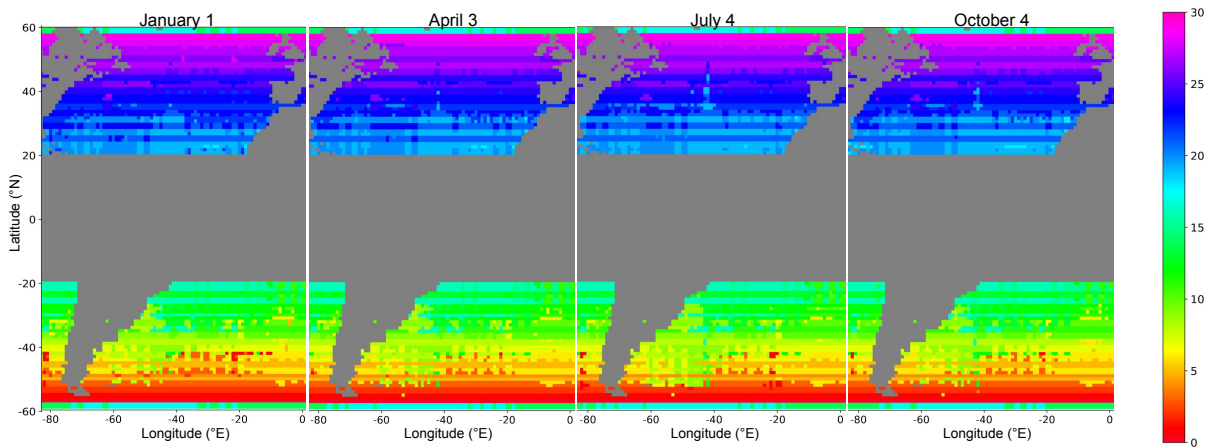
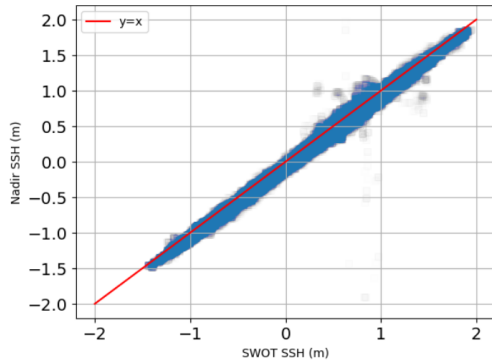
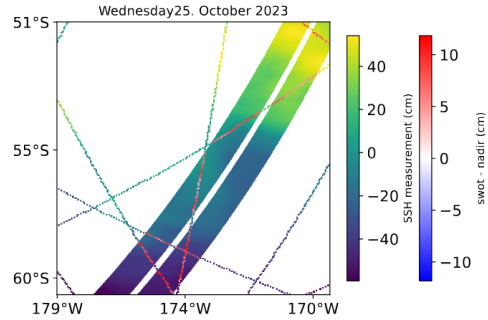


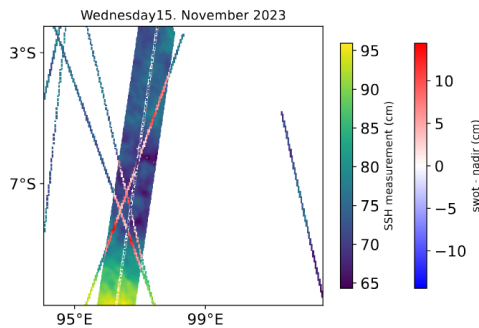
FIG. A3. Clustering of positional embeddings for longitudes 80°W to 0°E at 4 different seasons using k -means, where $k = 30$. Each colour represents a cluster learned by our model. We reorder the numbering and colours of the groups for clearer visualisation purposes. Land areas and the equator are in grey.




(a) KaRIn vs nadir SSH at overlapping measurements (on the same day at the same location)



(b) KaRIn-nadir crossings south-east of New-Zealand



(c) KaRIn-nadir crossings south-west of Singapore



(d) KaRIn-nadir crossings near the Gulf Stream

FIG. B1. We compare KaRIn and nadir altimetry across intersection points of both measures of SSH. We find that there is an underestimation of SSH from nadir compared to SSH measured by SWOT's KaRIn instrument, of around 5.26 ± 3.32 cm.

TABLE C1. Comparison of model performance using DUACS and NeurOST as targets during Stage 1 (S1) of training in the Gulf Stream region. Replacing DUACS with NeurOST improves performance in Stage 1 of training. However, the performance difference diminishes in Stages 2 and 3 (S2 and S3), where the models achieve similar results. See Section d for a description of our metrics. Values are rounded. We evaluate on drifter observations from 2023 with magnitude > 0.25 m/s. Best results are shown in bold font.

	Correct ang. %		Correct mag. %		MEVA (cm/s)	
	T+1	T+7	T+1	T+7	T+1	T+7
NeurOST (DT pers.)	86	63	72	65	25	37
S1	82	67	67	65	29	33
S1 on NeurOST	87	72	73	68	26	32
S2	86	72	76	72	25	32
S2, NeurOST for S1	87	72	76	71	25	31
S3	88	73	77	70	24	31
S3, NeurOST for S1	89	73	79	71	23	30

eddies and turbulent motions. It represents the kinetic energy associated with deviations from the mean flow and is particularly useful for comparing the dynamical activity of

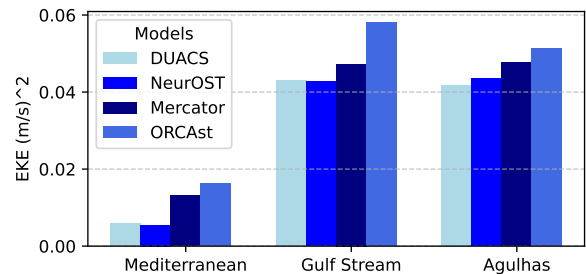


FIG. D1. Eddy Kinetic Energy (EKE) for our three zones of interest.

ocean models or observational datasets. For a given point, EKE is defined as: $EKE = \frac{1}{2} (u'^2 + v'^2)$, where u' and v' are the anomalies of the zonal and meridional velocities, respectively, relative to the mean flow. We compute the EKE using model outputs from 2023, with mean flows from 2022-2024. Figure D1 shows the EKE of DUACS, NeurOST, Mercator and ORCAst in our three regional domains. ORCAst tends to have a higher EKE than other models.

References

- GHRSSST NOAA/STAR ACSPO v2.81 0.02 degree L3S Daily Dataset from LEO Satellites, 2024. URL https://podaac.jpl.nasa.gov/dataset/L3S_LEO_DY-STAR-v2.81.
- A. E. Aouni, Q. Gaudel, C. Regnier, S. Van Gennip, M. Drevillon, Y. Drillet, and J.-M. Lellouche. Glonet: Mercator's end-to-end neural forecasting system. *arXiv preprint arXiv:2412.05454*, 2024.
- T. Archambault, A. Filoche, A. Charantonis, and D. Béréziat. Multi-modal unsupervised spatio-temporal interpolation of satellite ocean altimetry maps. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, Lisboa, Portugal, Feb. 2023.
- T. Archambault, A. Filoche, A. Charantonis, and D. Béréziat. Pre-training and Fine-tuning Attention Based Encoder Decoder Improves Sea Surface Height Multi-variate Inpainting. In *VISAPP 2024 - 19th International Conference on Computer Vision Theory and Applications*, Roma, Italy, Feb. 2024a. URL <https://hal.sorbonne-universite.fr/hal-04475205>.
- T. Archambault, A. Filoche, A. A. Charantonis, D. Béréziat, and S. Thiria. Learning Sea Surface Height Interpolation from Multi-variate Simulated Satellite Observations. *James, Journal of Advancing in Modeling Earth Syst*, 16(6):e2023MS004047, 2024b. doi: 10.1029/2023ms004047. URL <https://hal.sorbonne-universite.fr/hal-04568666>. Publisher: AGU.
- J. L. Callahan, K. Maeda, and S. L. Brunton. Robust flow reconstruction from limited measurements via sparse representation. *Physical Review Fluids*, 4(10):103907, Oct. 2019. doi: 10.1103/PhysRevFluids.4.103907. URL <https://link.aps.org/doi/10.1103/PhysRevFluids.4.103907>. Publisher: American Physical Society.
- L. Chen, X. Zhong, F. Zhang, Y. Cheng, Y. Xu, Y. Qi, and H. Li. FuXi: a cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, 6(1):1–11, Nov. 2023. ISSN 2397-3722. doi: 10.1038/s41612-023-00512-1. URL <https://www.nature.com/articles/s41612-023-00512-1>. Publisher: Nature Publishing Group.
- D. Ciani, C. Fanelli, and B. Buongiorno Nardelli. Estimating ocean currents from the joint reconstruction of absolute dynamic topography and sea surface temperature through deep learning algorithms. *EGU Sphere*, pages 1–25, Apr. 2024. doi: 10.5194/egusphere-2024-1164. URL <https://egusphere.copernicus.org/preprints/2024/egusphere-2024-1164/>. Publisher: Copernicus GmbH.
- Copernicus Climate Service. Global Ocean Gridded L4 Sea Surface Heights And Derived Variables Reprocessed 1993 On-going. Satellite observations, NetCDF-4 format, 11 2024. URL <https://doi.org/10.48670/moi-00148>. Product ID: SEALEVEL_GLO_PHY_L4_MY_008_047.
- Copernicus Marine Service. Global Ocean Along Track L3 Sea Surface Heights NRT. Satellite observations, NetCDF-4 format, 11 2023. URL <https://doi.org/10.48670/moi-00147>. Product ID: SEALEVEL_GLO_PHY_L3_NRT_008_044.
- Copernicus Marine Service. European Seas Gridded L4 Sea Surface Heights And Derived Variables NRT. Satellite observations, NetCDF-4 format, 11 2024. URL <https://doi.org/10.48670/moi-00142>. Product ID: SEALEVEL_EUR_PHY_L4_NRT_008_060.
- M. Drévillon, R. Bourdallé-Badie, C. Derval, Y. Drillet, J.-M. Lellouche, E. Rémy, B. Tranchant, M. Benkiran, E. Greiner, S. Guinehut, N. Verbrugge, G. Garric, C.-E. Testut, M. Laborie, L. Nouel, P. Bahurel, C. Bricaud, L. Crosnier, E. Dombrowsky, E. Durand, N. Ferry, F. Hernandez, O. L. Galloudec, F. Messal, and L. Parent. The GODAE/Mercator-Ocean global ocean forecasting system: results, applications and prospects. *Journal of Operational Oceanography*, 1(1):51–57, Feb. 2008.
- S. Elipot, R. Lumpkin, R. C. Perez, J. M. Lilly, J. J. Early, and A. M. Sykulski. A global surface drifter dataset at hourly resolution. *Journal of Geophysical Research: Oceans*, 121, 2016. doi: 10.1002/2016JC011716.
- S. Elipot, A. Sykulski, R. Lumpkin, L. Centurioni, and M. Pazos. Hourly location, current velocity, and temperature collected from Global Drifter Program drifters world-wide. Dataset, 2022. URL <https://doi.org/10.25921/x46c-3620>. Accessed on July 31, 2023.
- W. Emery, J. Brown, and Z. Nowak. AVHRR image navigation - Summary and review. *Photogrammetric Engineering & Remote Sensing*, 55, Sept. 1989.
- EUMETSAT for Copernicus. Ocean Colour Service: Product Data Format Specification - OLCI Level 1 & Level 2 Instrument Products, 2024. URL <https://data.eumetsat.int/product/EO:EUM:DAT:0407>. Platform: Sentinel-3; Sensor: OLCI (Optical); Data Level: Level 2; Collection ID: EO:EUM:DAT:0407; Parameters: Ocean, Ocean Colour, Level 2 Data.
- R. Fablet, B. Chapron, J. Le Sommer, and F. Sèvelec. Inversion of Sea Surface Currents From Satellite-Derived SST-SSH Synergies With 4DVarNets. *Journal of Advances in Modeling Earth Systems*, 16(6):e2023MS003609, 2024. ISSN 1942-2466. doi: 10.1029/2023MS003609. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2023MS003609>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2023MS003609>.
- A. Filoche, T. Archambault, A. Charantonis, and D. Béréziat. Statistics-free interpolation of ocean observations with deep spatio-temporal prior. In *ECML/PKDD Workshop on Machine Learning for Earth Observation and Prediction (MACLEAN)*, 2022.
- Z. Gao, C. Tan, L. Wu, and S. Z. Li. Simvp: Simpler yet better video prediction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3160–3170, 2022. doi: 10.1109/CVPR52688.2022.00317.
- M. Gurvan, R. Bourdallé-Badie, P.-A. Bouttier, C. Bricaud, D. Bruciferri, D. Calvert, J. Chanut, E. Clementi, A. Coward, D. Delrosso, C. Ethé, S. Flavoni, T. Graham, J. Harle, D. Iovino, D. Lea, C. Lévy, T. Lovato, N. Martin, S. Masson, S. Mocavero, J. Paul, C. Rousset, D. Storkey, A. Storto, and M. Vancoppenolle. NEMO ocean engine, Oct. 2017. URL <https://zenodo.org/records/3248739>. Publisher: Zenodo.
- Ifremer. Global Ocean In-Situ Near-Real-Time Observations. In-situ observations, NetCDF-4 format, 11 2024. URL <https://doi.org/10.48670/moi-00036>. Product ID: INSITU_GLO_PHYBGCWAV_DISCRETE.MYNRT_013_030.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2014.
- A. Kugusheva, H. Bull, E. Moschos, A. Ioannou, B. Le Vu, and A. Stegner. Ocean Satellite Data Fusion for High-Resolution Surface Current Maps. *Remote Sensing*, 16(7):1182, Mar. 2024. ISSN 2072-4292. doi: 10.3390/rs16071182. URL <https://www.mdpi.com/2072-4292/16/7/1182>.

- R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, and P. Battaglia. Learning skillful medium-range global weather forecasting. *Science (New York, N.Y.)*, page eadi2336, Nov. 2023. ISSN 1095-9203. doi: 10.1126/science.adi2336.
- F. Le Guillou, L. Gaultier, M. Ballarotta, S. Metref, C. Ubelmann, E. Cosme, and M.-H. Rio. Regional mapping of energetic short mesoscale ocean dynamics from altimetry: performances from real observations. *Ocean Science*, 19(5):1517–1527, Oct. 2023a. ISSN 1812-0784. doi: 10.5194/os-19-1517-2023. URL <https://os.copernicus.org/articles/19/1517/2023/>. Publisher: Copernicus GmbH.
- F. Le Guillou, B. Maxime, S. Metref, C. Ubelman, E. Cosme, and R. M.-H. The woc data challenges. <https://2024-dc-woc-esa.readthedocs.io/en/latest/index.html>, 2023b. Regional mapping of energetic short mesoscale ocean dynamics from altimetry: performances from real observations.
- S. A. Martin, G. E. Manucharyan, and P. 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*, 15(5):e2022MS003589, 2023. ISSN 1942-2466. doi: 10.1029/2022MS003589. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2022MS003589>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2022MS003589>.
- S. A. Martin, G. E. Manucharyan, and P. Klein. Deep Learning Improves Global Satellite Observations of Ocean Eddy Dynamics. *Geophysical Research Letters*, 51(17):e2024GL110059, 2024. ISSN 1944-8007. doi: 10.1029/2024GL110059. URL <https://onlinelibrary.wiley.com/doi/abs/10.1029/2024GL110059>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2024GL110059>.
- Mercator Océan International. Global Ocean Physics Analysis and Forecast, 2024. URL <https://doi.org/10.48670/moi-00016>. Product ID: GLOBAL_ANALYSISFORECAST_PHY_001_024.
- MET Norway. Mediterranean Sea - High Resolution and Ultra High Resolution L3S Sea Surface Temperature, 2023. URL <https://doi.org/10.48670/moi-00171>. Processing Level: 3; Spatial Coverage: Mediterranean Sea (Lat 30.25° to 46°, Lon -18.12° to 36.25°); Spatial Resolution: 0.01° × 0.01°; Temporal Coverage: 1 Jan 2008 to 17 Dec 2024; Update Frequency: Daily; Format: NetCDF-4.
- R. Morrow, D. Blurmstein, and G. Dibarbour. Fine-scale altimetry and the future swot mission. *New frontiers in operational oceanography*, pages 191–226, 2018.
- NASA/JPL and CNES. The SWOT_L3_LR_SSH product, derived from the L2 SWOT KaRIn low rate ocean data products [Dataset], 2024. <https://doi.org/10.24400/527896/A01-2023.018>.
- J. O’Reilly, S. Maritorea, B. Mitchell, D. Siegel, K. Carder, S. Garver, M. Kahru, and C. McClain. Ocean color chlorophyll algorithms for SeaWiFS. *Journal of Geophysical Research*, 103:937–953, Oct. 1998.
- A. M. Scott, E. M. Georgy, and K. Patrice. NeuroST Level 4 Sea Surface Height and Surface Geostrophic Currents Analysis Product (Version 2024.0). Daily mapped neural network product using Level 3 altimetry observations and MUR Level 4 SST, 08 2024. URL <https://doi.org/10.5067/NEURO-STV24>. Start/Stop Date: 2010-Jan-01 to 2024-Jun-15; Format: netCDF-4; Processing Level: 4.
- G. Taburet, A. Sanchez-Roman, M. Ballarotta, M.-I. Pujol, J.-F. Legeais, F. Fournier, Y. Faugere, and G. Dibarbour. DUACS DT2018: 25 years of reprocessed sea level altimetry products. *Ocean Science*, 15(5):1207–1224, Sept. 2019. ISSN 1812-0784. doi: 10.5194/os-15-1207-2019. URL <https://os.copernicus.org/articles/15/1207/2019/>. Publisher: Copernicus GmbH.
- B. D. Tapley, G. H. Born, and M. E. Parke. The seasat altimeter data and its accuracy assessment. *Journal of Geophysical Research: Oceans*, 87(C5):3179–3188, 1982.
- S. Thiria, C. Sorrow, T. Archambault, A. Charantonis, D. Bereziat, C. Mejia, J.-M. Molines, and M. Crépon. Downscaling of ocean fields by fusion of heterogeneous observations using Deep Learning algorithms. *Ocean Modelling*, 182:102174, Apr. 2023. ISSN 1463-5003. doi: 10.1016/j.ocemod.2023.102174. URL <https://www.sciencedirect.com/science/article/pii/S146350032300015X>.
- M. Tonani, M. Balmaseda, L. Bertino, E. Blockley, G. Brassington, F. Davidson, Y. Drillet, P. Hogan, T. Kuragano, T. Lee, A. Mehra, F. Paranathara, C. A. Tanajura, and H. Wang. Status and future of global and regional ocean prediction systems. *Journal of Operational Oceanography*, 8(sup2):s201–s220, Aug. 2015. ISSN 1755-876X. doi: 10.1080/1755876X.2015.1049892. URL <https://doi.org/10.1080/1755876X.2015.1049892>. Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/1755876X.2015.1049892>.
- C. Ubelmann, L. Carrere, C. Durand, G. Dibarbour, Y. Faugère, M. Ballarotta, F. Briol, and F. Lyard. Simultaneous estimation of ocean mesoscale and coherent internal tide sea surface height signatures from the global altimetry record. *Ocean Science*, 18(2):469–481, Apr. 2022. ISSN 1812-0784. doi: 10.5194/os-18-469-2022. URL <https://os.copernicus.org/articles/18/469/2022/>. Publisher: Copernicus GmbH.
- W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao. PVT v2: Improved baselines with Pyramid Vision Transformer. *Computational Visual Media*, 8(3):415–424, Sept. 2022. ISSN 2096-0662. doi: 10.1007/s41095-022-0274-8. URL <https://doi.org/10.1007/s41095-022-0274-8>.
- X. Wang, R. Wang, N. Hu, P. Wang, P. Huo, G. Wang, H. Wang, S. Wang, J. Zhu, J. Xu, J. Yin, S. Bao, C. Luo, Z. Zu, Y. Han, W. Zhang, K. Ren, K. Deng, and J. Song. XiHe: A Data-Driven Model for Global Ocean Eddy-Resolving Forecasting, Feb. 2024. URL <http://arxiv.org/abs/2402.02995>. arXiv:2402.02995 [physics].
- Y. Wu and K. He. Group Normalization. *International Journal of Computer Vision*, 128(3):742–755, Mar. 2020. ISSN 1573-1405. doi: 10.1007/s11263-019-01198-w. URL <https://doi.org/10.1007/s11263-019-01198-w>.
- W. Yu, M. Luo, P. Zhou, C. Si, Y. Zhou, X. Wang, J. Feng, and S. Yan. Metaformer is actually what you need for vision. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10809–10819, 2022. doi: 10.1109/CVPR52688.2022.01055.