

Computer Vision for Ocean Eddy Detection in Infrared Imagery

Evangelos Moschos*[†] Alisa Kugusheva*[†] Paul Coste*[‡] Alexandre Stegner*[†]

* AMPHITRITE

X-novation Center, École Polytechnique, Palaiseau, France

[†] École Polytechnique, LMD/CNRS
Avenue Coriolis, Palaiseau, France

[‡]IMT Atlantique
Avenue du Technopôle, Plouzané, France

Abstract

Reliable and precise detection of ocean eddies can significantly improve the monitoring of the ocean surface and subsurface dynamics, besides the characterization of local hydrographical and biological properties, or the concentration pelagic species. Today, most of the eddy detection algorithms operate on satellite altimetry gridded observations, which provide daily maps of sea surface height and surface geostrophic velocity. However, the reliability and the spatial resolution of altimetry products is limited by the strong spatio-temporal averaging of the mapping procedure. Yet, the availability of high-resolution satellite imagery makes real-time object detection possible at a much finer scale, via advanced computer vision methods. We propose a novel eddy detection method via a transfer learning schema, using the ground truth of high-resolution ocean numerical models to link the characteristic streamlines of eddies with their signature (gradients, swirls, and filaments) on Sea Surface Temperature (SST). A trained, multi-task convolutional neural network is then employed to segment infrared satellite imagery of SST in order to retrieve the accurate position, size, and form of each detected eddy. The EddyScan-SST is an operational oceanographic module that provides, in real-time, key information on the ocean dynamics to maritime stakeholders.

1. Introduction

Eddies, dynamical structures are to the oceans what weather systems are to the atmosphere. By transporting heat, momentum and mass from their regions of formation to distant areas, they affect their biological productivity [4], water transport [35], local hydrographic properties [7] and

the movement of pelagic species [19]. Mesoscale eddies, with radii of tens of kilometers and timescales on the order of months, amount for an oceanic energy partition on the order of the large ocean circulation [12, 35]. This makes their detection and characterization crucial, both for the study of the climate evolution as well as day-to-day, operational oceanography.

The advances in Satellite Altimetry, measuring the Sea Surface Height (SSH) has led to 40 years of development of altimetric eddy detection [13], revealing the prevalent role of eddies on ocean dynamics [5]. As a result, many altimetric eddy detection and tracking algorithms [14, 27, 20, 16, 28] have been recently developed. These algorithms, employ an objective function to detect eddies either on the SSH field or the derived geostrophic velocity field, which also stems from altimetry. Despite the importance of standard altimetric eddy detection, strong limitations have been evoked: by simulating satellite altimetry products, Amores *et al.* [1] showed that altimetric detection only captures 6 to 16 % of eddies in the North Atlantic Ocean and Mediterranean Sea respectively. In addition, they have calculated a constant bias of artificially larger detected eddies compared to their real size. In a study of the Mediterranean Sea, Stegner *et al.* [32] showed that altimetric detection has a Missed (False Negative) rate of 34% and a Ghost (False Positive) rate of 10% for large ($R > 20km$) eddies. These limitations of standard detection stem mainly from the altimetric observation, due to the strong spatio-temporal interpolation of SSH maps (15km resolution), creating strong uncertainties in areas not covered by satellite tracks. Thus, the error of the altimetric data is of a larger order of magnitude than the error of the standard detection algorithms *per se*. On Figure 1 (a) we plot the contours detected on a simulated altimetric geostrophic velocity field and compare them to those of the reference geostrophic velocity field at panel (b). Standard

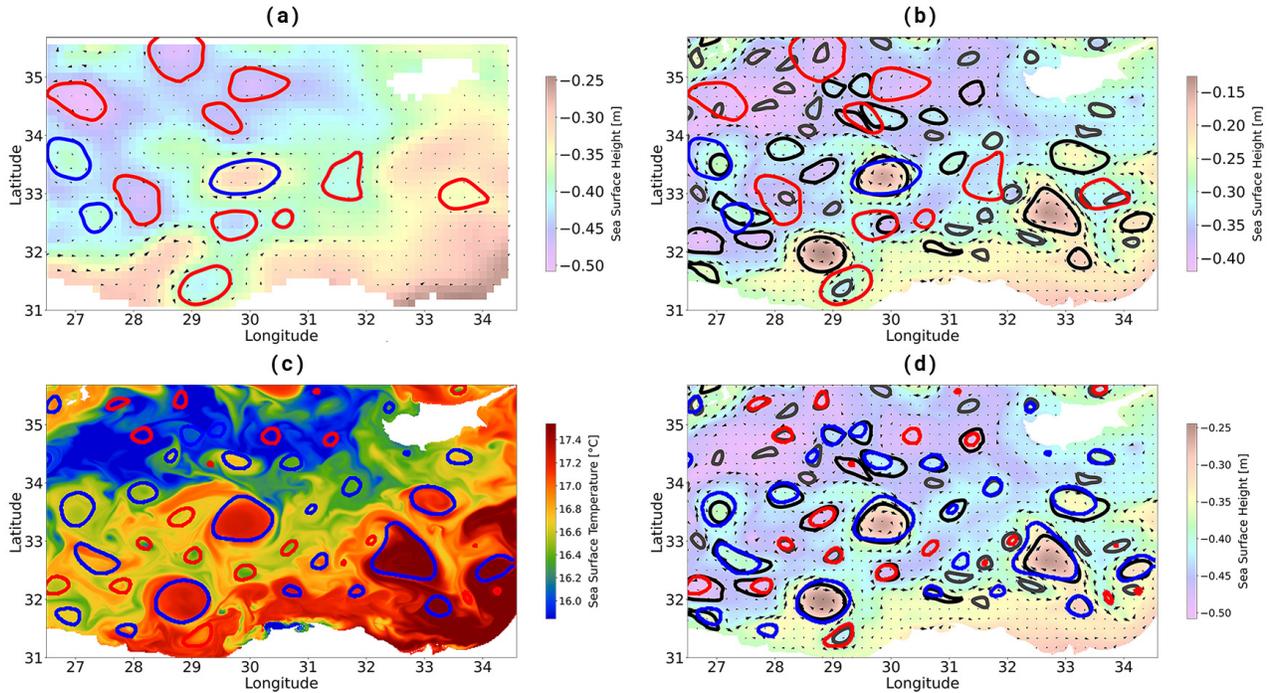


Figure 1. Examples of eddy detections on different satellite data. Blue (red) contours depict predicted anticyclones (cyclones). Black (grey) contours depict the mesoscale (submesoscale) reference geostrophic velocity eddies. (a) Standard eddy detection on simulated altimetric geostrophic velocity field (OSSE). (b) Reference contours and geostrophic velocity field with superimposed standard altimetric detections. (c) EddyScan CNN detections on SST data. (d) Reference contours and geostrophic velocity field with superimposed EddyScan detections.

detection struggles to capture small and submesoscale eddies, because of the low resolution of the altimetry fields, while also miss-detects many larger eddies in areas not covered by tracks. To surpass these limitations, other sources of satellite data should be considered: the Sea Surface Temperature (SST) are high-resolution ($1km$) observations on the infrared spectrum where the signature of eddies is portrayed through gradients, swirls, and filaments.

Convolutional Neural Networks (CNNs) have been successfully employed to detect eddies on satellite data. Most of the work using CNNs is applied to altimetric data to perform semantic segmentation of maps of eddies [10, 17, 9]. Yet, these successful implementations of Machine Learning stumble upon the inherent uncertainties of the altimetric observation data and the best of CNNs will only manage to replicate the best of standard eddy detection models. Two more studies [18, 15] have employed CNNs that fuse SSH and SST data to improve semantic segmentation characterizing eddies as warm or cold core. They showed that adding SST as an independent source of data can lead to an improvement of eddy detection. However, the SST training data used need to be hand labeled or inferred by altimetric detections resulting in sparse/noisy labeling. Finally, classification of eddy signatures has been performed successfully

by independently treating SST [23] and Synthetic Aperture Radar (SAR) data [8]. On Figure 1 (c) we show the detections on the SST field of the CNN proposed in this paper, EddyScan-SST on SST, and compare them with the reference velocity field. Due to the high-resolution and accurate signatures of eddies on the SST, we manage to capture correctly most of the mesoscale eddies and an important part of the submesoscale (black and grey contours respectively; Figure 1 (d)).

To perform an efficient operational eddy detection with CNNs two learning problems need to be treated successfully: the uncertainty of altimetric (SSH) eddy detections as ground truth, and the sparsity of ground truth in SST data. In this work we present EddyScan-SST, a pixel-wise segmentation CNN applied on SST data to detect with high precision the position, size, and form of eddies, without the need of an altimetry input. Our work provides several novelties in the task of eddy detection:

- Using CNNs, we infer from SST data the dynamical contours of eddies, which are local topological proxies of the field of velocities.
- As infrared satellite images are not linked with an accurate dynamical ground truth, we use a transfer learn-

ing schema, using ground truth from ocean numerical model simulations, to learn representations of eddy dynamical contours on SST data.

- We use a multi-task learning schema to perform contour detection with correct position, size and form on both numerical model data and satellite infrared images.
- By testing on numerical model data and satellite images, we find that the CNN applied on SST greatly outperforms standard altimetric eddy detection when infrared observations are available.
- We provide a hand-labeled dataset of satellite infrared images containing eddy signatures and their corresponding dynamical contours for use by the oceanographic and vision communities. The dataset can be found here.

2. Data: Simulation and Observation

2.1. The CROCO Numerical Model

The CROCO Numerical Model of the Ocean built upon the ROMS model [30] is a realistic numerical simulation of the ocean circulation, carried here on the domain of the Mediterranean Sea. CROCO is able to resolve very fine scales of ocean dynamics and their interactions with larger scales. The model solves the primitive equations on a grid, with a horizontal resolution of 2km in both longitudinal and latitudinal direction. We use in this study the SSH output of the numerical model as a reference (SSH REF) on which the Geostrophic Velocities are computed. We also use the output SST maps that correspond to the simulated dynamical field. These Numerical Model outputs serve as a reference ground truth for our experiments. Examples of the SSH and SST outputs of the CROCO Numerical Model are seen in panels (c) and (d) of Figure 1.

2.2. Observing System Simulation Experiment

The reference ground truth provided by the CROCO model simulations, as seen in Figure 2 (a), needs to be downsampled in an inhomogeneous manner, in order to reproduce the exact observation error found in the satellite altimetry. To do so, we perform an Observing System Simulation Experiment (OSSE) in a four-satellite configuration, composed of the reference mission Jason-3 and three other missions Sentinel3-A, Sentinel3-B, and Cryosat-2. Synthetic satellite tracks are reproduced through the SWOT simulator software [11], providing realistic measurement errors and noise. Example synthetic satellite tracks covering the CROCO model numerical field are depicted in Figure 2 (b). The resulting synthetically observed field is then processed to compute gridded fields with the same interpolation schema that is utilized for the production of gridded

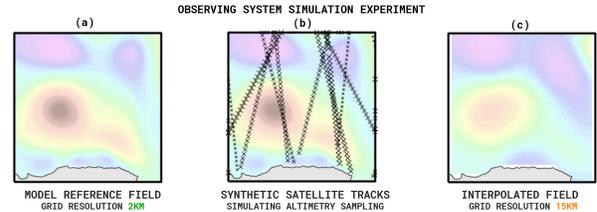


Figure 2. An OSSE consists of (a) Retrieving a high-resolution numerical model field of SSH (b) Sampling the field via synthetic satellite tracks which simulate observation by altimeters (c) Inhomogeneous spatio-temporal interpolation between the sampled points to receive the OSSE field. In our experiment grid resolution is degraded from 2km to 15km.

SSH satellite data [33]. It consists of an inhomogeneous spatiotemporal interpolation between the sampled points by the synthetic tracks, shown in Figure 2 (c). The resolution is thus downgraded from 2km of the initial numerical model field to 15km of the OSSE field. A preprocessing is also carried out on the CROCO simulation data to filter large scale, high-frequency signals derived from atmospheric forcing fields [3].

2.3. Satellite Data

By collating measurements made by multiple infrared sensors, high-resolution (1km) super-collated SST maps of the Mediterranean Sea are received from the Copernicus - Marine Environment Monitoring Service (CMEMS), Ultra High Resolution L3S SST Dataset, produced by the CNR - Italy and distributed by CMEMS. The process of supercollation uses SST measurements derived from the LSTR and AVHRR instruments on board Sentinel-3A/-3B and NOAA, VIIRS, MetOp-B, MODIS AQUA and TERRA, and SEVIRI on board the MSG satellite and are representative of nighttime SST values [26].

3. Methods and Learning

3.1. Standard Eddy Detection

The geostrophic velocity fields are derived from the SSH fields of the Reference Model simulation and the OSSE altimetric simulation, velocities being proportional to the gradient of the SSH. To represent dynamics on other non-dynamical variables such as SST (and therefore in imagery), we employ a topological proxy of the velocity field, the maximum velocity eddy contour. This contour is defined as the closed isoline around an eddy where its velocity is maximum (V_{max}). The plotted contours in this work correspond to the maximum velocity contour. To receive eddy contours and dynamical properties we use in this study the Angular Momentum Eddy Detection and Tracking Algorithm [16], which is employed on the calculated geostrophic velocity

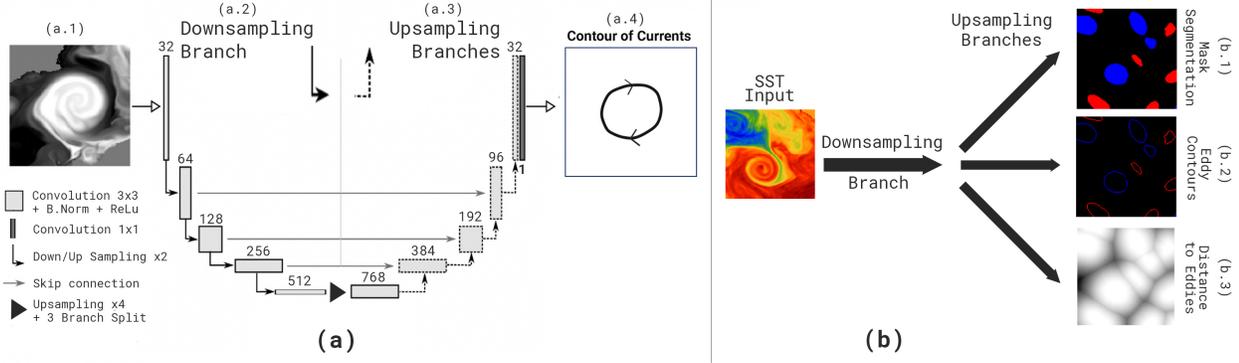


Figure 3. Schematic of the Neural Network used for Deep Eddy Scan. (a) The UNet-type architecture learns the mapping of an SST Input (a.1) via a downsampling branch (a.2) and three upsampling branches (a.3) to a dynamical contour (a.4). Skip connections are employed between the downsampling branch and each of the upsampling branches. (b) The SST Input follows a downsampling branch then leads to three upsampling branches predicting the (b.1) Mask Segmentation (b.2) Eddy Contours and (b.3) Distance to Eddies.

fields.

The identification of potential eddy centers by AMEDA is performed by computing the Local Normalised Angular Momentum (LNAM) [21] of the velocity field. Only eddy centers with at least one closed contour of the stream function of the velocity field are selected. A radial profile of the velocity for each detected eddy center is calculated by computing the average velocity and radius at each closed streamline around it:

$$\langle V \rangle = \frac{1}{L_p} \oint \vec{V} d\vec{l} \quad (1)$$

where \vec{V} is the local geostrophic velocity field and L_p is the streamline perimeter. The radius R of the characteristic contour is obtained by considering a circular contour of an equivalent area A :

$$\langle R \rangle = \sqrt{\frac{A}{\pi}} \quad (2)$$

The radii considered for sizing the eddies in this study correspond to the radius of the maximum velocity contour. Eddy centers and radius are important parameters used to retrieve SST patches for training and testing the Neural Network.

3.2. Convolutional Neural Network

We employ a CNN to learn the relation between the SST signature of an eddy (monochromatic image) with a corresponding maximum velocity dynamical contour. We treat this contour detection task, through a semantic segmentation of an image into regions of Anticyclones, Cyclones and No Eddies. Detected contours are subsequently extracted from the labeled regions.

UNET encoder-decoder architectures [29] have been successful in mapping low and high-resolution features of an

input image into a ground truth, due to skip connections between the downsampling and upsampling branches. Precisely, they have been successfully employed for the task of eddy detection [17, 15]. Multi-task learning approaches have proven robust on contour detection tasks [6, 34, 25]. By using additional output channels and corresponding loss function terms the CNN converges both faster and better to the task of contour detection. We employ an encoder-decoder CNN with one downsampling and three upsampling branches, following [25], each corresponding to a learning task. The encoder-decoder architecture is depicted in Figure 3 (a). The first and main learning task consists of learning the filled mask of each eddy, depicted in Figure 3 (b.1). The second task consists of learning a contour mask, which consists of the outline of each filled mask with a width of 2 pixels, depicted in Figure 3 (b.2). Learning contour masks has been reported to boost the detection of multiple contours in a single window [6]. A softmax activation is applied on the predictions before the calculation of the loss. For the first and the second task ($T1, T2$) we employ pixel-wise classification loss, for each pixel \mathbf{x} of a matrix:

$$L_{T1, T2} = \sum_{i,j} \log p(\mathbf{x}_{i,j}; \omega_c c(\mathbf{x}_{i,j})) \quad (3)$$

where $p(\mathbf{x}; l(\mathbf{x}))$ denotes the probability of predicting a class c between No Eddy, Anticyclone and Cyclone. A class weight ω_c is tuned so that Cyclones weight three times more than Anticyclones and No Eddy classes, as their representations are more complex to learn (see Figure 4). This weight tuning has been performed after experiments comparing validation metrics for anticyclones and cyclones, and seeking to boost the performance for the cyclonic class.

Finally, the third task consists of calculating a distance map, where filled masks containing eddies are labeled as zero and

for the non-eddy pixels we calculate the euclidean distance from the closest eddy contour, depicted in Figure 3 (b.3). Including a distance map in the multi-task learning has reportedly improved the smoothness of the retrieved contours, an important factor for realistic eddy detection [34]. A softmax activation function is performed to clamp the final values between zero and one before the calculation of the loss. For this third task (T_3) we calculate a pixel-wise MSE loss:

$$L_{T_3} = \sum_{i,j} (D'(\mathbf{x}_{i,j}) - D(\mathbf{x}_{i,j}))^2 \quad (4)$$

where $D(x)$ is the distance map of the ground truth and $D'(x)$ the predicted distance map. The three losses for L_{T_1} , L_{T_2} and L_{T_3} are summed up and weighted with weights $\lambda_{T_3} = 3 * \lambda_{T_{1,2}}$, in order to clamp values on the same order of magnitude.

3.3. Patch Creation on the fly

As the sampled domain of the Mediterranean Sea is large relative to the size of the eddies, we extract small windows, cropped randomly during the training process. To increase in the variance of the training samples while gaining in computational efficiency and memory load, we extract patches on the fly, during the training pipeline: First, patches of a size $(\pm 1.5) * 128px$ are cropped in a random location and day, having more than 80% of ocean pixel values. Then, these patches are interpolated to a constant size of 128px. Because of the multi-modality of the distribution of SST pixel values, patches are normalized and a mask containing the location of land values is added as a second input channel. We use one year of numerical model simulation of the Mediterranean sea as train data for the neural network and a second year as test data. This serves to avoid overlap between representations of the same eddy between the train and test datasets, as well as to guarantee a balance between representation of seasonal effects, which can affect significantly the signature of eddies on SST images [22]. Finally, the hand-labeled dataset provided here is too small (500 images) to use for direct supervised training and is only employed for validation.

3.4. Training and Validation

We use the above framework to generate patches on-the-fly as the network is trained. We define an epoch as an iteration of 1,000 training patches generated on the computer memory, passed to the GPU memory via batches of 16. After the end of each epoch, the CNN is validated on a constant set of 1,000 batches, extracted from the test test with an equal distribution between seasons. The evolution of the global train loss and the validation losses for the three tasks (equations 3 and 4) are shown in Figure 4 (a).

We define a global validation metric through the Intersection over Union (IoU), else called the Jaccard Index, which

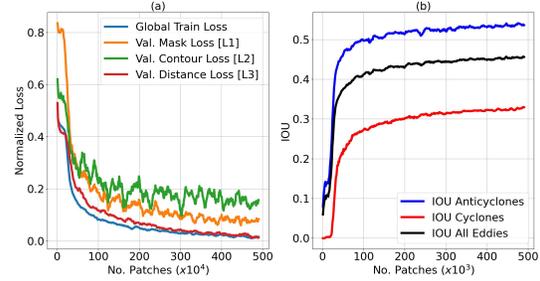


Figure 4. (a) Global train and validation losses for the Mask (L_{T_1}), Contour (L_{T_2}) and Distance (L_{T_3}). Loss values are standardized between zero and one. (b) Evolution of the global IOU metric on the validation set. On both x-axes is the number of patches iterated at training.

evaluates the global accuracy on the form of detected eddies. The IOU is calculated on each pixel-wise segmentation mask (T_1) as:

$$IoU(R, P) = \frac{|R \cap P|}{|R \cup P|} \quad (5)$$

where R is the reference and P is the predicted mask. Training with 500k patches takes 4 hours on 8GBs of GPU. However, we finally retrieved the trained model at 100k patches where validation loss gradients reduce significantly (Figure 4), in order to avoid overfit on numerical model data, as the final goal consists of an application on infrared satellite imagery. The asymmetry between the IOU scores for anticyclones and cyclones in the CNN validation can be explained by the more intrinsically complex signatures of cyclones on the SST images as well as their globally smaller size. These effects are linked with the different dynamical structure between anticyclones and cyclones as explained in [32].

We also evaluate the performance of the multi-task learning setup, by performing individual training runs for different combinations of learning tasks, for 100k patches each. Considering the three different tasks T_1 (Mask Segmentation), T_2 (Eddy Contours) and T_3 (Distance to Eddies) we train for the combinations T_1 , T_1+T_2 and $T_1+T_2+T_3$, the latter corresponding to the EddyNet-SST. We report the Global IOU scores of these three runs in Table 1, demonstrating the advantage of learning on both three tasks, in order to accurately detect the eddy form.

Training Task	T_1	T_1+T_2	$T_1+T_2+T_3$
IOU Score	0.4239	0.4257	0.4396

Table 1. Global IOU scores for different multi-task learning setups, training on combinations of tasks T_1 (Mask Segmentation), T_2 (Eddy Contours) and T_3 (Distance to Eddies)

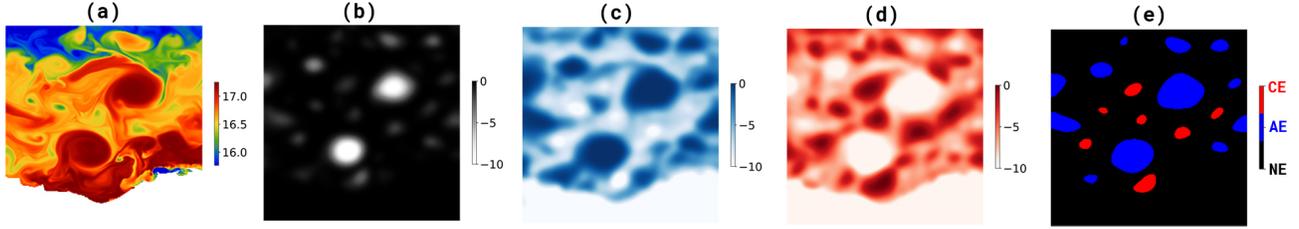


Figure 5. (a) An sample SST patch of the test set (size: 256^2 pixels) where the EddyScan CNN is applied. Negative likelihood heatmaps for the (b) No Eddy (c) Anticyclone and (d) Cyclone class. (e) Predicted pixel-wise segmentation mask (Task 1)

4. Results and Evaluation

4.1. Performance on Model Data

We first explore the network’s performance on the test dataset of the numerical model. An ensemble of 10,000 patches, of a size of 256^2 pixels each, is generated via the random window method from the numerical model simulation on the test year. By applying the EddyScan CNN on each SST patch - Figure 5 (a), we calculate the heatmaps of the output for task one (see Figure 3 (b.1)), for each of the three classes, as seen in Figure 5 (b)-(d). Segmentation is performed by retrieving the largest value, pixel-wise between the three heatmaps, as depicted in Figure 5 (e). On the pixel-wise classified image, we apply a contour detection algorithm to extract the boundaries of every eddy.

Additionally, a simple colocalization schema between predicted and reference eddies is performed on every patch. For each predicted eddy, we search for a barycenter of a reference eddy contour inside the predicted contour. If at least a reference eddy is found, we considered the predicted eddy as correctly detected. In this case, we calculate the position and size errors (see below) between the predicted eddies and all colocalized references and we match it with the one that is closest in position and size. If a predicted eddy has no colocalized reference eddy it is considered a *Ghost*. Finally, the eddy contours of the references which correspond to no predictions are named *Missed*.

As the eddy detection task is an object detection with underlying physical properties, we define relevant metrics to evaluate the performance of the CNN:

- An overall Precision and Recall metric.
- A metric on the error on the position of each eddy detection.
- A metric on the error of the size of each eddy detection.

The precision and recall scores are defined as:

$$\begin{aligned}
 Precision(c) &= \frac{|R_{i,c} \cap P_{i,c}|}{|P_{i,c}|} \\
 Recall(c) &= \frac{|R_{i,c} \cap P_{i,c}|}{|R_{i,c}|}
 \end{aligned} \tag{6}$$

The scores are calculated for each eddy object i and for each eddy class c , i.e. Anticyclones or Cyclones. $R_{i,c}$ and $P_{i,c}$ denote, respectively, reference and predicted eddies of each class. Through these metrics, we define the Ghost eddy rate, i.e. false positive detections $Ghost(c) = 1 - Precision(c)$ and the Missed Eddy rate i.e. false negative detections as $Missed(c) = 1 - Recall(c)$.

The precision and recall scores are shown in Figure 6 (a) and (b) respectively. We compare the scores of the EddyScan CNN, applied on SST images, with those of the Standard Eddy Detection applied on simulated altimetry fields. The latter, are received through an OSSE (see Figure 2) and our standard eddy detection experiment corresponds to those described by Amores et al. [1] and Stegner et al. [32]. The EddyScan CNN generally outperforms the standard eddy detection, as the eddy signatures on the SST are much closer to the reference dynamical field than those found in altimetric observations (Figure 1). The precision for Anticyclones is constantly higher for all eddy sizes, while cyclones have a lower ghost rate for radii up to 25km. On the recall metric, the CNN greatly outperforms standard eddy detection both on the mesoscale and submesoscale detections. For instance, the CNN missed rate 30% (40%) for small submesoscale Anticyclones (Cyclones) compared with 90% of missed small eddies by standard detection, due to the low resolution of the altimetry fields [1]. For large, mesoscale eddies with radii bigger than 25km, the CNN has a miss rate lower than 5% (10%) for Anticyclones (Cyclones). It is noted that the missed rate shown for the Standard Eddy Detection for large cyclones is overestimated: due to the colocalization schema followed, large cyclones of the altimetry corresponding to multiple smaller reference detections, are falsely labelled as correctly detected. This is portrayed in the increased error in size estimation for large cyclones in Figure 6 (d).

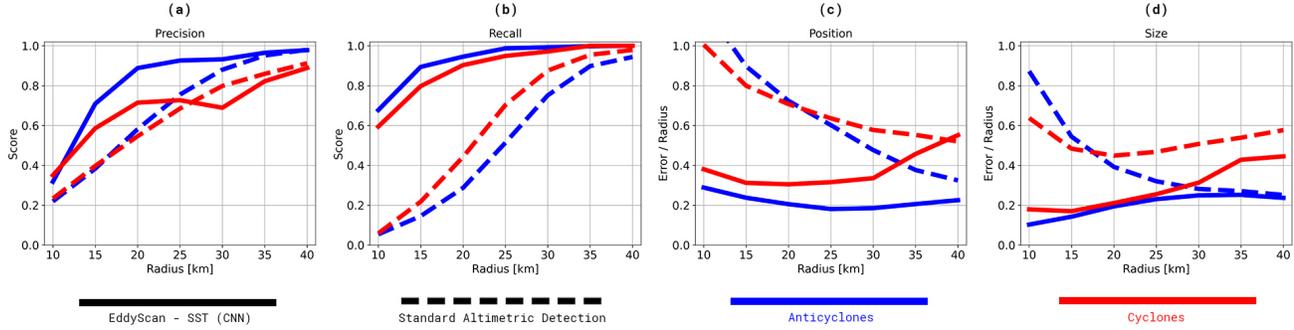


Figure 6. Performance indicators: (a) Precision (1-Ghost Rate) and (b) Recall (1 - Missed Rate). (c) Position (barycenter) and (d) Size (radius) mean errors normalized by the reference radius. Straight (dashed) lines correspond to the performance of the EddyScan-SST CNN (Standard Altimetric Detection). Blue (red) lines correspond to performance for anticyclones (cyclones).

We additionally quantify the error in the position and size of the correctly detected eddies. For the position error we consider the distance between the barycenters of two colocalized eddy contours, normalized by the radius of the reference eddy contour:

$$E_{pos} = \frac{\sqrt{(x_{pred}^{bar} - x_{ref}^{bar})^2 + (y_{pred}^{bar} - y_{ref}^{bar})^2}}{r_{ref}} \quad (7)$$

For the size error we consider the difference between the radii of the two colocalized eddy contours, normalized by the radius of the reference eddy contour:

$$E_{size} = \frac{|r_{pred} - r_{ref}|}{r_{ref}} \quad (8)$$

The mean errors, over all correctly detected eddies, on the position and the size are plotted in Figure 6 for both the EddyScan CNN and the Standard Altimetric Detection. On average, the EddyScan has an error of 20% of the radius when determining the center of an anticyclone, with an error of 30 – 50% for cyclones. Due to the degradation of the altimetry fields, the positioning of small eddies by standard methods reaches up to one radius of error. Size error for anticyclones does not exceed 20% for all sizes while the size of big cyclones tends to be underestimated by the CNN. In both cases, the EddyScan-SST outperforms the altimetric methods. Here again, cyclones prove more difficult to detect than their anticyclonic counterparts, due to their complex signatures, as discussed before.

4.2. Performance on Satellite Data

The EddyScan-SST serves as an operational oceanographic module, thus needing to be applied and evaluated on satellite observations of the ocean. In particular, infrared imagery is a proxy to obtain observations of the SST, which contains representations common to those learned by the CNN. However, the satellite data possesses some key differences from the numerical model data:

- The impact of noisy labels: As the altimetric observations of eddies often have low reliability, it is impossible to establish with certainty the link between a dynamic contour (altimetry) and the infrared satellite observation. Therefore there is a need for a transfer learning schema, using representations from models to train a CNN which will be applied to satellite data.
- The impact of cloud coverage: Infrared radiation cannot penetrate clouds, leading to a constant corruption of the observation by missing values. Furthermore, cloud presence corrupts the observed values of nearby pixels, creating noise in the data. Even though cloud coverage is out of the scope of this study, it has been demonstrated in [24] that CNNs can classify infrared eddy signatures even when impacted by strong cloud coverage of up to 80% .
- The need for hand labeling: stemming from the noisy label problem. As such we have extracted 500 patches, each one containing the ground truth of the dynamical contour of an Anticyclone (428 patches) or a Cyclone (72 patches). These patches are provided to the community with this paper and are used to evaluate the performance of the EddyScan on infrared observations.

In Figure 7 we depict several examples of eddy signatures found on infrared images: anticyclones and cyclones can have a coherent warm or cold anomaly in their cores depending on the season, as seen in panels (a)-(c). A typical case is a dipole structure as the one in panel (d), featuring an anticyclone rotating near a cyclone and constituting a dynamic ensemble. These pairs are often hard to detect on altimetry, as they are often formed by relatively small eddies, and have many times intense infrared signatures. Finally, as discussed, cloud coverage can hinder EddyScan prediction as seen in panel (e), which is why we have limited the maximum cloud coverage per patch at 10% for this dataset.

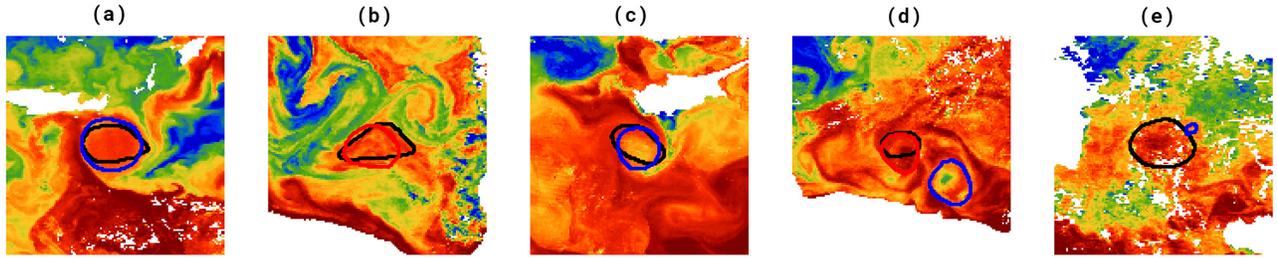


Figure 7. Samples from the Eddy Infrared Satellite Observations dataset released with this paper. Handlabeled contours are plotted with black. Predictions of the EddyScan CNN are plotted with blue for anticyclones and red for cyclones. Cases (a)-(c) portray coherent well-detected eddies. Case (d) portrays a typical structure of a Cyclone-Anticyclone dipole. Case (e) shows the impact of cloud coverage that can hinder prediction.

	ALL	AE	CE
IOU (eq. 5)	0.46	0.48	0.28
Miss Rate (eq. 6)	0.21	0.19	0.35
Position Error (eq. 7)	0.33	0.30	0.54
Size Error (eq. 8)	0.23	0.22	0.35

Table 2. EddyScan-SST performance on the hand labeled set of 500 satellite infrared images.

Applying the EddyScan on the 500 patches extracted, we receive the metric scores presented in Table 2. The overall IOU is calculated between the reference eddy and the corresponding predicted contour. All secondary contours are filtered out, as there is only one hand-labeled reference contour per image. Thus, ghost rate (precision) is not calculated here. The miss rate expresses how many of the patches had no CNN detection over the reference contour. The position error (equation 7), and the size error (equation 8) are expressed as a percentage of the radius of the reference contour. As with the numerical model data, performance on Anticyclones is significantly better than that on Cyclones, as the latter tend to be much smaller, and with more complex signatures.

5. Conclusion

In this work, we employ computer vision to perform eddy detection on infrared satellite imagery, surpassing the limitations of standard eddy detection methods on altimetry. Our CNN trained on patches of SST from a simulation of a high-resolution ocean numerical models, achieves important performance scores on the task of eddy detection with a global IOU of 0.45 and a Miss Rate of 0 – 20% for mesoscale eddies and 15%-40% for small and submesoscale eddies ($< 15km$). To compare the performance of the CNN with standard detection methods we perform a simulation of the satellite altimetry observation on numerical model fields. Standard detection methods have a miss rate of 10 – 80% for mesoscale eddies, while missing completely the small submesoscale eddies due to the low reso-

lution of altimetry data.

High-resolution satellite imagery, such as infrared measurements, remains largely unused for ocean structure detection, despite the rich amount of information contained in patterns of gradients, swirls and filaments. To extract the dynamical information from these complex representations in infrared imagery we utilize the topological information of eddy contours as a proxy of the surface dynamics of the ocean. Our method does not replace, but rather compliments standard eddy detection on altimetry, especially on spatio-temporal windows not covered by altimeter satellites, while also providing a reliability index.

We employ a transfer learning schema, using ground truth from a numerical model simulation of the ocean, with an almost perfect link between dynamics and temperature signature, applying afterwards the trained network on satellite observations. This way we avoid the noisy labeling of satellite data and the costly hand-labeling process. We treat the contour detection task as a multi-task learning that aids to retrieve information on the contours size, position and form. Finally, the CNN is invariant of the input size, meaning that it can be applied in any region of the globe.

The trained eddy-detecting neural network proves robust on infrared satellite imagery with a 20% miss rate of eddies and a mean error of 23% on their size and 33% on their position. We release through this repository a hand-labeled dataset containing eddy signatures on infrared observations along with their dynamical contours.

Including multi-modal satellite data such as satellite observations on the visible spectrum and synthetic aperture radar [8] as well as finetuning the neural network for cloud coverage [24] can boost the operational performance of the EddyScan module. Employing semi-supervised learning [31] can allow for learning with noisy-labeled satellite data. Furthermore, super-resolution neural networks, fusing satellite altimetry and infrared imagery [2] could work in conjunction with the eddy detection neural networks.

Precise and reliable eddy detection allows us to estimate, in real-time, the local hydrographic properties and the surface circulation in a given region, providing key information for many applications of maritime stakeholders.

References

- [1] Angel Amores, Gabriel Jordà, Thomas Arsouze, and Julien Le Sommer. Up to what extent can we characterize ocean eddies using present-day gridded altimetric products? *Journal of Geophysical Research: Oceans*, 123(10):7220–7236, 2018.
- [2] Bruno Buongiorno Nardelli, Davide Cavaliere, Elodie Charles, and Daniele Ciani. Super-resolving ocean dynamics from space with computer vision algorithms. *Remote Sensing*, 14(5):1159, 2022.
- [3] L Carrère and F Lyard. Modeling the barotropic response of the global ocean to atmospheric wind and pressure forcing-comparisons with observations (doi 10.1029/2002gl016473). *Geophysical Research Letters*, 30(6):8–8, 2003.
- [4] Dudley B Chelton, Peter Gaube, Michael G Schlax, Jeffrey J Early, and Roger M Samelson. The influence of nonlinear mesoscale eddies on near-surface oceanic chlorophyll. *Science*, 334(6054):328–332, 2011.
- [5] Dudley B Chelton, Michael G Schlax, and Roger M Samelson. Global observations of nonlinear mesoscale eddies. *Progress in oceanography*, 91(2):167–216, 2011.
- [6] Hao Chen, Xiaojuan Qi, Lequan Yu, and Pheng-Ann Heng. Dcan: deep contour-aware networks for accurate gland segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2487–2496, 2016.
- [7] Changming Dong, James C McWilliams, Yu Liu, and Dake Chen. Global heat and salt transports by eddy movement. *Nature communications*, 5(1):1–6, 2014.
- [8] Yanling Du, Wei Song, Qi He, Dongmei Huang, Antonio Liotta, and Chen Su. Deep learning with multi-scale feature fusion in remote sensing for automatic oceanic eddy detection. *Information Fusion*, 49:89–99, 2019.
- [9] Zijun Duo, Wenke Wang, and Huizan Wang. Oceanic mesoscale eddy detection method based on deep learning. *Remote Sensing*, 11(16):1921, 2019.
- [10] Katharina Franz, Ribana Roscher, Andres Milioto, Susanne Wenzel, and Jürgen Kusche. Ocean eddy identification and tracking using neural networks. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 6887–6890. IEEE, 2018.
- [11] Lucile Gaultier, Clément Ubelmann, and Lee-Lueng Fu. The challenge of using future swot data for oceanic field reconstruction. *Journal of Atmospheric and Oceanic Technology*, 33(1):119–126, 2016.
- [12] AE Gill, JSA Green, and AJ Simmons. Energy partition in the large-scale ocean circulation and the production of mid-ocean eddies. In *Deep sea research and oceanographic abstracts*, volume 21, pages 499–528. Elsevier, 1974.
- [13] Greg Holloway. Estimation of oceanic eddy transports from satellite altimetry. *Nature*, 323(6085):243–244, 1986.
- [14] Jordi Isern-Fontanet, Emilio García-Ladona, and Jordi Font. Identification of marine eddies from altimetric maps. *Journal of Atmospheric and Oceanic Technology*, 20(5):772–778, 2003.
- [15] Devyani Lambhate, Raghav Sharma, Jenifer Clark, Avijit Gangopadhyay, and Deepak Subramani. W-net: A deep network for simultaneous identification of gulf stream and rings from concurrent satellite images of sea surface temperature and height. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2021.
- [16] Briac Le Vu, Alexandre Stegner, and Thomas Arsouze. Angular momentum eddy detection and tracking algorithm (ameda) and its application to coastal eddy formation. *Journal of Atmospheric and Oceanic Technology*, 35(4):739–762, 2018.
- [17] Redouane Lguensat, Miao Sun, Ronan Fablet, Pierre Tandeo, Evan Mason, and Ge Chen. EddyNet: A deep neural network for pixel-wise classification of oceanic eddies. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 1764–1767. IEEE, 2018.
- [18] Yingjie Liu, Qunan Zheng, and Xiaofeng Li. Characteristics of global ocean abnormal mesoscale eddies derived from the fusion of sea surface height and temperature data by deep learning. *Geophysical Research Letters*, 48(17):e2021GL094772, 2021.
- [19] PS Lobel and AR Robinson. Transport and entrapment of fish larvae by ocean mesoscale eddies and currents in hawaiian waters. *Deep Sea Research Part A. Oceanographic Research Papers*, 33(4):483–500, 1986.
- [20] Evan Mason, Ananda Pascual, and James C McWilliams. A new sea surface height-based code for oceanic mesoscale eddy tracking. *Journal of Atmospheric and Oceanic Technology*, 31(5):1181–1188, 2014.
- [21] Nadia Mkhinini, Andre Louis Santi Coimbra, Alexandre Stegner, Thomas Arsouze, Isabelle Taupier-Letage, and Karine Béranger. Long-lived mesoscale eddies in the eastern mediterranean sea: Analysis of 20 years of aviso geostrophic velocities. *Journal of Geophysical Research: Oceans*, 119(12):8603–8626, 2014.
- [22] Evangelos Moschos, Alexandre Barboni, and Alexandre Stegner. Why do inverse eddy surface temperature anomalies emerge? the case of the mediterranean sea. *Remote Sensing*, 14(15):3807, 2022.
- [23] Evangelos Moschos, Olivier Schwander, Alexandre Stegner, and Patrick Gallinari. Deep-sst-eddies: A deep learning framework to detect oceanic eddies in sea surface temperature images. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4307–4311. IEEE, 2020.
- [24] Evangelos Moschos, Alexandre Stegner, Olivier Schwander, and Patrick Gallinari. Classification of eddy sea surface temperature signatures under cloud coverage. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13:3437–3447, 2020.
- [25] Balamurali Murugesan, Kaushik Sarveswaran, Sharath M Shankaranarayana, Keerthi Ram, Jayaraj Joseph, and Mohanasankar Sivaprakasam. Psi-net: Shape and boundary aware joint multi-task deep network for medical image segmentation. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 7223–7226. IEEE, 2019.
- [26] B Buongiorno Nardelli, C Tronconi, A Pisano, and R Santoleri. High and ultra-high resolution processing of satellite

sea surface temperature data over southern european seas in the framework of myocean project. *Remote Sensing of Environment*, 129:1–16, 2013.

- [27] Francesco Nencioli, Changming Dong, Tommy Dickey, Libe Washburn, and James C McWilliams. A vector geometry-based eddy detection algorithm and its application to a high-resolution numerical model product and high-frequency radar surface velocities in the southern california bight. *Journal of atmospheric and oceanic technology*, 27(3):564–579, 2010.
- [28] Cori Pegliasco, Antoine Delepoulle, Evan Mason, Rosemary Morrow, Yannice Faugère, and Gérald Dibarboure. Meta3.1exp: a new global mesoscale eddy trajectory atlas derived from altimetry. *Earth System Science Data*, 14(3):1087–1107, 2022.
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [30] Alexander F Shchepetkin and James C McWilliams. The regional oceanic modeling system (roms): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean modelling*, 9(4):347–404, 2005.
- [31] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020.
- [32] Alexander Stegner, Briac Le Vu, Franck Dumas, M Ali Ghannami, Amandine Nicolle, C Durand, and Y Faugere. Cyclone-anticyclone asymmetry of eddy detection on gridded altimetry product in the mediterranean sea. *Journal of Geophysical Research: Oceans*, 126(9):e2021JC017475, 2021.
- [33] Guillaume Taburet, Antonio Sanchez-Roman, Maxime Ballarotta, Marie-Isabelle Pujol, Jean-François Legeais, Florent Fournier, Yannice Faugere, and Gerald Dibarboure. Duacs dt2018: 25 years of reprocessed sea level altimetry products. *Ocean Science*, 15(5):1207–1224, 2019.
- [34] Chaowei Tan, Liang Zhao, Zhennan Yan, Kang Li, Dimitris Metaxas, and Yiqiang Zhan. Deep multi-task and task-specific feature learning network for robust shape preserved organ segmentation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1221–1224. IEEE, 2018.
- [35] Zhengguang Zhang, Wei Wang, and Bo Qiu. Oceanic mass transport by mesoscale eddies. *Science*, 345(6194):322–324, 2014.