KaRIn Noise Reduction Using a Convolutional Neural Network for the SWOT Ocean Products

Anaëlle Tréboutte, Elisa Carli, Maxime Ballarotta, Benjamin Carpentier, Yannice Faugère, Gérald Dibarboure

remote sensing

MDPI

Article

KaRIn Noise Reduction Using a Convolutional Neural Network for the SWOT Ocean Products

Anaëlle Tréboutte [1,*], Elisa Carli [2], Maxime Ballarotta [1], Benjamin Carpentier [1], Yannice Faugère [1] and Gérald Dibarboure [3]

[1] Collecte Localisation Satellites (CLS), 31520 Ramonville Saint-Agne, France
[2] Laboratoire d'Etudes en Géophysique et Océanographie Spatiale (LEGOS), 31400 Toulouse, France
[3] Centre National d'Etudes Spatiales (CNES), 31400 Toulouse, France
* Correspondence: atreboutte@groupcls.com

Abstract: The SWOT (Surface Water Ocean Topography) mission will provide high-resolution and two-dimensional measurements of sea surface height (SSH). However, despite its unprecedented precision, SWOT's Ka-band Radar Interferometer (KaRIn) still exhibits a substantial amount of random noise. In turn, the random noise limits the ability of SWOT to capture the smallest scales of the ocean's topography and its derivatives. In that context, this paper explores the feasibility, strengths and limits of a noise-reduction algorithm based on a convolutional neural network. The model is based on a U-Net architecture and is trained and tested with simulated data from the North Atlantic. Our results are compared to classical smoothing methods: a median filter, a Lanczos kernel smoother and the SWOT de-noising algorithm developed by Gomez-Navarro et al. Our U-Net model yields better results for all the evaluation metrics: 2 mm root mean square error, sub-millimetric bias, variance reduction by factor of 44 (16 dB) and an accurate power spectral density down to 10–20 km wavelengths. We also tested various scenarios to infer the robustness and the stability of the U-Net. The U-Net always exhibits good performance and can be further improved with retraining if necessary. This robustness in simulation is very encouraging: our findings show that the U-Net architecture is likely one of the best candidates to reduce the noise of flight data from KaRIn.

Keywords: SWOT; KaRIn; altimetry; noise reduction; neural networks

Citation: Tréboutte, A.; Carli, E.; Ballarotta, M.; Carpentier, B.; Faugère, Y.; Dibarboure, G. KaRIn Noise Reduction Using a Convolutional Neural Network for the SWOT Ocean Products. *Remote Sens.* **2023**, *15*, 2183. https://doi.org/10.3390/rs15082183

Academic Editor: Piotr Samczynski

Received: 7 February 2023
Revised: 7 April 2023
Accepted: 18 April 2023
Published: 20 April 2023

1. Introduction

Until now, observations of the dynamics of the ocean surface are done by nadir altimetric satellites. Nevertheless, the smallest structures observed by these altimeters are larger than 150 km [1,2]. To access the mesoscale and submesoscale dynamics of the ocean, a new satellite called SWOT (Surface Water and Ocean Topography) was launched in December, 2022 [3]. It is a joint mission between the National Aeronautics and Space Administration (NASA), the Centre National d'Études Spatiales (CNES) and the UK and Canadian space agencies. To achieve the objectives for oceanography, the Sea Surface Height (SSH) will be measured by the two KaRIn (Ka-band Radar Interferometer) instruments [4,5]. They will provide a 120 km wide swath of coverage with a resolution of 2 km.

SWOT will measure oceanic structures down to 15 km [6]. Hence, new phenomena will be studied more precisely such as kinetic energy, vertical transport, coastal ocean dynamics and internal tides [2,6]. Therefore, the first and second derivatives of the SSH will be needed to investigate quantities such as geostrophic velocities, vorticity and strain rate. However, despite its unprecedented resolution and precision (approximately five times less noise than a Jason-class altimeter), the KaRIn instrument has a significant amount of wave-modulated noise, hereafter referred to as KaRIn noise [7,8]. As expected from Chelton et al. [9], the first and the second derivatives of a noisy field at a 2 km resolution simply cannot be exploited as is. It is necessary to use noise-reduction algorithms—either generic image processing algorithms or SWOT-specific noise-mitigation algorithms [1,10].

Denoising data and especially images is a well-known topic in image processing. The main objective is to remove the noise from a signal while preserving its features of interest (e.g., textures, contours) without adding unwanted artefacts. In the context of SWOT ocean products, we want to reduce the random noise while leaving the ocean features intact on the SSH and its derivatives.

To achieve this, many techniques have been developed for different types of noise (additive white Gaussian noise, salt and pepper noise, Poisson noise ...) and for different databases (imageNet, Set12 ...) [11–13]. First, classical methods (i.e., not based on machine learning methods) have been developed such as the median filtering or the Lanczos smoother. The boxcar, Gaussian and Laplacian smoothers are also considered classical methods, but Gomez-Navarro et al. [1,10] show that the results are not satisfying and, therefore, they are not reused in this study.

With the development of artificial intelligence and convolutional neural networks (CNN), and particularly since the development of the AlexNet model [14], hundreds of image de-noising models [12,13] have been developed and outperform conventional methods [11]. These methods are used in many applications such as medical images [15,16] and Synthetic Aperture Radar (SAR) images [17].

In that context, the goal of this work is to develop a neural network model to remove the KaRIn noise on simulated SWOT products, and to compare its de-noising performance with various alternatives. The paper is organized as follows: Section 2 presents the generation of the SWOT data using the SWOT scientific simulator, then Section 3 presents the U-Net model and various other de-noising methods. Section 4 gives an overview of the performance of all noise mitigation algorithms tested in this study. Lastly, in Section 5, we test the robustness of the neural network model when it is used in non-perfect conditions.

2. Input Data

The inputs of the de-noising filter are based on the North Atlantic Ocean simulation at 1/60° grid resolution called eNATL60. This simulation is a configuration of the NEMO (Nucleus for a European Model of the Ocean) ocean model. The data used do not contain the effect of oceanic tides but include the effect of the wind and pressure (no dynamic atmospheric correction used, as in [18]) provided by an atmospheric forcing (ERA-interim). Please refer to [19,20] for source files, codes and more information. Several ocean models are available in the ocean community but the eNATL60 model provides one of the best resolutions: between 0.8 and 1.6 km depending on latitude.

SWOT swaths calculated from the eNATL60 are generated by the SWOT simulator [8] which is available at https://swot-simulator.readthedocs.io/ (accessed on 7 April 2023). The Sea Surface Height (SSH) variable is interpolated on the SWOT grid at a resolution of 2 km. Several errors are simulated: geophysical errors, including wet troposphere and sea state bias, as well as instrumental errors. The latter are roll errors, phase errors, baseline dilation errors, timing errors, orbital errors and KaRIn noise. Only the KaRIn noise is a decorrelated random error, and it impacts wavelengths approximately under 50 km. In this study, we will only focus on this random error. Other errors are described in Esteban-Fernandez [7] and Gaultier et al. [8]. The KaRIn noise is defined by a Gaussian zero-centered distribution. The associated standard deviation depends on the distance to the nadir and the Significant Wave Height (SWH), which varies between 0 and 8 m [7] as shown in Figure 1. The error increases with the SWH and is more important on the borders of the swath. In our simulations, real SWH from the global ocean reanalysis wave system of Météo-France (WAVERYS) with a resolution of 1/5° degree is specified as input. If no data is available at a particular location, the SWH is set at 2 m. More information about the WAVERYS product is available on https://doi.org/10.48670/moi-00022 (accessed on 7 April 2023).

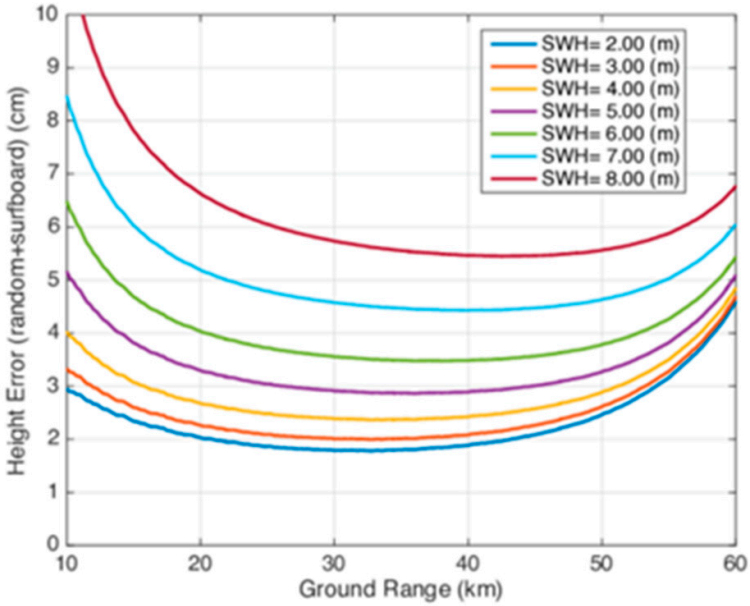

Figure 1. Height error (cm) of the KaRIn noise as a function of cross-track distance (km) from the nadir and the SWH (extracted from the documentation of the SWOT simulator). The KaRIn noise is more important at the borders of the swath and increases with the SWH. Thus, it varies geographically and temporally. These properties impact the performance of the classical smoothing.

In practice, the SWOT simulator provides the SSH from the eNATL60 model interpolated on the swath grid (hereafter referred to as true SSH), the values of each error and the SWH interpolated on the swath grid. These data are available from 1 July 2009 to 31 July 2010 over the North Atlantic Ocean.

The true SSH defines the ground truth. Noisy SSH is constructed as follows:

$$Noisy\ SSH = True\ SSH + KaRIn\ noise$$

where *KaRIn* noise is modulated by the waves. Only values above the ocean are considered; other values are set to zero to avoid training over land. As the neural network model described below requires data with a constant shape, each SWOT track is divided into 512 km long sections. Thus, input data images have a shape of 256 $\times$ 70 pixels.

To improve the performance of the model, two steps are performed. The first one is to remove the temporally averaged SSH model (calculated over the 13 months available) to work on Sea Level Anomalies (SLA). Second, the data is scaled: a standard scaler is used which centers the data by removing the mean and scaling to the unit variance of each patch of 256 $\times$ 70 pixels.

Data from the year 2009 are used and are divided into training and validation sets using, respectively, 75% and 25% of the dataset. Data from the year 2010 is used for the benchmark as the test set.

3. Noise-Reduction Algorithms

3.1. Neural Network Method

3.1.1. Convolutional Neural Network (CNN)

For several years, Convolutional Neural Networks (CNN) have been very successful and have been rapidly developed for various computer vision tasks [21]. A CNN consists of a succession of convolutional layers defined by:

$$h_{ij}^{k} = f\left(\left(W^{k} * x\right)_{ij} + b_{k}\right)$$

where h_{ij}^k is the output of layer k, commonly called the feature map, $*$ represents the convolutional operator, W^k is the weights and b_k the biases. f is a non-linear function called the activation function. In our case, the Rectifier Linear Unit is chosen and is defined by:

$$ReLU\ (x) = \max(0,\ x)$$

This function has several advantages that others such as the sigmoid or tanh do not have; it can speed up learning and it avoids vanishing gradients in the backpropagation [21,22].

The weights and biases are the parameters to be learned. The learning step consists of minimizing a loss function and updating the parameters using the backpropagation algorithm. For more information, please refer to Albawi et al. [22] and Li et al. [21].

3.1.2. U-Net Architecture for Denoising

CNNs are often used in an autoencoder framework. Their architecture is adapted to denoising problems [23]. Actually, they are composed of two parts: the encoder and the decoder. The encoder part subsequently downscales the input resolution and eventually creates a compact representation of the input data in a lower-dimensional space. The objective of the encoder is to keep the essential part of the data (i.e., data without noise in a denoising scenario). The decoder reconstructs the data from the low-dimensional space back to the original input data space [24].

One of the most successful encoder–decoder architectures is called U-Net. It was first proposed for biomedical image segmentation by Ronneberger et al. [25]. Its main feature is the use of skip connections between the encoder and the decoder to improve higher-resolution information flow. Many studies have used U-Net-like architectures to denoise images [16,26]. U-Net-like architectures also perform well for other applications such as segmentation [25,27] or nowcasting [28].

In our study, we designed a U-Net to remove the noise present in the SSH field. Figure 2 describes the U-Net architecture. The encoder part is made of three CNN blocks (each composed of two CNN layers) and, between each block, a max pooling layer. This layer divides the size of the feature map by two. The filter numbers of each CNN block are 16, 32 and 64. The decoder part is similar to the encoder, but the max pooling layer is replaced by transpose convolution layer. It multiplies the size of the feature map by two in order to reconstruct the data.

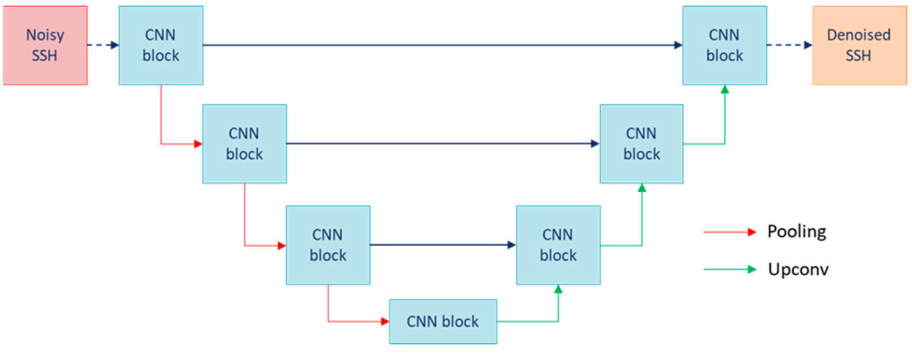

Figure 2. Architecture of the U-Net used for removing the KaRIn noise from a SSH field.

3.1.3. Experimental Setup

To train the U-Net model, the denoised SSH must be compared to the ground truth using a loss function. The two most used are L1-loss and L2-loss. Several studies [16,29] show that L1-loss provides the best results. Actually, L2-loss penalizes large errors more than small ones due to its definition. Therefore, L1-loss is chosen and is defined as follows:

$$L^{l_1} = \frac{1}{N} \sum_{p \in P} |x - y|$$

where x and y are the denoised SSH and the true SSH, respectively. It is calculated only on values above the ocean and on valid ocean values provided by a mask.

Data augmentation of the training dataset is one of the most common techniques to improve results [30]. In our study, the simplest data augmentation is the horizontal and/or vertical flip. It increases the training dataset by a factor of four. Please note that the learning is stopping after 15 epochs without improvement of the L1-loss of the validation dataset in order to avoid overfitting.

The network is coded using Pytorch [31] and Pytorch Lightning, two Python libraries. It is composed of 231,000 trainable parameters. The ADAM optimizer [32] is used with a learning rate of 0.001. The experiments were run on an NVIDIA Tesla V100 GPU available for high-performance computing provided by the CNES (more information on https://centredecalcul.cnes.fr/en/data-processing-centre, (accessed on 7 April 2023)). The hyperparameters (learning rate, batch size, number of filters . . .) are tuned with optuna, a python package [33].

For the training part, the computation time is around 2 h on GPU. Once it is done, the inference part is rapid: it takes around 30 s on CPU to denoise 100 swaths.

3.1.4. Data Post-Processing

The objective of this part is to reconstruct the swaths as the input data. As the swaths are divided every 512 km along-track, processing was performed to join them to avoid discontinuities. To do this, the first 10 pixels of 512 km along track swath are shared with the previous 512 km along-track swath. Similarly, the last 10 pixels are identical to the first 10 pixels of the next 512 km along-track swath. The pixels located at the ends of the swath are denoised twice by the inference. These pixels are weighted by the following function:

$$f(x) = \frac{1}{2}(\tanh(6x + 3) + 1)$$

where $x \in [0,\ 1]$. $x = 0$ corresponds to the pixel at the end of the swath, and $x = 1$ to the 10th pixel from the end. This function has these properties: $f(0) \approx 0$, $f(1) \approx 1$, $f(0.5) = 0.5$, f has rotational symmetry with respect to (0.5, 0.5) and f is infinitely differentiable. This last property is useful to avoid discontinuities when computing velocity and vorticity fields.

3.2. Other De-Noising Algorithms

In Section 4, the method above is compared to more "traditional" filtering/smoothing methods, i.e., algorithms that do not use a neural network. This is not an exhaustive list of all the "traditional" methods. They are frequently used in image denoising and/or developed and tested on SWOT data. Other "traditional" methods could provide better results than ones presented in Section 4 such as the method presented in [34].

Firstly, spatial-based filters are commonly used in image filtering [11]. In our experiments, we test two spatial-based filters: the median filter and the Lanczos smoother. Other filters used on SWOT data in Gomez-Navarro et al. [1] (Boxcar filter and Gaussian filter) and in Gómez-Navarro et al. [10] (Laplacian filter) do not provide satisfying results and, thus, are not reused in this study. The parameters of the median and the Lanczos smoother are window size and wavelength (l_x), respectively. Gaps are present in the data corresponding to land, islands, missing data or nadir gaps. As a better result of these filters is obtained on data without gaps, two operations are calculated on the data and on the mask associated with gaps with the same kernel. The filtered data is divided by the filtered mask. Finally, the mask is applied to filtered data. To implement these filters, the functions from scipy.ndimage (accessed on 7 April 2023) (a Python module) are used. The parameters of each filter are set as 7 for the window size and 5 for the wavelength.

Secondly, a variational filter was specifically developed for SWOT data by Gomez-Navarro et al. [1]. Its objective is to minimize the following cost function:

$$J(h) = \frac{1}{2}\|m \circ (h - h_{obs})^2\|^2 + \frac{\lambda_2}{2}\|\Delta h\|^2$$

where m is the mask, h the denoised image, h_{obs} the noisy image and $\|.\|$ the L2-norm. The parameter λ_2 must be optimized. Gomez-Navarro et al. [1] show that this parameter depends on the study area, the season and the ocean model. λ_2 varies between 1 and 54. Hereafter, we use $\lambda_2 = 10$ which is suitable for high-variability areas and for the eNATL60 model. This value is not optimal for low-variability areas; a new value of λ_2 for these areas should be set (not done in this study). Therefore, the results of this method presented in Section 4 could be improved. This variational problem is solved by using a gradient descent method. Hereafter, this filter will be referred to as the Gomez filter.

The computation time of these methods is longer than the U-Net in inference: about 45 min for 100 swaths.

3.3. Diagnostics for Evaluation

To quantify the denoising methods, a benchmark was developed. It provides the following metrics:

- Root Mean Square Error (RMSE): $RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(h_{noisyi} - h_{truei}\right)^2}$;
- Mean of SSH residuals: $Mean_{SSHres} = \frac{1}{N}\sum_{i=1}^{N}\left(h_{noisyi} - h_{truei}\right)$;
- Variance of SSH residuals: $Var_{SSHres} = \frac{1}{N}\sum_{i=1}^{N}\left(h_{noisyi} - h_{truei}\right)^2$;
- Resolved scale corresponding to the Signal-to-Noise-Ratio (SNR) equals to i.e., $\frac{PSD\left(h_{noisy} - h_{true}\right)}{PSD\left(h_{true}\right)} = 1$ [34];
- KaRIn noise reduction: $R = 10\log\left(\frac{RMSE^2_{before\ denoising}}{RMSE^2_{after\ denoising}}\right)$ (in dB).

h_{noisy} and h_{true} are, respectively, the denoised field and the true field, while PSD represents the Power Spectral Density. All these metrics must be as close to zero as possible. The mean of *SSH* residuals traduces the bias made by a filter while all the other metrics traduce the performance. The variance of SSH residuals and the RMSE reflect the performance of the models. The resolved scale corresponding to SNR = 1 provides information on the wavelength at which the noise dominates the oceanic signal.

The RMSE and the resolved scale are calculated on the SSH field, the geostrophic current field and the relative vorticity field. The geostrophic current and the relative vorticity are calculated as follows:

$$V = \frac{g}{f}\sqrt{\left(\frac{dh}{dx}\right)^2 + \left(\frac{dh}{dy}\right)^2} \text{ and } \zeta = \frac{1}{f}\left(\frac{d^2h}{dxdy} - \frac{d^2h}{dydx}\right)$$

with h representing the noisy or true SSH, f the Coriolis parameter which depends on the latitude and g the gravity. The derivatives are computed in the swath coordinate system. Hence, x and y correspond to the along-track and the across-track direction. dx and dy correspond to the distance between two consecutive pixels, i.e., the spatial resolution of the data (in our case, 2 km). The gradients are computed on all available pixels. Please note that all the metrics computed on the geostrophic velocity are similar, whatever the studied component.

4. Results

In this section, the U-Net model is trained on simulated products from 2009. To calculate the different scores, the model is applied to inference independent data from 2010 over the North Atlantic Ocean. The scores of the classical methods were also calculated for simulated data from 2010. All these scores are compared to the scores for the noisy field and the ground truth.

4.1. Examples of Denoised Swaths

In this section, two examples of swaths denoised by our method and by the classical methods are compared. They are 600 km long. The SWOT pass illustrated in Figures 3–5

is a part of pass 496 of cycle 17. It is situated in the Gulf Stream near the American coast. Figures 3–5 illustrate the SSH, the geostrophic current and the vorticity, respectively. On all these figures, the left column shows—from top to bottom—the true field, the noisy field before denoising and the field after denoising by the U-Net. The errors between the noisy field or the denoised field and the ground truth appear in the right column.

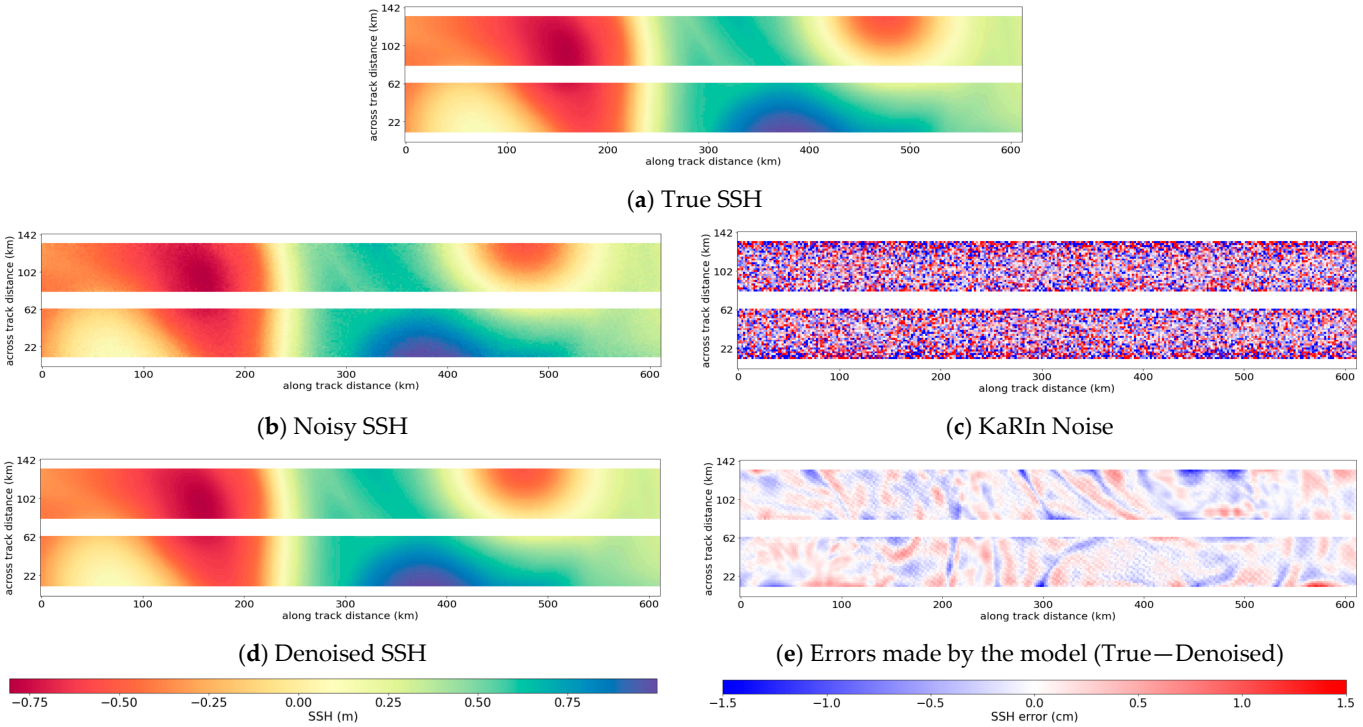

(**a**) True SSH

(**b**) Noisy SSH

(**c**) KaRIn Noise

(**d**) Denoised SSH

(**e**) Errors made by the model (True—Denoised)

Figure 3. (**a**) True SSH, (**b**) Noisy SSH, (**c**) the KaRIn noise applied, (**d**) SSH denoised by the U-Net and (**e**) its error (Pass 496, cycle 17 located in the Gulf Stream).

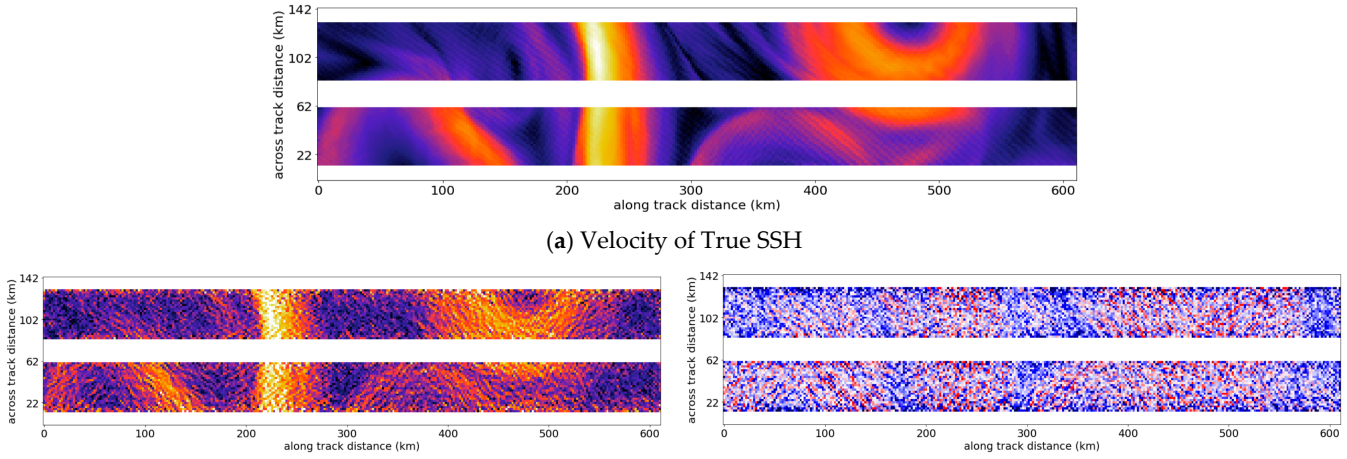

(**a**) Velocity of True SSH

(**b**) Velocity of Noisy SSH

(**c**) Errors made by the noisy field

Figure 4. *Cont.*

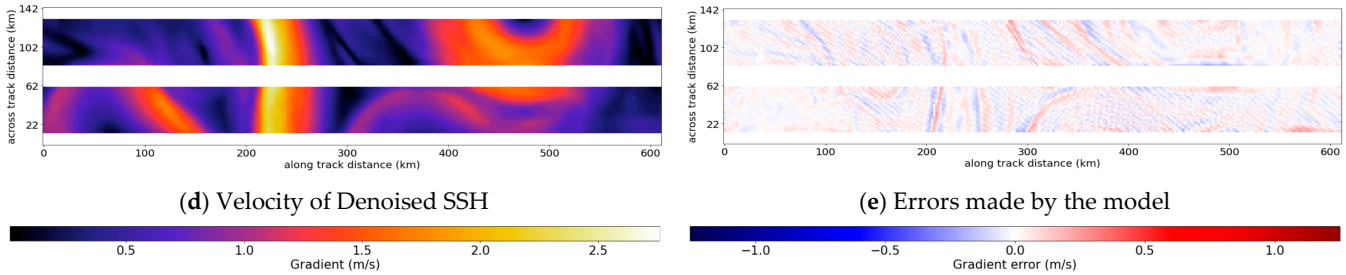

(**d**) Velocity of Denoised SSH (**e**) Errors made by the model

Figure 4. Velocity of (**a**) True SSH, (**b**) Noisy SSH, (**d**) SSH denoised by the U-Net and (**c**,**e**) their associated errors (Pass 496, cycle 17 located in the Gulf Stream).

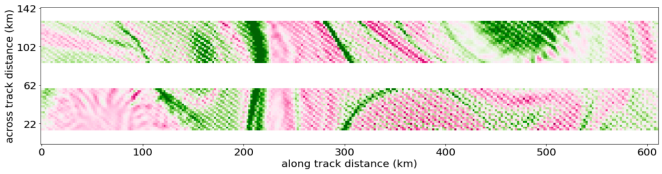

(**a**) Relative vorticity of true SSH

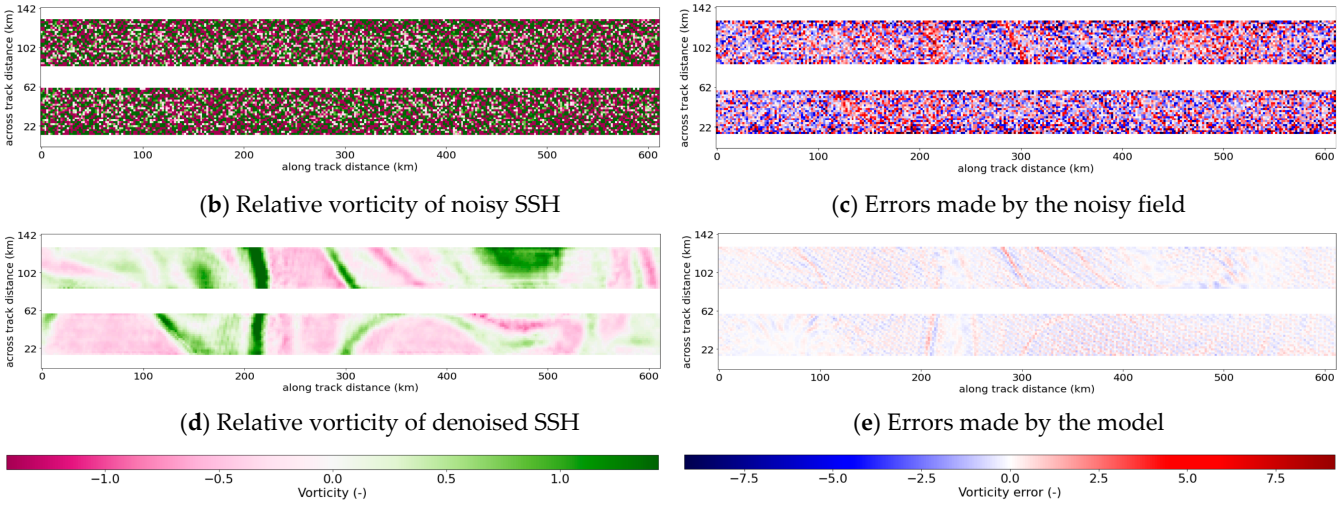

(**b**) Relative vorticity of noisy SSH (**c**) Errors made by the noisy field

(**d**) Relative vorticity of denoised SSH (**e**) Errors made by the model

Figure 5. Relative vorticity of (**a**) True SSH, (**b**) Noisy SSH, (**d**) SSH denoised by the U-Net and (**c**,**e**) their associated errors (Pass 496, cycle 17 located in the Gulf Stream).

The SSH is correctly restored in these two examples. The KaRIn noise is completely removed, and the maximum error made by the model is about 1 cm which is a reduction, by a factor of 3, in the KaRIn noise. However, the residual errors are correlated at wavelengths of around 10 km. The main structures of the geostrophic current calculated from the noisy SSH remain visible, but the associated errors reach 1 m/s. The U-Net reduces these errors by one order of magnitude and restores all the main structures. No information can be exploited from the vorticity field. Denoising is necessary, as shown in Chelton et al. [9]. The U-Net restores the main structures with the correct order of magnitude. Nevertheless, it is complicated to measure the error made due to some artefacts present on the true field. They are caused by the interpolation of the ocean model (which is on a regular grid) onto the swath grid in the SWOT simulator.

Figure 6 shows the same example, but the denoising is done by the Gomez filter which we have chosen as our reference for this study. The results are similar to the U-Net except for the vorticity field. A residual noise still remains. In low-variability areas, the main structures and the order of magnitude of the vorticity are not well restored (not shown). It is due to the choice of the value of λ_2. The swaths denoised by the other classical methods are presented in Appendix A.

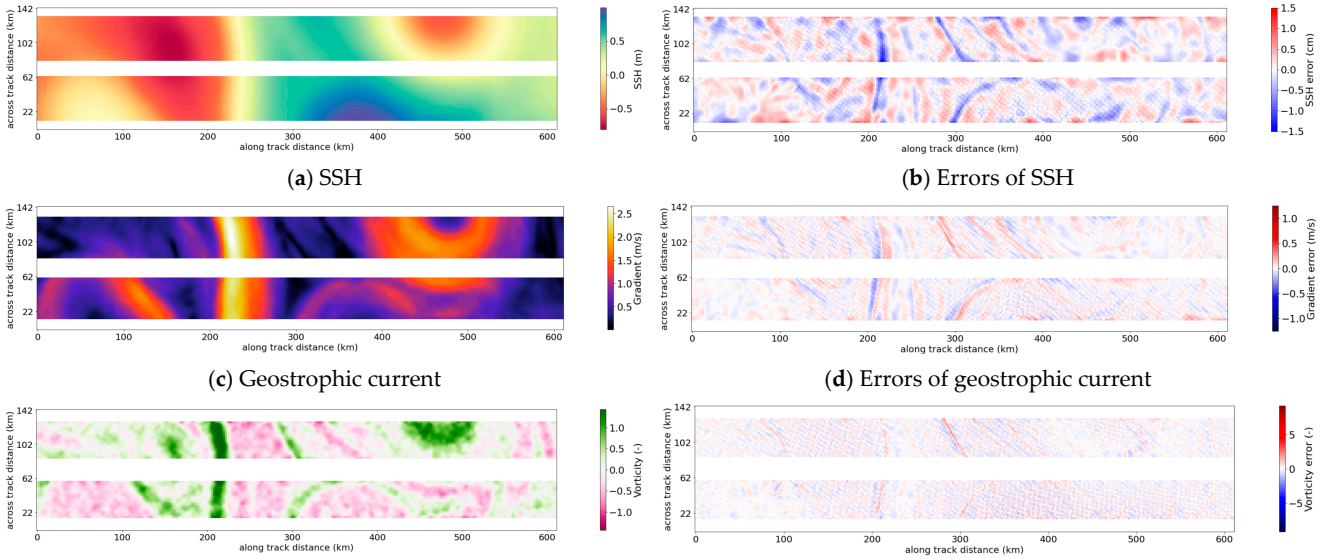

(**a**) SSH (**b**) Errors of SSH

(**c**) Geostrophic current (**d**) Errors of geostrophic current

(**e**) Relative vorticity (**f**) Errors of relative vorticity

Figure 6. SSH denoised by the Gomez filter and its derivatives in the left column. The associated errors are plotted in the right column. The ground truth of the SSH, the geostrophic current and the vorticity are in Figures 3–5, respectively (Pass 496, cycle 17 located in the Gulf Stream).

4.2. Statistics and Geographical Variations

The scores presented in Tables 1 and 2 are divided into two parts: the offshore scores which are calculated at more than 10 km from the coast and the coastal scores calculated at less than 10 km.

Table 1. Comparison of the offshore (>10 km to the coast) scores between the U-Net, the classical methods and with no filtering computing on the entire domain.

	U-Net	Gomez Filter	Median Filter	Lanczos Smoother	No Filter
$RMSE_{SSH}$ (cm)	0.19	0.24	0.31	0.44	1.27
KaRIn noise reduction	45 (16 dB)	28 (14 dB)	17 (12 dB)	8 (9 dB)	-
Mean of SSH residuals (mm)	<0.1	<0.1	<0.1	<0.1	0.
Variance of SSH residuals (cm^2)	0.04	0.07	0.10	0.20	1.63
$RMSE_V$ (m/s)	0.04	0.06	0.10	0.17	0.78
$RMSE_\zeta$ (-)	0.37	0.53	1.11	2.07	10.50

Table 2. Comparison of the coastal (<10 km to the coast) scores between the U-Net, the classical methods and with no filtering computing on the entire domain.

	U-Net	Gomez Filter	Median Filter	Lanczos Smoother	No Filter
$RMSE_{SSH}$ (cm)	0.32	0.84	0.87	0.81	1.21
KaRIn noise reduction	14 (11 dB)	2 (3 dB)	2 (3 dB)	2 (3 dB)	-
Mean of SSH residuals (cm)	0.01	0.04	0.03	0.01	<0.01
$RMSE_V$ (m/s)	0.09	0.14	0.14	0.21	0.69
$RMSE_\zeta$ (-)	0.67	0.80	1.10	1.88	7.92

First, the offshore scores show that all the methods remove a part—or almost all—of the noise, but the U-Net model performs best (Table 1). Indeed, its offshore RMSE of SSH is 0.19 cm, i.e., a reduction of 85% compared to noisy SSH. While the tenth of a millimeter difference is meaningless for real ocean data, it is nevertheless useful in a neural network context to evaluate the performance of the model. The KaRIn noise is reduced by a factor of

45, i.e., 16 dB. The offshore variance of SSH residuals is as much as 0.04 cm^2 and is reduced by 97%. The geostrophic current and the relative vorticity are restored: the associated RMSEs are reduced by 95%. The Gomez filter has the best performance among the classical methods. Nevertheless, it reduces the KaRIn noise by only 28 (14 dB). The offshore RMSE of the SSH and the variance of the SSH residuals are only reduced by 80% and by 94% compared with not using a filter. The RMSEs associated with the first two derivatives are also greater than for the neural network. For all the filters, the mean of the SSH residuals is lower than 1 mm which means that none of the filters used generate a bias.

However, the scores are worse near the coast and especially those of the classical methods (Table 2). The latter have an RMSE of SSH around 0.8 cm, low KaRIn noise reduction and a high RMSE of geostrophic current and relative vorticity. These methods are not suitable for edges and, in our case, for coasts. The U-Net gives the best coastal scores even though there are slightly greater RMSE and variance of SSH residuals than offshore. Actually, the RMSE and the variance of SSH residuals are only reduced by around 80%. It is also important to note that the interpolation of the ocean model on the swath grid creates some discontinuities, especially on jagged coastlines. The U-Net does not reproduce these discontinuities, which explains the poorer scores near the coast. As the Gomez filter has the best results among the classical filters, the comparisons will hereon be made between the U-Net and the Gomez filter only, treating the Gomez filter as the reference.

Figure 7 shows the variance of the SSH residuals, the RSMEs of the geostrophic current and relative vorticity of the U-Net filter and their comparison with the Gomez filter over North Atlantic Ocean. The RMSEs of the SSH are not shown because the results are similar to the variance of the SSH residuals.

The best results are obtained for areas of low variability. The improvement of the U-Net relative to the Gomez filter is higher than 50%. Near the coastline, the improvement approaches 100% due to the bad performance of the Gomez filter in these areas described above. However, in high-variability areas, the variance of SSH residuals is 0.1 cm^2, especially in the Gulf Stream, near the polar coast and near Florida. Near Florida, the bathymetry has important gradients which are neither observed elsewhere nor well-denoised by the model. Compared to the Gomez filter, the U-Net performs better: the improvement is higher than 25% except in the Gulf Stream where, for some pixels, the Gomez filter's performance is similar. This is because of the λ_2 parameter we used as a reference, which is optimal for high-variability regions [1]. Assuming that the λ_2 parameter of the Gomez filter is correctly optimized locally (for each longitude/latitude and wave-modulated SNR), it is likely that the Gomez filter would exhibit the same performance as the U-Net model. Nevertheless, an advantage of our U-Net filter is that a single training is needed to optimally denoise the entire domain, whereas a long and complex search for the optimal value of λ_2 at each pixel would be required with the Gomez filter.

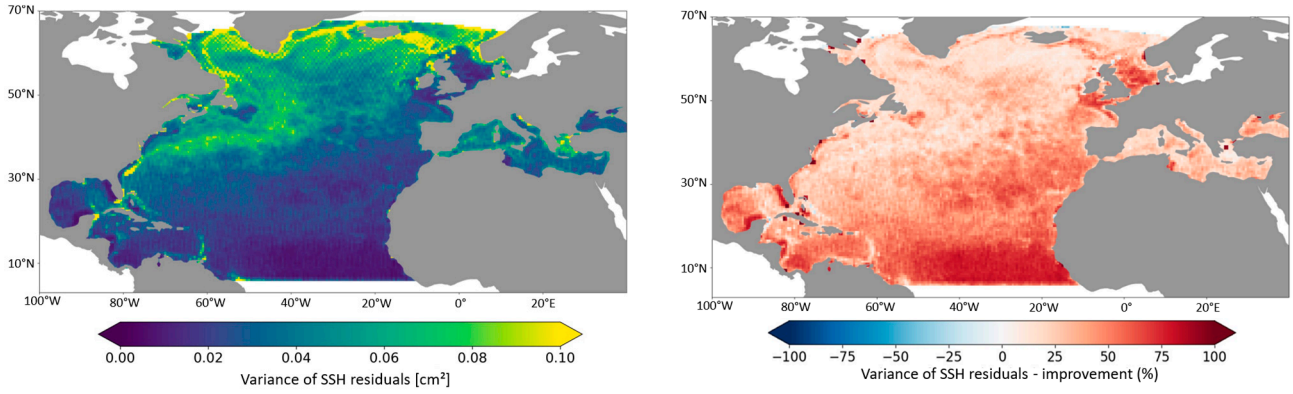

(**a**) (**b**)

Figure 7. *Cont.*

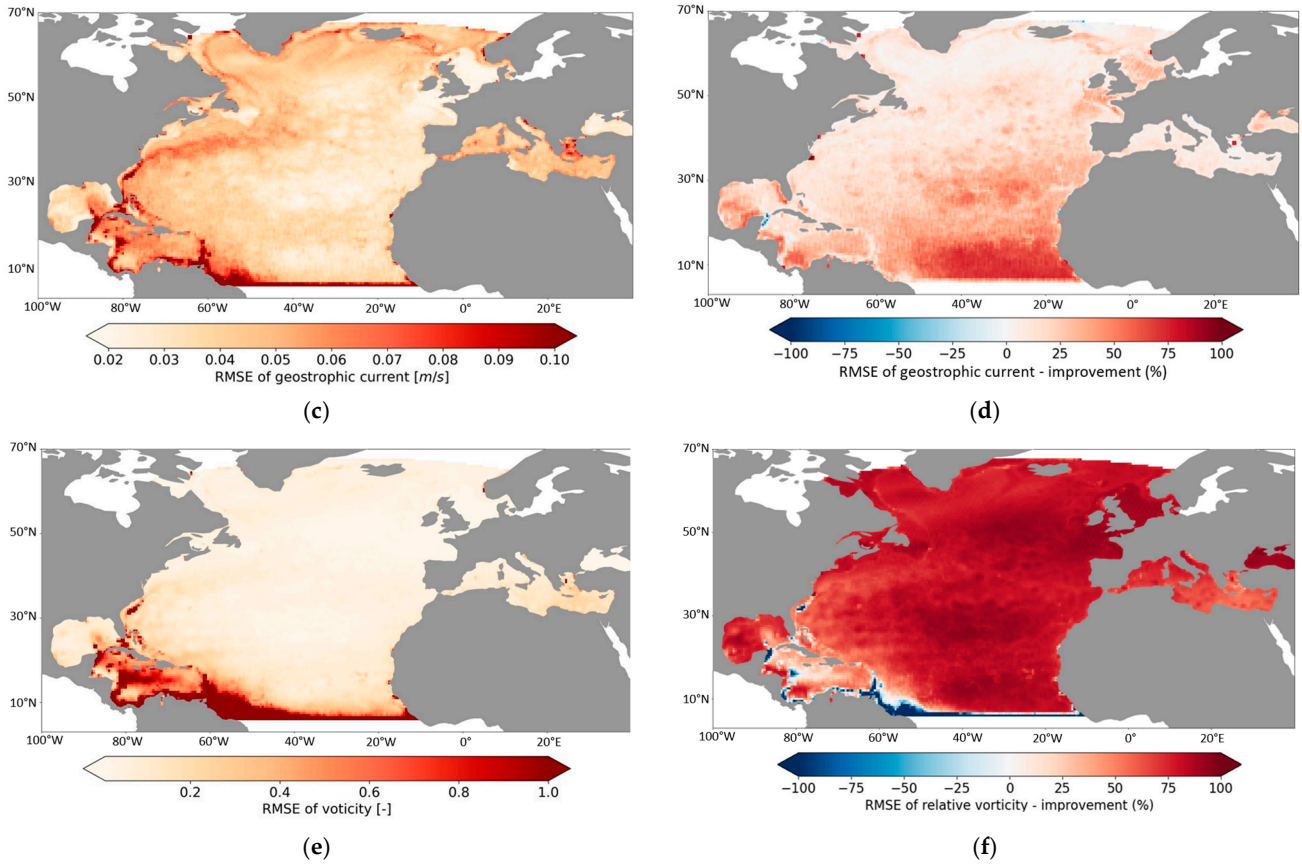

Figure 7. Variance of SSH residuals (**a**), RMSE of geostrophic current (**c**) and relative vorticity (**e**) and their comparison with the Gomez filter ((**b**), (**d**) and (**f**), respectively) over the North Atlantic Ocean. Red on the comparison map indicates that the U-Net performs better than the Gomez filter while blue indicates the opposite.

The SWOT tracks are visible in the variance of the SSH residuals (Figure 7a). It is caused by the definition of the KaRIn noise; more noise must be removed on the edges of the swaths (cf. Figure 1). Therefore, the error made by the denoising methods generally provides the highest error on the external edges of the swaths: 5.8 mm for the U-Net and 8.8 mm for the Gomez filter on average (cf. Figure 8). Whatever the along-track distance, the U-Net has the lowest error compared to the Gomez filter (reduction between 20 and 30%). In the middle of each half-swath, the error is the smallest: around 2.5 mm for the U-Net. On the inner edges, this error achieves 3.8 mm.

The RMSE of the geostrophic current is similar to the RMSE of the SSH. The RMSE of the relative vorticity is homogenous and is lower than 0.2 except in the Caribbean Sea and near the equator due to the definition of the Coriolis parameter. Except in these areas, the comparison with the Gomez filter shows an improvement superior to 75% over almost the entire domain as expected in the example of Figure 6. The U-Net better reproduces the second derivative.

4.3. Temporal Scores

Figure 9 shows the mean of SWH, the RMSE of SSH and the KaRIn noise reduction as a function of time. The computation is performed for each cycle and the date corresponding to the beginning of the cycle is also indicated. As previously, the U-Net model gives the best results whatever the season. Nevertheless, the performance is slightly correlated with the SWH. In winter, when the mean of SWH is the highest (2.75 m), the RMSE of SSH is equal to 0.21 cm. In summer, the SWH decreases to 1.50 m. In this case, the RMSE of SSH decreases to 0.16 cm. The KaRIn noise reduction is between 16 and 17 dB (i.e., a factor

of between 41 and 56). Thus, the waves have little impact on the results of the U-Net. Moreover, the observability after denoising is not affected by the season as with classical altimetry [35].

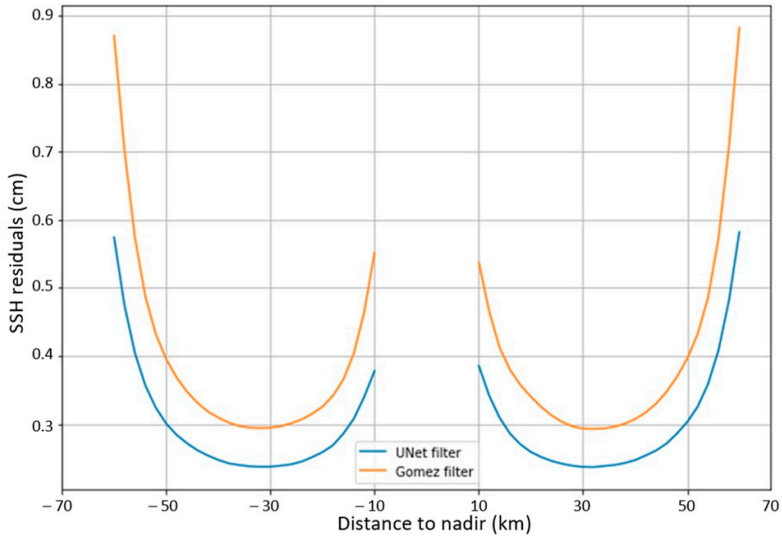

Figure 8. SSH residuals as function of the across-track distance from the nadir. Only data at more than 10 km from the coast are considered.

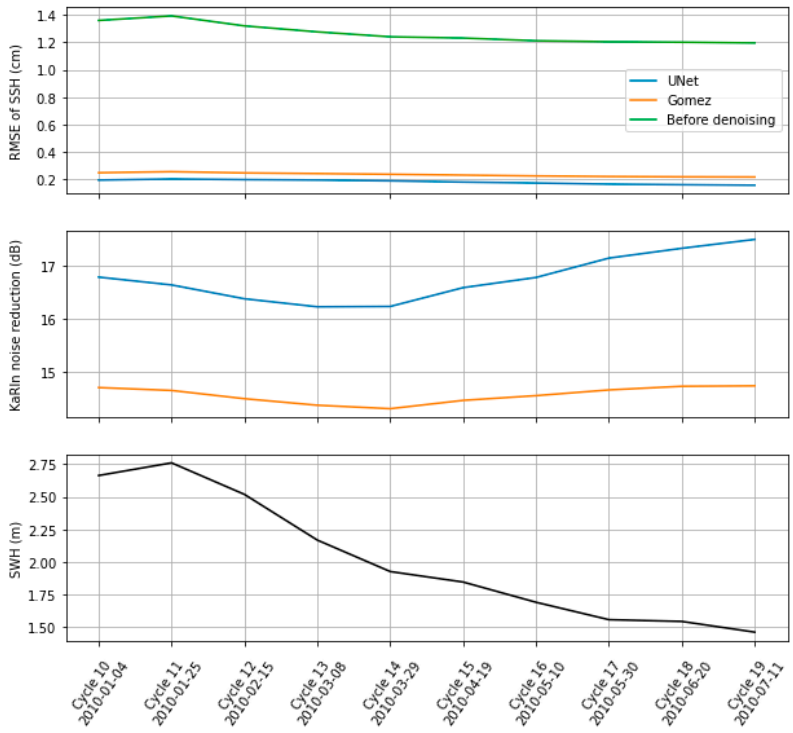

Figure 9. SWH, RMSE of SSH and KaRIn noise reduction as function of the cycle. The dates correspond to the beginning of each cycle. The results of the U-Net model are in blue and the Gomez filter in orange.

For the Gomez filter, the results are similar to the neural network model, but they are higher than before (around 20% whatever the period). Gomez-Navarro et al. [1] showed that a parametrization of λ_2 as a function of the season improved the results. On the other hand, performing several trainings of the U-Net depending on the season is not necessary even for improving a particular period.

4.4. Spectral Analysis

A spectral analysis was also performed on the denoised SSH over the eNATL60 domain. To do so, the PSDs were computed as follows. For each across-track distance, a PSD is computed with 512 km segments in the along-track direction. Then, all the obtained PSDs are averaged. The postprocessed data (i.e., reconstructed data after Section 3.1.4) are used.

Figure 10 shows the PSDs calculated for the SSH field, the geostrophic current field and the relative vorticity field in the along-track direction. The green spectrum represents the PSD of the noisy field. The noise impacts small structures in the SSH field and bigger structures in the geostrophic currents. The noise dominates the vorticity field. These results are consistent with the resolved scale given by SNR = 1. An SNR = 1 indicates that the PSD of the error made by the model intersects with the PSD of the ground truth, i.e., the spectrum of the denoised signal deviates from that of the ground truth (cf. Figure 11). In this case, there is as much noise/error as there is signal. The resolved scale defines the limit of observable structures in the ocean. If the spectrum of the denoised signal is lower than the spectrum of the reference, we assume that we cannot observe structures below 1000 km wavelength. The resolved scales associated with the noisy signal are 42 km for the SSH and 102 km for the geostrophic current, while no structures below a wavelength of 1000 km can be observed in the vorticity. These results are expected based on the examples in Section 4.1. Please note that the obtained spectra are similar to the ones shown in Appendix E of Chelton et al. [9].

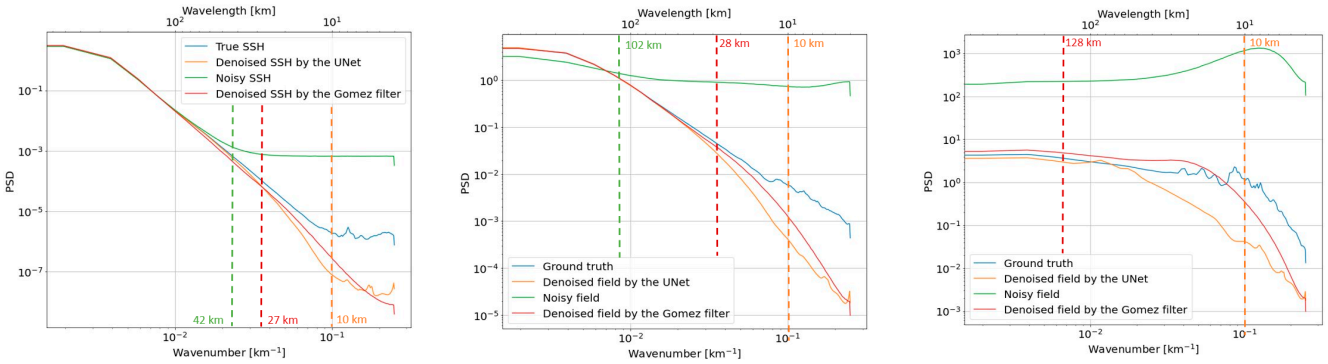

(**a**) (**b**) (**c**)

Figure 10. PSD of (**a**) true SSH (blue), noisy SSH (green), SSH denoised by the U-Net (orange) and SSH denoised by the Gomez filter (red) and the PSD of the associated first (**b**) and second (**c**) derivatives. The PSD is computed from the average of each along-track PSD over the eNATL60 domain and on the test dataset (year 2010). The vertical dotted lines are computed from Figure 11 and indicate the resolved wavelengths.

Turning now to the SSH spectra, the PSDs of the U-Net model and the Gomez filter correctly fit the PSD of the true field down to wavelengths of 20 km. Below 20 km, the spectrum of the U-Net deviates from the ground truth. Thus, too much energy is removed by the U-Net. On average, more energy is present in the small scales for the Gomez filter at these wavelengths. Nevertheless, this energy is not well localized as shown in Figure 11a. The resolved wavelengths reflect this; they are as small as 10 km for the U-Net and only 27 km for the Gomez filter, which is similar to the calculation by Gomez-Navarro et al. [1]. It is also important to note that, even at the longest wavelengths, the U-Net performs better than the Gomez filter.

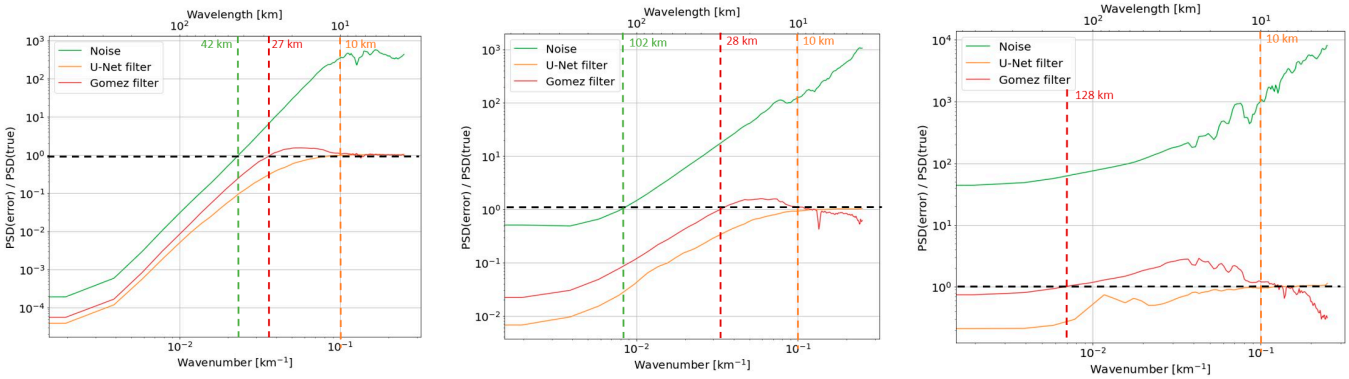

Figure 11. Examples of the method to determine the wavelength associated with the SNR = 1. It is the wavelength where the spectrum of the error over the spectrum of the ground truth crosses the value 1. If the spectrum is higher than 1 i.e., never crosses the value 1, the wavelength is setup as >1000 km like for the noisy field in (**c**). This example is calculated from the spectrums shown in Figure 10 i.e., (**a**) is computed from the SSH field, (**b**) from the first derivative and (**c**) from the second derivative over the eNATL60 domain and on the test dataset (year 2010). In green, the wavelengths of the noisy signal are, respectively, 42 km, 102 km and >1000 km. In red, the wavelengths of the SSH denoised by the Gomez filter are, respectively, 27 km, 28 km and 128 km. In orange, the wavelengths of the SSH denoised by the U-Net are 10 km.

The objective of SWOT is to observe oceanic structures below wavelengths of 15 km [6]. However, some assumptions are made. The first one is that SWH has a constant value of 2 m. This is very restrictive: the SWH varies between 0 and 8 m and the median SWH is around 3 m. Other assumptions are that a mean of SSH is calculated over 7.5 km $\times$ 7.5 km boxes and the 68% quantile spectrum is used. They are also restrictive and not used in this study: the 2 km resolution is maintained, and all the data is used to calculate the PSD. Despite breaking these assumptions, the U-Net model nonetheless produces the desired observability of 10 km. Among all the denoising methods tested, the U-Net is the only one to achieve this objective. Moreover, its spectrum adapts to all types of ocean variability.

The resolved scale is also calculated for the geostrophic current and for the relative vorticity. The U-Net provides the best results. Hence, structures down to 10 km can be observed in the geostrophic current field and in the relative vorticity field while, for the Gomez filter, these values are limited to 28 km and 128 km, respectively (cf. Figure 11b,c). Recall that, without denoising, these values are 102 km and higher than 1000 km, respectively. As with the spectra of the SSH fields, the spectrum of the velocity and vorticity field of the Gomez filter likewise seems to better fit the spectra of the ground truth (cf. Figure 11b,c), but the resolved wavelengths belie this. The energy of the Gomez filter is not as well-restored as the U-Net filter. As for the SSH field, the longest wavelengths are better restored by the U-Net.

5. Robustness of the Model

5.1. Method Applied to Test the Robustness

The data used in this study are simulated by the SWOT simulator. Even if it reproduces all the known errors, there is a possibility that the data are idealized: the KaRIn noise might not be exactly as described or might be transformed by several operations at level 1 or level 2 of the altimetric processing. The purpose of this section is to test the robustness of the U-Net model. Three scenarios are tested: a more intense KaRIn noise, a correlated KaRIn noise and a KaRIn noise impacted by the residuals of corrections of systemic errors. Another uncertainty is the description of the small structures present in the ocean. The ocean model used previously is regional and the phenomenon particular to other regions could impact the denoising.

To verify that our denoising model is effective in these various cases, the following method is applied:

1. A new dataset is generated either with the modified KaRIn noise or with another ocean model. As before, the year 2010 is used. The objective of this step is to simulate data which would have different properties from those of the SWOT simulator.
2. When the real SWOT data are available, the ground truth will not be accessible. Thus, the U-Net is applied in inference on this new dataset without training.
3. The scores are calculated from the output of the model.

This same method will also be applied to the first SWOT observations and the U-Net must provide good results for unseen data. Hereafter, the U-Net applied to the eNATL60 model and without modification of the noise is considered as our baseline.

5.2. Scenario: 50% More Noise

This scenario is the simplest. The objective is to ensure that more intense noise will be correctly removed. The KaRIn noise previously presented is multiplied by 1.5 to obtain the new KaRIn noise. Recall that the KaRIn noise depends on the SWH and the along-track distance.

The results are shown in Table 3. The scores are higher than the baseline: the RMSE of SSH is 0.24 cm, i.e., a noise reduction of 18 dB (factor 63). The RMSE of the geostrophic current is 0.05 m/s (i.e., 95% lower than before denoising) and the RMSE of the relative vorticity is 0.40 (i.e., 97% lower). The variance of SSH residuals is equal to 0.062 cm^2 (i.e., 98% lower). The high-variability areas have the highest variance. Nevertheless, these values remain very low even if the KaRIn noise is higher. Furthermore, the resolved scale remains unchanged regardless of the field. Thus, noisier data are correctly denoised by the U-Net and the elevated noise has little effect on the result.

Table 3. Offshore (>10 km to the coast) scores of four scenarios. The scores before denoising are also indicated for each scenario. The scores are calculated on the entire domain.

	50% More Noise		**Correlated Noise**		**XCal Residuals**		**Glorys Data**	
	No Filter	**U-Net**	**No Filter**	**U-Net**	**No Filter**	**U-Net**	**No Filter**	**U-Net**
RMSE (cm)	1.91	0.24	0.97	0.20	1.65	0.91	1.33	0.13
Mean of SSH residuals (mm)	<0.1	<0.1	<0.1	<0.1	<0.1	<0.1	<0.1	<0.1
Variance of SSH residuals (cm^2)	3.67	0.06	0.94	0.04	2.78	0.81	1.83	0.02
$RMSE_V$ (m/s)	1.22	0.05	0.60	0.04	0.85	0.13	0.84 *	0.02 *
$RMSE_\zeta$ (-)	15.75	0.40	8.35	0.39	10.82	0.58	11.72 *	0.12 *
Resolved scale of SSH (km)	56	10	42	22	46	10	57	22
Resolved scale of V (km)	>1000	10	102	20	171	10	>1000 *	21 *
Resolved scale of ζ (km)	>1000	10	>1000	85	>1000	>1000	>1000 *	46 *

* Values calculated on the entire domain of Glorys except between 6°N and 6°S.

5.3. Scenario: Correlated Noise

5.3.1. Generation of the Correlated Noise

SWOT consists of two antennas, each composed of one KaRIn instrument. Each KaRIn instrument is formed by nine beams that are reflected off the surface with an approximate angle of 0.1°. The approximate resolution is 250 m. For most oceanographic products, a two-dimensional smoothing will be applied to the nine-beam product. The final product will have a resolution of 2 km on an Earth-fixed grid as in this study. For more details, please refer to Appendix B of Chelton et al. [9]. As of yet, the filter of this smoothing has not been chosen by the SWOT community [36]. At the SWOT Science Team meeting in June 2022, Molero et al. [37] presented a method using a Hamming filter with a half-power filter cutoff wavelength of 4 km. This filter generates a two-dimensional correlated noise even with decorrelated waves as input.

To generate this correlated noise, a method similar to the one implemented in the SWOT simulator is used. The Gaussian zero-centered distribution is replaced by a random realization based on the Hamming spectrum with a variance of 1 cm^2. In this scenario, the

correlated noise also depends on the SWH to compare it to the baseline. The resolved scale of the SSH field remains the same i.e., 42 km. Figure 12 shows the PSD of the SSH with a correlated noise, with the white noise and the ground truth. The correlated noise impacts wavelengths under 20 km; the longest wavelengths remain unchanged.

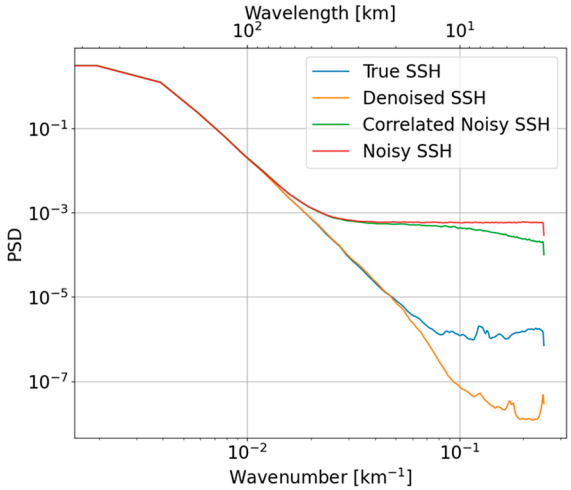

Figure 12. PSD of true SSH (in blue), noisy SSH (red), SSH with correlated noise (green) and denoising of the latter by U-Net (orange).

5.3.2. Results

The results of this second scenario are summarized in Table 3. As in the previous scenario, the scores are slightly higher than the baseline. The KaRIn noise reduction is 14 dB (factor 23). The RMSEs decreased by around 90% after the denoising by the U-Net. The variance of SSH residuals approaches 0.041 cm^2 (i.e., 96% lower than before denoising). The mean of the SSH residuals is higher than the U-Net with the classic KaRIn noise but remains lower than 0.1 mm. The resolved scale of the SSH field does not meet the requirement of SWOT observability [6]. The resolved scales computed for the first two derivatives are also higher than the baseline, especially for the relative vorticity. Note that if the U-Net is retrained on the correlated noise these values can be reduced and meet the requirements; the resolved scales of SSH field, geostrophic current, and relative vorticity become 12 km, 10 km, and 85 km, respectively. However, the model without retraining can be used for new types of noise even in inference as a first guess.

5.4. Scenario: XCal Residuals

5.4.1. XCal Residuals

Before filtering the KaRIn noise, another step consists of eliminating the systematic instrumental errors, i.e., the roll, the phase, the baseline dilation, the timing and the orbital errors. This step is called crossover calibration (referred to as XCal hereafter). Dibarboure et al. [38] developed a method based on the inversion of the swath crossings. In the fast-sampling phase, the time step between two swaths is small at crossovers. Therefore, the ocean state remains almost unchanged and the values of the systematic errors can be deduced from these crossings. By interpolating the result in time and space, the XCal correction is obtained by applying the standard crossover algorithm to the ground segment. However, some residuals of these errors will remain and will be present in the following steps such as the filtering. The objective of this scenario is to test our model on KaRIn noise to which XCal residuals are added. One of the objectives of the denoising is to only remove the small scales (i.e., KaRIn noise and high-frequency residuals), and to not modify the large scales such as bias or large, structured errors. To do so, a Lanczos filter is used to separate the small from the large scales with a cutoff of 50 km. Only the small scales are added to the KaRIn noise. The result of this processing is that vertical stripes appear and a

dipole is present between the borders and the center of a half-swath; these are shown in Figure 13.

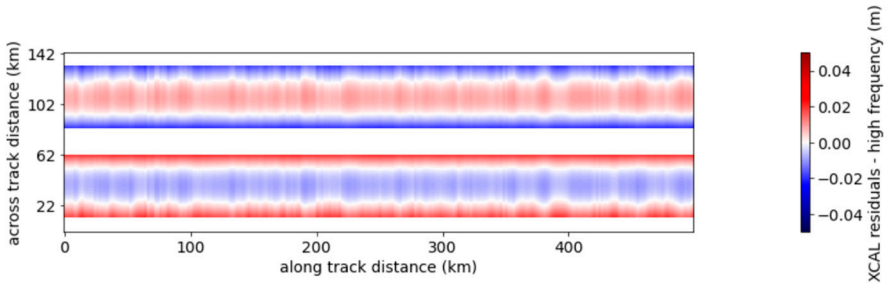

Figure 13. Example of high-frequency XCal residuals. This field is added to the KaRIn noise.

5.4.2. Results

The vertical stripes and the KaRIn noise are correctly removed by the model in this case. Nevertheless, the dipole between the edges and the center of a half-swath is still present after the denoising. That explains the high values of the RMSEs and the variance of SSH residuals (cf. Table 3). The RMSEs of the SSH field, of the geostrophic current and of the relative vorticity are 0.91 cm (i.e., 1.8 times lower than before denoising), 0.13 m/s (i.e., 6.5 times lower) and 0.58 (i.e., 1.5 times higher), respectively. The variance of the SSH residuals is equal to 0.83 cm^2 (i.e., 18 times lower). Bias is locally present in the map of the mean of SSH residuals (not shown) but is absent from the global scores. It is caused by the dipole. The resolved scales of the SSH and the geostrophic current fields approach 10 km like the baseline. However, these values are impacted by the computational method of the PSD.

Thus, the U-Net can remove small structures (the vertical stripes) even if they are not learnt by the model and do not look like noise. However, structures with the longest wavelengths are not removed by the model. Thus, the U-Net does not remove structures that could, because of their scale, potentially be oceanic information. To correct the dipole error, a more complex XCal was developed by Dibarboure et al. [38].

5.5. Generalization of the Model

5.5.1. Glorys Ocean Model

The Glorys 1/12° ocean model developed by Lellouche et al. [39] is used in the operational Copernicus Marine Service operated by Mercator Ocean International. This global model does not resolve the small-mesoscale, unlike eNATL60, but only large- and medium-mesoscale. The objective of this scenario is to test the robustness of the U-Net on other oceans where different phenomena and structures can be observed than in the North Atlantic Ocean. In practice, the daily SSH of Glorys is interpolated on a swath grid as before. The KaRIn noise generated by the SWOT simulator remains unchanged. The SWOT data are available from 12 April 2014 to 30 December 2015. The last six months are used for the benchmark.

5.5.2. Results

The results of this scenario are presented in Table 3. The RMSE of SSH is equal to 0.13 cm and the variance of SSH residuals is 0.022 cm^2. The RMSEs of the first and the second derivatives are 0.02 m/s and 0.12, respectively. Please note that these last two scores are calculated for the Glorys domain except between 6°N and 6°S. They are better than those calculated from the eNATL60 data. The smallest structures resolved by the Glorys ocean model are higher than the ones in the eNATL60 model and, therefore, the noise is more easily separable from the smallest oceanic structures. Figure 14 shows the variance of SSH residuals. All the oceans are correctly denoised. As before, the low-variability areas have the best scores and the error made by the U-Net model is less than 5 mm. The high-variability areas, especially in the Gulf Stream, have a higher variance of SSH

residuals and, thus, a higher error, which remains less than 2 cm. Other currents such as the Kuroshio Current or the Antarctic Circumpolar Current are well denoised even if they are not learnt by the model. The coasts are correctly denoised except near the Scandinavian coasts because of the jagged coastlines. The error is still less than 2 cm, however.

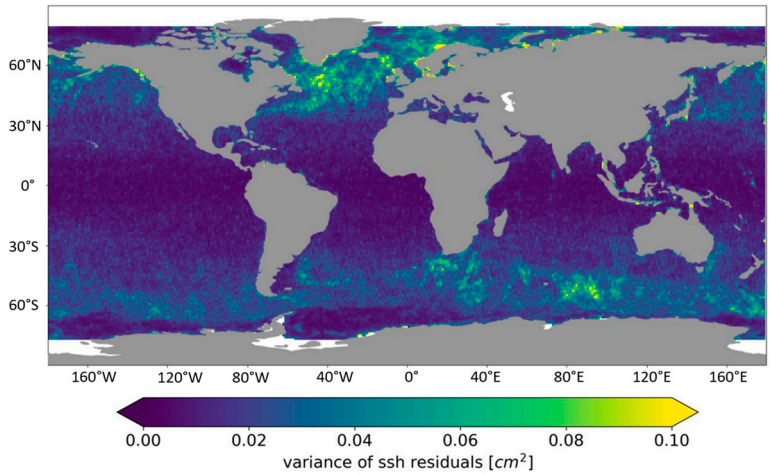

Figure 14. Variance of the SSH residuals (Glorys scenario).

A spectral analysis was also performed for the Glorys domain. The PSDs of U-Net calculated for the SSH field, the geostrophic current field and the relative vorticity field are similar to the PSD of the ground truth (cf. Figure 15). The resolved wavelength of the SSH field is 22 km which is higher than the wavelength required by the SWOT Requirement [6]. It is explained by the size of the smallest structures resolved by Glorys. The denoised spectrum is close to the ground truth up to the resolved scale. This is because the Glorys ocean model does not resolve small scales and does not model turbulence as well as the eNATL60 ocean model due to their respective resolutions. The resolved wavelengths of the geostrophic current field and of the relative vorticity field are 21 km and 46 km, respectively. Hence, the U-Net model trained on an ocean model and on a specific area can be used on another ocean model and on another area.

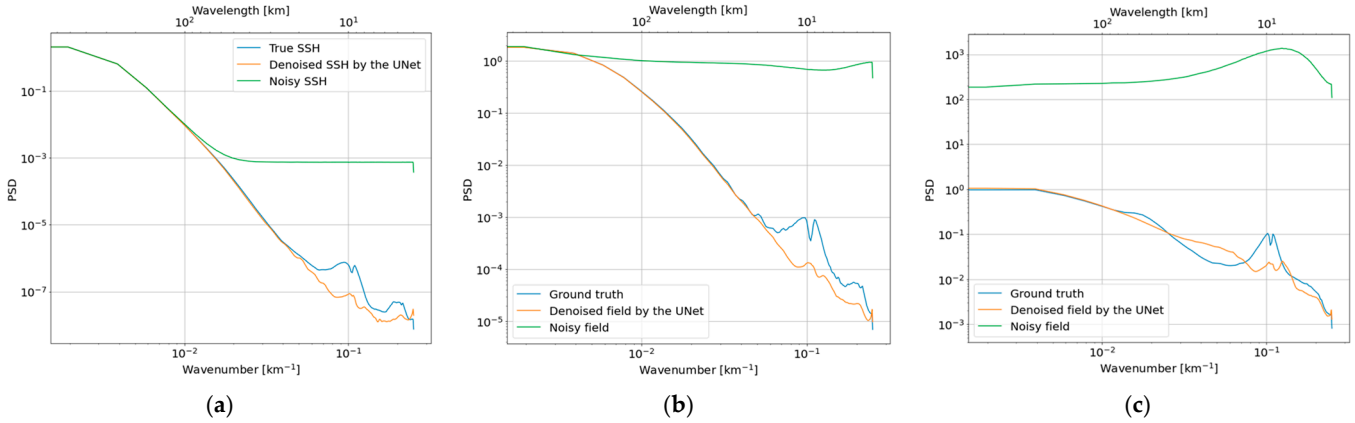

Figure 15. PSD of SSH (**a**), geostrophic current (**b**) and relative vorticity (**c**) of the Glorys model. The ground truth is in blue, the field denoised by the U-Net is in orange and the noisy field is in green. The PSD is computed from the average of each along-track PSD over the Glorys domain and on the test dataset (from the 1st July 2015 to the 31st December 2015).

5.6. Discussion about the Robustness of the U-Net

Different scenarios were proposed to test the robustness of the U-Net. It adapted well to all regions and to other types of oceanic data. This is principally due to the diversity of

the training dataset which is composed of high-variability areas (such as the Gulf Stream), low-variability areas and coastal areas. The chosen period—which covers the summer, the fall and the beginning of the winter—trains it on a large number of different situations. The training dataset is a crucial component of the neural network model (as, too, is the preprocessing of data). Some trainings are realized on Gulf Stream area and do not produce satisfying results by applying it outside its training area, particularly in low-variability areas (up to −75% of the variance of SSH residuals compared to the results presented in Section 4).

The data augmentation also has a major effect on the robustness of the model. The neural network interprets this data augmentation as a set of new situations to model. Thus, it improves the performance of the eNATL60 test dataset by reducing the variance of SSH residuals by 5%. However, for the Glorys test dataset (without retraining), the reduction is greater (around 10% for the variance of SSH residuals).

Another key to its robustness towards other types of noise is that KaRIn noise varies with SWH. This allows training on a wide range of noise intensities, which is why the model provides satisfying results on more or less noisy data. However, the U-Net removes some oceanic structures with less noisy data (not shown). The structure of the U-Net also has an important role because it allows it to comprehend other types of noise, such as correlated noise. Indeed, in the example presented above which featured noise correlated at 4 km, the first operations of the U-Net (pooling) allowed it to decrease the dimension of the input data to bring it back to a case already known to the model.

However, the U-Net does not provide the best possible results for the different scenarios. A retraining is possible for each scenario proposed. It would improve the results presented because it would better fit the properties (see Appendix B for an example on the first scenario) and/or avoid removing oceanic structures. The retraining of all the scenarios is not done in this study; the retraining and the retuning of the U-Net will be performed for real KaRIn noise if the properties are not as expected to improve the efficiency of the denoising. In addition, retraining can also be used for regional studies. Thus, a retraining on the region of interest improves results, but they are degraded outside of this area.

6. Interpretation of De-Noised Ocean Dynamics

The technical findings of this study are part of a wider set of open challenges concerning the understanding of fine-scale ocean dynamics and their role in global ocean variability and climate. They will largely help the scientific community in the use and interpretation of future SWOT data. While current satellite altimetry observations are limited to resolving scales down to around 80 to 70 km at best [35], a significant share of the observed SSH variability originates at scales shorter than 50 km [40–42]. Additionally, small-scale processes have been shown to significantly contribute to horizontal fluxes of energy, momentum and tracers such as heat, carbon and nutrients [2,43]. Our understanding of these dynamics relies today on either high-resolution simulations or scarce in situ measurements. Both highlight the important role of mesoscale and submesoscale processes when interpreting diagnostics such as vorticity or surface velocities, but these findings need to be supported by global and frequent observations.

In this context, SWOT will produce unprecedented two-dimensional swath images of ocean surface topography to characterize the ocean mesoscale and submesoscale from larger scales down to around 15 km wavelength [2,3]. However, SWOT instrumental noise will limit the actual observability, thereby limiting the physical interpretation of the observations. Figures 4b and 5b demonstrate this point perfectly. Clearly, the velocity and vorticity fields will not be directly detectable from a SWOT noisy SSH field. The effect of the speckles present in the SSH field is enhanced in higher-order diagnostics, such as the geostrophic velocity field (order 1) and vorticity (order 2). In line with the findings of Chelton et al. [9], the large-scale and most energetic traits of the velocity field are still detectable from the noisy data, but many of the smaller-scale noise features could be misinterpreted as a real signal without previous knowledge of the error-free field (Figure 4a).

On the other hand, it is impossible to extrapolate the vorticity features computed from the noisy SSH at any scale. Solving scales down to 10 km has powerful physical implications, particularly in regions of small Rossby radius where a significant part of the mesoscale variability is still out of reach for conventional altimetry (e.g., in the Antarctic Circumpolar Current or ACC). The potential of this study is also evident when comparing the result of resolving 10 km scales globally and independent of the sea state to previous findings by Wang et al. [34]. They show that SWOT observability degrades to 30–45 km at high latitudes due to the SWH-induced instrument noise, especially in the ACC, and that this correlation causes a large seasonality.

To date, the overall eddy field—a reservoir of kinetic energy—still cannot be balanced in terms of inputs and outputs [44]. In maps reconstructed with current altimetry observations, large-scale eddies seem to spontaneously emerge before suddenly disappearing, but the processes that generate and kill them, and the quantitative energy transfers between these and smaller scales, remain obscured [2,45]. The resolution of smaller scales will enlighten the energy transfers between the different scales, improving our understanding of the generation and extinction of mesoscale structures and energy sources and sinks in the ocean. The contribution of such high-resolution two-dimensional information is not limited to the ocean surface, but the estimates of the surface pressure could also be used as a proxy to gain insight into the vertical structure of the ocean velocity field [46] and its role as a heat pump for the climate system [43,47]. Having access to two-dimensional, de-noised 15 to 150 km SSH, geostrophic velocities and vorticity fields globally opens many new paths and research questions on the role of small scales in ocean dynamics and their interaction with larger-scale circulation—in terms of the energy budget, mixing and dissipation—and on the vertical exchanges between the surface and deeper layers.

7. Conclusions

Despite its unprecedented precision, SWOT's Ka-band radar interferometer yields an ocean surface topography with a non-zero random noise. A noise-reduction technique is therefore beneficial: to improve the observability of small to sub-mesoscales and to better investigate ocean dynamics at the smallest scales with SSH derivatives. Arguably the main challenge in noise mitigation is that KaRIn's noise is modulated by waves. Because ocean dynamics are also geographically and temporally variable, the signal-to-noise ratio is as well. The statistical properties of the denoising problem are local in space and time. In Section 4, the method above is compared to more "traditional" filtering/smoothing methods, i.e., algorithms that do not use a neural network.

In this paper, we used a U-Net neural network to reduce the noise on simulated SWOT products. The network was trained and tested on independent datasets based on a realistic ocean model and the SWOT scientific simulator. Its de-noising performance was compared to three other filters: the filter from Gomez-Navarro et al. (2020) which was the state of the art for SWOT at the beginning of our study, as well as a more classical median filter and a Lanczos filter. Our U-Net model yields better results for all the evaluation metrics: 2 mm root mean square error, sub-millimetric mean error, a factor of 44 of variance reduction (16 dB), and an accurate power spectral density down to 10–20 km wavelengths.

Furthermore, the U-Net is the only filter that correctly captures the first and second derivatives for smaller ocean scales of interest for SWOT. Only one training is necessary to denoise the entire North Atlantic Ocean for each period of the year, and the model can be used for the global ocean. More importantly, the training can be done without any kind of regional/local tuning: contrary to other filters, it is not necessary to find statistically optimal parameters. Another practical advantage of the model is that, once trained, the denoising is extremely fast.

Because SWOT's flight data might differ from our simulated training dataset, we tested the robustness of the U-Net noise reduction in various scenarios where the inference is done on a dataset with very different properties than the training data. For all the scenarios the U-Net residuals are still functioning as expected: random noise is mitigated

very efficiently, albeit with slightly larger residuals. Moreover, in the presence of a complex correlated error (namely, XCal residuals from Dibarboure et al. [38]), our filter is even able to capture a small fraction of said error (the highest frequencies) even if it is beyond the original denoising objectives. In addition to this solid capability, our simulations show that the U-Net can be retrained to better fit the actual noise properties that will be measured on flight data during the calibration/validation phase. To that extent, when SWOT's flight data become available, we expect that the U-Net algorithm will provide a solid and upgradable noise-mitigation baseline.

The literature about denoising is also very rich [26,48]; it is possible, if not likely, that other algorithms might yield even better results on SWOT products. To that extent, we set up a standardized open-source data challenge for any group who wants to test their algorithms in a controlled environment with predefined metrics (available at https://github.com/ocean-data-challenges/2022a_SWOT_karin_error_filtering, accessed on 7 April 2023).

Author Contributions: Conceptualization and methodology, A.T. and Y.F.; software, A.T. and B.C.; input data and analysis, A.T., M.B. and E.C.; writing, A.T. and E.C.; supervision, Y.F. and G.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research was financially and technically supported by the French Space Agency (CNES) (DUACS R&D 2022: 20221).

Acknowledgments: The authors would like to thank all the people including the members of SWOT project for their contributions in testing and/or improving the denoising model.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

This appendix presents the examples of denoising by the classical methods to illustrate the good performance of the U-Net filter. Hence, the Figures A1 and A2 represent the results of the median and the Lanczos smoothers, respectively. They can be compared to Figures 3–5. It is a part of the pass 496, cycle 17 located in the Gulf Stream. The figures are organized as follows. In the left column, the SSH, the geostrophic current and the vorticity are plotted. In the right column, the associated errors are plotted.

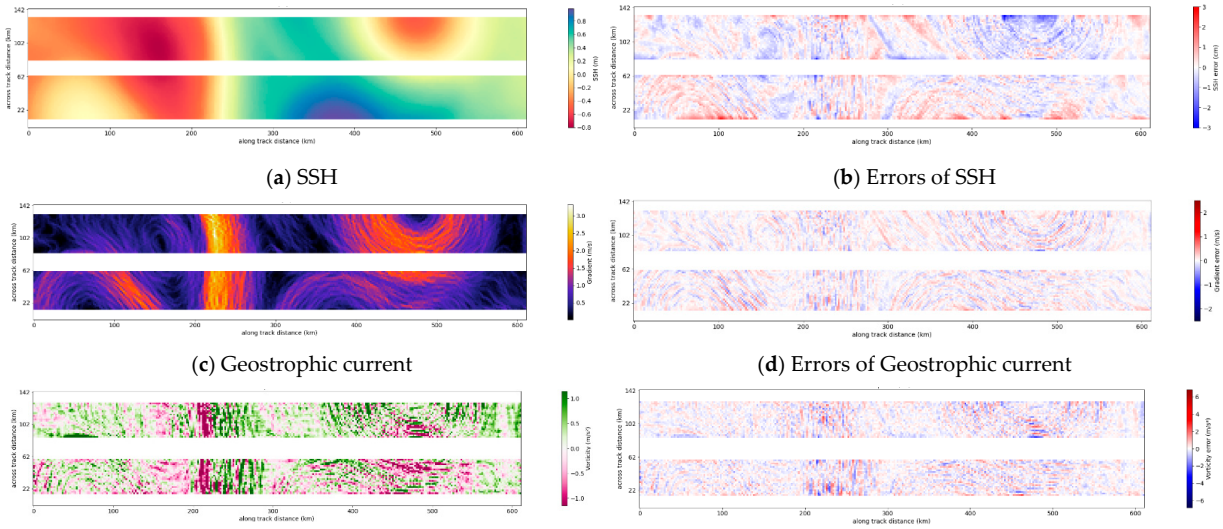

(**a**) SSH (**b**) Errors of SSH

(**c**) Geostrophic current (**d**) Errors of Geostrophic current

(**e**) Vorticity (**f**) Errors of Vorticity

Figure A1. SSH denoised by the Median smoother and its derivatives in the left column. The associated errors are plotted in the right columns. The ground truth of the SSH, the geostrophic current and the vorticity are in Figures 3–5, respectively (Pass 496, cycle 17 located in the Gulf Stream).

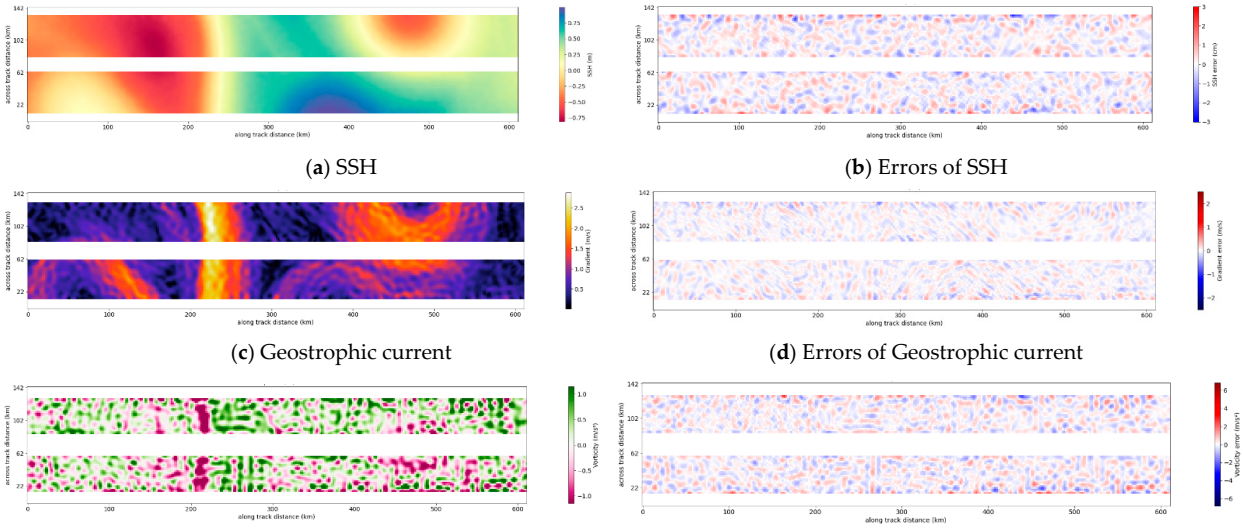

(**a**) SSH (**b**) Errors of SSH

(**c**) Geostrophic current (**d**) Errors of Geostrophic current

(**e**) Vorticity (**f**) Errors of Vorticity

Figure A2. SSH denoised by the Lanczos smoother and its derivatives in the left column. The associated errors are plotted in the right columns. The ground truth of the SSH, the geostrophic current and the vorticity are in Figures 3–5, respectively (Pass 496, cycle 17 located in the Gulf Stream).

Appendix B

This appendix presents the impacts of retraining the U-Net model. The example below shows a comparison before and after retraining on the scenario with 50% more noise presented in Section 5.2. The retraining is done on 2009 data. The hyperparameters can be modified to further improve the results but, in this example, the same ones are kept. The results are computed for 2010.

The scores are summarized in Table A1. All the scores with retraining are less than or equal to those without retraining. Actually, the RMSE of SSH are 17% lower. The KaRIn noise reduction increases by 1 dB with retraining and is equal to 19 dB. The variance of the SSH residuals, the RMSE of the geostrophic current and the vorticity decrease by 30%, 20% and 12%, respectively. The resolved scale remains unchanged. Therefore, the retraining improves the results and the model better fits the noise properties. Please note that a retuning of the hyperparameters of the U-Net will also improve the performance.

Table A1. Offshore (>10 km to the coast) scores with and without retraining for the scenario with 50% more noise.

	No Filter	Without Retraining	With Retraining
RMSE (cm)	1.91	0.24	0.20
Mean of SSH residuals (mm)	<0.1	<0.1	<0.1
Variance of SSH residuals (cm^2)	3.67	0.06	0.04
$RMSE_V$ (m/s)	1.22	0.05	0.04
$RMSE_\zeta$ (-)	15.75	0.40	0.35
Resolved scale of SSH (km)	56	10	10
Resolved scale of V (km)	>1000	10	10
Resolved scale of ζ (km)	>1000	10	10

References

1. Gómez-Navarro, L.; Cosme, E.; Le Sommer, J.; Papadakis, N.; Pascual, A. Development of an Image De-Noising Method in Preparation for the Surface Water and Ocean Topography Satellite Mission. *Remote Sens.* **2020**, *12*, 734. [CrossRef]
2. Morrow, R.; Fu, L.-L.; Ardhuin, F.; Benkiran, M.; Chapron, B.; Cosme, E.; D'ovidio, F.; Farrar, J.T.; Gille, S.T.; Lapeyre, G.; et al. Global Observations of Fine-Scale Ocean Surface Topography with the Surface Water and Ocean Topography (SWOT) Mission. *Front. Mar. Sci.* **2019**, *6*, 232. [CrossRef]

3. Fu, L.-L.; Alsdorf, D.; Morrow, R.; Rodriguez, E.; Mognard, N. *SWOT: The Surface Water and Ocean Topography Mission*; JPL Publication 12-05; JPL: Pasadena, CA, USA, 2012; p. 228.
4. Durand, M.; Fu, L.-L.; Lettenmaier, D.P.; Alsdorf, D.E.; Rodriguez, E.; Esteban-Fernandez, D. The Surface Water and Ocean Topography Mission: Observing Terrestrial Surface Water and Oceanic Submesoscale Eddies. *Proc. IEEE* **2010**, *98*, 766–779. [CrossRef]
5. Fu, L.-L.; Ubelmann, C. On the Transition from Profile Altimeter to Swath Altimeter for Observing Global Ocean Surface Topography. *J. Atmos. Ocean Technol.* **2014**, *31*, 560–568. [CrossRef]
6. Fu, L.-L. *SWOT Science Requirements Document*; JPL Publication D-61923; JPL: Pasadena, CA, USA, 2018; p. 29.
7. Esteban-Fernandez, D. SWOT Project Mission Performance and Error Budget 2014. Available online: https://swot.jpl.nasa.gov/system/documents/files/2178_2178_SWOT_D-79084_v10Y_FINAL_REVA__06082017.pdf (accessed on 7 April 2023).
8. Gaultier, L.; Ubelmann, C.; Fu, L.-L. The Challenge of Using Future SWOT Data for Oceanic Field Reconstruction. *J. Atmos. Ocean. Technol.* **2016**, *33*, 119–126. [CrossRef]
9. Chelton, D.B.; Samelson, R.M.; Farrar, J.T. The Effects of Uncorrelated Measurement Noise on SWOT Estimates of Sea-Surface Height, Velocity and Vorticity. *J. Atmos. Ocean. Technol.* **2022**, *39*, 72. [CrossRef]
10. Gómez-Navarro, L.; Fablet, R.; Mason, E.; Pascual, A.; Mourre, B.; Cosme, E.; Le Sommer, J. SWOT Spatial Scales in the Western Mediterranean Sea Derived from Pseudo-Observations and an Ad Hoc Filtering. *Remote Sens.* **2018**, *10*, 599. [CrossRef]
11. Fan, L.; Zhang, F.; Fan, H.; Zhang, C. Brief review of image denoising techniques. *Vis. Comput. Ind. Biomed. Art* **2019**, *2*, 1–12. [CrossRef]
12. Ilesanmi, A.E.; Ilesanmi, T.O. Methods for image denoising using convolutional neural network: A review. *Complex Intell. Syst.* **2021**, *7*, 2179–2198. [CrossRef]
13. Tian, C.; Fei, L.; Zheng, W.; Xu, Y.; Zuo, W.; Lin, C.-W. Deep learning on image denoising: An overview. *Neural Netw.* **2020**, *131*, 251–275. [CrossRef]
14. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. *Commun. ACM* **2017**, *60*, 84–90. [CrossRef]
15. Jifara, W.; Jiang, F.; Rho, S.; Cheng, M.; Liu, S. Medical image denoising using convolutional neural network: A residual learning approach. *J. Supercomput.* **2017**, *75*, 704–718. [CrossRef]
16. Kovacs, A.; Bukki, T.; Legradi, G.; Meszaros, N.J.; Kovacs, G.Z.; Prajczer, P.; Tamaga, I.; Seress, Z.; Kiszler, G.; Forgacs, A.; et al. Robustness analysis of denoising neural networks for bone scintigraphy. *Nucl. Instrum. Methods Phys. Res. Sect. A Accel. Spectrometers Detect. Assoc. Equip.* **2022**, *1039*, 167003. [CrossRef]
17. Wang, P.; Zhang, H.; Patel, V.M. SAR Image Despeckling Using a Convolutional Neural Network. *IEEE Signal Process. Lett.* **2017**, *24*, 1763–1767. [CrossRef]
18. Carrere, L.; Faugère, Y.; Ablain, M. Major improvement of altimetry sea level estimations using pressure-derived corrections based on ERA-Interim atmospheric reanalysis. *Ocean Sci.* **2016**, *12*, 825–842. [CrossRef]
19. Brodeau, L.; Sommer, J.L.; Albert, A. *Ocean-Next/eNATL60: Material Describing the Set-Up and the Assessment of NEMO-eNATL60 Simulations*; Zenodo: Geneva, Switzerland, 2020. [CrossRef]
20. Ajayi, A.; Le Sommer, J.; Chassignet, E.; Molines, J.; Xu, X.; Albert, A.; Cosme, E. Spatial and Temporal Variability of the North Atlantic Eddy Field From Two Kilometric-Resolution Ocean Models. *J. Geophys. Res. Ocean.* **2020**, *125*, e2019JC015827. [CrossRef]
21. Li, Z.; Liu, F.; Yang, W.; Peng, S.; Zhou, J. A Survey of Convolutional Neural Networks: Analysis, Applications, and Prospects. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *33*, 6999–7019. [CrossRef]
22. Albawi, S.; Mohammed, T.A.; Al-Zawi, S. Understanding of a convolutional neural network. In Proceedings of the 2017 International Conference on Engineering and Technology (ICET), Antalya, Turkey, 21–23 August 2017; pp. 1–6. [CrossRef]
23. Vincent, P.; Larochelle, H.; Lajoie, I.; Bengio, Y.; Manzagol, P.-A. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *J. Mach. Learn. Res.* **2010**, *11*, 3371–3408.
24. Guo, X.; Liu, X.; Zhu, E.; Yin, J. Deep Clustering with Convolutional Autoencoders. *Neural Inf. Process.* **2017**, *10635*, 373–382. [CrossRef]
25. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention 2015*; Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F., Eds.; Springer International Publishing: Cham, Switzerland, 2015; pp. 234–241.
26. Komatsu, R.; Gonsalves, T. Comparing U-Net Based Models for Denoising Color Images. *AI* **2020**, *1*, 465–487. [CrossRef]
27. Dong, H.; Yang, G.; Liu, F.; Mo, Y.; Guo, Y. Automatic Brain Tumor Detection and Segmentation Using U-Net Based Fully Convolutional Networks. *arXiv* **2017**, arXiv:1705.03820.
28. Agrawal, S.; Barrington, L.; Bromberg, C.; Burge, J.; Gazen, C.; Hickey, J. Machine Learning for Precipitation Nowcasting from Radar Images. *arXiv* **2019**, arXiv:1912.12132.
29. Zhao, H.; Gallo, O.; Frosio, I.; Kautz, J. Loss Functions for Image Restoration with Neural Networks. *IEEE Trans. Comput. Imaging* **2017**, *3*, 47–57. [CrossRef]
30. Taylor, V.; Nitschke, G. Improving Deep Learning using Generic Data Augmentation. *arXiv* **2017**, arXiv:1708.06020.
31. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *arXiv* **2019**, arXiv:1912.01703.
32. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2017**, arXiv:1412.6980.

33. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna. A Next-generation Hyperparameter Optimization Framework. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 25 July 2019; pp. 2623–2631. [CrossRef]
34. Wang, J.; Fu, L.-L.; Torres, H.S.; Chen, S.; Qiu, B.; Menemenlis, D. On the Spatial Scales to be Resolved by the Surface Water and Ocean Topography Ka-Band Radar Interferometer. *J. Atmos. Ocean Technol.* **2019**, *36*, 87–99. [CrossRef]
35. Dufau, C.; Orsztynowicz, M.; Dibarboure, G.; Morrow, R.; Le Traon, P.-Y. Mesoscale resolution capability of altimetry: Present and future. *J. Geophys. Res. Ocean.* **2016**, *121*, 4910–4927. [CrossRef]
36. Stiles, B.; Dubois. Algorithm Theoretical Basis Document for the Level 2 LR Sea Surface Height Science Algorithm. Technical Note Ref JPL-D56407 Press. 2022. Available online: https://podaac-tools.jpl.nasa.gov/drive/files/misc/web/misc/swot_mission_docs/pdd/D-56407_SWOT_Product_Description_L2_LR_SSH_20220902_RevA.pdf (accessed on 7 April 2023).
37. Molero, B.; Bohe, A.; Dubois, P. From 250 m to 2 km Posting: Implications of the L2B Averaging Step (Presented in the SWOT Meeting). 2022. Available online: https://swotst.aviso.altimetry.fr/programs/2022-swot-st-program (accessed on 7 April 2023).
38. Dibarboure, G.; Ubelmann, C.; Flamant, B.; Briol, F.; Peral, E.; Bracher, G.; Vergara, O.; Faugère, Y.; Soulat, F.; Picot, N. Data-Driven Calibration Algorithm and Pre-Launch Performance Simulations for the SWOT Mission. *Remote Sens.* **2022**, *14*, 6070. [CrossRef]
39. Jean-Michel, L.; Eric, G.; Romain, B.-B.; Gilles, G.; Angélique, M.; Marie, D.; Clément, B.; Mathieu, H.; Olivier, L.G.; Charly, R.; et al. The Copernicus Global 1/12° Oceanic and Sea Ice GLORYS12 Reanalysis. *Front. Earth Sci.* **2021**, *9*. [CrossRef]
40. Callies, J.; Ferrari, R.; Klymak, J.M.; Gula, J. Seasonality in submesoscale turbulence. *Nat. Commun.* **2015**, *6*, 6862. [CrossRef] [PubMed]
41. Richman, J.G.; Arbic, B.K.; Shriver, J.F.; Metzger, E.J.; Wallcraft, A.J. Inferring dynamics from the wavenumber spectra of an eddying global ocean model with embedded tides. *J. Geophys. Res. Atmos.* **2012**, *117*, C12012. [CrossRef]
42. Sasaki, H.; Klein, P.; Qiu, B.; Sasai, Y. Impact of oceanic-scale interactions on the seasonal modulation of ocean dynamics by the atmosphere. *Nat. Commun.* **2014**, *5*, 5636. [CrossRef] [PubMed]
43. Su, Z.; Wang, J.; Klein, P.; Thompson, A.F.; Menemenlis, D. Ocean submesoscales as a key component of the global heat budget. *Nat. Commun.* **2018**, *9*, 1–8. [CrossRef] [PubMed]
44. Ferrari, R.; Wunsch, C. Ocean Circulation Kinetic Energy: Reservoirs, Sources, and Sinks. *Annu. Rev. Fluid Mech.* **2009**, *41*, 253–282. [CrossRef]
45. Renault, L.; Marchesiello, P.; Masson, S.; McWilliams, J.C. Remarkable Control of Western Boundary Currents by *Eddy Killing* a Mechanical Air-Sea Coupling Process. *Geophys. Res. Lett.* **2019**, *46*, 2743–2751. [CrossRef]
46. Qiu, B.; Chen, S.; Klein, P.; Torres, H.; Wang, J.; Fu, L.-L.; Menemenlis, D. Reconstructing Upper-Ocean Vertical Velocity Field from Sea Surface Height in the Presence of Unbalanced Motion. *J. Phys. Oceanogr.* **2020**, *50*, 55–79. [CrossRef]
47. Siegelman, L.; Klein, P.; Rivière, P.; Thompson, A.F.; Torres, H.S.; Flexas, M.; Menemenlis, D. Enhanced upward heat transport at deep submesoscale ocean fronts. *Nat. Geosci.* **2019**, *13*, 50–55. [CrossRef]
48. Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; Zhang, L. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE Trans. Image Process.* **2017**, *26*, 3142–3155. [CrossRef]