
Filtering Internal Tides from Wide-Swath Altimeter Data Using Convolutional Neural Networks

Lguensat Redouane ¹, Fablet Ronan ², Le Sommer Julien ¹, Metref Sammy ¹, Cosme Emmanuel ¹, Ouenniche Kaouther ², Drumetz Lucas ², Gula Jonathan ³

¹ Univ Grenoble Alpes, CNRS, IRD, Grenoble INP,IGE, Grenoble, France.

² Univ Bretagne Loire, IMT Atlantique, LabSTICC, Brest, France.

³ IFREMER, LOPS, Brest, France.

Abstract :

The upcoming Surface Water Ocean Topography (SWOT) satellite altimetry mission is expected to yield two-dimensional high-resolution measurements of Sea Surface Height (SSH), thus allowing for a better characterization of the mesoscale and submesoscale eddy field. However, to fulfill the promises of this mission, filtering the tidal component of the SSH measurements is necessary. This challenging problem is crucial since the posterior studies done by physical oceanographers using SWOT data will depend heavily on the selected filtering schemes. In this paper, we cast this problem into a supervised learning framework and propose the use of convolutional neural networks (ConvNets) to estimate fields free of internal tide signals. Numerical experiments based on an advanced North Atlantic simulation of the ocean circulation (eNATL60) show that our ConvNet considerably reduces the imprint of the internal waves in SSH data even in regions unseen by the neural network. We also investigate the relevance of considering additional data from other sea surface variables such as sea surface temperature (SST).

Keywords : Internal Gravity Waves, Filtering, Deep Learning, Sea Surface Height, SWOT

1. INTRODUCTION

This study is conducted within the framework of the next-generation Surface Water Ocean Topography (SWOT) satellite altimetry mission. The SWOT altimeter will rely on its wide-swath capacities to provide unprecedented two-dimensional maps of Sea Surface Height (SSH) down to a 10 km effective resolution. These measurements are expected to drastically improve the quality of SSH mapping and therefore enhance our understanding of the mesoscale and submesoscale dynamics of the upper ocean [1, 2, 3]. SWOT




Fig. 1. Footprints of internal tides on the Laplacian of a SSH snapshot: (a) Original data (b) 24h-averaged data

data will also yield valuable information on tidal components of the SSH signal such as shelf tides, coastal tides and open-ocean internal tides. Separating the tidal and non-tidal components of the SSH data is a critical issue for the physical oceanography community to ensure a proper exploitation of SWOT data for studying mesoscale and submesoscale flows [4].

At the best of our knowledge, Torres et al. [5] have been the only to propose a method to extract Internal Gravity Waves (IGWs) from two-dimensional SSH snapshots. Their method is based on the identification of spectral slope discontinuities that separates IGWs at small scales from balanced motions at large scales. As such, it can recover the spectral content of the mesoscale and submesoscale flows only for spatial scales, where they strongly dominate over the IGWs, which may limit the use of this method. For instance, for this reason, only summer SSH data are investigated in [5].

Here, we state the filtering of tide signals in SSH fields as a supervised machine learning issue and explore deep learning techniques [6]. We present a case-study, which encompasses both summertime and wintertime data, based on high resolution oceanic numerical simulation data [7]. Our main contributions are as follows: i) investigating the extent to

which ConvNets can be relevant schemes for filtering IGWs, ii) illustrating the relevance of the proposed ConvNets for both summertime and wintertime sea surface dynamics, iii) evaluating the potential gain of considering multi-temporal data or multi-modal synergies.

This paper is organized as follows. Section 2 presents the dataset considered as a testbed in this study. We describe the proposed ConvNet schemes in Section 3. Numerical experiments are reported and discussed in Section 4, and finally a conclusion is drawn in Section 5.

2. DATA

2.1. Data preparation

Advancements in ocean numerical modeling have reached a state where realistic numerical simulation with Petabytes (10^{15} bytes) of data are available and represent a real opportunity for machine learning based techniques. In this work, we consider the high-resolution eNATL60 (North Atlantic, hourly temporal resolution, $1/60^\circ$ horizontal resolution) configuration of the NEMO (Nucleus for European Modelling of the Ocean) modelling system [7]. Within the framework of the upcoming SWOT mission, we degrade the resolution by a factor of 3. We here focus on a subregion of the North Atlantic, namely, the OSMOSIS region (44.821°N - 55.363°N , 20.016°W - 10.008°W). The OSMOSIS region has a weaker large-scale component compared to highly energetic regions such as the Gulf Stream [8]. This makes OSMOSIS a relevant region for assessing the ability of the investigated filtering schemes to recover small-scale dynamics which are the main component of interest awaited from SWOT.

We may point out that mesoscale and submesoscale upper ocean dynamics are known to be seasonally-dependent. Hereafter, we refer to "JAS" (July, August, September) for summer and "JFM" (January, February, March) for winter, when considering a season-related analysis. The one-year span of eNATL60 resorts in $24 \times 90 = 2160$ images for each dataset. We also split the region into several non-overlapping 64×64 patches. Splitting the dataset into train/val/test splits is done spatially which is arguably more challenging in our context than a temporal splitting, since oceanic patterns may strongly depend on the geographic area (different bathymetry, Rossby radius, etc.). We consider a spatial split according to the latitude. Overall, 5 boxes from 44.821°N - 52°N were used to train our models resulting in $5 \times 2160 = 10800$ patches, and 1 box from the northwestern area of the OSMOSIS region is considered as the test region.

2.2. Reference data used for supervised training

Within a supervised training framework, we exploit the hourly time sampling of the eNATL60 simulation to build a tide-free reference using a 24-hour time averaging. This straightforward approach is commonly used in the analysis of the out-

puts of oceanic numerical models [9]. More complex time filtering approaches could be considered to improve the generation of the tide-free samples.

An example of the Laplacian of a SSH snapshot from the considered dataset and its corresponding 24h filtered SSH is shown in Figure 1. we show Laplacian fields, which relate to the vorticity and clearly exhibit the footprints of the IGWs.

3. METHODS

3.1. Convolutional Neural Networks

Inspired by the connectivity patterns of neurons in animals' visual cortex, convolutional neural networks (hereinafter ConvNets) are one of the main and most important breakthroughs in neural networks literature [10]. They have rapidly become a key component of state-of-the-art deep learning architectures in numerous computer vision tasks. Mathematically speaking, ConvNets in their basic form consist of a cascade of convolutional layers where the output of layer k (consisting in so called feature maps) is a function of an affine transformation of the previous layer output:

$$h_{ij}^k = f\left((W^k * x)_{ij} + b_k\right), \quad (1)$$

where $*$ denotes the convolution operation, f is a nonlinear function (a common choice is the Rectifier Linear Unit $\text{ReLU}(x) = \max(0, x)$). Weights W^k and biases b_k are the parameters of the ConvNets to be inferred. The learning step comes to solve for a minimization issue with regard to the parameters of the ConvNet, for instance using the Stochastic Gradient Descent (SGD) with the backpropagation algorithm. The interested reader can refer to [6] and references therein for more reading about ConvNets and their mathematical properties.

Here, we focus on ResNet, a specific class of ConvNets which rely on residual blocks [11, 12]. Let \mathcal{G} be a series of neural operations (convolutions, nonlinearities, normalization, etc), a residual block takes an input and adds it to its transformed version by \mathcal{G} , i.e., $h^k = h^k + \mathcal{G}(h^k)$. The ease of training and the simple mathematical intuitions behind residual blocks made them ubiquitous in recent ConvNet architectures and led to successful application in several image and signal processing applications [13].

3.2. ConvNets for filtering IGWs

Here, we aim to design a ConvNet, denoted by \mathcal{F} , to reconstruct a tide-free SSH field, referred to as S_f from some *Input*, which may be given as a single or a series of SSH snapshots, possibly complemented by other observed sea surface fields, such as SST fields. Formally, the considering filtering issue comes to train \mathcal{F} such that $S_f = \mathcal{F}(\text{Input})$.

We consider the following architecture for ConvNet \mathcal{F} . It consists of a first Conv layer followed by N_r residual units as




Fig. 2. Illustration of the ConvNet based approach for filtering IGWs. A Residual block consists of two series of Batch Normalization (BN), LeakyReLU activation followed by a Conv layer. The notation $M@H*H$ means M filters of size $H*H$.

introduced in [11] then a final regression Conv layer. In this paper, we consider three choices for the input: i) **SSH**: Our ConvNet takes a single SSH snapshot as input, and outputs its filtered version. ii) **3SSH**: SWOT fast-sampling phase is expected to deliver SSH measurements on a daily cycle, we simulate this behaviour by considering as input SSH at time t concatenated with SSH at $t-24h$ and $t+24h$. iii) **3SSH-SST**: we complement the three successive daily SSH snapshots as above with a SST field. The later is motivated by known relationships between SSH and SST features for some dynamical modes [14].

To train our ConvNet, we consider a loss function, which combines a mean square error on the SSH and a mean absolute error on the Laplacian of the SSH:

$$\mathcal{L}(S_f, \hat{S}_f) = \|S_f - \hat{S}_f\|^2 + \alpha |\nabla^2 S_f - \nabla^2 \hat{S}_f|, \quad (2)$$

The rationale behind this choice is that SSH fields are rather smooth while vorticity fields contain a considerable amount of high frequency information that need to be preserved. In the machine learning community, it is known that the use of the MSE loss function leads to blurry images and that adding losses on image gradients or Laplacians can help sharpen the predictions [15].

4. NUMERICAL EXPERIMENTS

4.1. Experimental setup

Weights of the convolutional layers were initialized with the Kaiming initialization [16], biases to zero. To ease the optimization of the ResNets at early stages, we found it useful to use the "Zero γ " heuristic [17]. Slope of the LeakyReLUs is 0.2. The networks were coded using PyTorch, and trained using the ADAM optimizer with an initial learning rate of $3e-4$. The learning rate is divided by 10 every 100 epochs.

Early stopping is also used leading all the networks to converge after around 400 epochs. Experiments were run using a Nvidia Tesla V100 GPU with a batch size of 32. The hyper-parameters N_r and α were tuned using a grid search leading to $N_r = 10$ and $\alpha = 10^3$.

4.2. Results

We compare the performance of the variants of the proposed ConvNet. As baseline, we consider a linear 5×5 filtering of the SSH. The parameters of the linear filter are optimized with regard to loss function (2) to provide a fair comparison with the proposed schemes. This linear filtering approach is referred to as **Baseline**.

Top panels of Figure 3 show the power spectral density (PSD) functions of ∇^2 SSH resulting from the 4 competing methods. At large scale, the different approaches are close to the PSD of the reference data and differences start to appear at finer scales. For Summer data (left panel), the three proposed methods beat the baseline by a large margin, and unsurprisingly 3SSH and 3SSH-SST outperform 1SSH. For Winter data (right panel), the intensification of sea surface dynamics enhances non-wave processes at scales similar to IGWs, which makes the separation much more challenging. Still, the PSD shows that the three ConvNets outperform the baseline.

We further analyse the reconstruction performance according to ratio $R = 1 - \frac{E[(S_{ref} - S_f)^2]}{E[S_f^2]}$ that reflects the relative quality of the estimation as a function of the spatial frequency. R values close to 1 indicate a perfect estimation. The R values calculated on ∇^2 SSH are shown in the bottom panels of 3. For both JAS and JFM datasets we now clearly observe the differences between the introduced models. The baseline does not behave badly for a rough approximation, but is quickly outperformed by 1SSH and




Fig. 3. (Top) Power spectral density calculated on the Laplacian of SSH. (Bottom) Ratio R . Results on the JAS dataset are on the left, those on JFM on the right.

3SSH. The addition of SST introduces a slight improvement over 3SSH. Examples of SSH filtered maps and their Laplacians can be found in this following Github repository <https://github.com/CIA-Oceanix/DetideNet> along with the codes used for this work.

5. CONCLUSION

In this work, we present a new deep learning based method to filter internal gravity waves from SSH data. Aiming for a proof of concept, this paper considered a numerical simulation as a testbed and has shown the relevance of the use of domain knowledge in the derivation of the loss function used for training the ConvNets. Experiments performed on winter and summer datasets prove that our method can yield competitive results with regard to simple linear filters. Results as presented are promising and call for more thorough investigation of the effect of adding other sources of noise that contaminate satellite altimeters, especially in the context of SWOT. This is the main subject of future work.

6. ACKNOWLEDGMENTS

This research was funded through CNES OST/ST and the SWOT Science Team. S. Metref is funded by ANR through contract number ANR-17-CE01-0009-01. R. Fablet is supported by CNES (grant OSTST-MANATEE) and ANR (AI Chair Oceanix, EUR Isblue and Melody project).

7. REFERENCES

- [1] Michael Durand, Lee-Lueng Fu, Dennis P Lettenmaier, Douglas E Alsdorf, Ernesto Rodriguez, and Daniel Esteban-Fernandez, “The surface water and ocean topography mission: Observing terrestrial surface water and oceanic submesoscale eddies,” *Proceedings of the IEEE*, vol. 98, no. 5, pp. 766–779, 2010.
- [2] Lee-Lueng Fu, Douglas Alsdorf, Ernesto Rodriguez, Rosemary Morrow, Nelly Mognard, Juliette Lambin, Parag Vaze, and Thierry Lafon, “The swot (surface water and ocean topography) mission: spaceborne radar interferometry for oceanographic and hydrological applications,” in *OCEANOBS09 Conference*, 2009.
- [3] Lee-Lueng Fu and Raffaele Ferrari, “Observing oceanic submesoscale processes from space,” *Eos, Transactions American Geophysical Union*, vol. 89, no. 48, pp. 488–488, 2008.
- [4] Brian K Arbic, Florent Lyard, Aurelien Ponte, Richard D Ray, James G Richman, Jay F Shriver, Edward Zaron, and Zhongxiang Zhao, “Tides and the swot mission: Transition from science definition team to science team,” 2015.
- [5] H S Torres, P Klein, L Siegelman, B Qiu, S Chen, C Ubelmann, J Wang, D Menemenlis, and L.-L. Fu, “Diagnosing Ocean-Wave-Turbulence Interactions From Space,” *Geophysical Research Letters*, vol. 46, no. 15, pp. 8933–8942, 2019.
- [6] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436, 2015.
- [7] Gurvan Madec et al., “Nemo ocean engine,” 2015.
- [8] Christian E Buckingham, Alberto C Naveira Garabato, Andrew F Thompson, Liam Brannigan, Ayah Lazar, David P Marshall, AJ George Nurser, Gillian Damerell, Karen J Heywood, and Stephen E Belcher, “Seasonality of submesoscale flows in the ocean surface boundary layer,” *Geophysical Research Letters*, vol. 43, no. 5, pp. 2118–2126, 2016.
- [9] James G Richman, Brian K Arbic, Jay F Shriver, E Joseph Metzger, and Alan J Wallcraft, “Inferring dynamics from the wavenumber spectra of an eddying global ocean model with embedded tides,” *Journal of Geophysical Research: Oceans*, vol. 117, no. C12, 2012.
- [10] Yann LeCun, Yoshua Bengio, et al., “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, pp. 1995, 1995.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Identity mappings in deep residual networks,” in *European conference on computer vision*. Springer, 2016, pp. 630–645.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [13] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, MIT press, 2016.
- [14] Pierre Tandeo, Bertrand Chapron, Sileye Ba, Emmanuelle Autret, and Ronan Fablet, “Segmentation of mesoscale ocean surface dynamics using satellite sst and ssh observations,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 7, pp. 4227–4235, 2013.

- [15] Michael Mathieu, Camille Couprie, and Yann LeCun, “Deep multi-scale video prediction beyond mean square error,” *arXiv preprint arXiv:1511.05440*, 2015.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.
- [17] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li, “Bag of tricks for image classification with convolutional neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 558–567.