
Multi-scale decomposition of sea surface height snapshots using machine learning

Jingwen Lyu*

Applied Physics and Applied Mathematics
Columbia University
New York, USA
jl6811@columbia.edu

Yue Wang*

Applied Physics and Applied Mathematics
Columbia University
New York, USA
yw4236@columbia.edu

Christian Pedersen

Courant Institute
New York University
New York, USA
c.pedersen@nyu.edu

Spencer Jones

Department of Oceanography
Texas A&M University
Texas, USA
spencerjones@tamu.edu

Dhruv Balwada

Lamont Doherty Earth Observatory
Columbia University
New York, USA
dbalwada@ldeo.columbia.edu

Abstract

Knowledge of ocean circulation is important for understanding and predicting weather and climate, and managing the blue economy. This circulation can be estimated through Sea Surface Height (SSH) observations, but requires decomposing the SSH into contributions from balanced and unbalanced motions (BMs and UBM). This decomposition is particularly pertinent for the novel SWOT satellite, which measures SSH at an unprecedented spatial resolution. Specifically, the requirement, and the goal of this work, is to decompose instantaneous SSH into BMs and UBMs. While a few studies using deep learning (DL) approaches have shown promise in framing this decomposition as an image-to-image translation task, these models struggle to work well across a wide range of spatial scales and require extensive training data, which is scarce in this domain. These challenges are not unique to our task, and pervade many problems requiring multi-scale fidelity. We show that these challenges can be addressed by using zero-phase component analysis (ZCA) whitening and data augmentation; making this a viable option for SSH decomposition across scales.

1 Introduction

Oceans play a critical role in Earth's climate and weather, as they store and circulate large quantities of heat, carbon, and other biogeochemically relevant tracers (Cronin et al., 2019)(Dong et al., 2014). This circulation is composed of many different physical processes across a wide range of scales (Talley, 2011). Inferring circulation patterns plays a vital role in understanding climate variability, as well as management of vital sectors like shipping, fisheries, ocean energy, and marine carbon dioxide removal. Currently the only way to observationally estimate this flow globally and across a wide range of scales is through satellite altimetry, which measures sea surface height (SSH). SSH

corresponds to the pressure at the surface of the ocean, and is related to the surface flow. However, the relationship is not simple, and depends on the nature of the processes contributing to the flow. These processes can be broadly divided into two categories, balanced motions (BM, slowly evolving) and unbalanced motions (UBM, rapidly evolving) (Torres et al., 2018; Klein et al., 2019). It is essential to decompose the SSH into the contributions from BMs and UBMs before attempting to estimate flow from SSH, particularly when using high spatial resolution SSH data. Theoretically, BMs and UBMs can be decomposed by applying temporal or spatio-temporal filters to data (Jones et al., 2023), but this conventional approach requires data with a temporal resolution on the order of minutes to a few hours.

Traditional satellite altimeters are used to produce maps of SSH with spatial and temporal resolution of $O(100)\text{km}$ and $O(10)$ days. At these scales BMs are dominant, and no decomposition has been necessary. However, the recently launched Surface Water and Ocean Topography (SWOT) satellite provides data in a 2D swath at an unprecedented spatial resolution of 5-10 km (Fu et al., 2024; Chelton, 2024), where BMs and UBMs have comparable magnitudes. However, the SWOT satellite only passes over a region once every 21 days, which does not allow the use of conventional methods for UBM-BM separation that require data at high temporal resolution. *Thus, decomposition approaches that can work with only instantaneous SSH observations (snapshots) need to be developed, and this is the focus of our study.*

With the rapid advancement of machine learning techniques, SSH decomposition may be approached as an Image-to-Image translation problem. A few studies have now used convolutional neural networks (CNNs) to decompose the BMs and UBMs (Lguensat et al., 2020; Wang et al., 2022; Gao et al., 2024), with some success. However, significant challenges persist due to the multi-scale nature of these fields, implying that there is important detail in the SSH fields across a large range of scales. Often the amount of variance or signal present at different scales can differ by orders of magnitudes, as exemplified by the power law decay of variance across wavenumbers in power spectral density (PSD) plots (Vallis, 2017).

Pixel-wise mean squared error (MSE) loss fails to effectively capture patterns across the full range of scales for these multi-scale fields, because this choice of loss tends to prioritize learning the scales with the dominant signal (Rahaman et al., 2019). To address this issue, (Lguensat et al., 2020) and (Gao et al., 2024) augmented their loss function with gradients of the predicted field, which allowed the loss function to also focus on small scale information. While this approach improves the model’s performance at smaller scales, it introduces two significant drawbacks: (1) the method requires careful tuning of the gradient loss weight, and this weight likely needs to be adjusted for every new training dataset with different intensities of BMs and UBMs; (2) the gradient loss can force the model to focus on high-wavenumber information, and with insufficient training data this can easily lead to overfitting in the highest-wavenumber range.

To address the challenge of working with multi-scale data, we propose implementing ZCA transformation to whiten the UBM prior to processing. This approach offers several advantages: (1) It enhances information across a range of scales, thereby reducing the necessity for gradient loss. (2) It maintains compatibility with pixel-wise Mean Squared Error (MSE) calculations and mitigates the risk of overfitting at extremely high frequencies. (3) ZCA transformation also approximates the full covariance matrix of samples to an identity matrix (Kessy et al., 2018), which effectively reduces correlation between samples, leading to improved training stability and computational efficiency.

2 Methodology

Our task is to decompose the total SSH (η) into contributions from BMs (η_B) and UBMs (η_{UB}), where $\eta = \eta_B + \eta_{UB}$. We achieve this by training a ML model that can predict the UBM ($\hat{\eta}_{UB}$) for a given input η . The prediction for BM is consequently defined as $\hat{\eta}_B = \eta - \hat{\eta}_{UB}$.

Data: We utilize data from a high-resolution global ocean simulation, focusing on the Agulhas retroflection region (Jones et al., 2023). The dataset comprises 70 daily snapshots of total SSH, BM, and UBM from September to November 2011. This region presents challenges in distinguishing BM and UBM due to their amplitude disparity, diverse regional patterns, and the significance of UBM in gradient fields. Detailed information about the dataset and its processing is provided in Appendix B.

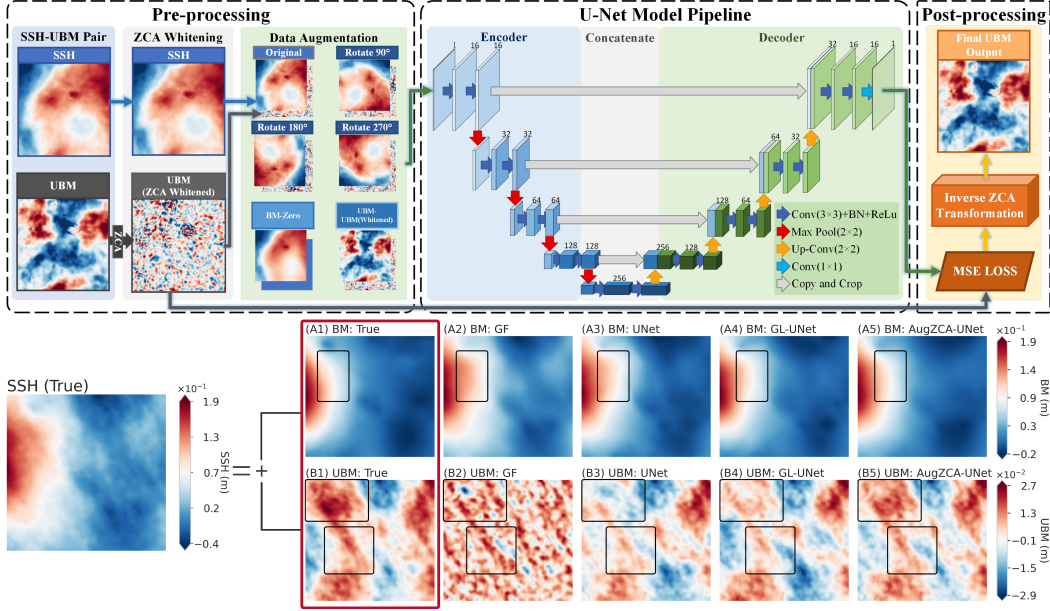


Figure 1: Top: Schematic representing model pipeline and architecture. Bottom: Visualization of UBM and BM for our best performing test sample using AugZCA-UNet. A1-A5: BM; B1-B5: UBM. Black squares are a visual aid.

To mitigate data scarcity, we employed two data augmentation techniques: (1) rotational augmentation (90°, 180°, and 270°) for promoting rotational invariance, and (2) synthetic sample generation of purely BM or UBM flows by adding pairs ($\eta_B, \eta_{UB} = 0$) and (η_{UB}, η_{UB}). As a final and crucial preprocessing step, we applied *ZCA whitening* to the output, η_{UB} . This technique, detailed in Appendix C, represents a key contribution of our methodology, as it enhances the model’s ability to detect and process both large-scale and fine-scale features in the data.

ML Models: To assess the effectiveness of our ML-based techniques and the impact of data augmentation and ZCA whitening, we systematically compared the following approaches: (1) Gaussian Filter (GF)—a non-ML baseline detailed in Appendix F; (2) UNet—an ML baseline without preprocessing or additional gradient loss terms, described in Appendix D; (3) UNet with gradient loss (GL-UNet)—incorporating gradient loss terms with varying weights ($\alpha = 0.5, 1, 5, 10$), with details provided in Appendix I; (4) UNet with data augmentation (Aug-UNet)—employing rotation and synthetic data generation techniques; (5) UNet coupled with ZCA whitening (ZCA-UNet); and (6) UNet coupled with both data augmentation and ZCA whitening (AugZCA-UNet). Notably, all ML models share the same UNet configuration as the ‘raw’ UNet, and models employing ZCA did not include gradient loss terms, allowing us to isolate the effects of each technique.

Model performance was rigorously evaluated using two complementary metrics: (i) Pixel-wise absolute error distribution of $\hat{\eta}_B$ and $|\nabla \hat{\eta}_B|$, which assesses local prediction accuracy across the spatial domain, and (ii) Power Spectral Density (PSD) of $\hat{\eta}_B$ and $\hat{\eta}_{UB}$, which evaluates the model’s ability to capture multi-scale characteristics across spatial scales, providing insight into the multi-scale fidelity of the predictions.

3 Results

A visual comparison of SSH decomposition methods into BM and UBM reveals that all ML approaches show skill in this task (Figure 1, additional samples in Figure App.3). The ML models effectively reduce fine-scale UBM features in the gradient field, which is crucial for flow estimation methods requiring accurate BM gradients. Unlike the GF, ML based models avoid over-smoothing the predicted BM and its gradients.

Quantitative evaluation using pixel-wise absolute error shows AugZCA-UNet consistently outperforms other methods for both BM and its gradient across all statistical measures (Table 1, Appendix E). Conversely, GF consistently performs the worst, indicating ML approaches can surpass traditional filters. GL-UNet’s performance varies parabolically with gradient weight, peaking at 1 and declining at lower or higher values, highlighting the need for precise gradient weight tuning. ZCA-UNet and AugZCA-UNet, requiring no gradient weight tuning, demonstrate superior performance due to their architectural adjustments. Data augmentation techniques also consistently enhance model performance.

Table 1: Pixel-wise Absolute Error Distribution Measures for $\hat{\eta}_B$ ($\times 10^{-2}$) and $|\nabla \hat{\eta}_B|$ ($\times 10^{-3}$)

Measures	GF	GL-UNet					Aug-UNet	ZCA-UNet	AugZCA-UNet	
		$\alpha = 0$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 5$	$\alpha = 10$				
$\hat{\eta}_B$	Median	0.406	0.404	0.420	0.397	0.398	0.406	0.384	0.393	0.372
	P95 (95%)	1.68	1.52	1.49	1.48	1.51	1.53	1.50	1.52	1.40
$ \nabla \hat{\eta}_B $	Median	0.505	0.464	0.447	0.450	0.450	0.453	0.460	0.450	0.427
	P95 (95%)	2.20	1.80	1.70	1.72	1.74	1.75	1.78	1.71	1.64

For many oceanographic and climate science applications, it is important that the ML models perform well across a wide range scales. To assess the prediction skill across spatial scales, we consider the wavenumber PSD of the predictions and the R^2 (defined in Appendix J), as shown in Figure 2. A more rigorous quantitative comparison is presented in Appendix H. GF, constrained by its fixed filter length scale, only works well on scales where BM or UBM are dominant, and is unable to perform at wavenumbers where they have comparable magnitudes or in cases where we are interested in predicting the sub-dominant signal. In contrast, the UNet models exhibit varying degrees of success across the wavenumber spectrum. At low wavenumbers (below 10^{-1} cpkm), all UNet models achieve relatively high accuracy in predicting the true PSD. However, at higher wavenumbers the performance characteristics of DL models diverge. The GL-UNet can be tuned to outperform a raw UNet, but often suffers from noisy prediction at the highest wavenumber that can severely degrade the skill in predicting gradient fields. In contrast the ZCA-UNet and AugZCA-UNet consistently exhibit superior performance across the full wavenumber spectrum, and are not plagued by noise at the highest wavenumbers.




Figure 2: Comparison of PSD and R^2 across various models. (a) Mean PSD over test samples for UBM (left) and BM (right). (b) R^2 for UBM (left) and BM (right).

4 Discussion and Outlook

Our novel AugZCA-UNet shows significant promise in skillfully decomposing SSH at finer scales with lower variance, while partially addressing data scarcity. It outperforms the previously proposed GL-UNet, demonstrating superior performance in both pixel-wise error (sensitive to dominant signals) and PSD accuracy (sensitive to multi-scale information). ZCA whitening is key to our model’s success, which enhances the loss function’s sensitivity to fine-scale variations that are otherwise overshadowed by dominant large-scale variations. This technique also has the potential to mitigate challenges across a broad spectrum of climate-related research where capturing multi-scale phenomena is crucial (Lai et al., 2024).

While our work marks an important step forward in applying ML to multi-scale systems, several challenges must be addressed before our decomposition approach can be effectively applied to actual SSH data from SWOT. Firstly, ZCA constrains the ML model to a specific image size, compromising

the fully convolutional nature of UNets. Therefore, pragmatic processing choices will be necessary when working with datasets that do not naturally align with the image sizes used during model training or when making predictions over large domains. Secondly, the memory cost of performing ZCA increases with image size, necessitating modifications to the standard ZCA algorithm for more efficient memory management in tasks requiring larger images. Thirdly, while our dataset included gappy data due to land and other numerical artifacts, we have not thoroughly assessed the model's performance near these gaps. This is particularly crucial for SWOT, where the satellite swath is 120 km wide but consists of two 50 km sections separated by a 20 km gap. Finally, the generalization of our model, particularly relative to spatial filters like GF, remains an open question. In future work, we plan to evaluate the model's generalization properties and explore whether flow-dependent normalization (Beucler et al., 2024) or transfer learning approaches (Xiao et al., 2023) can enhance its applicability.

References

- Joseph K Ansong, Brian K Arbic, Arin D Nelson, Matthew H Alford, Eric Kunze, Dimitris Menemenlis, Anna C Savage, Jay F Shriver, Alan J Wallcraft, and Maarten C Buijsman. Surface and sub-surface kinetic energy wavenumber-frequency spectra in global ocean models and observations. *Journal of Geophysical Research: Oceans*, 129(6):e2023JC020480, 2024.
- Brian K Arbic, Shane Elipot, Jonathan M Brasch, Dimitris Menemenlis, Aurélien L Ponte, Jay F Shriver, Xiaolong Yu, Edward D Zaron, Matthew H Alford, Maarten C Buijsman, et al. Near-surface oceanic kinetic energy distributions from drifter observations and numerical models. *Journal of Geophysical Research: Oceans*, 127(10):e2022JC018551, 2022.
- Tom Beucler, Pierre Gentine, Janni Yuval, Ankitesh Gupta, Liran Peng, Jerry Lin, Sungduk Yu, Stephan Rasp, Fiaz Ahmed, Paul A O’Gorman, et al. Climate-invariant machine learning. *Science Advances*, 10(6):eadj7250, 2024.
- Dudley B Chelton. A post-launch update on the effects of instrumental measurement errors on swot estimates of sea-surface height, velocity and vorticity. *Journal of Atmospheric and Oceanic Technology*, 2024.
- Meghan F Cronin, Chelle L Gentemann, James Edson, Iwao Ueki, Mark Bourassa, Shannon Brown, Carol Anne Clayson, Chris W Fairall, J Thomas Farrar, Sarah T Gille, et al. Air-sea fluxes with a focus on heat and momentum. *Frontiers in Marine Science*, 6:430, 2019.
- Changming Dong, James C. McWilliams, Yu Liu, and Dake Chen. Global heat and salt transports by eddy movement. *Nature Communications*, 5, 2014.
- Lee-Lueng Fu, Tamlin Pavelsky, Jean-Francois Cretaux, Rosemary Morrow, J. Thomas Farrar, Parag Vaze, Pierre Sengenès, Nadya Vinogradova-Shiffer, Annick Sylvestre-Baron, Nicolas Picot, and Gerald Dibarboure. The surface water and ocean topography mission: A breakthrough in radar remote sensing of the ocean and land surface water. *Geophysical Research Letters*, 51(4): e2023GL107652, 2024.
- Zhanwen Gao, Bertrand Chapron, Chunyong Ma, Ronan Fablet, Quentin Febvre, Wenxia Zhao, and Ge Chen. A deep learning approach to extract balanced motions from sea surface height snapshot. *Geophysical Research Letters*, 51(7):e2023GL106623, 2024.
- I. Grooms, N. Loose, R. Abernathey, J. M. Steinberg, S. D. Bachman, G. Marques, A. P. Guillaumin, and E. Yankovsky. Diffusion-based smoothers for spatial filtering of gridded geophysical data. *Journal of Advances in Modeling Earth Systems*, 13(9):e2021MS002552, 2021.
- C. Spencer Jones, Qiyu Xiao, Ryan P. Abernathey, and K. Shafer Smith. Using lagrangian filtering to remove waves from the ocean surface velocity field. *Journal of Advances in Modeling Earth Systems*, 15(4):e2022MS003220, 2023.
- Agnan Kessy, Alex Lewin, and Korbinian Strimmer. Optimal whitening and decorrelation. *The American Statistician*, 72(4):309–314, 2018.

- Patrice Klein, Guillaume Lapeyre, Lia Siegelman, Bo Qiu, Lee-Lueng Fu, Hector Torres, Zhan Su, Dimitris Menemenlis, and Sylvie Le Gentil. Ocean-scale interactions from space. *Earth and Space Science*, 6(5):795–817, 2019.
- Ching-Yao Lai, Pedram Hassanzadeh, Aditi Sheshadri, Maike Sonnewald, Raffaele Ferrari, and Venkatramani Balaji. Machine learning for climate physics and simulations. *arXiv preprint arXiv:2404.13227*, 2024.
- Redouane Lguensat, Ronan Fablet, Julien le Sommer, Sammy Metref, Emmanuel Cosme, Kaouther Ouenniche, Lucas Drumetz, and Jonathan Gula. Filtering internal tides from wide-swath altimeter data using convolutional neural networks. In *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, pages 3904–3907, 2020.
- Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks, 2019. URL <https://arxiv.org/abs/1806.08734>.
- Lynne D Talley. *Descriptive physical oceanography: an introduction*. Academic press, 2011.
- Hector S Torres, Patrice Klein, Dimitris Menemenlis, Bo Qiu, Zhan Su, Jinbo Wang, Shuiming Chen, and Lee-Lueng Fu. Partitioning ocean motions into balanced motions and internal gravity waves: A modeling study in anticipation of future space missions. *Journal of Geophysical Research: Oceans*, 123(11):8084–8105, 2018.
- HS Torres, Patrice Klein, Lia 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*, 46(15):8933–8942, 2019.
- Geoffrey K Vallis. *Atmospheric and oceanic fluid dynamics*. Cambridge University Press, 2017.
- Han Wang, Nicolas Grisouard, Hesam Salehipour, Alice Nuz, Michael Poon, and Aurélien L Ponte. A deep learning approach to extract internal tides scattered by geostrophic turbulence. *Geophysical Research Letters*, 49(11):e2022GL099400, 2022.
- Norbert Wiener. Generalized harmonic analysis. *Acta mathematica*, 55(1):117–258, 1930.
- Qiyu Xiao, Dhruv Balwada, C Spencer Jones, Mario Herrero-González, K Shafer Smith, and Ryan Abernathy. Reconstruction of surface kinematics from sea surface height using neural networks. *Journal of Advances in Modeling Earth Systems*, 15(10):e2023MS003709, 2023.

APPENDIX

A Data Availability Statement

The code repository for this work is at <https://github.com/Hamsterrrrrrrr/multiscale-ssh.git>. The data sets used for training and evaluation are available at <https://zenodo.org/records/7495109>.

B Comprehensive Overview of Dataset

Ocean models solve the fundamental equations of fluid flow, primitive equations, under realistic boundary conditions and forcing. At high-resolutions these simulations can generate a statistically well matched representation of the real world (Arbic et al., 2022; Ansong et al., 2024). Here we utilized the data generated by one of the highest resolution global ocean simulations that is currently available, referred to in the climate science community as LLC4320. While SSH observed by SWOT has not yet been decomposed into BMs and UBMs, we utilized a publicly available decomposed subset of the LLC4320 data (Jones et al., 2023). This was achieved using spatio-temporal filters, as the LLC4320 data is available at a high temporal resolution of 1 hour, making it directly amenable to conventional decomposition methods. The primary objective of our study is to assess the effectiveness of ML in performing this decomposition using only instantaneous/snapshot data, thereby addressing the temporal resolution limitations of SWOT satellite observations.

The spatial domain covers the Agulhas retroflection region (15°W to 29°E and 27°S to 57°S), located to the south and west of Africa. The temporal domain comprises 70 daily snapshots spanning from September 14, 2011 to November 22, 2011. We divided the dataset through temporal partitioning (days 0–60 for training, 61–65 for validation, and 66–70 for testing) and spatial tiling (400 non-overlapping 108×108 pixel patches from a 2160×2160 pixel domain), yielding 24,000 training images. The land regions (represented by NaN values) are masked when calculating the loss function, thereby reducing their influence on the model training.

In this spatiotemporal domain, accurately distinguishing between balanced and unbalanced components from snapshots presents three significant challenges:

1. **Scale disparity:** UBM has a magnitude approximately one-tenth that of total SSH. This disparity hinders CNNs in detecting UBM features without preprocessing.
2. **Diverse dynamics:** The dataset exhibits varied dynamical patterns across spatial regions, necessitating a robust scheme capable of isolating UBM across different oceanic dynamics.
3. **Gradient field complexity:** Despite their small magnitude, the UBM gradient fields exhibit amplitudes comparable to both BM and SSH. This similarity hinders the application of typical methods, such as geostrophic balance, to estimate the flow.

These challenges underscore the importance of effective UBM filtering for accurate ocean dynamics analysis.




Figure App.1: A snapshot of the whole region SSH, BM and UBM

C Zero-phase Component Analysis (ZCA)

Let $X_{\eta_{\text{UB}}} \in \mathbb{R}^{n \times d}$ represent the data matrix of UBM, where n is the number of training samples of UBM, and $d = 108 \times 108$ is the number of flattened features for a single UBM snapshot. The ZCA whitening process is implemented as follows:

1. **Data Centering:** Compute the centered data matrix $X_{\eta_{\text{UB}}}^c = X_{\eta_{\text{UB}}} - \mu_{\eta_{\text{UB}}}$ where $\mu_{\eta_{\text{UB}}}$ is the mean vector computed across all training UBM samples.
2. **Covariance Matrix Computation:** Calculate $\Sigma = \frac{1}{n-1} (X_{\eta_{\text{UB}}}^c)^T X_{\eta_{\text{UB}}}^c$.
3. **Eigendecomposition:** Decompose $\Sigma = U\Lambda U^T$, where U is the matrix of eigenvectors and Λ is the diagonal matrix of eigenvalues.
4. **Compute ZCA Matrix:** Calculate $W_{\text{ZCA}} = U(\Lambda + \epsilon I)^{-1/2} U^T$ where I is the identity matrix and ϵ is a small parameter added for numerical stability. This parameter controls the extent of whitening by slightly perturbing the eigenvalues. In our experiments, we set $\epsilon = 10^{-5}$. We found that performance was not highly sensitive to variations in ϵ , testing values from 10^{-3} to 10^{-7} .
5. **Apply Transformation:** Transform all UBM samples (training, validation, and test sets) using the same mean and ZCA matrix calculated from the training data: $X_{\text{ZCA}} = X_{\eta_{\text{UB}}}^c W_{\text{ZCA}}$

In the frequency domain, the covariance matrix is intrinsically linked to the Power Spectral Density (PSD) according to the Wiener-Khinchin theorem, which states that the PSD is the Fourier transform of the autocorrelation function (Wiener, 1930). Whitening the data through ZCA transforms the full covariance matrix into an identity matrix, enforcing a uniform power distribution across all frequencies. This process effectively results in a flatter PSD post-whitening transformation.

Figure App.2 illustrates an example of the impact of ZCA whitening on a UBM snapshot sample. The ZCA process notably enhances small-scale features, resulting in a more granular appearance in the spatial domain and a flatter spectrum in the frequency domain (Note that the whitened spectrum varies across about 3 orders of magnitude, while the original spectrum had a variation across about 7 orders of magnitude).




Figure App.2: Effect of ZCA whitening on the original UBM snapshot (a) and its PSD (b)

D UNet Architecture and Training Specifications

The UNet architecture employed in this study utilizes Rectified Linear Unit (ReLU) activation functions and the Adam optimizer. We implemented a dynamic learning rate schedule starting at an initial rate of 10^{-3} . This rate is adjusted during training with a reduction factor of 0.8, which is applied when the validation loss plateaus for three consecutive epochs. To guard against overfitting, we incorporated an early stopping mechanism with a patience of 20 epochs. This approach allows the model to continue training as long as it shows improvement, but halts the process if no progress is observed over an extended period.

The complete network configuration is illustrated in Figure 1. It's important to note that this configuration remained consistent across all UNet variants used in our experiments. This consistency

was maintained to ensure a fair comparison between different approaches and to isolate the effects of our proposed modifications.

E Visualization of Prediction Outcomes

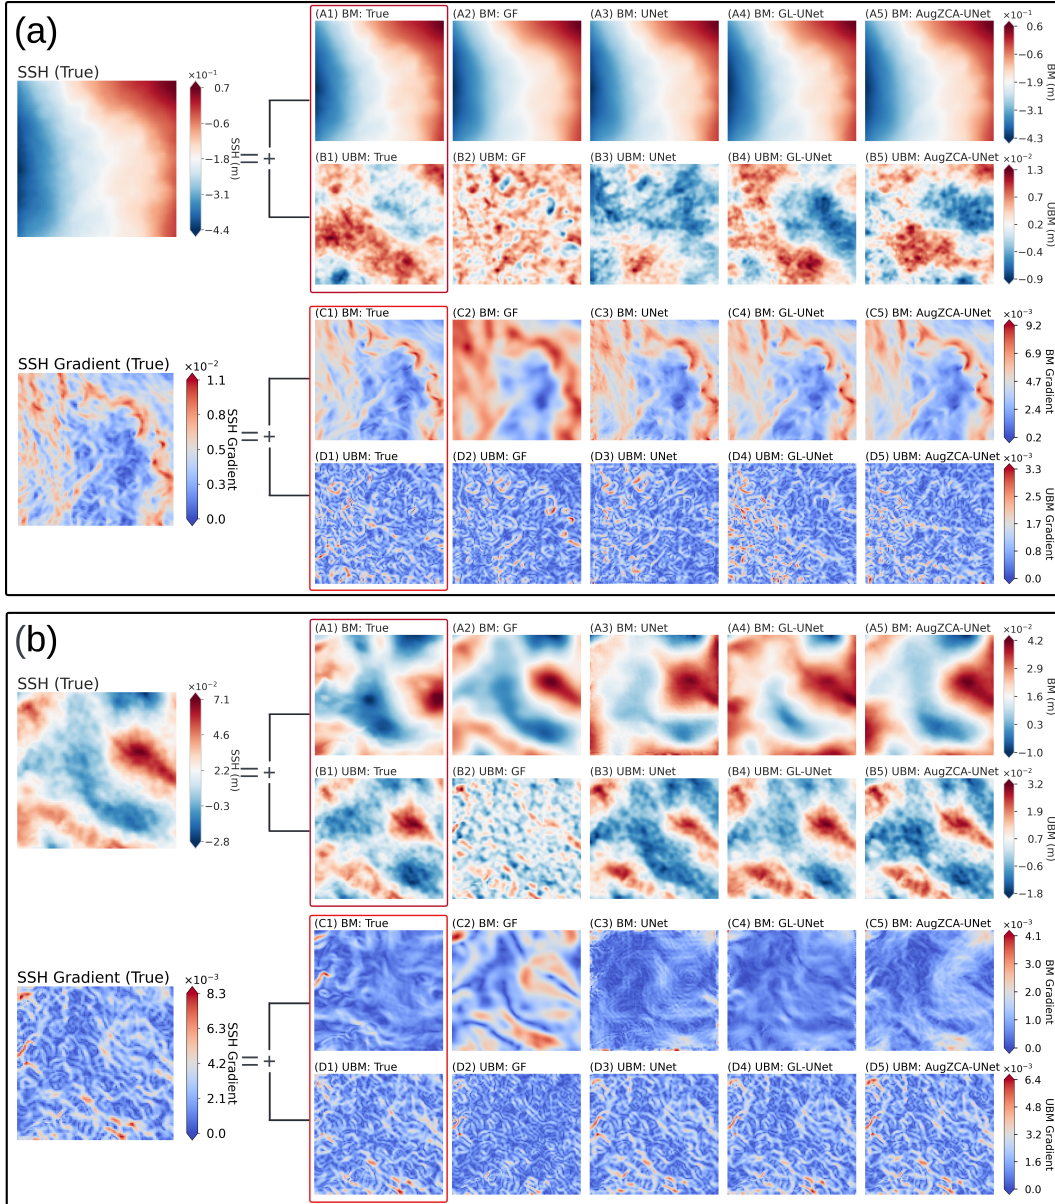


Figure App.3: Visualization of two test samples based on correlation coefficients between AugZCA-UNet predicted and true BM gradients: (a) Sample with median correlation (0.936); (b) Sample with lowest correlation (0.232). A1-A5: BM; B1-B5: UBM; C1-C5: BM Gradient; D1-D5: UBM Gradient.

F Gaussian Filter (GF)

Gaussian filters are fundamental tools in signal and image processing, prized for their isotropic and radially symmetric properties. They effectively attenuate high-frequency components while preserving low-frequency information, and have been used a lot in the filtering of gridded geophysical

data (Grooms et al., 2021). The filter is applied to an image or signal $f(x)$ through convolution, represented by the following integral:

$$(f * G)(x) = \int f(u) \cdot G(x - u) du$$

The GF approach, which acts as a spatial filter that effectively truncates the SSH power spectra at a specific wavenumber, effectively removing high-frequency components typically associated with UBM.

The scale of the filter is inferred from the power spectra of the total SSH, which often shows a change in spectral properties at scales where UBM start to dominate over BM (Torres et al., 2019). Here we used a single scale of 24km for the entire domain considered.

G Error Distribution

Table App.1 presents the pixel-wise absolute error distribution measures for $\hat{\eta}_B$ and $|\nabla\hat{\eta}_B|$, which are computed as follows:

$$\begin{aligned} \epsilon_{\eta_B} &= |\hat{\eta}_B - \eta_B| \\ \epsilon_{\nabla\eta_B} &= ||\nabla\hat{\eta}_B| - |\nabla\eta_B|| \end{aligned}$$

Table App.1: Pixel-wise Absolute Error Distribution Measures for $\hat{\eta}_B$ ($\times 10^{-2}$) and $|\nabla\hat{\eta}_B|$ ($\times 10^{-3}$)

Measures	GF	GL-UNet					Aug-UNet	ZCA-UNet	AugZCA-UNet	
		$\alpha = 0$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 5$	$\alpha = 10$				
$\hat{\eta}_B$	Mean	0.573	0.542	0.546	0.530	0.536	0.546	0.524	0.533	0.498
	Median	0.406	0.404	0.420	0.397	0.398	0.406	0.384	0.393	0.372
	Q1 (25%)	0.182	0.182	0.198	0.178	0.179	0.182	0.173	0.177	0.170
	Q3 (75%)	0.774	0.757	0.757	0.744	0.749	0.765	0.724	0.740	0.687
	P95 (95%)	1.68	1.52	1.49	1.48	1.51	1.53	1.50	1.52	1.40
$ \nabla\hat{\eta}_B $	Mean	0.795	0.656	0.626	0.636	0.640	0.642	0.651	0.634	0.604
	Median	0.505	0.464	0.447	0.450	0.450	0.453	0.460	0.450	0.427
	Q1 (25%)	0.229	0.213	0.205	0.207	0.206	0.208	0.211	0.207	0.196
	Q3 (75%)	0.947	0.849	0.811	0.819	0.824	0.828	0.841	0.816	0.775
	P95 (95%)	2.20	1.80	1.70	1.72	1.74	1.75	1.78	1.71	1.64

H Power Spectral Density (PSD) Analysis

Table App.2 presents the log-scale root mean square error (RMSE) of the Power Spectral Density (PSD), denoted as ϵ_{PSD} , calculated across wavenumbers for the mean PSD of test samples:

$$\epsilon_{\text{PSD}} = \sqrt{\frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} (\log \mathbb{E}[S_{\text{ref}}(k)] - \log \mathbb{E}[S_{\hat{\eta}}(k)])^2}$$

Here, $S_{\text{ref}}(k)$ and $S_{\hat{\eta}}(k)$ represent the true and predicted PSDs at wavenumber k , respectively, while $\mathbb{E}[\cdot]$ denotes the mean over test samples. This logarithmic approach allows for meaningful comparisons across the wide range of PSD magnitudes typical in spectral analysis.

To evaluate ML model performance at different scales, we calculate ϵ_{PSD} separately for small (< 0.2) and large (> 0.35) wavenumbers, corresponding to large-scale and small-scale features, respectively.

Table App.2: Log-Scale RMSE of PSD Across Wavenumber Ranges for ML Models

	GL-UNet					Aug-UNet	ZCA-UNet	AugZCA-UNet
	$\alpha = 0$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 5$	$\alpha = 10$			
<i>Small wavenumbers (< 0.2)</i>								
UBM	0.34	0.22	0.28	0.21	0.25	0.38	0.29	0.25
BM	0.41	0.67	0.61	0.69	0.58	0.39	0.27	0.71
<i>Large wavenumbers (> 0.35)</i>								
UBM	2.59	1.11	0.72	1.04	0.94	2.18	0.41	0.29
BM	2.00	2.48	1.82	2.41	2.11	2.20	0.43	0.76

I Gradient Loss

The loss function for the GL-UNet models involves the mean square error (MSE) of the UBM and its gradient:

$$\mathcal{L}(\eta_{\text{UB}}, \tilde{\eta}_{\text{UB}}) = \|\eta_{\text{UB}} - \tilde{\eta}_{\text{UB}}\|^2 + \alpha \|\nabla \eta_{\text{UB}} - \nabla \tilde{\eta}_{\text{UB}}\|^2$$

where η_{UB} is the ground truth UBM from training data set, $\tilde{\eta}_{\text{UB}}$ is the prediction from the network, and α is the weighting parameter. The choice of this loss function takes into account both the accuracy of the prediction and the preservation of high-frequency information for the UBM field.

J Scale-wise R^2

Figure 2 displays the averaged Power Spectral Density (PSD) for both UBM and BM in the first row, and the R^2 in the second row. The R^2 is calculated as follows:

$$R^2 = 1 - \frac{\text{MSE}(\text{Predicted PSDs})}{\text{Var}(\text{True PSDs})} = 1 - \frac{E[(S_{\text{ref}} - S_f)^2]}{E[S_{\text{ref}}^2]}$$

where S_{ref} and S_f are the true PSDs and predicted PSDs respectively. This formula reflects the relative quality of the estimation as a function of the wavenumber. Values of R^2 close to 1 indicate a perfect estimation.