MBARI Monterey Bay Aquarium Research Institute

Connecting the Surface to the Deep: Tracking Sinking Particle Fluxes Through a Machine Learning Approach in the California Current

Cate O'Bryon, UCLA

Mentors: Théo Picard and Monique Messié

Summer 2026

Keywords: Carbon Fluxes, Marine Snow, Lagrangian Particle Tracking, Machine Learning, Convolutional Neural Network

ABSTRACT

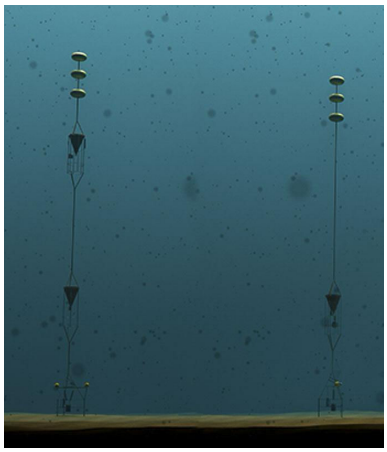
Sinking particle fluxes are essential for the transport of carbon within the biological carbon pump (BCP), however connecting particles collected at depth to their surface origins remains challenging due to the variability of physical advection. In the California Current, the surface origins of fluxes collected with sediment traps located at the deep abyssal Station M have yet to be thoroughly investigated. This study uses a machine learning approach through training a convolutional neural network (CNN) on modeled sea surface outputs and computational Lagrangian particle-tracking experiments to predict the surface origin areas of particles ending up at Station M. After training and validation, the CNN predicts a substantial amount of daily variability of particle origins when reproducing particle origin density maps from testing datasets. These findings support the use of machine learning methods in coordination with Lagrangian tracking of particle fluxes, and could ultimately be used with satellite-derived data to accurately predict particle origins from real-world sea surface conditions.

1 INTRODUCTION

Understanding the origins of sinking particles is essential to determine pathways of transport within the biological carbon pump (BCP). In the sunlit surface layer of the ocean, phytoplankton undergo photosynthesis to produce organic carbon, supporting higher trophic levels and diverse marine food webs. While a significant portion of this organic carbon is consumed near the surface, a central component of the BCP is the gravitational sinking of organic carbon particles from the surface to the deep ocean **[6]**. These sinking particles are commonly referred to as marine snow, and consist of decaying organic matter, including phytoplankton, zooplankton and fecal pellets.

As this organic matter descends through the water column, it can act as a food source, be remineralized into inorganic forms, or sequestered for long periods of time, influencing biological and chemical processes throughout the ocean. Through these functions, the BCP highlights key carbon transport pathways, and plays a critical role in supporting marine life. However, the mechanisms driving the BCP are not yet fully understood, and estimates of carbon export to the deep ocean range from approximately 5 to 15 Gt C yr^{-1} **[12]**. Additionally, with increased levels of atmospheric carbon dioxide from anthropogenic activity and climate change altering ocean temperature, chemistry, and circulation, the BCP may undergo significant change. Therefore, understanding the mechanisms and processes behind carbon transport is essential to assess the implications of these environmental changes on the BCP.

The study area for this work is centered on Station M (34°50' N, 123° W), an abyssal research station located off the coast of California. Through efforts by MBARI, physical data collection using moored sediment traps at 3,400 m was conducted at Station M for nearly 30 years **[14]**. Conical collecting devices were used to trap sediment in a preservative mixture, and the particulate organic carbon (POC) flux derived from the collected sediment data illustrates significant yearly variability (Fig. 1-3). These peaks may be associated with high primary production at the ocean surface; however, the specifics of regional productivity surrounding Station M cannot be resolved from collection at depth. One limitation to characterizing these deep carbon fluxes is the lack of known surface origins for collected particles. The surface source region of sinking particle fluxes, termed the "catchment area," has traditionally been assumed to encompass a circular area with a 100-km radius above a collection device at depth **[12]**. However, this assumption is oversimplified and does not account for advection due to surface dynamics. This study seeks to better characterize the spatial extent and daily variability of sinking particle catchment areas by accounting for the effects of advection.





Figures 1-2. Sediment trap mooring schematic and deployed traps, MBARI

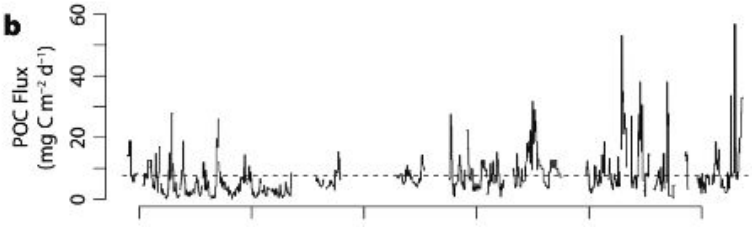


Figure 3. POC flux measurements at Station M (Smith et al. 2017)

This study uses a modeling-based approach to investigate the surface origins of sinking carbon particles in the California Current, as tracking the journey and evolution of these particles in the water column is important to understanding and quantifying large-scale carbon fluxes. In order to track particles as they are advected by ocean currents, a Lagrangian approach is adopted, which allows for observations to follow particle movement

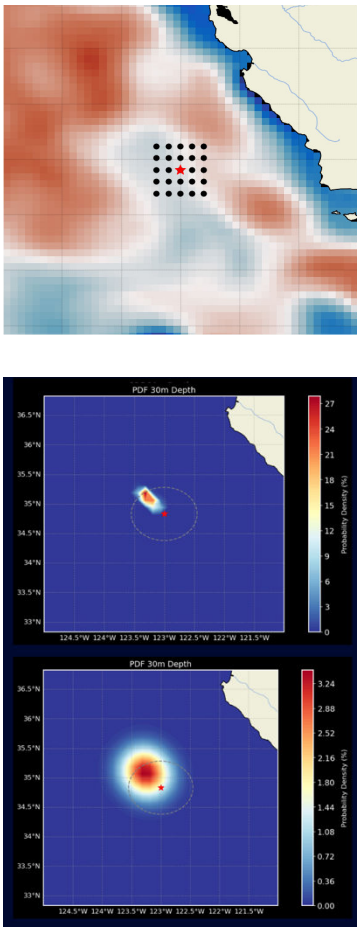
over time **[15]**. Lagrangian particle tracking is a computational method that aids in the analysis of circulation patterns on numerous scales, as well as the origins and destinations of particles. This approach was used in combination with a machine learning strategy to link particle flux origins to sea surface predictors using a convolutional neural network (CNN). Previous studies have demonstrated the ability of CNNs to learn spatial relationships through sea surface variables with less computational expense, with model performance depending strongly on the spatial extent of the input domain and architecture **[17]**. The workflow used in this study was inspired by the approach developed by Picard et al. (2024) in the North Atlantic at the Porcupine Abyssal Plain (PAP) Station, and this study investigates the application of this framework to the California Current Region **[10]**.

2 MATERIALS AND METHODS

The machine learning approach utilized in this study was adapted from the methodology of Picard et al. (2024) **[10].** The framework for the CNN was based upon a U-Net type of architecture, in which encoder layers learn to recognize important features of sea surface images, and decoder layers connect these features to spatial patterns in particle density maps **[11].** 30 years of modeled sea surface height data and particle trajectories were used for training, validation and testing of the model.

2.1 ROMS Simulation & Lagrangian Particle Tracking Scheme

The proposed machine learning approach required the creation of a dataset containing modeled sea surface conditions and corresponding sinking particle flux trajectories. In order to represent and simulate ocean conditions, outputs from a UCSC run Regional Ocean Modeling System (ROMS) model were extracted of the 8° x 8° region of interest, centered at the position of Station M. ROMS is a 3-D ocean circulation model that simulates ocean currents at a finely resolved spatiotemporal scale and contains multiple terrain-following



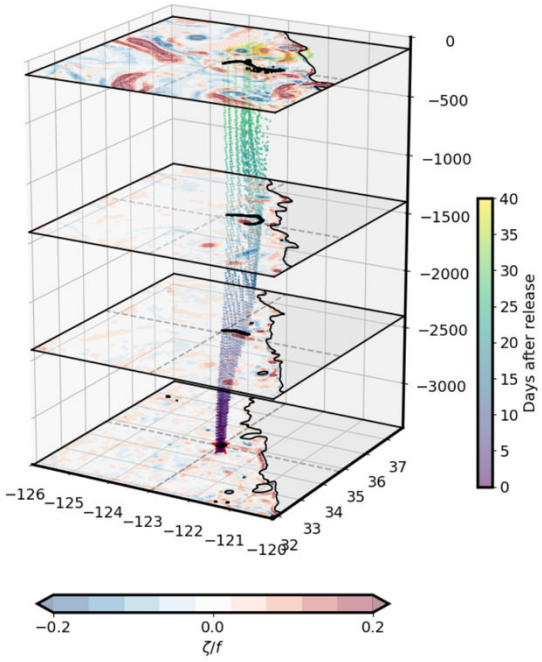


Figure 4-6 (clockwise starting top left). Station M, starred & 24 surrounding theoretical stations; depiction of Lagrangian particle tracking experiment originating at Station M (Picard, Pyticles); PDF before and after Gaussian filter application.

"sigma" depth levels **[13].** UCSC's 'WCR30' daily 3-km resolution ROMS model outputs with 40 sigma levels were used for sea surface height (SSH) data and zonal (U) and meridional (V) velocity fields.

Pyticles, a Python/Fortran hybrid parallelized code used for offline Lagrangian particle tracking, was used in conjunction with the ROMS U/V velocity fields to simulate sinking particle trajectories **[4]**. Particles were seeded at staggered levels around 3,400 m depth at theoretical Station M and backtracked to the surface with a constant sinking speed of 100 m/d, which represents dense, fast-sinking particle types (Fig. 5) **[16].** A 3-D "Runge-Kutta 4" scheme was used to advect the particles from depth to the surface, taking on average 30-40 days to complete their vertical passage. To create a more robust dataset of local dynamics, particles were additionally seeded at 24 other theoretical stations spaced 0.2° degrees apart in a 5x5 square surrounding Station M (Fig. 4). For each Lagrangian experiment, a 2-D snapshot of particle positions at 30 m depth was saved and converted into a probability density function (PDF) to represent the area of particle origination in the surface layer. A Gaussian filter was applied to each PDF to diffuse high densities of particles, and allow for the model to learn broader spatial patterns rather than focus on individual pixels (Fig. 6).

The training dataset consisted of PDF samples representing the target catchment area. Each sample was associated with 40 days of SSH data sampled at 5-day intervals over the period of particle sinking, which comprised the input to the network. Training data was extracted from the 30 years' dataset between January 1995 and December 2015; validation data was extracted from March 2016 to December 2019, and test data was extracted from March 2020 to December 2023. A two-month gap was intentionally added between training, validation and test sets to ensure independent samples and minimize overlap between the datasets. During training, data were processed in smaller batches of 32 samples. An additional input set up was created using U/V velocity fields at eight different depth levels along with sea surface height data to verify the model's performance. These depth levels were placed at every 500 m depth until 3,000 m, where an additional two layers at 3,200 and 3,400 m were included.

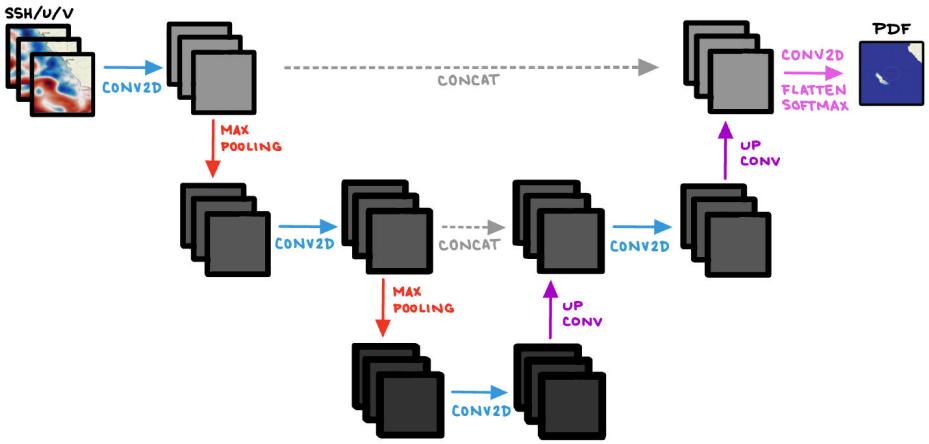


Figure 7. Schematic of U-Net neural network showing convolutional and pooling layers.

2.2 Machine Learning Approach

The CNN architecture chosen for this study was modeled after the U-Net, with batches of SSH images provided as input and paired PDF images provided as target outputs (Fig. 7). Both sets of images were converted to tensors with a spatial resolution of 64 x 64, and then

cropped to be centered on the station of interest, reducing the extent to 56 x 56 (approximately 0.1° x 0.1°). The contracting path of the U-Net, or encoder, used convolutional and pooling layers to increase the amount of identified feature maps and decrease the spatial resolution of the original SSH inputs. Within each convolutional layer, a 3x3 convolution, batch normalization and ReLU activation function were applied. The expanding path of the U-Net, or decoder, upsampled the compressed feature maps to produce an output PDF map. Up-convolutional layers complimentary to the encoder region were applied alongside skip connections with concatenation, and the last layer of the U-Net flattened and applied a softmax activation to the final feature maps to produce the predicted PDF. Dropout layers were additionally incorporated to reduce overfitting during training.

For training, a Bhattacharyya loss function and Wasserstein distance metric were used to evaluate the accuracy of the model in predicting particle PDFs. Bhattacharyya loss measures the distance between the predicted and actual particle PDF distributions via the calculation of the Bhattacharyya coefficient (Eq. 1) **[1].** A higher coefficient represents more overlap between the distributions. In order to express loss, the coefficient is subtracted from 1; therefore, a lower Bhattacharyya loss score is representative of a better matching predicted PDF (shown in Fig. 9). Wasserstein loss also measures distance between distributions through calculating how much mass must be moved to transform one distribution into the other (Eq. 2) **[2]**. Similar to Bhattacharyya loss, a lower Wasserstein loss score is representative of a more similar PDF.

$$D_B(P,Q) = -\ln\left(\int \sqrt{P(x)\cdot Q(x)}\,dx\right)$$

Equation 1. Bhattacharyya Coefficient

$$W_p(\mu,\nu) := \left(\inf_{\gamma\in\Gamma(\mu,\nu)}\int_{M\times M} d(x,y)^p\,\mathrm{d}\gamma(x,y)\right)^{1/p}$$

Equation 2. Wasserstein Distance Metric

An average loss value was calculated for the Bhattacharyya loss function by comparing all daily PDFs to the 20 years of training data (1995-2015) average PDF distribution. This average loss value was calculated to provide a baseline for prediction validity. The Bhattacharyya score and Wasserstein score were tracked for each epoch of training and validation to ensure a decreasing loss trend for both types of loss functions (Fig. 8), however only the Bhattacharyya loss function was used to train the model.

Training, validation and testing of the CNN were conducted through the PyTorch library. Training was conducted for 30 epochs with the Bhattacharyya loss function and the Adam optimization algorithm **[5]**. Additionally, a learning rate of 0.001 and a drop-out rate of 0.2 were used within the training process. The model weights from the epoch with the lowest validation loss were saved as a checkpoint and subsequently used to evaluate the model on the test data. Training was implemented using 4 different station configurations,

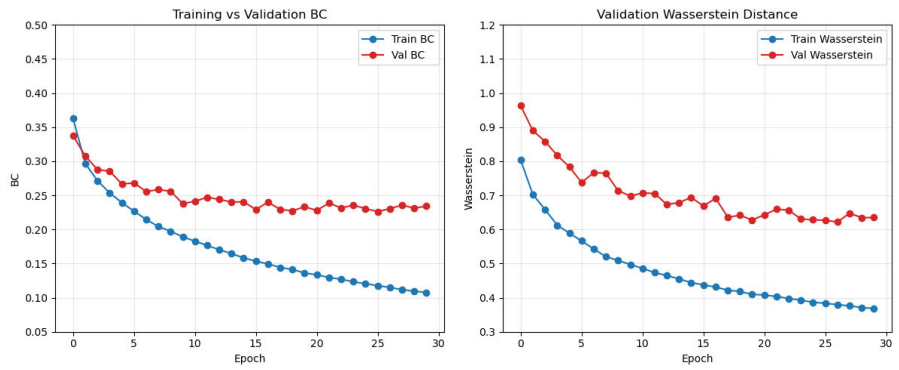


Figure 8. Decreasing Bhattacharyya and Wasserstein scores during example training and validation epochs.

consisting of particle releases from Station M alone, the immediate 8 stations surrounding Station M (9 total), an additional 8 outer stations (17 total), and all 25 stations. Testing was intentionally performed using PDF data from both 9 and 25 stations to assess the model's ability to predict particle distributions with different levels of spatial coverage. A bootstrapping method was additionally implemented by independently training the model 10 times for each particle release station configuration to account for randomness in the training process. The resulting PDF predictions were averaged during testing, and a final loss score was obtained.

3 RESULTS

The performance of each model configuration was evaluated against the previously described baseline Bhattacharyya loss score. The average Bhattacharyya loss score using 25 stations of PDF data was 0.437, and when using 9 stations, the score was 0.419. Testing scores below these specified baseline values were considered to be better predictions for each station configuration, and a Bhattacharyya score of under 0.3 corresponded to ~45% of PDFs being accurately predicted, representing a benchmark for stronger model performance. All model configurations resulted in scores below the 0.437 and 0.419 baseline for their respective configurations.

Configuration	**Bhattacharyya Loss** *(Test Dataset, 25 stations)*	**Bhattacharyya Loss** *(Test Dataset, 9 stations)*
Baseline	**0.437**	**0.419**
Station M, surface *767 experiments*	**0.417**	**0.397**
9 stations, surface *~7,000*	**0.372**	**0.332**
17 stations, surface *~13,000*	**0.313**	**0.311**
25 Stations, surface *~19,000*	**0.304**	**0.314**
25 stations, 8 depths *~19,000*	**0.189**	**0.170**

Table 1. Bhattacharyya scores for different station configurations

Table 1 shows a decreasing trend in Bhattacharyya loss scores during testing with the addition of more stations to the training configuration. Lower loss scores were obtained from the configurations with the largest amount of unique Lagrangian experiments (detailed underneath each configuration description). Test scores converge to approximately 0.3 for the maximum number of stations (25) using only surface data. With the addition of depth-specific U/V current maps to the 25-station configuration, the lowest loss scores of 0.189 (25-station testing) and 0.170 (9-station testing) were obtained.

4 DISCUSSION

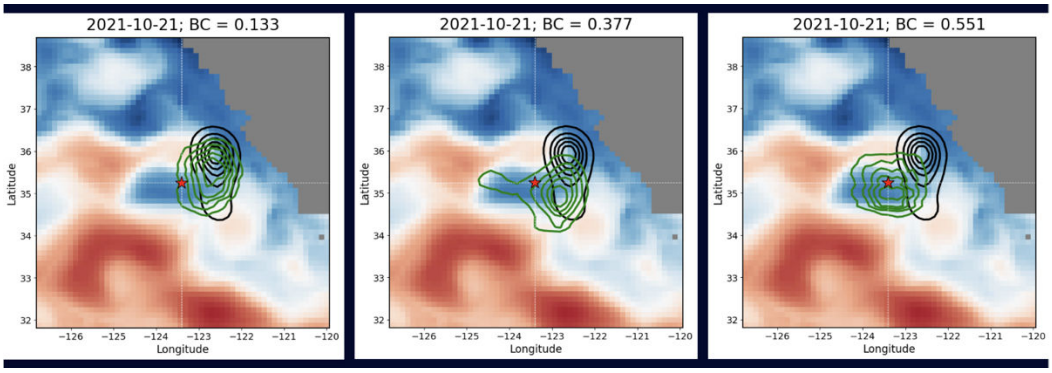


Figure 9. Increasing Bhattacharyya scores from left to right illustrate excellent, satisfactory, and poor model predictions (model depicted in green, true PDF depicted in black).

Taking sinking particle dynamics into account is essential to interpret deep carbon flux observations and understand the processes that influence carbon export. The results of this study illustrate effective retrieval of sinking particle source origins by training a CNN with SSH data and Lagrangian PDFs. This study demonstrates that the ability of the U-Net model to accurately predict surface catchment areas is dependent on the number of independent Lagrangian experiments and SSH/U/V input levels used in training. An increased number of Lagrangian experiments in training (implemented through the addition of 4 different station configurations) resulted in predictions closer to true PDFs, as seen through the decreasing trend observed in Bhattacharyya loss scores (Table 1). While all station configurations resulted in scores below the 0.4 baseline, it appears that using only SSH data as input was not a robust enough approach to achieve scores below 0.3, the second benchmark for performance where nearly half of PDFs are predicted accurately. Adding 8 depth levels of current information provided the most comprehensive view of dynamics impacting initial particle position, resulting in notably stronger model performance. This suggests that the higher loss scores observed for SSH-only configurations may be attributable to a lack of information at depth impacting the model's performance. However, this study concludes that a U-Net architecture used in coordination with ROMS surface outputs and modeled Lagrangian trajectories can predict a substantial amount of daily variability in surface catchment areas of particles collected at Station M.

This study serves as a validation and extension of the methodology developed by Picard et al. (2024) at the PAP Station in the North Atlantic by implementing a similar workflow to a new region of analysis **[10]**. The continuation of this study would involve the application of satellite sea surface products as input in place of ROMS. Following a similar approach used in Picard et al. (2025), satellite-derived SSH data would be input to a ROMS-trained model to predict particle catchment areas, which could then be combined with satellite-derived chlorophyll-a (*chl-a*) data to connect surface productivity to sinking fluxes **[9]**. While these findings are promising for the application of CNNs to connect particle fluxes to sea surface data, there are many limitations to an entirely modeling-based approach. Variations in particle size and speed are not accounted for within this methodology, and can introduce significant effects on particle trajectories in the water column. The use of a data-assimilated ROMS model could also have improved the actual daily accuracy of particle trajectories and

temporal patterns of transport pathways. Addressing these limitations is an important next step to improve model accuracy and the representation of real-world ocean conditions. Regardless, it is evident that the application of machine learning in combination with a Lagrangian perspective can greatly improve our understanding of particle fluxes, especially in diverse ecosystems such as the California Current. The ability to predict the origins of observed carbon fluxes from surface data may provide a valuable tool for elucidating the processes governing the BCP across spatiotemporal scales.

ACKNOWLEDGEMENTS

I would like to thank Théo Picard and Monique Messié for their guidance and support this summer, as well as George Matsumoto, Megan Bassett and Jessica Chapman for their organization of the MBARI intern program. The MBARI Summer Internship is generously supported by the David and Lucile Packard Foundation, with additional funding provided by Dalio Philanthropies and the Maxwell/Hanrahan Foundation.

REFERENCES

[1] Bhattacharyya, A. (1943). On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society, 35*, 99–110.

[2] Bonneel, N., Rabin, J., Peyré, G., & Pfister, H. (2015). Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision, 51*, 22–45. https://doi.org/10.1007/s10851-014-0532-4

[3] Breiman, L. (1996). Bagging predictors. *Machine Learning, 24*, 123–140. https://doi.org/10.1007/BF00058655

[4] Gula, J., & Collin, J. (2021). *Pyticles: A Python/Fortran hybrid parallelized code for 3D Lagrangian particles advection using ROMS/CROCO model data* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.4973786

[5] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*. https://doi.org/10.48550/arXiv.1412.6980

[6] Messié, M., Huffard, C. L., Stukel, M. R., & Ruhl, H. A. (2025). Spatial and temporal interplay between oceanic circulation and biological production in shaping carbon export off the California coast. *Geophysical Research Letters, 52*, e2024GL113707. https://doi.org/10.1029/2024GL113707

[7] Messié, M., Sancho-Gallegos, D. A., Fiechter, J., Santora, J. A., & Chavez, F. P. (2022). Satellite-based Lagrangian model reveals how upwelling and oceanic circulation shape krill hotspots in the California Current System. *Frontiers in Marine Science, 9*, 835813. https://doi.org/10.3389/fmars.2022.835813

[8] Nowicki, M., DeVries, T., & Siegel, D. A. (2022). Quantifying the carbon export and sequestration pathways of the ocean's biological carbon pump. *Global Biogeochemical Cycles, 36*(3), e2021GB007083. https://doi.org/10.1029/2021GB007083

[9] Picard, T., Baker, C. A., Gula, J., Fablet, R., Mémery, L., & Lampitt, R. (2025). Estimating the variability of deep-ocean particle flux collected by sediment traps using satellite data and machine learning. *Biogeosciences, 22*, 4309–4331. https://doi.org/10.5194/bg-22-4309-2025

[10] Picard, T., Gula, J., Fablet, R., Collin, J., & Mémery, L. (2024). Predicting particle catchment areas of deep-ocean sediment traps using machine learning. *Ocean Science, 20*, 1149–1165. https://doi.org/10.5194/os-20-1149-2024

[11] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015* (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28

[12] Ruhl, H. A., Bahr, F. L., Henson, S. A., Hosking, W. B., Espinola, B., Kahru, M., Daniel, P., Drake, P., & Edwards, C. A. (2020). Understanding the remote influences of ocean weather on the episodic pulses of particulate organic carbon flux. *Deep Sea Research Part II: Topical Studies in Oceanography, 173*, 104741. https://doi.org/10.1016/j.dsr2.2020.104741

[13] Shchepetkin, A. F., & McWilliams, J. C. (2005). The regional oceanic modeling system (ROMS): A split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modelling, 9*, 347–404. https://doi.org/10.1016/j.ocemod.2004.08.002

[14] Smith, K. L., Sherman, A. D., McGill, P. R., Henthorn, R. G., Ferreira, J., & Huffard, C. L. (2017). Evolution of monitoring an abyssal time-series station in the northeast Pacific over 28 years. *Oceanography, 30*(4), 72–81. https://doi.org/10.5670/oceanog.2017.421

[15] van Sebille, E., Griffies, S. M., Abernathey, R., Adams, T. P., Berloff, P., Biastoch, A., Blanke, B., Chassignet, E. P., Cheng, Y., Cotter, C. J., Deleersnijder, E., Döös, K., Drake, H. F., Drijfhout, S., Gary, S. F., Heemink, A. W., Kjellsson, J., Koszalka, I. M., Lange, M., ... Zika, J. D. (2018). Lagrangian ocean analysis: Fundamentals and practices. *Ocean Modelling, 121*, 49–75. https://doi.org/10.1016/j.ocemod.2017.11.008

[16] Wang, L., Gula, J., Collin, J., & Mémery, L. (2022). Effects of mesoscale dynamics on the path of fast-sinking particles to the deep ocean: A modeling study. *Journal of Geophysical Research: Oceans, 127*, e2022JC018799. https://doi.org/10.1029/2022JC018799

[17] Zavala-Romero, O., Bozec, A., Chassignet, E. P., & Miranda, J. R. (2025). Convolutional neural networks for sea surface data assimilation in operational ocean models: Test case in the Gulf of Mexico. *Ocean Science, 21*, 113–132. https://doi.org/10.5194/os-21-113-2025