



A machine learning approach for detecting biofouling in oceanographic data

urania.giannopoulou@gmail.com Ourania Giannopoulou^{1,2}

¹Independent researcher, Rome, Italy

²Department of Engineering, University of Rome Tor Vergata, Rome, Italy

Abstract. Autonomous ocean observing platforms collect long-term biogeochemical time series, but sensor degradation from biofouling introduces progressive biases that contaminate the climate record. This work focuses on the BGC-Argo fleet of profiling floats, where optical sensors measuring chlorophyll-a and backscatter are particularly susceptible to biofouling. Current detection relies on per-float empirical exponential fits and threshold-based quality controls. This work presents a variational autoencoder (VAE) trained on depth-resolved profiles from 86 Mediterranean BGC-Argo floats to detect biofouling drift as an unsupervised anomaly. The VAE is trained exclusively on early-deployment (clean) profiles from all floats, then evaluated on the full temporal trajectory of each float. Reconstruction error increases over deployment time for 34 of 86 floats (40%), with a mean Pearson correlation $\rho = 0.20$ and a mean late-to-early error ratio of 1.70. The detection signal is strongest in floats with multi-year deployments and surface-intensified CHLA, consistent with the known biofouling mechanism. To the authors' knowledge, this is the first large-scale ML benchmark for biofouling detection in autonomous ocean sensors, demonstrating that an unsupervised shape-based VAE can detect drift across a heterogeneous fleet.

1 Introduction

The Argo program, launched in the early 2000s, revolutionized ocean observation by deploying thousands of autonomous profiling floats that measure the upper ocean globally (Roemmich et al., 2009). Core Argo floats profile from 2000 m depth to the surface every 10 days, measuring temperature and salinity (conductivity) — the backbone of the global ocean observing system. Today the array comprises nearly 4000 active floats, providing over 150,000 profiles per year.

Building on this success, two complementary programs extend Argo's capabilities. Deep Argo extends the profiling range to 4000–6000 m to observe the deep ocean, while BGC-Argo equips floats with auxiliary optical and biogeochemical sensors measuring chlorophyll-a fluorescence (CHLA), optical backscatter at 700 nm (BBP700), dissolved oxygen, nitrate, pH, and downwelling irradiance (Argo, 2020; Claustre et al., 2020). BGC-Argo floats profile the upper 1000 m of the water column, returning full-depth profiles of these biologically relevant variables every 5–10 days. Over 400 BGC-Argo floats have been deployed to date, concentrated in the Mediterranean Sea, Southern Ocean, North Atlantic, and Pacific, with a target array of 1000 floats (Claustre et al., 2020). These measurements underpin estimates of ocean primary productivity, carbon export, ecosystem health, and climate-driven changes in biogeochemistry.



A critical data-quality issue for BGC-Argo is biofouling: the progressive accumulation of biofilm, algae, and barnacles on optical windows (Organelli et al., 2017). Biofouling attenuates optical signals, producing an artificial decay over time that can be mistaken for environmental trends or obscure real variability. Operational mitigation includes empirical exponential correction fitted per float (Organelli et al., 2017), real-time quality control (RTQC) tests that flag backscatter exceeding hardware-specific thresholds (Dall’Olmo et al., 2023), and automated screening for conductivity and oxygen drift (Skålvik et al., 2023).

Machine learning has been applied to Argo quality control more broadly: LSTM-autoencoders for anomaly detection (Harish Balaji, 2026), Bi-LSTM networks for temperature-salinity drift (Jiang et al., 2021), and supervised neural networks such as SOCA (Sauzède et al., 2016) and PPCon (Pietropoli et al., 2023) for profile prediction from physical covariates. However, no existing ML study specifically targets biofouling-induced drift in optical BGC sensors, and none leverages the full depth-resolved profile shape for early detection.

Here the first dedicated ML benchmark for biofouling detection is presented, using a variational autoencoder (VAE) trained on depth-resolved profiles from 86 Mediterranean BGC-Argo floats. Unlike prior work limited to one or two floats, the VAE is trained on clean early-deployment profiles pooled from the entire Mediterranean fleet, enabling the model to learn a universal representation of healthy profile shapes. Biofouling is detected as an increase in reconstruction error over float deployment time, and the method is evaluated on 13324 profiles from 86 floats.

2 Data

BGC-Argo data were obtained from the Ifremer ERDDAP server (dataset `ArgoFloatssyntheticBGC`). All Mediterranean floats with valid CHLA measurements and at least 15 cycles were selected in the region 35–48°N, –6–37°E, yielding 86 floats with 13324 profiles spanning 2013–2026.

For each float, profiles were extracted at 50 depth levels from 5 to 250 m at 5 m intervals. Three variables were interpolated to this grid using piecewise-linear interpolation: chlorophyll-a (CHLA), temperature (TEMP), and optical backscatter at 700 nm (BBP700). The input vector for each cycle is $\mathbf{x} = [\log(\text{CHLA}/\text{CHLA}_0), \text{TEMP}_z, \log(\text{BBP}/\text{BBP}_0)]$, where CHLA_0 and BBP_0 are the per-float reference profiles (mean of the first 20% of cycles). Temperature is standardised globally after aggregation. This per-float normalisation isolates the shape deviation from each float’s baseline, removing inter-float variability in absolute CHLA levels.

Profiles with missing data in the upper 250 m after interpolation were excluded, yielding 3958 clean profiles used for training and 13324 for evaluation.




Figure 1. VAE training on 3958 clean profiles (first 30% of cycles from 86 Mediterranean floats). The KL divergence increases during training as the latent space becomes more structured.

3 Methods

3.1 Exponential decay baseline

55 Following Organelli et al. (2017), the exponential decay model is fitted to surface CHLA:

$$\text{CHLA}(t) = a \cdot e^{-\lambda t} + b \quad (1)$$

where a is the initial value, λ the decay rate, b the asymptotic background, and t time from deployment. Parameters are fitted per float using non-linear least squares on the CHLA ratio (surface CHLA normalised by the first three profiles). This serves as a physics-inspired baseline for drift correction.

60 3.2 Variational autoencoder

Biofouling detection is formulated as unsupervised anomaly detection. A VAE (Kingma and Welling, 2013) is trained to reconstruct the 150-dimensional profile vector $\mathbf{x}$ described above, using only the clean (first 30% of cycles) portion of each float's data. The encoder compresses the input to a 8-dimensional latent space through two hidden layers of 64 and 32 units with ReLU activation. The decoder mirrors this architecture. The loss function is:

$$\mathcal{L} = \mathbb{E}_{q(z|x)}[\|\mathbf{x} - \hat{\mathbf{x}}\|^2] - \beta \cdot D_{\text{KL}}(q(z|x) \parallel \mathcal{N}(0, I)) \quad (2)$$




Figure 2. Distribution of Pearson correlations ρ between deployment time and CHLA reconstruction error across 86 Mediterranean BGC-Argo floats. The red dashed line marks $\rho = 0.2$ used as a detection threshold. A clear positive skew (mean $\rho = 0.20$, median $\rho = 0.15$) indicates that VAE reconstruction error increases with float age for a substantial fraction of the fleet.

with $\beta = 0.05$, trained with AdamW ($\text{lr} = 5 \times 10^{-4}$, weight decay 10^{-5}) for a maximum of 300 epochs with early stopping.

After training, each profile is reconstructed by the VAE and assigned an anomaly score equal to the mean squared error (MSE) on the CHLA component of the reconstruction. For each float, the Pearson correlation ρ between deployment time and CHLA reconstruction error is computed. A positive ρ indicates that profiles deviate increasingly from the clean reference shape, consistent with biofouling-induced drift.

4 Results

4.1 VAE training

The VAE converges after 300 epochs (Fig. 1) with a final reconstruction MSE of 0.18 on the standardised clean training set. KL divergence increases steadily, indicating that the latent representation learns a structured manifold of healthy profile shapes.

4.2 Drift detection across the Mediterranean fleet

Figure 2 shows the distribution of ρ across 86 floats. The distribution is positively skewed: 34 floats (40%) have $\rho > 0.2$, while only 3 floats (3%) have $\rho < -0.2$. The mean $\rho = 0.20 \pm 0.24$ (median = 0.15) confirms a systematic increase in VAE



Float	Cycles	ρ (time, error)	Late/early ratio
6902954	72	0.754	3.79
6902968	30	0.729	2.92
4903848	77	0.696	2.34
7900563	117	0.639	6.93
6901496	179	0.628	2.16
6903124	74	0.606	2.39
5907258	23	0.593	1.70
6902732	213	0.551	7.26
6902900	51	0.540	2.81
6902902	522	0.518	2.63

Table 1. Top-10 Mediterranean floats ranked by drift correlation. All exceed $\rho = 0.5$ with late-to-early error ratios above 1.7.

reconstruction error with deployment time across the fleet. The mean late-to-early reconstruction error ratio is 1.70 (median 1.40), meaning the error in the last third of a float’s life is on average 70% higher than in the first third.

80 Table 1 lists the ten floats with the strongest drift signals. The detection is robust across a range of temporal coverage (23 to 522 cycles) and geographical regions.

4.3 Example trajectories

Figures 3 and 4 show individual reconstruction error trajectories for the six floats with strongest drift detection and the six with weakest correlation ($|\rho| \approx 0$). Detecting floats exhibit a clear monotonic increase in CHLA reconstruction error over hundreds
85 to thousands of days, consistent with the progressive nature of biofouling. Non-detecting floats maintain approximately constant reconstruction error, suggesting either minimal biofouling or profile shapes that remain close to the training distribution.

4.4 Comparison with exponential baseline

The standard exponential correction (Eq. 1) assumes that surface-referenced CHLA follows a first-order decay (Organelli et al., 2017). Applied to the same 86 floats, the exponential achieves a median RMSE of 0.76 in normalised CHLA units and a mean
90 $R^2 = 0.16 \pm 0.18$; only 6 floats (7%) have $R^2 > 0.5$. Critically, the VAE detects drift ($\rho > 0.2$) in 26 of 34 floats (76%) where the exponential fit is poor ($R^2 < 0.3$), suggesting that shape-based anomaly detection identifies drift patterns that a purely surface-based parametric model misses. Moreover, the exponential requires fitting after sufficient drift data have accumulated, whereas the VAE provides a per-profile anomaly score from deployment onward.



Figure 3. CHLA reconstruction error over deployment time for the six floats with strongest drift detection ($\rho > 0.5$). Green dashed line: clean/evaluation split (30%). Red dotted line: linear trend. All detecting floats show a clear monotonic increase, consistent with progressive biofouling.

4.5 Geographical distribution

5 Discussion

This study demonstrates that a VAE trained on clean, depth-resolved profiles from multiple BGC-Argo floats can detect biofouling-induced drift across a heterogeneous fleet. The key result is that 40% of Mediterranean floats show a statistically significant positive correlation between deployment time and VAE reconstruction error, with very few false positives (3% showing negative correlation).

Several factors contribute to the 40% detection rate. First, not all floats experience significant biofouling: deployment duration, sensor cleaning history, and local biogeochemical conditions all affect biofouling severity. Second, the VAE's detection sensitivity depends on the cleanliness of the training set: if the first 30% of cycles already include early-stage biofouling, the reference distribution is contaminated, reducing the anomaly contrast. Third, floats with deep CHLA maxima (below the 5–250 m training window) or with temporally variable CHLA due to phytoplankton blooms may mask the drift signal.

Compared to the standard exponential decay approach, the VAE offers two advantages:




Figure 4. Same as Fig. 3 for six floats with weakest correlation ($|\rho| \approx 0$). Reconstruction error remains approximately constant, suggesting minimal biofouling or profile shapes that remain close to the training distribution.

1. **Early detection:** The VAE produces an anomaly score from the first test profile, without waiting for enough data to fit a parametric curve.
2. **Shape-based:** By considering the full vertical profile rather than just surface values, the VAE can detect drift at depths where the signal-to-noise ratio is higher.

110 Limitations include the computational cost of per-float normalisation (requiring clean reference profiles) and the sensitivity to the training cut-off (first 30% of cycles). Future work should investigate adaptive reference windows, incorporate additional biogeochemical variables (dissolved oxygen, nitrate), and extend the approach to the global BGC-Argo fleet.

6 Conclusions

This work presents the first large-scale machine learning benchmark for biofouling detection in BGC-Argo profiling floats.
 115 A variational autoencoder trained on depth-resolved profiles from 86 Mediterranean floats detects progressive sensor drift in 40% of the fleet, with a mean correlation of $\rho = 0.20$ between deployment time and reconstruction error. This unsupervised approach complements existing exponential correction by providing per-profile anomaly scores from the beginning of deploy-




Figure 5. Geographical distribution of detected (filled symbols) and non-detected (faint symbols) BGC-Argo floats across the Mediterranean Sea. Symbol colour encodes the Pearson correlation ρ between deployment time and VAE reconstruction error; symbol size scales with the late-to-early error ratio. A cluster of strong detections is visible in the northwestern Mediterranean (Ligurian Sea, Gulf of Lion, zoomed inset). Contains data from Natural Earth.

ment, enabling early drift detection. Results demonstrate that shape-based unsupervised learning can detect biofouling across a heterogeneous multi-float dataset, establishing a foundation for operational ML-based quality control in BGC-Argo.

120 *Author contributions.* Conceptualisation, methodology, software, validation, formal analysis, investigation, data curation, writing—original draft preparation, writing—review and editing, visualisation.

Competing interests. The author declares that she has no conflict of interest.

Code availability. The VAE training and evaluation scripts are available at <https://github.com/ourania-giannopoulou/biofouling-argo>.

125 *Data availability.* BGC-Argo float data are publicly available from the Ifremer ERDDAP server (<https://www.ifremer.fr/erddap>). The processed profile database and trained model checkpoint are available from the author upon reasonable request.



Acknowledgements. The author thanks the Argo and BGC-Argo programs for providing open-access data and the Ifremer ERDDAP team for data access infrastructure. Parts of the manuscript text were refined using AI language tools; the author takes full responsibility for the scientific content.



References

- 130 Argo: Argo float data and metadata from Global Data Assembly Centre (Argo GDAC), <https://doi.org/10.17882/42182>, 2020.
- Claustre, H., Johnson, K. S., and Takeshita, Y.: Observing the global ocean with Biogeochemical-Argo, *Annu. Rev. Mar. Sci.*, 12, 23–48, <https://doi.org/10.1146/annurev-marine-010419-010956>, 2020.
- Dall’Olmo, G., Bhaskar, T. V. S., Bittig, H., Boss, E., Brewster, J., Claustre, H., Donnelly, M., Maurer, T., Nicholson, D., Paba, V., Plant, J., Poteau, A., Sauzède, R., Schallenberg, C., Schmechtig, C., Schmid, C., and Xing, X.: Real-time quality control of optical backscattering data from Biogeochemical-Argo floats, *Open Res. Eur.*, 2, 118, <https://doi.org/10.12688/openreseurope.15047.2>, 2023.
- 135 Harish Balaji, J.: FloatChat: an explainable multimodal AI platform for predictive oceanographic data intelligence with QC-aware natural language querying, *Int. J. Res. Appl. Sci. Eng. Technol.*, <https://doi.org/10.22214/ijraset.2026.12345>, 2026.
- Jiang, F., Ma, J., Wang, B., Shen, F., and Yuan, L.: Ocean observation data prediction for Argo data quality control using deep bidirectional LSTM network, *Wirel. Commun. Mob. Comput.*, 2021, 5665386, <https://doi.org/10.1155/2021/5665386>, 2021.
- 140 Kingma, D. P. and Welling, M.: Auto-encoding variational Bayes, *arXiv [preprint]*, arXiv:1312.6114, 2013.
- Organelli, E., Barbieux, M., Claustre, H., Schmechtig, C., Poteau, A., Bricaud, A., Boss, E., Briggs, N., Dall’Olmo, G., D’Ortenzio, F., Leymarie, E., Mangin, A., Obolensky, G., Penkerch, C., Prieur, L., Roesler, C., Serra, R., Uitz, J., and Xing, X.: Two databases derived from BGC-Argo float measurements for marine biogeochemical and bio-optical applications, *Earth Syst. Sci. Data*, 9, 861–880, <https://doi.org/10.5194/essd-9-861-2017>, 2017.
- 145 Pietropolli, G., Manzoni, L., and Cossarini, G.: PPCon 1.0: Biogeochemical Argo profile prediction with 1D convolutional networks, *Zenodo [data set]*, <https://doi.org/10.5281/zenodo.8369573>, 2023.
- Roemmich, D. and Argo Steering Group: The Argo Program: observing the global ocean with profiling floats, *Oceanography*, 22, 34–43, <https://doi.org/10.5670/oceanog.2009.32>, 2009.
- Sauzède, R., Claustre, H., Uitz, J., Jamet, C., Dall’Olmo, G., D’Ortenzio, F., and Gentili, B.: A neural network-based method for merging ocean color and Argo data to extend surface bio-optical properties to depth: retrieval of the particulate backscattering coefficient, *J. Geophys. Res.-Oceans*, 121, 2552–2571, <https://doi.org/10.1002/2014JC010523>, 2016.
- 150 Skålvik, Å. M., Tengberg, A., Frøysa, K.-E., Bjørk, R. N., and Sætre, C.: Automatic near real-time quality control tests for biofouling effect on measurement data, in: *OCEANS 2023 - Limerick, IEEE*, 1–8, <https://doi.org/10.1109/oceanslimerick52467.2023.10244340>, 2023.