https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

FourierBGC 1.0: Fourier feature encoding of spatiotemporal inputs for biogeochemical Argo profile reconstruction

Samuel B. Phillips[1] and Xuyang Li[2]

[1]Department of Computer Science, University of North Carolina at Charlotte, Charlotte, NC, USA
[2]School of Construction, University of North Carolina at Charlotte, Charlotte, NC, USA

Correspondence: Samuel B. Phillips (sphill86@charlotte.edu)

Abstract. Reconstructing biogeochemical Argo profiles is essential to bridge the severe spatiotemporal coverage gaps between abundant physical observations and scarce biogeochemical data, providing the continuous ocean-state representations required for climate models and global carbon budget analyses. A recent convolutional method for profile reconstruction, PPCon, encodes day-of-year, year, latitude, and longitude through four independently learned multilayer perceptrons. These learned encoders add substantial parameter overhead and may yield sub-optimal spatiotemporal representations for the variables being reconstructed. To address this, we introduce FourierBGC, which replaces PPCon's learned encoders with a fixed Fourier-feature encoding of day-of-year, latitude, and longitude, combined with a lightweight learned projection, cutting encoder parameters by 88 %. By directly embedding periodic annual cycles and smooth spatial continuity into input space, this encoding provides strong multifrequency coordinate representations that facilitate lightweight yet highly expressive spatiotemporal representations. Through a controlled ablation holding PPCon's convolutional backbone fixed, FourierBGC outperforms PPCon on all three reconstructed variables (nitrate, chlorophyll-*a*, and bbp700), and the Fourier encoding achieves the lowest reconstruction error in every sub-basin and season. Disentangling the sources of error reduction reveals that beyond an optimized training procedure, the Fourier encoding yields a consistent, highly parameter-efficient performance boost across all variables. Furthermore, the learned projection's role exhibits variable- and season-specific behavior tied to upper-ocean physics: it is dispensable on nitrate, essential for capturing the complex profile of chlorophyll-*a*, and on bbp700 helps under summer stratification while costing accuracy in winter, a pattern consistent with the projection expressing vertical water-column structure only when it is present. Overall, FourierBGC demonstrates that structuring spatiotemporal inputs with explicit frequency priors offers a parameter-efficient strategy for reconstructing sparse geophysical and environmental fields.

1 Introduction

Argo is a global array of nearly 4,000 autonomous floats that profile temperature and salinity throughout the water column (Roemmich et al., 2019); a smaller Biogeochemical-Argo (BGC-Argo) extension equips a subset of these floats with additional sensors for oxygen, nitrate, chlorophyll, and particulate backscatter (Biogeochemical-Argo Planning Group, 2016). These biogeochemical variables underpin the observations on which marine ecosystems research depends: ecosystems that sustain

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

fisheries, regulate climate, and absorb a substantial share of anthropogenic carbon (Campbell et al., 2016), with nitrate and chlorophyll serving as core state variables in ocean biogeochemical models (Teruzzi et al., 2021).

Despite this importance, BGC-Argo coverage is severely uneven. Oxygen, the most mature and widely deployed biogeochemical sensor, has historically been measured several times more frequently than nitrate or chlorophyll (Wang et al., 2019), a gap compounded by the difficulty of maintaining additional sensors in a physically harsh, remote environment (Chai et al., 2020) and by limited long-term funding for sustained observation (Weller et al., 2019).

This asymmetry motivates a growing body of work that reconstructs the less-observed biogeochemical variables from more widely available physical and oxygen measurements. Reconstructed profiles are increasingly assimilated to constrain estimates of ocean productivity and nutrient cycling, and have been shown to reduce systematic bias in operational forecasting systems (Amadio et al., 2024).

Early reconstruction efforts focused predominantly on nutrient and carbonate-system variables rather than bio-optical properties. These methods were largely regression-based; for example, locally interpolated regressions (LIRs) estimated nitrate from temperature, salinity, and oxygen by fitting linear coefficients across spatial grids (Carter et al., 2018).

To better capture non-linear relationships, neural network models were introduced, beginning with CANYON (Sauzède et al., 2017). CANYON used multilayer perceptrons (MLPs) that directly ingested spatiotemporal coordinates (latitude, longitude, depth, and day of year) alongside physical properties (temperature, salinity, and oxygen) to estimate nutrient and carbonate variables. Subsequent variants, such as the Bayesian network CANYON-B (Bittig et al., 2018) and the Mediterranean-specific CANYON-MED (Fourrier et al., 2020), improved predictive robustness and regional accuracy while preserving this MLP-based direct-coordinate input strategy.

While Pietropolli et al. (2023) extended these MLP frameworks to chlorophyll-*a* using hand-crafted coordinate transforms, most of these MLP-based models process each depth level independently. This lack of vertical continuity inevitably yields noisy, discontinuous reconstructed profiles when applied to Argo floats.

To address this, PPCon (Pietropolli et al., 2024) introduced a 1D convolutional architecture operating on the entire profile as a vector, while extending targets to particulate backscattering (bbp700). However, PPCon replaced fixed coordinate transforms with four independently learned MLPs for day-of-year, year, latitude, and longitude. These encoders introduce substantial parameter overhead (158,800 parameters) and force the network to discover basic periodic and spatial continuity from scratch.

To resolve this limitation, we introduce FourierBGC, which retains PPCon's convolutional backbone while replacing its four independently learned encoders with a fixed, structured representation (Fig. 2). Day-of-year, latitude, and longitude are first expanded into sinusoidal components at dyadic frequencies, producing an 18-dimensional representation that embeds the periodicity of day-of-year and the smooth spatial continuity of latitude and longitude directly into the input space. These 18 features are projected jointly through a single MLP into a depth-varying channel; meanwhile, year, which is neither periodic over the study period nor spatially structured in the same sense, is projected separately through its own small MLP into a second channel. Together, these two projections add 19,152 parameters, an 88 % reduction from the 158,800 PPCon spends on four separately parameterized encoders. By replacing freely parameterized, brute-force learned transforms with explicit frequency priors, FourierBGC enables lightweight yet highly expressive spatiotemporal conditioning for profile reconstruction.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Comparing FourierBGC to PPCon alone would not isolate the encoding, since the two also differ in training procedure. We therefore evaluate a spine of models that share the same convolutional backbone and differ only in how the spatiotemporal coordinates are represented, spanning the design space from omitting them entirely, to supplying them as raw scalars, to PPCon's four independent learned encoders, to FourierBGC's shared structured encoding. Every model in that spine except PPCon itself is trained with a single, common procedure, so the encodings can be compared without the training recipe varying alongside them.

This separation yields three results. First, a substantial part of PPCon's reported error is attributable to its training configuration rather than to its encoder: the same four encoders, trained under the common procedure, improve on the published model on all three variables – nitrate, chlorophyll-*a*, and bbp700. Second, with the procedure held fixed, the Fourier encoding improves further on those encoders, substantially on nitrate and more modestly on chlorophyll-*a* and bbp700, at a fraction of their parameter count. On nitrate the effect does not require a learned encoder at all: a parameter-free variant of FourierBGC already outperforms PPCon's 158,800-parameter encoders, indicating that how the spatiotemporal coordinates are represented matters more than how much capacity is spent representing them. Third, the value of projecting that fixed encoding through a further learned step, rather than passing it unchanged to the backbone, is not uniform across variables: decisive on chlorophyll-*a*, dispensable on nitrate, and on bbp700 a net cost whose sign reverses between summer and winter. Taken together, some form of the Fourier encoding records the lowest error of any model in the spine on all three variables and in every sub-basin and season, though which form – fixed alone or fixed plus projection – differs by variable.

2 Data

The underlying dataset originates from the BGC-Argo program (Bittig et al., 2019), consisting of BGC-Argo float profiles for the Mediterranean Sea retrieved from the Coriolis Argo GDAC. The published splits used here span 2012–2020, with 2012 contributing only 13 profiles per target variable; this range differs from the 2013–2022 of the source dataset's title, and we inherit PPCon's partition as published. Amadio et al. (2023) applied a quality-control procedure to this data, retaining only profiles meeting specified quality-flag criteria for temperature, salinity, nitrate, oxygen, and chlorophyll, with additional checks specific to bbp700, and standardized the vertical resolution to 2 m with light smoothing. We inherit this quality-checked dataset unmodified via Pietropolli et al. (2024)'s published PPCon repository, including their published train and test splits, rather than reprocessing the raw Argo archive ourselves. QC methodology is described in Amadio et al. (2023) and Pietropolli et al. (2024).

Input variables are temperature, salinity, and dissolved oxygen profiles, each interpolated onto a fixed 200-point depth grid: 0–1000 m at 5 m resolution for nitrate, and 0–200 m at 1 m resolution for chlorophyll-*a* and bbp700. Target variables are nitrate, chlorophyll-*a*, and bbp700, reconstructed on the same depth grid as their respective inputs. Four spatiotemporal coordinates accompany each profile: day-of-year, year, latitude, and longitude.

Following PPCon's variable definitions, nitrate is expressed in mmol m^{-3}, chlorophyll-*a* in mg m^{-3}, and bbp700 in units of 10^{-4} m^{-1}. For context, observational uncertainty for BGC-Argo sensors has been estimated at approximately 0.25 mmol m^{-3}

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 1. Distinct profile counts per split and target variable, inherited unmodified from PPCon's published dataset. The Removed column reflects the number of *profiles* excluded during PPCon's quality-control step for that variable.

Variable	Train	Test	Removed
Nitrate	2497	625	0
Chlorophyll-*a*	3773	944	29
bbp700	3785	947	15

PPCon's loader counts each CSV's variable-name column as a profile, so evaluation iterates one duplicate per split; the 626/945/948 denominators quoted below reflect this.

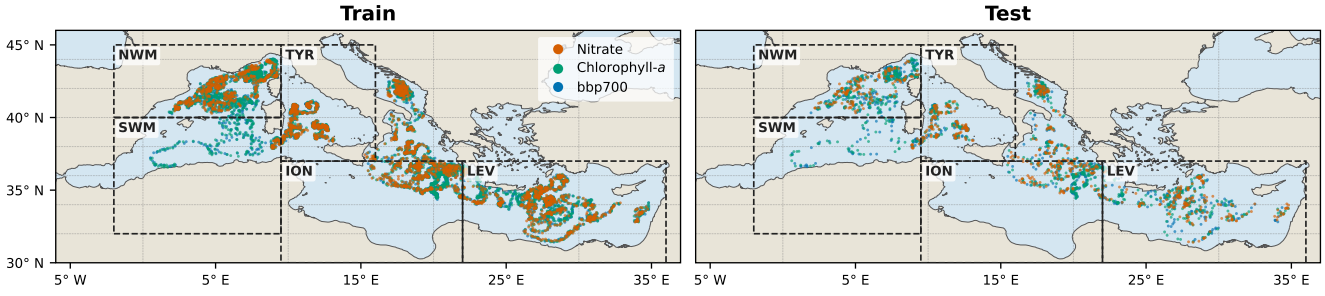


Figure 1. Spatial distribution of Mediterranean BGC-Argo profiles used in this study, inherited unmodified from PPCon's published train (left) and test (right) splits. A small random spatial jitter is applied so overlapping profiles remain visually distinguishable. Dashed boxes and labels indicate the five sub-regions used for regional evaluation in Sect. 5, following PPCon's own definitions: northwestern Mediterranean Sea (NWM), southwestern Mediterranean Sea (SWM), Tyrrhenian Sea (TYR), Ionian and southern Adriatic Sea (ION), and Levantine Sea (LEV).

for nitrate and 0.03 mg m^{-3} for chlorophyll-*a* (Mignot et al., 2019), providing a rough noise floor against which the RMSE improvements reported in Sect. 5 can be interpreted.

Regional evaluation follows PPCon's published sub-region definitions, implemented identically to their released evaluation code. We note a discrepancy between the boundaries reported in PPCon's Table 4 and those in their released evaluation code: Table 4 reports the Ionian Sea region as spanning 30–45 °N, while the code defines it as 30–37 °N. We follow the code's definitions throughout, since these are the boundaries that generated PPCon's published results. These five regions do not span the full extent of the Mediterranean test set; profiles falling outside all five regions are excluded from regional breakdowns in both PPCon's reported results and our own, so regional comparisons in Sect. 5 remain directly comparable (90/626 nitrate profiles, 116/945 chlorophyll-*a* profiles, and 117/948 bbp700 profiles fall outside all five regions).

Seasonal evaluation follows PPCon's published season definitions (Winter, Spring, Summer, Fall), implemented identically to their released evaluation code. PPCon's season boundaries are defined as integer cutoffs on day-of-year, while day-of-year itself is a continuous float; profiles whose value falls exactly between two integer bounds (for example, 91.99999961935734,

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

between Winter's upper bound of 91 and Spring's lower bound of 92) are silently excluded from both adjacent seasons by PPCon's own bucketing code. This is a property of PPCon's own evaluation code, not something introduced by this study, and applies identically to PPCon's own published seasonal results. It excludes a small number of profiles from seasonal breakdowns specifically, distinct from the regional exclusion described above: 4/626 nitrate profiles, 18/945 chlorophyll-*a* profiles, and 16/948 bbp700 profiles. Because their routine accumulates a profile only when it falls inside both a season and a region, PPCon's seasonal values are computed over the profiles common to both partitions, whereas ours are bucketed by season alone; placing PPCon on the season-only footing shifts its seasonal values by up to 13 % and reverses no ordering reported in Sect. 5.

3 Methods

3.1 Problem Setup and Notation

Bold symbols denote vector- or matrix-valued quantities and italic symbols scalars; a tilde denotes a normalized quantity and a hat a model estimate.

Every model in this study reconstructs the same target from the same inputs; they differ only in how the spatiotemporal coordinates are represented and where those representations enter the network. Each sample comprises three co-located input profiles, temperature, salinity, and dissolved oxygen, interpolated onto a fixed 200-point depth grid,

$$\mathbf{T},\ \mathbf{S},\ \mathbf{O}_2 \in \mathbb{R}^{200}, \tag{1}$$

together with four spatiotemporal coordinates, recorded once per profile rather than at each depth,

$$\begin{aligned} d &\in [0,\ 2\pi) && \text{day-of-year, in radians,} \\ t &\in \{2012,\ \ldots,\ 2020\} && \text{year,} \\ \varphi &\approx [31.5,\ 44.0]\,\mathrm{N} && \text{latitude,} \\ \lambda &\approx [0.6,\ 34.8]\,\mathrm{E} && \text{longitude,} \end{aligned} \tag{2}$$

where the latitude and longitude intervals are the observed extent of the dataset, rounded, rather than fixed bounds of the coordinate system or the normalization constants of Sect. 4. Collectively we write these spatiotemporal coordinates as $\mathbf{c} = (d,\ t,\ \varphi,\ \lambda)$ and the input profiles as $\mathbf{P} = (\mathbf{T}, \mathbf{S}, \mathbf{O}_2)$. The target profile $\mathbf{y} \in \mathbb{R}^{200}$ is one of nitrate, chlorophyll-*a*, or bbp700, on the depth grid corresponding to that variable. A separate model is trained for each target, so every architecture in the ablation is instantiated three times and evaluated independently per variable.

Each model assembles an input tensor $\mathbf{X} \in \mathbb{R}^{200 \times C}$, where C is the number of input channels, by concatenating the three profile channels with zero or more channels derived from $\mathbf{c}$; the shared convolutional backbone (Sect. 3.2) maps $\mathbf{X}$ to a reconstruction $\hat{\mathbf{y}} \in \mathbb{R}^{200}$. What occupies those additional channels, and at what depth in the network they are joined, is the axis of variation this study is concerned with (Table 2). One further difference is unavoidable: PPCon and PPCon-NoCoord are

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

trained with PPCon's own procedure while every model we train uses a single common one, so the training procedure varies between those two groups. Sect. 4.5 describes both, and Sect. 6 separates their effects from those of the encoding.

No normalization is applied at the dataset level: the profiles and coordinates are read as recorded, and $\mathbf{T}$, $\mathbf{S}$, and $\mathbf{O}_2$ enter the backbone unscaled in every model, following PPCon's original implementation; batch normalization within the backbone absorbs input scale from the first convolutional layer onward. Normalization of the spatiotemporal coordinates, where a model applies it, is part of that model's encoding and is described alongside it in Sect. 4.

3.2 PPCon Architecture

The convolutional backbone shared across this ablation is adopted unchanged from PPCon (Pietropolli et al., 2024): nine layers of 1D convolutions and transposed convolutions, each followed by SELU activation, batch normalization, and dropout, except the final projection layer, which maps the last hidden state directly to the single output channel. We follow the released code here; Pietropolli et al. (2024)'s Table 2 lists batch normalization and a SELU activation on that final layer, which the code does not apply. Kernel sizes, strides, and padding are given in the same table and are not reproduced here.

Before being passed to the convolutional network, PPCon gives every spatiotemporal coordinate its own encoder. Each of d, t, φ, and λ passes through a separate three-layer perceptron ($1 \to 80 \to 140 \to 200$, SELU), producing one depth-resolved channel per coordinate, so that

$$\mathbf{X} = [\mathbf{T},\, \mathbf{S},\, \mathbf{O}_2,\, \mathrm{MLP}_d(d),\, \mathrm{MLP}_t(t),\, \mathrm{MLP}_\varphi(\varphi),\, \mathrm{MLP}_\lambda(\lambda)] \in \mathbb{R}^{200\times 7}. \tag{3}$$

The four encoders are structurally identical, since Pietropolli et al. (2024) report no prior reason to treat the inputs differently, and each contributes 39,700 parameters, 158,800 in total. They receive their inputs unnormalized: PPCon applies no scaling at training or evaluation time, so the encoders operate on raw quantities such as $t = 2016$ and $\varphi = 35.29$.

Two aspects of PPCon are separable, and this study treats them separately. The first is the encoder architecture just described. The second is the procedure used to fit it, which differs from the one we use for the models we train ourselves in optimizer, learning-rate schedule, gradient clipping, and loss (Sect. 4.5). Because a comparison against PPCon alone would confound the two, the ablation varies each independently: the encoder of Eq. (3) reappears in Sect. 3.3 trained under the common procedure, so that its contribution can be assessed without the training procedure changing alongside it.

We use PPCon's original published checkpoint, obtained from the authors' released repository, as the external baseline against which every other model in this ablation is compared.

3.3 Models Overview

Every model in this study shares the backbone of Sect. 3.2 and differs only in how the spatiotemporal coordinates are represented and at what depth in the network they are joined. We refer to these models collectively as the *ablation spine*, because they form a chain in which consecutive pairs differ by a single design decision: removing the coordinate encoders, changing the training procedure, reintroducing the coordinates as raw values, replacing those with PPCon's learned per-coordinate encoders,

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

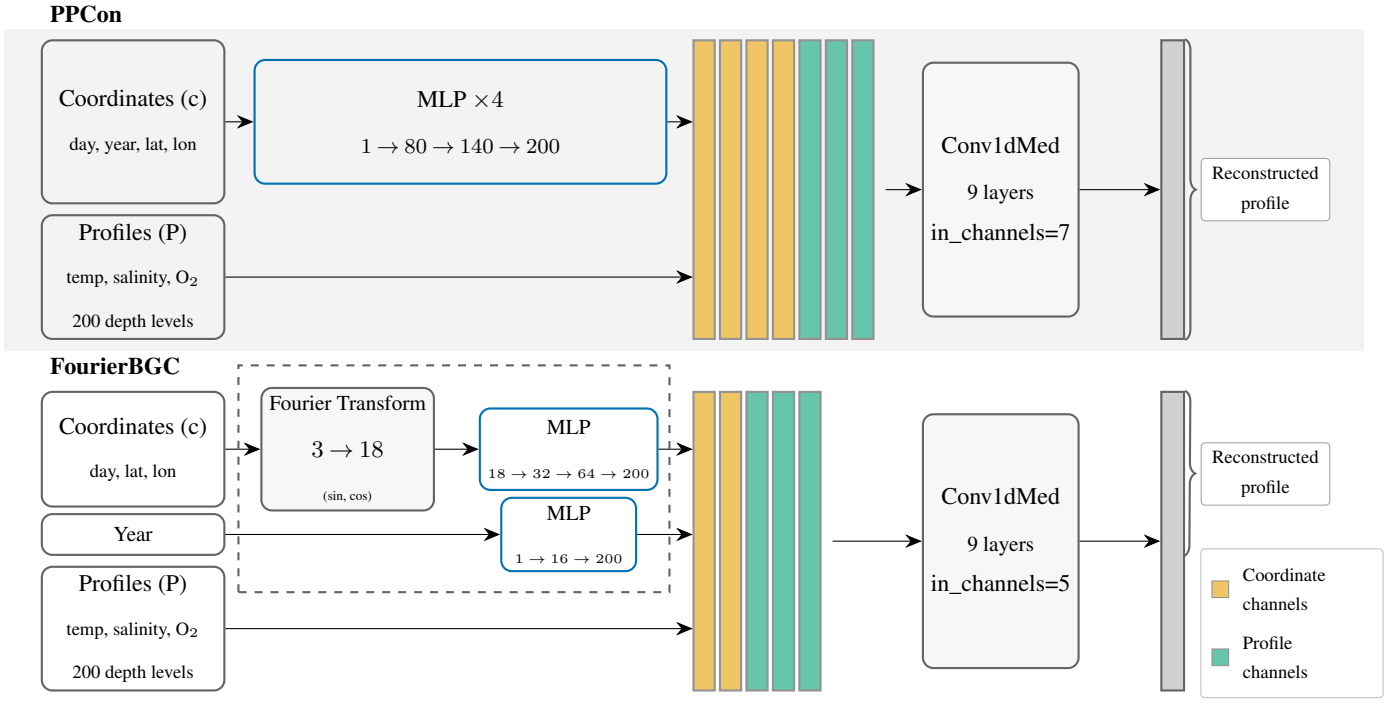


Figure 2. Architecture comparison between PPCon and FourierBGC, holding the convolutional backbone (Conv1dMed) fixed. The dashed box encloses FourierBGC's coordinate encoder, the component this study varies; Fig. 3 expands it in full.

replacing those in turn with a fixed Fourier encoding, and finally adding a learned projection on top of it. Figure 2 shows the two endpoints of that chain, PPCon and FourierBGC, in full.

PPCon is the external baseline, evaluated from the authors' published checkpoint without retraining ($C = 7$).

PPCon-NoCoord removes all four encoders and puts nothing in their place ($C = 3$), keeping PPCon's own training recipe. It converges poorly, so we read it as a control on that recipe rather than as a test of whether the coordinates matter.

CNN-NoCoord is the same architecture under the common recipe (Sect. 4.5). Its comparison with PPCon-NoCoord isolates the recipe, and it is the model that tests whether the coordinates are needed at all.

CNN-RawCoord reintroduces the four coordinates as raw z-scored values, joined late, immediately before the final projection rather than at the input. It tests whether making the coordinates available, untransformed, is sufficient.

CNN-MLPCoord restores PPCon's four learned encoders, Eq. (3), under the common recipe ($C = 7$). It is the model that makes the central comparison of this study unconfounded: against PPCon it isolates the training procedure with the encoder held fixed, and against FourierBGC it isolates the encoding with the procedure held fixed.

FourierBGC-Broadcast replaces the raw values with the fixed Fourier encoding of day-of-year, latitude, and longitude, plus the z-scored year, broadcast across depth with no learned parameters ($C = 22$).

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 2. Spatiotemporal coordinate representation and fusion point across the ablation spine.

Model	Day/lat/lon	Year	Fusion point	C	Encoder params
PPCon	3 separate MLPs (1 per input)	1 MLP	Input (early)	7	158,800
PPCon-NoCoord	None	None	—	3	0
CNN-NoCoord	None	None	—	3	0
CNN-RawCoord	Raw	Raw	Pre-output (late)	3	0
CNN-MLPCoord	3 separate MLPs (1 per input)	1 MLP	Input (early)	7	158,800
FourierBGC-Broadcast	Fourier	Raw	Input (early)	22	0
FourierBGC	Fourier + 1 shared MLP	1 MLP	Input (early)	5	19,152

C is the backbone input channel count. Two pairs share an encoder and differ only in training procedure: PPCon-NoCoord with CNN-NoCoord, and PPCon with CNN-MLPCoord (Sect. 4.5).

FourierBGC projects the same fixed encoding through a small learned network, with year projected separately ($C = 5$). Its comparison with FourierBGC-Broadcast isolates the learned projection; its comparison with CNN-MLPCoord is the headline contrast of this study.

This section describes the three encoding strategies compared in this ablation in more depth: the learned MLP encoding used by PPCon and CNN-MLPCoord, CNN-RawCoord's raw encoding, and FourierBGC's Fourier-feature encoding, of which FourierBGC-Broadcast is a variant. PPCon-NoCoord and CNN-NoCoord use no coordinate representation at all and are not discussed further here.

4 Coordinate Encoding Strategies

4.1 Learned MLP encoding (PPCon, CNN-MLPCoord)

PPCon represents each spatiotemporal coordinate with its own MLP, allowing the network to learn how much weight each should receive at each depth level rather than fixing that relationship in advance (Pietropolli et al., 2024). All four encoders share an identical architecture, since Pietropolli et al. (2024) report no prior reason to treat the inputs differently. This is a fully learned, dense encoding: nothing about periodicity, seasonality, or geographic structure is assumed or imposed, the coordinates are passed in unnormalized, and the entire representation, 158,800 parameters, is discovered during training. CNN-MLPCoord uses this encoder unchanged, including the unnormalized inputs, so that the two models differ only in training procedure and the encoding can be compared with that procedure held fixed.

4.2 CNN-RawCoord: raw encoding

CNN-RawCoord represents all four coordinates as their raw values, z-scored with fixed per-coordinate statistics computed on the pooled training splits, with no learned or nonlinear transform applied. Rather than concatenating them at the input

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

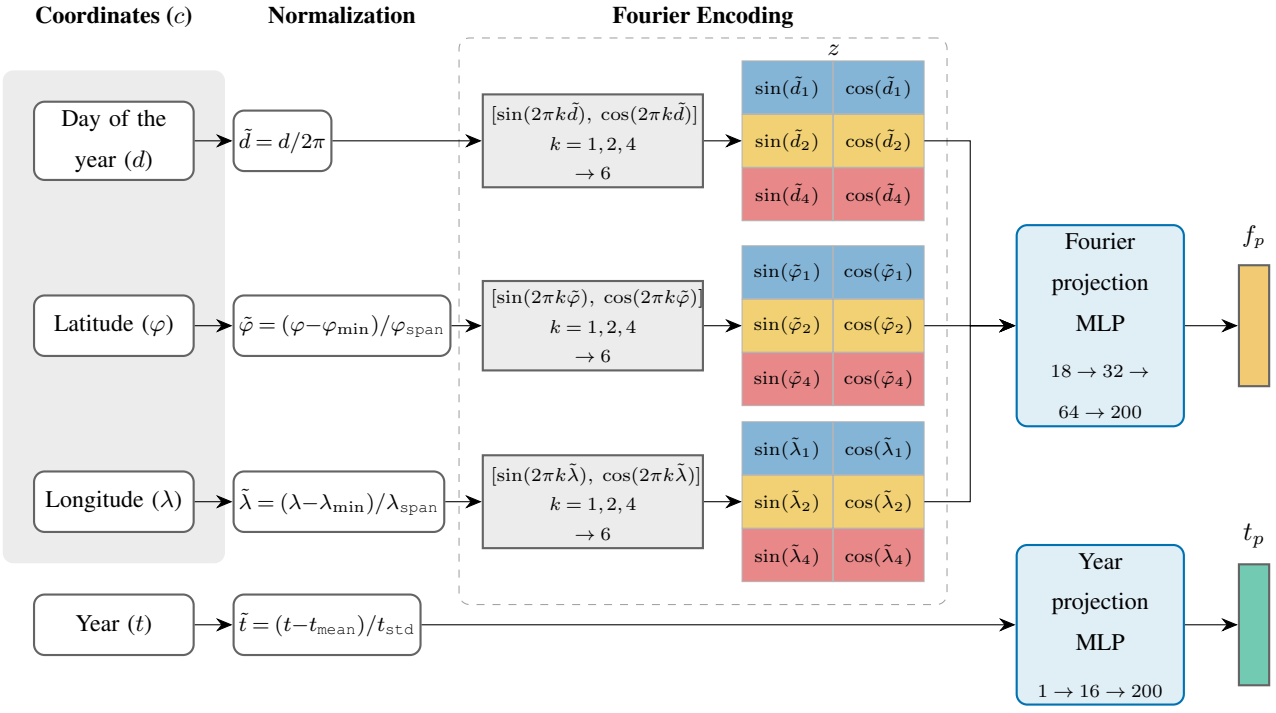


Figure 3. FourierBGC's spatiotemporal encoding, expanding the stage shown compressed in Fig. 2. Colour groups the sin/cos pair for each harmonic $k \in \{1, 2, 4\}$; subscripts denote k. f_p and t_p denote the combined Fourier projection and the year projection, respectively. The three coordinates are encoded independently and mix only at the concatenation into the shared projection MLP. FourierBGC-Broadcast omits both projection MLPs, broadcasting the fixed values across depth instead.

alongside $\mathbf{T}$, $\mathbf{S}$, and $\mathbf{O}_2$, as PPCon does, they are broadcast along depth and concatenated only after the full convolutional stack, immediately before the final projection layer, which is correspondingly widened from 32 to 36 input channels. This late fusion was adopted after concatenating raw-scale values at the input produced training instability. CNN-RawCoord isolates whether PPCon's learned encoding contributes anything beyond making the coordinates available to the network at all.

4.3 FourierBGC: Fourier-feature encoding with a learned projection

Day-of-year has an exact periodic structure: day 365 and day 1 are adjacent points in the annual cycle. Latitude and longitude, while not cyclical within this study's bounded region, are continuous spatial coordinates of the kind for which sinusoidal bases represent smoothly-varying structure well (Tancik et al., 2020; Mai et al., 2020). All three are first rescaled by fixed constants

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



onto a common unit scale,

$$\tilde{d} = d/2\pi, \tag{4}$$

$$\tilde{\varphi} = (\varphi - \varphi_{\min})/\varphi_{\text{span}}, \tag{5}$$

$$\tilde{\lambda} = (\lambda - \lambda_{\min})/\lambda_{\text{span}}, \tag{6}$$

with $(\varphi_{\min}, \varphi_{\text{span}}) = (31.5, 13.0)$ and $(\lambda_{\min}, \lambda_{\text{span}}) = (0.57, 38.0)$ degrees, a fixed bounding box on the Mediterranean. A small number of profiles fall marginally outside this box, so the rescaled values may lie slightly beyond $[0,1]$; the sinusoidal basis below is defined on all of $\mathbb{R}$ and accommodates this without modification. Day-of-year needs no bounding box: recorded in radians over $[0, 2\pi)$, it reaches the same scale by division alone. The spatial constants are fixed in advance rather than fitted to the training split.

Each rescaled coordinate is then expanded by a fixed sinusoidal basis at harmonics $\mathcal{K} = \{1, 2, 4\}$, a standard dyadic progression (Tancik et al., 2020),

$$\gamma(x) = [\sin(2\pi k x),\ \cos(2\pi k x)]_{k \in \mathcal{K}} \in \mathbb{R}^6, \tag{7}$$

and the three expansions are concatenated into a single fixed, parameter-free encoding (Fig. 3),

$$\mathbf{z} = \left[\gamma(\tilde{d}),\ \gamma(\tilde{\varphi}),\ \gamma(\tilde{\lambda})\right] \in \mathbb{R}^{18}. \tag{8}$$

Three properties motivate this choice. First, every component is bounded in $[-1, 1]$ by construction, which permits fusion at the network input rather than the late fusion CNN-RawCoord required. Second, the encoding is exactly periodic in d: $\gamma(\tilde{d})$ takes the same value on day 365 and day 1, so the annual wraparound is imposed rather than inferred. Third, the inner product between two encodings is stationary,

$$\gamma(x)^\top \gamma(x') = \sum_{k \in \mathcal{K}} \cos(2\pi k (x - x')), \tag{9}$$

depending only on the separation $x - x'$ and not on absolute position, with each harmonic contributing structure at its own scale: $k = 1$ resolves one cycle across the domain, basin-scale in space and annual in time, while $k = 4$ resolves four. A linear map on $\mathbf{z}$ can therefore express smooth multi-scale variation that the same map applied to a raw scalar cannot. Harmonic count and spacing were not tuned on this dataset.

The encoding is projected onto a single depth-varying channel by a three-layer network with SELU activations, and year is z-scored and projected separately by a second,

$$f_p = \text{MLP}_\theta(\mathbf{z}) \in \mathbb{R}^{200}, \tag{10}$$

$$t_p = \text{MLP}_\psi(\tilde{t}) \in \mathbb{R}^{200}, \tag{11}$$

with layer widths $18 \to 32 \to 64 \to 200$ and $1 \to 16 \to 200$, contributing 15,720 and 3,432 parameters respectively. The two channels are concatenated with the three profile channels to form the backbone input,

$$\mathbf{X} = [\mathbf{T},\ \mathbf{S},\ \mathbf{O}_2,\ f_p,\ t_p] \in \mathbb{R}^{200 \times 5}, \tag{12}$$

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

the same position in the network occupied by PPCon's four encoder outputs in Eq. (3).

The three coordinates share one projection rather than receiving separate networks: this is the simplest configuration that isolates the effect of the fixed encoding against the learned per-coordinate encoders of CNN-MLPCoord, and against CNN-RawCoord and FourierBGC-Broadcast, without introducing further un-ablated architectural choices. Whether splitting the projection across per-coordinate channels changes the result is left to future work.

Year is excluded from Eq. (8) because neither justification for the sinusoidal basis applies to it: it is not periodic over 2012–2020, and it indexes at most a coarse interannual trend rather than a smoothly-varying field, so Eq. (9) carries no meaning for it. It is retained as an input, rather than dropped, because every model in the spine receives the same four coordinates or none at all; removing year from FourierBGC alone would change the input set rather than the encoding, and the comparison would no longer isolate the encoding.

Together the two projections add 19,152 parameters, an 88 % reduction from PPCon's 158,800. Note that Eqs. (7)–(8) fix the basis in advance: structure not expressible in three harmonics must be recovered by MLP_θ or by the backbone, and we do not claim the chosen $\mathcal{K}$ is optimal for this dataset.

4.4 FourierBGC-Broadcast: fixed encoding without a learned projection

FourierBGC-Broadcast uses the same $\mathbf{z}$ and the same z-scored year as FourierBGC, but omits both projections: rather than passing either through a learned network, it broadcasts them directly along the depth dimension, using no learned parameters at all. Comparing the two isolates what the learned projection contributes beyond the fixed encoding alone.

4.5 Training Procedure

Two training recipes appear in this study, summarized in Table 3. PPCon and PPCon-NoCoord use PPCon's own, unchanged: Adadelta, with a three-term loss combining reconstruction error, weight regularization, and a smoothness penalty, and no gradient clipping or learning-rate schedule (Pietropolli et al., 2024). Adadelta is a widely used adaptive optimizer, not an outdated one; we retain it because our aim is to compare against PPCon's reported configuration rather than to re-tune it.

Every model we train ourselves uses a single common recipe, held fixed so that differences in reconstruction accuracy reflect the encoding under test rather than the training procedure. We optimize with Adam (Kingma and Ba, 2017) at a learning rate of 1×10^{-3}, its common default; Adam is also the optimizer used by Pietropolli et al. (2023), the model immediately preceding PPCon in this lineage. The remaining components are conventional rather than default choices: a five-epoch linear warmup followed by cosine decay of the learning rate. Gradient clipping at a maximum norm of 1.0 was adopted in response to an observed failure rather than by default: an early run without it showed pronounced instability, test-set error spiking to roughly 28 times its surrounding values before recovering over approximately 20 epochs, distinct from ordinary epoch-to-epoch noise. It was used throughout thereafter.

We do not attribute the resulting differences to any single component of this recipe. Establishing whether the optimizer, the schedule, or the clipping is responsible would require an ablation over the training procedure itself, which is outside the scope of a study whose subject is the coordinate encoding.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 3. Training configurations. Every model we train ourselves shares a single recipe; PPCon and PPCon-NoCoord use PPCon's own.

	PPCon's recipe	Common recipe
Models	PPCon, PPCon-NoCoord	CNN-NoCoord, CNN-RawCoord, CNN-MLPCoord, FourierBGC-Broadcast, FourierBGC
Optimizer	Adadelta	Adam
Learning rate	1.0	1×10^{-3}
Schedule	none	5-epoch linear warmup, cosine decay
Gradient clip	none	max norm 1.0
Epochs	100 / 150 / 125 [a]	100
Batch size	32	32
Dropout	0.2	0.2
Loss	three-term [b]	reconstruction only
Seeds	unseeded	5

[a] nitrate / chlorophyll-*a* / bbp700, from `dict.py` in the released code; Pietropolli et al. (2024)'s Table 3 reports 50/150/100. [b] reconstruction error, weight regularization, and a smoothness penalty.

Each model we train is run with five random seeds, varying weight initialization and data shuffling, and we report the mean and standard deviation across those runs. PPCon carries no such seed, since we evaluate the authors' published checkpoint directly rather than retraining it. Its released evaluation code exhibits run-to-run non-determinism of its own: an unseeded shuffle and an off-by-one indexing issue in its reconstruction routine cause repeated evaluations of the same unmodified checkpoint to give slightly different results. We therefore report the median of eleven independent re-runs of that code rather than a single run, using the median because the same non-determinism occasionally produces outlying evaluations.

The effect of the training procedure can be read from two pairs of models that differ in nothing else. Without coordinate encoders, PPCon-NoCoord converges poorly under PPCon's recipe, its final test-set error substantially worse than PPCon's despite leaving the convolutional backbone untouched, while CNN-NoCoord, identical in architecture, converges cleanly and improves on PPCon by a wide margin on nitrate, though not on chlorophyll-*a*. With the encoders present, the same contrast appears between PPCon and CNN-MLPCoord, which shares PPCon's encoder exactly and differs only in how it is trained. That the procedure matters in both cases indicates the difficulty is not confined to the configuration in which the encoders are absent. We accordingly treat PPCon-NoCoord's collapse as a property of the recipe rather than as a finding about the value of spatiotemporal information, and report it in full so that the contrast is verifiable rather than asserted.

We compare against PPCon's published configuration as its authors report it, and do not investigate whether further tuning of that recipe, independent of any architecture or encoding change, could narrow the gap.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 4. Overall pooled per-profile test RMSE across the ablation spine. Bold indicates the lowest value per column; parentheses are standard deviations across five seeded runs. PPCon is evaluated from its authors' released checkpoint, so it has no seed spread.

Model	Encoding	Nitrate (mmol m^{-3})	CHLA (mg m^{-3})	bbp700 (10^{-4} m^{-1})
PPCon	learned MLPs	0.5234	0.0802	2.153
PPCon-NoCoord	none	1.8265 (0.7839)	0.1489 (0.0002)	4.182 (1.332)
CNN-NoCoord	none	0.3404 (0.0036)	0.0954 (0.0006)	2.028 (0.005)
CNN-RawCoord	raw	0.3411 (0.0018)	0.0954 (0.0006)	2.033 (0.018)
CNN-MLPCoord	learned MLPs	0.3293 (0.0066)	0.0673 (0.0005)	1.761 (0.010)
FourierBGC-Broadcast	Fourier	0.2339 (0.0057)	0.0717 (0.0013)	**1.595** (0.027)
FourierBGC	Fourier + MLP	**0.2293** (0.0039)	**0.0641** (0.0003)	1.641 (0.018)

PPCon-NoCoord shows markedly larger seed-to-seed variance than every other model here on Nitrate and bbp700; see Sect. 6.

5 Results

Table 4 reports overall pooled test RMSE for all three variables. FourierBGC records the lowest error on nitrate and chlorophyll-*a*, and FourierBGC-Broadcast the lowest on bbp700. Every model trained with the common recipe improves on PPCon for nitrate and for bbp700; on chlorophyll-*a*, PPCon is better than the models that discard the coordinates or supply them without encoding, and is improved upon only once an encoding is introduced. PPCon-NoCoord is the weakest model on all three variables, by a wide margin on nitrate.

Two comparisons separate the contributions of the training procedure and the encoding. CNN-MLPCoord uses PPCon's encoder unchanged and differs from it only in training procedure; it improves on PPCon by 37.1 % on nitrate, 16.1 % on chlorophyll-*a* and 18.2 % on bbp700. FourierBGC differs from CNN-MLPCoord only in the encoding, the procedure being identical; it improves further by 30.4 %, 4.8 % and 6.8 % respectively. Both differences are large relative to seed variation: on nitrate a gap of 0.1000 mmol m^{-3} against standard deviations of 0.0066 and 0.0039, on chlorophyll-*a* a gap of 0.0032 mg m^{-3} against 0.0005 and 0.0003, and on bbp700 a gap of 0.120 against 0.010 and 0.018.

Supplying the coordinates without encoding them accounts for none of this. With the procedure held fixed, CNN-RawCoord and CNN-NoCoord are indistinguishable on all three variables: 0.3411 against 0.3404 on nitrate, 0.0954 on both for chlorophyll-*a*, and 2.033 against 2.028 on bbp700, differences well inside the seed spread of either model, although the two receive exactly the same four values. Encoding those same values instead improves on CNN-NoCoord by 32.6 %, 32.8 % and 19.1 % respectively.

The learned projection behaves differently on each of the three variables, and this is the one place where they diverge rather than agree. On nitrate, FourierBGC and FourierBGC-Broadcast differ by 0.0046 mmol m^{-3} against standard deviations of 0.0039 and 0.0057, so the projection lies within run-to-run variation; the parameter-free FourierBGC-Broadcast already improves on CNN-MLPCoord by 29.0 %. On chlorophyll-*a* the projection contributes 10.6 % (0.0641 against 0.0717), and it is required for the Fourier encoding to lead at all: FourierBGC-Broadcast alone is 6.5 % worse than CNN-MLPCoord on

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 5. Test RMSE by region. PPCon values are the median of 11 independent re-runs of PPCon's own evaluation code; CNN-MLPCoord and FourierBGC values are means across 5 seeded runs, standard deviation in parentheses. Bold indicates the lowest value per variable per region. Profiles falling outside all five region boxes are excluded (Sect. 2).

	Nitrate (mmol m^{-3})			CHLA (mg m^{-3})		
Region	PPCon	CNN-MLPCoord	FourierBGC	PPCon	CNN-MLPCoord	FourierBGC
NWM	0.6532	0.3357 (0.0093)	**0.2659** (0.0121)	0.1299	0.1063 (0.0015)	**0.0984** (0.0008)
SWM	0.5014	0.5900 (0.0727)	**0.2262** (0.0119)	0.1187	0.1104 (0.0014)	**0.1073** (0.0020)
TYR	0.4408	0.2438 (0.0142)	**0.1749** (0.0031)	0.0840	0.0735 (0.0011)	**0.0721** (0.0003)
ION	0.4121	0.2412 (0.0060)	**0.1912** (0.0027)	0.0445	0.0363 (0.0007)	**0.0333** (0.0006)
LEV	0.5143	0.3891 (0.0158)	**0.2260** (0.0029)	0.0468	0.0410 (0.0011)	**0.0408** (0.0006)

Table 6. Test RMSE by season. Conventions as in Table 5.

	Nitrate (mmol m^{-3})			CHLA (mg m^{-3})		
Season	PPCon	CNN-MLPCoord	FourierBGC	PPCon	CNN-MLPCoord	FourierBGC
Winter	0.5109	0.3072 (0.0090)	**0.2118** (0.0064)	0.0751	0.0689 (0.0004)	**0.0684** (0.0005)
Spring	0.5402	0.3248 (0.0054)	**0.2296** (0.0090)	0.1285	0.0958 (0.0009)	**0.0903** (0.0011)
Summer	0.4921	0.3611 (0.0090)	**0.2524** (0.0029)	0.0771	0.0600 (0.0009)	**0.0560** (0.0008)
Fall	0.5084	0.3233 (0.0096)	**0.2221** (0.0027)	0.0468	0.0416 (0.0007)	**0.0400** (0.0003)

this variable. On bbp700 the projection is instead a cost: FourierBGC-Broadcast improves on CNN-MLPCoord by 9.4 %, and adding the projection gives back part of that gain, raising RMSE from 1.595 to 1.641, a difference of roughly three times its standard error across seeds. The fixed encoding is therefore beneficial on all three variables, while the projection is decisive on one, immaterial on another, and mildly harmful on the third.

For nitrate, FourierBGC's pooled RMSE of 0.2293 mmol m^{-3} falls below the approximately 0.25 mmol m^{-3} observational uncertainty estimated for BGC-Argo nitrate sensors (Mignot et al., 2019), against roughly twice that figure for PPCon; the corresponding chlorophyll-*a* value of 0.0641 mg m^{-3} remains about twice the 0.03 mg m^{-3} estimate. The comparison is not extended to bbp700, for which we did not identify an equivalent published sensor-uncertainty estimate. Model error and observational uncertainty are in any case not directly commensurable, and Pietropolli et al. (2024) note that separating instrument error from model error in this setting remains open; we report the comparison to indicate the scale at which further gains would need to be interpreted.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table 7. bbp700 test RMSE by region and season (10^{-4} m^{-1}), for the four models that separate the training procedure, the encoding and the projection. Bold indicates the lowest value per row; parentheses are standard deviations across five seeded runs. PPCon is evaluated from its authors' released checkpoint and has no seed spread. Conventions otherwise as in Table 5.

Subset	PPCon	CNN-MLPCoord	FourierBGC-Broadcast	FourierBGC
NWM	2.423	1.976 (0.017)	**1.687** (0.036)	1.769 (0.017)
SWM	1.969	1.595 (0.044)	1.634 (0.044)	**1.593** (0.044)
TYR	2.299	1.714 (0.005)	**1.674** (0.037)	1.678 (0.037)
ION	1.579	1.269 (0.014)	1.230 (0.029)	**1.206** (0.031)
LEV	1.707	1.515 (0.010)	**1.446** (0.020)	1.491 (0.026)
Winter	2.390	2.033 (0.011)	**1.711** (0.041)	1.859 (0.036)
Spring	2.574	2.304 (0.011)	**2.115** (0.033)	2.136 (0.032)
Summer	1.432	1.281 (0.007)	1.277 (0.020)	**1.202** (0.017)
Fall	1.499	1.350 (0.030)	**1.210** (0.020)	1.300 (0.020)

5.1 Regional and Seasonal Variation

Following PPCon, we report regional and seasonal breakdowns, adding CNN-MLPCoord so that the encoding and the training procedure can be separated within each subset. FourierBGC-Broadcast is omitted from Tables 5 and 6, where its behaviour adds nothing to the pooled result: on chlorophyll-*a* the projection reduces error in all five regions and all four seasons, by between 2.3 % and 21.0 %, every difference exceeding three standard errors. On nitrate its pooled neutrality instead conceals partially offsetting subsets, the projection improving the Levantine Sea by 12.5 % and spring by 6.0 % while degrading the Ionian Sea by 4.3 %; we return to this in Sect. 6. Only on bbp700 does the projection change sign in a structured way, and we report that breakdown separately in Table 7.

FourierBGC records the lowest error in every region and every season for both variables, so the pooled margins of Table 4 are not produced by a small number of favourable subsets. Its margin over CNN-MLPCoord, however, is distributed very differently between the two variables. On nitrate the margin is large everywhere, ranging from 20.7 % in the Ionian Sea to 61.7 % in the southwestern Mediterranean, and from 29.3 % to 31.3 % across seasons.

On chlorophyll-*a* the margins are an order of magnitude smaller and not uniformly resolved. In five of the nine subsets, the northwestern Mediterranean and Ionian Sea and all of spring, summer and fall, the difference exceeds its standard error across seeds by a factor of four or more. In the southwestern Mediterranean and the Tyrrhenian Sea it is weaker. In the Levantine Sea (0.0408 against 0.0410) and in winter (0.0684 against 0.0689) the two models differ by less than the uncertainty in that difference, so their ordering in those subsets is not established, although FourierBGC is nominally lower in both.

The seasonal pattern is structured rather than arbitrary: the largest reductions relative to CNN-MLPCoord occur in spring and summer, the seasons in which the vertical distribution of chlorophyll is most strongly organised by the spring bloom and the deep chlorophyll maximum, and the smallest in winter, when deep mixing leaves that distribution more uniform. FourierBGC's

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

advantage is thus largest where seasonal structure is most pronounced, which is consistent with the encoding supplying that structure directly. We note this as a consistency check rather than a causal demonstration.

Two subsets warrant brief comment, both resting on sparse data. CNN-MLPCoord's nitrate error in the southwestern Mediterranean (0.5900) exceeds PPCon's (0.5014), the only case in which the common procedure does not improve on PPCon for this variable; that subset holds nine test profiles and carries the largest seed spread in either table (0.0727), too few to separate a basin-specific effect from sampling variability. In the Levantine Sea the chlorophyll-*a* margin over CNN-MLPCoord is not resolved, and it is one of the two sub-basins in which FourierBGC-Broadcast falls behind PPCon on that variable. We record both as open observations rather than findings.

On bbp700 both steps of the decomposition hold in all nine subsets, but the projection does not: FourierBGC-Broadcast records the lowest error in six of the nine, and the sign of its effect tracks season rather than region, raising error by 8.7 % in winter and lowering it by 5.9 % in summer, with the same sign in all five sub-regions within each. Spring is mixed, and in fall the projection costs error in four of the five sub-regions. We take this up in Sect. 6.

The breakdowns above are marginal, each averaging over the dimension it does not resolve. Regions and seasons partition the test set along independent axes, so their intersection yields a finer grid in which each cell holds the profiles from a single sub-basin in a single season. Of the sixty such cells across the three variables, fifty-six contain at least eight test profiles and are considered here; the four excluded are nitrate cells in the southwestern Mediterranean, which contributes nine nitrate profiles in total. In fifty of the fifty-six, the better of the two Fourier variants records the lowest error of any model in the spine. In four it is nominally higher than CNN-MLPCoord by less than the standard error of that difference, and in two it is beaten by more: chlorophyll-*a* in the southwestern Mediterranean in winter (0.0965 against 0.0865 $\mathrm{mg\,m^{-3}}$, $n = 10$) and bbp700 in the Levantine Sea in fall (0.843 against $0.794 \times 10^{-4}\,\mathrm{m^{-1}}$, $n = 60$). In all five of these cells the model matching or exceeding the Fourier variants is CNN-MLPCoord; PPCon, CNN-NoCoord and CNN-RawCoord record the lowest error in no cell of the grid.

6 Discussion

6.1 Decomposing the Improvement

FourierBGC improves on PPCon by 56.2 % on nitrate, 20.1 % on chlorophyll-*a* and 23.8 % on bbp700 (Table 4). None of these figures are attributable to the encoding alone, since PPCon and FourierBGC differ in training procedure as well. CNN-MLPCoord separates the two: it carries PPCon's encoder unchanged and differs from it only in training, so the step from PPCon to CNN-MLPCoord measures the procedure with the encoder fixed, and the step from CNN-MLPCoord to FourierBGC measures the encoding with the procedure fixed.

Both steps improve all three variables, and on all three the procedure accounts for the larger share. On nitrate, the change of procedure reduces error by 37.1 % and the change of encoding by a further 30.4 %; of the total reduction from PPCon to FourierBGC (0.2941 $\mathrm{mmol\,m^{-3}}$), 66.0 % is attributable to the procedure and 34.0 % to the encoding. On chlorophyll-*a* the

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

same steps give 16.1 % and 4.8 %, or 80.1 % and 19.9 % of a total reduction of 0.0161 mg m^{-3}. On bbp700 they give 18.2 % and 6.8 %, or 76.6 % and 23.4 % of a total reduction of 0.512×10^{-4} m^{-1}.

The encoding's contribution therefore differs substantially between variables, from 34.0 % of the total on nitrate to 19.9 % on chlorophyll-*a*, while the training procedure contributes on all three. A comparison against PPCon alone would attribute the whole of each margin to the encoding and would overstate it in every case, most severely on chlorophyll-*a*, where four fifths of the improvement follows from the procedure.

This also revises what can be inferred from the models that discard the coordinates entirely. CNN-NoCoord improves on PPCon by 35.0 % on nitrate and 5.8 % on bbp700 but is 19.0 % worse on chlorophyll-*a*, which might suggest the procedure helps two variables and harms the third. It does not: CNN-NoCoord differs from PPCon both in procedure and in having no encoder, and on chlorophyll-*a* the second of these is costly, since CNN-MLPCoord improves on CNN-NoCoord by 29.5 % on that variable, against 13.2 % on bbp700 and only 3.3 % on nitrate. Once the encoder is held fixed, the procedure improves all three.

The collapse of PPCon-NoCoord (1.8265 mmol m^{-3} on nitrate, 3.5 times PPCon's error, and 1.9 times on both chlorophyll-*a* and bbp700) remains best read as a property of PPCon's procedure rather than of the information removed. One plausible mechanism is that its loss applies weight regularization with a coefficient tuned per variable across all network weights (Pietropolli et al., 2024, Table 3); removing the four encoders removes 158,800 of 412,049 parameters, so that coefficient then regularizes a network 39 % smaller than the one it was tuned for. The seeded runs offer indirect support: PPCon-NoCoord is the only model in the spine with substantial seed-to-seed variance, at 43 % of its mean on nitrate and 32 % on bbp700 against 3 % or less for every other seeded model, and its recorded test loss is non-monotonic across seeds and in some cases diverges outright. That instability is what a mismatched regularization strength would produce, and it is absent on chlorophyll-*a* (0.1 %), where the same pipeline is stable. We offer this as a hypothesis consistent with their published configuration and with the instability observed here, not as a tested claim.

6.2 Representation Rather Than Availability

Holding the training procedure fixed, supplying the four spatiotemporal coordinates as raw values does not improve on supplying nothing at all. CNN-RawCoord is nominally 0.2 % worse than CNN-NoCoord on nitrate and 0.2 % worse on bbp700, and the two are identical to four decimal places on chlorophyll-*a*; all three differences fall inside the seed spread of either model. The information content of the two treatments is the same: the same four numbers per profile. A model given latitude, longitude, day-of-year and year as raw values therefore performs indistinguishably from one given no coordinates, so the values are present but unusable in that form.

Encoding them changes this on all three variables, but by different amounts and through different means. On chlorophyll-*a*, any learned encoding recovers most of the available gain: PPCon's per-coordinate encoders improve on CNN-NoCoord by 29.5 % and the Fourier encoding by 32.8 %. On bbp700 the two are further apart, at 13.2 % and, for the broadcast form, 21.4 %. On nitrate the specific form matters most of all: the same learned encoders improve by only 3.3 %, while the Fourier encoding

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

improves by 32.6 %. What the network can exploit is therefore not the availability of the coordinates but their representation, and on nitrate a fixed sinusoidal basis extracts structure that 158,800 learned parameters do not.

This is the one result that holds without qualification across the study. The Fourier encoding improves on both the coordinate-free and the raw-coordinate baseline on every variable, in every region and in every season.

6.3 The Role of the Learned Projection

FourierBGC and FourierBGC-Broadcast differ only in whether the fixed encoding is projected through a small learned network or broadcast unchanged along depth, which isolates the projection. Its effect is variable-dependent, and the three variables place it in three different roles.

On nitrate the projection is redundant overall, but that pooled neutrality conceals structure. FourierBGC and FourierBGC-Broadcast differ by 0.0046 mmol m^{-3} (0.2293 against 0.2339), smaller than either model's seed standard deviation, and the parameter-free FourierBGC-Broadcast already improves on CNN-MLPCoord by 29.0 %. Regionally the projection is not simply inert: it improves the Levantine Sea by 12.5 % and spring by 6.0 %, both resolved across seeds, while degrading the Ionian Sea by 4.3 %; these partially offset in the pooled figure. On this variable, then, a fixed encoding with no learned encoder at all already outperforms 158,800 learned parameters, and the projection redistributes error across sub-regions without changing the total.

On chlorophyll-*a* the projection is not optional but decisive. FourierBGC-Broadcast is 6.5 % *worse* than CNN-MLPCoord on this variable (0.0717 against 0.0673), so the fixed encoding alone does not improve on PPCon's learned encoders; only with the projection does FourierBGC lead, at 0.0641. The projection contributes 10.6 % over the broadcast form, more than ten times the standard error of that difference, and it is the sole reason the Fourier encoding wins on this variable at all. The regional and seasonal breakdowns qualify this lead rather than undermine it: FourierBGC records the lowest error in every subset, and its margin over CNN-MLPCoord exceeds four standard errors in five of the nine regions and seasons, is weaker in the southwestern Mediterranean and the Tyrrhenian Sea, and is not resolved at all in the Levantine Sea or in winter (Tables 5 and 6).

On bbp700 the projection is a cost. FourierBGC-Broadcast is the best model on this variable at 1.595×10^{-4} m^{-1}, and adding the projection raises error to 1.641, a difference of roughly three times its standard error; the broadcast form improves on CNN-MLPCoord by 9.4 %, and the projection returns about a third of that gain. We take this up separately below, since it is the one variable on which the projection's cost has a structured, rather than incidental, explanation.

Across all three variables, the projection is therefore the component of FourierBGC that generalizes least well: decisive on one variable, redistributive on a second, and a net cost on a third. This is a statement about the projection and not about the encoding – the fixed Fourier features improve on both baselines on every variable, and the disagreement is only over whether allowing them to vary with depth helps.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

6.4 Design Choices Not Established by Tuning

Several choices in this study were made on general grounds rather than selected by search, and we identify them so that they are not mistaken for tuned results. The harmonic set $\mathcal{K} = \{1, 2, 4\}$ follows a standard dyadic progression rather than a sweep on this dataset. Day-of-year, latitude and longitude share a single projection instead of receiving separate ones, which is the simplest configuration that isolates the encoding but is not shown to be the best; on chlorophyll-*a* this means the margin over CNN-MLPCoord reflects both the encoding and the change from four separate projections to one shared projection, which we do not separate. On nitrate and bbp700 that ambiguity does not arise, since FourierBGC-Broadcast has no projection at all and already accounts for the margin on both. The Mediterranean bounding box of Eqs. (5)–(6) is likewise fixed a priori. Each of these could plausibly be improved by tuning, which would raise FourierBGC's numbers relative to those reported here.

One aspect of PPCon's published configuration is worth noting in the same spirit. Its MLPs apply SELU activations, including on the output layer, where Pietropolli et al. (2024, Table 1) shows none, a convention FourierBGC's own projections follow; and their inputs are passed unnormalized, so that a year enters the network as 2016 and a latitude as 35.29. SELU's self-normalizing property is derived for inputs standardized to zero mean and unit variance (Klambauer et al., 2017), an assumption that raw coordinates of this magnitude do not satisfy. The evidence here suggests this is not the limiting factor: CNN-MLPCoord passes the same unnormalized values through the same encoders and improves on PPCon by 37.1 % on nitrate and 18.2 % on bbp700, so what those encoders can extract is constrained more by how they are trained than by the scale of their inputs.

6.5 bbp700

bbp700 was the variable on which PPCon reported its weakest reconstruction, and it is the variable on which the present study's two main claims come apart most cleanly.

It supports the central claim without qualification. Raw coordinates buy nothing (2.033 against 2.028, a difference well inside the seed spread of CNN-RawCoord), the training procedure contributes 18.2 %, and encoding the same four values improves on the coordinate-free baseline by 21.4 %. In this respect bbp700 behaves exactly as nitrate and chlorophyll-*a* do, and it is the third independent variable on which representation rather than availability determines what the network can use.

It is also the only variable on which the depth-varying projection is harmful, and Table 7 suggests why. In winter, the projection costs 0.149 in these units (1.711 for FourierBGC-Broadcast against 1.859 for FourierBGC); in summer it gains 0.076 (1.277 against 1.202). Both differences are several times their standard error, and the sign is consistent across all five sub-regions within each of those two seasons: the projection improves every region in summer and degrades every region in winter. Spring is mixed, and in fall the projection costs error in four of the five sub-regions. A regional sampling artefact would not reverse cleanly on season while holding in every basin.

A mechanism consistent with this is that the projection allows the coordinate encoding to modulate with depth, and that there is more for it to express when the water column is stratified than when it is mixed. The summer Mediterranean is strongly stratified, with subsurface particle structure that varies with location and time of year (Barbieux et al., 2019); winter convection

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

homogenizes the upper column, leaving a more uniform profile in which the additional freedom has little to represent and can instead be fitted to noise. We offer this as an interpretation consistent with the pattern, not as a tested mechanism: it rests on a subgroup analysis over five seeds, and the checkpoints compared are selected on test-set error, which warrants more caution at subset resolution than at the pooled level.

What bbp700 contributes to the study is therefore not a third confirmation but a boundary. The fixed encoding is the robust component and improves every variable; the projection is the tuned component, and this is the variable on which its cost becomes visible.

7 Conclusions

This study isolates the contribution of spatiotemporal coordinate encoding to CNN-based biogeochemical Argo reconstruction, holding PPCon's convolutional backbone fixed while varying both how the coordinates are represented and how the model is trained. FourierBGC encodes day-of-year, latitude and longitude through a fixed 18-dimensional Fourier basis and a small learned projection, with year projected separately, using 19,152 encoder parameters against PPCon's 158,800. It improves on PPCon by 56.2 % on nitrate and 20.1 % on chlorophyll-*a*, records the lowest error in every sub-basin and every season for both, and brings the pooled nitrate error below the observational uncertainty currently estimated for BGC-Argo nitrate sensors.

Training PPCon's own encoder under the common procedure separates the two effects that margin combines. With the encoder held fixed, the training procedure accounts for 37.1 % on nitrate and 16.1 % on chlorophyll-*a*; with the procedure held fixed, the encoding accounts for a further 30.4 % and 4.8 %. The procedure is the larger share on both. That a substantial part of a published model's reported error was attributable to how it was trained rather than to how it represented its inputs is worth establishing before further encoder complexity is added.

What the encoding contributes depends on its form rather than on the information it carries. Supplying the coordinates as raw values does not improve on supplying none at all, although the same four numbers are present in both cases; encoding them instead recovers roughly a third of the error on either variable. On nitrate the form matters most: a parameter-free Fourier encoding outperforms PPCon's 158,800 learned parameters by 29.0 %, and the learned projection adds nothing measurable. On chlorophyll-*a* the reverse holds, the projection being what places the encoding ahead at all, by a margin that is resolved in most sub-basins and seasons but not in every one. On bbp700 the fixed encoding again leads, but the projection is a net cost whose sign tracks the seasonal cycle of stratification.

That representation matters more than availability makes the open question not whether to encode the spatiotemporal coordinates but how, and this study samples three points in that space: raw values, learned per-coordinate transforms, and a fixed sinusoidal basis at dyadic frequencies. Alternatives remain untested, among them random rather than dyadic frequencies, learned rather than fixed frequencies, and encodings designed specifically for geographic position (Mai et al., 2020). Successive models in this lineage have increased the complexity of the coordinate encoder without comparing encoding families directly; an ablation across them, holding the backbone and training procedure fixed as here, would identify which properties of an encoding carry the improvement. The present study establishes that the choice is consequential, not which choice is best.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Table A1. Which model records the lowest error in each region-by-season cell, for the 56 cells with at least eight test profiles (Sect. 5). **F** = FourierBGC, **B** = FourierBGC-Broadcast, **M** = CNN-MLPCoord. A dash marks the four excluded nitrate/southwestern-Mediterranean cells ($n \leq 4$). * marks a CNN-MLPCoord value within one seed standard error of the better Fourier variant; † marks a margin exceeding it. PPCon values are a single deterministic evaluation rather than the median of 11 re-runs used in Tables 5–7; this affects none of its comparisons here, since it does not record the lowest error in any cell of the grid, nor does PPCon-NoCoord, CNN-NoCoord, or CNN-RawCoord.

	Winter	Spring	Summer	Fall
Nitrate				
NWM	F	F	F	B
SWM	–	–	–	–
TYR	B	F	B	B
ION	F	B	B	B
LEV	F	F	F	F
Chlorophyll-a				
NWM	B	B	F	F
SWM	M†	F	F	F
TYR	F	F	M*	M*
ION	F	F	F	F
LEV	F	F	F	F
bbp700				
NWM	B	B	F	B
SWM	B	F	F	M*
TYR	B	B	F	B
ION	B	F	F	B
LEV	B	M*	F	M†

Code and data availability. The code used to implement, train, and evaluate all models described in this paper (PPCon, PPCon-NoCoord, CNN-NoCoord, CNN-RawCoord, CNN-MLPCoord, FourierBGC-Broadcast, and FourierBGC), along with the scripts used to generate all figures and tables, is archived on Zenodo at https://doi.org/10.5281/zenodo.22650255 (Phillips, 2026) under the MIT license. The current development version is available at https://github.com/samphillips-ml/FourierBGC. The PPCon baseline architecture and training procedure follow Pietropolli et al. (2024), whose code is available at https://github.com/gpietrop/PPCon. The BGC-Argo Mediterranean float profiles used to train and evaluate all models in this study originate from Amadio et al. (2023), archived on Zenodo at https://doi.org/10.5281/zenodo.10391759 under the CC-BY 4.0 license. The specific train/test partition used follows Pietropolli et al. (2024).

Appendix A: Region- and season-resolved model comparison

Sect. 5 reports that a Fourier variant records the lowest error in 50 of the 56 region-by-season cells with at least eight test profiles. Table A1 gives the full grid.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Author contributions. S.B.P. developed the methodology, wrote the software, performed the formal analysis, and wrote the original draft. X.L. supervised the project and reviewed and edited the manuscript.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This research used resources of the University Research Computing (URC) at the University of North Carolina at Charlotte.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

References

Amadio, C., Teruzzi, A., Feudale, L., Bolzon, G., Di Biagio, V., Lazzari, P., Álvarez, E., Coidessa, G., Salon, S., and Cossarini, G.: Mediterranean Quality checked BGC-Argo 2013-2022 dataset, https://doi.org/10.5281/zenodo.10391759, 2023.

Amadio, C., Teruzzi, A., Pietropolli, G., Manzoni, L., Coidessa, G., and Cossarini, G.: Combining neural networks and data assimilation to enhance the spatial impact of Argo floats in the Copernicus Mediterranean biogeochemical model, Ocean Science, 20, 689–710, https://doi.org/10.5194/os-20-689-2024, 2024.

Barbieux, M., Uitz, J., Gentili, B., Pasqueron de Fommervault, O., Mignot, A., Poteau, A., Schmechtig, C., Taillandier, V., Leymarie, E., Penkerc'h, C., D'Ortenzio, F., Claustre, H., and Bricaud, A.: Bio-optical characterization of subsurface chlorophyll maxima in the Mediterranean Sea from a Biogeochemical-Argo float database, Biogeosciences, 16, 1321–1342, https://doi.org/10.5194/bg-16-1321-2019, 2019.

Biogeochemical-Argo Planning Group: The Scientific Rationale, Design, and Implementation Plan for a Biogeochemical-Argo Float Array, Tech. rep., Ifremer, https://doi.org/10.13155/46601, 2016.

Bittig, H., Wong, A., and Plant, J.: BGC-Argo synthetic profile file processing and format on Coriolis GDAC, https://doi.org/10.13155/55637, version 1.1, 2019.

Bittig, H. C., Steinhoff, T., Claustre, H., Fiedler, B., Williams, N. L., Sauzède, R., Körtzinger, A., and Gattuso, J.-P.: An Alternative to Static Climatologies: Robust Estimation of Open Ocean CO2 Variables and Nutrient Concentrations From T, S, and O2 Data Using Bayesian Neural Networks, Frontiers in Marine Science, 5, 328, https://doi.org/10.3389/fmars.2018.00328, 2018.

Campbell, L., Gray, N., Fairbanks, L., Silver, J., Gruby, R., Dubik, B., and Basurto, X.: Global Oceans Governance: New and Emerging Issues, Annual Review of Environment and Resources, 41, 517–543, https://doi.org/10.1146/annurev-environ-102014-021121, 2016.

Carter, B. R., Feely, R. A., Williams, N. L., Dickson, A. G., Fong, M. B., and Takeshita, Y.: Updated Methods for Global Locally Interpolated Estimation of Alkalinity, pH, and Nitrate, Limnology and Oceanography: Methods, 16, 119–131, https://doi.org/10.1002/lom3.10232, 2018.

Chai, F., Johnson, K. S., Claustre, H., Xing, X., Wang, Y., Boss, E., Riser, S., Fennel, K., Schofield, O., and Sutton, A.: Monitoring ocean biogeochemistry with autonomous platforms, Nature Reviews Earth & Environment, 1, 315–326, https://doi.org/10.1038/s43017-020-0053-y, 2020.

Fourrier, M., Coppola, L., Claustre, H., D'Ortenzio, F., Sauzède, R., and Gattuso, J.-P.: A Regional Neural Network Approach to Estimate Water-Column Nutrient Concentrations and Carbonate System Variables in the Mediterranean Sea: CANYON-MED, Frontiers in Marine Science, 7, 620, https://doi.org/10.3389/fmars.2020.00620, 2020.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, https://arxiv.org/abs/1412.6980, 2017.

Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S.: Self-Normalizing Neural Networks, https://arxiv.org/abs/1706.02515, 2017.

Mai, G., Janowicz, K., Yan, B., Zhu, R., Cai, L., and Lao, N.: Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells, https://arxiv.org/abs/2003.00824, 2020.

Mignot, A., D'Ortenzio, F., Taillandier, V., Cossarini, G., and Salon, S.: Quantifying Observational Errors in Biogeochemical-Argo Oxygen, Nitrate, and Chlorophyll a Concentrations, Geophysical Research Letters, 46, 4330–4337, https://doi.org/10.1029/2018GL080541, 2019.

Phillips, S.: samphillips-ml/FourierBGC: FourierBGC v1.0.0, https://doi.org/10.5281/zenodo.22650255, 2026.

Pietropolli, G., Manzoni, L., and Cossarini, G.: Multivariate Relationship in Big Data Collection of Ocean Observing System, Applied Sciences, 13, 5634, https://doi.org/10.3390/app13095634, 2023.

https://doi.org/10.5194/egusphere-2026-5646
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



Pietropolli, G., Manzoni, L., and Cossarini, G.: PPCon 1.0: Biogeochemical-Argo Profile Prediction with 1D Convolutional Networks, Geoscientific Model Development, 17, 7347–7364, https://doi.org/10.5194/gmd-17-7347-2024, 2024.

Roemmich, D., Alford, M. H., Claustre, H., Johnson, K., King, B., Moum, J., Oke, P., Owens, W. B., Pouliquen, S., Purkey, S., Scanderbeg, M., Suga, T., Wijffels, S., Zilberman, N., Bakker, D., Baringer, M., Belbeoch, M., Bittig, H. C., Boss, E., Calil, P., Carse, F., Carval, T., Chai, F., Conchubhair, D. Ó., d'Ortenzio, F., Dall'Olmo, G., Desbruyeres, D., Fennel, K., Fer, I., Ferrari, R., Forget, G., Freeland, H., Fujiki, T., Gehlen, M., Greenan, B., Hallberg, R., Hibiya, T., Hosoda, S., Jayne, S., Jochum, M., Johnson, G. C., Kang, K., Kolodziejczyk, N., Körtzinger, A., Le Traon, P.-Y., Lenn, Y.-D., Maze, G., Mork, K. A., Morris, T., Nagai, T., Nash, J., Naveira Garabato, A., Olsen, A., Pattabhi, R. R., Prakash, S., Riser, S., Schmechtig, C., Schmid, C., Shroyer, E., Sterl, A., Sutton, P., Talley, L., Tanhua, T., Thierry, V., Thomalla, S., Toole, J., Troisi, A., Trull, T. W., Turton, J., Velez-Belchi, P. J., Walczowski, W., Wang, H., Wanninkhof, R., Waterhouse, A. F., Waterman, S., Watson, A., Wilson, C., Wong, A. P. S., Xu, J., and Yasuda, I.: On the Future of Argo: A Global, Full-Depth, Multi-Disciplinary Array, Frontiers in Marine Science, 6, 439, https://doi.org/10.3389/fmars.2019.00439, 2019.

Sauzède, R., Bittig, H. C., Claustre, H., Pasqueron de Fommervault, O., Gattuso, J.-P., Legendre, L., and Johnson, K. S.: Estimates of Water-Column Nutrient Concentrations and Carbonate System Parameters in the Global Ocean: A Novel Approach Based on Neural Networks, Frontiers in Marine Science, 4, 128, https://doi.org/10.3389/fmars.2017.00128, 2017.

Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R.: Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains, https://arxiv.org/abs/2006.10739, 2020.

Teruzzi, A., Bolzon, G., Feudale, L., and Cossarini, G.: Deep chlorophyll maximum and nutricline in the Mediterranean Sea: emerging properties from a multi-platform assimilated biogeochemical model experiment, Biogeosciences, 18, 6147–6166, https://doi.org/10.5194/bg-18-6147-2021, 2021.

Wang, Z. A., Moustahfid, H., Mueller, A. V., Michel, A. P. M., Mowlem, M., Glazer, B. T., Mooney, T. A., Michaels, W., McQuillan, J. S., Robidart, J. C., Churchill, J., Sourisseau, M., Daniel, A., Schaap, A., Monk, S., Friedman, K., and Brehmer, P.: Advancing Observation of Ocean Biogeochemistry, Biology, and Ecosystems With Cost-Effective in situ Sensing Technologies, Frontiers in Marine Science, 6, 519, https://doi.org/10.3389/fmars.2019.00519, 2019.

Weller, R. A., Baker, D. J., Glackin, M. M., Roberts, S. J., Schmitt, R. W., Twigg, E. S., and Vimont, D. J.: The Challenge of Sustaining Ocean Observations, Frontiers in Marine Science, 6, 105, https://doi.org/10.3389/fmars.2019.00105, 2019.