

Global Biogeochemical Cycles

Supporting Information for

High resolution variability of the ocean carbon sink

Luke Gregor¹, Jamie Shutler², Nicolas Gruber¹

¹Environmental Physics, Institute of Biogeochemistry and Pollutant Dynamics, ETH Zurich, Zurich, Switzerland.

²Centre for Geography and Environmental Science, University of Exeter, Penryn, Cornwall, UK.

Contents of this file

Text S1: Data Preprocessing

Text S2: Machine Learning: OceanCarbNN details

Text S3: Evaluation

Text S4: Integration of Fluxes

Text S5: Computing the modes of variability

Figure S1: Train-test split diagram

Figure S2: Count and Variance time series

Figure S3: Resolution (1M vs 8D) influence on drivers

Figure S4: OceanCarbNN Δ^*CO_2 sensitivity to salinity

Table S1: Neural network architecture

Table S2: Regional metrics of OceanCarbNN fCO_2

Tables S3 – S5: Decomposition of modes of variability and drivers for
variances / standard deviation / percentages

Additional Supporting Information (Files uploaded separately)

Caption for Video S1

Introduction

The supplementary text contains information regarding the methodology of the main text. The figures and tables contain information regarding the results, except for Figure S1, and Table S1 which are both related to the methods.

Text S1: Data preprocessing

There is considerable preprocessing that is applied to each of the predictor variables before being passed to the regression step. The pre-processing can be separated into several steps: 1) gridding, 2) variable stacking and/or gap-filling, 3) separation of anomalies from climatologies, 4) include the previous time step as a predictor.

First, all variables are gridded to an 8-daily by $0.25^\circ \times 0.25^\circ$ resolution (abbreviated to 8D) where the first time-step of each year starts on January 1st. For daily datasets, this is easily achieved. However, for data with lower resolutions, data is weighted according to the contribution in time to that specific time step.

Variable stacking is applied to SSS, which is simply the use of a certain product if it is available, and if not, then use another lower-priority product. For SSS, we make use of three products with the following order of priority: ESA CCI v3.21 (Boutin et al., 2018), CMEMS-Multiobs (Droghei et al., 2016), and salinity from the Simple Ocean Data Assimilation (SODA 3.4.2; Carton et al., 2018). The ESA CCI data is purely satellite data, while the CMEMS-Multiobs data is a data-based product that incorporates satellite and float-based measurements of salinity. Lastly, the SODA salinity is data-assimilated model output that represents the upper 5 m of the ocean rather than the surface layer, which is the case of the other two products. The change from one data product to another impacts the long-term trends of $f\text{CO}_2$, but in the absence of a continuous time series over our period of interest (1982 onward), we resort to the variable stacking approach. Importantly, we provide a flag-variable for SSS indicating which product was used. Note that we also use SODA v3.4.2 density-based estimates of MLD as a predictor (Carton et al., 2018).

For CHL, we use the ESA-OCCCI v6 product for which there is no gap-filled version, meaning that cloud gaps are still present in the dataset. We use a three-step approach for filling these gaps. First, we use linear interpolation in both space and time, limited to ± 8 days and $\pm 0.25^\circ$. In the second step, we use a simplified version of the DINEOF approach (Alvera-Azcárate et al., 2011). We use a singular value decomposition (SVD) rather than an empirical orthogonal function (EOF) since this approach is functionally similar but computationally slightly more efficient. An 8-daily climatology of CHL is used as a first guess to speed up the optimization. Further, we fill the high latitudes of the 8-daily climatology with the long-term mean of CHL, which we then scale by photosynthetically available radiation so that the high latitudes have near-zero estimates of CHL-a during winter. The specific code for this function can be found in the [GitLab repository for OceanCarbNN](#).

The predictors are separated into climatological and anomalous components if time-series data are available for that variable (as also done by T. T. T. Chau et al., 2022; 2024). We create two types of climatological datasets: 1) the 8-daily climatologies ($X^{\text{clim}} = \bar{X}_j$), and 2) the difference of these 8-daily climatologies from the long-term mean (\bar{X}), which we call the seasonal anomaly ($X^{\text{season}} = \bar{X}_j - \bar{X}$). The anomalies are simply the difference between the predictor and its 8-daily climatology ($X^{\text{anom}} = X -$

\bar{X}_j). The purpose of this split is that the anomalies provide information of interannual variability when the data is available but is filled with zeros when there is no data. For example, CHL is only available from 1998 onward. Prior to this period, all CHL^{anom} data are set to zero. The climatological variables ($X^{\text{clim}}, X^{\text{season}}$) are used throughout the time series and capture seasonal cycle variability. The period over which the climatologies are calculated are specified in Figure 2. For SSH, we use absolute dynamic topography as the sea surface height and existing sea level anomalies from the Data Unification and Altimeter Combination System (DUACS) rather than calculating our own.

We also include information about the previous time step as a predictor, since this could add information about the rate of change of the predictor features, which could also impact $\Delta f\text{CO}_2$.

Finally, for training, we mask all regions where the ice fraction exceeds 85%, Hudson Bay, the Baltic Sea, Black Sea, and Red Sea. Masking the ice-covered regions may result in a loss of product coverage, but under-ice observations are scarce, meaning that we don't have a reliable way to test the predictions. These missing regions account for <1% of the global ocean.

Further, masking the marginal seas during training reduces the complexity of the target variable, where the marginal seas have complex marine carbonate chemistry due to sediment interaction, anoxic conditions, and injection of freshwater that is high in total alkalinity which often cannot be captured by the available predictors (Burt et al., 2016; Goyet et al., 1991; Müller et al., 2016). Note that these regions might not be masked during inference.

Text S2: Machine learning: OceanCarbNN details

Estimation of $\Delta f\text{CO}_2^{\text{clim}}$

Since we remove the atmospheric CO_2 mole fraction ($x\text{CO}_2$) rather than the fugacity, our target variable is $\Delta^* \text{CO}_2^{\text{clim}}$ averaged to an 8-daily climatology, since we assume that the atmospheric increase in $x\text{CO}_2$ represents the majority of long-term trend (Ma et al., 2023). We use Gradient Boosted Regression Trees (GBRT) from the LightGBM package in Python to perform the regression (Ke et al., 2017). The following non-default parameters are used:

- maximum depth for each tree = 5
- number of trees = 100

We include coordinates in the form of trigonometric transformations of latitude (lat), longitude (lon) and day of the year (d) after Gade (2010):

$$N = \begin{pmatrix} \sin\left(\frac{lat \cdot \pi}{180}\right) \\ \sin\left(\frac{lon \cdot \pi}{180}\right) \cdot \cos\left(\frac{lat \cdot \pi}{180}\right) \\ \cos\left(\frac{lon \cdot \pi}{180}\right) \cdot \cos\left(\frac{lat \cdot \pi}{180}\right) \end{pmatrix} \quad (S1)$$

$$D = \begin{pmatrix} \sin\left(\frac{d \pi}{365}\right) \\ \cos\left(\frac{d \pi}{365}\right) \end{pmatrix} \quad (S2)$$

Here N contains three vectors that represent the spatial coordinates, and D contains the temporal coordinates.

The estimate of $\Delta^* \text{CO}_2^{\text{clim}}$ is smoothed with a rolling mean with a two-month window (seven 8-day time steps) in the time dimension and a window width of 0.75° in the latitudinal and longitudinal dimensions. Note that we allow for a large temporal window, but narrow spatial window to preserve spatial gradients, particularly in the coastal regions. Importantly, the long rolling window in the time dimension means that our estimates are not overfit.

Train-test-validation split

To avoid over-fitting of our model, we split our data into training, testing, and validation datasets. We create test splits using the same approach as Bennington et al. (2022), where every seventh month is considered a test month. This amounts to roughly 15% of data being reserved for testing. We also set data aside for validation as every seventh month with a three-month starting offset.

Given that we perform our train-test split based on every seven months, we can create seven train-test-validation splits, each starting on a different month in 1982 (January through July). This is useful since we can create reconstruction of test uncertainties corresponding with the full dataset.

Neural Network Architecture and optimization

Recent studies have shown that GBMs may be better at predicting surface ocean CO_2 (Gloege et al., 2022). However, tree-based methods are not able to extrapolate outside the observed dataset. Recent work has shown that Neural Networks (NN) that are not too wide or deep are able to extrapolate beyond the training observations (Courtois et al., 2023). Further, it was shown that NNs are better able to capture the true relationship between a target variable and its predictors than random forests (Holder & Gnanadesikan, 2021). We thus opt to use NNs to perform the regression step.

We use an ensemble of FFNNs to predict $\Delta^* \text{CO}_2$, with five FFNNs per train-test-validation split, resulting in a total of 35 ensemble members. We use TensorFlow and Keras to construct our neural network ensemble (Martín Abadi et al., 2015).

Each ensemble member has the same architecture, with an overview shown in Table S1. First, all predictors are passed through a normalization layer, so that the mean = 0 and the standard deviation = 1 for each predictor. We use two fully connected hidden layers with 64 and 32 neurons. Given that we use 15 predictor variables, this amounts to ~3'000 trainable weights, in the order of that used in the SOMFFN approach (Landschützer et al., 2016). Grid search was used to find the optimal network architecture and hyperparameters. Batch normalization is applied after each hidden layer to speed up the convergence of the weights. A dropout layer reduces potential over-fitting by randomly dropping 10% of neurons during training. We then use a Rectified Linear Unit (ReLU) activation function, as it has been proven to be accurate and efficient for non-linear regression problems (Fukushima, 1969).

We use the Adam optimizer with a mean squared error (MSE) loss function to perform back-propagation using a batch size of 2048. We use an initial learning rate of 0.01 that decreases when no improvement is detected in the validation subset for 10 epochs. If no improvement is detected in the validation scores after 15 epochs, the training process is stopped.

Text S3: Evaluation

Metrics

We assess the performance of our model on $f\text{CO}_2$. Bias, mean absolute difference (MAD) root mean squared difference (RMSD) are used to assess our model.

$$Bias = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y) \quad (S3)$$

$$MAD = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (S4)$$

$$RMSD = \sqrt{\sum_{i=1}^n \frac{(y_i - \hat{y}_i)^2}{n}} \quad (S5)$$

where y_i is the observed (transformed) target at location i , \hat{y}_i is the estimated value for that location, and n is the number of train/test/validation observations. We provide assessment of these metrics for the global, open, and coastal ocean defined as the region within 300 km from the coast or the 1000 m isobath (Laruelle et al., 2017). Further, we provide metrics for three latitudinal bands: the high northern ($> 35^\circ\text{N}$) and southern latitudes ($> 35^\circ\text{S}$) and the bounded lower latitudes.

Note that we use the term *difference* instead of the typically used *error*, since the latter implies that we are predicting a quantity that is perfectly known. However, $f\text{CO}_2$ carries uncertainties associated with measurement and gridding (Bakker et al., 2016).

Evaluation data sets

We evaluate against autonomous open ocean mooring stations with $f\text{CO}_2$ from Sutton et al., (2019). Note that we do not apply quality control to these moorings. An important consideration is that these data are also used in the SOCAT database. However, as with the assessment of our model performance, we use only test predictions. Meaning that the mooring data have not been used to train the ensemble members that produce the test output.

Gradient analysis

To quantify the sharpness of $f\text{CO}_2$ gradients, we use the mean gradients over a domain:

$$\|\nabla f\text{CO}_2\| = \sqrt{\frac{\delta f\text{CO}_2^2}{\delta x} + \frac{\delta f\text{CO}_2^2}{\delta y}}$$

where $\frac{\delta f\text{CO}_2}{\delta x}$ is the zonal partial derivative (i.e., longitude) of $f\text{CO}_2$, and the equivalent applies to latitude (y) in the meridional dimension. Note that latitude and longitude are converted to kilometers. The normed sum of the derivatives is taken to represent the gradients in the units $\mu\text{atm km}^{-1}$.

Text S4: Integration of Fluxes

We calculate the integrated flux (F) with:

$$F = \sum_{xy} F_{xy} \cdot A_{xy} \quad (S6)$$

Where xy represents a pixel with a latitude and longitude coordinate, F_{xy} represents the flux per unit area for a specific pixel, and A_{xy} represents the area per pixel. Here the units of F_{xy} are grams of carbon per meter squared per day, and the units of A are meters squared. The units of the integrated fluxes are thus grams of carbon per day. This can easily be converted to the typical Peta grams per year with $F \cdot \frac{365}{10^{15}}$, where 365 is days per year and 10^{15} is grams per Peta gram.

Text S5: Computing the modes of variability

We use the Fast Fourier Transform (FFT) to compute the power spectral density, $S(x)$, of each pixel through time after data has been detrended over time. Using Parseval's Theorem, we can integrate over the frequencies of interest to approximate the variances:

$$\sigma^2 \approx 2 \int_{f_1}^{f_2} S(x) df \quad (S7)$$

where σ^2 is the variance in the time domain, f is a given frequency, and f_1 and f_2 represent the upper and lower frequencies. We thus separate the variability in the frequency space into three modes by integrating over the frequencies in the Fourier domain after Gu et al. (2023): sub-seasonal, < 3 month frequencies; seasonal, 3 – 15

month frequencies; and interannual, > 15 month frequencies. Later, we also separate the interannual variability into sub-decadal (15 months to 8 years) and decadal (> 8 years) variability.

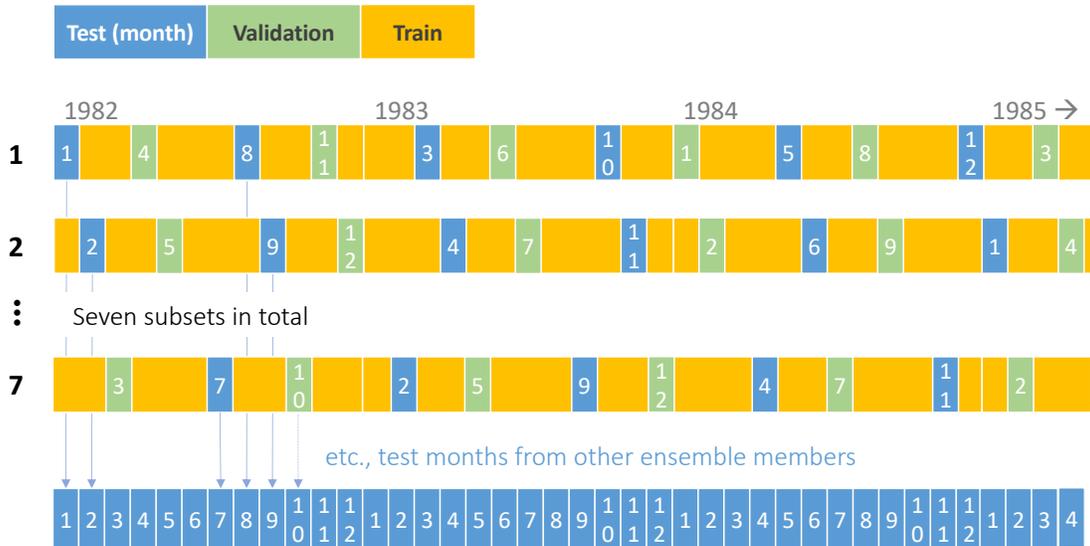


Figure S1. A schematic showing the concept of the full training dataset reconstruction. Since every seventh month is selected as a training month, we can create seven training subsets, where the starting month numbers from 1-7 (January to July). Gathering all the test-months allows for a full reconstruction consisting only of test estimates.

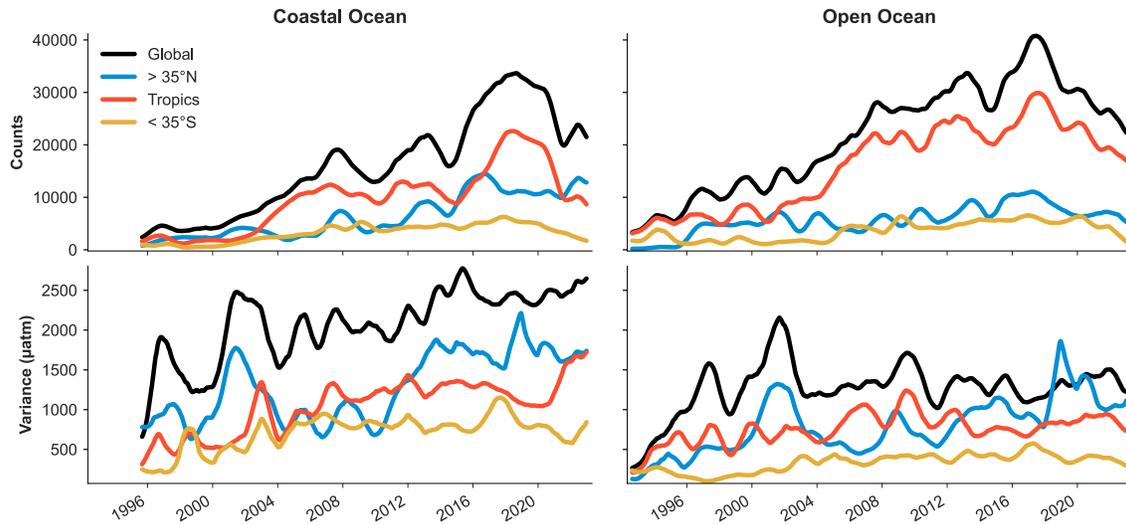


Figure S2: Increasing number of coastal ocean observations north of 35°N, which also leads to increased variance for that region. This contributes to decreasing RMSD estimates over time.

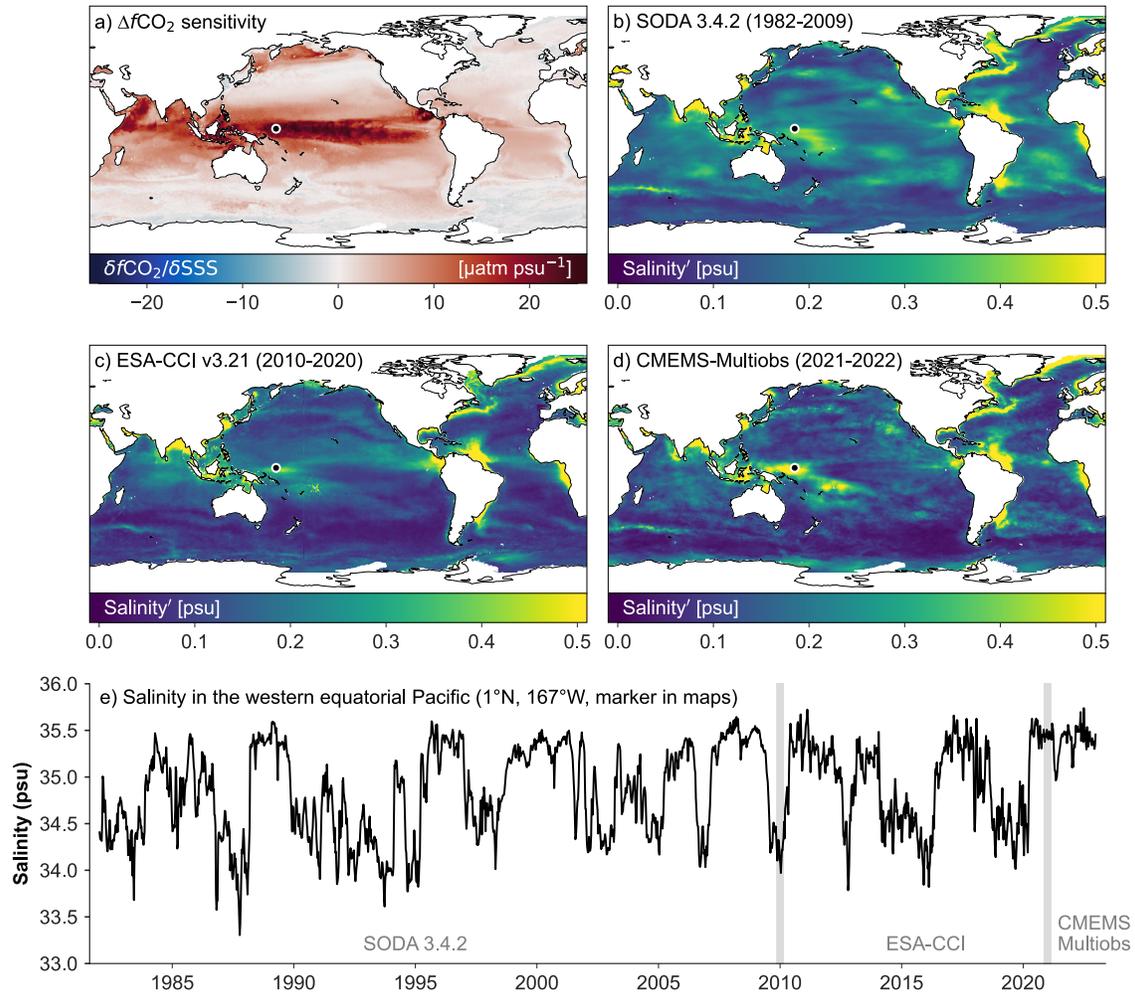


Figure S3: (a) The sensitivity of $\Delta f\text{CO}_2$ to salinity anomalies for a mean state ocean, i.e., all anomalies are set to 0 except for salinity anomalies. (b-d) Salinity anomalies are averaged over three different decades to show the variability of each salinity product. (e) Salinity time series in the western equatorial Pacific shown by the marker in (a-d). The vertical lines denote the change in the salinity products (aligned with first, second, and third rows) also indicated in Figure 2. Multiplying (a) by (b-d) gives the expected change in $\Delta f\text{CO}_2$.

Table S1. Table showing the structure of a feed-forward neural network ensemble member used to predict $\Delta^* \text{CO}_2^{\text{nonT}}$. We also show the number of trainable weights for the fully connected layers. This does not include the bias terms associated with the Activation functions.

Layer type	Description
Inputs	15 pre-processed variables for input
Normalization	Scales input variables so $\mu = 0$ and $\sigma = 1$
Fully connected	64 neurons, 960 weights
Batch normalization	Scales weights per batch to speed up training time
Dropout	10% of neurons dropped during training to prevent over-fitting
ReLu Activation	Rectified Linear Unit activation for non-linear responses
Fully connected	32 neurons, 2'048 weights
Batch normalization	Scales weights per batch to speed up training time
Dropout	10% of neurons dropped during training to prevent over-fitting
ReLu Activation	Rectified Linear Unit activation for non-linear responses
Fully connected	1 neuron, 32 weights

Table S2: Table showing unweighted metrics for different masks and regions. MAD = mean absolute difference, RMSD = root mean squared difference, r^2 = variance explained. All columns have units μatm , except r^2 . The mask used to define oceanic and coastal areas is from RECCAP2 (DeVries et al., 2023), the remaining regions are the biomes from Fay et al. (2014).

	Bias	MAD	RMSD	r^2 -score
Global	-0.13	11.3	19.2	0.81
Oceanic	0.07	8.81	13.1	0.87
Coastal	-0.4	14.8	25.3	0.76
NP-ICE	-0.18	31.2	43.5	0.62
NP-SPSS	-0.18	18.3	28.9	0.73
NP-STSS	0.64	7.9	11.5	0.88
NP-STPS	-0.12	6.4	8.8	0.91
PEQU-W	-0.07	8.3	11.1	0.72
PEQU-E	0.22	17.5	24.3	0.82
SP-STPS	0.13	6.9	9.9	0.92
NA-ICE	-1.7	23	30.5	0.58
NA-SPSS	-0.39	14	20.8	0.79
NA-STSS	-0.16	8.8	12.2	0.8
NA-STPS	-0.44	6.7	10.3	0.83
AEQU	-0.21	10.2	14	0.75
SA-STPS	-0.74	7.3	9.8	0.83
IND-STPS	0.9	6.8	9	0.85
SO-STSS	-0.04	9.9	15.3	0.72
SO-SPSS	0.4	9.3	13.5	0.68
SO-ICE	0.03	19.2	27.6	0.7

Table S3: Global area-weighted averages of the variance of fluxes ($\text{mmolC m}^{-2} \text{ day}^{-1}$). The first four columns show the variance for the different components for the different modes of variability. Σ shows the total variance. Note that the covariance is calculated as the total variance less the sum of the remaining columns.

		Wind	$f\text{CO}_2$	Cross-term	Covariance	Σ
Subseasonal	1M	0.15	0.08	0.1	-0.11	0.22
	8D	1.24	0.34	0.67	-0.71	1.54
Seasonal	1M	1.22	6.41	0.49	-1.52	6.59
	8D	1.26	6.52	0.53	-1.55	6.76
Subdecadal	1M	0.09	0.13	0.04	-0.06	0.21
	8D	0.09	0.14	0.04	-0.06	0.22
Decadal	1M	0.02	0.03	0.01	0	0.06
	8D	0.02	0.04	0.01	0	0.07

Table S4: Same as Table S3 but for standard deviation. Covariance is calculated as the square root of the absolute value of the variance.

		Wind	$f\text{CO}_2$	Cross-term	Covariance	Σ
Subseasonal	1M	0.39	0.28	0.32	-0.34	0.47
	8D	1.11	0.58	0.82	-0.84	1.24
Seasonal	1M	1.1	2.53	0.7	-1.23	2.57
	8D	1.12	2.55	0.73	-1.24	2.6
Subdecadal	1M	0.3	0.36	0.2	-0.24	0.46
	8D	0.31	0.38	0.21	-0.25	0.46
Decadal	1M	0.15	0.19	0.08	-0.02	0.25
	8D	0.15	0.19	0.08	-0.03	0.26

Table S5: Same as Table S3 but for percent contribution to total variance.

		Wind	$f\text{CO}_2$	Cross-term	Covariance
Subseasonal	1M	70	35	46	-51
	8D	81	22	44	-46
Seasonal	1M	18	97	7	-23
	8D	19	96	8	-23
Subdecadal	1M	44	63	20	-27
	8D	43	65	20	-28
Decadal	1M	35	55	11	-1
	8D	34	56	11	-2

Video S1. Air-sea CO_2 fluxes calculated at a 4-hourly temporal resolution using Eq. 3. The 8-daily $\Delta f\text{CO}_2$ ($f\text{CO}_2 - f\text{CO}_2^{\text{atm}}$) is linearly interpolated to a 4-hourly resolution. Similarly, daily SST, ICE and SSS are linearly interpolated to a 4-hourly resolution to calculate k_w and K_0 . In the calculation of k_w the winds are averaged from the original hourly resolution, to 4-hourly. Data was prepared by Luke Gregor; animation was created by Eike Köhn.

References

- Alvera-Azcárate, A., Barth, A., Sirjacobs, D., Lenartz, F., & Beckers, J. M. (2011). Data interpolating empirical orthogonal functions (DINEOF): A tool for geophysical data analyses. *Mediterranean Marine Science*, *12*(3), 5–11. <https://doi.org/10.12681/mms.64>
- Bakker, D. C., Pfeil, B., Landa, C. S., Metzl, N., O'Brien, K. M., Olsen, A., Smith, K., Cosca, C., Harasawa, S., Jones, S. D., Nakaoka, S., Nojiri, Y., Schuster, U., Steinhoff, T., Sweeney, C., Takahashi, T. T., Tilbrook, B., Wada, C., Wanninkhof, R. H., ... Xu, S. (2016). A multi-decade record of high-quality fCO₂ data in version 3 of the Surface Ocean CO₂ Atlas (SOCAT). *Earth System Science Data*, *8*(2), 383–413. <https://doi.org/10.5194/essd-8-383-2016>
- Bennington, V., Galjanic, T., & McKinley, G. A. (2022). Explicit Physical Knowledge in Machine Learning for Ocean Carbon Flux Reconstruction: The pCO₂-Residual Method. *Journal of Advances in Modeling Earth Systems*, *14*(10). <https://doi.org/10.1029/2021MS002960>
- Boutin, J., Vergely, J. L., Marchand, S., D'Amico, F., Hasson, A., Kolodziejczyk, N., Reul, N., Reverdin, G., & Vialard, J. (2018). New SMOS Sea Surface Salinity with reduced systematic errors and improved variability. *Remote Sensing of Environment*, *214*(August 2017), 115–134. <https://doi.org/10.1016/j.rse.2018.05.022>
- Burt, W. J., Thomas, H., Miller, L. A., Granskog, M. A., Papakyriakou, T. N., & Pengelly, L. (2016). Inorganic carbon cycling and biogeochemical processes in an Arctic inland sea (Hudson Bay). *Biogeosciences*, *13*(16), 4659–4671. <https://doi.org/10.5194/BG-13-4659-2016>
- Carton, J. A., Chepurin, G. A., & Chen, L. (2018). SODA3: A new ocean climate reanalysis. *Journal of Climate*, *31*(17), 6967–6983. <https://doi.org/10.1175/JCLI-D-17-0149.1>
- Chau, T. T. T., Gehlen, M., & Chevallier, F. (2022). A seamless ensemble-based reconstruction of surface ocean pCO₂ and air-sea CO₂ fluxes over the global coastal and open oceans. *Biogeosciences*, *19*(4), 1087–1109. <https://doi.org/10.5194/bg-19-1087-2022>
- Chau, T.-T.-T., Gehlen, M., Metzl, N., & Chevallier, F. (2024). CMEMS-LSCE: A global, 0.25°, monthly reconstruction of the surface ocean carbonate system. *Earth System Science Data*, *16*(1), 121–160. <https://doi.org/10.5194/essd-16-121-2024>
- Courtois, A., Morel, J. M., & Arias, P. (2023). Can neural networks extrapolate? Discussion of a theorem by Pedro Domingos. *Revista de La Real Academia de Ciencias Exactas, Fisicas y Naturales - Serie A: Matematicas*, *117*(2), 1–26. <https://doi.org/10.1007/s13398-023-01411-z>
- DeVries, T., Yamamoto, K., Wanninkhof, R., Gruber, N., Hauck, J., Müller, J. D., Bopp, L., Carroll, D., Carter, B., Chau, T.-T.-T., Doney, S. C., Gehlen, M., Gloege, L., Gregor, L., Henson, S., Kim, J. H., Iida, Y., Ilyina, T., Landschützer, P., ... Zeng, J. (2023). Magnitude, trends, and variability of the global ocean carbon sink from 1985–2018. *Global Biogeochemical Cycles*, *n/a*(n/a), e2023GB007780. <https://doi.org/10.1029/2023GB007780>

Droghei, R., Buongiorno Nardelli, B., & Santoleri, R. (2016). Combining in situ and satellite observations to retrieve salinity and density at the ocean surface. *Journal of Atmospheric and Oceanic Technology*, 33(6), 1211–1223. <https://doi.org/10.1175/JTECH-D-15-0194.1>

Fay, A. R., & McKinley, G. A. (2014). Global open-ocean biomes: Mean and temporal variability. *Earth System Science Data*, 6(2), 273–284. <https://doi.org/10.5194/essd-6-273-2014>

Fukushima, K. (1969). Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements. *IEEE Transactions on Systems Science and Cybernetics*, 5(4), 322–333. <https://doi.org/10.1109/TSSC.1969.300225>

Gade, K. (2010). A Non-singular Horizontal Position Representation. *Journal of Navigation*, 63(03), 395–417. <https://doi.org/10.1017/S0373463309990415>

Gloege, L., Yan, M., Zheng, T., & McKinley, G. A. (2022). Improved Quantification of Ocean Carbon Uptake by Using Machine Learning to Merge Global Models and pCO₂ Data. *Journal of Advances in Modeling Earth Systems*, 14(2), 1–19. <https://doi.org/10.1029/2021MS002620>

Goyet, C., Bradshaw, A. L., & Brewer, P. G. (1991). The carbonate system in the Black Sea. *Deep Sea Research Part A. Oceanographic Research Papers*, 38(Suppl. 2A), S1049–S1068. [https://doi.org/10.1016/S0198-0149\(10\)80023-8](https://doi.org/10.1016/S0198-0149(10)80023-8)

Gu, Y., Katul, G. G., & Cassar, N. (2023). Multiscale Temporal Variability of the Global Air-Sea CO₂ Flux Anomaly. *Journal of Geophysical Research: Biogeosciences*, 128(6). <https://doi.org/10.1029/2022JG006934>

Holder, C., & Gnanadesikan, A. (2021). Can machine learning extract the mechanisms controlling phytoplankton growth from large-scale observations?—A proof-of-concept study. *Biogeosciences*, 18(6), 1941–1970. <https://doi.org/10.5194/BG-18-1941-2021>

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T. Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems, 2017-Decem(Nips)*, 3147–3155.

Landschützer, P., Gruber, N., & Bakker, D. C. (2016). Decadal variations and trends of the global ocean carbon sink. *Global Biogeochemical Cycles*, 30(10), 1396–1417. <https://doi.org/10.1002/2015GB005359>

Laruelle, G. G., Landschützer, P., Gruber, N., Ti, J. L., Delille, B., Regnier, P. A. G., & Tison, J.-L. (2017). Global high-resolution monthly pCO₂ climatology for the coastal ocean derived from neural network interpolation. *Biogeosciences*, 14(19), 4545–4561. <https://doi.org/10.5194/bg-14-4545-2017>

Ma, D., Gregor, L., & Gruber, N. (2023). Four Decades of Trends and Drivers of Global Surface Ocean Acidification. *Global Biogeochemical Cycles*, 37(7), e2023GB007765. <https://doi.org/10.1029/2023GB007765>

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat,

Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Jia, Y., Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, ... Xiaoqiang Zheng. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. <https://www.tensorflow.org/>

Müller, J. D., Schneider, B., & Rehder, G. (2016). Long-term alkalinity trends in the Baltic Sea and their implications for CO₂-induced acidification. *Limnology and Oceanography*, *61*(6), 1984–2002. <https://doi.org/10.1002/lno.10349>

Sutton, A. J., Feely, R. A., Maenner-Jones, S., Musielwicz, S., Osborne, J., Dietrich, C., Monacci, N., Cross, J., Bott, R., Kozyr, A., Andersson, A. J., Bates, N. R., Cai, W. J., Cronin, M. F., De Carlo, E. H., Hales, B., Howden, S. D., Lee, C. M., Manzello, D. P., ... Weller, R. A. (2019). Autonomous seawater pCO₂ and pH time series from 40 surface buoys and the emergence of anthropogenic trends. *Earth System Science Data*, *11*(1), 421–439. <https://doi.org/10.5194/essd-11-421-2019>