

Reconstructing ocean carbon storage with CMIP6 models and synthetic Argo observations

Katherine E. Turner¹, Doug M. Smith², Anna Katavouta^{1,3}, and Richard G. Williams¹

¹Department of Earth, Ocean, and Ecological Sciences, School of Environmental Sciences University of Liverpool, Liverpool, United Kingdom

²UK Met Office, Exeter, United Kingdom

³National Oceanography Centre, Liverpool, United Kingdom

Correspondence: Katherine E. Turner (K.E.Turner2@liverpool.ac.uk)

Abstract. The ocean carbon store plays a vital role in setting the carbon response to emissions and variability in the carbon cycle. However, due to the ocean's strong regional and temporal variability, sparse carbon observations limit our understanding of historical carbon changes. We explore how ~~widespread~~ Argo temperature and salinity profiles can provide information to reconstruct ocean carbon inventories with Ensemble Optimal Interpolation. [Here](#), Ensemble Optimal Interpolation ~~draws upon~~ ~~first-order relationships between variables, and high-complexity climate models been used previously to reconstruct interior ocean heat content and salinity changes from Argo profiles. The CMIP6 model ensemble shows~~ [is used to reconstruct ocean carbon using synthetic Argo temperature and salinity observations, with examples for both the top 100m and top 2000m carbon inventories. When considering reconstructions of the top 100m carbon inventory,](#) coherent relationships between upper-ocean carbon, temperature, salinity, and atmospheric CO₂ ~~that result in~~ [result in optimal](#) solutions that reflect the controls of under-saturation, solubility, and alkalinity. ~~The Ensemble Optimal Interpolation method is tested with synthetic reconstructions of upper-ocean carbon fields using perfect information of temperature and salinity changes, along with atmospheric CO₂ concentrations. Sensitivity tests of~~ [Out-of-sample reconstructions of the top 100m](#) show that, in most regions, the trend in ocean carbon and over 60% of detrended variability can be reconstructed using local temperature and salinity measurements. ~~Expanding the synthetic reconstructions to include irregular sampling consistent with Argo profile locations results in a similar capacity to reconstruct ocean carbon variability, as the increased information provided from multiple sampling locations compensates for the propagation of errors within the CMIP6 covariance fields. Overall, our work indicates that,~~ [with only small changes when considering synthetic profiles consistent with irregular Argo sampling. Extending the method to reconstruct the upper 2000m reveals that model uncertainties at depth limit the reconstruction skill. The impact of these uncertainties on reconstructing the carbon inventory over the upper 2000m is small, and full reconstructions with historical](#) ~~Argo locations show that the method can reconstruct regional interannual and decadal variability. Hence, optimal interpolation based on model relationships combined with~~ hydrographic measurements can provide valuable information about ~~changes in the global~~ ocean carbon inventory. ~~The flexibility of optimal interpolation indicates a strong potential for new, independent, and mechanistic estimates of the historical carbon inventory that can be used to compare with other estimates of the ocean carbon system. changes.~~

25 1 Introduction

The global ocean plays a important role in the carbon cycle, being both a major reservoir of carbon and substantial sink of anthropogenic carbon transferred from geological storage to the coupled atmosphere-ocean-terrestrial system. ~~Recent studies estimate that the ocean has~~ The ocean is estimated to have taken up around 26% of anthropogenic emissions from fossil fuels and land use changes since 1850 ~~(?)~~ ~~(???)~~. In addition to the long-term uptake of anthropogenic carbon, ocean carbon uptake exhibits interannual and decadal variability on global and regional scales ~~(???)~~. This variability impacts the ability to detect trends in observations of the partial pressure of CO₂ in seawater (?) and, through oceanic transport, can lead to regionally enhanced acidification ~~(?)~~ ~~(?)~~. The ocean carbon inventory is thus an important ~~aspect of current~~ integral measure of climate change, and up-to-date estimates on its behaviour are vital for understanding its evolution and impacts on the climate system.

In order to characterise the ocean carbon system with its strong regional and temporal variability, extensive spatial and temporal coverage in observations are required. Whereas there are ~~millions of~~ sufficient surface pCO₂ measurements to provide a surface view of the ocean carbon system and air-sea carbon fluxes (?), complementary observations of ocean interior dissolved inorganic carbon (DIC) are limited by the logistics of ship-based bottle measurements. Campaigns for ocean interior DIC observations use repeat transects to provide high-quality observations for many regions (?). The transects resolve spatial variations in carbon but are limited in their ability to resolve temporal variability as they are repeated on decadal timescales. Alternatively, ocean DIC time series such as those found at the Bermuda Atlantic Time Series and the Hawaii Ocean Timeseries can resolve seasonal, interannual, and decadal variability, but are limited in how well they represent variability on larger spatial scales (?). The Bio-Argo programme has allowed for more autonomous sampling of ocean interior carbonate system variables such as pH (?), from which interior DIC can be estimated; currently, however, autonomous measurement methods for interior DIC remain in development.

By itself, the sparsity of interior observations hinders the ability to produce a coherent, global picture of recent ocean carbon changes from a storage perspective. A mapping technique is necessary to expand points of observations into more coherent spatial patterns of change and behaviour. Mapping techniques employ statistics to propagate information from observations to unobserved regions. Time-invariant climatologies of ocean carbon have been created using mapping procedure on data from repeat transects and other ship-based observations (?). ~~For observations of surface pCO₂, non-linear~~ Non-linear machine learning procedures have been able to reconstruct spatial patterns in interannual and decadal variability for surface pCO₂ ~~(???)~~ ~~, although these algorithms can lead to substantial biases if observations are too sparse for certain regions as well as the seasonal cycle for DIC (?), but the non-linearities can create substantial biases for regions with sparse observations (?).~~ Data assimilative methods that use observations to constrain model physics in forward experiments have also been expanded to include biogeochemistry (??).

For any mapping or data assimilative technique, it is imperative to use accurate statistics to avoid the erroneous propagation of information from observations. These statistical relationships can be parameterised or calculated from observations; however, parameterised statistics fail to reflect regional differences in the ocean, and for poorly observed variables, averaging over large length or timescales can lead to overly coarse covariance fields. State-of-the-art climate models provide complete

pseudo-data in both space and time and therefore may be used to calculate these statistics, though model biases may lead to errors in the covariances. Nevertheless, fully-coupled climate models have been used to reconstruct ocean heat content trends and variability from observations, with clear improvements in recent years due to the expansion of the Argo programme (???)

Ocean carbonate chemistry is controlled by both physical and biogeochemical processes. There are well-understood first-order principles that relate to ocean temperature and salinity to upper ocean carbon; using these relationships, the observations used in previous heat content mappings may be ~~useful~~ exploited in a similar manner for carbon mappings. Increases in temperature reduce the ocean's ability to take up CO₂ through gas solubility laws, while increases in salinity increase the ocean's ability to take up CO₂ by increasing alkalinity (?). Additionally, temperature and salinity provide constraints on ocean circulation, which alters the background vertical gradients of both heat and carbon (??). Anthropogenic carbon uptake in the high latitudes is constrained by salinity and stratification, which can be taken to be proxies for water mass formation (???). The relationships between temperature, salinity, and carbon are regionally dependent as ocean dynamics and biology can set different drivers of CO₂ uptake (?). If these relationships can be exploited, the increase in ocean observations from the Argo programme may provide valuable information that can help reconstruct ocean carbon fields alongside temperature and salinity fields. While the ocean carbonate system can also be approximated using other observations such as pH and salinity, for this study we focus on the ~~readily-available-and-widespread~~ potential benefit of observations present within the Argo dataset.

In this study, we apply an Ensemble Optimal Interpolation approach to reconstruct modelled upper-ocean dissolved inorganic carbon from synthetic temperature and salinity observations. ~~To calculate the covariance fields, we construct~~ Covariance fields are constructed using an ensemble of 6 Earth system models from the Climate Model Intercomparison Project Phase 6 (CMIP6) ~~and examine their relationships between upper-ocean carbon and other climate variables~~. In this "proof-of-concept" ~~study we aim to~~ proof-of-concept study we show the potential skill available in using model covariance fields and Argo-style ~~measurements~~. ~~We create synthetic~~ synthetic measurements to reconstruct carbon content between 0-2000m. Synthetic reconstructions of modelled ocean carbon ~~using perfect~~ are created by using pseudo-observations of temperature and salinity, similar to the ocean heat content synthetic reconstructions in ?. ~~We then compare the~~ The errors within these reconstructions ~~against~~ can then be compared with those from the climatological fields to see where and how the method best works in the model world.

~~We set out the~~ The work in this study is set out as follows. ~~In Section 2 we describe~~ Section 2 introduces the Ensemble Optimal Interpolation scheme for ocean DIC using an ensemble of 6 CMIP6 models ~~, with an overview of the experiments we conduct and the experiments used~~ to test the reconstruction skill. ~~In Section 3.1 we present~~ The method is first assessed by reconstructing DIC inventory changes over the top 100m, and then over the top 2000m. Section 3.1 presents the ensemble correlation fields between DIC over the top 100m and temperature and salinity at the same location, with a discussion as to how these correlations relate to physical controls on the ocean carbon response. In Sections 3.1 and 3.1 ~~we present sensitivity results~~ to discuss the reconstruction potential of different temperature and salinity sampling distributions is assessed for the upper 100m, ranging from perfect spatial coverage to coverage more representative of ~~ARGO observations, with global and regional comparisons of the different synthetic Argo reconstructions~~ Argo observations. The method is then extended to consider the full 2000m profiled by Argo floats in Section 4, and the potential to reconstruct recent global carbon changes is illustrated with

[a reconstruction using time-varying synthetic Argo observations](#) in Section 5. Lastly, in Section 6 we discuss the potential and shortcomings of this method, both in terms of the setup with temperature and salinity and [further expansions how the method may be expanded](#) with other oceanographic observations.

2 The Ensemble Optimal Interpolation method for ocean carbon

Optimal interpolation is a non-dynamical mapping approach that uses weights to propagate information from observations to regions without observations (??). The optimal interpolation method involves creating an analysis A at locations i and times t from the sum of a background state B_i and a weighted sum of the difference between the observed and background values, also known as the observation increments. The observations can be the same property as the final analysis, or can be of other properties, i.e. temperature and salinity observations can be used to reconstruct DIC if there is a physical link between them. The inclusion of these sorts of physically relevant data produces a *multivariate analysis*. For observations O at locations k within a sampling region K , the optimal interpolation method can be represented as

$$A_i = B_i + \sum_{k \in K} w_k (O_k - B_k). \quad (1)$$

The crux of the optimal interpolation problem is thus finding a suitable solution for the weights w_k . The optimal weights are those that minimise the expected analysis error at each gridpoint i , calculated as the root mean squared error (RMSE) between the analysis and the truth T_i :

$$RMSE(A_i, T_i) = \sqrt{\sum_t (A_i(t) - T_i(t))^2}. \quad (2)$$

In this work we explore how the relationships between DIC and observed ocean variables such as temperature and salinity can be used to reconstruct upper ocean carbon [stocks inventories down to a depth of 2000m](#). The use of optimal interpolation to reconstruct ocean carbon from available carbon measurements is inherently limited by the poor temporal and spatial coverage of existing DIC observations. However, temperature and salinity observations are far more plentiful, particularly in the Argo era, so we explore the extent to which these observations could be used to reconstruct DIC. To avoid the problems associated with sparse input data, we take a multivariate analysis approach to reconstruct carbon from extensive synthetic observations of ocean temperature T and salinity S consistent with observations from the Argo programme (?), as well as annual average atmospheric CO_2 concentrations. On annual and longer timescales, atmospheric CO_2 is well-mixed, and so the use of annual atmospheric CO_2 measurements can allow the analysis to capture longer-term DIC changes from changes in the global carbon budget. With background fields for DIC, T, and S taken to be their global climatologies (i.e., $DIC(i, t) = \overline{DIC}(i) + DIC'(i, t)$, where $\overline{DIC}(i)$ is the temporal average DIC concentration at location i over the period 1955-2014), the optimal interpolation scheme is formulated to calculate the residual DIC' from T' , S' within a region K :

$$DIC'(i, t) = \sum_{k \in K} (w_{C,k} \text{pCO}_2'(t) + w_{T,k} T'(k, t) + w_{S,k} S'(k, t)), \quad (3)$$

where $w_{C,k}$ is the local weighting for the annual global averaged atmospheric pCO₂ and $w_{T,k}$ and $w_{S,k}$ are the weightings for observed ocean temperatures and salinities within the sampling region K . These optimal weights thus describe how information is propagated from atmospheric pCO₂ and hydrographic observations to the ocean carbon system, taking into account the interdependencies between input variables.

2.1 Choices for calculating background error covariances

The optimal weights w_k in (1) and (3) are determined by the covariances between the background errors $B_i - T_i$. The covariance fields describe how information should be propagated from areas with observations to those without; the optimal solution will consider how much new information observations provide to the reconstruction and depend on local and larger-scale relationships within the climate system. Any ensemble optimal interpolation method must involve decisions made on how the background error covariances are calculated, and which covariances are included to solve for the weights w_k . In the following we describe how we have made these decisions using a CMIP6 multi-model ensemble and with various assumptions as to which observations are used for the multivariate DIC analyses.

2.1.1 Background covariances from a CMIP6 ensemble

Climate model outputs can be used to provide covariance fields for optimal interpolation solutions. By providing complete pseudo-data in both space and time, climate models avoid some of the errors that arise from the coarseness of observational or parameterised covariance fields. Background covariance fields from global climate models have been used to reconstruct ocean temperatures and salinities from observations as well as to initialise decadal forecasts (??). However, the model ensemble background covariance field will still contain errors that will need to be evaluated using sensitivity testing before the fields can be used with real-world observations.

To construct the background covariance fields, we obtained ocean potential temperature, ocean salinity, and ocean DIC model output from 6 CMIP6 Earth System models with a nominal horizontal resolution around 1° (Table 1). The output was taken from the historical experiments and covers the period of year 1955 to year 2014. This period was chosen as it has consistent behaviour in atmospheric CO₂ concentrations and is long enough to allow some multidecadal variability to be captured in the covariance fields. For each model, output from 5 realisations ~~was~~ were used so that the models' internal variability was captured by the ensemble.

For all variables, annual averages were calculated from monthly mean outputs, and outputs were bilinearly regridded from their native horizontal grids to a 1° × 1° grid using the Python package xESMF (<https://doi.org/10.5281/zenodo.1134365>). Oceanic variables were further integrated vertically ~~, with the results in this study focusing on the results from the integrated top-100m~~ layers that span the ocean surface to 2000m depth. The integrated layers were chosen to be 0m-100m, 100m-500m, 500m-1000m, and 1000-2000m.

Table 1. CMIP6 Earth System Models and realisations used for reconstruction

Model (Reference)	Realisations
ACCESS-ESM-1.5 (?)	r1i1p1f1, r2i1p1f1, r4i1p1f1, r5i1p1f1, r6i1p1f1
CanESM5 (?)	r10i1p1f1, r11i1p1f1, r12i1p1f1, r13i1p1f1, r14i1p1f1
CESM2 (?)	r1i1p1f1, r2i1p1f1, r3i1p1f1, r4i1p1f1, r5i1p1f1
IPSL-CM6A-LR (?)	r1i1p1f1, r2i1p1f1, r3i1p1f1, r4i1p1f1, r32i1p1f1
MPI-ESM1.2-LR (?)	r1i1p1f1, r2i1p1f1, r3i1p1f1, r4i1p1f1, r5i1p1f1
UKESM1 (?)	r1i1p1f2, r2i1p1f2, r3i1p1f2, r4i1p1f2, r8i1p1f2

2.1.2 Spatial limits on background covariances

The sampling region K for temperature and salinity observations in (3) can vary from co-located observations to global observations; the choice of sampling region requires balancing the extra information provided by additional observations with spurious propagation of information through errors in the covariance fields. In this study we use two idealised sampling methods to explore the limits of this first-order reconstruction of ocean carbon:

1. ~~Observations are taken from at each ocean grid cell at each year, therefore assuming perfect~~ The first reconstruction method assumes full global coverage of temperature and salinity profiles. The reconstruction is then set up to use local relationships between ocean carbon, temperature, and salinity observations. For these synthetic reconstructions, the ocean inputs to reconstruct ocean DIC are co-located model temperature and salinity anomalies, as well as globally averaged globally-uniform atmospheric pCO₂ anomalies. The resulting system thus has 3 input parameters to solve for ocean carbon at each grid cell. With this method, we explore how the local relationships between upper-ocean DIC, temperature, and salinity can be used to reconstruct carbon with perfect hydrographic knowledge. Since observations are complete in this test, the presence of errors can be attributed to poorly known relationships or nonlinear dynamics within the model ensemble.
2. ~~Observations~~ Synthetic observations are taken irregularly, consistent with the distribution of Argo profiles (?). While coverage of ocean temperature and salinity observations is significantly higher than those for ocean carbon, sampling remains irregular in space and time, so real-world reconstruction methods must account for this irregularity. We take Argo data from years ~~2002 and 2015 and 2002-2015 and~~ bin the profiles onto the $1^\circ \times 1^\circ$ horizontal grid. Any grid cell that has Argo observations for at least 6 months ~~of profiles in within~~ a given year (running from January to December) is taken to be sufficiently observed, and the modelled annual average profile there is used as an observation for the reconstruction. For year 2002, this leads to a distribution of synthetic profiles that is heavily concentrated in the Northern Hemisphere and the North Atlantic in particular (Fig. 1a). By year 2015, the scheme allows for sampling of most of the global ocean outside the Southern Ocean and the Arctic Ocean (Fig. 1b). Using this distribution of available profiles, we create reconstructions using time-varying profile locations and test how profile density impacts the reconstruction by setting the distribution as constant in time.

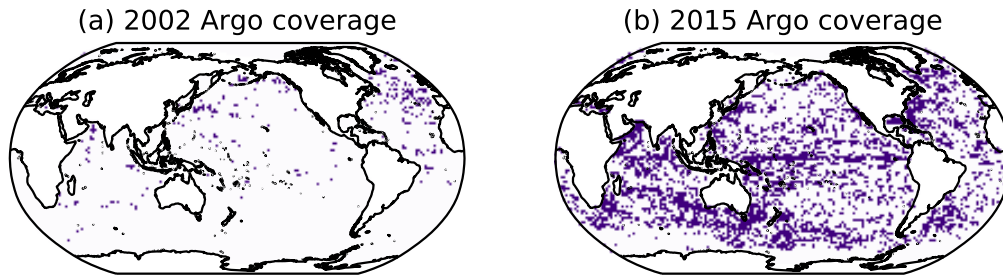


Figure 1. Argo profile distributions consistent with sampling for (a) year 2002 and (b) year 2015. Coloured regions indicate bins that have at least 6 months of observations in the given year.

As the Argo profiles have incomplete global coverage, we conduct additional sensitivity tests to see how nearby Argo observations can be used to construct near-global reconstructions. For these reconstructions, the reconstruction at any given point is made to be a linear combination of observations within a certain radius. ~~We choose radii~~ Radii of 1° , 2° , or 5° ~~are chosen~~: for a radius of 5° and current Argo coverage like that for year 2015, most regions of the ocean have observations within the search radius, allowing for nearly globally complete DIC reconstructions.

2.2 Synthetic reconstructions and tests

Evaluating how well an analysis reproduces variability requires a comparison with a field truth. As the ocean carbon field is not known to a high accuracy, ~~we conduct synthetic reconstructions in which we reconstruct other~~ synthetic reconstructions of models' ocean carbon fields are created using various distributions of modelled temperature and salinity (?). For these tests, the synthetic observations come from the models and therefore contain no errors outside of small errors possible from sampling and regridding. As the synthetic observations can be assumed to have little-negligible error, the weights w_k in (2) can be solved for through a least-squares algorithm, without having to consider observational errors (?).

The models used in the synthetic reconstructions are the same models used to calculate the ensemble covariance fields. The inclusion of a model in the ensemble can over-fit the results and produce spuriously accurate reconstructions; therefore, for a synthetic reconstruction of, for instance, UKESM1 DIC inventories, we eliminate all the UKESM1 realisations from the ensemble. ~~We then calculate the~~ The covariances and optimal weights are then calculated from the remaining out-of-sample ensemble members. Then, these optimal weights are used with the UKESM1 temperature and salinity profiles to create a DIC analysis.

To compare the reconstructions, ~~we calculate~~ the improvement in the RMSE, is calculated relative to the RMSE of the first-guess background using:

$$\varepsilon_i(A) = \frac{RMSE(B_i, T_i) - RMSE(A_i, T_i)}{RMSE(B_i, T_i)}. \quad (4)$$

Here the conventions follow those in (1), where $RMSE(B_i, T_i)$ is the RMSE between the modelled truth T and the background climatology field B (equivalent to the variance-standard deviation of T), and $RMSE(A_i, T_i)$ is the RMSE between the

modelled truth and the multivariate analysis A . Each RMSE is calculated for the same historical period used in the creation of the model ensemble, i.e. t is taken from output within the modelled years 1955 to 2014. The maximum value of 1 indicates a perfect reconstruction A , whereas values below 0 indicate that the errors are larger for the analysis than they are if the solution were to be taken as the climatological first-guess.

205 For the following analysis, synthetic reconstructions of DIC are first made over the upper 100m, and then over the upper 2000m, using both global and irregular temperature and salinity observations. Over the top 100m, the ocean carbon distribution is expected to reflect the controls of solubility and alkalinity on ocean carbon, due to the dominance of air-sea gas exchange on the carbon inventory. Over the upper 2000m, ocean circulation and ventilation are expected to play a larger role in determining the reconstructed carbon fields.

210 **3 ~~Model correlation fields between pCO₂, temperature, salinity, and~~ Reconstruction of ocean DIC within the top 100m**

Our reconstruction approach using the ensemble of 6 CMIP6 model runs is now applied to carbon inventory changes in the top 100m.

3.1 Model correlation fields between pCO₂, temperature, salinity, and DIC

215 The least-squares solution for the weights for temperature, salinity, and atmospheric CO₂ is a function of the covariances between the input variables and DIC. Thus, the structure of the covariance fields provides insight into how temperature and salinity can be used to reconstruct carbon. For simplicity, we illustrate these relationships through correlation fields taken for the entire model ensemble, which normalise the relationships using the variances of each input variable.

3.1.1 Correlations with atmospheric pCO₂

220 The correlation fields ~~of between atmospheric pCO₂ and DIC, temperature, and salinity with atmospheric pCO₂ or salinity~~ reflect the impact of emissions on the ~~upper-ocean-ocean mixed layer~~. Atmospheric pCO₂ is strongly positively correlated with ~~upper-ocean DIC, with strongest correlations (mixed layer DIC and reaches values~~ near 1 ~~)~~ in the mid-latitude ocean (Fig. 2a). This strong correlation reflects how the ocean takes up carbon under higher atmospheric pCO₂ due to gas disequilibrium. The correlation between atmospheric pCO₂ and ~~upper-ocean integrated temperature is weaker than that with DIC but still~~ positivealmost everywhere integrated temperature changes is widely positive, consistent with the long-term ocean uptake of both heat and carbon during carbon emissions (Fig. 2b). Lastly, the correlation between ~~atmospheric-atmospheric pCO₂ and upper-ocean integrated salinity is spatially heterogeneous and near-zero~~ integrated salinity changes is close to zero in most regions~~of the ocean~~, although smaller regions such as the Gulf Stream show stronger ~~positive~~ correlations.

~~Correlations between upper-ocean temperatureand DIC are positive-~~

230 **3.1.2 Correlations between DIC and temperature**

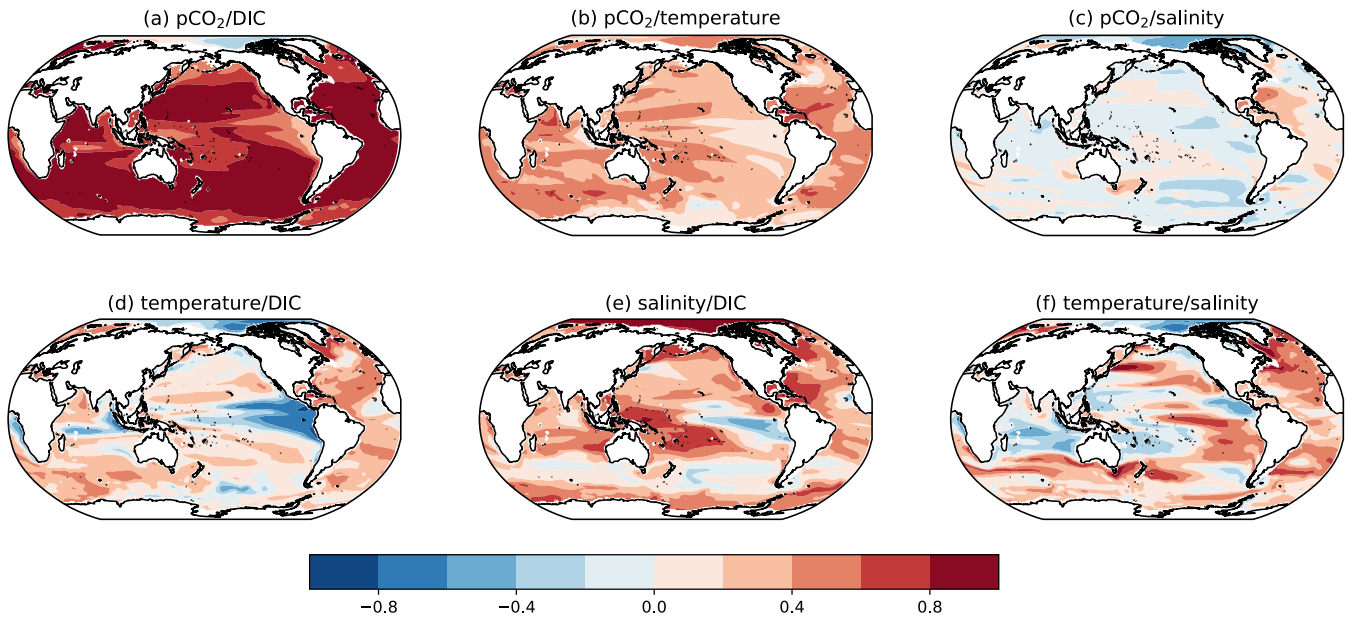


Figure 2. Correlations for the entire CMIP6 historical ensemble between changes in (a) atmospheric pCO₂ and upper-ocean (0-100m integrated) DIC, (b) atmospheric pCO₂ and upper-ocean temperature, (c) atmospheric pCO₂ and upper-ocean salinity, (d) upper-ocean temperature and DIC, (e) upper-ocean salinity and DIC, and (f) upper-ocean temperature and salinity.

Whereas atmospheric pCO₂ has little interannual variability, mixed layer temperature, salinity, and DIC will have variability on interannual and decadal timescales alongside long-term trends. The correlation fields between mixed layer DIC, temperature, and salinity reflect the combined response to external forcing from atmospheric pCO₂ and internal variability.

235 Mixed layer temperature and DIC are positively correlated in most of the ocean, with strong negative correlations in the eastern equatorial Pacific and Indian Oceans (Fig. 2d). ~~A positive correlation appears to contradict the negative correlations that would be~~ This positive correlation contradicts the negative correlation expected from the solubility control of temperature on solubility. ~~However, the residuals for temperature and DIC contain forced and climate variability responses. We can approximate the correlation as the sum of DIC.~~

240 To separate the correlation into terms relating to the external forcing and ocean variability, we can decompose temperature and DIC into terms proportional to atmospheric pCO₂ and non-pCO₂ components to separate the various responses:-

$$\rho(DIC', T') \approx \underbrace{\alpha\gamma \frac{\text{var}(pCO_2)}{\sigma_{DIC}\sigma_T}}_{\text{pCO}_2 \text{ term}} + \underbrace{\frac{\text{cov}(DIC_a, T_a)}{\sigma_{DIC}\sigma_T}}_{\text{non-pCO}_2 \text{ term}}.$$

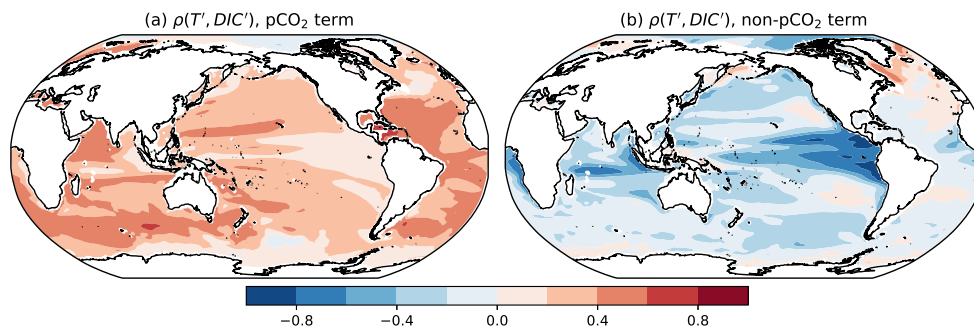


Figure 3. Breakdown of the correlation between upper-ocean DIC and temperature, $\rho(DIC', T')$, from Fig. 2d: (a) a term proportional to the variance of atmospheric pCO₂, and (b) a term consisting of the covariance of the residuals calculated after removing a linear fit against atmospheric pCO₂. For discussion of the breakdown of the correlation into terms and on the approximation, see Appendix A.

an anomaly term:

$$\underline{T'(x, y, t) = \alpha(x, y)pCO_2(t) + T_a(x, y, t)} \quad (5)$$

$$\underline{DIC'(x, y, t) = \gamma(x, y)pCO_2(t) + DIC_a(x, y, t)} \quad (6)$$

245 ~~where~~ The terms α, γ describe the respective spatial patterns of changes in T', DIC' with changes in pCO₂, ~~T_a, DIC_a are the~~
The anomalies of T', DIC' after a regression against pCO₂ is removed, are denoted as T_a and DIC_a .

With the decomposition in (5) and (6), the correlation between DIC' and T' can be approximated as the sum of a pCO₂ and a non-pCO₂ component, where σ_{DIC}, σ_T are the standard deviations of DIC' and T'.

$$\underline{\rho(DIC', T') \approx \underbrace{\alpha\gamma \frac{\text{var}(pCO_2)}{\sigma_{DIC}\sigma_T}}_{\text{pCO}_2 \text{ term}} + \underbrace{\frac{\text{cov}(DIC_a, T_a)}{\sigma_{DIC}\sigma_T}}_{\text{non-pCO}_2 \text{ term}}.} \quad (7)$$

250 This decomposition is accurate to first order. A more detailed derivation of this breakdown and a discussion of the approximation can be found in Appendix A.

The pCO₂ component of the ~~temperature/DIC correlation field~~ correlation field between temperature and DIC is positive almost everywhere, ~~which corresponds with the addition~~ (Fig. 3a). This relationship is consistent with the combined oceanic uptake of heat and carbon to the ocean system through from CO₂ emissions (Fig. 3a). Conversely, the anomaly term for
 255 the ~~temperature/DIC correlation field is negative almost everywhere~~ correlation field between temperature and DIC is broadly negative, which corresponds with the solubility argument-control for DIC (Fig. 3b). Thus the heterogeneity found in the overall correlation between DIC and temperature in Figure 2 can be understood as the sum of ~~two separate relationships that arise~~ relationships from the forced and natural variability in temperature and DIC.

3.1.3 Correlations between salinity and DIC or temperature

260 Unlike DIC, temperature, and atmospheric pCO₂, mixed-layer salinity changes have weak trends. Decomposing the correlations between salinity and temperature or DIC show little role for the pCO₂ term in (7) (Appendix A). Therefore we continue by exploring the full correlation fields with mixed layer salinity.

The correlations between upper-ocean salinity and upper-ocean DIC are positive in most regions, reflecting the alkalinity control on carbon solubility (Fig. 2e). Correlations are weak and slightly negative in the Southern Ocean and eastern equatorial Pacific, respectively. These correlations may be imprints of dynamical changes; for the Southern Ocean, the salinity in the frontal zone impacts the strength of mode and intermediate water formation (?), while for the equatorial Pacific enhanced wind-driven upwelling transports fresher, carbon-rich waters to the surface (?). Lastly, correlations between upper-ocean temperature and salinity are moderate and show strong regional variability (Fig. 2f).

Overall, the correlation fields show first-principle controls relating to the increase of ocean heat and carbon under emissions, as well as solubility and alkalinity controls of temperature and salinity anomalies on ocean DIC.

3.2 Ensemble optimal weights for pCO₂, temperature, and salinity

Translating covariance and correlation fields to optimal mapping parameters involves both the relationships between the input variables and the output variables as well as the relationships between input variables. Therefore, optimal solutions are non-trivial; for instance, regions with similar correlations between DIC and temperature and salinity (such as the equatorial Pacific) may have coefficients with different signs depending on how the system fits the linear model to the data. We thus continue by comparing the least square coefficients fit to atmospheric pCO₂, integrated ocean temperature, and integrated ocean salinity to the correlations in Section 3.1 and the first-principle controls on carbon solubility.

Increases in atmospheric pCO₂ are translated to increases in upper-ocean carbon (Fig. 4a). The weight magnitudes reach local maxima in the subtropics, in accordance with regional variability in the Revelle buffer factor, which describes the ratio between increases in DIC and increases in atmospheric pCO₂ (?) and thermocline ventilation. Temperature coefficients are negative almost everywhere (Fig. 4b), while salinity coefficients are positive everywhere (Fig. 4c). Thus, while the correlations between upper-ocean DIC, temperature, and salinity show regional variability, the ensemble ultimately reveals a consistent forced control of atmospheric CO₂, solubility control by temperature, and alkalinity control by salinity.

Thus, the least-squares solution for DIC as a combination of atmospheric pCO₂, temperature, and salinity changes is able to capture the controls of the Revelle buffer factor, solubility, and alkalinity on upper-ocean carbon. ~~While the correlations between input variables exhibit strong regional variability and structure in their cross-correlations,~~ The information from ocean temperature and salinity observations end-up-acting-in-complementary ways are complementary: regions with relatively high/low temperature coefficients correspond with low/high salinity coefficients. Thus the optimal weight solutions indicate that the information provided by these observations are consistent with our hypothesised first-order controls.




Figure 4. Ensemble coefficients for reconstructing upper-ocean (0-100m integrated) DIC with co-located observations. (a) Coefficients for annual average atmospheric CO₂ concentrations, in units mol C (ppm CO₂ m²)⁻¹, (b) coefficients for upper-ocean integrated temperature, in units mol C (°C m³)⁻¹, and (c) coefficients for upper-ocean integrated salinity, in units mol C (psu m³)⁻¹.

290 4 Reconstruction potential using co-located observations

3.1 Reconstruction potential using co-located observations

~~We continue by conducting the sensitivity tests to provide initial bounds for~~ Sensitivity tests are now conducted to estimate the ability to reconstruct ocean carbon, as well as explore how the ~~makeup of the ensemble~~ ensemble composition impacts the reconstruction. In total, 6 reconstructions were created, in which one model was removed from the ensemble, the covariance fields between ~~DIC, T, S, DIC', T', S'~~ DIC, T, S, DIC', T', S' and atmospheric pCO₂ are recalculated using the remaining models, and then ~~DIC~~ DIC' from the excluded model is reconstructed using the new covariance fields and its own ~~T and S~~ T' and S' observations (Section 2). The coefficients for atmospheric CO₂, temperature, and salinity are similar across the sensitivity setups with each model removed (Appendix B Supplementary Figures S1-S3), so we continue by comparing the RMSE improvements ε from (4) for each sensitivity test.

300 ~~We calculate the~~ The ensemble minimum, average, and maximum relative RMSE improvements ~~within the~~ are calculated for 6 reconstructions, ~~each~~ each created by eliminating ~~specific models a model~~ a model from the CMIP6 ensemble. The minimum and maximum error improvements $\min(\varepsilon)$ and $\max(\varepsilon)$ reflect the makeup of the ensemble. A negative minimum error improvement at any point indicates where one of the models in the ensemble has covariances that are substantially different than the others; therefore, including the model within the ensemble adds uncertainty to the reconstruction and pushes the reconstruction towards 305 the first-guess climatology field. A high maximum error improvement indicates that the solution weights for the sensitivity results are similar. This similarity can arise from strong physical constraints on the upper ocean carbon system but may also arise because the ensemble members are spuriously similar in their architecture or representations of climate processes.

~~We consider the~~ The RMSE improvement of the reconstruction relative to the climatological first-guess ~~as in (4)~~ is now considered for both the full upper-ocean carbon response and the detrended carbon response (Fig. 5). Over multiple decades, 310 we expect the ~~full-integrated-carbon-carbon inventory~~ response to be dominated by an upwards trend due to continuing carbon

emissions; thus, the improvement in the detrended response provides insight into how interannual and decadal variability in ocean carbon storage is reproduced. Areas where the RMSE increases relative to the climatological first guess are noted in red. For this analysis we focus on the open-ocean RMSE reductions, as the CMIP6 models have different coastlines after being regridded, leading to small areas of RMSE increases near the land/sea boundary.

315 Outside of some coastal and Arctic regions, each sensitivity test reduces the upper-ocean RMSE for upper-ocean carbon (Fig. 5a). This improvement is not due solely to the reproduction of the long-term increase in ocean carbon, as most regions show improvements in detrended ocean carbon as well; exceptions to this improvement can be found in the high latitudes and subtropical regions (Fig. 5d). The average reconstruction reduces the RMSE by between 60% to 90% on average, relative to the climatology first-guess RMSE (Fig. 5b) and reduces the detrended RMSE by 30% to 80% (Fig. 5e). These relative
320 improvements are equivalent to the method capturing over 60% of the detrended variability in most regions of the ocean. The high (close to 1) relative RMSE reductions found in the ensemble maximum statistics (Figs. 5c and e) suggest that there are models which are substantially similar to one another and thus able to reduce the RMSE of the full signal by over 80% in most regions and the RMSE of the detrended signal by over 60%.

For reconstructions of global DIC inventories using co-located temperature and salinity observations, the sensitivity experiments
325 show an average RMSE reduction of 93%. When considering detrended DIC inventories, the sensitivity experiments reduce the RMSE by 68% on average.

There are noticeable regional variations throughout all of the improvement statistics. The western low-latitude Pacific and subtropical Indian ocean show consistent local maxima in all of the improvement statistics. These regions of maximum relative RMSE reduction are characterised by their strong correlation between salinity and DIC, suggesting that the most constrained
330 responses within the CMIP6 ensemble may be related to the control of alkalinity on DIC. Conversely, the small regions that show potential degradation in the reconstructions due to the errors in the covariances are characterised by weak correlations against atmospheric $p\text{CO}_2$ changes (for the full carbon signal reconstruction, Fig. 2a-c) and weakly positive covariances between T_a and DIC_a (for the detrended carbon system, Fig. 3b).


Thus, when considering pointwise observations of temperature and salinity, alongside global average CO_2 concentrations,
335 our sensitivity experiments indicate that a substantial amount of upper-ocean carbon variability can be reconstructed. These estimates can provide an upper bound on the reconstruction potential as global coverage of temperature and salinity observations is theoretically ideal but difficult to accomplish, even with widespread autonomous observing tools. We therefore continue by exploring how irregular observations of temperature and salinity impact reconstructions of upper-ocean carbon.

4 Reconstructing carbon using irregular Argo observations

340 3.1 Reconstructing carbon using irregular Argo observations

The pointwise reconstructions of upper-ocean carbon from atmospheric $p\text{CO}_2$, temperature, and salinity reveal that the physical controls of solubility and alkalinity can explain a substantial amount of interannual variability in the upper-ocean carbon system. Although temperature and salinity profiles have a larger global coverage than ocean carbon profiles, coverage remains

Relative RMSE improvement ε , full signal



Relative RMSE improvement ε , detrended signal





Figure 5. Ensemble statistics for the relative RMSE reductions ε using co-located temperature and salinity [over the upper 100m](#): (a) ensemble minimum, (b) ensemble average, and (c) ensemble maximum. Ensemble statistics for the relative RMSE reductions, but considering only the detrended carbon signal: (d) ensemble minimum, (e) ensemble average, and (f) ensemble maximum. Red areas indicate regions where the sensitivity tests show a RMSE increase relative to the assumption of climatology.

irregular and incomplete. Thus, we continue by exploring how the irregular coverage in Argo profiles impacts carbon recon-
 345 structions by conducting further synthetic reconstructions with observations consistent with year ~~2002 and year~~ 2015 Argo
 coverage (Fig. 1b).

When considering irregular sampling, reconstructions with global coverage use background covariances to propagate infor-
 mation from Argo profile locations to the rest of the ocean. There is the potential for information from multiple Argo sites to be
 used to reconstruct ocean carbon at a given location. If the covariance fields are correct, additional information should improve
 350 the reconstruction and increase the relative error reduction; however, any errors in the covariance fields can propagate and de-
 crease the relative error reduction. Thus, the radius of influence for an observation should ideally balance the extra information
 and errors it provides in the reconstruction.

With the smaller radii of 1° and 2° , the method creates an incomplete reconstruction even under current-day Argo profile
 distributions. With a larger radius of 5° , most of the global ocean has nearby temperature and salinity profiles that can be used

Relative RMSE improvement ε , full signal



Relative RMSE improvement ε , detrended signal

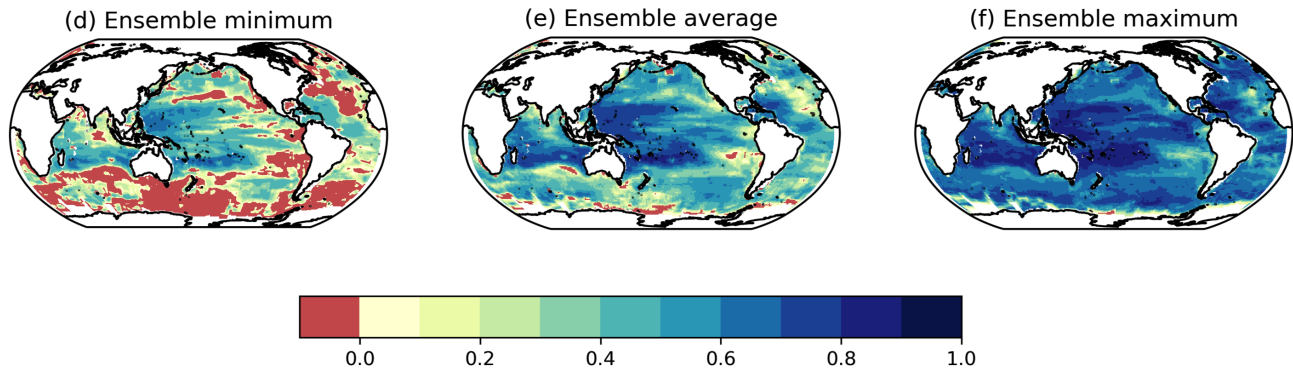


Figure 6. As for Fig. 5, but RMSE improvements over the upper 100m for the reconstruction-scheme that uses Argo 2015-type temperature and salinity observations within 5° of the reconstruction pointlocation.

355 for the reconstruction (see Fig. 6 coverage). Overall, the benefit of increased coverage outweighs the increases in reconstruction errors (Appendix C Supplementary Figure S4). Thus, in this section we focus on the 5° radius results for simplicity .

With the 5° search-cutoff radius for observations, the carbon reconstruction at each point will use between 0 and 81 ($= 1$ local observation + 80 observations within the cutoff radius) temperature and salinity inputs. Including the globally averaged atmospheric CO_2 inputs, DIC at each point can be reconstructed from at least 1 and up to 82 input variables (up to 81 temperature and salinity profiles + globally averaged atmospheric pCO_2).

360

Errors in spatial covariances result in some regions with lower RMSE improvements compared to reconstructions using co-located observations (Fig. 6a, d). Across the sensitivity tests, most regions show an-a relative RMSE reduction of at least 50%, and when considering the detrended carbon signal most regions show an RMSE reduction of at least 40%. Noticably, the ensemble minimum RMSE improvement for the detrend-detrended signal shows substantial regions where one of the sensitivity tests produces poor reconstructions. This sensitivity test is that for MPI-ESM1.2-LR, which exhibits more fine-scale structure in its annual DIC, temperature, and salinity. The errors arise because the correlation length scales for MPI-ESM1.2-LR are smaller than those for the other models, resulting in poor fits.

365

For most of the upper ocean the RMSE improvements are similar to those found in the reconstruction with co-located temperature and salinity measurements, and the regional variation in the RMSE improvements is similar regardless of whether
370 co-located or removed observations are used. Thus the additional information provided by nearby measurements counteracts the error propagation inherent in the scheme, suggesting that the use of irregular Argo profiles provides sufficient information to reproduce a substantial amount of ~~upper-ocean-carbon-variability~~ carbon variability in the mixed layer.

4 Reconstruction of ocean DIC within the top 2000m

Our reconstruction approach is now applied to carbon inventories within the upper 2000m. Our assessment of the CMIP6
375 models indicate that the relationships between carbon, temperature, and salinity may be used to reconstruct ocean carbon variability within the near-surface ocean. Anthropogenic carbon uptake and decadal variability are concentrated in the top 1000m (??). As Argo profiles can provide temperature and salinity information down to 2000m, the synthetic reconstructions are extended to cover this vertical extent. For these interior reconstructions, the carbon inventories are reconstructed as a set of layers: 100m-500m, 500m-1000m, and 1000m-2000m. For the interior layers, we conduct the same analysis as was used
380 for the top 100m reconstructions by first examining the uncertainties when reconstructing carbon using co-located temperature and salinity measurements, and then extending to reconstructions using irregular observations consistent with Argo profile locations.

4.1 Sensitivity of optimal weights to depth level

The optimal coefficients for carbon as a function of co-located temperature and salinity, along with atmospheric pCO₂ concentrations,
385 display distinct depth sensitivity (Figure 7). The optimal coefficients for atmospheric pCO₂ remain positive for most region and depth combinations. For the 100m-500m layer, pCO₂ coefficients reach their local maxima in the subtropical gyres, similarly to the optimal solution found for the upper 100m layer (Figure 4a). Below 500m, the optimal weights reflect regions of strong ventilation in the North Atlantic and Southern Oceans; below 1000m, the coefficients outside these regions are near zero.

The optimal coefficients for interior temperature changes remain negative for most regions above 1000m, with depth-dependent
390 structure (Figure 7d-e). For the 1000m-2000m layer, the optimal temperature coefficients exhibit a zonal asymmetry in the Pacific basin, where the coefficients are negative in the western Pacific and positive in the eastern Pacific. Salinity coefficients operate similarly to temperature coefficients; for the 100m-500m and 500m-1000m solutions salinity coefficients are largely positive (Figure 7g-h). Below 1000m, salinity coefficients exhibit the same zonal structure seen in the temperature coefficients. The optimal coefficients at this depth are spatially correlated, which suggests that the information provided by temperature and
395 salinities at this depth are less complementary than at shallower levels.

4.2 Carbon inventory reconstructions over 0-2000m using co-located observations

The skill of the Ensemble Optimal Interpolation reconstructions is now assessed for carbon inventory changes in the upper
2000m, following the prior assessment over the upper 100m.

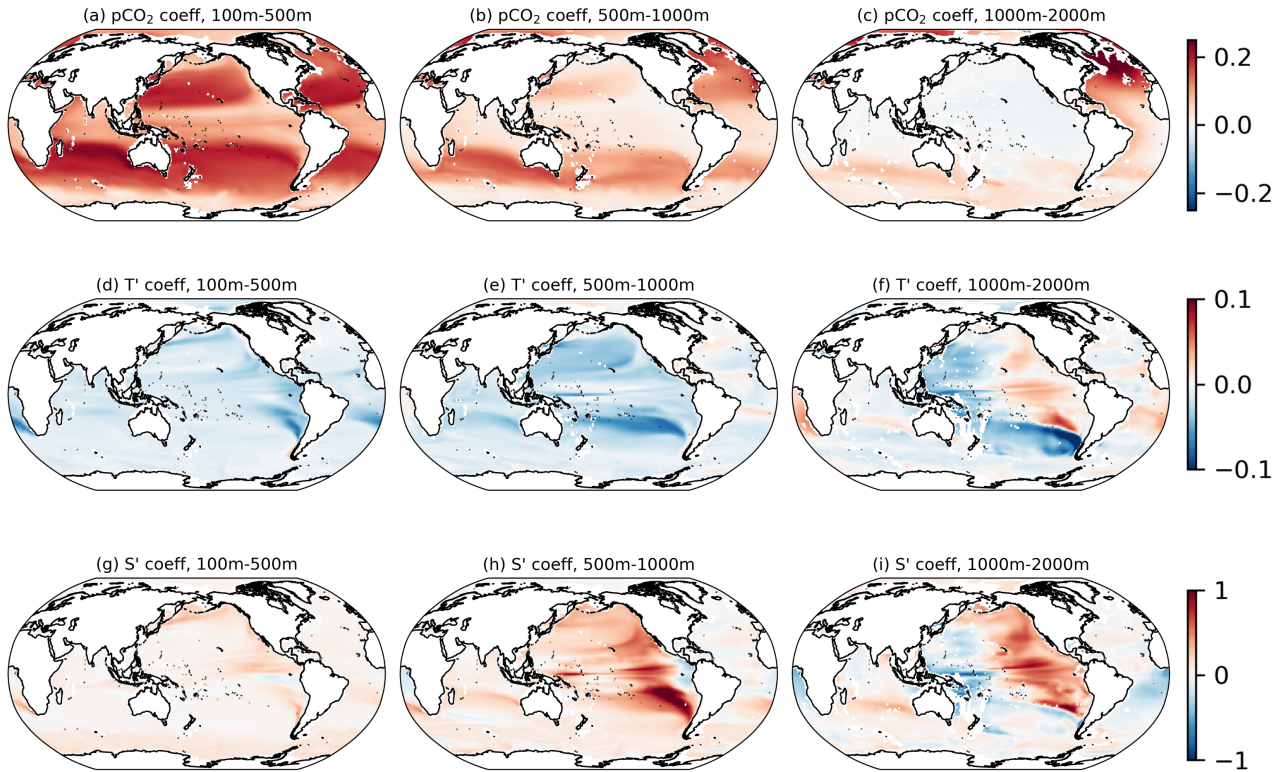


Figure 7. Optimal coefficients for interior DIC as a function of pCO₂ (in units mol C m⁻² ppm⁻¹, panels a-c), temperature (in units mol C m⁻² °C⁻¹, panels d-f), and salinity (in units mol C m⁻² psu⁻¹, panels g-i). Columns show coefficients for interior ocean layers: 100m-500m (a,d,g), 500m-1000m (b,e,h), and 1000m-2000m (c,f,i).

For both the full carbon signal and the detrended carbon signal, the reconstruction skill decreases with depth. For carbon within the top 100m of the water column, the average RMSE reduction is at least 60% in most regions, and for globally-integrated DIC reconstructions the average RMSE reduction is over 90% (Figure 5b). This RMSE reduction for the full DIC signal drops to 30%-80% for the 100m-500m layer (Figure 8b)

Below 500m, the sensitivity tests show large regions where the reconstruction increases the RMSE relative to a climatological first guess: for 500m-1000m the eastern North Pacific and equatorial Pacific show increased errors on average, whereas below 1000m most of the Pacific and Indian basins show increased errors (Figure 8c,d). The regions in which the reconstruction maintains skill with depth are the well-ventilated regions that show maximum optimal coefficients for pCO₂ (Figure 7c).

As with the reconstructions for the top 100m, reconstructions at depth also show a lower RMSE reduction for the detrended DIC signal (Figure 7f-h), and 0%-60%. For carbon between 100m-500m, the reconstruction reduces the detrended RMSE by

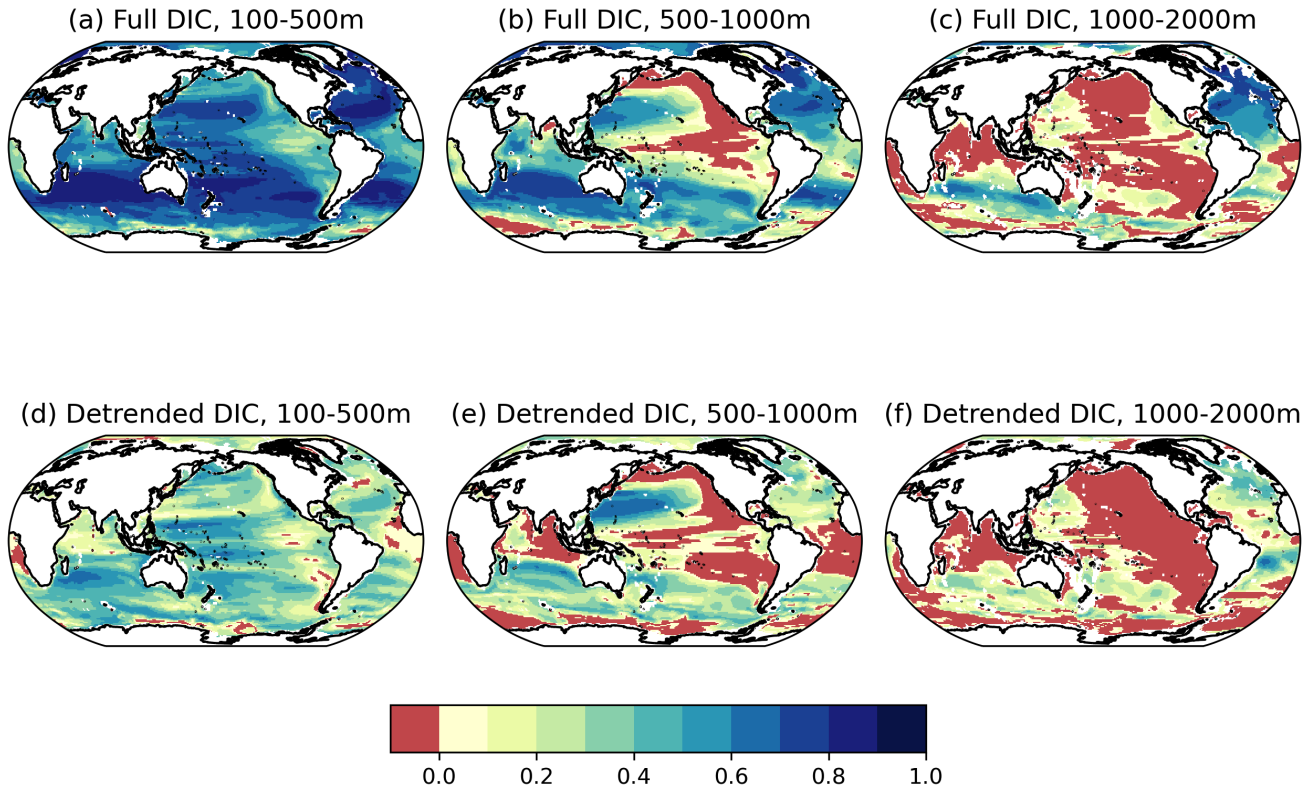


Figure 8. Average RMSE reduction ϵ , relative to a climatology, for sensitivity tests with out-of-sample models: (a-c) uses average ϵ for the full DIC' signal and (d-f) uses average ϵ for the detrended signal. The average RMSE reduction is calculated for individual depth levels: 100m-500m (a,d), 500m-1000m (b,e), and 1000m-2000m (c,f).

410 between 0-60% in most regions. Below 500m, the regions that exhibit increase errors in the detrended DIC signal are similar to those with larger errors in the full DIC signal.

While the sensitivity tests show poor reconstruction skill for large areas below 1000m, carbon changes below this depth horizon are generally small. Thus, the skill found in the 0-100m and 100-500m layers mitigates the impact of covariance errors in these deeper layers when considering full column reconstructions.

5 Potential of the method to provide global and regional carbon timeseries

415 To illustrate the potential for reconstructing upper-ocean carbon using temperature and salinity measurements, we reconstruct output from the Norwegian Earth System Model (NorESM2, ?) using the full ensemble covariance fields and temperature and salinity measurements at locations similar to Argo observation locations at year 2002 and year 2015-. This reconstruction

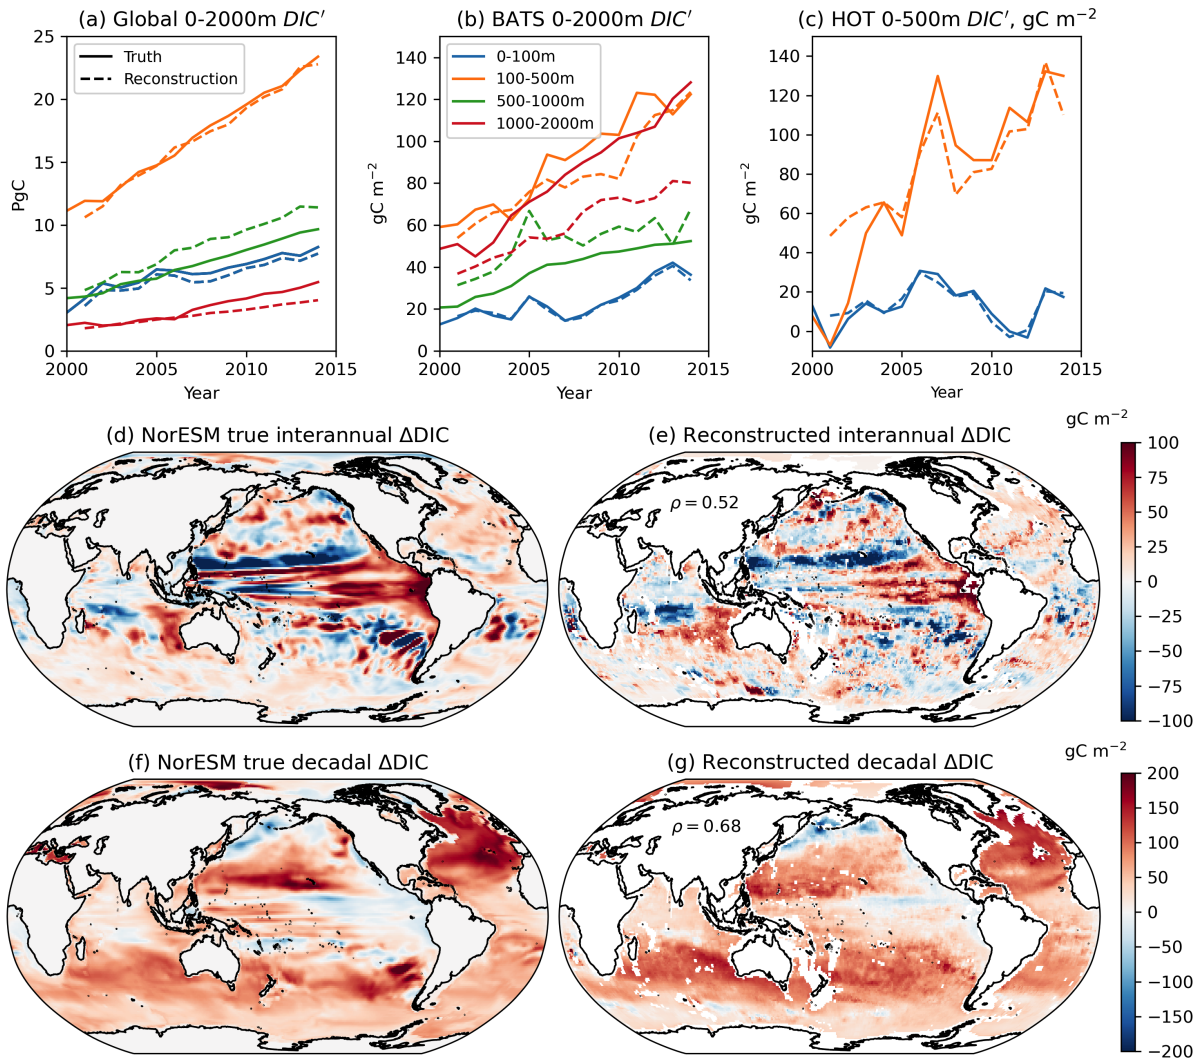


Figure 9. Reconstructions of NorESM upper-ocean 0-2000m carbon inventories, (black lines) using year 2002 Argo profile distributions Globally-integrated NorESM carbon inventories (blue dashed-solid lines) and year 2015-reconstructed carbon inventories using temporally-varying Argo profile distributions-coverage (red-dashed lines), in units Pg C. Colours indicate depths: 0-100m (blue) globally integrated carbon, in units PgC 100-500m (orange), 500-1000m (green), and 1000-2000m (red). (b) Integrated carbon-As for (a), but vertically-integrated DIC concentrations at BATS, in units gC m^{-2} and (c) Integrated carbon-As for (b), but vertically-integrated DIC concentrations at HOT to 500m, in units gC m^{-2} . Integrated carbon is calculated as a residual from the 1955-2014 average. Changes Snapshots of interannual changes in upper-ocean carbon in the NorESM0-2000m DIC, in units gC m^{-2} , between years 2001 and 2000: (d) modelled NorESM truth, and (e) reconstructed reconstruction. Interannual changes using are taken as the difference between year 2002 Argo profile distributions, 2013 and (f) reconstructed changes using year 2015 Argo profile distributions 2014 carbon fields. Correlations below Snapshot of near-decadal changes in 0-2000m DIC between periods 2001-2006 and 2009-2014, in units gC m^{-2} : (e) NorESM truth and (f) are spatial correlations calculated against the model truth (d) reconstruction.

is an additional out-of-sample reconstruction, similar to those made for the error reduction statistics in Sections 3.1 and 3.1, but uses covariance fields constructed from the entire CMIP6 model ensemble. We elect to use NorESM2 ~~was chosen~~ as an independent check as ~~its depth coordinates follow isopycnals rather than distance from the sea surface like z-level models~~ ~~it has only 3 available realisations and uses isopycnal depth coordinates~~; as all the models in the CMIP6 ensemble are z-level models, this choice allows us to ~~begin to consider some~~ consider potential errors from physical uncertainties ~~and the restrictions of model assumptions. The global upper ocean.~~

420 Within this reconstruction, the cutoff radii for our localization are depth-specific. The cutoff radii decrease with depth: 5° for 0-100m, 2° for 100-500m, 1° for 500-1000m, and only using co-located observations for 1000-2000m. These cutoff radii were chosen to limit the influence of errors in the covariance fields on the final reconstruction.

425 The global upper ocean carbon inventory is dominated by a positive, near-linear trend (Fig. 9a, black line). ~~Both Argo-style reconstructions reproduce this trend; however, due to having fewer observations, the year 2002 reconstruction (dashed blue line) reproduces less interannual variability than the year 2015 reconstruction (dashed red line).~~ across all depth levels (solid lines, Figure 9a). Most carbon changes occur within the top 500m (blue and orange lines); at these levels the reconstruction captures the long-term global behavior. Below 500m, the errors in the covariance fields result in biases in the reconstruction. Between 500-1000m, the reconstruction overestimates DIC accumulation, whereas below 1000m the reconstruction underestimates DIC accumulation. We note that, particularly below 1000m, the enhanced coverage of the Argo program over time results in a less accurate reconstruction due to the errors in the ensemble covariance fields.

430 To show how the reconstructions operate on more regional scales, we provide ~~the same reconstructions but for the~~ reconstructed modelled carbon inventory at the Bermuda Atlantic Time Series (BATS, 31°50'N 64°10'W) and the Hawaii Ocean Time series Series (HOT, 22°45'N 158°00'W) (Figure 9b,c). The multidecadal time series at BATS will be a useful validation tool for future reconstructions. With the localisation procedure, BATS has observations for a reconstruction for the top 2000m; at HOT a lack of observations restricts the analysis to the upper 500m. ~~These locations were chosen because they are the sites of multidecadal ocean DIC time series; thus, for future reconstructions the observations at these sites will can be used for validation. By providing reconstructions at these locations we can estimate the upper bound in how well the method works, ignoring any errors in the observations themselves. Modelled DIC at these two locations show varying levels of CO₂-driven increases and interannual variability to provide context as to how the method reconstructs both aspects.~~

445 At BATS, ~~the signal is again driven by~~ there is a long-term positive trend (Fig. 9b, black line). As the year 2002-style Argo distribution has no observations within 5°, the reconstruction uses atmospheric CO₂ measurements to reproduce the trend. For the reconstruction using a year 2015-style Argo distribution, the increase in observed locations results in a reconstruction that is able to reproduce the variability in the carbon inventory. ~~The reconstruction using the year 2015 observation locations for temperature and salinity inputs results in an RMSE that is 68% lower than that for the reconstruction that is based off of atmospheric pCO₂ changes alone in DIC in all layers (Figure 9b).~~ Reconstructed DIC captures the trend for carbon within the 0-100m and 100m-500m layers, but underestimates DIC in the 1000-2000m layer. DIC also exhibits interannual variability within the top 500m. While both the trend and variability are captured well in the top 100m, between 100m-500m the interannual variability is dampened.

At HOT, upper-ocean carbon exhibits more defined interannual and decadal variability (Fig. 9c, black line). The reconstruction using year 2002 Argo locations (blue dashed line) is able to reproduce the broad behaviour in ocean carbon, whereas the reconstruction using year 2015 Argo locations (red dashed line) shows an improvement in representation. DIC in the upper 100m is characterised by interannual variability, which the reconstruction captures. Between 100-500m DIC has a strong positive trend as well as interannual variability. The reconstruction captures both the trend and most of the interannual variability. By using multiple observations at year 2015 locations, the RMSE is reduced by 30% relative to the RMSE for the reconstruction using one observation. Thus the results at HOT provide further evidence that the inclusion of nearby Argo measurements provides skill in the reconstruction scheme, and emphasises the improvement that is already possible when using one nearby location for temperature and salinity observations. at this depth level.

Lastly, for a global view of the reconstruction we provide maps of the Norwegian ESM DIC change between years 2001 and 2000 (Fig. 9d) with the reconstructed difference using year 2002 Argo locations (Fig. 9e) and year 2015 Argo locations (Fig. 9f) changes and their reconstruction analogues. Reconstructed interannual carbon inventory changes, shown here to be the top 2000m DIC changes between year 2013 and year 2014, broadly match those in the model (Figure 9c-d). The truth field for DIC changes show regional variability, with increases in reconstruction captures the increase in carbon in the equatorial Pacific, the decrease in carbon in the subtropical North Pacific, and much of the western North Pacific and strong decreases in the equatorial upwelling zone. The reconstructed change using 2002 Argo locations is restricted by the number of observations: most regions show little change between the two years. The smaller-scale structure. Globally, the reconstruction has a spatial correlation of 0.45 using the relatively sparse temperature and salinity sampling procedure. For the reconstruction using 2015 Argo locations, most of DIC changes are reconstructed, and the reconstruction has a spatial correlation of 0.87 compared to the model truth pattern correlation of $\rho = 0.52$ for these interannual changes.

Longer-term carbon changes in the Argo period are marked by widespread increases, particularly in the Northern Hemisphere, and a DIC decrease in the Pacific (Figure 9e-f). The reconstruction captures the near-global increase in DIC and the decreases in DIC in the north and equatorial Pacific. The finer structures within the decadal DIC, such as the regions with the strongest DIC increases, are not fully captured by the reconstruction. The pattern correlation of 0.68 is an improvement on the interannual changes, likely as the role of atmospheric pCO₂ changes on this timescale are already more important in setting column carbon inventory changes.

6 Discussion and conclusions

In this study we present a new method for reconstructing upper-ocean carbon using observations of ocean temperatures and salinities is presented. While the ocean plays a large role in determining the partitioning of carbon in the Earth system, sparse observations inhibit a full characterisation of ocean interior carbon. Using synthetic profiles and creating mapped reconstructions of model truths, we have explored how wide-spread synthetic observations of temperature and salinity ~~from recent~~ representative of autonomous sampling programmes can provide global information about ocean carbon. Through reconstructions of CMIP6 model carbon fields, near-surface carbon can be reconstructed using synthetic Argo observations

due to consistent controls of solubility, alkalinity, and undersaturation on the carbonate system. The method retains skill for reconstructions down to 1000m. Uncertainties in the model covariance fields below 1000m reduce the skill in which irregular hydrographic measurements can be used to reconstruct carbon changes.

6.1 Near-surface carbon reconstructions

490 When considering the carbon system in the near-surface ocean, the correlation fields between atmospheric $p\text{CO}_2$, ocean temperature, ocean salinity, and ocean carbon reflect first-order controls (Fig. 2). Increases in atmospheric $p\text{CO}_2$ are correlated with increases in ocean heat and carbon due to the chemical and thermal disequilibria created by emissions. Increases in ocean salinity are broadly correlated with increases in ocean carbon due to the impact of alkalinity on solubility. Regional variations in the correlations between temperature and ocean carbon can be decomposed as a sum of a response related to added carbon
495 from emissions and a residual; the residual correlation reflects the impact of temperature on CO_2 solubility in seawater. Within all these correlation fields there are regional variations. The structure in the correlations between temperature and DIC and salinity and DIC attain their strongest values in different regions, and the cross-correlations between temperature and DIC are generally more moderate, indicating that the information provided by these observations are complementary.

The strength and structure in the covariance fields results in optimal weights for $p\text{CO}_2$, temperature, and salinity that reflect
500 global first-order controls of solubility and alkalinity (Fig. 4). For this linear model, using co-located temperature and salinity observations along with globally averaged atmospheric $p\text{CO}_2$ concentrations reduces the RMSE of the reconstructed DIC content by over 75% on average, relative to a reconstruction based purely on a time-invariant climatology (Fig. 5). In addition to replicating the long-term trend, the reconstructions are able to reproduce local interannual and decadal variability, capturing over 60% of the detrended upper-ocean carbon signal. When considering the detrended carbon inventory, the sensitivity tests
505 show an average relative RMSE reduction of 50% in most regions. The lower skill in reconstructing interannual and decadal variability is expected as the carbon system is nonlinear and controlled by other aspects such as circulation changes and ocean biology, which are not explicitly included in the variables used to calculate the covariance fields.

We have also explored whether the use of irregular temperature and salinity observations impacts the capacity to reconstruct ocean carbon. In theory, observations do not need to be co-located with the desired reconstruction, as the Ensemble Optimal
510 Interpolation scheme can include spatial covariances; however, there is the possibility of extra errors in the covariance fields from poorly represented spatial variability. In our sensitivity tests, we find that using irregular observations consistent with current Argo coverage can be combined with the CMIP6 covariance fields and provide similar levels of skill to the reconstruction. When taking the average RMSE reduction across our sensitivity tests, we find there is a reduction of over 60% in most regions when considering the full carbon signal and a reduction of over 40% in most regions when considering the detrended carbon
515 signal (Fig. 6b,d). However, ~~we find that there are~~ regions such as the Southern Ocean and the North Atlantic ~~which~~ show substantial error increases relative to the climatological first guess in some of the sensitivity experiments (Fig. 6a,ed). These increases in errors are found particularly when considering the detrended signal (i.e., the signal from interannual and decadal variability). The regions that show the largest error increases are those with high mesoscale ocean eddy activity, and are mainly in the sensitivity test that reconstructs ocean carbon within MPI-ESM1.2-LR. The horizontal resolution in MPI-ESM1.2-LR

520 model is higher than that for the other models in our ensemble, suggesting that the covariances' representation of mesoscale-scale ocean dynamics impacts the ability to translate nearby Argo observations into reconstructed carbon. Our sensitivity tests indicate that the correlation lengthscales ~~will be an important feature of~~ are an important factor for any real-world carbon reconstruction. ~~For future work, there are a variety of methods regarding both the ensemble selection and search radius selection that may be used to account for the model differences and model/truth discrepancy in spatial correlation lengthscales.~~

525 **6.2 Reconstructions of interior carbon**

The Argo programme measures temperature and salinity to 2000m, which potentially allows for the majority of ocean carbon changes to be captured by our reconstruction. However, within the interior, uncertainties in ocean circulation, ventilation pathways, and regenerated carbon pools can limit the skill of our reconstruction. We find that the optimal solutions for our interior carbon reconstructions are highly depth-dependent. Below 1000m, the use of only temperature and salinity as
530 oceanographic variables may be insufficient as the optimal coefficients become highly correlated.

Sensitivity experiments suggest that the heightened uncertainties in the CMIP6 models reduce the available skill in reconstructing carbon from only temperature and salinity observations at depths greater than 1000m. This increase in errors below 1000m arises from the ensemble construction, particularly the inclusion of the UKESM1. When the UKESM1 is excluded, the remaining models are more easily able to replicate each other's carbon fields. Before creating a real-world carbon reconstruction,
535 further investigation is necessary to determine whether the uncertainties provided by outlier models such as MPI-ESM and UKESM1 are physically-based or an artefact of model architecture or spin-up procedure.

When considering global reconstructions within the upper 2000m, we find that the top 500m is important for capturing both the trend and variability in ocean carbon. Within the top 500m the method is able to reproduce the trend and variability with skill, and synthetic tests with the NorESM indicate that patterns of interannual and decadal DIC changes can be captured.

540 **6.3 Flexibility of Ensemble Optimal Interpolation**

While we have opted to use only temperature and salinity as our ocean observations for this scheme, the ensemble optimal interpolation method is flexible and can take additional oceanographic variables such as oxygen, pH, nutrients, or chlorophyll. With increased coverage of these variables ~~thanks to from~~ campaigns such as the Southern Ocean Carbon and Climate Observations and Modeling programme (SOCCOM, ?) and Bio-Argo (?), there is the potential ~~for non-trivial improvements by~~
545 ~~including these variables.~~ to include these biologically-affected variables. For the high-latitude oceans where there are few Argo profiles, sea ice observations could lend additional information on upper-ocean carbon through the impact of sea ice on air-sea gas exchange.

With each new variable it is possible to quantify the amount of added information and test the ensemble optimal interpolation method at various stages by examining the covariance fields and uncertainties, and conducting sensitivity tests in a similar way
550 as has been done in this work. As more complex reconstruction schemes such as the machine learning algorithms of ? use these biogeochemical observations as well, including them in the Ensemble Optimal Interpolation scheme could enable a direct

comparison between linear and non-linear mapping methods and can help to quantify the merits of linear versus non-linear assumptions when reconstructing ocean interior carbon.

6.4 Usage of CMIP6 covariance fields

555 The Ensemble Optimal Interpolation scheme relies on some important assumptions. Firstly, the covariance fields assume that the processes relating ocean carbon to other variables are stationary. We have used historical CMIP6 experiments to calculate the covariances fields, so these covariances should be able to represent ocean carbon behaviour under current-day carbon emission forcing; however, under low or negative emissions the relationships between atmospheric $p\text{CO}_2$, ocean temperature, and ocean carbon changes will likely change due to hysteresis from continued heat uptake. Additionally, we have chosen to focus on annual average and depth-integrated carbon content in order to focus on the physical carbon response. It is also possible to consider seasonal or monthly variability by calculating covariance fields for each month or season, as is done for surface $p\text{CO}_2$ and heat content (??). In the current format, months or seasons that are poorly observed will have solutions that return towards climatological inventories. However, the system can be set up in such a way that lagged covariance fields propagate temperature and salinity information from observed to unobserved months as well as locations. The lack of wintertime carbon observations has created significant biases in machine learning products, as enhanced winter outgassing in polar waters is not present in the training data (?). As the response to gaps in data are sensitive to the reconstruction method, the Ensemble Optimal Interpolation reconstructions with seasonal data could provide added insight as to how information on the carbon system is best propagated to unobserved regions and seasons.

570 The use of a model ensemble rather than a single model allows for quantification of some of the uncertainties within the covariance fields. Our model ensemble has been constructed to consider both inter-model and intra-model uncertainty ~~when calculating the background error covariances~~. We have been able to explore the sensitivity of this method to the ensemble makeup ~~by reconstructing each model separately using out-of-sample covariances from the other models~~ through out-of-sample reconstructions, and assessing the minimum, mean and maximum ~~improvements~~ RMSE improvements. These statistics are model-specific; for outlier models such as the MPI-ESM for the upper 100m and the UKESM for the 1000m-2000m layer, the poor reconstructions reduce the average skill. Including outlier models should dampen the solution towards a climatology in regions where there are physical uncertainties. The ensemble makeup itself is limited only by the amount of CMIP6 model runs available, and so for future work we will include additional models with fewer realisations.

580 The improvements in the reconstructions indicate both that the models have regions with well-defined and strong correlations that lead to high improvements ~~such as the western tropical Pacific, as as~~ well as regions where the models are relatively uncertain in their relationships ~~such as the Southern Ocean when considering irregular spatial coverage~~ (Figs. 5, 6). While these sensitivity experiments and synthetic reconstructions allow for useful insight into the capacity of our reconstruction method, some level of uncertainty remains from the use of coarse-resolution CMIP6 models. Ocean models underestimate decadal variability in anthropogenic carbon (?) and CMIP6 models exhibit biases in the Revelle factor (?). More iterative Ensemble Optimal Interpolation techniques can be used to reduce the errors from the model ensemble by recalculating the covariance fields after observations have been included in the analysis (???)

~~As these synthetic tests~~

6.5 Moving from synthetic to real-world reconstructions of ocean DIC

Our tests using synthetic reconstructions have suggested that a large amount of ocean carbon variability can be reconstructed using only hydrographic observations, ~~future work will include expanding the method to account for observational errors and applying the method.~~ As such, this work provides a theoretical basis for applying linear reconstruction methods to real-world Argo profile data. ~~The reconstructions from Argo data can be cross-validated against independent time series sites and GLODAP data, and we will explore the added information from non-DIC biogeochemical observations. We will expand the method to consider how the covariances behave on various depth levels: as Argo profiles measure temperature and salinity down to 2000m depth, the vast majority of carbon changes can potentially be captured by the method, although for deeper~~ reconstructions there are few independent carbon observations for cross-validation. to create an independent estimate of the recent ocean carbon inventory. The transition from synthetic reconstructions to real-world reconstructions involves careful understanding of observational errors.

Unlike model output, which can be considered to be "perfect" outside of rounding and regridding errors, observational data have multiple sources of uncertainty, such as sensor errors, the representativeness of area- and time-averaged fields from Lagrangian observations, and errors in the background climatology fields. While there is rigorous quality control for Argo observations, the representation errors and climatology errors may contribute a substantial amount of uncertainty to the real-world reconstruction. Reconstructions of monthly salinity and temperature have used a parameterised observation error covariance fields (?). Future work towards a real-world reconstruction will include a biogeochemical extension of this parameterisation.

In conclusion, we have presented a method that draws upon first-principle relationships between temperature, salinity, and carbon to reconstruct ocean carbon changes from Argo style observations. We find that, for ~~the upper ocean~~ what can be considered the mixed layer, the CMIP6 models show consistent first principle constraints of solubility, alkalinity, and undersaturation, leading to accurate reproductions of trends and variability in the model world. When considering current-day distributions of Argo profiles, reconstructions that use multiple nearby observations are able to still reconstruct the trend and variability in carbon content, ~~and sensitivity tests using.~~ Below the mixed layer, the optimal coefficients become more spatially heterogeneous, and out-of-sample ~~temperature and salinity observations show strong spatial correlations of around 0.8 with the model truth and capture over 60% of the detrended variance~~ tests indicate lower skill in reconstructing carbon changes. However, because carbon uptake occurs at the atmosphere/ocean interface and the response within the top 500m dominates the upper-ocean response, reconstructions of the upper 2000m using Argo-type observations can recreate patterns of interannual and decadal variability. The results of our work provide a strong theoretical basis for using state-of-the-art autonomous hydrographic observations to supplement sparse interior carbon coverage. This proof-of-concept work shows that there is strong potential in using these measurements to create a new, independent estimate that can be used alongside reconstruction methods to better understand ocean carbon variability and its controls.

Code availability. The code to regrid, concatenate, and analyse CMIP6 model output is available in a Zenodo repository (DOI: 10.5281/zenodo.6669897).

Appendix A: Breakdown of covariance fields into pCO₂ and non-pCO₂ terms

It is possible to decompose the changes in ocean temperature, salinity, and DIC from their time-averaged states (T' , S' , DIC' from Section 2) into terms proportional to atmospheric pCO₂ and a residual:

$$T'(x, y, t) = \alpha(x, y) \text{pCO}'_2(t) + T_a(x, y, t) \quad (\text{A1})$$

$$625 \quad S'(x, y, t) = \beta(x, y) \text{pCO}'_2(t) + S_a(x, y, t) \quad (\text{A2})$$

$$DIC'(x, y, t) = \gamma(x, y) \text{pCO}'_2(t) + DIC_a(x, y, t), \quad (\text{A3})$$

where α, β, γ are the least-squares coefficients against changes in globally uniform atmospheric pCO₂ changes and T_a, S_a, DIC_a indicate the residuals for these decompositions.

While the decomposition into pCO₂ and non-pCO₂ terms is not inherently physically based, this breakdown may help us
630 understand how the complex structure in the correlation fields in Fig. 2 arise. Using covariance identities, the covariance between, for instance, T' and S' , is thus the sum of the covariances between the addends:

$$\text{cov}(T', S') = \text{cov}(\alpha \text{pCO}'_2(t), \beta \text{pCO}'_2(t)) + \text{cov}(\alpha \text{pCO}'_2(t), S'') + \text{cov}(T'', \beta \text{pCO}'_2(t)) + \text{cov}(T'', S''). \quad (\text{A4})$$

The first term is a function of the variance of pCO₂, scaled by the product $\alpha\beta$. If the term proportional to atmospheric pCO₂ were orthogonal to the residual term for each of these decompositions, the covariances $\text{cov}(\text{pCO}'_2, X_a) = 0$ for any variable
635 X . An ordinary least squares solution does not lead to orthogonal components, but the covariances between atmospheric pCO₂ and DIC_a, T_a , or S_a are small (shown as correlations in Fig. A1, panels d h and i) and can therefore be ignored to first order. Thus for these variables we ignore these terms as second-order components and can thus approximate the covariance as:

$$\text{cov}(T', S') \approx \alpha\beta \text{var}(\text{pCO}'_2) + \text{cov}(T_a, S_a). \quad (\text{A5})$$

With this decomposition of the covariance fields into pCO₂ and non-pCO₂ terms, we have the added benefit of approximating
640 the covariance fields between pCO₂ and one of T', S' , or DIC' through a global scaling of the coefficient terms α, β , or γ , respectively.

Scaling the decomposition by the product of the standard deviations allows us to explore how the correlation fields arise in Fig. 2. We compare the original correlation fields from Fig. 2 with their decompositions according to (A4) in Fig. A1. The approximation errors show how well the approximation in (A5) reflects the full correlation field.

645 For the correlations between T' and S' and S' and DIC' , the signal is dominated by the term associated with $\text{cov}(T_a, S_a)$ and $\text{cov}(S_a, DIC_a)$, respectively. We therefore understand these covariance and correlation fields to be set by the variability not forced by carbon emissions. Conversely, the correlation field between T' and DIC' appears to arise almost equally from the pCO₂ scaling and the $\text{cov}(T_a, DIC_a)$ terms. Thus we provide both terms in Section 3 to gain insight as to how the covariance fields reflect both the combined response to added carbon in the Earth system as well as the response to climate variability.




Figure A1. Correlation fields between non-pCO₂ variables T' , S' , and DIC' , with their terms from (A5) and approximation errors to measure the level of mismatch between (A4) and (A5).

650 **Appendix B: Coefficients for sensitivity tests**

Here we present the coefficients for reconstructing upper-ocean carbon using atmospheric CO_2 and co-located temperature and salinity measurements.

655 Optimal pCO_2 weights, in units $\text{mol C (ppm CO}_2 \text{ m}^2)^{-1}$, for solving for 0-100m integrated ocean DIC using a subset of models within the CMIP6 ensemble. The subsets are created by removing the runs of each model: (a) ACCESS, (b) CanESM, (c) CESM, (d) IPSL, (e) MPI-LR, and (f) UKESM

Optimal T weights, in units $\text{mol C (}^\circ\text{C m}^3)^{-1}$, for solving for 0-100m integrated ocean DIC using a subset of models within the CMIP6 ensemble. The subsets are created by removing the runs of each model: (a) ACCESS, (b) CanESM, (c) CESM, (d) IPSL, (e) MPI-LR, and (f) UKESM

660 Optimal S weights, in units $\text{mol C (psu m}^3)^{-1}$, for solving for 0-100m integrated ocean DIC using a subset of models within the CMIP6 ensemble. The subsets are created by removing the runs of each model: (a) ACCESS, (b) CanESM, (c) CESM, (d) IPSL, (e) MPI-LR, and (f) UKESM

Appendix B: Comparison of RMSE reductions with Argo profiles and various search radii

665 Here we present some reconstruction RMSE reductions for the Norwegian ESM model using the CMIP6 covariance fields, Argo profiles consistent with year 2015 sampling, and various search radii. We use the Norwegian ESM as it uses isopycnic depth coordinates as opposed to the z-level coordinates used in the models within the ensemble. We find that the coverage in the reconstructions increases to near-global coverage when allowing profiles within 5° to be used in the reconstruction and improves the detrended carbon reproduction; however, for regions that are reproduced with high skill using more restrictive search radii, the increase in search radius is accompanied by slight weakenings in the RMSE reductions.

670 Relative RMSE reduction for the NorESM full carbon (left column) and detrended carbon (right column) signals using year 2015 Argo distributions and varying radii of influence: (a,b) using only co-located temperature and salinity profiles, (c,d) using profiles within 1° latitude or longitude, (e,f) using profiles within 2° latitude or longitude, and (g,h) using profiles within 5° latitude or longitude. Regions of red indicate regions where the reconstruction increases the RMSE relative to a first-order climatology assumption.

675 *Author contributions.* All authors contributed towards the conceptualisation of the work. KET created the method and analysis with input from RGW, AK, and DMS. KT wrote the manuscript draft, and all authors reviewed and edited the manuscript. All authors signed off on the manuscript for submission.

Competing interests. The authors declare that they have no competing interests.

680 *Acknowledgements.* KET received support from the Leverhulme Trust via the Leverhulme Research Centre for Functional Materials Design. Authors KET, RGW, and AK received research support from the U.K. Natural Environmental Research Council (Grant number NE/T007788/1). DMS was supported by the Met Office Hadley Centre Climate Programme funded by BEIS and Defra. [The authors thank 2 anonymous referees for their constructive feedback.](#)