

Spatio-Temporal and Seasonal Modeling of the Ocean Heat Content Field.

Samuel Baugh and Karen McKinnon

October 30, 2022

Abstract

The Argo program has allowed for an unprecedented view of how global ocean heat content is changing, however, the complex properties of the ocean heat content field have made accurate quantification of the multi-year trend statistically challenging. Previous work has proposed a hierarchical Bayesian method for quantifying ocean heat content at a fixed time by modeling the spatial non-stationarity and anisotropic of the field. While this approach can infer the trend in ocean heat content over months, a more accurate picture requires a joint spatio-temporal model which takes into account both seasonality in the mean field and temporal correlation in the anomalies. In this paper, we propose a hierarchical Bayesian spatio-temporal model for ocean heat content that addresses these properties through a spatially-varying temporal range parameter field and a spatio-temporal Gaussian process to represent a seasonal structure that varies smoothly over space. This approach applied to ocean heat content observations from 2007 to 2021 is able to identify the ocean heat content trend with higher posterior confidence than would be expected just from the increase in the amount of data alone in the spatial-only approach. Due to the key role played by ocean heat content in regulating the atmospheric response to increasing greenhouse gas concentrations, these results have important implications for the estimation of transient climate sensitivity.

1 Introduction

Previous work [Baugh and McKinnon, 2022] introduced a hierarchical Bayesian Gaussian process framework for modeling the spatially non-stationary correlation structure of the ocean heat con-

tent field. This model was fit to a subset of the data such that seasonal variation in the mean-field and temporal correlations in the anomalies could be ignored, producing a posterior distribution for the year-to-year trend that takes into account sampling uncertainty, uncertainty from variability in the process, and uncertainty from parameter estimation. This paper extends the spatial-only model to the spatio-temporal setting in order to take advantage of the information presented by spatial correlation in the data as well as the additional data from incorporating all months. The treatment begins with updating the notation for representing the data in the spatio-temporal setting in Section 2. Section 3 introduces two improvements to the spatial-only model to account for temporal properties in the mean and anomaly fields. The first development in Section 3.2 is of a mean-field that models smooth seasonal variation simultaneously with spatial variation through a spatio-temporal Gaussian process prior with a “circular” correlation structure over time. Section 3.1 introduces a temporal dimension to the kernel-convolution process with spatially-varying temporal correlation length scale parameters analogous to the latitudinal and longitudinal convolutions presented in [Baugh and McKinnon, 2022].

This paper proceeds to apply the spatio-temporal framework to Argo data from all months over the years 2007-2021. Parameter fields for initializing the MCMC sampler are derived using an updated moving window approach in Section 4. In Section 5 the spatial Vecchia process used in Baugh and McKinnon [2022] as a computationally efficient approximation for the full Gaussian process is augmented to account for the spatio-temporal structure of the data. This requires a re-evaluation of the size of the conditioning sets needed to obtain a sufficiently accurate representation, which is done through a study evaluating the accuracy of the posterior distributions of the mean and trend fields for different conditioning set sizes. Finally, results from the posterior distribution are presented and discussed in Section 6, including the posterior correlation parameter fields, the spatially and seasonally varying mean-field, distributions for the ocean heat content values at each month, and the spatially-varying trend field. The resulting estimates for the trend in ocean heat content are consistent with the results in Baugh and McKinnon [2022] but with a much higher level of confidence due to the increase in data and the information provided by the temporal structure.

2 Spatio-Temporal Data

To extend the framework to data from all months, the notation must be adjusted to represent each location as a spatio-temporal point and to re-define heat content as a field varying continuously over time. Let $\mathbf{x} = (x_{\text{lat}}, x_{\text{lon}}, x_{\text{time}}) \in \mathbb{CL} \times \mathbb{R}$, where as before $\mathbb{CL} \equiv \mathbb{R} \times \mathbb{S}$ represents latitude and longitude in Euclidean and spherical coordinates respectively, and x_{time} represents time in days. Furthermore, let $z \in \mathbb{R}^+$ denote depth from the surface in meters. Define the three-dimensional temperature field at spatio-temporal location \mathbf{x} and depth z as $T(\mathbf{x}, z)$ as measured in degrees Celsius. Similarly to Baugh and McKinnon [2022], the spatio-temporal field is defined as heat content values integrated over depth as calculated using following formula:

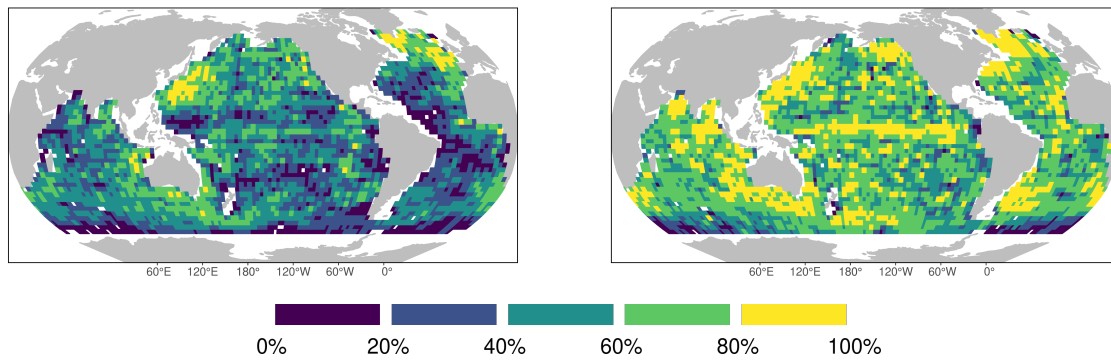
$$H(\mathbf{x}) = \int_{0\text{m}}^{2000\text{m}} \rho c_{\rho} T(\mathbf{x}, z) dz. \quad (1)$$

Here, ρ is seawater density in kg/m^3 and c_{ρ} is the specific heat of seawater in $\text{J}/(\text{kg}^{\circ} \text{C})$ [Dijkstra, 2008]. The units of $H(\mathbf{x})$ are J/m^2 . The total global ocean heat content or OHC is a spatially-integrated quantity that varies continuously over time:


$$\text{OHC}(t) = \int_{[x_{\text{lat}}, x_{\text{lon}}] \in S_{\text{mask}}} H([x_{\text{lat}}, x_{\text{lon}}, t]) d\mathbf{x} \quad (2)$$

where S_{mask} is the ocean’s surface masked following Roemmich and Gilson [2009]. Again, floats with maximum depths that are not within 100m of 2000m are excluded to avoid having to interpolate the heat content field over depth.

The data under consideration here range from January 2007 to November 2021, and after applying the mask and depth filters there are 1,111,023 total ocean heat content observations. These observations are not uniformly distributed over space, with some regions frequently seeing more floats than others, nor over time, with an increase in the number of floats deployed over time resulting in more recent years having more observations than earlier years. To examine the spatial distribution of observations and how data coverage has increased over recent years, the spatio-temporal domain was partitioned into $3^{\circ} \times 3^{\circ}$ grid-cells for each month within the time range. For each grid-cell-month pair, an indicator function was applied for whether or not it contained at least one observation; this is defined as the $3^{\circ} \times 3^{\circ}$ resolution “data coverage”. To visualize the spatial



(a) The percentage of months containing at least one observation within each $3^\circ \times 3^\circ$ grid-cell. Left panel is averaged over the years between 2007 and 2014, and the right panel is averaged over the years 2015 to 2021.



(b) The average number of $3^\circ \times 3^\circ$ grid-cells with at least one observation displayed over months.

Figure 1: Spatial distribution of data coverage at the $3^\circ \times 3^\circ$ grid-cell resolution over 2007 to 2014 and 2015 to 2021 (a), and the increase in data coverage over time (b).

distribution of the data coverage, the data was partitioned between subsets occurring between 2007 and 2014 and observations occurring between 2015 to 2021. For each of these subsets, the data coverage averaged over all months was calculated for each grid-cell. The resulting spatial distributions of data coverage are displayed in Figure 1a, with data coverage from 2014 and earlier on the left panel and 2015 and later on the right. It can be seen that in each of the time periods, data coverage is higher in certain areas than others, with the Pacific Ocean east of Japan and the North Atlantic exhibiting high data coverage over each time period. While these regions have over 80% data coverage in both time periods, grid-cells with less than 20% data coverage can be observed in most regions in the earlier time period. In the later time period, grid-cells with less than 20% data coverage occur rarely and most regions have greater than 60% data coverage. A notable exception is the far Southern Ocean, which exhibits low levels of data coverage even in the more recent time frame.

To get a further sense of the increase in data coverage over time, coverage was averaged spatially for each month in the full time period with results displayed in Figure 1b. It can be seen that data coverage increases continuously until about 2018, when it begins to stabilize at values nearing 70%. Coverage ranges from 30% in the first month of January 2007 to 65.67% to the last month of November 2021. In the dataset under consideration here, the number of unique floats with valid ocean heat content profiles ranges from 1,061 in January 2007 to 3,164 in November 2021.

3 Model Framework

In this section, the hierarchical Bayesian framework of Baugh and McKinnon [2022] is extended to model the spatio-temporal evolution of the Argo data. Let $H_{\text{obs}}(\mathbf{x})$ denote the observed spatio-temporal field of heat content values. This differs from the “true” unknown heat content field $H(\mathbf{x})$ since the former may include measurement error or small-scale variation. Specifically, let

$$H_{\text{obs}}(\mathbf{x}) \stackrel{\text{i.i.d.}}{\sim} N(H(\mathbf{x}), \sigma^2(\mathbf{x}))$$

where $\sigma^2(\mathbf{x})$ is the nugget variance. Under this general formulation, $\sigma^2(\mathbf{x})$ can be taken to be varying in both space and time, however the forthcoming construction will restrict this and the

other correlation parameters to vary only over space.

3.1 Temporal Correlation Model

Let $\theta_{\text{lat}}(\mathbf{x})$, $\theta_{\text{lon}}(\mathbf{x})$, and $\theta_{\text{time}}(\mathbf{x})$ represent spatially-varying latitudinal, longitudinal, and temporal range parameters respectively; these fields will be collectively referred to as $\boldsymbol{\theta} = (\theta_{\text{lat}}, \theta_{\text{lon}}, \theta_{\text{time}})$.

At a fixed spatio-temporal location \mathbf{x} , the kernel distance between spatio-temporal points \mathbf{x} and \mathbf{u} is defined as

$$d_{\text{st}}(\mathbf{x}, \mathbf{u}; \boldsymbol{\theta}) = \sqrt{\frac{d_{\text{euc}}(x_{\text{lat}}, u_{\text{lat}})^2}{\theta_{\text{lat}}(\mathbf{x})} + \frac{d_{\text{gc}}(x_{\text{lon}}, u_{\text{lon}})^2}{\theta_{\text{lon}}(\mathbf{x})} + \frac{d_{\text{euc}}(x_{\text{time}}, u_{\text{time}})^2}{\theta_{\text{time}}(\mathbf{x})}}.$$

Putting this distance function into the kernel convolution construction of Paciorek and Schervish [2006] yields the following covariance function between spatio-temporal points \mathbf{x} and \mathbf{y} :

$$k_{\text{st}}(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}, \phi) = \sqrt{\phi(\mathbf{x})\phi(\mathbf{y})} \int_{\mathbf{u} \in \text{CL} \times \mathbb{R}} \exp(-d_{\text{st}}(\mathbf{x}, \mathbf{u}; \boldsymbol{\theta})^2 - d_{\text{st}}(\mathbf{y}, \mathbf{u}; \boldsymbol{\theta})^2) d\mathbf{u}$$

where ϕ is the spatially-varying variance field. Since the squared distance terms d_{st} can be decomposed as the additive sum of distance terms in the three dimensions, the convolutions can be computed in each dimension separately. Specifically, convolutions in the latitudinal and temporal dimensions can be computed exactly using the formulas of Paciorek and Schervish [2006], and the convolutions in the longitudinal dimension can be computed using the Gaussian approximation method described in Baugh and McKinnon [2022].

The underlying heat content field $H(\mathbf{x})$ is modeled as a Gaussian process constructed from the kernel k :

$$H(\mathbf{x}) \sim GP(\mu(\mathbf{x}), k_{\text{st}}(\mathbf{x}, \mathbf{y}; \boldsymbol{\theta}, \phi)) \quad (3)$$

where μ will be defined in Section 3.2. Let $\boldsymbol{\rho} = \{\theta_{\text{lat}}, \theta_{\text{lon}}, \theta_{\text{time}}, \phi, \sigma\}$ denote the set of parameter fields excluding the mean-field. For each $\rho \in \boldsymbol{\rho}$, let

$$\log(\rho(\mathbf{x})) \sim GP(\mu_{\rho}, \phi_{\rho} \exp(-d_{\text{cyl}}(\mathbf{x}, \mathbf{y})/\theta_{\rho})) \quad (4)$$

where μ_{ρ} , ϕ_{ρ} , and θ_{ρ} are hyper-parameters which differ for each field. Note that d_{cyl} , defined in

Baugh and McKinnon [2022], is used as opposed to d_{st} ; since this distance metric only takes into account the spatial components of \mathbf{x} and \mathbf{y} , Equation (4) constrains the parameter fields ρ to be constant over time. These parameter fields are taken to vary only spatially primarily to avoid adding additional complexity to the model, particularly given the computational burden of the MCMC procedure with the large number of observations present here. There may be reasons to think that some of the parameters would vary seasonally, and if inference on spatio-temporally varying correlation parameters could be done in an efficient way this could be a potential extension for future work. However it is likely that such an extension would only have a minor effect on inference of the trend field.

3.2 Seasonally-Varying Mean-field

Unlike the correlation parameters, there is a stronger reason to think that the mean-field would vary seasonally, and given the stronger relationship between the mean-field and the trend field it is important for a spatio-temporal model of ocean heat content to take seasonality into account. Furthermore, the patterns of seasonality would not be expected to be constant over the domain; in particular one would expect heat content to peak in the late Northern Hemisphere summer or fall in the higher latitudes and peak in the late Northern Hemisphere winter or early spring in the Southern Hemisphere. To this aim, the spatially-varying mean-field of Baugh and McKinnon [2022] is extended to a field varying over both space and time, employing a “circular” correlation model over the temporal dimension in order to produce a field that is smooth at calendar-year endpoints.

To represent the temporal dimension, let $\text{doy}(x_{\text{time}}) \in [0, 366]$ denote the “day-of-year” for a particular time point x_{time} . Let $d_{\text{doy}}(x_{\text{time}}, y_{\text{time}}) = \min(|\text{doy}(x_{\text{time}}) - \text{doy}(y_{\text{time}})|, |365 - \text{doy}(x_{\text{time}}) + \text{doy}(y_{\text{time}})|)$ denote the “circular” distance over day-of-year. Then, analogously to d_{st} , the spatio-day-of-year distance can be defined as

$$d_{s\text{-doy}}(\mathbf{x}, \mathbf{y}; \theta_{\text{lat}}, \theta_{\text{lon}}, \theta_{\text{time}}) = \sqrt{\frac{d_{\text{euc}}(x_{\text{lat}}, y_{\text{lat}})^2}{\theta_{\text{lat}}} + \frac{d_{\text{gc}}(x_{\text{lon}}, y_{\text{lon}})^2}{\theta_{\text{lon}}} + \frac{d_{\text{doy}}(x_{\text{time}}, y_{\text{time}})^2}{\theta_{\text{time}}}}.$$

where unlike in the definition of d_{st} the spatial dependence on correlation length-scale parameters is omitted since this distance function will only be used in the context of a stationary Gaussian

process. Specifically, define the spatio-temporal mean-field as

$$\mu_{\text{st}}(\mathbf{x}) \sim \text{GP}(\mu_{\mu_{\text{st}}}, \phi_{\mu_{\text{st}}} \exp(-d_{\text{s-doy}}(\mathbf{x}, \mathbf{y}; \theta_{\mu;\text{space}}, \theta_{\mu;\text{space}}, \theta_{\mu;\text{time}})))$$

where $\theta_{\text{lat};\mu} \equiv \theta_{\text{lon};\mu} \equiv \theta_{\text{space};\mu}$ produces a process that is spatially isotropic with a separate correlation length-scale in the temporal dimension. While isotropic Gaussian processes are used for the parameter fields, the underlying model for the ocean heat content field remains anisotropic. This prior forces μ_{st} to be constant over each calendar year and to vary smoothly over the calendar-year endpoints.

Similar to Baugh and McKinnon [2022], the center of the Gaussian process for ocean heat content will be the sum of the year-constant mean-field and a spatially-varying trend field. Here, the trend-field β is taken to be spatially-varying only:

$$\beta(\mathbf{x}) \sim \text{GP}(\mu_{\beta}, \phi_{\beta} \exp(-d_{\text{cyl}}(\mathbf{x}, \mathbf{y})/\theta_{\beta})).$$

With these two components, the Gaussian process mean of Equation 3 is defined as

$$\mu(\mathbf{x}) = \mu_{\text{st}}(\mathbf{x}) + \beta(\mathbf{x})(x_{\text{time}} - x_{\text{origin}})$$

where unlike in Baugh and McKinnon [2022] the trend parameters are over units of days rather than years. Here, x_{origin} is defined as January 1, 2007.

The hyper-parameters in Equation 4 and above are estimated through obtaining a “first-guess” estimate of the parameter fields through a stationary moving-window approach analogous to the procedure of Baugh and McKinnon [2022]; the results of this procedure are described in Section 4. Also similarly to Baugh and McKinnon [2022] these fields will be represented and sampled through Gaussian process basis functions. Specifically, for each $\rho \in \boldsymbol{\rho}$, let $\mathbf{b}_{\rho} \sim N(0, I_{n_{\kappa;\text{space}}})$ denote basis values on a set of spatial knot locations denoted $\kappa_{\text{space}} \in \mathbb{R}^{n_{\kappa} \times 2}$. Define \mathbf{b}_{β} analogously, and for the mean-field μ_{st} , define $\mathbf{b}_{\mu_{\text{st}}} \sim N(0, I_{n_{\kappa_{\text{st}}}})$ where $\kappa_{\text{st}} = \kappa_{\text{space}} \times \kappa_{\text{time}}$. Here, κ_{st} is taken to consist of 253 knot points evenly spaced over the spatial mask at every 16 degrees in the longitudinal direction and every 8 degrees in the latitudinal direction, and κ_{time} consists of time points distributed bi-

monthly. For each $\rho \in \boldsymbol{\rho}$ let the following hold for each location \mathbf{x} :

$$\rho(\mathbf{x}) = \log(\mu_\rho + \phi_\rho \Sigma_{\mathbf{x}, \kappa_{\text{st}}; \theta_\rho} \Sigma_{\kappa_{\text{space}}, \kappa_{\text{space}}; \theta_\rho} \mathbf{b}_\rho). \quad (5)$$

For \mathbf{b}_β and $\mathbf{b}_{\mu_{\text{st}}}$, equations analogous to (5) hold by removing the logarithmic link function and substituting κ_{st} for μ_{st} .

To evaluate the posterior of the parameter fields, Markov Chain Monte Carlo (MCMC) with Metropolis-Hastings-within-Gibbs steps is used on \mathbf{b}_ρ for each $\rho \in \boldsymbol{\rho}$. The procedure is analogous to that described in Baugh and McKinnon [2022], to which the reader is directed for specific details.

4 Initial Configuration of Correlation Parameter Fields

To obtain the initial configuration for the sampler, and to estimate the hyper-priors for the Gaussian process priors assigned to each of the parameter fields, a moving-window approach is used to obtain estimated parameter fields to which Gaussian processes can be fit using maximum likelihood estimation. Specifically, for each grid-point of a $4^\circ \times 4^\circ$ latitude-longitude partition of the domain, parameters of a stationary Gaussian process were estimated from data within a $20^\circ \times 20^\circ$ window. This involves estimating parameters $\theta_{\text{lat}}, \theta_{\text{lon}}, \theta_{\text{time}}, \phi, \sigma, \beta$, and μ_{doy} from the data for each window, where μ_{doy} is an alteration of μ_{st} that is spatially stationary but contains the same structure of seasonal variation. The values of μ_{doy} at the origin, defined as January 1, 2007, is used to obtain the spatial hyper-parameters for μ_{st} . The temporal hyper-parameter for μ_{st} is taken to be fixed at $\theta_{\mu; \text{season}} \equiv 60^2$, which corresponds to a 0.05 effective correlation of 100 days, in order to achieve realistic smoothness in the estimated spatial pattern. This procedure yields a spatially-varying field of stationary parameter estimates for each of the parameters. To estimate the hyper-parameters $\{\mu_\rho, \phi_\rho, \theta_\rho\}_{\rho \in \boldsymbol{\rho}}$, Gaussian processes were fit to each of the respective gridded fields of log-transformed parameter estimates. For the mean and trend fields, spatial range parameters were estimated under the constraint $\theta_{\mu; \text{space}} \equiv \theta_\beta$ in order to ensure that the two fields that compose μ have the same spatial range. Furthermore, μ_β was constrained to be zero in order to avoid biasing the value of the overall trend. The estimated hyper-parameters obtained are displayed in Table 1.

| | Effective Range | $q_{.25}$ | $q_{.5}$ | $q_{.75}$ | Units |
|------------------------|-----------------|-----------|----------|-----------|------------------------|
| γ_{lat} | 27.09° | 2.82 | 3.61 | 4.61 | Degrees |
| γ_{lon} | 29.09° | 4.82 | 7.74 | 12.44 | Degrees |
| γ_{time} | 37.60° | 119.48 | 173.17 | 250.97 | Days |
| σ^2/ϕ | 24.39° | 0.04 | 0.07 | 0.16 | Unitless |
| $\sqrt{\phi}$ | 40.16° | 3.52 | 4.89 | 6.80 | GJ/m^2 |
| μ_{st} | 44.15° | 34.20 | 45.27 | 56.34 | GJ/m^2 |
| β | 44.15° | -0.10 | 0.00 | 0.10 | $(GJ/m^2)/\text{year}$ |

Table 1: Hyperparameters obtained from maximum likelihood estimation on fields of stationary parameters estimated from a $20^\circ \times 20^\circ$ moving window procedure. For each field the 5th, 50th, and 95th quantiles of the respective prior distributions are displayed. The range hyperparameter and the correlation length scale fields γ_{lat} , γ_{lon} , and γ_{time} are reported in terms of effective range, defined as the distance at which the correlation is 0.05.

The estimated hyper-parameters displayed in Table 1 were used along with standard kriging formulas to obtain the basis values at the spatial knot points κ_{space} . In order to visualize these configurations, these basis values were then kriged again onto a $1^\circ \times 1^\circ$ latitude-longitude grid on the domain, with results displayed in Figure 2. Here it can be seen longitudinal correlations are larger than latitudinal correlations in most areas, which is to be expected as the climate processes which influence ocean heat content vary more considerably over latitude than over longitude. Longitudinal, and to a lesser degree latitudinal, correlation lengths are noticeably larger around the Equatorial regions, particularly in the Pacific. The process standard deviation field $\sqrt{\phi}$ displays higher values around the major currents, in particular the Gulf Stream of the North Atlantic, the Kuroshio Current off the coast of Japan, the South Atlantic current off of Argentina, and the Arctic circumpolar current in the Southern Ocean. The nugget standard deviation also exhibits higher values in the area of these currents as well as noticeably high values in the South Pacific. The temporal effective correlation lengths are generally on the scale of several months. There is a notable area of temporal correlation lengths of less than a month in the western equatorial Pacific, and there are regions in the Northern Pacific and the Southern Pacific where temporal correlation lengths are more than a year.

This spatio-temporal initial configuration can be compared to the spatial-only initial configuration displayed in Baugh and McKinnon [2022]. The pattern of the process variance having higher values around major currents is consistent between the spatial-only and spatio-temporal models. This indicates that the variance in these regions persists when accounting for temporal correlations and

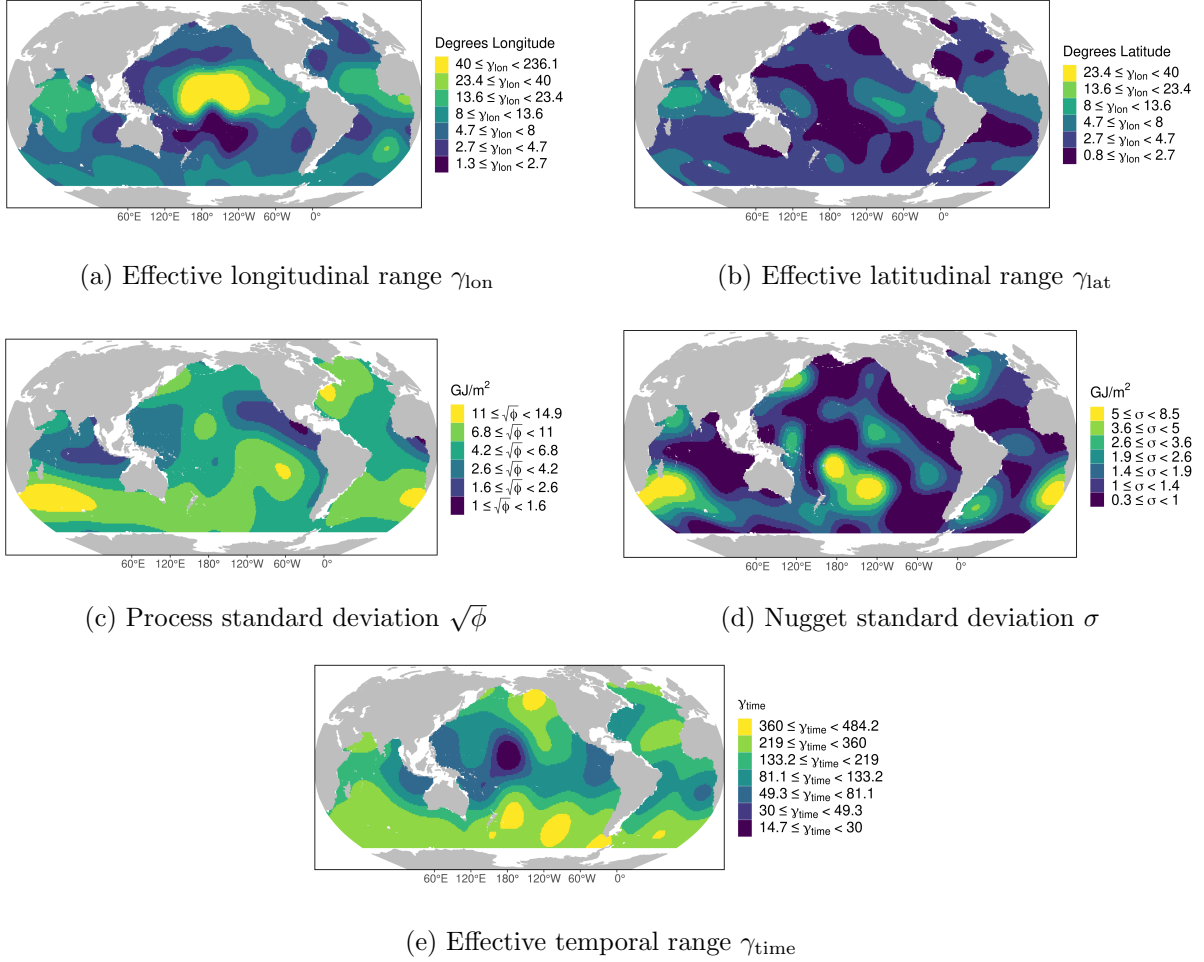


Figure 2: Initial parameter fields obtained from fitting Gaussian processes to estimated parameter fields obtained from a moving-window approach.

seasonal patterns in the larger dataset. While both models have higher longitudinal and latitudinal correlation length scales near the equator, the pattern is less robust in the spatio-temporal case, with the central equatorial Pacific having lower correlation scales than in the spatial-only initial configuration. Longitudinal correlations in the equatorial regions are even less prominent in the posterior distribution to be seen in Figure 3. This could perhaps be caused by climate phenomena with seasonal features that are not present in the January-only spatial treatment, but are present and modeled in the spatio-temporal context. An example would be ENSO, which causes spatially-coherent anomalies in the equatorial Pacific but is known to peak in the Northern Hemisphere winter [Glantz et al., 2001]. As such, the estimated correlation patterns in the spatial-only treatment may not be present in other months. Furthermore, such seasonal phenomenon would be modeled more precisely with the seasonally and spatially varying mean-field here, which could account for some of the correlation in the anomalies and lead to smaller ranges. Another discrepancy between the spatial-only and spatio-temporal initial conditions can be seen in the process and nugget variance fields, which exhibit high values in the south-eastern and south-western Pacific regions in the spatio-temporal configuration. These regions are unremarkable in the spatial-only configuration. This discrepancy could be caused by the same dynamics as the differences in the correlation length fields, as it is known that the variance and range parameters are closely related in Gaussian process likelihoods [Stein, 1999]. In other words, it would not be surprising that a change in the properties of the data when adding additional months would manifest in all of the correlation-length and variance fields. Precise identification of the climate feature causing this discrepancy may be of independent interest, and is left here as a suggestion for future research.

5 Selection of Vecchia Process Conditioning Set Size

Similar to Baugh and McKinnon [2022], a Vecchia process approximation to the full Gaussian process [Guinness and Katzfuss, 2018, Katzfuss et al., 2020a,b, Katzfuss and Guinness, 2021] will be used to render Bayesian inference tractable on the 1,111,023 observations in the dataset. To achieve a sufficiently accurate Vecchia process, the parameter m , which controls the size of the conditioning sets, needs to be re-selected within the new spatio-temporal context. In Baugh and McKinnon [2022] a value of $m = 50$ was justified by fitting Gaussian processes with the full Cholesky

likelihood to each of eleven regions partitioning the domain and then comparing the resulting integrated values to those obtained from fitting Vecchia processes with varying values of m to the same regions. With the spatio-temporal dataset, however, the full Gaussian process likelihood is too computationally intensive to fit on data restricted to the regions defined in Baugh and McKinnon [2022]. Further restricting the regions would lead to areas that would be insufficiently large for estimating the non-stationary parameter fields. Instead, to maintain the full correlation structure of the global model, the study presented here will focus on evaluating the posterior distributions for the mean and trend parameters conditioned on the values of the fully non-stationary correlation structure. This has the downside of assuming a fixed correlation structure rather than re-estimating the parameters, however, given the large amount of time required to fit the fully non-stationary posterior distribution, re-estimating the parameter fields for different values of m would not be practical.

| | μ_{post} | $\sigma_{\mu;\text{post}}$ | β_{post} | $\sigma_{\beta;\text{post}}$ |
|----------------|-----------------------|----------------------------|-----------------------|------------------------------|
| 5 | 3.12×10^{-4} | 0.049 | 8.30×10^{-3} | 0.046 |
| 10 | 1.84×10^{-4} | 0.027 | 4.53×10^{-3} | 0.025 |
| 20 | 1.33×10^{-4} | 0.014 | 6.95×10^{-3} | 0.013 |
| 30 | 8.67×10^{-5} | 9.76×10^{-3} | 5.34×10^{-3} | 9.41×10^{-3} |
| 40 | 7.11×10^{-5} | 7.46×10^{-3} | 4.36×10^{-3} | 7.31×10^{-3} |
| 50 | 6.85×10^{-5} | 5.93×10^{-3} | 5.62×10^{-3} | 5.89×10^{-3} |
| 60 | 4.61×10^{-5} | 4.95×10^{-3} | 3.66×10^{-3} | 4.87×10^{-3} |
| 70 | 3.58×10^{-5} | 4.15×10^{-3} | 3.10×10^{-3} | 4.10×10^{-3} |
| 80 | 3.05×10^{-5} | 3.62×10^{-3} | 2.92×10^{-3} | 3.55×10^{-3} |
| 90 | 2.68×10^{-5} | 3.20×10^{-3} | 3.14×10^{-3} | 3.09×10^{-3} |
| 100 | 2.51×10^{-5} | 2.83×10^{-3} | 2.73×10^{-3} | 2.77×10^{-3} |
| 200 | 2.40×10^{-6} | 1.17×10^{-3} | 6.60×10^{-4} | 1.12×10^{-3} |
| Reference: 500 | 0 | 0 | 0 | 0 |

Table 2: Fractional errors of posterior means and standard deviations for the mean and trend fields evaluated using Vecchia processes with increasing values of m . Errors are calculated against reference level of $m = 500$.

Define the posterior mean and standard deviation of the integrated mean-field conditioned on the set of correlation parameter fields $\boldsymbol{\rho}$ as

$$\left[\int_{x_{\text{lat}}, x_{\text{lon}}} \mu_{\text{st}}(x_{\text{lat}}, x_{\text{lon}}, x_{\text{origin}}) dx_{\text{lat}} dx_{\text{lon}} \middle| H_{\text{obs}}, \boldsymbol{\rho} \right] \equiv N(\mu_{\text{post}}, \sigma_{\mu;\text{post}})$$

where x_{origin} denotes January 1, 2007. For the trend field, analogously define

$$\left[\int_{x_{\text{lat}}, x_{\text{lon}}} \beta(x_{\text{lat}}, x_{\text{lon}}) dx_{\text{lat}} dx_{\text{lon}} \middle| H_{\text{obs}}, \boldsymbol{\rho} \right] \equiv N(\beta_{\text{post}}, \sigma_{\beta;\text{post}})$$

In order to select m before beginning the MCMC sampling procedure, $\boldsymbol{\rho}$ is taken to be the initial condition parameter fields obtained in Section 4.

The posterior distribution parameters defined above will be denoted with super-script (m) when calculated using the corresponding Vecchia process. Values of $\mu_{\text{post}}^{(m)}$, $\sigma_{\mu;\text{post}}^{(m)}$, $\beta_{\text{post}}^{(m)}$ and $\sigma_{\beta;\text{post}}^{(m)}$ were computed for $m \in \{5, 10, 20, \dots, 100, 200\}$ and $m = 500$. Since it is not computationally possible to evaluate the full Cholesky decomposition with over one million observations, $m = 500$ was the highest value of m evaluated; with this value, it took over eight hours with 30 cores on a cluster with 512GB of memory to evaluate the posterior distributions. For each value of m , fractional errors for the mean and slope fields were calculated using $m = 500$ as the reference level, yielding $\frac{|\mu_{\text{post}}^{(m)} - \mu_{\text{post}}^{(500)}|}{|\mu_{\text{post}}^{(500)}|}$, $\frac{|\sigma_{\mu;\text{post}}^{(m)} - \sigma_{\mu;\text{post}}^{(500)}|}{|\sigma_{\mu;\text{post}}^{(500)}|}$, $\frac{|\beta_{\text{post}}^{(m)} - \beta_{\text{post}}^{(500)}|}{|\beta_{\text{post}}^{(500)}|}$, and $\frac{|\sigma_{\beta;\text{post}}^{(m)} - \sigma_{\beta;\text{post}}^{(500)}|}{|\sigma_{\beta;\text{post}}^{(500)}|}$.

The results of this study are displayed in Table 2. It can be seen that the errors for the posterior standard deviation of both the mean and trend fields are higher than the errors for the posterior means. This is unsurprising as the values of the mean and trend fields are fairly constrained by the amount of data available, whereas the uncertainty values are more challenging to estimate accurately. For all four of the evaluated quantities, errors generally descend as the values of m increase. However, after about $m = 50$ the fractional errors are all less than 7.5×10^{-3} , with higher values of m seeing decreasing marginal improvements. As such, $m = 50$ is used for the MCMC results presented in the following section.

6 Posterior Distribution Results

To obtain the posterior distribution of the correlation parameters using MCMC, the sampler was run for increasing lengths of 1,000 iterations until convergence was achieved as indicated by the Heidelberger-Welch test applied to posterior densities of the parameter fields [Heidelberger and Welch, 1981]. This procedure required 3,000 iterations, of which the first 2,500 iterations were removed as a burn-in period. These 3,000 iterations took 528 hours (around 22 days) running on

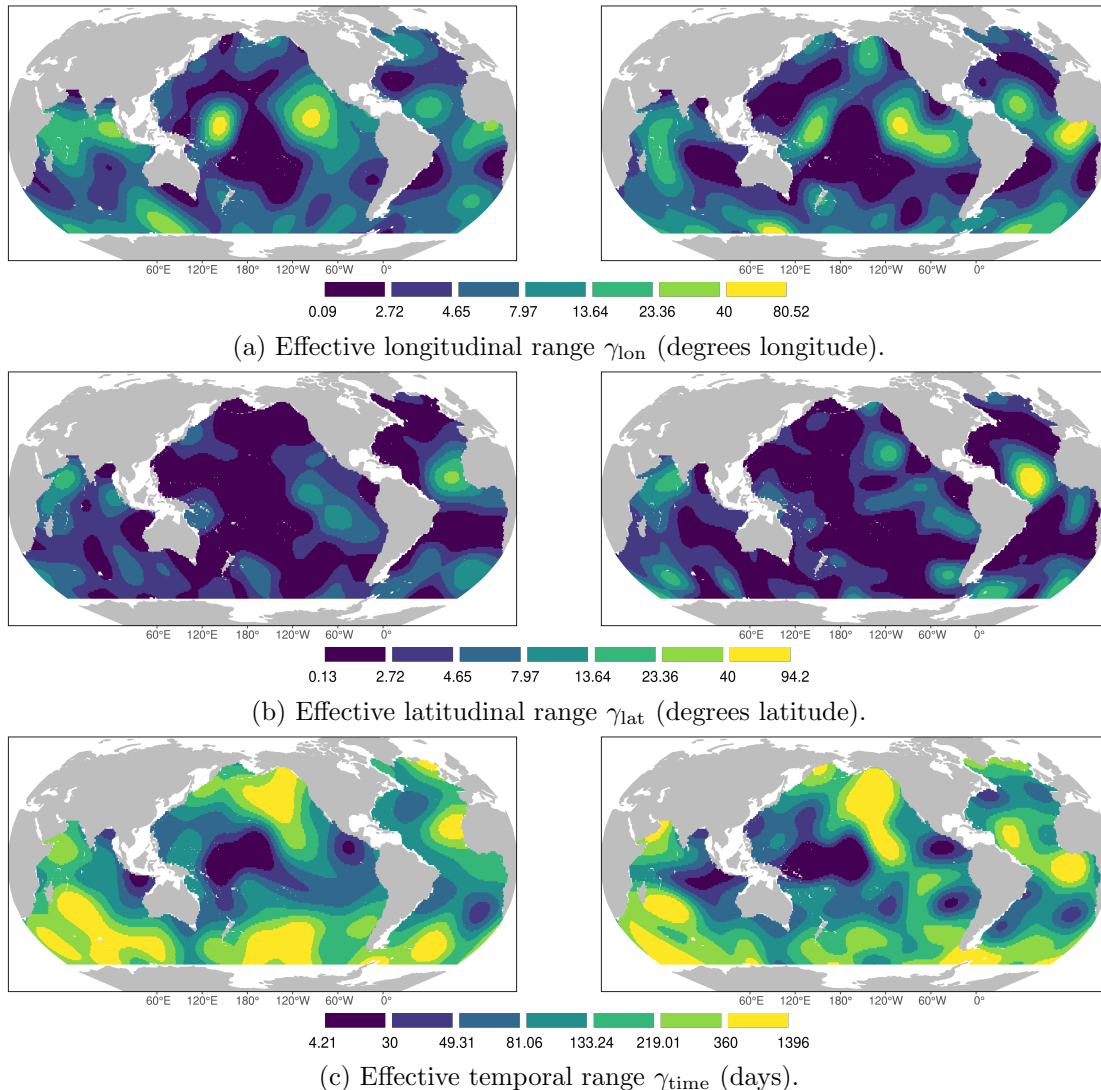


Figure 3: Samples corresponding to the first (left) and ninety-ninth (right) posterior quantiles for the effective longitudinal, latitudinal, and temporal ranges.

30 cores of a compute server with 528GB of memory. This section will proceed to present the results from the posterior distribution, including the covariance parameter fields, the spatially and seasonally varying mean-field, the ocean heat content values by month, and the spatially-varying trend field.

6.1 Posterior Parameter Fields

To visualize the posterior distribution, configurations corresponding to the first and ninety-ninth quantiles of the average parameter values were identified for each parameter field. The resulting

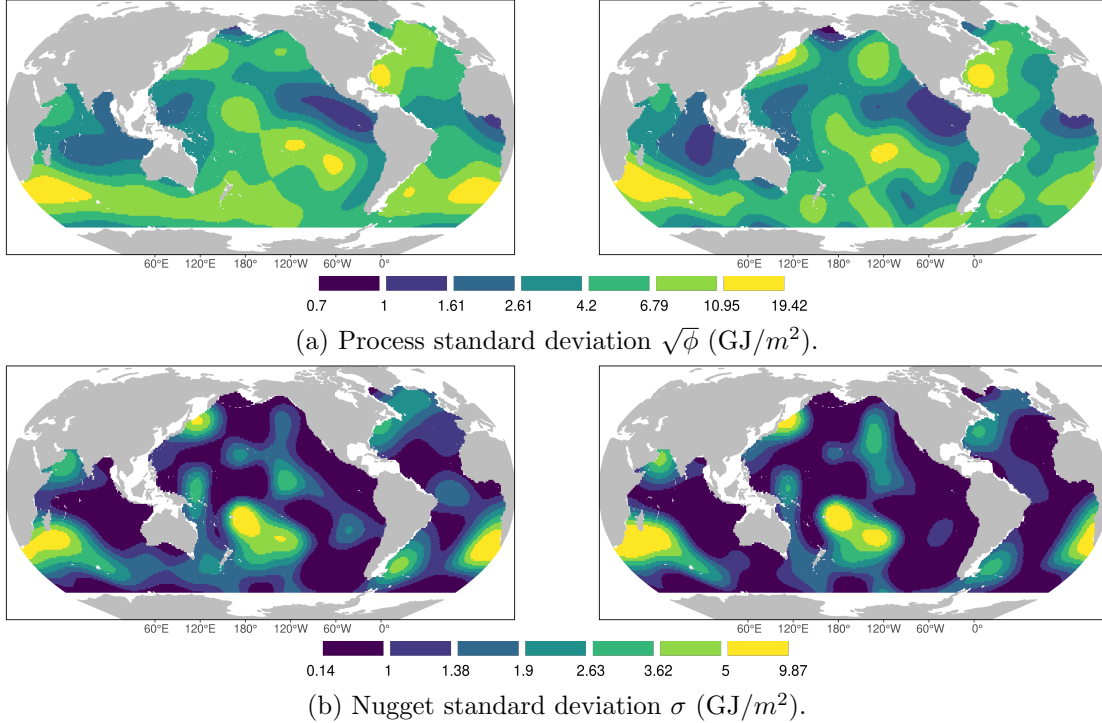


Figure 4: Samples corresponding to the first (left) and ninety-ninth (right) posterior quantiles for the process standard deviation and nugget standard deviation.

interpolated fields for longitudinal, latitudinal, and temporal ranges are displayed in Figure 3. For longitudinal ranges, it can be seen that the behavior of the high-range region in the equatorial Pacific has decreased in the center and moved both east and west in comparison to the initial configuration of Figure 2a. Variability in the posterior distribution for the longitudinal range field can be seen in this equatorial Pacific region, as well as in the equatorial Atlantic and in the Southern Ocean. For latitudinal ranges, Figure 3b exhibits higher values in the equatorial Atlantic when compared to the initial configuration, as well as noticeable variability. Variability in latitudinal ranges can also be observed in the eastern equatorial Pacific and the Southern Ocean. For temporal ranges, values in the posterior appear to be generally larger than those in the initial condition, with variability particularly high in the Southern Ocean.

The posterior distributions of process variance and nugget variance are displayed in Figure 4. For the process variance, Figure 4a shows general patterns similar to the initial configuration. Larger posterior variability is particularly notable around the location of the highly variable Antarctic Circumpolar current south-west of Africa. Overall, the process variance and nugget variance displayed

in Figure 4 appear to exhibit smaller levels of posterior variability than the other parameters. This is indicative of the fact that the variance parameters are fairly constrained by the amount of the data present, although larger posterior uncertainty is in general larger in the highly variable current regions. It should be noted that a common feature between all of the posterior distributions is the higher variability in the Southern ocean. This is not surprising when considering the spatial distribution of floats in Figure 1, where despite the large increase in float coverage over time, the data coverage rates in the Southern ocean remain low in later years.

6.2 Posterior Mean, Trend, and OHC Values

After obtaining the posterior distributions for the correlation parameter fields using MCMC, the posterior distribution of the spatio-temporal mean and trend fields can be evaluated directly. Figure 5 shows the mean-field for the MAP configuration evaluated at September 15. To get a sense of the spatial variation of seasonality, Figure 5(b) displays the difference between the September 15 mean-field minus the mean-field evaluated at March 15. As the ocean absorbs heat slowly throughout local summer, the warmest ocean heat content values would be expected to occur around the month of September in the Northern Hemisphere. Conversely, March would be expected to see the warmest ocean heat content values in the Southern Hemisphere. This intuition is confirmed in Figure 5(b), where it can be seen that September is generally warmer than March in the Northern Hemisphere and is generally colder than March in the Southern Hemisphere. Seasonal differences are most pronounced around currents, and in particular the Kuroshio current east of Japan and the North Atlantic current east of the United States.

Figure 6 displays the posterior distribution for globally-integrated ocean heat content evaluated at each month of the time range with 95% credible intervals. The distinction between confidence intervals and credible intervals are ignored here, as they are often difficult to distinguish in the full spatio-temporal setting. It can be seen that the width of the credible intervals for the ocean heat content values decrease over time as the number of floats increases. Also displayed is the posterior mean of the globally-integrated mean-field. The seasonal structure of ocean heat content is clearly visible in this figure, with the warmest values occurring in late March. This is due to the fact that March sees the warmest ocean heat content values in the Southern Hemisphere, which contains a

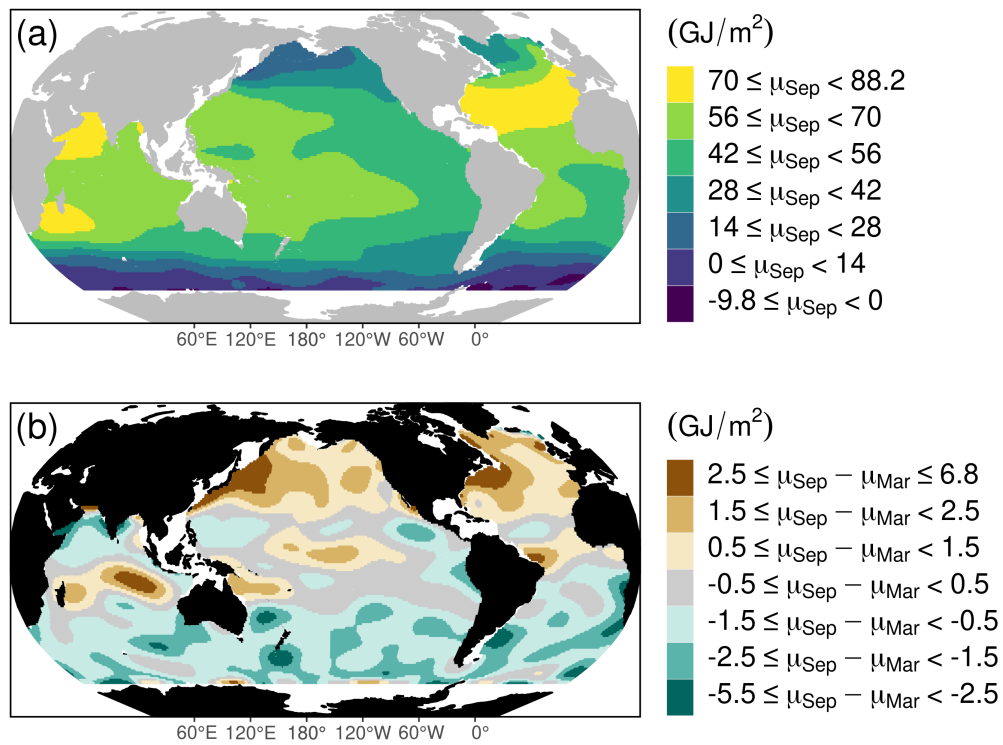


Figure 5: Mean-fields for September (a) and the difference between September and March (b) showing spatially-varying seasonality.




Figure 6: OHC by month with credible intervals. The 95% credible interval for the trend is $(12.59, 13.39) \times 10^{21}$ J/year.

larger percentage of the global ocean than the Northern Hemisphere.

It can be seen here that the massive increase in data produces a notably narrower credible interval for the trend than in Baugh and McKinnon [2022]. In particular, here the 95% credible interval for the globally integrated trend is $(12.59, 13.39)$ ZJ/year, whereas in the results of Baugh and McKinnon [2022] the interval was $(4.66, 16.03)$ ZJ/year. This corresponds to reduction in the posterior standard error by a factor of 15. This is larger than would be expected from the increase in the amount of data alone, as if the $1/\sqrt{n}$ convergence rate of the variance of a standard normal distribution is assumed, increasing the number of data points from 47,513 to 1,147,710 would suggest a reduction in the standard error by a factor of 5. This indicates that adding the spatial components to the model is adding additional confidence to the multi-year trend than would be expected from the amount of data alone. This can be attributed to three sources. The first and possibly the largest effect is that, with the addition of the temporal correlation structure, values at locations that may not have any nearby observations in a particular month now can use information from anomalies in nearby months. This will reduce the interpolation uncertainty over the single-month version whenever the effective temporal range is greater than a month, which as seen in Figure 3c is true for most of the ocean heat content domain. The second effect is that the addition of floats has increased the spatial coverage of the observations as can be seen in Figure 1. This increase in coverage will lead to a reduction of interpolation uncertainties in more recent years. The third effect is the inclusion of seasonal variation in the mean-field, which will reduce the values of the anomalies by modeling the albeit small amount of within-month seasonal variation.

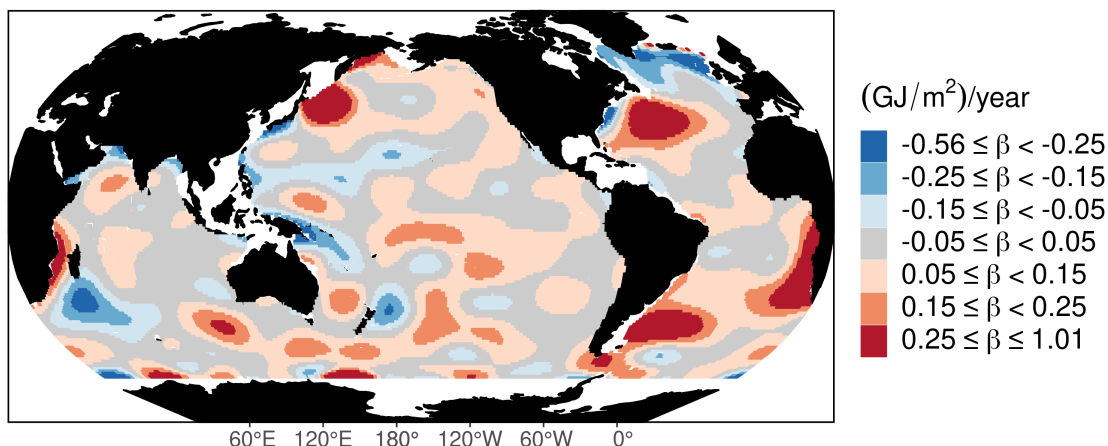


Figure 7: Posterior mean for the trend field β in units of $(\text{GJ}/\text{m}^2)/\text{year}$.

The posterior mean of the spatially-varying trend field is displayed in Figure 7. Since the posterior means are more than three standard deviations away from zero in 84% of the grid-cells, the posterior probability of a positive trend visualization is omitted in lieu of the map of posterior standard deviations. From Figure 8 it can be seen that the areas of the trend field with the highest uncertainty generally correspond to the high-variance regions displayed in Figure 4.

When comparing these results to Baugh and McKinnon [2022], some areas such as the Atlantic Ocean east of North America and south-east of South America have warming trends in both the January-only fit and in the all-years fit. The clear warming and cooling trends in the equatorial Pacific in the January-only fit are not as prominent in the all-years fit. This likely is due to the fact that ENSO's somewhat seasonal features has given it an outsized influence on the equatorial Pacific region in the January-only results from Baugh and McKinnon [2022].

7 Conclusion

The results in this chapter has shown that extending the January-only model from Baugh and McKinnon [2022] to all months and to additional years allows for ocean heat content to be estimated with a much higher degree of accuracy. In particular, the incorporation of temporal structure in the covariance structure of the anomalies and incorporating seasonality in the mean-field generates results with higher confidence than would be expected just from the increase in the amount of

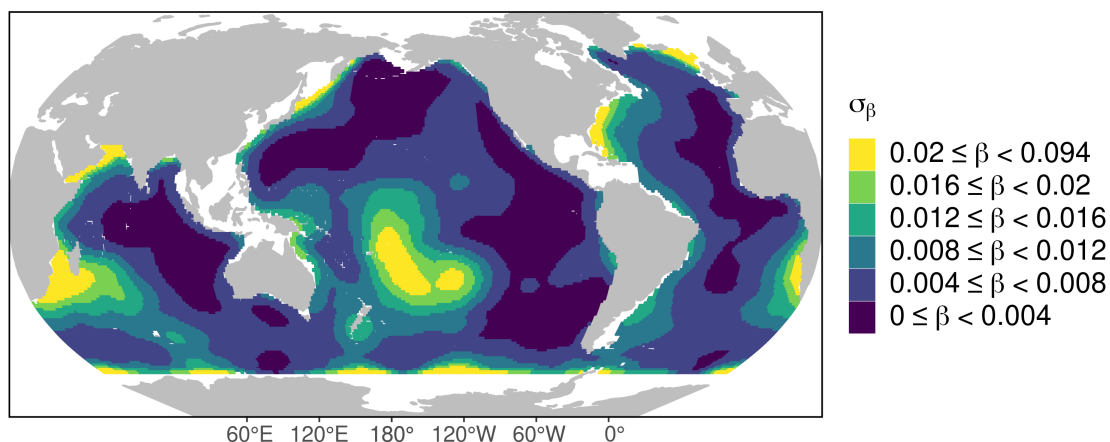


Figure 8: Posterior standard deviations of the trend field β in units of $(\text{GJ}/\text{m}^2)/\text{year}$.

data alone. Another notable change when extending the model to the spatio-temporal domain is the change in the posterior distributions for the latitudinal and longitudinal ranges as well as in the marginal variance and nugget variance parameters in the central Pacific. This suggests that allowing the correlation parameters to vary seasonally could present a meaningful improvement to the statistical modeled presented here. Since the estimation of the posterior distribution of ocean heat content is fairly constrained by the amount of information present in the data, further improvements in modeling the covariance structure of the ocean heat content field will likely have only a small effect on the strength of the ultimate conclusions. It should also be noted that the trend found here is broadly consistent with other entries in the literature [Cheng et al., 2019, Resplandy et al., 2019], although a more in-depth analysis of how the mean estimates and uncertainty levels compare between different approaches, and how these estimations influence estimates of Equilibrium climate sensitivity, would be an interesting extension in future work.

References

- Samuel Baugh and Karen McKinnon. Hierarchical bayesian modeling of ocean heat content and its uncertainty. *The Annals of Applied Statistics*, 16(4):2603–2625, 2022.
- Henk A. Dijkstra. *Dynamical Oceanography*. Springer, November 2008. ISBN 3540763759. URL <https://www.xarg.org/ref/a/3540763759/>.

- Dean Roemmich and John Gilson. The 2004–2008 mean and annual cycle of temperature, salinity, and steric height in the global ocean from the Argo program. *Progress in Oceanography*, 82(2):81 – 100, 2009. ISSN 0079-6611. URL <http://www.sciencedirect.com/science/article/pii/S0079661109000160>.
- Christopher J. Paciorek and Mark J. Schervish. Spatial modelling using a new class of non-stationary covariance functions. *Environmetrics*, 17(5):483–506, 2006. ISSN 1180-4009. URL <https://www.ncbi.nlm.nih.gov/pubmed/18163157>.
- Michael H Glantz, Michael H Glantz, et al. *Currents of change: impacts of El Nino and La Nina on climate and society*. Cambridge University Press, 2001.
- Michael L Stein. *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 1999.
- Jopseh Guinness and Matthias Katzfuss. GpGp: Fast Gaussian process computation using Vecchia’s approximation. *R package version 0.1. 0*, 2018.
- Matthias Katzfuss, Marcin Jurek, Daniel Zilber, Wenlong Gong, Joseph Guinness, Jingjie Zhang, and Florian Schäfer. GPvecchia: Fast Gaussian-process inference using Vecchia approximations. *R package version 0.1*, 3, 2020a.
- Matthias Katzfuss, Joseph Guinness, Wenlong Gong, and Daniel Zilber. Vecchia approximations of Gaussian-process predictions. *Journal of Agricultural, Biological and Environmental Statistics*, 25(3):383–414, 2020b.
- Matthias Katzfuss and Joseph Guinness. A general framework for Vecchia approximations of Gaussian processes. *Statistical Science*, 36(1):124–141, 2021.
- Philip Heidelberger and Peter D. Welch. A spectral method for confidence interval generation and run length control in simulations. *Communications of the ACM*, 24(4):233–245, 1981.
- Lijing Cheng, John Abraham, Zeke Hausfather, and Kevin E Trenberth. How fast are the oceans warming? *Science*, 363(6423):128–129, 2019.

Laure Resplandy, Ralph F Keeling, Yassir Eddebbar, Mariela Brooks, Rao Wang, Laurent Bopp, Matthew C Long, John P Dunne, Wolfgang Koeve, and Andreas Oschlies. Quantification of ocean heat uptake from changes in atmospheric o₂ and co₂ composition. *Scientific reports*, 9(1): 1–10, 2019.