Geoscience Data Journal

WILEY

Geoscience Data Journal RMetS

DATA SERVICES ARTICLE OPEN ACCESS

GloMarGridding: A Python Toolkit for Flexible Spatial Interpolation in Climate Applications

Richard C. Cornes[1] | Steven C. Chan[1] | Archie Cable[1] | Duo Chan[2] | Agnieszka Faulkner[1] | Elizabeth C. Kent[1] | Joseph T. Siddons[1]

[1]National Oceanography Centre, Southampton, UK | [2]University of Southampton, Southampton, UK

Correspondence: Richard C. Cornes (richard.cornes@noc.ac.uk)

Received: 21 August 2025 | **Revised:** 12 December 2025 | **Accepted:** 8 January 2026

ABSTRACT

Global surface temperature datasets are constructed through processing chains that inherently introduce structural uncertainty, arising from choices made both in the processing of input observations and in the spatial interpolation methods employed. Because these steps are often tightly integrated, it is difficult to isolate their individual contributions to uncertainty. Here, we introduce GloMarGridding, a Python package designed to support the evaluation of the component of structural uncertainty arising specifically from spatial interpolation. It provides tools to apply Gaussian Process Regression Modelling (GPRM), widely used in the production of global temperature datasets, enabling the generation of spatially complete temperature fields from grid-box average and point observations, along with estimation of uncertainty in those fields. GloMarGridding currently supports three spatial covariance parametrizations: fixed isotropic variograms, ellipse-based anisotropic, and empirically derived covariance matrices. It also allows for uncertainty propagation via error covariance matrices and conditional simulation from input ensembles. By decoupling spatial interpolation from earlier stages of dataset development—such as homogenization, quality control, and aggregation—this framework enables independent assessment of upstream processing choices and their impacts on gridded outputs.

1 | Introduction

Global surface temperature datasets—such as those used in the Intergovernmental Panel on Climate Change (IPCC) Assessment Report (IPCC 2021)—rely on a range of techniques to generate smoothed, infilled gridded fields from available observations (e.g., Lenssen et al. 2019; Kadow et al. 2020; Rohde and Hausfather 2020; Morice et al. 2021; Huang et al. 2022). The construction of these datasets involves many processing decisions, and the resulting variations in how each dataset represents the temperature field contribute to what is known as structural uncertainty (Thorne et al. 2005). This structural uncertainty has two primary sources: (1) differences in the selection, adjustment and processing of the input temperature measurements, and (2) differences in the spatial interpolation methods applied. Because these steps are often tightly integrated in the processing chains used by dataset producers, it can be difficult to determine their individual contributions. This was noted by Huang et al. (2016) in their evaluation of parametric uncertainties in the ERSSTv4 sea surface temperature (SST) dataset. Although the influence of parameter choice in ERSSTv4 could be evaluated by the dataset developers, this was not the case for the other SST datasets against which ERSSTv4 was compared. This hampered the evaluation of the causes of observed differences in uncertainty estimates across the datasets. Given the difficulty in accounting for the individual components of uncertainty in a consistent manner, recourse is usually made to the spread across datasets being taken as an estimate of structural uncertainty (Kennedy 2014). Such evaluations using the available "ensemble of opportunity" are valuable, but can provide false confidence if there are common biases across datasets (Sippel et al. 2024).

In this paper, we present GloMarGridding, a software package developed to facilitate the spatial interpolation component of structural uncertainty assessments. The package enables users to spatially interpolate grid-box average estimates and generate uncertainty estimates in those fields using Gaussian Process Regression Modelling (GPRM, often called Kriging: Rasmussen and Williams 2005; Cressie 1993). GloMarGridding allows two types of kriging to be performed (Ordinary and Simple) and three different methods of constructing the spatial covariance matrix: (1) using a selection of fixed/isotropic variograms[1]; (2) using an empirically-derived covariance matrix; and (3) generating a parametrized covariance matrix through the fitting of spatially-varying, anisotropic ellipses to the full covariance matrix after Karspeck et al. (2012). Further advantages of the code are that it allows propagation of correlated and uncorrelated uncertainty estimates from the source data through to the final gridded fields, and that it seamlessly accounts for the boundary at the −180°/180° longitude meridian; this is particularly important when interpolating global data. The grid-box estimate is a fundamental concept of the package since most climate datasets first form such estimates from point observations in order to reduce the computational burden of interpolation, and to allow careful management of representation uncertainty in the interpolated fields. Accordingly, although point observations can be supplied to GloMarGridding, these data will be aggregated to grid estimates prior to any further computation.

The techniques used in GloMarGridding build on established methods that are used for generating surface temperature fields (e.g., Lenssen et al. 2019; Rohde and Hausfather 2020; Morice et al. 2021). By decoupling spatial interpolation from earlier processing stages—such as homogenization, quality control, and aggregation to grid-cell averages—the software allows users to create spatially complete fields while independently assessing the effects of prior data processing choices.

2 | Methods & Algorithms

In this section, we describe the key components of GloMarGridding and its mathematical foundation and algorithms. A table of symbols and variables can be found in the Appendix (Table A1).

2.1 | The Kriging Module

The foundation of **`GloMarGridding`** is Gaussian Process Regression Modelling (GPRM), which when applied to geophysical quantities is commonly called *kriging* (Cressie 1993; Rasmussen and Williams 2005). Currently, **`GloMarGridding`**

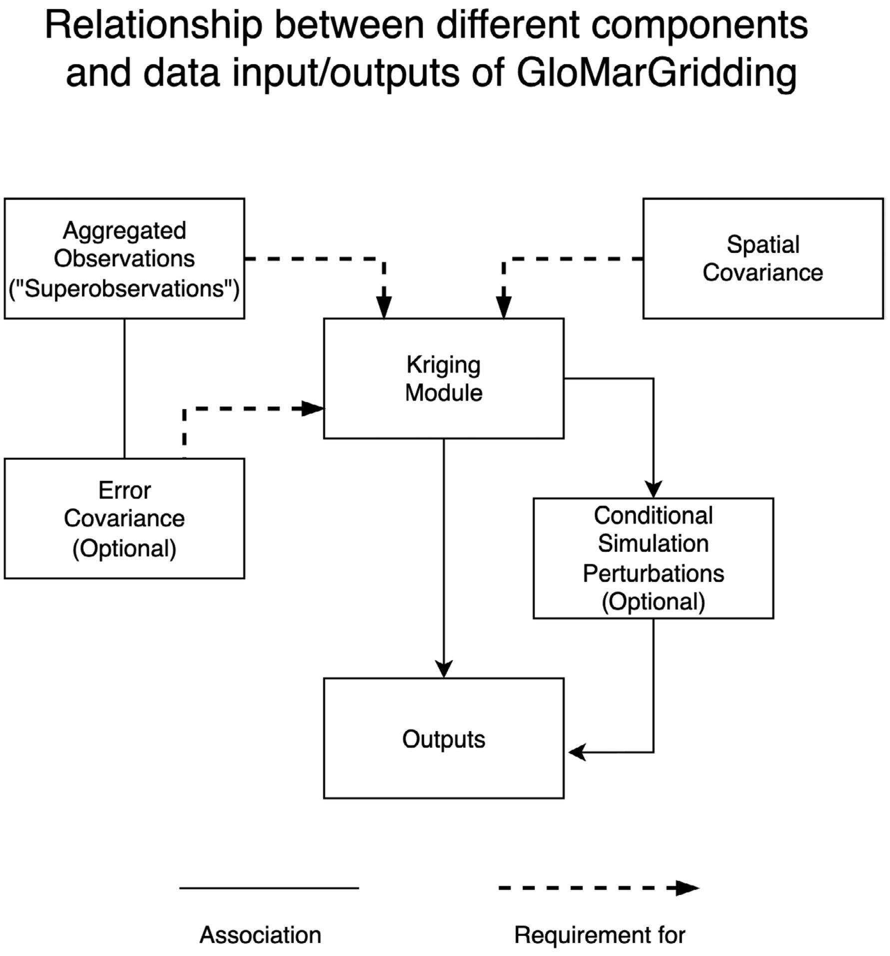

FIGURE 1 | The components of GloMarGridding and the relationship between its inputs and attributes.

can apply both *Simple* and *Ordinary* kriging to an input vector of observations ($\mathbf{y}$) to generate a smoothed/interpolated field. This is implemented using classes and methods in the **`kriging`** module (Figure 1).

Under *Simple Kriging*, an interpolated field ($\bar{\mathbf{g}}_{SK}$) is constructed as

$$\bar{\mathbf{g}}_{SK} = \mathbf{C}_{\text{cross}}^{\intercal}\left(\mathbf{C}_{\text{obs}} + \mathbf{E}\right)^{-1}\mathbf{y} + \mu, \tag{1}$$

by taking a constant known mean value (μ), the spatial covariances between grid cells that contain observations (C_{obs}), the cross-covariance between observational grid-cells and all grid cells (C_{cross}) and optionally an error covariance matrix (E). This is enabled by the **`SimpleKriging`** class, and associated methods, which assumes that the observations y—along with their associated spatial and error covariance matrices—are provided on a uniform grid. In climate data analysis based on in situ observations, these values typically represent grid-cell averages that have been derived from point measurements recorded in space and time. **`GloMarGridding`** includes tools to construct such grid-cell averages, or *super-observations*, from point data prior to spatial interpolation. This process is handled by the **`kriging.prep_obs_for_kriging`** function, which maps point observations onto the target grid and computes averages, optionally applying user-defined weights. Functions in the **`mask`** module ensure consistent mapping of observations to covariance matrices when regional masking is required. A common application is the global interpolation of air temperature anomalies. Because marine and terrestrial data differ in their pre-processing methods, error structures and spatial covariance characteristics, these two domains are often interpolated separately (see for example Morice et al. 2021).

Due to the requirement for the observations to be on a uniform grid, C_{obs} and C_{cross} are subsets from the spatial covariance across all grid cells (C). That covariance matrix may be supplied directly, for example if this has been derived empirically (e.g., Williams and Berry 2020), and can be estimated using spatially-stationary variograms or spatially varying ellipses (see Section 2.3). Because the covariance function exerts a fundamental control on the analysis, both the choice of method and the specification of its parameters must be made with care (see for example Cressie 1993; Rasmussen and Williams 2005). Of the three methods implemented in **`GloMarGridding`**, the stationary/isotropic model is the simplest, but it may not adequately capture the spatially varying structure of global surface temperature fields unless the domain is partitioned into regions, each with its own variogram. By contrast, the empirical covariance matrix offers a more detailed representation of spatial structure, but can introduce spurious features in data-sparse regions or periods. The spatially varying ellipse approach provides a practical compromise between these two extremes.

The covariance of the interpolated field is calculated as:

$$\mathbf{C}_{\bar{\text{g}},\text{SK}} = \mathbf{C} - \mathbf{C}_{\text{cross}}^{\intercal}\left(\mathbf{C}_{\text{obs}} + \mathbf{E}\right)^{-1}\mathbf{C}_{\text{cross}} \tag{2}$$

The **`solve`** method from the **`SimpleKriging`** class returns the interpolated field $\bar{g}_{SK}$, and the **`get_uncertainty`** method returns the uncertainty $\left(\sqrt{\text{diag}\left(\mathbf{C}_{\bar{\text{g}},\text{SK}}\right)}\right)$, which is widely taken to represent the uncertainty in $\bar{\mathbf{g}}_{SK}$ (Cressie 1993).

The optional error covariance matrix E can be provided as a pre-computed matrix, or calculated using functions contained in the **`error_covariance`** module. A common application of this error propagation is in the generation of interpolated fields across the ocean from ship data. The error in the observations is usually separated into a random component that is uncorrelated between observations and a systematic component that is correlated per ship, and hence may traverse grid cells (Kennedy 2014). The generation of matrices to quantify these components is enabled by the **`uncorrelated_components`** and **`correlated_components`** functions in the **`error_covariance`** module.

The *Simple Kriging* approach tends to have a rather limited application in the interpolation of climate fields, because of the requirement to supply a known, constant mean value (μ). In such applications, *Ordinary Kriging* is generally preferred because it ensures that the global mean value is calculated as the weighted mean of all observations. A disadvantage is that in data sparse regions the interpolated value will regress towards that average value. Indeed, The **`SimpleKriging`** class is mainly used in **`GloMarGridding`** in the estimation of sampling uncertainty through conditional simulation (where $\mu = 0$, see Section 2.5). However, *Simple Kriging* could be applied to the interpolation of global fields if a reliable global average value (μ) could be estimated independently. The flexibility of these functions would permit comparison of the two different approaches in that case.

Ordinary Kriging is applied in **`GloMarGridding`** via the **`OrdinaryKriging`** class, which has the same structure as the **`SimpleKriging`** class. This approach estimates an a priori unknown global mean value ($\bar{\beta}$) through the addition of a Lagrange multiplier (block matrices 0 and 1) to the *Simple Kriging* relationship shown in Equation 1:

$$\begin{bmatrix} \bar{\mathbf{g}}_{OK} \\ \bar{\beta} \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{\text{cross}} \\ 1 \end{bmatrix}^{\intercal} \begin{bmatrix} \mathbf{C}_{\text{obs}} + \mathbf{E} & \mathbf{1} \\ \mathbf{1}^{\intercal} & 0 \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{y} \\ 0 \end{bmatrix} \tag{3}$$

which has a posterior covariance matrix of

$$\mathbf{C}_{\bar{\text{g}},\text{OK}} = \mathbf{C}_{\bar{\text{g}},\text{SK}} + \mathbf{R}^{\intercal}\left(\mathbf{C}_{\text{obs}} + \mathbf{E}\right)\mathbf{R} \tag{4}$$

$$\boldsymbol{R} = \mathbf{1} - \left(\mathbf{C}_{\text{obs}} + \mathbf{E}\right)^{-1}\mathbf{C}_{\text{cross}} \tag{5}$$

2.2 | Masking of Highly Uncertain Values

In the generation of interpolated climate fields, it may be desirable to exclude grid-cells that have relatively high levels of uncertainty. These are often regions far away from observed gridboxes, and which therefore have estimates that are poorly constrained by observations. To enable this masking, the **`constrain_mask`** method from the kriging classes generate a measure of this relative uncertainty (α) following Morice et al. (2021, Note that Eq. A14 contains a typological error that is corrected here.):

$$\alpha = 1 - \frac{\mathrm{diag}\left(\mathbf{C}_{\overline{g}}\right)}{\mathrm{diag}(\mathbf{C})} \tag{6}$$

here, α represents the fraction of prior variance constrained by the kriging analysis at each location. For some applications, such as regional time series analysis, it may be desirable for grid cells with α below a threshold to be excluded (e.g., 0.25 was used for HadCRUT5). This becomes particularly important in data-sparse regions and periods, where the uncertainty in the interpolation is likely to be large relative to the expected variance at that location.

2.3 | Spatial Covariance Matrix Construction

2.3.1 | Spatially-Stationary Variograms

GloMarGridding contains functions and classes to calculate the spatial covariance from a selection of spatially-stationary variograms (Spherical, Gaussian, Exponential and Matérn; Cressie 1988). To initiate the classes for these models, the partial sill, nugget, effective range and range parameters need to be supplied. For the Matérn class the smoothing parameter (ν) and method of fitting (one of "sklearn", "gstat" or "karspeck") also need to be supplied. These classes do not contain functions to fit a variogram to the data, and if required the parameters need to be calculated outside of GloMarGridding, for example, using the scikit-gstat library. It should be stressed that in the classes/methods contained in **`kriging`** module, E (if provided) will be added to $\mathrm{C}_{\mathrm{obs}}$ within the function. Care should be taken therefore to ensure that the uncorrelated uncertainty is not included twice: via the diagonal elements of E and via the nugget variance argument in the variogram functions.

2.3.2 | Ellipse-Based Covariance Matrices

The **`ellipse`** module contains classes to construct a regionally-varying, non-isotropic representation of the spatial covariance. This method was developed by Paciorek and Schervish (2006) and was applied to climate data analysis by Karspeck et al. (2012), who demonstrated its application in the interpolation of sea-surface temperature (SST) data across the Atlantic Ocean. In that application, the method was used to represent mid-scale (order 1000km) SST variability, with the large-scale features of SST quantified via Empirical Orthogonal Functions (EOFs). The classes in this module allow the application of this method across the entire domain of interest (e.g., globally) and without the EOF reductions.

The key component of the spatial covariance quantified via the ellipse-based parametrization is the Matérn function $\widehat{C}$, and its spatial-scale kernel Σ. The Matérn function has a shape parameter ν; $\nu = 0.5$ is equivalent to the exponential variogram, while $\lim_{\nu \to \infty} \widehat{C}$ is a Gaussian-shaped variogram. This value is constant over the field.

Consider an origin x_i on a plane. Its observed regional covariance $\boldsymbol{C}\left(x_i, x'\right)$ with nearby points x' is modelled using $\widehat{\boldsymbol{C}}$. The spatially-varying standard deviations of each origin point is denoted $\sigma\left(x_i\right)$ and the nearby points $\sigma\left(x'\right)$:

$$\widehat{\boldsymbol{C}}\left(x_i, x', \boldsymbol{\Sigma}, \nu\right) = \frac{\sigma\left(\boldsymbol{x}_i\right)\sigma\left(\boldsymbol{x}'\right)}{\Gamma(\nu)2^{\nu-1}}\left(2\overline{\tau}\left(\boldsymbol{x}_i, \boldsymbol{x}', \boldsymbol{\Sigma}\right)\sqrt{\nu}\right)^{\nu} \boldsymbol{K}_{\nu}\left(2\overline{\tau}\left(\boldsymbol{x}_i, \boldsymbol{x}', \boldsymbol{\Sigma}\right)\sqrt{\nu}\right) \tag{7}$$

$$\boldsymbol{\tau}\left(\boldsymbol{x}_i, \boldsymbol{x}', \boldsymbol{\Sigma}\right) = \sqrt{\left(\boldsymbol{x}_i - \boldsymbol{x}'\right)^{\mathsf{T}} \boldsymbol{\Sigma}^{-1}\left(\boldsymbol{x}_i - \boldsymbol{x}'\right)} \tag{8}$$

$$\boldsymbol{\Sigma}\left(L_x, L_y, \theta\right) = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} L_x^2 & 0 \\ 0 & L_y^2 \end{bmatrix} \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}^{\mathsf{T}}, \tag{9}$$

in which Γ is the Gamma function and K_ν is the modified Bessel function of the second kind, with order ν. τ is the Mahalanobis distance between x_i and x' (Dodge 2008). This is the distance from the origin normalised by Σ and requires the plane-geometry vector displacements $(x_i - x')$ to be computed on a sphere. The scale of surface temperature variation are large enough that the curvature of the Earth is non-negligible and to approximate that, the **`compute_params`** method in the **`EllipseBuilder`** class includes two approximations (sinusoidal or cylindrical) that can be selected by the user:

$$x_i - x' = 111.12\frac{\mathrm{km}}{} \times \begin{cases} \dfrac{\cos\left(\mathrm{Lat}_i\right) + \cos\left(\mathrm{Lat}'\right)}{2}\left(\mathrm{Lon}_i - \mathrm{Lon}'\right)\hat{\mathbf{i}} + \left(\mathrm{Lat}_i - \mathrm{Lat}'\right)\hat{\mathbf{j}}, & \text{sinusoidal} \\ \left(\mathrm{Lon}_i - \mathrm{Lon}'\right)\hat{\mathbf{i}} + \left(\mathrm{Lat}_i - \mathrm{Lat}'\right)\hat{\mathbf{j}}, & \text{cylindrical} \end{cases} \tag{10}$$

The cylindrical projection is a simplification and hence an easier calculation, since the orthogonal space can be defined by lines of constant latitude and longitude. The sinusoidal projection is more realistic and is the default. However, the zonal displacement depends on whether the calculation is performed on the start or end latitude value. In the **`compute_params`** function, the average of the two latitudes is used. A maximum distance between x_i and x' for the training data is set via the **`max_dist`** parameter (default of 6000 km) in the **`EllipseCovarianceBuilder`** class, and is evaluated using Haversine distances. Alternatively, the distance value can be supplied in degrees to form a longitude/latitude box around the target location. This is often more convenient for constructing ellipses across small regions. However, for large areas this is not advised as spherical distortion may produce large differences in the sample sizes used to construct the ellipses.

Given $x_i - x'$, $\boldsymbol{C}\left(x_i, x'\right)$ and a user supplied fixed ν, $\Sigma\left(L_x, L_y, \theta\right)$ that minimises the square differences between the regional $\boldsymbol{C}\left(x_i, x'\right)$ and $\widehat{\boldsymbol{C}}\left(x_i, x'\right)$ is solved; technically, this is the minimisation of the negative log likelihood for $\widehat{\boldsymbol{C}}(\Sigma)$ in explaining $\boldsymbol{C}$. In **`GloMarGridding`**, fits are preformed against Fisher-transformed (i.e., inverse hyperbolic tangent, σ-normalised) correlation. As in Karspeck et al. (2012), **`GloMarGridding`** uses Nelder–Mead minimisation (Nelder and Mead 1965) by default, but any other **`scipy.optimise.minimise`** method can be selected by the user.

Once L_x, L_y and θ are estimated, the global covariance matrix $\mathbf{C}$ is constructed by combining the $\widehat{\boldsymbol{C}}$ s for each grid point by averaging the kernels between point pairs. This follows Equation 17 in Karspeck et al. (2012) in which element m and n of $\mathbf{C}$ is:

$$\mathbf{C}_{mn} = \frac{\sigma\left(x_m\right)\sigma\left(x_n\right)}{\Gamma(\nu)2^{\nu-1}} \frac{\det\left(\boldsymbol{\Sigma}_m\right)^{1/4}\det\left(\boldsymbol{\Sigma}_n\right)^{1/4}}{\det\left(\overline{\boldsymbol{\Sigma}}\right)^{1/2}}\left(2\overline{\tau}\left(x_m, x_n, \overline{\boldsymbol{\Sigma}}\right)\sqrt{\nu}\right)^{\nu} K_{\nu}\left(2\overline{\tau}\left(x_m, x_n, \overline{\boldsymbol{\Sigma}}\right)\sqrt{\nu}\right) \tag{11}$$

$$\bar{\tau}(x_m, x_n, \bar{\Sigma}) = \sqrt{(x_m - x_n)^{\intercal} \bar{\mathbf{\Sigma}}^{-1} (x_m - x_n)} \tag{12}$$

$$\bar{\mathbf{\Sigma}} = \frac{\mathbf{\Sigma}_m + \mathbf{\Sigma}_n}{2} \tag{13}$$

Optionally, the method **`uncompress_cov`** can be used to expand the covariance matrix to cover the entire domain of interest. This can be used to resolve discrepancies between training data coverage and interpolated field coverage. For example, training observations may not include lakes or transient/partial sea ice grid points, but observations may be present for such areas. However, this method should be used with caution since the covariance structure for those points will have been estimated from neighbouring points.

2.4 | Eigenvalue-Clipping of Covariance Matrices

To ensure numerical stability of the constructed covariance matrix, the **`covariance_tools`** module contains functions to check for definiteness and adjust the matrix if necessary. If the Earth is represented as either flat or cylindrical, the constructed spatial covariance should be positive definite and computationally stable (Paciorek and Schervish 2006; Kennedy et al. 2019). Geometry changes and having a large number of grid points relative to a small number of data per grid point would destabilise the covariance (Gneiting 2013; Higham et al. 2016; Huang et al. 2017). Destabilisation may also arise from the horizontal grid resolution being too high relative to range scales. This may result in the constructed covariance matrix containing negative eigenvalues, making it non-positive definite. While this could occur in a covariance matrix constructed by any of the three methods implemented in **`GloMarGridding`**, it is more common in the ellipse method due to the heterogeneity of the range parameters.

Eigenvalue clipping is implemented in **`GloMarGridding`** to allow users to effect a repair of non-positive definite covariance matrices (Bun et al. 2017). In the **`covariance_tools.explained_variance_clip`** function, we retain the largest eigenvalues up to a chosen threshold, with a default of the largest EOFs summing to 95% of the explained variance. The threshold chosen is a balance between ensuring the matrix is positive-definite and retaining the higher modes of variability. The user can explore the sensitivity of this choice by selecting different thresholds to understand the impact for their particular application, provided that the selected eigenvalues have a positive mean and capture a sufficient proportion of the total explained variance. For example, Wilks (2006) and Jolliffe (2005) suggest that a range of 70%–90% may be appropriate. Note, however, that thresholds of that order can lead to large changes in the covariance matrix.

In the **`covariance_tools.eigenvalue_clip`** function there is also the option to automatically estimate a suitable threshold. This approach uses the ratio of the time series length (T) to the number of spatial points (N) after it is standardised to a correlation matrix (Laloux et al. 2000).

$$Q = \frac{N}{T} \tag{14}$$

$$\lambda_{\text{threshold}} \approx \left(1 + \sqrt{Q}\right)^2 = \left(1 + Q + 2\sqrt{Q}\right) \tag{15}$$

The automatic method selects eigenvalues in a manner that aims to optimise the signal:noise rato represented by the EOFs. When $N \gg T$, only a small number of eigenvalues will be retained but note that even noisy or degenerate EOFs still represent components of the original sample covariance matrix (North et al. 1982). In contrast, the fixed-threshold selection method—particularly when using the 95% default threshold—simply aims to ensure positive definiteness, making only relatively small changes to the matrix. For example, in the SST interpolation example below (Sec. 3.2), 245 out of a possible 1484 eigenvalues are retained when a fixed threshold of 95% is chosen, compared to fewer than 10 retained using the automatic method.

Regardless of the method chosen to define an appropriate threshold, eigenvalues less than the threshold are replaced by the average of eigenvalues below the threshold. A new matrix is built with the modified eigenvalues and their original eigenvectors. If $\Lambda = \{\lambda_1, \ldots, \lambda_N\}$ are descending-ordered eigenvalues with corresponding eigenvectors V to the non-standardised covariance matrix $\mathbf{C}$:

$$\tilde{\Lambda} = \begin{cases} \lambda_1, \ldots, \lambda_M & \text{for } \sum_{i=1}^{M} \lambda_i < 0.95\text{Tr}(\mathbf{C})(\text{fixed})\text{or } \lambda > \lambda_{\text{threshold}}(\text{automatic}) \\ \boldsymbol{E}\left(\lambda_{M+1}, \ldots, \lambda_N\right) & \text{otherwise} \end{cases} \tag{16}$$

$$\tilde{\text{C}} = \text{Vdiag}(\tilde{\Lambda})\text{V}^{\intercal} \tag{17}$$

in which $\tilde{\text{C}}$ is a positive definite covariance matrix, assuming that the expected value of the discarded eigenvalues satisfies $\boldsymbol{E}\left(\lambda_{M+1}, \ldots, \lambda_N\right) > 0$. When the automatic threshold selection method (Equation 15) is used, the module applies a transformation involving $\sqrt{diag(\mathbf{C})}$ and its reciprocal where appropriate to convert between correlation and covariance matrices.

The total variance of the reconstructed matrix $\tilde{\mathbf{C}}$ is equal to that of the original $\mathbf{C}$; thus, the total variance is conserved. The eigenvectors that are not clipped retain their original explained variance, while the clipped eigenvectors are assigned new, adjusted values. Regardless of the threshold chosen, a full eigendecomposition is required to perform eigenvalue clipping.

2.5 | Conditional Simulation of Sampling Uncertainty

Although GloMarGridding produces uncertainty estimates as part of the kriging process—specifically through the extraction of kriging variance values (see Section 2.1)—the **`perturbation`** module also provides functionality for generating conditional simulations. This is enabled via the **`StochasticKriging`** class, which has the same structure and methods as the other kriging classes. The approach used in this class samples from the Gaussian process posterior distribution and is commonly used to quantify sampling uncertainty in climate data fields. Following the method described by Morice et al. (2021), **`GloMarGridding`** implements the procedure in six steps:

1. Stochastically generate a random zero-mean global field by drawing from the spatial covariance matrix (i.e., $\mathrm{MVN}(0, \mathbf{C})$).
2. Sub-sample the draw at observed locations to create a vector of pseudo-observations.
3. Draw from the error covariance matrix (i.e., $\mathrm{MVN}(0, \mathrm{E})$) and add the draw to perturb the pseudo-observations.
4. Generate an interpolated, zero-mean field using the perturbed pseudo-observations.
5. Compute the difference between the initial stochastic field and the interpolated zero-mean field to produce a single simulated field.
6. The simulated field is then added to the data-derived field to produce the final analysis field.

In the HadCRUT5 global temperature dataset, a different realisation is generated for each member of the 200-member input ensemble across the land and sea domains (Morice et al. 2021). As such, the uncertainty in the interpolated field arises from both the Gaussian process model, and the propagation of measurement uncertainty from the ensemble and error covariance matrix. However, this method treats each field independently and does not incorporate temporal propagation of uncertainty. In HadCRUT5, this limitation is addressed by fixing the stochastic field within each year, effectively assuming perfect temporal correlation over that period. A similar approach can be implemented with **`GloMarGridding`** when using a fixed covariance matrix, for example, derived from an isotropic variogram or a fixed set of ellipse parameters. However, this strategy cannot be applied when the covariance matrix varies throughout the year, for example, when a different set of ellipse parameters are used for each climatological month. Nonetheless, this provides a viable means of propagating uncertainty from data pre-processing through to the gridded fields. Using **`GloMarGridding`**, users can apply different spatial interpolation methods to these data, and by separating the interpolation stage from earlier processing steps, quantify the proportion of variance in the final dataset that arises specifically from spatial interpolation. It should be noted, however, that the uncertainty estimation functions currently implemented in **`GloMarGridding`** do not account for uncertainty in the covariance function parameters, as for example is enabled by the Maximum Likelihood methods described by Nychka et al. (2021).

2.6 | Unit Testing

The modules contained within the GloMarGridding library are supported by a testing suite covering 77% lines of code. These unit-tests form part of an automated CI/CD scheme run with GitHub Actions, and the tests are performed for python versions 3.11, 3.12, 3.13 and 3.14. The implementation of the *Ordinary Kriging* algorithm is tested against the implementation in **`GeoStats.jl`** (Hoffimann 2018) for a simple example, excluding error covariance. An expected output is calculated using the equivalent method in **`GeoStats.jl`**, and compared to the output from GloMarGridding. The Variogram classes and eigenvalue-clipping methods are tested to ensure they result in symmetric positive-definite matrices. The **`ellipse`** module is tested by checking that the full process is self-consistent. A covariance matrix is first generated using a known set of ellipse parameters. This covariance matrix is used to construct a dataset from which a new set of ellipse parameters are calculated (with the same configuration as used to generate the initial covariance matrix). These calculated parameters are then used to generate a new covariance matrix, which should be similar to the original covariance matrix. To determine that degree of similarity, the correlation matrix distance metric of Herdin et al. (2005) is used. For two arbitrary correlation matrices $\mathrm{R}_1, \mathrm{R}_2$, the metric is defined as:

$$d_{corr}\left(\mathbf{R}_1, \mathbf{R}_2\right) = 1 - \frac{\mathrm{Trace}\left(\mathbf{R}_1 \mathbf{R}_2\right)}{\|\mathbf{R}_1\|_f \|\mathbf{R}_2\|_f} \tag{18}$$

If the two correlation matrices are identical, d_{corr} would be 0. Maximum differences are indicated by a value of 1. An error is raised in the Unit Test if d_{corr} between the prescribed and generated matrix is greater than 10^{-3}.

3 | Example Applications of GloMarGridding

3.1 | Calculation of Kernel Parameters

To illustrate the parametrization of a covariance matrix using the ellipse method, we derived ellipse parameters for the global ocean and land domains using sea surface temperature data from the ESA Climate Change Initiative (ESA-CCI) (Embury et al. 2024) and 2m air temperature (t2m) data from ERA5 (Hersbach et al. 2020; Centre for Environmental Data Analysis 2024), respectively. The land domain was defined using the ESA-CCI land cover dataset (Lamarche et al. 2017) as any grid points with a land fraction > 0.01. In addition, any grid points that were excluded from the SST ellipse fits were also considered to be land. That included grid points that are sea-ice masked and larger lakes. Ellipses were fitted separately for each month of the year using samples of data over the period 1982–2022 for SST and 1979–2023 for t2m; the settings used for ellipse construction are listed in Table 1.

Prior to fitting the ellipses, the SST data were regridded to a 5° regular grid, from the 0.05° native resolution, and averaged to monthly anomalies relative to a 1982–2010 base period using the https://surftemp.net/ regridding service. Under that system the sea ice exclusion threshold was set to 15% sea ice concentration—a commonly used definition for the sea ice edge (Titchner and Rayner 2014). Grid-cells must be ice free according to that criterion in all of the 29 years for the ellipse parameters to be calculated. In the case of the ERA5 t2m data, a daily 30-year climatology was first computed at the native spatial resolution of 0.25°. These daily values per grid-cell were then temporally smoothed with a 10-day low-pass Lanczos filter (Duchon 1979), anomalies were computed and those values were upscaled to $5° \times 5°$ using the bilinear interpolation function from the Climate Data Operators package (Schulzweida 2023).

The convergence of the ellipse calculation was generally very good. One sea grid cell near Svalbard failed to converge because

TABLE 1 | Settings used in the fitting of the ellipse parameters. * 6667.2 km is the distance covered by 60° latitude of change (i.e., 3600 nautical miles).

Configuration	Sea	Land
Training Data	1982–2022 ESA-CCI level-4 SST anomalies (Embury et al. 2024) (baseline 1982–2010)	1979–2023 ERA5 HRES (Hersbach et al. 2020) 2 m air temperature anomalies (baseline 1981–2010)
Resolution	Averaged to 5° × 5° monthly means using University of Reading data delivery tool (SurfTemp.net 2026), excluding grid boxes with sea ice fraction ≥ 15 % as defined by the same dataset	0.25° × 0.25° hourly, regridded to 5° × 5° monthly averages (see Section 3.1)
Coverage	Marine grid boxes with no missing data on the same calendar month (see Section 3.1); lakes excluded except Caspian Sea	Any grid boxes with land fraction > 1% according to ESA-CCI land cover (Lamarche et al. 2017) plus grid boxes that are excluded from the sea analysis
Moving window region	Valid (ice-free ocean) grid boxes within 6667.2km*, using Haversine distance	Valid (land, lake, sea ice) grid boxes within 6667.2km*, using Haversine distance
ν	1.5	
L_x, L_y, θ limits	$[300\text{km}, 30000\text{km}], [-2\pi, 2\pi]$	
L_x, L_y, θ initial values	2000km, 2000km, 0	
Optimiser	Nelder–Mead	
Post hoc quality checks	If L_x becomes ≈ 30000km, L_x is set to L_y (see Section 3.1)	

L_x reached its maximum allowable limit. This issue arose due to a lack of valid ice-free sea grid cells within the 75°–80°N latitude band. An isolated terrestrial grid box over the Kerguelen Islands in the southern Indian ocean had a similar problem due data sparsity. In contrast, the L_y values at those locations did not have this problem. Therefore, at those locations, we set $L_x = L_y$ to ensure consistency.

Despite their simplicity, the ellipse-derived correlations offer a much more realistic parametrization compared to a stationary isotropic model. This is demonstrated in Figure 2 where ellipse-derived correlations for two grid points are shown alongside observed correlations and the isotropic, globally constant value used in HadCRUT5. Length scales at both locations are noticeably underestimated using the global constant value. The ellipses provide a much better approximation of the empirical length scales (panels a and c), and this is particularly the case when all global ellipses are combined to construct the full covariance matrix (panels b and d).

Beyond the two case studies shown in Figure 2, the distribution of values for December are shown in Figure 3. Panels a and b illustrate key contrasts between land and ocean. Over land, the spatial length scales typically range from 1000 to 3000 km, while the variability in standard deviation σ is large—from about 0.5°C to over 5.0°C—reflecting differences between low and mid-latitudes. In contrast, the ocean exhibits the reverse pattern: spatial scales are broader (ranging from short scales in boundary currents to ENSO-scale patterns exceeding 10,000 km; panel c), while σ values are more tightly distributed, with a maximum near 1.5°C which is close to the modal value over land (panel b). Anisotropy is typical across both domains (panels c and d); the majority of points lie below the $y = x$ line, and in some ocean regions L_x exceeds 10,000 km, while L_y remains near 2000 km. Maps of the Lx and Ly parameters for December are shown in Figure 4 separately for land and ocean.

On average, the ellipse-based results have longer spatial scales for both land and sea than those used in the HadCRUT5 analysis (Figure 5), with the geometric mean SST spatial scales being ≈ 80 % longer, with smaller differences over land. In addition, while our calculations of the global mean SST standard deviations are similar to HadCRUT5 (0.6°C), our estimates of the global mean terrestrial temperature standard deviation (1.5°C–2°C) are higher than the 1.2°C used in HadCRUT5. Furthermore, on an annual average basis the ellipses tend to be zonal in orientation, particularly over the land, although considerable monthly variability exists in these averages (Figure 5b). Notably, the boreal summer (JJA) land temperature estimates have globally averaged rotation angles of around −20°C. Given these improvements in the paramaterization of spatial covariance patterns, the technique has been used to interpolate the DCENT dataset (Chan et al. 2026).

3.2 | Interpolating HadSST4 for Two Sample Months

To demonstrate the capabilities of GloMarGridding for comparing different spatial interpolation methods, we present two example global SST fields generated using the HadSST4 dataset (version 4.1.1.0). The data are interpolated by *Ordinary Kriging* using: (1) an isotropic, globally-fixed variogram; and (2) a spatially varying ellipse model. The code to reproduce these examples is included in the **`GloMarGridding`** package as a Python notebook (**`Ellipse_to_grid.ipynb`**).

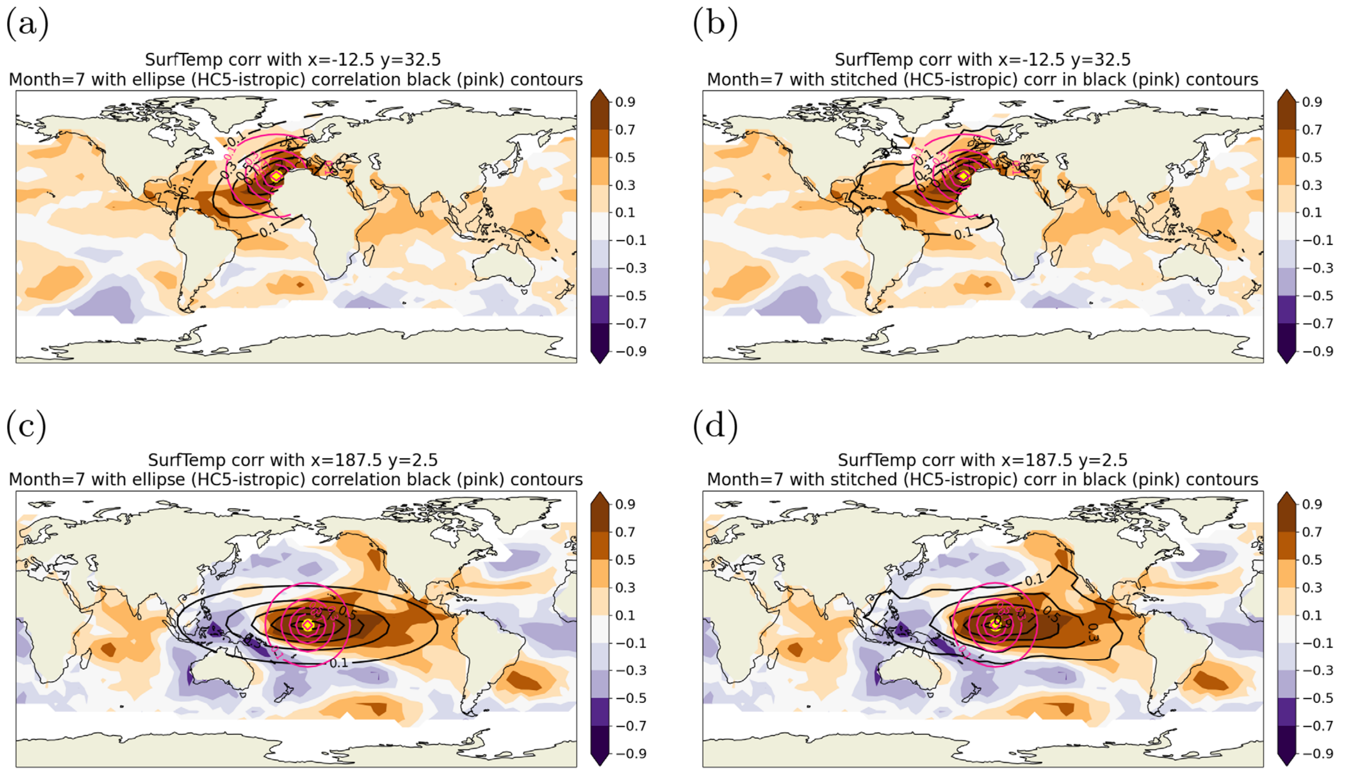

FIGURE 2 | Observed (colour) and statistically modelled (contoured) July global SST correlation from two grid points in (a, b) the north eastern Atlantic and (c, d) equatorial Pacific, calculated using ESA-CCI SST data (Embury et al. 2024). Panels a and c show the ellipse correlations at that grid point, while panels b and d show the correlations after they are blended with ellipse parameters elsewhere (Equation 11) and corrected for positive definiteness (Section 2.4). The magenta circles indicate the fixed spatial scale used in HadCRUT5. Due to notational differences in the Matérn covariance function (c.f. Equation 17 in Karspeck et al. (2012) with Equation 29 in Morice et al. (2021)), the quoted value of 1300 km in HadCRUT5 is equivalent to $1300\sqrt{2}$km ≈ 1840km following the notation used in Karspeck et al. (2012).

HadSST4 represents measurement-bias uncertainty in SST measurements through an ensemble of 200 realisations. In this example we interpolate SST data from the 94th ensemble member for March 1876, and from the 73rd ensemble member for March 2014. These cases illustrate two contrasting climate states: March 1876 follows the peak of the 1875 La Niña event (Singh et al. 2018), while March 2014 coincides with the onset of a major marine heatwave in the Northeast Pacific (Chen et al. 2023). They also highlight differences in data availability, with 1876 representing a data-sparse period and 2014 reflecting a relatively data-rich era.

The processing begins by constructing the grid specification via the **`grid_from_resolution`** function. In this case the grid-spacing is $5^\circ \times 5^\circ$ and the grid bounds are set to $(-87.5, 90)$ latitude and $(-177.5, 180)$ longitude to generate a global interpolation. In this example two versions of the monthly fields are constructed: one with a globally stationary covariance matrix and one with a non-stationary (ellipse-based) covariance matrix. The stationary covariance uses a Matérn function with the range = 1300 km and partial sill = 1.2°C. $\nu = 1.5$, which is considered to be an intermediate value that offers a good balance between smoothness and flexibility (Guttorp and Gneiting 2006). The nugget variance is set to 0.0°C because that measure of uncertainty is quantified via the incorporation of the error covariance matrix. This spatial covariance matrix is constructed in two steps: firstly the variogram is fitted via the **`MaternVariogram class`** and this is then supplied to the function **`variogram_to_covariance`** to construct the covariance matrix provided by HadSST4.

The ellipse-based covariance matrix is constructed as follows:

1. Define an **`EllipseModel`** which sets up default parameters, default values, and bounds for fitting.
2. Load data for Ellipse parameter estimation (in this case ESA CCI SST)
3. Mask the training data if required (in this example land points are masked and excluded)
4. Define an **`EllipseBuilder`** from the training data.
5. Loop through each grid-cell and compute Lx, Ly, θ, and σ values at each cell using the **`compute_params`** method
6. Use the output to estimate the covariance at ocean locations using the **`EllipseCovarianceBuilder`** class.
7. Extend the ocean covariance to full grid using the **`uncompress_cov`** method using a fill value of 1.2 along the diagonal of the matrix and 0 on the off-diagonal.
8. Ensure that the covariance is positive definite with the **`covariance_tools.eigenvalue_clip`** function. A

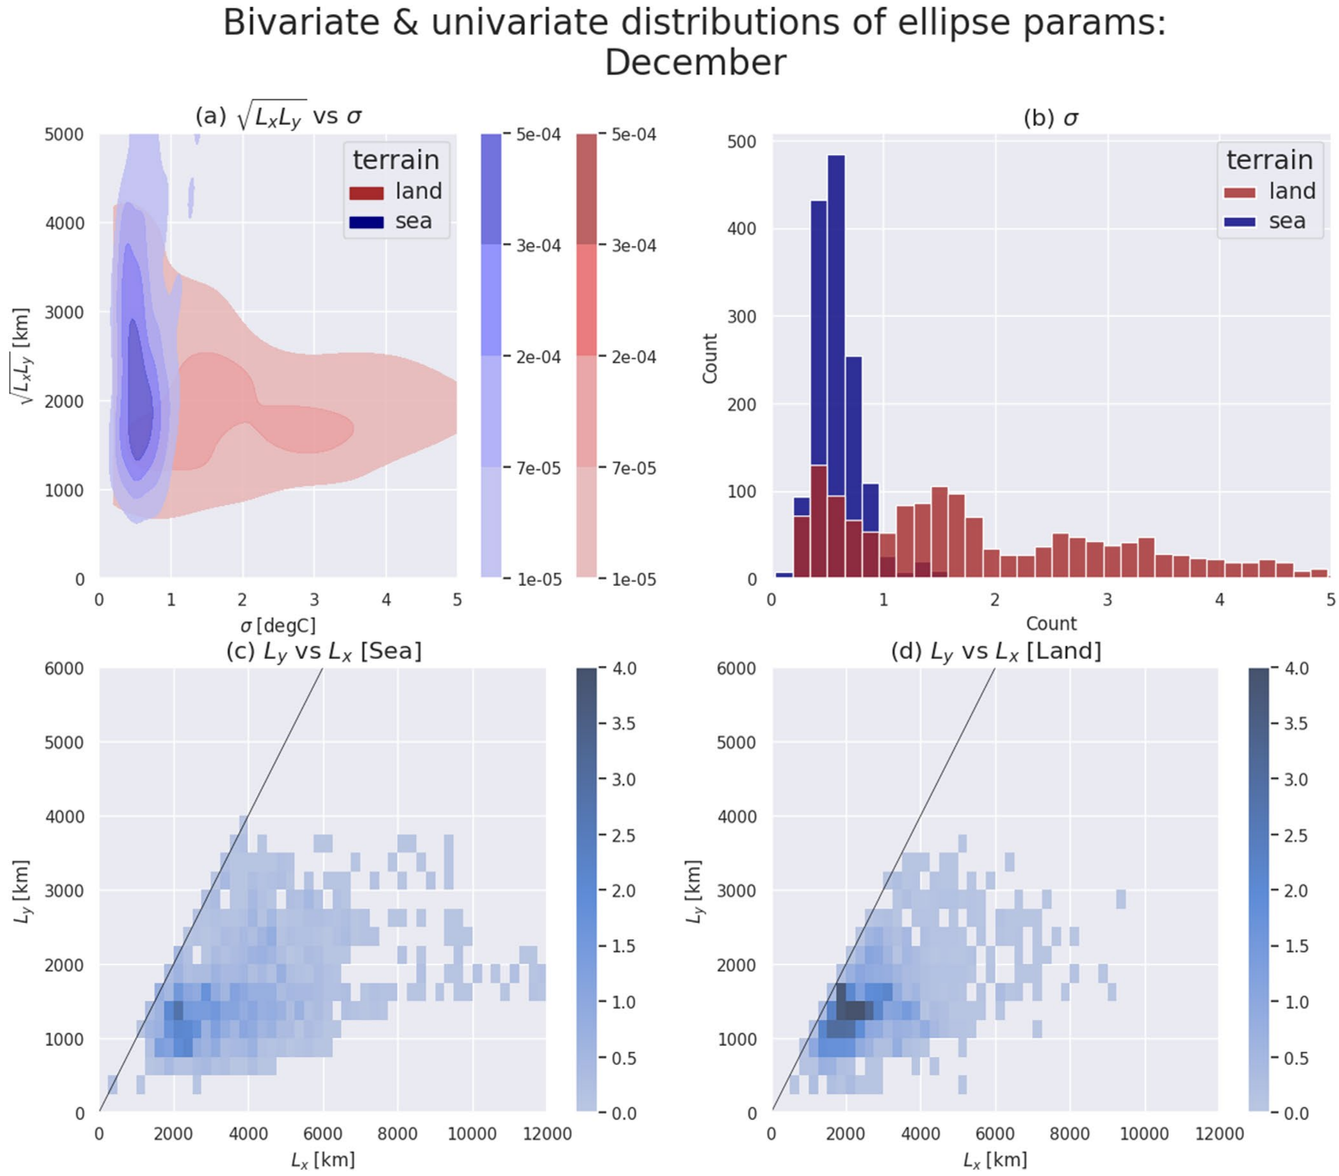

FIGURE 3 | The spatial distribution of L_x, L_y and σ for the month of December, separately for land and sea. Panel a shows the bivariate kernel density of $\sqrt{L_x L_y}$ against σ, red for land and blue for sea. Panel b shows the histogram of σ, following the same colour convention. Note that the different shades of red occur as a result of the overlap of the two histograms. Panels c and d show the bivariate histogram of L_y against L_x. For these panels, L_x always indicate the semi-major axis of the ellipses (i.e., the longer of the two). The black line indicates $y = x$ (i.e., isotropy).

fixed threshold of 95% is chosen, which retains 245 out of a possible 1484 eigenvalues.

Spatial interpolation then proceeds for the stationary covariance case by supplying the spatial covariance matrix to the **`OrdinaryKriging`** class, which implements Equations 3 and 4. The **`solve`** method from that class is then called with the SST data, grid index and error covariance matrix to generate the gridded field; uncertainty estimates are generated by calling the **`get_uncertainty`** method from the class. In the ellipse-based example, the same class and methods are invoked, with the key difference being the use of a covariance matrix derived from the ellipse parameters. These functions wrap globally in longitude. The longest ranges are comparable to $\pi R_{\oplus} \approx 20000$km, and are found in the ENSO region.

The fields generated from these interpolation exercises are plotted in Figure 6. Both analysis methods produce a similar result in terms of the large-scale features (c.f. Figure 6 a with b, and d with e). However, differences between these fields range between $\pm 1.5^{\circ}$C (Figure 6c,f), with the largest differences evident in the data-sparse regions of the Southern Ocean and across the south-east Pacific Ocean. This difference demonstrates the uncertainty that arises from the choice of analysis method.

The Python Notebook then generates a simulated field using the **`StochasticKriging`** class, which uses the method described in Section 2.5. The class is initiated by supplying the covariance matrix generated for the ellipse model and an error covariance matrix. The **`solve`** method from the class is then called to produce the perturbed field.

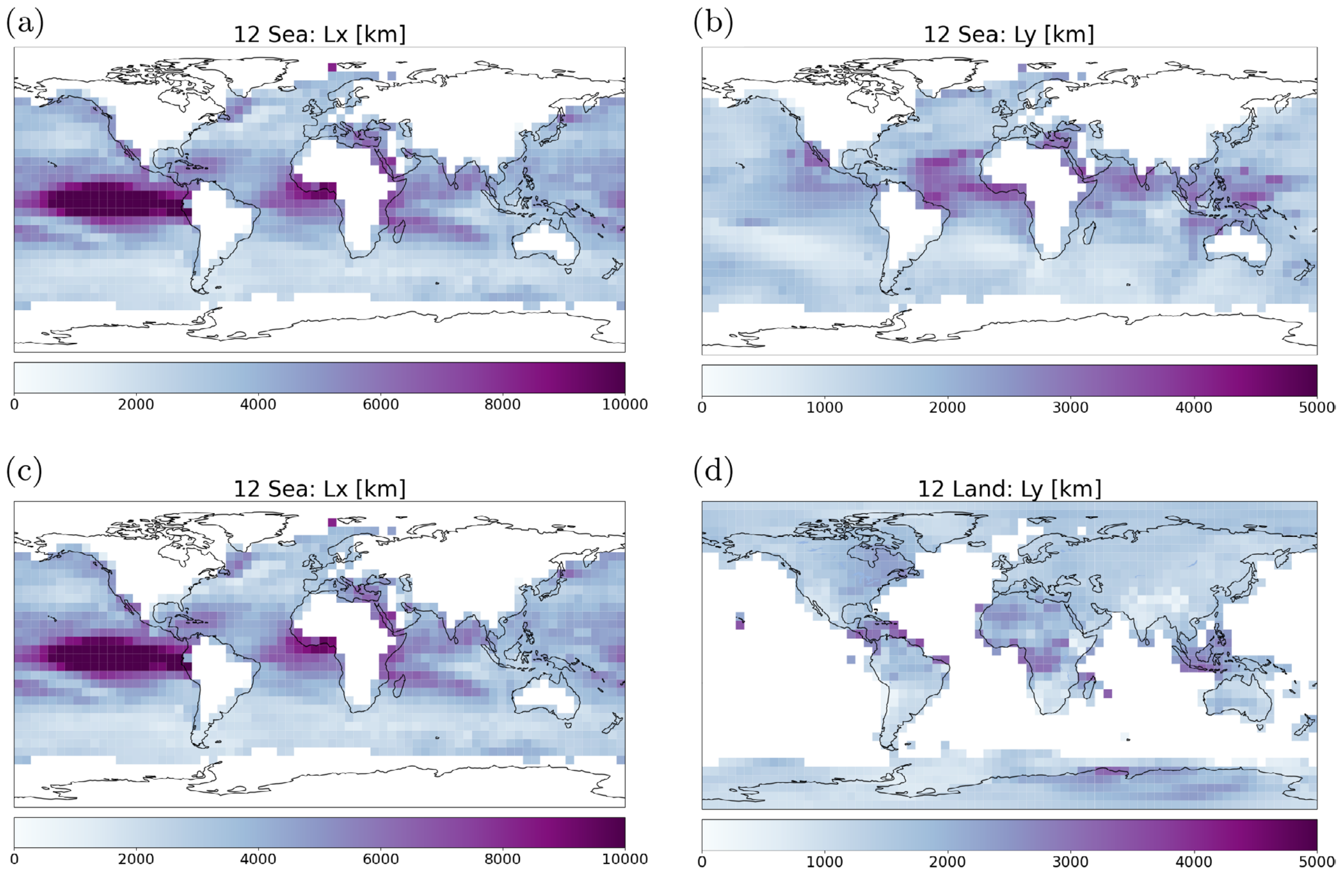

FIGURE 4 | Maps of December L_x and L_y kernel parameters for the sea (a and b) and land (c and d, including sea ice area) domains.

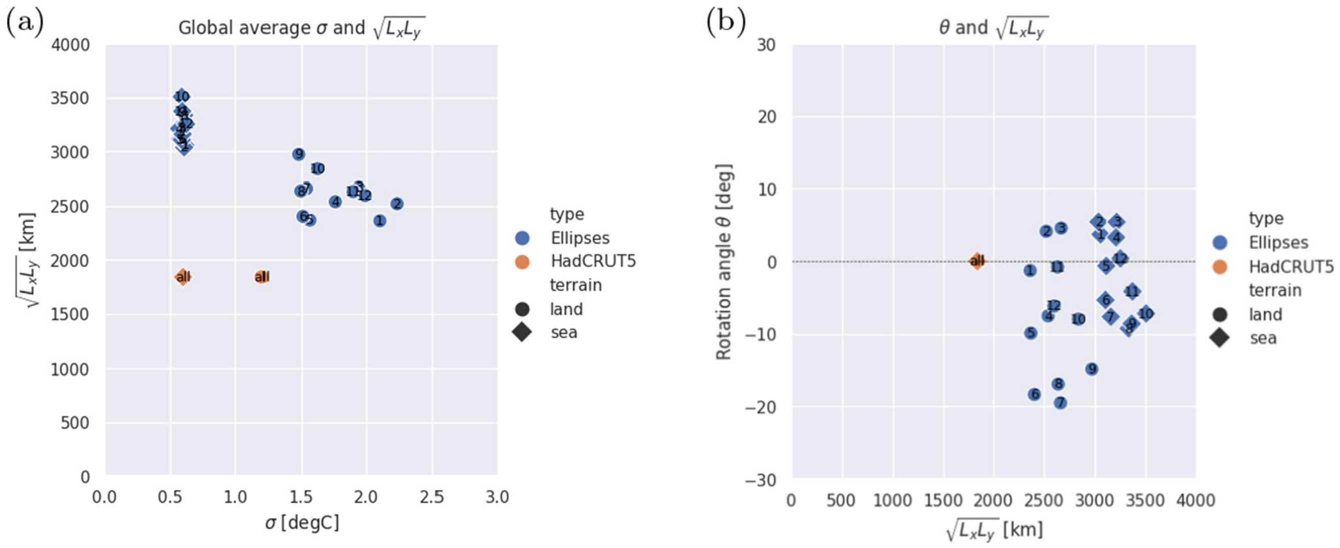

FIGURE 5 | Globally-average spatial scale parameters and standard deviations for SSTs and LSATs, diagnosed from the ellipse parameterization compared to values used in HadCRUT5. Month number is indicated by the number inside each shape. Note that the values used in HadCRUT5 have been converted to equivalent values for the Matérn functions used in the ellipses. The global averages are computed by averaging the kernel within the Matérn covariance function, and then applying an eigenvalue decomposition. To allow comparison with the isotropic values used in HadCRUT5, we take the geometric mean of major and minor axis lengths from our anisotropic analysis.

3.3 | Assessing the Effect of Sparse Sampling on an Interpolated Field

To further illustrate the value of comparing different interpolation methods in `GloMarGridding`, we have produced spatially complete fields of SST for December 1997 using the sampling of SST provided by ship observations in December 1877 (Figure 7). The field of SST values provided by the ESA CCI dataset (Embury et al. 2024) acts as the reference in this exercise (Figure 7a) and shows the large SST anomalies

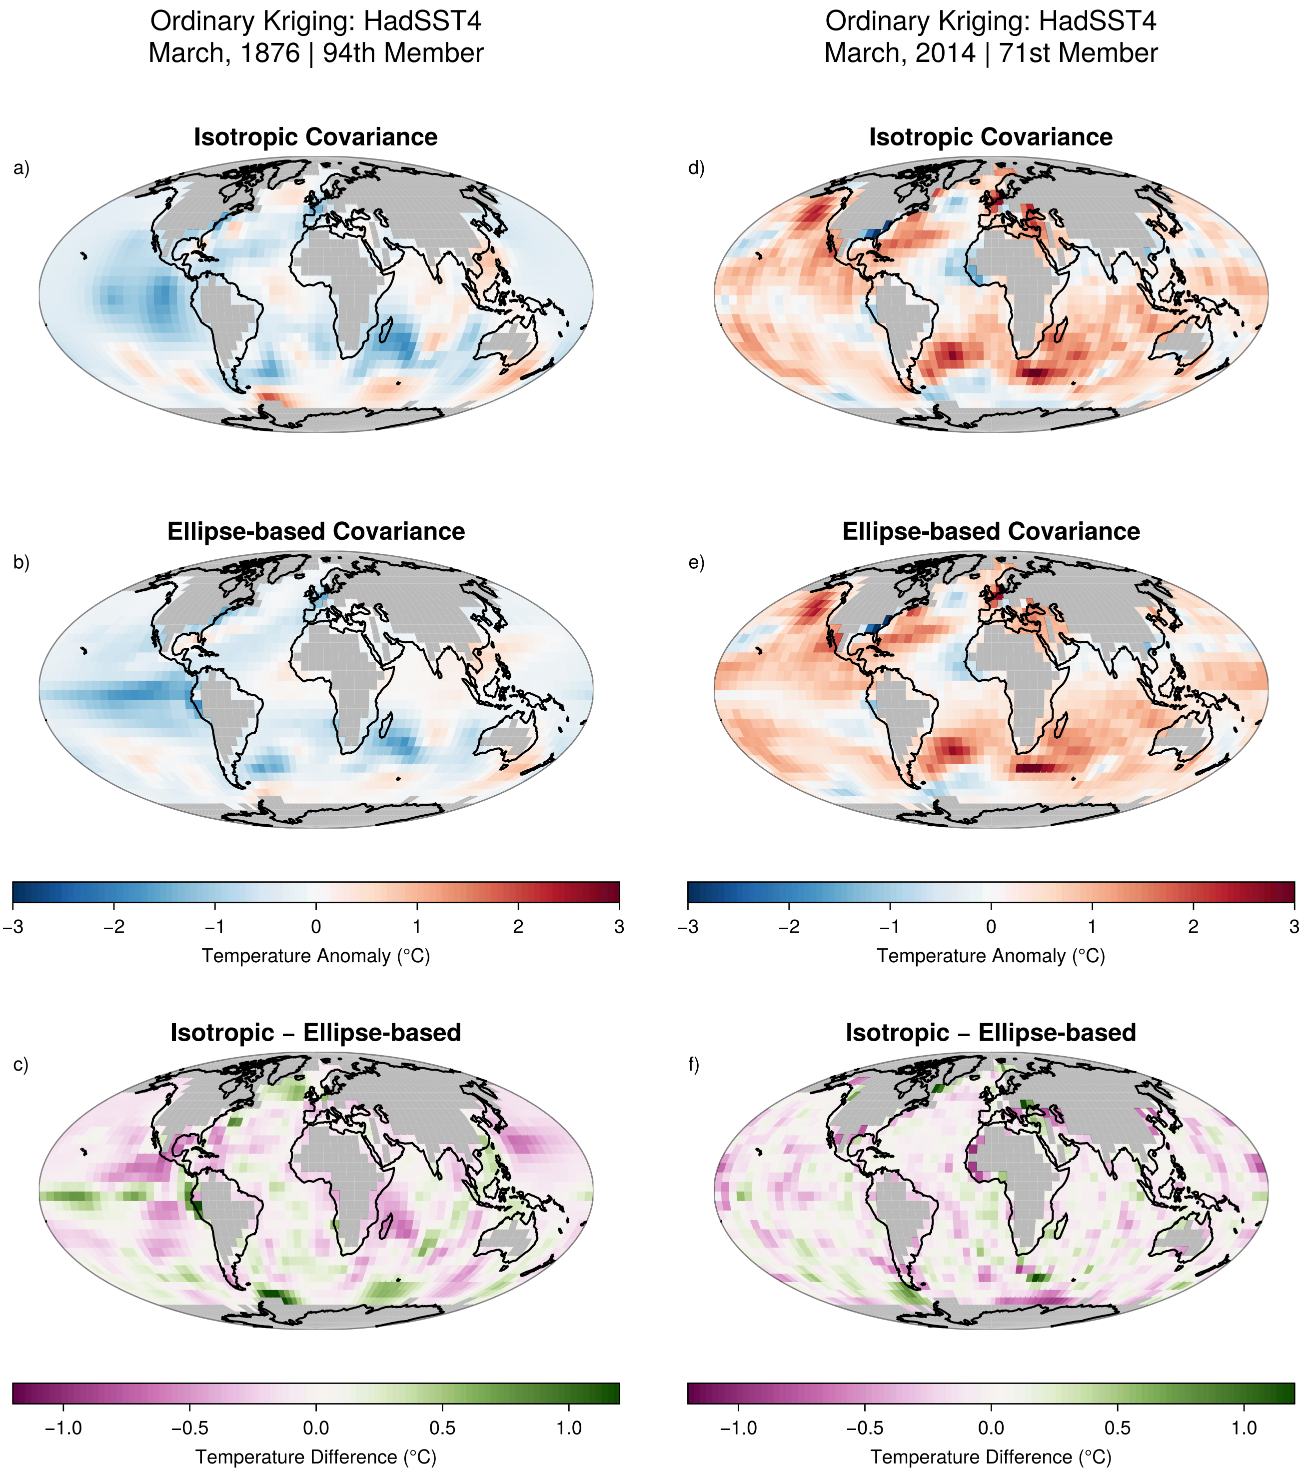

FIGURE 6 | Maps showing HadSST4 data for March 1876 (a–c, using the 94th ensemble member of HadSST4) and March 2014 (d–f, 71st member) interpolated using isotropic (a and d) and ellipse-based (b and e) covariance matrices. The differences (isotropic *minus* ellipse method) are shown in plots c and f.

that occurred in the central Pacific during this significant El Nino event. These data are shown at a 5°x5° resolution, and those values were subsampled to represent the coverage of ship-derived SST data in December 1877—another significant El Nino event—from the DCENT dataset (Chan et al. 2024). Those data were then spatially interpolated using `GloMarGridding` with spatial covariance matrices constructed using a non-stationary (ellipse-based) covariance matrix (Figure 7b) and an isotropic covariance matrix (Figure 7c). The calculation of these matrices is as described in Section 3.2, and the error covariance matrix from the DCENT dataset has been incorporated into the interpolation. The results clearly

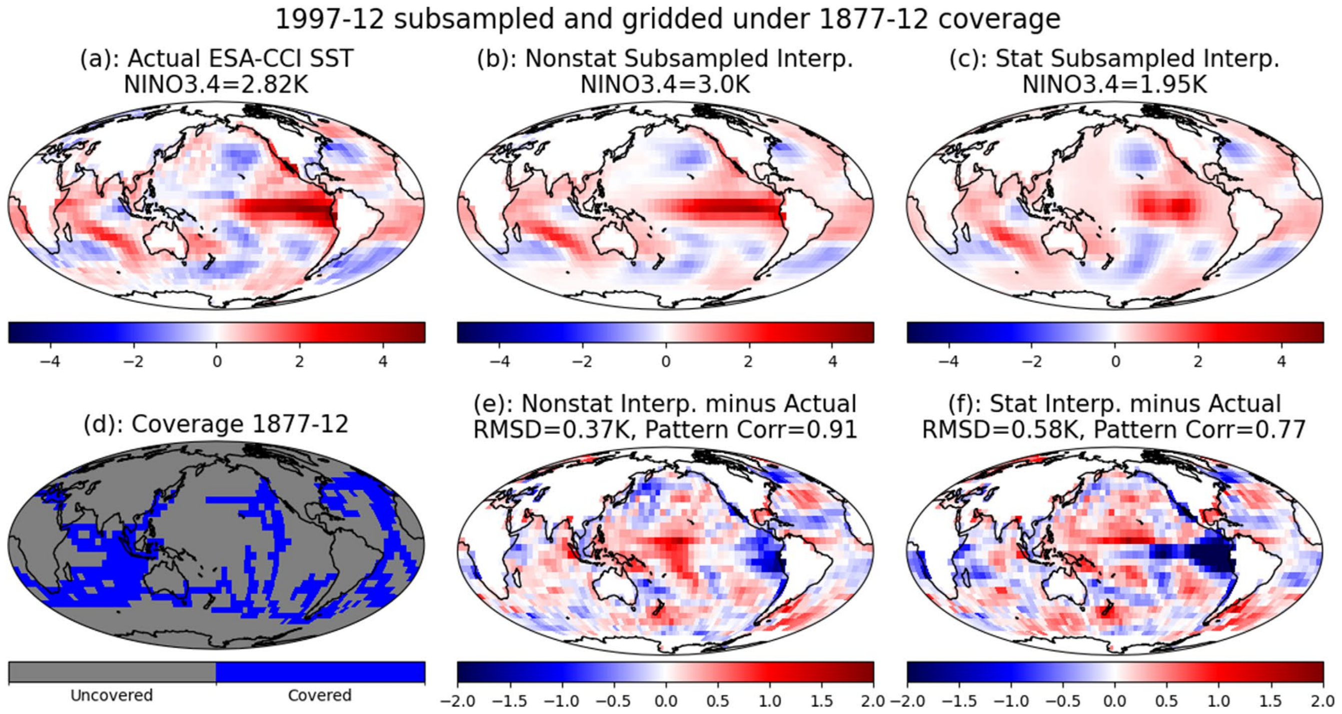

FIGURE 7 | The December 1997 ENSO event as represented by the ESA CCI SST data (a) and spatially interpolated using non-stationary (b) and stationary (c) covariance functions applied to ESA CCI SST data representative of ship data sampling for December 1877. The coverage of data (d) using in maps b and c is as defined by the DCENT dataset (Chan et al. 2024). The differences between the two subsampled interpolations relative to the full field of ESA CCI SST data are shown in (e) and (f). The corresponding NINO3.4 index values for the different grids are shown in a–c. In (e) and (f) the root mean square difference (RMSD) and pattern correlation against the full SST field are shown.

indicate the poor representation of the Pacific SST field when using a stationary, isotropic covariance matrix. This is indicated by the unrealistic roundness of the ENSO anomaly, the reduced Nino3.4 index, and the higher RMSD and lower pattern correlations relative to the ESA CCI data compared to the results from the ellipse-based analysis.

4 | FUTURE DEVELOPMENT OF GloMarGridding

`GloMarGridding` was originally conceived as a framework for analysing surface temperature fields, a focus reflected in this paper. However, the spatial interpolation techniques it implements are, in principle, applicable to other climate variables. For example, variables such as mean sea-level pressure or humidity may be handled effectively with the existing functionality. In contrast, more complex variables—such as wind speed or precipitation—are likely to require additional functions or preprocessing steps. It should be emphasised, however, that the current implementation has only been tested using SST and LSAT data.

It is envisaged that **`GloMarGridding`** will be extended to incorporate additional interpolation methods used in existing global surface temperature datasets. This will enable a more comprehensive evaluation of the contribution of spatial interpolation to total uncertainty. We also welcome collaboration to broaden the package's methodological capabilities or to enhance it in other ways. Kriging-based approaches are particularly amenable to integration into the package. For example, the Berkeley Earth dataset also utilises a kriging-based methodology (Rohde et al. 2013), while datasets such as GISTEMP apply inverse distance weighting for the spatial reconstruction of temperature fields (Lenssen et al. 2019). Empirical Orthogonal Function (EOF)-based analyses, such as Reduced Space Optimal Interpolation (RSOI), are also well-suited for inclusion in **`GloMarGridding`** (Kaplan et al. 1997). However, techniques such as the deep-learning methods demonstrated by Kadow et al. (2020), would present significant implementation challenges. We welcome feedback from users of the software package on its utility and ease of use and issues can be raised via the repository. We expect this open software initiative to be applicable to the generation of complete fields for a wide range of applications, especially in the development of climate data products.

Acknowledgements

The contributions of R.C., A.C., S.C., A.F., E.K. and J.S. were funded by UKRI/NERC via the National Oceanography Centre National Capability Program AtlantiS (grant number NE/Y005589/1). We acknowledge the ESA Climate Change Initiative and the SST CCI project for providing the data used in this paper. The authors acknowledge the ECMWF for producing and maintaining the ERA5 dataset. We are grateful for discussions with Dr. Colin Morice at the Met Office on the implementation of the HadCRUT5 gridding methodology, which has guided some aspects of our approach.

Funding

UKRI/NERC National Capability Program AtlantiS, Grant Number NE/Y005589/1.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The code developed in this paper is available from: https://github.com/NOCSurfaceProcesses/GloMarGridding and https://pypi.org/project/glomar_gridding/.

Endnotes

[1] Note that GloMarGridding does not provide functions to estimate variogram parameters. If required, these values would need to be calculated from another software package.

References

Bun, J., J. P. Bouchaud, and M. Potters. 2017. "Cleaning Large Correlation Matrices: Tools From Random Matrix Theory." *Physics Reports* 666: 1–109. https://doi.org/10.1016/j.physrep.2016.10.005.

Centre for Environmental Data Analysis. 2024. "European Centre for Medium-Range Weather Forecasts (2021): ECMWF ERA5: Surface Level Analysis Parameter Data." https://catalogue.ceda.ac.uk/uuid/c1145ccc4b6d4310a4fc7cce61041b63/.

Chan, D., S. C. Chan, J. T. Siddons, et al. 2026. "DCENT-I: A Globally Infilled Extension of the Dynamically Consistent ENsemble of Temperature Dataset." *Geoscience Data Journal* 13, no. 2: e70054. https://doi.org/10.1002/gdj3.70054.

Chan, D., G. Gebbie, P. Huybers, and E. C. Kent. 2024. "A Dynamically Consistent ENsemble of Temperature at the Earth Surface Since 1850 From the DCENT Dataset." *Scientific Data* 11, no. 1: 953. https://doi.org/10.1038/s41597-024-03742-x.

Chen, H. H., Y. Wang, P. Xiu, Y. Yu, W. Ma, and F. Chai. 2023. "Combined Oceanic and Atmospheric Forcing of the 2013/14 Marine Heatwave in the Northeast Pacific." *npj Climate and Atmospheric Science* 6, no. 1. https://doi.org/10.1038/s41612-023-00327-0.

Cressie, N. 1988. "Spatial Prediction and Ordinary Kriging." *Mathematical Geology* 20, no. 4: 405–421. https://doi.org/10.1007/BF00892986.

Cressie, N. A. C. 1993. *Statistics for Spatial Data*. 1st ed. Wiley.

Dodge, Y. 2008. "Mahalanobis Distance." In *The Concise Encyclopedia of Statistics*, 325–326. Springer New York.

Duchon, C. E. 1979. "Lanczos Filtering in One and Two Dimensions." *Journal of Applied Meteorology* 18, no. 8: 1016–1022. https://doi.org/10.1175/1520-0450(1979)018<1016:LFIOAT>2.0.CO;2.

Embury, O., C. J. Merchant, S. A. Good, et al. 2024. "Satellite-Based Time-Series of Sea-Surface Temperature Since 1980 for Climate Applications." *Scientific Data* 11, no. 1: 326. https://doi.org/10.1038/s41597-024-03147-w.

Gneiting, T. 2013. "Strictly and Non-Strictly Positive Definite Functions on Spheres." *Bernoulli* 19, no. 4. https://doi.org/10.3150/12-BEJSP06.

Guttorp, P., and T. Gneiting. 2006. "Studies in the History of Probability and Statistics Xlix: On the Matérn Correlation Family." *Biometrika* 93, no. 4: 989–995. https://doi.org/10.1093/biomet/93.4.989.

Herdin, M., N. Czink, H. Ozcelik, and E. C. Bonek. 2005. "Matrix Distance, a Meaningful Measure for Evaluation of Non-Stationary MIMO Channels." In *2005 IEEE 61st Vehicular Technology Conference*, vol. 1, 136–140. IEEE.

Hersbach, H., B. Bell, P. Berrisford, et al. 2020. "The ERA5 Global Reanalysis." *Quarterly Journal of the Royal Meteorological Society* 146, no. 730: 1999–2049. https://doi.org/10.1002/qj.3803.

Higham, N. J., N. Strabić, and V. Šego. 2016. "Restoring Definiteness via Shrinking, With an Application to Correlation Matrices With a Fixed Block." *SIAM Review* 58, no. 2: 245–263.

Hoffimann, J. 2018. "Geostats.Jl—High-Performance Geostatistics in Julia." *Journal of Open Source Software* 3, no. 24: 692. https://doi.org/10.21105/joss.00692.

Huang, B., P. W. Thorne, T. M. Smith, et al. 2016. "Further Exploring and Quantifying Uncertainties for Extended Reconstructed Sea Surface Temperature (Ersst) Version 4 (v4)." *Journal of Climate* 29, no. 9: 3119–3142. https://doi.org/10.1175/jcli-d-15-0430.1.

Huang, B., X. Yin, M. J. Menne, R. Vose, and H. M. Zhang. 2022. "Improvements to the Land Surface Air Temperature Reconstruction in NOAAGlobalTemp: An Artificial Neural Network Approach." *Artificial Intelligence for the Earth Systems* 1, no. 4: e220032. https://doi.org/10.1175/AIES-D-22-0032.1.

Huang, C., D. Farewell, and J. Pan. 2017. "A Calibration Method for Non-Positive Definite Covariance Matrix in Multivariate Data Analysis." *Journal of Multivariate Analysis* 157: 45–52. https://doi.org/10.1016/j.jmva.2017.03.001.

IPCC. 2021. *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. 1st ed. Cambridge University Press.

Jolliffe, I. 2005. "Principal Component Analysis." In *Encyclopedia of Statistics in Behavioral Science*, edited by B. S. Everitt and D. C. Howell, 1st ed. Wiley.

Kadow, C., D. M. Hall, and U. Ulbrich. 2020. "Artificial Intelligence Reconstructs Missing Climate Information." *Nature Geoscience* 13, no. 6: 408–413. https://doi.org/10.1038/s41561-020-0582-5.

Kaplan, A., M. A. Cane, Y. Kushnir, A. C. Clement, M. B. Blumenthal, and B. Rajagopalan. 1997. "Analysis of Global Sea Surface Temperature 1856–1991." *Journal of Geophysical Research: Oceans* 102, no. C13: 27835–27860. https://doi.org/10.1029/97JC01736.

Karspeck, A. R., A. Kaplan, and S. R. Sain. 2012. "Bayesian Modelling and Ensemble Reconstruction of Mid-Scale Spatial Variability in North Atlantic Sea-Surface Temperatures for 1850-2008." *Quarterly Journal of the Royal Meteorological Society* 138, no. 662: 234–248. https://doi.org/10.1002/qj.900.

Kennedy, J. J. 2014. "A Review of Uncertainty in In Situ Measurements and Data Sets of Sea Surface Temperature." *Reviews of Geophysics* 52, no. 1: 1–32. https://doi.org/10.1002/2013RG000434.

Kennedy, J. J., N. A. Rayner, C. P. Atkinson, and R. E. Killick. 2019. "An Ensemble Data Set of Sea Surface Temperature Change From 1850: The Met Office Hadley Centre HadSST.4.0.0.0 Data Set." *Journal of Geophysical Research: Atmospheres* 124, no. 14: 7719–7763. https://doi.org/10.1029/2018JD029867.

Laloux, L., P. Cizeau, M. Potters, and J. P. Bouchaud. 2000. "Random Matrix Theory and Financial Correlations." *International Journal of Theoretical and Applied Finance* 3, no. 3: 391–397. https://doi.org/10.1142/S0219024900000255.

Lamarche, C., M. Santoro, S. Bontemps, et al. 2017. "Compilation and Validation of SAR and Optical Data Products for a Complete and Global Map of Inland/Ocean Water Tailored to the Climate Modeling Community." *Remote Sensing* 9, no. 1: 36. https://doi.org/10.3390/rs9010036.

Lenssen, N. J. L., G. A. Schmidt, J. E. Hansen, et al. 2019. "Improvements in the GISTEMP Uncertainty Model." *Journal of Geophysical Research: Atmospheres* 124, no. 12: 6307–6326. https://doi.org/10.1029/2018JD029522.

Morice, C. P., J. J. Kennedy, N. A. Rayner, et al. 2021. "An Updated Assessment of Near-Surface Temperature Change From 1850: The

HadCRUT5 Data Set." *Journal of Geophysical Research: Atmospheres* 126, no. 3: e2019JD032361. https://doi.org/10.1029/2019JD032361.

Nelder, J. A., and R. Mead. 1965. "A Simplex Method for Function Minimization." *Computer Journal* 7, no. 4: 308–313.

North, G. R., T. L. Bell, R. F. Cahalan, and F. J. Moeng. 1982. "Sampling Errors in the Estimation of Empirical Orthogonal Functions." *Monthly Weather Review* 110, no. 7: 699–706.

Nychka, D., R. Furrer, J. Paige, and S. Sain. 2021. "fields: Tools for Spatial Data. R Package Version".

Paciorek, C. J., and M. J. Schervish. 2006. "Spatial Modelling Using a New Class of Nonstationary Covariance Functions." *Environmetrics* 17, no. 5: 483–506. https://doi.org/10.1002/env.785.

Rasmussen, C. E., and C. K. I. Williams. 2005. *Gaussian Processes for Machine Learning*. MIT Press.

Rohde, R., R. Muller, R. Jacobsen, et al. 2013. "A New Estimate of the Average Earth Surface Land Temperature Spanning 1753 to 2011." *Geoinformatics & Geostatistics: An Overview* 1, no. 1.

Rohde, R. A., and Z. Hausfather. 2020. "The Berkeley Earth Land/Ocean Temperature Record." *Earth System Science Data* 12, no. 4: 3469–3479. https://doi.org/10.5194/essd-12-3469-2020.

Schulzweida, U. 2023. "CDO User Guide (2.3.0)." *Zenodo*. https://doi.org/10.5281/zenodo.10020800.

Singh, D., R. Seager, B. I. Cook, et al. 2018. "Climate and the Global Famine of 1876–78." *Journal of Climate* 31, no. 23: 9445–9467. https://doi.org/10.1175/JCLI-D-18-0159.1.

Sippel, S., E. C. Kent, N. Meinshausen, et al. 2024. "Early-Twentieth-Century Cold Bias in Ocean Surface Temperature Observations." *Nature* 635, no. 8039: 618–624. https://doi.org/10.1038/s41586-024-08230-1.

SurfTemp.net. 2026. "Sea Surface Temperature Data—Re-Gridding Service." https://surftemp.net/regridding/.

Thorne, P. W., D. E. Parker, J. R. Christy, and C. A. Mears. 2005. "Uncertainties in Climate Trends: Lessons From Upper-Air Temperature Records." *Bulletin of the American Meteorological Society* 86, no. 10: 1437–1442. https://doi.org/10.1175/BAMS-86-10-1437.

Titchner, H. A., and N. A. Rayner. 2014. "The Met Office Hadley Centre Sea Ice and Sea Surface Temperature Data Set, Version 2: 1. Sea Ice Concentrations: HadISST.2.1.0.0 Sea Ice Concentrations." *Journal of Geophysical Research: Atmospheres* 119, no. 6: 2864–2889. https://doi.org/10.1002/2013jd020316.

Wilks, D. 2006. "Statistical Methods in the Atmospheric Sciences." In *National Geophysics*, 2nd ed. Elsevier.

Williams, S. D. P., and D. I. Berry. 2020. "ACSIS Atlantic Ocean Medium Resolution SST Dataset: Reconstructed 5-Day, ½-Degree, Atlantic Ocean SST (1950–2014)." *Geoscience Data Journal* 7, no. 2: 135–148. https://doi.org/10.1002/gdj3.94.

Appendix A

TABLE A1 | List of symbols and variable names used in this paper.

Symbol or variable name	Meaning
$\|\|\|\|_f$	Frobenius norm
$\mathbf{0}$	Block matrix of 0 s; as Lagrange multiplier, a single 0
$\mathbf{1}$	Column vector of 1 s
$\mathbf{C}$	Global spatial covariance
$\widehat{\boldsymbol{C}}$	Matérn covariance function
$\mathbf{C}_{\overline{g}}$	Covariance of Kriging estimates
$\mathbf{C}_{\text{cross}}$	Covariance between observed and prediction points, subsampled from $\mathbf{C}$
$\mathbf{C}_{\text{obs}}$	Covariance between observed points, subsampled from $\mathbf{C}$
det	Determinant (of matrix)
diag	Take diagonal of a matrix or make a vector into a diagonal matrix
$\mathbf{E}$	Error covariance
$\overline{\mathbf{g}}$	Kriging predictions
$\hat{\imath}, \hat{\jmath}$	Zonal and meridional unit vectors
K_ν	Modified Bessel function of the second kind of order ν
Lon, Lat	Longitude and latitude
MVN	Multi-variate normal distribution
Q	Automatic eigenvalue threshold for the repairing of correlation matrices
$R_\oplus$	Radius of the Earth
$\mathbf{S}$	Any covariance matrix including $\mathbf{C}$
sin, cos	Trigonometric sine and cosine
Subscript SK, OK	Simple and ordinary Kriging respectively
Superscript T	Transpose operator
Tr	Trace (of matrix)
V	(Right) eigenvectors
x_i	An arbitrary point (labelled i) on Earth surface
$x_i - x_j$	The displacement vector between points i and j
y	Vector of observations
α	Explained variance in the interpolated field
$\overline{\beta}$	Unknown global mean in ordinary Kriging
Γ	Gamma function
λ (or $\mathbf{\Lambda}$)	Eigenvalues (or vector of eigenvalues)
μ	Known global mean in simple Kriging
ν	Matérn covariance function shape parameter
$\Sigma\left(L_x, L_y, \theta\right)$	Kernel within $\widehat{\boldsymbol{C}}$ with length scales L_x, L_y and rotation angle θ
σ	Standard deviation (square root of the variance)
τ	Mahalanobis distance