

A FUNCTIONAL-DATA APPROACH TO THE ARGO DATA

BY DREW YARGER^a, STILIAN STOEV^b AND TAILEN HSING^c

Department of Statistics, University of Michigan, ^adyarger@umich.edu, ^bssstoev@umich.edu, ^cthhsing@umich.edu

The Argo data is a modern oceanography dataset that provides unprecedented global coverage of temperature and salinity measurements in the upper 2000 meters of depth of the ocean. We study the Argo data from the perspective of functional data analysis (FDA). We develop spatiotemporal *functional kriging* methodology for mean and covariance estimation to predict temperature and salinity at a fixed location as a smooth function of depth. By combining tools from FDA and spatial statistics, including smoothing splines, local regression, and multivariate spatial modeling and prediction, our approach provides advantages over current methodology that consider pointwise estimation at fixed depths. Our approach naturally leverages the irregularly-sampled data in space, time, and depth to fit a space-time functional model for temperature and salinity. The developed framework provides new tools to address fundamental scientific problems involving the entire upper water column of the oceans, such as the estimation of ocean heat content, stratification, and thermohaline oscillation. For example, we show that our functional approach yields more accurate ocean heat content estimates than ones based on discrete integral approximations in pressure. Further, using the derivative function estimates, we obtain a new product of a global map of the mixed layer depth, a key component in the study of heat absorption and nutrient circulation in the oceans. The derivative estimates also reveal evidence for density inversions in areas distinguished by mixing of particularly different water masses.

1. Introduction. The development of technology has vastly increased the amount and complexity of data available that monitor the Earth’s environment. We focus on one type of such data collected by the Argo project, an international collaboration that oversees more than 3800 devices, called floats, which measure the temperature and salinity of the oceans. Each float periodically ascends from two kilometers deep while collecting temperature and salinity measurements as a function of pressure, a proxy for depth, with one decibar (dbar) roughly corresponding to one meter of depth. These data, referred to as profiles, are transmitted over satellite to data processing centers along with the float’s coordinates and time stamps. The drifting floats collect approximately 100,000 profiles each year, resulting in a large and complex space-time dataset, indexed by longitude, latitude, time, and pressure; see [Argo \(2000\)](#) for more information.

The global coverage of the Argo data and the depth of measurements provide previously unavailable richness of oceanography data (see Figure 1). The data have begun to play a critical part in measuring sea level rise, currents, and the global distribution of temperature and salinity of the oceans. The oceans play a major role in the Earth’s climate; for example, [Roemmich et al. \(2015\)](#) uses Argo data to study the warming oceans which account for more than 90% of the net planetary energy increase. More than 1500 papers that use Argo data have been published in the past five years; recently, the Argo data have begun to see research in the statistics community. For example, [Kuusela and Stein \(2018\)](#) is the first such publication which enumerates some directions for future statistical research for the Argo data. To the best of our knowledge, none of the papers in this sizeable literature so far has fully taken into

account the dependence of the Argo data across location, time, and pressure. For instance, the inference of the spatial dependence of temperature and salinity has thus far been conducted using a pressure-level by pressure-level approach.

We consider this problem of temperature and salinity estimation, using data from all values of pressure simultaneously, under the framework of functional data analysis (FDA). The problem of spatial inference for functional data has only recently been addressed; see, for instance, Baladandayuthapani et al. (2008), Gromenko, Kokoszka and Sojka (2017), Zhang et al. (2016), Zhang and Li (2002), Zhou et al. (2010). A more thorough discussion of this area will be given in Section 2.3. One aspect of such inference is “functional kriging,” where the goal is to predict a function-valued variable at an unobserved location based on spatially correlated function-valued covariates. Here, we develop a functional kriging methodology, tailored to the challenges and complexities of the Argo data, and aimed at producing maps or spatiotemporal predictions of temperature and salinity as functions of pressure along with functional uncertainties. In the context of the Argo data, each profile can be considered functional data, with measurements observed as a function of pressure for a fixed time and location. In this framework we use nearby profiles in space and time to estimate temperature and salinity between the profile locations. This is done by using functional models for the mean and space-time covariance structure which also yields uncertainties and confidence sets for the functional kriging estimates. The FDA approach provides computational, scientific, and methodological advantages over current approaches that consider models for one pressure level at a time by linearly interpolating temperature and salinity onto that pressure (Kuusela and Stein (2018), Roemmich and Gilson (2009)). First, the FDA approach provides a principled way to share information in the irregularly-sampled measurements across pressure without perturbations (e.g., by linear interpolation). Second, the estimated functions capture the complex thermohaline structure in the oceans as a function of pressure that arises from the oceans’ stratification and mixing. The FDA approach also naturally yields estimates of derivatives and integrals over the entire pressure dimension which can provide new insight into key scientific problems.

We directly compare our FDA approach with current ones that, first, linearly interpolate each profile onto fixed pressure levels. While such an interpolation simplifies the data for the subsequent modeling compared to irregularly sampled pressures, it also introduces error or neglects data depending on whether the profiles observed are *sparse* or *dense* in pressure. Since Argo profiles typically range in number of observations from around 60 to 1000 measurements, the Argo data present a combination of such heterogeneous data. When sparse functional data are observed, that is, there are just a few measurements per profile, interpolating or presmoothing each curve can decrease accuracy in comparison to pooling data from profiles (Hall, Müller and Wang (2006), Li and Hsing (2010)). When dense functional data are available, only some observations are used to interpolate onto pressure levels, and the smaller features of the temperature and salinity in the pressure dimension will be undetected. The FDA approach both avoids the interpolation error for the sparsely-observed profiles and leverages all measurements from each profile, and thus it describes the pressure dimension in more intricate detail. Furthermore, when predicting at a large number of pressure levels (e.g., the 58 pressure levels or more in Roemmich and Gilson (2009)), the functional approach can considerably reduce computations by sharing information across pressure and providing functional predictions. Perhaps most notably, estimating at fixed levels limits one’s ability to predict derivative and integral functionals of the temperature and salinity, since these must be approximated from discrete predictions. On the other hand, derivatives and integral estimates, along with their uncertainties, are readily available in our functional kriging approach and can be leveraged for fundamental scientific problems, like the estimation of ocean heat content and mixed layer depth (see Section 5 below).

We first introduce our notation for the data and our model:

- *The data:* Denote the data for the i th profile as $s_i, d_i, y_i, (p_{i,j}, Y_{i,j})_{j=1}^{m_i}$ for $i = 1, \dots, n$, where j indexes the measurement, $s_i = (s_{i1}, s_{i2})$ is its location, (d_i, y_i) is its day of year and year, respectively, and $(p_{i,j}, Y_{i,j})_{j=1}^{m_i}$ is the pressure and response measurements. Here, $Y_{i,j}$ denotes temperature or salinity, depending on the context; in actuality, both are observed for each i and j . In this analysis, different floats are treated identically, and the various float characteristics are not used. Data can be viewed using an R Shiny Application (Yarger (2020a)).
- *The model:* We assume that

$$(1) \quad Y_{i,j} = \mu(s_i, d_i, y_i, p_{i,j}) + X(s_i, d_i, y_i, p_{i,j}) + \epsilon_{i,j},$$

where μ is a fixed mean function, X is a zero-mean stochastic process that captures the dependence of the data, and $\epsilon_{i,j}$ is measurement error. We assume that the distribution of $X(\cdot, \cdot, y, \cdot)$ is the same for all y and that X is weakly dependent in time so that $X(\cdot, d_i, y_i, \cdot)$ and $X(\cdot, d_j, y_j, \cdot)$ are independent for d_i near d_j and $y_i \neq y_j$. The $\epsilon_{i,j}$ are assumed to form a white noise process in space, time, and pressure, with mean zero and variance parameterized by $\kappa(s, d, p)$.

Our new approach to the estimation of the functional mean μ combines two established approaches in nonparametric statistics: smoothing splines and local polynomial regression; see Green and Silverman (1994) and Fan and Gijbels (1996) for more information on these methodologies, respectively. Specifically, we leverage irregularly sampled data in space and time using local regression to form a spline estimate of the mean function of pressure. This approach can model the strong vertical stratification in the oceans where water masses at different depths can have drastically different characteristics. Our mean estimation reflects the advantages of both of these approaches: computations are reduced by using univariate B-splines along the pressure dimension while the nonlinear features of the oceans in space and time are estimated in a statistically efficient manner by local polynomial regression. As a byproduct, our approach extends that of Fan and Gijbels (1996) to the case of function-valued data and provides new functional estimates of derivatives of the mean with respect to space and time.

After subtracting the functional mean, we model the covariance structure of the residuals in space, time, and pressure. First, we estimate the covariance between measurements in the same profile, decompose this estimate to form functional principal components (FPCs), and use the first K FPCs to estimate a space-time covariance structure. As in Kuusela and Stein (2018), locally-estimated space-time covariance models are used to perform kriging and to obtain the uncertainty in the estimates. This entails a unified and computationally tractable functional modeling and prediction framework that takes into account the dependence in space, time, and pressure.

Being able to fill data gaps in a principled manner to produce estimates of temperature and salinity continuously at all locations, times, and pressures is of tremendous value to ocean research. Some examples and references of traditional “mapping strategies,” or interpolation approaches, in oceanography can be found in Boyer and Levitus (1994) and Ishii and Kimoto (2009). As shown in Table 1 of Cheng and Zhu (2014), the resolution of available ocean data continues to improve, the Argo project being a contributing factor in the past 10 years. Our functional data approach is motivated by fully leveraging the benefit of the high-resolution Argo data, but generalizations to other similar high-resolution data should be straightforward. These approaches will potentially play an important role in ocean and climate research in general.

We mention some applications to demonstrate the advantages of our approach. The first application is the estimation of the integrated ocean heat content at each location which is related to the integral of the temperature curves. Traditionally, the integrated ocean heat content was studied through numerically interpolated data. The main focus of [Cheng and Zhu \(2014\)](#) is to address the bias in such estimates caused by data sparsity. In our opinion, however, this study and other traditional approaches do not always comprehensively consider the variability throughout the analysis. Using our approach, distributional properties of the estimated heat content can be easily obtained from the overall analysis. In particular, if the heat content at a location is estimated with sparse data, then the model-based estimated error will reflect that. More generally, by pooling data across space and time, our model-based approach can be used to identify statistically significant anomalies in the ocean heat content which is one important and active area of research ([Roemmich et al. \(2015\)](#)).

For the second application the functional predictions are used to estimate potential density, which provides valuable information about the vertical stratification of the oceans, a key factor in their ability to absorb heat ([Li et al. \(2020\)](#)). For example, we use potential density to estimate the depth of the *mixed layer* of the ocean, a region directly below the ocean surface where the ocean mixes uniformly and is characterized by near-constant ocean properties (Sections 4.2 and 7.4 of [Talley et al. \(2011\)](#)). The mixed layer drives the ocean-atmosphere interactions and thus influences heat and carbon flux of the ocean, ocean circulation, and biological processes dependent on light ([Holte et al. \(2017\)](#)). Our functional estimates provide mixed layer depth estimates over all open oceans for each day-year combination that have minimal discretization error in pressure. Employing the bootstrap, we obtain distributions of the mixed layer depths and use them to assess the within and between-year variations in mixed layer depth. Our analysis shows, in particular, that summer mixed layer depths generally have smaller within-year variations and, at some locations, slightly larger between-year variations than do winter mixed layer depths in proportion to their size. We also use the potential density estimates to evaluate the occurrence of nonmonotone features of potential density that indicate vertical instability in the water column, and we find evidence of such features. Our framework provides the means to identify such anomalies on a global scale, which can help oceanographers track the structural stability of the thermohaline oscillation, a fundamental driver of Earth's climate ([Li et al. \(2020\)](#), [Rahmstorf et al. \(2015\)](#)).

We outline the rest of the paper which loosely follows the structure of the [Introduction](#). In [Section 2](#) the Argo data are introduced in more detail. In [Section 3](#) we develop our approach for mean estimation and its computational implementation over the Argo data. After subtracting the mean from the data, we estimate the covariance of the residuals, predict using the estimated covariance, and assess the quality of our predictions in [Section 4](#). In [Section 5](#) a framework to use the functional estimates is developed, specifically applied to the examples outlined above. Throughout our analysis we provide the resulting estimates as data products to the community and introduce interactive R Shiny web applications for visualizing the results ([Yarger \(2020a\)](#)). We conclude and identify future research directions in [Section 6](#). Throughout the paper we refer to figures, tables, and text from the Supplementary Material.

2. Argo data and existing methodologies. In this section we give a more detailed overview of the Argo data, give an introduction to mapping methods, and situate our approach within the spatial FDA literature. While there is a variety of measurements of the oceans, including sea surface temperature and ship-based measurements, we only use data from the Argo project because it provides a natural comparison to existing approaches.

2.1. Data from the Argo project. The Argo program is an international collaboration that develops and manages floats, mechanical devices that collect measurements on the world's




FIG. 1. Argo data examples. (Left) Locations of profiles collected in February 2016, colored by the temperature of the measurement closest to the surface. (Right) Histogram of the number of measurements per profile.

oceans (Argo (2000)). The Argo project reached a goal of global coverage in late 2007 and has continued to increase the number of floats to nearly 4000 today; see, for example, Figures 1 and 2 or a Shiny application (Yarger (2020a)). In 10-day cycles, each float descends from its parking depth at 1000 dbar to a depth of 2000 dbar, then rises over the course of six hours to the surface, collecting measurements of pressure, temperature, and salinity. Upon surfacing, the float transmits the data via satellite. The pressure, temperature, and salinity data and its associated location and time for each cycle is called a profile.

The Argo program was designed to sample approximately one profile every 10 days in each 3 by 3 degree region of the open oceans. Before this relatively uniform sampling of the Argo program, sampling at greater depths was sparse and highly nonuniform in space and time, with fewer measurements in the Southern Hemisphere and during winter months (Roemmich and Gilson (2009)). In terms of depth, the pressures at which each float samples can vary from float to float as well as from profile to profile, due to varying data transmission technology, as seen in Figure 1. This heterogeneity in the sampling frequency is one important challenge addressed by our functional approach.

The Argo data is made publicly available after transmission through satellite and various data-quality control measures. For our analysis we use a preprocessed version of the Argo data which was formed and used in Kuusela and Stein (2018). The data span the years 2007 to 2016 based on the May 2017 snapshot of the Argo data. The data include more than 245 million total point measurements from 994,709 profiles, of which 551,536 have extended data quality (delayed-mode) checks. Throughout our analysis we, generally, use all profiles for temperature, while for salinity delayed-mode profiles are needed to ensure minimal drift or bias (Owens and Wong (2009)).

2.2. Argo mapping methodology. The problem of mapping irregularly-sampled spatial data onto a grid or unobserved location is a common problem in spatial statistics, oceanography, and the geosciences in general. The main methods to address this problem are similar in the different fields, though they may be referred to with different names. In statistics it is often called kriging or Gauss–Markov prediction, specifically referring to the conditional prediction of a Gaussian random vector based on a spatial covariance structure (cf. Cressie and Wikle (2015, Section 4.1)). In geology this method is also referred to as kriging (cf. Chilès and Delfiner (2012, Chapter 3)), while in oceanography this is usually called objective mapping or optimal interpolation (cf. Section 4.2 Thomson and Emery (2014), Barth et al. (2008)), and it primarily focuses on constructing gridded predictions. Each, in essence, involves specifying a mean and covariance structure, then using these to form a prediction. If the true mean and covariance structure is specified, then the resulting prediction minimizes mean-squared error over the class of linear predictors. These approaches generally require

inversion of the covariance matrix of size $n \times n$, where n is the number of observed spatial locations. In optimal interpolation the covariance structure is more often specified using subject-matter knowledge, rather than being estimated from the data.

In this framework we review approaches for mapping that specifically use the Argo data. We focus on the important work of Roemmich and Gilson (2009), who provide a methodology for mean estimation and analysis of anomalies, using the Argo data, as well as Kuusela and Stein (2018), who focus on covariance estimation and introduce maximum likelihood estimation for its model parameters in space and time. These are only two works in a wider array of temperature and salinity estimation works using Argo data. Other approaches used to form Argo data products include Gray and Riser (2015), who propose an iterative approach to estimating the covariance function, Li et al. (2017), Gaillard (2012), Hosoda, Ohira and Nakamura (2008), and Udaya Bhaskar, Ravichandran and Devender (2007). These focus on scalar data at a limited number of pressure levels, and each uses a Gaussian or exponential covariance function. We now turn to the Roemmich and Gilson product which is available as the standard in global oceanography analysis using the Argo data. This product provides estimates of the mean temperature and salinity separately as well as monthly anomalies from the mean over grids of different resolutions in space and fixed pressure levels. Before estimation the temperature and salinity for each profile is interpolated onto 58 fixed nonuniformly-spaced pressure levels. Throughout, they use a distance based on latitude, longitude, and the depth of the ocean floor at each location. The inclusion of the depth of the ocean floor better handles areas where ocean currents run along the shores of continents, like the western boundary currents (see, e.g., Section 7.8 of Talley et al. (2011)). To estimate the mean for each pressure level and grid point of space, they combine data from the years 2004–2016, using the 100 nearest profiles from each of the 12 months of the year. In addition, they only use the interpolated values at a pressure level as well as the two adjacent pressure levels. A weighted least squares approach based on distance from the grid point is used to fit a model of the form

$$(2) \quad \begin{aligned} & \beta_0 + \beta_1(s_{i1} - s_{01}) + \beta_2(s_{i2} - s_{02}) + \beta_3(s_{i1} - s_{01})^2 + \beta_4(s_{i2} - s_{02})^2 \\ & + \beta_5(p_i - p_0) + \beta_6(p_i - p_0)^2 + \sum_{k=1}^6 \gamma_k \sin\left(\frac{d_i 2\pi k}{365.25}\right) + \sum_{k=1}^6 \delta_k \cos\left(\frac{d_i 2\pi k}{365.25}\right), \end{aligned}$$

where s_{i1} and s_{i2} give the location of profile i , p_i the pressure level of the interpolated measurement, d_i is the day of the year profile i was observed, and β_k , γ_k , and δ_k are scalar coefficients. The coefficient β_0 represents the time-averaged mean, while the γ_k , and δ_k give the deviations from this mean at different times of the year. Overall, this approach is a form of local regression, where the time dimension is estimated using a fixed Fourier basis.

After subtracting the mean, Roemmich and Gilson then provide a field of anomalies for each month of each year that describes the variation away from the mean at a particular location. These are formed by computing the conditional mean at each grid point in space and pressure level, assuming Gaussianity and using a covariance of the form

$$(3) \quad C_{RG}(\Delta_{RG}) = 0.77 \cdot \exp(-(a^\top \Delta_{RG}/140)^2) + 0.23 \cdot \exp(-|a^\top \Delta_{RG}|/1111).$$

Here, $\Delta_{RG} = (\Delta_{s1}, \Delta_{s2}, \Delta_{dep})^\top$ denotes a vector of distances between two locations s and s' for the zonal direction (east-west), meridional direction (north-south), and the distance penalty for ocean depth described above. The vector a scales the relative directions and is $(1, 1, 1)$ above 20 degrees north and below 20 degrees south but changes linearly to $(0.25, 1, 1)$ at the equator which increases the covariance in the zonal direction in the tropics. This choice is supported by empirical estimates near the surface. The covariance in (3) is nonstationary, due to its dependence on a , though the covariance does not depend on time

or the pressure level. To form the final product, the anomalies over all months and years are averaged and added to the mean.

Kuusela and Stein (2018) employ the Roemmich and Gilson mean and study the covariance structure in more detail by proposing a space-time covariance model and fitting it using maximum likelihood. To address the nonstationarity of the data, they use the *locally stationary* assumption; that is, at each location, parameters of a stationary covariance are estimated using data nearby, and the local covariance estimates are used for prediction at that location. Data from different years are assumed independent, and one stationary covariance function for data observed in the same year they consider is

$$C_{KS}(\Delta_{KS}) = \phi \cdot \exp\left(-\sqrt{\frac{\Delta_{s_1}^2}{\theta_{s_1}^2} + \frac{\Delta_{s_2}^2}{\theta_{s_2}^2} + \frac{\Delta_d^2}{\theta_d^2}}\right) + \sigma^2 \cdot 1(\Delta_{KS} = 0),$$

where $\Delta_{KS} = (\Delta_{s_1}, \Delta_{s_2}, \Delta_d)^\top$ is the relevant distance between two locations and times in longitude, latitude, and day of the year, respectively. The estimated parameters are the process variance ϕ , nugget variance σ^2 , and three scale parameters θ subscripted by their direction. Thus, since the model is estimated at each pressure level, it can adapt to the large differences in the covariance structure at different depths. Furthermore, the model provides uncertainty for the estimates which are validated using cross validation for both Gaussian and t-distributed measurement errors. At many depths the residuals may have non-Gaussian features, as noted in Kuusela and Stein (2018). To further address this issue, Bolin and Wallin (2020) explore a class of multivariate non-Gaussian spatial models that offer some improvements in prediction on a limited analysis of Argo data.

To conclude this section, we recognize that some aspects of FDA are not altogether new to oceanography. For example, splines have been used as a smoothing approach to interpolate sparse observations in a profile; principal component analysis (PCA), known as empirical orthogonal functions (EOF) analysis, is a common dimension-reduction approach (Thomson and Emery (2014)). However, these are applied in somewhat limited ways that include little statistical considerations.

2.3. Spatial FDA literature. The extension of spatial prediction for scalar data to functional data has primarily been developed recently in the statistics discipline. For independent and identically distributed functional data, the literature has been well developed and presented, for example, in the books of Ramsay and Silverman (2013), Hsing and Eubank (2015), and Kokoszka and Reimherr (2017). For spatially-dependent functional data most of the literature has focused on the idealized regime where entire functions are observed. In particular, there are detailed reviews in Delicado et al. (2010), Aguilera-Morillo, Durbán and Aguilera (2017), Kokoszka and Reimherr (2019), and Martínez-Hernández and Genton (2020). We outline some of the work in this area.

Recent developments in spatial FDA have provided increasingly comprehensive approaches for complex spatiotemporal data. Most of the literature focuses on geostatistical (point-referenced) data, though approaches for areal data and point processes have been considered (Cronie et al. (2019), Delicado et al. (2010), Zhang et al. (2016)). In addition, methods for hierarchical spatial functional data have been developed through the work of Baladandayuthapani et al. (2008), Staicu, Crainiceanu and Carroll (2010), and Zhou et al. (2010). Ruiz-Medina (2011) and Zhang et al. (2016) extend spatial autoregressive or moving-average processes to functional data. Staicu et al. (2012) develop copula-based methods for skewed spatial functional data. Methods for clustering spatiotemporal functional data have been proposed in Jiang and Serban (2012) and Romano, Balzanella and Verde (2017), among others. Theory and methodology for spatial FDA have been explored in Zhang et al.

(2016), and Gromenko, Kokoszka and Sojka (2017) and Zhang and Li (2002). In particular, Gromenko, Kokoszka and Sojka (2017) propose an iterative approach for modeling the mean and covariance structure while also addressing inference on the mean function. Bayesian approaches to spatial FDA are proposed in Baladandayuthapani et al. (2008) and Song and Mallick (2019). The spatial FDA perspective has mostly been considered in applications to environmental data (King et al. (2018), Monestiez and Nerini (2008), Pauthenet et al. (2019), Rodríguez, Dunson and Gelfand (2009)) and medical applications, including neuroscience (Lynch and Chen (2018)) and a cancer study (Baladandayuthapani et al. (2008), Staicu, Crainiceanu and Carroll (2010), Zhou et al. (2010)). Here, we directly situate our approach within this rich literature. Most of the mentioned approaches propose basis expansions of mean and principal component functions and model the principal component scores as a spatial process (e.g., Matérn); our approach does as well. At the broad level our approach is similar to Gromenko, Kokoszka and Sojka (2017). Specifically, our two-stage approach, where we estimate the covariance only after estimating the mean and forming residuals, is similar to steps 1–3 of their Algorithm 3.1. Since there are a large number of parameters, this two-stage approach helps reduce the parameter space. This approach is common in FDA and is supported by the theoretical work in, for example, Li and Hsing (2010) or Yao, Müller and Wang (2005). Ideally, we would employ the iterative algorithm in Gromenko, Kokoszka and Sojka (2017) in the spirit of iteratively reweighted least squares, but we are limited by the computation. We justify this two-stage approach with an appeal to profile likelihood, where the mean is estimated, assuming a fixed within-profile covariance, after which the covariance is estimated assuming a fixed mean. We also extend their methodology by proposing a non-separable covariance structure, as discussed below. Some of the methodology and motivation is similar to King et al. (2018).

The Argo data call for more involved modeling than in the existing spatial FDA literature in a number of respects. The challenges include the addition of another dimension (pressure) to the space and time dimensions, irregularly-spaced data in each of these dimensions, the varying number of measurements per profile, the sparsity of data in space and time, and the large size of the data. Here, we detail a few aspects of our approach that address these complexities. First, the referenced approaches assume a constant mean in space, that is, $\mu(s, d, y, p) = \mu(p)$ for some function $\mu(p)$ as well as a stationary covariance. Due to the nonstationary nature of the Argo data, a constant mean would not be physically adequate. We allow the mean and covariance structure to change in space, providing a way to model nonstationary functional data in space and time, extending the local stationarity assumption of Kuusela and Stein (2018) for their setting of a fixed pressure level. This local approach for the mean and covariance also helps address the computational challenges with respect to the size of the data.

Next, most approaches in the literature depend on a basis representation of profiles. By projecting each profile onto a suitable basis before modeling, this simplifies the subsequent analysis, but such a step introduces systematic error. For such an interpolation approach to be justifiable, all profiles should be densely sampled (cf. Hall, Müller and Wang (2006), Li and Hsing (2010)). For many Argo profiles, in particular ones sampled 2007–2010 that have fewer measurements, an appropriate basis representation cannot be obtained. Our estimation and prediction methodology avoids this issue and naturally accommodates both sparse and dense data at their measured pressure.

Finally, in the Supplementary Material (Yarger, Stoev and Hsing (2022)) we illustrate the nonseparability of our covariance model, described in Section 4, and compare the model with existing literature. Our covariance model, by assuming separability for each principal component direction, allows for a varying space-time covariance structure as a function of depth. Such flexibility is necessary for the Argo data, since processes at the surface can be much different than those at greater depth.

3. Functional mean estimation for Argo data. In this section we introduce our functional approach for mean estimation in which we estimate a smooth function $\mu(s_0, d_0, y, p_{i,j})$ for a location s_0 . We assume that the mean function μ is smooth in terms of space, time, and pressure. Due to the functional nature of the data, we focus on and formalize the smoothness in pressure here. In particular, consider the class of functions

$$\mathbb{W}_2 = \left\{ f \mid f^{(2)} \text{ exists almost everywhere, and } \int_0^{2000} (f^{(2)}(p))^2 dp < \infty \right\},$$

where $f^{(k)}$ is the k th derivative of f . The space $\mathbb{W}_2$ is a Sobolev space of functions widely used for nonparametric inference, including problems in FDA (Hsing and Eubank (2015), Wahba (1990)). The size of $\|f^{(2)}\|_{\mathbb{L}_2}^2 = \int (f^{(2)}(p))^2 dp$ quantifies the smoothness of f , that is, if $\int (f^{(2)}(p))^2 dp = 0$, then f takes the form of a line.

3.1. A functional approach to mean estimation. We consider a mean estimated locally in space and day of the year which can be evaluated at any pressure in $[0, 2000]$. Our novel approach combines local regression (to smooth space and time) and smoothing splines (to smooth pressure) by estimating the function

$$\begin{aligned} & f_{\beta, s_0, d_0}(s_i, d_i, y_i, p) \\ (4) \quad &= \sum_{y=2007}^{2016} \beta_{0,y}(p) 1(y_i = y) + (s_{1i} - s_{10})\beta_1(p) + (s_{2i} - s_{20})\beta_2(p) \\ &+ (s_{1i} - s_{10})^2\beta_3(p) + (s_{2i} - s_{20})^2\beta_4(p) \\ &+ (s_{1i} - s_{10})(s_{2i} - s_{20})\beta_5(p) + (d_i - d_0)\beta_6(p) + (d_i - d_0)^2\beta_7(p), \end{aligned}$$

where $s_0 = (s_{10}, s_{20})$ is a fixed location and d_0 is a fixed day of the year. Here, the functions $\beta_{0,y}$ and β_k are specific to s_0 and d_0 , though we omit this notation for ease of writing. Assuming each function denoted with $\beta_{0,y}$ or β_k falls in the class $\mathbb{W}_2$, we include the standard smoothing spline penalty on the second derivative of each function,

$$\text{Pen}(\lambda) = \lambda_0 \sum_{y=2007}^{2016} \|\beta_{0,y}^{(2)}\|_{\mathbb{L}_2}^2 + \sum_{k=1}^7 \lambda_k \|\beta_k^{(2)}\|_{\mathbb{L}_2}^2,$$

where the λ_j are nonnegative smoothing parameters. This penalty controls the smoothness of the estimated functions. With this notation, for a fixed location s_0 we solve the optimization problem

$$(5) \quad \min_{\beta_k \in \mathbb{W}_2} (\ell_{s_0, d_0}(\beta) + \text{Pen}(\lambda)),$$

where

$$(6) \quad \ell_{s_0, d_0}(\beta) = \frac{1}{n} \sum_{i=1}^n \frac{K_{h_s, h_d}(s_i - s_0, d_i - d_0)}{m_i} \|\Sigma_i^{-\frac{1}{2}}(Y_i - f_{\beta, s_0, d_0, i})\|_2^2$$

and Y_i and $f_{\beta, s_0, d_0, i}$ are vectors with entries $\{Y_{i,j}\}_{j=1}^{m_i}$ and $\{f_{\beta, s_0, d_0}(s_i, d_i, y_i, p_{i,j})\}_{j=1}^{m_i}$, respectively. Here, K_{h_s, h_d} is a product of Epanechnikov kernels—the first based on the great-circle distance between s_i and s_0 with bandwidth h_s and the second based on the difference in day of the year between d_i and d_0 with bandwidth h_d . This type of kernel is commonly used for local regression (Fan and Gijbels (1996)). Also, Σ_i is a matrix that specifies the working correlation between measurements in the same profile; we address choosing its form in the next section. Dividing by m_i in (5) ensures that profiles with more measurements




FIG. 2. Data used from 2010 for first 350 dbar for -170.5° W 0.5° N, $d_0 = 45.25$ for (Left) temperature ($^\circ$ C) and (Right) salinity (practical salinity units, PSU). We plot the estimated mean functions $\beta_{2010}(p)$ and $\beta_{2013}(p)$ for two of the years and the year-averaged estimate $\bar{\beta}(p)$. In this figure as well as Figure 4, we plot according to the oceanography convention with pressure on the y axis in reference to depth in the ocean.

do not contribute in greater proportion to the loss function compared to profiles with fewer measurements. Once again, the resulting functions $\beta_{0,y}$ and β_k are estimated for each fixed location s_0 and time d_0 , omitted for simplicity in the notation. We propose this new general nonparametric approach of combining local regression and spline smoothing for estimating a spatially-varying functional mean.

The optimization problem (5) is solved for temperature and salinity separately. The functions $\beta_{0,y}$ for $y = 2007, \dots, 2016$ give a mean function estimated from each year. The function $\bar{\beta}(p) = \frac{1}{10} \sum_{y=2007}^{2016} \beta_{0,y}(p)$ is the year-averaged mean at s_0 and d_0 . The additional functions β_1 through β_7 are used to estimate the derivatives of the mean, with respect to space and time, for each pressure. Figure 2 gives results at one location in the Pacific Ocean for the first 350 dbar with $d_0 = 45.25$, corresponding the mid-February. The mean functions are able to capture the water column with constant temperature near the surface known as the mixed layer, which we address in more detail in Section 5. The reader can compute (5) for fixed smoothing parameters using an R Shiny application (Yarger (2020a)).

We motivate our functional approach by qualitatively comparing it to a multivariate local regression approach with respect to pressure, space, and time. Both approaches are nonparametric and should behave relatively similarly, given appropriate bandwidths and smoothing parameters. However, some key advantages of the functional approach are as follows. First, multivariate local regression can be challenged by the curse of dimensionality since no points are truly “local” (Fan and Gijbels (1996)), while our approach reduces this problem by using all data in pressure simultaneously. This provides a “middle-ground” nonparametric technique between local estimation (computationally manageable, using a limited amount of data) and multivariate/thin-plate splines (computationally intractable, using all of the data). Also, applying local regression in pressure can introduce new challenges of bandwidth selection, which we avoid. Due to the differences in variability and sampling in pressure, a constant

bandwidth in pressure would not be appropriate. Our approach can also save computation, since reestimation is not needed for any additional pressure measurement. Relatedly, derivatives and integrals of pressure are immediately available for the entire pressure dimension, while, for local regression, only derivatives are available at the points of computation. Our approach is tailored to and reflects the functional nature of the data so that profiles and estimated mean functions can be easily compared.

3.2. Computation and cross validation. In this section we give an overview of our approach for computation and how to choose smoothing parameters. More details are shown in the Supplementary Material (Yarger, Stoev and Hsing (2022)). The solution to (5) must be computed for each location of interest s_0 ; however, calculations for different s_0 do not rely on each other, so they can be easily made in parallel over multiple computer cores. By applying Theorem 6.6.9 of Hsing and Eubank (2015) to losses that include multiple functions in $\mathbb{W}_2$, we obtain that each function of the resulting solution to the infinite dimensional optimization problem (5) is a natural cubic spline with knots at each uniquely observed $p_{i,j}$ for each i such that $K_{h_s, h_d}(s_i - s_0, d_i - d_0) > 0$. Since the smoothness of each function is penalized in the objective function using GCV, having a large number of knots does not contribute to overfitting (Ruppert (2002)). On the other hand, placing a knot at each observed pressure value is prohibitively costly when a large number of profiles are included in each fit. Commonly in nonparametric regression, reducing the number of knots is done using the quantiles of p or equispaced knots (Ruppert (2002)). We adopt a similar strategy in this functional setting by employing penalized cubic B-splines bases with 200 equispaced knots in $[0, 2000]$. Our experiments indicated that the difference with the exact solution involving knots at all relevant pressures is small, and this basis provides knots at intervals near the size of Argo pressure uncertainties of ± 2.4 dbar. Due to the local nature of the B-splines, the relevant matrices needed to compute the solutions are sparse and banded which leads to further computational gains. In particular, the Cholesky decomposition of matrices is numerically efficient. To compute the B-spline basis functions and penalty, we use the `fd` package.

In addition to computing the solution, we also need to choose the smoothing parameters λ_j and bandwidths h_s and h_d . Smoothing parameters λ_j are currently chosen assuming h_s and h_d fixed. We set $h_s = 900$ kilometers for both temperature and salinity and $h_d = 45.25$ days. This provides nearly enough profiles for each grid point and year and uses data from three months of the year. Also, if fewer than 10 profiles were used for each year, h_s is increased so that there are at least 10 profiles used for each year. To choose λ_j , we use generalized cross-validation (GCV) for its favorable properties, ease of calculation, and ability to include a correlation structure in the observations (Wahba (1990)). The GCV score in the context of our problem is

$$\text{GCV}(\lambda) = \frac{(Y - \hat{Y})^\top \Sigma^{-1} (Y - \hat{Y})}{(1 - \text{tr}(A(\lambda))/n_{s_0})^2},$$

where $n_{s_0} = \sum_{i=1}^n 1(K_{h_s, h_d}(s_i - s_0, d_i - d_0) > 0)m_i$, $A(\lambda)$ is the “hat” matrix defined in Yarger, Stoev and Hsing (2022), Y are the observations for temperature or salinity, $\hat{Y}$ are predictions using smoothing parameters λ , and Σ^{-1} is the block-diagonal matrix

$$\Sigma^{-1} = \text{diag} \left[\frac{K_{h_s, h_d}(s_i - s_0, d_i - d_0)}{nm_i} \Sigma_i^{-1}, i = 1, \dots, n \right].$$

Considering approaches like variable bandwidth selection, jointly choosing bandwidths and smoothing parameters, and leave-one-profile-out cross-validation are methodological and computational challenges that can motivate further research. Computing the leverage scores for the calculation of GCV is the largest computational cost in the selection of λ , and we

detail how to compute them in [Yarger, Stoev and Hsing \(2022\)](#). For choosing multiple smoothing parameters, computing the GCV function on a two-dimensional or larger grid becomes prohibitively expensive. We have taken the approach of finding suitable fixed ratios $\eta_\ell = \lambda_\ell / \lambda_0$ for each ℓ , then using the smoothing parameters $a\eta$ and cross-validating on the single parameter $a > 0$. These ratios are chosen to balance the units of each of the covariates, and the quadratic terms require larger amounts of smoothing. In particular, we let $\eta = (1, 10^8, 10^8, 10^{13}, 10^{13}, 10^{13}, 10^9, 10^{13})$ and conduct standard 1-d optimization using `optimize` in R to search for $a \in (10^{-3}, 10^7)$.

The irregular sampling of Argo profiles over pressure can present challenges for naive spline estimation. This issue can be addressed using a working correlation structure in pressure. A simple choice employed here is Markovian-type dependence in continuous pressure. Specifically, we consider $(\Sigma_i)_{j,k} = \exp(-\tau |p_{i,j} - p_{i,k}|)$ with $\tau \in (0, \infty)$, resulting in a tridiagonal precision matrix Σ_i^{-1} . In practice, we have found that using this within-profile correlation with $\tau = 0.001$, which corresponds to a correlation of about 0.9512 for measurements 50 dbar apart, helps both the selection of λ as well as the quality of solution. In Section 4 the within-profile covariance is estimated, and in [Yarger, Stoev and Hsing \(2022\)](#) it is shown that the empirical covariance estimates generally match well with this choice. One could include a nonconstant working variance as well, though such benefit may be marginal.

We compute the solution to (5) in R on a one degree by one degree grid in space for mid-February ($d_0 = 45.25$) between -80° S and 80° N. This results in 47,938 and 46,023 grid points computed for temperature and salinity, respectively. For salinity we use only delayed-mode data. For each profile i from the first three months of the year, residuals were computed by using the mean estimate at the nearest grid point to profile i as $Y_{i,j} - \hat{f}_\beta(p_{i,j})$. The implicit assumption of computing these residuals is that the mean is represented well by a locally quadratic function of day of the year for these three months as in (4). In [Yarger, Stoev and Hsing \(2022\)](#) we compare with the February mean field estimates of [Roemmich and Gilson \(2009\)](#).

3.3. Functional derivatives. One novelty in our approach of combining local regression and spline smoothing is its estimation of functional derivatives. Namely, writing the mean averaged over years as $\bar{\mu}(s, d, p)$, the functions $(\hat{\beta}_1(p), \hat{\beta}_2(p))$ estimate

$$\left(\frac{\partial \bar{\mu}}{\partial s_1}(s, d, p), \frac{\partial \bar{\mu}}{\partial s_2}(s, d, p) \right),$$

the gradient consisting of the partial derivatives at $s = s_0$ and $d = d_0$ of the response with respect to zonal distance and meridional distance, respectively. Likewise, $(2 \cdot \hat{\beta}_3, 2 \cdot \hat{\beta}_4, \hat{\beta}_5)$ estimate the second-order derivatives

$$\left(\frac{\partial^2 \bar{\mu}}{\partial s_1^2}(s, d, p), \frac{\partial^2 \bar{\mu}}{\partial s_2^2}(s, d, p), \frac{\partial^2 \bar{\mu}}{\partial s_1 \partial s_2}(s, d, p) \right),$$

and $(\hat{\beta}_6, 2 \cdot \hat{\beta}_7)$ estimate

$$\left(\frac{\partial \bar{\mu}}{\partial d}(s, d, p), \frac{\partial^2 \bar{\mu}}{\partial d^2}(s, d, p) \right)$$

at the location and time $s = s_0$ and $d = d_0$. Collectively, these functions describe the local quadratic behavior of the mean near s_0 and d_0 . In Figure 3 the derivatives in latitude and time for temperature are given for a cross-section of the ocean for a fixed longitude. Also, the figure includes the direction and strength of the spatial gradient at a fixed pressure of 10 dbar for salinity. These derivatives can identify the direction of warming and cooling for each location and pressure as well as physical properties, including the exchange of salty and fresh waters near the Strait of Gibraltar. Our functional approach facilitates this detailed description of the ocean properties at any pressure.




FIG. 3. Estimates of derivatives of temperature in latitude (Top Left) and time (Bottom Left) for a fixed longitude 179.5 West. The derivative, with respect to latitude, reflects that in the middle latitudes (between (30, 50) and (−50, −30)) the temperature increases as one moves toward the equator near the surface. The derivatives also identify two separate areas of high temperature on either side of the equator near 250 dbar. The derivative with respect to time demonstrates that, in mid-February, the temperature is mostly increasing in the Southern Hemisphere and mostly decreasing in the Northern Hemisphere, as the time of year suggests. On the right, we give the gradient in space in PSU per 100 km for salinity at 10 dbar in mid-February (Right). The gradient points toward the salinity maximum in the Central North Atlantic identifies the flow of salty water from the Mediterranean Sea and shows accordance with Talley *et al.* (2011) Figure 4.15 that gives the distribution of sea-surface salinity in the oceans in January–March.

4. Covariance estimation. After subtracting the mean from the data, the spatial dependence structure of the residuals can be modeled to provide predictions and estimate uncertainties. Modeling the covariance in space, time, and pressure is a challenging task. For example, there are considerable differences in the spatial dependence structure and residual variances at different pressures and locations.

Our covariance estimation can be described in three steps. First, we estimate the functional principal components (FPCs) which explain the first few dimensions of variability in pressure (cf. Hsing and Eubank (2015, Chapter 9)). Next, each profile is summarized by these principal components, and the resulting scores are modeled. Lastly, we estimate the remaining variability not accounted for by the principal components. The implicit assumptions in this approach are that the covariance structure of temperature and salinity changes smoothly as a function of pressure, and only a small number of FPCs are needed to approximate the spatial and temporal structure in pressure. Since modeling the dependence between the raw measurements in space, time, and pressure simultaneously is not practical or appropriate due to the number of observations and flexibility of covariance models, our functional approach facilitates a dimension reduction strategy that shares information across pressure through the FPCs.

We develop this approach in mathematical notation first by assuming

$$(7) \quad Y_{i,j}^0 = X(s_i, d_i, y_i, p_{i,j}) + \epsilon_{i,j},$$

where $\{Y_{i,j}^0\}_{j=1}^{m_i}$ are the residuals for profile i formed by subtracting the mean estimate from the data, $X(s, d, y_i, \cdot)$ for $y_i = 2007, \dots, 2016$ are identically-distributed realizations of a functional random field with mean 0, and $\epsilon_{i,j} \stackrel{\text{ind}}{\sim} N(0, \kappa(p, s, d))$ is an independent measurement error noise with mean 0 and finite variance that may depend on pressure, location, and day of the year. If $X(s_i, d_i, y_i, \cdot) \in \mathbb{L}_2$ for each s_i and d_i , one can write

$$X(s_i, d_i, y_i, p) = \sum_{k=1}^{\infty} Z_k(s_i, d_i, y_i) \phi_k(p),$$

where ϕ_k are fixed orthonormal functions and the $Z_k(s, d, y)$ are scalar random fields that are weakly dependent in time that we refer to as *scores*. This is similar to the Karhunen–Loève expansion for zero-mean square-integrable stochastic processes, though the scores may be correlated across k due to their spatial dependence. In the subsequent development we simplify the notation by defining $Z_{i,k} = Z_k(s_i, d_i, y_i)$ where it does not cause confusion. For an adequate choice of ϕ_k , we would expect that $X(s_i, d_i, y_i, \cdot)$ can be approximated as

$$(8) \quad X(s_i, d_i, y_i, p) = \sum_{k=1}^{K_1} Z_{i,k} \phi_k(p)$$

for some small number K_1 . This effectively reduces the dimension of our problem. Here, each ϕ_k is a fixed function that has been estimated through some form of functional principal component analysis, with one such approach given in Section 4.1. For a choice of ϕ_k and a profile i , the scores are estimated by the least squares solution

$$(9) \quad Z_{i,\cdot} = (\Phi_i^\top \Phi_i)^{-1} \Phi_i^\top Y_i^0,$$

where Y_i^0 is the residuals for profile i and $\Phi_i \in \mathbb{R}^{m_i \times K_1}$ is the matrix with j, ℓ entry $\phi_\ell(p_{i,j})$. The principal component functions ϕ_k and the scores $Z_{i,\cdot}$ are only estimates and not the truth, though we use the same notation for convenience. We found that the alternative approach to estimating the scores proposed by Yao, Müller and Wang (2005) to give similar results, though this approach is computationally expensive to implement on a large scale.

We assume that the decomposition (8) of $X(s_i, d_i, y_i, p)$ holds locally with respect to both $Z_{i,k}$ and ϕ_k , similar to the locally stationary assumption of Kuusela and Stein (2018). That is, for a fixed location s_0 and time d_0 , the functions ϕ_k are estimated and used to form estimates of the $\{Z_{i,k}\}_{k=1}^{K_1}$ and the measurement error variance $\kappa(p) := \kappa(s_0, d_0, p)$ for all nearby profiles. Next, the joint distribution of the nearby scores is modeled. For different choices of (s_0, d_0) , the functions ϕ_k and resulting scores $\{Z_{i,k}\}_{k=1}^{K_1}$ and measurement error variance $\kappa(p)$ are different.

The model gives a clear approach to address the fundamental problem of functional kriging, that is, spatial prediction of functional data, using the conditional distribution at an unobserved location, given the data observed. For any set of data Y^0 to provide a prediction for the function-valued random field $X(s_*, d_*, y, \cdot)$ for an unobserved location s_* at time d_* , one has

$$(10) \quad \mathbb{E}\{X(s_*, d_*, y, p) | Y^0\} = \phi(p)^\top \mathbb{E}\{Z.(s_*, d_*, y) | Y^0\},$$

$$(11) \quad \text{Var}\{X(s_*, d_*, y, p) | Y^0\} = \phi(p)^\top \text{Var}\{Z.(s_*, d_*, y) | Y^0\} \phi(p),$$

where $\phi(p) = (\phi_1(p), \phi_2(p), \dots, \phi_{K_1}(p))^\top$ are the principal components and $Z.(s_*, d_*, y) = (Z_1(s_*, d_*, y), \dots, Z_{K_1}(s_*, d_*, y))^\top$ are the scores of $X(s_*, d_*, y, \cdot)$. Furthermore, for each residual point $Y_{i,j}^0$,

$$(12) \quad \mathbb{E}\{Y_{i,j}^0 | Y^0\} = \mathbb{E}\{X(s_i, d_i, y_i, p_{i,j}) | Y^0\},$$

$$(13) \quad \text{Var}\{Y_{i,j}^0 | Y^0\} = \text{Var}\{X(s_i, d_i, y_i, p_{i,j}) | Y^0\} + \kappa(p_{i,j}).$$

Thus, if one assumes that the field of $\{Z_k(s, d, y); k = 1, \dots, K_1; (s, d) \in \mathbb{R}^3\}$ is Gaussian, one only needs a spatiotemporal model of the scores $Z_k(s, d, y)$ for $k = 1, \dots, K_1$ using the conditional mean and variance as well as estimate $\kappa(p)$. We address the estimation of $\phi_k(p)$ in Section 4.1, the modeling of the scores $Z_{i,k}$ in Section 4.2, and the estimation of $\kappa(p)$ in Yarger, Stoev and Hsing (2022).

4.1. Marginal covariance estimation in pressure. In this section we focus on the estimation of ϕ_k in (8) which amounts to performing local functional principal component analysis (FPCA). A fixed set of basis functions may not be suitable for different locations or seasons, and the resulting decomposition would be suboptimal at most locations. We thus estimate ϕ_k locally in space and time as done with the mean to provide an optimal decomposition. At each location a local version of the approach given in Section 8.3 of Hsing and Eubank (2015) is used to estimate the entire within-profile covariance. Then, the covariance is decomposed to obtain the functional principal components. This approach uses data from both sparse and dense profiles and avoids needing a basis representation of each profile as in Ramsay and Silverman (2013). Also, it resembles our approach for mean estimation by treating the covariance as an expectation, and it provides advantages over other approaches, like thin plate splines by using B-splines that greatly reduce computations (Wahba (1990)).

For fixed s_0 and d_0 , we solve the optimization problem,

$$(14) \quad \min_{f_{s_0, d_0} \in \mathbb{W}_2 \otimes \mathbb{W}_2} (\ell_{s_0, d_0}(f_{s_0, d_0}) + \text{Pen}_{f_{s_0, d_0}}(\lambda)),$$

where

$$(15) \quad \ell_{s_0, d_0}(f_{s_0, d_0}) = \frac{1}{n} \sum_{i=1}^n \frac{K_{h_s, h_d}(s_i - s_0, d_i - d_0)}{m_i(m_i - 1)} \sum_{1 \leq j \neq k \leq m_i} (Y_{i,j}^0 Y_{i,k}^0 - f_{s_0, d_0}(p_{i,j}, p_{i,k}))^2.$$

In particular, f_{s_0, d_0} is restricted to be of the form

$$f_{s_0, d_0}(p_1, p_2) = \sum_{k_1=1}^M \sum_{k_2=1}^M \alpha_{k_1, k_2} \chi_{k_1}(p_1) \chi_{k_2}(p_2),$$

where $\{\alpha_{k_1, k_2}\}_{k_1, k_2=1}^M$ are scalar coefficients and $\{\chi_k(p)\}_{k=1}^M$ is a univariate B-spline basis over a fixed set of knots. As suggested in Wood (2006), the penalty used is

$$\text{Pen}_{f_{s_0, d_0}}(\lambda) = \lambda \text{vec}(\alpha)^\top (\Omega \otimes I_M + I_M \otimes \Omega) \text{vec}(\alpha),$$

where $\otimes$ is the standard Kronecker product, Ω is the univariate smoothing matrix for the B-splines used with k_1, k_2 entry $\int_0^{2000} \chi_{k_1}^{(2)}(p) \chi_{k_2}^{(2)}(p) dp$, and I_M is the $M \times M$ identity matrix. This penalty approximates

$$\lambda \int_0^{2000} \int_0^{2000} \left[\left(\frac{\partial^2 f_{s_0, d_0}}{\partial p_1^2} \right)^2 + \left(\frac{\partial^2 f_{s_0, d_0}}{\partial p_2^2} \right)^2 \right] dp_1 dp_2,$$

as given in Wood (2006). The computation is similar to the approach for mean estimation, with λ chosen by cross-validation and using a product kernel with $h_s = 550$ kilometers and $h_d = 45.25$; this smaller spatial bandwidth is possible since we pool together data from all years. We use $M = 102$ with equally spaced knots over $[0, 2000]$ for the basis χ_k . The overall size of the problem is M^2 , whose computational cost increases much faster compared to the mean estimation. This choice of knots is able to approximate the covariance operator reasonably well while ensuring the calculations are computationally manageable. The exclusion of points with $j = k$ ensures that the measurement error $\epsilon_{i,j}$ is not included in the




FIG. 4. Example of first two estimated functional principal components (Left) temperature (Right) salinity at Long. 90.5 W and Lat. -10.5 S. The principal components suggest much higher variance near the surface of the ocean, as expected.

estimates along the diagonal. This allows us to formally identify the variance parameter of the measurement error in (7), similar to Yao, Müller and Wang (2005).

The main goal of the covariance estimation is to obtain a basis of functional principal components for the space-time modeling; we detail how to obtain orthonormal principal component functions in Yarger, Stoev and Hsing (2022). At this point we estimate the FPCs for temperature and salinity separately. By working in the corresponding FPC bases for future modeling, we optimally reduce the infinite-dimensional kriging problem in pressure to a finite-dimensional one using principal components.

An example of the first two functional principal components for a location is shown in Figure 4. Similar plots for other locations can be viewed on an R Shiny application (Yarger (2020a)). These principal components can give descriptive information on the variance and dependence of temperature and salinity with respect to pressure. There is evidence that the covariance and the principal components for temperature and salinity exhibit considerably different structure.

REMARK. For any fixed location, the principal components are only identifiable up to a sign. For this reason, only one basis $\{\phi_k\}_{k=1}^K$ is used at one time, and the scores are only defined and predicted with respect to this fixed basis. Thus, the products of each score and principal component are invariant to the sign of the principal component. Local regression helps ensure, assuming a sensible smoothing parameter selection, that the estimated covariance varies smoothly as one moves in space. One interesting problem of future research is the estimation of the marginal covariance operator as a function of space and time. This will require a careful registration and alignment of the principal components and scores when moving from location to location.

4.2. Space-time modeling of scores. In this section we model the scores for spatiotemporal prediction. In standard FDA the scores are uncorrelated latent variables, and it is not always meaningful to model and predict them. For any two mean-zero, square-integrable random functions $X_i(p) = \sum_{k=1}^{\infty} Z_{i,k} \phi_k(p)$ for $i = 1, 2$, the covariance becomes $\text{Cov}(X_1(p_1), X_2(p_2)) = \sum_{k_1=1, k_2=1}^{\infty} \phi_{k_1}(p_1) \phi_{k_2}(p_2) \mathbb{E}(Z_{1,k_1} Z_{2,k_2})$. When X_1 and X_2 are independent, $\mathbb{E}(Z_{1,k_1} Z_{2,k_2}) = 0$, and one cannot leverage any dependence between the scores.

However, for spatially dependent functional data, $\text{Cov}(Z_{1,k_1}, Z_{2,k_2})$ may not be 0 and may depend on the distance between the location of profiles 1 and 2. This motivates our approach to model the dependence of the scores and utilize it for spatial prediction.

Focusing on the estimation of the random, mean-zero function in (8), we write

$$T^0(s_i, d_i, y_i, p) = \sum_{k=1}^{K_1} Z_{i,k} \phi_k(p) \quad \text{and} \quad S^0(s_i, d_i, y_i, p) = \sum_{k=1}^{K_2} W_{i,k} \psi_k(p),$$

where T^0 , $Z_{i,k} := Z_k(s_i, d_i, y_i)$ and ϕ_k denote the respective terms of (8) for temperature, and S^0 , $W_{i,k} := W_k(s_i, d_i, y_i)$ and ψ_k denote the terms for salinity. For the modeling we adopt the locally stationary assumption of Kuusela and Stein (2018). For each location, as described in (9), we use the ϕ_k estimated at that location to compute the temperature scores $Z_{i,k}$ for profiles within some radius of that location and, likewise, use the respective terms for salinity, ψ_k , to compute $W_{i,k}$. We exclude a small fraction of profiles that do not have sufficient measurements to compute scores. The goal of this section is to estimate a predictive distribution for the vector

$$\begin{pmatrix} Z_{*,\cdot} \\ W_{*,\cdot} \end{pmatrix} = (Z_{*,1}, Z_{*,2}, \dots, Z_{*,K_1}, W_{*,1}, W_{*,2}, \dots, W_{*,K_2})^\top$$

at an unobserved location to jointly model temperature and salinity. We first introduce our decorrelation step, as explained below, which is similar to Bachoc et al. (2020).

For the modeling of the resulting scores, let Σ_{scores} be a $(K_1 + K_2) \times (K_1 + K_2)$ marginal covariance matrix of $(Z_{i,\cdot}^\top, W_{i,\cdot}^\top)^\top$. This matrix Σ_{scores} is estimated by

$$\hat{\Sigma}_{\text{scores}} = \frac{1}{|D_{s_0}| - 1} \sum_{i \in D_{s_0}} \begin{pmatrix} Z_{i,\cdot} \\ W_{i,\cdot} \end{pmatrix} \begin{pmatrix} Z_{i,\cdot}^\top & W_{i,\cdot}^\top \end{pmatrix},$$

where D_{s_0} are the set of nearby delayed mode profiles. Then, consider the standard eigendecomposition

$$\hat{\Sigma}_{\text{scores}} = V \Gamma V^\top,$$

where Γ is a diagonal matrix, and define

$$(16) \quad \begin{pmatrix} \tilde{Z}_{i,\cdot} \\ \tilde{W}_{i,\cdot} \end{pmatrix} = V^\top \begin{pmatrix} Z_{i,\cdot} \\ W_{i,\cdot} \end{pmatrix}.$$

The resulting transformed scores $(\tilde{Z}_{i,\cdot}^\top, \tilde{W}_{i,\cdot}^\top)^\top$ are then approximately decorrelated with diagonal autocovariance matrix Γ .

Let $M(\nu, \Delta) = c_1 \Delta^\nu \mathcal{K}_\nu(\Delta)$ be the Matérn covariance with parameter ν at distance Δ with unit variance and scale, studied in, for example, Stein (1999). Here, c_1 is a constant so that $M(\nu, 0) = 1$, and $\mathcal{K}_\nu$ is the modified Bessel function of the second kind. The value of $\nu > 0$ governs the smoothness of the field of scores, where larger values give a smoother field. When $\nu = 1/2$, the Matérn model reduces to the exponential function. In our experiments the choice of ν had minimal effects on the resulting predictions, and we set it to the common choice $\nu = 1/2$ as in Kuusela and Stein (2018). For $\tilde{Z}_{i,k}$ and $\tilde{W}_{i,k}$ and each k , a Matérn model is fitted for the decorrelated scores of the form $\mathbb{E}(\tilde{Z}_{i,k} \tilde{Z}_{j,k}) = C_k(\Delta_{i,j})$ or $\mathbb{E}(\tilde{W}_{i,k} \tilde{W}_{j,k}) = C_{K_1+k}(\Delta_{i,j})$ if $y_i = y_j$ with

$$(17) \quad C_k(\Delta) = \gamma_k \cdot M\left(\nu, \sqrt{\left(\frac{\Delta_{s_1}}{\theta_{s_1,k}}\right)^2 + \left(\frac{\Delta_{s_2}}{\theta_{s_2,k}}\right)^2 + \left(\frac{\Delta_d}{\theta_{d,k}}\right)^2}\right) + \sigma_k^2 \cdot 1(\Delta = 0),$$

where $\Delta = (\Delta_{s_1}, \Delta_{s_2}, \Delta_d)$ is a vector of corresponding distances in space and time. The parameters $\theta_{s_1,k}$, $\theta_{s_2,k}$, and $\theta_{d,k}$ are scale parameters that specify the correlation ranges for each of the directions. Lastly, γ_k and σ_k^2 are parameters that describe the variance of the spatial process and the nugget, respectively. This space-time model is considered in [Kuusela and Stein \(2018\)](#) in their fixed pressure level analysis.

In summary, the resulting covariance of temperature and salinity is

$$(18) \quad \mathbb{E} \left\{ \begin{pmatrix} T^0(s_i, d_i, y_i, p_1) \\ S^0(s_i, d_i, y_i, p_1) \end{pmatrix} \begin{pmatrix} T^0(s_j, d_j, y_j, p_2) \\ S^0(s_j, d_j, y_j, p_2) \end{pmatrix}^\top \right\} = \Xi_{p_2}^\top V C(\Delta_{i,j}) V^\top \Xi_{p_1}$$

if $y_j = y_i$ and 0 otherwise, where $C(\Delta_{i,j}) \in \mathbb{R}^{(K_1+K_2) \times (K_1+K_2)}$ is the diagonal matrix with the k th element $C_k(\Delta_{i,j})$ and $\Xi_p = \begin{pmatrix} \phi^{(p)} & 0 \\ 0 & \psi^{(p)} \end{pmatrix} \in \mathbb{R}^{(K_1+K_2) \times 2}$. This model, by considering a nugget effect on each of the scores, also results in a kind of “functional nugget” for the process, as described in [\(Zhang and Li \(2002\)\)](#). This functional nugget has covariance

$$\Xi_p^\top V C_\sigma V^\top \Xi_p,$$

where $C_\sigma = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_{K_1+K_2}^2)$.

We estimate the spatial model for February at each location using data from January, February, and March. For each location, profiles within 1100 kilometers were used (similar to the size of the moving windows used in [Kuusela and Stein \(2018\)](#)). We set $K_1 = K_2 = 10$, which allows the profiles to be well represented by the principal components, though our experiments suggest that using more principal components may slightly improve predictions near the surface. We provide the reasoning of this choice in [Yarger, Stoev and Hsing \(2022\)](#), where we show that 10 components explain a large proportion of the variability in both temperature and salinity. Choosing the number of functional principal components under a smoothly-varying covariance structure in space could be developed based on [Li, Wang and Carroll \(2013\)](#). To estimate the parameters γ_k , $\theta_{s_1,k}$, $\theta_{s_2,k}$, $\theta_{d,k}$, and σ_k^2 for each k , we employ the same approach as [Kuusela and Stein \(2018\)](#) using maximum likelihood summarized below. Let $\tilde{Z}_y$ be the scores for one k for year y in each of the above models, and let $\text{Var}(\tilde{Z}_y) = \Sigma_y$ be a matrix specified by the parameters in (17) above. We assume that $\tilde{Z}_y$ are multivariate Gaussian so that the log likelihood of the data for all $y = 2007, \dots, 2016$ is

$$-\frac{1}{2} \left(\sum_{y=2007}^{2016} \log(\det(\Sigma_y)) + \tilde{Z}_y^\top \Sigma_y^{-1} \tilde{Z}_y + n_y \log(2\pi) \right),$$

where n_y is the number of observations used in year y . This likelihood treats data from different years as independent. To maximize the likelihood, we use the optimization L-BFGS-B algorithm, due to [Byrd et al. \(1995\)](#), implemented in the `optim` function in R.

One challenge is that quality control is essential for the salinity data; many Argo profiles that have high-quality temperature data may not have the same quality of salinity data. Instead of discarding such profiles, we offer a solution by using an established missing data approach via an expectation-maximization-type (EM) algorithm:

1. (E step) Using the temperature, delayed-mode salinity data, and the estimated parameters, form a prediction (the conditional expectation) for the real-time salinity scores. In the first iteration the prediction of the real-time salinity scores are 0.
2. (M step) Using all data (obtained from the E step) as if it were delayed-mode, estimate the model parameters via maximum likelihood.
3. Alternate between the E and M steps, and repeat until a convergence criterion is met.

This treats the real-time salinity scores as unobserved, latent variables. The differences between estimated parameters from consecutive steps decrease quickly after the first few steps. Therefore, at each grid point we decided to perform six iterations of the algorithm, and the parameter estimates from the final M step are used. This choice strikes a balance between computation time and statistical accuracy.

Estimating the joint dependence between temperature and salinity is not considered in Kuusela and Stein (2018) and Roemmich and Gilson (2009). While it requires additional computation, accounting for this dependence provides a more comprehensive analysis of the Argo data. In particular, this enables us to predict and provide uncertainty estimates for functionals of temperature and salinity, such as potential density and potential temperature. The estimated parameters can be viewed using an R Shiny application (Yarger (2020a)).

4.3. Predictions, uncertainties, and prediction bands. In this section we employ the estimated spatial covariance for functional kriging. Under the assumptions of our model, this provides an optimal functional prediction at an unobserved location. To detail this approach, let Σ_{y*} be covariance matrix of the (true) decorrelated scores $\tilde{Z}_{y*} = (\tilde{Z}_{i,k})_{i=1}^{n_{y*}}$ for a fixed k in an area around a fixed location for a year y_* . Notably, using the local stationarity assumption, profiles within 1100 kilometers are used as in the Matérn estimation step. This provides enough data for prediction while avoiding introducing data that may violate the locally-stationary assumption. The conditional distribution of $\tilde{Z}_{*,k} := \tilde{Z}_k(s_*, d_*, y_*)$ at an unobserved location given $\tilde{Z}_{y*}$ is

$$\tilde{Z}_{*,k} | \tilde{Z}_{y*} \sim N(\Sigma_{12}^\top (\Sigma_{y*})^{-1} \tilde{Z}_{y*}, \gamma_k + \sigma_k^2 - \Sigma_{12}^\top (\Sigma_{y*})^{-1} \Sigma_{12}),$$

where $\Sigma_{12} = \text{Cov}(\tilde{Z}_{*,k}, \tilde{Z}_{y*})$. However, in our prediction problem, $\tilde{Z}_{y*}$, Σ_{12} , and Σ_{y*} are unknown and are estimated by the approaches described in Section 4.2.

We similarly obtain the predictions for the decorrelated salinity scores $\tilde{W}_{*,k}$. From these estimated distributions of the $\tilde{Z}_{*,k}$ and $\tilde{W}_{*,k}$, using the relation that $\begin{pmatrix} Z_* \\ W_* \end{pmatrix} = V \begin{pmatrix} \tilde{Z}_* \\ \tilde{W}_* \end{pmatrix}$ described in (16), the conditional distribution of the original scores is

$$\begin{pmatrix} Z_* \\ W_* \end{pmatrix} | \tilde{Z}_{y*} = V \begin{pmatrix} \tilde{Z}_* \\ \tilde{W}_* \end{pmatrix} | \tilde{Z}_{y*} \sim N \left(V \mathbb{E} \left\{ \begin{pmatrix} \tilde{Z}_* \\ \tilde{W}_* \end{pmatrix} | \tilde{Z}_{y*} \right\}, V \text{Var} \left\{ \begin{pmatrix} \tilde{Z}_* \\ \tilde{W}_* \end{pmatrix} | \tilde{Z}_{y*} \right\} V^\top \right).$$

The conditional distribution of $T^0(s, d, y, p)$ and $S^0(s, d, y, p)$ can be found using (18) or (12) and (13), providing a prediction for any pressure. In Yarger, Stoev and Hsing (2022), we give an example prediction for one pressure.

We test the uncertainty estimates based on this model in a leave-one-profile-out manner. For each February profile, the profile is left out, and nearby profiles are used to predict at the location and time of the profile. Then, the left-out profile is compared with the predictions. For salinity, only delayed-mode profiles are compared. For each quantity, we use bounds of two standard deviations from the mean which corresponds to approximately a 95.4 percent prediction interval. For brevity, we develop uncertainties for temperature, and similar bounds are obtained for salinity. We consider both pointwise and uniform prediction bounds on the residual curves $Y_{i,j}^0 = \phi(p_{i,j})^\top Z_{*,\cdot} + \epsilon_{i,j}$. The pointwise $1 - \alpha$ interval for the residual at pressure $p_{i,j}$, based on (12) and (13), is

$$\phi(p_{i,j})^\top \mathbb{E}\{Z_{*,\cdot} | \tilde{Z}_{y*}\} \pm q_{1-\alpha/2} \sqrt{\phi(p_{i,j})^\top \text{Var}\{Z_{*,\cdot} | \tilde{Z}_{y*}\} \phi(p_{i,j}) + \hat{\kappa}(p_{i,j})},$$

where $q_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of $N(0, 1)$ and $\hat{\kappa}$ is estimated as described in Yarger, Stoev and Hsing (2022). In addition, we develop simultaneous predictions bands over pressure by using the approach of Choi and Reimherr (2018) reviewed in the Supplementary Material.

TABLE 1
Average pointwise coverages of intervals and bands over all pressures

Quantity	# Profiles	Pointwise coverage	Band coverage
Temperature	76,016	96.2	95.6
Salinity	45,188	98.2	96.6
Nominal level		95.4	95.4

In our empirical coverages in Table 1 (where $K_1 = K_2 = 10$), the intervals and bands show good coverage for both temperature and salinity. In Table 1 the band coverage refers to the proportion of profiles for which every observation of the left-out profile was covered by the estimated band. The pointwise coverages correspond to the proportion of all measurements covered by the intervals over all pressures. We also summarize the pointwise coverages by pressure in Yarger, Stoev and Hsing (2022), and the coverage is achieved for most of the pressure dimension, though typically the intervals in the range 20–200 dbar do not meet full coverage, due to more complex processes near the surface.

4.4. *Validation and comparison.* We compare our approach with the Roemmich and Gilson (RG) reference model and Model 5 of Kuusela and Stein (2018) (KS) which provide predictions only at fixed pressure levels. In the Supplementary Material we compare the differences between the KS and functional predictions at 10, 300, and 1500 dbar and find them to be generally comparable. Also, we can compare the predictive errors through the cross-validation approach described in the previous subsection. Our functional approach enables the prediction of temperature and salinity *without* interpolation onto fixed pressure levels. To provide a comparison with the fixed pressure levels of KS, we compute summaries of the residuals by breaking up the interval $[0, 2000]$ using the midpoints of the Roemmich and Gilson pressure levels. For example, the intervals $(6.25, 15]$, $(290, 310]$, and $(1456.25, 1550]$ correspond to the 10 dbar, 300 dbar, and 1500 dbar levels, respectively. Not all profiles are included in the comparison. For KS, profiles are removed in boundary seas and where the interpolation fails, that is, where there are no measurements either above or below the relevant pressure level, and we remove them in this comparison as well and only use profiles included in KS at any of 10, 300, and 1500 dbar. The prediction errors are evaluated by the root mean squared error (RMSE), defined as $\sqrt{\frac{1}{n} \sum_{i=1}^n (y_{i,p} - \hat{y}_{i,p})^2}$ and the 50% (median) and 75% (3rd quartile) quantiles of $|y_{i,p} - \hat{y}_{i,p}|$ where $y_{i,p}$ are the measurements corresponding to pressure level p , and $\hat{y}_{i,p}$ are the predictions for that measurement.

We show the results in Table 2 as well as in the Supplementary Material. Our method outperforms the Roemmich and Gilson-type reference model and has approximately the same the prediction error as KS. Notably, we suspect that avoiding interpolation onto pressure levels considerably improves our prediction error, especially at greater depths. For example, at 1500 dbar, the RMSE for the functional model outperforms KS, though it trails in the outlier-resistant measures of the median and third quartile. This is due to a small number of profiles that have sparse measurements at greater depths, leading to poor quality of interpolation in pressure. At 300 decibars our functional model improves upon KS for each of the metrics, and at 10 decibars the functional model is slightly worse. We explain a possible reason for this gap at 10 decibars. Mainly, the correlation lengths in space can decrease quickly when moving from 10–20 dbar to 40–50 dbar in some locations. Due to this effect, a pointwise approach as in KS can better model the surface pressure levels because the conditions in the small width of the interval near the surface are not easily isolated by the scores based on a limited number of principal components. This motivates future work on a new space-time

TABLE 2
Comparison of KS and functional approach prediction errors, temperature. RG residuals and functional residuals refer to the residuals after subtracting the respective mean

Pressure	Metric	RG residuals	Functional residuals	RG-type model	KS	Functional model
10	RMSE	0.8889	0.7540	0.6135	0.5072	0.5215
10	Q3	0.8670	0.6247	0.5026	0.3735	0.3940
10	Median	0.4750	0.3193	0.2556	0.1801	0.1961
300	RMSE	0.8149	0.8552	0.5782	0.5124	0.4968
300	Q3	0.6320	0.6845	0.4213	0.3684	0.3644
300	Median	0.3062	0.3494	0.1991	0.1740	0.1720
1500	RMSE	0.1337	0.1381	0.1014	0.0883	0.0857
1500	Q3	0.1043	0.1160	0.0736	0.0641	0.0689
1500	Median	0.0530	0.0620	0.0356	0.0311	0.0349

functional model that allows a scale parameter to change smoothly, but quickly, as a function of pressure, or, alternatively, an approach to adaptively choose the number of principal components in space (ref. Section 6).

In Yarger, Stoev and Hsing (2022), the computational costs of our approach and KS are roughly compared. We conclude that, when focusing on temperature, the FDA approach can provide similar predictions for all pressures in roughly the same amount of time it takes to compute a pointwise approach for 13 pressure levels. Thus, our approach can provide approximately a four to five times speedup when considering the 58 Roemmich and Gilson pressure levels.

5. Applications: Ocean heat content and potential density estimates. The procedures of Sections 3 and 4 result in estimated functions of temperature and salinity at each location. For these functions, derivatives and integrals can be easily calculated. Also, other oceanographic measures of interest, like potential density and conservative temperature, can be derived directly from the estimated temperature and salinity using TEOS-10 (e.g., in R, Kelley, Richards and WG127 SCOR/IAPSO (2017)). In the Supplementary Material we present a general framework for leveraging these estimates for other scientific problems and give specific examples in this section.

5.1. Ocean heat content. The amount of heat contained in the ocean is of great interest for global climate change and has been studied extensively, since the ocean absorbs the majority of the Earth’s excess heat. A nonexhaustive list includes Levitus et al. (2012), Roemmich et al. (2015), Lyman and Johnson (2013), Roemmich, Gould and Gilson (2012), and Johnson and Birnbaum (2017). While integrating temperature over pressure describes the heat content in the ocean, it is biased since the temperature of two volumes of water with the same amount of heat content at two different pressures is different. For this reason, conservative temperature is more commonly used for heat content estimates (McDougall (2003)). Conservative temperature can be calculated using the standard oceanographic toolbox McDougall and Barker (2011) which is implemented in R from Kelley, Richards and WG127 SCOR/IAPSO (2017). We use the delta method approach described in Yarger, Stoev and Hsing (2022) to estimate its distribution. Following Meyssignac et al. (2019), denote conservative temperature as a function of temperature, practical salinity, and pressure at a location as $\Theta(t, s, p)$ and the ocean heat content at a location as

$$Q = \int_0^{p^*} c_p \rho \Theta(T_{s,d,y}(p), S_{s,d,y}(p), p) dp,$$

where c_p and ρ are constants (the specific heat capacity and density of seawater, respectively) and $T_{s,d,y}$ and $S_{s,d,y}$ are the predicted temperature and salinity functions at location s on day d for year y .

For a fixed location and day of the year, we consider anomalies from the mean as the difference between the ocean heat content (OHC) using a mean averaged over all years and the conditional expectation of OHC for one year. Specifically, anomalies from the mean are computed for each year as $\mathbb{E}\{Q|\tilde{Z}_{y*}\} - \mathbb{E}\{Q\}$, where $\mathbb{E}\{Q\}$ is the ocean heat content, given by year-averaged mean $\bar{\beta}(p) = \frac{1}{10} \sum_{y=2007}^{2016} \beta_{0,y}(p)$ from the mean form described in (4), and an example of these estimates and standard deviations are shown in Figure 5 for February 2016, while similar plots for other years can be viewed in an R Shiny application (Yarger (2020a)). Such estimates are computed for each location and a fixed day of the year in mid-February. We compare the estimates for 0–700 dbar with the estimates available at NOAA NODC (2019) that employ the Levitus et al. (2012) approach to estimation of ocean heat content. The large-scale features of the fields are similar, though our integrated functions show finer-detail and smaller-scale features as well. We hypothesize that much of the difference in smoothness and features is due to different temporal windows. Our field, as a prediction for a fixed day in mid-February, estimates finer-scale activity, compared to the NOAA January–




FIG. 5. (Top Left) Ocean heat content anomaly estimates from the year-averaged mean, February 2016 functional estimate integrated for 0–700 dbar in zettajoules, (Top Right) NOAA estimate for 0–700 m, January–March 2016 average, (Bottom Left) our estimates of the standard deviation of OHC at each location, February 2016, and (Bottom Right) comparison of variance estimates on a coarse and fine grid, February 2016.

to-March average. The FDA approach, by modeling the dependence between different pressures, makes these uncertainty estimates possible between any two pressures in $[0, 2000]$ as a natural consequence of our functional kriging approach.

Our functional data approach can evaluate the level of error for any interpolation scheme when estimating integrals of the ocean properties. [Cheng and Zhu \(2014\)](#) have evaluated the levels of uncertainty in ocean heat content due to insufficient sampling in pressure; we evaluate here the amount of error one introduces by using a fixed number of pressure levels with respect to the integrated ocean heat content. We compute our estimates for 0–700 dbar on a fine (.5 dbar) grid as well as a coarse (10 dbar) grid that gives similar pressure gaps used in [Roemmich and Gilson \(2009\)](#) or [Li et al. \(2020\)](#). We evaluate our estimates of OHC as $Q = c_p \rho \sum_{m=1}^{M_1} (p_{m+1} - p_m) \Theta(T(p_m), S(p_m), p_m)$ for a grid of pressures. Based on the functional estimates, we derive the mean and variance of the estimates in the Supplementary Material for two different grids,

$$Q_{\text{fine}} \sim N(\mu_{\text{fine}}, \sigma_{\text{fine}}^2),$$

$$Q_{\text{coarse}} \sim N(\mu_{\text{coarse}}, \sigma_{\text{coarse}}^2).$$

For each location, $(\mu_{\text{coarse}} - \mu_{\text{fine}})/\mu_{\text{fine}}$ is negligible, suggesting there is little bias introduced by using a limited number of pressures. On the other hand, the differences in the estimates of variance $(\sigma_{\text{coarse}}^2 - \sigma_{\text{fine}}^2)/\sigma_{\text{fine}}^2$, plotted in [Figure 5](#), are larger and can reach 0.3% in some areas. This suggests that, if one uses oceanographic products at fixed levels to estimate the ocean heat content, the estimates may be practically unbiased but may be burdened with slightly higher variance. Our functional approach provides this comprehensive estimate of the mean and covariance in pressure which evaluates the consequences of specific discretization approaches in pressure.

5.2. Use of potential density estimates for mixed layer depth. Estimating the joint dependence of temperature and salinity for any pressure gives estimates of quantities, like potential density that give important information about the vertical structure of the oceans. More dense water sinks below less dense water, and thus potential density helps describe the stratification of the oceans: the larger the potential density gradient is in pressure, the more stratified the water is at this point ([Talley et al. \(2011\)](#)). Potential density can be computed directly from temperature, salinity, pressure, and location using [Kelley, Richards and WG127 SCOR/IAPSO \(2017\)](#). In this section we use potential density to estimate the depth of the mixed layer (which can be characterized by approximately constant potential density), and in the next section we evaluate the deviations from monotonicity of potential density.

To describe this first application, the mixed layer is a section of the ocean near the surface where the water mixes freely, giving near-uniform properties of temperature, salinity, and density. The mixed layer governs the interaction between the atmosphere and the ocean, and thus its study can reveal information about the carbon uptake and heat content of the ocean, among other features ([Holte et al. \(2017\)](#)). During the summer the temperature at the surface rises considerably, and the mixed layer is more shallow. During the winter the mixed layer deepens at a lower temperature, resulting in large seasonal changes of its depth.

Mixed layers are usually estimated using discretely observed profiles; see [Sections 4.2 and 7.4 of Talley et al. \(2011\)](#) and [Holte and Talley \(2009\)](#) for algorithms to estimate the mixed layer, and mixed-layer climatologies include [Schmidtke, Johnson and Lyman \(2013\)](#), [Holte et al. \(2017\)](#), and [Hosoda et al. \(2010\)](#). In comparison to these approaches, our functional approach offers two advantages. First, by basing the mixed layer estimates on entire functions predicted from pooled data, we avoid discretization error in the mixed layer estimates. Second, our estimates produce entire mixed layer distributions, even when few profiles have been observed nearby; these are robust to the skewed nature of mixed layers.

These estimates and their variability are assessed using the parametric bootstrap approach described in [Yarger, Stoev and Hsing \(2022\)](#). For each location and year, we simulate $B = 1000$ times from the distribution of February mixed layer depth, based on a year-averaged mean, and B times for each year from the conditional distribution of February mixed layer depth. We present results using the variable density threshold approach described in [Holte et al. \(2017\)](#), where the mixed layer depth is chosen as the first depth for which potential density decreases an amount corresponding to a temperature decrease of 0.2°C .

Our modeling framework allows us to examine the within and across year variability of MLD (mixed layer depth) estimates. This is an important step in quantifying significant anomalies and trends in the MLD, due perhaps to climate change, which is of fundamental scientific importance. Denoting D_j as the estimated mixed layer depth for the j th simulation using the year-averaged mean, and $D_{y,j}$ as the estimated mixed layer depth for the j th conditional simulation for year y , we define estimates based on summaries of the values: $\bar{D} = \frac{1}{10B} \sum_{y=2007}^{2016} \sum_{j=1}^B D_{y,j}$ estimates the overall mean, the conditional simulation median $\tilde{D}_y$ estimates the median for year y , and the year-averaged simulation median $\tilde{D}$ estimates the year-averaged median. Then, consider the estimates of the variation

$$\text{MAEY} = \frac{1}{10B} \sum_{y=2007}^{2016} \sum_{j=1}^B |D_{y,j} - \tilde{D}_y|,$$

$$\text{MAE} = \frac{1}{10B} \sum_{y=2007}^{2016} \sum_{j=1}^B |D_{y,j} - \tilde{D}|,$$

where D is distributed as the year-averaged mixed layer depth at a given location and the factor of 10 adjusts for the 10 years. That is, the mean absolute error with yearly estimates (MAEY) and the year-averaged mean absolute error (MAE) give estimates of the variation that are relatively robust to outliers. The value $p_{\text{year}} = \text{MAEY}/\text{MAE}$ gives an estimate of the relative sizes of errors with year-specific medians vs. a single group median which evaluate yearly variation in the MLD.

In Figure 6 we show selected summaries from the results. The algorithm picks out both shallow mixed layers during the summer in the Southern Hemisphere as well as deeper mixed layers during the winter for the Northern Hemisphere (top left). Winter mixed layers show more variation than summer mixed layers (bottom right). There are two main reasons for this. First, the distribution of the MLD is truncated near the surface, so distributions of MLD near the surface will show less variation. Also, during the summer there is more stratification and thus larger differences in ocean properties at the depth of the mixed layer, so the mixed layer is more consistent. On the other hand, in winter the mixed layer depth is less well defined, and absolute differences of the MLD from its median can be greater than 65 decibars. Finally, the MAE is generally about the same size as MAEY (bottom left) which indicates that there are not large differences in median mixed layer depths between years. Again, this pattern is more evident in the Northern Hemisphere, while, for some locations in the Southern Hemisphere, there is substantial between-year variation. These results indicate that differences in MLD from year-to-year may not be discernible based on the current Argo array alone.

5.3. Monotonicity of density. Potential density generally increases as a function of pressure as water becomes more dense, though this can often be violated for periods of a few hours ([Talley et al. \(2011\)](#)). Here, our estimates are used to evaluate the occurrence of these density “inversions,” where potential density becomes nonmonotone. [Kuusela and Stein \(2018\)](#) suggest that applying a monotonicity constraint on the density may improve estimation of the mean and covariance structure. [Talley et al. \(2011\)](#) suggest that these inversions occur




FIG. 6. February mixed layer depth in dbar by the variable threshold approach (Top Left) average $\overline{D}$, (Top Right) 80th percentile of the conditional MLD distribution, averaged over the ten years, February mixed layer depth (Bottom Left) p_{year} , (Bottom Right) MAEY.

only on the order of a few hours, as gravity removes the instability in the density. Though we have not imposed this constraint, we are in a position to evaluate how well this constraint is satisfied based on our estimates. We compare the amount of inversion in our predictions to the amount of density inversion in the raw Argo profile data, which we detail in [Yarger, Stoev and Hsing \(2022\)](#). Here, we use our functional uncertainty estimates to evaluate how consistent density inversions are at a fixed pressure.

To address the salience of the density constraints, we simulate from 1000 functions, using our conditional simulation approach, then compute the proportion of times a density inversion is shown at a particular pressure. In the bottom of Figure 7, we plot this at a pressure of 550 decibars for the year 2015 (right) and compare it to the estimated gradient from raw Argo profiles using finite differences (left). At this pressure, areas of density inversions are consistently shown where marginal seas mix with the open oceans as well as areas in the Antarctic Circumpolar Current, where there are stronger currents. These areas correspond to areas where negative or low density gradients occur in the profiles. Many of these areas are deep water formation regions, where fresh, cool water sinks, due to its high density. On the other hand, in most of the open oceans there is little evidence of density inversion at this pressure. We conclude that implementing a hard density constraint may not be appropriate, especially in areas of consistent ocean mixing near marginal seas and in the Southern




FIG. 7. (Left) *Estimated gradient using the two closest measurements to 550 decibars for February 2015 profiles*, (Right) *the proportion of conditional simulations that show increasing density at 550 decibars in 2015*.

Ocean. Moreover, the conditional simulations show that further study at all pressure levels can provide valuable statistical insights to the open scientific question on the change of ocean stratification and its effect on thermohaline circulation (Li et al. (2020)).

6. Conclusion and future directions. The Argo data are an exemplary modern dataset that motivates new statistical approaches, development of methodology, and appropriate statistical applications. In this paper we have provided the first comprehensive functional-data analysis of the Argo data which addresses methodological and computational challenges for mean estimation, covariance estimation, functional kriging, and estimation of functionals of the estimates. Our approach avoids the simplification of data in pressure via interpolation which limits other methods' ability to provide a comprehensive analysis. The predicted functions give powerful new tools to fully explore important scientific problems. Furthermore, our approach can decrease the computational burden of prediction by sharing information across pressure. Our estimates match well or outperform existing methodologies that estimate ocean properties at fixed pressure levels. Our analysis also introduces the local estimation of functions and represents a leap forward in the analysis of spatiotemporal functional data.

The methods we develop could be applied to scientific problems in neuroscience and spatial statistics. Here, we focus on spatial sensor networks. In the context of these applications, our methods are amenable to use time, instead of pressure, as the functional variable. Two specific case studies could include estimating air pollution (as in King et al. (2018)) and the Canadian weather data (Ramsay et al. (2018)) studied in Delicado et al. (2010) and Kokoszka and Reimherr (2019), among others. Instead of considering the annual cycle of temperature at only 35 locations in the Canadian weather data, one could provide high-granularity estimates using thousands of weather stations in North America. Notably, our mean estimation approach establishes a new, computationally-efficient, hybrid methodology that combines kernel estimation and smoothing splines.

Throughout our approach, there are areas for improvement. For mean estimation for Stage 1, one would want to select the amount of nearby data adaptively and allow for elliptical regions in space. This is especially important for areas in the Western boundary currents and other areas where changes in ocean properties are highly directional in space. One approach would be to extend algorithms from local regression that choose the bandwidth to this functional model. Using iteratively reweighted least squares or more careful smoothing parameter selection may also give improvements.

For our spatial covariance modeling for Stage 2, we employ a relatively simple model that successfully captures key features by jointly modeling temperature and salinity. In general, we are limited by computational challenges, which could be addressed with approximate models, for example, Vecchia's approximation (Guinness (2021)) or the SPDE approach (Lindgren, Rue and Lindström (2011)). More complexity should be explored in the models.

For example, a functional model that allows rapid changes in the scale parameter as a function of depth would likely improve upon our model. In addition, there is some evidence that the cross-covariance between vectors of principal component scores include nonreversible, that is, asymmetric dependence which is not available in the scalar Matérn-type multivariate models. In addition, one could explore non-Gaussian models, which could provide better coverage for prediction intervals, as demonstrated in Kuusela and Stein (2018) and Bolin and Wallin (2020). Also, we have only modeled the local spatial dependence, and, ideally, one would also like to combine estimates across space with uncertainty, for example, using an approach similar to Wiens, Nychka and Kleiber (2020). This would enable uncertainty estimates for global ocean heat content.

There is a wide variety of statistical research directions using the Argo data, many of which are noted in the conclusion of Kuusela and Stein (2018). For instance, one would want to consider integrating Argo data with other oceanographic data (e.g., satellite data) as well as using additional biogeochemical variables that a limited set of Argo floats measure. Although we have considered many standard approaches in FDA for use on the Argo data, there are more tools that could be applied, including clustering of profiles, functional regression, canonical correlation analysis, hypothesis testing, and data fusion with scalar data, like sea surface temperature data. Moreover, the Argo data calls for full methodological and theoretical development of the field of space-time functional data. For example, the large-sample properties of the methodology used in this paper could be explored under functional and spatial dependence.

Acknowledgments. We would like to thank the physical oceanography group at the Scripps Institution of Oceanography, including Sarah Gille, Lynne Talley, Matt Mazloff, Isabella Rosso, John Gilson, and Dean Roemmich; in addition, we want to thank Mikael Kuusela, Alison Gray, and Donata Giglio for helpful comments on our work. We thank Michael Stein for suggesting to look at the observed monotonicity properties of density. We would also like to thank Moritz Korte-Stapff for his work on the covariance estimation. We finally thank the reviewers and an Associate Editor for important comments and suggestions that substantially improved the quality of the paper.

These data were collected and made freely available by the International Argo Program and the national programs that contribute to it: <http://www.argo.ucsd.edu>. The Argo Program is part of the Global Ocean Observing System. This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) which is supported by National Science Foundation grant number ACI-1548562 (Towns et al. (2014)). Code for the analyses presented in this paper are available at Yarger (2020b).

Funding. The authors acknowledge support from grants DMS-1646108 and DGE-1841052 for Drew Yarger and DMS-1916226 for Stilian Stoev and Tailen Hsing.

SUPPLEMENTARY MATERIAL

Supplementary document PDF (DOI: [10.1214/21-AOAS1477SUPPA](https://doi.org/10.1214/21-AOAS1477SUPPA); .pdf). Description of further modeling choices, computational details, and additional figures.

Code for the described analyses (DOI: [10.1214/21-AOAS1477SUPPB](https://doi.org/10.1214/21-AOAS1477SUPPB); .zip). This code is available in a .zip file as Supplementary Material. The code is also available at <https://github.com/dyarger/argofda>. A link to data, results, and other larger files is also given at this site. The code is in folders that denote each step of the analysis.

REFERENCES

- AGUILERA-MORILLO, M. C., DURBÁN, M. and AGUILERA, A. M. (2017). Prediction of functional data with spatial dependence: A penalized approach. *Stoch. Environ. Res. Risk Assess.* **31** 7–22. <https://doi.org/10.1007/s00477-016-1216-8>
- ARGO (2000). Argo float data and metadata from global data assembly centre (Argo GDAC). *SEANOE*.
- BACHOC, F., GENTON, M. G., NORDHAUSEN, K., RUIZ-GAZEN, A. and VIRTA, J. (2020). Spatial blind source separation. *Biometrika* **107** 627–646. [MR4138980 https://doi.org/10.1093/biomet/asz079](https://doi.org/10.1093/biomet/asz079)
- BALADANDAYUTHAPANI, V., MALLICK, B. K., HONG, M. Y., LUPTON, J. R., TURNER, N. D. and CARROLL, R. J. (2008). Bayesian hierarchical spatially correlated functional data analysis with application to colon carcinogenesis. *Biometrics* **64** 64–73, 321–322. [MR2422820 https://doi.org/10.1111/j.1541-0420.2007.00846.x](https://doi.org/10.1111/j.1541-0420.2007.00846.x)
- BARTH, A., AZCARATE, A. A., JOASSIN, P., BECKERS, J.-M. and TROUPIN, C. (2008). Introduction to optimal interpolation and variational analysis. *GeoHydrodynamics and Environmental Research* **27**.
- BOLIN, D. and WALLIN, J. (2020). Multivariate type G Matérn stochastic partial differential equation random fields. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **82** 215–239. [MR4060983 https://doi.org/10.1111/rssb.12239](https://doi.org/10.1111/rssb.12239)
- BOYER, T. P. and LEVITUS, S. (1994). *Quality Control and Processing of Historical Oceanographic Temperature, Salinity, and Oxygen Data* **81**. US Department of Commerce, National Oceanic and Atmospheric Administration.
- BYRD, R. H., LU, P., NOCEDAL, J. and ZHU, C. Y. (1995). A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* **16** 1190–1208. [MR1346301 https://doi.org/10.1137/0916069](https://doi.org/10.1137/0916069)
- CHENG, L. and ZHU, J. (2014). Uncertainties of the ocean heat content estimation induced by insufficient vertical resolution of historical ocean subsurface observations. *J. Atmos. Ocean. Technol.* **31** 1383–1396. <https://doi.org/10.1175/JTECH-D-13-00220.1>
- CHILÈS, J.-P. and DELFINER, P. (2012). *Geostatistics: Modeling Spatial Uncertainty*, 2nd ed. *Wiley Series in Probability and Statistics*. Wiley, Hoboken, NJ. [MR2850475 https://doi.org/10.1002/9781118136188](https://doi.org/10.1002/9781118136188)
- CHOI, H. and REIMHERR, M. (2018). A geometric approach to confidence regions and bands for functional parameters. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 239–260. [MR3744720 https://doi.org/10.1111/rssb.12239](https://doi.org/10.1111/rssb.12239)
- CRESSIE, N. and WIKLE, C. K. (2015). *Statistics for Spatio-Temporal Data*. John Wiley & Sons. [MR2848400 https://arxiv.org/abs/1911.13142](https://arxiv.org/abs/1911.13142)
- CRONIE, O., GHORBANI, M., MATEU, J. and YU, J. (2019). Functional marked point processes—a natural structure to unify spatio-temporal frameworks and to analyse dependent functional data. [arXiv:1911.13142](https://arxiv.org/abs/1911.13142).
- DELICADO, P., GIRALDO, R., COMAS, C. and MATEU, J. (2010). Statistics for spatial functional data: Some recent contributions. *Environmetrics* **21** 224–239. [MR2842240 https://doi.org/10.1002/env.1003](https://doi.org/10.1002/env.1003)
- FAN, J. and GIJBELS, I. (1996). *Local Polynomial Modelling and Its Applications*. *Monographs on Statistics and Applied Probability* **66**. CRC Press, London. [MR1383587 https://doi.org/10.1002/9781118136188](https://doi.org/10.1002/9781118136188)
- GAILLARD, F. (2012). ISAS-tool version 6: Method and configuration.
- GRAY, A. R. and RISER, S. C. (2015). A method for multiscale optimal analysis with application to Argo data. *J. Geophys. Res.* **120** 4340–4356. <https://doi.org/10.1002/2014JC010208>
- GREEN, P. J. and SILVERMAN, B. W. (1994). *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. *Monographs on Statistics and Applied Probability* **58**. CRC Press, London. [MR1270012 https://doi.org/10.1002/9781118136188](https://doi.org/10.1002/9781118136188)
- GROMENKO, O., KOKOSZKA, P. and SOJKA, J. (2017). Evaluation of the cooling trend in the ionosphere using functional regression with incomplete curves. *Ann. Appl. Stat.* **11** 898–918. [MR3693551 https://doi.org/10.1214/17-AOAS1022](https://doi.org/10.1214/17-AOAS1022)
- GUINNESS, J. (2021). Gaussian process learning via Fisher scoring of Vecchia’s approximation. *Stat. Comput.* **31** Paper No. 25, 8. [MR4224950 https://doi.org/10.1007/s11222-021-09999-1](https://doi.org/10.1007/s11222-021-09999-1)
- HALL, P., MÜLLER, H.-G. and WANG, J.-L. (2006). Properties of principal component methods for functional and longitudinal data analysis. *Ann. Statist.* **34** 1493–1517. [MR2278365 https://doi.org/10.1214/009053606000000272](https://doi.org/10.1214/009053606000000272)
- HOLTE, J. and TALLEY, L. (2009). A new algorithm for finding mixed layer depths with applications to Argo data and Subantarctic mode water formation. *J. Atmos. Ocean. Technol.* **26** 1920–1939.
- HOLTE, J., TALLEY, L. D., GILSON, J. and ROEMMICH, D. (2017). An Argo mixed layer climatology and database. *Geophys. Res. Lett.* **44** 5618–5626. <https://doi.org/10.1002/2017GL073426>
- HOSODA, S., OHIRA, T. and NAKAMURA, T. (2008). A monthly mean dataset of global oceanic temperature and salinity derived from Argo float observations. *JAMSTEC Report of Research and Development* **8** 47–59. <https://doi.org/10.5918/jamstecr.8.47>
- HOSODA, S., OHIRA, T., SATO, K. and SUGA, T. (2010). Improved description of global mixed-layer depth using Argo profiling floats. *J. Oceanogr.* **66** 773–787.

- HSING, T. and EUBANK, R. (2015). *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators*. Wiley Series in Probability and Statistics. Wiley, Chichester. MR3379106 <https://doi.org/10.1002/9781118762547>
- ISHII, M. and KIMOTO, M. (2009). Reevaluation of historical ocean heat content variations with time-varying XBT and MBT depth bias corrections. *J. Oceanogr.* **65** 287–299.
- JIANG, H. and SERBAN, N. (2012). Clustering random curves under spatial interdependence with application to service accessibility. *Technometrics* **54** 108–119. MR2929427 <https://doi.org/10.1080/00401706.2012.657106>
- JOHNSON, G. C. and BIRNBAUM, A. N. (2017). As El Niño builds, Pacific warm pool expands, ocean gains more heat. *Geophys. Res. Lett.* **44** 438–445. <https://doi.org/10.1002/2016GL071767>
- KELLEY, D., RICHARDS, C. and WG127 SCOR/IAPSO (2017). gsw: Gibbs sea water functions. R package version 1.0-5.
- KING, M. C., STAICU, A.-M., DAVIS, J. M., REICH, B. J. and EDER, B. (2018). A functional data analysis of spatiotemporal trends and variation in fine particulate matter. *Atmos. Environ.* **184** 233–243. <https://doi.org/10.1016/j.atmosenv.2018.04.001>
- KOKOSZKA, P. and REIMHERR, M. (2017). *Introduction to Functional Data Analysis. Texts in Statistical Science Series*. CRC Press, Boca Raton, FL. MR3793167
- KOKOSZKA, P. and REIMHERR, M. (2019). Some recent developments in inference for geostatistical functional data. *Rev. Colombiana Estadíst.* **42** 101–122. MR3919840 <https://doi.org/10.15446/rce.v42n1.77058>
- KUUSELA, M. and STEIN, M. L. (2018). Locally stationary spatio-temporal interpolation of Argo profiling float data. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **474**.
- LEVITUS, S., ANTONOV, J. I., BOYER, T. P., BARANOVA, O. K., GARCIA, H. E., LOCARNINI, R. A., MISHONOV, A. V., REAGAN, J. R., SEIDOV, D. et al. (2012). World ocean heat content and thermocline sea level change (0–2000 m), 1955–2010. *Geophys. Res. Lett.* **39**. <https://doi.org/10.1029/2012GL051106>
- LI, Y. and HSING, T. (2010). Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *Ann. Statist.* **38** 3321–3351. MR2766854 <https://doi.org/10.1214/10-AOS813>
- LI, Y., WANG, N. and CARROLL, R. J. (2013). Selecting the number of principal components in functional data. *J. Amer. Statist. Assoc.* **108** 1284–1294. MR3174708 <https://doi.org/10.1080/01621459.2013.788980>
- LI, H., XU, F., ZHOU, W., WANG, D., WRIGHT, J., LIU, Z. and LIN, Y. (2017). Development of a global gridded Argo data set with Barnes successive corrections. *J. Geophys. Res.* **122**. <https://doi.org/10.1002/2016JC012285>
- LI, G., CHENG, L., ZHU, J., TRENBERTH, K. E., MANN, M. E. and ABRAHAM, J. P. (2020). Increasing ocean stratification over the past half-century. *Nat. Clim. Change* 1–8. <https://doi.org/10.1038/s41558-020-00918-2>
- LINDGREN, F., RUE, H. and LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **73** 423–498. With discussion and a reply by the authors. MR2853727 <https://doi.org/10.1111/j.1467-9868.2011.00777.x>
- LYMAN, J. M. and JOHNSON, G. C. (2013). Estimating global ocean heat content changes in the upper 1800 m since 1950 and the influence of climatology choice. *J. Climate* **27** 1945–1957. <https://doi.org/10.1175/JCLI-D-12-00752.1>
- LYNCH, B. and CHEN, K. (2018). A test of weak separability for multi-way functional data, with application to brain connectivity studies. *Biometrika* **105** 815–831. MR3877867 <https://doi.org/10.1093/biomet/asy048>
- MARTÍNEZ-HERNÁNDEZ, I. and GENTON, M. G. (2020). Recent developments in complex and spatially correlated functional data. *Braz. J. Probab. Stat.* **34** 204–229. MR4093256 <https://doi.org/10.1214/20-BJPS466>
- MCDUGALL, T. J. (2003). Potential enthalpy: A conservative oceanic variable for evaluating heat content and heat fluxes. *J. Phys. Oceanogr.* **33** 945–963. MR2001123 [https://doi.org/10.1175/1520-0485\(2003\)033<0945:PEACOV>2.0.CO;2](https://doi.org/10.1175/1520-0485(2003)033<0945:PEACOV>2.0.CO;2)
- MCDUGALL, T. J. and BARKER, P. M. (2011). Getting started with TEOS-10 and the Gibbs seawater (GSW) oceanographic toolbox. *SCOR/IAPSO WG127* **33** 28.
- MEYSSIGNAC, B., BOYER, T., ZHAO, Z., HAKUBA, M. Z., LANDERER, F. W., STAMMER, D., KÖHL, A., KATO, S., L'ECUYER, T. et al. (2019). Measuring global ocean heat content to estimate the Earth energy imbalance. *Front. Mar. Sci.* **6** 1–31. <https://doi.org/10.3389/fmars.2019.00432>
- MONESTIEZ, P. and NERINI, D. (2008). A cokriging method for spatial functional data with applications in oceanology. In *Functional and Operatorial Statistics. Contrib. Statist.* 237–242. Physica-Verlag/Springer, Heidelberg. MR2490355 https://doi.org/10.1007/978-3-7908-2062-1_36
- NOAA NODC (2019). Ocean climate laboratory, global ocean heat and salt content global anomaly fields. *National Oceanographic Data Center*.
- OWENS, W. B. and WONG, A. P. S. (2009). An improved calibration method for the drift of the conductivity sensor on autonomous CTD profiling floats by θ -S climatology. *Deep-Sea Res., Part 1, Oceanogr. Res. Pap.* **56** 450–457. <https://doi.org/10.1016/j.dsr.2008.09.008>

- PAUTHENET, E., ROQUET, F., MADEC, G., SALLÉE, J.-B. and NERINI, D. (2019). The thermohaline modes of the global ocean. *J. Phys. Oceanogr.* **49** 2535–2552. <https://doi.org/10.1175/JPO-D-19-0120.1>
- RAHMSTORF, S., BOX, J. E., FEULNER, G., MANN, M. E., ROBINSON, A., RUTHERFORD, S. and SCHAFERNICHT, E. J. (2015). Exceptional twentieth-century slowdown in Atlantic Ocean overturning circulation. *Nat. Clim. Change* **5** 475–480. <https://doi.org/10.1038/nclimate2554>
- RAMSAY, J. and SILVERMAN, B. W. (2013). *Functional Data Analysis. Springer Series in Statistics*. Springer Science & Business Media, New York. [MR2168993](https://doi.org/10.1007/978-1-4939-9738-9)
- RAMSAY, J., WICKHAM, H., GRAVES, S. and HOOKER, G. (2018). *fda: Functional data analysis*. R package version 2.4.8.
- RODRÍGUEZ, A., DUNSON, D. B. and GELFAND, A. E. (2009). Bayesian nonparametric functional data analysis through density estimation. *Biometrika* **96** 149–162. [MR2482141 https://doi.org/10.1093/biomet/asn054](https://doi.org/10.1093/biomet/asn054)
- ROEMMICH, D. and GILSON, J. (2009). The 2004–2008 mean and annual cycle of temperature, salinity, and steric height in the global ocean from the Argo program. *Prog. Oceanogr.* **82** 81–100.
- ROEMMICH, D., GOULD, W. J. and GILSON, J. (2012). 135 years of global ocean warming between the Challenger expedition and the Argo programme. *Nat. Clim. Change* **2** 425–428. <https://doi.org/10.1038/nclimate1461>
- ROEMMICH, D., CHURCH, J., GILSON, J., MONSELESAN, D., SUTTON, P. and WIJFFELS, S. (2015). Unabated planetary warming and its ocean structure since 2006. *Nat. Clim. Change* **5** 240–245. <https://doi.org/10.1038/nclimate2513>
- ROMANO, E., BALZANELLA, A. and VERDE, R. (2017). Spatial variability clustering for spatially dependent functional data. *Stat. Comput.* **27** 645–658. [MR3613590 https://doi.org/10.1007/s11222-016-9645-2](https://doi.org/10.1007/s11222-016-9645-2)
- RUIZ-MEDINA, M. D. (2011). Spatial autoregressive and moving average Hilbertian processes. *J. Multivariate Anal.* **102** 292–305. [MR2739116 https://doi.org/10.1016/j.jmva.2010.09.005](https://doi.org/10.1016/j.jmva.2010.09.005)
- RUPPERT, D. (2002). Selecting the number of knots for penalized splines. *J. Comput. Graph. Statist.* **11** 735–757. [MR1944261 https://doi.org/10.1198/106186002321018768](https://doi.org/10.1198/106186002321018768)
- SCHMIDTKO, S., JOHNSON, G. C. and LYMAN, J. M. (2013). MIMOC: A global monthly isopycnal upper-ocean climatology with mixed layers. *J. Geophys. Res.* **118** 1658–1672.
- SONG, J. J. and MALLICK, B. (2019). Hierarchical Bayesian models for predicting spatially correlated curves. *Statistics* **53** 196–209. [MR3900086 https://doi.org/10.1080/02331888.2018.1547905](https://doi.org/10.1080/02331888.2018.1547905)
- STAIU, A.-M., CRAINICEANU, C. M. and CARROLL, R. J. (2010). Fast methods for spatially correlated multilevel functional data. *Biostat.* **11** 177–194. <https://doi.org/10.1093/biostatistics/kxp058>
- STAIU, A.-M., CRAINICEANU, C. M., REICH, D. S. and RUPPERT, D. (2012). Modeling functional data with spatially heterogeneous shape characteristics. *Biometrics* **68** 331–343. [MR2959599 https://doi.org/10.1111/j.1541-0420.2011.01669.x](https://doi.org/10.1111/j.1541-0420.2011.01669.x)
- STEIN, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging. Springer Series in Statistics*. Springer, New York. [MR1697409 https://doi.org/10.1007/978-1-4612-1494-6](https://doi.org/10.1007/978-1-4612-1494-6)
- TALLEY, L. D., PICKARD, G. L., EMERY, W. J. and SWIFT, J. H. (2011). *Descriptive Physical Oceanography*, 6th ed. Academic Press, Boston, MA. <https://doi.org/10.1016/B978-0-7506-4552-2.10004-6>
- THOMSON, R. E. and EMERY, W. J. (2014). *Data Analysis Methods in Physical Oceanography*, 3rd edition. Elsevier, Boston, MA. <https://doi.org/10.1016/B978-0-12-387782-6.00003-X>
- TOWNS, J., COCKERILL, T., DAHAN, M., FOSTER, I., GAITHER, K., GRIMSHAW, A., HAZLEWOOD, V., LATHROP, S., LIFKA, D. et al. (2014). XSEDE: Accelerating scientific discovery. *Comput. Sci. Eng.* **16** 62–74. <https://doi.org/10.1109/MCSE.2014.80>
- UDAYA BHASKAR, T. V. S., RAVICHANDRAN, M. and DEVENDER, R. (2007). An operational objective analysis system at INCOIS for generation of Argo value added products.
- WAHBA, G. (1990). *Spline Models for Observational Data. CBMS-NSF Regional Conference Series in Applied Mathematics* **59**. SIAM, Philadelphia, PA. [MR1045442 https://doi.org/10.1137/1.9781611970128](https://doi.org/10.1137/1.9781611970128)
- WIENS, A., NYCHKA, D. and KLEIBER, W. (2020). Modeling spatial data using local likelihood estimation and a Matérn to spatial autoregressive translation. *Environmetrics* **31** A1E2652–14. [MR4151871 https://doi.org/10.1002/env.2652](https://doi.org/10.1002/env.2652)
- WOOD, S. N. (2006). Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics* **62** 1025–1036. [MR2297673 https://doi.org/10.1111/j.1541-0420.2006.00574.x](https://doi.org/10.1111/j.1541-0420.2006.00574.x)
- YAO, F., MÜLLER, H.-G. and WANG, J.-L. (2005). Functional data analysis for sparse longitudinal data. *J. Amer. Statist. Assoc.* **100** 577–590. [MR2160561 https://doi.org/10.1198/016214504000001745](https://doi.org/10.1198/016214504000001745)
- YARGER, D. (2020a). Shiny applications accompanying “A functional-data approach to the Argo data”. Available at <https://sites.google.com/a/umich.edu/argostatistics/home/fdaper>.
- YARGER, D. (2020b). Code accompanying “A functional-data approach to the Argo data”. Available at <https://github.com/dyarger/argofda>.
- YARGER, D., STOEVE, S. and HSING, T. (2022). Supplement to “A functional-data approach to the Argo data.” <https://doi.org/10.1214/21-AOAS1477SUPPA>, <https://doi.org/10.1214/21-AOAS1477SUPPB>

- ZHANG, H. and LI, Y. (2020). Unified principal component analysis for sparse and dense functional data under spatial dependency. [arXiv:2006.13489](https://arxiv.org/abs/2006.13489).
- ZHANG, L., BALADANDAYUTHAPANI, V., ZHU, H., BAGGERLY, K. A., MAJEWSKI, T., CZERNIAK, B. A. and MORRIS, J. S. (2016). Functional CAR models for large spatially correlated functional datasets. *J. Amer. Statist. Assoc.* **111** 772–786. [MR3538704 https://doi.org/10.1080/01621459.2015.1042581](https://doi.org/10.1080/01621459.2015.1042581)
- ZHOU, L., HUANG, J. Z., MARTINEZ, J. G., MAITY, A., BALADANDAYUTHAPANI, V. and CARROLL, R. J. (2010). Reduced rank mixed effects models for spatially correlated hierarchical functional data. *J. Amer. Statist. Assoc.* **105** 390–400. [MR2757206 https://doi.org/10.1198/jasa.2010.tm08737](https://doi.org/10.1198/jasa.2010.tm08737)