

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Spatial Statistics

journal homepage: www.elsevier.com/locate/spasta

Spatially penalized registration of multivariate functional data

Xiaohan Guo^a, Sebastian Kurtek^a, Karthik Bharath^{b,*}^a Department of Statistics, The Ohio State University, 1958 Neil Avenue, Columbus, OH 43210, USA^b School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, UK

ARTICLE INFO

Article history:

Received 14 February 2023

Received in revised form 25 April 2023

Accepted 7 June 2023

Available online 22 June 2023

Keywords:

Elastic metric

Functional random field

Phase trace-variogram

Warping invariance

ABSTRACT

Registration of multivariate functional data involves handling of both cross-component and cross-observation phase variations. Allowing for the two phase variations to be modelled as general diffeomorphic time warpings, in this work we focus on the hitherto unconsidered setting where phase variation of the component functions are spatially correlated. We propose an algorithm to optimize a metric-based objective function for registration with a novel penalty term that incorporates the spatial correlation between the component phase variations through a kriging prediction of an appropriate phase random field. The penalty term encourages the overall phase at a particular location to be similar to the spatially weighted average phase in its neighbourhood, and thus engenders a regularization that prevents over-alignment. Utility of the registration method, and its superior performance compared to methods that fail to account for the spatial correlation, is demonstrated through performance on simulated examples and two multivariate functional datasets pertaining to electroencephalogram signals and ozone concentration functions. The generality of the framework opens up the possibility for extension to settings involving different forms of correlation between the component functions and their phases.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

* Corresponding author.

E-mail addresses: gxh1280@gmail.com (X. Guo), kurtek.1@stat.osu.edu (S. Kurtek), Karthik.Bharath@nottingham.ac.uk (K. Bharath).

<https://doi.org/10.1016/j.spasta.2023.100760>

2211-6753/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Modern functional datasets, such as longitudinal records, medical imaging signals or geometric shapes, contain confounded amplitude and phase variations (Srivastava et al., 2011). Amplitude refers to the magnitude and shape of functions, e.g., number, order and scale of extrema, whereas phase refers to the timing of amplitude features. The adverse effects of ignoring phase variation during statistical analysis of functional data are now well-understood, including blurring of the structure in the data, e.g., during exploratory data analysis via visualization or summarization, inflation of variance measures, and potential for misleading conclusions based on statistical inference (Marron et al., 2015). The underlying issues are further exaggerated when the amplitude and/or phase components in a functional dataset are spatially dependent (Guo et al., 2022). For multivariate functional data $\mathbb{R} \supset I \mapsto F_i = (f_{i1}, \dots, f_{iK})^\top$ ($F_i(t) = (f_{i1}(t), \dots, f_{iK}(t))^\top \in \mathbb{R}^K \forall t \in I$) consisting of multiple correlated, univariate functional components $f_{ij} : I \rightarrow \mathbb{R}$, $j = 1, \dots, K$, the notion of phase variation can be decomposed into two types: (i) *cross-observation phase*, which is common across all K component univariate functions f_{ij} within an observation F_i , and (ii) *cross-component phase* within each observed unit. The goal of a registration procedure for multivariate functional data is to estimate the cross-observation and cross-component phase variations. Applying these phase estimates to the observed multivariate functional data results in time synchronization across observations and components (or alignment); the aligned data can then be used to quantify amplitude variation. In many cases, cross-observation phase can be assumed to be independent. However, since the components within a multivariate functional observation are often correlated, one must account for this additional structure in the registration procedure. Further, when the components are spatially indexed, the additional structure is engendered in spatial correlation among cross-component phase. In this case, each multivariate functional observation F_i can be viewed as an independent realization of a functional random field.

Registration procedures for multivariate functional data are driven by assumptions on the two types of temporal variation. At the two extremes are procedures which assume that cross-component variation within each F_i are either uncorrelated or identical. The former case is quite commonly considered by researchers in neuroimaging (Makeig et al., 2007; Tsai et al., 2014; Zhao et al., 2020), motivated mainly by the simplicity of independent componentwise registration. The latter case, referred to as universal registration, is common in shape analysis of K -dimensional curves $\{F_i\}$ (Kurtek et al., 2012), wherein one assumes that cross-component phase variation does not exist and a common warping function is estimated for all components within the same observation (Olsen et al., 2016). Thus, neither of the two extreme cases directly addresses the registration problem when cross-component variation is non-negligible.

Literature on registration methods for multivariate functional data that lie between the two extreme cases is sparse. Noteworthy exceptions within the statistics literature are recent work by Carroll et al. (2021), Carroll and Müller (2023) and Park and Ahn (2017), where specific forms of cross-component temporal variation were considered. The first two works restrict attention to the situation where *all* component functions f_{ij} for $j = 1, \dots, K$ are assumed to have the same shape: for each observation i , Carroll et al. (2021) assume that time-warped components $f_{ij}(t)$ arise through time shifts $t \mapsto (t - \tau)$ of the same function, say g , while Carroll and Müller (2023) extend this to the case where the time shift map is replaced by a general diffeomorphic warping $t \mapsto \gamma(t)$ of I ; a formal definition of diffeomorphic warping is given later in Section 3. Motivated mainly by clustering, Park and Ahn (2017) proposed a conditional observation-specific registration procedure to extract relevant features for clustering.

It is quite common to encounter multivariate functional datasets wherein the component functions do not possess the same shape and/or when correlations between component functions need to be explicitly incorporated into the registration procedure. As an example, consider data generated from electroencephalogram (EEG) signals, specifically pertaining to event-related potentials measured during EEG tests. Event-related potentials are small voltages that reflect the electrical activity on the scalp in response to specific stimuli received by study participants (Sur and Sinha, 2009). For the i th observation, EEG signals f_{ij} are densely recorded on a certain time domain I using a set of electrodes placed at K different locations on the scalp, and together comprise a multivariate

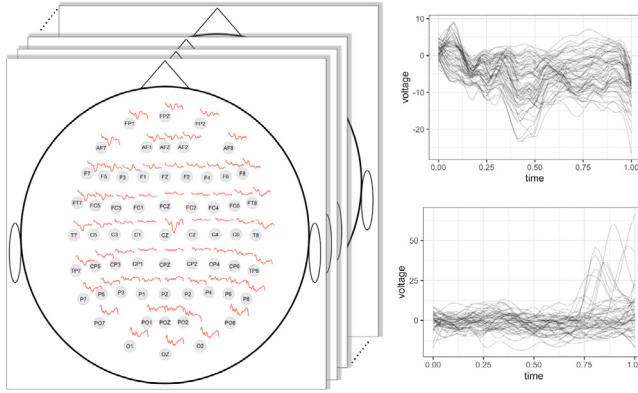


Fig. 1. Left: Example of a multi-trial EEG dataset with 61 electrodes from an alcoholism study. Each slide represents a single trial in the study, while the red functions plotted at the locations of electrodes (grey circles) constitute the components. Top Right: 61 signals from a single trial with considerable cross-component phase variation. Bottom Right: Signals collected at electrode AFZ across 50 trials with considerable cross-observation phase variation.

functional observation F_i . In the left panel of Fig. 1, we display the observed EEG signals at the (projected) two-dimensional electrode locations on a toy map of the scalp. The entire collection of signals recorded during a single EEG trial is shown as a single slide and corresponds to an observed multivariate functional data unit.

Evidently, the component functions f_{ij} and f_{ik} at electrode locations j and k during the i th trial are spatially correlated due to functional connectivity between brain regions. As such, each multivariate functional observation F_i can be viewed as a realization of a functional random field, wherein electrode locations define the spatial index for each component function within F_i . Latency of the brain’s responses to stimuli varies across trials for a single subject (and among different subjects) resulting in cross-observation phase variation (Wang et al., 2001). At the same time, different response lags of different brain regions to the presented stimuli (Stam et al., 2007) result in cross-component phase variation. The right two panels of Fig. 1 offer an illustration of cross-observation (bottom) and cross-component (top) phase variation among EEG signals. To account for the two different types of phase variation in multi-trial EEG, an appropriate registration procedure that accounts for spatially-correlated cross-component phase is required. Specifically, the procedure should enable simultaneous registration of all components f_{ij} , while taking into account the spatial dependence structure between the cross-component phases. This cannot be achieved by implementing independent registration of f_{1j}, \dots, f_{nj} for each component j , since such a procedure results in additional cross-component phase variation in the estimated average components due to each component being treated independently. We note this phenomenon in the right panel of Fig. 6 (and the associated discussion) in Section 6 during our detailed analysis of the EEG data.

Motivated by the type of registration task associated with the EEG data, we propose a registration procedure that exploits correlated phase variation between component functions $\{f_{ij}\}$ of multivariate functional data $\{F_i\}$ and offers a compromise between the two extreme cases of independent componentwise and universal registration methods alluded to earlier. Our focus is on multivariate functional data with component functions f_{ij} whose overall phase variation, represented through a time-warping function $\gamma_{ij} : I \rightarrow I$, is a combination $\gamma_{ij} = \xi_j \circ \alpha_i$ of cross-observation phase $\alpha_i : I \rightarrow I$ and spatially correlated cross-component phase $\xi_j : I \rightarrow I$ (\circ denotes function composition). The flexibility afforded, and desired, by such a general phase specification while registering multivariate functional data with correlated components, such as the EEG data, comes at a cost: the individual phase components α_i and ξ_j cannot be decoupled and estimated individually, and incorporating spatial information present in the latent $\{\xi_j\}$ when estimating γ_{ij} is hugely challenging. In view of this, our main contributions are as follows.

- (i) We propose a metric-based penalized registration method for multivariate functional data with a penalty term defined using spatial correlations between the cross-component phases $\{\xi_j\}$, using which the overall phases $\{\gamma_{ij}\}$ are estimated by eliminating cross-observation phase variation $\{\alpha_i\}$. Our framework utilizes the extended Fisher–Rao metric, which has desirable (and necessary) mathematical properties for the registration problem.
- (ii) The penalty term is defined using a spatially weighted combination of the (estimated) cross-component phases $\{\xi_j\}$, and is, owing to the invariant property of the metric to (simultaneous) time warping, impervious to the cross-observation phases $\{\alpha_i\}$; this allows us to entirely circumvent having to estimate the $\{\alpha_i\}$.
- (iii) With respect to cross-observation registration, we demonstrate clear superiority of our method over ones that fail to account for the spatial correlation under a variety of simulation settings and two real data examples.

The rest of this paper is organized as follows. Section 2 details the registration problem for multivariate functional data whose component functions have spatially correlated phase variation. Section 3 briefly reviews the geometric elastic functional data analysis framework, including univariate pairwise and multiple registration, and registration of multivariate functions with common cross-component phase. Section 4 details the proposed spatially penalized registration objective function that accounts for spatial correlation in cross-component phase, and provides an algorithm (Algorithm 1) for its optimization. Section 5 presents results of simulation studies; in Section 6, we apply the proposed registration approach to EEG and ozone data. Numerical assessment of convergence of the registration algorithm and comparative studies with an alternative penalized approach are available in the supplement.

2. The registration problem

We can view multivariate functional data $\{F_1, \dots, F_n\}$, where each $F_i(t) = (f_{i1}(t), \dots, f_{iK}(t))^T$, $t \in [0, 1]$, $i = 1, \dots, n$, has K correlated components, with cross-observation and spatially correlated cross-component temporal variation, as being generated from the model

$$F_i \circ \mathbf{w}_i(t) = \mu(t) + E_i(t), \quad t \in [0, 1], \quad i = 1, \dots, n, \tag{1}$$

where $\mu = (\mu_1, \dots, \mu_K)$, with $\mu_j : [0, 1] \rightarrow \mathbb{R}$, is a deterministic mean/template multivariate function, $E_i = (e_{i1}, \dots, e_{iK})^T$ are realizations of an error process with spatially correlated component functions $\{e_{ij} : [0, 1] \rightarrow \mathbb{R}\}$, and $\mathbf{w}_i = (\gamma_{i1}, \dots, \gamma_{iK})$ is a vector of monotone increasing, end-points preserving, time warping functions assuming values in $\Gamma := \{\gamma : [0, 1] \rightarrow [0, 1] \mid \gamma(0) = 0, \gamma(1) = 1, \dot{\gamma} > 0\}$, where $\dot{\gamma}$ is the time derivative of γ . Each \mathbf{w}_i consists of K random warping functions, accounting for the temporal variability across components in a multivariate functional observation F_i ; $\{\mathbf{w}_i\}$ define overall phase variation in the data. Then, $F_i \circ \mathbf{w}_i = (f_{i1} \circ \gamma_{i1}, \dots, f_{iK} \circ \gamma_{iK})^T$ represents componentwise warping, where \circ denotes function composition; $\{F_i \circ \mathbf{w}_i\}$ defines amplitude variation in the data. Note that Γ is a group with operation \circ , identity element $\gamma_{id}(t) = t$, and the function inverse as the group inverse. An essential implication of the group structure for warping functions is that the inverse of any $\gamma \in \Gamma$ is also an element of Γ ; this property will be used frequently in the development of the proposed registration framework. Further, it allows for a rigorous definition of the amplitude of a functional observation via an equivalence class as will be seen in subsequent sections. We will interchangeably refer to γ as a warping function or phase depending on the context.

The characterizing feature of the problem is that, for $i = 1, \dots, n$, $j = 1, \dots, K$, the warping functions $\gamma_{ij} = \xi_{\mathbf{s}_j} \circ \alpha_i$, where $\xi_{\mathbf{s}_1}, \dots, \xi_{\mathbf{s}_K}$ are spatially correlated warping functions at K spatial locations $\mathbf{s}_1, \dots, \mathbf{s}_K$ within a spatial domain $\mathcal{D} \subset \mathbb{R}^p$, with $\xi_{\mathbf{s}_j}, \alpha_i \in \Gamma$ for each $i = 1, \dots, n$, $j = 1, \dots, K$. The phase γ_{ij} of a component function f_{ij} is thus represented as a combination of warping functions corresponding to cross-observation and spatially correlated cross-component temporal variations. The index j will be used to denote the spatial location \mathbf{s}_j , i.e., ξ_j will be used in place of $\xi_{\mathbf{s}_j}$, and the two notations will be used interchangeably.

Under the above setup, we can interpret $\alpha_1, \dots, \alpha_n$ as representing cross-observation temporal variation within $\{F_i\}$ and ξ_1, \dots, ξ_K as representing spatially correlated cross-component temporal variation within $\{F_i\}$. The two-fold objective is to, jointly:

1. register F_1, \dots, F_n by estimating $\{\gamma_{ij} : i = 1, \dots, n, j = 1, \dots, K\}$ in a manner that ensures for every $i = 1, \dots, n$, the warping functions γ_{ij} and γ_{ik} in nearby spatial locations \mathbf{s}_j and \mathbf{s}_k in \mathcal{D} are similar, owing to similar cross-component warpings ξ_j and ξ_k ;
2. estimate components μ_1, \dots, μ_K of the template μ .

In the context of the EEG application, e.g., in the i th trial for a single subject, we desire signals f_{ij} and f_{ik} obtained from spatially proximate electrode locations \mathbf{s}_j and \mathbf{s}_k to have similar phase variation γ_{ij} and γ_{ik} , with respect to template components μ_j and μ_k . The requirement is compatible with findings by Stam et al. (2007) on different response lags of spatially disparate brain regions to presented stimuli.

When the template μ is unknown, as is typical in practice, for each pair (i, j) , the warping functions α_i and ξ_j are individually not identifiable in a nonparametric (infinite-dimensional) specification of the class Γ with error E_i without restrictions. It is hence not possible to estimate both α_i and ξ_j since their respective variabilities are confounded with the error E_i ; this implies that it is not possible to decompose the overall phase variation given by $\{\gamma_{ij}\}$ in $\{F_i\}$ into the cross-observation and cross-component phase variations.

In view of this, our registration procedure will first estimate the spatially correlated cross-component phase for each observation $i = 1, \dots, n$ and then use them to estimate γ_{ij} . Evidently then, the cross-component warping functions ξ_1, \dots, ξ_K will depend on i . The key challenge arises from the need to iterate between computing the template components μ_1, \dots, μ_K and warping functions $\{\gamma_{ij}\}$ by explicitly incorporating spatial dependence between the, unobserved but observation dependent, cross-component warping functions $\xi_{i1}, \dots, \xi_{iK}$. In Section 4.2, we consider a simplified geometry of Γ that will further clarify the dependence of the cross-component phases on the observations.

3. Overview of elastic functional data registration

The proposed algorithm is based on a componentwise spatially-penalized registration of f_{1j}, \dots, f_{nj} using the metric-based elastic functional data analysis framework (Srivastava et al., 2011; Srivastava and Klassen, 2016) of univariate functions. The metric enables us to propose a novel penalty that quantifies the spatial correlation between the unobserved cross-component phases $\{\xi_j\}$, unaffected by the cross-observation phase $\{\alpha_i\}$. We begin with a brief review of the elastic framework, and for later comparison with the proposed algorithm, also review universal registration of curves under this framework (Srivastava and Klassen, 2016), where all components are assumed to have identical temporal variation.

3.1. Univariate functions

We consider the representation space of univariate functional data objects to be $\mathcal{F} := \{f : [0, 1] \rightarrow \mathbb{R} \mid f \text{ is absolutely continuous}\}$. As stated earlier, the group of warping functions representing phase is $\Gamma = \{\gamma : [0, 1] \rightarrow [0, 1] \mid \gamma(0) = 0, \gamma(1) = 1, \dot{\gamma} > 0\}$. For any $f \in \mathcal{F}, \gamma \in \Gamma$, the warping of f by γ is given by the group action of composition, $f \circ \gamma$, i.e., $(f \circ \gamma)(t) = f(\gamma(t)) \forall t \in [0, 1]$. The group-theoretic formulation of phase further enables a definition of the amplitude of a function f as the equivalence class $[f] := \{f \circ \gamma \mid \gamma \in \Gamma\} \subset \mathcal{F}$, known as its orbit under the action of Γ ; thus, $f \circ \gamma \in [f]$ has the same amplitude as f for each $\gamma \in \Gamma$. In other words, all possible warpings of a function f are unified by a single equivalence class, which uniquely represents the function’s amplitude. The amplitude space then is the quotient space $\mathcal{F}/\Gamma := \{[f] \mid f \in \mathcal{F}\}$.

Separating amplitude and phase requires a metric on the amplitude space \mathcal{F}/Γ . A convenient way to define one is through a metric d on \mathcal{F} that is invariant to simultaneous warping: for every $\gamma \in \Gamma, d(f_1, f_2) = d(f_1 \circ \gamma, f_2 \circ \gamma)$. The standard \mathbb{L}^2 metric fails to be invariant and Srivastava et al. (2011) thus proposed to use the extended Fisher–Rao (eFR) metric. While direct use of this metric for amplitude-phase separation is difficult in practice, the square-root slope transform can be used to flatten the complicated eFR metric on \mathcal{F} to the standard \mathbb{L}^2 metric on the transformed space. The

transform maps $f \mapsto Q(f) = q := \dot{f}|\dot{f}|^{-1/2}$ (\dot{f} is the time derivative of f). Given $f(0)$, Q is bijective with inverse $Q^{-1}(q, f(0))(t) = f(t) = f(0) + \int_0^t q(u)|q(u)|du$. Henceforth, for any $f \in \mathcal{F}$, we will refer to $q = Q(f)$ as its square-root slope function (SRSF).

The transformed space $Q(\mathcal{F})$ is a subset of $\mathbb{L}^2([0, 1], \mathbb{R})$, and is denoted by \mathcal{Q} . Under Q , the eFR metric on \mathcal{F} maps to the standard \mathbb{L}^2 metric on \mathcal{Q} , and thus analysis of SRSFs can be carried out using standard Hilbert space machinery. Warping of $f \in \mathcal{F}$ by $\gamma \in \Gamma$ induces the warping action $q \odot \gamma := (q \circ \gamma)\dot{\gamma}^{1/2}$ on \mathcal{Q} equipped with the \mathbb{L}^2 metric, and the action is by isometries, i.e., for $q_1, q_2 \in \mathcal{Q}$ and $\gamma \in \Gamma$, $\|q_1 - q_2\| = \|q_1 \odot \gamma - q_2 \odot \gamma\|$ ($\|\cdot\|$ denotes the \mathbb{L}^2 norm); the action is also norm preserving, i.e., $\|q \odot \gamma\| = \|q\|$ for every $\gamma \in \Gamma, q \in \mathcal{Q}$.

3.1.1. Pairwise registration

Given two functions $f_1, f_2 \in \mathcal{F}$, amplitude-phase separation through pairwise alignment of f_2 to f_1 , or vice versa, is formulated as the determination of the relative phase of f_2 with respect to f_1 ($q_1 = Q(f_1), q_2 = Q(f_2)$):

$$\hat{\gamma} = \arg \min_{\gamma \in \Gamma} \|q_1 - q_2 \odot \gamma\|^2 = \arg \min_{\gamma \in \Gamma} \int_0^1 |q_1(t) - (q_2 \odot \gamma)(t)|^2 dt. \tag{2}$$

The minimization problem in (2) is typically solved using the dynamic programming algorithm. To further regularize pairwise registration, one can instead solve the penalized optimization problem given by

$$\hat{\gamma} = \arg \min_{\gamma \in \Gamma} \{ \|q_1 - q_2 \odot \gamma\|^2 + \lambda \|\sqrt{\dot{\gamma}} - 1\|^2 \}, \tag{3}$$

where λ is the regularization parameter. The penalty is defined as the squared \mathbb{L}^2 distance between the SRSFs of γ and the identity warping function γ_{id} , where $\gamma_{id}(t) = t \forall t$. It is evident that, depending on the magnitude of λ , this penalty forces the estimated phase to be close to the identity element, i.e., no warping.

3.1.2. Multiple registration

Amplitude-phase separation for a sample of functions $f_i, i = 1, \dots, n, n > 2$ is carried out with respect to a common template function that must also be estimated. Let $q_i, i = 1, \dots, n$ denote the corresponding SRSFs. A template for multiple registration is estimated via

$$\hat{\mu} = \arg \min_{\mu \in \mathcal{Q}} \sum_{i=1}^n \min_{\gamma_i \in \Gamma} \|\mu - q_i \odot \gamma_i\|^2. \tag{4}$$

Then, multiple registration is carried out via pairwise registration of each $q_i, i = 1, \dots, n$ with respect to $\hat{\mu}$ using (2), resulting in relative phases $\hat{\gamma}_i, i = 1, \dots, n$. The minimizer $\hat{\mu}$ of (4) is not unique, with any element of the orbit $[\hat{\mu}]$ resulting in the same value of the cost function due to the isometric action of Γ . Thus, for identifiability, we select $\hat{\mu} \in [\hat{\mu}]$ such that the relative phases $\hat{\gamma}_i, i = 1, \dots, n$ average to γ_{id} ; see [Srivastava et al. \(2011\)](#) for details.

3.2. Componentwise and universal registration

Let $\tilde{\mathcal{F}} = \{F : [0, 1] \rightarrow \mathbb{R}^K \mid F \text{ is absolutely continuous}\}$ denote the space of multivariate functional data. Each multivariate function F contains K univariate components $f_1, \dots, f_K \in \mathcal{F}$. Independent componentwise registration of multivariate functional data applies multiple registration given by (4) to functions in each component independently. This fails to account for correlation between γ_{ij} and γ_{ik} for every $i = 1, \dots, n$ and $j, k = 1, \dots, K$.

At the other end of the spectrum is universal registration, which treats each F_i as a parameterized curve $t \mapsto (f_{i1}(t), \dots, f_{iK}(t))^T$ in \mathbb{R}^K , and thus assumes the same relative phase for all components. Registration under such a setup is again available through the elastic framework under a suitable transformation. With a slight abuse in notation, let $Q(F) = q = \dot{F}|\dot{F}|^{-1/2}$ denote the SRSF of F , where \dot{F} is the componentwise time derivative of F and $|\cdot|$ is the Euclidean norm in \mathbb{R}^K ; each SRSF in this

case is a function $q : [0, 1] \rightarrow \mathbb{R}^K$ and the space of such SRSFs is denoted by $\tilde{\mathcal{Q}} \subset \mathbb{L}^2([0, 1], \mathbb{R}^K)$. Then, given $q_1 = Q(F_1)$ and $q_2 = Q(F_2)$, the relative phase of F_2 with respect to F_1 is given by

$$\hat{\gamma} = \arg \min_{\gamma \in \Gamma} \|q_1 - q_2 \odot \gamma\|^2 = \arg \min_{\gamma \in \Gamma} \int_0^1 |q_1(t) - (q_2 \odot \gamma)(t)|^2 dt, \tag{5}$$

where again $|\cdot|$ denotes the Euclidean norm in \mathbb{R}^K and $q \odot \gamma$ is applied componentwise. A penalized version of registration can also be implemented in this case by appropriately adapting the optimization problem defined in (3). Further, multiple registration, via estimation of a template $\hat{\mu} \in \tilde{\mathcal{Q}}$, and pairwise registration of each function to $\hat{\mu}$ via (5), follows the approach defined for univariate functional data.

4. Registration of multivariate functions with spatially dependent cross-component phase

We propose a penalized multiple registration method for multivariate functional data wherein cross-component phase in each observation is spatially correlated. The proposed algorithm is based on a *non-trivial* extension of the elastic functional data framework to account for the spatial correlations.

Let $F_i = (f_{i1}, \dots, f_{iK})^\top \in \tilde{\mathcal{F}}$, $i = 1, \dots, n$ denote a multivariate functional data sample; we assume that F_i , $i = 1, \dots, n$ are independent. Each component function f_{ij} , $i = 1, \dots, n$, $j = 1, \dots, K$ is assumed to be an element of \mathcal{F} , and the components f_{ij} , $j = 1, \dots, K$, for a fixed i , have spatially dependent phase. Further, let $q_{ij} = Q(f_{ij}) \in \mathcal{Q}$, $i = 1, \dots, n$, $j = 1, \dots, K$ denote the SRSFs of the component functions in each observation. Registration of multivariate functional data requires estimation of the overall phase (composition of cross-observation and spatially correlated cross-component phase) for each component in each observation, $\gamma_{ij} \in \Gamma$, $i = 1, \dots, n$, $j = 1, \dots, K$. This facilitates simultaneous synchronization of the component functions across i and j .

As a compromise between the two extreme settings of independent componentwise and universal registration described in Section 3.2, we propose a spatially penalized registration approach that takes spatial cross-component phase correlation into account, but allows each component to have its own phase. As described in Section 2, within each F_i , the dependence between the component functions f_{i1}, \dots, f_{iK} arises through spatial correlation in the cross-component phases ξ_1, \dots, ξ_K .

4.1. Spatially penalized componentwise registration

Our approach is to carry out cross-observation registration through a modification of the multiple registration procedure in (4) using a spatially-informed penalty term. For a fixed function component j , the registration of functions f_{ij} , $i = 1, \dots, n$, with SRSFs q_{ij} , $i = 1, \dots, n$, amounts to determination of γ_{ij} , $i = 1, \dots, n$. However, for each $i = 1, \dots, n$, we note that γ_{ij} is spatially correlated with γ_{il} , $l \neq j$, $l = 1, \dots, K$, through the spatial correlation between the latent ξ_j and ξ_l , $l \neq j$. We thus consider a penalized modification of (4) wherein the penalty term for each $i = 1, \dots, n$ depends on the phases $\gamma_{i1}, \dots, \gamma_{i(j-1)}, \gamma_{i(j+1)}, \dots, \gamma_{iK}$. Our choice is to define the penalty term using (an estimate of) the conditional mean of γ_{ij} given $\{\gamma_{il}\}_{l \neq j}$: we wish to discourage γ_{ij} from assuming values that are far away from its spatially weighted conditional mean. We accordingly consider the following optimization problem for each component $j = 1, \dots, K$:

$$(\hat{\gamma}_{1j}, \dots, \hat{\gamma}_{nj}, \hat{\mu}_j) = \arg \min_{\gamma_{1j}, \dots, \gamma_{nj} \in \Gamma, \mu_j \in \mathcal{Q}} \sum_{i=1}^n \{ \|\mu_j - q_{ij} \odot \gamma_{ij}\|^2 + \lambda d^2(\gamma_{ij}, \tilde{\gamma}_{ij}) \}, \tag{6}$$

where μ_j is the template for registration in component j , $\lambda > 0$ is a penalty parameter, and d is a distance on Γ ; the warping function $\tilde{\gamma}_{ij}$ is an estimate of the conditional mean of γ_{ij} given $\{\gamma_{il}\}_{l \neq j}$, which takes into account the spatial correlation among cross-component phases; $\tilde{\gamma}_{ij}$ is defined next in Section 4.2.

The first term in (6) provides a measure of synchronization for component j , across observations i , with respect to the template μ_j . The second term, a penalty on phase, measures the distance between the estimated phase γ_{ij} and a target $\tilde{\gamma}_{ij}$ determined by the spatially correlated phase in the

other components; the regularization penalty attempts to preserve the phase-induced correlation structure in the aligned components. If the function components are assumed to be independent, and $\tilde{\gamma}_{ij} = \gamma_{id}$ (identity warping) for all i and j , the proposed approach is equivalent to K independent univariate penalized registration problems, as specified in (3), with μ_j acting as the template for each component.

4.2. Penalty using spatial model for cross-component phase

We first discuss the choice of distance d on Γ in the penalty term in (6). Elements of Γ can be viewed as differentiable, increasing distribution functions of random variables on $[0, 1]$, and thus constitute a nonlinear, convex set. A simplified geometric structure compatible with the \mathbb{L}^2 norm preserving action under the operation \odot is available under the SRSF transform, where $\gamma \rightarrow Q(\gamma) = \dot{\gamma}^{1/2} =: \psi$. The SRSF map Q is bijective, and each ψ corresponds to a square-root probability density since $\int_0^1 \psi^2(t)dt = 1$. As a consequence, the set $\Psi := \{Q(\gamma) =: \psi : [0, 1] \rightarrow \mathbb{R}_+ \mid \gamma \in \Gamma\}$ can be identified with the positive orthant of the unit sphere in $\mathbb{L}^2([0, 1], \mathbb{R})$. While the natural candidate for distance d on Γ is the intrinsic arc-length distance on Ψ , we use the extrinsic distance

$$d(\gamma_1, \gamma_2) := \|Q(\gamma_1) - Q(\gamma_2)\| = \|\psi_1 - \psi_2\|,$$

where $\|\cdot\|$ is the standard \mathbb{L}^2 norm. Our choice is linked to the choice of $\tilde{\gamma}_{ij}$ in the penalty term. Observe that with d on Γ defined as above, we have that, for every $\alpha, \gamma_1, \gamma_2 \in \Gamma$,

$$d(\gamma_1 \circ \alpha, \gamma_2 \circ \alpha) = \|Q(\gamma_1 \circ \alpha) - Q(\gamma_2 \circ \alpha)\| = \|\psi_1 \odot \alpha - \psi_2 \odot \alpha\| = \|\psi_1 - \psi_2\|, \tag{7}$$

and Γ acts on itself by isometries under the SRSF map. We thus combine the SRSF of cross-component phase ξ_j and cross-observation phase α_i to obtain the SRSF ψ_{ij} of γ_{ij} :

$$\psi_{ij} := Q(\xi_j \circ \alpha_i) = Q(\xi_j) \odot \alpha_i. \tag{8}$$

As described in Section 2, our registration procedure will estimate ψ_{ij} by first estimating the spatially correlated cross-component phases for each observation $i = 1, \dots, n$, since α_i and ξ_j are individually not estimable; in other words, we wish to induce dependence in the generative model (1) between the confounded cross-component and cross-observation phases. We achieve this using the SRSF transform of γ_{ij} in the manner defined in (8).

Under the SRSF transform of Γ , the population conditional expectation $E[\psi_{ij} \mid \psi_1, \dots, \psi_{i(j-1)}, \psi_{i(j+1)}, \dots, \psi_{iK}]$ is an ideal choice for $\tilde{\psi}_{ij}$. As with spatially correlated real- or vector-valued observations in traditional statistics modelled using random fields, a viable estimate of the conditional expectation in this context is the spatial interpolant or *kriging prediction* at location \mathbf{s}_j of a functional random field $\{\psi_{\mathbf{s}} := Q(\xi_{\mathbf{s}}) \odot \alpha, \mathbf{s} \in \mathcal{D}\}$ assuming values in Ψ , conditioned on its values at locations $\mathbf{s}_1, \dots, \mathbf{s}_{j-1}, \mathbf{s}_{j+1}, \dots, \mathbf{s}_K$. In other words, the random field $\{\psi_{\mathbf{s}}\}$ is derived by time-warping the values assumed by another functional random field $\{\xi_{\mathbf{s}}, \mathbf{s} \in \mathcal{D}\}$ with a fixed $\alpha \in \Gamma$.

Computing the kriging prediction of $\{\psi_{\mathbf{s}}\}$ at \mathbf{s}_j requires estimation of $Q(\xi_{i1}), \dots, Q(\xi_{iK})$. Suppose that the components $f_{ij}, i = 1, \dots, n, j = 1, \dots, K$ are indexed by spatial locations $\mathbf{s}_j \in \mathcal{D}, j = 1, \dots, K$. Denote by ψ_{ij} the SRSFs of the warping functions γ_{ij} , which, for every i , are viewed as values assumed by the random field $\{\psi_{\mathbf{s}} : \mathbf{s} \in \mathcal{D}\}$ defined above at spatial locations $\mathbf{s}_1, \dots, \mathbf{s}_K$. Then, under assumptions of second-order stationarity and isotropy on $\{\psi_{\mathbf{s}}\}$, we can consider the *phase trace-variogram*, defined by Guo et al. (2022) as

$$V(h) = \frac{1}{2} \int_0^1 E[\psi_{\mathbf{s}}(t) - \psi_{\mathbf{s}'}(t)]^2 dt = \frac{1}{2} E(\|\psi_{\mathbf{s}} - \psi_{\mathbf{s}'}\|^2),$$

using Fubini's theorem, where $h = \|\mathbf{s} - \mathbf{s}'\|$ and $\|\cdot\|$ is the Euclidean norm on the spatial domain $\mathcal{D} \subset \mathbb{R}^p$. For a fixed observation i , the estimator of $V(h)$ is given by

$$\hat{V}_i(h) = \frac{1}{2|N(h)|} \sum_{a,b \in N(h)} \|\hat{\psi}_{ia} - \hat{\psi}_{ib}\|^2, \tag{9}$$

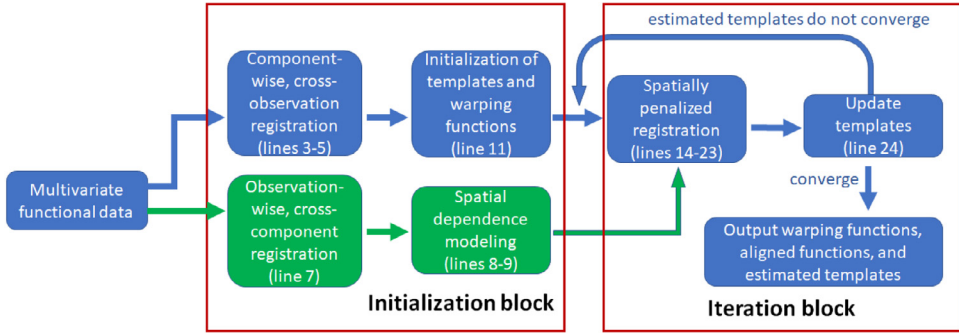


Fig. 2. High-level overview of Algorithm 1.

where $N(h) = \{(\mathbf{s}_a, \mathbf{s}_b) \mid a, b = 1, \dots, K, |\mathbf{s}_a - \mathbf{s}_b| = h\}$. For irregularly spaced data, $N(h)$ can be modified to $N_\epsilon(h) = \{(\mathbf{s}_a, \mathbf{s}_b) \mid |\mathbf{s}_a - \mathbf{s}_b| \in (h - \epsilon, h + \epsilon)\}$ for a small $\epsilon > 0$.

Despite the fact that the spatially correlated cross-component phases ξ_1, \dots, ξ_K are unobserved and confounded with the cross-observation phases $\alpha_1, \dots, \alpha_n$, we can access their correlation structure by estimating γ_{ij} using $V(h)$ and $\hat{V}(h)$. To see this, note that due to the isometry property in (7),

$$E(\|\psi_{\mathbf{s}} - \psi_{\mathbf{s}'}\|^2) = E(\|Q(\xi_{\mathbf{s}}) \odot \alpha - Q(\xi_{\mathbf{s}'}) \odot \alpha\|^2) = E(\|Q(\xi_{\mathbf{s}}) - Q(\xi_{\mathbf{s}'})\|^2),$$

and V thus equals the phase-trace variogram of the functional random field $\{Q(\xi_{\mathbf{s}})\}$; a similar argument applies to the estimate \hat{V} . We thus note that the phase-trace variogram is invariant to (simultaneous) warping (Guo et al., 2022, Lemma 1), which motivates our use of the extrinsic distance on Ψ and the functional random field $\{\psi_{\mathbf{s}}\}$ with values in Ψ .

We now move on to computing the kriging prediction $\tilde{\psi}_{ij}$ using the variogram estimate \hat{V} . For component j in observation i , given phase for the other components $\{\psi_{il}\}_{l \neq j}$, the conditional mean phase is given by the weighted average

$$\tilde{\psi}_{ij} = \sum_{l \neq j} \zeta_{ijl} \psi_{il}, \quad \text{where } \sum_{l \neq j} \zeta_{ijl} = 1, \zeta_{ijl} \geq 0. \tag{10}$$

The coefficient vector $\boldsymbol{\zeta}_{ij} = \{\zeta_{ijl}\}_{l \neq j}$ is implicitly defined as the minimizer of the expected prediction error functional. Guo et al. (2022) show that $\boldsymbol{\zeta}_{ij}$ can be estimated using a quadratic optimization problem that depends on the trace-variogram $V(h)$ and pairwise spatial distances $|\mathbf{s}_a - \mathbf{s}_b|$, for $a, b = 1, \dots, K$. The estimator $\tilde{\psi}_{ij}$ is subsequently normalized using $\tilde{\psi}_{ij} \rightarrow \tilde{\psi}_{ij} / \|\tilde{\psi}_{ij}\|$ to correspond to a valid warping function.

Summarily, for the i th observation, if $\tilde{\psi}_{ij}$ denotes the kriging prediction at \mathbf{s}_j , the penalty term in (6) becomes $d(\gamma_{ij}, \tilde{\gamma}_{ij}) = \|\psi_{ij} - \tilde{\psi}_{ij}\|$, and the optimization problem for registration in (6) with the spatial penalty assumes the specific form

$$(\hat{\gamma}_{1j}, \dots, \hat{\gamma}_{nj}, \hat{\mu}_j) = \arg \min_{\gamma_{1j}, \dots, \gamma_{nj} \in \Gamma, \mu_j \in \mathcal{Q}} \sum_{i=1}^n \{ \|\mu_j - q_{ij} \odot \gamma_{ij}\|^2 + \lambda \|\psi_{ij} - \tilde{\psi}_{ij}\|^2 \}. \tag{11}$$

4.3. Registration algorithm and implementation details

The optimization problem for registration in (11) is solved in an iterative fashion in Algorithm 1, where tools from Section 4.2 are used to model the spatially dependent cross-component phase within each observation. Fig. 2 provides a high-level diagrammatic representation of Algorithm 1. More specifically, given component functions $\{f_{ij}, i = 1, \dots, n, j = 1, \dots, K\}$ of multivariate functional data F_1, \dots, F_n , we first transform them to obtain their SRSFs $\{q_{ij}\}$. Then, for a fixed penalty parameter $\lambda > 0$, Algorithm 1 can be decomposed into two blocks:

Algorithm 1 Spatially penalized registration of multivariate functions

```

1: Input: SRSFs  $q_{ij}$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, K$  and regularization parameter  $\lambda$ .
2: Output: Estimated componentwise template functions  $\hat{\mu}_j$ ,  $j = 1, \dots, K$ , and warping functions  $\hat{\gamma}_{ij}$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, K$ .
3: for  $j = 1$  to  $K$  do
4:   Multiple alignment of  $\{q_{ij}, i = 1, \dots, n\}$  via (4) to initialize the template  $\hat{\mu}_j^{(0)}$ .
5: end for
6: for  $i = 1$  to  $n$  do
7:   Multiple alignment of  $\{q_{ij}, j = 1, \dots, K\}$  via (4) to estimate cross-component phase  $\{\hat{\xi}_{ij}, j = 1, \dots, K\}$ ;
8:   Estimation of trace-variogram  $\hat{V}_i(h)$  in (9) using  $\{\hat{\xi}_{ij}, j = 1, \dots, K\}$ ;
9:   Estimation of coefficient vector  $\zeta_{ij}$  using  $\hat{V}_i$  (Section 4.2).
10: end for
11: Set  $z = 0$ ,  $\epsilon_1 > 0$  (small) and  $\hat{\psi}_{ij}^{(0)}(t) = 1$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, K$  with  $\hat{\mu}_j^{(0)} = n^{-1} \sum_{i=1}^n q_{ij}$ ,  $j = 1, \dots, K$ .
12: while  $z < z_{max}$  and  $\sum_{j=1}^K \|\hat{\mu}_j^{(z)} - \hat{\mu}_j^{(z-1)}\| > \epsilon_1$  do
13:   for  $i = 1$  to  $n$  do
14:     Set  $k = 0$ ,  $\epsilon_2 > 0$  (small);
15:     while  $k < k_{max}$  and  $\sum_{j=1}^K \|\hat{\psi}_{ij}^{(k)} - \hat{\psi}_{ij}^{(k-1)}\|^2 > \epsilon_2$  do
16:       for  $j = 1$  to  $K$  do
17:         Estimation of  $\tilde{\psi}_{ij}$  using  $\zeta_{ij}$  and  $[\hat{\psi}_{i1}^{(k+1)}, \dots, \hat{\psi}_{i(j-1)}^{(k+1)}, \hat{\psi}_{i(j+1)}^{(k)}, \dots, \hat{\psi}_{iK}^{(k)}]$ ;
18:         Solve  $\hat{\gamma}_{ij} = \arg \min_{\gamma \in \Gamma} \left\{ \|\hat{\mu}_j^{(z)} - q_{ij} \odot \gamma\|^2 + \lambda \|\sqrt{\gamma} - \tilde{\psi}_{ij}\|^2 \right\}$ .
19:         Compute SRSF  $\hat{\psi}_{ij}^{(k+1)}$  of  $\hat{\gamma}_{ij}$ .
20:         Set  $k = k + 1$ .
21:       end for
22:     end while
23:     Set  $\hat{\gamma}_{ij} = \hat{\gamma}_{ij}^{(k)}$ 
24:   end for
25:   Set  $\hat{\mu}_j^{(z+1)} = \frac{1}{n} \sum_{i=1}^n (q_{ij} \odot \hat{\gamma}_{ij})$ ,  $j = 1, \dots, K$ ;
26:   Set  $z = z + 1$ .
27: end while

```

Initialization:

- (i) Obtain initial template component functions $\hat{\mu}_1^{(0)}, \dots, \hat{\mu}_K^{(0)}$ by performing independent componentwise registration of $\{q_{1j}, \dots, q_{nj}\}$ for each $j = 1, \dots, K$ given by (4);
- (ii) initialize the phase functions to the identity warping by initializing their SRSFs to $\psi_{ij}^{(0)} \equiv 1$, $i = 1, \dots, n$, $j = 1, \dots, K$;
- (iii) for each observation $i = 1, \dots, n$, implement multiple registration of $\{q_{i1}, \dots, q_{iK}\}$ using (4) to estimate cross-component phase $\xi_{i1}, \dots, \xi_{iK}$;
- (iv) for each observation $i = 1, \dots, n$, compute weights $\zeta_{i1}, \dots, \zeta_{iK}$ required for the kriging prediction using the estimated phase trace variogram \hat{V}_i .

Iteration: Starting with the initial values, obtain phases $\{\hat{\gamma}_{ij}\}$ and template components $\{\hat{\mu}_j\}$ by iterating and updating, until convergence.

- (i) Compute $\tilde{\psi}_{ij}$ in penalty using kriging coefficients $\{\zeta_{ij\ell}\}_{\ell \neq j}$ and phases $\{\hat{\gamma}_{i\ell}\}_{\ell \neq j}$;
- (ii) solve $\hat{\gamma}_{ij} = \arg \min_{\gamma \in \Gamma} \left\{ \|\hat{\mu}_j - q_{ij} \odot \gamma\|^2 + \lambda \|\sqrt{\gamma} - \tilde{\psi}_{ij}\|^2 \right\}$;
- (iii) compute template components $\{\hat{\mu}_j\}$ by averaging $\{q_{ij} \odot \hat{\gamma}_{ij}\}$ over i .

Algorithm 1 contains nested while loops (lines 12–27). The inner loop indexed by k (lines 15–22) updates the estimated warping functions given the templates for all components. The outer loop index by z (lines 12–27) updates the template for each component given the estimated warping

functions. The convergence of Algorithm 1 is assessed empirically using simulation studies in Section 1 in the supplement.

5. Simulation studies

We conduct simulation studies to assess the performance of the proposed multiple registration approach for multivariate functional data. Specifically, we estimate the componentwise template functions $\{\mu_j, j = 1, \dots, K\}$ using three different approaches: (1) the proposed method, which utilizes spatial cross-component phase correlation in penalized registration (Section 4 and Algorithm 1), (2) independent componentwise registration (Section 3.1 and Section 3.4 in Srivastava et al., 2011), and (3) universal registration wherein each component is assumed to have the same phase (Section 3.2). Section 2 in the supplement additionally compares the proposed approach to penalized independent componentwise registration via (3).

5.1. Data generating model

We consider multivariate functional data wherein dependence between function components is induced by spatial correlation. In this setting, each component of a simulated observation is indexed by a spatial coordinate $\mathbf{s} \in \mathbb{R}^p$; component j for observation i is denoted by f_{i,s_j} . Component functions of the data are generated using the model:

$$f_{i,s_j}(t) = (\mu_{s_j} + e_{i,s_j}) \circ (\xi_{i,s_j} \circ \alpha_i)(t), \quad t \in [0, 1], \quad i = 1, \dots, n, \quad j = 1, \dots, K.$$

In the model, $\mu_{s_j} \in \mathcal{F}$ denotes the template for component j indexed by spatial location s_j , which is the object that we wish to estimate through a registration procedure. The random functional error e_{i,s_j} and cross-component warping functions $\xi_{i,s_j} \in \Gamma$ are also indexed by the spatial location. The warping function $\alpha_i \in \Gamma$ denotes the observation specific warping function (common across all components).

We consider two simulation settings characterized by choice of the component template functions (see below). For both settings, the random phase components are composed of two warping functions: $\alpha_i \in \Gamma$, the phase that is common across components for observation i , and $\xi_{i,s_j} \in \Gamma$, the cross-component phase within observation i . Such a nested structure is similar to the structure in a mixed-effects model. The warping functions α_i are taken to be cumulative distribution functions (CDFs) of a beta distribution, $\text{Beta}(1, \exp(z_i))$, with random parameter $z_i \sim \text{Unif}[-Z, Z]$. The cross-component phase ξ_{i,s_j} is also the CDF of a beta distribution, $\text{Beta}(1, \exp(b_{i,s_j}))$, where $(b_{i,s_1}, \dots, b_{i,s_K})^\top$ follows the correlated uniform distribution on $[-B, B]$, independently for every i ; a sample from the correlated uniform distribution can be generated by transforming a correlated multivariate normal sample with mean $(0, \dots, 0)^\top$ and Matern covariance $C_{\text{Mat}}(\cdot, \cdot; 1, 0.5, \ell)$. The parameters Z and B control magnitudes of the cross-observation and cross-component phase variations; the range parameter is set to $\ell = 0.5 \times d_{\text{max}}$, where d_{max} is the maximum spatial distance between the simulated sites s_j .

5.1.1. Simulation setting 1

In this setting, the component template functions $\{\mu_{s_j}\}$ have a specific bimodal form. We consider $K = 20$ components where the spatial coordinates s_j are generated using a uniform distribution on $[-2, 2]^2 \subset \mathbb{R}^2$. The componentwise template functions are generated as $\mu_{s_j}(t) = a_{1s_j} \exp(-100(t - 1/3)^2) + a_{2s_j} \exp(-100(t - 2/3)^2)$, where $(a_{m,s_1}, \dots, a_{m,s_K})^\top$, $m = 1, 2$ are independently sampled from a multivariate normal distribution with mean vector $(3, \dots, 3)^\top$ and Matern covariance $C_{\text{Mat}}(\cdot, \cdot; \sigma_a^2, 0.5, \ell)$; here, σ_a^2 is the scale parameter, ℓ is the range, and the smoothing parameter is fixed to 0.5. The random errors $(e_{i,s_1}(t), \dots, e_{i,s_K}(t))^\top$ are generated independently for each observation i in a pointwise manner: for each value of t , we generate a sample from the multivariate normal distribution with mean vector $(0, \dots, 0)^\top$ and covariance $C_{\text{Mat}}(\cdot, \cdot; \sigma_e^2, 0.5, \ell)$. The other covariance parameters are set to $\sigma_a = 1$, $\sigma_e = 0.5$ and $\ell = 0.5 \times d_{\text{max}}$.

5.1.2. Simulation setting 2

In this setting, in order to replicate the structure of EEG data, the component template functions $\{\mu_{s_j}\}$ are chosen to possess the (typical) shape of EEG signals, i.e., each component is treated as a signal from an EEG electrode. We consider 16 electrode locations on the scalp with three-dimensional coordinates. To imitate the shape variation in EEG signals from different brain regions, we use a more flexible data generating model. The componentwise template functions are generated using 10 B-spline basis functions, B_m , $m = 1, \dots, 10$, of order 4 on the interval $[0, 1]$: $\mu_{s_j}(t) = \sum_{k=1}^{10} \beta_{k,s_j} B_k(t)$. The B-spline basis coefficients $(\beta_{m,s_1}, \dots, \beta_{m,s_K})^\top$, $m = 1, \dots, 10$ are generated independently from a multivariate normal distribution with mean vector $(0, \dots, 0)^\top$ and Matern covariance $C_{Mat}(\cdot, \cdot; \sigma_a^2, 0.5, \ell)$. Random functional errors, e_{i,s_j} , are generated in the same way as in Simulation setting 1, but under a lower signal-to-noise ratio: the covariance parameters are set to $\sigma_a = 2$, $\sigma_e = 0.5$ or 1 and $\ell = 0.5 \times d_{max}$.

5.2. Assessing performance of template estimation

Denote by \tilde{f}_{i,s_j} and \tilde{q}_{i,s_j} the aligned multivariate functions and their SRSFs obtained after registration. We define two performance metrics that quantify the accuracy of template estimation based on the registered functions:

$$MSE = \frac{1}{Kn} \sum_{j=1}^K \sum_{i=1}^n \|\tilde{f}_{i,s_j} - \mu_{s_j}\|^2, \quad QMSE = \frac{1}{Kn} \sum_{j=1}^K \sum_{i=1}^n \|\tilde{q}_{i,s_j} - Q(\mu_{s_j})\|^2,$$

where μ_{s_j} is the true template for component j . MSE, based on the \mathbb{L}^2 norm on \mathcal{F} , is sensitive to phase errors, while QMSE, based on the eFR metric on \mathcal{F} , is sensitive to shape differences (Srivastava and Klassen, 2016).

We use $n = 20, 20$ and $K = 20, 16$, respectively, for the first and second simulation settings described in Section 5.1; we replicate each simulation 50 times. When implementing the proposed method, we use 4-fold cross-validation to select a value for the regularization parameter λ . Denote by $\hat{\mu}_{s_j}^{[-k]}$ the estimated template function for component j based on data in all folds except the k th one; here, we use $\langle k \rangle$ to denote the set of observation indices in fold k . Then, the value for the regularization parameter λ is chosen by minimizing $\frac{1}{4Kn} \sum_{k=1}^4 \sum_{j=1}^K \sum_{i \in \langle k \rangle} \|\tilde{f}_{i,s_j} \circ \hat{\gamma}_{i,s_j} - \hat{\mu}_{s_j}^{[-k]}(\lambda)\|^2$, where $\hat{\gamma}_{i,s_j} = \arg \min_{\gamma \in \Gamma} \|Q(\hat{\mu}_{s_j}^{[-k]}(\lambda)) - Q(f_{i,s_j}) \odot \gamma\|^2$, i.e., we align functions in the validation set (fold k) to the template estimated using the training set (all folds except k). The selected value for the regularization parameter is then used for registration of all observations.

Results for each simulation setting described in Section 5.1 are summarized in the boxplots shown in Fig. 3. In all cases, we also report values of the error metrics when no registration is applied to the data; this serves as the baseline.

5.2.1. Simulation setting 1

The left panel of Fig. 3 shows that both componentwise registration and the proposed method clearly have smaller MSEs than universal registration; further, the mean MSE for the proposed method (0.083) is slightly lower than the mean MSE for the componentwise method (0.099). With respect to QMSE, the performance of componentwise registration is clearly worse than the proposed approach. Interestingly, universal registration outperforms componentwise registration in terms of median QMSE.

The improvement in registration quality with the proposed method over the componentwise approach is due to the addition of the regularization term. Componentwise registration tends to ‘overalign’ features of the component functions, e.g., peaks and valleys, that are due to the random error, i.e., the estimate of cross-component phase is too flexible. On the other hand, the universal approach performs poorly since all components within an observation are constrained to be identical, i.e., the estimate of cross-component phase is too restrictive. The compromise achieved by the proposed method between the componentwise and universal methods by introducing a regularization penalty that synthesizes spatial phase information from all components within an

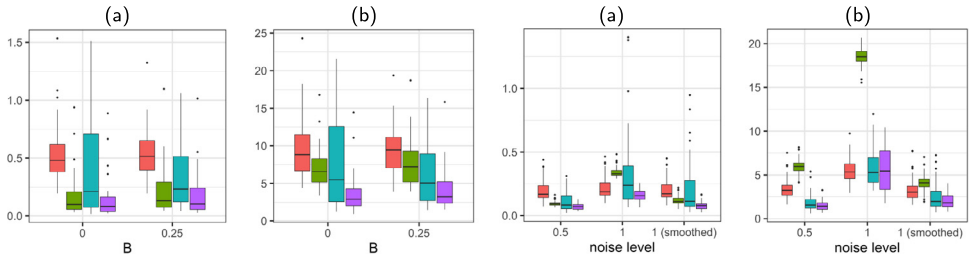


Fig. 3. Left: Simulation setting 1 with $Z = 0.5$ and $\sigma_e = 0.5$; scale parameter of cross-component phase variation is set to $B = 0$ or $B = 0.25$. Right: Simulation setting 2 with $Z = 0.5$ and $B = 0$; when $\sigma_e = 1$, we also present results after applying smoothing splines with a low value of the smoothing parameter. Boxplots of the (a) MSE and (b) QMSE for template functions estimated with no registration (red), componentwise registration (green), universal registration (cyan) and proposed penalized registration (purple). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Registration result for a single component in one simulation replicate under Simulation setting 2 with $\sigma_e = 0.5$. Panel 1: Simulated data in grey with cross-sectional average without registration in red. Panels 2-4: Aligned functions in grey with estimated template in red, generated using componentwise registration, universal registration, and the proposed method, respectively. The ground truth template function is shown in black. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

observation results in improved registration. Finally, we note that the performance of the proposed method is consistently better than the other two approaches, even if no cross-component phase variation exists in the data (the case when $B = 0$).

5.2.2. Simulation setting 2

Recall that the setting here is designed to replicate the structure in EEG data. The low signal-to-noise ratio presents significant challenges for all of the registration procedures. The right panel of Fig. 3 shows that the componentwise approach, expectedly, is the most sensitive to the random function errors due to its lack of regularization. Despite the good performance of this method in the first simulation setting, the low signal-to-noise ratio when $\sigma_e = 1$ significantly deteriorates the template estimates, especially with respect to QMSE. In fact, when $\sigma_e = 1$, the advantages of registration are not obvious since all of the procedures are essentially driven by random noise; regularization with the spatial penalty provides some help, but there is insufficient structure in $\{f_{i,s_j}\}$ to inform estimation of the conditional mean phase based on cross-component spatial correlation. In such situations, we recommend slight smoothing of the observed functions prior to registration using any off-the-shelf smoothing procedure (e.g., splines). To examine value in the pre-smoothing when $\sigma_e = 1$, we report additional results after applying smoothing splines to the simulated data, with a small value for the smoothing parameter. We see that the proposed registration method produces more accurate and stable template estimates than the other two registration approaches.

In Fig. 4, we provide results of template estimation for a single component in one simulation run when $\sigma_e = 0.5$. The simulated data for this component is shown in grey with the ground truth template in black. The estimated templates are shown in red. In Panel 1 of Fig. 4, it is evident that when no registration is performed, the resulting template underestimates relevant shape features, e.g., the two peaks and one valley between $t = 0.4$ and $t = 0.75$; in Panel 2, componentwise

registration results in a template that generates additional shape features due to overalignment of noise, e.g., the two small peaks and three valleys between $t = 0$ and $t = 0.3$; in Panel 3, universal registration results in a template that has the nearly correct shape, but that is out of phase with respect to the ground truth template, e.g., there is significant lag between $t = 0.7$ and $t = 1$. Finally, in Panel 4, we note that the proposed approach results in a template that is very similar in shape to the ground truth and is largely in phase.

6. Illustrations on real data

In this section, we examine performance of the registration method on two real-data examples pertaining to multivariate functional data arising from EEG and ozone studies.

6.1. Electroencephalogram data

We analyse an EEG dataset arising from a study that examined EEG correlates of genetic predisposition to alcoholism (Bache and Lichman, 2013). The study resulted in EEG measurements recorded at 64 electrodes placed on the subjects' scalps, with a sampling rate of 256 Hz (3.9-msec epoch) for 1 s. Each subject was exposed to a stimulus in each trial and completed multiple trials where different stimuli were shown. The study used standard 64-channel electrode placements as defined by the American Electroencephalographic Association. In our analysis, we use EEG signals from 50 trials for one subject, and treat each trial as an independent multivariate functional observation. We use 61 out of the 64 electrode locations as the components in each observation; the three electrode locations that are not used are very far away from the main scalp region. Thus, the data for this study is $F_i = (f_{i1}, \dots, f_{iK})^T$, $K = 61$, $i = 1, \dots, 50$. Prior to analysis, each EEG signal was smoothed using smoothing splines with a small parameter value of 1×10^{-5} . To compare the performance of different methods, we register the 50 trials using the componentwise, universal and proposed registration approaches. After registration, we average the aligned signals recorded by each electrode across trials to estimate the mean event related potentials. For the proposed method, we use 4-fold cross-validation to select a value for the regularization parameter λ .

For easy comparison across the three methods, we display the registration results for electrode AFZ in Fig. 5. As in the simulations, we also display the result of averaging across trials when no registration is applied to the data as a baseline (Panel 1). The electrode AFZ was chosen as an illustration since many of the EEG signals recorded there contain a pronounced activation peak between $t = 0.75$ and $t = 1$. We observe similar results in this real EEG data example, based on the different registration methods, as in the simulation studies. In Panel 1, when no alignment is applied, the pronounced activation peak is almost completely nonexistent in the average signal shown in red; thus, registration of the data is necessary prior to averaging. In Panel 2, the componentwise registration method tends to overalign many of the relatively small modes present in the signals that are likely due to noise. This results in an average that has many small fluctuations that are indistinguishable from the pronounced activation peak. In Panels 3 and 4, the universal and proposed methods are effective in capturing the large peak in the respective average signals and tend to produce fewer small fluctuations prior to $t = 0.75$. Essentially, these two methods mitigate the contribution of noise and result in average signals that reflect the most prominent feature in the given data. The magnitude of the pronounced activation peak is similar in the two averages, but it is stretched over a longer part of the domain in the average produced using the universal method.

In EEG data analysis, it is often of interest to study the relationship between average signals (across trials) at different electrode locations after registration. To show the benefits of the proposed approach in downstream EEG data analysis tasks, we explore the spatial correlation between estimated average EEG signals using the empirical trace-variogram (Giraldo et al., 2011), defined as

$$\hat{V}(h) = \frac{1}{2|N_\epsilon(h)|} \sum_{a,b \in N_\epsilon(h)} \|\hat{\mu}_{s_a} - \hat{\mu}_{s_b}\|^2, \tag{12}$$




Fig. 5. Registration and estimate $\hat{\mu}_{s_j}$ of template component μ_{s_j} at electrode AFZ (location s_j). Panel 1: Simulated data in grey with cross-sectional average without registration in red. Panels 2-4: Aligned functions in grey with estimated average signal in red, generated using componentwise registration, universal registration, and the proposed method, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Empirical trace-variograms \hat{V} (left) computed using estimated average signals $\{\hat{\mu}_{s_j}, j = 1, \dots, 61\}$ obtained using three registration methods: componentwise (red, solid), universal (blue, dashed), and proposed (green, dotted). Estimated average EEG signals $\{\hat{\mu}_{s_1}, \dots, \hat{\mu}_{s_{14}}\}$ at 14 (out of the total 61) electrodes placed in an area related to the parietal lobe, computed using the proposed (middle) and componentwise (right) registration approaches. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where $N_\epsilon(h) = \{(\mathbf{s}_a, \mathbf{s}_b) : \|\mathbf{s}_a - \mathbf{s}_b\| \in (h - \epsilon, h + \epsilon)\}$ for a small $\epsilon > 0$ and $\hat{\mu}_{s_a}$ denotes the estimated average signal for electrode located at \mathbf{s}_a , $a = 1, \dots, 61$. The empirical trace-variograms, computed based on average signals generated by the three different registration methods, are shown in the left panel of Fig. 6 (componentwise in red, universal in blue, proposed in green); the average signals estimated using the proposed approach, which were used to compute the green trace-variogram, are shown at their respective scalp locations in the left panel of Fig. 7. The variograms suggest very similar spatial correlation patterns between the average signals, but the proposed approach produces averages with smallest spatial variation. The universal approach results in more spatial variation due to its very restrictive assumption of common phase across all channels. The magnitude of spatial variation in this case is similar to the proposed method at small spatial distances. This is intuitive since cross-component phase variation at nearby electrodes is very small and the common phase assumption is reasonable. However, as the spatial distance increases, this assumption becomes unrealistic.

The componentwise method conducts separate alignment of EEG signals at each electrode, resulting in large phase variation across electrode locations. On the other hand, the proposed method avoids this issue by accounting for spatial phase correlation across electrodes via the regularization penalty. To demonstrate this, we consider 14 electrodes in the area corresponding to the parietal lobe. Since all 14 electrodes are related to the same brain region, we expect very little phase variation in the resulting average signals. The estimated averages are shown in the middle panel in Fig. 6 for the proposed method and the right panel for the componentwise method. In the middle panel, there is very little phase variation across the estimated averages, as expected. However, in the right panel, it is easy to see that considerable phase variation remains.

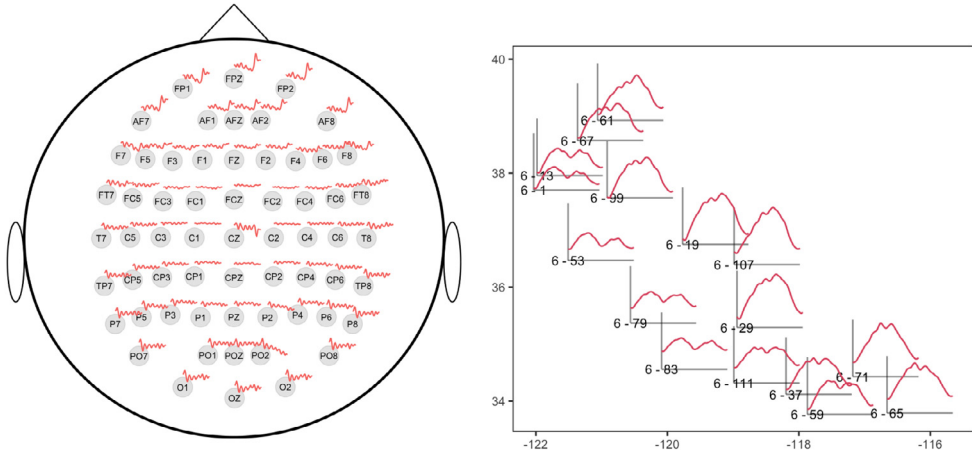


Fig. 7. Left: Average EEG signals estimated using the proposed method, $\{\hat{\mu}_{s_j}, j = 1, \dots, 61\}$, shown at the (projected) two-dimensional electrode locations on a toy map of the scalp. Right: Average ozone concentration functions estimated using the proposed method, $\{\hat{\mu}_{s_j}, j = 1, \dots, 16\}$, shown at the station locations in northern California.

6.2. Ozone concentration data

Ground-level ozone is a harmful pollutant that is monitored closely by the Environmental Protection Agency. The temporal trend of daily ozone concentration varies from year to year due to variations in various environmental conditions, including the weather. To explore average ozone concentration patterns at a certain location, we desire to first eliminate phase variation that exists across different observation years using a registration procedure. At the same time, we must account for cross-component phase variation across different locations within the same year that is due to different temporal patterns of pollutant spread. Thus, we view multi-year ozone concentration functions, observed at different spatial locations, as multivariate functional data with two different sources of phase variation; here, univariate ozone concentration functions observed at different spatial locations are treated as components of a full observation corresponding to a single year. In this analysis, we focus on a small area in northern California ($35^\circ \sim 39^\circ$ N, $120 \sim 123^\circ$ W) with $K = 16$ observation stations (components). Each station recorded daily average ozone concentration (in parts per million) from year 2000 to year 2019 (sample size $n = 20$). Thus, the data for this study is $F_i = (f_{i1}, \dots, f_{iK})^T$, $K = 16$, $i = 1, \dots, 20$. The data is publicly available on the air data website¹ of the United States Environmental Protection Agency. The data for the single year 2018, constituting one multivariate functional data observation, is shown in the left panel of Fig. 8. There is evidence of cross-component spatial correlation, i.e., ozone concentration patterns at nearby stations are similar. In the right panel of Fig. 8, we fix the spatial location to station 6–19, and present the observations at this single site across multiple years: 2000–2019. It is clear that there is considerable phase (or temporal) variation across these yearly observations, since the timing of features in the ozone concentration functions varies across years. This motivates application of the proposed approach for cross-year registration while accounting for the spatial structure across observation sites.

As in the previous real data analysis example that considered EEG data, we compare the performance of three registration procedures in this context: componentwise, universal and proposed. We display the registration results for a single location, corresponding to Alameda County (37.8 N, 122.3 W; location 6–1), in Fig. 9. We also display the result of averaging across trials when no registration is applied to the data as a baseline (Panel 1). We make several interesting observations based on these results. First, the proposed method yields an estimate of the average

¹ <https://www.epa.gov/outdoor-air-quality-data>




Fig. 8. Left: Daily average ozone concentration functions observed at 16 stations in northern California in 2018. The IDs mark the locations of stations on the map and red curves are smoothed concentration data during a year. Right: Daily ozone concentration functions at site 6-19 for each of the years 2000-2019; different colours indicate different years. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 9. Registration and estimate $\hat{\mu}_{s_j}$ of template component μ_{s_j} , where s_j is Alameda County (37.8 N, 122.3 W; ID 6-1 on map in left panel of Fig. 8). Panel 1: Observed data in grey with cross-sectional average without registration in red. Panels 2-4: Aligned functions in grey with estimated average signal in red, generated using componentwise registration, universal registration, and the proposed method. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

ozone concentration function (Panel 4) that contains clearer patterns than the average computed without alignment (Panel 1), e.g., the steep reduction in ozone concentration around day 200. The proposed approach produces an average that has very similar patterns to those computed using the componentwise and universal registration methods, but is smoother overall. The most noticeable difference occurs in early autumn where the proposed approach results in an average that has a single mode, while the other two methods result in averages with two smaller peaks. A possible reason for this phenomenon is the large regularization parameter value chosen via cross-validation; this restricts the complexity of estimated cross-component warping functions and forces the registration procedure to overlook small features that are potentially induced by noise.

As in the EEG data example, we further analyse the estimated average ozone concentration functions at the 16 spatial locations. Modelling of spatial correlation based on averages computed after registration is also of particular interest in this application. In the left panel of Fig. 10, we display the empirical trace-variograms, computed using (12) after estimating the average ozone concentration functions using the componentwise (red), universal (blue) and proposed (green) approaches; the average signals estimated using the proposed approach, which were used to compute the green trace-variogram, are shown at their respective station locations in the right panel of Fig. 7. The conclusions here are similar to those reached in the EEG data example. The proposed method yields componentwise averages that have smallest spatial variation. The variation gap between the universal and proposed methods appears bigger in this case, even for small spatial




Fig. 10. Empirical trace-variograms \hat{V} (left) computed using estimated average ozone concentration functions $\{\hat{\mu}_{s_j}, j = 1, \dots, 16\}$ at 16 locations in a small area in northern California obtained using three registration methods: componentwise (red, solid), universal (blue, dashed), and proposed (green, dotted). Estimated average ozone concentration functions $\{\hat{\mu}_{s_j}, j = 1, \dots, 16\}$ computed using the proposed (middle) and componentwise (right) registration approaches. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

distances. We further display the componentwise estimated average ozone concentration functions, for each of the 16 locations, generated using the proposed and componentwise registration methods in the middle and right panels of Fig. 10, respectively. It is obvious that the proposed approach (middle) results in estimated average functions that exhibit much less cross-component phase variation than the average functions generated using the componentwise registration method (right). Furthermore, most averages in the middle panel have two ozone concentration maxima around April/May and August, and a single ozone concentration minimum around June/July. In the right panel, there is much more variation in the number of extrema in the estimated averages as well as their timing.

7. Discussion

We have proposed a novel penalized registration framework for multivariate functional data wherein cross-component phase variation is spatially correlated. The spatial structure is incorporated into a penalty term that regularizes the estimated phase variation. Importance of the elastic metric and its invariance to time warping, which is made practically useful through the square-root slope function (SRSF) representation, cannot be overstated. The invariance property is used to great benefit in both the objective function and the novel spatially-informed penalty term through the kriging prediction $\hat{\psi}_{ij}$, and enables us to disregard estimating the cross-observation phases $\{\alpha_i\}$. We have shown the effectiveness of this approach using simulation studies, and via real data examples that consider multi-trial EEG data and daily ozone concentration functions observed across multiple years.

The proposed regularization penalty is limited by the strength of spatial correlation across components in a multivariate functional observation. Thus, a sufficient number of components (across locations that are not widely spread on the spatial domain) is needed for reliable estimation of the cross-component phase trace-variogram used to define the kriging prediction in the spatial penalty term. Furthermore, we only consider the case of dense multivariate functional data. For sparse functional data, an additional function estimation procedure needs to be incorporated into the framework.

The proposed penalized registration procedure can in principle be used to align multivariate functional data with component functions that are correlated in other ways. For example, when component functions are temporally correlated, we can replace the kriging prediction $\hat{\psi}_{ij}$ in the penalty term of the objective function (6) with a temporal interpolant or temporally weighted estimate (e.g., moving time-window average). There is much to be done in this direction, and the proposed method represents a promising initial foray.

The focus in this paper was restricted to multivariate functional data wherein *only* the cross-component phases are spatially correlated. The objective function (6) does not quantify and incorporate any spatial correlation between the cross-component amplitudes, and this may well

be of interest in certain applications. Extension of the algorithm to this case is possible using the *amplitude trace-variogram* proposed by Guo et al. (2022), in addition to the phase trace-variogram. This presents a fruitful line for future work.

There is potential for improvement in reducing the computational burden when implementing the registration procedure when the number n of functional observations with K components are both large. Registration in the elastic framework uses the dynamic programming algorithm since the class of warping functions Γ is infinite-dimensional and unconstrained (Srivastava and Klassen, 2016). Restricting attention to a smaller parametric class (e.g., parameterized class of distribution or quantile functions on $[0, 1]$) will reduce computing time considerably, but at the cost of flexibility in registration.

Acknowledgements

We acknowledge funding from the NSF DMS-2015374 and EPSRC EP/V048104/1 to KB; NSF CCF-1740761, DMS-2015226, CCF-1839252 to SK; and NIH R37-CA214955 to SK and KB.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.spasta.2023.100760>. The supplement includes an empirical assessment of convergence for Algorithm 1 and a comparison to penalized componentwise registration.

References

- Bache, K., Lichman, M., 2013. UCI machine learning repository. Irvine, CA, USA, <https://archive.ics.uci.edu/ml/datasets/EEG+Database>.
- Carroll, C., Müller, H.G., 2023. Latent deformation models for multivariate functional data and time-warping separability. *Biometrics* <http://dx.doi.org/10.1111/biom.13851>.
- Carroll, C., Müller, H.G., Kneip, A., 2021. Cross-component registration for multivariate functional data, with application to growth curves. *Biometrics* 77 (3), 839–851.
- Giraldo, R., Delicado, P., Mateu, J., 2011. Ordinary kriging for function-valued spatial data. *Environ. Ecol. Stat.* 18 (3), 411–426.
- Guo, X., Kurtek, S., Bharath, K., 2022. Variograms for kriging and clustering of spatial functional data with phase variation. *Spatial Stat.* 52, 100687.
- Kurtek, S., Srivastava, A., Klassen, E., Ding, Z., 2012. Statistical modeling of curves using shapes and related features. *J. Amer. Statist. Assoc.* 107 (499), 1152–1165.
- Makeig, S., Onton, J., Sejnowski, T., Poizner, H., 2007. Prospects for mobile, high-definition brain imaging: EEG spectral modulations during 3-D reaching. *Hum. Brain Mapp.*
- Marron, J.S., Ramsay, J.O., Sangalli, L.M., Srivastava, A., 2015. Functional data analysis of amplitude and phase variation. *Statist. Sci.* 30 (4), 468–484.
- Olsen, N., Markussen, B., Raket, L.L., 2016. Simultaneous inference for misaligned multivariate functional data. *J. R. Stat. Soc. Ser. C. Appl. Stat.* 67, 1147–1176.
- Park, J., Ahn, J., 2017. Clustering multivariate functional data with phase variation. *Biometrics* 73 (1), 324–333.
- Srivastava, A., Klassen, E.P., 2016. *Functional and Shape Data Analysis*. Springer.
- Srivastava, A., Wu, W., Kurtek, S., Klassen, E., Marron, J.S., 2011. Registration of functional data using Fisher-Rao metric. [arXiv:1103.3817](https://arxiv.org/abs/1103.3817).
- Stam, C.J., Nolte, G., Daffertshofer, A., 2007. Phase lag index: assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. *Hum. Brain Mapp.* 28 (11), 1178–1193.
- Sur, S., Sinha, V.K., 2009. Event-related potential: An overview. *Ind. Psychiatry J.* 18 (1), 70–73.
- Tsai, A.C., Jung, T.P., Chien, V.S., Savostyanov, A.N., Makeig, S., 2014. Cortical surface alignment in multi-subject spatiotemporal independent EEG source imaging. *NeuroImage* 87, 297–310.
- Wang, K., Begleiter, H., Porjesz, B., 2001. Warp-averaging event-related potentials. *Clin. Neurophysiol.* 112 (10), 1917–1924.
- Zhao, W., Xu, Z., Li, W., Wu, W., 2020. Modeling and analyzing neural signals with phase variability using Fisher-Rao registration. *J. Neurosci. Methods* 346, 108954.