

Bayesian semiparametric modelling of phase-varying point processes*

Bastian Galasso

*Department of Innovation and Digital Transformation,
Coca-Cola Embonor,
Chile*

e-mail: bgalasso@embonor.cl

Yoav Zemel

*Statistical Laboratory,
University of Cambridge,
United Kingdom*

e-mail: zemel@statslab.cam.ac.uk

and

Miguel de Carvalho

*School of Mathematics,
University of Edinburgh,
United Kingdom*

e-mail: miguel.decarvalho@ed.ac.uk

Abstract: We propose a Bayesian semiparametric approach for registration of multiple point processes. Our approach entails modelling the mean measures of the phase-varying point processes with a Bernstein–Dirichlet prior, which induces a prior on the space of all warp functions. Theoretical results on the support of the induced priors are derived, and posterior consistency is obtained under mild conditions. Numerical experiments suggest a good performance of the proposed methods, and a climatology real-data example is used to showcase how the method can be employed in practice.

Keywords and phrases: Bernstein–Dirichlet prior, Fréchet mean, phase variation, point processes, random Bernstein polynomials, Wasserstein distance.

Received June 2021.

1. Introduction

A prototypical characteristic in the analysis of a random function $X(t)$ —that distinguishes it from classical multivariate analysis—is that it potentially exhibits two distinct layers of stochastic variability. Amplitude variation is encap-

*BG was partially supported by the graduate scholarship 21140901 from the Chilean NSF (CONICYT), YZ was supported by Swiss National Science Foundation Early Postdoc Mobility Fellowship # 178220, and MdC was partially supported by FCT (Fundação para a Ciência e a Tecnologia, Portugal) through the projects PTDC/MAT-STA/28649/2017 and UID/MAT/00006/2019.

sulated in the fluctuations of $X \equiv X(t)$ around its mean function $\mu(t)$, and can be probed by linear tools, perhaps most prominently the covariance operator of X and the subsequent Karhunen–Loève expansion. Phase variation amounts to variability in the argument t , usually modelled by a random warp function T defined on the domain of definition of X , so that one observes realisations (discretised over some grid) from the random function $\tilde{X}(t) = X(T^{-1}(t))$ instead of $X(t)$. In short, phase variation is randomness in the t -axis, whereas amplitude variation pertains to stochasticity in the X -axis.

Typically, one is interested in inferring properties of the original function X , rather than those of \tilde{X} . In such situations phase variation can be thought of as a nuisance parameter, and failing to account for it may result in a severely distorted statistical analysis: the mean function and Karhunen–Loève expansion of \tilde{X} are smeared and less informative than those of X . Consequently, one needs to undo the warping effect of the phase variation by constructing estimators \hat{T} for the warp functions, and composing them with the observed realisations from \tilde{X} , a procedure known as registration, or alignment, of the functions. The registered functions $\tilde{X}_i \circ \hat{T}_i = X_i \circ T_i^{-1} \circ \hat{T}_i$ are then treated as distributed approximately as X , allowing for their use in probing the law of X . For a textbook treatment of phase variation, we refer to the books by [18, 19]; one may also consult the review articles [10] and [23].

In this paper, we propose a Bayesian method for registering phase-varying point processes. Our paper is aligned with recent developments focused on modelling phase and amplitude variation of complex objects that are not functional data per se, yet still carry infinite-dimensional traits. An intriguing example is that of point processes, appearing as spike trains in neural activity [e.g., 25], where phase variation can be viewed as smearing locations of peaks of activity. See Figure 1 for an example of such phase-varying point processes (and Section 3 for more details on the underlying processes). Such data can be transformed into functional data by smoothing and considering density functions [24], but can be also be dealt with directly, replacing the ambient space L^2 used for functional data by a space of measures. Indeed, [12] formalise the problem and show how the Wasserstein metric of optimal transport arises canonically in the point process version of the problem. Here we propose a Bayesian model that is flexible in being nonparametric, while at the same adapted to the warping problem in a point process context, in the sense that our priors for the warp functions obey the same classical phase variation assumption in functional data analysis (see Section 2.2). From a conceptual viewpoint, our model can be regarded as a semiparametric Bayesian version of [12], but by putting directly a prior on the space of all random measures on the unit interval it allows for straightforward inference from posterior outputs—both in terms of credible bands for warp functions, and credible intervals for registered points. By modelling the mean measure of each phase-varying point process with a random Bernstein polynomial [14, 15], we are able to show that the support of the induced priors for the warping functions and collections of registered points is ‘large’ in the sense made precise in Sections 2.3–2.4. Posterior consistency is established under a

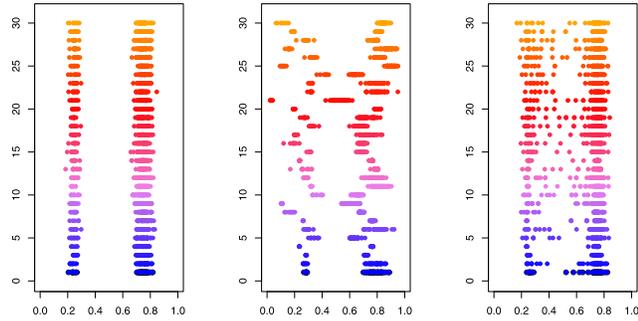


FIG 1. *Left: Simulated instances of the original point process. Middle: Phase-varying versions. Right: Registered versions obtained using the method proposed in the manuscript. Details on the underlying processes can be found in Section 3.*

proviso that is asymptotically equivalent to that of [12], but our large sample results only require the number of points in each process to increase.

Section 2 develops details of our approach, in Section 3 we report numerical experiments, and Section 4 includes a climatology real-data example. Concluding remarks are given in Section 5. To streamline the presentation, proofs, figures, simulations and implementation details are provided as appendices. An R implementation of our method is available at the following address:

<https://github.com/bgalasso/Rmpp>.

2. Random Bernstein polynomial-based registration of multiple point processes

2.1. Random Bernstein polynomials

Random Bernstein polynomials were introduced by [14, 15] and are defined as

$$B(x | k, G) = \sum_{j=0}^k G\left(\frac{j}{k}\right) \binom{k}{j} x^j (1-x)^{k-j}, \quad (2.1)$$

where G is a random function on $[0, 1]$ and k is a (positive) integer-valued random variable. The derivative of $B(x | k, G)$ is

$$b(x | k, G) = \sum_{j=1}^k w_{j,k} \beta(x | j, k-j+1), \quad (2.2)$$

where $w_{j,k} = G(j/k) - G((j-1)/k)$ and $\beta(x | a, b)$ is a Beta density function with parameters $a, b > 0$. When G is a distribution function (i.e., nondecreasing with $G(1) = 1$ and $G(0) = 0$), $(w_{1,k}, \dots, w_{k,k})$ is in the unit simplex $S_k = \{(w_1, \dots, w_k) \in [0, 1]^k : \sum_{j=1}^k w_j = 1\}$. It follows that $B(x | k, G)$ is a

distribution function on $[0, 1]$ with density $b(x | k, G)$. If G has a continuous density g , then $b(x | k, G)$ approximates g uniformly as $k \rightarrow \infty$ (see Lemma 3).

Definition 1 ([14, 15]). *The probability measure π induced by B in (2.1), on the set Δ of all continuous distribution functions defined on $[0, 1]$, is called Bernstein prior with parameters (k, G) . In symbols, $\pi \equiv \pi(k, G)$.*

Further details on random Bernstein polynomials can be found in [6, Section 5.5]. To avoid unnecessarily burdening notation, measure-theoretical considerations will be kept to a minimum (including the measures with respect to which expected values are defined).

2.2. Bayesian semiparametric inference for phase-varying point processes

Let Π be a point process in $[0, 1]$, with finite second moment ($E\{(\Pi[0, 1])^2\} < \infty$), and denote its mean measure by $\lambda(\cdot) = E\{\Pi(\cdot)\}$. Estimation of λ is straightforward when one has access to multiple realisations $\{\Pi_1, \dots, \Pi_n\}$ from Π , with $\hat{\lambda}$ asymptotically normal [7, Proposition 4.8]. Suppose, however, that one instead observes a sample $\{\tilde{\Pi}_1, \dots, \tilde{\Pi}_n\}$ with

$$\tilde{\Pi}_i = T_{i\#}\Pi_i,$$

where $T_{i\#}\Pi_i(\cdot) = \Pi_i\{T_i^{-1}(\cdot)\}$ denotes the push-forward of Π_i through T_i , for all i . In other words, if a given realisation of Π_i is the collection of points $\{x_{i,j}\}_{j=1}^{m_i}$, then one observes the deformed collection $\{\tilde{x}_{i,j}\}_{j=1}^{m_i} \equiv \{T_i(x_{i,j})\}_{j=1}^{m_i}$. Here, $\{T_1, \dots, T_n\}$ is a sequence of random warp functions, that is, increasing homeomorphisms on $[0, 1]$. A target of interest will be on learning about the warp functions, so to register the point processes. To achieve this goal we model the (conditional) mean measures of the phase-varying point processes with a Bernstein–Dirichlet prior, which induces a prior on the space of all warp functions. The conditional mean measure of the warped version $\tilde{\Pi}_i$ given T_i is denoted by $\Lambda_i(\cdot) = E\{\tilde{\Pi}_i(\cdot) | T_i\}$. We impose the rather standard assumptions that $E\{T_i(x)\} = x$ (unbiasedness) for all $x \in [0, 1]$, and that the collection $\{T_1, \dots, T_n\}$ is independent of $\{\Pi_1, \dots, \Pi_n\}$; the assumptions of unbiasedness and monotonicity of warp functions are *sine qua non* in the phase variation literature, often accompanied with additional conditions [e.g., 22, 23]. In words, unbiasedness is tantamount to requiring the average time change $E\{T(x)\}$ to be the identity: on average, the “objective” time-scale should be maintained, so that time is not sped up or slowed down. In fact, unbiasedness and monotonicity are key for identifiability.

To learn about $F_i(x) = \int_0^x \Lambda_i(dt)$, for $x \in [0, 1]$, we set the prior

$$F_i(x) = B(x | k_i, G_i), \quad x \in [0, 1], \tag{2.3}$$

where $\{k_1, \dots, k_n\}$ is a sequence of independent integer-valued random variables and $\{G_1, \dots, G_n\}$ is a sequence of independent random measures. As an example

of a concrete embodiment of (2.3) consider the following setup. Let $\{\tilde{x}_{i,j}\}_{j=1}^{m_i}$ be the points corresponding to $\tilde{\Pi}_i$, and for $i = 1, \dots, n$ set

$$\begin{aligned} \tilde{x}_{i,j} \mid F_i &\sim F_i, \quad j = 1, \dots, m_i, \quad F_i(x) = B(x \mid k_i, G_i), \\ G_i \mid \alpha &\sim \text{DP}(\alpha, G^*), \quad k_i \sim \text{Unif}\{1, \dots, k_{\max}\}. \end{aligned} \quad (2.4)$$

Here ‘DP’ stands for Dirichlet process [4], with precision parameter $\alpha > 0$ and centering distribution $G^* = E(G_i)$, and k_{\max} is a common hyperparameter. To complete the model specification we set $G^* = \text{Beta}(a_0^{G^*}, b_0^{G^*})$ and $\alpha \sim \text{Gamma}(a_0^\alpha, b_0^\alpha)$, for $i = 1, \dots, n$. Of course, more sophisticated versions of (2.4) are conceivable by e.g., specifying different precision and centering for the DP per each point process. For concreteness (2.4) should be regarded as a baseline instance of the model, and indeed similar $\text{Unif}\{1, \dots, k_{\max}\}$ -type specifications are often employed in practice [e.g. 15, p. 118]; yet, in line with related literature [e.g. 14, 15], theoretical claims on the flexibility of the induced priors to be made in Section 2.3 will require the support of the k_i to be the set of positive integers. Below, we assume that the $\{G_i\}$ and $\{k_i\}$ are independent. Moreover, by a slight abuse of notation we identify F_i with Λ_i and more generally, a measure μ with its distribution function $F_\mu(x) = \mu\{[-\infty, x]\}$.

Now, $\{F_1, \dots, F_n\}$, specified as in (2.3), can be used to induce a prior F on the mean measure λ of the random point process Π and on the warp maps T_i . The prior F will be centred around the structural mean λ in the Fréchet mean sense that λ is the closest to F in expectation, that is, $E_\lambda\{d^2(\lambda, F)\} \leq E_\lambda\{d^2(\gamma, F)\}$, for all diffuse measures γ on $[0, 1]$. An obvious question that arises is what metric d should one use, but the Wasserstein distance [20, 13] has been shown to be the canonical metric for phase-varying point processes by [12, Section 3]. It is defined by

$$d(\mu, \nu) = \inf_{Q \in \Gamma(\mu, \nu)} \sqrt{\int_0^1 \{Q(x) - x\}^2 \mu(\mathrm{d}x)}, \quad (2.5)$$

where $\Gamma(\mu, \nu)$ is the collection of functions $Q : [0, 1] \rightarrow [0, 1]$ such that $Q_{\#}\mu = \nu$. (If μ is not diffuse, then $\Gamma(\mu, \nu)$ may be empty and the definition of d needs to be modified, but we will only have to deal with diffuse measures in the sequel.) Since Fréchet averaging with respect to Wasserstein distance amounts to averaging of quantile functions [1], the prior on F is induced from the prior on $\{F_1, \dots, F_n\}$ as the probability law of

$$F(x) = \left(\frac{1}{n} \sum_{i=1}^n F_i^{-1} \right)^{-1}(x), \quad x \in [0, 1]. \quad (2.6)$$

The random Bernstein polynomial-induced prior on each T_i defines the optimal transport map of F onto F_i [20]:

$$T_i = F_i^{-1} \circ F. \quad (2.7)$$

Since F_1, \dots, F_n are independent, identically distributed and increasing distribution functions, it follows that the T_i are homeomorphisms with $E\{T_i(x)\} = x$.

Indeed, by construction it can be shown that $T_1(x) + \dots + T_n(x) = nx$ for every x , T_1, \dots, T_n are identically distributed given F and so, $E(T_i | F) = E(T_{i'} | F)$ for every $i \neq i'$, and taking expectation in both sides, we have that $E(T_i) = E(T_{i'})$; therefore,

$$nE\{T_i(x)\} = E\{T_1(x)\} + \dots + E\{T_n(x)\} = E\{T_1(x) + \dots + T_n(x)\} = nx, \tag{2.8}$$

and thus it follows that $E\{T_i(t)\} = t$, for $i = 1, \dots, n$.

The random Bernstein polynomial-induced priors on the registered point processes is constructed by pushing them forward through the registration maps

$$\Pi_i = T_{i\#}^{-1} \tilde{\Pi}_i, \quad i = 1, \dots, n. \tag{2.9}$$

The posterior sampling for the warping maps and registered points is then conducted as follows. Let $F_{i,[1]}, \dots, F_{i,[M]}$ be posterior samples from F_i , for $i = 1, \dots, n$, which can be obtained via a Gibbs sampler as described in Appendix C.1. Then, for each $j = 1, \dots, M$ we get $F_{[j]} = (\sum_{i=1}^n F_{i,[j]}^{-1}/n)^{-1}$ and so, $T_{i,[j]} = F_{i,[j]}^{-1} \circ F_{[j]}$ and $\Pi_{[j]} = T_{i,[j]\#}^{-1} \tilde{\Pi}_i$. Finally, pointwise estimation for mean measure, warp functions, and registered points are given by the posterior means,

$$\hat{F} = \frac{1}{M} \sum_{j=1}^M F_{[j]}, \quad \hat{T}_i = \frac{1}{M} \sum_{j=1}^M T_{i,[j]}, \quad \hat{\Pi}_i = \frac{1}{M} \sum_{j=1}^M \Pi_{i,[j]}. \tag{2.10}$$

Credible intervals or pointwise credible bands can be also directly obtained from the relevant quantiles of the corresponding posterior outputs.

2.3. Kolmogorov–Smirnov, Wasserstein, and Kullback–Leibler supports of induced priors

As it will be shown below, full support of the relevant parameters in our setup holds, under conditions on the support of the law of the k_i and on that of $w_{1,k_i}, \dots, w_{k_i,k_i} | k_i$. Extending the assumptions in [14], we assume that the prior probability mass function ρ_i of k_i is positive, that is $\rho_i(k) > 0$ for $i = 1, \dots, n$ and all k , and that $w_{1,k_i}, \dots, w_{k_i,k_i} | k_i$ has a family of conditional densities $l_i(w_{1,k_i}, \dots, w_{k_i,k_i} | k_i) > 0$, for every $(w_{1,k_i}, \dots, w_{k_i,k_i}) \in S_{k_i}$ and for every sequence of independent integer valued random variables $\{k_1, \dots, k_n\}$. Define the supremum norm

$$\|F - H\|_\infty = \sup_{x \in [0,1]} |F(x) - H(x)|.$$

Below, $\mathcal{F} \equiv (F_1, \dots, F_n)$ denotes the joint Bernstein prior and $N_i \equiv \Pi_i([0, 1]) > 0$ is the total number of points in the i th point process, for $i = 1, \dots, n$.

Theorem 1. *Let $F_1, \dots, F_n \stackrel{\text{iid}}{\sim} \pi$ with Fréchet–Wasserstein mean F , and with induced priors T_i and Π_i as defined in (2.7) and (2.9). For any continuous*

strictly increasing $\mathbb{F}_1, \dots, \mathbb{F}_n \in \Delta$, with Fréchet–Wasserstein mean \mathbb{F} , transport maps $\mathbb{T}_i = \mathbb{F}_i^{-1} \circ \mathbb{F}$, and registered discrete measures $P_i = \mathbb{T}_{i\#}^{-1} \tilde{\Pi}_i$, and for any $\varepsilon > 0$ the following events occur with positive probability:

- (a) $\{\mathcal{F} : \|F_j - \mathbb{F}_j\|_\infty < \varepsilon, j = 1, \dots, n\}$, (b) $\{\mathcal{F} : \|F - \mathbb{F}\|_\infty < \varepsilon\}$,
 (c) $\{\mathcal{F} : \|\mathbb{T}_i - \mathbb{T}_i\|_\infty < \varepsilon\}$, (d) $\{\mathcal{F} : d(\Pi_i/N_i, P_i/N_i) < \varepsilon\}$,

for $i = 1, \dots, n$.

Claims (a), (b), and (c) in Theorem 1 respectively state that the joint Bernstein prior, the Fréchet–Wasserstein mean, and the warp functions have large Kolmogorov–Smirnov support. Claim (d) states that the registered point processes have large Wasserstein support. The proof actually shows that the intersection of these four events (a)–(d) has positive probability. While the latter properties may not look surprising ex-post, as their proofs show they are not straightforward facts.

The characterisation of the Kullback–Leibler (KL) support is more challenging. By definition, a density f is said to possess the KL property relatively to a prior π if for any $\varepsilon > 0$ one has that $\pi\{H : \text{KL}(F, H) < \varepsilon\} > 0$, where

$$\text{KL}(F, H) = \int_0^1 h(x) \log \frac{h(x)}{f(x)} dx,$$

with F and H denoting the distribution functions respectively corresponding to f and h .

When π is a random Bernstein polynomial prior as per Definition 1, any continuous density f possesses the KL property [16, Theorem 2]. The following theorem inspects the permanence of the Kullback–Leibler property on the functionals of interest, and it shows that the property is preserved for the Fréchet–Wasserstein mean and the warping functions.

Theorem 2. *Let $F_1, \dots, F_n \stackrel{\text{iid}}{\sim} \pi$ with Fréchet–Wasserstein mean F and with transport maps $T_i = F_i^{-1} \circ F$ as defined in (2.7). For any $\varepsilon > 0$ and strictly increasing $\mathbb{F}_1, \dots, \mathbb{F}_n \in \Delta$ with densities \mathbb{f}_i that are continuous on $(0, 1)$, Fréchet–Wasserstein mean \mathbb{F} and transport maps $\mathbb{T}_i = \mathbb{F}_i^{-1} \circ \mathbb{F}$, \mathbb{F} also has a density \mathbb{f} and:*

- (a) *If $\int_0^1 \mathbb{f}(x) \log \mathbb{f}(x) dx < \infty$ then $\text{KL}(F, \mathbb{F}) < \varepsilon$ with positive probability.*
 (b) *If each \mathbb{f}_i is strictly positive on $(0, 1)$, then there is positive probability that $\text{KL}(T_i, \mathbb{T}_i) < \varepsilon$ for all $i = 1, \dots, n$.*

Remark 1. The densities \mathbb{f}_i can be unbounded or approach zero near 0 or 1. The condition $\int_0^1 \mathbb{f}(x) \log \mathbb{f}(x) dx < \infty$ in (a) is very weak and is satisfied when \mathbb{f} is a beta density with arbitrary (positive) parameters. This condition is, in fact, necessary; if it fails to hold, then $\text{KL}(F, \mathbb{F}) = \infty$ almost surely. The assumptions on the densities can be further relaxed to \mathbb{f}_i having finitely many discontinuity points on $[0, 1]$, and for part (b) \mathbb{f}_i may vanish on finitely many points on $[0, 1]$. We refrained from this level of generality for the purpose of clarity and because the current version includes the most important case of beta distributions.

Theorem 2 shows that under mild conditions, the Fréchet–Wasserstein mean and the warping functions possess the Kullback–Leibler property with respect to the prior on F induced from F_1, \dots, F_n via (2.6). We now study the large-sample behaviour of the posterior.

2.4. Posterior consistency

Contrarily to [12, Theorem 1], our asymptotic theory does not require $n \rightarrow \infty$; indeed we only require that $m_i \rightarrow \infty$, with $i = 1, \dots, n$, for any finite n . Yet note that the consequence is that under this assumption one is only able to approximate warping functions of the type $\mathbb{T}_i = \mathbb{F}_i^{-1} \circ \mathbb{F}$, for all i , where \mathbb{F} is the Fréchet–Wasserstein mean of $\mathbb{F}_1, \dots, \mathbb{F}_n$. This proviso is less and less restrictive as n increases, and it is asymptotically compatible with that of [12], as indeed if the \mathbb{T}_i are independent and identically distributed—rather than fixed as assumed in Theorem 2—then it follows that as $n \rightarrow \infty$,

$$\frac{1}{n} \sum_{i=1}^n \mathbb{T}_i(x) \xrightarrow{p} E\{\mathbb{T}_1(x)\} = x.$$

While the case $n \rightarrow \infty$ and $m \rightarrow \infty$ might also be natural from an asymptotic viewpoint, its study is not trivial and it entails an interplay between n and the m_i 's. Below, the posteriors induced by (2.6) and (2.7) should be understood respectively as the laws of F and T_i conditional on $\{\tilde{x}_{i,j}\}_{j=1}^{m_i}$. The following result holds.

Theorem 3. *Let $\mathbb{F}_1, \dots, \mathbb{F}_n \in \Delta$ be strictly increasing and possess densities f_i that are continuous on $(0, 1)$ with $\int_0^1 f_i(x) \log f_i(x) dx < \infty$. Let \mathbb{F} be their Fréchet–Wasserstein mean, and $\mathbb{T}_i = \mathbb{F}_i^{-1} \circ \mathbb{F}$ the corresponding transport maps. If $F_1, \dots, F_n \stackrel{\text{iid}}{\sim} \pi$ and $m_i \rightarrow \infty$ for $i = 1, \dots, n$, then the posteriors induced by (2.6) and (2.7) are respectively Kolmogorov consistent at \mathbb{F} and \mathbb{T}_i , for all i .*

This result closes the large sample properties of our methods; we next focus on assessing their finite-sample properties.

3. Numerical experiments and computing

We now present the main findings of a Monte Carlo simulation study. Some comments on MCMC and prior specification are in order. Throughout all computations in this paper, we use a burn-in period of 500 iterations and sample 5000 draws from the posterior; we set $(a_0^{G^*}, b_0^{G^*}) = (1, 1)$ for the parameters of the Beta centering, $(a_0^\alpha, b_0^\alpha) = (2, 2)$ for the hyperparameters of the precision parameter of the Dirichlet process prior, and $k_{\max} = 1000$.

3.1. Small n , large m

As our asymptotic theory does not require $n \rightarrow \infty$, we start by assessing performance of the proposed methods in a small n , large m setting. We generate

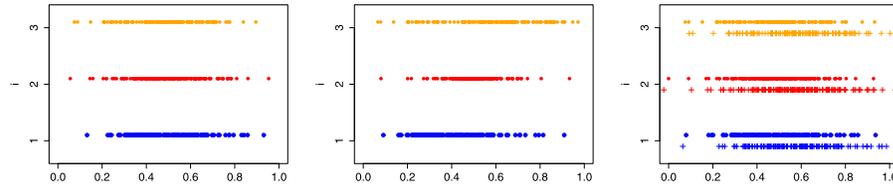


FIG 2. Left: Simulated instances of the original point process from the setup of Section 3.1 in the small n , large m regime. Middle: Corresponding phase-varying versions. Right: Registered versions via proposed approach (\bullet) and kernel ($+$).

random samples $x_{i,1}, \dots, x_{i,m_i} \mid m_i$, from

$$\lambda(x) = \Phi(x \mid 0.5, (0.15)^2), \quad m_i \sim \text{Poisson}(L),$$

for $i = 1, 2, 3$, with $L = 150$ and $\Phi(x \mid \mu, \sigma^2)$ denoting the Normal distribution function. Then the warped data $\tilde{x}_{i,j} = T_i(x_{i,j})$ are obtained using

$$\begin{cases} T_i(x) = x + (a_i - \frac{1}{2}) \sin(b_i x \pi) (b_i \pi)^{-1}, & i = 1, 2, \\ T_3(x) = 3x - T_1(x) - T_2(x), \end{cases} \quad (3.1)$$

where $a_1, a_2 \stackrel{\text{iid}}{\sim} \text{Unif}([0, 1/4] \cup [3/4, 1])$ and $b_1, b_2 \stackrel{\text{iid}}{\sim} \text{Unif}\{1, 2\}$. By construction these warp maps are in line with the model assumptions: each T_i is an increasing homeomorphism of $[0, 1]$, and $E\{T_i(x)\} = x$ follows from the fact that $E(a_i) = 1/2$. See Figure 2 (left and middle) for an instance of realisations of the original point process along with phase-varying versions obtained by warping the data as in (3.1).

The proposed semiparametric approach in Section 2 can be implemented with the aid of the R package Rmpp, which implements a version of the algorithm in [14, p. 383]; see Appendix C.1 for details. Figure 3 shows the estimators of each of the three warp maps through the posterior mean of the induced prior defined in (2.7), along with their credible bands and the true warp maps.

From Figure 3 it can be observed that our estimators are reasonably in line with the true warp functions. Thus, the method recovers quite well the original point processes, as can be seen when comparing the left and right panels of Figure 2.

A Monte Carlo study was conducted in this setting based on $B = 50$ simulated datasets. We apply our method to each, and then calculate the Monte Carlo L^2 -Wasserstein distance mean (WDM) by

$$\text{WDM} = \frac{1}{B} \sum_{b=1}^B \sum_{i=1}^n d(\hat{\Pi}_i^{[b]}, \Pi_i^{[b]}), \quad (3.2)$$

where the superscript $[b]$ denotes the corresponding object computed from the b th simulated dataset, in order to give a performance of our methods when

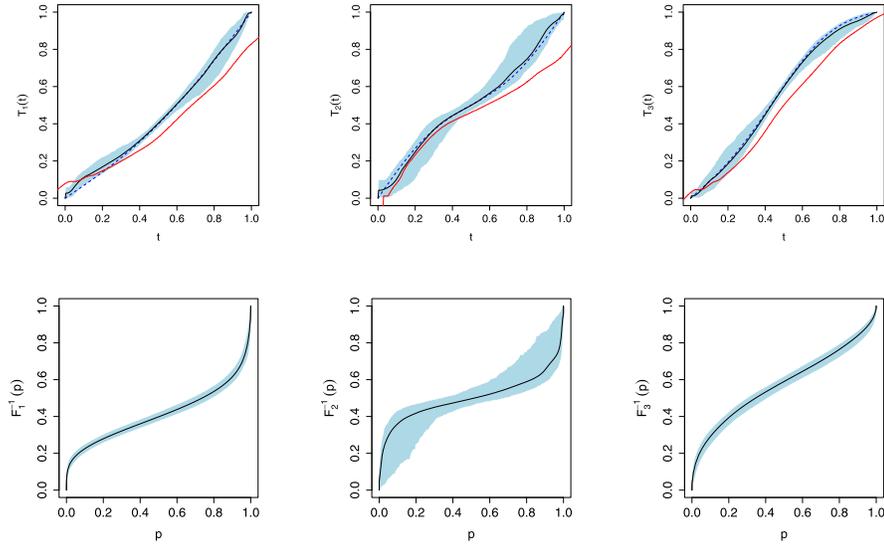


FIG 3. Fits from simulated data from the setup of Section 3.1 in the small n , large m regime. Top: True (dashed blue), kernel smoothing (solid red), and posterior Bernstein polynomial (solid black) warp maps along with credible bands. The Bayesian estimators are constructed as the posterior mean of the induced prior as (2.7). Bottom: Corresponding quantile function estimates along with credible bands.

n is small ($n = 3$) and the m_i 's are large. We obtained a WDM of 0.013. When taking $L = 75$ instead of 150 the WDM is 0.017, in accordance with the intuition that this value decreases with L . For the sake of comparison, if in (3.2) we use $\hat{\Pi}_i$ instead of $\hat{\Pi}_i$, the WDM becomes 0.079 and 0.080 with $L = 75$ and $L = 150$, respectively. Boxplots of $d(\hat{\Pi}_i^{[b]}, \Pi_i^{[b]})$ are given in Appendix D.1 for all i . In Appendix D.2, we also include an additional simulation study suggesting satisfactory performance of the methods under misspecification, with data being warped via biased warp maps (i.e., $E(T) \neq t$).

While the performance of both the Bernstein polynomial estimator and the kernel-based estimator is remarkable, there are situations where both estimators suffer from extrapolation issues on some subintervals; see the following subsection. Mitigating these nontrivial effects is an important avenue for future research.

3.2. Large n , small m

For comparison with [12] we now assess performance over a large n setup. We generate random samples $x_{i,1}, \dots, x_{i,m_i} \mid m_i$, from

$$\lambda(t) = 0.2 \phi(t \mid 0.25, 0.02^2) + 0.8 \phi(t \mid 0.75, 0.03^2), \quad m_i \sim \text{Poisson}(L),$$

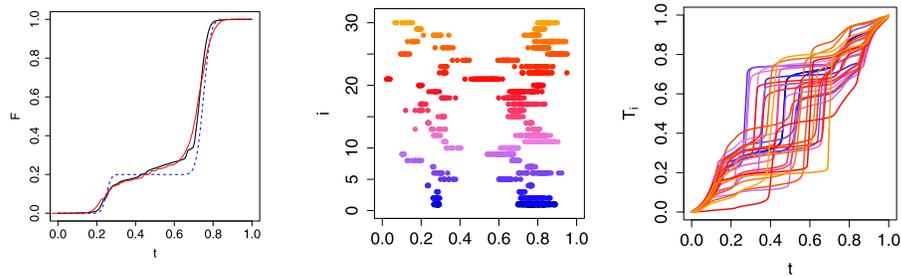


FIG 4. Fits from simulated data from the setup of Section 3.2 in the large n , small m regime. Left: Posterior Bernstein polynomial Fréchet mean (solid black) along with credible bands, kernel smoothing Fréchet mean (solid red), and original Fréchet mean (dashed blue). Middle: Phase-varying point process. Right: Posterior mean Bernstein polynomial warp functions coloured according to the same palette as in Figure 1.

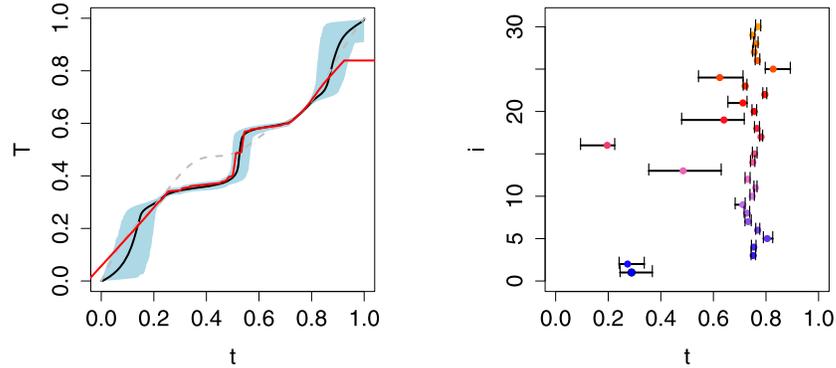


FIG 5. Further fits from simulated data from the setup of Section 3.2 in the large n , small m regime. Left: Posterior mean Bernstein polynomial warp function (solid black) and corresponding credible band, kernel smoothing warp function estimate (solid red), and original warp function (dashed grey) for $i = 5$. Right: Credible intervals for randomly selected registered points for each registered point process.

where $i = 1, \dots, n = 30$, $\phi(t \mid \mu, \sigma^2)$ is the normal density function, and $L = 50$. The warped data $\tilde{x}_{i,j} = T_i(x_{i,j})$ are obtained using

$$T_i(t) \stackrel{D}{=} U \zeta_{K_1}(t) + (1 - U) \zeta_{K_2}(t), \quad \zeta_k(t) = \begin{cases} t, & k = 0, \\ t - \frac{\sin(\pi tk)}{|k|\pi}, & \text{otherwise,} \end{cases}$$

where $U \sim \text{Unif}(0, 1)$, $K_j \stackrel{D}{=} V_1 V_2$ with $V_1 \sim \text{Poisson}(3)$ and $P(V_2 = -1) = P(V_2 = 1) = 1/2$.

We start by illustrating our method on this setup on a single run-experiment; a Monte Carlo study was also conducted in this setting along the same lines as in Section 3.1 and it will also be reported below. A realisation of the original point

process can be found in Figure 1. After estimating F_1, \dots, F_n using random Bernstein polynomials we obtain using (2.10) the posterior mean Fréchet mean depicted in Figure 4, along with pointwise credible bands that are so small that they are hardly visible; we comment further on the small size of such bands in Appendix D.

The posterior mean is quite similar to the kernel-based estimator of [12], and both are similar to the true Fréchet mean. Figure 4 also includes posterior inference for the warp functions. To examine the inference for warp functions in a greater level of detail Figure 5 presents the posterior mean Bernstein polynomial warp function along with credible bands for $i = 5$. As it can be observed from the latter figure, our estimator follows closely that of [12], and is reasonably in line with the original warp function; similar evidence holds for the remainder values of i (see Figure 10 in Appendix D.1). As expected, both estimators have however more difficulty in recovering the true value in the center of the unit interval but this is due to an extrapolation issue as there tends to be much less data on that region. Figure 5 also showcases that our method is more appropriate for bounded domains than the kernel as it takes full advantage of knowledge on the interval where the point processes lie.

While the theoretical claims in Section 3.2 extend those of [12]—in the sense that under extra conditions they support the use of the methods even under a small n large m setting—numerical experiments in Appendix D suggest that the pointwise performance of our methods is tantamount to that of [12]. Figure 5 presents additionally credible intervals for randomly selected registered points for each registered point process. Observe that wider intervals are associated to points falling on the interval separating the two ‘clusters’ of points.

4. Application: tracking phase variation of annual peak temperatures

We now showcase how our method can be used for tracking the phase variation of annual peak temperatures, that is, temperatures above or below a threshold. Throughout this section we will use the same MCMC setup and prior specification as the ones used in Section 3. Peaks of temperature are related with a variety of hazardous events—including heat-related mortality, destruction of crops, wildfires—and have a direct impact on a wealth of economic decisions—such as demand for fuel and electricity. A better understanding of the variation of the regularity of these peaks is thus of the utmost importance from an applied perspective. A main target of our analysis will be on assessing the variation of the onset of temperature peaks, as well as quantifying how atypical is a certain year’s pattern of such peaks. Our analysis has points of contact with the subject of shifts in seasonal cycles (e.g., late start of spring, or growing seasons), which is of wide interest in biology and climatology [e.g., 11, 21]. To illustrate how the method can be used for such purpose we gathered data from “*National Centers for Environmental Information of National Oceanic and Atmospheric Administration (NOAA)*” (<https://www.ncdc.noaa.gov/>), that consist of av-

erage daily air temperatures (in °F, rounded to the nearest integer) of Santiago (Chile) from April, 1990 to March, 2017. Let $\tilde{x}_{i,j}$ be the temperature on day i , year j . Below, we focus on the point processes of annual peaks over threshold, $\{\tilde{x}_{i,j}^+ \geq u_j^+\}$, and annual peaks below threshold, $\{\tilde{x}_{i,j}^- \leq u_j^-\}$; in practice we set the thresholds u_j^+ and u_j^- using the 95% and 5% quantiles of temperature over year j , and this results in m_1^+, \dots, m_n^+ and m_1^-, \dots, m_n^- ranging from 19 to 32. Appendix E includes a sensitivity analysis based on the 97.5% and 2.5% quantiles; the main empirical findings are tantamount to the ones presented here. In Figure 6 we present the point processes of interest along with the corresponding warping functions for peaks above the threshold (T_j^+) and peaks below the threshold (T_j^-). For the analysis of annual peaks over threshold, we fully support the warping functions between the minimum and maximum times corresponding to the pooled exceedances above the threshold; we proceed analogously for the analysis of annual peaks below the threshold. Here the interest is on the highest and lowest temperatures per year, which is tantamount to considering a fixed threshold per year (i.e. u_j^+ and u_j^-). Yet the analysis can be easily extended to the situation where the goal is on modeling conditionally high/low temperatures (i.e. taking into account seasonal variation) via a time-varying threshold that can be obtained, for example, via quantile regression [8].

To interpret Figure 6 we first focus on annual peaks below the threshold, for which there are at least two patterns of points that readily look unusual to the naked eye: 1991, for which there was an atypical cold weather event almost taking place in the summer; and 2010, given that lower temperatures peaked later on a concentrated period. The fact that these patterns of points look unusual agrees with what can be observed from the corresponding warping functions, that are among the ones that further deviate from the identity; cf. Figures 15 and 16 in Appendix E. In terms of peaks above the threshold, note how the antepenultimate pattern of points started much later than all the remainder, thus meaning that higher temperatures peaked much later than expected.

To assess how atypical is the climatological pattern of onset of peaks, we define the following measures to which we refer as scores of peak irregularity (SPI), and for temperatures above and below a threshold are respectively defined as

$$\text{SPI}^+ = \int_0^1 |T_j^+(t) - t| dt, \quad \text{SPI}^- = \int_0^1 |T_j^-(t) - t| dt; \quad (4.1)$$

to combine peaks over and below a threshold, we also define a global SPI = $(\text{SPI}^+ + \text{SPI}^-)/2$. Figure 7 depicts the scores of peak irregularity over time for peaks above and below a threshold; as can be seen from the latter figure, the fits from the proposed approach differ slightly from those of the kernel due to the fact that the kernel does not take into account that the support is the unit interval (see Figure 6). To shed light on interpretation of Figure 7 we note that if the climatological pattern of the onset of peaks above or below the threshold was always the same, then all SPIs would be equal to zero. The ranking of the SPIs, on the other hand, quantifies which patterns of onset of peaks are the most anomalous—with the largest SPI, for the most atypical year. Figure 7 is coherent

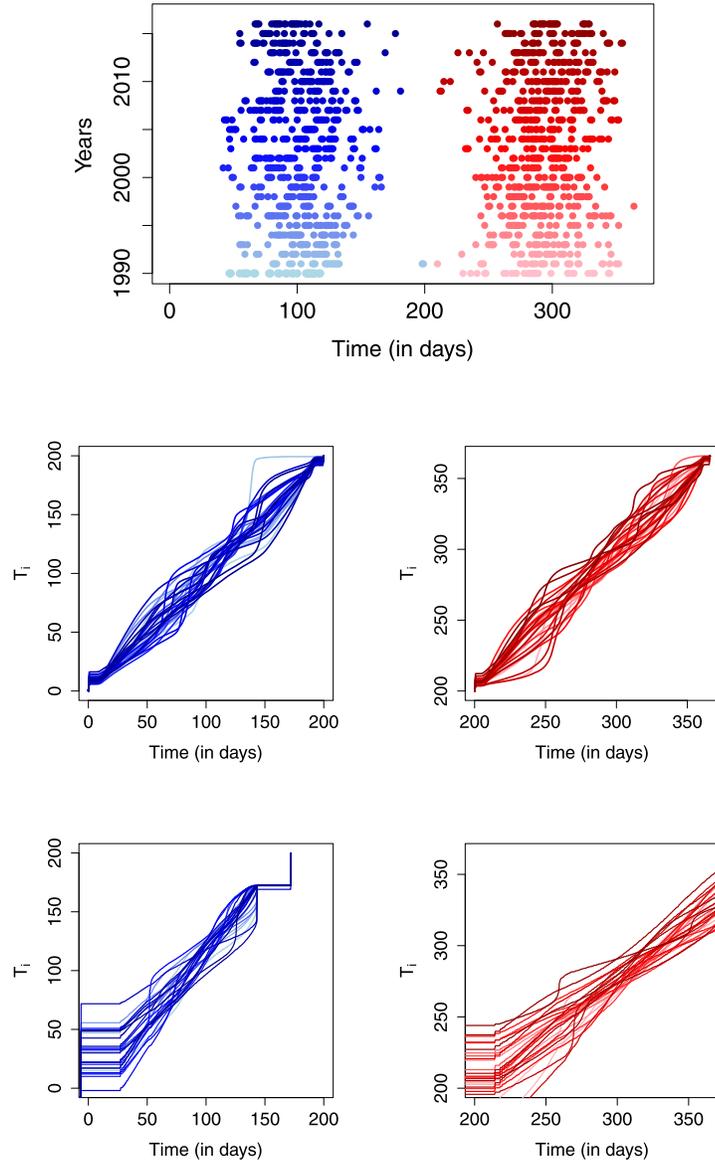


FIG 6. *Top: Point processes of annual peaks for peaks above (red) and below (blue) the temperature thresholds. Middle: Corresponding posterior Bernstein polynomial mean warp functions in the same palette of colours. Bottom: Corresponding kernel smoothing warp functions in the same palette of colours.*

with what was expected given the comments above surrounding Figure 6 on the patterns of points that looked immediately atypical, and on the shape of the corresponding warping functions.

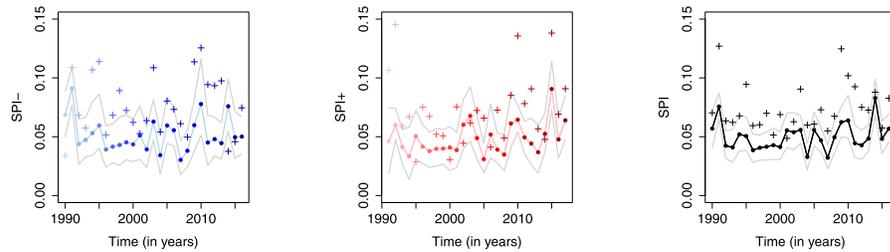


FIG 7. Posterior mean SPI (scores of peak irregularity, see (4.1)), along with credible intervals, below a cold temperature threshold (Left), above a hot temperature threshold (Middle), and global (Right). Corresponding SPI for kernel smoothing (+) in the same palette of colours.

5. Closing remarks

We propose a semiparametric Bayesian approach for the purpose of separating amplitude and phase variation in point process data. This paradigm has the advantage of providing a straightforward construction of credible sets via the posterior distribution, and in particular, we are able to quantify the uncertainty in learning not only the structural mean measure λ , but also the warping functions T_i and the latent point processes Π_i . The Bernstein–Dirichlet prior interweaves elegantly with the Wasserstein geometry of optimal transport. Indeed, its favourable support properties [as established by 16] carry over to the induced priors on the structural mean measure λ and all sufficiently regular warping functions, allowing to obtain Bayesian consistency in a genuinely infinite-dimensional setup.

An interesting question would be how to extend this work to the case of spatial point process supported on e.g., $[0, 1]^D$ with $D > 1$, as explored by [3] and [26]; a natural extension of our paper to this setup would entail modelling the mean measures of the corresponding spatial point processes via multivariate Bernstein polynomials [27]. The computation of the empirical Fréchet–Wasserstein mean can no longer however be done in closed form, requiring numerical schemes [17]. From a statistical viewpoint, another avenue for future research would be on modelling the phase variation of point processes conditionally on a covariate, by resorting to predictor-dependent versions of the Bernstein–Dirichlet prior [2].

Appendix A: Proofs of main results

We begin by stating a number of auxiliary lemmata that will be useful to deriving our main results. Lemma 1 is often known as Pólya’s theorem [9, Theorem 11.2.9]. Lemma 2 states that inversion is continuous in supremum norm [9, Lemma 11.2.1]. Lemma 3 discusses sufficient conditions for (local) uniform convergence of the Bernstein polynomial density; see [6, Lemma E.3] for a related result under further smoothness assumptions on f . As the proof of Lemma 3

shows, the uniform convergence holds on any set bounded away from the discontinuity points of f . For the sake of completeness, proofs are given in Appendix B.

Lemma 1. *Let \mathbb{F} be a continuous distribution function and let F_n be a sequence of distribution functions that converge weakly to \mathbb{F} . Then $\|F_n - \mathbb{F}\|_\infty \rightarrow 0$.*

Lemma 2. *Let $\mathbb{F} : [0, 1] \rightarrow [0, 1]$ be continuous, strictly increasing and with $F(0) = 0, F(1) = 1$. Then \mathbb{F}^{-1} is also continuous and strictly increasing, and for any $\varepsilon > 0$ there exists $\delta > 0$ such that for any continuous strictly increasing $H : [0, 1] \rightarrow [0, 1]$:*

1. *If $\|\mathbb{F} - H\|_\infty < \delta$, then $\|\mathbb{F}^{-1} - H^{-1}\|_\infty < \varepsilon$.*
2. *If $\|\mathbb{F}^{-1} - H^{-1}\|_\infty < \delta$, then $\|\mathbb{F} - H\|_\infty < \varepsilon$.*

Lemma 3. *Let $\mathbb{F} : [0, 1] \rightarrow \mathbb{R}$ be differentiable with derivative f that is continuous on $(0, 1)$. Then for any $a > 0, b(x | k, \mathbb{F})$ as defined in (2.2) converges to f uniformly on $[a, 1-a]$. If f is continuous on $[0, 1]$, then as $k \rightarrow \infty, b(x | k, \mathbb{F}) \rightarrow f$ uniformly on $[0, 1]$.*

Proof of Theorem 1.

- (a) The proof follows from Theorem 3 in [15], combined with the fact that by assumption $F_1, \dots, F_n \stackrel{\text{iid}}{\sim} \pi$. Indeed,

$$\pi^{(n)}\{\mathcal{F} : \|F_j - \mathbb{F}_j\|_\infty < \varepsilon, j = 1, \dots, n\} = \prod_{j=1}^n \pi\{F_j : \|F_j - \mathbb{F}_j\|_\infty < \varepsilon\},$$

which is strictly positive.

- (b) From (a) and Lemma 2 it follows that

$$\pi^{(n)}\{\mathcal{F} : \|F_i^{-1} - \mathbb{F}^{-1}\|_\infty < \eta, i = 1, \dots, n\} > 0, \quad \eta > 0. \quad (\text{A.1})$$

Also, note that

$$\|F^{-1} - \mathbb{F}^{-1}\|_\infty = \left\| \frac{1}{n} \sum_{i=1}^n F_i^{-1} - \mathbb{F}^{-1} \right\|_\infty \leq \frac{1}{n} \sum_{i=1}^n \|F_i^{-1} - \mathbb{F}^{-1}\|_\infty. \quad (\text{A.2})$$

From (A.2) and Lemma 2, it follows that to have $\|F - \mathbb{F}\|_\infty < \varepsilon$ it would suffice having $\|F_i^{-1} - \mathbb{F}^{-1}\|_\infty < \delta$ for all i , thus implying that

$$\pi^{(n)}\{\mathcal{F} : \|F - \mathbb{F}\|_\infty < \varepsilon\} \geq \pi^{(n)}\{\mathcal{F} : \|F_i^{-1} - \mathbb{F}^{-1}\|_\infty < \delta\} > 0.$$

- (c) Lemma 2 and the assumption \mathbb{F}_i is (uniformly) continuous on $[0, 1]$ imply that \mathbb{F}_i^{-1} is also uniformly continuous, for $i = 1, \dots, n$. Given $\eta > 0$, let $\delta > 0$ such that $|t - s| \leq \delta \Rightarrow |\mathbb{F}_i^{-1}(t) - \mathbb{F}_i^{-1}(s)| \leq \eta$, for $i = 1, \dots, n$. From (a) and (b) it respectively follows that

$$\pi^{(n)}\{\mathcal{F} : \|F_i^{-1} - \mathbb{F}_i^{-1}\|_\infty \leq \eta, i = 1, \dots, n\} > 0$$

and $\pi^{(n)}\{\mathcal{F} : \|F - \mathbb{F}\|_\infty \leq \delta\} > 0$. Thus, $\pi^{(n)}\{\mathcal{F} : |F(x) - \mathbb{F}(x)| \leq \delta, x \in [0, 1]\} > 0$, and this implies that the event

$$\begin{cases} F_i^{-1}(F(x)) \leq F_i^{-1}(\mathbb{F}(x) + \delta) \leq \mathbb{F}_i^{-1}(\mathbb{F}(x) + \delta) + \eta \leq \mathbb{T}_i(x) + 2\eta, \\ F_i^{-1}(F(x)) \geq F_i^{-1}(\mathbb{F}(x) - \delta) \geq \mathbb{F}_i^{-1}(\mathbb{F}(x) - \delta) - \eta \geq \mathbb{T}_i(x) - 2\eta, \end{cases} \tag{A.3}$$

occurs with positive probability, for $i = 1, \dots, n$. This thus yields that

$$\pi^{(n)}\{\mathcal{F} : \|T_i - \mathbb{T}_i\|_\infty \leq 2\eta\} > 0, \quad i = 1, \dots, n.$$

(d) The strategy of the proof is similar to that of [12, p. 798]. We start by noting that $T_i^{-1} \circ \mathbb{T}_i \in \Gamma(\Pi_i/N_i, P_i/N_i)$ as a consequence of

$$\Pi_i = T_i^{-1} \tilde{\Pi}_i = (T_i^{-1} \circ \mathbb{T}_i) \# P_i, \quad i = 1, \dots, n.$$

It thus follows that

$$\begin{aligned} d^2(\Pi_i/N_i, P_i/N_i) &\leq \int_0^1 \{(T_i^{-1} \circ \mathbb{T}_i)(x) - x\}^2 \frac{\Pi_i(dx)}{N_i} \\ &\leq \|\{T_i^{-1} \circ \mathbb{T}_i - x\}^2\|_\infty. \end{aligned}$$

To complete the proof just note that (c) implies that for all i

$$\pi^{(n)}\{\mathcal{F} : \|T_i^{-1} \circ \mathbb{T}_i - x\|_\infty < \varepsilon\} = \pi^{(n)}\{\mathcal{F} : \|T_i^{-1} - \mathbb{T}_i\|_\infty < \varepsilon\} > 0,$$

from where the final result follows.

Proof of Theorem 2. The derivatives of the induced priors (2.6) and (2.7) will be required for the proofs, and are respectively

$$f(x) = n \left(\sum_{i=1}^n \frac{1}{f_i(T_i(x))} \right)^{-1}, \quad T_i'(x) = \frac{f'(t)}{f_i(T_i(x))}, \quad i = 1, \dots, n, \quad f_i = F_i'.$$

(a) Let f_i be the density corresponding to \mathbb{F}_i and f that corresponding of \mathbb{F} . Then

$$|f(x) - \mathbb{f}(x)| = n \left| \left(\sum_{i=1}^n \frac{1}{f_i(T_i(x))} \right)^{-1} - \left(\sum_{i=1}^n \frac{1}{\mathbb{f}_i(\mathbb{T}_i(x))} \right)^{-1} \right|. \tag{A.4}$$

We first assume that $\inf f_i \geq 2l > 0$ for all i , and consequently $\inf f > 2l$ as well. For $g : [0, 1] \rightarrow \mathbb{R}$ and $1/2 > a > 0$ denote $\|g\|_{\infty, a} = \sup_{x \in [a, 1-a]} |g(x)|$. We shall show that the event

$$\Omega_{a, \rho} = \{f_i \geq l \ \& \ \|f_i - \mathbb{f}_i\|_{\infty, a} < \rho, i = 1, \dots, n\},$$

has positive probability for all $a, \rho > 0$. Let k_i be large so that $\|b(x | k_i, f_i) - \mathbb{f}_i\|_{\infty, a} < \rho/2$ (using Lemma 3), set $k = \max_i k_i$ and denote

$b(x \mid k, \mathbb{f}_i) = \sum_{j=1}^k w_{i,j} \beta(x \mid j, k - j + 1)$. The set of polynomials with slightly perturbed coefficients

$$\mathcal{P}_{i,\delta} = \left\{ p = \sum_{j=1}^k w'_{i,j} \beta(x \mid j, k - j + 1) : (w'_{i,1}, \dots, w'_{i,k}) \in S_k \right. \\ \left. \text{with } |w'_{i,j} - w_{i,j}| < \delta, \text{ for all } j \right\}$$

has positive probability under the Bernstein polynomial prior, for all $\delta > 0$, as a consequence of [16, p. 84] because the set where $(w'_{i,1}, \dots, w'_{i,k})$ lies is open in the unit simplex. Moreover, each $p \in \mathcal{P}_{i,\delta}$ satisfies

$$\|p - b(x \mid k, \mathbb{f}_i)\|_\infty \leq \delta k \max_{1 \leq j \leq k} \sup_x \beta(x \mid j, k - j + 1) < \infty$$

because $1 \leq j \leq k$. Thus for small enough δ , $\|p - b(x \mid k, \mathbb{F}_i)\|_{\infty,a} < \rho/2$. Since the F_i 's are independent, there is a positive probability that $f_i \in \mathcal{P}_{i,\delta}$ for all i , which implies that $\|f_i - \mathbb{f}_i\|_{\infty,a} < \rho$ for all i . Moreover, as $\mathbb{f}_i \geq 2l$, $w_{i,j} \geq 2l/k$ and if $\delta < l/k$ this yields $w'_{i,j} \geq l/k$ and thus $b(x \mid k, \mathbb{F}_i) \geq l$. Hence $\Omega_{a,\rho}$ has positive probability.

Fix $\epsilon > 0$; we wish to show that $\|F_i - \mathbb{F}_i\|_\infty \leq \epsilon$ holds on $\Omega_{a,\rho}$ for appropriate $a, \rho > 0$. Let $1/2 > a > 0$ such that $\mathbb{F}_i(a) < \epsilon/3$ and $\mathbb{F}_i(1-a) > 1 - \epsilon/3$, and let $\rho < \epsilon/3$. When $\Omega_{a,\rho}$ holds, we have

$$1 \geq F_i(1-a) = F_i(a) + \int_a^{1-a} f_i(x) dx \\ \geq F_i(a) + \int_a^{1-a} \mathbb{f}_i(x) dx - \rho(1-2a) \\ = F_i(a) + \mathbb{F}_i(1-a) - \mathbb{F}_i(a) - \rho(1-2a).$$

Thus $-\epsilon \leq F_i(a) - \mathbb{F}_i(a) \leq 1 - \mathbb{F}_i(1-a) + \rho(1-2a) < 2\epsilon/3$. For $x \leq a$ we have

$$-\epsilon \leq F_i(x) - \mathbb{F}_i(x) \leq F_i(a) - \mathbb{F}_i(x) \\ \leq 1 - \mathbb{F}_i(1-a) + \rho(1-2a) + \mathbb{F}_i(a) - \mathbb{F}_i(x) \leq \epsilon.$$

Thus $|F_i - \mathbb{F}_i| \leq \epsilon$ on $[0, a]$ and by a similar argument the same holds on $[1-a, 1]$. For $x \in [a, 1-a]$ observe that

$$|F_i(x) - \mathbb{F}_i(x)| \leq |F_i(a) - \mathbb{F}_i(a)| + \int_a^x |f_i(y) - \mathbb{f}_i(y)| dy \\ \leq |F_i(a) - \mathbb{F}_i(a)| + \rho < \epsilon.$$

Conclude that $\|F_i - \mathbb{F}_i\|_\infty \leq \epsilon$. As in the proof of Theorem 1 we have as a consequence that for sufficiently small a and ρ , on $\Omega_{a,\rho}$ $\|F_i^{-1} \circ F - \mathbb{F}_i^{-1} \circ \mathbb{F}\|_\infty < \epsilon$. Fix $a, \rho_2 \in (0, 1/2)$. Let $c_i = \min(\mathbb{F}_i^{-1}(\mathbb{F}(a)), 1 - \mathbb{F}_i^{-1}(\mathbb{F}(1-a)))$

and $a_1 = \min_i c_i/2$. Since \mathbb{f}_i is uniformly continuous on $[a_1, 1 - a_1]$, there exists $\delta_2 > 0$ such that $|\mathbb{f}_i(x) - \mathbb{f}_i(y)| \leq \rho_2$ for all $x, y \in [a_1, 1 - a_1]$ such that $|x - y| \leq \delta_2$; without loss of generality $\delta_2 \leq a_1$. Choose small $a_1 > a_2, \rho > 0$ such that on $\Omega_{a_2, \rho}$, $\|F_i^{-1} \circ F - \mathbb{F}_i^{-1} \circ \mathbb{F}\|_\infty < \delta_2$. Then on $\Omega_{a_2, \rho}$

$$\|f_i \circ F_i^{-1} \circ F - \mathbb{f}_i \circ \mathbb{F}_i^{-1} \circ \mathbb{F}\|_{\infty, a} \leq \|f_i - \mathbb{f}_i\|_{\infty, a_1} \leq \rho$$

and

$$\|\mathbb{f}_i \circ F_i^{-1} \circ F - \mathbb{f}_i \circ \mathbb{F}_i^{-1} \circ \mathbb{F}\|_{\infty, a} \leq \sup_{x, y \in [a_1, 1 - a_1], |x - y| \leq \delta_2} |\mathbb{f}_i(x) - \mathbb{f}_i(y)| \leq \rho_2.$$

This means that for any $\rho, \rho_2, a > 0$ there is positive probability that for all $i = 1, \dots, n$

$$\|f_i \circ F_i^{-1} \circ F - \mathbb{f}_i \circ \mathbb{F}_i^{-1} \circ \mathbb{F}\|_{\infty, a} \leq \rho + \rho_2,$$

and since $\Omega_{a_2, \rho}$ implies also that $f_i, \mathbb{f}_i \geq l$, it follows that for all $a, \epsilon > 0$ there is positive probability that $\|f - \mathbb{f}\|_{\infty, a} < \epsilon$. Now write

$$\text{KL}(F, \mathbb{F}) = \int_{x \in [a, 1 - a]} \mathbb{f} \log \frac{\mathbb{f}}{f} + \int_{x \notin [a, 1 - a]} \mathbb{f} \log \frac{\mathbb{f}}{f} = \text{KL}_1 + \text{KL}_2.$$

The definition of $\Omega_{a, \rho}$ implies that on this event $f \geq l$. Hence

$$\begin{aligned} \text{KL}_2 &= \int_{x \notin [a, 1 - a]} \mathbb{f} \log \mathbb{f} - \int_{x \notin [a, 1 - a]} \mathbb{f} \log f \\ &\leq \int_{x \notin [a, 1 - a]} \mathbb{f} \log \mathbb{f} - [1 - F(1 - a) + F(a)] \log l, \end{aligned}$$

and this vanishes as $a \rightarrow 0$ because $\int \mathbb{f} \log \mathbb{f} < \infty$. Hence we can pick $a > 0$ such that $\text{KL}_2 < \epsilon$. To bound KL_1 notice that when $\epsilon < l \leq \inf \mathbb{f}$, and $\|f - \mathbb{f}\|_{\infty, a} \leq \epsilon$, $\|\log \frac{\mathbb{f}}{f}\|_{\infty, a} \leq \log \frac{\inf \mathbb{f}}{\inf \mathbb{f} - \epsilon} \leq \log \frac{l}{l - \epsilon}$. Thus, for all $\epsilon > 0$ we have with positive probability

$$\text{KL}(F, \mathbb{F}) \leq \epsilon + \log \frac{l}{l - \epsilon}.$$

As this vanishes when $\epsilon \rightarrow 0$, the proof is complete under the assumption that $\inf \mathbb{f}_i > 0$ for all i . This assumption can be relaxed as in [16, p. 85][†]: take any \mathbb{f} as in the statement of the theorem and define $\mathbb{f}^a(x) = \max(\mathbb{f}(x), a)/A$, where $A = \int_0^1 \max(\mathbb{f}(x), a) dx \in [1, 1 + a]$. Then $\mathbb{f} \leq A\mathbb{f}^a$ and consequently $\int_0^1 \mathbb{f}^a(x) \log \mathbb{f}^a(x) dx < \infty$. Applying the theorem to $\mathbb{f}_1 = \dots = \mathbb{f}_n = \mathbb{f}^a$ we deduce the KL property for \mathbb{f}^a . Now, as $\mathbb{f} \leq A\mathbb{f}^a$ we have [5, Lemma 5.1]

$$\text{KL}\left(\int h, \int \mathbb{f}\right) \leq (A + 1) \log A$$

[†]beware that they denote $\text{KL}(F, \mathbb{F})$ by $\text{KL}(\mathbb{F}, F)$

$$+ A \left[\text{KL} \left(\int h, \int \mathbb{f}^a \right) + \sqrt{\text{KL} \left(\int h, \int \mathbb{f}^a \right)} \right].$$

As $a \searrow 0$, $A \searrow 1$. If we choose $a > 0$ such that $A < 2$ and $(A + 1) \log A < \varepsilon/3$, and then $\delta > 0$ such that $\delta + \sqrt{\delta} < \varepsilon/3$ then

$$\left\{ h : \text{KL} \left(\int h, \int \mathbb{f} \right) \leq \varepsilon \right\} \supseteq \left\{ h : \text{KL} \left(\int h, \int \mathbb{f}^a \right) \leq \delta \right\},$$

and the latter has positive prior probability. This completes the proof.

- (b) Again begin with the assumption that $\inf \mathbb{f}_i > 0$ for all i . Let $T'_i(x) = f(x)/f_i(T_i(x))$ and $\mathbb{T}'_i(x) = \mathbb{f}(x)/\mathbb{f}_i(\mathbb{T}_i(x))$, and note that

$$|T'_i(x) - \mathbb{T}'_i(x)| \leq \frac{|f(x) - \mathbb{f}(x)|}{f_i(T_i(x))} + \mathbb{f}(x) \left| \frac{1}{f_i(T_i(x))} - \frac{1}{\mathbb{f}_i(\mathbb{T}_i(x))} \right|.$$

For all $a, \epsilon > 0$, since \mathbb{f} is bounded on $[a, 1 - a]$, the same idea as in part (a) shows that with positive probability $\|T'_i - \mathbb{T}'_i\|_{\infty, a} < \epsilon$. Write again

$$\text{KL}(T_i, \mathbb{T}_i) = \int_{x \in [a, 1-a]} \mathbb{T}'_i \log \frac{\mathbb{T}'_i}{T'_i} + \int_{x \notin [a, 1-a]} \mathbb{T}'_i \log \frac{\mathbb{T}'_i}{T'_i} = \text{KL}_1 + \text{KL}_2.$$

These two terms can be made small as in part (a) because $\mathbb{T}'_i \leq n$.

To relax the condition $\inf \mathbb{f}_i > 0$ we use a similar idea as for part (a) but the argument is more subtle. Fix $a > 0$ and define

$$A_i = \int_0^1 \max(\mathbb{T}'_i(x), a) dx, \quad h_i^a(x) = \max(\mathbb{T}'_i(x), a)/A_i,$$

and $H_i^a(x) = \int_0^x h_i^a(t) dt$. For brevity we omit the dependence of h_i, H_i and A_i on a . Clearly H_i is strictly increasing, differentiable almost surely with derivative bounded below by a/A_i , $H_i(0) = 0$ and $H_i(1) = 1$. Moreover h_i is continuous and strictly positive on $(0, 1)$ because so is \mathbb{T}'_i . We shall view H_i as transport maps from a Fréchet mean to well-behaved measures; first we need to fix the issue that they do not necessarily average to the identity by adding another transport map that corrects the discrepancy. By assumption

$$A_i \leq \int_0^1 (\mathbb{T}'_i(x) + a) dx = \mathbb{T}_i(1) - \mathbb{T}_i(0) + a = 1 + a$$

and similarly $A_i \geq 1$. Thus we can choose $a > 0$ small such that $(1 + a)/A_i \leq 1 + 1/(2n)$ for all $i = 1, \dots, n$. Define the correction function

$$H_{n+1}(x) = (n + 1)x - \sum_{i=1}^n H_i(x).$$

Then $H_i, i = 1, \dots, n + 1$ average to the identity. Since $\mathbb{T}_i, i = 1, \dots, n$ average to the identity, whenever they are differentiable (that is, Lebesgue

almost everywhere since they are nondecreasing) we have $\sum_{i=1}^n \mathbb{T}'_i(x) = n$. Hence H_{n+1} is differentiable almost surely with derivative

$$\begin{aligned} n + 1 - \sum_{i=1}^n h_i(x) &\geq n + 1 - \sum_{i=1}^n \frac{\mathbb{T}'_i(x)}{A_i} - \frac{na}{A_i} \\ &\geq n + 1 - n \frac{1+a}{A_i} \geq n + 1 - n \left(1 + \frac{1}{2n}\right) = \frac{1}{2}. \end{aligned}$$

Now consider the distribution functions $\mathbb{G}_i = H_i^{-1}$, $i = 1, \dots, n + 1$ and let \mathbb{G} denote the identity. Then \mathbb{G}_i have Fréchet mean \mathbb{G} with densities bounded above by $\max(2, A_i/a)$ and below by $1/(n + 1)$. Therefore, by the previous part of the proof $H_i^a = \mathbb{G}_i^{-1} \circ \mathbb{G}$ is in the KL support of the induced Bernstein polynomial prior. Since $\mathbb{T}'_i \leq A_i h_i^a$ almost surely we have [5, Lemma 5.1]

$$\begin{aligned} \text{KL}(S, \mathbb{T}_i) &\leq (A_i + 1) \log A_i + A_i [\text{KL}(S, H_i^a) + \sqrt{\text{KL}(S, H_i^a)}] \\ &\leq (a + 2) \log(a + 1) + (a + 1) [\text{KL}(S, H_i^a) + \sqrt{\text{KL}(S, H_i^a)}]. \end{aligned}$$

As $(H_i^a)'$ is continuous and strictly positive on $(0, 1)$, $\text{KL}(S, H_i^a)$ can be made as small as we wish with positive probability. The fact that $a > 0$ is arbitrary completes the proof.

Proof of Theorem 2.4. Under the given assumptions the prior on F_i satisfies the Kullback–Leibler property at \mathbb{F}_i (use Theorem 2(a) with $n = 1$) and consequently the sequence of posteriors are weakly consistent for each \mathbb{F}_i . The operations

$$(F_1, \dots, F_n) \mapsto (F_1^{-1}, \dots, F_n^{-1}) \mapsto \left[F^{-1} = \frac{1}{n} \sum_{i=1}^n F_i^{-1} \right] \mapsto F,$$

are continuous in the supremum norm around $(\mathbb{F}_1, \dots, \mathbb{F}_n)$ by Lemma 2, (A.2) and again Lemma 2. Taking into account the equivalence of the supremum norm with weak convergence (Lemma 1), conclude that the operation $(F_1, \dots, F_n) \mapsto F$ is weakly continuous around $(\mathbb{F}_1, \dots, \mathbb{F}_n)$. Since each F_i is weakly consistent for \mathbb{F}_i , this yields that F is weakly (in fact, Kolmogorov) consistent for \mathbb{F} .

Weak (and Kolmogorov) consistency of T_i to \mathbb{T}_i follows in the same way, since in Equation (A.3) it has been established that

$$(F_i^{-1}, F) \mapsto F_i^{-1} \circ F$$

is continuous in supremum norm around $(\mathbb{F}_i, \mathbb{F})$.

Appendix B: Proofs of auxiliary lemmata

Proof of Lemma 1. Since \mathbb{F} is continuous $F_n \rightarrow \mathbb{F}$ pointwise. Let $\epsilon > 0$ and let $x < y$ such that $\mathbb{F}(x) \leq \epsilon$ and $\mathbb{F}(y) \geq 1 - \epsilon$. Since \mathbb{F} is uniformly continuous on

$[x, y]$ there exists a finite grid $x = x_1 < \dots < x_k = y$ with $\mathbb{F}(x_i) \geq \mathbb{F}(x_{i+1}) - \epsilon$ for all $i \leq k - 1$. For n large $|F_n(x_i) - \mathbb{F}(x_i)| \leq \epsilon$ for all i so that

$$\begin{aligned} \sup_{z \in [x_i, x_{i+1}]} F_n(z) - \mathbb{F}(z) &\leq F_n(x_{i+1}) - \mathbb{F}(x_i) \\ &\leq |F_n(x_{i+1}) - \mathbb{F}(x_{i+1})| + |\mathbb{F}(x_{i+1}) - \mathbb{F}(x_i)| \leq 2\epsilon. \end{aligned}$$

In the same way

$$\sup_{z \notin [x, y]} |F_n(z) - \mathbb{F}(z)| \leq 2\epsilon, \quad \sup_{z \in [x_i, x_{i+1}]} \mathbb{F}(z) - F_n(z) \leq \mathbb{F}(x_{i+1}) - F_n(x_i) \leq 2\epsilon,$$

and we conclude that $\|F_n - \mathbb{F}\|_\infty \leq 2\epsilon$ for n sufficiently large. □

Proof of Lemma 2. Since \mathbb{F} is bijective, it has an inverse \mathbb{F}^{-1} . The latter is nondecreasing and, being a bijection, must also be continuous and with $\mathbb{F}^{-1}(0) = 0$, $\mathbb{F}^{-1}(1) = 1$. Let $p \in (0, 1)$, and let $x \in (0, 1)$ such that $\mathbb{F}(x) = p$. For $\epsilon \in (0, 1 - p)$ we have $F_n(x + \epsilon) \rightarrow \mathbb{F}(x + \epsilon) > p$, which means that $x + \epsilon \geq F_n^{-1}(p)$ for n large. Similarly, $x - \epsilon \leq F_n^{-1}(p)$ for any $\epsilon \in (0, p)$ and all n large. This implies that $F_n^{-1}(p) \rightarrow x = \mathbb{F}^{-1}(p)$ for all $p \in (0, 1)$. Since

$$0 \leq F_n^{-1}(0) \leq F_n^{-1}(p) \xrightarrow{n \rightarrow \infty} \mathbb{F}^{-1}(p) \xrightarrow{p \rightarrow 0} \mathbb{F}^{-1}(0) = 0,$$

it also follows that $F_n^{-1}(0) \rightarrow \mathbb{F}^{-1}(0)$. Similarly $F_n^{-1}(1) \rightarrow \mathbb{F}^{-1}(1)$ and we conclude that $F_n^{-1} \rightarrow \mathbb{F}^{-1}$ pointwise on $[0, 1]$. By Lemma 1 the convergence is uniform. Convergence of sequences is equivalent to the statement of the lemma because the supremum norm defines a metric space.

Part b) is shown in the same way, since $(\mathbb{F}^{-1})^{-1} = \mathbb{F}$. There is a slight complication though because F_n^{-1} is only defined on $[F_n(0), F_n(1)]$ which may be a strict subinterval of $[0, 1]$. Let $x = \mathbb{F}^{-1}(p)$ for $x, p \in (0, 1)$. Then $F_n^{-1}(p - \epsilon) \rightarrow \mathbb{F}^{-1}(p - \epsilon) > x$ for $\epsilon > 0$ small, which means in particular that $F_n(0) \leq p - \epsilon$ and $F_n^{-1}(p - \epsilon)$ is defined, and also that $F_n(x) \leq p - \epsilon$ for n large. The inequality $F_n(x) \geq p + \epsilon$ is shown in the same way and we conclude that $F_n \rightarrow \mathbb{F}$ pointwise, and hence uniformly on $[0, 1]$ by Lemma 1. □

Proof of Lemma 3. Since \mathbb{F} is differentiable, there exists $x_j^* \in [j, j + 1]/k$ such that

$$b(x \mid k, \mathbb{F}) = B(x \mid k - 1, \mathbb{f}) + \sum_{j=0}^{k-1} \left[\mathbb{f}(x_j^*) - \mathbb{f}\left(\frac{j}{k-1}\right) \right] \binom{k-1}{j} x^j (1-x)^{k-1-j}.$$

Notice that $|x_j^* - j/(k-1)| \leq j/(k-1) - j/k \leq 1/(k-1) \rightarrow 0$ uniformly in j and as \mathbb{f} is uniformly continuous on $[a - 1/k, 1 - a + 1/k]$ for all $k > 1/a$, the sum at the right-hand side vanishes uniformly in $x \in [a, 1 - a]$ as $k \rightarrow \infty$. If \mathbb{f} is continuous on $[0, 1]$ then it is uniformly continuous there and the sum at the right-hand side vanishes uniformly in $x \in [0, 1]$. Since $B(x \mid k - 1, \mathbb{f})$ converge to \mathbb{f} uniformly, this completes the proof. □

Appendix C: Posterior sampling and computing

C.1. Gibbs sampling

Posterior sampling for each conditional mean measure, $\Lambda(\cdot) = E\{\tilde{\Pi}(\cdot) \mid T\}$, was conducted according to the following hierarchical structure:

$$k \sim \rho, \quad G \sim \text{DP}(\alpha, G^*), \quad y \mid k, G \sim G, \quad \tilde{x} \mid k, G, y \sim p(\cdot \mid k, y),$$

where $\tilde{x} = \{\tilde{x}_j\}_{j=1}^m$ are the raw warped data, $y = \{y_j\}_{j=1}^m$ are auxiliary latent indicators, and

$$p(x \mid k, y) = \sum_{j=1}^k \beta(x \mid j, k - j + 1) \mathbb{1}_{\{(j-1)/k < y_i \leq j/k\}}.$$

The posterior distribution can be computed using a Gibbs sampler with full conditionals being given by [6, Section 5.5]:

1. For k :

$$\rho(k \mid \tilde{x}, y) \propto \rho(k) \prod_{i=1}^m \beta(\tilde{x}_i \mid z(y_i, k), k - z(y_i, k) + 1),$$

where $z(y, k) = j$ if $(j - 1)/k < y \leq j/k$.

2. For y_i :

$$y_i \mid k, \tilde{x}, y_{-i} \sim \sum_{j \neq i} q_{i,j} \delta_{y_j} + q_{i,0} G_{b,i},$$

where $y_{-i} = (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_m)$ and

$$q_{i,j} \propto \begin{cases} \alpha b(\tilde{x}_i \mid k, G), & j = 0, \\ \beta(\tilde{x}_i \mid z(y_j, k), k - z(y_j, k) + 1), & 1 \leq j \leq k, j \neq i, \end{cases}$$

with

$$dG_{b,i}(y \mid k, Y_i) \propto g(y) \beta(\tilde{x}_i \mid z(y, k), k - z(y, k) + 1).$$

C.2. Parallel computing on the cloud

The simulation studies reported in Section 3 were conducted using a virtual machine instance on the Google Cloud Platform (cloud.google.com) running Linux SO with 8 vCPU and 32 GB RAM. Parallel computing was implemented with the R package *parallel* so to speed up the computations. To give an idea on how long it takes to execute the proposed methods on a regular computer, we note that running a single-run experiment, as in Section 3.1, on a machine with 4 Cores (2.3 Ghz) and 8 Gb RAM takes 338 seconds with a single-core (and 205 seconds when we use parallel computing).

Appendix D: Further numerical experiments

D.1. Supporting outputs

Here we present some figures which are derived from the simulation studies conducted in Section 3. In detail, Figure 8 refers to results in the simulation study in Section 3.1, Figure 9 refers to the comparison conducted in Section 3.2, and Figure 10 corresponds to Figure 5 (left), but for all warp maps.

D.2. Simulation study under misspecification

Here we analyse a simulation scenario similar to that in Section 3.1, but this time using warp maps T_i which not satisfy $E[T_i(t)] = t$; the goal will be to illustrate the performance of the proposed registration method under misspecification. We generate random samples $x_{i,1}, \dots, x_{i,m_i} \mid m_i$, from

$$f(t) = 0.45 \{ \phi(t \mid 0.25, 0.02^2) + \phi(t \mid 0.75, 0.03^2) \} + 0.1 \beta(t \mid 1.5, 1.5),$$

where $m_i \sim \text{Poisson}(L)$, for $i = 1, 2, 3$, with $L = 150$; here, $\phi(t \mid \mu, \sigma^2)$ denotes the Normal density function and $\beta(t \mid a, b)$ denotes the Beta density. The warped data $\tilde{x}_{i,j} = T_i(x_{i,j})$ are obtained using

$$T_i(t) = \int_0^t \beta(y \mid a, b) dy, \quad (i = 1, 2), \quad T_3(t) = 3t - T_1(t) - T_2(t), \quad a, b \stackrel{\text{iid}}{\sim} \text{Unif}[1, 3].$$

Figure 11 shows the estimators of each of the three warp maps through the posterior mean of the induced prior defined in Section 2.2, along with credible bands and the true warp maps over a one shot experiment.

From Figure 11 it can be noticed that—even under misspecification—our estimators are reasonably in line with the true warp functions, and as a consequence, the method recovers quite well the original point processes, as can be seen when comparing the left and right panels of Figure 12. Given (2.7) the width of the credible bands of the warps maps depicted in Figure 11 is determined by the width of the credible bands of the quantile function and of the Fréchet mean. To put differently, the fact that the bands of the warp maps are narrow on some regions is simply a consequence of the fact that the credible bands for the corresponding quantile functions are themselves narrow, on some other region.

A Monte Carlo study was conducted based on $B = 50$ simulated datasets. The WDM (Monte Carlo L^2 -Wasserstein distance mean) defined in Equation (3.2) was at this time 0.0417—which is of the same order of magnitude as the value obtained in Section 3.1 under a well-specified setting. For the sake of comparison, the WDM computed using $\tilde{\Pi}_i$ instead of $\hat{\Pi}_i$, is 0.151, nearly four times larger than that based on $\hat{\Pi}_i$.

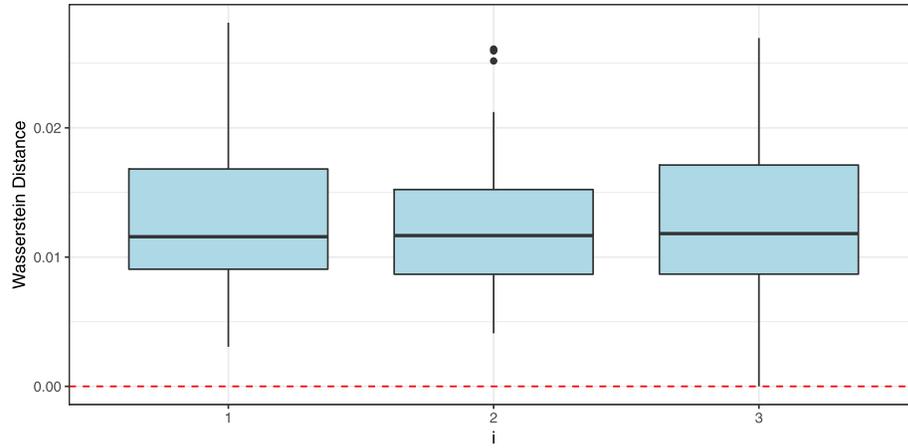


FIG 8. Boxplots of the L^2 -Wasserstein distance between the original processes $\Pi_i^{[b]}$ and the registered ones $\hat{\Pi}_i^{[b]}$ for the simulation study in Section 3.1. Here b ranges from 1 to $B = 50$ and $i = 1, 2, 3$ correspond to the three panels.

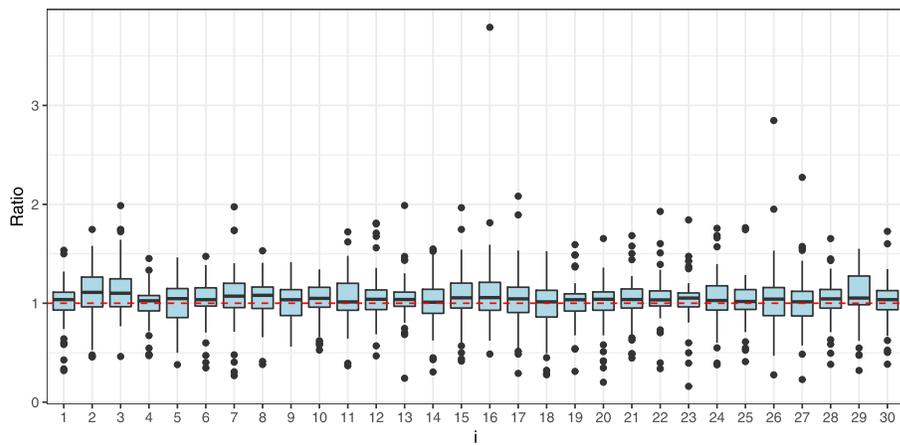


FIG 9. Comparison of proposed Bayesian registration with kernel-based registration of [12] according to numerical experiments from Section 3.2. Each boxplot contains the ratio $d(\hat{\Pi}_i^{[b, Bayes]}, \Pi_i^{[b]}) / d(\hat{\Pi}_i^{[b, Kernel]}, \Pi_i^{[b]})$ for all $i \in \{1, \dots, 30\}$.

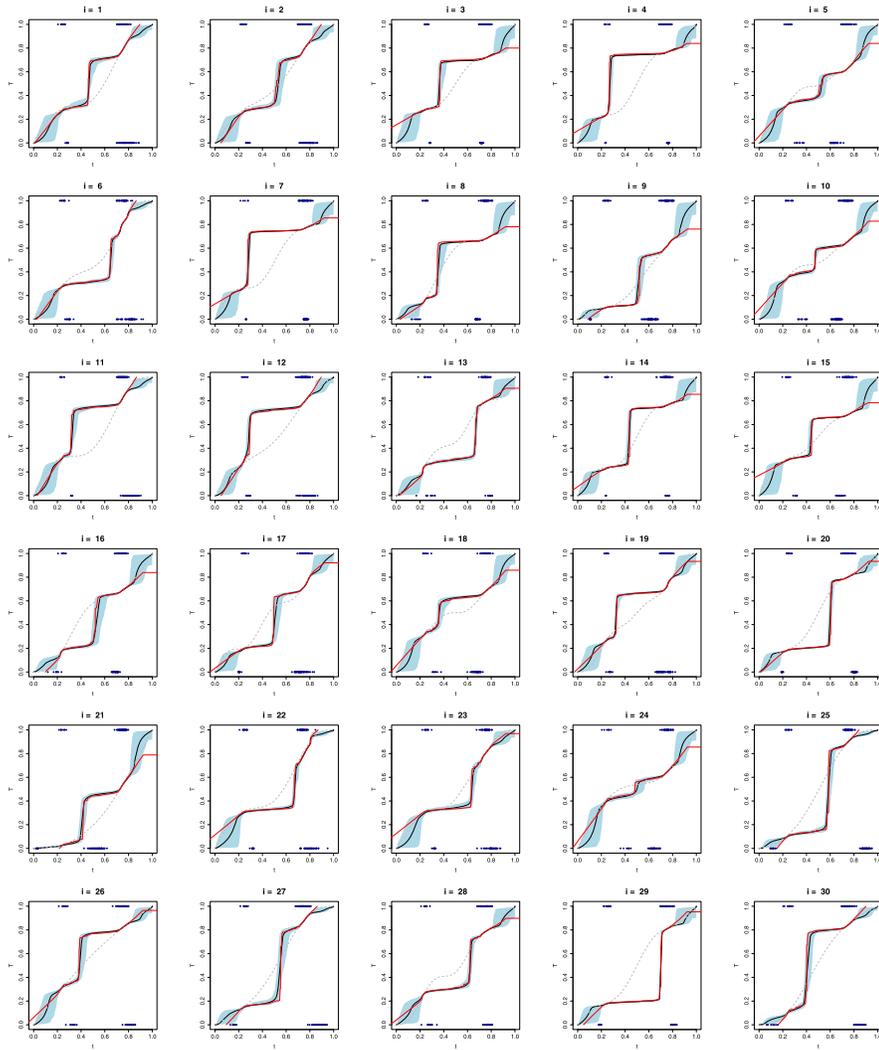


FIG 10. Further fits from simulated data from the setup of Section 3.2 in the large n , small m regime. 30 posterior mean Bernstein polynomial warp functions (solid black) and corresponding credible bands, with their kernel-based counterparts (solid red) and the original warp functions (dashed grey). Warped and original data are in the bottom and top, respectively.

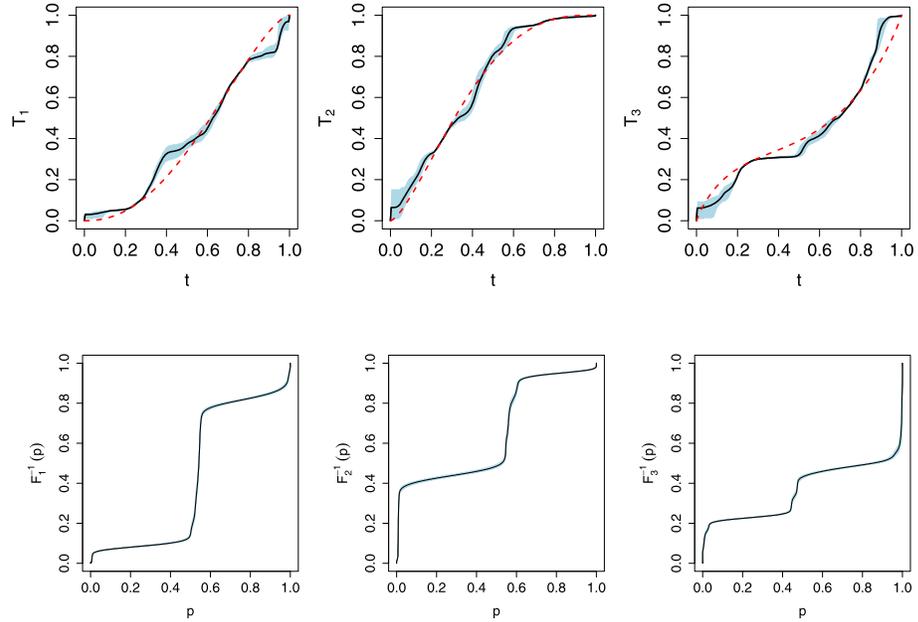


FIG 11. Fits from simulation study under misspecification. Top: True (dashed red) and estimated (solid black) warp functions along with credible bands. The estimators are constructed as the posterior mean of the induced prior. Bottom: Corresponding quantile function estimates along with credible bands.

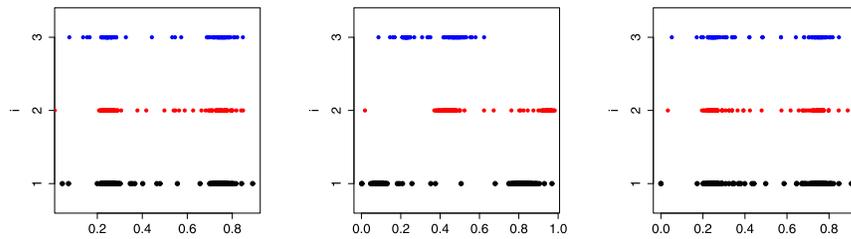


FIG 12. Further fits from simulation study under misspecification. Left: Simulated instances of the original point process. Middle: Their corresponding phase-varying point process. Right: Their corresponding registered versions.

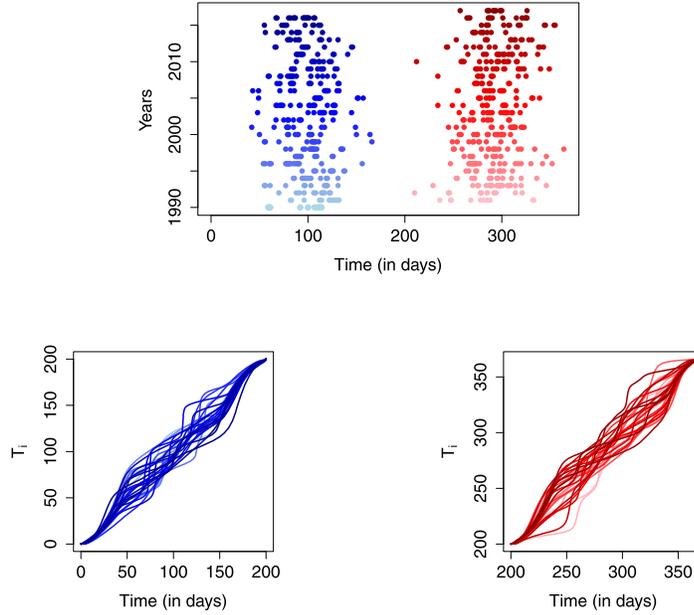


FIG 13. Top: Point processes of annual peaks for peaks above (red) and below (blue) temperature thresholds (97.5% and 2.5% quantiles). Bottom: Corresponding posterior Bernstein polynomial mean warp functions in the same palette of colours.

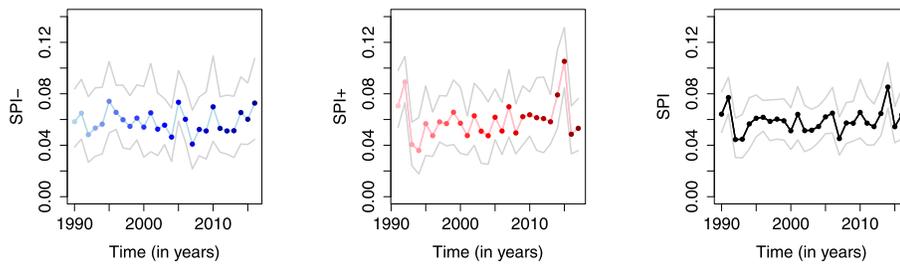


FIG 14. Posterior mean SPI (scores of peak irregularity, see (4.1)), along with credible intervals, below a cold temperature threshold (2.5% quantile; left), above a hot temperature threshold (97.5% quantile; middle), and global (right).

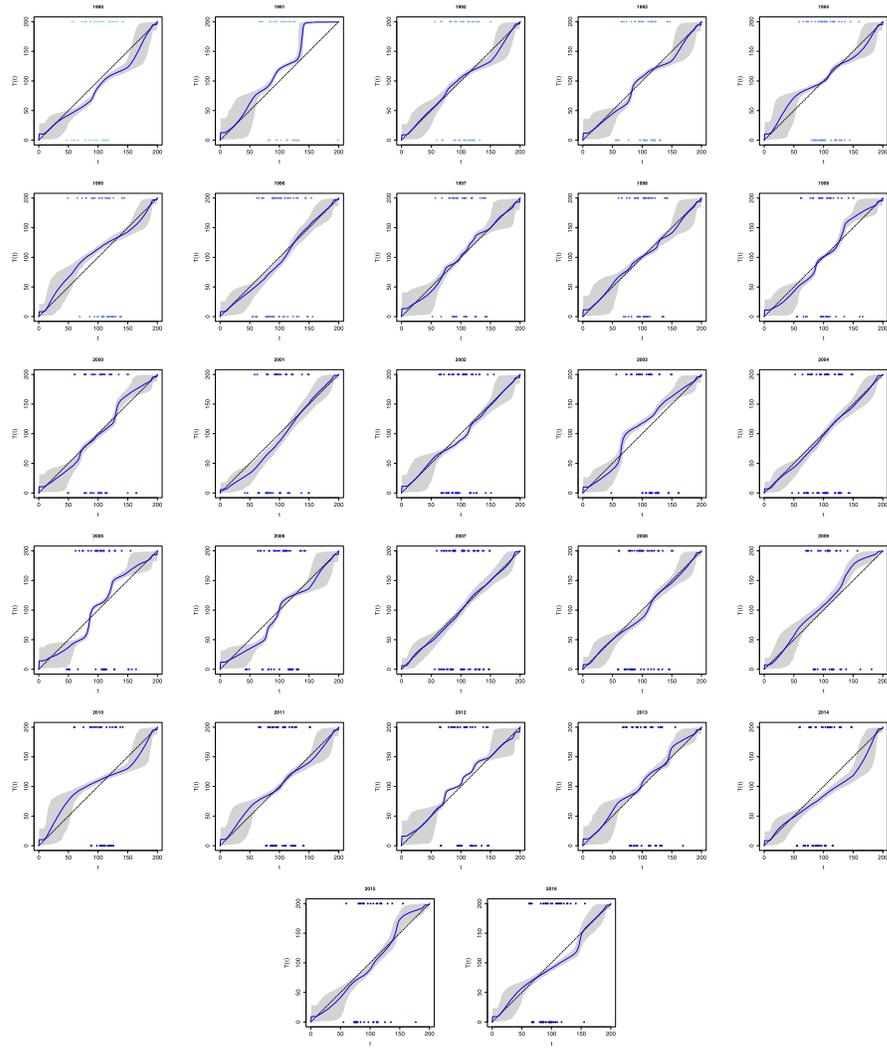


FIG 15. Yearly posterior mean Bernstein polynomial warp functions of low-temperatures (solid blue), along with credible bands, plotted with raw data (bottom), registered points (top) and the identity function (dashed black).

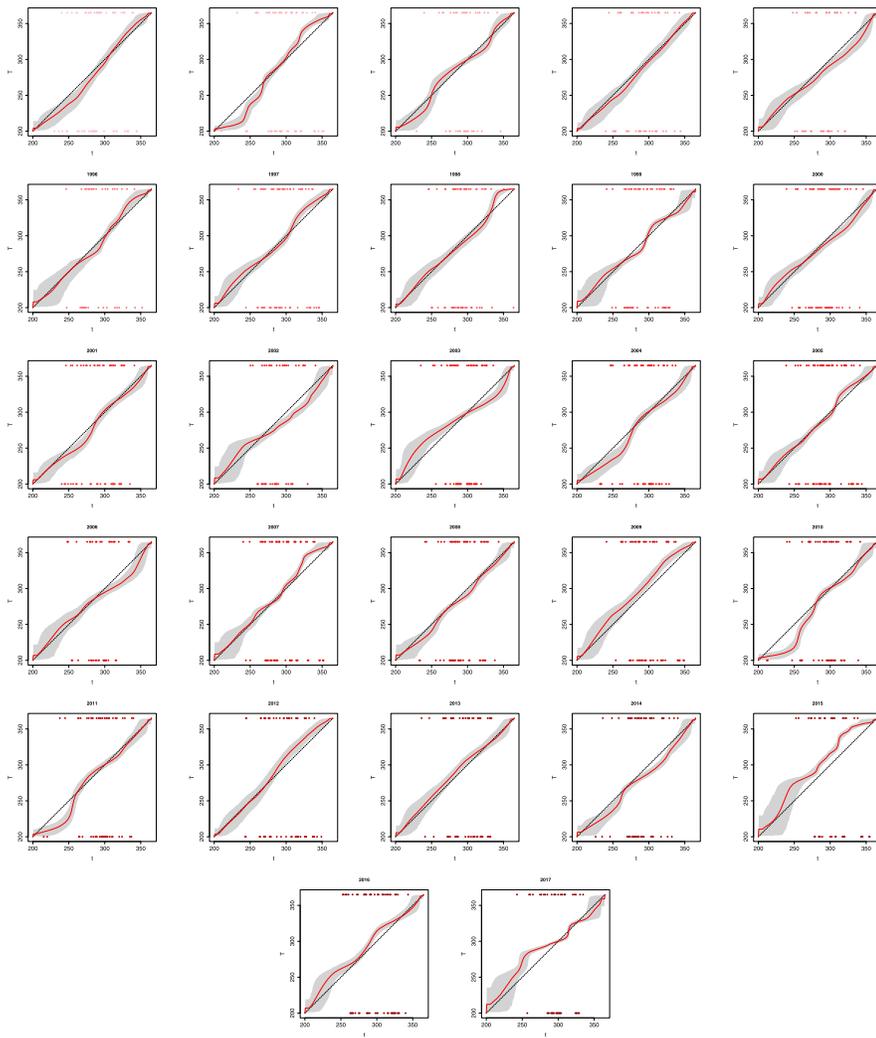


FIG 16. Yearly posterior mean Bernstein polynomial warp functions of high-temperatures (solid red), along with credible bands, plotted with raw data (bottom), registered points (top), and the identity function (dashed black). Here the year refers to that of onset of summer.

Appendix E: Additional outputs from application

As in Section 4, we analyse the annual peaks over threshold, $\{\tilde{x}_{i,j}^+ \geq u_j^+\}$, and annual peaks below threshold, $\{\tilde{x}_{i,j}^- \leq u_j^-\}$; we set the thresholds u_j^+ and u_j^- using the 97.5% and 2.5% quantiles of temperature over year j , and this results in m_1^+, \dots, m_n^+ ranging from 10 to 18 and m_1^-, \dots, m_n^- ranging from 10 to 20.

Some comments on the results reported above are in order. Figures 13 and 14 are the analogues of Figures 6 and 7, respectively. Figures 15 and 16 depict the fits of the warp maps from the data application in Section 4.

Acknowledgments

We are grateful to the associate editor and two referees for constructive feedback that led to substantial improvement of the paper.

References

- [1] Agueh, M. & Carlier, G. (2011). Barycenters in the Wasserstein space. *Soc. Ind. Appl. Math.* **43**, 904–924. [MR2801182](#)
- [2] Barrientos, A. F., Jara, A. & Quintana, F. A. (2017). Fully nonparametric regression for bounded data using dependent Bernstein polynomials. *J. Am. Statist. Assoc.* **112**, 806–825. [MR3671772](#)
- [3] Boissard, E., Le Gouic, T. & Loubes, J.-M. (2015). Distribution’s template estimate with Wasserstein metrics. *Bernoulli* **21**, 740–759. [MR3338645](#)
- [4] Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Stat.* **1**, 209–230. [MR0350949](#)
- [5] Ghosal, S., Ghosh, J. K. & Ramamoorthi, R. (1999). Consistent semiparametric Bayesian inference about a location parameter. *J. Statist. Plann. Inf.* **77**, 181–193. [MR1687955](#)
- [6] Ghosal, S. & van der Vaart, A. W. (2017). *Fundamentals of Nonparametric Bayesian Inference*. Cambridge University Press, Cambridge. [MR3587782](#)
- [7] Karr, A. F. (1991). *Point Processes and Their Statistical Inference*. Probability: Pure and Applied. New York: Dekker, 2nd ed. [MR1113698](#)
- [8] Koenker, R. (2005). *Quantile Regression*. Cambridge, MA: Cambridge University Press. [MR2268657](#)
- [9] Lehmann, E. L. & Romano, J. P. (2006). *Testing Statistical Hypotheses*. New York: Springer. [MR2135927](#)
- [10] Marron, J. S., Ramsay, J. O., Sangalli, L. M. & Srivastava, A. (2015). Functional data analysis of amplitude and phase variation. *Statist. Science* **30**, 468–484. [MR3432837](#)
- [11] Menzel, A. & Fabian, P. (1999). Growing season extended in Europe. *Nature* **397**, 659.
- [12] Panaretos, V. M. & Zemel, Y. (2016). Amplitude and phase variation of point processes. *Ann. Stat.* **44**, 771–812. [MR3476617](#)

- [13] Panaretos, V. M. & Zemel, Y. (2019). Statistical aspects of Wasserstein distances. *Annu. Rev. Stat. Appl.* **6**, 405–431. [MR3939527](#)
- [14] Petrone, S. (1999). Bayesian density estimation using Bernstein polynomials. *Canadian J. Statist.* **27**, 105–126. [MR1703623](#)
- [15] Petrone, S. (1999). Random Bernstein polynomials. *Scandinavian J. Statist.* **26**, 373–393. [MR1712051](#)
- [16] Petrone, S. & Wasserman, L. (2002). Consistency of Bernstein polynomial posteriors. *J. Royal Statist. Soc., Ser. B* **64**, 79–100. [MR1881846](#)
- [17] Peyré, G. & Cuturi, M. (2019). Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning* **11**, 355–607.
- [18] Ramsay, J. O. & Silverman, B. W. (2002). *Applied Functional Data Analysis: Methods and Case Studies*, New York: Springer. [MR1910407](#)
- [19] Ramsay, J. O. & Silverman, B. W. (2005). *Functional Data Analysis*. New York: Springer, 2nd ed. [MR2168993](#)
- [20] Santambrogio, F. (2015). *Optimal Transport for Applied Mathematicians*. Basel: Birkhäuser. [MR3409718](#)
- [21] Schwartz, M. D., Ahas, R. & Aasa, A. (2006). Onset of spring starting earlier across the Northern Hemisphere. *Global Change Bio.* **12**, 343–351.
- [22] Tang, R. & Müller, H.-G. (2008). Pairwise curve synchronization for functional data. *Biometrika* **95**, 875–889. [MR2461217](#)
- [23] Wang, J.-L., Chiou, J.-M. & Müller, H.-G. (2016). Functional data analysis. *Annu. Rev. Stat. Appl.* **3**, 257–295.
- [24] Wu, S., Müller, H. & Zhang, Z. (2013). Functional data analysis for point processes with rare events. *Statistica Sinica* **23**, 1–23. [MR3076156](#)
- [25] Wu, W. & Srivastava, A. (2014). Analysis of spike train data: Alignment and comparisons using the extended Fisher–Rao metric. *Electron. J. Stat.* **8**, 1776–1785. [MR3273594](#)
- [26] Zemel, Y. & Panaretos, V. M. (2019). Fréchet Means and Procrustes Analysis in Wasserstein Space. *Bernoulli* **25**, 932–976. [MR3920362](#)
- [27] Zheng, Y., Zhu, J. & Roy, A. (2009). Nonparametric Bayesian inference for the spectral density function of a random field. *Biometrika* **97**, 238–245. [MR2594432](#)