A spectral mean for random closed curves

M.N.M. van Lieshout [*]

CWI, P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands
University of Twente, P.O. Box 217, NL-7500 AE Enschede, The Netherlands

ARTICLE INFO

ABSTRACT

We propose a spectral mean for closed sets described by sample points on their boundaries subject to mis-alignment and noise. We derive maximum likelihood estimators for the model and noise parameters in the Fourier domain. We estimate the unknown mean boundary curve by back-transformation and derive the distribution of the integrated squared error. Mis-alignment is dealt with by means of a shifted parametric diffeomorphism. The method is illustrated on simulated data and applied to photographs of Lake Tana taken by astronauts during a Shuttle mission.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Many geographical or biological objects are observed in image form. The boundaries of such objects are seldom crisp, for example due to measurement error and discretisation, or because the boundaries themselves are intrinsically indeterminate (Burrough and Frank, 1996). Moreover, the objects may not be static in the sense that if multiple images are taken, the objects may have been deformed.

One attempt to model natural objects under uncertainty is fuzzy set theory (Zimmermann, 2001). However, the underlying axioms are too poor to handle topological properties of the shapes to be modelled and cannot deal with correlation. Similarly, the belief functions that lie at the heart of the Dempster–Shafer theory (Dempster, 1967; Shafer, 1976) do not necessarily correspond to the containment function of a well-defined random closed set (Molchanov, 2005).

Proper stochastic geometric models for natural objects are scarce. Indeed, the Boolean model (Molchanov, 1997), defined as the union of independent copies of a simple geometric object, to this date dominates the random set literature. It owes its popularity to the fact that it is analytically

* Correspondence to: University of Twente, P.O. Box 217, NL-7500 AE Enschede, The Netherlands. Tel.: +31 20 5924008.
E-mail address: M.N.M.van.Lieshout@cwi.nl.

tractable, but its realisations do not match the irregular shapes observed in natural phenomena. The modification in which objects may be subject to transformation is known as a deformable template model (Grenander and Miller, 2007).

Apart from the dearth of models, another complication is that, since the family of closed sets in the plane is not linear, it is not straightforward to define what is meant by a mean set. Indeed, many different suggestions have been made, often with a specific context in mind. For instance, the radius vector mean is taylor-made for star-shaped objects, the Vorob'ev expectation for fuzzy sets. This context dependence has obvious disadvantages. More specifically, the radius vector mean does not apply to sets that are not star shaped, the Vorob'ev expectation is a function of the coverage function only so that it cannot take into account spatial coherence and fine detail. The Fréchet expectation (Fréchet, 1948) on the other hand can be defined quite generally as a distance minimiser in a metric space and – in contrast to the previously mentioned definitions – is extendable to higher orders, but the computational burden may be large. For further details and examples, see Molchanov (2005).

In this paper, we extend an interesting approach suggested by Jónsdóttir and Jensen (2005) for star-shaped objects deformed by noise to objects that are not necessarily star-shaped by modelling their boundary as a closed curve. Additionally, we propose a spectral mean for such random sets and carry out inference in the Fourier domain. The approach is easy to implement and, moreover, the integrated squared error can be computed in closed form.

A complication is that the curves need to be well-aligned in the sense that their parametrisations must be in correspondence. Similar problems arise in shape registration (Dryden and Mardia, 1998), deformable template matching (Grenander and Miller, 2007), signal extraction (Bigot and Gendre, 2013) or image warping (Glasbey and Mardia, 2001). Although the terminology varies, the objective is always the same, namely to find a transformation between objects so that they resemble one another. The context determines the class of transformations and the criterion for evaluating the resemblance. Examples of the former include rigid motions or time shifts; resemblance can be based on Euclidean distance between landmarks (Besl and McKay, 1992), Hausdorff or image distances (Baddeley and Molchanov, 1998), or on intrinsic properties of the object (e.g. arc length and curvature) (Sebastian et al., 2003) or the group of transformations (Kendall, 1984). Further discussion and examples can be found in Davies et al. (2008).

Since we work on a space of smooth cyclic functions, we shall use the (expected) L_2 distance on this function space to measure the quality of alignment and we follow Younes (2010) in reparametrising a curve by diffeomorphisms. Note that, in contrast to function estimation in one dimension, a reparametrisation does not affect the appearance of the curve.

This paper is organised as follows. We begin by recalling basic facts about planar curves, cyclic Gaussian random processes and spectral analysis. Then we formulate a model for sampling noisy curves, carry out inference in the Fourier domain and quantify the error. We present a simulation study to assess performance. Finally, we describe how to deal with mis-alignment between curves and illustrate the approach on images of a lake in Ethiopia.

2. Noisy curves

In this section we recall basic facts about planar curves, Fourier bases and cyclic Gaussian random processes.

2.1. Planar curves

Throughout this paper, we model the boundary of a random object of interest by a smooth closed curve.

Consider the class of functions $\Gamma : I \to \mathbb{R}^2$ from some interval I to the plane. Define an equivalence relation $\sim$ on the function class as follows: Two functions Γ and Γ' are equivalent, $\Gamma \sim \Gamma'$, if there exists a strictly increasing function φ from I onto another interval I' such that $\Gamma = \Gamma' \circ \varphi$. Note that φ is a homeomorphism. The relation defines a family of equivalence classes, each of which is called a curve. Its member functions are called parametrisations. Since the images of two parametrisations of

the same curve are identical, we shall, with slight abuse of notation, use the symbol Γ for a specific parametrisation, for a curve and for its image.

A curve is said to be continuous if it has a continuous parametrisation $\Gamma : [p, q] \to \mathbb{R}^2$, in which case all parametrisations are continuous; it is closed (or cyclic) if $\Gamma(p) = \Gamma(q)$. Since closed curves have neither a 'beginning' nor an 'end', a rooted parametrisation is provided by a point on the curve together with a cyclic parametrisation from that point in a given orientation. For convenience, we shall often rescale the definition interval to $[-\pi, \pi]$.

In an optimisation context, it is natural to assume a curve to be parametrised by some function Γ that is continuously differentiable (C^1) and the same degree of smoothness to hold for the functions φ that define the equivalence relation between parametrisations. In effect, φ should be a diffeomorphism. See Younes (2010, Chapter 1) for further details.

2.2. Fourier representation

Let $\Gamma : [-\pi, \pi] \to \mathbb{R}^2$ be a C^1 function with $\Gamma(-\pi) = \Gamma(\pi)$. Since the family $\{\cos(j\theta), \sin(j\theta) : j \in \mathbb{N}_0, \theta \in [-\pi, \pi]\}$ forms an orthogonal basis for $L_2([-\pi, \pi])$, the space of all square integrable functions on $[-\pi, \pi]$, Γ can be approximated by a trigonometric polynomial of the form

$$\sum_{j=0}^{J} \left[\mu_j \cos(j\theta) + \nu_j \sin(j\theta)\right].$$

The vectors μ_j and ν_j are called the Fourier coefficients of order j and satisfy

$$\begin{cases} \mu_0 = \dfrac{1}{2\pi} \displaystyle\int_{-\pi}^{\pi} \Gamma(\theta)d\theta \\[2mm] \mu_j = \dfrac{1}{\pi} \displaystyle\int_{-\pi}^{\pi} \Gamma(\theta) \cos(j\theta)d\theta \\[2mm] \nu_j = \dfrac{1}{\pi} \displaystyle\int_{-\pi}^{\pi} \Gamma(\theta) \sin(j\theta)d\theta \end{cases}$$

for $j \in \mathbb{N}$. Moreover, Parseval's identity

$$\frac{1}{\pi} \int_{-\pi}^{\pi} \|\Gamma(\theta)\|^2 d\theta = 2\|\mu_0\|^2 + \sum_{j=1}^{\infty} \left[\|\mu_j\|^2 + \|\nu_j\|^2\right] \tag{1}$$

relates the squared L_2 norm of Γ to a sum of squared Fourier coefficients.

2.3. Stationary cyclic Gaussian random processes

As in Jónsdóttir and Jensen (2005, Proposition 2.1), let $N = (N_1, N_2)$ be a stationary cyclic Gaussian random process on $[-\pi, \pi]$ with values in $\mathbb{R}^2$ of the form

$$N(\theta) = \sum_{j=0}^{\infty} \left[A_j \cos(j\theta) + B_j \sin(j\theta)\right], \tag{2}$$

where the components of A_j and B_j are mutually independent zero-mean Gaussian random variables with variances σ_j^2 that are small enough for the series $\sum_j \sigma_j^2$ to converge. Then N has independent components with zero mean and covariance function

$$\rho(\theta) = \sum_{j=0}^{\infty} \sigma_j^2 \cos(j\theta).$$

For the existence of a continuous version, further conditions are needed. Indeed, Theorem 25.10 in Rogers and Williams (1994) implies that if

$$\sum_{j=1}^{\infty} j^{2k+\epsilon} \sigma_j^2 < \infty \qquad (3)$$

for $k \in \mathbb{N} \cup \{0\}$, $\epsilon > 0$, there exists a version of N that is k times continuously differentiable. From now on, we shall always assume (3) for $k = 1$.

A convenient model is the generalised p-order model of Hobolth et al. (2003), see also Aletti and Ruffini (in press) and Jónsdóttir and Jensen (2005), in which $\sigma_j^{-2} = \alpha + \beta j^{2p}$, $j \geq 2$, for parameters $\alpha, \beta > 0$. The parameter p determines the smoothness. By (3), a continuous version exists for all $p > 1/2$, a C^1 version for $p > 3/2$.

3. Data model

In this paper, we are interested in samples of objects with uncertain boundaries. Such data arise in, for example, medical imaging where delineation of organs or tumours is required. Another application domain is that of the earth sciences (e.g. the paper by Zhao et al. in this volume). Indeed, objects such as arctic glaciers or fresh water lakes are rarely crisp. Moreover, they tend to shrink and grow according to the seasons and climate. To analyse such changes, the object of interest must be mapped and the uncertainty in its outline quantified.

More formally, suppose the data consist of multiple observations X_t, $t = 0, \ldots, T$, of an object of interest in discretised form as a list of finitely many points $(X_t^l)_{l=1,\ldots,n}$ on its boundary, either explicitly (cf. Fig. 1) or implicitly in the form of an image as in Fig. 4. In other words, the lists trace some unknown closed curve Γ affected by noise.

As for the noise N_t, in the absence of systematic errors, we assume that $\mathbb{E} N_t(\theta) = 0$ for all $\theta \in [-\pi, \pi]$ and that the correlation between errors $N_t(\theta)$ and $N_t(\eta)$ depends only on the absolute difference $|\theta - \eta|$. Thus, we model the noise by independent mean-zero stationary cyclic Gaussian random processes of the form (2) on $[-\pi, \pi]$. Finally, the roots and parametrisations of the noisy curves must be well-aligned. This is taken care of by shift parameters $\alpha_t \in [-\pi, \pi]$ for the root and diffeomorphisms $\varphi_t : [-\pi, \pi] \to [-\pi, \pi]$ for the reparametrisation.

To summarise, we arrive at the following model for points X_t^l, $l = 1, \ldots, n$, sampled along the boundaries of observation $t \in \{0, \ldots, T\}$ of curve Γ.

Model 1. Let $\Gamma : [-\pi, \pi] \to \mathbb{R}^2$ be a C^1 function with $\Gamma(-\pi) = \Gamma(\pi)$. Let N_t be independent stationary cyclic Gaussian random processes on $[-\pi, \pi]$ of the form (2) with variances σ_j^2 for which (3) holds with $k = 1$. Then, for $\alpha_t \in [-\pi, \pi]$ and diffeomorphisms $\varphi_t : [-\pi, \pi] \to [-\pi, \pi]$, $\theta_l = -(n+1)\pi/n + 2\pi l/n$, $l = 1, \ldots, n$, and $t = 0, \ldots, T$, set

$$X_t^l = X_t(\theta_l) = \Gamma(\varphi_t(\theta_l - \alpha_t)) + N_t(\varphi_t(\theta_l - \alpha_t)),$$

interpreted cyclically modulo 2π.

Note that the independence of the N_t implies that the random vectors X_t of observations are independent. For fixed t, though, the X_t^l may well be correlated.

We set ourselves the goal of estimating Γ and the noise variance parameters σ_j^2. For the moment, assume that all $\alpha_t \equiv 0$ and that each φ_t is the identity operator. (We shall return to the issue of estimating the parametrisations later). Then Model 1 reduces to

$$X_t(\theta) = \Gamma(\theta) + N_t(\theta), \qquad (4)$$

which is observed at θ_l to give $X_t^l = X_t(\theta_l)$.

4. Parameter estimation

The objective of this section is to estimate the unknown curve Γ in model (4), as well as the parameters of the covariance function of the N_t based on the observations X_t^l. We propose to do this in the Fourier domain. Since the noise is defined by its random Fourier coefficients (cf. (2)), the estimation problem then reduces to inference for a multivariate normal distribution. A simple back-transformation to the spatial domain completes the procedure.

To be precise, let, for $j \in \mathbb{N}$ and $\theta_l = -(n+1)\pi/n + 2\pi l/n$, $l = 1, \ldots, n$,

$$\begin{cases} F_{0,n}^t = \dfrac{1}{n} \sum_{l=1}^{n} X_t^l = \dfrac{1}{n} \sum_{l=1}^{n} [\Gamma(\theta_l) + N_t(\theta_l)] \\[2mm] F_{j,n}^t = \dfrac{2}{n} \sum_{l=1}^{n} X_t^l \cos(j\theta_l) = \dfrac{2}{n} \sum_{l=1}^{n} [\Gamma(\theta_l) \cos(j\theta_l) + N_t(\theta_l)\cos(j\theta_l)] \\[2mm] G_{j,n}^t = \dfrac{2}{n} \sum_{l=1}^{n} X_t^l \sin(j\theta_l) = \dfrac{2}{n} \sum_{l=1}^{n} [\Gamma(\theta_l) \sin(j\theta_l) + N_t(\theta_l)\sin(j\theta_l)] \end{cases} \tag{5}$$

be the Riemann approximations of the random Fourier coefficients of X_t, $t = 0, \ldots, T$. We shall write $\mu_{j,n}$, respectively $\nu_{j,n}$ for the deterministic parts of $F_{j,n}^t$ and $G_{j,n}^t$ in (5). For example, for $j = 0$, $\mu_{j,n} = \sum_l \Gamma(\theta_l)/n$. Furthermore, let

$$\sigma_{j,n}^2 = \frac{2}{n} \sum_{l=1}^{n} \rho(\theta_l) \cos(j\theta_l)$$

for $j \geq 1$ and $\sigma_{0,n}^2 = \sum_l \rho(\theta_l)/n$ be the Riemann approximations of σ_j^2.

From now on, assume that $n \geq 3$ is odd. In this case, the sequence θ_l contains 0 and is symmetric around zero. We estimate $\mu_{j,n}$ and $\nu_{j,n}$ by their empirical counterparts

$$\begin{cases} \hat{\mu}_{j,n} = \dfrac{1}{T+1} \sum_{t=0}^{T} F_{j,n}^t, & j \in \mathbb{N} \cup \{0\} \\[2mm] \hat{\nu}_{j,n} = \dfrac{1}{T+1} \sum_{t=0}^{T} G_{j,n}^t, & j \in \mathbb{N} \end{cases}$$

and transform back to the spatial domain to obtain

$$\hat{\Gamma}_n(\theta) = \hat{\mu}_{0,n} + \sum_{j=1}^{J} \left[\hat{\mu}_{j,n} \cos(j\theta) + \hat{\nu}_{j,n} \sin(j\theta) \right] = \frac{1}{T+1} \sum_{t=0}^{T} \sum_{l=1}^{n} X_t^l S_l(\theta), \tag{6}$$

where we use the notation $S_l(\theta) = 1/n + 2\sum_{j=1}^{J} \cos(j(\theta - \theta_l))/n$. In words, $\hat{\Gamma}_n(\theta)$ is a weighted average of the observed points X_t^l with the weight determined by the difference between θ and θ_l.

Finally, the variances σ_j^2 are estimated by

$$\hat{\sigma}_{j,n}^2 = \frac{1}{4(T+1)} \sum_{t=0}^{T} \left[\|F_{j,n}^t - \hat{\mu}_{j,n}\|^2 + \|G_{j,n}^t - \hat{\nu}_{j,n}\|^2 \right] \tag{7}$$

for $j \in \mathbb{N}$ and

$$\hat{\sigma}_{0,n}^2 = \frac{1}{2(T+1)} \sum_{t=0}^{T} \|F_{0,n}^t - \hat{\mu}_{0,n}\|^2$$

for $j = 0$.

Note that when $J < n/2$, the number of Fourier coefficients to estimate does not exceed the number of observed boundary points.

4.1. Distributional results

In this section, we derive the distribution of all model parameters as well as that of the integrated squared error of the estimated curve. Our first result considers the estimators for the Fourier coefficients – hence for the unknown curve – and noise parameters. Its proof is relegated to the Appendix.

Theorem 1. *Let $J < n/2$ and suppose that $n \geq 3$ is odd. Consider model (4). Then the following hold.*

1. *$\hat{\mu}_{j,n}$, $\hat{v}_{j,n}$ and $\hat{\sigma}_{j,n}^2$ are maximum likelihood estimators for, respectively, $\mu_{j,n}$, $v_{j,n}$ and $\sigma_{j,n}^2$.*
2. *$\hat{\mu}_{j,n}$ and $\hat{v}_{j,n}$ are mutually independent. They are normally distributed with independent components, mean vectors $\mu_{j,n}$ and $v_{j,n}$, respectively, and component variances $\sigma_{j,n}^2/(T+1)$.*
3. *For $j \in \mathbb{N}$, $4(T+1)\hat{\sigma}_{j,n}^2/\sigma_{j,n}^2$ is χ^2 distributed with $4T$ degrees of freedom, whereas $2(T+1)\hat{\sigma}_{0,n}^2/\sigma_{0,n}^2$ is χ^2 distributed with $2T$ degrees of freedom.*
4. *$\hat{\Gamma}_n(\cdot)$ is a cyclic Gaussian random process with independent components and mean vector*

$$\mathbb{E}\hat{\Gamma}_n(\theta) = \sum_{l=1}^{n} \Gamma(\theta_l)S_l(\theta) = \mu_{0,n} + \sum_{j=1}^{J} \left[\mu_{j,n}\cos(j\theta) + v_{j,n}\sin(j\theta)\right];$$

the covariance function $\sum_{j=0}^{J} \sigma_{j,n}^2 \cos(j\theta)/(T+1)$ of both components is translation invariant.

Next, we consider the distribution of the integrated squared error of the curve estimator. Again, the proof can be found in the Appendix.

Theorem 2. *Consider the estimator (6) for model (4) and assume $J < n/2$ for odd $n \geq 3$. Then*

$$\frac{1}{\pi} \int_{-\pi}^{\pi} \|\hat{\Gamma}_n(\theta) - \Gamma(\theta)\|^2 d\theta = \sum_{j=J+1}^{\infty} \left[\|\mu_j\|^2 + \|v_j\|^2\right] + \tilde{Z},$$

where

$$\tilde{Z} = \frac{1}{T+1}\left(2\sigma_{0,n}^2 Z_{0,n} + \sum_{j=1}^{J} \sigma_{j,n}^2 Z_{j,n}\right)$$

and the $Z_{j,n}$ are independent χ^2 distributed random variables with four degrees of freedom for $j \geq 1$, two for $j = 0$, and non-centrality parameters $(T+1)c_{j,n}/\sigma_{j,n}^2$ with

$$c_{j,n} = \|\mu_{j,n} - \mu_j\|^2 + \|v_{j,n} - v_j\|^2$$

for $j = 1, \ldots, J$ and $c_{0,n} = \|\mu_{0,n} - \mu_0\|^2$ for $j = 0$.

Under the assumptions of Theorem 2, the expected integrated squared error of $\hat{\Gamma}_n$ reads

$$\sum_{j=J+1}^{\infty} \left[\|\mu_j\|^2 + \|v_j\|^2\right] + \frac{4}{T+1} \sum_{j=0}^{J} \sigma_{j,n}^2 + 2c_{0,n} + \sum_{j=1}^{J} c_{j,n}. \tag{8}$$

The first term in (8) is the bias caused by taking into account only a finite number of Fourier coefficients. The second term corresponds to the variance, and the last two terms describe the discretisation error in the Fourier coefficients. Indeed, if there would be no discretisation ($n = \infty$), the non-centrality parameters would be zero and $\sigma_{j,n}^2 = \sigma_j^2$.

To conclude this paragraph, note that the smoothness assumptions on Γ and the covariance function ρ imply that both $\sum_j(\sigma_{j,n}^2 - \sigma_j^2)$ and $\sum_j c_{j,n}$ are of the order J^3/n^2. Better rates might be obtained if J were adapted to n and T.

Theorem 3. *Under the assumptions of Theorem 2, (8) converges to zero as $(T, J, n) \to \infty$ in such a way that J is of the order $n^{2/3-\epsilon}$ for some $\epsilon > 0$.*

The proof is given in the Appendix.

4.2. Simulated example

The left-hand panel in Fig. 1 shows a hundred contours consisting of points sampled at $\theta_l = -(n+1)\pi/n + 2\pi l/n, l = 1, \ldots, n$ for $n = 129$, along a nested quintic curve (Keren, 2004) degraded by noise according to the generalised p-order model Hobolth et al. (2003) with $p = 2, \alpha = 1.0$ and $\beta = 10.0$, truncated at ten Fourier coefficients. Note that the sample paths are almost surely continuously differentiable.
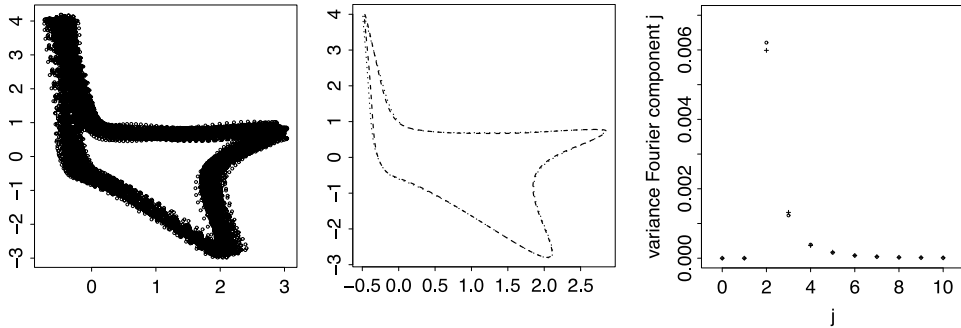
Fig. 1. Left-most panel: Data points sampled along 100 realisations of a nested quintic curve degraded by noise. Middle panel: Estimated curves based on data points sampled along 100 (dashed line) and 10 (stippled line) curves. Right-most panel: Estimated variance $\hat{\sigma}_{j,n}^2$ of the Fourier coefficients plotted against j (crosses) compared to their true value σ_j^2 (circles).

We use Eq. (6) to estimate the true curve Γ, which cannot be distinguished by eye from Γ. The effect of reducing the number of data images is small as shown in the middle panel of Fig. 1. We also estimate the variances σ_j^2 for $j = 0, \ldots, 10$ according to (7). These are shown as crosses in the right-most panel of Fig. 1. For comparison, the true values are plotted too (the circles in the right-most panel of Fig. 1). Note that the estimates are most variable for $j = 2$ and $j = 3$.

Recall that the result depends on the choices of J and n. To investigate the effect of the choice of J, we apply (6) with $J \in \{1, \ldots, 128\}$, keeping the values of n and T as they were. The fit is excellent for J in the range of about 7 up to 119. In Fig. 2, we show the estimated curves for $J = 3, 5$ and 121. For $J = 3$ and 5 there is over-smoothing with the larger bias for $J = 3$; for $J = 121$, over-fitting occurs.
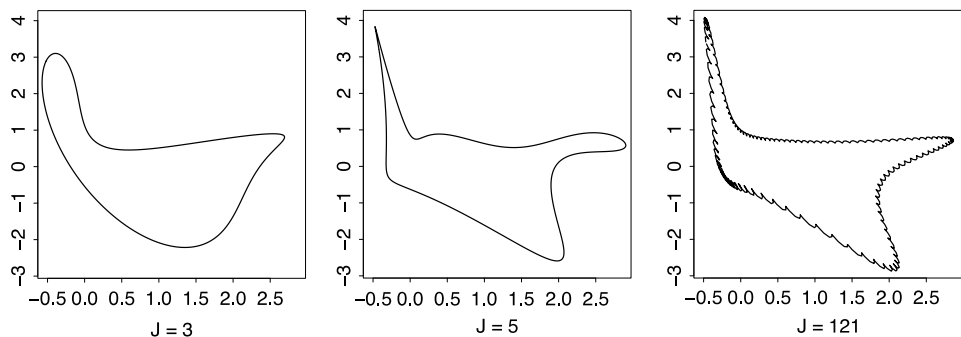
Fig. 2. Estimated curves for $J = 3$ (left-most panel), $J = 5$ (middle panel) and $J = 121$ (right-most panel) for the data displayed in the left-most panel of Fig. 1.

Next, we consider the effect of discretisation by sub-sampling along each of the data curves. In the two left-most panels of Fig. 3, the result for sampling every fourth and sixth point is displayed. It can be seen that the quality of the estimate is still very satisfying, even though 33 respectively 21 rather than 129 points are used in the Riemann approximations (5). This is no longer the case when
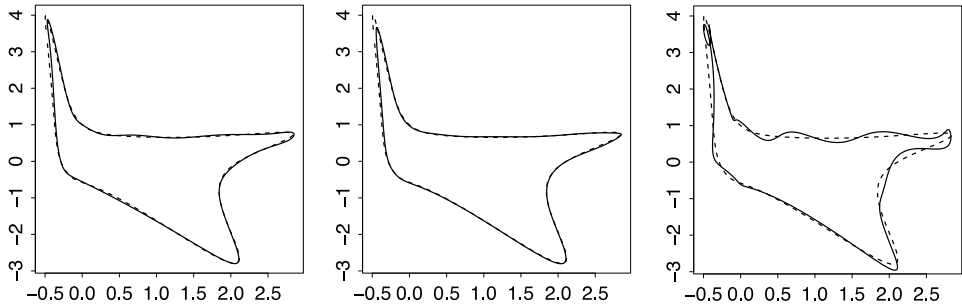
Fig. 3. Estimated (solid line) and true curve (dashed line) for the data displayed in Fig. 1 sub-sampled to a fraction four (left hand panel), six (middle panel) and eight (right hand panel).

we sub-sample further down to only 17 points along each curve, as can be seen from the right hand panel of Fig. 3. In the latter case, however, the number of Fourier coefficients to estimate ($J = 10$) is larger than $n/2 = 8.5$.

5. Parametrisation

As we saw previously, given a root, any parametrisation Γ of a closed C^1 curve can be written as a composition $\Gamma' \circ \varphi$ of a fixed parametrisation Γ' with a diffeomorphism φ. Thus, given two curves parametrised by, say, Γ and Γ_1, alignment of Γ_1 to Γ amounts to finding a shift α and a diffeomorphism φ such that $\Gamma_1(\theta) \approx \Gamma(\varphi(\theta - \alpha))$ interpreted cyclically. Without loss of generality, we consider diffeomorphisms φ from $[-\pi, \pi]$ onto itself.

Parametric diffeomorphisms can be constructed as the flow of differential equations (Younes, 2010, Chapter 8). In our context, it is convenient to consider the differential equation

$$x'(t) = f_w(x(t)), \quad t \in \mathbb{R}, \tag{9}$$

with initial condition $x(0) = \theta \in [-\pi, \pi]$. Heuristically, consider a particle whose position at time 0 is θ. If the particle travels with speed governed by the function f_w, then $x(t)$ is its position at time t. To emphasise the dependence on the initial state we shall also write $x_\theta(t)$.

Let f_w be a trigonometric polynomial, that is, a linear combination of Fourier basis functions with pre-specified values w_i at equidistant $x_i \in [-\pi, \pi]$ under the constraint that $f_w(-\pi) = f_w(\pi) = 0$. More precisely, let $-\pi = x_0 < x_1 < \cdots < x_{2m} < \pi$, $w_0 = 0$, and define $f_w(x) = \sum_{j=0}^{2m} w_j r_j(x)$ where

$$r_j(x) = \frac{\prod\limits_{\substack{j \neq k=0}}^{2m} \sin\left(\frac{x - x_k}{2}\right)}{\prod\limits_{\substack{j \neq k=0}}^{2m} \sin\left(\frac{x_j - x_k}{2}\right)}$$

for arbitrary $w_1, \ldots, w_{2m}$ and $m \geq 1$. Note that f_w vanishes at $-\pi$. Now, f_w is a C^1 function on $(-\pi, \pi)$ and

$$r_j'(x) = \frac{\sum\limits_{\substack{j \neq k=0}}^{2m} \cos\left(\frac{x - x_k}{2}\right) \prod\limits_{\substack{j,k \neq i=0}}^{2m} \sin\left(\frac{x - x_i}{2}\right)}{2 \prod\limits_{\substack{j \neq k=0}}^{2m} \sin\left(\frac{x_j - x_k}{2}\right)}$$

is Lipschitz. Therefore, by Coddington and Levinson (1955, Chapter 1.7), the function

$$\theta \mapsto x_\theta(1) = \theta + \int_0^1 \sum_{j=0}^{2m} w_j r_j(x_\theta(t)) dt,$$

the solution of (9) at time 1, is a diffeomorphism, cf. Younes (2010, Theorem 8.7). This function is known as the flow of the differential equation and denoted by $\varphi(\theta) = x_\theta(1)$. As it depends on the weights, we shall also write $\varphi_w(\theta)$ to emphasise this fact. Note that in total, there are $2m+1$ alignment parameters, $2m$ for the diffeomorphism and one for the shift in root.

5.1. Inference

Return to Model 1, that is, consider a curve

$$X_t(\theta) = \Gamma(\varphi_{w_t}(\theta - \alpha_t)) + N_t(\varphi_{w_t}(\theta - \alpha_t))$$

which is observed at $\theta_l = -(n+1)\pi/n + 2\pi l/n, l = 1, \ldots, n$, and extended to $[-\pi, \pi]$ by trigonometric interpolation. The latter is valid under our assumption that n is odd.

For each choice of w_t and α_t, the idea is to apply the inverse parametrisation to obtain curves

$$Y_t(\theta) = X_t(\varphi_{-w_t}(\theta) + \alpha_t)$$

that are well-aligned and apply the theory developed so far to the $Y_t, t = 0, \ldots, T$. Then, by (6), $\hat{\Gamma}_n(\theta) = \sum_{t=0}^{T} \hat{\Gamma}_{t,n}(\theta)/(T+1)$, where

$$\hat{\Gamma}_{t,n}(\theta) = \sum_{l=1}^{n} Y_t(\theta_l)S_l(\theta) = \sum_{l=1}^{n} X_t(\varphi_{-w_t}(\theta_l) + \alpha_t)S_l(\theta).$$

We pick the 'best' $\hat{\alpha}_t, \hat{w}_t$ by minimising

$$\sum_{t=0}^{T} \sum_{l=1}^{n} \|\hat{\Gamma}_{t,n}(\theta_l) - \hat{\Gamma}_n(\theta_l)\|^2, \tag{10}$$

the Riemann sum approximation to the total L_2-distance between the smoothed data curves and the estimated 'true' curve after alignment, over a compact cube containing the origin under the constraint $\alpha_0 = w_0 = 0$. Consequently, $\hat{\Gamma}_n$ can be interpreted as a Fréchet mean. Note that, alternatively, we could have constrained the average $\sum_t w_t$ to zero.

Criteria such as (10) are well known in the shape recognition literature. For an overview, see for example Davies et al. (2008); asymptotic properties can be found in Bigot and Gendre (2013).

5.2. An application

This section demonstrates our approach on geo-science data consisting of three images of a lake with fuzzy borders. The aims of the analysis are to summarise these images by a few parameters (estimated Fourier coefficients of the border curve), thus saving on storage, and to quantify the uncertainty.

Fig. 4 shows three 512×512 images of Lake Tana, the largest lake in Ethiopia and the source of the Blue Nile. It is approximately 84×66 km large and located near the centre of the high Ethiopian plateau. Clearly visible is Dek island, site of historic monasteries, in the south-central portion of the lake, which we shall use as the centre of our coordinate system. The three images were downloaded from NASA's 'The Gateway to Astronaut Photography of Earth' website (NASA, 2001) (frames 23, 24, 25). They were taken on February 17th, 2001, at one second intervals by astronauts on the STS098 mission from a space craft altitude of 383 km. The centre is at latitude 12.0 and longitude 37.5°. The cloud cover is about 25%.

Fig. 5 zooms in to part of the border of the middle image of Fig. 4. Note that it is rather fuzzy, resulting in a low image gradient, degraded even further by the substantial cloud cover. Therefore, for each of the three images shown in Fig. 4, $n = 73$ points were traced along the border using Gimp. The results are shown as, respectively, circles, triangles and crosses in the left-most panel in Fig. 6.

We start by considering shift parameters only. In other words, $w_t = 0$ for $t = 0, 1, 2$. Using $J = 20$ Fourier coefficients and $\alpha_0 = 0$, the optimal parameter values are $\hat{\alpha}_1 = -0.44$ and $\hat{\alpha}_2 = -2.33$ radians. The value of the optimisation function is 1195.048 square pixels, corresponding to
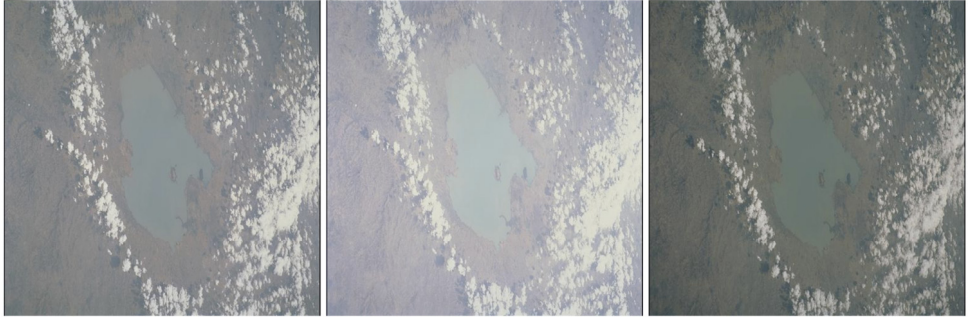
Fig. 4. Photographs of Lake Tana taken by astronauts on the STS098 mission. Images courtesy of the Image Science & Analysis Laboratory, NASA Johnson Space Center.

Fig. 5. Detail of the middle panel of Fig. 4.
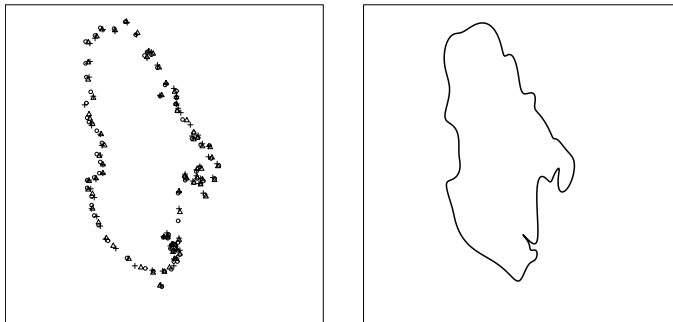
Fig. 6. Left panel: Sampled boundary curves for the images shown in Fig. 4. Circles trace the boundary of Lake Tana in the left-most panel, triangles correspond to the middle panel, and crosses trace the lake boundary in the right-most panel of Fig. 4. Right panel: Estimated border.

an average error of 2.34 pixels. The result can be improved by including diffeomorphic changes in speed. Optimising w_1, w_2 for vectors w_t, $t = 1, 2$, in $\mathbb{R}^{2m}$ with $m = 5$, we obtain a value of 568.0997 corresponding to an average error of 1.61 pixels, which confirms the visual impression from Fig. 5. Finally, the estimated curve $\hat{\Gamma}_n$ is plotted in the right-most panel in Fig. 6.

6. Discussion

In this paper, we formulated a model for objects with uncertain boundaries by generalising concepts from Jónsdóttir and Jensen (2005). The unknown boundary was estimated as a spectral mean by carrying out maximum likelihood estimation in the Fourier domain and transforming the results back to the spatial domain. We considered the integrated squared error and demonstrated how to deal with mis-alignment of the data. Finally, we applied the methods to simulated and real data.

The approach may be generalised to periodic change models. Indeed, write τ for the period. Then we may formulate the model

$$X_{j+t\tau}^l = X_{j+t\tau}(\theta_l) = \Gamma^{(j)}(\varphi_{j+t\tau}(\theta_l - \alpha_{j+t\tau})) + N_{j+t\tau}^{(j)}(\varphi_{j+t\tau}(\theta_l - \alpha_{j+t\tau})) \quad (11)$$

for $t \in \mathbb{N}_0$. Here the $\Gamma^{(j)}$ are unknown template curves at $j = 0, \ldots, \tau - 1$ steps into the period. The $N_{j+t\tau}^{(j)}$ are independent stationary cyclic Gaussian noise processes, the parameters of which may depend on j. In other words, (11) splits into τ sub-models of the form discussed in this paper.

Finally, it is worth noting that, although they are prevalent in shape analysis, diffeomorphisms have not been studied much in stochastic geometry. In this paper, they have been used in different roles: for curve modelling and for alignment. It seems to the author that there is scope for further research concerning the modelling of random compact sets by means of their boundary curves in light of the Jordan–Schönflies theorem (Keren, 2004).

Acknowledgements

Appendix. Proofs of the theorems in Section 4

Proof of Theorem 1.
Correlations
Recall the Lagrange identities stating that, for $\alpha \in (0, 2\pi)$,

$$\sum_{l=1}^{n} \sin(l\alpha) = \frac{1}{2} \cot\left(\frac{\alpha}{2}\right) - \frac{\cos\left(\left(n + \frac{1}{2}\right)\alpha\right)}{2\sin\left(\frac{\alpha}{2}\right)};$$

$$\sum_{l=1}^{n} \cos(l\alpha) = -\frac{1}{2} + \frac{\sin\left(\left(n + \frac{1}{2}\right)\alpha\right)}{2\sin\left(\frac{\alpha}{2}\right)}.$$

These can be used to show that – under the stated assumptions – the family $e^{ij\theta_l}$ as a function of θ_l, $l = 1, \ldots, n$, is orthogonal for $j = -(n-1)/2, \ldots, (n-1)/2$. Hence, using the Euler formulas $\cos x = (e^{ix} + e^{-ix})/2$ and $\sin x = (e^{ix} - e^{-ix})/(2i)$, we have that, for $j_1, j_2 \in \{1, \ldots, J\}$,

$$\sum_{l=1}^{n} \cos(j_1\theta_l)\cos(j_2\theta_l) = \sum_{l=1}^{n} \sin(j_1\theta_l)\sin(j_2\theta_l) = \frac{n}{2}\mathbf{1}\{j_1 = j_2\} \quad (12)$$

and

$$\sum_{l=1}^{n} \cos(j_1\theta_l)\sin(j_2\theta_l) = 0.$$

Write $\tilde{F}_{j,n}^t$ for the first component of the random 2-vector $F_{j,n}^t$. Now, for $j_1, j_2 \in \{1, \ldots, J\}$,

$$\mathrm{Cov}(\tilde{F}_{j_1,n}^t, \tilde{F}_{j_2,n}^t) = \frac{4}{n^2} \sum_{l=1}^{n} \sum_{k=1}^{n} \rho(\theta_k - \theta_l)\cos(j_1\theta_l)\cos(j_2\theta_k)$$

$$= \frac{4}{n^2} \sum_{l=1}^{n} \cos(j_1\theta_l) \left[\sum_{k=1}^{n} \rho(\theta_k - \theta_l) \cos(j_2(\theta_k - \theta_l + \theta_l)) \right],$$

which equals

$$\frac{4}{n^2} \sum_{l=1}^{n} \cos(j_1\theta_l) \left[\sum_{k=1}^{n} \rho(\theta_k - \theta_l) \left\{ \cos(j_2(\theta_k - \theta_l)) \cos(j_2\theta_l) - \sin(j_2(\theta_k - \theta_l)) \sin(j_2\theta_l) \right\} \right].$$

Note that the family $\{\theta_k - \theta_l : k = 1, \dots, n\}$ interpreted cyclically on $[-\pi, \pi]$ for fixed l is equal to the family $\{\theta_k : k = 1, \dots, n\}$ under the assumptions of Theorem 1. Therefore

$$\mathrm{Cov}(\tilde{F}_{j_1,n}^t, \tilde{F}_{j_2,n}^t) = \frac{4}{n^2} \sum_{l=1}^{n} \cos(j_1\theta_l) \cos(j_2\theta_l) \sum_{k=1}^{n} \rho(\theta_k) \cos(j_2\theta_k)$$

$$- \frac{4}{n^2} \sum_{l=1}^{n} \cos(j_1\theta_l) \sin(j_2\theta_l) \sum_{k=1}^{n} \rho(\theta_k) \sin(j_2\theta_k).$$

By the properties of the discrete Fourier basis, we conclude that $\mathrm{Cov}(\tilde{F}_{j_1,n}^t, \tilde{F}_{j_2,n}^t)$ is equal to $2 \sum_{k=1}^{n} \rho(\theta_k) \cos(j_2\theta_k)/n$ if $j_1 = j_2$ and zero otherwise. Similar calculations hold for the components of $G_{j,n}^t$ and $F_{0,n}^t$ as well as for the cross-covariances. We conclude that the (components of) $F_{j,n}^t$ and $G_{j,n}^t$ are mutually uncorrelated and normally distributed with means $\mu_{j,n}$, respectively $\nu_{j,n}$, and variance $\sigma_{j,n}^2$.

Log likelihood

By the preceding discussion, upon observation of $F_{j,n}^t = f_{j,n}^t$, $G_{j,n}^t = g_{j,n}^t$, the log likelihood after ignoring constants is

$$-(T+1) \left[\log \sigma_{0,n}^2 + \sum_{j=1}^{J} 2 \log \sigma_{j,n}^2 \right] - \frac{1}{2\sigma_{0,n}^2} \sum_{t=0}^{T} \| f_{0,n}^t - \mu_{0,n} \|^2$$

$$- \sum_{t=0}^{T} \sum_{j=1}^{J} \frac{\left[\| f_{j,n}^t - \mu_{j,n} \|^2 + \| g_{j,n}^t - \nu_{j,n} \|^2 \right]}{2\sigma_{j,n}^2}$$

so that $\hat{\mu}_{j,n}$, $\hat{\nu}_{j,n}$ and $\hat{\sigma}_{j,n}^2$ are maximum likelihood estimators.

Distributions of estimators

Recalling the first part of the proof, the distributions of $\hat{\mu}_{j,n}$, $\hat{\nu}_{j,n}$ and $\hat{\sigma}_{j,n}^2$ follow from classic multivariate statistics.

Since the X_t^l are normally distributed with independent components, $\hat{\Gamma}_n$ is a Gaussian random process with independent components. The expression for the mean vector follows from that for the mean vectors of $\hat{\mu}_{j,n}$ and $\hat{\nu}_{j,n}$. Finally,

$$\mathrm{Cov}(\hat{\Gamma}_n(\theta), \hat{\Gamma}_n(\eta)) = \mathrm{Cov}(\hat{\mu}_{0,n}) + \sum_{j=1}^{J} \left[\mathrm{Cov}(\hat{\mu}_{j,n}) \cos(j\theta) \cos(j\eta) + \mathrm{Cov}(\hat{\nu}_{j,n}) \sin(j\theta) \sin(j\eta) \right]$$

by the independence of the $\hat{\mu}_{j,n}$ and $\hat{\nu}_{j,n}$. Plugging in the expressions for the covariance matrices, the proof is complete. $\quad\square$

Proof of Theorem 2. By Parseval's identity (1),

$$\frac{1}{\pi} \int_{-\pi}^{\pi} \| \hat{\Gamma}_n(\theta) - \Gamma(\theta) \|^2 d\theta = 2\| \hat{\mu}_{0,n} - \mu_0 \|^2 + \sum_{j=1}^{\infty} \left[\| \hat{\mu}_{j,n} - \mu_j \|^2 + \| \hat{\nu}_{j,n} - \nu_j \|^2 \right],$$

since $\hat{\mu}_{j,n}$ and $\hat{\nu}_{j,n}$ are the Fourier coefficients of (6). Due to the truncation of (6) at J, $\hat{\mu}_{j,n} = \hat{\nu}_{j,n} = 0$ for $j \geq J + 1$.

By Theorem 1, the random vectors $\hat{\mu}_{j,n} - \mu_j$ and $\hat{\nu}_{j,n} - \nu_j$ are independent and normally distributed with mean vectors $\mu_{j,n} - \mu_j$ and $\nu_{j,n} - \nu_j$, respectively, and with covariance matrices whose diagonal

entries are $\sigma_{j,n}^2/(T+1)$. We conclude that, for $j = 1, \ldots, J$, $\|\hat{\mu}_{j,n} - \mu_j\|^2 + \|\hat{v}_{j,n} - v_j\|^2$ multiplied by $(T+1)/\sigma_{j,n}^2$ is the sum of four independent squared normals with unit variance but different means, that is, a non-central χ^2 distributed random variable with four degrees of freedom and non-centrality parameter $(T+1)c_{j,n}/\sigma_{j,n}^2$ with $c_{j,n} = \|\mu_{j,n} - \mu_j\|^2 + \|v_{j,n} - v_j\|^2$. For $j = 0$, $\|\hat{\mu}_{j,n} - \mu_j\|^2$ multiplied by $(T+1)/\sigma_{0,n}^2$ is the sum of two squared normals, hence a non-central χ^2 distributed random variable with two degrees of freedom and non-centrality parameter $(T+1)c_{0,n}/\sigma_{0,n}^2$ such that $c_{0,n} = \|\mu_{0,n} - \mu_0\|^2$. $\quad\square$

Proof of Theorem 3. Recall that we assumed that Γ is a C^1 function so that $\|\Gamma(\theta)\|$ and $\|\Gamma'(\theta)\|$ are bounded, say by γ_0 and γ_1. Furthermore,

$$v_{j,n} - v_j = \frac{1}{\pi} \int_{-\pi}^{\pi} \left(h_{j,n}(\theta) - h_j(\theta) \right) d\theta$$

for $h_j(\theta) = \Gamma(\theta)\sin(j\theta)$ and $h_{j,n}(\theta) = \Gamma(\theta_l)\sin(j\theta_l)$ if $|\theta - \theta_l| < \pi/n$. By the mean value theorem,

$$\|h_{j,n}(\theta) - h_j(\theta)\| \leq \frac{\pi H_j}{n}$$

on each interval $|\theta - \theta_l| < \pi/n$ where $H_j = \max_\theta \|h_j'(\theta)\|$. Hence $\|v_{j,n} - v_j\| \leq 2\pi H_j/n$. Consequently,

$$\sum_{j=1}^{J} \|v_{j,n} - v_j\|^2 \leq \sum_{j=1}^{J} \frac{4\pi^2 H_j^2}{n^2} \leq \sum_{j=1}^{J} \frac{4\pi^2(\gamma_1^2 + j^2\gamma_0^2)}{n^2}.$$

Thus, if $J^3/n^2 \to 0$, then $\sum_{j=1}^{J} \|v_{j,n} - v_j\|^2 \to 0$. A similar argument for the $\|\mu_{j,n} - \mu_j\|$ implies that $2c_{0,n} + \sum_{j=1}^{J} c_{j,n} \to 0$ provided $J^3/n^2 \to 0$ as n and J tend to infinity.

As for the variance term, note that since we assume (3) with $k = 1$, the covariance function is twice continuously differentiable (Rogers and Williams, 1994, p. 63). Therefore the second order derivative of $\rho(\theta)\cos(j\theta)$ is bounded by $j^2 R$ for some $R > 0$. Furthermore,

$$\sum_{j=1}^{J} \sigma_{j,n}^2 \leq \sum_{j=1}^{J} |\sigma_{j,n}^2 - \sigma_j^2| + \sum_{j=1}^{J} \sigma_j^2.$$

The Riemann middle sum error on $|\sigma_{j,n}^2 - \sigma_j^2|$ is bounded by Cj^2/n^2 for some $C > 0$, so that

$$\sum_{j=1}^{J} \sigma_{j,n}^2 \leq \sum_{j=1}^{J} \frac{Cj^2}{n^2} + \sum_{j=1}^{J} \sigma_j^2.$$

The second term goes to zero as J^3/n^2 tends to zero. Since the series $\sum_j \sigma_j^2$ converges and $1/(T+1) \to 0$ as $T \to \infty$, the proof is complete. $\quad\square$

References

Aletti, G., Ruffini, M., 2015. Is the Brownian bridge a good noise model on the circle? Ann. Inst. Statist. Math. (in press).
Baddeley, A.J., Molchanov, I.S., 1998. Averaging of random sets based on their distance functions. J. Math. Imaging Vision 8, 79–92.
Besl, P.J., McKay, N.D., 1992. A method for registration of 3D shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14, 239–256.
Bigot, J., Gendre, X., 2013. Minimax properties of Fréchet means of discretely sampled curves. Ann. Statist. 41, 923–956.
Burrough, P., Frank, A., 1996. Geographic Objects with Indeterminate Boundaries. Taylor & Francis, London.
Coddington, E.A., Levinson, N., 1955. Theory of Ordinary Differential Equations. McGraw–Hill, New York.
Davies, R., Twining, C., Taylor, C., 2008. Statistical Models of Shape: Optimisation and Evaluation. Springer, London.
Dempster, A.P., 1967. Upper and lower probabilities induced by a multivalued mapping. Ann. Statist. 38, 325–329.
Dryden, I.L., Mardia, K.V., 1998. Statistical Shape Analysis. Wiley, Chichester.
Fréchet, M., 1948. Les éléments aléatoires de nature quelconque dans un espace distancié. Ann. Inst. H. Poincaré Probab. Statist. 10, 235–310.
Glasbey, C.A., Mardia, K.V., 2001. A penalized likelihood approach to image warping. J. R. Stat. Soc. Ser. B Stat. Methodol. 63, 465–514.

Grenander, U., Miller, M.I., 2007. Pattern Theory: From Representation to Inference. Oxford, Oxford University Press.
Hobolth, A., Pedersen, J., Jensen, E.B.V., 2003. A continuous parametric shape model. Ann. Inst. Statist. Math. 55, 227–242.
Jónsdóttir, K.Y., Jensen, E.B.V., 2005. Gaussian radial growth. Image Anal. Stereol. 24, 117–126.
Kendall, D.G., 1984. Shape manifolds, Procrustean metrics and complex projective spaces. Bull. Lond. Math. Soc. 16, 81–121.
Keren, D., 2004. Topologically faithful fitting of simple closed curves. IEEE Trans. Pattern Anal. Mach. Intell. 26, 118–123.
Molchanov, I.S., 1997. Statistics of the Boolean Model for Practitioners and Mathematicians. Wiley, Chichester.
Molchanov, I.S., 2005. Theory of Random Sets. Springer, London.
NASA 2001. The gateway to astronaut photography of earth http://eol.jsc.nasa.gov/scripts/sseop/photo.pl?mission=STS098&roll=711&frame.
Rogers, L.C.G., Williams, D., 1994. Diffusions, Markov Processes, and Martingales. Volume One: Foundations, Second ed. Wiley, Chichester.
Sebastian, T.B., Klein, P.N., Kimia, B.B., 2003. On aligning curves. IEEE Trans. Pattern Anal. Mach. Intell. 25, 116–125.
Shafer, G., 1976. Mathematical Theory of Evidence. Princeton University Press, Princeton.
Soetaert, K., Petzoldt, T., Woodrow Setzer, R., 2010. Solving differential equations in R: Package deSolve. J. Stat. Softw. 33, 1–25.
Younes, L., 2010. Shapes and Diffeomorphisms. Springer, Berlin.
Zimmermann, H.J., 2001. Fuzzy Set Theory and its Applications, Fourth ed. Kluwer, Dordrecht.