

The Shape of Data and Probability Measures

Diego H. Díaz Martínez^a, Facundo Mémoli^b, Washington Mio^a

^a*Department of Mathematics, Florida State University, Tallahassee, FL 32306-4510 USA*

^b*Department of Mathematics, Ohio State University, Columbus, OH 43210-1174 USA*

Abstract

We introduce the notion of multiscale covariance tensor fields (CTF) associated with Euclidean random variables as a gateway to the shape of their distributions. Multiscale CTFs quantify variation of the data about every point in the data landscape at all spatial scales, unlike the usual covariance tensor that only quantifies global variation about the mean. Empirical forms of localized covariance previously have been used in data analysis and visualization, but we develop a framework for the systematic treatment of theoretical questions and computational models based on localized covariance. We prove strong stability theorems with respect to the Wasserstein distance between probability measures, obtain consistency results, as well as estimates for the rate of convergence of empirical CTFs. These results ensure that CTFs are robust to sampling, noise and outliers. We provide numerous illustrations of how CTFs let us extract shape from data and also apply CTFs to manifold clustering, the problem of categorizing data points according to their noisy membership in a collection of possibly intersecting, smooth submanifolds of Euclidean space. We prove that the proposed manifold clustering method is stable and carry out several experiments to validate the method.

Keywords: shape of data, multiscale data analysis, covariance fields, Fréchet functions, manifold clustering

1. Introduction

Probing, analyzing and visualizing the shape of complex data are challenges that are magnified by the intricate dependence of their structural properties, as basic as dimensionality, on location and scale (cf. [1]). As such, resolving and integrating the geometry and topology of data across scales are

problems of foremost importance. In this paper, we develop the notion of multiscale covariance tensor fields (CTF) associated with Euclidean random variables and show that many properties of the shape of their distributions become accessible through CTFs, which provide stable representations that can be estimated reliably from data.

For a random vector $y \in \mathbb{R}^d$, scale dependence is controlled by a kernel function $K(x, y, \sigma) \geq 0$, where $x, y \in \mathbb{R}^d$ and $\sigma > 0$ is the scale parameter. The idea is that from the standpoint of x , at scale $\sigma > 0$, the kernel masks the distribution by attributing weight $K(x, y, \sigma)$ to data located at y , creating a windowing effect. More simply put, $K(x, y, \sigma)$ quantifies how well an observer at x sees data at y at scale σ . Covariation of the weighted data is measured relative to every point $x \in \mathbb{R}^d$, not just about the mean as is common practice, thus giving rise to a multiscale covariance field. Special cases of these covariance fields were introduced in [2], targeting applications to such problems as detection of local scales and feature rich points in shapes. Here we present a more systematic treatment that includes a broader formulation of multiscale CTFs, stability theorems that ensure that properties of probability measures derived from multiscale CTFs are robust, as well as consistency results and convergence rates for empirical CTFs. We prove stability of CTFs with respect to the Wasserstein distance between probability measures, a metric that is finding uses in an ever expanding landscape of problems and whose origins are in optimal transport theory [3, 4]. Since Wasserstein distance metrizes weak convergence of probability measures, we obtain a strong stability result that ensures that if two probability distributions are similar in a weak sense, then their multiscale CTFs are uniformly close over the entire domain. Convergence rates are derived from the stability theorems and results by Fournier and Guillin [5] and García-Trillos and Slepčev [6] on convergence of empirical measures. The standard covariance tensor of a random vector $y \in \mathbb{R}^d$ quantifies covariation of y about the mean, but may be extended to a full covariance field by considering covariation about arbitrary points. Nonetheless, this field provides no information about the organization of the data other than that already contained in the covariance about the mean. Thus, a localized formulation is essential for gaining additional insight into the shape of data.

The trace of a multiscale CTF is a scalar field that gives a multiscale analogue of the classical Fréchet function $V(x) = \mathbb{E}[\|y - x\|^2]$ of a random variable y with finite second moment. The Fréchet function provides a more geometric interpretation of the mean as the unique minimizer of V ; that is,

the point $\mu \in \mathbb{R}^d$ with respect to which the spread of y is minimal. Similarly, the local extrema and other properties of the multiscale Fréchet function provide a wealth of information about the distribution of y . In fact, we show that the distribution of any random vector may be fully recovered from the multiscale Fréchet function associated with the Gaussian kernel.

Several variants of empirical localized or weighted covariance have been used in various problems in data analysis, but we develop a framework for the formulation and systematic treatment of such problems. Allard et al. have developed a computational model termed geometric multi-resolution analysis for multiscale data analysis based on covariance localized to hierarchies of dyadic cubes [7]. In computer graphics, localized covariance is commonly used in the estimation of normals to surfaces from point-cloud data [8] for surface reconstruction; see also [9] and references therein. In computer vision, tensor voting by Medioni et al. [10] has been applied to multiple image analysis and processing tasks. Brox et al. have used empirical covariance weighted by the isotropic Gaussian kernel in non-parametric density estimation targeting applications in motion tracking [11].

The paper includes several illustrations and applications of CTFs to data analysis. For example, to illustrate how geometric information can be extracted from CTFs, we show that the curvature of plane curves and the principal curvatures of surfaces in \mathbb{R}^3 can be calculated from the spectrum of multiscale CTFs. Thus, multiscale covariance tensors give a way of extending these infinitesimal measures of geometric complexity to all scales and general probability distributions, not just those supported on smooth submanifolds. We also apply multiscale CTFs to manifold clustering, the problem of clustering Euclidean data that are organized along a finite union of possibly intersecting smooth submanifolds. Fig.1 shows three such examples. The

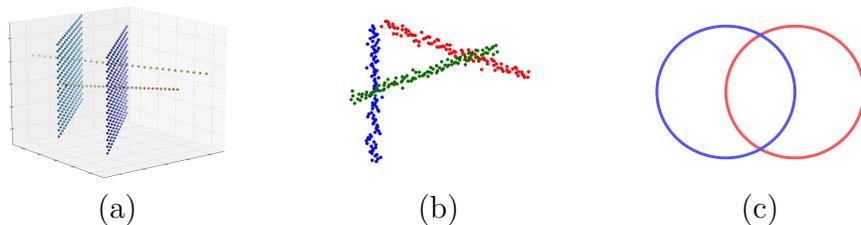


Figure 1: Examples of data clustered along intersecting manifolds.

special case of affine linear subspaces, known as subspace clustering, has

been addressed in the machine learning and computer vision literature by many authors using a variety of techniques (cf. [12, 13, 14, 15, 16, 17, 18]). More general manifold clustering has been considered in [19]. In our approach, we exploit the fact that localized covariance tensors encode rich information about the tangential structure of the submanifolds that underlie the data. Combined with information about the (relative) positions of the data points, they yield an effective data representation for manifold clustering. Although several different clustering techniques could be applied to the “tensorized” data, we use the single linkage hierarchical method because it produces provenly stable dendrograms. In conjunction with the stability and consistency results for covariance fields, this ensures that the manifold clustering method is stable at all steps. Dendrogram stability is analyzed in the framework of [20].

The main goals of the paper are: (i) to establish the foundations for analysis, visualization and management of data with methods based on multiscale covariance tensor fields, and (ii) to describe applications that characterize the usefulness of CTFs in data analysis. In Section 2, we formulate the notion of multiscale CTFs for a broad class of kernels and give examples that illustrate how CTFs reveal the geometry of data. In Section 3, we show that the curvature of a plane curve and the principal curvatures of a surface in \mathbb{R}^3 can be recovered from small-scale covariance. Section 4 is devoted to the main theoretical developments. We prove stability and consistency theorems for multiscale covariance tensor fields under mild regularity assumptions on the kernel, and also analyze rates of convergence that are important for applications in data analysis. Since some discontinuous kernels are of practical interest, we also investigate convergence results for such kernels, including a pointwise central limit theorem. Multiscale Fréchet functions are discussed in Section 5 and manifold clustering in Section 6. We close with a summary and some discussion.

2. Covariance Tensor Fields

2.1. Preliminaries

To define covariance tensor fields, we introduce some notation. Elements of the tensor product $\mathbb{R}^d \otimes \mathbb{R}^d$ may be identified with bilinear forms $B: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ through the Euclidean inner product. More precisely, a pure tensor $x \otimes y$ corresponds to the bilinear form

$$x \otimes y(u, v) = \langle x, u \rangle \cdot \langle y, v \rangle, \quad (1)$$

$\forall u, v \in \mathbb{R}^d$, where $\langle \cdot, \cdot \rangle$ denotes Euclidean inner product. Bilinear forms associated with more general elements of $\mathbb{R}^d \otimes \mathbb{R}^d$ can be described by linear extension. In Euclidean coordinates, we abuse notation and also write the coordinate vectors of $x, y \in \mathbb{R}^d$ as x and y . With this convention, letting A be the $d \times d$ matrix $A = xy^T$, we have

$$x \otimes y(u, v) = \langle u, Av \rangle, \quad (2)$$

where the superscript T denotes transposition. In this manner, using Euclidean coordinates, an element of $\mathbb{R}^d \otimes \mathbb{R}^d$ also can be identified with a $d \times d$ matrix by linear extension of the correspondence $x \otimes y \leftrightarrow A$. Through these identifications, we refer to an element $\Sigma \in \mathbb{R}^d \otimes \mathbb{R}^d$ interchangeably as a tensor, a bilinear form or a matrix. We equip $\mathbb{R}^d \otimes \mathbb{R}^d$ with the inner product defined on pure tensors by

$$\langle x_1 \otimes y_1, x_2 \otimes y_2 \rangle = \langle x_1, x_2 \rangle \langle y_1, y_2 \rangle \quad (3)$$

and extended linearly to $\mathbb{R}^d \otimes \mathbb{R}^d$. Thus, the corresponding norm satisfies

$$\|x \otimes y\| = \|x\| \|y\|, \quad (4)$$

for any $x, y \in \mathbb{R}^d$. In matrix representation, this is the Frobenius norm.

Throughout the paper, we view \mathbb{R}^d as a measurable space equipped with the Borel σ -algebra for the Euclidean metric. Let y be an \mathbb{R}^d -valued random variable distributed according to the probability measure α . Suppose that y has expected value $\mathbb{E}[y] = \mu \in \mathbb{R}^d$ and finite second moment. As a motivation for the definition of multiscale CTFs, recall that the covariance tensor of y is defined as

$$\Sigma_\alpha(\mu) = \mathbb{E}[(y - \mu) \otimes (y - \mu)] = \int_{\mathbb{R}^d} (y - \mu) \otimes (y - \mu) \alpha(dy) \in \mathbb{R}^d \otimes \mathbb{R}^d. \quad (5)$$

In matrix notation,

$$\Sigma_\alpha(\mu) = \int_{\mathbb{R}^d} (y - \mu)(y - \mu)^T \alpha(dy). \quad (6)$$

The bilinear form associated with $\Sigma_\alpha(\mu)$ clearly is symmetric and positive semi-definite.

Covariation of y may be measured with respect to any $x \in \mathbb{R}^d$, not just μ . Thus, $\Sigma_\alpha(\mu)$ may be extended to a global covariance tensor field $\Sigma_\alpha: \mathbb{R}^d \rightarrow \mathbb{R}^d \otimes \mathbb{R}^d$ given by

$$\Sigma_\alpha(x) = \int_{\mathbb{R}^d} (y - x) \otimes (y - x) \alpha(dy). \quad (7)$$

Note, however, that

$$\Sigma_\alpha(x) = \Sigma_\alpha(\mu) + (\mu - x) \otimes (\mu - x), \quad (8)$$

for any $x \in \mathbb{R}^d$. Thus, for $x \neq \mu$, $\Sigma_\alpha(x)$ does not reveal any information about the distribution of y other than that already contained in $\Sigma_\alpha(\mu)$. In contrast, as we shall see below, multiscale analogues are rich in information about the shape of α .

2.2. Multiscale Covariance Tensor Fields

We adopt the notation ν_d for the volume of the unit ball in \mathbb{R}^d and ω_{d-1} for the “surface area” of the unit sphere $\mathbb{S}^{d-1} \subset \mathbb{R}^d$, $d \geq 1$. Recall that $\omega_{d-1} = 2\pi^{d/2}/\Gamma(d/2)$, where $\Gamma(\cdot)$ is the usual Gamma function, and $\omega_{d-1} = d\nu_d$. We make the convention that $\nu_0 = 1$.

Let y be an \mathbb{R}^d -valued random variable with distribution α and let K be a multiscale kernel; that is, a measurable function $K: \mathbb{R}^d \times \mathbb{R}^d \times (0, \infty) \rightarrow \mathbb{R}$ such that $K(x, y, \sigma) \geq 0$, for any $x, y \in \mathbb{R}^d$ and $\sigma > 0$.

Definition 1. The *multiscale covariance tensor field* (CTF) of y associated with the kernel K is the one-parameter family of tensor fields, indexed by $\sigma \in (0, \infty)$, given by

$$\Sigma_\alpha(x, \sigma) := \int_{\mathbb{R}^d} (y - x) \otimes (y - x) K(x, y, \sigma) \alpha(dy), \quad (9)$$

provided that the integral converges for each $x \in \mathbb{R}^d$ and $\sigma > 0$.

Remark 1. Note that Σ_α depends only on the probability measure α , not on y . For this reason, we refer to Σ_α interchangeably as the multiscale CTF of the random variable y or the probability measure α .

$\Sigma_\alpha(x, \sigma)$ measures the covariation of y about x with probability mass at y weighted by $K(x, y, \sigma)$. It is simple to verify that the bilinear form $\Sigma_\alpha(x, \sigma)$ is symmetric and positive semi-definite. Note that if K is bounded

for each $\sigma > 0$, that is, $\exists M_\sigma > 0$ such that $K(x, y, \sigma) \leq M_\sigma, \forall x, y \in \mathbb{R}^d$, then $\Sigma_\alpha(x, \sigma)$ is well defined for any random variable y with finite second moment. In particular if $K \equiv 1$, $\Sigma_\alpha(x, \sigma) = \Sigma_\alpha(x), \forall x \in \mathbb{R}^d$. However, as our primary goal is to study the organization of data and random variables at scales ranging from local to global, we consider kernels in \mathbb{R}^d that satisfy additional decay conditions as they produce a windowing effect. The kernels are constructed as follows.

Definition 2. Let d be a positive integer and $f: [0, \infty) \rightarrow \mathbb{R}$ a bounded and measurable function satisfying:

- (a) $f(r) \geq 0, \forall r \in [0, \infty)$;
- (b) $M_d = \int_0^\infty r^{\frac{d}{2}-1} f(r) dr < \infty$;
- (c) There is $C > 0$ such that $rf(r) \leq C, \forall r \in [0, \infty)$.

The multiscale kernel $K: \mathbb{R}^d \times \mathbb{R}^d \times (0, \infty) \rightarrow \mathbb{R}$ associated with f is defined as

$$K(x, y, \sigma) := \frac{1}{C_d(\sigma)} f\left(\frac{\|y - x\|^2}{\sigma^2}\right), \quad (10)$$

where $C_d(\sigma) = \frac{1}{2}\sigma^d M_d \omega_{d-1}$.

Condition (b) in the definition implies that the normalizing constant $C_d(\sigma)$ is well defined. The normalization is adopted so that $\int K(x, y, \sigma) dy = 1, \forall x \in \mathbb{R}^d$ and $\forall \sigma > 0$. Condition (c) guarantees that the integral in (9) is convergent for any probability measure α . Henceforth, for convenience, we assume that $\sup f = 1$. This is not restrictive since scaling f does not change the kernel K because of the normalization.

Whereas we investigate properties of multiscale CTFs in a more general setting, our examples and experiments focus on two special kernels:

- (i) The isotropic Gaussian kernel

$$G(x, y, \sigma) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|y - x\|^2}{2\sigma^2}\right), \quad (11)$$

which is associated with the function $f(x) = e^{-x/2}$;

(ii) The truncation kernel

$$T(x, y, \sigma) = \frac{1}{\sigma^d \nu_d} \chi\left(\frac{\|y - x\|^2}{\sigma^2}\right) \quad (12)$$

associated with the characteristic function $\chi: [0, \infty) \rightarrow \mathbb{R}$ of the unit interval $[0, 1]$. In measuring covariation of random variables about x , the kernel T attributes a uniform weight to mass at points within the closed ball of radius σ centered at x and weight zero to mass elsewhere.

Remark 2. The kernel K defined in (10) is homogeneous and isotropic; that is, for any isometry $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$, $K(\varphi(x), \varphi(y), \sigma) = K(x, y, \sigma)$, $\forall x, y \in \mathbb{R}^d$ and $\sigma > 0$. Moreover, if we write $\varphi(x) = Ux + b$, with $U \in O(d)$ and $b \in \mathbb{R}^d$, then

$$U \Sigma_\alpha(x, \sigma) U^T = \Sigma_{\varphi_*\alpha}(\varphi(x), \sigma), \quad (13)$$

for any $(x, \sigma) \in \mathbb{R}^d \times (0, \infty)$. Here $O(d)$ is the group of $d \times d$ orthogonal matrices and $\varphi_*(\alpha)$ is the pushforward of α under φ .

Remark 3. Multiscale covariance tensor fields can be defined for any positive Borel measure α that satisfies

$$\int_{\mathbb{R}^d} \|z\|^2 f(\|z\|^2) \alpha(dz) < \infty, \quad (14)$$

not just for probability measures. In particular, if f has compact support, covariance fields are defined for any locally finite Borel measure α ; that is, measures for which every point $p \in \mathbb{R}^d$ has an open neighborhood U_p such that $\alpha(U_p) < \infty$.

We conclude this section with examples that support our contention that multiscale covariance tensor fields are rich in information about the shape of data.

Example 1. This example shows that the spectrum of multiscale covariance tensors allow us to estimate the dimensionality of data in a scale dependent manner. We consider the data points y_1, \dots, y_n in \mathbb{R}^2 , shown in Figure 2, and calculate Σ_{α_n} centered at one of the data points for the Gaussian kernel at scales $\sigma = 0.1$ and $\sigma = 2$. Here α_n denotes the empirical measure $n^{-1} \sum_{i=1}^n \delta_{y_i}$. The covariance tensors are depicted as ellipses whose principal axes are in the direction of the eigenvectors of the covariance matrix and

principal radii are proportional to $\sqrt{\lambda_1}$ and $\sqrt{\lambda_2}$, where $0 \leq \lambda_1 \leq \lambda_2$ are the eigenvalues of the covariance. At scale $\sigma = 0.1$, $\lambda_1/\lambda_2 = 0.908$, showing that the covariance tensor is nearly isotropic, indicating that the “dimension” of the data is 2. At $\sigma = 2$, the ratio of the eigenvalues is 0.025, giving a highly anisotropic covariance tensor, from which we infer that the dimension is 1.

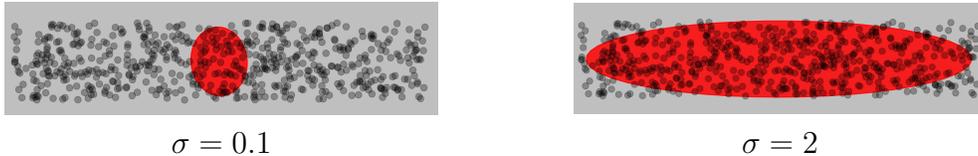


Figure 2: Estimating data dimensionality at different scales through multiscale covariance.

Example 2 (A linear subspace of \mathbb{R}^d). Let $v_1, \dots, v_r \in \mathbb{R}^d$, $1 \leq r \leq d$, be orthonormal vectors and consider the subspace $H = \langle v_1, \dots, v_r \rangle$ that they span. Let α denote the singular measure supported on H induced by the volume form on H . The measure α clearly is locally finite. We calculate multiscale covariance fields at points $x \in H$ to show that H may be recovered from $\Sigma_\alpha(x, \sigma)$. By Remark 2, we may assume that $x = 0$. A calculation shows that for the Gaussian kernel,

$$\Sigma_\alpha(0, \sigma) = \frac{r}{(\sqrt{2\pi})^{d-r} \sigma^{d-r-2}} \sum_{i=1}^r v_i \otimes v_i. \quad (15)$$

For the truncation kernel,

$$\Sigma_\alpha(0, \sigma) = \lambda_r \sum_{i=1}^r v_i \otimes v_i, \quad (16)$$

where

$$\lambda_r = \frac{1}{\sigma^{d-r-2}} \frac{\nu_{r-1}}{\nu_d} \int_{-\pi/2}^{\pi/2} \sin^2 \theta \cos^r \theta d\theta. \quad (17)$$

For $r = 1$, this expression simplifies to $\lambda_1 = 2/(3\sigma^{d-3}\nu_d)$. Thus, for both kernels, the orthogonal complement of H is the null space of $\Sigma_\alpha(0, \sigma)$ and H is the eigenspace associated with the positive eigenvalue λ_r .

Example 3 (Wedge of n segments). Consider the wedge (one-point union) W of n segments L_1, \dots, L_n in \mathbb{R}^d attached at the origin, as depicted in Fig. 3.

Each segment L_i is determined by its length $\ell_i > 0$ and a unit direction vector v_i . We assume that $v_i \neq v_j$, for any $1 \leq i < j \leq n$. Let α be the singular measure on \mathbb{R}^d that is supported on W and agrees with the measure induced by arc length on each segment L_i . We consider the multiscale covariance field of α associated with the truncation kernel. For $x \in L_i$, $x \neq 0$, as in the case $r = 1$ in Example 2, we have that $\Sigma_\alpha(x, \sigma) = (2/3\sigma^{d-3}\nu_d) v_i \otimes v_i$ at small enough scales. Thus, $\Sigma_\alpha(x, \sigma)$ has rank one. However, at the origin,

$$\Sigma_\alpha(0, \sigma) = \frac{1}{3\sigma^d\nu_d} \sum_{i=1}^n (\min(\sigma, \ell_i))^3 v_i \otimes v_i. \quad (18)$$

for any $\sigma > 0$. Thus, for $\sigma \leq \min\{\ell_i, 1 \leq i \leq n\}$,

$$\Sigma_\alpha(0, \sigma) = \frac{1}{3\sigma^{d-3}\nu_d} \sum_{i=1}^n v_i \otimes v_i. \quad (19)$$

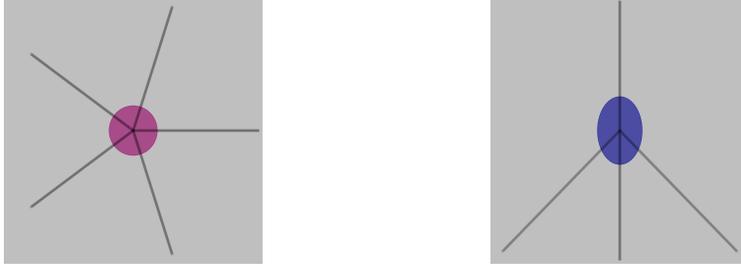


Figure 3: Covariance at the one-point union of line segments.

3. Geometry of Curves and Surfaces

In this section, we show how multiscale CTFs associated with the truncation kernel extract precise local geometric information from plane curves and surfaces in \mathbb{R}^3 .

3.1. Plane Curves

Example 4. We begin with the special case of a circle. Let $C_R \subset \mathbb{R}^2$ be the circle of radius R centered at the origin in \mathbb{R}^2 and α the singular measure

supported on C_R induced by arc length. For any $x \in \mathbb{R}^2$, we denote $r = \|x\|$. If x is such that $|r - R| > \sigma$ then $\Sigma_\alpha(x, \sigma) = 0$. Assume that $x \in \mathbb{R}^2$ and $0 < \sigma < R$ are such that $r \in [R - \sigma, R + \sigma]$. In this case, in the coordinate system given by the directions $n = x/\|x\|$ and $t = n^\perp$, a calculation shows that $\Sigma_\alpha(x, \sigma)$ is diagonal with entries

$$\begin{aligned}\lambda_n(x, \sigma) &= \frac{1}{\pi\sigma^2} [R\phi (R^2 + 2r^2) + R^2(R \cos \phi - 4r) \sin \phi] \\ \lambda_t(x, \sigma) &= \frac{R^3}{\pi\sigma^2} (\phi - \sin \phi \cos \phi),\end{aligned}\tag{20}$$

where $\phi = \arccos\left(\frac{R^2+r^2-\sigma^2}{2rR}\right)$. Thus, the normal and tangential vectors, n and t , are eigenvectors with eigenvalues λ_n and λ_t , respectively. Fig. 4 shows the eigenvalues as functions of r , $0.9 \leq r \leq 1.1$, for $\sigma = 0.1$.

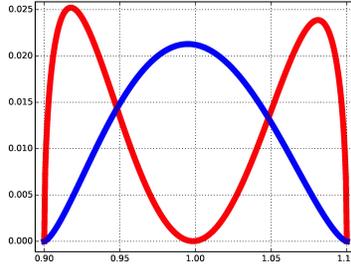


Figure 4: Tangential (blue) and normal (red) eigenvalues as a function of r , $0.9 \leq r \leq 1.1$, at $\sigma = 0.1$, of the multiscale CTF associated with the truncation kernel for the singular measure induced by arc length, supported on the unit circle in \mathbb{R}^2 .

Now we consider a general smooth curve $C \subset \mathbb{R}^2$, that is, a 1-dimensional, smooth, properly embedded submanifold of \mathbb{R}^2 . Let α be the singular measure on \mathbb{R}^2 supported on C and induced by arc length. This measure is locally finite because the embedding is proper. We calculate the small-scale covariance at points on C for the truncation kernel and show that the curvature can be recovered from the eigenvalues of Σ_α . Let $x \in C$ be fixed. The arc-length parametrization of C near x may be written as

$$X(s) = s - \frac{\kappa^2 s^3}{6} + O(s^4) \quad \text{and} \quad Y(s) = \frac{\kappa s^2}{2} + \frac{\kappa_s s^3}{6} + O(s^4), \tag{21}$$

where $X(s)$ and $Y(s)$ are coordinates along the tangent and normal to C at x , respectively [21]. Here, the curvature κ and its derivative κ_s are evaluated at x . A calculation yields:

Proposition 1. *Let $\sigma > 0$ be small. If C is a smooth plane curve and $x \in C$, then in the coordinates specified above we have*

$$\Sigma_\alpha(x, \sigma) = \begin{pmatrix} \frac{2\sigma}{3\pi} - \frac{\kappa^2\sigma^3}{20\pi} + O(\sigma^4) & \frac{\kappa_s\sigma^3}{15\pi} + O(\sigma^4) \\ \frac{\kappa_s\sigma^3}{15\pi} + O(\sigma^4) & \frac{\kappa^2\sigma^3}{10\pi} + O(\sigma^4) \end{pmatrix}. \quad (22)$$

Proposition 1 implies that, for $\sigma > 0$ small, the eigenvalues of Σ_α are

$$\lambda_1 = \frac{2\sigma}{3\pi} - \frac{\kappa^2\sigma^3}{20\pi} + O(\sigma^4) \quad \text{and} \quad \lambda_2 = \frac{\kappa^2\sigma^3}{10\pi} + O(\sigma^4), \quad (23)$$

so that

$$\text{tr } \Sigma_\alpha(x, \sigma) = \frac{2\sigma}{3\pi} + \frac{\kappa^2\sigma^3}{20\pi} + O(\sigma^4). \quad (24)$$

Thus, the curvature at $x \in C$ may be recovered, up to a sign, as

$$\kappa = \pm \lim_{\sigma \rightarrow 0} \frac{\sqrt{20\pi}}{\sigma^{3/2}} \left(\text{tr } \Sigma_\alpha(x, \sigma) - \frac{2\sigma}{3\pi} \right)^{1/2}. \quad (25)$$

3.2. Surfaces in \mathbb{R}^3

Example 5. Let S_R be the sphere of radius R centered at the origin in \mathbb{R}^3 . For $x \in \mathbb{R}^3$, we let $r = \|x\|$. If x is such that $|r - R| > \sigma$, then $\Sigma_\alpha(x, \sigma) = 0$. Assume that $x \neq (0, 0, 0)$ and $\sigma > 0$ are such that $r \in [R - \sigma, R + \sigma]$. In the coordinate system given by the vector $n = x/\|x\|$, and any orthonormal basis $\{t_1, t_2\}$ of the orthogonal complement n^\perp , a direct calculation shows that $\Sigma_\alpha(x, \sigma)$ is a 3×3 diagonal matrix with entries

$$\begin{aligned} \lambda_{t_1}(x, \sigma) &= \lambda_{t_2}(x, \sigma) = \frac{R^4}{\sigma^3} \sin^4 \left(\frac{\phi}{2} \right) (\cos \phi + 2) \\ \lambda_n(x, \sigma) &= \frac{R^2}{2\sigma^3} (1 - \cos \phi) (R^2 + R \cos \phi (R \cos \phi + R - 3r)) \\ &\quad + \frac{R^2}{2\sigma^3} (1 - \cos \phi) (-3Rr + 3r^2) \end{aligned} \quad (26)$$

where $\phi = \arccos \left(\frac{R^2 + r^2 - \sigma^2}{2Rr} \right)$. In particular, this means that λ_n is the eigenvalue corresponding to the eigenvector n along the normal direction to

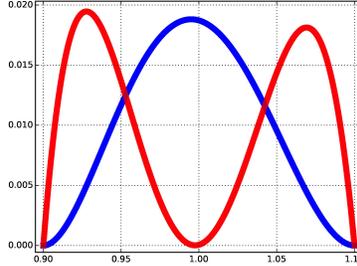


Figure 5: Tangential (blue) and normal eigenvalues (red) of the multiscale CTF associated with the truncation kernel as a function of r , $0.9 \leq r \leq 1.1$, at $\sigma = 0.1$, for the singular measure induced by surface area, supported on the unit sphere in \mathbb{R}^3 .

the sphere at $x/\|x\|$, and $\{t_1, t_2\}$ span the eigenspace along the tangent directions with eigenvalue $\lambda_{t_1} = \lambda_{t_2}$. Fig. 5 shows a plot of the eigenvalues as a function of r , $0.9 \leq r \leq 1$, for $\sigma = 0.1$.

Now we consider a general smooth compact surface $S \subset \mathbb{R}^3$. Let α be the singular measure on \mathbb{R}^3 supported on S and induced by the area measure on S . We calculate the small-scale covariance at points on S for the truncation kernel and show that the principal curvatures may indeed be recovered from the spectrum of Σ . Given a non-umbilic point $p \in S$, one can choose a Cartesian coordinate system centered at p so that the x -axis is along the direction of maximal curvature at p , the y -axis is along the direction of minimal curvature at p , and the z -axis is along the normal to S at p .

Proposition 2. *Let $\sigma > 0$ be small, $p \in S$ be non-umbilic, and α be the surface area measure on S . In the coordinate system described above, the covariance tensor for the truncation kernel is given by*

$$\Sigma_\alpha(p, \sigma) = \begin{bmatrix} A_{t_1} & O(\sigma^4) & O(\sigma^5) \\ O(\sigma^4) & A_{t_2} & O(\sigma^5) \\ O(\sigma^5) & O(\sigma^5) & A_n \end{bmatrix}, \quad (27)$$

where

$$\begin{aligned}
A_{t_1} &= \frac{3\sigma}{16} + \frac{1}{256}(-3\kappa_1^2 - 6\kappa_1\kappa_2 + \kappa_2^2)\sigma^3 + O(\sigma^4), \\
A_{t_2} &= \frac{3\sigma}{16} + \frac{1}{256}(\kappa_1^2 - 6\kappa_1\kappa_2 - 3\kappa_2^2)\sigma^3 + O(\sigma^4), \\
A_n &= \frac{3\kappa_1^2 + 2\kappa_1\kappa_2 + 3\kappa_2^2}{128}\sigma^3 + O(\sigma^4),
\end{aligned} \tag{28}$$

and $\kappa_1 > \kappa_2$ are the principal curvatures of S at p .

It follows from this result that, for $\sigma > 0$ small,

$$\begin{aligned}
\text{tr } \Sigma_\alpha(p, \sigma) &= \frac{3}{16}\sigma + \frac{1}{64}(\kappa_1 - \kappa_2)^2\sigma^3 + O(\sigma^4) \text{ and} \\
\det \Sigma_\alpha(p, \sigma) &= (3\kappa_1^2 + 2\kappa_1\kappa_2 + 3\kappa_2^2) \frac{\pi^2}{2048}\sigma^{11} + O(\sigma^{12}).
\end{aligned} \tag{29}$$

As a consequence, κ_1 and κ_2 can be recovered from the spectrum of $\Sigma_\alpha(p, \sigma)$ as a function of σ . Indeed, from the small scale asymptotics of the trace and determinant of $\Sigma_\alpha(p, \sigma)$, we can extract the values of $(\kappa_1 - \kappa_2)^2$ and $3\kappa_1^2 + 2\kappa_1\kappa_2 + 3\kappa_2^2$ from which we can determine the values of κ_1 and κ_2 .

Proof of Proposition 2. Using cylindrical coordinates in the chosen reference system, we can parametrize the patch $S \cap B(p, \sigma)$ as $(\rho \cos \phi, \rho \sin \phi, z(\rho, \phi))$, for $\phi \in [0, 2\pi]$, $\rho \in [0, \rho_\sigma(\phi)]$, where $\rho_\sigma(\phi) = \sigma - \frac{1}{8}(\kappa_1(\cos \phi)^2 + \kappa_2(\sin \phi)^2)^2\sigma^3 + O(\sigma^4)$, and $z(\rho, \phi) = \frac{\rho^2}{2}(\kappa_1(\cos \phi)^2 + \kappa_2(\sin \phi)^2) + O(\sigma^3)$. The area element on the patch is given by

$$dA = \left(\rho + \frac{\rho^3}{2}(\kappa_1^2(\cos \phi)^2 + \kappa_2^2(\sin \phi)^2) + O(\rho^5) \right) d\rho d\phi. \tag{30}$$

Now we have all the ingredients needed to compute $\Sigma_\alpha(p, \sigma)$. For example, to calculate the (1, 1)-entry, we express $\iint_{S \cap B(0, \sigma)} x^2 dA$ as

$$\int_0^{2\pi} \int_0^{\rho_\sigma(\phi)} \left[\rho^3 \cos^2 \phi + \frac{\rho^5 \cos^2 \phi}{2} (\kappa_2^2 \sin^2 \phi + \kappa_1^2 \cos^2 \phi) + O(\rho^6) \right] d\rho d\phi, \tag{31}$$

which after a simple but tedious calculation yields the desired result. The computation of other entries of the matrix follows similar steps. \square

4. Stability and Consistency

For each $p \in [1, \infty)$, let $\mathcal{P}_p(\mathbb{R}^d)$ denote the collection of all Borel probability measures α on \mathbb{R}^d whose p th moment $M_p(\alpha) = \int \|z\|^p \alpha(dz)$ is finite. We adopt the notation $m_p(\alpha) = M_p^{1/p}(\alpha)$. For $p = \infty$, we let $\mathcal{P}_\infty(\mathbb{R}^d)$ be the collection of all Borel probability measures on \mathbb{R}^d with bounded support and $m_\infty(\alpha) = \sup\{\|z\|, z \in \text{supp}[\alpha]\}$.

Definition 3. For $p \in [1, \infty]$, we consider two subsets of $\mathcal{P}_p(\mathbb{R}^d)$:

- (i) For $c > 0$, $\mathcal{P}_p^c(\mathbb{R}^d)$ is defined as the subset comprising all measures for which $m_p(\alpha) \leq c$.
- (ii) For $c, \lambda > 0$, $\mathcal{P}_p^{(c, \lambda)}(\mathbb{R}^d)$ denotes the subset of all $\alpha \in \mathcal{P}_p^c(\mathbb{R}^d)$ such that $\alpha(A) \leq \lambda \mathcal{L}(A)$, for all measurable sets A , where \mathcal{L} stands for Lebesgue measure.

By Jensen's inequality, if $1 \leq q \leq p \leq \infty$, then $\mathcal{P}_p^c(\mathbb{R}^d) \subset \mathcal{P}_q^c(\mathbb{R}^d)$ and

$$m_q(\alpha) \leq m_p(\alpha) \leq c, \quad (32)$$

for any $\alpha \in \mathcal{P}_p^c(\mathbb{R}^d)$.

Example 6. If $\alpha \in \mathcal{P}_p(\mathbb{R}^d)$ is absolutely continuous with respect to the Lebesgue measure with density function $f \in \mathbb{L}^\infty(\mathbb{R}^d)$ satisfying $\|f\|_\infty \leq \lambda$, then $\alpha \in \mathcal{P}_p^{(c, \lambda)}(\mathbb{R}^d)$.

Let us recall the definition of the p -Wasserstein distance $W_p(\alpha, \beta)$ between $\alpha, \beta \in \mathcal{P}_p(\mathbb{R}^d)$. Let $\Gamma(\alpha, \beta)$ be the collection of all couplings of α and β ; that is, probability measures μ on $\mathbb{R}^d \times \mathbb{R}^d$ such that $(\pi_1)_* \mu = \alpha$ and $(\pi_2)_* \mu = \beta$, where $\pi_1, \pi_2: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ denote projections onto the first and second components, respectively.

Definition 4. For $p \in [1, \infty)$, the p -Wasserstein distance between $\alpha, \beta \in \mathcal{P}_p(\mathbb{R}^d)$ is given by

$$W_p(\alpha, \beta) := \inf_{\mu \in \Gamma(\alpha, \beta)} \left(\iint \|z_1 - z_2\|^p \mu(dz_1 \times dz_2) \right)^{1/p},$$

and the ∞ -Wasserstein distance between $\alpha, \beta \in \mathcal{P}_\infty(\mathbb{R}^d)$ by

$$W_\infty(\alpha, \beta) := \inf_{\mu \in \Gamma(\alpha, \beta)} \sup \{ \|z_1 - z_2\|, (z_1, z_2) \in \text{supp}[\mu] \}.$$

Remark 4.

- (i) For any $\alpha, \beta \in \mathcal{P}_p(\mathbb{R}^d)$, $p \in [1, \infty]$, there exists a coupling that realizes the infimum in the definition of $W_p(\alpha, \beta)$ (cf. [22]).
- (ii) It is a standard result that, for each $p \in [1, \infty)$, W_p defines a metric on $\mathcal{P}_p(\mathbb{R}^d)$ that is compatible with weak convergence of probability measures [3].
- (iii) For any $c > 0$ and $p \in [1, \infty]$, the metric space $(\mathcal{P}_p^c(\mathbb{R}^d), W_p)$ has diameter $\leq 2c$.
- (iv) If $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is an isometry, then $W_p(\alpha, \beta) = W_p(\varphi_*(\alpha), \varphi_*(\beta))$, for any $\alpha, \beta \in \mathcal{P}_p(\mathbb{R}^d)$.

4.1. Smooth Kernels

For $p, q \in [1, \infty]$ satisfying $1/p + 1/q = 1$, we write $\ell = \max\{p, 2q\} = \max\{p, 2p/(p-1)\}$. Note that $3 \leq \ell \leq \infty$ and the minimum value of ℓ is achieved for $\ell = p = 2q = 3$.

Theorem 1 (Stability for Smooth Kernels). *Let $f: [0, \infty) \rightarrow \mathbb{R}$ be as in Definition 2 with multiscale kernel K , $p \in [1, \infty)$, and $\ell = \max\{p, 2p/(p-1)\}$. Suppose that f is differentiable and $\sqrt{r}|f'(r)| \leq A$, $\forall r \in [0, \infty)$ and some constant $A > 0$. If $\alpha, \beta \in \mathcal{P}_\ell^c(\mathbb{R}^d)$, then the multiscale covariance tensor fields for α and β associated with the kernel K satisfy*

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \frac{2c}{C_d(\sigma)} (1 + Ac/\sigma) W_p(\alpha, \beta),$$

for any $\sigma > 0$. Here $\|\cdot\|$ is the norm associated with the inner product defined in (3).

Theorem 1 shows that multiscale covariance fields yield a robust representation of probability measures that make their geometric properties more readily accessible, as illustrated in our examples. In Section 5, we show that not only is $\Sigma_\alpha(\cdot, \sigma)$ stable, but all the information contained in the probability measure α is fully absorbed into the multiscale CTF associated with the Gaussian kernel. In fact, α may be recovered from the multiscale scalar field given by $V_\alpha(x, \sigma) = \text{tr} \Sigma_\alpha(x, \sigma)$, $x \in \mathbb{R}^d$ and $\sigma > 0$.

The following lemma will be used in the proof of the stability theorem for smooth kernels. Let $K_\sigma: \mathbb{R}^d \rightarrow \mathbb{R}$ be given by

$$K_\sigma(x) = K(x, 0, \sigma) = \frac{1}{C_d(\sigma)} f\left(\frac{\|x\|^2}{\sigma^2}\right). \quad (33)$$

Note that $0 \leq K_\sigma(x) \leq 1/C_d(\sigma)$, since we are assuming that $\sup f = 1$.

Lemma 1. *If f is differentiable and $\sqrt{r}|f'(r)| \leq A$, $\forall r \in [0, \infty)$, then*

$$|K_\sigma(z_1) - K_\sigma(z_2)| \leq \frac{2A}{\sigma C_d(\sigma)} \|z_1 - z_2\|,$$

for any $z_1, z_2 \in \mathbb{R}^d$ and $\sigma > 0$; that is, K_σ is Lipschitz continuous with constant $2A/\sigma C_d(\sigma)$.

Proof. Let $z(t) = tz_1 + (1-t)z_2$, $0 \leq t \leq 1$. Then,

$$\begin{aligned} |K_\sigma(z_1) - K_\sigma(z_2)| &= \left| \int_0^1 \frac{d}{dt} K_\sigma(z(t)) dt \right| \leq \int_0^1 \left| \frac{d}{dt} K_\sigma(z(t)) \right| dt \\ &= \int_0^1 |\nabla K_\sigma(z(t)) \cdot (z_1 - z_2)| dt \\ &\leq \|z_1 - z_2\| \int_0^1 \|\nabla K_\sigma(z(t))\| dt. \end{aligned} \quad (34)$$

Since $\nabla K_\sigma(z) = \frac{2}{\sigma^2 C_d(\sigma)} z f'\left(\frac{\|z\|^2}{\sigma^2}\right)$, we obtain

$$\|\nabla K_\sigma(z)\| = \frac{2}{\sigma C_d(\sigma)} \frac{\|z\|}{\sigma} \left| f'\left(\frac{\|z\|^2}{\sigma^2}\right) \right| \leq \frac{2A}{\sigma C_d(\sigma)}. \quad (35)$$

The lemma follows from (34) and (35). \square

Proof of Theorem 1. In proving the theorem, by Remarks 2 and 4, we may assume without loss of generality that $x \in \mathbb{R}^d$ is the origin. To simplify notation, let $Q_\sigma(z) = (z \otimes z)K_\sigma(z)$, with K_σ as in (33). Then,

$$\Sigma_\alpha(x, \sigma) = \int Q_\sigma(z_1) \alpha(dz_1) \quad \text{and} \quad \Sigma_\beta(x, \sigma) = \int Q_\sigma(z_2) \beta(dz_2). \quad (36)$$

Given $\eta > 0$ satisfying $W_p(\alpha, \beta) < \eta$, there is a coupling $\mu \in \Gamma(\alpha, \beta)$ such that

$$\iint \|z_1 - z_2\|^p \mu(dz_1 \times dz_2) < \eta^p. \quad (37)$$

Since μ has marginals α and β , we may write

$$\Sigma_\alpha(x, \sigma) = \iint Q_\sigma(z_1) \mu(dz_1 \times dz_2) \quad (38)$$

and similarly for $\Sigma_\beta(x, \sigma)$. Thus,

$$\|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \iint \|Q_\sigma(z_1) - Q_\sigma(z_2)\| \mu(dz_1 \times dz_2). \quad (39)$$

Writing

$$\begin{aligned} Q_\sigma(z_1) - Q_\sigma(z_2) &= (z_1 - z_2) \otimes z_1 K_\sigma(z_1) + z_2 \otimes (z_1 - z_2) K_\sigma(z_1) \\ &\quad + (z_2 \otimes z_2) (K_\sigma(z_1) - K_\sigma(z_2)), \end{aligned} \quad (40)$$

it follows that

$$\begin{aligned} \|Q_\sigma(z_1) - Q_\sigma(z_2)\| &\leq \|z_1 - z_2\| \|z_1\| K_\sigma(z_1) + \|z_1 - z_2\| \|z_2\| K_\sigma(z_1) \\ &\quad + \|z_2\|^2 |K_\sigma(z_1) - K_\sigma(z_2)|. \end{aligned} \quad (41)$$

Using Lemma 1 and the fact that $K_\sigma(z) \leq 1/C_d(\sigma)$, $\forall z \in \mathbb{R}^d$, we have that

$$\begin{aligned} \|Q_\sigma(z_1) - Q_\sigma(z_2)\| &\leq \frac{\|z_1 - z_2\| \|z_1\|}{C_d(\sigma)} + \frac{\|z_1 - z_2\| \|z_2\|}{C_d(\sigma)} \\ &\quad + \frac{2A}{\sigma C_d(\sigma)} \|z_1 - z_2\| \|z_2\|^2. \end{aligned} \quad (42)$$

Thus, (39) and (42) imply that

$$\begin{aligned} \|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| &\leq \frac{1}{C_d(\sigma)} \iint \|z_1 - z_2\| \|z_1\| \mu(dz_1 \times dz_2) \\ &\quad + \frac{1}{C_d(\sigma)} \iint \|z_1 - z_2\| \|z_2\| \mu(dz_1 \times dz_2) \\ &\quad + \frac{2A}{\sigma C_d(\sigma)} \iint \|z_1 - z_2\| \|z_2\|^2 \mu(dz_1 \times dz_2) \end{aligned} \quad (43)$$

For $p \in (1, \infty)$, applying Hölder's inequality to (43) and using (37), we can conclude that

$$\begin{aligned} \|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| &\leq \frac{\eta}{C_d(\sigma)} \left(\iint \|z_1\|^q \mu(dz_1 \times dz_2) \right)^{1/q} \\ &\quad + \frac{\eta}{C_d(\sigma)} \left(\iint \|z_2\|^q \mu(dz_1 \times dz_2) \right)^{1/q} \\ &\quad + \frac{2A\eta}{\sigma C_d(\sigma)} \left(\iint \|z_1\|^{2q} \mu(dz_1 \times dz_2) \right)^{1/q}. \end{aligned} \quad (44)$$

Since μ has marginals α and β , (32) and (44) imply that

$$\|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \left(\frac{2c}{C_d(\sigma)} + \frac{2Ac^2}{\sigma C_d(\sigma)} \right) \eta. \quad (45)$$

For $p = 1$, in (43) we may assume that $\|z_1\| \leq c$ and $\|z_2\| \leq c$ because $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$. Thus, (45) also holds in this case. Since (45) holds for any $\eta > W_p(\alpha, \beta)$, we conclude that

$$\|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \frac{2c}{C_d(\sigma)} \left(1 + \frac{Ac}{\sigma} \right) W_p(\alpha, \beta), \quad (46)$$

as claimed. \square

In what follows, given random vectors $y_1, \dots, y_n \in \mathbb{R}^d$ with the same distribution α , the associated empirical measure is denoted $\alpha_n = n^{-1} \sum_{i=1}^n \delta_{y_i}$.

Corollary 1 (Consistency for Smooth Kernels). *Let K be a multiscale kernel as in Theorem 1. If $\alpha \in \mathcal{P}_3(\mathbb{R}^d)$ and y_1, \dots, y_n are i.i.d. random variables with distribution α , then, for fixed $\sigma > 0$,*

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_\alpha(x, \sigma)\| \xrightarrow{n \uparrow \infty} 0$$

almost surely.

Proof. It follows from Theorem 1 that

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_\alpha(x, \sigma)\| \leq \frac{2c}{C_d(\sigma)} \left(1 + \frac{Ac}{\sigma} \right) W_3(\alpha_n, \alpha). \quad (47)$$

The conclusion follows from the fact that W_3 metrizes weak convergence of probability measures in $\mathcal{P}_3(\mathbb{R}^d)$ and Varadarajan's Theorem [23] about convergence of empirical measures on Polish spaces that ensures that α_n converges weakly to α almost surely. \square

Corollary 1 guarantees the asymptotic consistency of empirical CTFs. However, in applications, it is important to have estimates of the rate of convergence, which we derive from the stability theorem and a result of Fournier and Guillin [5, Theorem 1].

Theorem 2 (Fournier and Guillin, [5]). *Let $\alpha \in \mathcal{P}_s(\mathbb{R}^d)$, where $s > 1$. If y_1, \dots, y_n are i.i.d. random variables with distribution α and $p \in [1, s)$, then there exists a constant $\beta_p > 0$ such that*

$$\mathbb{E} [W_p(\alpha, \alpha_n)] \leq \beta_p m_s^p(\alpha) \cdot \begin{cases} n^{-\frac{s-p}{s}} + n^{-\frac{1}{2}} & \text{if } p > d/2 \text{ and } s \neq 2p; \\ n^{-\frac{s-p}{s}} + n^{-\frac{1}{2}} \log(1+n) & \text{if } p = d/2 \text{ and } s \neq 2p; \\ n^{-\frac{s-p}{s}} + n^{-\frac{p}{d}} & \text{if } p \in [1, d/2) \text{ and } s \neq d/(d-p), \end{cases}$$

for any $n \geq 1$.

Corollary 2. *Let $\sigma, c > 0$, $\alpha \in \mathcal{P}_4^c(\mathbb{R}^d)$ and let y_1, \dots, y_n , $n \geq 1$, be i.i.d. random variables with distribution α . If f is differentiable and $\sqrt{r} |f'(r)| \leq A$, $\forall r \in [0, \infty)$ and some constant $A > 0$, then*

$$\mathbb{E} \left[\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_{\alpha}(x, \sigma)\| \right] \leq \frac{2c^3}{C_d(\sigma)} \left(1 + \frac{Ac}{\sigma}\right) \beta_2 \cdot \begin{cases} n^{-\frac{1}{3}} + n^{-\frac{1}{2}} & \text{if } d = 1, 2, 3; \\ n^{-\frac{1}{3}} + n^{-\frac{1}{2}} \log(1+n) & \text{if } d = 4; \\ n^{-\frac{1}{3}} + n^{-\frac{2}{d}} & \text{if } d > 4. \end{cases}$$

Proof. Since $\alpha \in \mathcal{P}_4^c(\mathbb{R}^d)$, $m_3(\alpha) \leq m_4(\alpha) \leq c$. The conclusion follows by invoking Theorem 1 with $p = q = 2$ and the result of Fournier and Guillin with $p = 2$ and $s = 3$. \square

In Corollary 2, it is possible to relax the assumption on finiteness of the fourth moment of α to order $3 + \epsilon$, $0 < \epsilon < 1$, at the expense of obtaining a weaker upper bound. Conversely, if α has finite higher moments, then the upper bound can be sharpened.

The stability and consistency results guarantee that for kernels satisfying $\sqrt{r} |f'(r)| < A$, $\forall r > 0$, multiscale covariance fields are robust to sampling, noise and outliers. However, the results do not apply to some discontinuous kernels of practical interest. Nonetheless, we prove a stability theorem for the truncation kernel, as well as pointwise convergence results for more general kernels.

4.2. The Truncation Kernel

We begin our discussion of covariance fields associated with the truncation kernel with a stability theorem with respect to the ∞ -Wasserstein metric. In preparation for the proof of the theorem, we introduce some notation. For $0 \leq a < b$, let $R_d(a, b)$ be the annulus $\{y \in \mathbb{R}^d : a \leq \|y\| \leq b\}$. Let $s_d(a, b)$ denote the radial moment of inertia of $R_d(a, b)$ given by $s_d(a, b) = \int_{R_d(a, b)} \|y\|^2 dy$. We use the fact that for any $B \geq b$, the inequality

$$s_d(a, b) \leq (b - a) \frac{\omega_{d-1}}{d+2} \frac{B^{d+2}}{B-a} \quad (48)$$

holds. Indeed,

$$\begin{aligned} s_d(a, b) &= \frac{\omega_{d-1}}{d+2} (b^{d+2} - a^{d+2}) = \frac{\omega_{d-1}}{d+2} a^{d+2} ((b/a)^{d+2} - 1) \\ &= \frac{\omega_{d-1}}{d+2} a^{d+2} \frac{(b-a)}{a} \left(1 + \left(\frac{b}{a}\right) + \dots + \left(\frac{b}{a}\right)^{d+1} \right) \\ &\leq \frac{b-a}{B-a} \frac{\omega_{d-1}}{d+2} a^{d+2} \frac{B-a}{a} \left(1 + \left(\frac{B}{a}\right) + \dots + \left(\frac{B}{a}\right)^{d+1} \right) \\ &\leq (b-a) \frac{\omega_{d-1}}{(d+2)} \frac{B^{d+2} - a^{d+2}}{B-a} \leq (b-a) \frac{\omega_{d-1}}{d+2} \frac{B^{d+2}}{B-a}. \end{aligned} \quad (49)$$

Let $\mathcal{P}_\infty^c(\mathbb{R}^d)$ and $\mathcal{P}_\infty^{(c, \lambda)}(\mathbb{R}^d)$ be as in Definition 3.

Theorem 3 (Stability for the Truncation Kernel). *If $\alpha \in \mathcal{P}_\infty^{(c, \lambda)}(\mathbb{R}^d)$ and $\beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$, then the multiscale covariance tensor fields for α and β associated with the truncation kernel T satisfy*

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \left(\frac{2c}{\sigma^d \nu_d} + \frac{\lambda d (\sigma + 2c)^{d+2}}{2c(d+2)\sigma^d} \right) W_\infty(\alpha, \beta).$$

for any $\sigma > 0$. Here $\|\cdot\|$ is the norm associated with the inner product defined in (3).

Proof. Without loss of generality, we may assume that $x \in \mathbb{R}^d$ is the origin. Write $\eta = W_\infty(\alpha, \beta)$ and let $\mu \in \Gamma(\beta, \alpha)$ be a coupling that realizes η . Denote by $\mathbb{I}_B: \mathbb{R}^d \rightarrow \mathbb{R}$ the characteristic function of the closed ball of radius

1 centered at the origin and let y and y' be random variables distributed according to α and β , respectively. Using

$$y \otimes y = (y - y') \otimes y + y' \otimes (y - y') + y' \otimes y', \quad (50)$$

we may write

$$\begin{aligned} \Sigma_\beta(0, \sigma) &= \frac{1}{\sigma^d \nu_d} \iint (y \otimes y) \mathbb{I}_B\left(\frac{y}{\sigma}\right) \mu(dy \times dy') \\ &= \frac{1}{\sigma^d \nu_d} \iint (y - y') \otimes y \mathbb{I}_B\left(\frac{y}{\sigma}\right) \mu(dy \times dy') \\ &\quad + \frac{1}{\sigma^d \nu_d} \iint y' \otimes (y - y') \mathbb{I}_B\left(\frac{y}{\sigma}\right) \mu(dy \times dy') \\ &\quad + \frac{1}{\sigma^d \nu_d} \iint (y' \otimes y') \mathbb{I}_B\left(\frac{y}{\sigma}\right) \mu(dy \times dy'). \end{aligned} \quad (51)$$

The last term in (51), in turn, may be expressed as

$$\begin{aligned} &\frac{1}{\sigma^d \nu_d} \iint (y' \otimes y') \left(\mathbb{I}_B\left(\frac{y}{\sigma}\right) - \mathbb{I}_B\left(\frac{y'}{\sigma + \eta}\right) \right) \mu(dy \times dy') \\ &\quad + \frac{1}{\sigma^d \nu_d} \iint (y' \otimes y') \mathbb{I}_B\left(\frac{y'}{\sigma + \eta}\right) \mu(dy \times dy') \\ &\quad - \frac{1}{\sigma^d \nu_d} \iint (y' \otimes y') \mathbb{I}_B\left(\frac{y'}{\sigma}\right) \mu(dy \times dy') \\ &\quad + \frac{1}{\sigma^d \nu_d} \iint (y' \otimes y') \mathbb{I}_B\left(\frac{y'}{\sigma}\right) \mu(dy \times dy'). \end{aligned} \quad (52)$$

The first term in (52) vanishes because we may assume that $y \in B_\sigma(0)$ and $\|y' - y\| \leq \eta$. Since the last term in (52) is $\Sigma_\alpha(0, \sigma)$, combining (51) and

(52), we get

$$\begin{aligned}
\Sigma_\beta(0, \sigma) - \Sigma_\alpha(0, \sigma) &= \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} (y - y') \otimes y \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} y' \otimes (y - y') \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{B_{\sigma+\eta}(0)} (y' \otimes y') \alpha(dy') - \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} (y' \otimes y') \alpha(dy') \\
&= \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} (y - y') \otimes y \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} y' \otimes (y - y') \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{R_d(\sigma, \sigma+\eta)} (y' \otimes y') \alpha(dy').
\end{aligned} \tag{53}$$

Thus,

$$\begin{aligned}
\|\Sigma_\alpha(0, \sigma) - \Sigma_\beta(0, \sigma)\| &\leq \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} \|y - y'\| \|y\| \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{B_\sigma(0)} \int_{\|y'-y\| \leq \eta} \|y'\| \|y - y'\| \mu(dy \times dy') \\
&+ \frac{1}{\sigma^d \nu_d} \int_{R_d(\sigma, \sigma+\eta)} \|y'\|^2 \beta(dy') \\
&\leq \frac{\eta}{\sigma^d \nu_d} \int \|y\| \alpha(dy) + \frac{\eta}{\sigma^d \nu_d} \int \|y\| \alpha(dy) \\
&+ \frac{1}{\sigma^d \nu_d} \int_{R_d(\sigma, \sigma+\eta)} \|y'\|^2 \beta(dy').
\end{aligned} \tag{54}$$

Since $\alpha \in \mathcal{P}_\infty^{(c, \lambda)}(\mathbb{R}^d)$, we have that

$$\int_{R_d(\sigma, \sigma+\eta)} \|y'\|^2 \alpha(dy') \leq \lambda \int_{R_d(\sigma, \sigma+\eta)} \|y'\|^2 dy' = \lambda s_d(\sigma, \sigma + \eta). \tag{55}$$

The fact that $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$ implies that $\eta \leq 2c$. Using (48) with $a = \sigma$, $b = \sigma + \eta$ and $B = \sigma + 2c$, it follows that

$$s_d(\sigma, \sigma + \eta) \leq \eta \frac{\omega_{d-1}}{d+2} \frac{(\sigma + 2c)^{d+2}}{2c} = \eta \nu_d \frac{d}{d+2} \frac{(\sigma + 2c)^{d+2}}{2c}. \tag{56}$$

From (54), (55), and (56), we can conclude that

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_\alpha(x, \sigma) - \Sigma_\beta(x, \sigma)\| \leq \left(\frac{2c}{\sigma^d \nu_d} + \frac{\lambda d (\sigma + 2c)^{d+2}}{2c(d+2)\sigma^d} \right) W_\infty(\alpha, \beta). \quad (57)$$

as claimed. \square

We now derive a consistency result and precise estimates for the rate of convergence of empirical approximations to multiscale covariance fields. The following result is a W_∞ -counterpart to the theorem by Fournier and Guillin stated above.

Theorem 4 (García-Trillos and Slepčev [6]). *Let $\Omega \subset \mathbb{R}^d$ be a bounded connected open subset with Lipschitz boundary. Let α be a probability measure on Ω with density $f_\alpha: \Omega \rightarrow (0, \infty)$ such that there exists $\lambda \geq 1$ with $\lambda^{-1} \leq f_\alpha(x) \leq \lambda$, for all $x \in \Omega$, and let y, \dots, y_n be i.i.d. random variables with distribution α . Then, there exist constants $c_1, C_1, C_2 > 0$, depending only on Ω and λ , such that for all $n \in \mathbb{N}$ and $p > 1$,*

$$\mathbb{P}(W_\infty(\alpha, \alpha_n) \leq (C_1 + C_2 \sqrt{p}) r_d(n)) \geq 1 - c_1 n^{-p},$$

where $r_2(n) = \frac{\ln(n)^{3/4}}{n^{1/2}}$ and $r_d(n) = \frac{\ln(n)^{1/d}}{n^{1/d}}$, for $d \geq 3$.

Corollary 3 (Consistency for the Truncation Kernel). *Let α be a probability measure on \mathbb{R}^d with density f_α and let Ω_α be the interior of the support of α . Assume that Ω_α is bounded and connected with Lipschitz boundary $\partial\Omega_\alpha$. Furthermore, assume that there exists $\lambda \geq 1$ such that $\lambda^{-1} \leq f_\alpha(z) \leq \lambda$, for all $z \in \Omega_\alpha$. If y_1, \dots, y_n , $n \geq 1$, are i.i.d. random variables with distribution α , then, for any $p > 1$, there are constants $C = C(\Omega_\alpha, \lambda, p) > 0$ and $c_1 = c_1(\Omega_\alpha, \lambda)$ such that*

$$\mathbb{P}\left(\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_\alpha(x, \sigma)\| \leq C r_d(n)\right) \geq 1 - c_1 n^{-p}.$$

Proof. We use Theorem 4 and write $C' = C_1 + \sqrt{p} C_2$. Theorem 3 implies that there is a constant $C'' = C''(\Omega_\alpha, \lambda) > 0$ such that

$$\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_\alpha(x, \sigma)\| \leq C'' W_\infty(\alpha, \alpha_n). \quad (58)$$

Thus,

$$\begin{aligned} \mathbb{P} \left(\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_{\alpha}(x, \sigma)\| \leq C' C'' r_d(n) \right) &\geq \\ &\geq \mathbb{P}(W_{\infty}(\alpha, \alpha_n) \leq C' r_d(n)) \geq 1 - c_1 n^{-p}. \end{aligned} \quad (59)$$

The claim follows by setting $C = C' C''$. \square

Corollary 4. *Let $\sigma > 0$ and $p > 1$. Under the assumptions of Corollary 3, for the truncation kernel, there exist $N = N(\sigma, \Omega_{\alpha}, \lambda) \in \mathbb{N}$ and a constant $A = A(\sigma, \Omega_{\alpha}, \lambda) > 0$ such that*

$$\mathbb{E} \left[\sup_{x \in \mathbb{R}^d} \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_{\alpha}(x, \sigma)\| \right] \leq A r_d(n),$$

for all $n \geq N$.

Proof. We apply to $W_{\infty}(\alpha, \alpha_n)$ the identity $\mathbb{E}(Z) = \int_0^{\infty} \mathbb{P}(Z > t) dt$ that is valid for any non-negative random variable Z with finite first moment. Since $W_{\infty}(\alpha, \alpha_n) \leq D = \text{diam}(\Omega_{\alpha})$, we get

$$\mathbb{E}[W_{\infty}(\alpha, \alpha_n)] = \int_0^D \mathbb{P}(W_{\infty}(\alpha, \alpha_n) > t) dt. \quad (60)$$

Theorem 4 implies that

$$\mathbb{P}(W_{\infty}(\alpha, \alpha_n) > (C_1 + \sqrt{p} C_2) r_d(n)) \leq n^{-p}. \quad (61)$$

Let $t_0 = \min\{D, (C_1 + \sqrt{p} C_2) r_d(n)\}$. From (60) and (61),

$$\begin{aligned} \mathbb{E}[W_{\infty}(\alpha, \alpha_n)] &= \int_0^{t_0} \mathbb{P}(W_{\infty}(\alpha, \alpha_n) > t) dt + \int_{t_0}^D \mathbb{P}(W_{\infty}(\alpha, \alpha_n) > t) dt \\ &\leq t_0 + D \mathbb{P}(W_{\infty}(\alpha, \alpha_n) > (C_1 + \sqrt{p} C_2) r_d(n)) \\ &\leq (C_1 + \sqrt{p} C_2) r_d(n) + D n^{-p}. \end{aligned} \quad (62)$$

Fixing p , say $p = 2$, for n sufficiently large, the dominant term on this last expression is the one involving $r_d(n)$. Thus, the claim follows from (62) and Theorem 3 applied to α and $\beta = \alpha_n$. \square

Remark 5. We carry out an experiment to test the convergence rates obtained in Corollary 4. We consider the probability measure α supported on the unit circle $\mathbb{S}^1 \subset \mathbb{R}^2$ induced by the normalized arc length element $(2\pi)^{-1}ds$. In this case, for the truncation kernel, Σ_α was calculated explicitly in Example 4. We consider sets of i.i.d. samples of size n , $10 \leq n \leq 10^6$. For each n , thirty sets of samples are taken. For each such set, we compute Σ_{α_n} and estimate the “error” as $\max \|\Sigma_{\alpha_n}(x, \sigma) - \Sigma_\alpha(x, \sigma)\|$, for $\sigma = 0.6$, where the maximum is taken over gridpoints on a 24×24 grid on the square $[-1.5, 1.5] \times [-1.5, 1.5]$. We let ε_n be the average error over all thirty sets of samples. Figure 6 shows a plot (in blue) of ε_n in log-log scale. To compare ε_n with the predicted rates, we use a least-squares fit, in log-log scale, of the form $\varepsilon = Cr_2(n) = C \frac{\ln(n)^{3/4}}{n^{1/2}}$, also shown in Figure 6 (in red). The discrepancy between the predicted and observed rates suggests that Corollary 4 might not be optimal. A curve of the form $\varepsilon = Cn^{-1/2}$, shown in green, produces a tighter fit to the data, suggesting that the optimal bound might be $O(n^{-1/2})$.

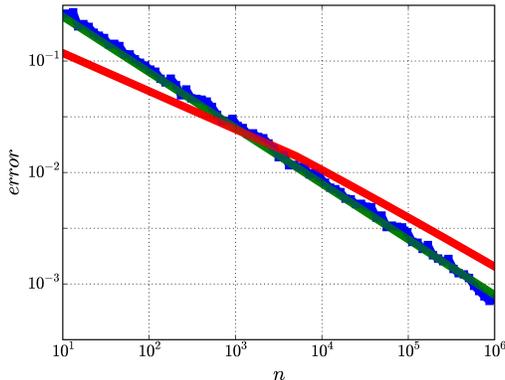


Figure 6: Log-log plots of experimental error rates (in blue) for empirical covariance fields, rates predicted by Corollary 4 (in red), and a least-squares fit of order $n^{-1/2}$ (in green).

4.3. General Kernels

We conclude the discussion of convergence of empirical CTFs with a pointwise central limit theorem (CLT) that holds for kernels in the full generality of Definition 2. One may think of it as a CLT for each entry of the matrix $\Sigma_\alpha(x, \sigma)$. If e_1, \dots, e_d is an orthonormal basis of \mathbb{R}^d , the (i, j) -entry of

the covariance matrix in this coordinate system is given by $\Sigma_\alpha(x, \sigma)(e_i, e_j)$, the bilinear form $\Sigma_\alpha(x, \sigma)$ evaluated at (e_i, e_j) . In matrix notation, this is the same as $\langle e_i, \Sigma_\alpha(x, \sigma)e_j \rangle$. More generally, for fixed $u, v, x \in \mathbb{R}^d$ and $\sigma > 0$, we consider

$$\Sigma_\alpha(x, \sigma)(u, v) = \int (y - x) \otimes (y - x)(u, v) K(x, y, \sigma) \alpha(dy). \quad (63)$$

Consider the random variable

$$z_{uv}(y) = (y - x) \otimes (y - x)(u, v) K(x, y, \sigma), \quad (64)$$

where y has distribution α . Clearly,

$$\mathbb{E}[z_{uv}] = \Sigma_\alpha(x, \sigma)(u, v). \quad (65)$$

Theorem 5 (Central Limit). *If f is as in Definition 2, then z_{uv} has finite variance σ_{uv}^2 . Moreover, if z_1, \dots, z_n are i.i.d. random variables with the same distribution as z_{uv} , then*

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n z_i - \Sigma_\alpha(x, \sigma)(u, v) \right) \xrightarrow{d} \mathcal{N}(0, \sigma_{uv}^2),$$

as $n \rightarrow \infty$, where convergence is in distribution and $\mathcal{N}(0, \sigma_{uv}^2)$ is normally distributed with mean zero and variance σ_{uv}^2 .

Proof. We show that z_{uv} has finite second moment. From (64) and (4),

$$\begin{aligned} \int z_{uv}^2 \alpha(dy) &\leq \|u\| \|v\| \int \|y - x\|^4 K^2(x, y, \sigma) \alpha(dy) \\ &\leq \frac{\sigma^4 \|u\| \|v\|}{C_d^2(\sigma)} \int \frac{\|y - x\|^4}{\sigma^4} f^2 \left(\frac{\|y - x\|^2}{\sigma^2} \right) \alpha(dy) \\ &\leq \frac{\sigma^4 \|u\| \|v\|}{C_d^2(\sigma)} C^2. \end{aligned} \quad (66)$$

The last inequality follows from condition (c) in Definition 2 that ensures that $r^2 f^2(r) < C^2$, for any $r > 0$. The theorem now follows from a direct application of the classical CLT. \square

Remark 6. Note that (66) implies that if $\|u\| = \|v\| = 1$, then

$$\sigma_{uv}^2 = \int z_{uv}^2 \alpha(dy) - (\mathbb{E}[z_{uv}])^2 \leq \frac{C^2 \sigma^4}{C_d^2(\sigma)}, \quad (67)$$

giving a uniform bound on the variance of z_{uv} over $x \in \mathbb{R}^d$ and $u, v \in \mathbb{S}^{d-1}$.

5. Multiscale Fréchet Functions

The mean of a random vector $y \in \mathbb{R}^d$ is a simple and yet oftentimes informative, “one-element” summary of the distribution of y . If y has finite second moment and is distributed according to the probability measure α , then the mean may be characterized more geometrically as the unique minimizer of the Fréchet function

$$F_\alpha(x) = \mathbb{E} [\|y - x\|^2] = \int \|y - x\|^2 \alpha(dy), \quad (68)$$

which measures the spread of y about $x \in \mathbb{R}^d$. The mean, however, is not as effective for complex distributions of practical interest such as multimodal distributions or those supported in nonlinear subspaces. In this section, we introduce a multiscale analogue of the Fréchet function that is rich in information about the shape of the distribution of y . At each fixed scale, the local minima of the function may be viewed as localized analogues of the mean, as illustrated in examples below. However, instead of just focusing on the local extrema, we take the view that it is more informative to investigate the behavior of the full multiscale Fréchet function, as this lets us uncover more information about the distribution of y .

Definition 5. Let $f: [0, \infty) \rightarrow \mathbb{R}$ be as in Definition 2 with associated kernel $K: \mathbb{R}^d \times \mathbb{R}^d \times (0, \infty) \rightarrow \mathbb{R}$. The *multiscale Fréchet function* $V_\alpha: \mathbb{R}^d \times (0, \infty) \rightarrow \mathbb{R}$ is defined as

$$V_\alpha(x, \sigma) := \int \|y - x\|^2 K(x, y, \sigma) \alpha(dy).$$

Proposition 3. For each $\sigma > 0$, the multiscale Fréchet function satisfies

$$V_\alpha(x, \sigma) = \text{tr } \Sigma_\alpha(x, \sigma).$$

Proof. Let $\{e_1, \dots, e_d\} \subset \mathbb{R}^d$ be an orthonormal basis. Then,

$$\|y - x\|^2 = \sum_{i=1}^d \langle y - x, e_i \rangle^2 = \sum_{i=1}^d (y - x) \otimes (y - x)(e_i, e_i). \quad (69)$$

Hence,

$$\begin{aligned}
V_\alpha(x, \sigma) &= \sum_{i=1}^d \int (y-x) \otimes (y-x)(e_i, e_i) K(x, y, \sigma) \alpha(dy) \\
&= \sum_{i=1}^d \left(\int (y-x) \otimes (y-x) K(x, y, \sigma) \alpha(dy) \right) (e_i, e_i) \quad (70) \\
&= \sum_{i=1}^d \Sigma_\alpha(x, \sigma)(e_i, e_i) = \text{tr} \Sigma_\alpha(x, \sigma),
\end{aligned}$$

as claimed. \square

Corollary 5 (Stability). *Let $p \in [1, \infty)$, $\ell = \max\{p, 2p/(p-1)\}$, and suppose that $\alpha, \beta \in \mathcal{P}_\ell^c(\mathbb{R}^d)$. If f is differentiable and $\sqrt{r} |f'(r)| \leq A$, $\forall r \in [0, \infty)$ and some constant $A > 0$, then*

$$\sup_{x \in \mathbb{R}^d} |V_\alpha(x, \sigma) - V_\beta(x, \sigma)| \leq \frac{2dc}{C_d(\sigma)} \left(1 + \frac{Ac}{\sigma} \right) W_p(\alpha, \beta),$$

for any $\sigma > 0$.

Proof. The result follows from Proposition 3, Theorem 1 and the fact that for any $d \times d$ matrix X , $|\text{tr} X| \leq d \|X\|$, where $\|X\|$ is the Frobenius norm of X . \square

Similarly, Corollary 1 and Proposition 3 yield the following consistency result for multiscale Fréchet functions.

Corollary 6 (Consistency). *Suppose that $\alpha \in \mathcal{P}_3(\mathbb{R}^d)$. Let y_1, \dots, y_n be i.i.d. random variables with distribution α and K a multiscale kernel as in Theorem 1. Then, for fixed $\sigma > 0$,*

$$\sup_{x \in \mathbb{R}^d} |V_\alpha(x, \sigma) - V_{\alpha_n}(x, \sigma)| \xrightarrow{n \uparrow \infty} 0$$

almost surely.

The following result about convergence of multiscale Fréchet functions are immediate consequences of Corollary 2 and Corollary 5.

Corollary 7. *Let $\sigma, c > 0$ and f differentiable satisfying $\sqrt{r}|f'(r)| \leq A$, $\forall r \in [0, \infty)$ and some $A > 0$. If $\alpha \in \mathcal{P}_4^c(\mathbb{R}^d)$ and y_1, \dots, y_n , $n \geq 1$, are i.i.d. random variables with distribution α , then*

$$\mathbb{E} \left[\sup_{x \in \mathbb{R}^d} |V_\alpha(x, \sigma) - V_{\alpha_n}(x, \sigma)| \right] \leq \frac{2d c^3}{C_d(\sigma)} \left(1 + \frac{Ac}{\sigma} \right) \beta_2 \cdot \begin{cases} n^{-\frac{1}{3}} + n^{-\frac{1}{2}} & \text{if } d = 1, 2, 3; \\ n^{-\frac{1}{3}} + n^{-\frac{1}{2}} \log(1+n) & \text{if } d = 4; \\ n^{-\frac{1}{3}} + n^{-\frac{2}{d}} & \text{if } d > 4. \end{cases}$$

Remark 7. Analogous stability and consistency results for the truncation kernel follow from Theorem 3, Corollary 3 and Corollary 4.

For more general kernels, the following pointwise central limit theorem holds. For fixed $x \in \mathbb{R}^d$ and $\sigma > 0$, let

$$t(y) = \|y - x\|^2 K(x, y, \sigma), \quad (71)$$

whose expected value is $\mathbb{E}[t] = V_\alpha(x, \sigma)$. As in Theorem 5, the variance of t is finite and denoted σ_t^2 .

Theorem 6 (Central Limit). *Let f be as in Definition 2. If t_1, \dots, t_n are i.i.d. random variables with the same distribution as t , then*

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n t_i - V_\alpha(x, \sigma) \right) \xrightarrow{d} \mathcal{N}(0, \sigma_t^2),$$

as $n \rightarrow \infty$, where convergence is in distribution and $\mathcal{N}(0, \sigma_t^2)$ is normally distributed with mean zero and variance σ_t^2 .

Multiscale Fréchet functions not only give stable representations of probability measures, but any probability measure α may be fully recovered from its multiscale Fréchet function associated with the Gaussian kernel, as the following result shows.

Proposition 4. *Let $\sigma > 0$ be fixed. Any probability measure α is completely determined by the Fréchet function $V_\alpha(\cdot, \sigma)$ associated with the Gaussian kernel at scale σ .*

Proof. Let $h_\sigma: \mathbb{R}^d \rightarrow \mathbb{R}$ be given by

$$h_\sigma(x) = \frac{\|x\|^2}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|x\|^2}{2\sigma^2}\right). \quad (72)$$

Then, for the Gaussian kernel, we may express the multiscale Fréchet function as the convolution $V_\alpha(x, \sigma) = (h_\sigma * \alpha)(x)$. Under Fourier transform, for each fixed $\sigma > 0$, we obtain

$$\widehat{V}_\alpha(\xi, \sigma) = \widehat{h}_\sigma(\xi) \phi_\alpha(-2\pi\xi), \quad (73)$$

where ϕ_α is the characteristic function of α defined as $\phi_\alpha(\xi) = \int e^{i\langle x, \xi \rangle} \alpha(dx)$. Therefore,

$$\phi_\alpha(\xi) = \widehat{V}_\alpha(-\xi/2\pi, \sigma) / \widehat{h}_\sigma(-\xi/2\pi) \quad (74)$$

provided that $\widehat{h}_\sigma(-\xi/2\pi) \neq 0$. A calculation shows that

$$\widehat{h}_\sigma(-\xi/2\pi) = \sigma^2 \left(d - \frac{\sigma^2 \|\xi\|^2}{\pi} \right) \exp\left(-\frac{\sigma^2 \|\xi\|^2}{2\pi}\right), \quad (75)$$

which only vanishes at points ξ on the sphere of radius $\rho_\sigma = \sqrt{\pi d}/\sigma$ about the origin. Thus, (74) implies that we can recover $\phi_\alpha(\xi)$ from $\widehat{V}_\alpha(\cdot, \sigma)$, if $\|\xi\| \neq \sqrt{\pi d}/\sigma$. By continuity, we can recover $\phi_\alpha(\xi)$, for any ξ . The claim now follows from the fact that the characteristic function ϕ_α determines α [23]. \square

The following examples illustrate how information about the shape of data can be extracted from multiscale Fréchet functions.

Example 7. We consider $n = 400$ data points distributed into two clusters of 200 points, each sampled from a Gaussian of variance 0.36 centered at different points. The data points are plotted in blue in Fig. 7(a), which also shows the empirical Fréchet function V_n at scale $\sigma = 3$. The local minima of V_n captures what is perceived as the “centers” of the two clusters at that scale. However, more information about the data distribution can be uncovered from V_n . For example, the local minima may be viewed as attractors of the (negative) gradient field $-\nabla V_n$, indicated by the arrows in the figure. The stable manifold of each attractor, which comprises points that move toward the attractors under the associated flow may be viewed as clusters inferred from the data at that scale. These clusters are delimited by the repellers of the system, which correspond to the local maxima of V_n .

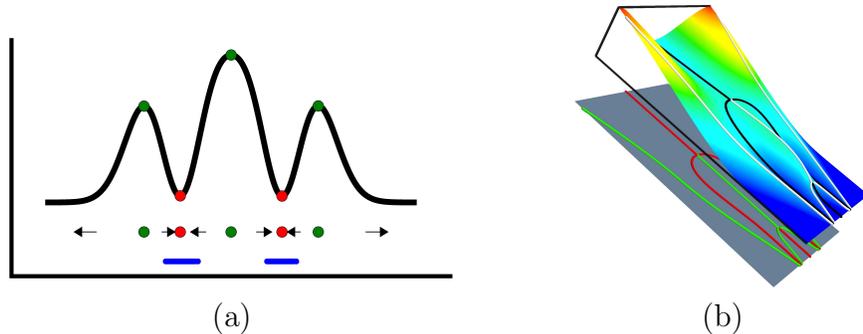


Figure 7: (a) Fréchet function for data on the line (highlighted in blue) computed with the Gaussian kernel at scale $\sigma = 3$; (b) Fréchet function across scales.

Fig. 7(b) shows how V_n varies across scales, highlighting the bifurcation of the attractors (in red) and repellers (in green) as σ changes. In data analysis, such bifurcation diagrams may find several applications. For example, if the data represent the distribution of some phenotypic trait for two species that have evolved from a single group, the multiscale Fréchet function and the associated bifurcation diagram let us create an evolutionary model for the trait from the observed data.

Example 8. Here we consider the dataset in \mathbb{R}^2 shown in panel (a) of Fig. 8. Panels (b)–(h) show the Fréchet function for the Gaussian kernel calculated at increasing scales. The gradient field $-\nabla V_n$ at scale $\sigma = 2.25$ is depicted in panel (a) of Fig. 9 along with the two attractors p_1 and p_2 , and their stable manifolds that were estimated numerically. The stable manifolds may be viewed as estimations at scale $\sigma = 2.25$ of clusters of the underlying probability measure α from which the data was sampled. Panel (b) shows the gradient field and the covariance tensors at the attractors depicted as ellipses with principal radii proportional to the square root of the eigenvalues of the covariance matrix. This may be viewed as a localized analogue of principal component analysis (PCA) that is able to uncover geometry that is not detectable with standard PCA. Analysis of the spectra of $\Sigma_n(p_i, \sigma)$, $i = 1, 2$, suggests that the data is organized around two one-dimensional clusters, whereas standard PCA is not sensitive to the local dimensionality because of the orientation of the clusters.

These examples are intended as proof-of-concept illustrations. Topological and other methods will be explored in forthcoming work for extraction

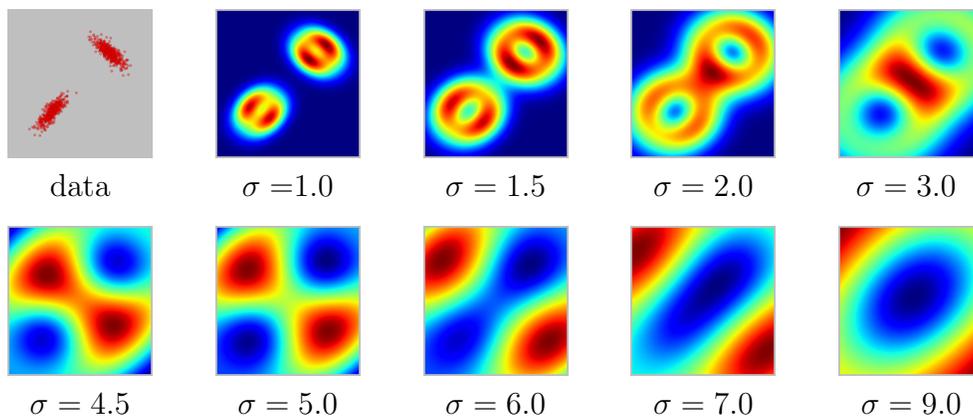


Figure 8: Heat maps of the multiscale Fréchet function for 2D data at increasing scales computed with the Gaussian kernel.

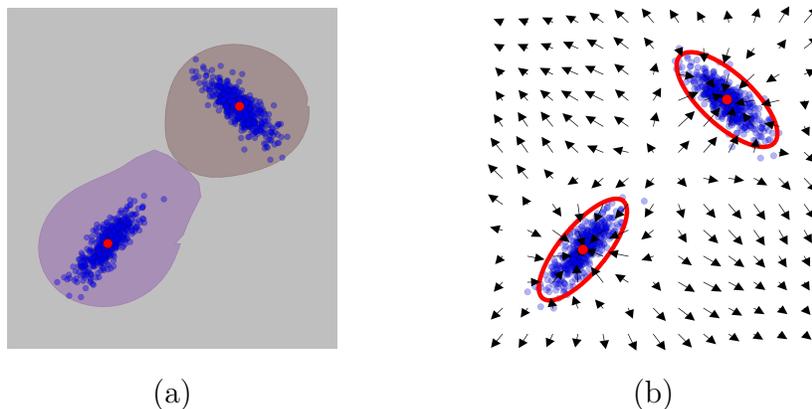


Figure 9: (a) 2D data, attractors and their stable manifolds at a fixed scale ($\sigma = 2.25$); (b) gradient vector field and covariance tensors at the attractors.

of structural information from V_n .

6. Hierarchical Manifold Clustering

Clustering is a central theme in pattern analysis with a rich history; cf. [24]. One of the most studied forms of the problem is that of partitioning a dataset into various subsets if there is some form of spatial separation of the data into subgroups. Motivated by problems in such areas as computer vision

and video analysis, cf. [25], there has been a growing interest in clustering data that are organized as a finite union of possibly intersecting subspaces that have some special geometric structure. As illustrated in Fig. 1, the data may consist of noisy samples from an arrangement of (affine) linear subspaces of a Euclidean space such as a collection of lines in a plane, or an arrangement of lines and planes in \mathbb{R}^3 . More generally, the clusters may comprise a finite collection of possibly non-linear, smooth submanifolds of a Euclidean space that intersect transversely. Here we propose an approach to manifold clustering based on CTFs. The basic idea is to use covariance fields to incorporate directional information at each data point. Formally, this is achieved via a section of the tensor bundle $\mathbb{R}^d \times (\mathbb{R}^d \otimes \mathbb{R}^d)$ over \mathbb{R}^d , as follows. Given a probability measure α and a multiscale kernel, let $\Sigma_\alpha(x, \sigma)$ be the associated CTF. For each $\sigma > 0$, consider the section $\iota_{\alpha; \sigma}: \mathbb{R}^d \rightarrow \mathbb{R}^d \times (\mathbb{R}^d \otimes \mathbb{R}^d)$ given by $x \mapsto (x, \Sigma_\alpha(x, \sigma))$. On the total space of the tensor bundle, define the metric

$$\|(x, \Sigma) - (x', \Sigma')\|_\gamma = (\|\Sigma - \Sigma'\|^2 + \gamma^2 \|x - x'\|^2)^{1/2}, \quad (76)$$

where $x, x' \in \mathbb{R}^d$, $\Sigma, \Sigma' \in \mathbb{R}^d \otimes \mathbb{R}^d$, and $\gamma \geq 0$ is a parameter that balances the contributions of the spatial and tensor components. Note that $\|\cdot\|_0$ only defines a pseudo-metric since $\|\cdot\|_0$ disregards “horizontal” distances.

For any subset $X \subseteq \mathbb{R}^d$, we denote by $\mathbb{X}_{\alpha; \gamma, \sigma}$ the metric space $(X, d_{\alpha; \gamma, \sigma})$, where

$$d_{\alpha; \gamma, \sigma}(x, x') = \|\iota_{\alpha; \sigma}(x) - \iota_{\alpha; \sigma}(x')\|_\gamma. \quad (77)$$

For a dataset $A = \{a_1, \dots, a_n\} \subset \mathbb{R}^d$, the proposed clustering method is based on the single-linkage method [26] applied to the finite metric space $\mathbb{A}_{\alpha_n; \gamma, \sigma}$ associated with the empirical measure $\alpha_n = n^{-1} \sum_{i=1}^n \delta_{a_i}$. Equivalently, clustering is based on the $n \times n$ affinity matrix D whose (i, j) -entry is

$$d_{ij} = d_{\alpha_n; \gamma, \sigma}(a_i, a_j). \quad (78)$$

Recall that single linkage on a finite metric space $\mathbb{A} = (A, d_A)$ starts from n clusters, each a singleton $\{a_i\}$, $1 \leq i \leq n$, sequentially merging the closest clusters until all data points coalesce into a single cluster. Closeness of two clusters, say $A_1, A_2 \subset A$, is measured by the inter-cluster distance

$$d_{sl}(A_1, A_2) = \min_{a \in A_1, a' \in A_2} d_A(a, a'), \quad (79)$$

We choose single linkage because it yields stable dendrograms, as expounded below, under assumptions on the probability measure from which the data is sampled that are not very restrictive. Combined with our stability and consistency results for covariance fields, this guarantees that the manifold clustering method is stable at all stages.

6.1. Dendrogram Stability

We denote a metric space by $\mathbb{X} = (X, d_X)$. An ultrametric space is a pair (X, u_X) , where $u_X: X \times X \rightarrow \mathbb{R}^+$ is a metric on X that satisfies the *strong triangle inequality*

$$u_X(x, x') \leq \max \{u_X(x, x''), u_X(x'', x')\} , \quad (80)$$

for all $x, x', x'' \in X$. Any such function u_X is called an *ultrametric* on X .

As proved in [20], dendrograms over a finite set X are in structure-preserving, bijective correspondence with ultrametries on X . In this formulation, a hierarchical clustering method can be regarded as a map $\mathcal{H}: \mathcal{M} \rightarrow \mathcal{U}$ from finite metric spaces into finite ultrametric spaces. Henceforward, \mathcal{H} will denote the map given by single linkage hierarchical clustering. It is known [20] that if $\mathbb{X} = (X, d_X) \in \mathcal{M}$, then $\mathcal{H}(\mathbb{X}) = (X, u_X)$ is given by

$$u_X(x, x') = \min_{x=x_0, \dots, x_r=x'} \max_i d_X(x_i, x_{i+1}) . \quad (81)$$

The minimum above is taken over all finite ordered sequences x_0, x_1, \dots, x_r of points in X such that $x_0 = x$ and $x_r = x'$. If $x, x' \in X$, then $u_X(x, x')$ may be interpreted as the dendrogram level at which the clusters containing x and x' first merge. This is known as the *cophenetic* distance between x and x' .

The main goal of this section is to formulate and prove stability of the map $\mathcal{M} \ni \mathbb{X} \mapsto \mathcal{H}(\mathbb{X}) \in \mathcal{U}$. The question of stability of single linkage clustering can be approached using ideas related to the Gromov-Hausdorff distance [27], as follows. A correspondence R between two sets X and Y is a subset of $X \times Y$ such that $\pi_1(R) = X$ and $\pi_2(R) = Y$, where π_1 and π_2 denote projections onto the first and second factors. Given X and Y , we denote by $\mathcal{R}(X, Y)$ the set of all correspondences between X and Y .

Definition 6. Let \mathbb{X} and \mathbb{Y} be compact metric spaces.

(i) The *distortion* of a correspondence R between \mathbb{X} and \mathbb{Y} is defined by

$$\text{dis}(R; \mathbb{X}, \mathbb{Y}) := \max_{(x,y),(x',y') \in R} |d_X(x, x') - d_Y(y, y')|.$$

(ii) The *Gromov-Hausdorff* distance between \mathbb{X} and \mathbb{Y} is given by

$$d_{GH}(\mathbb{X}, \mathbb{Y}) := \frac{1}{2} \inf_R \text{dis}(R; \mathbb{X}, \mathbb{Y}),$$

where the infimum is taken over all correspondences between \mathbb{X} and \mathbb{Y} .

The following stability result is a generalization of [20, Proposition 26].

Proposition 5. *For any $\mathbb{X}, \mathbb{Y} \in \mathcal{M}$ and any correspondence $R \in \mathcal{R}(X, Y)$,*

$$\text{dis}(R; \mathcal{H}(\mathbb{X}), \mathcal{H}(\mathbb{Y})) \leq \text{dis}(R; \mathbb{X}, \mathbb{Y}).$$

As a consequence, $d_{GH}(\mathcal{H}(\mathbb{X}), \mathcal{H}(\mathbb{Y})) \leq d_{GH}(\mathbb{X}, \mathbb{Y})$.

Remark 8. The claim of the proposition may be written, equivalently, as follows. If u_X and u_Y denote the ultrametrics produced by single linkage hierarchical clustering on \mathbb{X} and \mathbb{Y} , then

$$|u_X(x, x') - u_Y(y, y')| \leq \max_{(x,y),(x',y') \in R} |d_X(x, x') - d_Y(y, y')|, \quad (82)$$

for any correspondence R between X and Y and all $(x, y), (x', y') \in R$.

Proof of Proposition 5. We prove (82). Given a correspondence $R \in \mathcal{R}(X, Y)$ and $(x, y), (x', y') \in R$, let $x = x_0, x_1, \dots, x_n = x'$ in X be such that $\max_i d_X(x_i, x_{i+1}) = u_X(x, x')$. Let $y_0 = y, y_n = y'$ and choose $y_1, \dots, y_{n-1} \in Y$ such that $(x_i, y_i) \in R$ for all $i = 1, \dots, n-1$. This is possible since any correspondence R satisfies $\pi_1(R) = X$. Notice that

$$\begin{aligned} u_Y(y, y') &\leq \max_i d_Y(y_i, y_{i+1}) \\ &\leq \max_i (d_X(x_i, x_{i+1}) + |d_X(x_i, x_{i+1}) - d_Y(y_i, y_{i+1})|) \\ &\leq \max_i d_X(x_i, x_{i+1}) + \max_{(x,y),(x',y') \in R} |d_X(x, x') - d_Y(y, y')| \\ &= u_X(x, x') + \max_{(x,y),(x',y') \in R} |d_X(x, x') - d_Y(y, y')|. \end{aligned} \quad (83)$$

The claim follows since (83) also holds if we reverse the roles of X and Y . \square

Lemma 2. *Let $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$ and $\sigma > 0$. If a kernel satisfies the conditions of Lemma 1, then*

$$\sup_{(a,b) \in R_\mu} \|\Sigma_\alpha(a, \sigma) - \Sigma_\beta(b, \sigma)\| \leq \frac{8c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma}\right) \sup_{(a,b) \in R_\mu} \|a - b\|,$$

for any coupling $\mu \in \Gamma(\alpha, \beta)$, where $R_\mu := \text{supp}[\mu]$.

Proof. Set $\zeta = \sup_{(a,b) \in \text{supp}[\mu]} \|a - b\|$. Let $\mu \in \Gamma(\alpha, \beta)$ and $(y, y'), (a, b) \in \text{supp}[\mu]$. Since $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$, we have that $\|y - a\| \leq 2c$ and $\|y' - b\| \leq 2c$. Using the notation in the proof of Theorem 1, setting $z_1 = y - a$ and $z_2 = y' - b$, (45) yields

$$\begin{aligned} \|Q_\sigma(y - a) - Q_\sigma(y' - b)\| &\leq (\|y - y'\| + \|a - b\|) \frac{4c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma}\right) \\ &\leq \frac{8c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma}\right) \zeta, \end{aligned} \quad (84)$$

where in the last inequality above we used $\|y - y'\| \leq \zeta$ and $\|a - b\| \leq \zeta$. Since

$$\|\Sigma_\alpha(a, \sigma) - \Sigma_\beta(b, \sigma)\| \leq \iint \|Q_\sigma(y - a) - Q_\sigma(y' - b)\| \mu(dy \times dy'), \quad (85)$$

the lemma follows. \square

Lemma 3 (Lemma 2.2 of [28]). *Let $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$ for some $c > 0$. Then, for any coupling $\mu \in \Gamma(\alpha, \beta)$, $R_\mu = \text{supp}[\mu]$ gives a correspondence between $A = \text{supp}[\alpha]$ and $B = \text{supp}[\beta]$.*

Theorem 7. *Let $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$, $A = \text{supp}[\alpha]$, $B = \text{supp}[\beta]$, $\sigma > 0$ and $\gamma \geq 0$. Then, for any kernel satisfying the conditions of Lemma 1,*

$$d_{GH}((A, d_{\alpha; \sigma, \gamma}), (B, d_{\beta; \sigma, \gamma})) \leq \left(\frac{8c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma}\right) + \gamma\right) W_\infty(\alpha, \beta).$$

Proof. Let $\mu \in \Gamma(\alpha, \beta)$ be a coupling that realizes $W_\infty(\alpha, \beta)$. By Lemma 3,

$R_\mu = \text{supp} [\mu]$ is a correspondence between A and B . Thus,

$$\begin{aligned}
d_{GH}(\mathbb{A}_{\alpha;\sigma,\gamma}, \mathbb{B}_{\beta;\sigma,\gamma}) &\leq \frac{1}{2} \text{dis}(R_\mu; A, B) \\
&= \frac{1}{2} \sup_{(a,b),(a',b') \in R_\mu} |d_{\alpha;\sigma,\gamma}(a, a') - d_{\beta;\sigma,\gamma}(b, b')| \\
&\leq \frac{1}{2} \sup_{(a,b),(a',b') \in R_\mu} \left(\|\Sigma_\alpha(a, \sigma) - \Sigma_\beta(b, \sigma)\| + \right. \\
&\quad \left. + \|\Sigma_\alpha(a', \sigma) - \Sigma_\beta(b', \sigma)\| + \gamma \|a - b\| + \gamma \|a' - b'\| \right) \\
&\leq \sup_{(a,b) \in R_\mu} \|\Sigma_\alpha(a, \sigma) - \Sigma_\beta(b, \sigma)\| + \gamma \sup_{(a,b) \in R_\mu} \|a - b\| \\
&\leq \left(\frac{8c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma} \right) + \gamma \right) \sup_{(a,b) \in R_\mu} \|a - b\|,
\end{aligned} \tag{86}$$

where the last step follows from Lemma 2. The conclusion follows since $W_\infty(\alpha, \beta) = \sup_{(a,b) \in R_\mu} \|a - b\|$. \square

Combining Proposition 5 and Theorem 7, we obtain:

Corollary 8 (Stability of Manifold Clustering). *Let $\alpha, \beta \in \mathcal{P}_\infty^c(\mathbb{R}^d)$ be probability measures with finite support, $A = \text{supp} [\alpha]$, $B = \text{supp} [\beta]$, $\sigma > 0$ and $\gamma \geq 0$. Then,*

$$d_{GH}(\mathcal{H}(\mathbb{A}_{\alpha;\gamma,\sigma}), \mathcal{H}(\mathbb{B}_{\alpha;\gamma,\sigma})) \leq \frac{8c}{C_d(\sigma)} \left(1 + \frac{2Ac}{\sigma} \right) W_\infty(\alpha, \beta).$$

6.2. Examples and Applications

Let $X = \{x_1, \dots, x_n\}$ be a dataset in \mathbb{R}^d . For $\gamma, \sigma > 0$, we apply the single linkage method to the metric space $\mathbb{X}_{\alpha_n;\gamma,\sigma} = (X, d_{\alpha_n;\gamma,\sigma})$, where α_n is the empirical measure associated with X and $d_{\alpha_n;\gamma,\sigma}$ is the distance defined in (77). The ultrametric associated with $\mathcal{H}(\mathbb{X}_{\alpha_n;\gamma,\sigma})$ is abbreviated $u_{\alpha_n;\gamma,\sigma}$.

In this setting, analyzing informative dendrogram cutoff levels often is an important task, which can be approached in different ways, depending on the nature of the problem. For example, a cutoff level h may be based on a pre-assigned number of clusters, be learned from training data, or be more exploratory. We give examples that illustrate all three viewpoints.

Example 9 (Lines and planes). In this experiment we consider the unlabeled point cloud in Fig. 1(a) that represents an arrangement of two parallel planes

and two lines that intersect the planes transversely. Each plane contains 225 points on a uniform grid and each line contains 30 equally spaced points. Cutting the dendrogram at four clusters, our method finds the four affine linear subspaces accurately with the Gaussian kernel at $\sigma = 0.6$. In this case, it is important to choose $\gamma \neq 0$ since the spatial component of (78) is needed to discriminate the parallel planes. In Fig. 1(a), the points are colored according to cluster membership. In this case, the covariance tensors at data points on the planes that are away from the cluster intersections have two dominating eigenvalues, whereas for points on the lines they have only one such eigenvalue. Thus, an analysis of the spectrum of the covariance tensors at data points let us infer the dimension of each cluster.

Example 10 (A Line Arrangement). In this example, the point-cloud data represents three intersecting lines in \mathbb{R}^2 , as shown in Fig. 10(a). Each line segment is sampled at 200 equally spaced points. Since the slopes of the lines are different, we expect the covariance matrices to be able to cluster the points without the aid of additional spatial information. Thus, we set $\gamma = 0$ in (78) and $\sigma = 0.4$. The number of clusters was set to six to test the ability of the algorithm to detect not only the lines, but also the three intersections. Fig. 10(b) shows the single-linkage dendrogram, highlighting each of the six clusters. The data points are colored according to cluster membership. As

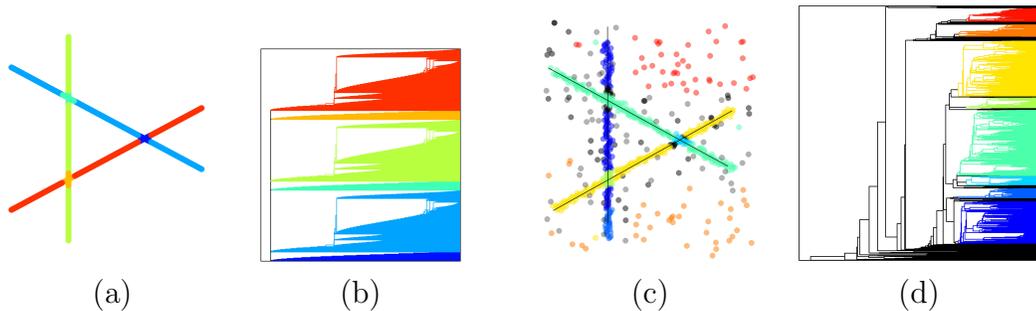


Figure 10: (a) an arrangement of three lines and (b) clustering dendrogram; (c) noisy lines with outliers and (d) clustering dendrogram.

expected, well delineated clusters are detected away from the intersection points because the covariance matrices are highly anisotropic with principal axes that align well with the corresponding line segments. Although the covariance matrices are not as anisotropic near the intersection points, there

are enough differences in their behavior near the three intersection loci for the algorithm to be able to place them into different clusters.

The next two examples are of a more exploratory nature in that dendrogram cutoff was chosen through experimentation with the data.

Example 11 (Noisy Lines with Outliers). This is a noisy version of Example 10, as shown in Fig. 10(c). As before, each line is represented by 200 points, but we have added Gaussian noise of width 0.015 to the data, as well as 180 outliers sampled from the uniform distribution on a rectangle containing the lines. Because of the nature of the data, the number of clusters was set to $m = 80$ so that the three main clusters did not get merged because of the outliers. The figure also shows a line fitted to each of the three largest clusters using principal component analysis. The method was able to sharply recover the three lines, even in the presence of noise and outliers. The majority of the 80 clusters are singletons of outliers and these are colored black in the figure. We remark that the choice of $\sigma = 0.51$ is crucial when dealing with data contaminated by noise. In this case, it was also important to set $\gamma \neq 0$ to better cope with noise.

Example 12 (Floor cracks). We apply the clustering method to segmentation of two images of concrete floor cracks. Panel (a) of Fig. 11 shows the original images, whereas panel (b) shows binary images obtained from an edge detection algorithm. We cluster the foreground pixels of the binary images. As in Example 11, it is important to allow a fairly large number of clusters so that the clusters that detect the main cracks do not get merged because of the noisy pixels. Panels (c) and (d) show the outputs (not to scale) of the clustering algorithm.

To further test the ability of the method to cluster intersecting manifolds, we experimented with synthetic data comprising multiple arrangements such as the intersecting lines in Fig. 1(b).

Example 13. We consider three synthetic datasets of point clouds representing random arrangements of: (i) three line segments in \mathbb{R}^2 ; (ii) four curves in \mathbb{R}^2 that are either line segments or arcs of parabolas; and (iii) three patches of planes in \mathbb{R}^3 . Each of these datasets contains a total of 250 point clouds, 50 used for training the algorithm and 200 test samples. The points in each point cloud are labeled to allow quantification of the accuracy of the output of the algorithm. Fig. 12 shows a few samples from each of these datasets. Parameter values that optimize classification performance are learned from

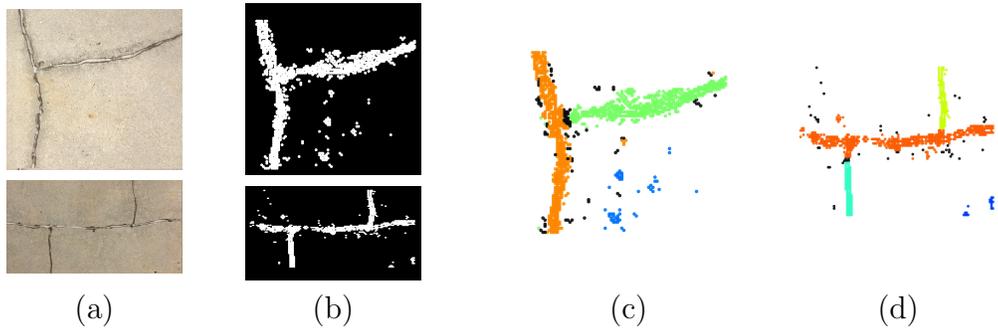


Figure 11: (a) original and (b) processed images of floor cracks; (c) and (d) show clustering based on CTFs.

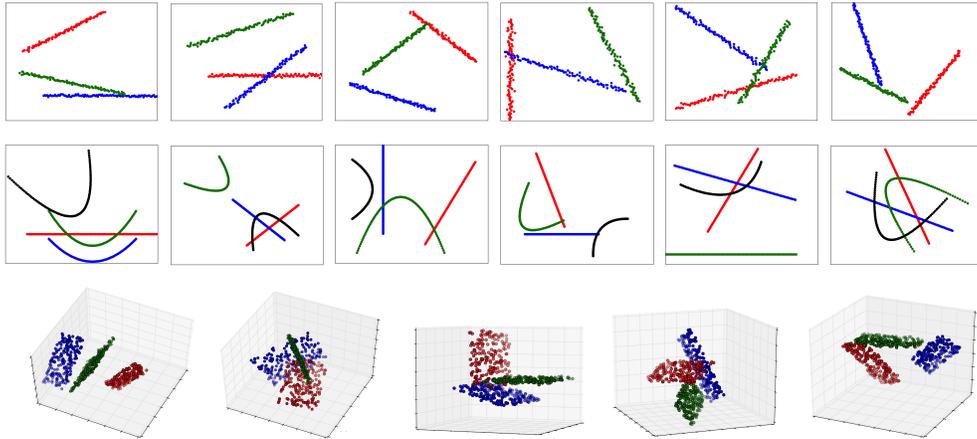


Figure 12: Random arrangements of line segments (row 1), segments of lines and parabolas (row 2), and plane patches (row 3).

the training samples. Note, however, that even though the number k of clusters is known, specifying a height h that yields precisely k clusters may yield undesirable results. For example, important clusters representing different components of an arrangement of manifolds may get merged due to the presence of outliers or the behavior near the intersections. Thus, it is often preferable to choose a lower cutoff level before this phenomenon occurs at the expense of getting a larger number ℓ of clusters. In such situations, we select the largest k clusters and assign each point in the remaining $(\ell - k)$ clusters to the closest of the top k clusters. Experiments indicate that a good

baseline for the cutoff level h is the mean cophenetic distance, which for a point cloud $\{x_1, \dots, x_n\} \subset \mathbb{R}^d$ is given by

$$h_0 = \frac{2}{n(n-1)} \sum_{i < j} u_{\alpha_n; \gamma, \sigma}(x_i, x_j). \quad (87)$$

In the learning phase, we typically search for h in a neighborhood of h_0 whose width is determined by the variance of the distribution of the cophenetic distances.

With the learned parameter values, the algorithm performs well in all three cases. For each point cloud we count the number of misclassified points and calculate the average error (AE) and the mean error (ME) rates over all test samples obtaining:

- (i) arrangements of lines: 9.59% (AE) and 4.17% (ME);
- (ii) lines and parabolas: 9.93% (AE) and 3.38% (ME);
- (iii) arrangements of planes: 7.00% (AE) and 2.42% (ME).

As expected, a closer inspection of the results reveals that most of the errors occur at points near the intersections of the clusters, where the covariance tensors are not as informative for clustering purposes.

7. Concluding Remarks

We introduced the notion of multiscale covariance tensor fields associated with Euclidean random variables and developed a framework for the systematic study of the shape of data using localized covariance tensors. We investigated foundational questions such as stability and consistency of multiscale CTFs, provided illustrations of how CTFs let us uncover geometry underlying data, and applied the methods to manifold clustering. We also introduced multiscale Fréchet functions, which are scalar fields derived from CTFs that fully capture the distributions of random vectors. Multiscale Fréchet functions are particularly well suited for extension of the methods of this paper to non-Euclidean random variables, a problem that is receiving ever increasing attention in data science. In this setting, the goal is to devise methods that can cope with random variables taking values in spaces such as Riemannian manifolds and more general metric spaces. Unless restrictive assumptions are imposed on the sample space and the distributions, CTFs may

be difficult to define in this nonlinear realm. In contrast, the Fréchet function formulation can be easily extended to metric spaces supporting a diffusion kernel [29]. In forthcoming work, we will investigate theoretical and computational aspects of such extensions, including the accessibility of information residing in multiscale Fréchet functions, a problem that poses computational challenges even in the case of high-dimensional Euclidean random variables.

In this paper, we only considered radial basis kernels; however, many results extend easily to more general kernels. We emphasized the multiscale formulation largely because of the questions that motivated this work. Nevertheless, the majority of the results apply to kernels that are not scale dependent.

Covariance tensor fields also suggest ways of formalizing the notion of shape of Euclidean data and probability measures. For example, for a distribution α with the property that the covariance tensor field $\Sigma_\alpha(\cdot, \sigma)$ associated with a smooth kernel (such as the Gaussian kernel) is non-singular for every $x \in \mathbb{R}^d$, $\Sigma_\alpha^{-1}(\cdot, \sigma)$ defines a metric tensor with close ties to α . This poses the problem of uncovering relationships between Riemannian metrics derived from CTFs, such as $\Sigma_\alpha^{-1}(\cdot, \sigma)$, and the shape of α .

In a different direction, for a fixed point $x \in \mathbb{R}^d$, an interesting problem is that of capturing the values of σ for which $\Sigma_\alpha(x, \sigma)$ exhibits a “jump” in behavior. This study, in the context of images, gives rise to notions of *local scales*. Knowledge of local scales for each point x leads to criteria for selecting important, salient points in the spirit of SIFT [30, 2]. The concept of local scales arose first in the context of images [31] and was later extended to probability distributions [32]. The notions of local scales in [31, 32] were isotropic. Thus, future developments related to characterizing shape using CTFs are suggested by the possibility of constructing notions of local scales on general shapes [32, 2] which —by exploiting the tensorial nature of $\Sigma_\alpha(x, \sigma)$ — become sensitive to direction.

Data Accessibility

The synthetic data used in the manifold clustering experiments are available at <https://bitbucket.org/diegodiaz-math/ctf-files/>.

Acknowledgements

This research was supported in part by NSF grants IIS-1422400, DMS-1418007 and DBI-1262351, and by the Erwin Schrödinger Institute in Vienna.

We thank Dejan Slepčev for a discussion about the results of [6].

- [1] A. Little, Y.-M. Jung, M. Maggioni, Estimation of intrinsic dimensionality of samples from noisy low-dimensional manifolds in high dimensions with multiscale SVD, in: IEEE/SP 15th Workshop on Statistical Signal Processing, 2009, pp. 85–88.
- [2] D. Díaz Martínez, F. Mémoli, W. Mio, Multiscale covariance fields, local scales, and shape transforms, in: F. Nielsen, F. Barbaresco (Eds.), Geometric Science of Information, Vol. 8085 of Lecture Notes in Computer Science, Springer, 2013, pp. 794–801.
- [3] C. Villani, Topics in Optimal Transportation, Vol. 58 of Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, 2003.
- [4] C. Villani, Optimal Transport, Old and New, Springer, 2009.
- [5] N. Fournier, A. Guillin, On the rate of convergence in Wasserstein distance of the empirical measure, Probability Theory and Related Fields.
- [6] N. García-Trillos, D. Slepčev, On the rate of convergence of empirical measures in ∞ -transportation distance, Canadian Journal of Mathematics, doi : 10.4153/CJM-2014-044-6.
- [7] W. Allard, G. Chen, M. Maggioni, Multi-scale geometric methods for data sets II: Geometric multi-resolution analysis, Appl Comput Harmonic Anal 32 (2012) 435–462.
- [8] J. Berkmann, T. Caelli, Computation of surface geometry and segmentation using covariance techniques, IEEE Trans Pattern Anal Mach Intell 16 (1994) 1114–1116.
- [9] R. Rusu, Semantic 3D object maps for everyday manipulation in human living environments, Ph.D. thesis, Technische Universität München (2009).
- [10] G. Medioni, M.-S. Lee, C.-H. Tang, A Computational Framework for Segmentation and Grouping, Elsevier Science, 2000.

- [11] T. Brox, B. Rosenhahn, D. Cremers, H.-P. Seidel, Nonparametric density estimation with adaptive, anisotropic kernels for human motion tracking, in: *Human Motion – Understanding, Modeling, Capture and Animation*, Vol. 4814 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 152–165.
- [12] R. Vidal, Y. Ma, S. Sastry, Generalized principal component analysis (GPCA), *IEEE Trans Pattern Anal Mach Intell* 27 (12) (2005) 1945–1959.
- [13] V. Govind, A tensor decomposition for geometric grouping and segmentation, in: *IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1, 2005, pp. 1150–1157.
- [14] G. Chen, G. Lerman, Spectral curvature clustering (SCC), *Int J Comput Vis* 81 (3) (2009) 317–330.
- [15] G. Lerman, T. Zhang, Robust recovery of multiple subspaces by geometric ℓ_p minimization, *Ann Stat* 39 (5) (2011) 2686–2715.
- [16] T. Zhang, A. Szlam, Y. Wang, G. Lerman, Hybrid linear modeling via local best-fit flats, *Int J Comput Vis* 100 (3) (2012) 217–240.
- [17] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Ma, Robust recovery of subspace structures by low-rank representation, *IEEE Trans Pattern Anal Mach Intell* 35 (1) (2013) 171–184.
- [18] E. Elhamifar, R. Vidal, Sparse subspace clustering: algorithm, theory, and applications, *IEEE Trans Pattern Anal Mach Intell* 35 (11) (2013) 2765–2781.
- [19] R. Souvenir, R. Pless, Manifold clustering, in: *International Conference on Computer Vision*, 2005, pp. 648–653.
- [20] G. Carlsson, F. Mémoli, Characterization, stability and convergence of hierarchical clustering methods, *J Mach Learn Res* 11 (2010) 1425–1470.
- [21] J. Bruce, P. Giblin, *Curves and singularities: a geometrical introduction to singularity theory*, Cambridge University Press, 1992.
- [22] C. Givens, R. Shortt, A class of Wasserstein metrics for probability distributions, *Michigan Math J* 31 (2) (1984) 231–240.

- [23] R. M. Dudley, *Real Analysis and Probability in Cambridge Studies in Advanced Mathematics*, Vol. 74, Cambridge University Press, Cambridge, England, 2002.
- [24] A. Jain, R. Dubes, *Algorithms for Clustering Data*, Prentice Hall, Englewood Cliffs, NJ, 1988.
- [25] R. Tron, R. Vidal, A benchmark for the comparison of 3-D motion segmentation algorithms, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8.
- [26] R. Sibson, SLINK: an optimally efficient algorithm for the single-link cluster method, *Comp J* 16 (1) (1973) 30–34.
- [27] D. Burago, Y. Burago, S. Ivanov, *A Course in Metric Geometry*, Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, 2001.
- [28] F. Mémoli, Gromov-Wasserstein distances and the metric approach to object matching, *Foundations of Computational Mathematics* (2011) 1–71.
- [29] R. Coifman, S. Lafon, Diffusion maps, *Appl Comput Harmonic Anal* 21 (1) (2006) 5–30.
- [30] D. G. Lowe, Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision* 60 (2) (2004) 91–110.
- [31] P. W. Jones, T. M. Le, Local scales and multiscale image decompositions, *Applied and Computational Harmonic Analysis* 26 (3) (2009) 371–394.
- [32] T. Le, F. Mémoli, Local scales on curves and surfaces, *Applied and Computational Harmonic Analysis* 33 (3) (2012) 401 – 437.