

Geodesic Exponential Kernels: When Curvature and Linearity Conflict

Aasa Feragen
DIKU, University of Copenhagen
Denmark
aasa@diku.dk

François Lauze
DIKU, University of Copenhagen
Denmark
francois@diku.dk

Søren Hauberg
DTU Compute
Denmark
sohau@dtu.dk

Abstract

We consider kernel methods on general geodesic metric spaces and provide both negative and positive results. First we show that the common Gaussian kernel can only be generalized to a positive definite kernel on a geodesic metric space if the space is flat. As a result, for data on a Riemannian manifold, the geodesic Gaussian kernel is only positive definite if the Riemannian manifold is Euclidean. This implies that any attempt to design geodesic Gaussian kernels on curved Riemannian manifolds is futile. However, we show that for spaces with conditionally negative definite distances the geodesic Laplacian kernel can be generalized while retaining positive definiteness. This implies that geodesic Laplacian kernels can be generalized to some curved spaces, including spheres and hyperbolic spaces. Our theoretical results are verified empirically.

1. Introduction

Standard statistics and machine learning tools require input data residing in a Euclidean space. However, many types of data are more faithfully represented in general non-linear metric spaces (e.g. Riemannian manifolds). This is, for instance, the case when analyzing shapes [10, 16, 26, 33, 41, 56, 58, 66], DTI images [25, 46, 49, 64], motion models [14, 61], symmetric positive definite matrices [12, 50, 62], illumination-invariance [13], human poses [32, 47], tree structured data [20, 22, 23], metrics [28, 31], probability distributions [2]; for general manifold learning metrics [60] or in general for data invariant to a group action [42]. The underlying metric space captures domain specific knowledge, e.g. non-linear constraints, which is available *a priori*. The intrinsic geodesic metric encodes this knowledge, often leading to improved statistical models.

A seemingly straightforward approach to statistics in metric spaces is to use kernel methods [54], designing kernels $k(x, y)$ which only rely on geodesic distances $d(x, y)$ between observations [15]:

$$k(x, y) = \exp(-\lambda(d(x, y))^q), \quad \lambda, q > 0. \quad (1)$$

For $q = 2$ this gives a geodesic generalization of the *Gaussian kernel*, and $q = 1$ gives the geodesic *Laplacian kernel*.

Kernel	Extends to general	
	Metric spaces	Riemannian manifolds
Gaussian ($q = 2$)	No (only if flat)	No (only if Euclidean)
Laplacian ($q = 1$)	Yes, iff metric is CND	Yes, iff metric is CND
Geodesic exp. ($q > 2$)	Not known	No

Table 1. Overview of results: For a geodesic metric, when is the geodesic exponential kernel (1) positive definite for all $\lambda > 0$?

While this idea has an appealing similarity to familiar Euclidean kernel methods, we show that it is highly limited if the metric space is curved.

Positive definiteness of a kernel k is critical for the use of kernel methods such as support vector machines or kernel PCA, as it ensures the existence of a reproducing kernel Hilbert space where these methods act [54]. In this paper, we analyze exponential kernels on geodesic metric spaces and show the following results, summarized in Table 1.

- The *geodesic Gaussian kernel* is positive definite (PD) for all $\lambda > 0$ only if the underlying metric space is flat (Theorem 1). In particular, when the metric space is a Riemannian manifold, the geodesic Gaussian kernel is PD for all $\lambda > 0$ if and only if the manifold is Euclidean (Theorem 2). This negative result implies that Gaussian kernels cannot be generalized to any non-trivial Riemannian manifolds of interest.
- The *geodesic Laplacian kernel* is PD if and only if the metric is conditionally negative definite (Theorem 4). This condition is not generally true for metric spaces, but it holds for a number of spaces of interest. In particular, the geodesic Laplacian kernel is PD on spheres, hyperbolic spaces, and Euclidean spaces (Table 2).
- For any Riemannian manifold (M, g) , the kernel (1) will never be PD for all $\lambda > 0$ if $q > 2$ (Theorem 3).

Generalization of geodesic kernels to metric spaces is motivated by the general lack of powerful machine learning techniques in these spaces. In that regard, our first results are disappointing as they imply that generalizing Gaussian kernels to metric spaces is *not* a viable direction forward. Intuitively, this is not surprising as kernel methods embed the data in a linear space, which cannot be expected to capture the curvature of a general metric space. Our second result is therefore a positive surprise: it allows the Laplacian kernel to be applied in *some* metric spaces, although this has strong implications for their geometric properties. This gives hope




Figure 1. Path length in a metric space is defined as the supremum of lengths of finite approximations of the path.

that other kernels can be generalized, though our third result indicates that the geodesic exponential kernels (1) have limited applicability on Riemannian manifolds.

The paper is organized as follows. We state our main results and discuss their consequences in Sec. 2, postponing proofs until Sec. 3, which includes a formal discussion of the preliminaries. This section can be skipped in a first reading of the paper. Related work is discussed in detail in Sec. 4, where we also review recent approaches which do not conflict with our results. Sec. 5 contains empirical experiments confirming and extending our results on manifolds that admit PD geodesic exponential kernels.

2. Main results and their consequences

Before formally proving our main theorems, we state the results and provide hints as to *why* they hold. We start with a brief review of *metric geometry* and the notion of a *flat space*, both of which are fundamental to the results.

In a general metric space (X, d) with distance metric d , the *length* $l(\gamma)$ of a path $\gamma: [0, L] \rightarrow X$ from x to y is defined as the smallest upper bound of any finite approximation of the path (see Fig. 1)

$$l(\gamma) = \sup_{0=t_0 < t_1 < \dots < t_n=1, n \in \mathbb{N}} \sum_{i=1}^n d(t_{i-1}, t_i).$$

A path $\gamma: [0, L] \rightarrow X$ is called a *geodesic* [9] from x to y if $\gamma(0) = x$, $\gamma(L) = y$ and $d(\gamma(t), \gamma(t')) = |t - t'|$ for all $t, t' \in [0, L]$. In particular, $l(\gamma) = d(x, y) = L$ for a geodesic γ . In a Euclidean space, geodesics are straight lines. A geodesic from x to y will always be the shortest possible path from x to y , but geodesics with respect to a given metric do not always exist, even if shortest paths do. An example is given later in Fig. 3.

A metric space (X, d) is called a *geodesic space* if every pair $x, y \in X$ can be connected by a geodesic. Informally, a geodesic metric space is merely a space in which distances can be computed as lengths of geodesics, and data points can be interpolated via geodesics.

Riemannian manifolds are a commonly used class of metric spaces. Here distances are defined locally through a smoothly changing inner product in the tangent space. Intuitively, a Riemannian manifold can be thought of as a smooth surface (e.g. a sphere) with geodesics corresponding to shortest paths on the surface. A geodesic distance metric corresponding to the Riemannian structure is defined explicitly as the length of the geodesic joining two points. Whenever a Riemannian manifold is complete, it is a geodesic space. This is the case for most manifolds of interest.




Figure 2. If any geodesic triangle in (X, d) can be isometrically embedded into some Euclidean space, then X is *flat*. Note in particular that when a geodesic triangle is isometrically embedded in a Euclidean space, it is embedded onto a Euclidean triangle — otherwise the geodesic edges would not be isometrically embedded.

Many efficient machine learning algorithms are available in Euclidean spaces; their generalization to metric spaces is an open problem. *Kernel methods* form an immensely popular class of algorithms including *support vector machines* and *kernel PCA* [54]. These algorithms rely on the specification of a kernel $k(x, y)$, which embeds data points x, y in a linear Hilbert space and returns their inner product. Kernel methods are very flexible, as they only require the computation of inner products (through the kernel). However, the kernel is only an inner product if it is PD, so kernel methods are only well-defined for kernels which are PD [54].

Many popular choices of kernels for Euclidean data rely only on the Euclidean distance between data points; for instance the widely used Gaussian kernel (given by (1) with $q = 2$). Kernels which only rely on distances form an obvious target for generalizing kernel methods to metric spaces, where distance is often the only quantity available.

2.1. Main results

In Theorem 1 of this paper we prove that **geodesic Gaussian kernels on metric spaces are PD for all $\lambda > 0$ only if the metric space is flat**. Informally, a metric space is flat if it (for all practical purposes) is Euclidean. More formally:

Definition 1. A geodesic metric space (X, d) is *flat in the sense of Alexandrov* if any geodesic triangle in X can be isometrically embedded in a Euclidean space.

Here, an embedding $f: X \rightarrow X'$ from a metric space (X, d) to another metric space (X', d') is *isometric* if $d'(f(x), f(y)) = d(x, y)$ for all $x, y \in X$ ¹. A *geodesic triangle* abc in X consists of three points a, b and c joined by geodesic paths γ_{ab} , γ_{bc} and γ_{ac} . The concept of flatness essentially requires that all geodesic triangles are *identical* to Euclidean triangles; see Fig. 2.

With this, we state our first main theorem:

Theorem 1. *Let (X, d) be a geodesic metric space, and assume that $k(x, y) = \exp(-\lambda d^2(x, y))$ is a PD geodesic Gaussian kernel on X for all $\lambda > 0$. Then (X, d) is flat in the sense of Alexandrov.*

¹The metric space definition of *isometric embedding* [9], which is used when distances are in focus, should not be confused with the definition of isometric embedding from Riemannian geometry, preserving *Riemannian metrics* which are not distances, but tangent space inner products.

This is a negative result, in the sense that most metric spaces of interest are not flat. In fact, the motivation for generalizing kernel methods is to cope with data residing in non-flat metric spaces.

As a consequence of Theorem 1, we show that **geodesic Gaussian kernels on Riemannian manifolds are PD for all $\lambda > 0$ only if the Riemannian manifold is Euclidean.**

Theorem 2. *Let M be a complete, smooth Riemannian manifold with its associated geodesic distance metric d . Assume, moreover, that $k(x, y) = \exp(-\lambda d^2(x, y))$ is a PD geodesic Gaussian kernel for all $\lambda > 0$. Then the Riemannian manifold M is isometric to a Euclidean space.*

These two theorems have several consequences. The first and main consequence is that defining geodesic Gaussian kernels on Riemannian manifolds or other geodesic metric spaces has limited applicability as most spaces of interest are not flat. In particular, on Riemannian manifolds the kernels will generally only be PD if the original data space is Euclidean. In this case, nothing is gained by treating the data space as a Riemannian manifold, as it is perfectly described by the well-known Euclidean geometry, where many problems can be solved in closed form. In Sec. 4 we re-interpret recent work which does, indeed, take place in Riemannian manifolds that turn out to be Euclidean.

Second, *this result is not surprising*: Curvature cannot be captured by a flat space, and Schönberg’s classical theorem (see Sec. 3.1) indicates a strong connection between PD Gaussian kernels and linearity of the employed distance measure. This is made explicit by Theorems 1 and 2.

While this paper was in print, a result similar to Theorem 2 appeared in [39]. However, the authors do not note that as a consequence, Gaussian RBF kernels that use the geodesic distance only apply to Riemannian manifolds that are Euclidean spaces, where they coincide with the standard Gaussian kernels [54]. In order to apply Gaussian kernels to non-Euclidean spaces they are forced to replace the Riemannian structure by a Euclidean chordal metric.

The obvious next question is the extent to which these negative results depend on the choice $q = 2$ in (1), which results in a Gaussian kernel. A recent result by Istas [35] implies that for Riemannian manifolds, passing to a higher power $q > 2$ will *never* lead to a PD kernel for all $\lambda > 0$:

Theorem 3. *Let M be a Riemannian manifold with its associated geodesic distance metric d , and let $q > 2$. Then there is some $\lambda > 0$ so that the kernel (1) is not PD.*

The existence of a $\lambda > 0$ such that the kernel is not PD may seem innocent; however, as a consequence, the kernel bandwidth parameter cannot be learned.

In contrast, the choice $q = 1$ in (1), giving a geodesic Laplacian kernel, leads to a more positive result: The geodesic Laplacian kernel will be positive definite if and only if the distance d is *conditionally negative definite* (CND). CND metrics have linear embeddability properties analogous to those of PD kernels; see Sec. 3.1 for formal




Figure 3. The chordal metric on $\mathbb{S}^2 \subset \mathbb{R}^3$ is measured directly in \mathbb{R}^3 , while the geodesic metric is measured along \mathbb{S}^2 . Shortest paths with respect to the two metrics coincide, but the chordal metric is *not* a geodesic metric, and the shortest path is *not* a geodesic for the chordal metric, because the shortest path between two points is longer than their chordal distance.

definitions and properties. This provides a PD kernel framework which, for several popular Riemannian data manifolds, takes advantage of the geodesic distance.

Theorem 4. *i) The geodesic distance d in a geodesic metric space (X, d) is CND if and only if the corresponding geodesic Laplacian kernel is PD for all $\lambda > 0$.*

ii) In this case, the square root metric $d_{\sqrt{\cdot}}(x, y) = \sqrt{d(x, y)}$ is also a distance metric, and $(X, d_{\sqrt{\cdot}})$ can be isometrically embedded as a metric space into a Hilbert space H .

iii) The square root metric $d_{\sqrt{\cdot}}$ is not a geodesic metric, and $d_{\sqrt{\cdot}}$ corresponds to the chordal metric in H , not the intrinsic metric on the image of X in H .

In Theorem 4, for $\phi: X \rightarrow H$, the chordal metric $\|\phi(x) - \phi(y)\|_H$ measures distances directly in H rather than intrinsically in the image $\phi(X) \subset H$, see also Fig. 3.

In Sec. 4 we discuss several popular data spaces for which geodesic Laplacian kernels are PD (see Table 2); examples include spheres, hyperbolic spaces and more. Nevertheless, we see from part *ii*) of Theorem 4 that any geodesic metric space whose geodesic Laplacian kernel is always PD must necessarily have strong linear properties: Its square root metric is isometrically embeddable in a Hilbert space.

This illustrates an intuitively simple point: A PD kernel has no choice but to linearize the data space. Therefore, its ability to capture the original data space geometry is deeply connected to the linear properties of the original metric².

3. Proofs of main results

In this section we prove the main results of the paper; this section may be skipped in a first reading of the paper. In the first two subsections we review and discuss classical geometric results on kernels, manifolds and curvature, which we will use to prove the main results.

3.1. Kernels

A modern and comprehensive treatment of the classical results on PD and CND kernels referred to here, can be found in [5, Appendix C].

²Another curious connection between kernels and curvature is found in [11], which shows that Gaussian and polynomial kernels on \mathbb{R}^n and \mathbb{R}^2 , respectively, have flat feature space images $\phi(\mathbb{R}^n)$ and $\phi(\mathbb{R}^2)$.

Definition 2. A positive definite (PD) kernel on a topological space X is a continuous function $k: X \times X \rightarrow \mathbb{R}$ such that for any $n \in \mathbb{N}$, any elements $x_1, \dots, x_n \in X$ and any numbers $c_1, \dots, c_n \in \mathbb{R}$, we have

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j k(x_i, x_j) \geq 0.$$

Definition 3. A conditionally negative definite (CND) kernel on a topological space X is a continuous function $\psi: X \times X \rightarrow \mathbb{R}$ which satisfies

- i) $\psi(x, x) = 0$ for all $x \in X$
- ii) $\psi(x, y) = \psi(y, x)$ for all $x, y \in X$
- iii) for any $n \in \mathbb{N}$, any elements $x_1, \dots, x_n \in X$ and any real numbers c_1, \dots, c_n with $\sum_{i=1}^n c_i = 0$, we have

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \psi(x_i, x_j) \leq 0.$$

Example 1. If $d: H \times H \rightarrow \mathbb{R}$ is the metric induced by the norm on a Hilbert space H , then the map $d^2: H \times H \rightarrow \mathbb{R}$ given by $d^2(x, y) = (d(x, y))^2$ is a CND kernel [5].

The following two theorems are key to understanding the connection between distance metrics and their corresponding exponential kernels.

Theorem 5 (Due to Schönberg [55], Theorem C.3.2 in [5]). *If X is a topological space and $\psi: X \times X \rightarrow \mathbb{R}$ is a continuous kernel on X with $\psi(x, x) = 0$ and $\psi(y, x) = \psi(x, y)$ for all $y, x \in X$, then the following are equivalent:*

- ψ is a CND kernel
- the kernel $k(x, y) = e^{-\lambda\psi(x, y)}$ is PD for all $\lambda \geq 0$. \square

Theorem 6 (Part of Theorem C.2.3 in [5]). *If $\psi: X \times X \rightarrow \mathbb{R}$ is a CND kernel on a topological space X , then there is a real Hilbert space H and a continuous mapping $f: X \rightarrow H$ such that $\psi(x, y) = \|f(x) - f(y)\|_H^2$ for all $x, y \in X$. \square*

From the above, it is straightforward to deduce:

Corollary 1. *If the geodesic Gaussian kernel is PD, then there is a mapping $f: X \rightarrow H$ into some Hilbert space H such that*

$$d(x, y) = \|f(x) - f(y)\|_H$$

for each $x, y \in X$. Note that this mapping f is not necessarily related to the feature mapping $\phi: X \rightarrow V$ such that $k(x, y) = \langle \phi(x), \phi(y) \rangle_V$.

3.2. Curvature

While curvature is usually studied using differential geometry, we shall access curvature via a more general approach that applies to general geodesic metric spaces. This notion of curvature, originating with Alexandrov and Gromov, operates by comparing the metric space to spaces whose geometry we understand well, referred to as *model*




Figure 4. Left: A geodesic triangle, right: the corresponding comparison triangles in hyperbolic space \mathbb{H}^2 , the plane \mathbb{R}^2 and the sphere \mathbb{S}^2 , respectively.

spaces. The model spaces M_κ are *spheres* (of positive curvature $\kappa > 0$), the *Euclidean plane* (flat, curvature $\kappa = 0$) and *hyperbolic space* (negative curvature $\kappa < 0$). Since metric spaces can be pathological, curvature is approached by bounding the curvature of the space at a given point from above or below. The bounds are attained by comparing geodesic triangles in the metric space with triangles in the model spaces, as expressed in the $CAT(\kappa)$ condition:

Definition 4. Let (X, d) be a geodesic metric space X . Let abc be a geodesic triangle of perimeter $< 2D_\kappa$, where D_κ is the diameter of M_κ , that is, $D_\kappa = \infty$ for $\kappa \leq 0$, and $D_\kappa = \frac{\pi}{\sqrt{\kappa}}$ for $\kappa > 0$. There exists a triangle $\bar{a}\bar{b}\bar{c}$ in the model space M_κ with vertices \bar{a}, \bar{b} and \bar{c} and with geodesic edges $\bar{\gamma}_{\bar{a}\bar{b}}, \bar{\gamma}_{\bar{b}\bar{c}}$ and $\bar{\gamma}_{\bar{a}\bar{c}}$, whose lengths are the same as the lengths of the edges γ_{ab}, γ_{bc} and γ_{ac} in abc . This is an M_κ -comparison triangle for abc (see Fig. 4).

For any point x sitting on the segment γ_{bc} , there is a corresponding point \bar{x} on the segment $\bar{\gamma}_{\bar{b}\bar{c}}$ in the comparison triangle, such that $d_{M_\kappa}(\bar{x}, \bar{b}) = d(x, b)$. If we have

$$d(x, a) \leq d_{M_\kappa}(\bar{x}, \bar{a}) \quad (2)$$

for every such x , and similarly for any x on γ_{ab} or γ_{ac} , then the geodesic triangle abc satisfies the $CAT(\kappa)$ condition.

The metric space X is a $CAT(\kappa)$ space if any geodesic triangle abc in X of perimeter $< 2D_\kappa$ satisfies the $CAT(\kappa)$ condition given in eq. 2. Geometrically, this means that triangles in X are *thinner* than triangles in M_κ . The metric space X has *curvature $\leq \kappa$ in the sense of Alexandrov* if it is locally $CAT(\kappa)$.

While curvature in the $CAT(\kappa)$ sense allows the study of curvature through the relatively simple means of geodesic distances alone, it is a weaker concept of curvature than the standard sectional curvature used in Riemannian geometry. Nevertheless, the two concepts are related, as captured by the following theorem due to Cartan and Alexandrov:

Theorem 7 (Theorem II.1A.6 [9]). *A smooth Riemannian manifold M is of curvature $\leq \kappa$ in the sense of Alexandrov if and only if the sectional curvature of M is $\leq \kappa$. \square*

The proof of the main theorem will, moreover, rely on the following theorem characterizing manifolds of constant zero sectional curvature:

Theorem 8 (Part of Theorem 11.12 [45]). *Let M be a complete, simply connected m -dimensional Riemannian manifold with constant sectional curvature $C = 0$. Then M is isometric to \mathbb{R}^m . \square*

We are now ready to prove our main theorems.

3.3. Geodesic Gaussian kernels on metric spaces: Proof of Theorem 1

As in the statement of Theorem 1, assume that the metric space (X, d) is a geodesic space as defined in Sec. 2, and that $k(x, y) = e^{-\lambda d^2(x, y)}$ is a PD geodesic Gaussian kernel on X for all $\lambda > 0$. An important consequence of Theorem 6 is that the map $f: X \rightarrow H$ must take geodesic segments to geodesic segments, which in H are straight line segments.

Lemma 1. *If $\gamma: [0, L] \rightarrow X$ is a geodesic of length L from $a = \gamma(0)$ to $b = \gamma(L)$ in X , then $f(\gamma([0, L]))$ is the straight line from $f(a)$ to $f(b)$ in H , and*

$$f(\gamma(t)) = f(a) + \frac{t}{L} (f(b) - f(a)) \quad (3)$$

for all $t \in [0, L]$.

Proof. Since $\gamma: [0, L] \rightarrow X$ is a geodesic, it contains every point $\gamma(t)$ for all $t \in [0, L]$, and since γ is a geodesic of length L , we have $d(\gamma(0), \gamma(t)) = t$ for each $t \in [0, L]$, so

$$\|f(\gamma(0)) - f(\gamma(t))\| = d(\gamma(0), \gamma(t)) = t.$$

This is only possible if $f \circ \gamma$ is the straight line from $f(a)$ to $f(b)$ in H . Equation (3) follows directly, as it is the geodesic parametrization of a straight line from $f(a)$ to $f(b)$. \square

This enables us to prove Theorem 1:

Proof of Theorem 1. Let $a, b, c \in X$ be three points in X and form a geodesic triangle spanned by their joining geodesics γ_{ab}, γ_{bc} and γ_{ca} . Then the points $f(a), f(b)$ and $f(c)$ in H are connected by straight line geodesics $f \circ \gamma_{ab}, f \circ \gamma_{bc}$ and $f \circ \gamma_{ca}$ by Lemma 1. These points and geodesics in H span a 2-dimensional linear subspace of H in which they form a Euclidean comparison triangle.

Without loss of generality, pick any two points x and y on the geodesic triangle and measure the distance $d(x, y)$. The corresponding distance in the comparison triangle is $\|f(x) - f(y)\|$, and by the definition of f we know that $d(x, y) = \|f(x) - f(y)\|$, so the geodesic triangle is isometrically embedded into the comparison triangle. Hence, X is flat in the sense of Alexandrov. \square

Corollary 2. *The metric space X is contractible, and hence simply connected.*

Proof. By Theorem 1, X must necessarily be a $CAT(0)$, and contractible by [9, Corollary II.1.5]. \square

3.4. Geodesic Gaussian kernels on Riemannian manifolds: Proof of Theorem 2

We prove that for a complete, smooth Riemannian manifold M with associated geodesic distance metric d , if the geodesic Gaussian kernel $k(x, y) = e^{-\lambda d^2(x, y)}$ is PD for all $\lambda > 0$, then M is isometric to a Euclidean space.

Proof of Theorem 2. We start out by showing that the sectional curvature of M is 0 everywhere.

By Theorem 1, M is a $CAT(0)$ space, so in particular it has curvature ≤ 0 in the sense of Alexandrov. Therefore, by Theorem 7, the sectional curvature of M is ≤ 0 .

To prove the claim, we need to show that M does not have any points with negative sectional curvature. To this end, assume that there is some point $p \in M$ such that the sectional curvature of M at p is $\kappa < 0$. Then, since sectional curvature on smooth Riemannian manifolds is continuous, there exists some neighborhood U of p and some $\kappa' < 0$ such that the sectional curvature in U is $\leq \kappa' < 0$. But then, by Theorem 7, U also has curvature $\leq \kappa'$ in the sense of Alexandrov, which cannot hold due to Theorem 1. It follows that the sectional curvature of M at p cannot be $\kappa < 0$; hence, the sectional curvature of M must be everywhere 0.

Since M is simply connected by Corollary 2, we apply Theorem 8 to conclude that M must be isometric to \mathbb{R}^m . \square

3.5. The case $q > 2$

Proof of Theorem 3. This is a direct consequence of [35, Theorem 2.12]. \square

3.6. Geodesic Laplacian kernels: Proof of Theorem 4

Another consequence of Schönberg's Theorem 5 is that the geodesic Laplacian kernel defined by (1) with $q = 1$ is PD if and only if the distance d is CND. This provides a PD kernel framework which, for several popular Riemannian data manifolds, utilizes the geodesic distance.

Proof of Theorem 4. i) By Theorem 5, d is CND if and only if the Laplacian kernel $k(x, y) = e^{-\lambda d(x, y)}$ is PD for all $\lambda > 0$.

ii) By Theorem 6, there exists a real Hilbert space H and a continuous map $f: X \rightarrow H$ such that

$$d(x, y) = \|f(x) - f(y)\|_H^2 \text{ for all } x, y \in X. \quad (4)$$

That is, $d_{\sqrt{\cdot}}(x, y) = \|f(x) - f(y)\|_H$ for all $x, y \in X$. The map f must be injective, because if $f(x) = f(y)$ for $x \neq y$ then by (4), $0 = \|f(x) - f(y)\|_H = d(x, y) > 0$, which is false. Therefore, $d_{\sqrt{\cdot}}$ coincides with the restriction to $f(X)$ of the metric on H induced by $\|\cdot\|$. Since the restriction of a metric to a subset is a metric, $d_{\sqrt{\cdot}}$ is a metric, and by definition, f is an isometric embedding of $(X, d_{\sqrt{\cdot}})$ into H .

iii) Since f is an isometric embedding as metric spaces, $d_{\sqrt{\cdot}}$ must correspond to the chordal metric in H .

Assume that $d_{\sqrt{\cdot}}$ is a geodesic metric on X , then by Lemma 1, f maps geodesics in $(X, d_{\sqrt{\cdot}})$ to straight line segments in H . Focusing on a single geodesic segment $\gamma: [0, L] \rightarrow X$, we obtain

$$d_{\sqrt{\cdot}}(\gamma(t), \gamma(t')) = \|f \circ \gamma(t) - f \circ \gamma(t')\| = |t - t'|$$

for all $t, t' \in [0, L]$. Since $d = d_{\sqrt{\cdot}}^2$ is a metric by assumption, the square $d_{\gamma}(t, t') = |t - t'|^2 =$

$d(\gamma(t), \gamma(t'))$ is a metric on $[0, L]$. But this is not true, as the triangle inequality fails to hold.

Therefore, $d_{\sqrt{\cdot}}$ cannot be a geodesic metric on X . \square

As noted in Table 2 below, for a number of popular Riemannian manifolds, the geodesic distance metric is *CND*, meaning that geodesic Laplacian kernels are *PD*.

Remark 5. For a *CND* distance metric $d: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, a *PD* kernel $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ can also be constructed through the formula $k(x, x') = d(x, x') - d(x, x_0) - d(x_0, x')$ [7, 53], where $x_0 \in \mathcal{X}$ is any point. For other distance-based kernels, *e.g.* the rational-quadratic kernel, little is known.

4. Implications for popular manifolds and related work

Many popular data spaces appearing in computer vision are *not* flat, meaning that their geodesic distances are not *CND* and their geodesic Gaussian kernels will not be *PD*. Table 2 lists known results on *CND* status of some popular data spaces. In particular, the classical intrinsic metrics on \mathbb{R}^n , \mathbb{H}^n and \mathbb{S}^n are all *CND*³. As the Fisher information metric on 1-dimensional normal distributions defines the hyperbolic geometry \mathbb{H}^2 [2], it will give a *CND* geodesic metric. For projective space, on the other hand, [51] provides an example showing that the classical intrinsic metric is *not* *CND*. As Grassmannians are generalizations of projective spaces, their geodesic metrics are therefore also not generally *CND*.

Symmetric, positive definite ($d \times d$) matrices form another important data manifold, denoted Sym_d^+ . While the popular Frobenius and Log-Euclidean [3] metrics on Sym_d^+ are actually Euclidean, little is known theoretically about whether the geodesic distance metrics of non-Euclidean Riemannian metrics on Sym_d^+ are *CND*. In Sec. 5 we show empirically that neither the affine-invariant metric [49] nor the Fisher information metric on the corresponding fixed-mean multivariate normal distributions [2, 4] induce a *CND* geodesic metric. Note how the qualitatively similar affine-invariant and Log-Euclidean metrics differ in whether they generate *PD* exponential kernels.

Non-manifold data spaces are also popular, *e.g.* the *edit distance* on strings was shown not to be *CND* by Cortes *et al.* [17]. As tree- and graph edit distances generalize string edit distance, the same holds for these. For this reason, *PSD* graph kernels are often similarity-based [8, 21], not distance-based. The metric along a metric tree, on the other hand, is *CND*. In Sec. 5 we show empirically that this does not generalize to the shortest path metric on a geometric graph, such as the k NN or ϵ -neighborhood graphs often used in manifold learning [1, 6, 52, 59].

4.1. Relation to previous work

Several *PD* kernels on manifolds have appeared in the literature, some of them even Gaussian kernels based on distance metrics on manifolds such as spheres or Grassman-

nian manifolds, which we generally consider as curved manifolds. The reader might wonder how this is possible given the above presented results. The explanation is that the distances used in these kernels are *not* geodesic distances and, in many cases, have little or nothing to do with the Riemannian structure of the manifold. We discuss a few examples.

Example 2. In [37], a *PD* kernel is defined on Sym_d^+ by using a geodesic Gaussian kernel with the log-Euclidean metric [49]. The log-Euclidean metric is defined by pulling the (Euclidean) Frobenius metric on Sym_d back to Sym_d^+ via the diffeomorphic matrix logarithm. Equivalently, data in Sym_d^+ is mapped into the Euclidean Sym_d via the diffeomorphic log map, and data is analyzed there. The geodesic Gaussian kernel is *PD* because the Riemannian manifold is actually a Euclidean space. In such cases, the Riemannian framework only adds an unnecessary layer of complexity.

Example 3. In [38], radial kernels are defined on spheres \mathbb{S}^n by restricting kernels on \mathbb{R}^{n+1} to \mathbb{S}^n , giving radial kernels with respect to the chordal metric on \mathbb{S}^n . Due to the symmetry of \mathbb{S}^n , any kernel which is radial with respect to the chordal metric, will also be radial with respect to the geodesic metric on \mathbb{S}^n . This result is next used to define *PD* radial kernels on the Grassmannian manifold \mathcal{G}_n^r and on the Kendall shape space \mathcal{SP}^n . However, these kernels are not radial with respect to the usual Riemannian metrics on these spaces, but with respect to the projection distance and the full Procrustes distance, respectively, both of which are *not* geodesic distances with respect to any Riemannian metric on \mathcal{G}_n^r and \mathcal{SP}^n , respectively⁴. These kernels, thus, have little to do with the Riemannian geometry of \mathcal{G}_n^r and \mathcal{SP}^n .

Example 4. In [18] it is noted that since the feature map ϕ corresponding to a Euclidean Gaussian kernel maps data onto a hypersphere \mathbb{S} in the reproducing kernel Hilbert space V [54], it might improve classification to consider the geodesic distance on \mathbb{S} rather than the chordal distance from V . This is, however, done by projecting each $\phi(x) \in V$ onto the tangent space $T_{\phi(\tilde{x})}\mathbb{S}$ at a fixed base point $\phi(\tilde{x})$, where the linear kernel in V is employed. This explains why the resulting kernel $k_{\tilde{x}}$ is *PD*: the kernel linearizes the sphere and, thereby, discards the spherical geometry.

Example 5. In [34] and [40], geodesic Laplacian kernels are defined on spheres; as shown above, these are *PD*.

Example 6. In [36], a kernel is defined on a general sample space \mathcal{X} by selecting a generating probability distribution P_θ on \mathcal{X} and defining a *Fisher kernel* on \mathcal{X} . Denote by M_Θ the Riemannian manifold defined by a parametrized family of probability distributions P_θ , $\theta \in \Theta$, on \mathcal{X} endowed with the Fisher information metric. The kernel $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$

⁴Assume that either of these metrics were a Riemannian geodesic distance metric. The family of *PD* radial kernels defined in [38] on both \mathcal{G}_n^r and \mathcal{SP}^n include Gaussian kernels with the projection distance and the full Procrustes distance, respectively. By our previous results, if these were geodesic distances with respect to some Riemannian metric, this Riemannian metric would define a Euclidean structure on \mathcal{G}_n^r and \mathcal{SP}^n , respectively. This is impossible, since these manifolds are both compact.

³As a curious side note, this implies that $\sqrt{\|x - y\|}$ is a metric on \mathbb{R}^n .

Space	Distance metric	Geodesic metric?	Euclidean? metric?	CND metric?	PD Gaussian kernel?	PD Laplacian kernel?
\mathbb{R}^n [54, 55]	Euclidean metric	✓	✓	✓	✓	✓
$\mathbb{R}^n, n > 2$ [35]	l_q -norm $\ \cdot\ _q, q > 2$	✓	÷	÷	÷	÷
Sphere \mathbb{S}^n [35]	classical intrinsic	✓	÷	✓	÷	✓
Real projective space $\mathbb{P}^n(\mathbb{R})$ [51]	classical intrinsic	✓	÷	÷	÷	÷
Grassmannian	classical intrinsic	✓	÷	÷	÷	÷
Sym_d^+	Frobenius	✓	✓	✓	✓	✓
Sym_d^+	Log-Euclidean	✓	✓	✓	✓	✓
Sym_d^+	Affine invariant	✓	÷	÷	÷	÷
Sym_d^+	Fisher information metric	✓	÷	÷	÷	÷
Hyperbolic space \mathbb{H}^n [35]	classical intrinsic	✓	÷	✓	÷	✓
1-dimensional normal distributions	Fisher information metric	✓	÷	✓	÷	✓
Metric trees [63], [35, Thm 2.15]	tree metric	✓	÷	✓	÷	✓
Geometric graphs (e.g. k NN)	shortest path distance	✓	÷	÷	÷	÷
Strings [17]	string edit distance	✓	÷	÷	÷	÷
Trees, graphs	tree/graph edit distance	✓	÷	÷	÷	÷

Table 2. For a set of popular metric and manifold data spaces and metrics, we record whether the metric is a geodesic metric, whether it is a Euclidean metric, whether it is a CND metric, and whether its corresponding Gaussian and Laplacian kernels are PD.

is defined by mapping samples in \mathcal{X} to the tangent space $T_\theta M_\Theta$ and applying the Riemannian metric at $P_\theta \in M_\Theta$. This is PD because the kernel is an inner product on data mapped into a Euclidean tangent space. Again, the statistical manifold is linearized and the resulting kernel does not fully respect its geometry.

In several of these examples the data space is linearized by mapping to a tangent space or into a linear ambient space, which always gives a PD kernel. It should, however, be stressed that the resulting kernels neither respect the distances nor the constraints encoded in the original Riemannian structure. Thus, the linearization will inevitably remove the information that the kernel was aiming to encode.

In general, whenever a data space is embedded into a Euclidean/Hilbert space, and the chordal metric is used in (1), the exponential kernel on a dataset coincides with an exponential kernel on the dataset embedded in the Euclidean/Hilbert ambient space. This therefore gives a PD kernel, and by the Whitney embedding theorem [45], universal kernels can thus be defined on any manifold. These kernels will, however, disregard any constraints encoded by the geodesic distance.

It is tempting to refer to the Nash theorem [48], which states that any Riemannian manifold can be isometrically embedded into a Euclidean space. Here, however, "isometric embedding" refers to a *Riemannian* isometry, which preserves the *Riemannian metric* (the smoothly changing inner product) — not to be confused with a distance metric! Therefore, in a Riemannian isometric embedding $f: \mathcal{X} \rightarrow \mathbb{R}^n$ we typically have $d(x, y) \neq \|f(x) - f(y)\|$. A kernel based on chordal distances in a Nash embedding will, thus, not generally be related to the geodesic distance.

Note, moreover, that the Nash theorem does not guarantee a unique embedding; in fact there are viable embeddings generating a wide range of distance metrics inherited from the ambient Euclidean space. Therefore, an exponential kernel based on the chordal metric will typically have little to do with the intrinsic Riemannian structure of the manifold.

There exist PD kernels that take full advantage of Riemannian geometry without relying on geodesic distances:

Example 7. Gong *et al.* [29] design a PD kernel for *domain adaptation* using the geometry of the Grassmann manifold: Let S_1 and S_2 be two low-dimensional subspaces of \mathbb{R}^n estimated with PCA on two related data sets. This gives two points x_1, x_2 on the Grassmann manifold. A test point can be projected into all possible subspaces along the Grassmann geodesic connecting x_1 and x_2 , giving an infinite dimensional feature vector in a Hilbert space. Gong *et al.* [29] show how to compute inner products in this Hilbert space in closed-form, thereby providing a PD kernel which takes geometry into account without relying on geodesic distances.

5. Experiments

We now validate our theoretical results empirically. First, we generate 500 randomly drawn symmetric PD matrices of size 3×3 . We compute the Gram matrix of both the Gaussian and Laplacian kernels under both the affine-invariant metric [49] and the Fisher information metric on the corresponding fixed-mean multivariate normal distributions [2, 4]. Fig. 5a shows the eigenspectrum of the four different Gram matrices. All four kernels have negative eigenvalues, which imply that none of them are positive definite. This empirically proves that neither the affine-invariant metric nor the Fisher information metric induce CND geodesic distance metrics in general, although we know this to hold for the Fisher information metric on $Sym_1^+ = \mathbb{R}_+$.

Next, we consider kernels on the unit sphere. We generate data from salient points in the 1934 painting *Etude de femmes* by *Le Corbusier*. At each salient point a HOG [19] descriptor is computed; as these descriptors are normalized they are points on the unit sphere. Fig. 5b shows the eigenspectrum of the Gram matrix of the geodesic Gaussian and Laplacian kernels. While the geodesic Gaussian kernel has negative eigenvalues, the geodesic Laplacian does not. This verifies our theoretical results from Sec. 4 and Table 2.

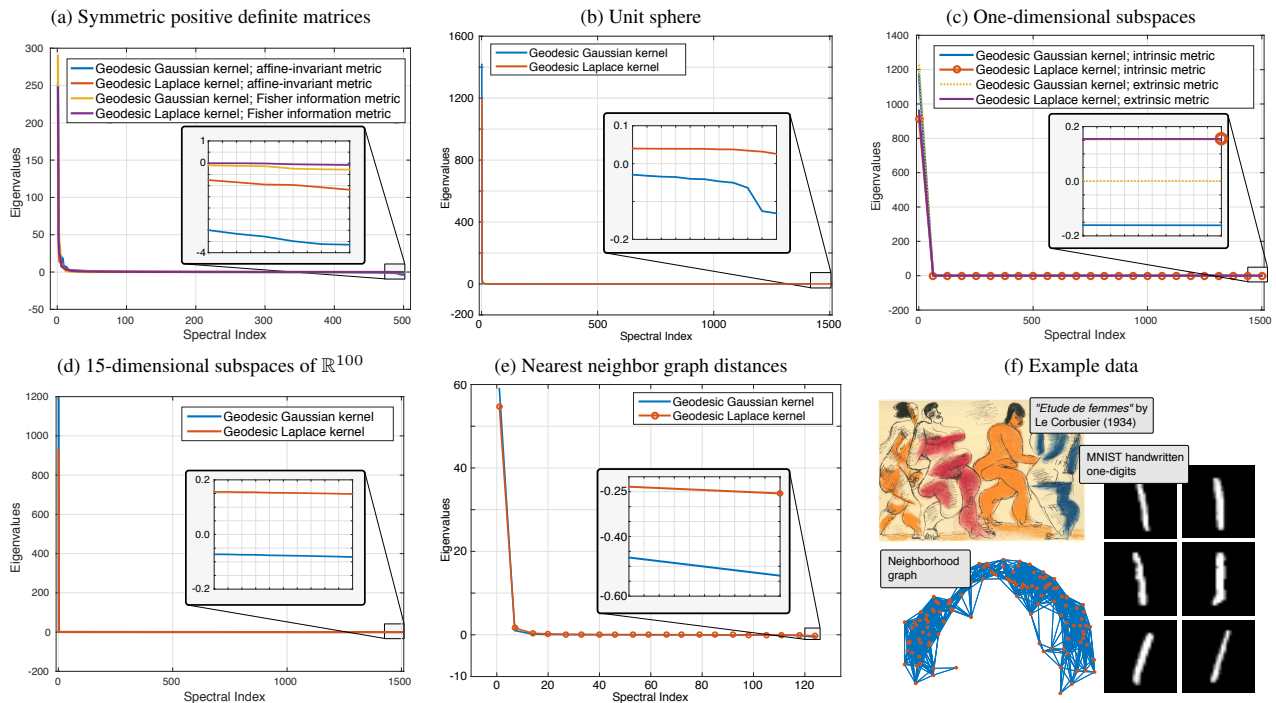


Figure 5. (a)–(e): Eigenspectra of the Gram matrices for different geodesic exponential kernels on different manifolds. (f) Data used in panels b and e.

We also consider data on the Grassmann manifold. First, we consider one-dimensional subspaces as spanned by samples from a 50-dimensional isotropic normal distribution. We again consider both the Gaussian and the Laplacian kernel; here both under the usual intrinsic metric, but also under the extrinsic metric [30]. Fig. 5c shows the eigenspectra of the different Gram matrices. Only the Gaussian kernel under the intrinsic metric appears to have negative eigenvalues, while the remaining have strictly positive eigenvalues.

Next, we consider 15-dimensional subspaces of \mathbb{R}^{100} drawn from a uniform distribution on the corresponding Grassmannian. We only consider kernels under the intrinsic metric, and the eigenspectra are shown in Fig. 5d. The Gaussian kernel has negative eigenvalues, while the Laplacian kernel does not. Note that this does not prove that the Laplacian kernel is PD on the Grassmannian; in fact, we know theoretically from [51] that it is generally *not*.

Finally, we consider shortest-path distances on nearest neighbor graphs as commonly used in manifold learning. We take 124 one-digits from the MNIST data set [44], project them into their two leading principal components, form a ϵ -neighborhood graph, and compute shortest path distances. We then compute the eigenspectrum of both the Gaussian and Laplacian kernel; Fig. 5e show these spectra. Both kernels have negative eigenvalues, which empirically show that the shortest-path graph distance is *not* CND.

6. Discussion and outlook

We have shown that exponential kernels based on geodesic distances in a metric space or Riemannian manifold will only be positive definite if the geodesic metric sat-

isfies strong linearization properties:

- for Gaussian kernels, the metric space must be flat (or Euclidean).
- for Laplacian kernels, the metric must be conditionally negative definite. This implies that the square root metric can be embedded in a Hilbert space.

With the exception of select metric spaces, these results show that geodesic exponential kernels are not well-suited for data analysis in curved spaces.

This does, however, *not* imply that kernel methods can never be extended to metric spaces. Gong *et al.* [29] provide an elegant kernel based on the geometry of the Grassmann manifold, which is well-suited for domain adaptation. This kernel is not a geodesic exponential kernel, yet it strongly incorporates the geodesic structure of the Grassmannian. As an alternative, the Euclidean Gaussian kernel is a *diffusion kernel*. Such kernels are positive definite on Riemannian manifolds [43], and might provide a suitable kernel. However, these kernels generally do not have closed-form expressions, which may hinder their applicability.

Most existing machine learning tools assume a linear data space. Kernel methods only encode non-linearity via a non-linear transformation between a data space and a linear feature space. Our results illustrate that such methods are limited for analysis of data from non-linear spaces. Emerging generalizations of learning tools such as regression [24, 33, 57] or transfer learning [27, 65] to nonlinear data spaces are encouraging. We believe that learning tools that operate directly in the non-linear data space, without a linearization step, is the way forward.

Acknowledgements

We thank Nicolas Charon, Oswin Krause, Michael Schober and Bernhard Schölkopf for insightful discussions. A.F. is funded in part by the Danish Council for Independent Research (DFF), Technology and Production Sciences. S.H. is funded by the DFF, Natural Sciences.

References

- [1] M. Alamgir and U. von Luxburg. Shortest path distance in random k-nearest neighbor graphs. In *Proceedings of the 29th International Conference on Machine Learning, ICML*, 2012. [6](#)
- [2] S. Amari and H. Nagaoka. *Methods of information geometry*. Translations of mathematical monographs; v. 191. American Mathematical Society, 2000. [1](#), [6](#), [7](#)
- [3] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Fast and simple calculus on tensors in the log-Euclidean framework. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 115–122. Springer, 2005. [6](#)
- [4] C. Atkinson and A. F. Mitchell. Rao’s distance measure. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 345–365, 1981. [6](#), [7](#)
- [5] A. V. B. Bekka, P. de la Harpe. *Kazhdan’s Property (T)*. New Mathematical Monographs, 2008. [3](#), [4](#)
- [6] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. [6](#)
- [7] C. Berg, J. P. R. Christensen, and P. Ressel. *Harmonic Analysis on Semigroups*. Springer, 1984. [6](#)
- [8] K. M. Borgwardt and H. Kriegel. Shortest-path kernels on graphs. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM)*, pages 74–81, 2005. [6](#)
- [9] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*. Springer, 1999. [2](#), [4](#), [5](#)
- [10] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff framework with diffusion geometry fortotopologically-robust non-rigid shape matching. *International Journal of Computer Vision*, 89(2-3):266–286, 2010. [1](#)
- [11] C. J. C. Burges. Geometry and Invariance in Kernel Based Methods. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods: Support Vector Learning*. MIT Press, 1999. [3](#)
- [12] J. Carreira, R. Caseiro, J. Batista, and C. Sminchisescu. Free-form region description with second-order pooling. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1–1, 2014. [1](#)
- [13] R. Caseiro, J. F. Henriques, P. Martins, and J. Batista. Semi-intrinsic mean shift on Riemannian manifolds. In *Computer Vision - ECCV - 12th European Conference on Computer Vision*, pages 342–355, 2012. [1](#)
- [14] H. Çetingül and R. Vidal. Intrinsic mean shift for clustering on Stiefel and Grassmann manifolds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1896–1902, June 2009. [1](#)
- [15] O. Chapelle, P. Haffner, and V. N. Vapnik. Support vector machines for histogram-based image classification. *Neural Networks, IEEE Transactions on*, 10(5):1055–1064, 1999. [1](#)
- [16] N. Charon and A. Trouvé. The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM Journal on Imaging Sciences*, 6(4):2547–2580, 2013. [1](#)
- [17] C. Cortes, P. Haffner, and M. Mohri. Rational kernels: Theory and algorithms. *Journal of Machine Learning Research*, 5:1035–1062, 2004. [6](#), [7](#)
- [18] N. Courty, T. Burger, and P.-F. Marteau. Geodesic analysis on the Gaussian RKHS hypersphere. In *Machine Learning and Knowledge Discovery in Databases*, volume 7523 of *Lecture Notes in Computer Science*, pages 299–313. Springer, 2012. [6](#)
- [19] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 886–893, 2005. [7](#)
- [20] A. Feragen, S. Hauberg, M. Nielsen, and F. Lauze. Means in spaces of tree-like shapes. In *IEEE International Conference on Computer Vision (ICCV)*, pages 736–746, 2011. [1](#)
- [21] A. Feragen, N. Kasenburg, J. Petersen, M. de Bruijne, and K. Borgwardt. Scalable kernels for graphs with continuous attributes. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 216–224, 2013. [6](#)
- [22] A. Feragen, P. Lo, M. de Bruijne, M. Nielsen, and F. Lauze. Toward a theory of statistical tree-shape analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(8):2008–2021, 2013. [1](#)
- [23] A. Feragen, M. Owen, J. Petersen, M. Wille, L. Thomsen, A. Dirksen, and M. de Bruijne. Tree-space statistics and approximations for large-scale analysis of anatomical trees. In *Information Processing in Medical Imaging*, volume 7917 of *Lecture Notes in Computer Science*, pages 74–85. Springer, 2013. [1](#)
- [24] P. T. Fletcher. Geodesic regression and the theory of least squares on Riemannian manifolds. *International Journal of Computer Vision*, 105(2):171–185, 2013. [8](#)
- [25] P. T. Fletcher and S. C. Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis, ECCV 2004 Workshops CVAMIA and MMBIA, Prague, Czech Republic, May 15, 2004, Revised Selected Papers*, pages 87–98, 2004. [1](#)
- [26] O. Freifeld and M. Black. Lie bodies: A manifold representation of 3D human shape. In A. Fitzgibbon et al. (Eds.), editor, *European Conference on Computer Vision (ECCV)*, Part I, LNCS 7572, pages 1–14. Springer-Verlag, 2012. [1](#)
- [27] O. Freifeld, S. Hauberg, and M. J. Black. Model transport: Towards scalable transfer learning on manifolds. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 1378–1385, 2014. [8](#)
- [28] O. Gil-Medrano and P. W. Michor. The Riemannian manifold of all Riemannian metrics. *Quarterly Journal of Mathematics*, 42(2):183–202, 1991. [1](#)
- [29] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2066–2073, 2012. [7](#), [8](#)
- [30] S. Hauberg, A. Feragen, and M. J. Black. Grassmann averages for scalable robust PCA. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. [8](#)

- [31] S. Hauberg, O. Freifeld, and M. Black. A geometric take on metric learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2033–2041, 2012. [1](#)
- [32] S. Hauberg, S. Sommer, and K. S. Pedersen. Natural metrics and least-committed priors for articulated tracking. *Image and Vision Computing*, 30(6-7):453–461, 2012. [1](#)
- [33] J. Hinkle, P. Fletcher, and S. Joshi. Intrinsic polynomials for regression on Riemannian manifolds. *Journal of Mathematical Imaging and Vision*, 50(1-2):32–52, 2014. [1](#), [8](#)
- [34] P. Honeine and C. Richard. The angular kernel in machine learning for hyperspectral data classification. In *Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS), 2010 2nd Workshop on*, pages 1–4. IEEE, 2010. [6](#)
- [35] J. Istas. Manifold indexed fractional fields. *ESAIM: Probability and Statistics*, 16:222–276, 1 2012. [3](#), [5](#), [7](#)
- [36] T. Jaakkola and D. Haussler. Exploiting generative models in discriminative classifiers. In *Advances in Neural Information Processing Systems (NIPS)*, pages 487–493, 1998. [6](#)
- [37] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel methods on the Riemannian manifold of symmetric positive definite matrices. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 73–80, 2013. [6](#)
- [38] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Optimizing over radial kernels on compact manifolds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3802–3809, 2014. [6](#)
- [39] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel Methods on Riemannian Manifolds with Gaussian RBF Kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2015. [3](#)
- [40] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. T. Harandi. Combining multiple manifold-valued descriptors for improved object recognition. In *International Conference on Digital Image Computing: Techniques and Applications, DICTA*, pages 1–6, 2013. [6](#)
- [41] D. Kendall. Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16(2):81–121, 1984. [1](#)
- [42] F. J. Király, A. Ziehe, and K. Müller. Learning with algebraic invariances, and the invariant kernel trick. *CoRR*, abs/1411.7817, 2014. [1](#)
- [43] J. D. Lafferty and G. Lebanon. Diffusion kernels on statistical manifolds. *Journal of Machine Learning Research (JMLR)*, 6:129–163, 2005. [8](#)
- [44] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [8](#)
- [45] J. M. Lee. *Riemannian manifolds: An introduction to curvature*. Graduate Texts in mathematics. Springer, 1997. [4](#), [7](#)
- [46] C. Lenglet, R. Deriche, and O. Faugeras. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In *European Conference on Computer Vision (ECCV)*, volume 3024 of *Lecture Notes in Computer Science*, pages 127–140, 2004. [1](#)
- [47] R. Murray, Z. Li, and S. Sastry. *A mathematical introduction to robotic manipulation*. CRC, 1994. [1](#)
- [48] J. Nash. The imbedding problem for Riemannian manifolds. *Annals of Mathematics*, 63:20–63, 1956. [7](#)
- [49] X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *International Journal of Computer Vision (IJCV)*, 66(1):41–66, 2006. [1](#), [6](#), [7](#)
- [50] F. Porikli, O. Tuzel, and P. Meer. Covariance tracking using model update based on Lie algebra. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 728–735, 2006. [1](#)
- [51] A. Robertson. Crofton formulae and geodesic distance in hyperbolic spaces. *Journal of Lie Theory*, 8:163–172, 1998. [6](#), [7](#), [8](#)
- [52] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. [6](#)
- [53] B. Schölkopf. The kernel trick for distances. In *Advances in Neural Information Processing Systems (NIPS)*, pages 301–307, 2001. [6](#)
- [54] B. Schölkopf and A. Smola. *Learning with kernels : support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning. MIT Press, 2002. [1](#), [2](#), [3](#), [6](#), [7](#)
- [55] I. J. Schönberg. Metric spaces and completely monotone functions. *Annals of Mathematics*, 39(4):811–841, 1938. [4](#), [7](#)
- [56] A. Srivastava, S. H. Joshi, W. Mio, and X. Liu. Statistical shape analysis: Clustering, learning, and testing. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 27(4):590–602, 2005. [1](#)
- [57] F. Steinke, M. Hein, and B. Schölkopf. Nonparametric regression between general Riemannian manifolds. *SIAM J. Imaging Sciences*, 3(3):527–563, 2010. [8](#)
- [58] S. Taheri, P. Turaga, and R. Chellappa. Towards view-invariant expression analysis using analytic shape manifolds. In *IEEE International Conference on Automatic Face & Gesture Recognition*, pages 306–313, 2011. [1](#)
- [59] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319, 2000. [6](#)
- [60] A. Tosi, S. Hauberg, A. Vellido, and N. D. Lawrence. Metrics for probabilistic geometries. In *The Conference on Uncertainty in Artificial Intelligence (UAI)*, Quebec, Canada, July 2014. [1](#)
- [61] P. K. Turaga, A. Veeraraghavan, A. Srivastava, and R. Chellappa. Statistical computations on Grassmann and Stiefel manifolds for image and video-based recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(11):2273–2286, 2011. [1](#)
- [62] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. *European Conference on Computer Vision (ECCV)*, pages 589–600, 2006. [1](#)
- [63] A. Valette. Les représentations uniformément bornées associées un arbre rel. *Bulletin of the Belgian Mathematical Society*, 42:747–760, 1990. [7](#)
- [64] Y. Wang, H. Salehian, G. Cheng, and B. C. Vemuri. Tracking on the product manifold of shape and orientation for tractography from diffusion MRI. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3051–3056. IEEE, 2014. [1](#)
- [65] Q. Xie, S. Kurtsek, H. Le, and A. Srivastava. Parallel transport of deformations in shape space of elastic surfaces. In *International Conference on Computer Vision (ICCV)*, pages 865–872, Dec 2013. [8](#)

- [66] L. Younes. Spaces and manifolds of shapes in computer vision: An overview. *Image and Vision Computing*, 30(6):389–397, 2012. [1](#)