

Intrinsic Topological Transforms via the Distance Kernel Embedding

Clément Maria

INRIA Sophia Antipolis-Méditerranée, Valbonne, France
clement.maria@inria.fr

Steve Oudot

INRIA Saclay, Palaiseau, France
steve.oudot@inria.fr

Elchanan Solomon

Department of Mathematics, Duke University, Durham, NC USA
yitzchak.solomon@duke.edu

Abstract

Topological transforms are parametrized families of topological invariants, which, by analogy with transforms in signal processing, are much more discriminative than single measurements. The first two topological transforms to be defined were the Persistent Homology Transform (PHT) and Euler Characteristic Transform (ECT), both of which apply to shapes embedded in Euclidean space. The contribution of this paper is to define topological transforms for abstract metric measure spaces. Our proposed pipeline is to pre-compose the PHT or ECT with a Euclidean embedding derived from the eigenfunctions and eigenvalues of an integral operator. To that end, we define and study an integral operator called the distance kernel operator, and demonstrate that it gives rise to stable and quasi-injective topological transforms. We conclude with some numerical experiments, wherein we compute and compare the eigenfunctions and eigenvalues of our operator across a range of standard 2- and 3-manifolds.

2012 ACM Subject Classification Mathematics of computing → Algebraic topology

Keywords and phrases Topological Transforms, Persistent Homology, Inverse Problems, Spectral Geometry, Algebraic Topology, Topological Data Analysis

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.56

Related Version <https://arxiv.org/abs/1912.02225>

1 Introduction

One way of viewing the success of convolutional neural networks in the classification and analysis of large, multi-channel images is that these neural networks learn an optimal set of coordinates for representing sets of images in Euclidean space. However, when the data consist of shapes whose underlying topologies may vary, it becomes apparent that new tools and techniques must be brought to bear. One approach is to use topological transforms to represent these shapes as collections of topological summaries.

The topological summaries we use are persistence diagrams. Given a real-valued filter function f on a space X , the associated persistence diagram $\text{Diag}(f)$ describes how the topology of the sublevel sets $X_\alpha = \{f(x) \leq \alpha \mid x \in X\}$ evolves as α increases. When the filter-function f measures something about the geometry of X , such as its curvature, the resulting persistence diagram contains both topological and geometric information. If we use not one, but a family of filter functions, the resulting collection of persistence diagrams is called a topological transform. Although the space of persistence diagrams (with the Bottleneck of Wasserstein distances) is not an inner product space, there are many methods



© Clément Maria, Steve Oudot, and Elchanan Solomon;
licensed under Creative Commons License CC-BY

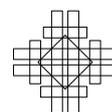
36th International Symposium on Computational Geometry (SoCG 2020).

Editors: Sergio Cabello and Danny Z. Chen; Article No. 56; pp. 56:1–56:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



for mapping persistence diagrams to a Hilbert space. The composition of any such map with our topological transform results in a collection of vectors that can be concatenated and fed into any standard machine learning model.

Turner et al. [20] first introduced topological transforms for shapes embedded in Euclidean space. Their work, and the work in various subsequent papers, demonstrates that this transform completely captures the geometric structure of its defining shape. Our goal in this article is to extend these topological transforms to intrinsic metric spaces by pre-composing them with a suitable Euclidean embedding. For this pipeline to be successful in practice, the chosen Euclidean embedding must preserve the geometric structure of our metric spaces. In this paper, we make the case that a particular integral operator, the distance kernel operator, is well suited to this task. We show that the eigenfunctions and eigenvalues of this operator are stable with respect to discretization and perturbations of the underlying shape, that they define a Euclidean embedding which encodes the large-scale geometry of our metric space, and that the resulting topological transforms enjoys favorable stability and quasi-injectivity properties.

Related Work. In [20], Turner et al. defined the Persistent Homology Transform (PHT) and Euler Characteristic Transform (ECT). These transforms take as input sufficiently regular subsets S of Euclidean space \mathbb{R}^d , and associate to every vector $v \in \mathbb{S}^{d-1}$ of the sphere in \mathbb{R}^d the persistence diagram, or Euler characteristic curve, of the sublevel-set filtration of S induced by the function $f_v : S \rightarrow \mathbb{R}: f_v(x) = v \cdot x$. It was subsequently proven in [9] and [12] that these topological transforms are injective in all dimensions¹. Moreover, it was shown in [5] and [9] that, for certain families of embedded shapes, these topological transforms can be computed in finitely many steps². Complimenting these theoretical results, Crawford et al. [8] demonstrated how to use these topological transforms to build an improved classifier for glioblastoma patient outcomes.

In [16], Oudot and Solomon defined a topological transform for intrinsic metric spaces (X, d_X) . This transform associates to every basepoint $x_0 \in X$ the extended persistence diagram of the function $f_{x_0} : X \rightarrow \mathbb{R}: f_{x_0}(x) = d_X(x_0, x)$. The resulting invariant, called the Intrinsic Persistent Homology Transform (IPHT), is the collection of all persistence diagrams arising from basepoints in X . By computing Euler characteristic curves instead of persistence diagrams, one obtains the Intrinsic Euler Characteristic Transform (IECT). This invariant was first studied, in the case of metric graphs, by Dey, Shi, and Wang in [10], where they proved stability and computability results and ran some experiments. The main result of [16] demonstrated that these invariants are injective³ on an appropriately generic subset of the space of metric graphs. For a detailed survey on related problems in applied topology, we refer the reader to [17]. Another line of research, at the intersection between persistent homology and spectral geometry, can be found in the work of Polterovich et al. [18], where they study and bound various functionals on persistence diagrams arising from Laplacian eigenfunctions on compact surfaces.

As this paper is concerned with both applied topology and spectral geometry, let us now consider some results, both classical and modern, in the latter field. To begin, the data of a weighted graph can be encoded via its adjacency matrix, and the spectral theory of

¹ By “injective”, we mean that two subsets of Euclidean space have the same transform if and only if they are identical. Thus, the transform is injective on the space of admissible subsets.

² That is, finitely many directions determine the entire transform, and these directions can be identified with finitely many geometric computations.

³ As with the PHT and ECT, this means that two graphs having the same transform must be isometric.

these matrices is deep and of great utility, seeing application in, e.g., graph clustering and Google's PageRank algorithm. Another matrix associated to a graph is its Laplacian, whose eigendecomposition forms the basis for the Laplacian Eigenmaps technique studied by Belkin and Niyogi in [2, 3, 4], as well as the diffusion maps of Coifman and Lafon [7]. Spectral analysis of the Gram matrix of the distances gives rise to the classical Multi-Dimensional Scaling embedding and its extension by Tenenbaum et al. [19] to non-linear embeddings: IsoMap. Lastly, the X-ray transform of [14] takes as input a continuous, compactly supported function f on \mathbb{R}^d , and outputs a function on the space of lines in \mathbb{R}^d that encodes the corresponding line integral of f .

Contributions and Outline. The structure of this article is as follows. In Section 2, we introduce a general framework for producing topological transforms on intrinsic metric objects via embedding these shapes in Euclidean space, and then applying the extrinsic topological transforms of [20]. To that end, we are charged with identifying a Euclidean embedding whose associated topological transforms have desirable stability and inverse properties. We observe that the eigenfunctions of the Laplacian are not well suited to this task, and so, in Section 3, we define a new operator, called the distance kernel operator, which we prove gives rise to a Euclidean embedding (which we call distance kernel embedding, or DKE). In Section 4, we show that the DKE is sufficiently regular to allow for the computation of topological invariants, in addition to being stable under discrete sampling and perturbation of the metric. We also show that the regularity of the DKE implies stability results for its associated topological transforms. In Section 5, we prove inverse results for the DKE and its topological transforms. We conclude, in Section 6, with a range of experiments illustrating the discriminative power of the DKE for discrete samples of various 2- and 3-manifolds.

2 Intrinsic Topological Transforms from Compact Operators

In this section we introduce a general framework for combining the existing extrinsic topological transforms with Euclidean embeddings of intrinsic metric spaces via compact operators.

2.1 Extrinsic Topological Transforms

► **Definition 1.** Let f be a real-valued function on a topological space X . We write $PH(X, f)$ to denote the graded sublevel set persistence diagram of (X, f) , which contains the sublevel set persistence diagrams of (X, f) for each homological degree. We write **GrDiag** to refer to the space of such graded persistence diagrams. For a fixed degree k , we write $\beta_k(X, f)$ to denote the corresponding Betti curve, which is the sum of the indicator functions of intervals in the persistence diagram. Lastly, we write $\chi(X, f)$ to denote the Euler curve, which is the alternating sum of the Betti curves in all degrees.

► **Definition 2.** Let \mathbb{S}^k be the k -dimensional sphere, and $L(\mathbb{R}^{k+1}, \mathbb{R})$ the space of linear maps from \mathbb{R}^{k+1} to \mathbb{R} . Define the map $\Theta : \mathbb{S}^k \rightarrow L(\mathbb{R}^{k+1}, \mathbb{R})$ which sends $v \in \mathbb{S}^k$ to the map $x \mapsto \langle x, v \rangle$.

► **Definition 3** ([9]). Let $X \subset \mathbb{R}^d$ be a compact, definable set⁴. For every $v \in \mathbb{S}^{d-1}$, the sublevel set persistence diagram and Euler curve of the pair $(X, \Theta(v))$ exist. The Persistent Homology Transform is the map $PHT(X) : \mathbb{S}^{d-1} \rightarrow \mathbf{GrDiag}$ defined by:

$$PHT(X)(v) = PH(X, \Theta(v)).$$

If one computes Euler curves instead of persistence diagrams, one obtains the Euler Characteristic Transform $ECT(X)$.

Intuitively, the PHT and ECT probe an embedded subset of Euclidean space like a multi-directional MRI scanner, recording how the topology evolves along each direction. The following injectivity result demonstrates the rich geometric content of these transforms.

► **Theorem 4** ([9, 12]). The PHT and ECT are injective for all k . That is, for any definable sets $X, Y \subset \mathbb{R}^d$, if $PHT(X) = PHT(Y)$ or $ECT(X) = ECT(Y)$, then $X = Y$ as sets.

2.2 Compact Operators and their Embeddings

► **Definition 5.** Let H be a Hilbert space. A linear operator $T : H \rightarrow H$ is bounded if there exists a constant M such that, for all $v \in H$, $\|Tv\| \leq M\|v\|$. A bounded operator is compact if the image of any bounded subset of H under T is relatively compact. A bounded operator T is self-adjoint if T is equal to its adjoint T^* ; equivalently, $\langle Tx, y \rangle = \langle x, Ty \rangle$ for all $x, y \in H$.

The spectral theorem for compact, self-adjoint operators on a Hilbert space asserts that these operators can be diagonalized.

► **Theorem 6** (Spectral Theorem). Let T be a compact, self-adjoint operator on a Hilbert space H . Then H admits a finite or countably infinite basis $\{\phi_i\}$ of eigenvectors of T with real eigenvalues $\{\lambda_i\}$, where $\lim_{i \rightarrow \infty} \lambda_i = 0$.

Given a compact metric (Borel) measure space (X, d_X, μ_X) , we can consider compact, self-adjoint operators on the Hilbert space $L^2(X)$. The eigenfunctions $\{\phi_i\}$ and eigenvalues $\{\lambda_i\}$ arising from the spectral theorem can then be used to define embeddings of X into Euclidean space. This requires the adoption of various conventions and generic assumptions:

1. The spectral theorem asserts the existence of the eigenfunctions ϕ_i , but it does not guarantee their uniqueness. Indeed, the choice is never unique. If the eigenvalue λ_i has geometric multiplicity one, then there are two choices of unit norm eigenfunctions: $\{\phi_i, -\phi_i\}$. If the eigenvalue has geometric multiplicity greater than one, then there are infinitely many choices. In the rest of the paper, we make the generic assumption that all the eigenvalues have multiplicity one⁵.
2. We adopt the convention of dropping the eigenfunctions in the zero-eigenspace, and, to fix the choice of sign, we pick ϕ_i such that $\langle \phi_i, |\phi_i| \rangle > 0$ for all i ⁶.
3. We order the eigenvalues (and hence eigenfunctions) in decreasing order of absolute value, $|\lambda_1| \geq |\lambda_2| \geq |\lambda_3| \geq \dots$, breaking the tie between positive-negative pairs by listing the positive eigenvalue first.

⁴ The notion of definability is always understood to be relative to a choice of o-minimal structure on \mathbb{R}^d , which is an algebra of sets satisfying certain membership conditions. Examples include the collection of semi-algebraic or analytic subsets of \mathbb{R}^d . See [9] §2 for details.

⁵ We can always infinitesimally perturb our space to make this true.

⁶ We make the generic assumption that this dot product is nonzero.

4. To take advantage of many useful results in operator theory, we restrict ourselves to operators that are *Hilbert-Schmidt*, which means that $\sum_{i=1}^{\infty} \lambda_i^2 < \infty$. Every Hilbert-Schmidt operator on $L^2(X)$ can be represented as an integral operator with square integrable kernel $K(\cdot, \cdot)$ ⁷, i.e., an operator of the form:

$$T: L^2(X) \rightarrow L^2(X) \quad (Tf)(x) = \int_X K(x, y)f(y)d\mu_X(y).$$

We thus assume our operators are of this form.

► **Definition 7.** *Given a compact metric (Borel) measure space (X, d_X, μ_X) , let T be a compact, self-adjoint operator on $L^2(X)$ with spectral decomposition $\{\phi_i, \lambda_i\}$, following the conventions above. We define coordinate functions on X as follows: $\alpha_i(x) = \sqrt{\lambda_i}\phi_i(x)$. Note that the eigenvalue λ_i may be negative (we have not assumed that the operator is positive definite), so the coordinate function takes values in \mathbb{C} . When λ_i is negative, we adopt the convention of taking the square root with positive imaginary part. By identifying \mathbb{C} with \mathbb{R}^2 , we can also think of this coordinate function as taking a pair of real values.*

Our rationale for scaling the eigenfunctions by the square root of their eigenvalues is that, for an integral operator T with kernel $K(\cdot, \cdot)$, the sum $\sum_{k=1}^{\infty} \lambda_k \phi_k(x)\phi_k(x') = \sum_{k=1}^{\infty} (\sqrt{\lambda_k}\phi_k(x))(\sqrt{\lambda_k}\phi_k(x')) = \sum_{i=1}^{\infty} \alpha_i(x)\alpha_i(x')$ converges to $K(x, x')$ in $L^2(X, X)$ ⁸. Using these coordinates, we define a kernel embedding⁹:

► **Definition 8.** *Let (X, d_X, μ_X) , T , and $\{\phi_i, \lambda_i\}$ be as in Definition 7. For $k \geq 1$, we define $\Phi_k : X \rightarrow \mathbb{C}^k \cong \mathbb{R}^{2k}$ to be the map sending a point $x \in X$ to $(\alpha_1(x), \dots, \alpha_k(x)) \in \mathbb{C}^k \cong \mathbb{R}^{2k}$. Setting $k = \infty$ gives us a map $\Phi : X \rightarrow \mathbb{C}^{\infty} \cong \mathbb{R}^{\infty}$. When Φ is continuous and injective, the image of Φ (resp. Φ_k) is called the kernel embedding (resp. truncated kernel embedding) associated to T .*

2.3 Topological Kernel Transforms

By post-composing this embedding with the PHT or the ECT, we obtain topological transforms that are defined intrinsically on metric measure spaces¹⁰.

► **Definition 9.** *Let X and Φ be as in Definition 8. For k finite, the embedded persistence kernel transform $e-PKT_k(X)$ is the PHT applied to the image of the embedding $\Phi_k(X) \subset \mathbb{R}^{2k}$, which takes as input vectors in \mathbb{S}^{2k-1} and takes values in **GrDiag**. Using Euler curves in place of persistent diagrams gives rise to the embedded Euler kernel transform $e-EKT_k$.*

The following meta-theorem motivates the constructions and results to follow.

► **Theorem 10.** *Fix a positive integer k . Let \mathcal{M} be a class of metric measure spaces with integral kernels $\{K^M\}_{M \in \mathcal{M}}$, giving definable embeddings $\{\Phi_k^M\}_{M \in \mathcal{M}}$. If $\Phi_k^M(M) \neq \Phi_k^{M'}(M')$ for any pair of non-isometric spaces $M \neq M' \in \mathcal{M}$, then the $e-PKT_k$ and $e-EKT_k$ are injective on \mathcal{M} .*

Proof. This is an immediate consequence of Theorem 4. ◀

⁷ This is the Hilbert-Schmidt Kernel Theorem. See Appendix B of [13], where the result is proven for subsets of \mathbb{R} . The proof for compact metric (Borel) measure spaces is identical.

⁸ This convergence is not necessarily uniform, unless $K(\cdot, \cdot)$ is positive semidefinite (Mercer's theorem).

⁹ The use of the term *embedding* comes from the injectivity properties of this map, proved in §5.1

¹⁰ We implicitly assume here that these persistence diagrams and Euler curves exist. Later on in this article, we verify this explicitly for the integral operator of interest.

For most classes \mathcal{M} of interest, such as the set of Riemannian manifolds of a given dimension, there are no known integral operators whose associated embeddings in finite dimensions are injective in the sense of Theorem 10. However, if we relax the assertion of injectivity, and ask only that $\Phi_k^M(M) = \Phi_k^{M'}(M')$ implies a bound on the Gromov-Hausdorff distance between M and M' , we can construct such an integral operator. Before defining this operator, which is the subject of the next section, we note that the diffusion operator, which is related to the *local* geometry of a space, does not enjoy *global* guarantees of this kind. Indeed, although one can recover the metric on a Riemannian manifold from the exact knowledge of all its Laplacian eigenfunctions and eigenvalues (cf. [21]), this reconstruction is asymptotic, cannot be approximated with finitely many eigenfunctions and eigenvalues, and is unstable to noise.

3 The Distance Kernel Operator and its Embedding

We now define our proposed integral operator, and prove that it is compact and self-adjoint. Recall that a compact metric measure space (X, d_X, μ_X) has finite volume $\mu_X(X) < +\infty$. We write $\text{Vol}(X) := \mu_X(X)$.

► **Definition 11.** Let (X, d_X, μ_X) be a compact metric measure space¹¹. We define the following operator D^X on $L^2(X)$, called the distance kernel operator (DKO):

$$(D^X f)(x) = \int_X f(y) d_X(x, y) d\mu_X(y).$$

► **Proposition 12.** D^X is a self-adjoint operator.

Proof. By convention, μ_X is Radon. Since X is compact, this implies that $\mu_X(X) < \infty$, and hence (X, μ_X) is σ -finite. We can thus apply Fubini's theorem, and the symmetry of the distance function d_X , to observe that, for two integrable functions f and g ,

$$\begin{aligned} \langle D^X f, g \rangle &= \int_X \left(\int_X f(y) d_X(x, y) d\mu_X(y) \right) g(x) d\mu_X(x) \\ &= \int_X \int_X f(y) g(x) d_X(x, y) d\mu_X(x) d\mu_X(y) \\ &= \int_X f(y) \left(\int_X g(x) d_X(y, x) d\mu_X(x) \right) d\mu_X(y) = \langle f, D^X g \rangle, \end{aligned}$$

demonstrating self-adjointness. ◀

► **Proposition 13.** D^X is a compact operator.

Proof. Let $f_n \in L^2(X)$ be a bounded sequence of functions, $\|f_n\|_{L^2} \leq C$ for all n .

For all $d_X(x, x') \leq \varepsilon$ and all n ,

$$\begin{aligned} |D^X f_n(x) - D^X f_n(x')| &= \left| \int_X (d_X(x, y) f_n(y) - d_X(x', y) f_n(y)) d\mu_X(y) \right| \\ &\leq \int_X |d_X(x, y) - d_X(x', y)| |f_n(y)| d\mu_X(y) \\ &\text{(Cauchy-Schwarz)} \leq \|d_X(x, y) - d_X(x', y)\|_{L^2(y)} \cdot \|f_n(y)\|_{L^2(y)} \\ &\text{(triangle inequality)} \leq \|\varepsilon\|_{L^2(y)} \cdot \|f_n(y)\|_{L^2(y)} \\ &\leq \varepsilon \sqrt{\text{Vol}(X)} \cdot C, \end{aligned}$$

¹¹ For the rest of the paper, we assume that μ_X is a Radon measure.

where $\text{Vol}(X)$ is finite. Thus, $D^X f_n$ is an equicontinuous family of functions on X , so, by the Arzelà-Ascoli theorem, it contains a uniformly convergent, and hence L^2 -convergent, subsequence. This demonstrates compactness. ◀

Note that, as a consequence of the proof of this proposition, the eigenfunctions of the distance kernel operator are always continuous. We can thus define an embedding as in Definition 8, which we call the *distance kernel embedding* (DKE)¹². From now on, we write Φ and Φ_k to exclusively denote the distance kernel embedding, and the e-PKT and e-EKT likewise refer exclusively to the resulting transforms. In the following sections, we study the stability and inverse properties of these embeddings and transforms.

4 Stability for the DKE and its Topological Transforms

Let (X, d_X, μ_X) be a compact metric (Borel) measure space. The eigenfunctions of the distance kernel operator with nonzero eigenvalue are Lipschitz continuous, with the Lipschitz constant being inversely proportional to the absolute value of the eigenvalue.

▶ **Lemma 14.** *For every $i \in \mathbb{N}_{>0}$, The function $\lambda_i \phi_i$ is $\sqrt{\text{Vol}(X)}$ -Lipschitz. Hence, if $\lambda_i \neq 0$, ϕ_i is $(\sqrt{\text{Vol}(X)} / |\lambda_i|)$ -Lipschitz. Note that X being compact, $\text{Vol}(X) < +\infty$ and these Lipschitz constants are indeed finite.*

Proof. Let $x, y \in X$ and $\varepsilon = d_X(x, y)$. By the fact that $\lambda_i \phi_i = D^X \phi_i$, we have

$$\begin{aligned} |\lambda_i \phi_i(x) - \lambda_i \phi_i(y)|^2 &= |(D^X \phi_i)(x) - (D^X \phi_i)(y)|^2 \\ &= \left| \int_X (d_X(x, z) - d_X(y, z)) \phi_i(z) d\mu_X(z) \right|^2 \\ \text{(Cauchy-Schwarz)} &\leq \int_X \underbrace{(d_X(x, z) - d_X(y, z))^2}_{\leq d_X(x, y) = \varepsilon} d\mu_X(z) \cdot \underbrace{\int_X \phi_i^2(z) d\mu_X(z)}_{=1} \\ &\leq \varepsilon^2 \text{Vol}(X). \end{aligned}$$

Thus, $|\lambda_i \phi_i(x) - \lambda_i \phi_i(y)| \leq \varepsilon \sqrt{\text{Vol}(X)}$, so $\lambda_i \phi_i$ is $\sqrt{\text{Vol}(X)}$ -Lipschitz. ◀

This regularity result on eigenfunctions has many implications for our topological transforms, which are given below (the proofs can be found in the full version of the paper [15], but are omitted here due to their complexity and length). The result implies in particular that persistence diagrams exist and are well-defined (which is not the case for an arbitrary continuous function on a compact topological space), and, under the additional assumption that the space X implies bounded degree- q total persistence¹³, that Euler curves exist:

▶ **Proposition 15.** *Let (X, d_X, μ_X) be a compact metric measure space homeomorphic to the geometric realization of a finite simplicial complex. Then, any finite linear combination $f = \sum_{i=1}^n c_i \phi_i$ of eigenfunctions of the distance kernel D^X has a well-defined sublevel set graded persistence diagram $PH(X, f)$. Now, suppose further that X implies bounded degree- q total persistence. Let $p = 1/q$. Then for any homological degree k , the sum defining $\beta_k(X, f)$ converges in L^p . Moreover, the sum defining $\chi(X, f)$ is finite, so the Euler curve exists as a function in L^p .*

¹²The matter of injectivity will be established in Lemma 20.

¹³Intuitively, this technical condition means that, for any graded persistence diagram $PH(X, f)$, the sum of the q th powers of the persistences of the points across all degrees is finite. The bilipschitz image of a finite dimensional Euclidean simplicial complex has bounded degree- q total persistence for q sufficiently large. See the full version of this paper for details.

In addition, we obtain the following stability result for the resulting topological transforms:

► **Theorem 16.** *Suppose X is homeomorphic to the geometric realization of a finite simplicial complex. If we equip **GrDiag** with the graded bottleneck distance and the sphere \mathbb{S}^{2k-1} with the ℓ^1 distance, then the $e\text{-PKT}_k$ is Lipschitz continuous. Now suppose further that X implies bounded degree- q total persistence for some $q > 0$, and that there is a uniform bound on the number of points in the persistence diagrams obtained when evaluating the $e\text{-PKT}_k$ at an arbitrary vector $v \in \mathbb{S}^{2k-1}$. If we equip the sphere \mathbb{S}^{2k-1} with the ℓ^1 distance, and the space of Euler curves with the $L^{1/q}$ distance, then the $e\text{-EKT}_k$ is q -Hölder continuous on \mathbb{S}^{2k-1} .*

We also have two more stability results for the DKE, for which we do not yet have analogues for the topological transforms. The first result asserts that the distance kernel embedding of a discrete sample of a space converges almost surely to the distance kernel embedding of the underlying space:

► **Theorem 17.** *Let (X, d_X, μ_X) be a compact metric measure space with (a, b) -standard Borel measure¹⁴. For an i.i.d sample X_n of X of size n , call $\hat{\Phi}_k(X_n)$ the empirical DKE defined on the metric measure space (X_n, d_X, μ_n) with uniform measure $\mu_n(\hat{x}) = \mu_X(X)/n$ for all $\hat{x} \in X_n$. Writing $d_H^{L^2}$ for the Hausdorff distance for the L^2 norm in \mathbb{C}^k , we have:*

$$d_H^{L^2}(\Phi_k(X), \hat{\Phi}_k(X_n)) \xrightarrow{a.s.} 0 \text{ as } n \rightarrow +\infty.$$

In addition, the distance kernel embedding is stable on the space of Riemannian manifolds. The following is a simplified version of the result contained in the full version of the paper, which gives an explicit form for the function F .

► **Theorem 18.** *Let (X, d_X, μ_X) and (Y, d_Y, μ_Y) be compact finite-dimensional Riemannian manifolds equipped with their volume measures, such that $\mu_X(X) = \mu_Y(Y)$. Let $\varepsilon = d_{G\bar{P}}(X, Y)$ be the modified Gromov-Prokhorov distance¹⁵ between X and Y , and let $|\lambda_1| > \dots > |\lambda_k| > 0$ and $|\nu_1| > \dots > |\nu_k| > 0$ be the k largest (in absolute value) eigenvalues of the distance kernel operators of X and Y respectively, all non-zero with distinct absolute values. Let $\Phi_k(X)$ and $\Psi_k(Y)$ be the induced DKE. If we write Λ for the set $\{\lambda_1, \dots, \lambda_k, \nu_1, \dots, \nu_k\}$, then there is a function $F(\Lambda, \varepsilon)$, depending on the magnitude of the elements of Λ and the gaps $|\lambda_i^2 - \nu_i^2|$ between corresponding eigenvalues, such that:*

$$d_H^{L^2}(\Phi_k(X), \Psi_k(Y)) \leq F(\Lambda, \varepsilon) \text{ and } \lim_{\varepsilon \rightarrow 0} F(\Lambda, \varepsilon) = 0.$$

5 Injectivity for the DKE and its Topological Transforms

We now demonstrate some inverse results for the distance kernel embedding and transforms. We stress that these results apply specifically when the integral kernel is taken to be $d_X(\cdot, \cdot)$.

5.1 Injectivity of Φ

Our first result, in Corollary 21 below, is that, under the mild hypothesis of strict positivity (defined below), the infinite-dimensional embedding Φ is a homeomorphism of the metric measure space onto its image in \mathbb{C}^∞ .

¹⁴This ensures a lower bound on the volume of metric balls. See the full paper for a precise definition.

¹⁵This is a slight modification of the Gromov-Prokhorov distance, introduced by Burago et al. in §8 of [6].

► **Definition 19.** For a topological space X equipped with its Borel σ -algebra, we call a measure μ_X strictly positive if the measure of any nonempty open set is strictly positive.

► **Lemma 20.** Let (X, d_X, μ_X) be a compact, strictly positive metric measure space. Then the map $\Phi : X \rightarrow \mathbb{C}^\infty$ is injective.

Proof. Suppose that there are $x \neq y \in X$ such that $\Phi(x) = \Phi(y)$. This implies that $\alpha_i(x) = \alpha_i(y)$ and, in turn, $\lambda_i \phi_i(x) = \lambda_i \phi_i(y)$ for all i . Let d_x and d_y be the distance functions associated to x and y respectively. Using the L^2 -convergence of the eigenfunction expansion, we know that:

$$\|d_x - \sum_{i=1}^n \lambda_i \phi_i(x) \phi_i\|_{L^2} \xrightarrow{n \rightarrow \infty} 0 \quad \text{and} \quad \|d_y - \sum_{i=1}^n \lambda_i \phi_i(y) \phi_i\|_{L^2} \xrightarrow{n \rightarrow \infty} 0.$$

Since $\sum_{i=1}^n \lambda_i \phi_i(x) \phi_i = \sum_{i=1}^n \lambda_i \phi_i(y) \phi_i$ for all n , the triangle inequality implies that $\|d_x - d_y\|_{L^2} = 0$. Let now $r = d_X(x, y)/3 > 0$, and let U be the open neighborhood of radius r around x . The function $|d_x - d_y|$ is bounded below by r on U , and since U is not empty (it contains x), it has strictly positive measure. This then implies $\|d_x - d_y\|_{L^2} > 0$, a contradiction. Thus, $\Phi(x) \neq \Phi(y)$ for $x \neq y$. ◀

► **Corollary 21.** By Lemma 14, every component of the map Φ is continuous. Meanwhile, any metric on \mathbb{C}^∞ gives it a Hausdorff topological structure. Thus, for any such choice of metric, Φ is a continuous injection from a compact space to a Hausdorff space. Hence, Φ is a homeomorphism.

We also have the following injectivity result, where the domain of interest is the space of compact metric measure spaces. As a consequence of Corollary 21, it suffices to consider pairs of metric measure spaces that are defined on a common topological space.

► **Theorem 22.** Fix a compact topological space Z . Let μ and μ' be strictly positive measures for the Borel σ -algebra on Z , with μ absolutely continuous with respect to μ' , and d and d' metrics on X , both consistent with the topology on Z , making $X = (Z, d, \mu)$ and $X' = (Z, d', \mu')$ metric measure spaces. If $\Phi(X) = \Phi(X')$, then $d = d'$.

Proof. By assuming that both metric measure spaces induce the same topology, we can work with a single σ -algebra: their common Borel σ -algebra. This will prove essential in the following proof, where we take various unions and complements of measurable sets for μ and μ' , respectively. Next, let D and D' be the integral operators with kernels d and d' , respectively. The equality $\Phi(X) = \Phi(X')$ implies that D and D' have the same scaled eigenfunctions α_i . The distance functions d, d' thus have the same eigenfunction expansion:

$$(x_1, x_2) \mapsto \sum_{i=1}^{\infty} \alpha_i(x_1) \alpha_i(x_2).$$

This converges to d in $L_2(\mu \otimes \mu)$ and to d' in $L_2(\mu' \otimes \mu')$ to d' . Let us denote by S_n the partial sums of this expansion:

$$S_n = \sum_{i=1}^n \alpha_i(x_1) \alpha_i(x_2).$$

It is a standard result in measure theory that any L^2 -convergent sequence admits a subsequence that converges pointwise a.e.¹⁶ Thus, one can extract a subsequence S_{n_k} that converges to d pointwise on $(Z \times Z) \setminus N_1$, for some set $N_1 \subset Z$ such that $(\mu \otimes \mu)(N_1) = 0$. We can then extract a further subsequence $S_{n_{k_j}}$ that converges pointwise to d' on $((Z \times Z) \setminus N_1) \setminus N_2$, where $(\mu' \otimes \mu')(N_2) = 0$. Since μ is absolutely continuous to μ' , if we set $N = N_1 \cup N_2$ then $(\mu \otimes \mu)(N) = 0$. Since μ is strictly positive, the set N cannot contain any open sets, hence N^c is dense in $Z \times Z$. We see then that $d = d'$ on a dense subset of $Z \times Z$; since these functions are both continuous in the same topology Z , they are equal everywhere. ◀

5.2 Quasi-Injectivity of Φ_k

While the truncated embedding Φ_k may not be injective, we can get control over the diameter of its fibers (Corollary 26). The bounds are expressed in terms of the error of the approximation of the metrics by its truncated expansion:

► **Definition 23.** For a compact metric measure space (X, d_X, μ_X) and a positive integer k , we define the error function $E_{X,k}$, which measures the pointwise distance between d_X and its truncated eigenfunction expansion:

$$E_{X,k}(x, x') = \left| \sum_{i=1}^k \alpha_i(x) \alpha_i(x') - d_X(x, x') \right|.$$

► **Theorem 24.** Let (X, d_X, μ_X) and (Y, d_Y, μ_Y) be compact metric measure spaces, with eigenvalues $\{\lambda_i\}$ and $\{\nu_i\}$. Let $k \in \mathbb{N}_{>0}$ be an integer, and $\varepsilon := d_H^{L^2}(\Phi_k(X), \Phi_k(Y))$. Then:

$$d_{GH}(X, Y) \leq 2\varepsilon \min \left\{ \max_{x \in X} \|\Phi_k(x)\|_2, \max_{y \in Y} \|\Phi_k(y)\|_2 \right\} + \|E_{X,k}\|_\infty + \|E_{Y,k}\|_\infty + \varepsilon^2.$$

In the special case where X and Y are finite metric measure spaces, Theorem 24 assumes a more precise form:

► **Theorem 25.** Let (X, d_X, μ_X) and (Y, d_Y, μ_Y) , with eigenvalues $\{\lambda_i\}$ and $\{\nu_i\}$, and let $\theta = \min\{\min_{x \in X} \mu_X(x), \min_{y \in Y} \mu_Y(y)\}$. Take $k \leq |X|, |Y|$, and suppose that $d_H^{L^2}(\Phi_k(X), \Phi_k(Y)) \leq \varepsilon$. Then,

$$d_{GH}(X, Y) \leq 2\varepsilon \frac{\min(\sqrt{|\lambda_1|}, \sqrt{|\nu_1|})}{\theta} + \varepsilon^2 + \frac{|\lambda_{k+1}| + |\nu_{k+1}|}{\theta}.$$

We thus obtain the following bound on the diameter of the fibers of the DKE.

► **Corollary 26.** Let (X, d_X, μ_X) and (Y, d_Y, μ_Y) be compact metric measure spaces, with eigenvalues $\{\lambda_i\}$ and $\{\nu_i\}$. Let $k \in \mathbb{N}$ be a positive integer, and suppose that $\Phi_k(X) = \Phi_k(Y)$. Then, $d_{GH}(X, Y) \leq \|E_{X,k}\|_\infty + \|E_{Y,k}\|_\infty$. If X and Y are finite metric spaces, and we set $\theta = \min\{\min_{x \in X} \mu_X(x), \min_{y \in Y} \mu_Y(y)\}$, then $d_{GH}(X, Y) \leq \frac{1}{\theta} (|\lambda_{k+1}| + |\nu_{k+1}|)$.

The effectiveness of these results depends on the magnitude of the quantities $\|E_{X,k}\|_\infty$ and the decay of the eigenvalues λ_i . It remains to identify general metric and measure-theoretic criteria that imply bounds on these spectral statistics. The rest of this section is devoted to

¹⁶ Chebyshev's inequality proves that L^2 convergence implies convergence in measure. See Theorem 2.15(c) in [11] for the implication that convergence in measure implies the existence of pointwise a.e. convergent subsequence.

the proof of Theorem 24. Theorem 25 follows from Theorem 24, and we refer the reader to the full version of the paper for the details of this derivation [15]. The proof of Theorem 24 requires us to define the following algebraic operation, and to prove some technical lemmas regarding it.

► **Definition 27.** For vectors $v, w \in \mathbb{C}^k$, define the following bilinear form:

$$[v, w] = \sum_{i=1}^k v_i w_i \in \mathbb{C}.$$

This form is symmetric but not a dot product.

The utility of the bilinear form $[\cdot, \cdot]$ comes from the following equality:

$$[\Phi_k(x), \Phi_k(x')] = \sum_{i=1}^k \alpha_i(x) \alpha_i(x') = \sum_{i=1}^k \left(\sqrt{\lambda_i} \phi_i(x) \right) \left(\sqrt{\lambda_i} \phi_i(x') \right) = \sum_{i=1}^k \lambda_i \phi_i(x) \phi_i(x').$$

That is, when applied to the distance kernel embedding, $[\cdot, \cdot]$ gives the first k terms of the eigenfunction expansion of the distance function d_X .

► **Lemma 28.** The bilinear form $[\cdot, \cdot]$ satisfies the following Cauchy-Schwarz inequality:

$$|[v, w]| \leq \|v\|_2 \|w\|_2.$$

Proof Sketch. By the triangle inequality for complex numbers, we have $|[v, w]| \leq [\tilde{v}, \tilde{w}] = \langle \tilde{v}, \tilde{w} \rangle$, where \tilde{v}, \tilde{w} are obtained from v and w by replacing each coordinate with its modulus. The result then follows by applying the standard Cauchy-Schwarz inequality. ◀

The following lemma asserts that pairs of nearby vectors have similar bilinear products.

► **Lemma 29.** Let $v_1, v_2, w_1, w_2 \in \mathbb{C}^k$ be such that $\|v_1 - w_1\|_2 \leq \varepsilon$ and $\|v_2 - w_2\|_2 \leq \varepsilon$. Then

$$|[v_1, v_2] - [w_1, w_2]| \leq \varepsilon \min \{ \|v_1\|_2 + \|v_2\|_2, \|w_1\|_2 + \|w_2\|_2 \} + \varepsilon^2.$$

Proof. By bilinearity,

$$[w_1, w_2] = [v_1, v_2] + [v_1, (w_2 - v_2)] + [(w_1 - v_1), v_2] + [(w_1 - v_1), (w_2 - v_2)].$$

Thus,

$$|[v_1, v_2] - [w_1, w_2]| \leq |[v_1, (w_2 - v_2)]| + |[(w_1 - v_1), v_2]| + |[(w_1 - v_1), (w_2 - v_2)]|.$$

By a symmetric argument, switching v_1 and v_2 with w_1 and w_2 , one obtains:

$$|[v_1, v_2] - [w_1, w_2]| \leq |[w_1, (v_2 - w_2)]| + |[(v_1 - w_1), w_2]| + |[(v_1 - w_1), (v_2 - w_2)]|.$$

The result then follows by applying the Cauchy-Schwarz inequality to each term on the right-hand sides of both inequalities, and by taking the minimum of the two sums. ◀

We can now prove Theorem 24:

Proof. Let C be an optimal Hausdorff correspondence between $\Phi_k(X)$ and $\Phi_k(Y)$. Let $(x, x') \in X \times X$ and $(y, y') \in Y \times Y$ with $(\Phi_k(x), \Phi_k(y)), (\Phi_k(x'), \Phi_k(y')) \in C$. Lemma 29, together with the bounds $\|\Phi_k(x) - \Phi_k(y)\|_{L^2} \leq \varepsilon$ and $\|\Phi_k(x') - \Phi_k(y')\|_{L^2} \leq \varepsilon$, gives

$$|[\Phi_k(x), \Phi_k(x')] - [\Phi_k(y), \Phi_k(y')]| \leq 2\varepsilon \min \left\{ \max_{x \in X} \|\Phi_k(x)\|_2, \max_{y \in Y} \|\Phi_k(y)\|_2 \right\} + \varepsilon^2.$$

Using the triangle inequality, we can replace $[\Phi_k(x), \Phi_k(x')]$ with $d_X(x, x')$ and $[\Phi_k(y), \Phi_k(y')]$ with $d_Y(y, y')$, at the cost of adding an additive error of at most $\|E_{X,k}\|_\infty$ and $\|E_{Y,k}\|_\infty$ respectively, giving the following inequality, from which the result follows:

$$|d_X(x, x') - d_Y(y, y')| \leq 2\varepsilon \min \left\{ \max_{x \in X} \|\Phi_k(x)\|_2, \max_{y \in Y} \|\Phi_k(y)\|_2 \right\} + \|E_{X,k}\|_\infty + \|E_{Y,k}\|_\infty + \varepsilon^2.$$

◀

5.3 Quasi-Injectivity of the $e\text{-PKT}_k$ and $e\text{-EKT}_k$

Corollary 26, taken together with Theorem 4, implies the following result, which bounds the diameter of the fibers of the topological transforms. For general metric measure spaces, the diameter depends on the error functions $E_{X,k}$ and $E_{Y,k}$. As k goes to infinity, these functions go to zero in the L^2 norm, but we do not have any general guarantees that this also holds in the L^∞ norm¹⁷. For finite metric spaces, the diameter does indeed go to 0 as k goes to infinity.

► **Theorem 30.** *Let (X, d_X, μ_X) and (Y, d_Y, μ_Y) be compact metric measure spaces, with eigenvalues $\{\lambda_i\}$ and $\{\nu_i\}$ respectively, giving rise to definable distance kernel embeddings. Let $k \in \mathbb{N}$ be a positive integer, and suppose that $e\text{-PKT}_k(X) = e\text{-PKT}_k(Y)$ or $e\text{-EKT}_k(X) = e\text{-EKT}_k(Y)$. Then $d_{GH}(X, Y) \leq \|E_{X,k}\|_\infty + \|E_{Y,k}\|_\infty$. If X and Y are finite spaces, and we set $\theta = \min\{\min_{x \in X} \mu_X(x), \min_{y \in Y} \mu_Y(y)\}$, then $d_{GH}(X, Y) \leq \frac{1}{\theta} (|\lambda_{k+1}| + |\nu_{k+1}|)$.*

The condition that the DKEs be definable is always satisfied when the spaces are finite. It remains to work out the correct hypotheses to ensure definability more generally; this is work in progress.

6 Experiments

The goal of this section is to illustrate the results of Sections 4 and 5. In the following experiments, we compute the DKE for a variety of discrete samples on the torus and 2-sphere, with metric induced by their embedding in Euclidean space, on the 3-sphere, and on the Lens spaces $L(7, 1)$ and $L(7, 4)$, with spherical geometry. The measures on these samples are uniform. These spaces have distinct integer homology, except for the two Lens spaces that have the same homotopy type but are not homeomorphic, and therefore not isometric. This makes $L(7, 1)$ and $L(7, 4)$ difficult to distinguish by purely topological methods. We see that the DKE (and, therefore, the resulting topological transforms) is capable of distinguishing these Lens spaces.

Spectra of various manifolds. In Figure 1a, we have plotted the first 8 eigenvalues of five discrete metric spaces, sampled from each of these five manifolds, normalized by the number of points in each sample. We can observe the following: (1) the two Lens spaces have relatively similar eigenvalues, (2) the 2- and 3-sphere have many similar eigenvalues, but their first and fourth eigenvalues are significantly different, and (3) the torus has the most distinct spectrum.

¹⁷ However, experimental studies suggest that this is the case for a variety of manifolds [15].

Spectra of Lens spaces for various samples. In Figure 1b, we compare the spectra of a number of different random i.i.d. samples of the two Lens spaces $L(7, 1)$ and $L(7, 4)$. To be precise, for each Lens space we compute the spectra of two distinct random samples with 2000 points, and a third sample with 5000 points. The spectra for the different samples of the same Lens space are virtually impossible to distinguish, and only two curves – one for the spectrum of $L(7, 1)$, and one of the spectrum of $L(7, 4)$ – are visible in Figure 1b. This attests to the stability of the eigenvalues of the distance kernel operator under random i.i.d. sampling, in line with Theorem 17. Notably, the two Lens spaces are distinguished by the first, third, and fourth eigenvalues of their distance kernel operators. $L(7, 1)$ and $L(7, 4)$ having same homotopy type, this illustrates the ability of the operator to capture geometric information and distinguish between non-isometric spaces.

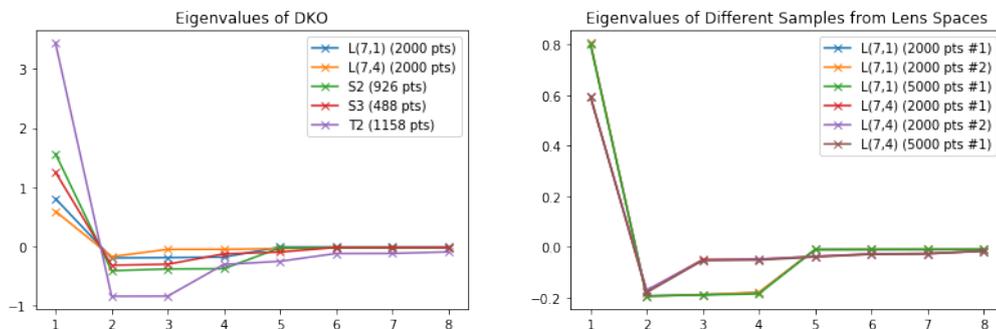
Hausdorff distance between DKEs. Finally, in Figure 1c, we compare the Hausdorff distances between various pairs of distance kernel embeddings. We observe the following: (1) The two samples of the same size coming from the $L(7, 1)$ Lens space are the closest in Hausdorff distance, and that distance is close to zero up to dimension $k = 4$. Indeed, if we had taken samples of sufficiently high resolution, we would see the Hausdorff distances going to zero for larger values of k , as proven in Theorem 17. (2) The second closest pair of spaces are the Lens spaces $L(7, 1)$ and $L(7, 4)$, that have same homotopy type and have both spherical geometry. (3) The third closest pair of spaces are the Lens space $L(7, 1)$ and the 3-Sphere, both with spherical geometry (Lens spaces are constructed as quotients of 3-Spheres). (4) The manifold that appears to be most distinct from the rest is the torus. (5) For all pairs of manifolds, the Hausdorff distance stabilizes at around $k = 10$, after which eigenvalues are close to 0.

In conclusion, these experiments illustrate that the spectra and embedding of the distance kernel operator can be approximated by finite samples, as predicted by Theorem 17. Moreover, by combining the DKE with the Hausdorff metric on Euclidean space, we obtain a pseudo-metric on the space of compact metric measure spaces that succeeds in distinguishing a variety of diverse manifolds.

7 Open Problems

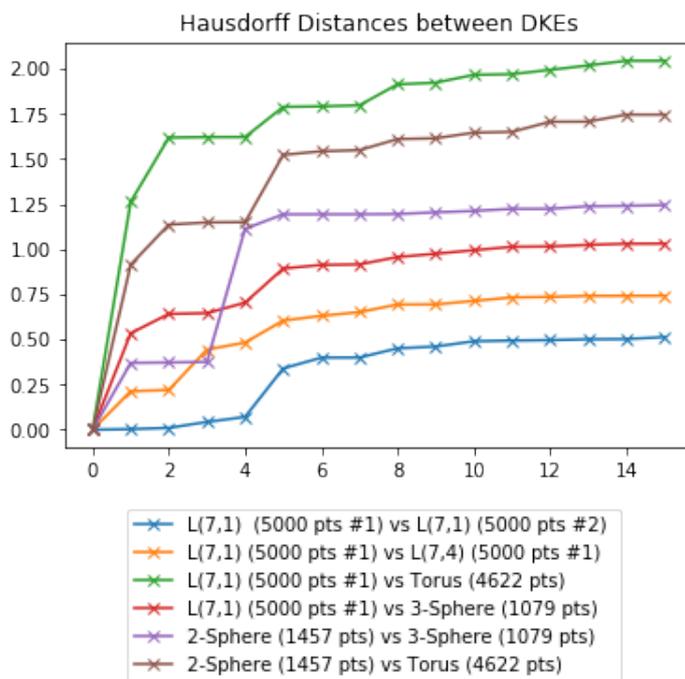
This work introduces new techniques, at the crossroads of topological data analysis and spectral geometry, to study general metric measure spaces. It also raises a number of interdisciplinary questions in persistence theory, optimal transport, spectral geometry, and o-minimal geometry, that, due to their specialized and technical nature, have not been resolved in this article:

- Using the sampling and stability results for the distance kernel operator (Theorems 17 and 18) to provide analogous results for the topological transforms.
- Proving that the truncated distance kernel embedding is an injection for k sufficiently large. This for example the case for Laplacian eigenfunctions on manifolds, as shown by Bates [1], whose proof relies on deeper results in spectral geometry.
- Providing general hypotheses that ensure the definability of the DKE.
- Obtaining experimental results for these topological transforms in line with the distance kernel embedding experiments of Section 6. These experiments will hinge on a principled method for choosing which vector directions should be used for the computation of topological invariants; this is a question of interest in the TDA community, and we expect some heuristics and theoretical guarantees to emerge on this topic in the near future.



(a) Eigenvalues of the DKO for a variety of spaces, normalized by the number of points in the sample.

(b) A comparison of the eigenvalues of various samples, at different resolutions, of these two Lens spaces.



(c) A comparison of Hausdorff distances between various samples of 2- and 3-manifolds.

■ **Figure 1** Spectra and DKE for samples of various manifolds. In subfigures (a) and (b), the x -represents the index of the eigenvalues in the sorted sequence of eigenvalues. In subfigure (c), the x -axis represents the embedding dimension (over \mathbb{C}).

References

- 1 Jonathan Bates. The embedding dimension of Laplacian eigenfunction maps. *Applied and Computational Harmonic Analysis*, 37(3):516–530, 2014.
- 2 Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in neural information processing systems*, pages 585–591, 2002.
- 3 Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- 4 Mikhail Belkin and Partha Niyogi. Convergence of laplacian eigenmaps. In *Advances in Neural Information Processing Systems*, pages 129–136, 2007.
- 5 Robin Lynne Belton, Brittany Terese Fasy, Rostik Mertz, Samuel Micka, David L Millman, Daniel Salinas, Anna Schenfisch, Jordan Schupbach, and Lucia Williams. Learning simplicial complexes from persistence diagrams. *arXiv preprint*, 2018. [arXiv:1805.10716](https://arxiv.org/abs/1805.10716).
- 6 Dmitri Burago, Sergei Ivanov, and Yaroslav Kurylev. A graph discretization of the laplace-beltrami operator. *arXiv preprint*, 2013. [arXiv:1301.2222](https://arxiv.org/abs/1301.2222).
- 7 Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006.
- 8 Lorin Crawford, Anthea Monod, Andrew X Chen, Sayan Mukherjee, and Raúl Rabadán. Topological summaries of tumor images improve prediction of disease free survival in glioblastoma multiforme. *arXiv preprint*, 2016. [arXiv:1611.06818](https://arxiv.org/abs/1611.06818).
- 9 Justin Curry, Sayan Mukherjee, and Katharine Turner. How many directions determine a shape and other sufficiency results for two topological transforms. *arXiv preprint*, 2018. [arXiv:1805.09782](https://arxiv.org/abs/1805.09782).
- 10 Tamal K Dey, Dayu Shi, and Yusu Wang. Comparing graphs via persistence distortion. *arXiv preprint*, 2015. [arXiv:1503.07414](https://arxiv.org/abs/1503.07414).
- 11 Gerald B Folland. *A Guide to Advanced Real Analysis*. Number 37 in Dolciani Mathematical Expositions. MAA, 2009.
- 12 Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and euler integral transforms. *Journal of Applied and Computational Topology*, 2(1-2):55–60, 2018.
- 13 Christopher Heil. Compact and hilbert–schmidt operators. In *A Basis Theory Primer*, pages 481–490. Springer, 2011.
- 14 Fritz John et al. The ultrahyperbolic differential equation with four independent variables. *Duke Mathematical Journal*, 4(2):300–322, 1938.
- 15 Clément Maria, Steve Oudot, and Elchanan Solomon. Intrinsic topological transforms via the distance kernel embedding. *arXiv preprint*, 2019. [arXiv:submit/2957368](https://arxiv.org/abs/submit/2957368).
- 16 Steve Oudot and Elchanan Solomon. Barcode embeddings for metric graphs. *arXiv preprint*, 2017. [arXiv:1712.03630](https://arxiv.org/abs/1712.03630).
- 17 Steve Oudot and Elchanan Solomon. Inverse problems in topological persistence. *arXiv preprint*, 2018. [arXiv:1810.10813](https://arxiv.org/abs/1810.10813).
- 18 Iosif Polterovich, Leonid Polterovich, and Vukašin Stojisavljević. Persistence barcodes and laplace eigenfunctions on surfaces. *Geometriae Dedicata*, 201(1):111–138, 2019.
- 19 Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- 20 Katharine Turner, Sayan Mukherjee, and Doug M Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 2014.
- 21 Sathamangalam R Srinivasa Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2):431–455, 1967.