

A Stable Multi-Scale Kernel for Topological Machine Learning

Jan Reininghaus, Stefan Huber
IST Austria

Ulrich Bauer
IST Austria, TU München

Roland Kwitt
University of Salzburg, Austria

Abstract

Topological data analysis offers a rich source of valuable information to study vision problems. Yet, so far we lack a theoretically sound connection to popular kernel-based learning techniques, such as kernel SVMs or kernel PCA. In this work, we establish such a connection by designing a multi-scale kernel for persistence diagrams, a stable summary representation of topological features in data. We show that this kernel is positive definite and prove its stability with respect to the 1-Wasserstein distance. Experiments on two benchmark datasets for 3D shape classification/retrieval and texture recognition show considerable performance gains of the proposed method compared to an alternative approach that is based on the recently introduced persistence landscapes.

1. Introduction

In many computer vision problems, data (*e.g.*, images, meshes, point clouds, etc.) is piped through complex processing chains in order to extract information that can be used to address high-level inference tasks, such as recognition, detection or segmentation. The extracted information might be in the form of low-level appearance descriptors, *e.g.*, SIFT [20], or of higher-level nature, *e.g.*, activations at specific layers of deep convolutional networks [18]. In recognition problems, for instance, it is then customary to feed the consolidated data to a discriminant classifier such as the popular support vector machine (SVM), a kernel-based learning technique.

While there has been substantial progress on extracting and encoding discriminative information, only recently have people started looking into the *topological structure* of the data as an additional source of information. With the emergence of *topological data analysis (TDA)* [4], computational tools for efficiently identifying topological structure have become readily available. Since then, several authors have demonstrated that methods of TDA can capture characteristics of the data that other methods often fail to reveal, *cf.* [26, 19].

Along these lines, studying persistent homology [11] is

a particularly popular method for TDA, since it captures the birth and death times of topological features, *e.g.*, connected components, holes, etc., at multiple scales. This information is summarized by the *persistence diagram*, a multiset of points in the plane. The key feature of persistent homology is its stability: small changes in the input data lead to small changes in the Wasserstein distance of the associated persistence diagrams [10]. Considering the discrete nature of topological information, the existence of such a well-behaved summary is perhaps surprising.

Note that persistence diagrams together with the Wasserstein distance only form a metric space. Thus it is not possible to directly employ persistent homology in the large class of machine learning techniques that require a Hilbert space structure, like SVMs or PCA. This obstacle is typically circumvented by defining a kernel function on the domain containing the data, which in turn defines a Hilbert space structure implicitly. While the Wasserstein distance itself does not naturally lead to a valid kernel (see supplementary material for details), we show that it is possible to define a kernel for persistence diagrams that is stable w.r.t. the 1-Wasserstein distance. This is the main contribution of this paper.

Contribution. We propose a (positive definite) multi-scale kernel for persistence diagrams (see Fig. 1). This kernel is defined via an L_2 -valued feature map, based on ideas from scale space theory [16]. We show that our feature map is Lipschitz continuous with respect to the 1-Wasserstein distance, thereby maintaining the stability property of persistent homology. The scale parameter of our kernel controls its robustness to noise and can be tuned to the data. We investigate, in detail, the theoretical properties of the kernel, and demonstrate its applicability on shape classification/retrieval and texture recognition benchmarks.

2. Related work

Methods that leverage topological information for computer vision or medical image analysis can roughly be grouped into two categories. In the first category, we identify previous work that *directly* utilizes topological information to address a specific problem, such as topology-guided segmentation. In the second category, we identify

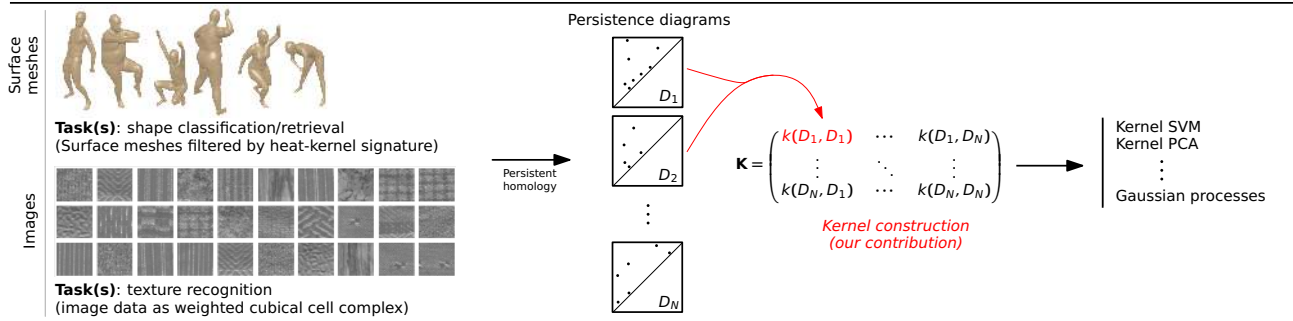


Figure 1: Visual data (e.g., functions on surface meshes, textures, etc.) is analyzed using persistent homology [11]. Roughly speaking, persistent homology captures the birth/death times of topological features (e.g., connected components or holes) in the form of *persistence diagrams*. Our contribution is to define a *kernel for persistence diagrams* to enable a theoretically sound use of these summary representations in the framework of kernel-based learning techniques, popular in the computer vision community.

approaches that *indirectly* use topological information. That is, information about topological features is used as input to some machine-learning algorithm.

As a representative of the first category, Skraba *et al.* [26] adapt the idea of persistence-based clustering [6] in a segmentation approach for surface meshes of 3D shapes, driven by the topological information in the persistence diagram. Gao *et al.* [12] use persistence information to restore so called *handles*, i.e., topological cycles, in already existing segmentations of the left ventricle, extracted from computed tomography images. In a different segmentation setup, Chen *et al.* [7] propose to directly incorporate topological constraints into random-field based segmentation models.

In the second category of approaches, Chung *et al.* [8] and Pachauri *et al.* [22] investigate the problem of analyzing cortical thickness measurements on 3D surface meshes of the human cortex in order to study developmental and neurological disorders. In contrast to [26], persistence information is not used directly, but rather as a *descriptor* that is fed to a discriminant classifier in order to distinguish between normal control patients and patients with Alzheimer’s disease/autism. Yet, the step of training the classifier with topological information is typically done in a rather ad hoc manner. In [22] for instance, the persistence diagram is first rasterized on a regular grid, then a kernel-density estimate is computed, and eventually the vectorized discrete probability density function is used as a feature vector to train a SVM using standard kernels for \mathbb{R}^n . It is however unclear how the resulting kernel-induced distance behaves with respect to existing metrics (e.g., bottleneck or Wasserstein distance) and how properties such as stability are affected. An approach that directly uses well-established distances between persistence diagrams for recognition was recently proposed by Li *et al.* [19]. Besides bottleneck and Wasserstein distance, the authors employ persistence landscapes [3] and the corresponding distance in their experiments.

Their results expose the complementary nature of persistence information when combined with traditional bag-of-feature approaches. While our empirical study in Sec. 5.2 is inspired by [19], we primarily focus on the development of the kernel; the combination with other methods is straightforward.

In order to enable the use of persistence information in machine learning setups, Adcock *et al.* [1] propose to compare persistence diagrams using a feature vector motivated by algebraic geometry and invariant theory. The features are defined using algebraic functions of the birth and death values in the persistence diagram.

From a conceptual point of view, Bubenik’s concept of *persistence landscapes* [3] is probably the closest to ours, being another kind of feature map for persistence diagrams. While persistence landscapes were not explicitly designed for use in machine learning algorithms, we will draw the connection to our work in Sec. 5.1 and show that they in fact admit the definition of a valid positive definite kernel. Moreover, both persistence landscapes as well as our approach represent computationally attractive alternatives to the bottleneck or Wasserstein distance, which both require the solution of a matching problem.

3. Background

First, we review some fundamental notions and results from persistent homology that will be relevant for our work.

Persistence diagrams. *Persistence diagrams* are a concise description of the topological changes occurring in a growing sequence of shapes, called *filtration*. In particular, during the growth of a shape, holes of different dimension (i.e., gaps between components, tunnels, voids, etc.) may appear and disappear. Intuitively, a k -dimensional hole, born at time a and filled at time b , gives rise to a point (a, b) in the k^{th} persistence diagram. A persistence diagram is thus




Figure 2: A function $\mathbb{R} \rightarrow \mathbb{R}$ (left) and its 0^{th} persistence diagram (right). Local minima create a connected component in the corresponding sublevel set, while local maxima merge connected components. The pairing of birth and death is shown in the persistence diagram.

a multiset of points in \mathbb{R}^2 . Formally, the persistence diagram is defined using a standard concept from algebraic topology called *homology*; see [11] for details.

Note that not every hole has to disappear in a filtration. Such holes give rise to *essential* features and are naturally represented by points of the form (a, ∞) in the diagram. Essential features therefore capture the topology of the final shape in the filtration. In the present work, we do not consider these features as part of the persistence diagram. Moreover, all persistence diagrams will be assumed to be finite, as is usually the case for persistence diagrams computed from data.

Filtrations from functions. A standard way of obtaining a filtration is to consider the *sublevel sets* $f^{-1}(-\infty, t]$ of a function $f: \Omega \rightarrow \mathbb{R}$ defined on some domain Ω , for $t \in \mathbb{R}$. It is easy to see that the sublevel sets indeed form a filtration parametrized by t . We denote the resulting persistence diagram by D_f ; see Fig. 2 for an illustration.

As an example, consider a grayscale image, where Ω is the rectangular domain of the image and f is the grayscale value at any point of the domain (*i.e.*, at a particular pixel). A sublevel set would thus consist of all pixels of Ω with value up to a certain threshold t . Another example would be a piecewise linear function on a triangular mesh Ω , such as the popular heat kernel signature [27]. Yet another commonly used filtration arises from a point cloud P embedded in \mathbb{R}^n , by considering the distance function $d_P(x) = \min_{y \in P} \|x - y\|$ on $\Omega = \mathbb{R}^n$. The sublevel sets of this function are unions of balls around each point in P . Computationally, they are usually replaced by equivalent constructions called *alpha shapes*.

Stability. A crucial aspect of the persistence diagram D_f of a function f is its stability with respect to perturbations of f . In fact, only stability guarantees that one can infer information about the function f from its persistence diagram D_f in the presence of noise.

Formally, we consider $f \mapsto D_f$ as a map of metric spaces and define *stability* as Lipschitz continuity of this map. This requires choices of metrics both on the set of functions and

the set of persistence diagrams. For the functions, the L_∞ metric is commonly used.

There is a natural metric associated to persistence diagrams, called the *bottleneck distance*. Loosely speaking, the distance of two diagrams is expressed by minimizing the largest distance of any two corresponding points, over all bijections between the two diagrams. Formally, let F and G be two persistence diagrams, each augmented by adding each point (t, t) on the diagonal with countably infinite multiplicity. The *bottleneck distance* is

$$d_B(F, G) = \inf_{\mu} \sup_{x \in F} \|x - \mu(x)\|_\infty, \quad (1)$$

where μ ranges over all bijections from the individual points of F to the individual points of G . As shown by Cohen-Steiner *et al.* [9], persistence diagrams are stable with respect to the bottleneck distance.

The bottleneck distance embeds into a more general class of distances, called *Wasserstein distances*. For any positive real number p , the p -*Wasserstein distance* is

$$d_{W,p}(F, G) = \left(\inf_{\mu} \sum_{x \in F} \|x - \mu(x)\|_\infty^p \right)^{\frac{1}{p}}, \quad (2)$$

where μ again ranges over all bijections from the individual elements of F to the individual elements of G . Note that taking the limit $p \rightarrow \infty$ yields the bottleneck distance, and we therefore define $d_{W,\infty} = d_B$. We have the following result bounding the p -Wasserstein distance in terms of the L_∞ distance:

Theorem 1 (Cohen-Steiner *et al.* [10]). *Assume that X is a compact triangulable metric space such that for every 1-Lipschitz function f on X and for $k \geq 1$, the degree k total persistence $\sum_{(b,d) \in D_f} (d-b)^k$ is bounded above by some constant C . Let f, g be two L -Lipschitz piecewise linear functions on X . Then for all $p \geq k$,*

$$d_{W,p}(D_f, D_g) \leq (LC)^{\frac{1}{p}} \|f - g\|_\infty^{1 - \frac{k}{p}}. \quad (3)$$

We note that, strictly speaking, this is not a stability result in the sense of Lipschitz continuity, since it only establishes Hölder continuity. Moreover, it only gives a constant upper bound for the p -Wasserstein distance when $p = 1$.

Kernels. Given a set X , a function $k: X \times X \rightarrow \mathbb{R}$ is a *kernel* if there exists a Hilbert space \mathcal{H} , called *feature space*, and a map $\Phi: X \rightarrow \mathcal{H}$, called *feature map*, such that $k(x, y) = \langle \Phi(x), \Phi(y) \rangle_{\mathcal{H}}$ for all $x, y \in X$. Equivalently, k is a kernel if it is symmetric and positive definite [24]. Kernels allow to apply machine learning algorithms operating on a Hilbert space to be applied to more general settings, such as strings, graphs, or, in our case, persistence diagrams.

A kernel induces a pseudometric $d_k(x, y) = (k(x, x) + k(y, y) - 2k(x, y))^{1/2}$ on \mathcal{X} , which is the distance $\|\Phi(x) - \Phi(y)\|_{\mathcal{H}}$ in the feature space. We call the kernel k *stable* w.r.t. a metric d on \mathcal{X} if there is a constant $C > 0$ such that $d_k(x, y) \leq C d(x, y)$ for all $x, y \in \mathcal{X}$. Note that this is equivalent to Lipschitz continuity of the feature map.

The stability of a kernel is particularly useful for classification problems: assume that there exists a separating hyperplane H for two classes of data points with margin m . If the data points are perturbed by some $\epsilon < m/2$, then H still separates the two classes with a margin $m - 2\epsilon$.

4. The persistence scale-space kernel

We propose a stable *multi-scale* kernel k_σ for the set of persistence diagrams \mathcal{D} . This kernel will be defined via a feature map $\Phi_\sigma : \mathcal{D} \rightarrow L_2(\Omega)$, with $\Omega \subset \mathbb{R}^2$ denoting the closed half plane above the diagonal.

To motivate the definition of Φ_σ , we point out that the set of persistence diagrams, *i.e.*, multisets of points in \mathbb{R}^2 , does not possess a Hilbert space structure per se. However, a persistence diagram D can be uniquely represented as a sum of Dirac delta distributions¹, one for each point in D . Since Dirac deltas are functionals in the Hilbert space $H^{-2}(\mathbb{R}^2)$ [17, Chapter 7], we obtain a canonical Hilbert space structure for persistence diagrams by adopting this point of view.

Unfortunately, the induced metric on \mathcal{D} does *not* take into account the distance of the points in the diagrams or to the diagonal, and therefore cannot be robust against perturbations of the diagrams. Motivated by scale-space theory [16], we address this issue by using the sum of Dirac deltas as an initial condition for a heat diffusion problem with a Dirichlet boundary condition on the diagonal. The solution of this partial differential equation is an $L_2(\Omega)$ function for any chosen scale parameter $\sigma > 0$. In the following paragraphs, we will

- 1) define the persistence scale space kernel k_σ ,
- 2) derive a simple formula for evaluating k_σ , and
- 3) prove stability of k_σ w.r.t. the 1-Wasserstein distance.

Definition 1. Let $\Omega = \{x = (x_1, x_2) \in \mathbb{R}^2 : x_2 \geq x_1\}$ denote the space above the diagonal, and let δ_x denote a Dirac delta centered at the point x . For a given persistence diagram D , we now consider the solution $u : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$, $(x, t) \mapsto u(x, t)$ of the partial differential equation²

$$\Delta_x u = \partial_t u \quad \text{in } \Omega \times \mathbb{R}_{>0}, \quad (4)$$

$$u = 0 \quad \text{on } \partial\Omega \times \mathbb{R}_{\geq 0}, \quad (5)$$

$$u = \sum_{y \in D} \delta_y \quad \text{on } \Omega \times \{0\}. \quad (6)$$

¹A Dirac delta distribution is a functional that evaluates a given smooth function at a point.

²Since the initial condition (6) is not an $L_2(\Omega)$ function, this equation is to be understood in the sense of distributions. For a rigorous treatment of existence and uniqueness of the solution, see [17, Chapter 7].

The feature map $\Phi_\sigma : \mathcal{D} \rightarrow L_2(\Omega)$ at scale $\sigma > 0$ of a persistence diagram D is now defined as $\Phi_\sigma(D) = u|_{t=\sigma}$. This map yields the persistence scale space kernel k_σ on \mathcal{D} as

$$k_\sigma(F, G) = \langle \Phi_\sigma(F), \Phi_\sigma(G) \rangle_{L_2(\Omega)}. \quad (7)$$

Note that $\Phi_\sigma(D) = 0$ for some $\sigma > 0$ implies that $u = 0$ on $\Omega \times \{0\}$, which means that D has to be the empty diagram. From linearity of the solution operator it now follows that Φ_σ is an injective map.

The solution of the partial differential equation can be obtained by extending the domain from Ω to \mathbb{R}^2 and replacing (6) with

$$u = \sum_{y \in D} \delta_y - \delta_{\bar{y}} \quad \text{on } \mathbb{R}^2 \times \{0\}, \quad (8)$$

where $\bar{y} = (b, a)$ is $y = (a, b)$ mirrored at the diagonal. It can be shown that restricting the solution of this extended problem to Ω yields a solution for the original equation. It is given by convolving the initial condition (8) with a Gaussian kernel:

$$u(x, t) = \frac{1}{4\pi t} \sum_{y \in D} \exp\left(-\frac{\|x - y\|^2}{4t}\right) - \exp\left(-\frac{\|x - \bar{y}\|^2}{4t}\right). \quad (9)$$

Using this closed form solution of u , we can derive a simple expression for evaluating the kernel explicitly:

$$k_\sigma(F, G) = \frac{1}{8\pi\sigma} \sum_{\substack{y \in F \\ z \in G}} \exp\left(-\frac{\|y - z\|^2}{8\sigma}\right) - \exp\left(-\frac{\|y - \bar{z}\|^2}{8\sigma}\right). \quad (10)$$

We refer to the supplementary material for the elementary derivation of (10) and for a visualization of the solution (9). Note that the kernel can be computed in $O(|F| \cdot |G|)$ time, where $|F|$ and $|G|$ denote the cardinality of the multisets F and G , respectively (or approximated with bounded error in $O(|F| + |G|)$ using the *Fast Gauss Transform* [14]).

Theorem 2. The kernel k_σ is 1-Wasserstein stable.

Proof. To prove 1-Wasserstein stability of k_σ , we show Lipschitz continuity of the feature map Φ_σ as follows:

$$\|\Phi_\sigma(F) - \Phi_\sigma(G)\|_{L_2(\Omega)} \leq \frac{1}{2\sqrt{\pi}\sigma} d_{W,1}(F, G), \quad (11)$$

where F and G denote persistence diagrams that have been augmented with points on the diagonal. Note that augmenting diagrams with points on the diagonal does not change the values of Φ_σ , as can be seen from (9). By definition of the Wasserstein distance, for any $\delta > d_{W,1}(F, G)$ there exists a matching μ between F and G with $\sum_{y \in F} \|y - \mu(y)\|_\infty \leq \delta$.

Writing $N_y(x) = \frac{1}{4\pi\sigma} \exp\left(-\frac{\|x - y\|^2}{4\sigma}\right)$, we have

$$\|N_y - N_z\|_{L_2(\mathbb{R}^2)} = \frac{1}{\sqrt{4\pi}\sigma} \cdot \sqrt{1 - \exp\left(-\frac{\|y - z\|^2}{8\sigma}\right)}.$$

The definitions of Φ_σ from (9) and of N_y , the Minkowski inequality, the inequality $1 - \exp(-\xi) \leq \xi$, and the inequality $\|x\|_2 \leq \sqrt{2}\|x\|_\infty$ finally yield

$$\begin{aligned}
& \|\Phi_\sigma(F) - \Phi_\sigma(G)\|_{L_2(\Omega)} \\
&= \left\| \sum_{y \in F} (N_y - N_{\bar{y}}) - \sum_{y \in G} (N_{\mu(y)} - N_{\bar{\mu(y)}}) \right\|_{L_2(\mathbb{R}^2)} \\
&\leq \left\| \sum_{y \in F} (N_y - N_{\mu(y)}) \right\|_{L_2(\mathbb{R}^2)} + \left\| \sum_{y \in F} (N_{\bar{y}} - N_{\bar{\mu(y)}}) \right\|_{L_2(\mathbb{R}^2)} \\
&= 2 \sum_{y \in F} \|N_y - N_{\mu(y)}\|_{L_2(\mathbb{R}^2)} \\
&= \frac{1}{\sqrt{\pi}\sigma} \sum_{y \in F} \sqrt{1 - \exp\left(-\frac{\|y - \mu(y)\|^2}{8\sigma}\right)} \\
&\leq \frac{1}{\sqrt{8\pi}\sigma} \sum_{y \in F} \|y - \mu(y)\| \leq \frac{1}{2\sqrt{\pi}\sigma} \delta
\end{aligned}$$

for any $\delta > d_{W,1}(F, G)$, and the claim follows. \square

We refer to the left-hand side of (11) as the *persistence scale space distance* d_{k_σ} . Note that the right hand side of (11) decreases as σ increases. Adjusting σ accordingly allows to counteract the influence of noise in the input data, which causes an increase in $d_{W,1}(F, G)$. We will see in Sec. 5.3 that tuning σ to the data can be beneficial for the overall performance of machine learning methods.

A natural question arising from Theorem 2 is whether our stability result extends to p -Wasserstein distances for $p > 1$. To answer this question, we first note that our kernel is *additive*: we call a kernel k on persistence diagrams additive if $k(E \cup F, G) = k(E, G) + k(F, G)$ for all $E, F, G \in \mathcal{D}$. By choosing $F = \emptyset$, we see that if k is additive then $k(\emptyset, G) = 0$ for all $G \in \mathcal{D}$. We further say that a kernel k is *trivial* if $k(F, G) = 0$ for all $F, G \in \mathcal{D}$. The next theorem establishes that Theorem 2 is sharp in the sense that *no* non-trivial additive kernel can be stable w.r.t. $d_{W,p}$ when $p > 1$.

Theorem 3. *A non-trivial additive kernel k on persistence diagrams is not stable w.r.t. $d_{W,p}$ for any $1 < p \leq \infty$.*

Proof. By the non-triviality of k , it can be shown that there exists an $F \in \mathcal{D}$ such that $k(F, F) > 0$. We prove the claim by comparing the rates of growth of $d_k(\bigcup_{i=1}^n F, \emptyset)$ and $d_{W,p}(\bigcup_{i=1}^n F, \emptyset)$ w.r.t. n . We have

$$d_k\left(\bigcup_{i=1}^n F, \emptyset\right) = n \sqrt{k(F, F)}.$$

On the other hand,

$$d_{W,p}\left(\bigcup_{i=1}^n F, \emptyset\right) = d_{W,p}(F, \emptyset) \cdot \begin{cases} \sqrt[p]{n} & \text{if } p < \infty, \\ 1 & \text{if } p = \infty. \end{cases}$$

Hence, d_k can not be bounded by $C \cdot d_{W,p}$ with a constant $C > 0$ if $p > 1$. \square

5. Evaluation

To evaluate the kernel proposed in Sec. 4, we investigate conceptual differences to persistence landscapes in Sec. 5.1, and then consider its performance in the context of shape classification/retrieval and texture recognition in Sec. 5.2.

5.1. Comparison to persistence landscapes

In [3], Bubenik introduced *persistence landscapes*, a representation of persistence diagrams as functions in the Banach space $L_p(\mathbb{R}^2)$. This construction was mainly intended for statistical computations, enabled by the vector space structure of L_p . For $p = 2$, we can use the Hilbert space structure of $L_2(\mathbb{R}^2)$ to construct a kernel analogously to (7). For the purpose of this work, we refer to this kernel as the *persistence landscape kernel* k^L and denote by $\Phi^L: \mathcal{D} \rightarrow L_2(\mathbb{R}^2)$ the corresponding feature map. The kernel-induced distance is denoted by d_{k^L} . Bubenik shows stability w.r.t. a weighted version of the Wasserstein distance, which for $p = 2$ can be summarized as:

Theorem 4 (Bubenik [3]). *For any two persistence diagrams F and G we have*

$$\begin{aligned}
& \|\Phi^L(F) - \Phi^L(G)\|_{L_2(\mathbb{R}^2)} \leq \\
& \inf_{\mu} \left(\sum_{y \in F} \text{pers}(y) \|y - \mu(y)\|_\infty^2 + \frac{2}{3} \|y - \mu(y)\|_\infty^3 \right)^{\frac{1}{2}}, \quad (12)
\end{aligned}$$

where $\text{pers}(y) = b - a$ denotes the persistence of $y = (a, b)$, and μ ranges over all bijections from F to G .

For a better understanding of the stability results in Theorems 2 and 4, we discuss two thought experiments.

For the first experiment, let $F_\lambda = \{-\lambda, \lambda\}$ and $G_\lambda = \{-\lambda + 1, \lambda + 1\}$ be two diagrams with one point each and $\lambda \in \mathbb{R}_{\geq 0}$. The two points move away from the diagonal with increasing λ , while maintaining the same Euclidean distance to each other. Consequently, $d_{W,p}(F_\lambda, G_\lambda)$ and $d_{k_\sigma}(F_\lambda, G_\lambda)$ asymptotically approach a constant as $\lambda \rightarrow \infty$. In contrast, $d_{k^L}(F_\lambda, G_\lambda)$ grows in the order of $\sqrt{\lambda}$ and, in particular, is unbounded. This means that d_{k^L} emphasizes points of high persistence in the diagrams, as reflected by the weighting term $\text{pers}(u)$ in (12).

In the second experiment, we compare persistence diagrams from data samples of two fictive classes A (*i.e.*, diagrams F, F') and B (*i.e.*, diagram G), illustrated in Fig. 4. We first consider $d_{k^L}(F, F')$. As we have seen in the previous experiment, d_{k^L} will be dominated by variations in the points of high persistence. Similarly, $d_{k^L}(F, G)$ will also be dominated by these points as long as λ is sufficiently large. Hence, instances of classes A and B would be inseparable




Figure 3: Examples from SHREC 2014 [23] (left, middle) and OuTeX Outex_TC.000000 [21] (right).




Figure 4: Two persistence diagrams F, F' from class A and one diagram G from class B. The classes only differ in their points of low-persistence (*i.e.*, points closer to the diagonal).

in a nearest-neighbor setup. In contrast, d_B , $d_{w,p}$ and d_{k_σ} do *not* over-emphasize points of high persistence and thus allow to distinguish classes A and B.

5.2. Empirical results

We report results on two vision tasks where persistent homology has already been shown to provide valuable discriminative information [19]: *shape classification/retrieval* and *texture image classification*. The purpose of the experiments is *not* to outperform the state-of-the-art on these problems – which would be rather challenging by exclusively using topological information – but to demonstrate the advantages of k_σ and d_{k_σ} over k^L and d_{k^L} .

Datasets. For shape classification/retrieval, we use the SHREC 2014 [23] benchmark, see Fig. 3. It consists of both *synthetic* and *real* shapes, given as 3D meshes. The synthetic part of the data contains 300 meshes of humans (five males, five females, five children) in 20 different poses; the real part contains 400 meshes from 40 humans (male, female) in 10 different poses. We use the meshes in full resolution, *i.e.*, without any mesh decimation. For classification, the objective is to distinguish between the different human models, *i.e.*, a 15-class problem for SHREC 2014 (synthetic) and a 40-class problem for SHREC 2014 (real).

For texture recognition, we use the Outex_TC.000000 benchmark [21], downsampled to 32×32 pixel images. The benchmark provides 100 predefined training/testing splits and each of the 24 classes is equally represented by 10 images during training and testing.

Implementation. For shape classification/retrieval, we compute the classic *Heat Kernel Signature* (HKS) [27] over a range of ten time parameters t_i of increasing value. For each specific choice of t_i , we obtain a piecewise linear function on the surface mesh of each object. As discussed in Sec. 3, we then compute the persistence diagrams of the induced filtrations in dimensions 0 and 1.

For texture classification, we compute CLBP [15] descriptors, (*cf.* [19]). Results are reported for the rotation-invariant versions of the CLBP-Single (CLBP-S) and the CLBP-Magnitude (CLBP-M) operator with $P = 8$ neighbors and radius $R = 1$. Both operators produce a scalar-valued response image which can be interpreted as a weighted cubical cell complex and its lower star filtration is used to compute persistence diagrams; see [28] for details.

For both types of input data, the persistence diagrams are obtained using DIPHA [2], which can directly handle meshes and images. A standard soft margin C -SVM classifier [24], as implemented in LIBSVM [5], is used for classification. The cost factor C is tuned using ten-fold cross-validation on the training data. For the kernel k_σ , this cross-validation further includes the kernel scale σ .

5.2.1 Shape classification

Tables 1 and 2 list the classification results for k_σ and k^L on SHREC 2014. All results are averaged over ten cross-validation runs using random 70/30 training/testing splits with a roughly equal class distribution. We report results for 1-dimensional features only; 0-dimensional features lead to comparable performance.

On both real and synthetic data, we observe that k_σ leads to consistent improvements over k^L . For some choices of t_i , the gains even range up to 30%, while in other cases, the improvements are relatively small. This can be explained by the fact that varying the HKS time t_i essentially varies the smoothness of the input data. The scale σ in k_σ allows to compensate—at the classification stage—for unfavorable smoothness settings to a certain extent, see Sec. 4. In contrast, k^L does not have this capability and essentially relies on suitably preprocessed input data. For some choices of t_i , k^L does in fact lead to classification accuracies close to k_σ .

HKS t_i	k^L	k_{σ}	Δ
t_1	68.0 ± 3.2	94.7 ± 5.1	+26.7
t_2	88.3 ± 3.3	99.3 ± 0.9	+11.0
t_3	61.7 ± 3.1	96.3 ± 2.2	+34.7
t_4	81.0 ± 6.5	97.3 ± 1.9	+16.3
t_5	84.7 ± 1.8	96.3 ± 2.5	+11.7
t_6	70.0 ± 7.0	93.7 ± 3.2	+23.7
t_7	73.0 ± 9.5	88.0 ± 4.5	+15.0
t_8	81.0 ± 3.8	88.3 ± 6.0	+7.3
t_9	67.3 ± 7.4	88.0 ± 5.8	+20.7
t_{10}	55.3 ± 3.6	91.0 ± 4.0	+35.7

Table 1: Classification performance on SHREC 2014 (synthetic).

HKS t_i	k^L	k_{σ}	Δ
t_1	45.2 ± 5.8	48.8 ± 4.9	+3.5
t_2	31.0 ± 4.8	46.5 ± 5.3	+15.5
t_3	30.0 ± 7.3	37.8 ± 8.2	+7.8
t_4	41.2 ± 2.2	50.2 ± 5.4	+9.0
t_5	46.2 ± 5.8	62.5 ± 2.0	+16.2
t_6	33.2 ± 4.1	58.0 ± 4.0	+24.7
t_7	31.0 ± 5.7	62.7 ± 4.6	+31.7
t_8	51.7 ± 2.9	57.5 ± 4.2	+5.8
t_9	36.0 ± 5.3	41.2 ± 4.9	+5.2
t_{10}	2.8 ± 0.6	27.8 ± 5.8	+25.0

Table 2: Classification performance on SHREC 2014 (real).

However, when using k^L , we have to carefully adjust the HKS time parameter, corresponding to changes in the input data. This is undesirable in most situations, since HKS computation for meshes with a large number of vertices can be quite time-consuming and sometimes we might not even have access to the meshes directly. The improved classification rates for k_{σ} indicate that using the additional degree of freedom is in fact beneficial for performance.

5.2.2 Shape retrieval

In addition to the classification experiments, we report on shape retrieval performance using standard evaluation measures (see [25, 23]). This allows us to assess the behavior of the kernel-induced distances $d_{k_{\sigma}}$ and d_{k^L} .

For brevity, only the nearest-neighbor performance is listed in Table 3 (for a listing of all measures, see the supplementary material). Using each shape as a query shape once, nearest-neighbor performance measures how often the top-ranked shape in the retrieval result belongs to the same class as the query. To study the effect of tuning the scale σ , the column $d_{k_{\sigma}}$ lists the *maximum* nearest-neighbor performance that can be achieved over a range of scales.

As we can see, the results are similar to the classification experiment. However, at a few specific settings of the HKS time t_i , d_{k^L} performs on par, or better than $d_{k_{\sigma}}$. As noted in Sec. 5.2.1, this can be explained by the changes in the smoothness of the input data, induced by different HKS times t_i . Another observation is that nearest-neighbor performance of d_{k^L} is quite unstable around the top result w.r.t.

HKS t_i	d_{k^L}	$d_{k_{\sigma}}$	Δ	d_{k^L}	$d_{k_{\sigma}}$	Δ
t_1	53.3	88.7	+35.4	24.0	23.7	-0.3
t_2	91.0	94.7	+3.7	20.5	25.7	+5.2
t_3	76.7	91.3	+14.6	16.0	18.5	+2.5
t_4	84.3	93.0	+8.7	26.8	33.0	+6.2
t_5	85.0	92.3	+7.3	28.0	38.7	+10.7
t_6	63.0	77.3	+14.3	28.7	36.8	+8.1
t_7	65.0	80.0	+15.0	43.5	52.7	+9.2
t_8	73.3	80.7	+7.4	70.0	58.2	-11.8
t_9	73.0	83.0	+10.0	45.2	56.7	+11.5
t_{10}	51.3	69.3	+18.0	3.5	44.0	+40.5
Top-3 [23]	99.3 – 92.3 – 91.0			68.5 – 59.8 – 58.3		

Table 3: Nearest-neighbor retrieval performance. *Left:* SHREC 2014 (synthetic); *Right:* SHREC 2014 (real).

t_i . For example, it drops from 91% (at t_2) to 53.3% and 76.7% on SHREC 2014 (synthetic) and from 70% (at t_8) to 45.2% and 43.5% on SHREC 2014 (real). In contrast, $d_{k_{\sigma}}$ exhibits stable performance around the optimal t_i .

To put these results into context with existing works in shape retrieval, Table 3 also lists the top three entries (out of 22) of [23] on the same benchmark. On both real and synthetic data, $d_{k_{\sigma}}$ ranks among the top five entries. This indicates that topological persistence alone is a rich source of discriminative information for this particular problem. In addition, since we only assess one HKS time parameter at a time, performance could potentially be improved by more elaborate fusion strategies.

5.3. Texture recognition

For texture recognition, all results are averaged over the 100 training/testing splits of the Outex_TC_000000 benchmark. Table 4 lists the performance of a SVM classifier using k_{σ} and k^L for 0-dimensional features (*i.e.*, connected components). Higher-dimensional features were not informative for this problem. For comparison, Table 4 also lists the performance of a SVM, trained on normalized histograms of CLBP-S/M responses, using a χ^2 kernel.

First, from Table 4, it is evident that k_{σ} performs better than k^L by a large margin, with gains up to $\approx 11\%$ in accuracy. Second, it is also apparent that, for this problem, topological information alone is not competitive with SVMs using simple orderless operator response histograms. However, the results of [19] show that a *combination* of persistence information (using persistence landscapes) with conventional bag-of-feature representations leads to state-of-the-art performance. While this indicates the complementary nature of topological features, it also suggests that kernel combinations (*e.g.*, via multiple-kernel learning [13]) could lead to even greater gains by including the proposed kernel k_{σ} .

To assess the stability of the (customary) cross-validation strategy to select a specific σ , Fig. 5 illustrates classification performance as a function of the latter. Given the smooth-

CLBP Operator	k^L	k_{σ}	Δ
CLBP-S	58.0 ± 2.3	69.2 ± 2.7	+11.2
CLBP-M	45.2 ± 2.5	55.1 ± 2.5	+9.9
CLBP-S (SVM- χ^2)	76.1 ± 2.2		
CLBP-M (SVM- χ^2)	76.7 ± 1.8		

Table 4: Classification performance on Outex_TC_000000.




Figure 5: Texture classification performance of a SVM classifier with (1) the kernel k_{σ} as a function of σ , (2) the kernel k_{σ} with σ cross-validated and (3) the kernel k^L are shown.

ness of the performance curve, it seems unlikely that parameter selection via cross-validation will be sensitive to a specific discretization of the search range $[\sigma_{\min}, \sigma_{\max}]$.

Finally, we remark that tuning k^L has the same drawbacks in this case as in the shape classification experiments. While, in principle, we could smooth the textures, the CLBP response images, or even tweak the CLBP operators, all those strategies would require changes at the beginning of the processing pipeline. In contrast, adjusting σ in k_{σ} is done at the end of the pipeline during classifier training.

6. Conclusion

We have shown, both theoretically and empirically, that the proposed kernel exhibits good behavior for tasks like shape classification or texture recognition using a SVM. Moreover, the ability to tune a scale parameter has proven beneficial in practice.

One possible direction for future work would be to address computational bottlenecks in order to enable application in large scale scenarios. This could include leveraging additivity and stability in order to approximate the value of the kernel within given error bounds, in particular, by reducing the number of distinct points in the summation of (10).

While the 1-Wasserstein distance is well established and has proven useful in applications, we hope to improve the understanding of stability for persistence diagrams w.r.t. the Wasserstein distance beyond the previous estimates. Such a result would extend the stability of our kernel from persistence diagrams to the underlying data, leading to a full stability proof for topological machine learning.

In summary, our method enables the use of topological information in all kernel-based machine learning methods. It will therefore be interesting to see which other application

areas will profit from topological machine learning.

References

- [1] A. Adcock, E. Carlsson, and G. Carlsson. The Ring of Algebraic Functions on Persistence Bar Codes. arXiv, available at <http://arxiv.org/abs/1304.0530>, 2013.
- [2] U. Bauer, M. Kerber, and J. Reininghaus. Distributed computation of persistent homology. In *ALLENEX*, 2014.
- [3] P. Bubenik. Statistical topological data analysis using persistence landscapes. arXiv, available at <http://arxiv.org/abs/1207.6437>, 2012.
- [4] G. Carlsson. Topology and data. *Bull. Amer. Math. Soc.*, 46:255–308, 2009.
- [5] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM TIST*, 2(3):1–27, 2011.
- [6] F. Chazal, L. Guibas, S. Oudot, and P. Skraba. Persistence-based clustering in Riemannian manifolds. In *SoSG*, 2011.
- [7] C. Chen, D. Freedman, and C. Lampert. Enforcing topological constraints in random field image segmentation. In *CVPR*, 2013.
- [8] M. Chung, P. Bubenik, and P. Kim. Persistence diagrams of cortical surface data. In *IPMI*, 2009.
- [9] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete Comp. Geom.*, 37(1):103–120, 2007.
- [10] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have L_p -stable persistence. *Found. Comput. Math.*, 10(2):127–139, 2010.
- [11] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. AMS, 2010.
- [12] M. Gao, C. Chen, S. Zhang, Z. Qian, D. Metaxas, and L. Axel. Segmenting the papillary muscles and the trabeculae from high resolution cardiac CT through restoration of topological handles. In *IPMI*, 2013.
- [13] M. Gönen and E. Alpaydin. Multiple kernel learning algorithms. *J. Mach. Learn. Res.*, 12:2211–2268, 2011.
- [14] L. Greengard and J. Strain. The fast Gauss transform. *SIAM J. Sci. Statist. Comput.*, 12(1):79–94, 1991.
- [15] Z. Guo, L. Zhang, and D. Zhang. A completed modeling of local binary pattern operator for texture classification. *IEEE TIP*, 19(6):1657–1663, 2010.
- [16] T. Iijima. Basic theory on normalization of a pattern (in case of typical one-dimensional pattern). *Bulletin of Electrical Laboratory*, 26:368–388, 1962.
- [17] R. J. J. Iorio and V. de Magalhães Iorio. *Fourier analysis and partial differential equations*. Cambridge Stud. Adv. Math., 2001.
- [18] A. Krizhevsky, I. Sutskever, and G. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [19] C. Li, M. Ovsjanikov, and F. Chazal. Persistence-based structural recognition. In *CVPR*, 2014.
- [20] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [21] T. Ojala, T. Mäenpää, M. Pietikäinen, J. Viertola, J. Kyllönen, and S. Huovinen. OuTeX – new framework for empirical evaluation of texture analysis algorithms. In *ICPR*, 2002.
- [22] D. Pachauri, C. Hinrichs, M. Chung, S. Johnson, and V. Singh. Topology-based kernels with application to inference problems in Alzheimers disease. *IEEE TMI*, 30(10):1760–1770, 2011.
- [23] Pickup, D. et al.. SHREC '14 track: Shape retrieval of non-rigid 3d human models. In *Proceedings of the 7th Eurographics workshop on 3D Object Retrieval*, EG 3DOR '14. Eurographics Association, 2014.
- [24] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [25] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The Princeton shape benchmark. In *Shape Modeling International*, 2004.
- [26] P. Skraba, M. Ovsjanikov, F. Chazal, and L. Guibas. Persistence-based segmentation of deformable shapes. In *CVPR Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, 2010.
- [27] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and probably informative multi-scale signature based on heat diffusion. In *SGP*, 2009.
- [28] H. Wagner, C. Chen, and E. Vućini. Efficient computation of persistent homology for cubical data. In *Topological Methods in Data Analysis and Visualization II*, Mathematics and Visualization, pages 91–106. Springer Berlin Heidelberg, 2012.