

Convolutional Persistence Transforms

Yitzchak Elchanan Solomon^{1,3*} and Paul Bendich^{1,2,3}

^{1*}Mathematics, Duke University, Durham, NC, USA.

²Geometric Data Analytics, Durham, NC, USA.

³The authors were partially supported by the Air Force Office of Scientific Research under the grant “Geometry and Topology for Data Analysis and Fusion”, AFOSR FA9550-18-1-0266. They would also like to extend their thanks to Alexander Wagner, for helpful conversations..

*Corresponding author(s). E-mail(s):

elchanansolomon@gmail.com;

Contributing authors: paul.bendich@duke.edu;

Abstract

In this paper, we consider topological featurizations of data defined over simplicial complexes, like images and labeled graphs, obtained by convolving this data with various filters before computing persistence. Viewing a convolution filter as a local motif, the persistence diagram of the resulting convolution describes the way the motif is distributed across the simplicial complex. This pipeline, which we call *convolutional persistence*, extends the capacity of topology to observe patterns in such data. Moreover, we prove that (generically speaking) for any two labeled complexes one can find *some* filter for which they produce different persistence diagrams, so that the collection of all possible convolutional persistence diagrams is an injective invariant. This is proven by showing convolutional persistence to be a special case of another topological invariant, the Persistent Homology Transform. Other advantages of convolutional persistence are improved stability, greater flexibility for data-dependent vectorizations, and reduced computational complexity for certain data types. Additionally, we have a suite of experiments showing that convolutions greatly improve the predictive power of persistence on a host of classification tasks, even if one uses random filters and vectorizes the resulting diagrams by recording only their total persistences.

Keywords: Persistent Homology, Topological Transform, Machine Learning, Convolutions

1 Overview

Persistent homology is a method of assigning multiscale topological descriptors to parametric families of shapes. In *functional* persistence, the object of study is a real-valued function $f : X \rightarrow \mathbb{R}$ defined over a space X , and the parametric family of shapes are the sublevel-sets $X_\alpha = \{x \in X : f(x) \leq \alpha\}$. It is similarly possible to consider superlevel-sets, which is equivalent to negating the *filter function* f . One crucial feature of this construction is that X_α is a subset of X_β for $\alpha \leq \beta$, so that the sublevel-sets are naturally nested. The output of persistent homology is a collection of intervals (equivalently, a collection of points), called a *barcode* (or *persistence diagram*, using the point representation). The space of barcodes is not a vector space, even approximately (Bubenik and Wagner, 2020; Wagner, 2021), but there do exist multiple *vectorizations* (Adams et al, 2017; Bubenik et al, 2015; Monod et al, 2019; Di Fabio and Ferri, 2015; Carrière et al, 2015) that transform barcodes into vectors suitable for machine learning and data analysis.

A very general setting for functional persistence is that of *simplicial complexes*, i.e. shapes obtained by gluing together points, edges, triangles, tetrahedra, etc. For example, if X is a triangulation of a 2D rectangular grid, a function $f : X \rightarrow [0, 1]$ can be viewed as a greyscale image. Persistent homology can then be understood as a feature extraction method for such images, either for supervised or unsupervised learning. Example applications include removing image noise (Chung and Day, 2018; Chung et al, 2022), parameter estimation for PDEs (Calcina and Gameiro, 2021; Adams et al, 2017), segmentation (Hu et al, 2019), flow estimation (Suzuki et al, 2021), tumor analysis (Crawford et al, 2020), cell immune micronenvironment (Aukerman et al, 2020), and materials science (Hiraoka et al, 2016).

It should be understood that purely topological methods do not provide state-of-the-art predictive accuracy on most machine learning tasks, and are not to be considered as *alternatives* to more general methods like neural nets and kernel methods. Rather, topological methods provide *principled* and *interpretable* features that are different from those extracted by other methods, and can help improve the performance and utility of the entire pipeline, especially in settings where data is limited, see (Khrantsova et al, 2022).

To that end, it is important to understand the properties of functional persistence and their role in machine learning. Here we consider some of the most salient properties:

- (Computational complexity) For a simplicial complex K with N simplices, computing persistence is $O(N^\omega)$, where ω is the matrix multiplication constant (Milosavljević et al, 2011). This means that such persistence calculations scale poorly in the resolution of data, especially high-dimensional data, where doubling the resolution in \mathbb{R}^d results in a 2^d -fold increase in

the number of simplices. Still, in many settings, persistence calculations are on average much faster than the pessimistic worst-case complexity, see e.g. (Giunti et al, 2021).

- (Stability) Persistent homology is stable to small perturbations of the input data, in that the distance between the barcodes for two functions f and g on a common space X is bounded by $\|f - g\|_\infty$. However, persistence is not stable to outliers, so images that look similar outside of a small fraction of pixels can produce wildly different barcodes.
- (Flexibility) There is a single persistence diagram to be associated with each pair (X, f) of space X and real-valued function f . This lack of additional parameters makes applying persistence straightforward, but can also be limiting in contexts where data-dependent featurizations are desired.
- (Invertibility) There exist many distinct space-function pairs (X, f) producing identical barcodes. Thus, persistence is not invertible as feature map, and this loss of information may hinder the capacity of persistence-based methods to identify patterns or distinguish distinct images.

Our goal in this paper is to introduce a modification to persistence of certain simplicial complexes that is (1) often faster to compute, (2) more robust to outliers, (3) allows for data-driven tuning, (4) is provably invertible, and (5) is more informative in practice. Essentially, this technique consists of passing multiple convolutional filters over the data before computing persistence, and so we call it *convolutional persistence*. A crucial insight in this paper is that convolutional persistence can generically be transformed into a special case of the Persistent Homology Transform (PHT) (Turner et al, 2014) in high-dimensions, implying that it shares the theoretical properties enjoyed by that invariant.

To get an intuitive sense for why convolutions might be of value in applied topology, consider Figure 1, which shows how convolution with a simple filter makes geometric information in an image more available to topological methods.

1.1 Results

The major theoretical result of the paper is that convolutional persistence is generically injective (Theorems 8 and 10). We also provide stability results for convolutional persistence (Propositions 5 and 6). The experimental results of Section 5 show that convolutional persistence significantly outperforms ordinary persistence at a collection of classification tasks. Interestingly enough, random filters perform well for the classification tasks considered, even when only the total persistences of the resulting diagrams are taken as feature vectors.




Fig. 1 The original image is shown on the left. The image has interesting topological structure but the sublevel-set persistence is trivial. On the right, we show the result of convolving with a 2×2 -filter with a value of 1 in each pixel. The sublevel-set persistence of the convolved image is more informative; at height zero, there are three connected components, corresponding to three regions in the image where the convolution is zero. These regions then merge at an immediate value, due to those parts of the image where the local patches are more similar to the filter.

1.2 Organization

The remainder of the paper is organized as follows. Section 2 provides a thorough, non-technical survey on persistent homology. Section 3 defines two versions of *convolutional persistence* and reviews prior related work. Section 4 contains the main theoretical results of the paper. Section 5 compares ordinary and convolutional persistence on a host of image datasets, showing the capacity of convolutional persistence to produce features well-suited for image classification. Finally, Section 6 discusses outstanding questions and potential generalizations of this work. Code for running experiments with convolutional persistence can be found at <https://github.com/yesolomon/convpers>.

2 Background

The content of this paper assumes familiarity with the concepts and tools of persistent homology. Interested readers can consult the articles of Carlsson (Carlsson, 2009) and Ghrist (Ghrist, 2008) and the textbooks of Edelsbrunner and Harer (Edelsbrunner and Harer, 2010) and Oudot (Oudot, 2015). We include the following primer for readers interested in a high-level, non-technical summary.

2.1 Persistent Homology

Persistent homology records the way topology evolves in a parametrized sequence of spaces. In the case of functional persistence, we consider an ambient space X equipped with a real-valued function $f : X \rightarrow \mathbb{R}$. The sublevel-sets $X_\alpha = \{x \in X : f(x) \leq \alpha\}$ of f are naturally nested, in that X_α is a subset of X_β for $\alpha \leq \beta$, and these form a *filtration* of X . Persistence records how the topology of this filtration evolves as a function of the parameter α .

Simplicial complexes admit particularly simple filtrations. To every simplex σ we can associate a real value $f(\sigma)$ that encodes the parameter value at which it appears in the filtration of X . The only restriction on the function f is the following consistency condition: if σ is a sub-simplex of a higher-dimensional simplex τ (i.e. an edge which sits at the boundary of a triangle), we must have $f(\sigma) \leq f(\tau)$, ensuring that simplices do not appear before any of their faces.

Given a filtered simplicial complex, as the sequence of sublevel-sets evolves, the addition of certain edges or higher-dimensional simplices alters the *topological type* of the space. A precise way of quantifying topology is *homology*, which measures the number of connected components (zero-dimensional homology), cycles (one-dimensional homology), or voids (higher-dimensional homology) in a space. Thus, homology can change when two connected components merge or a new cycle is formed. Simplices responsible for such topological changes are called *critical*. Persistent homology records the parameter values at which critical simplices appear, notes the dimension in which the homology changes, and pairs critical values by matching the critical value at which a new homological feature appears to the critical value at which it disappears. This information is then organized into a structure called a *barcode*, which is simply a collection of intervals. Figure 2 shows the computation of the zero- and one-dimensional barcodes for a simple simplicial complex.

A barcode can also be encoded as a collection of points, simply by mapping each interval $[a, b]$ to the point $(a, b) \in \{(a, b) \mid a < b, a, b \in \mathbb{R} \cup \infty\}$. This collection of points is called a *persistence diagram*, and besides for certain subtleties, such as dealing with open or closed endpoints and keeping track of points with multiplicity, it contains the exact same information as the original barcode. However, for certain analyses and vectorizations of persistence features, persistence diagrams are more useful than barcodes, as will be shown below.

When working with simplicial complexes, it is possible to limit the maximal dimension of simplices allowed in the construction. Given an arbitrary simplicial complex K , we write K^m to denote the subcomplex consisting of all simplices of dimension at most m ; this is called the *m -skeleton of K* . Note that K and K^m have the same homology in dimensions less than m , but may differ in dimension m , and K^m has no homology in dimension greater than m .

2.2 Image Cubical Complexes

A common alternative to simplicial complexes are *cubical complexes*, i.e. shapes that are built from vertices, edges, squares, cubes, etc. Such complexes are more natural when working with data built on top of a grid, such as pixel or voxel data, see (Kaczynski et al, 2004) for details. A cubical complex can

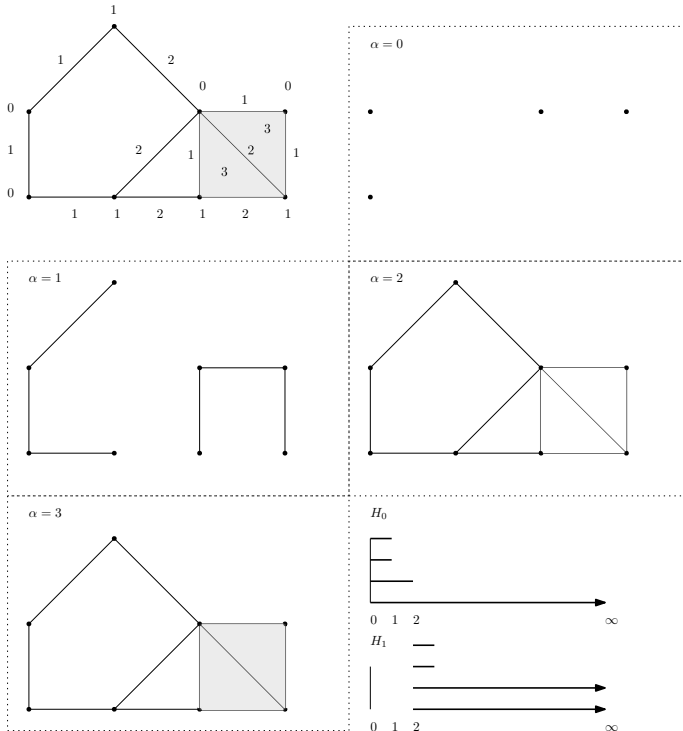


Fig. 2 Top-left: A simplicial complex with filtration values attached to vertices, edges, and squares. Top-right through bottom-left: sublevel-sets associated with different threshold values. Bottom-right: barcodes in dimensions zero and one. The zero-dimensional homology H_0 barcode contains four bars, since at $\alpha = 0$ there are four connected components. Two bars die at $\alpha = 1$, since at that threshold value there are only two connected components. Finally, $\alpha = 2$ sees the merger of these connected component, so another bar dies at $\alpha = 2$ and the last persists to infinity. In one-dimensional homology H_1 , three bars are born at $\alpha = 2$, when three loops appear in the sublevel-set, and one of these bars dies at $\alpha = 3$, when that loop is killed off by the introduction of a square.

always be converted into a simplicial complex by cutting up cubes into simplicial pieces: a canonical way of doing this is the *Freudenthal triangulation*. For a square, for example, this simply amounts to inserting a diagonal edge that produces two triangles. Given a function on a cubical complex, there is a canonical way to extend this function to its Freudenthal triangulation without changing its persistence; in the square example, in which the Freudenthal triangulation incorporates a diagonal edge and two triangles, we simply assign to these three simplices the filtration value assigned to the original square in the original cubical complex. Thus, from the perspective of persistence, it does not fundamentally matter if we work with cubical or simplicial complexes, although there is a difference vis-a-vis computations.

Given a d -dimensional grayscale image, there are two ways of turning this data into a cubical complex. One is to view the voxels as being vertices, and




Fig. 3 A 2×2 image can be turned into a complex in one of two ways. Left: A complex with four top-dimensional cubes, with function values, indicated using color, extended to vertices and edges via the upper-* rule. Right: A complex with four vertices, with function values extended to the edges and interior square via the lower-* rule.

higher-dimensional cubes as coming from voxels adjacencies, so that pairs of adjacent pixels form an edge and squares come from four voxels in a square formation, etc. There is a canonical way of extending the function f from the voxels (vertices) to the entire complex: given a cube τ , define $f(\tau)$ to be $\max_{\sigma} f(\sigma)$, where the max is taken over all vertices $\sigma < \tau$. Thus, a square appears precisely once all its constituent vertices appear; this is called the *lower-** filtration.

Alternatively, one can also view the voxels as being d -dimensional cubes, and have the lower-dimensional cubes be the faces of these voxels. As before, there is a canonical way of extending the function value from the voxels (top-dimensional cubes) down to entire complex: for a cube σ , define $f(\sigma)$ to be $\min_{\tau} f(\tau)$, where the min is taken over all voxels τ that contain σ . Thus, a cube appears precisely when at least one of the voxels in which it participates does; this is called the *upper-** filtration.

Consult figure 3 for an illustration of these two images complexes. Generally, these complexes will differ, and the resulting persistence barcodes will be different. However, there exists a formula for reading the barcodes for one construction from the barcodes of the other, see (Bleile et al, 2021).

2.3 Comparing Persistence Diagrams

As multi-sets of points in the plane, persistence diagrams are not vectors. However, there exist multiple metrics for comparing persistence diagrams. The most common approach is to view persistence diagrams as discrete distributions on the plane \mathbb{R}^2 , and use techniques from optimal transport theory, such as p -Wasserstein metrics W_p , to compare them, see (Oudot, 2017) Chapter 3 and (Villani, 2021). This analogy is complicated by the fact that persistence diagrams do not all have the same number of points, and that points in a persistence diagram near the diagonal line $y = x$ correspond to transient homological features, dying shortly after they are born, which ought not




Fig. 4 Two persistence diagram D_1 and D_2 , one in black and the other in blue. An optimal matching between these diagrams is shown, containing both diagonal and non-diagonal pairings. The resulting p -Wasserstein distance is then $W_p(D_1, D_2) = \sqrt[p]{d_1^p + d_2^p + d_3^p + d_4^p + d_5^p}$.

to play an important role in dictating similarity of diagrams. These problems are ameliorated by modifying the optimal transport protocol to allow pairing points in one diagram either with points in the other diagram or with the diagonal line $y = x$, the latter incurring a cost proportional to the distance of the given point from the diagonal. This *diagonal pairing* allows us to define an optimal transport distance between any pair of diagrams, regardless of how many points they have. In practice, the p -Wasserstein metrics in use are either $p = 1$, $p = 2$, or $p = \infty$, the latter of which is called the *Bottleneck distance* in persistence theory, and is written d_B . See Figure 4 for a visualization of a matching between persistence diagrams.

2.4 Computational Complexity

Persistence calculations are $O(N^\omega)$, where N is the number of simplices in the filtration and ω is the matrix multiplication constant (Milosavljević et al, 2011). Though this bound is in practice quite pessimistic, it is accurate in reflecting the poor scaling of persistence calculations in the resolution of an image (Otter et al, 2017).

The bottleneck distance between persistence diagrams can be computed exactly using the well-known Hungarian algorithm, whose complexity is $O(n^3)$ for pairs of diagrams with at most n points. The number of points in the persistence diagram of filtration is always bounded by the total number N of simplices in the filtration, and may be much smaller. If one is interested in *approximating* the bottleneck, or more generally, p -Wasserstein distances between persistence diagrams, there are significantly faster methods than the Hungarian algorithm. One such method is based on *entropic approximation* techniques from optimal transport (Lacombe et al, 2018), and provided that the norms of the points in the diagrams are uniformly bounded by some constant C , a transport plan within ϵ of the optimal matching can be produced

in $O(\frac{n \log(n)C}{\epsilon^2})$ iterations of the Sinkhorn algorithm, which is itself linear in n . Moreover, recent work in optimal transport theory provide an improved asymptotic convergence rates for Sinkhorn's algorithm of $\tilde{O}(n^2/\epsilon)$, where \tilde{O} indicates that logarithmic factors have been hidden, see (Pham et al, 2020).

2.5 Properties of Persistent Homology

Persistence theory guarantees that a small modification to the filtration of a space produces only small changes in its persistence diagram.

Theorem 1 ((Cohen-Steiner et al, 2007)¹) *Let $f, g : X \rightarrow \mathbb{R}$ be two filtrations on a simplicial complex. Then the Bottleneck distance between their persistence diagrams is bounded by $\|f - g\|_\infty$.*

Corollary 1 *Let $f, g : K^0 \rightarrow \mathbb{R}$ be two functions on the vertex set of a simplicial complex, with lower-* extensions \hat{f} and \hat{g} , respectively. Then $\|\hat{f} - \hat{g}\|_\infty \leq \|f - g\|_\infty$, so that the Bottleneck distance between the persistence diagrams of \hat{f} and \hat{g} is bounded by $\|f - g\|_\infty$.*

Proof Consider a simplex $\sigma \in K$. By definition, $\hat{f}(\sigma) = \max_{v \in \sigma^0} f(v)$ and $\hat{g}(\sigma) = \max_{v \in \sigma^0} g(v)$. Let v_f and v_g be the vertices of σ realizing these maxima for f and g , respectively. Then $f(v_f) \leq g(v_f) + \|f - g\|_\infty \leq g(v_g) + \|f - g\|_\infty$, so that $f(\sigma) \leq g(\sigma) + \|f - g\|_\infty$. A symmetric argument shows that $g(\sigma) \leq f(\sigma) + \|f - g\|_\infty$, so that $|f(\sigma) - g(\sigma)| \leq \|f - g\|_\infty$. \square

This implies that a small error in the filter function produces only small distortion in the resulting persistence diagram. However, persistent homology is not at all stable to *outliers*, i.e. a small subset of the data having large error (Buchet et al, 2014).

More recent work (Skraba and Turner, 2020) has established analogous stability results for W_p metrics between persistence diagrams, *although we caution readers that this paper has yet be to published*. The latter of the following two results is particularly useful for the subject matter of this paper.

Definition 1 ((Skraba and Turner, 2020), Definition 4.1) The L^p norm of a function on a simplicial complex $f : K \rightarrow \mathbb{R}$ is:

$$\|f\|_p^p = \sum_{\sigma \in K} |f(\sigma)|^p.$$

¹The main theorem of (Cohen-Steiner et al, 2007) is much more general, applying to tame functions on triangulable spaces, conditions which are automatically satisfied in the simplicial setting.



Fig. 5 Two functions on the same grid with the same persistence, using either sublevel- or superlevel-set filtrations. Color scheme: {black: 0, grey: 0.5, white: 1}. All persistence diagrams consist of the single point $(0, 1)$ in dimension 0.

Theorem 2 ((Skraba and Turner, 2020), Theorem 4.8) *Let $f, g : K \rightarrow \mathbb{R}$ be two monotone functions on a fixed simplicial complex. Then for all $p \geq 1$,*

$$W_p(\text{Diag}(f), \text{Diag}(g)) \leq \|f - g\|_p.$$

Theorem 3 ((Skraba and Turner, 2020), Theorem 5.1) *Let $f, g : P \rightarrow \mathbb{R}$ be the grayscale functions for two images defined over the same grid of pixels P . Let \hat{f} and \hat{g} be the corresponding extensions to the d -dimensional cubical complex built on top of P , by either method described in Section 2.2. Then we have the stability result:*

$$W_p(\text{Diag}(\hat{f}), \text{Diag}(\hat{g})) \leq \left(\sum_{i=1}^d 2^{d-i} \binom{d}{i} \right) \|f - g\|_p = (3^d - 2^d) \|f - g\|_p.$$

Another important feature of persistent homology is that the pipeline mapping a function f on a simplicial complex K , thought of as a vector in \mathbb{R}^K , to its associated persistence diagram is *almost everywhere differentiable* (Gameiro et al, 2016; Poulenard et al, 2018). This is because the coordinates of the points in the diagram correspond to the values of f on critical simplices, and this correspondence is generically locally stable. However, large perturbations of f will change the location of critical simplices, as well as their pairings (Bendich et al, 2020), and non-differentiable nonlinearities can be found in those degenerate cases when the critical simplex data is not unique and locally unstable. This almost everywhere differentiability of persistent homology allows us to use gradient descent methods for topological optimization (Carriere et al, 2021). Applications of topological optimization to machine learning models, and especially to deep learning, were considered in (Gabrielsson et al, 2020).

It is also important to note that persistent homology is not an injective invariant, meaning that different functions on the same space, or on different spaces, can have the same persistence diagrams (Curry, 2018; Leygonie and Tillmann, 2022; Leygonie and Henselman-Petrusek, 2021), see Figure 5. Thus, some information is lost in the process of converting a pair (X, f) to a persistence diagram.

We conclude the background section with an important lemma about the persistence of simplicial complexes embedded in Euclidean spaces. This lemma

equates the persistence of two filtrations on such complexes, both extended from a function on the vertex set, one discretely (via the lower-* filtration) and the other continuously (via linear interpolation).

Definition 2 Let K be a simplicial complex, and $f : K^0 \rightarrow \mathbb{R}$ a function on the vertices of K . We write \hat{f} to denote the lower-* extension of f to all of K , defined by $\hat{f}(\sigma) = \max_{v \in \sigma^0} f(v)$ for σ a simplex with vertex set σ^0 .

Proposition 4 Let K be an embedded simplicial complex in \mathbb{R}^d with vertex set K^0 . For a function $f : K^0 \rightarrow \mathbb{R}$ with lower-* extension \hat{f} , we can also extend f to all of K via linear interpolation, producing a function $f_* : K \rightarrow \mathbb{R}$. Then these two extensions produce identical persistence diagrams, i.e. $\text{Diag}(K, f_*) = \text{Diag}(K, \hat{f})$.

Proof This follows from Lemma 2.3 of (Bestvina and Brady, 1997) and Section VI.3 of (Edelsbrunner and Harer, 2010). \square

3 Convolutional Persistence

We now introduce two modifications to functional persistence of complexes, one based on techniques in *image processing* and applying to cubical complexes, and another based on *graph convolutions* and applying to either cubical or simplicial complexes (although, without loss of generality, we focus on the latter – see Section 2.2 for how a cubical complex can be converted into a simplicial complex). We call the former *Image Convolutional Persistence* and the latter *Simplicial Convolutional Persistence*.

3.1 Image Convolutional Persistence

Let $P \subset \mathbb{Z}^d$ be a rectangle inside of the integer lattice, and let $f : P \rightarrow \mathbb{R}$ be a function defined on P . P can be viewed as the vertex set of an m -dimensional cubical complex K_P^m , and f can be extended to this complex via the lower-* rule, in which $f(\sigma) = \max_{p \in \sigma} f(p)$, as in Section 2.2.

Definition 3 Let $B \subset \mathbb{Z}^d$ be another rectangle, and let $g : B \rightarrow \mathbb{R}$ be a function on this rectangle; the pair (B, g) acts as a *convolutional filter*. Fix a vector $k = (k_1, \dots, k_d)$ with $k_i \in \mathbb{N}_{>0}$, corresponding to the *stride* of the convolution. For $v \in \mathbb{Z}^d$ with $B + (v \odot k) \subseteq P$,² define:

$$(f * g)(v) = \sum_{p \in B} g(p)f(p + v \odot k).$$

Let $R \subset \mathbb{Z}^d$ be the collection of values r such that $B + (r \odot k) \subseteq P$, which is necessarily also a rectangle. The pair $(R, f * g)$ is the output of our convolution. See Figure 6.

²The symbol \odot refers to the Hadamard product, which performs componentwise multiplication between vectors.

We will generally assume that R is not empty, meaning that *some* translate of B is a subset of P .

Remark 1 If B consists of a single element b , and $g(b) = 1$, then, up to translation, $P = R$ and $f = f * g$.

Remark 2 Translating P or B has no substantial impact on the above construction, resulting only in a translation of the corresponding region R . Since all of the invariants we study here are invariant to translation of the domain, it would be possible to work entirely in the quotient space of grid functions $f : P \rightarrow \mathbb{R}$ defined up to translation. Equivalently, we might work with rectangles P whose lattice-minimal element (i.e. bottom-left corner) is the zero vector. Ultimately, we have chosen to eschew these additional formalisms in the interest of keeping our definitions as concrete as possible, but encourage the reader to adopt whichever perspective they find most insightful.

Definition 4 As a generalization of the above definition, we can allow n -channel images $f : P \rightarrow \mathbb{R}^n$ and filters $g : B \rightarrow \mathbb{R}^n$, and replace the product in the formula above with a dot product:

$$(f * g)(v) = \sum_{p \in B} g(p) \cdot f(p + v \odot k).$$

Thus, the result of the convolution is real-valued (not vector valued), and ordinary persistent homology can be computed. In principle, it is possible to have the image and filter dimensions differ, and to compute the persistent of multi-channel images via *multiparameter persistence*, see (Botnan and Lesnick, 2022), but this is beyond the scope of this work and presents its own set of theoretical and computational challenges.

R may be much smaller than the support of f , since the various translates of B covering this support are not required to overlap much, if at all. In the degenerate setting where P and B are the same shape, R will consist of a single vertex. For the purposes of computing persistence, we will think of R , like P , as being the vertex set of a cubical complex K_R^m , and we extend functions on R to the entirety of the cubical complex using the lower-* rule.

Notation 1 For a function f defined on the vertex set of a complex K , we write \hat{f} to denote its lower-* extension to all of K .

We propose that the collection of persistence diagrams of the form $PH(K_R^m, \widehat{f * g_i})$, for some set of filter functions $\mathcal{G} = \{g_i\}$, is of greater general utility than $PH(K_P^m, \hat{f})$. The computational advantages are immediate: when the stride k is large, so that R is much smaller than P , we have replaced a single, very expensive calculation with multiple, significantly faster calculations that can be performed in parallel. Moreover, we claim that this approach,




Fig. 6 f is a function on a 4×3 grid P , with function values indicated by subscripts. g is a filter function on a 2×2 grid B . Using a stride vector $k = (1, 2)$ means that the rectangle B is translated horizontally one unit but vertically two units; there are four possible translates of B that sit inside P , corresponding to the 2×2 grid R . Each vertex of R corresponds to a translate of B , and the value of $f * g$ at that vertex is obtained by lining up the values of g and f on that translate, multiplying them element-wise, and then summing them.

which we deem *convolutional persistence* has superior inverse properties. Proofs of stability and injectivity can be found in the following Section 4.

The flexibility of convolutional persistence comes from allowing the collection of filters \mathcal{G} to be curated for the task at hand. Indeed, there are many ways for this choice to be made:

1. Take \mathcal{G} to consist of a collection of popular filters in image processing, like blurring, sharpening, and boundary detection.
2. Take \mathcal{G} to consist of eigenfilters identified via PCA on the set of patches of images in the training set.
3. Picking filters *at random*.
4. Incorporate convolutional persistence in a deep learning pipeline and learn \mathcal{G} to minimize a chosen loss function.

In Section 5, we compare some of these different choices on a range of data sets and learning tasks. We also demonstrate the stability of computational persistence experimentally.

We conclude this section with a simple complexity analysis. Given a d -dimensional grid P with p pixels of resolution in each direction, the total number of voxels $|P|$ is equal to p^d . Each voxel participates in a uniformly bounded number of higher-dimensional simplices, so the total number of simplices in the resultant complex is $O(p^d)$. Note that this is different from simplicial complexes built on top of point clouds, in which each point is allowed to create simplices with every other point; these non-local connections cause a further combinatorial explosion in the number of simplices. Thus, for functional persistence on grids, the total complexity of computing the persistence of a function on such a complex is $O(p^{d\omega})$, where ω is the matrix multiplication constant. When moving to convolutional persistence, the resolution shrinks when the stride is bigger than one. For a stride vector $k = (2, \dots, 2)$, the total number of voxels in the resulting grid R is $O((\frac{p}{2})^d)$, and hence the persistence calculations are $O((\frac{p}{2})^{d\omega})$, so that the complexity bound has decreased by a factor of $2^{d\omega}$. More generally, setting $\kappa = \prod_i k_i$, computing convolutional persistence for M filters has complexity $O(M \left(\frac{|P|}{\kappa}\right)^\omega)$. Additionally, increasing the size of B decreases the size of R , as fewer translates of B sit inside P ; to be precise, increasing the length of B in a given coordinate shrinks the resulting region R by an identical amount in the same coordinate. Altogether, this means that, even computing convolutional persistence for a large class of filters, the speed-ups afforded by downsampling are significant, especially when using large stride vectors. This was verified experimentally in (Solomon et al, 2021b), in the context of optimization of topological functionals.

3.2 Simplicial Convolutional Persistence

Let K be a connected simplicial complex, and $f : K^0 \rightarrow \mathbb{R}^d$ a function on the vertices of K , which can be extended via the lower-* filtration to a function \hat{f} on all of K . If $|K^0| = n$, we can encode the function f by a $n \times d$ matrix X by imposing an order on the vertices of K and setting the i th row of X to be the value of f on the i th vertex. We can also encode the 1-skeleton K^1 of the complex as a graph with adjacency matrix A . Following (Kipf and Welling, 2016), for a weight matrix W of shape $d \times k$, we can transform the data X via the map $X \rightarrow AXW$, which amounts to transforming the data X by first multiplying by W , and then defining the value at a vertex to be the average of the values at its neighbors (adding a self-loop lets us include the vertex itself in this set). When $W = w$ has shape $d \times 1$, AXW encodes a map from K^0 to \mathbb{R} . We now define simplicial convolutional persistence by generalizing this construction and incorporating persistent homology.

Definition 5 Let K be a simplicial complex and $f : K^0 \rightarrow \mathbb{R}^d$ a function on the vertices of K , encoded by an $n \times d$ matrix X . Let A be an arbitrary, but fixed, $n \times n$ matrix, not necessarily the adjacency matrix of K^1 , and $W = w$ a $d \times 1$ weight matrix. We define the *simplicial convolutional persistence* of (K, A, f) on w to be

$\text{PH}(K, \widehat{f * w})$, where $f * w$ is the function on K^0 associated to AXw , and $\widehat{f * w}$ is its lower-* extension to all of K .

We take A to be an arbitrary matrix to allow for graph structures with self-loops, weights, and asymmetric distances. It might seem from the definition above that there is no need to separate the data matrix X from the generalized adjacency matrix A , and so we might consider only the product AX , as this is what appears in all calculations. However, distinguishing X from A has conceptual value, as X encodes the data defined on the simplicial complex, and A defines the complex itself. Thus, for example, noise in our measurements ends up in matrix X rather than the matrix A , and this is important for understanding the stability of convolutional persistence.

Remark 3 As simplicial convolutional persistence does not shrink the domain of the function, it does not exhibit the computational speedups of image convolutional persistence.

3.3 Prior Work

Convolutional persistence is a type of *topological transform*, in which a parametrized family of topological invariants is associated to a fixed input. The first topological invariants studied were the *persistent homology transform* (PHT) and *Euler characteristic transform* (ECT) (Turner et al, 2014). In the same paper in which the PHT and ECT were defined, inverse theorems were proven; these theorems were generalized and extended by later work (Curry et al, 2018; Ghrist et al, 2018). We will later show that convolutional persistence is, generically, a special case of the PHT, allowing us to take advantage of the inverse theory developed for that invariant. Invertible topological transforms have also been developed for weighted simplicial complexes (Jiang et al, 2020), metric graphs (Oudot and Solomon, 2017), metric spaces (Solomon et al, 2022), and metric measure spaces (Maria et al, 2019). Consult (Oudot and Solomon, 2020) for a more thorough survey of inverse problem in applied topology.

Prior research has also studied the interaction of persistent homology and image convolutions. In (Solomon et al, 2021b), the authors use convolutions to stabilize and speed up topological optimization; the approach taken there can be viewed as a special case of convolutional persistence, in which the filter set \mathcal{G} consists of many random approximations of the box-smoothing filter. Unlike in this work, the goal of (Solomon et al, 2021b) is to produce many downsampled images with persistence similar to the original, permitting the computation of robust topological gradients. Another paper that considers both persistence and convolutions is (Kim et al, 2020), where the authors develop a topological layer for deep learning models; one such model they consider is a convolutional neural network, and the authors design the network to

compute topological features both before and after computing convolutions. However, (Kim et al, 2020) do not study the interaction of convolutions and persistence in any generality.

4 Theoretical Results

In what follows, we prove two versions of each result, one for each flavor of convolutional persistence.

4.1 Stability

Remark 4 We remind readers that Theorem 3 (Skraba and Turner, 2020) is an unpublished result, and that this Theorem is needed to establish the $p \neq \infty$ cases of the following two propositions.

We begin with image convolutional persistence. Fix an input domain P , filter region B and stride vector k , and a skeleton dimension m . Focusing on single-channel images, this determines the domain R for any convolutions of $f : P \rightarrow \mathbb{R}$ and $g : B \rightarrow \mathbb{R}$.

Proposition 5 *Let $f_1, f_2 : P \rightarrow \mathbb{R}$ be two functions on P , and $g : B \rightarrow \mathbb{R}$ a filter function. Then, computing persistence over the complex K_R^m , for any $p \geq 1$ we have:*

$$\begin{aligned} W_p(\text{Diag}(\widehat{f_1 * g}), \text{Diag}(\widehat{f_2 * g})) &\leq \\ (3^m - 2^m) \|g\|_1 \|f_1 - f_2\|_p &\leq \\ |B|(3^m - 2^m) \|g\|_\infty \|f_1 - f_2\|_p. \end{aligned}$$

When $p = \infty$, the bound can be made tighter:

$$\begin{aligned} W_\infty(\text{Diag}(\widehat{f_1 * g}), \text{Diag}(\widehat{f_2 * g})) &\leq \\ \|g\|_1 \|f_1 - f_2\|_\infty &\leq \\ |B| \|g\|_\infty \|f_1 - f_2\|_\infty. \end{aligned}$$

Proof Young's convolutional inequality³ states that for $1 \leq p, q, r \leq \infty$ with $\frac{1}{p} + \frac{1}{q} = \frac{1}{r} + 1$,

$$\|f * g\|_r \leq \|g\|_q \|f\|_p.$$

Setting $r = p$ forces $q = 1$, and setting $f = f_1 - f_2$, we obtain the bound $\|f_1 * g - f_2 * g\|_p \leq \|g\|_1 \|f_1 - f_2\|_p$. Since g is supported on a set of $|B|$ elements, we have $\|g\|_1 \leq |B| \|g\|_\infty$, so that $\|f_1 * g - f_2 * g\|_p \leq |B| \|g\|_\infty \|f_1 - f_2\|_p$. The bound then follows from Theorem 3 in general, and from Corollary 1 when $p = \infty$. \square

We now consider simplicial convolutional persistence.

³See (Hewitt and Ross, 2012), Theorem 20.18 on page 296, for a proof of Young's inequality for locally compact groups, which in particular includes \mathbb{Z}^d .

Definition 6 Let X be an $n \times k$ matrix. For $1 \leq p \leq \infty$, we define the p -operator norm of X as follows:

$$\|A\|_p^* = \sup_{x \neq 0} \frac{\|Ax\|_p}{\|x\|_p}.$$

It is a standard result that operator norms are submultiplicative, meaning that $\|X_1 X_2\|_p^* \leq \|X_1\|_p^* \|X_2\|_p^*$, where X_1 and X_2 are any two matrices that can be multiplied.

Proposition 6 Let K be a simplicial complex and $f_1, f_2 : K^0 \rightarrow \mathbb{R}^d$ two functions on its vertex set, encoded by matrices X_1 and X_2 . Write $n = |K^0|$, and take A to be the $n \times n$ matrix used in convolutional persistence. Take a weight vector $w \in \mathbb{R}^d$. Then, computing persistence over the complex K^m , for any $p \geq 1$ we have:

$$W_p(\text{Diag}(\widehat{f_1 * w}), \text{Diag}(\widehat{f_2 * w})) \leq (3^m - 2^m) \|A\|_p^* \|X_1 - X_2\|_p^* \|w\|_p.$$

When $p = \infty$, the bound can be made tighter:

$$W_\infty(\text{Diag}(\widehat{f_1 * w}), \text{Diag}(\widehat{f_2 * w})) \leq \|A\|_\infty^* \|X_1 - X_2\|_\infty^* \|w\|_\infty.$$

Proof The difference $f_1 * w - f_2 * w$ can be written as $A(X_1 - X_2)w$. Submultiplicativity of the p -operator norm shows that $\|A(X_1 - X_2)w\|_p^* \leq \|A\|_p^* \|(X_1 - X_2)w\|_p^*$. For column vectors v , it is easy to see that $\|v\|_p^* = \|v\|_p$, i.e. the p -operator norm is equal to the usual p -norm. Thus we can write

$$\|f_1 * w - f_2 * w\|_p = \|A(X_1 - X_2)w\|_p \leq \|A\|_p^* \|(X_1 - X_2)w\|_p.$$

As before, the bound then follows from Theorem 3 in general, and from Corollary 1 when $p = \infty$. \square

As all matrix norms are equivalent, the quantity $\|X_1 - X_2\|_p^*$ is bounded above and below by a constant multiple of $\|X_1 - X_2\|_p$, the matrix p -Frobenius norm, providing for stability in this norm too.

4.2 Inverse Theory

We begin this section by recalling the definition of the persistent homology transform and some of its injectivity properties. We then prove that the persistent homology transform can be extended to functions defined on abstract simplicial complexes, and that the resulting transform also has strong inverse properties. Finally, we use this new theorem to prove inverse results for our two convolutional persistence transforms.

Definition 7 ((Turner et al, 2014)) Let $M \subset \mathbb{R}^d$ be a finite simplicial complex. For every vector $v \in \mathbb{S}^{d-1}$, we can define the function $f_v(x) = \langle x, v \rangle$. The persistent homology transform $PHT(M)$ is then defined as the map from \mathbb{S}^{d-1} to the space of persistence diagrams that send a vector v to the persistence diagram of (S, f_v) . One can define a similar invariant using Euler curves instead of persistence diagrams, called the Euler characteristic transform (ECT).

The PHT was shown to be injective for $d = 2, 3$ in (Turner et al, 2014). Later work (Ghrist et al, 2018; Curry et al, 2018) proved injectivity in all dimensions, for both the PHT and ECT, for a very general class of subsets that includes simplicial complexes.

Theorem 7 ((Turner et al, 2014; Ghrist et al, 2018; Curry et al, 2018)) *Let $K, L \subset \mathbb{R}^d$ be two embedded simplicial complexes. If $PHT(K) = PHT(L)$ or $ECT(K) = ECT(L)$ then $K = L$.*

Now, given a complex (simplicial or cubical) K with a function $f : K^0 \rightarrow \mathbb{R}^d$, we can define an analogous topological transform.

Definition 8 Let K be a complex (simplicial or cubical) with a function $f : K^0 \rightarrow \mathbb{R}^d$ defined on its vertices, extending to a function \hat{f} on K via the lower-* filtration. For $v \in \mathbb{S}^{d-1}$, we can define the function $\hat{f} \cdot v : K \rightarrow \mathbb{R}$. We then define the PHT of (K, f) to be the map on \mathbb{S}^{d-1} that sends a vector $v \in \mathbb{S}^{d-1}$ to the persistence diagram of $(K, \hat{f} \cdot v)$. We define the ECT similarly.

We now prove an injectivity result for the PHT and ECT of such complex-function pairs (K, f) . This result shows that the PHT and ECT are injective up to a simple equivalence relation which we now define.

Definition 9 For a complex K with a function $f : K^0 \rightarrow \mathbb{R}^d$, we write f^Δ to denote the function $f^\Delta : \text{geom}(K) \rightarrow \mathbb{R}^d$ defined on the geometric realization of K via linear interpolation. This differs from \hat{f} , which is defined via the discrete lower-* rule. Note that we use the notation $\text{geom}(K)$ as opposed to the more traditional $|K|$ to avoid confusion with the cardinality of K as a set.

Definition 10 Let (K, f) and (L, g) be two complexes with \mathbb{R}^d -valued functions defined on their vertex sets. We will say that (K, f) and (L, g) are *re-discretization equivalent* if there is a homeomorphism $h : \text{geom}(K) \rightarrow \text{geom}(L)$ with $f^\Delta = g^\Delta \circ h$. In other words, (K, f) and (L, g) induce the same geometric object and locally linear map.

Theorem 8 *Suppose that (K, f) and (L, g) are two complexes with \mathbb{R}^d -valued functions defined on their vertex sets, and suppose further that f^Δ and g^Δ are injective embeddings of $\text{geom}(K)$ and $\text{geom}(L)$ into \mathbb{R}^d . Then $PHT(K, f) = PHT(L, g)$ or $ECT(K, f) = ECT(L, g)$ implies that $f^\Delta(\text{geom}(K)) = g^\Delta(\text{geom}(L))$, i.e. (K, f) and (L, g) are re-discretization equivalent, with $h = g^\Delta \circ (f^\Delta)^{-1}$.*

Proof For a weight vector $w \in \mathbb{S}^{d-1}$, the functions $\hat{f} \cdot w$ and f^Δ are identical on the vertex set K^0 , but differ in their extensions to K or $\text{geom}(K)$, as the former

is discrete and the latter continuous. However, Proposition 4 shows that they give the same persistence diagrams (and hence, also, Euler characteristic curves). It is, moreover, easy to see that the filtration on $\text{geom}(K)$ given by $f^\Delta \cdot w$ is identical to the filtration on $f^\Delta(\text{geom}(K))$ given by $f_w(x) = x \cdot w$, identifying $\text{geom}(K)$ with $f^\Delta(\text{geom}(K))$ by dint of the fact that f^Δ is a homeomorphism. Thus, the PHT of (K, f) contains the same data as the PHT of $f^\Delta(\text{geom}(K))$ (and the same is true for the ECT), and so we can apply Theorem 7. \square

In order to use this theorem to prove inverse results for convolutional persistence, we need one more technical lemma.

Definition 11 A collection of points S in \mathbb{R}^d are in *general position* if no hyperplane intersects S in more than d points.

Remark 5 It is easy to see that the property of being in general position is invariant under translation. Moreover, if a set S in general position in \mathbb{R}^d contains the origin as a point $s_0 = \vec{0} \in S$, then any set of d vectors in $S \setminus s_0$ is linearly independent.

Remark 6 Being in general position is a generic property of point sets, meaning that the subset of $(\mathbb{R}^d)^S$ corresponding to subsets in general position is open and dense.

Lemma 1 *Let S be the vertex set of an m -dimensional simplicial complex K^m . Let $\iota : S \hookrightarrow \mathbb{R}^M$ be an embedding with $\iota(S)$ in general position. Then if $M \geq 2m + 1$, the embedding extends by interpolation to an injective embedding $\iota^\Delta : K^m \hookrightarrow \mathbb{R}^M$. Similarly, let S be vertex set of an m -dimensional cubical complex L^m and $\iota : S \hookrightarrow \mathbb{R}^M$ an embedding with $\iota(S)$ in general position. Then if $M \geq 2^{m+1} - 1$, the embedding extends by interpolation to an injective embedding $\iota^\Delta : L^m \hookrightarrow \mathbb{R}^M$.*

Proof We begin with the simplicial complex. A nontrivial intersection of two faces $\iota(\sigma_1)$ and $\iota(\sigma_2)$ for $\sigma_1, \sigma_2 \leq K^m$ implies that the union of their vertex sets $\iota(\sigma_1^0 \cup \sigma_2^0)$ is not linearly dependent (this includes the case when the two faces are identical, i.e. a self-intersection), and in fact this dependence relation has at least two nonzero terms, as a dependence relation with only one nonzero term corresponds to a vertex sitting at the origin, which does not reflect any intersection of faces. Translating one of the vertices in $\iota(\sigma_1^0 \cup \sigma_2^0)$ to the origin and dropping it from the dependence relation, we are still left with a nontrivial dependence relationship among the remaining vertices. Now, a face of K^m has at most $m + 1$ vertices, so the union of two faces involves at most $2m + 2$ vertices. If $M \geq 2m + 1$, translating one of these vertices to the origin means that the remaining vertices are linearly independent, which contradicts the possibility of such a dependence relation.

The proof is similar with the cubical complex. A face of L^m contains at most 2^m vertices, so the union of two faces contains at most 2^{m+1} vertices. If $M \geq 2^{m+1} - 1$, translating one of these 2^{m+1} points to the origin makes the remaining $2^{m+1} - 1$ linearly independent, and hence no intersection can occur. \square

We now prove our inverse results for convolutional persistence.

Simplicial Convolutional Persistence

Recall that for a function $f : K^0 \rightarrow \mathbb{R}^d$ defined on the vertex set of simplicial complex and encoded by an $|K^0| \times d$ matrix X , an arbitrary fixed $|K^0| \times |K^0|$ matrix A , and a $d \times 1$ weight vector w , the simplicial convolutional persistence of the triple (K, A, f) on w is defined by first writing $f * w$ for the function on K^0 defined by AXw and then computing the persistent homology of the lower-star extension $\widehat{f * w}$ on K . We now define a topological transform based on this convolution.

Definition 12 Fix $f : K^0 \rightarrow \mathbb{R}^d$ and A as in the definition of simplicial convolutional persistence. Define the convolutional persistence transform to be the mapping $CPT(K, f, A) : \mathbb{S}^{d-1} \rightarrow \mathbf{Diagrams}$ that sends a weight vector $w \in \mathbb{S}^{d-1}$ to the persistence diagram $\text{PH}(K, \widehat{f * w})$. One can similarly define an *Euler characteristic curve* version of this invariant, the *Convolutional Euler Characteristic Transform* (CECT), by computing Euler characteristic curves instead of persistence diagrams.

Definition 13 We will say that a property is *generically* true if it is true when restricted to some open, dense set of parameters. For example, with simplicial convolutional persistence, the parameters are the function f , encoded by the matrix X , and the generalized adjacency matrix A .

Proposition 9 Assume that $d \geq \dim(K) + 1$. Then f^Δ is generically an injective embedding of $\text{geom}(K)$, and so the CPT and CECT of (K, f, A) determine (K, f) up to re-discretization.

Proof It is a generic property of the matrix AX that its rows are in general position in \mathbb{R}^d . Lemma 1 then implies that f^Δ is injective on the geometric realization. We can then apply Theorem 8. \square

Image Convolutional Persistence

We begin by defining our topological transform.

Definition 14 Fix $f : P \rightarrow \mathbb{R}$. The *Convolutional Persistence Transform* is the mapping $CPT(f) : \mathbb{S}^{|B|-1} \rightarrow \mathbf{Diagrams}$ that sends an L^2 -normalized function $g : B \rightarrow \mathbb{R}$ to the persistence diagram of $f * g$ on K_R^m . One can similarly define an *Euler characteristic curve* version of this invariant, the *Convolutional Euler Characteristic Transform* (CECT), by computing Euler characteristic curves instead of persistence diagrams. Using multi-channel images $f : P \rightarrow \mathbb{R}^n$, we can define an analogous CPT and CECT. In this setting, g ranges over all normalized functions $B \rightarrow \mathbb{R}^n$, which can be identified with $\mathbb{S}^{|B|n-1}$.




Fig. 7 The function f is defined on the 4×4 grid on the top left. The box B is 3×3 , and using a $(1, 1)$ -stride there are four ways of laying B over the domain of f , so that R is a 2×2 grid. We can map the vertices of R into \mathbb{R}^9 by associating each vertex of R with its corresponding translate of B , and then taking as coordinates the values of f in that translate. This extends via interpolation to a map ι_f from the complex K_R^2 into \mathbb{R}^9 , which here is shown to be an embedding.

Remark 7 Taking a trivial convolution where B consists of a single vertex b and $g(b) = 1$, we can recover the original persistence diagram of f . Thus, the CPT is a strictly more general construction than ordinary persistence.

Next, we show how to associate (P, f) with a Euclidean embedding of K_R .

Definition 15 For a fixed function $f : P \rightarrow \mathbb{R}^n$, we obtain a mapping ι_f of the rectangle R into $\mathbb{R}^{|B|n}$ by sending every point $r \in R$ to the vector $\{f(b + k \odot r) \mid b \in B\}$. In other words, r corresponds to a translate of B , and $\iota_f(r)$ records the values of f restricted to this translate. This is technically a set, rather than a vector, but it becomes a vector after fixing an order on the elements of B . Such a mapping can be extended via the lower-* rule to a map $\widehat{\iota}_f : K_R \rightarrow \mathbb{R}^{|B|n}$ on the entire cubical complex K_R built on top of R . See Figure 7.

It is easily seen that, fixing the rectangles P and B and the stride vector k , that $(K_R, \widehat{\iota}_f)$ uniquely determines f , and moreover that $CPT(f) = PHT(K_R, \widehat{\iota}_f)$. To obtain an inverse result in this context, we must show that the continuous interpolation ι_f^Δ is generically injective on $\text{geom}(K_R^m)$. See Figure 8.




Fig. 8 The grid P has shape 2×4 , whereas B has shape 1×2 . Using a convolutional of stride $(2, 1)$, there convolutional grid R has shape 2×2 . The values of the function f are indicated in the grid P . This induces a mapping ι_f from the vertices of R into \mathbb{R}^2 , where the top-left vertex of R gets sent to $(1, 2)$, the top-right to $(3, 2)$, the bottom-left to $(1, -1)$ and the bottom-right to $(3, -1)$. Extending this via interpolation to the complex K_R^2 results in some self-intersections, so that ι_f^Δ is not a homeomorphism on to its image.

Lemma 2 Let $\kappa = \prod_i k_i$, where k_i is the i th entry of the stride vector k , and suppose that $n\kappa \geq 2^{m+1} - 1$. Then a generic function f has the property that ι_f^Δ is injective on the m -skeleton K_R^m , so that $\iota_f^\Delta(\text{geom}(K_R^m))$ has the structure of a cubical complex isomorphic to K_R^m .

Proof Let $B^* \subseteq B$ consist of those elements in the top-left $k_1 \times k_2 \times \dots \times k_d$ corner of B . The various translates $B^* + (k \odot r)$ are all disjoint subsets of P . To show that ι_f^Δ is injective, it suffices to show that it is injective when composed with the coordinate projection $\pi : \mathbb{R}^{|B|n} \rightarrow \mathbb{R}^{|B^*|n}$. However, since the translates of B^* are disjoint, a choice of $\pi \circ \iota_f$ is equivalent to choosing an arbitrary vector in $\mathbb{R}^{|B^*|n}$ for each pixel in R (a collection which will generically be in *general position*), and the map $\pi \circ \iota_f^\Delta = (\pi \circ \iota_f)^\Delta$ is obtained by linearly interpolating this mapping on the interior of the higher-order simplices of K_R^m . By Lemma 1, if $|B^*|n \geq 2^{m+1} - 1$, we know that $\pi \circ \iota_f^\Delta$ is injective, and hence ι_f^Δ is injective. \square

Theorem 10 Fix P, B, k , and assume that $\kappa = \prod_i k_i \geq 2^{m+1} - 1$. Then, generically, $CPT(f) = CPT(g)$ implies $f = g$, and the same is true for the CECT.

Proof Lemma 2 and Theorem 8 imply that $CPT(f)$ or $CECT(f)$ determines $\iota_f^\Delta(\text{geom}(K_R^m))$, which uniquely determines f . \square

Remark 8 What we have just shown is that having a large stride vector, in addition to providing computational speedups by lowering the resolution of the resultant grid, also provides generic injectivity for persistence of higher-dimensional data complexes.

The following corollary shows that to apply Theorem 10 one can ignore those filter functions orthogonal to image patches in our dataset. As (Carlsson et al, 2008) observed, the space of local patches in natural images has very high codimension in the space of all possible patches, and hence in practice the CPT may not require a very high-dimensional collection of filter functions to be effective at distinguishing images.

Definition 16 For a given function $f : P \rightarrow \mathbb{R}^d$, we say that a filter function $g : B \rightarrow \mathbb{R}^d$ is *convolutionally orthogonal* to f if $f * g = 0$, i.e. $(f * g)(r) = 0$ for all $r \in R$. We write f^\ddagger to indicate all functions convolutionally orthogonal to f .

Corollary 2 Let $f_1, f_2 : P \rightarrow \mathbb{R}^d$ be two functions, and write $\mathcal{G} = (f_1^\ddagger \cap f_2^\ddagger)^\perp$, the orthogonal complement of the subspace of all functions convolutionally orthogonal to both f_1 and f_2 . If $\text{Diag}(\widehat{f_1 * g}) = \text{Diag}(\widehat{f_2 * g})$ for all $g \in \mathcal{G}$, then $CPT(f_1) = CPT(f_2)$, and the same is true replacing persistence diagrams with Euler characteristic curves.

Proof If f_1 and f_2 have the same persistence diagram when convolving with a function $g \in \mathcal{G}$, they also have the same diagram after convolving with $g + h$ for $h \in \mathcal{G}^\perp$, due to the linearity of convolution. \square

5 Experiments

In this section, we consider how convolutional persistence compares with ordinary persistence in machine learning applications. We focus on image classification, which involves image convolutional persistence, leaving experiments with simplicial convolutional persistence for future work. Our aim is not to argue that persistence-based methods are superior to other methods for the tasks at hand. Nor we do try to demonstrate the computational advantages of downsampling, as this has already been shown in (Solomon et al, 2021b), so we use convolutions where the original image and the convolved image have the same dimensions, i.e. a unit stride vector $(1, 1)$. In all our computations, we build a simplicial complex wherein the pixels are the top-dimensional simplices, rather than the vertices; this is the second construction in Section 2.2, and is in line with how the Giotto toolkit computes cubical persistence

(Tauzin et al, 2021). Code for running experiments with convolutional persistence can be found at <https://github.com/yesolomon/convpers>.

Classification Tasks

In what follows, we consider five datasets:

1. The UCI digits dataset (Alpaydin and Kaynak, 1998). The images are 8×8 . This dataset has 5620 images.
2. The MNIST dataset (Deng, 2012). The images are 28×28 , and 5000 random images were chosen for the dataset.
3. An MNIST-like dataset of Chinese digits (Nazarpour and Chen, 2017). There are 15 classes, for a baseline accuracy of $1/15 \approx 0.066$. See Figure 9 for examples. To speed up computations, a 2×2 max pooling is applied before computing persistence. The original images are 64×64 , and are 32×32 after max pooling. 5000 random images were chosen the dataset.
4. An MNIST-like dataset of characters in the Devanagari script (Acharya et al, 2015). There are 46 classes, for a baseline accuracy of $1/46 \approx 0.0217$. See Figure 10 for examples. The images are 32×32 , and 5000 random images were chosen for the dataset.
5. A dataset of solutions to the *Kuramoto-Sivashinsky* PDE. This is a PDE of relevance in many systems exhibiting pattern formation (Cuerno and Barabási, 1995; Motta et al, 2012; Villain, 1991; Wolf, 1991; Golovin and Davis, 1998). This PDE was studied in the context of topological machine learning by Adams et al. (Adams et al, 2017), in which they consider the anisotropic form of the PDE:

$$u_t = -\nabla^2 u - \nabla^2 \nabla^2 u + r(u_x)^2 + (u_y)^2.$$

Starting from random initial conditions, we solve these PDEs out to time $t = 15$ for a range of r values, $r \in \{1, 1.25, 1.5, 1.75, 2\}$. The classification task is then parameter estimation: giving a 2D image of the solution to the PDE, guess the r parameter that produced it. See Figure 11 for examples. The images are 50×50 , and 500 images were generated for each r value, giving a dataset of 2500 points.

In all of our datasets, the images are small enough that the convolutions are defined to preserve the resolution, rather than downsample. The goal here is to ignore the computational advantages of convolutional persistence and focus only on the gains in discriminative power. For each dataset, we consider five versions of the CPT:

- Using a trivial 1×1 filter [1], which is equivalent to computing the persistence of the original image.
- Adding, in addition to the trivial filter, three more “standard” filters: sharpening, blur, and Gaussian kernels, all 3×3 .

Chinese Digit Examples




Fig. 9 Examples of chinese digit figures.

Devanagari Script Examples




Fig. 10 Examples of Devanagari script figures.

Kuramoto Sivanshinsky Examples




Fig. 11 Examples of Kuramoto-Sivanshinsky figures. Note how the cells become vertically stretched for larger values of r .




Fig. 12 Zero- and one-dimensional persistence diagram for a chinese digit corresponding to “2”. The sublevelset filtrations begins by adding dark regions and then includes regions with increasingly larger values, indicated with brighter colors. Even with the 10x10 resolution, we see clearly that the dark region has one connected component and two persistent cycles.

- Applying PCA to the space of 3×3 image patches to obtain the top-three principal component patches, and then generating five kernels as random normalized linear combinations of these patches. This approach is motivated by Corollary 2, which says that, for the purposes of injectivity, one can ignore filters orthogonal to those found in the patches of the dataset. We call the resulting kernels *eigenfilters*.
- Taking 5 random, normalized 3×3 filters.
- Taking 25 random, normalized 3×3 filters.

For each version of the CPT, we consider four vectorizations:

- Vectorize using persistence images (default setting, 10×10 resolution), concatenating across both homological dimensions and across filters. Thus, if there are B filters, the resulting vector has length $200B$. See Figure 12 for a visualization of the persistence images of an image in the Chinese digits dataset.
- Vectorize using persistence images, as above, but average along filters rather than concatenating, giving a vector of length 200.
- Vectorize using total persistence, concatenating across both homological dimensions and across filters. Thus, if there are B filters, the resulting vector has length $2B$.
- Vectorize using total persistence, as above, but average along filters rather than concatenating, giving a vector of length 2.

Finally, for each vectorization, we test the topological features on the classification task using the following three models:

- A 3-nearest-neighbors classifier.
- Gradient boosting trees, 10 estimators.
- A neural network model with two hidden layers with 100 units each, leakyRELU activations on the hidden layers and softmax on the final layer, and adam optimizer. Trains for 50 epochs with batch size 4.

Though this series of experiments does not fully explore all the choices, applications, or features of the CPT, it is sufficiently diverse to allow for some interesting insights and the observation of recurrent patterns. Moreover, the experimental pipeline chosen is very simple and does not rely on any careful engineering of features or models.

In Table 5, we indicate the time it takes to compute the persistence diagrams for all the images in the data set and transform the results into concatenated persistence image vectors. Other steps in the computational pipeline, like performing convolutions, are extremely fast, so the computing time for convolutional persistence using 10 filters is very close to 10 times the values found in the table. Distributing the persistence calculations across different processors, as well as convolving with larger stride vectors (to shrink image resolution) both provide major speedups (as discussed earlier).

	UCI Digits	MNIST	Chinese Digits	Devanagari Script	KS PDEs
processing time	2.4s	28.6s	22.8s	21.1s	21.6s

Table 1 Computing Time for Persistence Calculations

As many components of the experimental set-up are stochastic, such as the random filters used and the initializations of the neural nets, we perform 10 simulations for each combination of experimental hyperparameters, setting aside 20% of the data as a testing set to compute model accuracy. For each choice of filters and vectorization, the model with the highest average is chosen, and the standard deviation of its accuracy across the 10 simulations is shown in the bar plot. *NN* indicates *nearest neighbors*, *Tree* is *gradient boosting trees*, and *DL* is *deep learning*, i.e. neural networks.

Results

- Digits: Convolutional persistence provides great improvements in predictive accuracy. The best results come from concatenating features rather than averaging them. Total persistence is very powerful when many random filters are used. When using only five filters, eigenfilters outperforms random filters. For this dataset, the best average accuracies were achieved using deep learning. See Figure 13.

- MNIST: Convolutional persistence against boosts predictive accuracy by a large margin. Similar trends to the digits dataset, although the best models are not always neural networks. See Figure 14.
- Chinese digits: For this dataset, five eigenfilters are significantly better than five random filters. Moreover, the best vectorization is given by concatenating total persistence scores. See Figure 15.
- Devanagari: For this dataset, taking five random filters performs about as well as taking five eigenfilters. The best average performance is given by concatenating persistence images, but similar accuracy with lower standard deviation is given by a nearest neighbor classifier on the concatenated total persistence vectors. See Figure 16.
- Kuramoto-Sivashinsky: The advantages of convolutional persistence are evident but less dramatic. As in the prior two experiments, the concatenated total persistence vectors for 25 random filters is significantly more informative than any vectorizations of ordinary persistence. Better performance (over 95%!) is possible using neural networks on the concatenated persistence image vector, but this is less stable to some of the experimental hyperparameters. See Figure 17.

Observations & Discussion

We conclude this experimental section with a few comments.

- Random filters consistently work well at providing informative features, although there is little theoretical justification for why this should be the case. It is possible that this is related to Johnson-Lindenstrauss theory (Lindenstrauss, 1984), which shows that random projections are good at preserving the geometry of high-dimensional point clouds.
- It is also surprising that total persistence is such an effective vectorization across multiple experiments, despite being so reductive.
- Given that eigenfilters tend to outperform random filters, it would be interesting (and easy) to compare 25 random linear combinations of eigenfilters with 25 random filters.
- In our eigenfilters experiments, we opt for taking multiple random combinations of the top eigenfilters, rather than using the eigenfilters themselves, and taking more of them. This was based on some theoretical heuristics and small-scale experiments, but there is otherwise no reason to suggest why one approach is generally superior to the other.
- We did not apply any feature engineering before applying our machine learning models. For example, the concatenated persistence images vectors were fairly large, 5000-dimensional in the case of 25 random filters, and performance might be improved by pre-applying a dimensionality reduction method like PCA.
- Similarly, we did not do any engineering of our machine learning models, fixing a simple neural network architecture across all experiments. For the




Fig. 13 Results for the digits dataset.




Fig. 14 Results for the MNIST dataset.

Kuramoto-Sivashinsky PDE dataset, the high variance of the neural network models suggests that significantly improved performance might be achievable using the right architecture and training regime.

- Another hyper-parameter we did not test extensively is the size of the filters. Some preliminary experiments suggest that 4×4 filters outperform the 3×3 filters used here, but it remains to be seen if this phenomenon is robust.




Fig. 15 Results for the chinese digits dataset.




Fig. 16 Results for the Devanagari dataset.

6 Conclusion

In this paper, we have proposed a novel approach to topological analysis of images and simplicial complexes, based on combining convolutions and persistent homology. Packaging this invariant as a topological transform, we have shown that this transform is injective in a generic sense by viewing it as a special case of the persistent homology transform. We have also provided diverse experiments showing that convolutions are effective at improving the predictive power of persistent homology for classification tasks. Looking forward, there are multiple directions for future research.




Fig. 17 Results for the Kuramoto-Sivashinsky dataset.

- The injectivity theory for the PHT has been extended by results bounding the number of vector directions needed to obtain injectivity for certain classes of shapes, see (Curry et al, 2018; Belton et al, 2018; Fasy et al, 2019; Belton et al, 2020). It would be interesting to try and obtain similar, specialized results for the embedded cubical complexes appearing in this paper.
- As of yet, there are no *inverse stability* results for the PHT, as were obtained for the topological transform in (Solomon et al, 2021a). That is, the regularity of the inverse map from topological features to shapes has not been investigated. It is possible that the restricted, discrete structure of the CPT provides an easier setting for studying this question.
- Both versions of convolutional persistence explored in this paper are discrete. It would be interesting to study convolutional persistence in the continuous setting. The major obstacle to obtaining an inverse result in this setting is that the inverse theory for the persistent homology transform, as derived in (Turner et al, 2014; Curry et al, 2018; Ghrist et al, 2018), is that one must work within an O -minimal structure of definable sets, so that one must prove that all Euclidean embeddings used are definable in a precise sense. This is immediate when working with discrete objects but presents a novel technical challenge otherwise, and so falls outside of the scope of this paper.
- There are many choices to be made in the implementation of convolutional persistence, and it is not clear which give the best result. For example, whether using a big stride and many filters or a small stride and few filters is preferable. Another experiment is to consider whether it is best to work with eigenfilters directly, rather than random normalized linear combinations thereof, as we have done here.
- We did not consider a pipeline in which the convolutional filters are trainable parameters learned using the dataset. This is entirely possible, see (Carriere

et al, 2021; Solomon et al, 2021b), but presents multiple computational challenge, as topological optimization is expensive and unstable, and the topological energy landscape has many poor local minima. Still, this is an interesting direction for future work.

- We considered persistence images and total persistence as vectorizations in our experimental pipeline. Other possible vectorizations include persistence landscapes (Bubenik et al, 2015) and Euler characteristic curves.
- The datasets explored here, though not synthetic, are fairly simplistic, and such classification tasks have been *solved* by deep learning methods which surpasses human accuracy (Russakovsky et al, 2015). A natural next step is to consider more challenging settings, particularly in the data-starved regime, where topological features have proved useful (Khramtsova et al, 2022) for obtaining start-of-the-art performance, and see if convolutional persistence can further boost the accuracy of such models.
- Finally, it would be interesting to experiment with multi-channel images (e.g. RGB α), 3D images (as often arise in medical imaging), as well as datasets of simplicial complexes (like graphs, meshes).

References

- Acharya S, Pant AK, Gyawali PK (2015) Deep learning based large scale handwritten devanagari character recognition. In: 2015 9th International conference on software, knowledge, information management and applications (SKIMA), IEEE, pp 1–6
- Adams H, Emerson T, Kirby M, et al (2017) Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research* 18
- Alpaydin E, Kaynak C (1998) Optical recognition of handwritten digits data set. UCI Machine Learning Repository
- Aukerman A, Carrière M, Chen C, et al (2020) Persistent homology based characterization of the breast cancer immune microenvironment: a feasibility study. In: 36th International Symposium on Computational Geometry (SoCG)
- Belton RL, Fasy BT, Mertz R, et al (2018) Learning simplicial complexes from persistence diagrams. arXiv preprint arXiv:1805.10716
- Belton RL, Fasy BT, Mertz R, et al (2020) Reconstructing embedded graphs from persistence diagrams. *Computational Geometry* 90:101,658
- Bendich P, Bubenik P, Wagner A (2020) Stabilizing the unstable output of persistent homology computations. *Journal of Applied and Computational Topology* 4(2):309–338

- Bestvina M, Brady N (1997) Morse theory and finiteness properties of groups. *Inventiones mathematicae* 129(3):445–470
- Bleile B, Garin A, Heiss T, et al (2021) The persistent homology of dual digital image constructions. arXiv preprint arXiv:210211397
- Botnan MB, Lesnick M (2022) An introduction to multiparameter persistence. arXiv preprint arXiv:220314289
- Bubenik P, Wagner A (2020) Embeddings of persistence diagrams into hilbert spaces. *Journal of Applied and Computational Topology* 4(3):339–351
- Bubenik P, et al (2015) Statistical topological data analysis using persistence landscapes. *J Mach Learn Res* 16(1):77–102
- Buchet M, Chazal F, Dey TK, et al (2014) Topological analysis of scalar fields with outliers. arXiv preprint arXiv:14121680
- Calcina SS, Gameiro M (2021) Parameter estimation in systems exhibiting spatially complex solutions via persistent homology and machine learning. *Mathematics and Computers in Simulation* 185:719–732
- Carlsson G (2009) Topology and data. *Bulletin of the American Mathematical Society* 46(2):255–308
- Carlsson G, Ishkhanov T, De Silva V, et al (2008) On the local behavior of spaces of natural images. *International journal of computer vision* 76(1):1–12
- Carrière M, Oudot SY, Ovsjanikov M (2015) Stable topological signatures for points on 3d shapes. In: *Computer graphics forum*, Wiley Online Library, pp 1–12
- Carriere M, Chazal F, Glisse M, et al (2021) Optimizing persistent homology based functions. In: *International conference on machine learning*, PMLR, pp 1294–1303
- Chung YM, Day S (2018) Topological fidelity and image thresholding: A persistent homology approach. *Journal of Mathematical Imaging and Vision* 60(7):1167–1179
- Chung YM, Day S, Hu CS (2022) A multi-parameter persistence framework for mathematical morphology. *Scientific reports* 12(1):1–25
- Cohen-Steiner D, Edelsbrunner H, Harer J (2007) Stability of persistence diagrams. *Discrete & computational geometry* 37(1):103–120
- Crawford L, Monod A, Chen AX, et al (2020) Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis.

Journal of the American Statistical Association 115(531):1139–1150

- Cuerno R, Barabási AL (1995) Dynamic scaling of ion-sputtered surfaces. *Physical review letters* 74(23):4746
- Curry J (2018) The fiber of the persistence map for functions on the interval. *Journal of Applied and Computational Topology* 2(3):301–321
- Curry J, Mukherjee S, Turner K (2018) How many directions determine a shape and other sufficiency results for two topological transforms. *arXiv preprint arXiv:180509782*
- Deng L (2012) The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine* 29(6):141–142
- Di Fabio B, Ferri M (2015) Comparing persistence diagrams through complex vectors. In: *International Conference on Image Analysis and Processing*, Springer, pp 294–305
- Edelsbrunner H, Harer J (2010) *Computational topology: an introduction*
- Fasy BT, Micka S, Millman DL, et al (2019) Persistence diagrams for efficient simplicial complex reconstruction. *arXiv preprint arXiv:191212759* 1:5
- Gabrielsson RB, Nelson BJ, Dwaraknath A, et al (2020) A topology layer for machine learning. In: *International Conference on Artificial Intelligence and Statistics*, PMLR, pp 1553–1563
- Gameiro M, Hiraoka Y, Obayashi I (2016) Continuation of point clouds via persistence diagrams. *Physica D: Nonlinear Phenomena* 334:118–132
- Ghrist R (2008) Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society* 45(1):61–75
- Ghrist R, Levanger R, Mai H (2018) Persistent homology and euler integral transforms. *Journal of Applied and Computational Topology* 2(1):55–60
- Giunti B, Houry G, Kerber M (2021) Average complexity of matrix reduction for clique filtrations. *arXiv preprint arXiv:211102125*
- Golovin AA, Davis SH (1998) Effect of anisotropy on morphological instability in the freezing of a hypercooled melt. *Physica D: Nonlinear Phenomena* 116(3-4):363–391
- Hewitt E, Ross KA (2012) *Abstract Harmonic Analysis: Volume I Structure of Topological Groups Integration Theory Group Representations*, vol 115. Springer Science & Business Media

- Hiraoka Y, Nakamura T, Hirata A, et al (2016) Hierarchical structures of amorphous solids characterized by persistent homology. *Proceedings of the National Academy of Sciences* 113(26):7035–7040
- Hu X, Li F, Samaras D, et al (2019) Topology-preserving deep image segmentation. *Advances in neural information processing systems* 32
- Jiang Q, Kurtek S, Needham T (2020) The weighted euler curve transform for shape and image analysis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp 844–845
- Kaczynski T, Mischaikow KM, Mrozek M (2004) *Computational homology*, vol 3. Springer
- Khramtsova E, Zuccon G, Wang X, et al (2022) Rethinking persistent homology for visual recognition. *arXiv e-prints* pp arXiv-2207
- Kim K, Kim J, Zaheer M, et al (2020) Pllay: Efficient topological layer based on persistent landscapes. *Advances in Neural Information Processing Systems* 33:15,965–15,977
- Kipf TN, Welling M (2016) Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:160902907*
- Lacombe T, Cuturi M, Oudot S (2018) Large scale computation of means and clusters for persistence diagrams using optimal transport. *Advances in Neural Information Processing Systems* 31
- Leygonie J, Henselman-Petrusek G (2021) Algorithmic reconstruction of the fiber of persistent homology on cell complexes. *arXiv preprint arXiv:211014676*
- Leygonie J, Tillmann U (2022) The fiber of persistent homology for simplicial complexes. *Journal of Pure and Applied Algebra* p 107099
- Lindenstrauss WJJ (1984) Extensions of lipschitz maps into a hilbert space. *Contemp Math* 26(189-206):2
- Maria C, Oudot S, Solomon E (2019) Intrinsic topological transforms via the distance kernel embedding. *arXiv preprint arXiv:191202225*
- Milosavljević N, Morozov D, Skraba P (2011) Zigzag persistent homology in matrix multiplication time. In: *Proceedings of the twenty-seventh Annual Symposium on Computational Geometry*, pp 216–225
- Monod A, Kalisnik S, Patino-Galindo JÁ, et al (2019) Tropical sufficient statistics for persistent homology. *SIAM Journal on Applied Algebra and Geometry* 3(2):337–371

- Motta FC, Shipman PD, Bradley RM (2012) Highly ordered nanoscale surface ripples produced by ion bombardment of binary compounds. *Journal of Physics D: Applied Physics* 45(12):122,001
- Nazarpour K, Chen M (2017) Handwritten Chinese Numbers <https://doi.org/10.17634/137930-3>, URL https://data.ncl.ac.uk/articles/dataset/Handwritten_Chinese_Numbers/10280831
- Otter N, Porter MA, Tillmann U, et al (2017) A roadmap for the computation of persistent homology. *EPJ Data Science* 6:1–38
- Oudot S, Solomon E (2017) Barcode embeddings for metric graphs. arXiv preprint arXiv:171203630
- Oudot S, Solomon E (2020) Inverse problems in topological persistence. In: *Topological Data Analysis*. Springer, p 405–433
- Oudot SY (2015) Persistence theory: from quiver representations to data analysis, vol 209. American Mathematical Society Providence
- Oudot SY (2017) Persistence theory: from quiver representations to data analysis, vol 209. American Mathematical Soc.
- Pham K, Le K, Ho N, et al (2020) On unbalanced optimal transport: An analysis of sinkhorn algorithm. In: *International Conference on Machine Learning*, PMLR, pp 7673–7682
- Poulenard A, Skraba P, Ovsjanikov M (2018) Topological function optimization for continuous shape matching. In: *Computer Graphics Forum*, Wiley Online Library, pp 13–25
- Russakovsky O, Deng J, Su H, et al (2015) Imagenet large scale visual recognition challenge. *International journal of computer vision* 115(3):211–252
- Skraba P, Turner K (2020) Wasserstein stability for persistence diagrams. arXiv preprint arXiv:200616824
- Solomon E, Wagner A, Bendich P (2021a) From geometry to topology: Inverse theorems for distributed persistence. arXiv preprint arXiv:210112288
- Solomon E, Wagner A, Bendich P (2022) From Geometry to Topology: Inverse Theorems for Distributed Persistence. In: Goao X, Kerber M (eds) 38th International Symposium on Computational Geometry (SoCG 2022), Leibniz International Proceedings in Informatics (LIPIcs), vol 224. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp 61:1–61:16, <https://doi.org/10.4230/LIPIcs.SoCG.2022.61>, URL <https://drops.dagstuhl.de/opus/volltexte/2022/16069>

- Solomon Y, Wagner A, Bendich P (2021b) A fast and robust method for global topological functional optimization. In: International Conference on Artificial Intelligence and Statistics, PMLR, pp 109–117
- Suzuki A, Miyazawa M, Minto JM, et al (2021) Flow estimation solely from image data through persistent homology analysis. *Scientific reports* 11(1):1–13
- Tauzin G, Lupo U, Tunstall L, et al (2021) giotto-tda: A topological data analysis toolkit for machine learning and data exploration. *J Mach Learn Res* 22(39):1–6
- Turner K, Mukherjee S, Boyer DM (2014) Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA* 3(4):310–344
- Villain J (1991) Continuum models of crystal growth from atomic beams with and without desorption. *Journal de physique I* 1(1):19–42
- Villani C (2021) *Topics in optimal transportation*, vol 58. American Mathematical Soc.
- Wagner A (2021) Nonembeddability of persistence diagrams with p_i 2 wasserstein metric. *Proceedings of the American Mathematical Society* 149(6):2673–2677
- Wolf DE (1991) Kinetic roughening of vicinal surfaces. *Physical review letters* 67(13):1783