

Optimizing persistent homology based functions

Mathieu Carrière* Frédéric Chazal† Marc Glisse† Yuichi Ike‡
Hariprasad Kannan

February 18, 2021

Abstract

Solving optimization tasks based on functions and losses with a topological flavor is a very active, growing field of research in data science and Topological Data Analysis, with applications in non-convex optimization, statistics and machine learning. However, the approaches proposed in the literature are usually anchored to a specific application and/or topological construction, and do not come with theoretical guarantees. To address this issue, we study the differentiability of a general map associated with the most common topological construction, that is, the persistence map. Building on real analytic geometry arguments, we propose a general framework that allows us to define and compute gradients for persistence-based functions in a very simple way. We also provide a simple, explicit and sufficient condition for convergence of stochastic subgradient methods for such functions. This result encompasses all the constructions and applications of topological optimization in the literature. Finally, we provide associated code, that is easy to handle and to mix with other non-topological methods and constraints, as well as some experiments showcasing the versatility of our approach.

1 Introduction

Persistent homology is a central tool in Topological Data Analysis that allows to efficiently infer relevant topological features of complex data in a descriptor called persistence diagram. It found many applications in Machine Learning (ML) where it initially played the role of a feature engineering tool through the direct use of persistence diagrams, or through dedicated ML architectures to handle them—see, e.g., [HKNU17, Ume17, CCI⁺20, DUC20, KKZ⁺20]. For the last few years, a growing number of works have been successfully using persistence theory in different ways in order to, for instance, better understand, design and improve neural network architectures—see, e.g., [RTB⁺18, MHRB19, CG20, GC19]—or design regularization and loss functions incorporating topological terms and penalties for various ML tasks—see, e.g., [CNBW19, HKND19, CBO⁺20]. These new use cases of persistence generally involve minimizing functions that depend on persistence diagrams. Such functions are in general highly non-convex and not differentiable, and thus their theoretical and practical minimization can be difficult. In some specific cases, persistence-based functions can be designed to be differentiable and/or some effort have to be made to compute their gradient, so that standard gradient descent techniques can be used to minimize them—see e.g., [WLSC20, PSO18, BGGSSG20]. In the general case, recent attempts have been made to better understand their differential structures [LOT19]. Moreover, building on powerful tools provided by software libraries such as PyTorch or TensorFlow, practical methods allowing to encode and optimize a large family of persistence-based functions have been proposed and experimented [GNDS20, SWB20]. However, in all these cases, the algorithms used to minimize these functions do not come with theoretical guarantees of convergence to a global or local minimum.

Contributions and organization of the article. The aim of this article is to provide a general framework that includes almost all persistence-based functions from the literature, and for which stochastic subgradient descent algorithms are easy to implement and come with convergence guarantees.

More precisely, we first observe that the persistence map, converting a filtration over a given simplicial complex¹ into a persistence diagram, can be thought of as a map between Euclidean spaces (Section 2). This observation allows us to prove that the persistence map is semi-algebraic and, using classical arguments from o-minimal geometry, to study the differentiability of the persistence of parametrized families of filtrations (Section 3). Then, building on the recent work [DDKL20], we consider the minimization problem of persistence-based functions and show that under mild assumptions, stochastic subgradient descent algorithms applied to such functions converge almost surely to a critical point (Section 4). We also provide a simple corresponding Python implementation² for minimizing functions of persistence, and we illustrate it with several examples from the literature (Section 5).

2 Filtrations and persistence diagrams

In this section, we show that the persistence map is nothing but a permutation of the coordinates of a vector containing the filtration values.

*DataShape, Inria Sophia-Antipolis, Biot, France

†DataShape, Inria Saclay, Palaiseau, France

‡Fujitsu Laboratories, AI Lab, Tokyo, Japan

¹The presentation is restricted to simplicial complexes for simplicity, but this generalizes to other complexes as well. We present an example with cubical complexes in Appendix.

²<https://github.com/MathieuCarriere/difftda.git>

2.1 Simplicial complexes and filtrations

Recall that given a set V , a (finite) *simplicial complex* K is a collection of finite subsets of V that satisfies (1) $\{v\} \in K$ for any $v \in V$, and (2) if $\sigma \in K$ and $\tau \subseteq \sigma$ then $\tau \in K$. An element $\sigma \in K$ with $|\sigma| = k + 1$ is called a k -simplex.

Given a simplicial complex K and a subset R of \mathbb{R} , a *filtration* of K is an increasing sequence $(K_r)_{r \in R}$ of subcomplexes of K with respect to the inclusion, i.e., $K_r \subseteq K_s$ for any $r \leq s$, and such that $\bigcup_{r \in R} K_r = K$.

To each simplex $\sigma \in K$, one can associate its *filtering index* $\Phi_\sigma = \inf\{r \in R : \sigma \in K_r\}$. Thus, when K is finite, a filtration of K can be conveniently represented as a filtering function $\Phi: K \rightarrow \mathbb{R}$. Equivalently, it can be represented as a $|K|$ -dimensional vector $\Phi = (\Phi_\sigma)_{\sigma \in K}$ in $\mathbb{R}^{|K|}$ whose coordinates are the indices of the simplices of K and that satisfies the following condition: if $\sigma, \tau \in K$ and $\tau \subseteq \sigma$, then $\Phi_\tau \leq \Phi_\sigma$. As a consequence, if the vectorized filtration Φ depends on a parameter, the corresponding family of filtrations can be represented as a map from the space of parameters to $\mathbb{R}^{|K|}$ in the following way.

Definition 2.1. Let K be a simplicial complex and A a set. A map $\Phi: A \rightarrow \mathbb{R}^{|K|}$ is said to be a *parametrized family of filtrations* if for any $x \in A$ and $\sigma, \tau \in K$ with $\tau \subseteq \sigma$, one has $\Phi_\tau(x) \leq \Phi_\sigma(x)$.

2.2 Persistence computation from filtrations

We briefly recall how the computation of the persistence diagram of a filtered simplicial complex decomposes into: (i) a purely combinatorial part only relying on the order on the simplices induced by the filtration, and (ii) a part relying on the filtration values. A detailed introduction to persistent homology and its computation can be found in, e.g., [EH10, BCY18].

Let K be a simplicial complex endowed with a filtration and corresponding filtering function $\Phi \in \mathbb{R}^{|K|}$, where $|K|$ is the number of non-empty simplices of K .

First part: combinatorial part (persistence pairs). The filtering function Φ induces a total preorder on the elements of K as follows: $\tau \preceq \sigma$ if $\Phi_\tau \leq \Phi_\sigma$. This preorder can be refined into a total order by breaking ties in some fairly arbitrary way, as long as it is consistent with the face relation, i.e., if $\tau \subseteq \sigma$, then $\tau \preceq \sigma$. One way to break ties is to sort simplices that have the same filtration value by dimension, and then order the ones that are still equivalent according to some arbitrary indexing of the simplices. A different way is to index the vertices, represent simplices by their decreasing list of vertices, and sort equivalent simplices using the lexicographic order on those lists. In the following, we will assume that the total order is a function of the preorder, in particular it is deterministic and does not depend on the exact values of Φ . Note that while different orders may yield different pairings, they all translate to the same persistence diagram in the second part. The basic algorithm to compute persistence iterates over the ordered set of simplices $\sigma_1 \preceq \dots \preceq \sigma_{|K|}$ according to Algorithm 1 below—see Section 11.5.2 in [BCY18] for a detailed description of the algorithm.

```

Input: Ordered sequence of simplices  $\sigma_1 \preceq \dots \preceq \sigma_{|K|}$ 
 $K_0 = \emptyset$ 
Pairs $_0 = \text{Pairs}_1 = \dots = \text{Pairs}_{d-1} = \emptyset$ 
for  $j = 1$  to  $|K|$  do
   $k = \dim \sigma_j$ 
   $K_j = K_{j-1} \cup \sigma_j$ 
  if  $\sigma_j$  does not create a new  $k$ -dimensional homology class in  $K_j$  then
    a  $(k - 1)$ -dimensional homology class created in  $K_{l(j)}$  by  $\sigma_{l(j)}$  for some  $l(j) < j$  becomes homologous
    to 0 in  $K_j$ .
    Pairs $_{k-1} \leftarrow \text{Pairs}_{k-1} \cup \{(\sigma_{l(j)}, \sigma_j)\}$ ;
  end if
end for
Output: Persistence pairs in each dimension Pairs $_0, \text{Pairs}_1, \dots, \text{Pairs}_{d-1}$ 

```

Algorithm 1: Persistence pairs computation (sketch)

Note that for each dimension k , some k -dimensional simplices may remain unpaired at the end of the algorithm; their number is equal to the k -dimensional Betti number of K .

Second part: associated filtration values. The persistence diagram of the filter function Φ is now obtained by associating to each persistent pair $(\sigma_{l(j)}, \sigma_j)$ the point $(\Phi_{\sigma_{l(j)}}, \Phi_{\sigma_j})$. Moreover, to each unpaired simplex σ_l is associated the point $(\Phi_{\sigma_l}, +\infty)$.

If p is the number of persistence pairs and q is the number of unpaired simplices, then $|K| = 2p + q$ and the persistence diagram $D(\Phi)$ of the filtration Φ of K is made of p points in \mathbb{R}^2 (counted with multiplicity) and q points (also counted with multiplicity) with infinite second coordinate. Choosing the lexicographical order on $\mathbb{R} \times (\mathbb{R} \cup \{+\infty\})$, the persistence diagram $D(\Phi)$ can be represented as a vector in $\mathbb{R}^{|K|}$ and the output of the persistence algorithm can be simply seen as a permutation of the coordinates of the input vector Φ . Moreover, this permutation only depends on the order on the simplices of K induced by Φ .

Definition 2.2. The subset of points of a persistence diagram D with finite coordinates (resp. infinite second coordinate) is called the *regular part* (resp. *essential part*) of D and denoted by D_{reg} (resp. D_{ess}).

With the notations defined above, D_{reg} and D_{ess} can be represented as vectors in \mathbb{R}^{2p} and \mathbb{R}^q , respectively.

Note that, in practice, the above construction is usually done dimension by dimension to get a persistence diagram for each dimension in homology, by restricting to the subset of simplices of dimension k and $k + 1$. Without loss of generality, and to avoid unnecessary heavy notation, in the following we consider the whole persistence diagram, made of the union of the persistence diagrams in all dimensions k .

3 Differentiability of functions of persistence

o-minimal geometry provides a well-suited setting to describe the parametrized families of filtrations encountered in practice and to exhibit interesting differentiability properties of their composition with the persistence map.

3.1 Background on o-minimal geometry

In this section, we recall some elements of o-minimal geometry, which are needed in the next sections—see, e.g., [Cos00] for a more detailed introduction.

Definition 3.1 (o-minimal structure). An *o-minimal structure* on the field of real numbers \mathbb{R} is a collection $(S_n)_{n \in \mathbb{N}}$, where each S_n is a set of subsets of \mathbb{R}^n such that:

- (1) S_1 is exactly the collection of finite unions of points and intervals;
- (2) all algebraic subsets³ of \mathbb{R}^n are in S_n ;
- (3) S_n is a Boolean subalgebra of \mathbb{R}^n for any $n \in \mathbb{N}$;
- (4) if $A \in S_n$ and $B \in S_m$, then $A \times B \in S_{n+m}$;
- (5) if $\pi: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ is the linear projection onto the first n coordinates and $A \in S_{n+1}$, then $\pi(A) \in S_n$.

An element $A \in S_n$ for some $n \in \mathbb{N}$ is called a *definable set* in the o-minimal structure. For a definable set $A \subseteq \mathbb{R}^n$, a map $f: A \rightarrow \mathbb{R}^m$ is said to be *definable* if its graph is a definable set in \mathbb{R}^{n+m} .

Definable sets are stable under various geometric operations. The complement, closure and interior of a definable set are definable sets. The finite unions and intersections of definable sets are definable. The image of a definable set by a definable map is itself definable. Sums and products of definable functions as well as compositions of definable functions are definable—see Section 1.3 in [Cos00]. In particular, the max and min of finite sets of real-valued definable functions are also definable. An important property of definable sets and definable maps is that they admit a finite Whitney stratification. This implies that (i) any definable set $A \subseteq \mathbb{R}^n$ can be decomposed into a finite disjoint union of smooth submanifolds of \mathbb{R}^n and (ii) for any definable map $\Phi: A \rightarrow \mathbb{R}^m$, A can also be decomposed into a finite union of smooth manifolds such that the restriction of Φ on each of these manifolds is a differentiable function.

The simplest example of o-minimal structures is given by the family of semi-algebraic subsets⁴ of \mathbb{R}^n ($n \in \mathbb{N}$). Although most of the classical parametrized families of filtrations are semi-algebraic, the o-minimal framework actually allows to consider larger families. In particular, the result of [Wil96] says that the family of images of the sublevel sets of functions in $\mathbb{R}[x_1, \dots, x_N, \exp(x_1), \dots, \exp(x_N)]$ for some $N \in \mathbb{N}$ under linear projections is an o-minimal structure, which allows us to mix exponential functions with semi-algebraic functions.

3.2 Persistence diagrams of definable parametrized families of filtrations

Let K be a simplicial complex and $\Phi: A \rightarrow \mathbb{R}^{|K|}$ be a parametrized family of filtrations that is definable in a given o-minimal structure. If for any $x, x' \in A$, the preorders induced by $\Phi(x)$ and $\Phi(x')$ on the simplices of K are the same, i.e., for any $\sigma_1, \sigma_2 \in K$, $\Phi_{\sigma_1}(x) \leq \Phi_{\sigma_2}(x)$ if and only if $\Phi_{\sigma_1}(x') \leq \Phi_{\sigma_2}(x')$, then the pairs of simplices $(\sigma_{i_1}, \sigma_{j_1}), \dots, (\sigma_{i_p}, \sigma_{j_p})$, and the unpaired simplices $\sigma_{i_{p+1}}, \dots, \sigma_{i_{p+q}}$ that are computed by the persistence Algorithm 1 are independent of x . As a consequence, the persistence diagram $D = D(\Phi(x))$ of $\Phi(x)$ is

$$D = \bigcup_{k=1}^p (\Phi_{\sigma_{i_k}}(x), \Phi_{\sigma_{j_k}}(x)) \cup \bigcup_{k=1}^q (\Phi_{\sigma_{i_{p+k}}}(x), +\infty), \quad (1)$$

where $|K| = 2p + q$.

Given the lexicographic order on $\mathbb{R} \times (\mathbb{R} \cup \{+\infty\})$, the points of any finite multi-set $D \subseteq \mathbb{R} \times (\mathbb{R} \cup \{+\infty\})$ with p points in \mathbb{R}^2 and q points in $\mathbb{R} \times \{+\infty\}$ can be ordered in non-decreasing order, and D can be represented as a vector in \mathbb{R}^{2p+q} . As a consequence, denoting by Filt_K the set of vectors in $\mathbb{R}^{|K|}$ defining a filtration on K , the persistence map $\text{Pers}: \text{Filt}_K \rightarrow \mathbb{R}^{|K|}$ that assigns to each filtration of K its persistence diagram consists in a permutation of the coordinates of $\mathbb{R}^{|K|}$. This permutation is constant on the set of filtrations that define the same preorder on the simplices of K . This leads to the following statement.

Proposition 3.2. *Given a simplicial complex K , the map $\text{Pers}: \text{Filt}_K \subseteq \mathbb{R}^{|K|} \rightarrow \mathbb{R}^{|K|}$ is semi-algebraic, and thus definable in any o-minimal structure. Moreover, there exists a semi-algebraic partition of Filt_K such that the restriction of Pers to each element of this partition is a Lipschitz map.*

Proof. See Appendix. □

Since there exists a finite semi-algebraic partition of Filt_K on which Pers is a locally constant permutation, the subdifferential (see Section 4 for the definition) of Pers is well-defined and obvious to compute: each coordinate in the output (i.e., the persistence diagram) is a copy of a coordinate in the input (i.e., the filtration values of the simplices). This implies that every partial derivative is either 1 or 0. The output can be seen as a reindexing of the input, and this is indeed how we implement it in our code, so that automatic differentiation frameworks (PyTorch, TensorFlow, etc.) can process the function Pers directly and do not need explicit gradient formulas—see Section 5. Note that the subdifferential depends on the arbitrary refinement of the preorder in Subsection 2.2.

³Recall that an algebraic set is the 0-level set of a polynomial.

⁴It is the family of all finite unions and intersections of level sets and sublevel sets of polynomials [BR91].

Corollary 3.3. *Let K be a simplicial complex and $\Phi: A \rightarrow \mathbb{R}^{|K|}$ be a definable (in a given \mathfrak{o} -minimal structure) parametrized family of filtrations. The map $\text{Pers} \circ \Phi: A \rightarrow \mathbb{R}^{|K|}$ is definable.*

Note that according to the remark following Proposition 3.2, if Φ is differentiable, the subdifferential of $\text{Pers} \circ \Phi$ can be easily computed in terms of the partial derivatives of Φ using, for example, Equation (1).

It also follows from standard finiteness and stratifiability properties of definable sets and maps that $\text{Pers} \circ \Phi$ is differentiable almost everywhere. More precisely:

Proposition 3.4. *Let K be a simplicial complex and $\Phi: A \rightarrow \mathbb{R}^{|K|}$ a definable parametrized family of filtrations, where $\dim A = m$. Then there exists a finite definable partition of A , $A = S \sqcup O_1 \sqcup \dots \sqcup O_k$ such that $\dim S < \dim A := m$ and, for any $i = 1, \dots, k$, O_i is a definable manifold of dimension m and $\text{Pers} \circ \Phi: O_i \rightarrow \mathcal{D}$ is differentiable.*

3.3 Examples of definable families of filtrations

Vietoris-Rips filtrations. The family of Vietoris-Rips filtrations built on top of sets of n points $x_1, \dots, x_n \in \mathbb{R}^d$ is the semi-algebraic parametrized family of filtrations

$$\Phi: A = (\mathbb{R}^d)^n \rightarrow \mathbb{R}^{|\Delta_n|} = \mathbb{R}^{2^n - 1},$$

where Δ_n is the simplicial complex made of all the faces of the $(n-1)$ -dimensional simplex, defined, for any $x = (x_1, \dots, x_n) \in A$ and any simplex $\sigma \subseteq \{1, \dots, n\}$, by

$$\Phi_\sigma(x) = \max_{i,j \in \sigma} \|x_i - x_j\|.$$

One easily checks that the permutation induced by Pers is constant on the connected components of the complement of the union of the subspaces $S_{i,j,k,l} = \{(x_1, \dots, x_n) : \|x_i - x_j\| = \|x_k - x_l\|\}$ over all the 4-tuples (i, j, k, l) such that at least 3 of the 4 indices i, j, k, l are distinct. This example naturally extends to Vietoris-Rips-like filtrations in the following way. Let $A \subset \mathcal{M}_n(\mathbb{R})$ be the set of $n \times n$ symmetric matrices with non-negative entries and 0 on the diagonal. This is a semi-algebraic subset of the space of n -by- n matrices $\mathcal{M}_n(\mathbb{R}) \simeq \mathbb{R}^{n^2}$, of dimension $m = (n-1)(n-2)/2$. The map $\Phi: A \rightarrow \mathbb{R}^{|\Delta_n|} = \mathbb{R}^{2^n}$ defined by $\Phi_\sigma(M) = \max_{i,j \in \sigma} m_{i,j}$ for any $M = (m_{i,j})_{1 \leq i,j \leq n} \in A$, is a semi-algebraic family of filtrations. Note that the set S of Proposition 3.4 can be chosen to be the set of matrices with at least 2 entries that are equal.

Weighted Rips filtrations. Given a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, the family of weighted Rips filtrations $\Phi: A = (\mathbb{R}^d)^n \rightarrow \mathbb{R}^{|\Delta_n|} = \mathbb{R}^{2^n}$ associated to f is defined, for any $x = (x_1, \dots, x_n) \in A$ and any simplex $\sigma \subseteq \{1, \dots, n\}$, by

- $\Phi_\sigma(x) = 2f(x_j)$ if $\sigma = [j]$;
- $\Phi_\sigma(x) = \max(2f(x_i), 2f(x_j), \|x_i - x_j\| + f(x_i) + f(x_j))$, if $\sigma = [i, j]$, $i \neq j$;
- $\Phi_\sigma(x) = \max(\Phi_{[i,j]}(x), i, j \in \sigma)$ if $|\sigma| \geq 3$.

Since Euclidean distances and max function are semi-algebraic, this family of filtrations is definable as soon as the weight function f is definable.

This example easily extends to the case where the weight function depends on the set of points $x = (x_1, \dots, x_n)$: the weight at vertex y is defined by $f(x, y)$ with $f: (\mathbb{R}^d)^n \times \mathbb{R}^d \rightarrow \mathbb{R}$. A particular example of such a family is given by the so-called DTM filtration [ACG⁺20], where $f(x, y)$ is the average distance from y to its k -nearest neighbors in x . In this case, f is semi-algebraic, and the family of DTM filtrations is semi-algebraic.

The \mathfrak{o} -minimal framework also allows us to consider weight functions involving exponential functions [Wil96], such as, for instance, kernel-based density estimates with Gaussian kernels.

Sublevel sets filtrations. Let K be a simplicial complex with n vertices v_1, \dots, v_n . Any real-valued function f defined on the vertices of K can be represented as a vector $(f(v_1), \dots, f(v_n)) \in \mathbb{R}^n$. The family of sublevel sets filtrations $\Phi: A = \mathbb{R}^n \rightarrow \mathbb{R}^{|K|}$ of functions on the vertices of K is defined by $\Phi_\sigma(f) = \max_{i \in \sigma} f_i$ for any $f = (f_1, \dots, f_n) \in A$ and any simplex $\sigma \subseteq \{1, \dots, n\}$. This filtration is also known as the *lower-star filtration* of f . The function Φ is obviously semi-algebraic, and for Proposition 3.4 it is sufficient to choose $S = \bigcup_{1 \leq i < j \leq n} \{f = (f_1, \dots, f_n) \in A : f_i = f_j\}$.

4 Minimization of functions of persistence

Using the same notation as in the previous section, recall that the space of persistence diagrams associated to a filtration of K is identified with $\mathbb{R}^{|K|} = (\mathbb{R}^2)^p \times \mathbb{R}^q$, where each point in the p copies of \mathbb{R}^2 is a point with finite coordinates in the persistence diagram and each coordinate in \mathbb{R}^q is the x -coordinate of a point with infinite persistence.

Definition 4.1. A function $E: \mathbb{R}^{|K|} = (\mathbb{R}^2)^p \times \mathbb{R}^q \rightarrow \mathbb{R}$ is said to be a *function of persistence* if it is invariant to permutations of the points of the persistence diagram: for any $(p_1, \dots, p_p, e_1, \dots, e_q) \in (\mathbb{R}^2)^p \times \mathbb{R}^q$ and any permutations α, β of the sets $\{1, \dots, p\}$ and $\{1, \dots, q\}$, respectively, one has

$$E(p_{\alpha(1)}, \dots, p_{\alpha(p)}, e_{\beta(1)}, \dots, e_{\beta(q)}) = E(p_1, \dots, p_p, e_1, \dots, e_q).$$

It follows from this permutation invariance and Proposition 3.2 that if a function of persistence $E: \mathbb{R}^{2p+q} = \mathbb{R}^{|K|} \rightarrow \mathbb{R}$ is locally Lipschitz, then the composition $E \circ \text{Pers}$ is also locally Lipschitz. Moreover, if E is definable in an o-minimal structure, then for any definable parametrized family of filtrations $\Phi: A \subseteq \mathbb{R}^d \rightarrow \mathbb{R}^{|K|}$, the composition $\mathcal{L} = E \circ \text{Pers} \circ \Phi: A \rightarrow \mathbb{R}$ is also definable. As a consequence, \mathcal{L} has a well-defined Clarke subdifferential $\partial\mathcal{L}(z) := \text{Conv}\{\lim_{z_i \rightarrow z} \nabla\mathcal{L}(z_i) : \mathcal{L} \text{ is differentiable at } z_i\}$, since it is differentiable almost everywhere thanks to Proposition 3.4.

4.1 Stochastic gradient descent

To minimize \mathcal{L} , we consider the differential inclusion

$$\frac{dz}{dt} \in -\partial\mathcal{L}(z(t)) \quad \text{for almost every } t,$$

whose solutions $z(t)$ are the trajectories of the subgradient of \mathcal{L} . They can be approximated by the standard stochastic subgradient algorithm given by the iterations of

$$x_{k+1} = x_k - \alpha_k(y_k + \zeta_k), \quad y_k \in \partial\mathcal{L}(x_k), \quad (2)$$

where the sequence $(\alpha_k)_k$ is the learning rate and $(\zeta_k)_k$ is a sequence of random variables. In [DDKL20], the authors prove that under mild technical conditions on these two sequences, the stochastic subgradient algorithm converges almost surely to a critical point of \mathcal{L} as soon as \mathcal{L} is locally Lipschitz.

More precisely, consider the following assumptions, which correspond to Assumption C in [DDKL20]:

- (1) for any k , $\alpha_k \geq 0$, $\sum_{k=1}^{\infty} \alpha_k = +\infty$ and, $\sum_{k=1}^{\infty} \alpha_k^2 < +\infty$;
- (2) $\sup_k \|x_k\| < +\infty$, almost surely;
- (3) denoting by \mathcal{F}_k the increasing sequence of σ -algebras $\mathcal{F}_k = \sigma(x_j, y_j, \zeta_j, j < k)$, there exists a function $p: \mathbb{R}^d \rightarrow \mathbb{R}$ which is bounded on bounded sets such that almost surely, for any k ,

$$\mathbb{E}[\zeta_k | \mathcal{F}_k] = 0 \quad \text{and} \quad \mathbb{E}[\|\zeta_k\|^2 | \mathcal{F}_k] < p(x_k).$$

These assumptions are standard and not very restrictive. Assumption 1 depends on the choice of the learning rate by the user and is easily satisfied, e.g., taking $\alpha_k = 1/k$. Assumption 2 is usually easy to check for most of the functions \mathcal{L} encountered in practice. Assumption 3 is a standard condition, which states that, conditioned upon the past, the variables ζ_k have zero mean and controlled moments; e.g., this can be achieved by taking a sequence of independent and centered variables with bounded variance that are also independent of the x_k 's and y_k 's.

Under these assumptions, the following result is an immediate consequence of Corollary 5.9 in [DDKL20].

Theorem 4.2. *Let K be a simplicial complex, $A \subseteq \mathbb{R}^d$, and $\Phi: A \rightarrow \mathbb{R}^{|K|}$ a parametrized family of filtrations of K that is definable in an o-minimal structure. Let $E: \mathbb{R}^{|K|} \rightarrow \mathbb{R}$ be a definable function of persistence such that $\mathcal{L} = E \circ \text{Pers} \circ \Phi$ is locally Lipschitz. Then, under the above assumptions 1, 2, and 3, almost surely the limit points of the sequence $(x_k)_k$ obtained from the iterations of Equation (2) are critical points of \mathcal{L} and the sequence $(\mathcal{L}(x_k))_k$ converges.*

The above theorem provides explicit conditions ensuring the convergence of stochastic subgradient descent for functions of persistence. The main criterion to be checked is the local Lipschitz condition for \mathcal{L} . From the remark following Definition 4.1, it is sufficient to check that Φ and E are Lipschitz. Regarding Φ , it is obvious for the examples of Subsection 3.3. This is not the case for some other examples, such as the so-called alpha-complex filtration that can be made locally Lipschitz using a simple technical trick—see Appendix.

4.2 Examples of definable locally Lipschitz functions of persistence

Total persistence. Let E be the sum of the distances of each point of a persistence diagram with finite coordinates to the diagonal: given a persistence diagram represented as a vector in \mathbb{R}^{2p+q} , $D = ((b_1, d_1), \dots, (b_p, d_p), e_1, \dots, e_q)$,

$$E(D) = \sum_{i=1}^p |d_i - b_i|.$$

Then E is obviously semi-algebraic, and thus definable in any o-minimal structure. It is also Lipschitz.

Wasserstein and bottleneck distance Given a persistence diagram D^* , the bottleneck distance between the regular part of a diagram D and the regular part of D^* (see Definition 2.2) is given by

$$E(D) = d_B(D_{\text{reg}}, D_{\text{reg}}^*) = \min_m \max_{(p, p^*) \in m} \|p - p^*\|_{\infty},$$

where, denoting $\Delta = \{(x, x) : x \in \mathbb{R}\}$ the diagonal in \mathbb{R}^2 , m is a partial matching between D_{reg} and D_{reg}^* , i.e., a subset of $(D_{\text{reg}} \cup \Delta) \times (D_{\text{reg}}^* \cup \Delta)$ such that every point of $D_{\text{reg}} \setminus \Delta$ and $D_{\text{reg}}^* \setminus \Delta$, appears exactly once in m . One can easily check that the map E is semi-algebraic, and thus definable in any o-minimal structure. It is also Lipschitz. This property also extends to the case where the bottleneck distance is replaced by the so-called Wasserstein distance W_p with $p \in \mathbb{N}$ [CSEHM10], or its approximation, the Sliced Wasserstein distance [CCO17]. Optimization of these functions E and other functions of bottleneck and Wasserstein distances have been used, for example, in shape matching in [PSO18]. See also the example on 3D shape in Appendix.

Persistence landscapes [Bub15] To any given point $p = (x, y) \in \mathbb{R}^2$ with $x = \frac{b+d}{2}$ and $y = \frac{d-b}{2}$, associate the function $\Lambda_p: \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$\Lambda_p(t) = \begin{cases} t - b & (t \in [b, \frac{b+d}{2}]) \\ d - t & (t \in (\frac{b+d}{2}, d]) \\ 0 & (\text{otherwise}). \end{cases}$$

Given a persistence diagram D , the persistence landscape of D is a summary of the arrangement of the graphs of the functions Λ_p , $p \in D$:

$$\lambda_D(k, t) = k\text{-max}_{p \in D} \Lambda_p(t), \quad t \in [0, T], k \in \mathbb{Z}^+,$$

where k -max is the k th largest value in the set, or 0 when the set contains less than k points. Given a positive integer k , a finite set $\{t_1, \dots, t_n\} \subset \mathbb{R}$, and a simplicial complex K , the map that associates the vector $(\lambda_D(k, t_1), \dots, \lambda_D(k, t_n))$ to each persistence diagram D of a filtration of K is Lipschitz [Bub15] and clearly semi-algebraic.

Other classical ways to vectorize persistence diagrams are the *linear representations* [CD18] which are also definable in o-minimal structures, such as, e.g., persistence images [AEK⁺17]—see Appendix. In [DL19], the authors give explicit conditions for such representations to be locally Lipschitz.

5 Numerical illustrations

We showed in Sections 3 and 4 that the usual stochastic gradient descent procedure of Equation (2) enjoys some convergence properties for persistence-based functions. This means in particular that the algorithms available in standard libraries such as TensorFlow and PyTorch, which implement stochastic gradient descent (among other optimization methods), can be leveraged and used as is for differentiating persistence diagrams, while still ensuring convergence. The purpose of this section is to illustrate that our code, which implements the general gradient defined in Proposition 3.4 for persistence-based functions, and which is based on Gudhi⁵ and TensorFlow, can be readily used for studying several different persistence optimization tasks. On the way, we also suggest regularization terms that one can add to topological losses in order to avoid unwanted behaviors. We only present a few applications due to lack of space, and we refer the interested reader to Appendix for more examples.

Point cloud optimization. A toy example in persistence optimization is to modify the positions of the points in a point cloud so that its homology is maximized [GNDS20, GH016]. In this experiment, we start with a point cloud X sampled uniformly from the unit square $S = [0, 1]^2$, and then optimize the point coordinates so that the loss $\mathcal{L}(X) = P(X) + T(X)$ is minimized. Here $T(X) := -\sum_{p \in D} \|p - \pi_\Delta(p)\|_\infty^2$ is a topological penalty, D is the 1-dimensional persistence diagram associated to the Vietoris-Rips filtration of X , π_Δ stands for the projection onto the diagonal Δ , and $P(X) := \sum_{x \in X} d(x, S)$ is a penalty term ensuring that the point coordinates stay in the unit square. The topological penalty $T(X)$ was used in [BGND⁺19], and ensures that holes are created within the point cloud so that the cardinality of the persistence diagram D is as large as possible. However, we point out that if one uses $T(X)$ alone without the penalty $P(X)$, as in [BGND⁺19], then convergence is very difficult to reach since inflating the point cloud with dilations can make the topological penalty $T(X)$ arbitrarily small. In contrast, using our second term $P(X)$ in addition to $T(X)$ constrains the points to stay in a fixed region S of the Euclidean plane. Another effect of the penalty $P(X)$ is to flatten the boundary of the created holes along the boundary of S . See Figure 1 for an illustration.

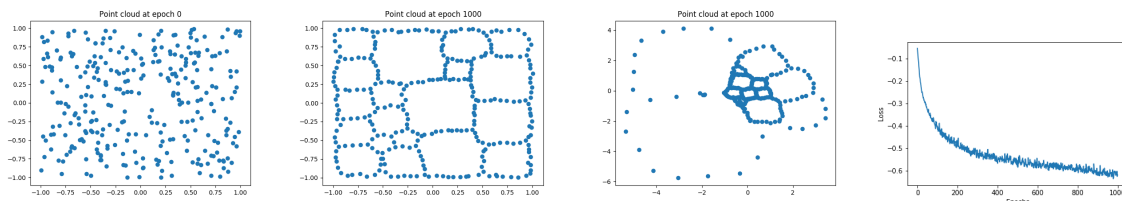


Figure 1: Illustration of point cloud optimization. We initialize with a random point cloud (left), and we show the optimized point cloud when optimization is done with topological and regularization losses (middle left). When only topological loss is used, the optimized point cloud inflated some loops to minimize the loss (middle right). Note how the coordinates are now much larger. We also show the convergence of the total loss (right).

Dimension reduction. In this experiment, we show how our general setup can be used to reduce dimension with the so-called *topological autoencoders* introduced in [MHRB19]. In this family of autoencoders, a topological loss $T(X, Z)$ between the input space X and latent space Z is used in addition to the usual loss $\mathcal{D}(X, Z) = \sum_i \|x_i - z_i\|_2^2$. This topological loss was computed in [MHRB19] by (i) computing the permutations induced by the persistence map (see Subsection 3.2) of the Vietoris-Rips complexes built from the input space X and the latent space Z , (ii) computing, for each simplex in these permutations, the corresponding edge that induces its filtration value, and (iii) measuring, for all those edges, the differences between the edge lengths in X and the same edge lengths in Z . To sum up, the loss function is defined as

$$\mathcal{L}(X, Z) = \|M_X[\pi_X] - M_Z[\pi_X]\|_2^2 + \|M_X[\pi_Z] - M_Z[\pi_Z]\|_2^2,$$

⁵<https://gudhi.inria.fr/>

where M_X, M_Z are the distance matrices of the input and the latent spaces respectively, and where π_X, π_Z denote the indices of the entries in M_X, M_Z that are picked by the permutation induced by the persistence map to generate the Vietoris-Rips persistence diagrams of X and Z . Note that \mathcal{L} is obviously semi-algebraic and thus fits in our framework. Moreover, in our setup we can directly use the bottleneck and Wasserstein distances between the Vietoris-Rips persistence diagrams of the input and latent spaces as the topological loss. This is relevant since in [MHRB19] the authors pointed out that looking at homology in dimension larger than 1 was not adding anything for their loss, and stuck to 0-dimensional homology. We show in Figure 2 an example in which 1-dimensional homology is also important, that is, a point cloud in \mathbb{R}^3 that is comprised of two nested circles, which is then non-linearly embedded in \mathbb{R}^9 by converting each point $p = (x, y, z)$ into the exponential of the 3×3 anti-symmetric matrix whose coefficient are x, y and z . We then train an autoencoder made of four fully-connected layers with 32 neurons and ReLU activations, using the usual loss, the usual plus the topological loss described above, and the usual plus a topological loss computed as $\mathcal{L}(X, Z) = W_1(D_X, D_Z)$, i.e., the 1-Wasserstein distance between the 1-dimensional Vietoris-Rips persistence diagrams of the input and latent spaces. It can be seen from Figure 2 that autoencoders without the Wasserstein loss cannot embed the point cloud in the plane perfectly, while using the Wasserstein loss between the 1-dimensional Vietoris-Rips persistence diagrams improves on the result by separating better the two intrinsic circles.

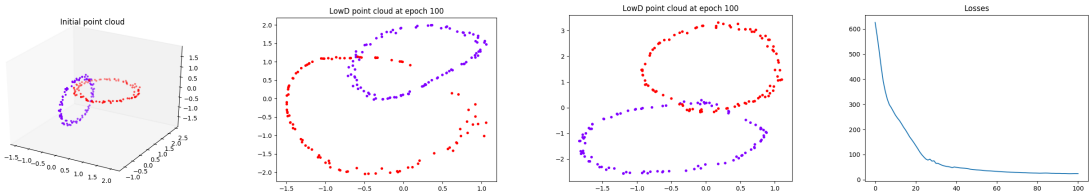


Figure 2: Example of dimension reduction with autoencoders. An initial point cloud made of two circles (left) is embedded in \mathbb{R}^9 , and then fed to autoencoders that either do not use topology, or only use the distances induced by the persistence maps in dimension 0. The resulting embeddings (middle left, we only show one but the two are similar) cannot separate the circles, while using 1-dimensional topology induces a better embedding (middle right). Convergence of the loss function is also provided (right).

Dataset	Baseline	Before	After	Difference	Dataset	Baseline	Before	After	Difference
vs01	100.0	61.3	99.0	+37.6	vs26	99.7	98.8	98.2	-0.6
vs02	99.4	98.8	97.2	-1.6	vs28	99.1	96.8	96.8	0.0
vs06	99.4	87.3	98.2	+10.9	vs29	99.1	91.6	98.6	+7.0
vs09	99.4	86.8	98.3	+11.5	vs34	99.8	99.4	99.1	-0.3
vs16	99.7	89.0	97.3	+8.3	vs36	99.7	99.3	99.3	-0.1
vs19	99.6	84.8	98.0	+13.2	vs37	98.9	94.9	97.5	+2.6
vs24	99.4	98.7	98.7	0.0	vs57	99.7	90.5	97.2	+6.7
vs25	99.4	80.6	97.2	+16.6	vs79	99.1	85.3	96.9	+11.5

Table 1: Accuracy scores obtained from persistence diagrams before and after performing our optimization over the filtration. Note that the difference between the scores is almost always positive, i.e., there is almost always improvement after our optimization process. Scores do not have standard deviations since we use the train/test splits of the `mnist.load_data` function in TensorFlow 2.

Filter selection. In this experiment, we address a very common issue in Topological Data Analysis, filter selection. Indeed, when computing persistence diagrams in order to generate topological features from a data set for further data analysis, the filter function that is being used to filter the data set always has to be specified a priori. Here, we provide a very simple heuristic to tune it if it comes from a parametrized family \mathcal{F} of filters and if the learning task is supervised, which is the case in, e.g., classification. We simply start from a random guess in \mathcal{F} and then optimize the following criterion, which is inspired from [ZW19]:

$$\mathcal{L}(f) = \sum_{l=1}^N \frac{\sum_{i,j:y_i=y_j=l} W_p(D_i(f), D_j(f))}{\sum_{i,j:y_i=l} W_p(D_i(f), D_j(f))}, \quad (3)$$

which amounts to minimize the distances between persistence diagrams that share the same label, and increase the distances between persistence diagrams with different labels. Note that the batch size that we use in this optimization process has a big influence on the computation time, since the larger the batch size, the more Wasserstein distances we will have to compute in our cost. To cope with this issue, we actually used the Sliced Wasserstein distance SW [CCO17] instead of W_p , which, since it is computed with projections onto lines, can be defined entirely with matrix operations that are usually available in any library with autodifferentiation. This drastically improves on computation time, while remaining in our framework since the Sliced Wasserstein distance is also a semi-algebraic function.

We classify images from the MNIST data set. We assign values to the pixels using a height function given by a direction (parametrized by an angle in the Euclidean plane), and we use 0-dimensional persistence diagrams computed after optimizing this direction using loss (3). See Figure 3.

We then compute the accuracy scores obtained with a random forest classifier for the (binary) classification tasks digit x vs. digit y for all pairs $0 \leq x, y \leq 9$, using the first five persistence landscapes with resolution

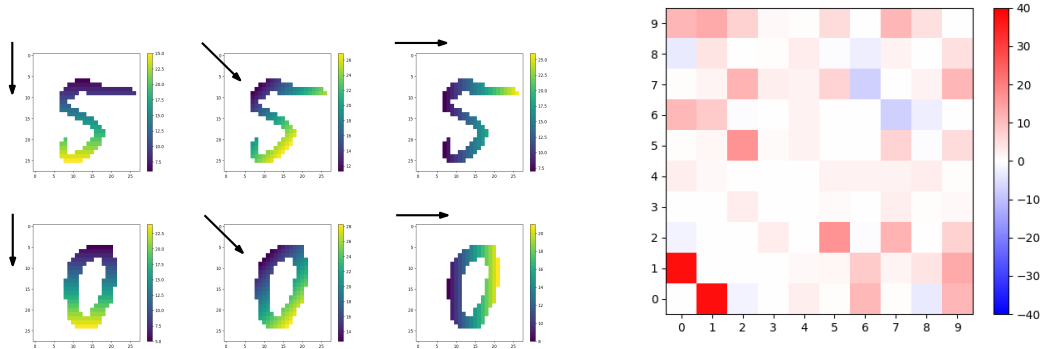


Figure 3: **Left:** Example of images and directions inducing different height functions. Different directions generate different height functions and filtrations and thus different persistence diagrams. In this experiment, we optimize over the direction so that the persistence diagrams are the most efficient for image classification. **Right:** Pairwise differences between the accuracy scores (%) of optimized and non optimized filtrations for classifying digit x vs. digit y . One can see that the difference is almost always positive.

100 associated to the persistence diagrams before and after optimization. Even though our primary goal is to demonstrate that optimizing the filter almost always lead to an improvement, we also add a baseline score obtained by training a random forest classifier directly on the images for proper comparison. Some of the scores are displayed in Table 1 (the full table can be found in Appendix), and all pairwise scores are shown in Figure 3. Interestingly, when starting with a random direction, scores can be much worse than the baseline, but our optimization process is then able to select the best direction that induces the best persistence diagrams (with respect to the classification task) without prior knowledge on the data set.

6 Conclusion

In this article we introduced a theoretical framework that encompasses most of the previous methods for optimizing topology-based functions. In particular, we obtained convergence results for very general classes of functions with topological flavor computed with persistence theory, and provided corresponding code that one can use to reproduce previously introduced topological optimization tasks. For future work, we are planning to further investigate tasks related to classifier regularization in ML [CNBW19], and to improve on computation time using, e.g., vineyards [CSEM06].

References

- [ACG⁺20] Hirokazu Anai, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hiroya Inakoshi, Raphaël Tinarrage, and Yuhei Umeda. DTM-based Filtrations. In *Topological Data Analysis*, pages 33–66. Springer, 2020.
- [AEK⁺17] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *The Journal of Machine Learning Research*, 18(1):218–252, 2017.
- [AGH⁺09] Dominique Attali, Marc Glisse, Samuel Hornus, Francis Lazarus, and Dmitriy Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. *TOPOINVIS*, 9:23–24, 2009.
- [BCY18] Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and topological inference*, volume 57. Cambridge University Press, 2018.
- [BE17] Ulrich Bauer and Herbert Edelsbrunner. The morse theory of čech and delaunay complexes. *Transactions of the American Mathematical Society*, 369(5):3741–3762, 2017.
- [BGGSSG20] Rickard Brüel-Gabrielsson, Vignesh Ganapathi-Subramanian, Primoz Skraba, and Leonidas J Guibas. Topology-aware surface reconstruction for point clouds. In *Computer Graphics Forum*, volume 39, pages 197–207. Wiley Online Library, 2020.
- [BGND⁺19] Rickard Brüel-Gabrielsson, Bradley J Nelson, Anjan Dwaraknath, Primoz Skraba, Leonidas J Guibas, and Gunnar Carlsson. A topology layer for machine learning. *arXiv preprint arXiv:1905.12200*, 2019.
- [BR91] Riccardo Benedetti and Jean-Jacques Risler. *Real algebraic and semialgebraic sets*. Hermann, 1991.
- [Bub15] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015.
- [CBO⁺20] James Clough, Nicholas Byrne, Ilkay Oksuz, Veronika A Zimmer, Julia A Schnabel, and Andrew King. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [CCI⁺20] Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. Perslay: A neural network layer for persistence diagrams and new graph topological signatures. In *International Conference on Artificial Intelligence and Statistics*, pages 2786–2796. PMLR, 2020.
- [CCO17] Mathieu Carrière, Marco Cuturi, and Steve Oudot. Sliced Wasserstein kernel for persistence diagrams. In *34th International Conference on Machine Learning (ICML 2017)*, volume 70, pages 664–673. JMLR.org, 2017.
- [CD18] Frédéric Chazal and Vincent Divol. The density of expected persistence diagrams and its kernel based estimation. In *34th International Symposium on Computational Geometry (SoCG 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [CdSGO16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. SpringerBriefs in Mathematics. Springer, 2016.
- [CG20] Gunnar Carlsson and Rickard Brüel Gabrielsson. Topological approaches to deep learning. In *Topological Data Analysis*, pages 119–146. Springer, 2020.
- [CNBW19] Chao Chen, Xiuyan Ni, Qinxun Bai, and Yusu Wang. A topological regularizer for classifiers via persistent homology. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2573–2582, 2019.
- [Cos00] Michel Coste. *An introduction to o-minimal geometry*. Istituti editoriali e poligrafici internazionali Pisa, 2000.
- [CSEH09] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.
- [CSEHM10] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have L_p -stable persistence. *Foundations of computational mathematics*, 10(2):127–139, 2010.
- [CSEM06] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In Nina Amenta and Otfried Cheong, editors, *22nd Annual Symposium on Computational Geometry (SoCG 2006)*, pages 119–126. Association for Computing Machinery, 2006.
- [DDKL20] Damek Davis, Dmitriy Drusvyatskiy, Sham M. Kakade, and Jason D. Lee. Stochastic subgradient method converges on tame functions. *Found. Comput. Math.*, 20(1):119–154, 2020.

- [DL19] Vincent Divol and Théo Lacombe. Understanding the topology and the geometry of the persistence diagram space via optimal partial transport. *arXiv preprint arXiv:1901.03048*, 2019.
- [DUC20] Meryll Dindin, Yuhei Umeda, and Frédéric Chazal. Topological data analysis for arrhythmia detection through modular neural networks. In *Canadian Conference on Artificial Intelligence*, pages 177–188. Springer, 2020.
- [Ede93] Herbert Edelsbrunner. The union of balls and its dual shape. In *Proceedings of the ninth annual symposium on Computational geometry*, pages 218–231, 1993.
- [EH10] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.
- [GC19] Rickard Brüel Gabrielsson and Gunnar Carlsson. Exposition and interpretation of the topology of neural networks. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 1069–1076. IEEE, 2019.
- [GHO16] Marcio Gameiro, Yasuaki Hiraoka, and Ipei Obayashi. Continuation of point clouds via persistence diagrams. *Physica D: Nonlinear Phenomena*, 334:118–132, 2016.
- [GNDS20] Rickard Brüel Gabrielsson, Bradley J Nelson, Anjan Dwaraknath, and Primoz Skraba. A topology layer for machine learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1553–1563, 2020.
- [HKND19] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Mandar Dixit. Connectivity-optimized representation learning via persistent homology. In *International Conference on Machine Learning*, pages 2751–2760. PMLR, 2019.
- [HKNU17] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl. Deep learning with topological signatures. In *Advances in Neural Information Processing Systems*, pages 1634–1644, 2017.
- [KKZ⁺20] K. Kim, J. Kim, M. Zaheer, J. Kim, F. Chazal, and L. Wasserman. Efficient topological layer based on persistent landscapes. In *NeurIPS 2020 (to appear)*, 2020.
- [LOT19] Jacob Leygonie, Steve Oudot, and Ulrike Tillmann. A framework for differential calculus on persistence barcodes. *arXiv preprint arXiv:1910.00960*, 2019.
- [MHRB19] Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. *arXiv preprint arXiv:1906.00722*, 2019.
- [PSO18] Adrien Poulenard, Primoz Skraba, and Maks Ovsjanikov. Topological function optimization for continuous shape matching. In *Computer Graphics Forum*, volume 37, pages 13–25. Wiley Online Library, 2018.
- [RTB⁺18] Bastian Rieck, Matteo Togninalli, Christian Bock, Michael Moor, Max Horn, Thomas Gumbsch, and Karsten Borgwardt. Neural persistence: A complexity measure for deep neural networks using algebraic topology. *arXiv preprint arXiv:1812.09764*, 2018.
- [SWB20] Elchanan Solomon, Alexander Wagner, and Paul Bendich. A fast and robust method for global topological functional optimization. *arXiv preprint arXiv:2009.08496*, 2020.
- [Ume17] Yuhei Umeda. Time series classification via topological data analysis. *Information and Media Technologies*, 12:228–239, 2017.
- [Wil96] Alex J Wilkie. Model completeness results for expansions of the ordered field of real numbers by restricted pfaiffian functions and the exponential function. *Journal of the American Mathematical Society*, 9(4):1051–1094, 1996.
- [WLSC20] Fan Wang, Huidong Liu, Dimitris Samaras, and Chao Chen. Topogan: A topology-aware generative adversarial network. In *European Conference on Computer Vision(ECCV)*, 2020.
- [ZW19] Qi Zhao and Yusu Wang. Learning metrics for persistence-based summaries and applications for graph classification. In Hanna Wallach, Hugo Larochelle, Alina Beygelzimer, Florence D’Alché-Buc, Emily Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pages 9855–9866. Curran Associates, Inc., 2019.

Proof of Proposition 3.4

Proposition 3.4. *Given a simplicial complex K , the map $\text{Pers}: \text{Filt}_K \subseteq \mathbb{R}^{|K|} \rightarrow \mathbb{R}^{|K|}$ is semi-algebraic, and thus definable in any o-minimal structure. Moreover, there exists a semi-algebraic partition of Filt_K such that the restriction of Pers to each element of this partition is a Lipschitz map.*

Proof. As K is finite, the number of preorders on the simplices of K is finite. Let \preceq be a preorder on simplices of K induced by some equalities and inequalities between the coordinates of $\mathbb{R}^{|K|}$. Then, the set of filtrations $F \in \text{Filt}_K$ such that F gives rise to a preorder equal to \preceq is a semi-algebraic set. Thus, Filt_K is a semi-algebraic set that admit a semi-algebraic partition such that the restriction of the persistence map Pers to each component is a constant permutation. As a consequence, on each open element of this partition, the partial derivatives of Pers are equal constant equal to 0 or 1.

Now, from the stability theorem for persistence [CdSGO16], the persistence modules induced by two filtrations $F_1, F_2 \in \text{Filt}_K$ are ϵ -interleaved where ϵ is the sup norm of the vector $F_2 - F_1$. As a consequence, the restriction of Pers to each component of the above semi-algebraic partition of Filt_K is 1-Lipschitz with respect to the sup norm in $\mathbb{R}^{|K|}$. \square

Other examples of definable families of filtrations

The Čech and alpha-complex filtrations The Čech complex filtration built on top of sets of n points $x_1, \dots, x_n \in \mathbb{R}^d$ is the semi-algebraic parametrized family of filtrations

$$\Phi: A = (\mathbb{R}^d)^n \rightarrow \mathbb{R}^{|\Delta_n|} = \mathbb{R}^{2^n - 1},$$

where Δ_n is the simplicial complex made of all the faces of the $(n - 1)$ -dimensional simplex, defined, for any $(x_1, \dots, x_n) \in A$ and any simplex $\sigma \subseteq \{1, \dots, n\}$, by

$$\Phi_\sigma(x) = \min \left\{ r \geq 0 : \bigcap_{j \in \sigma} B(x_j, r) \neq \emptyset \right\}$$

where $B(x_j, r) = \{x \in \mathbb{R}^d : \|x - x_j\| \leq r\}$ is the ball of center x_j and radius r . The filtrations naturally satisfy the Lipschitz property.

Note that the Čech complex filtration is closely related to the so-called alpha-complex filtration which is a filtration of the Delaunay triangulation of the set of points x_1, \dots, x_n —see, e.g., Chapter 6 in [BCY18] for a definition. The simplicial complex on which the alpha-complex filtration is built depends on the points x_1, \dots, x_n . However, if A is a connected component of the open semi-algebraic subset of $(\mathbb{R}^d)^n$ of the points x_1, \dots, x_n that are in general position⁶, then all the points in A have the same Delaunay triangulation. In that case the alpha-complex defines a semi-algebraic parametrized family of filtrations.

Moreover, the persistence diagram of the alpha-complex filtration built on top of $x_1, \dots, x_n \in \mathbb{R}^d$ is the same as the persistence diagram of the Čech complex filtration built on the same set of points [Ede93, BE17] if we ignore points on the diagonal.

Cubical complexes While this presentation was restricted to simplicial complexes for simplicity, the same properties apply for more general complexes. Indeed, as long as the boundary maps are well-defined and associate to each cell of dimension d a chain of dimension $d - 1$, the same persistence algorithm can be run, and its output is still a permutation of its input. The most common non-simplicial complexes are the so-called cubical complexes, where cells are cubes. They are particularly well suited to represent images, or (discretized) functions on \mathbb{R}^d .

Other examples of locally Lipschitz functions of persistence

Persistence images [AEK⁺17] Given a weight function $w: \mathbb{R}^2 \rightarrow \mathbb{R}^+$ and a real number $\sigma > 0$, the persistence image (also called the persistence surface) associated to a persistence diagram D is the function $\rho_D: \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by

$$\rho_D(q) = \sum_{p \in D} w(p) \exp\left(-\frac{\|q - p\|^2}{2\sigma^2}\right).$$

The definition of ρ_D only involves algebraic operations, the weight function w , and exponential functions. Thus, it follows from [Wil96] that if w is a semi-algebraic function, K is a simplicial complex, and $\{q_1, \dots, q_n\}$ is a finite set of points in \mathbb{R}^2 , then the map that associates the vector $(\rho_D(q_1), \dots, \rho_D(q_n))$ to each persistence diagram D of a filtration of K is definable in some o-minimal structure. In [AEK⁺17], the authors provide explicit conditions under which the map from persistence diagrams to persistence surfaces is Lipschitz.

⁶The points $x_1, \dots, x_n \in \mathbb{R}^d$ are in general position if any subset of size at most $d + 1$ is a set of affinely independent points, and if no subset of $d + 2$ points lies on the same $(d - 1)$ -dimensional sphere

A remark about the locally Lipschitz condition in Theorem 4.2

The convergence Theorem 4.2 provides explicit conditions ensuring the convergence of stochastic gradient descent for functions of persistence. The main criterion to be checked is the local Lipschitz condition for \mathcal{L} . From the remark following Definition 4.1, it is sufficient to check that Φ and E are Lipschitz. Regarding Φ , it is obvious for Vietoris-Rips, Čech, DTM, etc. filtrations on finite point clouds, but wrong for the alpha-complex filtration, which is not a locally Lipschitz function of the coordinates of the points. Indeed, one can take 3 points within a bounded region of the plane that are almost aligned and whose circumradius is arbitrarily large, and this circumradius is the filtration value of the triangle. However, by comparing with the Čech complex, we know that such large values always correspond to diagonal points of the persistence diagram. In our example of 3 points, the longest edge has the same filtration as the triangle and they are paired by the persistence algorithm. To handle alpha-complex filtrations, we need to restrict to functions of persistence that are defined for various numbers of points, ignore points on the diagonal (the image of a diagram is the same if we add or remove points on the diagonal), and are still locally Lipschitz. This is the case for all the functions of persistence presented in this paper. Composed with such a function, the difference between Čech and alpha-complex filtrations disappears, it becomes an implementation detail and all the differentiation and optimization properties proved for the first apply to the second.

More numerical experiments

3D shape processing. In [PSO18], persistence optimization is used for modifying functions defined on 3D shapes. More specifically, given a 3D shape S , one is interested in optimizing a function $F: V(S) \rightarrow \mathbb{R}$ defined on the vertices $V(S)$ of S , so that the Wasserstein distance between the persistence diagram associated to F and D^* is minimized, where D^* is a target persistence diagram (which either comes from another function $G: S \rightarrow \mathbb{R}$, or is defined by the user). In this experiment, we minimize the loss $\mathcal{L}(F) = T(F)$, where $T(F) := W_2(D, D^*)^2$, that is, the Wasserstein distance between the 0-dimensional persistence diagram D associated to the sublevel set of a function F —initialized with the height function, see Figure 4 (up)—of a human 3D shape, and a target persistence diagram D^* which only contains a single point, with the same coordinates than the point (in the persistence diagram of the height function of the shape S) corresponding to the right leg. This makes the function values to force the two points in D corresponding to the hands of S to go to the diagonal, by creating paths between the hands and the hips (middle). It is worth noting that these path creations come from the fact that we only used a topological penalty in the loss: in [PSO18], the authors ensure smoothness of the resulting function by forcing it to be a linear combination of a first few eigenfunctions of the Laplace-Beltrami operator on the 3D shape. We also display the sequence of optimized persistence diagrams in Figure 4 (lower row), from which it can be observed that the optimization is piecewise-linear, which reflects the fact that the persistence map has an associated semi-algebraic partition, as per Proposition 3.2.

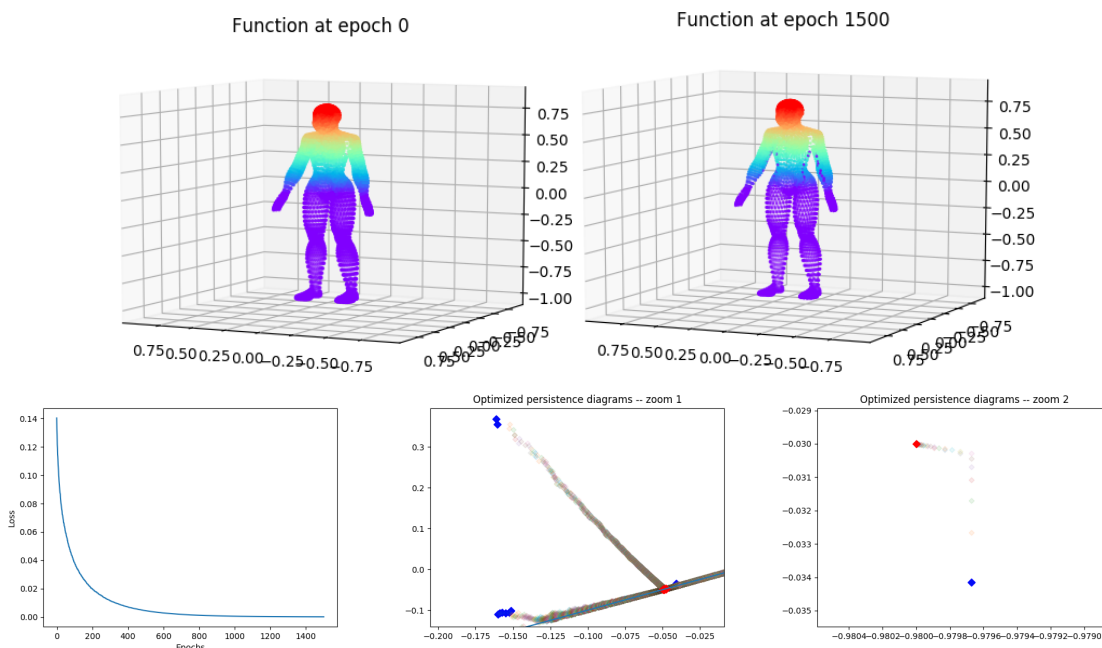


Figure 4: 3D shape before (upper left) and after (upper right) optimization, and corresponding loss function (lower left). Note how paths of low function values were created between the hips and the hands. We also show the sequence of persistence diagrams (lower middle and right) with blue points being the initial persistence diagram, and red ones being the fully optimized persistence diagram. Note how optimization trajectories look piecewise-linear.

Image processing. Another task presented in [GNDS20] is related to image processing. In this experiment, we optimize the 0-dimensional homology associated to the pixel values of a digit binary image I with noise (see Figure 5, upper left). Since noise can be detected as unwanted small connected components, we use the loss $\mathcal{L}(I) = P(I) + T(I)$, where $T(I) := \sum_{i=1}^p |d_i - b_i|$ is the total persistence penalty, $D_{\text{reg}} = \{(b_1, d_1), \dots, (b_p, d_p)\}$ is the finite 0-dimensional persistence diagram of the cubical complex associated to I ,

and $P(I) := \sum_{p \in I} \min\{|p|, |1 - p|\}$ is a penalty forcing the pixel values to be either 0 or 1. As can be seen from Figure 5, using the two penalties is essential: if only $P(I)$ is used, the noise is amplified (lower left), and if only $T(I)$ is used, the noise does not quite disappear, and paths are created between the noise and the central digit to ensure the corresponding connected components are merged right after they appear (lower middle). Note that this funny behavior that appears when penalizing topology alone is similar to what was observed in experiments where persistence was used to simplify functions [AGH⁺09]. Using both penalties completely remove the noise (lower right) in two steps: firstly topology is optimized, and then the paths created by optimizing $T(I)$ are removed by optimizing $P(I)$. This two step process can also be observed on the loss function (upper middle) and the sequence of optimized persistence diagrams of the image (upper right), where a bifurcation point can be observed in the optimization process.

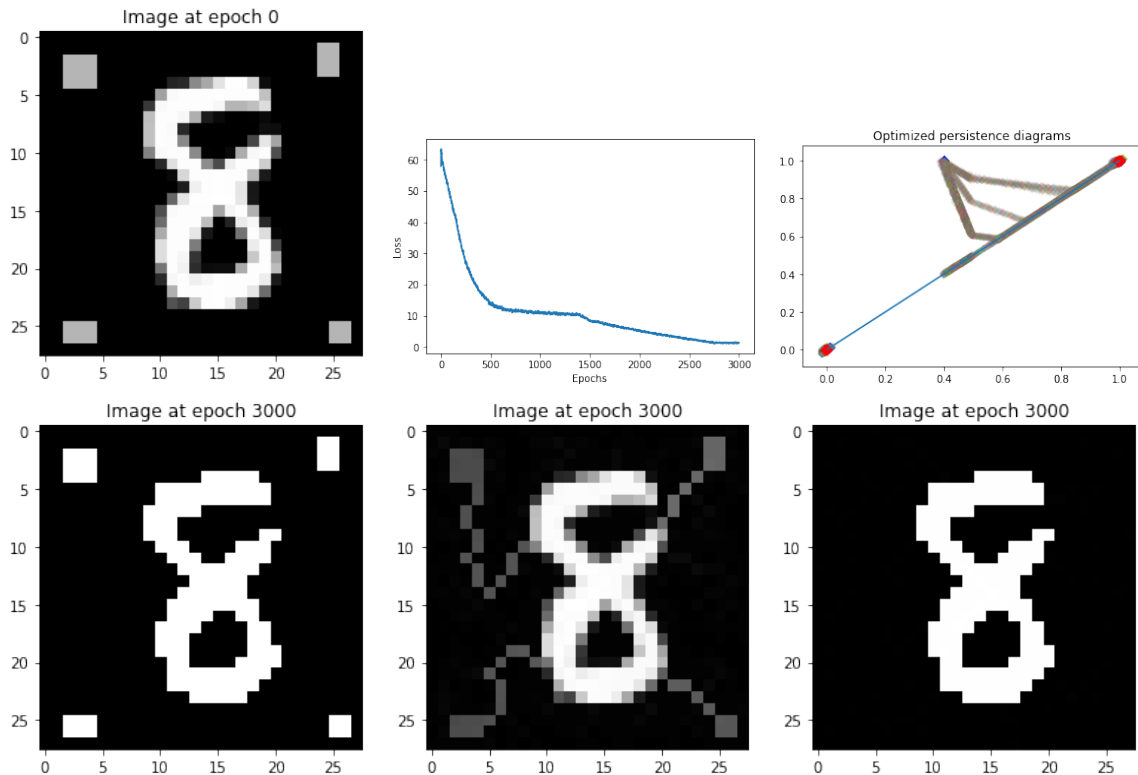


Figure 5: Image before (upper left) and after optimization for various penalties. When both $T(I)$ and $P(I)$ are used (lower right), we also show the corresponding loss function (upper middle) and the sequence of persistence diagrams (upper right) with blue points being the initial persistence diagram, and red ones being the fully optimized persistence diagram.

Linear regression. This experiment comes from [GNDS20], in which the authors use persistence optimization in a linear regression setting. Given some dataset $X \in \mathbb{R}^{n \times p}$ and ground-truth associated values $Y \in \mathbb{R}^n$ computed as $Y = X \cdot \beta^* + \epsilon$, where $\beta^* \in \mathbb{R}^p$ is the vector of ground-truth coefficients and ϵ is some high-dimensional Gaussian noise, one can leverage some prior knowledge on the shape of β to penalize the coefficients with bad topology. In particular, when using β^* with three peaks, as in Figure 6 (left), we use the loss $\mathcal{L}(\beta) = P(\beta) + TV(\beta) + T(\beta)$, where $T(\beta) := \sum_{i=1}^p |d_i - b_i|$, $\hat{D} = \{(b_1, d_1), \dots, (b_p, d_p)\}$ is the 0-dimensional persistence diagram of the sublevel sets of β , minus the three most prominent points, $TV(\beta) = \sum_i |\beta_{i+1} - \beta_i|$ is the usual total variation penalty (which can also be interpreted as a topological penalty as it corresponds to the total persistence of the so-called *extended persistence* [CSEH09] of the signal), and $P(\beta) := \sum_i |x_i \cdot \beta - y_i|^2$ is the usual mean-square error (MSE). We optimized β with the MSE alone, then MSE plus total variation, then MSE plus total variation plus topological penalty, and we generated new MSE values from new randomly generated test data, see Figure 6 (right). It is interesting to note that using all penalties lead to the best result: using MSE alone leads to overfitting, and adding total variation smooths the coefficients a lot since β is initialized with random values. Using all penalties ends up being a right combination of minimizing the error, smoothing the signal, and getting to the right shape of β .

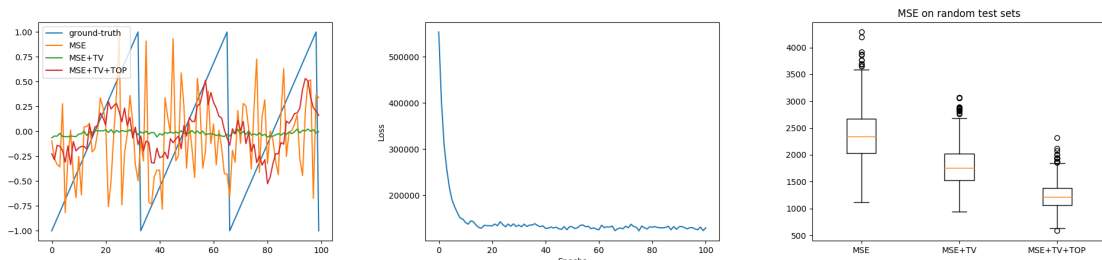


Figure 6: Regression coefficients after optimization for various penalties (left), and corresponding loss function when all penalties are used (middle). Generalization performance is increased by using all penalties, since the MSE on various randomly generated test sets is largely decreased (right).

Noisy point cloud. We perform another point cloud optimization (like in the article), but now we start with a noisy sample X of the circle with three outliers and use the loss $\mathcal{L}(X) = W_2(D, D^*)^2$, where W_2 stands for the Wasserstein distance, D is the 0-dimensional persistence diagram associated to the Vietoris-Rips filtration of X , and D^* is the 0-dimensional persistence diagram associated to the Vietoris-Rips filtration of a clean (i.e., with no noise nor outliers) sample of the circle. See Figure 7. Note that when one does not use extra penalties, optimizing only topological penalties can lead to funny effects: as one can see on the middle of Figure 7, the circle got disconnected, and one of the outliers created a small path out of the circle during optimization.

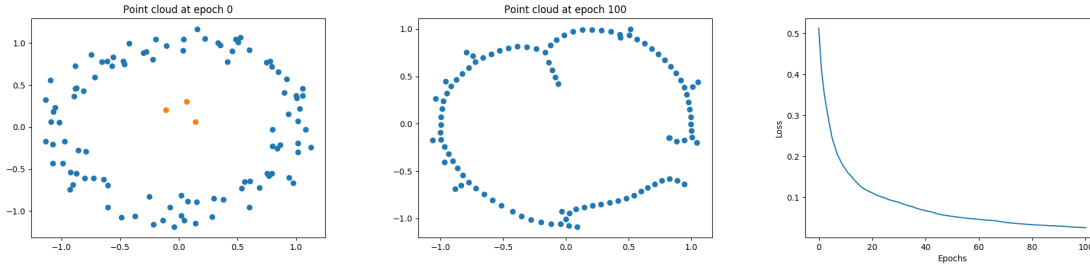


Figure 7: Noisy circle initialization with outliers, before (left) and after (middle) optimization, and corresponding loss function (right).

Filter optimization. In addition to classifying digits (as presented in the article), we also run filter selection on a second family of data sets, which is comprised of graphs. For each graph G , we compute its normalized Laplacian $L(G)$ and we define a family of filtrations by consider all filter functions that can be written as linear combinations of the eigenvectors of $L(G)$. Then we use persistence landscapes and random forest classifiers (as in the article) to generate accuracy scores. We also train a baseline random forest classifier directly on the eigenvalues of the graph Laplacians. The results are displayed (as well as the full results for the MNIST experiment) in Table 2.

It should be noted now that even though accuracy scores are sometimes significantly lower than the baseline, optimizing the filter prior to classification often improves the results, and if the score does not improve (which can be due to the fact that the batch size, taken as 150 so that computation runs in a reasonable amount of time, is too small to properly ensure convergence), it does not decrease by a significant margin either.

Details on experiments

All experiments were run on Intel dual-Xeon SP with 10 cores and 9.6GB/core (RAM). The Classification results for the experiments on graphs were generated with 10-fold cross-validation: each data set was divided into ten 90-10 train-test splits, and results were averaged over these splits. The Classification results for the experiments on images were generated using train-test splits provided in `Tensorflow`. We also added a multi-class classification task, called `all`, which consists in jointly classifying all images (and not only digit x vs. digit y). All the optimization processes (including those presented in the article) were done with stochastic gradient descent as implemented in `Tensorflow` 2.4.1. More specifically, we used the `SGD` optimizer class with `InverseTimeDecay` learning rate (in order to satisfy Assumptions 1, 2 and 3 of Section 4.1 in the article). We also used `Adam` optimizer with `ExponentialDecay` learning rate for some experiments since we noticed empirically that, even though Assumption 1 was not satisfied, the results were not very different, and convergence was slightly smoother. Parameter initialization was done randomly, and the batch sizes + numbers of epochs were taken sufficiently large so that convergence was reached in each illustration—see <https://github.com/MathieuCarriere/difftda.git> for exact values. For instance, filter selection was done with initial learning rate equal to 0.001 and batch size equal to 150.

Dataset	Baseline	Before	After	Difference
a11	96.9	65.1	80.6	+15.5
vs01	100.0	61.3	99.0	+37.6
vs02	99.4	98.8	97.2	-1.6
vs03	99.8	99.1	99.2	+0.1
vs04	99.9	96.0	98.8	+2.8
vs05	99.6	95.7	96.3	+0.6
vs06	99.4	87.3	98.2	+10.9
vs07	99.8	97.4	98.0	+0.6
vs08	99.4	90.4	87.0	-3.4
vs09	99.4	86.8	98.3	+11.5
vs12	99.6	98.3	98.5	+0.2
vs13	99.7	98.9	99.1	+0.2
vs14	100.0	97.1	98.3	+1.2
vs15	99.8	96.7	98.0	+1.3
vs16	99.7	89.0	97.3	+8.3
vs17	99.7	96.8	98.6	+1.8
vs18	99.8	91.7	96.0	+4.3
vs19	99.6	84.8	98.0	+13.2
vs23	99.1	95.2	98.0	+2.9
vs24	99.4	98.7	98.7	0.0
vs25	99.4	80.6	97.2	+16.6
vs26	99.7	98.8	98.2	-0.6
vs27	98.6	80.1	91.9	+11.8
vs28	99.1	96.8	96.8	0.0
vs29	99.1	91.6	98.6	+7.0
vs34	99.8	99.4	99.1	-0.3
vs35	99.2	93.5	94.3	+0.8
vs36	99.7	99.3	99.3	-0.1
vs37	98.9	94.9	97.5	+2.6
vs38	99.0	98.3	98.8	+0.6
vs39	98.8	96.8	97.8	+1.0
vs45	99.9	96.5	98.4	+1.9
vs46	99.6	94.1	96.0	+1.9
vs47	99.7	97.2	99.3	+2.1
vs48	99.2	90.4	93.4	+3.0
vs49	98.4	93.7	94.2	+0.5
vs56	99.0	96.9	97.1	+0.2
vs57	99.7	90.5	97.2	+6.7
vs58	98.9	92.7	92.3	-0.4
vs59	99.4	90.0	95.4	+5.5
vs67	99.7	98.4	91.0	-7.4
vs68	98.7	92.2	89.5	-2.7
vs69	99.7	87.0	86.7	-0.3
vs78	98.9	95.7	97.6	+1.9
vs79	99.1	85.3	96.9	+11.5
vs89	98.7	84.2	89.1	+4.9
PROTEINS	73.6 ± 3.23	68.7 ± 2.38	69.8 ± 3.42	+1.1
MUTAG	85.1 ± 7.06	76.1 ± 6.31	81.3 ± 6.25	+5.2
COX2	78.6 ± 1.73	77.5 ± 2.29	76.9 ± 3.04	-0.6
DHFR	78.8 ± 4.12	61.6 ± 4.85	60.3 ± 5.26	-1.3
BZR	84.9 ± 2.08	78.8 ± 3.73	77.5 ± 3.93	-1.2
FRANKENSTEIN	69.7 ± 1.40	63.3 ± 2.26	63.1 ± 2.20	-0.3
IMDB-MULTI	49.3 ± 3.26	40.6 ± 2.74	39.7 ± 3.21	-0.9
IMDB-BINARY	72.8 ± 4.45	60.1 ± 3.99	60.0 ± 4.36	-0.1
NCI1	74.3 ± 1.81	60.3 ± 1.27	61.4 ± 1.78	+1.0
NCI109	72.5 ± 1.69	60.2 ± 2.33	61.2 ± 2.87	+1.0

Table 2: All accuracy scores for graphs and MNIST data sets.