
Non-isotropic Persistent Homology

Vincent P. Grande¹ Michael T. Schaub¹

Abstract

Persistent Homology is a widely used topological data analysis tool that creates a concise description of the topological properties of a point cloud based on a specified filtration. Most filtrations used for persistent homology depend (implicitly) on a chosen metric, which is typically agnostically chosen as the standard euclidean metric on \mathbb{R}^n . Recent work has tried to uncover the “true” metric on the point cloud using distance-to-measure functions, in order to obtain more meaningful persistent homology results. Here we propose an alternative look at this problem: we posit that information on the point cloud is lost when restricting persistent homology to a single (correct) distance function. Instead we show how by varying the distance function on the underlying space and analysing the corresponding shifts in the persistence diagrams, we can extract additional topological and geometrical information. Finally, we numerically show that non-isotropic persistent homology can extract information on orientation, orientational variance, and scaling of randomly generated point clouds with good accuracy.

1. Introduction

Over the last decades, topological data analysis (TDA) has proven to provide a valuable toolkit for extracting information out of complex data sets. Most notably, persistent homology (PH) provides a straight-forward way to extract topological information across different scales from a point cloud. The resulting persistence diagrams and persistence barcodes form a metric spaces and have been used for many interesting applications, see for example (Edelsbrunner et al., 2008).

In order to form the simplicial filtration used for comput-

¹Department of Computer Science, RWTH Aachen University, Aachen, Germany. Correspondence to: Vincent P. Grande <grande@cs.rwth-aachen.de>, Michael T. Schaub <schaub@cs.rwth-aachen.de>.

Presented at the 2nd Annual Workshop on Topology, Algebra, and Geometry in Machine Learning (TAG-ML) at the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA. 2023. Copyright 2023 by the author(s).

ing PH, we need to specify a distance function on the point cloud, which is typically chosen as a metric on \mathbb{R}^n ; most often the standard ℓ_2 metric on \mathbb{R}^n is used. As in algebraic topology virtually all reasonable metrics on \mathbb{R}^n are equivalent and induce the same topology, it would seem that the choice of metric for PH is irrelevant. However, this is not true: Because the metric controls the birth-time of simplices, changing the metric alters the birth and death times of the topological features of the constructed simplicial complexes. A change in metric may even eradicate certain topological features or introduce new topological features. The role of the underlying metric for PH is often not further investigated, even though it has been acknowledged in recent work such as (Anai et al., 2020). However, this influence is treated by the authors as a problem that they attempt to fix by introducing an improved euclidean distance making PH more robust to outliers. Stated differently, the aim is to construct a single metric that leads to the “best possible” simplicial filtration that provides an accurate topological description of the underlying space (point cloud).

In this work, we take a different perspective on the influence of the choice of a metric on PH. Rather than trying to find the best possible metric, we ask: how does PH change as we vary the metric, and can we exploit the induced changes in PH to extract additional information from the data? We name this approach non-isometric persistent homology (NIPH) — see Figure 1. Note that NIPH not only provide us with a means to assess how robust (sensitive) the results of PH are to changes in the metric. In addition, non-isometric persistent homology (NIPH) harvests the rich information of how the persistence diagram of a point cloud changes when we change the underlying metric, to extract new information that is not apparent from any single PH analysis. This may be seen as analogous to cases in physics, where taking the derivative of one important physical quantity often yields another physical quantity of interest. In the next section, we will demonstrate how these rich features can be used to extract information on the preferred orientation, local orientational variance, and scaling factors of different dimensions of point clouds.

1.1. Related Work

Anai et al. (2020) introduced distance-to-measure filtrations to adapt the distance function of the point cloud to reduce

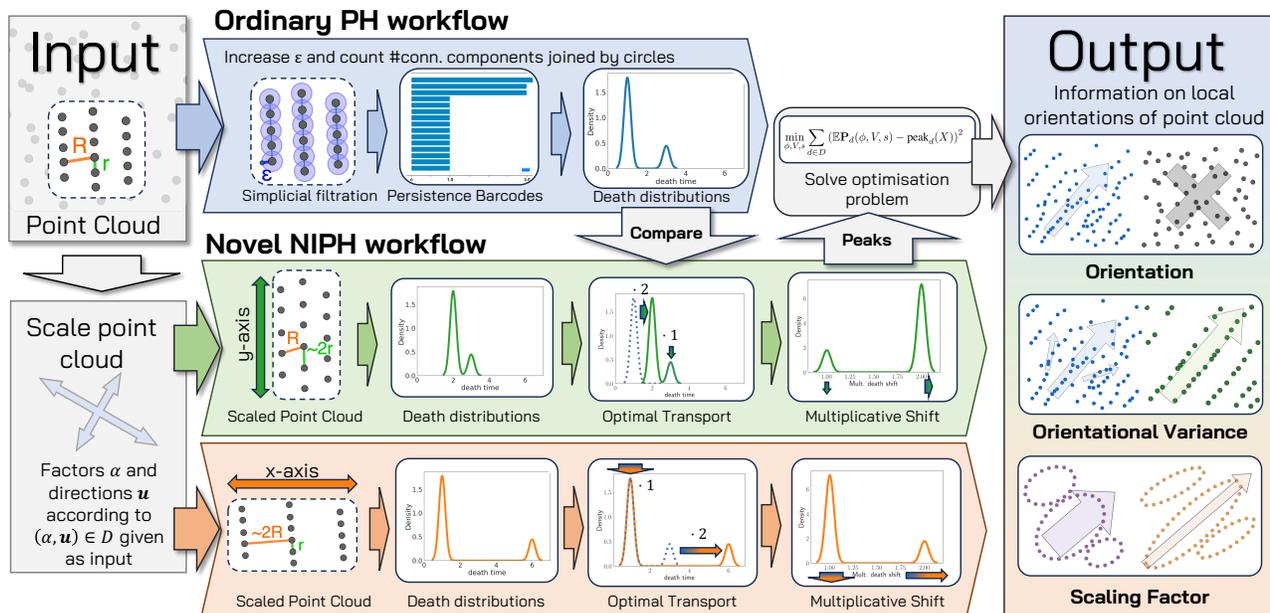


Figure 1. Schematic of Non-Isotropic Persistent Homology (NIPH). **Step 1.** Produce different versions of the input point cloud by applying directional scaling in direction \mathbf{u} and scaling-factor α according to elements in input set D . **Step 2.** Compute PH of desired degree (the diagram displays 0th persistent homology) for input point cloud and each of the scaled point clouds. **Step 3.** Compute optimal transport between weighted death time distributions of input point cloud and each of the scaled point clouds. **Step 4.** Compute multiplicative factor of shift for each death time. Extract maxima and use optimisation problem to compute preferred orientations, scaling-factor, and orientational variance of point cloud, as seen in the illustrations on the right. **Orientation.** Two point clouds with similar 0-dimensional PH. *Left:* The point cloud has additional structure: a preferred orientation. *Right:* There is no preferred orientation in the point cloud. NIPH can detect the differences in the amount of orientation in both point clouds across arbitrary scales. **Orientational variance.** Two point clouds with similarly structured 0-dimensional PH. Both have additional structure and a preferred dimension. *Left:* There is a comparably large variance in the local preferred direction. *Right:* The alignment of the point data with the preferred direction is very strong. NIPH can detect the differences in the amount of variance of orientation in both point clouds. **Scaling factor.** Two point clouds with similarly structured 1-dimensional PH. Both have additional structure and are scaled along the same preferred dimension. The ellipses of the *left* point cloud are scaled differently than the ellipses on the *right*. NIPH can detect and distinguish different scaling-factor along different dimensions. The above three properties cannot be extracted by using ordinary PH alone.



Figure 2. Data features not captured by standard PH *Left:* Persistent homology will not distinguish circles and ellipses by their death time ϵ . *Right:* Persistent homology cannot distinguish orientations and orientational variances of the data set by death times ϵ .

the effect of outliers. There has also been work on extracting geometric information using persistent homology, such as the persistent homology transform (Turner et al., 2014). However, this work focussed on extracting shapes of $2d$ surfaces and $3d$ objects in $3d$ space, whereas we focus on local geometric information encoded in point clouds. In (Hofer et al., 2019) and (Carriere et al., 2021), the authors come up with notions of differentiating persistent homology diagrams. However, their goal is to utilize this differentiability to make PH accessible to machine learning tasks, rather than to extract geometrical information on the point cloud.

2. Methods

An important motivation for our work is to combine the robust topological descriptors PH provides of (point-cloud) data with more refined geometric notions of orientation and preferred dimensions.

2.1. Introductory Example

As an introductory example, we can consider the following class of metrics on \mathbb{R}^2 :

Definition 2.1. For real $\alpha, \beta > 0$, we define the associated metric $d_{\alpha, \beta}$ on \mathbb{R}^2 as follows:

$$d_{\alpha, \beta} : \left(\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \right) \mapsto \sqrt{\alpha(x_1 - y_1)^2 + \beta(x_2 - y_2)^2},$$

which is the standard euclidean distance for $\alpha = \beta = 1$. Picking α and β amounts to (de-)prioritising the x - and y -axis for our metric. When choosing $\alpha \gg \beta$, the distance between two points is almost entirely determined by the distance of their x -values. When choosing $\beta \gg \alpha$, the

Algorithm 1 Non-isotropic persistent homology (NIPH)

Input: Point cloud X , list of directions and scale-factors D
 Compute persistent homology diagram \mathbf{P} of X
for $d \in D$ **do**
 Compute scaled data set X_d according to $d = (\mathbf{u}, \alpha)$.
 Compute persistent homology \mathbf{P}_d of X_d .
 Solve optimal transport from \mathbf{P} to \mathbf{P}_d .
 Compute mult. shift diagram \mathbf{MS}_d associated to d
end for
 Solve optimisation problem for best orientation ϕ , scaling-factor s , and orientational variance V matching the maxima of \mathbf{MS}_d .
Output: ϕ, s, V

distance between two points is almost entirely determined by the distance of their y -values, and differences in the x -coordinate have almost no influence.

To build some intuition, assume that we compute ordinary 1-dimensional PH with $\alpha = \beta = 1$ first and compare it with PH associated to $d_{\alpha, \beta}$ for $\alpha = 0.5$ and $\beta = 1$. On a data set consisting of points sampled from circles with radius $\varepsilon/2$, the death times would remain almost constant: although points are now farther apart in the x -axis, simplices with edge length of ε are enough to cover all circles in the y -direction. On the other hand, if the data set consisted of ellipses stretched along the y -axis the persistent homology would change with the change in metric. Because the change of the metric is orthogonal to the orientation of the ellipses, their death times will double (Figure 2).

This simple example motivates why looking at the change of PH under different metrics can be a powerful tool to enrich the standard tools of persistent homology with notions of orientation and preferred directions. We give a more rigorous account of our mathematical models below.

2.2. General Method

NIPH is sensitive to the presence or absence of a **preferred orientation, variance of the orientation**, and the **scaling factor** along the different preferred dimensions, see the right part of Figure 1. NIPH takes as input a point cloud $X \in \mathbb{R}^n$ and a set of scaling directions (see Algorithm 1).

The distance function For every direction of n -dimensional space given by a unit vector \mathbf{u} we can pick an orthonormal basis with first basis vector \mathbf{u} . Let U be the associated base transformation matrix. For a scaling-factor $\alpha > 0$ we can define the scaling function

$$S_{\mathbf{u}, \alpha}: x \mapsto U^{-1} \text{diag}((\alpha, 1, \dots, 1))Ux.$$

Definition 2.2. We compute the distance function $d_{\mathbf{u}, \alpha}$

associated to the pair (\mathbf{u}, α) as follows:

$$d_{\mathbf{u}, \alpha}: (x_1, x_2) \mapsto \|S_{\mathbf{u}, \alpha}(x_1) - S_{\mathbf{u}, \alpha}(x_2)\|_2.$$

We note that this definition generalises Definition 2.1 given in the introduction.

Definition 2.3. Let X be a point cloud in n -dimensional space $X \subset \mathbb{R}^n$, $\mathbf{u} \in \mathbb{R}^n$ a unit vector and $\alpha > 0$ a real scaling-factor. We denote by $X_{\mathbf{u}, \alpha}$ the scaled data-set

$$X_{\mathbf{u}, \alpha} := \{S_{\mathbf{u}, \alpha}(x) \mid x \in X\}.$$

It is easy to see that $(X_{\mathbf{u}, \alpha}, d)$ is isometric to $(X, d_{\mathbf{u}, \alpha})$ where d is the euclidean distance on \mathbb{R}^n .

Computing persistence diagrams We compute the persistence diagrams \mathbf{P} and $\mathbf{P}_{\mathbf{u}, \alpha}$ on the point clouds X and $(X, d_{\mathbf{u}, \alpha})$ for all $(\mathbf{u}, \alpha) \in D$ using a Vietoris-Rips filtration on X and $X_{\mathbf{u}, \alpha}$.

Death distributions In many cases, the death of homology classes carries more structural information than their birth time. This is because birth time is mainly controlled by point densities, whereas death times measure the distance between clusters or the size of loops. Hence we transform the persistence diagrams to density plots over the death time. In the case of 1-dimensional homology, we can weigh the points according to the difference or quotient of death and birth time, i.e., longer lived loops carry more weight. We denote by \mathbf{D} and $\mathbf{D}^{(\mathbf{u}, \alpha)}$ the vectors of death times.

Computing the multiplicative shift We are interested in how persistent homology and the death distributions change when changing the metric of the underlying data set. Hence we propose to compute an optimal transport based matching between the death distribution diagrams of X and of the $X_{\mathbf{u}, \alpha}$. We can then compute for every point in the original persistence diagram the factor by which the death time was multiplied in the optimal transport matching. We can then present these multiplicative death shift factors as another density plot (See Figure 3):

Definition 2.4 (Multiplicative Death Shift). Let $T^{(\mathbf{u}, \alpha)}$ denote the solution transport matrix of optimal transport between \mathbf{P} and $\mathbf{P}_{(\mathbf{u}, \alpha)}$. We denote by the i -th multiplicative death shift $\mathbf{ms}_i^{\mathbf{u}, \alpha}$ associated to scaling \mathbf{u}, α for $1 \leq i \leq |\mathbf{P}|$ the following expression

$$\mathbf{ms}_i^{\mathbf{u}, \alpha}(X) := \exp \left(\sum_j T_{i,j}^{\mathbf{u}, \alpha}(X) \ln \mathbf{D}_j^{\mathbf{u}, \alpha}(X) / \mathbf{D}_i(X) \right).$$

Now let w denote a weight vector for the vector of homological features of X in PH. Then we denote by $\mathbf{MS}_{\mathbf{u}, \alpha}$ the multiplicative death shift, given by the density diagram

$$\mathbf{MS}_{\mathbf{u}, \alpha}(X) := \text{density} \{(\mathbf{ms}_i^{\mathbf{u}, \alpha}(X), w_i) \mid 1 \leq i \leq |\mathbf{P}|\}.$$

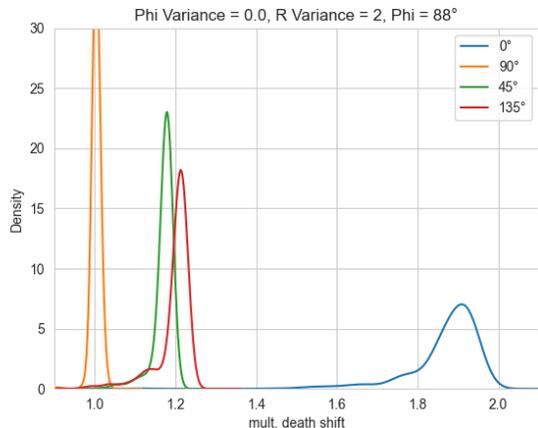


Figure 3. The multiplicative death shift diagram. The underlying point cloud consists of points sampled from ellipses with a scaling-factor of 2, where the size of the large half-axis varies between 0.2 and 2 with an orientation of 88° , and $\alpha = 2$. The orange curve represents scaling parallel to the directional scaling of the ellipses. Thus there is no change in the death times and the peak of the curve is at ~ 1.0 . The blue curve represents scaling in a direction almost orthogonal to the scaling of the data points, hence the death times are multiplied by the factor of the scaling. This is represented by the peak of the curve almost reaching 2.0. The red and green curve represent scaling roughly at a direction of 47° and 43° to the original scaling in the point cloud.

We furthermore denote the x -value of the maximum of $\text{MS}_{\mathbf{u},\alpha}(X)$ by $\text{peak}_{\mathbf{u},\alpha}(X)$.

Extracting the orientations The multiplicative death shift density diagrams obtained in the previous step are an interesting object in their own right. However, we here simply concentrate on extracting information on orientations and preferred dimensions from these shift diagrams

Definition 2.5 (Expected Peak). Let $\mathbb{E}\mathbf{P}_{\mathbf{u},\alpha}(\phi, V, s)$ be a function such that for a point cloud Y sampled from rectangles with scaling factor s , orientation ϕ and orientational variance V we have that $\text{MS}_{\mathbf{u},\alpha}(Y)$ takes its maximum value at $\mathbb{E}\mathbf{P}_{\mathbf{u},\alpha}(\phi, V, s)$.

$\mathbb{E}\mathbf{P}$ has an analytic description, but can be determined via sampling as well. NIPH solves the optimisation problem

$$\min_{\phi, V, s} \sum_{d \in D} (\mathbb{E}\mathbf{P}_d(\phi, V, s) - \text{peak}_d(X))^2$$

to obtain an estimated orientation ϕ , orientational variance V , and scaling factor α of the point cloud X where D is the set of sampling directions and scalings given as an input.

3. Experiments

We ran experiments on synthetic data to verify that NIPH can indeed infer information on the orientation, scaling,

$\text{std}(\phi)$	$\sqrt{\text{MSE}} \phi$	$\sqrt{\text{MSE}} \text{scaling}$	$\sqrt{\text{MSE}} \text{std}(\phi)$
0°	0.3°	0.09	0.26
5°	0.8°	0.11	0.18
10°	1.6°	0.15	0.12
15°	2.3°	0.18	0.08
20°	3.0°	0.19	0.07
25°	4.0°	0.19	0.07
30°	5.1°	0.19	0.08
35°	6.2°	0.20	0.08
40°	8.0°	0.20	0.08

Table 1. Quantitative performance analysis of NIPH. We have run NIPH on a point cloud sampled from 200 oriented rectangles with different orientational variances. We show the root of the mean squared error of the orientation (ϕ), orientational variance ($\text{std}(\phi)$), and scaling factor as predicted by NIPH depending on the orientational variance of the point cloud.

and variance inside a data set, see Table 1. We note that even with high orientational variance in the synthetic data, NIPH is able to predict orientation, scaling and orientational variance with good accuracy.

4. Conclusion

We have introduced NIPH, a novel method building on top of persistent homology. NIPH extracts additional topological and geometrical information by varying the distance function on the underlying space and analysing the corresponding shifts in the persistence diagrams. We have verified the performance of NIPH on a synthetic data set.

References

- Anai, H., Chazal, F., Glisse, M., Ike, Y., Inakoshi, H., Tinarage, R., and Umeda, Y. Dtm-based filtrations. In *Topological Data Analysis: The Abel Symposium 2018*, pp. 33–66. Springer, 2020.
- Carriere, M., Chazal, F., Glisse, M., Ike, Y., Kannan, H., and Umeda, Y. Optimizing persistent homology based functions. In *International conference on machine learning*, pp. 1294–1303. PMLR, 2021.
- Edelsbrunner, H., Harer, J., et al. Persistent homology—a survey. *Contemporary mathematics*, 453(26):257–282, 2008.
- Hofer, C., Kwitt, R., Niethammer, M., and Dixit, M. Connectivity-optimized representation learning via persistent homology. In *International Conference on Machine Learning*, pp. 2751–2760. PMLR, 2019.
- Turner, K., Mukherjee, S., and Boyer, D. M. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4): 310–344, 2014.