



Topology Applied to Machine Learning: From Global to Local

Henry Adams and Michael Moy*

Department of Mathematics, Colorado State University, Fort Collins, CO, United States

Through the use of examples, we explain one way in which applied topology has evolved since the birth of persistent homology in the early 2000s. The first applications of topology to data emphasized the *global* shape of a dataset, such as the three-circle model for 3×3 pixel patches from natural images, or the configuration space of the cyclo-octane molecule, which is a sphere with a Klein bottle attached via two circles of singularity. In these studies of global shape, short persistent homology bars are disregarded as sampling noise. More recently, however, persistent homology has been used to address questions about the *local* geometry of data. For instance, how can local geometry be vectorized for use in machine learning problems? Persistent homology and its vectorization methods, including persistence landscapes and persistence images, provide popular techniques for incorporating both local geometry and global topology into machine learning. Our meta-hypothesis is that the short bars are as important as the long bars for many machine learning tasks. In defense of this claim, we survey applications of persistent homology to shape recognition, agent-based modeling, materials science, archaeology, and biology. Additionally, we survey work connecting persistent homology to geometric features of spaces, including curvature and fractal dimension, and various methods that have been used to incorporate persistent homology into machine learning.

OPEN ACCESS

Edited by:

Umberto Lupo,
École Polytechnique Fédérale de
Lausanne, Switzerland

Reviewed by:

Vasileios Maroulas,
The University of Tennessee, Knoxville,
United States
Ashleigh Thomas,
Georgia Institute of Technology,
United States

*Correspondence:

Michael Moy
michael.moy@colostate.edu

Specialty section:

This article was submitted to
Machine Learning and Artificial
Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

Received: 16 February 2021

Accepted: 15 April 2021

Published: 14 May 2021

Citation:

Adams H and Moy M (2021) Topology
Applied to Machine Learning: From
Global to Local.
Front. Artif. Intell. 4:668302.
doi: 10.3389/frai.2021.668302

Keywords: persistent homology, topological data analysis, machine learning, local geometry, applied topology

1. INTRODUCTION

Applied topology is designed to measure the *shape* of data—but what is shape? Early examples in applied topology found low-dimensional structures in high-dimensional datasets, such as the three circle and Klein bottle models for grayscale natural image patches. These models are global: they parameterize the entire dataset, in the sense that most of the data points look like some point in the model, plus noise. In more recent applications, however, the shape that is being measured is not global, but instead local. Local features include texture, small-scale geometry, and the structure of noise.

Indeed, for the first decade after the invention of persistent homology, the primary story was that significant features in a dataset corresponded to long bars in the persistence barcode, whereas shorter bars generally corresponded to sampling noise. This story has evolved as applied topology has become incorporated into the machine learning pipeline. In machine learning applications, many researchers have independently found (as we survey in sections 4–6) that the short bars are often the most discriminating—the shape of the noise, or of the local geometry, is what often enables high classification accuracy. We want to emphasize that short bars do matter. Indeed,

the short bars in persistent homology are currently one of the best out-of-the-box methods for summarizing local geometry for use in machine learning. Though humans may not be able to interpret short persistent homology bars on our own (there may be too many short bars for the human eye to count), machine learning algorithms can be trained to do so. In this way, persistent homology has greatly expanded in scope during the second decade after its invention: persistent homology has important applications as a descriptor not only of global shape, but also of local geometry.

In this perspective article, we begin by outlining some of the most famous early applications of persistent homology in the global analysis of data, in which short bars were disregarded as noise. Our meta-hypothesis, however, is that short bars do matter, and furthermore, they matter crucially when combining topology with machine learning. As a partial defense for this claim, we provide a selected survey on the use of persistent homology in measuring texture, noise, local geometry, fractal dimension, and local curvature. We predict that the applications of persistent homology to machine learning will continue to advance in number, impact, and scope, as persistent homology is a mathematically motivated out-of-the-box tool that one can use to summarize not only the global topology but also the local geometry of a wide variety of datasets.

2. POINT CLOUD AND SUBLVEL SET PERSISTENT HOMOLOGY

What is a persistent homology bar? The homology of a space, roughly speaking, records how many holes that a space has in each dimension. A 0-dimensional hole is a connected component, a 1-dimensional hole is a loop, a 2-dimensional hole is a void enclosed by a surface like a sphere or a torus, etc. Homology becomes persistent when one is instead given a *filtration*, i.e., an increasing sequence of spaces. Each hole is now represented by a bar, where the start (resp. end) point of the bar corresponds to the first (resp. last) stage in the filtration where the topological feature is present (Edelsbrunner et al., 2000). Short bars correspond to features with short lifetimes, which are quickly filled-in after being created. By contrast, long bars correspond to more *persistent* features.

Perhaps the two most frequent contexts in which persistent homology is applied are point cloud persistent homology and sublevel set persistent homology. In *point cloud persistent homology*, the input is a finite set of points (a point cloud) residing in Euclidean space or some other metric space (Carlsson, 2009). For any real number $r > 0$, we consider the union of all balls of radius r centered at some point in our point cloud (see Figure 1). This union of balls provides our filtration as the radius r increases¹. A typical interpretation of the resulting persistent homology, from the global perspective, is that the long persistent homology bars recover the homology of the “true” underlying space from which the point cloud was sampled (Chazal and Oudot, 2008). A more modern but

increasingly utilized perspective is that the short persistent homology bars recover the local geometry—i.e., the texture, curvature, or fractal dimension of the point cloud data.

In *sublevel set persistent homology*, the input is instead a real-valued function $f: Y \rightarrow \mathbb{R}$ defined on a space Y (Cohen-Steiner et al., 2007). For example, Y may be a Euclidean space of some dimension. The filtration arises by considering the sublevel sets $\{y \in Y \mid f(y) \leq r\}$. As the threshold r increases, the sublevel sets grow. One can think of f as encoding an *energy*, in which case sublevel set persistent homology encodes the shape of low-energy configurations (Mirth et al., 2021). The length of a bar then measures how large of an energy barrier must be exceeded in order for a topological feature to be filled-in: a short bar corresponds to a feature that is quickly filled-in by exceeding a low energy barrier, whereas a long bar corresponds to a topological feature that persists over a longer range of energies (see Figure 2). Sublevel set persistent homology is frequently applied to grayscale image data or matrix data, where a real-valued entry of the image or matrix is interpreted as the value of the function f on a pixel.

We remark that the “union of balls” filtration for point cloud persistent homology can be viewed as a version of sublevel set persistent homology: a union of balls of radius r is the sublevel set at threshold r of the distance function to the set of points in the point cloud.

Persistent homology can be represented in two equivalent ways: either as a persistence barcode or as a persistence diagram (see Figure 3). Each interval in the persistence barcode is represented in the persistence diagram by a point in the plane, with its birth coordinate on the horizontal axis and with its death coordinate on the vertical axis². As the death of each feature is after its birth, persistence diagram points all lie above the diagonal line $y = x$. Short bars in the barcode correspond to persistence diagram points close to the diagonal, and long bars in the barcode correspond to persistence diagram points far from the diagonal.

3. EXAMPLES MEASURING GLOBAL SHAPE

The earliest applications of topology to data measured the global shape of a dataset. In these examples, the long persistent homology bars represented the true homology underlying the data, whereas the small bars were ignored as artifacts of sampling noise.

What do we mean by “global shape”? Consider, for example, conformations of the cyclo-octane molecule C_8H_{16} , which consists of a ring of eight carbon atoms, each bonded to a pair of hydrogen atoms (see Figure 4, left). The locations of the carbon atoms in a conformation approximately determine the locations of the hydrogen atoms via energy minimization, and hence each molecule conformation can be mapped to a point in $\mathbb{R}^{24} = \mathbb{R}^{8 \cdot 3}$, as the location of each carbon atom can be specified by three

¹In practice, the union of balls is stored or approximated by a *simplicial complex*, for example a Čech or Vietoris-Rips complex (Chazal et al., 2014).

²Barcodes allow for open or closed endpoints of intervals. This information can be also be encoded in a *decorated* persistence diagram (Chazal et al., 2016).

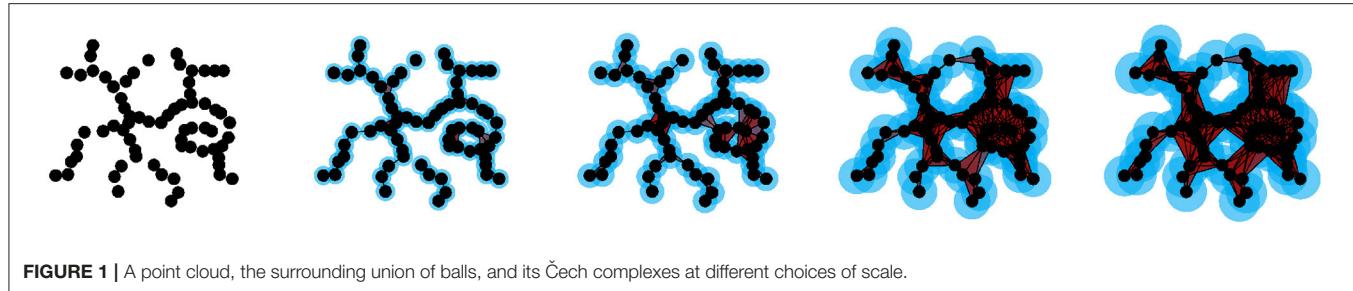


FIGURE 1 | A point cloud, the surrounding union of balls, and its Čech complexes at different choices of scale.

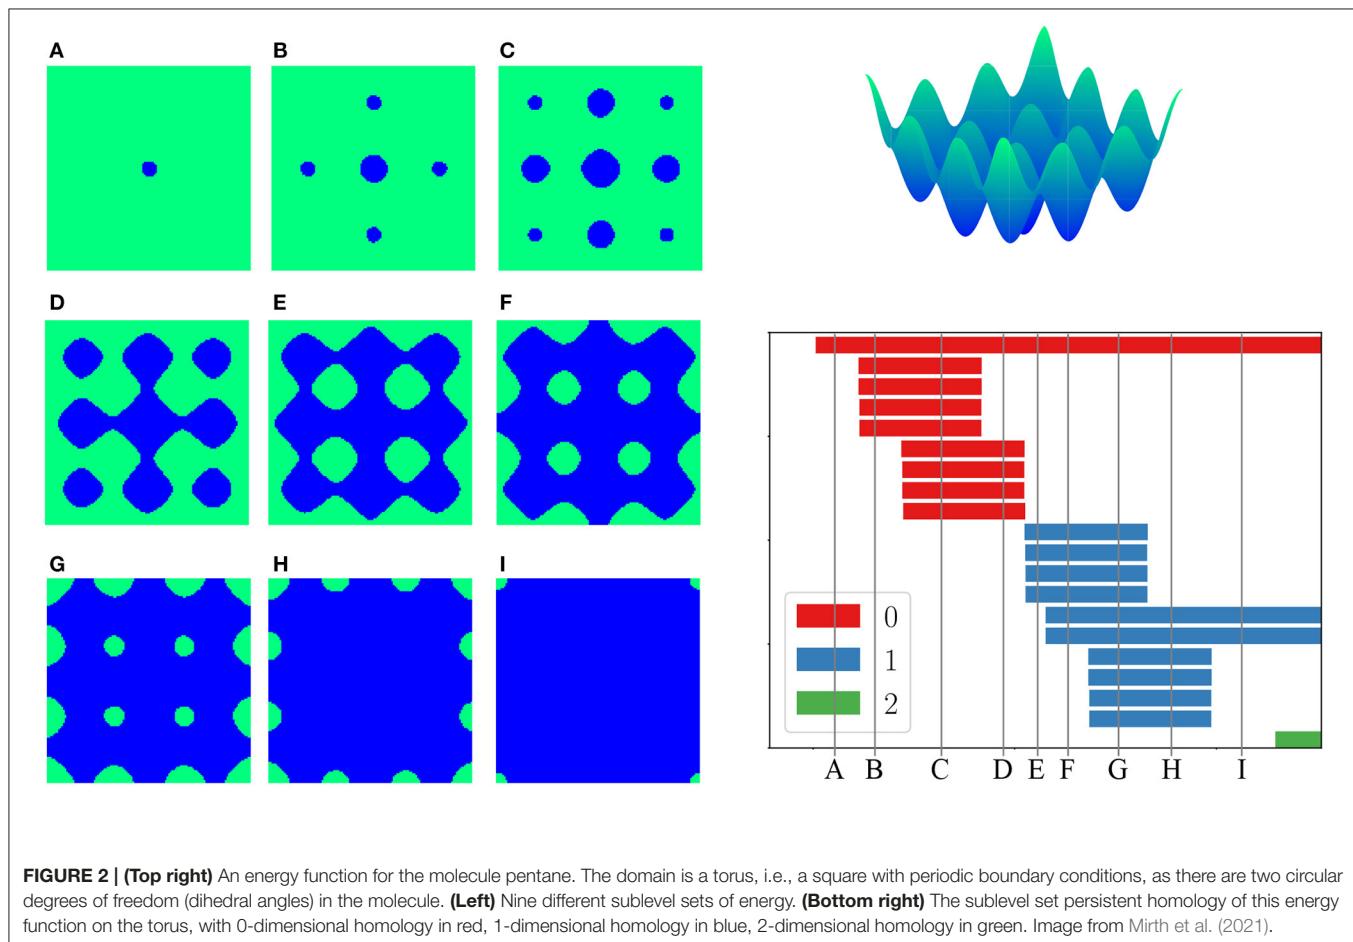


FIGURE 2 | (Top right) An energy function for the molecule pentane. The domain is a torus, i.e., a square with periodic boundary conditions, as there are two circular degrees of freedom (dihedral angles) in the molecule. **(Left)** Nine different sublevel sets of energy. **(Bottom right)** The sublevel set persistent homology of this energy function on the torus, with 0-dimensional homology in red, 1-dimensional homology in blue, 2-dimensional homology in green. Image from Mirth et al. (2021).

coordinates. This map realizes the conformation space of cyclo-octane as a subset of \mathbb{R}^{24} , and then we mod out by rigid rotations and translations. Topologically, the conformation space of cyclo-octane turns out to be the union of a sphere with a Klein bottle, glued together along two circles of singularities (see **Figure 4**, right). This model was obtained by Martin et al. (2010), Martin and Watson (2011), and Brown et al. (2008), who furthermore obtain a triangulation of this dataset (a representation of the dataset as a union of vertices, edges, and triangles).

A Klein bottle, like a sphere, is a 2-dimensional manifold. Whereas, a sphere can be embedded in 3-dimensional space, a Klein bottle requires at least four dimensions in order to be embedded without self-intersections. When a sphere and Klein

bottle are glued together along two circles, the union is no longer a manifold. Indeed, near the gluing circles, the space does not look like a sheet of paper, but instead like the tail of a dart with four fins, i.e., the letter “X” crossed with the interval $[0, 1]$. However, the result is still a 2-dimensional stratified space. In **Figure 5**, we compute the persistent homology of a point cloud dataset of 1,000,000 cyclo-octane molecule configurations. The short bars are interpreted as noise, whereas the long bars are interpreted as attributes of the underlying shape. We obtain a single connected component, a single 1-dimensional hole, and two 2-dimensional homology features. These homology signatures agree with the homology of the union of a sphere with a Klein bottle, glued together along two circles of singularities.

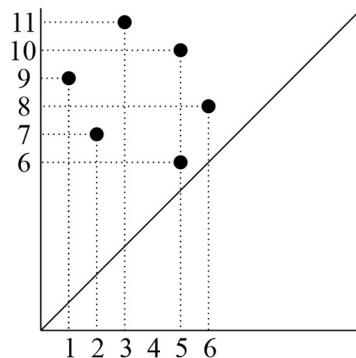


FIGURE 3 | (Left) A persistent homology barcode, with the birth and death scale of each bar indicated on the horizontal axis. **(Right)** Its corresponding persistence diagram, i.e., a collection of points in the first quadrant above the diagonal, with birth coordinates on the horizontal axis and death coordinates on the vertical axis.

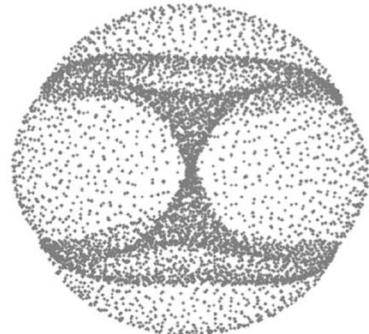
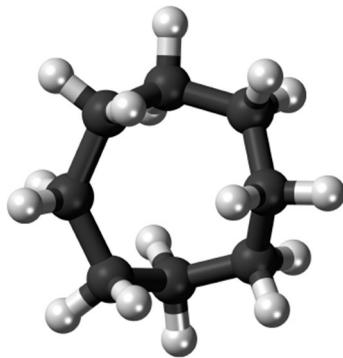


FIGURE 4 | (Left) The cyclo-octane molecule consists of a ring of 8 carbon atoms (black), each bonded to a pair of hydrogen atoms (white). **(Right)** A PCA projection of a dataset of different conformations of the cyclo-octane molecule; this shape is a sphere glued to a Klein bottle (the “hourglass”) along two circles of singularity. The right image is from Martin et al. (2010).

One of the first applications of persistent homology was to measure the global shape of a dataset of image patches (Carlsson et al., 2008). This dataset of natural 3×3 pixel patches from black-and-white photographs from indoor and outdoor scenes in fact has three different global shapes! The most common patches lie along a circle of possible directions of linear gradient patches (varying from black to gray to white). The next most common patches lie along a three circle model, additionally including a circle’s worth of horizontal quadratic gradients, and a circle’s worth of vertical quadratic gradients. At the next level of resolution, the most common patches in some sense lie along a Klein bottle. All three of these models—the circle, the three circles, and the Klein bottle—are global models, summarizing the global shape of the dataset at different resolutions.

4. EXAMPLES MEASURING LOCAL GEOMETRY

Though a single long bar in persistent homology may carry a lot of information, a single small bar typically does not. However, together a collection of small bars may unexpectedly carry a large

amount of geometric content. A long bar is a trumpet solo—piercing through to be heard over the orchestra with ease. The small bars are the string section—each small bar on its own is relatively quiet, but in concert the small bars together deliver a powerful message. We survey several modern examples where small persistent homology bars are now the signal, instead of the noise.

Birds, fish, and insects move as flocks, schools, and hordes in a way which is determined by *collective motion*: each animal’s next motion is a random function of the location of its nearby neighbors. In a flock of thousands of birds, there is an impressively large amount of time-varying geometry, including for example all $\binom{n}{2}$ pairwise distances, where n is the number of birds (see Figure 6). How can one summarize this much geometric content for use in machine learning tasks, say to predict how the motion of the flock will vary next, or to predict some of the parameters in a mathematical model approximately governing the motion of the birds? Persistent homology has been used in Topaz et al. (2015), Ulmer et al. (2019), Bhaskar et al. (2019), Adams et al. (2020b), and Xian et al. (2020) to reduce a large collection of geometric content down to a concise summary. These datasets of animal swarms do not lie along

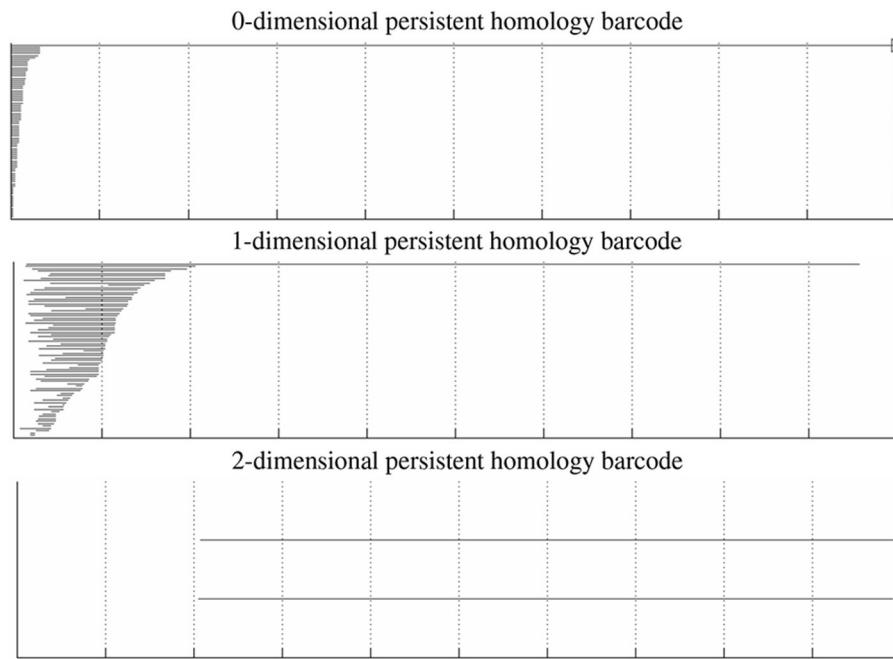


FIGURE 5 | 0-, 1-, and 2-dimensional persistent homology barcodes for the cyclo-octane dataset. The horizontal axis corresponds to the birth and death scale of the bars, and the vertical axis is an arbitrary ordering of the bars (here by death scale).



FIGURE 6 | A large amount of local and global geometric information is contained in a flock of birds.

beautiful manifolds (global shapes), but nevertheless there is a wealth of information in the local geometry as measured by the short persistent homology bars. For example, Ulmer et al. (2019) show via time-varying persistent homology³ that a control model for aphid motion, in which aphids move independently at random, does not fit experimental data as well as a model incorporating social interaction (distances to nearby neighbors) between the aphids.

Other recent work has used persistent homology to characterize the complexity of geometric objects. Bendich

et al. (2016b) apply sublevel set persistent homology to the study of brain artery trees, examining the effects of age and sex on the barcodes generated from artery trees. While younger brains have artery trees containing more local twisting and branching, older brains are sparser with fewer small branches and leaves. The authors use the 100 longest bars in dimensions 0 and 1 in their analysis, and they further examine which lengths of bars give the highest correlation with age and sex. For instance, when examining age, they find it is not the longest bars, but instead the bars of medium length (roughly the 21st through 40th longest bars) that are the most discriminatory.

In other datasets where points are nearly evenly spaced, barcodes will consist of bars with mostly similar birth and death times. Consider for instance the point cloud persistent homology for a square grid of points in the plane: all non-infinite 0-dimensional bars are identical and adding a small amount of noise to the points will result in a small change to the bars. The same is true for 1-dimensional bars. With this in mind, Motta et al. (2018) use persistent homology to measure the order, or regularity, of lattice-like datasets, focusing on hexagonal grids formed by ion bombardment of solid surfaces (see Figure 7). The authors' techniques use the variance of 0-dimensional homology bar lengths, and the sum of the lengths of 1-dimensional homology bars, as well as a particular linear combination of the two especially suited to hexagonal lattices. Their results suggest that techniques based on persistent homology can provide useful measures of order that are sensitive to both large scale and small scale defects in lattices. Point cloud persistence has also been used to summarize the local order and randomness in other materials science and chemistry

³In particular, the *crocker plot* (Topaz et al., 2015).

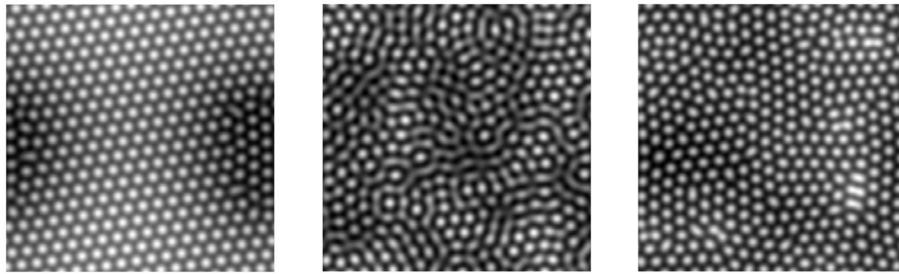


FIGURE 7 | Hexagonal lattices, of varying degree of regularity, created by ion bombardment. Figures from Motta et al. (2018).

contexts, including amorphous solids and glass (Nakamura et al., 2015; Hiraoka et al., 2016; Hirata et al., 2020), nanoporous materials used in gas adsorption (Krishnapriyan et al., 2020), crystal structure (Maroulas et al., 2020), and protein folding (Xia and Wei, 2014; Cang and Wei, 2018).

Though the above examples focus on point cloud persistence, sublevel set persistent homology has also been used to detect the local geometry of functions. Kramár et al. (2016) use sublevel set persistence to summarize the complicated spatio-temporal patterns that arise from dynamical systems modeling fluid flow, including turbulence (Kolmogorov flow) and heat convection (Rayleigh-Bénard convection). With sublevel set persistence, Zeppelzauer et al. (2016) improve 3D surface classification, including on an archaeology task of segmenting engraved regions of rock from the surrounding natural rock surface. In a task of tracking automobiles, Bendich et al. (2016a) use the sublevel set persistent homology of driver speeds in order to characterize driver behaviors and prune out improbable paths from their multiple hypothesis tracking framework.

5. THEORY OF HOW PERSISTENT HOMOLOGY MEASURES LOCAL GEOMETRY

Recent work has begun to formalize the idea that persistent homology measures local geometry. Bubenik et al. (2020) explore the effect of the curvature of a space on the persistent homology of a sample of points, focusing on disks in spaces with constant curvature. Their work includes theoretical results about the persistence of triangles in these spaces, and they are also able to demonstrate experimentally that persistent homology in dimensions 0 and 1 can be used to accurately estimate the curvature given a random sample of points. Since the disks in spaces with different curvature are homeomorphic, the differences in persistent homology cannot be due to topology, but rather result from the geometric features of the spaces.

Fractal dimension is another measure of local geometry, and indeed some of the earliest applications of persistent homology in Vanessa Robins' Ph.D. thesis were motivated as a way to capture the fractal dimension of an infinite set in Euclidean space (Robins, 2000; MacPherson and Schweinhart, 2012). Can this also be applied to datasets, i.e., to random collections of

finite sets of points? Given a random sample of points from a measure, Adams et al. (2020a) use persistent homology to detect the fractal dimension of the support of the measure. This notion of *persistent homology fractal dimension* agrees with the Hausdorff/box-counting dimension for 0-dimensional persistent homology and a restricted class of measures; see Schweinhart (2019, 2020) for further theoretical developments.

A related line of work studies what can be proven about the topology of random point clouds, typically as the number of points in the point cloud goes to infinity (Kahle, 2011; Adler et al., 2014; Bobrowski and Kahle, 2014; Bobrowski et al., 2017). The magnitude (Leinster, 2013) and magnitude homology (Hepworth and Willerton, 2017; Leinster and Shulman, 2017) of a metric space measure both local and global properties; recent and ongoing work is being done to connect magnitude with persistent homology (Otter, 2018; Govc and Hepworth, 2021). See also Weinberger (2019) for connections between sublevel set persistent homology and the geometry of spaces of functions, including Lipschitz constants of functions. We predict that much more work demonstrating how local geometric features can be recovered from persistent homology barcodes will take place over the next decade.

6. MACHINE LEARNING

Because persistent homology gives a concise description of the shape of data, it is not surprising that recent work has incorporated persistent homology into machine learning. When might one consider using persistent homology in concert with machine learning, as opposed to other more classical machine learning techniques measuring shape such as clustering (Xu and Wunsch, 2005) or nonlinear dimensionality reduction (Roweis and Saul, 2000; Tenenbaum et al., 2000; Kohonen, 2012; McInnes et al., 2018)? We recommend persistent homology when one desires either (i) a quantitative reductive summary of local geometry, (ii) an estimate of the number or size of more global topological features in a dataset, or (iii) a way to explore if either local geometry or global topology may be discriminatory for the machine learning task at hand. Researchers have taken at least three distinct approaches: persistence barcodes have been adapted to be input to machine learning algorithms, topological methods have been used to create new algorithms, and persistent homology has been used to analyze machine learning algorithms.

Perhaps the most natural of these approaches is inputting persistence data into a machine learning algorithm. Though the persistent homology bars provide a summary of both local geometry and global topology, for a quantitative summary to be fully applicable it needs to be amenable for use in machine learning tasks. The space of persistence barcodes is not immediately appropriate for machine learning. Indeed, averages of barcodes need not be unique (Mileyko et al., 2011), and the space of persistence barcodes does not coarsely embed into any Hilbert space (Bubenik and Wagner, 2020). These limitations have initiated a large amount of research on transforming persistence barcodes into more natural formats for machine learning. From barcodes, Bubenik (2015) creates *persistence landscapes*, which live in a Banach space of functions⁴. Persistence landscapes are created by rotating a persistence diagram on its side—so that the diagonal line $y = x$ becomes as flat as the horizon—and then using the persistence diagram points to trace out the peaks in a mountain landscape profile. A landscape can then be discretized by taking a finite sample of the function values, allowing it to be used in machine learning tasks (see for instance Kovacev-Nikolic et al. (2016)). From barcodes, Adams et al. (2017) create *persistence images*, a Euclidean vectorization enabling a diverse class of machine learning tools to be applied (see also Chen et al., 2015; Reininghaus et al., 2015). A persistence image is created by taking a sum of Gaussians, one centered on each point in a persistence diagram, and then pixelating that surface to form an image. By analogy, recall that in point cloud persistent homology, one “blurs their vision” when looking at a dataset by replacing each data point with a ball—this is similar to the process of “blurring one’s vision” when looking at a persistence diagram in order to create a persistence image.

Persistence landscapes were defined as part of an effort to give a firm statistical foundation to persistent homology. In fact, Bubenik (2015) proves a strong law of large numbers and a central limit theorem for persistence landscapes. This allows one to discuss hypothesis testing with persistent homology. Another approach to hypothesis testing is given by Robinson and Turner (2017). Other statistical approaches include Fasy et al. (2014), which describes confidence intervals and a statistical approach to distinguishing important features from noise, Divol and Polonik (2019) and Maroulas et al. (2019), which consider probability density functions for persistence diagrams, and Maroulas et al. (2020), which describes a Bayesian framework. See Wasserman (2018) for a review of statistical techniques in the context of topological data analysis.

Persistence landscapes and images are only two of the many different methods that have recently been invented in order to transform persistence barcodes into machine learning input. Algorithms that require only a distance matrix, such as many clustering or dimensionality reduction algorithms, can be applied on the bottleneck or Wasserstein distances between persistence barcodes (Cohen-Steiner et al., 2007; Mileyko et al., 2011; Kerber et al., 2017). Other techniques for vectorizing

⁴In practice, a different metric is sometimes chosen to map landscapes into a Hilbert space, though the restrictions of Bubenik and Wagner (2020) apply.

persistence barcodes involve heat kernels (Carriére et al., 2015), entropy (Merelli et al., 2015; Atienza et al., 2020), rings of algebraic functions (Adcock et al., 2016), tropical coordinates (Kališnik, 2019), complex polynomials (Di Fabio and Ferri, 2015), and optimal transport (Carriére et al., 2017), among others. Some of these techniques, including those by Zhao and Wang (2019) and Divol and Polonik (2019), allow one to learn the vectorization parameters that are best suited for a machine learning task on a given dataset. Others allow one to plug persistent homology information directly into a neural network (Hofer et al., 2017). Recent research on incorporating persistence as input for machine learning is vast and varied, and the above collection of references is far from complete.

As for the creation of new algorithms, persistent homology has recently been applied to regularization, a technique used in machine learning that penalizes overly complicated models to avoid overfitting. Chen et al. (2019) propose a “topological penalty function” for classification algorithms, which encourages a topologically simple decision boundary. Their method is based on measuring the relative importance of various connected components of the decision boundary via 0-dimensional persistent homology. They show how the gradient of such a penalty function can be computed, which is important for use in machine learning algorithms, and demonstrate their method on several examples. Similar work using topological methods to examine a decision boundary can also be found in Varshney and Ramamurthy (2015) and Ramamurthy et al. (2019).

Finally, other recent work has used persistent homology to analyze neural networks. Naitzat et al. (2020) provide experimental evidence that neural networks operate by simplifying the topology of a dataset. They examine the topology of a dataset and its images at the various layers of a neural network performing classification, finding that the corresponding barcodes become simpler as the data progresses through the network. Additionally, they observe the effects of different shapes of neural networks and different activation functions. They find that deeper neural networks have a tendency to simplify the topology of a dataset more gradually than shallow networks, and that networks with ReLU activation tend to simplify topology more in the earlier layers of a network than other activation functions.

7. CONCLUSION

Topological tools are often described as being able to stitch local data together in order to describe global features: from local to global. The history of applied topology, however, has in some sense gone in the reverse direction—from global to local—as surveyed above! Applied topology was developed in part to summarize global features in a point cloud dataset, as in the examples of the conformations of the cyclo-octane molecule or the collection of 3×3 pixel patches from images. If global shapes are the focus, long persistent homology bars are interpreted as the relevant features, while small bars are often disregarded as sampling artifacts or noise. However, in more recent applications, and in particular when using applied topology in concert with

machine learning, it is often many short persistent homology bars that *together* form the signal. One of the biggest benefits of applied topology is that one need not choose a scale beforehand: persistent homology provides a useful summary of both the local and global features in a dataset, and this summary has been made accessible for use in machine learning tasks.

We have seen how the short bars can be a measure of local geometry, texture, curvature, and fractal dimension; their sensitivity to various features of datasets leads to the wide variety of applications surveyed here. Because persistent homology provides a concise, reductive view of the geometry of a dataset, for instance in the examples studying brain artery trees or hexagonal grids, it is not hard to imagine the potential applications to machine learning problems. This has led to recent techniques that turn barcodes into machine learning input, exemplified by persistence landscapes and persistence images. We hope that this

wealth of recent work, which has shifted more attention to short persistent homology bars and the geometric information they summarize, will inspire further research at the intersection of applied topology, local geometry, and machine learning.

AUTHOR CONTRIBUTIONS

HA originally presented this material in conference talks. HA and MM reviewed literature and contributed to the writing of the article. All authors contributed to the article and approved the submitted version.

FUNDING

This material was based upon work supported by the National Science Foundation under Grant Number 1934725.

REFERENCES

Adams, H., Aminian, M., Farnell, E., Kirby, M., Mirth, J., Neville, R., et al. (2020a). “A fractal dimension for measures via persistent homology,” in *Topological Data Analysis*, eds N. A. Baas, G. E. Carlsson, G. Quick, M. Szymik, and M. Thaule (Cham: Springer International Publishing), 1–31. doi: 10.1007/978-3-030-43408-3_1

Adams, H., Chepushtanova, S., Emerson, T., Hanson, E., Kirby, M., Motta, F., et al. (2017). Persistence images: A vector representation of persistent homology. *J. Mach. Learn. Res.* 18, 1–35.

Adams, H., Ciocanel, M. V., Topaz, C. M., and Ziegelmeier, L. (2020b). Topological data analysis of collective motion. *SIAM News* 53, 1–4. Available online at: <https://sinews.siam.org/Details-Page/topological-data-analysis-of-collective-motion>

Adcock, A., Carlsson, E., and Carlsson, G. (2016). The ring of algebraic functions on persistence bar codes. *Homot. Homol. Appl.* 18, 341–402. doi: 10.4310/HHA.2016.v18.n1.a21

Adler, R. J., Bobrowski, O., and Weinberger, S. (2014). Crackle: The homology of noise. *Discr. Comput. Geometry* 52, 680–704. doi: 10.1007/s00454-014-9621-6

Atienza, N., González-Díaz, R., and Sorianó-Trigueros, M. (2020). On the stability of persistent entropy and new summary functions for TDA. *Pattern Recognit.* 107:107509. doi: 10.1016/j.patcog.2020.107509

Bendich, P., Chin, S. P., Clark, J., Desena, J., Harer, J., Munch, E., et al. (2016a). Topological and statistical behavior classifiers for tracking applications. *IEEE Trans. Aerosp. Electron. Syst.* 52, 2644–2661. doi: 10.1109/TAES.2016.160405

Bendich, P., Marron, J. S., Miller, E., Pieloch, A., and Skwerer, S. (2016b). Persistent homology analysis of brain artery trees. *Ann. Appl. Stat.* 10:198. doi: 10.1214/15-AOAS886

Bhaskar, D., Manhart, A., Milzman, J., Nardini, J. T., Storey, K. M., Topaz, C. M., et al. (2019). Analyzing collective motion with machine learning and topology. *Chaos* 29:123125. doi: 10.1063/1.5125493

Bobrowski, O., and Kahle, M. (2014). Topology of random geometric complexes: a survey. *J. Appl. Comput. Topol.* 1, 331–364. doi: 10.1007/s41468-017-0010-0

Bobrowski, O., Kahle, M., and Skraba, P. (2017). Maximally persistent cycles in random geometric complexes. *Ann. Appl. Probab.* 27, 2032–2060. doi: 10.1214/16-AAP1232

Brown, M. W., Martin, S., Pollock, S. N., Coutsias, E. A., and Watson, J. P. (2008). Algorithmic dimensionality reduction for molecular structure analysis. *J. Chem. Phys.* 129:064118. doi: 10.1063/1.2968610

Bubenik, P. (2015). Statistical topological data analysis using persistence landscapes. *J. Mach. Learn. Res.* 16, 77–102.

Bubenik, P., Hull, M., Patel, D., and Whittle, B. (2020). Persistent homology detects curvature. *Inverse Probl.* 36:025008. doi: 10.1088/1361-6420/ab4ac0

Bubenik, P., and Wagner, A. (2020). Embeddings of persistence diagrams into hilbert spaces. *J. Appl. Comput. Topol.* 4, 339–351. doi: 10.1007/s41468-020-00056-w

Cang, Z., and Wei, G.-W. (2018). Integration of element specific persistent homology and machine learning for protein-ligand binding affinity prediction. *Int. J. Numer. Methods Biomed. Eng.* 34, e2914. doi: 10.1002/cnm.2914

Carlsson, G. (2009). Topology and data. *Bull. Am. Math. Soc.* 46, 255–308. doi: 10.1090/S0273-0979-09-01249-X

Carlsson, G., Ishkhanov, T., de Silva, V., and Zomorodian, A. (2008). On the local behavior of spaces of natural images. *Int. J. Comput. Vis.* 76, 1–12. doi: 10.1007/s11263-007-0056-x

Carriére, M., Cuturi, M., and Oudot, S. (2017). “Sliced wasserstein kernel for persistence diagrams,” in *International Conference on Machine Learning, PMLR* (Sydney, VIC), 664–673.

Carriére, M., Oudot, S. Y., and Ovsjanikov, M. (2015). “Stable topological signatures for points on 3d shapes,” in *Computer Graphics Forum*, Vol. 34 (Wiley Online Library), 1–12. doi: 10.1111/cgf.12692

Chazal, F., de Silva, V., Glisse, M., and Oudot, S. (2016). *The Structure and Stability of Persistence Modules*. Springer. doi: 10.1007/978-3-319-42545-0

Chazal, F., de Silva, V., and Oudot, S. (2014). Persistence stability for geometric complexes. *Geometr. Dedic.* 173, 193–214. doi: 10.1007/s10711-013-9937-z

Chazal, F., and Oudot, S. (2008). “Towards persistence-based reconstruction in Euclidean spaces,” in *Proceedings of the 24th Annual Symposium on Computational Geometry* (College Park, MD: ACM), 232–241. doi: 10.1145/1377676.1377719

Chen, C., Ni, X., Bai, Q., and Wang, Y. (2019). “A topological regularizer for classifiers via persistent homology,” in *Proceedings of Machine Learning Research*, Vol. 89, eds K. Chaudhuri and M. Sugiyama (Naha), 2573–2582.

Chen, Y.-C., Wang, D., Rinaldo, A., and Wasserman, L. (2015). Statistical analysis of persistence intensity functions. *arXiv preprint arXiv:1510.02502*.

Cohen-Steiner, D., Edelsbrunner, H., and Harer, J. (2007). Stability of persistence diagrams. *Discret. Comput. Geomet.* 37, 103–120. doi: 10.1007/s00454-006-1276-5

Di Fabio, B., and Ferri, M. (2015). “Comparing persistence diagrams through complex vectors,” in *International Conference on Image Analysis and Processing* (Genoa: Springer), 294–305. doi: 10.1007/978-3-319-23231-7_27

Divilov, V., and Polonik, W. (2019). On the choice of weight functions for linear representations of persistence diagrams. *J. Appl. Comput. Topol.* 3, 249–283. doi: 10.1007/s41468-019-00032-z

Edelsbrunner, H., Letscher, D., and Zomorodian, A. (2000). “Topological persistence and simplification,” in *41st Annual Symposium on Foundations of Computer Science, 2000* (Redondo Beach, CA: IEEE), 454–463. doi: 10.1109/SFCS.2000.892133

Fasy, B. T., Lecci, F., Rinaldo, A., Wasserman, L., Sivaraman, B., and Singh, A. (2014). Confidence sets for persistence diagrams. *Ann. Stat.* 42, 2301–2339. doi: 10.1214/14-AOS1252

Govc, D., and Hepworth, R. (2021). Persistent magnitude. *J. Pure Appl. Algeb.* 225:106517. doi: 10.1016/j.jpaa.2020.106517

Hepworth, R., and Willerton, S. (2017). Categorifying the magnitude of a graph. *Homol. Homotopy Appl.* 19, 31–60. doi: 10.4310/HHA.2017.v19.n2.a3

Hiraoka, Y., Nakamura, T., Hirata, A., Escolar, E. G., Matsue, K., and Nishiura, Y. (2016). Hierarchical structures of amorphous solids characterized by persistent homology. *Proc. Natl. Acad. Sci. U.S.A.* 113, 7035–7040. doi: 10.1073/pnas.1520877113

Hirata, A., Wada, T., Obayashi, I., and Hiraoka, Y. (2020). Structural changes during glass formation extracted by computational homology with machine learning. *Commun. Mater.* 1, 1–8. doi: 10.1038/s43246-020-00100-3

Hofer, C., Kwitt, R., Niethammer, M., and Uhl, A. (2017). “Deep learning with topological signatures,” in *Advances in Neural Information Processing Systems*, eds I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Long Beach, CA: Curran Associates, Inc), 30, 1634–1644.

Kahle, M. (2011). Random geometric complexes. *Discr. Comput. Geometry* 45, 553–573. doi: 10.1007/s00454-010-9319-3

Kališnik, S. (2019). Tropical coordinates on the space of persistence barcodes. *Found. Comput. Math.* 19, 101–129. doi: 10.1007/s10208-018-9379-y

Kerber, M., Morozov, D., and Nigmetov, A. (2017). Geometry helps to compare persistence diagrams. *ACM. J. Exp. Algorithms.* 22, 1–20. doi: 10.1145/3064175

Kohonen, T. (2012). *Self-Organizing Maps*, Vol. 30. Berlin: Springer Science & Business Media.

Kovacev-Nikolic, V., Bubenik, P., Nikolic, D., and Heo, G. (2016). Using persistent homology and dynamical distances to analyze protein binding. *Stat. Appl. Genet. Mol. Biol.* 15, 19–38. doi: 10.1515/sagmb-2015-0057

Kramár, M., Levanger, R., Tithof, J., Suri, B., Xu, M., Paul, M., et al. (2016). Analysis of Kolmogorov flow and Rayleigh-Bénard convection using persistent homology. *Phys. D* 334, 82–98. doi: 10.1016/j.physd.2016.02.003

Krishnapriyan, A. S., Montoya, J., Hummelshøj, J., and Morozov, D. (2020). Persistent homology advances interpretable machine learning for nanoporous materials. *arXiv preprint arXiv:2010.00532*.

Leinster, T. (2013). The magnitude of metric spaces. *Doc. Math.* 18, 857–905.

Leinster, T., and Shulman, M. (2017). Magnitude homology of enriched categories and metric spaces. *arXiv preprint arXiv:1711.00802*.

MacPherson, R., and Schweinhart, B. (2012). Measuring shape with topology. *J. Math. Phys.* 53:073516. doi: 10.1063/1.4737391

Maroulas, V., Mike, J. L., and Oballe, C. (2019). Nonparametric estimation of probability density functions of random persistence diagrams. *J. Mach. Learn. Res.* 20, 1–49.

Maroulas, V., Nasrin, F., and Oballe, C. (2020). A Bayesian framework for persistent homology. *SIAM J. Math. Data Sci.* 2, 48–74. doi: 10.1137/19M1268719

Martin, S., Thompson, A., Coutsias, E. A., and Watson, J.-P. (2010). Topology of cyclo-octane energy landscape. *J. Chem. Phys.* 132:234115. doi: 10.1063/1.3445267

Martin, S., and Watson, J. P. (2011). Non-manifold surface reconstruction from high-dimensional point cloud data. *Comput. Geometry* 44, 427–441. doi: 10.1016/j.comgeo.2011.05.002

McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*. doi: 10.21105/joss.00861

Merelli, E., Rucco, M., Sloot, P., and Tesei, L. (2015). Topological characterization of complex systems: Using persistent entropy. *Entropy* 17, 6872–6892. doi: 10.3390/e17106872

Mileyko, Y., Mukherjee, S., and Harer, J. (2011). Probability measures on the space of persistence diagrams. *Inverse Probl.* 27:124007. doi: 10.1088/0266-5611/27/12/124007

Mirth, J., Zhai, Y., Bush, J., Alvarado, E. G., Jordan, H., Heim, M., et al. (2021). Representations of energy landscapes by sublevelset persistent homology: an example with n-alkanes. *J. Chem. Phys.* 154:114114.

Motta, F. C., Neville, R., Shipman, P. D., Pearson, D. A., and Bradley, R. M. (2018). Measures of order for nearly hexagonal lattices. *Phys. D* 380, 17–30. doi: 10.1016/j.physd.2018.05.005

Naitzat, G., Zhitnikov, A., and Lim, L.-H. (2020). Topology of deep neural networks. *J. Mach. Learn. Res.* 21, 1–40.

Nakamura, T., Hiraoka, Y., Hirata, A., Escolar, E. G., and Nishiura, Y. (2015). Persistent homology and many-body atomic structure for medium-range order in the glass. *Nanotechnology* 26:304001. doi: 10.1088/0957-4484/26/30/304001

Otter, N. (2018). Magnitude meets persistence. Homology theories for filtered simplicial sets. *arXiv preprint arXiv:1807.01540*.

Ramamurthy, K. N., Varshney, K., and Mody, K. (2019). “Topological data analysis of decision boundaries with application to model selection,” in *Proceedings of the 36th International Conference on Machine Learning, Vol. 97 of Proceedings of Machine Learning Research*, eds K. Chaudhuri and R. Salakhutdinov (Long Beach, CA), 5351–5360.

Reininghaus, J., Huber, S., Bauer, U., and Kwitt, R. (2015). “A stable multi-scale kernel for topological machine learning,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Boston, MA), 4741–4748. doi: 10.1109/CVPR.2015.7299106

Robins, V. (2000). *Computational topology at multiple resolutions: Foundations and applications to fractals and dynamics* (Ph.D. thesis). Boulder, CO: University of Colorado.

Robinson, A., and Turner, K. (2017). Hypothesis testing for topological data analysis. *J. Appl. Comput. Topol.* 1, 241–261. doi: 10.1007/s41468-017-0008-7

Roweis, S. T., and Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 2323–2326. doi: 10.1126/science.290.5555.02323

Schweinhart, B. (2019). Persistent homology and the upper box dimension. *Discr. Comput. Geometry* 65, 331–364. doi: 10.1007/s00454-019-00145-3

Schweinhart, B. (2020). Fractal dimension and the persistent homology of random geometric complexes. *Adv. Math.* 372:107291. doi: 10.1016/j.aim.2020.107291

Tenenbaum, J. B., De Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2323. doi: 10.1126/science.290.5500.2319

Topaz, C. M., Ziegelmeier, L., and Halverson, T. (2015). Topological data analysis of biological aggregation models. *PLoS ONE* 10:e0126383. doi: 10.1371/journal.pone.0126383

Ulmer, M., Ziegelmeier, L., and Topaz, C. M. (2019). A topological approach to selecting models of biological experiments. *PLoS ONE* 14:e0213679. doi: 10.1371/journal.pone.0213679

Varshney, K. R., and Ramamurthy, K. N. (2015). “Persistent topology of decision boundaries,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Processing* (Brisbane, QLD), 3931–3935. doi: 10.1109/ICASSP.2015.7178708

Wasserman, L. (2018). Topological data analysis. *Annu. Rev. Stat. Appl.* 5, 501–532. doi: 10.1146/annurev-statistics-031017-100045

Weinberger, S. (2019). Interpolation, the rudimentary geometry of spaces of Lipschitz functions, and geometric complexity. *Found. Comput. Math.* 19, 991–1011. doi: 10.1007/s10208-019-09416-0

Xia, K., and Wei, G.-W. (2014). Persistent homology analysis of protein structure, flexibility, and folding. *International J. Numer. Methods Biomed. Eng.* 30, 814–844. doi: 10.1002/cnm.2655

Xian, L., Adams, H., Topaz, C. M., and Ziegelmeier, L. (2020). Capturing dynamics of time-varying data via topology. *arXiv preprint arXiv:2010.05780*.

Xu, R., and Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Trans. Neural Netw.* 16, 645–678. doi: 10.1109/TNN.2005.845141

Zeppelzauer, M., Zieliński, B., Juda, M., and Seidl, M. (2016). “Topological descriptors for 3d surface analysis,” in *International Workshop on Computational Topology in Image Context* (Springer), 77–87. doi: 10.1007/978-3-319-39441-1_8

Zhao, Q., and Wang, Y. (2019). “Learning metrics for persistence-based summaries and applications for graph classification,” in *Advances in Neural Information Processing Systems*, 9859–9870.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Adams and Moy. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.