Of Shapes and Spaces

Geometry, Topology, and Machine Learning

Bastian Rieck

Setting the stage

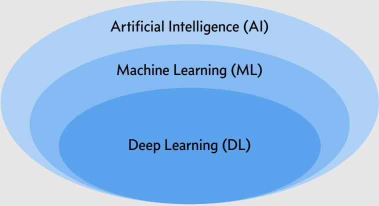

What is a neural network?

Perspective I: Neural networks as universal function approximators

Theorem (Universal function approximation)

Let $\sigma \in C(\mathbb{R}, \mathbb{R})$ *be a non-polynomial activation function. For every* $n, m \in \mathbb{N}$*, every compact subset* $K \subseteq \mathbb{R}^n$*, every function* $f \in C(K, \mathbb{R}^m)$ *and* $\epsilon > 0$*, there exist* $k \in \mathbb{N}$*,* $\mathbf{A} \in \mathbb{R}^{k \times n}$*,* $b \in \mathbb{R}^k$*, and* $\mathbf{C} \in \mathbb{R}^{m \times k}$ *such that*

$$\sup_{x \in K} \|f(x) - g(x)\| < \epsilon,$$

where $g(x) = \mathbf{C}\sigma(\mathbf{A}x + b)$*.*

A. Pinkus, *Approximation theory of the MLP model in neural networks*, *Acta Numerica* 8, 1999, pp. 143–195

What is a neural network?

Perspective II: Neural networks as computational building blocks

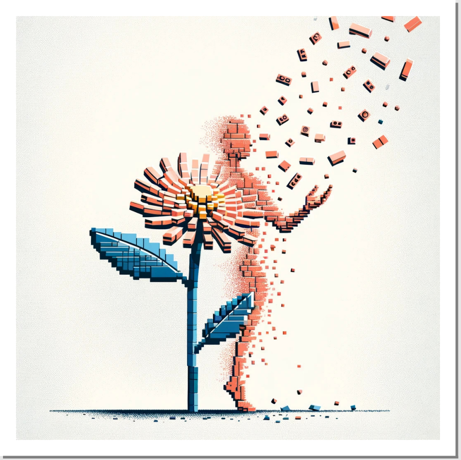

Maths is a continent

(Prof. Dr. Franka Miriam Brückler)

Topology and some of its fiefdoms

Point-set topology

What do *continuity*, *compactness*, and *connectivity* mean for general spaces?

Algebraic topology

What are characteristic properties of spaces that we can capture using algebra?

Differential topology

What are characteristic properties of spaces that we can capture using calculus?

What is algebraic topology?

- Develop invariants that classify topological spaces up to homeomorphism.

What is algebraic topology?

- Develop invariants that classify topological spaces up to homeomorphism.
- Use tools from algebra to study topological spaces.

What is algebraic topology?

- Develop invariants that classify topological spaces up to homeomorphism.
- Use tools from algebra to study topological spaces.
- **Understand shapes through calculations.**

A first taste

Seven Bridges of Königsberg

Is there a walk through the city that crosses every bridge *exactly* once?

A first taste

Seven Bridges of Königsberg

Is there a walk through the city that crosses every bridge *exactly* once?

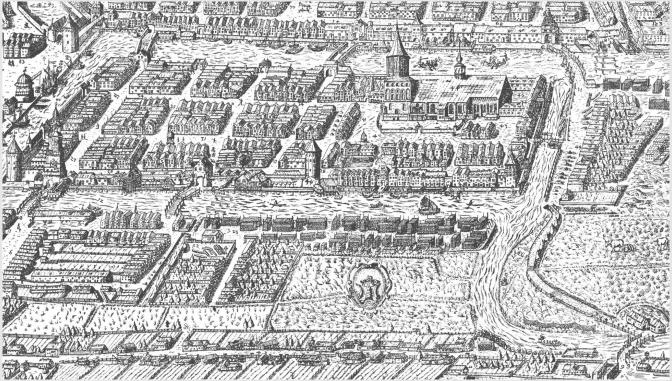

A first taste

Seven Bridges of Königsberg

Is there a walk through the city that crosses every bridge *exactly* once?

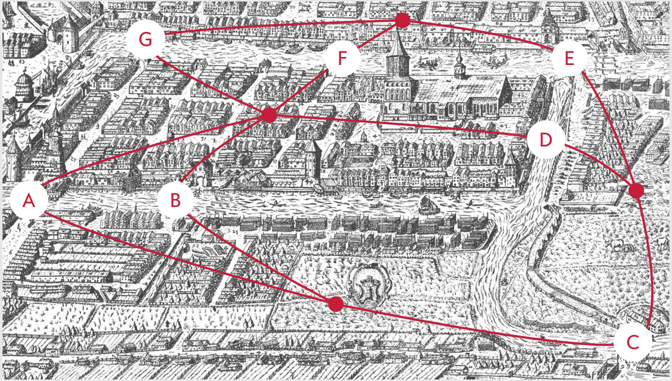

A first taste

Seven Bridges of Königsberg

Is there a walk through the city that crosses every bridge *exactly* once?

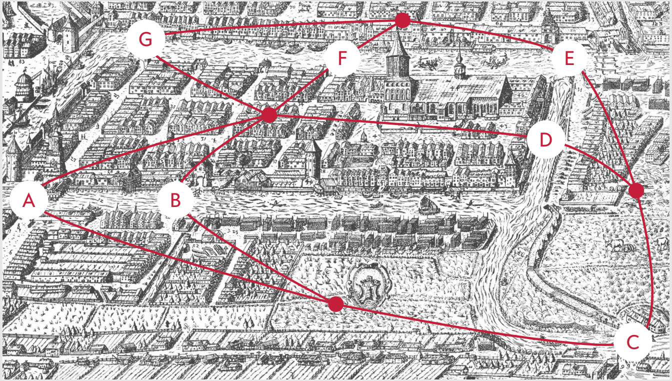

No such walk can exist because there are more than two vertices with *odd* degree!

A second taste

Euler characteristic

Definition

The *Euler characteristic* of a polyhedron is defined as $\chi := V - E + F$, where V is the number of vertices, E is the number of edges, and F is the number of faces, respectively.

A second taste

Euler characteristic

Definition

The *Euler characteristic* of a polyhedron is defined as $\chi := V - E + F$, where V is the number of vertices, E is the number of edges, and F is the number of faces, respectively.

Theorem

The Euler characteristic of every Platonic solid is $\chi = 2$.

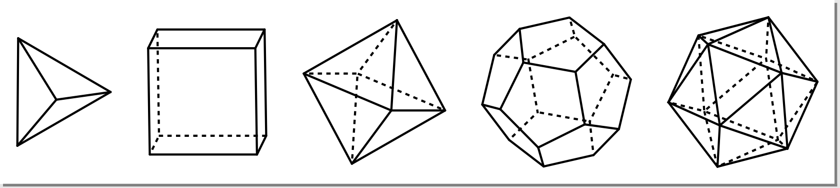

A third taste

Betti numbers

Space	β_0	β_1	β_2

The d^{th} Betti number counts the number of d-dimensional holes. It can be used to distinguish between spaces.

- $d = 0$: connected components
- $d = 1$: cycles
- $d = 2$: voids

A third taste

Betti numbers

Space		β_0	β_1	β_2
Point		1	0	0

The d^{th} Betti number counts the number of d-dimensional holes. It can be used to distinguish between spaces.

- $d = 0$: connected components
- $d = 1$: cycles
- $d = 2$: voids

A third taste

Betti numbers

The d^{th} Betti number counts the number of d-dimensional holes. It can be used to distinguish between spaces.

$d = 0$: connected components

$d = 1$: cycles

$d = 2$: voids

Space		β_0	β_1	β_2
Point		1	0	0
Cube	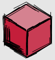	1	0	1

A third taste

Betti numbers

The d^{th} Betti number counts the number of d-dimensional holes. It can be used to distinguish between spaces.

- $d = 0$: connected components
- $d = 1$: cycles
- $d = 2$: voids

Space		β_0	β_1	β_2
Point		1	0	0
Cube	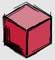	1	0	1
Sphere		1	0	1

A third taste

Betti numbers

The d^{th} Betti number counts the number of d-dimensional holes. It can be used to distinguish between spaces.

- $d = 0$: connected components
- $d = 1$: cycles
- $d = 2$: voids

Space		β_0	β_1	β_2
Point		1	0	0
Cube	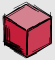	1	0	1
Sphere		1	0	1
Torus	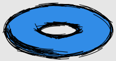	1	2	1

Reality is often messy…

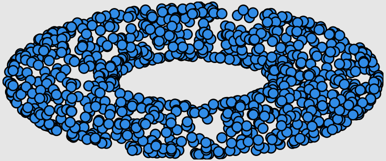

Reality is often messy…

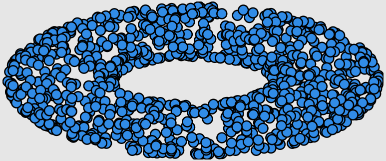

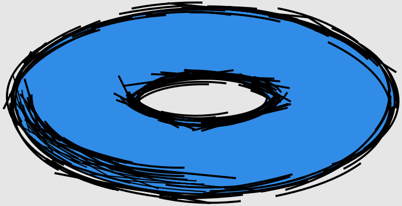

Representing spaces

Triangulations

Representing spaces

Triangulations

Representing spaces

Triangulations

Theorem

Every smooth manifold can be triangulated.

- S. S. Cairns, *Triangulation of the Manifold of Class One*, *Bulletin of the American Mathematical Society* 41.8, 1935, pp. 549–552
- J. H. C. Whitehead, *On C^1-Complexes*, *Annals of Mathematics* 41.4, 1940, pp. 809–824

Persistent homology

‘Points cross scales like clouds cross the sky’

Approximate a point cloud at different scales and observe how topological features appear and disappear.

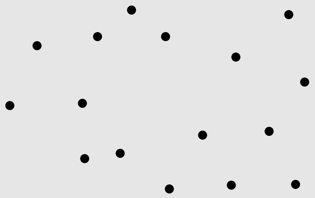

Persistent homology

‘Points cross scales like clouds cross the sky’

Approximate a point cloud at different scales and observe how topological features appear and disappear.

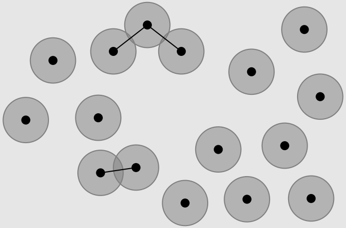

Persistent homology

'Points cross scales like clouds cross the sky'

Approximate a point cloud at different scales and observe how topological features appear and disappear.

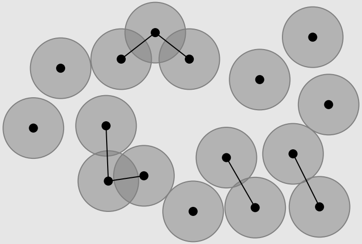

Persistent homology

‘Points cross scales like clouds cross the sky’

Approximate a point cloud at different scales and observe how topological features appear and disappear.

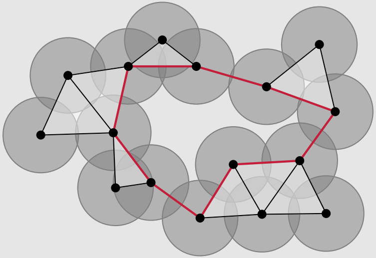

Persistent homology

'Points cross scales like clouds cross the sky'

Approximate a point cloud at different scales and observe how topological features appear and disappear.

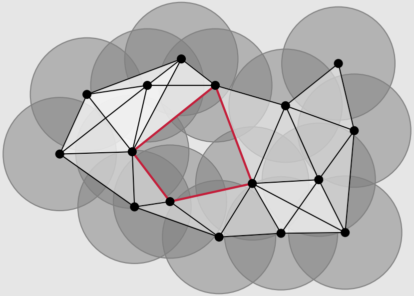

Persistent homology

‘Points cross scales like clouds cross the sky’

Approximate a point cloud at different scales and observe how topological features appear and disappear.

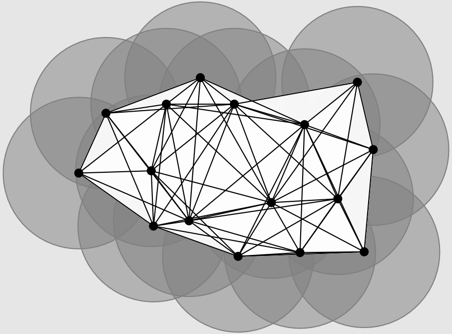

Persistent homology

One formalisation

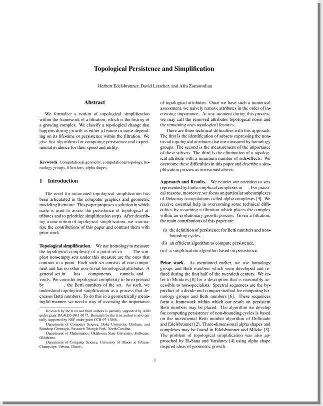

Topological Persistence and Simplification

Herbert Edelsbrunner, David Letscher, and Afra Zomorodian

Abstract

We formalize a notion of topological simplification within the framework of a filtration, which is the history of a growing complex. We classify a topological change that happens during growth as either a feature or noise depending on its life-time or persistence within the filtration. We give fast algorithms for computing persistence and experimental evidence for their speed and utility.

Keywords. Computational geometry, computational topology, homology groups, fi ltrations, alpha shapes.

1 Introduction

The need for automated topological simplification has been articulated in the computer graphics and geometric modeling literature. This paper proposes a solution in which scale is used to assess the persistence of topological attributes and to prioritize simplification steps. After describing a new notion of topological simplification, we summarize the contributions of this paper and contrast them with prior work.

Topological simplification. We use homology to measure the topological complexity of a point set in . The simplest non-empty sets under this measure are the ones that contract to a point. Each such set consists of one component and has no other non-trivial homological attributes. A general set in has components, tunnels, and voids. We consider topological complexity to be expressed by , the Betti numbers of the set. As such, we understand topological simplification as a process that decreases Betti numbers. To do this in a geometrically meaningful manner, we need a way of assessing the importance of topological attributes. Once we have such a numerical assessment, we naively remove attributes in the order of increasing importance. At any moment during this process, we may call the removed attributes topological noise and the remaining ones topological features.

There are three technical difficulties with this approach. The first is the identification of subsets expressing the non-trivial topological attributes that are measured by homology groups. The second is the measurement of the importance of these subsets. The third is the elimination of a topological attribute with a minimum number of side-effects. We overcome these difficulties in this paper and describe a simplification process as envisioned above.

Approach and Results. We restrict our attention to sets represented by finite simplicial complexes in . For practical reasons, moreover, we focus on particular subcomplexes of Delaunay triangulations called alpha complexes [3]. We receive essential help in overcoming some technical difficulties by assuming a filtration which places the complex within an evolutionary growth process. Given a filtration, the main contributions of this paper are:

(i) the definition of persistence for Betti numbers and non-bounding cycles,

(ii) an efficient algorithm to compute persistence,

(iii) a simplification algorithm based on persistence.

Prior work. As mentioned earlier, we use homology groups and Betti numbers which were developed and refined during the first half of the twentieth century. We refer to Munkres [8] for a description that is reasonably accessible to non-specialists. Spectral sequences are the by-product of a divide-and-conquer method for computing homology groups and Betti numbers [6]. These sequences form a framework within which our result on persistent Betti numbers may be placed. The algorithm we develop for computing persistence of non-bounding cycles is based on the incremental Betti number algorithm of Delfinado and Edelsbrunner [2]. Three-dimensional alpha shapes and complexes may be found in Edelsbrunner and Mücke [3]. The problem of topological simplification was also approached by El-Sana and Varshney [4] using alpha shape inspired ideas of geometric growth.

Research by the fi rst and third authors is partially supported by ARO under grant DAAG55-98-1-0177. Research by the fi rst author is also partially supported by NSF under grant CCR-97-12088.

Department of Computer Science, Duke University, Durham, and Raindrop Geomagic, Research Triangle Park, North Carolina.

Department of Mathematics, Oklahoma State University, Stillwater, Oklahoma.

Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois.

1

We formalize a notion of topological simplification within the framework of a filtration, which is the history of a growing complex. We classify a topological change that happens during growth as either a feature or noise depending on its life-time or persistence within the filtration. We give fast algorithms for computing persistence and experimental evidence for their speed and utility.

H. Edelsbrunner, D. Letscher and A. J. Zomorodian, *Topological persistence and simplification*, *Discrete & Computational Geometry* 28.4, 2002, pp. 511–533

Other formulations

On résiste à l'invasion des armées; on ne résiste pas à l'invasion des idées.

(Victor Hugo)

- P. Frosini, *A Distance for Similarity Classes of Submanifolds of a Euclidean Space*, *Bulletin of the Australian Mathematical Society* 42.3, 1990, pp. 407–415
- S. A. Barannikov, *The Framed Morse Complex and its Invariants*, *Advances in Soviet Mathematics* 21, 1994, pp. 93–115
- F. Cagliari, M. Ferri and P. Pozzi, *Size Functions from a Categorical Viewpoint*, *Acta Applicandae Mathematica* 67.3, 2001, pp. 225–235

A generic topology-driven machine learning pipeline

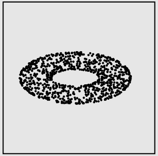

Point cloud

A generic topology-driven machine learning pipeline

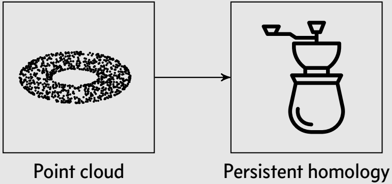

A generic topology-driven machine learning pipeline

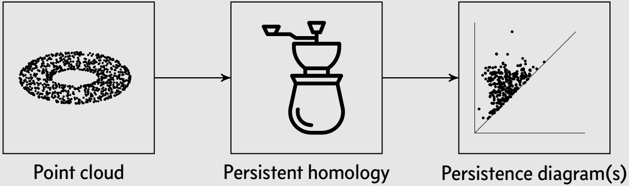

A generic topology-driven machine learning pipeline

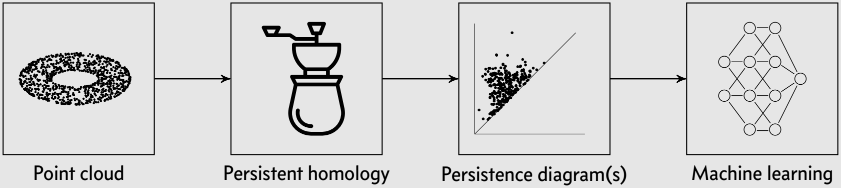

A generic topology-driven machine learning pipeline

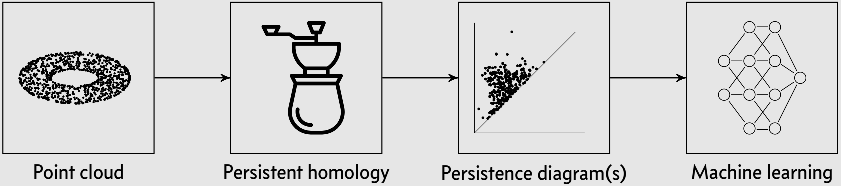

Crucial insight

It is possible to *backpropagate information through this pipeline*, i.e. a gradient exists under certain mild conditions.

- M. Carrière, F. Chazal, M. Glisse, Y. Ike, H. Kannan and Y. Umeda, *Optimizing persistent homology based functions*, Proceedings of the 38th International Conference on Machine Learning (ICML), 2021, pp. 1294–1303
- M. Moor*, M. Horn*, B. Rieck† and K. Borgwardt†, *Topological Autoencoders*, Proceedings of the 37th International Conference on Machine Learning, 2020, pp. 7045–7054
- A. Poulenard, P. Skraba and M. Ovsjanikov, *Topological Function Optimization for Continuous Shape Matching*, *Computer Graphics Forum* 37.5, 2018, pp. 13–25

Some success stories of this pipeline

- Topological attention improves time-series forecasting.[1]
- Topological autoencoders result in *faithful* low-dimensional representations.[2]
- Topological graph neural networks outperform other methods in structure-based graph-classification tasks.[3]
- Topological regularisation helps in 3D shape reconstruction.[4]
- Topological regularisation improves segmentation results.[5]

[1] S. Zeng, F. Graf, C. Hofer and R. Kwitt, *Topological Attention for Time Series Forecasting*, Advances in Neural Information Processing Systems, vol. 34, Curran Associates, Inc., 2021, pp. 24871–24882.

[2] M. Moor*, M. Horn*, B. Rieck† and K. Borgwardt†, *Topological Autoencoders*, Proceedings of the 37th International Conference on Machine Learning, 2020, pp. 7045–7054.

[3] M. Horn*, E. De Brouwer*, M. Moor, Y. Moreau, B. Rieck† and K. Borgwardt†, *Topological Graph Neural Networks*, International Conference on Learning Representations, 2022.

[4] D. J. E. Waibel, S. Atwell, M. Meier, C. Marr and B. Rieck, *Capturing Shape Information with Multi-Scale Topological Loss Terms for 3D Reconstruction*, Medical Image Computing and Computer Assisted Intervention, 2022.

[5] X. Hu, Y. Wang, L. Fuxin, D. Samaras and C. Chen, *Topology-Aware Segmentation Using Discrete Morse Theory*, International Conference on Learning Representations, 2021.

But there’s more than ‘just’ TDA!

A brief attempt at categorising things

TDA, TML, and TDL

- TDA (Topological data analysis): *Describe* data using geometrical-topological features.
- TML (Topological machine learning): Use *any* variant of topology in or for machine-learning algorithms.
- TDL (Topological deep learning): Use *any* variant of topology in or for deep-learning algorithms.

Advantages of this categorisation

In this categorisation, TDL is a subfield of TML, just like DL is a subfield of ML, but TML comprises more approaches, in particular the use of topological techniques as 'hand-crafted features for a machine-learning algorithm.

Examples of topological machine learning

- P. Bubenik, *Statistical Topological Data Analysis using Persistence Landscapes*, *Journal of Machine Learning Research* 16.3, 2015, pp. 77–102
- H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta and L. Ziegelmeier, *Persistence Images: A Stable Vector Representation of Persistent Homology*, *Journal of Machine Learning Research* 18.8, 2017, pp. 1–35
- E. J. Amézquita, M. Y. Quigley, T. Ophelders, E. Munch and D. H. Chitwood, *The shape of things to come: Topological data analysis and biology, from molecules to organisms*, *Developmental Dynamics* 249.7, 2020, pp. 816–833

Examples of topological deep learning

- 2017: Zixuan Cang and Guo-Wei Wei coin the term *topological deep learning* for the first time in a published work: Z. Cang and G.-W. Wei, *TopologyNet: Topology based deep convolutional and multi-task neural networks for biomolecular property predictions*, *PLOS Computational Biology* 13.7, 2017, pp. 1–27.
- 2021: Ephy R. Love et al. propose *topological deep learning* with a focus on convolutional neural networks: E. R. Love, B. Filippenko, V. Maroulas and G. Carlsson, *Topological Deep Learning*, Preprint, 2021, arXiv: `2101.05778 [cs.LG]`.[6]
- 2024: Papamarkou et al. outline the future of *topological deep learning*: T. Papamarkou, T. Birdal, M. Bronstein, G. Carlsson, J. Curry, Y. Gao, M. Hajij, R. Kwitt, P. Liò, P. D. Lorenzo, V. Maroulas, N. Miolane, F. Nasrin, K. N. Ramamurthy, B. Rieck, S. Scardapane, M. T. Schaub, P. Veličković, B. Wang, Y. Wang, G.-W. Wei and G. Zamzmi, *Position: Topological Deep Learning is the New Frontier for Relational Learning*, Proceedings of the 41st International Conference on Machine Learning, Proceedings of Machine Learning Research 235, PMLR, 2024, pp. 39529–39555

[6] This work was published in 2023 under a different name.

Research directions in topological deep learning

In the early stages, the term 'TDL' was often used to refer to the incorporation of features generated by persistent homology within the input pipeline of a deep neural network. […] However, this paper uses the term 'TDL' to refer to the collection of ideas and methods related to the use of topological concepts in deep learning.

- Point-set topology: Extend existing deep-learning paradigms to topological spaces (*neighbourhoods*).
- Algebraic topology: Improve deep-learning models by incorporating novel geometrical-topological inductive biases (*homology*).
- Differential topology: Study topological spaces from a function-space perspective (*cohomology*).

Example of point-set topology in TDL

C. Bodnar, F. Frasca, Y. Wang, N. Otter, G. F. Montufar, P. Lió and M. Bronstein, *Weisfeiler and Lehman Go Topological: Message Passing Simplicial Networks*, Proceedings of the 38th International Conference on Machine Learning, vol. 139, 2021, pp. 1026–1037

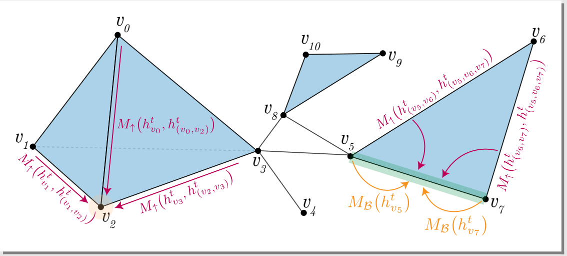

Example of operators in TDL

S. Ebli, M. Defferrard and G. Spreemann, *Simplicial Neural Networks*, ‘Topological Data Analysis and Beyond’ Workshop at NeurIPS, 2020, URL: `https://openreview.net/forum?id=nPCt39DVIfk`

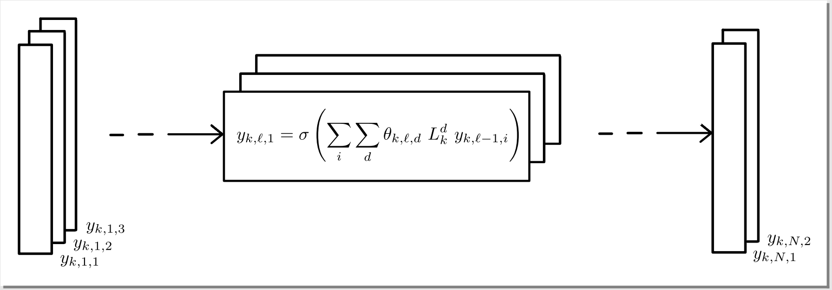

Example of topological features in TDL

M. Horn*, E. De Brouwer*, M. Moor, Y. Moreau, B. Rieck† and K. Borgwardt†, *Topological Graph Neural Networks*, International Conference on Learning Representations, 2022

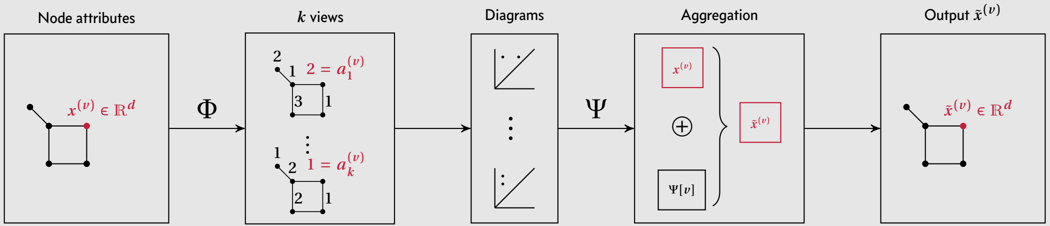

- Use a node map $\Phi\colon \mathbb{R}^d \to \mathbb{R}^k$ to create k different filtrations of the graph.
- Use a coordinatisation function Ψ to create *compatible* representations of the node attributes.

Example of smooth geometrical-topological features in TDL

K. Maggs, C. Hacker and B. Rieck, *Simplicial Representation Learning with Neural k-forms*, International Conference on Learning Representations, 2024

$$\underbrace{\Omega^k(\mathbb{R}^n) \xrightarrow{\int} C^k_{\text{sing}}(\mathbb{R}^n; \mathbb{R})}_{\text{de Rham map}} \underbrace{\xrightarrow{\phi^*} C^k_{\text{simp}}(S; \mathbb{R})}_{\text{Restriction map}}$$

Dramatis personæ

- $\Omega^k(\mathbb{R}^n)$: k-forms on $\mathbb{R}^n$
- $C^k_{\text{sing}}(\mathbb{R}^n; \mathbb{R})$: singular cochains
- $C^k_{\text{simp}}(S; \mathbb{R})$: simplicial cochains
- $\phi^*\gamma(\sigma) = \gamma(\phi|_\sigma)$ for a singular cochain $\gamma \in C^k_{\text{sing}}(\mathbb{R}^n; \mathbb{R})$

Slogan

We replace the evaluation of feature cochains by the *integration* of differential feature forms.

But what about geometry?

Geometry and topology are *dual*

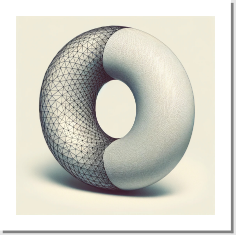

Geometry	=	fine details	+	quantitative answers
Topology	=	fundamental properties	+	qualitative answers

Geometry and topology are *dual*

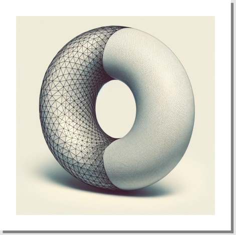

Geometry	=	fine details	+	quantitative answers
Topology	=	fundamental properties	+	qualitative answers

Data has shape, shape has meaning, and meaning begets understanding.

(Gunnar Carlsson, paraphrased)

What is curvature?

Motivation

Characterise how ‘curved’ an object (a surface, a manifold, a topological space, …) is.
Curvature can be *extrinsic* or *intrinsic*.

What is curvature?

Motivation

Characterise how ‘curved’ an object (a surface, a manifold, a topological space, …) is. Curvature can be *extrinsic* or *intrinsic*.

Gaussian curvature

Gaussian curvature K is the product of the *principal curvatures* κ_1, κ_2. It is an intrinsic property of a surface and does not depend on a specific embedding.

What is curvature?

Motivation

Characterise how 'curved' an object (a surface, a manifold, a topological space, …) is. Curvature can be *extrinsic* or *intrinsic*.

Gaussian curvature

Gaussian curvature K is the product of the *principal curvatures* κ_1, κ_2. It is an intrinsic property of a surface and does not depend on a specific embedding.

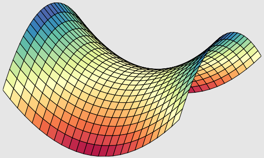

$K < 0$

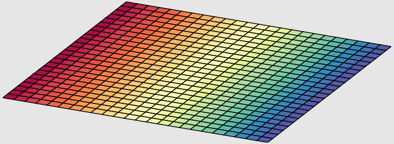

$K = 0$

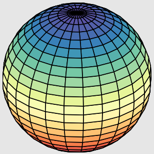

$K > 0$

Ollivier–Ricci curvature

Let G be a graph with its shortest-path metric d and μ_v be a probability measure on G for node $v \in V$. The *Ollivier–Ricci curvature* of a pair of nodes $i \neq j \in V$ is then defined as

$$\kappa_{\mathsf{OR}}(i, j) := 1 - \frac{\mathrm{W}_1(\mu_i, \mu_j)}{\mathrm{d}(i, j)}, \tag{1}$$

where W_1 refers to the first *Wasserstein distance* between μ_i and μ_j.

History

First introduced by Ollivier for metric (measure) spaces, this notion of curvature was quickly adopted for use in the graph setting.

Y. Ollivier, *Ricci curvature of Markov chains on metric spaces*, *Journal of Functional Analysis* 256.3, 2009, pp. 810–864

How to pick μ_i?

It is common practice to define a version of μ_i based on lazy random walks. Given a laziness parameter $\alpha \in [0, 1]$, we set

$$\mu_i(j) := \begin{cases} \alpha & \text{if } i = j \\ \frac{1-\alpha}{\deg(i)} & \text{if } i \neq j \text{ and } i \sim j \\ 0 & \text{otherwise} \end{cases}, \tag{2}$$

where $\deg(i)$ refers to the degree of node i.

Ollivier–Ricci curvature

Different regimes

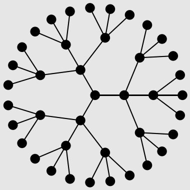

$\kappa_{\mathrm{OR}} < 0$

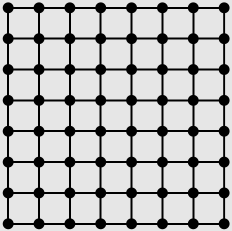

$\kappa_{\mathrm{OR}} = 0$

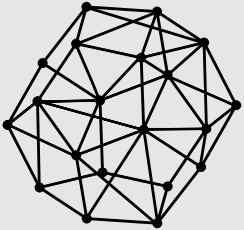

$\kappa_{\mathrm{OR}} > 0$

(figure inspired by K. Devriendt and R. Lambiotte, *Discrete curvature on graphs from the effective resistance*, *Journal of Physics: Complexity* 3.2, 2022, p. 025008)

Ollivier–Ricci curvature

Examples

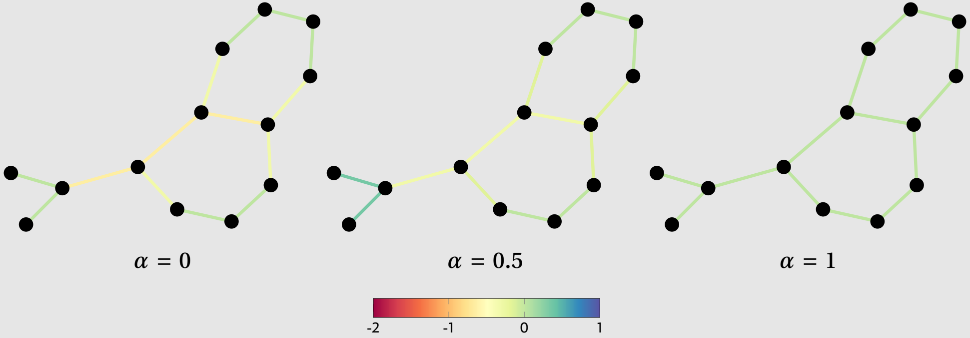

Useful properties of κ_{OR}

Lower bound

It is sufficient to know the values of κ_{OR} for each edge (i, j). If $\kappa_{\text{OR}}(i, j) \geq K$ for edges $(i, j) \in E$, then $\kappa_{\text{OR}}(k, l) \geq K$ for *all* pairs of vertices (k, l).

Curvature characterises graphs

If $\kappa_{\text{OR}}(i, j) \geq K > 0$ for all edges $(i, j) \in E$, then for $i, j \in V$, we have

$$\mathrm{d}(i, j) \leq \frac{\mathrm{W}_1(\delta_i, \mu_i) + \mathrm{W}_1(\delta_j, \mu_j)}{\kappa_{\text{OR}}(i, j)}, \tag{3}$$

where δ_i, δ_j refer to Dirac probability measures centred at node i and j, respectively. As a direct consequence, we obtain a *diameter bound* via

$$\mathrm{diam}(G) \leq \frac{\sup_i \mathrm{W}_1(\delta_i, \mu_i)}{K}. \tag{4}$$

How to leverage this in theory?

Success rate (↑) of distinguishing pairs of graphs in the 'BREC' data set when using different probability measures in the OR curvature calculation.

Method		Basic (56)	Regular (50)	STR (50)	Extension (97)	CFI (97)
1-WL		0.00	0.00	0.00	0.00	0.00
3-WL		1.00	1.00	0.00	1.00	0.59
κ_{OR}	$\mu_i^{(1)}$	1.00	0.96	0.06	0.87	0.00
	$\mu_i^{(2)}$	1.00	1.00	0.14	0.97	0.01
	$\mu_i^{(3)}$	1.00	1.00	0.14	0.99	0.04
	$\mu_i^{(4)}$	1.00	1.00	0.14	1.00	0.09
	$\mu_i^{(5)}$	1.00	1.00	0.14	1.00	0.19

Ollivier–Ricci curvature with *learnable* probability measures promises to lead to structural insights.

How to leverage this in practice?

Use curvature to evaluate *graph generative models*.
J. Southern*, J. Wayland*, M. Bronstein and B. Rieck, *Curvature Filtrations for Graph Generative Model Evaluation*, Advances in Neural Information Processing Systems, vol. 36, Curran Associates, Inc., 2023, pp. 63036–63061, arXiv: `2301.12906 [cs.LG]`

Use curvature as an efficient *graph descriptor* for graph learning tasks.
L. O'Bray*, B. Rieck* and K. Borgwardt, *Filtration Curves for Graph Representation*, Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, New York, NY, USA: Association for Computing Machinery, 2021, pp. 1267–1275

Back to TDL now!

Challenges in topological deep learning

Y. Eitan, Y. Gelberg, G. Bar-Shalom, F. Frasca, M. M. Bronstein and H. Maron, *Topological Blindspots: Understanding and Extending Topological Deep Learning Through the Lens of Expressivity*, International Conference on Learning Representations, 2025, URL: `https://openreview.net/forum?id=EzjsoomYEb`

> *We then use this criterion to prove HOMP's [higher-order message passing] inability to differentiate between complexes based on several fundamental topological and metric invariants, including diameter, orientability, planarity, and homology groups. These limitations are particularly noteworthy, as TDL's main goal is to leverage topological structure in data.*

Challenges in topological deep learning

Y. Eitan, Y. Gelberg, G. Bar-Shalom, F. Frasca, M. M. Bronstein and H. Maron, *Topological Blindspots: Understanding and Extending Topological Deep Learning Through the Lens of Expressivity*, International Conference on Learning Representations, 2025, URL: `https://openreview.net/forum?id=EzjsoomYEb`

> *We then use this criterion to prove HOMP's [higher-order message passing] inability to differentiate between complexes based on several fundamental topological and metric invariants, including diameter, orientability, planarity, and homology groups. These limitations are particularly noteworthy, as TDL's main goal is to leverage topological structure in data.*

Lesson

TDL $\neq$ HOMP or TDL $\supset$ HOMP

Implication for our terminology

Instead of referring to them as *topological neural networks*, we should refer to HOMP models as *combinatorial complex neural networks*.

New data set

R. Ballester*, E. Röell*, D. B. Schmid*, M. Alain*, S. Escalera, C. Casacuberta and B. Rieck, *MANTRA: The Manifold Triangulations Assemblage*, International Conference on Learning Representations, 2025, URL: `https://openreview.net/forum?id=X6y5CC44HM`

A noteworthy aspect of MANTRA is the conspicuous absence of any intrinsic vertex or edge features such as coordinates or signals. We argue that this absence renders tasks more topological, as models can only rely on topology, instead of non-topological information contained in features. Moreover, as manifold triangulations are directly related to the topological structure of the underlying manifold, we study to which extent higher-order models are invariant to triangulation transformations that preserve the topological structure of the associated manifold.

New data set

Some results

		Accuracy				*AUROC*	
Dataset	Model Family	β_0	β_1	β_2	β_3	Homeo. Type	Orientability
$2\text{-}\mathcal{M}^0$	$\mathcal{G}$	**1.00 ± 0.00**	0.50 ± 0.00	0.50 ± 0.00		0.47 ± 0.01	0.50 ± 0.00
	$\mathcal{T}$	0.73 ± 0.39	**0.68 ± 0.16**	**0.59 ± 0.10**		**0.69 ± 0.18**	**0.56 ± 0.07**
$2\text{-}\mathcal{M}^0_H$	$\mathcal{G}$	**1.00 ± 0.00**	0.21 ± 0.00	0.50 ± 0.00		0.49 ± 0.01	0.50 ± 0.00
	$\mathcal{T}$	0.57 ± 0.44	**0.25 ± 0.03**	**0.52 ± 0.02**		**0.66 ± 0.13**	**0.52 ± 0.02**
$2\text{-}\mathcal{M}^1_H$	$\mathcal{G}$	**0.47 ± 0.51**	0.22 ± 0.00	0.50 ± 0.00		0.49 ± 0.04	0.50 ± 0.00
	$\mathcal{T}$	0.21 ± 0.38	**0.25 ± 0.02**	**0.51 ± 0.01**		**0.60 ± 0.10**	**0.51 ± 0.01**
$3\text{-}\mathcal{M}^0$	$\mathcal{G}$	**1.00 ± 0.00**	0.23 ± 0.00	0.12 ± 0.00	0.14 ± 0.00		0.14 ± 0.00
	$\mathcal{T}$	0.78 ± 0.41	**0.25 ± 0.04**	**0.13 ± 0.03**	**0.16 ± 0.03**		**0.15 ± 0.02**

In a nutshell

'Topological' models show only slight improvements over much simpler graph-based models. Additional inductive biases are needed!

Conclusion

- Topological deep learning is an exciting nascent field.

Conclusion

- Topological deep learning is an exciting nascent field.
- The use of geometrical-topological concepts provides a novel perspective for many machine-learning applications…

Conclusion

- Topological deep learning is an exciting nascent field.
- The use of geometrical-topological concepts provides a novel perspective for many machine-learning applications…
- …but the use of geometrical and topological features should be *justified* and carefully assessed.

Conclusion

- Topological deep learning is an exciting nascent field.
- The use of geometrical-topological concepts provides a novel perspective for many machine-learning applications…
- …but the use of geometrical and topological features should be *justified* and carefully assessed.
- The best time to join is *now*!

Conclusion

- Topological deep learning is an exciting nascent field.
- The use of geometrical-topological concepts provides a novel perspective for many machine-learning applications…
- …but the use of geometrical and topological features should be *justified* and carefully assessed.
- The best time to join is *now*!

AIDOS Lab

Our research