HELMHOLTZ MUNICH
AIH Institute of AI for Health

A Good Scale is Hard To Find

Shape Analysis Using Topology

Bastian Rieck (@Pseudomanifold)

Getting Acquainted

What is computational topology?

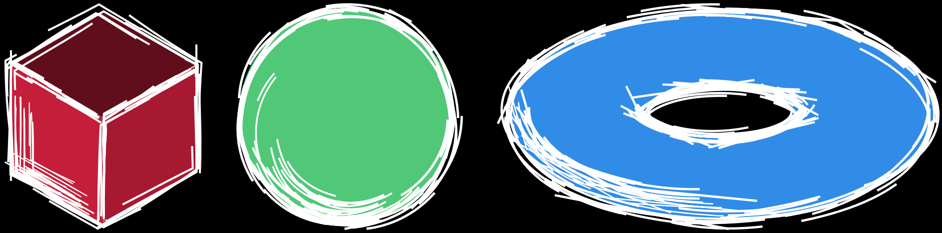

What is computational topology?

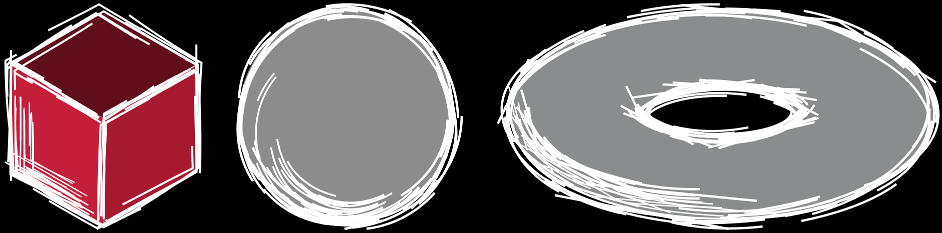

What is computational topology?

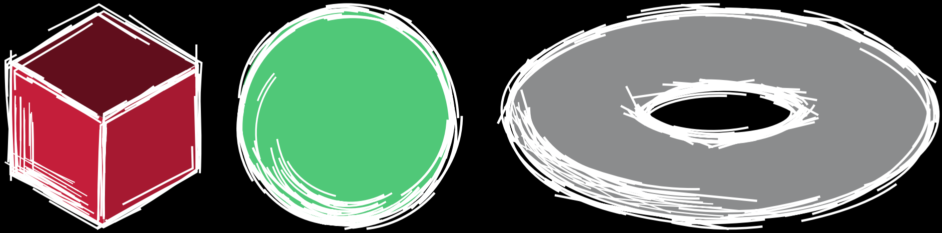

What is computational topology?

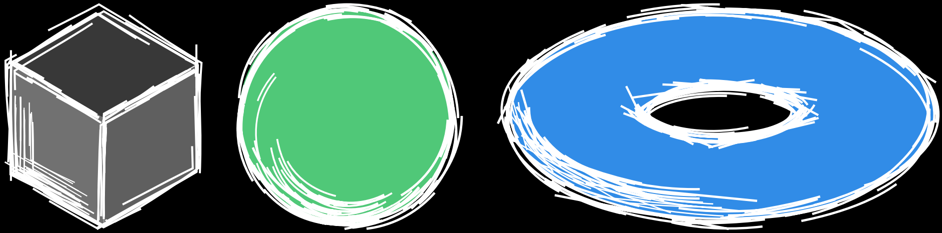

Why is a sphere *not* the same as a torus?

Betti numbers

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

β_0	Connected components
β_1	Tunnels
β_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

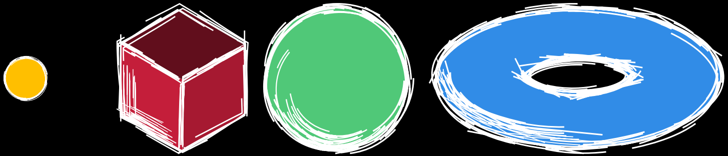

How to handle real-world data?

Calculating simplicial complexes from data

Vietoris–Rips complex

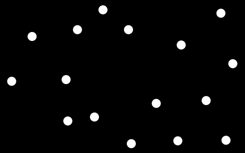

$$\mathcal{V}_\epsilon := \{\{x_1, x_2, \ldots\} \mid \operatorname{dist}\left(x_i, x_j\right) \leq \epsilon \text{ for all } i \neq j\}$$

Calculating simplicial complexes from data

Vietoris–Rips complex

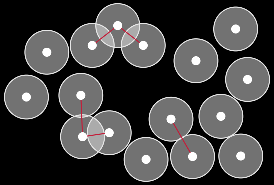

$$\mathcal{V}_\epsilon := \{\{x_1, x_2, \ldots\} \mid \operatorname{dist}\left(x_i, x_j\right) \leq \epsilon \text{ for all } i \neq j\}$$

Calculating simplicial complexes from data

Vietoris–Rips complex

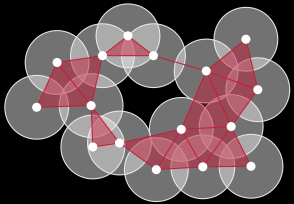

$$\mathcal{V}_\epsilon := \{\{x_1, x_2, \ldots\} \mid \mathrm{dist}(x_i, x_j) \leq \epsilon \text{ for all } i \neq j\}$$

Calculating simplicial complexes from data

Vietoris–Rips complex

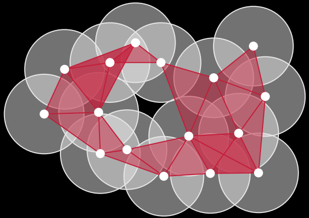

$$\mathcal{V}_\epsilon := \{\{x_1, x_2, \ldots\} \mid \operatorname{dist}\left(x_i, x_j\right) \leq \epsilon \text{ for all } i \neq j\}$$

Calculating simplicial complexes from data

Vietoris–Rips complex

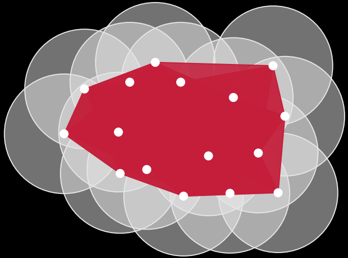

$$\mathcal{V}_\epsilon := \{\{x_1, x_2, \dots\} \mid \mathrm{dist}(x_i, x_j) \leq \epsilon \text{ for all } i \neq j\}$$

How to pick ϵ?

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

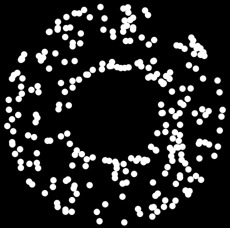

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on ‘when’ it was created and ‘when’ it was destroyed. Store these features in a *persistence diagram*.

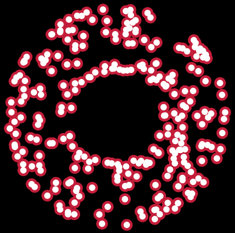

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

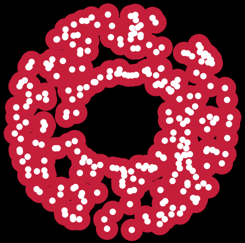

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

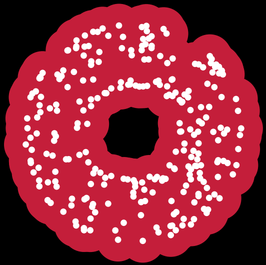

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

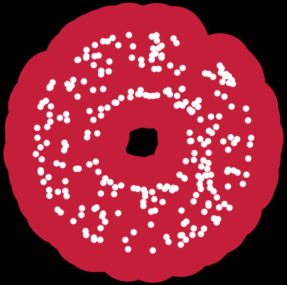

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

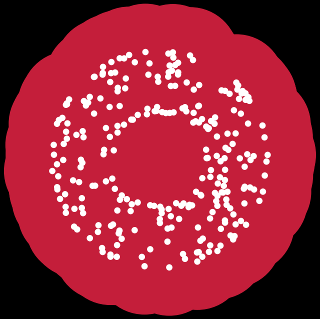

Persistent homology

Calculate simplicial complexes for *every* value of ϵ, while watching how topological features change. Assign each feature a duration, depending on 'when' it was created and 'when' it was destroyed. Store these features in a *persistence diagram*.

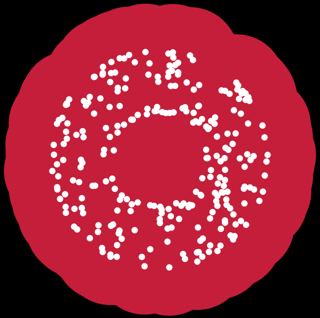

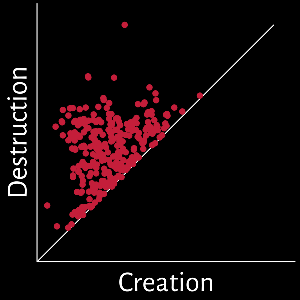

Distances between persistence diagrams

Bottleneck distance

Given two persistence diagrams $\mathscr{D}$ and $\mathscr{D}'$, their *bottleneck* distance is defined as

$$\mathrm{W}_\infty(\mathscr{D}, \mathscr{D}') := \inf_{\eta\colon \mathscr{D} \to \mathscr{D}'} \sup_{x \in \mathscr{D}} \|x - \eta(x)\|_\infty,$$

where $\eta\colon \mathscr{D} \to \mathscr{D}'$ denotes a bijection between the point sets of $\mathscr{D}$ and $\mathscr{D}'$ and $\|\cdot\|_\infty$ refers to the L_∞ distance between two points in $\mathbb{R}^2$.

Wasserstein distance

$$\mathrm{W}_p(\mathscr{D}_1, \mathscr{D}_2) := \left(\inf_{\eta\colon \mathscr{D}_1 \to \mathscr{D}_2} \sum_{x \in \mathscr{D}_1} \|x - \eta(x)\|_\infty^p \right)^{\frac{1}{p}}$$

Differences between the two distances

Bottleneck distance

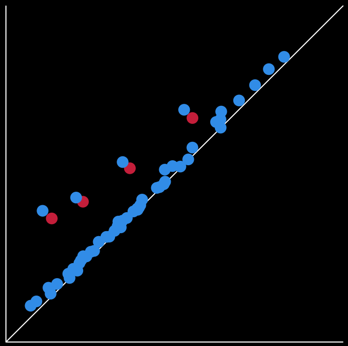

Wasserstein distance

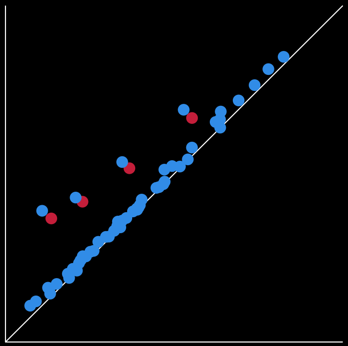

Differences between the two distances

Bottleneck distance

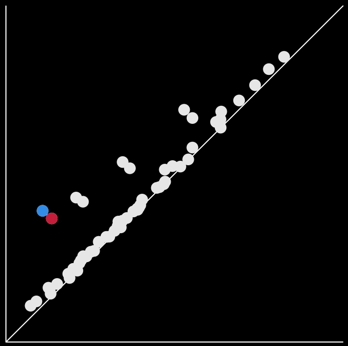

Wasserstein distance

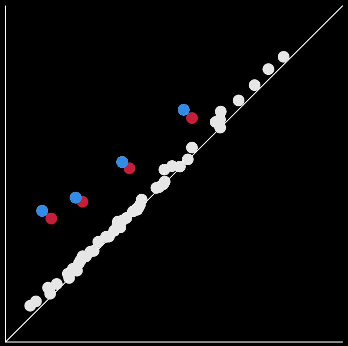

Persistence images

Multi-scale descriptors

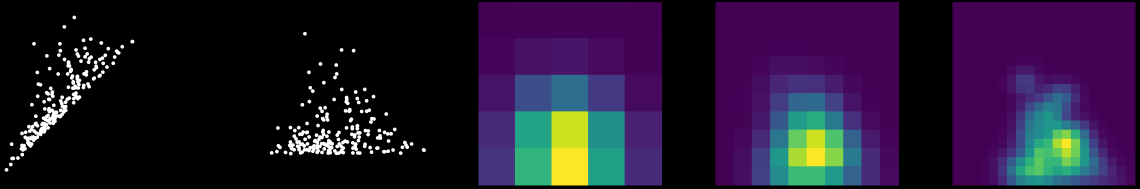

Algorithm

Use $\Psi\colon \mathbb{R}^2 \to \mathbb{R}$ to turn a diagram $\mathscr{D}$ into a surface via $\Psi(z) := \sum_{x,y\in\mathscr{D}} \mathrm{w}(x,y)\Phi(x,y,z)$, where $\mathrm{w}(\cdot)$ is a fixed piecewise linear weight function and $\Phi(\cdot)$ denotes a probability distribution, which is typically chosen to be a normalised symmetric Gaussian. By discretising Ψ (using an $r \times r$ grid), a persistence diagram is transformed into a *persistence image*.

Publication

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*, 2017.

Part I: Analysing the shape of fMRI data

fMRI data

Our approach

- Consider the BOLD activation function f to be a time-varying function on a manifold $\mathbb{M}$
- Calculate topological features of $\mathbb{M}$ 'measured' via f
- Obtain stable topological summaries at different resolutions

Main advantage of this approach

Working on the 'raw' data; no auxiliary representations necessary! In particular, no *atlas* required (fewer modelling choices in total).

Publication

B. Rieck*, T. Yates*, C. Bock, K. Borgwardt, G. Wolf, N. Turk-Browne† and S. Krishnaswamy†, 'Uncovering the Topology of Time-Varying fMRI Data using Cubical Persistence', *Advances in Neural Information Processing Systems (NeurIPS)*, 2020, arXiv: `2006.07882 [q-bio.NC]`.

Working with volume data

A volume is a special type of topological complex, a *cubical complex*. With minor modifications, persistent homology works in this setting as well, and a provided likelihood function f directly leads to a filtration.

Publication

See H. Wagner, C. Chen and E. Vuçini, 'Efficient Computation of Persistent Homology for Cubical Data', *Topological Methods in Data Analysis and Visualization II: Theory, Algorithms, and Applications*, 2012. for more information.

Intuition

Workflow

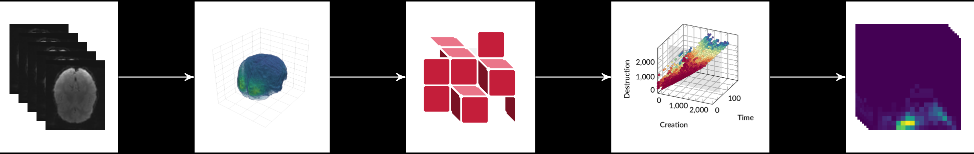

fMRI stack → fMRI volume → Cubical complex → Persistence diagram → Persistence images

Our data set

- 155 (122 children, 33 adults) participants are being shown the film ‘Partly Cloudy‘
- *Continuous* stimulation of participants
- *No additional information about participants has been provided on purpose*

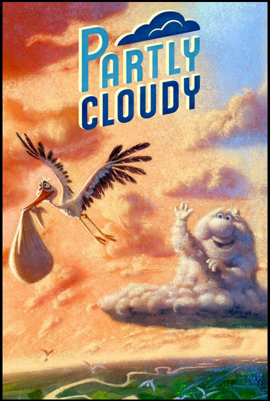

Summary statistics of a persistence diagram

Norms of a persistence diagram

$$\|\mathscr{D}\|_\infty := \max_{x,y\in\mathscr{D}} \operatorname{pers}(x,y)^p \quad \text{and} \quad \|\mathscr{D}\|_p := \sqrt[p]{\sum_{x,y\in\mathscr{D}} \operatorname{pers}(x,y)^p},$$

These norms are stable and highly useful in obtaining simple descriptions of time-varying persistence diagrams!

Age prediction based on summary statistics

Method	BM	OM	XM
`baseline-tt`	0.09	0.02	0.24
`baseline-pp`	0.41	0.40	0.40
`tt-corr-tda`	0.17	0.11	0.23
`pp-corr-tda`	0.25	0.27	0.23
`srm`	0.44		
$\|\mathscr{D}\|_1$	0.46	0.67	0.48
$\|\mathscr{D}\|_\infty$	**0.61**	**0.77**	**0.73**

Brain state trajectories

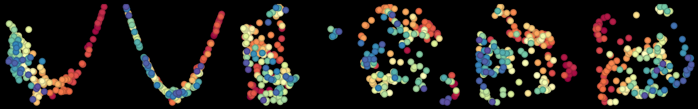

Whole-brain mask

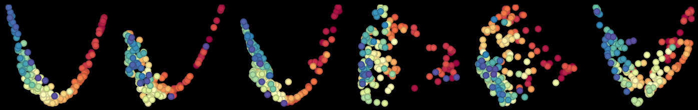

Occipital-temporal mask

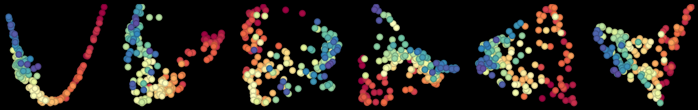

XOR mask

Part II: Predicting the shape of cells

Cell shape prediction

- Use confocal fluorescence microscopy to obtain images of cells.
- What is the 3D shape of a cell?
- Morphological analysis is crucial for certain pathologies!

  *When used properly, RBC [red blood cell] morphology can be a **key tool** for laboratory hematology professionals to recommend appropriate clinical and laboratory follow-up and to select the best tests for definitive diagnosis.* *(J. Ford, 'Red blood cell morphology', International Journal of Laboratory Hematology, 2013.)*

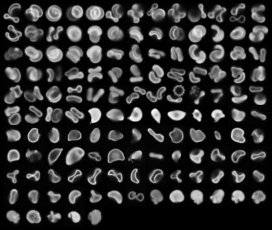

SHAPR

Overview

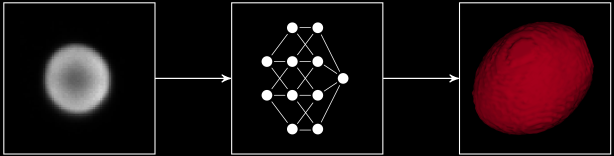

Input image Model Output image

We are trying to solve a complicated *inverse problem*, going from 2D to 3D. This is an ill-defined problem with a large number of potential solutions.

SHAPR

Architecture

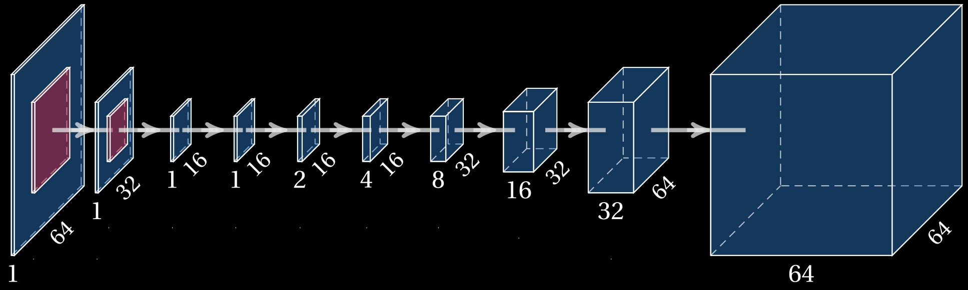

We are learning a *likelihood function* $f : \mathbb{R}^3 \to \mathbb{R}$. Formally, f ‘lives’ on a voxel grid, assigning each voxel x a value that indicates the likelihood of x being part of the ‘true’ volume.

SHAPR

Loss function

$$\mathscr{L}_{\mathsf{G}}(f, f') := \frac{2\,\mathscr{L}_{\mathsf{Dice}}(f, f') + \mathscr{L}_{\mathsf{BCE}(f, f')}}{2}$$

$$\mathscr{L}_{\mathsf{Dice}}(f, f') := \frac{2\left|\mathrm{Vol}_f \cap \mathrm{Vol}_{f'}\right|}{\left|\mathrm{Vol}_f\right| + \left|\mathrm{Vol}_{f'}\right|} = \frac{2\mathsf{TP}}{2\mathsf{TP} + \mathsf{FP} + \mathsf{FN}}$$

Intuition

Compare *geometry* of the resulting volumes on a per-voxel basis. Is the reconstructed volume well-aligned with the ground truth one?

SHAPR goes topological

$$\mathscr{L}_{\mathsf{T}}(f, f', q) := \sum_{i=0}^{d} \mathrm{W}_q\left(\mathscr{D}_f^{(i)}, \mathscr{D}_{f'}^{(i)}\right) + \mathrm{pers}\left(\mathscr{D}_{f'}^{(i)}\right)$$

SHAPR goes topological

$$\mathscr{L}_{\mathsf{T}}(f, f', q) := \sum_{i=0}^{d} \mathrm{W}_q\left(\mathscr{D}_f^{(i)}, \mathscr{D}_{f'}^{(i)}\right) + \mathrm{pers}\left(\mathscr{D}_{f'}^{(i)}\right)$$

Loss components

- Aligning the ground truth likelihood f and the predicted likelihood function f'.
- Reducing the geometrical–topological variation of the predicted likelihood function f'.

SHAPR goes topological

$$\mathscr{L}_{\mathsf{T}}(f, f', q) := \sum_{i=0}^{d} \mathrm{W}_q\left(\mathscr{D}_f^{(i)}, \mathscr{D}_{f'}^{(i)}\right) + \mathrm{pers}\left(\mathscr{D}_{f'}^{(i)}\right)$$

Loss components

- Aligning the ground truth likelihood f and the predicted likelihood function f'.
- Reducing the geometrical–topological variation of the predicted likelihood function f'.

We obtain a *combined loss* by choosing $\lambda \in \mathbb{R}_{>0}$ and calculating:

$$\mathscr{L} := \mathscr{L}_{\mathsf{G}} + \lambda \mathscr{L}_{\mathsf{T}}$$

Quantitative results

		Red blood cell ($n = 825$)		Nuclei ($n = 887$)	
Metric	$\mathscr{L}_T$	Median	$\mu \pm \sigma$	Median	$\mu \pm \sigma$
1-IoU	✗	0.48	0.49 ± 0.12	0.62	0.62 ± 0.11
	✓	**0.46**	**0.47 ± 0.10**	**0.61**	**0.61 ± 0.11**
Volume	✗	0.31	0.35 ± 0.31	0.34	0.48 ± 0.47
	✓	**0.21**	**0.25 ± 0.24**	**0.32**	**0.43 ± 0.42**
Surface area	✗	0.20	0.24 ± 0.20	0.21	0.27 ± 0.25
	✓	**0.13**	**0.18 ± 0.16**	**0.18**	**0.25 ± 0.24**
Surface roughness	✗	0.35	0.36 ± 0.24	**0.17**	**0.18 ± 0.12**
	✓	**0.24**	**0.29 ± 0.22**	0.18	0.19 ± 0.13

Summary

- Topology can provide useful inductive biases for shape reconstruction tasks.
- Persistence diagrams encode geometrical *and* topological properties of data.
- Integration into 'standard' machine learning models is possible!

Publications

- F. Hensel, M. Moor and **B. Rieck**, 'A Survey of Topological Machine Learning Methods', *Frontiers in Artificial Intelligence*, 2021.
- 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 (MICCAI)*, 2022, arXiv: `2203.01703 [cs.CV]`, in press.

Software

`https://github.com/aidos-lab/pytorch-topological`

❤ Acknowledgements

My co-authors, in particular Carsten, Dominik, Felix, and Michael.