Foundations of Computational Geometry and Topology
Poster Session Abstracts
May 20, 2026
4:00-5:00pm

Computing the Intrinsic Delaunay Triangulation of a Closed Polyhedral Surface

Loïc Dubois, Notre Dame University

Every surface that is intrinsically polyhedral can be represented by a portalgon: a collection of polygons in the Euclidean plane with some pairs of equally long edges abstractly identified. While this representation is arguably simpler than meshes (flat polygons in R3 forming a surface), it has unbounded happiness: a shortest path in the surface may visit the same polygon arbitrarily many times. This pathological behavior is an obstacle towards efficient algorithms. On the other hand, Löffler, Ophelders, Staals, and Silveira (SoCG 2023) recently proved that the (intrinsic) Delaunay triangulations have bounded happiness. In this paper, given a closed polyhedral surface S, represented by a triangular portalgon T, we provide an algorithm to compute the Delaunay triangulation of S whose vertices are the singularities of S (the points whose surrounding angle is distinct from 2pi). The time complexity of our algorithm is polynomial in the number of triangles and in the logarithm of the aspect ratio r of T. Within our model of computation, we show that the dependency in log(r) is unavoidable. Our algorithm can be used to pre-process a triangular portalgon before computing shortest paths on its surface, and to determine whether the surfaces of two triangular portalgons are isometric.

Topology-Aware Graph Diffusion: Harnessing Discrete Morse Theory

Jennifer Rozenblit, University of Texas, Austin

Denoising diffusion models have driven recent advances in generative modeling but struggle when transferred to graphs; the data becomes discrete, the forward process destroys global topological invariants, and standard edge-factorized denoisers supply no direct supervisory signal for the higher-order structure that distinguishes plausible graphs from noise. We address this by conditioning graph diffusion and graph generation on a topological invariant extracted via discrete Morse theory, a tool from low-dimensional topology whose generative modeling potential remains largely unexplored. Given a graph G, a degree-based vertex potential induces a discrete Morse function whose gradient pairing partitions the edge set into a topology-aware decomposition. We feed the two views derived from this decomposition into a mixture-of-experts conditioning module that soft-routes between them based on context, and use the fused representation as the only graph-aware signal seen by the denoiser. On probabilistic spatio-temporal forecasting the resulting model attains the best CRPS on every benchmark, while on graph generation it achieves the lowest Betti and Wasserstein errors among strong baselines. We view this as a first step in a broader program of using discrete Morse theory graph machine learning.

TopoCAM: ROI-Driven Topological Signatures in 3D Medical Imaging

Brighton Nuwagira, University of Texas at Dallas

Accurate classification of 3D medical images is challenging due to the high dimensionality of volumetric data and the scarcity of well-annotated clinical datasets. We propose a hybrid framework that couples explainable deep learning with topological data analysis (TDA). First, we compute layer-weighted Grad-CAM across multiple network layers, upsample and normalize the maps to the input grid, and threshold them to produce a binary region-of-interest (ROI) mask. We then apply this mask to the input volume to obtain a segmented image that suppresses irrelevant anatomy

while preserving clinically salient structures. Within these attention-derived ROIs and segmented images, we compute cubical persistent homology to derive compact topological descriptors that capture diagnostically meaningful features. Across both 3D volumes and 2D medical imaging benchmarks, this segmentation-guided TDA pipeline surpasses strong 3D CNN and Transformer baselines, yielding higher accuracy and improved robustness in limited-data settings while providing localized, interpretable evidence for clinical decision support.

Multiparameter Persistent Homology for Spatial Transcriptomics

Kylie Savoye, University of Birmingham

Spatial transcriptomics enables measurement of gene expression while preserving tissue architecture but extracting structure from joint spatial–molecular organisation remains challenging. We apply Multi-Parameter Persistent Homology (MPH), a topological data analysis framework based on multiparameter persistence modules, to identify joint tissue and gene structures not captured by conventional methods. MPH constructs a two-parameter filtration by combining spatial proximity with gene expression thresholds to track emergence and persistence of topological features across scales. As complete discrete invariants are not available in the multiparameter setting, we instead use computable module summaries to enable stable comparison. This work highlights MPH as a practical framework for analysing structured biological data, illustrated on colorectal cancer tissue, and motivates further development of multiparameter methods for interpretable biological applications.

Unlinking, Unknotting, and Extrinsic Topology in Deep Learning

Junyu Ren, University of Chicago

Under the manifold hypothesis, data classes lie near submanifolds embedded in an ambient space. We study such data through *extrinsic* topology: classification of linked class manifolds becomes an unlinking problem, while representation learning for a knotted object becomes an unknotting problem. We prove that width-3 feedforward networks with coordinate-wise monotonic activations preserve linking number in $\mathbb{R}^3$ and therefore cannot separate Hopf-linked data regardless of depth. Related arguments yield unknotting obstructions for point-separating embeddings, higher-dimensional obstructions for linked manifolds $M^m, N^n \subset \mathbb{R}^{m+n+1}$, and multi-component obstructions detected by Milnor's $\bar\mu$-invariants. We then identify the geometric mechanisms by which modern architectures escape these barriers: non-monotonic activations break the homotopy argument, ResNet skip connections synthesize coordinate-wise folds such as $|x| = x + 2\text{ReLU}(-x)$, and attention yields input-dependent non-monotonic transformations. Complementing the theory, we give a Gauss-integral algorithm for detecting linking in point clouds via spatial graphs and cycle bases, and illustrate it on synthetic links and CIFAR-10 witness cycles. The result is a computational and theoretical perspective on deep learning grounded in classical extrinsic low-dimensional topology, complementary to intrinsic TDA.

A theoretical framework for chromatic TDA: the constrained Gromov-Hausdorff distance

Nicolò Zava, Institute of Science and Technology Austria (ISTA)

Chromatic TDA focuses on developing invariants of point clouds where points have additional features or information attached, often represented as colours. This line of research is receiving growing attention due to its applications in biological image analysis. In this context, the need for a unifying framework to study the stability of those invariants emerges. In this poster, I present the constrained Gromov-Hausdorff distance, modelled after the usual Gromov-Hausdorff distance to include colours and hierarchies among them. Furthermore, we rephrase known stability results, Dowker persistence diagrams, and present new ones, six-packs of persistence diagrams, with respect to this distance. This poster is based on a joint project with Sophie Rosenmeier and Ondřej Draganov (arXiv:2507.17994).

From Frames to Features: Fast Zigzag Persistence for Binary Videos

David Lanners, Durham University

Zigzag persistence is a powerful tool for tracking topological changes in time-dependent data, such as video streams, but its practical use has been limited by significant computational and memory costs. In this work, we present a reformulation of zigzag persistence for binary image sequences as a graph problem. This perspective allows us to exploit the near-linear time algorithm of Tamal K. Dey and Yusu Wang Hou, dramatically improving efficiency. By further incorporating Alexander duality, we obtain both connected components (H_0) and loop structures (H_1) at essentially the same computational cost. This unified approach significantly accelerates the extraction of homological features from dynamic data. Our method brings zigzag persistence closer to real-time applicability, opening new possibilities for analyzing evolving systems. As a motivating direction, we highlight potential applications to pattern-forming partial differential equations, such as the Gray-Scott model, where fast topological summaries can provide insight into complex spatiotemporal behavior.

Topoformer: Topology-Infused Transformer for Medical Imaging

Sayoni Chakraborty, The University of Texas at Dallas

Deep learning has transformed 2D medical imaging, but scaling to 3D volumes remains difficult because of high computational cost, limited annotations, and loss of global context in patch-based pipelines. We present Topoformer, a transformer framework for data and compute-efficient 3D classification through topological priors. First, we introduce a sliding-band cubical filtration that replaces a single global persistent-homology pass with overlapping intensity bands, producing an ordered sequence of Betti tokens representing connected components, tunnels, and cavities. These tokens serve as transformer inputs, enabling multi-scale topological reasoning while avoiding early filtration saturation. Second, we propose Topological Supervised Contrastive Learning (TopoSupCon), which treats each image and its label-preserving topological representation as complementary views, reducing reliance on brittle geometric or generative augmentations. A lightweight TopoGate module further allows the image encoder to softly weight multiple band widths on a case-by-case basis. On 3D brain MRI tumor grading and chest CT classification benchmarks under low-data regimes, Topoformer consistently improves over strong 3D CNN and ViT baselines, with gains of up to 12 AUC points and 8 accuracy points. These results suggest that sequential topology-aware representations provide an effective inductive bias for volumetric medical image analysis.

Skeletal Homology: Singular Homology at Scale

Ivy Dey, University of Tennessee, Knoxville

Simplicial homology at scale is analogous to the homology of Vietoris–Rips complexes. In a similar spirit, skeletal homology provides an analogue of singular homology at scale for metric spaces. We show that, for any metric space and fixed scale, skeletal homology agrees with the Vietoris–Rips homology at that scale. In addition, Gromov-Hausdorff "close" metric spaces have "close" skeletal homology at corresponding scales. This stability gives motivation for reducing the number of data points in computations under some conditions.

Wasserstein Stability of the Nerve Map

Kenneth McCabe, Northeastern University

Given a finite set of landmarks in a metric space, the nerve map sends a neighborhood of the landmarks into the geometric realization of the associated Čech complex via barycentric coordinates. We introduce a weighted nerve map, where landmarks are equipped with a probability measure, and study its stability under perturbations of the weights.

Our main result shows that, under a uniform nondegeneracy assumption, the weighted nerve map is Lipschitz stable with respect to a localized optimal transport metric. This metric induces the same topology as the 1-Wasserstein metric, yielding Wasserstein stability as a corollary. We also show that our estimate is sharp.

Symmetries in the Euler Characteristic Transform and Separable Shape Tensors

Evgeniya Lagoda, Southeastern Universities Research Association; NIST PREP

The action of the general linear group on Euclidean space induces an action on the Euler Characteristic Transform (ECT). Prior work has established injectivity, finite-direction sufficiency, and rotation/reflection-invariant summaries of the ECT; in particular, Curry-Mukherjee-Turner (2022) showed that the pushforward of the Lebesgue measure on the sphere under the ECT is an orthogonal-group invariant. In this poster, we aim to extend this perspective by considering the action of $\mathrm{GL}(d)$ on the ECT and investigating possible decompositions of this group action and invariants that can be extracted from the transform. After considering the general setting, we specialize to planar curves and analyze the relationship between ECT-based invariants and Separable Shape Tensors (SST), a product representation of a discrete curve introduced by Grey, Doronina, and Glaws (2023). SST, initially applied in aerodynamic design, represents a discrete parametrized curve in the product of the Grassmannian and symmetric positive definite cone, offering a natural decomposition into $\mathrm{GL}(d)$-invariant information (undulations) for comparison with ECT.

Low-Shot Graph Learning with Topological and Spectral Embeddings

Surbhi Kumar, University of Texas at Dallas

Deep graph learning has achieved remarkable success, but its reliance on abundant labeled data limits use in many scientific domains, where each label may require costly experiments or simulations. We revisit graph classification in the low-shot regime and ask whether explicit, theory-grounded descriptors can provide a reliable foundation and how they can be combined with modern architectures. We study two families of label-free embeddings: *topological vectors* from persistent homology that capture multiscale connectivity, and *spectral vectors* from the Laplacian density of states that summarize diffusion geometry. To harness their complementary strengths, we introduce *prototype embeddings*, which project graphs onto class-level prototypes in the joint topological–spectral space, and *STAMP*, a lightweight controller that conditions GNN and GT backbones on these descriptors through layer-wise modulation. Across ten TU benchmarks and label budgets $K \in \{1, 5, 10, 25, 50\}$, prototype embeddings with simple classifiers consistently outperform strong baselines in the extreme low-label setting, while *STAMP* achieves the best overall rank and smallest accuracy gap once $K \geq 10$. Our results demonstrate that explicit structural priors offer a powerful and complementary route to label-efficient graph learning, closing much of the gap to larger deep models without pretraining.