Foundations of Computational Geometry and Topology
Poster Session Abstracts
May 19, 2026
4:00-5:00pm

A novel topological characterization of brain-wide co-fluctuation patterns over time reveals brain-behavior links across spontaneous and evoked activity

Chunyin Siu, Stanford University

Traditional analyses of brain functional connectivity using fMRI data rely on Pearson correlation between brain regions: a static, time-averaged measure. While informative, this approach overlooks rich temporal dynamics. Here, we leverage the edge time series—the framewise product of fMRI activity between brain region pairs, whose temporal mean recovers the Pearson correlation—to capture fine-grained neural coordination (Faskowitz et al., 2020; Zamani Esfahlani et al., 2020; Owen et al., 2021). To our knowledge, we are the first to apply the Mapper algorithm—a topological data analysis (TDA) method (Singh et al., 2007)—to edge time series, enabling a visualization of brain-wide co-fluctuation patterns over time, which clinicians can use directly. This approach fills a key methodological gap in capturing the shape of edge-level dynamics across spontaneous (resting) and task-evoked conditions. We also find that mesoscale properties of the resulting Mapper graphs are significantly associated with individual differences in personality traits and self-reported psychopathology.

Dynamic persistent homology: Implementing computation

Jānis Lazovskis, University of Latvia

Persistent homology has become a central tool in topological data analysis. In many applications, data is inherently dynamic: evolving over time, admitting a hierarchical structure, or presenting filtrations that reflect inherent or interpreted relationships within the data. Despite this, most computational approaches to dynamic data treat it as a sequence of independent static inputs, recomputing persistence from scratch at each stage. While highly optimized algorithms and software for persistent homology have proliferated, comparatively few implementations explicitly address dynamic updates to the underlying data. Existing approaches include considering dynamics as a zigzag filtration, developed by Dey--Hou and implemented in FastZigZag; using sequences of swaps to describe all dynamics, introduced by Cohen-Steiner--Edelsbrunner--Morozov and implemented in Dionysus 1 and GUDHI; and restoring representative cycles after simplex removal, recently introduced by Giunti--Lazovskis with the SiRUP algorithm and implemented as a fork of PHAT. This poster surveys existing computational strategies for defining and handling dynamics in persistent homology. We examine how different softwares represent evolving filtrations, update boundary matrix factorizations, and track or update representative cycles underlying the persistence barcode.

Homomorphisms and chromatic numbers of Borsuk Graphs

Sucharita Mallick, University of Florida

The Borsuk graph Bor[S^k;r] is an infinite graph with vertex set the k-sphere and with edges between two vertices when they are at least r distance apart in the geodesic metric. For k>n, when does a graph homomorphism f: Bor[S^k;r] --> Bor[S^n;t] exist? Using topological obstructions, we prove there is no such graph homomorphism f when r<\pi and t > 2\pi/3. We give further bounds when n=1 (using chromatic numbers of graphs) or when f is measurable (using independence ratios of graphs). Additionally, we produce explicit graph homomorphisms f when r is large and t is small compared to sphere coverings and packings. Further, we investigate how many points need to be uniformly sampled from a sphere until the chromatic number of the resulting random Borsuk graph reaches a certain threshold.

Computing a Metric on the Shape Space of Tori

Yuan Luo, University of California, Davis

This paper presents a new optimization-based framework for comparing genus-one shapes, namely tori. We compute a metric structure on the shape space of tori with Riemannian metrics. While metric structures for 2-spheres are well-established, computing distances between tori is significantly more difficult due to their non-trivial topology. Drawing inspiration from Teichmüller theory, we formulate an algorithm for shape comparison that minimizes an energy functional over the space of diffeomorphisms within a fixed homotopy class. This homotopy class is canonical for a pair of tori embedded in $\mathbb{R}^3$. We derive a mathematically rigorous metric on the shape space of tori, and therefore produce a single number that measures the geometric similarity of the two tori. Our method also yields optimal, landmark-free correspondences between pairs of surfaces. We propose a novel computational framework to compute this metric. Experiments on flat tori, tori of revolution, and real-world scanned surfaces demonstrate that our approach captures natural geometric properties and enables robust comparison of genus-one shapes.

Principal Component Analysis on Manifolds

Wenwen Li, Florida State University

TBD

A glimpse into the space of connection matrices

Michał Lipiński, Institute of Science and Technology Austria (ISTA)

A connection matrix is an algebraic invariant capturing information about connecting orbits in a dynamical system and can be seen as a generalization of a Morse complex. This project is motivated by two recurring questions in the study of connection matrices: whether it is feasible to find all possible connection matrices for a given dynamical system as well as how to classify those matrices. Building upon the recent algorithms for computing connection matrices for combinatorial dynamical systems, a connection matrix can be obtained by finding the Morse complex corresponding to a Forman combinatorial vector field inscribed in the original multivector field. However, there exist connection matrices that cannot be obtained directly by applying the algorithm. To investigate the source of this discrepancy we construct a topological space AP(V) of acyclic subpartitions of the original combinatorial dynamical system V. In particular, this space is inhabited by all the Forman vector fields leading to connection matrices that can be obtained with the algorithm. Subsequently, using the notion of continuation of Morse decompositions, we subdivide the space into smaller regions corresponding to different gradient dynamics. With this preparation, we study the transfer morphisms connecting two neighboring Forman vector fields and the corresponding connection matrices, and analyze how the topological constraints of AP(V) affect the composition of those morphisms. The project is still at an early stage, but the observations give hope for a better understanding of the nature of connection matrices. One of the examples we explore was originally studied by Reineck (1990). It consists of an attracting periodic orbit surrounded by three repellers and three saddles. We observed that for a combinatorial model of that system the constructed space of acyclic partitions takes the form of a Möbius strip. This project is a joint work with Thomas Wanner

Parallel Computation of Critical Sets

Arnur Nigmetov, Lawrence Berkeley National Laboratory

Critical sets are chains whose filtration values should be modified to move a point of a persistence diagram. They make the convergence of topological optimization much more efficient in terms of the number of iterations, but their computation is more expensive than simple persistence pairing. We propose an algorithm to compute them in parallel and provide an interface to our implementation that makes topological optimization easy and intuitive.

TopoFormer: Topology Meets Attention for Graph Learning

Md Joshem Uddin, The University of Texas at Dallas

We introduce TopoFormer, a lightweight and scalable framework for graph representation learning that encodes topological structure into attention-friendly sequences. At the core of our method is Topo-Scan, a novel module that decomposes a graph into a short, ordered sequence of topological tokens by slicing over node or edge filtrations. These sequences capture multi-scale structural patterns, from local motifs to global organization, and are processed by a Transformer to produce expressive graph-level embeddings. Unlike traditional persistent homology pipelines, Topo-Scan is parallelizable, avoids costly diagram computations, and integrates seamlessly with standard deep learning architectures. We provide theoretical guarantees on the stability of our topological encodings and demonstrate state-of-the-art performance across graph classification and molecular property prediction benchmarks. Our results show that TopoFormer matches or exceeds strong GNN and topology-based baselines while offering predictable and efficient compute. This work opens a new path for parallelizable and unifying approaches to graph representation learning that integrate topological inductive biases into attention frameworks.

Lower-bounding the Gromov-Hausdorff Distance Between Balls

Kushagri Sharma, University of Florida

What is the Gromov-Hausdorff distance between Euclidean unit balls of different dimensions, $d_\gh(B^m,B^n)$ for $m>n$? Note that the lower bound coming from the stability of persistent homology is zero, is zero since all balls possess trivial persistent homology. To lower bound these values we exploit the Borsuk-Ulam theorem, namely that any continuous map between a sphere and a ball of appropriate dimensions identifies antipodal points. We prove that $d_\gh(B^m,B^n)\ge \frac{\sqrt{n+1}}{\sqrt{n+1}+\sqrt{n}} > \frac{1}{2}$ for $m>n\ge 1$.

Topological Representations and Dynamics of Hunger and Thirst

Trong-Thuc Trang, Florida Atlantic University

Uncovering meaningful representations and dynamics within neuronal signals and associating those findings with specific behaviors is a foundational task in computational neuroscience. In a joint project with the Stern Lab at the Max Planck Florida Institute for Neuroscience, we performed various topological data analyses on neural and behavioral datasets obtained from the insular cortex of hungry and thirsty mice. Our goal was to reveal how these internal states interact with eating and drinking behaviors. Our findings indicate that the brain does not utilize a single, unified "ingestion" circuit for both behaviors, but rather employs distinct neural mechanisms.