Abstracts

Contents

Updated: 30 September 2024

Pierre Alliez
Inria Sophia-Antipolis, France

Quadric Error Metrics for Variational Reconstruction and Neural Mesh Representation

Inspired by the strengths of quadric error metrics initially designed for mesh decimation, we propose a concise mesh reconstruction approach for 3D point clouds. Our approach proceeds by clustering the input points enriched with quadric error metrics, where the generator of each cluster is the optimal 3D point for the sum of its quadric error metrics. This approach favors the placement of generators on sharp features, and tends to equidistribute the error among clusters.

As a follow-up work, I will present a novel learnable mesh representation through a set of local 3D sample Points and their associated Normals and Quadric error metrics (QEM) w.r.t. the underlying shape, which we denote PoNQ. A global mesh is directly derived from PoNQ by efficiently leveraging the knowledge of the local quadric errors. Besides marking the first use of QEM within a neural shape representation, our contribution guarantees both topological and geometrical properties by ensuring that a PoNQ mesh does not self-intersect and is always the boundary of a volume.

Back to Contents Page →

Chandrajit Bajaj
The University of Texas at Austin, USA

DPO: Differential reinforcement learning with application to optimal geometric and topological configuration search

Reinforcement learning (RL) with continuous state and action spaces is arguably one the most challenging problems within the field of machine learning. Most current learning methods focus on integral identities such as value (Q) functions to derive an optimal strategy for the learning agent. In this talk we present the dual form of the original RL formulation to propose the first differential RL framework that can handle settings with limited training samples and short-length episodes. Our approach introduces Differential Policy Optimization (DPO), a pointwise and stage-wise iteration method that optimizes policies encoded by local-movement operators. We prove a pointwise convergence estimate for DPO and provide a regret bound comparable with the best current theoretical derivation. Such pointwise estimate ensures that the learned policy matches the optimal path uniformly across different steps. We then apply DPO to a class of practical RL problems with continuous state and action spaces, and which search for optimal geometric and topological configurations with Lagrangian rewards. DPO is easy to implement, scalable, and shows competitive results on benchmarking experiments against several popular RL methods.

This is joint work with Minh Nguyen.

Back to Contents Page →

Xavier Bresson
National University of Singapore, Singapore

Graph Transformers and Developments

Graph Neural Networks (GNNs) have shown great potential in the field of graph representation learning. Standard GNNs define a local message-passing mechanism which propagates information over the whole graph domain by stacking multiple layers. This paradigm suffers from two major limitations, over-squashing and poor long-range dependencies, that can be solved using global attention but significantly increases the computational cost to quadratic complexity. In this work, we propose an alternative approach to overcome these structural limitations by leveraging the ViT/MLP-Mixer architectures introduced in computer vision. We introduce a new class of GNNs, called Graph MLP-Mixer/ViT, that holds three key properties. First, they capture long-range dependency as demonstrated on the long-range LRGB datasets and mitigate the over-squashing issue on the TreeNeighbour dataset. Second, they offer memory and speed efficiency, surpassing related techniques. Third, they show high expressivity in terms of graph isomorphism as they can distinguish at least 3-WL isomorphic graphs. As a result, this novel architecture provides significantly better results over standard message-passing GNNs for molecular datasets.

Back to Contents Page →

Cheng Cheng
Sun Yat-sen University, China

Graph Fourier Transforms on Directed Graphs

The Graph Fourier transform (GFT) is a fundamental tool in graph signal processing. Based on singular value decomposition of the Laplacian, we introduce a novel definition of GFT on directed graphs, and use the singular values of the Laplacian to carry the notion of graph frequencies. We show that the proposed GFT has its frequencies and frequency components evaluated by solving some constrained minimization problems with low computational cost, and it could represent graph signals with different modes of variation efficiently. Moreover, the proposed GFT is consistent with the conventional GFT in the undirected graph setting, and on directed circulant graphs, it is the classical discrete Fourier transform, up to some rotation, permutation and phase adjustment.

Back to Contents Page →

Chuan-Shen Hu
Nanyang Technological University, Singapore

Quotient Complex (QC)-based Learning Models for Material Data Analysis

Materials innovation and design offer the ultimate solution for humanity's pressing challenges of sustainable energy and climate change. Efficient representation and characterization of material crystal structures, in particular, their periodic and high-order information, are of key importance for material data analysis. In this talk, we will discuss the proposed quotient complex (QC), which is a novel material representation that can characterize both material periodic and high-order information. We also develop QC-based learning models, including QC-GBT and QCformer. In QC-GBT, material features are generated from QC-based topological data analysis (TDA) models and further combined with GBT. In QCformer, a QC-based simplex attention module is proposed and incorporated into topological deep learning architecture. Our QC-based learning models are extensively trained and validated on datasets such as NMSE, Material Project, and JARVIS. Our results demonstrate that QC-based learning models outperform state-of-the-art methods, offering a powerful new approach to material data analysis and property prediction.

Back to Contents Page →

Jiang Jian
Wuhan Textile University, China

Harnessing Topology and AI for Virtual Screening in Drug Design

Drug addiction is a global public health crisis, and the design of antiaddiction drugs remains a major challenge due to intricate mechanisms. Since experimental drug screening and optimization are too time-consuming and expensive, there is urgent need to develop innovative artificial intelligence (AI) methods for addressing the challenge. We tackle this challenge by topology-inferred drug addiction learning (TIDAL) built from integrating multiscale topological Laplacians, deep bidirectional transformer, and ensemble-assisted neural networks (EANNs). Multiscale topological Laplacians are a novel class of algebraic topology tools that embed molecular topological invariants and algebraic invariants into its harmonic spectra and nonharmonic spectra, respectively. These invariants complement sequence information extracted from a bidirectional transformer. We validate the proposed TIDAL framework on 22 drug addiction related, 4 hERG, and 12 DAT data sets, which suggests that the proposed TIDAL is a state-of-the-art framework for the modeling and analysis of drug addiction data. We carry out cross-target analysis of the current drug addiction candidates to alert their side effects and identify their repurposing potentials. Our analysis reveals drug-mediated linear and bilinear target correlations. Finally, TIDAL is applied to shed light on relative efficacy, repurposing potential, and potential side effects of 12 existing antiaddiction medications. Our results suggest that TIDAL provides a new computational strategy for pressingly needed antisubstance addiction drug development.

Back to Contents Page →

Parvaneh Joharinad
University of Leipzig, Germany

Locally distorted metric spaces and their representations

In this talk, we present a novel approach to learning the manifold underlying a dataset initially represented as a discrete metric space, with a particular focus on non-uniform distributions, an assumption that can prove overly restrictive in various standard manifold learning-based dimension reduction techniques. We will show how, motivated by conformal transformations, local distortions of the distance function, will address the problem, and then present a method for merging this family of metric spaces, thereby yielding a global unified metric on the dataset.

Back to Contents Page →

Ye Ke
Academy of Mathematics and Systems Science, CAS, China

Kempe's Universality Theorem and rational curves on real classical groups

Kempe's Universality theorem is one of the most important theorems in theoretical mechanism. Roughly speaking, it states that every algebraic plane curve can be realized by some linkage. In this talk, we generalize the theorem to rational curves on real classical groups. We will discuss the decomposition and classification of such curves. As an application, we will explain how our results can be applied to mechanism.

Back to Contents Page →

Stephan Klaus
Mathematisches Forschungsinstitut Oberwolfach, Germany

A topological perspective on the space of all samplings from a manifold

Hyperbolic Neural Networks (HNNs) have recently found successful in representing hierarchical and complex data. However, unlike other domains, the exploration of hyperbolic neural operations and the development of effective hyperbolic representations remains limited. We delve into developing HNN architectures and addressing the stability and robustness issues. In this talk, we discuss some of our recent findings in the following aspects: first, we investigate representation learning through hyperbolic convolutions with provable properties; second, we enhance representations with Gromov-Wasserstein regularization; third, we improve stability and robustness by designing hyperbolic operations and regularization techniques. Finally, we demonstrate the practical applications of these advancements in tasks such as few-shot image classification, graph classification, and anomaly detection.

Back to Contents Page →

Christian Kuehn
Technical University of Munich, Germany

Dynamics on Neural Networks

In this talk, I shall the discuss several techniques that naturally can get transported from dynamical systems to the analysis of machine learning algorithms, particularly in the context of deep neural nets. First, I am going to sketch results on mean-field/continuum limits to link discrete networks to classical continuous-time models. In the

main part of the talk, I am going to use a normal form type approach to classify various neural network and neural ODE architectures in their ability to represent Morse functions, which provides a direct link to geometric analysis and singularity theory. Time permitting, I will also briefly discuss the embedding problem for neural ODEs as well as a rigorous proof of high complexity dynamics using computer assistance.

Back to Contents Page →

Ming Li
Zhejiang Normal University, China

From Heterophilous Graph Learning to Heterophilous Hypergraph Learning: Exploring New Frontiers

This talk begins by presenting recent advances in heterophilous graph learning, a field that addresses the challenges posed by heterophilous data in graph neural networks (GNNs). I will discuss notable works that have improved our understanding of how to learn effective representations in heterophilous graph settings. Moving forward, I will introduce the emerging topic of heterophilous hypergraph learning, which expands the scope of GNNs to hypergraphs that naturally model complex, higher-order relationships. As part of this new direction, I will present a recent work that showcases the potential of heterophilous hypergraph neural networks (HHGNNs) to capture richer structural information. The talk will highlight new metrics, essential concepts, and newly developed benchmark datasets that are driving this exciting field forward, along with future research opportunities.

Back to Contents Page →

Pietro Lio
University of Cambridge, UK

Actionable and responsible graph representation learning in medicine

---

In this talk I will focus on how to build a digital patient twin using graph and hypergraph representation learning and considering physiological (cardiovascular), clinical (inflammation) and molecular variables (multi omics and genetics). I will consider patient disease trajectories through the use of neural graph ODEs. I will discuss how these approaches could also keep the clinicians in the loop to avoid excessive automatisation using logic and network explainer frameworks.

Back to Contents Page →

Lin Liu
Shanghai Jiao Tong University, China

Statistical Inference of GLMs and Causal Effects in Observational Studies under Proportional Asymptotics

TBA n this talk, we consider the statistical problem that even undergraduates know about – inference on regression coefficients in GLMs. This problem is almost completely addressed in the low-dimensional setting ($p \ll \sqrt{n}$) and in the traditional high-dimensional sparse setting ($p \gg n$ but $s \ll n$). Here, we consider the the "modern" high-dimensional setting with $p/n \to c \in (0, \infty)$. We first survey the current state-of-the-art methods based on mainly Approximate Message Passing (AMP) and Stein's lemma under Gaussian designs. Interestingly, we show that, despite the success of AMP, we can instead construct extremely simple, moment-based estimators that have almost the same theoretical and empirical guarantees as the AMP-based estimators, as long as a particular diffeomorphism between the parameter space and the moment space holds. These estimators do not require any AMP machinary and can be easily generalized to non-Gaussian settings. Numerical experiments support our theoretical findings. Finally, if time permitting, we will discuss how the moment approach can also boost (1) AMP-based estimators and (2) recent proposals by statistical physicists on empirical Bayes methods for high-dimensional GLMs.

Back to Contents Page →

Stephen Marron
University of North Carolina at Chapel Hill, USA

Object Oriented Data Analysis

The rapid change in computational capabilities has made Big Data a major modern statistical challenge. Less well understood is the rise of Complex Data as a perhaps greater challenge. Object Oriented Data Analysis (OODA) is a framework for addressing this, in particular providing a general approach to the definition, representation, visualization and analysis of Complex Data. The notion of OODA generally guides data analysis, through providing a useful terminology for interdisciplinary discussion of the many choices typically needed in modern complex data analyses. The main ideas are illustrated through several OODA contexts, including shapes, trees (in the sense of graph theory), covariance matrices and nonnegative curves as data objects.

Back to Contents Page →

Konrad Polthier
Freie Universität Berlin, Germany

Vibrations of Geometric Shapes

Vibrations of musical strings are well understood by Fourier analysis while the vibrations of geometric shapes exhibit surprising properties triggered by careful choices of differential geometric energies. We will review solved problems and introduce novel approaches with applications in biology, computer graphics and crystallography, and in topology.

Back to Contents Page →

Areejit Samal
The Institute of Mathematical Sciences (IMSc), India

Geometry-inspired measures with diverse applications in network science

---

In this talk, I will present our work on the development of geometry-inspired measures for the edge-based characterization of real-world complex networks. In particular, we were first to introduce a discretization of the classical Ricci curvature proposed by R. Forman to the domain of real-world complex networks. Forman-Ricci curvature is an attractive tool in network science due to the following reasons. Firstly, most traditional graph-theoretic measures such as degree and clustering coefficient are vertex-specific, while Forman-Ricci curvature is edge-specific. Secondly, the mathematical formula of the Forman-Ricci curvature elegantly allows for the analysis of weighted and unweighted graphs. Thirdly, we have also extended the definition of Forman-Ricci curvature to the realm of directed graphs. Fourthly, an important distinguishing feature of the Forman-Ricci curvature, in contrast to the other well-known discretization, namely, Ollivier-Ricci curvature, is its simplicity and suitability from a computational perspective for analysis of very large networks. Fifthly, we have developed an augmented version of the Forman-Ricci curvature which is suitable for analysis of higher-order networks. In this talk, I will mainly focus on the successful applications of Forman-Ricci curvature to real-world networks across different domains including life science and finance. Specifically, I will present application of discrete Ricci curvatures to: (a) brain functional connectivity networks constructed from resting-state fMRI data, and (b) time-series of financial networks

Back to Contents Page →

Emil Saucan
Braude College of Engineering Karmiel, Israel

Discrete Geometric Invariants for Textures Classification

We propose a discrete geometric approach to the intelligence of complex networks and to the classification of textures in images, with special emphasis on natural ones. We first make appeal to a number of discrete notions of curvature eminently suited for this task, namely both the graph and the full Forman-Ricci curvatures, the Ollivier- Ricci curvature and the Menger curvature measure. Furthermore, we consider a different type of geometric network measure, inspired by the early work of Duffin on electrical networks and stemming from Complex Function Theory, namely the network modulus. In addition, we propose a number of metric invariants introduced by Grove and Markvorsen that encode the essential global geometry of the given structure. Combining these geometric measures with comparison methods developed originally for the study of networks, we are able to distinguish and classify various types of textures, with a special focus on stochastic textures.

Back to Contents Page →

Sho Sonoda
RIKEN, Japan

Deep Ridgelet Transform: Harmonic Analysis for Deep Neural Network

The ridgelet transform has been developed to study neural network parameters, and it can describe the distribution of parameters. Mathematically, it is defined as a pseudo-inverse operator of neural networks. Namely, given a function f, and network $NN[\gamma]$with parameter γ, the ridgelet transform $R[f]$ for the network NN satisfies the reconstruction formula $NN[R[f]]=f$. For depth-2 fully-connected networks on a Euclidean space, the ridgelet transform has been discovered up to the closed-form expression, thus we could describe how the parameters are distributed. However, for a variety of modern neural network architectures, the closed-form expression has not been known. In this talk, I will introduce a systematic method to induce the generalized neural networks and their corresponding ridgelet transforms from group equivariant functions, and present an application to deep neural networks.

Back to Contents Page →

Xian Wei
East China Normal University, China

Geometric Transformer Learning for Point Clouds

Due to the outstanding competence in capturing long-range relationships, self-attention mechanism based Transformers have achieved remarkable progress in modeling point clouds, such as LiDAR point clouds, high dimensional image data points, movement trajectories, and Molecules. Nevertheless, point cloud object often has complex non-Euclidean spatial structures, with the behavior changing dynamically and unpredictably. Most current self-attention modules highly rely on the dot product multiplication and matrix transformations in Euclidean space, which cannot capture internal non-Euclidean structures of point cloud objects, especially the long-range relationships along the curve of the implicit manifold surface represented by point cloud objects. To address this problem, we will introduce our several progresses on Geometric Transformer learning models, based on the manifold assumption.

Back to Contents Page →

Tailin Wu
Westlake University, China

Learning adaptive and compositional models for multi-resolution simulation and inverse design

In science and engineering, multi-resolution simulation and inverse design are two universal tasks. For the multi-resolution challenge in simulation, where a small fraction of the system is extremely dynamic while other places are changing slowly, we introduce a LAMP architecture that consists of two Graph Neural Networks (GNNs), one for learning the forward evolution and one for learning the policy of spatial refinement and coarsening. Together, they jointly learn to reduce prediction error and devote more computation to the highly dynamic regions. For the inverse design task, we introduce a generative method that learns the energy function representing the joint probability distribution of state trajectory and design parameters. During inference, the method generates new designs by optimizing the learned energy function alongside the design objective. We demonstrate our methods in challenging 1D and 2D PDE systems, for example, designing complex airfoil shapes to enhance lift-to-drag ratios.

Back to Contents Page →

Rex Zhitao Ying
Yale University, USA

Building Foundation Models via Hyperbolic Representation Learning

When considering ubiquitous and important relational information between entities such as text documents, images and nodes in complex networks, they often exhibit non-Euclidean properties, where the relationships are reflected by distances in non-Euclidean space. Here we explore the use of curved, non-Euclidean embedding space in capturing relationships between tokens or entities of different modalities in self-supervised learning. Firstly, we demonstrate a generic, fully hyperbolic Transformer backbone for pre-training architecture. Such building blocks then allow us to create foundation models in retrieval and graph learning. Finally, we show the effectiveness of a hyperbolic adapter to convert existing foundation models including large language models into non-Euclidean space, achieving state-of-the-art performance among LLMs on a diverse range of reasoning tasks.

Back to Contents Page →

Kun Zhan
Lanzhou University, China

Graph neural estimators

The topic of my discussion revolves around graph neural estimators, typically involving an inequality. On one side of the inequality, there is a constant representing the desired outcome, while on the other side, there is a variable that can be modeled using a neural network.

The goals in machine learning are pretty similar to the idea of the shortest code length in information theory. The shortest code length is what we call information entropy, and you can think of machine learning as basically trying to get as close as possible to this entropy concept.

Back to Contents Page →

Bohang Zhang
Peking University, China

A Quantitative Framework for Graph Neural Network Expressiveness via Graph Homomorphism

Designing expressive Graph Neural Networks (GNNs) is a fundamental topic in the graph learning community. So far, GNN expressiveness has been primarily assessed via the Weisfeiler-Lehman (WL) hierarchy. However, such an expressivity measure has notable limitations: it is inherently coarse, qualitative, and may not well reflect practical requirements (e.g., the ability to encode substructures). This talk will introduce a novel framework for quantitatively studying the expressiveness of GNN architectures, addressing all the above limitations. Specifically, we introduce a fundamental expressivity measure termed homomorphism expressivity, which quantifies the ability of GNN models to count graphs under homomorphism. Homomorphism expressivity offers a complete and practical assessment tool: the completeness enables direct expressivity comparisons between GNN models, while the practicality allows for understanding concrete GNN abilities such as subgraph counting. By examining four classes of prominent GNNs as case studies, we can derive simple, unified, and elegant descriptions of their homomorphism expressivity for both invariant and equivariant settings. These results provide novel insights into a series of previous work, unify the landscape of different subareas in the community, and settle several open questions.

Back to Contents Page →

Dingxuan Zhou
University of Sydney, Australia

Mathematical theory of structured deep neural networks

Deep learning has been widely applied and brought breakthroughs in speech recognition, computer vision, natural language processing, and many other domains. The involved deep neural network architectures and computational issues have been well studied in machine learning. But there is much less theoretical understanding about the modelling, approximation or generalization abilities of deep learning models with network architectures. An important family of structured deep neural networks is deep convolutional neural networks (CNNs) induced by convolutions. The convolutional architecture gives essential differences between deep CNNs and fully-connected neural networks, and the classical approximation theory for fully-connected networks developed around 30 years ago does not apply. This talk describes approximation and generalization analysis of deep CNNs and related structured deep neural networks.

Back to Contents Page →

Xiaosheng Zhuang
City University of Hong Kong, Hong Kong SAR

Permutation Equivariant Graph Framelets for Heterophilous Graph Learning

The nature of heterophilous graphs is significantly different from that of homophilous graphs, which causes difficulties in early graph neural network models and suggests aggregations beyond the 1-hop neighborhood. In this talk, we discuss a new way to implement multi-scale extraction via constructing Haar-type graph framelets with desired properties of permutation equivariance, efficiency, and sparsity, for deep learning tasks on graphs. We design a graph framelet neural network model PEGFAN (Permutation Equivariant Graph Framelet Augmented Network) based on our constructed graph framelets. The experiments are conducted on a synthetic dataset and 9 benchmark datasets to compare performance with other state-of-the-art models. The result shows that our model can achieve the best performance on certain datasets of heterophilous graphs (including the majority of heterophilous datasets with relatively larger sizes and denser connections) and competitive performance on the remaining. This is joint work with Jianfei Li (CityUHK), Ruigang Zheng (CityUHK), Han Feng (CityUHK), and Ming Li (Zhejiang Normal University).

Back to Contents Page →

Dongmian Zou
Duke Kunshan University, China

Exploring Effective Representations Using Hyperbolic Neural Networks

Hyperbolic Neural Networks (HNNs) have recently found successful in representing hierarchical and complex data. However, unlike other domains, the exploration of hyperbolic neural operations and the development of effective hyperbolic representations remains limited. We delve into developing HNN architectures and addressing the stability and robustness issues. In this talk, we discuss some of our recent findings in the following aspects: first, we investigate representation learning through hyperbolic convolutions with provable properties; second, we enhance representations with Gromov-Wasserstein regularization; third, we improve stability and robustness by designing hyperbolic operations and regularization techniques. Finally, we demonstrate the practical applications of these advancements in tasks such as few-shot image classification, graph classification, and anomaly detection.

Back to Contents Page →