Cai Zhou

Massachusetts Institute of Technology – Cambridge, 02135, MA, USA
caiz428@mit.edu • Homepage

First year PhD at MIT EECS. Advisors: Tommi Jaakkola, Stephen Bates.

Education Background

- **Massachusetts Institute of Technology** — **Cambridge, MA**
  PhD in EECS, CSAIL & LIDS — *2024-2029*
- **Tsinghua University** — **Beijing**
  Undergraduate in Department of Automation, Class of General Artificial Intelligence — *2021-2024*
  GPA: 3.93/4.0, Ranking: Top 2%
- **Tsinghua University** — **Beijing**
  Undergraduate in Statistics, Joint Degree — *2021-2024*
- **Tsinghua University** — **Beijing**
  Undergraduate in School of Economy and Management — *2020-2021*

Research Interest

My research interest lies broadly in **theoretical and applied machine learning**. I aim to understand the foundations of machine learning, with a special focus on its probabilistic and geometric nature. My current research focuses on the intersection of generative models and AI4Science. I'm also interested in application areas including computer vision, natural language processing and computational biology.

Publications and Research Experience

Conference and Journal Papers

- Unifying Generation and Prediction on Graphs with Latent Graph Diffusion
  Cai Zhou, Xiyuan Wang, Muhan Zhang.
  Thirty-eighth Conference on Neural Information Processing Systems (NeurIPS, 2024) [PDF]

  TL;DR: Propose Latent Graph Diffusion, a generative framework for both generation and prediction tasks utilizing latent graph diffusion. Derive theoretical guarantees for solving regression and classification tasks with diffusion models. Strong performance in both generation and prediction tasks for molecules and general graphs.

- On the Theoretical Expressive Power and Design Spaces of High Order Graph Transformers
  Cai Zhou, Rose Yu, Yusu Wang
  Twenty-seventh International Conference on Artificial Intelligence and Statistics (AISTATS, 2024)

  TL;DR: Theoretically analyze the expressive power and approximation power of high order graph transformers; propose scalable and powerful high order graph transformers and simplicial transformers.

- Facilitating Graph Neural Networks with Random Walk on Simplicial Complexes
  Cai Zhou, Xiyuan Wang, Muhan Zhang.
  Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS, 2023) [PDF]

  TL;DR: Design positional and structural encodings for GNNs based on random walk on simplicial complexes, specifically EdgeRWSE and Hodge1Lap for 1-simplices (edge-level) that can provably improve expressive power and practical performance.

- From Relational Pooling to Subgraph GNNs: A Universal Framework for More Expressive Graph Neural Networks
  Cai Zhou*, Xiyuan Wang*, Muhan Zhang.
  Fortieth International Conference on Machine Learning (ICML, 2023) [PDF]

  TL;DR: Propose k, l-WL algorithm (running k-WL on a graph with l labels) and theoretically establish a strict expressive power hierarchy; Incorporate a wide range of GNNs including relational pooling and subgraph GNNs.

- Locally Supervised Deep Learning by Maximizing Information Propagation
  Yulin Wang, Zanlin Ni, Yifan Pu, **Cai Zhou**, Shiji Song, Gao Huang
  International Journal of Computer Vision (IJCV)

  TL;DR: Apply information theory to model locally supervised learning, propose InfoPro loss to alleviate information collapse in locally trained deep networks and verify the effectiveness in computer vision tasks.

Preprints

- Geometric Representation Condition Improves Equivariant Molecule Generation
  Zian Li, **Cai Zhou**, Xiyuan Wang, Xingang Peng, Muhan Zhang.
  Under review [PDF]

  TL;DR: Propose GeoRCG, a general framework to enhance the performance of molecular generative models by integrating geometric representation conditions. Provide theoretical guarantees and achieve SOTA performance on several molecule generation tasks.

- Towards Stable, Globally Expressive Graph Representations with Laplacian Eigenvectors
  Junru Zhou, **Cai Zhou**, Xiyuan Wang, Pan Li, Muhan Zhang.
  Under review [PDF]

  TL;DR: Propose a novel method exploiting Laplacian eigenvectors to generate stable and globally expressive graph representations, utilizing learnable $O(p)$-invariant representations for each Laplacian eigenspace and applying a smooth function for generalization.

Academic Services

- Reviewer for AISTATS 2024, NeurIPS 2024, ICLR 2025, AISTATS 2025
- Teaching Assistant for General Artificial Intelligence System Practice, 2023-Summer, Tsinghua University

Crucial Honors and Awards

Comprehensive Awards

- Outstanding Graduate by Beijing Ministry of Education and by Tsinghua University
- Outstanding Undergraduate Thesis by Beijing Ministry of Education and by Tsinghua University, 2024
- National Scholarship (Highest honor for undergraduates in China), 2023
- Comprehensive Excellence Award of Tsinghua University (Highest honor in THU and the Dept. Automation), 2023 & 2021
- Research and Innovation Excellence Award of Tsinghua University, 2022
- Academic Excellence Award of Tsinghua University, 2022
- Several prizes in national English speech contests and proficiency competitions.

Mathematics and Physics

- First prize of National Physics Competition for Undergraduates, 2021

Technical and General Skills

Technical Skills

- **Mathematics:** Calculus, Algebra, Discrete and Combinatorial Mathematics, Differential Geometry, Algebraic Topology, Complex Analysis, Fourier Analysis, Functional Analysis, Operation Research, etc.
- **Statistics:** Probability Theory, Statistical Inference, Stochastic Process, Multivariate Statistics and Regression, Computational Statistics, Time Series Analysis, Causal Inference, Bayesian Statistics, Biostatistics, etc.
- **Programming Skills:** Proficient in Python, Pytorch, PyG, DGL, C, C++, R, Linux, Latex, Markdown, Git, etc. Familiar with Tensorflow, Keras, C#, HTML, MATLAB, etc.

General Skills

Fluent in English; Good writing and presentation skills; Outstanding leadership and cooperation skills.