




**Conference on the
Mathematical Theory of Deep
Neural Networks**

November 6-7, 2025
University of Michigan

www.deepmath-conference.com

Contents

About DeepMath	3
Code of Conduct	4
Important Information	5
Schedule	6
Thursday, Nov 6	6
Friday, Nov 7	7
Invited Talks	8
Accepted Posters	9

About DeepMath

Recent advances in deep neural networks (DNNs), combined with open, easily-accessible implementations, have made DNNs a powerful, versatile method used widely in both machine learning and neuroscience. These advances in practical results, however, have far outpaced a formal understanding of these networks and their training. The dearth of rigorous analysis for these techniques limits their usefulness in addressing scientific questions and, more broadly, hinders systematic design of the next generation of networks. Recently, long-past-due theoretical results have begun to emerge from researchers in a number of fields. The purpose of this conference is to give visibility to these results, and those that will follow in their wake, to shed light on the properties of large, adaptive, distributed learning architectures, and to revolutionize our understanding of these systems.

Code of Conduct

All attendees, speakers, sponsors and volunteers at DeepMath are required to agree with the following code of conduct. Organizers will enforce this code throughout the event, which includes both the main conference and any conference-associated events. We expect cooperation from all participants to help ensure a safe environment for everybody.

DeepMath is dedicated to providing a harassment-free conference experience for everyone, regardless of gender, gender identity and expression, age, sexual orientation, disability, physical appearance, body size, race, ethnicity, religion (or lack thereof), education, mental health, or technology choices. Harassment (whether related to the above characteristics or not) includes offensive verbal comments, deliberate intimidation, stalking, following, harassing photography or recording, sustained disruption of talks or other events, inappropriate physical contact, and unwelcome sexual attention. Behavior that skirts the specific examples described in the code of conduct but is nonetheless commonly understood to be harassment (i.e. it violates the spirit but not the letter of the code of conduct), is a violation of the code of conduct.

We do not tolerate harassment of conference participants and staff in any form. Participants asked to stop any harassing behavior are expected to comply immediately. If a participant engages in harassing behavior, the conference organizers may take any action they deem appropriate, including warning the offender or expulsion from the conference with no refund.

If you are being harassed, notice that someone else is being harassed, or have any other concerns, please contact a member of conference staff or organizers immediately. Conference staff can be identified as they'll be wearing specifically marked badges.

If you make a report, conference staff will work with you to understand the situation and assist you to feel safe and welcomed at DeepMath. You will not be intentionally identified, and any action that may incidentally identify you will be discussed with you before we take it. In addition to corrective action against the violator of the code of conduct, organizers can help participants contact venue security or local law enforcement, provide escorts, or otherwise assist those experiencing harassment to feel safe for the duration of the conference.

Important Information

- **The conference venue.** Talks will be held at the Johnson Room at the Lurie Engineering Center on the University of Michigan campus.
- **The poster session.** The poster session will be held will be held on Nov 7, 2022 4-7pm in the ECE building on Beal Ave.

Questions / Comments / Concerns? Find an organizer in person, or email deepmath.conference@gmail.com.

Schedule

Thursday, Nov 6

Day 1: Nov 6	
Time	What
8:00 AM	Breakfast
9:00 AM	Intro
9:30	Invited talk 1: Laura Balzano
10:30	Break
11:00	Contributed talk: Compositional Reasoning with Transformers, RNNs, and Chain of Thought
11:20	Contributed talk: A Scalable Measure of Loss Landscape Curvature for Analyzing the Training Dynamics of LLMs
11:40	Contributed talk: Generalization of Diffusion Models Arises from a Regularized Representation Space
12:00	Lunch
2:00	Invited talk 2: Robert Nowak
3:00	Contributed talk: Data Augmentation: A Fourier Analysis Perspective
3:20	Contributed talk: Domain generalization: A Tale of Two ERMs
3:50	Break
4:00	Invited talk 3: Surbhi Goel
5:00	Panel: AI 4 math? Math 4 AI

Friday, Nov 7

Day 2: Nov 7	
Time	Who/what
8:00 AM	Breakfast
9:30	Invited talk 4: Jelena Diakonikolas
10:30	Break
11:00	Contributed talk: Absorb and Converge: Provable Convergence Guarantee for Absorbing Discrete Diffusion Models
11:20	Contributed talk: Understanding In-Context Learning's Out-of-Distribution Generalization Abilities via Low-Dimensional Subspaces
11:40	Contributed talk: Multi-head Transformers Provably Learn Symbolic Multi-step Reasoning via Gradient Descent
12:00	Lunch
2:00	Invited talk 5: Thomas Icard
3:00	Invited talk 6: Anqi Liu
4:00	Poster session

Invited Talks

Compression for Deep Overparameterized Learning and Fine-Tuning

Laura Balzano

University of Michigan

Gradient Alignment, Learning, and Optimization

Jelena Diakonikolas

University of Wisconsin–Madison

Human-AI Collaboration: Agreement and Alignment

Surbhi Goel

University of Pennsylvania

Causal Abstraction and Interpretability

Thomas Icard

Stanford University

Towards Trustworthy Decision Making: Calibration and Alignment under Distribution Shift

Anqi Liu

Johns Hopkins University

Function Space Perspectives on Neural Networks

Robert D Nowak

University of Wisconsin–Madison

ID 1: From Compression to Expansion: A Layerwise Analysis of In-Context Learning

Jiachen Jiang; Yuxin Dong; Jinxin Zhou; Zhihui Zhu

ID 2: A Scalable Measure of Loss Landscape Curvature for Analyzing the Training Dynamics of LLMs

Dayal Singh Kalra; Jean-Christophe Gagnon-Audet; Andrey Gromov; Ishita Mediratta; Kelvin Niu; Alexander H Miller; Michael Shvartsman

ID 3: On the Role of Transformer Feed-Forward Layers in Nonlinear In-Context Learning

Haoyuan Sun; Ali Jadbabaie; Navid Azizan

ID 4: etrov-Galerkin Networks: A Koopman Operator Architecture for Interpretable Forecasting

Alexander Richardson; Jeova Farias Sales Rocha Neto
Affil1; Affil2;

ID 5: Compositional Reasoning with Transformers, RNNs, and Chain of Thought

Noah Amsel; Gilad Yehudai; Joan Bruna

ID 6: DuoLens: A Framework for Robust Detection of Machine-Generated Multilingual Text and Code

Shriyansh Agrawal; Aidan Lau; Sanyam Shah; Ahan M R

ID 7: Exploring the Inner Workings and Internal Representations of Predictive Coding Networks in Comparison to Usual Feedforward Neural Networks

Aslan Satary Dizaji

ID 8: Decipher Deep Math: Numeric Rounding Behaviors in LLMs

Chiung-Yi Tseng; Maisha thasin; Danyang Zhang; Blessing Effiong; Somshubhra Roy

ID 9: What Happens During the Loss Plateau? Understanding Abrupt Learning in Transformers

Pulkit Gopalani; Wei Hu

ID 10: Gradient Confusion in Continual Learning

Adrian Candocia; Aamod Varma; Raaghav Agarwal; Omer Mustafa Inan; Mark A. Davenport

ID 11: Exploring the Manifold of Neural Networks using Diffusion Geometry

Andrew J Steindl; Elliott Abel; Selma Mazioud; Ellie L. Schueler; Folu Ogundipe, Ellen Zhang; Yvan Grinspan; Kristof Reimann; Peyton Crevasse; Dhananjay Bhaskar; Siddharth Viswanath; Yanlei Zhang; Tim G. J. Rudner; Ian Adelstein; Smita Krishnaswamy

ID 12: Generalization of Diffusion Models Arises from a Regularized Representation Space

Zekai Zhang; Xiang Li; Xiao Li; Molei Tao; Qing Qu

ID 13: Deep Networks Need Sleep – Just Like You

Dashiell R. Fairborn; Yuqing Zhu

ID 14: Curse of Dimensionality in Neural Network Optimization

Sanghoon Na; Haizhao Yang

ID 15: Understanding Representation Dynamics of Diffusion Models via Low-Dimensional Modeling

Xiao Li; Zekai Zhang; Xiang Li; Siyi Chen; Zhihui Zhu; Peng Wang; Qing Qu

ID 16: An Information Theoretic Lens on Language Model Planning

Muhammed Ustaomeroglu; Baris Askin; Carlee Joe-Wong; Gauri Joshi; Guannan Qu

ID 17: From Finite to Infinite Groups: A Polynomial-Time Algorithm for Learning with Exact Invariances

Ashkan Soleymani; Behrooz Tahmasebi; Patrick Jaillet; Stefanie Jegelka

ID 18: Data Augmentation: A Fourier Analysis Perspective

Behrooz Tahmasebi; Melanie Weber; Stefanie Jegelka

ID 19: Domain generalization: A Tale of Two ERM

Yilun Zhu; Naihao Deng; Naichen Shi; Aditya Gangrade; Clayton Scott

ID 20: Understanding In-Context Learning's Out-of-Distribution Generalization Abilities via Low-Dimensional Subspaces

Soo Min Kwon; Alec S. Xu; Can Yaras; Laura Balzano; Qing Qu

ID 21: AURORA: Global GP-UCB Controllers for Test-Time Compute with Calibrated Verifiers ($\mu(x) + \beta\sigma(x)$ over Random Fourier Features)

Jasraj Budigam

ID 22: From Cortex to Compression

Panayiotis Ketonis; Smita Krishnaswamy

ID 23: Multi-head Transformers Provably Learn Symbolic Multi-step Reasoning via Gradient Descent

Tong Yang; Yu Huang; Yingbin Liang; Yuejie Chi

ID 24: On the Transferability of Collapsed Features in Unsupervised Representation Learning

Assaf Izhar; George Attallah; Altay Unal; Akshay Rangamani

ID 25: Diffusion models learn low-dimensional distributions via subspace clustering

Peng Wang; Huijie Zhang; Zekai Zhang; Siyi Chen; Yi Ma; Qing Qu

ID 26: Understanding Generalization in Diffusion Models via Probability Flow Distance

Huijie Zhang; Zijian Huang; Siyi Chen; Jinfan Zhou; Zekai Zhang; Peng Wang; Qing Qu

ID 27: Mixture-of-Transformers Learn Faster: A Theoretical Study on Classification Problems

Hongbo Li; Qinhang Wu; Sen Lin; Yingbin Liang; Ness Shroff

ID 28: Absorb and Converge: Provable Convergence Guarantee for Absorbing Discrete Diffusion Models

Yuchen Liang; Renxiang Huang; Lifeng Lai; Ness Shroff; Yingbin Liang

ID 29: Learning to Choose or Choosing to Learn: Best-of-N vs. Supervised Fine-Tuning for Bit String Generation

Seamus Somerstep; Vinod Raman; Unique Subedi; Yuekai Sun

ID 30: KPFlow: Uncovering Implicit Constraints Imposed by Parameters in Recurrent Gradient Descent Learning

James Hazelden; Laura Driscoll; Eli Shlizerman; Eric Shea-Brown

ID 31: A Closer Look at Model Collapse: From a Generalization-to-Memorization Perspective

Lianghe Shi; Meng Wu; Huijie Zhang; Zekai Zhang; Molei Tao; Qing Qu

ID 32: Why are the logits of trained models distorted? A theory of overfitting for imbalanced classification

Jingyang Lyu; Kangjie Zhou; Yiqiao Zhong

ID 33: Beyond Sharpness - A PAC-Bayes View on the Triple Factors of Generalization

Ziyu Cheng; Xitong Zhang; Rongrong Wang