David Koplow

dkoplow@mit.edu davidkoplow.com DavidKoplow

Research Interests

Biologically inspired learning systems: What allows humans to generalize beyond the benchmark?

Education

Massachusetts Institute of Technology (MIT) February 2025 to May 2026 (expected)
MEng in Artificial Intelligence and Decision Making Cambridge, MA

- **GPA: 5.0 / 5.0**.
- Advised by Tomaso Poggio. Program fully funded by the Center for Brains, Minds and Machines and the Schwarzman College of Computing.

Massachusetts Institute of Technology (MIT) September 2020 to February 2024
BS in Artificial Intelligence (Course 6-4) and Neuroscience (Course 9) Cambridge, MA

- **GPA: 5.0 / 5.0**. Completed both majors in three and a half years.
- Relevant advanced coursework: Computational Cognitive Science (graduate), Intro to Neural Computation, Computer Vision (graduate), Parallel Processing and Scientific Machine Learning (graduate), Real Analysis (graduate), Artificial Intelligence (graduate), Natural Language Processing.

Honors and Awards

- Neo Scholar 2023, chosen as one of thirty students nationally after technical evaluation (1% acceptance).
- Phi Beta Kappa 2024, elected for top GPA and faculty recommendation.
- Tau Beta Pi 2023, inducted as top ten percent of engineering students.
- Eta Kappa Nu 2023, selected for ranking in the top quarter of EECS and department contributions.
- Member, Schwarzman College of Computing Undergraduate Advisory Group, one of two student reps for the major.
- Interact Fellow, selected for cross-disciplinary work in tech, policy, and society ($< 10\%$ acceptance).
- Brain and Cognitive Sciences Excellence Award 2023 for outstanding performance in neuroscience.
- Fully funded MEng by the Center for Brains, Minds and Machines and the Schwarzman College of Computing.
- Accepted to Y Combinator, Prod, MIT Sandbox, and MIT DHIVE for former startup, Tau Clinical.

Manuscripts

Submitted and In Preparation

Homeostatic Ubiquity of Hebbian Dynamics in Regularized Learning Rules Submitted to ICLR 2026
D. Koplow, Z. Liu, T. Poggio Manuscript

- Available on ArXiv.
- Shows a surprising similarity between deep learning and the rb. For a broad class of regularized optimization algorithms with weight decay, the effective local synaptic updates near stationarity align with classical Hebbian and anti-Hebbian rules.
- Combines analytical results with experiments on convolutional networks, transformers, and multilayer perceptrons trained on vision and language tasks on MIT OpenMind slurm cluster.
- Finds that from gradient and weight decay signals alone, it is difficult to distinguish which learning rule generated the updates, and that individual neurons often display predominantly Hebbian or anti-Hebbian dynamics.
- Presented as posters at two ICML 2025 workshops on learning dynamics and biologically inspired learning.

Zero Knowledge Learning Preprint, targeting ICML 2026
D. Koplow, S. Layton Manuscript

- Early draft available at davidkoplow.com/zkl
- Introduces a gradient-free post-training framework for large language models that uses evolution strategies to perturb parameters, evaluate scalar feedback, and apply sparse structural updates.
- Implements the method as a PyTorch-compatible library with Ray-based distributed sampling and minibatched evolution strategy updates that can run during inference without noticeable slowdown.

- Improves sample efficiency over GRPO and speeds up wall clock time by 10-20 times compared with prior evolution-strategy baselines on reasoning and alignment benchmarks while preserving base-task performance.
- Enables privacy-preserving post-training under zero data retention agreements by storing only scalar rewards rather than user inputs or generations.

Empirical Bounds on "Superintelligence" Targeting ICML 2026

D. Koplow, T. Galanti, T. Poggio Manuscript

- Designs Boolean circuit benchmark families with controlled depth, fan-in, and compositional structure to study when language models construct useful intermediate computations.
- Builds a generator and evaluation pipeline where models predict both circuit outputs and internal node values using vLLM and OpenRouter to compare frontier models.
- Develops checks that verify whether partial solutions are on the right track and separate tasks solvable by scaling search and compositional prompting from those that require new architectures or supervision.
- Aims to place empirical bounds on what scaling alone can achieve on reasoning-heavy tasks and to clarify where current learning paradigms fall short.

Technical Reports and Student Papers

"No, GPT4 can't ace MIT" Rebuttal Summer 2023

R. Chowdhuri, N. Deshmukh, D. Koplow Technical report

- Wrote a viral rebuttal (**3.4M** views) to "Exploring the MIT Mathematics and EECS Curriculum Using Large Language Models" after it went viral on Twitter. It led to the withdrawal of the paper and an article in the Chronicle of Higher Education. The pre-print was originally co-authored by **15** researchers, including **4** faculty from MIT.
- I was asked to speak about this to over **100** researchers and educators at a conference on LLMs in education run by MIT. List of other speakers. I also gave this talk to Tomaso Poggio's lab and was a panelist for MIT Generative AI Week.

Optimizing Eta Kappa Nu Tutor Assignments 2023

D. Koplow, R. Chowdhuri Project Report

- Available on my website.
- Developed an algorithm for matching tutors and students in MIT's official EECS tutoring service that raised the match rate by about 40 percent over historical levels.
- Used structured representations of student needs and tutor expertise, augmented by simple language model-based tools, to increase coverage of demonstrated demand.

Automatic Shadow-Based Satellite Feed Verification 2022

D. Koplow, A. Cai Project Report

- Available on my website.
- Proposed a pipeline that uses LiDAR and satellite imagery to verify satellite position, image location, and capture time by analyzing cast shadows.
- Demonstrated a dataset of major cities that the method can recover the image time of day within about one hour under varied conditions, and can detect certain inconsistencies in real satellite imagery.

Exploration of the Graph Isomorphism Problem 2021

D. Koplow Technical report

- Available on my website.
- Proposed and empirically evaluated an $O(n^4)$ algorithm for graph isomorphism that succeeded on several graph families.
- Constructed challenging instances of strongly regular graphs and analyzed failure modes under the guidance of Professor Virginia Williams.

Talks and Presentations

- *Ubiquity of Hebbian Dynamics in Complex Learning Rules.* Poster, ICML 2025 High-dimensional Learning Dynamics (HiLD) and Methods and Opportunities at Small Scale (MOSS) Workshops.
- *No, GPT 4 Has Not Aced MIT.* Invited talk at MIT conference on large language models in education, 2023.
- Talk on GPT 4 curriculum performance to the Poggio Lab at the Center for Brains, Minds and Machines, MIT, 2023.
- Panelist, MIT Generative AI Week, 2023.

Research Experience

Center for Brains, Minds and Machines (MIT) February 2025 to Present

MEng Researcher and Teaching Assistant Cambridge, MA

- Work with Tomaso Poggio and Ziyin Liu on the theory of learning dynamics in deep networks and their relation to biological learning rules.
- Developed the theoretical framework behind *Ubiquity of Hebbian Dynamics in Complex Learning Rules*, demonstrating that classical Hebbian dynamics emerge generically in regularized optimization near stationarity.
- Designed and ran empirical studies across architectures and datasets to validate the theory, including experiments on convolutional networks and transformers with various forms of noise and weight decay.
- Built a biologically-inspired post-training framework for large language models using evolution strategies that operate with only scalar feedback, enabling privacy-preserving updates without gradient access and without storing user inputs.
- Constructed a Boolean circuit benchmark suite to test when large language models produce consistent and useful intermediate reasoning steps as circuit depth and complexity grow, including tasks that separate partial and complete solutions.

MIT Bear Lab Spring 2022

Student Researcher Cambridge, MA

- Developed a Julia based toolkit for constructing and training spiking neural networks, modeling neurons and synapses via coupled differential equations.
- Implemented biologically-inspired neuron models capable of sub-threshold resonance and surrogate gradient-based learning and studied their role in circuit-level computation.
- Linked automatic differentiation and numerical methods to questions in neural computation, connecting earlier work in the Julia Lab to models of learning in the brain.

Independent Research with Professor Virginia Williams Summer 2021 to Fall 2021

Student Researcher Cambridge, MA

- Investigated the graph isomorphism problem after encountering it in an algorithms course and pursued the project independently alongside a summer internship under the guidance of Professor Virginia Williams.

MIT Julia Lab Spring 2019 to Spring 2020

Student Researcher Cambridge, MA

- Contributed to research on neural differential equations and pharmacometric modeling under the mentorship of Chris Rackauckas.
- Implemented reverse mode automatic differentiation and vector Jacobian products in Julia for components of deep nonlinear mixed effects models and differential equation solvers.
- Improved numerical stability and efficiency of gradient-based fitting routines used in pharmacometric simulations that informed later production code in the Pumas Julia package.

Professional and Industry Experience

Reducto AI September 2024 to January 2025

Resident Machine Learning Research Engineer San Francisco, CA

- Joined as the second employee shortly after the seed round and contributed through a subsequent Series A.
- Trained state-of-the-art vision language models on internet-scale corpora of complex enterprise documents to improve long-tail parsing accuracy.
- Authored early versions of the split endpoint and agentic extract, which have become two of Reducto's core offerings.
- Designed and implemented a training orchestration system that improved compute efficiency on AWS and Kubernetes and presented it at a board meeting.
- Supported enterprise customers, including a Fortune 10 client, to prioritize features, triage issues, and deliver production systems.

Tau Clinical January 2023 to May 2024

Co-Founder Cambridge, MA

- Co-founded a pre-seed company focused on automated clinical data harmonization and standards conversion for FDA submissions.
- Conducted market research with more than one hundred stakeholders across pharmaceutical companies and regulatory bodies to validate product requirements and constraints.
- Developed prototype software and a mock clinical trial pipeline used to demonstrate potential time savings and error reductions to prospective customers.
- Secured acceptance to Y Combinator, Prod, MIT Sandbox, and MIT DHIVE.

BrainQ Technologies June 2022 to August 2022
Machine Learning Engineer Intern Jerusalem, Israel

- Applied transformers, variational autoencoders, and related methods to EEG data from stroke patients treated with a neurorehabilitation device.
- Performed data augmentation and feature engineering with MNE and Braindecode to improve robustness of treatment effect detection.
- Worked with clinical and scientific teams to turn model outputs into biomarkers that informed clinical trial design.

Pumas AI June 2021 to August 2021
Machine Learning Engineer Intern Cambridge, MA

- Implemented reverse mode automatic differentiation and vector Jacobian products in the Pumas Julia package for deep nonlinear mixed effects models.
- Integrated additional pharmacometric models into the clinical test suite and improved numerical robustness in typical trial scenarios.

Stem Development Inc. June 2020 to August 2020
Software Developer (Consultant) Boston, MA

- Built more than 140 interactive JavaScript simulations in physics, biology, and chemistry for a digital textbook series used to teach all high schoolers in a Middle Eastern country as part of its national curriculum.
- Refactored the simulation architecture to support rapid curriculum changes and helped the company retain a seven-figure national contract.

Teaching and Leadership

Massachusetts Institute of Technology 2025 to Present
Teaching Assistant, Natural Language Processing Cambridge, MA

- Teaching assistant for MIT Natural Language Processing (Fall 2025) for a class of about 300 students.
- Lead weekly office hours, support lecture material, and help create, refine, and grade assignments and projects.
- Selected to serve as a teaching assistant for Representation, Inference and Reasoning in AI in Spring 2026.

AI@MIT, Massachusetts Institute of Technology 2020 to 2024
Co-President and Labs Lead Cambridge, MA

- Managed organizational logistics and communication for a few hundred students and dozens of active project teams.
- Led AI@MIT Labs, mentoring roughly 40 students through a project incubator focused on projects in computer vision, generative AI.
- Developed sponsorship relationships with industry partners, including OpenAI's first student group sponsorship.
- Designed and built the AI@MIT website in Next.js with integrated scheduling and calendar tools.
- Started Permeate, an initiative that matches MIT students with local non-profits to solve technical problems.

Eta Kappa Nu (HKN) Honor Society, MIT Chapter 2022 to 2024
Co-President Cambridge, MA

- Oversaw MIT's official EECS honor society and tutoring service, including the largest student-run repository of class evaluations.
- Wrote a matching algorithm that improved tutor and student match rates by about 40 percent on historical data.

Technical Skills

Python Used in projects, especially Numpy, Pandas, PyTorch, Scikit-learn, OpenCV, and Plotly libraries.
Julia Language of choice, experience with DifferentialEquations, Flux, Zygote, and Diffractor packages.
TypeScript Experience with React/Next.js, Flask, HTML, and CSS to build a number of websites.