brando9@stanford.edu
Stanford, CA 94305

Brando Miranda

Gates Bldg., 353 Jane Stanford Way

Google Scholar | Website | X (formerly Twitter) | LinkedIn | Stanford Profile | MIT CBMM Profile

Research Interests

AI for formal mathematics, Lean 4 theorem proving, theorem generation, autoformalization, and end-to-end formal verification of software and mathematics; data-centric machine learning for foundation and frontier models; alternative architectures/methods/algorithms (e.g. energy-based models) toward Artificial General Intelligence (AGI).

Education

Ph.D. in Computer Science *2022-2026 (expected)*
Stanford University GPA: 4.045/4.0
Advisor: Prof. Sanmi Koyejo, Stanford Trustworthy AI Research Group (STAIR).

Master of Engineering in Electrical Engineering and Computer Science *2014-2016*
Massachusetts Institute of Technology GPA: 4.8/5.0
Advisor: Prof. Tomaso Poggio, Center for Brains, Minds and Machines (CBMM).
Thesis Title: *Function Approximation with Deep Neural and Gaussian Networks.*

Bachelor of Science, Computer Science and Engineering *2010-2014*
Massachusetts Institute of Technology Minors: Mathematics, Music

Professional Experience

Stanford University - Stanford, CA *September 2022 - 2026 (expected)*
Ph.D. Student in Computer Science. Advisor: Professor Sanmi Koyejo
Research in AI for formal mathematics, Lean 4 verification benchmarks, contamination-resistant LLM evaluation, and data-centric machine learning for foundation models; mentored a Stanford CS CURIS undergraduate research intern and student collaborators on autoformalization and formal verification.

Amazon Web Services (AWS) - Cupertino, CA *June 2024 - September 2024*
Applied Scientist Intern
Applied LLMs to Lean 4 theorem proving and formal verification workflows.

Morph Labs - Remote *October 2023 - December 2023*
Machine Learning Research Scientist Consultant
Key contributor to Morph Prover v0 7b and Moogle.ai; designed and built the embedding-based vector database powering Moogle's semantic search over Mathlib.

Wise Agents - Stanford Spin-out *2023*
AI Research Consultant
Advised a Stanford spin-out on AI-agent workflows for sales-performance applications.

IBM Research - Yorktown Heights, NY *May 2022 - August 2022*
Graduate Research Intern
Conducted machine-learning research for program synthesis.

University of Illinois Urbana-Champaign - Urbana-Champaign, IL *September 2018 - May 2022*
Graduate Student. Advisor: Professor Sanmi Koyejo
Research on meta-learning, task and data diversity; mentored undergraduate and graduate collaborators on data quality, meta-learning, and evaluation methodology.

IBM Research - Yorktown Heights, NY *May 2021 - August 2021*
Graduate Research Intern
Conducted machine-learning research in program synthesis and code generation.

MIT CBMM (Center for Brain Minds & Machines) - Cambridge, MA *June 2015 - September 2018*
Research Assistant. Advisor: Professor Tomaso Poggio
Deep-learning theory and generalization research, including IJAC and ICML workshop publications; mentored undergraduate collaborators through MIT CBMM's Engineering of Intelligence team.

Rackspace - San Antonio, TX *June 2014 - August 2014*
Software Engineering Intern
Built a concurrent Redis proxy with additional security functionality.

Adobe - San Jose, CA *June 2013 - August 2013*
Software Engineering Intern — ML/NLP for Spanish Sentiment Analysis
Built Spanish sentiment-analysis NLP methods for negation handling and automatic training-data extraction.

Research Artifacts & Systems

- **VeriBench**: End-to-end Lean 4 benchmark for evaluating AI-generated code with formal verification proofs; includes a public Colab tutorial and Harbor Hub release for containerized, held-out evaluation.
- **VeriBench-FTP**: Formal theorem-proving benchmark in Lean 4 for code verification.
- **VeriBench-DT**: Trustworthy agentic autoformalization benchmark verifying original-code equivalence via differential testing across Python source, gold Lean reference, and model output (Three-Way Semantic Triangulation).
- **VeriBench-Deps**: Repository-level autoformalization benchmark for multi-file Python programs into Lean 4; introduces the Axiom Trust Boundary (ATB) metric and a 2×2 provenance/effect dependency taxonomy.
- **lean-ebm**: Energy-Based Models for Lean 4 theorem proving — Stanford AI for Lean Club project unifying policy and value via Langevin-dynamics tactic search; pairs a synthetic-data track with an EBM-based MCTS proof-search fallback.
- **Putnam-AXIOM**: Functional and static benchmark for higher-level LLM mathematical reasoning.
- **Putnam-AXIOM-Grading**: Human-graded 1,000-solution benchmark for Putnam-style partial-credit mathematical evaluation.
- **Pantograph**: Machine-to-machine interface for Lean 4 theorem proving, reasoning, and data extraction.
- **Morph Prover v0 7b and Moogle.ai**: Lean 4 proof model and verified-code search engine developed with Morph Labs; designed and built the embedding-based vector database powering Moogle's semantic search over Mathlib.
- **AlphaApollo**: Agentic reasoning framework for tool-integrated reasoning, agentic post-training, and self-evolving verification workflows.
- **AlphaDiana**: Harness-aware evaluation system for open agents on verifiable reasoning tasks with trajectory logging.
- **ultimate-utils**: Reusable ML and research-engineering utility library for experiment workflows.

Publications

Google Scholar: 3,014+ citations; h-index 17; i10-index 22; as of May 14, 2026.

Refereed Publications

(1) **B. Miranda**, S. Daruru, E. S. Hersch, Z. Zhou, A. Nie, D. Amrollahi, L. Aniva, I. Mlauzi, et al., *VeriBench: End-to-End Formal Verification Benchmark for AI Coding Agents in Lean 4.*
(**Preprint. 2026** & **Run Tutorial (Colab). 2026** & **Harbor Hub. 2026** & **Blog. 2026** & **2nd Workshop on AI for Math at International Conference on Machine Learning (ICML). 2025**)

(2) Z. Zhou, X. Lu, C. Cao, **B. Miranda**, T. Liu, B. Han, S. Koyejo, *CoDaPO: Confidence and Difficulty-Adaptive Policy Optimization for LLM Reasoning.*
(**International Conference on Machine Learning (ICML), Main Track. 2026** & **International Conference on Learning Representations (ICLR) Workshop on Lifelong Agents: Learning, Aligning, Evolving. 2026** & **2nd Workshop on AI for Math at International Conference on Machine Learning (ICML). 2025**)

(3) R. Schaeffer, N. Levi, **B. Miranda**, S. Koyejo, *Pretraining Scaling Laws for Generative Evaluations of Language Models.*
(**International Conference on Learning Representations (ICLR), Main Track. 2026**)

(4) E. Chen, A. Gulati, **B. Miranda**, Z. Tang, S. Koyejo, *Rethinking LLM Judges: Chain-of-Thought and Multi-Step Pipelines for Math Grading.*
(**International Conference on Learning Representations (ICLR) Workshop on Logical Reasoning of Large Language Models. 2026**)

(5) R. Schaeffer, H. Schoelkopf, **B. Miranda**, G. Mukobi, V. Madan, A. Ibrahim, H. Bradley, S. Biderman, S. Koyejo, *Why Has Predicting Downstream Capabilities of Frontier AI Models with Scale Remained Elusive?.*
(**International Conference on Machine Learning (ICML). 2025** & **International Conference on Machine Learning (ICML) Workshop on Trustworthy Multi-modal Foundation Models and AI Agents (TiFA). 2024**)

(6) **B. Miranda***, A. Gulati*, E. Chen*, E. Xia*, K. Fronsdal*, B. de Moraes Dumont, S. Koyejo, *Putnam-AXIOM: A Functional & Static Benchmark for Measuring Higher Level Mathematical Reasoning in LLMs.*
(**International Conference on Machine Learning (ICML), Main Track. 2025** & **Neural Information Processing Systems (NeurIPS) Workshop on Mathematics and AI (MATH-AI). 2024**) (*equal contribution)

(7) L. Aniva, C. Sun, **B. Miranda**, C. Barrett, S. Koyejo, *Pantograph: A machine-to-machine interaction interface for advanced theorem proving, high level reasoning, and data extraction in Lean 4.*
(**International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). 2025**)

(8) R. Schaeffer, **B. Miranda**, J. Kazdan, K. Z. Liu, A. M. Ahmed, N. Mireshghallah, S. Koyejo, *Causally Quantifying the Effect of Test Set Contamination on Generative Benchmarks.*
(**Neural Information Processing Systems (NeurIPS) Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling. 2025**)

(9) R. Schaeffer, J. Kazdan, Y. Denisov-Blanch, **B. Miranda**, M. Gerstgrasser, et al., *Position: Machine Learning Conferences Should Establish a "Refutations and Critiques" Track.*
(**Advances in Neural Information Processing Systems 38 (NeurIPS), Position Paper Track Oral. 2025**)

(10) S. Barkallah, S. Daruru, **B. Miranda**, L. Aniva, A. Nie, S. Koyejo, *VeriBench-FTP: A Formal Theorem Proving Benchmark in Lean 4 for Code Verification.*
(**5th Neural Information Processing Systems (NeurIPS) Workshop on Mathematical Reasoning and AI. 2025**)

(11) R. Schaeffer, K. Liu, **B. Miranda**, A. M. Ahmed, N. Mireshghallah, S. Koyejo, *The Contamination Paradox: Why Test Set Leakage Can Be Both Potent and Negligible.*
(**Neural Information Processing Systems (NeurIPS) Workshop on Evaluating the Evolving LLM Lifecycle. 2025**)

(12) R. Schaeffer, D. Valentine, L. Bailey, J. Chua, et al., **B. Miranda**, et al., S. Koyejo, E. Perez, *Failures to Find Transferable Image Jailbreaks Between Vision-Language Models.*
(**International Conference on Learning Representations (ICLR). 2024** & **Neural Information Processing Systems (NeurIPS) Workshop on Red Teaming Generative AI. 2024**)

(13) R. Schaeffer, M. Khona, S. Chandra, M. Ostrow, **B. Miranda**, S. Koyejo, *Does Maximizing Neural Regression Scores Teach Us About The Brain?.*
(**Neural Information Processing Systems (NeurIPS) Workshop on Unifying Representations in Neural Models (UniReps), 2nd Edition. 2024**)

(14) K. Chawla, A. Sahai, M. DePavia, S. Sundar, **B. Miranda**, *Quantifying the Importance of Data Alignment in Downstream Model Performance.*
(**International Conference on Learning Representations (ICLR) Workshop on Data-Centric Machine Learning Research (DMLR). 2024**)

(15) A. Gulati, D. Ladsaria, S. Mishra, J. Sidhu, **B. Miranda**, *An Evaluation Benchmark for Autoformalization in Lean4.*
(**International Conference on Learning Representations (ICLR) Tiny Papers Track, Second Edition. 2024**)

(16) K. Chawla, A. Sahai, M. DePavia, **B. Miranda**, *A Systematic Study of the Role of Data Quality and Alignment for Fine-tuning LLMs for Enhanced Autoformalization.*
(**International Conference on Learning Representations (ICLR) Tiny Papers Track, Second Edition. 2024**)

(17) R. Schaeffer, **B. Miranda**, S. Koyejo, *Are Emergent Abilities of Language Models a Mirage?.*
(**Neural Information Processing Systems (NeurIPS), Main Track. 2023**)

(18) **B. Miranda***, A. Lee*, P. Yu, S. Koyejo, *Beyond Scale: the Diversity Coefficient as a Data Quality Metric Demonstrates LLMs are Pre-trained on Formally Diverse Data.*
(**International Conference on Machine Learning (ICML) Workshop on Data-Centric Machine Learning. 2023** & **International Conference on Machine Learning (ICML) Workshop on Deployable Generative AI. 2023**) (*equal contribution)

(19) **B. Miranda**, P. Yu, Y. Wang, S. Koyejo, *The Curse of Zero Task Diversity: the Failure of Transfer Learning to Outperform MAML and their Empirical Equivalence.*
(**Neural Information Processing Systems (NeurIPS) Workshop on Meta-Learning. 2022**)

(20) T. Poggio, A. Banburski, Q. Liao, **B. Miranda**, L. Rosasco, J. Hidary, *Weight and Batch Normalization implement Classical Generalization Bounds.*
(**International Conference on Machine Learning (ICML) Workshop. 2019**)

(21) D. Kita, **B. Miranda**, H. Lin, J. Michon, D. Favela, J. Hu, *High-resolution on-chip digital Fourier transform spectroscopy.*
(**Conference on Lasers and Electro-Optics (CLEO): Science and Innovations. 2018**)

(22) D. Kita, **B. Miranda**, D. Favela, D. Bono, J. Michon, H. Lin, T. Gu, J. Hu, *High-performance and scalable on-chip digital Fourier transform spectroscopy.*
(**Nature Communications 9 (1), 4405. 2018**)

(23) T. Poggio, H. Mhaskar, L. Rosasco, **B. Miranda**, Q. Liao, *Why and when can deep-but not shallow-networks avoid the curse of dimensionality: A review.*
(**International Journal of Automation and Computing (IJAC). 2017**)

Preprints and Technical Reports

(1) D. Amrollahi, M. Karimi, **B. Miranda**, L. Aniva, C. Sun, C. Barrett, S. Koyejo, *AI Coding Benchmarks Need Proofs, Not Just Tests.*
(**Preprint 2026**)

(2) E. Obbad, **B. Miranda**, D. L. W. Hall, R. Schaeffer, S. Koyejo, P. Liang, *Curating High Quality Pretraining Data for Language Models via Compression Ratios.*
(**Preprint 2026**)

(3) Z. Zhou, C. Cao, X. Feng, X. Li, Z. Li, X. Lu, J. Yao, W. Huang, T. Cheng, et al., **B. Miranda**, *AlphaApollo: A System for Deep Agentic Reasoning.*
(**Preprint 2025, arXiv:2510.06261**)
(4) R. Schaeffer, N. Levi, A. Kirsch, T. Guenais, **B. Miranda**, E. Obbad, S. Koyejo, *Evaluating the Robustness of Chinchilla Compute-Optimal Scaling.*
(**Preprint 2025, arXiv:2509.23963**)

(5) W. Chan, M. Souliman, J. Nordhagen, **B. Miranda**, E. Obbad, S. Koyejo, *Lean-ing on Quality: How High-Quality Data Beats Diverse Multilingual Data in Autoformalization.*
(**Preprint 2025, arXiv:2502.15795**)

(6) K. Selva, S. Vittayaareekul, **B. Miranda**, *Exploring the Efficacy of Meta-Learning: Unveiling Superior Data Diversity Utilization of MAML Over Pre-training.*
(**Preprint 2025, arXiv:2501.08506**)

(7) E. Obbad, I. Mlauzi, **B. Miranda**, R. Schaeffer, K. Obbad, S. Bedi, S. Koyejo, *ZIP-FIT: Embedding-Free Data Selection via Compression-Based Alignment.*
(**Preprint, arXiv 2024**)

(8) J. Rotella, Z. Qin, A. Z. H. Yang, **B. Miranda**, M. E. A. Seddik, J. Zuo, H. Hacid, et al., *Synthetic Theorem Generation in Lean.*
(**Preprint 2024**)

(9) **B. Miranda** (ML Research Scientist Consultant) & Morph Labs Team, *Morph Prover v0 7b: The 1st Frontier Model for the Lean 4 Formal Verification Programming Language.*
(**Technical Report and Model Card. 2024**)

(10) **B. Miranda**, P. Yu, S. Goyal, Y. Wang, S. Koyejo, *Is Pre-training Truly Better Than Meta-Learning?.*
(**Preprint 2023**)

(11) **B. Miranda**, A. Shinnar, V. Pestun, B. Trager, *Transformer Models for Type Inference in the Simply Typed Lambda Calculus: A Case Study in Deep Learning for Code.*
(**Preprint 2023**)

(12) **B. Miranda**, Y. Wang, S. Koyejo, *Does MAML Only Work via Feature Re-use? A Data Set Centric Perspective.*
(**Preprint 2021**)

(13) **B. Miranda**, *An Empirical Study of Meta-Learning: a step towards rigorously understanding meta-learning algorithms.*
(**Preprint 2020**)

(14) A. Banburski, Q. Liao, **B. Miranda**, L. Rosasco, B. Liang, J. Hidary, T. Poggio, *Theory III: Dynamics and Generalization in Deep Networks.*
(**Preprint 2019**)

(15) Q. Liao, **B. Miranda**, A. Banburski, J. Hidary, T. Poggio, *A Surprising Linear Relationship Predicts Test Performance in Deep Networks.*
(**Preprint 2018**)

(16) C. Zhang, Q. Liao, A. Rakhlin, **B. Miranda**, N. Golowich, T. Poggio, *Theory of Deep Learning IIb: Optimization Properties of SGD.*
(**Preprint 2018**)

(17) T. Poggio, Q. Liao, **B. Miranda**, A. Banburski, X. Boix, J. Hidary, *Theory IIIb: Generalization in Deep Networks.*
(**Preprint 2018**)

(18) T. Poggio, K. Kawaguchi, Q. Liao, **B. Miranda**, L. Rosasco, X. Boix, J. Hidary, M. Mhaskar, *Theory of Deep Learning III: explaining the non-overfitting puzzle.*
(**Preprint 2017**)

Awards & Honors

- **Inaugural Veritas Scholar, Math Inc.** (offered, declined; 2026) — selected for the inaugural Veritas Scholar offer, recognizing frontier work in AI-assisted formalization and verified mathematics.
- **ICML 2026 Silver Reviewer** (May 2026) — top reviewer recognition signed by the ICML 2026 Program Chairs.
- **Oral Recommendation (Top 1)**, 2nd AI for Math Workshop @ ICML 2025 — for *VeriBench: End-to-End Formal Verification Benchmark for AI Code Generation in Lean 4* (Miranda et al.); reviewer recommendation "Accept (Oral, Top 1)".
- **Pear AI Researchers Circle** (July 2025) — selective AI-researcher network convened by Pear VC; invited following the final round (R3) of the Pear AI Researcher Grant program.
- **ICML Workshop on Trustworthy Multi-modal Foundation Models and AI Agents (TiFA) — Outstanding Paper Award** (July 2024) — for "Why Has Predicting Downstream Capabilities of Frontier AI Models with Scale Remained Elusive?" (Schaeffer, Schoelkopf, Miranda, et al.).
- **NeurIPS Outstanding Main Track Paper Award** (December 2023) — top 0.4% of NeurIPS submissions; only 2 main-track papers selected. For "Are Emergent Abilities of Large Language Models a Mirage?" (Schaeffer, Miranda, Koyejo).
- **EDGE Scholar**, Stanford University (September 2022) — Stanford fellowship supporting first-generation / low-income PhD students with additional mentorship and stipend.
- **Stanford School of Engineering Fellowship** (September 2022) — multi-year departmental fellowship for incoming Stanford engineering PhD students.
- **Honorable Mention, Ford Foundation Predoctoral Fellowship** (2020, 2021) — national fellowship recognizing PhD applicants with potential to diversify the U.S. professoriate.
- **Best Research Project Award**, UIUC graduate course CS 598 "Learning to Learn" (Prof. Y. Wang, December 2020) — course-level award for the top research project.
- **HSF Scholar**, Hispanic Scholarship Fund (2020) — competitive national merit fellowship for Hispanic graduate students.
- **Computer Science Excellence Saburo Muroga Endowed Fellowship**, UIUC (2019-2020) — top departmental fellowship for outstanding incoming CS PhD students.

- **Most Cited Paper Certificate**, International Journal of Automation & Computing (IJAC, December 2019) — for "Why and when can deep- but not shallow-networks avoid the curse of dimensionality: a review" (Poggio, Mhaskar, Rosasco, Miranda, Liao).
- **Sloan Scholar**, Alfred P. Sloan Foundation Minority Ph.D. (MPHD) Program (2018-2019) — prestigious national fellowship supporting underrepresented STEM PhD students.
- **Grainger Engineering SURGE Fellowship**, UIUC (2018-2019) — multi-year college-level fellowship for diverse engineering PhD students.
- **MIT Mitchell B. Kaufman Memorial Scholarship** (2012-2013, 2013-2014) — MIT undergraduate scholarship support.
- **MIT Eugene and Margaret McDermott Scholarship** (2012-2013, 2013-2014) — MIT undergraduate scholarship support.
- **Chopper Trading, LLC, Best Strategy Report Award**, MIT BattleCode AI competition (2013) — top finisher among ∼300 teams.
- **Greengates Scholarship** (30% 2007, 50% 2008, 100% 2009, 100% 2010) — merit scholarship that increased to full-tuition support.
- **High Achievement Prize Award**, Greengates School (2007, 2008, 2009, 2010) — equivalent to Valedictorian.
- **Best All-Round Student Award**, Greengates School (2010) — school-wide award recognizing the top all-around student.
- **Achievement Prize Award**, Greengates School (2006) — school recognition for academic achievement.
- **Certificate for Progress Award**, Greengates School (2005) — school recognition for academic progress.

Media Coverage

- **Hacker News / Y Combinator (January 2025)**: *Putnam-AXIOM* benchmark posted on Hacker News — front-page attention on Y Combinator's tech-news forum
- **Aran Komatsuzaki / @arankomatsuzaki on X (June 2024)**: shared the downstream-capabilities paper on X to his AI-research audience
- **AIhub (April 2024)**: interview with Brando Miranda on the NeurIPS Outstanding Paper Award-winning emergent-abilities paper
- **American Scientist (March-April 2024)**: "Is There an AI Metrics Mirage?"
- **White House Economic Report of the President (March 2024)**: cited Miranda et al.'s work on emergent abilities in the White House's annual economic-policy report
- **Andrew Ng (March 2024)**: Endorsed the emergent-abilities paper as evidence for smooth, predictable AGI development
- **Quanta Magazine (February 2024)**: "How Quickly Do Large Language Models Learn Unexpected Skills?"
- **Stanford AI Lab Blog (ICML 2023)**: featured the *Beyond Scale* and *Is Pre-training Truly Better Than Meta-Learning?* papers
- **AK / @_akhaliq on X (June 2023)**: shared the *Beyond Scale* paper on X to a large AI-research audience (68.9K views)
- **The New York Times (June 2023)**: "Silicon Valley Confronts the Idea That the 'Singularity' Is Here"
- **The Register (May 2023)**: "LLM emergent behavior written off as 'a mirage' by study"
- **Stanford Institute for Human-Centered Artificial Intelligence (HAI) (May 2023)**: "AI's Ostensible Emergent Abilities Are a Mirage"
- **Y Combinator News (May 2023)**: "Are emergent abilities of large language models a mirage?"
- **Forbes (May 2023)**: "AI 'Emergent Abilities' Are A Mirage, Says AI Researcher"
- **Vice / Motherboard (May 2023)**: "Scary 'Emergent' AI Abilities Are Just a 'Mirage' Produced by Researchers, Stanford Study Says"
- **Additional coverage**: Medium, Hacker News, NeurIPS blog, Reddit, and other community discussions

Invited & Contributed Talks

Applications of Trustworthy Machine Reasoning with AI Coding Agents
AAAI 2026 Tutorial "Trustworthy Machine Reasoning with Foundation Models" (Part IV, 50 min). Co-presented with B. Han, Z. Zhou, C. Cao, P. Lu, S. Koyejo. *Singapore, January 20, 2026.*

Emergent Abilities in Large Language Models: Mirage or an Elusive Predictive Frontier?
Hong Kong Baptist University, Department of Computer Science Seminar (2026 Series). *Hong Kong, January 19, 2026.*

Applications of Trustworthy Machine Reasoning with AI Coding Agents
Hong Kong Baptist University, Department of Computer Science Seminar (invited). *Hong Kong, January 2026.*

VeriBench: End-to-End Formal Verification Benchmark for AI Code Generation in Lean 4
Lawrence Livermore National Laboratory Reading Group (invited). *July 31, 2025.*

Intro to Lean 4, VeriBench, and Trustworthy Testing with Theorems
Prof. Azalia Mirhoseini's Lab Meeting, Stanford (invited). *2025.*

Are Emergent Abilities of Large Language Models a Mirage? — Why Has Predicting Downstream Capabilities Remained Elusive?
Lawrence Livermore National Laboratory (invited talk). *2025.*

Are Emergent Abilities of Large Language Models a Mirage?
Stanford IEEE Invited Talk. *Stanford, CA, 2023.*

Emergent Abilities of Large Language Models
Amazon Research (invited talk). *2023.*

The Curse of Low Task Diversity: On the Failure of Transfer Learning to Outperform MAML and Their Empirical Equivalence
NeurIPS Meta-Learning Workshop, Contributed Talk; slides. *New Orleans, LA, December 2022.*

Selected Writing

- **Metallica Goes Jazz: A 2010 High School Arrangement, Reanimated by AI** *May 2026*
- **Embrace AI or Be Left Behind — By People, Not Machines** *April 2026*
- **Asking the Right Question: Formal Methods as Scalable Oversight** *April 2026*
- **What I Learned Building a Correctness-Gated Multi-Agent Workflow for Research** *April 2026*

Academic Service

- Reviewer for NeurIPS 2026, ICLR 2026, ICML 2026 (**Silver Reviewer** – top reviewer recognition), TMLR 2026, ICLR 2025, ICML 2025 AI4MATH Workshop, ICLR 2024 DMLR Workshop, NeurIPS 2023 MATH-AI Workshop, ICLR 2020, and JMLR 2018

Service & Outreach

- Graduate advisor for Latinos in Computer Science (LCS) at UIUC, 2019-Present
- Undergraduate and outreach research mentor through Stanford CS CURIS (Summer 2025), DREU at UIUC, UROP at MIT, and MIT CBMM's Engineering of Intelligence Team

Teaching experience

Stanford University - Stanford, CA *2023 - 2024*
Course Assistant (3 quarters) & Instructor of Record (Spring)

- CS 197 *How to Do CS Research* with Prof. Michael Bernstein — TA'd 3 quarters; served as **Instructor of Record** for the Spring offering
- Mentored undergraduate research projects from the course that produced **2 workshop publications** (including ICML workshop tracks)

UIUC - Urbana-Champaign, IL *August 2020 - December 2020*
Graduate Teaching Assistant

- CS 446 Machine Learning
- Responsible for problem sets and exam design. Held weekly office hours.

MIT EECS and CBMM - Cambridge, MA *September 2014 - December 2016*
Graduate Teaching Assistant

- Statistical Learning Theory & Applications (9.520/6.860), Fall 2016
- Introduction to Algorithms (6.006), Spring 2016
- Design & Analysis of Algorithms (6.046), Fall 2015
- Introduction to Machine Learning (6.036), Spring 2015
- Mathematics for Computer Science (6.042), Fall 2014

MIT - Cambridge, MA *September 2012 - June 2013*
National Honor Society for Computer Science and Electrical Engineering (Eta Kappa Nu)

- Algorithms tutor for undergraduates at MIT

Leadership

Stanford AI for Lean (Lean AI Club) - Stanford, CA *2025 - Present*
Co-Founder & President — Stanford AI for Lean, a research community advancing AI for Lean theorem proving and formalizing mathematics; created public Learning Lean 4 tutorial videos with a companion code repository.

Stanford Bachata Sensual & Brazilian Zouk (SBSBZ) - Stanford, CA *September 2021 - Present*
Founder & President — founded the first Bachata Sensual & Zouk student organization at Stanford; 200+ members; 83-video YouTube lesson playlist with 25K+ total video views; Instagram.

UIUC Latinos in Computer Science (LCS) - Urbana-Champaign, IL *December 2019 - Present*
Graduate Advisor — founded the LCS Professional & Wellness Development Colloquium speaker series.

UIUC Bachata Sensual & Zouk (Latin Dance Group) - Urbana-Champaign, IL *January 2019 - December 2021*
Founder & President — founded the first Bachata Sensual & Zouk group at UIUC. Bachata Sensual Excellence Diploma, Korke & Judith *World Mastery* (Beginner, Intermediate, Advanced), 2019.

Certifications

- Coursera: *Learning How to Learn* (with Honors), UC San Diego, and *Mindshift* (with Honors), McMaster University; both by Barbara Oakley, Terrence Sejnowski, and Linda Walker
- World Mastery: Certified dance instructor; completed online training by Korke & Judith