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, autoformalization, and end-to-end formal verification of software and mathematics; contamination-resistant evaluation of LLM mathematical reasoning; data-centric machine learning for foundation and frontier models; alternative architectures (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.
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

Awards & Honors

- **Inaugural Veritas Scholar, Math Inc.** (offered, declined; 2026) — selected for 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 scholars.
- **Stanford School of Engineering Fellowship** (September 2022) — multi-year departmental fellowship for incoming engineering PhDs.
- **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" (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 Fellow**, UIUC (2019-2020) — top departmental fellowship for incoming CS PhDs.
- **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".
- **Sloan Scholar**, Alfred P. Sloan Foundation Minority Ph.D. (MPHD) Program (2018-2019) — national fellowship for underrepresented STEM PhD students.
- **Grainger Engineering SURGE Fellowship**, UIUC (2018-2019) — multi-year college-level fellowship for diverse engineering PhDs.

Selected Publications

Google Scholar: full publication list.

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) 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**)
(6) 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**)

(7) **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)

(8) 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**)

(9) 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**)

(10) 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**)

(11) 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**)

(12) 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**)

(13) 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**)

(14) 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**)

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

(16) **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)

(17) 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) 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**)

(3) 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 2024, arXiv:2410.18194**)

Research Artifacts & Systems

- **VeriBench** and **VeriBench-FTP**: Lean 4 verification benchmarks for AI-generated code and theorem-proving code verification; VeriBench includes a public Colab tutorial and Harbor Hub release
- **VeriBench-DT** and **VeriBench-Deps**: Lean 4 autoformalization benchmarks — differential-testing trustworthiness and repository-level Python→Lean with the Axiom Trust Boundary metric
- **lean-ebm**: Energy-Based Models for Lean 4 theorem proving (Stanford AI for Lean Club project)
- **Putnam-AXIOM** and **Putnam-AXIOM-Grading**: LLM mathematical reasoning benchmark and human-graded Putnam-style partial-credit benchmark
- **Pantograph**: Lean machine-to-machine theorem-proving interface
- **Morph Prover v0 7b / 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** and **AlphaDiana**: agentic reasoning and harness-aware evaluation systems for verifiable reasoning
- **ultimate-utils**: reusable ML and research-engineering utility library

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; mentored a Stanford CS CURIS intern and student collaborators on Lean 4 verification research.

Amazon Web Services (AWS) - Cupertino, CA *June 2024 - September 2024*
Applied Scientist Intern

Morph Labs - Remote *October 2023 - December 2023*
Machine Learning Research Scientist Consultant

Wise Agents - Stanford Spin-out *2023*
AI Research Consultant

IBM Research - Yorktown Heights, NY *May 2022 - August 2022*
Graduate Research Intern

University of Illinois Urbana-Champaign - Urbana-Champaign, IL *September 2018 - May 2022*
Ph.D. Student in Computer Science. Advisor: Professor Sanmi Koyejo
Research on meta-learning and data diversity; mentored student collaborators on data quality and evaluation methodology.

IBM Research - Yorktown Heights, NY *May 2021 - August 2021*
Graduate Research Intern

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

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), Singapore, January 2026
- **Emergent Abilities in Large Language Models: Mirage or an Elusive Predictive Frontier?** – Hong Kong Baptist University, Department of Computer Science Seminar, Hong Kong, January 2026
- **Applications of Trustworthy Machine Reasoning with AI Coding Agents** – Hong Kong Baptist University, Department of Computer Science Seminar, invited, January 2026
- **VeriBench: End-to-End Formal Verification Benchmark for AI Code Generation in Lean 4** – Lawrence Livermore National Laboratory Reading Group, invited, July 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** – NeurIPS Meta-Learning Workshop, Contributed Talk; slides, New Orleans, LA, December 2022

Teaching Experience

Stanford University *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; **Instructor of Record** for the Spring offering
- Mentored undergraduate research that produced **2 workshop publications** (including ICML workshop tracks)

University of Illinois Urbana-Champaign *August 2020 - December 2020*

Graduate Teaching Assistant

- CS 446 Machine Learning
- Designed problem sets and exams; held weekly office hours

Massachusetts Institute of Technology *2014-2016*

Graduate Teaching Assistant

- Statistical Learning Theory & Applications (9.520/6.860); Introduction to Machine Learning (6.036)
- Introduction to Algorithms (6.006); Design & Analysis of Algorithms (6.046)

Leadership & Service

Research Mentorship *2016 - Present*

- Mentored undergraduate and graduate students at Stanford, UIUC, and MIT CBMM, including a Stanford CS CURIS intern
- Led research teams on data quality metrics, meta-learning, and formal reasoning

Academic Service *2018 - Present*

- Co-founder and president of Stanford AI for Lean; created public Learning Lean 4 tutorial videos with a companion code repository
- 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, JMLR 2018
- Graduate advisor for Latinos in Computer Science (LCS) at UIUC
- Founded "Stanford Bachata Sensual & Brazilian Zouk" and "UIUC Bachata Sensual & Zouk" official student organizations; Stanford group has an 83-video YouTube lesson playlist with 25K+ total video views and Instagram
- Undergraduate and outreach research mentorship through Stanford CS CURIS (Summer 2025), DREU at UIUC, UROP at MIT, and MIT CBMM's Engineering of Intelligence Team