Haoze (Andrew) Wu

Email: hwu@amherst.edu

Website: https://wu-haoze.github.io/

Education

- **Ph.D.**, Stanford University, **2024**, Computer Science. Advisor: Clark Barrett.
  Thesis Title: Bridging the Gap Between Automated Logical Reasoning and Machine Learning
- **M.S.**, Stanford University, **2020**, Computer Science.
- **B.S.**, Davidson College, **2018**, Mathematics and Philosophy. *Magna Cum Laude.*

Appointment

- Assistant Professor of Computer Science, Amherst College (9/2024–)
- Adjunct Assistant Professor of Computer Science, UMass Amherst (5/2026–)
- Affiliated Researcher, Broadcom (5/2025–)

Research Overview

Automated reasoning and machine learning have both driven remarkable technological progress. However, both fields also face substantial challenges: automated reasoners struggle with scalability and abstract high-level reasoning, while deep-learning-enabled systems are not yet stable and interpretable enough to deploy in high-stakes applications. A promising avenue towards addressing these challenges, and creating safer and smarter computer systems, is to bridge the gap between these two fields. To this end, I work on **AI-AR co-design**: integrating data-driven learning and deductive reasoning to produce systems that are both powerful and reliable.

System Under Development

SpotIt: A bounded-verification-based Text-to-SQL evaluation tool.
Available at https://github.com/ai-ar-research/SpotIt-plus

Luna: A bound propagation tool for neural network verification.
Available at https://github.com/ai-ar-research/luna

Marabou: An open-source verifier for neural networks.
Available at https://github.com/NeuralNetworkVerification/Marabou

Selected Publications

1. Rocky Klopfenstein*, Yang He*, Andrew Tremante*, Yuepeng Wang, Nina Narodytska, Haoze Wu. **SpotIt: Evaluating Text-to-SQL Evaluation with Formal Verification**. *The Fourteenth International Conference on Learning Representations (ICLR'26).*
2. Haoze Wu, Clark Barrett, Nina Narodytska. **Cubing for Tuning.** *The 40th Annual AAAI Conference on Artificial Intelligence (AAAI'26).*
3. Haoze Wu, Omri Isac, Aleksandar Zeljić, Teruhiro Tagomori, Matthew Daggitt, Wen Kokke, Idan Refaeli, Guy Amir, Kyle Julian, Shahaf Bassan, Pei Huang, Ori Lahav, Min Wu, Min Zhang, Ekaterina Komendantskaya, Guy Katz, Clark Barrett. **Marabou 2.0: A Versatile Formal Analyzer of Neural Networks**. *The 37th International Conference on Computer Aided Verification (CAV'24).*
4. Haoze Wu, Clark Barrett, Nina Narodytska. **Lemur: Integrating Large Language Models in Automated Program Verification**. *The Twelfth International Conference on Learning Represen-*

tations (ICLR'24). (Oral presentation at *MATH-AI'23, BayLearn'24*)

5. Min Wu, Haoze Wu, Clark Barrett. **VeriX: Towards Verified Explainability of Deep Neural Networks**. *The 37th Conference on Neural Information Processing Systems (NeurIPS'23).*
6. Haoze Wu, Christopher Hahn, Florian Lonsing, Makai Mann, Raghuram Ramanujan, Clark Barrett. **Lightweight Online Learning for Sets of Related Problems in Automated Reasoning**. *The 23rd International Conference on Formal Methods in Computer-Aided Design (FMCAD'23).* (Best paper award nomination)
7. Haoze Wu, Clark Barrett, Mahmood Sharif, Nina Narodytska, Gagandeep Singh. **Scalable Verification of GNN-based Job Schedulers**. *Proceedings of the ACM on Programming Languages (OOPSLA'22).*
8. Haoze Wu, Aleksandar Zeljic, Guy Katz, Clark Barrett. **Efficient Neural Network Analysis with Sum-of-Infeasibilities**. *The 28th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS' 22).*

All Publications

1. Anjiang Wei, Tianran Sun, Tarun Suresh, Haoze Wu, Ke Wang, Alex Aiken. **Quokka: Accelerating Program Verification with LLMs via Invariant Synthesis**. *The Conference on Language Modeling (COLM'26).*
2. Henry LeCates, Haoze Wu. **The Luna Bound Propagator for Formal Analysis of Neural Networks**. *The Static Analysis Symposium (SAS'26).*
3. Liam Davis, Duo Zhou, Huan Zhang, Guy Katz, Clark Barrett, Haoze Wu. **Lookahead Branching for Neural Network Verification**. *The 35th International Joint Conference on Artificial Intelligence (IJCAI'26).*
4. Ido Pinto, Yizhak Yisrael Elboher, Haoze Wu, Nina Narodytska, Guy Katz. **Not All Invariants Are Equal: Curating Training Data to Accelerate Program Verification with SLMs**. *The 43rd International Conference on Machine Learning (ICML'26).*
5. Rocky Klopfenstein*, Yang He*, Andrew Tremante*, Yuepeng Wang, Nina Narodytska, Haoze Wu. **SpotIt: Evaluating Text-to-SQL Evaluation with Formal Verification**. *The Fourteenth International Conference on Learning Representations (ICLR'26).*
6. Haoze Wu, Clark Barrett, Nina Narodytska. **Cubing for Tuning.** *The 40th Annual AAAI Conference on Artificial Intelligence (AAAI'26).*
7. Pei Huang, Dennis Wei, Omri Isac, Haoze Wu, Min Wu, Clark Barrett. **Parameterized Abstract Interpretation for Transformer Verification.** *The 40th Annual AAAI Conference on Artificial Intelligence (AAAI'26).*
8. Min Wu, Xiaofu Li, Haoze Wu, Clark Barrett. **Efficiently Computing Compact Formal Explanations.** *The 40th Annual AAAI Conference on Artificial Intelligence (AAAI'26).* (Oral Presentation)
9. Omri Isac, Idan Refaeli, Haoze Wu, Clark Barrett, Guy Katz. **Proof Minimization in Neural Network Verification.** *The 27th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'26).*
10. Kenan Wood, Runtian Zhou, Haoze Wu, Hammurabi Mendes, Jonad Pulaj. Satisfiability Modulo Theories for Verifying MILP Certificates. *Journal of Symbolic Computation.*
11. Amalee Wilson, Nina Narodytska, Clark Barrett, and Haoze Wu. **Per-Instance Subproblem Generation for Strategy Selection in SMT.** *The 25th International Conference on Formal Methods in Computer-Aided Design (FMCAD'25).*

12. Yizhak Yisrael Elboher, Omri Isac, Guy Katz, Tobias Ladner, Haoze Wu. **Abstraction-Based Proof Production in Formal Verification of Neural Networks.** *The 8th International Symposium on AI Verification (SAIV'25).*

13. Lucas C. Cordeiro, Matthew L. Daggitt, Julien Girard-Satabin, Omri Isac, Taylor T. Johnson, Guy Katz, Ekaterina Komendantskaya, Augustin Lemesle, Edoardo Manino, Artjoms Šinkarovs, Haoze Wu. **Neural Network Verification is a Programming Language Challenge.** *The 34th European Symposium on Programming (ESOP'25).*

14. Udayan Mandal, Guy Amir, Haoze Wu, Ieva Daukantas, Fletcher Lee Newell, Umberto Ravaioli, Baoluo Meng, Michael Durling, Kerianne Hobbs, Milan Ganai, Tobey Shim, Guy Katz, Clark Barrett. **Safe and Reliable Training of Learning-Based Aerospace Controllers.** 2024 AIAA DATC/IEEE 43rd Digital Avionics Systems Conference (DASC'24).

15. Huang, P., Yang, Y., Wu, H., Daukantas, I., Wu, M., Jia, F., & Barrett, C. (2024, July). **Parallel Verification for δ-Equivalence of Neural Network Quantization**. *International Symposium on AI Verification (SAIV'24).*

16. Udayan Mandal, Guy Amir, Haoze Wu, Ieva Daukantas, Fletcher Lee Newell, Umberto J Ravaioli, Baoluo Meng, Michael Durling, Milan Ganai, Tobey Shim, Guy Katz, Clark Barrett. **Formally Verifying Deep Reinforcement Learning Controllers with Lyapunov Barrier Certificates.** *The 24th International Conference on Formal Methods in Computer-Aided Design (FMCAD'24).*

17. Haoze Wu, Omri Isac, Aleksandar Zeljić, Teruhiro Tagomori, Matthew Daggitt, Wen Kokke, Idan Refaeli, Guy Amir, Kyle Julian, Shahaf Bassan, Pei Huang, Ori Lahav, Min Wu, Min Zhang, Ekaterina Komendantskaya, Guy Katz, Clark Barrett. **Marabou 2.0: A Versatile Formal Analyzer of Neural Networks**. *The 37th International Conference on Computer Aided Verification (CAV'24).*

18. Haoze Wu, Clark Barrett, Nina Narodytska. **Lemur: Integrating Large Language Models in Automated Program Verification**. *The Twelfth International Conference on Learning Representations (ICLR'24).* (Oral presentation at *MATH-AI'23, BayLearn'24*)

19. Pei Huang, Haoze Wu, Yuting Yang, Ieva Daukantas, Min Wu, Yedi Zhang, Clark Barrett. **Towards Efficient Verification of Quantized Neural Networks.** *The 38th Annual AAAI Conference on Artificial Intelligence (AAAI'24).* (Oral presentation)

20. Min Wu, Haoze Wu, Clark Barrett. **VeriX: Towards Verified Explainability of Deep Neural Networks**. *The 37th Conference on Neural Information Processing Systems (NeurIPS'23).*

21. Haoze Wu, Christopher Hahn, Florian Lonsing, Makai Mann, Raghuram Ramanujan, Clark Barrett. **Lightweight Online Learning for Sets of Related Problems in Automated Reasoning**. *The 23rd International Conference on Formal Methods in Computer-Aided Design (FMCAD'23).* (Best paper award nomination)

22. Haoze Wu, Min Wu, Dorsa Sadigh, Clark Barrett. **Soy: An Efficient MILP Solver for Piecewise-Affine Systems**. *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'23).*

23. Dennis Wei, Haoze Wu, Min Wu, Pin-Yu Chen, Clark Barrett, Eitan Farchi. **Convex Bounds on the Softmax Function with Applications to Robustness Verification**. *The 26th International Conference on Artificial Intelligence and Statistics (AISTATS'23).*

24. Haoze Wu*, Teruhiro Tagomori*, Alexander Robey*, Fengjun Yang*, Corina Pasareanu, Clark Barrett, Nikolai Matni, Hamed Hassani, George J. Pappas. **Toward Certified Robustness against Real-World Distribution Shifts**. *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML'23).*

25. Haoze Wu, Clark Barrett, Mahmood Sharif, Nina Narodytska, Gagandeep Singh. **Scalable Ver-**

ification of GNN-based Job Schedulers. *Proceedings of the ACM on Programming Languages (OOPSLA'22).*

26. Tom Zelazny, Haoze Wu, Clark Barrett, Guy Katz. **On Optimizing Back-Substitution Methods for Neural Network Verification.** *The 22nd International Conference on Formal Methods in Computer-Aided Design (FMCAD'22).*

27. Abhishek Nair, Saranyu Chattopadhyay, Haoze Wu, Alex Ozdemir, Clark Barrett. **Proof-Stitch: Proof Combination for Divide and Conquer SAT Solvers**. *The 22nd International Conference on Formal Methods in Computer-Aided Design (FMCAD'22).*

28. Haoze Wu, Aleksandar Zeljic, Guy Katz, Clark Barrett. **Efficient Neural Network Analysis with Sum-of-Infeasibilities**. *The 28th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS' 22).*

29. Colin Paterson, Haoze Wu, John Grese, Radu Calinescu, Corina S. Pasareanu, Clark W. Barrett. **DeepCert: Verification of Contextually Relevant Robustness for Neural Network Image Classifier**. *The 40th International Conference on Computer Safety, Reliability and Security (SafeComp' 21).*

30. Alex Ozdemir*, Haoze Wu*, Clark Barrett. **SAT-solving in the Serverless Cloud**. *The 21st International Conference on Formal Methods in Computer-Aided Design (FMCAD'21).*

31. Guy Amir, Haoze Wu, Clark Barrett, Guy Katz. **An SMT-Based Approach for Verifying Binarized Neural Networks**. *The 27th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'2021).*

32. Christopher A. Strong, Haoze Wu, Aleksandar Zeljic, Kyle D. Julian, Guy Katz, Clark Barrett, and Mykel J. Kochenderfer. **Global Optimization of Objective Functions Represented by ReLU Networks**. *Machine Learning Journal.*

33. Ahmed Irfan, Kyle D. Julian, Haoze Wu, Clark Barrett, Mykel J. Kochenderfer, Baoluo Meng, and James Lopez. **Towards Verification of Neural Networks for Small Unmanned Aircraft Collision Avoidance**. *The 39th Digital Avionics Systems Conference (DASC'20).*

34. Haoze Wu, Alex Ozdemir, Aleksandar Zeljic, Kyle Julian, Ahmed Irfan, Divya Gopinath, Sadjad Fouladi, Guy Katz, Corina Pasareanu, and Clark Barrett. **Parallelization Techniques for Verifying Neural Networks**. *The 20th International Conference on Formal Methods in Computer-Aided Design (FMCAD'20).*

35. Jiaxuan You*, Haoze Wu*, Clark Barrett, Raghuram Ramanujan, Jure Leskovec. **G2SAT: Learning to Generate SAT Formulas**. *The 33rd Conference on Neural Information Processing Systems (NeurIPS'19).*

36. Haoze Wu, Raghuram Ramanujan. **Learning to Generate Industrial SAT Instances**. *International Symposium on Combinatorial Search (SoCS'19).*

37. Guy Katz, Derek A Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljić, David L Dill, Mykel J Kochenderfer, Clark Barrett. **The Marabou Framework for Verification and Analysis of Deep Neural Networks**. *The 32nd International Conference on Computer-Aided Verification (CAV'19).*

Others:

1. Konstantin Kaulen, Tobias Ladner, Stanley Bak, Christopher Brix, Hai Duong, Thomas Flinkow, Taylor T. Johnson, Lukas Koller, Edoardo Manino, ThanhVu H Nguyen, Haoze Wu. **The sixth International Verification of Neural Networks Competition (VNN-COMP 2025): Summary and Results**. *arXiv preprint.* 2025.

2. Christopher Brix, Stanley Bak, Taylor T Johnson, Haoze Wu. **The Fifth International Verification of Neural Networks Competition (VNN-COMP 2024): Summary and Results**. *arXiv preprint.* 2024.

Grants

1. *Towards a Practical Framework for Trustworthy LLM Code Generation*, VMware University Research Fund. PI. $300,000. April 2026.

Selected Talks

1. "Evaluating Text-to-SQL Evaluation with Formal Verification." DREAM Lab, UMass Amherst. Amherst, MA. Mar. 2026.
2. "AI meets AR." *Cornell University*. Ithaca, NY. Feb. 2026.
3. "AI meets AR." *Brown University*. Providence, RI. Dec. 2025.
4. "AI meets AR." *UT Austin*. Austin, TX. Nov. 2025.
5. "AI meets AR." *Duke University*. Raleigh, NC. Nov. 2025.
6. "AI meets AR." *Princeton University*. Princeton, NJ. Nov. 2025.
7. "AI meets AR." *Johns Hopkins University.* Baltimore, MD. Nov. 2025.
8. "AI meets AR." *Yale University*. New Haven, CT. Oct. 2025.
9. "AI meets AR." *New York University*. Manhattan, NY. Oct. 2025.
10. "AI meets AR." *PL Seminar. Harvard University*. Cambridge, MA. Oct. 2025.
11. "Lemur: Integrating LLM's into Program Verification." *Guest lecture for ECE/CS 584: Verification of Embedded and Cyberphysical Systems.* UIUC. Remote talk. Feb. 2025.
12. "Marabou 2.0: A Versatile Formal Analyzer of Neural Networks." *Machine Learning & Friends Lunch Seminar.* UMass Amherst. Amherst, MA. Nov. 2024.
13. "Lemur: Integrating LLM's into Program Verification." *BayLearn Symposium*. Sunnyville, CA. Oct. 2024.
14. "Marabou 2.0: A Versatile Formal Analyzer of Neural Networks." *GE Aerospace Research LEAD Symposium 2024*. Niskayuna, NY. Oct. 2024
15. "Formal Verification of Neural Networks." *Guest lecture for CS120: Introduction to AI Safety*. Stanford, CA. May 2024.
16. "Lemur: Integrating LLM's into Program Verification." *Amazon Web Service*. Santa Clara, CA. March 2024.
17. "Lemur: Integrating LLM's into Program Verification." *Nvidia.* Cuptertino, CA. December 2023.
18. "Soy: an efficient MILP solver for piecewise-affine systems." *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'23).* Detroit, MI. October 2023.
19. "A framework for Integrating LLM's into Program Verification." *VMware Research.* Palo Alto, CA. September 2023.
20. "Lightweight Online Learning for Sets of Related Problems in Automated Reasoning." *AHA (Agile Hardware) Retreat*. Half Moon Bay, CA. August 2023.
21. "Scalable Verification of GNN-Based Job Schedulers." *The 43rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'23).* Orlando, FL. June 2023.

22. "Toward Certified Robustness against Real-World Distribution Shifts." *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML'23).* Raleigh, NC. February 2023.
23. "Toward Certified Robustness against Real-World Distribution Shifts." *Davidson College*. Davidson, NC. February 2023.
24. "Scalable Verification of GNN-Based Job Schedulers." *The 5th Workshop on Formal Methods for ML-Enabled Autonomous Systems (FoMLAS'23).* Haifa, Israel. July 2022.
25. "Efficient Neural Network Analysis with Sum-of-Infeasibilities." *The 28th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'22).* Munich, Germany. April 2022.
26. "Formal Verification of Neural Networks." Guest lecture for CS357: Formal Methods for Computer Systems. Stanford, CA. November 2021.
27. "SAT-Solving in the Serverless Cloud." *The 22nd International Conference on Formal Methods in Computer-Aided Design (FMCAD'21)*. Virtual event. October 2021.
28. "Verifying Fairness Properties of the Decima Job Schedulers." *VMware Research.* Palo Alto. September 2020.
29. "Parallelization Techniques for Verifying Neural Networks." *The 21st International Conference on Formal Methods in Computer-Aided Design (FMCAD'20)*. Virtual event. October 2020.

Professional Services

Program Committees

CAV: 2025, FASE: 2026, FMCAD: 2025, 2026, NFM: 2024, TACAS: 2025.

Referee

- **Formal Methods Conferences**: CAV: 2019, 2021, 2022; FMCAD: 2020, 2023; NFM: 2019; SAT: 2023; TACAS: 2023; VSSTE: 2022.
- **AI/ML Conferences**: AAAI: 2023–2026; AISTATS: 2023–2024; NeurIPS: 2021–2023; ICLR: 2022–2024; ICML: 2024–25; IJCAI: 2024.
- **Other Conferences**: ICASSP'2024.
- **Journals**: Journal of Artificial Intelligence Research, Machine Learning, Machine Intelligence Research, Formal Methods for System Design.
- **Workshop**: DAV: 2023, WFVML'23.

Other service

1. External reviewer for Advanced Research + Invention Agency (ARIA), 2025, 2026
2. Co-organizer. The 5th International Verification of Neural Networks Competition (VNN-COMP'25), 2025.
3. Co-organizer. Verification Mentoring Workshop at 22nd International Symposium on Automated Technology for Verification and Analysis (ATVA'24), Kyoto, 2024.
4. Co-organizer. The 5th International Verification of Neural Networks Competition (VNN-COMP'24), 2024.
5. Session Chair. The IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Detroit, 2023.

Teaching

1. Spring 2026, Instructor. COSC-241: Artificial Intelligence. Amherst College.
2. Spring 2026, Instructor. COSC-111: Introduction to Computer Science I. Amherst College.
3. Fall 2025, Instructor. COSC-227: Neural Safety Net. Amherst College.
4. Fall 2025, Instructor. COSC-111: Introduction to Computer Science I. Amherst College.
5. Spring 2025, Instructor. COSC-345: Automated Reasoning. Amherst College.
6. Spring 2025, Instructor. COSC-111: Introduction to Computer Science I. Amherst College.
7. Fall 2024, Instructor. COSC-241: Artificial Intelligence. Amherst College.
8. Fall 2024, Instructor. COSC-111: Introduction to Computer Science I. Amherst College.
9. Fall 2023, Co-instructor. CS 257: Introduction to Automated Reasoning. Stanford University.
10. Fall 2022, Co-instructor. CS 257: Introduction to Automated Reasoning. Stanford University.
11. Fall 2021, TA. CS 357S: Formal Methods for Computer Systems. Stanford University.
12. Winter 2020, TA. CS 243: Program Analysis. Stanford University.

Awards and Honors

1. Centennial Teaching Assistant Award, Stanford, 2024.
2. Leadership in Inclusive Teaching (LIT) Fellow, Stanford University, 2023.
3. Best Paper Nominee, Formal Methods in Computer-Aided Design (FMCAD), 2023
4. 2nd Place Overall, International Verification of Neural Networks Competition, 2023
5. Enlight Foundation Graduate Fellowship, 2018-2019.
6. Davidson Computer Science Award, Davidson College, 2018.
7. DAAD RISE Scholarship, 2017.
8. R. Bruce Jackson, Jr., Mathematics Award, Davidson College, 2017.
9. Meritorious winner, COMAP'S Mathematical Contest in Modeling (MCM) 2016.

Professional Experience

2018 – 2024	Ph.D. Candidate in Computer Science Advisor: Prof. Clark Barrett	Stanford University Stanford, CA
2023	Graduate Research Intern Mentor: Nina Narodytska	VMware Research Group Palo Alto, CA
2021	Graduate Research Intern Mentors: Nina Narodytska, Mahmood Sharif, Gagandeep Singh	VMware Research Group Palo Alto, CA
2020	Graduate Research Intern, CyLab Mentor: Corina Păsăreanu, Divya Gopinath	Carnegie Mellon University Mountain View, CA
2017	Research Intern Mentors: Marko Doko, Viktor Vafeiadis	MPI-SWS Kaiserslautern, Germany

2017	Research Intern Mentors: Gopalan Nadathur, Mary Southern	U. of Minnesota, Twin Cities Minneapolis, MN