

Haoze Wu

Assistant Professor of Computer Science, Amherst College

Amherst College
Computer Science Department
25 East Drive, Amherst, MA, 01002

Phone: (650)680-6363
Email: hwu@amherst.edu
Home: amherst.edu/people/facstaff/hwu

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. Advisor: Clark Barrett.

B.S.: Davidson College, 2018, Mathematics and Philosophy. *Magna Cum Laude*.
Advisors: Raghuram Ramanujan, David Robb

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 neural networks are not 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, my research has two main threads: one studies automated reasoning techniques for formally verifying learning-enabled systems; the other studies machine learning techniques for enhancing automated reasoning tools.

System Under Development

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

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.

Publications

Refereed:

1. 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).*
2. 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).*
 3. 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*)
 4. 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)
 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, 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).*
 8. 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).*
 9. 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).*
 10. 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).*
 11. 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).*
 12. 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).*
 13. 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).*
 14. 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).*
 15. 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).*

16. 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).*
17. 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.*
18. 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).*
19. 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).*
20. 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).*
21. Haoze Wu, Raghuram Ramanujan. **Learning to Generate Industrial SAT Instances.** *International Symposium on Combinatorial Search (SoCS'19).*
22. 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).*

Selected Talks

1. "Formal Verification of Neural Networks." Guest lecture for CS120: Introduction to AI Safety. Stanford, CA. May 2024.
2. "Lemur: Integrating LLM's into Program Verification." *Amazon Web Service.* Santa Clara, CA. March 2024.
3. "Lemur: Integrating LLM's into Program Verification." *Nvidia.* Cuptertino, CA. December 2023.
4. "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.
5. "A framework for Integrating LLM's into Program Verification." *VMware Research.* Palo Alto, CA. September 2023.
6. "Lightweight Online Learning for Sets of Related Problems in Automated Reasoning." *AHA (Agile Hardware) Retreat.* Half Moon Bay, CA. August 2023.
7. "Scalable Verification of GNN-Based Job Schedulers." *The 43rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'23).* Orlando, FL. June 2023.
8. "Toward Certified Robustness against Real-World Distribution Shifts." *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML'23).* Raleigh, NC. February 2023.
9. "Toward Certified Robustness against Real-World Distribution Shifts." *Davidson College.* Davidson, NC. February 2023.
10. "Scalable Verification of GNN-Based Job Schedulers." *The 5th Workshop on Formal Methods for ML-Enabled Autonomous Systems (FoMLAS'23).* Haifa, Israel. July 2022.
11. "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.

12. “Formal Verification of Neural Networks.” Guest lecture for CS357: Formal Methods for Computer Systems. Stanford, CA. November 2021.
13. “SAT-Solving in the Serverless Cloud.” *The 22nd International Conference on Formal Methods in Computer-Aided Design (FMCAD’21)*. Virtual event. October 2021.
14. “Verifying Fairness Properties of the Decima Job Schedulers.” *VMware Research*. Palo Alto. September 2020.
15. “Parallelization Techniques for Verifying Neural Networks.” *The 21st International Conference on Formal Methods in Computer-Aided Design (FMCAD’20)*. Virtual event. October 2020.

Teaching

1. Fall 2024, Instructor. COSC-241: Artificial Intelligence. Amherst College
2. Fall 2024, Instructor. COSC-111: Introduction to Computer Science I. Amherst College
3. Fall 2023, Co-instructor. CS 257: Introduction to Automated Reasoning. Stanford University.
4. Fall 2022, Co-instructor. CS 257: Introduction to Automated Reasoning. Stanford University.
5. Fall 2021, TA. CS 357S: Formal Methods for Computer Systems. Stanford University.
6. Winter 2020, TA. CS 243: Program Analysis. Stanford University.

Professional Services

Program Committees

AAAI: 2023–2025; DAV: 2023, NFM: 2024, IJCAI: 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:** AISTATS: 2023, 2024; NeurIPS: 2021–2023; ICLR: 2022–2024; ICML: 2024.
- **Other Conferences:** ICASSP’2024.
- **Journals:** Journal of Artificial Intelligence Research, Machine Learning, Machine Intelligence Research.
- **Workshop:** WFMV’23.

Other service

1. Co-organizer. Verification Mentoring Workshop at 22nd International Symposium on Automated Technology for Verification and Analysis (ATVA’24), Kyoto, 2024.
2. Co-organizer. The 5th International Verification of Neural Networks Competition (VNN-COMP’24), 2024.
3. Session Chair. The IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Detroit, 2023.

Mentorship

1. 2023–present, Daniel Zhou’24, Davidson undergraduate
2. 2023–2024, Fletcher Newell’24, Stanford undergraduate
3. 2022–2024, Udayan Mandal, Stanford MS student
4. 2022, Luna Yang’21, Stanford undergraduate

5. 2022, Joseph Guman'24, Stanford undergraduate
6. 2022, Tobey Shim'25, Stanford undergraduate
7. 2022, Abhishek Nair, Stanford MS student
8. 2022, Sanjana Sarda, Stanford MS student
9. 2021–2022 Tom Zelazny, HUJI MS student
10. 2020, Matei Budiu, high school student

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