Ravi Mangal

Research Interests

Trustworthy AI, Machine Learning, Formal Methods, Program Verification

Education

Georgia Institute of Technology, Atlanta, Georgia, USA
Ph.D., Computer Science, Dec 2020
Advisor: Dr. Alessandro Orso

Georgia Institute of Technology, Atlanta, Georgia, USA
M.S., Computer Science, May 2012

Veermata Jijabai Technological Institute, Mumbai, India
B.Tech., Information Technology, May 2010

Work Experience

CyLab, Carnegie Mellon University, Silicon Valley, California, USA
Postdoctoral Researcher with Dr. Corina Păsăreanu **Jan 2021 - Present**

- Drove research on techniques for improving the local robustness of neural networks via training and run-time certification as part of the DARPA Guaranteeing AI Robustness Against Deception (GARD) program.
- Designed approaches for automatically extracting high-level descriptions explaining the internal behavior of neural networks and using these descriptions for formal analysis of neural networks. This research is being extended in collaboration with SRI as part of the DARPA Assured Neuro Symbolic Learning and Reasoning (ANSR) program.
- Developed techniques for formally analyzing the safety of closed-loop autonomous systems that use neural networks for perception in collaboration with researchers at University of York as part of the Assured Autonomy International Program.
- Designed algorithms for repairing neural network outputs at run-time in order to ensure compliance with user-provided safety specifications.
- Initiated collaboration with researchers at VMware Research to develop techniques for verifiable personalization of ML models in the context of federated learning.
- Authored 10 research papers (published in venues such as CAV, FASE, ICLR, RV, and TMLR), raised ∼$100K grant money, and mentored PhD and Masters students.

Georgia Institute of Technology, Atlanta, Georgia, USA
Graduate Research Assistant **Jan 2012 - Dec 2020**

- Developed new theoretical frameworks for constructing scalable, precise static program analyses.
- Designed a new approach for interactive, user-guided static program analyses by combining formal methods with probabilistic techniques.
- Developed algorithms for analyzing robustness properties of neural networks.
- Published 10 research papers in top academic conferences including AAAI, ESEC/FSE, ESOP, ICSE, OOPSLA, PLDI, POPL, and SAT.

Microsoft Research, Redmond, Washington, USA
Research Intern **May 2016 - Aug 2016**
Developed a tool to help pen-testers perform security analysis of Android apps using a new algorithm for probabilistic and interactive information-flow analysis of programs with Dr. Patrice Godefroid and Marina Polishchuk.

Google, Mountain View, California, USA
Research Intern **May 2014 - Aug 2014**
Contributed in the design and development of an industry-strength static program analysis framework for analyzing security properties of Android apps with Dr. Jayanthkumar Kannan and Dr. Domagoj Babic.

Nvidia, Santa Clara, California, USA
Software Intern **May 2011 - Aug 2011**
Built a software simulator of DisplayPort devices for stress testing GPU device drivers as a member of the GPU Resource Manager team.

Microsoft, Hyderabad, India
Software Development Engineer in Test Intern **May 2009 - Jul 2009**
Built a test status dashboard that featured real-time updates from multiple sources of software testing data for the team developing the Data Protection Manager product.

Indian Institute of Technology-Bombay, Mumbai, India
Undergraduate Researcher **May 2008 - Jul 2008**
Worked on automated speech recognition algorithms in the Digital Audio Processing lab with Dr. Preeti Rao.

Research Articles

(* indicates equal contribution, **(**α**)** indicates alphabetical ordering)

Preprints

Chi Zhang, Zifan Wang, **Ravi Mangal**, Matt Fredrikson, Limin Jia, and Corina Păsăreanu. Transfer attacks and defenses for large language models on coding tasks. *arXiv:2311.13445*, 2023

Ravi Mangal*, Klas Leino*, Zifan Wang*, Kai Hu*, Weicheng Yu, Corina Păsăreanu, Anupam Datta, and Matt Fredrikson. Is certifying ℓ_p robustness still worthwhile? *arXiv:2310.09361*, 2023

(α) Radu Calinescu, Calum Imrie, **Ravi Mangal**, Genaína Nunes Rodrigues, Corina Păsăreanu, Misael Alpizar Santana, and Gricel Vázquez. Discrete-event Controller Synthesis for Autonomous Systems with Deep-learning Perception Components. *arXiv:2202.03360*, 2022

Conference Publications

Corina Păsăreanu, **Ravi Mangal**, Divya Gopinath, and Huafeng Yu. Assumption generation for the verification of learning-enabled autonomous systems. In *International Conference on Runtime Verification*. Springer, 2023

Corina S Păsăreanu, **Ravi Mangal**, Divya Gopinath, Sinem Getir Yaman, Calum Imrie, Radu Calinescu, and Huafeng Yu. Closed-loop analysis of vision-based autonomous systems: A case study. In *International Conference on Computer Aided Verification*, pages 289–303. Springer, 2023

Ravi Mangal*, Zifan Wang*, Chi Zhang*, Klas Leino, Corina Păsăreanu, and Matt Fredrikson. On the Perils of Cascading Robust Classifiers. In *International Conference on Learning Representations*, ICLR '23, 2023

(α) Divya Gopinath, Luca Lungeanu, **Ravi Mangal**, Corina Păsăreanu, Siqi Xie, and Huafeng Yu. Feature-guided Analysis of Neural Networks. In *Fundamental Approaches to Software Engineering*, FASE'23. Springer, 2023

Klas Leino*, Chi Zhang*, **Ravi Mangal***, Matt Fredrikson, Bryan Parno, and Corina Păsăreanu. Degradation Attacks on Certifiably Robust Neural Networks. *Transactions on Machine Learning Research*, 2022

Ravi Mangal, Kartik Sarangmath, Aditya V. Nori, and Alessandro Orso. Probabilistic Lipschitz Analysis of Neural Networks. In *International Static Analysis Symposium*, SAS '20. Springer, 2020

Ravi Mangal, Aditya V. Nori, and Alessandro Orso. Robustness of Neural Networks: A Probabilistic and Practical Approach. In *Proceedings of the 41st International Conference on Software Engineering: New Ideas and Emerging Results*, ICSE-NIER '19, 2019

Sulekha Kulkarni, **Ravi Mangal**, Xin Zhang, and Mayur Naik. Accelerating Program Analyses by Cross-program Training. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, OOPSLA '16, 2016

Ravi Mangal, Xin Zhang, Aditya Kamath, Aditya V. Nori, and Mayur Naik. Scaling Relational Inference Using Proofs and Refutations. In *Thirtieth AAAI Conference on Artificial Intelligence*, AAAI '16, 2016

Xin Zhang, **Ravi Mangal**, Aditya V. Nori, and Mayur Naik. Query-guided Maximum Satisfiability. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '16, 2016

Ravi Mangal, Xin Zhang, Aditya V. Nori, and Mayur Naik. Volt: A Lazy Grounding Framework for Solving Very Large MaxSAT Instances. In *International Conference on Theory and Applications of Satisfiability Testing*, SAT '15, 2015

Ravi Mangal, Xin Zhang, Aditya V. Nori, and Mayur Naik. A User-guided Approach to Program Analysis. In *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering*, ESEC/FSE '15, 2015

Xin Zhang, **Ravi Mangal**, Mayur Naik, and Hongseok Yang. Hybrid Top-down and Bottom-up Interprocedural Analysis. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, 2014

Xin Zhang, **Ravi Mangal**, Radu Grigore, Mayur Naik, and Hongseok Yang. On Abstraction Refinement for Program Analyses in Datalog. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, 2014

Ravi Mangal, Mayur Naik, and Hongseok Yang. A Correspondence Between Two Approaches to Interprocedural Analysis in the Presence of Join. In *Proceedings of the 23rd European Symposium on Programming Languages and Systems - Volume 8410*, ESOP '14, 2014

Workshop Papers and Abstracts

Ravi Mangal and Corina Păsăreanu. A Cascade of Checkers for Run-time Certification of Local Robustness. In *5th Workshop on Formal Methods for ML-Enabled Autonomous Systems*, 2022

Klas Leino, Aymeric Fromherz, **Ravi Mangal**, Matt Fredrikson, Bryan Parno, and Corina Păsăreanu. Self-correcting Neural Networks for Safe Classification. In *5th Workshop on Formal Methods for ML-Enabled Autonomous Systems*, 2022

Ravi Mangal, Aditya V. Nori, and Alessandro Orso. Checking Probabilistic Properties of Neural Networks via Symbolic Methods and Sampling. In *First ICSE Workshop on Testing for Deep Learning and Deep Learning for Testing*, DeepTest '19, 2019

Ravi Mangal, David Devecsery, and Alessandro Orso. On Optimally Combining Static and Dynamic Analyses for Intensional Program Properties. In *The Southeast Regional Programming Languages Seminar*, SERPL '19, 2019

Technical Reports

Ravi Mangal, Xin Zhang, Mayur Naik, and Aditya V. Nori. Solving Weighted Constraints with Applications to Program Analysis. Technical report, Georgia Institute of Technology, 2015

Honors and Awards

Invited to Dagstuhl Seminar on Resilience and Antifragility of Autonomous Systems	**2024**
Invited to attend the DARPA AI Forward workshop	**2023**
Invited to Dagstuhl Seminar on Machine Learning and Logical Reasoning: The New Frontier	**2022**
Distinguished paper award at ESEC/FSE	**2015**
Distinguished paper award at PLDI	**2014**
Best paper award nominee at ESOP	**2014**

Grants

LLM Self-Defense Against Adversarial Attacks for Coding Tasks
CyLab Future Enterprise Security Initiative
PIs: Corina Păsăreanu, Limin Jia, Ravi Mangal, USD 75,000 **2023**

Verifiable Personalization for Federated Learning
CyLab Future Enterprise Security Initiative
PIs: Corina Păsăreanu, Ravi Mangal, USD 60,000 **2022**

Research Talks

Invited Talks

Safety Analysis of Vision-based Autonomous Systems
CMU CyLab Partners Conference **Oct 2023**

Feature-Guided Engineering of Neural Networks
CMU CyLab Partners Conference **Oct 2022**

The Necessity of Run-time Techniques for Safe ML
Dagstuhl Seminar on Machine Learning and Logical Reasoning: The New Frontier **July 2022**

Repairing Neural Classifiers at Run-time: Safety for Free
CMU CyLab Partners Conference **Oct 2021**

A User-Guided Approach to Program Analysis
Microsoft Research India **Aug 2015**

Conference and Workshop Presentations and Posters
CAV'23, FoMLAS'22, SAS'20, ICSE-NIER'19, DeepTest'19, SERPL'19, AAAI'16, SAT'15, FSE'15, ESOP'14

Teaching Experience

Teaching Assistant

CS6340 Software Analysis and Testing (Online)	**Spring 2016**
CS6340 Software Analysis and Testing	**Spring 2016**
CS8803 Foundations of Programming Languages	**Fall 2013**
CS8803 Foundations of Programming Languages	**Fall 2012**
CS6340 Software Analysis and Testing	**Fall 2011**

Mentoring Experience

Haoran Wang (MS at CMU)	**Jan 2023 - July 2023**
Siqi Xi (MS at CMU)	**Aug 2022 - July 2023**
Chi Zhang (PhD at CMU)	**Mar 2021 - Present**
Kartik Sarangmath (BS/MS at Georgia Tech)	**Aug 2019 - Dec 2022**
SIGPLAN-M	**Oct 2020 - Present**

Academic Service

Program Committee
ICSE'25, NASA Formal Methods'24, CAIN'24, CAV'24

Reviewer
NSF Panel, International Journal of Information Security, TOSEM, ICLR'24, NeurIPS'23, NeurIPS'22, ICLR'22

Sub-Reviewer
NFM'23, POPL'22, CAIN'22, Oakland'22, PLDI'21, ICSE'20, ISSTA'20, FSE'19, FSE'18, ISSTA'18, SPIN'17, RV'17, JCST'17, ESSOS'17, CAV'14, HVC'14, ICSE-SRC'14

Artifact Evaluation Committee Member
OOPSLA'24, POPL'23, POPL'20, ISSTA'18, OOPSLA'17, OOPSLA'16

Member of Graduate Student Council for College of Computing at Georgia Tech 2017