ASHISH HOODA

https://pages.cs.wisc.edu/~hooda

EDUCATION

University of Wisconsin-Madison, Madison, WI — Aug 2019 - Present
Doctoral Student, Computer Science

Indian Institute of Technology, Delhi, India — July 2014 - May 2018
B.E. in Electrical Engineering (*Minor in Computer Science*)

INTERESTS

Security & Privacy, Computer Vision, Large Language Models, Graph Learning

WORK EXPERIENCE

Research Intern @ Google — Sep 2024 - Dec 2024
Working on robustness of embedding models for code retrieval.

Research Intern @ Google — Jul 2023 - Nov 2023
Supervisor: Mihai Christodorescu, Miltiadis Allamanis
Worked on evaluating program semantics understanding of Large Language Models for Code. Project resulted in a paper at ICML 2024.

Applied Scientist Intern @ Amazon AWS — Jun 2022 - Sep 2022
Supervisor: Ali Torkamani
Developed an efficient Graph Neural Network training framework that scales to billion node scale graphs. Utilized residual quantization to reduce codebook size without sacrificing precision. Demonstrated memory and compute efficiency on the largest Open Graph Benchmark dataset - ogbn-papers100M.

Software Engineer @ Microsoft India R&D — Jun 2018 - Jul 2019
Worked on Omnichannel Engagement Hub in the Dynamics CRM team; Created a Microsoft Azure Service-Fabric based service for configuring presence of a user. Proposed and Implemented a probabilistic distribution model for agent assignment with real-time feedback.

INVITED TALKS

Counterfactual Analysis for Code Predicates — *JetBrains Research*, Oct 2024
Is Detection A Viable Defense For against Attacks? — *Visa Research*, June 2024
Do Code LLMs understand program semantics? — *Google ML4Code Team*, Nov 2023
Do Stateful Defenses Work Against Black-Box Attacks? — *Google AI Red Team*, Oct 2023
Deepfake Detection Against Adaptive Attackers — *Google AI Red Team*, Aug 2023

PUBLICATIONS

* : CO FIRST AUTHORS

PolicyLR: A LLM compiler for Logic based Representation for Privacy Policies
Ashish Hooda, Rishabh Khandelwal, Prasad Chalasani, Kassem Fawaz, Somesh Jha
NeurIPS 2024 Workshop *(Safe & Trustworthy Agents Workshop)* *[Paper]*

PRP: Propagating Universal Perturbations to Attack LLM Guard-Rails
Neal Mangaokar*, Ashish Hooda*, Jihye Choi, Shreyas Chandrashekaran, Kassem Fawaz, Somesh Jha, Atul Prakash
ACL 2024 *(Association for Computational Linguistics)* *[Paper][Code]*

Do Large Code Models Understand Programming Concepts? Counterfactual Analysis for Code Predicates
Ashish Hooda, Mihai Christodorescu, Miltiadis Allamanis, Aaron Wilson, Kassem Fawaz, Somesh Jha
ICML 2024 *(International Conference on Machine Learning)* *[Paper]*

D4: Detection of Adversarial Diffusion Deepfakes Using Disjoint Ensembles
Ashish Hooda*, Neal Mangaokar*, Ryan Feng, Kassem Fawaz, Somesh Jha, Atul Prakash
WACV 2024 (*IEEE/CVF Winter Conference on Applications of Computer Vision*) *[Paper][Code]*

Theoretically Principled Trade-off for Stateful Defenses against Query-Based Black-Box Attacks
Ashish Hooda*, Neal Mangaokar*, Ryan Feng, Kassem Fawaz, Somesh Jha, Atul Prakash
ICML 2023 Workshop (*2nd AdvML Frontiers Workshop*) *[Paper]*

Stateful Defenses for Machine Learning Models Are Not Yet Secure Against Black-box Attacks
Ryan Feng*, Ashish Hooda*, Neal Mangaokar*, Kassem Fawaz, Somesh Jha, Atul Prakash
CCS 2023 *(ACM Conference on Computer and Communications Security)* *[Paper][Code]*

Experimental Analyses of Physical Surveillance Risks in Client-Side Content Scanning
Ashish Hooda, Andrey Labunets, Tadayoshi Kohno, Earlence Fernandes
NDSS 2024 (*Network and Distributed System Security Symposium*) *[Paper]*

SkillFence: A Systems Approach to Mitigating Voice-Based Confusion Attacks
Ashish Hooda, Matthew Wallace, Kushal Jhunjhunwalla, Earlence Fernandes, Kassem Fawaz
IMWUT 2022 (*ACM Interactive, Mobile, Wearable and Ubiquitous Technologies*) *[Paper]*

Invisible Perturbations: Physical Adv Examples Exploiting the Rolling Shutter Effect
Athena Sayles*, Ashish Hooda*, Mohit Gupta, Rahul Chatterjee, Earlence Fernandes
CVPR 2021 (*Conference on Computer Vision and Pattern Recognition*) *[Paper][Code]*

PREPRINTS
* : CO FIRST AUTHORS

Functional Homotopy: Smoothing Discrete Optimization Via Continous Parameters for LLM Jailbreak Attacks
Zi Wang*, Divyam Anshuman*, Ashish Hooda, Somesh Jha
Preprint *[Paper]*

Synthetic Counterfactual Faces
Guruprasad V Ramesh, Harrison Rosenberg, Ashish Hooda, Kassem Fawaz
Preprint *[Paper]*

TECHNICAL

Languages: Python, Java, C++, C, MATLAB

Frameworks/Libraries: PyTorch, Tensorflow, Apache Spark, Deep Graph Library

SERVICE

- Reviewer: ICML 2024, ICLR 2025, ICLR ME-FoMO Workshop ('23, '24)
- Artifact Evaluation Committee Member: USENIX Security Symposium '22
- External Reviewer: USENIX Security Symposium ('19, '20, '21, '22, '23, '24), IEEE S&P ('19, '20, '21, '22, '23, '24), IEEE SaTML ('24)
- Mentor at Individualized Cybersecurity Research Mentoring (iMentor) Workshop 2023

AWARDS & ACHIEVEMENTS

- Accepted for NDSS Travel Grant 2024.
- WACV Doctoral Consortium Award 2024.
- Runner up in CS Research Symposium, 2022 (UW Madison).
- Qualified for regionals at ACM International Collegiate Programming Contest (ICPC), 2017.
- Runner-up at Microsoft CODE-FUN-DO Hackathon, 2015.
- Secured **All India Rank 4** in Central Board of Secondary Education (CBSE) Board Examination given by over 2 million students.
- Secured **All India Rank 17** in Joint Entrance Exam (JEE) given by over 1 million students.
- Selected for Special Class Railway Apprentice (SCRA) (Top 100 out of over 0.1 million applicants).
- Awarded the Junior Science Talent Search Examination (JSTSE) Scholarship.