ANSHUMAN SURI

github/linkedin/twitter: iamgroot42 ⋄ website: https://anshumansuri.me ⋄ google scholar
+1-434-242-9346 ⋄ anshuman@virginia.edu

EDUCATION

University of Virginia *2019 - Present*
Ph.D. in Computer Science

Indraprastha Institute of Information Technology, Delhi *2014 - 2018*
B.Tech (Hons.) in Computer Science

RESEARCH EXPERIENCE

Oracle Research Labs *Sep 2021 - Dec 2021*
Research Intern

· Worked on studying subject-level membership inference in federated learning.

Microsoft *June 2018 - July 2019*
Applied Scientist

· Part of Bing STCI team: Project Personality Chat, Microsoft Icecaps, and QnA Maker.
· Worked on language modeling systems, increasing their qualitative performance and latency.

Image Analysis and Biometrics Lab, IIITD/IITJ *Jan 2018 - May 2019*
Undergraduate Researcher

· Worked on active learning and domain adaptation for face identification under multiple covariates.

PreCog, IIITD *May 2016 - March 2019*
Undergraduate Researcher

· Worked on various projects, including multi-modal content analysis on Online Social Media (OSM).
· Collaborated (remotely) with IBM Research Labs, Bangalore for a project on adversarial defense.

Microsoft *May 2017 - July 2017*
SDE Intern, Bing Team

· Worked on developing AI-powered game bots as part of Bing STCI.

RIGOROUSLY REVIEWED PUBLICATIONS

SoK: Pitfalls in Evaluating Black-Box Attacks. IEEE Conference on Secure and Trustworthy Machine Learning (SaTML) 2024.
Fnu Suya, ***Anshuman Suri****, Tingwei Zhang, Scott Hong, Yuan Tian, David Evans*

SoK: Let The Privacy Games Begin! A Unified Treatment of Data Inference Privacy in Machine Learning. IEEE Symposium on Security and Privacy (S&P) 2023.

Ahmed Salem, Giovanni Cherubin, David Evans, Boris Köpf, Andrew Paverd, ***Anshuman Suri****, Shruti Tople, Santiago Zanella-Bguelin*

Manipulating Transfer Learning for Property Inference. The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023.
Yulong Tian, Fnu Suya, ***Anshuman Suri****, Fengyuan Xu, David Evans*

Dissecting Distribution Inference. IEEE Conference on Secure and Trustworthy Machine Learning (SaTML) 2023.
Anshuman Suri*, Yifu Lu, Yanjin Chen, David Evans*

Formalizing and Estimating Distribution Inference Risks. Privacy Enhancing Technologies Symposium (PETS) 2022.
Anshuman Suri*, David Evans*

Model-Targeted Poisoning Attacks: Provable Convergence and Certified Bounds. International Conference on Machine Learning (ICML) 2021.
Fnu Suya, Saeed Mahloujifar, ***Anshuman Suri****, David Evans, Yuan Tian*

A2-LINK: Recognizing Disguised Faces via Active Learning and Adversarial Noise Based Inter-Domain Knowledge. IEEE Transactions on Biometrics, Identity and Behavior (T-BIOM) 2020
Anshuman Suri*, Mayank Vatsa, Richa Singh*

PREPRINTS

Do Membership Inference Attacks Work on Large Language Models? arXiv, 2024.
*Michael Duan**, ***Anshuman Suri****, *Niloofar Mireshghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, Hannaneh Hajishirz*

SoK: Memorization in General-Purpose Large Language Models. arXiv, 2023.
Valentin Hartmann, ***Anshuman Suri****, Vincent Bindschaedler, David Evans, Shruti Tople, Robert West*

Subject Membership Inference Attacks in Federated Learning. arXiv, 2022.
Anshuman Suri*, Pallika Kanani, Virendra J Marathe, Daniel W Peterson*

TEACHING EXPERIENCE

UVA: Introduction to Computer Vision, Vision and Language, Computational Biology / Biological Computing (2020-2022)

IIITD: Deep Learning, Machine Learning, Linear Algebra (2016-2018)

Responsibilities: Designing and grading assignments, holding office hours, leading weekly discussion sessions (Linear Algebra), and taking guest lectures (Computational Biology).

SOFTWARE

· mimir Python package for evaluating MIA in LLMs.

- distribution_inference Python package for distribution inference.
- popcll_torch: CUDA-level implementation for popcll operation in PyTorch.
- PyTorch implementation for permutation-invariant networks, with extension for CNNs.
- Course project, Cyber-Physical Systems (UVA): Analysis of potential issues in AI-aided autopilot.
- Course project, Deep Learning for Visual Recognition (UVA): Semi-supervised method to learn independent, redundant features as "concepts".

ACHIEVEMENTS

Received Endowed Graduate Fellowship, SEAS, UVA for 2023-24	*2023*
Awarded bounty for identifying vulnerabilities, New Bing Research Challenge, MSRC	*2023*
Received John A. Stankovic Graduate Research Award, UVA	*2023*
Second position in MICO challenge, CIFAR-track (co-located with SaTML)	*2023*
Outstanding Reviewer: ICLR, ICCV	*2021*
Invited to give talk on Project Personality Chat, Intelligent Cloud Conference	*2019*
Exceptionally fast entry-level promotion: Level 59 $\rightarrow$ 60 in 6 months, Microsoft IDC	*2018*
On Dean's List for academic excellence, IIITD	*2016*

REVIEWER DUTIES

NeurIPS, ICLR, ICML, ACL, TPAMI, ARR, CVPR, EMNLP, ICCV, ECCV, ACL-IJCNLP

OTHER ACTIVITIES

Student Advisory Council Member (representing UVA), ACTION Institute.	*2023 - Present*
Social Chair, Computer Science Department Graduate Student Group (CSGSG)	*2023 - 2024*
Co-President, Animal Justice Advocates (AJA)	*2020 - 2022*