VARUN CHANDRASEKARAN

PERSONAL INFORMATION

e-mail (MSR): varuncha@microsoft.com
e-mail (UIUC): varunc@illinois.edu

EMPLOYMENT

Microsoft Research, Redmond, WA
Post-doctoral Researcher, Adaptive Systems & Interaction Group — Aug 2022 onwards

University of Illinois Urbana-Champaign, Champaign-Urbana, IL
Adjunct Assistant Professor in ECE with a courtesy appointment in CS — Aug 2022 onwards

EDUCATION

University of Wisconsin-Madison, Madison, WI
Ph.D, Computer Science — Sep 2016 - July 2022

New York University, New York, NY
M.S, Computer Science — Aug 2014 - June 2016

Anna University, Chennai, India
B.E, Computer Science and Engineering — Aug 2010 - May 2014

AWARDS AND HONORS

Long Talk at ICML (Top 3% of all papers submitted), 2021
Landweber NCR Fellowship in Distributed Systems (at UW-Madison), 2020
M.S. Thesis/Research Fellowship (at NYU), 2016

CONFERENCE PUBLICATIONS

Bibliometrics (Google scholar, October 27, 2022): citations: 529, h-index: 11

* DENOTES JOINT CONTRIBUTION
AR DENOTES ACCEPTANCE RATE

[**C11**] Unrolling SGD: Understanding Factors Influencing Machine Unlearning
Anvith Thudi, Gabriel Deza, **VC**, Nicolas Papernot
In the 7th IEEE European Symposium of Security & Privacy (EuroS&P), June 2022 *(AR: 29.7%)*

[**C10**] CONFIDANT: A Privacy Controller for Social Robots
Brian Tang, Dakota Sullivan, Bengisu Cagiltay, **VC**, Kassem Fawaz, Bilge Multu
In the 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI), March 2022 *(AR: 24.8%)*

[**C9**] PowerCut & Obfuscator: An Exploration of the Design Space for Privacy-Preserving Interventions for Smart Speakers
VC, Suman Banerjee, Bilge Multu, Kassem Fawaz
In the 17th USENIX Symposium on Usable Privacy and Security (SOUPS), August 2021 *(AR: 26.4%)*

[**C8**] A General Framework For Detecting Anomalous Inputs to DNN Classifiers
Jayaram Raghuram*, **VC***, Somesh Jha, Suman Banerjee
Long Talk (Top 3% of submitted papers) In the 38th International Conference on Machine Learning (ICML), July 2021 *(AR: 21.5%)*

[**C7**] Proof-of-Learning: Definitions and Practice
Hengrui Jia, Mohammad Yaghini, Christopher A. Choquette-Choo, Natalie Dullerud, Anvith Thudi, **VC**, Nicolas Papernot
In the 42nd IEEE Symposium on Security and Privacy (S&P), May 2021 *(AR: 12.08%)*

[**C6**] Entangled Watermarks as a Defense against Model Extraction
Hengrui Jia, Christopher A. Choquette-Choo, **VC**, Nicolas Papernot
In the 30th USENIX Security Symposium, August 2021 *(AR: 19%)*

[**C5**] Face-Off: Adversarial Face Obfuscation
VC*, Chuhan Gao*, Brian Tang, Kassem Fawaz, Somesh Jha, Suman Banerhee
In the 21st Proceedings on Privacy Enhancing Technologies (PoPETS), July 2021 *(AR: 19%)*

[**C4**] Machine Unlearning
Lucas Bourtoule*, **VC***, Christopher A. Choquette-Choo*, Hengrui Jia*, Adelin Travers*, Baiwu Zhang*, David Lie, Nicolas Papernot
In the 42nd IEEE Symposium on Security and Privacy (S&P), May 2021 *(AR: 12.08%)*

[**C3**] Exploring Connections Between Active Learning and Model Extraction
VC, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, Songbai Yan
In the 29th USENIX Security Symposium, August 2020 *(AR: 16.1%)*

[**C2**] A Framework for Analyzing Spectrum Characteristics in Large Spatio-Temporal Scales
Yijing Zeng, **VC**, Suman Banerjee, Domenico Giustiniano
In the 25th Annual International Conference on Mobile Computing and Networking (MobiCom), October 2019 *(AR: 17.3%)*

[**C1**] Alphacodes: Usable, Secure Transactions with Untrusted Providers using Human Computable Puzzles
Ashlesh Sharma, **VC**, Fareeha Amjad, Dennis Shasha, Lakshminarayanan Subramanian
In the 7th Annual Symposium on Computing for Development (DEV), November 2016 *(AR: 32%)*

Workshop Publications

* denotes joint contribution

[**W3**] Analyzing And Improving Neural Networks By Generating Semantic Counterexamples Through Differentiable Rendering
Lakshya Jain, **VC**, Uyeong Jang, Somesh Jha, Sanjit A. Seshia
Workshop on Uncertainty & Robustness in Deep Learning at ICML 2021
`Long Version (URL):` https://arxiv.org/abs/1910.00727

[**W2**] Causally Constrained Data Synthesis For Private Data Release
VC, Darren Edge, Somesh Jha, Amit Sharma, Cheng Zhang, Shruti Tople
Workshop on Distributed and Private Machine Learning at ICLR 2021
`Long Version (URL):` https://arxiv.org/abs/2105.13144

[**W1**] Traversing the Quagmire that is Privacy in your Smart-Home
Chuhan Gao*, **VC***, Kassem Fawaz, Suman Banerjee
Workshop on IoT Security and Privacy at SIGCOMM 2018

Manuscripts

* denotes joint contribution

[**P8**] Verifiable and Provably Secure Machine Unlearning
Thorsten Eisenhofer, Doreen Riepel, **VC**, Esha Ghosh, Olga Ohrimenko, Nicolas Papernot
`URL:`https://arxiv.org/pdf/2210.09126

[**P7**] On the Fundamental Limits of Formally (Dis) Proving Robustness in Proof-of-Learning
Congyu Fang, Hengrui Jia, Anvith Thudi, Mohammad Yaghini, Christopher Choquette-Choo, Natalie Dullerud, **VC**, Nicolas Papernot
`URL:`https://arxiv.org/pdf/2208.03567

[**P6**] Generative Extraction of Audio Classifiers for Speaker Identification
Tejumade Afonja, Lucas Bourtoule, **VC**, Sageev Oore, Nicolas Papernot
`URL:`https://arxiv.org/pdf/2207.12816

[**P5**] Hierarchical Federated Learning with Privacy
VC, Suman Banerjee, Diego Perino, Nicolas Kourtellis
`URL:`https://arxiv.org/pdf/2206.05209

[**P4**] SoK: Machine Learning Governance
VC*, Hengrui Jia*, Anvith Thudi*, Adelin Travers*, Mohammad Yaghini*, Nicolas Papernot
`URL:`https://arxiv.org/abs/2109.10870

[**P3**] On the Exploitability of Audio Machine Learning Pipelines to Surreptitious Adversarial Examples
Adelin Travers, Lorna Licollari, Guanghan Wang, **VC**, Adam Dziedzic, David Lie, Nicolas Papernot
`URL:`https://arxiv.org/abs/2108.02010

[**P2**] On the Effectiveness of Mitigating Data Poisoning Attacks with Gradient Shaping
Sanghyun Hong, **VC**, Yigitcan Kaya, Tudor Dumitras, Nicolas Papernot
`URL:`https://arxiv.org/abs/2002.11497

[**P1**] Rearchitecting Classification Frameworks For Increased Robustness
VC, Brian Tang, Nicolas Papernot, Kassem Fawaz, Somesh Jha, Xi Wu
`URL:`https://arxiv.org/abs/1905.10900

Internship Experience

Research Intern, **Lacework**, Santa Clara, CA, USA — Aug 2021 - Nov 2021
Supervisors: Dr. Ting-Fang Yen, Dr. Ulfar Erlingsson

Research Intern, **Telefonica**, Barcelona, Spain — March 2021 - June 2021
Supervisors: Dr. Nicolas Kourtelis, Dr. Diego Perino

Research Intern, **Microsoft Research**, Cambridge, UK — Nov 2020 - Jan 2021
Supervisors: Dr. Shruti Tople, Dr. Cheng Zhang

Visiting Scholar, **University of Toronto**, Toronto, Canada — Sept 2019 - Jan 2020
Supervisor: Prof. Nicolas Papernot

Research Intern, **Microsoft Research**, Bengaluru, India — Feb 2017 - June 2017
Supervisors: Dr. Ranjita Bhagwan, Dr. Ramachandran Ramjee

Research Intern, **IBM Research**, Yorktown Heights, NY, USA — July 2016 - Sep 2016
Supervisor: Dr. John Tracey

Intern, **AT&T Labs & Research**, Middletown, NJ, USA — June 2015 - Aug 2015
Managers: Vatsal Parikh, Matt Szela

Patents

∘WARF ref. P200291US01 "METHOD AND APPARATUS USING BLENDED BIOMETRIC DATA", Filed 2/16/2021, Application 17/177080 (Pending)
∘WARF ref. P210077US01 "APPARATUS FOR BANDWIDTH EFFICIENT VIDEO COMMUNICATION", Filed 7/14/2020, Application 16/928690 (Pending)
∘ "METHOD AND SYSTEM FOR PROVIDING DIFFERENTIAL PRIVACY USING FEDERATED LEARNING", EU Patent Office, by Telefonica I+D, Spain (Pending)

Mentoring

∘ Fan Wu, UIUC (joint with David Forsyth)
∘ Brian Tang, UW-Madison (now PhD student at University of Michigan) **[C5, P1, C10]**
∘ Mohammad Yaghini, University of Toronto **[C7, P4]**
∘ Chris A. Choquette-Choo, Google Brain **[C6, C7, P4]**
∘ Hengrui Jia, University of Toronto **[C4, C6, C7]**
∘ Adelin Travers, University of Toronto (now senior pen-tester at Verizon, Japan) **[C4, P3]**
∘ Anvith Thudi, University of Toronto **[C7, C11, P4]**

Service

∘ **(External+Sub) Reviewer for:**

1. European Symposium on Research in Computer Security (**ESORICS**), 2020
2. IEEE Symposium on Security & Privacy (**S&P**), 2019-2021
3. **USENIX Security** Symposium (2020-2021)
4. Privacy Enhancing Technologies Symposium (**PETS**) 2021-2022
5. ACM Conference on Computer and Communication Security (**CCS**), 2019-2020
6. Interactive, Mobile, Wearable and Ubiquitous Technologies (**IMWUT**), 2021
7. ACM Conference On Mobile Computing And Networking (**MobiCom**), 2017-2021
8. ACM/IEEE Symposium on Edge Computing (**SEC**), 2017-2018, 2021
9. ACM Conference On Emerging Networking Experiments And Technologies (**CoNEXT**), 2017
10. USENIX Symposium on Networked Systems Design and Implementation (**NSDI**), 2021
11. ACM Conference on Embedded Networked Sensor Systems (**SenSys**), 2018
12. International Conference on Machine Learning (**ICML**), 2019
13. Neural Information Processing Systems (**NeurIPS**), 2020-2021
14. Society for Artificial Intelligence and Statistics (**AISTATS**), 2021
15. ACM The Web Conference (**WWW**), 2022

∘ **PC Member/Reviewer:**

1. NeurIPS Workshop on Security in Machine Learning, 2018
2. DSN Workshop on Dependable and Secure Machine Learning (**DSML**), 2019
3. ICML Workshop on Security and Privacy of Machine Learning, 2019
4. ICLR Workshop on Trustworthy ML, 2020
5. **DLS** (co-located with S&P), 2020-2022

6. International Conference on Learning Representations (**ICLR**), 2022
7. **AISTATS, CCS, ICML, NeurIPS**, 2022
8. **USENIX Security, S&P, PETS, AISTATS**, 2023
9. IEEE Secure & Trustworthy ML (**SaTML**), 2023

◦ **Organizer:**
DSML, 2020 onwards (co-chair in 2020, 2022)