VARUN CHANDRASEKARAN

Personal Information

1210 West Dayton Street
Madison, WI 53706
e-mail: chandrasekaran@cs.wisc.edu
Phone: +1 (845) 978 0909
http://pages.cs.wisc.edu/~chandrasekaran/

Education

University of Wisconsin-Madison, Madison, WI — Sep 2016 - May 2022 (expected)
Doctoral Candidate, Computer Science

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

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

Research Areas

- Security & Privacy
- Machine Learning

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, November 2021): citations: 239, h-index: 8

* denotes joint contribution
AR denotes acceptance rate

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

[**C9**] PowerCut & Obfuscator: An Exploration of the Design Space for Privacy-Preserving Interventions for Smart Speakers; **Varun C**, 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*, **Varun C***, 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, **Varun C**, 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, **Varun C**, Nicolas Papernot; In the 30th USENIX Security Symposium, August 2021 *(AR: 19%)*

[**C5**] Face-Off: Adversarial Face Obfuscation; **Varun C***, 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*, **Varun C***, 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; **Varun C**, 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, **Varun C**, 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, **Varun C**, 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, **Varun C**, 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; **Varun C**, 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*, **Varun C***, Kassem Fawaz, Suman Banerjee; Workshop on IoT Security and Privacy at SIGCOMM 2018

Manuscripts

* denotes joint contribution

[**P6**] Murfe: Multi-modal Biometric Fuzzy Extractor; Xiaohan Fu*, **Varun C***, Lakshminarayanan Subramanian, Jin-Yi Cai, Rahul Chaterjee *(Under review at USENIX Security 2022)*

[**P5**] SoK: Machine Learning Governance; **Varun C***, Hengrui Jia*, Anvith Thudi*, Adelin Travers*, Mohammad Yaghini*, Nicolas Papernot; `URL:`https://arxiv.org/abs/2109.10870

[**P4**] Unrolling SGD: Understanding Factors Influencing Machine Unlearning; Anvith Thudi, Gabriel Deza, **Varun C**, Nicolas Papernot; `URL:`https://arxiv.org/abs/2109.13398 *(Under review at EuroS&P 2022)*

[**P3**] On the Exploitability of Audio Machine Learning Pipelines to Surreptitious Adversarial Examples; Adelin Travers, Lorna Licollari, Guanghan Wang, **Varun C**, 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, **Varun C**, Yigitcan Kaya, Tudor Dumitras, Nicolas Papernot; `URL:`https://arxiv.org/abs/2002.11497

[**P1**] Rearchitecting Classification Frameworks For Increased Robustness; **Varun C**, 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
- I worked on techniques to detect *colored subgraph isomorphisms* using graph neural networks.

Research Intern, **Telefonica**, Barcelona, Spain — March 2021 - June 2021
Supervisors: Dr. Nicolas Kourtelis, Dr. Diego Perino
- I worked on improving privacy vs. utility trade-offs in *federated learning*.

Research Intern, **Microsoft Research**, Cambridge, UK — Nov 2020 - Jan 2021
Supervisors: Dr. Shruti Tople, Dr. Cheng Zhang
- I worked on the implications of *causal information* on the *privacy of generative models* [**W2**].

Visiting Scholar, **University of Toronto**, Toronto, Canada — Sept 2019 - Jan 2020
Supervisor: Prof. Nicolas Papernot
- I worked on designing techniques for effective *data removal* from trained ML models [**C4**].
- I also worked on designing defense mechanisms against *data poisoning* [**P2**].

Research Intern, **Microsoft Research**, Bengaluru, India — Feb 2017 - June 2017
Supervisors: Dr. Ranjita Bhagwan, Dr. Ramachandran Ramjee
- I worked on identifying *homomorphic encryption* schemes for neural network inference.

Research Intern, **IBM Research**, Yorktown Heights, NY, USA — July 2016 - Sep 2016
Supervisor: Dr. John Tracey
- Analyzed the last *decade of data center research* to identify problems for the next decade.

Intern, **AT&T Labs & Research**, Middletown, NJ, USA — June 2015 - Aug 2015
Manager(s): Vatsal Parikh, Matt Szela
- I performed a feasibility study, PoC testing & bench-marking focused on NFV of packet routing & forwarding functions.

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

∘ Brian Tang, UW-Madison (now PhD student at University of Michigan) **[C5, P1, C10]**
∘ Mohammad Yaghini, University of Toronto **[C7, P5]**
∘ Chris A. Choquette-Choo, University of Toronto (now AI resident at Google) **[C6, C7, P5]**
∘ Hengrui Jia, University of Toronto (now MEng student at University of Toronto) **[C4, C6, C7]**
∘ Adelin Travers, University of Toronto (now senior pen-tester at Verizon, Japan) **[C4, P3]**
∘ Anvith Thudi, Maple Ridge Secondary (now BSc student at University of Toronto) **[C7, P4, P5]**

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), 2020-2021
6. International Conference on Learning Representations (**ICLR**), 2022
7. **AISTATS**, 2022

∘ **Organizer:**
DSML, 2020-2021 (https://dependablesecureml.github.io)

∘ **Student Service:**
Vice President of the Students of the ACM, UW-Madison Chapter, 2018-2019

Graduate Teaching Experience

Networks and Mobile Systems, Fall 2016
Computer Systems Organization, Fall 2015
Technology and Economic Development, Spring 2015

Suman Banerjee
Professor, Computer Sciences Department
University of Wisconsin-Madison
1210 West Dayton Street
Madison, WI 53706, USA
+1 (608) 262 7387
suman@cs.wisc.edu

Somesh Jha
Professor, Computer Sciences Department
University of Wisconsin-Madison
1210 West Dayton Street
Madison, WI 53706, USA
+1 (608) 262 9519
jha@cs.wisc.edu

Nicolas Papernot
Assistant Professor, Department of Electrical & Computer Engineering
University of Toronto
10 Kings College Road, Room SFB540
Toronto, ON M5S 3C4, Canada
+1 (416) 978 7600
nicolas.papernot@utoronto.ca

Patrick McDaniel
Professor, Computer Science and Engineering
Pennsylvania State University
W329 Westgate Building
University Park, PA 16802, USA
+1 (814) 863 3599
mcdaniel@cse.psu.edu

Sanjit A. Seshia
Professor, Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
253 Cory Hall #1770
Berkeley, CA 94720, USA
+1 (510) 643 6968
sseshia@eecs.berkeley.edu

Hamed Haddadi
Professor, Dyson School of Design Engineering
Imperial College London
Level 3, Dyson Building
25 Exhibition Road, South Kensington
London SW7 2DB, UK
+44 (0)20 7594 2584
h.haddadi@imperial.ac.uk