ARUNESH SINHA

Assistant Professor		Cell Phone:	+65-90990641
School of Information Systems		Official Email:	aruneshs@smu.edu.sg
Singapore Management University		Alternate Email:	aruneshsinha@gmail.com
80 Stamford Road, Singapore 178902		Website:	Link to Webpage

RESEARCH INTERESTS

My interest lies in the area of Artificial Intelligence. I have utilized techniques from game theory and machine learning to address varied problems involving adversarial reasoning such as privacy audits, airport screening, and learning adversary behavior models in games.

EDUCATION

Carnegie Mellon University, Pittsburgh, PA Aug 2008 – Aug 2014

- Ph.D. in Electrical and Computer Engineering, CGPA: 4.0/4.0
- Thesis: Audit Games, Advisor: Anupam Datta

Indian Institute of Technology, Kharagpur, India June 2000 – June 2004

- B.Tech in Electrical Engineering, with minor in Electronics Engineering

WORK EXPERIENCE

Singapore Management University, Singapore Aug 2019 – Present

- Assistant Professor

University of Michigan, MI Aug 2016 – Jul 2019

- Assistant Research Scientist

University of Southern California, Los Angeles, CA Aug 2014 – Jul 2016

- Postdoctoral Researcher with Prof. Milind Tambe

Trilogy Software, Bangalore, India July 2004 – July 2008

- Software Engineer: Worked on various software projects.

INTERNSHIP EXPERIENCE

Intel Labs, Hillsboro, OR June 2013 – Aug 2013

- Research Intern: Machine learning based detection of activity of mobile-phone users.

Microsoft Research, Redmond, WA June 2012 – Aug 2012

- Research Intern: Designed and formally verified a TPM based logging scheme that guarantees log integrity.

Bhabha Atomic Research Center, Mumbai, India June 2003 – July 2003

- Undergrad Intern: Programmed a micro-controller for tasks to be executed inside a nuclear reactor.

HONORS AND AWARDS

- Distinguished Program Committee member award in International Joint Conference on AI (IJCAI) 2019.
- Best innovative application paper runner-up in the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016).
- 2012 Bertucci Fellowship, College of Engineering, Carnegie Mellon University: Awarded for accomplished Ph.D. students pursuing novel research. Awarded on recommendation of two faculty.
- Winner of BobbleHead Award in Trilogy India: Awarded in recognition of outstanding contribution to software delivery.
- Best Undergraduate Project at department level at IIT, Kharagpur. The project used signal processing to predict the movement of the boundary between two immiscible liquids flowing in a vertical tube.
- Attended the INMO (Indian National Math's Olympiad) Camp in Mumbai in 1999. Thirty students are chosen (from all over India) for this camp after a series of regional and national Mathematics exams.
- Won NTSE (National Talent Search) scholarship in 1998: Awarded to 1000 students from all over India, based on a national examination.

PUBLICATIONS

Journal Papers with Rigorous Reviews:

1. F. Fang, T. H. Nguyen, **A. Sinha**, S. Gholami, A. Plumptre, L. Joppa, M. Tambe, M. Driciru, F. Wanyama, A. Rwetsiba, R. Critchlow, C. M. Beale. Predicting Poaching for Wildlife Protection, in IBM Journal of Research and Development, 2017.
2. C. Zhang, S. Gholami, D. Kar, **A. Sinha**, M. Jain, R. Goyal, M. Tambe. Keeping Pace with Criminals: An Extended Study of Designing Patrol Allocation against Adaptive Opportunistic Criminals, in Games Journal, 2016.
3. **A. Sinha**, T.H. Nguyen, D. Kar, M. Brown, M. Tambe. A. X. Jiang, From Physical Security to Cyber Security, in Journal of Cybersecurity, 2015.

Conference Papers with Rigorous Reviews:

4. S. Jecmen, **A. Sinha**, Z. Li, L. Tran-Thanh. Bounding Regret in Empirical Games, in Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI), 2020 (to appear)
5. S. Shah, **A. Sinha**, P. Varakantham, A. Perrault, M. Tambe. Solving Online Threat Screening Games using Constrained Action Space Reinforcement Learning, in Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI), 2020 (to appear)
6. J. Li, X. Wang, Y. Lin, **A. Sinha**, M. P. Wellman. Generating Realistic Stock Market Order Streams, in Proceedings of 34rd AAAI Conference on Artificial Intelligence (AAAI), 2020 (to appear)
7. P. Naghizadeh*, **A. Sinha***. Adversarial Contract Design for Private Data Commercialization, in Proceedings of the 20th ACM Conference on Economics and Computation (EC-19), 2019 (to appear). [*=*Joint lead authors, Names Ordered Alphabetically*].
8. **A.Sinha**, M. P. Wellman. Incentivizing Collaboration in a Competition, in Proceedings of 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2019.

9. T. H. Nguyen, Y. Wang, **A. Sinha**, M. P. Wellman. Deception in Finitely Repeated Security Games, in Proceedings of 33rd AAAI Conference on Artificial Intelligence (AAAI), 2019.
10. **A. Sinha**, A. Schlenker, D. Dmello, M. Tambe. Scaling-up Stackelberg Security Games Applications using Approximations, in Proceedings of Conference on Decision and Game Theory for Security (GameSec), Oct 2018.
11. L. Nguyen, S. Wang, **A. Sinha**. A Learning and Masking Approach to Secure Learning, in Proceedings of Conference on Decision and Game Theory for Security (GameSec), Oct 2018.
12. **A. Sinha**, F. Fang, B. An, C. Kiekintveld, M. Tambe. Stackelberg Security Games: Looking Beyond a Decade of Success, in Proceedings of 27th International Joint Conference on Artificial Intelligence (IJCAI), July 2018.
13. S. M. Mc Carthy, C. M. Laan, K. Wang, P. Vayanos, **A. Sinha**, M. Tambe. The Price of Usability: Designing Operationalizable Strategies for Security Games, in Proceedings of 27th International Joint Conference on Artificial Intelligence (IJCAI), July 2018.
14. N. Papernot, P. McDaniel, **A. Sinha**, M. P. Wellman. Towards the Science of Security and Privacy in Machine Learning, in Proceedings of 3rd IEEE European Symposium on Security and Privacy (EuroS&P), April 2018.
15. A. Schlenker, H. Xu, M. Guirguis, C. Kiekintveld, **A. Sinha**, M. Tambe, S. Sonya, D. Balderas, N. Dunstatter. Don't Bury your Head inWarnings: A Game-Theoretic Approach for Intelligent Allocation of Cyber-security Alerts, in Proceedings of 26th International Joint Conference on Artificial Intelligence (IJCAI), Aug 2017.
16. S. M. Mc Carthy, **A. Sinha**, M. Tambe, P. Manadhata. Data Exfiltration Detection and Prevention: Virtually Distributed POMDPs for Practically Safer Networks, in Proceedings of Decision and Game Theory for Security (GameSec), 2016.
17. A. Schlenker, M. Brown, **A. Sinha**, M. Tambe, R. Mehta. Get Me to My GATE On Time: Efficiently Solving General-Sum Bayesian Threat Screening Games, in Proceedings of 22nd European Conference on Artificial Intelligence (ECAI), Aug 2016.
18. N. Haghtalab, F. Fang, T. Nguyen, **A. Sinha**, A. D. Procaccia, M. Tambe. Three strategies to success: Learning adversary models in security games, in Proceedings of 25th International Joint Conference on Artificial Intelligence (IJCAI), July 2016.
19. B. Ford, M. Brown, A. Yadav, A. Singh, **A. Sinha**, B. Srivastava, C. Kiekintveld, M. Tambe. Protecting the NECTAR of the Ganga River through Game-Theoretic Factory Inspections, in Proceedings of International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS), June 2016.
20. **A. Sinha**, D. Kar, M. Tambe. Learning Adversary Behavior in Security Games: A PAC Model Perspective, in Proceedings of 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2016.
21. T. Nguyen, **A. Sinha**, S. Gholami, A. Plumptre, L. Joppa, M. Tambe, M. Driciru, F. Wanyama, A. Rwetsiba, R. Critchlow, C. Beale. CAPTURE: A New Predictive Anti-Poaching Tool for Wildlife Protection, in Proceedings of 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2016. [**Best application paper runner-up**].
22. C. Zhang, V. Bucarey, A. Mukhopadhyay, **A. Sinha**, Y. Qian, Y. Vorobeychik, M. Tambe. Using abstractions to solve opportunistic crime security games at scale, in Proceedings of 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2016.

23. M. Brown*, **A. Sinha***, A. Schlenker, M. Tambe. One Size Does Not Fit All: A Game-Theoretic Approach for Dynamically and Effectively Screening for Threats, in Proceedings of 30th AAAI Conference on Artificial Intelligence (AAAI), 2016. [*=*Joint Lead Authors*]
24. A. Carbonara, A. Datta, **A. Sinha**, Y. Zick. Incentivizing Peer Grading in MOOCs: An Audit Game Approach, in Proceedings of International Joint Conference on Artificial Intelligence (IJCAI), July 2015. [*Names Ordered Alphabetically*]
25. H. Xu, A. X. Jiang, **A. Sinha**, Z. Rabinovich, S. Dughmi, M. Tambe. Security Games with Information Leakage: Modeling and Computation, in Proceedings of International Joint Conference on Artificial Intelligence (IJCAI), July 2015.
26. A. Datta, D. Garg, D. Kaynar, D. Sharma, **A. Sinha**. Program Actions as Actual Causes: A Building Block for Accountability, in Proceedings of 28th IEEE Computer Security Foundations Symposium (CSF), July 2015. [*Names Ordered Alphabetically*]
27. Y. D. Abbasi, M. Short, **A. Sinha**, N. Sintov, C. Zhang, M. Tambe. Human Adversaries in Opportunistic Crime Security Games: Evaluating Competing Bounded Rationality Models, in Proceedings of 3rd Conference on Advances in Cognitive Systems (ACS), May 2015.
28. C. Zhang, **A. Sinha**, M. Tambe. Keeping pace with criminals: Designing patrol allocation against adaptive opportunistic criminals, in 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2015.
29. J. Blocki, N. Christin, A. Datta, A. D. Procaccia, **A. Sinha***. Audit Games with Multiple Defender Resources, in Proceedings of 29th AAAI Conference on Artificial Intelligence (AAAI), Jan 2015. [*=*Lead Author*, *Names Ordered Alphabetically*]
30. **A. Sinha**, J. Lia, P. England, J. Lorch. Continuous Tamper-proof Logging Using TPM 2.0, in Proceedings of 7th International Conference on Trust & Trustworthy Computing (TRUST), June 2014.
31. J. Blocki, N. Christin, A. Datta, **A. Sinha**. Adaptive Regret Minimization in Bounded-Memory Games, in Proceedings of 4th Conference on Decision and Game Theory for Security (GameSec), November 2013. [*Names Ordered Alphabetically*]
32. J. Blocki, N. Christin, A. Datta, A. D. Procaccia, **A. Sinha***. Audit Games, in Proceedings of 23rd International Joint Conference on Artificial Intelligence (IJCAI), August 2013. [*=*Lead Author*, *Names Ordered Alphabetically*]
33. J. Blocki, N. Christin, A. Datta, **A. Sinha***. Audit Mechanisms for Provable Risk Management and Accountable Data Governance, in 3rd Conference on Decision and Game Theory for Security (GameSec), November 2012. [*=*Lead Author*, *Names Ordered Alphabetically*]
34. A. Datta, D. Sharma, **A. Sinha**. Provable De-anonymization of Large Datasets with Sparse Dimensions, in Proceedings of Principles of Security and Trust (POST), March 2012. [*Names Ordered Alphabetically*]
35. A. Datta, J. Blocki, N. Christin, H. DeYoung, D. Garg, L. Jia, D. Kaynar, **A. Sinha**. Understanding and protecting privacy: Formal semantics and principled audit mechanisms, in Information Systems Security, pp. 1-27. Springer Berlin Heidelberg, 2011.
36. J. Blocki, N. Christin, A. Datta, **A. Sinha***. Regret Minimizing Audits, A Learning-Theoretic Basis for Privacy Protection, in Proceedings of 24th IEEE Computer Security Foundations Symposium (CSF), June 2011. [*=*Lead Author*, *Names Ordered Alphabetically*]
37. P. Bhallamudi, S. Tilley, **A. Sinha**. Migrating a Web-Based Application to a Service-Based System

– An Experience Report, in Proceedings of the 11th IEEE International Symposium on Web Systems Evolution (WSE 2009)

Book Chapters:

38. S. Mc Carthy, **A. Sinha**, M. Tambe, P. Manadhata. Decision Theory for Network Security: Active Sensing for Detection and Prevention of Data Exfiltration. In Applied Risk Analysis for Guiding Homeland Security Policy and Decisions. John Wiley & Sons Inc. 2017.
39. S. M. Mc Carthy, **A. Sinha**, M.Tambe, Game Theoretic Defense for Maritime Security, in Book on Challenges in Maritime Security (CCICADA Department of Homeland Security), 2016.
40. B. An, M. Tambe, **A. Sinha**. Stackelberg Security Games (SSG) Basics and Application Overview, in Improving Homeland Security Decisions, edited by A. Abbas, M. Tambe, D. Von Winterfeldt, Cambridge University Press, 2016.
41. T.H. Nguyen, D. Kar, M. Brown, **A. Sinha**, A. X. Jiang, M. Tambe. Towards a Science of Security Games, in Interdisciplinary Mathematical Research and Applications, edited by B. Toni, Springer, 2016.

Workshop Papers and Demos:

42. J. Li, X. Wang, Yaoyang Lin, **A. Sinha**, M. P. Wellman. Generating Realistic Stock Market Order Streams, in ICML-19 Workshop on AI in Finance: Applications and Infrastructure for Multi-Agent Learning, June 2019.
43. S. Jecmen, E. Brinkman, **A. Sinha**. Bounding Regret in Simulated Games, in ICML-18 Workshop on Exploration in RL, July 2018.
44. T. H. Nguyen, M. P. Wellman, **A. Sinha**. Deceitful Attacks in Security Games, in the AICS workshop at AAAI 2018.
45. S. Gholami, C. Zhang, **A. Sinha**, M. Tambe. An extensive study of Dynamic Bayesian Network for patrol allocation against adaptive opportunistic criminals, in 2015 IJCAI Workshop on Behavioral, Economic and Computational Intelligence for Security.
46. Y. D. Abbasi, M. Short, **A. Sinha**, N. Sintov, C. Zhang, M. Tambe. Human Adversaries in Opportunistic Crime Security Games: How Past success (or failure) affects future behavior, in 2015 IJCAI Workshop on Behavioral, Economic and Computational Intelligence for Security.
47. D. Kar, F. Fang, F. Delle Fave, N. Sintov, **A. Sinha**, A. Galstyan, B. An, M. Tambe. Learning Bounded Rationality Models of the Adversary in Repeated Stackelberg Security Games, In: AAMAS Adaptive Learning Agents (ALA) Workshop, May 2015.
48. C. Zhang, M. Jain, R. Goyal, **A. Sinha**, M. Tambe. Keeping pace with criminals: Learning, Predicting and Planning against Crime: Demonstration Based on Real Urban Crime Data (Demonstration), in 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2015.
49. **A. Sinha**, Y. Li, L. Bauer. What you want is not what you get: Predicting sharing policies for text-based content on Facebook, in Proceedings of the 6th ACM Workshop on Artificial Intelligence and Security (AISec), November 2013.
50. J. Blocki, N. Christin, **A. Sinha**, A. Sinha. Audit Mechanisms for Privacy Protection in Healthcare Environments, Position paper in USENIX Workshop on Health Security and Privacy (HealthSec '11).

PATENT

- A. Sinha, C. Zhang, M. Tambe. Keeping pace with criminals: Designing patrol allocation against

adaptive opportunistic criminals, US Provisional Application 62/155,315, 2015

SCIENTIFIC COMMUNITY ACTIVITIES

Program Co-chair:

- AAAI workshop on AI and cyber-security (AICS 2016, 2017, 2018, 2019, 2020, co-located with AAAI)
- ACM workshop on AI and Security (AISec 2015, 2016, 2017, co-located with CCS)

Reviewer:

- Senior PC Member:
  - International Joint Conference on AI (IJCAI) 2017
- PC Member:
  - Neural Information Processing Systems (NeurIPS), 2019
  - International Conference on Machine Learning (ICML), 2019
  - AAAI Conference on Artificial Intelligence (AAAI), 2017, 2018, 2019, 2020
  - International Joint Conference on AI (IJCAI) 2015, 2016, 2018, 2019
  - Autonomous Agents and Multiagent Systems (AAMAS) 2015, 2016, 2017, 2018, 2019, 2020
  - IRMAS - Intelligent Robotics and Multi-Agent Systems track of SAC 2016
  - Decision and Game Theory for Security (GameSec) 2016, 2017, 2018, 2019
- External Reviewer (Conference):
  - Network and Distributed System Security Symposium (NDSS) 2016
  - ACM-SIAM Symposium on Discrete Algorithms (SODA) 2016
  - Privacy Enhancing Technologies Symposium (PETS) 2015, 2018, 2019
  - Uncertainty in Artificial Intelligence (UAI) 2014
  - IEEE Security and Privacy 2013
  - ACM Conference on Computer and Communications Security (CCS) 2012
  - ACM AsiaCCS 2013
  - IEEE Computer Security Foundations (CSF) 2012
  - European Symposium on Research in Computer Security (ESORICS) 2012
  - Workshop on Privacy in the Electronic Society (WPES) 2013
- External Reviewer (Journals):
  - Journal of Artificial Intelligence Research
  - ACM Transactions on Economics and Computation
  - Journal of Autonomous Agents and Multi-Agent Systems

INVITED TALKS AND PANELS

- ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems, George Mason University, Fairfax, VA, April 23, 2019

- Workshop on Competitive Economics of Cybersecurity, Sandia National Laboratory, Albuquerque, NM, Nov 2018
- Panel on Adversarial Machine Learning in GameSec Conference, October 2018
- ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems, George Mason University, Fairfax, VA, Feb 27-28, 2018
- Dagstuhl Seminar on Game Theory in AI, Logic, and Algorithms, March 2017
- ICSI seminars, International Computer Science Institute, Berkeley, Feb 2017
- 2016-2017 Homeland Security Symposium Series No. 6, University of Texas at El Paso, Dec 2016
- Max Planck Institute for Software Systems, Kaiserslautern Germany, March 2016
- Research Transition Advisory Committee Meeting, CREATE, University of Southern California (USC): Science of Security Games, July 2015
- CREATE Research Seminar, USC: One Size Does Not Fit All: A Game-Theoretic Approach for Dynamically and Effectively Screening for Threats, May 2015
- Advanced Development for Security Applications Workshop (ADSA), Northeastern University : Science of Security Games, Nov 2014
- Computer Science colloquium, USC: Audit games, March 2014

SUCCESSFUL GRANT PROPOSALS INVOLVED AS PI AND INVESTIGATOR

- Singapore Ministry of Education Academic Research Fund (AcRF) Tier 1 grant on Threat Screening Games, 2019, S$100,000 – **Funded and PI**
- DARPA Spectrum Collaboration Challenge Scoring Methodology, 2017, $125, 000, DARPA – **Funded and co-PI**
- Multi-Scale Network Games of Collusion and Competition, Department of Defense Multidisciplinary University Research Initiatives (MURI), awarded 2018, $6,250,000 – **Funded and Investigator**

GRANT PROPOSALS HELPED ON AS POSTDOC AND PHD

- Towards a Science of Cyber-Security Games, 7/15/15-7/15/18, $750, 000 Army Research Office – **Funded**
- Hardening a Distributed CPS through Rational and Dynamic Decision-Making Among Multiple Stakeholders, Submitted to National Science Foundation Frontier
- Adversarial Cyber Security Game among Attackers, Defenders, and Users, Submitted to National Science Foundation SaTC program
- Towards OPtimized SECurity at Ports of Entry (TOP-SEC), Submitted to National Science Foundation SaTC program
- Hardening the Power Grid: Rational and Dynamic Decision-Making Among Multiple Stakeholders, DHS Broad Agency Announcement Call HSHQDC-14-R-B0016
- Adaptive Defense Mechanisms for Adversarial Environments, Department of Defense Multi-disciplinary Research Program of the University Research Initiative

PARTICIPATION IN THESIS COMMITTEES

- *Erik Brinkmann*: Ph.D. thesis committee member (2017), Computer Science and Engineering, University of Michigan.

TEACHING AND MENTORING

I advised three master student in the University of Michigan:

- *Yongzhao Wang*: I mentored and worked with Yongzhao on adversarial deception in security games and published one conference paper with him.
- *Linh Nguyen*: I mentored and worked with Linh on adversarial learning and published one conference paper with him.
- *Junyi Li*: I mentored and worked with Junyi on the aspects of manipulation in complex systems such as stock markets. I submitted a paper with him.

I also worked with two undergraduate students in the University of Michigan: Sky Wang and Yaoyang Lin. Besides these, I am also involved in research mentoring of under-represented minorities as part of the "Explore CS Research" at the University of Michigan.

I have mentored many Ph.D. students in my post doc as well as when I was a senior Ph.D. student. They include students at University of Southern California

- *Aaron Schlenker*: I mentored and worked with Aaron on TSA project of optimally screening passengers at airports and its follow up on screening of cyber alerts.
- *Chao Zhang*: I mentored and worked with Chao on the problem of predicting opportunistic crime and planning optimal patrols against the same.
- *Matthew Brown*: I mentored and worked with Matthew on the TSA project of optimally screening passengers at airports.
- *Haifeng Xu*: I helped Haifeng on the work on information leakage in security games, and also helped him with writing skills.
- *Yasaman Abbasi*: I mentored Yasaman with the work on human behavior models of the criminal in opportunistic crime setting.
- *Shahrzad Gholami*: I mentored 1^{st} year Ph.D. student Shahrzad to catch up on the technical material and contribute to the project on opportunistic crimes.

and students at Carnegie Mellon University

- *Divya Sharma*: I mentored Divya and helped her with her first paper on linking attacks and also provided directions for her Ph.D. thesis on causality.
- *Alejandro Carbonara*: I mentored (then) 1^{st} year Ph.D. student Alejandro to catch up on the technical material and helped him with his first paper published in IJCAI.

Teaching experience:
Courses taught at Singapore Management University:

- Computational Thinking (2019 Term 1)
- Heuristic Search and Optimization (2020 Term 2)

Other lectures:

- Guest lecture on Game Theory in "EECS 492: Introduction to Artificial Intelligence" at University of Michigan

- Taught two lecture on PAC learning for course "Security Games" (2016) by Milind Tambe at University of Southern California
- Teaching assistant for course "Foundations of Privacy" (2013) by Anupam Datta at Carnegie Mellon University
  - Designed and graded assignments for this new course
  - Taught all recitation classes, teaching additional material not covered in class
  - Taught classes when Anupam Datta was out of town
- Teaching assistant for course "Foundations of Security and Privacy" (2011) by Anupam Datta at Carnegie Mellon University
  - Taught all recitation classes and graded assignments
  - Taught classes when Anupam Datta was out of town