Thanh H. Nguyen

Associate Professor
Computer Sciene Department
University of Oregon
https://pages.uoregon.edu/tnguye11/

Deschutes Hall
Eugene, OR 97403-1202
Phone: +1 541-346-3974
Email: tnguye11@uoregon.edu

Research Interests: Artificial Intelligence, Multi-Agent Systems, Machine Learning, Reinforcement Learning, Game Theory, Optimization.

EDUCATION

University of Southern California, Los Angeles, California. August 2011 – May 2016.
Ph.D. Student, Department of Computer Science.
Adviser: Prof. Milind Tambe.

Hanoi University of Science and Technology, Hanoi, Vietnam. August 2005 – July 2010.
B.Sc., Center for Training of Excellent Students.

AWARDS

- **Deployed Application Award, Innovative Applications of Artificial Intelligence (IAAI), 2016:** The award is presented to Fei Fang, **Thanh H. Nguyen**, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe and Andrew Lemieux for the deployed application paper titled "Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security."
- **Runner-up of the Best Innovative Application Paper Award at the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2016:** This is presented to **Thanh H. Nguyen**, Arunesh Sinha, Shahrzad Gholami, Andrew Plumptre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, and Colin Beale for the paper titled "CAPTURE: A New Predictive Anti-Poaching Tool for Wildlife Protection."

EXPERIENCE

- **University of Oregon, Eugene, Oregon,** September 2025 – Present.
Associate Professor in Computer Science. Develop novel models and algorithms in Reinforcement Learning, Machine Learning, and Game Theory to address challenges in multi-agent decision-making under uncertain and adversarial environments, and apply these techniques to interdisciplinary problems spanning fields such as physics and healthcare.
- **University of Oregon, Eugene, Oregon,** September 2018 – September 2025.
Assistant Professor in Computer Science. Focus on applying techniques from Artificial Intelligence, Multi-agent Systems, Machine Learning, and Game Theory for addressing real-world challenges in the areas of Public Safety and Security, Public Health, and Conservation.
- **University of Michigan, Ann Arbor, Michigan,** July 2016 – August 2018.
Postdoctoral Researcher with Prof. Michael P. Wellman and Prof. Satinder Singh. Focus on game-theoretic modeling and solution techniques for cybersecurity problems. Develop scalable algorithms for computing optimal defensive strategies to protect computer networks from cyber-attacks.

- **University of Southern California, Los Angeles, California,** August 2011 – May 2016.
  Research Assistant with TEAMCORE Research Group. Focus on game-theoretic modeling, robust solution techniques, and behavioral modeling for real-world security applications. Develop algorithms to compute the optimal resource allocation strategies in Stackelberg games with uncertainties. Build new behavioral models of adversaries based on real-world security data.

- **VNG Corporation, Vietnam,** July 2010 – July 2011.
  Researcher and Developer. Worked on developing algorithms for recommender systems such as friend suggestion or feed ranking in the largest Vietnamese online social network, Zing Me.

- **High Performance Computing Center, Hanoi University of Science and Technology, Vietnam,** September 2009 – June 2010.
  Research Intern. Worked with Prof. Thuy Thanh Tran on developing algorithms for document matching in a grid-based system.

- **Image Processing Research Group, FPT Software Company, Vietnam,** April 2008 – August 2009.
  Research Intern. Worked on developing algorithms, i.e., image compression and image content retrieval under cooperation with the Fujifilm company in Japan.

REAL-WORLD APPLICATIONS

- **PAWS.** I have contributed to developing the wildlife-protection application, PAWS (Protection Assistant for Wildlife Security). I led the project on wildlife protection in Indonesia in 2015 and participated in extending the application for protecting tigers in Malaysia in 2016. In this work, I collaborated with different NGOs, including World Wildlife Fund, Panthera, Rimba, and Wildlife Conservation Society. My work has led to new research extending PAWS for wildlife protection in Uganda in 2017. PAWS has been extensively tested and deployed in both Malaysia and Uganda.

FUNDED PROPOSALS

- *[Co PI] Knowledge Graph-Regularized Reinforcement Learning for Tumor Microenvironment Spatial Omics*, Total Funding: $\$100,000$ (PI portion: $\$25,000$), CBDS UO-OHSU Collaborative Project Awards, 2025 - 2027.

- *[Co PI] Visualizing Defects and Dopants in Three-Dimensions using Ptychography and Physics-Informed Machine Learning*, Total Funding: $\$185,762$. Intel Corporation, 2025 - 2026.

- *[Lead PI] Adversarial Reasoning: Tackling Sequential and Coordinated Attacks in Security Domains with Real-time Information*, approximately $\$343,000$. Army Research Office, Grant W911NF-20-1-0344, 2020 - 2024.

LIST OF PUBLICATIONS

Refereed Journal Articles

1. **Thanh H. Nguyen**, Amulya Yadav. A Complete Analysis on the Risk of Using Quantal Response: When Attacker Maliciously Changes Behavior under Uncertainty. In *the Games Journal, 2022, 13(6), 81. doi: 10.3390/g13060081.*

2. **Thanh H. Nguyen**, Mason Wright, Michael P. Wellman, Satinder Singh. Multi-Stage Attack Graph Security Games: Heuristic Strategies, with Empirical Game-Theoretic Analysis. In *the Journal of the Security and Communication Networks, 2018.*

3. (Joint Lead Author) Arunesh Sinha, **Thanh H. Nguyen**, Debarun Kar, Matthew Brown, Milind Tambe, Albert Xin Jiang. From Physical Security to Cyber Security. In *Journal of Cyber Security 2015; 1 (1): 19-35. doi: 10.1093/cybsec/tyv007.*

4. Jason Tsai, **Thanh H. Nguyen**, Nicholas Weller, Milind Tambe. Game-Theoretic Target Selection in Contagion-based Domains. In *The Computer Journal, 57(6):893–905, 2014*.

Refereed Conference Papers

5. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. SamaDICE: Safe Multi-Agent Reinforcement Learning with Stationary Distribution Correction Estimation. In *Proceedings of 40th Conference on Neural Information Processing Systems (NeurIPS), December 2026.*

6. Hai An Vu, Phung Gia Huy, Minh-Phuc Truong, Linh Ngo Van, **Thanh H. Nguyen**, Thanh-Toan Do, Trung Le. SUGAR: Spectral and Geometry-Aware Alignment for Matryoshka Representation Distillation. In *Proceedings of 31st Conference on Empirical Methods in Natural Language Processing (EMNLP), October 2026.*

7. Pham Khanh Chi, Quoc Phong Dao, Thuat Nguyen, Linh Ngo Van, Trung Le, **Thanh H. Nguyen**. MTA: Multi-Granular Trajectory Alignment for Large Language Model Distillation. In *Proceedings of 64th Annual Meeting of the Association for Computational Linguistics (ACL), July 2026.*

8. Phung Gia Huy, Hai An Vu, Minh-Phuc Truong, Thang Duc Tran, Linh Ngo Van, **Thanh H. Nguyen**, Trung Le. MIPIC: Matryoshka Representation Learning via Self-Distilled Intra-Relational and Progressive Information Chaining. In *Findings of the Association for Computational Linguistics (ACL), July 2026.*

9. Truong Nguyen, Phi Van Dat, Ngan Nguyen, Linh Ngo Van, Trung Le, **Thanh H. Nguyen**. CTPD: Cross Tokenizer Preference Distillation. In *Proceedings of 40th Annual AAAI Conference on Artificial Intelligence (AAAI), Special Track on AI Alignment, January 2026.*

10. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. MisoDICE: Multi-Agent Imitation from Mixed-Quality Demonstrations. In *Proceedings of 39th Conference on Neural Information Processing Systems (NeurIPS), December 2025.*

11. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. O-MAPL: Offline Multi-Agent Preference Learning. In *Proceedings of 42nd International Conference on Machine Learning (ICML), July 2025.*

12. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. ComaDICE: Offline Cooperative Multi-Agent Reinforcement Learning with Stationary Distribution Shift Regularization. In *Proceedings of The 13th International Conference on Learning Representations (ICLR), April 2025.*

13. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. Inverse Factorized Soft Q-Learning for Cooperative Multi-agent Imitation Learning. In *Proceedings of 38th Conference on Neural Information Processing Systems (NeurIPS), December 2024.*

14. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. Mimicking To Dominate: Imitation Learning Strategies for Success in Multiagent Games. In *Proceedings of 38th Conference on Neural Information Processing Systems (NeurIPS), December 2024.*

15. Tien Mai, Avinandan Bose, Nalini Saligram, Arunesh Sinha, **Thanh H. Nguyen**, and Ayushman Kumar-Singh. Tackling Stackelberg Network Interdiction against a Boundedly Rational Adversary. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI), August 2024.*

16. Roman Belaire, Pradeep Varakantham, **Thanh H. Nguyen**, David Lo. Regret-based Defense in Adversarial Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2024.*

17. Changyu Chen, Ramesha Karunasena, **Thanh H. Nguyen**, Arunesh Sinha, Pradeep Varakantham. Generative Modelling of Stochastic Actions with Arbitrary Constraints in Reinforcement Learning. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS), December 2023.*

18. Hangzhi Guo, **Thanh H. Nguyen**, Amulya Yadav. CounterNet: End-to-End Training of Prediction Aware Counterfactual Explanations. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), August 2023.*

19. Sarah Kinsey, Jack Wolf, Nalini Saligram, Varun Ramesan, Meeta Walavalkar, Nidhi Jaswal, Sandhya Ramalingam, Arunesh Sinha, **Thanh H. Nguyen**. Building a Personalized Messaging System for Health Intervention in Underprivileged Regions Using Reinforcement Learning. In *Proceedings of the 32th International Joint Conference on Artificial Intelligence (IJCAI), AI for Good Track, August 2023.*

20. The Viet Bui, Tien Mai, **Thanh H. Nguyen**. Imitating Opponent to Win: Adversarial Policy Imitation Learning in Two-player Competitive Games. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2023.*

21. Wai Tuck Wong, Sarah Kinsey, Ramesha Karunasena, **Thanh H. Nguyen**, Arunesh Sinha. Beyond NaN: Resiliency of Optimization Layers in the Face of Instability. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI), February 2023.*

22. **Thanh H. Nguyen**, Arunesh Sinha. Behavioral Learning in Security Games: Threat of Multi-Step Manipulative Attacks. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI), February 2023.*

23. Chenghan Zhou, **Thanh H. Nguyen**, Haifeng Xu. Algorithmic Information Design in Multi-Player Games: Possibilities and Limits in Singleton Congestion. In *Proceedings of the 23th ACM Conference on Economics and Computation (EC), July 2022.*

24. Chenghan Zhou, Andrew Spivey, Haifeng Xu, **Thanh H. Nguyen**. Information Design for Multiple Independent and Self-Interested Defenders: Work Less, Pay Off More. In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI), August 2022.*

25. **Thanh H. Nguyen**, Haifeng Xu. When Can the Defender Effectively Deceive Attackers in Security Games?. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI), February 2022 (Acceptance rate: 15%).*

26. **Thanh H. Nguyen**, Amulya Yadav. The Risk of Attacker Behavioral Learning: Can Attacker Fool Defender under Uncertainty? In *Proceedings of the 13th Conference on Decision and Game Theory for Security (GameSec), October 2022.*

27. Sarah Kinsey, Wong Wei Tuck, Arunesh Sinha, **Thanh H. Nguyen**. An Exploration of Poisoning Attacks on Data-based Decision Making. In *Proceedings of the 13th Conference on Decision and Game Theory for Security (GameSec), October 2022.*

28. Ryan Butler, Arunesh Sinha, **Thanh H. Nguyen**. Countering Attacker Data Manipulation in Security Games. In *Proceedings of the 12th Conference on Decision and Game Theory for Security (GameSec), October 2021.*

29. **Thanh H. Nguyen**, Arunesh Sinha, He He. Partial Adversarial Behavior Deception in Security Games. In *Procceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI), January 2021 (Acceptance rate: 12.6%).*

30. Petr Tomášek, Branislav Bošanský, **Thanh H. Nguyen**. Using One-Sided Partially Observable Stochastic Games for Solving Zero-Sum Security Games with Sequential Attacks. In *Proceedings of the 11th Conference on Decision and Game Theory for Security (GameSec), November 2020.*

31. **Thanh H. Nguyen**, Nam Vu, Amulya Yadav, Uy Nguyen. Decoding the Imitation Security Game: Handling Attacker Imitative Behavior Deception. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI), August 2020.*

32. **Thanh H. Nguyen**, Andrew Butler, Haifeng Xu. Tackling Imitative Attacker Deception in Repeated Bayesian Stackelberg Security Games. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI), August 2020.*

33. Yebo Feng, Jun Li, **Thanh H. Nguyen**. Application-Layer DDoS Defense with Reinforcement Learning. In *Proceedings of the IEEE/ACM 28th International Symposium on Quality of Service (IWQoS), June 2020.*

34. **Thanh H. Nguyen**, Amulya Yadav, Branislav Bosansky, Yu Liang. Tackling Sequential Attacks in Security Games. In *Proceedings of the 10th Conference on Decision and Game Theory for Security (GameSec), November 2019.*

35. Sarah Cooney, Kai Wang, Elizabeth Bondi, **Thanh H. Nguyen**, Phebe Vayanos, Hailey Winetrobe, Edward A. Cranford, Cleotilde Gonzalez, Christian Lebiere, Milind Tambe. Learning to Signal in the Goldilocks Zone: Improving Adversary Compliance in Security Games. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECMLPKDD), September 2019 (Acceptance rate: 17.7%).*

36. **Thanh H. Nguyen**, Haifeng Xu. Imitative Attacker Deception in Stackelberg Security Games. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), August 2019 (Acceptance rate: 17.9%).*

37. Sarah Cooney, Phebe Vayanos, **Thanh H. Nguyen**, Cleotilde Gonzalez, Christian Lebiere, Edward A. Cranford, Milind Tambe. Warning Time: Optimizing Strategic Signaling for Security Against Boundedly Rational Adversaries (Extended Abstract). In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2019.*

38. **Thanh H. Nguyen**, Yongzhao Wang, Arunesh Sinha, Michael P. Wellman. Deception in Finitely Repeated Security Games. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), January 2019 (Acceptance rate: 16.2%).*

39. **Thanh H. Nguyen**, Michael P. Wellman, Satinder Singh. A Stackelberg Game Model for Botnet Data Exfiltration. In *Proceedings of the 7th Conference on Decision and Game Theory for Security (GameSec), October 2017.*

40. Anjon Basak, Fei Fang, **Thanh H. Nguyen**, Christopher Kiekintveld. Combining Graph Contraction and Strategy Generation for Green Security Games. In *Proceedings of the 7th Conference on Decision and Game Theory for Security (GameSec), November 2016.*

41. Nika Haghtalab, Fei Fang, **Thanh H. Nguyen**, Arunesh Sinha, Ariel Procaccia, Milind Tambe. Three strategies to success: Learning adversary models in security games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), July 2016.*

42. **Thanh H. Nguyen**, Arunesh Sinha, Shahrzad Gholami, Andrew Plumptre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, Colin Beale. CAPTURE: A New Predictive Anti-Poaching Tool for Wildlife Protection. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2016.* **Runner-up of the Best Innovative Application Paper Award in AAMAS-16.**

43. Fei Fang, **Thanh H. Nguyen**, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, Andrew Lemieux. Deploying PAWS: Field Optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the 28th Innovative Applications of Artificial Intelligence Conference (IAAI), February 2016.* **Winner of Deployed Application Award in IAAI-16.**

44. **Thanh H. Nguyen**, Francesco M. Delle Fave, Debarun Kar, Aravind D. Lakshminarayanan, Amulya Yadav, Milind Tambe, Noa Agmon, Andrew J. Plumptre, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba. Making the most of Our Regrets: Regret-based Solutions to Handle Payoff Uncertainty and Elicitation in Green Security Games. In *Proceedings of the 6th Conference on Decision and Game Theory for Security (GameSec), November 2015.*

45. Benjamin Ford, **Thanh H. Nguyen**, Nicole Sintov, Milind Tambe, Francesco Delle Fave. Beware the Soothsayer: Evaluating the Reliability of Attack Predictions in Stackelberg and Network Security Games. In *Proceedings of the 6th Conference on Decision and Game Theory for Security (GameSec), November 2015.*

46. **Thanh H. Nguyen**, Amulya Yadav, Bo An, Milind Tambe, Craig Boutilier. Regret-based Optimization and Preference Elicitation for Stackelberg Security Games with Uncertainty. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI), July 2014.*

47. **Thanh H. Nguyen**, Albert Xin Jiang, Milind Tambe. Stop the Compartmentalization: Unified Robust Algorithms for Handling Uncertainties in Security Games. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2014.*

48. Albert Xin Jiang, **Thanh H. Nguyen**, Milind Tambe, Ariel D. Procaccia. Monotonic Maximin: A Robust Stackelberg Solution Against Boundedly Rational Followers. In *Proceedings of the Conference on Decision and Game Theory for Security (GameSec), 2013.*

49. **Thanh H. Nguyen**, Rong Yang, Amos Azaria, Sarit Kraus, Milind Tambe. Analyzing the Effectiveness of Adversary Modeling in Security Games. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI), July 2013.*

50. **Thanh H. Nguyen**, Amos Azaria, James Pita, Rajiv Maheswaran, Sarit Kraus, Milind Tambe. Modeling Human Adversary Decision Making in Security Games: An Initial Report. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (Short paper), May 2013.*

51. Jason Tsai, **Thanh H. Nguyen**, Milind Tambe. Security Games for Controlling Contagion. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI), July 2012.*

Refereed Technical Magazine Articles

52. Fei Fang, **Thanh H. Nguyen**, Arunesh Sinha, Shahrzad Gholami, Andrew Plumptre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, Colin Beale. Predicting Poaching for Wildlife Protection. In *IBM Journal of Research and Development, 61(6):3:1-3:12, 2017.*

53. Nicole Sintov, Debarun Kar, **Thanh H. Nguyen**, Fei Fang, Kevin Hoffman, Arnaud Lyet, Milind Tambe. Keeping it Real: Using Real-World Problems to Teach AI to Diverse Audiences. In *AI Magazine, 38 (2):35-47, 2017.*

54. Fei Fang, **Thanh H. Nguyen**, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, Andrew Lemieux. PAWS – A Deployed Game-Theoretic Application to Combat Poaching. In *AI Magazine, 38(1):23-36, 2017.*

Book Chapters

55. Haifeng Xu, **Thanh H. Nguyen**. Be Careful When Learning Against Adversaries: Imitative Attacker Deception in Stackelberg Security Games. In *Adversary-Aware Learning Techniques and Trends in Cybersecurity*. Springer, 2020. DOI:10.1007/978-3-030-55692-1_3.

56. Massimiliano Albanese, Sushil Jajodia, Sridhar Venkatesan, George Cybenko, **Thanh H. Nguyen**. Adaptive Cyber Defenses for Botnet Detection and Mitigation. In *Adversarial and Uncertain Reasoning for Adaptive Cyber Defense*. Springer, 2019. DOI:10.1007/978-3-030-30719-6_2.

57. Michael P. Wellman, **Thanh H. Nguyen**, Mason Wright. Empirical Game-Theoretic Methods for Adaptive Cyber-Defense. In *Adversarial and Uncertain Reasoning for Adaptive Cyber Defense*. Springer, 2019. DOI:10.1007/978-3-030-30719-6_2.

58. Debarun Kar, **Thanh H. Nguyen**, Fei Fang, Matthew Brown, Arunesh Sinha, Milind Tambe, Albert Xin Jiang. Recent Trends and Applications in Security Games. In *Handbook on Dynamic Game Theory (Edited by Tamar Basar and Georges Zaccour)*. Springer, 2017. DOI:10.1007/978-3-319-27335-8_27-1.

59. (Joint Lead Author) Tracy Cui, **Thanh H. Nguyen**, James Pita, Richard S. John. Methods for Addressing the Unpredictable Human Element in Security. In *Improving Homeland Security Decisions, CREATE, 2017*.

60. **Thanh H. Nguyen**, Debarun Kar, Matthew Brown, Arunesh Sinha, Albert Xin Jiang, Milind Tambe. Towards a Science of Security Games. In *New Frontiers of Multidisciplinary Research in STEAM-H, 2016*.

Symposium Publications

61. Nicole Sintov, Debarun Kar, **Thanh H. Nguyen**, Fei Fang, Kevin Hoffman, Arnaud Lyet, Milind Tambe. From the Lab to the Classroom and Beyond: Extending a Game-Based Research Platform for Teaching AI to Diverse Audiences. In *Proceedings of the Symposium on Educational Advances in Artificial Intelligence (EAAI), Feb 2016.*

62. Arjun Tambe, **Thanh H. Nguyen**. Robust Resource Allocation in Security Games and Ensemble Modeling of Adversary Behavior. In *Proceedings of the ACM Symposium on Applied Computing (ACM SAC) Track, April 2015.*

63. **Thanh H. Nguyen**, Jason Tsai, Albert Jiang, Emma Bowring, Rajiv Maheswaran, Milind Tambe. Security Games on Social Networks. In *Proceedings of the AAAI Fall Symposium, November 2012.*

Refereed Workshop Papers

64. **Thanh H. Nguyen**, Arunesh Sinha. The Art of Manipulation: Threat of Multi-Step Manipulative Attacks in Security Games. In *Artificial Intelligence for Cyber Security (AICS) held at the 36th AAAI Conference on Artificial Intelligence (AAAI), February 2022.*

65. Ryan Butler, Wong Wai Tuck, Arunesh Sinha, **Thanh H. Nguyen**. Poisoning Attacks on Data-based Decision Making: A Preliminary Study. In *3rd Autonomous Agents for Social Good (AASG) held at the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), May 2022.*

66. Chenghan Zhou, **Thanh H. Nguyen**, Haifeng Xu. Algorithmic Information Design for Singleton Congestion Games. In *2nd Workshop on Graphs and More Complex Structures for Learning and Reasoning (GCLR) held at the 36th AAAI Conference on Artificial Intelligence (AAAI), February 2022.*

67. Hangzhi Guo, **Thanh H. Nguyen**, Amulya Yadav. CounterNet: End-to-End Training of Counterfactual Aware Predictions. In *ICML Workshop on Algorithmic Recourse held at the Thirty-eighth International Conference on Machine Learning (ICML), July 2021.* **Runner-Up of the Best Paper Award.**

68. **Thanh H. Nguyen**, Arunesh Sinha. Sequential Manipulative Attacks in Security Games. In *International Workshop on Autonomous Agents for Social Good (AASG) held at the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2021.*

69. Stephanie Milani, Amulya Yadav, Fei Fang, **Thanh H. Nguyen**, Zhou Fan and Saurabh Gulati. Intelligent Tutoring Strategies for Students with Autism Spectrum Disorder: A Reinforcement Learning Approach. In *Workshop on Artifical Intelligence for Education held at the 34th AAAI Conference on Artificial Intelligence (AAAI), February 2020.*

70. **Thanh H. Nguyen**, Michael P. Wellman, Arunesh Sinha. Deceitful Attacks in Security Games. In *Workshop on Artificial Intelligence for Cyber Security held at the 32nd AAAI Conference on Artificial Intelligence (AAAI), February 2018.*

71. **Thanh H. Nguyen**, Mason Wright, Michael P. Wellman, Satinder Singh. Multi-Stage Attack Graph Security Games with Empirical Game Theoretic Analysis. In *ACM Workshop on Moving Target Defense (MTD), October 2017.*

72. **Thanh H. Nguyen**, Michael P. Wellman, Satinder Singh. A Stackelberg Game Model for Botnet Traffic Exfiltration. In *Workshop on Artificial Intelligence for Cyber Security held at the 31st AAAI Conference on Artificial Intelligence (AAAI), February 2017.*

73. **Thanh H. Nguyen**, Arunesh Sinha, Milind Tambe. Addressing Behavioral Uncertainty in Security Games: An Efficient Robust Strategic Solution for Defender Patrols. In *IEEE Workshop on Parallel and Distributed Processing for Computational Social Systems, May 2016.*

74. **Thanh H. Nguyen**, Arunesh Sinha, Shahrzad Gholami, Andrew Plumptre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, Colin Beale. Protecting Wildlife under Imperfect Observation. In *Workshop on Computer Poker and Imperfect Information Games held at the 30th AAAI Conference on Artificial Intelligence (AAAI), February 2016.*

75. Fei Fang, **Thanh H. Nguyen**, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, Andrew Lemieux. Deploying PAWS in the Field: Designing Efficient Patrols to Combat Poaching. In *Workshop on Computer Poker and Imperfect Information Games held at the 30th AAAI Conference on Artificial Intelligence (AAAI), February 2016.*

76. Amulya Yadav, **Thanh H. Nguyen**, Francesco Delle Fave, Milind Tambe, Noa Agmon, Manish Jain, Widodo Ramono, Timbul Batubara. Handling Payoff Uncertainty with Adversary Bounded Rationality in Green Security Domains. In *Workshop on Algorithmic Game Theory (AGT) held at the 24th International Joint Conference on Artificial Intelligence (IJCAI), July 2015.*

77. Fei Fang, **Thanh H. Nguyen**, Benjamin Ford, Nicole Sintov, Milind Tambe. Introduction to Green Security Games. In *Workshop on Cognitive Knowledge Acquisition and Applications (Cognitum) held at the 24th International Joint Conference on Artificial Intelligence (IJCAI) (Extended Abstract), July 2015.*

78. Fei Fang, **Thanh H. Nguyen**, Bo An, Milind Tambe, Rob Pickles, Wai Lam, Gopalasamy Clements. Towards Addressing Challenges in Green Security Games in the Wild. In *Workshop on Behavioral, Economic and Computational Intelligence for Security (BECIS) held at the 24th International Joint Conference on Artificial Intelligence (IJCAI), July 2015.*

79. **Thanh H. Nguyen**, Amulya Yadav, Francesco Delle Fave, Milind Tambe, Noa Agmon, Manish Jain, Richard Van Deventer. Behavioral Minimax Regret for Security Games and Its Application for UAV Planning. In *Workshop on Optimisation in Multi-Agent Systems (OptMAS) held at the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2015.*

80. Fei Fang, **Thanh Nguyen**, Rob Pickles, Lam Wai Yee, Milind Tambe. Challenges of Green Security Games in the Wild. In *Workshop on Issues with Deployment of Emerging Agent-based Systems (IDEAS) held at the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2015.*

81. **Thanh H. Nguyen**, Rong Yang, Amos Azaria, Sarit Kraus, Milind Tambe. The Power of Modeling Human Adversary Behaviors in Security Games. In *Workshop on Human-Agent Interaction Design and Models (HAIDM) held at the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2013.*

82. Jason Tsai, **Thanh H. Nguyen**, Milind Tambe. Game-Theoretic Target Selection in Contagion-based Domains. In *Workshop on Optimization in Multiagent Systems (OPTMAS) held at 11th the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), July 2012.*

Student Abstracts

83. **Thanh H. Nguyen**, Arunesh Sinha, Milind Tambe. Conquering Adversary Behavioral Uncertainty in Security Games: An Efficient Modeling Robust based Algorithm. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI) (Student Abstract), February 2016.*

Demos

84. Fei Fang, **Thanh H. Nguyen**, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe. Deploying PAWS to Combat Poaching: Game-theoretic Patrolling in Areas with Complex Terrains. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI) (Demonstration), February 2016.*

TALKS

- *School of Electrical Engineering and Computer Science, Oregon State University:* Robust Decision-Making with Limited and Adversarial Data. December 2025.
- *Computer Science and Engineering Department, Washington University in St. Louis:* Data-Driven Decision Making in Adversarial Environments: Challenges and Applications. April 2024.
- *Computer Science Department, George Mason University:* Data-Driven Decision Making in Adversarial Environments: Challenges and Applications. April 2024.
- *Computer Science Department, Rensselaer Polytechnic Institute:* Data-Driven Decision Making in Adversarial Environments: Challenges and Applications. April 2024.
- *Department of Computer Science, University of Chicago:* Data-Driven Decision Making in Adversarial Environments: Challenges and Applications. April 2024.
- *Software and Societal Systems Department, School of Computer Science, Carnegie Mellon University:* Data-Driven Decision Making in Adversarial Environments: Challenges and Applications. March 2024.
- *School of Information Systems, Singapore Management University, Singapore:* AI for Social Impact: Data-Driven Decision Making. July 2023.
- *Department of Mathematics, Oregon State University, United States:* Data-Centered Decision Making. January 2023.
- *School of Computer Science and Engineering, Nanyang Technological University, Singapore:* Vulnerabilities in Data-Centered Decision Making. July 2022.

PROFESSIONAL ACTIVITIES

- **Senior Area Chair** of the *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*: 2027.
- **Area Chair** of the *International Conference on Learning Representations (ICLR)*: 2027.
- **Area Chair** of the *International Conference on Machine Learning (ICML)*: 2027.
- **Senior Program Committee member** of the *AAAI Conference on Artificial Intelligence (AAAI)*: 2027.
- **Program Committee member** of the *International Conference on Learning Representations (ICLR)*: 2026.
- **Program Committee member** of the *International Conference on Machine Learning (ICML)*: 2025 – 2026.
- **Program Committee member** of the *Annual Meeting of the Association for Computational Linguistics (ACL)*: 2026.
- **Program Committee member** of the *Annual Conference on Neural Information Processing Systems (NeurIPS)*: 2025.
- **Senior Program Committee member** of the *International Joint Conference on Artificial Intelligence (IJCAI) (Special Track on AI and Social Good)*: 2026.
- **Senior Area Chair** of the *International Joint Conference on Artificial Intelligence (IJCAI)*: 2025.
- **Co-Area Chair** of the *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS) (Innovative Applications area/track)*: 2025.
- **Senior Program Committee member** of the *AAAI Conference on Artificial Intelligence (AAAI) (Social Impact Track)*: 2020 – 2022, 2024 – 2026.
- **Program Committee member** of the *AAAI Conference on Artificial Intelligence (AAAI) (Main Track)*: 2017 – 2021, 2023 – 2026.
- **Program Committee member** of the *International Joint Conference on Artificial Intelligence (IJCAI)*: 2015 – 2017, 2019 – 2024.
- **Program Committee member** of the *Conference on Uncertainty in Artificial Intelligence (UAI)*: 2024 – 2025.
- **Program Committee member** of the *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*: 2016, 2018 – 2023.
- **Program Committee member** of the *European Conference on Artificial Intelligence (ECAI)*: 2024.
- **Program Committee member** of the *Conference on Decision and Game Theory for Security (GameSec)*: 2017, 2021, 2022, 2024.
- **Program Committee member** of the *IEEE/WIC International Conference on Intelligence Agent Technology (IAT)*: 2015.
- **Co-chair** of *the 3rd Workshop on Artificial Intelligence for Social Good (AI4SG)*, held at IJCAI 2021.
- **Co-chair** of *the Workshop on Data Science for Social Good (DSSG)*, held at KDD 2021.

- **Co-chair** of the *Strategic Reasoning for Societal Challenges (SRSC) Workshop*, at 18th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS), May 2019.
- **Co-chair** of the *Workshop on Computer Poker and Imperfect Information Games*, at 31st AAAI Conference on Artificial Intelligence (AAAI), February 2017.
- **Co-chair** of the *Workshop on Computer Poker and Imperfect Information Games*, at 30th AAAI Conference on Artificial Intelligence (AAAI), February 2016.
- **Co-organizer** of the *Conference on Conservation, Criminology and Computation (C4)*, at Washington DC office of the University of Southern California, June 2015.
- **Co-organizer** of the *Workshop on Wildlife Crime: An Interdisciplinary Perspective*, at Washington DC office of the University of Southern California, July 2014.