

Kai Wang

☎ 626-588-8197 | ✉ kaiwang.gua@gmail.com | 🏠 guaguakai.com | 📱 [guaguakai](https://www.instagram.com/guaguakai) | 📺 [guaguakai](https://www.youtube.com/guaguakai) | 🐦 [kaiwang_gua](https://twitter.com/kaiwang_gua)

Academic Appointments

Georgia Institute of Technology

Assistant Professor in the School of Computational Science and Engineering

2024 Jan-
GA, United States

Massachusetts Institute of Technology

Postdoctoral Researcher

Hosted by Marzyeh Ghassemi (CSAIL) and Suvrit Sra (LIDS)

2023 Aug - 2023 Dec
MA, United States

Education

Harvard University

Ph.D. in Computer Science

Advised by Professor Milind Tambe

2017 - 2023
MA, United States

National Taiwan University

B.S. in Electrical Engineering and Mathematics

2012 - 2017
Taipei, Taiwan

Honors

Siebel Scholars

Scholarship (awarded to a small fraction of students at the “world’s leading graduate schools”)

2023

Taiwan Ministry of Education

Scholarship (awarded to outstanding graduate students doing exceptional and innovative research)

2021

Association for the Advancement of Artificial Intelligence (AAAI)

Best Paper Runner Up

2021

Asian Pacific Alumni Association

Scholarship (awarded to students who have significant impact in society)

2018

International Mathematical Olympiad

Silver Medal

2010, 2011

Conference Publications

18. Sonja Johnson-Yu, Jessica Finocchiaro, Arunesh Sinha, **Kai Wang**, Yevgeniy Vorobeychik, Aparna Taneja, Milind Tambe “Characterizing and Improving the Robustness of Predict-Then-Optimize Frameworks”, *Under Review at NeurIPS 2023*
17. Shresth Verma, Aditya Mate, **Kai Wang**, Neha Madhiwalla, Aparna Hegde, Aparna Taneja, Milind Tambe “Restless Multi-Armed Bandits for Maternal and Child Health: Results from Decision-Focused Learning”, *International Conference On Autonomous Agents and Multi-Agent Systems, AAMAS 2023*
16. **Kai Wang***, Lily Xu*, Aparna Taneja, Milind Tambe “Optimistic Whittle Index Policy: Online Learning for Restless Bandits”, *Association for the Advancement of Artificial Intelligence, AAAI 2023*
15. **Kai Wang***, Shresth Verma*, Aditya Mate, Sanket Shah, Aparna Taneja, Neha Madhiwalla, Aparna Hegde, and Milind Tambe “Decision-Focused Learning in Restless Multi-Armed Bandits with Application to Maternal and Child Care Domain”, *Association for the Advancement of Artificial Intelligence, AAAI 2023*

14. **Kai Wang**, Zhao Song, Georgios Theochar, and Sridhar Mahadevan
“Smoothed Online Combinatorial Optimization Using Imperfect Predictions”, *Association for the Advancement of Artificial Intelligence, AAAI 2023*
13. Sanket Shah, **Kai Wang**, Bryan Wilder, Andrew Perrault, Milind Tambe “Decision-Focused Learning without Decision-Making: Learning Locally Optimized Decision Losses”, *Conference on Neural Information Processing Systems, NeurIPS 2022*
12. **Kai Wang**, Lily Xu, Andrew Perrault, Michael K. Reiter, and Milind Tambe
“Coordinating Followers to Reach Better Equilibria: End-to-End Gradient Descent for Stackelberg Games”, *Association for the Advancement of Artificial Intelligence, AAAI 2022*
11. **Kai Wang**, Sanket Shah, Haipeng Chen, Andrew Perrault, Finale Doshi-Velez, and Milind Tambe
“Learning MDPs from Features: Predict-Then-Optimize for Sequential Decision Problems by Reinforcement Learning”, *Conference on Neural Information Processing Systems, NeurIPS 2021 (spotlight)*
10. Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, **Kai Wang**, and Milind Tambe
“Dual-Mandate Patrols: Multi-Armed Bandits for Green Security”, *Association for the Advancement of Artificial Intelligence, AAAI 2021 (best paper runner up)*
9. **Kai Wang**, Bryan Wilder, Andrew Perrault, and Milind Tambe
“Automatically Learning Compact Quality-aware Surrogates for Optimization Problems”, in *Conference on Neural Information Processing Systems, NeurIPS 2020 (spotlight)*.
8. Ayan Mukhopadhyay, **Kai Wang**, Andrew Perrault, Mykel Kochenderfer, Milind Tambe, and Yevgeniy Vorobeychik
“Robust Spatial-Temporal Incident Prediction”, in *Conference on Uncertainty in Artificial Intelligence, UAI 2020*.
7. **Kai Wang**, Andrew Perrault, Aditya Mate, and Milind Tambe.
“Scalable Game-Focused Learning of Adversary Models: Data-to-Decisions in Network Security Games”, in *International Conference On Autonomous Agents and Multi-Agent Systems, AAMAS 2020*.
6. Nitin Kamra, Umang Gupta, **Kai Wang**, Fei Fang, Yan Liu, Milind Tambe.
“DeepFP for Finding Approximate Nash Equilibrium in Continuous Action Spaces”, *Conference on Decision and Game Theory for Security, GameSec 2019*
5. Sarah Cooney, **Kai Wang**, Elizabeth Bondi, Thanh Nguyen, Phebe Vayanos, Hailey Winetrobe, Edward Cranford, Cleotilde Gonzalez, Christian Lebiere, and Milind Tambe.
“Learning to Signal in the Goldilocks Zone: Improving Adversary Compliance in Security Games”, *European Conference on Machine Learning, ECML 2019*.
4. Nitin Kamra, Umang Gupta, **Kai Wang**, Fei Fang, Yan Liu, Milind Tambe.
“Deep Fictitious Play for Games with Continuous Action Spaces”, extended abstract in *International Conference On Autonomous Agents and Multi-Agent Systems, AAMAS 2019*.
3. Sara Marie Mc Carthy, Corine Laan, **Kai Wang**, Phebe Vayanos, Milind Tambe, and Arunesh Sinha.
“The Price of Usability: Designing Operationalizable Strategies for Security Games”, in *International Joint Conference on Artificial Intelligence, IJCAI 2018*.
2. **Kai Wang**, Qingyu Guo, Phebe Vayanos, Milind Tambe, and Bo An.
“Equilibrium Refinement in Security Games with Arbitrary Scheduling Constraints”, in *International Conference On Autonomous Agents and Multi-Agent Systems, AAMAS 2018*.
1. Haifeng Xu, **Kai Wang**, Phebe Vayanos, and Milind Tambe.
“Strategic Coordination of Human Patrollers and Mobile Sensors with Signaling for Security Games”, in *Association for the Advancement of Artificial Intelligence, AAAI 2018*.

Workshop Publications

7. Shresth Verma, Aditya Mate, **Kai Wang**, Aparna Taneja, and Milind Tambe “Case Study: Applying Decision Focused Learning in the Real World”, *NeurIPS 2022 Workshop on Trustworthy and Socially Responsible Machine Learning (TSRML)*
6. **Kai Wang**, Jeffrey Brantingham, and Milind Tambe.
“Learning Opportunistic Adversarial Model on Global Wildlife Trade”, *AAMAS 2021 Autonomous Agents for Social Good Workshop (AASG)*
5. **Kai Wang**.
“Balance Between Scalability and Optimality in Network Security Games”, *AAMAS 2020 Doctoral Consortium*
4. Han Ching Ou, **Kai Wang**, Finale Doshi-Velez, Milind Tambe “Active Screening on Recurrent Diseases Contact Networks with Uncertainty: a Reinforcement Learning Approach”, *AAMAS 2020 International Workshop on Multi-Agent Systems and Agent-Based Simulation*
3. **Kai Wang**, Bryan Wilder, and Milind Tambe.
“Adversarial Machine Learning with Double Oracle”, *IJCAI 2019 Doctoral Consortium*
2. **Kai Wang**, Bryan Wilder, Sze-chuan Suen, Milind Tambe, and Bistra Dilkina.
“Improving GP-UCB Algorithm by Harnessing Decomposed Feedback”, *ECML 2019 SoGood Workshop*
1. **Kai Wang**, Hong-Jyun Wang, and Ho-Lin Chen.
“Routing Games with Priorities”, in *Asian Association for Algorithms and Computation, AAAC 2016 (Oral Presentation)*

Industry Experience

Google Research India Research Intern (mentored by Aparna Taneja)	Dec 2021 - May 2022 <i>Remote</i>
Adobe Research Research Intern (mentored by Georgios Theodorou, Sridhar Mahadevan)	May 2021 - Dec 2021 <i>Remote</i>
BravoAI Cofounder	Mar 2017 - Jul 2017 <i>Taipei, Taiwan</i>
Mixerbox Incorporation Software Engineer Intern	Apr 2016 - Mar 2017 <i>Taipei, Taiwan</i>

Patent Application

Active Control System for Data System Allocation US Patent Application No. 17/671,075 (Under Review)	Feb 2022
--	----------

Professional Service

IJCAI 2022 Tutorial Presenter Differentiable Optimization: Integrating Structural Information into Training Pipeline	July 2022
AAMAS 2022 Workshop Organizer Autonomous Agents for Social Good Workshop	May 2022
Organizer and Moderator Panel Discussion on Career Journey in Artificial Intelligence and Biotechnology	May 2022
Organizer and Moderator The Future of AI for Social Impact Panel at Harvard Center for Research on Computation and Society (CRCS)	April 2022

Organizer and Moderator

Seminar Series at Harvard Center for Research on Computation and Society (CRCS)

Jan - May 2022

Conference Reviewer AAAI 2023, 2022, 2021, 2020; NeurIPS 2022, 2021; ICML 2022, 2021; IJCAI 2021; JMLR 2022; AAMAS 2020; EC 2021; EAAMO 2022

Workshop Reviewer IJCAI AI4SG 2021; AAMAS AASG 2021, 2020, 2019; IJCAI AIWC 2018

Teaching Experience

Teaching Fellow

Harvard CS 127/227 Cryptography

Jan - May 2020

Taiwan National Team Coach

The 56th International Mathermatic Olympiad

April - July 2015

Lecturer of Competition and College Math for K12 Students

National Tsing Hua University College of Education

2014 - 2016

President and Lecturer

IMO Summer Camp

Aug 2013, Aug 2015

Invited Talks

University of Maryland

Oct 2022

Harvard Theory TGINF Seminar

Oct 2022

INFORMS General Session: Decision Analysis for Security and Defense: Models and Analytics

Oct 2022

Google Research India

Aug 2022

AI Research Student Conference at Nanyang Technology University

May 2022

Singapore Management University

Feb 2022

The Pennsylvania State University

Feb 2022

Google Research India

Feb 2022

Adobe Research

Mar 2021

Harvard EconCS Seminar

Oct 2020

Guest Lecture at Harvard CS182 Artificial Intelligence

Sep 2020

Leadership Experience

Vice President

Harvard Taiwanese Student Association

2022

Research Mentor

I am fortunate to work with the following students:

- London Lowmanstone '21 (mentored from 2020 - 2021): currently PhD student at the University of Minnesota
- Preston Yi Jie Ching '21 (mentored from 2020 - 2021): currently Assistant Head - National Research Foundation Singapore (NRF)
- Caroline Abel '23 (mentored from 2021 - 2022)

2020 - 2022

Captain

National Taiwan University Volleyball Varsity Team

Jun 2014 - Jun 2015