

Tonghan Wang

Science and Engineering Complex
Harvard School of Engineering and Applied Sciences
150 Western Avenue
Cambridge, MA 02138
twang1@g.harvard.edu
Google Scholar Homepage

Education

Harvard University	2025 (Expected)
PhD in Computer Science <i>Advisors:</i> David C. Parkes, Milind Tambe; <i>Thesis Committee:</i> Peter Stone, Kianté Brantley	
Tsinghua University	2021
M.E. in Computer Science, Institute of Interdisciplinary Information Sciences (IIIS) National Scholarship (top 1%) <i>Advisor:</i> Chongjie Zhang	
Shandong University	2018
B.E. in Computer Science, Special Honors Program, Taishan Academy National Scholarship; Principal Scholarship (top 0.05%)	

Honors and Awards

AAAI Oral	2025
Oral presentation, AAAI Conference on Artificial Intelligence, 2025	
EC Best Paper Award	2024
🏆 AI Track, ACM Conference on Economics and Computation, 2024	
NeurIPS Spotlight	2022
Spotlight paper, Conference on Neural Information Processing Systems, 2022	
ICLR Spotlight	2022
Spotlight paper, International Conference on Learning Representations, 2022	
ICLR Outstanding Reviewer Award	2021
International Conference on Learning Representations, 2021	
National Scholarship at Tsinghua University	2020
Awarded to top 1% students at IIIS, Tsinghua University	
NeurIPS Spotlight	2020
Spotlight paper, Conference on Neural Information Processing Systems, 2020	
ICLR Spotlight	2020
Spotlight paper, International Conference on Learning Representations, 2020	
Outstanding Undergraduates at Shandong University	2017
Awarded to top 10 (0.05%) students at Shandong University	

National Scholarship at Shandong University Awarded to top 1% students at Shandong University	2016
Principal Scholarship Awarded to top 18 of 40000 (0.05%) students at Shandong University	2016
First Class Scholarship Top 1 student at Taishan Academy, a special honors program at Shandong University	2015 – 2017

Publications

Peer-Reviewed Conference Papers

- [C24] **Tonghan Wang***, Heng Dong*, Yanchen Jiang, David C. Parkes, Milind Tambe. On Diffusion Models for MultiAgent Partial Observability: Shared Attractors, Error Bounds, and Composite Flow. *International Conference on Autonomous Agents and Multiagent Systems (AA-MAS)*. 2025. **Oral**.
- [C23] Yunfan Zhao*, **Tonghan Wang***, Dheeraj Mysore Nagaraj, Aparna Taneja, Milind Tambe. The Bandit Whisperer: Communication Learning for Restless Bandits. *AAAI Conference on Artificial Intelligence (AAAI)*. 2025. **Oral**.
- [C22] **Tonghan Wang***, Yanchen Jiang*, David C. Parkes. GemNet: Menu-Based, Strategy-Proof Multi-Bidder Auctions Through Deep Learning. *ACM Conference on Economics and Computation (EC)*. 2024. 🏆 **Best Paper, AI Track**.
- [C21] Safwan Hossian*, **Tonghan Wang***, Tao Lin*, Yiling Chen, David C. Parkes, Haifeng Xu. Multi-Sender Persuasion: A Computational Perspective. *International Conference on Machine Learning (ICML)*. 2024.
- [C20] Edwin Zhang, Sadie Zhao, **Tonghan Wang**, Safwan Hossain, Henry Gasztowtt, Stephan Zheng, David C. Parkes, Milind Tambe, Yiling Chen. Position: Social Environment Design. *International Conference on Machine Learning (ICML)*. 2024.
- [C19] **Tonghan Wang**, Paul Dütting, Dmitry Ivanov, Inbal Talgam-Cohen, David C. Parkes. Deep Contract Design via Discontinuous Piecewise-Affine Neural Networks. *Conference on Neural Information Processing Systems (NeurIPS)*. 2023.
- [C18] Heng Dong, Junyu Zhang, **Tonghan Wang**, Chongjie Zhang. Symmetry-Aware Robot Design with Structured Subgroups. *International Conference on Machine Learning (ICML)*. 2023.
- [C17] Heng Dong*, **Tonghan Wang***, Jiayuan Liu, Chongjie Zhang. Low-Rank Modular Reinforcement Learning via Muscle Synergy. *Conference on Neural Information Processing Systems (NeurIPS)*. 2022.
- [C16] Yipeng Kang*, **Tonghan Wang***, Qianlan Yang, Xiaoran Wu, Chongjie Zhang. Non-Linear Coordination Graphs. *Conference on Neural Information Processing Systems (NeurIPS)*. 2022. **Spotlight**.
- [C15] Weijun Dong, Qianlan Yang, Zhizhou Ren, Jianhao Wang, **Tonghan Wang**, Chongjie Zhang. Self-Organized Polynomial-time Coordination Graphs. *International Conference on Machine Learning (ICML)*. 2022.

- [C14] **Tonghan Wang***, Liang Zeng*, Weijuan Dong, Qianlan Yang, and Chongjie Zhang. Context-Aware Sparse Deep Coordination Graphs. *International Conference on Learning Representations (ICLR)*. 2022. **Spotlight**.
- [C13] Chenghao Li, **Tonghan Wang**, Chengjie Wu, Yiqin Yang, Qianchuan Zhao, Chongjie Zhang. Celebrating Diversity in Shared Multi-Agent Reinforcement Learning. *Conference on Neural Information Processing Systems (NeurIPS)*. 2021.
- [C12] **Tonghan Wang**, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang. RODE: Learning Roles to Decompose Multi-Agent Tasks. *International Conference on Learning Representations (ICLR)*. 2021.
- [C11] Yihan Wang*, Beining Han*, **Tonghan Wang***, Heng Dong, and Chongjie Zhang. DOP: Off-Policy Multi-Agent Decomposed Policy Gradients. *International Conference on Learning Representations (ICLR)*. 2021.
- [C10] Yipeng Kang, **Tonghan Wang**, Gerard de Melo. Incorporating Pragmatic Reasoning Communication into Emergent Language. *Conference on Neural Information Processing Systems (NeurIPS)*. 2020. **Spotlight**.
- [C9] **Tonghan Wang**, Heng Dong, Victor Lesser, and Chongjie Zhang. ROMA: Multi-Agent Reinforcement Learning with Emergent Roles. *International Conference on Machine Learning (ICML)*. 2020.
- [C8] **Tonghan Wang***, Jianhao Wang*, Yi Wu, and Chongjie Zhang. Influence-Based Multi-Agent Exploration. *International Conference on Learning Representations (ICLR)*. 2020. **Spotlight**.
- [C7] **Tonghan Wang***, Jianhao Wang*, Chongyi Zheng, and Chongjie Zhang. Learning Nearly Decomposable Value Functions with Communication Minimization. *International Conference on Learning Representations (ICLR)*. 2020.
- [C6] Xinliang Song, **Tonghan Wang**, and Chongjie Zhang. Convergence of Multi-Agent Learning with a Finite Step Size in General-Sum Games. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2019.
- [C5] **Tonghan Wang**, Xueying Qin, Fan Zhong, Baoquan Chen, and Ming C. Lin. Compact Object Representation of a Non-Rigid Object for Real-Time Tracking in AR Systems. *International Symposium on Mixed and Augmented Reality (ISMAR)*. 2018.
- [C4] Lingkai Kong, Haichuan Wang, Yuqi Pan, Cheol Woo Kim, Mingxiao Song, Alayna Nguyen, **Tonghan Wang**, Haifeng Xu, Milind Tambe. Robust Optimization with Diffusion Models for Green Security. *Conference on Uncertainty in Artificial Intelligence (UAI)*. 2025. To appear.

Journal Papers

- [C3] Michael J. Curry, Zhou Fan, Yanchen Jiang, Sai Srivatsa Ravindranath, **Tonghan Wang**, David C. Parkes. Automated Mechanism Design: A Survey. *ACM SIGecom Exchanges*. 2025.
- [C2] Rongjun Qin*, Feng Chen*, **Tonghan Wang***, Lei Yuan, Xiaoran Wu, Zongzhang Zhang, Chongjie Zhang, Yang Yu. Multi-Agent Policy Transfer via Task Relationship Modeling. *Science China Information Sciences (SCIS)*. 2024.

- [C1] Chenghao Li, **Tonghan Wang**, Chengjie Wu, Qianchuan Zhao, Jun Yang, Chongjie Zhang. Celebrating Diversity With Subtask Specialization in Shared Multiagent Reinforcement Learning. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*. 2023.

Preprints and Peer-Reviewed Workshop Papers

- [P5] Siyang Wu*, **Tonghan Wang***, Xiaoran Wu, Jingfeng Zhang, Yujing Hu, Changjie Fan, Chongjie Zhang. Model and Method: Training-Time Attack for Cooperative Multi-Agent Reinforcement Learning. *Deep Reinforcement Learning Workshop at NeurIPS*. 2022.
- [P4] Rongjun Qin*, Feng Chen*, **Tonghan Wang***, Lei Yuan, Xiaoran Wu, Zongzhang Zhang, Chongjie Zhang, Yang Yu. Multi-Agent Policy Transfer via Task Relationship Modeling. *Deep Reinforcement Learning Workshop at NeurIPS*. 2022.
- [P3] **Tonghan Wang**, Yanchen Jiang, David C. Parkes. BundleFlow: Deep Menus for Combinatorial Auctions by Diffusion-Based Optimization. 2025.
- [P2] Xinyi Yang, Liang Zeng, Heng Dong, Chao Yu, Xiaoran Wu, Huazhong Yang, Yu Wang, Milind Tambe, **Tonghan Wang**. Policy-to-Language: Train LLMs to Explain Decisions with Flow-Matching Generated Rewards. 2025.
- [P1] Lingkai Kong*, Haichuan Wang*, **Tonghan Wang***, Guojun Xiong, Milind Tambe. Composite Flow Matching for Reinforcement Learning with Shifted-Dynamics Data. 2025.

Teaching Assistant

AM 220: Geometric Methods for Machine Learning, Harvard University	Spring 2024
Deep Reinforcement Learning, Tsinghua University	Spring 2020
Artificial Intelligence: Principles and Techniques, Tsinghua University	Fall 2020
Artificial Intelligence: Principles and Techniques, Tsinghua University	Fall 2019
Linear Algebra, Shandong University	Fall 2016

Professional Services

Conference Reviewer

International Conference on Learning Representations (ICLR)	2021 –
Conference on Neural Information Processing Systems (NeurIPS)	2021 –
International Conference on Machine Learning (ICML)	2021 –
AAAI Conference on Artificial Intelligence (AAAI)	2020 –
International Joint Conference on Artificial Intelligence (IJCAI)	2020 –

Journal Reviewer

Autonomous Agents and Multi-Agent Systems	2025
Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2024, 2025