

Yi Wu

Email: jxwuyi@gmail.com

Homepage: <http://jxwuyi.weebly.com/>

Google Scholar: <http://scholar.google.com/citations?user=dusV5HMAAAAI>

EMPLOYMENT

Assistant Professor, Institute of Interdisciplinary Information Sciences (IIIS), Tsinghua University, Beijing, China 09/2020 – present

- Research interests include: Deep Reinforcement Learning, Human-AI Interaction, Robot Learning, Language Model Grounding and Alignment

Visiting Researcher, OpenAI Inc., San Francisco, USA 03/2019 – 07/2019

EDUCATION

Ph.D. in Artificial Intelligence (AI), University of California, Berkeley, USA 08/2014 – 09/2019

- Advised by Prof. Stuart Russell

B.E. in Computer Science, Special CS Pilot Class, Institute of Interdisciplinary Information Sciences (IIIS), Tsinghua University, Beijing, China 08/2010 – 07/2014

PUBLICATIONS

1. Weihua Du*, Jinglun Zhao*, Chao Yu, Xingcheng Yao, Zimeng Song, Siyang Wu, Ruifeng Luo, Zhiyuan Liu, Xianzhong Zhao, **Yi Wu**, Automatics Truss Design with Reinforcement Learning, *International Joint Conference on Artificial Intelligence (IJCAI) 2023*
2. Chao Yu*, Jiaxuan Gao*, Weilin Liu, Botian Xu, Hao Tang, Jiaqi Yang, Yu Wang, **Yi Wu**, Learning Zero-Shot Cooperation with Humans, Assuming Humans Are Biased, *International Conference on Learning Representations (ICLR) 2023*
3. Yixuan Mei*, Jiaxuan Gao*, Weirui Ye, Shaohuai Liu, Yang Gao, **Yi Wu**, SpeedyZero: Mastering Atari with Limited Data and Time, *International Conference on Learning Representations (ICLR) 2023*
4. Yunfei Li*, Chaoyi Pan*, Huazhe Xu, Xiaolong Wang, **Yi Wu**, Efficient Bimanual Handover and Rearrangement via Symmetry-Aware Actor-Critic Learning, *IEEE International Conference on Robotics and Automation (ICRA) 2023*
5. Zelai Xu, Yancheng Liang, Chao Yu, Yu Wang, **Yi Wu**, Fictitious Cross-Play: Learning Nash Equilibrium in Mixed Cooperative-Competitive Games, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2023*
6. Chao Yu*, Xinyi Yang*, Jiaxuan Gao*, Jiayu Chen, Yunfei Li, Jijia Liu, Yunfei Xiang, Ruixin Huang, Huazhong Yang, **Yi Wu**, Yu Wang, Asynchronous Multi-Agent Reinforcement Learning for Efficient Real-Time Multi-Robot Cooperative Exploration, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2023*
7. Jing Wang*, Meichen Song*, Feng Gao*, Boyi Liu, Zhaoran Wang, **Yi Wu**, Differentiable Arbitrating in Zero-sum Markov Games, *International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2023*
8. Shusheng Xu, Yancheng Liang, Yunfei Li, Simon Shaolei Du, **Yi Wu**, Beyond Information Gain: An

-
- Empirical Benchmark for Low-Switching-Cost Reinforcement Learning, *Transactions on Machine Learning Research (TMLR)*, ISSN 2835-8856, 2023
9. Rui Zhao, Jinming Song, Yufeng Yuan, Hu Haifeng, Yang Gao, **Yi Wu**, Zhongqian Sun, Yang Wei, Maximum Entropy Population-Based Training for Zero-Shot Human-AI Coordination, *The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI) 2023*
 10. Chao Yu*, Akash Velu*, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, **Yi Wu**, The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games, *Advances in Neural Information Processing Systems (NeurIPS), Benchmark Track, 2022, (citation 508)*
 11. Shusheng Xu, Huaijie Wang, **Yi Wu**, Grounded Reinforcement Learning: Learning to Win the Game under Human Commands, *Advances in Neural Information Processing Systems (NeurIPS) 2022*
 12. Zhecheng Yuan*, Zhengrong Xue*, Bo Yuan, Xueqian Wang, **Yi Wu**, Yang Gao, Huazhe Xu, Pre-Trained Image Encoder for Generalizable Visual Reinforcement Learning, *Advances in Neural Information Processing Systems (NeurIPS) 2022*
 13. Chao Yu*, Xinyi Yang*, Jiaxuan Gao*, Huazhong Yang, Yu Wang, **Yi Wu**, Learning Efficient Multi-Agent Cooperative Visual Exploration, *European Conference on Computer Vision (ECCV) 2022*
 14. Yunfei Li*, Tian Gao*, Jiaqi Yang, Huazhe Xu, **Yi Wu**, Phasic Self-Imitative Reduction for Sparse-Reward Goal-Conditioned Reinforcement Learning, *International Conference on Machine Learning (ICML) 2022*
 15. Wei Fu, Chao Yu, Zelai Xu, Jiaqi Yang, **Yi Wu**, Revisiting Some Common Practices in Cooperative Multi-Agent Reinforcement Learning, *International Conference on Machine Learning (ICML) 2022*
 16. Zihan Zhou*, Wei Fu*, Bingliang Zhang, **Yi Wu**, Continuously Discovering Novel Strategies via Reward-Switching Policy Optimization, *International Conference on Learning Representations (ICLR) 2022*
 17. Yunfei Li, Tao Kong, Lei Li, **Yi Wu**, Learning Design and Construction with Varying-Sized Materials via Prioritized Memory Resets, *IEEE International Conference on Robotics and Automation (ICRA) 2022*
 18. Shusheng Xu, Xingxing Zhang, **Yi Wu**, Furu Wei, Sequence Level Contrastive Learning for Text Summarization, *The Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI) 2022*
 19. Jiayu Chen, Yuanxin Zhang, Yuanfan Xu, Huimin Ma, Huazhong Yang, Jiaming Song, Yu Wang, **Yi Wu**, Variational Automatic Curriculum Learning for Sparse-Reward Cooperative Multi-Agent Problems, *Advances in Neural Information Processing Systems (NeurIPS) 2021*
 20. Tianjun Zhang, Huazhe Xu, Xiaolong Wang, **Yi Wu**, Kurt Keutzer, Joseph E. Gonzalez, Yuandong Tian, NovelD: A Simple yet Effective Exploration Criterion, *Advances in Neural Information Processing Systems (NeurIPS) 2021*
 21. Shusheng Xu*, Yichen Liu*, Xiaoyu Yi, Siyuan Zhou, Huizi Li, **Yi Wu**, Native Chinese Reader: A Dataset Towards Native-Level Chinese Machine Reading Comprehension, *Advances in Neural Information Processing Systems (NeurIPS), Benchmark Track, 2021*
 22. Yunfei Li, Tao Kong, Lei Li, Yifeng Li, **Yi Wu**, Learning to Design and Construct Bridge without Blueprint, *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2021*
 23. Weizhe Chen*, Zihan Zhou*, **Yi Wu**, Fei Fang, Temporal Induced Self-Play for Stochastic Bayesian Games, *International Joint Conference on Artificial Intelligence (IJCAI) 2021*
 24. Zhenggang Tang*, Chao Yu*, Boyuan Chen, Huazhe Xu, Xiaolong Wang, Fei Fang, Simon Shaolei

-
- Du, Yu Wang, **Yi Wu**, Discovering Diverse Multi-Agent Strategic Behavior via Reward Randomization, *International Conference on Learning Representations (ICLR) 2021*
25. Yunfei Li, Yilin Wu, Huazhe Xu, Xiaolong Wang, **Yi Wu**, Solving Compositional Reinforcement Learning Problems via Task Reduction, *International Conference on Learning Representations (ICLR) 2021*
26. Yining Wang, **Yi Wu**, Simon Shaolei Du, Near-Linear Time Local Polynomial Nonparametric Estimation with Box Kernels, *INFORMS Journal on Computing* 33(4):1339-1353, 2021
27. Ruihan Yang, Huazhe Xu, **Yi Wu**, Xiaolong Wang, Multi-Task Reinforcement Learning with Soft Modularization, *Advances in Neural Information Processing Systems (NeurIPS) 2020*
28. Shusheng Xu, Xingxing Zhang, **Yi Wu**, Furu Wei, Ming Zhou, Unsupervised Extractive Summarization by Pre-training Hierarchical Transformers, *Conference on Empirical Methods in Natural Language Processing (EMNLP), Findings of EMNLP 2020*
29. Bowen Baker*, Ingmar Kanitscheider*, Todor Markov*, **Yi Wu***, Glenn Powell*, Bob McGrew*, Igor Mordatch*, Emergent Tool Use From Multi-Agent Autocurricula, *International Conference on Learning Representations (ICLR) 2020* (* team project), (**spotlight**, [citation 637](#))
30. Tonghan Wang*, Jianhao Wang*, **Yi Wu**, Chongjie Zhang, Influence-Based Multi-Agent Exploration, *International Conference on Learning Representations (ICLR) 2020*, (**spotlight**)
31. Qian Long*, Zihan Zhou*, Abhinav Gupta, Fei Fang, **Yi Wu†**, Xiaolong Wang†, Evolutionary Population Curriculum for Scaling Multi-Agent Reinforcement Learning, *International Conference on Learning Representations (ICLR) 2020* († equal advising)
32. **Yi Wu**, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, Yuandong Tian, Bayesian Relational Memory for Semantic Visual Navigation, *IEEE/CVF International Conference on Computer Vision (ICCV) 2019*
33. Shihui Li, **Yi Wu**, Xinyue Cui, Honghua Dong, Fei Fang, Stuart Russell, Robust Multi-Agent Reinforcement Learning via Minimax Deep Deterministic Policy Gradient, *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI) 2019*, (**Oral**)
34. Yufei Wang, Zheyuan Ryan Shi, Lantao Yu, **Yi Wu**, Rohit Singh, Lucas Joppa, Fei Fang, Deep Reinforcement Learning for Green Security Games with Real-Time Information, *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI) 2019*, (**Oral**)
35. Tongzhou Wang, **Yi Wu**, David Moore, Stuart Russell, Meta-Learning MCMC Proposals, *Advances in Neural Information Processing Systems (NIPS) 2018*
36. **Yi Wu**, Siddharth Srivastava, Nicholas Hay, Simon Du, Stuart Russell, Discrete-Continuous Mixtures in Probabilistic Programming: Generalized Semantics and Inference Algorithms, *International Conference on Machine Learning (ICML) 2018*
37. **Yi Wu**, Yuxin Wu, Georgia Gkioxari, Yuandong Tian, Building Generalizable Agents with a Realistic and Rich 3D Environment, *International Conference on Learning Representations (ICLR), Workshop Track, 2018*
38. Ryan Lowe*, **Yi Wu***, Aviv Tamar, Jean Harb, Pieter Abbeel, Igor Mordatch, Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments, *Advances in Neural Information Processing Systems (NIPS) 2017* (* equal contribution) ([citation 3700](#))
39. **Yi Wu**, David Bamman, Stuart Russell, Adversarial Training for Relation Extraction, *Conference on*

40. Yusuf B. Erol*, **Yi Wu***, Lei Li, Stuart Russell, A Nearly-Black-Box Online Algorithm for Joint Parameter and State Estimation in Temporal Models, *The Thirty-First AAAI Conference on Artificial Intelligence (AAAI) 2017* (* equal contribution), (**Oral**)
41. Aviv Tamar, **Yi Wu**, Garrett Thomas, Sergey Levine, Pieter Abbeel, Value Iteration Network, *Advances in Neural Information Processing Systems (NIPS) 2016*, (**Best Paper Award**)
42. **Yi Wu**, Lei Li, Stuart Russell, Rastislav Bodik, Swift: Compiled Inference for Probabilistic Programming Languages, *International Joint Conference on Artificial Intelligence (IJCAI) 2016*
43. **Yi Wu**, David Wipf, Jeong-Min Yun, Understanding and Evaluating Sparse Linear Discriminant Analysis, *International Conference on Artificial Intelligence and Statistics (AISTATS) 2015*, (**Oral**)
44. David Wipf, **Yi Wu**, Dual Space Analysis on the Sparse Linear Model, *Advances in Neural Information Processing Systems (NIPS) 2012*

SELECTED WORKSHOP PAPERS

1. Zhiyu Mei*, Wei Fu*, Guangju Wang, Huanchen Zhang, **Yi Wu**, SRL: Scaling Distributed Reinforcement Learning to Over Ten Thousand Cores, *International Conference on Machine Learning (ICML) 2023, Efficient Systems for Foundation Models Workshop* (**Oral, top 5**)
2. Yingying Wu, Shusheng Xu, Shing-Tung Yau, **Yi Wu**, PhyloTransformer: A Self-supervised Discriminative Model for SARS-CoV-2 Viral Mutation Prediction Based on a Multi-head Self-attention Mechanism, *International Conference on Learning Representations (ICLR) 2023, Workshop on Machine Learning and Global Health*
3. Runlong Zhou, Yuandong Tian, **Yi Wu**, Simon Shaolei Du, Understanding Curriculum Learning in Policy Optimization for Online Combinatorial Optimization, *Advances in Neural Information Processing Systems (NeurIPS) 2020, Optimization for Machine Learning Workshop*
4. Zihan Zhou, Zheyuan Ryan Shi, **Yi Wu**, Fei Fang. Approximated Temporal-Induced Neural Self-Play for Finitely Repeated Bayesian Games. *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI) 2020, Reinforcement Learning in Games Workshop* (**Oral**)
5. **Yi Wu**, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, Yuandong Tian, Learning and Planning with a Semantic Model, *European Conference on Computer Vision (ECCV) 2018, Workshop on Visual Learning and Embodied Agents in Simulation Environments*
6. Brian Milch, **Yi Wu**, Lei Li, Stuart Russell, BLOG: a probabilistic programming language for open-universe contingent Bayesian networks, *International Conference on Machine Learning (ICML) 2015, 5th Workshop on Machine Learning Open Source Software* (**Oral**)
7. **Yi Wu**, Lei Li, Stuart Russell, BFiT: From Possible-World Semantics to Random-Evaluation Semantics in Open Universe, *Advances in Neural Processing Systems (NIPS) 2014, 3rd Workshop on Probabilistic Programming* (**Spotlight**)

AWARDS AND HONORS

Academic Awards

- NIPS 2016, Best Paper Award 12/2016

Algorithm and Programming

- North America Champion, Silver Medal (ranked 6th in the world) in ACM/ICPC 05/2015

World Finals 2015, Marrakech, Morocco

- Bronze Medal (ranked 12th in the world) in **ACM/ICPC World Finals 2013**, Saint Petersburg, Russia 06/2013
- Two Champions and Four Second Places in ACM/ICPC Regional Contests 2010 – 2014
- Silver Medal (ranked 6th) in 22nd **International Olympiad in Informatics (IOI2010)**, Waterloo, Canada 08/2010

Scholarship

- Baidu Scholarship 2016
– Awarded to 10 PhD students all over the world
- Department Fellowship, EECS, UC Berkeley 2014
- Excellent Graduate Student in Tsinghua University 2014
– Awarded to top 10% of all the undergraduate students in Tsinghua
- Principal Scholarship in Institute of Interdisciplinary Information Sciences 2011, 2013
– Awarded to only 3 students in 2011 and only 6 students in 2013
- Tsinghua University Principal Scholarship for Freshman 2010
– Awarded to only 53 of 3300 freshmen in the whole Tsinghua University

PROFESSIONAL EXPERIENCES

- Engineer Intern, Recommendation Team, Bytedance Inc., Beijing, China** 05/2018 – 08/2018
 - Mentored by Mu Qiao (team manager)
- Research Intern, Facebook AI Research (FAIR), Menlo Park, CA, USA** 05/2017 – 05/2018
 - Mentored by Yuandong Tian and Georgia Gkioxari
- Research Intern, Toutiao AI Lab, Bytedance Inc., Beijing, China** 05/2016 – 08/2016
 - Mentored by Lei Li (AI Lab director)
- Research Intern, Visual Computing Group, Microsoft Research Asia (MSRA), Beijing, China** 10/2012 – 07/2013
11/2011 – 06/2012
 - Mentored by David Wipf (lead researcher)
- Software Engineering Intern, Search Team, Facebook Inc. Headquarters, Menlo Park, CA, USA** 06/2012 – 09/2012
 - Mentored by Gintaras Woss