

# Zongzhang Zhang

163 Xianlin Ave  
Nanjing, Jiangsu, China, 210023

<https://www.lamda.nju.edu.cn/zhangzz/>  
zzzhang@nju.edu.cn

## Employment

---

<b>Associate Professor</b> <i>School of Artificial Intelligence, Nanjing University</i>	Jul, 2019 – Present Nanjing, Jiangsu, China
<b>Associate Professor</b> <i>School of Computer Science and Technology, Soochow University</i>	Jul, 2014 – Jun, 2019 Suzhou, Jiangsu, China
<b>Research Fellow</b> <i>Department of Computer Science, National University of Singapore</i>	Nov, 2012 – Jun, 2014 Singapore, Singapore
<b>Research Engineer</b> <i>Noah's Ark Lab, Huawei Company</i>	Aug, 2012 – Oct, 2012 Shenzhen, Guangdong, China

## Education

---

<b>University of Science and Technology of China</b> <i>Ph.D. in Computer Science</i>	Sep, 2007 – Jun, 2012 Hefei, Jiangsu, China
<b>Central South University</b> <i>B.S. in Mathematics</i>	Sep, 2003 – Jun, 2007 Changsha, Hunan, China

## Research Experience

---

<b>Visiting Scholar</b> <i>Department of Aeronautics and Astronautics, Stanford University</i>	Sep, 2018 – Mar, 2019 Stanford, California, USA
<b>Research Visiting Student</b> <i>Department of Computer Science, Rutgers University</i>	Oct, 2010 – Oct, 2011 Piscataway, New Jersey, USA

## Research Interests

---

My research interests mainly include artificial intelligence and machine learning. Now I am working on

**Reinforcement Learning (RL)**, including deep RL, transfer RL, data-driven RL, and visual RL  
**Multi-agent systems**, e.g., multi-agent learning, multi-agent communication, and multi-agent coordination  
**Probabilistic planning**, particularly in partially observable Markov decision processes  
**Imitation learning** based on generative adversarial nets

## Conference Publications

---

- Xinyu Zhang, Wenjie Qiu, Yi-Chen Li, Lei Yuan, Chengxing Jia, Zongzhang Zhang, and Yang Yu. Debiased Offline Representation Learning for Fast Online Adaptation in Non-stationary Dynamics. In: *Proceedings of the 41st International Conference on Machine Learning (ICML-2024)*, Vienna, Austria, 2024. (To Appear)
- Xiong-Hui Chen, Junyin Ye, Hang Zhao, Yi-Chen Li, Xu-Hui Liu, Haoran Shi, Yu-Yan Xu, Zhihao Ye, Si-Hang Yang, Yang Yu, Anqi Huang, Kai Xu, and Zongzhang Zhang. Deep Demonstration Tracing: Learning Generalizable Imitator for Runtime One-Shot Imitation. In: *Proceedings of the 41st International Conference on Machine Learning (ICML-2024)*, Vienna, Austria, 2024. (To Appear)
- Rui Kong, Chenyang Wu, Chen-Xiao Gao, Zongzhang Zhang, and Ming Li. Efficient and Stable Offline-to-online Reinforcement Learning via Continual Policy Revitalization. In: *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI-2024)*, Jeju Island, South Korea, 2024. (To Appear)

4. Zican Hu, Zongzhang Zhang, Huaxiong Li, Chunlin Chen, Hongyu Ding, and Zhi Wang. Attention-Guided Contrastive Role Representations for Multi-agent Reinforcement Learning. In: *Proceedings of the 12th International Conference on Learning Representations (ICLR-2024)*, Vienna, Austria, 2024.
5. Chengxing Jia, Chenxiao Gao, Hao Yin, Fuxiang Zhang, Xiong-Hui Chen, Tian Xu, Lei Yuan, Zongzhang Zhang, Yang Yu, and Zhi-Hua Zhou. Policy Rehearsing: Training Generalizable Policies for Reinforcement Learning. In: *Proceedings of the 12th International Conference on Learning Representations (ICLR-2024)*, Vienna, Austria, 2024.
6. Jing-Cheng Pang, Pengyuan Wang, Kaiyuan Li, Xiong-Hui Chen, Jiacheng Xu, Zongzhang Zhang, and Yang Yu. Language Model Self-improvement by Reinforcement Learning Contemplation. In: *Proceedings of the 12th International Conference on Learning Representations (ICLR-2024)*, Vienna, Austria, 2024.
7. Feng Chen, Fuguang Han, Cong Guan, Lei Yuan, Zhilong Zhang, Yang Yu, and Zongzhang Zhang. Stable Continual Reinforcement Learning via Diffusion-based Trajectory Replay. In: *ICLR 2024 Workshop on Generative Models for Decision Making*, Vienna, Austria, 2024.
8. Cong Guan, Feng Chen, Ke Xue, Chunpeng Fan, Lichao Zhang, Ziqian Zhang, Pengyao Zhao, Zongzhang Zhang, Chao Qian, Lei Yuan, and Yang Yu. One by One, Continual Coordinating with Humans via Hyper-Teammate Identification. In: *ICLR 2024 Workshop on Generative Models for Decision Making*, Vienna, Austria, 2024.
9. Chao Chen, Dawei Wang, Feng Mao, Jiacheng Xu, Zongzhang Zhang, and Yang Yu. Deep Anomaly Detection via Active Anomaly Search. In: *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2024)*, pages 308–316, Auckland, New Zealand, 2024.
10. Chengxing Jia, Fuxiang Zhang, Yi-Chen Li, Chenxiao Gao, Xu-Hui Liu, Lei Yuan, Zongzhang Zhang, and Yang Yu. Disentangling Policy from Offline Task Representation Learning via Adversarial Data Augmentation. In: *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2024)*, pages 944–953, Auckland, New Zealand, 2024.
11. Chao Chen, Jiacheng Xu, Weijian Liao, Hao Ding, Zongzhang Zhang, Yang Yu, and Rui Zhao. Focus-Then-Decide: Segmentation-Assisted Reinforcement Learning. In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 11240–11248, Vancouver, Canada, 2024.
12. Chenxiao Gao, Chenyang Wu, Mingjun Cao, Rui Kong, Zongzhang Zhang, and Yang Yu. ACT: Empowering Decision Transformer with Dynamic Programming via Advantage Conditioning. In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 12127–12135, Vancouver, Canada, 2024.
13. Renzhe Zhou, Chenxiao Gao, Zongzhang Zhang, and Yang Yu. Generalizable Task Representation Learning for Offline Meta-Reinforcement Learning with Data Limitations. In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 17132–17140, Vancouver, Canada, 2024.
14. Yujian Zhu, Hao Ding, and Zongzhang Zhang. Multi-Expert Distillation for Few-Shot Coordination (Student Abstract). In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 23717–23719, Vancouver, Canada, 2024.
15. Rui Kong, Chenyang Wu, and Zongzhang Zhang. Generalizable Policy Improvement via Reinforcement Sampling (Student Abstract). In: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI-2024)*, pages 23546–23547, Vancouver, Canada, 2024.
16. Jiacheng Xu, Chao Chen, Fuxiang Zhang, Lei Yuan, Zongzhang Zhang, and Yang Yu. Internal Logical Induction for Pixel-Symbolic Reinforcement Learning. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD-2023)*, pages 2825–2837, Long Beach, CA, USA, 2023.
17. Yuhang Ran, Yi-Chen Li, Fuxiang Zhang, Zongzhang Zhang, and Yang Yu. Policy Regularization with Dataset Constraint for Offline Reinforcement Learning. In: *Proceedings of the 40th International Conference on Machine Learning (ICML-2023)*, pages 28701–28717, Honolulu, Hawaii, USA, 2023.

18. Guoqing Liu, Di Xue, Shufang Xie, Yingce Xia, Austin Tripp, Krzysztof Maziarczyk, Marwin Segler, Tao Qin, Zongzhang Zhang, and Tie-Yan Liu. Retrosynthetic Planning with Dual Value Networks. In: *Proceedings of the 40th International Conference on Machine Learning (ICML-2023)*, pages 22266-22276, Honolulu, Hawaii, USA, 2023.
19. Fuxiang Zhang, Chengxing Jia, Yi-Chen Li, Lei Yuan, Yang Yu, and Zongzhang Zhang. Discovering Generalizable Multi-agent Coordination Skills from Multi-task Offline Data. In: *Proceedings of the 11th International Conference on Learning Representations (ICLR-2023)*, Kigali, Rwanda, 2023.
20. Xuhui Liu, Feng Xu, Xinyu Zhang, Tianyuan Liu, Shengyi Jiang, Ruifeng Chen, Zongzhang Zhang, and Yang Yu. How To Guide Your Learner: Imitation Learning with Active Adaptive Expert Involvement. In: *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, pages 1276-1284, London, United Kingdom, 2023.
21. Weijian Liao, Zongzhang Zhang, and Yang Yu. Policy-Independent Behavioral Metric-Based Representation for Deep Reinforcement Learning. In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 8746-8754, Washington, DC, USA, 2023.
22. Aoran Wang, Hongyang Yang, Feng Mao, Zongzhang Zhang, Yang Yu, and Xiaoyang Liu. Anti-Drifting Feature Selection via Deep Reinforcement Learning (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16356-16357, Washington, DC, USA, 2023. **(Best Student Abstract Award – Honorable Mention)**
23. Chao Chen, Dawei Wang, Feng Mao, Zongzhang Zhang, and Yang Yu. Deep Anomaly Detection and Search via Reinforcement Learning (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16180-16181, Washington, DC, USA, 2023.
24. Feng Chen, Chenghe Wang, Fuxiang Zhang, Hao Ding, Qiaoyong Zhong, Shiliang Pu, and Zongzhang Zhang. Towards Deployment-Efficient and Collision-Free Multi-Agent Path Finding (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16182-16183, Washington, DC, USA, 2023.
25. Fuguang Han and Zongzhang Zhang. Expert Data Augmentation in Imitation Learning (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16220-16221, Washington, DC, USA, 2023.
26. Yi-Chen Li, Wen-Jie Shen, Boyu Zhang, Feng Mao, Zongzhang Zhang, and Yang Yu. Learning Generalizable Batch Active Learning Strategies via Deep Q-Networks (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16258-16259, Washington, DC, USA, 2023.
27. Renzhe Zhou, Zongzhang Zhang, and Yang Yu. Model-based Offline Weighted Policy Optimization (Student Abstract). In: *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 16392-16393, Washington, DC, USA, 2023.
28. Chenyang Wu, Tianci Li, Zongzhang Zhang, and Yang Yu. Bayesian Optimistic Optimization: Optimistic Exploration for Model-based Reinforcement Learning. In: *Advances in Neural Information Processing Systems 35 (NeurIPS-2022)*, pages 14210-14223, New Orleans, USA, 2022.
29. Cong Guan, Feng Chen, Lei Yuan, Chenghe Wang, Hao Yin, Zongzhang Zhang, and Yang Yu. Efficient Multi-agent Communication via Self-supervised Information Aggregation. In: *Advances in Neural Information Processing Systems 35 (NeurIPS-2022)*, pages 1020-1033, New Orleans, USA, 2022.
30. Ke Xue, Jiacheng Xu, Lei Yuan, Miqing Li, Chao Qian, Zongzhang Zhang, and Yang Yu. Multi-agent Dynamic Algorithm Configuration. In: *Advances in Neural Information Processing Systems 35 (NeurIPS-2022)*, pages 20147-20161, New Orleans, USA, 2022.

31. Rongjun Qin, Feng Chen, Tonghan Wang, Lei Yuan, Xiaoran Wu, Yipeng Kang, Zongzhang Zhang, Chongjie Zhang, and Yang Yu. Multi-Agent Policy Transfer via Task Relationship Modeling. In: *Deep RL Workshop on Neural Information Processing Systems*, New Orleans, USA, 2022.
32. Di Xue, Lei Yuan, Zongzhang Zhang, and Yang Yu. Efficient Multi-Agent Communication via Shapley Message Value. In: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 578-584, Vienna, Austria, 2022.
33. Lei Yuan, Chenghe Wang, Jianhao Wang, Fuxiang Zhang, Feng Chen, Cong Guan, Zongzhang Zhang, Chongjie Zhang, and Yang Yu. Multi-Agent Concentrative Coordination with Decentralized Task Representation. In: *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 599-605, Vienna, Austria, 2022.
34. Lei Yuan, Jianhao Wang, Fuxiang Zhang, Chenghe Wang, Zongzhang Zhang, Yang Yu, and Chongjie Zhang. Multi-Agent Incentive Communication via Decentralized Teammate Modeling. In: *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI-2022)*, pages 9466-9474, Virtual Conference, 2022.
35. Fan-Ming Luo, Shengyi Jiang, Yang Yu, Zongzhang Zhang, and Yi-Feng Zhang. Adapt to Environment Sudden Changes by Learning a Context Sensitive Policy. In: *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI-2022)*, pages 7637-7646, Virtual Conference, 2022.
36. Chenyang Wu, Guoyu Yang, Zongzhang Zhang, Yang Yu, Dong Li, Wulong Liu, and Jianye Hao. Adaptive Online Packing-guided Search for POMDPs. In: *Advances in Neural Information Processing Systems 34 (NeurIPS-2021)*, pages 28419-28430, Virtual Conference, 2021.
37. Xiong-Hui Chen, Shengyi Jiang, Feng Xu, Zongzhang Zhang, and Yang Yu. Cross-Modal Domain Adaptation for Cost-Efficient Visual Reinforcement Learning. In: *Advances in Neural Information Processing Systems 34 (NeurIPS-2021)*, pages 12520-12532, Virtual Conference, 2021.
38. Feng Xu, Shengyi Jiang, Hao Yin, Zongzhang Zhang, Yang Yu, Ming Li, Dong Li, and Wulong Liu. Enhancing Context-Based Meta-Reinforcement Learning Algorithms via An Efficient Task Encoder (Student Abstract). In: *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 15937-15938, Virtual Conference, 2021.
39. Chenyang Wu, Rui Kong, Guoyu Yang, Xianghan Kong, Zongzhang Zhang, Yang Yu, Dong Li, and Wulong Liu. LB-DESPOT: Efficient Online POMDP Planning Considering Lower Bound in Action Selection (Student Abstract). In: *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 15927-15928, Virtual Conference, 2021.
40. Cong Fei, Bing Wang, Yuzheng Zhuang, Zongzhang Zhang, Jianye Hao, Hongbo Zhang, Xuewu Ji, and Wulong Liu. Triple-GAIL: A Multi-Modal Imitation Learning Framework with Generative Adversarial Nets. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI-2020)*, pages 2929-2935, Yokohama, Japan, 2020.
41. Tianpei Yang, Jianye Hao, Zhaopeng Meng, Zongzhang Zhang, Weixun Wang, Yujing Hu, Yingfeng Chen, Changjie Fan, Wulong Liu, Zhaodong Wang, and Jiajie Peng. Efficient Deep Reinforcement Learning via Adaptive Policy Transfer. In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI-2020)*, pages 3094-3100, Yokohama, Japan, 2020.  
This work also appeared in: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2020)*, pages 2053-2055, as an extended abstract.
42. Chong Jiang, Zongzhang Zhang, Zixuan Chen, Jiacheng Zhu, and Junpeng Jiang. Third-person Imitation Learning via Image Difference and Variational Discriminator Bottleneck (Student Abstract). In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 13819-13820, New York, USA, 2020.
43. Jiacheng Zhu, Jiahao Lin, Meng Wang, Yingfeng Chen, Changjie Fan, Chong Jiang, and Zongzhang Zhang. Generative Adversarial Imitation learning from Failed Experiences (Student Abstract). In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 13997-13998, New York, USA, 2020.

44. Linjing Zhang and Zongzhang Zhang. Double Replay Buffers with Restricted Gradient. In: *Proceedings of the 27th International Conference on Neural Information Processing (ICONIP-2020)*, pages 295-306, Bangkok, Thailand, 2020.
45. Zhen Wu, Zongzhang Zhang, and Xiaofang Zhang. Recency-Weighted Acceleration for Continuous Control through Deep Reinforcement Learning. In: *Proceedings of the 27th International Conference on Neural Information Processing (ICONIP-2020)*, pages 604-615, Bangkok, Thailand, 2020.
46. Xiaobai Ma, Katherine R. Driggs-Campbell, Zongzhang Zhang, and Mykel J. Kochenderfer. Monte-Carlo Tree Search for Policy Optimization. In: *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*, pages 3116-3122, Macao, China, 2019.
47. Linjing Zhang, Zongzhang Zhang, Zhiyuan Pan, Yingfeng Chen, Jiangcheng Zhu, Zhaorong Wang, Meng Wang, and Changjie Fan. A Framework of Dual Replay Buffer: Balancing Forgetting and Generalization in Reinforcement Learning. In: *IJCAI-2019 Workshop on Scaling Up Reinforcement Learning*, Macao, China, 2019.
48. Zixuan Chen and Zongzhang Zhang. Deep Recurrent Policy Networks for Planning under Partial Observability. In: *Proceedings of the 28th International Conference on Artificial Neural Networks (ICANN-2019)*, pages 598-610, Munich, Germany, 2019.
49. Yisheng Wang and Zongzhang Zhang. Experience Selection in Multi-Agent Deep Reinforcement Learning. In: *Proceedings of the 31st International Conference on Tools with Artificial Intelligence (ICTAI-2019)*, pages 864-870, Portland, USA, 2019.
50. Yan Zheng, Zhaopeng Meng, Jianye Hao, Zongzhang Zhang, Tianpei Yang, and Changjie Fan. A Deep Bayesian Policy Reuse Approach Against Non-Stationary Agents. In: *Advances in Neural Information Processing Systems 31 (NeurIPS-2018)*, pages 960-970, Montreal, Canada, 2018.
51. Jiahao Lin and Zongzhang Zhang. ACGAIL: Imitation Learning about Multiple Intentions with Auxiliary Classifier GANs. In: *Proceedings of the 15th Pacific Rim International Conference on Artificial Intelligence (PRICAI-2018)*, pages 321-334, Nanjing, China, 2018.
52. Yan Zheng, Zhaopeng Meng, Jianye Hao, and Zongzhang Zhang. Weighted Double Deep Multiagent Reinforcement Learning in Stochastic Cooperative Environments. In: *Proceedings of the 15th Pacific Rim International Conference on Artificial Intelligence (PRICAI-2018)*, pages 421-429, Nanjing, China, 2018.
53. Zhiyuan Pan, Zongzhang Zhang, and Zixuan Chen. Asynchronous Value Iteration Network. In: *Proceedings of the 25th International Conference on Neural Information Processing (ICONIP-2018)*, pages 169-180, Siem Reap, Cambodia, 2018.
54. Zongzhang Zhang, Zhiyuan Pan, and Mykel J. Kochenderfer. Weighted Double Q-learning. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pages 3455-3461, Melbourne, Australia, 2017.
55. Zongzhang Zhang and Quan Liu. Covering Number: Analyses for Approximate Continuous-state POMDP Planning (Extended Abstract). In: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*, pages 1293-1294, Singapore, Singapore, 2016.
56. Jianwei Zhai, Quan Liu, Zongzhang Zhang, Shan Zhong, Haijun Zhu, Peng Zhang, and Cijia Sun. Deep Q-learning with Prioritized Sampling. In: *Proceedings of the 23rd International Conference on Neural Information Processing (ICONIP-2016)*, pages 13-22, Kyoto, Japan, 2016. (Finalist of Best Student Paper Award)
57. Weisheng Qian, Quan Liu, Zongzhang Zhang, Zhiyuan Pan, and Shan Zhong. Policy Graph Pruning and Optimization in Monte Carlo Value Iteration for Continuous-State POMDPs. In: *Proceedings of the 2016 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (IEEE ADPRL-2016)*, Athens, Greece, 2016.

58. Zongzhang Zhang, David Hsu, Wee Sun Lee, Zhan Wei Lim, and Aijun Bai. PLEASE: Palm Leaf Search for POMDPs with Large Observation Spaces. In: *Proceedings of the 25th International Conference on Automated Planning and Scheduling (ICAPS-2015)*, pages 249-257, Jerusalem, Israel, 2015.  
This work also appeared in: *Proceedings of the 8th Annual Symposium on Combinatorial Search (SoCS-2015)*, pages 238-239, as a two-page abstract.
59. Yicheng Zhou, Quan Liu, Qiming Fu, and Zongzhang Zhang. Trajectory Sampling Value Iteration: Improved Dyna Search for MDPs (Extended Abstract). In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 1685-1686, Istanbul, Turkey, 2015.
60. Shuhua You, Quan Liu, Zongzhang Zhang, Hui Wang, and Xiaofang Zhang. Intelligent Model Learning Based on Variance for Bayesian Reinforcement Learning. In: *Proceedings of the 27th International IEEE Conference on Tools with Artificial Intelligence (ICTAI-2015)*, pages 170-177, Salerno, Italy, 2015.
61. Zongzhang Zhang, David Hsu, and Wee Sun Lee. Covering Number for Efficient Heuristic-Based POMDP Planning. In: *Proceedings of the 31st International Conference on Machine Learning (ICML-2014)*, pages 28-36, Beijing, China, 2014.
62. Aijun Bai, Feng Wu, Zongzhang Zhang, and Xiaoping Chen. Thompson Sampling based Monte-Carlo Planning in POMDPs. In: *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS-2014)*, pages 28-36, Portsmouth, USA, 2014.
63. Zongzhang Zhang, Michael L. Littman, and Xiaoping Chen. Covering Number as a Complexity Measure for POMDP Planning and Learning. In: *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-2012)*, pages 1853-1859, Toronto, Ontario, Canada, 2012.
64. Zongzhang Zhang and Xiaoping Chen. FHHOP: A Factored Hybrid Heuristic Online Planning Algorithm for Large POMDPs. In: *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI-2012)*, pages 934-943, Catalina Island, United States, 2012.
65. Zongzhang Zhang and Xiaoping Chen. Accelerating Point-Based POMDP Algorithms via Greedy Strategies. In: *Proceedings of International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN-2010)*, pages 545-556, Darmstadt, Germany, 2010.

#### *Journal Publications*

---

1. Lei Yuan, Tao Jiang, Lihe Li, Feng Chen, Zongzhang Zhang, and Yang Yu. Robust Cooperative Multi-agent Reinforcement Learning via Multi-view Message Certification. *SCIENCE CHINA Information Sciences*, 2024, 67(4): 142102.
2. Lei Yuan, Feng Chen, Zongzhang Zhang, and Yang Yu. Communication-Robust Multi-Agent Learning by Adaptable Auxiliary Multi-Agent Adversary Generation. *Frontiers of Computer Science*, 2024, 18(6): 186331.
3. Chengxing Jia, Fuxiang Zhang, Tian Xu, Jing-Cheng Pang, Zongzhang Zhang, and Yang Yu. Model Gradient: Unified Model and Policy Learning in Model-based Reinforcement Learning. *Frontiers of Computer Science*, 2024, 18(4): 184339.
4. Chenyang Wu and Zongzhang Zhang. Surfing Information: The Challenge of Intelligent Decision-Making. *Intelligent Computing*, 2023, 2: Article 0041.
5. Rongjun Qin, Feng Chen, Tonghan Wang, Lei Yuan, Xiaoran Wu, Yipeng Kang, Zongzhang Zhang, Chongjie Zhang, and Yang Yu. Multi-Agent Policy Transfer via Task Relationship Modeling. *SCIENCE CHINA Information Sciences*, 2023. [Online]
6. Tian Chang, Zongzhang Zhang, and Yang Yu. Stochastic Ensemble Policy Transfer. *Journal of Frontiers of Computer Science and Technology*, 2022, 16(11): 2531-2536. (In Chinese with English Abstract).

7. Yan Zheng, Jianye Hao, Zongzhang Zhang, Zhaopeng Meng, Tianpei Yang, Yanran Li, and Changjie Fan. Efficient Policy Detecting and Reusing for Non-Stationarity in Markov Games. *Autonomous Agents and Multi-Agent Systems*, 2021, 35(2): 1-29.
8. Zixuan Chen, Zongzhang Zhang, Zhiyuan Pan, and Linjing Zhang. Planning Network Model Based on Generalized Asynchronous Value Iteration. *Journal of Software*, 2021, 32(11): 3496-3511. (In Chinese with English Abstract).
9. Chong Jiang, Zongzhang Zhang, Zixuan Chen, Jiacheng Zhu, and Junpeng Jiang. Data Efficient Third-person Imitation Learning Method. *Computer Science*, 2021, 48(2): 238-244. (In Chinese with English Abstract).
10. Yan Zheng, Jianye Hao, Zongzhang Zhang, Zhaopeng Meng, and Xiaotian Hao. Efficient Multiagent Policy Optimization Based on Weighted Estimators in Stochastic Environments. *Journal of Computer Science and Technology*, 2020, 35(2): 268-280.
11. Jiahao Lin, Zongzhang Zhang, Chong Jiang, and Jianye Hao. A Survey of Imitation Learning Based on Generative Adversarial Nets. *Chinese Journal of Computers*, 2020, 43(2): 326-351. (In Chinese with English Abstract).
12. Shan Zhong, Quan Liu, Zongzhang Zhang, and Qiming Fu. Efficient Reinforcement Learning in Continuous State and Action Spaces with Dyna and Policy Approximation, *Frontiers of Computer Science*, 2019, 13(1): 106-126.
13. Jin Xu, Quan Liu, Zongzhang Zhang, Bin Liang, and Qian Zhou. Asynchronous Deep Reinforcement Learning with Multiple Gating Mechanisms. *Chinese Journal of Computers*, 2019, 42(3): 636-653. (In Chinese with English Abstract).
14. Song Chen, Xiaofang Zhang, Zongzhang Zhang, Quan Liu, Jinjin Wu, and Yan Yan. Deep Double Q-Network Based on Linear Dynamic Frame Skip. *Chinese Journal of Computers*, 2019, 42(11): 2561-2573. (In Chinese with English Abstract).
15. Quan Liu, Jianwei Zhai, Zongzhang Zhang, Shan Zhong, Qian Zhou, Peng Zhang, and Jin Xu. A Survey on Deep Reinforcement Learning. *Chinese Journal of Computers*, 2018, 41(1): 1-27. (In Chinese with English Abstract). **(2017-2021's Best Paper Award)**
16. Quan Liu, Jianwei Zhai, Shan Zhong, Zongzhang Zhang, Qian Zhou, and Peng Zhang. A Deep Recurrent Q-Network Based on Visual Attention Mechanism. *Chinese Journal of Computers*, 2017, 40(6): 1353-1366. (In Chinese with English Abstract).
17. Zongzhang Zhang, Qiming Fu, Xiaofang Zhang, and Quan Liu. Reasoning and Predicting POMDP Planning Complexity via Covering Numbers. *Frontiers of Computer Science*, 2016, 10(4): 726-740.
18. Xiaofang Zhang, Zongzhang Zhang, Xiaoyuan Xie, and Yicheng Zhou. An Approach of Iterative Partition Testing Based on Priority Sampling. *Chinese Journal of Computers*, 2016, 39(11): 2307-2323. (In Chinese with English Abstract).
19. Shan Zhong, Quan Liu, Qiming Fu, Zongzhang Zhang, Fei Zhu, and Shengrong Gong. A Heuristic Dyna Optimizing Algorithm Using Approximate Model Representation. *Journal of Computer Research and Development*, 2015, 52(12): 2764-2775. (In Chinese with English Abstract).
20. Zongzhang Zhang and Xiaoping Chen. Hybrid Heuristic Online Planning for POMDPs. *Journal of Software*, 2013, 24(7): 1589-1600. (In Chinese with English Abstract).

### *Selected Book Chapters*

---

1. Zongzhang Zhang and Yang Yu. *Single-Agent Reinforcement Learning*. Distributed Artificial Intelligence, B. An, Y. Gao, and Y. Yu Eds., Electronic Industry Press, 2022, 207-246. (In Chinese).
2. Zongzhang Zhang and Mykel J. Kochenderfer. *Decision-Theoretic Planning in Partially Observable Environments*. Interactions in Multiagent Systems, J. Hao and H. Leung Eds., World Scientific, 2018, 65-90.

### *Reports*

---

1. Hongyao Tang, Jianye Hao, Tangjie Lv, Yingfeng Chen, Zongzhang Zhang, Hangtian Jia, Chunxu Ren, Yan Zheng, Changjie Fan, and Li Wang. Hierarchical Deep Multiagent Reinforcement Learning with Temporal Abstraction. *arXiv preprint* arXiv: 1809.09332v2.
2. Zongzhang Zhang. Complexity Theory and Planning Algorithms in Partially Observable Markov Decision Processes. *Ph.D. Thesis*, University of Science and Technology of China, School of Computer Science and Technology, 2012. (In Chinese).

### *Patents*

---

1. Zongzhang Zhang, Haoran Chen, Yishen Wang, and Yongliang Shen. Recognition System for Security Check and Control Method Thereof. US Patent of Intention, 2023, No. US11574152 B2.
2. Zongzhang Zhang, Zhiyuan Pan, and Hui Wang. Large Area Surveillance Method and Surveillance Robot Based on Weighted Double Deep Q-learning. US Patent of Intention, 2022, No. US11224970 B2.
3. Zongzhang Zhang, Yang Yu, Zhi-Hua Zhou, Chenyang Wu, and Guoyu Yang. Partially Observable Driving Planning Method Based on Adaptive Particles and Belief Packing. China Patent of Intention, 2022, No. ZL202110410291.1.
4. Zongzhang Zhang, Yang Yu, Zhi-Hua Zhou, Yafei Hu, and Feng Xu. Vehicle Adaptive Automatic Driving Decision-Making Method and System Based on Meta-Reinforcement Learning. China Patent of Intention, 2022, No. ZL202110356309.4.
5. Zongzhang Zhang, Weijiang Liao, Yang Yu, Ming Li, and Zhi-Hua Zhou. Partially Observed Intersection Autonomous Merging Method Based on Particle Attention Deep Q-Learning. China Patent of Intention, 2022, No. ZL202110337809.3.
6. Zongzhang Zhang, Yang Yu, and Chong Jiang. Robot Arm Action Learning Method and System Based on Third-Person Imitation Learning. China Patent of Intention, 2022, No. ZL202010040178.4.
7. Zongzhang Zhang, Yang Yu, Zhi-Hua Zhou, Yishen Wang, and Junpeng Jiang. Automated Driving Decision-Making Method and System Based on Partial Observable Transfer Reinforcement Learning. China Patent of Intention, 2021, No. ZL201911373375.1.
8. Zongzhang Zhang, Haoran Chen, and Yishen Wang. Recognition System for Security Check and Control Method Thereof. China Patent of Intention, 2021, No. ZL201810048208.9.
9. Jiacheng Zhu and Zongzhang Zhang. Automatic Parking Method and System Based on Generative Adversarial Imitation Learning. China Patent of Intention, 2021, No. ZL202010260031.6.
10. Linjing Zhang and Zongzhang Zhang. Intelligent Control System of Manipulator Based on Reservoir Sampling and Double Replay Buffers. China Patent of Intention, 2021, No. ZL202010202667.5.
11. Zongzhang Zhang, Jiahao Lin, Yingfeng Chen, and Changjie Fan. End-to-end Game Robot Generation Method and System Based on Multi-Category Imitation Learning. China Patent of Intention, 2020, No. ZL201810498479.4.



12. Zixuan Chen and Zongzhang Zhang. Planning Method for Autonomous Driving Systems. China Patent of Intention, 2020, No. ZL201811622988.X.
13. Chong Jiang and Zongzhang Zhang. Takeover Cruise Method and System Based on Automatic Reasoning Mechanism. China Patent of Intention, 2020, No. ZL201811168888.4.
14. Chong Jiang and Zongzhang Zhang. Decision-Making Method for Partially Observable Autonomous Driving Based on Constrained Online Planning. China Patent of Intention, 2020, No. ZL201810595164.1.
15. Yishen Wang, Zongzhang Zhang, and Haoran Chen. Recurrent Network Human-Robot Dialogue Method Based on Actor-Critic Reinforcement Learning. China Patent of Intention, 2020, No. ZL201811124263.8.
16. Quan Liu, Fei Zhu, Weisheng Qian, and Zongzhang Zhang. Robot Optimal Path Planning Method Based on Partially Observable Markov Decision Processes. China Patent of Intention, 2020, No. ZL201810102240.0.
17. Zongzhang Zhang, Zhiyuan Pan, and Hui Wang. Large Area Surveillance Method and Surveillance Robot Based on Weighted Double Deep Q-learning. China Patent of Intention, 2019, No. ZL201710329549.9.
18. Xiaofang Zhang, Qian Zhou, Zongzhang Zhang, and Yicheng Zhou. An Approach and System of Iterative Partition Testing. China Patent of Intention, 2018, No. ZL201610160340.X.

#### *Grants and Contracts*

---

1. National Natural Science Foundation of China (No. 62276126): “Research on Cooperative Multi-Agent Deep Reinforcement Learning Based on Knowledge Transfer”, 2013.01 – 2026.12, Principal Investigator.
2. National Natural Science Foundation of China (No. 61876119): “Deep Reinforcement Learning Based on Partially Observable Models: Theory and Applications”, 2019.01 – 2022.12, Principal Investigator.
3. National Natural Science Foundation of China (No. 61502323): “Covering Number for Theories and Methods of Uncertainty Planning in Partially Observable Environments”, 2016.01 – 2018.12, Principal Investigator.
4. Natural Science Foundation at Jiangsu Province (No. BK20221442): “Research on Transfer Reinforcement Learning in Cooperative Multi-Agent Environments”, 2022.07 – 2025.06, Principal Investigator.
5. Natural Science Foundation at Jiangsu Province (No. BK20181432): “Research on Theory and Methods of Planning and Reinforcement Learning in Partially Observable Environments”, 2018.07 – 2021.06, Principal Investigator.
6. Natural Science Foundation of Universities at Jiangsu Province (No. 16KJB520041): “Research on Cooperative Multi-agent Planning Methods Based on the Covering Number”, 2016.09 – 2018.06, Principal Investigator.
7. Tencent: “Sample-Efficient Reinforcement Learning for Complex Robot Control Scenarios”, 2023.7 – 2024.12, Principal Investigator.
8. China Electronics Technology Group Corporation: “Robust Reinforcement Learning Theory and Methods of Human-Machine Collaboration in Open Heterogeneous Dynamic Environments”, 2023.9 – 2024.6, Principal Investigator.
9. Ministry of Education - Tencent Industry-University Cooperative Education Project: “Course Construction on Multi-Agent Systems”, 2023.09-2024.09, Principal Investigator.
10. Hikvision Inc.: “Research on Efficient Collaboration Methods in Large-Scale Multi-Agent Systems”, 2021.11 – 2022.12, Principal Investigator.
11. Alibaba Group: “Research on Human-in-the-Loop Reinforcement Learning Methods for Risk Knowledge Acquisition”, 2021.07 – 2022.06, Principal Investigator.

12. Huawei Company: “Research on Data Efficient Reinforcement Learning Methods for Uncertain Environment Perception”, 2019.11 – 2020.11, Principal Investigator.
13. Nanjing University’s Original Cross Program (No. 022114380010): “Research on Application of Reinforcement Learning in Energy Storage Systems”, 2021.01-2023.12. Principal Investigator.
14. Nanjing University’s Initiative Research Project (No. 022114380005): “Research on Deep Reinforcement Learning Methods in Uncertain Environments”, 2020.01-2020.12. Principal Investigator.

### *Professional Services*

---

- **Editorial Board Member:** Intelligent Computing (AAAS/Science Partner Journal, 2022-2024)
- **Young Associate Editor:** Frontiers of Computer Science (2019-2025)
- **Area Chair/Senior Program Committee Member:** NeurIPS 2024; AAMAS 2024; ECAI 2020, 2024; IJCAI 2020-2021; AAAI 2019; ICAPS 2021
- **Member of the Novel Program Committee Board:** IJCAI 2022-2024
- **Program Committee Member/Reviewer:** AAAI 2018, 2020, 2022-2023; ICML 2019-2023; IJCAI 2013, 2017-2019; NeurIPS 2018-2023; ICLR 2021-2022; RLC 2024; AISTATS 2022; DAI 2019-2023; AAMAS 2021; ICAPS 2020; ECML-PKDD 2020; CoRL 2020; IJCNN 2020; CCDM 2020; ACML 2017-2019; PRICAI 2018-2019; ICA 2017-2019; ADPRL 2018; SSCI 2019; CCFAI 2019
- **Journal Reviewer:** Transactions on Pattern Analysis and Machine Intelligence; Artificial Intelligence; Journal of Artificial Intelligence Research; IEEE Transactions on Neural Networks and Learning Systems; IEEE Transactions on Cybernetics; ACM Transactions on Intelligent Systems and Technology; Autonomous Agents and Multi-Agent Systems; Machine Learning; Pattern Recognition; IEEE Computational Intelligence Magazine; Robotics and Autonomous Systems; Information Sciences; Frontiers of Computer Science; Neurocomputing; Knowledge-Based Systems; Applied Intelligence; Expert Systems with Applications; SCIENCE CHINA Information Sciences; Chinese Journal of Computers; Journal of Software; IEEE/CAA Journal of Automatica Sinica; Journal of Computer Research and Development
- **Workshop Co-chair:** Asian Workshop on Reinforcement Learning (AWRL) 2016-2018; PRICAI 2018’s Workshop on Methods and Applications of Reinforcement Learning
- **Local Organizing Committee Chair:** DAI 2020; MLA 2020, 2022
- **Professional Organization Membership:** CCF Senior Member; AAAI Member; IEEE Member
- **Reviewer Award:** ICLR 2021’s Outstanding Reviewer; NeurIPS 2019, 2022’s Top Reviewer
- **Consultant/Visiting Scholar:** Polixir Technologies (2019-now); Alibaba Group (2021-2022); Net-ease (2017-2020)

### *Teaching Experience*

---

<b>Multi-Agent Systems</b>	for undergraduate students
<i>Nanjing University</i>	<i>Spring 2021 – 2024</i>
<b>Big Data, Large Model, and Decision Intelligence</b>	for undergraduate students
<i>Nanjing University</i>	<i>Spring 2024</i>

<b>Control Theory and Methods</b> <i>Nanjing University</i>	for undergraduate and graduate students <i>Fall 2020 – 2023</i>
<b>Reinforcement Learning</b> <i>Nanjing University, Co-taught with Prof. Yang Yu</i>	for undergraduate and graduate students <i>Fall 2020 – 2023</i>
<b>Introduction to Artificial Intelligence</b> <i>Nanjing University, Co-taught with Prof. Yang Yu</i>	for undergraduate students <i>Fall 2021 – 2023</i>
<b>Intelligent Systems: Design and Application</b> <i>Nanjing University</i>	for undergraduate and graduate students <i>Spring 2020 – 2021</i>
<b>Intelligent Application Modeling</b> <i>Nanjing University, A Summer Course Co-constructed with Tencent</i>	for undergraduate students <i>July 2019</i>
<b>Introduction to Software Engineering</b> <i>Soochow University</i>	for undergraduate students <i>Fall 2015 – 2018</i>
<b>Computer Aided Software Engineering</b> <i>Soochow University</i>	for undergraduate students <i>Spring 2014 – 2017</i>

---

### *Research Advisor*

#### **Current Ph.D. Students**

*Nanjing University*

*Nanjing, Jiangsu, China*

- 2020 – : Feng Xu (co-advised), Weijian Liao (co-advised)
- 2022 – : Aoran Wang (co-advised)
- 2023 – : Chenyang Wu, Di Xue

#### **Current Master Students**

*Nanjing University*

*Nanjing, Jiangsu, China*

- 2021 – : Fuxiang Zhang, Chenghe Wang, Tianchi Li, Yichen Li, Yuhang Ran, Jiacheng Xu
- 2022 – : Feng Chen, Hao Ding, Chenxiao Gao, Fuguang Han, Rui Kong, Renzhe Zhou, Chao Chen, Xinyu Zhang
- 2023 – : Ningjing Chao, Tao Jiang, Tianyuan Liu, Tianyi Zhang, Wenjie Qiu, Yujian Zhu, Lei Zhang

#### **Past Master Students**

*Nanjing University*

*Nanjing, Jiangsu, China*

- 2020 – 2023: Chenyang Wu (PhD, Nanjing University), Di Xue (PhD, Nanjing University), Guoyu Yang (Xuanyuan Investment), Dongyu Guo (ByteDance), Quan He (Huawei, Noah Ark's Lab), Yafei Hu (Ali Fliggy)
- 2019 – 2022: Tian Chang (Education Department of Guangdong Province), Yue Chen (Tencent), Xi-anghan Kong (Huatai Securities), Yupeng Zhang (co-advised, Netease)

#### **Past Master Students**

*Soochow University*

*Suzhou, Jiangsu, China*

- 2018 – 2021: Jiarun Cai (Didi), Zhen Wu (Xuzhou Public Security Bureau), Jiacheng Zhu (Baidu), Lin-jing Zhang (PetroChina)
- 2017 – 2020: Zixuan Chen (Ph.D. Candidate in Nanjing University), Chong Jiang (BIRENTECH), Yishen Wang (Haixin Key Laboratory)
- 2016 – 2019: Zhiyuan Pan (China Telecom), Jiahao Lin (Polixir Technologies)