Zizhao **Wang**

zizhao.wang@utexas.edu · https://wangzizhao.github.io/

Research focus

World model

Reinforcement learning

Causality

Robot learning

Service

- *Co-organizer*, workshop on Causality for Robotics: Answering the Question of Why, IROS 2023
- *Program Committee*, workshop on Generalization in Planning, NeurIPS 2023
- *Reviewer*, ICML, NeurIPS, ICLR, ICRA, IROS, RA-L

zizhao.wang@utexas.edu
wangzizhao.github.io
google scholar
@duke_zzwang

Education

2020–	**PhD**, Electrical and Computer Engineering advisor: **Peter Stone**	**University of Texas at Austin**
2018 - 2019	**MS**, Computer Science	**Columbia University**
2016 - 2018	**BS**, Computer Engineering (dual degree program)	**University of Michigan**
2014 - 2018	**BS**, Electrical and Computer Engineering	**Shanghai Jiao Tong University**

Work Experience

2024	**Research Intern** Enhance **world model** with object-centric representation.	**Microsoft Research**
2024	**Research Intern** Improve the robustness of motion prediction in **autonomous driving** by reasoning about the **causal relationships** between vehicles.	**Honda Research Institute**

Selected Research Experience

World Model + Causality

Reason about causal relationships between different state factors (e.g., how objects depend on each other).
Improve world model's generalization and learning efficiency.

Unsupervised Reinforcement Learning (RL) + Causality

Reason about interactions between different state factors.
Propose novel intrinsic motivation and skill discovery algorithms to improve RL sample efficiency.

Selected Publications

See google scholar for a complete list of publications.

1. SkiLD: Unsupervised Skill Discovery Guided by Local Dependencies, *In Review*
   Zizhao Wang*, Jiaheng Hu*, Caleb Chuck*, Stephen Chen, Roberto Martín-Martín, Amy Zhang, Scott Niekum, Peter Stone.
2. Disentangled Unsupervised Skill Discovery for Efficient Hierarchical Reinforcement Learning, *In Review*
   Jiaheng Hu, **Zizhao Wang**, Peter Stone, Roberto Martín-Martín
3. Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning (**oral**), *AAAI 2024*
   Zizhao Wang*, Caroline Wang, Xuesu Xiao, Yuke Zhu, and Peter Stone.
4. ELDEN: Exploration via Local Dependencies, *NeurIPS 2023*
   Zizhao Wang*, Jiaheng Hu*, Roberto Martín-Martín, and Peter Stone.
5. Causal Dynamics Learning for Task-Independent State Abstraction (**oral**), *ICML 2022*
   Zizhao Wang, Xuesu Xiao, Zifan Xu, Yuke Zhu, and Peter Stone.
6. Learning to Correct Mistakes: Backjumping in Long-horizon Task and Motion Planning, *CoRL 2022*
   Yoonchang Sung*, **Zizhao Wang***, and Peter Stone.
7. Task-Independent Causal State Abstraction, *NeurIPS 2021, robot learning workshop*
   Zizhao Wang, Xuesu Xiao, Yuke Zhu, and Peter Stone.
8. CLAMGen: Closed-Loop Arm Motion Generation via Multi-view Vision-Based RL, *IROS 2021*
   Iretiayo Akinola*, **Zizhao Wang***, and Peter Allen.
9. From Agile Ground to Aerial Navigation: Learning from Learned Hallucination, *IROS 2021*
   Zizhao Wang, Xuesu Xiao, Alexander J Nettekoven, Kadhiravan Umasankar, Anika Singh, Sriram Bommakanti, Ufuk Topcu, and Peter Stone.
10. APPLE: Adaptive Planner Parameter Learning from Evaluative Feedback, *RAL 2021*
    Zizhao Wang, Xuesu Xiao, Garrett Warnell, and Peter Stone.
11. APPLI: Adaptive Planner Parameter Learning from Interventions, *ICRA 2021*
    Zizhao Wang, Xuesu Xiao, Bo Liu, Garrett Warnell, and Peter Stone.
12. Variational Objectives for Markovian Dynamics with Backward Simulation, *ECAI 2020*
    Antonio Khalil Moretti*, **Zizhao Wang***, Luhuan Wu*, Iddo Drori, and Itsik Pe'er.