

Zhihong Deng

Australian Artificial Intelligence Institute, University of Technology Sydney

dengzh0551@gmail.com | [Google Scholar](#) | [LinkedIn](#) | [Github](#) | [Homepage](#)

My current research focuses on **improving the reliability of autonomous agents from a causal perspective**. This not only allows agents to gain a deeper understanding of the world and make more informed decisions, but also fosters trust and transparency by elucidating the causal relationships behind their decisions, not just correlations. Ultimately, the goal of my research is to advance the progress of **next-generation AI agents** that are not only intelligent but also reliable, with enhanced **robustness, interpretability, fairness, and safety**.

EDUCATION

University of Technology Sydney, Sydney, Australia

Faculty of Engineering and Information Technology

Advisor: Jing Jiang and Chengqi Zhang

Ph.D. student

Mar. 2021–Now

Southern University of Science and Technology, Shenzhen, China

Department of Computer Science and Engineering

Visiting Scholar

Sep. 2020–Feb. 2021

Sun Yat-sen University, Guangzhou, China

School of Data and Computer Science

Advisor: Chang-Dong Wang

M.Eng.

Sep. 2018–Jun. 2020

Sun Yat-sen University, Guangzhou, China

School of Data and Computer Science

Advisor: Chang-Dong Wang

B.Eng.

Sep. 2013–Jun. 2017

PUBLICATION

What Hides behind Unfairness? Exploring Dynamics Fairness in Reinforcement Learning

[\[paper\]](#)[\[code\]](#)

Zhihong Deng, Jing Jiang, Guodong Long, Chengqi Zhang

International Joint Conference on Artificial Intelligence, 2024 (**CORE A***)

- A novel causal fairness notion called dynamics fairness, which is the first notion defined on the underlying mechanisms governing the environment, filling in the missing piece in studying long-term fairness in reinforcement learning.
- A general method to quantitatively evaluate dynamics fairness with theoretical guarantees that does not rely on making parametric assumptions.
- Provide a principled approach to monitor and improve fairness in RL systems.

False Correlation Reduction for Offline Reinforcement Learning [\[paper\]](#)[\[code\]](#)

Zhihong Deng, Zuyue Fu, Lingxiao Wang, Zhuoran Yang, Chenjia Bai, Tianyi Zhou, Zhaoran Wang, Jing Jiang

IEEE Transactions on Pattern Analysis and Machine Intelligence, 2023 (**CORE A***, **IF=24.31**)

- A simple yet effective approach with theoretical guarantees for offline RL without OOD sampling and large-scale ensemble.
- Deploys an annealing behavior cloning regularizer to help produce high-quality estimations of uncertainty, which is crucial for eliminating false correlations in offline RL.
- Achieves state-of-the-art performance with 3.1x speedup on D4RL.

Causal Reinforcement Learning: A Survey [\[paper\]](#)[\[repo\]](#)[\[tutorial\]](#)

Zhihong Deng, Jing Jiang, Guodong Long, Chengqi Zhang

Transactions on Machine Learning Research, 2023 (**Survey Certification**)

- The first comprehensive survey on causal reinforcement learning.
- A problem-oriented taxonomy and analyses of existing methods and their limitations.
- Shed light on promising research directions including advancing theoretical understanding, establishing benchmarks, and tackling specific problem settings.

Pessimistic Bootstrapping for Uncertainty-Driven Offline Reinforcement Learning

[[paper](#)][[code](#)]

Chenjia Bai, Lingxiao Wang, Zhuoran Yang, **Zhihong Deng**, Animesh Garg, Peng Liu, Zhaoran Wang
International Conference on Learning Representations, 2022 (**CORE A***, **Spotlight Paper**)

- An uncertainty-driven model-free algorithm for offline RL.
- A novel approach to quantify uncertainty in offline RL using OOD samples, allowing refined conservatism and better control of OOD actions.

G³SR: Global Graph Guided Session-Based Recommendation [[paper](#)][[code](#)]

Zhihong Deng, Chang-Dong Wang, Ling Huang, Jian-Huang Lai, Philip S. Yu

IEEE Transactions on Neural Networks and Learning Systems, 2022 (**CORE A***, **IF=10.4**)

- An innovative graph-based framework for session-based recommendation that incorporates a unsupervised pre-training stage to extract global item-to-item relational information.
- A novel parameter-free readout mechanism to aggregate item representations.

DeepCF: A Unified Framework of Representation Learning and Matching Function Learning in Recommender System [[paper](#)][[code](#)][[talk](#)]

Zhihong Deng, Ling Huang, Chang-Dong Wang, Jian-Huang Lai, Philip S. Yu

Annual AAAI Conference on Artificial Intelligence, 2019 (**CORE A***, **Oral Paper**)

- A general framework for recommender systems that unifies deep learning-based representation learning and matching function learning.
- A novel model that has great flexibility to learn the complex matching function while being efficient to learn low-rank relations between users and items.

Serendipitous recommendation in e-commerce using innovator-based collaborative filtering

[[paper](#)]

Chang-Dong Wang, **Zhihong Deng**, Jian-Huang Lai, Philip S. Yu

IEEE Transactions on Cybernetics, 2019 (**CORE A**, **IF=11.8**)

- One of the earliest attempts to solve the serendipity problem in recommender systems.
- Introduces the Rogers' Innovation Adoption Curve from sociology to identify innovators.
- A recommendation algorithm that captures users' potential interests through a group of innovators who are skilled at discovering new and niche items.

HONOURS & AWARDS

Student Best Paper Award (Three Ph.D. students per year) 2023

Australian Artificial Intelligence Institute

International Research Training Program Scholarship 2021

University of Technology Sydney

First-class Scholarship for Graduate Students (Top 5% students) 2020

Sun Yat-sen University

National Scholarship for Graduate Students (Top 1% students) 2020

Sun Yat-sen University

INVITED TALKS & PRESENTATIONS

Causal Reinforcement Learning: Empowering Agents with Causality [[slides](#)] 2024

Zhihong Deng, Jing Jiang, Chengqi Zhang 2023

International Joint Conference on Artificial Intelligence (IJCAI). Jeju, South Korea

International Joint Conference on Neural Networks (IJCNN). Yokohama, Japan

International Conference on Advanced Data Mining and Applications (ADMA). Shenyang, China

From Predict to Control & From RL to Offline RL [[slides](#)] 2020

Zhihong Deng

Southern University of Science and Technology. Shenzhen, China

A Unified Framework of Representation Learning and Matching Function Learning in Recommender System [[slides](#)]

2019

Zhihong Deng

Graduate Student Academic Forum of The Guangdong-Hong Kong-Macao Greater Bay.
Guangzhou, China.

PROFESSIONAL SERVICE

Conference Reviewer

ICML, NeurIPS, IJCAI, KDD, CIKM, ICDM, ADMA

Journal Reviewer

IEEE TNNLS, IEEE TCYB, ESWA

PERSONAL INTERESTS

I am passionate about documenting and analyzing insights from academic papers, a hobby that has led me to amass over 10,000 followers on Zhihu, a popular knowledge-sharing platform in China. This activity allows me to exchange ideas and grow alongside individuals from diverse backgrounds. Additionally, I enjoy swimming and photography, hobbies I have enjoyed since childhood. These activities not only provide me with relaxation but also enhance my creativity and focus.
