Junzhe Zhang

111 College Pl
Syracuse, NY 13210
☐ +1 (734) 546-9634
✉ jzhan403@syr.edu
🌐 junzhez.com
Last updated November 28, 2024

Research Interests

Causal Inference, Reinforcement Learning, Algorithmic Fairness.

Education

2015 - 2023   **PhD. in Computer Science**, *Columbia University*, New York
Dissertation Title: *Towards Causal Reinforcement Learning*
Advisor: Elias Bareinboim
Committee Members: Daniel Hsu, Nathan Kallus, Vishal Misra, Lihong Li
- 2019 - 2023. Columbia University, New York
- 2015 - 2019. Purdue University, West Lafayette (Transferred)

2019 - 2021   **M.Sc. in Computer Science**, *Columbia University*, New York

2011 - 2014   **B.Sc. in Computer Science Honor and Math Science**, *University of Michigan*, Ann Arbor

Work Experience

Aug. 2024 - ...   **Assistant Professor**, *Syracuse University*, New York, Department of Electrical Engineering and Computer Science

2023 - 2024   **Postdoctoral Research Scientist**, *Columbia University*, New York, Causal AI Lab

2019 - 2023   **Graduate Research Assistant**, *Columbia University*, New York

2015 - 2019   **Graduate Research Assistant**, *Purdue University*, West Lafayette

2014 - 2015   **Software Engineer**, *Microsoft*, Redmond, Azure

Selected Publications

Citations via Google Scholar: h-index of 12+, total citation count of 1100+. **Bold*** means I am the first, co-first, or corresponding author.

17. Causal Imitation for Markov Decision Processes: a Partial Identification Approach
Kangrui Ruan, **Junzhe Zhang***, Xuan Di, Elias Bareinboim.
Proceedings of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

16. Causally Aligned Curriculum Learning
Mingxuan Li, Junzhe Zhang, Elias Bareinboim.
Proceedings of the 12th Twelfth International Conference on Learning Representations (ICLR), 2024.
(Acceptance rate = 31%)

15. Towards Safe Policy Learning under Partial Identifiability: A Causal Approach
Shalmali Joshi, **Junzhe Zhang***, Elias Bareinboim.
Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence (AAAI), 2024.
(Acceptance rate = 23.75%)

14. Scores for Learning Discrete Causal Graphs with Unobserved Confounders
Alex Bellot, Junzhe Zhang, Elias Bareinboim.
Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence (AAAI), 2024.
(Acceptance rate = 23.75%)

13. Causal Imitation Learning via Inverse Reinforcement Learning
Kangrui Ruan, **Junzhe Zhang***, Xuan Di, Elias Bareinboim.
Proceedings of the 11th Eleventh International Conference on Learning Representations (ICLR), 2023.
(Acceptance rate = 31.8%)

12. Online Reinforcement Learning for Mixed Policy Scopes
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.
(Acceptance rate = 25.6%)

11. Partial Counterfactual Identification from Observational and Experimental Data
Junzhe Zhang*, Jin Tian, Elias Bareinboim.
Proceedings of the 39th International Conference on Machine Learning (ICML), 2022.
(Acceptance rate $< 20\%$)

10. Can Humans Be Out of the Loop?
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 1st Conference on Causal Learning and Reasoning (CleaR), 2022.

9. Sequential Causal Imitation Learning with Unobserved Confounders
Daniel Kumor, Junzhe Zhang, Elias Bareinboim.
Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate $= 1\%$ (Oral Presentation))

8. Bounding Causal Effects on Continuous Outcomes
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), 2021.
(Acceptance rate $= 21\%$)

7. Causal Imitation Learning with Unobserved Confounders
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate $= 1\%$ (Oral Presentation))

6. Designing Optimal Dynamic Treatment Regimes: A Causal Reinforcement Learning Approach
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 37th International Conference on Machine Learning (ICML), 2020.
(Acceptance rate $= 21.8\%$)

5. Near-Optimal Reinforcement Learning in Dynamic Treatment Regimes
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 33rd Annual Conference on Neural Information Processing Systems(NeurIPS), 2019.
(Acceptance rate $= 21\%$)

4. Equality of Opportunity in Classification: A Causal Approach
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 31st Annual Conference on Neural Information Processing Systems(NeurIPS), 2018.
(Acceptance rate $= 21\%$)

3. Non-Parametric Path Analysis in Structural Causal Models
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI), 2018.
(Acceptance rate $= 9\%$ (Plenary Presentation))

2. Fairness in Decision-Making – The Causal Explanation Formula
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018.
(Acceptance rate $= 21\%$)

1. Transfer Learning in Multi-Armed Bandits: A Causal Approach
Junzhe Zhang*, Elias Bareinboim.
Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), 2017.
(Acceptance rate $= 26\%$)

Teaching Experience

Spring, 2023	**Causal Inference II**, *Columbia University*, New York, Guest Instructor
Spring 2020 - 2022	**Causal Inference II**, *Columbia University*, New York, Teaching Assistant
Fall 2020	**Causal Inference I**, *Columbia University*, New York, Teaching Assistant
Spring 2018	**Advanced Machine Learning**, *Purdue University*, West Lafayette, Teaching Assistant
Spring 2016	**Software Engineering Senior Project**, *Purdue University*, West Lafayette, Teaching Assistant
Fall 2015	**Computer Architecture**, *Purdue University*, West Lafayette, Teaching Assistant

Invited Talks & Presentations

October 2023	INFORMS, Phoenix, AZ. *Could Humans be out of the Loop?*
October 2023	INFORMS, Phoenix, AZ. *Bayesian Approach to Bounding Counterfactual Fairness Measure* (with S. Rho and E. Bareinboim)

April 2023   Causality Discussion Group, Virtual

Tutorials

March 2021   **Causal Fairness Analysis**, *ACM Conference on Fairness, Accountability, and Transparency (FaccT)*, Virtual, with E. Bareinboim and D. Plecko

August 2019   **Causal Reinforcement Learning**, *International Joint Conference on Artificial Intelligence (IJCAI)*, Macau China, with E. Bareinboim and S. Lee

Community Services

Area Chair-Conferences:

2025   ICML

Program Committee-Conferences:

2024   NeurIPS, ICML, AAAI, ICLR, IJCAI

2023   NeurIPS, ICML, AAAI, ICLR, IJCAI, UAI, AISTATS

2022   NeurIPS, ICML, AAAI, ICLR, IJCAI, UAI, AISTATS

2021   NeurIPS, ICML, AAAI, ICLR, UAI, AISTATS

2020   NeurIPS, ICML, AAAI, UAI, AISTATS

2019   NeurIPS

Journal Reviewing:

2023   Transactions on Machine Learning Research

2022   Transactions on Machine Learning Research, Journal of Causal Inference

2021   Journal of Artificial Intelligence Research, IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Neural Networks and Learning Systems

2020   Machine Learning Journal, Artificial Intelligence Journal

Academic References

○ **Elias Bareinboim**, Associate Professor of Computer Science, Columbia University, E-mail: eb@cs.columbia.edu

○ **Jin Tian**, Professor of Computer Science, Iowa State University, E-mail: jtian@iastate.edu

○ **Xuan Sharon Di**, Associate Professor of Civil Engineering and Engineering Mechanics, E-mail: sharon.di@columbia.edu

○ **Shalmali Joshi**, Assistant Professor of Biomedical Informatics, E-mail: shalmali.joshi@columbia.edu