

JUNZHE ZHANG

500 W 120th St (Mudd Building), Columbia University,
New York, NY 10027
(734)546-9634 ◊ junzhez@cs.columbia.edu ◊ www.cs.columbia.edu/~junzhez/

RESEARCH INTERESTS

Causal Inference; Reinforcement Learning; Fairness Analysis.

EDUCATION

Doctor of Philosophy, Computer Science *August 2019 - Present*
Columbia University, New York

Doctor of Philosophy, Computer Science *September 2015 - July 2019*
Purdue University, West Lafayette *(Transferred)*

Bachelor of Science, Computer Science Honor and Math Science *September 2011 - April 2014*
University of Michigan, Ann Arbor

Bachelor of Engineering, Software Engineering *September 2009 - June 2011*
Fudan University, Shanghai, China *(Transferred)*

SELECTED PUBLICATIONS

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

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))

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%)

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%)

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%)

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))

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% (Oral Presentation))

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% (Oral Presentation))

PREPRINTS AND OTHERS

Partial Identification of Counterfactual Distribution

Junzhe Zhang, Jin Tian, Elias Bareinboim.

Columbia CausalAI Laboratory, Technical Report (R-78), May, 2021.

Characterizing the Limits of Autonomous Systems

Junzhe Zhang, Elias Bareinboim.

Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2018.

(Acceptance rate = 47% (Extended Abstract Paper))

TUTORIALS

Causal Reinforcement Learning (with Elias Bareinboim, Drago Plecko)

ACM Conference on Fairness, Accountability, and Transparency (FaccT), Virtual, Mar, 2021

Causal Reinforcement Learning (with Elias Bareinboim, Sanghack Lee)

International Joint Conference on Artificial Intelligence (IJCAI), Macau, China, August 2019.

RESEARCH EXPERIENCE

Columbia University, Prof. Elias Bareinboim

Research Assistant

New York, NY, USA

September 2019 - Present

Purdue University, Prof. Elias Bareinboim

Research Assistant

West Lafayette, IN, USA

September 2015 - July 2019

University of Michigan, Prof. Peter Chen

Research Assistant

Ann Arbor, MI, USA

October 2013 - April 2014

University of Michigan, Prof. Robert Dick

Research Assistant

Ann Arbor, MI, USA

April 2012 - September 2012

INDUSTRIAL EXPERIENCE

Microsoft, Azure

Software Engineer

Redmond, WA, USA

June 2014 - June 2015

Qualcomm

Interim Engineering Intern

San Diego, CA, USA

May 2013 - August 2013

PROFESSIONAL SERVICE

Program Committee-Conferences:

- Neural Information Processing Systems (NeurIPS), 2019, 2020.
- International Conference on Machine Learning (ICML), 2020, 2021.
- Association for the Advancement of Artificial Intelligence Conference (AAAI), 2020, 2021.
- International Conference on Learning Representations (ICLR), 2021.
- Conference on Uncertainty in Artificial Intelligence (UAI), 2020, 2021.
- Artificial Intelligence and Statistics (AISTATS), 2020, 2021.

Reviewer-Journals:

- 2021: Journal of Artificial Intelligence Research.
- 2020: Machine Learning Journal.
- 2020: Artificial Intelligence Journal.