Elias Bareinboim

Email: eb@cs.columbia.edu — **Webpage:** https://causalai.net — **X:** @eliasbareinboim
Address: 530 W 120th St (Schapiro), Columbia University, New York, NY, 10027.

Research Interests

- Causal Inference: Theory and Applications.
- Causal Data Science; Causal Fairness Analysis; Causal Reinforcement Learning.
- Artificial Intelligence, Machine Learning, Statistics.
- Cognitive Science, Philosophy of Science.

Education

- Ph.D. in Computer Science – University of California, Los Angeles (UCLA), 2014.
  Title: *Generalizability in Causal Inference: Theory and Algorithms*.
  Advisor: *Judea Pearl*.
- B.Sc., M.Sc. in Computer Science – Federal University of Rio de Janeiro (UFRJ), 2007.
  Title: *Descents and nodal load in scale-free networks*.
  Advisor: *Valmir C. Barbosa*.

Academic Positions

- Associate Professor (tenured), Computer Science, Columbia University, Summer/2019-now.
  - Director, Causal Artificial Intelligence Laboratory.
  - Member, Data Science Institute.
  - Member, NSF National AI Institute for Artificial and Natural Intelligence.
  - Member, Program for Mathematical Genomics.
  - Member, Advisory Board, Columbia-DreamSports AI Innovation Center.
- Assistant Professor, Computer Science, Purdue University, Fall/2015-Spring/2019.
  - Director, Causal Artificial Intelligence Laboratory.
  - Assistant Professor, courtesy appointment, Statistics.
  - Faculty Affiliate, Regenstrief Center for Healthcare Engineering.
- Postdoctoral Scholar, Cognitive Systems Lab/UCLA, Judea Pearl, Fall/2014-Summer/2015.
- Research Assistant, Cognitive Systems Lab/UCLA, Judea Pearl, Fall/2009-Summer/2014.

Awards and Honors

- 2023 DARPA Young Faculty Award.
- 2022 ONR Young Investigator Award.
- 2021 JP Morgan Faculty Research Award.
- 2020 Amazon Research Award.
- 2019 UAI Best Paper Award (1 out of 450 papers).

- 2018 NSF Faculty Early Career Development (CAREER) Award.
- 2018 Adobe Data Science Research Award.
- 2018 UAI Best Student Paper Award (1 out of 337 papers).
- 2018 AAAI Outstanding Paper Award Honorable Mention (2 out of 3800 papers).
- 2017 IBM Open Collaborative Award.
- 2016 IEEE AI's 10 to Watch, Intelligent Systems.
- 2015 ACM Notable Paper, 19th Annual Best of Computing, Computing Reviews.
- 2014 UCLA Edward K. Rice Outstanding Doctoral Student Award (given to a single PhD student in all engineering and applied sciences majors), School of Engineering and Applied Sciences, UCLA.
- 2014 AAAI Outstanding Paper Award (1 out of 1406 papers).
- 2014 UCLA Outstanding Graduating PhD Student (commencement award), Computer Science.
- 2014 Google Outstanding Graduate Research Award, Computer Science, UCLA.
- 2014 Dan David Scholar, Future Dimension: Artificial Intelligence ($15,000), Dan David Foundation.
- 2013 UCLA Dissertation Year Fellowship (DYF) (~$35,000).
- 2012 Yahoo! Key Scientific Challenges Award, area Machine Learning & Statistics ($5,000).
- 2008 UCLA Ph.D.'s Fellowship (~$45,000).
- 2008 Top 10 award – National contest of M.Sc. thesis (2007), Brazilian Computer Society.
- 2008-2012 Ph.D.'s Fellowship, Fulbright – U.S. Dep. of State / CAPES-MEC, declined.
- 2003-2007 Undergraduate's and Master's Fellowships, Brazilian Research Council CNPq.

Publications

131. Counterfactual Image Editing with Disentangled Causal Latent Space.
Yushu Pan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-137), May/2025.

130. Epidemiology of LLMs: A Benchmark for Observational Distribution Knowledge.
Drago Plecko, Patrik Okanovic, Torsten Hoefler, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-136), May/2025.

129. Causal Explanations through Counterfactual Variable Attributions.
Kai-Zhan Lee, Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-135), May/2025.

128. Less Greedy Equivalence Search.
Adiba Ejaz, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-134), May/2025.

127. Adapting, Fast and Slow: A Causal Approach to Few-Shot Sequence Learning.
Kasra Jalaldoust, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-133), May/2025.

126. Confounding Robust Deep Reinforcement Learning: A Causal Approach.
Mingxuan Li, Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-132), May/2025.

125. Causal Generative Modeling for Confounding Robust Treatment Evaluation.
Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-131), May/2025.

124. A Hierarchy of Graphical Models for Counterfactual Inferences.
Hongshuo Yang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-130), May/2025.

123. Learning Invariances for Causal Abstraction Inference.
Paul Kroeger, Kevin Xia, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-129), May/2025.

122. Causal Discovery over Clusters of Variables in Markovian Systems.
Tara Anand, Adèle H. Ribeiro, Jin Tian, George Hripcsak, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-128), forthcoming.

121. From Black-box to Causal-box: Towards Building More Interpretable Models.
Inwoo Hwang, Yushu Pan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-127), May/2025.

120. Sample Complexity of Few-Shot Learning: a Causal Perspective.
Julia Kostin, Kasra Jalaldoust, Elias Bareinboim, Fanny Yang, Samory Kpotufe
Columbia CausalAI Laboratory, Technical Report (R-126), forthcoming.

119. Counterfactual Rationality: A Causal Approach to Game Theory.
Aurghya Maiti, Prateek Jain, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-125), Jan/2025.

118. Causal Abstraction Inference under Lossy Representations.
Kevin Xia, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-124), Jan/2025.
Proceedings of the 42nd International Conference on Machine Learning (ICML), 2025.
(Acceptance rate = 26.9%)

117. Automatic Reward Shaping from Confounded Offline Data.
Mingxuan Li, Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-123), Jan/2025.
Proceedings of the 42nd International Conference on Machine Learning (ICML), 2025.
(Acceptance rate = 26.9%)

116. Structural Causal Bandits under Markov Equivalence.
Min Woo Park, Andy Arditi, Elias Bareinboim, Sanghack Lee
Columbia CausalAI Laboratory, Technical Report (R-122), Jan/2025.

115. An Algorithmic Approach for Causal Health Equity: A Look at Race Differentials in Intensive Care Unit (ICU) Outcomes.
Drago Plecko, Paul Secombe, Andrea Clarke, Amelia Fiske, Samarra Toby, Donisha Duff, David Pilcher, Leo Celi, Rinaldo Bellomo, Elias Bareinboim (2025)
Columbia CausalAI Laboratory, Technical Report (R-121), Jan/2025.

114. Beyond the back-door: Probabilities of Identification
D. Plecko, D. Bradac, M. Bucic, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-118), May/2025.

113. Testing Causal Models with Hidden Variables in Polynomial Delay via Conditional Independencies
Hyung Jung, Adiba Ejaz, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-117), Aug/2024.
*Proceedings of the 39th AAAI Conference on Artificial Intelligence (*AAAI), 2025.
(Acceptance rate < 5% (oral))

112. Counterfactual Identification Under Monotonicity Constraints.
Aurghya Maiti, Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-116), Aug/2024.
*Proceedings of the 39th AAAI Conference on Artificial Intelligence (*AAAI), 2025.
(Acceptance rate = 23.4%)

111. Counterfactual Graphical Models: Constraints and Inference.
Juan Correa, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-115), Aug/2024.
Proceedings of the 42nd International Conference on Machine Learning (ICML), 2025.
(Acceptance rate, 5.2% (spotlight))

110. Characterizing and Learning Multi-domain Causal Structures from Observational and Exp. Data.
Adam Li, Amin Jaber, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-114), Aug/2024.

109. Counterfactual Realizability.
Arvind Raghavan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-113), May/2024.
Proc. of the 13th International Conference on Learning Representations (ICLR), 2025.
(Acceptance rate < 5.2% (spotlight))

108. Unified Covariate Adjustment for Causal Inference.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-112), May/2024.
Proc. of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

107. Disentangled Representation Learning in Non-Markovian Causal Systems.
Adam Li, Yushu Pan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-110), May/2024.
Proc. of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

106. Mind the Gap: A Causal Perspective on Bias Amplification in Prediction & Decision-Making.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-108), May/2024.
Proc. of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

105. Fairness-Accuracy Trade-Offs: A Causal Perspective.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-107), May/2024.
*Proceedings of the 39th AAAI Conference on Artificial Intelligence (*AAAI), 2025.
(Acceptance rate = 23.4%)

104. Eligibility Traces for Confounding Robust Off-Policy Evaluation.
Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-105), May/2024.
*Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence (*UAI), 2025.
(Acceptance rate = 30.7%)

103. Causal Imitation for Markov Decision Processes: a Partial Identification Approach.
Darren Kangrui, Junzhe Zhang, Sharon Di, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-104), May/2024.
Proc. of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

102. Counterfactual Image Editing.
Yushu Pan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-103), Dec/2023.
Proceedings of the 40th International Conference on Machine Learning (ICML), 2024.
(Acceptance rate = 27.5%)

101. Causally Aligned Curriculum Learning.
Mingxuan Li, Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-102), Oct/2023.
Proc. of the 12th International Conference on Learning Representations (ICLR), 2024.
(Acceptance rate = 31%)

100. Neural Causal Abstractions.
Kevin Xia, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-101), Dec/2023.
*Proceedings of the 38th AAAI Conference on Artificial Intelligence (*AAAI), 2024.
(Acceptance rate = 23.7%)

99. Transportable Representations for Out-of-distribution Generalization.
Kasra Jalaldoust, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-99), May/2023.
*Proceedings of the 38th AAAI Conference on Artificial Intelligence (*AAAI), 2024.
(Acceptance rate = 23.7%)

98. Causal discovery from observational and interventional data across multiple environments.
Adam Li, Amin Jaber, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-98), May/2023.
Proc. of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS), 2023.
(Acceptance rate = 26%)

97. Estimating Causal Effects Identifiable from a Combination of Observations and Experiments.
Yonghan Jung, Ivan Diaz, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-97), May/2023.
Proc. of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS), 2023.
(Acceptance rate = 26%)

96. Towards Safe Policy Learning under Partial Identifiability: A Causal Approach.
Shalmali Joshi, Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-96), May/2023.
*Proceedings of the 38th AAAI Conference on Artificial Intelligence (*AAAI), 2024.
(Acceptance rate = 23.75%)

95. Causal Fairness for Outcome Control.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-95), May/2023.
Proc. of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS), 2023.
(Acceptance rate = 26%)

94. Nonparametric Identifiability of Causal Representations from Unknown Interventions.
Julius von Kügelgen, Michel Besserve, Wendong Liang, Luigi Gresele, Armin Kekic,
Elias Bareinboim, David Blei, Bernhard Schölkopf
Columbia CausalAI Laboratory, Technical Report (R-94), June/2023.
Proc. of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS), 2023.
(Acceptance rate = 26%)

93. A Causal Framework for Decomposing Spurious Variations.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-93), May/2023.
Proc. of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS), 2023.
(Acceptance rate = 26%)

92. Reconciling Predictive and Statistical Parity: A Causal Approach.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-92), February/2023.
*Proceedings of the 38th AAAI Conference on Artificial Intelligence (*AAAI), 2024.
(Acceptance rate = 23.75%)

91. Estimating Joint Treatment Effects by Combining Multiple Experiments.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-91), Apr/2023.
Proceedings of the 39th International Conference on Machine Learning (ICML), 2023.
(Acceptance rate = 27%)

90. Causal Fairness Analysis.
Drago Plecko, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-90), July/2022.
Foundations and Trends in Machine Learning, Vol. 17: No. 3, pp 304-589, 2024.

89. Causal Imitation Learning via Inverse Reinforcement Learning.
Darren Kangrui, Junzhe Zhang, Sharon Di, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-89), May/2022.
Proc. of the 11th International Conference on Learning Representations (ICLR), 2023.
(Acceptance rate = 31%)

88. Partial Transportability for Domain Generalization.
Kasra Jalaldoust, Alexis Bellot, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-88), May/2023.
Proc. of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS), 2024.
(Acceptance rate = 25.8%)

87. Neural Causal Models for Counterfactual Identification and Estimation.
Kevin Xia, Yushu Pan, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-87), May/2022.
Proc. of the 11th International Conference on Learning Representations (ICLR), 2023.
(Acceptance rate = 31%)

86. Causal Identification under Markov equivalence: Calculus, Algorithm, and Completeness.
Amin Jaber, Adele Ribeiro, Jiji Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-86), May/2022.
Proc. of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.
(Acceptance rate < 2% (highlight))

85. Finding and Listing Front-door Adjustment Sets.
Hyunchai Jeong, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-85), Sep/2022.
Proc. of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.
(Acceptance rate = 26%)

84. Online Reinforcement Learning for Mixed Policy Scopes.
Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-84), May/2022.
Proc. of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.
(Acceptance rate = 26%)

83. Scores for Learning Discrete Causal Graphs with Unobserved Confounders.
Alexis Bellot, Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-83), May/2022.
*Proceedings of the 38th AAAI Conference on Artificial Intelligence (*AAAI), 2024.
(Acceptance rate = 23.75%)

82. Counterfactual Transportability: A Formal Approach.
Juan Correa, Sanghack Lee, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-82), May/2022.
Proceedings of the 38th International Conference on Machine Learning (ICML), 2022.
(Acceptance rate = 21%)

81. On Measuring Causal Contributions via do-Interventions.
Yonghan Jung, Shiva Kasiviswanathan, Jin Tian, Dominik Janzing, Elias Bareinboim (2022)
Columbia CausalAI Laboratory, Technical Report (R-81), May/2022.
Proceedings of the 38th International Conference on Machine Learning (ICML), 2022.
(Acceptance rate = 21%)

80. The Causal-Neural Connection: Expressiveness, Learnability, and Inference.
Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-80), May/2021.
Proc. of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate = 26%)

79. Nested Counterfactual Identification from Arbitrary Surrogate Experiments.
Juan Correa, Sanghack Lee, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-79), May/2021.
Proc. of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate = 26%)

78. Partial Counterfactual Identification from Observational and Interventional Data.
Junzhe Zhang, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-78), May/2021.
Proceedings of the 38th International Conference on Machine Learning (ICML), 2022.
(Acceptance rate = 21%)

77. Effect Identification in Causal Diagrams with Clustered Variables.
Tara Anand, Adele Ribeiro, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-77), May/2021.
*Proceedings of the 37th AAAI Conference on Artificial Intelligence (*AAAI), 2023.
(Acceptance rate = 19.6%)

76. Sequential Causal Imitation Learning with Unobserved Confounders.
Daniel Kumor Justin Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-76), May/2021.
Proc. of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate < 1% (oral))

75. Double Machine Learning Density Estimation for Local Treatment Effects with Instruments.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-75), May/2021.
Proc. of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate < 3% (spotlight))

74. Causal Transportability for Neural Representations.
C. Mao, K. Xia, J. Wang, H. Wang, J. Yang, E. Bareinboim, C. Vondrick
Columbia CausalAI Laboratory, Technical Report (R-74), forthcoming.
Proc. of IEEE/CVF Conference on Computer Vision & Pattern Recognition (CVPR), 2022.
(Acceptance rate = 25%)

73. Causal Inference and Data Fusion: Towards an Accelerated Process of Scientific Discovery.
Adele Ribeiro, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-73), Apr/2022.
Organisation for Economic Co-operation and Development (OECD),
Volume "AI and the productivity of science", forthcoming.

72. Non-Parametric Methods for Partial Identification of Causal Effects.
Junzhe Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-72), Feb/2021.

71. Estimating Identifiable Causal Effects on Markov Equiv. Class through Double Machine Learning.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-71), Feb/2021.
Proceedings of the 37th International Conference on Machine Learning (ICML), 2021.
(Acceptance rate = 21%)

70. Causal Identification with Matrix Equations.
Sanghack Lee, Elias Bareinboim
CausalAI Laboratory, Technical Report (R-70), Jun/2021.
Proc. of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS), 2021.
(Acceptance rate < 1% (oral))

69. Estimating Identifiable Causal Effects through Double Machine Learning.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-69), Jun/2020.
*Proceedings of the 35th AAAI Conference on Artificial Intelligence (*AAAI), 2021.
(Acceptance rate = 21%)

68. Bounding Causal Effects on Continuous Outcomes.
Junzhe Zhang and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-61).
*Proceedings of the 35th AAAI Conference on Artificial Intelligence (*AAAI), 2021.
(Acceptance rate = 21%)

67. General Transportability of Soft Interventions: Completeness Results.
Juan Correa and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-68).
Proc. of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate = 20%)

66. Causal Discovery from Soft Interventions with Unknown Targets: Characterization & Learning.
Amin Jaber, Murat Kocaoglu, Karthikeyan Shanmugam, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-67).
Proc. of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate = 20%)

65. Causal Imitation Learning with Unobserved Confounders.
Junzhe Zhang, Daniel Kumor, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-66).
Proc. of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate < 1% (oral))

64. An Introduction to Causal Reinforcement Learning.
Elias Bareinboim, Sanghack Lee, Junzhe Zhang (2020)
Columbia CausalAI Laboratory, Technical Report (R-65).

63. Can Humans Be Out of the Loop?
Junzhe Zhang and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-64), Jun/2020.
Proc. of the 1st Conference on Causal Learning and Reasoning (CLeaR), 2022.

62. Characterizing Optimal Mixed Policies: Where to Intervene, What to Observe.
Sanghack Lee and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-63).
Proc. of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate = 20%)

61. Learning Causal Effects via Empirical Risk Minimization.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-62).
Proc. of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS), 2020.
(Acceptance rate = 20%)

60. On Pearl's Hierarchy and the Foundations of Causal Inference.
Elias Bareinboim, Juan Correa, Duligur Ibeling, Thomas Icard
Columbia CausalAI Laboratory, Technical Report (R-60), 2020.
ACM Special Turing Series, Vol. "Probabilistic and Causal Inference: The Works of Judea Pearl".

59. Efficient and Doubly Robust Estimation of Causal Effects.
Yonghan Jung, Yuhao Wang, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-59), 2020, forthcoming.

58. Causal Effect Identifiability under Partial-Observability
Sanghack Lee and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-58), 2020.
Proceedings of the 37th International Conference on Machine Learning (ICML), 2020.
(Acceptance rate = 21.8%)

57. Designing Optimal Dynamic Treatment Regimes: A Causal RL Approach.
Junzhe Zhang and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-57), 2020.
Proceedings of the 37th International Conference on Machine Learning (ICML), 2020.
(Acceptance rate = 21.8%)

56. Efficient Identification in Linear Structural Causal Models with Auxiliary Cutsets.
Daniel Kumor, Carlos Cinelli, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-56), 2020.
Proceedings of the 37th International Conference on Machine Learning (ICML), 2020.
(Acceptance rate = 21.8%)

55. A Calculus For Stochastic Interventions: Causal Effect Identification and Surrogate Experiments.
Juan Correa and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-55), 2020.
*Proceedings of the 34th AAAI Conference on Artificial Intelligence (*AAAI), 2020.
(Acceptance rate = 20.6%)

54. Estimating Causal Effects Using Weighting-Based Estimators.
Yonghan Jung, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-54), 2020.
*Proceedings of the 34th AAAI Conference on Artificial Intelligence (*AAAI), 2020.
(Acceptance rate = 20.6%)

53. Generalized Transportability: Synthesis of Experiments from Heterogeneous Domains.
Sanghack Lee, Juan Correa, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-53), 2020.
*Proceedings of the 34th AAAI Conference on Artificial Intelligence (*AAAI), 2020.
(Acceptance rate = 20.6%)

52. Identifiability from a Combination of Observations and Experiments.
Sanghack Lee, Juan Correa, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-52), 2020.
*Proceedings of the 34th AAAI Conference on Artificial Intelligence (*AAAI), 2020.
(Acceptance rate = 20.6% [best paper award — sister's conference track])

51. Causal Inference and Data-Fusion in Econometrics.
P. Hunermund and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-51), 2019.
The Econometrics Journal, 2023 in press.

50. Identification of Conditional Causal Effects under Markov Equivalence.
Amin Jaber, Jiji Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-50), 2019.
Proc. of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), 2019.
(Acceptance rate = 2.5% (spotlight))

49. Efficient Identification in Linear Structural Causal Models with Instrumental Cutsets.
Daniel Kumor, Bryant Chen, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-49), 2019.
Proc. of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), 2019.
(Acceptance rate = 21%)

48. Near-Optimal Reinforcement Learning in Dynamic Treatment Regimes.
Junzhe Zhang and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-48), 2019.
Proc. of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), 2019.
(Acceptance rate = 21%)

47. Characterization and Learning of Causal Graphs with Latent Variables from Soft Interventions.
Murat Kocaoglu, Amin Jaber, Karthikeyan Shanmugam, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-47), 2019.
Proc. of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), 2019.
(Acceptance rate = 21%)

46. General Identifiability with Arbitrary Surrogate Experiments.
Sanglack Lee, Juan Correa, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-46), 2019.
Proceedings of the 35th Uncertainty in Artificial Intelligence (UAI), 2019.
Best Paper Award (1 out of 450 papers).
(Acceptance rate = 26%)

45. From Statistical Transportability to Estimating the Effects of Stochastic Interventions.
Juan Correa and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-45), 2019.
Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), 2019.
(Acceptance rate = 17.8%)

44. On Causal Identification under Markov Equivalence.
Amin Jaber, Jiji Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-44), 2019.
Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), 2019.
(Acceptance rate = 17.8%)

43. Adjustment Criteria for Generalizing Experimental Findings.
Juan Correa, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-43), 2019.
Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
(Acceptance rate = 22.5%)

42. Causal Identification under Markov Equivalence: Completeness Results.
Amin Jaber, Jiji Zhang, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-42), 2019.
Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
(Acceptance rate = 22.5%)

41. Sensitivity Analysis of Linear Structural Causal Models.
Carlos Cinelli, Daniel Kumor, Bryant Chen, Judea Pearl, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-41), 2019.
Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
(Acceptance rate = 22.5%)

40. On Structural Causal Bandits with Non-manipulable Variables.
Sanghack Lee and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-40), 2019.
*Proceedings of the 33th AAAI Conference on Artificial Intelligence (*AAAI), 2019.
(Acceptance rate = 16.2%)

39. Counterfactual Randomization: Rescuing Experimental Studies from Obscured Confounding.
Andrew Forney and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-39), 2019.
*Proceedings of the 33th AAAI Conference on Artificial Intelligence (*AAAI), 2019.
(Acceptance rate = 16.2%)

38. Identification of Causal Effects in the Presence of Selection Bias.
Juan Correa, Jin Tian, Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-38), 2019.
*Proceedings of the 33th AAAI Conference on Artificial Intelligence (*AAAI), 2019.
(Acceptance rate = 16.2%)

37. Equality of Opportunity in Classification: A Causal Approach.
Junzhe Zhang and Elias Bareinboim
Proc. of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS), 2018.
(Acceptance rate = 21%)

36. Structural Causal Bandits: Where to intervene?
Sanghack Lee and Elias Bareinboim
Proc. of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS), 2018.
(Acceptance rate = 21%)

35. Causal Identification under Markov Equivalence.
Amin Jaber, Jiji Zhang, Elias Bareinboim
Proceedings of the 34th Uncertainty in Artificial Intelligence (UAI), 2018.
Best Student Paper Award (1 out of 337 papers).
(Acceptance rate = 9% (plenary))

34. Non-Parametric Path Analysis in Structural Causal Models.
Junzhe Zhang and Elias Bareinboim
Proceedings of the 34th Uncertainty in Artificial Intelligence (UAI), 2018.
(Acceptance rate = 9% (plenary))

33. Budgeted Experimental Design for Causal Structural Learning.
Amiremad Ghassami, Saber Salehkaleybar, Negar Kiyavash, Elias Bareinboim
Proceedings of the 35th International Conference on Machine Learning (ICML), 2018.
(Acceptance rate = 25%)

32. A Graphical Criterion for Effect Identification in Equivalence Classes of Causal Diagrams.
Amin Jaber, Jiji Zhang, Elias Bareinboim
Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), 2018.
(Acceptance rate = 20%)

31. A note on "Generalizability of Study Results (Lesko et al., 2017)"
Judea Pearl and Elias Bareinboim
Epidemiology, v. 30(2), pp. 186-188, Mar/2019.

30. Fairness in Decision-Making — The Causal Explanation Formula.
Junzhe Zhang and Elias Bareinboim
*Proceedings of the 32nd AAAI Conference on Artificial Intelligence (*AAAI), 2018.
(Acceptance rate = 24%)

29. Generalized Adjustment under Confounding and Selection Biases.
Juan Correa, Jin Tian, Elias Bareinboim
*Proceedings of the 32nd AAAI Conference on Artificial Intelligence (*AAAI), 2018.
Outstanding Paper Honorable Mention (2 out of 3800 papers).
(Acceptance rate = 24%)

28. Experimental Design for Learning Causal Graphs with Latent Variables.
Murat Kocaoglu, Karthikeyan Shanmugam, Elias Bareinboim
Proc. of the 31st Annual Conference on Neural Information Processing Systems (NIPS), 2017.
(Acceptance rate = 21%)

27. Identification and Model Testing in Linear Structural Equation Models using Auxiliary Variables.
Bryant Chen, Daniel Kumor, Elias Bareinboim
Proceedings of the 34th International Conference on Machine Learning (ICML), 2017.
(Acceptance rate = 24%)

26. Counterfactual Data-Fusion for Online Reinforcement Learners.
Andrew Forney, Judea Pearl, Elias Bareinboim
Proceedings of the 34th International Conference on Machine Learning (ICML), 2017.
(Acceptance rate = 24%)

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

24. Causal Effect Identification by Adjustment under Confounding and Selection Biases.
Juan Correa and Elias Bareinboim
*Proceedings of the 31th AAAI Conference on Artificial Intelligence (*AAAI), 2017.
(Acceptance rate = 25%)

23. Markov Decision Processes with Unobserved Confounders: A Causal Approach.
Junzhe Zhang and Elias Bareinboim
Columbia CausalAI Laboratory, Technical Report (R-23), Dec/2016.

22. Incorporating Knowledge into Structural Equation Models using Auxiliary Variables.
Bryant Chen, Judea Pearl, Elias Bareinboim
Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), AAAI Press, pp. 3577-3583, 2016.
(Acceptance rate = 25%)

21. Causal Inference and the Data-Fusion Problem.
Elias Bareinboim and Judea Pearl
Proceedings of the National Academy of Sciences (PNAS), v. 113(27), 2016.

20. Comment on "Causal Inference using invariance prediction: identification and confidence intervals by Peters, Buhlmann and Meinshausen".
Elias Bareinboim
Journal of the Royal Statistical Society, Series B.

19. Bandits with Unobserved Confounders: A Causal Approach.
Elias Bareinboim, Andrew Forney, Judea Pearl
Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS), pp. 1342-1350, 2015.
(Acceptance rate = 21.9%)

18. Recovering Causal Effects From Selection Bias.
Elias Bareinboim and Jin Tian
*Proceedings of the 29th AAAI Conference on Artificial Intelligence (*AAAI), pp. 3475-3481, 2015.
(Acceptance rate = 26.7%)

17. External Validity: From do-calculus to Transportability across Populations.
Judea Pearl and Elias Bareinboim
Statistical Science, v. 29(4), pp. 579-595, 2014.

16. Transportability from Multiple Environments with Limited Experiments: Completeness Results.
Elias Bareinboim and Judea Pearl
Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS), pp. 280-288, 2014.
(Acceptance rate = 24.7%.)

15. Recovering from Selection Bias in Causal and Statistical Inference.
Elias Bareinboim, Jin Tian, Judea Pearl
*Proceedings of the 28th AAAI Conference on Artificial Intelligence (*AAAI), pp. 2410-2416, 2014.
Outstanding Paper Award (1 out of 1406 papers).
(Acceptance rate = 28%.)

14. A General Algorithm for Deciding Transportability of Experimental Results.
Elias Bareinboim and Judea Pearl
Journal of Causal Inference, v. 1(1), pp. 107-134, 2013.

13. Transportability from Multiple Environments with Limited Experiments.
Elias Bareinboim, Sanghack Lee, Vasant Honavar, Judea Pearl
Proceedings of the 27th Annual Conference on Neural Information Processing Systems (NIPS), pp. 136-144, 2013.
(Acceptance rate = 25%)

12. Causal Transportability with Limited Experiments.
Elias Bareinboim and Judea Pearl
Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI), pp. 95-101, 2013.
(Acceptance rate = 29%)

11. Meta-transportability of Causal Effects: A Formal Approach.
Elias Bareinboim and Judea Pearl
Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS), JMLR, pp. 135-143, 2013.
(Acceptance rate = 11% (plenary))

10. Causal Inference by Surrogate Experiments (or, *z*-Identifiability).
Elias Bareinboim and Judea Pearl
Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI), AUAI Press, pp. 113-120, 2012.
(Acceptance rate = 31%)

9. Transportability of Causal Effects: Completeness Results.
Elias Bareinboim and Judea Pearl
Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI), pp. 698-704, 2012.
(Acceptance rate = 26%)

8. Controlling Selection Bias in Causal Inference.
Elias Bareinboim and Judea Pearl
Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS), JMLR, *pp. 100-108, 2012.*
(Acceptance rate = 33%)

7. Local characterizations of Causal Bayesian Networks.
Elias Bareinboim, Carlos Brito, Judea Pearl
Lecture Notes in Artificial Intelligence, v. 7205, Springer-Verlag, pp. 1-17, 2012.

6. Transportability across studies: A formal approach.
   Judea Pearl and Elias Bareinboim
   Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI), pp. 247-254, 2011.
   (Acceptance rate = 24.8%)

5. External Validity and Transportability: A formal approach.
   Judea Pearl and Elias Bareinboim
   Proceedings of the Joint Statistical Meetings, American Statistical Association, pp. 157-171, 2011.

4. A statistical approach for analyzing marginal cases in shotgun proteomics.
   Paulo Carvalho, J. Fischer, J. Perales, J. Yates, V. C. Barbosa, Elias Bareinboim
   Bioinformatics, v. 27(2), 2011.

3. Local characterizations of Causal Bayesian Networks.
   Elias Bareinboim, Carlos Brito, Judea Pearl
   Proceedings of Graph Structures for Knowledge Representation and Reasoning — IJCAI, 2011.

2. Descents and nodal load in scale-free networks.
   Elias Bareinboim and Valmir C. Barbosa
   Physical Review E, v. 77(4), American Physical Society, 2008.

1. Grammatical inference applied to linguistic modeling of biological networks.
   Elias Bareinboim, Ana T. R. Vasconselos, Joao C. P. Silva
   E. Journal of Communication, Information & Innovation in Health, v.1, pp. 329-333, 2007.

Team / Mentoring

— Ph.D. students

- Hyun Chai Jeong (Fall/18-)
- Kevin Xia (Spring/20-)
- Kasra Jalaldoust (Fall/21-)
- Kai-Zhan Lee (Fall/21-)
- Mingxuan Li (Fall/21-)
- Yushu Pan (Fall/21-)
- Tara Anand (Fall/21-; co-advised w/ Prof. George Hripcsak @DBMI)
- Aurghya Maiti (Fall/22-)
- Arvind Raghavan (Spring/23-)
- Adiba Ejaz (Fall/23-)
- Hongshuo Yang (Fall/23-)
- Jeffrey Wu (Fall/24-)
- Shreyas Havaldar (Fall/24-)

— Postdoctoral Scholars

- Inwoo Hwang (Spring/25-)

— Alumni

- Juan David Correa (PhD: Fall/16-Summer/21)
  Title: A Computational Perspective of Causal Inference and the Data Fusion Problem
  Current: Assistant Professor, Universidad Autónoma de Manizales, Colombia.
- Daniel Kumor (PhD: Fall/16-Spring/21)
  Title: Effect Algorithms for Identification in Linear Systems and Imitation Learning
  Current: Researcher, Amazon.
- Amin Jaber (PhD: Fall/16-Fall/22)
  Title: Causal Identification in Equivalence Classes
  Current: Researcher, Synlico.
- Junzhe Zhang (PhD: Fall/16-Summer/23; Postdoc: Fall/23-Summer/24)
  Title: Towards Causal Reinforcement Learning
  Current: Assistant Professor, Syracuse University.
- Yonghan Jung (PhD: Fall/18-Summer/25)
  Title: Estimation of Causal Effects and Other Quantities
  Current: Assistant Professor, University of Illinois Urbana-Champaign (UIUC) (Fall-25).
- Alexis Bellot (Postdoc: Summer/21-Spring/22)
  Current: Researcher, Google DeepMind, UK.
- Adele Ribeiro (Postdoc: Fall/19-Summer/22)
  Current: Postdoctoral Scholar, Philipps-Universität Marburg.
- Sanghack Lee (Postdoc: Spring/18-Spring/21)
  Current: Assistant Professor, Seoul National University, Korea.
- Adam Li (Postdoc: Spring/22-Fall/24)
  Current: Senior Applied Scientist, Amazon.
- Drago Plecko (Postdoc: Fall/22-Summer/25)
  Current: Assistant Professor, University of California, Los Angeles (UCLA) (Fall-25).

— Visiting Scholars

- Inwoo Hwang (Fall/24)
- Luigi Gresele (Nov/24)
- Audrey Morgane Poinsot (Summer/24)
- Prof. Juan Correa (Summer/22, Summer/23)
- Christoffer Riis (Fall/22-Spring/23)
- Julius von Kügelgen (Jun/22)

- Drago Plecko (Fall/21-Spring/22)
- Prof. Jin Tian (Sabbatical; Fall/20-Summer/21)

— M.Sc. students

- Prateek Jain (Spring/23-Fall/24)
- Yusuf Efe (Fall/22-Spring/24)

— Undergraduate

- Noah Rouleau (Fall/15)
- Mahimna Kelkar (Fall/17)
  Current: PhD student, Cornell University

— PhD Committees

- Dustin Train (defense: 8/20)
  Advisor: David Blei
- Tyler Joseph (defense: 3/21)
  Advisor: Itsik Pe'er

Teaching

— At Columbia (instructor)

- CS 4775 (graduate), Causal Inference I: Spring/2020, Fall/2020, Fall/2021, Fall/2022, Fall/2023, Fall/2024.
- CS 4995 (graduate), Causal Inference II: Spring/2021, Spring/2022, Spring/2023, Spring/2025.
- CS 6995 (graduate), Causal Trustworthy AI: Fall/2023.

— At Purdue (instructor)

- CS 47100 (undergraduate), Artificial Intelligence, Spring/2017, Spring/2018.
- CS 57800 (graduate) Machine Learning, Fall/2015.
- CS 59000-AI (graduate), Artificial Intelligence, Fall/2016, Fall/2018.
- CS 59000-AML (graduate), Causal Inference / Advanced Machine Learning, Spring/2016, Fall/2017, Spring/2019.

— Before

- CS 262Z (graduate), Causal Inference, instructor with J. Pearl and J. Tian, UCLA, Spring/2013.
- CS 262Z (graduate), Causal Inference, teaching assistant, UCLA, Spring/2010, Spring/2011.
- MAB 525 (undergrad), Special Topics in Artificial Intelligence, instructor with J. C. P. Silva, Federal University of Rio de Janeiro (UFRJ), Spring/2007.

Tutorials & Short Courses

- "Causal Fairness Analysis" (with D. Plecko)
  European Conference on Artificial Intelligence (ECAI), Santiago de Compostela, Oct/2024.

- "Causal Fairness Analysis" (with D. Plecko)
  Association for Advancement of Artificial Intelligence (AAAI), Vancouver, Feb/2024.

- "Causal Fairness Analysis" (with D. Plecko)
  International Conference on Machine Learning (ICML), Baltimore, Jul/2022.

- "Causal Inference and the Data-Fusion Problem" (with A Ribeiro)
  Lisbon Machine Learning School (LxML), Jul/2022.

- "An Introduction to Causal Inference"
  Bellairs Invitational Workshop on Causal Inference & Representation Learning, Barbados, Mar/2022.

- "Causal Inference and the Data-Fusion Problem" (with A Ribeiro)
  Lisbon Machine Learning School (LxML), Jun/2021.

- "Causal Fairness Analysis" (with D. Plecko, J. Zhang)
  ACM Conference on Fairness, Accountability, and Transparency (FaccT), Mar/2021.

- "Causal Inference and the Data-Fusion Problem" (with A Ribeiro)
  Annual Deming Conference on Applied Statistics , NY, Dec/2020.

- "Causal Inference in the Health Sciences" (with M. Adibuzzaman, A. Ribeiro).
  American Medical Informatics Association Annual Symposium (AMIA), Nov/2020.

- "Causal Reinforcement Learning"
  International Conference on Machine Learning (ICML), Jul/2020.

- "Causal Reinforcement Learning" (with S. Lee, J. Zhang)
  International Joint Conference on Artificial Intelligence (IJCAI), Macau, China, Aug/2019.

- "An Introduction to Causal Inference"
  Machine Learning Research School (MLRS), Bangkok, Thailand, Aug/2019.

- "Causal Reinforcement Learning"
  Uncertainty in Artificial Intelligence (UAI), Tel Aviv, Israel, Jul/2019.

- "Causal Inference and the Data-Fusion Problem"
  International Conference on Autonomous Agents and Multi-agent Systems (AAMAS), Sao Paulo, Brazil, May/2017.

- "An Introduction to Causal Inference"
  West Coast Experiments Conference (Graphical Models in Economics), Los Angeles, CA, Apr/2017.

- "Causal Inference and the Data-Fusion Problem"
  Association for Advancement of Artificial Intelligence (AAAI), San Francisco, CA, Feb/2017.

- "Causal Inference and the Data-Fusion Problem"
  Department of Computing Science, University of Alberta, Edmonton, Canada, August/2016.

- "Causes and Counterfactuals: Concepts, principles, and tools" (with J. Pearl)
  Neural Information Processing (NIPS), Lake Tahoe, Nevada, December/2013.

- "Causality and Big Data"
  EMC² Summer School on Big Data, Rio de Janeiro, Brazil, February/2013.

- "An Introduction to Causal Inference"
  The Second IEEE Conference on Healthcare Informatics and Systems Biology (Analyzing Big Data For Healthcare and Biomedical Sciences), UCSD, La Jolla, California, September/2012.

Invited Talks, Lectures, Panels

- 2025 UChicago, Integrating Normative Considerations into Inequality Measurement, forthcoming.
- 2025 European Conference on Machine Learning (ECML PKDD), keynote, forthcoming.
- 2025 Workshop on Causal Neuro-symbolic AI, forthcoming.
- 2025 Annual Conference of Advanced Quantitative Methods and Analytics for Public Policy Support (AQMAPPS), forthcoming.
- 2025 Uncertainty in Artificial Intelligence Conference (UAI), keynote, forthcoming.
- 2025 RLC Workshop on Causal Reinforcement Learning, forthcoming.
- 2025 Conference on Causal Learning and Reasoning (CleaR).
- 2025 Duke Causation Group Workshop.
- 2025 Flatiron "Learning Meets Geometry, Graphs, and Networks" Workshop.
- 2025 University of Toronto, Data Science Institute.
- 2025 Harvard Causal Inference Working Group.
- 2025 Lemann Dialogues, Institute of Latin American Studies, Columbia edition.
- 2025 Columbia AI Summit.
- 2025 AFOSR AI roadmap workshop.
- 2025 SoCal Causal Inference Workshop (UCI).

- 2024 NeurIPS Workshop on Causality & Language Models.
- 2024 Brazilian Conference on Intelligent Systems (BRACIS).
- 2024 Yale AI in Medicine Interest Group.
- 2024 Joint Statistical Meetings.
- 2024 Workshop on Applied Algorithms for Machine Learning (Future of Computation).
- 2024 Columbia Engineering Alumni Reunion.
- 2024 Universitatea Politehnica Timișoara, Tech Talks.

- 2024 Workshop on Causal Discovery in Semiconductor Manufacturing (NSF/NIST).
- 2024 Columbia-Dream Sports Research Center (kickoff).

- 2023 Yale Research Initiative on Innovation and Scale Annual Meeting.
- 2023 Columbia Economics Department.
- 2023 NSF NAII (kickoff), The Artificial and Natural Intelligence Institute.
- 2023 National Academy of Sciences.
- 2023 IROS workshop on Causality for Robotics.
- 2023 Causality for Ethics and Society Workshop, LMU Munich.
- 2023 CVPR Workshop on Compositionality, Prompts and Causality.
- 2023 Bloomberg's Quant seminar series.
- 2023 UChicago Booth Econometrics and Statistics seminar.
- 2023 UIUC Causal Inference Workshop: Current Trends and the Future of Research.
- 2023 Vanderbilt Biostatistics seminar.
- 2023 Columbia Department of Biomedical Informatics.
- 2023 UMass Computer Science seminar.

- 2022 Boeing Aerospace & Autonomy Center.
- 2022 Bloomberg's Data Science Speaker Series.
- 2022 Oregon State University, AI seminar.
- 2022 ICLR Workshop on "Privacy, Accountability, Interpretability, Robustness, Reasoning on Structured Data" (PAIR2Struct).
- 2022 1st International Workshop on Interactive Causal Learning.
- 2022 MIT IDSS Distinguished Speaker Seminar.
- 2022 UC Berkeley/Simons Institute Workshop on "Learning from Interventions".

- 2021 NeurIPS Workshop on Algorithmic Fairness thr. the Lens of Causality & Robustness.
- 2021 NeurIPS Workshop on Causality in Sequential Decision Making.
- 2021 MIT-Harvard Economics seminar.
- 2021 ICML Algorithmic Recourse Workshop.
- 2021 ICAPS Workshop on Planning and Reinforcement Learning.
- 2021 JPMorgan Chase Faculty Research Meeting.
- 2021 OECD workshop on AI & the Productivity of Science.
- 2021 Society for Imprecise Probability (SIPTA) Annual Meeting Keynote.
- 2021 Inria Workshop "Leveraging Observational Data with Machine Learning".
- 2021 23rd Japanese Workshop on Information-Based Induction Sciences (IBIS).
- 2021 Seoul National University, Data Science Seminar.
- 2021 Columbia Data Science Institute.

- 2020 Machine Learning in Science and Engineering Conference (MLSE).
- 2020 Society for Epidemiologic Research (SER) Annual Meetings.
- 2020 AMIA Causal Inference from Observational Healthcare Data.
- 2020 CMU Machine Learning Department Seminar.
- 2020 MICCAI Causality in Medical Computing.
- 2020 AFOSR Understanding in the Human and the Machine Workshop.

- 2020 KDD Workshop on Causal Discovery.
- 2020 Microsoft Research Frontiers of Machine Learning.

- 2019 Max Planck Institute (Intelligent Systems), Tübingen, Germany.
- 2019 Mailman School of Public Health, Columbia University, NY.
- 2019 Data Council New York City, NY.
- 2019 INFORMS Annual Meeting, Seattle, WA.
- 2019 Stanford Graduate School of Business, CA.
- 2019 MIT-IBM Watson AI Lab - workshop on "Bridging causal inference, reinforcement learning and transfer learning (CRT)", MA.
- 2019 MIT workshop on "Graphical Models, Causality, Exchangeable Models, Graphons", MA.
- 2019 Technion - Israel Institute of Technology, Haifa, Israel.
- 2019 Hebrew University of Jerusalem, Jerusalem, Israel.
- 2019 Oberwolfach Research Institute for Mathematics, "Foundations and New Horizons for Causal Inference", Germany.
- 2019 Foundations of Data Science, Purdue University, Lafayette, IN.
- 2019 FDA / DIA Statistics Forum, Washington DC.
- 2019 Computer Science, Columbia University, NY.
- 2019 Harvard Medical School, Boston, MA.
- 2019 UIC Department of Information & Decision Sciences, Chicago, IL.
- 2019 DARPA CausalX-World Modelers' meeting, Los Angeles, CA.
- 2019 AI Roadmap: Learning and Robotics, Computing Community Consortium (CCC), CA.

- 2018 NeurIPS-18 Workshop "Causal Learning", Montreal, Canada.
- 2018 School of Medicine, Indiana University, Indianapolis, IN.
- 2018 NIH Division of Cancer Biology, National Cancer Institute (NCI), Rockville, MD.
- 2018 UAI-18 Workshop on Causal Inference, Monterey, CA.
- 2018 Adobe Research, San Jose, CA.
- 2018 RSS-18 Workshop "Causal Imitation in Robotics", Pittsburgh, PA.
- 2018 Atlantic Causal Inference Conference (ACIC), Pittsburgh, PA.
- 2018 TTI Vanguard Conference (Intelligence: Natural and Artificial), New York, NY.

- 2017 CVPR-17 Workshop "Functionality, Physics, Intentionality, and Causality", Honolulu, HI.
- 2017 Statistical Society of Canada Annual Meeting, Winnipeg, Canada.
- 2017 School of Engineering, University of São Paulo (USP), São Paulo, Brazil.
- 2017 Institute of Computing, University of Campinas (UNICAMP), Campinas, Brazil.
- 2017 Workshop on Causal Analysis in the Social Sciences, UCLA, CA.
- 2017 NSF Workshop: Advancing the Science of Transportation Demand Modeling, UC Berkeley, CA.
- 2017 Computer Science, University of Wisconsin, Madison, WI.
- 2017 Computer Science, ISI / University of Southern California (USC), CA.

- 2016 NeurIPS-16 Workshop "Inference and Learning of Hypothetical and Counterfactual Interventions in Complex Systems", Barcelona, Spain.
- 2016 AAAI-16 Fall Symposium on Accelerating Science: A Grand Challenge for AI, Arlington, VA.
- 2016 Department of Public Health Sciences, University of Chicago, Chicago.

- 2016 54th Allerton Conference on Communication, Control, and Computing, UIUC, IL.
- 2016 Department of Computing Science, University of Alberta, Edmonton, Canada.
- 2016 International Conference on Thinking (ICT), Providence, RI.
- 2016 Joint Statistical Meetings (JSM), Chicago, IL.
- 2016 Workshop on Statistical Causal Inference and its Applications to Genetics, Centre de Recherches Mathématiques (CRM), Montreal, Canada.
- 2016 Frontiers of Engineering Symposium (US-JP), National Academy of Engineering (NAE), CA.
- 2016 Max Planck Institute (Empirical Inference Dept.), Tübingen, Germany.
- 2016 Department of Computer Science and Mathematics, University of Passau, Germany.
- 2016 Munich Workshop on Causal Inference and Information Theory (MCI), Munich, Germany.
- 2016 Statistics Colloquium, Purdue University, West Lafayette, IN.

- 2015 Computer Science, Purdue University, West Lafayette, Indiana.
- 2015 Biostatistics and Computer Science, Johns Hopkins University, Baltimore, Maryland.
- 2015 Computer Science Division, University of California, Berkeley, California.
- 2015 Department of Computer Science, University of Southern California (USC), CA.
- 2015 School of Information and Computer Science, University of California, Irvine, CA.
- 2015 Department of Computer Science, Cornell University, New York.
- 2015 Department of Statistics, Stanford University, California.
- 2015 60th World Congress of Statistics, International Statistics Institute (ISI), Brazil.

- 2014 Department of Economics, University of Chicago, Chicago.
- 2014 Kyoto International Conference on Modern Statistics, Kyoto.
- 2014 International Workshop on Causal Inference and its related topics, Tokyo.
- 2014 ACM-SIGKDD-14 Workshop on Discovery Informatics, New York.
- 2014 UAI-14 Workshop on Causality: Learning and Prediction, Quebec City, Canada.
- 2014 NICTA, Sydney, Australia.
- 2014 Institute of Mathematical Statistics (IMS) Annual Meeting, Sydney, Australia.
- 2014 MURI, Office of Naval Research (ONR), UCLA, Los Angeles, California.
- 2014 Atlantic Causal Inference Conference, Brown University, Providence, RI.
- 2014 Joint Mathematics Meetings, American Mathematical Society, Baltimore, Maryland.

- 2013 NeurIPS-13 Workshop "Causality: Large-scale Experimental Design", Lake Tahoe, NV.
- 2013 MURI, Office of Naval Research (ONR), UCLA, Los Angeles, California.

- 2012 Graduate School of Engineering, Federal University of Rio de Janeiro (UFRJ), Brazil.
- 2012 Computer Science Colloquium, Federal University of Rio de Janeiro (UFRJ), Brazil.
- 2012 MURI, Office of Naval Research (ONR), UCLA, Los Angeles, California.

- 2011 International Workshop on Mining Multiple Information Sources, International Conference on Data Mining (ICDM), Vancouver, Canada.
- 2011 58th World Congress of Statistics, International Statistics Institute (ISI), Dublin.
- 2011 DERI/National University of Ireland (NUI), Galway, Ireland

Funding (Bareinboim's share > $8M)

- Columbia Center for AI and Responsible Financial Innovation (CAIRFI), PI
  Title: Trustworthy Lending: Explaining Group Differences, 7/2025-6/2026.
  Amount: $100,000 (=100% of total).

- NSF CISE: Large: Causal Foundations of Decision Making and Learning, PI
  Title: Causal Decision-Making, 10/2023-09/2028.
  Amount: $1,672,312 (=33% of total).

- Defense Advanced Research Projects Agency (DARPA), Young Faculty Award, PI
  Title: Causal Reinforcement Learning, 9/2023 - 8/2025.
  Amount: $500,000 (=100% of total).

- Columbia-Amazon Center of AI Technology, PI
  Title: Algorithmic Fairness through a Causal Lens, 7/2023 - 6/2024.
  Amount: $100,000 (=100% of total).

- NSF Eager, Robust Intelligence/IIS, PI
  Title: Causal Decision-Making, 9/2022-8/2023.
  Amount: $150,000 (=50% of total).

- Computing Research Association, PI
  Title: Computing Innovation Fellows, 1/2022-5/2024.
  Amount: $321,288 (=100% of total).

- Air Force Office of Scientific Research (AFOSR), PI
  Title: Causal Reinforcement Learning: Discovery and Decision Making, 09/2022-08/2025.
  Amount: $825,000 (=100% of total).

- Office of Naval Research (ONR), Young Investigator Program (YIP), PI
  Title: Causal Reinforcement Learning: Theory, Algorithms, & Applications, 05/2022-04/2025.
  Amount: $510,000 (=100% of total).

- The Alfred P. Sloan Foundation Award, PI
  Title: The Mathematics of Fair Decision-Making, 08/2021-07/2023.
  Amount: $564,726 (=100% of total).

- Amazon, Research Award (gift)
  Title: Approximate Causal Inference, cycle 2021; awarded 2022.
  Amount: $140,000 (=100% of total).

- JP Morgan, Research Award (gift)
  Title: Causal Reinforcement Learning, cycle 2021.
  Amount: $120,000 (=100% of total).

- Carnegie Mellon University, Software Engineering Institute (SEI)
  Title: Investigating the Maturation of Determining the Limits of AI Robustness, 09/2021-08/2022.
  Amount: $200,000 (=20% of total).

- Department of Energy (DoE), ASCR (thr. UCSD); CU PI: Gentine
  Title: Discovering Physically Meaningful Structures from Climate Extreme Data, 09/2021-08/2024.
  Amount: $300,000 (=25% of total).

- Columbia University, SIRS/STAR Program, PI
  Title: Causal Data Science: Towards an Accelerated Process of Cancer Translation Research, (cycles: 2021-22 and 2022-23). Co-PI: Prof. Anil Rustgi (Medical School).
  Amount: $170,000 (=100% of total).

- Amazon, Research Award (gift)
  Title: Off-policy Evaluation through Causal Models, cycle 2020; awarded 2021.
  Amount: $90,000 (=100% of total).

- NSF, Robust Intelligent/IIS, PI
  Title: Towards Causal Fair Decision-Making, 04/2021 - 03/2023.
  Amount: $270,000 (=35% of total).

- Columbia-Amazon Center of AI Technology, PI
  Title: Counterfactual Reinforcement Learning for Personalized Decision-Making, 1/2021 - 12/2021.
  Amount: $150,000 (=100% of total).

- NIH, R01, PI: George Hripcsak
  Title: NLM:Discovering and Applying Knowledge in Clinical Databases. 09/2020 - 08/2021.
  Amount: $74,485 (=5% of total)

- NSF, CAREER, PI
  Title: Approximate Causal Inference, 04/2018 - 03/2023.
  Amount: $499,712 (=100% of total).

- NSF, Robust Intelligence, Medium, PI
  Title: Causal Inference: Identification, Learning, and Decision-Making, 10/2017 - 09/2020.
  Amount: $536,515 (=50% of total).

- Purdue, Integrative Data Science Initiative, PI
  Title: Causally-driven Healthcare Science, 06/2018 - 05/2020.
  Amount: $200,000 (=75% of total).

- Adobe, Data Science Research Award (gift)
  Title: Optimal Decision-making under Causal Constraints, 2018.
  Amount: $50,000 (100% of total).

- IBM, Open Collaborative Research Award (gift)
  Title: Machine Learning and Causal Inference, 2017.
  Amount: $50,000 (100% of total).

- DARPA, Fundamental Limits of Learning (FunLol), co-PI
  Title: Fundamental Limits of Learning Concepts and Models for Complex Systems, 10/2016-12/2017.
  Amount: $125,000 (=16.6% of total).

Community Service

- Editor-in-Chief, Journal of Causal Inference (JCI), 2023-now.
- Action Editor, Journal of Machine Learning Research (JMLR), 2022-now.

- Tutorial Chair, Neural Information Processing Systems (NeurIPS), 2025.
- Reviewer, Israel National Science Foundation, 2023.
- Editorial Board, Journal of Causal Inference (JCI), 2017-2023.
- Chair (with J. Pearl, B. Schölkopf, Y. Bengio, T. Sejnowski), NeurIPS-21 workshop, "WHY-21 Causal Inference and Machine Learning: Why now?", 2021.
- Editor (w/ Mark V. D. Laan), Journal of Causal Inference Special issue on "Integrating Observational Studies with Randomized Trials", 2021-2022.
- Reviewer, National Science Foundation (NSF), area: IIS, 2019, 2020, 2021, 2023.
- Co-chair (w/ B. Schölkopf, K. Zhang, B. Huang et al), NeurIPS Workshop on Causal Discovery, 2020.
- Chair (w/ J. Pearl, B. Schölkopf, C. Szepesvari, S. Mahadevan, P. Tadepalli), AAAI-SS-19, "WHY-19 Beyond Curve Fitting: Causation, Counterfactuals, and Imagination-based AI", 2019.
- Chair (with K. Zhang, C. Uhler, J. Zhang, D. Janzing), 7th UAI Causality Workshop, 2017.
- Co-chair (with K. Zhang, J. Li, L. Liu), KDD Workshop on Causal Discovery, 2016.
- Co-chair (with F. Eberhardt, R. Silva, J. Mooij, M. Maathuis), UAI Causality Workshop, 2016.
- Guest Editor (with J. Pearl, B. Schölkopf, K. Zhang, J. Li), Special Issue on Causality, ACM Transactions on Intelligent Systems and Technology (TIST), 2015.
- Co-chair (with B. Schölkopf, K. Zhang, J. Zhang), ICML 2014 Workshop on Causal Modeling and Machine Learning, 2014.
- Reviewer, National Science Foundation (NSF), area: Methodology, Measurement, and Statistics, 2014.
- Area Chair / Senior PC-Conferences (* Senior AC):
  - 2025: NeurIPS*, ICML*.
  - 2024: NeurIPS*, ICML*.
  - 2023: NeurIPS*, AAAI*, ICLR.
  - 2022: NeurIPS*, ICML, AAAI, AIStats, ICLR[1], CLeaR.
  - 2021: NeurIPS*, ICML, AAAI, AIStats, ICLR, UAI, IJCAI.
  - 2020: NeurIPS, ICML, AAAI, AIStats, UAI, IJCAI.
  - 2019: NeurIPS, AAAI.

[1] Selected as "Highlighted Area Chair".

- Program Committee-Conferences:
  - 2020: FODS (Foundations of Data Science).
  - 2019: UAI, IJCAI, ICML.
  - 2018: NeurIPS, UAI, AAAI, IJCAI, ICML.
  - 2017: NeurIPS, UAI, AAAI, AISTATS.
  - 2016: NeurIPS, UAI, AAAI, IJCAI, ECAI.
  - 2015: NeurIPS, UAI, AAAI, AISTATS, UAI-Causality.
  - 2014: UAI, ICML, AISTATS, KDD-DI.
  - 2013: UAI, AAAI, IJCAI, ICML, NeurIPS-Causality, IEEE-BigData, UAI-Causality.
  - 2012: UAI, ICML.
  - 2011: NeurIPS, UAI, IJCAI, ICDM-MMIS.
  - 2010: KR (rev).
- Reviewer-Journals:
  - 2023: Statistics in Medicine.
  - 2022: J. of Machine Learning Research (JMLR), Statistical Science, Journal of the ACM (JACM).
  - 2021: J. of Machine Learning Research (JMLR), Statistical Science, PloS Medicine, Epidemiology, Am. J. of Epidemiology.
  - 2020: J. of Machine Learning Research (JMLR), Statistical Science, Statistics in Medicine.
  - 2019: J. of Machine Learning Research (JMLR), Statistical Science, Statistics in Medicine.
  - 2018: J. of Machine Learning Research (JMLR), Artificial Intelligence Journal (AIJ), Statistics in Medicine, Peer J (Computer Science).
  - 2017: J. of Machine Learning Research (JMLR), J. of Causal Inference.
  - 2016: Biometrika, Bayesian Analysis, J. Causal Inference, Epidemiology, Behaviormetrika.
  - 2015: Artificial Intelligence Journal (AIJ), Biometrics, J. of Causal Inference, Epidemiology.
  - 2014: Statistical Science, The British Journal for the Philosophy of Science, Annals of Applied Statistics.
  - 2013: J. of Machine Learning Research (JMLR), Scandinavian Journal of Statistics, Annals of Applied Statistics, J. of Causal Inference, Statistics in Medicine, Statistics.
  - 2012: J. of Machine Learning Research (JMLR), IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Statistics in Medicine, Bioinformatics, J. of Proteome Research.
  - 2011: J. of Causal Inference.
  - 2009: J. of Proteomics, Bioinformatics, Physica A.

Departmental & University Service

- Member, Interdisciplinary Recruitment Committee on Causal Inference, Columbia University, 2025.
- Member, Task Force for AI Initiative, School of Engineering, Columbia University, since 2021.
- Member, Data Science in Health Initiative (DASHI), Data Science Institute, Columbia University, since 2021.
- Columbia University (CS Department):
  - Member, Graduate Admissions Committee, since Fall/2020.
  - Member, Student Nominations Committee, since Fall/2019.

- Purdue University (CS Department):
  - Member, Graduate Committee, Fall/2017-Spring/2019.
  - Member, Graduate Admissions Committee, cycle: Fall/2016, Fall/2017.
- UCLA (CS Department):
  - Reviewer, Graduate Admissions Committee, 2013-2014;
  - Mentor for 3 PhD students, 2010-2013.

Software & Infrastructure

- LLM Observatory (link: https://llm-observatory.org/)
  Platform for monitoring, understanding, and mapping large language models
- Causal Fusion (link: https://causalfusion.net/)
  Platform for data scientists to perform causal analysis.
- Other packages: http://github.com/CausalAILab/

Industrial Experience

- Software Engineer, Intern (Systems/Data Mining), Google, Mountain View/CA, USA, Summer 2009.
- Software Engineer, Programare Software Factory, Brazil, Feb/2008 – Aug/2008.
- CTO and Co-Founder, Linux Solutions Ltda, Brazil, 1999 – 2004.

Professional Associations

- Association for the Advancement of Artificial Intelligence (AAAI), since 2011.
- Association for Computing Machinery (ACM), since 2011.
- Brazilian Computer Society (SBC), since 2004.

Media coverage

- MIT Technology Review (featured), "What AI still can't do", Feb/2020 (link).
- NewScientist (featured), "Correlation or causation? Mathematics can finally give us an answer", Apr/2020 (link).
- Communication of ACM (featured),"Solving for Why", Vol. 65(2), p. 11-13, Feb/2022 (link).
- The State of Sao Paulo (in Portuguese), "Brazilians in AI: Elias Bareinboim", Oct/2023 (link).