

## Alexander N. D’Amour

**Office**  
Google Cambridge  
355 Main St, Floor 5  
Cambridge, MA 02142

alexdamour@google.com  
(650)793-4899  
alexdamour.com

**Home**  
104 Churchill Ave  
Arlington, MA 02476

### Appointments

- 2023–Present **Staff Research Scientist, Google DeepMind, Cambridge, MA.**
- 2020–2023 **Senior Research Scientist, Google Brain, Cambridge, MA.**
- 2018–2020 **Research Scientist, Google Brain, Cambridge, MA.**
- 2016–2018 **Neyman Visiting Assistant Professor of Statistics, UC Berkeley, Berkeley, CA.**

### Education

- 2010–2017 **PhD in Statistics, Harvard University, Cambridge, MA.**  
Dissertation: Superpopulation Generalization in Social Network Analysis.
- 2007–2008 **SM in Applied Mathematics, Harvard University, Cambridge, MA.**
- 2004–2008 **AB in Applied Mathematics, Harvard University, Cambridge, MA.**

### Research

#### Selected Papers and Preprints

- 2024 Zihao Wang, Chirag Nagpal, Jonathan Berant, Jacob Eisenstein, Alex D’Amour, Sanmi Koyejo, and Victor Veitch. “Transforming and Combining Rewards for Aligning Large Language Models.” *arXiv preprint arXiv:2402.00742*, 2024.
- 2024 Katherine Tsai, Stephen R Pfohl, Olawale Salaudeen, Nicole Chiou, Matt Kusner, Alexander D’Amour, Sanmi Koyejo, and Arthur Gretton. “Proxy Methods for Domain Adaptation.” In “International Conference on Artificial Intelligence and Statistics,” pages 3961–3969. PMLR, 2024.
- 2024 Qingyao Sun, Kevin P Murphy, Sayna Ebrahimi, and Alexander D’Amour. “Beyond Invariance: Test-Time Label-Shift Adaptation for Addressing” Spurious” Correlations.” *Advances in Neural Information Processing Systems*, 36, 2024.
- 2023 Jiajing Zheng, Jiayi Wu, Alexander D’Amour, and Alexander Franks. “Sensitivity to Unobserved Confounding in Studies with Factor-structured Outcomes.” *Journal of the American Statistical Association*, pages 1–12, 2023.
- 2023 Jacob Eisenstein, Chirag Nagpal, Alekh Agarwal, Ahmad Beirami, Alex D’Amour, DJ Dvijotham, Adam Fisch, Katherine Heller, Stephen Pfohl, Deepak Ramachandran, et al. “Helping or herding? reward model ensembles mitigate but do not eliminate reward hacking.” *arXiv preprint arXiv:2312.09244*, 2023.
- 2023 Maggie Makar and Alexander D’Amour. “Fairness and robustness in anti-causal prediction.” *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
- 2022 Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma, Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D Sculley. “Underspecification presents challenges for credibility in modern machine learning.” *Journal of Machine Learning Research*, 2022.
- 2022 Ibrahim Alabdulmohsin, Nicole Chiou, Alexander D’Amour, Arthur Gretton, Sanmi Koyejo, Matt J Kusner, Stephen R Pfohl, Olawale Salaudeen, Jessica Schrouff, and Katherine Tsai. “Adapting to Latent Subgroup Shifts via Concepts and Proxies.” In “The 26th International Conference on Artificial Intelligence and Statistics,” 2023.
- 2022 Jessica Schrouff, Natalie Harris, Oluwasanmi O Koyejo, Ibrahim Alabdulmohsin, Eva Schneider, Krista Opsahl-Ong, Alexander Brown, Subhrajit Roy, Diana Mincu, Chrsitina Chen, Awa Dieng, Yuan Liu, Vivek Natarajan, Alan Karthikesalingam, Katherine A Heller, Silvia Chiappa, and Alexander D’Amour. “Diagnosing failures of fairness transfer across distribution shift in real-world medical settings.” In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, “Advances in Neural Information Processing Systems,” 2022.
- 2022 Andrew C Miller, Lauren A Hannah, Joseph Futoma, Nicholas J Foti, Emily B Fox, Alexander D’Amour, Mark Sandler, Rif A Saurous, and Joseph A Lewnard. “Statistical deconvolution for inference of infection time series.” *Epidemiology (Cambridge, Mass.)*, 33(4):470, 2022.
- 2022 Jiajing Zheng, Jiayi Wu, Alexander D’Amour, and Alexander Franks. “Sensitivity to Unobserved Confounding in Studies with Factor-structured Outcomes.” *arXiv preprint arXiv:2208.06552*, 2022.
- 2022 Jiajing Zheng, Alexander D’Amour, and Alexander Franks. “Bayesian Inference and Partial Identification in Multi-Treatment Causal Inference with Unobserved Confounding.” In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, “Proceedings of The 25th International Conference on Artificial Intelligence and Statistics,” volume 151 of *Proceedings of Machine Learning Research*, pages 3608–3626. PMLR, 28–30 Mar 2022.

- 2022 Maggie Makar, Ben Packer, Dan Moldovan, Davis Blalock, Yoni Halpern, and Alexander D’Amour. “Causally motivated shortcut removal using auxiliary labels.” In “International Conference on Artificial Intelligence and Statistics,” pages 739–766. PMLR, 2022.
- 2021 Victor Veitch, Alexander D’Amour, Steve Yadlowsky, and Jacob Eisenstein. “Counterfactual invariance to spurious correlations in text classification.” *Advances in neural information processing systems*, 34:16196–16208, 2021.
- 2021 Shira Mitchell, Eric Potash, Solon Barocas, Alexander D’Amour, and Kristian Lum. “Algorithmic fairness: Choices, assumptions, and definitions.” *Annual Review of Statistics and Its Application*, 8:141–163, 2021.
- 2021 Alexander D’Amour, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. “Overlap in observational studies with high-dimensional covariates.” *Journal of Econometrics*, 221(2):644–654, 2021.
- 2021 Alexander D’Amour. “Revisiting Rashomon: A Comment on ‘The Two Cultures’.” *Observational Studies*, 7(1):59–63, 2021.
- 2021 Steve Yadlowsky, Taedong Yun, Cory Y McLean, and Alexander D’Amour. “SLOE: A faster method for statistical inference in high-dimensional logistic regression.” *Advances in Neural Information Processing Systems*, 34:29517–29528, 2021.
- 2020 Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. “Fairness is not static: deeper understanding of long term fairness via simulation studies.” In “Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency,” pages 525–534. 2020.
- 2020 David Madras, James Atwood, and Alexander D’Amour. “Detecting Extrapolation with Local Ensembles.” In “International Conference on Learning Representations,” 2020.
- 2019 Alexander D’Amour. “Comment: Reflections on the deconfounder.” *Journal of the American Statistical Association*, 114(528):1597–1601, 2019.
- 2019 Alexander D’Amour. “On multi-cause causal inference with unobserved confounding: Counterexamples, impossibility, and alternatives.” In “The 22nd International Conference on Artificial Intelligence and Statistics,” pages 3478–3486. PMLR, 2019.
- 2019 Alexander M Franks, Alexander D’Amour, and Avi Feller. “Flexible sensitivity analysis for observational studies without observable implications.” *Journal of the American Statistical Association*, 2019.
- 2017 A. Miller, N. Foti, **A. D’Amour**, and R. Adams. “Reducing Reparameterization Gradient Variance.” In “Advances in Neural Information Processing Systems 30,” 2017. <https://arxiv.org/abs/1705.07880>.
- 2017 A. Franks, **A. D’Amour**, D. Cervone, and L. Bornn. “Meta-Analytics: Tools for Understanding the Statistical Properties of Sports Metrics.” *Journal of Quantitative Analysis in Sports*, 2017.
- 2016 D. Cervone, **A. D’Amour**, L. Bornn, and K. Goldsberry. “A Multiresolution Stochastic Process Model for Predicting Basketball Possession Outcomes.” *Journal of the American Statistical Association*, 111(514):585–599, 2016.
- 2014 D. Cervone, **A. D’Amour**, L. Bornn, and K. Goldsberry. “POINTWISE: Predicting Points and Valuing Decisions in Real Time with NBA Optical Tracking Data.” *MIT Sloan Sports Analytics Conference*, 2014.
- 2014 G. Li, R. Lai, **A. D’Amour**, D. Doolin, Y. Sun, V. Torvik, A. Yu, and L. Fleming. “Disambiguation and Co-authorship Networks of the U.S. Patent Inventor Database.” *Research Policy*, 43(6):941–955, 2014.
- 2012 M. Lipsitch, O. Abdullani, **A. D’Amour**, W. Xie, D. Weinberger, E. Tchetgen, and J. Scott. “Estimating Rates of Carriage Acquisition and Clearance and Competitive Ability for Pneumococcal Serotypes in Kenya With a Markov Transition Model.” *Epidemiology*, 23(4):510–519, 2012.
- 2008 R. Acharya, A. Ahmed, **A. D’Amour**, H. Lu, C. Morris, B. Oglevee, A. Peterson, and R. Swift. “Improving Major League Baseball Park Factor Estimates.” *Journal of Quantitative Analysis in Sports*, 4(2), 2008.

#### Invited Talks (through 2018)

- 2018 “Model-Free Statistical Assessment of Population Overlap.” Atlantic Causal Inference Conference at CMU., May 2018. Slides available at: <https://goo.gl/1cdXB9>.
- 2018 “Overlap in High Dimensions.” European Causal Inference Meeting in Florence, IT., April 2018. **Best early career talk award**. Slides available at: <https://goo.gl/cK34mc>.
- 2017 “Overlap in Observational Studies with High-Dimensional Covariates.” Berkeley Division of Biostatistics Seminar, October 2017.
- 2017 “Advances in Basketball Analytics Using Player-Tracking Data.” Invited talk at Berkeley Consortium for Data Analytics in Risk, October 2017.
- 2017 “Overlap in High Dimensions.” Invited talk at Atlantic Causal Inference Conference at UNC, May 2017.
- 2017 “Prediction is Not Enough: Designing decision-support statistics for causal inference.” Invited talk at Clarify Health Solutions in San Francisco, CA, May 2017.
- 2016 “Advances in Basketball Analytics Using Player-Tracking Data.” Invited talk at Boston ML Meetup, July 2016.
- 2015 “Prediction is Not Enough: Designing decision-support statistics for causal inference and attribution.” Invited talk at Lumos Labs in San Francisco, CA, October 2015.
- 2014 “POINTWISE: Predicting Points and Valuing Decisions in Real Time with NBA Optical Tracking Data.” Research Paper Competition Finalist Presentation at MIT Sloan Sports Analytics Conference 2014, February 2014.
- 2013 “Sparse is Different: Covariate Effect Estimation on Sparse Networks.” Invited talk presented in the PED Seminar series at MIT Lincoln Laboratory in Lexington, MA, November 2013.
- 2013 “Multi-Concept Item Response Theory.” Talk given at Knewton, Inc. in New York, NY, April 2013.
- 2012 “Analysis of Sparsity: An Observation Model for Interaction Data.” Invited talk presented to the ISR and Tactical Systems Division of MIT Lincoln Laboratory in Lexington, MA, June 2012.
- 2009 “Dataverse Network Patent Network Database Project.” Invited talk and workshop given at the University of Trento X Summer School in Networks and Innovation in Trento, IT, July 2009.

## Media

- 2020 Will Douglas Heaven. "The way we train AI is fundamentally flawed." MIT Technology Review. [technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-health-nlp-computer-vision/](https://technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-health-nlp-computer-vision/).
- 2014 K. Goldsberry. "Behind DataBall: A Discussion on the Methodology of Expected Possession Value." Grantland. <http://grantland.com/the-triangle/behind-databall-a-discussion-on-the-methodology-of-expected-points-value/>, 2014.
- 2014 K. Goldsberry. "DataBall." Grantland. <http://grantland.com/features/expected-value-possession-nba-analytics/>, 2014.
- 2013 C. Duffy. "Bayesian." Interview for radio show "You're the Expert", recorded at Oberon Theater, Cambridge, MA, January 2013.

## Teaching

### Awards

- Fall 2014 **2014 Pickard Memorial Teaching Fellow.** Departmental award for sustained excellence in teaching.
- Fall 2011, 2012 Spring 2013, 2014 **Harvard University Certificate of Distinction in Teaching.** University commendation for receiving excellent student evaluations.

### Courses

- 2017-2018 **STAT 278B: Causal Inference Reading Group.**  
Co-instructor of reading group studying contemporary issues in causal inference. Topics have included high-dimensional causal inference, interference, and optimization approaches.
- Spring 2017, Fall 2017 **STAT 153: Introduction to Timeseries.**  
Upper-division undergraduate course introducing time-domain and frequency-domain approaches to timeseries analysis.
- Summer 2017 **STAT 199: Supervised Independent Study and Research.**  
Summer independent study course for several undergraduates interested in experimental and observational causal inference.
- Fall 2016 **STAT 88: Probability and Mathematical Statistics in Data Science.**  
Introductory course to provide technical depth for students taking Foundations of Data Science.
- Spring 2014 **Teaching Fellow for Statistics 225: Spatial Statistics.**  
Graduate course introducing core topics in spatio-temporal statistical methods, covering both theoretical approaches and computational methods.
- Fall 2013, 2014 **Teaching Fellow for Computer Science 109/Statistics 121: Data Science.**  
Introductory course concerning the diverse set of skills necessary for modern data science. Focus on prediction, visualization, Bayesian methods, and empirical model validation.
- Spring 2013 **Teaching Fellow for Statistics 221: Statistical Computation and Visualization.**  
Graduate course on computational methods and visualization for moderns statistical problems. Topics include building probabilistic models, EM algorithm, MCMC techniques, and visualization using d3.
- Fall 2011, 2012 **Teaching Fellow for Statistics 220: Bayesian Data Analysis.**  
Core graduate course on Bayesian approaches to model building, model selection, inference, and computation.
- Spring 2012 **Teaching Fellow for Statistics 107: Financial Statistics.**  
Intro-level course applying basic statistical ideas to trading strategies in financial markets.
- Summer 2011 **Teaching Fellow for Statistics S100: Introduction to Statistics.**  
Intro-level summer course for college and advanced high school students.

## Professional Experience

- 2017-2018 **Advisor at Clarify Health Solutions, Inc.**  
Advise the Data Science team on goals and implementation of key analytical components in Clarify's care management and patient engagement platform.
- 2015-2020 **Founding Partner at XY, LLC.**  
Sports analytics consultancy, with a focus on player-tracking data. Clients include Philadelphia 76ers, LA Dodgers, Arsenal, San Antonio Spurs, Dallas Mavericks.
- 2015-2018 **Founding Partner at Damyata, LLC.**

- Data science consultancy. Founded with two partners with deep experience as data scientists in the tech space, Damyata focuses on enabling teams of any size to effectively exploit data. We offer advice and solutions over the full stack of technical, organizational, and analytical challenges that come with ramping up data science operations. Clients include Blue Apron and Demand Signals.
- 2013–2017 **Data Science Consultant at FirstAccess, Inc.**  
 Team lead on statistical modeling and experimental design for construction of real-time credit scores for microlenders and development of lending strategies that reduce default rates while expanding access to credit.
- 2014–2016 **Quantitative Analytics Consultant at Legendary Pictures.**  
 Advisor for quant analytics unit for Legendary Pictures. Developed a market research tool to predict audience response to potential film properties in the pre-greenlighting stage of production. Currently supervising a project to estimate causal effects of marketing engagement with the intention of producing publications.
- 2012–2014 **Statistical Consultant at Knewton, Inc.**  
 Developed statistical tools for an EdTech company’s core adaptive teaching platform. Simultaneously assesses student proficiency and measures question difficulty across multiple concept areas.
- 2009–2011 **Statistical Consultant at Harvard School of Public Health.**  
 Implemented statistical models of infection dynamics for a variety of contexts. Work featured in an *Epidemiology* publication and a software package prepared for the CDC.
- 2009–2010 **Statistical Programmer at Harvard Institute for Quantitative Social Science.**  
 Lead architect on NSF-funded project to construct and distribute the full collaboration network of inventors who have held patents in the United States since 1975.
- 2008 **Senior Analyst at Lehman Brothers, Fixed Income Analytics.**  
 Implemented financial models in the Fixed Income Capital Markets Division.