

# Alan E. Mishler

New York, NY • [contact@alanmishler.com](mailto:contact@alanmishler.com) • [github.com/amishler](https://github.com/amishler)  
Google Scholar: [bit.ly/g\\_s\\_alanmishler](https://bit.ly/g_s_alanmishler) • [www.linkedin.com/in/alanmishler](https://www.linkedin.com/in/alanmishler)

---

## Research interests

Causal inference · Sequential inference · Algorithmic fairness · Uncertainty quantification

---

## Education

08/2021	PhD in Statistics, Carnegie Mellon University
05/2017	MS in Statistics, Carnegie Mellon University
08/2016	BS in Mathematics, University of Maryland
04/2009	BA in Linguistics, <i>summa cum laude</i> , University of Michigan

---

## Employment

06/21–present	<b>Vice President/AI Research Lead (02/24–present), Senior Research Associate (06/21–02/24):</b> JP Morgan AI Research Designing and leading basic research in causal inference, sequential decision making, algorithmic fairness, and other AI-related topics.
06/19–08/19	<b>Data Scientist Intern:</b> Google Developed improved methodology to estimate ads lift/incrementality using combined experimental and observational data, with both empirical and theoretical results.
06/18–08/18	<b>Data Scientist Intern:</b> Google Built a machine learning pipeline to probabilistically match entities in text with entries in a database.
05/17–08/17	<b>Data Science Intern, Business Analytics:</b> Box, Inc. Built a machine learning pipeline to automatically identify new marketing and sales leads.
09/10–08/16	<b>Senior Faculty Research Specialist:</b> University of Maryland Center for Advanced Study of Language Designed, conducted, and analyzed experiments in areas such as psycholinguistics, speech perception, and signal detection theory. Presented research milestones and results to U.S. Government clients. First author on 10 technical reports to the U.S. Government.
8/09–9/10	<b>Post-Baccalaureate Research Fellow:</b> University of Maryland Department of Linguistics Designed and conducted research on syntactic processing during reading and listening, using eye tracking and EEG.

---

## Research

### *Preprints and manuscripts under review*

- 2025 Harit Vishwakarma, **Alan Mishler**, Thomas Cook, Niccolò Dalmaso, Natraj Raman, Sumitra Ganesh. (Under review). Monty hall and score optimization in conformal prediction to improve LLM performance in multi-choice question answering. <https://arxiv.org/abs/2501.00555>.
- 2022 Kwangho Kim, **Alan Mishler**, José Zubizarreta. (2022). Counterfactual mean-variance optimization. <https://arxiv.org/abs/2209.09538>.
- 2022 Raphael Sonabend, Florian Pfisterer, **Alan Mishler**, Moritz Schauer, Lukas Burk, Sumantrak Mukherjee, Sebastian Vollmer. (2022). Flexible Group Fairness Metrics for Survival Analysis. <https://arxiv.org/abs/2206.03256>.
- 2022 **Alan Mishler**, Niccolò Dalmaso. (2022). Fair when trained, unfair when deployed: observable fairness measures are unstable in performative prediction settings. <https://arxiv.org/abs/2202.05049>.
- 2021 **Alan Mishler**, Niccolò Dalmaso. (2021). When the Oracle Misleads: Modeling the Consequences of Using Observable Rather than Potential Outcomes in Risk Assessment Instruments. <https://arxiv.org/abs/2104.01921>.

### *Publications*

- 2025 **Alan Mishler**, Mohsen Ghassemi, Niccolò Dalmaso, Luhao Zhang, Vamsi K. Potluru, Tucker Balch, Manuela Veloso. (2025, in press). Auditing and enforcing conditional fairness via optimal transport. *Proceedings of the AAAI Conference on Artificial Intelligence*. <https://arxiv.org/abs/2410.14029>. [\*Authors 1-4 contributed equally.]
- 2024 Thomas Cook, **Alan Mishler**, Aaditya Ramdas. (2024). Semiparametric efficient inference in adaptive experiments. *Proceedings of the Third Conference on Causal Learning and Reasoning* 236. : 1033–1064. <https://proceedings.mlr.press/v236/cook24a.html>.
- 2024 Akshaj Kumar Veldanda, **Alan Mishler**, Ivan Brugere, Sanghamitra Dutta, Siddharth Garg. (2024). Hyper-parameter tuning for fair classification without sensitive attribute access. *Transactions on Machine Learning Research*. <https://openreview.net/forum?id=ZSWKdRi2cU>. [Authors 2-4 contributed equally.]
- 2024 Zikai Xiong, Niccolò Dalmaso, **Alan Mishler**, Vamsi K. Potluru, Tucker Balch, Manuela Veloso. (2024). Fairwasp: fast and optimal fair wasserstein pre-processing. *Proceedings of the AAAI Conference on Artificial Intelligence* 38(14): 16120–16128. <https://doi.org/10.1609/aaai.v38i14.29545>.
- 2023 **Alan Mishler**, Mohsen Ghassemi, Alec Koppel, Sumitra Ganesh. (2023). Active learning with missing not at random outcomes. *NeurIPS 2023 Workshop on Adaptive Experimental Design and Active Learning in the Real World*. <https://openreview.net/forum?id=CzdSCF0G1n>.
- 2023 Akshaj Kumar Veldanda, **Alan Mishler**, Ivan Brugere, Sanghamitra Dutta, Jiahao Chen, Siddharth Garg. (2023). Fairness via in-processing in the over-parameterized regime: a cautionary tale with mindiff loss. *Transactions on Machine Learning Research*. <https://openreview.net/pdf?id=f4VyYhkRvi>. [Authors 2-5 contributed equally.]

- 2022 **Alan Mishler**, Edward H. Kennedy. (2022). FADE: Fair double ensemble learning for observable and counterfactual outcomes. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. <https://doi.org/10.1145/3531146.3533167>.
- 2021 **Alan Mishler**, Matteo Bonvini, Edward H. Kennedy. (2021). Comment on “Statistical Modeling: The Two Cultures” by Leo Breiman [invited commentary]. *Observational Studies*. <https://muse.jhu.edu/article/799726/pdf>. [Authors 1-2 contributed equally.]
- 2021 Alessandro Fabris, **Alan Mishler**, Stefano Gottardi, Mattia Carletti, Matteo Daicampi, Gian Antonio Susto, Gianmaria Silvello. (2021). Algorithmic audit of italian car insurance: Evidence of unfairness in access and pricing. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. <https://doi.org/10.1145/3461702.3462569>.
- 2021 **Alan Mishler**, Edward H. Kennedy, Alexandra Chouldechova. (2021). Fairness in risk assessment instruments: post-processing to achieve counterfactual equalized odds. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. <https://dl.acm.org/doi/10.1145/3442188.3445902>.
- 2021 Ashley I. Naimi, **Alan E. Mishler**, Edward H. Kennedy. (2021). Challenges in obtaining valid causal effect estimates with machine learning algorithms. *American Journal of Epidemiology*. <https://academic.oup.com/aje/article/192/9/1536/6322279>.
- 2020 Amanda Coston, **Alan Mishler**, Edward H. Kennedy, Alexandra Chouldechova. (2020). Counterfactual risk assessments, evaluation, and fairness. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. <https://doi.org/10.1145/3351095.3372851>.
- 2019 **Alan Mishler**, Edward H. Kennedy, Amanda Coston, Alexandra Chouldechova. (2019). Modeling risk and achieving algorithmic fairness using potential outcomes. *Proceedings of the 2019 AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*. <https://doi.org/10.1145/3306618.3314323>.
- 2018 Erica B. Michael, Petra Bradley, **Alan Mishler**, Lelyn D. Saner, Brook Hefright, Ann Zeng, Joseph H. Danks. (2018). Human use of machine translation to extract information from texts. In Isabel Lacruz (ed.), *Innovation and expansion in translation process research*. Amsterdam: John Benjamins Publishing Company. <https://www.jbe-platform.com/content/books/9789027264756-ata.18.09mic>.
- 2017 **Alan Mishler**, Kevin Wonus, Wendy Chambers, Michael Bloodgood. (2017). Filtering tweets for social unrest. *Proceedings of the 2017 IEEE 11th International Conference on Semantic Computing (ICSC)*. <https://doi.org/10.1109/ICSC.2017.75>.
- 2016 J. Isaiah Harbison, **Alan Mishler**, Thomas Wallsten. (2016). Modeling triage decision making. *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*. <https://escholarship.org/uc/item/1461m8q3>.
- 2016 Erika K. Hussey, J. Isaiah Harbison, Susan E. Teubner-Rhodes, **Alan Mishler**, Kayla Velnoskey, Jared M. Novick. (2016). Memory and language improvements following cognitive control training. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. <https://dx.doi.org/10.1037/xlm0000283>.
- 2016 Susan Teubner-Rhodes, **Alan Mishler**, Ryan Corbett, Llorenc Andreu, Monica Sanz-Torrent, John Trueswell, Jared Novick. (2016). The bilingual advantage: conflict monitoring, cognitive control, and garden-path recovery. *Cognition* 150: 213–231. <https://doi.org/10.1016/j.cognition.2016.02.011>.

- 2013 Brian Dillon, **Alan Mishler**, Shayne Sloggett, Colin Phillips. (2013). Contrasting interference profiles for agreement and anaphora: experimental and modeling evidence. *Journal of Memory and Language* 69(2): 85–103. <http://dx.doi.org/10.1016/j.jml.2013.04.003>.
- 2012 Charles Chang, **Alan Mishler**. (2012). Evidence for language transfer leading to a perceptual advantage for non-native listeners. *Journal of the Acoustical Society of America* 132(4): 2700–10. <http://dx.doi.org/10.1121/1.4747615>.

---

### *Other conference presentations*

- 2024 Harit Vishwakarma, **Alan Mishler**, Thomas Cook, Niccolò Dalmaso, Natraj Raman, Sumitra Ganesh. (2024). Monty Hall and Optimized Conformal Prediction to Improve Decision-Making with LLMs. Poster presented at the NeurIPS Workshop on Statistical Frontiers in LLMs and Foundation Models.
- 2023 Thomas Cook, **Alan Mishler**, Aaditya Ramdas. (2023). Anytime-Valid Inference in Adaptive Experiments. Paper presented at the NeurIPS Workshop on Adaptive Experimental Design and Active Learning in the Real World.
- 2023 **Alan Mishler**, Mohsen Ghassemi, Alec Koppel, Sumitra Ganesh. (2023). Model Robustness and Active Learning with Missing-Not-At-Random Outcomes. Paper presented at the *UAI Workshop on Epistemic Uncertainty in Artificial Intelligence*.
- 2022 Luhao Zhang, Mohsen Ghassemi, **Alan Mishler**, Ivan Brugere, Niccolò Dalmaso, Vamsi Potluru, Tucker Balch, Manuela Veloso. (2022). Conditional Demographic Parity Through Optimal Transport. Paper presented at the NeurIPS Workshop on Algorithmic Fairness through the Lens of Causality and Privacy. [Authors 3-5 contributed equally.]
- 2022 Raphael Sonabend, Florian Pfisterer, **Alan Mishler**, Moritz Schauer, Lukas Burk, Sumantrak Mukherjee, Sebastian Vollmer. (2022). Flexible Group Fairness Metrics for Survival Analysis. Paper presented at the KKD Workshop on Applied Data Science for Healthcare.
- 2022 Alessandro Fabris, **Alan Mishler**, Stefano Gottardi, Mattia Carletti, Matteo Daicampi, Gian Antonio Susto, Gianmaria Silvello. (2022). Algorithmic Audit of Italian Car Insurance: Evidence of Unfairness in Access and Pricing. Poster presented at the *Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EEAMO)*.
- 2022 Alessandro Fabris, **Alan Mishler**, Stefano Gottardi, Mattia Carletti, Matteo Daicampi, Gian Antonio Susto, Gianmaria Silvello. (2022). Algorithmic Audit of Italian Car Insurance: Evidence of Unfairness in Access and Pricing. Talk given at the *European Workshop on Algorithmic Fairness*.
- 2021 **Alan Mishler**, Niccolò Dalmaso. (2021). Fair When Trained, Unfair When Deployed: Observable Fairness Measures are Unstable in Performative Prediction Settings. Poster presented at the *2022 NeurIPS workshop on Algorithmic Fairness through the Lens of Causality and Robustness*.
- 2020 **Alan Mishler**, Edward H. Kennedy, Alexandra Chouldechova. (2020). Fairness in Risk Assessment: Post-Processing to Achieve Counterfactual Equalized Odds. Talk given at the NeurIPS workshop *Algorithmic Fairness through the Lens of Causality and Interpretability*.
- 2020 **Alan Mishler**, Edward H. Kennedy. (2020). Counterfactually Fair Risk Assessment Instruments. Talk given at the Joint Statistical Meetings. <https://ww2.amstat.org/meetings/jsm/2020/index.cfm>.

- 2020 **Alan Mishler.** (2020). Achieving Counterfactual Fairness in Risk Assessment Instruments. Talk given at the CMU Symposium on AI and Social Good. <https://www.youtube.com/watch?v=V9NUCCrRZac>.
- 2019 **Alan Mishler, Niccolo Dalmasso.** (2019). When the Oracle Misleads: Modeling the Consequences of Using Observable Rather than Potential Outcomes in Risk Assessment Instruments. Paper/poster presented at NeurIPS Workshop *“Do the right thing”: machine learning and causal inference for improved decision making*. Vancouver, Canada. [https://cpb-us-w2.wpmucdn.com/sites.coecis.cornell.edu/dist/a/238/files/2019/12/Id\\_92\\_final.pdf](https://cpb-us-w2.wpmucdn.com/sites.coecis.cornell.edu/dist/a/238/files/2019/12/Id_92_final.pdf).
- 2019 **Alan Mishler, Edward H. Kennedy.** (2019). Counterfactual Fairness in the Actual World. Poster presented at the *Atlantic Causal Inference Conference*. Montreal, Canada. [amishler.github.io/files/ACIC\\_2019\\_Poster.pdf](https://amishler.github.io/files/ACIC_2019_Poster.pdf).
- 2018 **Alan Mishler, Rebecca Nugent.** (2018). Clustering students and inferring skill set profiles with skill hierarchies. Poster presented at the *2018 Classification Society Annual Meeting*. Stony Brook, NY. [amishler.github.io/files/Classification\\_Society\\_2018\\_poster.pdf](https://amishler.github.io/files/Classification_Society_2018_poster.pdf).
- 2018 **Alan Mishler, Edward H. Kennedy.** (2018). Counterfactual prediction and fairness in risk assessment tools. Poster presented at the *2018 Atlantic Causal Inference Conference*. Pittsburgh, PA. [amishler.github.io/files/ACIC\\_2018\\_poster.pdf](https://amishler.github.io/files/ACIC_2018_poster.pdf).
- 2015 **Alan Mishler, Erin Smith Crabb, Susannah Paletz, Brook Hefright, Ewa Golonka.** (2015). Using structural topic modeling to detect events and cluster Twitter users in the Ukrainian crisis. Poster presented at the *17th International Conference on Human-Computer Interaction*. Los Angeles, CA.
- 2015 Erin Smith Crabb, **Alan Mishler, Susannah Paletz, Brook Hefright, Ewa Golonka.** (2015). Reading between the lines: A prototype model for detecting Twitter sockpuppet accounts using language-agnostic processes. Poster presented at the *17th International Conference on Human-Computer Interaction*. Los Angeles, CA.
- 2015 Erin Smith Crabb, **Alan Mishler, Susannah Paletz, Brook Hefright, Ewa Golonka.** (2015). Detecting and clustering sockpuppets in Ukrainian social media: Two computational linguistic proof-of-concepts. Talk given at the *2015 INGroup [Interdisciplinary Network for Group Research] Conference*. Pittsburgh, PA.
- 2013 Erika Hussey, Isaiah Harbison, **Alan Mishler, Susan Teubner-Rhodes, Kayla Velnoskey, Jared Novick.** (2013). Memory and language improvements following cognitive control training. Poster presented at the *54th Annual Meeting of the Psychonomic Society*. Toronto, Canada.
- 2013 Erika Hussey, Susan Teubner-Rhodes, **Alan Mishler, Isaiah Harbison, Jared Novick.** (2013). Trainability and selective transferability of conflict resolution skills to parsing and non-parsing domains. Paper presented at the *26th CUNY Sentence Processing Conference*. Columbia, SC.
- 2012 Erika Hussey, J. Isaiah Harbison, **Alan Mishler, Susan Teubner-Rhodes, Jared M. Novick.** (2012). Trainability and selective transferability of conflict resolution skills. Poster presented at the *53rd Psychonomic Society Annual Meeting*. Minneapolis, MN.
- 2012 Charles B. Chang, **Alan Mishler.** (2012). The role of canonicity in shaping perceptual biases in speech processing. Paper presented at the *13th Conference on Laboratory Phonology*. Stuttgart, Germany.

- 2012 Lelyn D. Saner, J. Isaiah Harbison, DeLaura Jansen, **Alan Mishler**. (2012). Crowd-sourced ratings yield the same empirical results as trained raters. Poster presented at the *24th Annual Convention of the Association for Psychological Science*. Chicago, IL.
- 2012 Susan E. Teubner-Rhodes, **Alan Mishler**, Ryan Corbett, Llorenç A. Barrachina, Monica Sanz-Torren, John C. Trueswell, Jared M. Novick. (2012). The bilingual advantage: Conflict monitoring, cognitive control, and garden-path recovery. Poster presented at the *25th Annual CUNY Conference on Human Sentence Processing*. New York, NY.
- 2012 **Alan Mishler**. (2012). Domain-initial strengthening in voiceless stops cues word boundary perception in Japanese. Paper presented at the *86th Annual Meeting of the Linguistic Society of America*. Portland, OR.
- 2011 Brian Dillon, **Alan Mishler**, Shayne Sloggett, Colin Phillips. (2011). Contrasting interference profiles for agreement and anaphora: Experimental and modeling evidence. Poster presented at the *24th Annual CUNY Human Sentence Processing Conference*. Stanford, CA.

---

## Invited Talks

- 2024 **Alan Mishler**. (2024). Why counterfactuals matter for fairness. Invited talk given to the Nokia Bell Labs Responsible AI group.
- 2021 **Alan Mishler**. (2021). Fair and accurate risk assessment for healthcare decision making. Invited talk given at the *14th International Conference of the ERCIM Working Group on Computational and Methodological Statistics*.
- 2020 **Alan Mishler**. (2020). Achieving Counterfactual Equalized Odds in Risk Assessment Instruments. Invited talk given at the *Carnegie Mellon Symposium on AI and Social Good*.

---

## Teaching at Carnegie Mellon

- 2021 Spring TA, StatDS 401: Modern Regression
- 2020 Summ. TA, StatDS 225: Introduction to Probability Theory
- 2020 Spring TA, StatDS 731: Foundations of Causal Inference
- 2019 Spring TA, StatDS 226: Introduction to Statistical Inference
- 2018 Fall TA, StatDS 617: Applied Linear Models
- 2018 Spring TA, StatDS 326: Honors Introduction to Statistical Inference
- 2017 Fall TA, StatDS 350: Statistical Computing
- 2017 Spring TA, StatDS 402: Advanced Methods for Data Analysis
- 2016 Fall TA, StatDS 350: Statistical Computing

---

## Selected Awards and Honors

- 2023 Paper “Semiparametric Efficient Inference in Adaptive Experiments” nominated for Best Student Paper award at NeurIPS Workshop on Adaptive Experimental Design and Active Learning in the Real World.
- 2019 Travel grant, Conference on Neural Information Processing Systems (NeurIPS).
- 2019 Travel grant, AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society.

- 2018 Best poster award: “Optimized Random Partition Tree-based Kernels,” Convex Optimization Mini-Conference, Carnegie Mellon, Pittsburgh, PA. Co-authors: Benjamin LeRoy, Niccolò Dalmaso.
- 2017 Winner of the Fall 2017 Citadel Data Open at Carnegie Mellon. (550 student applications, 125 selected to compete; \$20,000 prize). Teammates: Niccolò Dalmaso, Kwangho Kim, Chirag Nagpal.
- 2017 Best short talk: “Subspace Clustering,” Statistical Machine Learning Mini-Conference, Carnegie Mellon, Pittsburgh, PA.

---

## Selected Service

- 2022–2025 Reviewer: Transactions on Machine Learning Research (TMLR; journal).
- 2022–2025 Reviewer: ACM Conference on Fairness, Accountability, and Transparency (FAccT).
- 2024 Reviewer: “Algorithmic Fairness through the Lens of Metrics and Evaluation” (NeurIPS workshop).
- 2023 Reviewer: Journal of Biostatistics.
- 2023 Reviewer: European Workshop on Algorithmic Fairness (EWAf).
- 2022 Reviewer: “Algorithmic Fairness through the Lens of Causality and Privacy” (NeurIPS workshop).
- 2021 Reviewer: “Algorithmic Fairness through the Lens of Causality and Robustness” (NeurIPS workshop).
- 2021 Reviewer: NeurIPS Datasets and Benchmarks track.
- 2020 Reviewer: American Journal of Epidemiology.
- 2019–2021 Reviewer: “Machine Learning for the Developing World (ML4D)” (NeurIPS workshop).
- 2019 Reviewer: Annals of Applied Statistics.
- 2009–2016 **Telluride Association:** Board member and committee chair for an educational nonprofit devoted to providing opportunities to extraordinary and underprivileged high school and college students.

Last updated: April 16, 2025