

Srikar Katta

PHD CANDIDATE · APPLE SCHOLAR IN AI/ML · CAUSAL INFERENCE · INTERPRETABLE MACHINE LEARNING

☎ (+1) 484-343-7381 | ✉ srikar.katta@duke.edu | 🏠 <https://kattasa.github.io/> | 📱 kattasa | 🌐 srikar-katta

Summary

Complex data from social media environments and wearable devices are becoming increasingly popular but are inherently unstructured, high-dimensional, and noisy. To handle these complexities and aid high-stakes decision-making, I develop machine learning methods that are trustworthy *by design* without sacrificing accuracy [8, 3]. In many high-stakes settings, however, decision-makers must contend with models they did not design, such as leading generative AI and machine learning systems with opaque reasoning processes. To address this, I also develop benchmarking and explainability tools to audit the capabilities and limitations of these black boxes [1, 5]. My work has been recognized by multiple student paper awards from the American Statistical Association, spotlights at leading machine learning conferences, and the Apple Scholar in AI/ML Fellowship.

Education

Duke University

PHD IN COMPUTER SCIENCE

- Advisors: Cynthia Rudin & Alexander Volfovsky

Durham, NC

August 2022 - PRESENT

Temple University

BS IN MATHEMATICS & COMPUTER SCIENCE

- Minor in Economics

Philadelphia, PA

August 2018 - May 2022

GPA: 3.96/4.00

Honors & Awards

- 2026 **Risk Analysis Section Early Career Paper Award Honorable Mention**, American Statistical Association
 - Awarded for quantifying variable importance even in the presence of unobserved variables
- 2025 **Apple Scholar in AI/ML Fellowship**, Apple
 - Invitation-only, selective fellowship supporting emerging leaders in machine learning research
- 2024 **Biometrics Section Early Career Paper Award**, American Statistical Association
 - Awarded for designing interpretable estimators of causal effects from wearables and sensor data
- 2023 **1st Place**, AAAI W3PHIAI Hackathon on Aging
 - Awarded for highlighting the distribution shift issues arising from biological age predictors
- 2023 **Open Data Award**, Scientific Research Network on Decision Neuroscience and Aging
 - Awarded for creating high-quality, publicly available data for studying financial exploitation

Publications

PUBLISHED WORK

* indicates joint authorship

- [1] J. Donnelly*, **S. Katta***, E. Borgonovo, and C. Rudin, "Variable importance with unobserved confounding and the rashomon effect," *International Conference on Artificial Intelligence and Statistics*, 2026, **Risk Analysis Section Early Career Paper Award Honorable Mention**, American Statistical Association.
- [2] H. Jeong*, **S. Katta***, W. K. Wang, A. Volfovsky, and J. Dunn, "Impact of daylight saving time on physical activity patterns," *Nature Health*, 2026, **Covered by U.S. News and World Report**.
- [3] **S. Katta**, H. Parikh, C. Rudin, and A. Volfovsky, "Interpretable causal inference for analyzing wearable, sensor, and distributional data," *International Conference on Artificial Intelligence and Statistics*, pp. 3340–3348, 2024, **Biometrics Section Early Career Paper Award**, American Statistical Association.

- [4] C. Rudin, C. Zhong*, L. Semenova*, M. Seltzer*, R. Parr*, J. Liu*, **S. Katta***, J. Donnelly*, H. Chen*, and Z. Boner*, “Amazing things come from having many good models,” *International Conference of Machine Learning*, 2024, **Spotlight Paper (Top 3% of Submissions)**.
- [5] J. Donnelly*, **S. Katta***, C. Rudin, and E. Browne, “The Rashomon importance distribution: Getting rid of unstable, single model-based variable importance,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 6267–6279, 2023, **Spotlight Paper (Top 3% of Submissions)**.
- [6] Y. J. Yang, K. Hackett, **S. Katta**, R. M. Ludwig, J. Jarcho, T. Giovannetti, D. S. Fareri, and D. V. Smith, “Psychological, social, and health-related factors predict risk for financial exploitation,” *Communications Psychology*, 2023, **SRNDNA Open Data Award**.

SUBMITTED WORK

- [7] A. van Loon, **S. Katta**, C. A. Bail, D. S. Hillygus, and A. Volfovsky, “Designing social media to promote productive political dialogue on a new research platform,” *Revise and Resubmit at Scientific Reports*, 2025.
- [8] G. Tierney*, **S. Katta***, C. Bail, S. Hillygus, and A. Volfovsky, “A new experimental design for causal inference with text: Can a language model be too big?” *Revise and resubmit at JRSSA*, 2025.

IN PREPARATION

- [9] **S. Katta**, H. Parikh, C. Rudin, A. Volfovsky, H. Sun, and B. Westover, *Why machine learning your biological age fails*, **1st place prize in AAAI W3PHIAI Hackathon on Aging**, 2025+.

Technical Experience

Apple, Inc

New York City, NY

MACHINE LEARNING RESEARCH SCIENTIST INTERN

May 2025 – Aug. 2025

- Developed longitudinal aging biomarker combining foundation models for wearables with random effects, improving disease diagnosis AUC by 20%+ over cross-sectional baselines.
- Predicted onset of adverse heart events with same accuracy as invasive longitudinal biomarkers using genetics or neuroimaging data from non-invasive wearables data
- Built scalable pipeline for processing over 149 million participant-days of wearables data using random effects models

Amazon

Seattle, WA

APPLIED SCIENTIST INTERN

May 2024 – Aug. 2024

- Derived asymptotic behavior of new estimator, enabling efficient computation of p-values and confidence intervals
- Automated analyses of customer-level experimental data at scale, reducing runtime by at least 10x
- Analyzed *hundreds* of customer-level datasets with over 9 billion observations to generate causal insights under unobserved confounding

Teaching

2025	Causal Inference in Data Analysis , Guest Lecturer	<i>Duke</i>
2023	Data Science , Teaching Assistant	<i>Duke</i>
2020	Applied Statistics and Data Science , Teaching Assistant	<i>Temple</i>
2020	Elements of Statistics and Data Science , Teaching Assistant	<i>Temple</i>

Invited Presentations

INVITED TALKS

2026	Bocconi University , Unobservables, Inference, and Variable Importance Using Rashomon Sets	<i>Milan, IT</i>
2025	INRIA , Interpretable causal inference for analyzing wearable, sensor, and distributional data	<i>Paris, FR</i>
2025	Statistical Learning and Data Science , Interpretable causal inference for analyzing wearable, sensor, and distributional data	<i>Newport Beach, CA</i>
2024	Joint Statistics Meeting , Interpretable causal inference for analyzing wearable, sensor, and distributional data	<i>Portland, OR</i>
2024	Blue Cross Blue Shield of North Carolina , Interpretable causal inference via almost matching exactly	<i>Virtual</i>
2024	INFORMS , Interpretable causal inference for analyzing wearable, sensor, and distributional data	<i>Phoenix, AZ</i>

POSTER PRESENTATIONS

2025	Experimental Design in the Age of Generative AI , A new experimental design for causal inference with text: Can a language model be too big?	<i>Berkeley, CA</i>
2024	Temple-Wells Fargo Data Science Conference , Interpretable causal inference for analyzing wearable, sensor, and distributional data	<i>Philadelphia, PA</i>

Service

2026	Session Organizer , When Does Explainable AI Actually Open the Black Box	<i>SLDS</i>
2026	Session Chair , Adaptive treatments for optimal long-term medical care	<i>JSM</i>
2026	Area Chair , NeurIPS	
2026	Reviewer , Nature	
2025	Session Chair , Interpretable Machine Learning	<i>INFORMS</i>
2025	Session Chair , Variable Importance for Scientific Discovery	<i>INFORMS</i>
2024	Session Chair , Machine Learning Aided Causal Inference	<i>INFORMS</i>
2025	Reviewer , European Journal of Operational Research, Neural Information Processing Systems	
2024	Reviewer , Machine Learning for Health, International Conference of Machine learning	
2023	Reviewer , INFORMS Journal on Computing, International Symposium on Information Theory	