

ALEXANDER KOUJIANOS GOLDBERG

✉ akgoldbe@andrew.cmu.edu
🌐 <https://akgoldberg.github.io/>

EDUCATION

2021 - Present	PhD Candidate in Computer Science Department Advised by Giulia Fanti and Nihar B. Shah	Carnegie Mellon University
2014 - 2018	A.B. in Applied Mathematics, Magna Cum Laude with Highest Honors Secondary in History, GPA: 3.91	Harvard University

RESEARCH

A Principled Approach to Randomized Selection under Uncertainty

Alexander Goldberg, Giulia Fanti, Nihar B. Shah
NeurIPS, 2025 (Spotlight)

Usefulness of LLMs as an Author Checklist Assistant for Scientific Papers: NeurIPS'24 Experiment

Alexander Goldberg, Ihsan Ullah, Thanh Gia Hieu Khuong, Benedictus Kent Rachmat, Zhen Xu, Isabelle Guyon, Nihar B. Shah
arXiv, 2024 (Under submission)

Benchmarking Fraud Detectors on Private Data

Alexander Goldberg, Giulia Fanti, Nihar B. Shah, and Steven Wu
ACM KDD 2025

What Can Natural Language Processing Do for Peer Review?

I. Kuznetsov, O. Afzal, Koen Dercksen, N. Dyck, *Alexander Goldberg*, T. Hope, D. Hovy, J. Kummerfeld, A. Lauscher, K. Leyton-Brown, S. Lu, M. Mieskes, A. Nèvéol, D. Pruthi, L. Qu, R. Schwartz, N.A. Smith, T. Solorio, J. Wang, X. Zhu, A. Rogers, N.B. Shah, I. Gurevych
arXiv, 2024

Peer Reviews of Peer Reviews: A Randomized Controlled Trial and Other Experiments

Alexander Goldberg, Ivan Stelmakh, Kyunghyun Cho, Alice Oh, Alekh Agarwal, Danielle Belgrave, Nihar B. Shah
PLOS ONE, 2025

Benchmarking Private Population Data Release Mechanisms: Synthetic Data vs. TopDown

Aadyaa Maddi, Swadhin Routray, Alexander Goldberg, Giulia Fanti
PPAI Workshop @ AAAI 2024

Batching of Tasks by Users of Pseudonymous Forums: Anonymity Compromise and Protection

Alexander Goldberg, Giulia Fanti, and Nihar B. Shah
ACM SIGMETRICS, 2023

Towards Differentially Private Inference on Network Data

Alexander Goldberg (advised by Salil Vadhan)
Undergraduate Thesis, 2018

AWARDS

2023	NSF Graduate Research Fellowship
2018	Thomas Temple Hoopes Prize Awarded by Harvard College for outstanding work on my senior thesis titled "Towards Differentially Private Inference on Network Data" advised by Professor Salil Vadhan.
2018	Phi Beta Kappa, Harvard College
2018	Certificate of Distinction in Teaching Awarded by the Derek Bok Center at Harvard to reflect excellence in teaching as a Teaching Fellow in Computer Science 181 (Machine Learning).
2017	History Department Undergraduate Essay Prize Awarded for the best work of original scholarship produced before the senior year in a history department course at Harvard for my paper entitled "'The House in China was Obligated to Succumb': The Decline of American Traders in China and the Rise of a Chinese Merchant Class, 1873-1883."

WORK EXPERIENCE

2018-2021	Facebook - Data Scientist <ul style="list-style-type: none"> • Worked on the Ads Integrity team, leveraging statistical analysis and machine learning to ensure that Facebook's ad ecosystem is safe and trustworthy. • Designed improvements to Facebook's Political Ad Library to improve transparency into political advertising in the lead-up to US 2020 elections. • Improved measurement and classification of counterfeits and job scams in online ads.
2017	Facebook - Data Science Intern Worked on Messenger analytics team to drive strategy around increasing teenagers' engagement with Messenger using data.
2016	Microsoft - Software Engineering Intern As part of the Microsoft Garage program, built and shipped a conversational chat bot for developers that integrated with Visual Studio Team Services and Slack.

TEACHING EXPERIENCE

Fall 2025	Information Security, Privacy, and Policy (17-331) Instructors: Joseph A. Calandrino and Hana Habib	CMU
Fall 2023	Foundations of Privacy (18-734), Teaching Assistant Instructors: Steven Wu and Giulia Fanti	CMU
Spring 2018	Machine Learning (CS 181), Teaching Fellow Instructor: Finale Doshi-Velez	Harvard
Spring 2016	Abstraction and Design in Computation (CS 51), Teaching Fellow Instructor: Stuart M. Shieber	Harvard
Fall 2015	Introduction to Computer Science (CS 50), Teaching Fellow Instructor: David J. Malan	Harvard

TALKS

September 2025	Peer Reviews of Peer Reviews: A Randomized Controlled Trial and Other Experiments International Congress on Peer Review and Scientific Publication, 2025 (Podium Presentation)
November 2024	Usefulness of LLMs as an Author-Checklist Assistant for Scientific Papers: NeurIPS '24 Experiment METRICS International Forum, Stanford University
January 2024	Evaluating the peer review process Reviewer No. 2: Old and New Problems in Peer Review (Dagstuhl Seminar)
April 2023	Batching of Tasks by Users of Pseudonymous Forums: Anonymity Compromise and Protection CMU Theory Lunch (video)

SKILLS

Technical skills: Python, NumPy/SciPy, SQL, LaTeX

Research areas: Distributed human evaluations, differential privacy, machine learning, statistics, graph mining