

Prakhar Ganesh

COMPUTER SCIENCE · PHD STUDENT

McGill University | Quebec AI Institute (Mila)

✉ prakhar.ganesh@mila.quebec | 🌐 <https://prakharg24.github.io/>

Research Interests: I am interested in the challenges that arise in AI when moving from the sandbox to the real world, primarily at the intersection of multiplicity, fairness, and privacy, with a focus on uncovering the learning dynamics of neural networks.

Education

McGill University | Quebec AI Institute (Mila) | Montreal, Canada

PHD COMPUTER SCIENCE (THESIS) | ADVISOR: PROF GOLNOOSH FARNADI

Jan 2024 - Present

National University of Singapore (NUS) | Singapore

MASTER OF COMPUTING (AI SPECIALISATION) | ADVISOR: PROF REZA SHOKRI

Aug 2021 - Mar 2023

Indian Institute of Technology Delhi (IITD) | New Delhi, India

BACHELOR OF TECHNOLOGY IN COMPUTER SCIENCE AND ENGINEERING

Aug 2015 - May 2019

Work Experience

APPLE | AI/ML - MACHINE LEARNING INTERN | CUPERTINO, USA

Apr 2026 - Sep 2026

NATIONAL UNIVERSITY OF SINGAPORE (NUS) | RESEARCH INTERN | SINGAPORE

Jan 2023 - Apr 2023

ALVISUAL | AI ENGINEER (INTERN) | SINGAPORE

Jan 2022 - Mar 2022

ADVANCED DIGITAL SCIENCES CENTER (ADSC) | RESEARCH ENGINEER | SINGAPORE

Jun 2019 - Jul 2021

KRIDA.AI | ENGINEER (INTERN) | NEW DELHI, INDIA

Jan 2019 - May 2019

WEALTHNET ADVISORS | ANALYST (INTERN) | NEW DELHI, INDIA

May 2018 - Jul 2018

Publications

Rethinking Hallucinations: Correctness, Consistency, and Prompt Multiplicity

[ORAL] [RABAT, MOROCCO] EUROPEAN CHAPTER OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS (EACL), 2026

[NON-ARCHIVAL] [SINGAPORE] WORKSHOP ON BUILDING TRUST IN LANGUAGE MODELS AND APPLICATIONS, ICLR 2025

Prakhar Ganesh, Reza Shokri, Golnoosh Farnadi

Say It Another Way: A Framework for User-Grounded Paraphrasing

[ORAL] [RABAT, MOROCCO] EUROPEAN CHAPTER OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS (EACL), 2026

Cléa Chataigner*, Rebecca Ma*, Prakhar Ganesh, Yuhao Chen, Afaf Taïk, Elliot Creager, Golnoosh Farnadi

Data as a Lever: A Neighbouring Datasets Perspective on Predictive Multiplicity

[NON-ARCHIVAL] [SINGAPORE] WORKSHOP ON NAVIGATING MODEL UNCERTAINTY AND THE RASHOMON EFFECT (MURE), AAAI 2026

Prakhar Ganesh, Hsiang Hsu, Golnoosh Farnadi

SoK: Data Minimization in Machine Learning

[MUNICH, GERMANY] IEEE CONFERENCE ON SECURE AND TRUSTWORTHY MACHINE LEARNING (SATML), 2026

Robin Staab*, Nikola Jovanovic*, Kimberly Mai, Prakhar Ganesh, Martin Vechev, Ferdinando Fioretto, Matthew Jagielski

Systemizing Multiplicity: The Curious Case of Arbitrariness in Machine Learning

[MADRID, SPAIN] AAAI/ACM CONFERENCE ON AI, ETHICS, AND SOCIETY (AIES), 2025

Prakhar Ganesh, Afaf Taïk, Golnoosh Farnadi

From Model Multiplicity to Prompt Multiplicity: Emerging Arbitrariness Concerns in the Age of GenAI

[STUDENT ABSTRACT] [MADRID, SPAIN] AAAI/ACM CONFERENCE ON AI, ETHICS, AND SOCIETY (AIES), 2025

Prakhar Ganesh

Towards More Realistic Extraction Attacks: An Adversarial Perspective

TRANSACTIONS OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS (TACL), 2025

[NON-ARCHIVAL] [BANGKOK, THAILAND] WORKSHOP ON PRIVACY IN NATURAL LANGUAGE PROCESSING (PRIVATENLP), ACL 2024

Yash More*, Prakhar Ganesh*, Golnoosh Farnadi

The Data Minimization Principle in Machine Learning

[ATHENS, GREECE] ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY AND TRANSPARENCY (FACCT), 2025

[CONTRIBUTED TALK] [NON-ARCHIVAL] [VANCOUVER, CANADA] WORKSHOP ON REGULATABLE ML, NEURIPS 2024

[NON-ARCHIVAL] [VIENNA, AUSTRIA] WORKSHOP ON GENERATIVE AI + LAW (GENLAW), ICML 2024

Prakhar Ganesh, Cuong Tran, Reza Shokri, Ferdinando Fioretto

Different Horses for Different Courses: Comparing Bias Mitigation Algorithms in ML

[SPOTLIGHT] [VANCOUVER, CANADA] WORKSHOP ON ALGORITHMIC FAIRNESS THROUGH THE LENS OF METRICS AND EVALUATION, NEURIPS 2024

Prakhar Ganesh*, Usman Gohar*, Lu Cheng, Golnoosh Farnadi

Model Multiplicity for Responsible AI

[STUDENT ABSTRACT] [SAN JOSE, USA] AAAI/ACM CONFERENCE ON AI, ETHICS, AND SOCIETY (AIES), 2024

Prakhar Ganesh

An Empirical Investigation into Benchmarking Model Multiplicity for Trustworthy Machine Learning: A Case Study on Image Classification

[HAWAII, USA] IEEE/CVF WINTER CONFERENCE ON APPLICATIONS OF COMPUTER VISION (WACV), 2024

Prakhar Ganesh

On The Impact of Machine Learning Randomness on Group Fairness

[BEST PAPER AWARD] [CHICAGO, USA] ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY AND TRANSPARENCY (FACCT), 2023

Prakhar Ganesh, Hongyan Chang, Martin Strobel, Reza Shokri

YOLO-ReT: Towards High Accuracy Real-time Object Detection on Edge GPUs

[HAWAII, USA] IEEE/CVF WINTER CONFERENCE ON APPLICATIONS OF COMPUTER VISION (WACV), 2022

Prakhar Ganesh*, Yao Chen*, Yin Yang, Deming Chen, Marianne Winslett

HiKonv: High Throughput Quantized Convolution With Novel Bit-wise Management and Computation

[VIRTUAL] ASIA AND SOUTH PACIFIC DESIGN AUTOMATION CONFERENCE (ASP-DAC), 2022

Yao Chen*, Xinheng Liu*, Prakhar Ganesh, Junhao Pan, Jinjun Xiong, Deming Chen

Compressing Large-Scale Transformer-Based Models: A Case Study on BERT

TRANSACTIONS OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS (TACL), 2021

[PRESENTED AT] [DUBLIN, IRELAND] ASSOCIATION FOR COMPUTATIONAL LINGUISTICS (ACL), 2022

Prakhar Ganesh, Yao Chen, Xin Lou, Mohd. Ali Khan, Yin Yang, Hassan Sajjad, Preslav Nakov, Deming Chen, Marianne Winslett

Learning-based Detection and Characterization of Time Delay Attack in Cyber-Physical Systems

IEEE TRANSACTIONS ON SMART GRID (TSG), 2021

Prakhar Ganesh, Xin Lou, Yao Chen, Rui Tan, David K.Y. Yau, Deming Chen, Marianne Winslett

Assessing and Mitigating the Impact of Time Delay Attack: Case Studies for Power Grid Controls

IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (JSAC), 2020

Xin Lou, Cuong Tran, Rui Tan, David K.Y. Yau, Zbigniew T. Kalbarczyk, Ambarish Kumar Banerjee, Prakhar Ganesh

Invited Talks

- May 2026 **Summer School in RAI and Human Rights, Mila (Montreal, Canada)**, Technical Mitigation Strategies
- Jan 2026 **AAAI 2026 MURE Workshop (Singapore)**, Recentering Multiplicity Around Developer Decisions
- June 2025 **AI Perspective Student Program, IVADO (Montreal, Canada)**, Algorithmic Fairness: A Hands-on Introduction
- May 2025 **Summer School in RAI and Human Rights, Mila (Montreal, Canada)**, Technical Mitigation Strategies
- Oct 2024 **Responsible AI Week, Mila (Montreal, Canada)**, The Curious Case of Arbitrariness in Machine Learning
- Feb 2024 **FATE Reading Group, Mila**, On The Impact of Machine Learning Randomness on Model Fairness
- Jan 2024 **Saarland University (Virtual)**, On The Impact of Machine Learning Randomness on Model Fairness
- May 2023 **EQUAL Lab, Mila (Virtual)**, On The Impact of Machine Learning Randomness on Model Fairness

Awards & Scholarships

- 2026 **Gold Reviewer (Top 25%)**, International Conference on Machine Learning (ICML)
- 2025 **TA Excellence Award**, School of Computer Science - McGill University
- 2025 **Winner**, Upper Bound - AI Research Pitch Competition
- 2024 **Doctoral Research Scholarship**, Fonds de recherche du Québec – Nature et Technologies (FRQNT)
- 2024 **Excellence Scholarship - EDI in Research Category**, Mila EDI Scholarship Program
- 2024 **Graduate Excellence Award**, McGill University
- 2023 **Graduate Scholarship**, Prof. Golnoosh Farnadi, Quebec AI Institute (Mila)
- 2023 **Best Paper Award**, ACM Conference on Fairness, Accountability and Transparency (FACCT)
- 2023 **Travel Grant**, ACM Conference on Fairness, Accountability and Transparency (FACCT)
- 2022 **Winner**, 8-Bit Bias Bounty Competition - Supervised Learning Track
- 2021 **Top 3 Teams**, Asia and South Pacific Design Automation Conference (ASP-DAC) System Design Contest

Travel Grants

- 2026 **Graduate Research Enhancement And Travel (GREAT) Award**, McGill University
- 2026 **Travel Grant**, AAAI Conference on Artificial Intelligence (AAAI) - Doctoral Consortium
- 2025 **Graduate Research Enhancement And Travel (GREAT) Award**, McGill University
- 2025 **Travel Grant**, AAAI/ACM Conference on AI, Ethics, and Society (AIES) - Student Program

- 2025 **Travel Grant**, ACM Conference on Fairness, Accountability and Transparency (FACCT) - Doctoral Consortium
- 2024 **Travel Grant**, IAPP Privacy Security Risk (PSR) - Student Scholar program
- 2024 **Travel Grant**, AAAI/ACM Conference on AI, Ethics, and Society (AIES) - Student Program
- 2024 **Travel Grant**, ACM Conference on Fairness, Accountability and Transparency (FACCT) - Doctoral Consortium
- 2024 **Graduate Research Enhancement And Travel (GREAT) Award**, McGill University
- 2023 **Travel Grant**, ACM Conference on Fairness, Accountability and Transparency (FACCT)

Teaching and Mentorship Experience

- 2026 **AAAI Education AI Videos**, Multiplicity: The Curious Case of Arbitrariness in Machine Learning
- 2026 **Bridge Mentor**, CSGS x CSUS Bridge Mentorship Program, McGill University
- 2025 **NeurIPS Education Showcase**, An Interactive Demo of Multiplicity in ML
- 2025 **Course Lecturer**, (Co-teaching w/ Prof Jackie Cheung) Fall-COMP-424: Artificial Intelligence, McGill University
- 2025 **Instructor**, Indigenous Pathfinders in AI, Mila
- 2025 **Teaching Assistant**, AI4Good Lab: Summer ML Bootcamp, Mila
- 2025 **Teaching Assistant**, Winter-COMP-598: Responsible AI, McGill University
- 2025 **Bridge Mentor**, CSGS x CSUS Bridge Mentorship Program, McGill University
- 2024 **Teaching Assistant**, Fall-COMP-370: Introduction to Data Science, McGill University
- 2024 **Instructor**, Indigenous Pathfinders in AI, Mila
- 2024 **Teaching Assistant**, AI4Good Lab: Summer ML Bootcamp, Mila
- 2024 **Coding Mentor**, Girls Who Code Hackathon, McGill University
- 2024 **Teaching Assistant**, Winter-COMP-553: Algorithmic Game Theory, McGill University

Organizational Experience

- 2026 **ACM FACCT Tutorial (Montreal)**, How to Cheat Responsible AI Audits (and Prevent Cheating)
- 2026 **ICLR Workshop (Rio de Janeiro)**, Algorithmic Fairness Across Alignment Procedures and Agentic Systems
- 2026 **Embedded Ethics @ McGill (Montreal)**, Working Group (All Courses) and Module Developer (Winter-COMP-303)
- 2026 **AAAI Tutorial and Lab Forum (Singapore)**, The Many Faces of Multiplicity in Machine Learning
- 2025 **Embedded Ethics @ McGill (Montreal)**, Working Group (All Courses) and Module Developer (Fall-COMP-424)
- 2025 **ACM FACCT Tutorial and Off-Session Social (Athens)**, The Many Faces of Multiplicity in Machine Learning

Volunteer Experience

- 2026 **Student Volunteer (Montreal)**, ACM Conference on Fairness, Accountability and Transparency (FACCT)
- 2026 **Student Volunteer (Rabat)**, European Chapter of the Association for Computational Linguistics (EACL)
- 2025 **Student Volunteer (Athens)**, ACM Conference on Fairness, Accountability and Transparency (FACCT)
- 2025 **Student Volunteer (Singapore)**, International Conference on Learning Representations (ICLR)
- 2024 **Student Volunteer (Rio De Janeiro)**, ACM Conference on Fairness, Accountability and Transparency (FACCT)

Professional Service

Area Chair

EAAMO 2026; ACL ROLLING REVIEWS - DEC 2024

Peer Reviewer

NEURIPS 2026, 2025, 2024, 2023; ICLR 2026, 2025; ICML 2026, 2025; FACCT 2026, 2025, 2024; CVPR 2026; AISTATS 2026, 2025; AAAI 2026, 2021; JAIR 2026; AIES 2025; WACV 2025, 2024, 2022, 2021; ACL ROLLING REVIEWS - FEB 2024; MURE WORKSHOP @ AAAI 2026; ACA WORKSHOP @ NEURIPS 2025; WINLP WORKSHOP @ EMNLP 2025; BUILDINGTRUST WORKSHOP @ ICLR 2025; AFME WORKSHOP @ NEURIPS 2024; GENAI WORKSHOP @ CVPR 2024; AFT WORKSHOP @ NEURIPS 2023

Ethics Reviewer

ACL ROLLING REVIEWS - JUNE 2024