

Dhruv Agarwal

agdhruv.github.io • da399@cornell.edu

RESEARCH INTERESTS

Natural Language Processing (NLP), Human-AI Interaction, AI Alignment

EDUCATION

Cornell University Ithaca, USA

PHD, COMPUTER SCIENCE | 2022–Present

Advisor: Aditya Vashistha

Ashoka University India

BS, COMPUTER SCIENCE | 2016–2020

CGPA: **3.84** / 4.0 • Major: **3.98** / 4.0 • **Magna cum laude**

Thesis: Incentivizing Drivers to Collect Off-Route Pollution Data

EXPERIENCE

Microsoft Research Redmond

RESEARCH INTERN, ALIGNMENT | Jun–Aug 2025

- Built a multi-agent system to systematize a concept for robust AI evaluation
- Improved math reasoning in LLMs using internal dialogue

Google DeepMind New York

STUDENT RESEARCHER | Feb–May 2025

- Investigated safety & capabilities of multimodal models for captioning images of humans

Accenture Labs San Francisco

AI RESEARCH INTERN | Jun–Aug 2024

- Built an LLM-powered interface to help general readers understand scientific papers
- Investigated racial bias in AI models using logit-level evaluations of GPT, LLaMA, Mistral

Microsoft Research Bangalore, India

RESEARCH FELLOW | Jun 2020–Jul 2022

- Built Turn2Earn to improve data coverage using vehicular crowdsensing
- Built a mixed-effects model using wind dispersion models and NASA pollutant data

RESEARCH INTERN | May–Aug 2019

- Designed a greedy algorithm to maximize city coverage for drive-by sensing.
- Library open-sourced by Microsoft: github.com/microsoft/Modulo

Mirabilis Design Gurgaon, India

SOFTWARE ENGINEERING INTERN | Dec 2018–Jan 2019

- Generated execution traces of ARM executables from QEMU source code for VisualSim software.

TECHNICAL SKILLS

Python • PyTorch • JAX • Distributed Pretraining
AI Agents • Reinforcement Learning • AI evals

JavaScript • React/Next.js • Multimodal Safety
RAG + Vector DBs • Human-AI Experiments

PUBLICATIONS

1. PLURAL: A Preference Dataset for Global Value Alignment. Under review (COLM).
D. Agarwal, A. Shukla, T. Goyal, A. Vashistha.

2. AI-Assisted Systematization for Evaluating GenAI Systems. Under review (COLM).
D. Agarwal*, E. Sheng*, ..., A. Chouldechova, S. Barocas, H. Wallach. [Product Impact](#)
3. Fluent but Foreign: Even Regional LLMs Lack Cultural Alignment. Under review (ACL).
D. Agarwal, A. Shukla, S. Sitaram, A. Vashistha.
4. AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances. ACM CHI 2025.
D. Agarwal, M. Naaman, A. Vashistha. [Press Coverage](#)
5. Steering AI-driven Personalization of Scientific Text for General Audiences. ACM CSCW 2025.
T.Kim, D. Agarwal, J. Ackerman, M. Saha.
6. ‘One Style Does Not Regulate All’: Moderation Practices in Public and Private WhatsApp Groups. ACM CSCW 2025.
F. Shahid, D. Agarwal, A. Vashistha.
7. Conversational Agents to Facilitate Deliberation on Harmful Content in WhatsApp Groups. ACM CSCW 2024.
D. Agarwal, F. Shahid, A. Vashistha.
8. “If it is easy to understand then it will have value”: Examining Perceptions of Explainable AI with Community Health Workers in Rural India. ACM CSCW 2024.
D. Agarwal*, C. T. Okolo*, N. Dell, A. Vashistha.
9. PollutionMapper: Identifying Global Air Pollution Sources. ACM COMPASS 2023.
D. Agarwal, P. Kumar, S. Iyengar.
10. India’s “Aadhaar” Biometric ID: Structure, Security, and Vulnerabilities. Financial Crypto 2022.
D. Agarwal*, P.R. Tiwari*, P. Jain, S. Dasgupta, D. Gupta. (21% acceptance rate)
11. Understanding Driver-Passenger Interactions in Vehicular Crowdsensing. ACM CSCW 2021.
D. Agarwal, S. Agarwal, V. Singh, S. Iyengar, M. Jain. [Methods Recognition Award](#)
12. Modulo: Drive-by Sensing at City-scale on the Cheap. ACM COMPASS 2020.
D. Agarwal, S. Iyengar, M. Swaminathan, et al. [Best Paper Award](#)
13. Poster: An Alternative to India’s Reservation Policy. ACM COMPASS 2020.
D. Sinha, O. Sahasrabudhe, D. Agarwal, D. Gupta.
14. Poster: System for Vehicle Selection in Drive-by Sensing. ACM SenSys 2019.
D. Agarwal, S. Iyengar, M. Swaminathan.

AWARDS AND RECOGNITION

LinkedIn Fellowship (\$75k)

Quad Fellowship (\$50k)

Cornell University Fellow (\$75k)

Winner – Edge Computing Track

Undergraduate Research Excellence Award (top 2)

Student Life Excellence Award (< 1%)

Dean’s List – All 8 Semesters

Outstanding Graduating Student

LinkedIn 2025
Schmidt Futures 2023
Cornell University 2022
Microsoft Global Hackathon 2021
CS@Ashoka, 2020
Ashoka University, 2020
Ashoka University, 2016–2020
High School, 2016

SERVICE

Reviewer – COLM

Reviewer – IASEAI

Reviewer – ACM CHI

Reviewer – ACM CSCW

Reviewer – ACM COMPASS

2026

2026

2023–2026

2023–2025

2021–2024

MEDIA COVERAGE

The Atlantic

The Great Language Flattening

MSN

How AI tools can threaten cultural diversity

The Hindu

Ola, Microsoft to analyse street-level air quality

The Business Insider

Ola and Microsoft come together to measure air quality on the streets of India's national capital

The Economic Times

Ola partners with Microsoft Research to study street-level air quality in Delhi-NCR

SELECTED COURSEWORK

AI/ML Intro to ML • Text Processing • Grad NLP • Grad ML • Advanced LLMs
Math Discrete Math • Probability & Statistics • Linear Algebra
Systems Computer Organization • OS • PLDI (PLs) • Networks • Grad OS
Theory Algorithms • Theory of Computation • Information Networks
Security Computer Security • Blockchain
HCI Data Viz • CS Ethics • ICTD

TA Intro to ML • Discrete Math • HCI

PAST PROJECTS

Exploring Human Bias in FIFA Ratings

Prediction of football players' future FIFA ratings from their previous ratings and on-ball actions.

Fake News Challenge

An NLP approach to the Stance Detection problem using a novel "staged" technique.

Beating Facial Obfuscation with ML

Convolutional autoencoder to recreate blurred, pixelated or blocked facial images.

COVID-19 Multilingual Knowledgebase

FAQ website to combat misinformation surrounding the crisis in local Indian languages.

VAHN: Vehicle Accident Hazard Notifier

A real-time smart city system for automatic detection and reporting of road accidents.

Ashoka Menu and Shuttle Bot

Chatbot to inform students about mess menu, shuttle timings. [800 users](#) and 190k messages.

Energy Sensing System for Dorms

Hardware prototype to monitor real-time energy wastage expressed in monetary loss.