

Anjishnu Mukherjee

PhD in Computer Science | George Mason University

🌐 iamshnoo.github.io @ amukher6@gmu.edu 🌐 github.com/iamshnoo 🎓 Google Scholar

Research Interests

• NLP • LLMs • Human-centered AI • Multilinguality • Multimodality • Alignment • Society • Culture

I work on Responsible AI, specifically in the evaluation and mitigation of cultural biases within multimodal, multilingual applications with a focus on LLMs. I have published recently at EMNLP, NAACL, ECIR, WACV and AIES.

Education

Present Aug 2022	George Mason University Ph.D. in Computer Science, <i>GPA: 3.93/4</i> Advisors: Antonios Anastasopoulos, Ziwei Zhu	Fairfax, USA
Dec 2025 Aug 2022	George Mason University M.S. in Computer Science, <i>GPA: 3.93/4</i>	Fairfax, USA
Jun 2021 Jul 2017	Indian Institute of Engineering Science and Technology (IIST) Bachelor of Technology (First class) in Computer Science, <i>CGPA: 9.63/10</i>	Kolkata, India

Experience

Aug 2025 May 2025	Amazon <i>Applied Scientist Intern Mentor: Wenqin Hu</i> Project: Predicting product return trends from historical data using LLMs for A/B testing experiments	Seattle, USA
Aug 2024 May 2024	Amazon <i>Applied Scientist Intern Mentor: Wenqin Hu</i> Project: Grounded explanations for black-box Size Recommendation systems using LLMs	Seattle, USA
Aug 2022 May 2022	Google Summer of Code Google People + AI Research <i>Developer Mentor: Ryan Mullins</i> Project: Tabular Feature Attributions and Saliency Methods for Language Interpretability using LIT	Remote
May 2022 July 2021	Wells Fargo <i>Program Associate Mentor: Aswathy P S</i> Project: Interactive visualizations to track adoption metrics for Microsoft Teams and Zoom, Automatic detection and editing of system configurations for video conferencing devices	Bangalore, India
Aug 2020 Jun 2020	Google Summer of Code mlpack <i>Student Developer Mentors: Ryan Curtin, Sreenik Seal</i> Project: Computer vision functionality and loss functions for mlpack v4	Remote
Aug 2020 Mar 2020	University of Bremen <i>Research Intern Advisor: Gabriel Zachmann</i> Project: Heuristic Time of Impact algorithm for Continuous Collision Detection for CollDet.	Bremen, Germany

Selected Publications

- [S.1] **Metadata Conditioned Large Language Models for Localization** [PDF | Code]
Anjishnu Mukherjee, Ziwei Zhu, Antonios Anastasopoulos
Under Review [ArXiv]
- [C.7] **Crossroads of Continents: Automated Artifact Extraction for Cultural Adaptation with Large Multimodal Models** [PDF | Code]
Anjishnu Mukherjee, Ziwei Zhu, Antonios Anastasopoulos
IEEE/CVF Winter Conference on Applications of Computer Vision [WACV'25]
- [C.6] **BiasDora: Exploring Hidden Biased Associations in Vision-Language Models** [PDF | Code]
Chahat Raj, Anjishnu Mukherjee, Aylin Caliskan, Antonios Anastasopoulos, Ziwei Zhu
Conference on Empirical Methods in Natural Language Processing [EMNLP Findings'24]

- [C.5] **Breaking Bias, Building Bridges: Evaluation and Mitigation of Social Biases in LLMs via Contact Hypothesis** [PDF | Code]
 Anjishnu Mukherjee*, Chahat Raj*, Aylin Caliskan, Antonios Anastasopoulos, Ziwei Zhu
 AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society [AIES'24]
- [C.4] **Global Gallery: The Fine Art of Painting Culture Portraits through Multilingual Instruction Tuning** [PDF | Code]
 Anjishnu Mukherjee, Aylin Caliskan, Ziwei Zhu, Antonios Anastasopoulos
 Conference of the North American Chapter of the Association for Computational Linguistics [NAACL'24]
- [C.3] **SALSA: Saliency-Based Switching Attack for Adversarial Perturbations in Fake News Detection Models** [PDF | Code]
 Anjishnu Mukherjee*, Chahat Raj*, Hemant Purohit, Antonios Anastasopoulos, Ziwei Zhu
 European Conference on Information Retrieval [ECIR'24]
- [C.2] **Global Voices, Local Biases: Socio-cultural Prejudices across Languages** [PDF | Code]
 Anjishnu Mukherjee*, Chahat Raj*, Ziwei Zhu, Antonios Anastasopoulos
 Conference on Empirical Methods in Natural Language Processing [EMNLP'23]
- [C.1] **True and Fair: Robust and Unbiased Fake News Detection via Interpretable Machine Learning** [PDF | Code]
 Chahat Raj, Anjishnu Mukherjee, Ziwei Zhu
 AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society [AIES'23]

Talks

- “Crossroads of Continents: Automated Artifact Extraction for Cultural Adaptation with Large Multimodal Models”
 > WACV March 2025 (Arizona, USA)
- “Tutorial: Diffusion Models”
 > CS678, Advanced NLP - Guest Lecture @ GMU Nov 2024 (Fairfax, USA)
- “Multilingual socio-cultural biases”
 > Multilinguality and Language Diversity - Invited Speaker Series @ UToronto Sep 2024 (Remote)
- “Global Gallery: The Fine Art of Painting Culture Portraits through Multilingual Instruction Tuning”
 > MASC-SLL @ JHU May 2024 (Baltimore, USA)
- “Cross-Lingual Biases and Cultural Understanding in LLMs”
 > NL+, the Natural Language Processing Lunch Seminar @ University of Notre Dame Feb 2024 (Remote)

Awards

- Outstanding Reviewer** EMNLP 2024
- Outstanding GTA Award @ GMU, 2024** For Advanced NLP, CS 678.
- Summer Graduate Research Award @ GMU, 2023** For evaluating socio-cultural biases in multilingual language models.
- Scholarship @ DAAD WISE, 2020** Awarded USD 3K by the German Academic Exchange Service (1 of 130 across India).

Teaching

- CS 678 Advanced Natural Language Processing, GMU** Teaching Assistant Fall'23, Spring'24
 > Responsible for aiding in developing new assignments, 1:1 student tutoring and assignment grading.
- CS 112 Introduction to Python Programming, GMU** Teaching Assistant Fall'22, Spring'23
 > Responsible for teaching labs, 1:1 student tutoring, creating lab modules, assignment grading, monitoring exams.

Academic Service

- Reviewer** ACL ARR - '23-'26 (EMNLP, ACL, NAACL, EACL)
- Sub-Reviewer** AIES - '23-'25; COLM - '24-'25
- Reviewer (Workshops)** AfricaNLP - '22-'26; New In ML - NeurIPS'23; WinNLP - EMNLP'25,'24; LTEDI - EACL'24