

Anjishnu Mukherjee

PhD in Computer Science | George Mason University

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

Research Interests

• Multilingual & Multicultural AI • Locale-Aware Text & Image Generation • Social Bias Mitigation

I study how to build AI systems that work equitably beyond dominant languages and cultures. My research spans understanding multilingual and multicultural AI, adapting text and image generation to local contexts, and mitigating social bias. I have published at EMNLP, COLM, NAACL, ICML, ECIR, WACV, and AIES.

Education

May 2027 Aug 2022	George Mason University Ph.D. in Computer Science, GPA: 3.93/4 Advisors: Antonios Anastasopoulos , Ziwei Zhu	Fairfax, VA
Dec 2025 Aug 2022	George Mason University M.S. in Computer Science, GPA: 3.93/4; <i>Distinguished Academic Achievement Award</i>	Fairfax, VA
Jun 2021 Jul 2017	Indian Institute of Engineering Science and Technology (IIST) B.Tech. in Computer Science, <i>First Class</i> ; CGPA: 9.63/10	Kolkata, India

Experience

Aug 2026 May 2026	Amazon <i>Applied Scientist Intern</i> Mentor: Brian Balagot Project: Grounded summary generation for product attributes discussed in returns and reviews	Seattle, USA
Aug 2025 May 2025	Amazon <i>Applied Scientist Intern</i> Mentor: Wenqin Hu 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</i> Mentor: Wenqin Hu 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</i> Mentor: Ryan Mullins Project: Tabular Feature Attributions and Salience Methods for Language Interpretability using LIT	Remote
May 2022 July 2021	Wells Fargo <i>Program Associate</i> Mentor: Aswathy P S 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</i> Mentors: Ryan Curtin , Sreenik Seal Project: Computer vision functionality and loss functions for mlpack v4	Remote

Peer-Reviewed Publications

- [C.12] **Double Trouble: Bilingual Pretraining Leaves Language-Conditioned Effects in Shared-Language Representations** [PDF | Code]
[Anjishnu Mukherjee](#), Ziwei Zhu, Antonios Anastasopoulos
Conference on Empirical Methods in Natural Language Processing [EMNLP Main '26]
- [C.11] **MAPLE: Metadata Conditioned LLM Pretraining for Locale-Aware Question Answering** [PDF | Code]
[Anjishnu Mukherjee](#), Ziwei Zhu, Antonios Anastasopoulos
Conference on Empirical Methods in Natural Language Processing [EMNLP Main '26]

- [C.10] **Lost in the Tower of Babel: The Adverse Effects of Incidental Multilingualism in LLMs** [PDF | Code]
Anjishnu Mukherjee, Chutong Meng, Antonios Anastasopoulos
Conference on Language Modeling [COLM '26]
- [C.9] **Knowing Bias, Doing Better: Mitigating Social Bias in LLMs via Know-Bias Neuron Enhancement** [PDF | Code]
Jinhao Pan, Chahat Raj, Anjishnu Mukherjee, Sina Mansouri, Bowen Wei, Shloka Yada, Ziwei Zhu
International Conference on Machine Learning; Poster [ICML '26]
- [C.8] **Measuring South Asian Biases in Large Language Models** [PDF | Code]
Mamnuya Rinki, Chahat Raj, Anjishnu Mukherjee, Ziwei Zhu
Workshop on Trustworthy Natural Language Processing [TrustNLP at ACL '26]
- [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; Oral [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
Findings of the Conference on Empirical Methods in Natural Language Processing; Poster [EMNLP Findings '24]
- [C.5] **Breaking Bias, Building Bridges: Evaluation and Mitigation of Social Biases in LLMs via Contact Hypothesis** [PDF | Code]
Chahat Raj*, Anjishnu Mukherjee*, Aylin Caliskan, Antonios Anastasopoulos, Ziwei Zhu
AAAI/ACM Conference on AI, Ethics, and Society; Oral [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
North American Chapter of the Association for Computational Linguistics; Poster [NAACL '24]
- [C.3] **SALSA: Salience-Based Switching Attack for Adversarial Perturbations in Fake News Detection Models** [PDF | Code]
Chahat Raj, Anjishnu Mukherjee, 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 Main '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 AI, Ethics, and Society [AIES '23]

Teaching

-
- > **CS 678 Advanced Natural Language Processing, GMU** *Teaching Assistant; Outstanding GTA Award* Fall '23, Spring '24
 - > **CS 112 Introduction to Python Programming, GMU** *Teaching Assistant* Fall '22, Spring '23

Student Mentorship

-
- > **Lauren Chen** [Syosset HS, Summer 2026] Multilingual and multi-attribute cultural steering for LLMs
 - > **Hasini Jasthi** [TJHSST, 2026–2027] Accent and dialect disparities in voice assistants, focusing on South and Southeast Asian English
 - > **Anubhav Maity, Som Shekhar Sharma, and Akshat Sharma** [IIT Roorkee, Spring 2026] Cultural adaptation for Indian food contexts; final-year project
 - > **Mamnuya Rinki** [M.S. thesis, GMU, 2024–2025] South Asian bias in multilingual, open-ended LLM generations; thesis defended Nov. 2025
 - > **Sharanya Garika and Aksh Patel** [GMU, Fall 2024] Intersectional bias across Indo-Aryan languages; Advanced NLP course project
 - > **Samruddhi Deshmukh and Abhi Krishna** [GMU, Fall 2025] Auditing cultural representation in text-to-image models; AI, Power & Society course project

Awards

- > **Distinguished Academic Achievement Award, GMU '25** M.S. in Computer Science
- > **Outstanding Reviewer, EMNLP '24** Recognition for peer review
- > **Outstanding GTA Award, GMU '24** Advanced NLP (CS 678)
- > **Summer Graduate Research Award, GMU '22** \$8,500; multilingual socio-cultural bias evaluation
- > **DAAD WISE Scholarship '20** \$3,000; real-time collision-detection research in computer graphics

Talks

- > **Crossroads of Continents** [\[Slides\]](#) WACV oral presentation Mar. 2025, Tucson, AZ
- > **Tutorial: Diffusion Models** [\[Slides\]](#) Advanced NLP guest lecture, GMU Nov. 2024, Fairfax, VA
- > **Multilingual Socio-cultural Biases** [\[Slides\]](#) University of Toronto invited speaker series Sep. 2024, Remote
- > **Global Gallery** [\[Slides\]](#) MASC-SLL, Johns Hopkins University May 2024, Baltimore, MD
- > **Cross-Lingual Biases and Cultural Understanding in LLMs** [\[Slides\]](#) Notre Dame NLP seminar Feb. 2024, Remote

Academic Service

- > **Reviewer** COLM and AIES '26; ACL Rolling Review '23–'26 (ACL, EMNLP, NAACL, EACL)
- > **Sub-reviewer** COLM and AIES '24–'25
- > **Workshop reviewer** AfricaNLP, DravidianLangTech, EACL SRW, NewInML, WiNLP, and LT-EDI ('22–'26)
- > **Journal reviewer** ACM Transactions on Intelligent Systems and Technology (TIST), '26