Anjishnu Mukherjee

PhD in Computer Science | George Mason University

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

RESEARCH INTERESTS

My research focuses on natural language processing, particularly multilingual and multicultural large language models. I study how pretraining data shapes the cultural knowledge of LLMs, and develop methods to evaluate and mitigate cultural and social biases in language and vision models. My work has been published at EMNLP, COLM, NAACL, ICML, WACV, and AIES.

EDUCATION

May 2027 – Aug 2022 | **George Mason University** | **Fairfax, VA**
Ph.D. in Computer Science, *GPA: 3.93/4*
Advisors: Antonios Anastasopoulos, Ziwei Zhu

Dec 2025 – Aug 2022 | **George Mason University** | **Fairfax, VA**
M.S. in Computer Science, *GPA: 3.93/4* Distinguished Academic Achievement Award

Jun 2021 – Jul 2017 | **Indian Institute of Engineering Science and Technology (IIEST)** | **Kolkata, India**
B.Tech. in Computer Science, *CGPA: 9.63/10* Convocation Medal (1st in class)

EXPERIENCE

Aug 2026 – May 2024 | **Amazon** · 3 internships | **Seattle, USA**
Applied Scientist Intern | Mentors: Brian Balagot (2026), Wenqin Hu (2024, 2025)
- 2026 *Grounded summaries of product attributes discussed in returns and reviews*
- 2025 *Predicting product return trends with LLMs for A/B testing experiments*
- 2024 *Grounded explanations for black-box size recommendation systems using LLMs*

Aug 2022 – Jun 2020 | **Google Summer of Code** · 2 summers | **Remote**
Developer | Mentors: Ryan Mullins (2022), Ryan Curtin and Sreenik Seal (2020)
- 2022 *Feature attributions and salience methods for Google PAIR's Language Interpretability Tool*
- 2020 *Computer vision functionality and loss functions for mlpack v4*

May 2022 – Jul 2021 | **Wells Fargo** | **Bangalore, India**
Program Associate
- *Adoption dashboards for Teams and Zoom; automated configuration checks for conferencing devices*

SELECTED PUBLICATIONS

[C.10] **Double Trouble: Bilingual Pretraining Leaves Language-Conditioned Effects in Shared-Language Representations**
Anjishnu Mukherjee, Ziwei Zhu, Antonios Anastasopoulos
EMNLP 2026 (Main) Outstanding Paper Award

[C.9] **MAPLE: Metadata Conditioned LLM Pretraining for Locale-Aware Question Answering**
Anjishnu Mukherjee, Ziwei Zhu, Antonios Anastasopoulos
EMNLP 2026 (Main) Oral

[C.8] **Lost in the Tower of Babel: The Adverse Effects of Incidental Multilingualism in LLMs**
Anjishnu Mukherjee, Chutong Meng, Antonios Anastasopoulos
COLM 2026

[C.7] **Knowing Bias, Doing Better: Mitigating Social Bias in LLMs via Know-Bias Neuron Enhancement**
Jinhao Pan, Chahat Raj, **Anjishnu Mukherjee**, Sina Mansouri, Bowen Wei, Shloka Yada, Ziwei Zhu
ICML 2026

[C.6] **Purdah and Patriarchy: Evaluating and Mitigating South Asian Biases in Open-Ended Multilingual LLM Generations**
Mamnuya Rinki, Chahat Raj, **Anjishnu Mukherjee**, Ziwei Zhu
TrustNLP at ACL 2026

[C.5] **Crossroads of Continents: Automated Artifact Extraction for Cultural Adaptation with Large Multimodal Models**
Anjishnu Mukherjee, Ziwei Zhu, Antonios Anastasopoulos
WACV 2025 Oral

[C.4] **BiasDora: Exploring Hidden Biased Associations in Vision-Language Models**
Chahat Raj, **Anjishnu Mukherjee**, Aylin Caliskan, Antonios Anastasopoulos, Ziwei Zhu
Findings of EMNLP 2024

[C.3] **Global Gallery: The Fine Art of Painting Culture Portraits through Multilingual Instruction Tuning**
Anjishnu Mukherjee, Aylin Caliskan, Ziwei Zhu, Antonios Anastasopoulos
NAACL 2024

[C.2] **Breaking Bias, Building Bridges: Evaluation and Mitigation of Social Biases in LLMs via Contact Hypothesis**
Chahat Raj, **Anjishnu Mukherjee**, Aylin Caliskan, Antonios Anastasopoulos, Ziwei Zhu
AIES 2024 Oral

[C.1] **Global Voices, Local Biases: Socio-cultural Prejudices across Languages**
Anjishnu Mukherjee*, Chahat Raj*, Ziwei Zhu, Antonios Anastasopoulos
EMNLP 2023 (Main)

AWARDS

Outstanding Paper Award	*EMNLP, for Double Trouble*	2026
Dissertation Completion Grant	*GMU, $12,500, Spring 2027*	2026
Honorable Mention, Oral Presentation	*GMU School of Computing PhD Research Symposium*	2026
Distinguished Academic Achievement Award	*M.S. in Computer Science, GMU*	2026
Outstanding Reviewer	*EMNLP*	2024
Outstanding GTA Award	*GMU, CS 678 Advanced NLP*	2024
Convocation Medal	*1st in B.Tech. Computer Science and Technology, IIEST*	2022
Summer Graduate Research Award	*GMU, $8,500*	2022
DAAD WISE Scholarship	*$3,000*	2020

TEACHING

CS 678 Advanced NLP	*Teaching Assistant, GMU*	Fall 2023, Spring 2024
CS 112 Introduction to Python	*Teaching Assistant, GMU*	Fall 2022, Spring 2023

STUDENT MENTORSHIP

Fatima Fellowship	*Mentor, Fatima Institute for Global AI Research*	Fall 2026
Lauren Chen	*Cultural steering for LLMs (high school)*	Summer 2026
Hasini Jasthi	*Accent and dialect disparities in voice assistants (high school)*	2026–2027
A. Maity, S. S. Sharma, A. Sharma	*Cultural adaptation for Indian food (undergraduate)*	Spring 2026
Samruddhi Deshmukh, Abhi Krishna	*Text-to-image cultural representation (M.S.)*	Fall 2025
Mamnuya Rinki	*South Asian bias in multilingual LLMs (M.S. thesis)*	2024–2025
Sharanya Garika, Aksh Patel	*Intersectional bias (M.S.)*	Fall 2024
Diwita Banerjee	*Cuisine biases and chart hallucinations in LLMs (M.S.)*	Summer 2024

TALKS

MAPLE	*Lightning talk, GMU SoC PhD Symposium (15 of 100+ selected)*	Nov 2026
LLM Agents	*NLP (CS 478) guest lecture, GMU*	Oct 2026
MAPLE	*EMNLP oral presentation*	Oct 2026
Measuring and Modeling Cultural Representations	*Advanced NLP guest lecture, GMU*	Sep 2026
Crossroads of Continents	*WACV oral presentation*	Mar 2025
Tutorial: Diffusion Models	*Advanced NLP guest lecture, GMU*	Nov 2024
Multilingual Socio-cultural Biases	*University of Toronto invited speaker series*	Sep 2024
Global Gallery	*MASC-SLL, Johns Hopkins University*	May 2024
Cross-Lingual Biases and Cultural Understanding in LLMs	*Notre Dame NLP seminar*	Feb 2024

ACADEMIC SERVICE

Ethics reviewer	*ACL Rolling Review, October 2026 cycle*	2026
Reviewer	*ACL Rolling Review (ACL, EMNLP, NAACL, EACL), COLM, AIES*	2023–2026
Sub-reviewer	*COLM, AIES*	2024–2025
Workshop reviewer	*AfricaNLP, DravidianLangTech, EACL SRW, NewInML, WiNLP, LT-EDI*	2022–2026
Journal reviewer	*ACM Transactions on Intelligent Systems and Technology*	2026