

SAGNIK MUKHERJEE

Computer Science, University of Illinois Urbana-Champaign

sagnikm3@illinois.edu | [Website](#) | [Scholar](#)

RESEARCH INTERESTS

My research studies how large language models can learn to make decisions scalably and efficiently through reinforcement learning. On the *computational* side, I uncover emergent structure in on-policy training, such as sparse, predictable parameter updates^{[2][5]} and a largely preserved spectrum^[3], and turn it into sample-efficient optimizers^{[1][3]} and efficient serving. On the *algorithmic* side, I design inductive biases inspired by human learning that make learning more sample-efficient, from single-step reasoning^{[4][7]} to long-horizon decision making that requires foresight and information seeking.

EDUCATION

University of Illinois Urbana–Champaign

Ph.D. Student in Computer Science, GPA: 4.0/4.0

Advisors: Prof. Dilek Hakkani-Tür, Prof. Hao Peng

Urbana, USA

Aug. 2024 – Present

Indian Institute of Technology, Kanpur

B.Tech in Electrical Engineering, GPA: 9.5/10

India

Aug. 2017 – Aug. 2021

SELECTED AWARDS AND FELLOWSHIPS

ICML Oral (Top 0.7%)	2026
NeurIPS Scholar Award	2025
Academic Excellence Award, IIT Kanpur (all years)	2017–2021
IUSSTF-Viterbi Research Fellowship, USC (top 15 from India)	2021
SURGE Research Fellowship, IIT Kanpur	2020
JBNSTS Junior Scholar for Basic Sciences	2016
Top 1%, National Standard Examination in Physics	2016
Rank 33 out of 150,000 WBEE	2017

SELECTED PUBLICATIONS

* Denotes Equal Contribution.

- [1] S. Mukherjee, L. Yuan, P. Jayasinha, D. Hakkani-Tür, H. Peng, “Do We Need Adam? Surprisingly Strong and Sparse RL with SGD in LLMs”, ICML ’26, Oral (top 0.7%).
- [2] S. Mukherjee, L. Yuan, D. Hakkani-Tür, H. Peng, “Reinforcement Learning Finetunes Small Subnetworks in Large Language Models”, [NeurIPS 2025].
- [3] H. Zhu*, W. Cong*, Z. Sha, S. Mukherjee, X. Song, . . . , Z. Wang, “ISO: An RLVR-Native Optimization Stack”, Preprint.
- [4] S. Mukherjee, T. Chakraborty, D. Hakkani-Tür, H. Jhamtani, J. Andreas, “VarPO: Variational Policy Optimization with Privileged Information”, In progress.
- [5] Y. Zhang, S. Mukherjee, H. Peng, “Understanding LLM Parameter Update Sparsity through the Lens of Fisher”, Submitted to ICLR 2027.
- [6] S. Shekhar, S. Mukherjee, J.Y. Zhang, T. Zhang, “VASR: Variance-Aware Systematic Resampling for Reward-Guided Diffusion”, [NeurIPS 2026].
- [7] S. Mukherjee*, A. Chinta*, T. Kim, T.A. Sharma, D. Hakkani-Tür, “Premise-Augmented Reasoning Chains Improve Error Identification in Math Reasoning with LLMs”, [ICML 2025].
- [8] R.G. Reddy*, S. Mukherjee*, J. Kim*, Z. Wang*, D. Hakkani-Tür, H. Ji, “Infogent: An Agent-based Framework for Web Information Aggregation”, [NAACL 2025].
- [9] S. Mukherjee*, M.F. Adilazuarda*, S. Sitaram, K. Bali, A.F. Aji, M. Choudhury, “Cultural Conditioning or Placebo? On the Effectiveness of Socio-Demographic Prompting”, [EMNLP 2024].
- [10] M.F. Adilazuarda*, S. Mukherjee*, P. Lavania, S. Singh, A. Dwivedi, A.F. Aji, J. O’Neill, A. Modi, M. Choudhury, “Towards Measuring and Modeling ‘Culture’ in LLMs: A Survey”, [EMNLP 2024].

- [11] S.R. Nangi, A. Tyagi, J. Mundra, S. Mukherjee, S. Raj, A. Garimella, N. Chhaya, “AUTO-SUMM: Automatic Model Creation for Text Summarization”, [EMNLP 2022].
- [12] J. Yan, Y. Xiao, S. Mukherjee, B.Y. Lin, R. Jia, X. Ren, “On the Robustness of Reading Comprehension Models to Entity Renaming”, [NAACL 2022].
- [13] N.R. Shirude*, S. Mukherjee*, T. Shandhilya, A. Mukherjee, A. Modi, “Classification for Lexical Complexity Regression Task”, [SemEval 2021].
- [14] J. Mundra*, R. Gupta*, S. Mukherjee*, “Multi-task Learning and Transformer Finetuning for Emotion Classification & Empathy Prediction”, [WASSA, EACL].

Patent

Auto-Creation of Custom Models for Text Summarization [\[link\]](#)

INVITED TALKS

Optimization Geometry in LLM post-training, host: Tong Zhang, UIUC *Sept 2026*
ICML Oral Presentation *July 2026*

INTERNSHIP EXPERIENCE

Research Intern, Microsoft Frontier Tuning *2026*
Hosts: Jacob Andreas and Harsh Jhamtani.

Research Intern, Microsoft Research Redmond *May 2025 – Aug. 2025*
Hosts: Peeyush Kumar and Jennifer Neville.
Studied how LLMs fail under varying degrees of underspecification in multi-turn conversations. Paper in progress.

Research Associate, MBZUAI *Mar. 2024 – Jul. 2024*
Cultural awareness in LLMs and efficacy of socio-demographic prompting. 2 first-authored papers at EMNLP 2024.

Research Intern, Adobe Research *May 2020 – Jul. 2020*
Hosts: Niyati Chhaya.
Combined AutoML techniques with NLP for near-SOTA text summarization via auto-generated models.

Research Intern, USC INK Lab *May 2021 – Aug. 2021*
Hosts: Xiang Ren.
Robustness of reading comprehension models under entity replacement.

INDUSTRY EXPERIENCE

Data and Applied Scientist, Microsoft Inc. (Bing Ads) *Nov. 2022 – Feb. 2024*
Query understanding, dense and generative retrieval. Shipped 2.5% revenue lift on EMEA and APAC markets.

Associate ML Scientist I, Wadhvani AI *Jul. 2022 – Nov. 2022*
AI for healthcare—TB detection from cough sound.

Member of Technical Staff, Adobe *Jul. 2021 – Jul. 2022*
App recommendation system for Adobe Desktop App.

PROFESSIONAL ACTIVITY

Conference Reviewer

ARR (2024–2025), ICLR (2025), TMLR (2026–Present)

Student Volunteer

TACL (Sept. 2025–Present)

Teaching

- TA, *Advanced Topics in NLP* (Prof. Dilek Hakkani-Tür), UIUC, Fall 2025

- Tutor, *Fundamentals of Computing* (ESC101), IIT Kanpur, 2020–2021
- TA, *Machine Learning for Signal Processing* (EE698V), IIT Kanpur