Ruhma Mehek Khan

(404)496-0699 | mehek.ruhma@gmail.com | www.ruhmakhan.com | linkedin.com/in/ruhmakhan

Education

Georgia Institute of Technology, Atlanta Aug. 2024 - May. 2026
Master's in Computer Science, Specializing in Machine Learning & Artificial Intelligence *CGPA: 4.0/4.0*
Teaching Assistant: CSE6140–Computer Science & Engineering Algorithms, CS3751– Intro to UI Design

IIIT Delhi — *Dean's list for academic excellence* Aug. 2018 - May. 2022
BTech (with Honors) in Computer Science and Social Sciences *CGPA: 3.95/4.0, Department rank 1*
Teaching Assistant: CSE343– Machine Learning, CSE222– Algorithm Design and Analysis,

Technical Skills

Languages & Tools: Java, Python, JavaScript, TypeScript, C, SQL, R, React, Node.js, Flask, Django, PyTorch, TensorFlow, LangChain, Git, Docker

ML/AI: Large Language Models, AI Agents, RAG, Reinforcement Learning, Deep Learning, NLP, Multimodal Learning, Model Training, Evaluation, Data Visualization, Convex Optimization

Experience

Amazon Sunnyvale, California, USA
Applied Scientist Jul. 2026 – Present

Applied Science Intern *May. 2025 – Aug. 2025*
- Built an end-to-end **AI agent** using AWS Bedrock to translate causal ML model outputs into actionable financial insights, reducing manual analysis from days to minutes with **100% numeric accuracy**
- Designed modular **LLM-powered workflows** with programmatic reasoning and parallel processing; engineered a **RAG-based grounding pipeline** with semantic chunking and automated remediation, and built automated ground-truth validation with an **LLM-as-judge** evaluation framework

Microsoft Research Redmond, USA
AI Research Engineer Sep. 2023 – Jul. 2024
- Evaluated LLMs for the hard-constrained task of **program synthesis by example (PBE)** and designed techniques to improve robustness to prompt sensitivity
- Developed a neuro-symbolic pipeline to **repair inaccuracies in LLM-generated code**, solving over **30%** of previously unsolved benchmark tasks across Python and Excel
- Khan, R., Gulwani, S., Le, V., Radhakrishna, A., Tiwari, A., Verbruggen, G. *LLM-Guided Compositional Program Synthesis.* arXiv:2503.15540, 2025. [Link]

Design & Intelligence Lab, Georgia Tech Atlanta, Georgia, USA
Graduate Research Assistant Aug. 2024 – Dec. 2025
- Developed a **TMK-based self-model** for an AI system for scientific modeling, enabling LLMs to generate **metacognitive self-explanations** grounded in an AI agent's task, method, and knowledge structure
- Designed an **iterative reasoning pipeline** composing explanations across TMK components with traceable evidence; work resulted in a **Springer book chapter**

Adobe Noida, India
Software Development Engineer Jun. 2022 – Sep. 2023
- Spearheaded the revamping of commenting functionality for a web-based document review workflow used by **100+ enterprise customers**, including live synchronization of comments in multi-user environments
- Built an **open-source customization plugin** for Adobe Experience Manager, enabling enterprise users to customize the platform without legacy-code knowledge; received a **Spot Award**

Adobe Research Bangalore, India
Research Intern May. 2021 – Aug. 2021
- **Filed a patent** and improved F1 by **50%** for Adobe Acrobat's accessibility-tagging AI pipeline by developing an algorithm to incorporate and propagate human feedback

Selected Research Projects

SImpL: Self-Improving LLMs via Implication Learning [Report] 2025
- Built a self-improving LLM framework that distills context into reusable, query-independent implications, matching SEAL's accuracy with **20×+ lower training time**
- Improved reasoning accuracy by up to **10.5 percentage points** across Llama and Qwen models; explored LoRA-based test-time training and GRPO/DRGRPO for implication learning

Risk-Aware Actor-Critic Reinforcement Learning [Report] 2025
- Implemented two risk-aware extensions of **TD3** using distributional critics and **CVaR** objectives to trade off expected performance against worst-case outcomes
- Evaluated quantile-regression and Gaussian variants across seven MuJoCo environments; Gaussian TD3 achieved the highest return in **4/7 environments**

Awards & Achievements

JN Tata Scholarship, awarded to ∼100 students across India 2024, 2026
Adobe India WiT Scholarship, 6 winners pan India (among 30k+ applicants) 2021
All Round Award and Institute Medal, awarded on graduation for best academic and all round performance 2022
Grace Hopper Student Scholar, awarded by Anita B.Org, USA 2021

Leadership & Community Involvement

Program Committee Member, Agentic & Generative AI for E-Commerce Workshop, ACM RecSys Aug. 2026
Mentor, Princeton NVIDIA Open Hackathon Jun. 2026
Director – Career Development, MTC Georgia Tech. Leading career building programs Jul. 2024 – Apr. 2025
Mentor, Anita B.Org. Helped sophomores secure internships and build profiles Oct. 2022 – Jan. 2023
Core Team, Women in Machine Learning and Data Sciences. Organised first ever mentorship program Feb. 2021 – Jan. 2022
Co-Founder & President, Lean In IIITD - Scaled mentorship community to 180+ members in a year Jul. 2020 – Jun. 2022