

Aman Gupta

+1 415-378-8237 | amangupt@cs.cmu.edu | linkedin.com/amangupta27 | Pittsburgh, PA

EDUCATION

Carnegie Mellon University – School of Computer Science <i>Master of Machine Learning (MSML)</i> Research Advisor: Prof. Graham Neubig Relevant Coursework: Advanced Machine Learning, Probabilistic Graphical Models, ML Systems, Reinforcement Learning	August 2023 - December 2024 4.0/4.0
Indian Institute of Technology Delhi <i>Bachelor of Computer Science & Engineering</i> Relevant Coursework: NLP, Deep Learning, Operating System, Distributed Computing, Software Design Practices	June 2019 - May 2023 9.75/10

PROFESSIONAL EXPERIENCE

Amazon <i>Applied Scientist Intern</i>	Seattle, USA May 2024-August 2024
<ul style="list-style-type: none">Built a scalable multi-agent system with a central orchestrator LLM, enhancing Amazon's Customer ServiceEvaluated the system by applying retrieval-based augmentations to the dialog system on MultiWOZ benchmarkImplemented an ensemble of fine-tuned Mistral-7B, Flan-T5-large, and Prompted Claude Sonnet-3.0 for the dialog state tracking and response generation pipelines, achieving SOTA 96% Inform and 88% Success rates	
Microsoft <i>Cloud & AI Intern</i>	Hyderabad, India June 2022-July 2022
<ul style="list-style-type: none">Improved restore and backup cycles of user data in the fundamental backend layer of Azure cloudDesigned dynamic batch downloads supporting scalability, resilience to failed API calls, and crash recoverability	
LimeChat <i>NLP Intern</i>	Bangalore, India June 2021-July 2021
<ul style="list-style-type: none">Analyzed e-commerce chatbot-customer conversations to identify model errors and generate augmented dataDesigned a Redis-based pipeline to enable clients to adjust bot responses for specific customer requests	

RESEARCH EXPERIENCES & PROJECTS

Personalizing Chatbots Prof. Graham Neubig & Snap Research	August 2023 - May 2024
<ul style="list-style-type: none">Developed retrieval-augmented methods for personalized dialogue generation, analyzed performance gainsEvaluated the impact of extended conversation history on LLM-generated responses by assessing coherence and personalization using advanced metrics like BERTScore, UniEval, and chrF. Work under review in COLING 2025	
LLM Enhanced Graphical Learning Prof. Andrej Risteski	February 2024 - May 2024
<ul style="list-style-type: none">Used embeddings enriched by encoder language models and LLMs to demonstrate performance improvementsBenchmarked our approach on 4 different test-attributed graph datasets and conducted detailed ablation studies	
Natural Language to Machine Instructions Prof. Mausam	February 2023 - June 2023
<ul style="list-style-type: none">Fine-tuned several language models (GPT-2, BERT, T5) to develop a Task-Oriented Dialog (TOD) system, transforming natural language inputs into precise user intents and related key-value pair dataSystematically hyper-tuned weight and generation parameters leading to a match accuracy of 0.89	
Parallel & Distributed Computing Prof. Subodh Kumar	January 2022 - March 2022
<ul style="list-style-type: none">Designed and implemented parallel, distributed C++ programs for MergeSort, Twitter's RWR (PageRank), and HNSW algorithms. Used OpenMP and MPI constructs for threads and tasks.	

SKILLS & CERTIFICATES

Academic Achievements: Expert Rated on Codeforces, Top-10 Institute Rank in IIT Delhi on Final GPA, Institute Merit Award - IIT Delhi (1st in Freshmen), Gold Medal in National Chemistry Olympiad (Top-35 in India), All India Rank-14 KVPY
Publications: ACM Sigmoid - 2024, ACM Compass - 2023 (Co-First Author), Nature Journal - 2022 (Co-First Author)
Programming & Software: Python, C++, SQL, Pytorch, PySpark, CUDA, Hadoop, Ray, Docker, Redis, AWS, Azure
Mathematics: Probability & Stochastic Processes, Statistics, Linear Algebra, Theory of Computation, Adv. Calculus