Akshita Gupta

akshita24680@gmail.com | (412) 214-2506 | https://www.linkedin.com/in/aks-gupta/ | https://github.com/aks-gupta

---

Education

Carnegie Mellon University — Pittsburgh, Pennsylvania
Master of Computational Data Science (CGPA: 4.07/4.0) — **December 2025**
Relevant Coursework: Foundations of Computational Data Science, Advanced NLP, Introduction to Machine Learning, Interactive Data Science, Multimodal Machine Learning, AI Engineering, Cloud Computing, Search Engines, Generative AI, Deep Reinforcement Learning

R. V. College of Engineering — Bangalore, India
Bachelor of Engineering, Computer Science and Engineering (CGPA: 9.27/10) — **June 2023**
Relevant Coursework: Data Structures, Advanced Algorithms, Artificial Neural Networks, AIML, Mathematical Modelling

---

Skills

Programming Languages: Python, C++, C, Java, SQL, React
Tools and Frameworks: Pytorch, scikit-learn, Pandas, HuggingFace, OpenCV, Elasticsearch, Git, Flask, Tableau, AWS, Azure, GCP, MLFlow
Concepts: NLP, LLMs, Generative AI, Multimodal ML, CV, Deep Learning, Responsible AI, Recommender Systems

---

Experience

Amazon Web Services — New York, New York
AI/ML Software Development Engineering Intern — **May 2025 – August 2025**
- Designed and developed a Gen-AI application to internally dogfood Automated Reasoning checks for AWS Bedrock Guardrails.
- Built a corrective pipeline to reduce hallucinations in LLM responses using formal methods.
- Delivered actionable insights from application, through prompt tuning, and tracking user struggles in using AR checks.

Carnegie Mellon University — Pittsburgh, Pennsylvania
Teaching Assistant – Advanced Natural Language Processing (PhD) — **January 2025 - Present**
- Assisted in designing, grading, and mentoring students on several graduate level NLP assignments and projects.

Cisco — Bangalore, India
Software Engineer — **August 2023 – July 2024**
- Proposed a proof-of-concept leveraging eBPF and Machine Learning to obtain router packet path insights for Cisco ASR9k.
- Scaled the MAP-T feature on Cisco ASR9k switches, increasing capacity by 10x from 200 to 2k entries.
- Resolved several critical sanity failures in the nightly test suite, improving nightly test pass rate by 12%.

Technical Undergraduate Intern — **January 2023 – June 2023**
- Created a data pipeline to transmit network telemetry data from routing tables to Elasticsearch via the Kafka Queueing service.
- Designed optimized Elasticsearch queries to analyze over 100,000 synchronous entries, enabling differential insights through time-series forecasting models.

Samsung Research — Bangalore, India (Remote)
Research Intern — **June 2021 – January 2022**
- Translated the ATIS dataset with 4633 data points to "Hinglish" Code-Mixed Language for Intent Classification and Slot Tagging.
- Validated the dataset with ten-fold cross-validation and obtained an intent accuracy of 99.6% on BERT.

Projects

Theory of Mind for Explainable AI (Ongoing**)**
- Working on understanding theory of mind in LLMs and applying chain-of-thought for faithfulness and commonsense in biased data.

Exploring the use of test-time scaling for Multimodal Reasoning Tasks
- Compared scaling methods like chain-of-thought reasoning and synthetic data generation in multimodal datasets like MathVista.

Debiasing Large Language Models through Casual-Guided Active Learning
- Improved performance on novel evaluation metrics by 5% through prompt finetuning for position bias in MT-Bench dataset.

Retrieval Augmented Generation (RAG) Model using LLMs
- Developed an end-to-end question-answering RAG Model on Llama 3.2 for Pittsburgh events, using Langchain and Ollama.

Selected Publications

"Multimodality in Online Education: A Comparative Study." Multimedia Tools and Applications (Q1, 2024)
- Proposed a decision-level majority fusion model for multiple modalities - Facial Expressions, Posture, Speech, Eye-tracking.

"Joint Intent Classification and Slot Tagging on Agricultural Dataset for Indic Languages." ICACCS (2023)
- Created a QA dataset for farmers with over 2400 entries, obtaining an accuracy of 93.89% and 98.32% on LSTM and BERT.

"Comparison of Perplexity Scores of Language Models for Telugu Data Corpus in the Agricultural Domain." ICICCS (2024)
- Compared the perplexity scores of 1 million "Telugu" web-scraped tokens across n-gram, GPT-2, and LSTM models.

Leadership and Activities

Student Placement Coordinator, RV College of Engineering — September 2022 – August 2023
Member, Coding Club, RV College of Engineering — August 2019 – June 2023
Director and Actor, Circle of Acting at RV College of Engineering — January 2023 – May 2023