Aashay Pinkesh Mehta

aashaypm@cs.cmu.edu • Aashay Mehta • aashay-mehta

Experience

Capital One **New York, NY**
Principal Data Scientist, AI Foundations *July 2025 – Present*
- Building an SDK to standardize LLM agent development across the company, streamlining integration with internal infrastructure and providing capabilities such as observability and secure code execution.
- Leading a team of six engineers to deliver iterative improvements to Capital One's first customer-facing AI product.
- Leveraged supervised fine-tuning (SFT) and preference optimization to distill a 70B-parameter LLM into a 10B-parameter model, achieving a 3x reduction in inference latency and a 10x increase in product throughput.

Senior Data Scientist, AI Foundations *June 2024 – July 2025*
- Built and deployed multi-agent LLM systems, including travel assistants and banking support bots.
- Simulated multi-turn conversations to generate synthetic test sets for benchmarking chatbot performance.
- Designed and implemented a scalable, configurable multi-agent AI architecture that enabled rapid prototyping and component reuse.

Carnegie Mellon University **Pittsburgh, PA**
Research Assistant, ML Department *June 2023 – May 2024*
- Studied autoregressive sampling techniques for verifier-assisted constrained generation in LLMs.
- Established a statistical perspective on the optimal masking ratio for training BERT-like language models.

Research Assistant, Auton Lab *Sept 2022 – May 2023*
- Combined supervised prediction with reinforcement learning for trading in inefficient markets with high friction costs.

EdgeVerve Systems **Bengaluru, India**
Member of Technical Staff *Nov 2020 – May 2022*
- Enhanced digital payment capabilities by developing REST APIs, and created a UI in React to demonstrate their use.
- Implemented a two layer caching mechanism (with synchronization via Redis), reducing API response times by 85%.

Max Planck Institute for Intelligent Systems **Tuebingen, Germany**
Bachelor's Student Intern *Jan 2020 – Aug 2020*
- Developed reinforcement learning algorithms robust to erroneous descriptions of environments, leading to up to 40% increase in cumulative return over their non-robust counterparts.

Education

Carnegie Mellon University - School of Computer Science **Pittsburgh, PA**
Master of Science in Machine Learning | GPA: 4.08 *Dec 2023*
Selected Coursework: Visual Learning and Recognition, Probabilistic Graphical Models, Statistical Machine Learning, Convex Optimization, Probability and Statistics, Deep Reinforcement Learning, Machine Learning in Practice

Birla Institute of Technology and Science (BITS), Pilani **Pilani, India**
Bachelor of Engineering in Computer Science | GPA: 9.55/10 *July 2020*

Skills

Languages Python, Java, C, JavaScript/TypeScript
Tools & Frameworks PyTorch, LangGraph, vLLM, HuggingFace, scikit-learn, NumPy, Pandas, Git, Tensorflow

Selected Publications

- Botta, E.*, Li, Y.*, Mehta, A.*, Ash, J., Zhang, C., Risteski, A. (2025). On the Query Complexity of Verifier-Assisted Language Generation. *International Conference on Machine Learning (ICML)*.
- Li, Y., Kirchmeyer, A.*, Mehta, A.*, Qin, Y.*, et al. (2024). Promises and Pitfalls of Generative Masked Language Modeling: Theoretical Framework and Practical Guidelines. *International Conference on Machine Learning (ICML)*.
- Mehta, A.*, Jain, Y. R.*, Kemtur, A., et al. (2022). Leveraging Machine Learning to Automatically Derive Robust Decision Strategies from Imperfect Knowledge of the Real World. *Computational Brain & Behavior*.
- Kemtur, A.*, Jain, Y. R.*, Mehta, A., et al. (2020). Leveraging Machine Learning to Automatically Derive Robust Planning Strategies from Biased Models of the Environment. *CogSci 2020*.