

Aditya Oke

Pittsburgh, PA | aoke@cs.cmu.edu | linkedin.com/in/aditya-oke-99b74717a | github.com/oke-aditya

RESEARCH STATEMENT

My research focuses on **LLM safety, robustness, and interpretability** for **agentic and decision-making systems**. I develop **evaluation and training methods** that improve **factual reliability, reasoning consistency, and behavior in real world**.

EDUCATION

Carnegie Mellon University <i>Master of Science in Computational Data Science; CGPA: 3.86/4.0</i>	Pittsburgh, PA Aug 2025 – May 2027
– Coursework: LLM Methods and Applications, Machine Learning, Cloud Computing, LLM Systems	

Vellore Institute of Technology <i>Bachelor of Technology in Computer Science and Engineering; CGPA: 9.20/10.0</i>	Vellore, India Aug 2018 – May 2022
– Focus: Artificial Intelligence, Image Processing, Operating Systems, Distributed Systems	

RESEARCH EXPERIENCE

Amazon Nova Challenge <i>PyTorch, LLM Agents, AI Safety</i>	Feb 2026 – Present
– Built red-teaming LLM agents to detect vulnerable code generation in production tasks.	
– Creating safety benchmarks for adversarial coding prompts , evaluating exploitability and model robustness.	
Graduate Researcher: Li Lab (Advisor: Prof. Lei Li) <i>Carnegie Mellon University</i>	Sep 2025 – Present Pittsburgh, PA
– Investigating hallucination, attribution, and robustness of large language models including GPT-4, Gemini, and DeepSeek under distribution shift.	
– Developing automatic and human-aligned evaluation frameworks for factuality, reasoning consistency, and failure analysis in scientific text generation.	
Probing Reasoning in Large Language Models <i>PyTorch, Hugging Face, LLaMA 3</i>	Nov 2025 – Dec 2025
– Performed knowledge distillation of LLMs into smaller models while preserving chain-of-thought and multi-step reasoning .	
– Analyzed generalization, calibration, and faithfulness of reasoning traces across model scales.	
Torchvision Contributor <i>PyTorch, Python</i>	Aug 2020 – Jun 2023
– Contributed to PyTorch Torchvision , building reproducible vision benchmarks and model evaluation utilities .	
– Co-Contributed SwinTransformer3D for large-scale video research used by academic and industry labs.	

PUBLICATIONS

- D Andre, Tufts B, Oke, A, et al. **Sem-Detect: Semantic Level Detection of AI Generated Peer-Reviews**. Submitted, International Conference on Machine learning (ICML) 2026
- Oke, A. et al. **A Novel Approach to Encrypt Data Using Deep Neural Networks**. Springer ICADML, 2022.
- Oke, A. et al. **Real-time Face Mask Detection on Edge IoT Devices**. IET Computer Engineering, 2021.
- **Deep Learning with PyTorch Lightning**. Technical Reviewer, Packt Publishing, 2022.

PROFESSIONAL EXPERIENCE

Software Engineer II <i>JPMorgan Chase & Co.</i>	Jan 2022– Aug 2025 Bengaluru, India
– Designed and deployed distributed Spark pipelines processing multi-terabyte financial datasets for real-time analytics.	
– Built Python services and AWS workflows to support compliance, reporting, and internal AI-driven tools.	
– Migrated legacy systems to AWS cloud infrastructure , improving system reliability and reducing operating costs.	

TECHNICAL SKILLS

Languages: Python, C++, Java, SQL
Machine Learning: PyTorch, Hugging Face Transformers, LoRA, PEFT, Fine-Tuning, Model Evaluation
Data & Systems: Apache Spark, AWS, Databricks, Docker, Kubernetes, Linux

AWARDS

Corporate Tech Champion Award <i>JPMorgan Chase</i>	Apr 2025
Early Career Recognition Award <i>JPMorgan Chase</i>	Nov 2024
Engineering Excellence Award <i>JPMorgan Chase</i>	Aug 2023
Innovation Hackathon Winner <i>JPMorgan Chase</i>	Mar 2023