Indiana University
USA
zoher.kachwala@gmail.com
zoher15.github.io/
www.linkedin.com/in/zoher-kachwala/
github.com/Zoher15/

Zoher Kachwala

PhD Candidate & Student Researcher in AI/ML

Education

May 2026 **PhD in Computer Science**, *Indiana University*, USA
GPA: 3.8/4.0

May 2022 **MS in Data Science**, *Indiana University*, USA

May 2019 **MS in Computer Science**, *Indiana University*, USA

May 2017 **B.Tech in Computer Engineering**, *NMIMS*, India

Research Focus & Expertise

PhD candidate advancing the capabilities of frontier **Large Language Models** to solve large-scale, real-world challenges. My collaborative research drives scientific advancement in foundational AI systems through:

- **Scalable LLM Systems**: Designing novel decoding strategies and training methodologies for improved reasoning capabilities, with focus on robust deployment across diverse applications and large-scale inference scenarios.
- **Multimodal AI & Evaluation**: Developing breakthrough prompt-based methods for Vision-Language Models and creating comprehensive evaluation frameworks for next-generation intelligent systems.
- **AI Safety at Scale**: Building production-ready moderation systems through collaborative fine-tuning and deployment of specialized LLMs, addressing real-world safety challenges across hundreds of online communities.

Experienced in the full research-to-production lifecycle, including large-scale distributed training, collaborative development, and applying state-of-the-art techniques to advance foundational research with measurable impact.

Publications & Research Impact

2024 **Advanced Heuristics for LLM Decoding Improve Chain-of-thought Reasoning**, *In Progress*
Proposes breakthrough decoding strategies that enhance **chain-of-thought reasoning** in large-scale LLM deployment, advancing foundational capabilities in complex reasoning tasks with measurable performance improvements.

2024 **Fine-Tuning Specialized LLMs for Large-Scale Community Content Moderation**, *In Progress*
Developing a production-scale moderation framework through collaborative fine-tuning and deployment of specialized LLMs; addresses real-world safety challenges across 500+ communities, demonstrating practical AI system deployment.

2024 **Task-Aligned Prompting Improves Detection of AI-Generated Images in VLMs**, *Under Review* – **NeurIPS 2025**
Introduced breakthrough zero-shot prompting methodology (zero-shot-s^2) achieving 29% F1 improvement in VLM performance, demonstrating scalable alternatives to supervised approaches for next-generation multimodal systems.

2023 **REMATCH: Robust and Efficient Knowledge Graph Matching**, **NAACL 2024**
Developed novel semantic similarity metric with 5x computational efficiency improvement and 5% accuracy gains, alongside comprehensive benchmark (RARE) advancing knowledge graph research infrastructure.

2023 **Midterm Elections: Social Media Dataset**, **ICWSM 2023**
Constructed and released large-scale 30M-post research dataset enabling breakthrough research in information systems, narrative analysis, and misinformation detection across the research community.

2023 **The Inexplicable Efficacy of Language Models**, **ACM XRDS**
Comprehensive survey of LLM capabilities across domains, analyzing production trade-offs and generalization principles for advancing foundational AI research.

Research Projects & Collaborative Work

2023–Present **Large-Scale Content Moderation with Specialized LLMs**, *Indiana University*
Leading collaborative development of production-scale moderation framework using fine-tuned LLMs; addressing real-world safety challenges across 500+ online communities through innovative AI system design.

2019–2023 **Text2Graph: Scalable Knowledge Extraction**, *Indiana University*
Designed collaborative system for extracting structured knowledge representations from text at scale, enabling advanced knowledge-grounded reasoning in large language models.

2021 **Transformer Architecture Analysis**, *Indiana University*
Conducted comprehensive literature review and analysis of foundational Transformer architectures, contributing to understanding of core mechanisms driving breakthrough AI capabilities.

Industry Experience

Summer 2018 **Technology Consultant**, *PwC*
Collaborated on enterprise-scale data integration projects, focusing on automation of validation pipelines and scalable audit-ready logic modeling for large-scale systems.

Collaborative Teaching & Mentorship

2022–2024 **Introduction to Network Science**, *Indiana University*
Led collaborative Python-based tutorials on large-scale network analysis; guided students through graph neural network concepts and dynamic analytics for real-world applications.

2019–2021 **Elements of Artificial Intelligence**, *Indiana University*
Designed scalable autograding systems and provided collaborative support on foundational AI concepts including search, logic, and reasoning for 300+ students annually.

Fall 2020 **Applied Machine Learning**, *Indiana University*
Collaborated on PyTorch-based model development curriculum; emphasized reproducible research practices, scalable inference validation, and deployment-ready system design.

Technical Expertise

Research Areas **Large Language Models, Generative AI, Multimodal Systems, AI Safety, Scalable ML**

ML Frameworks **PyTorch, Hugging Face**, TensorFlow, Scikit-learn, Distributed Training

Programming **Python (expert)**, C++, SQL, Bash, CUDA

Systems & Infrastructure Git, Docker, Linux, Google Cloud Platform, Large-scale Computing

Research Methods Collaborative Research, Model Evaluation, Benchmark Development, Reproducible Science, Large-scale Experimentation