

TOHEED AKHTAR

+91 ***** , *****@gmail.com , linkedin.com/in/toheedakhtar01 , github.com/toheedakhtar

Education

Rajasthan Technical University, Kota, India

B.Tech. in Computer Science & Engineering

2022 – 2026

8.0/10 CGPA

Work Experience

Lossfunk

Mar 2026 – Present

Research Intern (AI Safety)

Bengaluru, Karnataka

- Developing an architecture-agnostic framework for metacognition by studying invariants across frontier LLMs
- Investigating factual correctness in 8B-scale LLMs for the manuscript *Is Factual Correctness Represented in LLMs? A Cross-Dataset Probing Study*.

Chirpn AI

Mar 2025 – Oct 2025

AI/ML Intern

Pune, Remote

- Created **coding, GitHub, and Jira agents** from scratch using CrewAI.
- Owned workflow development and reduced agent token usage by **85%**.
- Implemented **A2A and MCP protocols** and custom tools for GitHub, Jira, and file operations.
- Built and deployed production agentic workflows using **LangGraph and Docker**.
- Fine-tuned Llama and open-source vision models for handwritten-text recognition, achieving **95% accuracy**.

buildspace

Aug 2024 – Sep 2024

Full Stack Developer

San Francisco Bay Area, Remote

- Selected for Nightz & Weekends, a six-week program funded by **Y Combinator and a16z**.
- Built a technology product alongside a global community.
- Independently built and scaled an MVP to more than 1,000 visits.

Research Experience

Is Factual Correctness Represented in LLMs? A Cross-Dataset Probing Study : Studying whether factual correctness is represented consistently across datasets in LLM internal representations. *Manuscript in preparation*.

Higher-Order Alleys [Code] : Found exploratory evidence that a Qwen3.6-27B latent tracks hidden-process degradation under fixed visible output; causal mediation to later judgments remains unresolved.

Uncertainty Circuits : Investigated epistemic uncertainty, including its representation and causal role in model behavior.

\$5 LLM : Replaced the Newton–Schulz iteration with Polar Express in the MuON optimizer; credited on the project's 1B-token marathon leaderboard.

Projects

LLM-from-Scratch: Built a 127M-parameter GPT-2-style model in PyTorch

[Source Code](#)

- Implemented the complete pipeline from **tokenization** and **pretraining** through **fine-tuning** and **evaluation**.
- Trained the **127M-parameter** model on **CPU**.
- Reached **46.8%** of Llama 3 8B's score under the same evaluation setup.

Gemma3-GenZ: Fine-tuned Gemma 3 270M for Gen-Z-style conversations

[Hugging Face](#)

- Curated a synthetic dataset of more than 120 Gen-Z college conversations with slang and casual dialogue.
- Fine-tuned Gemma 3 270M using **PEFT, LoRA, TRL, and SFTTrainer**.
- Produced a model that generates contextually relevant, natural Gen-Z-style responses.

Research Implementations: Research-paper and ML/DL implementations

[1](#) [2](#)

- Implemented research papers and ML/DL concepts from scratch in Python and PyTorch.
- Includes **MHA, GPT, RoPE, BPE, RNN, and AlexNet**.

Achievements

NIT Goa 0xGenIgnite Winner: Won a national-level Web3 hackathon at NIT Goa.

Logithon '25 IIT Bombay Finalist: Placed in the top five at the national logistics hackathon.

World Computer Hacker League '25: National qualifier.