Abhishu Oza

abhishu.oza@gmail.com | (732) 880-1364 | abhishu-oza | abhishuoza

Summary

MS Computer Science student at Rutgers University specializing in machine learning and AI. Experience with LLM fine-tuning and deep reinforcement learning (LoRA, RLHF/GRPO). Published in IEEE Access on deep learning for time series forecasting. Interested in alignment research and scalable oversight. Seeking ML engineering roles and research opportunities.

Education

Rutgers University - New Brunswick Jan 2025 to present
Master of Science, Computer Science GPA: 4.0/4.0
Coursework: Machine Learning I & II, Artificial Intelligence, Massive Data Mining, Computational Geometry, Topics in Computers in Biomedicine, MS Thesis

Ahmedabad University Oct 2020 to May 2024
BTech in Computer Science and Engineering GPA: 3.2/4.0
Coursework: Machine Learning, Computer Vision, Probabilistic Graphical Models, Design and Analysis of Algorithms, Data Structures, Applied Linear Algebra, Probability and Stochastic Processes, Operating Systems

Experience

Machine Learning Intern, AIML.com Dec 2025 to Jan 2026

- Developed coding tutorials implementing modern architectures from scratch in PyTorch (GPT, Vision Transformer, Contrastive Learning, Distributed Optimization) including mathematical formulations, attention visualizations, and benchmarks on CIFAR-10 and TinyShakespeare datasets.
- Authored technical articles synthesizing research papers into accessible tutorials.

Student Programmer and Part Time Lecturer, Rutgers University May 2025 to Dec 2025

- Engineered a coding assignment system to convert DAGs and YAML specifications into Jupyter notebooks for automated assignment grading. Contributed to backend infrastructure using Flask, AWS, and PostgreSQL to serve 100s of students and teachers at the Department of Computer Science.
- Worked as a Part Time Lecturer and TA for the Introduction to Data Science course under Professor James Abello.

Machine Learning Intern, Reliance Jio Platforms May 2023 to Jul 2023

- Engineered a computer vision classification model utilizing VGG19 CNN for the project "Agriculture - Crop Selection and Planning". Automated satellite imagery geofencing and plant anomaly detection.
- Developed a geospatial web interface using the Google Maps API, enabling visual analytics and interaction with large-scale geospatial datasets.

Undergraduate Researcher, Ahmedabad University Sep 2022 to May 2024

- Researched long-term time series forecasting techniques (DLinear, TimesNet, iTransformer) for undergraduate thesis.
- Proposed a novel fusion 1DCNN-LSTM based architecture for 24 hour energy load forecasting of 100 households, work published in IEEE Access.

Publications

[1] **A. Oza**, D. K. Patel and B. J. Ranger, "Fusion ConvLSTM-Net: Using Spatiotemporal Features to Increase Residential Load Forecast Horizon," in *IEEE Access*, vol. 13, pp. 12190-12202, 2025, doi: 10.1109/ACCESS.2025.3528072.

Projects

Documentation RAG based coding tool

- Built an RAG-based coding tool using LangChain, ChromaDB, and sentence-transformers that scrapes live documentation, indexes it into a vector store, and generates contextual code via OpenAI, Anthropic, or local HuggingFace models.
- Designed a thin-client/fat-server architecture with FastAPI and Typer CLI with remote mode via httpx, Docker deployment, optional dependency groups via pyproject.toml, and a comprehensive test suite using pytest.

GRPO Math Reasoning with Reinforcement Learning

- Applied Group Relative Policy Optimization (GRPO) to Qwen2.5-1.5B-Instruct on GSM8K via TRL's GRPOTrainer. Achieved +16.7 pt correctness and +47.0 pt formatting accuracy improvement on the full 1,319-example test set.
- Separated generation and training across 8 GPUs using vLLM as a dedicated inference server with DDP training, improving training time by 5x (from 90 hours to 18 hours). Added Liger kernel for memory optimization.

Legal Domain LoRA Finetuning using PEFT

- Finetuned Qwen3-8B on 100,000 court opinions from CourtListener using Low-Rank Adaptation (LoRA) via the PEFT library. Conducted rank ablation studies for $r \in \{4, 8, 16\}$.
- Evaluated text completion perplexity on held out court opinions and downstream reasoning performance on LegalBench.

Connect 4 AlphaZero

- Implemented AlphaZero from scratch in PyTorch to master Connect 4 through self-play reinforcement learning. Designed a ResNet with dual policy and value heads, guided by Monte Carlo Tree Search (MCTS) with 400 simulations per move.
- Trained using 500 self-play games per iteration with batched GPU inference across 32 parallel games, a 200K replay buffer with augmentation, and a playable GUI via pygame.

Uncertainty Direction in LLM activation space

- Investigated if uncertainty can be represented as a linear direction in the activation space of gpt2-small via mechanistic interpretability techniques. Performed activation caching and causal intervention using the TransformerLens library.

Skills

Machine Learning: PyTorch, TensorFlow, NumPy/Pandas, Scikit-learn, HuggingFace, vLLM, LangChain, PEFT, TRL, OpenCV

Programming: Python, C/C++, Java, SQL, FastAPI, Flask, PostgreSQL, Kubernetes, Docker

Tools: Git, Linux Shell, Slurm, MATLAB, Google Maps API, AutoCAD, Fusion360

Activities

Co-Founder and Chairman – Ahmedabad University ACM Student Chapter — Aug 2022 to Aug 2023

- Co - Founded and became chairman of the ACM student chapter at Ahmedabad University.
- Hosted events and conducted competitions for foundational concepts in Machine Learning and Blockchain.

Silver Certificate – Technothlon — Jul 2019

- Secured an All-India Rank of 62 in Technothlon, a school championship organized by IIT Guwahati.

Volunteer – Prabhat Foundation — Sep 2021 to Dec 2021

- Completed 30 hours of volunteer work comprised of creating disability awareness.