

Ayush Morbar

+91 (6266) 506 736 | amorbar@andrew.cmu.edu | Gwalior, India | linkedin.com/in/ayushmorbar | github.com/ayushmorbar

SUMMARY

Deep learning / ML engineer (LLMs, NLP) | New Grad (B.Tech 2026). Experienced in training and fine-tuning models (BERT, YOLO) and building agentic systems and optimization-backed decision engines. Teaching Assistant for CMU 11-785 Deep Learning.

EDUCATION

Carnegie Mellon University Pittsburgh, PA (Remote)
Non-Degree Coursework: 11-785 Deep Learning (CMU) Jan 2026 – May 2026

Rustamji Institute of Technology Gwalior, India
Bachelor of Technology in Information Technology Oct 2022 – Jul 2026

EXPERIENCE

Teaching Assistant (11-785: Introduction to Deep Learning) May 2026 – Present
Carnegie Mellon University (CMU) Remote (Pittsburgh, PA)

- Supporting the upcoming Fall 2026 offering led by Prof. Bhiksha Raj by preparing recitation/assignment materials and helping validate course infrastructure.
- Running model training and evaluation sanity checks for assignment reference implementations, including porting selected notebooks from PyTorch to JAX across multiple deep learning architectures.

Software Engineering Fellow Jul 2024 – Sep 2024
Headstarter AI Remote (Long Island City, NY)

- Built LLM-driven prototypes using Next.js and FastAPI, including RAG-style demos with OpenAI APIs and Pinecone.

Head of AI/ML (Google Developer Groups on Campus) Jul 2024 – May 2026
Google Developer Groups (GDG) on Campus RJIT Gwalior, India (Hybrid)

- Led campus AI/ML workshops and study jams (200+ students reached), including GenAI sessions using Google Cloud and Gemini APIs.

PROJECTS

OptiBuild (GitHub | Demo) | *Google ADK, MCP, OR-Tools (CP-SAT), Pydantic* Jun 2026 – Jul 2026

- Designed an agent-to-solver architecture that turns natural language goals into a formal optimization model (decision variables, objectives, constraints), solved selection problems with OR-Tools CP-SAT, and added multi-objective ranking via TOPSIS.

Resurrect (GitHub) | *FastAPI, Docker, Multi-Agent Systems, LiteLLM, GitHub API* Sep 2025 – Present

- Built an autonomous code-repair system that runs tests in a Docker sandbox, applies multi-agent fixes, and can commit changes/open PRs; validated on a small repo.

Blueprint Door & Window Detection (GitHub) | *YOLOv8, Computer Vision* 2025

- Trained a YOLOv8 detector for doors/windows in architectural blueprints; built the end-to-end pipeline (preprocessing, augmentation, training/validation, inference) and evaluated with mAP + error analysis to guide dataset cleanup.

SKILLS

Languages: Python, SQL, TypeScript/JavaScript

ML/LLMs: PyTorch, JAX, Hugging Face Transformers, Fine-tuning, Evaluation

Agents & Tooling: Google Agent Development Kit (ADK), MCP, LiteLLM, RAG, Embeddings, Vector DBs (Pinecone)

Backend/Web: FastAPI, Next.js, React

DevOps & Systems: Docker, Linux, Git, CI/CD

Cloud: GCP (course/research testing), AWS (basic)