Sanyam Wadhwa

Ludhiana, Punjab, India | +91-87280-67646 | sanyamwadhwa.in@gmail.com
sanyamwadhwa.tech | GitHub | LinkedIn | LeetCode

Education

Thapar Institute of Engineering and Technology — Patiala, Punjab
B.E. Computer Engineering, CGPA 8.31 — *Aug. 2023 – Expected 2027*

Experience

Software Engineering Intern — Feb. 2025 – Nov. 2025
TIET Examination Platform, exams.thapar.edu — *Patiala, Part Time*

- Built the seating and invigilation workflow of TIET's 5-developer Flask exam platform that seats all **12,000+ students** for every paper in a cycle, cutting each plan from **3–4 days** of manual work to **about 40 minutes**
- Rewrote **2 of the platform's 4 seating algorithms**, AB and ABCD, so when a course runs out mid-room, waiting courses take its seats in order of students left, and rooms down to 2 courses switch to a diagonal 2x2 layout
- Shipped the exam office's change tools to block unavailable rooms, swap rooms and invigilators and print per-teacher duty PDFs, and moved their inputs from browser sessions into MySQL so **all admins share one saved state**

AI Research Intern — Jun. 2025 – Jul. 2025
Centre of Excellence - UQ, TIET — *Patiala, On Site*

- Designed a mental-health chatbot's safety path as plain rules instead of a model, with a keyword crisis gate that returns helpline numbers **before any LLM call** and override rules that route every other message
- Built the message pipeline that runs fine-tuned BERT emotion and sentiment classifiers and several mental-health tests in parallel, then has the LLaMA responder write **a report** from their combined results
- Served the responder as a 4-bit LLaMA via llama.cpp, so the prototype ran locally **without a paid API**

Projects

PaperMind | *Python, FastAPI, LangGraph, React, PostgreSQL, pgvector, Docker, GitHub Actions* — GitHub

- Built a research-paper reader that turns an arXiv ID into the paper's key findings with their reported numbers, plus its methods, datasets and figures, reading the **full text** instead of only the abstract
- Achieved figure-extraction F1 of **0.972** on 16 arXiv papers, recovering **all 122** vector and raster figures by growing each crop upward from its caption with PyMuPDF
- Kept paper runs alive through API outages by failing over from Gemini to Groq to Ollama; a paper too large for the fallback's free-tier limit is reprocessed section by section **instead of failing**

Onramp | *Python, Agent Skills, GitHub Actions* — Live | GitHub

- Built an agent skill for Claude Code, Codex, Copilot, Cursor and **60+** other coding agents that turns an unfamiliar repository into a guided tour of its code, sized to the reader's time
- Tied every claim in a guide to the exact source lines it comes from, with a checker that caught **all 525** planted citation errors and raised no false alarms on 200 real ones
- Measured the guides in a blind fact check where **97%** of sampled claims held up against the code

Technical Skills

AI	Hugging Face, LangGraph, llama.cpp, pgvector, PyTorch, RAG
Languages	C, C++, JavaScript, Python, SQL, TypeScript
Backend & Web	FastAPI, Flask, Node.js, SQLAlchemy, REST APIs, WebSockets, React, Next.js, Tailwind CSS
Data & Infra	Docker, GitHub Actions, Linux, MySQL, PostgreSQL

Achievements and Leadership

1st Place, Clinical/Engineering Track, PGIMER & AOPN Advances in Pediatric Epilepsy, for NeoSpasm — Apr. 2026
Winner, Samsung PRISM Web-Agent Hackathon, for KONTA, on-device tab search built by a team of 4 — Mar. 2026
Top 50 of 1,000+ on Eris, shipd.ai's Kaggle-style ML platform, with 10 gold, 8 silver and 9 bronze medals — Ongoing
Top 200 in 2025 and **Top 300** in 2026, Amazon ML Challenge — 2025, 2026
Technical Coordinator, Microsoft Learn Student Chapter, TIET — Aug. 2023 – Apr. 2025