Anubhab Pradhan

Bengaluru, Karnataka (open to relocate) | +91 63622 02547 | pradhananubhab25@gmail.com
github.com/ANUBprad | linkedin.com/in/anubhabpradhan | leetcode.com/u/Anubhab25 | anubhab99.vercel.app

Experience

Applied AI Engineer Intern Jan. 2026 – May 2026
Lamatic.ai *Remote*

- Tested live AI agents and found 3 failure modes: fabricated plans, wrong tool choices, and reasoning loops. Built tracing tools to step through each run and find where it failed
- A/B tested prompt and routing changes on task completion and retention, shipped the versions that won, and built multi-step workflows that call REST and GraphQL APIs

AI Engineer Intern Aug. 2026 – Present
Hasprana Health Care Solutions *Bengaluru*

- Audited a 62K-line detector and found its confidence score was computed but never used. Fixed two calibration bugs and added 251 automated tests
- Retrained a U-Net on clinical eye data (mean IoU 0.888 to 0.970). Quantized it with ONNX and INT8, cutting model size 75% with no accuracy loss

Projects

RedOps | *Python, FastAPI, React, Next.js, TypeScript, PostgreSQL, Temporal, Docker* GitHub

- Built an LLM and agent evaluation platform that scores safety, groundedness, and cost. A FastAPI backend powers a React dashboard, with one provider layer for OpenAI, Anthropic, and Groq
- Ran evaluations as Temporal workflows that retry transient failures and stop cancelled runs. Verified 3,243 passing backend tests (3 skipped), 31 frontend tests, and 24/24 algorithmic metric checks

Kairos | *Next.js, React, TypeScript, Python, PostgreSQL, pgvector, Tailwind CSS* GitHub

- Built a RAG research workspace that ingests PDFs, URLs, and YouTube videos, then answers with inline citations and a trace of retrieved chunks and scores
- Evaluated retrieval with 12+ metrics and confidence intervals, and covered the work with 1,010 automated portal tests across 206 suites

Open Source Contributions

LangChain (Python): PR #31802 (merged) fixed a missing `else` branch in project evaluation logic

Copperhead (TypeScript): PR #368 (open) addressed a ReDoS risk in file search. Regex evaluation now runs in a worker with a 2-second budget and a 1,000-character pattern limit, with timeout, cancellation, and regression tests

Cognee (Python): PR #5583 (open) closed a tool-handler security gap. Dynamic imports are limited to approved namespaces, unauthorized handlers are filtered during discovery, and custom schemas cannot impersonate reserved types

VS Code (TypeScript): PR #339463 (open) fixed a promise that could stay pending forever when a terminal's process had already exited. Live terminals are unchanged; 219 terminal tests passed

Ollama (Go): PR #18677 and PR #18682 (both open) proposed fixes for tool-call markup lost in streamed output. One keeps partial opening tags in the DeepSeek3 parser, and the other stops an OLMo3 tool call in the final chunk from bypassing the parser

Technical Skills

Languages: Java, JavaScript, TypeScript, Python, SQL, Go
Frontend: React, Next.js, Tailwind CSS
Backend & Data: FastAPI, REST APIs, GraphQL, PostgreSQL, pgvector, Temporal, Docker
Testing & Quality: Unit, Integration, and API Testing, JUnit, pytest, Vitest, Git, GitHub Actions, Linux
AI & LLM: LLM Output Evaluation, RAG, Tool Calling, Agentic Workflows, A/B Testing, OpenAI API, Anthropic API, Groq, ONNX, INT8 Quantization

Achievements

Agentic AI Day 2025: **Finalist**, built a smart city planner agent for daily tasks across Bengaluru
Amazon ML Challenge 2026: Team lead; ranked in the top 20% with an $F_{0.5}$ score of 0.9561
Smart India Hackathon 2024: Finalist at national finals (Chennai), built JAL SENTINEL, a flood alert platform

Education

CMR Institute of Technology Bengaluru, Karnataka
B.E. in AI and Data Science, CGPA: 7.9 *Sep 2023 – June 2027*