

Sundar Suhrith Adavi

+91 9381519387 | suhrithwork@gmail.com | github.com/suhrith13 | linkedin.com/in/sundaradavi

TECHNICAL SKILLS

Languages: Python, TypeScript/JavaScript, SQL, Bash

Backend & APIs: FastAPI, RESTful API design, secure API design, PostgreSQL, MongoDB, MySQL

Frontend: React, TypeScript

Distributed Systems & Cloud: Docker, Kubernetes (pods, deployments, services), AWS (EC2, S3, Lambda), Cloudflare Workers/Pages/D1/KV, Terraform (IaC fundamentals), Linux systems administration

CI/CD & Engineering Practice: GitHub Actions, GitLab CI concepts, pipeline automation, TDD, OOP, dependency scanning, RBAC fundamentals, data governance

AI/LLM & Agent Tooling: Claude API, Gemini API, RAG pipelines, agent-based automation, structured output integration

PROFESSIONAL EXPERIENCE

Software Engineer Intern | DRDO, India

Jan 2023 – May 2023

- Developed Python-based predictive models achieving 90% target position accuracy through ML techniques and automated testing.
- Optimized analytical workflows by 40% using MATLAB and Python for distributed applications with design patterns for modularity/scalability.
- Enhanced model reliability via refactoring and automated tests for multi-threaded data processing on Linux, establishing governance/data quality practices.
- Facilitated cross-platform deployment via Docker for Python/C++ apps across distributed teams.
- Collaborated cross-functionally to translate analytical requirements into scalable data pipelines using TDD and OOP principles.

INDEPENDENT PRODUCT DEVELOPMENT

SponsorPath — Full-Stack AI Job Matching Engine

sponsorpath.pages.dev

- Delivered a full-stack application end-to-end — Cloudflare Workers backend paired with a React/TypeScript frontend — fetching and quality-scoring data at runtime, with results cached in the user's browser to cut redundant calls.
- Built an automated data pipeline that ingests, normalizes, and quality-scores 10K+ employer records from live Greenhouse and Lever job board APIs, applying validation rules that improved match accuracy 35% over baseline keyword matching.
- Designed a two-part AI service layer (Claude API) for resume tailoring and sponsorship-likelihood scoring, with structured output feeding directly into the downstream data model.
- Implemented client-side caching to cut redundant LLM API calls, reducing cost and latency for repeat user queries.

InterviewIQ — AI Mock Interview Platform

interviewiq.pages.dev

- Designed, built, and deployed a full-stack LLM-powered platform end-to-end — from system design through implementation and production support — now serving 150+ users across 400+ sessions.
- Architected the backend on Cloudflare Workers and KV for low-latency, serverless request handling, with per-IP rate limiting to keep the platform stable under production traffic.
- Implemented a RAG-based question-generation and evaluation pipeline (Gemini API) with role-aware question banks and personalized scoring from structured model output.

PR Merge Predictor — Self-Verifying ML Backend Service

github.com/suhrith13/PR-Merge-Predictor

- Built a backend scoring service that predicts GitHub pull request merge likelihood in real time, training a logistic regression model on 18 open-source repos to reach 0.923 AUC and 89.4% accuracy vs. a 52.8% majority-class baseline.
- Deployed inference as a dependency-free scorer embedded in a Cloudflare Worker, backed by a D1 data store and a Cron-triggered job that re-checks unresolved predictions every 6 hours against real outcomes.
- Shipped a public live-accuracy dashboard computed from resolved, real-world outcomes, paired with a static Cloudflare Pages frontend for one-click scoring.

EDUCATION

Master of Science, Computer and Information Science | Purdue University, Indiana

May 2025

Bachelor of Technology, Information Technology | Manipal University, Jaipur, India

May 2023