

Vidith Salla

(703) 861-1056 | vidithsalla@gmail.com | [LinkedIn](#) | [GitHub](#)

EDUCATION

Bachelor of Science in Computer Science

The Pennsylvania State University, University Park

GPA: 3.52/4.00

May 2025

EXPERIENCE

Software Engineer, CYMA Systems, Remote, US

Jun 2026 – Present

AI-native engineer to project matching platform

- Co-designed an AI-native engineer-to-project matching system with a senior data architect, structuring historical staffing decisions, company policies, and engineer experience into a retrieval and ranking architecture.
- Own the Graph RAG retrieval workflow on a 5-engineer team, combining deterministic eligibility filters, graph-based skill and work-history relationships, and semantic retrieval of historical context.
- Designed the retrieval layer to separate hard staffing constraints from qualitative evidence, keeping deterministic eligibility rules outside downstream AI-based match reasoning and explanations.

Full-Stack Software Engineer, WinWinLabs, Remote, US

Sep 2025 – Jun 2026

SensorWatch: IoT Monitoring SaaS

- Built and deployed an IoT telemetry and alerting platform with a 3 person engineering team, working directly with the founders, for a pilot across 12 franchise locations using ESP32 → MQTT → PostgreSQL → Next.js, with RBAC, device onboarding, deduplication, and offline detection.
- Reduced API latency from 8–15s to 1–3s by implementing adaptive time-series bucketing in PostgreSQL, compressing 24h telemetry from 86K points (2MB) to ~200 points (10KB).
- Built an LLM-based compliance copilot that generates audit-ready incident reports from alert data, using structured outputs, prompt guardrails, and redaction workflows for store and regional stakeholders.

Software Engineer, Tride (Part-time)

Nov 2025 – Present

Construction ERP SaaS with AI document ingestion for India and the US

- Independently built and rapidly iterated a multi-tenant ERP from initial concept to MVP and live use by 2 companies, using Next.js, TypeScript, and Supabase/PostgreSQL with org-scoped tenancy and server-side RBAC across 5 roles.
- Built a Gemini multimodal pipeline to extract invoices, delivery challans, and attendance registers, routing results by confidence and reusing approved corrections in vendor-specific prompts.
- Designed transaction-ledger inventory workflows for GRNs, issues, transfers, and adjustments using atomic PostgreSQL transactions and reconciliation tooling.
- Diagnosed a production connection deadlock between the PostgreSQL driver and Supabase transaction pooler, shipped a session-mode fix, and added regression coverage.

Undergraduate Researcher, Penn State ReSense Lab, University Park, PA

Jan 2024 – May 2025

- Built and deployed FastAPI backend services integrated with the OpenAI API, supporting 3 production courses and scaling to 11,000+ users.
- Designed a multi-session, multi-profile architecture to support concurrent users and persistent course-specific interactions.
- Redesigned relational data models and REST APIs for secure multi-session access and GDPR/FERPA-aligned data handling.
- Collaborated in a 6-person team to ship production features and improve the UX based on student and instructor feedback.

Data Engineer Intern, American Chemical Society, Washington, DC

May 2024 – Aug 2024

- Automated data cleanup across 15+ GCP environments using Python and Cloud Scheduler, improving system reliability and reducing storage usage by 22%.
- Improved data pipeline performance by 15% by optimizing BigQuery and Cloud Data Fusion workflows, reducing data load and query latency.
- Automated test result logging into Jira using Python and Selenium, eliminating manual QA steps and accelerating bug triage.
- Supported GA4 migration to BigQuery by adapting event schemas and writing py scripts for batch processing & validation.

Software Engineering Intern, Tsaro Labs, Farmington, MI

May 2023 – Jul 2023

- Built React + Chart.js dashboards for real-time KPI tracking, integrating with Spring Boot REST APIs to support a production tool used daily by 25+ sales and finance users.
- Automated commission calculations by building SQL pipelines and REST endpoints, reducing payout processing time from 2–3 days to ~1 hour.
- Contributed to KPI data pipelines using PostgreSQL and Airflow, supporting ETL workflows and improving reliability of reporting systems used by internal teams.

PROJECTS (github.com/vidithsalla)

Trace: Grounded Legal AI Copilot (Next.js, TypeScript, FastAPI, PostgreSQL/pgvector, Redis, Docker) | [GitHub](#)

- Built a production style AI system for legal document analysis with multi-format ingestion (PDF/DOCX/XLSX/CSV), async processing, and matter-scoped RBAC with audit logging.
- Implemented hybrid retrieval combining Postgres full-text search and pgvector embeddings with reciprocal-rank fusion, enabling accurate search across large document sets.
- Developed grounded AI workflows (Q&A, chronologies, entity summaries, contradiction detection) with deterministic citation validation and RAG pipeline (parse → chunk → embed → retrieve → generate → validate) with eval harness and 50+ API tests.

InboxPort: MCP Access Manager for Gmail & Outlook (Next.js, TypeScript, OAuth 2.1) | [GitHub](#) | [Live Demo](#)

- Built a remote, multi-tenant MCP gateway connecting multiple Gmail and Outlook accounts (Gmail API and Microsoft Graph).
- Enforced revocable, server-side permissions per AI client and account, fetching mail on demand without retaining any mail data.
- Implemented human-approved sending with pre-dispatch draft re-hashing to block modified messages; validated authorization and send safety across 227 unit and security tests.

TECHNICAL SKILLS

Languages: Python, TypeScript, Java, C, SQL; **Frameworks/Tools:** Next.js, React, FastAPI, Docker, Git, GitHub Actions, Selenium, MQTT, Supabase; **Databases/Cloud:** PostgreSQL, BigQuery, Cloud Data Fusion, Snowflake, AWS, GCP