Prajwal Hosahalli Dayananda

California, US | 260-999-0600 | hdprajwalgowda@gmail.com | linkedin/in/hdprajwal | github/hdprajwal | hdprajwal.dev

Summary

Software Engineer with 3+ years building scalable full-stack web and mobile applications. Shipped distributed systems handling 100,000+ cloud resources with 99% uptime and $100K+ monthly cost savings. Skilled in TypeScript, React, Next.js, Node.js, Python, PostgreSQL, and Kubernetes across AWS, GCP, and Azure. MS in Computer Science, Purdue (GPA 3.88). Open to relocation.

Education

Purdue University **Jan 2024 - Dec 2025**
Master of Science, Computer Science *Fort Wayne, IN*

Visvesvaraya Technological University **Aug 2017 - Jul 2021**
Bachelor of Engineering, Computer Science & Engineering *Bengaluru, India*

Experience

Translation Commons
Software Engineer *CA* · **Feb 2026 - Present**

- Designed and shipped 4 data-navigation features across 5 merged PRs, establishing a reusable filter-context pattern on an open-source platform backed by UNESCO, Unicode CLDR, and Ethnologue.
- Built a Playwright visual regression suite (4 pages, OS-agnostic baselines) with a label-driven CI workflow that regenerates snapshots from PRs for 16+ contributors, enabling faster detection of UI regressions and reducing manual review effort
- Introduced Amplitude product analytics (page views + export events) and Cloudflare Pages per-PR preview deployments with automatic teardown, giving maintainers the platform's first usage telemetry and reviewers zero-setup environments.

Opslyft
Cloud Engineer *Bangalore, India* · **Mar 2022 - Nov 2023**

- Built a cloud cost optimization platform (Next.js + Python) that analyzed spending patterns across AWS, GCP, and Azure - saving one enterprise client $100K+/month by automatically identifying and flagging underused resources.
- Developed the company's marketing website using Next.js and Tailwind CSS with responsive design, animations, and dynamic blog content integration via WordPress API, serving as the primary customer acquisition channel.
- Developed React + GraphQL dashboards enabling real-time visibility and management of 100,000+ cloud resources across multi-cloud environments for enterprise customers.

Software Development Engineer *Bangalore, India* · **Apr 2021 - Mar 2022**

- Redesigned and rebuilt the Instance Scheduler service (Node.js GraphQL API + React frontend), fixing critical reliability bugs and adding input validation, retry logic, and error recovery - enabling fault-tolerant scheduling of thousands of cloud resources across customer accounts.
- Achieved 99% service uptime by designing fault-tolerant error handling that isolated individual failures, ensuring one problem never disrupted other customers' automated workflows.
- Built a full-stack cost tracking system with a Python data pipeline, PostgreSQL database, Node.js GraphQL API, and dashboard UI - giving customers complete per-resource accountability over their data transfer spending.

Gradspace
Software Engineer Intern *Bangalore, India* · **Sep 2020 - Feb 2021**

- Reduced API response times by 75% (2s to 500ms) by restructuring PostgreSQL schema and optimizing queries, directly improving perceived app performance for 500+ beta users.
- Containerized the backend with Docker and built a GitHub Actions CI/CD pipeline, eliminating manual deployments and accelerating release cycles.

Projects

Agentic AI Coding Assistant - Quackcode | Github **Jan 2026 - Present**

- Architected QuackCode, a Linux-native Electron + TypeScript coding agent unifying Claude, Gemini, and Codex behind one interface with thread-level provider locking and a local SQLite event log chosen for concurrent scheduler writes; daily driver for production coding, including most of QuackCode's own development.
- Designed provider-agnostic event replay, git worktree isolation for bad agent edits, and an idempotent scheduler; profiled and fixed a renderer hotspot by batching streaming chunks per frame, eliminating UI lag.

Multimodal Malware Detection - M.S. Thesis Research | Paper **Jan 2025 - Dec 2025**

- Investigated whether fusing visual and tabular representations of Android malware could outperform single-modality approaches; achieving ROC AUC of 0.9633 with LightGBM on EMBER 2024, surpassing a ViT baseline trained on 1.2M MalNet samples.
- Designed a multimodal fusion framework combining tabular static features with visual byte-plot representations, evaluating gradient-boosted tree and Vision Transformer baselines independently before applying LightGBM + DINOv2 ensemble fusion; improving TPR by 6 percentage points at FPR=0.001.

Technical Skills

- **Languages**: Python, TypeScript, Go, Java
- **Frontend**: React, Next.js, React Native, Tailwind CSS, Vue.js
- **Backend**: FastAPI, Express, Node.js, Gin, REST API, GraphQL
- **Databases & Caching**: PostgreSQL, MySQL, MongoDB, Redis, DynamoDB, BigQuery
- **Cloud & DevOps**: GCP, AWS, Azure, Docker, Kubernetes, Terraform, Cloudflare, GitHub Actions
- **ML & AI**: Multi-provider LLM APIs, Langchain, AI Agents, Weights & Biases, PyTorch, TensorFlow, scikit-learn

Publications

- Hosahalli Dayananda, P; Chen, Z.Re-Evaluating Android Malware Detection: Tabular Features, Vision Models, and Ensembles.*Compares tabular static features, vision-based models, and ensemble strategies for Android malware classification.*