

Divyam Ahuja

Systems & Software Engineer

☎ +91 7015600547 ✉ ahujadivyam@gmail.com 🌐 divyam.dev 📄 in/ahujadivyam

🔗 git.divyam.dev 📁 gh/ahujadivyam

Experience

Microsoft

Software Engineer

Noida

Mar 2025 - Present

Tech Stack: Python, TypeScript, React, Azure OpenAI, Claude Code, Copilot, Playwright

- Owned the end-to-end Responsible AI (RAI) evaluation pipeline, executing critical harms assessments to gate model releases. Engineered custom retry mechanisms and architectural optimizations, reducing average pipeline execution time by **50%**.
- Engineered an automated synthetic data generation framework using LLMs, generating **70,000+** diverse queries and multi-step interaction trajectories across **20,000+** real-world Excel workbooks to power large-scale model evaluation pipelines.
- Served as on-call Shield Engineer, owning incident triage and partner bug resolution across AI evaluation pipelines—maintaining **>95%** uptime by engineering comprehensive telemetry dashboards with an adjacent team.
- Accelerated engineering velocity by building LLM-based agents for automated bug triage, PR summarization, and reviewer assignment.
- Shipped high-impact performance optimizations and accessibility fixes directly to the core Microsoft Excel product.

Advanced Micro Devices (AMD)

System Software Designer

Bangalore

Aug 2022 - Feb 2025

Tech Stack: C/C++, Java (OpenJDK), Python, Django, React, SQLite

- Engineered a visualization dashboard and RESTful API for OpenJDK telemetry, increasing internal data accessibility by **40%**.
- Automated nightly Java benchmarking and tuning pipelines using machine learning models, achieving a **30%** reduction in manual effort and earning the **Spotlight Award**.
- Optimized JVM flags, GC policies, and runtime parameters for OpenJDK on High Performance Computing (HPC) servers, achieving a **2%** throughput improvement on SPECjbb benchmarks.

Chromium

Google Summer of Code Contributor (Open Source Developer)

Remote

May 2022 - Aug 2022

Tech Stack: C++, JavaScript/TypeScript, WebAssembly, WebAudio, WebWorkers

- Engineered a high-performance, lock-free C++ ring buffer library for real-time audio processing, establishing optimal design patterns to integrate the WebAudio API with WebAssembly and WebWorkers.

Skills

Programming Languages: C/C++, Python, Go (Golang), Java, JavaScript/TypeScript, Bash

AI & Data: PyTorch, Azure OpenAI, Antigravity, Pandas, NumPy, Claude Code

Backend & Web: Node.js, Django/Flask, React/Next.js, gRPC, Playwright

Infrastructure & Systems: Kubernetes, Docker, Linux/POSIX, WebAssembly, WireGuard, CI/CD

Databases: PostgreSQL, Redis, MySQL, SQLite, MongoDB

Projects

fastbin: a code sharing webapp | HTMX, Go (Golang), gRPC, PostgreSQL, Docker, Kubernetes

- Pastebin-style app with a Go backend, gRPC service layer, and Kubernetes deployment.
- git.divyam.dev/divyam/fastbin, fastbin.lab.divyam.dev

Homelab | Raspberry Pi 5, WireGuard, Kubernetes, Rancher, MetalLB

- Self-hosted Kubernetes cluster with WireGuard tunnel to a public VPS for secure external access.

Education

Deenbandhu Chhotu Ram University of Science and Technology

Bachelor of Technology (B.Tech.) Computer Science

Haryana, India

Aug 2018 - Aug 2022

Certifications

- MITx 6.86x — Machine Learning with Python: Linear Models to Deep Learning
- CS50AI: CS50's Introduction to Artificial Intelligence with Python