Dhruv Bansal

San Francisco, CA | (602)-790-9929 | dbansa11@asu.edu | linkedin.com/in/dhruvb26 | github.com/dhruvb26

Summary

Final-year computer science undergraduate with experience in backend software engineering, distributed systems, and shipping production software.

Education

Bachelor of Science in Computer Science *May 2026*
Arizona State University **GPA 4.0**
Relevant Coursework: Software Engineering, Data Structures, Social Media Mining, Operating Systems, Deep Learning
Achievements: NAMU Scholar ($14,500 p.a.), FACE Program @ ASU, Dean's List, 3x Hackathon Winner

Professional Experience

Software Engineer Intern *June 2025 – Dec 2025*
Chunkr San Francisco, CA

- Developed high-throughput data discovery pipelines over millions of images and improved by handling complex layouts, malformed inputs, and extraction edge cases at scale.
- Engineered a concurrent, high-performance Rust extraction service, reducing end-to-end processing latency by >90% (from days to hours) and unblocking large-scale document pipelines.
- Implemented core Rust extraction and chunking logic, optimizing task scheduling, bounding box integration and segment-level processing in a high-volume production environment.

Software Engineer Intern *Jan 2025 – Apr 2025*
Oats Overnight Co. Phoenix, AZ

- Built internal customer-insights and visualization tooling using Mapbox, GeoJSON, and Deepmerge to aggregate and render geospatial and behavioral data, enabling data-driven analysis of customer patterns at scale.
- Refactored reusable user-interface components using Storybook, reducing frontend code duplication by ~30% and improving release consistency across 5+ production web applications.
- Implemented internal systems for real-time inventory tracking and automated replacement-order workflows, reducing manual intervention by 43% and improving fulfillment reliability across high-volume ecommerce operations.

Software Engineer Intern *June 2024 – Aug 2024*
PrivateBlok Remote

- Engineered data ingestion pipelines for the platform, redesigning legacy Python scripts to support ingestion of 100,000+ private market signals per day, improving upstream model training.
- Automated extraction of 1,500+ web pages daily, reducing manual collection time by ~70% and improving research dataset freshness for Copilot features.

Projects

Bluecast - The LinkedIn Growth Tool | *https://www.bluecast.ai/*

- Built and operated a scalable LinkedIn growth platform with automated content optimization and engagement scoring pipelines, driving $40k+ in monetized platform usage.

Paperal - Research Writing & Knowledge Management Platform | *https://paperal.com/*

- Built a research writing platform for long-form documents with context-aware retrieval, citation grounding, and real-time autocomplete.

Multi-Agent Code Generation System | *https://github.com/dhruvb26/CSE475-Project*

- Built a multi-agent code generation and coordination framework, incorporating GRPO-style reinforcement learning to study cooperative agent behaviors and improve solution quality on programming tasks.

Causal Feature Selection | *https://github.com/dhruvb26/CSE472-blanket-challenge*

- Designed a causal feature selection pipeline using Markov blanket discovery and TabPFN evaluation to test robustness and generalization under distribution shifts.

Technical Skills

Languages: Rust, Python, TypeScript, Go, Swift, Java, C++, C, Assembly
Systems & Foundations: Distributed Systems, Concurrent Programming, High-Throughput Data Pipelines, Fault Tolerance, Data Ingestion, Web Automation, Reinforcement Learning, Multi-Agent Systems
Frameworks & Tools: Linux/Unix, Docker, AWS, PostgreSQL, PyTorch, GraphQL, LangChain, CrewAI, React, Next.js, Vue.js, SwiftUI