ISHAANI JAIN

U.S. Citizen | ishaani.jain1@gmail.com | Seattle, WA & Irvine, CA | Linkedin | Github | Portfolio

EDUCATION

University of California, Irvine - B.S. Computer Science & Physics | Class of 2028
Relevant Coursework: Computer Architecture, Operating Systems, Data Structures & Algorithms, Discrete Mathematics, C++ Programming, Java Programming, Informatics, Multi-variable Calculus, Mechanics, Linear Algebra
Research & Leadership: SPAR AI Safety Fellow (15% acceptance) | UROP Research Fellow | Naval Sea Cadet Petty Officer 2nd Class (1 of 10 statewide)

TECHNICAL SKILLS

Systems Programming: C++, CUDA, Python, Java | Low-level optimization, memory management, performance profiling
ML Infrastructure: PyTorch (C++ extensions), TensorFlow, CUDA kernel development, NCCL, MPI, TensorRT
Hardware & Architecture: Cache simulation, performance analysis (nsight-compute, perf), Roofline modeling, memory hierarchy optimization
Cloud & Tools: AWS (Lambda, CDK, DynamoDB, S3, API Gateway, Cognito), Docker, Linux, Git, Jest

EXPERIENCE

Founding Engineer - Stealth Aviation Tech Startup | Sept 2025 – Present

- Architected ML extraction pipeline achieving **90%+ accuracy** on aviation compliance docs using AWS Bedrock; optimized **Lambda 65% (8s→2.8s);** serving 2 international clients processing **500+** documents with **90%+ test coverage across 15K LOC**

Software Development Intern - GoFlyy | Feb 2026 – Present

- Selected through **4% acceptance** rate (3/70); building CV pipeline for automated garment condition assessment with **90%+ defect classification accuracy**
- Optimized inference latency **67% (2.1s→680ms)** via TensorRT INT8 quantization; building virtual try-on system to reduce return rates

AI Safety Research Fellow - UC Irvine SPAR | Jan 2026 – Present

- Competitively **selected (15% acceptance rate**) for 8-week fellowship researching reward misspecification, deceptive alignment, and instrumental convergence
- Investigating reward hacking in vision-language models; documenting failure modes where models exploit proxy metrics while violating intended behavior
- **Red-teaming GPT-4, Claude 3.5, Gemini Pro** for deceptive capabilities and refusal inconsistency; contributing to emerging safety benchmarks
- Exploring mesa-optimization risks and scalable oversight techniques including debate-based and recursive reward modeling

Outreach Officer - AWS Cloud Club, UC Irvine (Founding Member) | Sept 2024 – Present

- **Competitively selected for Amazon-funded AWS** Cloud Club; grew from 15 founding members to 80+ students; coordinate technical workshops **(40+ attendance)**, speaker recruitment, and certification initiatives

Software Engineering Team Lead - C2S Technologies | July 2025 – Sept 2025

- Won **1st place among 11 engineering teams,** securing **$10K seed funding**; led **3 engineers** building investor-startup matchmaking platform
- Built full-stack app (Next.js, Firebase, MongoDB) and mobile app (React Native); implemented AI matching engine with semantic search
- Platform enabled **50+** mock investors to discover startups for micro-investment crowdfunding campaigns

PROJECTS

High-Performance CUDA Kernels for Matrix Operations | C++, CUDA, Python, PyTorch

- **3x+ throughput** via shared memory tiling, warp primitives, bank-conflict elimination; 40%+ bandwidth reduction through register blocking/coalescing; integrated as **PyTorch C++ extension**; validated via nsight-compute

Custom PyTorch Data Loader Optimization | Python, PyTorch

- Profiled data loading bottlenecks in training pipeline; implemented parallel data augmentation with multiprocessing; added smart batching for variable-length sequences; reduced data loading time from **40% to 10%** of training time

Image Classification with CNNs | Python, PyTorch

- Trained **ResNet-18 on CIFAR-10 (60K images, 10 classes**) with data augmentation; implemented training loop with learning rate scheduling and early stopping; achieved **90% test accuracy**; deployed model with ONNX for inference, learning computer vision fundamentals and transfer learning