Alex Woelkers

📞 858-353-3997   ✉ amwoelkers@icloud.com   in linkedin.com/in/alex-woelkers   🔗 alexwoelkers.com

Technical Skills

Languages: Python, Java, TypeScript, C (Proficient), C++ (Proficient).
Infrastructure & Systems: Kubernetes, Docker, AWS (EKS, EC2, S3, DynamoDB, ECR, CloudWatch), RabbitMQ, Redis
Backend & Tooling: Spring Boot, REST APIs, JUnit, Jenkins, GitHub Actions (CI/CD), Prometheus, Grafana.
ML (Applied): PyTorch, Tensorflow, OpenCV, LLM integration.

Relevant Experience

Esri
Software Development Engineer - Intern

June 2025 - Sept 2025
Redlands, CA

- Enabled standards-based GIS interoperability across European enterprise deployments by designing and shipping **9** OGC-compliant Java REST endpoints, reducing reliance on proprietary interfaces and achieving **99.6%** uptime in integration testing.
- Validated end-to-end cross-platform client compatibility (Leaflet, QGIS), ensuring open-source ecosystem support without unnecessary backend changes.
- Reduced CI build times from ~**2m30s** to ~**25s** by customizing and optimizing JUnit-based validation workflows within Jenkins, improving developer iteration speed and reducing feedback latency during deployments.

Inter-networking Research Group
Undergraduate Research Intern

July 2024 - Present
Santa Cruz, CA

- Redesigned a socket-based messaging layer by introducing a RabbitMQ broker with per-client queues, eliminating a single-threaded routing bottleneck and improving multi-client packet throughput by **17%** in benchmarks.
- Designed failure-handling for broker unavailability, daemon timeouts, and Redis-backed large-payload storage, improving reliability in distributed IoT and autonomous simulation environments.
- Built edge-based AI prediction models within a federated learning framework (**200** simulated devices), enabling real-time wildfire detection experimentation.

Projects

GridAI | *Personal Project*

- Designed and deployed an hourly Kubernetes-based ingestion pipeline on AWS EKS, processing **5K-7K** datapoints per run across **300+** grid metrics, persisting raw data to S3 and normalized timeseries to DynamoDB.
- Reduced manual grid analysis time from **3-4** hours to under **5** seconds by automating ingestion, storage, and reporting with CronJobs, CloudWatch monitoring, and structured LLM summarization.
- Engineered failure isolation (non-blocking LLM stage, retry policies, IAM debugging) to ensure pipeline continuity despite partial dependency failures.

Inlyne | *Capstone Project*

- Reduced developer onboarding time by **35%** by designing a real-time documentation platform with Java (Spring Boot) backend and PostgreSQL, deployed via CI/CD pipelines to cloud infrastructure.
- Led backend development within a 6-member team, defining API contracts and deployment workflows, enabling real-time editing at **100 edits/collaborator**, tested with **7,500** concurrent users.

Image Classification Model | *Class Project*

- Achieved top performance (**94.98%** accuracy, top team of 102) in image classification competition by optimizing GPU training, hyperparameters, and batch scheduling with PyTorch and OpenCV to improve baseline accuracy by **34%** compared to class 60% average.

Education

University of California, Santa Cruz
Bachelor of Science in Computer Science (GPA: 3.92 / 4.00)

Sept 2022 - June 2026
Santa Cruz, California

- **Relevant Coursework:** Data Structures and Algorithms (C++, Python), Computer Systems Design (C), Distributed Systems (C, Python), Operating Systems*, Embedded Systems*, Computer Networking (C, Python), Computer Architecture, Linear Algebra, Databases, Computer Vision