Aman Nindra

US Citizen | (925) 605-9936 | amannindra@gmail.com | linkedin.com/in/amannindra | github.com/amannindra

Education

University of California, San Diego — La Jolla, CA

Bachelor of Science, Data Science — *Expected Spring 2028*

Relevant Coursework: Data Structures, Advanced Programming, Computer Organization and Assembly Language, Algorithm Design and Analysis Course, Linear Algebra, Calculus I–III, Probability & Statistics

Experience

Data Science Challenge Intern — July 2026 – Present

Lawrence Livermore National Laboratory — *Livermore, CA*

- Developing **agentic AI pipelines** built on foundational **vision and language models** to automate scientific workflows for **materials discovery and additive manufacturing**, coordinating specialized agents across data exploration, visualization, computer vision, and analytical tasks.
- Building task-specific AI agents that process real-world materials and manufacturing datasets, perform structured analysis and visual interpretation, and collaborate within multi-agent workflows to generate scientifically meaningful insights for researchers.

Undergraduate Researcher — May 2026 – Present

University of California, Merced

- Built a real-time ego-lane detection pipeline for autonomous vehicles using **LaneATT** trained on the **CULane** dataset, achieving **0.79 F1** (**0.84** precision, **0.75** recall) on the validation set; trained on a **SLURM**-managed **NVIDIA H100** cluster.
- Engineered a **COCO** preprocessing pipeline distilling 80 classes down to the 4 relevant to driving (person, vehicle, traffic light, stop sign) to train a **YOLOv11** detector for obstacle and traffic-control awareness.
- Integrated lane detection and object detection into a single video-inference script with frame pre/post-processing, producing annotated driving footage for qualitative evaluation of both models.

Software Engineering Intern — May 2026 – Present

Valley Institute for Sustainability, Technology & Agriculture

- Architected a machine learning training platform on a **3-node Kubernetes cluster**, including a control-plane node that manages cluster orchestration, **PostgreSQL** user data, and **70TB** of persistent storage, along with **2 GPU worker nodes** used for training. The platform enables **40+ professors** and **50+ PhD researchers** to train computer vision and NLP models.
- Built **Terraform** workspace templates that let users configure and save custom machine learning environments, including PyTorch, TensorFlow, and Conda. Users can also select hardware resources such as CPU cores, RAM, and up to **2x NVIDIA A30 GPUs** per training job. Each configuration provisions a dedicated development environment accessible through browser-based code-server or through the user's local VS Code using the **Coder** remote extension.

Software Engineering Intern — Aug. 2025 – Jan. 2026

Arrowz (Mental Health Startup) — *Remote*

- Restructured a **24,000-line** offline-first **Flutter/Dart** codebase into feature-based modules, eliminating duplicated legacy screens and consolidating navigation into a single reusable drawer component shared by **6** core screens.
- Built the app's complete first-run experience: a **10-step** onboarding questionnaire with per-page input validation and **Isar** local persistence, **Firebase**-authenticated routing to gate new users, and an interactive **ShowcaseView** walkthrough with database-backed progress tracking across **5+** screens.
- Shipped a goals-tracking feature end-to-end (**Isar** data model, **Provider** state management, home-screen UI) and built **4** custom date-range selector widgets (day/week/month/year) for the emotion-analytics dashboard.

Projects

Self-Hosted Distributed AI Training Platform | *PyTorch, Kubernetes, FastAPI, React.js, Linux* — Feb. 2026 – Present

- Building a self-hosted AI training platform that lets users upload datasets (S3-style object storage), configure hyperparameters, and train models through a web interface.
- Implementing GPU-aware job scheduling with **Kubernetes** to dynamically allocate training workloads across multiple Linux machines based on real-time VRAM and compute availability.
- Enabling concurrent model training on a single GPU when memory permits, maximizing hardware utilization through resource isolation and process-level management.
- Deployed a **Raspberry Pi** as a lightweight control node hosting the website, managing job queues, and persisting uploaded datasets to attached external storage.

Technical Skills

Languages: Python, JavaScript/TypeScript, C++, Dart, SQL, HTML/CSS, TailwindCSS
AI/ML: PyTorch, Scikit-Learn, OpenCV, NumPy, Pandas, YOLO, Transformers, Natural Language Processing, Computer Vision, TensorRT, ONNX Runtime, AI Automation,
Cloud & Tools: AWS, Google Cloud Kubernetes, Terraform, Slurm, Firebase, Flutter, Flask, FastAPI, Next.js, React.js, Docker, Git, Red Hat Enterprise Linux, Ubuntu