

Aman Nindra

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

EDUCATION

University of California, San Diego

La Jolla, CA

Bachelor of Science, Data Science

June 2028

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

TECHNICAL SKILLS

Languages: Python, C++, C, JavaScript/TypeScript, SQL, Bash/Shell

Deep Learning & Computer Vision: PyTorch, TensorRT, ONNX Runtime, Computer Vision, Natural Language Processing, Transformers, YOLO, NumPy, Pandas, Matplotlib, Scikit-learn, OpenCV

GPU & Systems Optimization: NVIDIA Jetson, JetPack, ROS 2, Real-Time Inference, GPU Performance Profiling, Distributed Training

Cloud & HPC: AWS, Google Cloud, Linux, Docker, Kubernetes, Slurm, Terraform

Frameworks & Tools: React.js, Next.js, Flask, FastAPI, Three.js, Git, Firebase, MCP, Claude Code, Codex, Copilot

EXPERIENCE

Data Science Intern

July 2026 – August 2026

Lawrence Livermore National Laboratory

[Poster Link](#)

- Built the core of a multi-agent scientific AI system for detecting defects in additively manufactured metal lattices, aligning volumetric X-ray CT scans with nominal 3D design geometry and automatically measuring and classifying structural components.
- Automatically analyzed all **18,468** struts in a lattice scan against its nominal design, identifying **118** missing and **256** broken candidate struts with results traceable to the original CT volume.
- Extended an open-source 3D scientific-computing pipeline with image segmentation, geometric alignment, centerline analysis, and skeletonization, and built an interactive Streamlit and Three.js interface for reviewing defect candidates.

Undergraduate Researcher at R&N Lab

May 2026 – Present

University of California, Merced

[Research Link](#)

- Built a real-time autonomous-vehicle perception pipeline using LaneATT with a ResNet-34 backbone, improving CULane lane-detection F1 from **0.77** to **0.79** through custom data augmentation.
- Trained a YOLOv11n object detector for pedestrians, vehicles, traffic lights, and stop signs by adapting the COCO dataset, achieving **0.650 mAP50** and **0.465 mAP50-95**.
- Profiled and optimized deep-learning inference on an NVIDIA Jetson Orin Nano Super, benchmarking ONNX Runtime with CUDA against TensorRT FP16 and reducing LaneATT latency by **4.55x** (**125.7 ms** to **27.6 ms**) and YOLO latency by **1.83x** (**30.3 ms** to **16.6 ms**).

Software Engineering Intern

May 2026 – August 2026

Valley Institute for Sustainability, Technology & Agriculture

Remote

- Architected a GPU-accelerated machine-learning platform on a **3-node Kubernetes cluster** with **70 TB** of persistent storage and **2 GPU worker nodes**, enabling **40+ professors** and **50+ PhD researchers** to train computer vision and NLP models.
- Built Terraform-based development environments that allowed researchers to configure CPU, RAM, and NVIDIA GPU allocations and provision reproducible PyTorch, TensorFlow, and Conda workloads through browser-based or remote VS Code environments.

Machine Learning Intern

Jan 2026

Machine Learning Club

[GitHub](#)

- Fine-tuned BERT for multi-label emotion classification across **28 categories** on **206K Reddit comments**, reaching **0.712 Macro F1** and **97.0%** label-wise accuracy over 25 training epochs.
- Developed per-class threshold calibration using grid search across **0.01–0.99** for each output label, replacing a single global decision threshold and improving prediction quality for minority emotion classes.
- Built an evaluation pipeline measuring Macro F1, Jaccard similarity, label-wise accuracy, and exact-set accuracy, achieving **0.688 Jaccard** similarity and predicting the complete correct emotion set for **47.7%** of samples.

Machine Learning Intern

Nov 2025 – Feb 2026

Data Science Society

[GitHub](#)

- Placed **4th** on DrivenData's wildlife camera-trap leaderboard, reducing competition log loss from **2.18** to **0.67** by fine-tuning a **300M-parameter EVA-Large vision transformer** to classify **8 African species** across **16,488** images.
- Trained EVA-Large using PyTorch (DDP) and mixed-precision training across 4 NVIDIA Tesla V100 GPUs on AWS SageMaker, achieving **93.3% validation accuracy**.
- Built a per-class evaluation harness measuring precision, recall, F1, and top-3 accuracy every epoch, identifying blank-frame misclassification as the dominant error source and the only class remaining below **0.80 F1** across all four evaluated architectures.