

Aditya Hegde

San Jose, CA | 408-396-6553 | aditya.hegde@sjsu.edu | [linkedin.com/in/aditya-hegde712/](https://www.linkedin.com/in/aditya-hegde712/) | github.com/AdityaHegde712

Education

San Jose State University <i>Master of Science in Artificial Intelligence (GPA: 4.0)</i>	<i>San Jose, CA</i> <i>Aug 2025 – May 2027 (Expected)</i>
Woxsen University <i>Bachelor of Technology in Computer Science Engineering (GPA: 3.73)</i> <i>Dean's List: 8 semesters</i>	<i>Telangana, India</i> <i>Sep 2021 – May 2025</i>

Technical Skills

Languages: Python, HTML/CSS, JavaScript
ML/Data: Machine Learning, Deep Learning, PyTorch, Tensorflow, Transformers, Scikit-Learn, OpenCV, Keras
Frameworks: Flask, Django, ReactJS, ExpressJS, Tailwind CSS, Node.js
LLM/AI Tools: HuggingFace, OpenAI API, Gemini, Claude, LLM Integration
Tools: Git, Docker, GCP, AWS (Lambda, EC2), PyCharm, VS Code, Linux
Backend: REST APIs, FastAPI, Flask, SQL

Experience

Data Scientist Intern <i>SUHORA Technologies Pvt. Ltd.</i>	<i>Mar 2024 – Jul 2024</i> <i>Uttar Pradesh, India</i>
- Engineered an end-to-end maritime vision pipeline improving detection performance by 40%+ mAP via architecture customizations, curated SAR/Optical datasets (40,000+ images), and hyperparameter optimization, sustaining 0.93 Mean Average Precision. - Reduced inference latency from 7 minutes to <1 minute by implementing a hybrid C# + Python GPU-accelerated preprocessing and parallel worker architecture, enabling near-real-time geospatial surveillance. - Designed an event-driven NAS ingestion system processing 10,000 km² high-resolution satellite images, integrating geometric correction, normalization, uint8 compression, and tiling for distributed inference at scale. - Architected an in-memory reconstruction and overlap-aware postprocessing pipeline generating GIS-compatible GeoJSON outputs, supporting seamless downstream analytics integration. - Stress-tested distributed GPU workloads on 15B-pixel batches, achieving 45s processing per 10,000 km² image with runtime memory guards, overflow controls, and structured logging for production stability.	
Research Assistant <i>Woxsen University</i>	<i>Aug 2023 – Feb 2024</i> <i>Telangana, India</i>
- Optimized video frame interpolation models, increasing inference speed by 28.5% while improving PSNR and SSIM through architectural refinements and performance tuning. - Evaluated federated learning frameworks across 10+ edge devices, reducing convergence time by 15% and cutting communication overhead by 50% versus centralized training without accuracy degradation. - Engineered scalable Python-based data pipelines processing 50K+ multi-domain records (medical, satellite, environmental), enabling reproducible experimentation across distributed training setups. - Developed modular PyTorch and OpenCV training frameworks with reusable preprocessing, validation, and evaluation components, accelerating iterative experimentation and deployment readiness.	

Projects

Marine Debris Detector <i>PyTorch, Computer Vision, Git, Docker</i>	<i>Apr 2025</i>
- Engineered an end-to-end debris detection pipeline processing 1000+ Sentinel-2 and PlanetScope images, training CNN models to achieve 90%+ mAP over baselines, improving large-scale marine monitoring reliability. - Designed a hybrid detection framework integrating NDWI, NDVI, and FDI rule-based geospatial fallbacks, reducing compute cost while maintaining detection robustness in low-confidence regions. - Optimized batch inference and structured JSON serialization workflows in Python, enabling scalable downstream automation and programmatic geospatial analysis. - Architected and containerized a Flask-based inference API with GIS-compatible vector export endpoints (QGIS/ArcGIS), deploying reproducible Dockerized pipelines for production-ready ML serving.	

Seg2Box Annotation Tool | *PyTorch, Algorithm Design, Git*

Nov 2024

- Developed a deterministic multi-stage annotation conversion pipeline executing in <1 ms per JSON, accelerating dataset preprocessing for large-scale vision training workflows.
- Replaced conventional IoU matching with IoMin-based box alignment, improving annotation precision by 13% and reducing preprocessing latency by 40%, increasing downstream training efficiency.
- Packaged as a reusable Python module, enabling automated integration into computer vision data pipelines.

Denoising Autoencoder | *PyTorch, Git, Python, OpenCV*

Dec 2025

- Engineered a lightweight CNN autoencoder achieving 35+ PSNR and 0.95+ SSIM, incorporating model pruning, quantization, and optimized data pipelines to deliver sub-30 ms inference per 512×512 image.
- Directed a 12-member team across model development and experimentation tracks, standardizing Git-based workflows and CI (GitHub Actions) to ensure reproducible training and rapid iteration.
- Presented project outcomes to AI/ML faculty panel, earning 2nd Best Project recognition for performance-driven architecture and deployment readiness.

Campus Involvement and Leadership

Chief Projects Officer, AI&ML Club

Feb 2026 – Present

San Jose State University

San Jose, CA

- Directed a portfolio of 17+ AI/ML projects spanning 70 active members and multiple GitHub repositories, instituting structured development standards and oversight to ensure consistent model experimentation and delivery.
- Recruited and led a team of 8 project officers, formalizing documentation, activity logging, and reproducible workflow practices to improve cross-project transparency and execution velocity.
- Standardized a reusable GitHub project template embedding best practices for ML experimentation, version control, and repository structure, elevating code quality and reproducibility across the organization.
- Secured and coordinated 5 industry-sponsored AI/ML projects, aligning academic teams with real-world problem statements and production-oriented development cycles.

PUBLICATIONS AND AWARDS

- A. Baggu, A. Hegde, H. Morayya, S. A. Yallapragada, and H. Mazumdar, “*Machine Learning-Based Space Risk Management: Asteroid and Solar Flare Prediction*,” 2024 IEEE 3rd International Conference on Electrical Power and Energy Systems (ICEPES), Bhopal, India, 2024, pp. 1–6, doi: 10.1109/ICEPES60647.2024.10653497.
- Baggu, A., Hegde, A., Morayya, H., Kumar, L., Sattarapu, P., Yallapragada, S.A. (2026). “*Deep Learning Based Dementia Detection on MRI Data*.” In: Pal, S., Malhotra, S., Gupta, I., Kumar, A. (eds) Emerging Technology and Sustainable Solutions. ICETSS 2024. Communications in Computer and Information Science, vol 2610. Springer, Cham. doi: 10.1007/978-3-032-11488-4_15