HIMANSHU DHURVE

+91-9699538032 himanshudhurve96@gmail.com linkedin.com/in/himanshu-dhurve-92057b183

SUMMARY

ML engineer with experience across model development, evaluation, and deployment, from fine-tuning and post-training alignment to building pipelines that go from raw data to production. I am drawn to problems where the gap between a working model and a useful system is still large, and where getting the evaluation right matters as much as the model itself. Looking to contribute to research that is grounded in real-world constraints and builds towards systems that are reliable, not just accurate.

EDUCATION

San Diego State University — **San Diego, USA**
Master of Science in Aerospace Engineering, GPA: 3.52/4.0 — *2021 – 2024*

Dr. Babasaheb Ambedkar Technological University — **Raigad, India**
Bachelor of Technology in Mechanical Engineering, GPA: 8.57/10.0 — *2014 – 2018*

PROFESSIONAL EXPERIENCE

AI/ML Research Intern, Vizuara AI Labs, India — *Jan 2026 – March 2026*

- Designed and deployed a **multimodal RAG evaluation framework** for technical documentation (OpenFOAM); introduced a **six-dimension LLM-as-a-judge benchmark** across groundedness, citation correctness, factual recall, and coverage, achieving overall scores **4.6+/5.0**; this **document processing pipeline** is directly applicable to form extraction and structured output retrieval from technical corpora. Submitted to *RAG4Rep Workshop, ACL 2025*. [openfoam-ai-guide]
- Developed a **signal processing pipeline** for SAR (Synthetic Aperture Radar) image analysis: fine-tuned **RF-DETR** (transformer-based detector) on a custom multi-class dataset achieving mAP **0.614**.
- Built an **agentic vision-language model (VLM) pipeline** converting research PDFs into structured digests by autonomously generating narrative summaries, AI-illustrated diagrams, and executable Jupyter notebooks via coordinated **VLM reasoning**.

Aerodynamics Research Assistant, San Diego State University, USA — *Aug 2021 – Dec 2023*

- Designed **large-scale data acquisition and analysis pipelines** on an HPC cluster (SLURM); processed high-dimensionality flow field datasets in Python, cutting simulation wall-time by **75%** through parallelization.

TECHNICAL SKILLS

Core Competencies	Vision-Language Models (VLMs), LLM Evaluation & Benchmarking, Post-Training Alignment, Agentic AI
Languages & Tools	Python, C++, PyTorch, TensorFlow, Scikit-learn, OpenCV, Hugging Face, Git
MLOps & Infra	Databricks, MLFlow, DVC, Docker, CI/CD, OpenAI SDK, CrewAI, FastAPI, uv
Research Methods	Dataset Curation, Benchmark Design, Human-in-the-Loop Assessment, Model Evaluation, Quantization

PROJECTS

End-to-End MLOps Pipeline for Customer Churn Prediction — *Apr 2026*

- Built a fully modular production pipeline on 243k rows: **XGBoost** with **Optuna** tuning (30 trials, recall-optimised), **MLflow** experiment tracking and model registry, **FastAPI + Gradio** serving, **multi-stage Docker** build, and **GitHub Actions** CI/CD, achieving **93% recall** at a tuned threshold of 0.30 on an imbalanced churn dataset. [repo]

LLM Evaluation via Multi-Signal Reward Engineering (Medical MCQ) — *Nov 2025*

- Fine-tuned **Qwen3-1.7B** with **GRPO** on structured medical QA using **multi-signal reward functions** (binary correctness, regex format validation, length-calibrated explanation scoring); managed full **SFT + GRPO pipeline** on constrained Kaggle GPUs with 4-bit quantization (LoRA rank 32) and **vLLM-accelerated sampling**, analysed failure modes and output distributions to identify model weaknesses. [repo]

Temporal Signal Stabilization for Video OCR (License Plate Recognition) — *Sep 2025*

- Built a real-time detection and recognition pipeline using fine-tuned **YOLO11** and **EasyOCR**; addressed **temporal signal noise** (OCR flickering) with a majority-voting algorithm over a 12-frame rolling buffer. [repo]

Video Instance Segmentation for Scene Analysis — *Sep 2025*

- Implemented per-frame object detection and pixel-level segmentation using PyTorch's pre-trained **Mask R-CNN** on video streams, producing per-instance coloured masks for audio-visual scene understanding and spatial activity localisation. [repo]

Multimodal Benchmarking, Student Misconception Detection (Kaggle MAP) *Aug – Oct 2025*

- Fine-tuned **Qwen2.5-Math-1.5B-Instruct** with **Unsloth** and a compact **MLP classification head** on curated math reasoning misconception benchmarks; improved MAP leaderboard score from 0.885 to **0.925** (+4.5%) by combining LLM semantic understanding with structured output evaluation. [repo]

PUBLICATIONS

- **Himanshu C. Dhurve** et al. "Decompose, Retrieve, Cite: A RAG Pipeline for Structured Report Generation from Technical Documentation." *RAG4Report Workshop, ACL 2026* (First Author).
- Joseph Katz, **Himanshu C. Dhurve**. "Effect of Surface Texture on the Lift and Drag of Small Spinning Balls." *AIAA Aviation Forum and Ascend 2025*. [Paper Link]

HONORS & CERTIFICATIONS

- Machine Learning Specialization, Stanford University & DeepLearning.AI *Jan 2025*
- Recipient of San Diego State University Master's Research Scholarship *May 2022*

Public Repositories: GitHub **Publications:** AIAA