DURGESH GOWDA C

AI/ML Engineer - Computer Vision, NLP & Generative AI

Bengaluru, India | durgeshcgowda@gmail.com | +91 98869 94431

linkedin.com/in/durgeshgowdac | github.com/durgeshgowdac | leetcode.com/durgeshgowdac

AI/ML engineer with an MTech in Machine Intelligence (CGPA 9.41), specializing in computer vision, NLP and generative AI systems that perform reliably under real-world conditions. Delivered a 10x faster satellite-data retrieval tool on ISRO's operational servers and a real-time multi-camera re-identification pipeline on NVIDIA DeepStream at the Indian Navy's INICAI. Drawn to the gap between AI research and what actually survives production: reliability and real impact first, elegance second.

TECHNICAL SKILLS

Languages: Python, C++, Java, Bash

AI / ML: PyTorch, TensorFlow, scikit-learn, Hugging Face Transformers, OpenCV, spaCy | Computer vision (detection, tracking, Re-ID), NLP, RAG, LLMs, multimodal models

Systems & Data: NVIDIA DeepStream, CUDA, GStreamer, FAISS, FastAPI, Streamlit, PostgreSQL/pgvector, MySQL, MongoDB

Cloud & Tools: Google Cloud Platform, Modal, Linux, Git/GitHub

EXPERIENCE

AI/ML Engineer Intern - Indian Navy Incubation Centre for AI (INICAI) *Oct 2025 – Dec 2025*

- Developed a real-time multi-camera person re-identification system for live CCTV feeds on the NVIDIA DeepStream SDK, running GPU-accelerated analytics across distributed camera nodes.
- Developed end-to-end inference pipelines on DeepStream's GStreamer framework in Python and C++, combining detection, multi-object tracking and appearance-based Re-ID in one low-latency streaming graph.
- Designed a multi-stage identity resolution pipeline with robust distance aggregation that keeps global IDs consistent across camera handoffs, occlusions and re-entries.
- **Tech:** Python, C++, PyTorch, DeepStream, CUDA, OpenCV, GStreamer.

AI/ML Project Intern - ISRO Telemetry, Tracking & Command Network (ISTRAC) *Feb 2023 – Jul 2023*

- Designed and delivered a full-stack NLP application for real-time satellite telemetry retrieval and visualization, improving mission control data access speed by **10x**; deployed on ISRO's operational servers.
- Fine-tuned RoBERTa (Hugging Face Transformers) for named entity recognition on satellite-mission vocabulary to extract technical parameters from natural-language commands.
- Integrated OpenAI Whisper with the NER pipeline, enabling scientists to query servers by voice or text.
- Built a semantic search engine (Whoosh + spaCy) replacing keyword-only lookup, and shipped it as a Tkinter desktop application in ISRO's live environment.
- **Tech:** Python, Tkinter, spaCy, PyTorch, Whoosh, OpenAI Whisper.

Open-Source Contributor - Simple Icons *Jan 2024 – Present*

- Contributed several brand icons and updates to a widely used open-source icon library. Maintained JSON schema compliance via standard GitHub PR workflows. Continue to contribute occasionally.

EDUCATION

M.Tech, Computer Science (Machine Intelligence) *2024 – 2026*
PES University | CGPA 9.41 | Prof. CNR Rao Scholar

B.Tech, Computer Engineering (Data Science) *2019 – 2023*
Presidency University | CGPA 9.19

SELECTED PROJECTS

FoundYou: Multi-Camera Person Re-Identification - github.com/durgeshgowdac/FoundYou
Python, PyTorch, YOLO, ByteTrack, OSNet-512, FAISS | *Preprint:* doi.org/10.5281/zenodo.23178297

- Built solo a pipeline that assigns one persistent global ID per person across camera feeds: per-camera YOLO and ByteTrack workers, OSNet-512 embeddings and FAISS-based reacquisition.
- Added robust gallery matching, an optional face veto, a track-merge pass and a runtime ID validator (SWAP, DUPLICATE, UNSTABLE), designed to limit false matches and fragmented IDs.
- On a synthetic two-camera benchmark: **MOTA 0.93, IDF1 0.89, 18–24 fps** on an Apple M1 laptop.
- Co-authored the research paper with Dr. Chaitra N (PES University).

GlyphScholar: Multimodal RAG over PDFs - github.com/durgeshgowdac/GlyphScholar
Python, FastAPI, Chainlit, Next.js, pgvector, Modal, MinerU | *Live demo:* glyphscholar.vercel.app

- Built layout-aware PDF ingestion (MinerU) that preserves text, tables, formulas and figures, with vision-language captions for uncaptioned crops.
- Designed hybrid retrieval (dense, BM25, image) with reranking and a full-page fallback, so answers can cite the source page.
- Made inference swappable (Ollama locally, Modal GPUs in production), traced requests with Opik, and deployed on Vercel, Render and Supabase.

RoleSense: AI Resume Screener - github.com/durgeshgowdac/RoleSense
Python, scikit-learn, TF-IDF, Streamlit | *Live demo:* rolesense.streamlit.app

- Compared five classifiers on a cleaned resume dataset; the best (Random Forest) reached about **82% accuracy across 43 job categories**.
- Built a hybrid score blending classifier confidence with TF-IDF/cosine job-description similarity, in a Streamlit app for batch screening of PDF, DOCX and TXT resumes.

MLXnarrator: Offline Narration Pipeline - github.com/durgeshgowdac/MLXnarrator
Python, MLX, Kokoro-82M, whisper.cpp, FFmpeg

- Built a tool that turns a text script into narrated, subtitled screen recordings on Apple Silicon, running offline after a one-time model download.

HACKATHONS & COMPETITIONS

Event	Project	Role	Result
CIDECODE Cyber Hackathon, 2025	**NoFakeAI**: deepfake detection	Team Lead	**6th of 341 teams;** top 40 nationally
TexPedition Hackathon, 2025	**AdComply**: ad compliance	Team Lead	**Finalist**
Smart India Hackathon, 2025	**FixMyNagara**: civic issue reporting	Team Lead	**Selected**
Tata Elxsi, 2026	**Code Bunnies**	Mentor	Mentored team

HONORS & CERTIFICATIONS

- Google Cloud Skills Diamond League (Gemini, Multimodal RAG, MLOps on Vertex AI, Kubernetes)
- SMOPS-2023 (Spacecraft Mission Operations) participant
- 600+ LeetCode problems
- 6,000+ Hyperskill problems