Abhijeet Dhupia

+919901656836 | hello@abhijeetdhupia.com | linkedin.com/in/abhijeetdhupia | github.com/abhijeetdhupia

Experience

Senior Machine Learning Engineer Jan. 2025 – Present
Bazaarvoice *Bangalore, India*

- Owned production migration of vector search from PostgreSQL/pgvector to Qdrant: designed cutover/backfill strategy, built ingestion + validation tooling, and scaled to **2.8B+ vectors** by migrating **4+ TB** of embeddings with operational monitoring and rollback readiness.
- Authored a production-focused benchmark (Aurora pgvector vs Qdrant) to drive the platform decision: compared indexing/ingestion cost, resource usage, and latency; validated HNSW + product quantization retrieval quality (**Precision@1 0.98**).
- Contributed to AI Review Summaries and Shopping Assistant initiatives: improved reliability of the vectorization → clustering → summarization pipeline serving **4 languages** and **1000+ clients**, and built an LLM RAG assistant over product descriptions + syndicated reviews for pilot rollout.

Senior Computer Vision Engineer March 2024 – Sep. 2024
Sentryc *Berlin, Germany (Remote)*

- Built YOLO-based detection pipelines for copyright/brand enforcement across **8 ECA clubs**, **5 NBA teams**, and **14 kit manufacturers**, driving a **99%** reduction in image processing costs.
- Reduced monthly costs by **$5k+** by developing in-house models and processing **∼3M images/month**; improved scalability by deploying on **AWS ECS (Fargate)** and using SQL for data management.

Computer Vision Engineer Oct. 2021 – March 2024
Big Vision *Bangalore, India*

- Delivered end-to-end ML systems across CV/NLP for multiple clients (collectibles pricing, OCR/text understanding, taxonomy/classification), translating ambiguous requirements into production pipelines and measurable model improvements.
- Built a document fraud detection pipeline achieving **90%+** accuracy; improved OCR throughput by **60%** via optimized preprocessing, batching, and inference performance tuning.
- Shipped production features powering iOS/Android experiences; contributed Prompt Engineering content to OpenCV's Generative AI courses backed by **$150k+** Kickstarter funding.

Deep Learning Research Assistant Aug. 2019 – Aug. 2021
Indian Institute of Science, Bangalore *Bangalore, India*

- Developed deep learning models for medical imaging (endoscopy segmentation; OCT pathology classification) in collaboration with IAF Command Hospital and Carl Zeiss India; partnered with QpiAI to integrate models into a cloud-based healthcare AI workflow.

Education

Manipal Institute of Technology Manipal, India
Bachelor of Technology in Electrical and Electronics, Minor in Data Science *July 2015 – June 2019*

Projects

Odometer Reading Detection | *Python, PyTorch, OpenCV, Git* *[GitHub]*
Built a YOLOv5-nano + OCR pipeline to read odometers; achieved **90.75%** accuracy on **3,460** test images after training on **6,000** labeled images; improved robustness to glare/blur via targeted augmentations and post-processing.

Publications

A. Dhupia, J. R. Harish Kumar, J. Andrade, and K. V. Rajagopal, "Automatic segmentation of lumen intima layer in longitudinal mode ultrasound images", in 2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC), 2020, pp. 2125–2128.

Technical Skills

Languages: Python, SQL, C++, Bash, LaTeX
ML/Infra: PyTorch, TensorFlow, OpenCV, FastAPI, Flyte, Docker, AWS, GCP, Git