

Abdul Manan

Islamabad, Pakistan | +92 327 0941379 | mananbyte@gmail.com | linkedin.com/in/mananbyte | github.com/mananbyte

EDUCATION

National University of Sciences and Technology (NUST)

Islamabad, Pakistan

Bachelor of Science in Computer Science

Graduation: 2027

- Relevant Coursework: Machine Learning, Deep Learning, Data Structures, Algorithms, Big Data, Software Engineering

EXPERIENCE

Computer Vision Intern

June – July 2025

Infix Systems

Islamabad, Pakistan (Full Time)

- Contributed to early-stage development of a vehicle speed estimation and license plate recognition system
- Annotated industrial safety data using Labelme
- Built backend for a Windows virtual camera app, streaming real-time UDP video via FFmpeg as a low-latency system-level source for Zoom, Teams, and Meet

Research Intern — Cephalometric Retrospective

June – August 2024

MachVIS Lab

Islamabad, Pakistan

- Learned the complete research workflow, including literature review and identifying research gaps
- Executed research-driven data workflows including academic paper retrieval, exploratory data analysis, and structured documentation using spreadsheets for literature synthesis

PROJECTS

AI Research Analysis Platform | *FastAPI, Qdrant, AWS (S3, Bedrock, SageMaker)*

- Building a multi-tenant SaaS platform parsing scientific PDFs via Docling into citation-grounded structured evidence, with AWS S3 storage
- Engineering hybrid retrieval (BM25 + Qdrant dense search) with BGE-M3 embeddings and cross-encoder reranking on AWS SageMaker
- Designing an agentic query router with sub-100ms intent routing and a Corrective RAG loop for low-confidence queries
- Enforcing zero-hallucination output via a RAGAS faithfulness gate validating every claim against retrieved chunks

HandCraft-Path: Classical ML for Nuclei Segmentation | *Python, RAPIDS (cuML, CuPy), XGBoost, LightGBM*

- Built a GPU-accelerated classical ML pipeline segmenting nuclei on 7,904 PanNuke images using 93 handcrafted features, without any neural network
- Designed a mmap-based streaming architecture with dynamic batching and crash-resilient checkpointing for memory-safe processing at scale
- Devised Active Boundary Mining, a GPU distance-transform sampling strategy prioritizing hard boundary pixels to cut edge false positives
- Reduced 93 features to 25 via GPU-accelerated RFE (12x speedup) and trained an RF/XGBoost/LightGBM ensemble achieving 0.8802 Macro-F1

Scalable Academic Policy QA System | *Python, LangChain, RAG, LLMs*

- Built a RAG system to answer academic policy queries from institutional documents
- Built ingestion, chunking, embedding, and retrieval pipeline for accurate context
- Integrated LLMs with retrieved context to reduce hallucinations and ground responses

Cepharix | *Next.js, React, TypeScript, Tailwind CSS* | cepharix.vercel.app

- Built a full clinical workflow prototype covering patient/case management, X-ray upload, AI landmark detection, interactive landmark refinement, and Ricketts cephalometric analysis with report generation
- Implemented an interactive drag-and-drop landmark editor with undo/redo and clinician verification, plus a model training center simulating active learning from clinician corrections

TECHNICAL SKILLS

Languages: Python, C, C++, Java, JavaScript/TypeScript, SQL

AI & ML: Machine Learning, Deep Learning, Computer Vision, Medical Image Analysis, RAG, LLMs, Agentic Pipelines

Frameworks: TensorFlow, Keras, Scikit-learn, OpenCV, YOLO, LangChain, FastAPI, React.js

Data & GPU Computing: NumPy, Pandas, RAPIDS AI (cuML, cuDF, CuPy), XGBoost, LightGBM

LLM Infra & Cloud: AWS (Bedrock, SageMaker, S3, App Runner), Qdrant, Docling, RAGAS, Instructor

Developer Tools: Git, GitHub, Linux, VS Code, Jupyter Notebook, Google Colab, Docker

Additional Tools: LabelMe, FFmpeg, Antigravity CLI, Microsoft Office, Google Workspace