

Mohammad Alhan Bellary

AI Systems & Backend Software Engineer

Bengaluru, India | +91 8867634490 | alhanbellary123@gmail.com | linkedin.com/in/alhan-bellary | github.com/alhaannn | alhan.me

PROFESSIONAL SUMMARY

AI Systems & Backend Engineer pairing algorithmic rigor (NeetCode 150) with production Django experience and IIT Dharwad AI training. Built agentic LangGraph systems, FastAPI/Celery backends, and Kubernetes MLOps pipelines. Targeting junior GenAI, backend, applied ML, and MLOps roles.

TECHNICAL SKILLS

Languages & Backend: Python 3.13, SQL, Bash, Django, FastAPI, Celery, Redis, Pydantic, SQLAlchemy, PostgreSQL

GenAI & Agentic AI: LangChain, LangGraph, LiteLLM, Agentic & Vectorless RAG, Ragas, DeepEval, Guardrails

ML & Deep Learning: PyTorch, TensorFlow, Scikit-learn, XGBoost, OpenCV, YOLOv8, ONNX Runtime, Pandas

DevOps, MLOps & Cloud: Docker, Kubernetes, Jenkins, GitHub Actions, MLflow, Prometheus, Grafana, AWS, Azure

CORE PROJECTS

Enterprise Multi-Agent Research & Knowledge Assistant | *LangGraph, FastAPI, Chroma/Qdrant, Ragas, Docker*

- Architected a cyclical LangGraph pipeline (Planner, Retriever, Fact-Checker) that cut context hallucination by 34%, as measured by Ragas/DeepEval faithfulness gates, by replacing single-pass RAG with verify-then-synthesize loops.
- Secured the pipeline, as validated by DeepEval regression evals, by implementing PII-scrubbing and prompt-injection guardrails with FastAPI SSE streaming.

High-Throughput AI Gateway & Task Processing Engine | *FastAPI, Celery, Redis, PostgreSQL, GitHub Actions*

- Engineered an asynchronous gateway sustaining 250+ concurrent requests at sub-50ms queue latency, by moving blocking LLM calls onto a Celery/Redis task queue.
- Eliminated single-vendor dependency, as verified by PyTest coverage above 80%, by implementing OpenAI-to-Groq fallback, token-bucket rate limiting, and LRU caching in a GitHub Actions CI/CD pipeline.

Multimodal Defect & Threat Inspection System | *PyTorch, YOLOv8, ONNX, TensorFlow Lite, FastAPI*

- Developed a defect-detection pipeline reaching 94.2% mAP across five defect categories, by fine-tuning YOLOv8 and fusing image embeddings with tabular metadata.
- Quantized the model with ONNX Runtime and TensorFlow Lite to reach 42 FPS on CPU/edge hardware, served via async FastAPI with anomalies logged to SQL.

PROFESSIONAL EXPERIENCE & TRAINING

AI Technocrat Industry Training (Mandatory OJT, 120 Hours) | **IIT DHARWAD** | Dharwad | Sep 27 – Oct 22, 2027

- Engineered Jenkins and GitHub Actions CI/CD pipelines, reducing manual releases, by adding automated rollback triggered by anomalous deployment telemetry.
- Containerized Python inference services with Docker and orchestrated multi-container Kubernetes clusters with Horizontal Pod Autoscaling.
- Instrumented Prometheus and Grafana dashboards tracking cluster resources and system uptime.

Junior Software Developer (Internship), Heagl Technologies Pvt Ltd | Hubballi | Feb – May 2026

- Delivered 5 features across 2 production Django codebases, cutting feature turnaround by ~25%, by collaborating with senior developers and code reviews.
- Resolved 15+ bugs, reducing debugging time by ~30%, by systematically diagnosing defects with AI coding tools (Claude Code, Cursor).

EDUCATION & CERTIFICATIONS

AI Technocrat, NSQF Level 5.0 Professional Certification | **IIT DHARWAD** | Oct 22, 2026 – Oct 22, 2027

BCA, Information Technology, IBMR Business School | Hubballi | Sep 2023 – Jul 2026 | CGPA 7.66/10

Certifications: Oracle Certified Data Science Professional, Oracle, 2025 | AI Literacy, IBM, 2025 | Certified Product Manager (CPM), GeeksforGeeks

Awards: 1st Place, CheTech Hackathon (Chetan College), 2026 | Empathy Encryption Hackathon (UnsaidTalks), 2025