

# Aadarsh Praveen Selvaraj Ajithakumari

Boston, MA — +1 781 408 7807 — praveen.aadarsh@gmail.com — LinkedIn — GitHub

## Summary

AI/ML Engineer and Data Scientist with **3+ years** of experience building production-grade machine learning systems, RAG pipelines and data engineering solutions. Proven track record across enterprise domains real estate, municipal government and financial services deploying end-to-end AI systems from vector search infrastructure to LLM orchestration on **AWS** and **Azure**. **2nd place, HSIL GrandHack 2026** (Intersystems Challenge).

## Skills

- **Generative AI & LLM:** OpenAI GPT-4o, LangChain, LlamaIndex, RAG Architectures, Agentic AI, Multi-Agent Systems, Prompt Engineering, RAGAS, Vertex AI Embeddings
- **ML & Deep Learning:** PyTorch, TensorFlow, Scikit-learn, HuggingFace Transformers, NLP, Transfer Learning, Clustering
- **Search & Retrieval:** Pinecone, FAISS, Qdrant, BM25, Hybrid Search, Dense Retrieval, Reranking, Semantic Search
- **Cloud & MLOps:** AWS (S3, EC2, Lambda, SageMaker), Azure (AI Foundry, App Service, Key Vault, SQL), GCP, Docker, FastAPI, MLflow, Airflow, GitHub Actions
- **Data Engineering:** Apache Spark, Polars, Pandas, PostgreSQL, MongoDB, DVC, ETL Pipelines
- **Languages & Tools:** Python, SQL, Pydantic, PyTest, OpenTelemetry, Git, Power BI, Tableau, Streamlit

## Work Experience

### AI/ML Engineer

Feb 2026 – Present

*IpsierLab LLC, Fort Worth, TX (Remote) — Smart Pantry, AI-Powered Meal Planning Platform*

- Building an LLM-powered meal recommendation engine with OCR-based grocery bill scanning, expiry-aware inventory prioritization, and auto-generated shopping lists using **Python, FastAPI, and React**.
- Designing LLM prompt pipelines incorporating dietary constraints (vegan, allergies, macros) and architecting full-stack infrastructure with **Supabase Auth** (OTP/OAuth), **Neon PostgreSQL**, and **Supabase Storage**.
- Conducting ML model evaluation with precision/recall/F1 metrics; supporting mobile deployment via **React Native/Expo** across iOS and Android.

### AI Developer

Jan 2025 – Jun 2025

*EasyBee AI, Boston, MA — Real Estate & Municipal Government AI Agents*

- Designed and maintained production **RAG workflows** for enterprise knowledge access over property records, policy documents, and municipal FAQs, deployed on **AWS (S3, Lambda, EC2)**.
- Built automated RAG evaluation framework using **Python asyncio** pipelines, achieving **85% accuracy** with **0.52s avg latency** across thousands of QA pairs using RAGAS metrics.
- Engineered **Pinecone** vector search with namespace isolation, metadata enrichment, and bulk upserts; built modular **FastAPI + Docker** service layers with **Apache Airflow** orchestration.

### Application Development Analyst

Oct 2021 – Jul 2023

*Accenture, Chennai, India — Financial Services & Enterprise Data Engineering*

- Developed and optimized **ETL pipelines** using **Python + Apache Spark**, improving data integration throughput by **30%** for enterprise financial services clients.
- Applied **KMeans and hierarchical clustering** to employee behavioral data, improving engagement metrics by **15%** through data-driven segmentation and personalized recommendations.
- Built **Power BI and Tableau** dashboards; mentored junior developers through code reviews and enforced Agile/Scrum delivery standards.

## Projects

### VabGenRx — Clinical Drug Safety & Decision Support Platform | *Live Demo*

Mar 2026

- Architected a multi-agent clinical decision system on **Azure AI Foundry** (Microsoft Agent Framework) coordinating 5 agents — safety analysis, contraindication, dosing, counselling, orchestration — generating reports in **under 90 seconds**. **2nd place, MIT GrandHack 2026 Intersystems Challenge**.
- Deployed on Azure (App Service, Static Web Apps, SQL, Key Vault, Blob Storage); configured **9-alert Application Insights** suite with **OpenTelemetry** tracing; implemented **HIPAA-compliant** RBAC, SHA-256 patient IDs, and 6-year PHI log retention.
- Achieved **100% evaluation pass rate** (relevance, coherence, fluency, groundedness) on 15/15 test cases; integrated **100+ language** patient counselling via GPT-4o with 30-day composite-key caching.

### AI Semantic Recommendation Engine | *Live Demo*

Dec 2025

- Built a production recommendation system over **6.7M Amazon reviews / 31K products**, achieving **85.4% NDCG@10** via hybrid retrieval (BM25 + BGE embeddings + cross-encoder reranker) with **sub-1s latency**.
- Deployed **Pegasus** summarization on **AWS SageMaker (4xA10G GPUs)**, generating 32K+ summaries; built **Qdrant Cloud** backend with 4-level LRU cache achieving **10–15x speedup** (545ms → 50ms). **\$25/month** operating cost, **99.5% uptime**.

### RAG Financial Analysis Chatbot

Dec 2024

- Production RAG system for financial Q&A using **LlamaIndex, LangChain, Vertex AI** (768-dim) embeddings on **GCS**; containerized with **Docker**, CI/CD via GitHub Actions, **DVC** versioning, and Slack alerting for pipeline failures.

## Education

### Master of Science in Data Analytics Engineering

Sep 2023 – Dec 2025

Northeastern University, Boston, USA

GPA: 3.8/4

### Bachelor of Engineering in Electronics and Communication Engineering

Sep 2017 – Jun 2021

RMK Engineering College, Chennai, India

GPA: 3.4/4