Amruth Pai Thukaram

469-465-4829 | amruthpaiuni@gmail.com | linkedin.com/in/amruthpai | github.com/Immortal-Pi

Summary

ML/AI Engineer specializing in agentic AI frameworks, multi-agent orchestration, and production ML systems across energy, manufacturing, and enterprise domains. 5+ years delivering scalable inference infrastructure and measurable business impact from hypothesis-generation engines to real-time forecasting pipelines. Translates ambiguous business problems into deployable, data-driven solutions.

Education

University of Texas at Dallas **Aug 2024 – May 2026**
Master of Science — Business Analytics & Artificial Intelligence *GPA 3.89*

Nitte Meenakshi Institute of Technology, India **Jun 2015 – Aug 2019**
Bachelor of Engineering — Computer Science Engineering *GPA 3.65*

Professional Experience

Data Science & AI Intern | *Carlisle Construction Materials, Pennsylvania* **Jan 2026 – May 2026**

- Architected an AI Ideation Engine for R&D and material science teams with dynamic agent spawning, orchestrating multi-agent reasoning (research, synthesis, critique, validation) using LLMs and tool-based workflows; reduced hypothesis-generation cycles from weeks to hours
- Implemented agent-level validation and critique loops with structured confidence scoring, grounding generated concepts against the material science team's property database for factual accuracy; reduced low-confidence proposals by 40–50%
- Scaled concept evaluation throughput by 10× and surfaced key priority risks in implementing generated concept strategies, enabling material science and R&D leadership to make informed go/no-go decisions on the innovation roadmap

AI Engineer Intern | *Pegasus Knowledge Solution Inc., Texas (Remote)* **Aug 2025 – Oct 2025**

- Designed and deployed an on-premise RAG system over 10K+ internal policy documents, enabling natural language querying with sub-2-second latency using vector search and contextual retrieval
- Built natural-language-to-SQL agents using a 1B-parameter LM, generating optimized queries against production databases; orchestrated multi-agent workflows via LangGraph integrating SQL analytics with Plotly visualization pipelines
- Deployed models on vLLM inference infrastructure, benchmarked throughput/latency trade-offs, and established serving patterns for scalable, low-latency AI workloads

Data Science Intern | *Van Brunt & Associates, Texas (Remote)* **May 2025 – Aug 2025**

- Designed an ML forecasting pipeline for ERCOT electricity load prediction, ingesting weather, grid, battery, and market price signals across 500K+ daily data points; achieved 18% improvement in forecast accuracy
- Automated ETL workflows using Apache Airflow with scheduled DAGs for data ingestion, feature engineering, model retraining, and prediction. Deployed on Heroku with model artifacts versioned on AWS S3
- Delivered 87% precision on 4CP event prediction and 30% reduction in peak load inefficiencies; generated automated daily/monthly reports that enabled advisors to proactively alert clients on billing events

Senior Software Engineer | *Infosys Limited, Bangalore, India* **Sep 2019 – Apr 2024**

- Built an ML-powered EOD tracking dashboard using classification models on workflow metadata (20K+ daily records) to predict task completion risk, improving cross-functional team efficiency by 10%
- Led a team of 3 engineers to migrate legacy Excel-macro reporting to automated SQL pipelines, reducing manual effort by 45% and eliminating recurring support tickets for financial operations
- Engineered ETL pipelines integrating upstream APIs with real-time dashboards, processing 100K+ daily financial records with automated data validation that reduced inconsistencies by 25%
- Drove integration of Java APIs with IBM iSeries systems for automated financial close-price retrieval, cutting client report delivery time by 60%; presented architecture to stakeholders and mentored 2 junior developers

Projects

CareerGuidanceAI: Serverless Agentic Career Platform github.com/Immortal-Pi/career-guidance-agent

- Built a fully serverless GenAI platform on AWS (Lambda, S3, API Gateway), combining agentic orchestration with vector-indexed university course catalogs and external job-market APIs to deliver personalized career guidance

SurgeSense: Production ML Forecasting System github.com/Immortal-Pi/SurgeSense

- Designed an ML pipeline for real-time surge price prediction (96% R^2), with Hyperopt tuning, MLflow experiment tracking, Docker containerization, and CI/CD via GitHub Actions deployed on AWS EC2

PlanMyTrip: Multi-Agent AI Travel Planner github.com/Immortal-Pi/PlanMyTrip

- Developed a multi-agent travel planning system using FastAPI and LangGraph, coordinating weather, budgeting, and itinerary agents with external API integrations for cost-optimized, personalized trip recommendations

Technical Skills

- **ML & AI:** Scikit-Learn, PyTorch, TensorFlow, Keras, PySpark, Hugging Face, LangChain, LangGraph, vLLM
- **LLM & GenAI:** Prompt Engineering, RAG, Fine-Tuning, Agentic AI, Multi-Agent Orchestration, Azure OpenAI, Claude
- **Cloud & MLOps:** AWS (EC2, S3, Lambda, SageMaker), GCP, Azure, Docker, CI/CD, MLflow, Airflow, DVC, W&B
- **Data & Analytics:** Python, R, SQL, NoSQL, MongoDB, Pandas, NumPy, Tableau, Power BI, Matplotlib, Seaborn
- **Core Competencies:** Statistical Modeling, Predictive Analytics, NLP, Computer Vision, Time Series (LSTM), A/B Testing, Experiment Design, ETL, Stakeholder Communication

Certifications

- AWS Certified Solutions Architect – Associate | IBM Professional Data Science | DeepLearning.AI TensorFlow Developer