

MANJUNATH POPURI

+1 (917) 815-9580 | manjunathpopuri2@gmail.com | [linkedin.com/in/manjunathpopuri](https://www.linkedin.com/in/manjunathpopuri) | New Hyde Park, NY

TECHNICAL SKILLS

- **Programming and Scripting:** Python, SQL, R, Java, C/C++
- **ML and AI Frameworks:** Torch, Tensorflow, Hugging Face Transformers, Scikit-Learn, XGBoost, OpenCV, Deep Learning (CNNs, RNNs, LSTM, GRU)
- **LLMs and NLP:** LLM Fine-tuning (FLAN-T5, Phi-3, ClinicalBERT), RAG, PEFT/LoRA, Context Engineering, Prompt Engineering, NLP Pipelines
- **Data and Analytics:** Feature Engineering, Statistical Modeling, Experimental Design, A/B Testing, Snowflake, dbt, Airflow, Pandas, Numpy, SciPy, PySpark, Power BI, Tableau
- **Cloud and MLOps:** AWS (SageMaker, EC2, S3, Lambda), Azure ML, GCP (BigQuery, Dataflow), Docker, Kubernetes, MLflow, CI/CD

PROFESSIONAL EXPERIENCE

Binghamton University, NY, USA

Jul 2025 - Present

AI/ML Engineer

- Designed and deployed end-to-end production ML pipelines on **large-scale healthcare datasets**, reducing clinical inference errors by 15% and enabling real-time decision support for clinical teams.
- Fine-tuned **LLMs (FLAN-T5, Phi-3)** with **human-in-the-loop** feedback loops and **PEFT/LoRA**; improved model accuracy by 15% on prior authorization and clinical classification benchmarks.
- Engineered **distributed ETL pipelines** processing **1M+ EHR and claims records**, applying advanced feature engineering and statistical analysis to surface actionable clinical insights.
- Designed and executed experimental workflows - **hypothesis formulation**, controlled data collection, evaluation metrics, and result interpretation - ensuring scientific rigor across all **ML iterations**.
- Built **model monitoring dashboards** to track performance drift, data quality, and inference reliability, enabling proactive detection of model degradation in production environments.

Binghamton University, NY, USA

Aug 2024 - May 2025

Graduate Research Assistant

- Built **clinical NLP pipelines** using **ClinicalBERT** to extract diagnostic and billing entities from unstructured EHR notes and insurance claims (MIMIC-III, CMS Open Claims); ensured **HIPAA-compliant** data handling throughout.
- Fine-tuned **quantized LLMs (Phi-3)** and trained ensemble classifiers (**XGBoost, Random Forest**) for prior authorization classification - approved, denied, flagged - optimizing for accuracy and clinical relevance.
- Performed **feature engineering** and **statistical analysis** on structured clinical data to identify high-signal predictive features, directly improving downstream model performance.
- Deployed **Power BI dashboards** to surface model predictions, classification trends, and evaluation metrics, enabling data-driven decision support for clinical operations.

Cognizant, Hyderabad, India

Jul 2021 - Jun 2023

Machine Learning Engineer

- Built production **fraud and anomaly detection** models (XGBoost, Random Forest, Gradient Boosting) on **10M+ healthcare claims**, improving detection accuracy by 12% over baseline.
- Developed predictive **readmission risk** models using **PyTorch and TensorFlow**, reducing patient readmission rates by 8% in clinical testing through advanced feature engineering and model calibration.
- Engineered and optimized **SQL & PySpark ETL pipelines** for real-time model training and feature store updates, cutting pipeline runtime by 20% and increasing data throughput.
- Managed end-to-end **model deployment lifecycle**, ensuring robust model maintenance and continuous **performance monitoring** while implementing **runtime optimization strategies** that improved inference efficiency in production systems.

EDUCATION

Binghamton University, State University of New York, Binghamton, NY, USA

Aug 2023 - May 2025

Master of Science - Computer Science (AI Specialization) | CGPA: 3.54/4.0

Courses: Machine Learning, Artificial Intelligence, Data Mining, Human Computer Interaction, Natural Language Processing

PROJECTS

AI-Powered Medical Assistant using RAG & LLMs

Aug 2025 - Mar 2026

- Architected a **clinical AI assistant** leveraging **Med-PaLM and ClinicalBERT** fine-tuned with PEFT/LoRA, integrating a multi-stage RAG pipeline to handle real-time clinical queries, patient summarization, and drug interaction checks - **reducing hallucinations** by grounding responses in curated medical knowledge bases.

Gen AI Clinical Summarization & Compliance

Jan 2025 - Apr 2025

- Developed a **production-grade NLP system** combining FLAN-T5, Phi-3, and ClinicalBERT with a **PPO-driven reinforcement feedback loop** to automate **ICD coding** compliance and clinical note summarization across **structured and unstructured claims data**, enforcing **HIPAA-aligned toxicity constraints** and factual consistency at inference time.

Medical Imaging Classification

Sep 2024 - Dec 2024

- Achieved 92% accuracy classifying X-ray/MRI abnormalities using **ResNet and MobileNetV3**; optimized preprocessing with **PySpark** and deployed as a real-time inference service via **MLflow and Docker**.

CERTIFICATIONS

- Generative AI for LLMs, AWS Cloud Foundations & Architecting, Microsoft Certified: Azure Data Scientist Associate