

Name: ACHYUTHA PRANAVI K

Role: Senior AI/ML Engineer | Senior Data Analyst

Email address: achyuthapranavi268@gmail.com

Contact: +1 (347) 921-0136

LinkedIn: <http://www.linkedin.com/in/achyutha-p-89a3b5216>



Professional Summary

- Results-driven **Senior AI/ML Engineer and Senior Data Analyst** with **11 years of experience** in **data science, machine learning, advanced analytics, SAS programming, and AI solutions** across **finance, healthcare, retail, insurance, and telecom** sectors, delivering measurable business impact through cutting-edge implementations.
- Expert in **cloud-native AI/ML solutions** on **Microsoft Azure (Azure ML, Databricks, Synapse Analytics, Cognitive Search, Azure OpenAI)** and **AWS (Redshift, S3, Lambda, Athena, Glue, QuickSight)**, architecting and deploying **scalable, production-ready ML/AI systems**.
- Skilled in **Python, TensorFlow, PyTorch, R, SQL, and cloud-based MLOps workflows (AWS SageMaker, Azure ML, Databricks)** for **developing, optimizing, and deploying ML/AI models**, including **real-time inference pipelines**.
- Advanced practitioner in **Large Language Models (LLMs), Generative AI, Cursor AI, and Agentic Workflows**, building **retrieval-augmented generation (RAG) pipelines** and **localization-focused AI solutions** using **LangChain, Semantic Kernel, FAISS, Azure Cognitive Search, LlamaIndex, Spring AI, and Agno**, enhancing accuracy and reducing hallucinations.
- Experienced in **advanced AI architecture patterns** including **RAG, Agentic RAG, MCP (Model Context Protocol), Function Calling, and Agent-to-Agent (A2A) communication** within **Azure environments**.
- Developed **transformer-based NLP models (BERT, GPT, Azure OpenAI)** for **document summarization, sentiment analysis, entity/risk detection, and clinical note processing**, enabling **automated compliance and research workflows**.
- Skilled in **fine-tuning, reinforcement learning (RLHF), prompt-tuning, and post-processing** to improve model alignment, accuracy, and relevance.
- Designed and deployed **multi-agent systems** using **LangChain, LangGraph, CrewAI, AutoGen, and Cursor AI** for **intelligent automation in finance, healthcare, and compliance reporting**.
- Applied **time-series forecasting models (ARIMA, Prophet, LSTM, Neural Prophet)** for **financial trend prediction, demand forecasting, and risk assessment** in lending, trading, and retail operations.
- Built comprehensive **MLOps and CI/CD pipelines** using **MLflow, Azure DevOps, Docker, and Kubernetes**, ensuring **reproducibility, monitoring, governance, and compliance** in high-stakes deployments.
- Engineered **feature engineering and ETL pipelines** with **PySpark, pandas, Azure Data Factory, AWS Glue, Delta Lake, and Alteryx**, delivering **high-quality datasets** that significantly improved model performance.
- Delivered **enterprise analytics and BI solutions** using **Power BI, Tableau, and QuickSight** for **real-time KPI tracking and executive reporting** across multiple industries.
- Implemented **responsible AI practices** with **SHAP, LIME, fairness metrics, and Azure AI Content Safety**, ensuring **explainability, transparency, and regulatory compliance (HIPAA, financial services)**.
- Designed **predictive models** for **fraud detection, claims management, and patient stratification**, achieving **30% reduction in false claims and 18% improvement in fraud detection accuracy**.
- Applied **knowledge graph models** with **Neo4j, NetworkX, and graph embeddings** to map **complex relationships in fraud detection and compliance analytics**.
- Conducted **advanced statistical analysis, A/B testing, cohort analysis, and segmentation**, contributing to **23% campaign targeting improvement and 12% revenue uplift**.
- Built **scalable data pipelines** processing **financial transactions, healthcare records, and retail data** using **Snowflake, Alteryx, and distributed computing frameworks**.
- Leveraged **self-supervised learning and contrastive learning** for **domain-specific pretraining** on financial and healthcare text corpora, improving **embedding quality** for downstream ML tasks.

- Developed **synthetic data generation workflows** with **Gretel.ai** and internal tools to address **imbalanced datasets** while ensuring **privacy, security, and compliance**.
- Proven expertise in **SQL query optimization and data warehousing** using **Amazon Redshift, Snowflake, SQL Server, MySQL, Oracle**, delivering **near real-time reporting** and **30% system performance improvements**.
- Conducted **sales funnel, conversion rate, and clickstream data analysis** using **Python (pandas, NumPy)** and **R**, identifying **growth opportunities across 1,800+ retail locations**.
- Integrated **external datasets** (e.g., **Nielsen, social media sentiment, market data, social determinants of health**) with internal data for **enriched analytical models and competitive insights**.
- Implemented **data governance, lineage tracking, and metadata management** using **AWS Lake Formation, Unity Catalog, and dbt**, ensuring **standardization and regulatory compliance**.
- Automated **analytical workflows** with **Python scripting, AWS Lambda, and Azure Functions**, reducing **manual workload by 40%** and enabling **real-time insights**.
- Experienced in **private equity AI platform development and integration**, leveraging **Azure infrastructure, Flask-based API services, and secure microservices design**.
- Delivered **cross-functional leadership and mentorship**, conducting **knowledge-sharing sessions on MLOps, Generative AI, Azure AI Studio**, and translating **complex AI models into actionable business insights**.

Technical Skills

Programming & Core Technologies	Python, Java, JavaScript, SQL, SAS, PySpark, Bash, R, Flask, Pydantic, MSSQL, Linux/Unix
Cloud Platforms	Microsoft Azure (Azure ML, Azure Databricks, Azure Synapse, Azure Cognitive Search / Azure AI Search, Azure OpenAI, Azure ML Studio), AWS (SageMaker, Redshift, S3, Lambda, Athena, Glue, QuickSight), Snowflake, Databricks
Data Visualization & BI	Power BI, Tableau, Amazon QuickSight, Matplotlib, Seaborn
Big Data & ETL	PySpark, pandas, Azure Data Factory, AWS Glue, SQL, Alteryx, Delta Lake, Apache Airflow
Databases & Warehousing	Amazon Redshift, Snowflake, SQL Server, MySQL, Oracle, Data Modeling, Query Optimization
Statistical Analysis & Business Intelligence	Advanced Statistics, A/B Testing, Cohort Analysis, Customer Segmentation, Market Research, KPI Development, Forecasting & Demand Planning
AI/ML Frameworks & Advanced Learning	PyTorch, TensorFlow, scikit-learn, Keras, XGBoost, Hugging Face Transformers, Cursor AI, CrewAI, AutoGen, LangChain, LangGraph, Semantic Kernel, Spring AI, Agno, Databricks, Self-Supervised Learning, Contrastive Learning, Pretraining (Financial Text), Prompt Engineering, Context Engineering, Chunking Strategies for RAG Pipelines, Functional Calling
NLP & LLMs	BERT, GPT, Azure OpenAI, RAG (FAISS, Vector DBs), Text Summarization, Sentiment Analysis, Entity Recognition, Localization QA, Multi-lingual NLP, Translation Consistency Checks
Time-Series Forecasting	ARIMA, Prophet, LSTM, Neural Prophet
Vector & Graph DBs	FAISS, Azure Cognitive Search / Azure AI Search, Neo4j, NetworkX, Graph Embeddings, Knowledge Graphs
MLOps & DevOps	MLflow, Azure DevOps, AWS SageMaker, Azure ML, Databricks, Docker, Kubernetes, Prometheus, Azure Monitor, CI/CD for ML, Real-Time Model Inference Pipelines, Unit Testing, Postman
Responsible AI	SHAP, LIME, Fairness Metrics, Azure AI Content Safety
Synthetic Data & Privacy	Gretel.ai, Differential Privacy, Data Augmentation

Educational Details

- Master of Science in Computer Science - University of Central Missouri (Aug 2011 - Dec 2012)
- Bachelor of Science in Computer Science - Lovely Professional university (Aug 2007 - Jun 2011)

Certifications

- Microsoft Certified: Azure AI Engineer Associate - 2023
- AWS Certified Machine Learning - Specialty - 2021

Work Experience

Client: Jefferies Financial Group Inc, New York, NY

May 2024 - Present

Role: Senior AI/ML Engineer

Responsibilities:

- Designed and implemented **AI/ML solutions on Microsoft Azure** using **Azure ML, Azure Databricks, and Azure Synapse Analytics** for scalable, secure model training, deployment, and governance.
- Led **end-to-end ML pipeline development** with **Python, Azure ML SDK, and MLflow**, integrating **CI/CD** and governance for reproducibility and operational excellence.
- Built and deployed **NLP/LLM models** for **financial document summarization, sentiment analysis, risk entity detection, and localization QA**, leveraging **BERT, GPT, Azure OpenAI**, and custom transformers.
- Architected **agentic AI workflows** using **LangChain, LangGraph, LlamaIndex, Spring AI, Agno, CrewAI, AutoGen, and Cursor AI**, integrating **Pydantic** for schema validation, functional calling, and **A2A (Agent-to-Agent) communication**.
- Developed **RAG pipelines** with **FAISS, Azure Cognitive Search (AI Search)**, applying **embedding, chunking, grounding, prompt engineering, and context engineering strategies** to reduce hallucinations and improve retrieval accuracy.
- Enhanced **LLM/NLP performance** through **prompt tuning, PEFT (parameter-efficient fine-tuning), RLHF, and post-processing optimizations**.
- Built **multi-agent systems** for **financial research automation, compliance reporting, and earnings analysis**, reducing analyst effort and improving accuracy.
- Applied **time-series forecasting models (ARIMA, Prophet, LSTM)** for **multi-asset financial trend prediction and risk forecasting** in lending and trading use cases.
- Engineered **feature pipelines** using **PySpark, pandas, Azure Data Factory, and Databricks Delta Lake**, enabling **advanced transformations and temporal joins**.
- Applied **knowledge graph modeling (Neo4j, NetworkX)** to map customer–transaction–instrument relationships for **fraud detection and compliance analytics**.
- Developed **REST APIs (Flask, Python)** and **hybrid Python-Java microservices** integrated with AI APIs for financial analytics; validated using **unit tests & Postman**.
- Applied **design patterns (MVC, Singleton, Factory)** for scalable AI microservices and enforced **application security (SSL/TLS, JWT, RBAC)**.
- Implemented **real-time inference pipelines** for high-frequency financial streams, reducing latency by **35%** and enabling **sub-second decision-making**.
- Built **synthetic data generation workflows** with **Gretel.ai and internal tools** to augment imbalanced datasets while protecting sensitive financial data.
- Created **monitoring dashboards** with **Power BI, Azure Application Insights, Prometheus** to track **model drift, prediction accuracy, and anomalies**.
- Researched **self-supervised and contrastive learning** for pretraining on proprietary financial corpora to improve domain embeddings.
- Applied **data mesh principles** for scalable ML architectures across **global financial domains**, enabling decentralized innovation.
- Integrated **Responsible AI practices** with **SHAP, LIME, and Azure AI Content Safety** to ensure transparency, fairness, and compliance with financial regulations.
- Participated in **Agile/Scrum ceremonies** (Planning, Standups, Retrospectives), collaborating across engineering, compliance, and trading teams.
- Created **executive reporting suites** in **Power BI and Azure Synapse Analytics**, delivering real-time insights on **trading volumes, market volatility, and portfolio performance**.

- Mentored engineers and analysts, leading **training on MLOps, Generative AI, Azure AI Studio, and foundation model adaptation**, establishing best practices for **production-scale AI in finance**.

Client: HCA Healthcare Inc, Nashville, TN

Nov 2022 - April 2024

Role: AI/ML Engineer

Responsibilities:

- Designed and deployed **ML models for patient risk prediction and stratification**, improving proactive care delivery and clinical decision support.
- Partnered with **clinical and engineering teams** to translate healthcare goals into **data-driven ML solutions** using **Azure Machine Learning Studio, Azure Synapse Analytics, and Azure Databricks**.
- Built **ETL/data pipelines** with **Azure Data Factory, PySpark, and SQL**, ensuring accurate ingestion and transformation of **large-scale healthcare datasets**.
- Developed **fraud detection models** for claims, achieving a **30% reduction in false claims** using supervised ML.
- Created **deep learning models** with **TensorFlow and PyTorch** for **classification, anomaly detection, and time-series forecasting** in healthcare operations.
- Integrated **explainable AI (SHAP, LIME)** to ensure transparency, compliance, and clinician adoption in predictive models.
- Built **NLP models** using **BERT, spaCy, and early Azure OpenAI APIs** for clinical note classification, entity extraction, and document summarization.
- Explored **early Retrieval-Augmented Generation (RAG) designs** with **LangChain + Azure Cognitive Search** to enable **clinical Q&A and compliance document retrieval**.
- Leveraged **Azure AutoML** for hyperparameter tuning and accelerated model experimentation.
- Implemented **MLOps pipelines** using **Azure DevOps, MLflow, Docker, and Kubernetes (AKS)** for reproducible and scalable deployments.
- Developed **monitoring dashboards** with **Power BI and Azure Monitor** for drift detection, accuracy tracking, and compliance reporting.
- Integrated **social determinants of health (SDoH)** datasets with internal records on **Azure Data Lake**, enhancing predictive models through feature engineering.
- Applied **transfer learning** to diagnostic imaging use cases, improving anomaly detection accuracy across medical scans.
- Supported **annotation workflows and active learning** for medical imaging and text datasets, reducing labeling overhead.
- Contributed to **data privacy & HIPAA compliance**, embedding security and governance controls across the ML lifecycle.
- Mentored junior engineers on **Azure ML workflows, Git version control, and MLflow tracking**.
- Delivered **presentations to clinicians and leadership**, demonstrating AI's role in **value-based care** and operational efficiency.

Client: Target Corp, Minneapolis, MN

Jan 2019 - Oct 2022

Role: Senior Data Analyst

Responsibilities:

- Led the design and implementation of **scalable AWS data pipelines** using **Amazon Redshift, S3, Glue, and Athena** to process and store **high-volume retail transaction data** across 1,800+ locations.
- Delivered **actionable insights** with **complex SQL queries, stored procedures, and SAS reports**, improving reporting speed by **35%** and enabling deeper customer behavior and product trend analysis.
- Created **executive dashboards** in **Power BI, Tableau, and Amazon QuickSight**, providing **real-time visibility** into **inventory optimization, supply chain efficiency, and sales KPIs**.
- Collaborated with **marketing, merchandising, and e-commerce teams** to align analytics with business goals, improving **campaign targeting** by **23%**.
- Conducted **sales funnel, conversion rate, and clickstream analysis** using **Python (pandas, NumPy)**, identifying friction points in the customer journey.
- Managed **data modeling projects** with **dbt and Snowflake**, ensuring consistency in dimensions/metrics and strengthening data integrity across teams.

- Designed **ETL workflows** in **Alteryx** to automate ingestion and transformation of financial and operational data.
- Leveraged **Snowflake data warehousing** to prepare **model-ready datasets**, laying groundwork for **future RAG/LLM use cases** in retail knowledge bases.
- Applied **advanced statistical models** and **demand forecasting (Prophet, ARIMA, regression models)**, contributing to a **12% revenue uplift in 2021** through optimized pricing and inventory.
- Automated recurring reporting with **Python scripting and AWS Lambda**, reducing manual workload by **40%** and improving accuracy.
- Led **data governance initiatives** with **AWS Lake Formation**, implementing **data quality checks, lineage tracking, and access controls** to meet compliance standards.
- Integrated **external datasets** (Nielsen, social media sentiment, market data) with internal sources, enriching analytics for **competitive intelligence and localized merchandising**.
- Partnered with DevOps to deploy **monitoring solutions via CloudWatch**, improving reliability and proactive issue detection.
- Conducted **churn and cohort analysis** with SQL and Python, helping identify high-value customers and shaping retention strategies.
- Implemented **CI/CD pipelines and version control** with **Git, AWS CodeCommit, and Jenkins** for analytics and reporting workflows.
- Performed **A/B testing, statistical significance testing, and confidence interval analysis**, enabling **data-backed marketing and product feature decisions**.
- Built **executive-level financial reporting dashboards** for C-suite, tracking **revenue, margins, ROI**, and supporting **quarterly board reviews**.
- Acted as a **liaison between business stakeholders and technical teams**, ensuring analytical findings were translated into **strategic business outcomes**.
- Delivered **training sessions** for junior analysts on **AWS tools, SQL, visualization best practices**, fostering a data-driven culture.

Client: Allstate, Northbrook, IL

Sept 2015 - Dec 2018

Role: Data Analyst

Responsibilities:

- Conducted **data analysis on policyholder behavior, claims trends, and underwriting data**, supporting **insurance pricing and risk assessments**.
- Leveraged **AWS services (S3, Redshift, Athena, Glue)** for large-scale insurance data processing in a secure cloud environment.
- Built and maintained **dashboards in Tableau, Power BI, and SAS Visual Analytics**, improving visibility into KPIs, claim volumes, and customer retention.
- Designed **ETL pipelines** using **AWS Glue, SQL, and Python**, automating ingestion from **CRM systems and actuarial datasets**.
- Collaborated with actuaries to develop **predictive models in SAS and Python (pandas, scikit-learn)** for **churn prediction and fraud detection**, reducing fraud by **18%**.
- Utilized **SAS statistical procedures (PROC REG, PROC LOGISTIC, PROC GLM)** to support actuarial pricing models and dynamic customer segmentation.
- Optimized **SQL queries in Redshift** to improve reporting performance by **30%** and deliver near real-time insights.
- Conducted **cohort analysis and customer segmentation** with SAS and SQL, directly improving cross-sell campaign effectiveness.
- Standardized **structured & unstructured data integration** across policy, claims, and customer interaction systems using **AWS Lambda + Python**.
- Participated in **data governance initiatives**, ensuring HIPAA and regulatory compliance with **SAS data quality checks and metadata management**.
- Applied **statistical anomaly detection** in **SAS and SQL** to flag unusual claim submissions.
- Partnered with DevOps for secure deployment of analytics workflows using **AWS IAM, EC2, and CloudWatch**.

- Delivered **ad hoc SAS/SQL reports** to support sales, marketing, and claims leadership in aligning business goals with analytics.
- Maintained **data dictionaries & metadata repositories** for transparency and cross-team usage.
- Regularly evaluated **new BI and ML tools**, helping Allstate evolve from **SAS-heavy workflows to cloud-native Python/AWS analytics**.
- Provided **training on SAS, SQL, and Tableau** for junior analysts, strengthening enterprise adoption of analytics practices.

Client: Ooma Inc, Sunnyvale, CA

Feb 2013 - Aug 2015

Role: Data Analyst

Responsibilities:

- Partnered with **engineering, operations, and marketing teams** to analyze **large-scale telecom datasets**, improving operational efficiency by **20%+**.
- Built and maintained **dashboards and reports in Tableau and Power BI**, providing executives with real-time visibility into **call quality, churn, and customer KPIs**.
- Executed **end-to-end data analysis projects**: data extraction (SQL Server, Oracle, MySQL), ETL transformations, and statistical evaluations.
- Conducted **churn and customer behavior analysis** using **Python (pandas, NumPy) and R**, leading to data-driven retention and loyalty strategies.
- Developed **predictive models** (logistic regression, decision trees) to **forecast call drop rates and optimize network resource allocation**, improving service quality.
- Implemented **data validation and cleansing routines**, improving reporting accuracy and consistency by **30%**.
- Performed **market segmentation analysis**, uncovering key demographics and usage trends to support **pricing and customer acquisition strategies**.
- Supported **telecom feature experimentation** by designing **A/B testing frameworks** and analyzing product adoption outcomes.
- Conducted **VoIP/call quality analytics**, identifying network bottlenecks and recommending improvements that enhanced **service reliability and customer satisfaction**.
- Designed **ETL pipelines** for structured/unstructured data to standardize telecom analytics workflows.
- Assisted with **database migration projects**, validating data integrity and ensuring compliance with industry standards.
- Produced **trend analyses for monthly/quarterly business reviews**, highlighting growth opportunities and operational risks.
- Delivered **training sessions** to stakeholders on interpreting BI dashboards and reports, fostering a **data-driven culture** across the company.
- Maintained **documentation of data models, processes, and business logic**, ensuring scalability and knowledge transfer across teams.