

# Viswas Puli

+1-718-223-3045

puliviswas1998@gmail.com

<http://www.linkedin.com/in/viswaspuli>

## SUMMARY

Results-driven Data Scientist and Production Support Specialist with 4+ years of experience building scalable ML pipelines, delivering intelligent automation solutions, and supporting enterprise-grade systems. Proficient in AWS SageMaker, Splunk, Python, and big data tools. Proven track record in anomaly detection, chatbot development, customer engagement, and application performance monitoring. Experienced in creating real-time dashboards and alerts using Splunk, automating root cause detection and incident reporting. Skilled in ITIL-based production support, large-scale data handling (5TB+), and cross-functional collaboration to translate technical insights into business outcomes.

## TECHNICAL SKILLS:

- **Programming Languages:** Python, R, Java, C++
- **Machine Learning & AI:** TensorFlow, PyTorch, Scikit-learn, Keras, XGBoost, PyCaret, SageMaker Canvas, Amazon Bedrock, LLMs, RAG workflows, Llama, Claude AI
- **Data Analysis & Visualization:** SQL, Power BI, Tableau, Splunk (Dashboards, Alerts, Reports), Streamlit, Matplotlib, Toad
- **Big Data & ETL:** PySpark, Apache Spark, Apache Kafka, Hadoop, AWS Glue  
DevOps & Tools: JupyterHub, Git, Docker, Kubernetes, Jenkins, MLflow, DBeaver, TensorBoard, Bitbucket, AutoSys
- **Databases:** MySQL, MongoDB, PostgreSQL, SQL Server, AWS Redshift
- **Cloud Platforms:** AWS (S3, EC2, Lambda, RDS, SageMaker, Athena, Redshift, Bedrock), GCP, Azure
- **Version Control & CI/CD:** Git, GitHub, Bitbucket, Jenkins, Docker, Kubernetes
- **Operating Systems:** Linux, Windows

## EDUCATION

Master's in information technology and management Jan 2021  
– Dec 2023 *University of Texas at Dallas*

Bachelor's in computer science and engineering May  
2016 – Jul 2020  
*Sathyabama University, Chennai, Tamil Nadu, India*

## WORK EXPERIENCE

### Production Support Engineer

Bank of America, Plano, TX Jun 2025  
– Present

At Bank of America, I serve as a Level 3 Production Support Engineer, applying my academic training in Information Technology and Management to support enterprise systems. My role involves monitoring application performance, troubleshooting production issues, and delivering IT solutions in alignment with ITIL best practices. I work extensively with Splunk, Dynatrace, SQL, and Linux while managing automation tasks and real-time diagnostics.

### Responsibilities:

- Supported critical Java Spring-based enterprise applications by analyzing backend code, identifying defects, and resolving production incidents.
- Created and optimized Splunk dashboards to visualize system performance metrics, user activity trends, and application health KPIs.

- Built real-time alerts in Splunk to proactively monitor anomalies, system errors, and performance degradation, reducing MTTR.
- Automated Splunk-based reporting workflows to deliver daily/weekly system health summaries to engineering and leadership teams.
- Designed custom Splunk queries to trace errors, filter logs, and extract root causes during high-severity incident investigations.
- Utilized Dynatrace for application monitoring, performance profiling, and detecting system bottlenecks in production workloads.
- Wrote advanced SQL queries for on-demand data validation, performance monitoring, and post-incident analysis.
- Managed and improved shell scripts for Linux-based system tasks, log archival, job monitoring, and alerting integration.
- Collaborated with SREs and infrastructure teams to support 24/7 high-availability production environments.
- Documented standard operating procedures and knowledge articles in alignment with ITIL practices to support consistent issue resolution and team onboarding.
- Participated in major incident reviews and worked on corrective action plans to prevent the recurrence of system failures.

### **Data Scientist**

*Verizon, Tampa, FL*  
– May 2025

Jul 2024

Led AI/ML initiatives to enhance Wi-Fi router performance, automate anomaly detection, and drive customer retention strategies. Leveraged Amazon SageMaker, Bedrock, Streamlit, Redshift, Splunk, and JupyterHub across large-scale AWS infrastructure.

### **Responsibilities**

- Built end-to-end ML pipelines in JupyterHub using AutoML packages (e.g., PyCaret) to detect anomalies in router performance data extracted from Amazon Redshift.
- Trained an unsupervised model to surface unknown patterns and anomalies; collaborated with SMEs to define labeled thresholds and develop a supervised classification model for known issue detection.
- Combined both models into a meta-model framework to enhance anomaly detection accuracy, identify top contributing features, and quantify percentage change leading to anomalies.
- Developed and deployed a chatbot using Amazon Bedrock, integrating Retrieval-Augmented Generation (RAG) workflows with LLMs for intelligent, automated customer support.
- Built LLM-powered RAG applications to resolve dynamic queries using internal knowledge bases and contextual embeddings.
- Utilized SageMaker Canvas to create no-code models for anomaly detection, pattern change identification, and churn prediction, improving proactive network management.
- Queried and processed 5TB+ of telemetry data from S3 using PySpark and supporting packages within SageMaker notebooks, enabling efficient large-scale data retrieval and preprocessing.
- Implemented a K-Means-based sampling strategy to create a balanced and representative 40K-row training dataset while avoiding data overlapping across multiple prediction datasets.
- Validated model performance across multiple months (July–October) to detect and address feature distribution drift, improving long-term model generalization.
- Optimized Parquet file extraction in S3 by selectively loading columns like rowkey and ts, reducing memory overhead and improving Spark job performance.
- Performed sentiment analysis on JIRA support tickets to identify customer pain points, enabling insights into emerging issues and support trends.
- Created interactive Streamlit dashboards to visualize router health, signal strength trends, and station-level analytics in real time for technical teams.
- Executed complex SQL queries in Amazon Redshift and DBeaver to support telemetry analysis and model evaluation.
- Designed and deployed Splunk dashboards for network KPI tracking, failure diagnostics, and predictive maintenance signal visualization.

### **Machine Learning Engineer**

At Omegasoft International Inc., I led the Next-Generation Customer Engagement initiative for the retail banking division, focusing on improving customer interactions through predictive modeling and personalized banking services. I developed and deployed machine learning models using AWS SageMaker, leveraging deep learning and NLP techniques to analyze customer behaviors across multiple channels. My work resulted in enhanced customer engagement and more tailored financial products, while also ensuring efficient data handling and compliance with industry standards.

**Responsibilities:**

- Led the development of machine learning models using AWS SageMaker for the Next-Generation Customer Engagement initiative, focusing on predicting customer financial behaviors and preferences.
- Designed and deployed real-time models, enabling personalized banking services based on customer interactions.
- Integrated NLP techniques to analyze customer feedback from digital platforms, enhancing the personalization of product offerings.
- Automated data collection and preprocessing workflows with Apache Spark, improving data readiness and accuracy.
- Developed a multi-channel data integration framework to create comprehensive customer profiles, combining data from mobile apps, web portals, and in-branch services.
- Collaborated with data engineers to build robust data pipelines using Apache Kafka, ensuring real-time data flow and updates.
- Monitored and optimized model performance using MLflow and TensorBoard, continuously refining models for better outcomes.
- Mentored junior data scientists in deep learning and NLP model development.
- Managed data lakes on AWS S3, ensuring efficient storage and access to large volumes of financial and interaction data.
- Worked closely with product management and marketing teams to translate complex models into actionable insights.
- Ensured compliance with data privacy regulations through advanced data governance frameworks.
- Utilized AWS Lambda for real-time data processing, enabling dynamic model adaptation.

**Data Scientist**

*TechMahindra, Chennai*  
- Nov 2021

Jun 2020

At TechMahindra, I played a key role in the E-commerce Personalization Initiative, a project designed to improve user engagement and sales through a recommendation system. I developed and deployed this system using advanced machine-learning algorithms on AWS EC2 and Lambda. The project required extensive data analysis, feature engineering, and the creation of ETL processes to ensure high-quality data for model training. I also worked on NLP projects for text classification and sentiment analysis, and implemented reinforcement learning algorithms for supply chain optimization, leading to reduced operational costs. My role involved close collaboration with engineering teams to ensure smooth integration of ML models into production systems.

**Responsibilities:**

- Played a key role in the E-commerce Personalization Initiative, developing and deploying a recommendation system to enhance user engagement and increase sales.
- Created and maintained ETL processes to support data collection, cleaning, and transformation for ML model training.
- Performed extensive data analysis and feature engineering, optimizing inputs for machine learning models.
- Developed and deployed NLP models for text classification and sentiment analysis, improving customer interaction and feedback analysis.
- Implemented reinforcement learning algorithms for supply chain optimization, resulting in reduced operational costs.
- Deployed machine learning models on AWS EC2 and Lambda, ensuring scalability and real-time processing.

- Collaborated with engineering teams to integrate ML models into production systems, maintaining consistent performance and reliability.
- Conducted A/B testing and model performance evaluation to ensure accuracy and effectiveness of predictions.
- Contributed to the development of custom data analysis tools to streamline data exploration and preprocessing tasks.
- Supported the implementation of CI/CD pipelines for seamless updates and deployment of ML models in production.
- Provided detailed reports to stakeholders, highlighting key insights and metrics derived from ML models.
- Implemented data encryption and secure data handling processes within AWS, ensuring compliance with industry standards and protecting sensitive customer information.
- Utilized advanced feature selection methods, such as Recursive Feature Elimination (RFE), to improve model performance by focusing on the most relevant data features.
- Monitored and addressed model drift over time, ensuring that deployed models continued to perform accurately as new data was introduced.