Kushal Maktal

+1 (469) 249-2839 | maktalkushal47@gmail.com | www.linkedin.com/in/kushal-maktal

Professional Summary

AI/ML Engineer with **8+ years of experience** designing, developing, and deploying **large-scale machine learning systems, LLM-powered applications, and Generative AI (GenAI) solutions**. Expertise in building **NLP-driven search engines, recommendation systems, and fraud detection models** leveraging **Hugging Face Transformers, GPT, BERT, and multimodal AI**. Adept at **fine-tuning LLMs, retrieval-augmented generation (RAG), and vector search** for AI-driven personalization. Experienced in **MLOps, CI/CD, and scalable data engineering** on **AWS, GCP, and Azure**. Enthusiastic about **solving complex real-world problems** using cutting-edge AI techniques, **explainable AI (XAI), AutoML, and Reinforcement Learning**.

Technical Skills

- **Programming & Frameworks:** Python (TensorFlow, PyTorch, Hugging Face, FastAPI, Scikit-learn), R, SQL, Java, Scala, C++
- **Machine Learning & AI:** Large Language Models (LLMs), Generative AI (GPT, T5, BERT), Retrieval-Augmented Generation (RAG), Supervised & Unsupervised Learning (XGBoost, Random Forests, K-Means), Deep Learning (CNNs, RNNs, GANs), Reinforcement Learning
- **NLP & Information Retrieval:** Named Entity Recognition (NER), Sentiment Analysis, Semantic Search (FAISS, Pinecone), Hugging Face Transformers, Embeddings, Text-to-Image Generation
- **Big Data & Data Engineering:** Hadoop, Apache Spark, Apache Kafka, AWS Glue, GCP Dataflow, Redshift, PostgreSQL, Snowflake, SQL Server
- **Cloud Platforms & Deployment**: AWS (S3, SageMaker, Lambda, Redshift), GCP (BigQuery, Vertex AI, Dataflow), Azure (Data Factory, Databricks, Synapse Analytics), Kubernetes, Docker
- **MLOps & DevOps:** CI/CD Pipelines, Airflow, MLflow, Model Drift Detection, AutoML, Explainable AI (SHAP, LIME), Auto-scaling, Cloud Monitoring (AWS CloudWatch, GCP Monitoring)
- **Visualization & Analytics:** Tableau, Power BI, Looker, Matplotlib, Seaborn

Professional Experience

AI/ML ENGINEER

RetailMeNot - Austin, TX **JULY 2023 – Present**

- Developed a **personalized recommendation system** using **LLMs and Generative AI (GenAI)** to improve deal relevance and user engagement.
- Integrated **Hugging Face Transformers (BERT, GPT, T5)** to enhance deal categorization, sentiment analysis, and search relevance.
- Built a **retrieval-augmented generation (RAG) system** leveraging **FAISS and Pinecone** for fast and accurate search results.
- Fine-tuned **GPT models** on transaction and clickstream data to generate dynamic deal descriptions.
- Designed an **NLP-based chatbot** using **Hugging Face pipelines** to provide real-time deal recommendations.
- Deployed AI models on **Vertex AI** with optimized latency, achieving a **sub-200ms response time** for deal suggestions.
- Implemented **multi-modal learning** to process text, images, and structured data for a richer personalization experience.
- Applied **SHAP & LIME** for **Explainable AI (XAI)** to enhance deal ranking and recommendation logic transparency.
- Built **auto-scaling inference pipelines** using **Kubernetes, FastAPI, and TensorFlow** to handle peak loads during

sales events.

- Automated **A/B testing** for AI-generated recommendations, leading to a **15% increase in user conversion rates**.
- Engineered **real-time data pipelines (GCP Dataflow, BigQuery, Apache Kafka)** to ingest and process millions of user interactions daily.
- Integrated **Reinforcement Learning (RL)** algorithms to dynamically adjust deal offers based on user behavior patterns.
- Leveraged **Hugging Face's Diffusers** library to experiment with **text-to-text and text-to-image generative models** for deal visualization.
- Applied **zero-shot and few-shot learning** techniques with **Hugging Face Transformers** for adaptive content recommendations.
- Built **LLM-powered embeddings for similarity search** to improve product and deal recommendations.
- Developed **synthetic data generation pipelines** using **GANs and Hugging Face datasets** to enrich training data.
- Implemented **MLOps best practices (CI/CD, model drift detection, AutoML)** to continuously refine AI model performance.
- Designed **AI-driven fraud detection models** using **Anomaly Detection, Autoencoders, and Hugging Face Transformers** to flag suspicious transactions.

Tech Stack: Python, TensorFlow, PyTorch, Hugging Face, GPT, BERT, T5, Vertex AI, GCP (BigQuery, Dataflow), FAISS, Pinecone, FastAPI, Airflow, MLOps, Reinforcement Learning, SHAP, LIME

ML ENGINEER

State of MD - Baltimore, MD **JAN 2022 – MAY 2023**

- Designed and deployed an **ML-based Medicaid fraud detection system**, reducing fraudulent claims and improving financial oversight.
- Developed a **Random Forest** and **XGBoost** model to classify fraudulent vs. non-fraudulent Medicaid claims with improved accuracy.
- Built an **Anomaly Detection** model using **Autoencoders** to flag unusual claims for further investigation.
- Designed **data preprocessing pipelines** using **AWS Glue** and **Redshift**, ensuring scalable processing of large transactional datasets.
- Conducted **Exploratory Data Analysis (EDA)** and **Feature Engineering** to improve fraud classification model performance.
- Integrated **real-time fraud alerts** into the state's financial monitoring system to assist auditors in identifying suspicious activities.
- Created a **Tableau-based analytics dashboard** to visualize fraud detection trends and track investigation outcomes.
- Worked with **government agencies and compliance teams** to ensure **Machine Learning (ML) models** aligned with regulatory requirements.
- Designed **automated model retraining pipelines** to adapt to evolving fraud patterns in Medicaid transactions.
- Developed a **Natural Language Processing (NLP) module** to extract insights from healthcare provider claims and flag inconsistencies.
- Reduced **false positive rates by 18%**, improving efficiency for fraud investigation teams.
- Implemented **privacy-preserving ML techniques** to ensure **HIPAA and state data regulations compliance**.
- Collaborated with **Data Engineers** and **Legal Teams** to establish the best **fraud detection model deployment** practices.
- Contributed to **policy-making discussions** on the adoption of **AI-driven fraud prevention** techniques.
- Conducted **training workshops** for government auditors and compliance officers on using **ML insights for investigations**.
- Integrated **graph-based ML techniques** to identify hidden relationships between fraudulent providers and suspicious claims.
- Developed a **real-time monitoring system** to flag high-risk providers and insurance claims, preventing fraudulent transactions before processing.
- Applied **Reinforcement Learning (RL) techniques** to optimize audit case prioritization, ensuring high-value fraud

cases were investigated first.

- Designed a **self-learning fraud detection model** that continuously adapts based on new patterns detected in Medicaid claims.
- Enhanced **geospatial analytics integration**, helping law enforcement track and predict fraud hotspots based on historical claim locations.

Tech Stack: Python, Scikit-Learn, XGBoost, Autoencoders, AWS Glue, Redshift, Tableau, PostgreSQL, FastAPI, Airflow, SHAP, NLTK, CI/CD Pipelines, Lambda, Bitbucket

DATA/ML ENGINEER

Stryker Corporation - Dallas, TX **OCT 2019 – DEC 2021**

- Developed a **Machine Learning-driven Quality Control System** to detect **manufacturing defects** in medical devices.
- Built **Classification Models (Random Forest, XGBoost)** to predict defective products using **sensor data, material logs, and production insights**.
- Designed an **automated data pipeline** to process **real-time data** from **IoT-enabled manufacturing equipment**.
- Implemented **Anomaly Detection Algorithms (Isolation Forests, Autoencoders)** to flag deviations in material consistency and machine operation.
- Developed a **Predictive Maintenance Model** using **Time-Series Forecasting** to reduce machine downtime.
- Integrated **ML models with Manufacturing Execution Systems (MES)** to improve **production line efficiency**.
- Deployed an **ML-driven Root Cause Analysis System**, identifying patterns leading to defects and optimizing **manufacturing parameters**.
- Created an **interactive dashboard in Power BI** to track **real-time defect rates and production insights** for process engineers.
- Reduced **defect rates by 25%** by proactively detecting inconsistencies in raw materials and production cycles.
- Lowered **production costs by 20%** by optimizing **manufacturing workflows** through **ML-driven process automation**.
- Conducted **feature selection** and **dimensionality reduction** to improve **model accuracy** while maintaining inference speed.
- Developed and deployed **FastAPI microservices** to integrate **ML-driven quality control insights** into existing systems.
- Collaborated with **Quality Assurance Teams** to implement **ML-backed Decision Support Systems for Batch Testing**.
- Ensured **compliance with FDA and ISO standards** by integrating **Machine Learning techniques** for **regulatory quality control**.
- Provided **ML training and mentorship** to **junior Data Engineers** and **Quality Control Specialists** to improve **AI/ML adoption** in manufacturing.
- Designed a **Computer Vision-based Defect Detection System** using **Convolutional Neural Networks (CNNs)** to enhance **product inspections**.
- Built a **Predictive Analytics Tool** that analyzes **supplier quality history** and predicts **potential material failures before manufacturing**.
- Developed an **Automated Root Cause Analysis System** that correlates **machine sensor anomalies** with **defect trends** to optimize **machine configurations**.
- Applied **Unsupervised Learning Techniques** to **cluster production failures** and identify new potential **defect patterns**.
- Created a **real-time quality feedback loop**, allowing **ML models** to self-adjust parameters based on **live factory conditions**.

Tech Stack: Python, Scikit-Learn, XGBoost, Random Forest, TensorFlow, AWS Redshift, Snowflake, FastAPI, PostgreSQL, Autoencoders, Docker, Jenkins, Airflow, IoT Sensors, Kubernetes, Power BI

DATA SCIENTIST

Sailpoint - Austin, TX **MAR 2019 – SEP 2019**

- Built **scalable ETL workflows on AWS** for **data extraction, preprocessing, and ingestion**, improving efficiency by **40%**.
- Developed **Predictive Maintenance Models** using **Machine Learning**, reducing **equipment downtime by 25%**.
- Applied **Clustering Algorithms (K-Means)** to segment **customer behavior**, enabling **better-targeted marketing**

campaigns.

- Built **APIs for real-time integration** of **Machine Learning models** into **Stryker's data platforms**.
- Automated **Time-Series Forecasting** for **Supply Chain Optimization**, leading to a **15% reduction in inventory costs**.
- Designed **Power BI dashboards** to provide **actionable insights** into **key operational metrics**.
- Conducted **Feature Engineering** and **Model Tuning for Gradient-Boosted Machines (GBMs)**, improving performance by **20%**.
- Led the **implementation of NLP models** for **document classification**, automating **compliance workflows**.
- Migrated **on-premises data systems to AWS**, enabling **cost savings and improved scalability**.
- Developed **Anomaly Detection Systems** to identify **inefficiencies in manufacturing processes**.
- Conducted **Root Cause Analysis** to resolve **data discrepancies** in **manufacturing datasets**.
- Designed **data pipelines** to preprocess **IoT sensor data** for **Machine Learning models**, improving **prediction accuracy by 30%**.
- Collaborated with **DevOps teams** to establish **CI/CD pipelines for model deployment**, reducing update downtime.
- Implemented **Explainable AI (XAI) techniques** for compliance, ensuring **transparent predictions in regulatory audits**.
- Conducted **Time-Series Analysis on Equipment Telemetry Data** to detect **early signs of wear and tear**.

Tech Stack: Python, TensorFlow, AWS (Glue, Redshift, SageMaker), Power BI, K-Means, Gradient-Boosted Machines (GBMs), FastAPI, Airflow, CI/CD Pipelines, PostgreSQL, Anomaly Detection, Time-Series Forecasting, NLP (Document Classification)

DATA ENGINEER

Cuspyd - Hyderabad, India **JULY 2015 – OCT 2018**

- Designed and implemented **scalable ETL pipelines** for processing data from **relational and NoSQL databases**, improving **data processing efficiency by 30%**.
- Migrated **legacy systems to the Hadoop ecosystem**, enabling the processing of **large datasets with high availability and scalability**.
- Automated **SQL query optimizations** for **data aggregation workflows**, reducing execution times by **45%**.
- Developed **data pipelines** to integrate **CRM tools with PostgreSQL**, ensuring **seamless reporting for marketing and sales teams**.
- Built **recommendation systems** using **Collaborative Filtering Techniques** to improve **product suggestions and drive customer engagement**.
- Conducted **data validation** and implemented **error-handling mechanisms** to improve **pipeline reliability**.
- Created **real-time dashboards using Power BI** to monitor **operational metrics**, providing **actionable insights to business leaders**.
- Applied **Clustering Algorithms** to segment **customer data**, enabling **targeted marketing campaigns** and improving **ROI by 20%**.
- Developed **MapReduce programs** to analyze **unstructured data** and identify **trends**, enhancing **decision-making processes**.
- Conducted **performance tuning on Hadoop jobs**, improving the **throughput of batch processing pipelines by 3x**.
- Collaborated with **cross-functional teams** to establish **data modeling standards**, ensuring **project consistency**.
- Wrote and maintained **detailed documentation for ETL workflows**, facilitating **knowledge transfer and scalability**.
- Developed **Anomaly Detection Scripts** for **real-time system performance monitoring**, preventing **data processing delays**.
- Conducted **Root Cause Analysis** of **data inconsistencies** and implemented **solutions to ensure data quality and integrity**.
- Integrated **data from third-party APIs** to enhance **reporting capabilities**, enabling **more comprehensive business insights**.
- Created and implemented a **Metadata Management System** to simplify **data lineage tracking and governance**.
- Trained **junior team members** on **Hadoop, Hive, and other Big Data technologies**, fostering a **collaborative learning environment**.
- Designed **scalable data schemas** to accommodate the **rapid growth of business data requirements**.
- Implemented **Data Lake Solutions using AWS S3**, optimizing **storage and retrieval of structured and unstructured**

data.

- Participated in **weekly sprint meetings** to prioritize and deliver **high-impact data engineering solutions aligned with business objectives**.

Tech Stack: Python, Hadoop, Hive, MapReduce, PostgreSQL, Power BI, SQL Server, AWS S3, Data Validation Pipelines, NoSQL (MongoDB, Cassandra), API Integration, Metadata Management