

NABEEL MUKHTAR SIDDIQUI

SUMMARY

- 20+ years of experience in architecture and development of large-scale enterprise applications
- 10+ years of experience in production ML/AI systems combining deep ML fundamentals (anomaly detection, NLP, clustering, deep learning) with cloud-native architecture expertise at scale.
- Currently expanding into the modern LLM/Agentic AI stack — RAG pipelines, vector databases, multi-agent orchestration, LLM fine-tuning, RLHF, and LLM serving and observability

SKILLS & ABILITIES

- **Machine Learning/Deep Learning:** PyTorch • Keras • scikit-learn • Hugging Face Transformers • spaCy • Anomaly Detection • NLP / NER • Recommender Systems
- **LLM & Agentic Stack:** RAG Pipelines • Vector Databases (pgvector / Chroma / Weaviate) • LLM Eval Design (RAGAS / DeepEval) • Agent Orchestration (LangGraph) • LLM Fine-tuning (unsloth) and RLHF • LLM Serving (vLLM / AWS Bedrock) • MCP • LLM Observability (LangSmith/PostHog) • Prompt Engineering
- **MLOps & Data Engineering:** AWS SageMaker/AWS Bedrock • vLLM • Docker/Kubernetes • MLflow • CI/CD for ML (GitHub Actions)
- **Cloud & Infrastructure:** Docker, Kubernetes, AWS (EKS/EC2/Lambda/SQS/SNS), Azure, Terraform, OpenTelemetry, Prometheus/Grafana
- **Backend & Architecture:** Java • Python • NestJS/Node.js • Spring Boot • Go • Microservices • Event-Driven Architecture • REST/GraphQL/gRPC
- **Databases:** PostgreSQL, MongoDB, Elasticsearch, Neo4J, Redis, Snowflake, MySQL, Oracle, SQL Server

CERTIFICATIONS

- AI-Native Data Engineering Professional Certificate – IBM (Coursera)
- Deep Learning Specialization - Coursera
- Machine Learning Engineering for Production (MLOps) Specialization – Coursera
- Software Architecture Principles and Patterns - SEI

EXPERIENCE

PRINCIPAL SOFTWARE ENGINEER, XYMBIA VENTURES INC

April 2023 – Present

- Architecting and developing Xymbya Planner, a calendar and productivity SaaS platform with workspace collaboration, real-time event streams (SSE).
- Leading Backend and ML/AI integration layer: NestJS microservices on AWS EKS, MongoDB, and Spring Boot microservices for Calendar and Search.
- Implementing structured observability (Grafana/Loki, Prometheus) for Kubernetes-hosted services and evolving toward OpenTelemetry-based distributed tracing.

PRINCIPAL SOFTWARE ENGINEER, QLIK

August 2018 – February 2023

- Worked as an architect in Catalog and Cognitive Data Services team at Qlik
- Participated in architecture review meetings with Qlik Architecture Council (QAC)
- Developed code and unit tests for cloud native micro services using Spring Boot/Go and Kubernetes
- Monitored performance metrics on stage and production environments
- Investigated and resolved production incidents and attended postmortem meetings with SRE team

DISTINGUISHED SOFTWARE ENGINEER, CREATIVE CHAOS

April 2015 – April 2023

- Developed the backend for Podium Data (www.podiumdata.com) application for Big Data Management on multiple platforms like Cloudera, BigInsights, MapR and Hortonworks. The application is also supported on single server platforms.
- Developed the backend for Podium Data Flow Engine that runs on MapReduce, Spark and Tez.
- Wrote the module for Podium Data Publish to multiple targets like FTP, SFTP, S3, HDFS etc. in multiple formats including PARQUET, AVRO, ORC and TEXT.
- Developed the Data Conductor module that allows user to control the levels of data management while retaining a rich and complete metadata catalog for the enterprise.
- Worked on the recommendation engine for Podium that uses a rule-based approach to show design, performance and usage tips to the user while designing a data flow.
- Added the ability to compute custom metrics and KPI scores for entities managed by Podium.

SENIOR CONSULTANT (ARCHITECTURE & MACHINE LEARNING), GFK ETILIZE INC.

August 2014 – April 2015

- Architected and developed the search and indexing engine for products.
- Worked on full text search of product rich content (documents and images) using Elasticsearch.
- Researched techniques for using Named Entity Recognition for product attribute extraction.
- Implemented kNN product classification using Elasticsearch MLT (More Like This) API.

- Researched Sentiment Analysis techniques for product reviews.
- Developed the ETL process for loading data from various heterogeneous data sources into a multi-tenant Elasticsearch data store.
- Developed the REST/HATEOAS API for product search using Spring REST.

PRINCIPAL DATA SCIENTIST, ELASTICA INC.

August 2013 – August 2014

- Architected and developed the machine learning workflow for anomaly detection subsystem of Elastica CloudSOC (<https://www.elastica.net/>).
- Implemented feature extraction and scaling from user logs related to session, geography, categorical distributions and work habits.
- Created methods for detecting outliers using K-means clustering.
- Implemented rule-based detection of outliers using Drools rule engine.
- Build exploratory models for other clustering/outlier detection algorithms including 1-Class Classification, ORCLUS, DBSCAN and LOF.
- Created visualizations for clusters and categorical distributions for exploratory data analysis.
- Developed sensitivity analysis method for finding sensitive outlier features.
- Implemented bootstrapping for generating synthetic data.
- Wrote functions for generating activity graphs from logs for graph analysis.
- Investigated supervised learning algorithms like Support Vector Machines and Deep Learning.

SOFTWARE ARCHITECT, CREATIVE CHAOS

April 2008 – August 2013

- Developed the anomaly detection engine of Compuware Outage Analyzer application for streaming data.
- Architected the data analytics and reporting platform for the Bradford application.
- Worked as a developer on the data collection module of the application and developed both pull-based and push-based data collection workflows.
- For BEZNext application (<http://www.beznext.com/>), developed model service to align and aggregate various metrics and characterize workloads into discrete classes.
- Developed the data collection and the database advisor modules for Microsoft SQL Server for BEZNext.
- Developed the auto discovery module for VMWare.
- Developed the workload management module for Teradata.
- Architected the backend of the face2face application.
- Implemented the social network analysis services using graph database.

SENIOR SOFTWARE ENGINEER, GFK ETILIZE INC.

October 2003 – April 2008

- Developed the taxonomy editor for the ConQuire CMS that allows user to create parameters, display attributes and templates and group them into various categories
- Involved in the development of multi-market/multilingual module for ConQuire

- Worked on the cross-sell/up-sell module of the application that suggested products similar to a product.
- Worked on the data synchronization module for synchronizing the data between various geographically dispersed servers

EDUCATION

N.E.D. UNIVERSITY OF ENGINEERING & TECHNOLOGY - KARACHI

Bachelor of Engineering (Computer Systems Engineering)
