

Meezaan Chishty

Full Stack Engineer • AI/ML Engineer • Backend Developer

syedmeezan905@gmail.com • +91 7426917505 • Udaipur, Rajasthan, India

linkedin.com/in/syed-mohammed-meezaan-chishty • github.com/Mizan-muf • mizanchishty.ddns.net

PROFESSIONAL SUMMARY

Full Stack and AI/ML engineer with hands-on experience building production systems — from FastAPI microservices and React/Next.js frontends to LLM pipelines, RAG architectures, and real-time data platforms. I have shipped systems used by real clients, fine-tuned models, and written the infrastructure code to keep them running. Currently pursuing an MCA at Jamia Hamdard University (GPA 8.95/10) while working full-time; looking for a team where engineering craft and real-world impact both matter.

EDUCATION

Jamia Hamdard University (*Jan 2025 – Present*)

Master of Computer Applications (MCA), Computer Science

GPA: 8.95 / 10

Manipal University Jaipur (*Oct 2021 – Apr 2024*)

Bachelor of Computer Applications (BCA)

GPA: 8.88 / 10 (approx. 3.55 / 4.00)

Dean's List Awardee — Highest GPA in cohort

EXPERIENCE

Full Stack AI/ML Engineer (*Feb 2025 – Present*)

FBSP — India – Udaipur

- Architected and shipped **FBSPAI** — a production-grade AI insurance sales platform with **3 Next.js portals** and **10 FastAPI microservices**, serving live clients; full stack ownership from database schema to UI components.
- Built an ACORD document extraction pipeline using computer vision, OCR, and visual document models — replacing a largely manual process and feeding structured data into downstream reporting systems.
- Delivered a distributed real-time monitoring platform across **700+ endpoints** with sub-**200 ms** WebSocket updates, giving operations teams live visibility they previously lacked.
- Currently leading development of an insurance-domain LLM fine-tuned on proprietary data, using Transformer architecture, RAG, and LangChain — **45%** complete.
- Introduced asynchronous task queues (Celery + Redis) and containerised all services with Docker, significantly improving deployment reliability and developer iteration speed.

Python Developer (*Mar 2024 – Feb 2025*)

NeoPanda Tech — China – Shenzhen (Remote)

- Built and maintained RESTful APIs and backend services in Python; integrated AI/ML capabilities into production data workflows
- Built a robust web crawler and scraper suite with proxy pooling, retry logic, and browser fallback (Playwright), ingesting **5,000+ daily articles** from **50+ regional news sources** into a SQLite pipeline.
- Built a RAG-based news intelligence system with vector embeddings and an autonomous AI agent, achieving **99.9%** deduplication and **85%** reduction in manual newsletter curation time.

AI/ML Engineering Intern (*Jan 2024 – Feb 2024*)

IIT Roorkee — India

- Trained and evaluated supervised ML models (Scikit-Learn, Pandas, NLP) across three tasks: sentiment analysis, emotion classification, and small-business sales forecasting; achieved **>85%** classification accuracy on validation sets.
- Benchmarked Small Language Models (SLMs) against larger counterparts to guide practical model selection decisions for resource-constrained deployments.

Lead Product Engineer (*Apr 2023 – Mar 2024*)

Zume, Inc. — USA – California (Remote)

- Led AI application and ML feature development across an Agile/Scrum workflow; owned end-to-end delivery for **3+ product features** with cross-functional stakeholders.
- Designed OpenAI-powered voice analytics pipelines, LightGBM sales forecasting, and a RESTful backend for **Vocalytics-V0.3** — a voice-driven BI tool for sales teams.

TECHNICAL SKILLS

Languages: Python, JavaScript, TypeScript, Java, C#, SQL, Go, Bash, HTML5, CSS3

Frontend / Full-Stack: React.js, Next.js 14, Tailwind CSS, Vite, Turborepo, Radix UI, WebSockets, REST APIs, GraphQL, Responsive Design

Backend & APIs: FastAPI, Flask, Node.js, Express.js, Django, Socket.IO, Microservices Architecture, JWT Authentication, RBAC, Rate Limiting, Celery, Pydantic

AI / ML & Data Science: Large Language Models (LLM), Transformer Architecture, Model Fine-Tuning, Retrieval-Augmented Generation (RAG), LangChain, LangGraph, Prompt Engineering, OpenAI API, Google Gemini, PyTorch, TensorFlow, Scikit-Learn, Computer Vision, OCR, NLP, Pandas, NumPy, LightGBM

Vector Databases: pgvector, Qdrant, Weaviate, Pinecone, Milvus, Elasticsearch

Databases: PostgreSQL, MongoDB, MySQL, Redis, SQLite

DevOps / Cloud / MLOps: AWS (EC2, S3, ECS, CloudWatch), Docker, Kubernetes, Nginx, Git, Linux, CI/CD Pipelines, Agile, Scrum

PROJECTS

FBSPLAI — AI-Powered Insurance Sales Platform (2025)

- *Source code proprietary / NDA. Live at fbsplai.com.*
- Full-stack AI insurance sales platform: **3 Next.js 14 portals** (client, admin, AI chat) sharing a Turborepo monorepo, Radix UI component library, and Tailwind CSS design system.
- **10 FastAPI microservices**: document extraction (LangChain/LangGraph agents), Quoteflow engine, loss-run reporting, AI chat, document renaming, quote comparison, and an API gateway — orchestrated with Docker Compose, Celery, MongoDB, and Redis.
- Multi-provider LLM integration (OpenAI + Google Gemini + Cloud Vision OCR) for document intelligence, quote analysis, and conversational AI
- Deployed on AWS EC2 behind Nginx with JWT/RBAC, rate limiting, and async task queues

Vocalytics-V0.3 (2024)

- Voice-driven sales analytics platform: converts spoken commands into Plotly BI workflows and natural-language summaries using OpenAI GPT-4.
- Implemented secure AST-validated code execution, LightGBM forecasting, and CSV ingestion with encoding detection; packaged for Windows desktop with cx_Freeze.
- **Stack**: Python, Flask, OpenAI API, LightGBM, Pandas, Plotly, SpeechRecognition. **GitHub**: Vocalytics-V0.3

Intelli-News v1.1 — RAG News Intelligence Pipeline (2024)

- Crawled and ingested **5,000+ daily articles** from **50+ regional sources** with proxy pooling, retry logic, and Playwright browser fallback.
- Built a RAG-based vector embedding system with an autonomous AI agent achieving **99.9% deduplication** and **85% reduction** in manual curation effort.
- **Stack**: Python, LangChain, RAG, Vector Databases, SQLite, Scrapy, Playwright. **GitHub**: Newsletter-Intelli-v1.1

ACORD Forms Extraction Pipeline (2025)

- Layout-aware document extraction pipeline using computer vision, visual document models, and OCR for structured text, tables, checkboxes, and form fields from ACORD insurance documents
- Self-learning iteration loop with visual-model fine-tuning and version switching for safe rollback
- **Stack**: Python, Computer Vision, OCR, FastAPI, PostgreSQL, Pydantic

More projects: mizanchishty.ddns.net/#projects

CERTIFICATIONS

AI / Machine Learning: Hands-On AI: Building LLM-Powered Apps, Introduction to Artificial Intelligence, Deep Learning with TensorFlow, Deep Learning: Model Optimization & Tuning, Google Cloud Platform for ML, R for Data Science

Generative AI: Career Essentials in Generative AI (Microsoft & LinkedIn), What Is Generative AI?, Generative AI: Evolution of Online Search, Ethics in the Age of Generative AI

Backend & Web: Django Essential Training, Node.js Essential Training, Introduction to HTML, Introduction to CSS, TCP/IP and Advanced Topics

OOP & Java: Core Java Specialization (*LearnQuest*), Intro to Java, Intro to OOP with Java, OOP Hierarchies in Java, Java Class Library, Fundamentals of Data Warehousing

Game Dev & C#: C# for Unity Specialization (*Univ. of Colorado*), Intro to C# and Unity, More C# and Unity, Intermediate OOP for Unity, Intro to Game Design (*CalArts*)

Productivity & Tools: Streamlining Work with Microsoft Copilot, Learning Microsoft 365 Copilot for Work