

MATEEH ULLAH

AI SYSTEMS / FULL-STACK / CLOUD

AUSTRIA - WORLDWIDE	mateehullah89@gmail.com	+43 676 985 2028
	linkedin.com/in/mateeh-ullah	github.com/MateehUllah

Software engineer with 3+ years building production LLM workflows, scalable backends, and full-stack products - from benchmark-driven evaluation to dependable delivery.

3+	7%	20%	99.9%
YEARS ENGINEERING	HALLUCINATION REDUCTION	EFFICIENCY IMPROVEMENT	CONNECTOR UPTIME

01 EXPERIENCE

Software Engineer @ Footmetrics MAR 2025 - FEB 2026

Built data pipelines and APIs powering customer analytics, focused on latency, throughput, and operational efficiency.

- Built 12+ Django REST endpoints for customer analytics ingestion, processing, and delivery.
 - Reduced operational latency by 40% and server load by 30% through caching and query optimization.
- DJANGO / PYTHON / POSTGRESQL / REDIS / REST

LLM Engineer @ Turing SEP 2025 - FEB 2026

Built benchmark-driven workflows and evaluation pipelines for agentic AI systems in fund management.

- Reduced hallucinations by 7% while improving task success across fund-management workflows.
 - Designed investor-onboarding and audit scenarios that achieved 20% efficiency gains; risk assessment cut false positives by 10%.
- LLMS / SFT / RLHF / PYTHON / EVALUATION

AI Engineer @ Micro1 SEP 2025 - JAN 2026

Validated and stress-tested AI connector integrations across major SaaS platforms, automating QA into CI/CD.

- Validated Zoho CRM, Freshdesk, and Salesforce connectors at 99.9% uptime.
 - Improved test coverage by 30% with synthetic data and saved 10+ hours per week through CI automation.
- PYTHON / ZOH0 / SALESFORCE / FRESHDESK / CI/CD

Associate to Software Engineer @ Devsinc JUN 2023 - MAR 2025

Built internal tools, reusable component systems, and customer-engagement integrations supporting 1,000+ users.

- Improved Twilio response rates by 70% and reduced page load times by 40% across user-facing workflows.
 - Shipped 20+ reusable components, cut frontend bugs by 15%, and automated 100+ regression tests.
- REACT / VUE / TYPESCRIPT / DJANGO / TWILIO / DOCKER

02SELECTED WORK

- 01

SereniMind

25% fewer hallucinations

Empathetic AI companion with contextual memory, secure auth, and conversation history using Deepseek-R1 through Ollama.

REACT / FASTAPI / DEEPSEEK-R1 / OLLAMA
- 02

Enhancing Infrastructure & Accessibility

Lablab.ai finalist

Combined XGBoost, Random Forest, K-Means, and Mistral-7B recommendations over John Snow Labs and World Bank datasets.

XGBOOST / MISTRAL-7B / REACT / FLASK
- 03

pyERP

45% fewer critical errors

Migrated a monolith to Django microservices and ETL pipelines, reducing redundancy by 30% and improving the first delivery cycle.

DJANGO / VUE / MICROSERVICES / ETL
- 04

BotBlogR

70% faster generation

Re-architected Flask to FastAPI and React, lowering latency 40%; integrated the OpenAI API and tuned PostgreSQL for 30% more RPS.

FASTAPI / REACT / OPENAI / POSTGRESQL
- 05

IntelPort

99.9% auth uptime

Built a location-based analytics platform with Django, React, and PostgreSQL; REST API work reduced latency by 30%.

DJANGO / REACT / POSTGRESQL / REST

03CAPABILITIES

- AI / ML

LLMs, SFT, RLHF, Hugging Face, Scikit-learn, XGBoost, K-Means
- BACKEND

Python, Django, FastAPI, NestJS, REST APIs, Redis
- PRODUCT

React, Next.js, Vue, TypeScript, Tailwind CSS
- DATA

PostgreSQL, MongoDB, MySQL, SQLite
- CLOUD

AWS, GCP, Docker, Kubernetes, Lambda, S3, EC2
- QUALITY

Robot Framework, Browser Library, CI/CD, Git

04EDUCATION & RECOGNITION

- B.S. Computer Science - FAST-NUCES**

LAHORE / 2019 - 2023

Dean's List of Honors - Jun 2023

Finalist - AI for Connectivity Hackathon, Lablab.ai