TUNA KOMURCU

AI Engineer | LLM & Machine Learning Developer

Ankara, Turkey | +90 555 097 69 29 | tunakomurcu@gmail.com | github.com/TunaKomurcu | linkedin.com/in/tunakomurcu

EXPERIENCE

AI Engineering Intern / Developer — Mikrogrup *May 2026 – Aug 2026 (4 months)*

- Developed **MomTest AI** end-to-end to automate structured customer discovery workflows for business validation.
- Designed an agent loop using Function Calling and Structured Outputs, with strict guardrails to prevent leading-question bias.
- Integrated dynamic Markdown-based rule files for methodology enforcement; deployed and operated the production environment on AWS (EC2, S3, IAM, CloudWatch) and built unit/integration test suites.
- Owned model training and fine-tuning for **GİB CAPTCHA OCR**, raising recognition accuracy from **41% to 93%**.

AI Engineering Intern / Developer — Başarı Mobile *Sep 2025 – Jan 2026 (4 months)*

- Designed and built a database-agnostic **Text2SQL Agent** end-to-end, from architecture to prompt engineering.
- Built a Hybrid Retrieval (Dense + BM25) architecture with Qdrant (HNSW) for dynamic schema retrieval and selection.
- Achieved **~95%** accuracy on simple queries and **~85%** on medium-complexity queries.
- Load-tested with Locust and monitored performance and latency (p50/p95/p99) via Prometheus & Grafana.

AI Engineering Intern — Havelsan Technology Radar *Sep 2024 – Jan 2025 (4 months)*

- Designed and trained an **Edge-Attentive Residual CNN** from scratch, using attention mechanisms to denoise raw radar imagery.
- Implemented hybrid activation and loss functions to preserve high-frequency structural details.
- Applied ensemble learning across diverse model architectures and managed the end-to-end training pipeline, validating results against accuracy metrics.

KEY PROJECTS

Text2SQL Agent github.com/TunaKomurcu/Text2SQL-Agent

- **Architecture & Retrieval:** Architected a database-agnostic Text-to-SQL agent using Qdrant (HNSW) hybrid retrieval (Dense + BM25) to index schema metadata, table descriptions, and sample values for context-aware schema selection.
- **Inference & Execution:** Integrated automated SQL validation and repair loops against target databases, deploying quantized LLMs via llama.cpp (GGUF) and optimized embedding pipelines with ONNX Runtime.
- **Performance & Observability:** Achieved **~95%** accuracy on simple and **~85%** on medium-complexity queries, validating system stability under load with Locust and monitoring p50/p95/p99 latency via Prometheus and Grafana.

GİB CAPTCHA OCR *(private)*

- Fine-tuned a TrOCR transformer model for CAPTCHA recognition, raising accuracy from **41% to 93%** on 20,000 combined training samples.

MomTest AI github.com/TunaKomurcu/momtest-ai | momtest-demo.online

- **Agent Orchestration:** Engineered a two-stage customer discovery platform with an automated agent-loop engine, utilizing OpenAI Function Calling to enforce type-safe, structured outputs during automated interviews.
- **Quality & Validation:** Developed real-time AI guardrails and validation logic to detect vagueness, prevent leading questions, and eliminate conversation bias during live user interviews and post-analysis.
- **Cloud Deployment:** Deployed and operated the production environment on AWS (EC2, S3, IAM, CloudWatch, SNS, Secrets Manager, Route 53), implementing automated database backups, resource monitoring/alerting, and secretless credential management via IAM roles.

Edge-Attentive Residual CNN (Image Denoising) github.com/TunaKomurcu/Edge-Attention-Residual-CNN

- Attention-based CNN for image denoising using hybrid activation/loss functions and ensemble learning to maximize accuracy.

Spotify Music Features Analysis github.com/TunaKomurcu/Spotify-Music-Features-Analyzing

- EDA and feature engineering on Spotify data to uncover track popularity trends and data-driven insights.

Origami: Modular microUI Framework *| Team Project, 3 members* github.com/OrigamiUI/origami

- Lightweight, decoupled microUI library built from scratch, applying advanced design patterns for scalability.

SKILLS

Programming Languages: Python, Java, SQL

Machine Learning & Deep Learning: PyTorch, TensorFlow, scikit-learn, Fine-Tuning / Transfer Learning, LoRA / PEFT, Computer Vision, OpenCV, ONNX / Model Export & Optimization

NLP, LLMs & Vector Search: RAG, LangChain, Hugging Face, Prompt Engineering, LangGraph, OpenAI API Integration, Function Calling / Structured Output, Guardrails, LLM Evaluation, Multi-Agent Systems / Agent Orchestration, Model Quantization & Local Inference (GGUF, llama.cpp), Vector Databases (Qdrant, FAISS, Chroma), Sentence Transformers, HNSW / ANN Search, Hybrid Retrieval (Dense + BM25)

Data Analysis & Engineering: NumPy, Pandas, PostgreSQL, EDA, Feature Engineering, Data Visualization (Matplotlib, Seaborn), Jupyter

MLOps, Cloud & Deployment: AWS (EC2, EBS, S3, IAM, CloudWatch, SNS, Secrets Manager, Route 53), Docker, Podman, Nginx, Git, CI/CD, GitHub Actions, Unit Testing (pytest), FastAPI, Vercel, Latency & Throughput Benchmarking (Locust, Prometheus, Grafana)

EDUCATION

TOBB University of Economics and Technology, Ankara — B.S. in Artificial Intelligence Engineering *2021 – 2026*

- Focus: Machine Learning, Deep Learning, NLP, and Software Architecture

Graduation Project — Stack Underflow: AI Legal & Tender Analyst *| Team Project, 4 members (private)*

- Developed an 8-stage multi-agent LLM system with RAG pipelines to automate public tender analysis and antitrust legal workflows.
- Built a legal RAG module that semantically retrieves past Competition Board precedents matching each case's violation type and market structure.
- Orchestrated autonomous AI agents to analyze evidence and auto-generate structured investigation reports, defense rebuttals, and decision documents.