Enis Furkan Kırmızı

Istanbul, Türkiye | enisfurkankirmizi@gmail.com | +90 543 362 39 32
LinkedIn | GitHub

EDUCATION

Istanbul Technical University — Istanbul, Türkiye
BSc in Artificial Intelligence & Data Engineering — June 2026

- **GPA:** 3.33 / 4.00
- **Relevant Coursework:** Deep Learning, Natural Language Processing, Computer Vision, Data Mining, Analysis of Algorithms

TECHNICAL SKILLS

Languages: Python, SQL, C++

AI & ML: PyTorch, TensorFlow, scikit-learn

Specializations: Computer Vision, NLP, Knowledge Distillation, Model Compression

Data & Infrastructure: PySpark, Hadoop, Docker, REST APIs, Git

EXPERIENCE

Ziraat Teknoloji A.Ş. — *Istanbul, Türkiye*
Artificial Intelligence Intern — *August 2025 – September 2025*

- Built a multimodal Retrieval-Augmented Generation (RAG) pipeline using CLIP (ViT-B/32) to align text and image embeddings for context-aware responses.
- Developed an automated document ingestion pipeline using PyMuPDF and Tesseract OCR to extract, chunk, and embed PDF data into a ChromaDB vector database.
- Implemented a real-time face recognition system using OpenCV, MTCNN, and FaceNet (InceptionResnetV1) for video stream processing.

PROJECTS

Model Compression Architecture (PURSUhInT Extension) — *2025 – June 2026*
Senior Graduation Project

- Extended the PURSUhInT knowledge distillation framework with CP and Tucker tensor decompositions to train compressed student models.
- Built a distributed training pipeline using DistributedDataParallel (DDP) and Automatic Mixed Precision (AMP) to improve training efficiency across CIFAR and ImageNet.
- Achieved a 13.7x parameter reduction (21.3M to 1.5M) and 12.3x fewer FLOPs when compressing a ResNet-34 into a ResNet-18 with only a 3% drop in Top-1 accuracy.

Multimodal Idiomaticity Representation (AdMIRe 2.0)

- Designed a multimodal architecture combining vision-language encoders (SigLIP2) and BGE-M3 embeddings to model multilingual idiomatic expressions.
- Implemented dynamic idiom paraphrasing techniques via Small Language Models to mitigate literal visual bias in complex image-text ranking tasks.
- Built a modular pipeline to evaluate semantic alignment between text and images, streamlining the evaluation process for custom validation sets.

Big Data Analysis Pipeline

- Built a distributed data pipeline using PySpark and Hadoop on Google Cloud Dataproc to process a 40 GB Steam user review dataset.
- Performed exploratory data analysis by implementing efficient sampling techniques to extract insights on gamer profiling and engagement trends.
- Evaluated system performance by benchmarking single-node versus multi-node cluster configurations, demonstrating the efficacy of in-memory distributed processing.