

Alper Kaplan

Senior AI/ML Engineer/Data Scientist | 8+ Years of Experience | GenAI | Scalable ML Systems

📍 Luxembourg ✉ alperkaplan@outlook.com ☎ +352 691 177 633 🌐 alperkaplan.co in alperkaplan

Innovative and results-driven Machine Learning Engineer with almost a decade of experience designing and deploying scalable ML systems across NLP, LLMs, and Generative AI. Proven track record delivering production-grade solutions from data ingestion to real-time inference, including pipelines that reduced processing time by 90%+ and served millions of assets. Strong in software engineering, system design, and model optimization — with deep expertise in Python, PyTorch, AWS (SageMaker, Lambda, EFS), and MLOps tooling. Known for end-to-end ownership, bias for action, and translating complex ML challenges into measurable business outcomes.

Experience

Senior Machine Learning Engineer

Docler Holding Luxembourg, Byborg Enterprises

Kirchberg, Luxembourg (Hybrid)

Mar 2025 – Present

- Built a scalable LoRA training pipeline (kohya_ss) with prompt optimization and distributed GPU orchestration, improving multi-frame character consistency.
- Deployed 20+ ComfyUI-based LoRA models, cutting training time by 30% via modular preprocessing and minimal-step tuning.
- Designed and integrated a production-grade video generation API, automating 100+ theme-based character videos and accelerating storytelling workflows.
- Collaborated with product and design teams to define KPIs, optimize model success metrics, and ship features quickly through iterative cycles.

Senior Data Scientist

ATP Tech

Istanbul, Turkey (Hybrid)

Jun 2023 – Feb 2025

AI-Powered Resume Parsing & Ranking System:

- Architected and deployed a scalable ML pipeline for resume parsing and ranking using transformers, LLMs, and traditional ML; cut candidate screening time from 3 days to 2 hours per 100 applicants, boosting HR throughput.
- Led system design end-to-end, including tokenizer implementation, JSON output schemas, vector-based scoring models, and modular deployment architecture for integration across HR platforms.

Chatbot for Branch Manager Support (Burger King Turkey):

- Built and productionized a customer service chatbot to automate branch-level support; reduced manual support workload by 40%, handling hundreds of monthly inquiries.
- Delivered full-stack solution with Flask, MongoDB, ChromaDB, LangChain, and OpenAI on AWS; added CI/CD pipelines and latency-optimized retrieval.

Applied ML Research: Co-authored multiple applied ML papers on LLM evaluation and benchmarking; integrated research outputs into live systems, improving inference pipeline reliability and decision transparency.

Senior Machine Learning (Founding) Engineer

Ginoa

Amsterdam, Netherlands (Remote)

Sep 2022 – May 2023

NFT Price Prediction:

- Led end-to-end ML system design for NFT price forecasting, combining collection-level and global models to handle asset variability and noise.
- Achieved 90%+ T+1 prediction accuracy on millions of assets by optimizing batch inference and reducing latency for real-time pricing APIs.
- Mentored junior/mid-level engineers, driving experimentation, code reviews, and best practices across modeling and data workflows.
- Partnered with founders to translate model outputs into investor-facing KPIs, directly contributing to a \$500K seed round.

Research Engineer

Huawei

Istanbul, Turkey (Hybrid)

Apr 2020 – Aug 2022

Natural Language Understanding Projects:

- Architected and deployed production-scale NLP models for sentiment analysis, topic clustering, and multilingual query understanding (Arabic, English, Spanish, Russian), powering search intelligence for Huawei AppGallery at multi-million-user scale.
- Designed modular, zero-downtime APIs and lightweight pipelines to integrate NLU features into AppGallery search infra; collaborated cross-functionally with PMs and infra teams to ensure reliability and extensibility.
- Delivered measurable growth by boosting quarterly app downloads by 5–8%, driven by ML-powered personalization and fine-grained relevance optimization.

Research Assistant

Yeditepe University

Istanbul, Turkey (On-site)

Jan 2018 - Mar 2020

Preprocessed EEG data and built ML models for cognitive signal analysis; supported graduate-level cognitive science instruction.

Software Engineer

Freelance

Istanbul, Turkey (Remote)

Jan 2018 - Jan 2019

Built and published a real-time Android chat app (4.4/5.0, 1,000+ downloads) with IRC backend and custom emoji support.

Software Engineer

bidolubaski.com

Istanbul, Turkey (On-site)

Jun 2017 - Mar 2020

Developed an image matching system using deep learning, reducing matching time by 91%; co-authored a related peer-reviewed paper.

Jr. Software Engineer

OBSS

Istanbul, Turkey (On-site)

Jun 2013 - Jun 2014

Developed an educational Android app for children featuring animations, level-based progression, and SQLite integration.

Research & Publications

Enhancing Product Bundling with Large Language Models

Mar 2025

AT Bayrak, **A Kaplan**

[10.1109/ICICIP64458.2025.10898122](https://arxiv.org/abs/10.1109/ICICIP64458.2025.10898122) [↗](#)

Data Generation with ChatGPT: A Review on Sentiment Analysis Classification

Jun 2024

AT Bayrak, E Begen, SC Öner, **A Kaplan**, M Demirdağ, İU Sayan

[10.1109/HORA61326.2024.10550586](https://arxiv.org/abs/10.1109/HORA61326.2024.10550586) [↗](#)

Response performance evaluations of ChatGPT models on large language model frameworks

Jun 2024

A Kaplan, İU Sayan, H Şahan, E Begen, AT Bayrak

[10.1109/SIU61531.2024.10601055](https://arxiv.org/abs/10.1109/SIU61531.2024.10601055) [↗](#)

Image to Music: Cross-Modal Melody Generation through Image Captioning

Jun 2022

A Kaplan, D Goularas

Openaccess: [Pre-print](#) [↗](#)

A Hybrid Framework for Matching Printing Design Files to Product Photos

Apr 2020

A Kaplan, E Akagunduz

[10.17694/bajece.677326](https://arxiv.org/abs/10.17694/bajece.677326) [↗](#)

Skills

Core ML: LLMs · GenAI (Generative AI) · NLP · Chatbots · Transfer Learning · Fine-Tuning · Price Prediction · Learning to Rank · Transformers · LangChain · LlamaIndex · scikit-learn · PyTorch · TensorFlow/Keras · LoRA

MLOps: MLFlow · Weights & Biases · Model Serving · Real-Time & Batch Inference · Model Monitoring · Drift Detection · Model Retraining · Docker · Kubernetes · CI/CD

Data: SQL · NoSQL · Pandas · Polars · Numpy · Airflow · Spark · Hadoop · matplotlib · seaborn · GeoPandas

Cloud: AWS (SageMaker · Lambda · Bedrock · S3 · EFS · EBS) · Azure (ML Studio · Databricks)

Languages: English (fluent) | French (A2/B1) | German (A1) | Turkish (native)

Education

Yeditepe University

M.A. in Cognitive Science

June 2022

GPA: 3.71/4.0

Dissertation: [Image to Music: Cross-Modal Melody Generation through Image Captioning](#) [↗](#)

Yeditepe University

B.A. in Computer Science & Engineering

June 2016

Extracurricular Activities

Age group triathlete, powerlifter | Guitar, mandolin, and harmonica player | Former columnist on literature and music