

BERKANT KOCABAY

M.Sc. Business Informatics · AI / ML · Data Science

Freudenstadt, Germany · +49 170 4861783 · berkant.kocabay22@gmail.com

LinkedIn: [linkedin.com/in/berkant-kocabay](https://www.linkedin.com/in/berkant-kocabay) · GitHub: github.com/berkantky

Website: berkantkocabay.com

PROFILE

Master's student in Business Informatics (current GPA 1.1) with a strong interest in applying AI and data-driven systems to real-world-problems. Through projects ranging from driver emotion recognition and demand forecasting to LLM knowledge integration, I have developed experience across machine learning, software development, and system evaluation. What motivates me most is not AI itself, but understanding where technology can create measurable value and translating complex problems into practical solutions. I enjoy working at the intersection of technology, business, and human factors, combining technical implementation with a broader view of how systems are used in practice.

TECHNICAL SKILLS

Programming: Python (advanced), TypeScript / JavaScript, Java, SQL, HTML / CSS

Machine Learning & AI: TensorFlow, Keras, Scikit-learn, NumPy, Pandas; CNNs, time-series forecasting (ANN), clustering (K-Means), feature engineering, model evaluation; LLM integration — RAG, fine-tuning, prompt engineering

Backend & Frameworks: FastAPI, Node.js, Next.js, React, REST APIs, microservices

Databases: PostgreSQL, MySQL / MariaDB, MongoDB; SQL analytics, indexing, buffer management

Cloud & DevOps: Docker, Kubernetes, VMware, Linux (RHEL, Ubuntu)

BI & Analytics: Power BI, Grafana, RapidMiner, R Studio, SQL reporting

Enterprise Systems (SAP): Cloud ALM, Solution Manager, HANA-Studio, BPMN 2.0 (Order-to-Cash, Order-to-Pay)

Tools & Methods GitHub, Jira, Confluence, Figma, Agile / Scrum, scientific writing

SELECTED PROJECTS

RAG vs. Fine-Tuning on Financial Filings — Master's Thesis

2026

Designed and implemented an end-to-end evaluation framework to compare Retrieval-Augmented Generation (RAG) and LoRA fine-tuning for knowledge integration in local LLMs. Built a benchmark of 237 financial QA tasks from 60 SEC 10-K reports covering six major US technology companies (FY2016–2025) and evaluated seven system variants on accuracy, temporal knowledge adaptation, robustness, and confidence calibration. Fine-tuned Llama 3.1 models on an HPC cluster and automated the complete retrieval, inference, and evaluation pipeline.

Stack: Python · Llama 3.1 (8B/70B) · LoRA/QLoRA · FAISS · Sentence Transformers · Hugging Face · SLURM/HPC · SEC EDGAR API

RoboAdvisor — Year-Long Master Project

2024 – 2025

Scrum Master and Backend Developer for a portfolio optimization platform using real financial market data. Built a modular backend system for portfolio management, transaction processing, and automated rebalancing using FastAPI and PostgreSQL. Implemented core financial logic including allocation, cash handling, and periodic portfolio rebalancing, and integrated external market data APIs with rate-limiting and historical data ingestion. Coordinated a 4-person team in an agile development setup.

Stack: Python · FastAPI · PostgreSQL · Docker · TypeScript · Node.js · Jira · Confluence

Real-Time Emotion Detection in Vehicles via Deep Learning — Bachelor's Thesis

2024-2025

Explored how computer vision can support driver safety by detecting emotional states (anger, fear, sadness) under the real-time and resource constraints of in-vehicle hardware. The real question wasn't which model scores highest — it was which one can actually run in a car. Built a transfer learning pipeline across four pre-trained CNNs on FERPlus and applied model optimization techniques to make them embedded-ready. Achieved 92% faster inference and 89% smaller model size with negligible accuracy loss.

Stack: Python · TensorFlow/Keras · LiteRT (TF Lite) · Transfer Learning · Post-Training Quantization

Demand-Driven Supply Chain Simulator

2025

Built an interactive dashboard that simulates how demand swings hit a manufacturer's production, material consumption and costs. Modeled a fictional transmission company with 6 products across two production lines, where the tool schedules daily production and triggers material replenishment automatically based on buffer levels. Lets you run different demand scenarios and immediately see the impact on machine utilization, service level and procurement costs.

Stack: Python · Dash · Plotly · pandas

Demand Forecasting with Neural Networks

2024

Built MLP time-series models to forecast 12-month demand for products, categories, and reseller revenue on the AdventureWorks data warehouse. Engineered a full data-to-model pipeline (SQL extraction, preprocessing, training, forecasting) with revenue forecasts landing within 2% of actual values.

Stack: Python · TensorFlow/Keras · SQL Server

EDUCATION

M.Sc. Business Informatics — Hochschule Reutlingen

10/2024 – 09/2026 (expected)

Current GPA: 1.1 · Focus areas: AI / ML, Data Science, Business Intelligence, Cloud Computing, Demand-Driven SCM, Management Consulting

B.Sc. Business Informatics — Hochschule Reutlingen

08/2021 – 02/2025

GPA: 1.3 · Award for Top of Cohort (best student of the semester)

Exchange Semester — Marmara University, Istanbul

02/2024 – 06/2024

GPA: 1.0 · Enterprise Systems, Data Mining, Decision Support Systems, BI, Supply Chain

Abitur — Business Gymnasium Eduard-Spranger-Schule

08/2018 – 06/2021

GPA 1.6 · Award for Excellence in Computer Science · Award for Outstanding Achievement

PROFESSIONAL EXPERIENCE

SAP Basis Intern — IT Operations & Core Services

03/2023 – 08/2023

Arburg GmbH & Co. KG · Lossburg, Germany

- Set up a Red Hat Linux 8 VM in VMware and configured core services (Apache, NTP, rsyslog, NFS) plus a MariaDB container in Podman
- Installed and administered an SAP HANA database in HANA-Studio, including data imports and backup & recovery
- Built a central Grafana dashboard consolidating system and application KPIs across all ERP systems, with threshold-based alerting
- Evaluated SAP Solution Manager vs. SAP Cloud ALM across monitoring, BPMN process management, and testing
- Created Smartsheet dashboards for two IT teams to track progress, workload, and capacity

LANGUAGES

German (native) · **Turkish** (native) · **English** (fluent) · **Spanish** (B1)