

Kerem Delikoyun, Dr.-Ing.

AI Engineer

Contact

📍 Singapore

☎ +65 8732 3073

✉ kerem.delikoyun@gmail.com

🔗 [GitHub](https://github.com/krmidel)
github.com/krmidel

🔗 [LinkedIn](https://www.linkedin.com/in/keremdelikoyun)
linkedin.com/in/keremdelikoyun

🔗 [Personal Webpage](https://keremdelikoyun.com)
keremdelikoyun.com

🔗 [Company Webpage](https://www.tum-create.edu.sg/research/cellface)
https://www.tum-create.edu.sg/research/cellface

Awards

SMART IG 2.0 Commercialization Funding Grant of SGD 800,000

SMART & NRF, Singapore
Sep 2025

1st Place in TextArena Agent Hackathon

AITinkerers, Singapore
Mar 2025

4th Place in National Patent Competition

Turkish Patent Institute, Türkiye
Oct 2020

Certifications

AWS Certified AI Practitioner

Amazon Web Services
Oct 2024

AWS Certified Cloud Practitioner

Amazon Web Services
Oct 2024

LLMOps

DataTalksClub
Oct 2024

MLOps

DataTalksClub
Sep 2024

Patents

Quantitative Analysis of Immune Cell and Platelet Aggregate Activation Using Realtime and End-To-End Image Processing Pipeline for Digital Holography Microscopy (pending)

10202500263R - 2025 | Singapore

Real-Time And End-To-End Image Processing Framework for Digital Holographic Microscopy ([link](#))

10202401892U - 2024 | Singapore

Languages

English

Full professional proficiency

Turkish

Native

Summary

AI Engineer and Team Lead at TUMCREATE with 5+ years building production-grade AI systems on cloud-native platforms. Doctor of Engineering (Dr.-Ing.) from TU Munich in LLM-based multi-agent systems, agentic OS architectures, and AI-native development tools.

Leading an international team across Singapore and Germany to deliver a patented health AI platform and secured SGD 800,000 commercialization funding. Currently developing an Agentic OS, integrating fine-tuned LLMs, skills, plugins, and MCP libraries. Looking to apply this full-stack AI engineering expertise to build scalable AI systems for data-driven decision-making in high-stakes environments.

Experience

TUMCREATE | Research Fellow – Team Lead

Jan 2026 – Present | Singapore

- Leading a multidisciplinary team of engineers and researchers across Singapore and Germany, managing the end-to-end project lifecycle of a health AI platform: defining technical roadmaps, translating product requirements into detailed technical specifications and development plans, tracking deliverables, facilitating cross-functional collaboration using agile practices, and reporting progress to management to ensure on-time milestone delivery and stakeholder alignment.
- Architecting and developing a production-grade Agentic OS for automated knowledge discovery and data analysis on a full-stack platform with RAG on RDBMS (PostgreSQL and Vector DBs), short/long-term memory, and skills/plugins/MCP library; implemented unit, integration, and E2E tests and CI/CD pipelines to ensure code reliability and repeatable deployments across environments.
- Improving domain-specific medical reasoning by multi-GPU distributed post-training of Qwen 3.5 on clinical datasets and integrating the fine-tuned model into the Agentic OS to enable automated "AI doctor" workflows with monitoring and evaluation dashboards for model and system performance.
- Delivering a 12h/6-week hands-on workshop series on multi-agent systems, post-training, RAG, and evaluation frameworks to drive adoption of agentic AI at TUMCREATE ([link](#)).

TUMCREATE | Research Associate

Aug 2021 – Dec 2025 | Singapore

- Reduced manual research effort by >95% (~3 days → 15 min) with exceptional performance on grounded evidence findings: topic adherence (55.7%) and faithfulness (0.42), by building end-to-end LLM-based multi-agent system for automated data analysis and deep research with evaluation and monitoring frameworks (LLM-as-a-Judge using RAGAS and Opik), achieving inference at ~500 s and <\$1 per run on AWS ([link](#)).
- Engineered a health AI pipeline (CNN, GAT, ViT) with optimized and quantized neural network architectures on multi-GPU infrastructure using PyTorch, achieving >70× throughput improvement (2.5 h → 2 min) and >99% storage reduction per sample (30 GB → 15 MB) across 100,000+ data points (>15 TB); delivered real-time analysis at <10 ms/frame and ~\$0.01 cost/inference in production on-prem with observability ([link](#)).
- Secured commercialization funding award of SGD 800,000 by SMART & NRF for the patented health AI deep-tech solution, through customer discovery interviews, business plan, and articulating technical vision and investment pitching to stakeholders ([link](#)).

Izmir Institute of Technology | Research Assistant

Sep 2016 – Jul 2021 | Türkiye

- Enabled automated analysis of holographic microscopy imaging data by building a CNN-based cell recognition AI pipeline with end-to-end data processing workflows from raw sensor data through data curation, model training, selection and deployment ([link](#)).
- Implemented feature extraction and model evaluation workflows to support real-time cell viability recognition and cell typing, earning 4th place in the National Patent Competition for the image processing algorithms and backend architecture.

Skillset

Technical: Python • Flask • FastAPI • JavaScript/React | PyTorch • Scikit-learn • H2O AutoML • HuggingFace • Unsloth • vLLM • Ray | LangChain • LangGraph • AutoGen • MCP • RAGAS | PostgreSQL • Vector DB (Pinecone, Chroma, FAISS) • dbt • Spark • Kafka | MLflow • Opik/CometML • Prefect • Grafana | AWS (EC2, S3, Lambda, SageMaker, Bedrock) • Docker • Kubernetes • Terraform • CI/CD (GitHub Actions) | AI-native Dev Tools & Agentic OS (Claude Code, Cursor, OpenClaw)

Domain: Agentic AI • LLM-based Multi-Agent Systems • RAG • Embeddings & Prompt-based Systems • Deep Research • SFT/DPO/RLHF/PEFT • Multi GPU/Distributed Post-Training • Automated Knowledge Discovery • LLM Evaluation Frameworks • Data Extraction/Transformation/Integration • MLOps/LLMOps/AgentOps • Feature Engineering • Unit/Integration/E2E Testing • Experiment Tracking & Observability

Education

Technical University of Munich, Germany

Electrical and Computer Engineering

Technical University of Munich, Germany

Physics, graduated with the highest distinction

Gaziantep University, Türkiye

Engineer Physics, graduated as salutatorian

Doctor of Engineering (Dr.-Ing.)

2023 – 2025

Master of Science (M.Sc.)

2022 – 2023

Bachelor of Science (B.Sc.)

2009 – 2015