

Koray Ulsan

Machine Learning Engineer

✉ koray.ulsan@uni-tuebingen.de | 📍 Tübingen, Germany | 🌐 📧 in 🏠



EDUCATION

- M.Sc. Machine Learning** | University of Tübingen **GPA: 1.9 (Current)** — Tübingen, Germany Oct 2024 – Mar 2027
Focus: Generalist, GenAI, Synthetic Data, Diffusion Models
- B.Sc. Computer Science** | University of Tübingen **GPA: 1.8 / 1.0** — Tübingen, Germany Oct 2020 – Sep 2024
Final Thesis: Generating Professional Portraits from Amateur Photos via Personalization of Stable Diffusion
Received Deutschlandstipendium

PREVIOUS WORK

- Startup Engineer & Researcher (HiWi)** | University of Tübingen — Tübingen, Germany Nov 2024 – Jul 2025
 - Architected **full-stack GenAI platform** as the foundational engineer to productionize bachelor thesis.
 - Orchestrated k-shot inference on a **HPC Slurm cluster** (8xL40S, 16xRTX 3090) using Redis and WebSockets to manage 100k+ jobs with estimated wait time report, having a peak throughput of **6,000 generations per hour**.
 - Developed a generative augmentation framework using InstantID and FLUX.1 Kontext, while improving facial resemblance via FaceNet, leading to a **first-author publication** at a [CVPR 2025 Workshop](#).
 - Optimized end-to-end generation latency to 14s by identifying bottlenecks via E2E profiling and implementing model pre-loading with a **“warm-node” queue management** system to eliminate cold-start latency.
 - Sustained a **49% activation rate** (visit-to-generation) for 2,500+ users, supporting high-retention cohorts with average 40+ monthly generations and 14–22 minute sessions for power users.*Stack: Python (FastAPI, PyTorch, wandb), SQLite3, Redis, Slurm, ComfyUI, Vue.js, Apache2, Linux (Self-hosted); GitHub Actions*

SELECTED PROJECTS

- Delayed Auditory Feedback Online** | Real-Time DSP Web App ([Micro-SaaS](#)) Feb 2025 – Present
 - Identified an underserved niche in speech therapy, built a low-latency audio engine to compete with native apps and physical DAF hardware, scaling to 500+ monthly users and 100h+ of session time with 65% returning traffic.
 - Directed SEO strategy to reach a **2.4 weighted average position** for core keywords, capturing 85% of organic traffic from Top 3 results with a 75% CTR on primary search intent.*Stack: JavaScript (Web Audio API), HTML, CSS, Formspree; Google Search Console, Google Analytics*
- Spurious Feature Annotation Platform** | Distributed Dataset Auditing System Apr 2023 – Jul 2023
 - Architected a modular **HITL tool** to identify spurious features with multi-source API dataset enrichment, enabling researchers to label datasets via a centralized backend to offload model inference.
 - Steered a 4-person Agile team as **Technical Lead**, managing the end-to-end SDLC via Scrum.*Stack: Python (PyTorch, timm, torchvision, HuggingFace, gradio, NumPy), JavaScript, Conda; GitHub Actions/Projects*

PUBLICATIONS

Ulsan, K., & Kiefer, B. (2025). [Generating Synthetic Data via Augmentations for Improved Facial Resemblance in DreamBooth and InstantID](#). CVPR 2025 Workshop “Synthetic Data for Computer Vision”.

COURSEWORK

- Lectures:** Reinforcement Learning · Trustworthy ML · AI Safety · Benchmarks · Non-convex Optimization · NLP · LMs
Gaussian Processes · Recurrent & Generative NNs · CUDA · C++ · DBMS (PostgreSQL/DuckDB) · SQL
- Selected Coursework Projects:
- RL Paper: [Implementing TD3 with PER / Reward Shaping / Truncated n-Step Return Using SB3](#)
 - Paper: [Understanding Disagreement in Peer Review](#)
 - Code: Training a VAE on 64x64 images to achieve latent space vector disentangling
 - Presentation: [Synthetic Image Generation with a focus on GANs](#)

Feb 2025

Feb 2025

Dec 2024

Dec 2023

LANGUAGES: Turkish (Native) · English (Professional, IELTS 7.0) · German (Advanced, DSD B2/C1)

HOBBIES: Homelab · 3D Printing · Adobe Photoshop · Stained Glass Art · Karate (Brown Belt)