

Kaan Kuzu

kaankuzu122@outlook.com • +31 6 29 17 52 80 • [linkedin.com/in/kaankkuzu](https://www.linkedin.com/in/kaankkuzu) • github.com/kaankuzu1 • kaankuzu1.github.io

EDUCATION

Eindhoven University of Technology (TU/e)

Jul 2029 (expected)

B.Sc. Computer Science and Engineering | Eindhoven, Netherlands

Availability: part-time during the academic year (8–12 h/week) from Nov 2026; full-time from Jul 2027.

Lycée Saint-Benoît

2026

High school diploma, grade 87/100 | Istanbul, Turkey (French-Turkish bilingual program)

Honors and Awards: AP Calculus BC score 5/5; IELTS 7.0; DALF C1

Additional Coursework: Sabancı University Summer School (2024): Artificial Intelligence and Machine Learning, Introduction to Web Programming

TECHNICAL SKILLS

Programming Languages: Python, TypeScript/JavaScript, SQL

Data and ML: NumPy, PyTorch, XGBoost, librosa, Faster-Whisper, pyannote, Hugging Face Transformers, LLM APIs

Frameworks and Development Tools: FastAPI, Next.js/React, PostgreSQL, Docker, Git, ffmpeg, Modal, Vercel, Claude Code

Languages: Turkish (native), English (IELTS 7.0), French (DALF C1)

PROJECTS

AudienceLabs | *Multimodal audio/video analysis pipeline (Python, FastAPI, PyTorch, Modal)*

Apr 2026 - Present

Link: audiencelabs.net

- Built a pipeline that ingests recordings of up to 4 hours, extracts audio with ffmpeg, transcribes it with word-level timestamps (Faster-Whisper) and aligns speaker diarization (pyannote 3.1) and audio-energy/emotion signals (librosa, wav2vec2) into a per-second feature timeline.
- Designed a two-pass LLM analysis (segmentation, then event detection) followed by deterministic validation rules (minimum duration, truncation and mid-sentence checks, sentence-boundary snapping) and IoU-based deduplication, turning noisy model output into consistent results.
- Implemented a pluggable scoring interface with a heuristic baseline and an XGBoost regressor, plus a weekly retraining job with seeded 80/20 split and a versioned model registry with automatic fallback; released the scoring stage as an open-source package with tests and CI (github.com/kaankuzu1/clip-virality-scorer).
- Runs as fail-soft, idempotent GPU jobs on Modal with per-stage status in PostgreSQL and a Next.js dashboard; documented with ADRs. Sole developer, 120+ commits; open for beta requests.

Multi-agent content automation for a law practice | *TypeScript, Claude API, Trigger.dev, Supabase*

Sep 2026 - Present

- Building a daily workflow, now in final testing, in which six specialized LLM agents (research with legal-database tools, writing, design, captioning) prepare an informational post, with every step gated by the lawyer's approval over WhatsApp.
- Added output validation per step, event and cost logging in PostgreSQL, and a watchdog that checks every 5 minutes for stalled or timed-out runs and resumes them from the last completed step.

LEADERSHIP & PROFESSIONAL DEVELOPMENT

E.S.A.I.V. Serpentine | Member, TU/e AI association and student team

2026 - Present

GEWIS | Member, study association for Mathematics and Computer Science

2026 - Present

FIRST Robotics Competition (FRC Turkey), Team ROBIN | Software team member

2023 - 2025

"Move on Sui Mini Hackathon" Bootcamp, RiseIn | Participant, Sui blockchain ecosystem

2023

INTERESTS

Theater (Théâtre Choc; competitions in Ankara and Istanbul, performed Andromaque in Lille on an exchange), volunteering (Hatay earthquake relief), football, tennis (local league), shot put