Mădălin Tatarciuc

Applied AI Researcher · LLMs, speech and video models

madalintat@gmail.com | +40 754 461 974 | cetusian.com | Bucharest, Romania · open to remote
linkedin.com/in/madalin-tatarciuc | github.com/madalintat | huggingface.co/cetusian

PROFILE

Applied AI researcher who trains small, specialised models and takes them all the way to production: data, training, evaluation, a fast runtime and the product around it. I learn by rebuilding: take a paper, implement it without the reference code, and write each experiment down with its bar before it runs.

EXPERIENCE

Invergent · *Applied AI Researcher* · surogate.ai — Mar 2026 – Present

- **Speech recognition:** led jackrabbit-110m-ro, Romanian ASR at 5.69% WER on FLEURS: ahead of Whisper large-v3 at a fourteenth of its size, with a streaming variant for live calls.
- **Text-to-speech:** led amami-357m-ro, three Romanian voices in one model, running natively on CPU, CUDA and Metal.
- **Data and evaluation:** collected and cleaned the Romanian speech corpus, with one normaliser so every run stays comparable.
- **Training:** built the terminal dashboard to launch, monitor and compare training runs (including GRPO with an LLM judge) on cloud GPUs.
- **Agents:** built the parts of Surogates where an agent meets a customer: phone calls over SIP on my own speech models, Discord, Teams, WhatsApp and website chat, and a simulator where an AI plays the customer before an agent goes live.

Salestools · *Engineering Team Lead* — Jun 2025 – Nov 2025

- Led the team building an LLM-driven outbound sales platform on multi-step agents (research → personalise → write).
- Designed agent pipelines combining retrieval, reasoning and validated JSON output, with confidence scoring and failure routing.

Plant an App & Bazinga AI · *AI Automation Engineer* — Jul 2024 – Jun 2025

- Fine-tuned and deployed models for sales and support automation; integrated them into production APIs and internal tools.

.worlds · *Co-founder, CEO* — Jun 2024 – Jun 2025

- Built text-to-image and image-to-video pipelines that turned a song into a music video. buildspace S5.

Omdena · *Romania Chapter Lead · BD Partner* — Apr 2023 – Apr 2024

- Led open-source applied ML projects, including computer vision for environmental monitoring; brought Omdena into the Romanian market.

Earlier: Skillab · Gameloft · Coreto.io

OWN RESEARCH

ipporu · *Video understanding for kendo* · ipporu.com — Oct 2026 – Present

- **Models:** a chain of specialist video models on Modal GPUs: player tracking, a referee-flag reader, strike detection, shinai keypoints and a who-struck model on V-JEPA 2.
- **Method:** every experiment logged with its gate before it runs; when clean labels disagreed with my own result, I retracted it and rebuilt the method.
- **Product:** live, with the record of 139k kendoka, scouting, video analysis and live scoring on top of the models.

Open models and tools · huggingface.co/cetusian

- **Chess LLM:** SFT datasets from Lichess games and puzzles for training a chess-playing model.
- **Tools:** PS3 Eye camera driver for Apple silicon; vecdb, a disk-backed vector database in C; LlamaRanch for local models.

SKILLS

Research: LLM fine-tuning (SFT, LoRA), speech (ASR, TTS), video understanding, evaluation, quantisation (GGUF)
Engineering: Python, PyTorch, NeMo, transformers, Go, TypeScript, Rust, C; FastAPI, Postgres, Redis
Runtime: llama.cpp, CPU and streaming inference, Modal GPUs

EDUCATION AND MORE

Education: buildspace, AI Builder Cohort (2024)
Certifications: Deep Learning, Machine Learning, Computer Vision, Hugging Face Deep RL
Languages: Romanian, English · Kendo: Romanian national team, 1 Dan