Szymon **Chirowski**

Applied Data Science & AI Student · MLOps

Breda, Netherlands

+31 6 85 28 55 98 | szymon@szymonchirowski.com | szymonchirowski.com | szymonchirowski242621 | szymonchirowski

Summary

Applied Data Science & AI student specializing in MLOps and production ML — building systems that hold up outside the notebook, from gated training pipelines and model registries to CI/CD and containerized deployment across cloud and edge. Work that has shipped under real constraints: an adversarially-hardened RAG assistant for university administration, vision models running real-time on a Raspberry Pi at 10% of commercial cost, and reinforcement-learning control hitting sub-millimeter precision 21× faster than a human. First-author on a peer-reviewed paper probing the limits of text-only emotion classification. I optimize for the number that proves it works and the uptime that keeps it working.

Education

B.Sc. Applied Data Science & Artificial Intelligence — *Breda, Netherlands*

Breda University of Applied Sciences — *2024 - 2028 (expected)*

IT Technician (Technical Secondary School) — *Katowice, Poland*

Technikum Lotnicze — *2019 - 2024*

Work Experience

Breda University of Applied Sciences — *Breda, Netherlands*

Research Assistant — *Feb. 2026 - Present*

- Extended a completed arousal-inference research app to a new client need by integrating it with the GRASS GIS / Tangible Landscape stack, so stakeholders shape terrain with physical sand and read arousal results directly on the model instead of a screen.
- Built the multi-CRS integration layer linking live 3D-scan ingestion to the pipeline and parallelised viewshed extraction for a 2× speedup.

Breda University of Applied Sciences — *Breda, Netherlands*

AI Chatbot Developer · Student Assistant — *Apr. 2025 - Present*

- Architected and built (sole developer) a production RAG assistant (FastAPI, LangChain, Claude, ChromaDB) for university administration, covering all high-traffic Student Desk / UniBuddy topics — programmes, admissions, fees, housing, visas — now in final sign-off before launch.
- Engineered the surrounding system: Selenium + PDF ingestion of 100+ documents into a chunked vector store on a 24h refresh, an LLM-supervisor moderation layer, async PostgreSQL, and a 4-service Dockerised deployment serving responses at 1.6s P50 / 7s P99.
- Hardened the assistant against adversarial use (social-pressure erosion, false-authority injection, tone manipulation), reaching 81.6% functional and 90.0% adversarial pass rates across a 900-case, 26-tester QA round.

Selected Projects

PyTorch, FastAPI, Azure ML, MLflow, Docker, GitHub Actions

Real-Time Sign Language Recognition — End-to-End MLOps — *2026*

- Served a 24-class Dutch Sign Language (NGT) model (MediaPipe + EfficientNet-B0) behind a JWT-authenticated FastAPI API over async PostgreSQL at 85% accuracy, 0.64ms P50 / 4.53ms P99.
- Gated Azure ML training on accuracy/F1 so only passing models register and the API auto-pulls the latest on startup, shipping the pipeline with GitHub Actions CI/CD, multi-target Docker, and 90% test coverage.

PyTorch, SegFormer, Stable-Baselines3, Weights & Biases, Opentrons

Automated Root Phenotyping & Inoculation — *2026*

- Built a CV + RL system (NPEC) segmenting plant roots/shoots (F1 0.81 / 0.93) and driving a liquid-handling robot to inoculate them at 0.18mm accuracy in 1.2s/plant — beating a PID baseline (0.32mm, 3s) and running 21× faster per dish than a skilled scientist.
- Tracked every run in W&B on a baseline-first, leakage-aware workflow for reproducible model decisions.

TensorFlow, Keras, MobileNetV2, EfficientNetB0, Grad-CAM, Raspberry Pi

Edge AI — Manufacturing Defect Detection — *2025*

- Trained a transfer-learning QC vision model running real-time on a Raspberry Pi 4 with Grad-CAM, reaching 100% accuracy on a 150-image held-out test set across four part types at 10% the cost of commercial systems.

Skills

ML & Modeling	PyTorch, TensorFlow, scikit-learn, XGBoost; computer vision, NLP, reinforcement learning, geospatial
MLOps & Deployment	MLflow, W&B, model registries, gated training; Docker, GitHub Actions, GHCR, Azure ML; FastAPI
LLM Systems	Claude API, LangChain, RAG, adversarial hardening
Languages & Infra	Python, SQL, Bash, JS/TS; Linux, PostgreSQL, Git, Tailscale/WireGuard
Spoken	Polish (native), English (professional), Dutch (B1)

Publications

2026 **First Author,** When Text Is Not Enough: Limits of Text-Only Emotion Classification — *39th Bled eConference*

doi.org/10.18690/um.fov.4.2026.44