

Michał Maik

dr.mike.maik@gmail.com

+48 537-429-174

Kraków, Poland

[LinkedIn](#) · [GitHub](#)

PROFILE

Senior AI/ML Engineer with a Physics PhD and 7+ years of hands-on machine learning experience spanning LLM-powered agentic systems, NLP, computer vision, and quantitative finance. Builds AI products from scratch. Most recently led R&D for a real-time NPC AI system using RAG, LangGraph, and fine-tuning pipelines. Strong foundation in probability, statistics, and quantitative modelling. Conference speaker and patent holder.

TECHNICAL SKILLS

LLM & Agentic: LangGraph, LLM APIs (OpenAI, Anthropic, Groq, Gemini), RAG (vector stores, rerankers, agentic RAG), context engineering, fine-tuning (Unsloth, Tinker), synthetic data generation, ModernBERT validators, MCP servers, LangSmith

ML & Deep Learning: Python, SQL, PyTorch, TensorFlow/Keras, Scikit-learn, NumPy, SciPy, Pandas, Matplotlib, Seaborn

NLP: Text classification (BOW, TF-IDF), NER, sentiment analysis, LLM-as-judge evaluation, perplexity-based detection, chatbot and AI assistant development

Computer Vision: OpenCV, PIL, Scikit-image, real-time pose estimation (YOLO, MediaPipe), object detection

EXPERIENCE

Foregamer — Head of Research & Development

Jan 2025 – Mar 2026

Built an LLM-powered NPC system from scratch and took it to production. Led a team of 4 engineers; owned architecture, hiring, product roadmap, and delivery.

- Designed hierarchical RAG pipeline with semantic tagging, reranking, and Hybrid vector search for character knowledge retrieval across complex game lore
- Built LangGraph-based agentic state machines for multi-NPC orchestration: concurrent AI characters with independent goals, shared world state, and character-specific persistent memory
- Implemented ModernBERT validation layer for hallucination reduction and output consistency; engineered techniques to override “helpful assistant” defaults for in-character behaviour (deceptive, hostile, evasive NPCs)
- Created synthetic data pipeline with multi-layer deduplication (simhash, Jaccard, semantic) and model distillation for fine-tuning; designed evaluation framework for internal benchmarks
- Developed the AI system powering Immersey.ai. This is a character platform where users create, share, and converse with AI characters across genres and go on adventures with them

Ecential, LLC — Machine Learning Engineer (Freelance)

Jan 2025 – Present

- Building agentic competitive analysis platform from scratch: multi-step web search with source grounding, structured data extraction, and LLM-generated strategic reports
- Designed end-to-end pipeline from raw web data to formatted deliverables; integrated multi-provider LLM routing for cost and quality optimisation

Zerapy — Chief Research Officer

Aug 2020 – Jan 2025

- Architected computer-vision-based physical therapy assistant: real-time pose estimation, repetition counting, and biomechanical form feedback
- Designed personalised progress reporting system tracking patient improvement over time; contributed to provincial patent for biomechanical tracking
- Led R&D strategy and research direction across a 4.5-year engagement; interfaced with clinical stakeholders to translate medical requirements into technical specs

Oanda — Data Scientist → Quantitative Analyst

Jun 2020 – Jun 2022, Feb – Dec 2024

Two stints at a global FX brokerage, progressing from data science to quantitative analysis across risk and client analytics.

- Implemented ML-driven hedging strategies for institutional FX risk management, modelling exposure across multiple currency pairs to improve execution quality
- Built high-value client classification model using gradient boosting; designed customer segmentation pipelines for targeting, retention, and lifetime value estimation

EARLIER EXPERIENCE

- Nauto** — *Deep Learning Engineer* Apr 2019 – May 2020
Person and license plate detection/blurring for GDPR compliance; time-to-collision estimation for collision warning; CARLA simulation for scenario generation
- Avalanche** — *Computer Vision Engineer* Jan 2018 – Oct 2019
Automated background cleaning and retouching pipeline for e-commerce product images
- Fideltronik Poland** — *AI Specialist* Jul 2018 – Mar 2019
R&D in computer vision: TensorFlow object detection API, transfer learning, Keras model training, image processing pipelines
- inFakt** — *Data Scientist* Jul 2017 – Mar 2018
Predictive churn model, NLP-based transaction category classification for invoices
- Azimo** — *Technical Analyst* Jan 2015 – Jun 2017
Real-time cashflow diagrams integrating 20+ payment systems; transactional profitability modelling; EBICS banking integration; ML delivery time estimation

SELECTED PROJECTS & TALKS

- Conference Speaking** — “Are LLMs the Future of NPCs?” — Game Industry Conference (GIC), Poznań & Data Science Summit, Warsaw
- GOTCHA CAPTCHA** — Novel motion-perception CAPTCHA exploiting human temporal vision vs. AI frame-averaging; built as working proof-of-concept
<https://github.com/DrMikeMaik/GOTCHA>
- Polish Text Normaliser** — Open-source TTS preprocessing library: expands numbers, currencies, dates, and abbreviations into grammatically correct spoken Polish. Pure Python, 10 chained modules
<https://github.com/DrMikeMaik/polish-text-normalizer>
- Featured** — AI Gamechangers blog — profiled for work on LLM-powered NPCs and applied AI in gaming:
["Natural voice conversations are the next step in NPC evolution"](#)

EDUCATION

- Jagiellonian University, Kraków** — PhD in Physics 2014
Included 18-month research placement at ICFO, Barcelona
- UCCS, Colorado Springs** — MSc in Physics 2008

ACHIEVEMENTS & LANGUAGES

- Provincial Patent** (Biomechanical Tracking & Feedback System)
- MinFinTech** First Place, Technical Category
- MPD** International Science Scholarship

Polish (fluent) · English (fluent) · Spanish (beginner)