

Krzysztof WasiucioneK

About me: <https://kwasiucioneK.cytr.us>

Wrocław, Poland · +48 667 366 100 · kwasiek@gmail.com

www.github.com/kwasiucioneK

<https://www.linkedin.com/in/krzysztof-wasiucioneK-724209269>



SUMMARY

IT professional with 14 years of experience, primarily focused on maintaining backend and frontend systems. For the past few years, I have been actively developing and practically applying Machine Learning and AI skills: training Polish-language NLP models (NER), building RAG systems, and integrating language models (LLMs) into real-world business applications.

PROFESSIONAL EXPERIENCE

Krzysztof WasiucioneK AI Solutions (sole proprietorship)

Sep 2026 – present

Founder, AI & Python Developer

AI solutions for businesses and public institutions: document RAG systems, data extraction from documents (OCR + LLMs), on-premise language models, API integrations and process automation.

Neurosoft Sp. z o.o., Wrocław

Feb 2012 – present

IT Operations Specialist

- Maintain and monitor the Court Judgments Portal (<https://orzeczenia.ms.gov.pl>) for the Polish Ministry of Justice — ensuring stability, availability and issue diagnosis.
- Provide support and maintenance for the NeuroCar system (licence-plate & vehicle recognition, red-light enforcement) - <https://neurocar.pl>
- Built internal automation and monitoring tools in Python.
- Contributed to the training and deployment of a new version of the document anonymiser (NER-based) for the Court Judgments Portal.
- Built internal proof-of-concept RAG systems for legal and data-protection document search (UODO RAG, LexSearch), now under further in-house development: <https://lexsearch.cytr.us/>
- Built internal automation tools in Python, including a Django application for testing new anonymiser versions (for the Court Judgments Portal and UODO): <https://github.com/kwasiucioneK/anonymizer>

EDUCATION

Higher School of Vocational Education (WSKZ)

2023 – 2024

Postgraduate studies: Back-end Python Development

University of Wrocław

2004 – 2009

Information Science & Library Studies (specialisation: Digital Libraries)

COURSES & CERTIFICATIONS

NVIDIA-Certified Associate: Generative AI LLMs

Sep 2026

Certificate: https://www.credly.com/badges/aac6a87f-1126-4b85-8622-1abf9d3795e0/linked_in?t=tlgg23

AI_devs 4 Builders

Mar – May 2026

Building AI agents: Context Engineering, Observability, Evals.

Certificate: <https://credsverse.com/credentials/5974d298-4be8-4999-9148-c23c241f227a>

AI_devs 3 Agents Reloaded

2025

Building AI agents: LLM, RAG, MCP.

Certificate: <https://credsverse.com/credentials/2cb780a4-23d5-4c2f-b122-4bea2193be42>

KEY SKILLS

Programming: Python (FastAPI, Django, Flask, asyncio, Pydantic, Celery, Redis), REST API design, JSON/XML, WebSocket, Server-Sent Events (SSE)

AI / Machine Learning: PyTorch, Transformers (Hugging Face), sentence-transformers, XGBoost, ONNX, Unsloth; model training & fine-tuning; generative AI, agentic workflows, function calling; LLM APIs and local models (Azure OpenAI, OpenAI API, Ollama)

NLP / RAG: Agentic RAG, hybrid search (BM25 + vectors), reranking, structure-aware chunking, time-aware retrieval (document versioning), RAG evaluation; NER (BIO tagging); OCR (Tesseract); Polish language models (Bielik, HerBERT, polish-RoBERTa)

Data & Databases: PostgreSQL/PostGIS, pgvector, SQLite (FTS5), Qdrant, OpenSearch, Azure AI Search, MongoDB, Redis, Solr; ETL & data pipelines, data validation and deduplication

Cloud & Notebooks: Microsoft Azure — basic (Azure OpenAI, AI Search, AKS, Key Vault); GCP and AWS — working exposure (free tier); JupyterLab, Google Colab (extensive — GPU model training)

Infrastructure & Ops: Docker, Kubernetes & Helm (basic, AKS), infrastructure as code (Bicep, basic), Linux (Debian/Ubuntu, SLES), Git, CI/CD (GitHub Actions, GitLab CI), systemd, nginx, Ansible; production monitoring (Icinga, Nagios)

Collaboration & tooling: Jira, Confluence, Redmine, GitLab, BitBucket

Data tooling: NumPy, pandas, xarray, conda, uv

PERSONAL AI / ML PROJECTS

- **Skryba (document processing & RAG)** — Self-hosted Django archive automating document processing with OCR and LLM-based metadata extraction, plus RAG search over Polish documents: <https://github.com/kwasiucione/skryba>
- **LeasePilot** — RAG assistant for leasing documents on Azure (FastAPI, Azure OpenAI, AI Search, AKS, Bicep): cited answers from the document version valid on the contract date, refusal without grounding, CI/CD with an evaluation gate: <http://leasepilot.57.174.151.186.nip.io/?t=a6e83c9ede6bc58c040d6afb432e9f56>
- **TaxPilot (RAG platform)** — Agentic RAG over tax law (acts, official guidance, rulings). Django + Celery + Redis + OpenSearch (hybrid BM25 + kNN), SSE streaming, and a self-updating ETL pipeline from the ELI API detecting legal amendments: <https://github.com/kwasiucione/taxpilot>
- **Sentinel Land Cover Analyzer (deep learning / computer vision)** — Land-cover classification from Sentinel-2 satellite imagery using a SegFormer transformer model; PostGIS; comparison of deep-learning vs LLM-vision approaches.
- **TenderBot (data integration & ETL)** — Automated monitor integrating two public procurement sources (BZP, TED); scheduled ingestion, normalisation, deduplication and alerting: <https://tenderbot.cytr.us/>

LANGUAGES

Polish - native · English - B1 (intermediate; technical reading & writing) · Russian - basic · German - basic

ADDITIONAL

Driving licence (category B) · Registered disability with a mild degree

I consent to the processing of my personal data for the purposes of this recruitment process in accordance with Regulation (EU) 2016/679 (GDPR).