

MIKOLAJ JANUSZ

contact@janumiko.com | github.com/janumiko | linkedin.com/in/janumiko

TECHNICAL PROFILE

Software Engineer focused on building reliable AI systems, including LLM-based agents. Experienced in developing high-scale distributed systems at frontier technology companies. Supported by an academic research background in Machine Learning, co-authoring publications on Explainable AI (XAI) and developing new methods with PyTorch.

EXPERIENCE

Snowflake Warsaw, Poland
Software Engineer Feb. 2026 – current

- Building the LLM-based orchestration layer behind Snowflake Intelligence, Cortex Agents and other AI offerings.
- Improving agent's visualisation generation of the agent through evaluation, observability and harness engineering.

Google Warsaw, Poland
Software Engineer Intern Jun. 2025 – Oct. 2025

- Shipped two projects ahead of schedule for the Sovereign Cloud team; received a full-time return offer.
- Built a Go testing framework for containerized (gRPC) services that was quickly adopted by 5+ teams.
- Designed a release validation system (Go, GCP) which became a core part of the team's CI/CD pipeline.

Jagiellonian University Krakow, Poland
Student Researcher Sep. 2024 – Jun. 2025

- Led student-grant funded research on neural network pruning (optimization), with findings accepted at ECAI 2025.
- Developed and published novel explainable AI methods for computer vision models (Vision Transformers, CNN).
- Built a research codebase (PyTorch) for model training and experiments on Poland's largest HPC cluster PLGrid.

Quadrature Capital London, United Kingdom
Platform Engineer Intern Jul. 2024 – Sep. 2024

- Enhanced monitoring capabilities of a Kubernetes-based platform supporting the automated trading system.
- Delivered multiple developer productivity tools under 10 weeks, achieving adoption across engineering teams.
- Coordinated development across two simultaneous, independent projects on the platform components.

Amazon Gdansk, Poland
Software Development Engineer Intern Jul. 2023 – Sep. 2023

- Automated a data ingestion workflow using Python and AWS, saving 10+ hours of manual effort weekly.
- Optimized execution time of a (Python) data processing pipeline for forecasting models by over 90%.

Meta London, United Kingdom
Production Engineer Intern Jul. 2022 – Sep. 2022

- Implemented a performance testing service for Scribe, a critical distributed log server that handles TB/s of data.
- Expanded CI/CD pipeline with a stage to stress-test Scribe against hardware failure and heavy load scenarios.

SKILLS

Programming Languages: Python, Go, C++, SQL, x86 assembly (basic)

Frameworks & Libraries: PyTorch, Transformers, PydanticAI, Hydra, wandb

Infrastructure & Platform: Kubernetes, Prometheus, ArgoCD, Linux, Docker, Bazel, Git, Slurm

Technical Focus: Machine Learning, AI Agents, Distributed Systems, Large Reasoning Models

EDUCATION

Jagiellonian University Krakow, Poland
Master of Science in Computer Science; Graduated With Distinction Oct. 2023 – Oct. 2025
Bachelor of Science in Computer Science Oct. 2020 – Jul. 2023