

Adam Pardyl

ML Researcher & Engineer — Computer Vision, Vision-Language Models, Large-Scale Training

✉ adam@pardyl.com · 🌐 adam.pardyl.com · 📍 Warsaw / Kraków, Poland

🔗 Google Scholar · 🐙 github.com/apardyl · 🔗 linkedin.com/in/apardyl

SUMMARY

Machine learning researcher and engineer building training systems, evaluation infrastructure and production software. Experience spans distributed VLM training, embodied AI benchmarks and Google engineering teams. PhD thesis on efficient visual exploration and search for embodied agents submitted in 2026.

TECHNICAL SKILLS

Research areas: Vision-language models, reinforcement learning, self-supervised learning, active visual exploration, embodied AI and interpretable computer vision.

Programming: Python, C++, Java, JavaScript, SQL.

ML systems: PyTorch, PyTorch Lightning, Hugging Face Transformers/TRL, FSDP/DDP, DeepSpeed ZeRO-2, Accelerate and vLLM for distributed training and efficient model inference.

Tools & infrastructure: Slurm HPC, Triton training kernels, Linux administration (Debian/Arch), Docker, FastAPI, Weights & Biases, Git, OpenCV and Unreal Engine 5.

EXPERIENCE

Jagiellonian Center for Artificial Intelligence Kraków, Poland
Research staff member — Jagiellonian University Oct. 2026 – present

Max Planck Institute for Informatics Saarbrücken, Germany
Visiting Researcher — Computer Vision Group Jun. 2026 – Sep. 2026

- Studied self-supervised concept discovery with Prof. Dr. Bernt Schiele. Extended DINOv2 with an attention-based concept pathway and VAE quantisation to make model predictions interpretable through visual concepts learned without concept labels.

IDEAS NCBR (now IDEAS Research Institute) Warsaw, Poland
ML Researcher & PhD Candidate — joint doctoral programme with Jagiellonian University Sep. 2022 – Sep. 2026

- Built the GRPO fine-tuning and evaluation stack for VLMs, owning RL environments, training code and distributed runs of up to **11B parameters on 32 H100 GPUs** (typically 4–8 GPUs), using Transformers/TRL, DeepSpeed ZeRO-2, vLLM, Accelerate and Slurm for workload scheduling on multi-node GPU clusters.
- Fine-tuned Qwen3.5 9B exceeded the zero-shot GPT-4o and Gemini 2.5 Flash baselines on the FlySearch FS-1 evaluation.
- Designed FlySearch's Unreal Engine 5 environments and evaluation infrastructure with collaborators: three difficulty tiers, nine VLMs and repeatable forest/city scenarios. Supervised an MSc student on the evaluation harness.
- Contributed merged changes to Hugging Face TRL. Mentored two MSc students, both subsequently admitted to PhD programmes.

AGH University of Krakow Kraków, Poland
Machine Learning Engineer Jul. 2021 – Dec. 2022

- Designed and trained the CT malignancy classifier for CompLung in the X-rAI project. The pipeline achieved AUC 0.90 versus 0.83 for DeepLung on LIDC-IDRI in the paper's setting with malignancy scores > 3.
- Developed ML components later integrated into the RayGenic radiology platform as part of the National Centre for Research and Development X-rAI project; responsibility focused on the ML components.

Google Zürich (remote), Switzerland
Software Engineering Intern — YouTube Content ID Jul. 2020 – Sep. 2020

- Analysed repetitive video fragments and designed an ML detector that reduced false-positive Content ID early matches by **over 82% in offline evaluation**, on an experimental instance indexing 25% of the most-watched YouTube videos.

Google Zürich, Switzerland
Software Engineering Intern — Android Frameworks Jul. 2019 – Sep. 2019

- Built on-device binary logging for Android Window Manager using compile-time Java code generation; logging calls are stripped from production builds. Released in AOSP; debug builds retain verbose analytics in a compact binary format.
- Built the Winscope ADB proxy for browser-based debugging of Android window transitions.

Google London, UK
Software Engineering STEP Intern — Android Enterprise Jul. 2017 – Sep. 2017

- Launched a security feature for the G Suite Mobile Management platform letting enterprise administrators monitor and mitigate potentially harmful apps on corporate devices.

SELECTED PUBLICATIONS

FlySearch: Exploring how vision-language models explore. **Pardyl, A.**; Matuszek, D.; Przebieracz, M.; et al. **NeurIPS** 2025.
Introduced a benchmark for goal-directed VLM exploration in photorealistic environments. On FS-2 City, humans reached 60.8% success versus at most 6.5% for evaluated VLMs. Project: flysearch.gmum.net.

AdaGlimpse: Active Visual Exploration with Arbitrary Glimpse Position and Scale. **Pardyl, A.**; Wronka, M.; et al. **ECCV** 2024.

Learned where to look and at what scale using reinforcement learning; matched STAM accuracy with over 40% fewer pixels in the ImageNet early-stopping setting. At equal pixel budgets, accuracy was 77.54% versus 76.13% for STAM.

Beyond Grids: Exploring Elastic Input Sampling for Vision Transformers. **Pardyl, A.**; et al. **WACV**, 2025.

Studied flexible patch sampling for vision transformers under limited observation budgets, with changes to model architecture and training to support non-grid inputs.

Active Visual Exploration Based on Attention-Map Entropy. **Pardyl, A.**; et al. **IJCAI**, 2023.

Used attention-map entropy to select informative observations without a separate learned policy; evaluated against active exploration baselines on reconstruction quality.

Full publication record: [Google Scholar](https://scholar.google.com/citations?user=...).

EDUCATION

Jagiellonian University — GMUM

Kraków, Poland

PhD candidate in Computer Science

Oct. 2022 – present

- Thesis submitted in 2026: “*Learning to look and search: Efficient visual exploration for embodied agents*”. Advisor: Prof. Bartosz Zieliński. GMUM Group of Machine Learning Research; joint doctoral programme with IDEAS NCBR. GPA 5.0/5.

Jagiellonian University

Kraków, Poland

MSc in Computer Science, graduated with honours

Oct. 2020 – Jul. 2022

- Specialisation: Machine Learning; GPA 4.9/5.

Jagiellonian University

Kraków, Poland

BSc in Computer Science — IT Analyst

Oct. 2016 – Jul. 2019

- Specialisation: Theoretical Computer Science; GPA 4.0/5.

RECOGNITION & SERVICE

Awards: FNP START scholarship from the Foundation for Polish Science (2026), one of 100 recipients nationally and five in computer science; finalist, Polish Olympiad in Informatics (2016).

Research funding: Principal investigator, NCN Preludium; co-investigator on national and EU projects, including ELIAS.

Patent application: EP24461637.1 (pending): A computer implemented method for identifying a microorganism in blood and a data processing system therefor (2024).

Leadership: Co-managed a \$500,000+ annual budget as Vice Chairman of the Jagiellonian University Student Council (2021–2022); streamlined financing, increasing funded student projects by 29%. Previously Vice President, JU Computer Science Students’ Association (2018–2020).

Teaching: Laboratory instruction in Advanced Python Programming and Computer Vision & Pattern Recognition.

Community: Peer review for NeurIPS, IJCV, WACV and KDD; helped organise the Machine Learning Summer School in Kraków and the ELLIS Doctoral Symposium on Robust AI.

Languages: Polish (native); English (full professional proficiency).