

Adam Pardyl

AI Researcher — Machine Learning, Vision-Language Models, Embodied AI

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

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

SUMMARY

AI researcher working on visual perception, vision-language models and embodied agents. PhD thesis submitted in 2026. First author of five papers at NeurIPS, ECCV, WACV, IJCAI and ECAI; principal investigator of an NCN Preludium project on uncertainty-guided active visual exploration.

RESEARCH EXPERIENCE

Jagiellonian Center for Artificial Intelligence

Research staff member — Jagiellonian University

Kraków, Poland

Oct. 2026 – present

Max Planck Institute for Informatics

Visiting Researcher — Computer Vision Group

Saarbrücken, Germany

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, using vision-only self-supervised training on ImageNet.
- Preliminary experiments retained baseline classification accuracy while introducing an explainable concept bottleneck for visual-prototype-based explanations on ImageNet.

IDEAS NCBR (now IDEAS Research Institute)

PhD Candidate & Researcher — joint doctoral programme with Jagiellonian University

Warsaw, Poland

Sep. 2022 – Sep. 2026

- Led visual exploration research in Prof. Tomasz Trzciński's Computer Vision group. Designed FlySearch's Unreal Engine 5 forest and city environments and evaluation infrastructure: three difficulty tiers and nine VLMs, with analysis of perception, scene understanding and planning failures.
- Built the GRPO fine-tuning and evaluation stack for VLMs, owning RL environments, training code and distributed experiments with models up to 11B parameters on 32 H100 GPUs (typically 4–8 GPUs), using Transformers/TRL, DeepSpeed ZeRO-2, vLLM, Accelerate and Slurm.
- Fine-tuned Qwen3.5 9B exceeded the zero-shot GPT-4o and Gemini 2.5 Flash baselines in the FlySearch FS-1 evaluation.
- Mentored two MSc students, including supervision of the FlySearch evaluation harness; both subsequently entered PhD programmes. Contributed merged changes to Hugging Face TRL.

Jagiellonian University — Collegium Medicum

Researcher

Kraków, Poland

Nov. 2024 – Mar. 2026

- Developed computer vision methods for pathogen identification with clinical collaborators: ROC AUC 0.97 for bacteria and 0.88 for fungi, evaluated on 16,637 Gram-stained images covering 14 bacterial species and three fungi.
- Combined image segmentation and multiple-instance learning to support microscopy-based identification. Delivered a proof-of-concept deployment evaluated by clinicians at a Kraków university hospital.

AGH University of Krakow

Machine Learning Engineer

Kraków, Poland

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.

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 achieved 60.8% success versus at most 6.5% for the evaluated models. The benchmark exposes failures in perception, scene understanding and action planning. 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. Project resources and implementation: io.pardyl.com/AdaGlimpse.

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

Defined input elasticity and evaluated how vision transformers handle patches with flexible positions, scales and coverage. Studied architectural and training changes that improve robustness to sparse, non-grid observations.

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

Introduced attention-map entropy as a signal for choosing informative observations, without an additional learned glimpse policy. Evaluated reconstruction quality under limited observation budgets. Project: io.pardyl.com/AME.

CompLung: Comprehensive Computer-Aided Diagnosis of Lung Cancer. **Pardyl, A.**; Rymarczyk, D.; et al. **ECAI** 2023.

Connected lung segmentation, nodule analysis and patient-level prediction in an integrated CT pipeline; released radiologist-annotated segmentation masks.

SpaRRTa: A Synthetic Benchmark for Evaluating Spatial Intelligence in Visual Foundation Models. Kargın, T.; Jasiński, W.; **Pardyl, A.**; et al. **IJCV**, accepted, 2026.

AI-Driven Rapid Identification of Bacterial and Fungal Pathogens in Blood Smears of Septic Patients. Sroka-Oleksiak, A.; **Pardyl, A.**; et al. **Computers in Biology and Medicine**, 2025.

ProtoMIL: Multiple Instance Learning with Prototypical Parts for Whole-Slide Image Classification. Rymarczyk, D.; **Pardyl, A.**; et al. **ECML PKDD**, 2022.

Full publication record: [Google Scholar](#).

EDUCATION

Jagiellonian University — GMUM

PhD candidate in Computer Science

Kraków, Poland

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

MSc in Computer Science, graduated with honours

Kraków, Poland

Oct. 2020 – Jul. 2022

- Specialisation: Machine Learning; GPA 4.9/5. Thesis: “*ProtoMIL: Multiple Instance Learning with Prototypical Parts for Fine-Grained Interpretability*”. Research focused on interpretable multiple-instance learning.

Jagiellonian University

BSc in Computer Science — IT Analyst

Kraków, Poland

Oct. 2016 – Jul. 2019

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

INDUSTRY EXPERIENCE

Google

Zürich (remote), Switzerland

Software Engineering Intern — YouTube Content ID

Jul. 2020 – Sep. 2020

- Designed an ML-based detector that reduced false-positive Content ID early matches by over **82%** in offline evaluation on an experimental instance indexing 25% of most-watched YouTube videos; analysed YouTube-scale matching data.

Google

Zürich, Switzerland

Software Engineering Intern — Android Frameworks

Jul. 2019 – Sep. 2019

- Implemented a high-performance on-device binary logging solution for the Android Window Manager using compile-time Java code generation that strips logging calls from production builds (now in AOSP).

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.

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.

Spoken languages: Polish (native); English (full professional proficiency).

AWARDS, GRANTS & PATENTS

Awards: **BNP START 2026 scholarship** — one of 100 young scientists selected nationally and one of only 5 in computer science. Finalist, Polish Olympiad in Informatics (2016).

Principal Investigator: National Science Center *Preludium* — “Where to look next: guiding active visual exploration with internal model uncertainty”, 2024–2025 (no. 2023/49/N/ST6/02465).

Co-investigator: European Commission *Horizon Europe* — European Lighthouse of AI for Sustainability (ELIAS), 2024–2027, PI Massimiliano Pontil. National Science Center *Sonata Bis* — sustainable computer vision for autonomous machines, 2024–2027, PI Bartosz Zieliński. Foundation for Polish Science *Proof of Concept* — AI for rapid microbiological diagnosis of sepsis, 2024–2025, PI Monika Brzywczy-Włoch. Foundation for Polish Science *Team-Net* — bio-inspired artificial neural networks, 2022–2023, PI Jacek Tabor.

Patent: (pending) A computer implemented method for identifying a microorganism in blood and a data processing system therefor. European Patent Office, 2024, application no. EP24461637.1.

SERVICE & ACTIVITIES

Reviewer: NeurIPS 2026; IJCV 2026; CVPR (CVDD workshop) 2025; WACV 2025; KDD 2024; WACV 2024.

Teaching: Laboratory instructor for *Advanced Python Programming* and *Computer Vision & Pattern Recognition*, as part of doctoral teaching duties at Jagiellonian University.

Leadership: Vice Chairman of the Board, Jagiellonian University Student Council (2021–2022) — co-managed an organisation with a \$500,000+ annual budget and streamlined project financing (+29% financed student projects). Earlier Vice President, JU Computer Science Students’ Association (2018–2020).

Research community: Helped organise the Machine Learning Summer School in Kraków and the ELLIS Doctoral Symposium on Robust AI.