Dmitry Lvov

d.lvov@innopolis.ru | lvovdmitro.github.io | GitHub

Research Profile

I study how earlier exams affect generated reports and how models learn from predicted gaze.

Education

ITMO University | MSc with honours in Applied Mathematics and Computer Science 2024–2026
Big Data and Machine Learning. **ITMO Best Graduate 2026.**
Kazan Federal University | BSc in Applied Mathematics 2020–2024

Experience

Innopolis University | Researcher, AI in Medicine Lab October 2023 – present

- Design studies with radiologists and test which X-ray findings the models miss or carry over.
- Teach students to design experiments and interpret results. Review their code and manuscripts.
- Wrote the EHR model selection paper and designed the experiments on prenatal referral workload.

ITMO University | Part-time engineer, Mind Lab December 2024 – present

- Built a pipeline that detects page elements, links figures to captions and describes figures with a vision-language model. The work led to my MSc thesis and an IEEE Access paper.
- Compared full pages and selected regions to test the effect of localisation on answer quality.

Selected Publications

9 papers listed

Main-track conference papers (2)

- **NeurIPS 2026.** Source-Causal Control of Historical Context in Longitudinal Radiology Report Generation. Developed SCFR. Led an audit of 18,716 pairs, varying history while keeping each current X-ray fixed.
- **NeurIPS 2025.** From Human Attention to Diagnosis: Semantic Patch-Level Integration of Vision-Language Models in Medical Imaging.
  Tested whether gaze predicted from patch-level features helped train a chest X-ray classifier.

Journal articles (4)

- **Journal of Pathology Informatics, 2026 (100763).** From whole-slide prediction to compact visual evidence: Spatial and functional assessment of attention-selected patches.
  Measured confidence with patches kept or removed, using equal-size random sets from the same slide.
- Multimodal document segmentation for technical QA. IEEE Access 14 (2026), 12733–12743.
- Uncertainty in cardiac landmark detection and diagnosis. IEEE Access 13 (2025), 87229–87240.
- Automated cardiometry on chest X-rays. IEEE Access 13 (2025), 10702–10712.

Workshop papers (3)

- Similar Scores, Different High-Risk Episodes. NeurIPS 2026, Can We Trust AI Evaluation?
- Same Top Fraction, Different Workload. NeurIPS 2026, Can We Trust AI Evaluation?
- LogitGaze: Predicting Human Attention. ICLR 2025, Reasoning and Planning for LLMs.

Technical Skills and Qualifications

Academic activity. Main-track reviewer, NeurIPS 2026. Recorded presentations for NeurIPS 2025 and 2026.
Tools. Python, PyTorch, Transformers, scikit-learn, Docker, ClearML, Git and Linux.
English. IELTS Academic 7.5 (Nov 2025). **Track and field.** University team.