

Viacheslav Iablochnikov

Abu Dhabi, UAE | slavawertlon@gmail.com | [linkedin](#) | github.com/slava-qw

Research Interests: Generative models for image and video, vision-language models, scientific machine learning, Computer Vision.

Education

Mohamed bin Zayed University of Artificial Intelligence (MBZUAI)

M.Sc. in Computer Vision (**GPA: 3.73/4.00**), supervisor: [Ivan Laptev](#), Abu Dhabi, UAE

Aug 2025 – Present

Expected graduation: **06/2027**

Core Coursework: Deep Learning, Probabilistic Machine Learning, Computer Vision; *PhD-level courses:* ML801: Advanced Machine Learning (by [Eric Moulines](#)); ML815: Advanced Distributed and Parallel Machine Learning Systems (by [Qirong Ho](#)); ML8103: Sequential Decision Making (by [Salem Lahlou](#)). CV806: Advanced Topics in Vision and Language (by [Ivan Laptev](#))

National Research University Higher School of Economics (HSE)

B.Sc. in Physics (**GPA: 8.33/10**), Moscow, Russia

Sep 2021 – Jun 2025

Thesis: Diffusion models for astronomical image differencing. *Core coursework:* Linear Algebra, Calculus, Data structures and Algorithms, Probability and Statistics, Optimization, Machine Learning, Deep Learning

Research Experience

Computer Vision Researcher, MBZUAI, Abu Dhabi, UAE

Sep 2025 – Present

- Investigating dual diffusion generation utilizing 2D/3D conditioning to improve visual fidelity and enforce structural consistency, ensuring accurate object placement via layout generation.
- Develop physically informed video diffusion models for controllable object interactions, co-supervised by [Fabio Pizzati](#). Implemented and trained a Wan2.2 5B and Cosmos-Transfer2 2B image-to-video pipelines with ControlNet: set up data curation/conditioning, mixed-precision multi-GPU training and evaluation tooling, achieving more physically plausible videos than strong baselines and extending the [KineMask](#) framework (aimed for ECCV'26, [project page](#))
- Collaborating with [Senmao Li](#) under [Fahad Khan](#) on profiling and reducing inference time of transformer-based diffusion models (DiT/MMDiT) via caching/acceleration (aimed for ECCV'26). Investigating training-free inference adaptation of pretrained diffusion models for execution in low-dimensional spaces.

Research Intern, Institute of Smart Systems & AI (ISSAI), Nazarbayev University, Astana, Kazakhstan

Jun 2025 – Aug 2025

- Optimized [OYLAN2](#)'s next-generation text-to-image diffusion model, focusing on culturally faithful Kazakh imagery from multilingual prompts.
- Curated a targeted prompt-image corpus and trained LoRAs with [diffusers](#) to improve prompt adherence and cultural fidelity, shipping clean, reproducible training scripts and experiment logs to support reruns and handoffs.
- Full-fine-tuned the Lumina model to achieve the best multi-concept Gemini mean score of **88.29** (up from 53.60 for the base model, **65%** relative improvement).

Research Intern, GenAI Laboratory (AIRI & HSE), Moscow, Russia

Oct 2024 – May 2025

- Implemented techniques for face and hair swapping with generative models, building on GAN baselines such as [StyleGAN2](#) and [MegaFS](#).
- Applied attention mechanisms and rotary encoder in the FS+ latent space, cutting LPIPS by **40%** and FID1k by **87%** compared to baselines.

Research Intern (VK Company), Moscow, Russia

Apr 2024 – Oct 2024

- Conducted research on adversarial attacks on multimodal models and implemented attacks on visual and textual encoders for image retrieval and LLM-vision systems.
 - Verified and evaluated patch attacks for CLIP, [LLaVA](#)-family models and video retrieval using [LLaVA-OneVision](#).
 - Main author of "[Attacks on multimodal models](#)" paper, proposing new attack variants that reach up to **86%** success rate in forcing models to generate restricted content.
-

Projects

DreamBooth Lumina2 training with SingLoRA via MeanFlow

2025

- Implemented [DreamBooth](#)-style [Lumina2](#) fine-tuning with three key improvements: [SingLoRA](#) as a compact LoRA replacement, [MeanFlow](#) for stable one-step sampling, and [dispersive regularization](#) to reduce mode collapse.
- Optimized training/inference scripts for single- and multi-GPU setups, reducing VRAM usage and iteration time while maintaining sample quality and diversity. [Project page](#).

DiffDiff: diffusion for transient detection via image differencing (Bachelor Thesis, HSE)

2024 – 2025

- Designed an end-to-end astronomy pipeline for discovering *transient* events from sky survey images using diffusion models.
 - Trained a conditional diffusion model to produce clean difference images used for flux and coordinate estimation with uncertainty calibration for follow-up prioritization.
 - Achieved a **5.6**× speedup over the existing pipeline and up to **93%** detection accuracy. [Project page](#).
-

Technical Skills

Programming: Algorithms and Data Structures, Model Optimization, Inference, Python, Numpy, Scipy, Pandas, SQL, SQLite, MongoDB, Flask, FastAPI. ML Frameworks: PyTorch, Torch Profiler, JAX/Flax, Triton, HuggingFace `transformers/diffusers`, PEFT, `accelerate`, Megatron, DeepSpeed, XGBoost. Systems: Linux, SLURM, Git, CUDA, DVC, Docker, Kubernetes, PySpark, MLflow.

Teaching Experience: Teaching assistant for BSc courses “Algorithms & Data Structures” and “Deep Learning” at HSE University, 2023–2024.

Languages: English (C1), German (A2), Russian (native)

Publications, Workshops and Talks

- David Romero, **Viacheslav Iablochnikov**, Ariana Bermudez, Hao Li, Fabio Pizzati, Ivan Laptev *Learning to Generate Rigid Body Interactions with Video Diffusion Models* [Project page](#)
 - Nikita Ligostaev, Vladimir Latypov, **Viacheslav Iablochnikov**, Maria Nesterova, Ramil Khafizov, Sergei Kholkin, Alexander Korotin. *Categorical Iterative Proportional Fitting via Discrete Diffusion Models*. [Pre-print](#).
 - **V. Iablochnikov**, A. Rogachev. *Attacks on multimodal models*. [arXiv:2412.01725](#).
- School of Machine Learning (SMILES-2025), Harbin Institute of Technology (HIT), China
 - AIRI Summer School Poster, 2024 (**top 50**/1100);
 - Member of the “Embodied Perception” research reading group at MBZUAI;
 - Gave lecture talks about discrete diffusion models ([D3PM](#)), FlashAttention 1 and 2 ([FA1](#), [FA2](#)), and pose estimation ([PoseBH](#)), visual reasoning ([Corvid](#))

Honors and Awards

- Scholarship for both MSc years (on a competitive basis), 2025-2027
- Scholarship for excellent performance for all four BSc years (on a competitive basis), 2021-2025
- Prize-winner of the Deep Learning University Olympiad at Skoltech, 2025 (**top 3** out of ~1000)
- Prize-winner of the Deep Learning University Olympiad at MIPT, 2024 (**top 10** out of 300 participants)