

KIRILL STRUMINSKY

Research Scientist | 3D Vision · Generative Models · Differentiable Rendering

 k.struminsky@gmail.com · [Google Scholar](#) · [GitHub](#)

SUMMARY

Research scientist with 10+ years of experience in machine learning and computer vision, specialising in differentiable rendering, 3D reconstruction, and generative models. Currently working at the intersection of 3D vision, computer graphics, and generative modelling at Yandex Research, with recent work spanning physically-based relighting, neural rendering, and text-to-texture synthesis. Broadly interested in building generative world models for 3D scene understanding, photorealistic simulation, and interactive content creation.

RESEARCH INTERESTS

- **3D Vision & Synthesis:** Differentiable rendering, inverse rendering, physically-based relighting, 3D reconstruction, novel view synthesis, neural radiance fields, Gaussian splatting
 - **Generative Models:** Latent variable models, diffusion models (including video and multi-view diffusion), structured probabilistic models, approximate inference
 - **World Models & Simulation:** Leveraging generative modelling for scene-level understanding, compositional scene reconstruction, and physically grounded simulation
-

EDUCATION

Ph.D. in Computer Science — HSE University, Moscow (2015–2024)

Research Group for Bayesian Methods in Machine Learning

Thesis: *Learning Guarantees and Efficient Inference for Structured Prediction* | Advisor: Dmitry Vetrov | GPA: 5.0/5.0

B.A. + M.S. in Mathematics — Lomonosov Moscow State University, Faculty of Mechanics and Mathematics (2010–2015)

Specialisation: Kolmogorov complexity and computational complexity | GPA: 4.8/5.0 (with honours)

PROFESSIONAL EXPERIENCE

RESEARCH FELLOW — YANDEX RESEARCH LAB @ HSE UNIVERSITY (2024 – PRESENT)

HSE University hosts one of the strongest computer science programmes in Russia, with research groups backed by leading technology companies.

Conducting research at the intersection of 3D vision, computer graphics, and generative modelling. The overarching goal is to leverage generative modelling within scene reconstruction pipelines to empower downstream applications such as relighting, asset synthesis, and photorealistic simulation.

- Built a relightable reconstruction pipeline for 3D car assets, enabling scene-consistent relighting for driving scene simulation; also released [MADCars](#), a large-scale multi-view car dataset with rich annotations (MADrive, CVPR 2026 Findings)
- Developed a cascaded text-to-texture synthesis pipeline using explicit texture maps and physically-based shading (CasTex, WACV 2026)

RESEARCH FELLOW — BAYESGROUP @ HSE UNIVERSITY (2015 – 2024)

PhD student, then independent researcher. Began working on probabilistic machine learning under the supervision of Dmitry Vetrov, building a deep grounding in structured probabilistic models, approximate inference, and generative modelling. Over time, developed an increasingly independent research agenda, progressively moving toward 3D vision and neural rendering —

fields where the same probabilistic intuition translates naturally into differentiable, uncertainty-aware approaches to scene understanding. Supervised 5 graduate students and led several industry collaboration projects.

- Developed reparameterized volume sampling for differentiable rendering, enabling gradient-based optimisation through volumetric scenes (AISTATS 2024)
- Designed recursive Gumbel-Max estimators for approximate inference over combinatorial (graph/tree-structured) spaces (NeurIPS 2021)
- Proposed low-variance gradient estimators for ranking distributions, enabling differentiable learning over permutations (AAAI 2020, oral)

Industry Collaborations:

- *Huawei, Moscow (2021–2022)*: Built a generative data augmentation pipeline for image classification, validated on internal benchmarks; also conducted research on RL-based network distillation for improving model efficiency
- *Sber AI Lab, Moscow (2020–2021)*: Consulted on approaches to interpretable ML for GDPR-style compliance, focusing on causal learning methods for model transparency
- *IBM, Moscow (2018–2019)*: Designed an HMM-based pipeline for automated well log standardisation and anomaly detection, developed in collaboration with domain engineers and validated on production log data (published at Petroleum Geostatistics 2019)

SELECTED PUBLICATIONS

Publications marked with * denote equal contribution.

3D Vision & Rendering

- P. Karpikova*, D. Selikhanych*, **K. Struminsky***, R. Musaev, M. Golitsyna, D. Baranchuk. *MADrive: Memory-Augmented Driving Scene Modeling*. Findings of CVPR 2026.
- M. Aliev, D. Baranchuk, **K. Struminsky**. *CasTex: Cascaded Text-to-Texture Synthesis via Explicit Texture Maps and Physically-Based Shading*. WACV 2026.
- N. Morozov, D. Rakitin, O. Desheulin, D. Vetrov, **K. Struminsky**. *Differentiable Rendering with Reparameterized Volume Sampling*. AISTATS 2024 & Neural Fields workshop at ICLR 2023.

Probabilistic Models & Inference

- **K. Struminsky***, A. Gadetsky*, D. Rakitin*, D. Karpushkin, D. Vetrov. *Leveraging Recursive Gumbel-Max Trick for Approximate Inference in Combinatorial Spaces*. NeurIPS 2021.
- A. Gadetsky*, **K. Struminsky***, C. Robinson, N. Quadrianto, D. Vetrov. *Low-variance black-box gradient estimates for the Plackett-Luce distribution*. AAAI 2020 (oral) & BDL Workshop at NeurIPS 2019 (spotlight).
- A. Atanov*, A. Ashukha*, **K. Struminsky**, D. Vetrov, M. Welling. *The Deep Weight Prior: Modeling a prior distribution for CNNs using generative models*. ICLR 2019.
- **K. Struminsky**, S. Lacoste-Julien, A. Osokin. *Quantifying Learning Guarantees for Convex but Inconsistent Surrogates*. NeurIPS 2018.
- **K. Struminsky**, D. Vetrov. *A Simple Method to Evaluate Support Size and Non-uniformity of a Decoder-Based Generative Model*. AIST 2019.
- **K. Struminsky** et al. *Well Log Data Standardization, Imputation, and Anomaly Detection Using Hidden Markov Models*. Petroleum Geostatistics 2019.
- M. Figurnov, **K. Struminsky**, D. Vetrov. *Noise-robust method for training of variational autoencoder*. Intellectual Systems: Theory and Applications, 2017.
- M. Figurnov, **K. Struminsky**, D. Vetrov. *Robust Variational Inference*. NIPS Workshop on Advances in Approximate Bayesian Inference, 2016.
- **K. Struminsky**, S. Kruglik, D. Vetrov, I. Oseledets. *A New Approach for Sparse Bayesian Channel Estimation in SCMA Uplink Systems*. WCSP 2016.

TEACHING EXPERIENCE

Lecturer, 3D Computer Vision — HSE University (2024 – present) · [Course repository](#)

Designed and deliver a graduate course covering 3D reconstruction methods (SfM, radiance fields, Gaussian splatting, multi-view foundational transformers), inverse graphics, discriminative models on 3D data, and generative models for 3D. Several lectures from the course were also delivered at the Yandex School of Data Analysis.

Lecturer, Deep Generative Models — Yandex Data School, Tel-Aviv (Spring 2019)

Teaching Fellow — Deep Bayes Summer School (2017–2019)

Teaching Assistant, Bayesian Methods in ML / Bayesian Deep Learning — MSU & Yandex School of Data Analysis (2017–2022)

Seminar Organiser — Research Seminar in ML and Applications, HSE University (2016–2022)

Organised a weekly research seminar series for undergraduate and graduate students.

TECHNICAL SKILLS

- **Languages & Frameworks:** Python, PyTorch, JAX, C++ (basic)
 - **3D & Rendering:** Differentiable rendering, physically-based shading, volumetric rendering, NeRF-based methods, Gaussian splatting, texture synthesis
 - **ML Expertise:** Generative models (VAEs, diffusion), approximate Bayesian inference, structured prediction, gradient estimation
 - **Tools:** CUDA, Git, Linux
-

REVIEWING

Conferences: AISTATS, ICML (top 33%, 2020), NeurIPS (top 50%, 2019), CVPR, ICCV, ICLR

Journals: IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Journal of Machine Learning Research (JMLR)