

ALEKSEI PETRENKO

Email: apetrenko1991@gmail.com
Based in: Mountain View, CA, US

Phone: +1 (323) 961-17-12

Website: <https://alex-petrenko.github.io>
Github: <https://github.com/alex-petrenko>

Reinforcement learning expert (esp. policy gradient/trust region methods). Versatile researcher and programmer with deep expertise in C++, Python, software architecture, optimization, high-throughput systems, parallel computing.

EDUCATION & RESEARCH

2018-2022, USA:

PhD in Computer Science at University of Southern California

Advisor: Prof. Gaurav Sukhatme, USC Robotic Embedded Systems Lab

Research areas: deep reinforcement learning, machine learning systems, high-throughput learning and simulation, distributed RL, multi-agent learning and self-play, robotics, sim-to-real.

2008-2014, Russia:

M. Sc. in Computer Science at Nizhny Novgorod State Technical University (GPA 5.0/5.0)

B. Sc. in Computer Science at Nizhny Novgorod State Technical University (GPA 4.95/5.0)

RECENT PUBLICATIONS

2022:

- **A Petrenko**, A Allshire, G State, A Handa, V Makoviychuk. “DexPBT: Scaling up Dexterous Manipulation for Hand-Arm Systems with Population Based Training.” Submitted to RSS 2023. [Website](#)
- A Handa*, A Allshire*, V Makoviychuk*, **A Petrenko***, R Singh*, J Liu*, D Makoviichuk, K Van Wyk, A Zhurkevich, B Sundaralingam, Y Narang, J Lafleche, D Fox, G State. “DeXtreme: Transfer of Agile In-hand Manipulation from Simulation to Reality.” In ICRA, 2023. [Arxiv](#), [Website](#)

2021:

- **A Petrenko**, E Wijmans, B Shacklett, V Koltun. “Megaverse: Simulating Embodied Agents at One Million Experiences per Second.” In ICML, 2021. [Arxiv](#), [Github](#), [Website](#)
- S Batra*, Z Huang*, **A Petrenko***, T Kumar, A Molchanov, G Sukhatme. “Decentralized Control of Quadrotor Swarms with End-to-end Deep Reinforcement Learning”. In CORL, 2021. [Arxiv](#), [Github](#), [Website](#)
- S Hegde, A Kanervisto, **A Petrenko**. “Agents that Listen: High-Throughput Reinforcement Learning with Multiple Sensory Systems”. In IEEE Conference on Games, 2021. [Arxiv](#), [Github](#), [Website](#)
- B Shacklett, E Wijmans, **A Petrenko**, M Savva, D Batra, V Koltun, K Fatahalian. “Large Batch Simulation for Deep Reinforcement Learning.” In ICLR, 2021. [Arxiv](#), [Github](#)

2020:

- **A Petrenko**, Z Huang, T Kumar, G Sukhatme, V Koltun. “Sample Factory: Egocentric 3D Control from Pixels at 100,000 FPS with Asynchronous Reinforcement Learning”. In ICML, 2020. [Arxiv](#), [Github](#), [Website](#)

SKILLSET

Reinforcement learning:

- I dedicated my PhD to studying deep reinforcement learning and its applications. I am an expert in high-throughput learning systems and parallel simulation who authored multiple open-source libraries.
- Studied and/or extended numerous RL and simulation frameworks: Sample Factory, RLLIB, Baselines, IMPALA (scalable-agent), Seed RL, rlyt, OpenAI Gym, Isaac Gym, DeepMind Lab, VizDoom, etc.

Programming languages:

- C++ and Python (advanced).
- C#, Java, Javascript (extensive experience in industry, proficient with some catching up).

Other skills: software development and architecture, deep learning, algorithms, computer vision, clusters, 3D.

EMPLOYMENT HISTORY

Sr. Research Scientist at Apple

2023-present, Cupertino, USA

I work in Vladlen Koltun's group on embodied AI and robotic autonomy.

Robotics/AI Research Intern at [NVIDIA](#)

2021-2022, Los Angeles, USA

- Created GPU-accelerated vectorized robotic simulations featuring extensive domain randomization for zero shot sim-to-real transfer.
- Developed large-scale reinforcement learning and population-based training algorithms to learn in-hand manipulation, as well as object grasping and placement with hand-arm systems.
- This work resulted in two publications: "DeXtreme" (in-hand manipulation with real robotic hand) and "DexPBT" (dexterous robotic control with high-DoF hand-arm systems), both in-review @ ICRA 2023.
- Video demonstrations for [DexPBT](#) and [DeXtreme](#) (coming soon).

Technologies: Python, PyTorch, PhysX, Isaac Gym, Docker, WandB.

PhD Research Intern at [Intel Labs](#)

2019-2020, Santa Clara & Los Angeles, USA

- Developed a new asynchronous reinforcement learning system that can achieve up to 150000 FPS of training throughput on a single GPU in immersive environments with visual observations.
- Using this framework, achieved new SOTA in VizDoom, including challenging multi-agent and self-play scenarios. Project titled "[Sample Factory](#)" was presented at ICML 2020.
- Created a high-throughput simulator "[Megaverse](#)" capable of rendering 1,000,000 observations per second on a single machine (presented at ICML 2021).

Technologies: C++17, Python, PyTorch, TensorFlow, RLLib, Vulkan, OpenGL, OpenAI Gym, Slurm.

Tech Lead at [itSeez3D](#)

2015-2018, Nizhny Novgorod, Russia

Sr. Software Engineer at [itSeez3D](#)

2013-2015, Nizhny Novgorod, Russia

ItSeez/itSeez3D was a company behind the computer vision library OpenCV where I worked on computer vision, 3D scanning, geometric deep learning, and virtual reality applications.

- Developed backend (algorithms & server-side applications) and frontend (mobile/desktop apps) for the first portable consumer [3D scanner](#).
- Lead algorithms and systems developer of [Avatar SDK](#), an AI-based human digitization engine powered by geometric deep learning.
- R&D for the automatic generator of full-body characters (skeleton animation, rendering, deep learning)
- Optimized deep neural networks for deployment (quantization, compression, etc.)
- Worked on various applications for the avatar engine: mobile, desktop, and virtual reality. Developed and deployed interactive demos for major conferences and exhibitions: SIGGRAPH, GDC, VRX.

My avatar automatically created from a selfie by the Avatar SDK engine (circa 2016): <https://skfb.ly/6FUYY>

Technologies: C++11/14, OpenCV, Python, Django, TensorFlow, Keras, OpenGL, Unity, Unreal, Oculus VR.

Software Engineer at [Tecom](#)

2010-2013, Nizhny Novgorod, Russia

Developed a big real-time (frame perfect) massively parallel broadcast automation system with many levels of redundancy. My responsibilities ranged from R&D to maintenance with emphasis on software architecture and optimization.

Technologies: C++/Qt, C#, Java, Javascript, Python, Linux.

SELECTED OPEN-SOURCE PROJECTS

- <https://github.com/alex-petrenko/sample-factory> at the time of release the fastest open-source PPO implementation.
- <https://github.com/alex-petrenko/megaverse> embodied simulation for AI research at 1,000,000 FPS.
- <https://github.com/alex-petrenko/faster-fifo> faster alternative to built-in Python's multiprocessing queue.
- <https://github.com/alex-petrenko/signal-slot> Python implementation of Qt-like signal & slot asynchronous

programming paradigm with focus on multiprocessing applications.

- <https://github.com/alex-petrenko/rl-experiments> RL algorithms (Double DQN, A2C) for experimental gridworld RL environment called MicroTbs.
- <https://github.com/alex-petrenko/4dvideo> volumetric video recorder and player for Intel RealSense and Google Tango. Capture and playback “4D” clips like [this](#) in real time with a mobile device. Features a very fast Delaunay triangulation [algorithm](#) based on modified Guibas-Stolfi method.

OTHER INTERESTS

- Futurology, transhumanism, singularity. Authors: G. Egan, R. Kurzweil, K. Eric Drexler, V. Vinge, E. Yudkowsky, N. Bostrom.
- Scientific discovery: physics, cosmology, evolutionary biology, human longevity.
- Hobbies: motor racing, go-karting, travel, diving, chess.