Anwar Ibrahim

Machine Learning Researcher | Machine Learning Engineer

Moscow, Russia ▪ anwar96ibrahim@gmail.com

anwar9ibrahim.github.io ▪ github.com/Anwar9Ibrahim ▪ linkedin.com/in/anwar-ibrahim-7116b4205

HSE University profile: hse.ru/en/org/persons/494515691

Professional Summary

Machine Learning Researcher and Machine Learning Engineer with experience developing reinforcement learning (RL), research software, production-oriented machine learning, and agentic large language model (LLM) systems. Expertise includes Python, PyTorch, continuous-control methods, simulation environments, benchmarking against classical optimization, model evaluation, deployment, and monitoring. First author of three peer-reviewed journal articles and lecturer in generative models and practical machine learning.

Work Experience

Junior Research Fellow
HSE University — LAMBDA, Moscow, Russia | **Oct 2023–Present**

- Design and implement reinforcement-learning methods and Python research software for accelerator beamline control.
- Build Gymnasium-compatible simulation environments that integrate learning agents with Elegant and Self Describing Data Sets (SDDS) accelerator tools.
- Train and benchmark Deep Deterministic Policy Gradient (DDPG), Soft Actor-Critic (SAC), and Twin Delayed DDPG (TD3) agents against differential-evolution and Bayesian-optimization baselines.
- Authored three peer-reviewed journal articles as first author on simulation-based accelerator control and present results at international conferences and workshops.

Lecturer — Generative Models in Machine Learning
HSE University, Moscow, Russia | **Jan 2025–Present**

- Teach *Generative Models in Machine Learning* in English, covering variational autoencoders, generative adversarial networks (GANs), autoregressive models, and normalizing flows.
- Deliver practical machine-learning instruction and contribute to a reinforcement-learning competition at the All-Russian High-Energy Physics School.

AI Engineer — Agentic Systems
Eolos AI | **Dec 2025–Oct 2026**

- Developed components for private agentic artificial intelligence systems built around large language models and multi-step LangGraph workflows.
- Designed and implemented human-in-the-loop review and feedback paths, with application data managed in Convex.
- Used LangSmith tracing and evaluation to support the observability and reliability of LLM workflows.

Machine Learning Engineer
EXPentech, Damascus, Syria | **Jun 2020–Sep 2023**

- Developed and evaluated machine-learning models and prepared datasets through exploratory analysis.
- Supported model deployment and monitoring in production-oriented machine-learning workflows.

Education

PhD Candidate, Artificial Intelligence and Machine Learning
HSE University, Moscow, Russia | **Nov 2023–Present**

- Doctoral research: reinforcement-learning-based control of accelerator beamlines.

Master's in Data Science
HSE University, Moscow, Russia | **2021–Jul 2023**

- Degree programme in Applied Mathematics and Information Science.

Bachelor's in Computer Software Engineering
Higher Institute for Applied Sciences and Technology, Damascus, Syria | **2015–2020**

Peer-Reviewed Publications

1. **Anwar Ibrahim**, Fedor Ratnikov, Maxim Kaledin, Alexey Petrenko, Denis Derkach. *RL-ABC: Reinforcement Learning for Accelerator Beamline Control.*
   Computer Physics Communications, 329, 110407 (2026).
2. **Anwar Ibrahim**, Alexey Petrenko, Maxim Kaledin, Ehab Suleiman, Fedor Ratnikov, Denis Derkach. *Reinforcement Learning for Accelerator Beamline Control: A Simulation-Based Approach.*
   International Journal of Modern Physics E, 35(07) (2026).
3. **Anwar Ibrahim**, Denis Derkach, Alexey Petrenko, Fedor Ratnikov, Maxim Kaledin. *Optimization of the Accelerator Control by Reinforcement Learning: A Simulation-Based Approach.*
   Physics of Particles and Nuclei, 56(6), 1476–1481 (2025).

Selected Conference Contributions

- **CHEP 2026, Bangkok — Talk:** RL-ABC: Reinforcement Learning for Accelerator Beamline Control.
- **ACAT 2025 — Oral contribution:** Optimisation of the Accelerator Control by Reinforcement Learning: A Simulation-Based Approach.
- **NUCLEUS 2025 — Talk:** Reinforcement Learning for Accelerator Beamline Control.
- **CERN Workshop 2025:** Reinforcement Learning for Accelerator Beamline Control.
- **MMCP 2024:** Optimization of Accelerator Control by Reinforcement Learning.

Selected Projects

RL-ABC — GitHub repository
Research framework that converts Elegant accelerator lattice configurations into reinforcement-learning environments, with continuous-control agents and classical optimization baselines.

Implicit Q-Learning on Hopper — GitHub repository
PyTorch implementation of Implicit Q-Learning on the D4RL Hopper expert dataset, including actor, value, twin-critic, and target networks.

Interpretable Reinforcement Learning for Pong — GitHub repository
Analysis toolkit using saliency maps, Q-value inspection, paddle-position analysis, and an interactive dashboard.

Multilingual Contexto — GitHub repository
Semantic word-guessing game built with BERT-family models and Streamlit, supporting Arabic, English, and Russian.

Technical Skills

Machine Learning: Reinforcement learning, generative models, diffusion models, natural language processing (NLP), computer vision, human-in-the-loop systems

Programming: Python, C#, SQL

Frameworks: PyTorch, TensorFlow, Gymnasium, Stable-Baselines3, Streamlit, OpenCV

Agentic AI: LangChain, LangGraph, LangSmith, Convex, large language model (LLM) workflow orchestration

Scientific Tools: Elegant, SDDS, simulation-based optimization, experiment tracking

Engineering: Git, model evaluation, deployment, monitoring

Languages: Arabic, English, Russian