

ALEXANDER SHMAKOV**ARTIFICIAL INTELLIGENCE SCIENTIST**

I am a current Ph.D. student at the UCI Institute for Genomics and Bioinformatics with training in deep learning and artificial intelligence. I specialize in reinforcement learning, scientific computing, neural network design, and optimization using *Python*, *C++*, *PyTorch*, and *Jax*.

-
- EDUCATION** **Undergraduate study at University of California Irvine** *September 2015 - March 2020*
- Computer Science B.S.** *cum laude*. 3.92 Major GPA.
 - Mathematics B.S.** *magna cum laude*. *Pi Mu Epsilon*. 3.88 Major GPA.
 - Campuswide Honors Collegium**. 13 Dean's List awards.
- Graduate study at University of California Irvine** *June 2020 - June 2025 (Expected)*
- Computer Science Ph.D.** 3rd year. 4.0 GPA.
 - Primary Advisor** Distinguished Professor Pierre Baldi.
 - NSF Machine Learning and Physical Sciences (UCI MAPS)** training program fellow.
-

- EXPERIENCE** **Deep Learning Research, UCI Institute for Genomics and Bioinformatics** *January 2018 - Present*
- Scientific deep learning with applications to high-energy physics, chemistry, astronomy, and biology.
 - Inventing deep learning methods which integrate physical symmetries and domain knowledge directly into network architectures, incorporating cross-disciplinary knowledge from domain experts to produce state-of-the-art results for experimental high-energy physics at the Large Hadron Collider.
 - Designing and implementing deep learning algorithms and experiments using *Python*, *C++*, and *PyTorch*.
- Reinforcement Learning Research, UCI Intelligent Dynamics Lab** *September 2020 - Present*
- Deep reinforcement learning for applications to robotics, planning, and multiplayer games.
 - Produced headlining work to design the first agent that could solve the Rubik's Cube with no human knowledge.
 - Developing reinforcement learning agents to incorporate long-term memory using *Python*, *C++*, *PyTorch*, and *Jax*.
- Artificial Intelligence Research Associate, HPE Labs** *June 2021 - Present*
- Deep multi-agent reinforcement learning for applications to industrial control and green energy.
 - Realizing safe and trustworthy AI solutions to reduce leveled cost of wave-powered green energy.
 - Designing and implementing reinforcement learning methods and training algorithms to maximize performance while ensuring trustworthiness using *Python*, *PyTorch*, and *RLlib*.
- Neurophysiology Research, Stanford Hansen Experimental Physics Laboratory** *Summer 2019*
- Brain-retina signal modeling with deep learning for artificial retinal implants to restore vision.
 - Designed recurrent convolutional networks to accurately predict the behavior of complex retinal spiking timeseries.
 - Achieved state-of-the-art accuracy on modeling ganglion cell activity in response to stochastic visual stimulus.
- Undergraduate Research Program, UCI Institute for Genomics and Bioinformatics** *Summer 2018*
- Distributed Monte Carlo Tree Search for reconstructing protein tertiary structure using deep reinforcement learning.
 - Designed unique convolutional approach to model grid-based spatial protein folding.
 - Implemented a distributed computing system to efficiently run large protein simulation and neural network training across multiple CPUs and GPUs using *Python*, *C++*, and *TensorFlow*.
- R&D Internship, Pivotal Systems - Fremont** *Summer 2016 & Summer 2017*
- Developed automated methods for calibrating embedded systems using uncertainty-aware statistical methods.
 - Improved critical sensor accuracy in a variety of conditions including extreme temperatures and pressures.
 - Architected automated testing procedures and communicated with production engineers to deploy guided testing in production systems.
-

- PUBLICATIONS** **Solving the Rubik's Cube with Deep Reinforcement Learning and Search**
 Forest Agostinelli†, Stephen McAleert†, Alexander Shmakov†, Pierre Baldi. († equal contribution)
Nature Machine Intelligence, 2019, www.nature.com/articles/s42256-019-0070-z
- SPANet: Generalized Permutationless Set Assignment for Particle Physics using Symmetry Preserving Attention**
 Alexander Shmakov†, Michael Fenton†, Ta-Wei Ho, Shih-Chieh Hsu, Daniel Whiteson, Pierre Baldi
SciPost, 2022, www.scipost.org/SciPostPhys.12.5.178
-

Skip Training for Multi-Agent Reinforcement Learning Controller for Industrial Wave Energy Converters

S. Sarkar, V. Gundecha, S. Ghorbanpour, A. Shmakov, A. R. Babu, A. Pichard, M. Cocho

IEEE CASE, 2022, arxiv.org/abs/2209.05656

Best Application Paper Award

Rxn Hypergraph: A Hypergraph Attention Model for Chemical Reaction Representation

Mohammadamin Tavakoli[†], Alexander Shmakov[†], Francesco Ceccarelli, Pierre Baldi († equal contribution)

NeurIPS Machine Learning for Physical Sciences Workshop, 2021, arxiv.org/abs/2201.01196.pdf

Deep Learning Models of the Discrete Component of the Galactic Interstellar Gamma-Ray Emission

A. Shmakov, M. Tavakoli, P. Baldi, C. M. Karwin, A. Broughton, S. Murgia

In Review, 2022, arxiv.org/abs/2206.02819.pdf

Decoding Network-Mediated Retinal Response to Electrical Stimulation: Implications for Fidelity of Prosthetic Vision

Elton Ho, Alexander Shmakov, Daniel Palanker.

Journal of Neural Engineering, 2020, iopscience.iop.org/article/10.1088/1741-2552/abc535

ColosseumRL: A Framework for Multiagent Reinforcement Learning in N-Player Games

Alexander Shmakov, John Lanier, Stephen McAleer, Rohan Achar, Christina Lopes, Pierre Baldi

COMARL AAAI, 2020, arxiv.org/abs/1912.04451

Deep Learning for Drug Discovery and Cancer Research: Automated Analysis of Vascularization Images

Urban G., Bache K., Phan D.T.T., Sobrino A., Shmakov A., Hache S.J., Hughes C.C.W., Baldi P.

IEEE / ACM, 2019, ieeexplore.ieee.org/document/8368281

PROJECTS SPANet, Symmetry Preserving Attention Networks github.com/Alexanders101/SPANet

- An open-source python library implementing my paper on symmetry preserving attention networks for event reconstruction at the Large Hadron Collider. Used by teams at the LHC for performing event reconstruction.
- Implemented a general group-theoretic approach to dynamically create transformer attention networks specially designed to account for the symmetries and invariances in high energy physics classifications.

Fermilab NOvA, Event Detection and Classification novaexperiment.fnal.gov

- Implementing attention-based neural networks for interpretable event reconstruction on the NOvA experiment at Fermilab, improving the characterization of neutrino oscillations through more accurate classification.
- Designing deep neural networks and working with a diverse group of inter-disciplinary scientists from Fermilab to improve discovery potential from neutrino decays for long-running experiment.

TorchSpread, PyTorch Distributed RL Library github.com/Alexanders101/TorchSpread

- Building a python library to support distributed reinforcement learning development by providing utilities for distributed prediction with dynamic batching, transporting tensors between processes or machines, and reinforcement learning data structures.

ColosseumRL, Multi-Agent RL Framework colosseumrl.igb.uci.edu

- Oversaw a school-wide competition for reinforcement learning to gain an understanding of multiplayer games and the best way to train agents for them.
- Mentoring students in a project-oriented course, I helped design reinforcement learning algorithms and fostered interest in artificial intelligence.

EmojiNet, Project Class in Artificial Intelligence

- Created Emoji-Net, an application that can read hand gestures from a phone camera and detect the presence of a hand-based emoji in real time.
- Implemented a convolution network that was pretrained on a provided dataset of hand gestures to minimize required data collection.

SKILLS Programming Languages: Python, C/C++, C#

Machine Learning Frameworks: PyTorch, Jax, Tensorflow, RLLib, DeepMind ACME, Scikit-Learn, OpenCV

Artificial Intelligence Techniques: Attention and Transformers, Physics-informed Machine Learning, Computer Vision, Reinforcement learning, Bioinformatics, Computational Chemistry, Scientific Computing
