John Banister Lanier

Irvine, California (USA) | (205) 994-0128 | jblanier@uci.edu
https://jblanier.net

Computer Science Ph.D. student at UC Irvine specializing in deep reinforcement learning for sim-to-real robotics and multi-agent systems.

EXPERIENCE

University of California Irvine
Ph.D. Student, (09/2020)-*Present*
Advised by: Roy Fox and Pierre Baldi

Specialization in Deep Reinforcement Learning. Published multiple top-conference papers in efficiently solving two-player competitive games with populations of agents. Currently working on simulation-to-real transfer for vision-based robotics with a focus on creating better model-based RL methods for fast adaptation.

OffWorld, Inc., Pasadena, California
Machine Learning & Robotics Research Intern, (03/2020)-(09/2020)

Developed, tested, and maintained publicly available sim-to-real deep reinforcement learning robotics environment with ROS Gazebo simulation and real-world automated rover environment controlled through web API and frontend. Developed robot computer vision and navigation systems for object retrieval. Generated synthetic computer vision training data using simulated environments and domain randomization. Built simulated rover navigation, SLAM, and object-detection vision systems.

University of California Irvine
M.S. Research in Artificial Intelligence, (01/2018)-(03/2020)
Advised by: Pierre Baldi

Designed novel deep reinforcement learning algorithms to solve sparse reward goal-oriented robotics environments. Developed large-scale distributed deep reinforcement learning applications with Tensorflow. Created online multi-agent reinforcement learning environment framework for competitions and research. Developed SOTA method for deep reinforcement learning in large-scale imperfect-information zero-sum games.

Novacoast, Inc., Santa Barbara, California
Software Engineering Intern, (06/2017)-(09/2017)

Developed public key infrastructure management system for Motorola to securely distribute certificates to untrusted environments with fault-tolerant transfer and storage. Frontend developed with Ruby on Rails, backend applications in Java with MongoDB and PostgreSQL. Designed NoSQL database model. Microservices managed with Docker.

University of California Santa Barbara
Undergraduate Research in Augmented Reality, (01/2017)-(06/2017)
Advised by: Tobias Höllerer and Matthew Turk

Developed HoloLens augmented reality real-time application. Built a remote YOLO object detection inference server to label objects in the user's view and provide real-time AR annotations for their labels in other languages.

EDUCATION

Ph.D. Computer Science, University of California Irvine *(Ongoing)*
M.S. Computer Science, 2019, University of California Irvine
B.S. Computer Science, 2017, University of California Santa Barbara

PUBLICATIONS

Toward Optimal Policy Population Growth in Two-Player Zero-Sum Games.
ICLR 2024.
Stephen McAleer, **JB Lanier**, Kevin Wang, Pierre Baldi, Tuomas Sandholm, Roy Fox.

Selective Perception: Learning Concise State Descriptions for Language Model Actors.
NAACL 2024.
Kolby Knottingham, Yasaman Razeghi, Kyungmin Kim, **JB Lanier**, Pierre Baldi, Roy Fox, Sameer Singh.

Feasible Adversarial Robust Reinforcement Learning for Underspecified Environments.
NeurIPS 2022 Deep RL Workshop.
JB Lanier, Stephen McAleer, Pierre Baldi, Roy Fox.

Self-Play PSRO: Toward Optimal Populations in Two-Player Zero-Sum Games.
arxiv (2022).
Stephen McAleer, **JB Lanier**, Kevin Wang, Pierre Baldi, Roy Fox, Tuomas Sandholm.

Anytime PSRO for Two-Player Zero-Sum Games.
arxiv (2022).
Stephen McAleer, Kevin Wang, **JB Lanier**, Marc Lanctot, Pierre Baldi, Tuomas Sandholm, Roy Fox.

XDO: A Double Oracle Algorithm for Extensive-Form Games.
NeurIPS 2021.
Stephen McAleer, **JB Lanier**, Kevin Wang, Roy Fox, Pierre Baldi.

Improving Social Welfare While Preserving Autonomy via a Pareto Mediator.
arxiv (2021).
Stephen McAleer, **JB Lanier**, Michael Dennis, Pierre Baldi, Roy Fox.

OffWorld Gym: Open-Access Physical Lunar Analog Environment for Reinforcement Learning and Robotics Research.
43rd COSPAR Scientific Assembly (2021).
Ashish Kumar, Toby Buckley, **JB Lanier**, Qiaozhi Wang, Alicia Kavelaars, Ilya Kuzovkin.

Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games.
NeurIPS 2020.
Stephen McAleer*, **JB Lanier***, Roy Fox, Pierre Baldi (*equal contribution).

ColosseumRL: A Framework for Multiagent Reinforcement Learning in N-Player Games.
COMARL AAAI 2020.
Alex Shmakov, **JB Lanier**, Stephen McAleer, Rohan Archar, Cristina Lopes, Pierre Baldi.

Curiosity-Driven Multi-Criteria Hindsight Experience Replay.
NeurIPS 2019 Deep RL Workshop.
JB Lanier, Stephen McAleer, Pierre Baldi.

TECHNICAL SKILLS

- **Languages:** Python | C++ | Bash | Javascript | HTML | CSS | Java | C | C#
- **Platforms/APIs**: JAX | PyTorch | Tensorflow | Linux | ROS | Unreal Engine
- **Software Development Skills**: Reinforcement Learning | Distributed ML | Sim-to-Real | Robotics | Game-Dev | Web-Dev