

# John Banister Lanier

Irvine, California | (205) 994-0128 | jblanier@uci.edu

<https://jblanier.net>

Computer Science Ph.D. student at UC Irvine specializing in Deep Reinforcement Learning.

## EDUCATION

Ph.D. Computer Science, University of California Irvine (*Ongoing*)

M.S. Computer Science, 2019, University of California Irvine, GPA 3.946

B.S. Computer Science, 2017, University of California Santa Barbara, GPA 3.70

## EXPERIENCE

### University of California Irvine

*Ph.D. Student, (09/2020)-Present*

*Advised by: Pierre Baldi and Roy Fox*

Focus on Deep Reinforcement Learning, using multi-agent methods to elicit complex and robust emergent behavior for competition, cooperation, and human-computer interaction.

### 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 full-stack PKI management system with multiple microservices to enable key management and distribution to untrusted environments with secure and 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 object-annotation application. Created streaming solution for server-side inference using YOLO CNN to detect and label objects in client device's field of view. 2D labels are projected as 3D annotations alongside respective real-world objects by raycasting virtual image positions onto the HoloLens's spatial map. Server application developed in C, HoloLens application developed in C# with Unity game engine.

## PUBLICATIONS

Feasible Adversarial Robust Reinforcement Learning for Underspecified Environments.

**JB Lanier**, Stephen McAleer, Pierre Baldi, Roy Fox. *In submission*.

Self-Play PSRO: Toward Optimal Populations in Two-Player Zero-Sum Games.

Stephen McAleer, **JB Lanier**, Kevin Wang, Pierre Baldi, Roy Fox, Tuomas Sandholm. *In submission*.

Anytime PSRO for Two-Player Zero-Sum Games.

Stephen McAleer, Kevin Wang, **JB Lanier**, Marc Lanctot, Pierre Baldi, Tuomas Sandholm, Roy Fox. *In submission*.

XDO: A Double Oracle Algorithm for Extensive-Form Games.

Stephen McAleer, **JB Lanier**, Kevin Wang, Roy Fox, Pierre Baldi. *NeurIPS 2021*.

Improving Social Welfare While Preserving Autonomy via a Pareto Mediator.

Stephen McAleer, **JB Lanier**, Michael Dennis, Pierre Baldi, Roy Fox. *arXiv preprint arXiv:2106.03927 (2021)*.

OffWorld Gym: Open-Access Physical Lunar Analog Environment for Reinforcement Learning and Robotics Research.

Ashish Kumar, Toby Buckley, **JB Lanier**, Qiaozhi Wang, Alicia Kavelaars, Ilya Kuzovkin. *43rd COSPAR Scientific Assembly (2021)*.

Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games.

Stephen McAleer\*, **JB Lanier\***, Roy Fox, Pierre Baldi. *NeurIPS 2020*.

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

Alex Shmakov, **JB Lanier**, Stephen McAleer, Rohan Archar, Cristina Lopes, Pierre Baldi. *COMARL AAAI 2020*.

Curiosity-Driven Multi-Criteria Hindsight Experience Replay.

**JB Lanier**, Stephen McAleer, Pierre Baldi. *NeurIPS 2019 Deep RL Workshop*.

\* denotes equal contribution.

## TECHNICAL SKILLS

- **Languages:** Python | Java | Javascript | C | C++ | C# | Bash | HTML | CSS | SQL
- **Platforms/APIs:** Tensorflow | PyTorch | Jupyter Notebook | Git | Docker | Linux | ROS | MongoDB | Node.js | Unity | Unreal Engine
- **Software Development Skills:** Deep Reinforcement Learning | Distributed ML | Full-stack web development