

Stephen McAleer

Education

- 2017–2021 **Ph.D. Computer Science**, *University of California, Irvine*.
Advisor: Pierre Baldi
- 2017–2019 **M.S. Computer Science**, *University of California, Irvine*.
- 2013–2017 **B.S. Mathematics, Economics**, *Arizona State University*.

Experience

Research

- 2022 **Postdoctoral Researcher**, *Carnegie Mellon University*, Pittsburgh, PA.
Advisor: Tuomas Sandholm
- 2021 **Research Scientist Intern**, *DeepMind*, Remote.
Advisors: Laurent Orseau and Marc Lanctot
- Summer 2019 **Research Scientist Intern**, *Intel AI*, San Diego, CA.
Research multiagent deep reinforcement learning
- 2017–Present **Graduate Researcher**, *Baldi Lab, UCI*, Irvine, CA.
Research deep reinforcement learning and applications in science.
- 2016–2017 **Undergraduate Researcher**, *Active Perception Group, ASU*, Tempe, AZ.
Research deep generative models and computer vision under Yezhou Yang.
- 2016–2017 **Undergraduate Researcher**, *Data Mining and Machine Learning Lab, ASU*, Tempe, AZ.
Research deep learning and computational biology under Huan Liu.

Professional

- Summer 2016 **Investment Banking Summer Analyst**, *Goldman Sachs*, New York, NY.
Worked in the Financial Institutions Group covering financial technology and insurance.
- Summer 2015 **Investment Management Summer Analyst**, *Goldman Sachs*, New York, NY.
Rotational program between the Private Equity Group and the Real Estate Group.
- Fall 2013 **Software Engineer**, *ASU Biodesign Institute*, Tempe, AZ.
Developed iOS application allowing users to compare evolutionary trees of different species.
- Summer 2013 **Software Engineer**, *Integrated Engineering Solutions*, Pullman, WA.
Developed iPhone application and implemented new audio codec for use in streaming.

Teaching

- Spring 2019 **Teaching Assistant**, *University of California, Irvine*, Computer Science 175: Deep Reinforcement Learning.
- Fall 2018 **Teaching Assistant**, *University of California, Irvine*, Statistics 7: Basic Statistics.

- Spring 2018 **Teaching Assistant**, *University of California, Irvine*, Statistics 67: Introduction to Probability and Statistics for Computer Science.
- Winter 2018 **Teaching Assistant**, *University of California, Irvine*, Statistics 7: Basic Statistics.
- Fall 2017 **Teaching Assistant**, *University of California, Irvine*, Statistics 7: Basic Statistics.
[Service](#)
- Summer 2017 **International Lead**, *Harvard University*, Adama, Ethiopia.
Organized and led an effort to implement CS50 throughout universities in Ethiopia

Papers

Journal Publications

- 2021 **Stephen McAleer**, Alex Fast, Yuntian Xue, Magdalene Seiler, William Tang, Mihaela Balu, Pierre Baldi, Andrew W Browne. Deep learning-assisted multiphoton microscopy to reduce light exposure and expedite imaging in tissues with high and low light sensitivity. *Translational Vision Science and Technology*.
- 2019 Forest Agostinelli*, **Stephen McAleer***, Alexander Shmakov*, Pierre Baldi. Solving the Rubik's Cube with Deep Reinforcement Learning and Search. *Nature Machine Intelligence*.
- 2018 Siyu Shao, **Stephen McAleer**, Ruqiang Yan, Pierre Baldi. Highly-Accurate Machine Fault Diagnosis Using Deep Transfer Learning. *IEEE Transactions on Industrial Informatics*.

Conference Publications

- 2022 Litian Liang, Yaosheng Xu, **Stephen McAleer**, Dailin Hu, Alexander Ihler, Pieter Abbeel, Roy Fox. Reducing Variance in Temporal-Difference Value Estimation via Ensemble of Deep Networks. *International Conference on Machine Learning (ICML) 2022*.
- 2022 Eser Ayygün, Laurent Orseau, Ankit Anand, Xavier Glorot, **Stephen McAleer**, Vlad Firoiu, Lei Zhang, Doina Precup, Shibl Mourad. Proving Theorems using Incremental Learning and Hindsight Experience Replay. *International Conference on Machine Learning (ICML) 2022*.
- 2022 Roy Fox, **Stephen McAleer**, Will Overman, Ioannis Panageas. Independent Natural Policy Gradient Always Converges in Markov Potential Games. *International Conference on Artificial Intelligence and Statistics (AISTATS) 2022*.
- 2021 **Stephen McAleer**, John Lanier, Pierre Baldi, Roy Fox. XDO: A Double Oracle Algorithm for Extensive Form Games. *Neural Information Processing Systems (NeurIPS) 2021*
- 2021 Xidong Feng, Oliver Slumbers, Yaodong Yang, Ziyu Wan, Bo Liu, **Stephen McAleer**, Ying Wen, Jun Wang. Discovering Multi-Agent Auto-Curricula in Two-Player Zero-Sum Games. *Neural Information Processing Systems (NeurIPS) 2021*
- 2020 **Stephen McAleer***, John Lanier*, Roy Fox, Pierre Baldi. Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games. *Neural Information Processing Systems (NeurIPS) 2020*.

2020 Shauharda Khadka, Somdeb Majumdar, Santiago Miret, **Stephen McAleer**, Kagan Tumer. Evolutionary Reinforcement Learning for Sample-Efficient Multiagent Coordination. *International Conference on Machine Learning (ICML) 2020*.

2019 **Stephen McAleer***, Forest Agostinelli*, Alexander Shmakov*, Pierre Baldi. Solving the Rubik's Cube with Approximate Policy Iteration. *International Conference on Learned Representations (ICLR) 2019*.

Workshop Publications and Arxiv Preprints

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

2021 Forest Agostinelli, **Stephen McAleer**, Alexander Shmakov, Roy Fox, Marco Valtorta, Biplav Srivastava, Pierre Baldi. Obtaining Approximately Admissible Heuristic Functions through Deep Reinforcement Learning and A* Search.

2021 Forest Agostinelli, Alexander Shmakov, **Stephen McAleer**, Roy Fox, Pierre Baldi. A* Search Without Expansions: Learning Heuristic Functions with Deep Q-Networks.

2020 Astrid Anker, Pierre Baldi, Steven Barwick, et. al. White Paper: ARIANNA-200 high energy neutrino telescope.

2019 Alexander Shmakov, John Lanier, **Stephen McAleer**, Rohan Achar, Cristina Lopes, Pierre Baldi. ColosseumRL: A Framework for Multiagent Reinforcement Learning in N-Player Games. *AAAI Spring Symposium Series: Challenges and Opportunities for Multi-Agent Reinforcement Learning 2019*.

2019 John Lanier, **Stephen McAleer**, Pierre Baldi. Curiosity-Driven Multi-Criteria Hind-sight Experience Replay. *NeurIPS 2019 Deep Reinforcement Learning Workshop*.

Awards and Honors

2022 **NSF Computing Innovation Fellow**.

2019 **UC Irvine Nominee for the Google PhD Fellowship**.

2019 **Machine Learning and Physical Sciences (MAPS) NSF Fellow**.

2017 **Bidstrup Undergraduate Fellow**.

2015 **Kakehashi Fulbright Fellow**.

2013 **National Merit Scholar**.

Invited Talks

2022 Tsinghua University: *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

2022 Workshop on Algorithms for Learning and Economics (WALE): *Deep RL for Large Games*

2021 DeepMind: *XDO: A Double Oracle Algorithm for Extensive-Form Games*

2021 ASU Design Informatics Lab: *RL for Game Theory and Game Theory for RL*

2020 Facebook AI Research: *Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games*

- 2020 DeepMind: *Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games*
- 2019 Intel AI: *Solving the Rubik's Cube Without Human Knowledge*
- 2019 Cylance Inc.: *Solving the Rubik's Cube Without Human Knowledge*
- 2018 UCI AI/ML Seminar: *Solving the Rubik's Cube Without Human Knowledge*

Guest Lectures

- 2020 CS 295: Reinforcement Learning, UCI: *Multiagent Reinforcement Learning*
- 2019 PSYCH 293: Neural Networks and Reinforcement Learning, UCI: *Trust Region Policy Optimization*
- 2019 CS 274C: Neural Networks and Deep Learning, UCI: *Reinforcement Learning*
- 2018 CS 274C: Neural Networks and Deep Learning, UCI: *GANs in the Context of Reinforcement Learning*

Leadership

- 2018–2020 **Coordinator**, *Data Science Initiative, UCI, Irvine, CA.*
Coordinate and develop data science training workshops
- 2017–2020 **Organizer**, *Machine Learning Reading Group, UCI, Irvine, CA.*
Lead weekly discussion of machine learning papers for machine learning PhD students
- 2015–2016 **Portfolio Manager**, *ASU Student Investment Management Fund, Tempe, AZ.*
Lead team of 15 students through investment program which oversees 750,000 dollars
- 2013–2016 **President**, *Investment Banking Industry Scholars, Tempe, AZ.*
Lead group of 16 students through highly selective investment banking training program.
- 2016 **Founder**, *Investment Banking and Consulting Club, Tempe, AZ.*
Founded club which provides professional advice and currently has over 200 members.

Mentoring and Advising

Masters Research

J.B. Lanier (PhD at UC Irvine)

Undergraduate Research

Kevin Wang (PhD at Brown)

Alexander Shmakov (PhD at UC Irvine)

Siqi Tang (Masters at Johns Hopkins)

Outreach

- Summer 2019 **Organizer and Instructor**, *UCI Artificial Intelligence Summer Institute, Irvine, CA.*
Taught machine learning classes to high school and college students from China
- Summer 2018 **Organizer and Instructor**, *UCI Artificial Intelligence Summer Institute, Irvine, CA.*
Taught machine learning classes to high school and college students from China

Professional Activities

- 2022 **Reviewer**, *Neural Information Processing Systems (NeurIPS)*.
- 2022 **Reviewer**, *International Conference on Machine Learning (ICML)*.
- 2021 **Reviewer**, *International Conference on Learning Representations (ICLR)*.
- 2021 **Reviewer**, *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- 2021 **Reviewer**, *Web and Internet Economics (WINE)*.
- 2021 **Reviewer**, *Neural Information Processing Systems (NeurIPS)*.
- 2021 **Reviewer**, *Nature Machine Intelligence*.
- 2020 **Program Committee Member**, *AAAI Reinforcement Learning in Games Workshop*.
- 2019 **Reviewer**, *Science Advances*.
- 2019 **Reviewer**, *International Journal of Information Technology*.
- 2019 **Reviewer**, *IEEE Transactions on Industrial Informatics*.
- 2019 **Reviewer**, *Artificial Intelligence*.
- 2018 **Reviewer**, *Neural Networks*.

Selected Press Coverage

- 2019 **AI Solves Rubik's Cube in One Second**, *BBC*.
- 2019 **How quickly can AI solve a Rubik's Cube? In less time than it took you to read this headline.**, *Washington Post*.
- 2019 **This AI Can Solve a Rubik's Cube Super Fast**, *Forbes*.
- 2019 **One more thing artificial intelligence can beat you at: Solving a Rubik's cube**, *Popular Science*.
- 2019 **Rubik's Cube Solved in 'Fraction of a Second' by Artificial Intelligence Machine Learning Algorithm**, *Newsweek*.
- 2019 **AI teaches itself to complete the Rubik's cube in just 20 MOVES - and gets it spot on every single time**, *Daily Mail*.
- 2019 **AI solves Rubik's Cube in 1.2 seconds**, *The Register*.
- 2019 **Self-Taught AI Masters Rubik's Cube Without Human Help**, *Gizmodo*.
- 2019 **Rubik's Cube solved by deep learning algorithm in fraction of a second**, *Fox News*.
- 2019 **AI solves Rubik's cube in under a second**, *Mashable*.
- 2019 **This AI Solves the Rubik's Cube Way Better Than You**, *Discover Magazine*.
- 2018 **A machine has figured out Rubik's Cube all by itself**, *MIT Technology Review*.
- 2018 **A machine taught itself to solve Rubik's Cube without human help, UC Irvine researchers say**, *LA Times*.
- 2018 **Machine Learning Finally Tackles the Rubik's Cube**, *Popular Mechanics*.
- 2018 **Self-Taught AI Masters Rubik's Cube in Just 44 Hours**, *Gizmodo*.