Stephen McAleer

Professional Appointment

2022–Present **Postdoctoral Fellow**, *Carnegie Mellon University*, Pittsburgh, PA.
Advisor: Tuomas Sandholm

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

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

2017–2021 **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

Teaching

Fall 2023 **Instructor**, *Carnegie Mellon University*, Computer Science 15-888: Computational Game Solving.

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.

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

Papers

Machine Learning Journal Publications

Science 2022 Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T. Connor, Neil Burch, Thomas Anthony, **Stephen McAleer**, Romuald Elie, Sarah H. Cen, Zhe Wang, Audrunas Gruslys, Aleksandra Malysheva, Mina Khan, Sherjil Ozair, Finbarr Timbers, Toby Pohlen, Tom Eccles, Mark Rowland, Marc Lanctot, Jean-Baptiste Lespiau, Bilal Piot, Shayegan Omidshafiei, Edward Lockhart, Laurent Sifre, Nathalie Beauguerlange, Remi Munos, David Silver, Satinder Singh, Demis Hassabis, Karl Tuyls. *Mastering the Game of Stratego with Model-Free Multiagent Reinforcement Learning.*

TMLR 2021 Le Cong Dinh, Yaodong Yang, **Stephen McAleer**, Nicolas Perez Nieves, Oliver Slumbers, Zheng Tian, David Henry Mguni, Haitham Bou Ammar, Jun Wang. *Online Double Oracle.*

Nature MI 2019 Forest Agostinelli*, **Stephen McAleer***, Alexander Shmakov*, Pierre Baldi. *Solving the Rubik's Cube with Deep Reinforcement Learning and Search.*.

Machine Learning Conference Publications

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

ICLR 2024 Ted Moskovitz, Aaditya K Singh, DJ Strouse, Tuomas Sandholm, Ruslan Salakhutdinov, Anca Dragan, **Stephen McAleer**. *Confronting Reward Model Overoptimization with Constrained RLHF.* ***(Spotlight)***

ICLR 2024 Yongyuan Liang, Yanchao Sun, Ruijie Zheng, Xiangyu Liu, Benjamin Eysenbach, Tuomas Sandholm, Furong Huang, **Stephen McAleer**. *Game-Theoretic Robust Reinforcement Learning Handles Temporally-Coupled Perturbations.*

ICLR 2024 Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, **Stephen McAleer**, Albert Q. Jiang, Jia Deng, Stella Biderman, Sean Welleck. *Llemma: An Open Language Model for Mathematics.*

ICLR 2024 Tim Franzmeyer, **Stephen McAleer**, Joao F. Henriques, Jakob Nicolaus Foerster, Philip Torr, Adel Bibi, Christian Schroeder de Witt. *Illusory Attacks: Detectability Matters in Adversarial Attacks on Sequential Decision-Makers.* ***(Spotlight)***

AAAI 2024 Michael Curry, Vinzenz Thoma, Darshan Chakrabarti, **Stephen McAleer**, Christian Kroer, Tuomas Sandholm, Niao He, Sven Seuken. *Automated Design of Affine Maximizer Mechanisms in Dynamic Settings.*

AAMAS 2024 Pengdeng Li, Shuxin Li, Xinrun Wang, Jakub Cerny, Youzhi Zhang, **Stephen McAleer**, Hau Chan, Bo An. *Grasper: A Generalist Pursuer for Pursuit-Evasion Problems.*

NeurIPS 2023 Geunwoo Kim, Pierre Baldi, **Stephen McAleer**. *Language Models can Solve Computer Tasks.*

NeurIPS 2023 **Stephen McAleer**, Gabriele Farina, Gaoyue Zhou, Mingzhi Wang, Yaodong Yang, Tuomas Sandholm. *Team-PSRO for Learning Approximate TMECor in Large Team Games via Cooperative Reinforcement Learning.*

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

NeurIPS 2023 Jian Yao, Weiming Liu, Haobo Fu, Yaodong Yang, **Stephen McAleer**, Qiang Fu, Wei Yang. *Policy Space Diversity for Non-Transitive Games.*

NeurIPS 2023 Brian Hu Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, **Stephen McAleer**, Andreas Alexander Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, Tuomas Sandholm. *Computing Optimal Equilibria and Mechanisms via Learning in Zero-Sum Extensive-Form Games.*

EC 2023 Ioannis Anagnostides, Fivos Kalogiannis, Ioannis Panageas, Emmanouil-Vasileios Vlatakis-Gkaragkounis, **Stephen McAleer**. *Algorithms and Complexity for Computing Nash Equilibria in Adversarial Team Games.*

ICML 2023 Xiaohang Tang, Le Cong Dinh, **Stephen McAleer**, Yaodong Yang. *Regret-Minimizing Double Oracle for Extensive-Form Games.*

ICML 2023 David Henry Mguni, Taher Jafferjee, Haojun Chen, Jianhong Wang, Long Fei, Xidong Feng, **Stephen McAleer**, Feifei Tong, Jun Wang, Yaodong Yang. *MANSA: Learning Fast and Slow in Multi-Agent Systems.*

ICML 2023 Oliver Slumbers, David Henry Mguni, **Stephen McAleer**, Stefano B Blumberg, Yaodong Yang, Jun Wang. *A Game-Theoretic Framework for Managing Risk in Multi-Agent Systems.*

ICLR 2023 **Stephen McAleer**, Gabriele Farina, Marc Lanctot, Tuomas Sandholm. *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret.*

NeurIPS 2022 Yuanpei Chen, Yaodong Yang, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuang Jiang, **Stephen McAleer**, Hao Dong, Zongqing Lu, Song-Chun Zhu. *Towards Human-Level Bimanual Dexterous Manipulation with Reinforcement Learning.*

ICML 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.*

ICML 2022 Eser Aygün, 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.*

AISTATS 2022 Roy Fox, **Stephen McAleer**, Will Overman, Ioannis Panageas. *Independent Natural Policy Gradient Always Converges in Markov Potential Games.*

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

NeurIPS 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.*

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

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

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

ICLR 2019 **Stephen McAleer***, Forest Agostinelli*, Alexander Shmakov*, Pierre Baldi. *Solving the Rubik's Cube with Approximate Policy Iteration.*

Applications in Science

2022 Christian Glaser, **Stephen McAleer**, Sigfrid Stjärnholm, Pierre Baldi, Steven W Barwick. Deep-Learning-Based Reconstruction of the Neutrino Direction and Energy for In-Ice Radio Detectors. *Astroparticle Physics.*

2022 ARIANNA Collaboration. Triboelectric Backgrounds to Radio-Based UHE Neutrino Exeperiments. *Astroparticle Physics.*

2022 ARIANNA Collaboration. Measuring the Polarization Reconstruction Resolution of the Arianna Neutrino Detector With Cosmic Rays. *Journal of Cosmology and Astroparticle Physics.*

2022 ARIANNA Collaboration. Improving sensitivity of the ARIANNA detector by rejecting thermal noise with deep learning. *Journal of Instrumentation.*

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*.

2021 ARIANNA Collaboration. Polarization Reconstruction of Cosmic Rays with the ARIANNA Neutrino Radio Detector. *37th International Cosmic Ray Conference.*

2021 ARIANNA Collaboration. TAROGE-M: Radio Observatory on Antarctic High Mountain for Detecting Near-Horizon Ultra-High Energy Air Showers. *37th International Cosmic Ray Conference.*

2021 Christian Glaser, **Stephen McAleer**, Pierre Baldi, Steven W Barwick. Deep Learning Reconstruction of the Neutrino Energy with a Shallow Askaryan Detector. *37th International Cosmic Ray Conference.*

2021 ARIANNA Collaboration. A Novel Trigger Based on Neural Networks for Radio Neutrino Detectors. *37th International Cosmic Ray Conference.*

2021 ARIANNA Collaboration. Capabilities of ARIANNA: Neutrino Pointing Resolution and Implications for Future Ultra-high Energy Neutrino Astronomy. *37th International Cosmic Ray Conference.*

2021 ARIANNA Collaboration. Development of an In-Situ Calibration Device of Firn Properties for Askaryan Neutrino Detectors. *37th International Cosmic Ray Conference.*

2021 ARIANNA Collaboration. Science Case and Detector Concept for ARIANNA High Energy Neutrino Telescope at Moore's Bay, Antarctica. *37th International Cosmic Ray Conference.*

2020 ARIANNA Collaboration. White Paper: ARIANNA-200 High Energy Neutrino Telescope.

2018 Siyu Shao, **Stephen McAleer**, Ruqiang Yan, Pierre Baldi. Highly-Accurate Machine Fault Diagnosis Using Deep Transfer Learning. *IEEE Transactions on Industrial Informatics.*

Workshop Publications and Arxiv Preprints

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

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

2022 **Stephen McAleer**, JB Lanier, Kevin Wang, Pierre Baldi, Roy Fox, Tuomas Sandholm. Toward Optimal Policy Population Growth in Two-Player Zero-Sum Games.

2021 Yaosheng Xu, Dailin Hu, Litian Liang, **Stephen McAleer**, Pieter Abbeel, Roy Fox. Target Entropy Annealing for Discrete Soft Actor-Critic. *NeurIPS 2021 Deep Reinforcement Learning Workshop*.

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. *Bridging the Gap Between AI Planning and Reinforcement Learning Workshop at ICAPS*

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 Hindsight 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

2023 NeurIPS Multi-Agent Security Workshop: *Panel Discussion*

2023 MIT Computational Sensorimotor Learning Seminar: *Toward General Virtual Agents*

2023 Harvard EconCS Seminar: *Game-Theoretic Reinforcement Learning*

2023 Boston University Machine Learning Symposium: *Toward General Virtual Agents*

2023 Stanford (Dorsa Sadigh's Group): *Toward General Virtual Agents*

2023 UC Berkeley (Pieter Abbeel's Group): *Toward General Virtual Agents*

2023 Tsinghua University (Jie Tang's Group): *Language Models can Solve Computer Tasks*

2023 SAP Innovation Sessions: *Language Models can Solve Computer Tasks*

2022 NeurIPS Deep RL Workshop Spotlight: *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

2022 University of Edinburgh (Stefano Albrecht's Group): *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

2022 Brown University (Amy Greenwald's Group): *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

2022 UC Berkeley MARL Seminar: *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

2022 University of Maryland (Furong Huang's Group): *ESCHER: Eschewing Importance Sampling in Games by Computing a History Value Function to Estimate Regret*

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*

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

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

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

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

Summer 2017 **International Lead**, *Harvard University*, Adama, Ethiopia.
Organized and led an effort to implement CS50 throughout universities in Ethiopia

Service

2023 **Organizer**, *Foundation Models for Decision Making workshop at NeurIPS*.

2018-Present **Reviewer**, *ICML, NeurIPS, ICLR, and other conferences*.

Selected Press Coverage

2023 **Meet LLEMMA, the math-focused open source AI that outperforms rivals.**, *VentureBeat*.

2022 **Here's how a new AI mastered the tricky game of Stratego**, *Popular Science*.

2022 **Now AI can outmaneuver you at both Stratego and Diplomacy**, *TechCrunch*.

2022 **DeepMind's New AI Uses Game Theory to Trounce Humans in 'Stratego'**, *Gizmodo*.

2022 **DeepMind AI topples experts at complex game Stratego**, *Nature*.

2022 **DeepMind AI uses deception to beat human players in war game Stratego**, *New Scientist*.

2022 **Two new AI systems beat humans at complex games**, *Axios*.

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 **AI solves the Rubik's cube**, *Nature*.

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 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*.

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/

2018 **Machine Learning Finally Tackles the Rubik's Cube**, *Popular Mechanics*.
2018 **Self-Taught AI Masters Rubik's Cube in Just 44 Hours**, *Gizmodo*.

smcaleer@cs.cmu.edu
https://www.andrew.cmu.edu/user/smcaleer/