Steven Morad

Reinforcement Learning · Deep Sequence Modeling · Robotics

Right to work in the USA, EU, and UK

steven.morad@gmail.com | www.dangersteve.com | scholar | smorad

Education

University of Cambridge *Cambridge, UK*

PhD Computer Science *2021 - 2024*

- Advisor: Prof. Amanda Prorok
- Thesis: Deep Memory Models and Efficient Reinforcement Learning under Partial Observability

University of Arizona *Tucson, AZ, USA*

MS Aerospace Engineering *2017 - 2019*

- Advisor: Prof. Jekan Thanga
- Thesis: The Spinning Projectile Extreme Environment Robot

University of California, Santa Cruz *Santa Cruz, CA, USA*

BS Honors Computer Science *2011 - 2015*

Professional Experience

University of Macau

Assistant Professor *2024-*

Research track, joint appointments in Computer and Information Science and the Centre for Artificial Intelligence and Robotics

University of Cambridge

Adjunct Lecturer *2024*

Instructor for R255: Deep Reinforcement Learning

Anyscale

Contract Software Engineer *2022-2023*

Development of RLlib, an open source reinforcement learning library

Toshiba Research Europe

Graduate Research Intern *2020*

Vision-only navigation for drones and ground vehicles using RL

NASA Jet Propulsion Lab

Robotics Research Intern *2018, 2019*

Visual inertial odometry and robot dynamics for icy moons (Enceladus and Europa)

Meta *2015-2017*

Production Engineer

Software development and systems engineering for Meta's distributed compute engine

Selected Publications and Patents

Morad, S.D., Lu, C., Kortvelesy, R., Liwicki, S., Foerster, J., Prorok, A. (2025) *Recurrent Reinforcement Learning with Memoroids.* **Neural Information Processing Systems**.

Blumenkamp, J., **Morad, S.D.**, Gielis, J., Prorok, A. (2025) *CoViS-Net: A Cooperative Visual Spatial Foundation Model for Multi-Robot Applications.* **Conference on Robot Learning.**

Morad, S.D., Mecca R., Poudel, R., Liwicki, S., Cipolla, R. (2024). *Task Performing Agent Systems and Methods.* **US Patent US12085947B2, UK Patent GB2598758B, Japan Patent JP7225292B2**

Morad, S.D., Kortvelesy, R., Liwicki, S., Prorok, A. (2023) *Reinforcement Learning with Fast and Forgetful Memory.* **Neural Information Processing Systems**.

Kortvelesy, R., **Morad, S.D.**, Prorok, A. (2023) *Generalised f-Mean Aggregation for Graph Neural Networks.* **Neural Information Processing Systems**.

Morad, S.D., Kortvelesy, R., Liwicki, S., Prorok, A. (2023) *POPGym: Benchmarking Partially Observable Reinforcement Learning.* **The International Conference on Learning Representations**.

Kortvelesy, R., **Morad, S.D.**, Prorok, A. (2023) *Permutation-Invariant Set Autoencoders with Fixed-Size Embeddings for Multi-Agent Learning.* **The International Conference on Autonomous Agents and Multiagent Systems**.

Blumenkamp, J., **Morad, S.D.**, Gielis, J., Li, Q., and Prorok, A. (2022) *A Framework for Real-world Multi-robot Systems Running Decentralized GNN-based Policies.* ***The International Conference on Robotics and Automation.***

Morad, S.D., Liwicki, S., Kortvelesy, R., Mecca, R., Prorok, A. (2022). *Modeling Partially Observable Systems using Graph-Based Memory and Topological Priors.* **Learning for Dynamics and Control**.

Morad, S.D., Mecca, R., Poudel, R., Liwicki, S., Cipolla, R. (2020). *Embodied Visual Navigation with Automatic Curriculum Learning in Real Environments.* **Robotics and Automation Letters**.

Morad, S.D., Nash, J., Higa, S., Smith, R., Parness, A., and Barnard, K. (2019). *Improving Visual Feature Extraction in Glacial Environments.* **Robotics and Automation Letters**.

Morad, S.D., Dailey, T., Vance, L.D., and Thangavelautham, J. (2019). *A Spring Propelled Extreme Environment Robot for Off-world Cave Exploration.* **IEEE Aerospace Conference**.

Awards, Fellowships, & Grants

2025	**SRG**, University of Macau	*MOP 150,000*
2023	**Postgraduate Research Fund**, Jesus College	*£ 650*
2023	**2nd Best Research Talk**, Jesus College Graduate Conference, University of Cambridge	-
2021	**Graduate Research Studentship**, Toshiba Research	*£ 149,953*
2015	**Cum Laude**, University of California, Santa Cruz	-

Invited Talks

2025	**Reinforcement Learning with Memory**, R271 Guest Lecture - University of Cambridge	*Cambridge*
2024	**Value-Based Reinforcement Learning**, Cambridge Ellis Unit Summer School on Probabilistic Machine Learning	*Cambridge*
2024	**Reinforcement Learning under Partial Observability**, University of Macau	*Macau*
2023	**An Introduction to Reinforcement Learning**, Toshiba Research Seminar	*Cambridge*

Teaching Experience

2025	**Special Topics in AI: Decision Making**, Lecturer	*Macau*
2024	**Introduction to Deep Learning**, Lecturer	*Macau*
2024	**Advanced Topics: Deep Reinforcement Learning**, Lecturer	*Cambridge*
2023	**Programming in C/C++**, Supervisor	*Jesus College*
2022	**Introduction to Robotics**, Teaching Assistant	*Cambridge*
2021	**Mobile Robot Systems**, Teaching Assistant	*Cambridge*

Thesis Supervision

2024-	**Wang Zekang**, MS	*Macau*
2024-	**He Zhe**, MS	*Macau*
2022-2023	**Dulhan Jayalath**, MPhil	*Cambridge*
2023-2024	**Mark Li**, Part II	*Cambridge*
2021-2022	**James Read**, Part II	*Cambridge*

Outreach & Professional Development

Service and Outreach

2023 **UC Santa Cruz**, Mentor *Remote*
2018 **Boys and Girls Club**, Volunteer *Tucson*

Peer Review

NeurIPS
ICLR
ICML
AAAI
ICRA/RA-L
RSS

Tools and Frameworks

Python, PyTorch, JAX, Numpy, C/C++, SolidWorks, Linux, Machine Learning, ROS, ROS2, MATLAB, Git

Project Media

(a) (NASA/JPL internships) Evaluating mobility and vision in icy environments

(b) (MS Thesis) Sensor node using two-axis spin stabilization (precession) for rocket motor exhaust vectoring, achieving soft touchdown at 10cm/s after an 8.7m drop. The system was designed for Lunar lava tube descent and mapping.

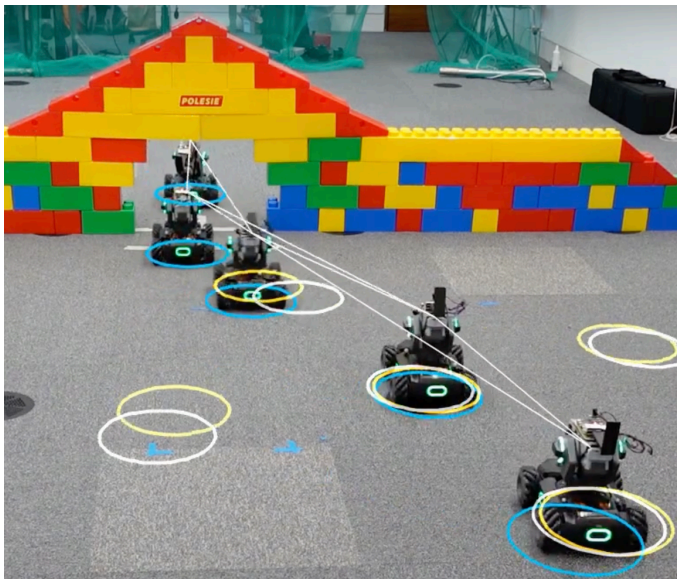

Figure 2: (PhD) A fully distributed multirobot system, running policies trained using reinforcement learning, on the robot computers in real-time. The goal is for the agents to negotiate passage through a chokepoint via local communication, without explicit guidance. White lines denote the dynamic communication topology, white circles denote goals, blue circles denote the real-world agent state, and yellow circles denote the agent state when running the same episode in simulation. The difference between whie and yellow circles visualizes the simulation to reality gap.