Steven Morad

Reinforcement Learning · Sequence Modeling · Robotics

Cambridge, UK · 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 (Est.)*

- Advisor: Prof. Amanda Prorok
- Thesis: Efficient Reinforcement Learning in Partially Observable Domains

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

2022-2023	**Anyscale (RLlib)**, Part-Time Software Engineer
2020	**Toshiba Research Europe**, Graduate Research Intern
2018-2019	**NASA Jet Propulsion Lab**, Robotics Research Intern
2017-2019	**Space and Terrestrial Robotic Exploration Lab**, Research Assistant
2015-2017	**Meta**, Production (Software/Systems) Engineer

Selected Publications

Published

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**.

Under Review

Jayalath, D., **Morad, S.D.**, and Prorok, A. (2024) *Generalising Multi-Agent Cooperation Through Task-Agnostic Communication.* **The International Conference on Learning Representations**.

Awards, Fellowships, & Grants

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

2023 **An Introduction to Reinforcement Learning**, Toshiba Research Seminar

Teaching Experience

2024 Est.	**Advanced Topics: Deep Reinforcement Learning**, Adjunct 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

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
ICRA/RA-L

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.

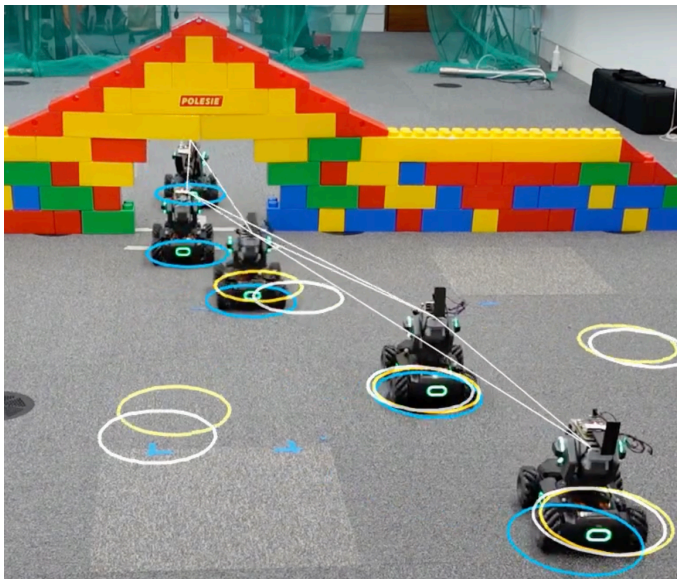

Figure 2: (PhD) A fully distributed multirobot system, running policies trained using reinforcement learning, on the robot computers in real-time. The goal is for the agents to negotiate passage through a chokepoint via local communication, without explicit guidance. White lines denote the dynamic communication 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.