

Aditya Kapoor

ELLIS PhD Student | University of Manchester

aditya.kapoor@postgrad.manchester.ac.uk

Website | Google Scholar | GitHub

I study how learning agents assign credit, communicate and coordinate. My research spans multi-agent reinforcement learning, communication-efficient learning and language-model-based planning, with emerging work in probabilistic generative modelling.

Education & supervision

University of Manchester | PhD in Computer Science, ELLIS programme
November 2024 - December 2028 (expected)

Supervisors: Omar Rivasplata (primary); Samuel Kaski (secondary); Jan Peters (secondary and ELLIS co-supervisor).

BITS Pilani, Goa Campus | B.E. (Hons.), Computer Science and Information Systems
August 2017 - August 2021 | CGPA: 8.43/10

Selected publications

Assigning Credit with Partial Reward Decoupling in Multi-Agent Proximal Policy Optimization

Aditya Kapoor, Benjamin Freed, Howie Choset, Jeff Schneider

RLC / Reinforcement Learning Journal, 2024 | First author

Learns which teammates are relevant to an agent's updates to improve credit assignment in cooperative MARL.

Learning Cooperative Multi-Agent Policies with Partial Reward Decoupling

Benjamin Freed*, Aditya Kapoor*, Ian Abraham, Jeff Schneider, Howie Choset

IEEE Robotics and Automation Letters, 2022; presented at ICRA 2022 | Equal contribution

Uses learned relevance between agents and rewards to reduce interference in cooperative policy learning.

Long-Horizon Planning for Multi-Agent Robots in Partially Observable Environments

Siddharth Nayak et al. (including Aditya Kapoor)

NeurIPS 2024, main conference | Co-author

LLaMAR: language-model-based multi-agent planning with a plan-act-correct-verify loop and evaluation under partial observability.

Selected preprint

Learning what to say and how precisely: Efficient Communication via Differentiable Discrete Communication Learning

Aditya Kapoor, Yash Bhisikar, Benjamin Freed, Jan Peters, Mingfei Sun

arXiv preprint, 2025 | First author

Studies end-to-end learning of discrete messages and their precision for communication-constrained multi-agent learning.

Research in progress

Communication precision in multi-agent reinforcement learning

Investigating stochastic quantisation and dithering for discrete inter-agent messages. Controlled comparisons examine message precision, communication cost and coordination performance against fixed-precision baselines.

Measuring contributions across agents and time

Developing controlled diagnostic studies of credit assignment in cooperative MARL, distinguishing predictive contribution scores from evidence about how interventions change team outcomes.

Emerging direction: probabilistic generative modelling

Studying structured and randomised decoding schedules for frozen masked language models, separating schedule structure from the effect of learning mixture weights at matched model-call budgets.

Research experience

Tata Consultancy Services, Research & Innovation | Predoctoral Research Fellow

October 2021 - July 2024

Research on sparse-reward multi-agent credit assignment and autonomous composition of vision pipelines for robotic perception. Worked with Mayank Baranwal, Harshad Khadilkar and colleagues.

Carnegie Mellon University, Biorobotics Lab | Research collaboration

May 2020 - May 2022

Cooperative MARL and partial reward decoupling, with Howie Choset and Benjamin Freed.

Aston University / RoboComp | Research collaboration

April 2020 - November 2020

Reinforcement learning for social navigation, with Luis Manso and Pilar Bachiller.

Additional relevant publications

SocNavGym: A Reinforcement Learning Gym for Social Navigation

IEEE RO-MAN, 2023 | Equal contribution

Low-Rank Agent-Specific Adaptation (LoRASA) for Multi-Agent Policy Learning

DAI 2025 | Co-author

Redistributing Rewards Across Time and Agents for Multi-Agent Reinforcement Learning

arXiv preprint, 2025 | First author

Technical skills & service

Programming and tools: Python, PyTorch, C/C++, Git, Docker, LaTeX; robotics simulation with PyBullet and Webots.

Open source and mentoring: Google Summer of Code mentor, RoboComp (2022); multi-agent RL contributor/lead, GenRL (2020-2021).

Selected reviewing: NeurIPS 2024; RLC 2024-2025; IJCAI 2023-2024; ICRA 2022.