Aditya Kapoor

ELLIS PhD Student · Multi-agent learning & probabilistic generative modelling
aditya.kapoor@postgrad.manchester.ac.uk | Website | Google Scholar | GitHub

Education

University of Manchester Nov 2024 - present
PhD in Computer Science · ELLIS programme · Expected completion: December 2028
Supervisors: Omar Rivasplata, Samuel Kaski and Jan Peters.

BITS Pilani, Goa Aug 2017 - Aug 2021
B.E. (Hons.), Computer Science and Information Systems · CGPA 8.43/10

Research Experience

University of Manchester Nov 2024 - present
PhD research

- Introduced Differentiable Discrete Communication Learning (DDCL) for learning discrete messages with shared dithering and a differentiable bit-cost penalty at fixed grid spacing; developed the method and implementation.
- Developed TAR^2 for reward redistribution across agents and time, combining method design, theoretical analysis and evaluation on SMACLite and Google Research Football.

TCS Research & Innovation Oct 2021 - Jul 2024
Predoctoral Research Fellow
Mentors: Harshad Khadilkar, Mayank Baranwal and Jayavardhana Gubbi.

- Developed reward-redistribution methods for sparse-feedback multi-agent reinforcement learning, alongside robotic perception and automatic algorithm selection for retail environments.
- Co-led research development and manuscript writing for LLaMAR with Siddharth Nayak (NeurIPS 2024). Its feedback-driven multi-robot planner achieved 66% task success versus 34% for ReAct on two-agent MAP-THOR with GPT-4V.

Carnegie Mellon University, Biorobotics Lab May 2020 - May 2022
Research collaboration / undergraduate thesis
With Benjamin Freed, Howie Choset and Jeff Schneider.

- Undergraduate thesis on credit assignment in cooperative multi-agent reinforcement learning. Co-developed PRD and led its implementation and evaluation; published in IEEE RA-L (2022) and presented at ICRA 2022.
- Developed PRD-MAPPO with soft relevance weighting and support for shared team rewards (RLC 2024). Reduced critic calls for advantage estimation from quadratic to linear in team size.

Aston University / RoboComp Apr - Nov 2020
Research collaboration
Co-developed a social-navigation dataset toolkit with Luis Manso and Pilar Bachiller.

BITS Pilani, Goa 2018 - 2019
Spiking neural networks research project
Worked on spiking neural networks and speech synthesis on SpiNNaker with Basabdatta Sen Bhattacharya. Also implemented the anti-Hebbian learning rule in collaboration with the Human Brain Project.

Research Interests

Learning and coordination in multi-agent systems
Learning communication precision, assigning agent-level credit that supports the team objective, and coordinating simultaneous policy updates through measured interactions.

Probabilistic and generative modelling
Studying decoding-order persistence in masked language models, reward-guided diffusion on curved spaces, and dithered corruption for discrete generative models.

Selected Publications

Assigning Credit with Partial Reward Decoupling in Multi-Agent Proximal Policy Optimization
A. Kapoor, B. Freed, H. Choset, J. Schneider
RLC / Reinforcement Learning Journal 2024 · Paper · Code

Long-Horizon Planning for Multi-Agent Robots in Partially Observable Environments
S. Nayak, A. M. Orozco, M. ten Have, V. Thirumalai, J. Zhang, D. Chen, **A. Kapoor**, E. Robinson, K. Gopalakrishnan, J. Harrison, B. Ichter, A. Mahajan, H. Balakrishnan
NeurIPS 2024 · Paper · Code

Learning Cooperative Multi-Agent Policies with Partial Reward Decoupling
B. Freed, **A. Kapoor**, I. Abraham, J. Schneider, H. Choset
IEEE RA-L 2022; presented at ICRA 2022 · Equal contribution · Paper · Code

Preprints - Under Review

Redistributing Rewards Across Time and Agents for Multi-Agent Reinforcement Learning
A. Kapoor, K. Tessera, H. Khadilkar, M. Baranwal, J. Peters, S. Albrecht, M. Sun
TAR² · arXiv 2025 · Paper · Code

Learning what to say and how precisely: Efficient Communication via Differentiable Discrete Communication Learning
A. Kapoor, Y. Bhisikar, B. Freed, J. Peters, M. Sun
DDCL · arXiv 2025 · Paper · Code

Additional Publications

Low-Rank Agent-Specific Adaptation (LoRASA) for Multi-Agent Policy Learning
B. Zhang, **A. Kapoor**, M. Sun · **DAI 2025** · Equal contribution · Paper · Code

SocNavGym: A Reinforcement Learning Gym for Social Navigation
A. Kapoor, S. Swamy, P. Bachiller, L. Manso · **RO-MAN 2023** · Equal contribution · Paper · Code

Open Source & Mentoring

Google Summer of Code / RoboComp: Contributor (2021): Improvement of a Social Navigation component. Mentor (2022): Extending the Social Navigation component to realistic scenarios; proposed and guided the project that became SocNavGym (RO-MAN 2023).

CORD.ai: Founding member; mentored students in research design, implementation and manuscript development, including projects that led to coauthored papers.

Technical Skills & Service

Tools: Python, PyTorch, C/C++, Git, Docker, LaTeX; PyBullet, Webots, SpiNNaker.

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