An Investigation of Reward Machines for Robotics

January 4, 2026

Research Question

How useful are Reward Machines in typical robotics applications?

Background

Many tasks we might want an autonomous system to perform involve sequences of actions and dependencies we are already aware of. For instance we may know that a robot needs to pick up a box before it can take to the location it is supposed to be in. Traditional reinforcement learning [4] uses a simple reward function (for instance delivering a rewards once the box is in the right location) and the learning process may then use up considerable resource before it learns that in order of this to be achieved the intermediate step of picking up the box must be completed.

Reward machines [1, 2], allow a reward function to be structured, so this information is available and the learning process can be streamlined.

Approach

The purpose of this project would be to implement and evaluate reward machines on a typical robotic inspection or manufacturing task. If successful this could be extended to consider the issue or multi-agent reinforcement learning [3]. This will involve identifying a suitable use case to investigate the value of reward machines, possibly drawing on examples from the existing literature and serving to validate their results. A simulation of the use case will need to be implemented and reinforcement learning both using a standard reward function and a reward machine applied so that the performance of the two approaches can be compared.

Milestones

1. Implementation of a simulation of a robotic system task with multiple stages, suitable for the application of reward machines.
2. Implementation of a standard reinforcement learning system to learn to operate successfully in the chosen simulation.
3. Implementation of a reward machine based reinforcement system to learn to operation successfully in the chosen simulation.
4. Comparison of the performance of the two implementation

References

[1] Icarte, R. T., Klassen, T. Q., Valenzano, R. A., and McIlraith, S. A. Using reward machines for high-level task specification and decomposition in reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (2018), J. G. Dy and A. Krause, Eds., vol. 80 of *Proceedings of Machine Learning Research*, PMLR, pp. 2112–2121.

[2] Icarte, R. T., Klassen, T. Q., Valenzano, R. A., and McIlraith, S. A. Reward machines: Exploiting reward function structure in reinforcement learning. *J. Artif. Intell. Res. 73* (2022), 173–208.

[3] Neary, C., Xu, Z., Wu, B., and Topcu, U. Reward machines for cooperative multi-agent reinforcement learning. In *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021* (2021), F. Dignum, A. Lomuscio, U. Endriss, and A. Nowé, Eds., ACM, pp. 934–942.

[4] Sutton, R. S., and Barto, A. G. *Reinforcement learning - an introduction, 2nd Edition*. MIT Press, 2018.