Symbolic Task Inference in Deep Reinforcement Learning (Abstract Reprint)

Hosein Hasanbeig[1], Natasha Yogananda Jeppu[2], Alessandro Abate[2], Tom Melham[2], Daniel Kroening[3]

[1]Microsoft Research
[2]Department of Computer Science, University of Oxford
[3]Amazon

Abstract Reprint. This is an abstract reprint of the journal article by Hasanbeig, Jeppu, Abate, Melham, and Kroening (2024).

Abstract

This paper proposes DeepSynth, a method for effective training of deep reinforcement learning agents when the reward is sparse or non-Markovian, but at the same time progress towards the reward requires achieving an unknown sequence of high-level objectives. Our method employs a novel algorithm for synthesis of compact finite state automata to uncover this sequential structure automatically. We synthesise a human-interpretable automaton from trace data collected by exploring the environment. The state space of the environment is then enriched with the synthesised automaton, so that the generation of a control policy by deep reinforcement learning is guided by the discovered structure encoded in the automaton. The proposed approach is able to cope with both high-dimensional, low-level features and unknown sparse or non-Markovian rewards. We have evaluated DeepSynths performance in a set of experiments that includes the Atari game Montezumas Revenge, known to be challenging. Compared to approaches that rely solely on deep reinforcement learning, we obtain a reduction of two orders of magnitude in the iterations required for policy synthesis, and a significant improvement in scalability.

References

Hasanbeig, H.; Jeppu, N. Y.; Abate, A.; Melham, T.; and Kroening, D. 2024. Symbolic Task Inference in Deep Reinforcement Learning. *Journal of Artificial Intelligence Research*, 80: 1099–1137.