

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Action decision congruence between human and deep reinforcement learning agents during a coordinated action task

Permalink

<https://escholarship.org/uc/item/91b6b7ht>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Patil, Gaurav
Bagala, Phillip
Nalepka, Patrick
et al.

Publication Date

2023

Peer reviewed

Action decision congruence between human and deep reinforcement learning agents during a coordinated action task

Gaurav Patil

Macquarie University, Sydney, NSW, Australia

Phillip Bagala

Macquarie University, Sydney, NSW, Australia

Patrick Nalepka

Macquarie University, Sydney, NSW, Australia

Michael Richardson

Macquarie University, Sydney, NSW, Australia

Rachel W. Kallen

Macquarie University, Macquarie Park, NSW, Australia

Abstract

Deep reinforcement learning (DRL) is capable of training agents that exceed human-levels of performance in multi-agent tasks. However, the behaviors exhibited by these agents are not guaranteed to be human-like or human-compatible. This poses a problem if the goal is to design agents capable of collaborating with humans or augmenting human actions in cooperative or team-based tasks. Indeed, recommender systems designed to augment human decision-making need to not only recommend actions that align with the task goal, but which also maintain coordinative behaviors between agents. The current study simultaneously explored skill learning performance of human learners when working alongside different artificial agents (AAs) during a collaborative problem-solving task, as well as evaluated the effectiveness of the same AAs as action decision recommender systems to aid learning. The action decisions of the AAs were either modelled by a heuristic model based on human performance or by a deep neural network trained by reinforcement learning using self-play. In addition to evaluating skill learning performance, the current study also tested the congruence between the decisions of the AAs with actual decisions made by humans. Results demonstrate that the performance of humans was significantly worse when working alongside the DRL AA compared to the heuristic AA. Additionally, the action decisions participants made also showed less alignment with the recommendations made by the DRL AA.