

Play the Reinforcement Learning Agent

Plisnier, Helene; Fasano, Alessandro Antonio; Nowe, Ann

Publication date:
2021

Document Version:
Submitted manuscript

[Link to publication](#)

Citation for published version (APA):
Plisnier, H., Fasano, A. A., & Nowe, A. (2021). *Play the Reinforcement Learning Agent*. Paper presented at 33rd Benelux Conference on Artificial Intelligence and 30th Belgian-Dutch Conference on Machine Learning, Luxembourg.

Copyright

No part of this publication may be reproduced or transmitted in any form, without the prior written permission of the author(s) or other rights holders to whom publication rights have been transferred, unless permitted by a license attached to the publication (a Creative Commons license or other), or unless exceptions to copyright law apply.

Take down policy

If you believe that this document infringes your copyright or other rights, please contact openaccess@vub.be, with details of the nature of the infringement. We will investigate the claim and if justified, we will take the appropriate steps.

Play the Reinforcement Learning Agent

Hélène Plisnier, Alessandro Fasano, and Ann Nowé

Vrije Universiteit Brussels
helene.plisnier@vub.be
<https://youtu.be/CP4HPBzsmtU>

Abstract. In the past few decades, artificial intelligence has gained an increasing amount of interest from the general public. Accompanying this interest, comes expectations of how sophisticated AI methods and their abilities are, often without a proper understanding of how they actually work. This demonstration is meant to give non-expert participants an idea of the view an RL agent has of its environment. We invite a volunteer to take the place of a standard RL agent and try learning the task solely based on information that would be available in a typical RL setting. The purpose of this demonstration is to illustrate how unintuitive an RL agent’s perspective of its environments is from a human point of view, and hence how limited its understanding of the task it is learning is. By establishing this idea in non-experts minds, we hope to debunk certain inaccurate assumptions people may have about AI technologies, specifically RL in this case.

Keywords: Reinforcement Learning · Transparency · Volunteer-Driven Demonstration

1 Introduction

Reinforcement Learning [3] (RL) is an Artificial Intelligence method, in which an agent learns to perform a task from scratch by repetitively interacting with its environment. At each timestep, the agent chooses and executes one of the actions at its disposal based on some state information, then receives a reward or punishment and goes on to the next state. RL methods consist in a promising approach to make robots easier to deploy in human-populated spaces, such as industries, offices and homes.

However, for robots executing RL algorithms to be accepted and allowed in human spaces, and to comply to (current and future) AI Transparency [1] and Explainability regulations ¹ their decision-making processes must be made clear for human users [6, 4, 2, 5]. An important component of the RL process is the way RL agents “perceive” and “understand” their environment. That kind of information is often restricted to experts in RL, who are used to design RL algorithms and test them on different environments. Our demonstration has been

¹ Such as the General Data Protection Regulation: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

developed with the intent to give non-expert participants a *feel* of the reasoning used by an RL agent. Such demonstration can help users level their expectations of the capabilities of AI technologies, and help them shape an informed mental representation of these technologies. In this paper, we present a demonstration involving a volunteer, whom is put in the position of a reinforcement learner, with access to a number of actions, some state information, and reward signals in the form of colored lights (see Figure 1). Without knowing it, the volunteer is controlling a simulated hook used to move a container onto a ship; this graphical representation of the task is not shown to the volunteer during the demonstration, but is visible to accompanying people, allowing them to follow the volunteer’s progress.

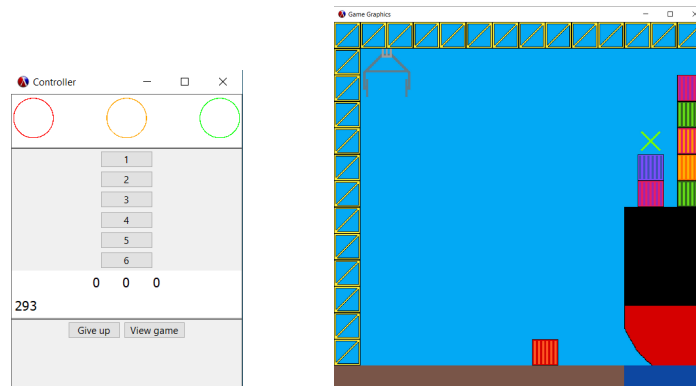


Fig. 1. *Left:* Controller window provided to the volunteer. *Right:* Graphical representation of the task. The volunteer is not shown this window during the demonstration.

Acknowledgement

The first author is funded by the Science Foundation of Flanders (FWO, Belgium) as 1SA6619N Applied Researcher.

References

1. Felzmann, H., Villaronga, E.F., Lutz, C., Tamò-Larrieux, A.: Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society* **6**(1), 2053951719860542 (2019)
2. Hagaras, H.: Toward human-understandable, explainable ai. *Computer* **51**(9), 28–36 (2018)
3. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An Introduction*. MIT Press, Cambridge (2018)

4. Thomaz, A.L., Breazeal, C.: Transparency and socially guided machine learning. In: 5th Intl. Conf. on Development and Learning (ICDL) (2006)
5. Vallor, S., Bekey, G.A.: Artificial intelligence and the ethics of self-learning robots (2017)
6. Weller, A.: Challenges for transparency (2017)