

reinforcement learning (RL)

Authors

Alexander Jung*, Department of Computer Science, Aalto University, Espoo, Finland

*Corresponding author: alex.jung@aalto.fi

Abstract

Reinforcement learning (RL) is an online learning setting in which, at each time step t , the loss can be evaluated only for the single hypothesis $h^{(t)}$ actually used. The hypothesis maps the feature vector $\mathbf{x}^{(t)}$ of the current data point to an action, and the usefulness of that action is reported by a reward $r^{(t)}$, which acts as a noisy estimate of the negative loss. In general the reward depends on earlier actions as well, and the goal is to choose the hypotheses so that the possibly discounted cumulative reward is maximized. Unlike supervised learning, RL never observes the correct label for a data point, only the reward of the action taken.

Definition

RL refers to an online learning setting where it is only possible to evaluate the usefulness of a single hypothesis (i.e., a specific choice of model parameters) at each time step t . In particular, RL methods apply the current hypothesis $h^{(t)}$ to the feature vector $\mathbf{x}^{(t)}$ of the newly received data point to predict the next action. The usefulness of the resulting prediction $h^{(t)}(\mathbf{x}^{(t)})$ is quantified by a reward signal $r^{(t)}$ (see Fig. 1).




Figure 1: Three consecutive time steps $t, t+1, t+2$ with corresponding loss functions $L^{(t)}, L^{(t+1)}, L^{(t+2)}$. During time step t , an RL method can evaluate the loss function only for one specific hypothesis $h^{(t)}$, resulting in the reward signal $r^{(t)} = -L^{(t)}(h^{(t)})$

In general, the reward depends also on the previous predictions $h^{(t')}(\mathbf{x}^{(t')})$ for $t' < t$. The goal of RL is to learn $h^{(t)}$, for each time step t , such that the (possibly discounted) cumulative reward is maximized (Jung, 2022; Sutton and Barto, 2018).

Cross References

reward, loss function, ML

References

1. Jung (2022). Machine Learning: The Basics. Springer Nature, Singapore, Singapore.
2. Sutton and Barto (2018). Reinforcement Learning: An Introduction. 2nd ed., MIT Press, Cambridge, MA, USA.