

## RL Theory

[Online RL](#) / 22. Introduction

# 22. Introduction

[PDF Version](#)

Online learning in reinforcement learning refers to the idea that a learner is placed in an (initially) *unknown* MDP. By interacting with the MDP, the learner collects data about the unknown transition and reward function. The learner's goal is to collect as much reward as possible, or output a near-optimal policy. The difference to planning is that the learner does *not* have access to the true MDP. Unlike in batch RL, the learner gets to decide what actions to play. Importantly, this means the learner's action affect the data that is available to the learner (sometimes referred to as "closed loop").

The fact that the learner needs to create its own data leads to an important decision: Should the learner sacrifice reward to collect more data that will improve decision making in the future? Or should it act according to what seems currently best? Clearly, too much exploration will be costly if the learner chooses actions with low reward too often. On the other hand, playing actions that appear optimal with limited data comes at the risk of missing out on even better rewards. In the literature, this is commonly known as *exploration-exploitation dilemma*.

The exploration-exploitation dilemma is not specific to the MDP setting. It already arises in the simpler (multi-armed) bandit setting (i.e. an MDP with only one state and stochastic reward).

In the following, we focus on finite-horizon episodic (undiscounted) MDPs

. The learner interacts with the MDP for episodes of length . At the beginning of each episode , an initial state is sampled from the initial distribution . The data collected during the episode is

where is the action chosen by the learner at step , is the next state and is the (possibly stochastic) reward.

This model contains some important settings as a special case. Most notably,

- $\mathcal{C}$  recovers the contextual bandit setting, where the “context”  $c$  is sampled from the distribution  $\mathcal{C}$
- $\mathcal{A}$  and  $\mathcal{C}$  is the finite multi-armed bandit setting.

## Sample complexity and regret: How good is the learner?

The goal of the learner is to collect as much reward as possible. We denote

$r_t$  as the reward collected by the learner in episode  $t$ . The total reward is  $\sum_{t=1}^T r_t$ . For the analysis it will be useful to introduce a normalization: Instead of directly arguing about the total reward, we compare the learner to the value  $v^*$  of the best policy in the MDP. This leads to the notation of *regret* defined as follows:

A learner has sublinear expected regret if  $R_T \leq o(T)$  as  $T \rightarrow \infty$ . Sublinear regret means that the average reward of the learner approaches the optimal value  $v^*$  as the number of episodes increases. Certainly that is a desirable property!

Before we go on to construct learners with small regret, we briefly note that there are also other objectives. The most common alternative is PAC – which stands for *probably approximately correct*. A learner is said to be  $\epsilon$ -PAC if upon termination in episode  $T$ , it outputs a policy such that  $v_T \geq v^* - \epsilon$  with probability at least  $1 - \delta$ . We have discussed PAC bounds already in the context of planning.

The difference to bounding regret is that in the first  $T_0$  episodes, the learner does not ‘pay’ for choosing suboptimal actions. This is sometimes called a *pure exploration problem*. Note that a learner that achieves sublinear regret can be converted into a PAC learner (discussed in the notes). However, this may lead to a suboptimal (large)  $T_0$  in the PAC framework.

### –greedy

There exist many ideas on how to design algorithms with small regret. We first note that a “greedy” agent can easily fail: Following the best actions according to some empirical

estimate can easily get you trapped in a suboptimal policy (think of some examples where this can happen!).

A simple remedy is to add a small amount of “forced” exploration: With (small) probability  $\epsilon$ , we choose an action uniformly at random. Thereby we eventually collect samples from all actions to improve our estimates. With probability  $1 - \epsilon$  we follow the “greedy” choice, that is the action that appears best under our current estimates. This gives rise to the name  $\epsilon$ -greedy.

It is often possible to show that  $\epsilon$ -greedy converges. By carefully choosing the exploration probability  $\epsilon$ , we may show that in finite MDPs, the regret is at most  $O(\frac{1}{\epsilon})$ . As we will discuss later, there are multiple algorithms that achieve a regret of only  $O(\sqrt{T})$ . Thus,  $\epsilon$ -greedy is not the best algorithm to minimize regret.

Not unexpectedly, this type of exploration can be quite sub-optimal. It is easy to construct examples, where  $\epsilon$ -greedy takes exponential time (in the number of states) to reach an optimal policy. Can you find an example (Hint: construct the MDP such that each time the agent explores a suboptimal action, the agent is reset to the starting state)?

On the upside,  $\epsilon$ -greedy is very simple and can easily be used in more complex scenarios. In fact, it is a popular choice when using neural network function approximations, where theoretically grounded exploration schemes are much harder to obtain.

## Optimism Principle

A popular technique to construct regret minimizing algorithms is based on *optimism in the face of uncertainty*. To formally define the idea, let  $\mathcal{E}$  be the set of possible environments (e.g. finite MDPs). We make the realizability assumption that the true environment  $\mathcal{E}^*$  is in this set. After obtaining data in rounds  $1:n$ , the learner uses the observations to compute a set of plausible models  $\mathcal{E}_n$ . The plausible model set is such that it contains the true model with high probability. Although this is not always required, it is useful to think of a decreasing sequence of sets  $\mathcal{E}_1 \supseteq \mathcal{E}_2 \supseteq \dots$ . This simply means that as more data arrives, the learner is able to exclude models that are statistically unlikely to produce the observation data.

The optimism principle is to act according to the policy that achieves the highest reward among all plausible models, i.e.

At this point it is not clear why this leads to an efficient learning algorithm (with small

regret). The idea is that the learner systematically obtains data about the environment. For example, if data contradicts the optimistic model, then  $\theta_t$  is excluded from the set of plausible models in the future. Consequently, the learner chooses a different policy in the next round.

On the other hand, the learner ensures that  $\theta_t$  is close to  $\theta^*$  with high probability. In this case, it is often possible to show that the gap  $\mathbb{E}[R_T] - R^*$  is small (more specifically,  $\mathbb{E}[R_T] - R^* \leq C \sqrt{T}$ ), where  $C$  is a constant that depends on the “size” of  $\mathcal{H}$ .  $\theta_t$  behaves like a statistical estimation error of order  $\frac{1}{\sqrt{T}}$  with a leading constant that depends on the “size” of  $\mathcal{H}$ .

One should also ask if the optimization problem  $\min_{\theta \in \mathcal{H}} \mathbb{E}[R_T]$  can be solved efficiently. This is far from always the case. Often one needs to rely on heuristics to implement the optimistic policy, or use other exploration techniques such as Thompson sampling (see below).

How much regret the learner has of course depends on the concrete setting at hand. In the next lecture we will see how we can make use of optimism to design (and analyze) an online learning algorithm for finite MDPs. The literature has produced a large amount of papers with algorithms that use the optimism principle in many settings. This however does not mean that optimism is a universal tool. More recent literature has also pointed out limitations of the optimism principle, and in lieu proposed other design ideas.

## Notes

### Other Exploration Techniques

Some other notable exploration strategies are:

- Phased-Elimination and Experimental Design
- Thompson Sampling
- Information-Directed Sampling (IDS) and Estimation-To-Decisions (E2D)

## References

The paper showing the details behind how to convert between Regret and PAC bounds.

- Dann, C., Lattimore, T., & Brunskill, E. (2017). Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 30. [\[link\]](#)

