KOENIG
step forward

Introduction to Reinforcement Learning

Introduction

In this pre-course module, you'll be introduced to your instructors, get a flavour of what the course has in store for you, and be given an in-depth roadmap to help make your journey through this specialization as smooth as possible.

An Introduction to Sequential Decision-Making

you will learn how to understand the exploration-exploitation trade-off in sequential decision-making, implement incremental algorithms for estimating action-values, and compare the strengths and weaknesses to different algorithms for exploration.

Markov Decision Processes

When you're presented with a problem in industry, the first and most important step is to translate that problem into a Markov Decision Process (MDP). The quality of your solution depends heavily on how well you do with this translation. You will learn the definition of MDPs, you will understand goal-directed behaviour and how this can be obtained from maximizing scalar rewards, and you will also understand the difference between episodic and continuing tasks.

Value Functions & Bellman Equations

Once the problem is formulated as an MDP, finding the optimal policy is more efficient when using value functions. You will learn the definition of policies and value functions, as well as Bellman equations, which is the key technology that all of our algorithms will use.

Dynamic Programming

You will learn how to compute value functions and optimal policies, assuming you have the MDP model. You will implement dynamic programming to compute value functions and optimal policies and understand the utility of dynamic programming for industrial applications and problems. Further, you will learn about Generalized Policy Iteration as a common template for constructing algorithms that maximize reward.