

---

**Exercise: Reinforcement Learning – Your ideas (ca. 2-3h)**

*“Science advances one funeral at a time”  
(Max Planck)*

*“The future depends on some graduate student who is deeply suspicious of everything I have said.”  
(Geoffrey Hinton)*

## **1. Reinforcement Learning - What are your ideas?**

In the lecture, I introduce the standard reinforcement framework which can be used to learn “good” actions, in the sense that these actions help to maximize the expected rewards. We will practice some of the standard Reinforcement Learning algorithms later.

In this first exercise of Reinforcement Learning I do not want to bias your ideas towards some direction. Instead, you shall develop some own ideas and try them out!

## **2. The task**

Assume a simple robot that lives in a 2D world and build up situations in which the robot gets a positive reward and situations in which it gets negative rewards. Assume that the robot has some sensors and that it can choose from a small set of actions in each simulation step.

Now develop and try out some idea such that the robot learns to choose “good” actions while exploring the world such that it comes more often to states in which it receives a positive reward. Or in other words:

Think about and implement some idea, such that the robot learns actions that lead to positive rewards.

In the exercise, first explain your idea, then your implementation and then show some simulations.