

# ACAI Reinforcement Learning Summer School: Assignments Intro<sup>1</sup>

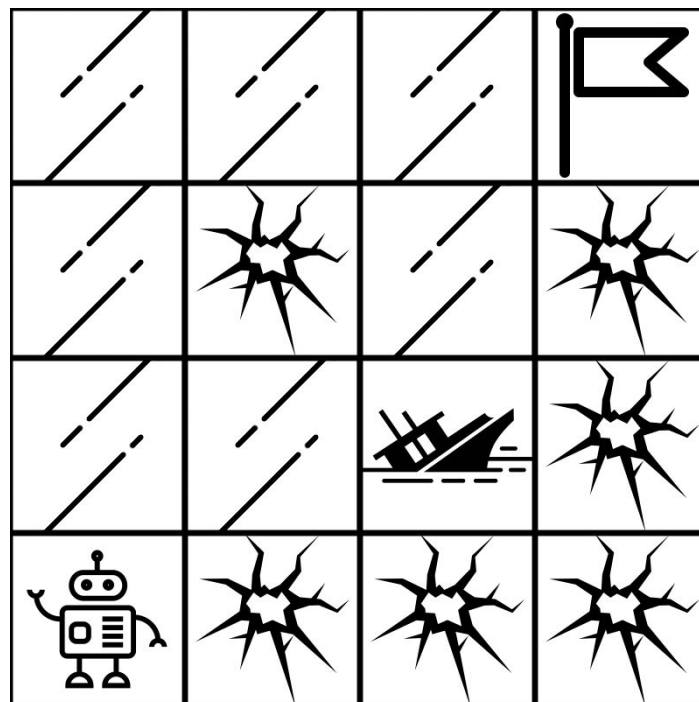
This document describes a simple (yet challenging) environment for which you will develop reinforcement learning agents. The first section presents the environment, the next one contains questions and suggestions about what to do next.

## World Description

The ice world is a small, 4x4 icy grid. A robot is tasked to cross through the ice to reach the opposite side of the small world, while, if possible, collecting some treasure from a shipwreck embedded in the ice.

Unfortunately, the ice has been melting, is slippery, and is already cracked in several places. The robot must avoid falling through the cracks! Due to the slippery nature of the ice, the robot risks sliding on it at each step.

The world can be represented as follows:



(icons from [thenounproject.com](http://thenounproject.com))

<sup>1</sup> <http://roijers.info/acaiss.html>

The robot tries to take one step on the ice, one cell at a time, vertically or horizontally (not diagonally). At each time-step, the robot has a 5% chance of slipping on the ice, and go all the way to the side of the environment. For instance, if the robot tries to go up from any cell in the first column, it may slide across the environment and end up on the top-left cell (or in any crack it may encounter while sliding). The robot cannot move outside the environment, so trying to go up when in the first row has no effect.

Reaching the goal rewards the robot with 100 points. Passing on the shipwreck will allow the robot to collect some of the treasure inside: each time the robot passes there it will get an additional 20 points reward. If the robot falls through the cracks is destroyed (ending the episode), and so it's penalized with -10 points.

More formally, the environment has a total of 16 states (one for each possible position of the robot in the ice world). The robot has 4 actions available ('UP', 'RIGHT', 'DOWN' and 'LEFT'). Each action moves the robot in its direction with probability 0.95. When the robot tries to move outside of the grid, the action will have no effect with probability 1.

Slipping will happen with probability 0.05, and it results in the robot moving in the direction specified by the action up to either the grid border, or to the first crack (into it). (NB: it is possible for both slipping and not slipping as a result of an action to result in the same state, i.e., when the robot is one step away from the edge in the direction of movement, or when the agent is already at the edge and cannot move in that direction any further. Therefore, the transition function for states near the edge (for a given action/direction) looks different from states further away from the edge.)

## Questions

The first two questions cover basic aspects of Reinforcement Learning, and should be implemented for you to be properly trained at RL. The other questions are optional, as they illustrate challenges and ways to improve RL algorithms.

Implementations of the environment described above are provided in Python, Java and Matlab. You can also reimplement the environment yourself in any language you want, but we strongly advise you to use our implementation.

1. Implement Value Iteration<sup>2</sup> on the environment. Our implementation of the environment provides transition probabilities and rewards, so you can focus on the VI algorithm without having to fully understand what the environment does. Your implementation should output the optimal deterministic policy for this environment, along with the value of all 16 states.

---

<sup>2</sup> <http://incompleteideas.net/sutton/book/ebook/node44.html>

2. Implement Q-Learning<sup>3</sup> on the environment. Compare the values output by Value Iteration to your Q-Values. Does Q-Learning lead to the optimal policy? Does it find the same values as Value Iteration?
3. Optional: Implement Policy Iteration<sup>4</sup> and compare the output and runtime of Value Iteration and Policy Iteration.
4. Optional: Implement SARSA<sup>5</sup> and compare the performance (i.e., cumulative reward) during learning of Q-learning and SARSA. Is there a difference, and if so, how can this be explained?
5. Optional: Extend your Q-Learning implementation with Experience Replay<sup>6</sup>. Does it still find the optimal policy? How many episode does your agent need in order to learn (with and without ER)?
6. Optional: Extend Q-Learning with Eligibility Traces<sup>7</sup>. Combining eligibility traces with experience replay is difficult (so don't do it at first). Why, and how would you do it?
7. Optional: Do you see any other way to increase the sample-efficiency of Q-Learning. Why is sample-efficiency so important in real-world applications of reinforcement learning ?

## Plotting

Debugging a reinforcement learning agent is quite difficult. The first and most useful step consists of plotting learning curves. Those curves show the cumulative reward per episode obtained over time, and should go upwards. Plotting those curves is extremely easy using **gnuplot**.

1. Your program should write cumulative rewards in a file. For instance, after each episode, write the episode number and cumulative reward obtained during this episode in a file. The file looks like this:

```
1,1.283764
2,0.938478
3,1.383632
4,1.029384
5,0.29474
```

The file is a valid CSV file, that you can open in Excel/Calc/etc to plot, if you are familiar with one of these programs. Or you can use gnuplot to produce plots very quickly.

2. Install **gnuplot** (gnuplot and gnuplot-qt on most Linux distributions), an Open-Source plotting software available on every major operating system. Gnuplot is actually a tool

---

<sup>3</sup> <http://incompleteideas.net/sutton/book/ebook/node65.html>

<sup>4</sup> <http://incompleteideas.net/sutton/book/ebook/node43.html>

<sup>5</sup> <http://incompleteideas.net/sutton/book/ebook/node64.html>

<sup>6</sup> <http://busoniu.net/files/papers/smcc11.pdf>

<sup>7</sup> <http://incompleteideas.net/sutton/book/ebook/node72.html>

you should discover, if you are not yet using it. It is extremely quick at plotting big datasets, and produces Nature-ready, state-of-the-art-looking plots (with some configuration).

3. Launch Gnuplot, by typing `gnuplot` on the command line.
4. In gnuplot, type these commands (the first one allows Gnuplot to read CSV files)

```
set datafile separator ","  
plot "file.csv" using 2 smooth bezier with lines title  
"Learning Curve"
```

Translated to English, this long command tells Gnuplot to plot data from "file.csv", with the data in the second column, smoothing the data so that it is easier to read, using lines instead of dots or points, with the given title. Our reference implementation produces the following plot:

