
Assignment 1

csce774, Spring 2022

Assigned: January 24

Due: February 7, 9:30am

In this assignment, you will study discrete planning algorithms by implementing some software. You may use any language and computing platform you prefer.

Problem description

The state space is a specific subset of a 2D integer grid. Let W , an integer, denote the *width* of the grid. To make things easier in the definition and the software, assume that W is divisible by 3. The state space X is the set of all (i, j) such that

- both i and j are integers, with $1 \leq i \leq W$ and $1 \leq j \leq W$, and
- at least one of these four inequalities holds: $i \leq W/3$, $i > 2W/3$, $j \leq W/3$, and $j > 2W/3$.

These conditions yield an integer grid in which the middle 1/9 of the points are missing. For the initial state, use $x_I = (1, 1)$. Let the goal X_G be the set of all states (i, j) for which $i > 2W/3$ and $j > 2W/3$.

Let $A = \{0, 1, 2, 3, 4\}$ be a set of actions, which denote:

- 0 stay in the same location
- 1 move right one unit
- 2 move up one unit
- 3 move left one unit
- 4 move down one unit

Let $U = A \cup \{u_T\}$. For variations involving nature (see below), let $\Theta = A$. The state transition function f is formed by applying the effect of both u_k and θ_k . If the resulting value is outside the state space, then the state remains unchanged. For problems that do not involve nature, one can assume that each $\theta_k = 0$, which means that nature does

not interfere with the outcome. Note that the state does not change after u_T is applied, regardless of nature's actions.

For the cost function, let $l(x_k, u_k) = 1$, unless $u_k = u_T$ (in this case, $l(x_k, u_k) = 0$). For the final cost, use the usual $l_F(x_F) = 0$ if $x_F \in X_G$ and $l_F(x_F) = \infty$ otherwise.

Assume that K is *not* given.

Tasks

Your task is to create software that enables you to answer the following questions:

1. Assume that there are no nature effects and that $W = 15$. Use value iteration to compute an optimal action sequence. Give both the action sequence and the final resulting cost. How many iterations were required?
2. Assume that there is nondeterministic uncertainty and $W = 15$. Explain what happens in this case. (Hint: This is easy.)
3. Assume there is probabilistic uncertainty and $W = 15$. Assume that the following probabilities hold (except when $u_k = u_T$, in which case the state never changes):
 - $P(\theta_k = 0) = 1/2$
 - $P(\theta_k = 1) = 1/8$
 - $P(\theta_k = 2) = 1/8$
 - $P(\theta_k = 3) = 1/8$
 - $P(\theta_k = 4) = 1/8$

Use value iteration to compute the plan that minimizes the expected cost to within some reasonable accuracy. Show, on a 15×15 grid, the optimal cost-to-go values for each state. On another 15×15 grid, show robot's optimal action for each state under this strategy.

4. Conduct a simulation of the execution of the optimal strategy for the previous scenario. Generate at least 1000 sample paths by making nature decisions, according to the given distribution, using a pseudo-random number generator. For each path, compute the cost incurred. Make a histogram with cost values along the horizontal axis, and the frequency of their occurrence along the y -axis. What is the average cost incurred? How does this compare to computed expected cost for this strategy?
5. Have some fun with the code. How large can you make W and still get it to finish? What happens if $P(\theta_k = 0)$ is chosen to be very small? If $P(\theta_k = 0)$ is close to one, do you get results similar to that from part 1?

Submitting your solution

Send by email:

- A report, in PDF format, answering the questions raised above.
- Any source code you wrote to complete the assignment.