



**Post-Graduation in  
Electrical and Computer Engineering**

**TOPICS ON ADVANCED ROBOTICS  
2020/21**

**Homework 2**

**Hand-out:** 19 April 2021

**Due:** 10 May 2021

The objective of this homework is to provide the course students the opportunity to get familiarized with the practical aspects of single robot planning under uncertainty, taking a probabilistic approach. For this purpose, students must

- i) download an appropriate toolbox to solve and execute (PO)MDPs
- ii) download a benchmark POMDP problem
- iii) solve the POMDP problem in ii) using three different approaches:
  - a. assume full state observability and implement the value iteration algorithm to solve the corresponding MDP
  - b. assume full state observability and implement the *Q-learning* algorithm to solve the corresponding MDP
  - c. use the POMDP toolbox to solve the problem under partial observability, i.e., the full problem

and compare results obtained.

Students are also encouraged to play with parameters of the different (PO)MDP solving algorithms and of the (PO)MDP models themselves, and report on the effect of such modifications.

The homework will be evaluated taking into account the quality of reporting on the work done, the results presented and the ability to explain them. The report should be no longer than 10 pages A4, 12pt, 1.5pt spacing.

The (PO)MDP solver toolbox (**mandatory**) can be chosen from<sup>1</sup>

- the Perseus approximate POMDP solver implementation for Matlab [1] (courtesy M. Spaan, TU Delft) - <http://www.st.ewi.tudelft.nl/~mtjspaan/software/approx/> - *easier to use*
- the Symbolic Perseus approximate POMDP solver implementation for Matlab - <https://cs.uwaterloo.ca/~ppoupart/software.html#symbolic-perseus> - *more efficient and can handle factored POMDPs*

Students must develop their own implementation of value iteration and Q-learning algorithms (e.g., in Matlab) to be used in approaches iii) a. and iii) b., respectively.

The **mandatory** problem to be solved (see ii) above) is the “*Hallway (57 state navigation)*” problem, described in [2] (downloadable PostScript at <http://www.cs.brown.edu/research/ai/pomdp/papers/cs-95-11.ps.gz>). The problem description is conveniently coded in the input description file `hallway.POMDP` at <http://www.cs.brown.edu/research/ai/pomdp/examples/index.html>. A Matlab parser for the problem description in POMDP file format, suitable to be used with the Perseus toolbox, is available at <http://www.st.ewi.tudelft.nl/~mtjspaan/software/pomdp/>.

## References

[1] Matthijs T. J. Spaan and Nikos Vlassis. “Perseus: Randomized Point-based Value Iteration for POMDPs”. *Journal of Artificial Intelligence Research*, 24:195-220, 2005

---

<sup>1</sup> Students can also use MADP (<http://www.fransoliehoek.net/fb/index.php?fuseaction=software.madp>), a toolbox for representing multiagent problems and performing planning and learning in multiagent systems, including a set of POMDP solvers, which includes Perseus and enables simulating the obtained policies.

[2] Michael Littman, Anthony Cassandra, and Leslie Kaelbling. "Learning policies for partially observable environments: Scaling up". Brown University, Department of Computer Science Technical Report CS-95-11

---