

CS422 Spring 2005

Project

Due: 2:00pm, Tuesday 5 April 2005.

- This project can be done in a group of any size.
- The main deadlines are
 - Proposal due: Tuesday 8 March
 - Final report due: Tuesday April 5
 - Project presentations: April 5, 7.

Scope of the project

This project is to test one extension of the simple framework of decision-theoretic planning / reinforcement learning that we have been using. In particular, the current version assumes:

- explicit state space
- single agent
- completely observable state
- no model built (for reinforcement learning)
- no generalization

In this project, you are to extend the framework in one direction and investigate how well it works. Some ideas you could investigate are:

- DTP/RL for some real non-trivial domain
- multiple agents (e.g., a two player game) - the theory doesn't need to change; all you need to do is for each agent to maximize for themselves. This works as long as the world is completely observable (and the agents take turns).
- partially observable state, when you only get some perception of the state. The theory says that it shouldn't work in this case, but for some examples, it may.
- investigate methods for reinforcement learning that build a model (a representation of $P(s'|s, a)$ and $R(s, a, s')$ that can be used in value iteration). An example of this is called *prioritized sweeping*.
- a method to learn a Q-function that is not a big array of the state and the action, but a function of the features of the state and the action. This allows you to represent problems where the state space may be huge or infinite.
- different exploration strategies.
- any other idea you may have.

You will need to implement something, but only enough to test how well the idea you are investigating works. Your implementation should be the minimal amount you need to test your extension. You will need to be able to generate some data to say how well it works. You do not need to hand in your code (but may be asked for it).

Note that the emphasis of the first idea (using a non-trivial domain) is on modelling: what is the state space and what are the actions, etc. The emphasis of the other domains is on testing an algorithm. Your implementation and your report must be tailored to the purpose of your project.

Proposal

You need to hand in a 1-page proposal that includes

- all members of the group (including email addresses)
- what you intend to investigate
- what you hope to learn
- what you will actually do

Final Report

Your final report will outline what you have learned. You need to be clear about what you tested and present the evidence that leads you to your conclusion. For example, it needs to state “we investigated whether X works. We found that it works well / badly for this domain. The evidence that we base our conclusion on is ...”.

It needs to have an abstract, an introduction (what is the problem, why is it interesting), a description of what you did (enough so that someone can reproduce what you did), a description of your test methodology, your results a conclusion, acknowledgments and a list of references.

Project Presentations

In the last two lectures of the class each person will give a 3-4 minute presentation about what they learned doing this project. You should see this as an opportunity to explain something interesting to your classmates.

Plagiarism

Unlike assignments, discussing the project with your classmates is encouraged, as long as:

- Everyone you discuss your papers with is listed in the acknowledgements section of your paper
- You are explicit about where the ideas in your project came from.
- What you hand in must be your own work
- If you want to quote something that someone else wrote, put it in quotations and attribute it to that person.
- You must reference every external source (book, paper, website) used.

Please read the bulletin board for more details.