

CS 520 — Artificial Intelligence
Homework #3
Out: April 11, 2001
Due: April 25, 2001

Problem 1. This homework aims to provide a fun setting to explore reinforcement learning and planning in uncertain domains. If you seem to be getting frustrated, come see me or send email soon—you might be doing something wrong. (This is the first year I've used this assignment, and I'd like to make sure it gets any improvement it needs.)

For this homework, you should download the five further files that you will find in the list of links where you got the assignment text itself. By making some small additions to one of these files, `learner.c`, you will be able to construct and visualize a system that learns to route packets dynamically in a network by keeping track of its success in routing packets. The files are:

learner.c Skeleton code for the simulator, including detailed instructions about what to fill in to build a learning agent.

router.tcl A tcl-tk script that calls the simulator and displays the results graphically in a pair of windows. (I've had this call `wish-4.0` to work on Paul, but if you're working elsewhere you may want to change the top line of `router.tcl` to point to a good version of `wish` for that system.)

Makefile The makefile to compile your code.

6x6.net, lana.net Two sample network data files.

To complete the assignment, write the functions specified in `learner.c`, compile the program `learner`, run `./router.tcl graphname xxx.net` and watch. Just hand in the new functions that you write.

In preparing the program, you can consult Justin Boyan and Michael Littman's paper about this program, as well as material in the chapters from Russell and Norvig that are on reserve. In fact, Russell and Norvig suggest some fancier ways of implementing reinforcement learning in addition to the basic Q-learning and extended Q-learning that the assignment asks you to implement. You're welcome to explore these for extra credit...