CMPSCI 383, Fall 2011 Homework 4 Due in class or in the main office of the Computer Science building by 4:00 PM, December 6, 2011

Problem 1: (10 points) Exercise 14.1 on page 558
Problem 2: (10 points) Exercise 14.4 on page 559
Problem 3: (10 points) Exercise 14.8 on page 561
Problem 4: (20 points) Exercise 16.5 on page 641
Problem 5: (10 points) Exercise 17.2 on page 688
Problem 6: (15 points) Exercise 17.4 on page 688

Programming Assignment: (25 points)

For this programming assignment, you will implement the *value iteration* algorithm for a 5×5 gridworld with no walls and a terminal goal in the bottom right corner. Use $\gamma = 0.9$. The agent has four possible actions, *up*, *down*, *left*, *right*. Each action achieves the intended effect with probability 0.8, but the rest of the time, the action moves the agent at right angles to the intended direction (as in Figure 17.1). If the movement would take the agent into a wall, the agent does not move.

Your program should read an input file, *in.txt*, which contains the reward function, R(s). It should then run *value iteration* and print out the final utilities of each state, an optimal policy (any one of them is ok), and the number of iterations required for convergence. The utilities and policy should be printed in a 5×5 grid matching the orientation of the input file. An example *in.txt* is available here: *http://www.psthomas.com/Data/HW4/in.txt*. Assume that the terminal state always transitions to an absorbing state with reward 0, i.e., the utility of the goal state is always equal to R(goal).

You should submit:

- Your source code should be uploaded to the Edlab machines. You should not submit a hard copy of your code. You should provide instructions for compiling and executing your code on the Edlab machines.
- Your program's output for the provided *in.txt*.