Assignment A6: Policy Iteration

CS 6380
Spring 2020

Assigned: 24 March 2020

Due: 31 March 2020

For this problem, we examine the *policy iteration* algorithm. The agent is to learn a policy for the following Wumpus world:

```
0 0 0 G
0 0 P 0
0 0 W 0
0 0 P 0
```

For this assignment, assume the actions available to the agent are

$$A = \{RIGHT, LEFT, DOWN, UP\}$$

where these are movements with probabilistic outcomes as described in the text (i.e., 0.8 probability of going the direction selected, 0.1 of going to either side). This requires development of a transition probability table for the 16 cells for the 4 actions. You are to implement the *policy iteration* algorithm given on p. 653 of the text and show:

- comparable results to those given in Fig. 17.2 (a) and (b), p. 648 of R&N for the 16 utilities learned. In part (b) show results for several values of R so that some lead to the *gold* and some to *death*.
- what values you obtain comparable to those in Fig. 17.3, p. 651 in R&N.
- your results for the plot shown in Fig. 17.5 (a), p. 653 in R&N. Show this plot for γ values of 0.9, 0.99, 0.999, 0.9999, 0.99999 and 0.999999.

- for rewards: movement = -0.4, gold = 10, pit = -5, Wumpus = -10, do 1,000 trials and determine the mean and vairance of the scores, and the mean success rate.

Note that you must implement a policy generator function separately from the utility calcuation. You should handin the source code used in the study. The code should conform to the style requested in the class materials.

Here are the function descriptions of what you are to create:

```
function [policy,U,Ut] = CS6380_MDP_policy_iteration(S,A,P,R,k,gamma)
% CS6380_MDP_policy_iteration - policy iteration
% On input:
%     S (vector): states (1 to n)
%     A (vector): actions (1 to k)
%     P (nxk array): transition model
%     R (vector): state rewards
%     k (int): number of iterations for linear solver
%     gamma (float): discount factor
% On output:
%     policy (1xn vector): actions for each state
%     U (vector): state utilities
%     U_trace (iter x n): trace of utility values during iteration
% Call:
%
%   [S,A,R,P,U,Ut] = CS6380_run_value_iteration(0.999999,1000);
%   p = CS6380_MDP_policy_iteration(S,A,P,R,10,0.999)
%
% Author:
%     T. Henderson
%     UU
%     Spring 2020
%
```