

Assignment 12

Introduction to Machine Learning

Prof. B. Ravindran

1. If $\{A, B, C\}$, $\{B, C\}$ and $\{A\}$ are the only frequent closed itemsets, then which of the following statements are true? (Assume that A , B and C are the only items in our universe. Note that multiple statements could be true)
 - (a) B and C always occur together in a transaction
 - (b) $\{A, B\}$ and $\{A, C\}$ have the same frequency
 - (c) $\{A, B\}$ is not frequent
 - (d) B has a greater frequency than $\{B, C\}$
 - (e) There are 2 maximal frequent itemsets
2. For a transactional database with 7 different items (A, B, C, D, E, F, G), how many possible itemsets would need to be examined if one were to determine frequent itemsets by exhaustive enumeration (brute force)?
 - (a) 255
 - (b) 128
 - (c) 127
 - (d) 256
3. Modelling a living organism as an RL agent, the environment encompasses
 - (a) everything external to the organism
 - (b) some portions internal to the organism as well
4. Suppose we want an RL agent to learn to play the game of golf. For training purposes, we make use of a golf simulator program. Assume that the original reward distribution gives a reward of +10 when the golf ball is hit into the hole and -1 for all other transitions. To aide the agents learning process, we propose to give an additional reward of +3 whenever the ball is within a 1 metre radius of the hole. Is this additional reward a good idea or not? Why?
 - (a) Yes. The additional reward will help speed-up learning.
 - (b) Yes. Getting the ball to within a metre of the hole is like a sub-goal and hence, should be rewarded.
 - (c) No. The additional reward may actually hinder learning.
 - (d) No. It violates the idea that a goal must be outside the agents direct control.
5. You face a particularly challenging RL problem, where the reward distribution keeps changing with time. In order to gain maximum reward in this scenario, does it make sense to stop exploration or continue exploration?
 - (a) Stop exploration

- (b) Continue exploration
6. This question is related to the one discussed in class. Recall the temporal difference learning approach to the tic-tac-toe problem. Suppose that the probability of winning at a particular state is 0.6, the max probability value in the next set of states is 0.8, and based on our exploration policy, we choose a next state which has probability value 0.4. Should you backup the current state's probability value based on this choice of next state (i.e., move probability value 0.6 closer to 0.4) or not, given that the agent never stops exploring (i.e., the agent always makes an exploratory move some fraction of the time)?
- (a) Backup the value
 - (b) Do not backup the value
7. Consider a transactional database, where every item (or 1 itemset e.g A or $\{A\}$) is frequent. The height of the FP Tree constructed from this database is then equal to
- (a) The number of transactions in the database
 - (b) One greater than the number of frequent itemsets in the database
 - (c) One greater than the maximum number of distinct items occurring in a transaction
 - (d) One greater than the average number of items per transaction in the database
 - (e) One greater than the number of items in the database
8. How many passes does the FP-Tree algorithm perform over the transactional database?
- (a) 1
 - (b) 2
 - (c) 3
 - (d) 4