

COMPSCI 295, Fall 2019

Assignment 4

Instructor: Rina Dechter

Due (by email): Friday, Nov 15

We had 4 presentations of papers in the last class. Below are a set of questions, one from each presentation. Please answer 2 of the 4 questions. You should send me your homework as well as to each of the presenters of the questions you selected since they will grade your answer. Precede each question with a few sentences of summary of the paper

1. (Paper: Efficient Solution Algorithms for Factored MDPs, Presenter: Yasaman Razeghi). Please formulate the 4 machine instances problem of the SysAdmin problem with Linear Programming. Show the four steps of 1. Problem representation, 2. Selection of Basis Functions, 3. Backprojection, 4. LP construction. Then show the steps of solving the LP with variable elimination. There is no need to write the optimization process.
2. (Paper: Incremental Learning of Planning Actions in Model-Based Reinforcement Learning, Presenter: Hieu Le) : What do high reliability and low reliability mean? How does each case affect learning of new rules?
3. (presenter: Sri Krishna Priya Dhulipala) As discussed in section 2.1 of the paper "R-Max: A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning", describe what stochastic games are and how do they relate to Markov Decision Process.
4. (Paper: Greedy Algorithms for Sparse Reinforcement Learning, Presenter: Michael Tao-Yi Lee) What are the advantages and disadvantages of OMP-BRM over OMP-TD? *Hint*: Exercise both algorithms listed in the paper with the 5-state Markov chain described below, assuming discount factor $\gamma = 0.7$. How many iterations did you spend to reach stopping criterion? What are the resulting index sets \mathcal{I} ?

