Fall 2024 - CSC545/645 Artificial Intelligence - Assignment 10

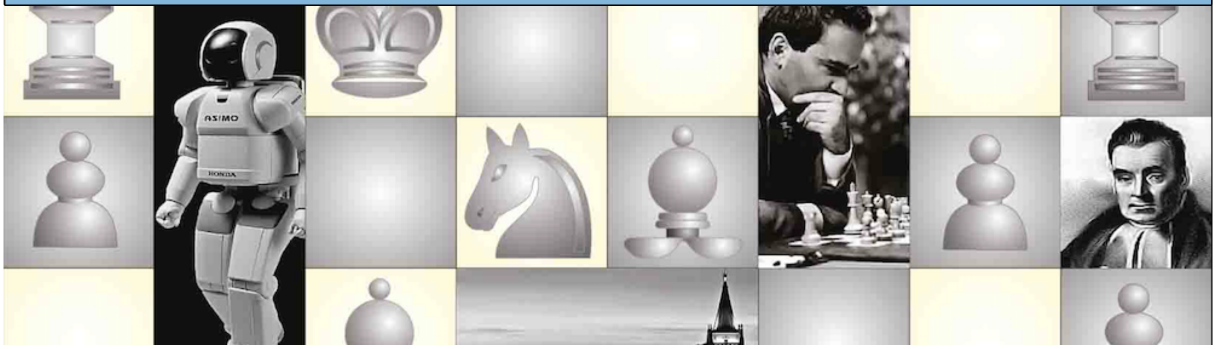

Due date: Thursday, November 21, 2024, 2:00 pm. Please create a folder called `assignment10` **in your local working copy of the repository and place all files and folders necessary for the assignment in this folder. Once done with the assignment, add the files and folders to the repo with** `svn add` *`files, folders`* **and then commit with** `svn ci -m "SOME USEFUL MESSAGE"` *`files, folders`*.

Exercise 10.1 [20/25 points (CSC545/645)]
Read chapter "Probabilistic Reasoning over Time" of the textbook (chapter 15 in the 3rd edition, chapter 14 in the 4th edition).

1. A professor is investigating the sleep habits of students. Each day, the professor observes whether students are sleeping during class and whether they exhibit signs of sleep deprivation, such as red eyes. The professor's domain theory is as follows:

   (a) The prior probability of getting enough sleep, with no observations, is 0.7.

   (b) The probability of getting enough sleep on night t is 0.8 given that the student got enough sleep the previous night, and 0.3 if not.

   (c) The probability of having red eyes is 0.2 if the student got enough sleep, and 0.7 if not.

   (d) The probability of sleeping in class is 0.1 if the student got enough sleep, and 0.3 if not.

   Formulate this information as a dynamic Bayesian network (DBN) that the professor could use to filter or predict from a sequence of observations. Then reformulate it as a hidden Markov model that has only a single observation variable. Give the complete probability tables for the model.
   [10 points]

2. For the DBN specified in Exercise 10.1 and for the evidence values
   $\mathbf{e}_1$ = not red eyes, not sleeping in class
   $\mathbf{e}_2$ = red eyes, not sleeping in class
   $\mathbf{e}_3$ = red eyes, sleeping in class
   perform the following computations:

   (a) State estimation: Compute $P(EnoughSleep_t|\mathbf{e}_{1:t})$ for each of t = 1,2,3.

   (b) Smoothing: Compute $P(EnoughSleep_t|\mathbf{e}_{1:3})$ for each of t = 1, 2, 3.

   (c) Compare the filtered and smoothed probabilities for $t = 1$ and $t = 2$.

   [10 points]

3. Mandatory for CSC645 students, optional for others: Consider applying the variable elimination algorithm to the umbrella DBN unrolled for three slices, where the query is $\mathbf{P}(R_3|u_1, u_2, u_3)$. Show that the space complexity of the algorithm – the size of the largest factor – is the same, regardless of whether the rain variables are eliminated in forward or backward order.
[5 points]