

Sample questions for COMP-424 final exam

Doina Precup

These are examples of questions from past exams. They are provided without solutions. However, Doina and the TAs would be happy to answer questions about the solutions, if you try them. Note that the exam also has questions similar to those on the homeworks.

1. Search algorithms

(a) Suppose you have an admissible heuristic h . Is h^2 admissible? Is \sqrt{h} admissible? Would using any of these alternatives be better or worse than using h in the A^* algorithm?

(b) Suppose you were using a genetic algorithm and you have the following two individual, represented as strings of integers:

1324421 and 2751421

Show the result of performing crossover between the 3rd and 4th digit.

(c) Which of the following statements, contrasting genetic algorithms and simulated annealing, are true?

- i. Genetic algorithms are used for minimization problems while simulated annealing is used for maximization problems
- ii. Genetic algorithms maintain several possible solutions, whereas simulated annealing works with one solution.
- iii. Genetic algorithms maintain one solution, whereas simulated annealing maintains several possible solutions.
- iv. Simulated annealing is guaranteed to produce the best solution, while genetic algorithms do not have such a guarantee.

2. Machine learning

(a) Which of the following can learn the OR function (circle all that apply):

- i. linear perceptron
- ii. a single sigmoid neuron
- iii. a network of sigmoid neurons with one hidden layer
- iv. none of the above

(b) Which of the following can learn the XOR function (circle all that apply):

- i. linear perceptron
- ii. a single sigmoid neuron
- iii. a network of sigmoid neurons with one hidden layer
- iv. none of the above

3. Bayes nets

The Starfleet academy has decided to create a class of android students. 90% of these androids study hard for their exams. Out of the androids who study hard for an exam, 80% get an A. Out of the androids who do not study, only half get an A. Androids who study hard have a 75% probability of depleting their battery in less than a day. Androids who do not study hard have a longer battery life: only 10% of them deplete their batteries within the next day.

- (a) Draw a Bayes net describing the problem statement above.
- (b) You notice that your android has depleted its battery in less than a day. What is the probability that it will get an A on the exam it had yesterday?
- (c) Your friend does not believe that androids are much good at studying. He says he is willing to pay you \$10 if your android gets an A in the class. Recharging the battery costs \$5. Suppose that you could program your android to study or not to study at will (this is not very ethical, but it is technically feasible). What is the best course of action for you?

4. Logic

- (a) Translate the following sentences in first-order logic:
 - i. All citizens of Fredonia speak the same language.
 - ii. The Fredonese language has two dialects
 - iii. Each citizen of Fredonia speaks exactly one of the two dialects
- (b) Translate the knowledge base above into conjunctive normal form (using Skolem constants and functions as appropriate)

5. Naive Bayes


Data the android is about to play in a concert on the Enterprise and he wants to use a naive Bayes classifier to predict whether he will impress Captain Picard. He believes that the outcome depends on whether Picard has been reading Shakespeare or not for the three days before the concert. For the previous five concerts, Data has observed Picard and noted on which days he read Shakespeare. His observations look like this:

D1	D2	D3	LikedConcert
1	1	0	yes
0	0	1	no
1	1	1	yes
1	0	1	no
0	0	0	no

- (a) Show the Naive Bayes model that Data obtains, using maximum likelihood, from these instances.
- (b) If Picard reads Shakespeare only on day 1, how likely is he to enjoy Data's concert? What if he reads Shakespeare on days 1 and 3?

6. Markov Decision Processes

Consider the n -state MDP in the figure below. In state n there is just one action that collects a reward of $+10$, and terminates the episode. In all the other states there are two actions: float, which moves deterministically one step to the right, and reset, which deterministically goes back to state 1. There is a reward of $+1$ for a float and 0 for reset. The discount factor is $\gamma = \frac{1}{2}$.



- (a) What is the optimal policy?
- (b) What is the optimal value of state n , $V^*(n)$?
- (c) Compute the optimal value function, $V^*(k)$ for all $k = 1, \dots, n-1$.
- (d) Suppose you are doing value iteration to figure out these values. You start with all value estimates equal to 0. Show all the non-zero values after 1 and 2 iterations respectively.
- (e) Suppose that instead of knowing the model of the MDP, you decided to always float and observed the following trajectories:
- $n-1$, float, $+1$, n , float, $+10$.
 - $n-2$, float, $+1$, $n-1$, float, $+1$, n , float, $+10$.

What would be the Monte Carlo estimate for all states, based on this data? What would be the certainty equivalence estimate?

7. Applying AI methods

You have been hired by a large retail company who is having trouble managing its inventory. The company provides 100 kinds of products, and for each product there are two or three

suppliers. For each supplier, the company keeps track of their price, of whether they deliver on time or if not, how late they are on each order, and on how many broken products they get in each order. The company also knows at what prices it is able to sell each product, and how fast they can sell it. The company has past data about all the orders it placed or got in the last 5 years. The company also has limited space to store its stock of products. You are required to develop a software package which would decide as well as possible what kinds of products to keep in stock, and what orders to place in order to get them.

Choose **one AI technique** to solve this task. Explain in detail how it would be implemented, what are the pros and the cons.

8. Problem formulation

You have been hired by a scientific computing center to help them manage the requests they get. Presently, m scientists have each submitted a list of n computations that they would like to complete. The center has 3 machines: a parallel supercomputer, a quantum computer and a cluster of regular computers. Each scientist submits, with each job, a list of the computers that would be suitable to execute it. There are $m \leq k < mn$ days on which to perform computations before the center closes for Christmas. To be fair, at least one computation from each scientist's list has to be performed. Each computation takes exactly one day. Each computer can only do one computation on any given day.

- (a) Describe this as a constraint satisfaction problem. Clearly specify what are the variables, the domains of the variables and the constraints.
- (b) Your manager also wants to minimize the total time in which the machines are idle (keeping all the same constraints as above). What kind of problem do you have? Specify one suitable solution method and motivate your choice in one sentence.

9. Logic

- (a) Translate the following sentences in first-order logic.
 - i. Star Trek, Star Wars and The Matrix are science fiction movies.
 - ii. Every AI student loves Star Trek or Star Wars.
 - iii. Some AI students do not love Star Trek.
 - iv. All AI students who love Star Trek also love The Matrix.
 - v. Every AI student loves some science fiction movie.
 - vi. No science fiction movie is loved by all AI students.
 - vii. There is an AI student who loves all science fiction movies.
- (b) Based on the knowledge base above, prove formally that there exists some AI student who loves Star Wars.