Artificial Intelligence—Qual (Fall 2021)

Your name:

Please read instructions carefully before starting to write, work carefully, and circle your final answers (unless it is obvious what the final answer is). Do not worry if you cannot finish everything. Do not write down disorganized answers in the hope of getting partial credit; I will give partial credit only for clearly and correctly worked out parts, so it is better to do fewer questions completely right. You can use extra pages. There are 100 points in total. **The exam is closed book, no calculators.** Good luck!
–Vince

1 Problem 1 (25 points): Search, constraint satisfaction, constrained optimization.

It is time to schedule the qualifying exams in the Department of Silly Engineering. The following qualifying exams need to be scheduled:

- Adorable Airplanes (A)
- Bubbly Boats (B)
- Crinkly Computers (C)
- Delicious Dams (D)
- Evaporating Engines (E)

All exams must be scheduled on Monday, Tuesday, or Wednesday. Besides that, we face the following constraints:

1. There can be at most two exams on a day.
2. At least one of A and B must be scheduled on Wednesday.
3. C must be scheduled on the day immediately before A (e.g., if A is on Wednesday, then C must be on Tuesday).
4. B and D must be scheduled two days apart (so one on Monday and one on Wednesday).
5. E cannot be on the same day as A, and E cannot be on the same day as B.

We will solve this as a constraint satisfaction problem. For every one of the five exams, there is a variable that can take one of three values (days).

(a, 7 points.) **Use** depth-first search to find a feasible solution (returning the first feasible solution). **Show** your entire search tree. You must assign values to the variables in the order A, B, C, D, E, and you must start with the earlier values (first Monday, then Tuesday, then Wednesday). **Write** the complete current partial assignment in each search tree node. Do not do anything fancy such as filling in the values of variables that have only one feasible value left or arc consistency; just try to assign values to variables in the right order, and backtrack (don't generate children) if the current assignment already violates a constraint (i.e., there is no combination of values left for the variables mentioned in that constraint that would satisfy that constraint). (If you are familiar with the Backtracking-Search algorithm from the Russell and Norvig textbook, this algorithm is equivalent to that.)

(b, 6 points.) Do the same, but now in the order D, E, A, B, C. You still must start with the earlier values (first Monday, then Tuesday, then Wednesday).

Now, we add an objective function, turning this into a constrained optimization problem instance. Suppose our objective is to have as many of A and B on Wednesday as possible. So, whenever one of A and B is scheduled on another day, this comes at a cost of 1. The cost so far is the number of A and B that have already been scheduled on a different day than Wednesday.

(c, 6 points.) **Solve** the constrained optimization problem instance using depth-first branch-and-bound search. Recall that this means depth-first search, except we don't stop at the first solution found; instead, we continue with depth-first search, except we don't explore nodes that (purely based on the cost so far) can be no better than the best complete feasible solution found so far. This is continued until no nodes are left to be explored, at which point the best complete feasible solution found so far is returned. Use the order D, E, A, B, C. *Hint:* use/refer to your work from part **(b)** but be very clear about what is different now.

(d, 6 points.) **Solve** the constrained optimization problem instance using the A* algorithm. You must assign values to the variables in the order A, B, C, D, E. So, again, at each node there is a partial assignment, and children of this node will have a value assigned to the next variable in the order. **Write down** the total cost so far at each node (which in this case will be used as the $g + h$ value in A*).[1] You may break ties as you wish, but point out explicitly where you break ties.

[1] That is, g is the cost from the assignment so far, and h is just 0.

2 Problem 2 (25 points): First-order logic.

(a, 10 points.) **Translate** the following sentences into first-order logic.[2] For full credit, and to make the next part easier, use disjunctions, i.e., put everything in the form

$$(\forall x, y, \ldots) \ \ldots \vee \ldots \vee \ldots \ \ldots$$

The first one has been done for you, as an example. Also, you can use more than one proposition for one English sentence (since propositions just count as having an $\wedge$ between them). Note the difference between being *in* a country and being *from* a country.

1. Every student is from Lockdownia, or from Openland, or from Universitasia.

$$(\forall x) \ \neg S(x) \vee F(x, L) \vee F(x, O) \vee F(x, U)$$

2. If there is a virus that is active, then no student from Lockdownia has a visa. (Use $V(\cdot)$ for the property of being a virus, $A(\cdot)$ for the property of being active, and $HV(\cdot)$ for the property of having a visa.)

3. If there is a virus that is active, then every student from Openland has that virus. (Use $H(\cdot, \cdot)$ for the relation of having something like a virus; don't get confused with HV.)

4. If a student is in Universitasia, then that student is from Universitasia, or that student has a visa (or both). (Use $I(\cdot, \cdot)$ for the relation of being in a country.)

5. Alice is a student who is in Universitasia, and she does not have any virus. (Use A for Alice.)

6. The stinky hair virus is active. (Use SH for the stinky hair virus.)

(b, 15 points.) Prove, by resolution, that Alice is from Universitasia. Do this by assuming

7 Alice is not from Universitasia.

and deriving a contradiction. Each step should indicate which known propositions you use (please number all propositions) and give a new proposition that is unconditionally true (given everything above). E.g., write "By applying resolution to i and j, we obtain k: ..." For full credit, use resolution as much as you can; at most small amounts of partial credit will be given for non-resolution steps. (It is possible to prove it using only resolution.) Please try to make as many of your propositions about Alice as possible.

[2] Please note that I am not trying to downplay the continued seriousness of our own pandemic. Please be careful and I hope you are all doing well.

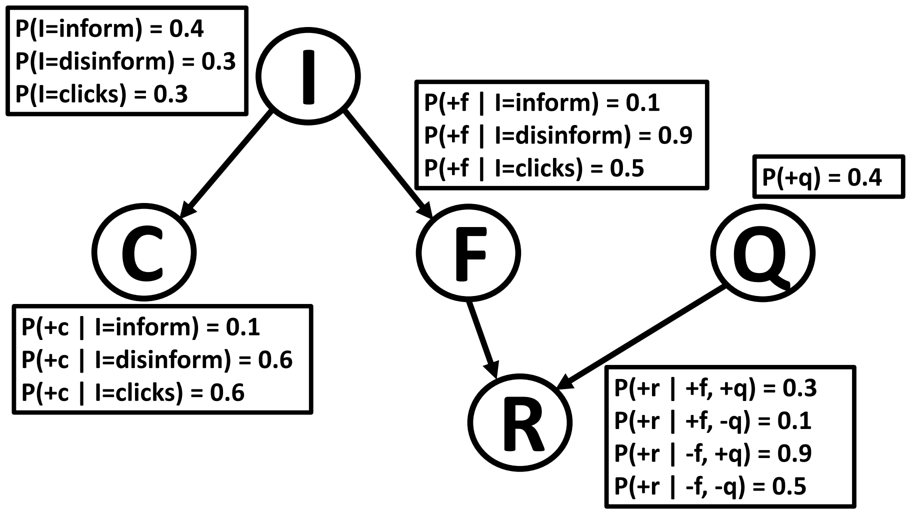

Figure 1: Fake news Bayesian network.

3 Problem 3 (25 points): Bayesian Networks.

We are interested in understanding which news articles are fake. After much research,[3] we have come up with the Bayesian network in Figure 1 to model the situation. There are five variables. **Intent** is the intent of the author, which can be to inform, to misinform, or to get clicks. **Charged** is whether the article uses politically charged (arousing strong emotions) words. **Fake** is whether the article is fake (untruthful). **References** is whether the article has good references to support its position. **Quality** is whether the quality of the writing is high. For the binary variables, we use a shorthand: e.g., $+f$ is shorthand for **Fake** = *true*. Note that, e.g., we don't have to specify $P(-f \mid I = \text{inform})$ explicitly because $P(-f \mid I = \text{inform}) = 1 - P(+f \mid I = \text{inform}) = 0.9$. (We didn't really have to specify $P(I = \text{clicks})$ either.)
(a, 5 points.) The conditional probability table (CPT) for **References** currently has 4 values. If we added an additional edge from **Intent** to **References**, **how many** probability values would the CPT for **References** have then?

(b, 10 points.) **Label** each of the following statements as true or false. (This can be done purely based on the *structure* of the network, and not the numbers in the CPTs.) When you label one as false, **state** in parentheses at which node the path between the variables is blocked (according to the concept of *d*-separation), or justify your statement in a different way. (This is for the *original* network, i.e., there is *no* edge from **Intent** to **References**.)

1. **Charged** and **References** are independent.
2. **Charged** and **Quality** are independent.
3. **Fake** and **Quality** are conditionally independent, given **Intent**.
4. **Charged** and **Quality** are conditionally independent, given **References**.
5. **Charged** and **Quality** are conditionally independent, given **References** and **Fake**.

[3] Not really—this is just a completely made up example for an exam question, based on no data whatsoever.

(c, 10 points.) Compute, using variable elimination, the probability that the article is fake, given that the references are poor, and the quality is low. That is, compute $P(+f \mid -r, -q)$. If some variables are irrelevant, you may drop them from the network for this question. First **write abstract expressions** without any concrete numbers, and simplify them as much as possible; only then use the specific numbers from the example. Give the answer as an exact fraction (you don't have to simplify).

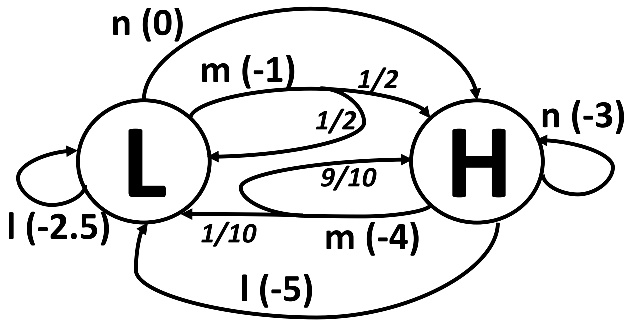

Figure 2: Virus MDP.

4 Problem 4 (25 points): Markov Decision Processes (MDPs).

The highly contagious stinky hair virus has reached Universitasia, which must decide on a policy. There are two states that Universitasia can be in: low caseload (few cases of the virus, L), and high caseload (H). In each state, it can choose to lock down (l), require masks (m), or have no requirements (n). A lockdown will definitely result in reaching L the next time, and no requirements will definitely result in reaching H the next time. Wearing masks in L will result in being in state L with probability 1/2 next period; wearing masks in H will result in being in state L with probability 1/10 next period. In the H state, the immediate reward is -3 for n, -4 for m, and -5 for l. In the L state, the immediate reward is 0 for n, -1 for m, and -2.5 for l. (People don't like mitigation measures, but they like getting sick even less.)[4] Figure 2 summarizes all this information, with rewards in parentheses. The discount factor is $\gamma = 4/5$.

Show that the optimal policy (for the long term, with that discount factor) is to require masks in L, and to lock down in H. To do so, first **compute (15 points)** the long-term discounted values for each state for this policy. (Compute these in exact fractions. *Hint:* the denominator should be 7.) Then, **do (10 points)** one round of policy iteration based on these values and show that the policy does not change (or, equivalently, **show** that the Bellman equation is satisfied for all states by these values)—thereby proving the policy is indeed optimal. For the case of taking action m in H, you don't have to compute the exact value as long as you show that the value is sufficiently low.

[4]Again, this is just an exam question; none of this is based on real data and the numbers are all made up.