# CS 188
## Fall 2008

## Introduction to
## Artificial Intelligence

# Final Exam

**INSTRUCTIONS**

- You have 180 minutes. 100 points total. Don't panic!

- The exam is closed book, closed notes except a two-page crib sheet, non-programmable calculators only.

- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a *brief* explanation. All short answer sections can be successfully answered in a few sentences at most.

- **Question 0: Fill out the following grid and write your name, SID, login, and GSI at the top of each subsequent page.** (-1 points if done incorrectly!)

| | |
|---|---|
| Last Name | |
| First Name | |
| SID | |
| Login | |
| GSI | |
| *All the work on this exam is my own.* **(please sign)** | |

**For staff use only**

| Q. 1 | Q. 2 | Q. 3 | Q. 4 | Q. 5 | Q. 6 | Total |
|------|------|------|------|------|------|-------|
| /17 | /13 | /18 | /15 | /19 | /18 | /100 |

THIS PAGE INTENTIONALLY LEFT BLANK

1. **(17 points.)  Search: A\* Variants**

   **Queuing variants:** Consider the following variants of the A\* *tree search* algorithm. In all cases, $g$ is the cumulative path cost of a node $n$, $h$ is a lower bound on the shortest path to a goal state, and $n'$ is the parent of $n$. Assume all costs are non-negative.

   (i) Standard A\*

   (ii) A\*, but we apply the goal test before enqueuing nodes rather than after dequeuing

   (iii) A\*, but prioritize $n$ by $g(n)$ only (ignoring $h(n)$)

   (iv) A\*, but prioritize $n$ by $h(n)$ only (ignoring $g(n)$)

   (v) A\*, but prioritize $n$ by $g(n) + h(n')$

   (vi) A\*, but prioritize $n$ by $g(n') + h(n)$

   **(a) (3 points)** Which of the above variants are complete, assuming all heuristics are admissible?

   **(b) (3 points)** Which of the above variants are optimal, again assuming all heuristics are admissible?

   **Upper Bounds:** A\* exploits lower bounds $h$ on the true completion cost $h^*$. Suppose now that we also have an *upper bound* $k(n)$ on the best completion cost (i.e. $\forall n, k(n) \geq h^*(n)$). We will now consider A\* variants which still use $g + h$ as the queue priority, but save some work by using $k$ as well. Consider the point at which you are inserting a node $n$ into the queue (fringe).

   **(c) (3 points)** Assume you are required to preserve optimality. In response to $n$'s insertion, can you ever delete any nodes $m$ currently on the queue? If yes, state a general condition under which nodes $m$ can be discarded, if not, state why not. Your answer should involve various path quantities $(g, h, k)$ for both the newly inserted node $n$ and other nodes $m$ on the queue.

   In a *satisficing* search, you are only required to find *some* solution of cost less than some threshold $t$ (if one exists). You need not be optimal.

   **(d) (3 points)** In the satisficing case, in response to $n$'s insertion, can you ever delete any nodes $m$ currently on the queue? If yes, state a general condition, if not, state why not. Your answer should involve various path quantities $(g, h, k)$ for both the newly inserted node $n$ and other nodes $m$ on the queue.

$\epsilon$-**Admissible Heuristics:** Suppose that we have a heuristic function which is not admissible, but $\epsilon$-*admissible*, meaning for some known $\epsilon > 0$,

$$h(n) \leq h^*(n) + \epsilon \qquad \text{for all nodes } n$$

where $h^*(n)$ is the optimal completion cost. In other words, $h$ is never more than $\epsilon$ from being optimal.

**(e) (1 point)** Is using A$^*$ with an $\epsilon$-admissible heuristic complete? Briefly justify.

**(f) (2 points)** Assuming we utilize an $\epsilon$ admissible heuristic in standard A$^*$ search, how much worse than the optimal solution can we get? I.e., if $c^*$ is the optimal cost for a search problem, what is the worst cost solution an $\epsilon$ admissible heuristic would yield? Justify your answer.

**(g) (2 points)** Suggest a modification to the A$^*$ algorithm which will be guaranteed to yield an optimal solution using an $\epsilon$-admissible heuristic with fixed, known $\epsilon$. Justify your answer.

**2. (13 points.)   CSPs: A Greater (or Lesser) Chain**

Consider the general less-than chain CSP below. Each of the $N$ variables $X_i$ has the domain $\{1 \ldots M\}$. The constraints between adjacent variables $X_i$ and $X_{i+1}$ require that $X_i < X_{i+1}$.



For now, assume $N = M = 5$.

**(a) (1 point)** How many solutions does the CSP have?

**(b) (1 point)** What will the domain of $X_1$ be after enforcing the consistency of *only* the arc $X_1 \rightarrow X_2$?

**(c) (2 points)** What will the domain of $X_1$ be after enforcing the consistency of *only* the arcs $X_2 \rightarrow X_3$ then $X_1 \rightarrow X_2$?

**(d) (2 points)** What will the domain of $X_1$ be after fully enforcing arc consistency?

Now consider the general case for arbitrary $N$ and $M$.

**(e) (3 points)** What is the minimum number of arcs (big-O is ok) which must be processed by AC-3 (the algorithm which enforces arc consistency) on this graph before arc consistency is established?

**(f) (4 points)** Imagine you wish to construct a similar family of CSPs which forces one of the two following types of solutions: either all values must be ascending *or* all values must be descending, from left to right. For example, if $M = N = 3$, there would be exactly two solutions: $\{1, 2, 3\}$ and $\{3, 2, 1\}$. Explain how to formulate this variant. Your answer should include a constraint graph and precise statements of variables and constraints.

3. **(18 points.)   RL and MDPs: Two-Armed Bandit**

Imagine you have two slot machine levers. You are playing a game where at each time step, you must pull exactly one lever. Lever A always pays a reward of 6. Lever B pays a reward of either 10 or 0. If B is a lucky lever (L=$\ell$), it pays 10 with probability 4/5. If it is an unlucky one (L=$\neg\ell$), it pays 10 with probability 1/5. B is equally likely to lucky or unlucky a priori. Assume $\gamma = 1$ (which is ok for finite games, not a trick).

**(a) (2 points)** If you can only pull a lever once, what is the MEU?

**(b) (1 point)** Which action(s) (A or B or both) give that MEU?

If you play this game multiple times, it becomes more difficult to figure out what actions to take. The game is formally a POMDP, but we can turn it into an MDP in which the states encode our past outcomes from lever B. In particular, a state will be of the form $(m, n)$, where $m$ is the number of times we pulled B and got 10 and $n$ is the number of times we pulled B and got 0. We begin in state $(0, 0)$. If we then pull lever B and get the outcome 10 we will go to state $(1, 0)$, while getting the 0 outcome puts us in state $(0, 1)$. Your actions are {A, B}, and the rewards are as described above.

**(c) (3 points)** If you will play exactly two rounds, draw the computation tree which represents the possible outcomes of the MDP. Clearly indicate which nodes are of which type (min, max, expectation, etc).

Note that if you pull lever B, the resulting payoff should change your beliefs about what kind of lever B is, and therefore what future payoffs from B might be. For example, if you get the 10 reward, your belief that B is lucky should increase.

**(d) (2 points)** If you are in state $(0, 1)$ and select action B, list the states you might land in and the probability you will land in them.

**(e) (2 points)** If you are in state $(1, 0)$ and select action B, list the states you might land in and the probability you will land in them.

**(f) (3 points)** On the computation tree in (c), clearly mark the probabilities on each branch of any chance nodes.

**(g) (3 points)** Again in this two-round setting, what is the MEU from the state state, and which first action(s) (A or B or both) give it?

**(h) (2 points)** If the number of plays $N$ is large enough, the optimal first action will eventually be to pull lever B. Explain why this makes sense using concepts from reinforcement learning.

### 4. (15 points.)   Bayes Nets: Snuffles

Assume there are two types of conditions: (S)inus congestion and (F)lu. Sinus congestion is is caused by (A)llergy or the flu.

There are three observed symptoms for these conditions: (H)eadache, (R)unny nose, and fe(V)er. Runny nose and headaches are directly caused by sinus congestion (only), while fever comes from having the flu (only). For example, allergies only cause runny noses indirectly. Assume each variable is boolean.



(i)                         (ii)                         (iii)                         (iv)

**(a) (2 points)** Consider the four Bayes Nets shown. Circle the one which models the domain (as described above) best.

**(b) (3 points)** For each network, if it models the domain exactly as above, write *correct*. If it has too many conditional independence properties, write *extra independence* and state one that it has but should not have. If it has too few conditional independence properties, write *missing independence* and state one that it should have but does not have.

(i)

(ii)

(iii)

(iv)

**(c) (3 points)** Assume we wanted to remove the Sinus congestion (S) node. Draw the minimal Bayes Net over the remaining variables which can encode the original model's marginal distribution over the remaining variables.

**(d) (2 points)** In the original network you chose, which query is more efficient to compute using variable elimination: $P(F|r, v, h, a, s)$ or $P(F)$? Briefly justify.

Assume the following samples were drawn from prior sampling:

$$a, s, r, \neg h, \neg f, \neg v$$
$$a, s, \neg r, h, f, \neg v$$
$$a, \neg s, \neg r, \neg h, \neg f, \neg v$$
$$a, \neg s, \neg r, h, f, \neg v$$
$$a, s, \neg r, h, \neg f, \neg v$$

**(e) (1 point)** Give the sample estimate of $P(f)$ or state why it cannot be computed.

**(f) (1 point)** Give the sample estimate of $P(f|h)$ or state why it cannot be computed.

**(g) (1 point)** Give the sample estimate of $P(f|v)$ or state why it cannot be computed.

**(h) (2 points)** For rejection sampling in general (not necessarily on these samples), which query will require more samples to compute to a certain degree of accuracy, $P(f|h)$ or $P(f|h, a)$? Justify your answer in general terms.

**5. (19 points.)   HMMs: Tracking a Jabberwock**

You have been put in charge of a Jabberwock for your friend Lewis. The Jabberwock is kept in a large tugley wood which is conveniently divided into an $N \times N$ grid. It wanders freely around the $N^2$ possible cells. At each time step $t = 1, 2, 3, \ldots$, the Jabberwock is in some cell $X_t \in \{1, \ldots, N\}^2$, and it moves to cell $X_{t+1}$ randomly as follows: with probability $1 - \epsilon$, it chooses one of the (up to 4) valid neighboring cells uniformly at random; with probability $\epsilon$, it uses its magical powers to teleport to a random cell uniformly at random among the $N^2$ possibilities (it might teleport to the same cell). Suppose $\epsilon = \frac{1}{2}$, $N = 10$ and that the Jabberwock always starts in $X_1 = (1, 1)$.

**(a) (2 points)** Compute the probability that the Jabberwock will be in $X_2 = (2, 1)$ at time step 2. What about $P(X_2 = (4, 4))$?

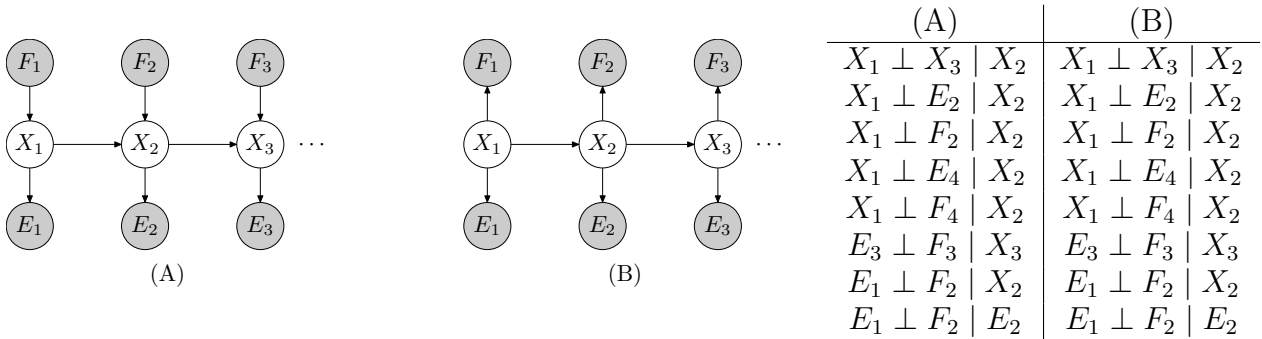$P(X_2 = (2, 1)) =$

$P(X_2 = (4, 4)) =$

At each time step $t$, you don't see $X_t$ but see $E_t$, which is the row that the Jabberwock is in; that is, if $X_t = (r, c)$, then $E_t = r$. You still know that $X_1 = (1, 1)$.

**(b) (4 points)** Suppose we see that $E_1 = 1, E_2 = 2, E_3 = 10$. Fill in the following table with the distribution over $X_t$ after each time step, taking into consideration the evidence. Your answer should be concise. *Hint: you should not need to do any heavy calculations.*

| $t$ | $P(X_t, e_{1:t-1})$ | $P(X_t, e_{1:t})$ |
|---|---|---|
| 1 | | |
| 2 | | |
| 3 | | |

You are a bit unsatisfied that you can't pinpoint the Jabberwock exactly. But then you remembered Lewis told you that the Jabberwock teleports only because it is frumious on that time step, and it becomes frumious independently of anything else. Let us introduce a variable $F_t \in \{0,1\}$ to denote whether it will teleport at time $t$. We want to to add these frumious variables to the HMM.

Consider the two candidates:



|  | (A) | (B) |
| --- | --- | --- |
|  | $X_1 \perp X_3 \mid X_2$ | $X_1 \perp X_3 \mid X_2$ |
|  | $X_1 \perp E_2 \mid X_2$ | $X_1 \perp E_2 \mid X_2$ |
|  | $X_1 \perp F_2 \mid X_2$ | $X_1 \perp F_2 \mid X_2$ |
|  | $X_1 \perp E_4 \mid X_2$ | $X_1 \perp E_4 \mid X_2$ |
|  | $X_1 \perp F_4 \mid X_2$ | $X_1 \perp F_4 \mid X_2$ |
|  | $E_3 \perp F_3 \mid X_3$ | $E_3 \perp F_3 \mid X_3$ |
|  | $E_1 \perp F_2 \mid X_2$ | $E_1 \perp F_2 \mid X_2$ |
|  | $E_1 \perp F_2 \mid E_2$ | $E_1 \perp F_2 \mid E_2$ |

**(c) (3 points)** For each model, circle the conditional independence assumptions above which are true in that model.

**(d) (2 points)** Which Bayes net is more appropriate for the problem domain here, (A) or (B)? Justify your answer.

For the following questions, your answers should be fully general for models of the structure shown above, not specific to the teleporting Jabberwock. For full credit, you should also simplify as much as possible (including pulling constants outside of sums, etc.).

**(e) (2 points)** For (A), express $P(X_{t+1}, e_{1:t+1}, f_{1:t+1})$ in terms of $P(X_t, e_{1:t}, f_{1:t})$ and the CPTs used to define the network. Assume the $E$ and $F$ nodes are all observed.

**(f) (2 points)** For (B), express $P(X_{t+1}, e_{1:t+1}, f_{1:t+1})$ in terms of $P(X_t, e_{1:t}, f_{1:t})$ and the CPTs used to define the network. Assume the $E$ and $F$ nodes are all observed.

Suppose that we don't actually observe the $F_t$s.

**(g) (2 points)** For (A), express $P(X_{t+1}, e_{1:t+1})$ in terms of $P(X_t, e_{1:t})$ and the CPTs used to define the network.

**(h) (2 points)** For (B), express $P(X_{t+1}, e_{1:t+1})$ in terms of $P(X_t, e_{1:t})$ and the CPTs used to define the network.

**6. (18 points.)  Classification and VPI: Cat Cravings**

Consider the following Naive-Bayes model for diagnosing whether your cat is (H)ungry. Signs of hunger include that the cat is (T)hin, (M)eowing, or (W)eak.

| $H$ | $P(H)$ |
|---|---|
| h | 0.5 |
| ¬h | 0.5 |

| $H$ | $T$ | $P(T\|H)$ |
|---|---|---|
| h | t | 0.6 |
| h | ¬t | 0.4 |
| ¬h | t | 0.4 |
| ¬h | ¬t | 0.6 |

| $H$ | $M$ | $P(M\|H)$ |
|---|---|---|
| h | m | 0.6 |
| h | ¬m | 0.4 |
| ¬h | m | 0.4 |
| ¬h | ¬m | 0.6 |

| $H$ | $W$ | $P(W\|H)$ |
|---|---|---|
| h | w | 0.5 |
| h | ¬w | 0.5 |
| ¬h | w | 0.0 |
| ¬h | ¬w | 1.0 |

**(a) (3 points)** If your cat is thin and meowing, but not weak, what is the probability that he is hungry?

**(b) (2 points)** Which of the following smoothing options *might* have been applied to produce the CPTs above from training data? Circle the best answer:

   (i)  Laplace smoothing only might have been applied

   (ii)  Linear interpolation only might have been applied

   (iii)  Neither could have been applied

   (iv)  Either might have been applied

**(c) (2 points)** Assume that no smoothing has been applied (so these are the maximum likelihood estimates). Compute the linear interpolation smoothed estimate of $P_{\text{LIN}}(w|h)$ using $\alpha = 0.5$.

**(d) (2 points)** In a single word, state why smoothing is necessary.

Imagine you cannot tell whether your cat is weak or not.

**(e) (2 points)** Is it correct to simply skip over any unobserved evidence variables when classifying in a Naive Bayes model? That is, will you get the same answer as if you had marginalized out the missing nodes? Briefly justify why or why not.

Now return to the original probabilities, reprinted here:

| H | P(H) |
|---|---|
| h | 0.5 |
| ¬h | 0.5 |

| H | T | P(T\|H) |
|---|---|---|
| h | t | 0.6 |
| h | ¬t | 0.4 |
| ¬h | t | 0.4 |
| ¬h | ¬t | 0.6 |

| H | M | P(M\|H) |
|---|---|---|
| h | m | 0.6 |
| h | ¬m | 0.4 |
| ¬h | m | 0.4 |
| ¬h | ¬m | 0.6 |

| H | W | P(W\|H) |
|---|---|---|
| h | w | 0.5 |
| h | ¬w | 0.5 |
| ¬h | w | 0.0 |
| ¬h | ¬w | 1.0 |

You can decide whether or not to give your cat a mega-feast (F) to counteract his (possible) hunger. Your resulting utilities are below:

| H | F | U(H, F) |
|---|---|---|
| h | f | 0 |
| h | ¬f | -100 |
| ¬h | f | 0 |
| ¬h | ¬f | 10 |

**(f) (2 points)** Draw the decision diagram corresponding to this decision problem.

If you do not know $W$, but wish to determine whether your cat is weak, you can apply the weak-o-meter test, which reveals the value of $W$.

**(g) (3 points)** In terms of *high-level quantities* (MEUs, EUs, conditional probabilities, or similar) and variables, give an expression for the maximum utility you should be willing pay to apply the weak-o-meter, assuming the cat is again thin and meowing?

**(h) (2 point)** What is the maximum utility you should be willing to pay, as a specific real number?