# CS 188 Summer 2020 — Introduction to Artificial Intelligence — Midterm 2

- You have approximately 110 minutes.
- The exam is open book, open calculator, and open notes.
- In the interest of fairness, we want everyone to have access to the same information. To that end, we will not be answering questions about the content or making clarifications.
- For multiple choice questions,
  - ☐ means mark **all options** that apply
  - ◯ means mark a **single choice**

| | |
|---|---|
| First name | |
| Last name | |
| SID | |

**For staff use only:**

| | | |
|---|---|---|
| Q1. | Plants vs. Zombies, Cont'd | /26 |
| Q2. | Potpourri | /15 |
| Q3. | Value of Stock Information | /11 |
| Q4. | Secret Tunnel | /21 |
| Q5. | Probability | /27 |
| | Total | /100 |

THIS PAGE IS INTENTIONALLY LEFT BLANK

# Q1. [26 pts] Plants vs. Zombies, Cont'd

**(a)** Zomboss is sending zombies to invade one of the grass lanes on your lawn! There could be a Snow Pea ($S = \pm s$) and/or a Wall-nut ($W = \pm w$) on the lane, and Zomboss may put a regular Zombie ($Z = \pm z$) and/or a Conehead zombie ($C = \pm c$) to try to get your brain ($B = \pm b$).
You come up with the following Bayes net to model the situation:

(Your BN)

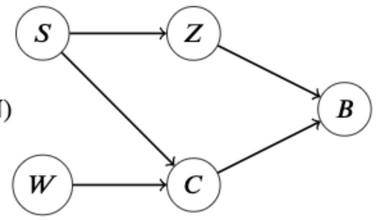


**(i)** [2 pts] Is $W \perp\!\!\!\perp B | C$ guaranteed?
○ Guaranteed ○ Not Guaranteed

**(ii)** [2 pts] Is $Z \perp\!\!\!\perp C | S$ guaranteed?
○ Guaranteed ○ Not Guaranteed

Pacman also comes up with a bunch of Bayes net representations, seen below:

(1)

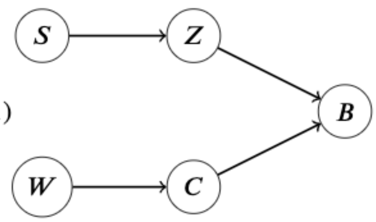


(3)

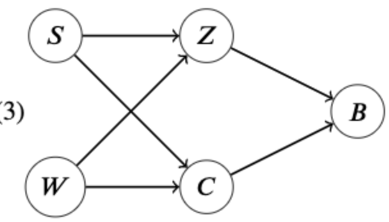


(2)

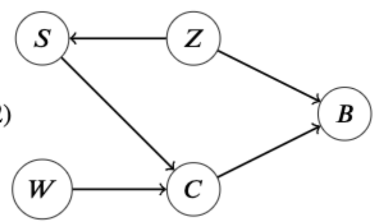


(4)

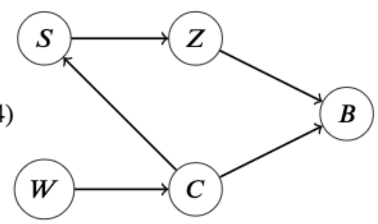


**(iii)** [3 pts] Which of Pacman's Bayes nets can represent at least one of the distributions that your Bayes net can represent?
☐ (1) ☐ (2) ☐ (3) ☐ (4) ○ None

**(iv)** [4 pts] Which of Pacman's Bayes nets can represent every distribution that your Bayes net can represent?
☐ (1) ☐ (2) ☐ (3) ☐ (4) ○ None

**(b)** [4 pts] Perform the first step of variable elimination on your Bayes net representation by eliminating variable $C$. Which probability tables and factors do you have now? Your Bayes net is repeated below for your convenience.

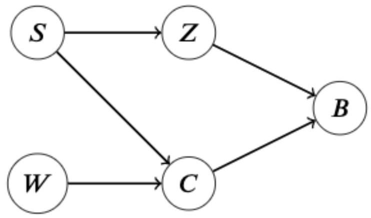


☐ P($S$) ☐ $f(B|S,Z,W)$ ☐ $f(B|Z,W,C)$ ☐ P($B|Z,C$) ○ None of these

☐ P($Z|S$) ☐ $f(B|Z)$ ☐ $f(C|S)$ ☐ P($B|Z,W$)

☐ P($W$) ☐ $f(B|Z,C)$ ☐ P($B|S,Z,W$) ☐ P($B|Z,W,C$)

☐ P($C|S,W$) ☐ $f(B|Z,W)$ ☐ P($B|Z$) ☐ P($C|S$)

**(c)**

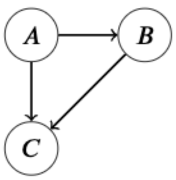


| A | P(A) |
|---|---|
| $a_1$ | 0.1 |
| $a_2$ | 0.9 |

| B | A | P(B\|A) |
|---|---|---|
| $b_1$ | $a_1$ | 0.3 |
| $b_2$ | $a_1$ | 0.7 |
| $b_1$ | $a_2$ | 0.8 |
| $b_2$ | $a_2$ | 0.2 |

| C | A | B | P(C\|A,B) |
|---|---|---|---|
| $c_1$ | $a_1$ | $b_1$ | 0.3 |
| $c_2$ | $a_1$ | $b_1$ | 0.7 |
| $c_1$ | $a_2$ | $b_1$ | 0.6 |
| $c_2$ | $a_2$ | $b_1$ | 0.4 |
| $c_1$ | $a_1$ | $b_2$ | 0.9 |
| $c_2$ | $a_1$ | $b_2$ | 0.1 |
| $c_1$ | $a_2$ | $b_2$ | 0.2 |
| $c_2$ | $a_2$ | $b_2$ | 0.8 |

You'd like to try out sampling to find $P(a_1, c_1|b_2)$. Recall from homework that we generate samples by assigning values to variables **(ordering: A, B, C)** with bins and random numbers in the range [0,1].

**You generate 200 samples with likelihood sampling.**

**(i)** [3 pts] How many random numbers did you use to generate these samples with likelihood sampling? Simplify your answer as much as possible.

**(ii)** [4 pts] How many random numbers would you approximately need to generate with prior sampling if you wanted to reach a similar accuracy as likelihood sampling from the previous part? Recall that we generated 200 samples with likelihood sampling. Simplify your answer as much as possible.

*(Hint: first, think about the* ***number of samples*** *you would need to generate with prior sampling if you wanted to reach a similar accuracy.)*

**(iii)** [4 pts] How many random numbers would you approximately need to generate with rejection sampling if you wanted to reach a similar accuracy as the previous part? Recall that we generated 200 samples with likelihood sampling. Also recall that we assign values to our variables in the order $A, B, C$. Simplify your answer as much as possible.

# Q2. [15 pts] Potpourri

**(a)** [4 pts] (1 pt each) Suppose we run variable elimination on the following Bayes net, where each random variable has domain size 2. We define the size of a factor to be its number of entries. For each variable elimination ordering, what is the size of the largest factor generated?

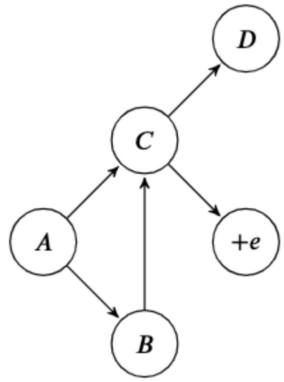


- $A, B, C, D$ ______________
- $C, B, A, D$ ______________
- $D, C, B, A$ ______________
- $B, C, D, A$ ______________

**(b)** [2 pts] Which of the following statements about decision networks are true?

☐ If the VPI of observing a variable is 0, then that variable is conditionally independent from the parents of the utility given current evidence.

☐ If a variable is conditionally independent from the parents of the utility given current evidence, then the VPI of observing that variable is 0.

☐ A Markov decision process can be represented as a decision network.

☐ The value of observing two evidence nodes is always greater than the value of observing one.

○ None of the above

**(c)** Recall the conventions from the lecture notes:

action nodes as rectangles ☐, chance nodes as ovals ◯, and utility nodes as diamonds ◇.

Consider the following decision network.

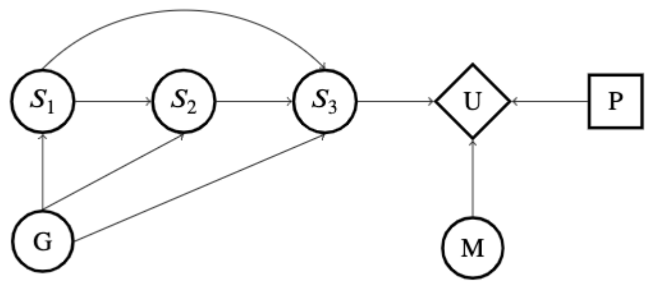


For each expression, select the choices that are possible:

**(i)** [1 pt] $VPI(G|S_3)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(ii)** [1 pt] $VPI(S_1|S_2)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(iii)** [1 pt] $VPI(G) - VPI(S_2)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(iv)** [1 pt] $VPI(G) - VPI(M)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(v)** [1 pt] $VPI(S_2) - VPI(S_1, S_2)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(vi)** [2 pts] $VPI(M, S_3) - VPI(M) - VPI(S_3)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

**(vii)** [2 pts] $VPI(S_2, S_3) - VPI(S_2) - VPI(S_3)$ ☐ $< 0$ ☐ $= 0$ ☐ $> 0$

# Q3. [11 pts] Value of Stock Information

A UC Berkeley professor is trying to maximize the value of her stock holdings. The situation is described as follows:

(1) The professor ($P$) can either sell ($s$), hold ($h$), or buy ($b$) more stock every day.
(2) A market gremlin ($G$) works for the professor's trading company, and has its own opinion on what its customers should be doing. The gremlin takes one of three positions: sell ($s$), hold ($h$), buy ($b$).
(3) The professor **always writes** ($W$) on a media platform. She varies between posting her decision on Reddit ($r$) or Twitter ($t$).
(4) There is a random internal variable representing the market ($M$), which either trends down ($m_0$), trends flat ($m_1$), or trends up ($m_2$).

**Utility function**: The professor only makes money (gains a utility ($U$) of \$100) when the gremlin matches her order, she posts on twitter ($t$), **and** the market is up or flat ($m_1, m_2$). If she selects hold ($h$) and the other conditions are not met, she gains a utility of \$0. For all other cases, her utility is $-\$50$.
Mathematically, this can be expressed as:

$$U(n) = \begin{cases} 100 & \text{if } (P = G) \wedge (W = t) \wedge (M = m_1 \vee M = m_2) \\ 0 & \text{if } (P = h) \wedge \neg((P = G) \wedge (W = t) \wedge (M = m_1 \vee M = m_2)) \\ -50 & \text{else.} \end{cases}$$

**(a)** Before selecting your actions, suppose that someone could tell you the value of $M$ or $G$. Calculate the **maximum expected utility (MEU)** given $M$, given $G$, and given $M$ and $G$ together. Use the following probability tables to compute your answers.

| $M$ | $P(M)$ |
|---|---|
| $m_0$ | $1/4$ |
| $m_1$ | $1/4$ |
| $m_2$ | $1/2$ |

| $G$ | $P(G)$ |
|---|---|
| $s$ | $1/3$ |
| $h$ | $1/3$ |
| $b$ | $1/3$ |

**(i)** [3 pts] What is $MEU(M)$? Simplify your answer as much as possible and round your answer to the nearest **tenth** (i.e.: 1.234 would round to 1.2).

**(ii)** [3 pts] What is $MEU(G)$? Simplify your answer as much as possible and round your answer to the nearest **tenth** (i.e.: 1.234 would round to 1.2).

**(iii)** [3 pts] What is $MEU(M, G)$? Simplify your answer as much as possible and round your answer to the nearest **tenth** (i.e.: 1.234 would round to 1.2).

**(iv)** [2 pts] Would you prefer to be told $G$, $M$, or *either*?

○ $G$ ○ $M$ ○ *either*

# Q4. [21 pts] Secret Tunnel

Aang and his friends are traveling to the Earth Kingdom, when they find themselves trapped in a secret maze of tunnels through the mountains!

They begin in tunnel $Z_1$ and estimate that any tunnel $Z_i$ will collapse with some probability that they can't observe directly; however, Aang's friend Toph can partially sense information about the tunnels, which directly influences her choices $C_i$. Aang also senses the existence of an enormous badgermole $B_i$ roaming around the mountains, shaking the tunnel Aang is currently in.

**(a)** Aang comes up with the following Hidden Markov Model (HMM), where shaded nodes are observable:

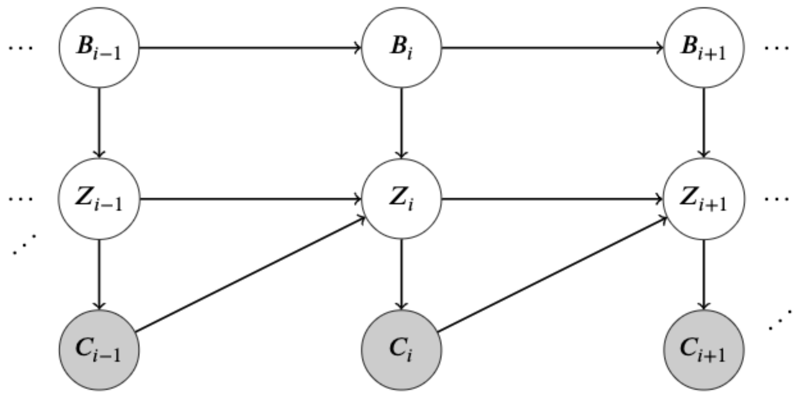


**(i)** [6 pts] Which of the following independence assumptions are enforced by Aang's model? Assume in this subpart that observable nodes **aren't observed** unless given in the independence assumption.

☐ $Z_{i-1} \perp\!\!\!\perp Z_{i+1} \mid Z_i$
☐ $C_{i-1} \perp\!\!\!\perp C_{i+1} \mid C_i$
☐ $B_{i-1} \perp\!\!\!\perp B_{i+1} \mid B_i$
☐ $C_{i-1} \perp\!\!\!\perp Z_{i+1} \mid Z_i$
☐ $C_{i-1} \perp\!\!\!\perp Z_{i+1} \mid Z_i, B_i$
☐ $C_{i-1} \perp\!\!\!\perp B_{i+1} \mid B_i$
☐ $Z_{i-1} \perp\!\!\!\perp B_{i+1} \mid B_{i-1}, C_i$

Aang now wants to derive an algorithm to compute the belief distribution, but he's not sure how to get started! You decide to help him out by providing him with the time elapse update equation.

**(ii)** [8 pts] Derive the time elapse update equation using the available terms going from tunnel $i-1$ to tunnel $i$. Select one choice per blank.

Time Elapse: **(1)** = **(2)** **(3)** **(4)** **(5)**

**(1)** ○ $P(Z_i, B_i)$ ○ $P\left(Z_{i-1}, B_{i-1}\right)$ ○ $P\left(Z_{i-1}, B_{i-1} | c_{1:i-1}\right)$ ○ $P\left(Z_i, B_i | c_{1:i-1}\right)$ ○ $P\left(Z_i, B_i | c_{1:i}\right)$

**(2)** ○ $\Sigma_{z_{i-1}}$ ○ $\Sigma_{b_{i-1}}$ ○ $\Sigma_{z_{i-1}} \Sigma_{b_{i-1}}$ ○ $\Sigma_{z_i} \Sigma_{b_i}$ ○ $\max_{z_i} \max_{b_i}$ ○ $\max_{z_{i-1}}$ ○ $\max_{b_{i-1}}$ ○ $\max_{z_{i-1}} \max_{b_{i-1}}$ ○ 1

**(3)** ○ $P(B_i | b_{i-1})$ ○ $P(B_{i+1} | b_i, z_i)$ ○ $P(Z_i, B_i | b_{i-1})$ ○ $P(b_i | B_{i-1})$ ○ $P(z_i, b_i | B_{i-1})$ ○ $P(B_{i+1} | B_i, Z_i)$ ○ 1

**(4)** ○ $P(Z_i | z_{i-1}, B_i)$ ○ $P(Z_i | z_{i-1}, B_i, c_{i-1})$ ○ $P(Z_i | z_{i-1}, b_{i-1}, c_{i-1})$ ○ $P(z_i | b_i, c_{i-1})$ ○ $P(z_i | Z_{i-1}, b_i)$ ○ $P(z_i | Z_{i-1}, b_i, c_{i-1})$ ○ 1

**(5)** ○ $P(z_{i-1}, b_{i-1})$ ○ $P(z_{i-1}, b_{i-1} | c_{1:i-1})$ ○ $P(z_{i-1}, b_{i-1} | c_{1:i})$ ○ $P(z_i, b_i | c_{1:i-1})$ ○ $P(z_i, b_i | c_{1:i})$ ○ $P(z_i, b_i)$ ○ 1

**(b)** Some wandering nomads have joined Aang's group of friends in the tunnels; however, their constant singing $S_i$ annoys Toph, influencing her choices $C_i$. Aang quickly revises his HMM to incorporate this disturbance, where shaded nodes are observable:

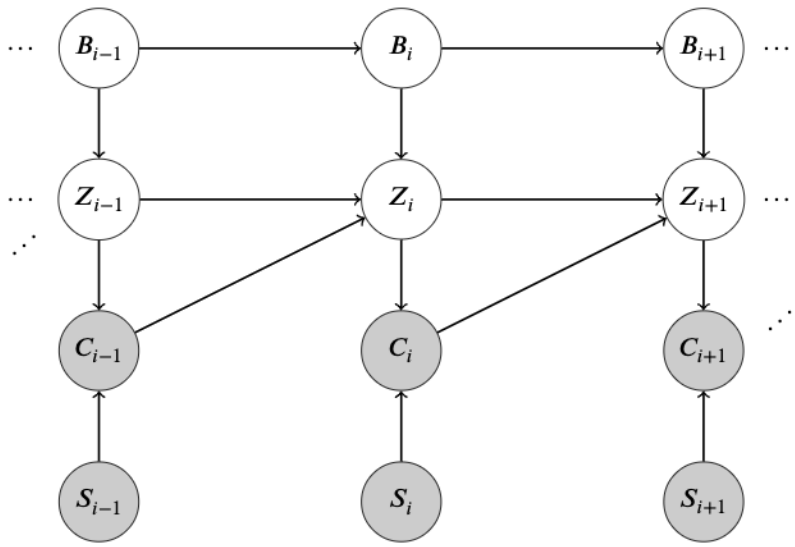


**(i)** [4 pts] Starting from scratch, Aang wants to derive an algorithm to recompute the belief distribution $P(Z_i, B_i|c_{1:i}, s_{1:i})$ from his revised model. He succeeds at getting the time elapse update equation, but can't remember how to proceed from there.

Complete the observation update equation using the available terms to incorporate evidence for tunnel $i$. Select one choice per row for each blank.

Observation: $P(Z_i, B_i \mid c_{1:i}, s_{1:i}) \propto$ **(1)** ______ **(2)** ______ **(Time Elapse Update)**

**(1)** ○ $\Sigma_{z_{i-1}}$ ○ $\Sigma_{b_{i-1}}$ ○ $\Sigma_{z_{i-1}}\Sigma_{b_{i-1}}$ ○ $\Sigma_{z_i}\Sigma_{b_i}$ ○ $\max_{c_i}$
○ $\max_{s_i}$ ○ $\max_{s_{1:i}} \max_{c_i}$ ○ $\max_{s_{1:i}}$ ○ 1

**(2)** ○ $P(c_i)$ ○ $P(c_i|Z_i)$ ○ $P(c_i|Z_i, s_{1:i})$ ○ $P(c_i|Z_i, s_i)$ ○ 1

**(ii)** [3 pts] Aang's friend Katara advises Aang to use particle filtering, but her brother Sokka disagrees with her.
For which of the following cases would particle filtering be a strictly better option over the forward algorithm? Select all that apply, or 'None of the above.'

- ☐ There is an infinite number of tunnels, so $i \to \infty$.
- ☐ All variables are observable.
- ☐ There is an exponentially large amount of unobservable states.
- ☐ The domains of some random variables are infinite.
- ☐ We want the most accurate possible result.
- ☐ None of the above

# Q5. [27 pts] Probability

**(a)** For each statement below about distributions over $W, X, Y$ and $Z$:

- If the statement is **always true** for any distribution, then write **'Always true.'**
- If the statement is **not always true**, write **EXACTLY ONE** conditional independence assumption which makes it true. For example, if the necessary assumption is that $W$ must be independent of $X$ given $Y$, please write $W \perp\!\!\!\perp X|Y$.
- If the statement is **never true**, or is only true with **TWO OR MORE** conditional independence assumptions then write **'No solution.'**

Assume that the distributions are **nontrivial**, meaning that you should not assume them to be always true or always false.

**(i)** [2 pts] $P(W, X|Y) = \frac{P(Y,W|X)P(X)}{P(Y)}$

**(ii)** [2 pts] $P(Y|X, Z) = \frac{P(X)P(Y|X)P(Z|X,Y)}{\sum_y P(X)P(Y|X)P(Z|Y,X)}$

**(iii)** [2 pts] $P(X, Z) = P(X|Y)P(Z)$

**(iv)** [2 pts] $P(X, Y, Z) = P(X, Z)P(Y|X, Z)$

**(v)** [2 pts] $P(W|X, Y) = \frac{P(W)P(X|W)P(Y|W)}{P(X|Y)P(Y)}$

**(vi)** [2 pts] $P(W, X, Y) = P(Y)P(W|Y)P(X|Y)$

**(vii)** [2 pts] $P(W, X|Y, Z) = P(W|Y, Z)P(X|W, Y, Z)$

**(viii)** [2 pts] $P(X|Y, Z) = \frac{P(X,Y)}{P(Y)P(Z)}$

**(b)** You recently flew back to Berkeley for the Online 2020 semester and have decided to use a simple Bayes nets to analyze the probability that you will enjoy your new college experience. However, this depends on whether your friends come back to Berkeley and if you have any early morning classes. You know that the probability of your friends returning (event $F = +f$) is 0.8 and the probability of you having to enroll in a morning class (event $M = +m$) is 0.7. Let random variables $F$ and $M$ represent each of these events and let $L$ take the value $+l$ if you like the new format and $-l$ if you don't.

The relationship between these variables is given by the following Bayes Net:

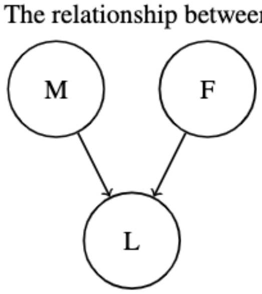


| $M$ | $F$ | $P(L = +l\|M, F)$ |
|---|---|---|
| $+m$ | $+f$ | 0.05 |
| $+m$ | $-f$ | 0.01 |
| $-m$ | $+f$ | 0.9 |
| $-m$ | $-f$ | 0.02 |

**(i)** [3 pts] Give the formula for the joint probability distribution induced by the above Bayes Net.

**(ii)** [2 pts] Compute the probability that your friends return and you don't have to take any morning classes. Simplify your answer as much as possible and round your answer to the nearest **hundredth** (i.e.: 1.234 would round to 1.23).

**(iii)** [2 pts] Compute the probability that your friends return *or* you don't have to take any morning classes. Simplify your answer as much as possible and round your answer to the nearest **hundredth** (i.e.: 1.234 would round to 1.23).

**(iv)** [2 pts] Compute the probability that your friends came back given that you like the semester. Simplify your answer as much as possible and round your answer to the nearest **hundredth** (i.e.: 1.234 would round to 1.23).

**(v)** [2 pts] Compute the probability that you had to take a morning class and none of your friends came back given that you do not like the semester. Simplify your answer as much as possible and round your answer to the nearest **hundredth** (i.e.: 1.234 would round to 1.23).