Practice Final G - Solutions

1 Spock's shorts

1. (a) (2 points) We are interested in whether some sentence α in first-order logic entails some other sentence β. Does the answer to that question depend on what we think the symbols in α and β meant when we wrote the sentences down? Explain.

   Solution: No. Entailment means that *all* models of α are also a model of β. The set of all models includes all the different possible meanings of the symbols.

   (b) (2 points) If we apply the propositional resolution rule to **two clauses** $\{P\}$ and $\{\neg P\}$, what clause results? Explain its consequences for a resolution-refutation proof.

   Solution: It results in False. This is used in the final step of resolution-refutation in proving $\alpha \models \beta$ to conclude that our assumptions $\alpha \wedge \neg\beta$ entail False. We can then conclude that $\alpha \models \beta$.

   (c) (3 points) Imagine we have the **single clause** $\{P, \neg P\}$ as one of our assumptions in a resolution refutation proof. When, if ever, would it be a good idea use this clause in a resolution step? Explain in technical detail with an example.

   Solution: It's never a good idea, since it never does anything. Intuitively, the clause $\{P, \neg P\}$ is equivalent to True, so doing a resolution step doesn't do anything.
   Concretely: if we apply resolution to the clauses $\{P, \neg P\}$ and $\{P, Q\}$, then we get $\{P, Q\}$, which we already had. It's similar if we had $\{\neg P, Q\}$.

(d) (2 points) We can't use the usual mechanisms of proof to show that some sentence α *does not* entail some sentence β. If we do want to prove that α doesn't entail β, what should we do? Explain your answer.

- ◯ Find an interpretation that is a model of α but not a model of β.
- ◯ Find an interpretation that is a model of β but not a model of α.
- ◯ Show that there are no interpretations that are models of both α and β.

Solution: We should find an interpretation that is a model of α but not a model of β. By definition, α entails β if and only if $M(\alpha) \subseteq M(\beta)$, so finding such an interpretation would imply that α does not entail β.

(e) (4 points) Provide an interpretation that is a model of

$$\alpha : \forall x.\exists y.R(x, y)$$

but not a model of

$$\beta : \exists x.\forall y.R(x, y)$$

Recall that to specify an interpretation for these sentences, you need to supply the universe U and interpretation of the predicate symbol R.

Solution: Let U be the set of all integers. Let $R(x, y)$ be $x > y$. This is a model of α because for all integers x, there exists a smaller integer y. This is not a model of β because there does not exist an integer x that is larger than all other integers y.

(f) (4 points) Let $\gamma : \forall z. \exists w. R(w, z)$. Let's use resolution refutation to prove that β from the previous question entails γ.

i. Convert β and $\neg\gamma$ to clausal form and write the clauses down here.

> **Solution:** 1. $R(f1, y)$
> 2. $\neg R(w, f2)$

ii. Complete the proof, indicating what lines you resolve and what substitutions are necessary.

> **Solution:** 3. False 1, 2 , $\{w : f1, y : f2\}$

(g) (8 points) Consider a variant of the Search and Rescue problem from HW4 with N cells labeled $1, \ldots, N$. We'll assume the world works as follows.

1. Each location has **exactly one** of {smoke, fire}. (Oh no!)
2. There is smoke at a location if and only if there is a fire in at least one of the adjacent locations.

We'll use first-order logic to formalize this problem, using the following symbols and intended meanings:

- constants $1, 2, \ldots, N$: location indices
- predicate $\mathit{Smoke}(i)$: whether location i has smoke
- predicate $\mathit{Fire}(i)$: whether location i has fire
- predicate $\mathit{Adj}(i, j)$: whether location i is adjacent to location j

Express each of the following axioms in first-order logic:

i. Location 1 has smoke.

Solution:

$$\mathit{Smoke}(1)$$

ii. Adjacency is symmetric. That is, two cells i and j are adjacent if and only if j and i are adjacent.

Solution:

$$\forall i. \forall j. \mathit{Adj}(i, j) \iff \mathit{Adj}(j, i)$$

iii. Each location has **exactly one** of {smoke, fire}.

Solution:

$$\forall i. (\mathit{Fire}(i) \lor \mathit{Smoke}(i)) \land (\neg \mathit{Fire}(i) \lor \neg \mathit{Smoke}(i))$$

iv. There is smoke at a location if and only if there is a fire in at least one of the adjacent locations.

Solution:

$$\forall j. (\mathit{Smoke}(j) \iff (\exists k. \mathit{Fire}(k) \land \mathit{Adj}(j, k))$$

2 Oration

2. (20 points) Your professor asks you to organize a seminar series with 4 talks on different dates (1, 2, 3, and 4).

- The following constraints must be satisfied
  - C1. No speaker can talk on more than one date.
  - C2. Each speaker can speak on a particular subset of the topics.
  - C3. No two consecutive seminars can be on the same topic.
  - C4. No two consecutive speakers can be from the same university.
  - C5. Each date has exactly one speaker.
  - C6. Each date has exactly one topic.
- The three topics are: theory, modeling, and implementation (T, M, I).
- The speakers are: Antonius, Bibulus, Claudius, Dracontius, and Emporius (A, B, C, D, E).
  - Antonius can speak about theory or modeling.
  - Bibulus can speak about theory, modeling, or implementation.
  - Claudius can speak about modeling or implementation.
  - Dracontius can speak about theory or implementation.
  - Emporius can speak about implementation.
  - Antonius and Emporius are from Romulus University.
  - Bibulus, Claudius, and Dracontius are from Remus University.

(a) Consider a formulation with variables S_d indicating the speaker for each date d, and variables T_d indicating the topic for each date d.

i. What is the value domain for the variables, taking any unary constraints into account?

> **Solution:** For S_d it's (A, B, C, D, E). For T_d it's (T, M, I).

ii. Which of the constraints are automatically satisfied because of the formulation? Mark all that apply. ○ C1 ○ C2 ○ C3 ○ C4 √ **C5** √ **C6**

(b) Consider a formulation with variables D_s indicating a date for each speaker and T_s indicating a topic for each speaker.

i. What are the value domains for the variables, taking any unary constraints into account?

> **Solution:** For D it's $(1, 2, 3, 4, None)$. For T_s it's different for each speaker: $(T, M), (T, M, I), (M, I), (T, I), (I)$.

ii. Which of the constraints are automatically satisfied because of the formulation? Mark all that apply. √ **C1** √ **C2** ○ C3 ○ C4 ○ C5 ○ C6

(c) Consider a formulation with variables ST_d indicating both the speaker and topic for each date.

i. What is the value domain for the variables, taking any unary constraints into account?

> **Solution:** $\{(A, T), (A, M), (B, T), (B, M), (B, I), (C, M), (C, I), (D, T), (D, I), (E, I)\}$

ii. Which of the constraints are automatically satisfied because of the formulation? Mark all that apply. ○ C1 √ **C2** ○ C3 ○ C4 √ **C5** √ **C6**

(d) Which of the three previous formulations do you think would be most effective for a standard CSP solver, and why?

> **Solution:** The last one because it has the smallest total set of assignments and the most constraints built in. Also accepted: the first one because the constraints are already binary.

(e) In the formulation from part (a), all the constraints are binary. Here is a constraint diagram. Label each type of arc in the diagram with the constraint (1, 2, 3, 4, 5, 6) it represents.

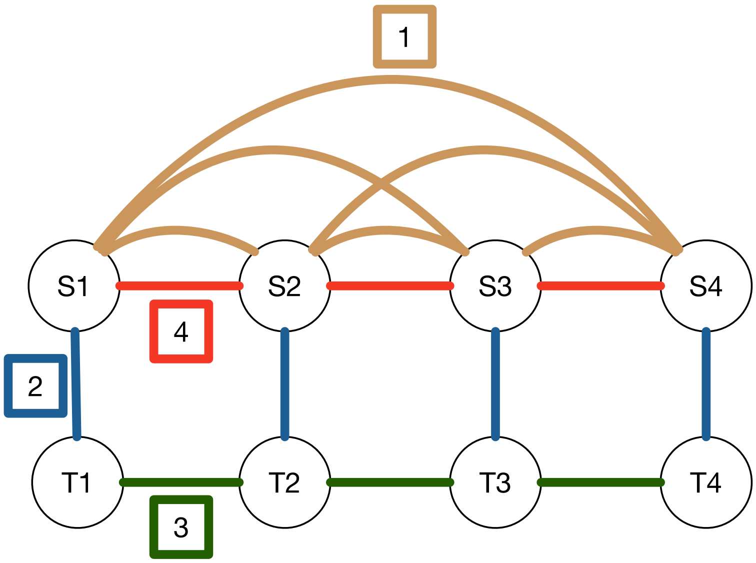

(f) Assume we are using the formulation from part (a). We have already chosen variable S_1 and assigned it to E. After doing forward checking, which variable will be chosen, using the most-constrained-variable heuristic, and what is its domain after forward checking?

> **Solution:** Variable T_1, which has unit domain $\{I\}$.

3 Catching the bus

3. (20 points) Busta Rhydes needs to catch a bus in an infinite 2-d grid (from $-\infty$ to $+\infty$ in both directions). They know the bus's schedule, and want to plan a path such that they end up in the same grid location at the same time as the bus.

   Assume that the search starts at time 0 with Busta in location (0,0). On each step, time increases by 1, and Busta can move to any of the **eight** neighboring locations, or remain in the same location.

   The position of the bus is specified by a function B which takes a time as input and returns the bus's location at that time as a tuple `(r,c)`.

   (a) For the following implementations of the bus's schedule, what is one sequence of locations (representing Busta's location at time 0, 1, 2...) that might result from running a breadth-first search in this domain? Enter `None` if no path is returned, either because BFS runs indefinitely, or because it returns without having found a path. Note that there may be more than one correct answer, but you need only find one.

   i. $B_1(t) = (1, t-3)$

   > **Solution:** Busta Path (t,r,c) = (0,0,0),(1,0,0),(2,1,-1) (Other solutions are possible)
   > Bus Path (t,r,c) = (0,1,-3),(1,1,-2),(2,1,-1)

   ii. $B_2(t)$ defines a counterclockwise loop, shaped like a square of size 10 with its lower-left corner at (0,0). Bus starts at (10,0) at time 0 and advances by **ten** grid squares each time step, i.e. (10,0), (10,10), (0, 10), ...

   > **Solution:** Busta Path (t,r,c) = (0,0,0),(1,0,0),(2,0,0),(3,0,0) (Other solutions are possible)
   > Bus Path (t,r,c) = (0,10,0),(1,10,10),(2,0,10),(3,0,0)

   (b) To formalize this type of problem for a heuristic search method, for any $B(t)$:

   i. What would be a minimal definition of the state?

   > **Solution:** $s = (t, r, c)$ where t is time and r, c is Busta's location. We can infer bus position as $br, bc = B(t)$.
   > Note that the bus position is not necessary and, also, the bus position at time t does not determine the bus position at time $t+1$ and so you need to remember the time in the state.

   ii. Specify a goal test $g(s)$ using your state representation that works for any $B(t)$.

   > **Solution:** $t, r, c = s$. If $br(t), bc(t)$ is bus location $B(t)$, then goal test is $br(t) = r$ and $bc(t) = c$.

   (c) What is a useful admissible heuristic function for this problem, given you know the bus schedule is (arbitrary) $B(t)$ and that the bus moves at most K grid squares per time step? For this question, we'll say the heuristic is *useful* if it is 0 only at the goal state and is greater than 1 at at least one state. Explain your answer.

   > **Solution:** Assuming state $s = (t, r, c)$ encodes Busta's location and $br(t), bc(t) = B(t)$ is the bus location, then $h(s) = \max(|br(t) - r|/(K+2), |bc(t) - c|/(K+2))$ is an admissible heuristic. The $K+2$ in the denominator is the maximum Manhattan displacement per step between the bus and Busta. We gave credit for different forms that had similar idea.

4 Frogger

4. (20 points) You are a small frog, in a 1D world, with states s in the interval $[0, 1]$. You can take one of four possible actions:

   - *big right (BR)* which moves you to a new location drawn from a Gaussian distribution with mean $s + .2$ and standard deviation 0.1.
   - *small right (SR)* which moves you to a new location drawn from a Gaussian distribution with mean $s + .1$ and standard deviation 0.01.
   - *big left (BL)* which moves you to a new location drawn from a Gaussian distribution with mean $s - .2$ and standard deviation 0.1.
   - *small left (SL)* which moves you to a new location drawn from a Gaussian distribution with mean $s - .1$ and standard deviation 0.01.

   However, any action result that would move you to a state less than 0.0 instead puts you at 0.0 and any action result that would move to a state greater than 1.0 instead puts you at 1.0.

   If you are in the *goal region* then any action yields a reward of 0 and the game terminates. All other actions have reward -1. The goal region will have slightly different definitions in the parts below.

   There is no discounting ($\gamma = 1.0$)

   (a) Easy Frogger: the *goal region* is $[0.8, 1.0]$.

   On the axes below, with the x axis corresponding to the state s and the y axis corresponding to the expected optimal value, plot $Q(s, \text{BR})$, $Q(s, \text{SR})$, $Q(s, \text{BL})$, and $Q(s, \text{SL})$.

   Clearly label the y axis and indicate which curve you draw corresponds to which action.

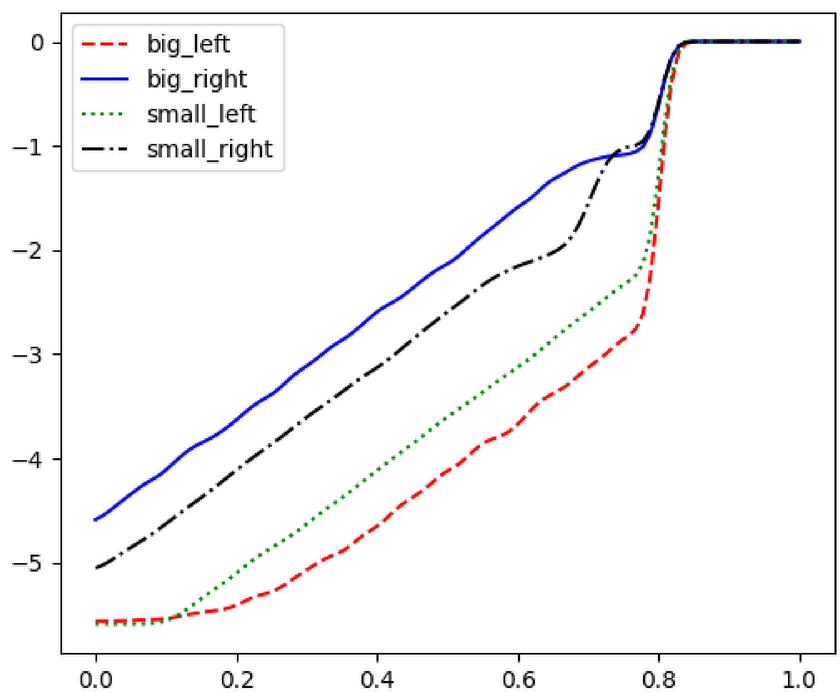

   (b) Medium Frogger: the *goal region* is now $[0.8, 0.9]$.

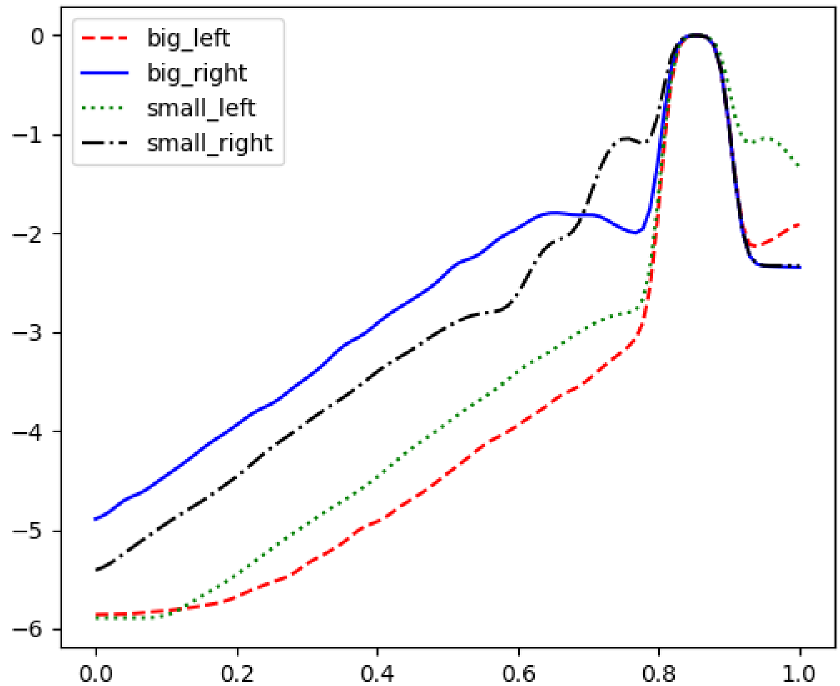

i. What is the best action to take at $s = 0.65$? √ **BR** ◯ SR ◯ BL ◯ SL ◯ all equivalent

ii. What is the best action to take at $s = 0.75$? ◯ BR √ **SR** ◯ BL ◯ SL ◯ all equivalent

iii. What is the best action to take at $s = 0.85$? ◯ BR ◯ SR ◯ BL ◯ SL √ **all equivalent**

iv. What is the best action to take at $s = 0.95$? ◯ BR ◯ SR ◯ BL √ **SL** ◯ all equivalent

v. Explain your answers.

Solution:

(c) Hard Frogger: the *goal region* is $[0.8, 0.9]$, as above, but we add a region of mud $[.2, .4]$. Inside the muddy region, no matter what action you select, the transition will be *as if* you had selected action BR.

The Q value functions are shown below.

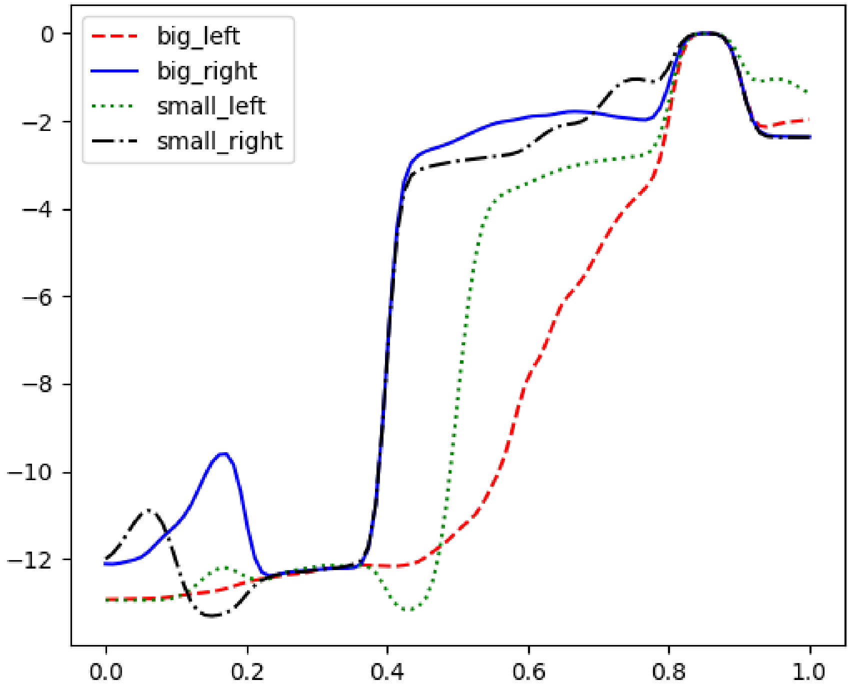

Specify the optimal policy that would be derived from these Q functions. You can write it as a math expression or an informal piece of code. Don't worry about getting the numeric parts exactly right.

Solution:

5 Flames or fortune

5. Consider the following gridworld containing a robot, a large positive reward (the money, worth +1000 points) and a large negative reward (the flames, worth -1000, points). The robot is trying to obtain the money and avoid the flames.

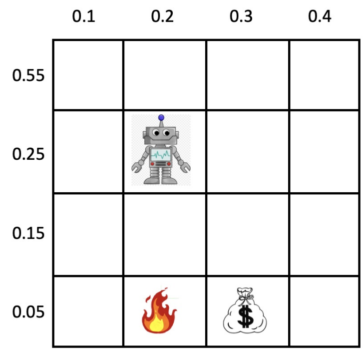

The robot's orientation is fixed, but it can move up, down, left and right, and each action that does not result in fortune or flame costs -1 points. However, it cannot directly observe its (x, y) position — it can only observe whether there are any walls next to it in the four compass directions. For instance, if the robot were in the top right corner, it would see a wall to the north and to the east. It would also know unambiguously that it is in that cell — no other cells have a wall to the north and to the east.

There are only walls around the boundary of the gridworld — there are no internal walls. The problem ends when the robot obtains the money or is consumed by flames.

Let us also assume the robot's motion dynamics are perfect — when it attempts to move in a given direction (e.g., move up), the motion occurs deterministically, unless there is a wall in the way. Let us also assume that the robot's sensing is perfect — it always sees the walls it is adjacent to, and never sees walls that aren't there. Recall that on any given step, *the robot moves then observes*, so it will start with a motion.

Finally, assume the robot's (unobservable) start location is the one drawn in the figure, at $(1, 2)$. (We will use the x, y coordinate system convention, so the bottom-left corner is $(0, 0)$.)

The robot will have to select an action based on its initial belief, *then* it will receive an observation, etc.

(a) (3 points) For the purposes of this part, assume the robot starts out with an initial belief that is uniform over all 16 locations. How many distinct belief states are reachable, including the start belief state? (Recall that a belief state is a distribution over the underlying states.)

Solution: 22 belief states. 1 uniform belief, 1 uniform belief over the four internal squares, 4 beliefs where one of the coordinates is known exactly and the other coordinate is distributed uniformly between two of the edge cells, and 16 beliefs where both coordinates are known exactly.

20 belief states is also acceptable if we don't count the belief that results when the robot moves into the positive or negative reward squares and the game ends.

Now, assume that when the robot wakes up, it doesn't know where it is, but it does have a prior on its location indicated by the marginal distribution on the columns $(.1, .2, .3, .4)$ and on the rows $(.55, .25, .15, .05)$ as indicated in the figure.

(b) (2 points) What action would the *most likely state* method select initially?

> **Solution:** Down

(c) (4 points) What would the posterior belief state be after the robot (in its actual location as shown in the figure) executed a *down* action and got the observation corresponding to its resulting actual state?

> **Solution:**
> ```
> 0 0 0 0
> 0 x y 0
> 0 z w 0
> 0 0 0 0
> ```
> (x, y, z, w) = normalized (.11, .165, .05, .075)

(d) (2 points) What action would the *most likely state* method select from this belief?

> **Solution:** down

Now let's consider the *most likely observation* (MLO) strategy, in which we search in a determinized version of belief space, assuming that we will always get the most likely observation after each action.

In all the following questions, assume we start with belief b_0:

$$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & .3 & .4 & 0 \\ 0 & .1 & .2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

(e) (2 points) What is the most likely observation after we move *left*? Explain why.

> **Solution:** no walls in any direction, because the .6 of our probability mass will be in states that generate that observation

(f) (2 points) In the MLO belief-space search, what is the one-step reward for moving *down* from belief b_0?

> **Solution:** .1 * -1000 + .2 * 1000 + .7 * -1

(g) (3 points) In the MLO belief-space search, what is an optimal open-loop plan from b_0, and its approximate value (to the nearest 100)?

> **Solution:** left, right, down, down
>
> value approx 1000

(h) (2 points) Given the robot's *actual* initial state, will the robot end up executing the whole action sequence found in the previous step? Explain.

> **Solution:** No. On its first step it will get an observation that is not the MLO, so it will need to replan.

6 Where was I?

6. You have been moving around a tiny world with four states: (0, 0), (0, 1), (1, 0), and (1, 1):

   We have the following probabilistic transition model:

S_t \ S_{t+1}	(0, 0)	(0, 1)	(1, 0)	(1, 1)
(0, 0)	0	$\frac{1}{2}$	$\frac{1}{2}$	0
(0, 1)	1	0	0	0
(1, 0)	$\frac{1}{3}$	0	$\frac{1}{3}$	$\frac{1}{3}$
(1, 1)	0	0	$\frac{1}{2}$	$\frac{1}{2}$

   Observation model:

S_t \ O_t	Red	Green	Blue	Gray
(0, 0)	$\frac{1}{2}$	$\frac{1}{2}$	0	0
(0, 1)	0	0	1	0
(1, 0)	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
(1, 1)	0	0	$\frac{1}{2}$	$\frac{1}{2}$

   And starting state distribution:

S_0	$P(S_0)$
(0, 0)	0
(0, 1)	0
(1, 0)	0
(1, 1)	1

   You moved four times, winding up in some state described by random variable S_4. During that time, you made two observations: $O_2 = \textit{red}$ and $O_4 = \textit{gray}$.

   You love red flowers, and really want to know where you were when you observed *red*. Now, using all the information you have, you would like to compute $P(S_2 \mid O_2 = \textit{red}, O_4 = \textit{gray})$.

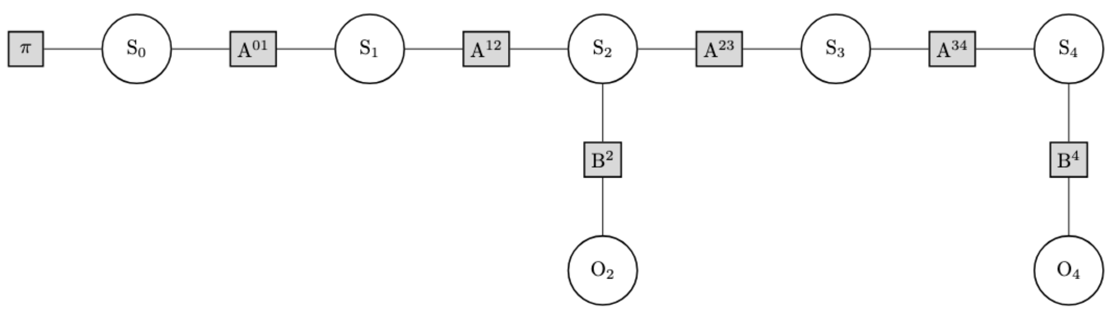

   We will think step by step, using sum-product message passing using the labeled figure above.

   (a) (4 points) Shakey the robot thinks that we have to pass messages from S_0 all the way over to S_4, and then back again, before we can compute the quantity of interest. Rosie the robot disagrees and says there's a more efficient method. Who is right, and why?

   ○ Shakey ○ Rosie

   Solution: Rosie is correct. Because we just want to compute $P(S_2 \mid O_2 = red, O_5 = gray)$, we can use sum-product message passing where the root node is S_2. This doesn't require us to pass messages from S_0 to S_4. Instead, we pass messages from S_0 to S_2 and from S_4 to S_2.

(b) (15 points) In sum-product message-passing on this graph, given that $O_2 = \textit{red}$ and $O_4 = \textit{grey}$, what are the following messages?

- Note that you may want to think about the observations as adding factors to the graph.
- Please answer with numerical expressions (but you don't have to do multiplication/addition/simplification).

We've filled in some for you.

i. $\mu_{S_0 \to A^{01}}$

Solution: $[0, 0, 0, 1]$

ii. $\mu_{A^{01} \to S_1}$

Solution: $[0, 0, 1/2, 1/2]$

iii. $\mu_{S_1 \to A^{12}}$

Solution: $[0, 0, 1/2, 1/2]$

iv. $\mu_{A^{12} \to S_2} = (1/6, 0, 5/12, 5/12)$

v. $\mu_{O_4 \to B^4}$

> **Solution:** $[0, 0, 0, 1]$

vi. $\mu_{B^4 \to S_4}$

> **Solution:** $[0, 0, 1/4, 1/2]$

vii. $\mu_{S_4 \to A^{34}} = \mu_{B^4 \to S_4}$

viii. $\mu_{A^{34} \to S_3}$

> **Solution:** $(1/8, 0, 1/4, 3/8)$

ix. $\mu_{S_3 \to A^{23}} = \mu_{A^{34} \to S_3}$

x. $\mu_{A^{23} \to S_2} = (1/8, 1/8, 1/4, 5/16)$

xi. $\mu_{B^2 \to S_2} = (1/2, 0, 1/4, 0)$

(c) (3 points) Express $P(S_2)$ (without conditioning on any observations) symbolically in terms of *one of* the messages above.

> **Solution:** $\mathit{normalize}(\mu_{A^{12} \to S_2})$

(d) (3 points) Express $P(S_2 \mid O_2 = \mathit{red},\ O_4 = \mathit{gray})$ symbolically in terms of some of the messages above.

> **Solution:**
> $$\mathit{normalize}(\mu_{A^{12} \to S_2} \cdot \mu_{A^{23} \to S_2} \cdot \mu_{B^2 \to S_2})$$