COMPSCI 370 **Instructor:** Ronald Parr

Homework 3 **Due:** Monday, April 7, 2025

Updated 3/29

1 Value of Information (15 points)

Consider the traffic information question from the slides. In the slides, we assume that we are told the true value of the traffic variable T. In greater generality, however, we may only get a traffic report R which is correlated with T. In this context, we can think of the case in the slides as corresponding to:

$$\begin{aligned} P(r|t) &= 1.0 \\ P(\overline{r}|\overline{t}) &= 1.0 \end{aligned}$$

Assume $p(r|t) = 1.0$, but let $P(\overline{r}|\overline{t})$ be some free variable $0 \leq x \leq 1$. Ax x decreases, you should expect the value of information to also decrease. Solve for the crossover point where the value of information is no longer positive and becomes 0. (Another way to think about this is to find the largest value of x where VOI is 0.)

2 MDPs (20 points)

Consider the 3-state MDP shown in the image below.

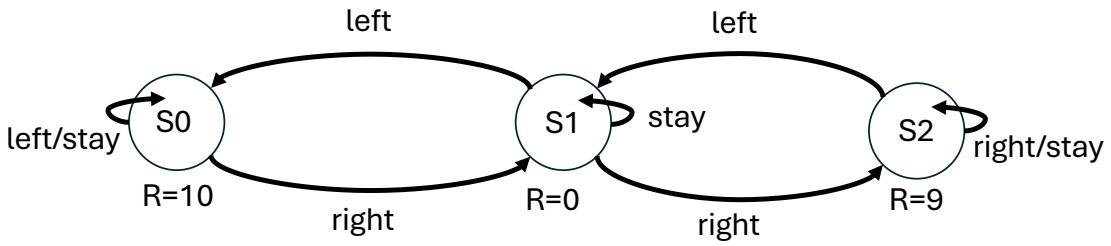

Each state has 3 possible actions, left, right, and stay. All actions are deterministic, so arcs are not annotated with probabilities. We will refer to polices using a sequence of three letters. For example, the SRL, would correspond to taking stay action in state 0, the right action in state 1, and the left action in state 2.

a) Using a discount factor of $\gamma = 0.95$, and an initial value function that is all zeros, show three steps of value iteration. (More specifically: Start with $V_0(S0) = V_0(S1) = V_0(S2) = 0$, and show

V_1, V_2, and V_3. (5 points)

b) Solve for the exact value of the policy SLL with discount factor $\gamma = 0.95$. (5 points)

c) Observe that left and stay are basically the same action in $S0$, and right and stay are basically the same action in $S2$. Assuming that ties are always broken in favor of the stay action, prove the SLL is the optimal policy for $\gamma = 0.95$. Hint: You can use the termination condition for policy iteration. (5 points)

d) It's easy to see that for any discount factor, the optimal action in $S0$ will always be S (equivalent to L) and that the optimal policy in $S1$ will always be L. Solve for the value of γ where the optimal policy for $S2$ becomes L. (5 points)

3 Bayes Nets (20 points)

Our Bayes net question will use the following network. Note that all variables are binary:

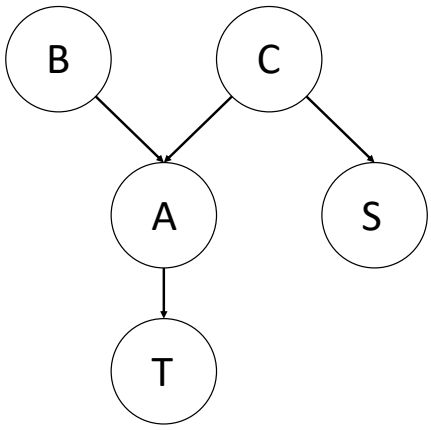

We motivate this structure with the following scenario: You have a crude burglar alarm installed at your house (A). Two things can cause the alarm to go off, a burglar (B) or vibration from construction in your neighborhood (C). If your alarm is ringing, you may get a text from your neighbor (T). If there is construction in your neighborhood, there might be a post about the construction on social media (S).

This network has the following CPTs. Note that the CPTs have probabilities that are convenient, not necessarily realistic:

$$P(b) = 1/8$$

$$P(c) = 1/4$$

B	C	$P(a\|BC)$
$\overline{b}$	$\overline{c}$	$1/8$
$\overline{b}$	c	$3/8$
b	$\overline{c}$	$7/8$
b	c	1

C	$P(s\|C)$
$\overline{c}$	$1/8$
c	$3/4$

A	$P(t\|A)$
$\overline{a}$	$1/8$
a	$3/4$

For the following questions, use variable elimination to compute the desired quantities. Show your work and any intermediate results computed explicitly.

a) Compute the marginal probability: $P(t)$. (10 points)

b) Compute $P(b|t)$, the probability of burglary given that you got a text. (You can do this by computing $P(bt)$). Note that you should expect this to be higher than the prior probability of burglary, $P(b) = 1/8$. (10 points)