1. Agents and Environments

True False

- ❑ ❑ There exists (at least) one environment in which every agent is rational.
- ❑ ❑ For every agent, there exists (at least) one environment in which the agent is rational.
- ❑ ❑ To solve the sliding-tile 15-puzzle, an optimal agent that searches will usually require less memory than an optimal table-lookup reflex agent.
- ❑ ❑ To solve the sliding-tile 15-puzzle, an agent that searches will always do better (=find shorter paths) than a table-lookup reflex agent.

2. A* Search

For heuristic function h and action cost 10 (per step), enter into each node the order (1,2,3,…) when the node is expanded (=removed from queue). Start with "1" at start state at the top. Enter "0" if a node will never be expanded.

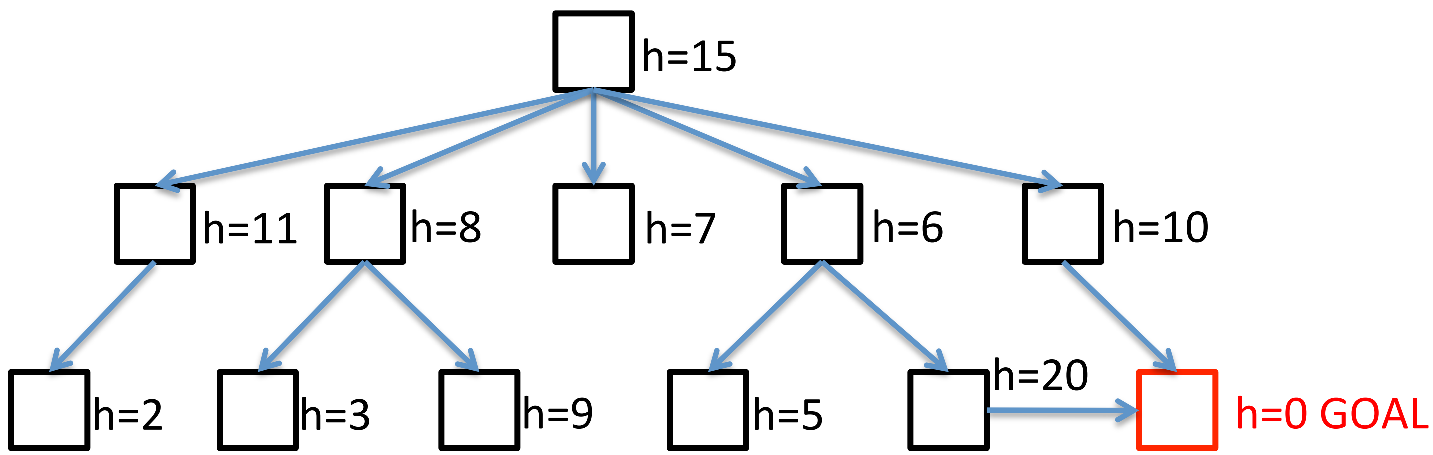

Is the heuristic h admissible? ○ Yes ○ No

3. Probability I

For a coin X, we know P(heads) = 0.3

What is P(tails)?

4. Probability II

Given a potentially loaded (=unfair) coin, which we flip twice. Say the probability for it coming up heads both times is 0.04. These are independent experiments with the same coin.

What is the probability it comes up tails twice, if we flip it twice?

5. Probability III

We have two coins, one fair P_1 (heads)=0.5 and one loaded with P_2 (heads)=1. We now pick a coin at random with 0.5 chance.

- We flip this coin, and see "heads". What is the probability this is the loaded coin?

- We now flip this coin again (the same coin) and see "heads". What is now the probability this is the loaded coin?

6. Bayes Network I

Consider the following Bayes network. True or false?

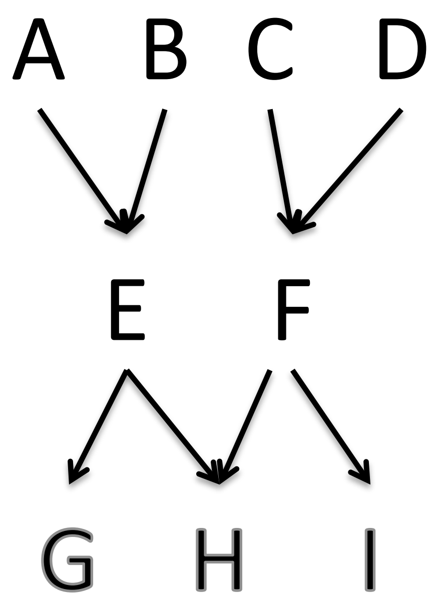

	TRUE	FALSE
$A \perp B$		
$A \perp B \mid E$		
$A \perp B \mid G$		
$A \perp B \mid F$		
$A \perp C \mid G$		

7. Bayes Network II

Given this Bayes network

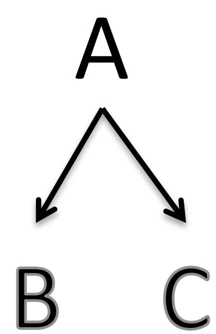

$$P(A) = 0.5$$

$$P(B \mid A) = 0.2$$

$$P(B \mid \neg A) = 0.2$$

$$P(C \mid A) = 0.8$$

$$P(C \mid \neg A) = 0.4$$

Calculate

$$P(B \mid C) = \square$$

$$P(C \mid B) = \square$$

8. Naïve Bayes with Laplacian Smoothing

We have two classes of movies, new and old

OLD	NEW
Top Gun	Top Gear
Shy People	Gun Shy
Top Hat	

Using Laplacian smoothing (k=1), compute

P(OLD) =	
P("Top" | OLD) =	
P(OLD | "Top") =	

Use a single dictionary for smoothing. Think of "Top" as a word and as a single-word new movie title

9. K-Nearest Neighbor

Given the following labeled data set

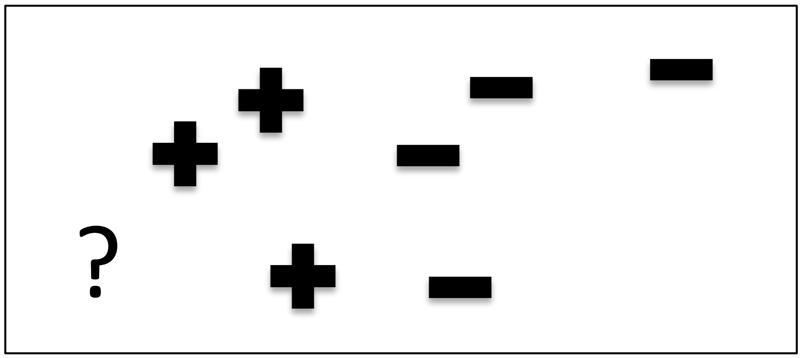

For what (minimal) value of k will the query point "?" be negative? Enter "0" if this is impossible. Ties are broken at random – try to avoid them.

10. Linear Regression

We have the following data:

x	y
1	2
3	5.2
4	6.8
5	8.4
9	14.8

Perform linear regression: $y = w_1x + w_0$

What is w_1 ? [ ] What is w_0 ? [ ]

11. K-Means Clustering

For the following data set (solid dots) with initial cluster centers C1, C2: What will be the final location of C1 after running K-Means (A, B,C or D).

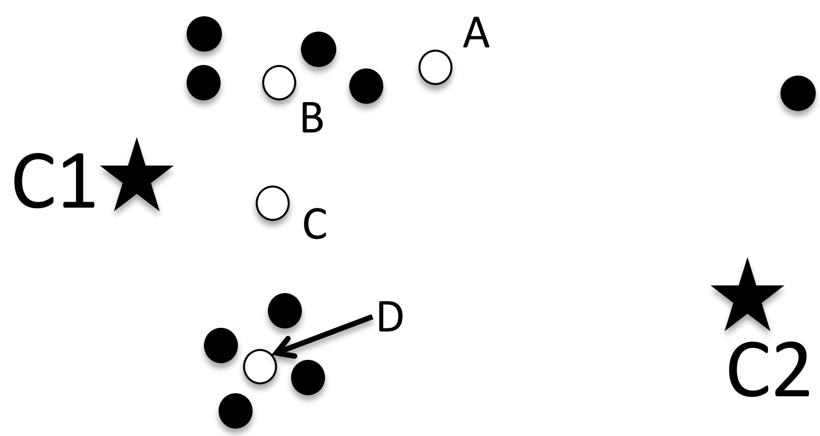

12. Logic

Mark each sentence as Valid (always true), Satisfiable but not Valid, or Unsatisfiable

Valid	Sat.	Unsat.	
○	○	○	$\neg A$
○	○	○	$A \lor \neg A$
○	○	○	$(A \land \neg A) \Rightarrow (B \Rightarrow C)$
○	○	○	$(A \Rightarrow B) \land (B \Rightarrow C) \land (C \Rightarrow A)$
○	○	○	$(A \Rightarrow B) \land \neg(\neg A \lor B)$
○	○	○	$((A \Rightarrow B) \land (B \Rightarrow C)) \Leftrightarrow (A \Rightarrow C)$

13. Planning

In the state space below, we can travel between locations S, A, B, D, and G along roads (SA means go from S to A). But the world is partially observable and stochastic: there may be a stop light (observable only at B) that prevents passing from B to G, and there may be a flood (observable only at A) that prevents passing from A to G. If the flood occurs it will always remain flooded; if stop light is on it will always switch off at some point in the future. For each plan, starting at S, click if it it ***A****lways reaches G in a bounded number of steps; always reaches G after an* ***U****nbounded number of steps; or* ***M****aybe reaches G and maybe fails.*

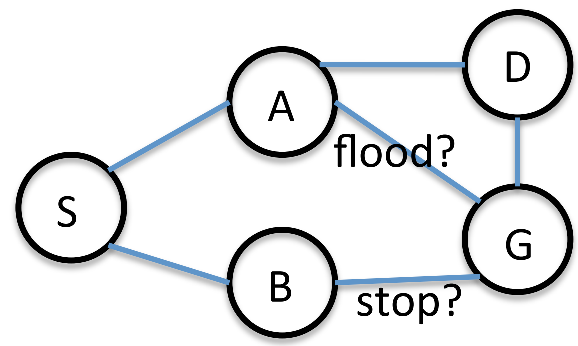

Always bounded	Always unbounded	Maybe (may fail)	
o	o	o	[SA, AG]
o	o	o	[SB, 2:(if stop: goto 2), BG]
o	o	o	[SA, 2:(if flood: goto 2), AG]
o	o	o	[SA, 2:(if flood: [AD, DG] else: AG)]
o	o	o	[SB, 2:(if stop: [BS, SA, AD, DG]), BG]
o	o	o	[SA, AD, DG]

14. Markov Decision Process (MDP)

Deterministic state transitions. Cost of motion is -5. Terminal state is 100. Actions are N/S/W/E. Shaded state can't be entered. Fill in the final values after value iteration.

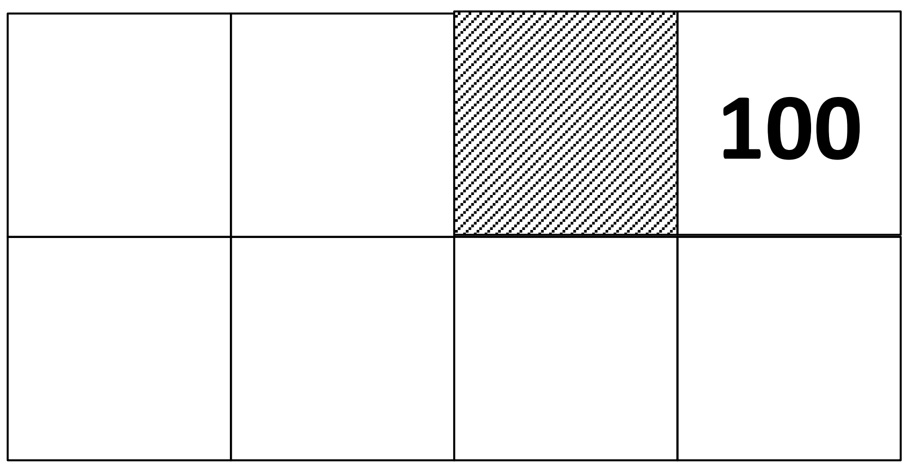

15. Markov Chains

Use Laplacian smoothing with k=1, to learn the parameters of this Markov Chain model from the observed state sequence:

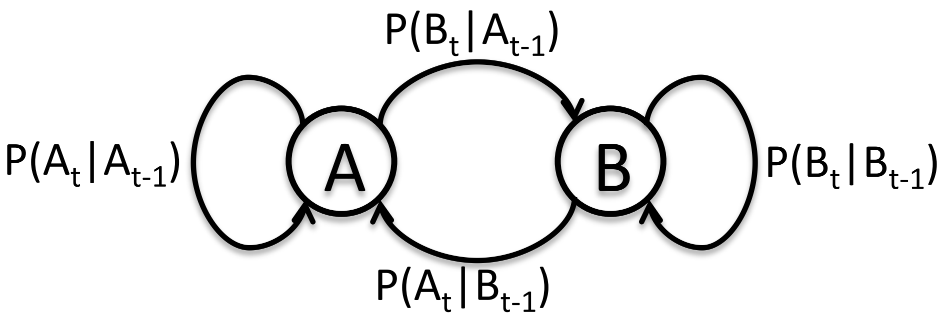

A A A A B

Initial state distribution	$P(A_0) =$	
Transition distribution	$P(A_t|A_{t-1}) =$	
	$P(A_t|B_{t-1}) =$	