CS221 Practice Midterm #2

Summer 2013

The following pages are excerpts from similar classes' midterms. The content is similar to our midterm – but I have opted to give you a document with more problems rather than one that follows the structure of the midterm precisely. See the midterm handout for more details on what the exam will look like. The midterm is 2 hours. It is open book and open computer but **closed** Internet.

For this practice midterm, the short answers have been omitted. You will have a set of short answers on the real midterm (like in practice midterm #1).

2. [Deterministic Search] Pacfamily (20 points)

Pacman is trying eat all the dots, but he now has the help of his family! There are initially k dots, at positions (f1, ...fk). There are also n Pac-People, at positions (p1, ..., pn); initially, all the Pac-People start in the bottom left corner of the maze. Consider a search problem in which all Pac-People move simultaneously; that is, in each step each Pac-Person moves into some adjacent position (N, S, E, or W, no STOP). Note that any number of Pac-People may occupy the same position.

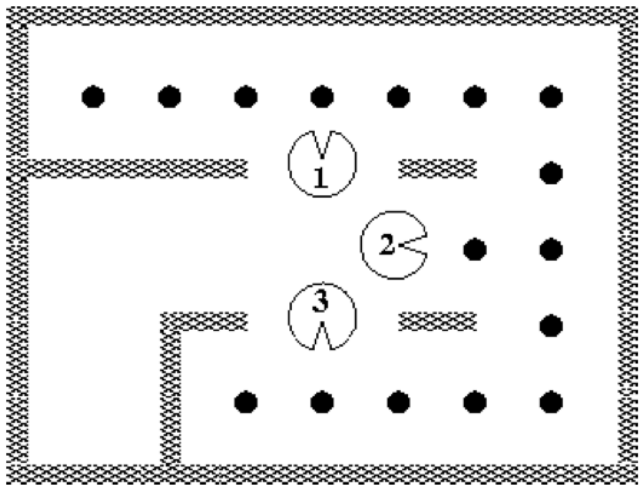

(a) Define the state space of the search problem.

(b) Give a reasonable upper bound on the size of the state space for a general *r* by *c* grid.

(c) What is the goal test?

(d) What is the maximum branching factor of the successor function in a general grid?

(e) Circle the admissible heuristics below (-1/2 point for each mistake.)

h1(s)=0 Solution

h2(s) = 1 Solution :

h3(s) = number of remaining food / n

h4(s) = maxi maxj manhattan(pi, foodj)

h5(s) = maxi minj manhattan(pi, foodj)

3. [Adversarial Search] Expectimax Car (15 points)

The goal of this problem is to extend the self driving car to reason about the future and use that reasoning to make a utility maximizing decision on how to act.

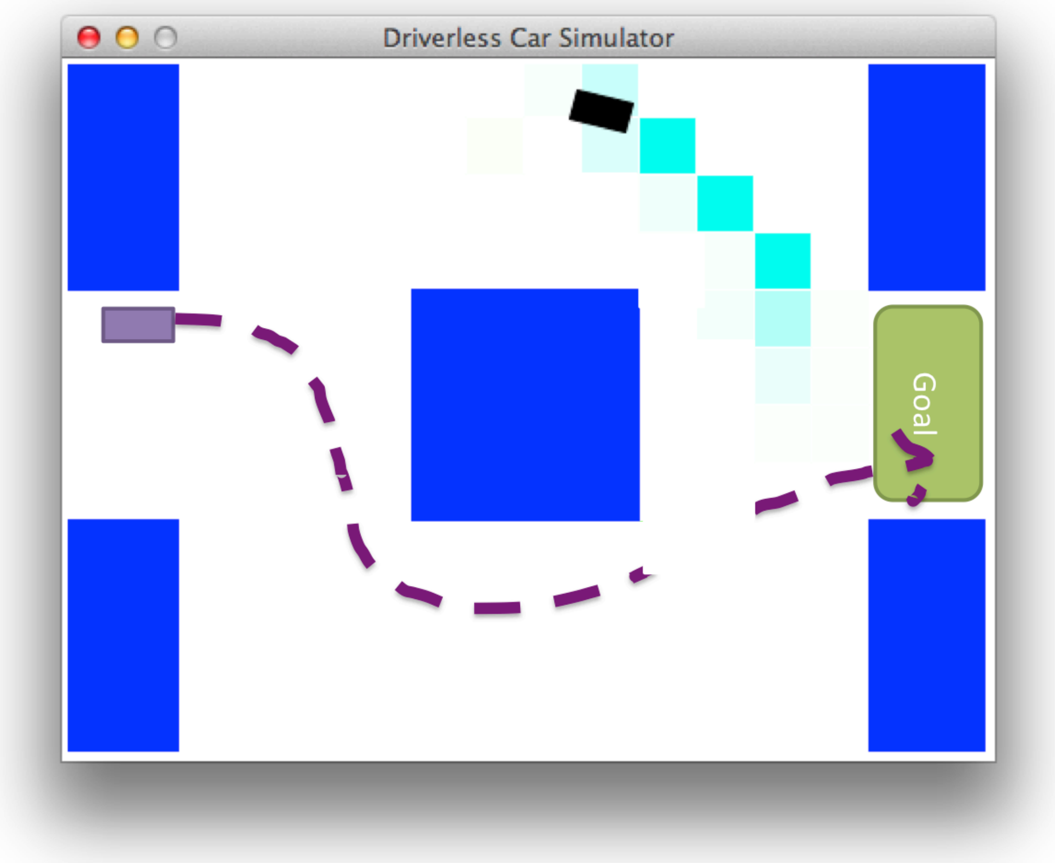

In this problem we are going to assume that the world is the same as in the Driverless Car programming problem. There is a single agent that exists in a closed world and wants to drive to a goal area. For this problem we are going to assume that there is only one other car.

Each heartbeat (there are twenty per second) each car can perform one of four actions {Accelerate, TurnLeft, TurnRight, Brake, None}. Assume that for each tile we have a probability distribution that the other car will take each of those actions. A car always has a wheel position and a velocity. Even though the car might not be taking the Accelerate action, if it has positive velocity it may continue to move forward.

[important] For this problem assume that cars are manufactured so precisely that if they take an action there is no uncertainty as to how the car will respond. Also assume that you know the start state of both your car and the other.

Formalize the problem of choosing an action as a markov decision problem:

(a) What variables make up a state:

(b) What is the start state?

(c) For each state what are the legal actions.

(d) What is a terminal condition for a given state *S*? For all terminal conditions specify a utility.

(e) Given a state *S* from which your agent took action *A*, what is the successor state distribution?

(f) We are going to solve this problem using expectimax. However, we may not want to expand the entire tree. Give a reasonable heuristic utility for a given state *S.*

4. [Bayes Net] Snuffles (20 points)

Assume there are two types of conditions: (S)inus congestion and (F)lu. Sinus congestion 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.

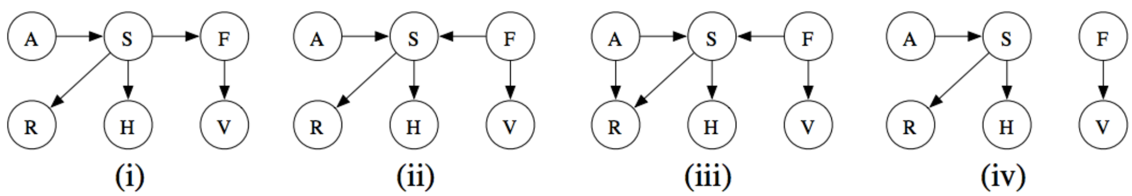

(a) Consider the four Bayes Nets shown. Circle the one which models the domain (as described above) best.

(b) 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.

(c) 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) In the original network you chose, which query is easier to answer: P (F | r, v, h, a, s) or P (F)? Briefly justify.

(e) Assume the following samples were drawn from prior sampling:

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

Give the sample estimate of P(f) or state why it cannot be computed.

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

5. [Temporal Models] The Jabberwok (25 points)

You have been put in charge of a Jabberwock for your friend Lewis. The Jabberwock is kept in a large tugleywood which is conveniently divided into an N × N grid. It wanders freely around the N^2 possible cells. At each time step t = 1,2,3,..., the Jabberwock is in some cell $X_t \in \{1,...,N\}^2$, and it moves to cell X_{t+1} randomly as follows: with probability 1 − ε, it chooses one of the (up to 4) valid neighboring cells uniformly at random; with probability ε, 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 ε = ½ , N = 10 and that the Jabberwock always starts in X1 = (1, 1).

(a) Compute the probability that the Jabberwock will be in X_2 = (2,1) at time step 2.

What about P (X_2 = (4, 4))?

(b) (4 points) 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).

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 Ft ∈ {0, 1} to denote whether it will teleport at time t. We want to add these frumious variables to the HMM. Consider the two candidates:

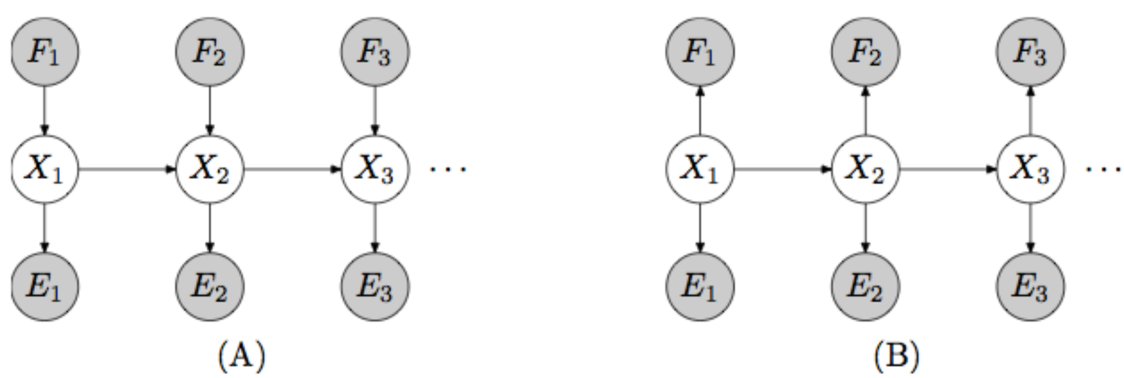

(A)	(B)
X_1 independent of $X_3 \mid X_2$	X_1 independent of $X_3 \mid X_2$
X_1 independent of $E_2 \mid X_2$	X_1 independent of $E_2 \mid X_2$
X_1 independent of $F_2 \mid X_2$	X_1 independent of $F_2 \mid X_2$
X_1 independent of $E_4 \mid X_2$	X_1 independent of $E_4 \mid X_2$
X_1 independent of $F_4 \mid X_2$	X_1 independent of $F_4 \mid X_2$
E_3 independent of $F_3 \mid X_3$	E_3 independent of $F_3 \mid X_3$
E_1 independent of $F_2 \mid X_2$	E_1 independent of $F_2 \mid X_2$
E_1 independent of $F_2 \mid E_2$	E_1 independent of $F_2 \mid E_2$

(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 a better representation of the described causal relationships.