CS 188 Spring 2002 — Introduction to AI — Stuart Russell — Final

You have 2 hours and 50 minutes. The exam is open-book, open-notes. 100 points total.

You will not necessarily finish all questions, so do your best ones first.
Write your answers in blue books. Check you haven't skipped any by accident. Hand them all in. Panic not.

HAND IN THE EXAM COPY AS WELL AS YOUR BLUE BOOKS.
DO NOT DISCLOSE ANY EXAM CONTENT OR DISCUSS WITH OTHER STUDENTS!!!!!

1. (12 pts.) True/False

Decide if each of the following is true or false. If you are not sure you may wish to provide a *brief* explanation to follow your answer.

(a) (2) The truth of any English sentence can be determined given a grammar and given semantic definitions for all the words.

(b) (2) Using dynamic Bayesian networks for speech recognition instead of HMMs does not necessarily change the complexity of the problem.

(c) (2) It is not always possible to determine the size of an object from a single image.

(d) (2) There is no clause that, when resolved with itself, yields (after factoring) the clause $(\neg P \vee \neg Q)$.

(e) (2) Every partial-order plan with no open conditions and no possible threats has a linearization that is a correct solution.

(f) (2) There exists a set S of Horn clauses such that the assignment in which every symbol is false is not a model of S.

2. (15 pts.) Logic

(a) (2) Translate into *good, natural* English (no xs and ys!):

$$\forall x, y, l \; SpeaksLanguage(x, l) \wedge SpeaksLanguage(y, l) \\ \Rightarrow Understands(x, y) \wedge Understands(y, x)$$

(b) (3) Translate into first-order logic the following sentences:

i. "If someone understands someone, then he is that someone's friend."

ii. "Friendship is transitive."

Remember to define all predicate, function, or constants and avoid the *LongPredicateNames* trap.

(c) (5) Suppose that Ann and Bob speak French and Bob and Cal speak German. Prove, using any first-order logical theorem-proving method you like, that Ann is Cal's friend, using as axioms the sentences from parts (a) and (b). Explain each step in detail, including any unifications required. You may abbreviate any symbols as necessary.

(d) (5) Give a formal proof that the sentence in (a) is entailed by the sentence

$$\forall x, y, l \; SpeaksLanguage(x, l) \wedge SpeaksLanguage(y, l) \\ \Rightarrow Understands(x, y)$$

3. (14 pts.) Games

Consider a two-player game featuring a board with four locations, numbered 1 through 4 and arranged in a line. Each player has a single token. Player A starts with his token on space 1, and player B starts with his token on space 4. Player A moves first.

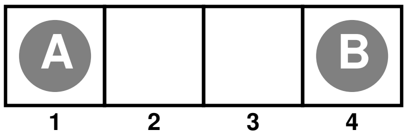

The two players take turns moving, and each player must move his token to an open adjacent space *in either direction*. If the opponent occupies an adjacent space, then a player may jump over the opponent to the next open space if any. (For example, if A is on 3 and B is on 2, then A may move back to 1.) The game ends when one player reaches the opposite end of the board. If player A reaches space 4 first, then the value of the game is $+1$; if player B reaches space 1 first, then the value of the game is -1.

(a) (5) On a fresh page, draw the complete game tree, using the following conventions:

- Write each state as (s_A, s_B) where s_A and s_B denote the token locations.
- Put the terminal states in square boxes, and annotate each with its game value in a circle.
- Put *loop states* (states that already appear on the path to the root) in double square boxes. Since it is not clear how to assign values to loop states, annotate each with a "?" in a circle.

(b) (4) Now mark each node with its backed-up minimax value (also in a circle). Explain in words how you handled the "?" values, and why.

(c) (5) Explain why the standard minimax algorithm would fail on this game tree and briefly sketch how you might fix it, drawing on your answer to (b). Does your modified algorithm give optimal decisions for all games with loops?

4. (12 pts.) MDPs and Games (ANSWER Q.3 FIRST)

Now we will take a different approach to the game in Q.3, viewing it in the framework of MDPs. Here is the state-space graph for the game, showing moves by A as solid lines and moves by B as dashed lines.

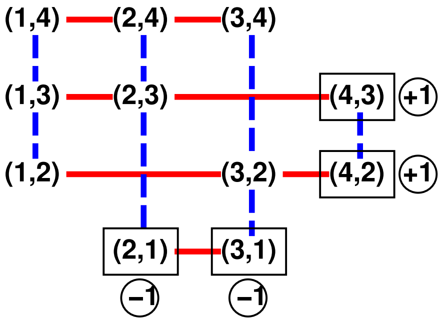

(a) (5) Consider a general zero-sum, turn-taking, stochastic MDP with two players A and B. Let $U_A(s)$ be the utility of state s when it is A's turn to move in s, and let $U_B(s)$ be the utility of state s when it is B's turn to move in s. Let $R(s)$ be the reward in s. All rewards and utilities are calculated from A's point of view (just as in a minimax game tree). Write down Bellman equations defining $U_A(s)$ and $U_B(s)$.

(b) (5) Briefly explain how to do two-player value iteration with these equations and apply value iteration to the game from Q.3, using the following table. We have initialized U_A and marked the terminal values as fixed. Your job is to complete the next two rows (in your blue book).

	(1,4)	(2,4)	(3,4)	(1,3)	(2,3)	(4,3)	(1,2)	(3,2)	(4,2)	(2,1)	(3,1)
U_A	0	0	0	0	0	**+1**	0	0	**+1**	**–1**	**–1**
U_B						**+1**			**+1**	**–1**	**–1**
U_A						**+1**			**+1**	**–1**	**–1**

(c) (2) Define a suitable termination condition for two-player value iteration.

5. (22 pts.) Statistical learning, Bayes nets In this question we will look at maximum likelihood learning, as discussed in the lecture on Chapter 19.

(a) (2) Consider a single Boolean random variable Y (the "classification"). Let the prior probability $P(Y = true)$ be π. Let's try to find π, given a training set $D = (y_1, \ldots, y_N)$ with N independent samples of Y. Furthermore, suppose p of the N are positive and n of the N are negative. Write down an expression for the likelihood of D (i.e., the probability of seeing this particular sequence of examples, given a fixed value of π) in terms of π, p, and n.

(b) (3) By differentiating the log likelihood L, find the value of π that maximizes the likelihood.

(c) (2) Now suppose we add in k Boolean random variables $X_1, X_2, \ldots, X_k$ (the "attributes") that describe each sample, and suppose we assume that the attributes are conditionally independent of each other given the goal Y. Draw the Bayes net corresponding to this assumption.

(d) (4) Write down the likelihood for the data including the attributes, using the following additional notation:

- α_i is $P(X_i = true | Y = true)$.
- β_i is $P(X_i = true | Y = false)$.
- p_i^+ is the count of samples for which $X_i = true$ and $Y = true$.
- n_i^+ is the count of samples for which $X_i = false$ and $Y = true$.
- p_i^- is the count of samples for which $X_i = true$ and $Y = false$.
- n_i^- is the count of samples for which $X_i = false$ and $Y = false$.

[*Hint*: consider first the probability of seeing a single example with specified values for $X_1, X_2, \ldots, X_k$ and Y.]

(e) (5) By differentiating the log likelihood L, find the values of α_i and β_i (in terms of the various counts) that maximize the likelihood and say in words what these values represent.

(f) (3) Let $k = 2$, and consider a data set with 4 examples as follows:

X_1	X_2	Y
0	0	0
0	1	1
1	0	1
1	1	0

Compute the maximum likelihood estimates of π, α_1, α_2, β_1, and β_2.

(g) (2) Given these estimates of π, α_1, α_2, β_1, and β_2, what are the posterior probabilities $P(Y = true | x_1, x_2)$ for each example?

(h) (2) Comment on the connection between this result and the capabilities of a single-layer perceptron.

6. (15 pts.) Natural language

The next page shows the lexicon and grammar rules for "wumpus pidgin" (slightly modified from the book).

(a) (3) Which of the following sentences are generated by the grammar (possibly more than one):

i. I see the gold but it is near the smelly wumpus.
ii. I shoot the breeze back east in Boston
iii. You that smell the wumpus that stinks I go and you kill it

(b) (4) Propose a modified rule for relative clauses that also allows the sentence "The wumpus that the dogs see stinks."

(c) (4) Show a parse tree of this sentence using your new rule.

(d) (4) In English it is also legal to say "The wumpus the dogs see stinks," omitting the word "that". It is not, however, legal to say "The wumpus the dogs I smell see stinks."[1] Make minimal adjustments to the grammar to allow the first sentence but not to allow the second.

[1]This is the result of removing two "that"s from "The wumpus that the dogs that I smell see stinks." Said carefully, the latter is really a sentence of English.

Noun	$\rightarrow$	***stench*** | ***breeze*** | ***glitter*** | ***nothing***
		| ***wumpus*** | ***pit*** | ***dogs*** | ***gold*** | ***east*** | ...
Verb	$\rightarrow$	***is*** | ***see*** | ***smell*** | ***shoot*** | ***feel*** | ***stinks***
		| ***go*** | ***grab*** | ***carry*** | ***kill*** | ***turn*** | ...
Adjective	$\rightarrow$	***right*** | ***left*** | ***east*** | ***south*** | ***back*** | ***smelly*** | ...
Adverb	$\rightarrow$	***here*** | ***there*** | ***nearby*** | ***ahead***
		| ***right*** | ***left*** | ***east*** | ***south*** | ***back*** | ...
Pronoun	$\rightarrow$	***me*** | ***you*** | ***I*** | ***it*** | ...
Name	$\rightarrow$	***John*** | ***Mary*** | ***Boston*** | ***Aristotle*** | ...
Article	$\rightarrow$	***the*** | ***a*** | ***an*** | ...
Preposition	$\rightarrow$	***to*** | ***in*** | ***on*** | ***near*** | ...
Conjunction	$\rightarrow$	***and*** | ***or*** | ***but*** | ...
Digit	$\rightarrow$	**0** | **1** | **2** | **3** | **4** | **5** | **6** | **7** | **8** | **9**

The lexicon for wumpus pidgin.

S	$\rightarrow$	*NP VP*	I + feel a breeze
	|	*S Conjunction S*	I feel a breeze + and + I smell a wumpus
NP	$\rightarrow$	*Pronoun*	I
	|	*Noun*	dogs
	|	*Article Noun*	the + wumpus
	|	*Digit Digit*	3 4
	|	*NP PP*	the wumpus + to the east
	|	*NP RelClause*	the wumpus + that is smelly
VP	$\rightarrow$	*Verb*	stinks
	|	*VP NP*	feel + a breeze
	|	*VP Adjective*	is + smelly
	|	*VP PP*	turn + to the east
	|	*VP Adverb*	go + ahead
PP	$\rightarrow$	*Preposition NP*	to + the east
RelClause	$\rightarrow$	***that*** *VP*	that + is smelly

The grammar for wumpus pidgin, with example phrases for each rule.

7. (10 pts.) Robotics

Consider the two-link robotic arm shown in the following figure. The arm rotates at the pivot in the center and its position is defined by two angles: θ_1 is the angle between the x-axis and the first link (ranging from 0 to 2π) and θ_2 is the angle between the first and second links (also ranging from 0 to 2π). The square block on the left and the vertical wall on the right represent obstacles. Both a start configuration (solid lines) and a goal configuration (dashed lines) are shown.

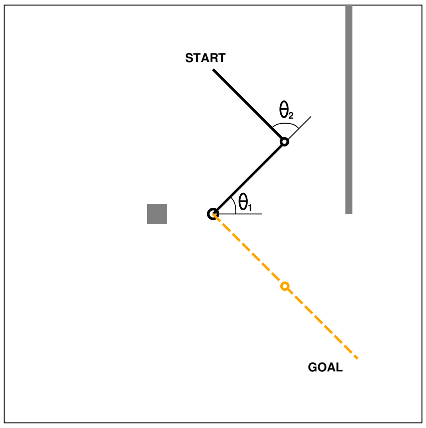

(a)

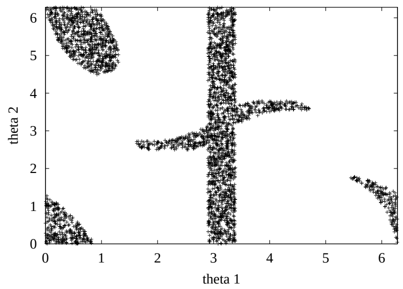

(b)

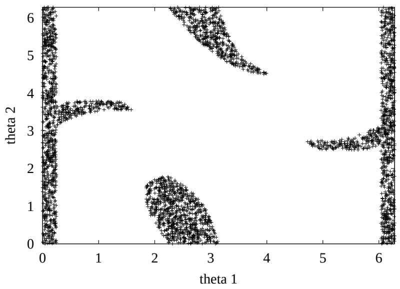

(c)

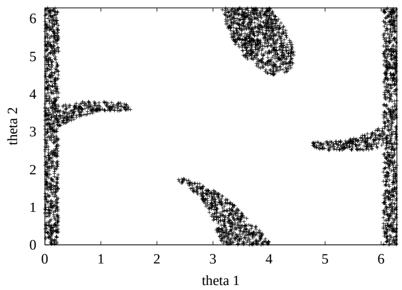

(d)

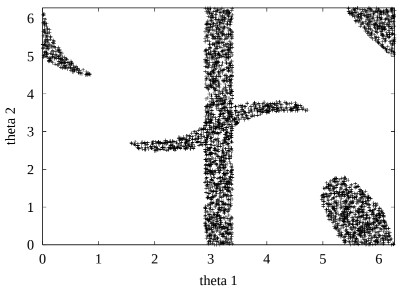

(a) (5) Choose the appropriate configuration space from above.

(b) (5) Copy the configuration space diagram, mark the start and goal configurations, and show an appropriate plan for the robot to move from the start to the goal.