Exam for IN4010TU Artificial Intelligence Techniques

26 January 2006

This exam will test your knowledge and understanding of Russell and Norvig, *Artificial Intelligence: A Modern Approach*. Using the book during the examination is *not* allowed. You will have **3** hours (from 9 till 12) to complete the exam. It has 17 questions, for a total of 80 points. For the **multiple-choke** questions, write down on a separate sheet of paper only the letter with the best answer. For the **open** questions, include a short sentence of explanation if the answer consists of only one item. Please don't include irrelevant information: you will be marked down for this. Before you hand in your answers, please check that you have put your name on top of every sheet you hand in.

Multiple-choice questions

Question 1 *2 points*

What does a *rational agent* do?

A. Gather as much information as possible before making a decision.

B. Maximise the expected performance.

C. Maximise the actual performance.

D. Use only *a priori* knowledge.

Question 2 *2 points*

What does it mean for a heuristic function $h(n)$ to be *admissable?*

A. $h(n) = 0$.

B. $h(n)$ never overestimates the cost to reach the goal.

C. $h(n)$ is by nature pessimistic.

D. $h(n)$ is the exact cost to reach the goal.

Question 3 *2 points*

What is *iterative deepening depth-first search?*

A. Depth-first search with a gradually increasing depth limit.

B. Depth-first search with a predetermined depth limit.

C. A search strategy where two searches are run simultaneously.

D. Depth-first search that expands the lowest-cost node first.

Question 4 *2 points*

What does it mean for an inference algorithin to be *sound?*

A. It derives all entailed sentences.

B. It derives only entailed sentences and does not make things up.

C. It derives all sentences, whether entailed or not.

D. None of the previous answers is correct.

Question 5 *2 points*

What is a *unit clause?*

A. A clause with one positive literal.

B. A clause with a single literal.

C. A clause with a single false literal.

D. A Horn clause.

Question 6 *2 points*

In situation calculus, what is meant by a *fluent*?

A. A sequence of actions that achieves the desired effect.

B. An agent that moves.

C. An atemporal predicate or function.

D. A function or predicate that varies from one situation to the next.

Question 7 *2 points*

What is the difference between using fuzzy logic and using probability theory for representing uncertain knowledge?

A. Fuzzy logic deals with the properties of hairy truths; probability theory deals with the properties of clean truths.

B. Fuzzy logic deals with degrees of ignorance; probability theory deals with degrees of knowledge.

C. Fuzzy logic deals with degrees of belief; probability theory deals with degrees of truth.

D. Fuzzy logic deals with degrees of truth; probability theory deals with degrees of belief.

Question 8 *2 points*

Say we have a dynamic probabilistic model. Which algorithm should be used for finding the most likely state sequence given an observation sequence: the smoothing algorithm, or the Viterbi algorithm?

A. The smoothing algorithm: Viterbi uses both the forward and the backward message so that it becomes intractable for any model of considerable size.

B. The Viterbi algorithm: it uses both the forward and the backward message whereas the smoothing algorithm uses only the forward message.

C. The smoothing algorithm: it finds a smoothed sequence that has the more likely transitions.

D. The Viterbi algorithm: it maximises the probability of the path whereas the smoothing algorithm maximises the sum of the probabilities of all paths to one state.

Question 9 *2 points*

What assumption do naive Bayesian models for document classification make?

A. A document can be classified only using the first N words.

B. Words in a document are independent from their neighbours.

C. No documents with explicit content, unsuitable for a naive audience, should be classified.

D. Common words like "the", "is", "are" have no semantic content and should not be taken into account.

Question 10 *2 points*

Suppose a knowledge base contains just one sentence:

$$\exists x \quad AsHighAs(x, Everest) \tag{1}$$

Is this sentence a legitimate result of applying Existential Instantiation:

$$AsHighAs(Kilimanjaro, Everest) \wedge AsHighAs(BenNevis, Everest) \tag{2}$$

A. Yes: any name that does not occur in the knowledge base may be filled in for x.

B. No: x may not be replaced by a new name, but only by a Skolem constant.

C. No: Universal Instantiation should be used.

D. No: Existential Instantiation may not be applied twice on one sentence.

Question 11 *2 points*

How long does it take to prove $KB \models \alpha$ using the DPLL algorithm for checking satisfiability when α is a literal already contained in KB?

A. The algorithm will find the literal immediately: $O(1)$.

B. Because of early termination, the algorithm has to iterate only over the clauses up to the literal α so it takes $O(n)$ times where n is the number of clauses.

C. Being based on backtracking search, the algorithm takes logarithmic time.

D. Checking satisfiability is NP-complete, so the algorithm does not necessarily terminate.

Question 12 *2 points*

In search problems, path costs are usually defined. Suppose that actions can have arbitrarily large negative costs; would this possibility force any optimal algorithm to explore the entire state space?

A. No. An optima1 algorithm is guaranteed to terminate in finite time; it will by definition only explore part of the state space.

B. No. Tlie state changes after every action that carries negative cost.

C. Yes. Every branch could lead to a sequence of negative-cost transitions or a negative-cost loop, so the algorithm must explore every branch.

D. Yes. Negative costs lead to an unobservable environment so that all actions must be tried until the belief state contains one world state.

Question 13 *2 points*

Why can a *leak node* be added for a Noisy-OR relation?

A. To make the inhibition probabilities add up to 1.

B. To cover miscellaneous causes that are not modelled explicitly.

C. To turn the Noisy-OR relation into Noisy-AND.

D. None of the above.

Question 14 *e points*

What is meant by *hidden variables* in a Bayesian network?

A. The nonevidence variables.

B. The evidence variables.

C. The variables whose existence has not been ensured yet.

D. The query variables.

Open questions

Question 15 *18 points*

Consider an agent trying its hand at a sudoku. Sudokus are puzzles that have recently become popular. Solving a sudoku is a matter of filling in the missing digits in a 9 x 9 tableau. Every row, every column, and the nine 3 x 3 squares should contain all digits $\{1 \ldots 9\}$. An example of a puzzle and its solution:

	9		3		1		5	
			9		2			
5	1						9	6
		7	4		3	9		
2								1
		8	2		9	4		
7	8		*				3	2
			8		4			
	6		5		7		4	

8	9	2	3	6	1	7	5	4
4	7	6	9	5	2	1	8	3
5	1	3	7	4	8	2	9	6
6	5	7	4	1	3	9	2	8
2	4	9	6	8	5	3	7	1
1	3	8	2	7	9	4	6	5
7	8	4	1	9	6	5	3	2
9	2	5	8	3	4	6	1	7
3	6	1	5	2	7	8	4	9

Please do not let yourself be fooled into believing this has got to do with numbers. The puzzle could have used letters $\{a \ldots i\}$ or nine different colours. It is an example of a constraint satisfaction problem.

(a) (z points) Is the environment deterministic?

(b) (2 points) By what variables is this problem defined and what do they stand for?

(c) (2 points) What is the *domain size* of the variables?

(d) (5 points) Give the initial state, successor function, and goal test for the search problem. Choose a formulation that is precise enough to be implemented.

(e) (4 points) Why would using plain depth-first search be a bad idea in general for solving sudokus? In what way would using backtracking search help and it what way would it be far from perfect?

(f) (3 points) In trying to solve the puzzle, the agent fills in a digit at the position of the asterisk *, because there are only two legal values for that square: {1,6). What is the name of the heuristic that leads to this strategy?

Question 16 15 *points*

The monkey-and-bananas problem is faced by a monkey in a laboratory with some bananas hanging out of reach from the ceiling. A box is available that will enable the monkey to reach the bananas if he climbs on it. Initially, the monkey is at A, the bananas at B, and the box at C. The monkey and box have height *Low*, but if the monkey climbs onto tlie box he will have height *High*, the same as the bananas. The actions available to the monkey include *Go* from one place to another, *Push* an object from one place to another, *ClimbUp* onto or *ClimbDown* from an object, and *Grasp* or *Ungrasp* an object. Grasping results in holding the object if the monkey and object are in the same place at the same height.

(a) (3 points) Write down the initial state description.

(b) (4 points) Write down STRIPS-style definitions of the six actions.

(c) (5 points) Suppose the monkey wants to fool tlie scientists, who are off to tea, by grabbing the bananas, but leaving the box in its original place. Write this as a general goal (i.e. not assuming that the box is necessarily at C) in the language of situation calculus. Can this goal be solved by a STRIPS-style system?

(d) (3 points) Your axiom for pushing is probably incorrect, because if the object is too heavy, its position will remain the same when the *Push* operator is applied. Is this an example of the ramification problem or the qualification problem? Fix your problem description to account for heavy objects.

Question 17 *19 points*

The following grammar is given:

S		*NP VP*
NP	→	*NP RelClause*
	|	*Pronoun*
	|	*Noun*
	|	*Adjective Noun*
	|	*Article Noun*
VP	→	*Verb*
	|	*Verb Adjective*
		Verb PP
PP	→	*Preposition NP*
RelClause	→	that *VP*

with this lexicon:

Noun	→	agents | people | chair | east | wumpus
Article	→	the | a | an
Adjective	→	smelly | evil
Verb	→	is | are | turn | turns | go | goes
Preposition	→	in | to

(a) (6 points) Consider the sentence

The agents that are smelly turn to the east. (3)

Draw the parse tree for this sentence.

(b) (5 points) Consider the sentence

The chair are smelly. (4)

This is a valid sentence according to the grammar above, but not a valid sentence of English, because "chair" is a singular noun and "are" is a plural verb. Show how we can augment the grammar and the lexicon above to capture the constraint in English that verbs must "agree" with their subjects.

(c) (4 points) Give one advantage and one disadvantage of using a bigram model over using the grammar above as a language model.

(d) (4 points) Will a chart parser run on sentence ***(3)*** find an edge from vertex i to vertex g in the process? That is, will it find a parse for the sequence "agents that are smelly turn to the east"?

End of exam

Please check again that you have included one sentence of explanation for all open questions for which the answer consists of only one item.