



Försättsblad till skriftlig tentamen vid Linköpings Universitet

(fylls i av ansvarig)

Datum för tentamen	2009-08-18
Sal	KÅRA
Tid	8-12
Kurskod	TDDC17
Provkod	TEN1
Kursnamn/benämning	Artificiell intelligens
Institution	IDA
Antal uppgifter som ingår i tentamen	8
Antal sidor på tentamen (inkl. försättsbladet)	8
Jour/Kursansvarig	Mariusz Wzorek
Telefon under skrivtid	070-3887122
Besöker salen ca kl.	10.00
Kursadministratör (namn + tfnr + mailadress)	Anna Grabska Eklund Ankn. 23 62, annek@idag.liu.se
Tillåtna hjälpmedel	Miniräknare/Hand calculators
Övrigt (exempel när resultat kan ses på webben, betygsgränser, visning, övriga salar tentan går i m.m.)	
Vilken typ av papper ska användas, rutigt eller linjerat	
Antal exemplar i påsen	31

Linköpings Universitet
Institutionen för Datavetenskap
Patrick Doherty

Tentamen
TDDC17 ¹Artificial Intelligence
18 August 2009 kl. 08-12

Points:

The exam consists of exercises worth 37 points.
To pass the exam you need 19 points.

Auxiliary help items:

Hand calculators.

Directions:

You can answer the questions in English or Swedish.
Use notations and methods that have been discussed in the course.
In particular, use the definitions, notations and methods in appendices 1-3.
Make reasonable assumptions when an exercise has been under-specified.
Begin each exercise on a new page.
Write only on one side of the paper.
Write clearly and concisely.

Jourhavande: Mariusz Wzorek 0703887122. Mariusz will arrive for questions around 10:00.

¹This exam should also be used for those students who have previously taken TDDA13 or TDDA58.

1. Suppose there are three blocks, a, b, c , where a is stacked on top of b and b is stacked on top of c . In addition, block a is green and block c is not. This can be described with the following theory, (where a, b, c are constants):

$$On(a, b), On(b, c), Green(a), \neg Green(c) \tag{1}$$

Suppose, one would like to prove that given the theory above, there is a green block directly on top of a non-green block (note that a is not directly on top of c). This can be described as the following query:

$$\exists x \exists y. On(x, y) \wedge Green(x) \wedge \neg Green(y) \tag{2}$$

- (a) Using Resolution, show that the query (2) above follows from the theory (1). [4p]
- Your answer should be structured using a resolution refutation tree (as used in the book).
 - Since the unifications are trivial, it suffices to simply show the binding lists at each resolution step.
 - Remember that formulas in the theory and the query must be transformed into CNF form. Use appendix 1 if required.
2. Constraint satisfaction problems consist of a set of variables, a value domain for each variable and a set of constraints. A solution to a CS problem is a consistent set of bindings to the variables that satisfy the constraints. A standard backtracking search algorithm can be used to find solutions to CS problems. In the simplest case, the algorithm would choose variables to bind and values in the variable's domain to be bound to a variable in an arbitrary manner as the search tree is generated. This is inefficient and there are a number of strategies which can improve the search. Describe the following three strategies:
- (a) Minimum remaining value heuristic (MRV). [1p]
- (b) Degree heuristic. [1p]
- (c) Least constraining value heuristic. [1p]
- Constraint propagation is the general term for propagating constraints on one variable onto other variables. Describe the following:
- (d) What is the Forward Checking technique? [1p]
- (e) What is arc consistency? [1p]
3. The following questions pertain to partial-order planning:
- (a) What is partial-order planning and how does it differ from STRIPS-based planning in terms of search space and output of the respective planning algorithms? [3p]
- (b) Describe the four basic components in a partial-order plan according to Russell/Norvig. [2p]

4. Consider the following example:

Aching elbows and aching hands may be the result of arthritis. Arthritis is also a possible cause of tennis elbow, which in turn may cause aching elbows. Dishpan hands may also cause aching hands.

- (a) Represent these causal links in a Bayesian network. Let *ar* stand for "arthritis", *ah* for "aching hands", *ae* for "aching elbow", *te* for "tennis elbow", and *dh* for "dishpan hands". [2p]
- (b) Given the independence assumptions implicit in the Bayesian network, write the formula for the full joint probability distribution over all five variables? [2p]
- (c) Compute the following probabilities using the formula for the full joint probability distribution and the probabilities below:
 - $P(ar | te, ah)$ [1p]
 - $P(ar, \neg dh, \neg te, ah, \neg ae)$ [1p]
 - Appendix 2 provides you with some help in answering these questions.

Table 1: probabilities for question 5.

$$\begin{aligned}P(ah | ar, dh) &= P(ae | ar, te) = 0.1 \\P(ah | ar, \neg dh) &= P(ae | ar, \neg te) = 0.99 \\P(ah | \neg ar, dh) &= P(ae | \neg ar, te) = 0.99 \\P(ah | \neg ar, \neg dh) &= P(ae | \neg ar, \neg te) = 0.00001 \\P(te | ar) &= 0.0001 \\P(te | \neg ar) &= 0.01 \\P(ar) &= 0.001 \\P(dh) &= 0.01\end{aligned}$$

5. Alan Turing proposed the Turing Test as an operational definition of intelligence.

- (a) Describe the Turing Test using your own diagram and explanations.[2p]
- (b) Do you believe this is an adequate test for machine intelligence? Justify your answer.[1p]

6. Modeling actions and change in incompletely represented, dynamic worlds is a central problem in knowledge representation. The following questions pertain to reasoning about action and change.

- (a) What is Temporal Action Logic? Explain by describing the ontology used in the formalism, that is, what is taken to exist, and what notation is used in the logical language to represent those things that are taken to exist. [2p]
- (b) What is the frame problem? Use the Wumpus world to provide a concrete example of the problem. Represent the problem by representing an initial timepoint, an action and the result of the action using the TAL notation (either with macros or without). [2p]
- (c) What is nonmonotonic logic? How can it be used to provide solutions to the frame problem? [2p]

7. A* search is the most widely-known form of best-first search. The following questions pertain to A* search:

- (a) Explain what an *admissible* heuristic function is using the notation and descriptions in (c). [1p]
- (b) Suppose a robot is searching for a path from one location to another in a rectangular grid of locations in which there are arcs between adjacent pairs of locations and the arcs only go in north-south (south-north) and east-west (west-east) directions. Furthermore, assume that the robot can only travel on these arcs and that some of these arcs have obstructions which prevent passage across such arcs.

The *Manhattan distance* between two locations is the shortest distance between the locations ignoring obstructions. Is the Manhattan distance in the example above an admissible heuristic? Justify your answer explicitly. [2p]

- (c) Let $h(n)$ be the estimated cost of the cheapest path from a node n to the goal. Let $g(n)$ be the path cost from the start node n_0 to n . Let $f(n) = g(n) + h(n)$ be the estimated cost of the cheapest solution through n .

Provide a general proof that A* using tree-search is optimal if $h(n)$ is admissible. If possible, use a diagram to structure the proof. [2p]

8. The following question pertains to Decision Tree Learning. Use the definitions and data table in Figure 2, Appendix 3 to answer this question. Figure 1 shows a partial decision tree for the Table in Figure 2 with target attribute *PlayTennis*.

- (a) What attribute should be tested in the box with the question mark on the right branch of the decision tree in Figure 1? Justify your answer by computing the information gain for the appropriate attributes in the Table in Figure 2 in Appendix 3. [3p]

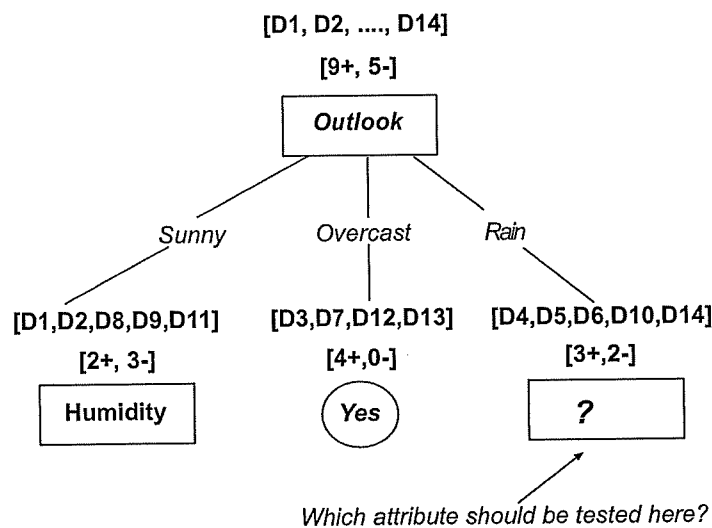


Figure 1: Partial Decision Tree for PlayTennis

Appendix 1

Converting arbitrary wffs to clause form:

1. Eliminate implication signs.
2. Reduce scopes of negation signs.
3. Standardize variables within the scopes of quantifiers (Each quantifier should have its own unique variable).
4. Eliminate existential quantifiers. This may involve introduction of Skolem constants or functions.
5. Convert to prenex form by moving all remaining quantifiers to the front of the formula.
6. Put the matrix into conjunctive normal form. Two useful rules are:
 - $\omega_1 \vee (\omega_2 \wedge \omega_3) \equiv (\omega_1 \vee \omega_2) \wedge (\omega_1 \vee \omega_3)$
 - $\omega_1 \wedge (\omega_2 \vee \omega_3) \equiv (\omega_1 \wedge \omega_2) \vee (\omega_1 \wedge \omega_3)$
7. Eliminate universal quantifiers.
8. Eliminate \wedge symbols.
9. Rename variables so that no variable symbol appears in more than one clause.

Appendix 2

A generic entry in a joint probability distribution is the probability of a conjunction of particular assignments to each variable, such as $P(X_1 = x_1 \wedge \dots \wedge X_n = x_n)$. The notation $P(x_1, \dots, x_n)$ can be used as an abbreviation for this.

The chain rule states that any entry in the full joint distribution can be represented as a product of conditional probabilities:

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i | x_{i-1}, \dots, x_1) \quad (3)$$

Given the independence assumptions implicit in a Bayesian network a more efficient representation of entries in the full joint distribution may be defined as

$$P(x_1, \dots, x_n) = \prod_{i=1}^n P(x_i | \text{parents}(X_i)), \quad (4)$$

where $\text{parents}(X_i)$ denotes the specific values of the variables in $\text{Parents}(X_i)$.

Recall the following definition of a conditional probability:

$$P(X | Y) = \frac{P(X \wedge Y)}{P(Y)} \quad (5)$$

The following is a useful general inference procedure:

Let X be the query variable, let \mathbf{E} be the set of evidence variables, let \mathbf{e} be the observed values for them, let \mathbf{Y} be the remaining unobserved variables and let α be the normalization constant:

$$P(X | \mathbf{e}) = \alpha P(X, \mathbf{e}) = \alpha \sum_{\mathbf{y}} P(X, \mathbf{e}, \mathbf{y}) \quad (6)$$

where the summation is over all possible \mathbf{y} 's (i.e. all possible combinations of values of the unobserved variables \mathbf{Y}).

Equivalently, without the normalization constant:

$$P(X | \mathbf{e}) = \frac{P(X, \mathbf{e})}{P(\mathbf{e})} = \frac{\sum_{\mathbf{y}} P(X, \mathbf{e}, \mathbf{y})}{\sum_{\mathbf{x}} \sum_{\mathbf{y}} P(\mathbf{x}, \mathbf{e}, \mathbf{y})} \quad (7)$$

Appendix 3

Definition 1

Given a collection S , containing positive and negative examples of some target concept, the entropy of S relative to this boolean classification is

$$Entropy(S) \equiv -p_{\oplus} \log_2 p_{\oplus} - p_{\ominus} \log_2 p_{\ominus},$$

where p_{\oplus} is the proportion of positive examples in S and p_{\ominus} is the proportion of negative examples in S .

Definition 2

Given a collection S , containing positive and negative examples of some target concept, and an attribute A , the information gain, $Gain(S, A)$, of A relative to S is defined as

$$Gain(S, A) \equiv Entropy(S) - \sum_{v \in Values(A)} \frac{|S_v|}{|S|} Entropy(S_v),$$

where $values(A)$ is the set of all possible values for attribute A and S_v is the subset of S for which the attribute A has value v (i.e., $S_v = \{s \in S \mid A(s) = v\}$).

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain	Mild	High	Weak	Yes
D5	Rain	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

Figure 2: Sample Table for Target Attribute PlayTennis

For help in converting from one logarithm base to another (if needed):

$$\log_a x = \frac{\log_b x}{\log_b a}$$

$$\ln x = 2.303 \log_{10} x$$

Note also that for the example, we define $0 \log 0$ to be 0.