Examination
Artificial Intelligence
(214030)
June 25, 2008

Introduction

This examination consists of 32 multiple-choice questions. You may only use the book "Artificial Intelligence. A Modern Approach" during the exam. You have 3 hours and 30 mins. At the end of the exam you must hand in this question paper and the answer form. Each multiple choice question counts for 4 points.
Tips:

- Read each question carefully keeping the possible answers covered.
- Try to answer the question yourself, before you look at the answers you are given to choose from. Make a note of your first thoughts and calculations on a scribbling-paper (kladpapier).
- All of the questions in this examination call for understanding and insight. You will frequently need to look up a term or a formula in the textbook. This will be necessary to enable you to make calculations or to build up an argumentation which you can then compare with the possible answers of the multiple-choice questions.
- Beware of double negations (negatives) as these can be confusing.
- Do not stay on any one question too long. If you do not know the answer and have spent more than 10 minutes on the question, move on to the next question and come back to this one later.
- If you have any time over at the end, check your answers.
- Fill in your answers on this question form first and transfer them to the answer form at the end

Good luck!

Multiple-choice questions

1. A search problem can be described by a directed finite search graph, cf. Chapter 3. Consider the following statements:

   (i) The Depth First Search Algorithm terminates on the search problem *if* the corresponding directed search graph contains no cycles.

   (ii) If the directed search graph contains no cycles then the Depth First Search Algorithm terminates on the corresponding search problem.

   Which of the following claims is true?

   (a) Both statements (i) and (ii) are false.

   (b) Only statement (ii) is true

   (c) Only statement (i) is true

   (d) Both statements (i) and (ii) are true.

2. Consider an A* search algorithm for which $h(n) = 0$. To which of the following search algorithms is this A^* equivalent?

   (a) Greedy best-first search

   (b) Depth-First Search

   (c) Uniform Cost Search

   (d) None of the above.

3. Consider the following Tower of Hanoi problem.

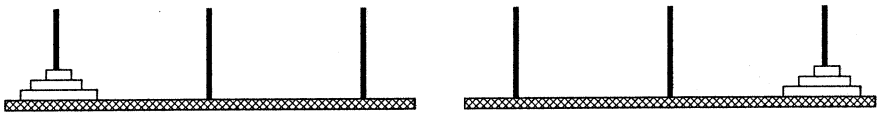

Left: Initial state Right: Goal state

   The cost of moving the small disk is 1, moving the middle sized disk is 2, and moving the large disk is 3. Hence the average cost is 2. Define the heuristic function h as follows: $2 \times$ *number of disks not on the rightmost peg.*
   What is the h value of the initial state?

   (a) 3

   (b) 4

   (c) 5

   (d) 6

4. Assume that we apply A* search to the above Tower of Hanoi problem. Which node will be the **second** one that will be expanded after the initial node (node corresponding to the start state) of the search tree?

   (a) The node corresponding to the state which arises from the initial state by moving disc 1 to the rightmost peg and then afterwards moving the middle sized disc to the empty peg.

   (b) The node corresponding to the state which arises from the initial state by moving the small disc to the middle peg and then afterwards moving the middle sized disc to the rightmost peg.

   (c) The node corresponding to the state initial state, i.e. moving the small disc to an empty peg and back.

   (d) The node corresponding to the state which arises from the initial state by moving the small disc first to the rightmost peg and afterwards moving the small disc to the middle peg.

5. Once again consider the above "cost to go" function h and the following statements about h:

   (i) h is admissible.

   (ii) h is consistent.

   Which of the above statements are true?

   (a) Only statement (i) is true.

   (b) Both statements (i) and (ii) are true.

   (c) Only statement (ii) is true.

   (d) Both statements (i) and (ii) are false.

6. Consider the following part of a two-player game tree.

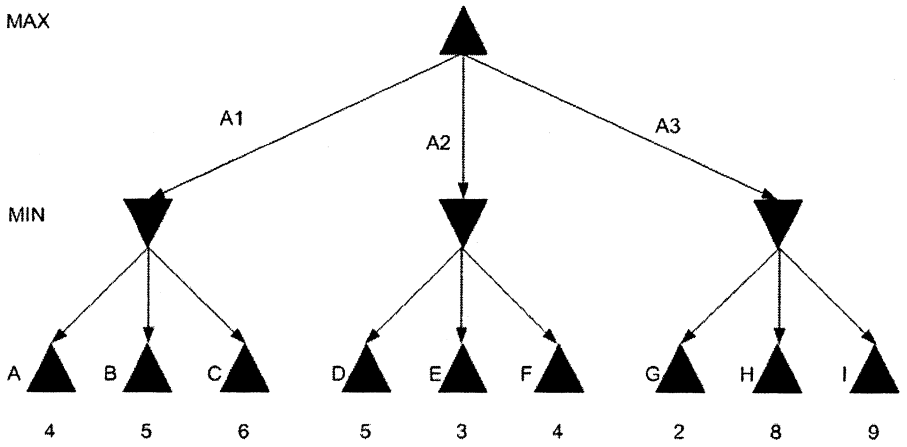

What will be the value of the top MAX node

(a) 4

(b) 6

(c) 8

(d) 9

7. Once again consider the two-player game tree of question 6.

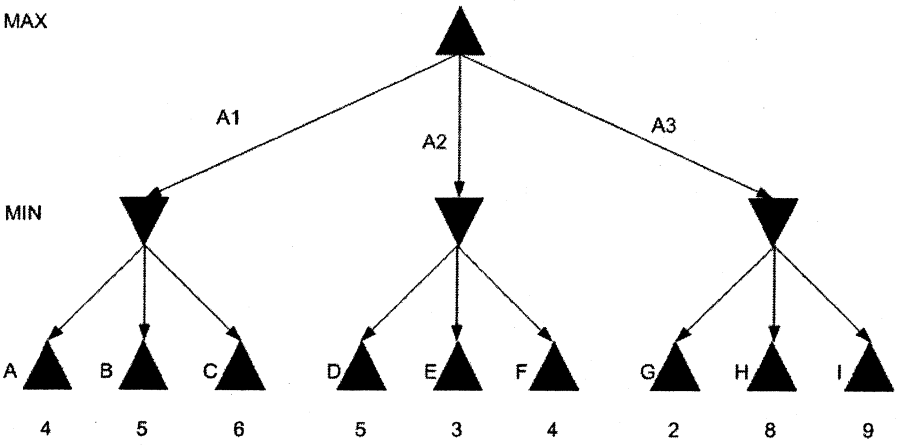

Assume one applies alpha-beta pruning. Which of the following collection of nodes will **all not** being explored?

(a) $\{A, D, G\}$

(b) $\{G, H, I\}$

(c) $\{C, F, I\}$

(d) $\{F, H, I\}$

8. Consider the following game tree for three players, player 1, 2 and 3.

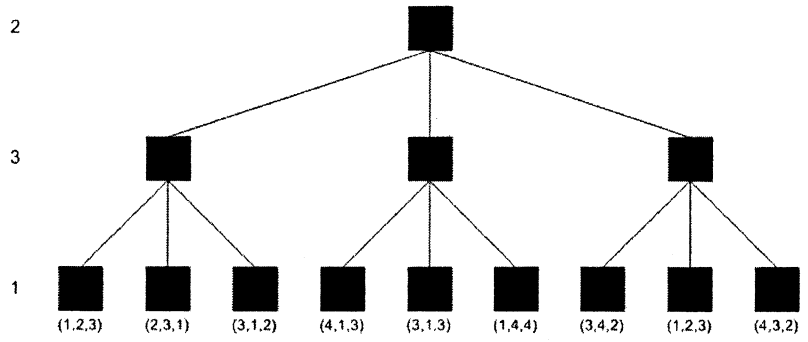

At the lowest nodes the evaluation value of the relevant node is given (the value of the evaluation function) for each of the players. In the i position the evaluation value of the node for player i is given, $i = 1, 2, 3$. What are the values of the nodes for player 2 in the above representation?

(a) (3,4,4)

(b) 4

(c) 1

(d) (1,4,4)

9. Consider the following two-player game tree in which the game has an element of chance, which is shown by the so-called probability nodes in the game tree.

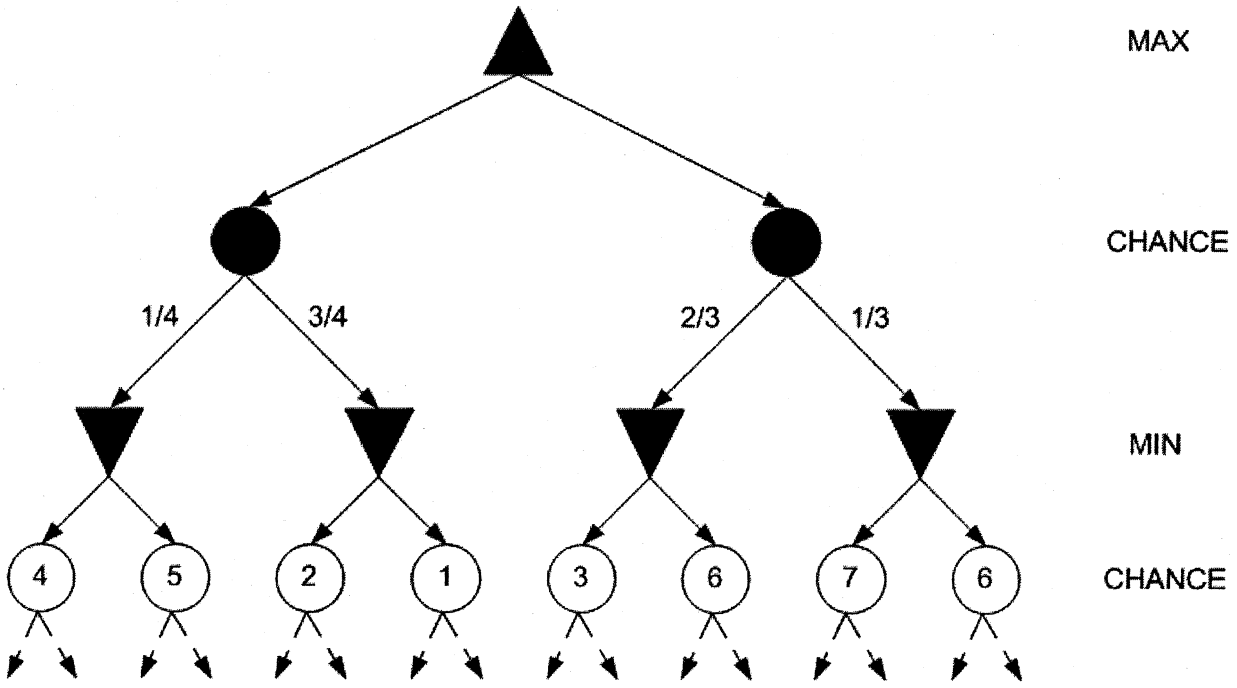

What is the correct value for the top MIN-node if one applies the expectiminimax algorithm?

(a) 6

(b) 2

(c) 4

(d) 7

10. Consider the following game tree with an element of chance.

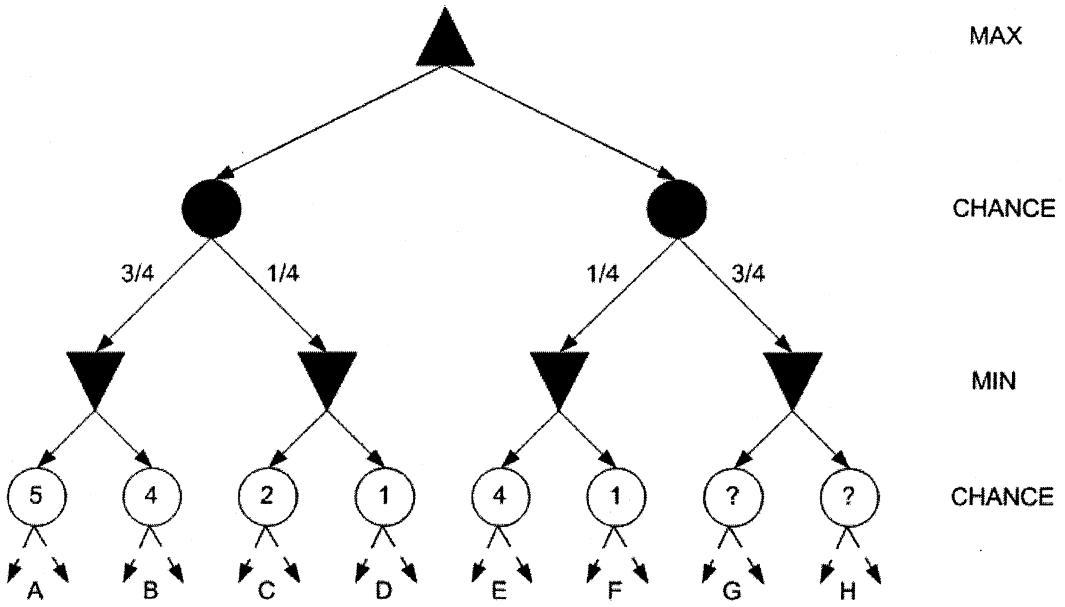

The letters under the bottom row of chance nodes are labels for the nodes just above the letter. One can also apply $\alpha - \beta$ pruning on this game tree. The numbers inside the chance nodes on the bottom row are the computed values of these chance nodes; ? indicates not computed yet. For which values of node G is it **not** necessary to expand node H (to compute the value of node H)?

(a) $\alpha - \beta$ pruning will **always** expand node H.

(b) If the value of node G is less or equal to 4 then $\alpha - \beta$ pruning will **not** expand node H.

(c) If the value of node G is greater then 4 then $\alpha - \beta$ pruning will **not** expand node H.

(d) $\alpha - \beta$ pruning will **never** expand node H.

11. A well-known Dutch folk song says that men who want to sail on a pirate ship have to have beards. We will try to formalise the somewhat simpler sentence "Men who are pirates have beards". Given are the following predicates:

    - $Man(x)$: x is a man.
    - $Pirate(x)$: x is a pirate.
    - $Beard(x)$: x is a beard.
    - $Has(x, y)$: x has y.

    What is the correct formalisation in first-order logic?

    (a) $\forall x\ Man(x) \wedge Pirate(x) \Rightarrow \exists y\ Beard(y) \wedge Has(x, y)$

    (b) $\forall x, y\ Man(x) \wedge Pirate(x) \Rightarrow Beard(y) \wedge Has(x, y)$

    (c) $\forall x\ \exists y\ Man(x) \wedge Pirate(x) \wedge Beard(y) \Rightarrow Has(x, y)$

    (d) $\exists x \exists y\ Man(x) \wedge Pirate(x) \wedge Beard(y) \Rightarrow Has(x, y)$

12. We will formalise the sentence "Neighbours have bikes of the same brand" in first-order logic. Given are the following predicates:

    - $Neighbour(x, y)$: x is the neighbour of y.
    - $HasBike(x, y)$: x has (owns) bike y.
    - $Brand(x, y, z)$: the brand of x is y.

    Further given are two formalisations of this sentence in first-order logic:

    I. $\forall x, y, z, v, w\ Neighbour(x, y) \wedge HasBike(x, v) \wedge Brand(v, z) \wedge HasBike(y, w) \Rightarrow Brand(w, z)$

    II. $\forall x, y, z, v, w\ Neighbour(x, y) \wedge HasBike(x, v) \wedge Brand(v, z) \Rightarrow [HasBike(y, w) \Rightarrow Brand(w, z)]$

    The question is whether these two formalisations are correct: both, one of them, or neither. Which of the following is true?

    (a) The two formalisations are inferentially equivalent and they are both correct.

    (b) The two formalisations are inferentially equivalent and they are both wrong.

    (c) Formalisation I is correct while formalisation II is wrong.

    (d) Formalisation I is wrong while formalisation II is correct.

13. Given are the following premisses:

    - $\neg[D \wedge \neg(E \vee B)]$
    - $C \Rightarrow (E \vee A)$

    The question is whether we can prove $\neg(\neg A \wedge \neg B) \vee \neg(C \vee D)$ from these premisses. Which of the following answers is correct?

    (a) Yes, the conclusion follows.

    (b) No, the conclusion does not follow, but if you add the premiss $C \Rightarrow E$ the conclusion can be derived.

    (c) No, the conclusion does not follow, but if you add the premiss $\neg(E \vee F)$ the conclusion can be derived.

    (d) No, the conclusion does not follow, but if you add the premiss $\neg A \vee \neg B$ the conclusion can be derived.

14. The unification algorithm with occurs-check will fail for exactly one of the following pairs of sentences. Which pair is that?

    (a) $Sibling(Father(x), y)$ and $Sibling(x, Father(y))$

    (b) $Sibling(Father(x), Father(y))$ and $Sibling(x, y)$

    (c) $Sibling(Father(x), Mother(y))$ and $Sibling(Mother(y), x)$

    (d) $Sibling(x, y)$ and $Sibling(v, z)$

15. Consider the problem of planning the assembly of bookshelves, such that every bookshelf needs 20 screws. The screws are modelled as a consumable resource in a STRIPS-formulation of the problem. There is also a *Refill* action that supplies screws if they run out. The problem is how to prevent POP from inserting *Refill* actions when in fact no such action is needed. Which of the following statements is true?

    (a) There is no need to take special measures, POP will not insert redundant *Refill* actions.

    (b) A special precondition has to be incorporated in the *Refill* action such that a *Refill* is only needed when there are too few screws.

    (c) POP cannot solve this problem, one needs a modified version of POP for this.

    (d) None of the above is true.

16. Consider the STRIPS definition of the planning problem of making tea in figure 1. In the planning graph for this problem, S_0 is identical to the *Init*-state of the formal problem definition. How many mutex links will there be in state S_1?

    (a) 0

    (b) 1

    (c) 2

    (d) 3

Action(*Init*,
  PRECOND: ,
  EFFECT: $ColdWater \wedge Kettle \wedge Teabag \wedge TeapotEmpty$)

Action(*Finish*,
  PRECOND: $TeaReady$,
  EFFECT:)

Action(*BoilWater*,
  PRECOND: $ColdWater \wedge Kettle$,
  EFFECT: $HotWater \wedge \neg ColdWater$)

Action(*MakeTea*,
  PRECOND: $Teabag \wedge TeapotFilled$,
  EFFECT: $TeaReady \wedge \neg TeapotFilled \wedge \neg Teabag$)

Action(*PourWater*,
  PRECOND: $HotWater \wedge TeapotEmpty$,
  EFFECT: $TeapotFilled \wedge \neg TeapotEmpty \wedge \neg HotWater$)

Figure 1: STRIPS representation of a plan for making tea, to be used in Question 16.

Action(*Start*,
PRECOND: ,
EFFECT: $HaveEggs \wedge HaveButter \wedge HavePan$)

Action(*Finish*,
PRECOND: *EggsFried*,
EFFECT:)

Action(*ScrambleEggs*,
PRECOND: *HaveEggs*,
EFFECT: $EggsScrambled \wedge \neg HaveEggs$)

Action(*ButterInPan*,
PRECOND: $HavePan \wedge HaveButter$,
EFFECT: $PanButtered \wedge \neg HavePan \wedge \neg HaveButter$)

Action(*FryEggs*,
PRECOND: $PanButtered \wedge EggsScrambled$,
EFFECT: $EggsFried \wedge \neg PanButtered \wedge \neg EggsScrambled$)

Figure 2: STRIPS-definition of the problem of planning the preparation of scrambled eggs. Belongs to question 17.

17. Consider the STRIPS definition of the problem of making scrambled eggs in figure 2. We use the POP algorithm to generate a partial-order plan. Which of the following statements is true of POP's output?

    (a) The plan is already completely linear.

    (b) The plan branches at precisely one point.

    (c) The plan branches at precisely two points.

    (d) The STRIPS-definition is not sufficient for constructing a complete and consistent partially ordered plan.

18. Non-consumable resources, such as a forklift, can be incorporated in STRIPS by means of an extra field RESOURCE: x. The question is why this is an extra field rather than an addition to the EFFECTS field. Which of the following statements is true?

    (a) There is no formal reason to keep resources and effects apart, but doing so makes the representation easier to read.

    (b) A non-consumable resource does not have to be modelled as a separate field provided it is specified in the PRECOND field.

    (c) A non-consumable resource is not an effect because once the action is done it is available again.

    (d) Neither of the above is true.

19. We have a potted plant that lives outside. Although a hardy type of plant, it can suffer from exposure to a frost. We have three possible actions that can be performed: do nothing, provide simple insulation for the pot, or place the plant in a greenhouse. The chance for survival given each of the actions are shown in table 1.

Action	Survive	
	Yes	No
Do Nothing	0.6	0.4
Simple Insulation	0.85	0.15
Greenhouse	0.95	0.05

Table 1: Probabilities of the potted plant surviving given the actions for providing insulation against frost.

    The utility is simply what the plant is worth to us, that is 80 euro if the plant survives. If the plant dies we have 0 euro. Since we don't possess any material or greenhouse we need to buy it. Cost of a greenhouse is 100 euro. Cost of insulation is 10 euro. What is the optimum action ?

    (a) do nothing

    (b) simple insulation

    (c) greenhouse

    (d) undecided

20. In the course of one year 8 elderly women died unexpected and without a clear medical cause, all in the same hospital. A jury uses the following argumentation to motivate why she considers the crime suspect, a hospital nurse, guilty of the murder of the 8 elderly women. The argument consists of the following 4 clauses.

    i) The suspect was on duty in each of the eight cases of death.

    ii) There are no medical explanations for the death of the women.

    iii) The probability that one out of 100 nurses is on duty in each of the eight cases is so very small (less than one in a million) that the death of the women cannot be by mere chance.

    iv) The nurse must have murdered the eight women.

    The general argument is something like: if there is a very unlikely coincidence between two events, then there is a causal relation between these two events. Which of the following statements is true?

    (a) The general argument is valid: there is a causal relation.

    (b) The general argument is not valid: there need not be a causal relation.

    (c) If there is a direct causal relation between event A and event B then the chance that A and B occur simultaneously is high.

    (d) If there is no causal relation between two events then the chance that that A and B occur simultaneously is low.

21. Consider the following simple-chain Bayesian network, in which the three nodes represent three boolean-valued stochastic variables.

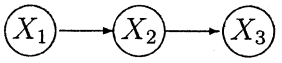

    The following probabilities hold in this network.

    - $P(X_1 = true) = 0.5$
    - $P(X_2 = true | X_1 = true) = 1.0$
    - $P(X_2 = true | X_1 = false) = 0.5$
    - $P(X_3 = true | X_2 = true) = 1.0$
    - $P(X_3 = true | X_2 = false) = 0.5$

    What is the value of $P(X_3 = true)$?

    (a) 1.0

    (b) 0.5

    (c) 0.875

    (d) 0.125

22. What is the value of $P(X_3 = true | X_1 = true)$ in the simple-chain network of question 21 above?

a) 1.0

b) 0.5

c) 0.875

d) 0.125

23. Consider the following Bayesian Network with boolean variables:

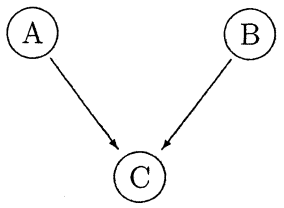

Figure 1

Consider the following statements:

i) A and B are independent (without any information about the value of C)

ii) A and B are dependent given C

Which of the above statements are true?

a) only i is true

b) i as well as ii is true

c) only ii is true

d) neither i nor ii is true

24. Consider the following Bayesian Network:

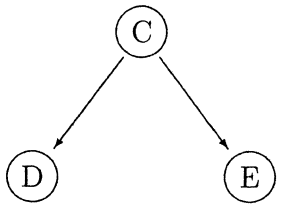

Figure 2

And consider the following statements about this Bayesian Network

(i) D and E are independent (without any information about the value of C)

(ii) D and E are dependent given C

Which of the above two statements about are true?

(a) only (i) is true

(b) (i) as well as (ii) is true

(c) only (ii) is true

(d) neither (i) nor (ii) is true

25. In the Bayesian Network below with three boolean variables the probabilities for P and M are: $P(M = true) = 0,1$ and $P(L = true) = 0.7$ and the conditional probabilities for variable V are as shown in the table.

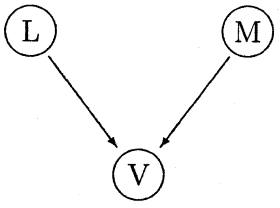

L	M	$P(V = true \mid L, M)$
true	true	0,9
true	false	0,5
false	true	0,3
false	false	0,05

What is the value of $P(V = true \mid M = true)$?

(a) 0.72

(b) 0.02

(c) 0.9

(d) 0.875

26. A en B are propositions and $P(A|B) = x$ (short for: $P(A = true|B = true) = x$). Which of the following statements are necessarily true?

(i) $P(B \Rightarrow A) = x$

(ii) $B \Rightarrow P(A) = x$

(iii) $P(A|\neg B) = 1 - x$

(iv) $P(\neg A|B) = 1 - x$

(a) only (i) en (ii)

(b) only (i) en (iii)

(c) only (iv)

(d) only (ii) en (iv)

27. A retailer wants for marketing purposes distinguish between costumers younger then 35 and customers older then 35. The following table summarizes the data set in the data base of the retailer in an abstract form. The relevant attributes, determined by domain knowledge, are for convenience denoted by A with values *a1*, *a2* and *a3*, B with values *b1* and *b2*, C with values *c1* and *c2* and D with values *d1* and *d2*

A	B	C	D	Number of Instances	
				Y	O
a1	b1	c1	d1	12	2
a2	b1	c1	d2	2	6
a3	b1	c1	d1	6	0
a1	b2	c1	d2	0	12
a2	b2	c1	d1	4	2
a3	b2	c1	d2	0	4
a1	b1	c2	d1	0	8
a2	b1	c2	d2	8	0
a3	b1	c2	d1	4	0
a1	b2	c2	d2	0	4
a2	b2	c2	d1	2	2
a3	b2	c2	d2	5	0

Assume that the retailer wants to use Decision Trees to classify the costumers in the class "young", denoted by Y, and "old", denoted by O.

What is the Information Gain of attribute D?

(a) 0

(b) 1

(c) 0.0664

(d) 0.9336

28. The retailer wants to learn the above classification problem, exercise 27, using decision trees. If he uses "information gain" as selection criteria what will be the first attribute for splitting the examples?

    (a) A

    (b) B

    (c) C

    (d) D

29. Consider an ensemble learning algorithm that uses simply majority vote among 3 learned hypothesis. Suppose that each hypothesis has error ε and that the errors made by each hypothesis are independent of each other.
    What will be the error of the ensemble algorithm?

    (a) 3ε

    (b) ε^3

    (c) $\varepsilon^3 + 3\varepsilon^2(1-\varepsilon)$

    (d) $3\varepsilon^2(1-\varepsilon)$

30. An agent interacts with a *deterministic* environment and it's current state is s which has estimated utility value 4. The for the agent reachable states from s with their corresponding utility value are:

state	*utility*
s_1	2
s_2	4
s_3	6
s_4	3

    Furthermore assume that $R(s) = 2$. What will be the next state of the agent and what will be the new utility value $U(s)$ of s if the agent applies temporal difference learning with learning parameter $1/2$ and discount factor 1, i.e $\alpha = \frac{1}{2}$ and $\gamma = 1$?

    (a) Next state will be s_1 and the new utility for s will be $U(s) = 2$

    (b) Next state will be s_2 and the new utility for s will be $U(s) = 4$

    (c) Next state will be s_3 and the new utility for s will be $U(s) = 5$

    (d) Next state will be s_3 and the new utility for s will be $U(s) = 6$

31. An agent uses Q-learning, to learn an optimal strategy for a probabilistic game. The current (internal) state of the game is s. In this state S the agent can do four actions; a, b, c and d. The Q-value for these state action pairs are given by:

action x	$Q(s,x)$
a	70
b	50
c	20
d	60

What will be the next action of the agent if the actions applies a greedy strategy?

(a) action a

(b) action b

(c) action c

(d) action d

32. Once again consider the above Assume that the agent decides to take the action b (maybe due to exploration) and receives a reward 10. Due to this action b the agent ends up in state s'. In this new state s' the agent can do actions d, e and f with the following Q-values:

action x	$Q(s',x)$
d	30
e	20
f	10

Assume that the agent applies Temporal Difference Learning with learning parameter and discount factor equal to 1. What will be the new Q-values for state s?

(a) $Q(s,a) = 40$, $Q(s,b) = 40$, $Q(s,c) = 40$ and $Q(s,d) = 40$.

(b) $Q(s,a) = 70$, $Q(s,b) = 50$, $Q(s,c) = 20$ and $Q(s,d) = 60$.

(c) $Q(s,a) = 40$, $Q(s,b) = 50$, $Q(s,c) = 20$ and $Q(s,d) = 60$.

(d) None of the above.