Examination
Artificial Intelligence
(214030)
June 30, 2009

Introduction

This examination consists of 37 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. Consider the following statements about agent architectures and decision trees:

   (i) A Decision Tree Algorithm is **not** programmable on a Simple Reflex Agent.

   (ii) A Decision Tree Algorithm is programmable on a Goal Based Agent.

   (iii) A Decision Tree Algorithm is programmable on a Learning Agent

   Which of the following claims is true?

   (a) All statements (i), (ii) and (iii) are false.

   (b) Only statement (iii) is true

   (c) Only statement (i1) and (iii) are true

   (d) All statements (i), (ii) and (iii) are true.

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

   (i) If Depth First Search Algorithm terminates on the search problem **then** 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.

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

   (a) Greedy best-first search

   (b) Depth-First Search

   (c) Uniform Cost Search

   (d) None of the above.

4. Consider the sliding puzzle below, with start and goal state as given.

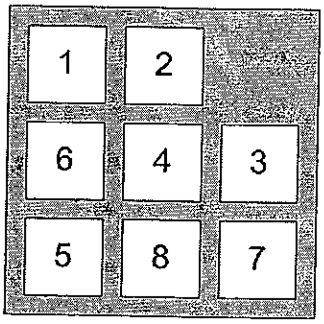

Start state

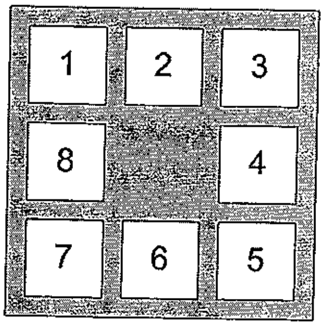

Goal state

The legal actions are the sliding of a tile to a neighboring empty square. The cost of moving tile with number i is i. For example the cost of sliding tile (with number) 2 is 2.

A corresponding heuristic "cost to go" function $h(n)$ is defined as follows. Let s_n be the state corresponding to node n. For each tile i in the state s_n let $d(i, s_n, s_{goal})$ denote the Manhattan distance between the position of tile i in state s_n and the position of tile i in the goal state s_{goal}. Then

$$h(n) = \sum_{i=1}^{i=8} i \times d(i, s_n, s_{goal})$$

What is the value $h(n)$ of nodes n corresponding to the initial state?

(a) 10

(b) 32

(c) 59

(d) 52

5. Assume that A* search is applied to the above sliding puzzle. Which of the following nodes will be the **third** node that will be expanded?

(a) The node corresponding to the state which arises from the initial state by moving first tile 3 and afterwards again moving tile 4 to the empty square.

(b) The node corresponding to the state which arises from the initial state by moving first tile 3 and afterwards moving tile 7 to the empty square.

(c) The node corresponding to the state which arises from the initial state by moving first tile 2 and afterwards moving 4 to the empty square.

(d) The node corresponding to the state which arises from the initial state by moving first tile 2 and afterwards moving 1 to the empty square.

6. Once again consider the above sliding puzzle including the "cost to go" function h and the function $f(n) = g(n) + h(n)$, with $g(n)$ the cost function of the path from the root of the search tree to the node n.
Consider the following statements about h:

(i) h is admissible.

(ii) h is consistent.

Which of the above statements are true?

(a) Only hypothesis (i) is true.

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

(c) Only hypothesis (ii) is true.

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

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

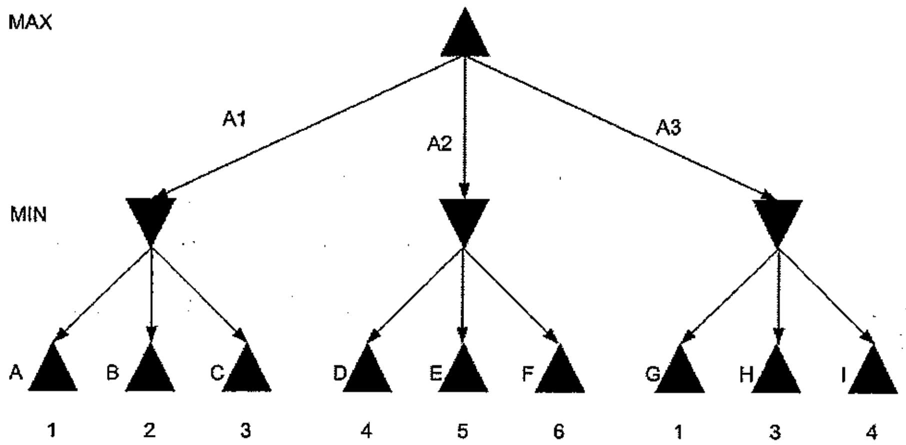

What will be the value of the top MAX node

(a) 6

(b) 5

(c) 4

(d) 3

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

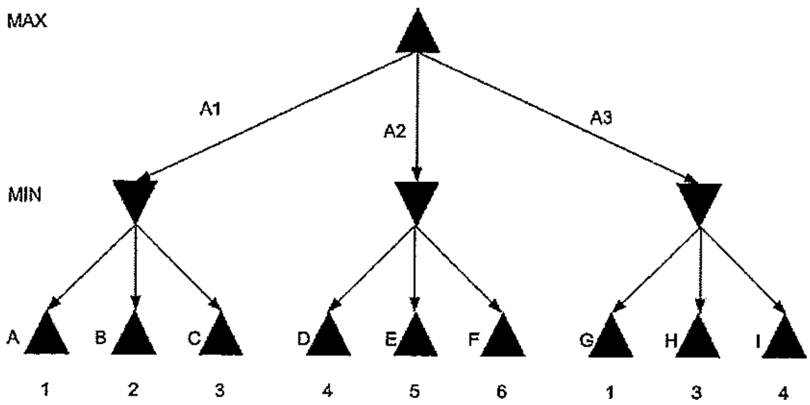

Assume one applies alpha-beta pruning. Which of the following collection of nodes will **all not** being explored?

(a) $\{E, H, I\}$

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

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

(d) None of the above is correct.

9. Consider the following **zero sum game** tree for three players, player 1, 2 and 3.

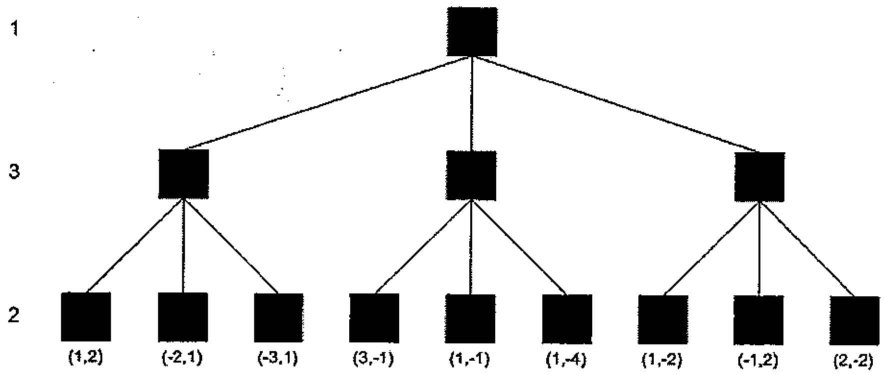

At the lowest nodes the evaluation value of the relevant node is given (the value of the evaluation function) for player 1 and 2. In the i position the evaluation value of the node for player i is given, $i = 1, 2$. What is the values of the nodes for player 1 in the above representation?

(a) $(3, -1)$

(b) $(-3, 1)$

(c) $(1, -1)$

(d) $(1, -4)$

10. 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.

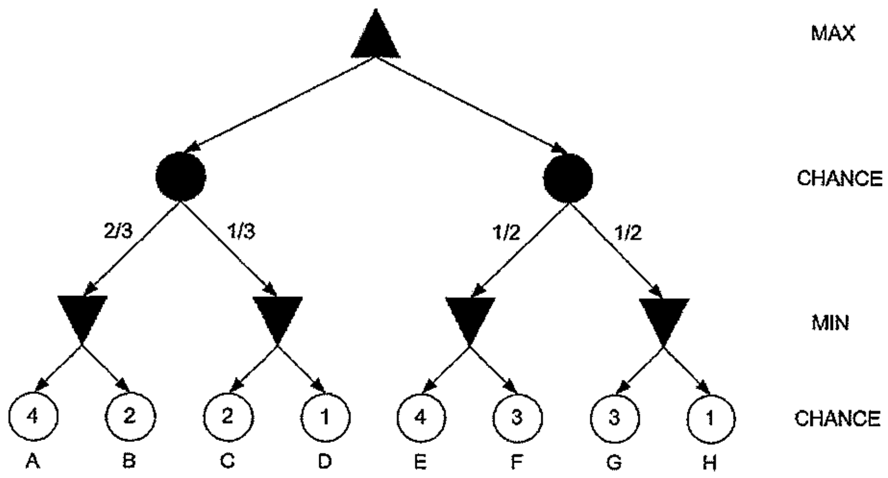

What is the correct value for the top MAX-node if one applies the expectiminimax algorithm?

(a) 4

(b) 3

(c) 2

(d) 1

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

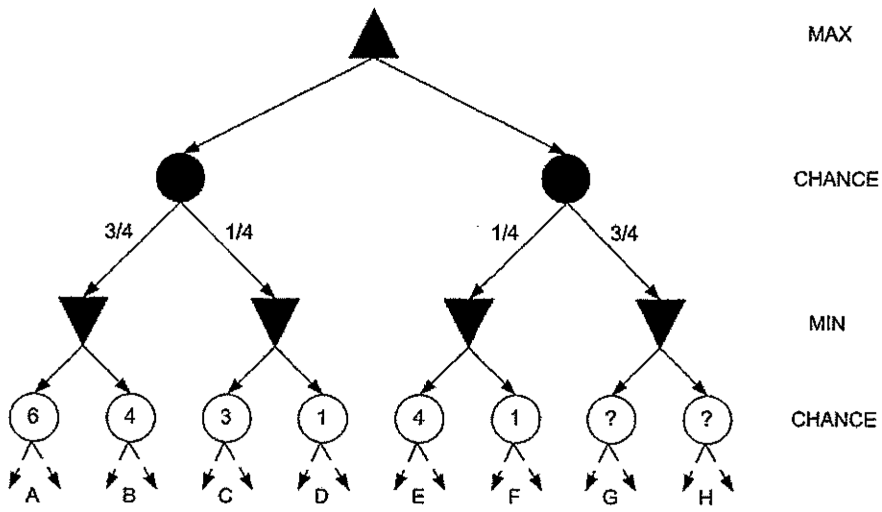

The letters under the bottom row of chance nodes are labels for the nodes just above the letter. One can also apply $\alpha - \beta$ pruning on this game tree. The numbers inside the chance nodes on the bottom row are the computed values of these chance nodes; ? indicates not computed yet. For which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

(a) If the value of node G is less then 3 then $\alpha - \beta$ pruning will **not** expand node H.

(b) If the value of node G is greater or equal to 3 then $\alpha - \beta$ pruning will **not** expand node H.

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

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

12. We are given an arbitrary sentence in first-order logic. Let's call that sentence p. Now consider the following assertion:

> "We can find out whether p is a contradiction or a tautology by applying resolution plus unification on p."

Which of the following alternatives is true?

(a) The assertion is true only for contradiction.

(b) The assertion is true only for tautology.

(c) The assertion is true for both contradiction and tautology.

(d) The assertion is false for both contradiction and tautology.

13. Given are the following premisses:

   - $R \Rightarrow (\neg A \wedge T)$
   - $B \vee \neg S$

   The question is whether we can prove $A \Rightarrow B$ 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 $T \vee A$ the conclusion can be derived.

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

   (d) No, the conclusion does not follow, but if you add the premiss $R \vee S$ 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) $Same(ColourOf(x),\ y)$ and $Same(Red,\ ColourOf(x))$

   (b) $Same(ColourOf(x),\ Red)$ and $Same(ColourOf(x),\ y)$

   (c) $Same(x,\ y)$ and $Same(y,\ x)$

   (d) $Same(ColourOf(x),\ y)$ and $Same(x,\ ColourOf(y))$

15. We want to formalise the sentence "No student who attends the Artificial Intelligence course is stupid" in first-order logic. Given are the following predicates:

   - $S(x)$: x is a student.
   - $A(x, y)$: x attends course y.
   - $D(x)$: x is stupid.

   We further introduce a term constant, AI, that stands for the Artificial Intelligence course. We are given two different formalisations:

   I. $\forall x\ S(x) \wedge A(x,\ AI) \Rightarrow \neg D(x)$
   II. $\exists x\ \neg (S(x) \wedge A(x,\ AI) \wedge D(x))$

   Which of the following assertions is correct?

   (a) Formalisation I is correct while formalisation II is incorrect.

   (b) Formalisation I is incorrect while formalisation II is correct.

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

   (d) The two formalisations are inferentially equivalent and they are both incorrect.

16. When will the GRAPHPLAN algorithm terminate?

   (a) When a further iteration does not add new literals.

   (b) When a state contains all literals that have to be true in the goal state, and no pair of those literals is connected by a mutex link.

   (c) Both (a) and (b) are correct. What counts is whichever comes first.

   (d) Neither (a) nor (b) is correct, there is another criterion for termination.

17. The planning problem representation languages STRIPS and ADL assume that all changes in the world during the execution of a plan are consequences of the actions defined in the representation of the planning problem. We will call this the *static-world assumption.* We all know that the static-world assumption is often naive because the world can very well change on its own. In this question, we will explore whether we can make these languages a little more expressive to incorporate changes in the world. Which of the following statements is correct:

   (a) This is impossible in STRIPS because the static-world assumption is essential for STRIPS, but it is possible in ADL.

   (b) This is impossible in either STRIPS and ADL.

   (c) This is possible by incorporating a feature in STRIPS, that allows one to define sensing actions of which the outcome depends on what the sensor senses.

   (d) Neither of the statements (a), (b), and (c) is correct.

18. In planning, the *frame problem* plays an important role. Briefly, the frame problem is about what changes in the world and what remains constant during plan execution. Which of the following assertions is true:

   (a) Planning in first-order predicate logic (situation calculus) cannot solve the frame problem.

   (b) The POP algorithm solves the frame problem in a meta-logical way.

   (c) Both (a) and (b) are correct.

   (d) Neither (a) nor (b) are correct.

Op(ACTION: *Fill*,
PRECOND: *BottleEmpty* ∧ *BottleUpright* ∧ *BottleOpen*,
EFFECT: ¬*BottleEmpty* ∧ *BottleFilled*)

Op(ACTION: *Finish*,
PRECOND: *BottleClosed*,
EFFECT:)

Op(ACTION: *CloseBottle*,
PRECOND: *BottleOpen* ∧ *BottleUpright* ∧ *BottleFilled*,
EFFECT: ¬*BottleOpen* ∧ *BottleClosed*)

Op(ACTION: *Start*,
PRECOND: ,
EFFECT: *BottleLying* ∧ *BottleOpen* ∧ *BottleEmpty*)

Op(ACTION: *Upright*,
PRECOND: *BottleLying* ∧ *BottleEmpty* ∧ *BottleOpen*,
EFFECT: *BottleUpright*)

Figure 1: The STRIPS definition of the planning process of the production of *Blarf Cola*. Belongs to question 19.

19. The company *Blarf* wants to plan the production process of its primary product, *Blarf Cola*. The STRIPS-representation of that plan can be found in figure 1. We will make a planning graph for this problem. Assume that the S_0 state corresponds to the *Init* action in the STRIPS representation. How many mutex links will there be in state S_1?

    (a) 0
    (b) 1
    (c) 2
    (d) 3

20. Let's suppose you end up in Hell. (don't take this personal. it's just an exercise.) After a few days the Devil appears with an offer. God has commanded that you be shown some mercy. So the Devil has hatched a plan. He will give you *one chance* to get out of Hell. You can toss a coin; if it comes down heads, you are out and go to Heaven (that is the place you want to be, not in Hell; that's also being supposed). If it comes down tails, you stay in Hell forever.
The coin is not a fair one, however, and the Devil has control over the odds. If you toss it *today*, the chance of heads is 1/2 (that is $1 - 1/2$). If you *wait till tomorrow*, the chances go up to 3/4 (that is $1 - 1/2^2$). Escaping has a very large positive value, 10^6 and staying in Hell has a very large negative value -10^6. Moreover, these values are the same today and tomorrow. Of course, if you decide to toss tomorrow you stay one day more in Hell, but that is negligible. Your chances are higher to escape tomorrow, and you only have one chance. The trouble is that the Devil comes to you tomorrow and says that if you wait till tomorrow the odds to escape are even higher 7/8 (that is $1 - 1/2^3$). And the odds get better day by day, as follows:

$$1 - 1/2, 1 - 1/2^2, 1 - 1/2^3, ...1 - 1/2^n, ...$$

Suppose you are rational, you have studied AI, and every day you make your decision based on the principle of maximum expected utility. What will be the outcome?

(a) You decide to take your chance on the first day

(b) You decide to take your chance after day 12

(c) You decide to take your chance on day 5

(d) Every day you decide to take your chance tomorrow, (and you stay in Hell for ever).

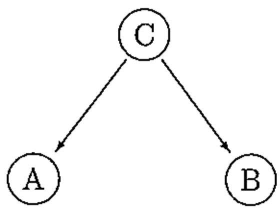

Figure 2: Bayesian Net One

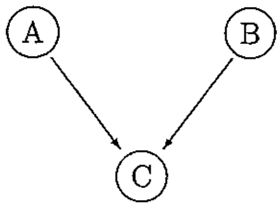

Figure 3: Bayesian Net Two

21. The inference

$$P(A \wedge B) = P(A).P(B) \Rightarrow P(A \wedge B|C) = P(A|C).P(B|C)$$

   (a) holds for the Bayesian network One in Figure 2 as well as for Two in Figure 3
   (b) holds neither for the Bayesian network One in Figure 2 nor for Two in Figure 3
   (c) holds for One, not for Two.
   (d) holds for Two, not for One.

22. Consider the (fully observable) environment as specified in Figure 17.1 of the book of Russel and Norvig (page 614). We apply the value iteration algorithm (see Figure 17.4) for the given MDP with discount factor $\gamma = 1$ for computing (an approximation of) an optimal policy. We start iteration with initial utilities $U_0(s) = 0$ for all states s. Below are for some states the values of U_1 (i.e. utilities computed after the first Bellman update). Which ones are correct?

   (a) $U_1(1,1) = 0.0$, $U_1(3,3) = 0.0$, $U_1(4,3) = 0.0$
   (b) $U_1(1,1) = 0.0$, $U_1(3,3) = 0.0$, $U_1(4,3) = 1.0$
   (c) $U_1(1,1) = -0.04$, $U_1(3,3) = -0.04$, $U_1(4,3) = -0.04$,
   (d) $U_1(1,1) = -0.04$, $U_1(3,3) = -0.04$, $U_1(4,3) = 1.0$,

23. In the previous exercise we stated that we take discount $\gamma = 1$. What about the outcome of the value iteration algorithm in this case?

(a) it will never stop and return nothing

(b) it will only stop for specific reward functions R

(c) it will stop after the first iteration and output U_1.

(d) it will always return a utility function for the given MDP.

24. For the world of Figure 17.1, what is the probability of ending in state $(2, 2)$ from state $(1, 1)$ by the sequence of actions $[Up, Up, Right]$?

(a) 0.064

(b) 0.128

(c) 0.129

(d) other

25. A full joint distribution for the *Toothache, Cavity, Catch World* is given by the table below, copied from Figure 13.3 in the book of Russel and Norvig.

	toothache		¬ toothache	
	catch	¬catch	catch	¬catch
cavity	0.108	0.012	0.072	0.008
¬cavity	0.016	0.064	0.144	0.567

A full joint distribution for the "Toothache, Cavity, Catch World"

How will a Bayesian network for this world look like, if we look at the conditional independencies that follow from the probabilities?

(a) The BN is a linear model with Toothache as child node of Catch and Catch as child node of Cavity

(b) like (a) but Catch is also parent of Toothache.

(c) Toothache has two parents Cavity and Catch, and there are no other arrows in the network.

(d) Catch and Toothache both have one parent node: Cavity.

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

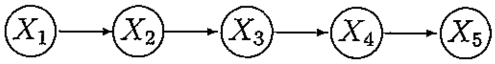

One of the following statements is true. Which one?

(a) For computing $P(X_3|X_1)$ we may remove only node X_5.

(b) For computing $P(X_3|X_1)$ we may remove node X_4 as well as node X_5.

(c) For computing $P(X_4|X_2)$ we may remove $X1$ as well as X_5.

(d) For computing $P(X_4|X_2)$ we need to take the complete network into account.

27. One of 4 undistinguishable boxes contains an amount of 1 million euro. You may buy one of the boxes for the price of a quarter of a million euros. Someone offers you the information about the contents of box number 1, telling you whether or not it contains the million euros. What is the maximum price that you should be willing to pay for this information based on the value of this information?

(a) 1 million euro

(b) 250.000 euro

(c) 500.000 euro

(d) other

28. 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.

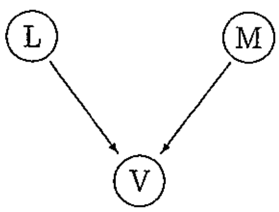

L	M	$P(V = true \mid L, M)$
true	true	0,9
true	false	0,5
false	true	0,3
false	false	0,05

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

(a) 0.075

(b) 0.54

(c) 0.46

(d) 0.925

29. A witness of an nighttime accident involving a taxi declares that the taxi was blue.

All taxis in town are blue or green. It is known that under dim lighting conditions discrimination between blue and green is 75% reliable; which means that $P(LB|B)$ as well as $P(\neg LB|\neg B)$ are 0.75, where B is the boolean variable for the predicate *"the taxi is blue"* and LB is the boolean variable for the predicate *"the taxi looked blue"*.

Suppose that 9 out of 10 taxis are actually green. Given the declaration of our witness what is the probability that the taxi is indeed blue?

(a) 0.25
(b) 0.375
(c) 0.9
(d) 0.75

30. Consider the following two logical sentences

(i) $(p \wedge q \wedge s) \vee (\neg r \wedge r \wedge s)$
(ii) $(p \vee q \vee s) \wedge (\neg r \vee r \vee s)$

Which of the following statements are true?

(a) Both sentences can be represented by a decision tree.
(b) Only sentence (i) can be represented by a decision tree.
(c) Only sentence (ii) can be represented by a decision tree.
(d) None of the sentences (i) and (ii) can be represented by a decision tree.

31. An analyst wants to automate a classification procedure. Therefore he has collected several classification examples, see the table below, the attributes are for simplicity denoted A, B, C, D.

nr	A	B	C	D	*classification*
1	0	0	0	1	0
2	0	0	1	1	0
3	0	1	0	0	0
4	0	1	1	0	1
5	0	1	0	1	1
6	1	1	1	1	1
7	1	0	0	0	1
8	1	1	0	0	1
9	1	0	1	0	1
10	1	0	0	1	0

What is the information provided by the attribute A?

(a) 1.0000

(b) 0.5000

(c) 0.9710

(d) 0.0290

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

(a) B

(b) D

(c) A

(d) C

33. Consider an ensemble learning algorithm that uses simply majority vote among 5 learned hypothesis. Assume that:

- each hypothesis has error ε
- the errors made by each hypothesis are independent of each other.

What will be the error of the ensemble algorithm?

(a) 5ε

(b) ε^5

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

(d) None of the above.

34. Assume an agent applies a learning algorithm to learn a Boolean function of n variables. How many mutually different training examples does the agent need in order to learn the complete function?

   (a) 2^{2^n}
   (b) 2^n
   (c) n
   (d) none of the above

35. Consider the following statements about Reinforcement Learning:

   (i) In a passive reinforcement learning context the policy of an agent is fixed.
   (ii) In a passive reinforcement learning context the utility function of an agent is fixed.
   (iii) In an active reinforcement learning context the policy can be adapted by the agent.
   (iv) In an active reinforcement learning context the utility function can be adapted by the agent.

   Which of the following statements are true?

   (a) Only statement (ii) is false
   (b) Only statement (iii) and (iv) are true.
   (c) Only statement (ii), (iii) and (iv) are true.
   (d) All statements (i), (ii), (iii) and (iv) are true.

36. Consider the 4×3 grid world described in Chapter 17 and Section 21.2 of the course book. We make the following assumptions:

   - the agent is in state $(2, 1)$.
   - the agent follows the policy π described in Fig. 21.1 (b)
   - The estimated utilities are $U^{\pi}(2,1) = 0.640$ and $U^{\pi}(1,1) = 0.920$

   Assume that action $\pi(2, 1)$ is successful, what will be the new value of $U^{\pi}(2,1)$ if the agent applies Temporal Difference Learning with $\alpha = 0.25$ and $\gamma = 1$?

   (a) 0.640
   (b) 0.710
   (c) 0.920
   (d) None of the above

37. 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	40
b	20
c	30
d	50

Assume that the agent decides to do some explorations and does the action b 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	80
e	40
f	70

Assume that the agent applies Temporal Difference Learning with learning parameter $\alpha = 0.5$ and discount factor $\gamma = 1$. What will be the new Q-values for state s?

(a) $Q(s, a) = 60$, $Q(s, b) = 50$, $Q(s, c) = 55$ and $Q(s, d) = 65$.

(b) $Q(s, a) = 40$, $Q(s, b) = 50$, $Q(s, c) = 30$ and $Q(s, d) = 50$.

(c) $Q(s, a) = 40$, $Q(s, b) = 80$, $Q(s, c) = 30$ and $Q(s, d) = 50$.

(d) None of the above.