26-01-2021 - Intelligent Interaction Design CS/BIT - AI. & Cyber Security Theory

Cursus: B-BIT-MOD06-1B-202001032 B-CS Intelligent Interaction Design CS/ BIT 202001032

Tijdsduur: 3 uur

Aantal vragen: 21

Gegenereerd op: 20 dec. 2021

Inhoud: Pagina's:

26-01-2021 - Intelligent Interaction Design CS/BIT - AI. & Cyber Security Theory

Cursus: B-CS Intelligent Interaction Design CS/BIT 202001032

This exam is closed book, you may only use a simple calculator (addition, substraction, multiplication, division and exponentiation) and a 1A4 cheat sheet.

This examination consists of multiple-choice questions, for which you have 3 hours.

Tips:
• Read each question carefully without looking at the possible answers.
• 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).
• Beware of double negations (negatives) as these can be confusing.
• Do not stay on any one question for 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 left over at the end, check your answers.
• As part of these instructions, and as part of the relevant question(s) of this exam, you can find a table with values for −p log 2 (p).

If you have questions about the content of the exam, you can contact the teachers via BigBlueButton. Make sure to address the teacher corresponding to the topic of the exam question that you want to ask something about, namely:

Birna van Riemsdijk: Q1-9

Gwenn Englebienne: Q10-15

Nacir Bouali: Q16-18

Andrea Continella: Q19-21

Please allow 10-15min for your question to be answered. If you have not yet heard from a teacher by then, please contact a TA or teacher in the room. Use the information button on the screen to see this info again while making the exam."

Good luck!

Aantal vragen: 21

In totaal zijn 24 punten voor deze toets te behalen, 15,9 punten zijn nodig om voor de toets te slagen.

1 The AI field comprises a range of computational techniques which each have advantages and disadvantages.
1 pt.

What can be a disadvantage of the use of machine reasoning?

a. Machine reasoning requires explicit knowledge representation.

b. Machine reasoning needs a lot of data.

c. Machine reasoning is explainable.

d. Machine reasoning requires use of many different techniques.

2 Consider the propositional logic formula $P \Rightarrow ((Q \vee \neg P) \Rightarrow \neg P)$.
1 pt.

How many models does this formula have?

a. 0

b. 3

c. 4

d. 8

3 Consider the propositional logic formula $P \Rightarrow (\neg Q \wedge R)$.
1 pt.

Which of the following formulas is an equivalent formula in Conjunctive Normal Form (CNF)?

a. $(\neg P \wedge \neg Q) \vee (\neg P \wedge R)$

b. $(\neg P \vee \neg Q) \wedge (\neg P \vee R)$

c. $(P \vee \neg Q) \wedge (P \vee R)$

d. none of the above

4 A knowledge base KB contains the following statement (i.e. considered to be true).
1 pt.

$$\neg W \Rightarrow (Q \vee T)$$

The question is whether we can prove $Q \vee T$ from this KB . Which of the following answers is correct?

a. Yes, we can prove $Q \vee T$.

b. No, we cannot derive $Q \vee T$, but if we add the premiss W to KB the statement $Q \vee T$ can be derived.

c. No, we cannot derive $Q \vee T$, but if we add the premiss $\neg W$ to KB the statement $Q \vee T$ can be derived.

d. None of the above.

5 Given are the following predicates:
1 pt.

- $family(x, y)$: x is a member of family y
- $present\text{-}for(z, x)$: z is a Christmas present for x

Consider the translation of the sentence "there is a Christmas present for everyone in the Jones family" into predicate logic.

Which of the following translations is correct?

a. $\forall x\ family(x, Jones) \Rightarrow \exists z\ present\text{-}for(z, x)$

b. $\forall x\ family(x, Jones) \wedge \exists z\ present\text{-}for(z, x)$

c. $\exists x\ family(x, Jones) \wedge \exists z\ present\text{-}for(z, x)$

d. none of the above

6 1 pt.

Consider the following two predicate logic formulas, where x, y are variables and A, B, C are constants.

- $P(F(x), G(B, y))$
- $P(F(G(C)), G(x, A))$

Does a unifier exist for these formulas?

a. no unifier exists

b. yes, the unifier is $\{x/G(C), x/B, y/A\}$

c. yes, the unifier is $\{x/G(C), y/A\}$

d. yes, the unifier is $\{x/B, y/A\}$

7 1 pt.

We want to Skolemise the following sentence in first-order logic:

$$\forall x\ \exists y\ \forall z\ A(y, z)\ \Rightarrow\ [B(y)\ \wedge\ (C(x, y)\ \vee\ D(y, z))]$$

Only one the following four substitutions, produces a correct Skolemisation. Which one?

a. $\{y/S(x)\}$

b. $\{y/S(x, z)\}$

c. $\{y/S(z)\}$

d. $\{y/C\}$ with C a Skolem constant.

8 Consider the graph below. We want to search for a path from startnode S to goal node G using
1 pt. a **breadth-first** search algorithm. We denote the order in which nodes are visited when a search is performed as a sequence consisting of the names of nodes. For example, A, B, C expresses that node A is visited first, then node B , and then node C .

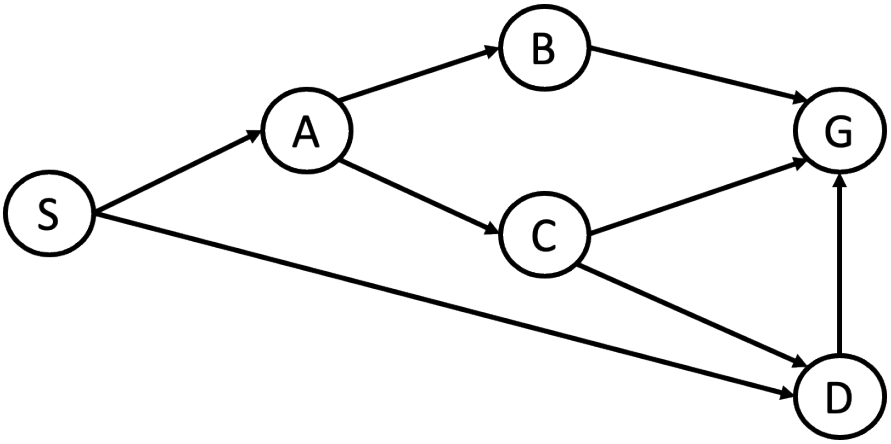

Which of the following orderings of visited nodes corresponds to a breadth-first search?

a. S,D,G

b. S,A,D,B,G

c. S,A,D,B,C,G

d. none of the above

9 Consider the following Tower of Hanoi problem. The game consists of three rods and a number of disks of different sizes, which can slide onto any rod. The puzzle starts with the disks in a neat stack in ascending order of size on one rod, the smallest at the top.

The objective of the puzzle is to move the entire stack to another rod, obeying the following simple rules:

1. Only one disk can be moved at a time.
2. Each move consists of taking the upper disk from one of the stacks and placing it on top of another stack or on an empty rod.
3. No larger disk may be placed on top of a smaller disk.

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.
We define the heuristic function *h* for a state/node as follows: 2 x *number of disks not on the rightmost peg*.

1 pt. **a.** What is the *h* value of the initial state?

a. 3

b. 4

c. 5

d. 6

1 pt. **b.** Assume that we apply ***A* tree*** search (i.e., without using a closed list) to the above Tower of Hanoi problem. Which node will **not** be in the frontier after two steps (expandig two nodes)?

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

10 Consider the following Bayesian network

1 pt.

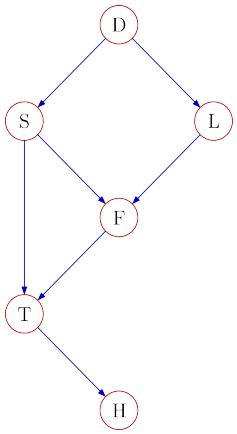

and corresponding statements:

I. Variables S and L are independent, no matter what other variables are observed in the network

II. Variables F and H are independent if T is observed

Which of the above statements is correct?

a. None of the statements

b. Both statements I and II

c. Only statement II

d. Only statement I

11
1 pt.

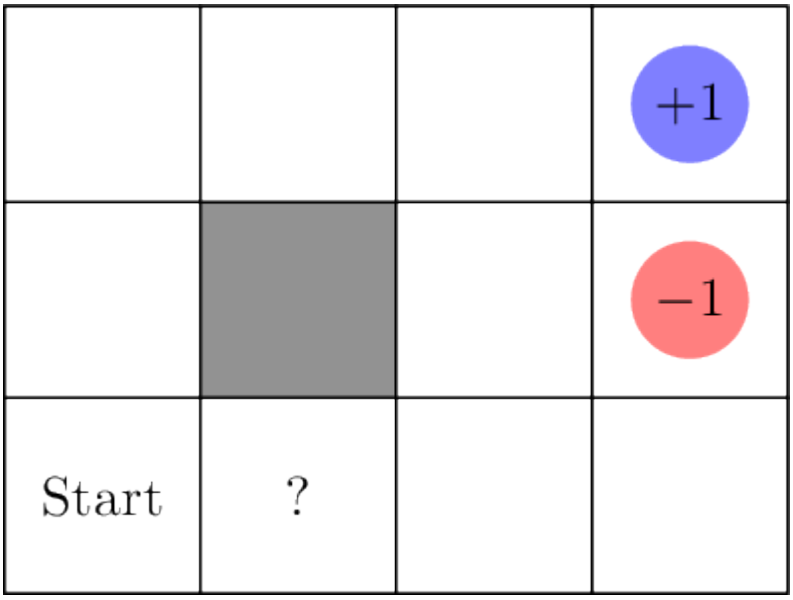

In the above grid world, the agent starts in the cell labelled "Start" and ends when it reaches one of the two terminal cells, which are labelled with a circle. It gets a reward of +1 in the cell with that label, and -1 in the cell the cell labelled "-1". There are four possible actions: *Up*, *Down*, *Left* and *Right*. These can be executed from any of the non-terminal states, and result in no movement if the agent hits a wall (these surround the world and surround the grayed-out cell). The world is stochastic, and the transition function has a probability of 0.8 of executing the desired action, and 0.1 of executing an action perpendicular to it. So, for example, if the agent is located in the cell marked with "?" and attempts to go "Up", it has 80% chance of ending up in the same cell marked "?", 10% chance of ending up in the cell to the right of the cell marked "?", and 10% chance of ending up in the cell marked "Start".

If all non-terminal cells get a reward of 0.01, what is the preferred action in the cell labelled "?" (as obtained by letting value-iteration run until convergeance):

a. All actions are equally desirable

b. Up

c. Left

d. Right

12 Consider the following statements:

1 pt.

(i) In order to apply the value iteration algorithm (Bellman update) the agent must know the reward function R and the successor state (transition) function δ .

(ii) In order to apply Q-learning the agent must know the successor state function δ but does not need to know the reward function R .

Which of the following claims is true?

a. Only statement (i) is true

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

c. Only statement (ii) is true

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

13 Consider the training a perceptron classifier with as activition function the **Threshold Function** (that is, **Heaviside step function**). Assume that the current linear classifier is given by the line $3 - 2x_1 + x_2 = 0$. The next feature point in our training set is given by $x = (-2, 2)$.

1 pt.

What will be the output of the classifier?

a. 9

b. 0

c. 1

d. None of the above.

14 Perceptrons are artificial neurons where the activation function is a threshold function (step function). Compared to later developments in artificial neural networks, they have a number of limitations. Which of the following is **not** a limitation of perceptrons?

1 pt.

a. If the perceptron learning algorithm does terminate, the resulting classifier may not be a good classifier (in the sense that it would perform well on previously unseen data)

b. The perceptron learning algorithm will never stop if the data is not linearly separable

c. Perceptrons can only learn linear classifiers, and combining perceptrons into a multi-layer network to deal with non-linear problems would be intractable (for reasonably large problems)

d. None of the above (All of the other answers are actually limitations of perceptrons)

15 Now assume that we have 4 urns which contain the following number of white, red and blue balls.

1 pt.

Urn	White	Red	Blue
H_1	5	4	1
H_2	2	5	3
H_3	5	2	3
H_4	3	3	4

One randomly selects an urn and draws 4 balls with replacement. Let the outcome of this chance experiment be $X = (X_w, X_r, X_b) = (1, 2, 1)$, i.e. 1 white ball, 2 red balls and 1 blue ball. What is the most likely urn selected for drawing the balls?

a. H_1

b. H_2

c. H_3

d. H_4

16 Consider the dataset below, with attribute A , which takes as value a1 and a2 and attribute B, which takes as value b1 and b2.
1 pt.

A	B	Number of Instances	
		Y	O
a1	b1	8	4
a2	b1	4	8
a1	b2	0	10
a2	b2	4	2

You can find a table of $-p \log2(p)$ values below.

p	$-p \log_2(p)$	p	$-p \log_2(p)$	p	$-p \log_2(p)$
0	0	1/8	0.38	1/10	0.33
1	0	2/8	0.50	2/10	0.46
1/2	0.50	3/8	0.53	3/10	0.52
1/3	0.53	4/8	0.50	4/10	0.53
2/3	0.39	5/8	0.42	5/10	0.50
1/4	0.50	6/8	0.31	6/10	0.44
2/4	0.50	7/8	0.17	7/10	0.36
3/4	0.31	1/9	0.35	8/10	0.26
1/5	0.46	2/9	0.48	9/10	0.14
2/5	0.53	3/9	0.53	1/11	0.31
3/5	0.44	4/9	0.52	2/11	0.45
4/5	0.26	5/9	0.47	3/11	0.51
1/6	0.43	6/9	0.39	4/11	0.53
2/6	0.53	7/9	0.28	5/11	0.52
3/6	0.50	8/9	0.15	6/11	0.48
4/6	0.39			7/11	0.42
5/6	0.22			8/11	0.33
1/7	0.40			9/11	0.24
2/7	0.51			10/11	0.13
3/7	0.52				
4/7	0.46				
5/7	0.35				
6/7	0.19				

what is entropy of the dataset above with respect to the class labels Y and O. Choose the alternative which is closest to your answer.

a. 0.77

b. 0.87

c. 0.97

d. 0.99

17 Consider the dataset below, with attribute A , which takes as value a1 and a2 and attribute B, which
1 pt. takes as value b1 and b2.

A	B	Number of Instances	
		Y	O
a1	b1	8	4
a2	b1	4	8
a1	b2	0	10
a2	b2	4	2

You can find a table of $-p \log2(p)$ values below.

p	$-p \log_2(p)$	p	$-p \log_2(p)$	p	$-p \log_2(p)$
0	0	1/8	0.38	1/10	0.33
1	0	2/8	0.50	2/10	0.46
1/2	0.50	3/8	0.53	3/10	0.52
1/3	0.53	4/8	0.50	4/10	0.53
2/3	0.39	5/8	0.42	5/10	0.50
1/4	0.50	6/8	0.31	6/10	0.44
2/4	0.50	7/8	0.17	7/10	0.36
3/4	0.31	1/9	0.35	8/10	0.26
1/5	0.46	2/9	0.48	9/10	0.14
2/5	0.53	3/9	0.53	1/11	0.31
3/5	0.44	4/9	0.52	2/11	0.45
4/5	0.26	5/9	0.47	3/11	0.51
1/6	0.43	6/9	0.39	4/11	0.53
2/6	0.53	7/9	0.28	5/11	0.52
3/6	0.50	8/9	0.15	6/11	0.48
4/6	0.39			7/11	0.42
5/6	0.22			8/11	0.33
1/7	0.40			9/11	0.24
2/7	0.51			10/11	0.13
3/7	0.52				
4/7	0.46				
5/7	0.35				
6/7	0.19				

What is the information gain with respect to the attribute B. Choose the alternative which is the closest to your answer.

a. 0.04

b. 0.06

c. 0.08

d. 0.1

18 A certain classifier was tested on a test set, resulting in the following confusion matrix:

		Predicted class		
		C_1	C_2	C_3
Actual Class	C_1	140	20	19
	C_2	11	120	9
	C_3	15	10	130

1 pt. **a.** What is recall for class C_2 ?

a. 0.86

b. 0.80

c. 1.00

d. 0.75

1 pt. **b.** What is the precision for class C_1

a. 0.92

b. 0.84

c. 0.78

d. 0.50

1 pt. **c.** What is the accuracy of this classifier?

a. 0.91

b. 0.72

c. 0.82

d. 0.89

19
1 pt.

Consider a one-class SVM defined by a training set D , a kernel k , and a regularization parameter v . Which of the following statements is correct?

a. If $v = 0.5$, then the hypersphere *always* encloses all the samples in the training set D .

b. When optimising the constrained problem $\min_{R,\mu,\xi} R^2 + \frac{1}{vn}\sum_{i=1}^{N}\xi_i$, selecting a smaller value for v decreases the squared radius of the hypersphere.

c. If $v > 0$, then some feature vectors that are *not* enclosed in the hypersphere can be predicted as “normal” (i.e., “not anomalous”).

d. If k is a linear kernel, then all the samples in the training set D are support vectors.

20
1 pt.

Consider a NIDS based on a one-class SVM model $\theta = (\mu_+)$, with $\mu_+ = (0.3, 0.3)$, and a prediction function $f_\theta(z) = ||z - \mu_+||_2 \geq \tau$, where $||z - \mu_+||_2$ indicates the Euclidean distance. Also, consider the following normalized feature vectors, together with the samples' labels, used for testing the learned model.

	Label	**Feature 1**	**Feature 2**
Sample 1	normal	0.1	0.3
Sample 2	normal	0.3	0.5
Sample 3	anomalous	0.9	1
Sample 4	anomalous	0.7	0.3

For which value of τ the NIDS achieves the best accuracy?

a. $\tau = 0.3$

b. $\tau = 0.1$

c. $\tau = 0$

d. $\tau = 0.5$

21
1 pt.

Why does a one-class SVM aim to *minimize* the volume of the hypersphere R^2 ?

a. To increase the number of false positives.

b. To reduce the number of true positives

c. To increase the number of true negatives.

d. To reduce the number of false negatives.

Naam:				Handtekening:
Datum:	- -	Geboortedatum:	- -	

Cursus: B-BIT-MOD06-1B-202001032 B-CS Intelligent Interaction Design CS/BIT 202001032 - Vragen: 26-01-2021 - Intelligent Interaction Design CS/BIT - AI. & Cyber Security Theory

This exam is closed book, you may only use a simple calculator (addition, substraction, multiplication, division and exponentiation) and a 1A4 cheat sheet.

This examination consists of multiple-choice questions, for which you have 3 hours.

Tips:
• Read each question carefully without looking at the possible answers.
• 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).
• Beware of double negations (negatives) as these can be confusing.
• Do not stay on any one question for 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 left over at the end, check your answers.
• As part of these instructions, and as part of the relevant question(s) of this exam, you can find a table with values for $-p \log 2\,(p)$.

If you have questions about the content of the exam, you can contact the teachers via BigBlueButton. Make sure to address the teacher corresponding to the topic of the exam question that you want to ask something about, namely:

Birna van Riemsdijk: Q1-9

Gwenn Englebienne: Q10-15

Nacir Bouali: Q16-18

Andrea Continella: Q19-21

Please allow 10-15min for your question to be answered. If you have not yet heard from a teacher by then, please contact a TA or teacher in the room. Use the information button on the screen to see this info again while making the exam."

Good luck!

1 1 pt. A B C D

2 1 pt. A B C D

3 1 pt. A B C D

4 1 pt. A B C D

5 1 pt. A B C D

6 1 pt. A B C D

7 1 pt. A B C D

8 1 pt. A B C D

9 2 pt. **a.** A B C D

b. A B C D

10 1 pt. A B C D

11 1 pt. A B C D

12 1 pt. A B C D

13 1 pt. A B C D

14 1 pt. A B C D

15 1 pt. A B C D

16 1 pt. A B C D

17 1 pt. A B C D

18 3 pt. **a.** A B C D

b. A B C D

c. A B C D

19 1 pt. A B C D

20 1 pt. A B C D

21 1 pt. A B C D

Correctiemodel

1. A
1 pt.

2. B
1 pt.

3. B
1 pt.

4. C
1 pt.

5. A
1 pt.

6. A
1 pt.

7. A
1 pt.

8. C
1 pt.

9. a. 1 pt. D
2 pt.
b. 1 pt. B

10. C
1 pt.

11. A
1 pt.

12. A
1 pt.

13. C
1 pt.

14. D
1 pt.

15. B
1 pt.

16. C
1 pt.

17. A
1 pt.

18. a. 1 pt. A
3 pt.
b. 1 pt. B
c. 1 pt. C

19. C
1 pt.

20. A
1 pt.

21. D
1 pt.

Cesuur

Toegepaste raadscore: 6 pt

Behaalde punten	Cijfer
24	10
23	9,4
22	8,9
21	8,3
20	7,8
19	7,2
18	6,7
17	6,1
16	5,6
15	5,1
14	4,6
13	4,2
12	3,7
11	3,3
10	2,8
9	2,4
8	1,9
7	1,5
6	1,0
5	1,0
4	1,0
3	1,0
2	1,0
1	1,0
0	1,0

Vraag-identificatiecodes

Deze identifiers kunnen worden gebruikt om de precieze vraag in de vragenbanken te identificeren. Gebruik deze code in combinatie met de documentcode wanneer u feedback doorgeeft, zodat precies duidelijk is op welke vraag en -versie uw feedback van toepassing is.

Documentidentificatiecode: 4310-9998

Vraagnummer	Vraag-identificatiecode	Versie-identificatiecode
1	44910	e89d8b7e-3cfa-3233-0b7f-70e137ab032d
2	44525	8a3b4210-04ef-73d0-a8e7-366d6dd7ae4c
3	44535	6766d6de-093c-2b94-812f-12c137437292
4	47330	cb832b39-4823-6a49-28dc-65afa3803578
5	44545	d5970f5e-3008-1f7f-f576-83f9cf54f8c5
6	44850	9fb4d9e5-10ed-5edf-ffc6-1ff8122ca612
7	47355	fea2d054-a394-2979-cf37-1d8f0acadb1d
8	44865	690edbe1-dc43-ae01-1a4d-27a5cd0ada86
9	47360	9102f5f3-e7ff-eab2-6b08-c9b41a3c5a02
10	44835	f4226c90-4e5f-8aa1-961d-776fa7678bb9
11	44930	89fcfe39-4c94-ed6c-4ec0-3ded99d415fd
12	20626	69a64866-139c-30cd-da58-6f8e9ff67502
13	47280	8ae00da4-625b-731a-4a25-41b2a256d7b9
14	47285	1e665aef-7eee-3ec9-eae9-e1d6f432f6ae
15	49395	21fa6028-709b-f290-d52d-57cdc2eaf258
16	47485	b5417d29-6698-1d79-01a7-db4bf680efff
17	47490	2dbde2bd-f616-eda5-ac9e-2dc3697b1d44
18	23326	2ff5cf2b-27a0-4c23-3078-0a5481f10313
19	44250	7763f6eb-27c6-e077-30cd-3def70f1ddcb
20	44285	7d5b8155-32f1-904d-1c26-a587fdd0aa65
21	44300	bb8ff610-1684-48e8-0efe-4af22ba474ff