Information about the exam

1. The exam has 6 sections of questions. They consist of both multiple choice questions and free text questions.
2. The total number of points for the exam is 48 and a passing grade is at most 25. Grade 4 is at most 34 points and Grade 5 is at most 41 points. The limits will be modify based on the results, but they will not be higher.
3. For multiple choice questions you will get a positive score for each correct choice and a negative score for each incorrect choice. They are both the same, usually 0.5 or 1. This means that it may be an idea to only check those choices you are certain about. You cannot get a negative score.
4. The free text questions are designed so that text should be enough to answer them. There are limited formatting facilities that can be used to format and structure your answers. For instance, there is a table facility and special character facility. These features are accessible via icons at the top of the answer textbox.
5. A scientific calculator is accessible for the Bayesian Network question and in the Resource section.
6. Both the course textbook and the lecture slides are accessible from the Resource section at the end of the exam. Simply click on a link and a new tab will be generated in the browser where you can read the selected material.
7. Please make reasonable assumptions if you believe an exercise is under specified and state those assumptions explicitly in your answer.
8. Your answers should be clear, concise and compact.
9. Fredrik will visit the exam rooms around 9.45 to provide clarifications if needed for the exam questions.

0 / 10000 Word Limit

1. Logic

1a) List all models of $\varphi = ((A \lor \neg B) \land (A \rightarrow (C \land \neg B)))$ over the propositions $\Sigma = \{A,B,C\}$. Does $\neg B$ logically follow from φ? Justify your answer.
Note that a truth table is not a model. (3p)

0 / 10000 Word Limit

1b) Use the equivalences introduced in the lecture to bring the formula $\psi = ((A \lor \neg B) \rightarrow C)$ into conjunctive normal form (CNF). Use only one equivalence per step. (2p)

0 / 10000 Word Limit

1c) Use DPLL to find a satisfying assignment for $\Delta = \{\{X, Y\}, \{\neg Y, Z\}, \{Y, Z\}, \{\neg X, \neg Y, \neg Z\}, \{\neg X, Y, \neg Z\}\}$.

Try variables in the order X, Y,Z and truth values in the order T,F.

Specify the list of clauses after every modification.

Clearly specify the recursive call structure of the algorithm, for example by using a nested list or labeling steps as “1a, 2a, 2b, 1b . . . ”. (5p)

0 / 10000 Word Limit

This space is available for comments and/or assumptions that you wish to state.

0 / 10000 Word Limit

2. Bayesian Networks

Consider the following problem 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. Let AR stand for "arthritis", AH for "aching hands", AE for "aching elbow", TE for "tennis elbow", and DH for "dishpan hands".

$P(ah|ar, dh) = P(ae|ar, te) = 0.2$
$P(ah|ar, \neg dh) = P(ae|ar, \neg te) = 0.8$
$P(ah|\neg ar, dh) = P(ae|\neg ar, te) = 0.8$
$P(ah|\neg ar, \neg dh) = P(ae|\neg ar, \neg te) = 0.1$
$P(te|ar) = 0.1$
$P(te|\neg ar) = 0.1$
$P(ar) = 0.02$
$P(dh) = 0.02$

2a) Which of the following Bayesian networks represent the causal links described in the problem example defined above? (1p)

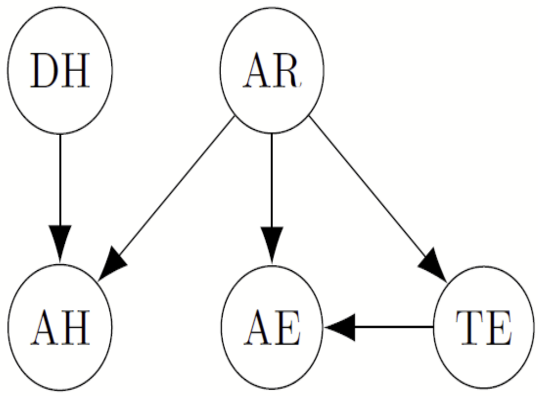

(a) Network 1

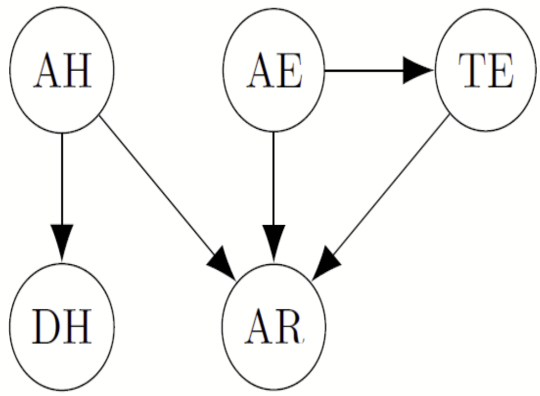

(b) Network 2

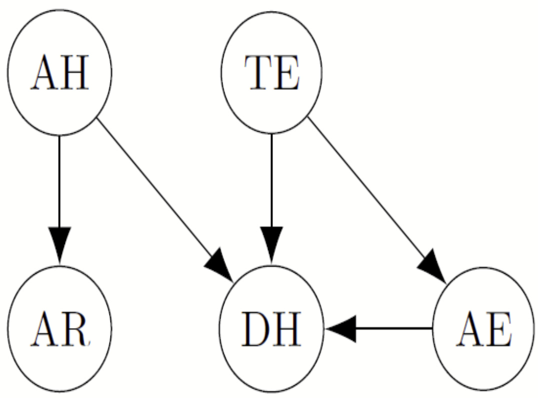

(c) Network 3

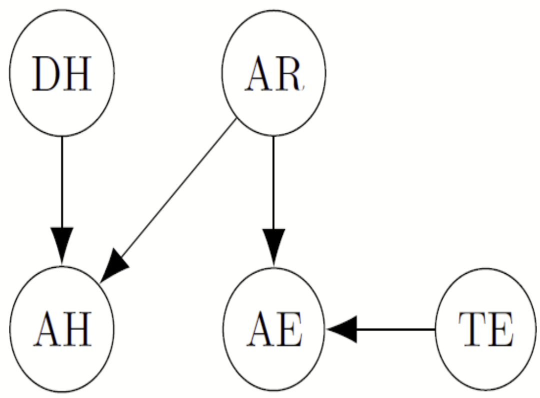

(d) Network 4

- ☐ Network 1 ✔
- ☐ Network 2
- ☐ Network 3
- ☐ Network 4

2b) Given the independence assumptions implicit in the Bayesian network, which of the formulas below represent the full joint probability distribution over all five variables, i.e. P(AR,AH,AE,TE,DH) =? (1p)

- ☐ P(AR) · P(AH) · P(AE) · P(TE) · P(DH)
- ☐ P(AR|AH,AE, TE) · P(AH) · P(AE) · P(TE|AE) · P(DH|AH)
- ☐ P(AR) · P(AH|AR,DH) · P(AE|AR, TE) · P(TE) · P(DH)
- ☐ P(AR) · P(AH|AR,DH) · P(AE|AR, TE) · P(TE|AR) · P(DH) ✔

2c) Using the formula for the full joint probability distribution and the probabilities given in the table above select statements which are True: (4p)

- ☐ P(¬ar, ¬dh,¬te, ah,¬ae) = 0.98 · 0.98 · 0.1 · 0.8 · 0.8 = 0.061466
- ☐ P(¬ar, ¬dh,¬te, ah,¬ae) = 0.98 · 0.98 · 0.9 · 0.1 · 0.9 = 0.07779 ✔
- ☐ P(ar|ah, te) = α ·$\sum_{DH,\ AE}$P(ar, ah, te,DH,AE), where α is the normalization factor. ✔
- ☐ P(ar|ah, te) = α ·$\sum_{AR,\ DH,\ AE}$P(AR, ah, te,DH,AE), where α is the normalization factor.
- ☐ P(ar|ah, te) = 0.1236 ✔
- ☐ P(ar|ah, te) = 0.0697
- ☐ P(ar, dh, ah, te) =$\sum_{AE}$P(ar, ah, te, dh,AE) ✔
- ☐ P(ar, dh, ah, te) = α ·$\sum_{AE}$P(ar, ah, te, dh, AE), where α is the normalization factor.

This space is available for comments and/or assumptions that you wish to state.

0 / 10000 Word Limit

Calculator

3. CSP

The following questions pertain to Constraint Satisfaction Problems (CSPs). CSPs 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.

The figure below shows a constraint graph with eight variables. The value domains for each variable are the integer numbers 1 to 8. The constraints state that adjacent/connected nodes can not have consecutive numbers and they must be different. For example, if node C is labeled 2, then nodes A, D, and G cannot be labeled with either 1 or 3 (consecutive numbers) or 2 (the same number).

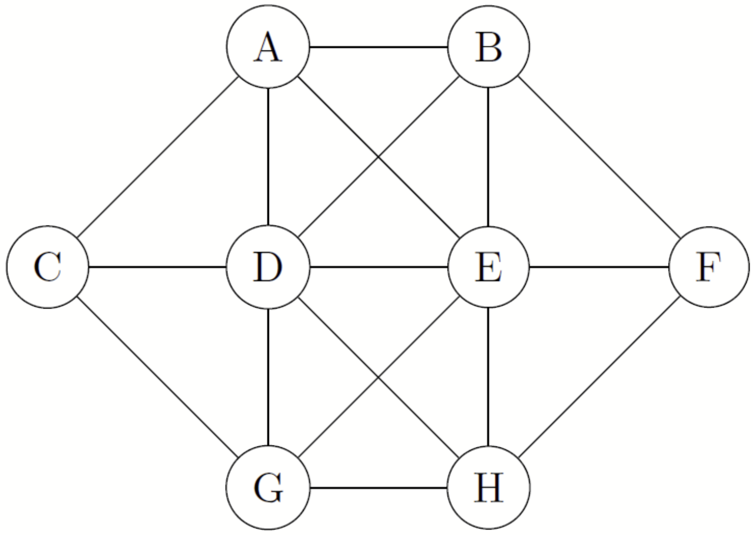

Constraint graph.

3a) Select statements which are True? (1p)

- ☐ Applying Degree Heuristic to a CSP selects a variable which is involved in the largest number of constraints on other unassigned variables. ✔
- ☐ Applying Degree Heuristic to a CSP selects a variable with the fewest possible bindings left.
- ☐ If we apply the Degree Heuristic to the constraint graph defined in the figure above, C and F nodes will be chosen as potential candidates for labeling.
- ☐ If we apply the Degree Heuristic to the constraint graph defined in the figure above, D and E nodes will be chosen as potential candidates for labeling. ✔

3b) Select statements which are True: (1p)

- [ ] Applying Least Constraining Value Heuristic to a CSP selects a value for chosen variable that rules out the fewest choices for the neighboring variables in the constraint graph. ✔
- [ ] Applying Least Constraining Value Heuristic to a CSP selects a value for chosen variable that yields the lowest number of consistent values in the neighboring variables in the constraint graph.
- [ ] Assuming a variable was chosen using the Degree Heuristic in the previous question, the Least Constraining Value Heuristic will select 2, 3, 4, 5, 6, and 7 as the potential candidate values.
- [ ] Assuming a variable was chosen using the Degree Heuristic in the previous question, the Least Constraining Value Heuristic will select 1 and 8 as the potential candidate values. ✔

3c) Suppose node D = 8, node E = 1 and nodes A,B,C,F,G,H are labeled 1,2,3,4,5,6,7,8.Which of the following graphs will be the result of applying the AC-3 algorithm which makes the constraint graph arc consistent. (2p)

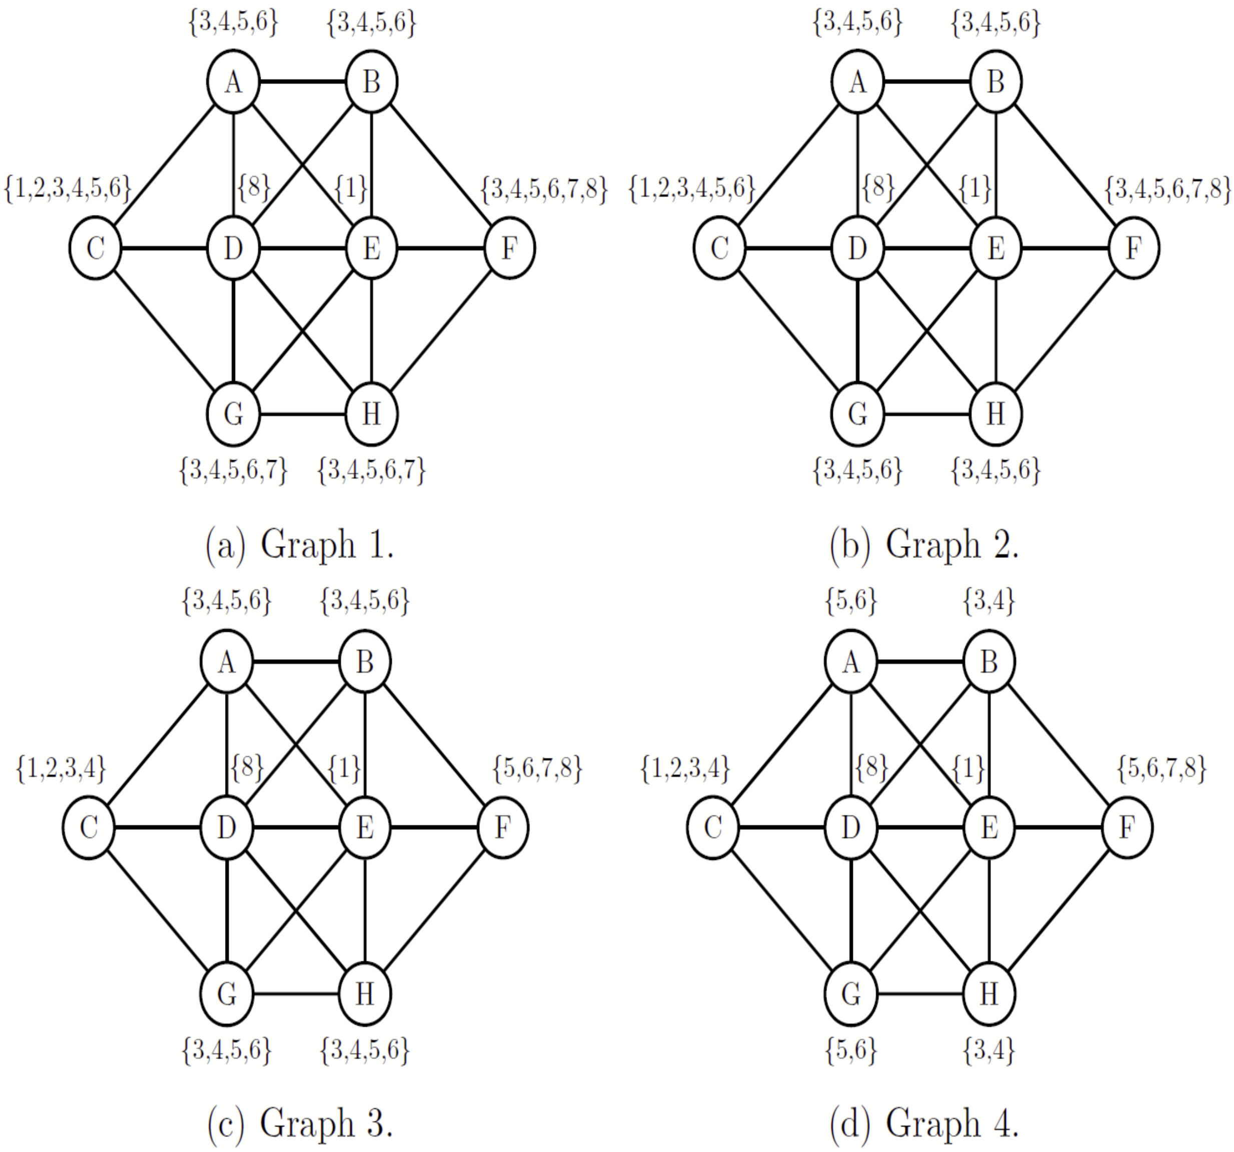

Potential constraint graphs after applying AC-3.

- ○ Graph 1
- ○ Graph 2 ✔
- ○ Graph 3
- ○ Graph 4
- ○ None of the graphs

3d) The figure below presents candidate solutions to the CSP defined. Which of the following are globally consistent solutions?: (2p)

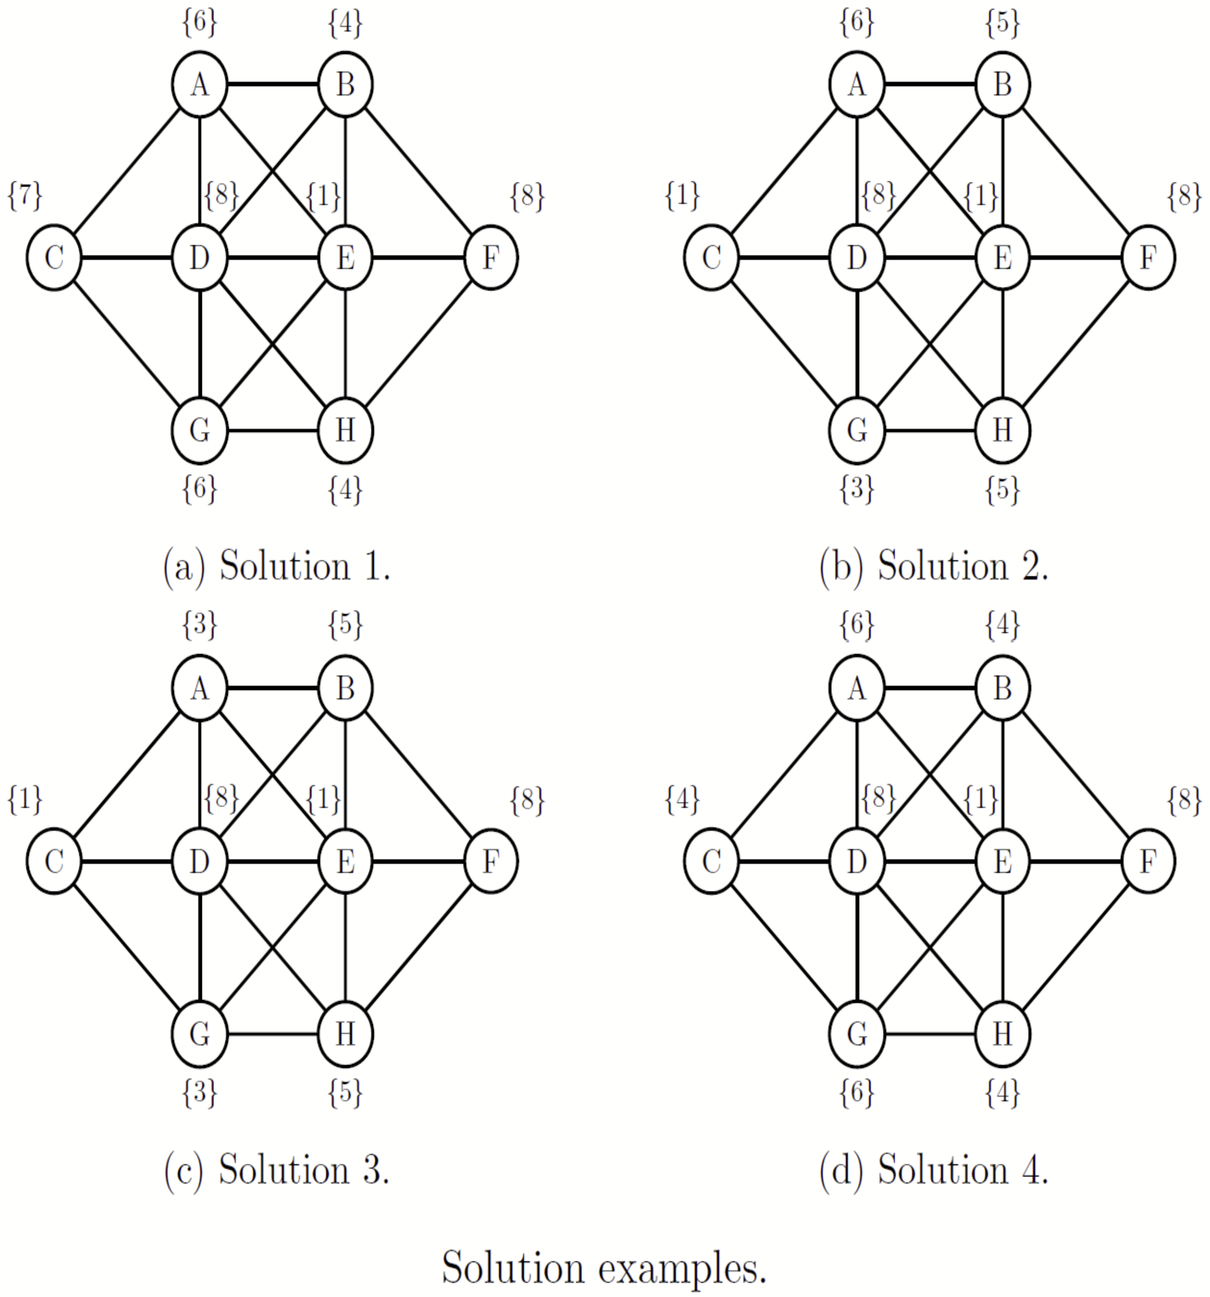

(a) Solution 1. (b) Solution 2. (c) Solution 3. (d) Solution 4.

Solution examples.

- [ ] Solution 1
- [ ] Solution 2
- [ ] Solution 3 ✔
- [ ] Solution 4 ✔

This space is available for comments and/or assumptions that you wish to state.

0 / 10000 Word Limit

4. Search

Consider the game tree in the figure below in which the leaf nodes show heuristic values and where all heuristic values are from the MAX players point of view. Assume search is in the left to right direction.

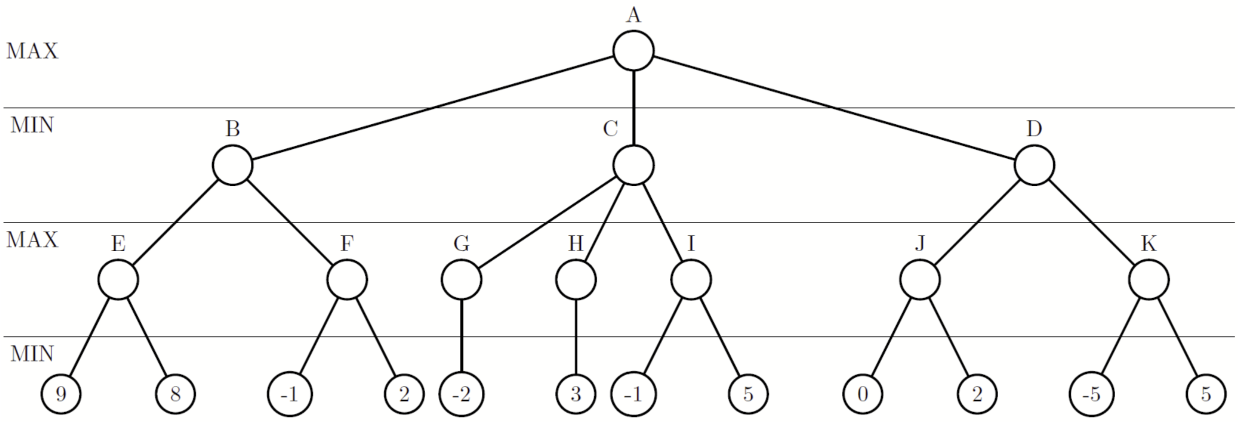

Figure 5: Two player game search

4a) Apply the MinMax algorithm to the game tree in the figure and state what move the first player (maximiser) would make. Provide heuristic values for each node in form of a table or text (e.g. A: value, B: value etc.). (2p)

0 / 10000 Word Limit

4b) In the game tree above, what nodes would not need to be examined using the alphabeta pruning procedure? Justify your answer in terms of the relevant α/β values in the nodes of the tree and why certain branches would be cutoff based on this evaluation.

Use annotation similar to the question above. To describe edges use one of the following notations: ”second edge below X” or “edge between X and Y”). (2p)

0 / 10000 Word Limit

4c) A* search is the most widely-known form of best-first search. 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. (2p)

0 / 10000 Word Limit

4d) Which of the following are true and which are false? Explain your answers.
I. Depth-first search always expands at least as many nodes as A* search with an admissible heuristic. (1p)
II. Heuristic function h(n) = 0 is an admissible heuristic for the 8-puzzle game. (1p)
III. Breadth-first search is complete even if zero step costs are allowed. (1p)

0 / 10000 Word Limit

5. Planning

Questions 5a - 5b gives 1 point for each correct answer. In total you can get 4 points for these.

5a) Which of the following statements about planning formalisms are true? (2p)

- [ ] The state-space induced by a planning task Π can be exponentially larger than the encoding of Π. ✔
- [ ] The main difference between STRIPS tasks and SAS+ tasks is the size of the variable domains. ✔
- [ ] In SAS+ conditions are represented by complete states.
- [ ] Some tasks require multi-valued variables and can thus only be encoded in SAS+.

5b) Which of the following statements about planning heuristics are true? (2p)

- [ ] For optimal planning, we want heuristics to be admissible. ✔
- [ ] Every relaxed plan is a plan.
- [ ] A PDB heuristic with a pattern containing n binary variables can be precomputed and stored in space linear in n.
- [ ] It is possible to create a PDB heuristic that is equal to the perfect heuristic h*. ✔

5c) Complete the definition of the following STRIPS planning task so that its optimal plan has length 3 and cost 5. (3p)
Π = ⟨V, I, G,A⟩ with V = {x, y}, I = ∅, G = {x, y}, A = {a1, a2}, and

a	$pre(a)$	$add(a)$	$del(a)$	$cost(a)$
a_1				
a_2				

0 / 10000 Word Limit

5d) Consider the abstraction of the following state space that is induced by the projection π{X} to the variable X.
Give the heuristic value of the pattern database heuristic h{X} for each state in the original state space.
A state label of x, y corresponds to state {X 7→ x, Y 7→ y} and all actions have a cost of 1. (3p)

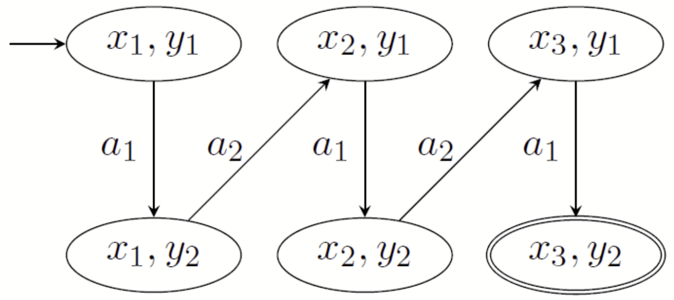

0 / 10000 Word Limit

6. Machine Learning

6a) Explain the main components of a neural network and how it is trained in a supervised manner. (3p)

0 / 10000 Word Limit

6b) What is the Q-function in reinforcement learning and how is it related to selecting the optimal action? (2p)

0 / 10000 Word Limit

6c) What is the XOR-problem in relation to the perceptron? (1p)

0 / 10000 Word Limit

6d) What is regularization in supervised learning? (1p)

0 / 10000 Word Limit

Resources

The chapters of the course book are available here:

- Preface: Preface.pdf
- Chapter 1 Introduction: Chapter 1.pdf
- Chapter 2 Intelligent Agents: Chapter 2.pdf
- Chapter 3 Solving Problems by Searching: Chapter 3.pdf
- Chapter 4 Search in Complex Environments: Chapter 4.pdf
- Chapter 5 Constraint Satisfaction Problems: Chapter 5.pdf
- Chapter 6 Adversarial Search and Games: Chapter 6.pdf
- Chapter 7 Logical Agents: Chapter 7.pdf
- Chapter 8 First-Order Logic: Chapter 8.pdf
- Chapter 9 Inference in First-Order Logic: Chapter 9.pdf
- Chapter 10 Knowledge Representation: Chapter 10.pdf
- Chapter 11 Automated Planning: Chapter 11.pdf
- Chapter 12 Quantifying Uncertainty: Chapter 12.pdf
- Chapter 13 Probabilistic Reasoning: Chapter 13.pdf
- Chapter 14 Probabilistic Reasoning over Time: Chapter 14.pdf
- Chpater 15 Making Simple Decisions: Chapter 15.pdf
- Chapter 16 Making Complex Decisions: Chapter 16.pdf
- Chpater 17 Multiagent Decision Making: Chapter 17.pdf
- Chapter 18 Probabilistic Programming: Chapter 18.pdf
- Chapter 19 Learning from Examples: Chapter 19.pdf
- Chapter 20 Knowledge in Learning: Chapter 20.pdf
- Chapter 21 Learning Probabilistic Models: Chapter 21.pdf
- Chapter 22 Deep Learning: Chapter 22.pdf
- Chapter 23 Reinforcement Learning (from 3rd ed): Chapter 23 3rd ed.pdf
- Appendix A: Appendix A.pdf
- Bibliography: Bibliography.pdf

0 / 10000 Word Limit

Lecture notes

- LE1 Course Introduction, History of AI: 2023-08-29-LE1-Introduction to AI.pdf
- LE2 Search I: 2023-08-31-LE2-Search I.pdf
- LE3 Search II: 2023-09-04-LE3-Search II.pdf
- LE4 Constraint Satisfaction: 2023-09-05-LE4-CSP.pdf
- LE5 Machine Learning I: 2023-09-08-LE5 ML I.pdf
- LE6 Machine Learning II: 2023-09-11-LE6 ML II.pdf
- LE7 Machine Learning III: 2023-09-12-LE7 ML III.pdf
- LE8 Knowledge Representation I: TDDC17_Le8_logic1.pdf
- LE9 Knowledge Representation II: TDDC17_Le9_logic2.pdf
- LE10 Knowledge Representation III: TDDC17_Le10_logic3.pdf
- LE11 Bayesian Networks: 2023-09-22-LE11-Bayesian Networks.pdf
- LE12 Planning I: TDDC17_Le12_planning1.pdf TDDC17_Le12_planning2.pdf
- LE13 Planning II: TDDC17_Le13_planning3.pdf TDDC17_Le13_planning4.pdf
- LE14 Planning III: TDDC17_Le14_planning5.pdf TDDC17_Le14_planning6.pdf
- LE15 Robotics/Perception I: TDDC17_Le15_robotics1.pdf
- LE16 Robotics/Perception II: TDDC17_Le16_robotics2.pdf
- LE17 Course Summary: 2023-10-06 LE17 Exam questions.pdf

0 / 10000 Word Limit

Calculator