Sektion 1

Information

1. The exam has 6 sections of questions. They consist of both multiple choice questions and free text questions.
2. 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 points. This means that it may be an idea to only check those choices you are certain about. You cannot get a negative score.
3. 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.
4. 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.
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.
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. A teacher will visit the exam rooms around 15.30 to provide clarifications and answer any questions.

1. Logics

1a) Which of the following statements about logical formulas are correct? (1p)

- [ ] Every valid formula is satisfiable.
- [ ] Every unsatisfiable formula is valid.
- [ ] There exist formulas that are both satisfiable and unsatisfiable.
- [ ] There exist formulas that logically entail each other.

1b) What properties does the logical formula $(A \to (A \to B))$ have? (1p)

- [ ] satisfiable
- [ ] falsifiable
- [ ] valid
- [ ] unsatisfiable

Let $\varphi = (A \lor \neg B) \land (A \lor B) \land (A \lor \neg C) \land (A \lor D) \land (C \lor \neg D)$ be a propositional logic formula over the propositions $\{A,\ B,\ C,\ D\}$.

1c) Write φ in set notation for CNF formulas. (1p)

B *I* U á

0 / 10000 Word Limit

1d) Show the full execution and result of the DPLL algorithm on φ. For every node in the DPLL tree, show the current formula and which rule is applied. For the splitting rule, consider the variables in alphabetical order and consider the truth value **F** before **T**. (6p)

B *I* U á

0 / 10000 Word Limit

1e) State three models for φ. (3p)

0 / 10000 Word Limit

2. Bayesian Networks

Consider the following problem statement:

> The fire alarm in a building can go off if there is a fire in the building or if the alarm is tampered with by vandals. If the fire alarm goes off, this can cause crowds to gather at the front of the building and fire trucks to arrive.

Let A stand for ”alarm sounds”, C for ”crowd gathers”, F for ”fire exists”, T for ”truck arrives”, and V for ”vandalism exists”.

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

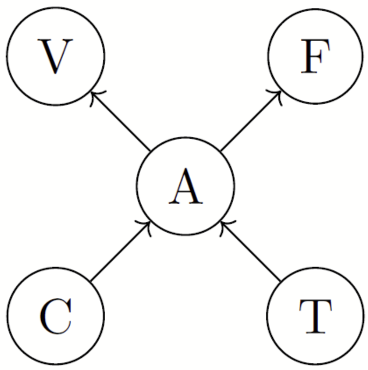

(a) Network 1

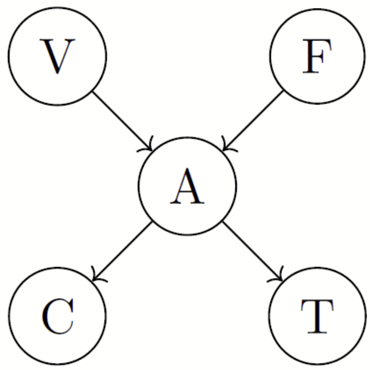

(b) Network 2

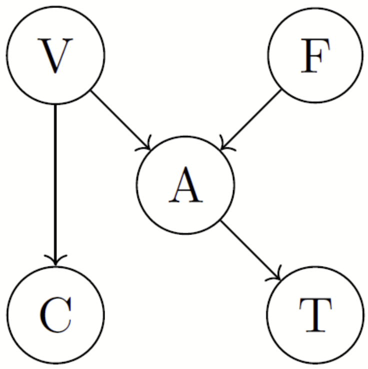

(c) Network 3

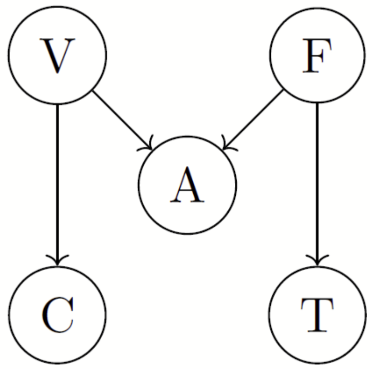

(d) Network 4

- ☐ 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(A, C, F, T, V) =? (1p)

- [ ] P(A|F, V) · P(C|A) · P(F) · P(T |A) · P(V)
- [ ] P(A | C, T) · P (C) · P(F | A) · P(T) · P(V | A)
- [ ] P(A) · P(C) · P(F) · P(T) · P(V)
- [ ] P(A | F, V) · P(C | V) · P(F) · P(T | F) · P(V)

2c) Using the formula for the full joint probability distribution and the probabilities given below, select which statements are true. (3p)

P(a | f, v) = 0.9
P(a | ¬f, v) = 0.5
P(a | f, ¬v) = 0.9
P(a | ¬f, ¬v) = 0.01
P(c | a) = 0.7
P(c | ¬a) = 0.1
P(f) = 0.1
P(t | a) = 0.9
P(t | ¬a) = 0.1
P(v) = 0.2

- [ ] P(a, c, ¬f, t, v) = 0.5 · 0.7 · 0.1 · 0.9 · 0.2 = 0.0063
- [ ] P(a, c, ¬f, t, v) = 0.5 · 0.7 · 0.9 · 0.9 · 0.2 = 0.0567
- [ ] P(f | c, t) = α · $\Sigma_{F,A,V}$P(A, c, F, t, V), where α is the normalization factor
- [ ] P(f | c, t) = α · $\Sigma_{A,V}$P(A, c, f, t, V), where α is the normalization factor
- [ ] P(f | c, t) ≈ 0.45
- [ ] P(f | c, t) = 0.1
- [ ] P(a, c, f, t) = Σ_{V}P(a, c, f, t, V)
- [ ] P(a, c, f, t) = α · Σ_{V}P(a, c, f, t, V), where α = $\frac{1}{\sum_A \sum_V P(A,c,f,t,V)}$

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

B I U á

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 domain for each variable is 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, B, D, G, and H cannot be labeled with either 1 or 3 (consecutive numbers) or 2 (the same number).

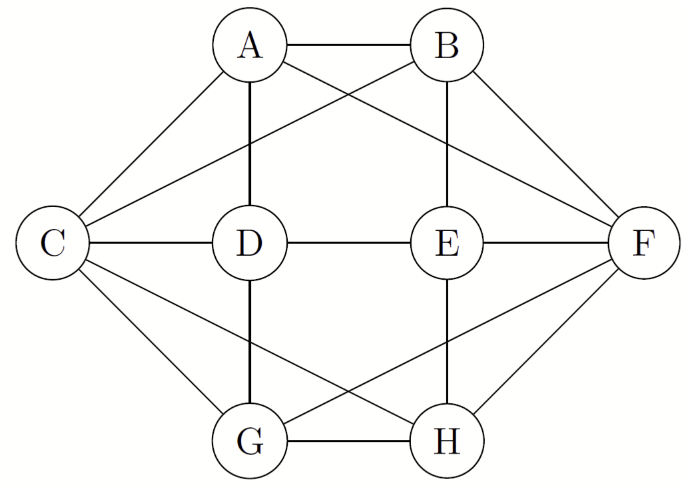

3a) Select which statements are True. (1p)

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

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

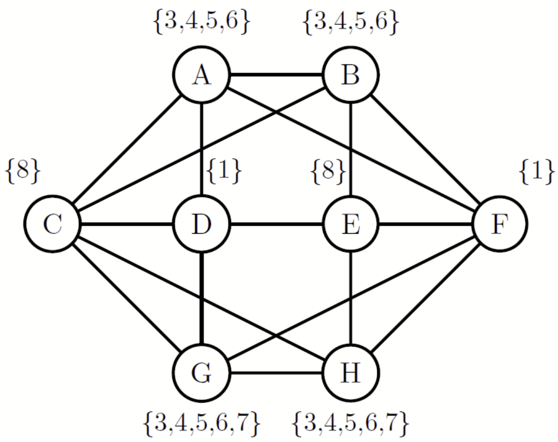

(a) Graph 1.

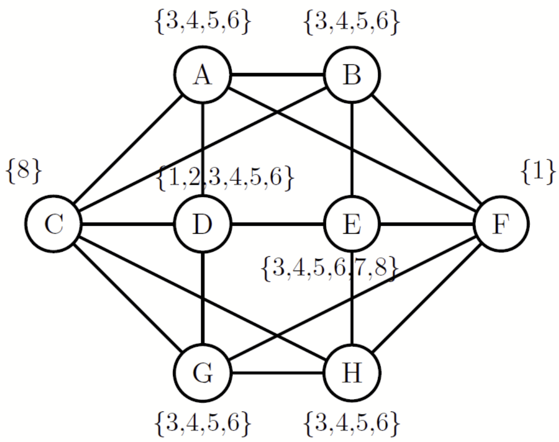

(b) Graph 2.

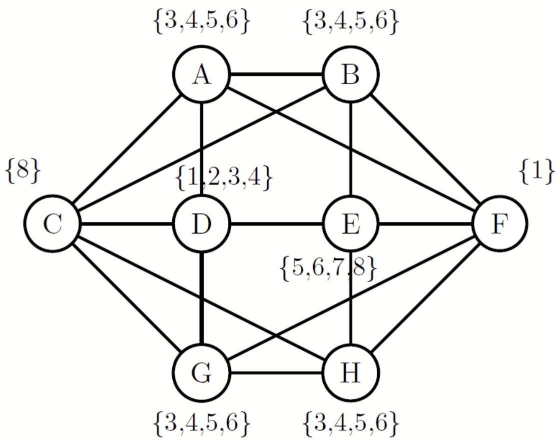

(c) Graph 3.

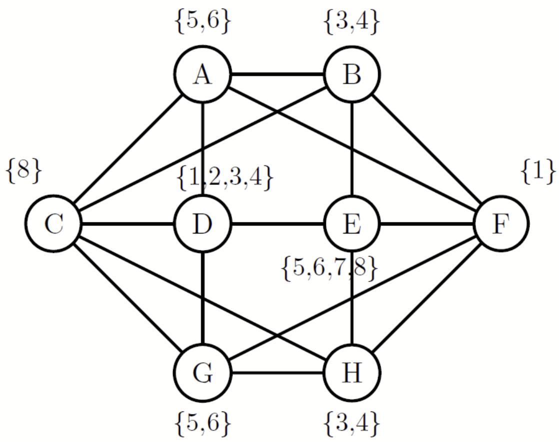

(d) Graph 4.

- [ ] Graph 1
- [ ] Graph 2
- [ ] Graph 3
- [ ] Graph 4

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

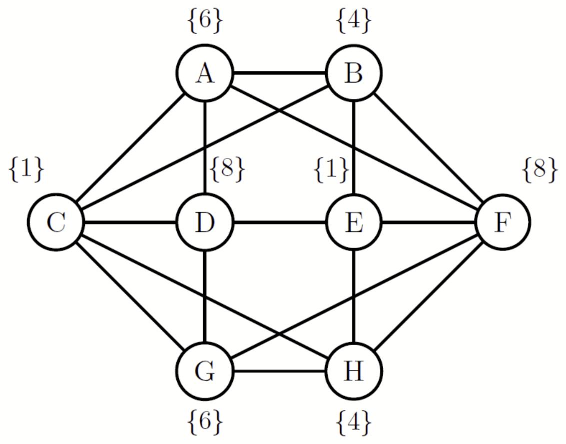

(a) Solution 1.

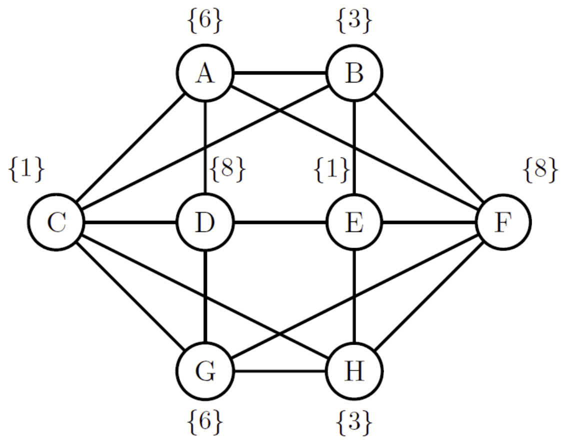

(b) Solution 2.

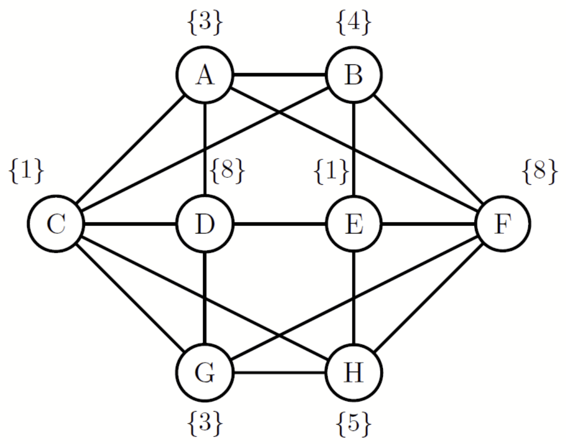

(c) Solution 3.

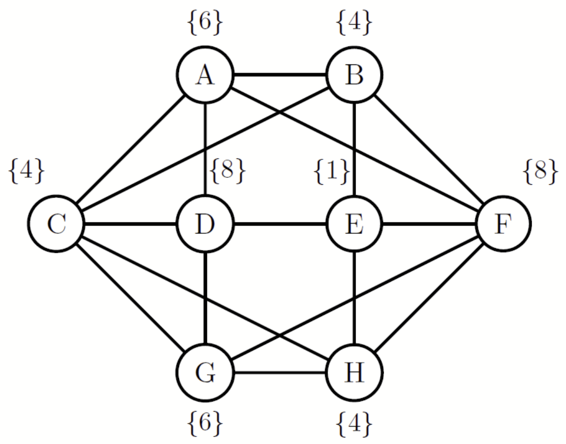

(d) Solution 4.

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

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

B *I* U ⁝☰ ☰ á

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.

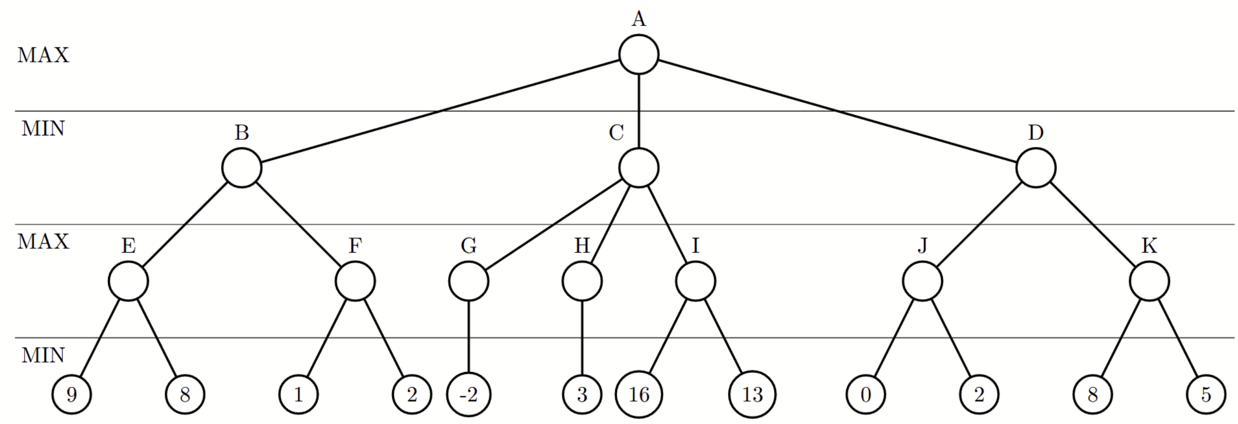

4a) Apply the MinMax algorithm to the game tree in the figure and state what move the first player (maximiser) would make. Provide heuristic values for each node in form of a table or text (e.g. A: value, B: value etc.). (2p)

B I U á

0 / 10000 Word Limit

4b) In the game tree above, what nodes would not need to be examined using the alpha-beta 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)

B I U á

0 / 10000 Word Limit

Let's consider the 8-puzzle. It consists of a 3x3 game board with 8 tiles (and one empty space) in the 9 slots. Each tile is numbered from 1 to 8. The start state of the game places the 8 numbered tiles arbitrarily on the board. The goal state is for the tiles to be numbered in numerical order from top to bottom, left to right. There is one action in the game. A tile can be moved from slot A to slot B if slot B is empty.

4c) The Hamming distance measures the number of misplaced tiles in a board configuration. Is this an admissible heuristic? Explain why this is the case. (1p)

B *I* U ⁝≡ 1≡ á

0 / 10000 Word Limit

4d) Provide an additional admissible heuristic for this problem and explain why it is an admissible heuristic. (1p)

B *I* U ⁝≡ 1≡ á

0 / 10000 Word Limit

5. Planning

5a) Which of the following combinations of search algorithm and heuristic for a planning task are guaranteed to lead to optimal solutions? (1p)

- [ ] Greedy best-first search with the FF heuristic.
- [ ] A* with the $h^{\max}$ heuristic.
- [ ] A* with reopening with the h^{add} heuristic.
- [ ] Weighted A* with a weight of 2 and the $h^{\max}$ heuristic.

5b) Consider a STRIPS planning task Π with 10 state variables. Which of the following statements must be true? (1p)

- [ ] Π has 2^{10} states.
- [ ] Π has 2^{10}reachable states.
- [ ] The delete relaxation of Π has 2^{10} states.
- [ ] Π has exactly one initial state.

Consider the delete-free planning task Π^+ in simplified notation with initial state {i}, goal {g} and the following actions:

$a_1 = \{i\} \xrightarrow{3} \{a\}$

$a_2 = \{i\} \xrightarrow{3} \{b\}$

$a_3 = \{i\} \xrightarrow{4} \{a, b\}$

$a_4 = \{a\} \xrightarrow{1} \{c\}$

$a_5 = \{b\} \xrightarrow{1} \{d\}$

$a_6 = \{c, d\} \xrightarrow{1} \{g\}$

Remember that in simplified notation, we write actions a with $\text{pre}(a) = \{p_1, \ldots, p_k\}$, $add(a) = \{a_1, \ldots, a_l\}$, $del(a) = \varnothing$ and $\cos t(a) = c$ as $\{p_1, \ldots, p_k\} \xrightarrow{c} \{a_1, \ldots, a_l\}$.

For (e) and (f), the state variables are $\{i, a, b, c, d, g\}$. The $h^{\max}$/h^{add} value of a state variable is its value in the last variable layer of the relaxed planning graph. You do not need to draw a relaxed planning graph or show computation steps.

5c) Give an optimal plan for Π^+. What is its cost? (2p)

B I U

0 / 10000 Word Limit

5d) Give a suboptimal plan for Π^+. What is its cost? (2p)

B I U

0 / 10000 Word Limit

5e) What is the $h^{\max}$ heuristic value for the initial state? State the overall heuristic value and the $h^{\max}$ values of all state variables. (3p)

B I U

0 / 10000 Word Limit

5f) What is the h^{add} heuristic value for the initial state? State the overall heuristic value and the h^{add} values of all state variables. (3p)

B I U á

6. Machine Learning

6a) What is the role of optimization in ML? Give at least one example. (1p)

B I U á

0 / 10000 Word Limit

6b) What are the two main differences between supervised and unsupervised learning? (2p)

B I U á

0 / 10000 Word Limit

6c) What is model-based reinforcement learning? Give at least one example of such a method. (1p)

B I U á

0 / 10000 Word Limit

6d) Explain the concepts of model capacity and error and their relationsship. (2p)

B I U á

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

B I U á

0 / 10000 Word Limit

Lecture notes

- LE1 Course Introduction, History of AI: 2022-08-30-LE1-Introduction to AI-part1.pdf 2022-08-30-LE1-Introduction to AI-part2.pdf
- LE2 Search I: 2022-09-01-LE2-Search I.pdf
- LE3 Search II: 2022-09-06-LE3-Search II_compressed.pdf
- LE4 Constraint Satisfaction: 2022-09-08-LE4-CSP.pdf
- LE5 Knowledge Representation I: 2022-09-09-LE5-KR I.pdf
- LE6 Knowledge Representation II: 2022-09-12-LE6-KR II.pdf
- LE7 Knowledge Representation III: 2022-09-13-LE7-KR III.pdf
- LE8 Bayesian Networks: 2022-09-16-LE8-Bayesian Networks.pdf
- LE9 Machine Learning I: 2022-09-19-LE9 ML I.pdf
- LE10 Machine Learning II: 2022-09-20-LE10 ML II.pdf
- LE11 Machine Learning III: 2022-09-23-LE11 ML III_compressed.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: 2022-10-07 LE17 Exam questions.pdf

B *I* U á

0 / 10000 Word Limit

Calculator