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. Mariuz will visit the exam rooms around 15.30 to provide clarifications if needed for the exam questions.

0 / 10000 Word Limit

1. Logic

Each question 1a) - 1e) gives 1 point. There is at least one correct choice for each question. It may be more than one.

1a) The propositional logic formula $(X \lor (\neg X \rightarrow Y))$ is... (1p)

- [ ] satisfiable ✔
- [ ] falsifiable ✔
- [ ] tautological
- [ ] unsatisfiable

1b) The propositional logic formulas φ and ψ are logically equivalent. Which of the following statements follow? (1p)

- [ ] φ |= ψ ✔
- [ ] φ is satisfiable
- [ ] φ is unsatisfiable
- [ ] (φ ∨ ¬ψ) is a tautology ✔

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

- [ ] We can test I |= φ in polynomial time ✔
- [ ] There exist formulas that are valid and unsatisfiable
- [ ] If φ ≡ ψ, then φ is satisfiable iff ψ is satisfiable ✔
- [ ] If φ and ψ have the same number of models, then φ ≡ ψ

1d) Which of the following statements about propositional logic formulas are true? (1p)

- [ ] If φ and ψ are satisfiable, then (φ ∧ ψ) is also satisfiable
- [ ] If φ and ψ are satisfiable, then (φ → ψ) is also satisfiable ✔
- [ ] If φ is satisfiable, then ¬φ is also satisfiable
- [ ] If φ is satisfiable, then ¬φ is unsatisfiable

1e) Which of the following statements about propositional logic are true? (1p)

- [ ] (A → (B ∨ C)) logically entails (¬A ∨ B)
- [ ] We can test Φ |= ψ by checking if $\neg(\bigwedge_{\varphi \in \Phi} \varphi \rightarrow \psi)$ is unsatisfiable ✔
- [ ] The resolvent of parent clauses {A,B} and {¬A,¬B} is the empty clause
- [ ] SAT instances far away from the phase transition are generally easy to solve ✔

1f) Transform the logical formula φ = ((A ∧ B) ↔ ¬C) ∧ (¬B → ¬C) into conjunctive normal form by applying logical equivalences. Which of the formulas below is the result? (Note that clauses and literals within clauses may be rearranged freely.) (3p)

- [ ] (A ∨ C) ∧ (A ∨ B ∨ ¬C) ∧ (B ∨ C) ∧ (¬B ∨ ¬C)
- [ ] (A ∨ C) ∧ (¬A ∨ ¬B ∨ ¬C) ∧ (B ∨ C) ∧ (B ∨ ¬C) ✔
- [ ] (A ∨ C) ∧ (¬A ∨ B ∨ ¬C) ∧ (B ∨ C) ∧ (¬B ∨ ¬C)
- [ ] (A ∨ C) ∧ (¬A ∨ ¬B ∨ ¬C) ∧ (B ∨ C) ∧ (¬B ∨ ¬C)

2. Bayesian Networks

Consider the following problem:

> *A company is trying to hire a recent college graduate. The company aims to hire an intelligent employee, but there is no way to test intelligence directly. The company does have information about a student's grade in a relevant course, the student's national SAT score, and the quality of a recommendation letter for the course.*
> *A student's **Grade** (High or Low) in a relevant course is dependent on the student's **Intelligence** (High or Low) and the **Difficulty** (Easy or Hard) of the course. A student's national **SAT score** (High or Low) is also dependent on the student's **Intelligence**. Additionally, the student's recommendation **Letter** (Good or Bad) for the course is dependent on the student's **Grade**.*

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

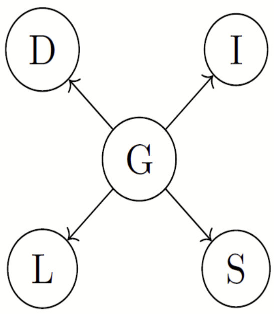

(a) Network 1

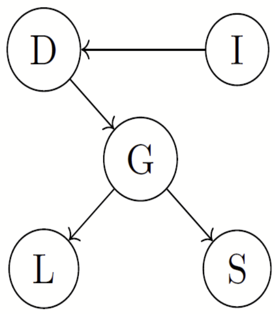

(b) Network 2

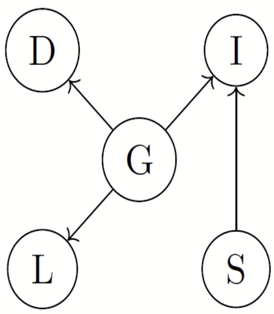

(c) Network 3

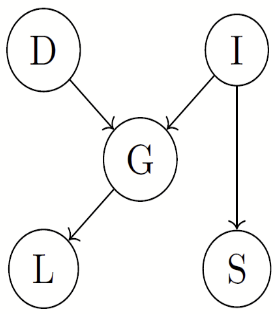

(d) Network 4

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

2b) Given the independence assumptions implicit in the Bayesian network, which of the formulas below represents the full joint probability distribution over all five variables, i.e. P (G, I, D, S, L) =? (1p)

- ☐ P(G) · P(I|G) · P(D|G) · P(S|G) · P(L|G)
- ☐ P(G|D) · P(I) · P(D|I) · P(S|G) · P(L|G)
- ☐ P(G) · P(I|G, S) · P(D|G) · P(S) · P(L|G)
- ☐ P(G|I,D) · P(I) · P(D) · P(S|I) · P(L|G) ✔

Probabilities for the Bayesian network:
P(d = Easy) = 0.4 P(d = Hard) = 0.6
P(i = High) = 0.7 P(i = Low) = 0.3
P(s = High | i = High) = 0.95 P(s = Low | i = High) = 0.05
P(s = High | i = Low) = 0.2 P(s = Low | i = Low) = 0.8
P(l = Good | g = High) = 0.9 P(l = Bad | g = High) = 0.1
P(l = Good | g = Low) = 0.4 P(l = Bad | g = Low) = 0.6
P(g = High | i = High, d = Hard) = 0.7 P(g = Low | i = High, d = Hard) = 0.3
P(g = High | i = High, d = Easy) = 0.95 P(g = Low | i = High, d = Easy) = 0.05
P(g = High | i = Low, d = Hard) = 0.1 P(g = Low | i = Low, d = Hard) = 0.9
P(g = High | i = Low, d = Easy) = 0.5 P(g = Low | i = Low, d = Easy) = 0.5

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(g = High, i = Low, d = Hard, s = High, l = Good) = 0.1·0.3·0.6·0.2·0.9 = 0.00324 ✔
- ☐ P(g = High, i = Low, d = Hard, s = High, l = Good) = 0.7·0.3·0.6·0.1·0.9 = 0.01134
- ☐ P(i = High | g = High, s = High, l = Good) = α · $\sum_D$ P(g, i, D, s, l), where α is the normalization factor ✔
- ☐ P(i = High | g = High, s = High, l = Good) = α · $\sum_{I,\,D}$ P(g, I, D, s, l), where α is the normalization factor
- ☐ P(i = High | g = High, s = High, l = Good) $\approx$ 0.97 ✔
- ☐ P(i = High | g = High, s = High, l = Good) $\approx$ 0.81
- ☐ P(g, i, d, s) = $\sum_L$ P(g, i, d, s, L) ✔
- ☐ P(g, i, d, s) = α · $\sum_L$ P(g, i, d, s, L), where α = $\frac{1}{\sum_{G,L} P(G,i,d,s,L)}$

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 2 to 8. The constraints state that adjacent/connected nodes cannot have consecutive numbers, and they must be different. For example, if node C is labeled 3, then nodes A, D, and G cannot be labeled with either 2 or 4 (consecutive numbers) or 3 (the same number).

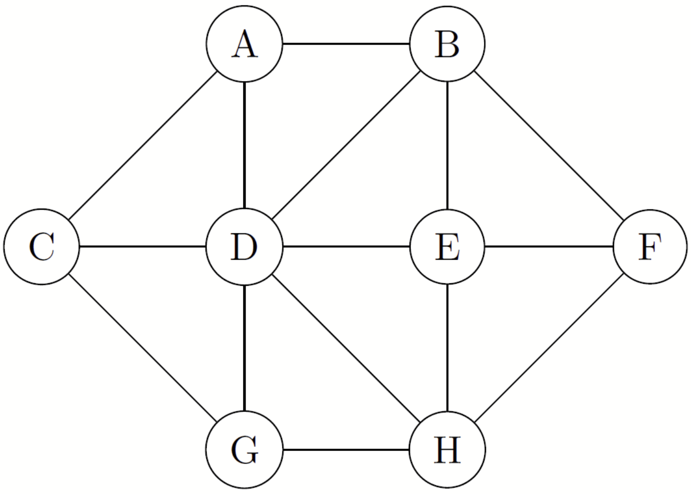

3a) Select the statements that are True: (1p)

- ☐ Applying the *most constraining variable order heuristic* to a CSP selects a variable with the fewest possible bindings left.
- ☐ Applying the *most constraining variable order heuristic* to a CSP selects a variable that is involved in the largest number of constraints on other unassigned variables. ✔
- ☐ If we apply the *most constraining variable order heuristic* to the constraint graph above, the C and F nodes will be chosen as potential candidates for labeling.
- ☐ If we apply the *most constraining variable order heuristic* to the constraint graph above, the D node will be chosen as potential candidates for labeling. ✔

3b) Select the statements that are True: (1p)

- ☐ Applying the *minimum conflicts heuristic* to a CSP selects a value for the chosen variable that yields the lowest number of consistent values in the neighboring variables in the constraint graph.
- ☐ Applying the *minimum conflicts heuristic* to a CSP selects a value for the chosen variable that rules out the fewest choices for the neighboring variables in the constraint graph. ✔
- ☐ Assuming a variable was chosen using the *most constraining variable order heuristic* in the previous question, the *minimum conflicts heuristic* will select 2 and 8 as the potential candidate values. ✔
- ☐ Assuming a variable was chosen using the *most constraining variable order heuristic* in the previous question, the *minimum conflicts heuristic* will select 3, 4, 5, 6, and 7 as the potential candidate values.

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

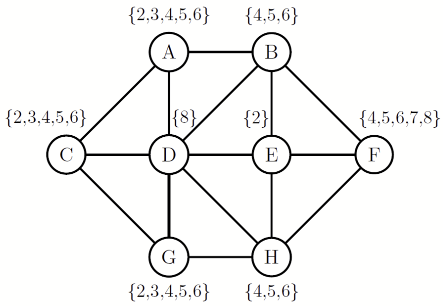

(a) Graph 1.

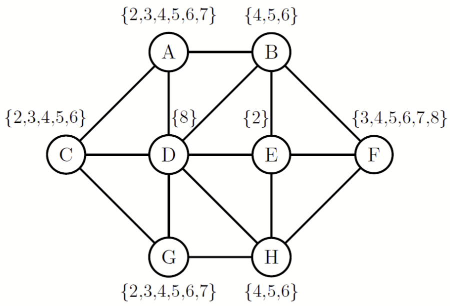

(b) Graph 2.

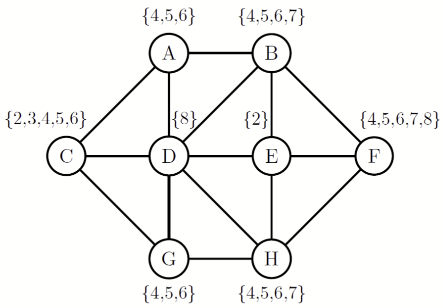

(c) Graph 3.

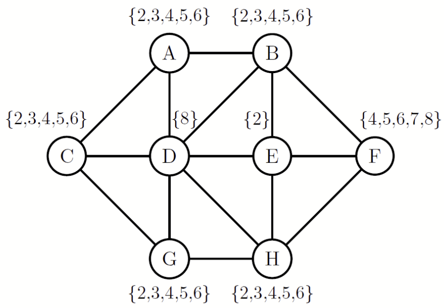

(d) Graph 4.

- ☐ 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)

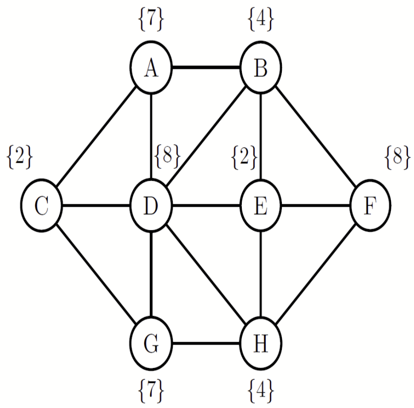

(a) Solution 1.

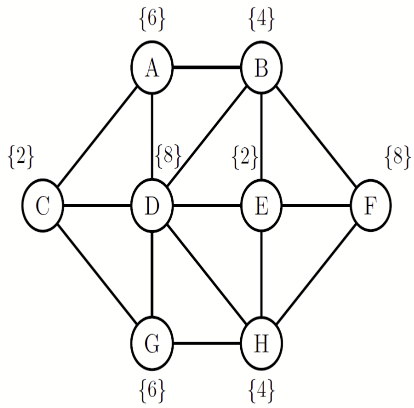

(b) Solution 2.

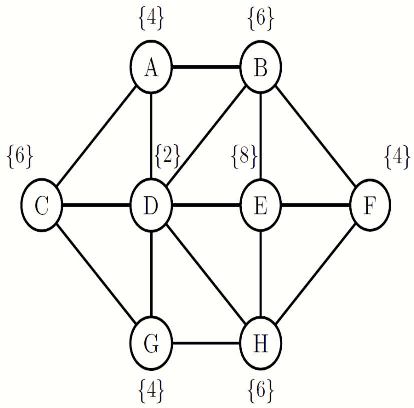

(c) Solution 3.

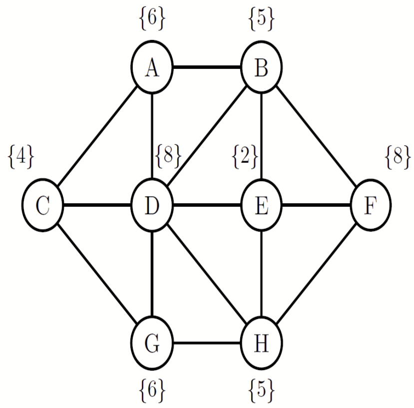

(d) Solution 4.

☐ Solution 1

☐ Solution 2 ✔

☐ Solution 3 ✔

☐ Solution 4

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

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.

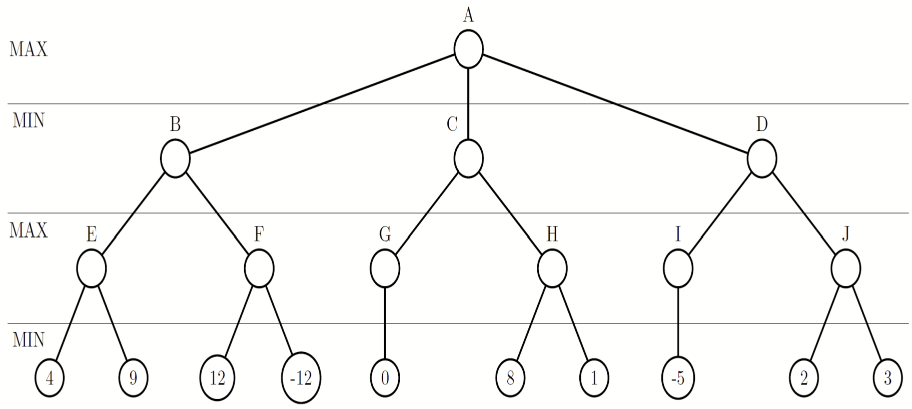

Figure 1: Two player game search

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) What is the difference between a world state, a state description, and a search node? Why is this distinction useful? (2p)

0 / 10000 Word Limit

A∗ search is the most widely-known form of best-first search.
4d) Explain what an admissible heuristic function is. (1p)

0 / 10000 Word Limit

4e) 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. Provide an admissible heuristic for this problem. Explain why it is an admissible heuristic and justify your answer explicitly. (2p)

0 / 10000 Word Limit

5. Planning

Questions 5a - 5e gives 0.5 points for each correct answer and -0.5 points for each incorrect answer. In total you can get 7 points for these.

5a) Which of the following statements about planning formalisms are true? (Each correct answer gives 0.5p, each incorrect answer gives -0.5

- [ ] 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 partial states. ✔
- [ ] Some tasks require multi-valued variables and can thus only be encoded in SAS+.

5b) Which of the following statements are true for all propositional planning tasks? (Each correct answer gives 0.5p, each incorrect answer gives -0.5p)

- [ ] The number of states grows exponentially in the number of state variables. ✔
- [ ] A solution to a planning task with n state variable has length at most n-1.
- [ ] We can determine in polynomial time if a solution of length at most 2 exists. ✔
- [ ] The number of operators is finite. ✔

5c) Which of the following statements about planning heuristics are true? (Each correct answer gives 0.5p, each incorrect answer gives -0.5p)

- [ ] 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 $h*$. ✔

5d) Which of these statements about delete relaxation for propositional planning tasks are correct? (Each correct answer gives 0.5p, each incorrect answer gives -0.5p)

- [ ] Delete-relaxation simplifies a planning task by reducing the number of states.
- [ ] If a planning task is solvable, then its delete relaxation is solvable. ✔
- [ ] The relaxed planning graph for a delete-relaxed planning task can be constructed in polynomial time. ✔
- [ ] The main advantage of hmax over h+ is that it is more efficiently computable. ✔

5e) Which of the following statements about probabilistic planning are true? (Each correct answer gives 0.5p, each incorrect answer gives -0.5p)

- [ ] Value iteration and policy iteration are both techniques that converge to an optimal policy. ✔
- [ ] Value iteration uses Bellman backups and policy iteration is based on Monte-Carlo backups.
- [ ] The Markov property states that the probability distribution for the next state only depends on the current state and the applied action. ✔
- [ ] The Markov property is true for MDPs and SSPs. ✔

5f) Consider the SSP $\mathcal{T} = \langle\{s0, s1, s2, s3\}, \{o1, \ldots, o4\}, c, T, s0, \{s3\}\rangle$ with cost function c(o1) = 2, c(o2) = 3, c(o3) = 1 and c(o4) = 6, and transition function T as follows:
T(s0, o1, s2) = 0.1
T(s0, o1, s3) = 0.9
T(s0, o2, s1) = 1
T(s1, o3, s0) = 0.5
T(s1, o3, s3) = 0.5
T(s2, o4, s2) = 0.5
T(s2, o4, s3) = 0.5
T(s, o, s′) = 0 for all other s, o, s′

Determine V_*(s2) using the Bellman equations. It may help you to draw a graphical representation of $\mathcal{T}$, but you do not get any points for it. Only enter the final numeric value, nothing else. (3p)

Correct answer:

12

6. Machine Learning

6a) Explain the concept of overfitting and a principled way to detect it. (2p)

0 / 10000 Word Limit

6b) Give two examples of the relation between the training data and the model learned from the data in supervised learning. (2p)

0 / 10000 Word Limit

6c) What is precision and recall and how are they related? (2p)

0 / 10000 Word Limit

6d) What is self-supervised learning? (1p)

0 / 10000 Word Limit

6e) What is the exploration exploitation trade-off in reinforcement learning? What is one way to deal with it? (2p)

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