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 15.30 to provide clarifications if needed for the exam questions.

0 / 10000 Word Limit

Sektion 2

1. Logic

Each multiple choice question 1a) - 1e) gives at most 2 points. Each correct choice gives 0.5 points. Each incorrect choice gives -0.5p. There is at least one correct choice for each question. It may be more than one.

1a) Consider the propositions P (“I go to the pool.”) and S (“It is summer.”). Which propositional logic formula captures the statement “I only go to the pool in summer.” best? (2p)

- ○ $P \rightarrow S$
- ○ $S \rightarrow P$
- ○ $S \leftrightarrow P$ ✔
- ○ $S \wedge P$

1b) The propositional logic formula $(X \rightarrow (\neg X \vee Y))$ is... (2p)

- ☐ satisfiable ✔
- ☐ falsifiable ✔
- ☐ tautological
- ☐ unsatisfiable

1c) Does the following hold? $(X \rightarrow Y) \rightarrow Z \equiv X \rightarrow (Y \rightarrow Z)$ (2p)

- ○ Yes, because both formulas have the same models.
- ○ No, because there is an interpretation that models the left formula but not the right formula (but not the other way around).
- ○ No, because there is an interpretation that models the right formula but not the left formula (but not the other way around). ✔
- ○ No, because there is an interpretation that models the left formula but not the right formula, and there is an interpretation that models the right formula but not the left formula.

1d) Let $\Delta = \{\{\neg A, \neg B, \neg C\}, \{A,B\}, \{\neg B,C\}, \{B,C\}\}$ and $\phi = (B \wedge C) \vee (A \wedge \neg B)$. Consider the statement ”$\Delta \models \phi$”. Which of the following hold? (1p)

- ○ The statement is correct because Δ and ϕ are logically equivalent.
- ○ The statement is incorrect because Δ and ϕ are not logically equivalent.
- ○ The statement is correct because Δ and the negation of ϕ is unsatisfiable. ✔
- ○ The statement is incorrect because Δ and the negation of ϕ is satisfiable.

1e) Which of the following statements about DPLL are true? (2p)

- [ ] DPLL can only handle formulas in conjunctive normal form. ✔
- [ ] DPLL can get stuck in local minima.
- [ ] DPLL can only be used for satisfiable formulas.
- [ ] DPLL can find a satisfying assignment for some formulas in polynomial time. ✔

1f) Local search methods for SAT... (2p)

- [ ] are usually complete
- [ ] can prove unsatisfiability for a formula
- [ ] usually benefit from randomization ✔
- [ ] can find models for extremely hard problems ✔

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

0 / 10000 Word Limit

2. Bayesian Networks

Consider the following random variables for a satellite monitoring problem: B (Battery Failure), S (Solar Panel Failure), E (Electrical System Failure), D (Trajectory Deviation), and C (Communication Loss). The domain of B, S, E, D, and C is {true, false}. The Bayesian network model for these variables has the following conditional table entries (the rest can be derived easily):

$P(s) = 0.1$
$P(b) = 0.1$
$P(e|b, s) = 0.9$
$P(e|b,\neg s) = 0.8$
$P(e|\neg b, s) = 0.8$
$P(e|\neg b,\neg s) = 0.1$
$P(d|e) = 0.9$
$P(d|\neg e) = 0.1$
$P(c|e) = 0.9$
$P(c|\neg e) = 0.15$

2a) Which of the following Bayesian networks represents the joint distribution P(B,S,E,D,C) constrained by the conditional table entries defined above? (2p)

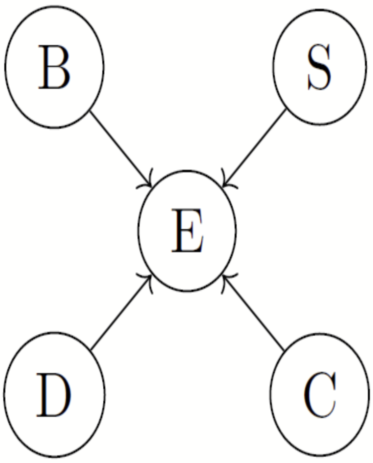

(a) Network 1

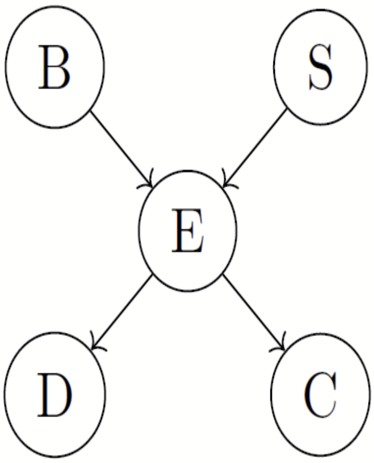

(b) Network 2

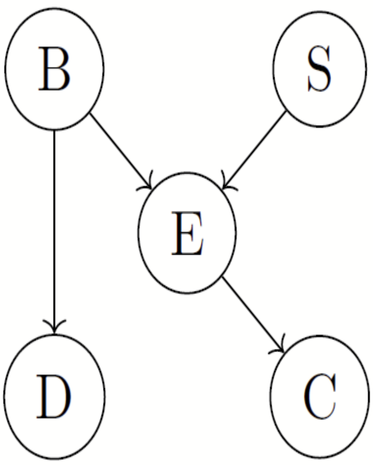

(c) Network 3

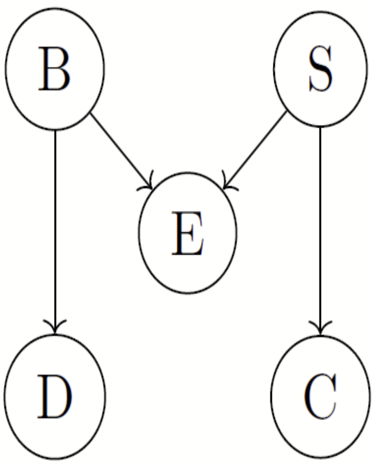

(d) Network 4

- [ ] 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(B, S,E,D,C) =? (1p)

- ☐ P(B) · P(S) · P(E|B, S) · P(D|E) · P(C|E) ✔
- ☐ P(B) · P(S) · P(E|B, S,D,C) · P(D) · P(C)
- ☐ P(B) · P(S) · P(E|B, S) · P(D|B) · P(C|E)
- ☐ P(B) · P(S) · P(E|B, S) · P(D|B) · P(C|S)

2c) Using the formula for the full joint probability distribution and the probabilities given in conditional probability table select statements which are True: (3p)

- ☐ P(¬b,¬s, ¬e,¬d, ¬c) = 0.9 · 0.9 · 0.9 · 0.9 · 0.15 ≈ 0.10.
- ☐ P(¬b,¬s, ¬e,¬d,¬c) = 0.9 · 0.9 · 0.9 · 0.9 · 0.85 ≈ 0.56. ✔
- ☐ P(d|b, s,¬c) = α · $\sum_{D.\ E}$P(b, s,E,D, ¬c), where α is the normalization factor.
- ☐ P(d|b, s,¬c) = α · $\sum_{E}$P(b, s,E, d,¬c), where α is the normalization factor. ✔
- ☐ P(d|b, s,¬c) ≈ 0.51. ✔
- ☐ P(d|b, s,¬c) ≈ 0.87.

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

0 / 10000 Word Limit

Calculator

Sektion 4

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 CSP is a consistent set of bindings to the variables that satisfy the constraints.

Problem 1: Suppose there are 5 territories T1, T2, T3, T4, and T5, each with a sensor that monitors the area associated with that territory. Each sensor has three possible radio frequencies: F1, F2, and F3. Sensors overlap if they are in adjacent areas. The adjacency relation between two territories is symmetric. Let Adj(x,y) represent the adjacency relation where Adj(T1,T2), Adj(T1,T3), Adj(T2,T3), and Adj(T3,T4). If two sensors overlap, they cannot use the same frequency.

Problem 2: The figure below depicts a constraint graph with 3 variables A, B, and C, each with the value domain {1, 2, 3, 4}. Binary constraints are associated with each arc.

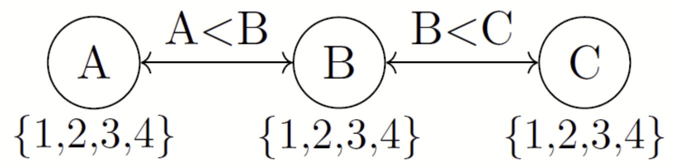

Constraint graph for Problem 2.

3a) Which of the following constraint graphs represent the CSP for **Problem 1**? (2p)

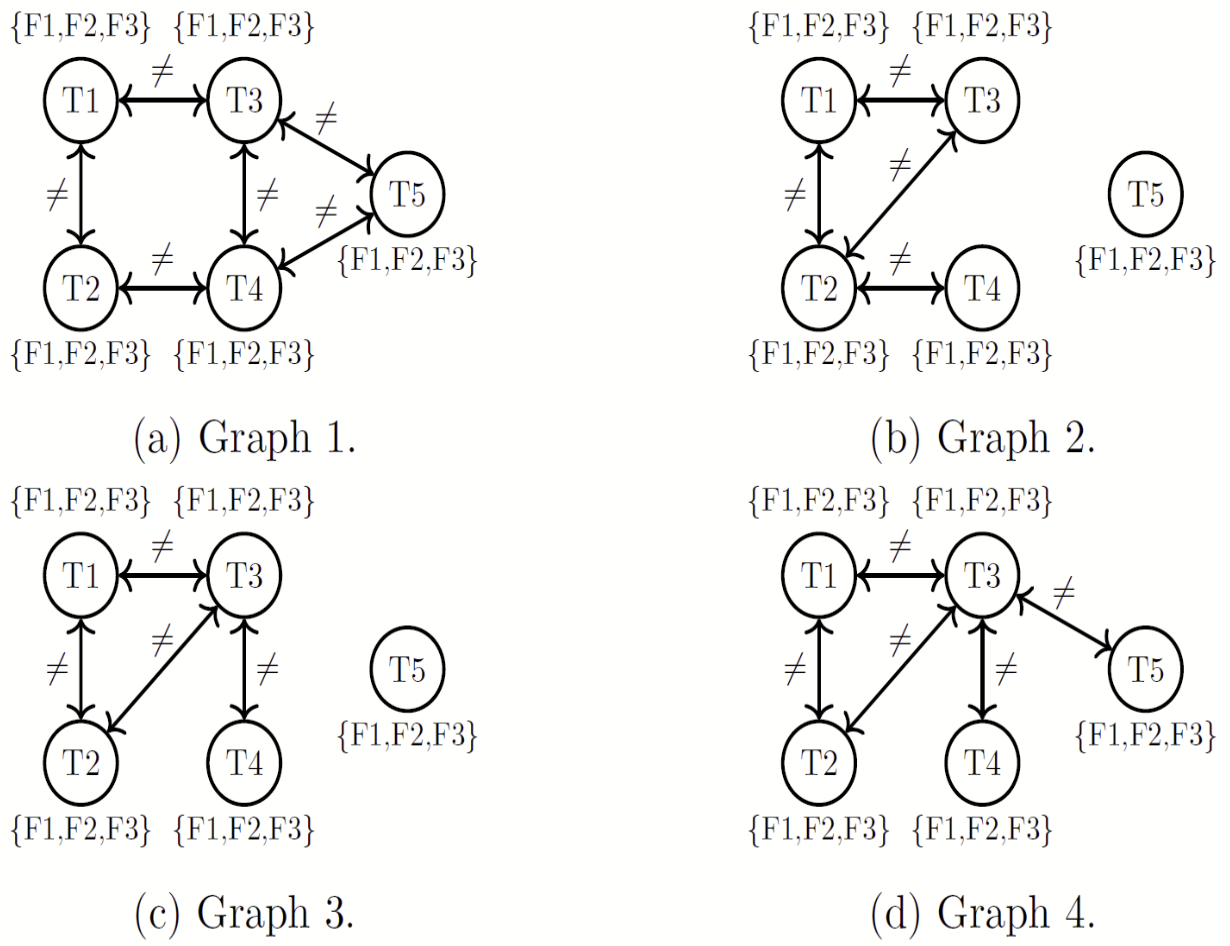

Potential constraint graphs for Problem 1.

- ◯ Graph 1
- ◯ Graph 2
- ◯ Graph 3 ✔
- ◯ Graph 4
- ◯ None of the graphs.

3b) Which of the following value assignments are globally consistent solutions to **Problem 1**?: (2p)

- ☐ T1={F1}, T2={F1}, T3={F2}, T4={F1}, T5={F1}
- ☐ T1={F1}, T2={F2}, T3={F2}, T4={F3}, T5={F2}
- ☐ T1={F1}, T2={F2}, T3={F3}, T4={F1}, T5={F1} ✔
- ☐ T1={F3}, T2={F2}, T3={F1}, T4={F3}, T5={F3} ✔

3c) Which of the following graphs will be the result of applying the AC-3 algorithm to **Problem 2**, making the constraint graph arc consistent? (2p)

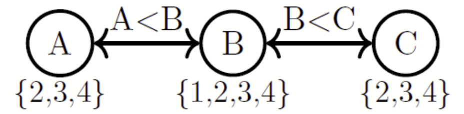

(a) Graph 1.

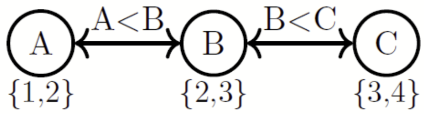

(b) Graph 2.

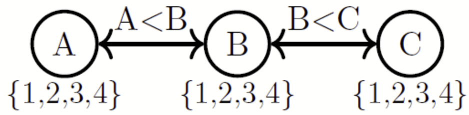

(c) Graph 3.

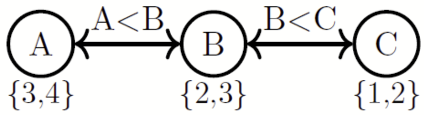

(d) Graph 4.

Potential constraint graphs after applying AC-3.

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

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.

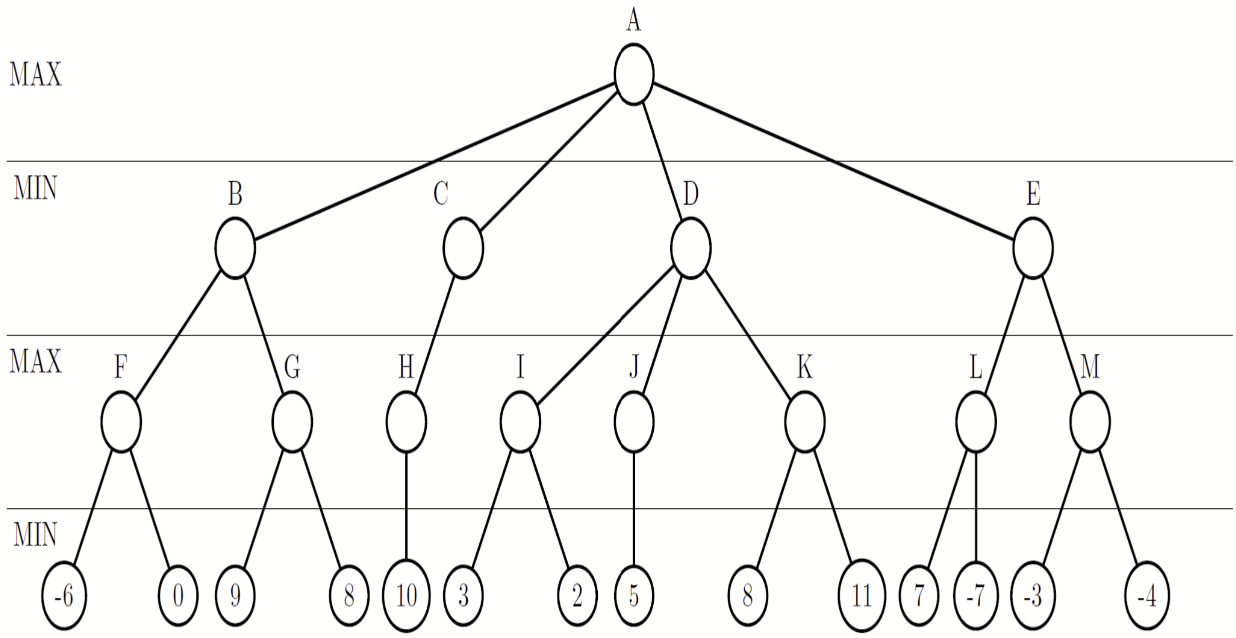

4a) Apply the MinMax algorithm to the game tree in the figure and state what move the first player (maximiser) would make. Provide heuristic values for each node in form of a table or text (e.g. A: value, B: value etc.). (2p)

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

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? Motivate your answer. (1p)

0 / 10000 Word Limit

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

0 / 10000 Word Limit

4e) 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

Sektion 6

5. Planning

Questions 5a - 5d gives 1 point for each correct answer and -1 points for each incorrect answer. In total you can get 8 points for these. You cannot get a negative score for any question.

5a) Which of the following statements are true for all propositional planning tasks? (2p)

- ☐ They have a single goal state.
- ☐ They have a finite number of state variables. ✔
- ☐ They have a finite number of plans.
- ☐ They have a finite number of states. ✔

5b) Consider a SAS+ planning task Π with 20 state variables, each with 10 values in its domain. Which of the following statements follow?

- ☐ Π has 10 · 20 states.
- ☐ Projecting Π to 5 state variables results in an abstract state space with 10^5 states. ✔
- ☐ Pattern database heuristics for Π are not safe.
- ☐ Pattern database heuristics for Π are admissible. ✔

5c) Which of the following statements related to patterns and PDBs are true?

- ☐ The estimates of a PDB heuristic correspond to path costs in an abstract transition system. ✔
- ☐ A projection maps concrete states to heuristic estimates.
- ☐ Adding a variable to a pattern can never make the resulting PDB heuristic more accurate.
- ☐ For any PDB heuristic h, state s and action a, it holds that h(s) ≤ cost(a) + h(s[[a]]), where s[[a]] is the state that results from applying a in s. ✔

5d) Let Π be a planning task. Which of the following statements about the delete relaxation Π+ are guaranteed? (2p)

- ☐ Given Π, Π+ can be computed in polynomial time in the size of Π. ✔
- ☐ Π and Π+ have the same number of states. ✔
- ☐ The costs of an optimal plan for Π and an optimal plan for Π+ are the same.
- ☐ h^FF computes the cost of an optimal plan for Π+.

6. Machine Learning

6a) In supervised learning, what is the purpose of the loss function? Give an example of a loss function and explain what type of solution it will favor. (3p)

0 / 10000 Word Limit

6b) Provide your best arguments for why reinforcement learning is a supervised learning method and why reinforcement learning is an unsupervised learning method? (2p)

0 / 10000 Word Limit

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

0 / 10000 Word Limit

6d) Is an autoencoder a supervised or unsupervised method? Motivate your answer. (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