TDDC17 Exam Reading List 2023

Fredrik Heintz

2023-10-21

It is important to read the pages explicitly specified in the bullet list for each chapter. It is also important to review the slides. Questions may be taken directly from them or from the book contents. It is definitely good if you have completed the labs before the exam although this is not obligatory.

The reading list is based on edition 4 of the Russell & Norvig course textbook (global edition).

The exam generally consists of around 8-10 questions. Most often, there will be questions on the exam chosen from the following (although this list is not necessarily complete):

1. Conceptual question about parts and types of intelligent agents. (ch. 2)
2. Search in general. Uninformed search. A* search and proof of optimality. Admissible heuristics. Other search algorithms such as local stochastic search, genetic algorithms, simulated annealing. (ch. 3-4)
3. MinMax search, alpha-beta pruning and Monte Carlo search. (ch. 6)
4. Constraint satisfaction. Different heuristics. Backtracking algorithm. (ch. 5)
5. Conceptual questions about logic and the DPLL algorithm. Possibly a question on Propositional resolution theorem proving. (ch. 7 (part of ch 8 and 9))
6. Automated Planning. (ch. 11)
7. Reasoning about uncertainty. General probabilistic reasoning. Naive Bayes. Bayes rule. Bayesian Networks. (ch. 12-13)
8. Supervised Learning, linear regression, classifiers, loss functions. (ch. 19)
9. Deep Learning: neural networks, back propagation, Layered neural networks, stochastic gradient descent, convolutional neural networks. (ch. 23)
10. Reinforcement learning: Markov decision Processes, Policies, Q-Learning, Deep reinforcement learning. (ch. 22)

More specific focuses are provided below.

Specific Reading Guidelines:

- Chapter 1, Introduction, pp.1-53.
  - Read, good to understand background and history.
  - Look at slides. Exam questions may be taken explicitly from the slides.
- Chapter 2, Intelligent Agents, pp.54-80.
  - Read, some questions are possible.
  - Draw diagram of different agent types, explain different agent types as well as pros and cons and application environments. Different task environment types, PEAS.
  - Look at the slides. Exam questions may be taken explicitly from the slides.
- Chapter 3, Solving Problems by Searching, pp.81-100, 102-108, 115-121.
  - Read sec.3.1 - 3.4, 3.4:1-4,6, 3.5:1-3, 3.6:1-4, some questions are possible.
  - Conceptual definitions, Measuring problem solving performance. Pros and cons of different search algorithms.
  - sec. 3.5.2-3: A* search and proof of optimality. Very important.
  - sec. 3.6.1-3: Be aware of heuristic functions, ways to construct them.
  - Problem solving as search, heuristic search.
  - Look at the slides. Exam questions may be taken explicitly from the slides.
- Chapter 4, Search in Complex Environments, pp.128-137.
  - Read sec. 4.1, some questions are possible.
  - Genetic algorithms, simulated annealing hill-climbing, local beam search.
  - Look at the slides. Exam questions may be taken explicitly from the slides.
- Chapter 5, Constraint Satisfaction Problems, pp.164-182.
  - Read sec. 5.1-4, some questions are possible.
  - AC3 algorithm, backtracking search for csps, arc, node and path consistency, forward checking, variable and value ordering, etc.
  - Look at the slides. Exam questions may be taken explicitly from the slides.
- Chapter 6, Adverserial Search and Games, pp. 192-210.
  - Read sec. 6.1-4, some questions are possible.
  - Understand Minimax search, alpha-beta pruning. Heuristic alpha-beta search, Monte Carlo Tree search.
  - Look at the slides. Exam questions may be taken explicitly from the slides.

- Chapter 7, Logical Agents pp.226-255.
  - Look at the slides. Slides are very important for logic questions.
  - Read sec. 7.1-6, some questions are possible.
  - Propositional Logic: Syntax, Semantics, Conjunctive Normal Form, Reasoning, Resolution, DPLL, Local Search (GSAT, WalkSAT) and Hardness
  - Understand validity, satisfaction, consistency, interpretation, models, inference, etc.
- Chapters 11 and 16, Automated Planning, pp. 362-374, 552-568.
  - Look at the slides. They are very important for automated planning questions.
  - Read sec. 11.1-11.2.1, 11.3, 11.3.2, 16.1-16.2.2. Some questions are possible.
  - Basics: Planning Formalisms, State Spaces, Planning as Forward State Space Search, Satisficing vs. Optimal Planning.
  - Heuristics: Abstractions, Pattern Databases, Delete Relaxation Heuristics
  - Probabilistic Planning: Stochastic Shortest Path Problems, Markov Decision Processes, Linear Programming for MDPs, Policy Iteration, Value Iteration, Rewards, discount factors, and (expected) utility. Bellmans principle of optimality.
- Chapter 12, Quantifying Uncertainty, pp. 403-422.
  - Read sec. 12.1-12.6.
  - This chapter provides the basis for understanding what probabilistic reasoning is about. It is a pre-requisite for the next chapter. You should be familiar with the material, in particular Bayes' rule, naive bayes, marginalization, conditionalization, product and chain rule., etc.
  - Look at the slides. Exam questions may be taken explicitly from the slides.

- Chapter 13, Probabilistic Reasoning, pp. 430-438, 445.
  - Read sec 13.1, 13.2:1, 13.3:intro. some questions are possible.
  - Be able to answer questions about and work with Bayesian Networks.
  - Look at the slides. Exam questions may be taken explicitly from the slides.
- Chapter 19, Learning from Example, pp. 669-90, 694-704
  - Read sec. 19.1-4, 19.6
  - Look at the slides! Exam questions may be taken explicitly from the slides.
  - Understand supervised learning, training, classification, regression, decision trees, linear models, neural networks, deep learning, overfitting, curse of dimensionality
- Chapter 22, Deep Learning, pp. 801-819, 823-26, 833-35
  - Read sec. 22.1-4, 22.6, 22.8
  - Look at the slides! Exam questions may be taken explicitly from the slides.
  - Neural Networks, Gradients, loss functions, convolutional NNs, recurrent NNs, unsupervised and transfer learning.
- Chapter 23, Reinforcement Learning, pp. 840-858
  - Read sec. 23.1-3, 23.4.1-3, (also ch 16: 16.1-2 for background)
  - Understand the Q-learning algorithm and concepts learnt from the lab.
  - Review Q-learning example from lecture
  - Look at the slides! Exam questions may be taken explicitly from the slides.
  - Understand reinforcement learning, utility, policy, Q-learning, exploration, curse of dimensionality