CSCI-UA 9472. Artificial Intelligence
Material for the Final

Augustin Cosse

December 6, 2021

1 Material covered

Introduction

1. You must be able to define the notions of intelligent agent, environment and action selection, and explain the various ways an agent can interact with its environment including the notions of Reflex and model based agent as well as goal based and utility based agent

Search Agent

2. You must be able to understand and explain the different search methods, in particular the differences between Informed and Uninformed search, and between those two approaches and Hill Climbing
3. You must be able to give the pseudo code for the Uninformed Search Methods (Breadth First and Depth First Search) as well as the Informed A^* Search.
4. You must be able to explain the notion of completeness in the framework of search methods.

Logical Agent

5. You must be able to describe the syntax and semantics of Propositional Logic and give the Truth table of each of the logical connectives.
6. You must be able to translate a simple logical expression from Propositional Logic into a conjunctive normal Form
7. You must be able to define Horn and definite clauses
8. You must be able to explain the interest of Horn clauses and definite clauses for inference and you must be able to turn such clauses into implications
9. You must be able to give the pseudo code for the forward and backward search algorithms in PL
10. You must be able to give the PL resolution rule and use it on simple examples
11. You must be able to give the pseudo code for the resolution algorithm in PL.

12. You must be able to define the notions of entailment, inference, completeness and soundness
13. You must be able to compare the Resolution, Forward and Backward Chaining algorithms in terms of their respective complexity and vis a vis the notion of completeness.
14. You must be able to explain the syntax and semantics of First Order Logic (in particular what changes with respect to Propositional Logic)
15. You must be able to use and explain the notions of existential and universal quantifiers
16. You must be able to explain and use the universal and existential instantiation rules (i.e quantifier elimination).
17. You must be able to use and explain the Generalized Modus Ponens rule in the framework of FOL inference.
18. You must be able to explain the result of Herbrand and the notion of semidecidability from Turing and Church
19. You must be able to explain the concept of Unification and give the corresponding pseudo code.
20. You must be able to give the pseudo code for the FOL forward chaining algorithm
21. You must be able to explain how to turn a FOL sentence into a conjunctive normal form (in particular the Skolemization and Standardization steps)
22. You must be able to state the FOL resolution rule and apply it on very simple examples
23. You must be able to explain the frame and qualification problems.

Learning Agent

24. You must be able to explain the difference between reasoning agents and learning agents
25. You must be able to define the concepts of supervised, unsupervised, semi-supervised and reinforcement learning and give an illustration for each.
26. You must be able to discuss the test training split and explain why such a split is important in learning.
27. You must be able to explain how to build a decision tree from a set of examples $\left\{\boldsymbol{x}^{(i)}, t^{(i)}\right\}$ where $t^{(i)}$ encodes a binary decision (Yes/No or True/False).
28. You must be able to explain how one can use the entropy to order the features when learning a decision tree
29. You must be able to understand and explain the notion of tradeoff between training accuracy and model complexity and its connection to regularization
30. You must be able to understand and explain the notion of overfitting through the example of polynomial features.
31. You must be able to explain how to learn a simple linear model on some data through the minimization of the residual sum of squares criterion.
32. You must be able to discuss the various models (logistic regression and perceptron) that can be derived, from the simpler linear model, by adding non linear activation functions. You must be able to know their associated activation function and preferred loss. You must be able to explain how those models are trained.

33. You must be able to give the general expression of a one hidden layer neural network and to provide the associated diagram

34. You must be able to explain backpropagation, give the main steps and the associated equations.

Markov Decision Processes and Reinforcement learning

35. You must be able to characterize a Markov Decision Process

36. You must be able to describe stationarity for preferences and explain the implication of this assumption on the definition of the utility.

37. You must be able to describe the notion of (Maximum Expected Utility) MEU policy

38. You must be able to define the Bellman equation

39. You must be able to describe Value Iteration (including the Bellman update).

40. You must be able to explain policy iteration

41. You must be able to explain the difference between active and passive learning

42. You must be able to understand and explain direct utility estimation

43. You must be able to characterize Adaptive Dynamic Programming (ADP) agents

44. You must be able to characterize the Temporal Difference (TD) learning agent

45. You must be able to describe the Multi-armed bandit problem and the corresponding Bandit Algorithm (including the tradeoff between exploration and exploitation)

46. You must be able to discuss the notion of action utility function and how such a function can be learned through Q-learning (both temporal difference Q-learning and its State-Action-Reward-State-Action (SARSA) extension). You must be able to characterize the Exploratory Q-learning agent

47. You must be able to explain how to extend simple TD learning and Q learning to parametric representations of the utility function and of the Q-table.

Biology-inspired computing

48. You must be able to describe the implementation of a simple Genetic Algorithm (including the initialization, the selection (Roulette wheel and rank), the notion of recombination/crossover and mutation)

49. You must be able to describe how to implement a simple Evolution Strategy (including recombinations and mutations) for both (μ, λ) and $(\mu + \lambda)$ selection.