

Detailed breakdown of topics and resources to support preparation for an admissions interview

Artificial Intelligence

1. Agent-based and multi-agent systems

Topics	Notes, textbooks
1. The definition of an agent, its components, and the concept of a rational agent. Specification and types of task environments. The internal structure and behavior of agents. The relationship between the search space and the fundamental types of agents.	[1] 2. Intelligent Agents, pg. 54-78
2. Problem-solving through search. Graph search algorithms. Breadth-first and depth-first traversal, depth-limited and iterative deepening search, bidirectional search. Issues of completeness, optimality, and time and space complexity.	[1] 3.1 Problem-Solving Agents, 3.3 Search Algorithms, 3.4 Uninformed Search Strategies pg. 81-102.
3. Heuristics and informed search.	[1] 3.5 Informed Search Strategies, 3.6 Heuristic Functions pg. 102-122.
4. Basics of utility theory, preferences, utility functions, simple decisions.	[1] 15.1 Combining Beliefs and Desires under Uncertainty, 15.2 The Basis of Utility Theory, 15.3 Utility Functions pg. 518-530.

II. Knowledge representation and inference

Topics	Notes, textbooks
1. Formalization of knowledge with logic. Boolean operators, propositional logic	[1] 7.1 Knowledge-Based Agents 7.3 Logic

	7.4 Propositional Logic pg. 225-240.
2. Logical inference and theorem proving. Resolution.	[1] 7.5 Propositional Theorem Proving 7.5.1-7.5.2 pg. 240-247.
3. Logical inference using Horn-clauses. Forward and backward chaining.	[1] 7.5.3 Horn clauses and definite clauses 7.5.4 Forward and backward chaining pg. 247-250.

III. Machine learning fundamentals

Topics	Notes, textbooks
1. Supervised learning, training and test sets, generalization. Bias-variance trade-off.	[1] 19.2 Supervised Learning pg. 671-675.
2. Learning decision trees. Classification, learning from samples. Attribute test, entropy, information gain.	[1] 19.3.1 Expressiveness of decision trees, 19.3.2 Learning decision trees from examples, 19.3.3 Choosing attribute tests pg. 675-680.
3. Linear regression. Univariate and multivariate linear regression, loss function.	[1] 19.6.1 Univariate linear regression, 19.6.3 Multivariable linear regression pg. 694-695, 697-700.
4. Linear classification with logistic regression.	[1] 19.6.5 Linear classification with logistic regression pg. 702-704.
5. Performance measures for binary classification. Sensitivity, specificity, accuracy.	https://en.wikipedia.org/wiki/Evaluation_of_binary_classifiers
6. Neural networks. A basic neuron and activation functions. Feed-forward neural network architectures, hidden layer, output layer.	[1] 22.1 Simple Feedforward Networks pg. 801-805.
7. Learning and gradients. Gradient descent, learning rate.	[1] 19.6.2 Gradient descent pg. 695-697. 22.1.2 Gradients and learning pg. 805-807.

Resources

1. *Stuart Russell, Peter Norvig: Artificial Intelligence: A Modern Approach (4th Edition)*.
Pearson 2020, ISBN 9780134610993, Global edition:
<https://aima.cs.berkeley.edu/global-index.html>