Expected knowledge from course
NAIL120 Introduction to Artificial Intelligence
Version 11.05.2025

Intelligent agents, environment, and structure of agents:
- Define an agent and a rational agent.
- Explain properties of environment (partial/full observability, deterministic/stochastic, episodic/sequential, static/dynamic, discrete/continuous, single/multi-agent).
- Explain and compare reflex and goal-based agent architectures and possible representations of states (atomic, factored, structured).

Problem solving and search:
- Formulate a well-defined problem and its (optimal) solution, give some examples.
- Explain and compare tree search and graph search, discuss memory vs time.
- Explain and compare node selection strategies (depth-first, breadth-first, uniform-cost).
- Informed (heuristic) search: explain evaluation function f and heuristic function h, define admissible and monotonous heuristics and prove their relation. Describe A* algorithm and prove its properties (optimality for tree search and graph search).

Constraint satisfaction:
- Define a constraint satisfaction problem (including the notion of constraint) and give some examples of CSPs.
- Apply problem solving techniques to CSPs, explain why a given technique is appropriate for CSPs.
- Explain principles of variable and value ordering and give examples of heuristics.
- Define arc consistency and show an algorithm to achieve it (AC-3).
- Define k-consistency and explain its relation to backtrack-free search; give an example of a global constraint.
- Explain forward checking and look ahead techniques.

Knowledge representation and propositional logic:
- Define a knowledge-based agent.
- Define a formula in propositional logic, describe conjunctive and disjunctive normal forms (and how to obtain them), define Horn clauses.
- Explain the notions of model, entailment, satisfiability, and their relations.
- Explain DPLL algorithm (including the notions of pure symbol and unit clause).
- Explain resolution algorithm (and how it proves entailment) and explain forward and backward chaining as its special cases.

Automated planning:

- Define planning domain and problem (representation of states, planning operator vs action, applicability and relevance of action, transition function, regression set).
- Explain progression/forward planning and regression/backward planning.
- Describe how planning can be realized by logical reasoning (situation calculus).

Knowledge representation and probabilistic reasoning:

- Define the core notions (event, random variable, conditional probability, full joint probability distribution, independence).
- Explain probabilistic inference (Bayes' rule, marginalization, normalization, causal direction, diagnostic direction).
- Define Bayesian network, explain its relation to full joint probability distribution; describe a method for constructing Bayesian networks (explain chain rule); describe inference techniques for Bayesian networks (enumeration, variable elimination, rejection sampling, likelihood weighting).

Probabilistic reasoning over time:

- Define transition and observation models and explain their assumptions.
- Define basic inference tasks (filtering, prediction, smoothing, most likely explanation) and show how they are solved.
- Define and compare Hidden Markov Model and Dynamic Bayesian Network.

Decision making:

- Formalize rationality via maximum expected utility principle (define expected utility and describe relation between utility and preferences).
- Define decision networks and show how the rational decision is done.
- Define a sequential decision problem (Markov Decision Process and its assumptions) and its solution (policy); describe Bellman equation and techniques to solve MDP (value and policy iteration); formulate Partially Observable MDP and show how to solve it.

Adversarial search and games:

- Explain core properties of environment and information needed to apply adversarial search.
- Explain and compare mini-max and alpha-beta search.
- Define an evaluation function and give some examples.
- Describe how stochastic games are handled (expected mini-max).
- Define single-move games, explain the notions of strategy, Nash equilibrium, Pareto dominance, explain Prisoner's dilemma; define maximin technique and show some strategies for repeated games.
- Mechanism design: explain classical auctions (English, Dutch, sealed bid) and tragedy of commons (and how it can be solved).

Machine learning:

- Define and compare types of learning (supervised, unsupervised, reinforcement), explain Ockham's Razor principle and difference between classification and regression.
- Define decision trees and show how to lean them (including the definition of entropy and information gain).
- Describe principles and methods for learning logical models (explain false negative and false positive notions, describe current-best-hypothesis and version-space learning).
- Describe linear regression, show its relation to linear classification (define linearly separable examples) and artificial neural networks (describe backpropagation technique).
- Explain parametric and non-parametric models, describe k-nearest neighbor methods and the core principles of Support Vector Machines (maximum margin separator, kernel function, support vector).
- Describe and compare Bayesian learning and maximum-likelihood learning; describe parameter learning for Bayesian networks including expectation-maximization (EM) algorithm (explain the notion of hidden variable).
- Formulate reinforcement learning problem, describe and compare passive and active learning; describe and compare methods of direct utility estimation, adaptive dynamic programming (ADP), and temporal difference (TD) for passive learning; explain the difference between model-based and model-free approaches; explain active version of ADP and the notions of greedy agent and exploration vs exploitation problem; describe exploration policies (random vs exploration function); describe and compare methods Q-learning (including Q-value and its relation to utility) and SARSA.