Expected knowledge from course NAIL70 Artificial Intelligence II.
Version 25.05.2025

Foundations of probabilistic reasoning:
- Define the core notions (sample space, event, random variable, conditional probability, full joint probability distribution, independence).
- Explain basic probabilistic inference concepts (marginalization, normalization, hidden variables, enumeration).
- Explain Bayes' rule and its usage for inference (explain and compare causal direction and diagnostic direction), explain Naïve Bayes model.

Bayesian networks:
- Define Bayesian network (including conditional probability tables and conditional independence within DBN), explain its relation to full joint probability distribution
- Describe and prove correctness of a method for constructing Bayesian networks (explain product and chain rule).
- Describe and compare exact inference techniques for Bayesian networks (enumeration, variable elimination).
- Describe and compare approximate inference techniques for Bayesian networks (rejection sampling, likelihood weighting, Markov Chain Monte Carlo); show why these methods are correct.

Probabilistic reasoning over time:
- Define transition and observation models and explain Markov assumptions.
- Define basic inference tasks (filtering, prediction, smoothing, most likely explanation) and show how they are solved (recursive formulas and how they are obtained).
- Define Hidden Markov Model and show how it can be used for smoothing, full smoothing, and smoothing with fixed time lag.
- Compare HMM and Dynamic Bayesian Network, explain particle filtering.
- Describe how transient and persistent failures of sensors are modeled.

Decision making:
- Formalize rationality via maximum expected utility principle (define expected utility and describe relation between utility and preferences).
- Explain human judgement notions: certainty effect, ambiguity aversion, framing effect, anchoring effect.
- Explain value of information and how it is used in decision making (information gathering).
- 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); how utility is defined (finite vs. infinite horizon, discounted vs additive rewards); describe Bellman equation and techniques to solve MDP (value and policy iteration).
- Formulate Partially Observable MDP and show how to solve it (belief states and game-theoretic approach).

Game theory:

- Define single-move games, explain the notions of strategy (pure vs. mixed), Nash equilibrium, Pareto dominance, explain Prisoner's dilemma; define maximin technique and show some strategies for repeated games.
- Mechanism design: explain and compare classical auctions (English, Dutch, sealed bid), their properties and strategies; explain tragedy of commons and how it can be solved (Vickrey-Clarks-Groves).

Supervised 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 use them (in practice) and how to learn them (including the definition of entropy and information gain); explain how overfitting is solved (pruning).
- Describe linear regression and show its relation to linear classification (define linearly separable examples).
- Explain parametric and non-parametric models, describe k-nearest neighbor methods (for classification and regression) and the core principles of Support Vector Machines (maximum margin separator, kernel function, support vector).
- Explain ensemble learning and boosting.

Learning knowledge models:

- Describe principles and methods for learning logical models (explain false negative and false positive notions); describe current-best-hypothesis, version-space learning, and inductive logic programming.
- Describe and compare Bayesian (statistical), MAP (maximum a posteriori hypothesis), MDL (minimum description length), and ML (maximum-likelihood) learning.
- Describe parameter learning for Bayesian networks including expectation-maximization (EM) algorithm (explain the notion of hidden variable).

Reinforcement learning:

- 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.