AI topics/questions

1. Intro
    a. The four approaches to AI.
    b. The Turing test
    c. Acting rationally. The rational agent.
2. Agents
    a. Agent function, agent program, agent types/architectures.
    b. Environment properties: Observable, deterministic, static, single-agent.
    c. The reflex agent architecture.
    d. The utility-based agent architecture.
3. Problem-solving with search
    a. Problem types.
    b. The single-state problem formalization.
    c. The general tree search algorithm (pseudo-code only).
    d. States vs. nodes
    e. The four evaluation metric/properties for search strategies: completeness, space-complexity, time-complexity, optimality (branching factor, diameter of the state space).
    f. Uninformed search.
        i. Breadth-first (concept, pseudocode, properties)
        ii. Depth-first (concept, pseudocode, properties)
        iii. Iterative deepening (depth-limited depth-first) search (concept, pseudocode, properties)
        iv. Comparison of properties.
4. Informed search
        i. Heuristic function
            1. admissibility
            2. dominance
            3. derivation with relaxing the problem
        ii. Greedy search (concept, pseudocode, properties)
        iii. A* (concept, pseudocode, properties) optimality with informal proof
5. Local search
    a. The hill-climbing algorithm (pseudocode)
    b. Problems with the hill-climbing algorithm
    c. Simulated annealing
6. Logic
    a. The concept of general purpose inference and domain specific knowledge-base.
    b. Logic: syntax and semantics (conceptualization).
    c. The syntax of propositional logic.
    d. The concept of models wrt KBs and the model-based definition of semantic inference: entailment.
    e. Implementation of entailment in propositional logic (pseudocode).
    f. (Syntactic) inference: elementary steps: modus ponens, resolution.

g. Relation between entailment and (syntactic) inference: soundness, completeness.
h. The forward-chaining proof method (pseudocode)
i. Conversion of a KB to CNF form.
j. The resolution-based proof method (pseudocode)
k. The first-order logic:
    i. Advantages
    ii. Quantifiers

7. Uncertainty
    a. The subjective interpretation of probability
    b. Decision theory: the binary decision problem (which action?)
    c. Probability theory
        i. Atomic events, composite events, joint distribution
        ii. Conditional probability, the chain rule
        iii. The Bayes rule,
            1. prior and posterior probabilities
            2. relevance: causal and diagnostic direction
        iv. Independence, conditional independence
    d. Inference by enumeration
    e. The naive Bayes model.
        i. The product form for the joint distribution.
        ii. Diagnostic inference
        iii. The structure.
8. The Bayesian networks.
    a. Syntax.
    b. Compactness (for binary random variables with max k parents).
    c. Global semantics (the product decomposition of the joint wrt the structure)
    d. Construction steps.
9. Inference in Bayesian networks.
    a. Tasks: simple query, composite query, relevance
    b. Inference by enumeration (pseudocode).
    c. Inference by stochastic simulation
        i. Sampling from an empty network (concept, pseudocode).
        ii. Gibbs sampling, Markov blanket.
10. Temporal probability models
    a. Definition of a Markov process (homogeneous).
    b. Definition of a Hidden Markov model (homogeneous).
        i. Inference tasks: definitions of filtering, most likely explanation, smoothing.
            1. Filtering (concept, derivation, pseudocode)
11. Decision theory
    a. Utility theory, preferences, the conditions for the existence of a utility function.
    b. The maximum expected utility principle.
    c. Decision network: elements and structure.
    d. Value of perfect information, formula

12. Learning
    a. The function approximation view of inductive learning.
    b. The Ockham principle, Hume's problem of induction.
    c. Bayesian learning
        i. Bayes rule
        ii. Posterior probability of a model/hypothesis
        iii. Prediction using averaging, MAP and ML approximations.
    d. The learning curve.
    e. The bias-variance dilemma
    f. Probably Approximately Correct (PAC) learning
        i. definition
        ii. the misclassification rate as loss
        iii. derivation of sample complexity of concept learning in i.i.d. context (independent identically distributed)
    g. Reinforcement learning
13. The decision tree representation.
    a. Expressivity
    b. Cardinality
    c. Learning method