

Intro to AI Exam Review

Introduction (1.1-1.3)

- ◆ What is AI? Thinking/Acting like Humans/Rationally
- ◆ Turing Test
- ◆ Rational Behavior
- ◆ Fields that influenced AI
- ◆ Kinds of early AI systems

Intelligent Agents (2.1-2.4)

- ◆ What is an Agent?
- ◆ Percept, Percept Sequence
- ◆ Agent Function, Agent Program
- ◆ Rational Agent
- ◆ Task Environment Specification and Properties
- ◆ General Structure of Agents

Search Problems (3.1-3.2)

- ◆ Formulating a search problem
 - State space, initial state, successor function, goal test, path cost
 - Validity and usefulness
- ◆ State space vs. search tree
- ◆ Assumptions in basic search
- ◆ Important parameters

Blind Search (3.3-3.4)

- ◆ Data structure of a search node
- ◆ Expanding a node
- ◆ The fringe of a search tree
- ◆ Search algorithm #1 and #2
- ◆ The search strategy
 - Breadth-First, Depth-First, Depth-Limited, Iterative Deepening, Uniform Cost
- ◆ Performance measures

Heuristic Search (4.1-4.2 ~4.3)

- ◆ Evaluation function and Best-First search
- ◆ Heuristic function
 - Admissible, consistent, accuracy
 - Finding an admissible heuristic
 - Some examples for 8-puzzle and robot navigation
- ◆ A* Search Algorithm

Constraint Satisfaction Probs. (5.1-5.2)

- ◆ Formulating a CSP
 - variables, domains and constraints
- ◆ Valid assignment, complete assignment
- ◆ CSP as a search problem
- ◆ Commutativity (why useful?)
- ◆ CSP Backtracking Algorithm
 - With Forward Checking
 - Pick what variable? Assign what value?

Adversarial Search (6.1-6.3)

- ◆ The specific setting for adversarial search or "game playing"
- ◆ Constructing an evaluation function
- ◆ Minimax Game Tree and Algorithm
- ◆ Alpha-beta Pruning and Algorithm

Propositional Logic (7.1-7.4 ~7.5)

- ◆ Representing the world, possible worlds
- ◆ PL Symbols, Syntax and Semantics
- ◆ Models, Knowledge Bases, Satisfiability, Entailment and Equivalence
- ◆ Inference, Modus Ponens, Modus Tollens
- ◆ Soundness and Completeness

First-Order Logic (8.1-8.3)

- ◆ Why FOL?
- ◆ Objects, Relations and Functions -> Constants, Predicates and Functions
- ◆ Atomic and Complex sentences
- ◆ Truth in FOL. Models and Interpretation
- ◆ Universal and Existential quantification
- ◆ Building a KB. Axioms vs. Theorems

Planning (11.1-11.2 ~11.4)

- ◆ STRIPS
 - Closed and Open world assumptions
- ◆ Goal, State and Action. Action schema
- ◆ Forward planning
 - Applicable action, consistent heuristic, planning graph, level cost of a goal
- ◆ Backward planning
 - Regression of a goal

Uncertainty (13.1-13.6)

- ◆ Sources of uncertainty
- ◆ Belief State
- ◆ Random Variables, Domains
- ◆ Atomic events
- ◆ Prior probability, probability distribution, full joint probability distribution
- ◆ Conditional / Posterior probability
- ◆ Bayes Rule

Bayesian Networks (14.1 ~14.2)

- ◆ Probabilistic Belief State
- ◆ Probabilistic Inference
- ◆ Conditional Probability
- ◆ Issues with inference using full joint probability distribution tables
- ◆ Independence and Bayesian networks
- ◆ Conditional Probability Tables (CPTs)

Machine Learning (slides)

- ◆ Supervised Learning
 - Approaches
 - ◆ Classification, Decision tree, Regression
 - Hypothesis, Overfitting / Underfitting
- ◆ Unsupervised Learning
 - Clustering (K-Means)
- ◆ Reinforcement Learning
 - Delayed reward

Embodied Conversational Agents (slides)

- ◆ Different uses of Embodied Conversational Agents (ECAs)
- ◆ Why is the visible behavior important as well as the AI "inside the box"?
- ◆ Two main aspects of conversational smarts: Propositional and Interactional
- ◆ Gandalf and REA
