

CSCI 446 – ARTIFICIAL INTELLIGENCE

EXAM 1 STUDY OUTLINE

Introduction to Artificial Intelligence

- I. Definitions of Artificial Intelligence
 - A. Acting Like Humans -- Turing Test
 - B. Thinking Like Humans -- Cognitive Modeling
 - C. Thinking Rationally -- Logicist Approach
 - D. Acting Rationally -- Rational Agents
 - 1. Rationality
- II. History of Artificial Intelligence
 - A. Gestation
 - B. Early Enthusiasm, Great Expectations
 - C. Dose of Reality
 - D. Knowledge Based Systems
 - E. AI Becomes and Industry
 - F. Return of Neural Networks
 - G. Recent Events
- III. Rational Agents
 - A. Percepts
 - B. Environment
 - C. Actions

Uninformed Search

- I. Planning Agents
 - A. Optimal vs. Complete
 - B. Planning vs. Replanning
- II. Search Problem Formulation
 - A. State Space
 - B. Successor Function
 - C. Start State
 - D. Goal Test
 - E. Solution / Plan
- III. State Space Graphs and Search Trees
 - A. Tree Search
 - 1. Completeness
 - 2. Time Complexity
 - 3. Space Complexity
 - 4. Optimality
 - B. Depth First Search
 - C. Breadth First Search
 - D. Iterative Deepening
 - D. Uniform Cost Search

Informed Search

- I. Heuristics
 - A. Admissible Heuristic

- B. Consistency or Monotonicity
 - C. Dominance
 - D. Creating Heuristics – Relaxed Problems
- II. Greedy Search
 - A. Heuristic $h(n)$
- III. A* Search
 - A. Actual Cost to Current Node + Heuristic -- $g(n) + h(n)$
- IV. Graph Search
 - A. Consistency of Heuristic

Constraint Satisfaction Problems (CSPs)

- I. CSP Problem Formulation
- II. Using Search in CSPs
- III. Improving Search
 - A. Backtracking Search
 - B. Filtering
 - 1. Forward Checking
 - 2. Constraint Propagation
 - C. Arc Consistency
 - C. Ordering
 - 1. Minimum Remaining Values
 - 1. Least Constraining Value
 - D. Problem Structure
- IV. Problem Structure and Decomposition
 - A. Independent Sub-problems
 - B. Tree-Structured CSPs
 - C. Nearly Tree Structured CSPs
 - 1. Cutset Conditioning
- V. Local Search
 - A. Iterative Improvement
 - B. Hill Climbing
 - C. Genetic Algorithms

Games (Adversarial Search)

- I. Overview
 - A. Deterministic Games
 - B. Zero-Sum Games
- II. Adversarial Search – Minimax (Perfect Play)
- III. Resource Limits
 - A. Evaluation Functions
- III. α - β Pruning

Expectimax Search and Utilities

- I. Uncertain Outcomes
- II. Expectimax
- III. Optimism vs. Pessimism
- IV. Utilities and Preferences
 - A. Lotteries
 - B. Rational Preferences

- C. MEU Principles
- D. Human Utilities
 - 1. Micromorts
 - 2. QALYs
 - 3. Money – not an exact utility

Markov Decision Processes

- I. Non-deterministic Search
 - A. MDP Formulation
 - B. Policies
 - C. MDP Search Trees
- II. Utilities of Sequences
 - A. Discounting (γ)
- III. Solving MDPs
 - A. Optimal Quantities
 - 1. $V^*(s)$
 - 2. $Q^*(s,a)$
 - 3. $\pi^*(s)$
 - B. Bellman Equations
- IV. Value Iteration
- V. Policy Methods
 - 1. Policy Evaluation
 - 2. Policy Extraction
 - 3. Policy Iteration

Reinforcement Learning

- I. Offline (MDPs) vs. Online (Reinforcement Learning)
 - A. Model-Based Learning
 - 1. Learn empirical MDP model
 - 2. Solve the learned MDP
 - B. Model-Free Learning
 - C. Passive Reinforcement Learning
 - 1. Policy Evaluation vs. Direct Evaluation
 - D. Temporal Difference Learning
 - D. Active Reinforcement Learning
- II. Exploration vs. Exploitation
 - A. ϵ -Greedy
 - B. Exploration Functions
 - C. Regret
- III. Approximate Q-Learning
 - A. Generalizing Across States – Feature Based Representations
- IV. Relationship to Least Squares
 - A. Minimizing Error
 - B. Overfitting
- V. Policy Search