


Artificial Intelligence


Review for Final Examination



Mountains & Minds 1

Time and Location


- Date: December 18, 2008
- Time: 8:00–9:50 am
- Location: EPS 108



Mountains & Minds 2

Rules for the Exam



- Comprehensive (i.e., covers material from entire semester).
- Closed book, closed notes.
- You are permitted one 8½ x 11 inch “cheat sheet” (both sides), to be turned in with the exam.
- No calculators, PDAs, cell phones, blackberries, or any other electronic device.
- Use the space provided (front and back). Extra sheets of paper should not be required.
- **Print** legibly.
- Show all work for partial credit.
- Time limit will be strict.
- There will be no break. Students are not permitted to leave the room until finished unless prior arrangements are made.



Mountains & Minds 3



Topics Covered

- The following are the main topics that may appear on the exam.
 - Agent architectures
 - Search (state-space and local)
 - Games (normal form and extensive form)
 - Logic (propositional and first-order)
 - Planning (total order and partial order)
 - Probability (unconditional and conditional)
 - Graphical models (Bayesian networks, decision networks, DBNs)
 - Learning (algorithms and experimental design)
- Note that all of the reading as well as class content is “fair game,” whether covered in this review or not.


Mountains  Minds 4


Agent Architectures

- Rationality
- Feedback loop
- Agent types
 - Reactive agents
 - Deliberative agents
 - Goal-based agents
 - Utility-based agents
- Strengths and weaknesses of each
- Properties of environments


Mountains  Minds 5

Search

- State-Space Search
 - Describing a state-space.
 - Describing operators/transitions in state-space.
 - Search strategies (completeness, admissibility, complexity)
 - Depth- and breadth-first search
 - Uniform-cost search
 - Best-first search
 - A* search and admissibility
 - Heuristic evaluation functions


Mountains  Minds 6


Search

- Local Search
 - How it differs from state-space search.
 - Applicability
 - Apply on state-space search problems?
 - Apply state-space search algorithms on local search problems?
 - Hillclimbing (pitfalls?)
 - Multiple restarts
 - Simulated annealing (relationship to hillclimbing?)
 - Evolutionary methods (representation, selection, operators, schemata)


7
Mountains & Minds


Games

- Normal-form games
 - Zero sum games
 - Non-zero sum games
 - Dominance
 - Pure vs. mixed strategies (when do you need which?)
 - Nash equilibria (always exist?)
- Extensive-form games
 - Minimax
 - Negamax
 - Expectiminimax
 - Alpha-beta pruning (how to do it, performance)
- Relationship between the two forms of games?


8
Mountains & Minds


Logic

- Propositional Logic
 - Propositions/Facts
 - Connectives
 - Truth tables
 - Inference rules (MP, MT, AE, AI, OI, Resolution)
- First-order Logic
 - Predicates, Functions, Variables, and Quantifiers
 - Inference rules (Generalized MP, UE, EE, EI)
 - Forward/backward chaining
 - Unification
 - Clause form, and Skolemization
 - Resolution refutation


9
Mountains & Minds


Planning

- Total-Order Planning
 - Situation-space search
 - STRIPS and STRIPS operators
 - Progression planning
 - Regression planning
- Partial-Order Planning
 - Least-commitment search
 - Sussman anomaly
 - Plan-space search
 - POP
 - Plan linearization
 - GraphPlan


10
Mountains & Minds


Probability

- The axioms of probability (Kolmogorov's axioms)
- Joint probability distributions
- Unconditional and conditional probabilities
- Marginalization and conditioning
- Bayes' Rule
- The chain rule of probabilities
- Marginal independence and conditional independence
- Inference over joint probability distributions


11
Mountains & Minds


Graphical Models

- Bayesian Networks
 - Factorizing joint distributions
 - Conditional independence and d -separation
 - Markov blankets
 - Causal, diagnostic, and inter-causal inference
 - Explaining away
 - Inference by enumeration
 - Inference by variable elimination
 - Pointwise-multiplication
 - Complexity with and without polytrees


12
Mountains & Minds


Graphical Models

- Dynamic Bayesian Networks
 - Markov chains
 - Relationship between the Markov assumption and conditional independence
 - Transition and emission models
 - Filtering, prediction, and smoothing (what are they?)
 - Hidden Markov models
 - Forward-backward algorithm (intuition)
 - Viterbi algorithm (intuition)
 - Baum-Welch algorithm (intuition)
 - The role of traditional inference algorithms on DBNs.


13
Mountains & Minds


Graphical Models

- Decision Networks
 - Bayesian decision theory
 - *MAP* hypotheses
 - Maximum expected utility
 - Expected monetary value
 - Multi-attribute utility theory
 - Strict and stochastic dominance
 - Chance, decision, and utility nodes
 - Evaluating decision networks


14
Mountains & Minds


Learning

- Experimental design
 - Hold-out studies
 - Cross-validation studies (LOO, *k*-fold, *k* x 2-cv, stratification)
 - Hypothesis testing (*t*-test, *z*-test)
 - Parameter setting (factorial, cross-validation)
 - Performance measures (accuracy, precision, recall, loss functions, true/false positives, true/false negatives)


15
Mountains & Minds


Learning

- Logic-based Models
 - Conjunctive models (Find-S, Version Spaces)
 - Decision trees
 - Entropy and expected entropy
 - Information gain
 - Gain ratio
 - Handling continuous values
 - Pruning and early stopping


16 Mountains & Minds


Learning

- Statistical Models
 - Nearest neighbor classification
 - Voronoi diagrams
 - Distance/similarity measures (L_p norms)
 - K -nearest neighbor
 - Naïve Bayes classification
 - Conditional independence assumption
 - MAP vs ML hypotheses
 - The m -estimate
 - Relationship to Bayesian network


17 Mountains & Minds

Learning

- Connectionist Models
 - McCulloch-Pitts neuron
 - Hard limiter vs linear threshold vs sigmoid activation
 - Perceptrons and linear separability
 - Multi-layer feedforward networks
 - Gradient descent
 - Backpropagation/Delta Rule
 - Momentum


18 Mountains & Minds

Learning

- Reinforcement Learning
 - Markov Decision Processes
 - Relationships between MDPs and decision theory
 - Finding optimal value functions
 - Finding optimal policies
 - Bellman equation
 - Model-based vs. model-free learning
 - Value iteration
 - Q-learning
