Final Review

CMPT 310

Review

What did we learn this semester?

Review

What did we learn this semester?

Search problems

- methods for solving

Review

What did we learn this semester?

Search problems

- methods for solving

Logic and probability

- logic can generalize, but not good for representing uncertainty

Review

What did we learn this semester?

Search problems

- methods for solving

Logic and probability

- logic can generalize, but not good for representing uncertainty

Learning

- how to build agents that learn from experience

What to do next?

- ♢ 411 - Knowledge Representation (Delgrande)
- ♢ 412/414 - Computer Vision (Funt, Drew, Li, Mori)
- ♢ 413 - Computational Linguistics (Sarkar, Popowich)
- ♢ 418 - Computational Cognitive Architecture (Hadley)
- ♢ 419 - Game Theory (Schulte)
- ♢ 419 - Machine Learning (Mori, Schulte, Sarkar)

Final Exam Format

Approximately 10 questions (with subparts)

- Same format as midterm, but longer
- Allowed cheat sheets, 2 single-sided pages (8.5" x 11" each)
- Bring a calculator

Two broad types:

- Apply algorithm A to problem P (denoted by A)
- Discuss tradeoffs / assumptions / relative merits of algorithms / problem formulations / models (short answer)

Exam is Sat. Dec. 6 12:00-15:00

– I will be away at a conference

– Prof. Richard Vaughan will proctor the exam

Intelligent Agents

- ♢ Definitions of AI
- ♢ Rationality
- ♢ Environment types
- ♢ Agent types

Search

♢ Understand different problem types and strategies for solving each of them

♢ Naive search algorithms, tradeoffs and advantages/disadvantages of each (A)

♢ Heuristic search algorithms, why and how they work, advantages/disadvantages (A)

♢ Iterative improvment algorithms (hill climbing)

Constraint Satisfaction Problems

- ♢ Relationship to search problems
- ♢ Heuristics for solving
- ♢ Apply heuristics to problems (A)
- ♢ Algorithm for arc consistency ($\neg$ A)

Game Playing

- Minimax search (A)
- α-β pruning (A)
- Apply these algorithms to game trees

Propositional Logic

- Models and entailment
- Inference algorithms
  - Enumeration
  - Forward/backward chaining ($\neg$ A)
  - Resolution proof
- Horn clauses

First Order Logic

- ♢ Models and interpretations
- ♢ Write sentences in FOL (A)
- ♢ Actions
  - Situation calculus
  - Frame problem
  - Successor-state axioms

Uncertainty

♢ Probabilities, conditional probabilities, random variables

♢ Normalization via α constant

♢ Independence and conditional independence: reducing size of joint probability distribution

♢ Bayes' rule

♢ If you can do question 16.11 that's all you need in terms of algorithms

Bayesian Networks

♢ Constructing Bayesian Networks

♢ Joint distribution represented by a BN

♢ How to represent conditional distributions: Gaussian, Linear Gaussian, Sigmoid

♢ Methods for inference: enumeration, stochastic simulation ($\neg$ A)

Temporal Probability Models

- ♢ Hidden Markov Models
- ♢ Filtering (A)
- ♢ Most likely explanation (Viterbi) (A)

Rational Decisions

- Utilities
- Value of information (A, not really an algorithm, but a formula)
- Again, knowing how to solve problem 16.11 is enough

Learning

- ♢ Generalizing from a set of examples (inductive learning)
- ♢ Issues: consistency vs. simplicity (Ockham's razor)
- ♢ What is a decision tree?
  - How do we build one from a set of examples? (A)