CSCI 4150 Introduction to Artificial Intelligence, Fall 2001
Topics covered & Final examination information

Final examination information

The final examination is scheduled for Friday December 14 from 11:30am until 2:30pm in Amos Eaton 214. You may feel free to bring food as long as you clean up after yourself.

The examination is closed book and closed notes. No calculators are necessary or allowed; you will probably have to do a little arithmetic, but it should be simple.

The examination will be designed to test:

- conceptual understanding — the ideas behind the algorithms, which algorithm to apply to a problem and what the tradeoffs are
- detailed understanding — the intricacies of how an algorithm works and issues in its implementation.

There will be questions involving factual recall or simple explanation of concepts or algorithms, but there will also be questions asking you to apply the course material to various problems and situations. You may be asked to extrapolate from course material or apply related concepts to new problems; I think one characteristic of a good examination is that students should learn something from it.

I will release on the web page the midterm examinations for this class from Fall 1999 and 2000. I will not be releasing previous years' final examinations. I highly recommend reviewing your quizzes.

Formulas provided on the final examination

- Bayes brute force classifier:
$$h = \operatorname{argmax}_{h_i \in H} \frac{P(D|h_i)P(h_i)}{P(D)}$$
- Bayes optimal classifier:
$$v = \operatorname{argmax}_{v_j \in V} \sum_{h_i \in H} P(v_j|h_i)P(h_i|D)$$
- Bayes naive classifier:
$$v = \operatorname{argmax}_{v_j \in V} P(v_j) \prod_i P(a_i|v_j)$$
- Perceptron learning rule (and the delta rule):
$$\vec{w} \leftarrow \vec{w} + \alpha \times \vec{I} \times err$$
- Backpropagation (for a three layer artificial neural network):
$$\vec{W}_i^{(2)} \leftarrow \vec{W}_i^{(2)} + \alpha\delta_i^{(2)}\vec{X}^{(1)} \qquad \delta_i^{(2)} = [f_i(1-f_i)](d_i - f_i)$$
$$\vec{W}_i^{(1)} \leftarrow \vec{W}_i^{(1)} + \alpha\delta_i^{(1)}\vec{X}^{(0)} \qquad \delta_i^{(1)} = [f_i^{(1)}(1-f_i^{(1)})]\sum_j \delta_j^{(2)} W_{i,j}^{(2)}$$
- Information
$$I(P_1, \ldots, P_n) = \sum_{\{i|P_i \neq 0\}} -P_i \log_2 P_i$$

Topics covered in class

Introduction	
What is AI?	Chapter 1
Overview of AI techniques	Chapters 2–6
Search	
Formulating search problems	Chapter 7
State space versus search tree	
Blind searches	Chapter 8
Breadth first, Depth first, iterative deepening	
Time and space complexity, optimality, and completeness	
Heuristic search	Chapter 9
Greedy and A* search (i.e. best first searches)	
Admissibility and monotonicity of heuristics	
Properties of A* (optimality, optimal efficiency, completeness)	
Memory bounded versions of A* (IDA* and SMA*)	
Miscellaneous search strategies	Sections 10.1–2
Island driven search	
Hierarchical	
Dealing with repeated states (i.e. cycles in state space)	
Constraint satisfaction search	Chapter 11
Assignment problems	
Applications of blind search	
Heuristics to improve blind search strategies	
Application of iterative improvement algorithms, min-conflicts heuristic	
Iterative improvement algorithms	Section 11.4
Hill climbing	
Potential problems with hill climbing	
Simulated annealing	
Game playing	Chapter 12
Minimax search	
Evaluation functions	
Alpha-beta pruning	
Probabalistic games (expectimax)	
Logic	
Framework for symbolic reasoning and knowledge representation	
Formal logic systems and knowledge representation	
Entailment, proof, soundness, and completeness	
Propositional logic	Chapter 13
First order logic	Chapters 15–16
Reasoning with logics	
Inference rules	
Normal forms (Horn normal form, conjunctive normal form)	
Forward and backward chaining on Horn databases	
Resolution	Chapter 14
Resolution strategies (e.g. set of support)	

Learning	
Decision trees	
(Heuristic) learning algorithm	
Overfitting	
Pruning (rule post pruning)	
Bayesian learning/classifiers	
Basic probability, conditional probability, Bayes' rule	
Bayesian classifiers (brute force, optimal, naive)	
Artificial neural networks	Chapter 3
Perceptrons	
Layered feedforward networks of sigmoid units	
Backpropagation	