CSCI 4150 Introduction to Artificial Intelligence, Fall 2002
Final examination information

Final examination information

The final examination is on Thunsday December 12 from 11:30am until 2:30pm in DCC 324. 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 both:

- 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:
$$h = \operatorname*{argmax}_{h_i \in H} P(D|h_i)P(h_i)$$
- 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):
$$\begin{array}{rcl} W_{j,i} & \leftarrow & W_{j,i} + \alpha a_j \Delta_i \\ W_{k,j} & \leftarrow & W_{k,j} + \alpha I_k \Delta_j \end{array} \qquad \begin{array}{rcl} \Delta_i & = & \mathrm{Err}_i g' \\ \Delta_j & = & g' \sum_i W_{j,i}\Delta_i \end{array}$$
- Sequential decision problems & reinforcement learning:
$$\begin{array}{rcl} U(i) & = & R(i) + \sum_j M_{ij}U(j) \\ U(i) & = & R(i) + \sum_j M^a_{ij}U(j) \\ U(i) & = & R(i) + \max_a \sum_j M^a_{ij}U(j) \end{array} \qquad U(i) \leftarrow U(i) + \alpha(R(i) + U(j) - U(i))$$
- Information
$$I(P_1, \ldots, P_n) = \sum_{\{i|P_i \neq 0\}} -P_i \log_2 P_i$$

Final examination topics

Introduction	
What is AI?	1.1, [1.2–4]

Search	
Blind search	
Formulating search problems	3.1–4
State space versus search tree	
Six blind searches	3.5
Optimality, completeness, time & space complexity	
Avoiding repeated states	3.6
Heuristic search	
Greedy search	4.1
A* search search	4.1
Heuristic functions	4.2
Admissibility	
Memory bounded A* algorithms (IDA* and SMA*)	4.3
Iterative improvement algorithms	4.4
Hill climbing	
Simulated annealing	
Constraint satisfaction problems	
Constraint satisfaction search	3.7, p. 114, pp. 104–5, slides
Assignment problems	
Constructive vs. repair methods	
Blind search approaches	
Backtracking, forward checking	
Heuristics to improve blind search strategies	
Application of iterative improvement algorithms	
Min-conflicts heuristic	
Constraint propagation	slides
Game playing search	
Minimax search	5.1–2
Evaluation functions	5.3
Alpha-beta pruning	5.4
Probabalistic games (expectimax)	5.5

Logic	
Knowledge representation & logical systems	6.1–3
Inference & entailment	
Soundness & completeness	
Propositional logic	6.4
Horn normal form	p. 270
Conjunctive & implicative normal forms	pp. 278–9
Inference in propositional logic	6.4–5, slides
Forward and backward chaining	
Resolution refutation proofs	
First order logic	7.1–2, [7.3–4]

Logic, continued	
Inference in first order logic	9.1–6
Dealing with quantifiers	
Conversion to CNF	
Unification	[pp. 302–3]
Generalized modus ponens & resolution	
Dealing with equality	
Resolution & resolution strategies	

Learning	
Introduction	18.1–2
Classification problems	
Decision trees	18.3–4
Overfitting	
Gain ratio	
χ^2 pruning	
Rule post pruning	slides
Bayesian learning/classifiers	slides, reserves, 19.6
Probability basics	14
Bayes rule, Bayesian updating	
Conditional independence	
Brute force classifier	
Optimal classifier	
Naive classifier	
m-estimates	
Reinforcement learning	
Utility	16.1–3
Sequential decision problems	17.1–3, 20.1
Value iteration	
Policy iteration	
Passive, Active, and Q- learning	20.2–6
LMS updating	
Adaptive dynamic programming	
Temporal differencing	
Exploration	
Neural networks	
Preliminaries, Perceptrons	19.1–3
Perceptron learning rule	
Delta learning rule	
Representational power of perceptrons	
Multilayer feed-forward networks	19.4
Sigmoid units	
Backpropagation	
Representational power	
PAC learning	18.6
Genetic algorithms	20.8, slides