CSCI 4150 Introduction to Artificial Intelligence, Fall 2000
Topics covered & Final examination information

Introduction	
What is AI?	1.1, [1.2–4]
Agent framework	2.1–4
Search	
Formulating search problems	3.1–4
Six blind searches	3.5
Avoiding repeated states	3.6
Constraint satisfaction search	3.7, p. 114, pp. 104–5
Informed searches (Greedy & A*)	4.1
Heuristic functions	4.2
Memory bounded A*	4.3
Iterative improvement algorithms	4.4
Game playing	
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
Propositional logic	6.4
Inference in propositional logic	6.4–5
First order logic	7.1–2, [7.3–4]
Inference in first order logic	
Dealing with quantifiers	9.1–2
Generalized modus ponens	9.3
Forward and backward chaining	9.4
Resolution	9.6
Completeness in first order logic	9.5
Rule based systems	
Expert system basics	(11.2–4)
Knowledge engineering	8.1–2, etc.
Uncertainty	
Probability and Bayes rule	14
Bayesian belief network basics	15.1
Learning	
Introduction	18.1–2
Decision trees	18.3–4
Rule post pruning	Mitchell handout
Bayesian learning/classifiers	19.6
Optimal classifier	Mitchell handout
Naive classifier	Mitchell handout
Reinforcement learning	
Sequential decision problems	17.1–3, 20.1
Passive, Active, and Q- learning	20.2–6
Neural networks	
Preliminaries, Perceptrons	19.1–3
Multilayer feed-forward networks	19.4
PAC learning	18.6
Genetic algorithms	20.8, in class

Final examination information

The final examination will be comprehensive, though there will be greater focus on the second half of the course. The examination is closed book and closed notes. No calculators are necessary or allowed; you will probably have to do a little computation, 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 some questions involving factual recall or simply explaining concepts or algorithms, but there will also be questions asking you to apply the course material to various problems and situations.

Formulas provided on the final examination

- Bayes optimal classifier

$$arg\max_{v_j} \sum_{h_i \in H} P(v_j|h_i)P(h_i|D)$$

- Bayes naive classifier

$$v_{NB} = arg\max_{v_j \in V} P(v_j) \prod_i P(a_i|v_j)$$

- Perceptron learning rule and the delta rule:

$$W_j \leftarrow W_j + \alpha \times I_j \times Err$$

- Neural network backpropagation:

$$\begin{array}{rcllcl} W_{j,i} & \leftarrow & W_{j,i} + \alpha \times a_j \times \Delta_i & \Delta_i & = & Err_i \times g'(in_i) \\ W_{k,j} & \leftarrow & W_{k,j} + \alpha \times I_k \times \Delta_j & \Delta_j & = & g'(in_j) \sum_i W_{j,i}\Delta_i \end{array}$$

- Utility for states in a known environment for reinforcement learning

$$U(i) = R(i) + \max_a \sum_j M^a_{i,j} U(j)$$

- Q-learning update rule:

$$Q[a,i] \leftarrow Q[a,i] + \alpha(r + \max_{a'} Q[a',j] - Q[a,i])$$

- Information

$$I(P_1, \ldots, P_n) = \sum_i -P_i \log_2 P_i \qquad \text{if } \forall_i P_i \neq 0$$

- PAC learning

$$m \geq \frac{1}{\epsilon}\left(\ln\frac{1}{\delta} + \ln|H|\right)$$