CS-540: Introduction to Artificial Intelligence Summer 2003 Final Exam Study Guide

I. Planning

a.

- a. Situation calculus
- b. STRIPS planning
- c. POP planning
- II. General Machine Learning
 - Inductive learning
 - i. Classification tasks
 - ii. Regression tasks
 - b. Feature vector representation
 - c. Experimental methodology
 - i. Train/test/tune sets
 - ii. Cross-validation
 - d. Overfitting and avoidance
 - e. Hypothesis spaces
 - f. Feature spaces
 - i. Feature selection
 - ii. Feature induction
 - g. Characteristics and tradeoffs of the various machine learning algorithms
- III. Decision Trees
 - a. ID3 algorithm
 - b. Information gain and entropy
 - c. Pruning
- IV. *k*-Nearest Neighbors
 - a. Distance metrics
 - b. Voting schemes
 - c. Tuning k
- V. Neural networks
 - a. Perceptrons
 - b. Perceptron training rule, gradient descent
 - c. Multi-layer networks, backpropagation
 - d. Expressiveness of perceptrons & ANNs
- VI. Bayesian Learning
 - a. General probability theory
 - b. Bayes' rule
 - c. Minimum description length principle
 - d. Bayesian classification
 - i. Naïve Bayes classifier
 - ii. Bayesian networks
 - 1. Structure
 - 2. Learning
 - 3. Inference

VII. Advanced Machine Learning Topics

[Note: you only need a qualitative understanding of these topics... not specific algorithms/calculations]

- a. Alternatives to feature vectors
 - i. Sequence data
 - ii. Relational data
 - b. Using unlabeled data
 - i. Unsupervised / semi-supervised / active learning
 - c. Ensemble methods

Note: this is a list of topics we've covered SINCE the midterm. Howeve, r the final is CUMULATIVE. Topics on the midterm study guide will comprise 25%-33% of the final exam questions, and the rest will be on topics listed in this study guide.

The exam will be closed-book. You will be allowed a 2-sided page of hand-written notes and a calculator (I highly recommend the calculator this time).