CSE 473 FINAL REVIEW LIST
OPEN BOOK, OPEN NOTES, IN-CLASS EXAM

1. Search
   - Be able to give a formal state-space model for a problem expressed in English. Formal means to specify S, s, A, f, g and c as sets or functions as appropriate.
   - Be able to specify what would be the dead states for a given problem.
   - Be able to generate part of a search tree for a given model, either depth-first or breadth-first.
   - Be able to answer questions about the general tree-search and general graph-search algorithms given in Chapter 3 and how they differ.
   - Be able to answer questions about the completeness and complexity of the various search variants given in Chapter 3.
2. Informed Search
   - Be able to explain the use of a heuristic function in a search or to give an example of one for a stated problem.
   - Be able to motivate the use of heuristic-search vs. blind search.
   - Be able to apply any of the following search methods to a well-stated problem and show a portion of the search.
     - greedy best-first search
     - A* algorithm
     - (steepest-ascent) hill climbing
   - Be able to answer questions about admissibility and consistency with respect to heuristic functions for A*.
   - Be able to describe the simulated annealing approach and its advantages/disadvantages and variants.
   - Be able to answer questions about complexity, completeness, and optimality for the above algorithms.
3. Constraint Satisfaction Problems
   - Be able to formalize a constraint satisfaction problem by specifying the sets of variables, possible values, and constraints.
   - Be able to explain or illustrate how a backtracking tree search for a constraint satisfaction problem would work: alone or with forward checking.

- Be able to answer questions about forward checking and arc-consistency.
- Be able to compare how a general heuristic search would compare with a constraint satisfaction search when both are applicable to a given problem.

4. Logic

   - Be able to interpret predicate calculus formulas in English.
   - Be able to answer questions about the normalization done by resolution theorem provers for predicate calculus in order to get the formulas into conjunctive normal form.
   - Be able to give the clause form equivalent (CNF) of a SIMPLE set of formulas.
   - Be able to show how to produce a resolvent on a SMALL set of SIMPLE formulas.
   - Be able to perform a given small resolution proof.

5. Game Playing

   - Be able to develop a utility function for a given game or show how a given one works.
   - Be able to show how a basic minimax search works for some given example.
   - Be able to show how the alpha-beta procedure works for some given example.
   - Be able to show how shallow search might be used to improve the alpha-beta procedure.
   - Be able to answer questions about how Samuel's checker player works.
   - Be able to show how minimax generalizes to games of chance.

6. Learning

   - Be able to use a given decision tree to classify a test vector.
   - Be able to construct the best decision tree for a given training set by
     (a) yourself, given the criteria for best
     (b) information gain
     (c) information content
   - Be able to answer questions about overfitting in decision trees and what can be done about it.
   - Be able to answer questions about the ensembles: bagging, boosting, stacking, and Chou's system for classifying pap smears.
   - Be able to show how a given perceptron classifies a test vector.
   - Be able to answer questions about how perceptrons learn their weights.

- Be able to answer questions about the EM algorithm and how it differs from K-means.

7. Probabilistic Inference

   - Be able to perform inference by enumeration given the full joint distribution of several variables including the following kind of queries:
     - $P(a \land b)$
     - $P(a \lor b)$
     - $P(a \mid b)$
   - Be able to normalize distributions (to all add up to 1).
   - Be able to answer questions about how independence and conditional independence help the inference process.
   - Be able to explain the semantics of a given small Bayesian network or to construct a tiny one, given the semantics.
   - Be able to decompose a query over a given Bayesian network as in Uncertainty 3, slide 4, into pieces that can be computed from the conditional probability tables of the network.
   - Be able to answer questions about the use of sampling for inference.
   - Be able to answer questions about the use of the Markov Chain Monte Carlo procedure for inference.

8. Computer Vision

   - Be able to answer questions about the 3 stages of computer vision.
   - Be able to answer questions about how Yi Li used the EM algorithm as a classifier to recognize objects in images.
   - Be able to answer questions about how Yi also used the EM algorithm for clustering, and used the clusters to produce feature vectors for training (discriminative) classifiers to recognize objects in images using multiple features.