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

1. Logic (from the handout of MY lecture, NOT the book stuff
   - 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 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 explain what is going on in a given small resolution proof.
2. 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 variant.
   - 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 support vectors, margin, and the kernel function in support vector machines. Be able to identify the support vectors given a diagram.
   - Be able to show how K-means clustering works on a very small data set.
   - Be able to answer questions about the EM algorithm and how it differs from K-means.
3. Uncertainty
   - Be able to interpret the meaning of a given Bayesian Network.
   - Be able to answer simple questions about the use of Markov Chain Monte Carlo methodology for inference.

4. Vision

- Be able to answer questions about the low, mid, and high-level aspects of vision.
- Be able to apply a mask to a (portion of an) input image to produce an output image.
- Be able to apply an edge operator, such as the Sobel or Prewitt operators, to a portion of an image to find gradient magnitude and/or angle.
- Be able to answer questions about the difference between these simple operators and the Canny edge operator.
- Be able to answer questions about how the K-means and EM algorithms can be used to find regions in a gray-tone or color image.
- Be able to show how to reduce a high-level computer vision problem to graph matching.
- Be able to show how relational indexing would work on a given problem.
- Be able to answer conceptual questions about the use of neural nets in Rowley's face detector.