CSE 473 WINTER 2023: FINAL REVIEW LIST
OPEN BOOK, NOTES, CALCULATOR, INTERNET, TAKE-HOME EXAM

1. Search
   - 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 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 apply the A* algorithm search method to a well-stated problem and show a portion of the search.
   - Be able to describe the simulated annealing approach and its advantages/disadvantages and variants.
   - * Be able to make comparisons between A* and simulated annealing.
   - Be able to answer questions about complexity, completeness, and optimality for the above algorithms.
3. 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 show how a search with chance nodes works for both basic minimax and when applying the alpha-beta procedure.
4. 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 backtracking, forward checking, and arc consistency.

5. Machine Learning

   - Be able to show how a decision tree is constructed using entropy on a simple example.
   - Be able to show how Adaboost would work on a simple example or answer questions about how it works.
   - * Be able to show how a neural net computes its result (ie. just going forward from input to output) for a simple example, using the sigmoid activation function g(x) $= 1/(1+e^{-x})$.
   - * Be able to answer questions on neural net learning with just one layer and on back propagation with multiple layers.
   - * Be able to show the the weights are updated in a single-layer neural net, using the same sigmoid activation function. NOTE: You will be able to use the fact that for this function $g'(x) = g(x)(1 - g(x))$.
   - * Be able to answer questions about deep learning architectures.
   - * Be able to answer questions about how SVMs work with support vectors.
   - * Be able to show how K-means would work on a simple 2D example.
   - * Be able to answer questions about how K-means generalizes to the EM clustering algorithm.

6. Computer Vision

   - * Be able to answer questions on color histograms, the LBP texture operator, and how they can be used to retrieve images, ie. HW 4.
   - * Be able to answer questions on how relational indexing works in the RIO system.
   - * Be able to answer questions on the pyramid approach to detecting faces in Rowley's neural net face detector.
   - * Be able to answer questions about the difference between the two EM-based methods that Yi Li developed: 1) the one that trained one Gaussian per object in color space. 2) the one that had a two-phase learning methodology.
   - * In the two-phase learning methodology, be able to explain where the feature vectors used by the neural net come from.
   - * Be able to answer questions about the HOG operator for pedestrian detection and how it was generalized in the Deformable Parts Model.
   - * Be able to answer questions about how CNNs differ from classical neural nets. Be able to define the different layers that CNNs can have and what they DO.
   - * Be able to simulate a simple CNN by performing convolutions, pooling and/or ReLU operations on small images.