CSE 455: MIDTERM REVIEW LIST
OPEN BOOK, OPEN NOTES, IN-CLASS EXAM

1. Basics
    - Be able to show how connected components works on a given image.
    - Be able to show how mathematical morphology operators work on a given image.
    - Be able to use the operations of thresholding, connected components, mathematical morphology, and logical operators (and, or, not) to perform given tasks.
    - Be able to represent an image by the (given) properties of its regions and by a region-adjacency graph.
    - Be able to show how to apply a mask (or filter) to an image.
    - Be able to show the results of mean filtering, Gaussian filtering, and median filtering on a given image or answer questions about them.
    - Be able to show how to computer gradient magnitude and gradient direction using the Sobel or Prewitt operators and the second derivative using the Laplacian operator.
    - Be able to answer questions about the improvement of the Canny operator over the simpler Sobel and Prewitt operators.
2. Color and Texture
    - Be able to show the color (or gray scale) histogram representation of a given small image.
    - Be able to compute the Swain/Ballard histogram distance between two such image histograms.
    - Be able to compute simple texture measures such as edge magnitude and direction histograms, local binary pattern histograms, and co-occurrence matrices and resultant features.
    - Be able to answer questions about the Laws texture measure.
3. Segmentation
    - Be able to explain the differences between the three approaches: region growing, split and merge, and clustering.
    - Be able to answer specific questions about
        - K-means Clustering Algorithm
        - EM Clustering Algorithm

  - Recursive Histogram-Based Clustering Algorithm
  - Shi's Normalized Graph Cuts Algorithm
- Be able to explain the parallels between K-Means and EM.
- Be able to show how the Hough transform works for lines or circles.
- Be able to answer questions about the Hough transform and its variant, the Burns line finder.

4. Interest Operators and their Applications

- Be able to explain the difference between a detector and a descriptor.
- Be able to answer questions about invariance requirements for detectors and descriptors.
- Be able to answer questions about the calculation of the Harris corner detector.
- Be able to show how cross-correlation matching works on a simple example.
- Be able to answer questions about how the RANSAC method can cut down on the correspondences found by cross-correlation or some other matching method.
- Be able to answer questions about the calculation of the SIFT operator (both the detector and the descriptor) and what makes it scale and rotation-invariant.
- Be able to answer questions about how Fergus got his "parts" in his recognition-by-parts algorithm. (You don't have to know how the Kadir operator works, just what it produces.)
- Be able to explain the ideas, but not the mathematics, behind his object representation.
- Be able to describe how Sivic gets his "words" in Video Google and how they are used in retrievals.

5. CBIR

- Be able to answer questions about or use the following as operators in a higher-level procedure for a given image retrieval task expressed in English.
  - color histogram distance
  - gridded color distance
  - LBP histogram texture distance
  - shape projection distance
  - tangent-angle histograms
  - Rowley's face finder
  - Fleck and Forsyth's skin detector
  - Jacobs' wavelet distance

  - Andy Berman's indexing system
  - Yi's consistent line clusters
- Be able to answer questions about the difference between Yi's EM Variant as used for object classification and the standard EM algorithm.
- Be able to explain the new representation for an image used by the Generative/Discriminative algorithm in its learning phase.
- Be able to compare how multiple different features are combined by the EM Variant and in the Generative/Discriminative Algorithm.

6. 2D Object Recognition

- Be able to answer questions about alignment in general.
- Be able to show how the local-feature-focus method works on a given small example.
- Be able to show how geometric hashing would work on a given small example and answer questions about how it works.
- Be able to answer questions about pose clustering and its similarities and differences to local-feature-focus, geometric hashing, and RANSAC.
- Be able to express a given problem in the consistent labeling formalism given on slide 33.
- Be able to show an interpretation tree seach for a given problem.
- Be able to express a given problem in the relational distance formalism of slides 42-4.
- Be able to explain how appearance-based object recognition using eigenvectors of a set of training images works and how it differs from the other object recognition paradigms we covered.
- Be able to compare the different approaches for some specific given problem.

7. Motion

- Be able to explain the brightness constraint and the small motion comstraint used in optical flow detection.
- Be able to indicate regions in an image where the assumptions don't hold.
- Be able to answer questions about how the pyramid structure is used.