

395T Visual Recognition: Outline of lecture for Sept 14, 2012

- I. Specific object recognition with local feature matching
 - a. Visual words for feature matching (last time)
 - i. Forming a visual vocabulary
 - ii. Inverted file index
 - iii. Bag-of-words representation for an image
 - iv. Good and bad things about bag of words representation
 - b. Spatial verification
 - i. RANSAC (review)
 - 1. Translation-only transformation
 - 2. Affine transformation
 - ii. Generalized Hough Transform
 - 1. Main idea of voting
 - 2. Line detection as an example
 - 3. Hough for SIFT matches, Lowe's approach
 - c. Additional insights from text retrieval
 - i. Scoring retrieval results
 - ii. Tf-idf weighting
 - iii. Query expansion
- II. Large-scale visual search
 - a. Indexing features
 - i. (Inverted files)
 - ii. Efficient search algorithms for low-d points – kd-trees
 - iii. Efficient approximate search algorithms for high-d points -- hashing
 - 1. Min Hash
 - 2. Random projections
 - 3. Hash tables and probabilities of collision
 - b. Injecting supervision into the search (brief)
 - i. Metric learning / kernel learning
 - ii. Semantic hashing
 - 1. Hamming space search
 - c. Summarizing large collections
 - i. Geometric Min Hash with connected components
 - ii. Visual rank with LSH
 - iii. Frequent item-set mining
 - iv. Iconoid shift

Reminder: Assignment 1 due Sept 21.