


**I. Short answer questions [20 points]**

1. Suppose we are given a disparity map indicating offset in the x direction for corresponding points. What does this imply about the layout of the epipolar lines in the two images?
2. Name and briefly explain two possible failure modes for dense stereo matching, where points are matched using local appearance and correlation search within a window.
3. What exactly does the value recorded in a single dimension of a SIFT keypoint descriptor signify?
4. If using SIFT with the Generalized Hough Transform to perform recognition of an object instance, what is the dimensionality of the Hough parameter space? Explain your answer.

**II. Programming problem: Video search with bags of visual words [80 points]**



For this problem, you will implement a video search method to retrieve relevant frames from a video based on the features in a query region selected from some frame.

## Data

We are providing image data and some starter code for this assignment. You can access pre-computed SIFT features here </filer4b/v44q010/a4data/sift/>. The associated images are stored here: </filer4b/v44q010/a4data/frames/>. The data takes about 6 G, so you should *\*not\** try to copy it to your home directory. Just point to the data files directly in your code.

Each `.mat` file in </filer4b/v44q010/a4data/sift/> corresponds to a single image, and contains the following variables, where  $n$  is the number of detected SIFT features in that image:

<code>descriptors</code>	<code>nx128</code>	<code>double</code>	// the SIFT vectors as rows
<code>imname</code>	<code>1x57</code>	<code>char</code>	// name of the image file that goes with this data
<code>numfeats</code>	<code>1x1</code>	<code>double</code>	// number of detected features
<code>orients</code>	<code>nx1</code>	<code>double</code>	// the orientations of the patches
<code>positions</code>	<code>nx2</code>	<code>double</code>	// the positions of the patch centers
<code>scales</code>	<code>nx1</code>	<code>double</code>	// the scales of the patches

## Provided code

These are the provided code files, available in the shared `code` directory.

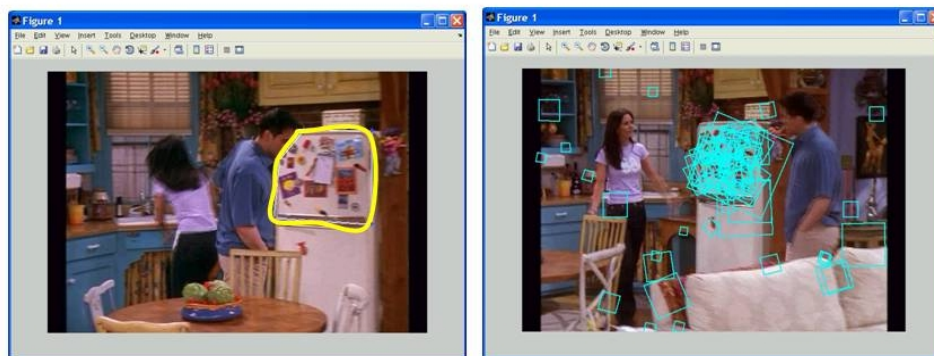
- `loadDataExample.m`: Run this first and make sure you understand the data format. It is a script that shows a loop of data files, and how to access each descriptor. It also shows how to use some of the other functions below.
- `displaySIFTPatches.m`: given SIFT descriptor info, it draws the patches on top of an image
- `getPatchFromSIFTParameters.m`: given SIFT descriptor info, it extracts the image patch itself and returns as a single image
- `selectRegion.m`: given an image and list of feature positions, it allows a user to draw a polygon showing a region of interest, and then returns the indices within the list of positions that fell within the polygon.
- `dist2.m`: a fast implementation of computing pairwise distances between two matrices for which each row is a data point
- `kmeansML.m`: a faster k-means implementation that takes the data points as columns

You are not required to use any of these functions, but you will probably find them helpful.

## What to implement and discuss in the writeup

Write one script for each of the following (along with any helper functions you find useful), and in your pdf writeup report on the results, explain, and show images where appropriate. Run/read the file `loadDataExample.m` to see how to access the provided data and draw query regions.

1. **Raw descriptor matching [20 pts]:** Allow a user to select a region of interest (see provided `selectRegion`) in one frame, and then match descriptors in that region to descriptors in the second image based on Euclidean distance in SIFT space. Display the selected region of interest in the first image (a polygon), and the matched features in the second image, something like the below example. Use the two images and associated features in the provided file `twoFrameData.mat` (in the shared `sift` directory) to demonstrate. Note, no visual vocabulary should be used for this one. Name your script `rawDescriptorMatches.m`.



2. **Visualizing the vocabulary [20 pts]:** Build a visual vocabulary. Display example image patches associated with two of the visual words. Choose two words that are distinct to illustrate what the different words are capturing, and display enough patch examples so the word content is evident (e.g., say 10-20 patches per word displayed). See provided helper function `getPatchFromSIFTParameters`. Explain what you see. Name your script `visualizeVocabulary.m`.
3. **Full frame queries [20 pts]:** Write code to do bag-of-words matching, using a linear scan through all the “database” frames. After testing your code for bag-of-words visual search, choose 2 different frames from the entire video dataset to serve as interesting successful queries, and 1 other frame to serve as an interesting failure case query. Display the  $M=5$  most similar frames to each of these queries (in rank order) based on the normalized scalar product between their bag of words histograms. Briefly explain the results. Name your script `fullFrameQueries.m`
4. **Region queries [20 pts]:** Select your favorite query *regions* (i.e. a user-specified part of the frame) from within 2 frames (which may be different than those used above) to demonstrate the retrieved frames when only a portion of the SIFT descriptors are used to form a bag of words. *Try to include example(s) where the same object is found in the most similar  $M$  frames but amidst different objects or backgrounds, and also include a failure case.* Explain the results, including possible reasons for the failure cases. Name your script `regionQueries.m`

## Tips: overview of framework requirements

The basic framework will require these components:

- **Form a visual vocabulary.** Cluster a large, representative random sample of SIFT descriptors from some portion of the frames using k-means. Let the  $k$  centers be the visual words. The value of  $k$  is a free parameter; for this data something like  $k=1500$  should work fine, but feel free to play with this parameter [see Matlab's `kmeans` function, or provided `kmeansML.m` code]
- **Map a raw SIFT descriptor to its visual word.** The raw descriptor is assigned to the nearest visual word. [see provided `dist2.m` code for fast distance computations]
- **Map an image's features into its bag-of-words histogram.** The histogram for image  $I_j$  is a  $k$ -dimensional vector:  $F(I_j) = [freq_{1,j}, freq_{2,j}, \dots, freq_{k,j}]$ , where each entry  $freq_{i,j}$  counts the number of occurrences of the  $i$ -th visual word in that image, and  $k$  is the number of total words in the vocabulary. In other words, a single image's list of  $n$  SIFT descriptors yields a  $k$ -dimensional bag of words histogram. [Matlab's `histc` is a useful function]
- **Compute similarity scores.** Compare two bag-of-words histograms using the normalized scalar product.
- **Sort the similarity scores** between a query histogram and the histograms associated with the rest of the images in the video. Pull up the images associated with the  $M$  most similar examples. [see Matlab's `sort` function]
- **Form a query from a region within a frame.** Select a polygonal region interactively with the mouse, and compute a bag of words histogram from only the SIFT descriptors that fall within that region. Weight it with tf-idf. [see provided `selectRegion.m` code]
- **Compute nearest raw SIFT descriptors.** Use the Euclidean distance between SIFT descriptors to determine which are nearest among two images' descriptors. That is, "match" features from one image to the other, without quantizing to visual words.

### III. OPTIONAL : Extra credit (up to 10 points each, max 20 points total)

- **Stop list and tf-idf.** Implement a stop list to ignore very common words, and apply tf-idf weighting to the bags of words. Discuss and create an experiment to illustrate the impact on your results.
- **Spatial verification.** Implement a spatial consistency check for an affine warp based on either RANSAC or the Generalized Hough Transform to post-process, and re-rank the initial shortlist produced based on the normalized scalar product scores. Demonstrate a query example where this improves the results. Make an output visualization showing the verified matches vs. the raw matches.
- **Shot break detection.** Design an approach to detect the shot breaks / scene changes in the video, and apply it to the provided data. Visualize the outputs by showing adjacent frames at the shot break and/or representative keyframes from each scene discovered. This approach may use the SIFT descriptors or not – any techniques from class are fair game.

Want to generate DoG+SIFT descriptors or apply other interest point operators for your own images? Check out implementations available online:

- <http://www.vlfeat.org/>
- <http://www.robots.ox.ac.uk/~vgg/research/affine/>

### **Submission instructions: what to hand in**

Please make a single zip file to submit on Canvas with:

- Your well-documented Matlab code .m files named as specified above.
- A pdf file containing
  - Your name and CS login ID at the top.
  - Your answers for Section I
  - Your image results and accompanying explanations for Section II.
  - (optional): Any results and descriptions for extra credit portions. For max credit explain what you implemented and also interpret the results. Please indicate clearly in the pdf what parts of the extra credit you did and which submitted code files are involved.