Assignment: Image stitching with RANSAC

Erik G. Learned-Miller and Cheni Chadowitz

April 15, 2014

Due: April 23, 2014 by 11:59pm by email to the TA

In this assignment, you will implement code that will automatically stitch pairs of images into simple panoramas using RANSAC.

NOTE: This assignment requires MATLAB and will NOT work with Octave. If you do not have MATLAB on your personal computer, you will need to SSH into the Edlab (use elnux1, 2 or 4) and use MATLAB remotely. You do not need to have X-forwarding set up to do this assignment. Any images normally displayed by the starter code will also be saved as PDF documents in the current directory, which you can then copy to your local machine to view. The final stitched image will also be saved as a MATLAB .mat file that you can load in Octave (or MATLAB) with

```
load stitched.mat; imagesc(stitched); colormap(gray);
```

For more information, check `http://people.cs.umass.edu/~cheni/370/matlab-octave.html#edlab`.

1. **Setup - Download the handout archive:**

   `http://people.cs.umass.edu/~cheni/370/CS370-Assignment4-StarterCode.zip`

   Extract the archive in your current working directory. If you are working on the Edlab, you can download and extract the starter code to your current working directory with

   ```
   wget http://people.cs.umass.edu/~cheni/370/CS370-Assignment4-StarterCode.zip
   unzip CS370-Assignment4-StarterCode.zip
   ```

2. **RANSAC**

   You will be implementing the main piece of the RANSAC algorithm. Open 'stitch.m' and 'ransac.m' and read the comments to understand the code. You will be filling in the loop in 'ransac.m' as described in the comments. The starter code takes care of loading the images, finding keypoints using MATLAB's Harris corner detector, and computing SIFT descriptors at the keypoints. It also computes a tentative correspondence between the SIFT descriptors in the first image to the SIFT descriptors in the second image, using Lowe's approach for thresholding.

   Your part is to fill in the loop in 'ransac.m' where specified. You should follow the algorithm as described in the comments. We will only handle transformations involving simple x,y translations, not full homographies. In other words, we are attempting to find a translation $[x, y]$ such that the descriptors in the second image D_2 best match the descriptors in the first image D_1, transformed by the given translation:

   $$D_2 = D_1 + [x, y]$$

You will need to pick an appropriate number of iterations to loop T. The value given in the starter code may suffice, but you made find that changing it produces better (or worse) results, depending on your input images.

Be sure to return the average translation estimate you find for the best representative feature, and the number of inliers for that feature.

To test your code, you can use the provided 'castle.mat', 'river.mat', and 'riverDiagonal.mat' data files. Load these into MATLAB with 'load filename.mat'. Each data file contains two variables: 'left' and 'right.' The starter code assumes that if your workspace contains these two variables, then the upper-left image in your panorama is saved in 'left', and the lower-right image in your panorama is saved in 'right'. The 'river.mat' data file contains a manually created panorama formed by splitting the 'river1.jpg' file into the left two thirds and the right two thirds. The translation for this is $[0, 342]$. The 'castle.mat' data file contains a similarly split image, but with a different translation, and the 'riverDiagonal.mat' data file has a translation in both the x and y directions.

Use these data files to test your code. The starter code will display some images along the way and will show the final stitched image. If you are using the Edlab without X-forwarding, these images will be saved to PDF files 'multiImages.pdf' and 'fullStitched.pdf' for viewing.

Lastly, try your code on a panorama of two images of your own. A simple way to do this is to take a photo and split it into upper-left and lower right pieces. Make sure the two pieces have a sufficient amount of overlap (first two thirds and last two thirds should be sufficient in most cases). You may also try using a set of photos, but make sure that the transformation between them is predominantly translational, and not rotational or otherwise.

What to turn in

Any image files should be in JPG format.

- Your modified, commented 'ransac.m' file.
- The final stitched images for the 'castle.mat' and 'riverDiagonal.mat' data files
- The original and final stitched images for your own panorama
- A single PDF named lastnameFirstname_assign4.pdf containing:
  1. your modified, commented 'ransac.m' code
  2. the number of iterations you chose T (this may remain the same)
  3. the translation you found for the 'castle.mat' and 'riverDiagonal.mat' data files
  4. the translation you found for your own panorama
  5. the final stitched images for the 'castle.mat' and 'riverDiagonal.mat' data files
  6. the original and final stitched images for your own panorama

Email these files as a zip named lastnameFirstname_assign4.zip to the TA.
Submissions will ONLY be accepted if they are ZIP archives, named as described.