

CMSC 426, Image Processing

Project 2: Myautopano!

Due on: 11:59:59PM on Sunday, Oct 16 2016

Prof. Yiannis Aloimonos,
Nitin J. Sanket and Kiran Yakkala

October 3, 2016

The aim of this project is to implement an end-to-end pipeline to do image panorama stitching. We all use the panorama mode on our smart-phones. You'll be implementing a simple pipeline which does the same. Aren't you excited?

In the next few sections, we'll detail how this has to be done along with the specifications of the functions for each part.

1 Capture Images

For this project, you need to capture multiple images of a scene, which you will use to stitch a panorama from. In general, you should limit your camera motion to purely translational, or purely rotational (around the camera center). Make sure you have about 30-50% overlap between consecutive images to have enough common features to be able to stitch panoramas.

You'll need creative ways to make it work for harder images. A set of sample images are shown in Fig. 1.



Figure 1: Sample input images used for panorama stitching.

2 Adaptive Non-maximal Suppression (ANMS) to find corners distributed equally over the image - 25Pts

The aim of this step is to detect corners spread all across the image to avoid weird artifacts in warping. Implement the following steps for ANMS:

Detect corner features in the image using `cornermetric` with the appropriate parameters. The output is a matrix of corner scores. The higher the value, the higher the score/probability of that pixel being a corner. Visualize the output using the MATLAB function `imagesc`.

Now, find the N_{strong} strongest corners using the MATLAB function `imregionalmax`. The ANMS algorithm is described next:

Input : Corner score Image (C_{img} obtained using `cornermetric`), N_{best} (Number of best corners needed)

Output: (x_i, y_i) for $i = 1 : N_{best}$

Find all local maxima using `imregionalmax` on C_{img} ;

Find (x, y) co-ordinates of all local maxima;

((x, y) for a local maxima are inverted row and column indices i.e., If we have local maxima at $[i, j]$ then $x = j$ and $y = i$ for that local maxima);

Initialize $r_i = \infty$ for $i = [1 : N_{strong}]$

```

for  $i = [1 : N_{strong}]$  do
    for  $j = [1 : N_{strong}]$  do
        if  $(C_{img}(y_j, x_j) > C_{img}(y_i, x_i))$  then
            |  $ED = (x_j - x_i)^2 + (y_j - y_i)^2$ 
        end
        if  $ED < r_i$  then
            |  $r_i = ED$ 
        end
    end
end

```

Sort r_i in descending order and pick top N_{best} points

Algorithm 1: ANMS algorithm.

Consider the case where we are stitching first 2 images in Fig. 1. The output of ANMS are shown in Fig. 2. To plot dots on the images you can do `imshow(..); hold on; plot(..); hold off;`. Note that plot uses x and y wherein x is the column number and y is the row number.

3 Feature Descriptor - 25Pts

In the previous step, you found the feature points (locations of the N_{best} best corners after ANMS are your feature points). You need to describe each feature point by a feature vector,



Figure 2: Output of ANMS on first 2 images shown in Fig. 1.

this is like encoding the information at each feature points by a vector. One of the easiest feature descriptor is described next.

Take a patch of size 40×40 centered (this is very important) around the keypoint. Now apply gaussian blur (feel free to play around with the parameters, for a start you can use MATLAB's default parameters in `fspecial` command, Yes! you are allowed to use `fspecial` command in this project). Now, sub-sample the blurred output (this reduces the dimension) to 8×8 . Then reshape to obtain a 64×1 vector. Standardize the vector to have zero mean and variance of 1 (Done by subtracting all values by mean and then dividing by the standard deviation). Standardization is used to remove bias and some illumination effect.

4 Feature Matching - 10Pts

In the previous step, you encoded each keypoint by a 64×1 feature vector. Now, you have similar feature points and feature descriptors for each point in the 2 images you are trying to stitch together. In computer vision terms, this step is called as finding feature correspondences within the 2 images. Pick a point in image 1, compute sum of square difference between all points in image 2. Take the ratio of best match (lowest distance) to the second best match (second lowest distance) and if this is above some ratio keep the matched pair or reject it. Repeat this for all points in image 1. You will be left with only the confident feature correspondences and these points will be used to estimate the transformation between the 2 images also called as *Homography*. Use the function `showMatchedFeatures` given to you to visualize feature correspondences (Sample output is shown in Fig. 3).



Figure 3: Output of `showMatchedFeatures` on first 2 images of Fig. 1.

5 RANSAC to estimate robust Homography - 15Pts

Note that not all the matches we found in the previous step will be right. We will use a robust method called RANSAC to compute homography.

Recall the RANSAC steps are:

- Select four feature pairs (at random), p_i, p_i^1 .
- Compute homography H (exact). Use the function `est_homography` given to you.
- Compute inliers where $SSD(p_i^1, Hp_i) < thresh$. Here, Hp_i computed using the `apply_homography` function given to you.
- Repeat the last three steps until you have exhausted N_{max} number of iterations (specified by user) or you found more than a percentage of inliers (say 90% for example).
- Keep largest set of inliers.
- Re-compute least-squares $\hat{H}$ estimate on all of the inliers. Use the function `est_homography` given to you.

6 Output

Produce a panorama by overlaying the pairwise aligned images to create the final output image. The following MATLAB functions should help you in this part, e.g `imtransform` and `imwarp`. If you want to implement `imwarp` (or similar function) by yourself, you should apply bilinear interpolation when you copy pixel values. The panorama output obtained for Fig. 1 is shown in Fig. 4.



Figure 4: Final Panorama output for Fig. 1 image.

7 Cylindrical Projection - 10Pts

When you try to stitch a lot of images with translation, a simple projective transformation (applying homography) will give bad results and the images will be stretched/shrunk to a large extent on the edges. A sample case is shown in Fig. 5.

Now, to overcome the distortion problems at edges, we will be using cylindrical projection on the images first before performing other operations. Essentially this is a pre-processing step.

The following equations transform from normal image co-ordinates to cylindrical co-ordinates.

$$x' = f \tan \left(\frac{x - x_c}{f} \right) + x_c$$

$$y' = \left(\frac{y - y_c}{\cos \left(\frac{x - x_c}{f} \right)} \right) + y_c$$

In the above equations, f is the focal length of the lens in pixels (feel free to experiment with values, generally values range from 100 to 500, however this totally depends on the camera and can lie outside this range). The original image co-ordinates are (x, y) and the transformed image co-ordinates (in cylindrical co-ordinates) are (x', y') . x_c and y_c are the image center co-ordinates. Note that, x is the column number and y is the row number in MATLAB.

You can use `meshgrid`, `ind2sub`, `sub2ind` to speed up this part. Using loops will TAKE FOREVER!

A sample input image and it's cylindrical projection output is shown in Fig. 6.

Note that, the above equations talk about pixel co-ordinates are not pixel values. The




Figure 5: Panorama stitched using projective transform showing bad distortion at edges.




Figure 6: Left: Original input image, Right: Cylindrical output image.

idea is you compute the co-ordinate transformation and copy paste pixel values to these new pixel co-ordinates (in all 3 channels, i.e., RGB). However, when you compute the values of (x', y') they might not be integers. A simple way to get around this is to use round or actually interpolate the values. If you decide to round the co-ordinates off you might be left with black pixels, fill them using some weighted combination of it's neighbours (gaussian works best). A trivial way to do this is to blur the image and copy paste pixel values on-to original image where there were pure black pixels. (You can also initialize pixels to NaN's instead of zeros to avoid removing actual zero pixels).

8 Blending images to get a seamless panorama - 15Pts

Come up with a logic to blend the common region between images while not affecting the regions which are not common. Here common means shared region, i.e., a part of first image and part of 2nd image should overlap in the output panorama. Describe what you did in your report. You CAN NOT use any built-in MATLAB function or third party code to do this.

9 Extra Credit

- Interesting selection of images. (5Pts)
- Use 8 or more images for stitching. (5Pts)
- Achieve rotation invariance (for feature matching) in the feature descriptor. Hint: Use dominant direction of image patch around the keypoints. (10Pts).

10 Starter Code

There is no starter code given for this project.

11 Submission Guidelines

We have given three sets of images in addition to which you need to beg/borrow/steal (or capture) two sets of images (each with ≥ 4 images). You need to submit the output panorama images for all these 5 sets in the report. Describe how you are blending images (while doing stitching) in your report. You should also include a detailed README file explaining how to run your code.

Submit your own test images (in Images/CustomSet1, Images/CustomSet2) and codes (.m files) with the naming convention YourDirectoryID_proj2.zip onto ELMS/Canvas (Please compress it to .zip and no other format). Your zip file should have the following things:

If your code does not comply with the above guidelines, you'll be given **ZERO** credit.

Also make a presentation if you want to present it in the class **Note: Every student should present AT LEAST once in the class (you can volunteer to present for more than one project) and can choose to do for any of the Projects from 1 to 4. Good presentations will receive cookie points.**

12 Allowed Matlab functions

`imfilter`, `conv2`, `imrotate`, `im2double`, `rgb2gray`, `fspecial`, `imtransform`, `imwarp`, `meshgrid`, `sub2ind`, `ind2sub` and all other plotting and matrix operation/manipulation functions are allowed.

13 Collaboration Policy

You are restricted to discuss the ideas with at most two other people. But the code you turn-in should be your own and if you **DO USE** (try not it and it is not permitted) other external codes/codes from other students - do cite them. For other honor code refer to the CMSC426 Fall 2016 website.

DON'T FORGET TO HAVE FUN AND PLAY AROUND WITH IMAGES!.

Acknowledgements

This fun project was inspired from 'Computer Vision & Computational Photography' (CIS 581) course of University of Pennsylvania (https://alliance.seas.upenn.edu/~cis581/wiki/index.php?title=CIS_581:_Computer_Vision_%26_Computational_Photography).