
ASSIGNMENT 2: PROJECTIVE GEOMETRY AND RANSAC

Please prepare a single, condensed and neatly edited document for submission. For each problem include a short description of what you did, results, and a brief discussion/interpretation. It is not necessary to include any code in your document (unless a snippet helps to explain your method). Hand in a printed copy **before 14:00** on the due date, and also e-mail me a zip-file containing your document and code.

You are free to use any programming language. I recommend Matlab or Python. If you are asked to implement a specific technique, the idea is that you do so from scratch; do not simply use a function from an image processing or computer vision library. Collaboration is restricted to the exchange of a few ideas. All code, results, and write-up that you submit must be your own work.

- Suppose we'd like to rotate an image about its centre by an angle θ , and scale the result by a factor s . This operation can be expressed as

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = s \left(R \left(\begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} \right) + \begin{bmatrix} x_0 \\ y_0 \end{bmatrix} \right), \quad \text{with } R = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

Here (x, y) are coordinates in the input image, (x', y') are corresponding coordinates in the output image, and (x_0, y_0) are the coordinates of the centre of the input image.

- Rewrite this expression in homogeneous coordinates, as $\underline{x}' = H\underline{x}$ where \underline{x} and \underline{x}' are elements of \mathbb{P}^2 and H is a 3×3 homogeneous matrix that you need to specify in terms of θ , s , x_0 and y_0 .
 - Transform a colour image with this homography, for two or three different choices of θ and s . You will find a document on the module's website detailing how we can use bilinear interpolation to apply a given homography to an image.
 - If the procedure in that document is followed, why would different values for the translation component in H (the last column's first two entries) have no effect on the output image? Also, if $\theta > 0$, why is the image rotated clockwise even though the form of R suggests counterclockwise rotation?
-

- Paste a poster of your favourite movie onto the side of the building in `griest.jpg`. More specifically, manually select the four corners of the building to which the corners of the poster image should map, calculate a homography (section 3.3 in the notes), apply the homography to the poster image, and overlay the result. Note that the transformed poster image will contain parts with no image data that we want to ignore, and its image origin may be shifted (why?). Explain how you addressed these two issues.



- The file `linedata.txt` contains 50 coordinate pairs (one per row). Your task is to recreate the three plots in Figure 4.1 of the notes, using this data. First plot the given data, then fit a line through all the points using least squares, and finally employ RANSAC to automatically identify a good set of inliers and fit a line through only those. State in your report the slope and y -intercept of each line.
-

-
4. In this problem the idea is to use a projective transformation, found from SIFT feature matches, for image stitching. Consider the two images `perold1.jpg` and `perold2.jpg`. The file `peroldsift12.txt` provides SIFT feature matches across these images, as an $n \times 4$ matrix. A row in this matrix is of the form $[x \ y \ x' \ y']$ and indicates that point (x, y) in the first image matches with point (x', y') in the second image. Note that these coordinates are (x, y) -pairs, not (row, column)-pairs.
- (a) Display all the given feature matches between `perold1.jpg` and `perold2.jpg`. Here it is advised to use the visualization illustrated in Figure 2.3(c) of the notes.
 - (b) Implement a RANSAC based procedure to find a set of inlier feature matches, and a homography calculated from them. Display the inlier matches thus obtained, in the same manner as you displayed matches in part (a).
 - (c) Use the homography determined in part (b) to transform one of the images, and paste this transformed image on top of the other. Of course, the way in which you define your homography will determine which image to transform. Also, just as in problem 2, careful consideration should be given to the location of the true origin of the transformed image (the values of `minx` and `miny` in your “apply homography” function can be very useful here).
-

Hand in: 21 August 2017