

Homework 5

Sravani Ramishetty(sramishe@purdue.edu)

Theory Question 1

How do we differentiate between the inliers and the outliers when using RANSAC for solving the homography estimation problem using the interest points extracted from two different photos of the same scene?

We differentiate between the inliers and the outliers using the decision threshold. For each data point, we compute its projection using the estimated homography matrix, if the distance between the actual data point and the projected point is less than the threshold, we call it an inlier else an outlier.

Mathematically, let \mathbf{x}_i and \mathbf{x}'_i be the corresponding points in the two images, and \mathbf{H} be the estimated homography matrix. We project \mathbf{x}_i using \mathbf{H} as:

$$\mathbf{x}'_i \sim \mathbf{H}\mathbf{x}_i$$

We then calculate the distance d_i between the actual point \mathbf{x}'_i and the projected point $\mathbf{H}\mathbf{x}_i$ as:

$$d_i = \|\mathbf{x}'_i - \mathbf{H}\mathbf{x}_i\|$$

If d_i is less than the threshold δ , the point is classified as an inlier:

$$d_i < \delta \implies \text{inlier}$$

Otherwise, it is classified as an outlier:

$$d_i \geq \delta \implies \text{outlier}$$

Theory Question 2

Explain in your own words how the Levenberg-Marquardt (LM) algorithm combines the best of GD and GN to give us a method that is reasonably fast and numerically stable at the same time.

LM can switch from GN to GD using the damping coefficient(μ). If the parameter is close to local minima, LM switches to GN and when it is far from local minima, it switches to GD.

- When the solution is far from the local minimum, the damping parameter μ is large, and the LM algorithm behaves like GD. In this regime, it takes smaller steps, focusing on stabilizing the optimization process.
- As the solution approaches the local minimum, μ becomes smaller, causing the LM algorithm to behave more like GN. In this case, it takes larger, more efficient steps that converge faster.

Task 1: Automated Homography Estimation using RANSAC and Nonlinear Least-Squares Refinement

The objective of this homework is to implement a fully automated approach for robust homography estimation between two images using RANSAC and refine it using a Nonlinear Least-Squares minimization approach, specifically the Levenberg-Marquardt (LM) algorithm.

The following steps outline the automation of homography estimation between two images:

Step 1: Interest Point Detection and Descriptor Extraction Interest points and their corresponding descriptors are extracted from both images and compared using a similarity criterion. The following techniques can be used for interest point extraction and descriptor computation:

- **Harris Corner Detector:** Detect interest points using the Harris corner detector. Then, for each interest point, compute a gray-level descriptor using a $W \times W$ window around the point. The similarity between corresponding points can be calculated using either:
 - Sum of Squared Differences (SSD)
 - Normalized Cross-Correlation (NCC) (preferred for improved robustness)
- **SIFT, SURF, or SuperPoint+SuperGlue:** Alternatively, we can use more advanced techniques such as SIFT, SURF, or SuperPoint+SuperGlue to extract both interest points and their descriptors. In these cases, the similarity between interest points is calculated using the Euclidean distance between the descriptor vectors (SuperPoint+SuperGlue gave best results in this homework).

Step 2: Linear least square minimization to obtain the initial homography

Given a point X in a planar scene and its corresponding pixel X' in the image plane, for most cameras, we can express the relationship between them as:

$$X' = HX$$

where both X and X' are expressed in homogeneous coordinates:

$$X = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}, \quad X' = \begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix}$$

and H is the 3×3 homography matrix, given by:

$$H = \begin{pmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{pmatrix}$$

Expanding the matrix multiplication for $X' = HX$, we obtain the following equations for the components of X' :

$$x'_1 = h_{11}x_1 + h_{12}x_2 + h_{13}x_3$$

$$x'_2 = h_{21}x_1 + h_{22}x_2 + h_{23}x_3$$

$$x'_3 = h_{31}x_1 + h_{32}x_2 + h_{33}x_3$$

Denoting the physical scene coordinates by (x, y) and the physical pixel coordinates by (x', y') , we have the following relationships:

$$x = \frac{x_1}{x_3}, \quad y = \frac{x_2}{x_3}$$

$$x' = \frac{x'_1}{x'_3}, \quad y' = \frac{x'_2}{x'_3}$$

Thus, we can express the physical coordinates of the image pixel as:

$$x' = \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}}$$

$$y' = \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}}$$

These expressions for the physical pixel coordinates can be rewritten in the homogeneous equation form:

$$xh_{11} + yh_{12} + h_{13} - xx'h_{31} - yx'h_{32} - x'h_{33} = 0$$

$$xh_{21} + yh_{22} + h_{23} - xy'h_{31} - yy'h_{32} - y'h_{33} = 0$$

Converting this into a matrix format and in-homogeneous form with $h_{33} = 1$ with 1 point gives below system of in-homogeneous equations:

$$\begin{bmatrix} x_1 & y_1 & 1 & 0 & 0 & 0 & -x_1x'_1 & -y_1x'_1 \\ 0 & 0 & 0 & x_1 & y_1 & 1 & -x_1y'_1 & -y_1y'_1 \end{bmatrix} \begin{pmatrix} h_{11} \\ h_{12} \\ h_{13} \\ h_{21} \\ h_{22} \\ h_{23} \\ h_{31} \\ h_{32} \end{pmatrix} = \begin{pmatrix} x'_1 \\ y'_1 \end{pmatrix}$$

This system can be represented in the general form as $Ah = b$, where A is a 2×8 matrix of known values based on the coordinates of the point (x_1, y_1) and the corresponding transformed point (x'_1, y'_1) , h is the vector of unknown homography parameters, and b is a 2×1 vector representing the coordinates of the transformed point.

System of Equations for Multiple Correspondences

For N point correspondences $(x_i, y_i) \leftrightarrow (x'_i, y'_i)$, where $i = 1, 2, \dots, N$, we can write a system of equations for each correspondence (X_i, X'_i) , represented as $A_i h = b_i$ for the i -th correspondence.

By stacking the equations for all N correspondences, we obtain a system of $2N$ inhomogeneous equations that can be expressed as:

$$Ah = b$$

where A is a $2N \times 8$ matrix of known values, b is a $2N \times 1$ vector of known values, and h is an 8×1 vector of unknown homography parameters.

Solving the Overdetermined System

In general, for $N \geq 4$, the system is overdetermined (i.e., more equations than unknowns), and we solve it using the pseudo-inverse of A . The solution is given by:

$$h = (A^T A)^{-1} A^T b$$

This provides the least-squares solution for the homography parameters, minimizing the error between the projected and actual points in the target image.




Figure 1: Given Input Images



Figure 2: Collected Input Images

Step 3: RANSAC Algorithm for Homography Estimation and outlier rejection(Pseudocode)**Algorithm 1:** RANSAC for Homography Estimation**Input:** Matched keypoints between two images, δ , ϵ , N , M **Output:** Best homography matrix H_{best} , list of inliers**Initialize:**

```

 $H_{\text{best}} \leftarrow \text{None}$ 
 $\text{best\_inliers} \leftarrow 0$ 
 $\text{best\_inliers\_list} \leftarrow \text{None}$ 
 $n_{\text{total}} \leftarrow \text{number of matched keypoints}$ 
 $p \leftarrow 0.99$ 
 $n \leftarrow \text{minimum number of points to estimate homography}$ 
 $\text{max\_iters} \leftarrow \frac{\log(1-p)}{\log(1-(1-\epsilon)^n)}$ 
 $M \leftarrow (1 - \epsilon) \times n_{\text{total}}$ 
 $\text{max\_iters\_all} \leftarrow \frac{n_{\text{total}}!}{n!(n_{\text{total}}-n)!}$ 

```

while $\text{iterations} < \text{max_iters_all}$ **do** **Step 1:** Randomly select n correspondences **Step 2:** Compute homography matrix H from the selected points **Step 3:** Compute inliers by calculating error for each correspondence **for** each correspondence j **do** Compute $p_1 \leftarrow (x_j, y_j, 1)^T$ Compute $p_2 \leftarrow (x'_j, y'_j, 1)^T$ error $\leftarrow \|p_2 - H \cdot p_1\|_2$ **if** $\text{error} < \delta$ **then**

Increment the inliers count

 Add correspondence j to inliers list **end** **end** **if** $\text{inliers count} > M$ **then** **if** $\text{inliers count} > \text{best_inliers}$ **then** Update $H_{\text{best}} \leftarrow H$ Update $\text{best_inliers_list} \leftarrow \text{current inliers list}$ **end** **Break** **end****end****Return:** H_{best} , best_inliers_list

In this implementation, the number of iterations exceeded the theoretical value $\frac{\log(1-p)}{\log(1-(1-\epsilon)^n)}$, as the keypoints detected by SIFT contained a significant number of outliers, which prevented the algorithm from finding the best homography matrix. In contrast, the iterations required for SuperPoint+SuperGlue were considerably fewer compared to SIFT. Additionally, the threshold δ used for SuperPoint+SuperGlue (25) was smaller than that for SIFT (75) in order to achieve the desired number of inliers. Once the best homography matrix H_{best} and the corresponding list of inliers were identified, the homography was refined using all the inliers found. The homography was then further optimized using the 'lm' from `least_squares` method from `scipy.optimize`.

Extracted correspondences using SuperPoint+SuperGlue between given input images



(a) Image 1 and 2 correspondences

(b) Image 2 and 3 correspondences

Figure 3: Correspondences using SuperPoint+SuperGlue

Sets of outliers(red) and inliers(green) for given input images




(a) outliers-inliers for Image 1 and 2


(b) outliers-inliers for Image 2 and 3

Figure 4: outliers-inliers after applying RANSAC

Panoramic view obtained after stitching given input images



(a) Panoramic view after applying RANSAC



(b) Panoramic view after applying RANSAC and LM

Extracted correspondences using SuperPoint+SuperGlue between collected input images




Figure 6: Correspondences using SuperPoint+SuperGlue

Sets of outliers (red) and inliers (green) for collected input images



Figure 7: outliers-inliers after applying RANSAC

Panoramic view obtained after stitching collected input images




Step 4: Extra Credit (Pseudocode)

Algorithm 1: Levenberg-Marquardt for Homography Refinement

Input: Initial homography matrix H , set of inliers**Output:** Refined homography matrix H_{refined} **Initialize:** $\epsilon_1, \epsilon_2, \epsilon_3 \leftarrow 10^{-15}$, $\text{max_iterations} \leftarrow 100$ $\nu \leftarrow 2, \tau \leftarrow 0.5$ $h \leftarrow H$ flattened into a 9D vector $\text{iterations} \leftarrow 0$ **while** $\text{iterations} < \text{max_iterations}$ **do** **Step 1:** Compute residuals and Jacobian matrix $J \leftarrow \text{Jacobian}(H, \text{inliers})$ Compute cost function $e \leftarrow \text{cost_function}(h, \text{inliers})$ **Step 2:** Compute gradient and approximate Hessian $g \leftarrow -J^T \cdot e$ $A \leftarrow J^T \cdot J$ **Step 3:** Check for convergence based on gradient **if** $\|g\|_\infty < \epsilon_1$ **then** **break** **end** **Step 4:** Compute λ and update step $I \leftarrow$ identity matrix of size 9 $\lambda \leftarrow \tau \times \text{max}(\text{diag}(A))$ **while** *True* **do** Solve for Δh : $\Delta h \leftarrow (A + \lambda \cdot I)^{-1} \cdot g$ **if** $\|\Delta h\| \leq \epsilon_2 \cdot (\|h\| + \epsilon_2)$ **then** **break** **end** $h_{\text{new}} \leftarrow h + \Delta h$ $e_{\text{new}} \leftarrow \text{cost_function}(h_{\text{new}}, \text{inliers})$ Compute ρ : $\rho \leftarrow \frac{\|e\|^2 - \|e_{\text{new}}\|^2}{\Delta h^T \cdot (\lambda \cdot \Delta h + g)}$ **if** $\rho > 0$ **then** $h \leftarrow h_{\text{new}}$ $\lambda \leftarrow \lambda \cdot \max\left(\frac{1}{3}, 1 - (2\rho - 1)^3\right)$ $\nu \leftarrow 2$ **break** **end** **else** $\lambda \leftarrow \lambda \cdot \nu$ $\nu \leftarrow 2 \cdot \nu$ **end** **end** **if** $\|e\|^2 \leq \epsilon_3$ **then** **break** **end** $\text{iterations} \leftarrow \text{iterations} + 1$ **end****return** h_{new}

| Images | Inliers | M | Iterations | Best Inliers |
|--------|---------|-----|------------|--------------|
| 0 1 | 224 | 202 | 1086 | 208 |
| 1 2 | 318 | 286 | 5627 | 292 |
| 2 3 | 258 | 232 | 4012 | 251 |
| 3 4 | 128 | 115 | 4844 | 117 |

Table 1: Comparison of Inliers, Minimum Inliers to Accept(M), Iterations, and Best Inliers between Image Pairs




(a) Panoramic view after applying RANSAC and Own LM

| Images | RANSAC Cost | Own LM Cost |
|--------|-------------|-------------|
| 0 1 | 2443 | 1128 |
| 1 2 | 1530 | 753 |
| 2 3 | 1776 | 871 |
| 3 4 | 846 | 390 |

Table 2: Comparison of RANSAC and LM Cost between given input Image Pairs

| Images | Inliers | M | Iterations | Best Inliers |
|--------|---------|-----|------------|--------------|
| 0 1 | 768 | 692 | 529 | 721 |
| 1 2 | 875 | 788 | 513 | 789 |
| 2 3 | 687 | 619 | 770 | 626 |
| 3 4 | 871 | 784 | 37 | 802 |

Table 3: Comparison of Inliers, Minimum Inliers to Accept(M), Iterations, and Best Inliers between Collected Pairs



(a) Panoramic view after applying RANSAC and Own LM

| Images | RANSAC Cost | Own LM Cost |
|--------|-------------|-------------|
| 0 1 | 56765 | 27667 |
| 1 2 | 47086 | 23287 |
| 2 3 | 78634 | 37355 |
| 3 4 | 58491 | 29108 |

Table 4: Comparison of RANSAC and LM Cost between collected input Image Pairs

Programming

```

1  import numpy as np
2  import cv2
3  import matplotlib.pyplot as plt
4  from skimage.feature import corner_harris, corner_peaks
5  from scipy.optimize import least_squares
6  import random
7  import os
8  import math

9
10 #SIFT to get the keypoints and descriptors
11 def sift_feature_matching(img1, img2):
12     # Convert images to grayscale
13     gray1 = cv2.cvtColor(img1, cv2.COLOR_BGR2GRAY)
14     gray2 = cv2.cvtColor(img2, cv2.COLOR_BGR2GRAY)

15
16     # Initialize the SIFT detector
17     sift = cv2.SIFT_create()

18
19     # Detect keypoints and compute descriptors for both images
20     keypoints1, descriptors1 = sift.detectAndCompute(gray1, None)
21     keypoints2, descriptors2 = sift.detectAndCompute(gray2, None)

22
23     # Use BFMatcher with default params (L2 norm, as it's good for SIFT)
24     bf = cv2.BFMatcher()

25
26     # Match descriptors using KNN (k=2 for ratio test)
27     matches = bf.knnMatch(descriptors1, descriptors2, k=2)

28
29     # Apply the ratio test to retain good matches (Lowe's ratio test)
30     good_matches = []
31     for m, n in matches:
32         if m.distance < 0.75 * n.distance:
33             good_matches.append(m)

34
35     # Draw the matches
36     img_matches = cv2.drawMatches(img1, keypoints1, img2, keypoints2,
37         good_matches, None, flags=cv2.
38         DrawMatchesFlags_NOT_DRAW_SINGLE_POINTS)

39
40     #save the output image
41
42     cv2.imwrite('sift_sample.jpg', img_matches)

```

```
41     print("done!")
42
43     return img_matches, keypoints1, keypoints2, good_matches
44
45
46 #Code to find solution using Linear Least-Squares Ah = b
47 #How to construct A? Construct A from the correspondences
48
49 def homography_eq(p1 , p2):
50     A_eq = np.zeros((2,8))
51     A_eq[0] = [p1[0], p1[1], 1, 0, 0, 0, -p1[0]*p2[0], -p1[1]*p2[0]]
52     A_eq[1] = [0, 0, 0, p1[0], p1[1], 1, -p1[0]*p2[1], -p1[1]*p2[1]]
53     return A_eq
54
55 #create a function to find the homography matrix H using the
56 #matched_keypoints where the matched points can be more than 4
57 def find_homography(matched_keypoints):
58     A = []
59     b = []
60     for i in range(len(matched_keypoints)):
61         A.append(homography_eq(matched_keypoints[i][0], matched_keypoints[
62             i][1]))
63         b.append(matched_keypoints[i][1][0])
64         b.append(matched_keypoints[i][1][1])
65     A = np.array(A)
66     b = np.array(b)
67     #Change A to 2n x 8
68     A = A.reshape(-1, 8)
69     #Change b to column vector
70     b = b.reshape(-1, 1)
71     #Solve for h using pseudo-inverse
72     h = np.linalg.pinv(A) @ b
73     H = np.reshape(np.append(h,1),(3,3))
74     return H
75
76 #find_homography(matched_keypoints)
77
78 #Implement RANSAC to find the best homography matrix
79 #parameters are delta, epsilon, N(number of trials to conduct) and M(
80 #    Minimum number of inliers to accept a model)
81
82 def ransac_homography(matched_keypoints, delta= 12, epsilon= 0.1 , n= 20):
83     # Initialize the best homography matrix
84     best_H = None
85     best_inliers = 0
86     best_inliers_list = None
87
88     # Number of correspondences
89     n_total = len(matched_keypoints)
90     #print("n_total:", n_total)
91
92     # Probability that atleast one of the max_iters will be free of
93     # outliers is p
94     p = 0.99
```

```
91      # The minimum set of correspondences used to estimate the homography
92      # matrix be n
93
94      # The maximum number of trials to conduct
95      max_iters = int((np.log(1-p))/(np.log(1-(1-epsilon)**n)))
96      #max_iters = 30
97      #print("max_iters:", max_iters)
98      # The minimum number of inliers to accept a model
99      M = int((1-epsilon)*n_total)
100     print("Minimum inliers to accept:", M)
101
102     max_iters_all = (math.factorial(n_total)/(math.factorial(n)*math.
103     factorial(n_total-n)))
104
105     # RANSAC
106     #for i in range(max_iters):
107     interations = 0
108     while interations < max_iters_all:
109         # Randomly select n correspondences
110         random_indices = random.sample(range(n_total), n)
111         #print("random_indices:", random_indices)
112         random_points = [matched_keypoints[i] for i in random_indices]
113         #print("random_points:", random_points)
114
115         # Compute the homography matrix
116         H = find_homography(random_points)
117
118         # Compute the inliers
119         inliers = 0
120         inliers_list = []
121         for j in range(n_total):
122             # Compute the error
123             p1 = np.append(matched_keypoints[j][0], 1)
124             p2 = np.append(matched_keypoints[j][1], 1)
125             error = np.linalg.norm(p2 - H @ p1)
126             #print("error:", error)
127             # Check if the error is less than the delta
128             if error < delta:
129                 inliers += 1
130                 inliers_list.append(j)
131
132             # ***** check if the number of inliers are greater than
133             # minimum set M *****
134             if inliers > M:
135                 if inliers > best_inliers:
136                     best_inliers = inliers
137                     best_inliers_list = inliers_list
138                     print("interations:", interations)
139                     break
140             else:
141                 interations += 1
142                 continue
```

```
142
143     print("best_inliers:", best_inliers)
144     #print("best_inliers_list:", best_inliers_list)
145
146     best_points = [matched_keypoints[i] for i in best_inliers_list]
147     best_H = find_homography(best_points)
148     #print("best_H:", best_H)
149
150     return best_H, best_inliers_list
151
152 # Define the cost function (updated)
153 def cost_function(h, points):
154     h = np.reshape(h, (3, 3))
155     p1 = np.array([point[0] + [1] for point in points]) # Homogeneous
156         coordinates
157     p2 = np.array([point[1] for point in points]) # Actual 2D points
158
159     # Apply the homography transformation
160     f = (h @ p1.T).T
161     f = f[:, :2] / f[:, 2][:, np.newaxis] # Normalize homogeneous
162         coordinates to get (x, y)
163
164     # Compute error for both x and y coordinates (2N length)
165     error = (p2 - f).flatten() # Flatten to make it a 1D array of length
166         2N
167     return error
168
169
170 # Implement Jacobian from best_points and the homography matrix
171 def Jacobian(h, points):
172     h = np.reshape(h, (3, 3))
173     J = []
174
175     # Calculate the error without using a for loop
176     p1 = np.array([point[0] + [1] for point in points])
177     p2 = np.array([point[1] + [1] for point in points])
178     X = p2
179     f = (h @ p1.T).T
180
181     for i in range(len(points)):
182         x = p1[i][0]
183         y = p1[i][1]
184         J.append([x, y, 1, 0, 0, 0, -x*f[i][0], -y*f[i][0], -f[i][0]])
185         J.append([0, 0, 0, x, y, 1, -x*f[i][1], -y*f[i][1], -f[i][1]])
186
187     J = np.array(J)
188     return J
189
190
191 def Leven_Marqua(H, inliers):
192     epsilon1 = 1e-15
193     epsilon2 = 1e-15
194     epsilon3 = 1e-15
195     max_iterations = 100
196     nu = 2
197     tau = 0.5
198     h = H.flatten()
199     iterations = 0
200
201     while iterations < max_iterations:
```

```
193
194     # Calculate residuals and Jacobian
195     J          = Jacobian(H, inliers)
196     error      = cost_function(h, inliers)
197
198     # Calculate gradient and approximate Hessian
199     g          = J.T @ error
200     A          = J.T @ J
201
202     # Check for convergence
203     if np.linalg.norm(g, ord=np.inf) < epsilon1:
204         break
205
206     # Compute lambda and update step
207     I          = np.eye(9)
208     lambda_    = tau * np.max(np.diag(A))
209
210     while True:
211         # Solve the augmented normal equations
212         delta_h = np.linalg.solve(A + lambda_ * I, g)
213
214         if np.linalg.norm(delta_h) <= epsilon2 * (np.linalg.norm(h) +
215             epsilon2):
216             break
217
218         h_new = h + delta_h
219         error_new = cost_function(h_new, inliers)
220
221         rho = (np.linalg.norm(error)**2 - np.linalg.norm(error_new)**2) / (delta_h.T @ (lambda_ * delta_h + g))
222
223         if rho > 0:
224             h = h_new
225             lambda_ *= max(1/3, 1 - (2*rho - 1)**3)
226             nu = 2
227             break
228         else:
229             lambda_ *= nu
230             nu *= 2
231
232         if np.linalg.norm(error)**2 <= epsilon3:
233             break
234
235         iterations += 1
236
237     return h_new
238
239 # Write code to read consecutive images of form 1.jpg, 2.jpg,...,n.jpg
240 # Extract the keypoints and descriptors for each pair of images using SIFT
241 # function sift_feature_matching above
242 # Use the matched_keypoints to find the homography matrix using
243 # ransac_homography function
244 # Refine the homography using levenberg-marquardt optimization
245 # project all the views onto common reference frame (middle image) using
```

```
    the homography matrix
243
244 # Read the images
245 images = []
246 for i in range(1, 6):
247     img = cv2.imread(f'pics/{i}.jpg')
248     images.append(img)
249
250 # Initialize the feature extractor
251 feature_extractor = "sift"
252 #feature_extractor = "SuperPoint_SuperGlue"
253
254 Homographies = []
255 ransac_error_list = []
256 error_list = []
257 Leven_Marqua_error_list = []
258 for i in range(len(images)-1):
259     print("images:", i, i+1)
260     img1 = images[i]
261     img2 = images[i+1]
262     if feature_extractor == "sift":
263         img_matches, keypoints1, keypoints2, good_matches =
264             sift_feature_matching(img1, img2)
265
266         # Extract the good_matches for each pair of images
267         good_matches_list = []
268         # Initialize the list of points
269         matched_keypoints = []
270         for match in good_matches:
271             # Get the matching keypoints for each of the images
272             img1_idx = match.queryIdx
273             img2_idx = match.trainIdx
274
275             # x - columns
276             # y - rows
277             (x1, y1) = keypoints1[img1_idx].pt
278             (x2, y2) = keypoints2[img2_idx].pt
279
280             matched_keypoints.append([[x1, y1], [x2, y2]])
281     if feature_extractor == "SuperPoint_SuperGlue":
282         #Write code to read matched_keypoints from .npy file for
283         # SuperPoint+SuperGlue
284         matched_keypoints_np = np.load(f'pics/{i+1}_and_{i+2}'
285             '_matched_points.npy', allow_pickle=True)
286         print(matched_keypoints_np.size)
287
288         #edit matched_keypoints_np to match the format of
289         # matched_keypoints
290         # Initialize the list of points
291         matched_keypoints = []
292         for match in matched_keypoints_np:
293             # Get the matching keypoints for each of the images
294             x1, y1 = match[0]
295             x2, y2 = match[1]
```

```
292         matched_keypoints.append([[x1, y1], [x2, y2]])
293
294
295 #Apply RANSAC to find the best homography matrix
296 best_H, best_inliers_list = ransac_homography(matched_keypoints,
297     delta= 25, epsilon= 0.1 , n= 8)
298
299 best_inliers = [matched_keypoints[i] for i in best_inliers_list]
300 outliers = [matched_keypoints[i] for i in range(len(
301     matched_keypoints)) if i not in best_inliers_list]
302
303 #Draw the lines connecting inliers as green and outliers as red in
304 # the same image
305 img_matches = np.concatenate((img1, img2), axis=1)
306 for inlier in best_inliers:
307     cv2.line(img_matches, (int(inlier[0][0]), int(inlier[0][1])), (int
308         (inlier[1][0] + img1.shape[1]), int(inlier[1][1])), (0, 255, 0
309         , 1))
310 for outlier in outliers:
311     cv2.line(img_matches, (int(outlier[0][0]), int(outlier[0][1])), (int
312         (outlier[1][0] + img1.shape[1]), int(outlier[1][1])), (0, 0,
313         255), 1)
314 cv2.imwrite(f'pics/{i}_and_{i+1}_ransac.jpg', img_matches)
315
316 #RANSAC error
317 error = cost_function(best_H.flatten() , best_inliers)
318 #print("error: ", error)
319 ransac_error = np.sum(error**2)
320 #print("ransac_error: ",np.sum(error)/len(error))
321 ransac_error_list.append(ransac_error)
322 #Refine the homography using levenberg-marquardt optimization
323 res = least_squares(cost_function, best_H.flatten(), method= 'lm',
324     args=(best_inliers,))
325 H = np.reshape(res.x, (3, 3))
326 Homographies.append(H)
327 #print("H:", H)
328 #print("lm_error:", res.cost)
329 error_list.append(res.cost)
330 #Use defined LavMer
331 Leven_Marqua_H = Leven_Marqua(best_H, best_inliers)
332 Leven_Marqua_error = cost_function(Leven_Marqua_H, best_inliers)
333 Leven_Marqua_error_list.append(np.sum(Leven_Marqua_error**2))
334
335 print("error_list:", error_list)
336 print("ransac_error_list:", ransac_error_list)
337 print("Leven_Marqua_error_list:", Leven_Marqua_error_list)
338
339
340
341
342
343
344 # Calculate the index of the middle image
345 middle_image_index = len(images)//2
346
347 # Initialize an identity matrix to accumulate transformations for images
```

```
    after the middle image
338 transform_to_middle = np.eye(3)
339
340 # Loop through homographies of images after the middle image
341 # The goal is to adjust these homographies to the middle image's reference
342 # frame
343 for i in range(middle_image_index, len(Homographies)):
344     # Get the homography matrix
345     H = Homographies[i]
346     # Compute the transformation to the middle image's reference frame
347     transform_to_middle = transform_to_middle @ np.linalg.inv(H)
348     Homographies[i] = transform_to_middle
349
350 # Initialize an identity matrix to accumulate transformations for images
351 # before the middle image
352 transform_to_middle = np.eye(3)
353
354 # Loop through homographies of images before the middle image
355 # The goal is to adjust these homographies to the middle image's reference
356 # frame
357 for i in range(middle_image_index-1, -1, -1):
358     # Get the homography matrix
359     H = Homographies[i]
360     # Compute the transformation to the middle image's reference frame
361     transform_to_middle = transform_to_middle @ H
362     Homographies[i] = transform_to_middle
363
364 # Insert an identity matrix at the index of the middle image, as it doesn't
365 # need any transformation
366 # The middle image is the reference frame, so its homography is the
367 # identity matrix
368 Homographies = np.insert(Homographies, middle_image_index, np.eye(3), axis
369 =0)
370 print("Homographies:", Homographies.shape)
371
372 # Now, all homographies are adjusted relative to the middle image,
373 # ensuring all images are aligned to a common reference frame
374 #define a function to find the max and min of the transformed image
375 def maxmin(h, w, H):
376     # Define the four corners of the image
377     corners = np.array([[0, 0, 1], [w, 0, 1], [0, h, 1], [w, h, 1]])
378
379     # Apply the homography matrix to the corners
380     transformed_corners = H @ corners.T
381
382     # Normalize the transformed corners
383     transformed_corners = transformed_corners / transformed_corners[-1]
384
385     # Find the maximum and minimum x and y coordinates
386     max_x = np.max(transformed_corners[0])
387     min_x = np.min(transformed_corners[0])
388     max_y = np.max(transformed_corners[1])
389     min_y = np.min(transformed_corners[1])
```

```
384     return max_x, min_x, max_y, min_y
385
386 #define a create_pano(panorama_image, image, homography_matrix) function
387 # to create the panorama image
388 def create_pano(curr_img, new_img, H):
389     # Get dimensions of the current and new images
390     h_curr, w_curr, _ = curr_img.shape
391     h_new, w_new, _ = new_img.shape
392
393     # Invert the homography matrix
394     H_inv = np.linalg.inv(H)
395
396     # Create grid of pixel coordinates (x, y) for the current image
397     x_coords, y_coords = np.meshgrid(np.arange(w_curr), np.arange(h_curr))
398     # Flatten the grids and create homogeneous coordinates (x, y, 1)
399     homogeneous_coords = np.stack([x_coords.ravel(), y_coords.ravel(), np.
400         ones_like(x_coords).ravel()], axis=1)
401
402     # Apply the inverse homography transformation to the pixel coordinates
403     transformed_coords = H_inv @ homogeneous_coords.T
404     transformed_coords /= transformed_coords[2, :] # Normalize by the
405         last row
406
407     # Extract the transformed x and y coordinates
408     x_transformed = transformed_coords[0, :].astype(int)
409     y_transformed = transformed_coords[1, :].astype(int)
410
411     # Filter out points that are outside the bounds of the new image
412     valid_idx = (x_transformed >= 0) & (x_transformed < w_new) & (
413         y_transformed >= 0) & (y_transformed < h_new)
414
415     # Get the valid pixel locations in both the current and new images
416     x_curr_valid = homogeneous_coords[valid_idx, 0].astype(int)
417     y_curr_valid = homogeneous_coords[valid_idx, 1].astype(int)
418     x_new_valid = x_transformed[valid_idx]
419     y_new_valid = y_transformed[valid_idx]
420
421     # Map the valid pixels from new_img to curr_img
422     curr_img[y_curr_valid, x_curr_valid] = new_img[y_new_valid,
423         x_new_valid]
424
425     return curr_img
426
427
428 # Initialize translation parameters to calculate total width and height
429 # offsets for images
430 total_translation_x = 0
431 total_translation_y = 0
432
433 for i in range(len(images)//2):
434     image = images[i]
435     total_translation_y += image.shape[0] # Accumulate height for
436         vertical alignment
```

```
431     total_translation_x += image.shape[1] # Accumulate width for
432     horizontal alignment
433 print("total_translation_x:", total_translation_x)
434 print("total_translation_y:", total_translation_y)

435 # Initial transformation matrix set as an identity matrix with x-
436 # translation offset
437 initial_homography = np.eye(3)
438 #print("initial_homography:", initial_homography)
439 initial_homography[0, 2] = total_translation_x # Apply translation in x
440     direction
441 #print("initial_homography:", initial_homography)

442 # Calculate the final dimensions of the combined panorama image
443 final_height = 0
444 final_width = 0
445 for idx in range(len(images)):
446     image = images[idx]
447     h, w, _ = image.shape
448     # Get the max and min of the transformed image
449     max_x, min_x, max_y, min_y = maxmin(h, w, Homographies[idx])
450     # Update the final height and width
451     final_height = max(final_height, int(np.ceil(max_y) - np.floor(min_y)))
452     )
453     final_width += w
454 print("final_height:", final_height)
455 print("final_width:", final_width)

456 panorama_image = np.zeros((int(final_height), int(final_width), 3), dtype=
457     np.uint8)
458 #print("panorama_image:", panorama_image.shape)
459 # Apply transformations to each image and blend them onto the blank canvas
460 for idx in range(len(images)):
461     image = images[idx]

462     # Compute the final homography matrix by combining translation and
463     # homography for the image
464     homography_matrix = initial_homography @ Homographies[idx]
465     #print("homography_matrix:", homography_matrix)
466     print("panorama_image:", panorama_image.shape)
467     # 'create_panorama()' is a custom function that handles warping and
468     # blending of images using the homography
469     panorama_image = create_pano(panorama_image, image, homography_matrix)

470 # Convert the final panorama to RGB format (from BGR) for correct color
471 # representation
472 final_panorama_rgb = cv2.cvtColor(panorama_image, cv2.COLOR_BGR2RGB)
473 # Write the final stitched panorama image
474 cv2.imwrite('panorama.jpg', final_panorama_rgb)
```

Listing 1: Python Code