

# HW3

February 7, 2018

## CSE 252B: Computer Vision II, Winter 2018 – Assignment 3

Instructor: Ben Ochoa

Due: Wednesday, February 21, 2018, 11:59 PM

### Instructions

- Review the academic integrity and collaboration policies on the course website.
- This assignment must be completed individually.
- This assignment contains both math and programming problems.
- All solutions must be written in this notebook
- Math problems must be done in Markdown/LATEX. Remember to show work and describe your solution.
- Programming aspects of this assignment must be completed using Python in this notebook.
- This notebook contains skeleton code, which should not be modified (This is important for standardization to facilitate efficient grading).
- You may use python packages for basic linear algebra, but you may not use packages that directly solve the problem. Ask the instructor if in doubt.
- You must submit this notebook exported as a pdf. You must also submit this notebook as an .ipynb file.
- You must submit both files (.pdf and .ipynb) on Gradescope. You must mark each problem on Gradescope in the pdf.
- It is highly recommended that you begin working on this assignment early.

### Problem 1 (Programming): Estimation of the camera pose - Outlier rejection (20 points)

Download input data from the course website. The file `hw3_points3D.txt` contains the coordinates of 60 scene points in 3D (each line of the file gives the  $\tilde{X}_i$ ,  $\tilde{Y}_i$ , and  $\tilde{Z}_i$  inhomogeneous coordinates of a point). The file `hw3_points2D.txt` contains the coordinates of the 60 corresponding image points in 2D (each line of the file gives the  $\tilde{x}_i$  and  $\tilde{y}_i$  inhomogeneous coordinates of a point). The corresponding 3D scene and 2D image points contain both inlier and outlier correspondences. For the inlier correspondences, the scene points have been randomly generated and projected to image points under a camera projection matrix (i.e.,  $x_i = PX_i$ ), then noise has been added to the image point coordinates.

The camera calibration matrix was calculated for a  $1280 \times 720$  sensor and  $45^\circ$  horizontal field of view lens. The resulting camera calibration matrix is given by

$$\mathbf{K} = \begin{bmatrix} 1545.0966799187809 & 0 & 639.5 \\ 0 & 1545.0966799187809 & 359.5 \\ 0 & 0 & 1 \end{bmatrix}$$

For each image point  $\mathbf{x} = (x, y, w)^\top = (\tilde{x}, \tilde{y}, 1)^\top$ , calculate the point in normalized coordinates  $\hat{\mathbf{x}} = \mathbf{K}^{-1}\mathbf{x}$ .

Determine the set of inlier point correspondences using the M-estimator Sample Consensus (MSAC) algorithm, where the maximum number of attempts to find a consensus set is determined adaptively. For each trial, use the 3-point algorithm of Finsterwalder (as described in the paper by Haralick et al.) to estimate the camera pose (i.e., the rotation  $\mathbf{R}$  and translation  $\mathbf{t}$  from the world coordinate frame to the camera coordinate frame), resulting in up to 4 solutions, and calculate the error and cost for each solution. Note that the 3-point algorithm requires the 2D points in normalized coordinates, not in image coordinates. Calculate the projection error, which is the (squared) distance between projected points (the points in 3D projected under the normalized camera projection matrix  $\hat{\mathbf{P}} = [\mathbf{R}|\mathbf{t}]$ ) and the measured points in normalized coordinates (hint: the error tolerance is simpler to calculate in image coordinates using  $\mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$  than in normalized coordinates using  $\hat{\mathbf{P}} = [\mathbf{R}|\mathbf{t}]$ ).

hint: this problem has codimension 2).

```
In [1]: import numpy as np
import matplotlib.pyplot as plt

x=np.loadtxt('hw3_points2D.txt').T
X=np.loadtxt('hw3_points3D.txt').T
print('x is', x.shape)
print('X is', X.shape)

K = np.array([[1545.0966799187809, 0, 639.5],
              [0, 1545.0966799187809, 359.5],
              [0, 0, 1]])

print('K =')
print(K)

def toHomo(x):
    # converts points from inhomogeneous to homogeneous coordinates
    return np.vstack((x,np.ones((1,x.shape[1]))))
def fromHomo(x):
    # converts points from homogeneous to inhomogeneous coordinates
    return x[:-1,:]/x[-1,:]

def proj(P,X):
    # projects 3d points X to 2d using projection matrix P
    return fromHomo(np.matmul(P,toHomo(X)))

def displayResults(R, t, K, x, X, inliers, iters, cost):
    print ('R = ')
    print (R)
```

```

print ('t = ')
print (t)
print ('cost = ', cost)
print ('iterations = ', iters)

x_proj = proj(np.matmul(K,np.hstack((R,t))),X)
plt.plot(x[0,:], x[1,:],'.k')
plt.plot(x_proj[0,:], x_proj[1,:],'.r')
for i in range(x.shape[1]):
    if i in inliers:
        line_style = '-'
    else:
        line_style = ':'
    plt.plot([x[0,i], x_proj[0,i]], [x[1,i], x_proj[1,i]], line_style+'r')

plt.show()

```

x is (2, 60)

X is (3, 60)

K =

```

[[ 1.54509668e+03  0.00000000e+00  6.39500000e+02]
 [ 0.00000000e+00  1.54509668e+03  3.59500000e+02]
 [ 0.00000000e+00  0.00000000e+00  1.00000000e+00]]

```

In [2]: def MSAC(x, X, K, max\_iters):

```

    """your code here"""

```

```

    R = np.eye(3) # estimated rotation matrix

```

```

    t = np.array([[0,0,100]]).T # estimated translation

```

```

    inliers = np.sort(np.random.choice(x.shape[1],35,replace=False)) # indices of inliers

```

```

    cost = np.inf # MSAC cost

```

```

    iters = 0 # number of MSAC iterations executed

```

```

    return R, t, inliers, cost, iters

```

```

# MSAC hyperparameters (add any additional hyperparameters necessary here. For example

```

```

# You should pass these hyperparameters as additional parameters to MSAC(...))

```

```

max_iters=1

```

```

R_MSAC, t_MSAC, inliers, cost_MSAC, iters_MSAC = MSAC(x, X, K, max_iters)

```

```

displayResults(R_MSAC, t_MSAC, K, x, X, inliers, iters_MSAC, cost_MSAC)

```

```

print('inliers: ', inliers)

```

```

print('inlier count: ', len(inliers))

```

R =

```

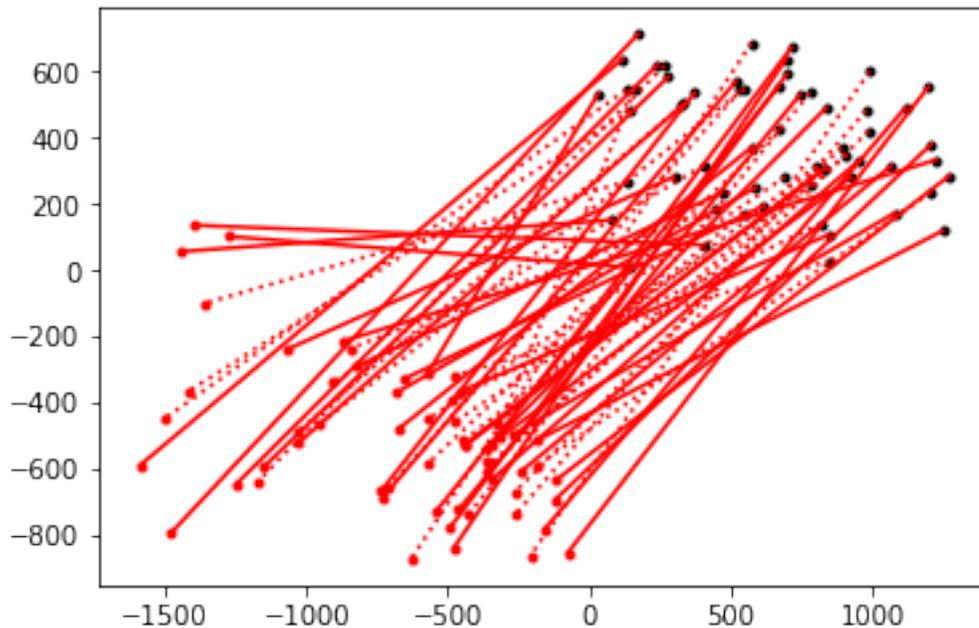
[[ 1.  0.  0.]
 [ 0.  1.  0.]
 [ 0.  0.  1.]]

```

```

t =
[[ 0]
 [ 0]
 [100]]
cost = inf
iterations = 0

```



```

inliers: [ 0  3  6  7  8 13 14 17 19 20 22 24 25 27 28 29 30 31 32 33 36 37 39 43 44
 45 46 48 50 51 55 56 57 58 59]
inlier count: 35

```

## Problem 2 (Programming): Estimation of the camera pose - Linear Estimate (30 points)

Estimate the normalized camera projection matrix  $\hat{P}_{\text{linear}} = [\mathbf{R}_{\text{linear}} | \mathbf{t}_{\text{linear}}]$  from the resulting set of inlier correspondences using the linear estimation method (based on the EPnP method) described in lecture.

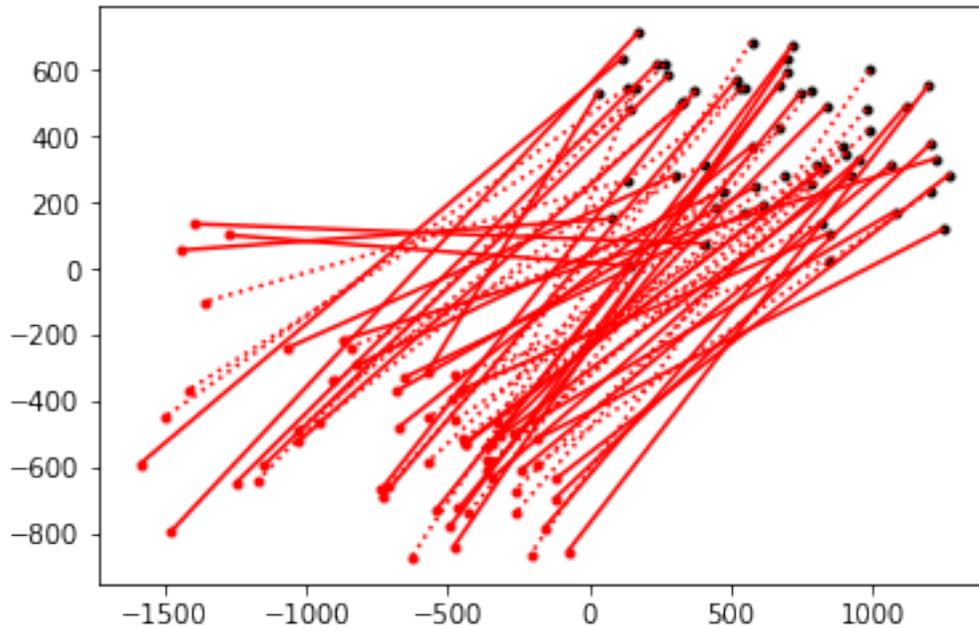
```

In [3]: def EPnP(x, X, K):
        """your code here"""
        R = np.eye(3) # estimated rotation matrix
        t = np.array([[0,0,100]]).T # estimated translation
        cost = np.inf # linear cost
        return R, t, cost

```

```
R_EPnP, t_EPnP, cost_EPnP = EPnP(x[:,inliers], X[:,inliers], K)
displayResults(R_EPnP, t_EPnP, K, x, X, inliers, 1, cost_EPnP)
```

```
R =
[[ 1.  0.  0.]
 [ 0.  1.  0.]
 [ 0.  0.  1.]]
t =
[[ 0]
 [ 0]
 [100]]
cost = inf
iterations = 1
```



### Problem 3 (Programming): Estimation of the camera pose - Nonlinear Estimate (30 points)

Use  $R_{\text{linear}}$  and  $t_{\text{linear}}$  as an initial estimate to an iterative estimation method, specifically the Levenberg-Marquardt algorithm, to determine the Maximum Likelihood estimate of the camera pose that minimizes the projection error under the normalized camera projection matrix  $\hat{P} = [R|t]$ . You must parameterize the camera rotation using the angle-axis representation  $\omega$  (where  $[\omega]_{\times} = \ln R$ ) of a 3D rotation, which is a 3-vector.

```
In [4]: def R2w(R):
        # given a rotation matrix R return the angle-axis representation
```

```

        """your code here"""
    return np.array([[0,0,0]])
def w2R(w):
    # given the angle-axis representation w return the rotation matrix
    """your code here"""
    return np.eye(3)

def LMstep(w, t, x, X, K, l, v):
    # inputs:
    # w current estimate of rotation in angle-axis representation
    # t current estimate of t
    # x 2D points
    # X 3D points
    # K camera calibration matrix
    # l LM lambda parameter
    # v LM change of lambda parameter
    # output:
    # R updated by a single LM step
    # t updated by a single LM step
    # cost
    # l accepted lambda parameter

    """your code here"""
    cost = np.inf
    return w, t, cost, l

# use linear estimate as an initialization for LM
w_LM = R2w(R_EPnP)
t_LM = t_EPnP

# LM hyperparameters
l=.001
v=10
max_iters=10

# LM optimization loop
for i in range(max_iters):
    w_LM, t_LM, cost_LM, l = LMstep(w_LM, t_LM, x[:,inliers], X[:,inliers], K, l, v)
    print ('iter %d cost %f'%(i+1, cost_LM))
    R_LM = w2R(w_LM)

displayResults(R_LM, t_LM, K, x, X, inliers, i+1, cost_LM)

iter 1 cost inf
iter 2 cost inf
iter 3 cost inf
iter 4 cost inf

```

```
iter 5 cost inf
iter 6 cost inf
iter 7 cost inf
iter 8 cost inf
iter 9 cost inf
iter 10 cost inf
R =
[[ 1.  0.  0.]
 [ 0.  1.  0.]
 [ 0.  0.  1.]]
t =
[[ 0]
 [ 0]
 [100]]
cost = inf
iterations = 10
```

