

HW2

October 21, 2021

1 CSE 252A Computer Vision I Fall 2021 - Assignment 2

1.1 Instructor: Ben Ochoa

- Assignment Published On: **Wed, Oct. 20, 2021.**
- Due On: **Wed, Nov. 3, 2021 11:59 PM (Pacific Time).**

1.2 Instructions

Please answer the questions below using Python in the attached Jupyter notebook and follow the guidelines below:

- This assignment must be completed **individually**. For more details, please follow the Academic Integrity Policy and Collaboration Policy on [Canvas](#).
- All the solutions must be written in this Jupyter notebook.
- After finishing the assignment in the notebook, please export the notebook as a PDF and submit both the notebook and the PDF (i.e. the `.ipynb` and the `.pdf` files) on Gradescope.
- You may use basic algebra packages (e.g. NumPy, SciPy, etc) but you are not allowed to use the packages that directly solve the problems. Feel free to ask the instructor and the teaching assistants if you are unsure about the packages to use.
- It is highly recommended that you begin working on this assignment early.

Late Policy: Assignments submitted late will receive a 15% grade reduction for each 12 hours late (i.e., 30% per day). Assignments will not be accepted 72 hours after the due date. If you require an extension (for personal reasons only) to a due date, you must request one as far in advance as possible. Extensions requested close to or after the due date will only be granted for clear emergencies or clearly unforeseeable circumstances.

1.3 Problem 1 Image filtering [15 pts]

1.3.1 Problem 1.1 Implementing Convolution[5 pts]

In this problem, you will implement the convolution filtering operation in NumPy.

As shown in the lecture, a convolution can be considered as a sliding window that computes a sum of the pixel values weighted by the flipped kernel. Your version will i) zero-pad an image, ii) flip the kernel horizontally and vertically, and iii) compute a weighted sum of the neighborhood at each pixel.

Problem 1.1.1 [1 pts] First you will want to implement the `zero_pad` function.

```
In [ ]: import numpy as np
        from time import time
        from skimage import io
        %matplotlib inline
        import matplotlib.pyplot as plt

In [ ]: def zero_pad(image, pad_top, pad_down, pad_left, pad_right):
        """ Zero-pad an image.

        Ex: a 1x1 image [[1]] with pad_top = 1, pad_down = 1, pad_left = 2, pad_right = 2

        [[0, 0, 0, 0, 0],
         [0, 0, 1, 0, 0],
         [0, 0, 0, 0, 0]]          of shape (3, 5)

        Args:
        image: numpy array of shape (H, W)
        pad_left: width of the zero padding to the left of the first column
        pad_right: width of the zero padding to the right of the last column
        pad_top: height of the zero padding above the first row
        pad_down: height of the zero padding below the last row

        Returns:
        out: numpy array of shape (H + pad_top + pad_down, W + pad_left + pad_right)
        """
        =====
        YOUR CODE HERE
        ===== """

        return out

# Open image as grayscale
img = io.imread('dog.jpg', as_gray=True)

# Show image
plt.imshow(img)
plt.axis('off')
plt.show()

pad_width = 20 # width of the padding on the left and right
pad_height = 40 # height of the padding on the top and bottom

padded_img = zero_pad(img, pad_height, pad_height, pad_width, pad_width)

# Plot your padded dog
```

```

plt.subplot(1,2,1)
plt.imshow(padded_img)
plt.title('Padded dog')
plt.axis('off')

# Plot what you should get
solution_img = io.imread('padded_dog.jpg', as_gray=True)
plt.subplot(1,2,2)
plt.imshow(solution_img)
plt.title('What you should get')
plt.axis('off')

plt.show()

```

Problem 1.1.2 [2 pts] Now implement the function `conv`, using at most 2 loops. This function should take an image f and a kernel h as inputs and output the convolved image ($f * h$) that has the same shape as the input image (use zero padding to accomplish this). Depending on the computer, your implementation should take around a second or less to run.

```

In [ ]: def conv(image, kernel):
        """ An efficient implementation of a convolution filter.

        This function uses element-wise multiplication and np.sum()
        to efficiently compute a weighted sum of the neighborhood at each
        pixel.

        Hints:
        - Use the zero_pad function you implemented above
        - You should need at most two nested for-loops
        - You may find np.flip() and np.sum() useful
        - You need to handle both odd and even kernel size

        Args:
        image: numpy array of shape (Hi, Wi)
        kernel: numpy array of shape (Hk, Wk)

        Returns:
        out: numpy array of shape (Hi, Wi)
        """
        Hi, Wi = image.shape
        Hk, Wk = kernel.shape
        out = np.zeros((Hi, Wi))

        """ =====
        YOUR CODE HERE
        ===== """

```

```

    return out

# Simple convolution kernel.
# Feel free to change the kernel and to see different outputs.
kernel = np.array(
    [
        [1,0,-1],
        [2,0,-2],
        [1,0,-1]
    ])

t1 = time()
out = conv(img, kernel)
t2 = time()
print("took %f seconds." % (t2 - t1))

# Plot original image
plt.subplot(2,2,1)
plt.imshow(img)
plt.title('Original')
plt.axis('off')

# Plot your convolved image
plt.subplot(2,2,3)
plt.imshow(out)

plt.title('Convolution')
plt.axis('off')

# Plot what you should get
solution_img = io.imread('convolved_dog.jpg', as_gray=True)
plt.subplot(2,2,4)
plt.imshow(solution_img)
plt.title('What you should get')
plt.axis('off')

plt.show()

```

Problem 1.1.3 [1 pt] Now let's filter some images! Here, you will apply the convolution function that you just implemented in order to bring about some interesting image effects. More specifically, we will use convolution to blur and sharpen our images.

First we will apply convolution for image blurring. To accomplish this, convolve the dog image with a 7x7 Gaussian filter for $\sigma = 1.0$. You can use the included function to obtain the Gaussian kernel.

```

In [ ]: def gaussian2d(filter_size=7, sig=1.0):
        """
        Creates 2D Gaussian kernel with side length `filter_size` and a sigma of `sig`.

```

```

    """
    ax = np.arange(-filter_size // 2 + 1., filter_size // 2 + 1.)
    xx, yy = np.meshgrid(ax, ax)
    kernel = np.exp(-0.5 * (np.square(xx) + np.square(yy)) / np.square(sig))
    return kernel / np.sum(kernel)

def blur_image(img):
    """Blur the image by convolving with a Gaussian filter."""
    blurred_img = np.zeros_like(img)
    """ =====
    YOUR CODE HERE
    ===== """

    return blurred_img

# Plot original image
plt.subplot(2,2,1)
plt.imshow(img)
plt.title('Original')
plt.axis('off')

# Plot blurred image
plt.subplot(2,2,2)
plt.imshow(blur_image(img))
plt.title('Blurred')
plt.axis('off')

plt.show()

```

Problem 1.1.4 [1 pt] Next, we will use convolution to sharpen the images. Convolve the image with the following filter to produce a sharpened result. For convenience, we have defined the filter for you:

```

In [ ]: sharpening_kernel = np.array([
    [1, 4, 6, 4, 1],
    [4, 16, 24, 16, 4],
    [6, 24, -476, 24, 6],
    [4, 16, 24, 16, 4],
    [1, 4, 6, 4, 1],
]) * -1.0 / 256.0

```

```

In [ ]: def sharpen_image(img):
    """Sharpen the image by convolving with a sharpening filter."""
    sharpened_img = np.zeros_like(img)
    """ =====
    YOUR CODE HERE
    ===== """

```

```

    return sharpened_img

# Plot original image
plt.subplot(2,2,1)
plt.imshow(img, vmin=0.0, vmax=1.0)
plt.title('Original')
plt.axis('off')

# Plot sharpened image
plt.subplot(2,2,2)
plt.imshow(sharpen_image(img), vmin=0.0, vmax=1.0)
plt.title('Sharpened')
plt.axis('off')

plt.show()

```

1.3.2 Problem 1.2: Convolution Theory [5 pts]

Problem 1.2.1 [2 pts] Consider (1) smoothing an image with a 3x3 box filter and then computing the derivative in the x-direction. Also consider (2) computing the derivative first, then smoothing. What is a single convolution kernel that will simultaneously implement both (1) and (2)? Try to give a brief justification for how you arrived at the kernel.

Use the x-derivative filter

$$[-1/2, 0, 1/2]$$

for this problem.

Problem 1.2.2 [3 pts] Certain 2D filters can be expressed as a convolution of two 1D filters. Such filters are called separable filters. Give an example of a 3x3 separable filter and compare the number of arithmetic operations it takes to convolve an $n \times n$ image using that filter before and after separation. Count both, the number of multiplication and addition operations in each case.

Assume that the convolution of the image and filter is performed in "valid" mode, i.e., the image is not padded before convolution.

1.3.3 Problem 1.3 Template Matching [5 pts]

Suppose that you are a clerk at a grocery store. One of your responsibilities is to check the shelves periodically and stock them up whenever there are sold-out items. You got tired of this laborious task and decided to build a computer vision system that keeps track of the items on the shelf.

Luckily, you have learned in CSE 252A (or are learning right now) that convolution can be used for template matching: a template g is multiplied with regions of a larger image f to measure how similar each region is to the template. Note that you will want to flip the filter before giving it to your convolution function, so that it is overall not flipped when making comparisons. You will also want to subtract off the mean value of the image or template (whichever you choose, subtract the same value from both the image and template) so that our solution is not biased toward higher-intensity (white) regions.

The template of a product (template.jpg) and the image of the shelf (shelf.jpg) is provided. We will use convolution to find the product in the shelf.

```
In [ ]: # Load template and image in grayscale
img = io.imread('shelf.jpg')
img_gray = io.imread('shelf.jpg', as_gray=True)
temp = io.imread('template.jpg')
temp_gray = io.imread('template.jpg', as_gray=True)

# Perform a convolution between the image and the template and store
# the result in the out variable
""" =====
YOUR CODE HERE
===== """

# Display product template
plt.figure(figsize=(20,16))
plt.subplot(3, 1, 1)
plt.imshow(temp_gray, cmap="gray")
plt.title('Template')
plt.axis('off')

# Display convolution output
plt.subplot(3, 1, 2)
plt.imshow(out, cmap="gray")
plt.title('Convolution output (white means more correlated)')
plt.axis('off')

# Display image
plt.subplot(3, 1, 3)
plt.imshow(img, cmap="gray")
plt.title('Result (blue marker on the detected location)')
plt.axis('off')

# Draw marker at detected location
plt.plot(x, y, 'bx', ms=40, mew=10)
plt.show()
```

1.4 Problem 2: Edge detection [21 pts]

In this problem, you will write a function to perform Canny edge detection. The following steps need to be implemented.

1.4.1 Problem 2.1 Smoothing [1 pt]

First, we need to smooth the images in order to prevent noise from being considered as edges. For this assignment, use a 9x9 Gaussian kernel filter with $\sigma = 1.4$ to smooth the images.

```

In [ ]: import numpy as np
        from skimage import io
        import matplotlib.pyplot as plt
        import matplotlib.cm as cm
        from scipy.signal import convolve
        %matplotlib inline

        import matplotlib
        matplotlib.rcParams['figure.figsize'] = [5, 5]

In [ ]: def gaussian2d(filter_size=None, sig=None):
        """Creates a 2D Gaussian kernel with
        side length `filter_size` and a sigma of `sig`."""
        ax = np.arange(-filter_size // 2 + 1., filter_size // 2 + 1.)
        xx, yy = np.meshgrid(ax, ax)
        kernel = np.exp(-0.5 * (np.square(xx) + np.square(yy)) / np.square(sig))
        return kernel / np.sum(kernel)

In [ ]: def smooth(image):
        """ =====
        YOUR CODE HERE
        ===== """

In [ ]: # Load image in grayscale
        image = io.imread('sio_pier.jpg', as_gray=True)
        assert len(image.shape) == 2, 'image should be grayscale; check your Python/skimage version'
        smoothed = smooth(image)
        print('Original:')
        plt.imshow(image, cmap=cm.gray)
        plt.show()

        print('Smoothed:')
        plt.imshow(smoothed, cmap=cm.gray)
        plt.show()

```

1.4.2 Problem 2.2 Gradient Computation [5 pts]

After you have finished smoothing, find the image gradient in the horizontal and vertical directions. Compute the gradient magnitude image as $|G| = \sqrt{G_x^2 + G_y^2}$. The edge direction for each pixel is given by $G_\theta = \tan^{-1}\left(\frac{G_y}{G_x}\right)$.

```

In [ ]: def gradient(image):
        """ =====
        YOUR CODE HERE
        ===== """

        return g_mag, g_theta

```



```
In [ ]: g_mag, g_theta = gradient(smoothed)
        print('Gradient magnitude:')
        plt.imshow(g_mag, cmap=cm.gray)
        plt.show()
```

1.4.3 Problem 2.3 Non-Maximum Suppression [7 pts]

We would like our edges to be sharp, unlike the ones in the gradient image. Use non-maximum suppression to preserve all local maxima and discard the rest. You can use the following method to do so:

- For each pixel in the gradient magnitude image:
 - Round the gradient direction θ to the nearest multiple of 45° (which we will refer to as ve).
 - Compare the edge strength at the current pixel to the pixels along the $+ve$ and $-ve$ gradient direction in the 8-connected neighborhood.
 - If the pixel does not have a larger value than both of its two neighbors in the $+ve$ and $-ve$ gradient directions, suppress the pixel's value (set it to 0). By following this process, we preserve the values of only those pixels which have maximum gradient magnitudes in the neighborhood along the $+ve$ and $-ve$ gradient directions.
- Return the result as the NMS response.

```
In [ ]: def nms(g_mag, g_theta):
        """ =====
        YOUR CODE HERE
        ===== """

        return nms_response
```

```
In [ ]: nms_image = nms(g_mag, g_theta)
        print('NMS:')
        plt.imshow(nms_image, cmap=cm.gray)
        plt.show()
```

1.4.4 Problem 2.4 Hysteresis Thresholding [8 pts]

Choose suitable values of thresholds and use the thresholding approach described in lecture 6. This will remove the edges caused by noise and color variations.

- Define two thresholds t_{min} and t_{max} .
- If $nms > t_{max}$, then we select that pixel as an edge.
- If $nms < t_{min}$, we reject that pixel.
- If $t_{min} < nms < t_{max}$, we select the pixel only if there is a path from to another pixel with $nms > t_{max}$. (Hint: Think of all pixels with $nms > t_{max}$ as starting points and run BFS/DFS from these starting points).

- The choice of value of low and high thresholds depends on the range of values in the gradient magnitude image. You can start by setting the high threshold to some percentage of the max value in the gradient magnitude image, e.g. `thres_high = 0.2 * image.max()`, and the low threshold to some percentage of the high threshold, e.g. `thres_low = 0.85 * thres_high`. And then you can tune those values however you want.

```
In [ ]: def hysteresis_threshold(image, g_theta, use_g_theta=False):
        """ =====
        YOUR CODE HERE
        ===== """

        return result
```

```
In [ ]: thresholded = hysteresis_threshold(nms_image, g_theta)
        print('Thresholded:')
        plt.imshow(thresholded, cmap=cm.gray)
        plt.show()
```

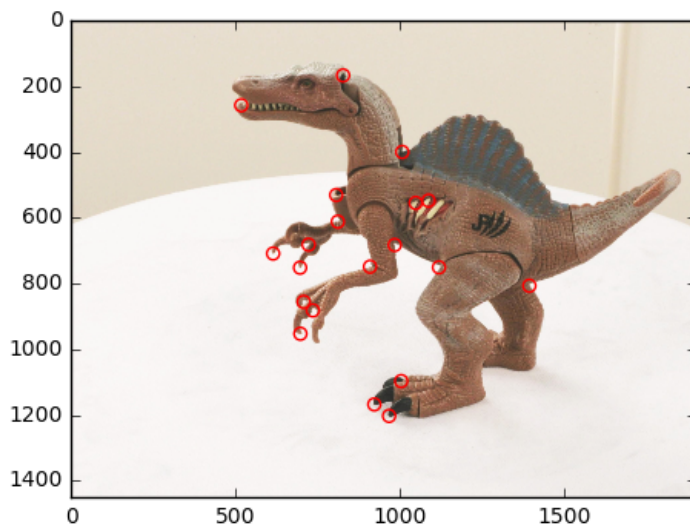
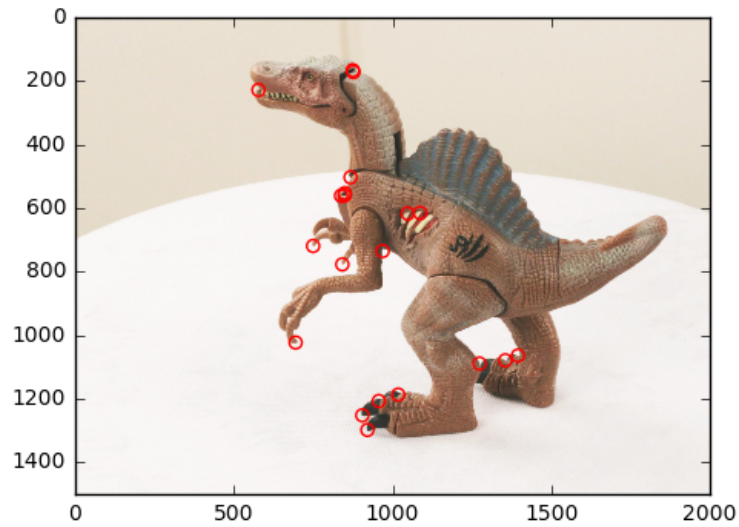
1.5 Problem 3 Corner detection [13 pts]

1.5.1 Problem 3.1 [12 pts]

In this problem, we are going to build a corner detector. This should be done according to the lecture slides. You should fill in the function `corner_detect` below, which takes as input `image`, `nCorners`, `smoothSTD`, `windowSize` -- where `smoothSTD` is the standard deviation of the smoothing kernel and `windowSize` is the window size for corner detector and non-maximum suppression. In the lecture, the corner detector was implemented using a hard threshold. Do not do that; instead, return the `nCorners` strongest corners after non-maximum suppression. This way you can control exactly how many corners are returned. Run your code on all four images (with `nCorners = 20`) and display outputs as shown below. You may find `scipy.ndimage.filters.gaussian_filter` helpful for smoothing.

In this problem, try the following different standard deviation (σ) parameters for the Gaussian smoothing kernel: 0.5, 1, 2 and 4. For a particular σ , you should take the kernel size to be $6 \times \sigma$ (add 1 if the kernel size is even). So for example if $\sigma = 2$, corner detection kernel size should be 13. This should be followed throughout all of the experiments in this assignment.

There will be a total of 16 images as outputs: 4 choices of `smoothSTD` \times 2 matrix images \times 2 warrior images.



```
In [ ]: import numpy as np
import matplotlib.pyplot as plt
from scipy.ndimage.filters import gaussian_filter
import imageio
from scipy.signal import convolve

def rgb2gray(rgb):
    """ Convert rgb image to grayscale.
    """
    return np.dot(rgb[...,:3], [0.299, 0.587, 0.114])
```

```
In [ ]: def corner_detect(image, nCorners, smoothSTD, windowSize):
    """Detect corners on a given image.
```

Args:

image: Given a grayscale image on which to detect corners.

nCorners: Total number of corners to be extracted.

smoothSTD: Standard deviation of the Gaussian smoothing kernel.

windowSize: Window size for corner detector and non-maximum suppression.

Returns:

*Detected corners (in image coordinate) in a numpy array (n*2).*

```
"""
```

```
""" =====
```

```
YOUR CODE HERE
```

```
===== """
```

```
return corners
```

```
In [ ]: def show_corners_result(imgs, corners):
fig = plt.figure(figsize=(8, 8))
ax1 = fig.add_subplot(221)
ax1.imshow(imgs[0], cmap='gray')
ax1.scatter(corners[0][:, 0], corners[0][:, 1], s=35, edgecolors='r', facecolors='r')

ax2 = fig.add_subplot(222)
ax2.imshow(imgs[1], cmap='gray')
ax2.scatter(corners[1][:, 0], corners[1][:, 1], s=35, edgecolors='r', facecolors='r')
plt.show()

for smoothSTD in (0.5, 1, 2, 4):
    windowSize = int(smoothSTD * 6)
    if windowSize % 2 == 0:
        windowSize += 1

    print('smooth stdev: %r' % smoothSTD)
    print('window size: %r' % windowSize)

nCorners = 20

# read images and detect corners on images

imgs_din = []
crns_din = []
imgs_mat = []
crns_mat = []
imgs_war = []
crns_war = []
```

```

for i in range(2):
    img_din = imageio.imread('dino/dino' + str(i) + '.png')
    imgs_din.append(rgb2gray(img_din))
    # downsize your image in case corner_detect runs slow in test
    # imgs_din.append(rgb2gray(img_din)[:2, :2])
    crns_din.append(corner_detect(imgs_din[i], nCorners, smoothSTD, windowSize))

    img_mat = imageio.imread('matrix/matrix' + str(i) + '.png')
    imgs_mat.append(rgb2gray(img_mat))
    # downsize your image in case corner_detect runs slow in test
    # imgs_mat.append(rgb2gray(img_mat)[:2, :2])
    crns_mat.append(corner_detect(imgs_mat[i], nCorners, smoothSTD, windowSize))

    img_war = imageio.imread('warrior/warrior' + str(i) + '.png')
    imgs_war.append(rgb2gray(img_war))
    # downsize your image in case corner_detect runs slow in test
    # imgs_war.append(rgb2gray(img_war)[:2, :2])
    crns_war.append(corner_detect(imgs_war[i], nCorners, smoothSTD, windowSize))

show_corners_result(imgs_din, crns_din)
show_corners_result(imgs_mat, crns_mat)
show_corners_result(imgs_war, crns_war)

```

1.5.2 Problem 3.2 [1 pts]

Comment on your results and observations. You don't need to comment per output; just discuss any trends you see for the detected corners as you change the windowSize and increase the smoothing w.r.t the two pairs of images (warrior and matrix). Also discuss whether you are able to find corresponding corners for the pairs of images.

1.6 Problem 4 Epipolar rectification and feature matching [43 pts]

1.6.1 4.1 Epipolar rectification [22 pts]

In this problem, we are going to perform epipolar rectification. Given calibrated stereo cameras (i.e., calibration matrices K_1 and K_2 , camera rotation matrices R_1 and R_2 , camera translation vectors t_1 and t_2), you are expected to determine the rotation matrix R and calibration matrix K of the virtual cameras. Your goal is to complete the function *epipolarRectification*, which determines the calibration matrix and rotation matrix of both cameras, the translation vector of each of the cameras, and matching planar transformations that epipolar rectify the two images acquired by the cameras. The destination virtual cameras have the same centers as the source real cameras.

4.1.1 Camera translation matrices and Projective Transformation matrices [6 pts] To calculate the camera translation from cameras with the same camera center, you will have to complete the *cameraTranslation* first. Another function you need to complete is *calcProjectiveTransformation*, which calculates the planar projective transformation from cameras with the same camera center. The camera calibration matrix (same for both cameras) will be calculated by *calcDestinateK*. This is provided for you. To get the rotation matrix R of the virtual camera, we usually interpolate

halfway between the two 3D rotations embodied by R_1 and R_2 . For simplicity, this will be also given to you.

```
In [ ]: from imageio import imread
import matplotlib.pyplot as plt
import numpy as np
import math
import pickle
from math import floor, ceil

In [ ]: def cameraTranslation(R1, t1, R2):
    '''
    Calculate the camera translation from cameras with the same camera center.

    Args:
    R1: The rotation matrix of the first camera.
    t1: The translation vector of the first camera.
    R2: The rotation matrix of the second camera.

    Returns:
    The translation vector of the second camera.
    '''
    """ =====
    YOUR CODE HERE
    ===== """

In [ ]: def calcProjectiveTransformation(K1, R1, K2, R2):
    '''
    Calculates the planar projective transformation from cameras with the same camera center.
    This function determines the planar projective transformation from the image of a point in the
    first camera to the image of the same point in the second camera.

    Args:
    K1: The calibration matrix of the first camera.
    R1: The rotation matrix of the first camera.
    K2: The calibration matrix of the second camera.
    R2: The rotation matrix of the second camera.

    Returns:
    The transformation matrix.
    '''
    """ =====
    YOUR CODE HERE
    ===== """

In [ ]: def calcDestinateK(srcK1, srcK2):
    '''
    Camera calibration matrix (same for both cameras)'''
```

```

'''
alpha = (srcK1[0][0] + srcK2[0][0] + srcK1[1][1] + srcK2[1][1]) // 4
x0 = (srcK1[0][2] + srcK2[0][2]) // 2
y0 = (srcK1[1][2] + srcK2[1][2]) // 2
dstK = np.zeros((3, 3))
dstK[0][0] = alpha
dstK[0][2] = x0
dstK[1][1] = alpha
dstK[1][2] = y0
dstK[2][2] = 1
return dstK

```

```

In [ ]: def epipolarRectification(srcK1, srcR1, src_t1,
                                srcK2, srcR2, src_t2, dstR):

```

```

'''

```

Given two calibrated cameras, this function determines the calibration matrix and

Args:

srcK1: The calibration matrix of the first source camera.

srcR1: The rotation matrix of the first source camera.

src_t1: The translation vector of the first source camera.

srcK2: The calibration matrix of the second source camera.

srcR2: The rotation matrix of the second source camera.

src_t2: The translation vector of the second source camera.

dstR: The rotation matrix of the destination cameras.

Returns:

dstK: The calibration matrix of the destination cameras.

dst_t1: The translation vector of the first destination camera.

dst_t2: The translation vector of the second destination camera.

H1, H2: The image rectification transformation matrices.

```

'''

```

```

dst_t1 = cameraTranslation(srcR1, src_t1, dstR)

```

```

dst_t2 = cameraTranslation(srcR2, src_t2, dstR)

```

```

dstK = calcDestinateK(srcK1, srcK2)

```

```

H1 = calcProjectiveTransformation(srcK1, srcR1, dstK, dstR)

```

```

H2 = calcProjectiveTransformation(srcK2, srcR2, dstK, dstR)

```

```

return dstK, dst_t1, dst_t2, H1, H2

```

Problem 4.1.2 Warp Image [10 pts] After calling *epipolarRectification*, we can get the projective transformation matrices $H1$ and $H2$. Next, we will geometrically transform (i.e., 'warp') the image so that the epipolar lines are image rows. You have to complete *warpImage* using the backward method in Lecture 7. Note the destination images are required to be the same size as the source

images.

```
In [ ]: def warpImage(image, H, out_height, out_width):
        """
        Performs the warp of the full image content.
        Calculates bounding box by piping four corners through the transformation.

        Args:
        image: Image to warp
        H: The image rectification transformation matrices.
        out_height, out_width: The shape of output image.

        Returns:
        Out: An inverse warp of the image, given a homography.
        min_x, min_y, max_x, max_y: The minimum/maximum of warped image bound.
        """
        """ =====
        YOUR CODE HERE
        ===== """

        return out, min_x, min_y, max_x, max_y

In [ ]: file_param = open('param.pkl', 'rb')
        param = pickle.load(file_param)
        file_param.close()
        srcK1, srcR1, src_t1 = param['srcK1'], param['srcR1'], param['src_t1']
        srcK2, srcR2, src_t2 = param['srcK2'], param['srcR2'], param['src_t2']
        dstR = param['dstR']

In [ ]: dstK, dst_t1, dst_t2, H1, H2 = epipolarRectification(srcK1, srcR1, src_t1,
                                                            srcK2, srcR2, src_t2, dstR)

In [ ]: src1 = imread('Sport0_OGO.bmp')
        plt.imshow(src1)
        print('Original image 1:')

In [ ]: height1, width1, _ = src1.shape
        rectified_im1_unbounded, min_x1, min_y1, max_x1, max_y1 = warpImage(src1, H1, height1,
        plt.imshow(rectified_im1_unbounded)
        print('Unbounded Rectified image 1:')

In [ ]: src2 = imread('Sport1_OGO.bmp')
        plt.imshow(src1)
        print('Original image 2:')

In [ ]: height2, width2, _ = src2.shape
        rectified_im2_unbounded, min_x2, min_y2, max_x2, max_y2 = warpImage(src2, H2, height2,
        plt.imshow(rectified_im2_unbounded)
        print('Unbounded Rectified image 2:')
```


4.1.3 Partial bounded rectification [3 pts] In the resulting images, although they are epipolar rectified, you should observe portions of the source images being transformed "out of bounds" of the destination images. To fix this problem, we can introduce a 2D transformation containing a translation (i.e., $T1$ and $T2$).

$$T1 = \begin{bmatrix} 1 & 0 & -\min_x1 - 0.5 \\ 0 & 1 & -\min(\min_y1, \min_y2) - 0.5 \\ 0 & 0 & 1 \end{bmatrix}$$

$$T2 = \begin{bmatrix} 1 & 0 & -\min_x2 - 0.5 \\ 0 & 1 & -\min(\min_y1, \min_y2) - 0.5 \\ 0 & 0 & 1 \end{bmatrix}$$

$H1$, $H2$ can be updated by left multiplying $T1$, $T2$, respectively. Again, geometrically transform the images under the updated $H1$, $H2$. The destination image is required to be the same size as the source images. In the resulting images, although they are (still) epipolar rectified, you should observe the portions of the source images being transformed are no longer "out of bounds" on the top and left of the destination images.

```
In [ ]: def partialboundedRectification(min_x1, min_y1, min_x2, min_y2, H1, H2):
        """
        Update the projective transformation matrices so that the rectified images are no longer out of bounds.
        """
        """ =====
        YOUR CODE HERE
        ===== """

        return H1_bounded, H2_bounded
```

```
In [ ]: H1_bounded, H2_bounded = partialboundedRectification(min_x1, min_y1, min_x2, min_y2, H1, H2)
        rectified_im1_bounded, min_x1_bounded, min_y1_bounded, max_x1_bounded, max_y1_bounded = H1_bounded
        plt.imshow(rectified_im1_bounded)
        print('The partial bounded rectified image 1:')
```

```
In [ ]: rectified_im2_bounded, min_x2_bounded, min_y2_bounded, max_x2_bounded, max_y2_bounded = H2_bounded
        plt.imshow(rectified_im2_bounded)
        print('The partial bounded rectified image 2:')
```

4.1.4 Completely bounded rectification [3 pts] Finally, determine the size of the destination images that completely bound the transformed images. Again geometrically transform the images under the updated 2D projective transformation matrices $H1$ and $H2$ (these are not updated a second time). You should complete the function *completelyBoundedRectification*. The destination images are required to be the size you just calculated. In the resulting images, you should observe the source images being transformed such that they are epipolar rectified and are completely bounded.

```
In [ ]: def completelyBoundedRectification(src1, src2, H1_bounded, H2_bounded,
        min_x1_bounded, max_x1_bounded,
        min_y1_bounded, max_y1_bounded):
```

```

'''
Determine the size of the destination images (same size for both) that completely
'''
''' =====
YOUR CODE HERE
===== '''

return rectified_im1_final, rectified_im2_final

```

```

In [ ]: rectified_im1_final, rectified_im2_final = completelyBoundedRectification(src1, src2,
min_x1_bounded, max_x1_bounded,
min_y1_bounded, max_y1_bounded)

```

1.6.2 Problem 4.2 Feature matching [3 pts]

4.2.1 SSD (Sum Squared Distance) Matching [1 pts] Complete the function `ssdMatch`:

$$\text{SSD} = \sum_{x,y} |W_1(x,y) - W_2(x,y)|^2$$

```

In [ ]: def ssdMatch(img1, img2, c1, c2, R):
'''Compute SSD given two windows.

Args:
img1: Image 1.
img2: Image 2.
c1: Center (in image coordinate) of the window in image 1.
c2: Center (in image coordinate) of the window in image 2.
R: R is the radius of the patch, 2 * R + 1 is the window size

Returns:
SSD matching score for two input windows.

'''
''' =====
YOUR CODE HERE
===== '''

return matching_score

```

```

In [ ]: # Here is the code for you to test your implementation
img1 = np.array([[1, 2, 3, 4], [4, 5, 6, 8], [7, 8, 9, 4]])
img2 = np.array([[1, 2, 1, 3], [6, 5, 4, 4], [9, 8, 7, 3]])
print(ssdMatch(img1, img2, np.array([1, 1]), np.array([1, 1]), 1))
# should print 20
print(ssdMatch(img1, img2, np.array([2, 1]), np.array([2, 1]), 1))
# should print 30
print(ssdMatch(img1, img2, np.array([1, 1]), np.array([2, 1]), 1))
# should print 46

```

Problem 4.2.2 NCC (Normalized Cross-Correlation) Matching [2 pts] Write a function `ncc_match` that implements the NCC matching algorithm for two input windows.

$$\text{NCC} = \sum_{i,j} \tilde{W}_1(i,j) \cdot \tilde{W}_2(i,j)$$

where $\tilde{W} = \frac{W - \bar{W}}{\sqrt{\sum_{k,l} (W(k,l) - \bar{W})^2}}$ is a mean-shifted and normalized version of the window and \bar{W} is the mean pixel value in the window W .

```
In [ ]: def normalize_window(window):
    _mean = np.mean(window)
    _stdev = np.sqrt(np.sum((window - _mean) ** 2))
    return (window - _mean) / (_stdev + 1e-6)

def ncc_match(img1, img2, c1, c2, R):
    """Compute NCC given two windows.

    Args:
        img1: Image 1.
        img2: Image 2.
        c1: Center (in image coordinate) of the window in image 1.
        c2: Center (in image coordinate) of the window in image 2.
        R: R is the radius of the patch, 2 * R + 1 is the window size

    Returns:
        NCC matching score for two input windows.

    """

    """ =====
    YOUR CODE HERE
    ===== """

In [ ]: # test NCC match
img1 = np.array([[1, 2, 3, 4], [4, 5, 6, 8], [7, 8, 9, 4]])
img2 = np.array([[1, 2, 1, 3], [6, 5, 4, 4], [9, 8, 7, 3]])

print (ncc_match(img1, img2, np.array([1, 1]), np.array([1, 1]), 1))
# should print 0.8546

print (ncc_match(img1, img2, np.array([2, 1]), np.array([2, 1]), 1))
# should print 0.8457

print (ncc_match(img1, img2, np.array([1, 1]), np.array([2, 1]), 1))
# should print 0.6258
```

1.6.3 Problem 4.3 Naive Matching [8 pts]

Equipped with the corner detector and the NCC matching function, we are ready to start finding correspondences. One naive strategy is to try and find the best match between the two sets of

corner points. Write a script that does this, namely, for each corner in image1, find the best match from the detected corners in image2 (or, if the NCC match score is too low, then return no match for that point). You will have to figure out a good threshold (NCCth) value by experimentation.

Write a function `naive_matching` and call it as below. Examine your results for 10, 20, and 30 detected corners in each image. Choose the number of detected corners to maximize the number of correct matching pairs. `naive_matching` will call your NCC matching code.

Properly label or mention which output corresponds to which choice of number of corners. The total number of outputs is 6 images: (3 choices of number of corners for each of `matrix` and `warrior`), where each figure might look like the following:

Number of corners: 10

```
In [ ]: def naive_matching(img1, img2, corners1, corners2, R, NCCth):
        """Compute NCC given two windows.

        Args:
            img1: Image 1.
            img2: Image 2.
            corners1: Corners in image 1 (n x 2)
            corners2: Corners in image 2 (n x 2)
            R: NCC matching radius
            NCCth: NCC matching score threshold

        Returns:
            NCC matching result a list of tuple (c1, c2),
            c1 is the 1 x 2 corner location in image 1,
            c2 is the 1 x 2 corner location in image 2.

        """

        """ =====
        YOUR CODE HERE
        ===== """

        return matching

In [ ]: # detect corners on warrior and matrix sets
        # you are free to modify code here, create your helper functions, etc.

        nCorners = 20 # do this for 10, 20 and 30 corners
        smoothSTD = 1
        windowSize = 17

        # read images and detect corners on images

        imgs_sport = []
        crns_sport = []

        for i in range(2):
```

```

img_sport = imageio.imread('Sport' + str(i) + '_OG0.bmp')
imgs_sport.append(rgb2gray(img_sport))
# downsize your image in case corner_detect runs slow in test
# imgs_mat.append(rgb2gray(img_mat)[:2, :2])
crns_sport.append(corner_detect(imgs_sport[i], nCorners, smoothSTD, windowSize))

```

```

In [ ]: # match corners
R = 16
NCCth = 0.6 # put your threshold here
matching_sport = naive_matching(imgs_sport[0]/255,
                                imgs_sport[1]/255,
                                crns_sport[0],
                                crns_sport[1],
                                R, NCCth)

```

```

In [ ]: # plot matching result
def show_matching_result(img1, img2, matching):
    """ =====
    YOUR CODE HERE
    ===== """

    print("Number of Corners:", nCorners)
    show_matching_result(imgs_sport[0], imgs_sport[1], matching_sport)

```

1.6.4 Problem 4.4 Matching using epipolar geometry [10 pts]

Next, we will use the epipolar geometry constraint on the rectified images and updated corner points to build a better matching algorithm. First, detect 10 corners in image1. Then, for each corner, do a line search along the corresponding parallel epipolar line in image2.

Evaluate the NCC score for each point along this line and return the best match (or no match if all scores are below the NCCth). R is the radius (size) of the NCC patch in the code below.

You do not have to run this in both directions. Show your result as in the naive matching part.

```

In [ ]: def display_correspondence(img1, img2, corrs):
    """Plot matching result on image pair given images and correspondences

    Args:
        img1: Image 1.
        img2: Image 2.
        corrs: Corner correspondence
    """

    """ =====
    YOUR CODE HERE
    You may want to refer to the `show_matching_result` function.
    ===== """

```

```

def correspondence_matching_epipole(img1, img2, corners1, R, NCcth):
    """Find corner correspondence along epipolar line.

    Args:
        img1: Image 1.
        img2: Image 2.
        corners1: Detected corners in image 1.
        F: Fundamental matrix calculated using given ground truth corner correspondences.
        R: NCC matching window radius.
        NCcth: NCC matching threshold.

    Returns:
        Matching result to be used in display_correspondence function
    """

    """ =====
    YOUR CODE HERE
    ===== """

    return matching

```

```

In [ ]: rectified_im1_final, rectified_im2_final = completelyBoundedRectification(src1, src2, H1_bounded, H2_bounded, min_x1_bounded, min_x2_bounded, max_x1_bounded, max_x2_bounded, min_y1_bounded, min_y2_bounded, max_y1_bounded, max_y2_bounded)

```

```

In [ ]: # replace black pixels with white pixels
_black_idxs = (rectified_im1_final[:, :, 0] == 0) & (rectified_im1_final[:, :, 1] == 0)
rectified_im1_final[:, :, _black_idxs] = [1.0, 1.0, 1.0]
_black_idxs = (rectified_im2_final[:, :, 0] == 0) & (rectified_im2_final[:, :, 1] == 0)
rectified_im2_final[:, :, _black_idxs] = [1.0, 1.0, 1.0]

nCorners = 20
# Choose your threshold and NCC matching window radius
NCcth = 0.6
R = 8
# detect corners using corner detector here, store in corners1
corners1 = corner_detect(rgb2gray(rectified_im1_final), nCorners, smoothSTD, windowSize)
corrs = correspondence_matching_epipole(rectified_im1_final, rectified_im2_final, corners1, R, NCcth)
display_correspondence(rectified_im1_final, rectified_im2_final, corrs)

```

```

In [ ]:

```