

1 CSE 152A Winter 2021 – Assignment 1

1.1 Instructor: Ben Ochoa

- Assignment Published On: **Wed, Jan 13, 2021**
- Due On: **Wed, Jan 27, 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: Geometry [10 points]

Consider a line in the 2D plane, whose equation is given by $a\tilde{x} + b\tilde{y} + c = 0$. This can equivalently be written as $\mathbf{l}^\top \mathbf{x} = 0$, where $\mathbf{l} = (a, b, c)^\top$ and $\mathbf{x} = (\tilde{x}, \tilde{y}, 1)^\top$. Noticing that \mathbf{x} is a homogeneous representation of $\tilde{\mathbf{x}} = (\tilde{x}, \tilde{y})^\top$, we can view \mathbf{l} as a homogeneous representation of the line $a\tilde{x} + b\tilde{y} + c = 0$. We see that the line is also defined up to a scale since $(a, b, c)^\top$ and $k(a, b, c)^\top$ with $k \neq 0$ represents the same line. All points (\tilde{x}, \tilde{y}) that lie on the line $a\tilde{x} + b\tilde{y} + c = 0$ satisfy the equation $\mathbf{l}^\top \mathbf{x} = 0$.

Statement 1: A point \mathbf{x} lies on the line \mathbf{l} if and only if $\mathbf{l}^\top \mathbf{x} = \mathbf{x}^\top \mathbf{l} = 0$

1. [6 points] Prove the following two statements that follow from **Statement 1**.
 - a). The cross product between two points gives us the line connecting the two points
 - b). The cross product between two lines gives us their point of intersection
2. [4 points] What is the line, in homogenous coordinates, joining the inhomogeneous points $(2, 9)$ and $(-3, 4)$.

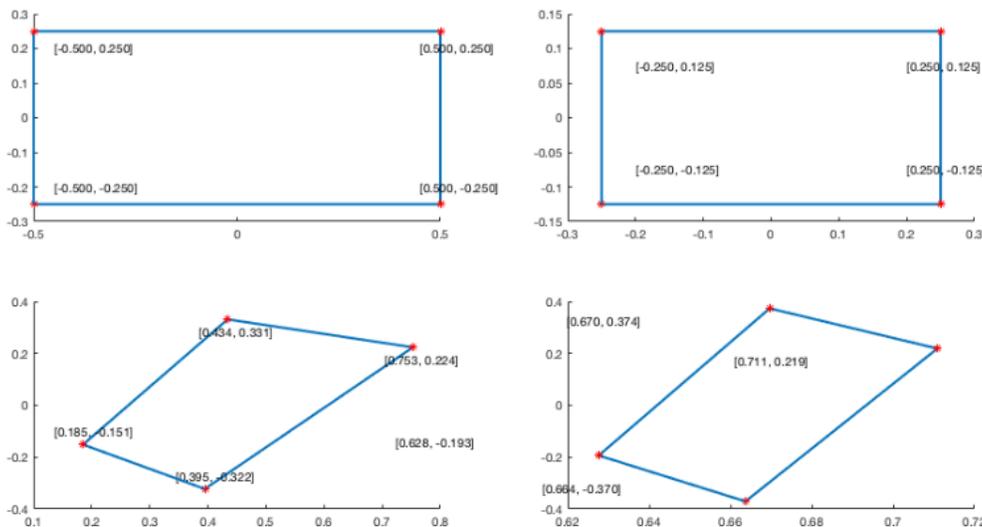
1.4 Problem 2: Image Formation and Rigid Body Transformations [20 points]

In this problem we will practice rigid body transformations and image formations through the projective camera model. The goal will be to photograph the following four points $\widetilde{\mathbf{X}}_1 = [-2 \ -4 \ 3]^T$, $\widetilde{\mathbf{X}}_2 = [2 \ -4 \ 3]^T$, $\widetilde{\mathbf{X}}_3 = [2 \ 4 \ 3]^T$, $\widetilde{\mathbf{X}}_4 = [-2 \ 4 \ 3]^T$ in the world coordinate frame. First, recall the following formula for rigid body transformation

$$\widetilde{\mathbf{X}}_{cam} = R\widetilde{\mathbf{X}} + \mathbf{t}$$

Where $\widetilde{\mathbf{X}}_{cam}$ is the point coordinate in the camera coordinate system. $\widetilde{\mathbf{X}}$ is a point in the world coordinate frame, and R and \mathbf{t} are the rotation and translation that transform points from the world coordinate frame to the camera coordinate frame. Together, R and \mathbf{t} are the *extrinsic* camera parameters. Once transformed to the camera coordinate frame, the points can be photographed using the 3×3 camera calibration matrix K , which embodies the *intrinsic* camera parameters, and the canonical projection matrix $[I|0]$. Given K , R , and \mathbf{t} , the image of a point $\widetilde{\mathbf{X}}$ is $\mathbf{x} = K[I|0]\mathbf{X}_{Cam} = K[R|\mathbf{t}]\mathbf{X}$, where the homogeneous points $\mathbf{X}_{Cam} = (\widetilde{\mathbf{X}}_{Cam}^\top, 1)^\top$ and $\mathbf{X} = (\widetilde{\mathbf{X}}^\top, 1)^\top$. We will consider four different settings of focal length, viewing angles and camera positions below.

- The extrinsic transformation matrix,
- Intrinsic camera matrix under the perspective camera assumption.
- Calculate the image of the four vertices and plot using the supplied `plot_points` function (see e.g. output in figure below).



- [No rigid body transformation]. Focal length = 1. The optical axis of the camera is aligned with the z-axis.
- [Translation]. $\mathbf{t} = [0 \ 0 \ 3]^T$. The optical axis of the camera is aligned with the z-axis.
- [Translation and Rotation]. Focal length = 1. R encodes a 45 degrees around the z-axis and then 45 degrees around the y-axis. $\mathbf{t} = [0 \ 0 \ 3]^T$.
- [Translation and Rotation, long distance]. Focal length = 7. R encodes a 60 degrees around the z-axis and then 30 degrees around the y-axis. $\mathbf{t} = [0 \ 0 \ -10]^T$.

We will not use a full intrinsic camera matrix (e.g. that maps centimeters to pixels, and defines the coordinates of the center of the image), but only parameterize this with f , the focal length. In other words: the only parameter in the intrinsic camera matrix under the perspective assumption is f .

For all the four cases, include a image like above. Note that the axis are the same for each row, to facilitate comparison between the two camera models. Note: the angles and offsets used to generate these plots may be different from those in the problem statement, it's just to illustrate how to report your results.

Also, Explain why you observe any distortions in the projection, if any, under this model.

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

# convert points from euclidian to homogeneous
def to_homog(points):
    """
    your code here
    """
    return points_homog

# convert points from homogeneous to euclidian
def from_homog(points_homog):
    """
    your code here
    """
    return points

# project 3D euclidian points to 2D euclidian
def project_points(P_int, P_ext, pts):
    """
    your code here
    """
    #return the 2d euclidean points
    pts_2d=np.zeros([2,1])
    return pts_2d

def camera1():
    """
    replace with your code
    """
    P_int_proj = np.eye(3,4)
    P_ext = np.eye(4,4)
    return P_int_proj, P_ext

def camera2():
    """
```

```

replace with your code
"""
P_int_proj = np.eye(3,4)
P_ext = np.eye(4,4)
return P_int_proj, P_ext

def camera3():
    """
    replace with your code
    """
    P_int_proj = np.eye(3,4)
    P_ext = np.eye(4,4)
    return P_int_proj, P_ext

def camera4():
    """
    replace with your code
    """
    P_int_proj = np.eye(3,4)
    P_ext = np.eye(4,4)
    return P_int_proj, P_ext

#####
# test code. Do not modify
#####

def plot_points(points, title='', style='.-r', axis=[]):
    inds = list(range(points.shape[1]))+[0]
    plt.plot(points[0,inds], points[1,inds],style)
    if title:
        plt.title(title)
    if axis:
        plt.axis('scaled')
        #plt.axis(axis)

def main():
    point1 = np.array([[-2,-4,3]]).T
    point2 = np.array([[2,-4,3]]).T
    point3 = np.array([[2,4,3]]).T
    point4 = np.array([[-2,4,3]]).T
    points = np.hstack((point1,point2,point3,point4))

    for i, camera in enumerate([camera1, camera2, camera3, camera4]):
        P_int_proj, P_ext = camera()
        plt.subplot(1, 2, 1)

```

```

        plot_points(project_points(P_int_proj, P_ext, points), title='Camera %d_
↪Projective'%(i+1), axis=[-1,1,-1,1])
        plt.show()

main()

```

1.5 Problem 3: Photometric Stereo [20 pts]

The goal of this problem is to implement Lambertian photometric stereo.

Note that the albedo is unknown and non-constant in the images you will use.

As input, your program should take in multiple images along with the light source direction for each image.

1.5.1 Data

You will use synthetic images as data. These images are stored in `.pickle` files which were graciously provided by Satya Mallick. Each `.pickle` file contains

- `im1, im2, im3, im4, ...` images.
- `l1, l2, l3, l4, ...` light source directions.

Implement the photometric stereo technique described in the lecture. Your program should have two parts:

1. Read in the images and corresponding light source directions, and estimate the surface normals and albedo map.
2. Reconstruct the depth map from the surface with the implementation of the Horn integration technique given below in `horn_integrate` function. Note that you will typically want to run the `horn_integrate` function with 10000 - 100000 iterations, meaning it will take a while.

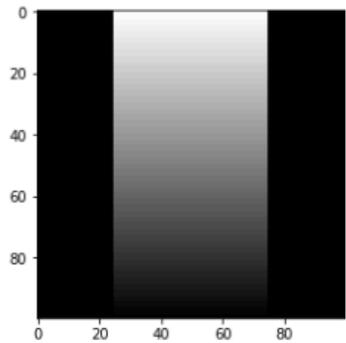
You will find all the data for this part in `synthetic_data.pickle`. Try using only `im1, im2` and `im4` first. Display your outputs as mentioned below.

Then use all four images (most accurate).

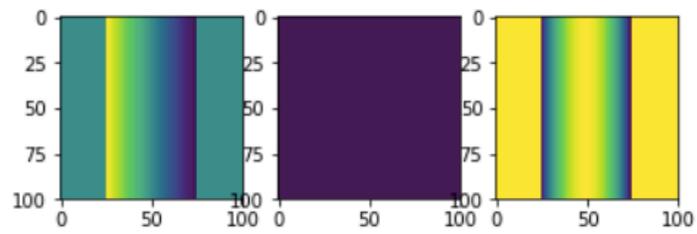
For **each** of the **two above cases** you must output:

1. The estimated albedo map.
2. The estimated surface normals by showing both
 1. Needle map, and
 2. Three images showing components of surface normal.
3. A wireframe of depth map.

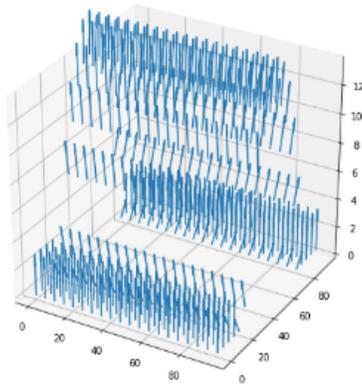
An example of outputs is shown in the figure below.



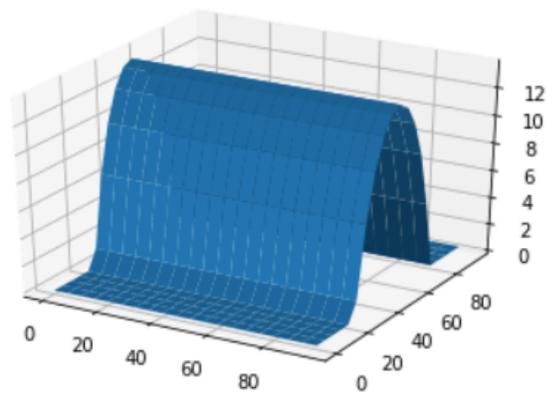
Albedo map



Normals as three separate channels



Needle map



Wireframe of depth map

```
[2]: ## Example: How to read and access data from a pickle
import pickle
import matplotlib.pyplot as plt
%matplotlib inline
pickle_in = open("synthetic_data.pickle", "rb")
data = pickle.load(pickle_in, encoding="latin1")

# data is a dict which stores each element as a key-value pair.
print("Keys: " + str(data.keys()))

# To access the value of an entity, refer it by its key.
print("Image:")
plt.imshow(data["im1"], cmap = "gray")
plt.show()

print("Light source direction: " + str(data["l1"]))

plt.imshow(data["im2"], cmap = "gray")
plt.show()
```

```
print("Light source direction: " + str(data["l2"]))

plt.imshow(data["im3"], cmap = "gray")
plt.show()

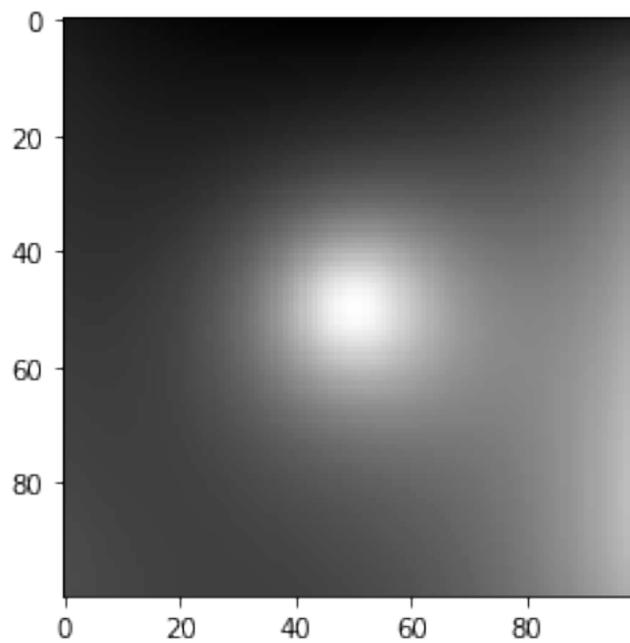
print("Light source direction: " + str(data["l3"]))

plt.imshow(data["im4"], cmap = "gray")
plt.show()

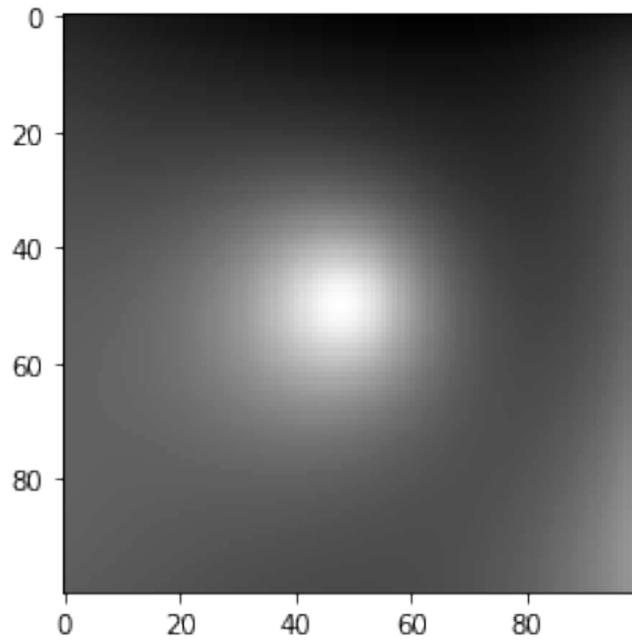
print("Light source direction: " + str(data["l4"]))
```

Keys: dict_keys(['__version__', 'l4', '__header__', 'im1', 'im3', 'im2', 'l2', 'im4', 'l1', '__globals__', 'l3'])

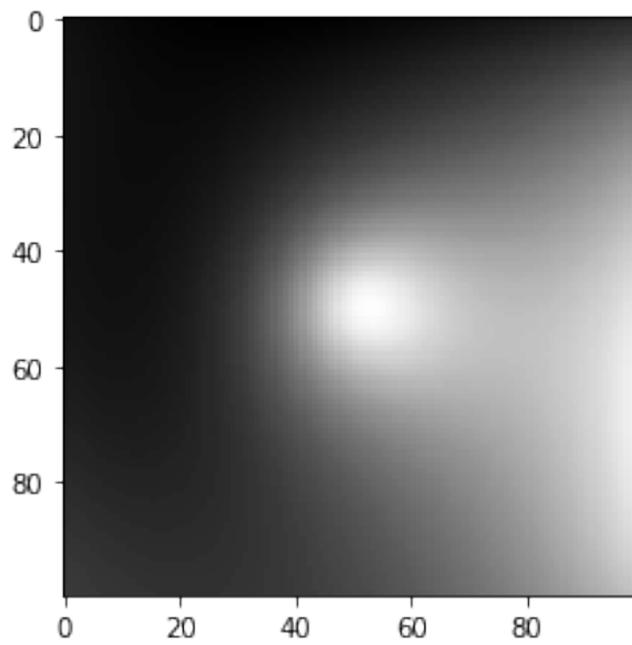
Image:



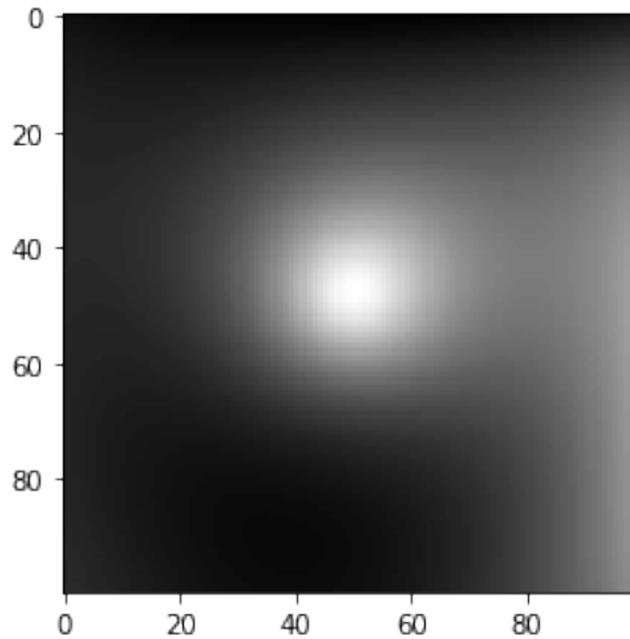
Light source direction: $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$



Light source direction: `[[0.2 0. 1.]]`



Light source direction: `[[-0.2 0. 1.]]`



Light source direction: `[[0. 0.2 1.]]`

```
[3]: import numpy as np
from scipy.signal import convolve
from numpy import linalg

def horn_integrate(gx, gy, mask, niter):
    """
    horn_integrate recovers the function g from its partial
    derivatives gx and gy.
    mask is a binary image which tells which pixels are
    involved in integration.
    niter is the number of iterations.
    typically 100,000 or 200,000,
    although the trend can be seen even after 1000 iterations.
    """
    g = np.ones(np.shape(gx))

    gx = np.multiply(gx, mask)
    gy = np.multiply(gy, mask)

    A = np.array([[0,1,0],[0,0,0],[0,0,0]]) #y-1
    B = np.array([[0,0,0],[1,0,0],[0,0,0]]) #x-1
    C = np.array([[0,0,0],[0,0,1],[0,0,0]]) #x+1
    D = np.array([[0,0,0],[0,0,0],[0,1,0]]) #y+1
```

```

d_mask = A + B + C + D

den = np.multiply(convolve(mask,d_mask,mode="same"),mask)
den[den == 0] = 1
rden = 1.0 / den
mask2 = np.multiply(rden, mask)

m_a = convolve(mask, A, mode="same")
m_b = convolve(mask, B, mode="same")
m_c = convolve(mask, C, mode="same")
m_d = convolve(mask, D, mode="same")

term_right = np.multiply(m_c, gx) + np.multiply(m_d, gy)
t_a = -1.0 * convolve(gx, B, mode="same")
t_b = -1.0 * convolve(gy, A, mode="same")
term_right = term_right + t_a + t_b
term_right = np.multiply(mask2, term_right)

for k in range(niter):
    g = np.multiply(mask2, convolve(g, d_mask, mode="same")) + term_right

return g

```

```

[3]: def photometric_stereo(images, lights, mask, horn_niter=25000):

    """mask is an optional parameter which you are encouraged to use.
    It can be used e.g. to ignore the background when integrating the normals.
    It should be created by converting the images to grayscale and thresholding
    (only using locations for which the pixel value is above some threshold).

    The choice of threshold is something you can experiment with,
    but in practice something like 0.05 tends to work well.
    """

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

    # note:
    # images : (n_ims, h, w)
    # lights : (n_ims, 3)
    # mask   : (h, w)

    albedo = np.ones(images[0].shape)
    normals = np.dstack((np.zeros(images[0].shape),
                        np.zeros(images[0].shape),
                        np.ones(images[0].shape)))

```

```
H_horn = np.ones(images[0].shape)
return albedo, normals, H_horn
```

```
[ ]: from mpl_toolkits.mplot3d import Axes3D

pickle_in = open("synthetic_data.pickle", "rb")
data = pickle.load(pickle_in, encoding="latin1")

lights = np.vstack((data["l1"], data["l2"], data["l4"]))

## Set-1: USING ONLY im1, im2, im4

images = []
images.append(data["im1"])
images.append(data["im2"])
images.append(data["im4"])
images = np.array(images)

mask = np.ones(data["im1"].shape)

albedo, normals, horn_depth = photometric_stereo(images, lights, mask)

# -----
# The following code is just a working example so you don't get stuck with any
# of the graphs required. You may want to write your own code to align the
# results in a better layout. You are also free to change the function
# however you wish; just make sure you get all of the required outputs.
# -----

def visualize(albedo, normals, horn_depth):
    # Stride in the plot, you may want to adjust it to different images
    stride = 15

    # showing albedo map
    fig = plt.figure()
    albedo_max = albedo.max()
    albedo = albedo / albedo_max
    plt.imshow(albedo, cmap="gray")
    plt.show()

    # showing normals as three separate channels
    figure = plt.figure()
    ax1 = figure.add_subplot(131)
    ax1.imshow(normals[... , 0])
    ax2 = figure.add_subplot(132)
    ax2.imshow(normals[... , 1])
```

```

ax3 = figure.add_subplot(133)
ax3.imshow(normals[... , 2])
plt.show()

# showing normals as quiver
X, Y, _ = np.meshgrid(np.arange(0,np.shape(normals)[0], 15),
                      np.arange(0,np.shape(normals)[1], 15),
                      np.arange(1))

X = X[... , 0]
Y = Y[... , 0]
Z = horn_depth[::stride,::stride].T
NX = normals[... , 0][::stride,::-stride].T
NY = normals[... , 1][::-stride,::stride].T
NZ = normals[... , 2][::stride,::stride].T
fig = plt.figure(figsize=(5, 5))
ax = fig.gca(projection='3d')
plt.quiver(X,Y,Z,NX,NY,NZ, length=10)
plt.show()

# plotting wireframe depth map

H = horn_depth[::stride,::stride]
fig = plt.figure()
ax = fig.gca(projection='3d')
ax.plot_surface(X,Y, H.T)
plt.show()

visualize(albedo, normals, horn_depth)

```

```

[1]: # Don't forget to run your photometric stereo code on TWO sets of images!
# (One being {im1, im2, im4}, and the other being {im1, im2, im3, im4}.)
""" =====
YOUR CODE HERE
===== """

```

1.6 Problem 4: Image Rendering [20 points]

In this exercise, we will render the image of a face with two different point light sources using a Lambertian reflectance model. We will use two albedo maps, one uniform and one that is more realistic. The face heightmap, the light sources, and the two albedo are given in `facedata.npy` for Python (each row of the 'lightsource' variable encode a light location). The data from `facedata.npy` is already provided to you.

Note: Please make good use out of subplot to display related image next to eachother.

Plot the face in 2-D [2 pts]

Plot both albedo maps using `imshow`. Explain what you see.

Plot the face in 3-D [2 pts]

Using both the heightmap and the albedo, plot the face using `plot_surface`. Do this for both albedos. Explain what you see.

Surface normals [8 pts]

Calculate the surface normals and display them as a quiver plot using `quiver` in `matplotlib.pyplot` in Python. Recall that the surface normals are given by

$$\left[-\frac{\delta f}{\delta x}, -\frac{\delta f}{\delta y}, 1\right]. \quad (1)$$

Also, recall, that each normal vector should be normalized to unit length.

Render images [8 pts]

For each of the two albedos, render three images. One for each of the two light sources, and one for both light-sources combined. Display these in a 2×3 subplot figure with titles. Recall that the general image formation equation is given by

$$I = a(x, y) \hat{\mathbf{n}}(x, y)^\top \hat{\mathbf{s}}(x, y) s_0 \quad (2)$$

where $a(x, y)$ is the albedo for pixel (x, y) , $\hat{\mathbf{n}}(x, y)$ is the corresponding surface normal, $\hat{\mathbf{s}}(x, y)$ the light source direction, s_0 the light source intensity. Let the light source intensity be 1 and make the 'distant light source assumption'. Use `imshow` with appropriate keyword arguments .

```
[ ]: import numpy as np
import matplotlib.pyplot as plt
from matplotlib.path import Path
import matplotlib.patches as patches
# Load facedata.npy as ndarray
face_data = np.load('facedata.npy', encoding='latin1', allow_pickle=True)
# Load albedo matrix
albedo = face_data.item().get('albedo')
# Load uniform albedo matrix
uniform_albedo = face_data.item().get('uniform_albedo')
# Load heightmap
heightmap = face_data.item().get('heightmap')
# Load light source
light_source = face_data.item().get('lightsource')
```