

Problem Set 2: Anaglyph Camera Obscura

Posted: Thursday, September 13, 2018

Due: Thursday, September 20, 2018

Submission Instructions: Please submit **two separate files:** 1) a report named `<your_kerberos>.pdf`, including your responses to all required questions with images and/or plots showing your results and 2) a file named `<your_kerberos>.zip`, containing relevant source code. **Submissions that do not adhere to these instructions are subject to an additional penalty.**

Late Submission Policy: We do not accept late submissions. The submission deadline has a 50-minute soft cut-off; after midnight Thursday, submissions are penalized 2% per minute late.

Collaborators: You are free to discuss problems with other students but all writing must be done individually. Please list all collaborators at the top of your report.

Make sure to read the **entire** problem set before you start working.

You can do this problem set in groups of 2 - 4 people, however each person should submit a separate report. Write the names of the students you collaborated with in the report. You can include the same set of images.

1 Building a digital camera obscura

1.1 Build the camera obscura

The first part of this project consists of building your own “camera obscura”, or pinhole camera. A camera obscura is a dark chamber with a single hole in one extreme that lets light in. The light gets projected into the opposite wall forming an image. In order to capture the projected image we will use a digital camera. One of the problems with pinhole cameras is that the amount of light that enters the device is very small. Therefore, we will need to use long exposure time to capture images with enough contrast. Many digital cameras with a manual mode allow setting the exposure time. We used a Canon G11 for the pictures in this problem set, but an iPhone (or other) camera should work fine. Please note: you may need a lot of light, so don’t start this project on Thursday night!

Let’s first start building the device:



Figure 1: Building the pinhole camera.

1. Take a box (we used a card box with $11 \times 9 \times 15$ inches).
2. Make a hole in the center of one of the faces. The hole has to be small (a diameter of 5 millimeters).
3. Make a second hole so that you can fit the lens of the camera inside.
4. Cover the inside walls with black paper to reduce inter-reflections.
5. Cover the inside wall facing the pinhole with white paper. The white paper will be the screen on which the image will be projected. The digital camera will be used to take pictures of the image that gets projected onto the white paper. Using a paper of standard size (e.g. A4), will assist you in calibrating the camera later on (Section 3).

Figure 1 shows some pictures of our construction. In this construction, the distance between the pinhole and the white wall is 11 inches. The big hole visible in the pictures is the hole we used to insert the digital camera. The pinhole is a very small hole in the center and it is hard to see in these pictures. It is important that the camera fits tightly so that no light enters via that aperture. Make sure that light enters only via the small hole in the center of the box.

Once you have built your camera obscura, take some pictures of your construction to put on the report.

1.2 Take some pictures with the camera

For this you will need a digital camera for which you can set the exposure time. We used a Cannon G11 with exposure time of 15 seconds and set the ISO to maximum sensibility. Figure 2 shows a few images taken with our setup.

If the images appear too blurry, you can increase the sharpness by reducing the size of the hole. But then you will need to increase the exposure time.

You need to take pictures under bright light (sunny day outdoors). But you can play first by shining a light at the hole. If you see nothing, then there might be something wrong with the camera settings. Think about the geometry of the construction to understand why the pictures appear slightly deformed.



Figure 2: Upper left: this is a picture of a bright light source. This picture makes the white screen on the back of the wall appear clearly. Upper right: a picture on a cloudy day. Bottom row: a couple of images we have taken during a sunny day.

2 Anaglyph camera obscura

One interesting aspect of this device is that you can build different types of cameras by playing with the shape of the pinhole. So, let's make a new kind of camera obscura. In particular, we can build a camera that produces anaglyph images in a single shot by making two pinholes instead of just one.

There are several ways of taking stereo pictures and viewing them. The most common mechanism is that two pictures are taken independently (with two cameras) and then they are combined to form anaglyph images, or viewed with special devices (e.g., LCD shutter glasses, polarized 3d glasses, stereoscopes,...) that show each image from the stereo pair to each of the two eyes. Anaglyph images are a particular method to produce 3D images. Anaglyph images are formed by superimposing a pair of stereo images into two different color channels. They provide a 3D effect when viewed with color glasses.

We will transform the camera obscura into an anaglyph camera obscura by making two holes and placing different color filters on each hole. This will produce an image that will be the superposition of two images taken from slightly different viewpoints. Each image will have a different color. Then, to see the 3D image we can look at the picture by placing the same color filters in front of our eyes. In order to do this, we used color glasses with blue and red filters:



We used one of the glasses to get the filters and placed them in front of the two holes. Figure 3 shows two pictures of the two holes (top left) and the filters positioned in front of each hole (top right). On the bottom of that figure is a resulting 3D image, with the two views superimposed. That image should give you a 3D percept when viewed with anaglyph glasses.

Suggestions: For this to work you will have to experiment with the distance between the two holes. We tried two distances: 3 inches and 1.5 inches. We found it easier to see the 3D for the 1.5 inches images. Try to take good quality images. Once you would have placed the color filters on each hole, you will need longer exposure times in order to get the same image quality than before. If the images are too blurry, it will be hard to perceive 3D. Also, as you increase the distance between the pinholes, you will increase the range of distances that will provide 3D information. However, the images will be farther apart which makes the 3D perception harder. If the images are too far apart, the visual system will not be able to fuse them and you will not see a 3D image when looking with the color glasses. You can also try viewing the images by changing the image size.

3 Calibration and image correction

The captured images contain distortion due to the angle between the digital camera and the principle axis of the pinhole camera. If the digital camera could be precisely calibrated (i.e. we knew the exact angle and its distance to the image plane, we could calculate an exact *homography* that transforms the image from the digital camera viewpoint to the viewpoint of the pinhole. Since we cannot measure this accurately in our setup, we will use a simplified calibration procedure.

At each time you setup your camera, take one picture with a strong light source that will make the four corners of the white sheet apparent (Figure 2 top left; this can also be done by reducing the contrast of a capture image). Write a small program

```
H = calibrate(im)
```

in MATLAB that allows you to mark those four corners manually on the image, and computes the (homogeneous) transformation matrix H that rectifies the camera distortion. To figure out the dimensions of the white sheet when viewed from the pinhole direction at a distance focal length away, take a picture of the white sheet with your digital camera at the same focal length and height as the pinhole (this need only be done once).




Figure 3: Anaglyph camera obscura.

This process will give you four point correspondences, from which you can compute an homography that transforms one set of points to the other, as discussed in class. Note that this mapping will need to be computed each time you move your digital camera. You'd might also want to add to your transformation matrix reflection about the x -axis. This will rectify the image and flip it with a single warp.

Please apply your function to the images you've taken with the camera, and submit the calibrated images.

Hints: you can use the MATLAB function `ginput` or `input` for entering the corner locations, and the functions `interp2` or `griddata` to warp the image. If your images are noisy, you can further play with different filtering techniques such as `medfilt2`.




Figure 4: Light at different distances as captured by the anaglyph camera

4 Anaglyph camera obscura: A device to measure distances to objects

We can all measure distances to an object by using a ruler. However, using a ruler is an intrusive procedure, as it requires physically placing one extreme of the ruler touching the object. Some objects do not like to be constantly approached by rulers. Stereovision provides a way of measuring distances between the camera and other objects in the world remotely, without touching them.

Figure 4 shows two example pictures of a light taken at two different distances. As you can see, a single light produces two spots in the picture. This is due to the fact that the light enters the camera obscura through two pinholes instead of one. Each spot has a different color because we placed two colored filters on each pinhole to form the anaglyph image. The interesting property is that the distance between the two spots, d , changes with respect to the distance, Z , between the camera and the light. Therefore, we could measure the distance Z to the light by just measuring d and if we had a perfect model of the geometry of the camera. We will do this in the next part.

4.1 Find the relation between the separation of the lights and distance

Place the light at various known distances from the pinhole and take a set of pictures.

You can decompose the anaglyph image into the images projected by each pinhole by visualizing the red and blue channels separately:

```
img = imread ('yourfilename.jpg');
figure
imshow([img(:,:,1); img(:,:,3)])
```

Plot the relationship between Z and d . You should think about the geometry of the camera in order to derive an analytical expression that relates the distance between the two spots in the picture with the distance of the light with respect to the camera and the parameters of the camera (distance to the back wall, separation between the two holes, etc.). Check that your predictions fit the experimental data.

4.2 A manual stereo algorithm (6.869 only)

In this part, the goal is to apply manually a stereo reconstruction algorithm. The goal is: given two images, to recover the full 3D structure of the image (we will only produce a coarse approximation to it). Automatic algorithms for measuring depths from stereo can be challenging. However, we will do it manually by following the next steps:

1. Take a picture with the anaglyph camera. Try to take a picture in which there are many clear elements that can be easily identified in the image. Also, select a scene with enough depth differences (with some elements very close to the camera and others far away).
2. Use MATLAB to select a number of matching points between the two color channels of the anaglyph image manually.
3. Use the function from the previous section that relates the horizontal distance between matching points to recover the distance between the camera and the selected points.
4. Now you will need to interpolate the distance in order to create a continuous image of depths. You can use the function `griddata`.

Figure 5 (top) shows one example showing the red and blue channels of the stereo picture. We used the top image as the reference image, and then computed depth at a set of selected points on the top image. Once you have a depth map (the value of the distance at each pixel in the image), you can do different things like trying to rotate the image in 3D. Figure 5 shows two different viewpoints of the car.

Add to your report the picture you took, the point correspondences, and different viewpoints of the scene.




Figure 5: Point correspondences (top) and viewpoint interpolation (bottom) of the car image.