Assignment: Cameras and Light

Erik G. Learned-Miller

April 5, 2011

For this assignment, I do not want you to do ANY collaboration whatsoever. It is important that you work through this assignment on your own. You are welcome to come to Manju or me for help. -Erik

Problem 1. (45 points) Building a 3-CCD camera.

You've joined the pinhole camera club, and you've become so obsessed with pinhole cameras you want to make the world's first pinhole camera with three charge-coupled devices (CCDs) instead of just one. CCDs are light sensitive devices that convert light energy into electrical signals that can be digitized into digital images.

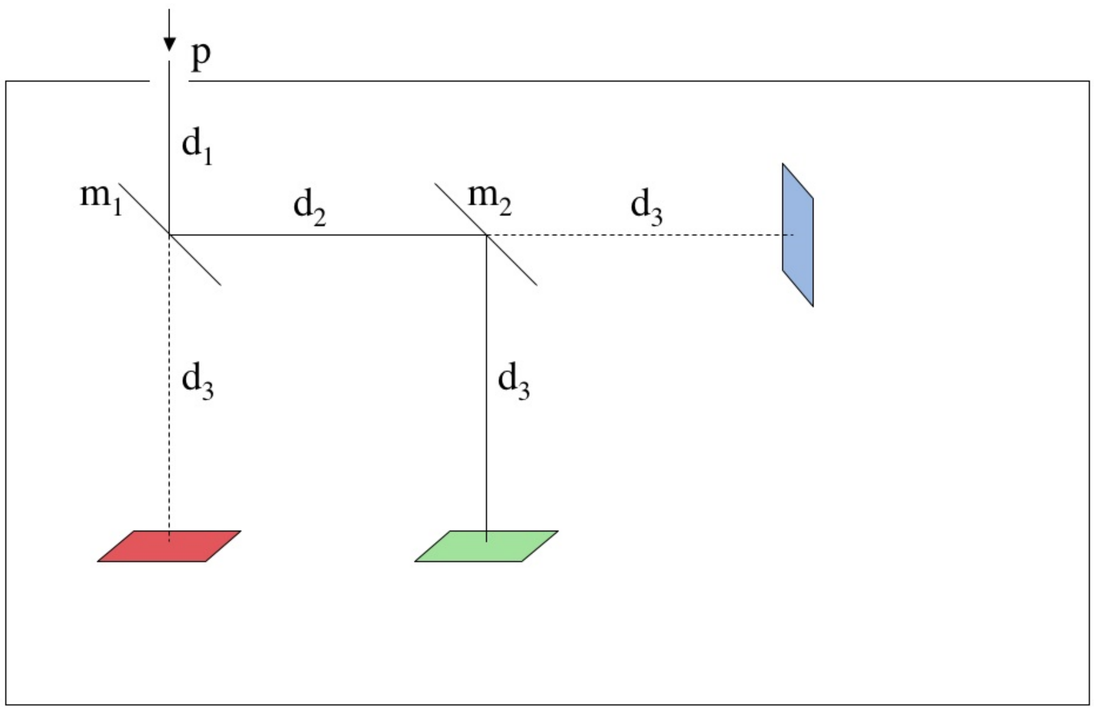

Figure 1: **Rough diagram of 3-CCD pinhole camera.** The pinhole where light enters is marked as p. Semi-transparent mirrors used to both reflect and transmit light are labelled m_1 and m_2. The distance between the pinhole and mirror 1 is d_1. The distance between m_1 and m_2 is d_2. The distances between the mirrors and the three CCDs are all d_3. The CCDs are shown in red, green, and blue to show the 3 different sensitivities of the detectors.

You have been given the rough design of a 3-CCD pinhole camera. The light comes in at the pinhole p and hits a semi-transparent mirror m_1 after travelling a distance d_1. Some of the light bounces off of the mirror continuing on to d_2, while some of the light continues on to hit the red CCD. Dotted lines show light that has passed through a mirror and solid lines show light that has bounced off of a mirror.

Your job is to complete the design of the pinhole camera, which involves the following constraints.

- You must order the proper mirrors from a manufacturer that sells a variety of mirror types. Each mirror comes with 3 numbers $\{a, r, t\}$, which stand for the absorption coefficient, the reflection coefficient, and the transmission coefficient for the mirror. They describe, for a particular mirror, how much light is absorbed, reflected, and transmitted. For every mirror, these numbers add up to 1.0. For example, a mirror might have $\{a = 0.2, r = 0.5, t = 0.3\}$. **Every mirror absorbs at least 10% of the light falling on it. In other words, every a is at least 0.1.** Your mirror specification should specify the values of a, r, and t for each mirror m_1 and m_2. They can be different or the same, depending upon what you want to achieve. All other things being equal, it is clearly better to have more light falling on a CCD rather than less light.

- You will also need to chose the distances d_1, d_2, and d_3 in the given design. All distances must be greater than 0.

- Note that no matter what these distances are, the green and blue CCDs will be further from p than the red CCD. **This means that the red CCD will be closer to p, along the path of the light, than the green and blue CCDs.** This has two separate implications in terms of the size of the image and the amount of light falling on each pixel of the CCD. Your solution should discuss the specifics of these implications.

- After taking an image with your 3 CCD camera, your colleague will write a program to make the images the same size by expanding the red image in memory (using the procedure we went over in class for transforming an image). However, he doesn't want to write code to change the brightness of the image. Therefore, he asks you to buy the appropriate mirrors and design the camera so that the brightness of the images falling on the red CCD matches the brightness falling on the other CCDs.

You now have everything you need to design the camera. Your job is to specify the mirrors to be purchased, and how to space them to achieve the desired goals. There may be more than one correct solution. Please give a detailed discussion of how you came up with the appropriate parameters.

Problem 2. (10 points) Number of materials. In class, we discussed various features that could be used to distinguish between materials of various types (brass, polished billiard balls, surface of a banana) even when you don't know the shape of an object.

I want you to think of a few features that could be used to get information about the material a surface is constructed from. Then, I want you to take an educated guess (I won't grade you on the accuracy of your guess for this part) about how many different values of that feature a human could distinguish. Using the number of distinguishable values of each feature, speculate about the number of different materials a human might be able to distinguish for unfamiliar objects.

Problem 3. (45 points) Bayer patterns

When a color camera is built from a single sensor, it typically uses a "Bayer pattern" to measure either red, green, or blue intensities at each point in the image. These cameras do NOT measure all three intensities at each pixel. It is later, during an image processing phase, that these cameras "fill in" the missing values for the color channels that were not measured directly. This process is usually some sort of interpolation. For example, a missing blue channel value can be replaced by the average of the nearby blue values.

Part A. Look up the specifics of the Bayer pattern on the internet. Write a matlab function that takes a regular RGB image, and produces the image as it would be seen by a Bayer pattern. Since a Bayer pattern only measures a single number for each position in the image, the output of this function should just be a two-dimensional array of values. In other words, for the first pixel in the image, you will want to record only the red value. For the second pixel, only the green value, and so forth.

Part B. Write a simple matlab function that takes an image generated by a Bayer filter (like the output of Part A) and recreates the original 3-channel image. Use averaging to replace missing channel values. (**THIS MUST BE YOUR OWN CODE, NOT CODE YOU FOUND OR ALTERED FROM THE WEB.**) Show some examples of images with a very small number of pixels (say, 25) that have been blown up to be quite large to show that your interpolation works.

Part C. Imagine you have an image showing a "perfect" American flag. That is, it has perfectly horizontal red and white stripes, and perfectly white stars on a perfectly blue background. Now imagine that you sample this image with a Bayer pattern and apply your interpolation algorithm. Where would you see problems? Where would the image be perfect?