

Assignment #3




Figure 1: Layout for 3-camera stereo. The distance between consecutive cameras is B . All three cameras are identical: they have the same focal length f and their image planes are parallel to the line of camera centers C_L, C_0, C_R . A point P is imaged at positions x_L, x_0, x_R in the left, central and right camera respectively.

1 Disparities

Consider the three-camera setup shown in Fig. 1. A point P is imaged at positions x_L, x_0, x_R in the left, central and right camera respectively. The *disparity* of P in the left camera (with respect to the central camera) is given by

$$d_L = x_L - x_0$$

Similarly, the disparity of P in the right camera (with respect to the central camera) is given by

$$d_R = x_R - x_0$$

Show that $d_L = -d_R$. Submit your answer in a Readme file. Any diagrams/drawings should be attached (scanned images or photographs of hand drawing are welcome).

2 3D Reconstruction

You are required to write code which given three images from cameras in the layout of Fig. 1., computes the disparity of each pixel in the central camera. Consider a pixel with coordinates (x_0, y_0) in the central camera. Suppose it has disparity d_0 in the left camera. Thus, it corresponds to the pixel $(x_0 + d_0, y_0)$ in the left camera and $(x_0 - d_0, y_0)$ in the right camera. Our job is to estimate d_0 for each pixel in the central camera image.

Let I_L, I_0, I_R be the images from the left, central and right cameras respectively. Assume that the disparity of each pixel is an integer and lies in a range $[D_{min}, D_{max}]$.

We will assume that corresponding pixels in different cameras have the same color. (This is a reasonable assumption when the objects in the scene are diffuse and the cameras are identical Typically, this is achieved by using a single camera which moves across three positions). Now suppose we have hypothesized that a pixel (x, y) has disparity d . How can we evaluate how good our hypothesis is ? We shall describe two ways, which you are required to implement.

2.1 Sum of Absolute Differences

We can use what is known as the sum-of-absolute distances (SAD) approach. For a pixel (x, y) in the central image and a hypothesized disparity d , we can define:

$$SAD(x, y, d) = \sum_{|x'|\leq W, |y'|\leq W} |I_0(x + x', y + y') - I_L(x + d + x', y + y')| + \sum_{|x'|\leq W, |y'|\leq W} |I_0(x + x', y + y') - I_R(x - d + x', y + y')|$$

This computes the sum of absolute color differences between a window of pixels around (x, y) in the central image I_0 and the corresponding windows in the left and right images. The size of the window used is $2W + 1 \times 2W + 1$. The idea is that if the disparity is correct, the windows are similar and $SAD(x, y, d)$ is small but if the disparity is not correct then $SAD(x, y, d)$ could be large since we are taking color differences between independent pixels. We can estimate the disparity as:

$$Disparity(x, y) = \arg \min_d SAD(x, y, d)$$

That is, for each point we choose the disparity which minimizes the sum of absolute differences.

Important Note: For color images, the difference of two color values is defined to be the sum of the absolute values of the differences in each color channel. If I, I' are two color images then

$$|I(x, y) - I'(x', y')| = |I_{Red}(x, y) - I'_{Red}(x', y')| + |I_{Green}(x, y) - I'_{Green}(x', y')| + |I_{Blue}(x, y) - I'_{Blue}(x', y')|$$

2.2 Median of Absolute Differences

Instead of adding all pixel color differences over a window, we could take the median difference. Specifically, we define

$$MAD(x, y, d) = \text{Median}_{|x'|\leq W, |y'|\leq W} \{|I_0(x + x', y + y') - I_L(x + d + x', y + y')|\} + \\ \text{Median}_{|x'|\leq W, |y'|\leq W} \{|I_0(x + x', y + y') - I_R(x - d + x', y + y')|\}$$

and can estimate the disparity as:

$$Disparity(x, y) = \arg \min_d MAD(x, y, d)$$

Even at the correct disparity all pixels in a window may not match across images, such as due to occlusion or depth discontinuities. Taking the median is an attempt to be robust to such a scenario.

2.3 Implementation

Write code to compute the disparity for pixels in the central image using *SAD* and *MAD*. W , which controls the size of the window used should be an input parameter to your code. You do not need to find the disparity for pixels near the boundary where the window crosses the image boundary. The disparity estimation program should output a *disparity map* which is grayscale image where the intensity of each pixel is equal to the 6 times the disparity value estimated for it. Test your code on the three images provided in the directory <http://cs223b.stanford.edu/homework/ass3/2/> using a disparity range of 0 to 42 pixels.

Submit:

1. Your code for disparity estimation using *SAD* and *MAD*.
2. The disparity map obtained using *SAD* with a 5×5 window ($W = 2$) as a .png image
3. The disparity map obtained using *MAD* with a 7×7 window ($W = 3$) as a .png image.