Homework 2

16-311: Introduction to Robotics

2018

Contents

1 Learning Objectives

1. Gain experience creating and using filters.
2. Improve understanding of convolution.
3. Practice geometric techniques for finding distance using similar triangles and stereo cameras.

2 Filters

In this section, we will be discussing filters. Filtering an image can be useful to average the value of a group of pixels and reduce noise or missing data in the image file.

Create a .pdf file with the written answers ALL THE SECTIONS named hw2.pdf.

1. Come up with a three-by-three filter that can take the average of the nine pixels it overlaps with.

2. Come up with a three-by-three filter that takes the average of the neighbors of the center pixel and the center pixel by the 4-point connectivity definition of neighbor.

3. If we convolve an image with either of these filters, will we lose information? Is it possible to recover the original image from the result of either of these convolutions?

4. Come up with a slightly more trusting three-by-three filter that gives a weighted average of the 9 pixels that it would overlap with, with more weight on the center pixel. Please make sure all the weights sum to 1.

3 Convolution

The goal of this section is to create your own implementation of a convolution function.

Create a MATLAB function (called convolve.m) that can be called by the following script to perform convolution on an image with a mask: `http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/16311/www/current/hw/hw2/convolutionScript.m`. This script uses this image: `http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/16311/www/current/hw/hw2/image2.pgm`.

Specifically, your convolution function must perform convolution on an image with double values between 0 and 1 with a mask of variable size. The original image and your convolved image must be the same height and width. You must develop some way of dealing with the edges that is not padding your image with 0s or 1s. Please develop some other method. This is a relatively brief function to write. But it is extremely versitile. It is more efficient to create a vectorized version of this function (i.e. limiting looping), but that is not required for submission. Do not worry about flipping the mask 90 degrees or normalizing it. Just leave it as written. You may place the result of your multiplication anywhere you choose as long as you are consistant. That is, if you have a 3 by 3 mask, you could elect to place the product at the top left of that mask and have a buffer at the bottom and right of

your image. Or place the product in the center and have a buffer on all sides. Keep in mind the masks are of different sizes.

In your hw2.pdf writeup, provide the three images produced by this script with your convolution function. Make sure you note which mask produced each image and comment on the general effect of that mask.

You may not use any im_ function besides imread. You may not use a funtion that does convolution for you. You may not use a function that pads your image for you.

4 Determining Distance Geometrically

Suppose you have a robot with a camera that you need to characterize. This robot has been in the lab for a while so there is little documentation available and the sensing box has been welded shut so you cannot measure the camera's geometry. You want to find the camera's focal length so you can use that for future experimentations to judge distance.

You take a picture of a one-inch square cube from a precise distance of one foot. You threshold and segment the image and count the number of pixels that represent the cube. You repeat this process 10 times and get an average of 47 pixels for the width of the one-inch cube. You read someone's well-documented code for displaying the images from this camera and see that the camera's resolution is 100 ppi (pixels per inch).

Using this information, what is the camera's focal length? Please include at least one symbolic equation with your answer.

5 Stereo Vision

The goal of this section is to practice implementing the geometry involved in determining depth from an image and position from stereo cameras.

The image below shows a robot with two forward-facing cameras. The cameras are 10 cm apart and have a focal length of 50 mm.

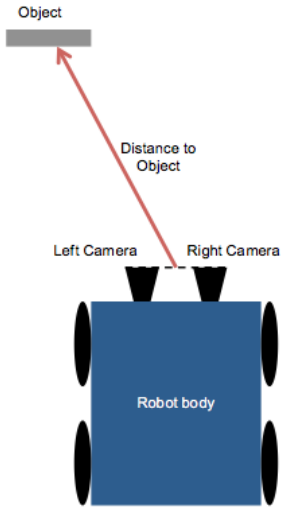

Figure 1: Top view of robot and object. Not to scale.

Two frames from the camera pair are shown below. The object is 30 pixels left of the center of the image taken by the left camera and 50 pixels to the left of the center of the image taken from the right camera. The resolution of the camera is 20 pixels per cm.

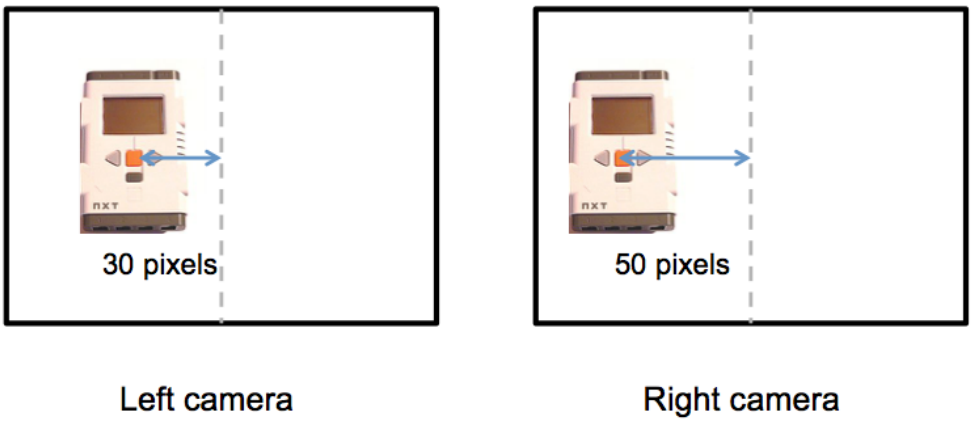

Figure 2: Image from left and right cameras.

In your writeup, find the distance from the center of the front of the lenses to

the center of the front of the object (realize that this is not just a single-dimensional problem). Show your work. At a minimum, write down two different symbolic equations showing the geometry used to determine the distance.

6 What To Submit

Sumbissions are due on Autolab by the date specified in the Syllabus.

1. A .pdf file with the written answers to Setion 2, the images and text to accompany Section 3 and the math and answer to Sections 4 and 5.
2. Your convolve.m function for verification. The requirements for this function are specified in 3.