Homework Set Four
ECE 161C
Department of Electrical and Computer Engineering
University of California, San Diego

Nuno Vasconcelos

1. In this question we study an example of why edge detection is such a hard problem. Figure 1 presents an imaging scenario, where a camera faces a wall at 90^0. The wall contains a door, and both the wall and the door can be considered Lambertian surfaces and infinite in extent (i.e. in the figure we only see the part of the surfaces that are projected on the image plane and, beyond that, can only assume that the surfaces extend infinitely). The door has albedo $b = 2a$ where a is the albedo of the wall. It is usually closed, as shown in the left, but sometimes moves to the open position (shown on the right). The angle between the open and closed positions is θ, as shown on the figure. To make everything simpler, we assume that the scene is lit by a point source at infinity, the light travels in the direction of the optical axis of the camera, i.e. the direction of vector $\mathbf{v}$, and has power $E = 1$.

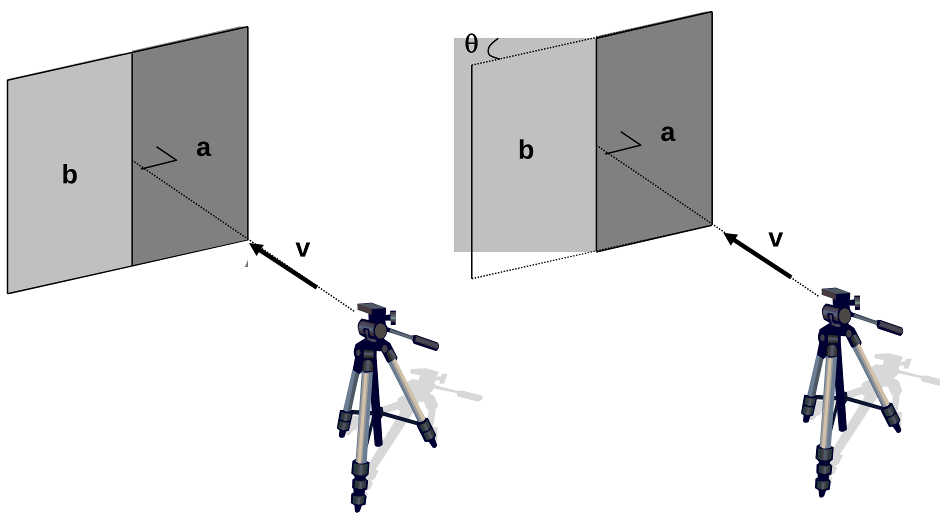

Figure 1: Left: Closed door. Right: Open door.

a) The image obtained, in both cases, will consist of two vertical rectangular regions side-by-side. What is the ratio between the image intensities in the region where the door is projected and the region where the wall is projected? Are there any changes from the scenario where the door is closed to that where it is open? If not, explain why. If yes, how does the ratio change from one situation to the other?

b) Consider the filter

$$\begin{bmatrix} -1 & 1 \\ -1 & 1 \\ -1 & 1 \end{bmatrix}.$$

What is the value of θ for which the convolution of the image of the open door with this filter is zero everywhere? What does that mean in terms of the performance of any edge detector on this scene?

(Note: assume that the imaging process introduces no noise, and disregard non-zero values that may be due to image processing artifacts around the image borders.)

c) we now show that this problem is not particular to the location of the light source that we have chosen. Assume that the light source was rotated by $\alpha = 30$ degrees as shown in Figure 2. What is the new value of θ for which the convolution with the filter of **b)** is zero?

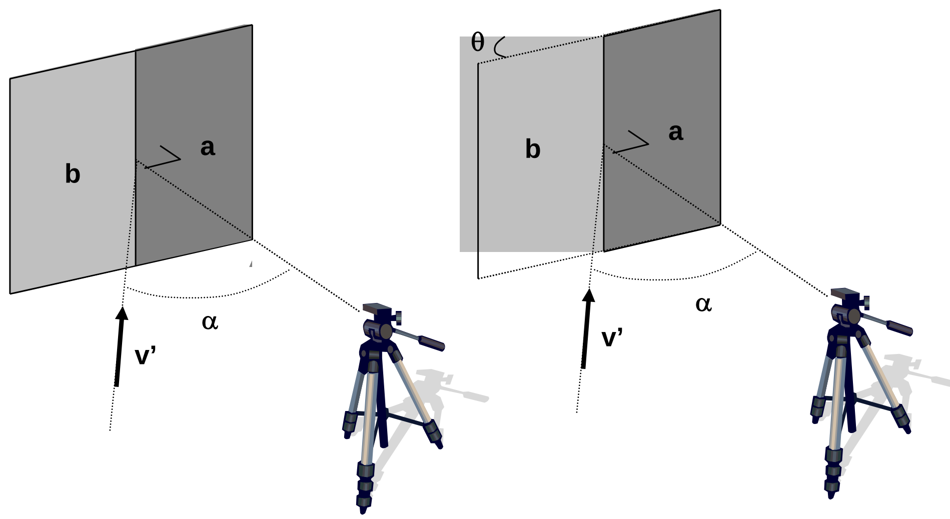

Figure 2: Left: Closed door. Right: Open door.

2. In this problem we explore the interpretation of filters as templates. For simplicity, we start by considering the 1D case. Suppose that we have two signals $x(k)$ and $h(k)$ of finite support (i.e. equal to zero for k outside an interval of length N) that are identical up to a shift (i.e. both x and h are different than zero only for a finite number of ks and $x(k) = h(k-n)$ for some n). One possibility to determine the shift n is to minimize the squared error between the signals, i.e. determine the value of n that minimizes

$$\mathcal{E}(n) = \sum_{k=-\infty}^{\infty} [x(k) - h(k-n)]^2.$$

a) show that this is equivalent to finding the n that maximizes the correlation of x and h

$$\mathcal{C}(n) = \sum_{k=-\infty}^{\infty} x(k)h(k-n).$$

When is this the same as maximizing the convolution of x and h?

b) suppose that x is a shifted and truncated replica of the step function, i.e.

$$x(k) = \begin{cases} 1, & n \leq k < n+l \\ 0, & \text{otherwise} \end{cases}$$

and we want to determine the shift n by convolving x with a filter h and locating the maximum of the convolution. Design an optimal detector h, i.e. a filter for which the maximum of the convolution indeed occurs at location n. Can you give a simple rule for designing the optimal h for an arbitrary signal x?

c) this rule generalizes to 2D. To check this convolve the image `square.bmp` with the following filters (use the function `conv2`)

$$h_1 = \begin{bmatrix} -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \end{bmatrix} \quad h_2 = \begin{bmatrix} -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \end{bmatrix} \quad h_3 = \begin{bmatrix} 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \end{bmatrix}.$$

Hand in plots of the filtered images. Can you explain their appearance? What are the filters detecting?

3. (MATLAB) in this problem we gain some familiarity with the basic operations required for edge detection. We will use the image `square.bmp`. and the following piece of code to synthesize a Gaussian filter of variance $\sigma^2 =$ `s`,

```
x = -32:32;
g = exp(-1/(2*s)*x'.^2)*exp(-1/(2*s)*x.^2); .
```

a) using `filter2`, filter the image with Gaussian filters of variance $\sigma^2 \in \{1, 2, 4, 16\}$. For each case, hand in a plot(`imagesc`) of both the filter and the filtered image. Explain the appearance of the latter.

b) using `gradient`, compute the derivatives of each filtered image. Hand in a plot of the derivatives for each case. Explain the appearance of the derivative images. Why do we have white and black bars? What are the differences between the derivatives obtained with different values of σ^2?

c) using `imagesc` and `quiver` create, for each σ^2, a plot containing: 1) the filtered image, and 2) the gradient field superimposed on it. Hand in the plots, and explain the observed gradient field.

d) hand in a plot of the magnitude of the gradient for each σ^2. In each case, find a threshold and threshold the gradient magnitude to produce an *edge map*. Hand in plots of the resulting edge maps, and comment on the role of smoothing with respect to the ability to obtain edge maps that are 1) spatially localized (i.e. such that each edge is only one pixel wide) and 2) continuous (i.e. without gaps along the recovered contours).