

COM S 327, Spring 2019

Programming Project 0

Image Processing

One of the most important operations in image processing and computer vision is edge detection. A very simple and effective edge detector is the Sobel Filter. The Sobel Filter is a pair of 3×3 matrices which are convolved with the input image separately then recombined. Specifically, the Sobel Filter is given by the pair of discrete convolutions:

$$O_x = \begin{bmatrix} -1 & 0 & +1 \\ -2 & 0 & +2 \\ -1 & 0 & +1 \end{bmatrix} \otimes I$$
$$O_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ +1 & +2 & +1 \end{bmatrix} \otimes I$$

where I is the input image, O_x and O_y are the piecewise output images, and \otimes is the convolution operator (described below). The x and y components are recombined with:

$$O = \sqrt{O_x^2 + O_y^2}$$

Given a convolution matrix (a *kernel*) of size $n \times n$ and a matrix M , the convolution $M' = K \otimes M$ is given by:

$$M'_{x,y} = \sum_{i=1}^n \sum_{j=1}^n K_{i,j} \times M_{x+(i-\lceil \frac{n}{2} \rceil), y+(j-\lceil \frac{n}{2} \rceil)}$$

That may look like some scary linear algebra, but it's actually very simple. Here's some pseudocode:

```
for each row r in M
  for each column c in M
    accumulator = 0
    for each row j in K
      for each column i in K
        accumulator = accumulator +
          K[j][i] * M[r + (j - ceil(n/2))][c + (i - ceil(n/2))]
    M'[r][c] = accumulator
```

Positions in the matrix where the kernel only partially covers the matrix (e.g., the edges) have to be handled specially. For our purposes, we'll ignore those cells and simply assign 0 (zero) to the output.

This YouTube video gives some visual examples of how convolution works¹: https://www.youtube.com/watch?v=C_zFhWdM4ic

I have provided source code that implements image reading and writing (and shows example usage). Starting with that code, write a program that takes the name of a PGM image on the command line, reads the image, applies a Sobel filter, and write the edge-detected image to disc with the file name `sobel.pgm`. All input files will be grayscale PGM images of size 1024×1024 .

PGM is a very, very simple image format. Tools that can display PGM images in UNIX and UNIX-like environments include *xv* and *gimp*. In Windows, *IrfanView* can do the job.

¹You can ignore the division step described in the video, because our kernels sum to zero so the division is undefined (and unnecessary).