Canny Edge Detection steps

The process of edge extraction is formed by the following (general) steps:

- Convolution of the image with edge enhancing masks in the x and y directions.

- Computation of image gradient (magnitude and direction)

- Thresholding of the gradient image

- Non-maximum edgel suppression

Image Gradient

The convolution of the image with derivative filters in the x and y directions yields the x and y components of the image gradient $\nabla I = [I_x I_y] = A(x, y)e^{-i(\theta(x,y))}$.

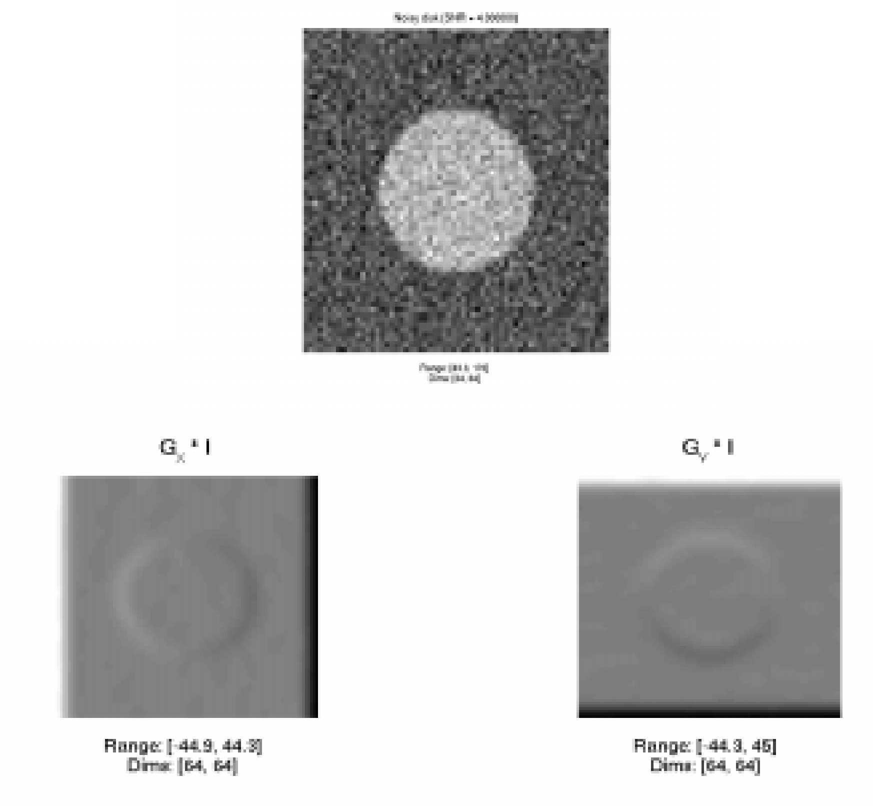

The amplitude and orientation of the gradient are computed directly from the image derivatives along x and y.

Gradient magnitude

The magnitude of the gradient at pixel $p(x, y)$ is given by $A(x, y) = \sqrt{I_x^2 + I_y^2}$, it gives an indication of the strength of a possible edge at $p(x, y)$.

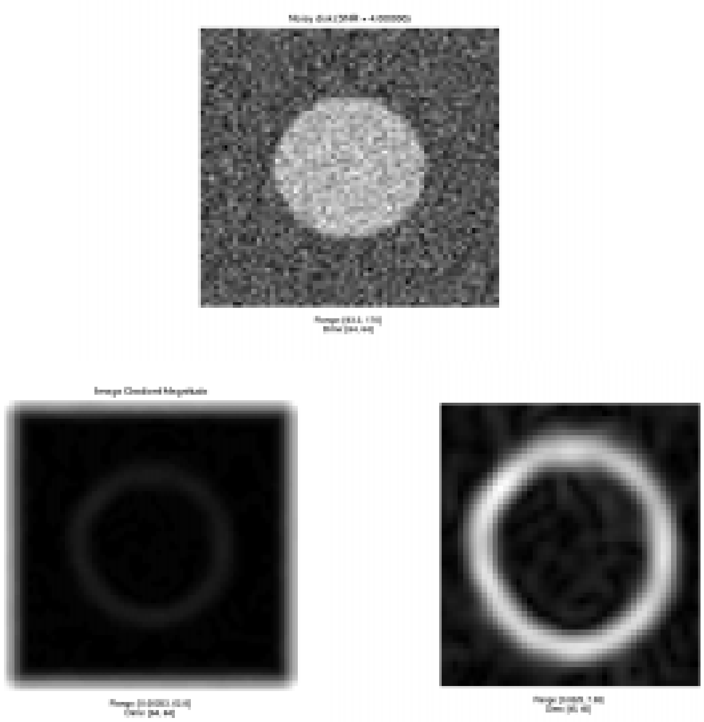

The figure on the lower-right shows a section of the gradient magnitude image on the lower-left, with pixel brightness rescaled to better show gradient magnitude around the circle.

Gradient thresholding

Once the gradient magnitude has been computed, a threshold is applied to remove all weak responses due to noise.

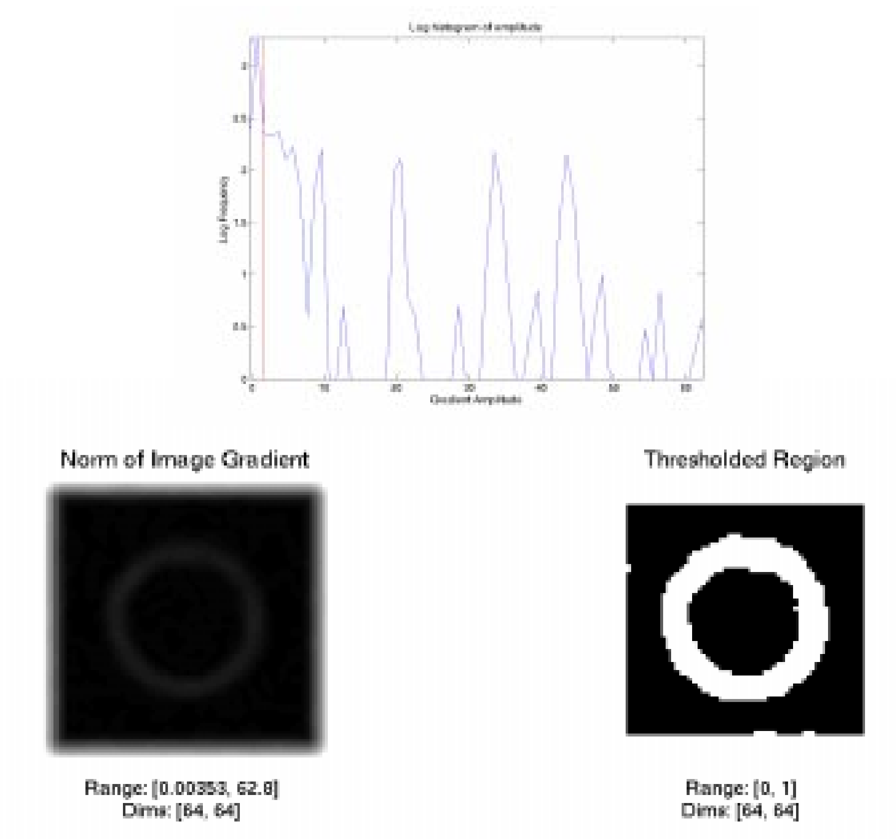

The resulting binary image is where the actual search for edgels takes place.

Gradient orientation

Besides the thresholded gradient magnitude map, we require an estimate of the gradient direction at each pixel, $\theta(x, y) = arctan(I_y/I_x)$. The gradient direction is always perpendicular to the direction of an edge passing through $p(x, y)$.

The gradient orientation is only computed at image locations that passed the gradient magnitude threshold.

Non-Maximal Suppression

Real edges correspond to places where the gradient magnitude is maximal, other locations along the gradient direction, with non-zero but not maximal responses must be discarded.

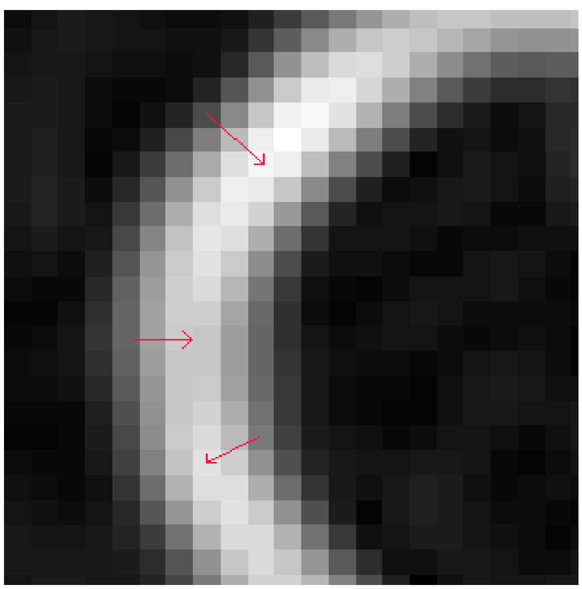

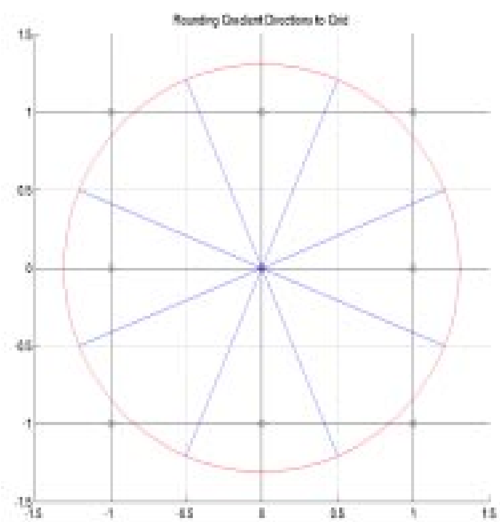

Non maximal suppression looks along the gradient direction, suppressing any edgel locations with non-maximal response.

The final edge image

Edgels that remain after non-maximal suppression make up the final edge map.

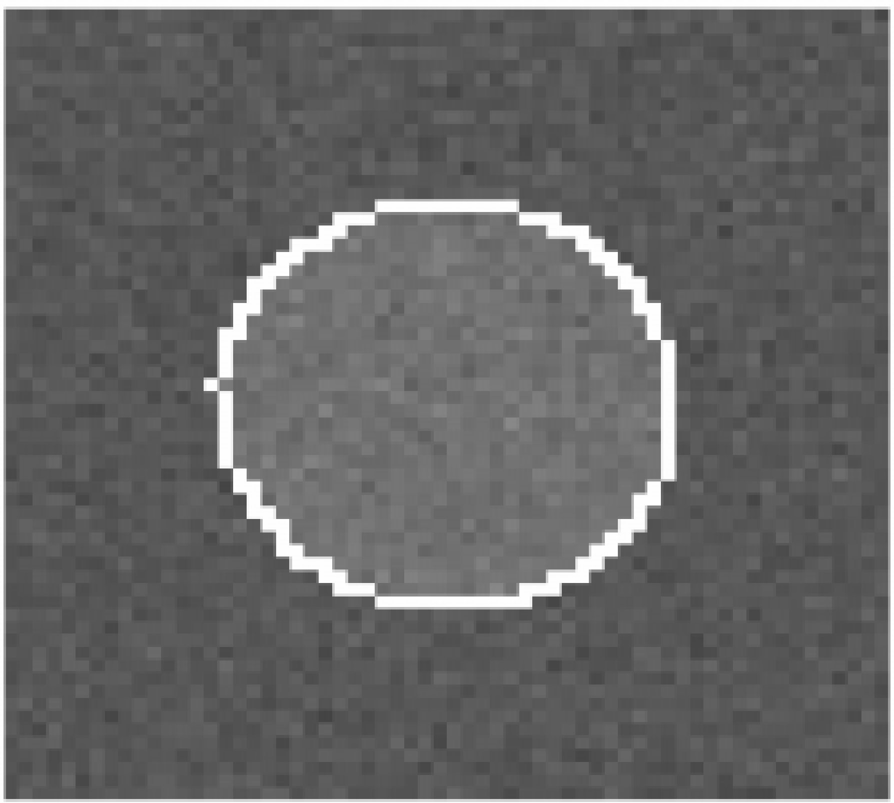

Choice of sigma and level of detail

Smaller sigma values cause the derivative filters to respond to smaller features, but also make the filters more sensitive to noise. Conversely, larger sigma values decrease localization accuracy.

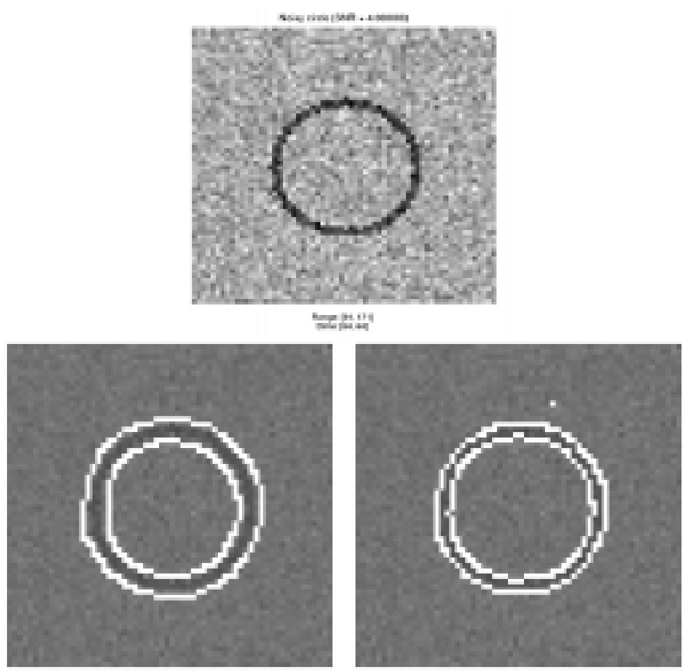

The edges on the lower-left correspond to $\sigma = 2$, the edges on the lower right correspond to $\sigma = 1$. The value of σ determines the scale of the edges that are detected.

Choice of sigma and level of detail

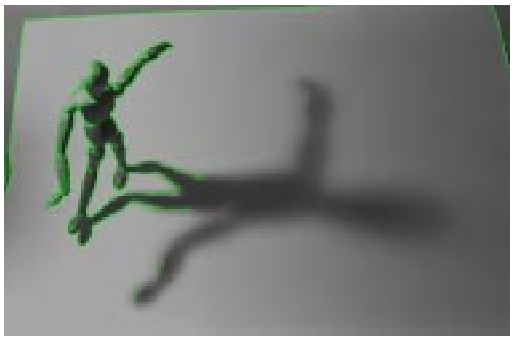

$\sigma=1$

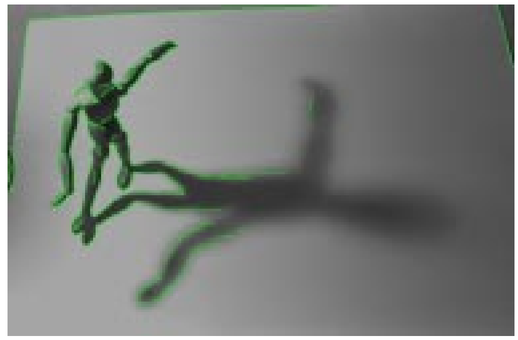

$\sigma=2$

$\sigma=4$

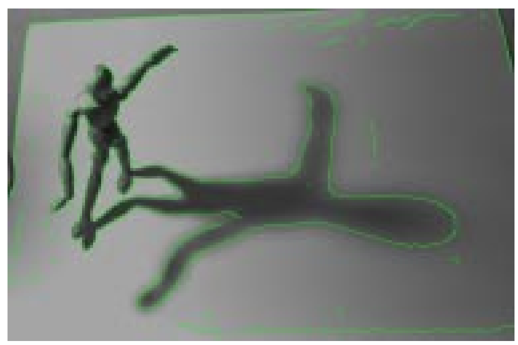

$\sigma=6$

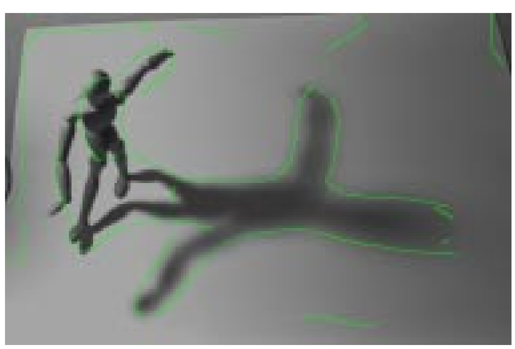

$\sigma=20$

Other edge extraction operators

Instead of using a Gaussian and its derivatives, we could use other standard filter masks such as those defined by Sobel or Roberts.

Canny

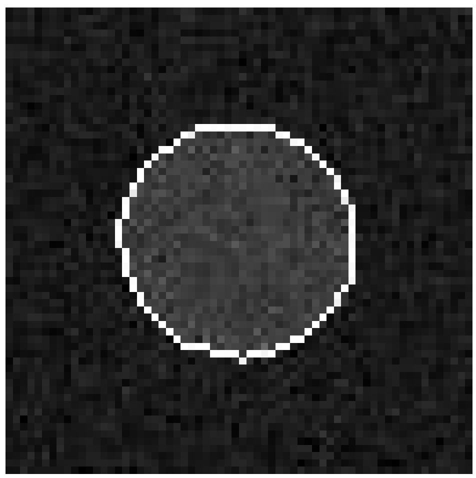

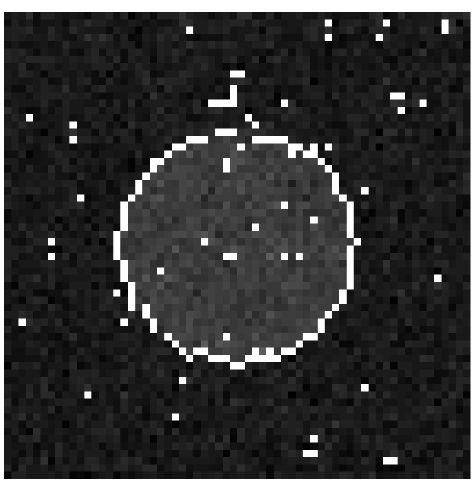

Sobel, [1 2 1; 0 0 0; −1 −2 −1] and its transpose

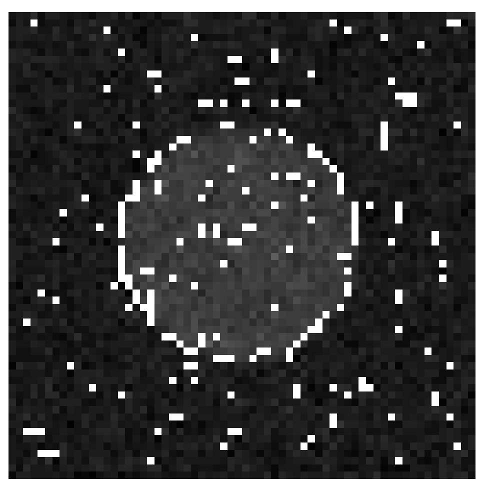

Roberts, [1 0; 0 −1] and its transpose

Common issues with Canny edges

- Threshold selection
- False positives
- False negatives
- Double edges
- How to determine edge saliency?

Threshold selection

A small change in threshold can make a large difference in the resulting edgel map.