

A Comprehensive Guide To Canny Edge Detection In Python

Tyler Jones - tyler.jones@byu.edu

December 26, 2019

Abstract

Canny Edge Detection is an algorithm to minimize noise and identify edges in a picture. Through this process one can reduce the amount of data in an image, while still maintaining its general shape and form. Edge detection is often just one of many filters used in image processing, and this paper will look to guide users in implementing their own Canny Edge Detection in Python.

Introduction & Prerequisites

Canny Edge Detection utilizes convolutions and intermediate level programming, as well as basic concepts in linear algebra. It is highly recommended that users are familiar with the Python Libraries *Numpy*, *Pillow*, and *Matplotlib*.

There are 5 primary steps to implement this algorithm: smoothing, estimating gradients, non-maximum suppression, double thresholding, and edge tracking. Each step will be described with visuals in detail below, along with code samples.

For those unfamiliar, a brief example of a convolution K being applied to a matrix A is given below. Note that with convolutions the resulting image is slightly smaller, just by definition of the convolution. Mathematically this would be written as:

$$A * K = A'$$

$$\underbrace{\begin{pmatrix} 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}}_A * \underbrace{\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}}_K = \underbrace{\begin{pmatrix} 1 & 4 & 3 & 4 & 1 \\ 1 & 2 & 4 & 3 & 3 \\ 1 & 2 & 3 & 4 & 1 \\ 1 & 3 & 3 & 1 & 1 \\ 3 & 3 & 1 & 1 & 0 \end{pmatrix}}_{A'}$$

Step 1 - Smoothing

The motivation for this first step is to remove small amounts of noise from our image, which should ensure that we only pick up on legitimate edges. To smooth or slightly blur the image we use a 5×5 Gaussian Kernel.

$$S = \frac{1}{159} \begin{pmatrix} 2 & 4 & 5 & 4 & 2 \\ 4 & 9 & 12 & 9 & 4 \\ 5 & 12 & 15 & 12 & 5 \\ 4 & 9 & 12 & 9 & 4 \\ 2 & 4 & 5 & 4 & 2 \end{pmatrix}$$

Original Image	Blurred Image
	

Figure 1: An example of a Gaussian Kernel is found on the left, and its implementation on a simple graphic is shown on the right.

```
1 def Gaussian_Smoothing(K, A):
2     """ Given a grayscale matrix A and kernel K, apply
3         and return the kernel.
4     """
5
6     # Get bounds of image and check kernel
7     n = K.shape[0]
8     assert n % 2 == 1
9
10    # Find bounds for our kernel
11    offset_ = int(np.floor(n/2))
12    lim_x, lim_y = A.shape
13    B = np.copy(A)
14
15    # Apply our kernel across our image
16    for i in range(offset_, lim_x-offset_):
17        for j in range(offset_, lim_y-offset_):
18            # Get the bounds of our convolution
19            x_l, x_r = i-offset_, i+offset_+1
20            y_l, y_r = j-offset_, j+offset_+1
21            A_ = A[x_l:x_r, y_l:y_r]
22            # Apply and store our convolution
23            L = np.multiply(A_, K)
24            B[i, j] = np.sum(L.flatten())
25
26    return B
```

Listing 1: Applying a Gaussian Kernel

Utilizing the example code above and any gray-scale image, the resulting image will be slightly more softened. Note that edges are smoothed, yet the general structure of the image still remains unchanged. This process can be done multiple times to further soften the image if needed.

Step 2 - Estimating Gradients

Next we look at estimating the strength of edges among each pixel. We do this by again using convolutions, but this time we use the *Sobel-operator*. At each pixel of our now smoothed image, use each respective kernel to estimate the gradient or edge strength in the x and y direction.

$$G_x = \begin{pmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{pmatrix} \quad G_y = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix}$$

Figure 2: Kernels used to find the estimated gradients in the x and y direction

At each pixel i,j , calculate $G_x^{i,j}$ and $G_y^{i,j}$, and then define the total gradient or strength of said pixel by the following equation:

$$G^{i,j} = \min(255, \sqrt{(G_x^{i,j})^2 + (G_y^{i,j})^2})$$

In addition to estimating the strength of the gradient at each pixel, we also store the direction in a separate matrix T . To do this we use the *atan2* function, and then round it to the nearest $\frac{\pi}{4}$ for simplicity. Thus the direction of the gradient at pixel i, j is defined as:

$$T_{i,j} = \text{atan2}\left(\frac{G_y^{i,j}}{G_x^{i,j}}\right)$$

Notice from the figures below that some noise may be present in both the pixel gradients or the gradient direction. This is to be expected, and can be accounted for in our next steps.

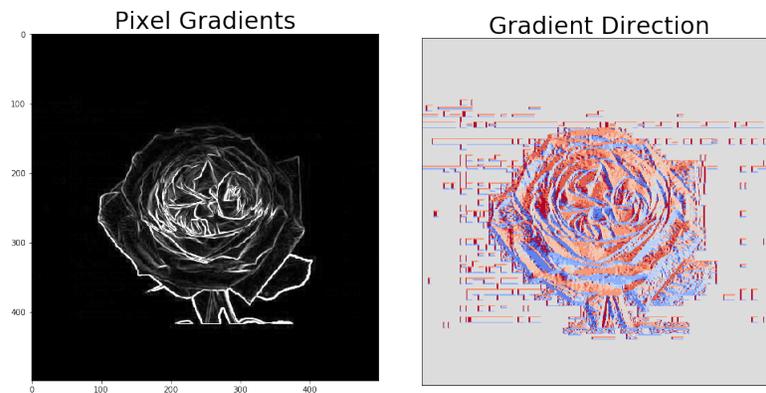


Figure 3: On the left we have our matrix G of gradient strengths, and on the right we see our matrix T of the direction of the gradient in radians.

```

1 def gradient_magnitudes(K_X, K_Y, A):
2     """ Given a matrix A and two kernels K_X, K_Y,
3         run the kernels on our matrix A.
4     """
5
6     # Get bounds of image and check kernel
7     n = K_X.shape[0]
8     assert n % 2 == 1
9
10    offset_ = int(np.floor(n/2))
11    lim_x, lim_y = A.shape
12
13    # Hold edge strengths and edge direction
14    G = np.zeros((A.shape[0]-offset_, A.shape[1]-offset_))
15    T = np.zeros((A.shape[0]-offset_, A.shape[1]-offset_))
16
17    # Get our 8 directions that our gradient can be going
18    neighborhood = [0, np.pi/4, np.pi/2, 3*np.pi/4, np.pi,
19                  -np.pi/4, -np.pi/2, -3*np.pi/4, -np.pi]
20
21    # Go through our whole image
22    for i in range(offset_, lim_x-offset_):
23        for j in range(offset_, lim_y-offset_):
24
25            # Get our kernel indices
26            x_l, x_r = i-offset_, i+offset_+1
27            y_l, y_r = j-offset_, j+offset_+1
28            A_ = A[x_l:x_r, y_l:y_r]
29
30            # Calculate our kernel and respective values
31            X_res = np.multiply(A_, K_X)
32            Y_res = np.multiply(A_, K_Y)
33            G_y = np.sum(Y_res.flatten())
34            G_x = np.sum(X_res.flatten())
35
36            # Calculate our edge strength
37            G[i, j] = min(np.sqrt(G_y**2+G_x**2), 255)
38
39            # Use atan2 to calculate our gradient direction
40            direc_ = math.atan2(G_y, G_x)
41
42            # Fit edge direction to closest pi/4 angle
43            closest = neighborhood[np.argmin([np.abs(direc_-d)
44                                             for d in neighborhood])]
45
46            T[i, j] = closest
47
48    return G, T

```

Listing 2: Example code to find the gradient magnitudes and gradient directions for each pixel in an image.

Step 3 - Non-Maximum Suppression

In order to reduce noise and keep only the strongest edges, we apply non-maximum suppression to our entire image. Given our gradient magnitude matrix G , and gradient directional matrix T we apply this step pixel by pixel to G .

Recall that $T_{i,j}$ contains the direction of the gradient of the i,j 'th pixel of G , rounded to the nearest $\frac{\pi}{4}$. To perform non-maximum suppression on the (i,j) pixel in G , follow the following steps:

1. Find the direction of the gradient of said pixel, given in $T_{i,j}$.
2. Compare the value of $G_{i,j}$ to the values of G immediately parallel along the direction of the gradient.
3. If $G_{i,j}$ is not the largest value, then it is not the strongest edge and can be discarded, set $G_{i,j} = 0$. Otherwise if it is the largest value, retain it.

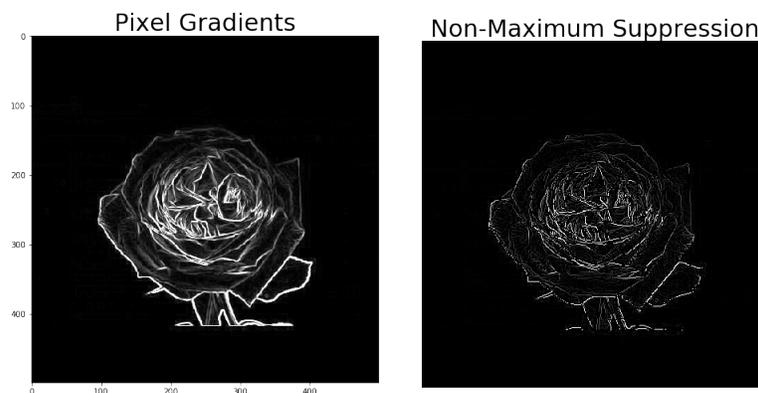


Figure 4: On the left we again have our matrix G of gradient strengths, and on the right we see our modified matrix G' where only the strongest edges are kept. This greatly reduces noisy edges and keeps only the strongest edge signals.

```
1 def non_max_suppression(G, T):
2     """ Given an image matrix A, along with a matrix of it's
3         magnitude direction we identify and keep only the
4         strongest edges.
5     """
6
7     assert G.shape == T.shape
8
9     # Iterate through our whole image
10    for i in range(1,G.shape[0]-1):
11        for j in range(1,G.shape[1]-1):
```

```

12
13     # Get direction of gradient in specific cell
14     direc_ = np.abs(T[i,j])
15     val_ = G[i,j]
16
17     if direc_ == np.pi/4:
18         """ Look in the upper right and bottom left
19             direction
20         """
21         if val_ > G[i-1, j+1] and val_ > G[i+1, j-1]:
22             pass
23         else:
24             G[i,j] = 0
25
26     elif direc_ == np.pi/2:
27         """ Look in the up/down direction """
28         if val_ > G[i-1, j] and val_ > G[i+1, j]:
29             pass
30         else:
31             G[i,j] = 0
32
33     elif direc_ == 3*np.pi/4:
34         """ Look in the top left and bottom right
35             direction
36         """
37         if val_ > G[i-1, j-1] and val_ > G[i+1, j+1]:
38             pass
39         else:
40             G[i,j] = 0
41
42     elif direc_ == np.pi or direc_ == 0:
43         """ Look in the left and right directions """
44         if val_ > G[i, j-1] and val_ > G[i, j+1]:
45             pass
46         else:
47             G[i,j] = 0
48
49     return G

```

Listing 3: Code to perform Non-Maximum Suppression on our target image G .

Step 4 - Double Thresholding

Now we have an array G where only the strongest edges are retained. However the strength of each of those pixels has a wide distribution. To remedy this we classify each pixel as either *strong*, *weak* or *noise*. This step is aptly named as to do this we simply use a double threshold with values θ_l and θ_h where $\theta_l < \theta_h$. To determine the value of $G_{i,j}$:

$$G_{i,j} = \begin{cases} 0 & G_{i,j} < \theta_l \\ 128 & \theta_l \leq G_{i,j} \leq \theta_h \\ 255 & \theta_h < G_{i,j} \end{cases}$$

The choices for our threshold θ_l and θ_h are hyper-parameters that must be chosen wisely for the image on hand. It is recommended to try many different values and identify what works best. Images with a lot of noise may require higher filters than average in order to best discern actual edges.

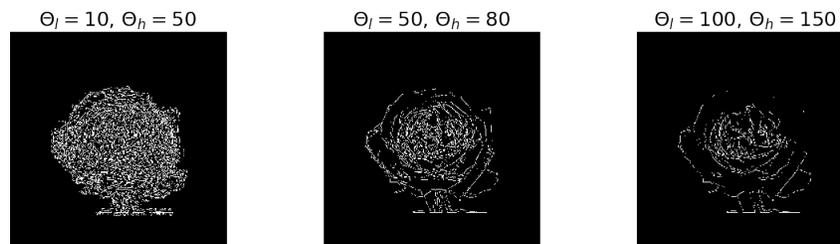


Figure 5: Example of how different threshold values alter the output of our image. Note that the lower θ_l is, the more noise is classified as an actual edge and kept in our image.

```

1 def double_thresh(H, L, G):
2     """ Given a matrix G, and hyper-parameters H and L,
3         performs double threshold on G.
4     """
5     A = np.copy(G)
6
7     A[A > H] = 255
8     A[A < L] = 0
9     A[(A > L) & (A < H)] = 128
10
11    return A

```

Listing 4: Code to perform a double threshold on our target image G .

Step 5 - Edge Tracking

For our final step we look to discern which weak edges to keep, and which to discard. To do this we retain all of the strong edges, and then look at every weak edge. If the weak edge is connected to a strong edge in its direct neighborhood, then we retain it. Otherwise we discard it as noise.

```

1 def edge_track(G):
2     """ Given a matrix G, remove all weak edges that are not
3         connected to a strong edge.
4     """
5
6     for i in range(1, G.shape[0]-1):
7         for j in range(1, G.shape[1]-1):
8             if G[i, j] == 128:
9                 # Examine the neighborhood of a weak edge

```

```
10         N = (G[i-1:i+1, j-1:j+1]).flatten()
11
12         if 255 not in N:
13             G[i,j] = 0
14     else:
15         pass
```

Listing 5: Code to implement edge tracking on our G .

Conclusion

Through the process of smoothing, estimating gradients, non-maximum suppression, double thresholding, and edge tracking, you are able to reduce the amount of data in an image and identify edges. Canny Edge Detection is just one of many edge detection methods, but understanding how to implement it is a great way to learn more about computer vision.

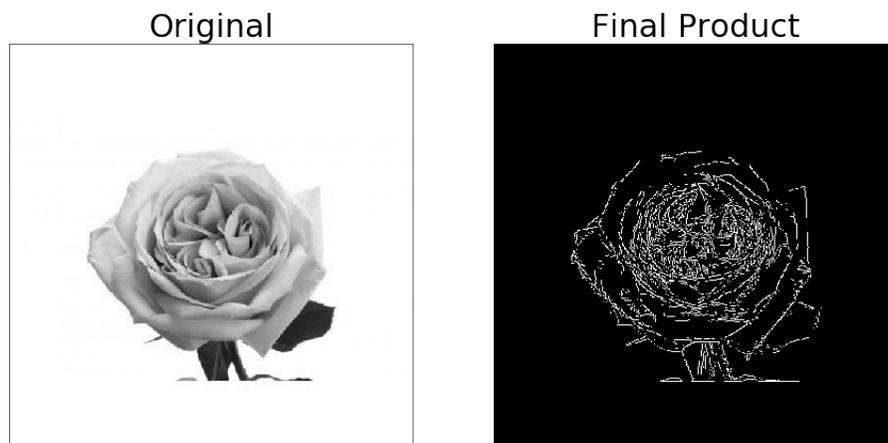


Figure 6: The full implementation of our Canny Edge Detection Algorithm