



# Image Filtering

Image filtering is used to:

- Remove noise
- Sharpen contrast
- Highlight contours
- Detect edges
- Other uses?

Image filters can be classified as linear or nonlinear.

Linear filters are also known as convolution filters as they can be represented using a matrix multiplication.

Thresholding and image equalisation are examples of nonlinear operations, as is the median filter.



# Median Filtering

Median filtering is a nonlinear method used to remove noise from images.

It is widely used as it is very effective at removing noise while preserving edges.

It is particularly effective at removing 'salt and pepper' type noise.

The median filter works by moving through the image pixel by pixel, replacing each value with the median value of neighbouring pixels.

The pattern of neighbours is called the "window", which slides, pixel by pixel, over the entire image.

The median is calculated by first sorting all the pixel values from the window into numerical order, and then replacing the pixel being considered with the middle (median) pixel value.



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Median Filtering example

The following example shows the application of a median filter to a simple one dimensional signal.

A window size of three is used, with one entry immediately preceding and following each entry.

Window for  $x[6] \rightarrow y[6]$

x =	3	9	4	52	3	8	6	2	2	9
-----	---	---	---	----	---	---	---	---	---	---

$$y[1] = \text{median}[3 \ 3 \ 9] = 3$$

$$y[2] = \text{median}[3 \ 4 \ 9] = 4$$

$$y[3] = \text{median}[4 \ 9 \ 52] = 9$$

$$y[4] = \text{median}[3 \ 4 \ 52] = 4$$

$$y[5] = \text{median}[3 \ 8 \ 52] = 8$$

$$y[6] = \text{median}[3 \ 6 \ 8] = 6$$

$$y[7] = \text{median}[2 \ 6 \ 8] = 6$$

$$y[8] = \text{median}[2 \ 2 \ 6] = 2$$

$$y[9] = \text{median}[2 \ 2 \ 9] = 2$$

$$y[10] = \text{median}[2 \ 9 \ 9] = 9$$

y =	3	4	9	4	8	6	6	2	2	9
-----	---	---	---	---	---	---	---	---	---	---



# Median Filtering

In the previous example, because there is no entry preceding the first value, the first value is repeated (as is the last value) to obtain enough entries to fill the window.

What effect does this have on the boundary values?

There are other approaches that have different properties that might be preferred in particular circumstances:

- Avoid processing the boundaries, with or without cropping the signal or image boundary afterwards.
- Fetching entries from other places in the signal. With images for example, entries from the far horizontal or vertical boundary might be selected.
- Shrinking the window near the boundaries, so that every window is full.


What effects might these approaches have on the boundary values?



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Median Filtering

On the left is an image containing a significant amount of salt and pepper noise. On the right is the same image after processing with a median filter.



Notice the well preserved edges in the image.



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

## Median Filtering example 2

2D Median filtering example using a 3 x 3 sampling window:

Sorted: 0,0,1,1,1,2,2,4,4

Input

1	4	0	1	3	1
2	2	4	2	2	3
1	0	1	0	1	0
1	2	1	0	2	2
2	5	3	1	2	5
1	1	4	2	3	0

Output

	1				



# Average Filtering

Average (or mean) filtering is a method of 'smoothing' images by reducing the amount of intensity variation between neighbouring pixels.

The average filter works by moving through the image pixel by pixel, replacing each value with the average value of neighbouring pixels, including itself.

There are some potential problems:

- A single pixel with a very unrepresentative value can significantly affect the average value of all the pixels in its neighbourhood.
- When the filter neighbourhood straddles an edge, the filter will interpolate new values for pixels on the edge and so will blur that edge. This may be a problem if sharp edges are required in the output.



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Average Filtering

The following example shows the application of an average filter to a simple one dimensional signal.

A window size of three is used, with one entry immediately preceding and following each entry.

Window for  $x[4] \rightarrow y[4]$

x=	3	9	4	52	3	8	6	2	2	9
----	---	---	---	----	---	---	---	---	---	---

$$y[1] = \text{round}((3+3+9)/3) = 5$$

$$y[2] = \text{round}((3+9+4)/3) = 5$$

$$y[3] = \text{round}((9+4+52)/3) = 22$$

$$y[4] = \text{round}((4+52+3)/3) = 20$$

$$y[5] = \text{round}((52+3+8)/3) = 21$$

$$y[6] = \text{round}((3+8+6)/3) = 6$$

$$y[7] = \text{round}((8+6+2)/3) = 5$$

$$y[8] = \text{round}((6+2+2)/3) = 3$$

$$y[9] = \text{round}((2+2+9)/3) = 4$$

$$y[10] = \text{round}((2+9+9)/3) = 7$$


y=	5	5	22	20	21	6	5	3	4	7
----	---	---	----	----	----	---	---	---	---	---





THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Filter Comparison



The graph above shows the 1D signals from the median and average filter examples.

What are the differences in the way the filters have modified the original signal?



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

## Average Filtering example 2

2D Average filtering example using a 3 x 3 sampling window:

$$\text{Average} = \text{round}(1+4+0+2+2+4+1+0+1)/9 = 2$$

Input

1	4	0	1	3	
2	2	4	2	2	
1	0	1	0	1	
1	2	1	0	2	
2	5	3	1	2	

Output

	2				


10



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Average Filtering

On the left is an image containing a significant amount of salt and pepper noise. On the right is the same image after processing with an Average filter.



What are the differences in the result compared with the Median filter?  
Is this a linear (convolution) or nonlinear filter?

# 3 by 3 Median filtering example

0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	100	100	100
0	0	0	100	100	100
0	0	0	100	100	100


# 3 by 3 average filtering example

0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	0	0	0
0	0	0	100	100	100
0	0	0	100	100	100
0	0	0	100	100	100




THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Gaussian Filtering


Gaussian filtering is used to blur images and remove noise and detail. In one dimension, the Gaussian function is:

$$G(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}$$

Where  $\sigma$  is the standard deviation of the distribution. The distribution is assumed to have a mean of 0.

Shown graphically, we see the familiar bell shaped Gaussian distribution.

Gaussian distribution with mean 0 and  $\sigma = 1$





THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND


# Gaussian Filtering

When working with images we need to use the two dimensional Gaussian function.

This is simply the product of two 1D Gaussian functions (one for each direction) and is given by:

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

A graphical representation of the 2D Gaussian distribution with mean(0,0) and  $\sigma = 1$  is shown to the right.





THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Gaussian Filtering





# Convolution

The convolution of two functions  $f$  and  $g$  is defined as:

$$(f * g)(x, y) = \sum_{v=-\infty}^{\infty} \sum_{u=-\infty}^{\infty} f(u, v)g(x-u, y-v)$$

Where  $f(x, y)$  is a function that represents the image and  $g(x, y)$  is a function that represents the kernel.

In practice, the kernel is only defined over a finite set of points, so we can modify the definition to:

$$(f * g)(x, y) = \sum_{v=y-h}^{y+h} \sum_{u=x-w}^{x+w} f(u, v)g(x-u, y-v)$$

Where  $2w+1$  is the width of the kernel and  $2h+1$  is the height of the kernel.

$g$  is defined only over the points  $[-w, w] \times [-h, h]$



# Convolution Pseudocode

Pseudocode for the convolution of an image  $f(x,y)$  with a kernel  $g(x,y)$  to produce a new image  $h(x,y)$ :

```
for y = 0 to ImageHeight do  
  for x = 0 to ImageWidth do  
    sum = 0  
    for v = (y - h) to (y + h) do  
      for u = (x - w) to (x + w) do  
        sum = sum + f(u,v) * g(x - u, y - v)  
      end for  
    end for  
    h(x,y) = sum  
  end for  
end for
```



# Convolution – Potential Problems

Summation over a neighbourhood might exceed the range and/or sign permitted in the image format:

- The data may need to be temporarily stored in a 16 – 32 bit integer representation.
- Then normalised back to the appropriate range (0-255 for an 8 bit image).

Another issue is how to deal with image borders:

- Convolution is not possible if part of the kernel lies outside the image.
- What is the size of image window which is processed normally when performing a Convolution of size  $m \times n$  on an original image of size  $M \times N$  ?



# Convolution Border Issues

Two potential methods for dealing with convolution at image borders:

## 1) Reflected Indexing

Mirror the image outside the borders.

For x coordinate of image with width M:

```
if x < 0 then
  x =
else if x >= M then
  x =
end if
```

## 2) Circular Indexing

Repeat image endlessly outside the borders.

For x coordinate of image with width M:

```
if x < 0 then
  x =
else if x >= M then
  x =
end if
```



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Convolution Border example

Find the corner and border specific kernel for:

1	1	1
1	1	1
1	1	1

-1	-1	-1
-1	8	-1
-1	-1	-1

-1	0	1
-1	0	1
-1	0	1



# Edge Detection

Edges in images are areas with strong intensity contrasts; a jump in intensity from one pixel to the next.

The process of edge detection significantly reduces the amount of data and filters out unneeded information, while preserving the important structural properties of an image.

There are many different edge detection methods, the majority of which can be grouped into two categories:

- Gradient,
- and Laplacian.

The gradient method detects the edges by looking for the maximum and minimum in the first derivative of the image.

The Laplacian method searches for zero crossings in the second derivative of the image .

We will look at two examples of the gradient method, Sobel and Prewitt.



# Sobel Filter

The Sobel filter is used for edge detection.

It works by calculating the gradient of image intensity at each pixel within the image. It finds the direction of the largest increase from light to dark and the rate of change in that direction.

The result shows how abruptly or smoothly the image changes at each pixel, and therefore how likely it is that that pixel represents an edge.

It also shows how that edge is likely to be oriented.

The result of applying the filter to a pixel in a region of constant intensity is a zero vector.

The result of applying it to a pixel on an edge is a vector that points across the edge from darker to brighter values.



# Sobel Filter

The sobel filter uses two 3 x 3 kernels. One for changes in the horizontal direction, and one for changes in the vertical direction. The two kernels are convolved with the original image to calculate the approximations of the derivatives.

If we define  $G_x$  and  $G_y$  as two images that contain the horizontal and vertical derivative approximations respectively, the computations are:

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

Where  $A$  is the original source image.

The  $x$  coordinate is defined as increasing in the right-direction and the  $y$  coordinate is defined as increasing in the down-direction.





# Sobel Filter

To compute  $G_x$  and  $G_y$  we move the appropriate kernel (window) over the input image, computing the value for one pixel and then shifting one pixel to the right. Once the end of the row is reached, we move down to the beginning of the next row.

The example below shows the calculation of a value of  $G_x$

$a_{11}$	$a_{12}$	$a_{13}$	...	
$a_{21}$	$a_{22}$	$a_{23}$	...	
$a_{31}$	$a_{32}$	$a_{33}$	...	
...	...	...	...	

Input image

$G_x =$

1	0	-1
2	0	-2
1	0	-1

$b_{11}$	$b_{12}$	$b_{13}$	...	
$b_{21}$	$b_{22}$	$b_{23}$	...	
$b_{31}$	$b_{32}$	$b_{33}$	...	
...	...	...	...	

Output image

$$b_{22} = (a_{11} * 1) + (a_{12} * 0) + (a_{13} * -1) + (a_{21} * 2) + (a_{22} * 0) + (a_{23} * -2) + (a_{31} * 1) + (a_{32} * 0) + (a_{33} * -1)$$



# Sobel Filter

At each pixel in the image, the gradient approximations given by  $G_x$  and  $G_y$  are combined to give the gradient magnitude, using:

$$G = \sqrt{G_x^2 + G_y^2}$$

The gradient's direction is calculated using:

$$\Theta = \arctan\left(\frac{G_y}{G_x}\right)$$

A  $\Theta$  value of 0 would indicate a vertical edge that is darker on the left side.



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Sobel Filter example



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Sobel Filter



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Sobel Filter



THE UNIVERSITY OF AUCKLAND  
NEW ZEALAND

# Prewitt Filter