

Lecture 03

Image Filtering

2026-03-06

Sébastien Valade



1. Introduction

2. Spatial domain filtering

3. Frequency domain filtering

Previous lecture: *Point Operations*

The image transformations discussed so far are based on the following expression:

$$g(x, y) = T[f(x, y)]$$

- $f(x, y)$ is an input image
- $g(x, y)$ is the output image
- T is an operator on f at point (x, y)

⇒ the operator T was applied to individual pixels ("Point Operations"), i.e. neighborhood = 1x1 pix
⇒ the function is an intensity transformation function, to change image brightness, etc.


Previous lecture: Point Operations

The image transformations discussed so far are based on the following expression:

$$g(x, y) = T[f(x, y)]$$

- $f(x, y)$ is an input image
- $g(x, y)$ is the output image
- T is an operator on f at point (x, y)

- ⇒ the operator T was applied to individual pixels ("Point Operations"), i.e. neighborhood = 1x1 pix
 ⇒ the function is an intensity transformation function, to change image brightness, etc.



Today: *Neighborhood Operations* $\Rightarrow$ **filtering**

$\Rightarrow$ Purpose: blur, sharpen, remove noise, filter frequencies, etc.

$\Rightarrow$ Approaches:

1. spatial domain filtering

- the neighborhood is >1 pixel ("Point Processing" $\rightarrow$ "Neighborhood Processing")
- spatial filtering modifies an image by replacing the value of each pixel by a function of the values of the pixel and its neighbors
- spatial filters are applied by convolution

2. frequency domain filtering

- the 2D direct Fourier transform is applied to extract image frequencies
- the amplitude spectrum can be band-passed to filter certain frequencies
- the inverse 2D direct Fourier transform is used to reconstruct the filtered image

Today: *Neighborhood Operations* $\Rightarrow$ **filtering**

$\Rightarrow$ Purpose: blur, sharpen, remove noise, filter frequencies, etc.

$\Rightarrow$ Approaches:

1. spatial domain filtering

- the neighborhood is >1 pixel ("Point Processing" $\rightarrow$ "Neighborhood Processing")
- spatial filtering modifies an image by replacing the value of each pixel by a function of the values of the pixel and its neighbors
- spatial filters are applied by convolution

2. frequency domain filtering

- the 2D direct Fourier transform is applied to extract image frequencies
- the amplitude spectrum can be band-passed to filter certain frequencies
- the inverse 2D direct Fourier transform is used to reconstruct the filtered image

Today: *Neighborhood Operations* $\Rightarrow$ **filtering**

$\Rightarrow$ Purpose: blur, sharpen, remove noise, filter frequencies, etc.

$\Rightarrow$ Approaches:

1. spatial domain filtering

- the neighborhood is >1 pixel ("Point Processing" $\rightarrow$ "Neighborhood Processing")
- spatial filtering modifies an image by replacing the value of each pixel by a function of the values of the pixel and its neighbors
- spatial filters are applied by convolution

2. frequency domain filtering

- the 2D direct Fourier transform is applied to extract image frequencies
- the amplitude spectrum can be band-passed to filter certain frequencies
- the inverse 2D direct Fourier transform is used to reconstruct the filtered image

1. Introduction

2. Spatial domain filtering

1. linear spatial filter
2. convolutions
3. kernels types and applications


3. Frequency domain filtering

Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t)$$

- $f(x, y)$ is a pixel of the input image at position (x, y)
- $g(x, y)$ is a pixel of the output image at position (x, y)
- w is the filter kernel of size (m, n)
where $a = (m - 1)/2$, $b = (n - 1)/2$




Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t)$$

- $f(x, y)$ is a pixel of the input image at position (x, y)
- $g(x, y)$ is a pixel of the output image at position (x, y)
- w is the filter kernel of size (m, n)
where $a = (m - 1)/2$, $b = (n - 1)/2$




Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t)$$

- $f(x, y)$ is a pixel of the input image at position (x, y)
- $g(x, y)$ is a pixel of the output image at position (x, y)
- w is the filter kernel of size (m, n)
where $a = (m - 1)/2$, $b = (n - 1)/2$




$$\begin{array}{r}
 (3 \cdot -1) \\
 + (0 \cdot 0) \\
 + (1 \cdot 1) \\
 + (2 \cdot -2) \\
 + (6 \cdot 0) \\
 + (2 \cdot 2) \\
 + (2 \cdot -1) \\
 + (4 \cdot 0) \\
 + (1 \cdot 1) \\
 \hline
 -3
 \end{array}$$

Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t)$$

- $f(x, y)$ is a pixel of the input image at position (x, y)
- $g(x, y)$ is a pixel of the output image at position (x, y)
- w is the filter kernel of size (m, n)
where $a = (m - 1)/2$, $b = (n - 1)/2$




Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

- kernel size (m,n) defines the *neighborhood* of operation on pixel at position (x,y)
- kernel coefficients define the *nature* of the filter

⇒ kernel *slides* across the input image to produce a *filtered* output image $g(x,y)$

- stride = sliding step (stride=1 ⇒ kernel will slide by 1 pixel per column/row at a time)




Linear spatial filter

⇒ sum-of-products operation between an input image $f(x,y)$ and a filter kernel w

- kernel size (m,n) defines the *neighborhood* of operation on pixel at position (x,y)
- kernel coefficients define the *nature* of the filter

⇒ kernel *slides* across the input image to produce a filtered output image $g(x,y)$


- stride = sliding step (stride=1 ⇒ kernel will slide by 1 pixel per column/row at a time)



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image


- padding = pad the image so the kernel can also operate on the edges ($\text{pad_size} = \text{kernel_size} // 2$)



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image


- padding = pad the image so the kernel can also operate on the edges ($\text{pad_size} = \text{kernel_size} / 2$)



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image

- padding = pad the image so the kernel can also operate on the edges ($\text{pad_size} = \text{kernel_size} // 2$)



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image

- padding types: modified after *Burger and Burge (2009)*

constant

⇒ *use constant value*



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image

- padding types: modified after *Burger and Burge (2009)*

constant

⇒ *use constant value*



clamp

⇒ *use edge values*



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image

- padding types: modified after *Burger and Burge (2009)*

constant

⇒ use constant value



clamp

⇒ use edge values



mirror

⇒ use mirrored edges



Linear spatial filter

⇒ input image needs *padding* in order for output image to keep the same size as the input image

- padding types: modified after *Burger and Burge (2009)*

constant

⇒ use constant value



clamp

⇒ use edge values



mirror

⇒ use mirrored edges



warp

⇒ use wrapped image



Linear spatial filter

⇒ the sum-of-products operation between the input image $f(x, y)$ and filter kernel w (eq.1) is the implementation of a spatial convolution (eq.2):

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t) \quad (1)$$

$$g = w * f \quad (2)$$

Linear spatial filter

⇒ the sum-of-products operation between the input image $f(x, y)$ and filter kernel w (eq.1) is the implementation of a spatial convolution (eq.2):

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t) \quad (1)$$

$$g = w * f \quad (2)$$

linear spatial filtering $\iff$ spatial convolution

Linear spatial filter

⇒ the sum-of-products operation between the input image $f(x, y)$ and filter kernel w (eq.1) is the implementation of a spatial convolution (eq.2):

$$g(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t) \quad (1)$$

$$g = w * f \quad (2)$$

linear spatial filtering $\iff$ spatial convolution

convolutions are the core operations used by Convolutional Neural Networks (CNN)

Kernel coefficients define the nature of the filter

⇒ vary kernels coefficients according to the desired filtering operation:

- **smoothing** filters

⇒ **low-pass** filters → retains low-frequency components of the image

- *averaging kernel (a.k.a. box filter)*
- *gaussian kernel*

- **sharpening** filters

⇒ **high-pass** filters → retains high-frequency components of the image

⇒ **edge detection** filters:

- *Sobel kernel, Prewitt kernel, etc.* → directional filters (sensitive to edge orientation)
- *Laplacian kernel* → isotropic filter (not sensitive to edge orientation)

⇒ **sharpening** filters: increase image contrast along edges

- **other**

- *Emboss filter* → appearance of the image being “embossed” or elevated from the background

Kernel coefficients define the nature of the filter

⇒ vary kernels coefficients according to the desired filtering operation:

- **smoothing** filters

⇒ **low-pass** filters → retains low-frequency components of the image

- *averaging* kernel (a.k.a. *box filter*)
- *gaussian* kernel

- **sharpening** filters

⇒ **high-pass** filters → retains high-frequency components of the image

⇒ **edge detection** filters:

- *Sobel* kernel, *Prewitt* kernel, etc. → directional filters (sensitive to edge orientation)
- *Laplacian* kernel → isotropic filter (not sensitive to edge orientation)

⇒ **sharpening** filters: increase image contrast along edges

- **other**

- *Emboss* filter → appearance of the image being “embossed” or elevated from the background

Kernel coefficients define the nature of the filter

⇒ vary kernels coefficients according to the desired filtering operation:

- **smoothing** filters

⇒ **low-pass** filters → retains low-frequency components of the image

- *averaging* kernel (a.k.a. *box filter*)
- *gaussian* kernel

- **sharpening** filters

⇒ **high-pass** filters → retains high-frequency components of the image

⇒ **edge detection** filters:

- *Sobel* kernel, *Prewitt* kernel, etc. → directional filters (sensitive to edge orientation)
- *Laplacian* kernel → isotropic filter (not sensitive to edge orientation)

⇒ **sharpening** filters: increase image contrast along edges

- **other**

- *Emboss* filter → appearance of the image being “embossed” or elevated from the background

Kernel coefficients define the nature of the filter

⇒ vary kernels coefficients according to the desired filtering operation:

- **smoothing** filters

⇒ **low-pass** filters → retains low-frequency components of the image

- *averaging* kernel (a.k.a. *box filter*)
- *gaussian* kernel

- **sharpening** filters

⇒ **high-pass** filters → retains high-frequency components of the image

⇒ **edge detection** filters:

- *Sobel* kernel, *Prewitt* kernel, etc. → directional filters (sensitive to edge orientation)
- *Laplacian* kernel → isotropic filter (not sensitive to edge orientation)

⇒ **sharpening** filters: increase image contrast along edges

- **other**

- *Emboss* filter → appearance of the image being “embossed” or elevated from the background


original image



identity

0	0	0
0	1	0
0	0	0

filtered image



original image



identity

0	0	0
0	1	0
0	0	0

⇒ no change!

filtered image




original image



average

0.1	0.1	0.1
0.1	0.1	0.1
0.1	0.1	0.1

filtered image



LOW-PASS FILTER

original image



average

0.1	0.1	0.1
0.1	0.1	0.1
0.1	0.1	0.1

filtered image



unweighted average, a.k.a. box filter
⇒ blurring effect


original image



gaussian

.0	.0	.0	.0	.0	.0	.0
.0	.0	.0	.1	.0	.0	.0
.0	.0	.1	.2	.1	.0	.0
.0	.1	.2	.4	.2	.1	.0
.0	.0	.1	.2	.1	.0	.0
.0	.0	.0	.1	.0	.0	.0
.0	.0	.0	.0	.0	.0	.0

filtered image



LOW-PASS FILTER

original image



gaussian

.0	.0	.0	.0	.0	.0	.0
.0	.0	.0	.1	.0	.0	.0
.0	.0	.1	.2	.1	.0	.0
.0	.1	.2	.4	.2	.1	.0
.0	.0	.1	.2	.1	.0	.0
.0	.0	.0	.1	.0	.0	.0
.0	.0	.0	.0	.0	.0	.0

weighted average

⇒ blurring effect with more weight on central pixel

filtered image




original image



laplacian

0	-1	0
-1	4	-1
0	-1	0

filtered image



HIGH-PASS FILTER


original image



laplacian

0	-1	0
-1	4	-1
0	-1	0

filtered image



⇒ edge detection (no orientation)


original image



sharpen

0	-1	0
-1	5	-1
0	-1	0

filtered image



HIGH-PASS FILTER

original image



sharpen

0	-1	0
-1	5	-1
0	-1	0

filtered image



identity kernel + highpass kernel
⇒ sharpening effect


original image



sobel x

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

filtered image



HIGH-PASS FILTER


original image



sobel x

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

filtered image



⇒ *vertical* edge detection (x-direction)


original image



sobel y

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

filtered image



HIGH-PASS FILTER


original image




sobel y

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

filtered image



⇒ *horizontal* edge detection (y-direction)



$\Rightarrow$ edges + magnitude

original image



emboss

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

filtered image



⇒ styling effect

Gaussian filters are a true low-pass filter for the image

⇒ we can retrieve the low-frequency in an image

⇒ we can retrieve the high-frequency in an image by subtracting the low-frequency from the original image

original



low frequency
(gaussian)



high frequency
(=original - gaussian)



reconstructed
(=low fq + high fq)



1. Introduction

2. Spatial domain filtering

3. Frequency domain filtering

1. 1D Fourier transform

2. 2D Fourier transform

3. Butterworth filter

Why filter in the frequency domain?

- ⇒ convolutions for spatial domain filtering is powerful, BUT it has high computational costs
- ⇒ frequency domain filtering offers computational advantages:
 - convolution in the time domain $\iff$ multiplication in the frequency domain

Why filter in the frequency domain?

- ⇒ convolutions for spatial domain filtering is powerful, BUT it has high computational costs
- ⇒ frequency domain filtering offers computational advantages:
 - convolution in the time domain $\iff$ multiplication in the frequency domain

3.1. 1D Fourier transform



Fourier theorem:

⇒ any periodic function can be expressed as the sum of sines and/or cosines of different frequencies

Fourier transform:





⇒ **Forward transform (FFT):** Time Domain → Frequency Domain

⇒ **Inverse transform (IFFT):** Frequency Domain → Time Domain



Fourier transform on images ?

- ⇒ an image can also be expressed as the sum of sinusoids of different frequencies and amplitudes
- ⇒ the appearance of an image depends on the frequencies of its sinusoidal components:
 - low frequencies → regions with intensities that vary slowly
 - high frequencies → edges and other sharp intensity transitions







NB: Fourier transform of a real function is symmetric about the origin; by convention frequency 0 is set at the center of image.

NB: When applying the Fourier transform to images, the term time domain is still used by convention, although the signal is defined in the spatial domain.

Fourier transform on images ?

- ⇒ an image can also be expressed as the sum of sinusoids of different frequencies and amplitudes
- ⇒ the appearance of an image depends on the frequencies of its sinusoidal components:
 - low frequencies → regions with intensities that vary slowly
 - high frequencies → edges and other sharp intensity transitions





NB: Fourier transform of a real function is symmetric about the origin; by convention frequency 0 is set at the center of image.

NB: When applying the Fourier transform to images, the term time domain is still used by convention, although the signal is defined in the spatial domain.



2D Fourier transform on SYNTHETIC images

- ⇒ “dots” symmetric about origin in amplitude spectrum
- ⇒ distance/direction from origin imply frequency in time domain





2D Fourier transform on SYNTHETIC images

- ⇒ “dots” symmetric about origin in amplitude spectrum
- ⇒ distance/direction from origin imply frequency in time domain




2D Fourier transform on SYNTHETIC images

- ⇒ “dots” symmetric about origin in amplitude spectrum
- ⇒ distance/direction from origin imply frequency in time domain




2D Fourier transform on REAL images



Credit: A. Zisserman



2D Fourier transform on REAL images



Credit: Alegre et al. 2016





2D Fourier transform on REAL images

⇒ let's try on our astronaut



2D Fourier transform on REAL images


⇒ let's try on our astronaut



2D Fourier transform on REAL images

⇒ band-pass image frequencies?





- low-pass filter → cut off high-frequencies
- high-pass filter → cut off low-frequencies



2D Fourier transform on REAL images


⇒ band-pass image frequencies?

- low-pass filter → cut off high-frequencies
- high-pass filter → cut off low-frequencies



2D Fourier transform on REAL images

⇒ image can be reconstructed from band-passed spectra using the 2D inverse Fourier transform (iFFT2)



2D Fourier transform on REAL images

- ⇒ the ideal low-pass filter (LPF) introduces artefacts:
- “ripples” near strong edges in the original image: ringing artifact
 - related to the sharp cut-off in ideal frequency domain

low-pass filtered image




ringing effect



2D Fourier transform on REAL images

⇒ the **ideal low-pass filter** (LPF) introduces artefacts:


- “ripples” near strong edges in the original image: ringing artifact
- related to the sharp cut-off in ideal frequency domain



- Ideal LPF has significant 'side-lobes' in the time domain

2D Fourier transform on REAL images

- ⇒ the **Butterworth** filter offers impulse response without side-lobes in the time domain ideal
→ no “ringing artifacts”, due to the absence of discontinuity in spectrum



- Impulse response without side-lobes in the time domain

2D Fourier transform on REAL images

- ⇒ the **Butterworth** filter offers impulse response without side-lobes in the time domain ideal
→ no “ringing artifacts”, due to the absence of discontinuity in spectrum

