Aaron Ponti

Microscopy Network Basel

Image processing course Linear shift-invariant systems





Signals

 A function containing information about some phenomenon of interest.



Time (s)

 A quantity exhibiting variation in time and/or space.



1D



Analog and digital signals (ID)





A/D conversion

- Analog-to-digital conversion is a 2step process:
 - Sampling: converts a continuous signal into a discrete one
 - Quantization: discretizes the amplitude of the signal.



Sampling



The continuous signal (blue) is measured at discrete time intervals (red dots).



The continuous signal (blue) at discrete time intervals is approximated to a fixed number of discrete values (magenta crosses).

Example: CD sound quality 44.1 kHz, 16bit, stereo



Compact disc

2 channels, each with:

Sampling rate 44.1 kHz $\rightarrow \Delta T = 1/rate \approx 2.3 \cdot 10^{-5} s$

 $\frac{\text{Quantization levels}}{16 \text{ bit} \rightarrow 2^{16} = 65536 \text{ levels}}$



Analog and digital signals (2D)





Digital image



In this example, signal intensity is approximated by **256 discrete values (8 bits)**.

Sampling -> spatial resolution



256 x 256

128 x 128



32 x 32

Quantization -> grayscale resolution



The grayscale resolution of an image is expressed as its **bit depth**. The maximum number of brightness (i.e. gray) levels in an **n-bit** image is **2**ⁿ.

Quantization \rightarrow grayscale resolution



256 levels (8 bit)



64 levels (6 bit)



8 levels (3 bit)

2 levels (1 bit)



Resolution summary



Digital Camera System

•° LSI SYSTEMS



Discrete signal (notation)





<u>Examples</u>:

 $y(n_1, n_2) = 255 - x(n_1, n_2)$ $y(n_1, n_2) = \text{median}(N(x(n_1, n_2)))$ A neighborhood of a given position (pixel) $x(n_1, n_2)$



T[] can be any sort of transformation (system) of the input signal $x(n_1, n_2)$. We will now consider a family of systems with following properties:

- Linearity
- Spatial (shift) invariance



Linear systems

lf

$\mathbf{T}[a_1x_1(n_1, n_2) + a_2x_2(n_1, n_2)] = a_1\mathbf{T}[x_1(n_1, n_2)] + a_2\mathbf{T}[x_2(n_1, n_2)]$

then T[] is linear.

The transformed version of a weighted sum of signals is the same as the weighted sum of the signals transformed individually.

(Alternatively, a linear system can be decomposed into constituents that are processed independently, and the result combined in the end.)

Shift-invariant systems

Given:

$$\mathbf{T}[x(n_1, n_2)] = y(n_1, n_2)$$

lf

$$\mathbf{T}[\mathbf{x}(\mathbf{n}_1 - \mathbf{k}_1, \, \mathbf{n}_2 - \mathbf{k}_2)] = \mathbf{y}(\mathbf{n}_1 - \mathbf{k}_1, \, \mathbf{n}_2 - \mathbf{k}_2)$$

then T[] is shift-invariant.

If the input is shifted by a given amount, the output will be shifted by the same amount.

(Or, the location of the origin of the coordinate system is irrelevant.)



Discrete Unit Impulse

$$\delta(n_1, n_2) = \begin{cases} 1, & \text{for } n_1 = n_2 = 0\\ 0, & \text{otherwise} \end{cases}$$











The system response to the unit impulse is all we need to fully describe the LSI system.

Convolution

$$x(n_1, n_2) \longrightarrow h(n_1, n_2) \longrightarrow y(n_1, n_2)$$

LSI systems can be described and efficiently implemented by the mathematical operation of **convolution**.

$$y(n_1, n_2) = x(n_1, n_2) \circledast h(n_1, n_2)$$

$$y(n_1, n_2) = x(n_1, n_2) \circledast h(n_1, n_2) = \sum_{k_1 = -\infty}^{\infty} \sum_{k_2 = -\infty}^{\infty} x(k_1, k_2) h(n_1 - k_1, n_2 - k_2)$$

Convolution (ID example)

$$y(n) = x(n) \circledast h(n) = \sum_{k=-\infty}^{\infty} x(k) h(n-k)$$

Х	=	[1	2	3	4	5]
h	=	[2	4	6]	

[1 2 3 4 <mark>5</mark>] [6 4 2]	5 * 6 = [30]
[1 2 3 4 5] [6 4 2]	6 * 4 + 5 * 4 = [44 30]
[1 2 3 4 5] [6 4 2]	3 * 6 + 4 * 4 + 5 * 2 = [44 44 30]
[1 2 3 4 5]	2 * 6 + 3 * 4 + 4 * 2 = [32 44 44 30]
[1 2 3 4 5] ← [6 4 2]	1 * 6 + 2 * 4 + 3 * 2 = [20 32 44 44 30
[1 2 3 4 5] ←── [6 4 2]	1 * 4 + 2 * 2 = [8 20 32 44 44 30]
[1 2 3 4 5] [6 4 2]	1 * 2 = [<mark>2</mark> 8 20 32 44 44 30]

Full convolution: $y(n) = [2 \ \underline{8} \ \underline{20} \ \underline{32} \ \underline{44} \ \underline{44} \ \underline{30}]$





Impulse signal

Impulse response function

 $x(n_1, n_2)$

 $y(n_1, n_2)$



 $y(n_1, n_2) = x(n_1, n_2) \circledast h(n_1, n_2)$

 $x(n_1, n_2)$

 $y(n_1, n_2)$



The output signal is formed as a linear combination (i.e. weighted sum) of spatially-shifted¹ impulse response functions.

¹ Or time-shifted in 1D.

Spatial filtering through convolution

Examples of convolution filters



 $x(n_1, n_2)$

3x3 average filter (poor) Low-pass filter



 $h = \begin{bmatrix} 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \end{bmatrix}$

3x3 Gaussian filter (better) Low-pass filter



 $h = \begin{bmatrix} 0.0751 & 0.1238 & 0.0751 \\ 0.1238 & 0.2042 & 0.1238 \\ 0.0751 & 0.1238 & 0.0751 \end{bmatrix}$





	-1	-1	-1
h =	-1	8	-1
	1	-1	-1

"kernels" h

$$y(n_1, n_2) = x(n_1, n_2) \circledast h(n_1, n_2) = \sum_{k_1 = -\infty}^{\infty} \sum_{k_2 = -\infty}^{\infty} x(k_1, k_2) h(n_1 - k_1, n_2 - k_2)$$



Spatial filtering through convolution

Examples of convolution filters





Spatial filtering



Low-pass filter Gaussian kernel Band-pass filter Difference of Gaussians kernel

High-pass filter I - Gaussian kernel

The convolution of a signal in the spatial domain has very specific effects on the frequency content of the signal.



What happens if we pass exponential sequences through an LSI system?

 $\begin{array}{c} e^{jw_1'n_1}e^{jw_2'n_2} \\ \textbf{x}(n_1, n_2) \end{array} \xrightarrow{\mathsf{LSI}} y(n_1, n_2) ? \end{array}$

Frequency response of a system $x(n_1, n_2) \xrightarrow{\text{LSI}} y(n_1, n_2)$?

We calculate the convolution of $x(n_1, n_2)$ with the impulse response $h(n_1, n_2)$:

$$y(n_1, n_2) = x(n_1, n_2) \otimes h(n_1, n_2)$$

The signal goes infough uniouched.

Frequency response of a system

$$y(n_{1}, n_{2}) = e^{jw_{1}^{'}n_{1}}e^{jw_{2}^{'}n_{2}} \sum_{k_{1}=-\infty}^{\infty} \sum_{k_{2}=-\infty}^{\infty} h(k_{1}, k_{2})e^{-jw_{1}^{'}k_{1}}e^{-jw_{2}^{'}k_{2}}$$
$$x(n_{1}, n_{2}) \qquad \qquad H(\omega_{1}^{'}, \omega_{2}^{'})$$

 $H\left(\omega_{1}^{\scriptscriptstyle \text{\tiny I}},\omega_{2}^{\scriptscriptstyle \text{\tiny I}}\right)\!:$

- is the frequency response of the system
- tells us how the LSI system reacted to the input frequencies
- is the **Fourier transform** of the impulse response $h(n_1, n_2)$
- has a magnitude and a phase



Exponential sequences

(Joseph Fourier, 1768 – 1830)

Exponential sequences are **building blocks** of <u>any signal</u> and so called **eigen-functions** of LSI systems.



LSI systems cannot produce frequencies that are not in the input.

Continuous Fourier Transform

- We consider the continuous Fourier transform of a discrete signal.
- The Fourier transform maps a signal to its frequency representation.

Continuous variables

$$X(\omega_{1},\omega_{2}) = \sum_{n_{1}=-\infty}^{\infty} \sum_{n_{2}=-\infty}^{\infty} x(n_{1},n_{2})e^{-j\omega_{1}n_{1}}e^{-j\omega_{2}n_{2}}$$
Discrete signal

• The Fourier transform is periodic with period 2π in the ω_1 and ω_2 directions (since the exponential sequences have the same periodicity).

Inverse Fourier transform

Given an LSI system and its impulse response $h(n_1, n_2)$, we want to calculate its (continuous) frequency response $H(\omega_1, \omega_2)$, i.e. the Fourier Transform of $h(n_1, n_2)$.



$$H(\omega_1, \omega_2) = \sum_{n_1 = -\infty}^{\infty} \sum_{n_2 = -\infty}^{\infty} h(n_1, n_2) e^{-j\omega_1 n_1} e^{-j\omega_2 n_2}$$

 $h\left(0,0\right) + h\left(-1,0\right)e^{j\omega_{1}} + h\left(1,0\right)e^{-j\omega_{1}} + h\left(0,-1\right)e^{j\omega_{2}} + h\left(0,1\right)e^{-j\omega_{2}} = 0$

$$\frac{1}{3} + \frac{1}{6}e^{j\omega_1} + \frac{1}{6}e^{-j\omega_1} + \frac{1}{6}e^{j\omega_2} + \frac{1}{6}e^{-j\omega_2} =$$

$$\frac{1}{3} + \frac{1}{6} \cdot 2\cos\omega_1 + \frac{1}{6} \cdot 2\cos\omega_2 = \frac{1}{3}\left(1 + \cos\omega_1 + \cos\omega_2\right)$$

Continuous and periodic.







Convolution theorem

The convolution theorem describes the input – output relationships of an LSI system in the **frequency** domain.

$$y(n_{1}, n_{2}) = x(n_{1}, n_{2}) \circledast h(n_{1}, n_{2})$$

$$\mathsf{FT} \qquad \qquad \mathsf{FT} \qquad \qquad$$

Multiplication in the frequency domain

Reverse is also true.



Sampling

Under which conditions can we reconstruct a continuous, band-limited signal from its discrete representation with no loss of information?

 $X_a(\Omega_1, \Omega_2)$



An alternative view: sampling can be modelled in direct space with a **multiplication** of the signal with a train of delta functions. The convolution theorem tells us that this corresponds to a **convolution** of the Fourier Transform of the signal with the Fourier Transform of the impulse train.



The <u>spectrum of the analog signal $X_{\underline{a}}$ is periodically extended with periods $2\pi/T_i$.</u> T_1 and T_2 define how far apart the replica of the spectrum will be.



This (base) band contains the spectrum of the <u>analog</u> signal with no loss of information.



Oversampling

 $X_{S}(\Omega_{1}T_{1},\,\Omega_{2}T_{2})$





Undersampling



Aliased spectrum of the analog signal



Aliasing: example



Properly sampled



Undersampled and aliased



A 2D sync function in spatial domain. Constant (with gain T_1T_2) on the support area.

Discrete Fourier Transform (DFT)

- The continuous Fourier Transform of a discrete signal is not computable
 - continuous (i.e. infinitely many) frequencies ω_1 and ω_2 .
- Sampling in the frequency domain results in periodic extension of the sampled spatial signal.



- One period in the frequency domain corresponds to one period of the spatial domain: this mapping is the **Discrete Fourier Transform (DFT)**.
- A sampled version of one period of the continuous Fourier transform is all is needed to reconstruct the analog signal.

$$\begin{aligned} \textbf{Discrete Fourier Transform (DFT)} \\ & \text{Continuous} \\ X(\omega_1, \omega_2) = \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} x(n_1, n_2) e^{-j\omega_1 n_1} e^{-j\omega_2 n_2} \\ & \text{Sampling in frequency space. We keep only one period.} \\ X(k_1, k_2) = X(\omega_1, \omega_2) \mid_{\omega_1 = \frac{2\pi}{N_1} k_1, \omega_2 = \frac{2\pi}{N_2} k_2} \begin{cases} k_1 = 0, \dots, N_1 - 1 \\ k_2 = 0, \dots, N_2 - 1 \end{cases} \\ & \text{N}_r \text{ samples } N_2 \text{ samples} \end{cases} \end{aligned}$$

Most properties of the continuous FT apply to the DFT with one exception: **linear shifts** become **circular shifts** ("wrap around").

Fast Fourier Transforms (FFTs)

DFT:
$$X(k_1, k_2) = \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} x(n_1, n_2) e^{-j\frac{2\pi}{N_1}n_1k_1} e^{-j\frac{2\pi}{N_2}n_2k_2}$$

For each
$$(k_1, k_2)$$
: $N_1 * N_2$ multiplications;

$$\begin{cases}
k_1 = 0, \dots, N_1 - 1 \\
k_2 = 0, \dots, N_2 - 1
\end{cases}$$

For $N_1 = N_2 = N$, a full DFT requires **N**⁴ multiplications.

Fast Fourier Transforms (FFTs) are a family of algorithms that impressively speed up calculation of the DFT.

Best runtime is in the order:

 $aN^2\log_2N$

with *a* < 1.

For a 1024 x 1024 image, the FFT is approximately 10^5 times faster that the DFT.





DFT examples











Low-pass filtering



In practice, one does not create sharp cut-offs in frequency domain, since this creates **ringing artifacts** that appear as spurious signals near sharp transitions in a signal, i.e. they appear as "rings" near edges.





In practice, one does not create sharp cut-offs in frequency domain, since this creates **ringing artifacts** that appear as spurious signals near sharp transitions in a signal, i.e. they appear as "rings" near edges.





In practice, one does not create sharp cut-offs in frequency domain, since this creates **ringing artifacts** that appear as spurious signals near sharp transitions in a signal, i.e. they appear as "rings" near edges.

Removing unwanted frequencies



Circular convolution

linear convolution

$$\begin{array}{c} y\left(n_{1},n_{2}\right)=x\left(n_{1},n_{2}\right)\circledast h\left(n_{1},n_{2}\right)\\ & \text{DFT} \end{array} \quad \begin{array}{c} \text{DFT} \end{array} \quad \begin{array}{c} \text{DFT} \end{array} \\ Y\left(k_{1},k_{2}\right)=X\left(k_{1},k_{2}\right)\cdot H\left(k_{1},k_{2}\right)\\ & \text{DFT} \end{array} \quad \begin{array}{c} \text{circular convolution}\\ y\left(n_{1},n_{2}\right) \end{array} \end{array}$$

The circular convolution is infinite-length and periodic whereas the linear convolution is finite length, therefore a trick is needed to calculate linear convolution in frequency domain.



Circular convolution



Inappropriate support \rightarrow aliasing

Appropriate support

Linear convolution in frequency domain (how to)



- Pad $x(n_1, n_2)$ and $h(n_1, n_2)$ with zeros to size $(N_1+M_1-1 \times N_2+M_2-1)$
- Calculate the FFT (DFT) of both
- Multiply the transforms together
- Calculate the inverse FFT of the result \rightarrow same result as linear convolution
- Carve out the result from the center of the result

Linear convolution in frequency domain (how to)



