



Lecture 3: Linear Filters

Professor Fei-Fei Li
Stanford Vision Lab

What we will learn today?

- Images as functions
- Linear systems (filters)
- Convolution and correlation
- Discrete Fourier Transform (DFT)
- Sampling and aliasing

Some background reading:

Forsyth and Ponce, Computer Vision, Chapter 7 & 8

Jae S. Lim, Two-dimensional signal and image processing, Chapter 1, 4, 5

Images as functions






- **An Image** as a function f from \mathbb{R}^2 to \mathbb{R}^M :

- $f(x, y)$ gives the **intensity** at position (x, y) $M=1$
- Defined over a rectangle, with a finite range:

$$f: [a,b] \times [c,d] \rightarrow [0,255] \quad \mathbb{R}^2 \rightarrow \mathbb{R}^1$$

Domain
support

range



Images as functions

- **An Image** as a function f from \mathbb{R}^2 to \mathbb{R}^M :
 - $f(x, y)$ gives the **intensity** at position (x, y)
 - Defined over a rectangle, with a finite range:

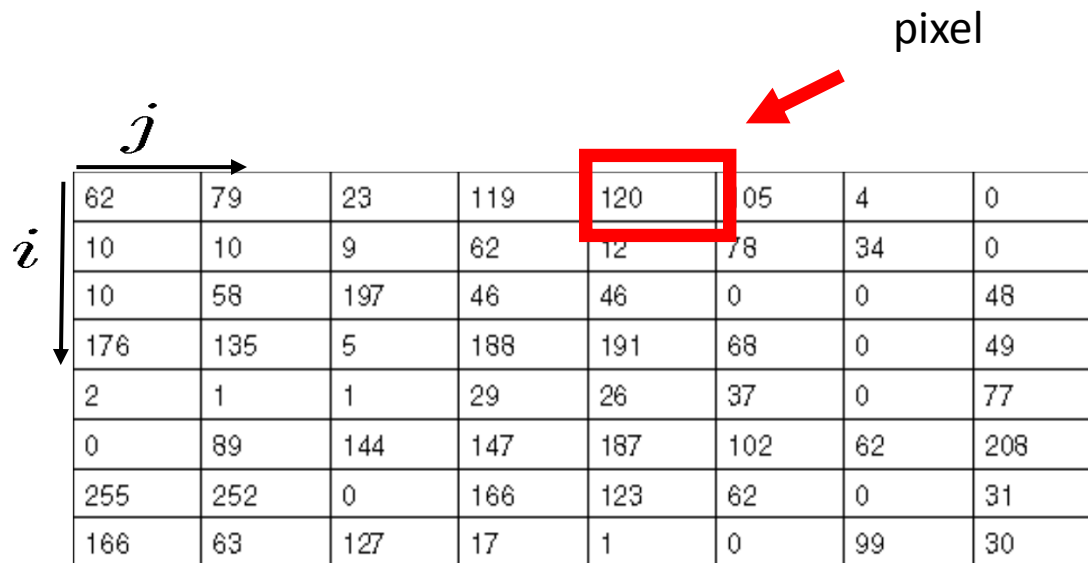
$$f: \underbrace{[a,b] \times [c,d]}_{\substack{\text{Domain} \\ \text{support}}} \rightarrow \underbrace{[0,255]}_{\text{range}}$$

- A color image: $f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$ $\mathbb{R}^2 \rightarrow \mathbb{R}^3$
color: RGB
HSV
Lab

Images as discrete functions

- Images are usually **digital (discrete)**:
 - **Sample** the 2D space on a regular grid
- Represented as a matrix of integer values

pixel



	j								
		62	79	23	119	120	05	4	0
i		10	10	9	62	12	78	34	0
		10	58	197	46	46	0	0	48
		176	135	5	188	191	68	0	49
		2	1	1	29	26	37	0	77
		0	89	144	147	187	102	62	208
		255	252	0	166	123	62	0	31
		166	63	127	17	1	0	99	30

Images as discrete functions

Cartesian coordinates

$f[n, m]$ =

↑
Notation for discrete functions

$$\begin{bmatrix} \ddots & & \vdots & & \\ & f[-1, 1] & f[0, 1] & f[1, 1] & \\ \dots & f[-1, 0] & \underline{f[0, 0]} & f[1, 0] & \dots \\ & f[-1, -1] & f[0, -1] & f[1, -1] & \\ & & \vdots & & \ddots \end{bmatrix}$$

Images as discrete functions

Array coordinates

$$\mathbf{A} = \begin{bmatrix} a_{11} & \dots & a_{1M} \\ \vdots & \ddots & \vdots \\ a_{N1} & \dots & a_{NM} \end{bmatrix}$$

Matlab notation

What we will learn today?

- Images as functions
- **Linear systems (filters)**
- Convolution and correlation
- Discrete Fourier Transform (DFT)
- Sampling and aliasing

Some background reading:

Forsyth and Ponce, Computer Vision, Chapter 7 & 8

Jae S. Lim, Two-dimensional signal and image processing, Chapter 1, 4, 5

Systems and Filters

- **Filtering:**

- Form a new image whose pixels are a combination original pixel values

Goals:

- Extract useful information from the images
 - Features (edges, corners, blobs...)
- Modify or enhance image properties:
 - super-resolution; in-painting; de-noising

De-noising



Salt and pepper noise

Super-resolution



In-painting



Bertamio et al

2D discrete-space systems (filters)

$$f[n, m] \rightarrow \boxed{\text{System } \mathcal{S}} \rightarrow g[n, m]$$

$$g = \mathcal{S}[f], \quad g[n, m] = \mathcal{S}\{f[n, m]\}$$

$$f[n, m] \xrightarrow{\mathcal{S}} g[n, m]$$

Filters: Examples

- 2D DS moving average over a 3×3 window of neighborhood

filtered im

$$g[n, m] = \frac{1}{9} \sum_{k=n-1}^{n+1} \sum_{l=m-1}^{m+1} f[k, l]$$
normalizing const

3x3 window

orig. im

filter

$$= \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 f[n-k, m-l]$$

		h	
1	1	1	
1	1	1	
1	1	1	

$\frac{1}{9}$

convolution

$$(f * h)[m, n] = \frac{1}{9} \sum_{k, l} f[k, l] h[m-k, n-l]$$

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

Courtesy of S. Seitz

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10							

$$(f * g)[m, n] = \sum_{k,l} f[k, l] g[m - k, n - l]$$

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20						

$$(f * g)[m, n] = \sum_{k,l} f[k, l] g[m - k, n - l]$$

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20	30					

$$(f * g)[m, n] = \sum_{k,l} f[k, l] g[m - k, n - l]$$

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

	0	10	20	30	30				

$$(f * g)[m, n] = \sum_{k,l} f[k, l] g[m - k, n - l]$$

Moving average

$F[x, y]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$G[x, y]$

"blur" or
"smooth"

		0	10	20	30	30	30	20	10
		0	20	40	60	60	60	40	20
		0	30	60	90	90	90	60	30
		0	30	50	80	80	90	60	30
		0	30	50	80	80	90	60	30
		0	20	30	50	50	60	40	20
		10	20	30	30	30	30	20	10
		10	10	10	0	0	0	0	0

$$(f * g)[m, n] = \sum_{k,l} f[k, l] g[m - k, n - l]$$

Source: S. Seitz

Moving average

In summary:

- Replaces each pixel with an average of its neighborhood.
- Achieve smoothing effect (remove sharp features)

$$\frac{1}{9} g[\cdot, \cdot]$$

1	1	1
1	1	1
1	1	1

Moving average



Filters: Examples

- Image segmentation based on a simple threshold:

$$g[n, m] = \begin{cases} 1, & f[n, m] > 10 \\ 0, & \text{otherwise.} \end{cases}$$

Classification of systems

- Amplitude properties

- Linearity

- Stability

- Invertibility

- Spatial properties

- Causality

- Separability

- Memory

- Shift invariance

- Rotation invariance

filter
↓
Linear Shift-Invariant Systems
LSI

Shift-invariance

- If $f[n, m] \xrightarrow{\mathcal{S}} g[n, m]$ then

$$f[n - n_0, m - m_0] \xrightarrow{\mathcal{S}} g[n - n_0, m - m_0]$$

for every input image $f[n, m]$ and shifts n_0, m_0

- Is the moving average shift invariant a system ?

Is the moving average system is shift invariant?

$$F[x, y]$$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$$G[x, y]$$

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

Is the moving average system is shift invariant?

$$f[n, m] \xrightarrow{\mathcal{S}} g[n, m] = \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 f[n - k, m - l]$$

$$f[n - n_0, m - m_0]$$

$$\begin{aligned} \xrightarrow{\mathcal{S}} g[n, m] &= \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 f[n - k, m - l] \\ &= \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 f[(n - n_0) - k, (m - m_0) - l] \\ &= g[n - n_0, m - m_0] \end{aligned}$$

Yes!

Linear Systems (filters)

$$f(x, y) \rightarrow \boxed{\mathcal{S}} \rightarrow g(x, y)$$

- Linear filtering:
 - Form a new image whose pixels are a weighted sum of original pixel values
 - Use the same set of weights at each point
- **S** is a linear system (function) iff it *S satisfies*

$$\mathcal{S}[\alpha f_1 + \beta f_2] = \alpha \mathcal{S}[f_1] + \beta \mathcal{S}[f_2]$$

superposition property

Linear Systems (filters)

$$f(x, y) \rightarrow \boxed{S} \rightarrow g(x, y)$$

- Is the moving average a linear system?

$$f(n, m) \xrightarrow{S} g(n, m) = \frac{1}{q} \sum_{k=-1}^1 \sum_{l=-1}^1 f(n-k, m-l)$$

$$S(\alpha f) = \frac{1}{q} \sum \sum \alpha f(n-k, m-l)$$

$$= \alpha \cdot \left[\frac{1}{q} \sum \sum f(n-k, m-l) \right]$$

$$= \alpha \cdot S(f)$$

- Is thresholding a linear system?

$$g(n, m) = \begin{cases} f(n, m) & f(n, m) > T \\ 0 & \text{otherwise} \end{cases}$$

$$- f_1[n, m] < T$$

$$- f_2[n, m] < T$$

No!

LSI (linear *shift invariant*) systems

Impulse response

delta func. $\delta[n, m] = \begin{cases} 1 & n=m=0 \\ 0 & \text{otherwise} \end{cases}$

$$\delta_2[n, m] \rightarrow \boxed{\mathcal{S}} \rightarrow h[n, m]$$

$$\delta_2[n - k, m - l] \rightarrow \boxed{\mathcal{S} \text{ (SI)}} \rightarrow h[n - k, m - l]$$

LSI (linear *shift invariant*) systems

Example: impulse response of the 3 by 3 moving average filter:

$$h[n, m] = \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 \delta_2[n - k, m - l]$$

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

$$\frac{1}{9} \begin{matrix} & & & \mathbf{h} \\ & & & \\ & & & \end{matrix} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

LSI (linear *shift invariant*) systems

An LSI system is completely specified by its impulse response.

sifting property of the delta function

$$f[n, m] = \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} f[k, l] \delta_2[n - k, m - l]$$

superposition

$$\delta_2[n, m] \rightarrow \boxed{\mathcal{S}} \rightarrow h[n, m] \quad \rightarrow \boxed{\mathcal{S} \text{ LSI}} \rightarrow \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} f[k, l] h[n - k, m - l]$$

Discrete convolution

$$= f[n, m] ** h[n, m]$$

What we will learn today?

- Images as functions
- Linear systems (filters)
- **Convolution and correlation**
- Discrete Fourier Transform (DFT)
- Sampling and aliasing



Some background reading:

Forsyth and Ponce, Computer Vision, Chapter 7 & 8

Jae S. Lim, Two-dimensional signal and image processing, Chapter 1, 4, 5

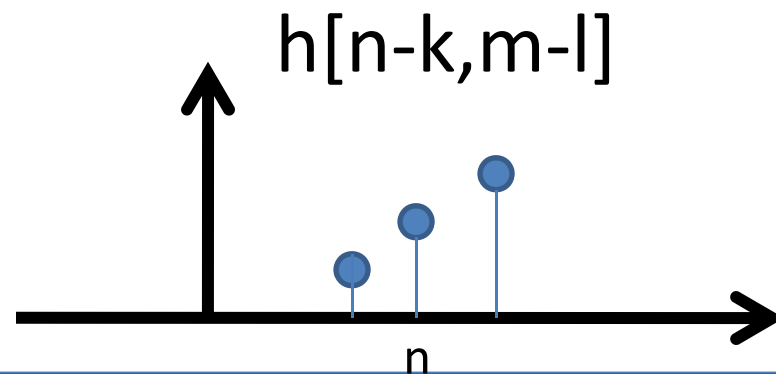

Discrete convolution

- Fold $h[n,m]$ about origin to form $h[-k,-l]$
- Shift the folded results by n,m to form $h[n-k,m-l]$
- Multiply $h[n-k,m-l]$ by $f[k,l]$
- Sum over all k,l
- Repeat for every n,m





Discrete convolution

- Fold $h[n,m]$ about origin to form $h[-k,-l]$
- Shift the folded results by n,m to form $h[n-k,m-l]$
- Multiply $h[n-k,m-l]$ by $f[k,l]$
- Sum over all k,l
- Repeat for every n,m



Discrete convolution

- Fold $h[n,m]$ about origin to form $h[-k,-l]$
- Shift the folded results by n,m to form $h[n-k,m-l]$
- Multiply $h[n-k,m-l]$ by $f[k,l]$
- Sum over all k,l
- Repeat for every n,m



Convolution in 2D - examples



Original

$**$

•0	•0	•0
•0	•1	•0
•0	•0	•0

$=$?

The diagram illustrates a 2D convolution operation. On the left is a grayscale image of a face. To its right is a 3x3 kernel matrix with a red 'h' above it. The kernel matrix has a central value of 1 and all other values are 0. To the left of the kernel matrix are two red asterisks (**). To the right of the kernel matrix is an equals sign followed by a large question mark.

Convolution in 2D - examples



Original

•0	•0	•0
•0	•1	•0
•0	•0	•0

=



Filtered
(no change)

Courtesy of D Lowe

Convolution in 2D - examples



Original

•0	•0	•0
•0	•0	•1
•0	•0	•0

=

?


Convolution in 2D - examples



Original

•0	•0	•0
•0	•0	•1
•0	•0	•0

=



Shifted right
By 1 pixel

Courtesy of D Lowe

Convolution in 2D - examples



Original


$$\frac{1}{9} \begin{bmatrix} \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \end{bmatrix} = ?$$

Convolution in 2D - examples



Original

$$\frac{1}{9} \begin{bmatrix} \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \end{bmatrix} =$$



Blur (with a box filter)

Courtesy of D Lowe

Convolution in 2D - examples



•0	•0	•0
•0	•2	•0
•0	•0	•0

-

$\frac{1}{9}$

•1	•1	•1
•1	•1	•1
•1	•1	•1

= ?

(Note that filter sums to 1)

Original

unchanged

“details of the image”

•0	•0	•0
•0	•1	•0
•0	•0	•0

+

•0	•0	•0
•0	•1	•0
•0	•0	•0

-

$\frac{1}{9}$


•1	•1	•1
•1	•1	•1
•1	•1	•1

Courtesy of D Lowe


- What does blurring take away?



-




=



- Let's add it back:




+ a



=



Convolution in 2D – Sharpening filter



Original

Sharpening filter: Accentuates differences with local average

Courtesy of D Lowe

Convolution properties

- **Commutative property:**

$$f ** h = h ** f$$

- **Associative property:**

$$(f ** h_1) ** h_2 = f ** (h_1 ** h_2)$$

- **Distributive property:**

$$f ** (h_1 + h_2) = (f ** h_1) + (f ** h_2)$$

The order doesn't matter! $h_1 ** h_2 = h_2 ** h_1$

Convolution properties

- **Shift property:**

$$f[n, m] ** \delta_2[n - n_0, m - m_0] = f[n - n_0, m - m_0]$$

- **Shift-invariance:**

$$g[n, m] = f[n, m] ** h[n, m]$$

$$\implies f[n - l_1, m - l_1] ** h[n - l_2, m - l_2]$$

$$= g[n - l_1 - l_2, m - l_1 - l_2]$$

Image support and edge effect

- A computer will only convolve **finite support signals**.

- That is: images that are zero for n, m outside some rectangular region

- MATLAB's `conv2` performs 2D DS convolution of finite-support signals.


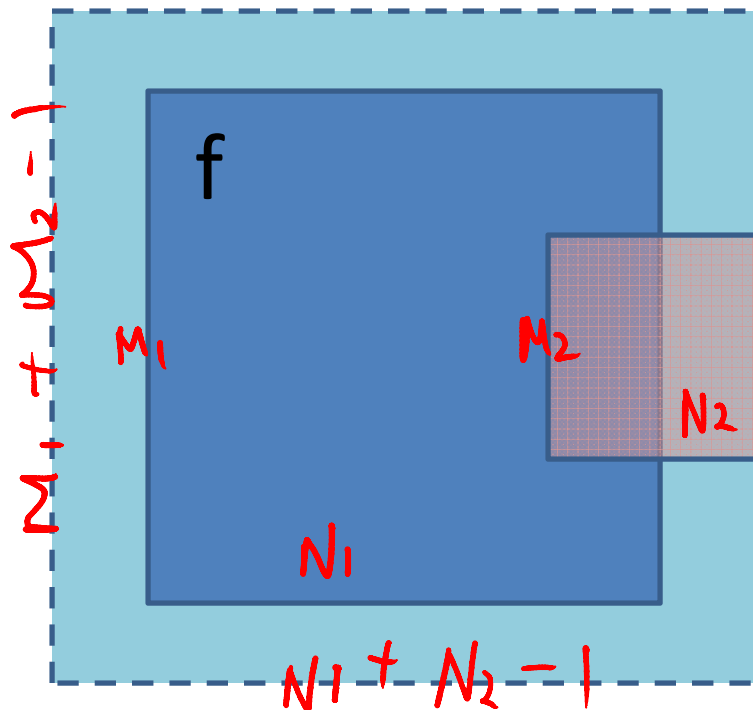


Image support and edge effect

- A computer will only convolve **finite support signals**.
- What happens at the edge?



h

- zero “padding”
- edge value replication
- mirror extension
- more (beyond the scope of this class)

-> Matlab conv2 uses zero-padding

Cross correlation

Cross correlation of two 2D signals $f[n,m]$ and $g[n,m]$

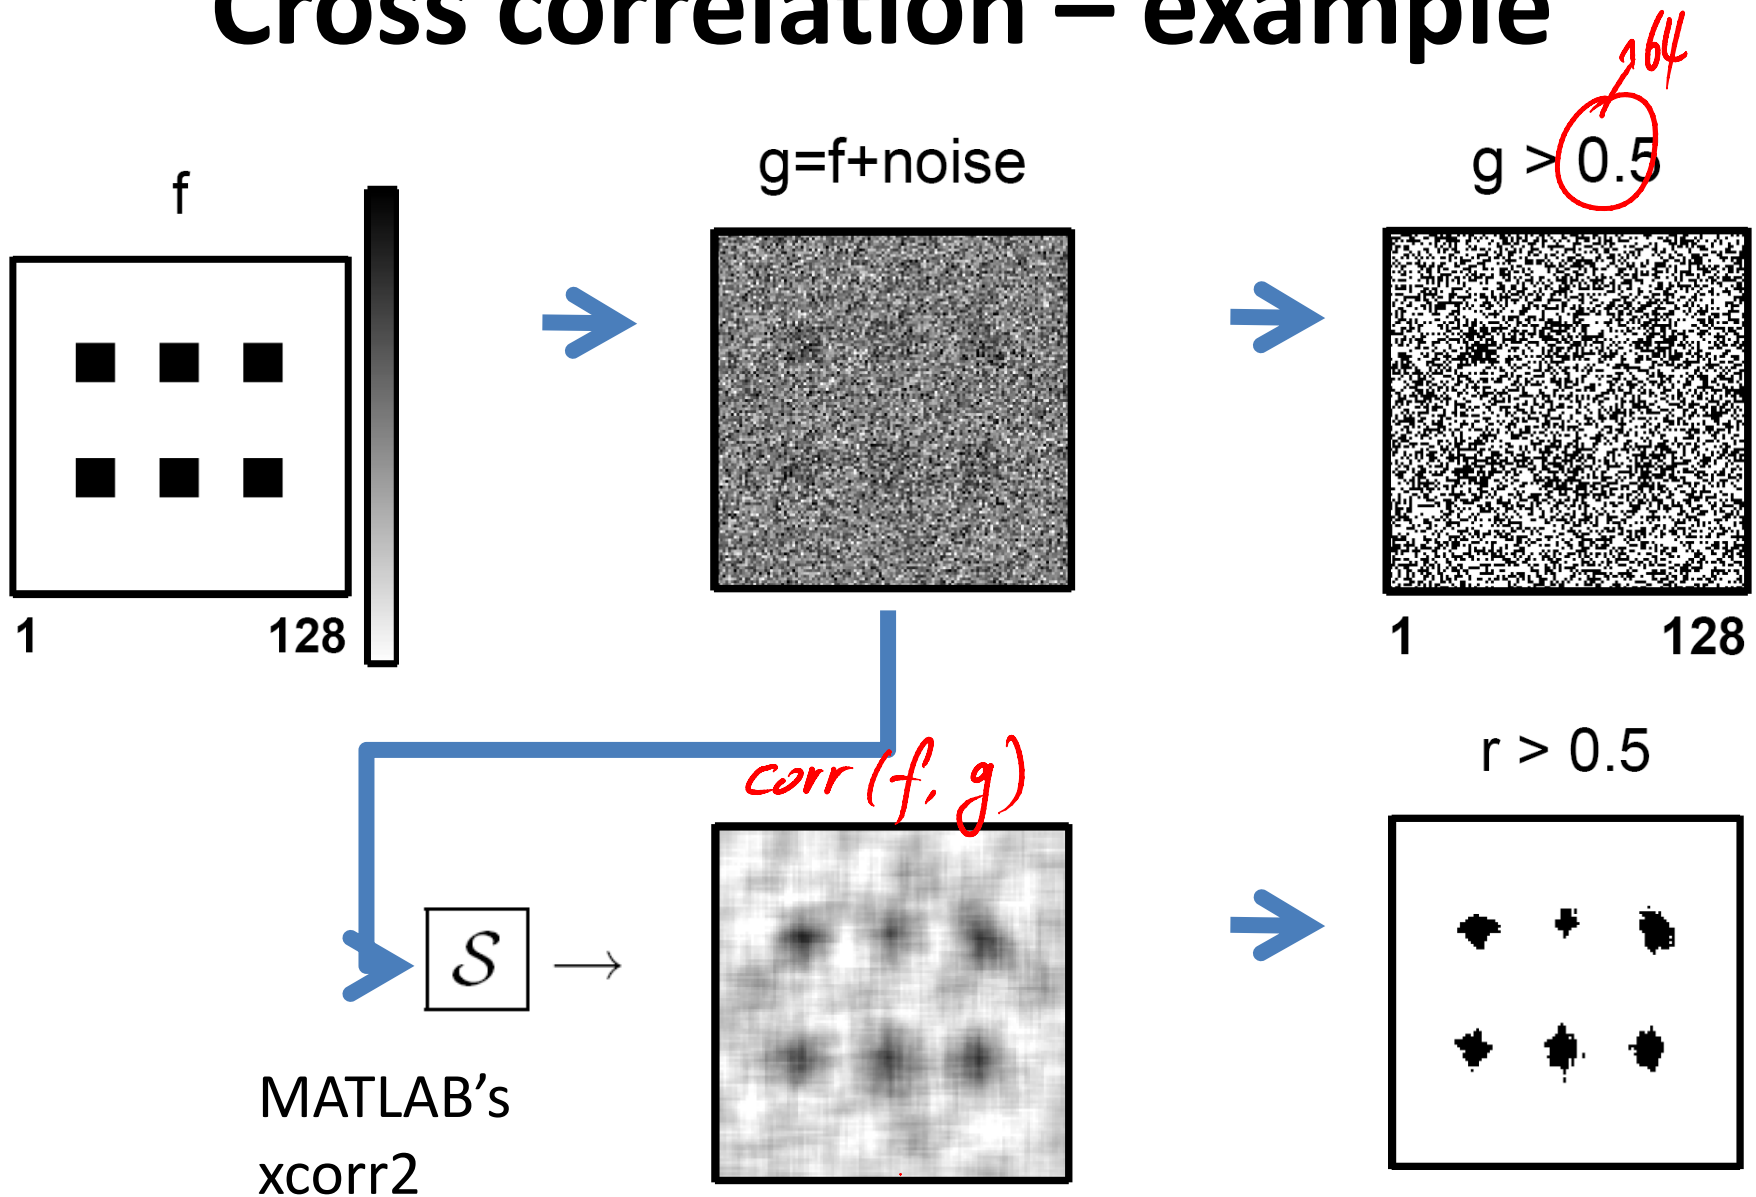
$$\begin{aligned} r_{fg}[k, l] &\triangleq \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} f[n, m] g^*[n - k, m - l] \\ &= \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} f[n + k, m + l] g^*[n, m], \quad k, l \in \mathbb{Z}. \end{aligned}$$

(k, l) is called the **lag**

- Equivalent to a convolution without the flip

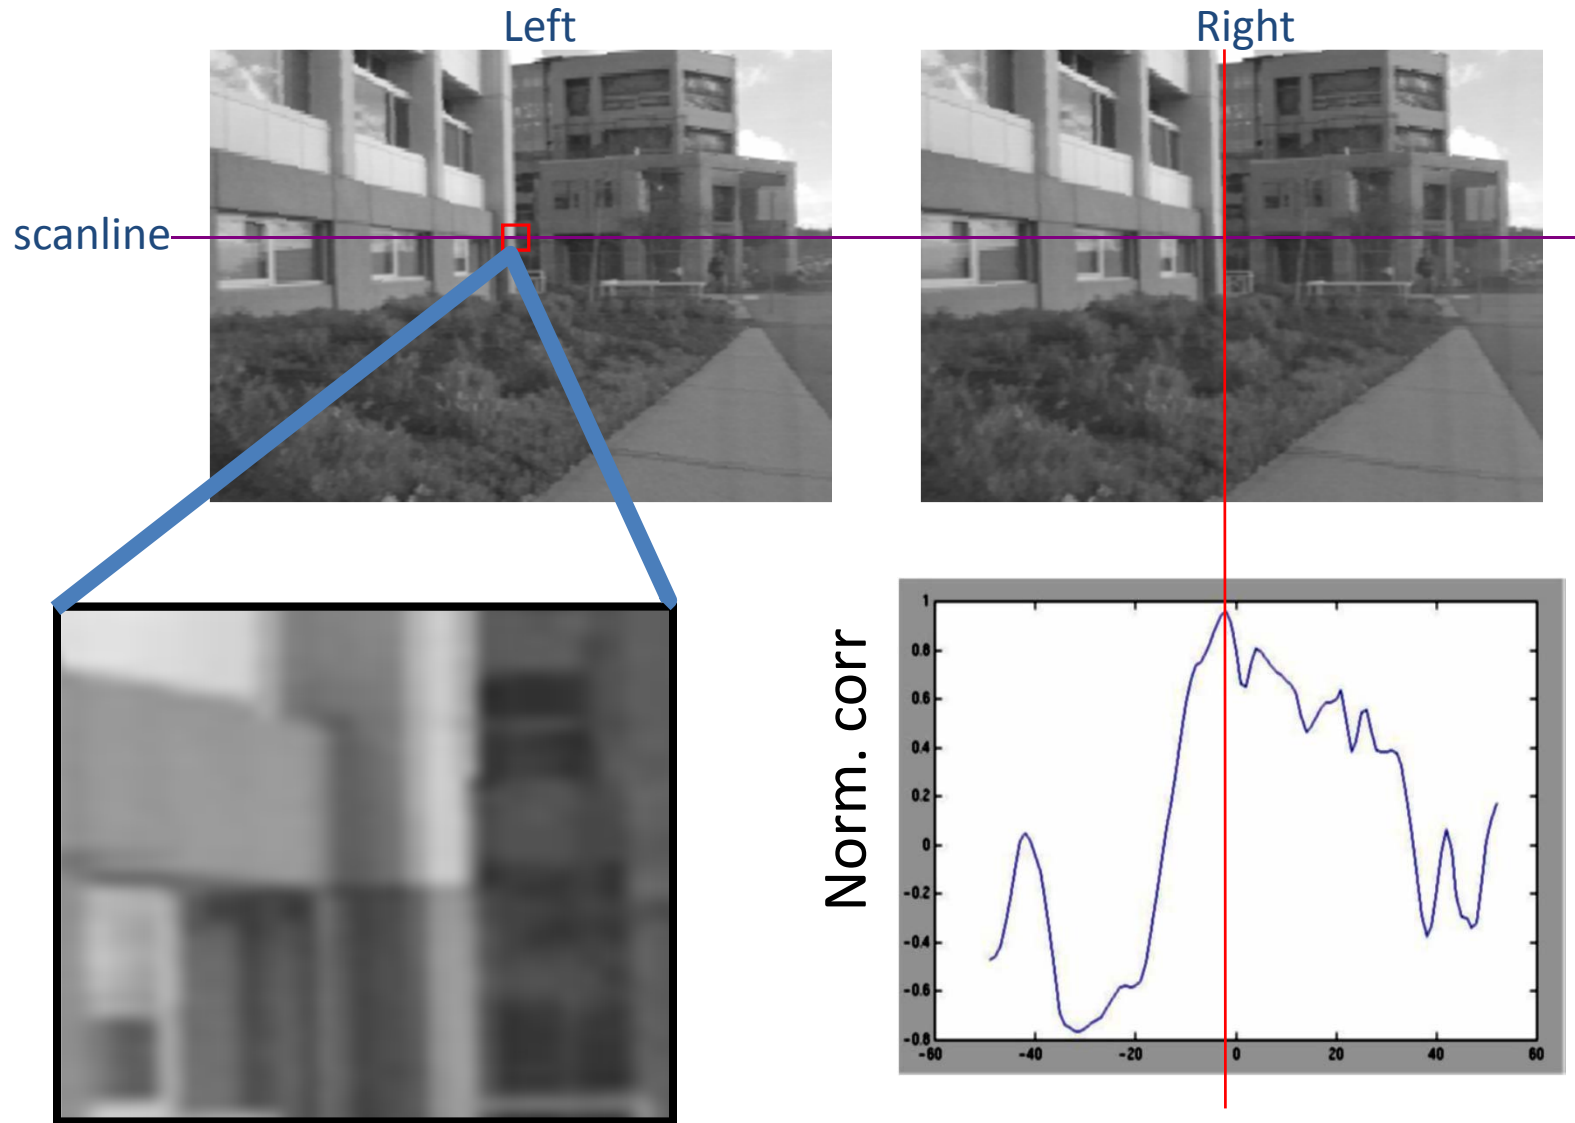
$$r_{fg}[n, m] = f[n, m] ** g^*[-n, -m]$$

Cross correlation – example



Courtesy of J. Fessler

Cross correlation – example



Convolution vs. Correlation

- A **convolution** is an integral that expresses the amount of overlap of one function as it is shifted over another function.
 - convolution is a filtering operation
- **Correlation** compares the *similarity of two sets of data*. Correlation computes a measure of similarity of two input signals as they are shifted by one another. The correlation result reaches a maximum at the time when the two signals match best .
 - correlation is a measure of relatedness of two signals

What we will learn today?

- Images as functions
- Linear systems (filters)
- Convolution and correlation
- **Discrete Fourier Transform (DFT)**
- Sampling and aliasing

Some background reading:

Forsyth and Ponce, Computer Vision, Chapter 7 & 8

Jae S. Lim, Two-dimensional signal and image processing, Chapter 1, 4, 5

2D Discrete-Signal Fourier Transform (DSFT)

2D discrete-signal Fourier transform (DSFT) of a 2D discrete-space signal $g[n,m]$:

$$G(\omega_X, \omega_Y) = \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} g[n, m] e^{-i(\omega_X n + \omega_Y m)}$$

frequency space *signal space*

Inverse 2D DSFT:


$$g[n, m] = \frac{1}{(2\pi)^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} G(\omega_X, \omega_Y) e^{i(\omega_X n + \omega_Y m)} d\omega_X d\omega_Y$$

For memory refresher: Forsyth and Ponce, Ch 8

Jean Baptiste Joseph Fourier




Fourier's Transform: log(Magnitude)



DSFT
↔

Fourier's Transform: Phase



DSFT – properties

Shift:

$$g[n - n_0, m - m_0] \longleftrightarrow \overset{\text{original FFT}(g)}{G(\omega_X, \omega_Y)} e^{-j(\omega_X n_0 + \omega_Y m_0)} \overset{\text{phase shifted}}{}$$

Convolution:

$$f[n, m] ** h[n, m] \longleftrightarrow F(\omega_X, \omega_Y) H(\omega_X, \omega_Y)$$

Delta function:

$$\delta_2[n, m] \longleftrightarrow 1$$

$$\delta_2|[n - n_0, m - m_0] \longleftrightarrow e^{-j(\omega_X n_0 + \omega_Y m_0)}$$

Example: DSFT of moving average filter




$$g[n, m] = \frac{1}{9} \sum_{k=-1}^1 \sum_{l=-1}^1 f[n - k, m - l]$$

$$(f * h)[m, n] = \frac{1}{9} \sum_{k,l} f[k, l] h[m - k, n - l]$$

$$\begin{aligned} & F(\omega_X, \omega_Y) H(\omega_X, \omega_Y) \\ &= \frac{1}{9} \sum_{n=-1}^1 \sum_{m=-1}^1 e^{-j\omega_X n} e^{-j\omega_Y m} \\ &= \frac{1}{9} [1 + 2 \cos \omega_X][1 + 2 \cos \omega_Y] \end{aligned}$$

h

	1	1	1
1	1	1	1
1	1	1	1

Function	Fourier transform
$g(x, y)$	$\int \int_{-\infty}^{\infty} g(x, y) e^{-i2\pi(ux+vy)} dx dy$
$\int \int_{-\infty}^{\infty} \mathcal{F}(g(x, y))(u, v) e^{i2\pi(ux+vy)} dudv$	$\mathcal{F}(g(x, y))(u, v)$
$\delta(x, y)$	1
$\frac{\partial f}{\partial x}(x, y)$	$u\mathcal{F}(f)(u, v)$
$0.5\delta(x + a, y) + 0.5\delta(x - a, y)$	$\cos 2\pi au$
$e^{-\pi(x^2+y^2)}$	$e^{-\pi(u^2+v^2)}$
 $box_1(x, y)$	$\frac{\sin u}{u} \frac{\sin v}{v}$ 
$f(ax, by)$	$\frac{\mathcal{F}(f)(u/a, v/b)}{ab}$
$\sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} \delta(x - i, y - j)$	$\sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} \delta(u - i, v - j)$ 
$(f * g)(x, y)$	$\mathcal{F}(f)\mathcal{F}(g)(u, v)$
$f(x - a, y - b)$	$e^{-i2\pi(au+bv)} \mathcal{F}(f)$
$f(x \cos \theta - y \sin \theta, x \sin \theta + y \cos \theta)$	$\mathcal{F}(f)(u \cos \theta - v \sin \theta, u \sin \theta + v \cos \theta)$

orig def!

✓

✓



✗✗

Why is DFT important?

- Perform efficient linear convolution as product of DFTs
- Each DFT can be implemented using the FFT (Fast Fourier Transform) [see appendix for details]

What we will learn today?


- Images as functions
- Linear systems (filters)
- Convolution and correlation
- Discrete Fourier Transform (DFT)
- **Sampling and aliasing**

Some background reading:

Forsyth and Ponce, Computer Vision, Chapter 7 & 8

Jae S. Lim, Two-dimensional signal and image processing, Chapter 1, 4, 5

Sampling



Throw away every other row and column
to create a 1/2 size image

Sampling

- Down-sampling operation:
(trivial form of image compression)

$$g[n, m] = f[2n, 2m] = \begin{bmatrix} \ddots & & \vdots & & \\ \dots & f[-2, 2] & f[0, 2] & f[2, 2] & \dots \\ & f[-2, 0] & f[0, 0] & f[2, 0] & \\ f[-2, -2] & f[0, -2] & f[2, -2] & & \\ & \vdots & & & \ddots \end{bmatrix}$$

Why is a multi-scale representation useful?

- Find template matches at all scales
 - e.g., when finding hands or faces, we don't know what size they will be in a particular image
 - Template size is constant, but image size changes
- Efficient search for correspondence
 - look at coarse scales, then refine with finer scales
- Examining all levels of detail
 - Find edges with different amounts of blur
 - Find textures with different spatial frequencies (levels of detail)

Slide credit: David Lowe (UBC)

Things can go wrong! -- Aliasing


256x256

128x128

64x64

32x32

16x16




Things can go wrong! -- Aliasing

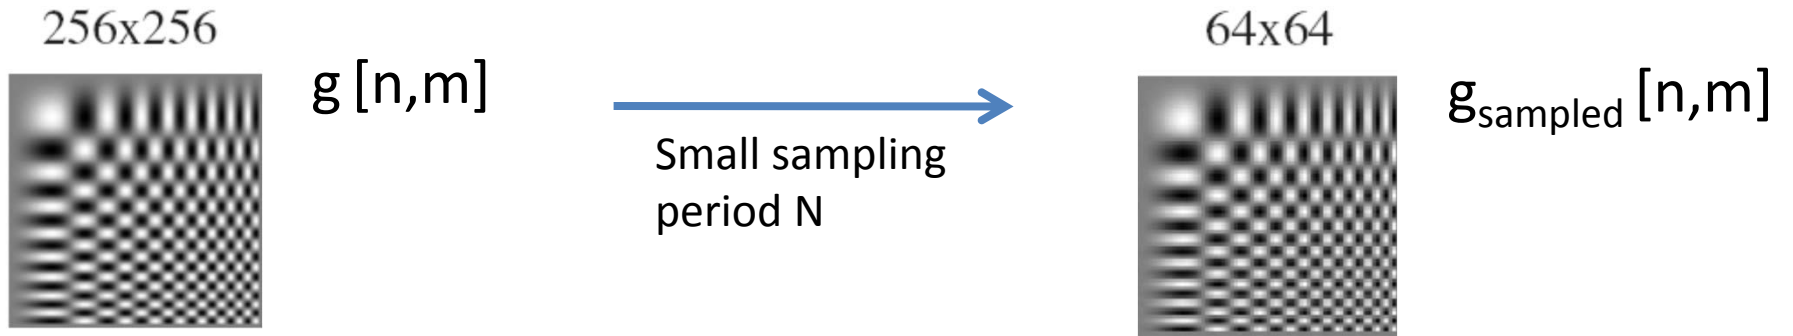




Sampling Theorem (Nyquist)


- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{\max}$
- f_{\max} = max frequency of the input signal.

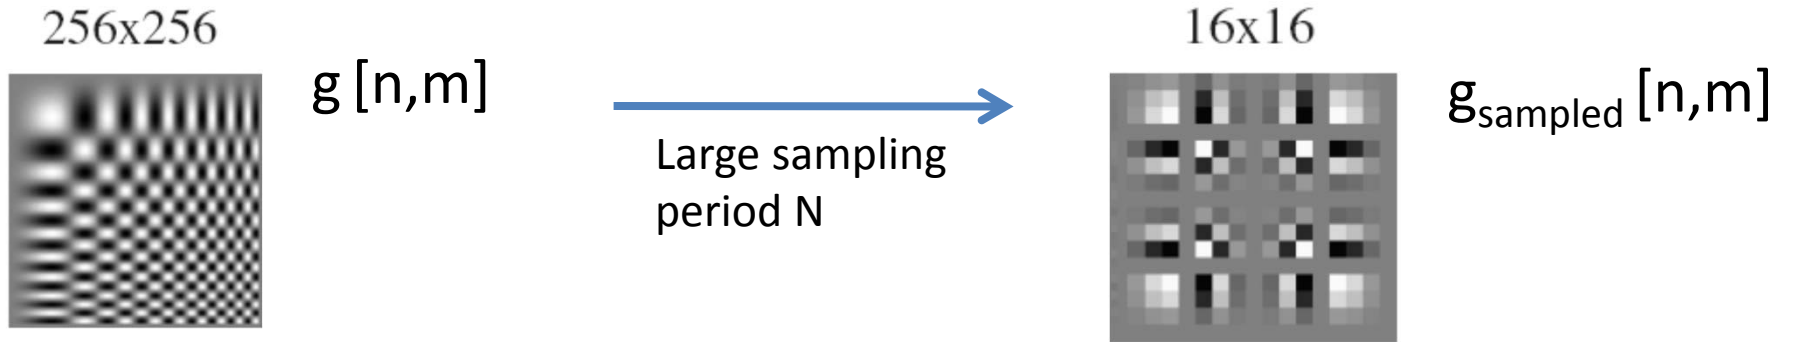




- N = sampling period
- T = periodicity of the replicas of the DFT (g)


$$T \sim 1/N$$





- N = sampling period
- T = periodicity of the replicas of the DFT (g)

$$T \sim 1/N$$




Anti-aliasing

Solutions:

- Sample more often
- Get rid of all frequencies that are greater than half the new sampling frequency
 - Will lose information - but it's better than aliasing
 - Apply a smoothing filter to remove high frequencies

Anti-aliasing


Apply a smoothing filter to remove high frequencies:



Sampling algorithm

Algorithm 7.1: Sub-sampling an Image by a Factor of Two


Apply a low-pass filter to the original image
(a Gaussian with a σ of between one
and two pixels is usually an acceptable choice).



Create a new image whose dimensions on edge are half
those of the old image

Set the value of the i, j 'th pixel of the new image to the value
of the $2i, 2j$ 'th pixel of the filtered image

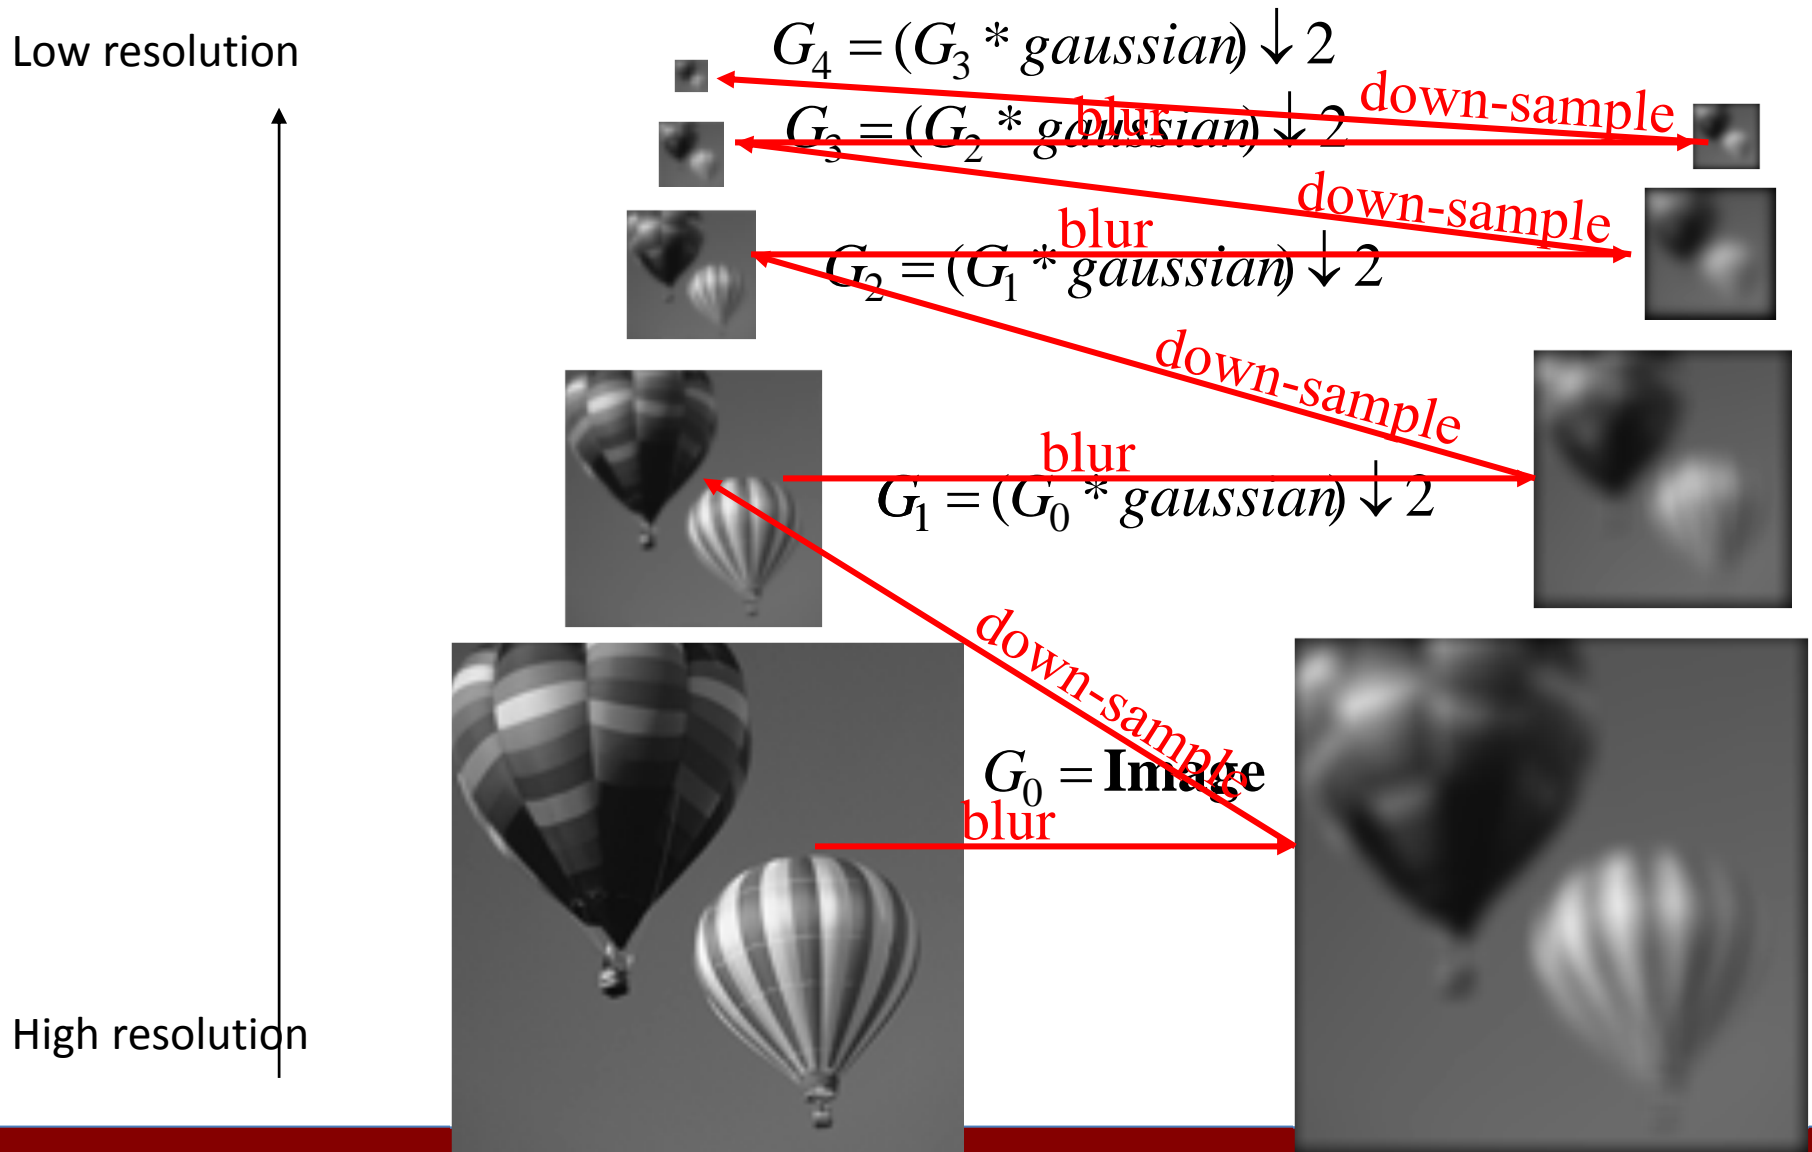
Resampling with Prior Smoothing




- Note: We cannot recover the high frequencies, but we can avoid artifacts by smoothing before resampling.

Image Source: Forsyth & Ponce

The Gaussian Pyramid



Gaussian Pyramid – Stored Information



Source: Irani & Basri

Summary: Gaussian Pyramid

- Construction: create each level from previous one
 - Smooth and sample
- Smooth with Gaussians, in part because
 - a Gaussian * Gaussian = another Gaussian
 - $G(\sigma_1) * G(\sigma_2) = G(\text{sqrt}(\sigma_1^2 + \sigma_2^2))$
- Gaussians are low-pass filters, so the representation is redundant once smoothing has been performed.
 - ⇒ There is no need to store smoothed images at the full original resolution.

Slide credit: David Lowe

Application: Vision system for TV remote control

- uses template matching

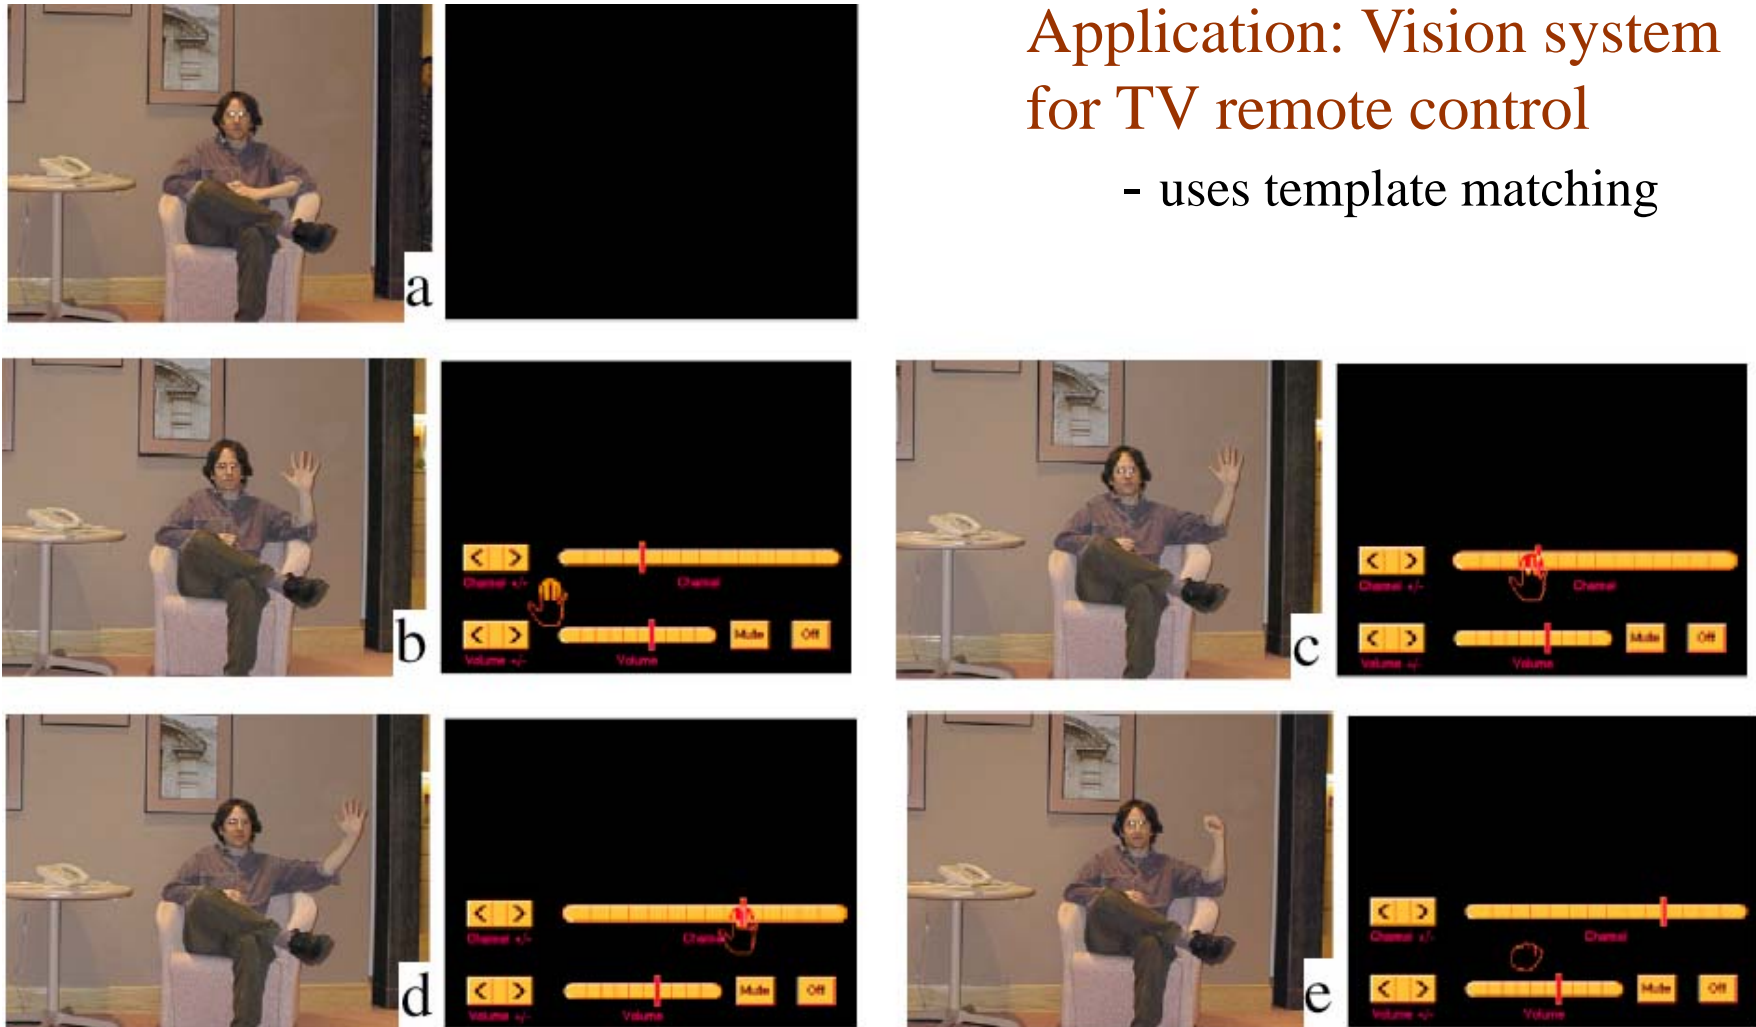


Figure from "Computer Vision for Interactive Computer Graphics," W.Freeman et al, IEEE Computer Graphics and Applications, 1998 copyright 1998, IEEE

What we have learned today?

- Images as functions
- Linear systems (filters)
- Convolution and correlation
- Discrete Fourier Transform (DFT)
- Sampling and aliasing

Appendix

Convergence

If g absolutely summable: $\sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} |g[n, m]| < \infty$

then

$$\lim_{N \rightarrow \infty} \sum_{n=-N}^N \sum_{m=-N}^N g[n, m] e^{-i(\omega_X n + \omega_Y m)} = G(\omega_X, \omega_Y)$$

If g is square summable
(energy signal):

$$E_g \triangleq \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} |g[n, m]|^2 < \infty$$

$$\int_{-\pi}^{\pi} \int_{-\pi}^{\pi} |G_N(\omega_X, \omega_Y) - G(\omega_X, \omega_Y)|^2 d\omega_X d\omega_Y \rightarrow 0$$



3-Oct-12

$$\tilde{g}[n, m]$$

Replicate and sum



$$g_d[n, m]$$

Truncate (space limited)



Impulses



$$g_s(x, y)$$

Impulse Sampling



$$g_a(x, y)$$

2D Sinc Convolution (bandlimited)



Sampling




2D Sinc Convolution (bandlimited)




2D Sinc Interpolation (bandlimited)




DTFS



DFT



DSFT




FT

80

$$G[k, l]$$

Sample



Dirichlet Interpolation (space limited)



$$G_d(\omega_X, \omega_Y)$$




$$\omega_X = 2\pi\nu_X L$$

$$\omega_Y = 2\pi\nu_Y L$$

$$G_s(\nu_X, \nu_Y)$$

Replicate and sum



Lowpass Filter (bandlimited)



$$G_a(\nu_X, \nu_Y)$$



Fast Fourier transforms (FFT)

- Brute-force evaluation of the 2D DFT would require $O((NM)^2)$ flops

$$X[k, l] \triangleq \begin{cases} \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} x[n, m] e^{-i2\pi(kn/N + lm/M)} \\ 0, \end{cases}$$

$k = 0, \dots, N - 1, l = 0, \dots, M - 1$
otherwise.

- DFT is a **separable operation**
→ we can reduce greatly the computation

Fast Fourier transforms (FFT)

- DFT is a **separable operation**:

$$\begin{aligned} X[k, l] &= \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} x[n, m] e^{-i2\pi(kn/N + lm/M)} \\ &= \sum_{n=0}^{N-1} e^{-i2\pi kn/N} \left[\sum_{m=0}^{M-1} x[n, m] e^{-i2\pi lm/M} \right] \end{aligned}$$

- Apply the 1D DFT to each column of the image, and then apply the 1D DFT to each row of the result.
- Use the **fast Fourier transform (1-D FFT)** for these 1D DFTs!

Fast Fourier transforms (FFT)

- FFT computational efficiency:
 - **inner set of 1D FFTs** : $N \mathcal{O}(M \log M)$
 - **outer set of 1D FFTs**: $M \mathcal{O}(N \log N)$
 - **Total**: $\mathcal{O}(MN \log MN)$ flops
 - A critical property of FFT is that $N = 2^k$ with $k = \text{integer}$
- If $x = 512 \times 512$ image \rightarrow
saving is a factor of 15000 relative to the brute-force 2D DFT!

Matlab: `fft2 = fft(fft(x).')`

FFT & Efficiency

- **General goal:** perform efficient linear convolution
- Perform convolution as product of DFTs
- **Pros:** DFT can be implemented using the FFT (fast fourier transform)
 - FFT is very efficient (fast!)
- **Cons:** DFT perform circular convolution
 - Compensate the wrap-around effect
- **Cons:** Online-memory storage
 - Use the overlap-add method or overlap-save method

FFT & Efficiency

- Suppose we wish to convolve a **256 × 256** image with a **17 × 17** filter.
- The result will be **272 × 272**.
- The smallest prime factors of 272 is **2**.
- So one could pad to a **512 × 512 image**
 - Note: only 28% of the final image would be the part we care about - the rest would be zero in exact arithmetic.
- Handling to a 512 × 512 image requires **much memory**

→ Use overlap-add method