


# Lecture 4

## Image Filtering II: Fourier Domain



How would you generate this function?




=

?


+

?


How would you generate this function?



=



+




$$f(x) = \sin(2\pi x) + \frac{1}{3} \sin(2\pi 3x)$$

$$\sin(2\pi x)$$

$$\frac{1}{3} \sin(2\pi 3x)$$


# Examples

How would you generate this function?




square wave


$\approx$



+




$=$




# Examples

How would you generate this function?




square wave


$\approx$



+




$=$




# Examples

How would you generate this function?




square wave


$\approx$



+




$=$




# Examples

How would you generate this function?




square wave

$\approx$



+




$=$




# Examples

How would you generate this function?




square wave


$\approx$



+




$=$



How would you express this mathematically?



# Examples



square wave


=

$$A \sum_{k=1}^{\infty} \frac{1}{k} \sin(2\pi kx)$$

infinite sum of sine waves

How would could you visualize this in the frequency domain?

# Examples




square wave

=

$$A \sum_{k=1}^{\infty} \frac{1}{k} \sin(2\pi kx)$$

infinite sum of sine waves

magnitude




# Fourier Series

$$A \sin(\omega x + \phi)$$

Diagram illustrating the components of the Fourier series equation  $A \sin(\omega x + \phi)$ :


- $A$ : amplitude
- $\sin$ : sinusoid
- $\omega$ : angular frequency
- $x$ : variable
- $\phi$ : phase

Fourier's claim:  
Add enough of these  
to get any *periodic* signal you  
want!



# Visualizing the frequency spectrum


Recall the temporal domain visualization




# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3} \sin(2\pi 3kx)$$



not visualizing the  
symmetric negative part



↑  
signal average (zero  
for a sine wave with  
no offset)



Need to understand this to  
understand the 2D version!

# Examples


Spatial domain visualization

Frequency domain visualization

1D



2D





?

# Examples



Spatial domain visualization

Frequency domain visualization

1D




2D




What do the three dots correspond to?

# Examples



+






=

?




# Examples



?

# Examples



# Fourier transform

Fourier transform

inverse Fourier transform

continuous

$$F(k) = \int_{-\infty}^{\infty} f(x) e^{-j2\pi kx} dx$$

$$f(x) = \int_{-\infty}^{\infty} F(k) e^{j2\pi kx} dk$$

discrete

$$F(k) = \frac{1}{N} \sum_{x=0}^{N-1} f(x) e^{-j2\pi kx/N}$$

$k = 0, 1, 2, \dots, N-1$

$$f(x) = \sum_{k=0}^{N-1} F(k) e^{j2\pi kx/N}$$

$x = 0, 1, 2, \dots, N-1$

# Fourier transform

Where is the connection to the ‘summation of sine waves’ idea?

$$f(x) = \sum_{k=0}^{N-1} F(k) e^{j2\pi kx/N}$$

Euler's formula

$$e^{j\theta} = \cos \theta + j \sin \theta$$

sum over frequencies

$$f(x) = \sum_{k=0}^{N-1} F(k) \left\{ \cos(2\pi kx) + j \sin(2\pi kx) \right\}$$

scaling parameter

wave components

# 2D Fourier Transform

## Definition

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-j2\pi(ux+vy)} dx dy,$$
$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) e^{j2\pi(ux+vy)} du dv$$


where  $u$  and  $v$  are spatial frequencies.

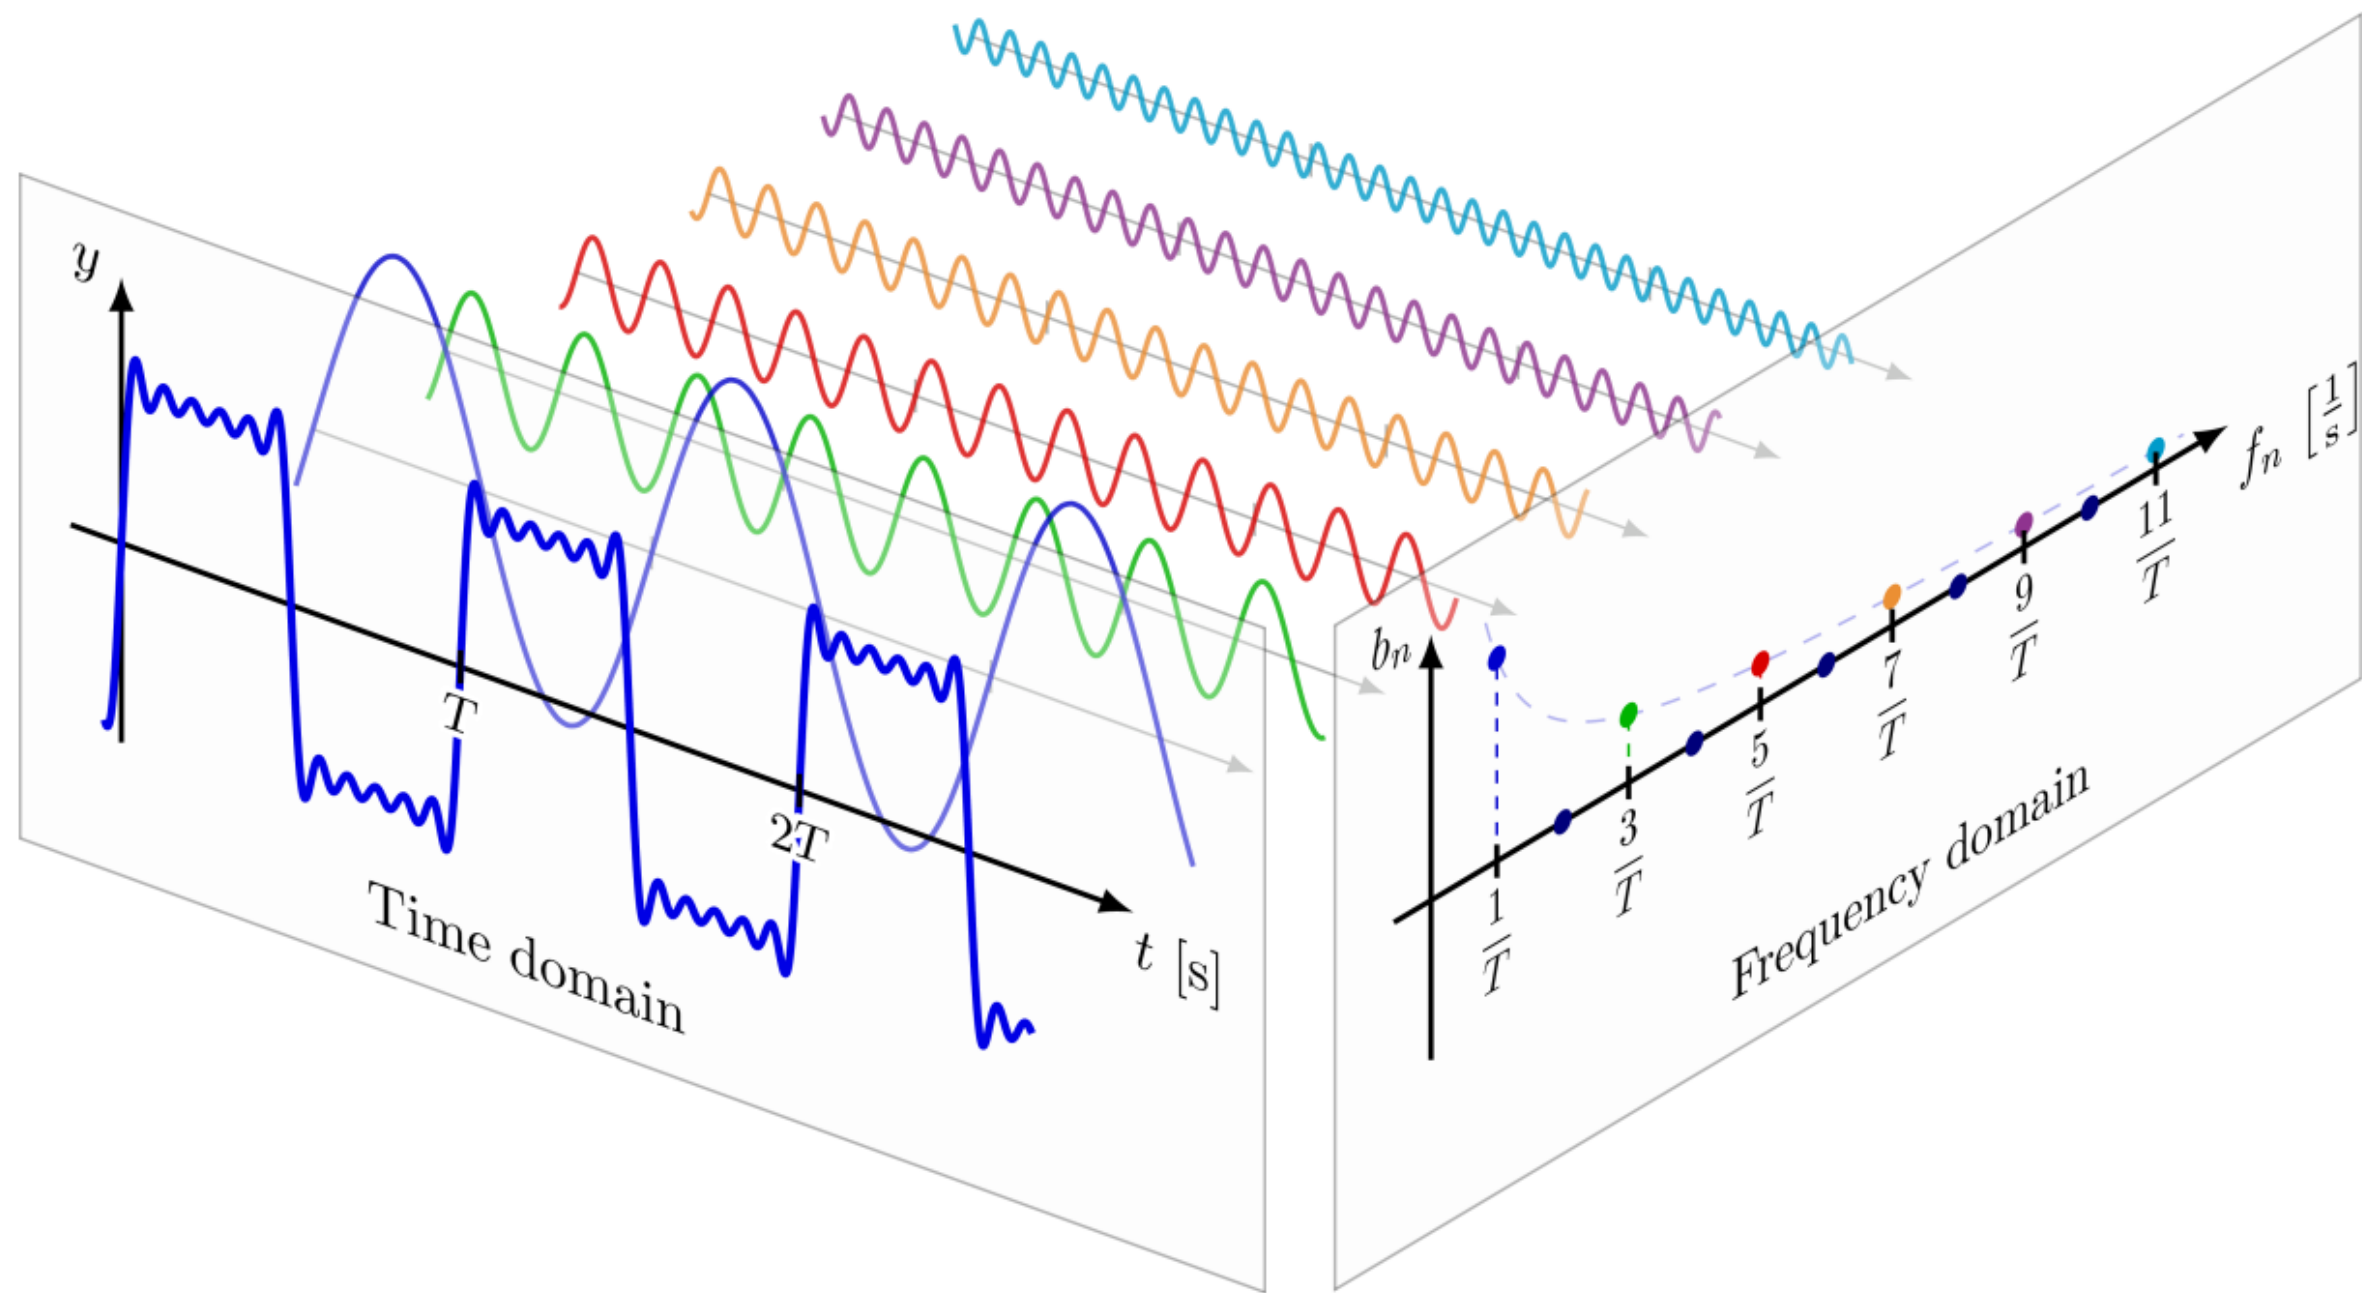
Also will write FT pairs as  $f(x, y) \Leftrightarrow F(u, v)$ .

- $F(u, v)$  is complex in general,

$$F(u, v) = F_R(u, v) + jF_I(u, v)$$

- $|F(u, v)|$  is the **magnitude** spectrum
- $\arctan(F_I(u, v)/F_R(u, v))$  is the **phase** angle spectrum.





# The Discrete Fourier transform

Discrete Fourier Transform (DFT) transforms a signal  $f[n]$  into  $F[u]$  as:

$$F[u] = \sum_{n=0}^{N-1} f[n] \exp \left( -2\pi j \frac{un}{N} \right)$$

$$e^{ix} = \cos x + i \sin x$$

Discrete Fourier Transform (DFT) is a linear operator. Therefore, we can write:

$$\mathbf{F} = \exp \left( -2\pi j \frac{un}{N} \right) \mathbf{f}$$

NxN array



# For images, the 2D DFT


1D Discrete Fourier Transform (DFT) transforms a signal  $f[n]$  into  $F[u]$  as:

$$F[u] = \sum_{n=0}^{N-1} f[n] \exp \left( -2\pi j \frac{un}{N} \right)$$

2D Discrete Fourier Transform (DFT) transforms an image  $f[n,m]$  into  $F[u,v]$  as:

$$F[u, v] = \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} f[n, m] \exp \left( -2\pi j \left( \frac{un}{N} + \frac{vm}{M} \right) \right)$$

# Simple Fourier transforms




$$\cos \left( 2\pi \left( \frac{u_0 n}{N} + \frac{v_0 m}{M} \right) \right)$$



$$\frac{1}{2} (\delta [u - u_0, v - v_0] + \delta [u + u_0, v + v_0])$$

# Simple Fourier transforms

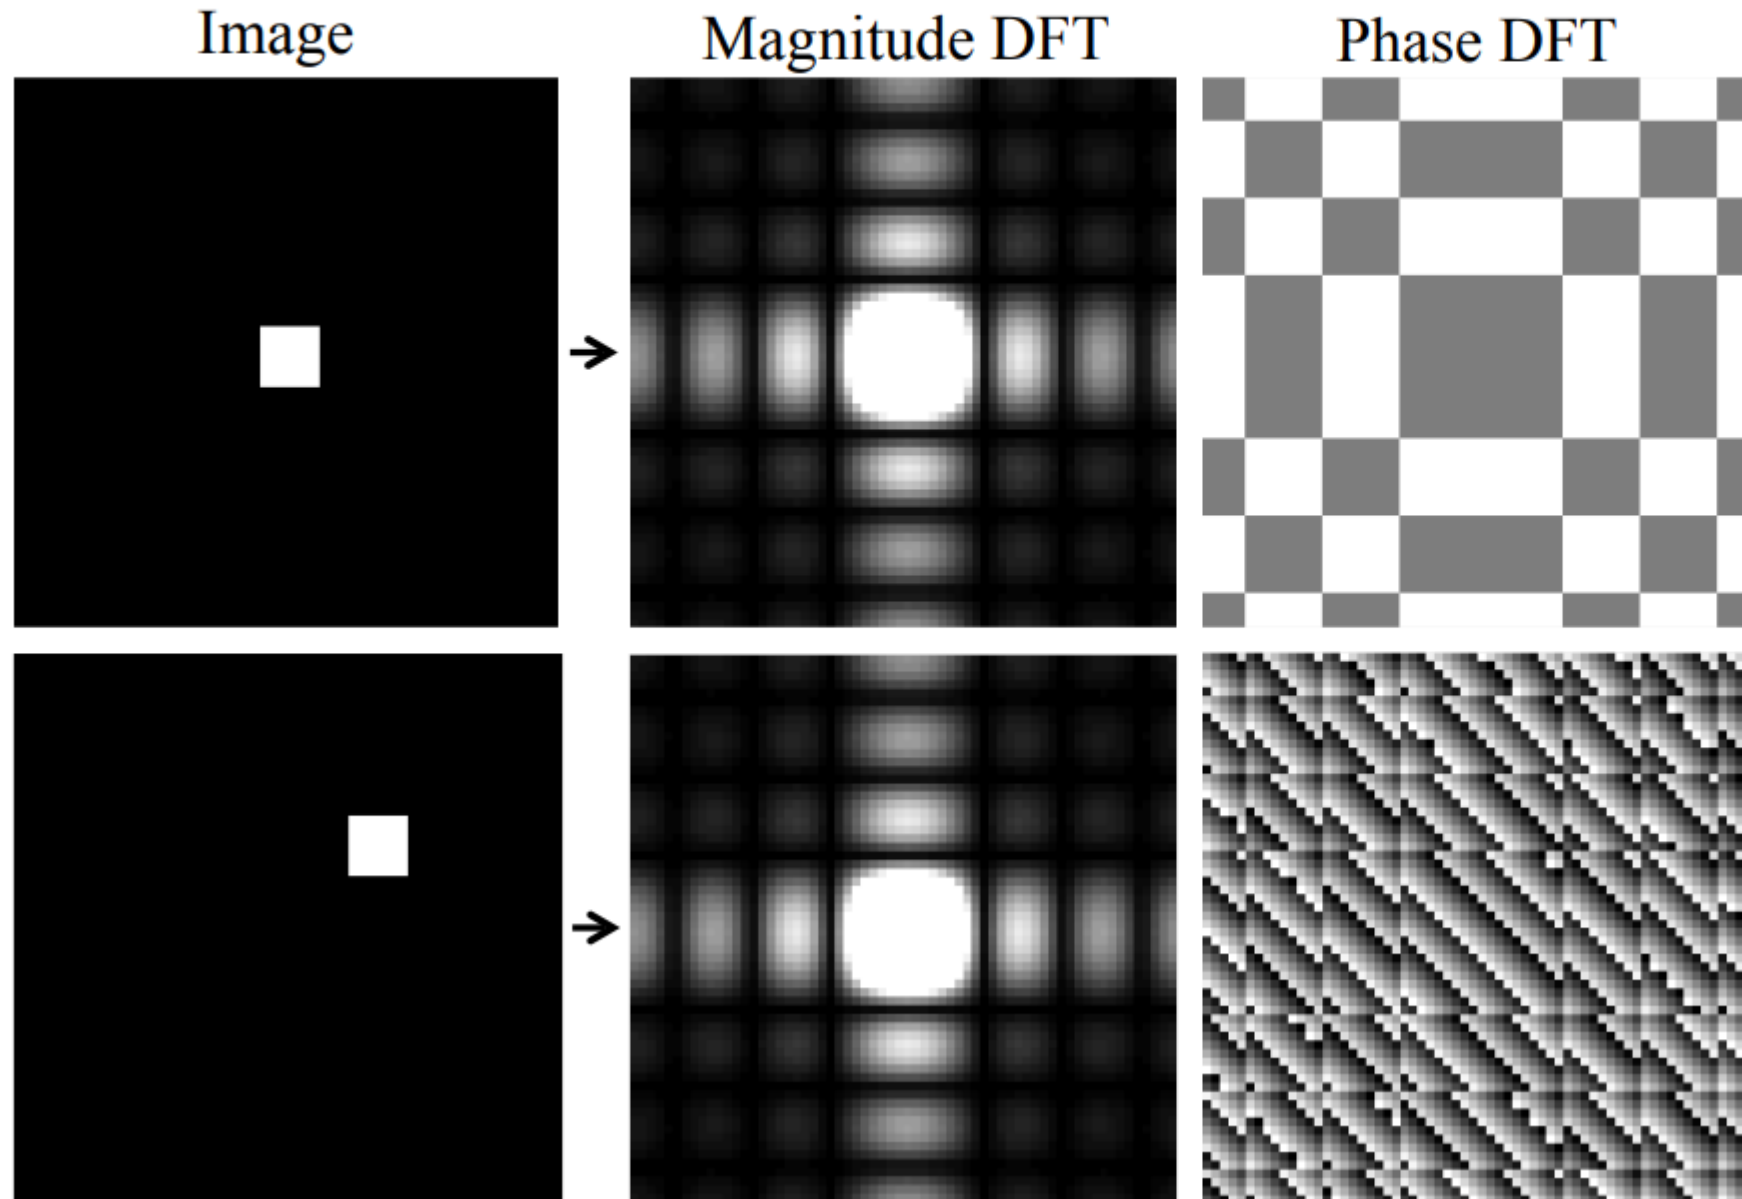
Image



DFT Amplitude



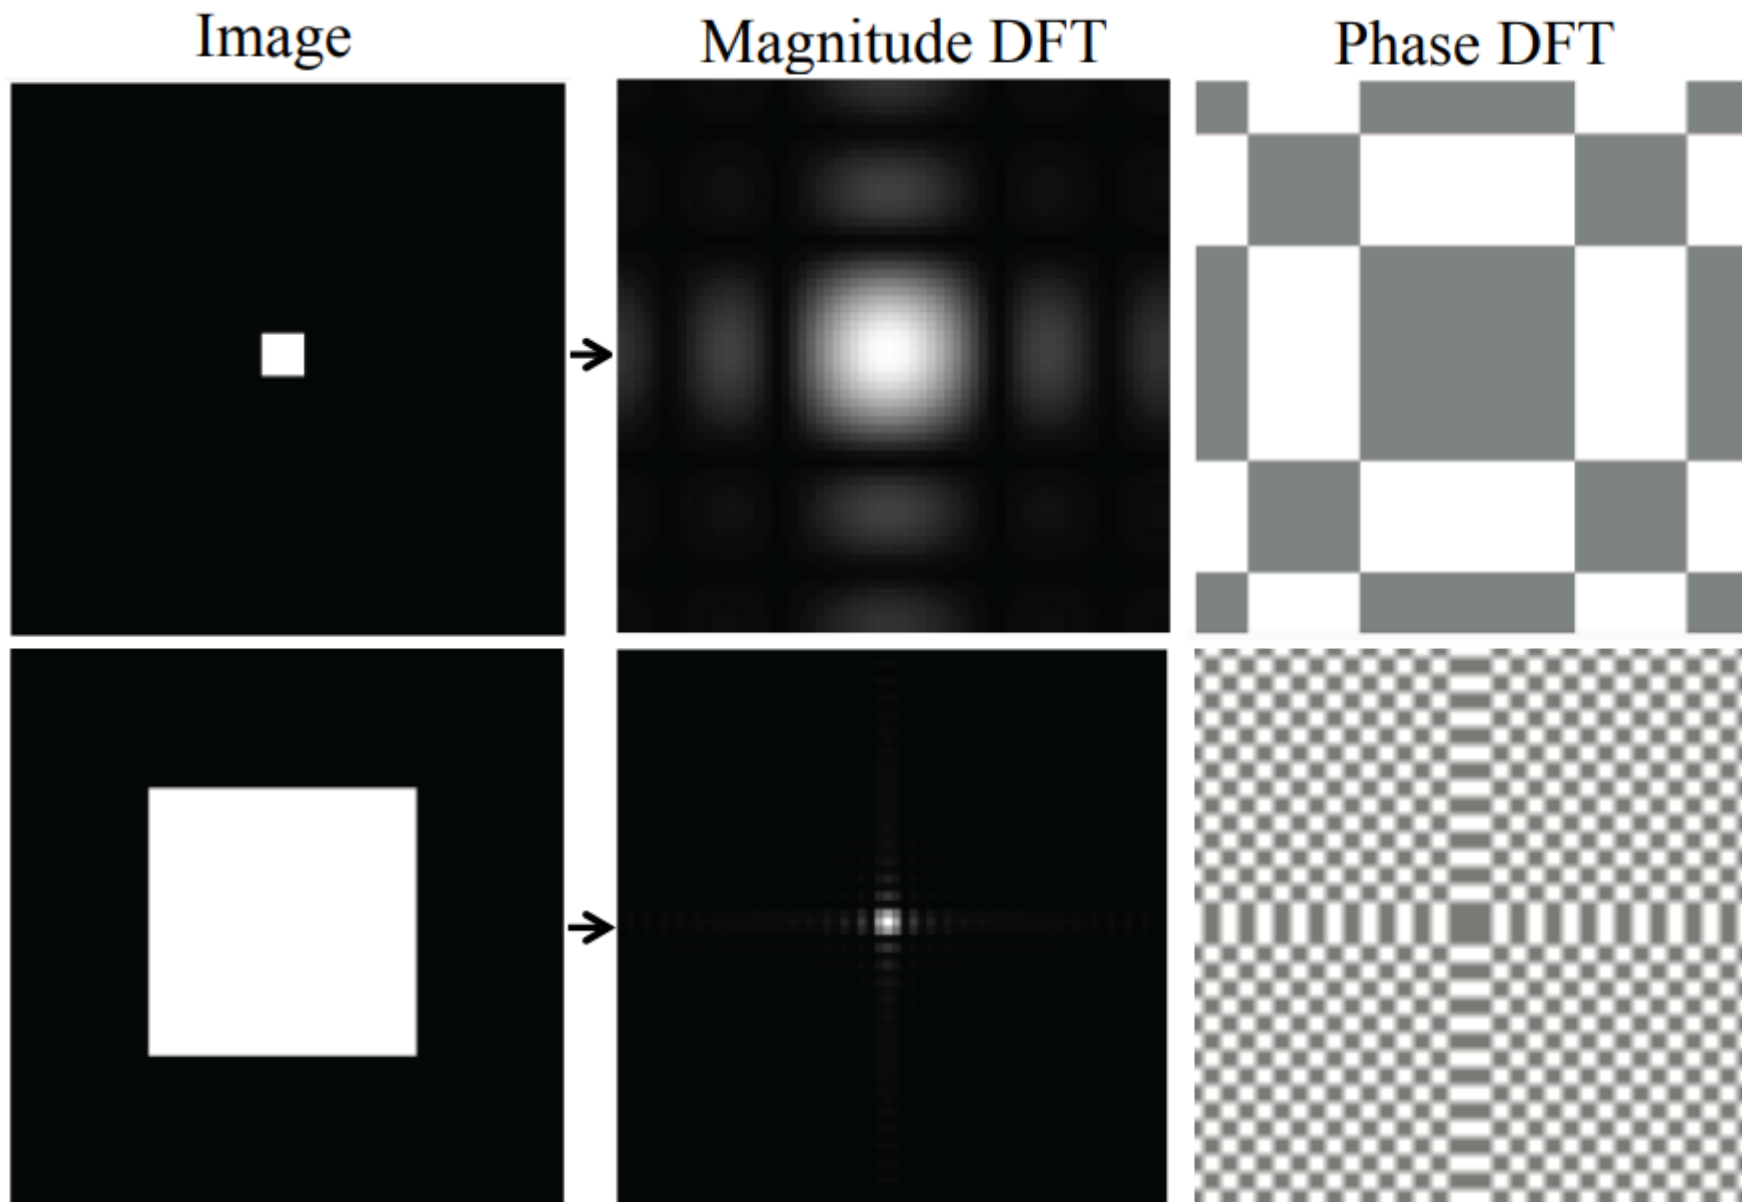
# Some important Fourier transforms



## Translation

Shifts of an image only produce changes on the phase of the DFT.

# Some important Fourier transforms




## Scale


Small image details  
produce content in high  
spatial frequencies

# Some important Fourier transforms



Image



Magnitude DFT



Phase DFT



## Orientation


A line transforms to a line oriented perpendicularly to the first.




# Magnitude & Phase

Magnitude encodes most of the color (intensity) information


Phase encodes most of the "location" information




Original



Magnitude



Phase





Magnitude Only




Phase Only

# The DFT Game: find the right pairs


Images




DFT  
magnitude



$f_x$ (cycles/image pixel size)



$f_x$ (cycles/image pixel size)





$f_x$ (cycles/image pixel size)




# The DFT Game: find the right pairs


Images




DFT  
magnitude



$f_x$ (cycles/image pixel size)







$f_x$ (cycles/image pixel size)



$f_x$ (cycles/image pixel size)

# Frequency Domain of the Image

## Fourier transforms of natural images






original



amplitude

phase

# More examples







mag



phase



## Fourier transforms of natural images



original

amplitude

phase

“We generally do not display phase images because most people who see them shortly thereafter succumb to hallucinogenics or end up in a Tibetan monastery” – John Brayer


# Magnitude only and phase only reconstructions



Reconstruction using  
magnitude only

Top Left Photo: Ralph's  
magnitude is the same,  
Phase = 0


Top Right Photo: Meg's  
magnitude is the same,  
Phase = 0



Reconstruction using  
phase only

Top Left Photo: Ralph's  
magnitude normalized to  
one, Phase is the same





Top Right Photo: Meg's  
magnitude normalized to  
one, Phase is the same





# Phase swapping

## Fourier transforms of natural images



What if we took the phase of each image, swapped it, and did the inverse Fourier transform?


original

amplitude

phase

# Phase Swapping

Image phase matters!



cheetah phase with zebra amplitude



zebra phase with cheetah amplitude

# The Convolution Theorem

The Fourier transform of the convolution of two functions is the product of their Fourier transforms:

$$\mathcal{F}\{g * h\} = \mathcal{F}\{g\}\mathcal{F}\{h\}$$

The inverse Fourier transform of the product of two Fourier transforms is the convolution of the two inverse Fourier transforms:


$$\mathcal{F}^{-1}\{gh\} = \mathcal{F}^{-1}\{g\} * \mathcal{F}^{-1}\{h\}$$

Convolution in spatial domain is equivalent to multiplication in frequency domain!




# Low-Pass Filter (Pixel Domain)

- low-pass filter: convolution in primal domain  $b = x * c$




# Low-Pass Filter (Frequency Domain)

- low-pass filter: multiplication in frequency domain  $F\{b\} = F\{x\} \cdot F\{c\}$




# High Pass Filter (Frequency Domain)






## Band-Pass filtering (Frequency Domain)



- edges with specific orientation (e.g., hat) are gone!



# The Convolution Theorem

The Fourier transform of the convolution of two functions is the product of their Fourier transforms:

$$\mathcal{F}\{g * h\} = \mathcal{F}\{g\}\mathcal{F}\{h\}$$

**Convolution** in the pixel domain = **multiplication** in the Fourier domain

The inverse Fourier transform of the product of two Fourier transforms is the convolution of the two inverse Fourier transforms:

$$\mathcal{F}^{-1}\{gh\} = \mathcal{F}^{-1}\{g\} * \mathcal{F}^{-1}\{h\}$$

Convolution in spatial domain is equivalent to multiplication in frequency domain!



Fourier Domain Filtering: Can be much faster for big filters because speed is independent of filter size

Convolution: Speed is proportional to filter size!

# Low Pass Filtering in Fourier Domain




$x$



$FT(x)$

# Low Pass Filtering in Fourier Domain



This is a low-pass filter in Fourier Domain


Look how it is centered around  $(0, 0)$  – it allows low frequencies and rejects high frequencies.

How can we apply this to the image?

Use the Convolution Theorem




# Low Pass Filtering in Fourier Domain



$FT(f)$

Fourier Transform of  
Low-Pass Filter


$\times$



$FT(x)$

Fourier Transform of Image

$=$




$$FT(x) \times FT(f) = FT(x * f)$$

Multiplication

Convolution

# Low Pass Filtering in Fourier Domain



$$FT(x) \times FT(f) = FT(x * f)$$

Multiplication



Convolution

# Low Pass Filtering in Fourier Domain



$$FT^{-1}[FT(x) \times FT(f)]$$

Low-Pass Filtered Image




$$FT(x) \times FT(f) = FT(x * f)$$


Multiplication

Convolution

# Low Pass Filtering in Fourier Domain




# High Pass Filtering in Fourier Domain



# Blurring / Smoothing


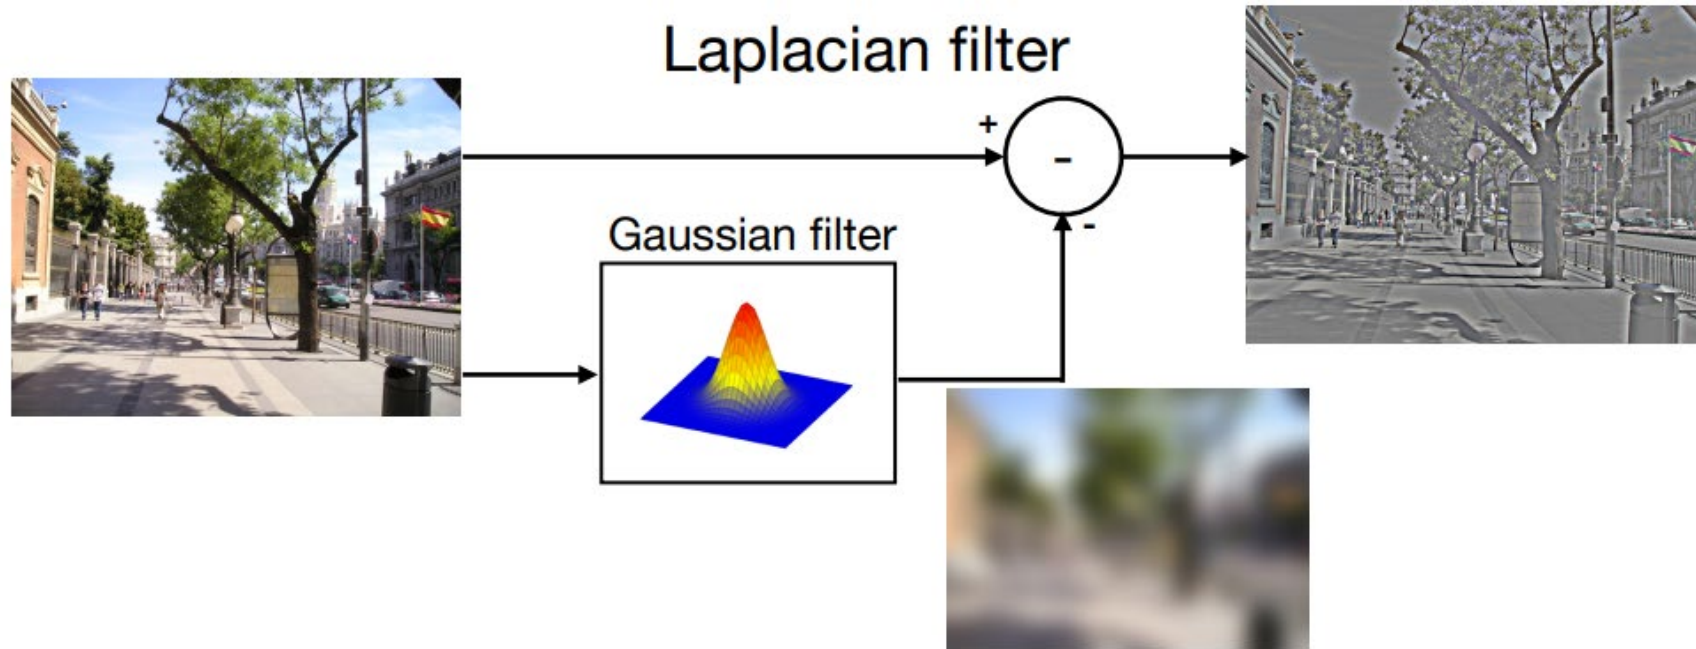


Gaussian filter





# Opposite of Blurring: Sharpening








# Gaussian Filter vs Laplacian Filter



Gaussian filter



Laplacian filter






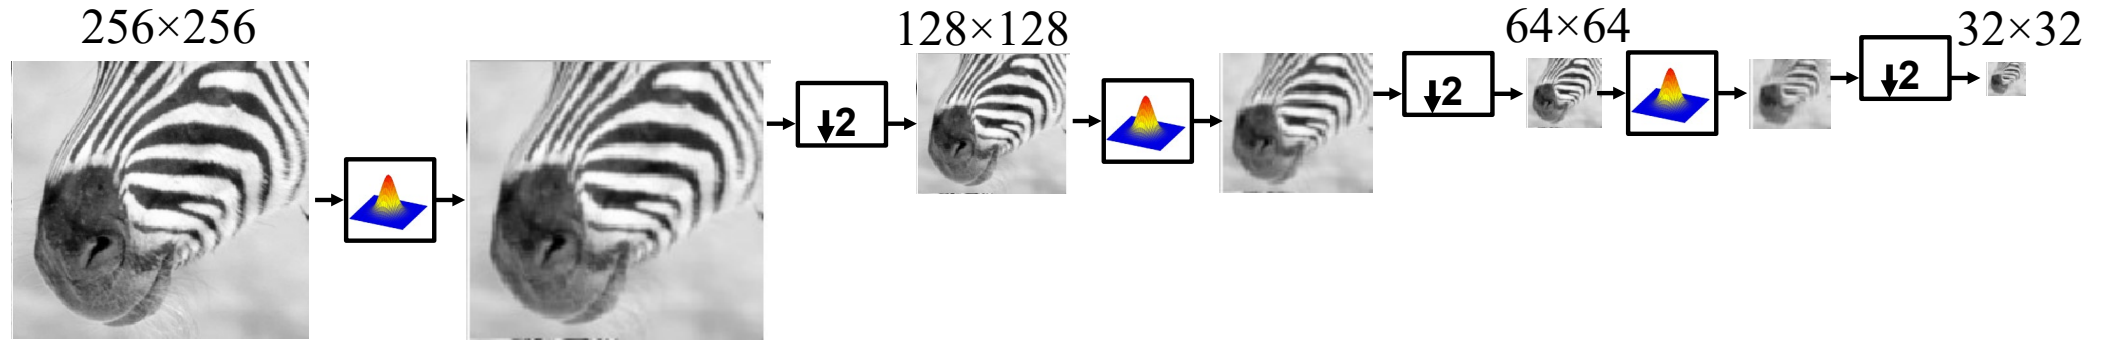
# The Gaussian pyramid

For each level

1. Blur input image with a Gaussian filter
2. Downsample image



# The Gaussian pyramid



# The Gaussian pyramid

512×512




(original image)

256×256



128×128



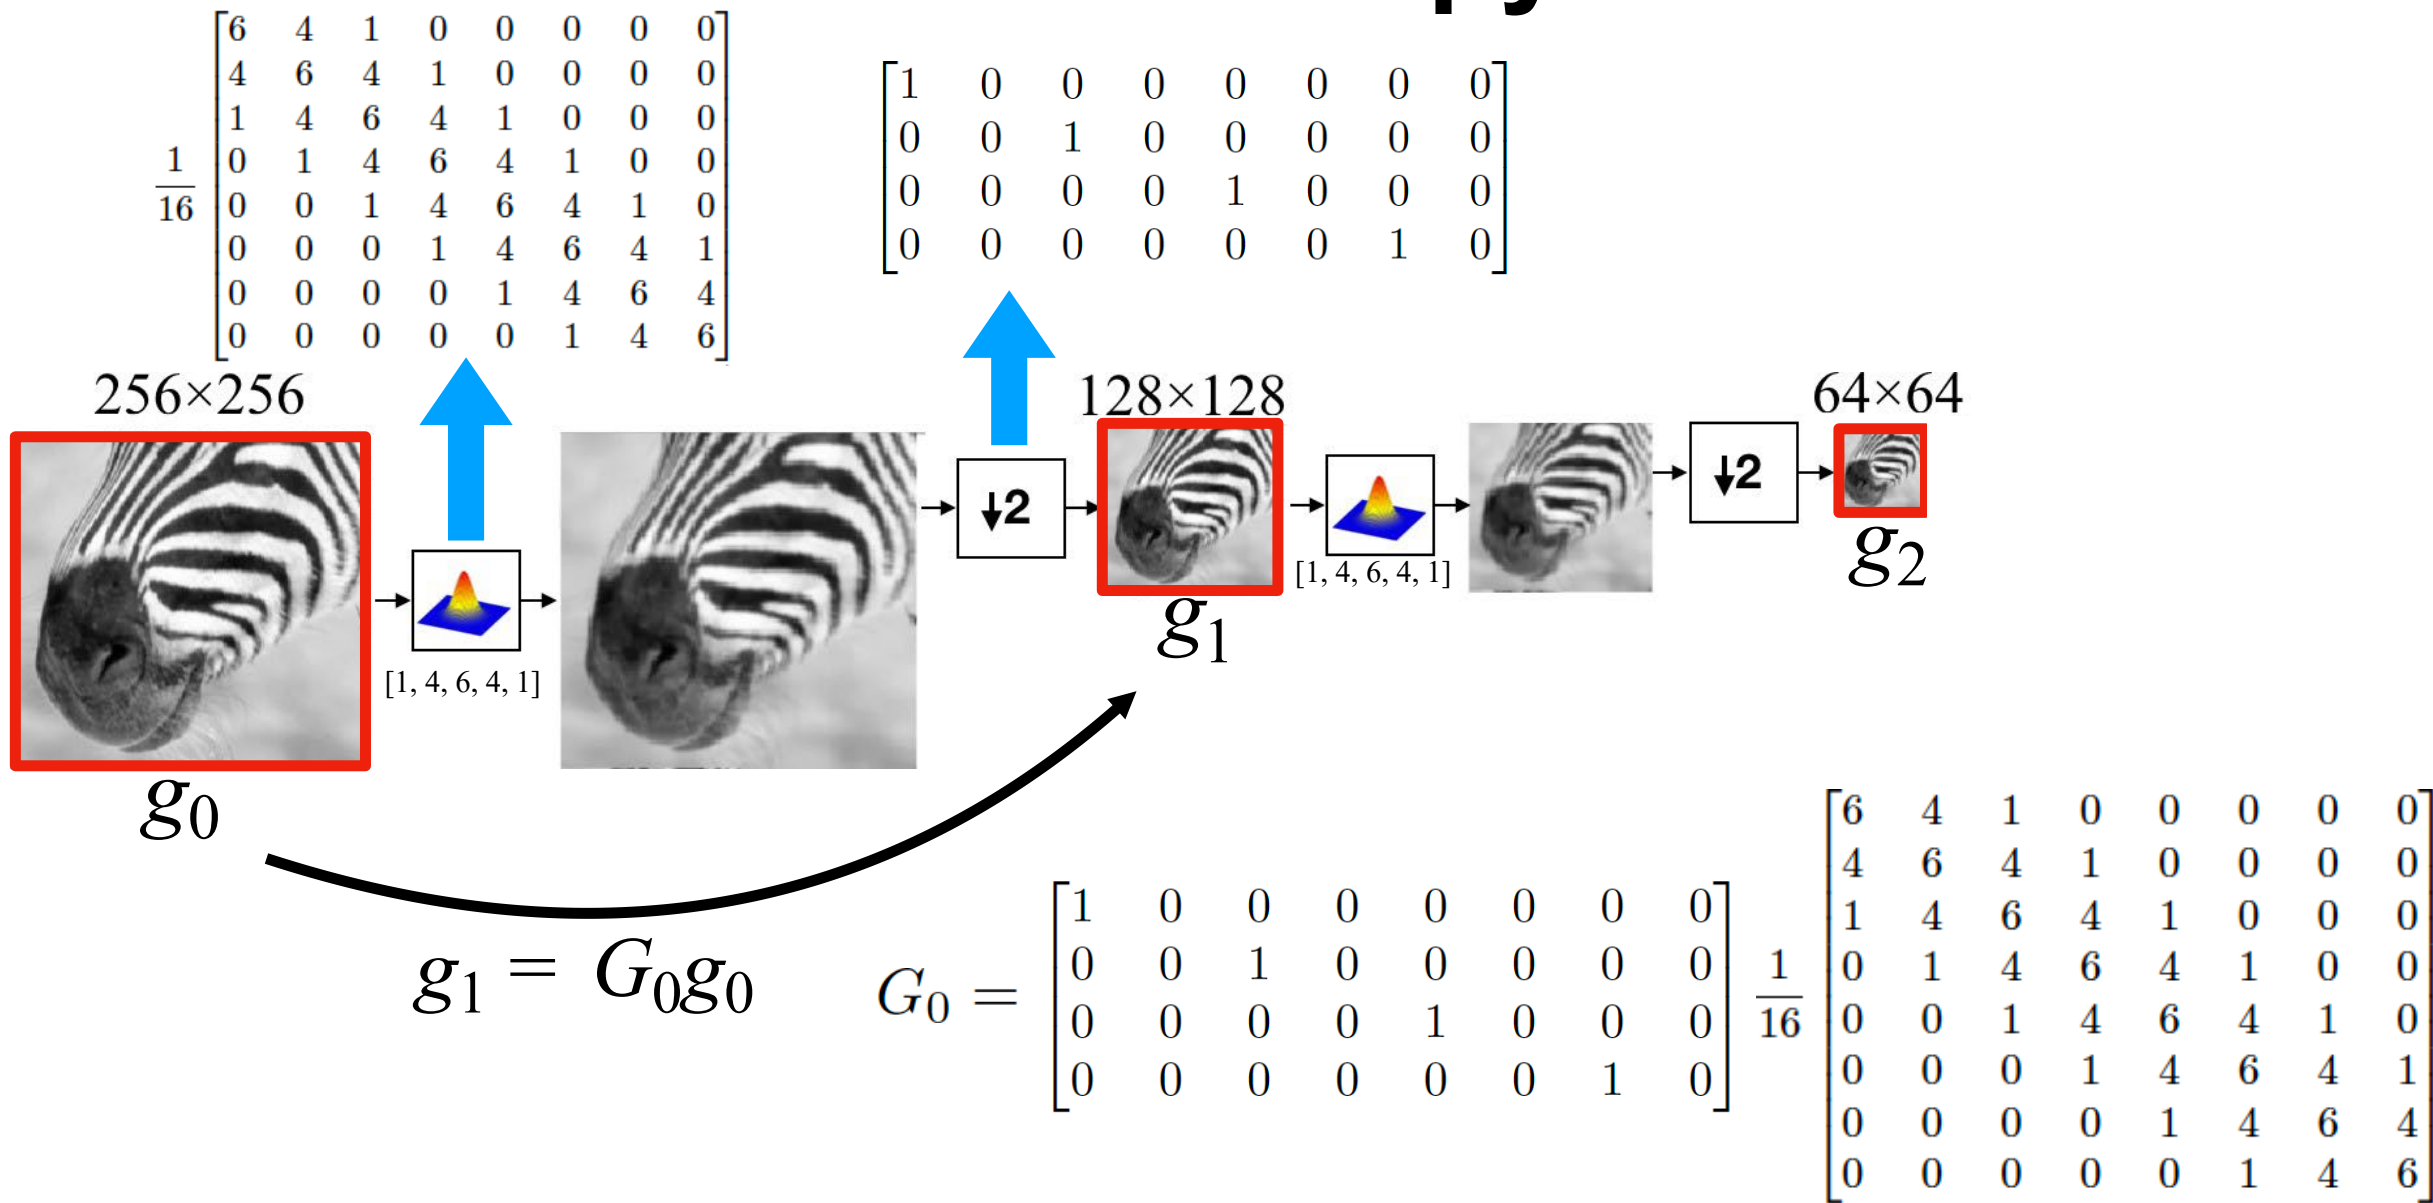
64×64




32×32



# The Gaussian pyramid



# The Gaussian pyramid

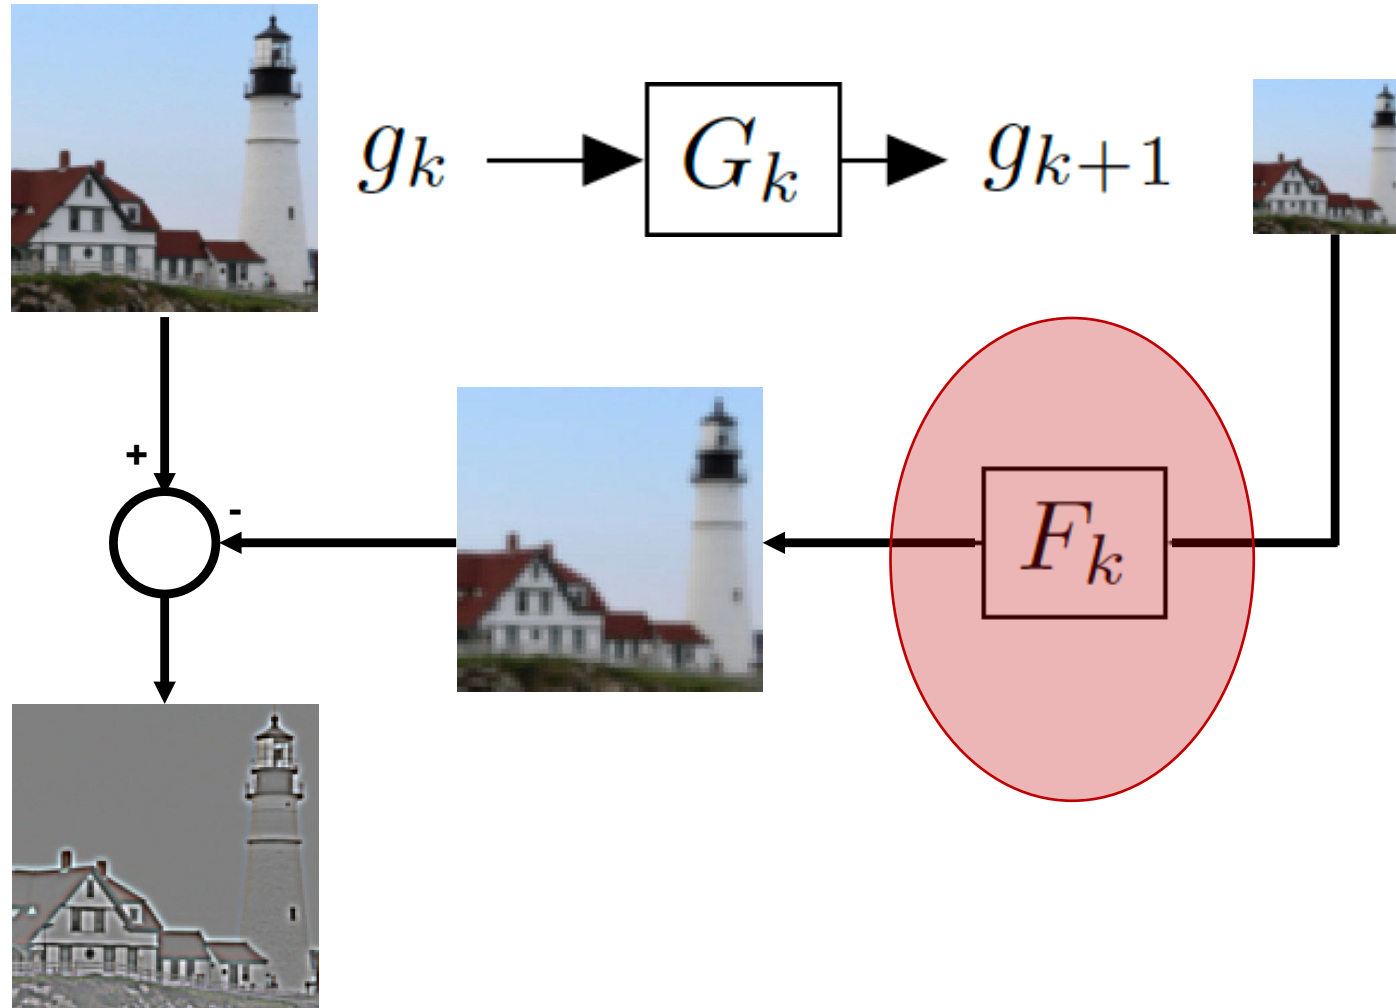


For each level

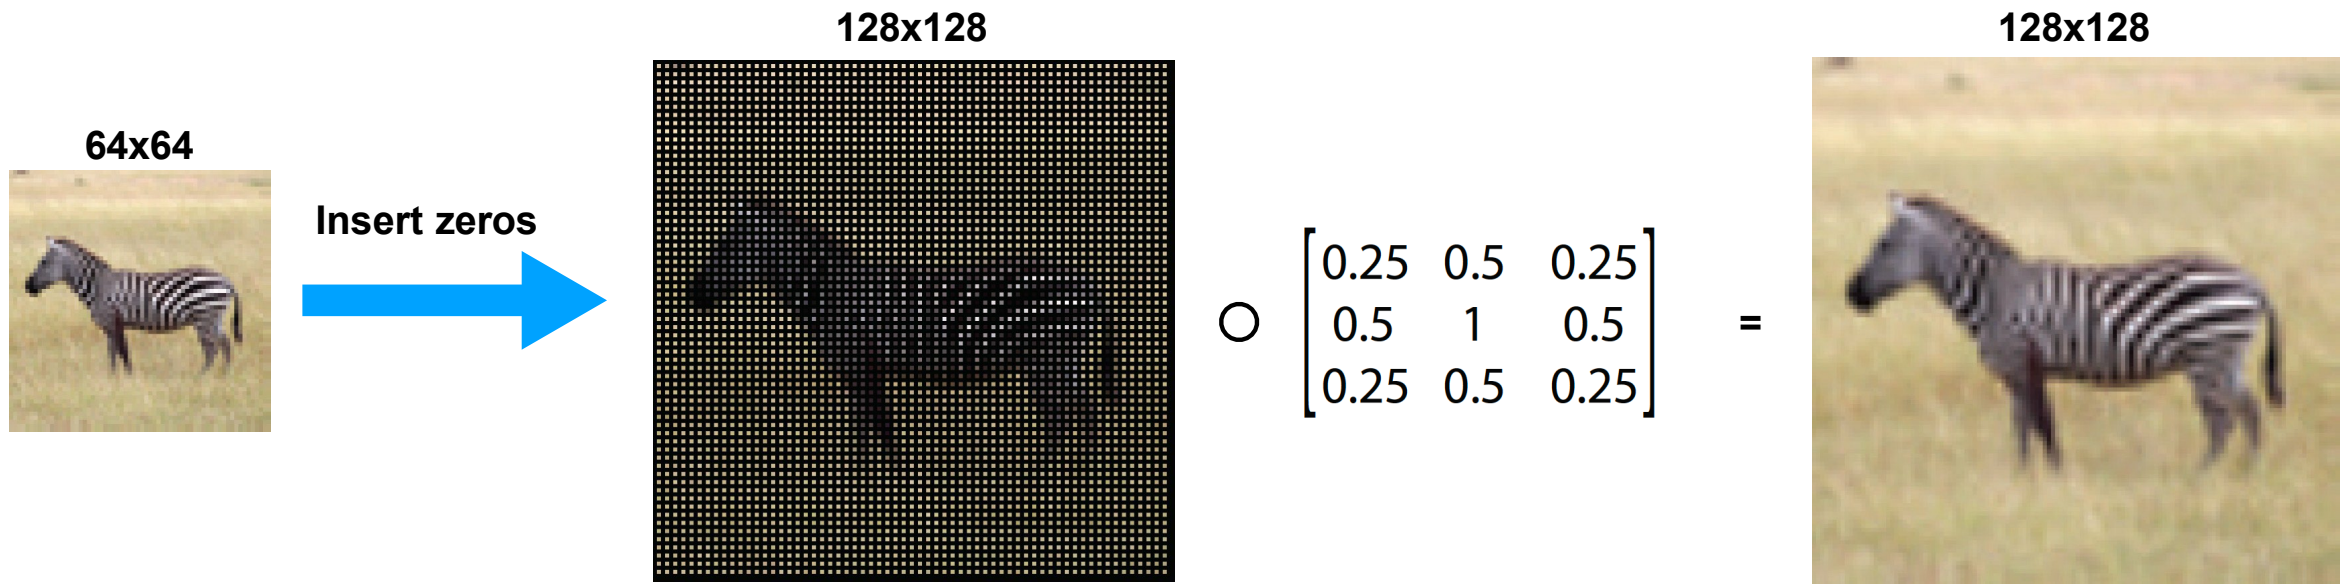
1. Blur input image with a Gaussian filter
2. Downsample image

# The Laplacian pyramid


Compute the difference between **upsampled** Gaussian pyramid level  $k+1$  and Gaussian pyramid level  $k$ . Recall that this approximates the blurred Laplacian.



# Upsampling

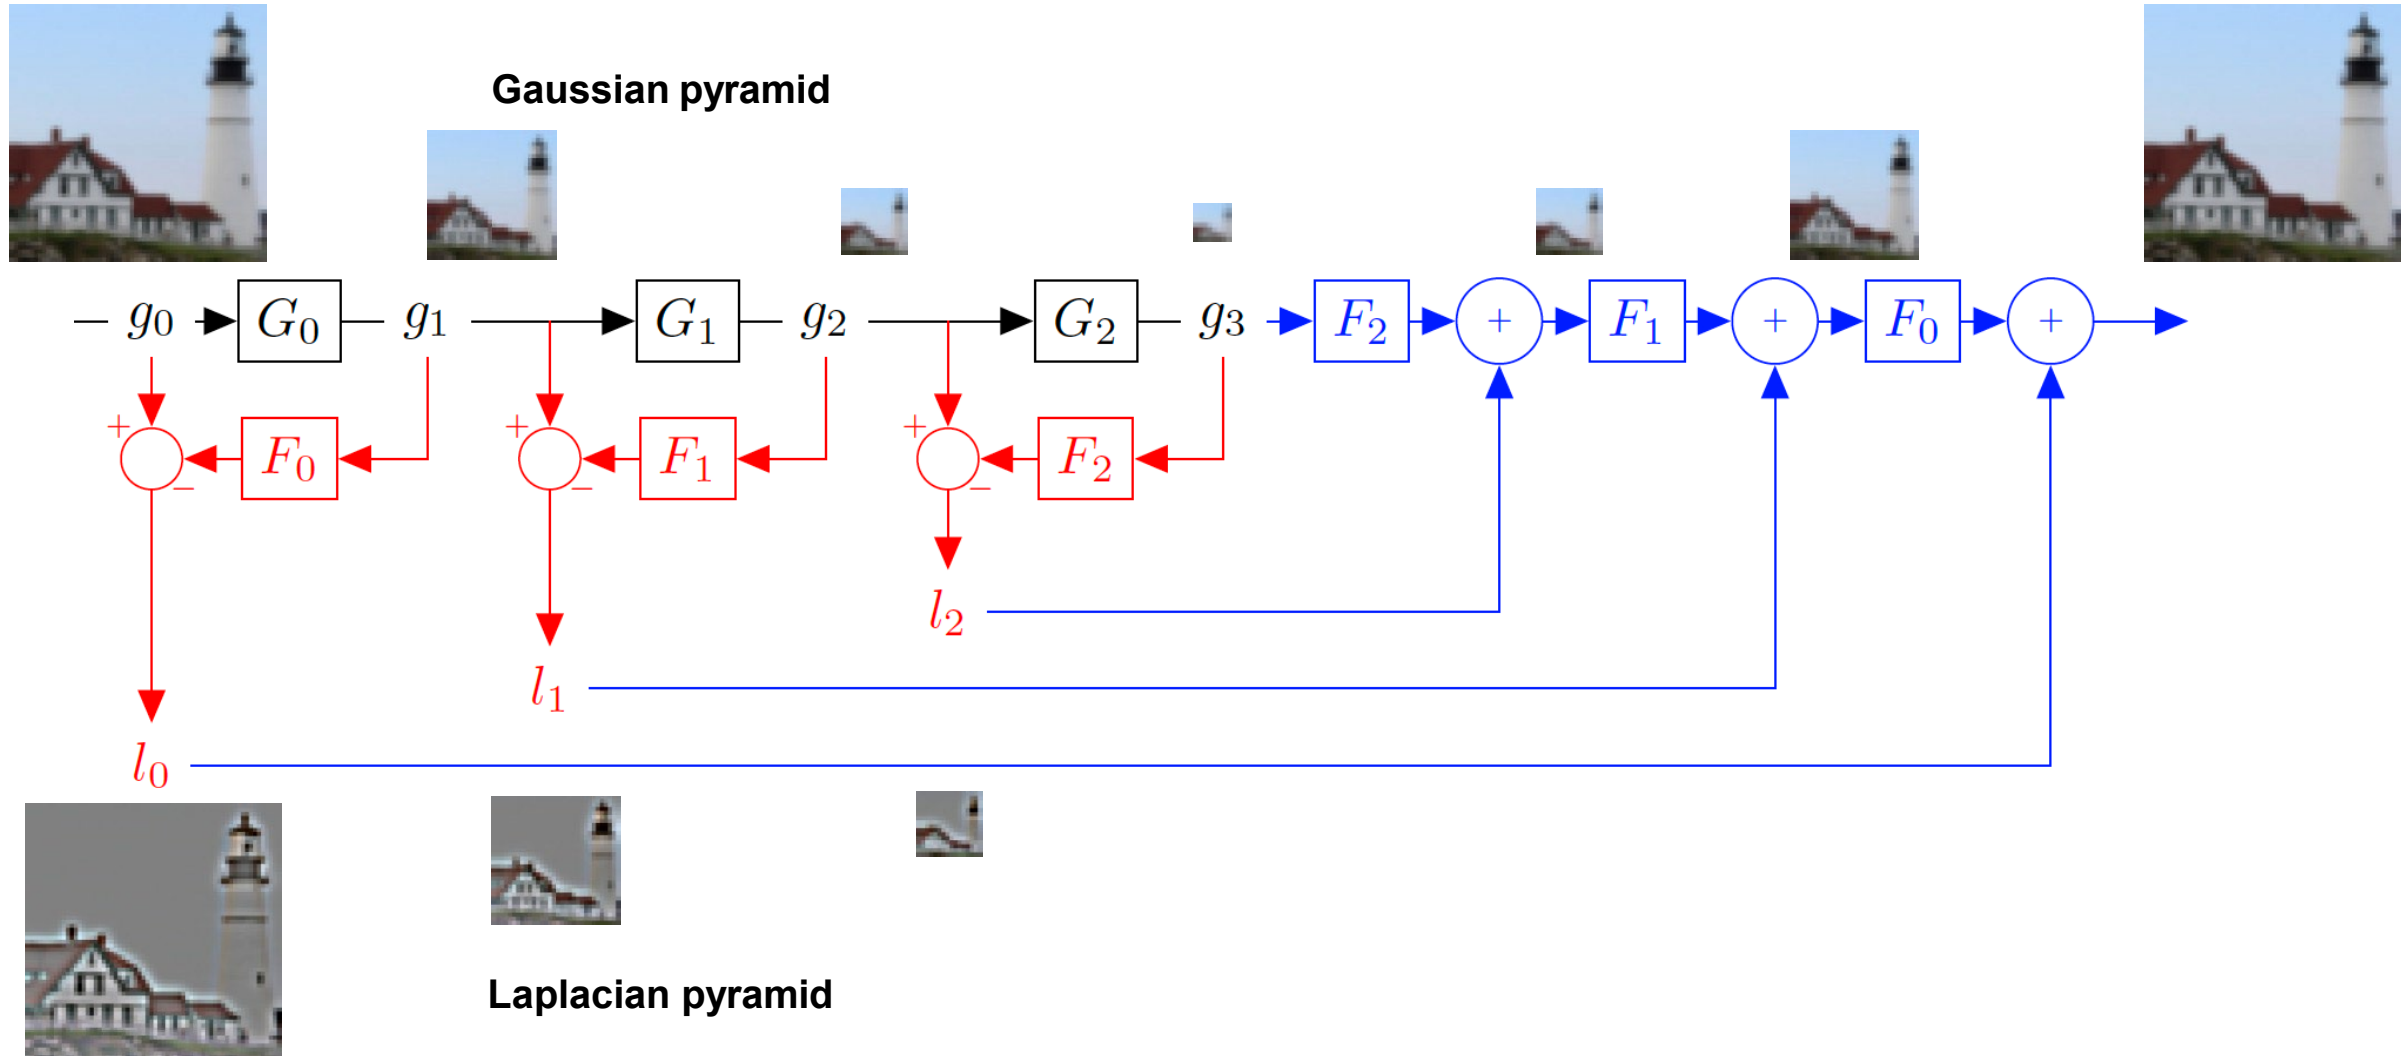


# The Laplacian pyramid



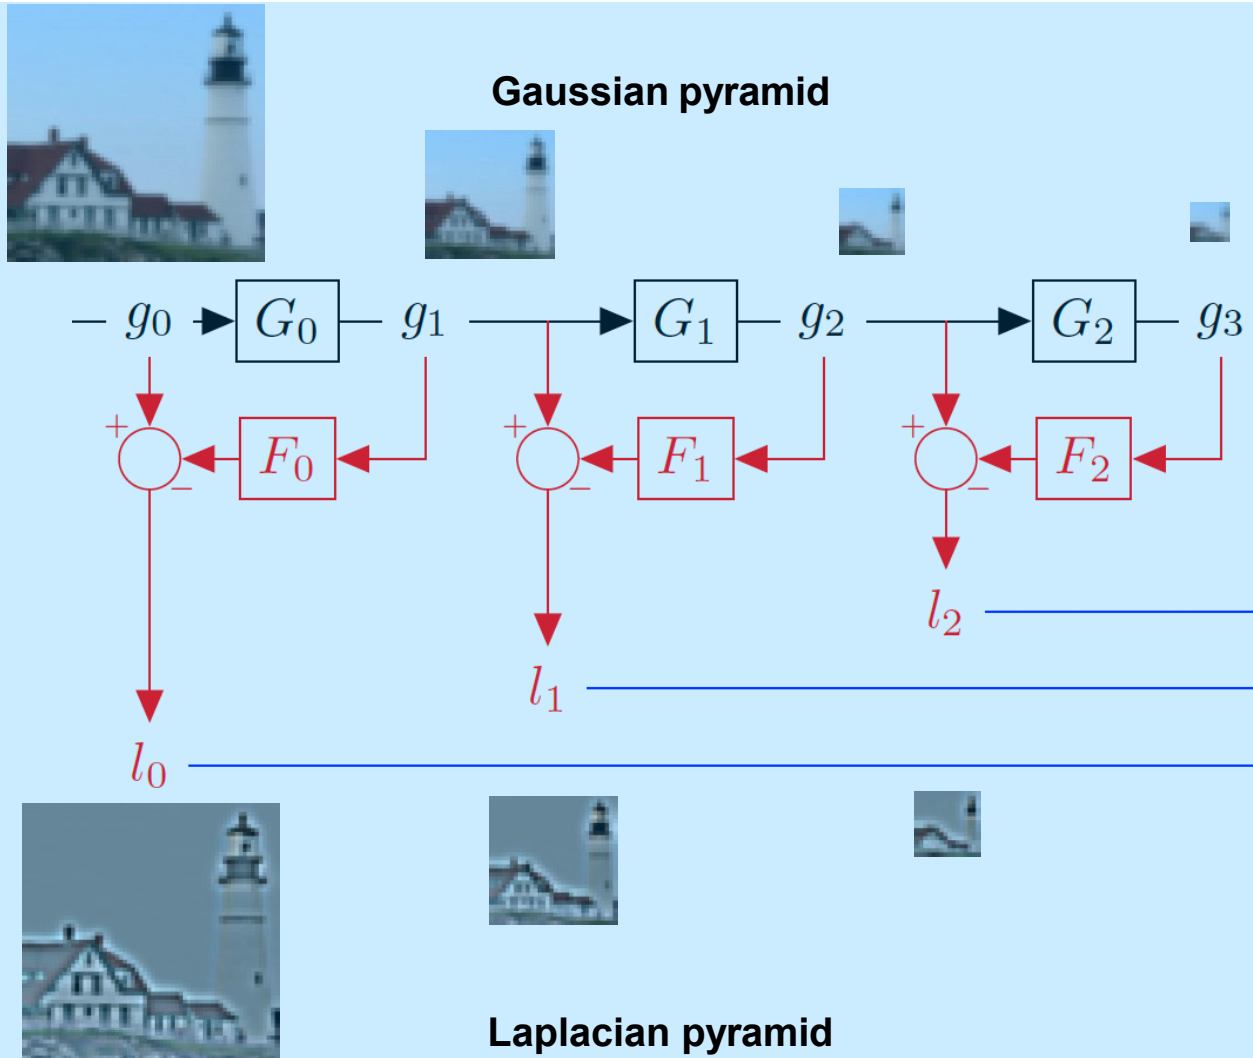


# Inverting the Laplacian Pyramid



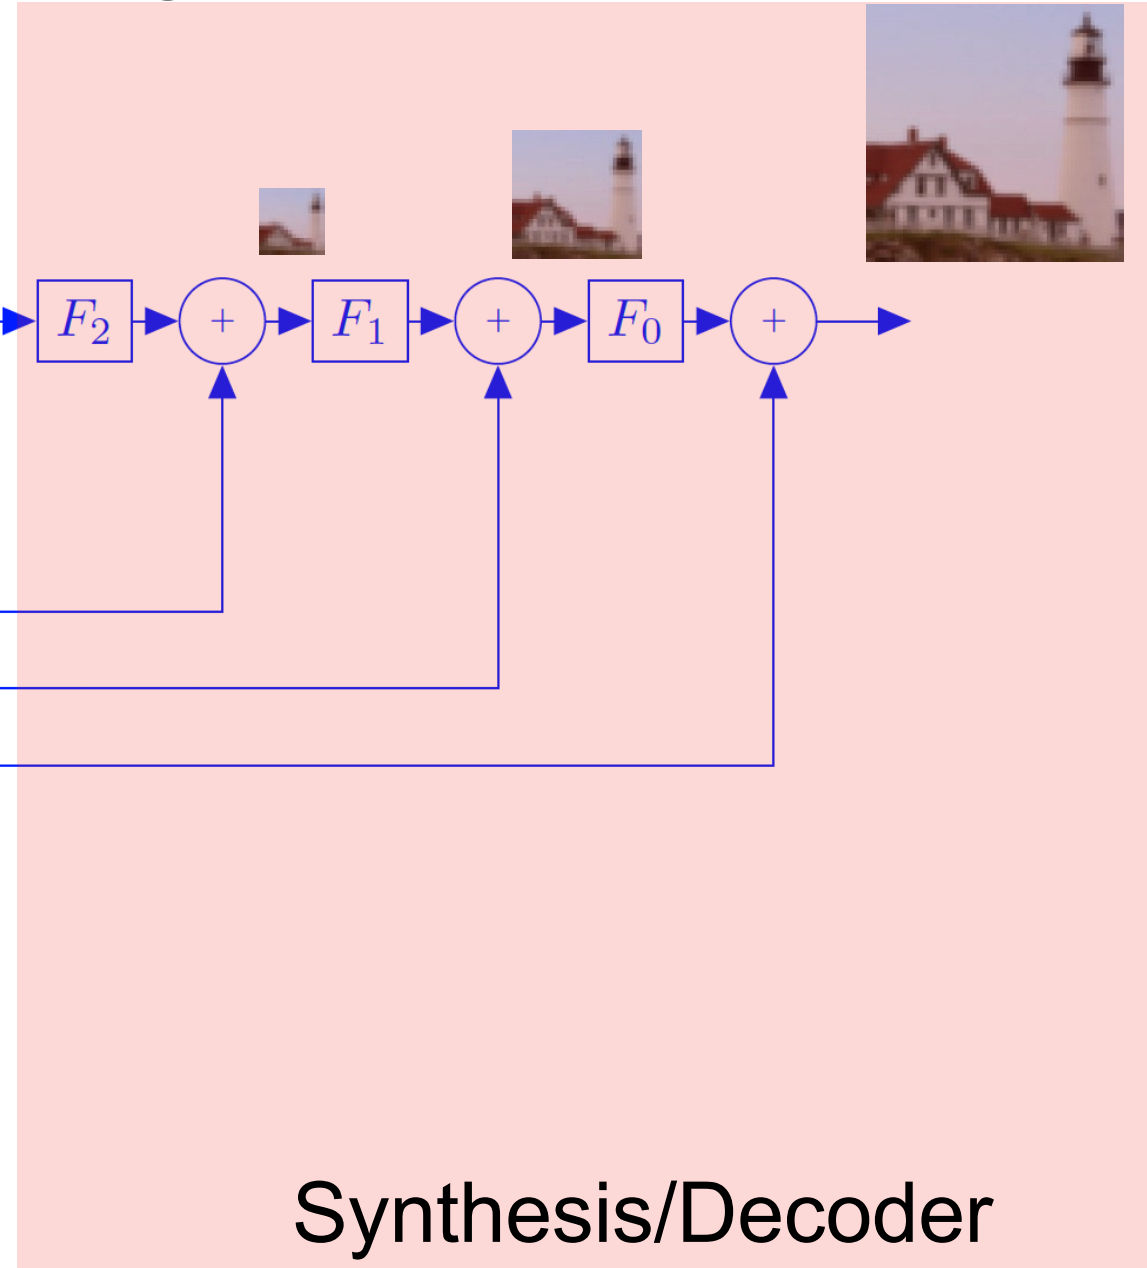
# The Laplacian pyramid

Gaussian pyramid



Laplacian pyramid

Analysis/Encoder



Synthesis/Decoder

# Applications of Laplacian Pyramid

- Image Blending
- Image Compression
- Noise Removal
- **IMAGE FEATURES → IMAGE CLASSIFICATION ...**



# Application 1: Image Blending

# Image Blending



# Image Blending






# Simplest (but far from the best) Solution



- How would you do this?
- Give me an equation







# Simplest (but far from the best) Solution



$$I = m * I^A + (1 - m) * I^B$$

# Image Blending with the Laplacian Pyramid



$$l_k = l_k^A * m_k + l_k^B * (1 - m_k)$$




# Simple Masked Summation vs. Laplacian Pyramid







+



=



Simple blend




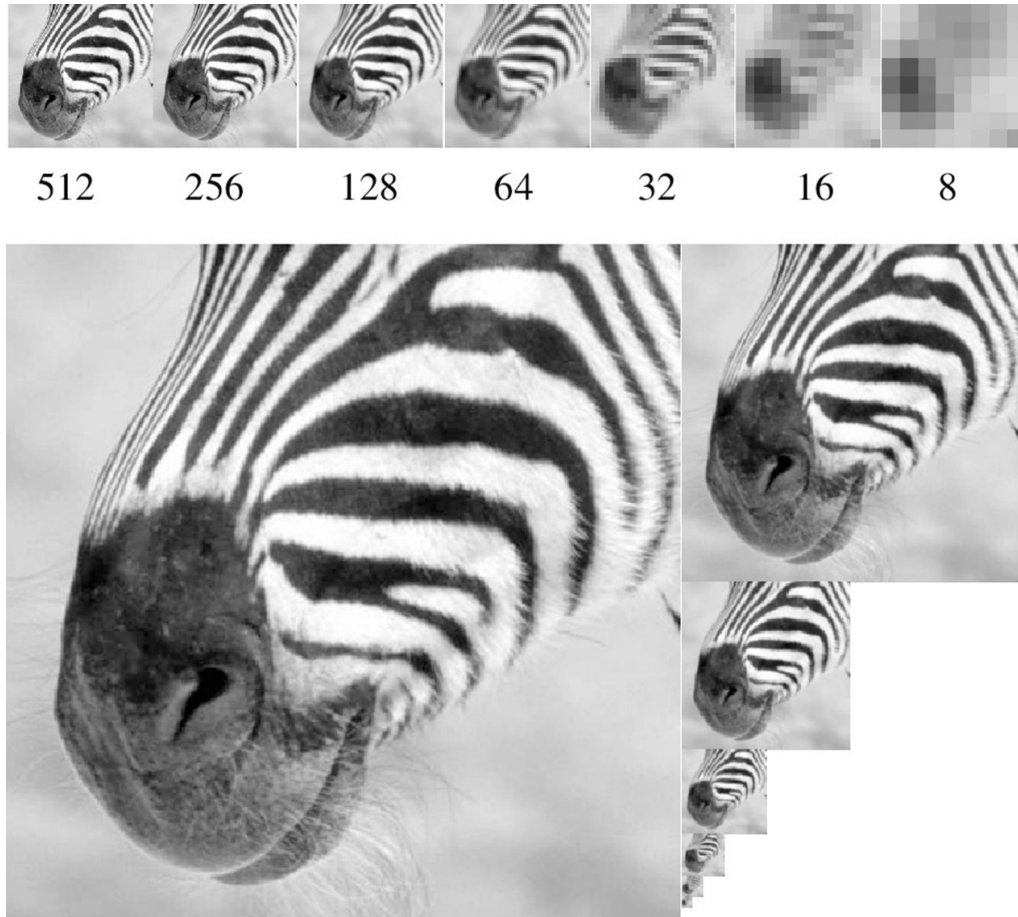
With Laplacian pyr.





Photo credit: Chris Cameron

# Image Pyramids



And many more: steerable filters, wavelets, ...  
convolutional networks!



