

# Lecture 4: Image pyramids

- PS1 due at midnight
- PS2 out, due next Tues.
- No Thursday office hours this week
- If you're on the waitlist, submit your PS1 via email to the course staff.

# Today

- Image pyramids
- Texture

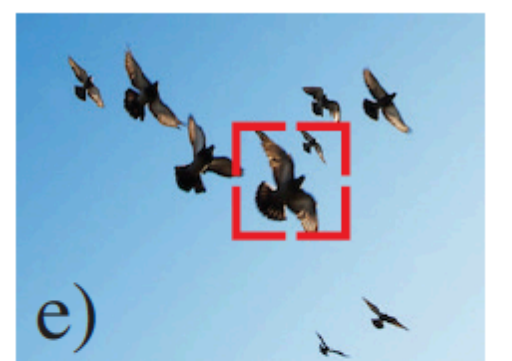



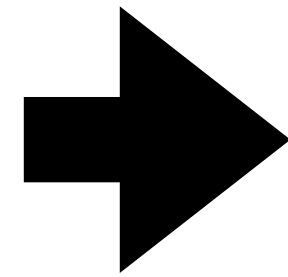




We want scale and translation invariance.




# Image pyramids



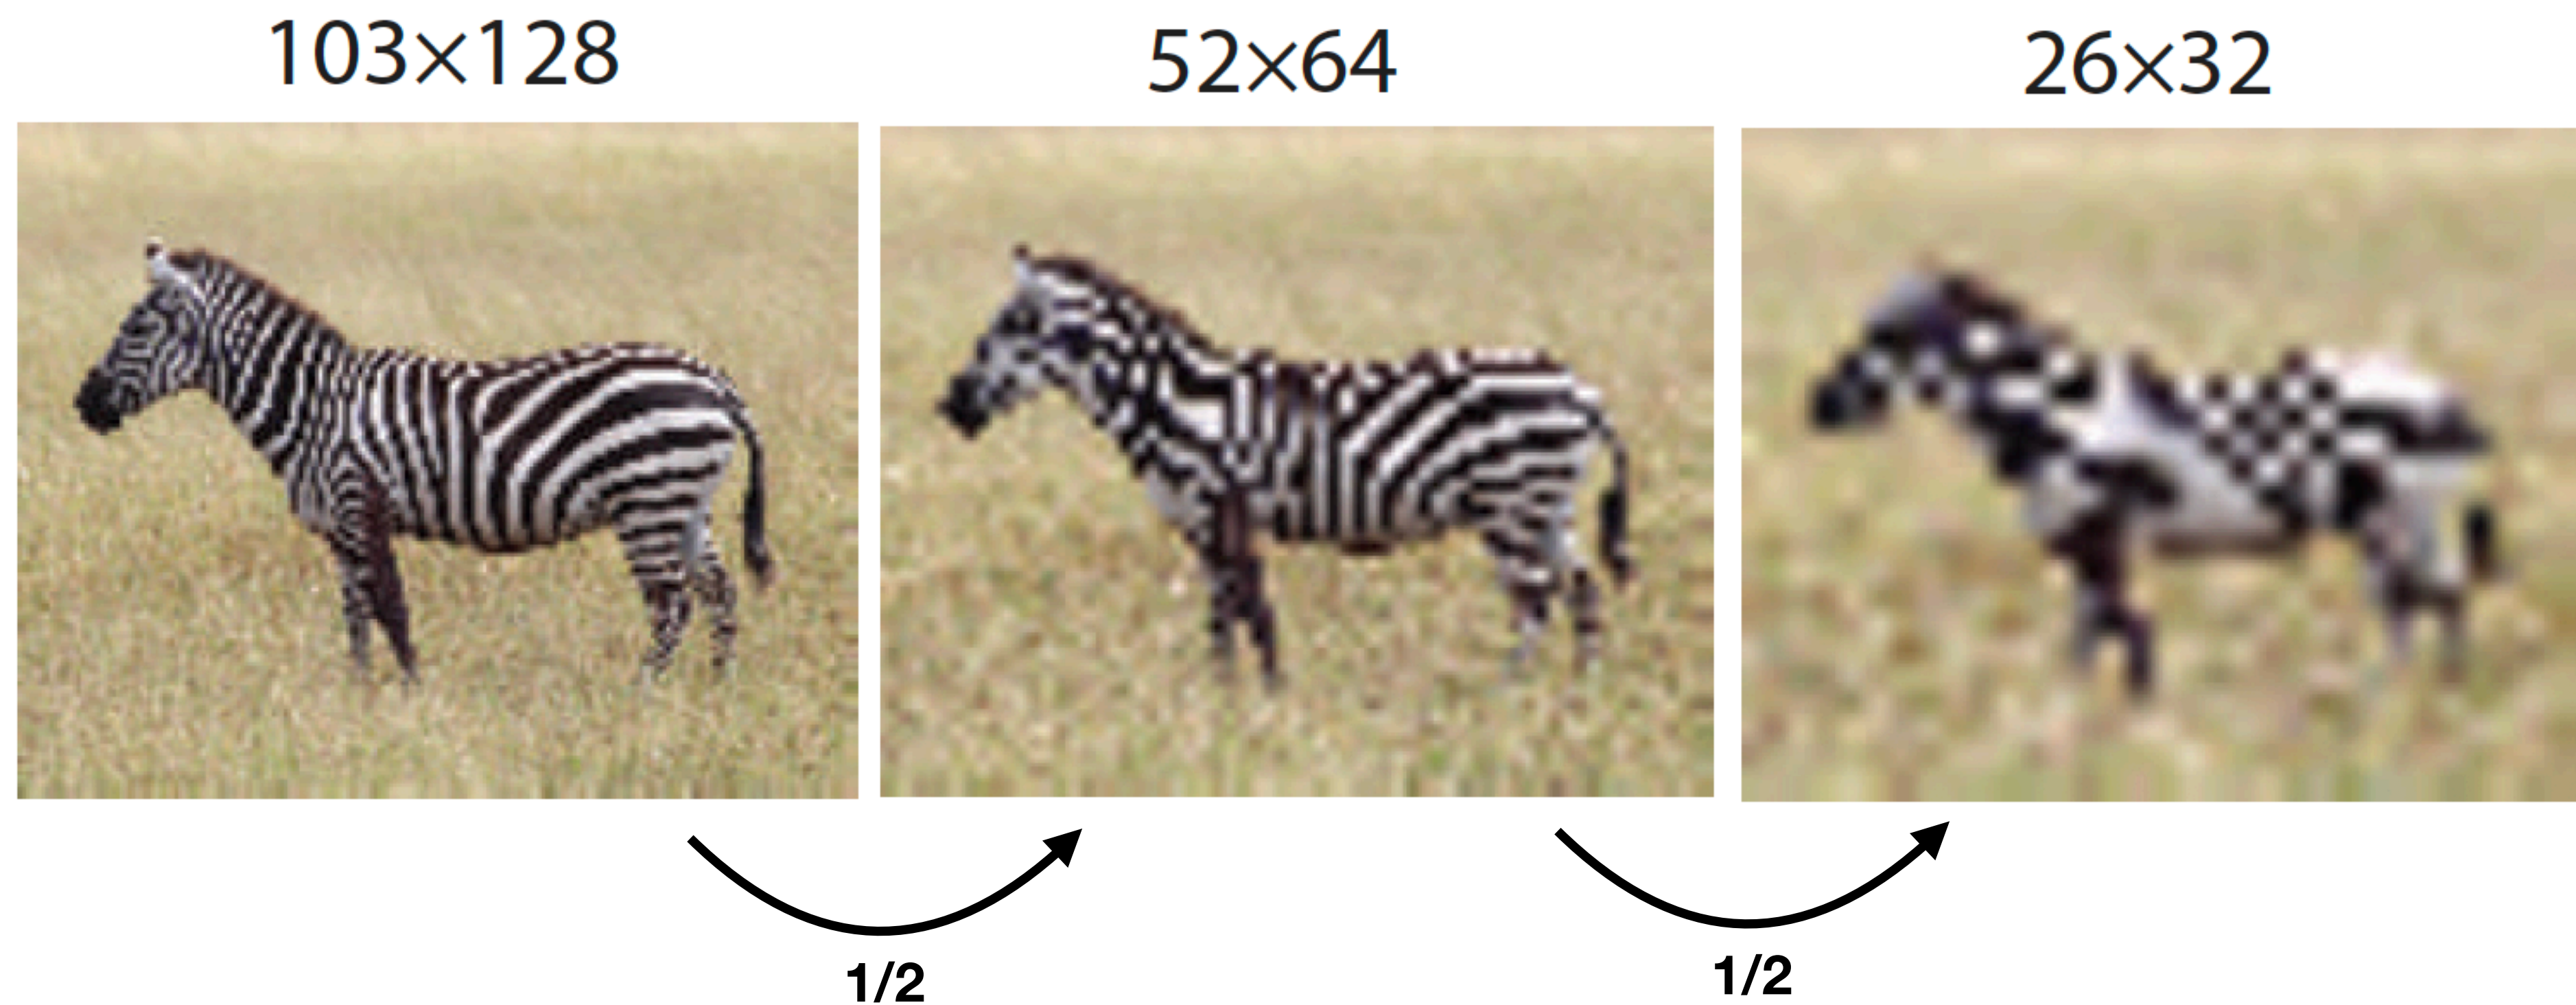


# Gaussian Pyramid






# Subsampling and aliasing



# Aliasing



Both waves fit the same samples. We “perceive” the red wave when the actual input was the blue wave.



# The Gaussian pyramid

For each level

1. Blur input image with a Gaussian filter

$[1, 4, 6, 4, 1]$




$$\bigcirc [1, 4, 6, 4, 1] \bigcirc \begin{matrix} [1 \\ 4 \\ 6 \\ 4 \\ 1] \end{matrix} \longrightarrow$$



# The Gaussian pyramid

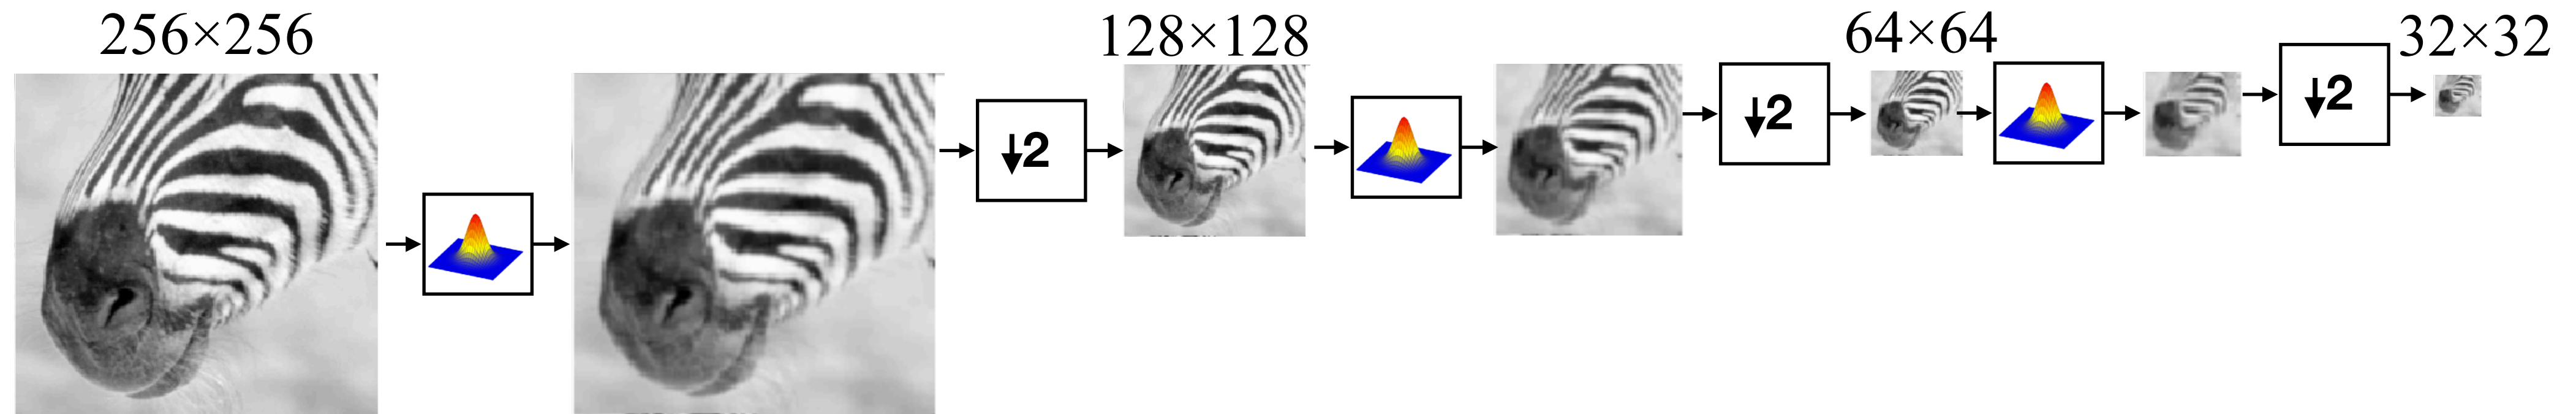
For each level

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

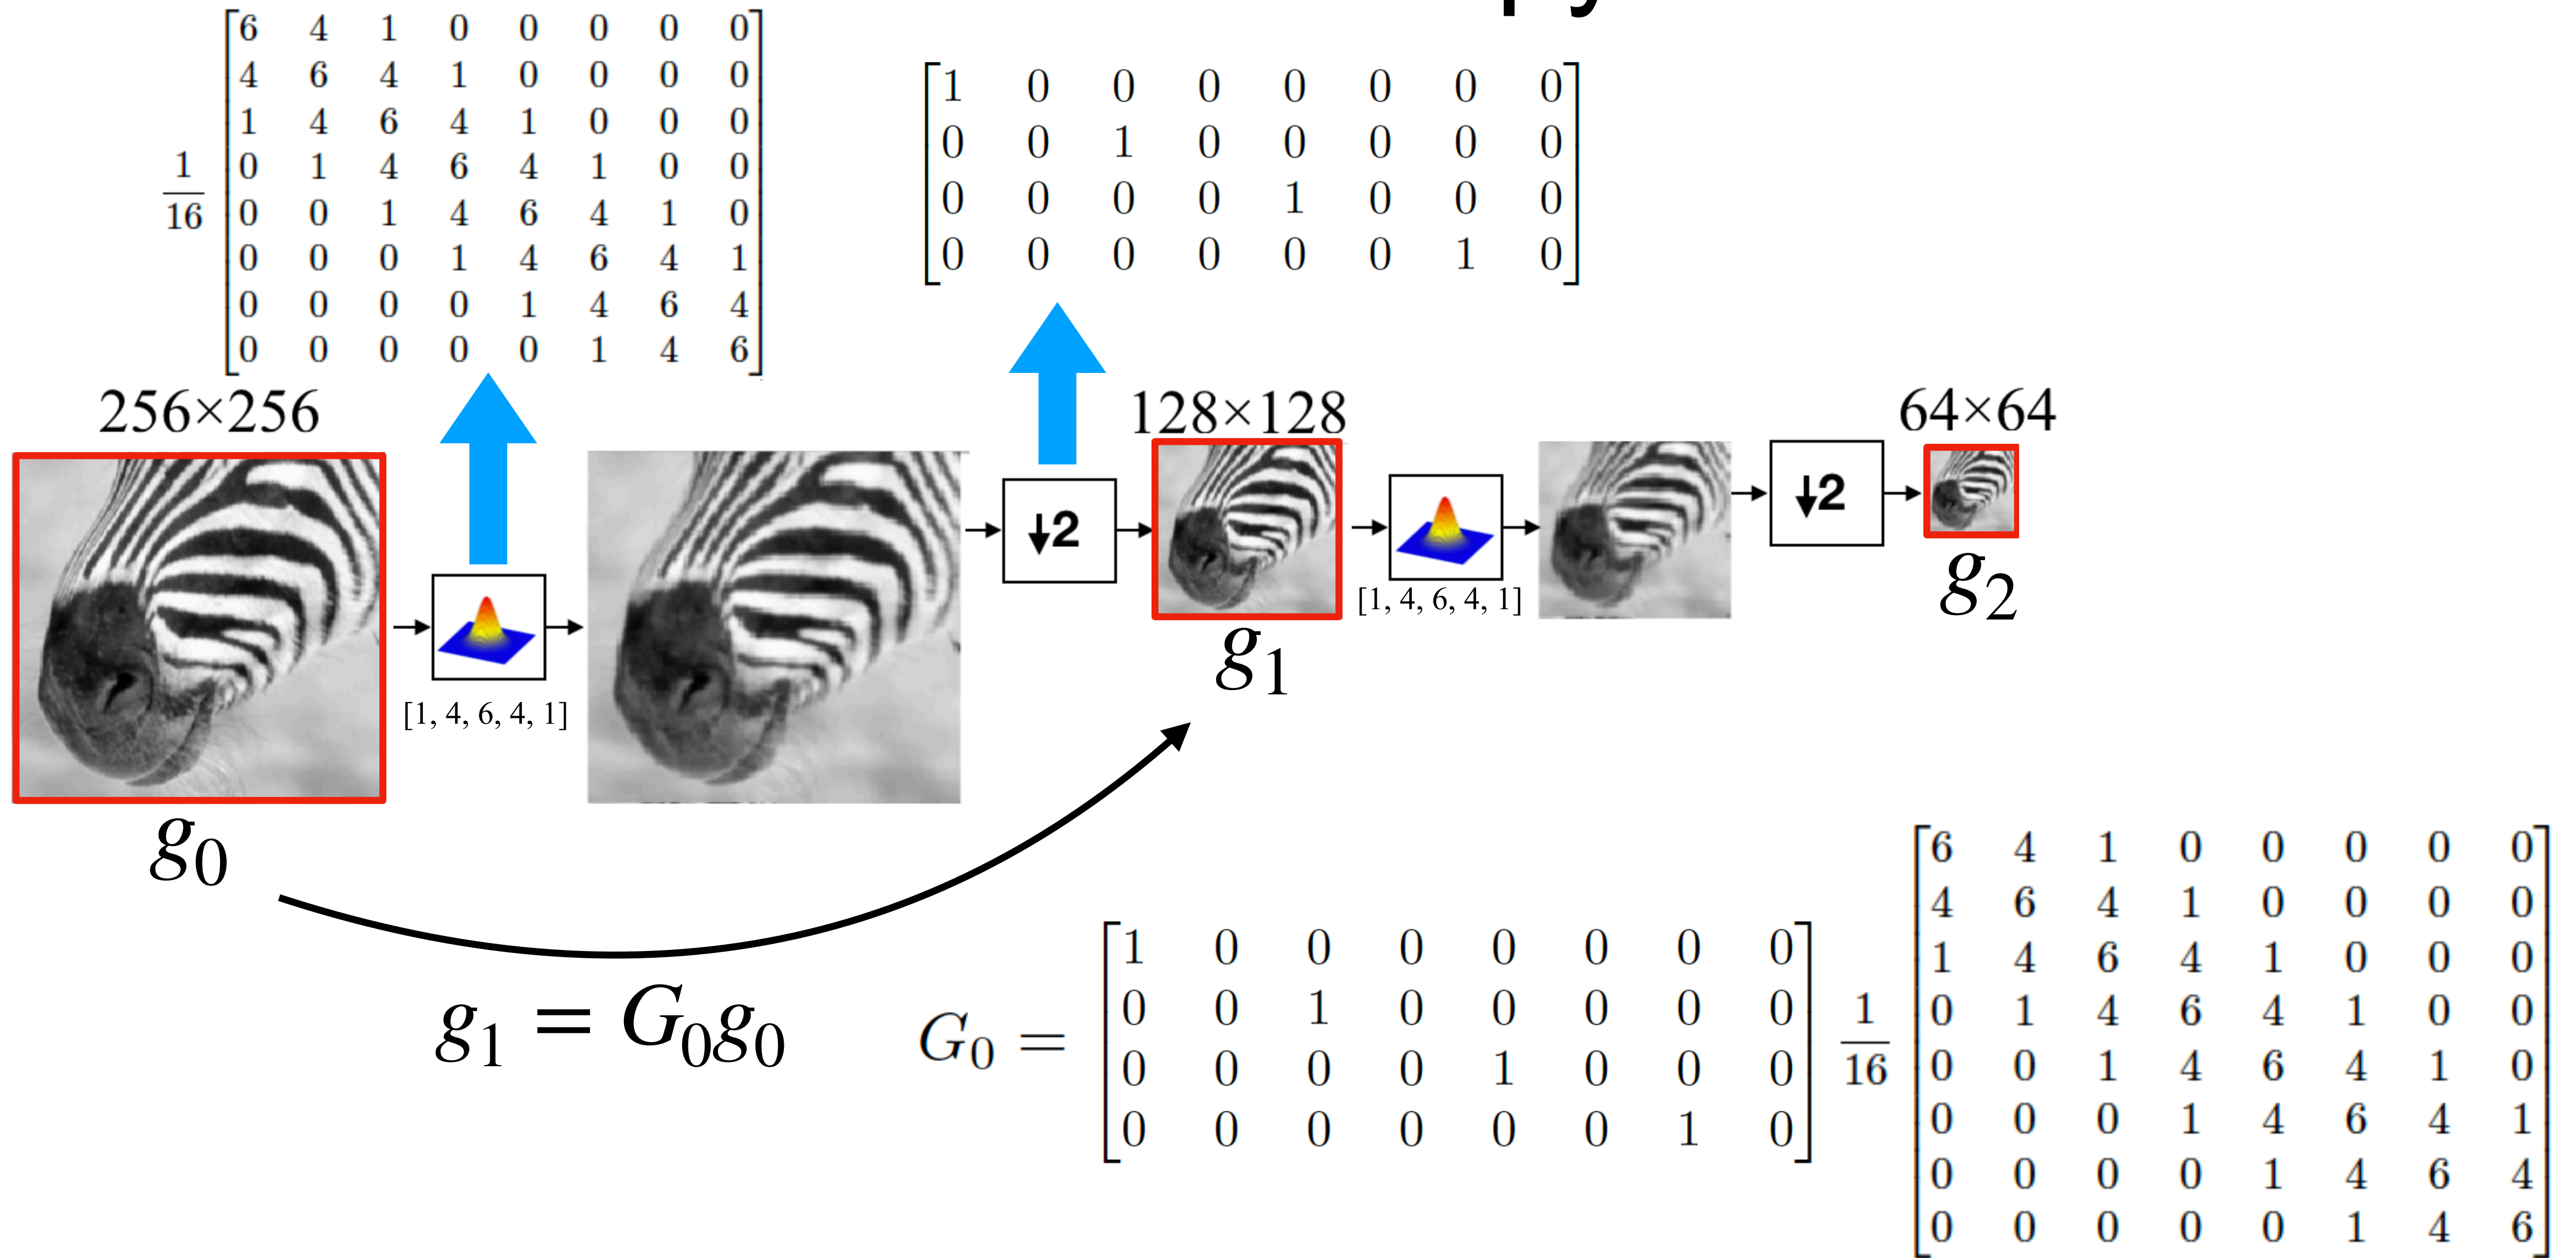


32×32

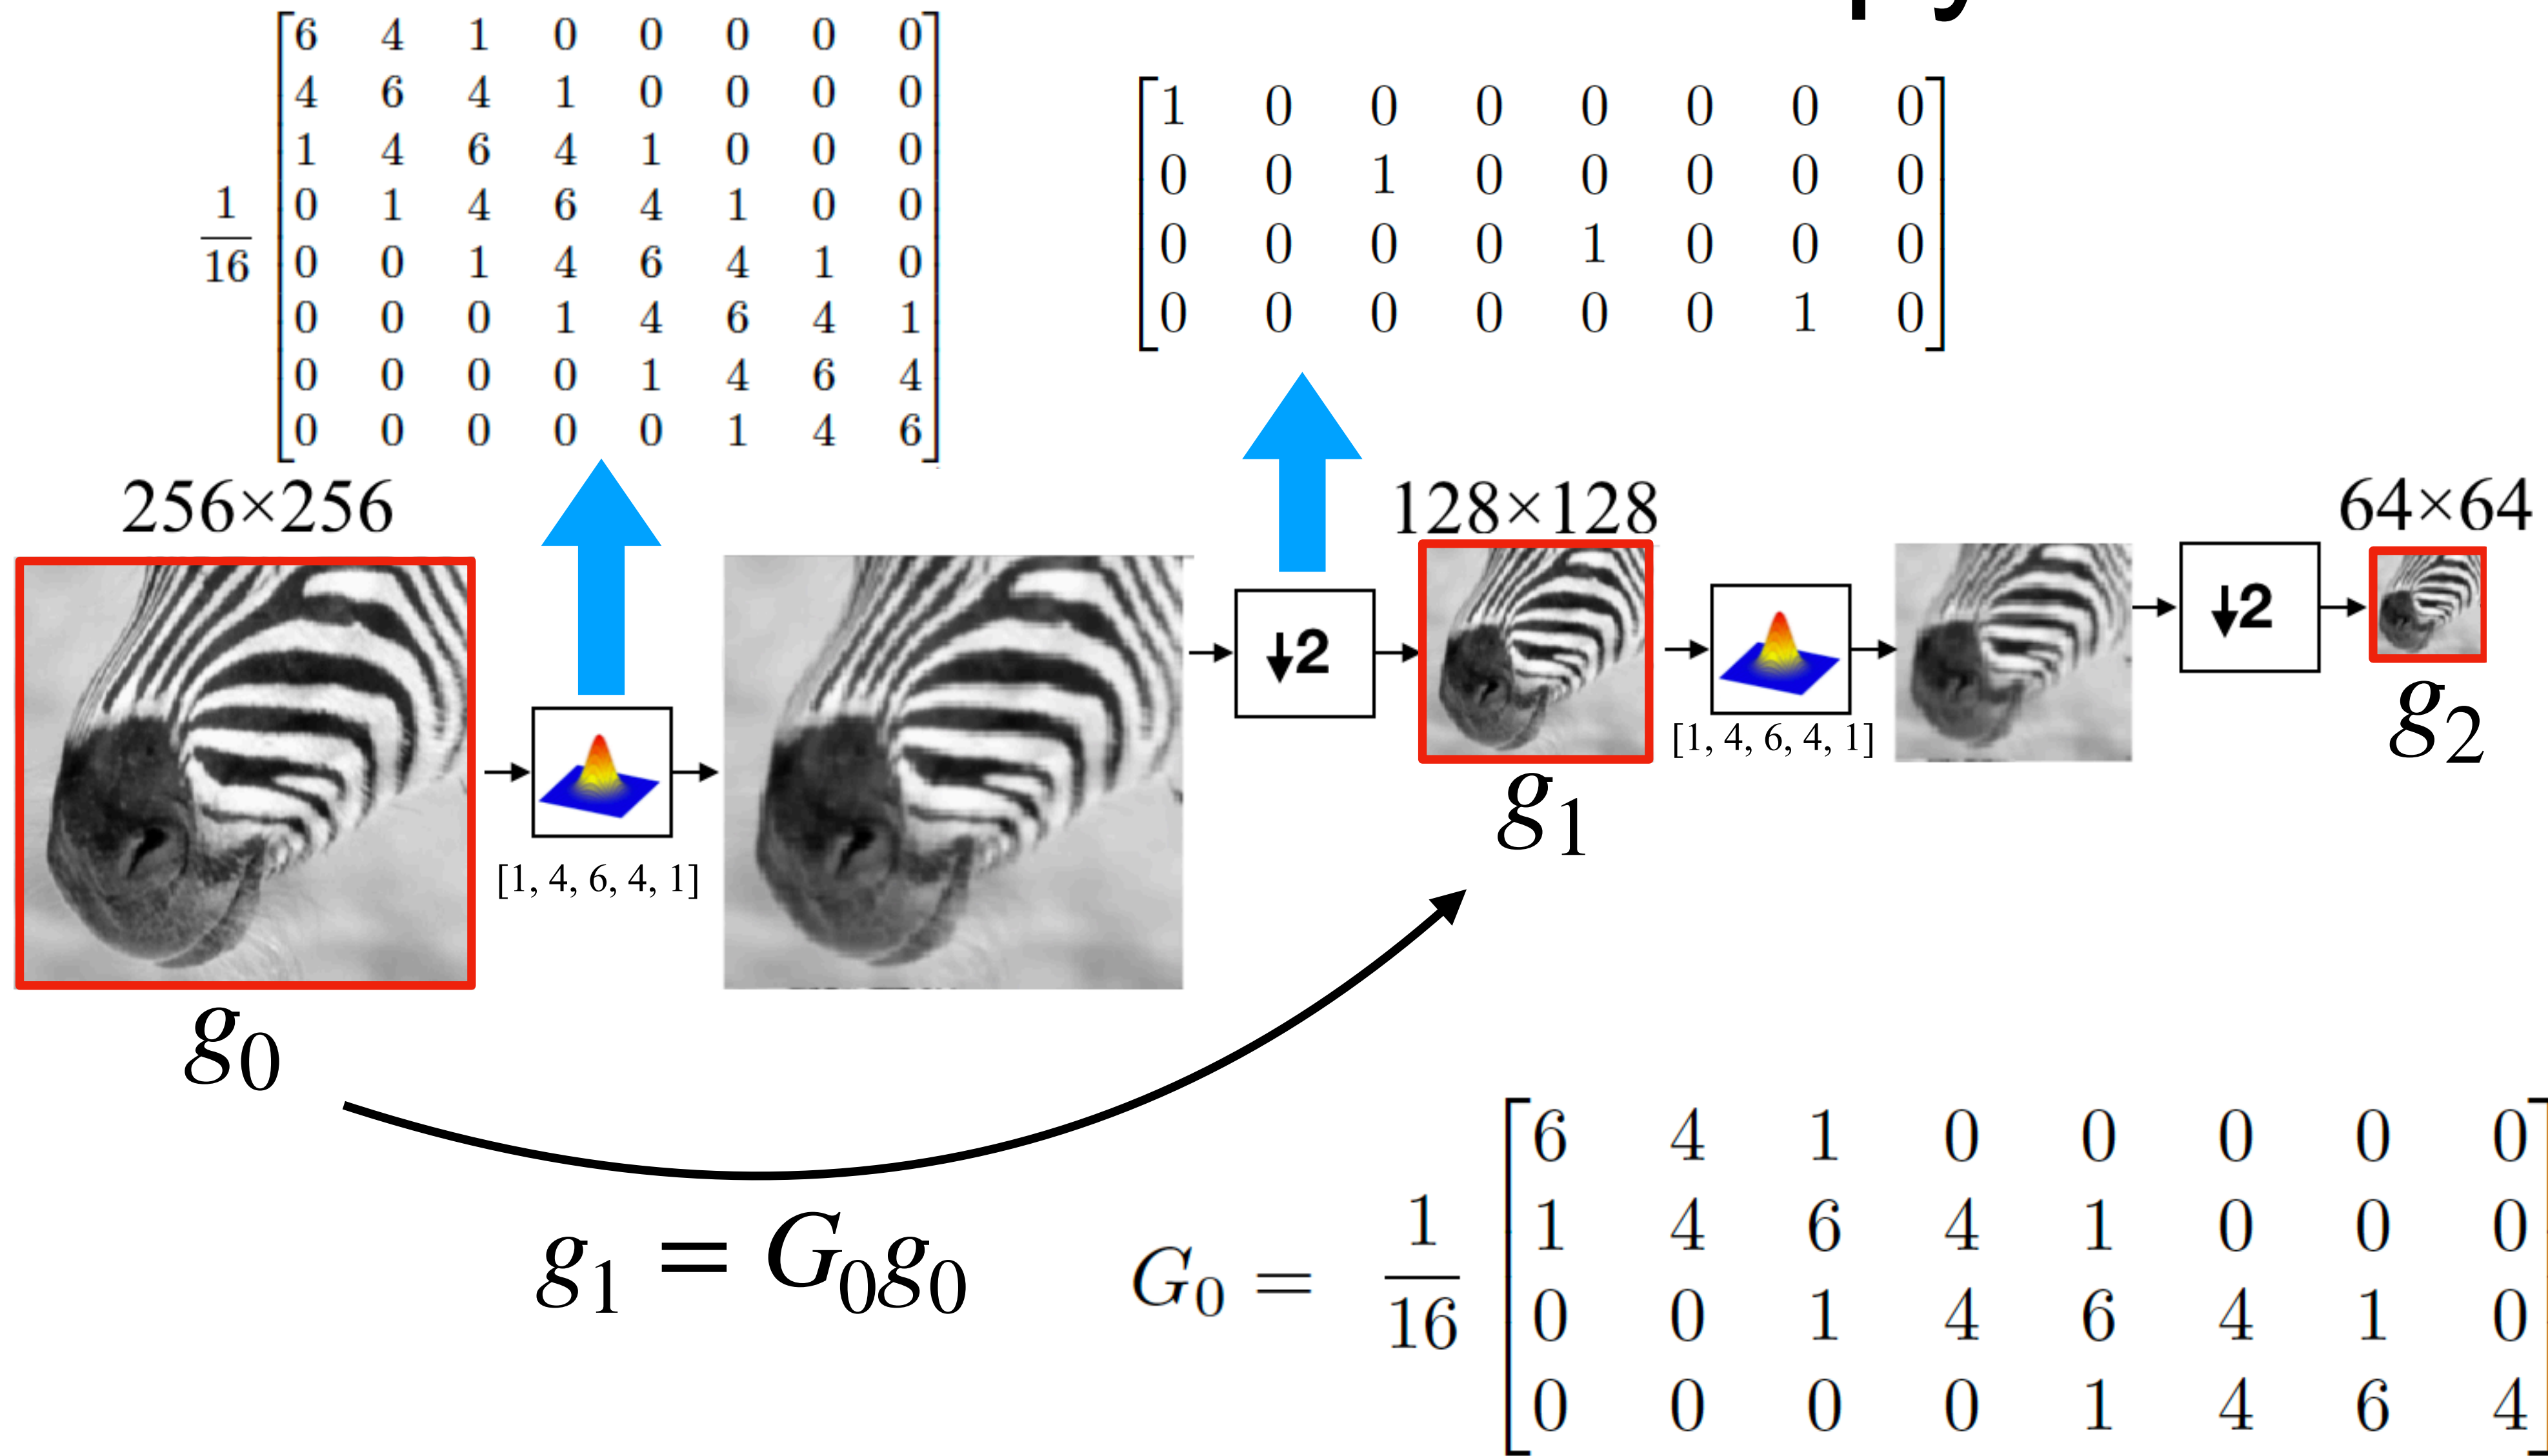




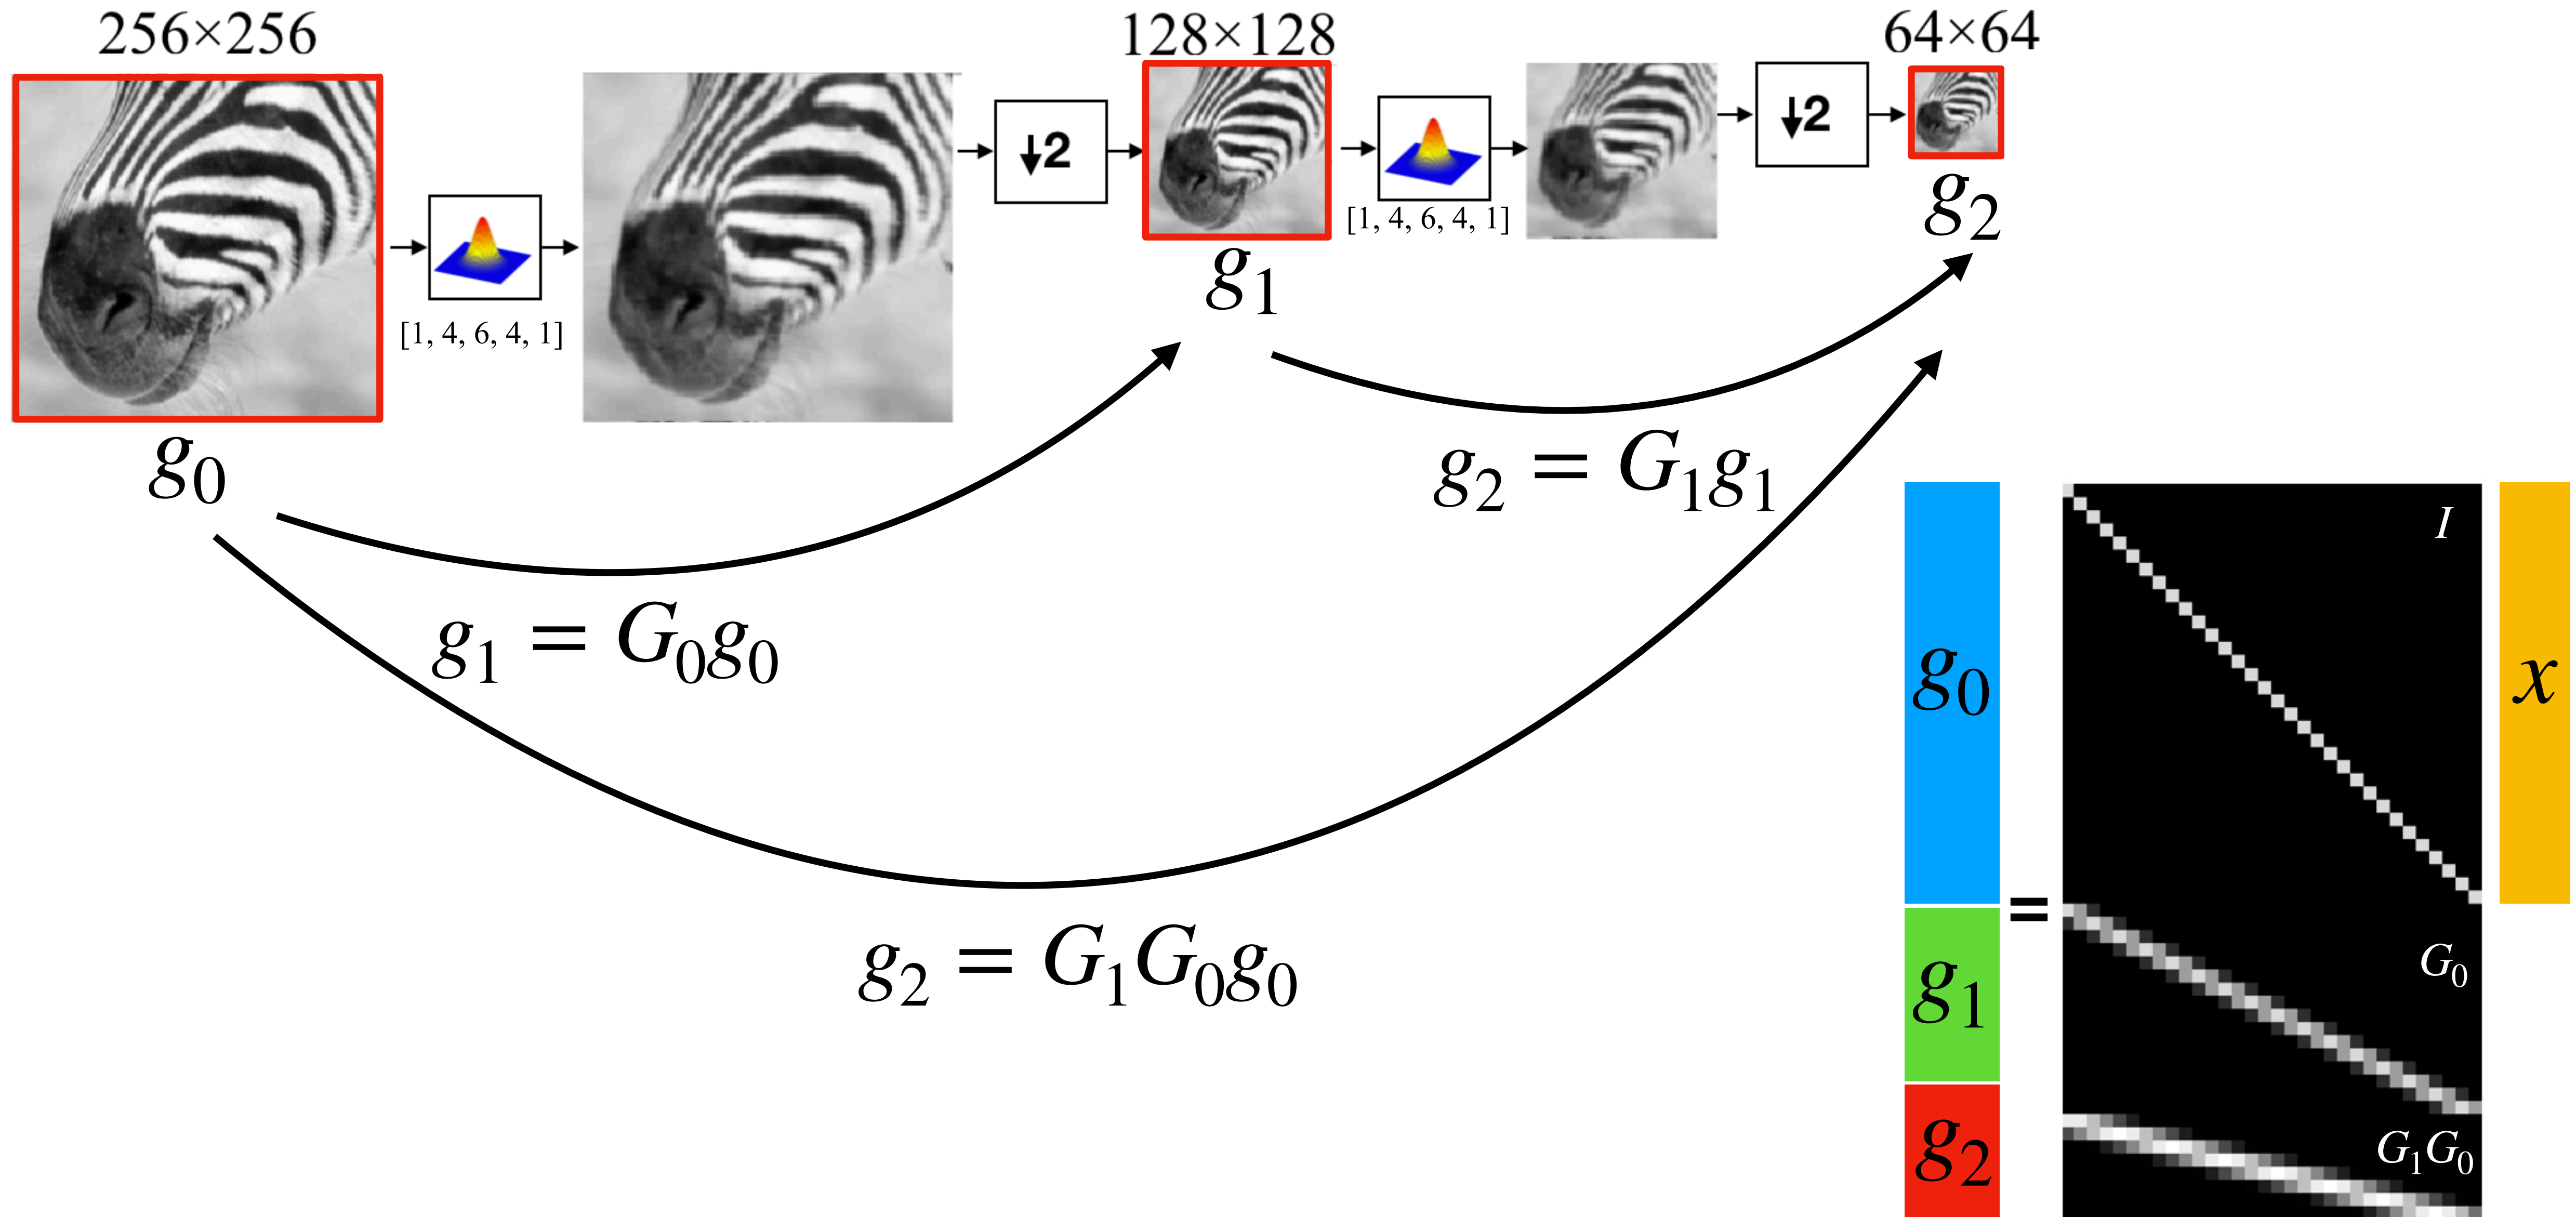
# The Gaussian pyramid



# The Gaussian pyramid




# The Gaussian pyramid





# The Gaussian pyramid

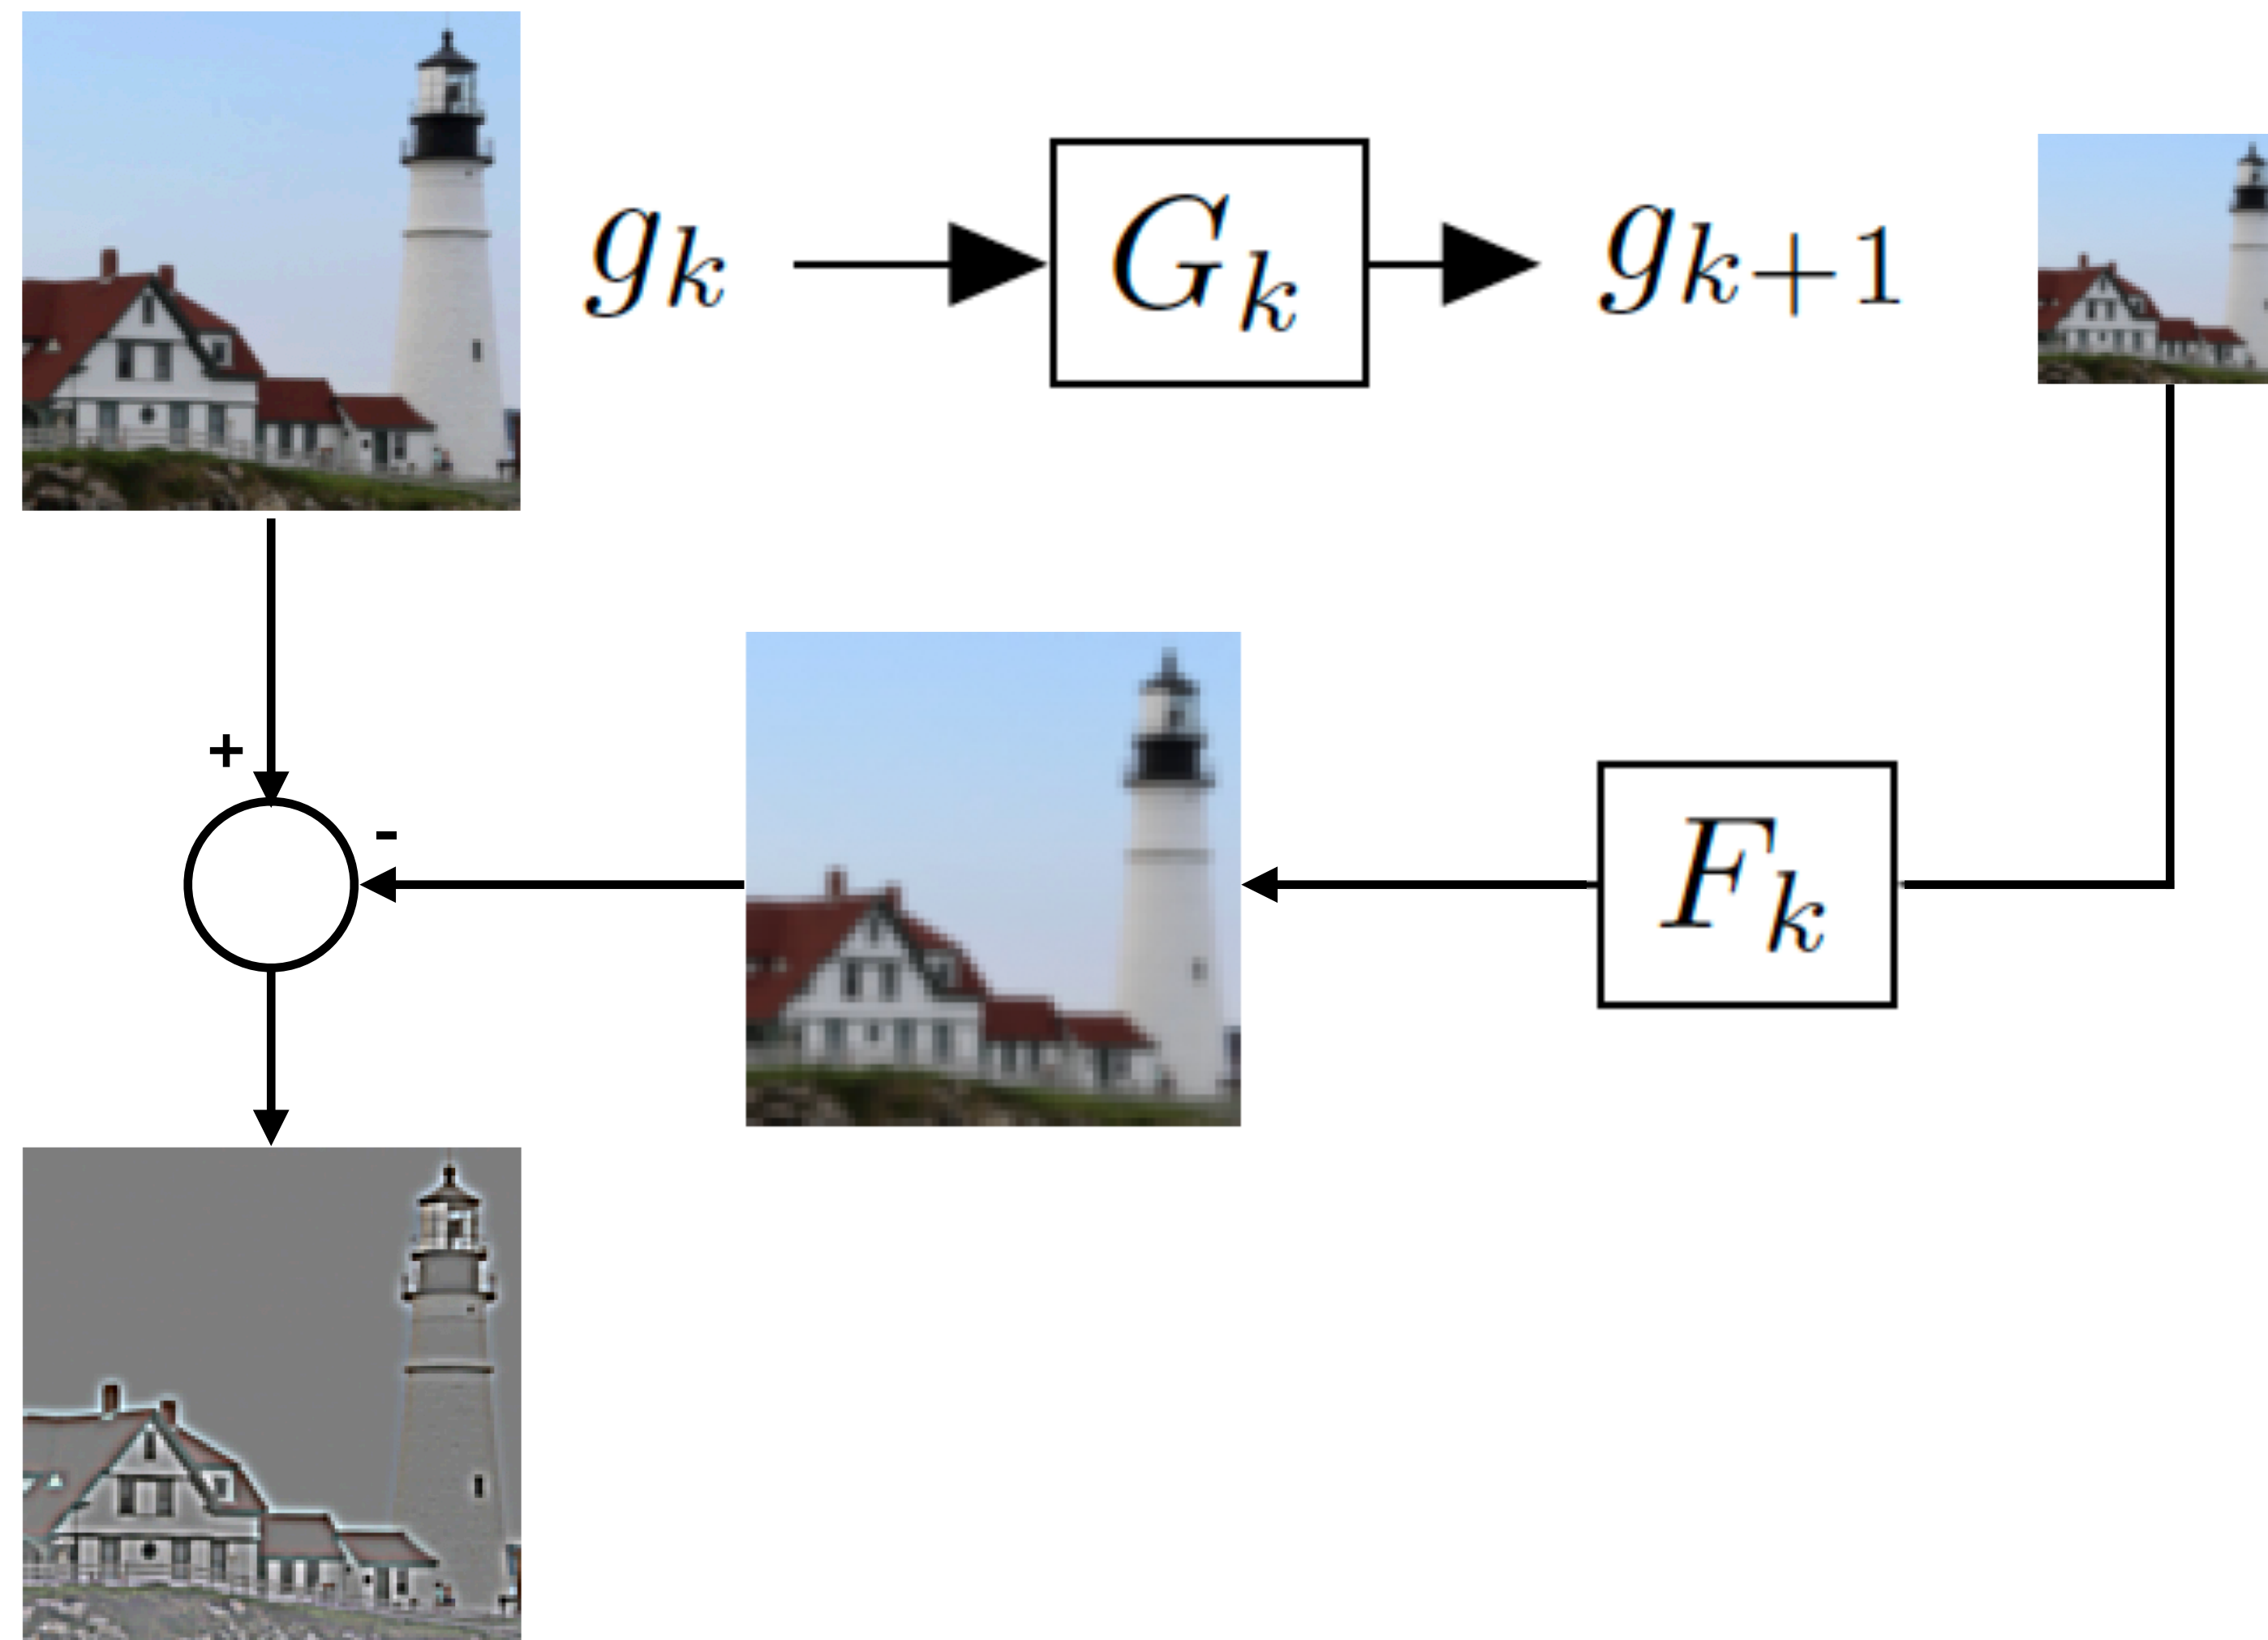


For each level

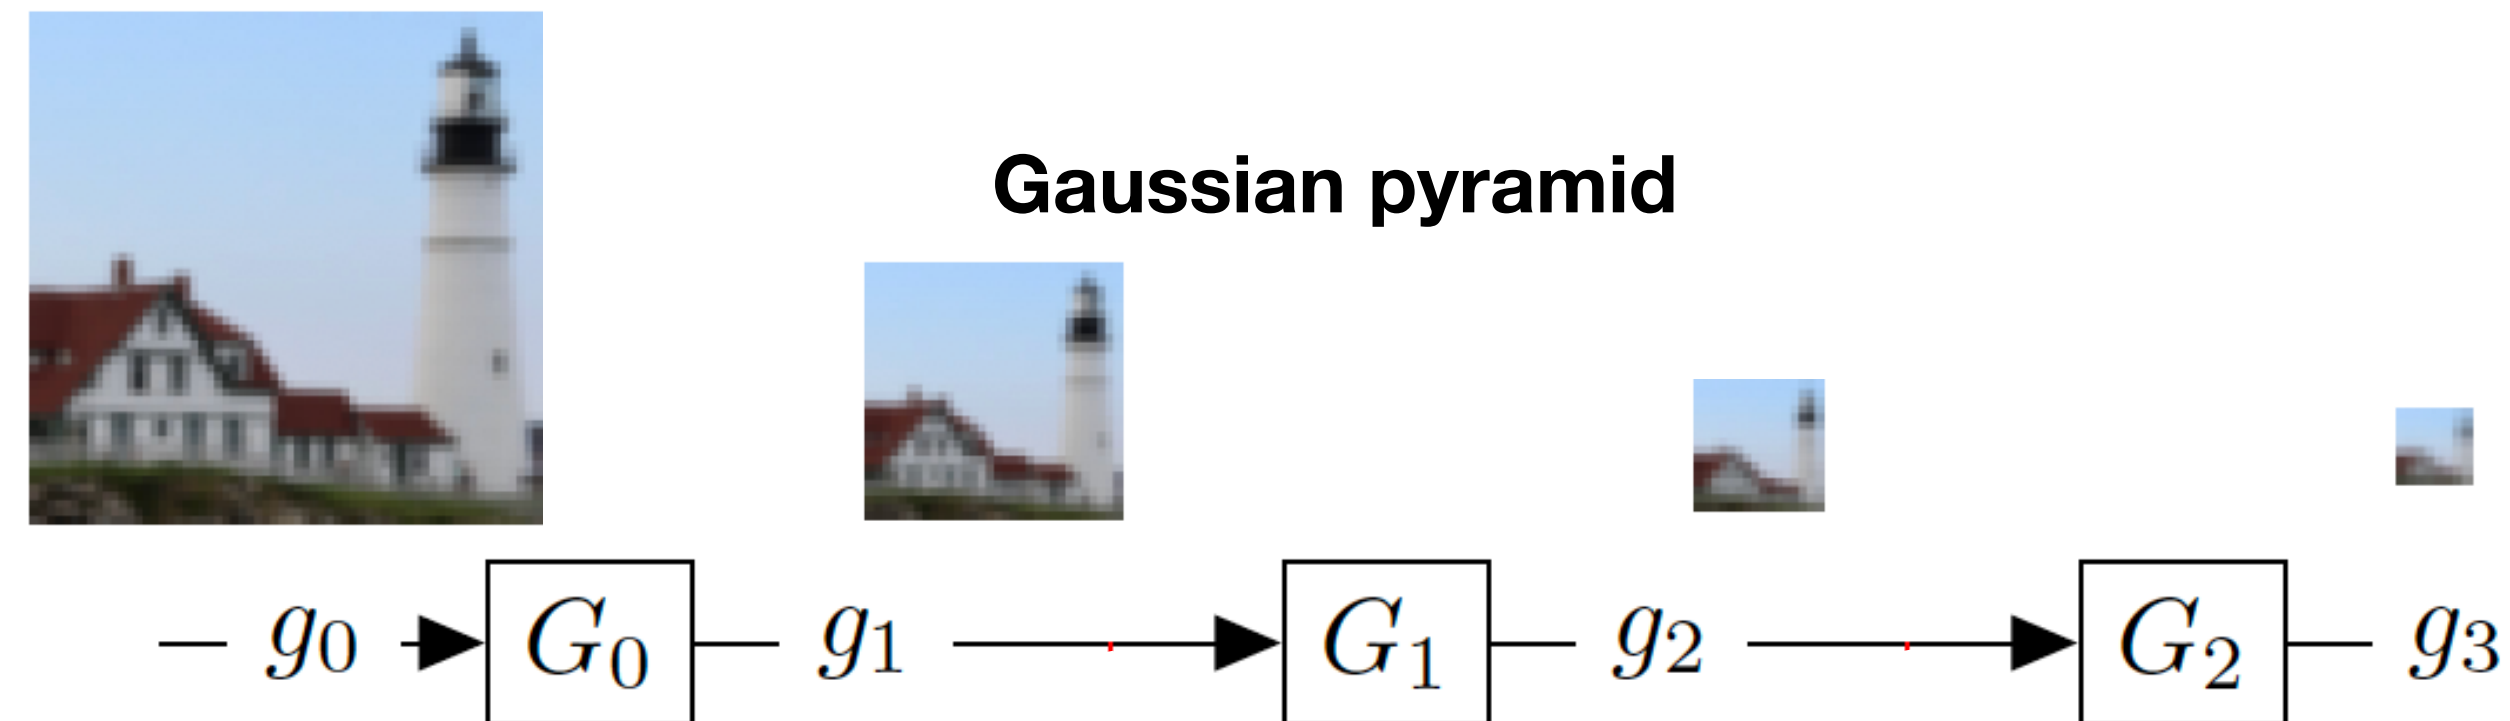
1. Blur input image with a Gaussian filter
2. Downsample image

# 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.

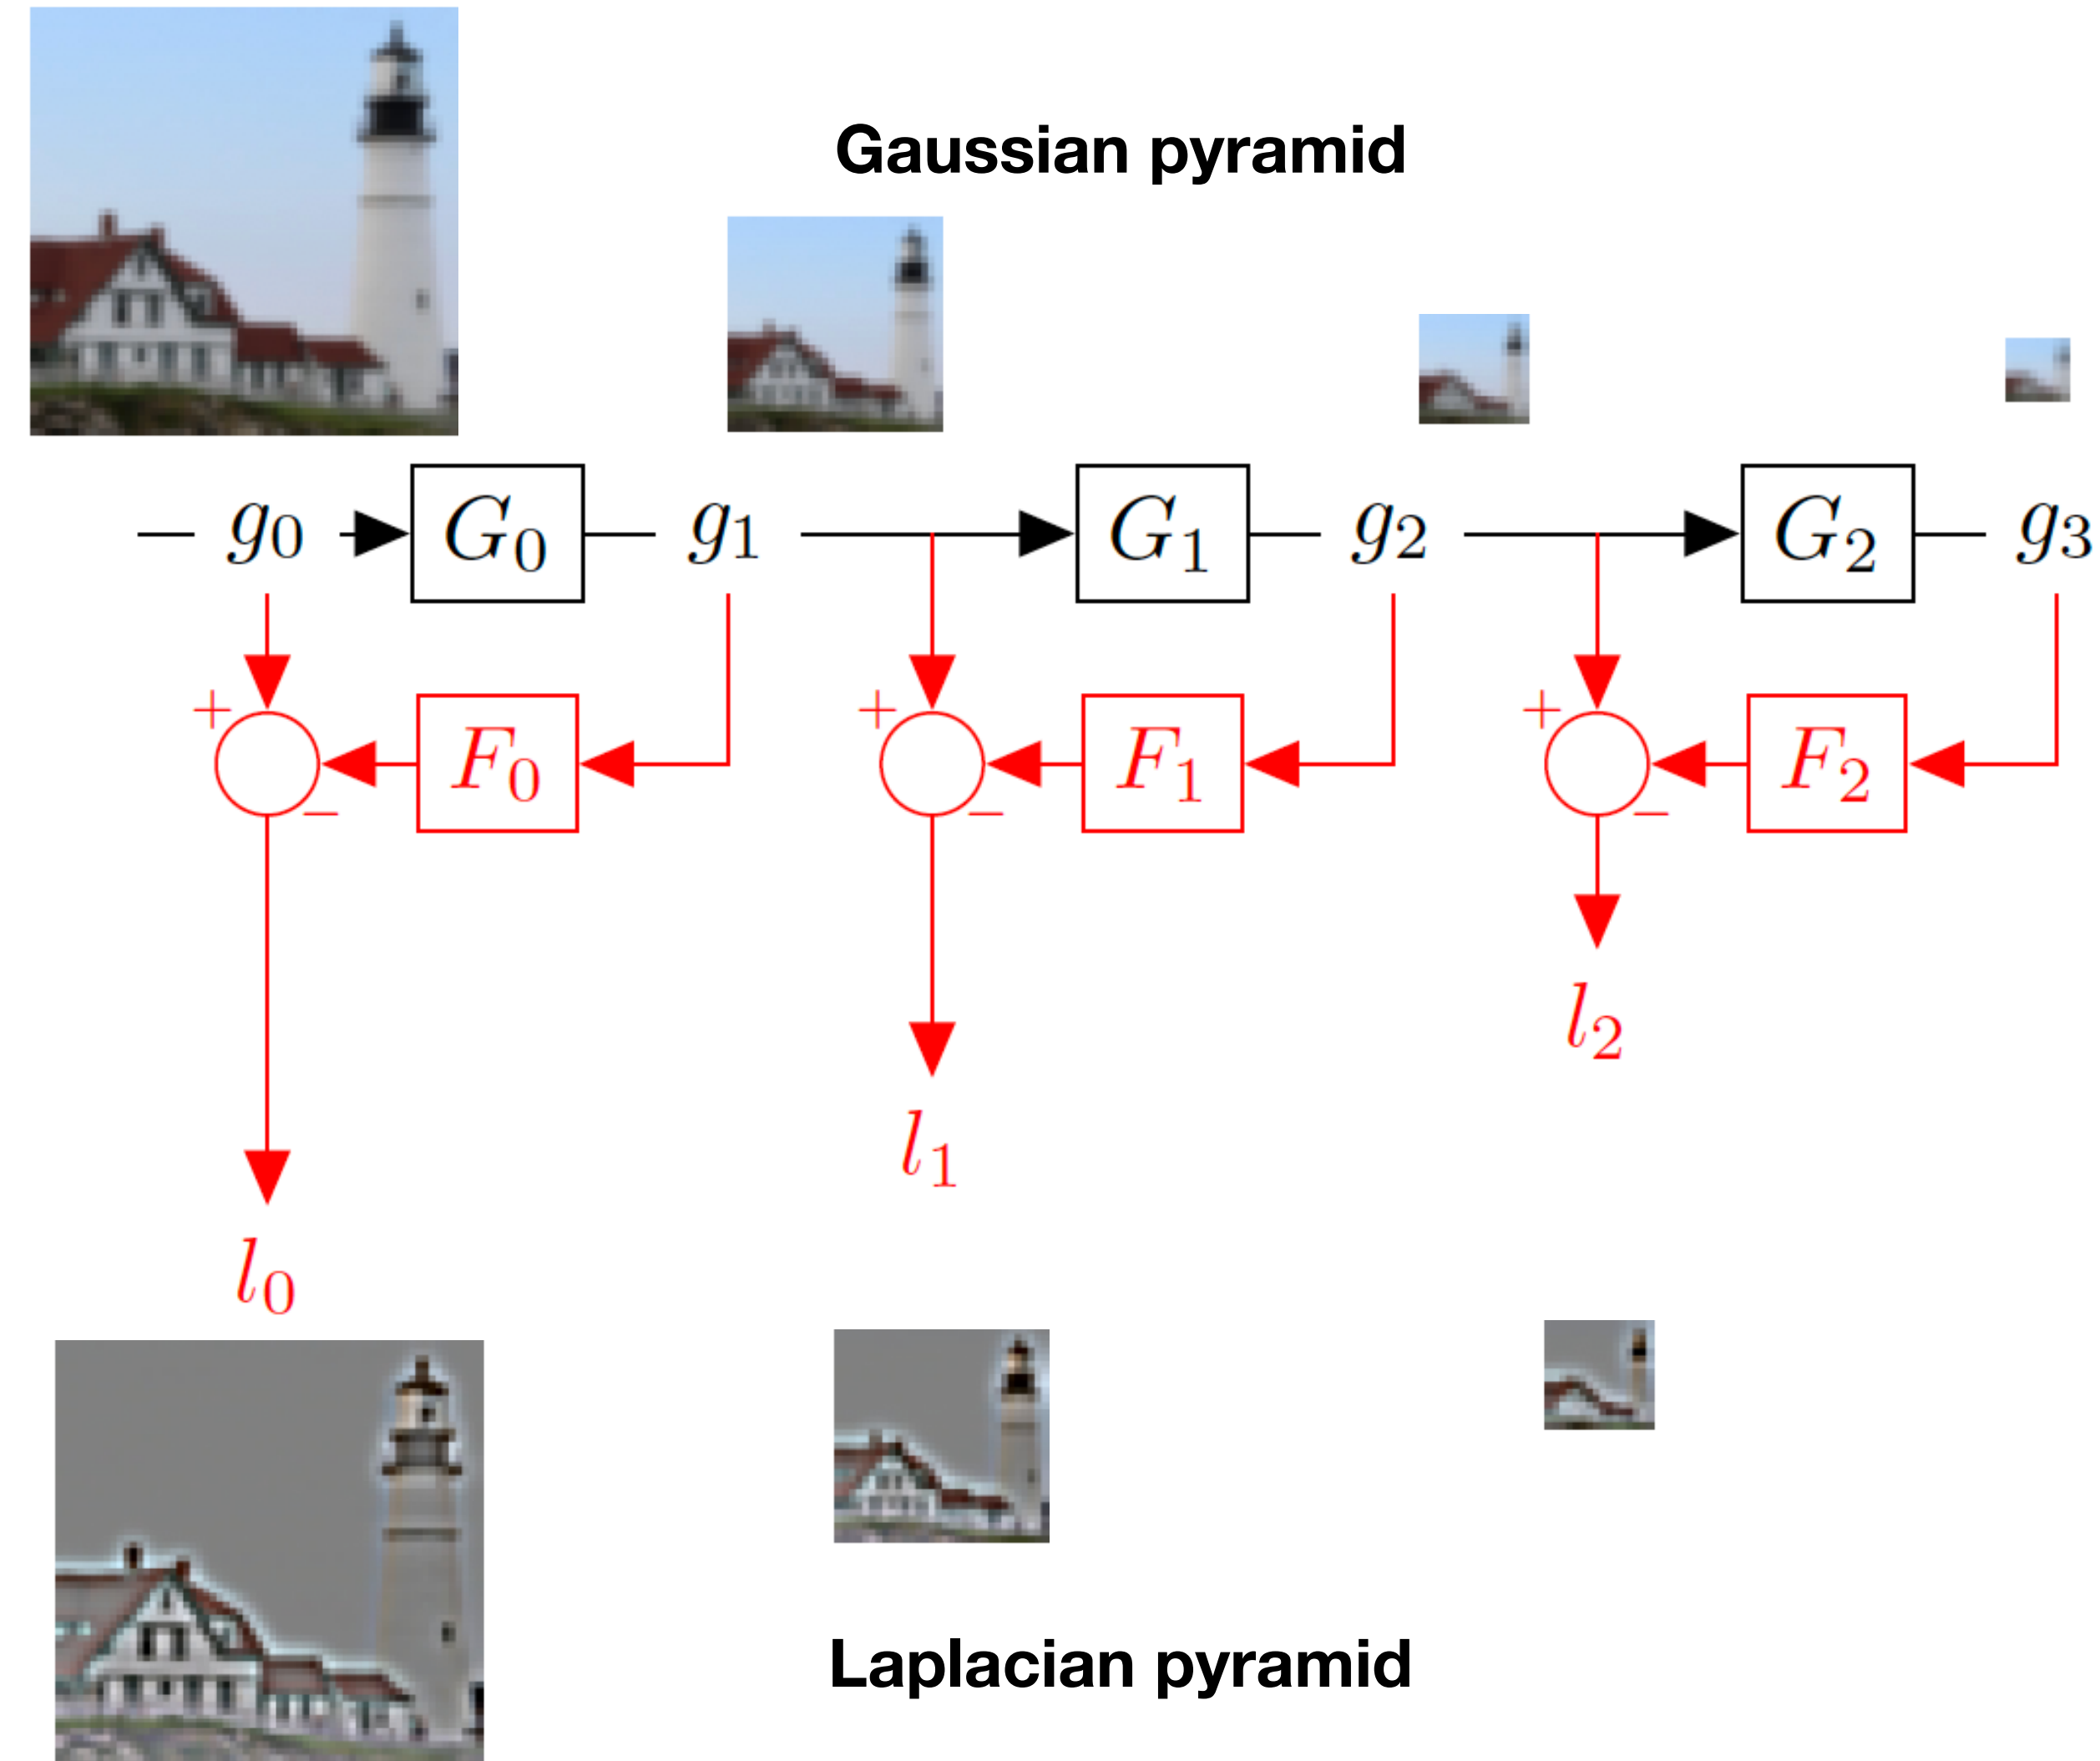


# The Laplacian Pyramid






# The Laplacian Pyramid



# The Laplacian Pyramid



**Blurring and downsampling:**

$$G_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \frac{1}{16}$$

(Downsampling by 2)

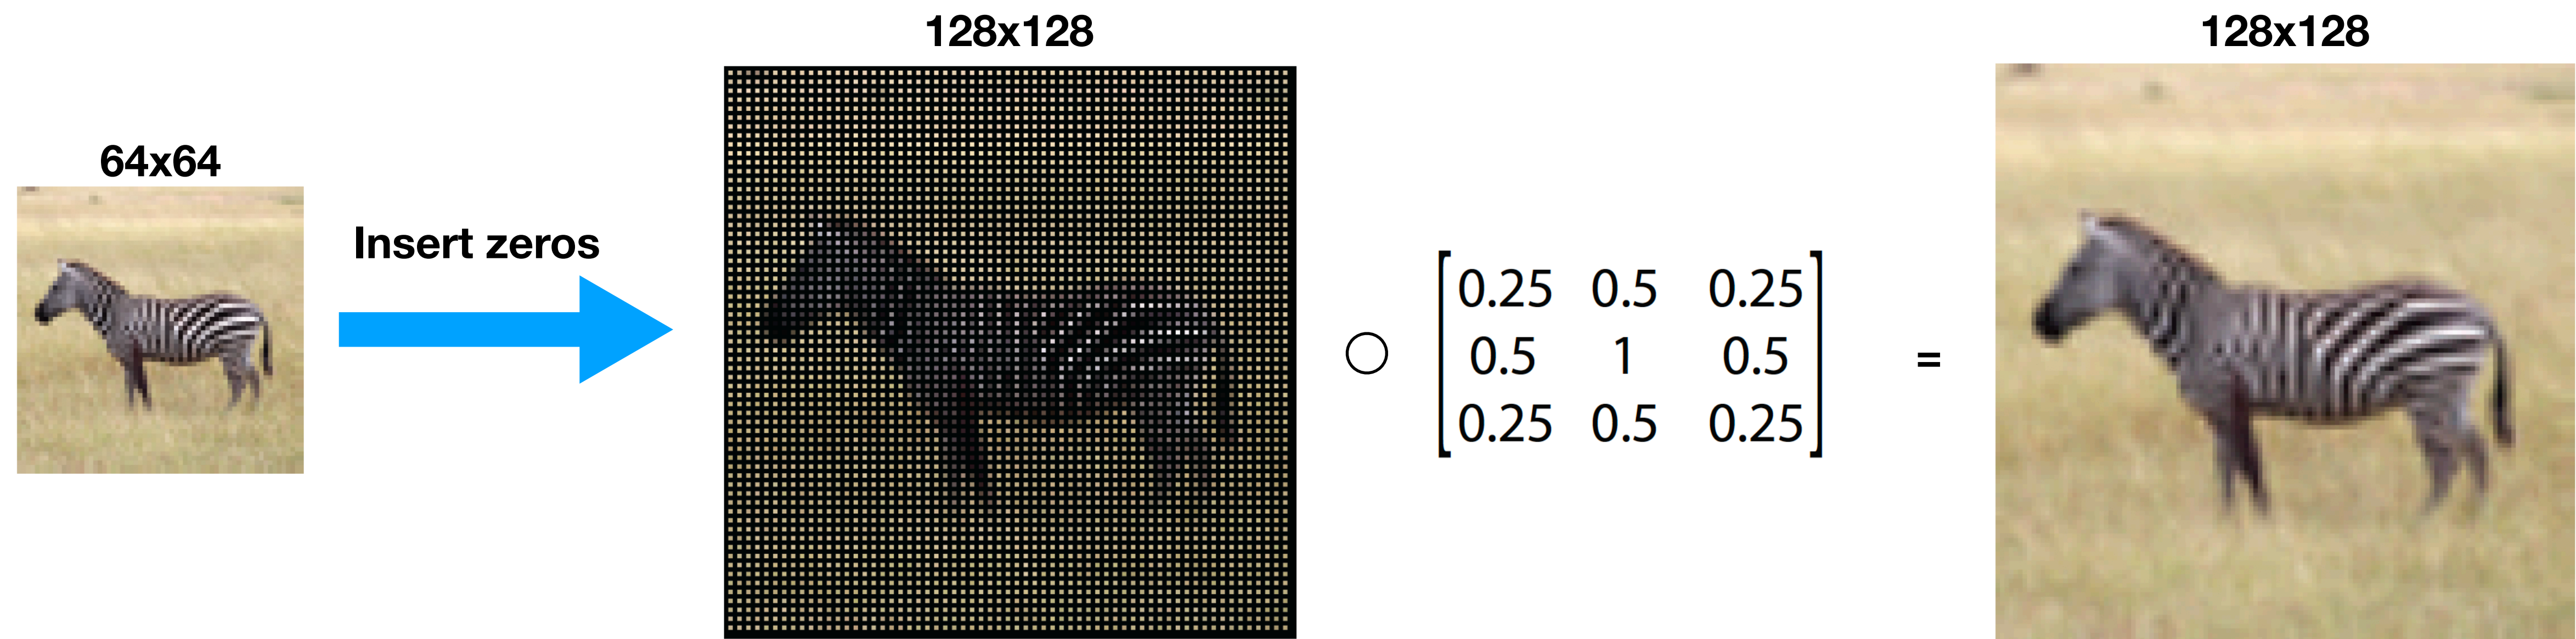
**Upsampling and blurring:**

$$F_0 =$$

$$\begin{bmatrix} 6 & 4 & 1 & 0 & 0 & 0 & 0 & 0 \\ 4 & 6 & 4 & 1 & 0 & 0 & 0 & 0 \\ 1 & 4 & 6 & 4 & 1 & 0 & 0 & 0 \\ 0 & 1 & 4 & 6 & 4 & 1 & 0 & 0 \\ 0 & 0 & 1 & 4 & 6 & 4 & 1 & 0 \\ 0 & 0 & 0 & 1 & 4 & 6 & 4 & 1 \\ 0 & 0 & 0 & 0 & 1 & 4 & 6 & 4 \\ 0 & 0 & 0 & 0 & 0 & 1 & 4 & 6 \end{bmatrix}$$

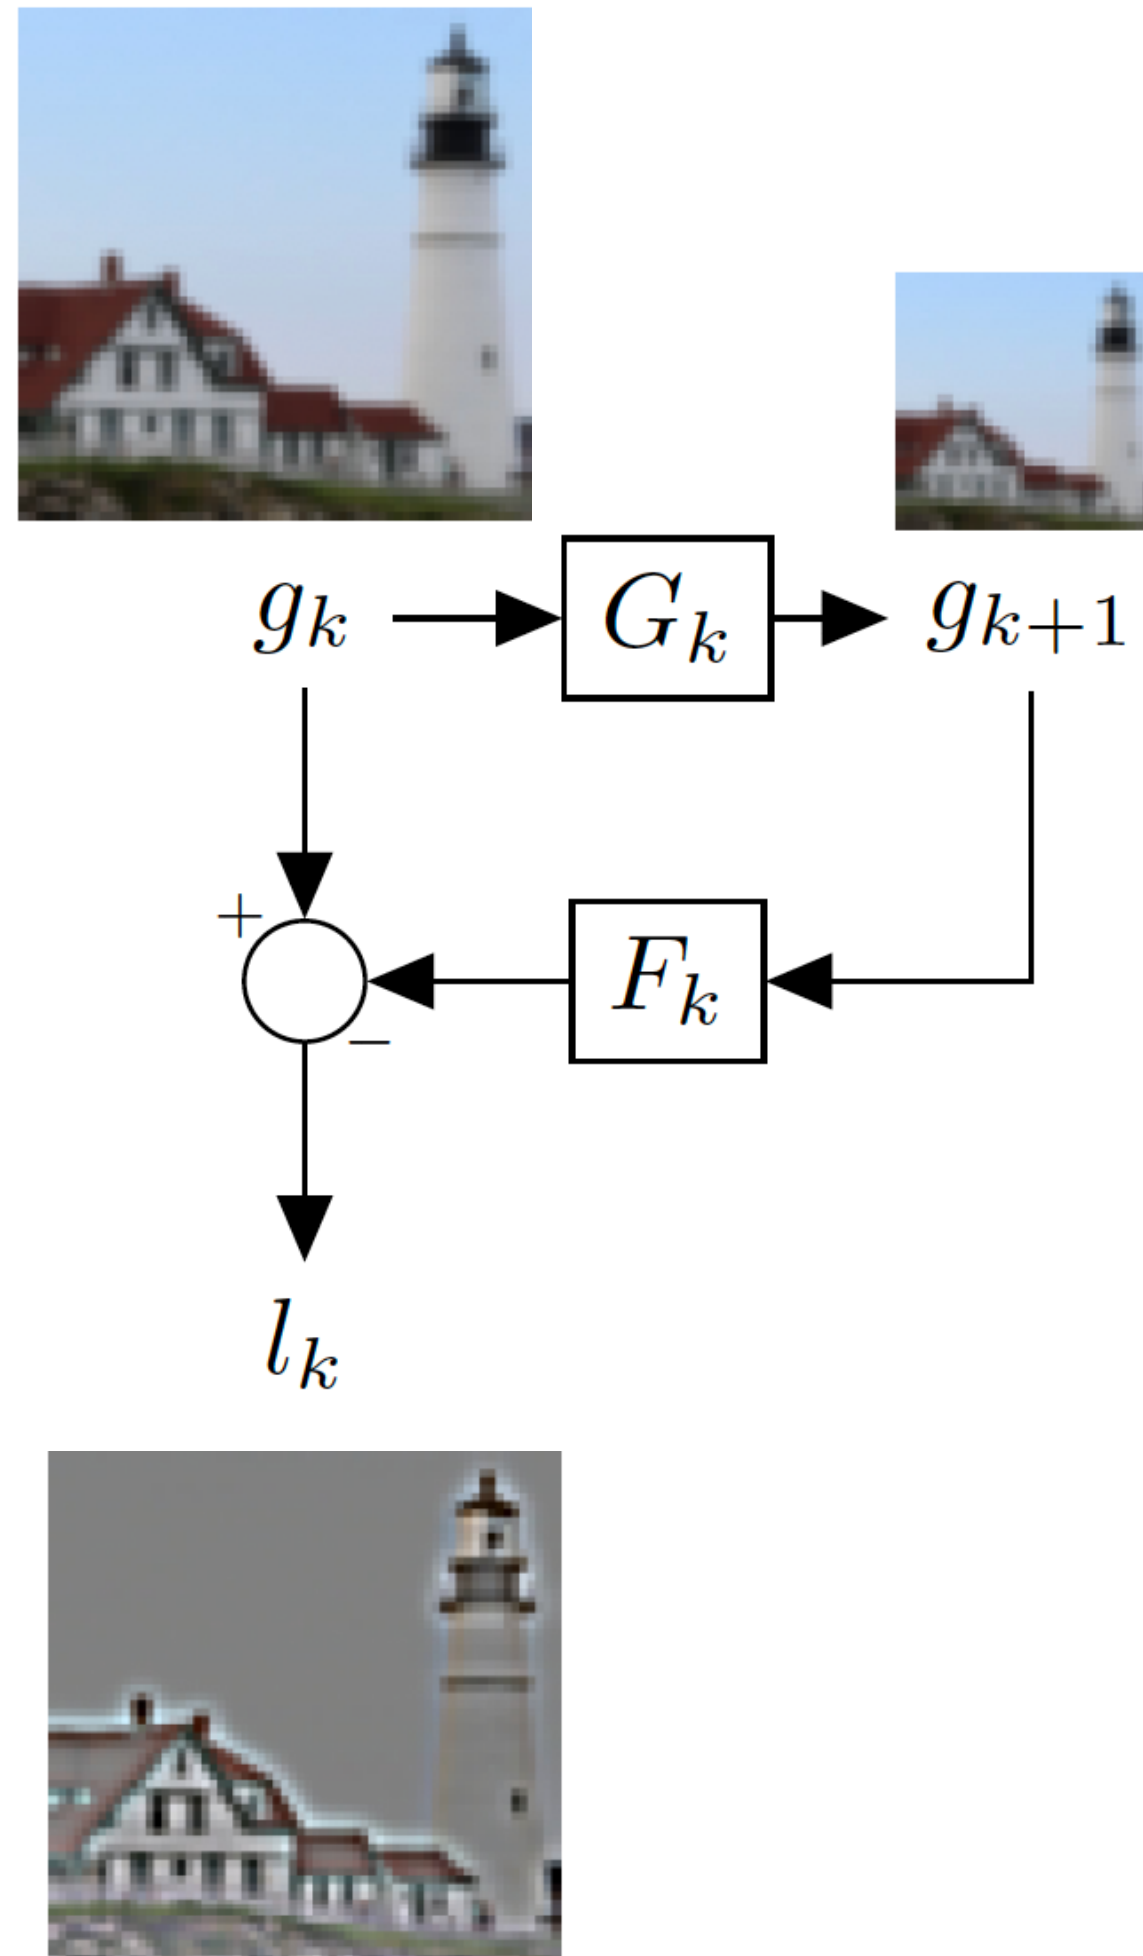
(blur)

# Upsampling





# The Laplacian Pyramid



**Blurring and downsampling:**

$$G_0 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \frac{1}{16} \begin{bmatrix} 6 & 4 & 1 & 0 & 0 & 0 & 0 & 0 \\ 4 & 6 & 4 & 1 & 0 & 0 & 0 & 0 \\ 1 & 4 & 6 & 4 & 1 & 0 & 0 & 0 \\ 0 & 1 & 4 & 6 & 4 & 1 & 0 & 0 \\ 0 & 0 & 1 & 4 & 6 & 4 & 1 & 0 \\ 0 & 0 & 0 & 1 & 4 & 6 & 4 & 1 \\ 0 & 0 & 0 & 0 & 1 & 4 & 6 & 4 \\ 0 & 0 & 0 & 0 & 0 & 1 & 4 & 6 \end{bmatrix}$$

(Downsampling by 2) (blur)

**Upsampling and blurring:**

$$F_0 = \frac{1}{8} \begin{bmatrix} 6 & 4 & 1 & 0 & 0 & 0 & 0 & 0 \\ 4 & 6 & 4 & 1 & 0 & 0 & 0 & 0 \\ 1 & 4 & 6 & 4 & 1 & 0 & 0 & 0 \\ 0 & 1 & 4 & 6 & 4 & 1 & 0 & 0 \\ 0 & 0 & 1 & 4 & 6 & 4 & 1 & 0 \\ 0 & 0 & 0 & 1 & 4 & 6 & 4 & 1 \\ 0 & 0 & 0 & 0 & 1 & 4 & 6 & 4 \\ 0 & 0 & 0 & 0 & 0 & 1 & 4 & 6 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$


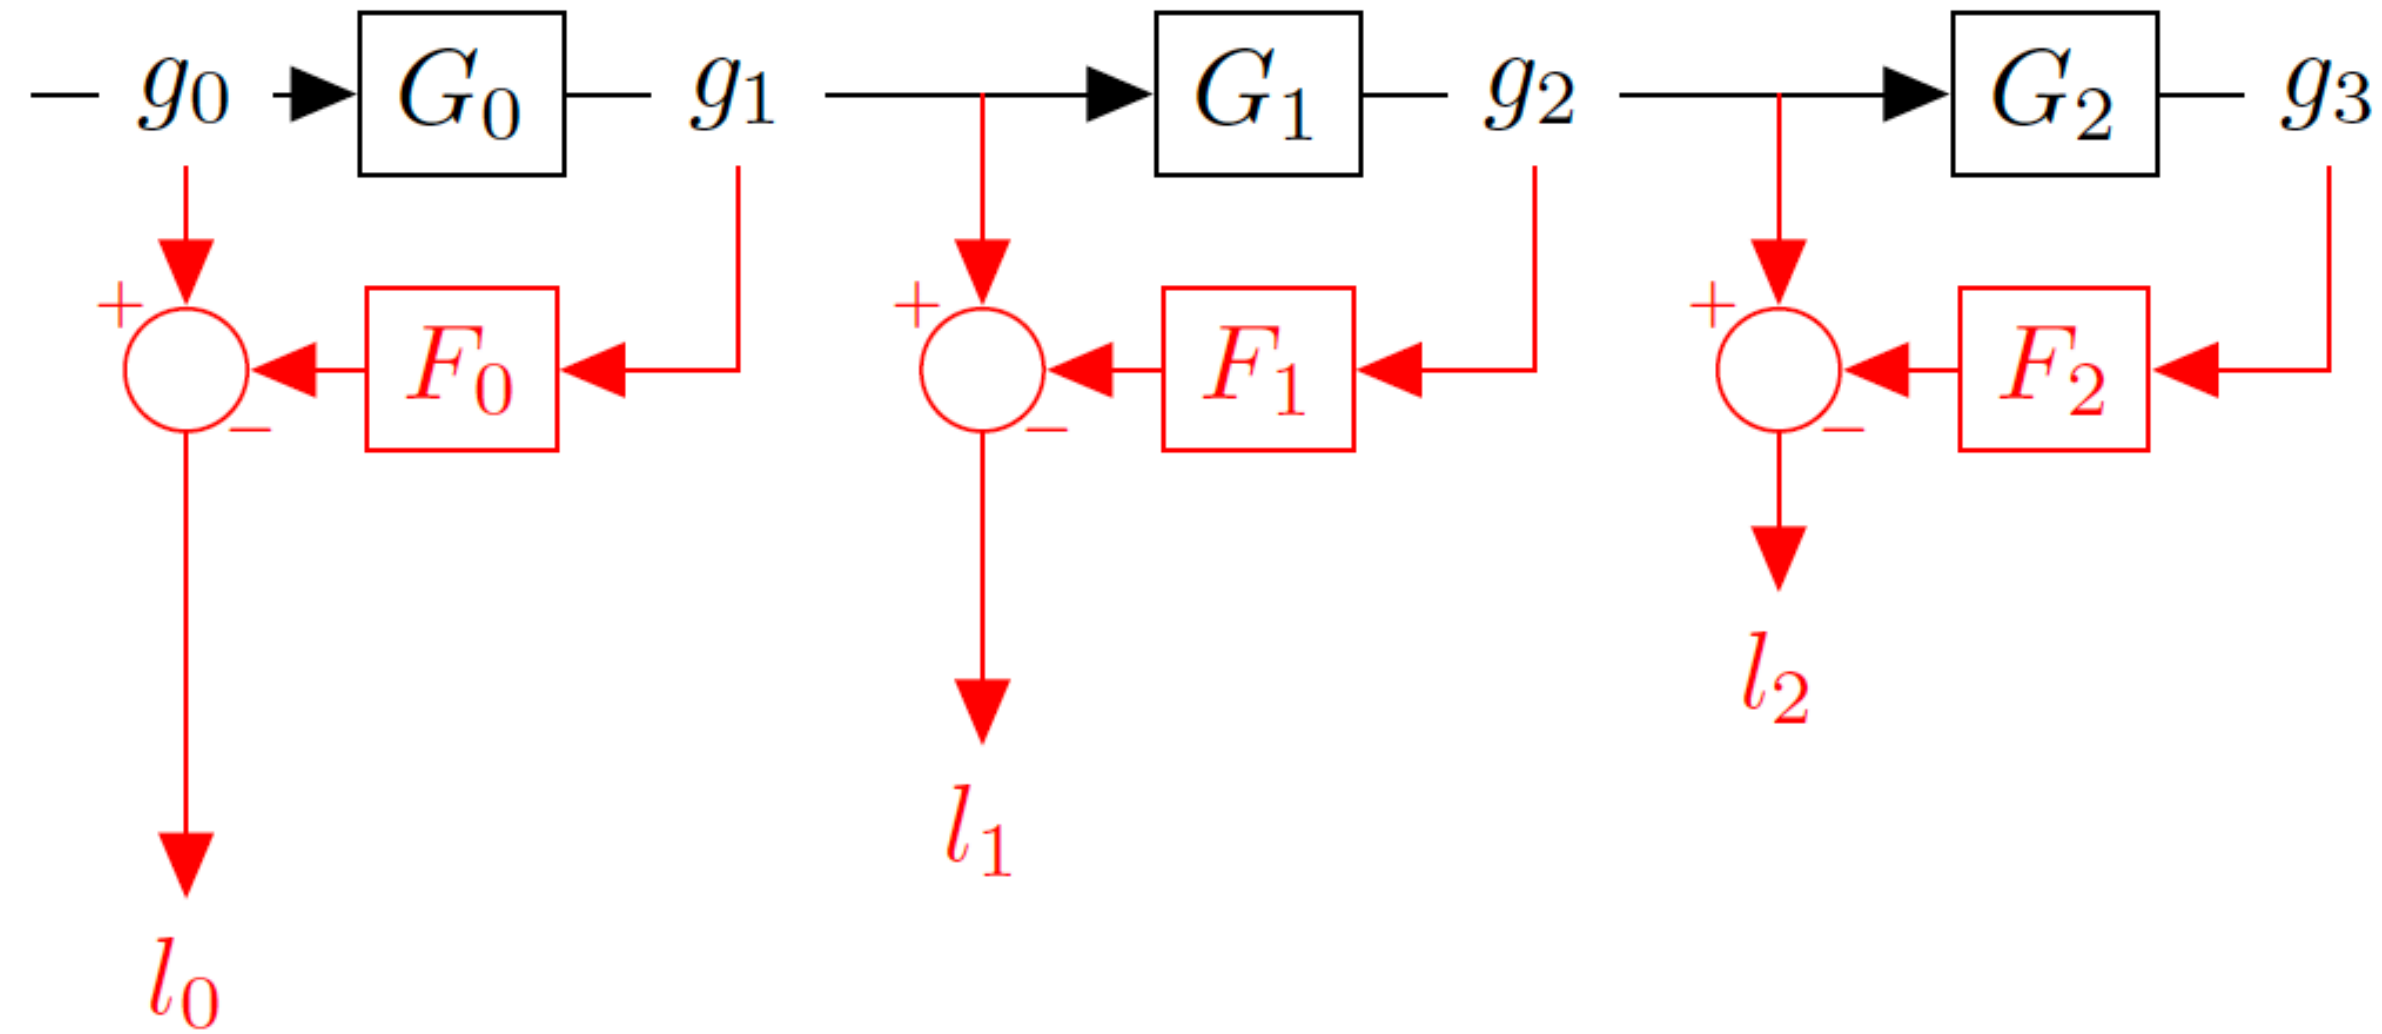

(blur) (Upsampling by 2)

$$l_0 = (I_0 - F_0 G_0) g_0$$


# The Laplacian Pyramid



Gaussian pyramid



Laplacian pyramid







# The Laplacian Pyramid

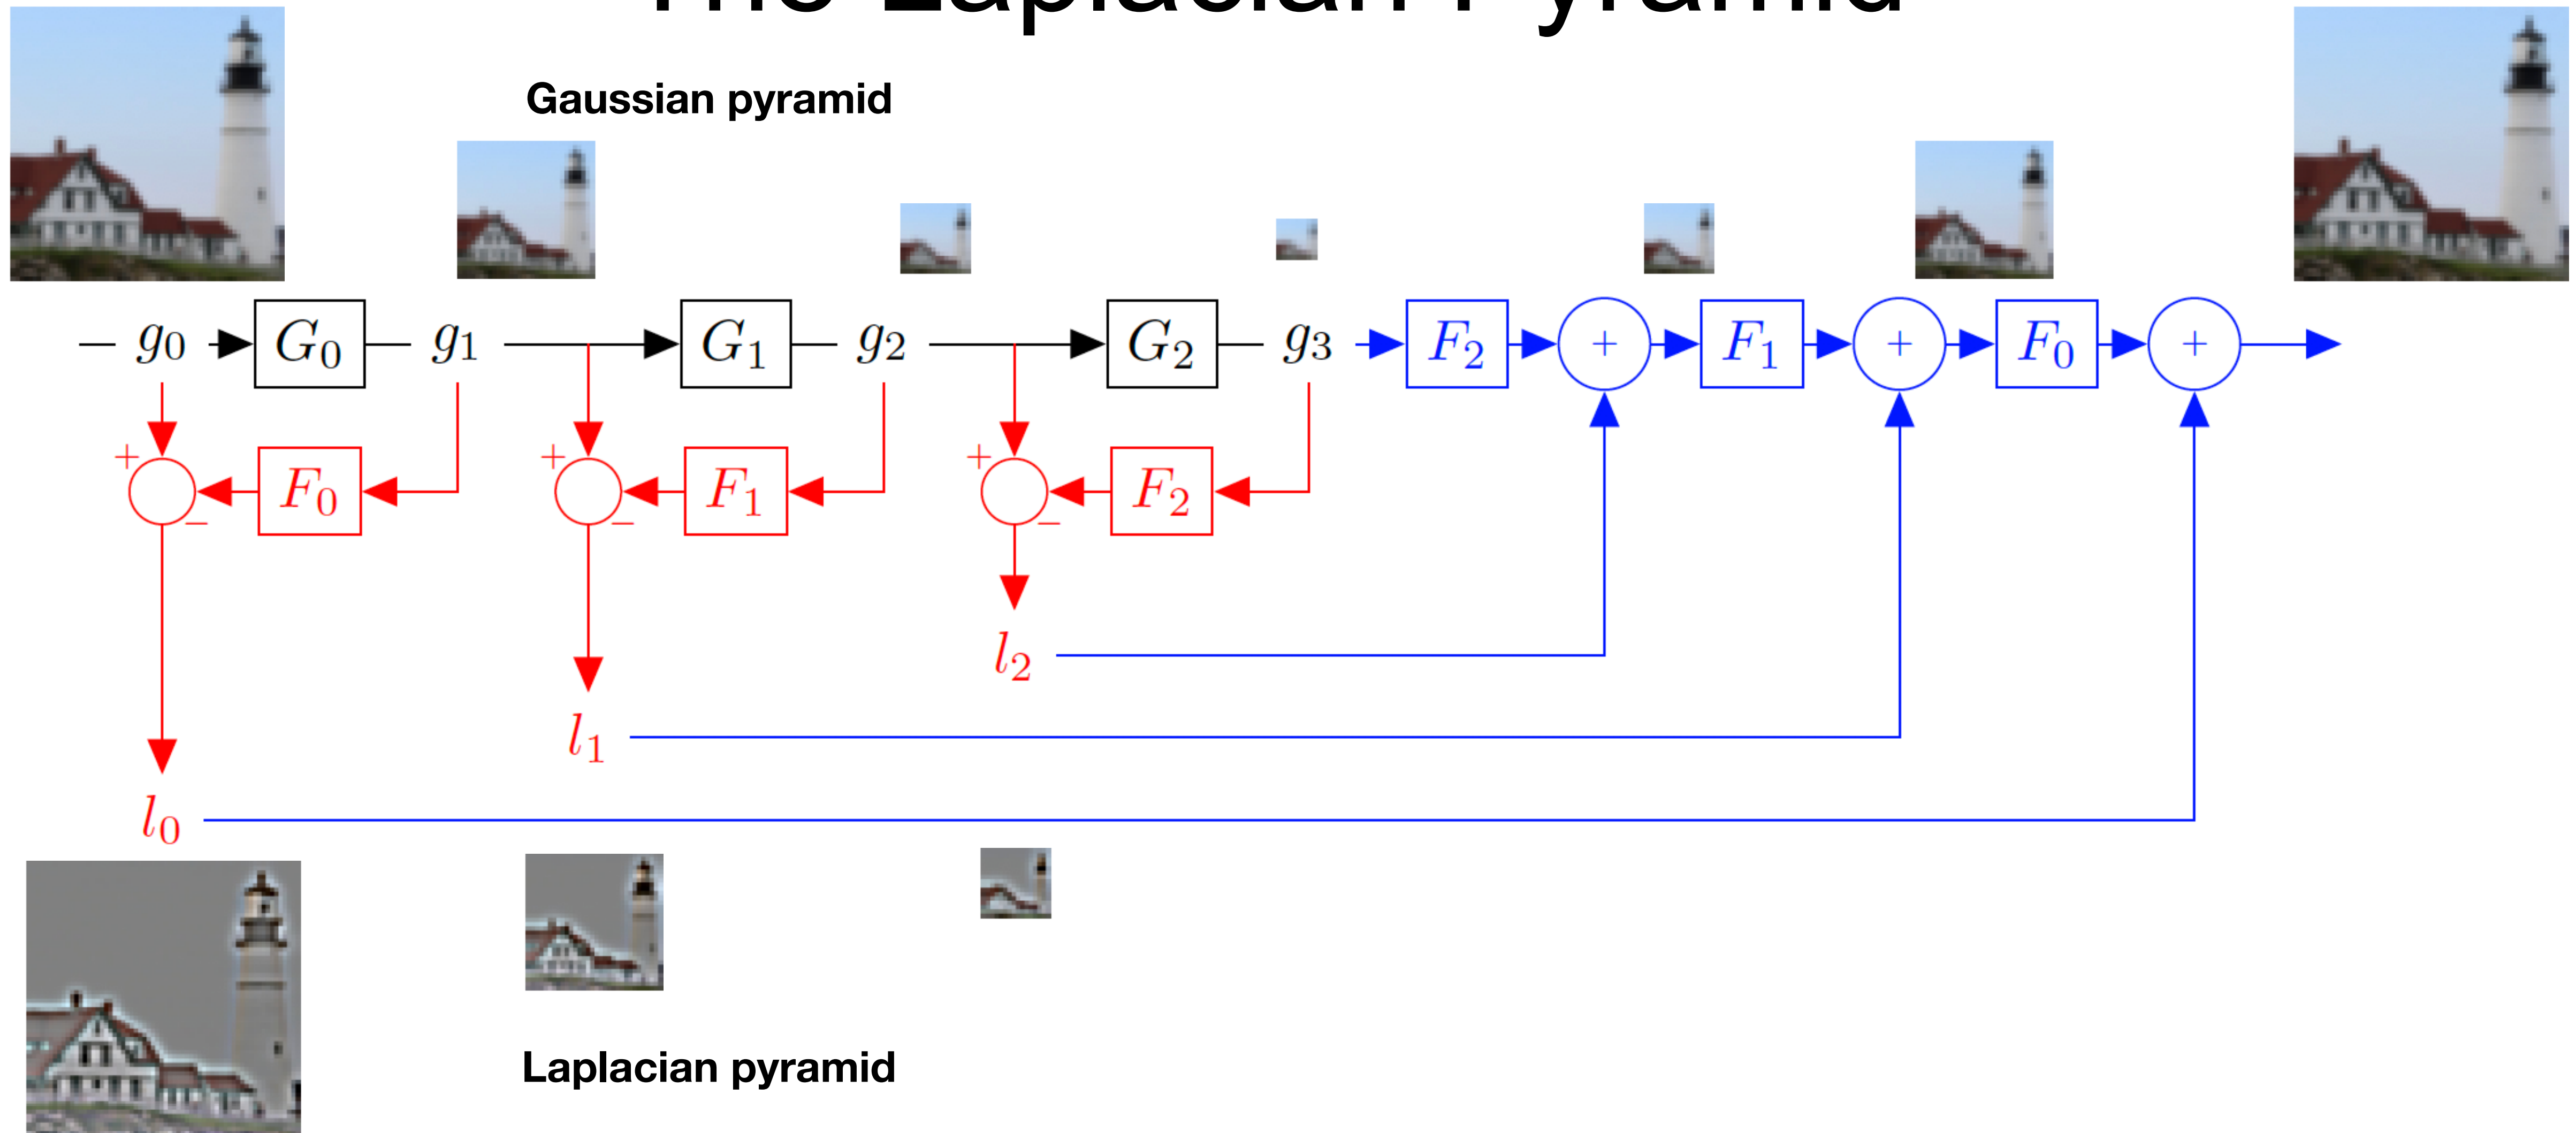


Laplacian pyramid



Can we invert the  
Laplacian Pyramid?

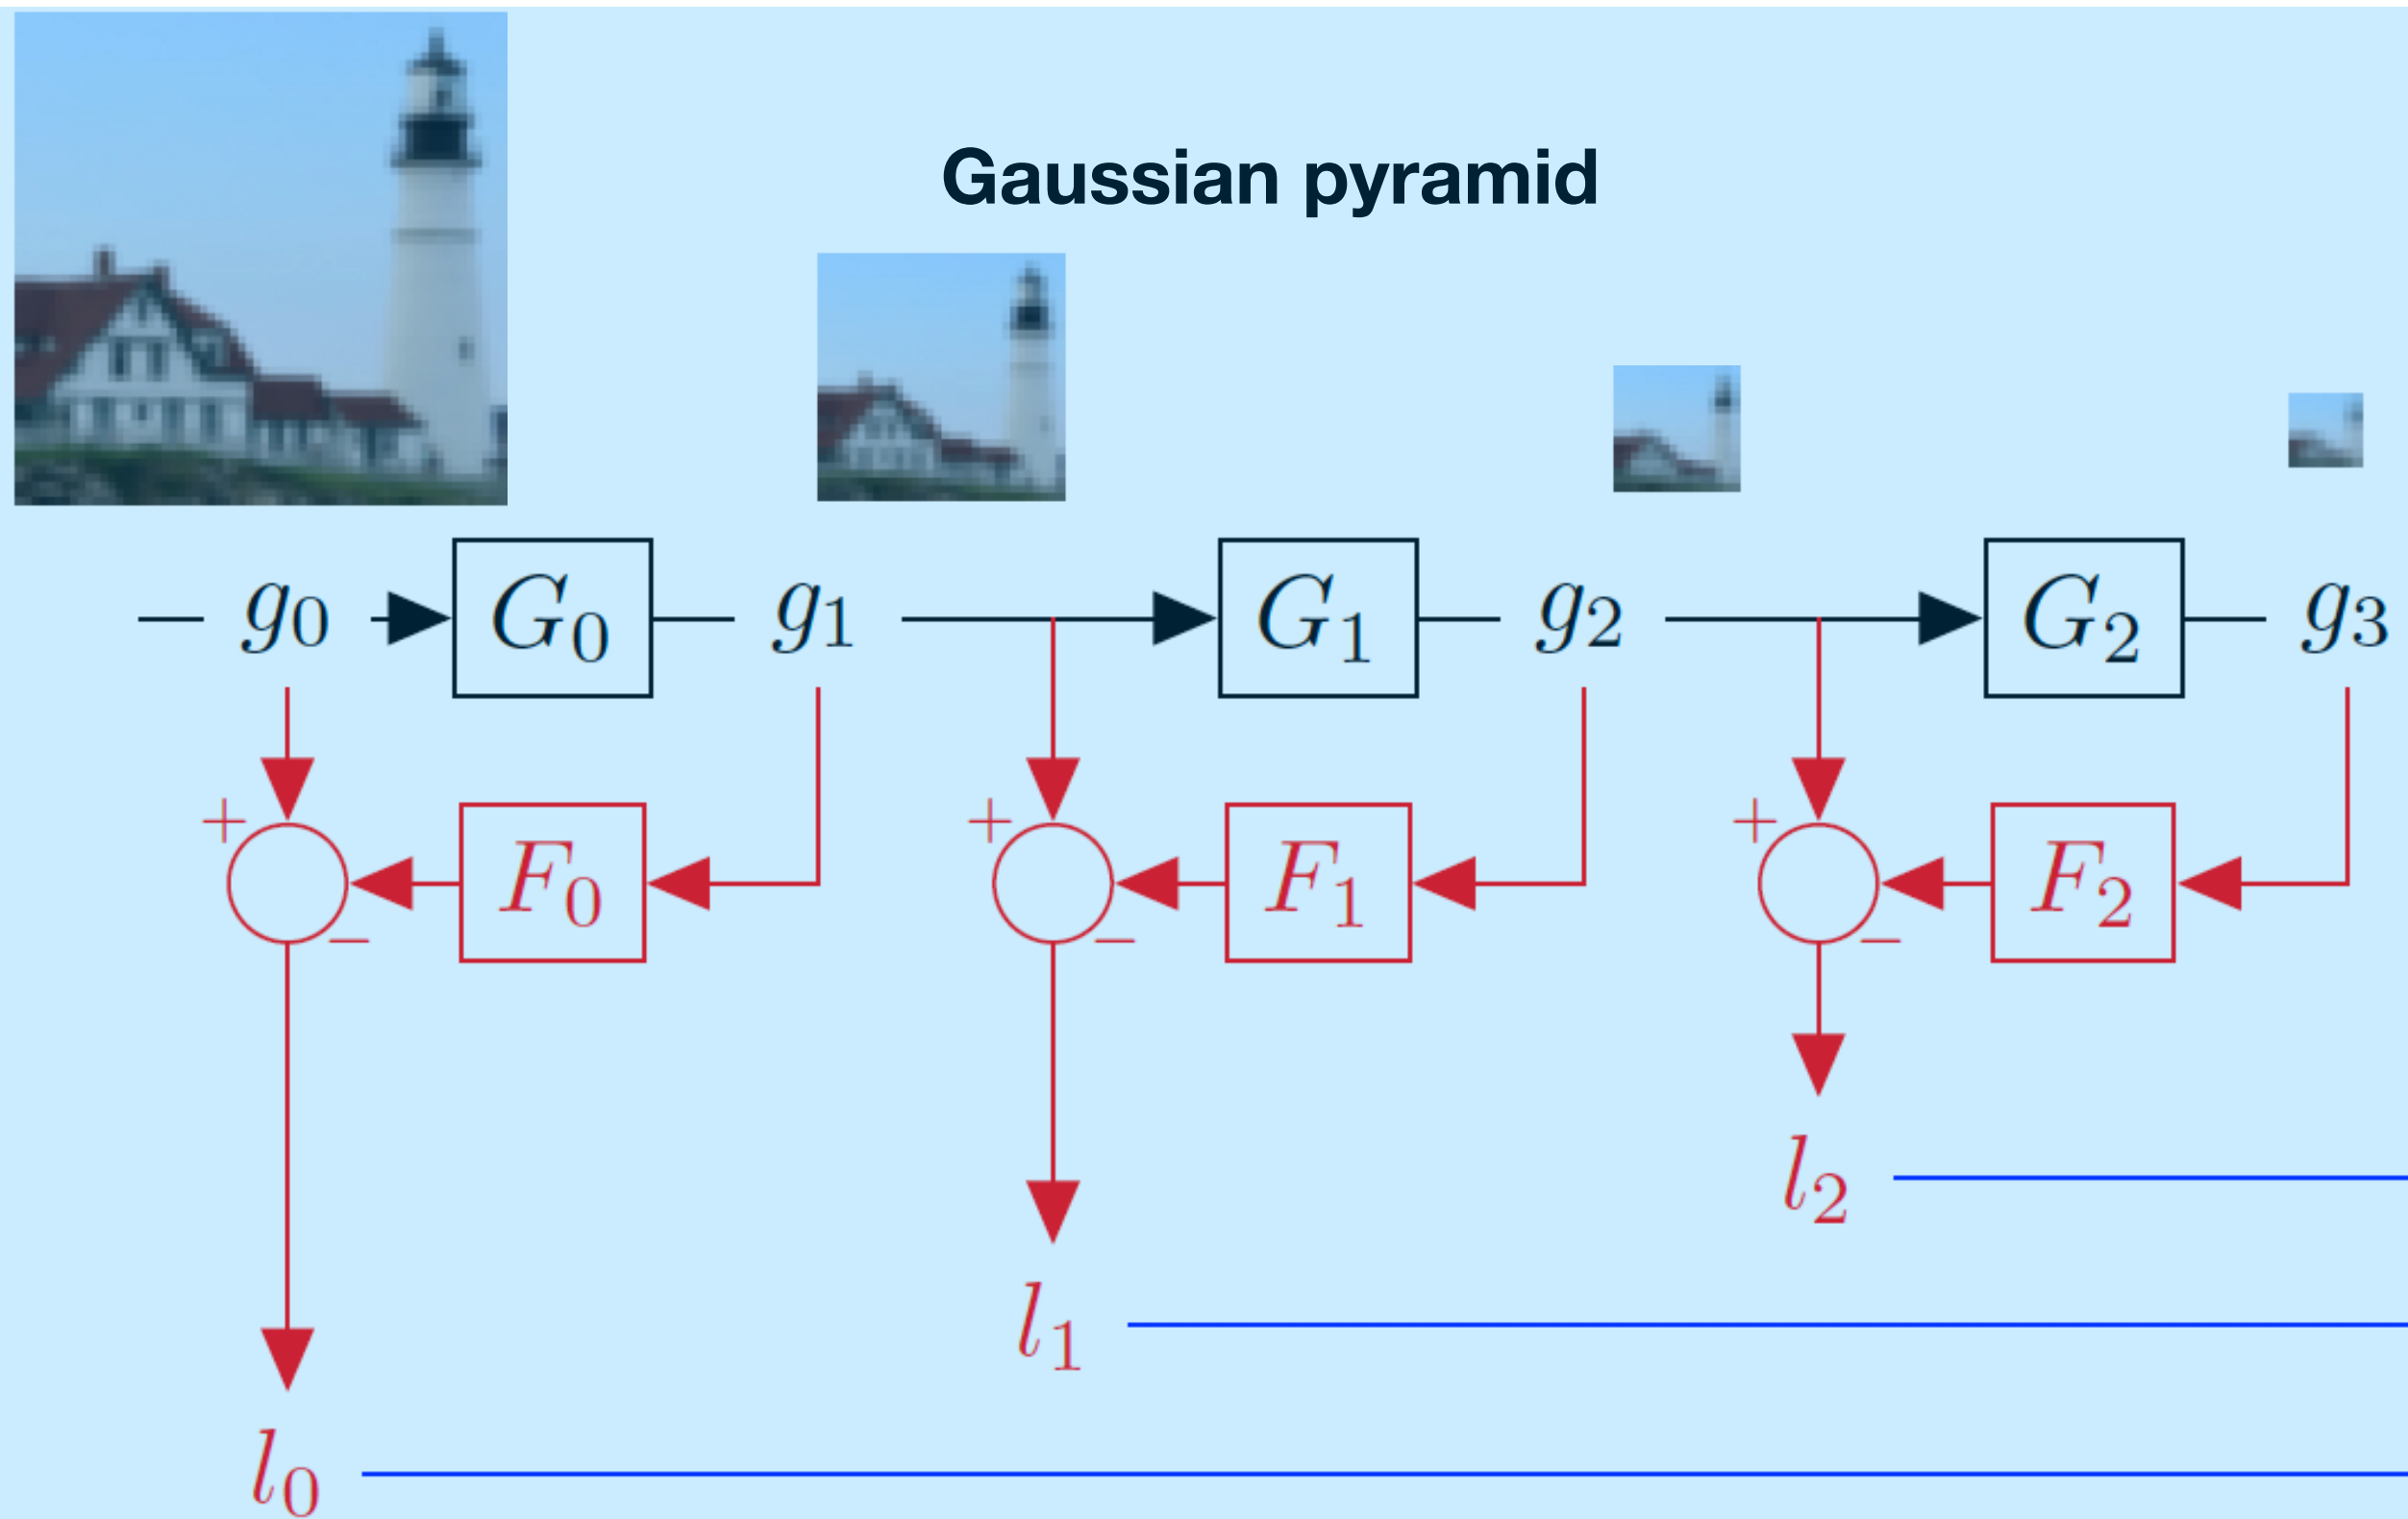
# The Laplacian Pyramid





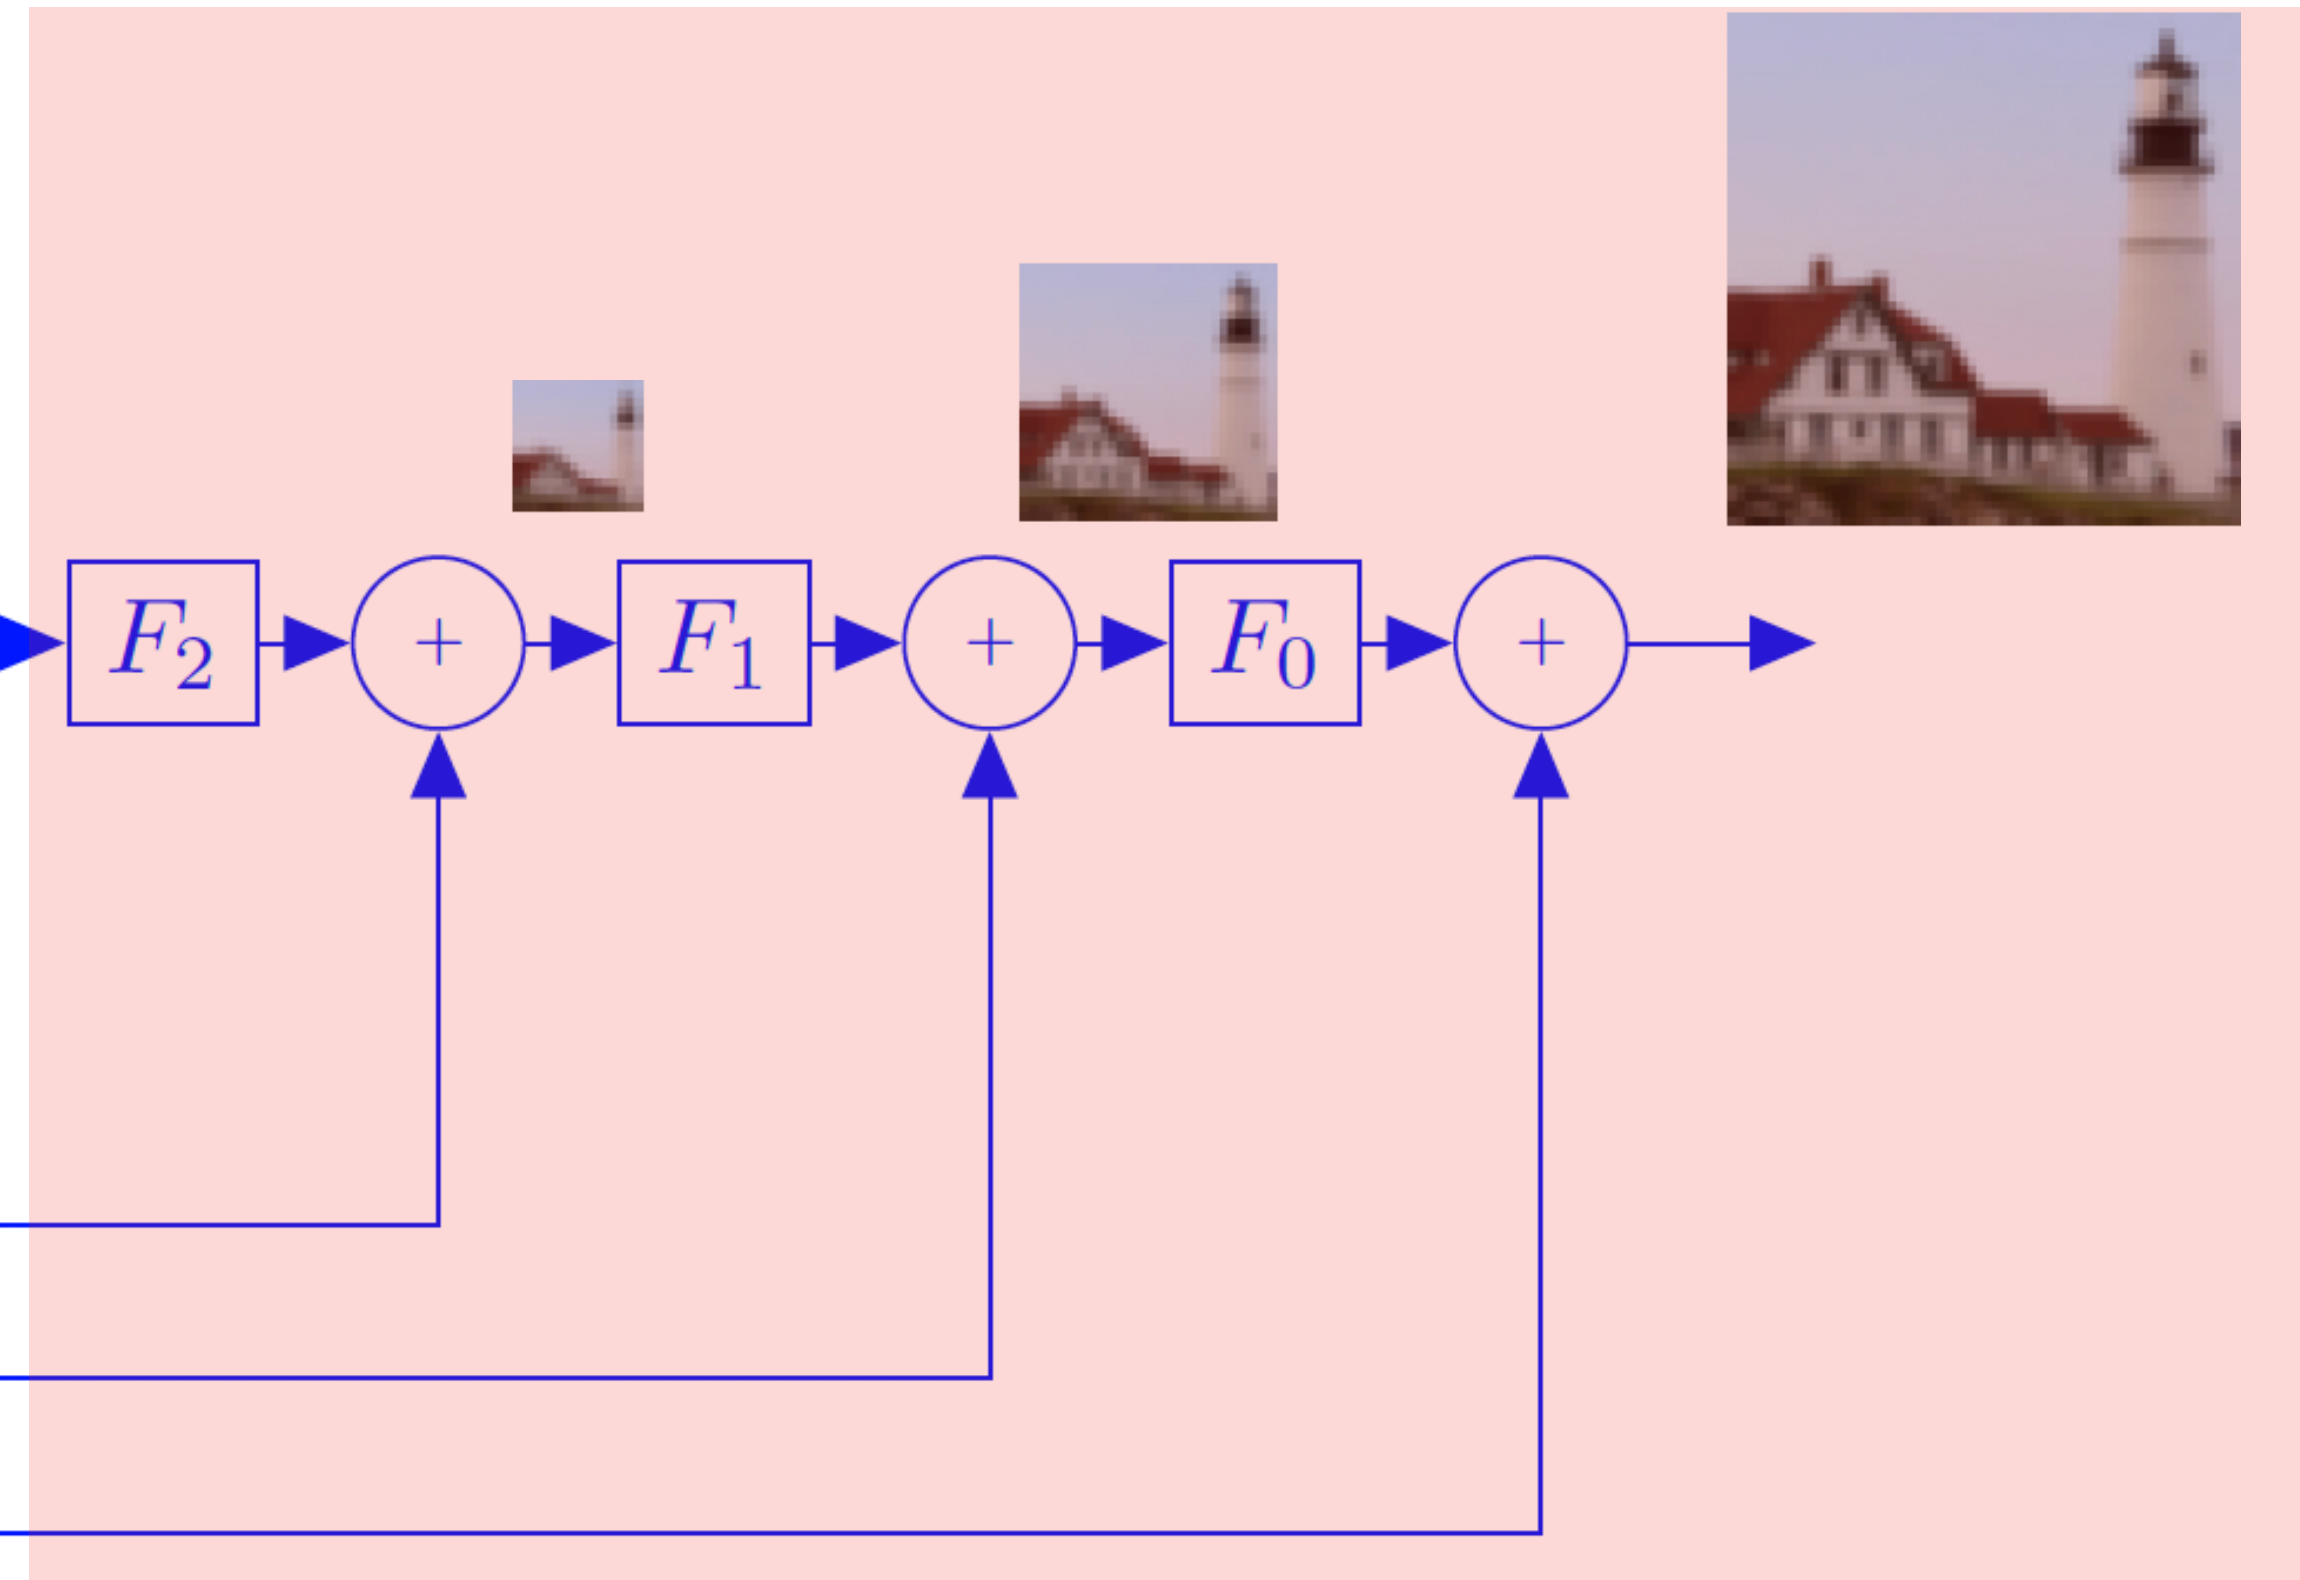
# The Laplacian Pyramid

Gaussian pyramid



Laplacian pyramid

Analysis/Encoder



Synthesis/Decoder

# Laplacian pyramid applications

- Texture synthesis
- Image compression
- Noise removal
- Computing image features (e.g., SIFT)



# Image Blending





# Image Blending






# Image Blending



$I^A$



$I^B$



$m$






















$I$

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



# Image Blending with the Laplacian Pyramid



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




# Image Blending with the Laplacian Pyramid









+



=



Simple blend



With Laplacian pyr.





Photo credit: Chris Cameron



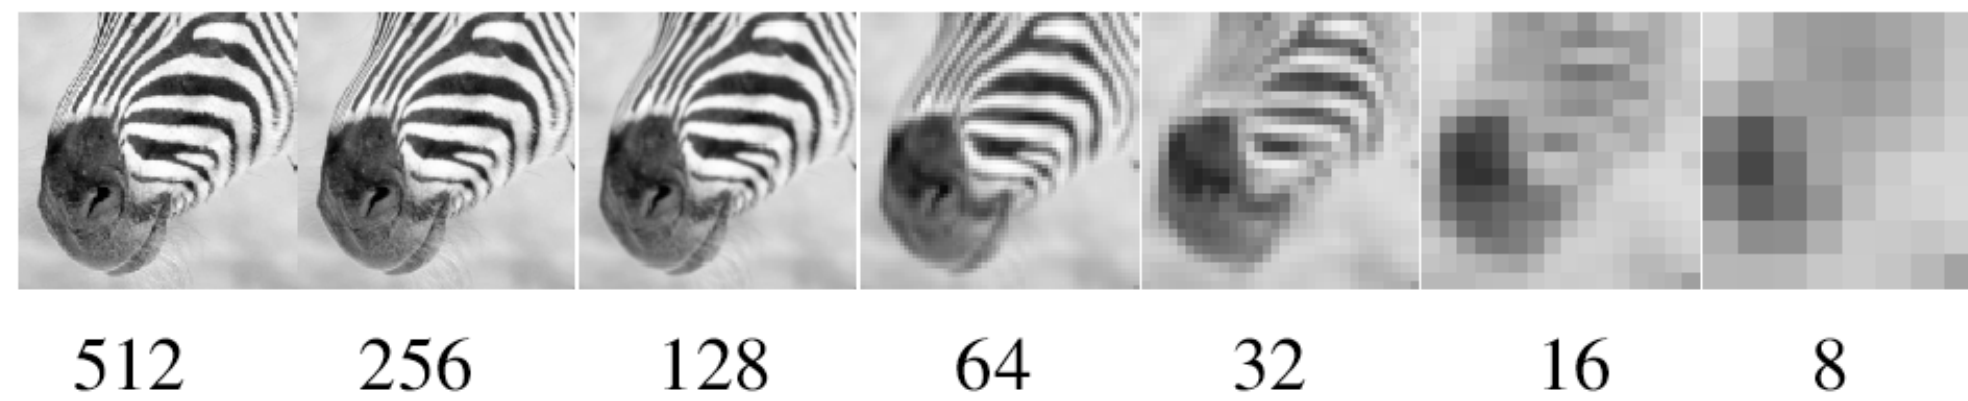


# Image Blending (PS2 problem)

- Build Laplacian pyramid for both images:  $L_A$ ,  $L_B$
- Build Gaussian pyramid for mask:  $G$
- Build a combined Laplacian pyramid:
- Collapse  $L$  to obtain the blended image



# Image pyramids



Gaussian Pyr.

Laplacian Pyr.

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








# Orientations






# Steerable Pyramid


Oriented gradient




# Linear Image Transforms



a)  
Fourier transform



Gaussian pyr.



Laplacian pyr.



Steerable pyr.







# Texture


Stationary

Stochastic






# Texture analysis




What we'd like: are they made of the same “stuff”. Are these textures similar?



# How do humans analyze texture?



Human vision is sensitive to the difference of some types of elements and appears to be “numb” on other types of differences.



Béla Julesz



# Pre-attentive texture discrimination



Bela Julesz, "Textons, the Elements of Texture Perception, and their Interactions". Nature 290: 91-97. March, 1981.

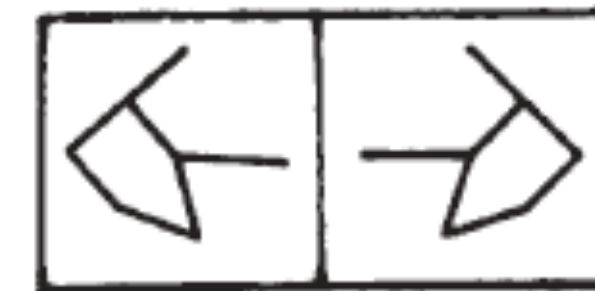
# Pre-attentive texture discrimination



Bela Julesz, "Textons, the Elements of Texture Perception, and their Interactions". Nature 290: 91-97. March, 1981.



# Pre-attentive texture discrimination






This texture pair is pre-attentively indistinguishable. Why?

Bela Julesz, "Textons, the Elements of Texture Perception, and their Interactions". Nature 290: 91-97. March, 1981.

Source: A. Efros


# Search Experiment I



The subject is told to detect a target element in a number of background elements.  
In this example, the detection time is independent of the number of background elements.



# Search Experiment II



Here detection time is proportional to the number of background elements, and thus suggests that the subject is doing element-by-element scrutiny.

# Heuristic

Julesz then conjectured the following:

Human vision operates in two distinct modes:

1. Preattentive vision

parallel, instantaneous ( $\sim 100\text{--}200\text{ms}$ ), without scrutiny,  
independent of the number of patterns, covering a large visual field.

2. Attentive vision

serial search by focal attention in 50ms steps limited to small aperture.




# Julesz Conjecture

*Textures cannot be spontaneously discriminated if they have the same first-order and second-order statistics and differ only in their third-order or higher-order statistics.*

(later proved wrong)

# 1<sup>st</sup> order statistics differ




5% white

20% white




# 2<sup>nd</sup> order statistics differ



10% white


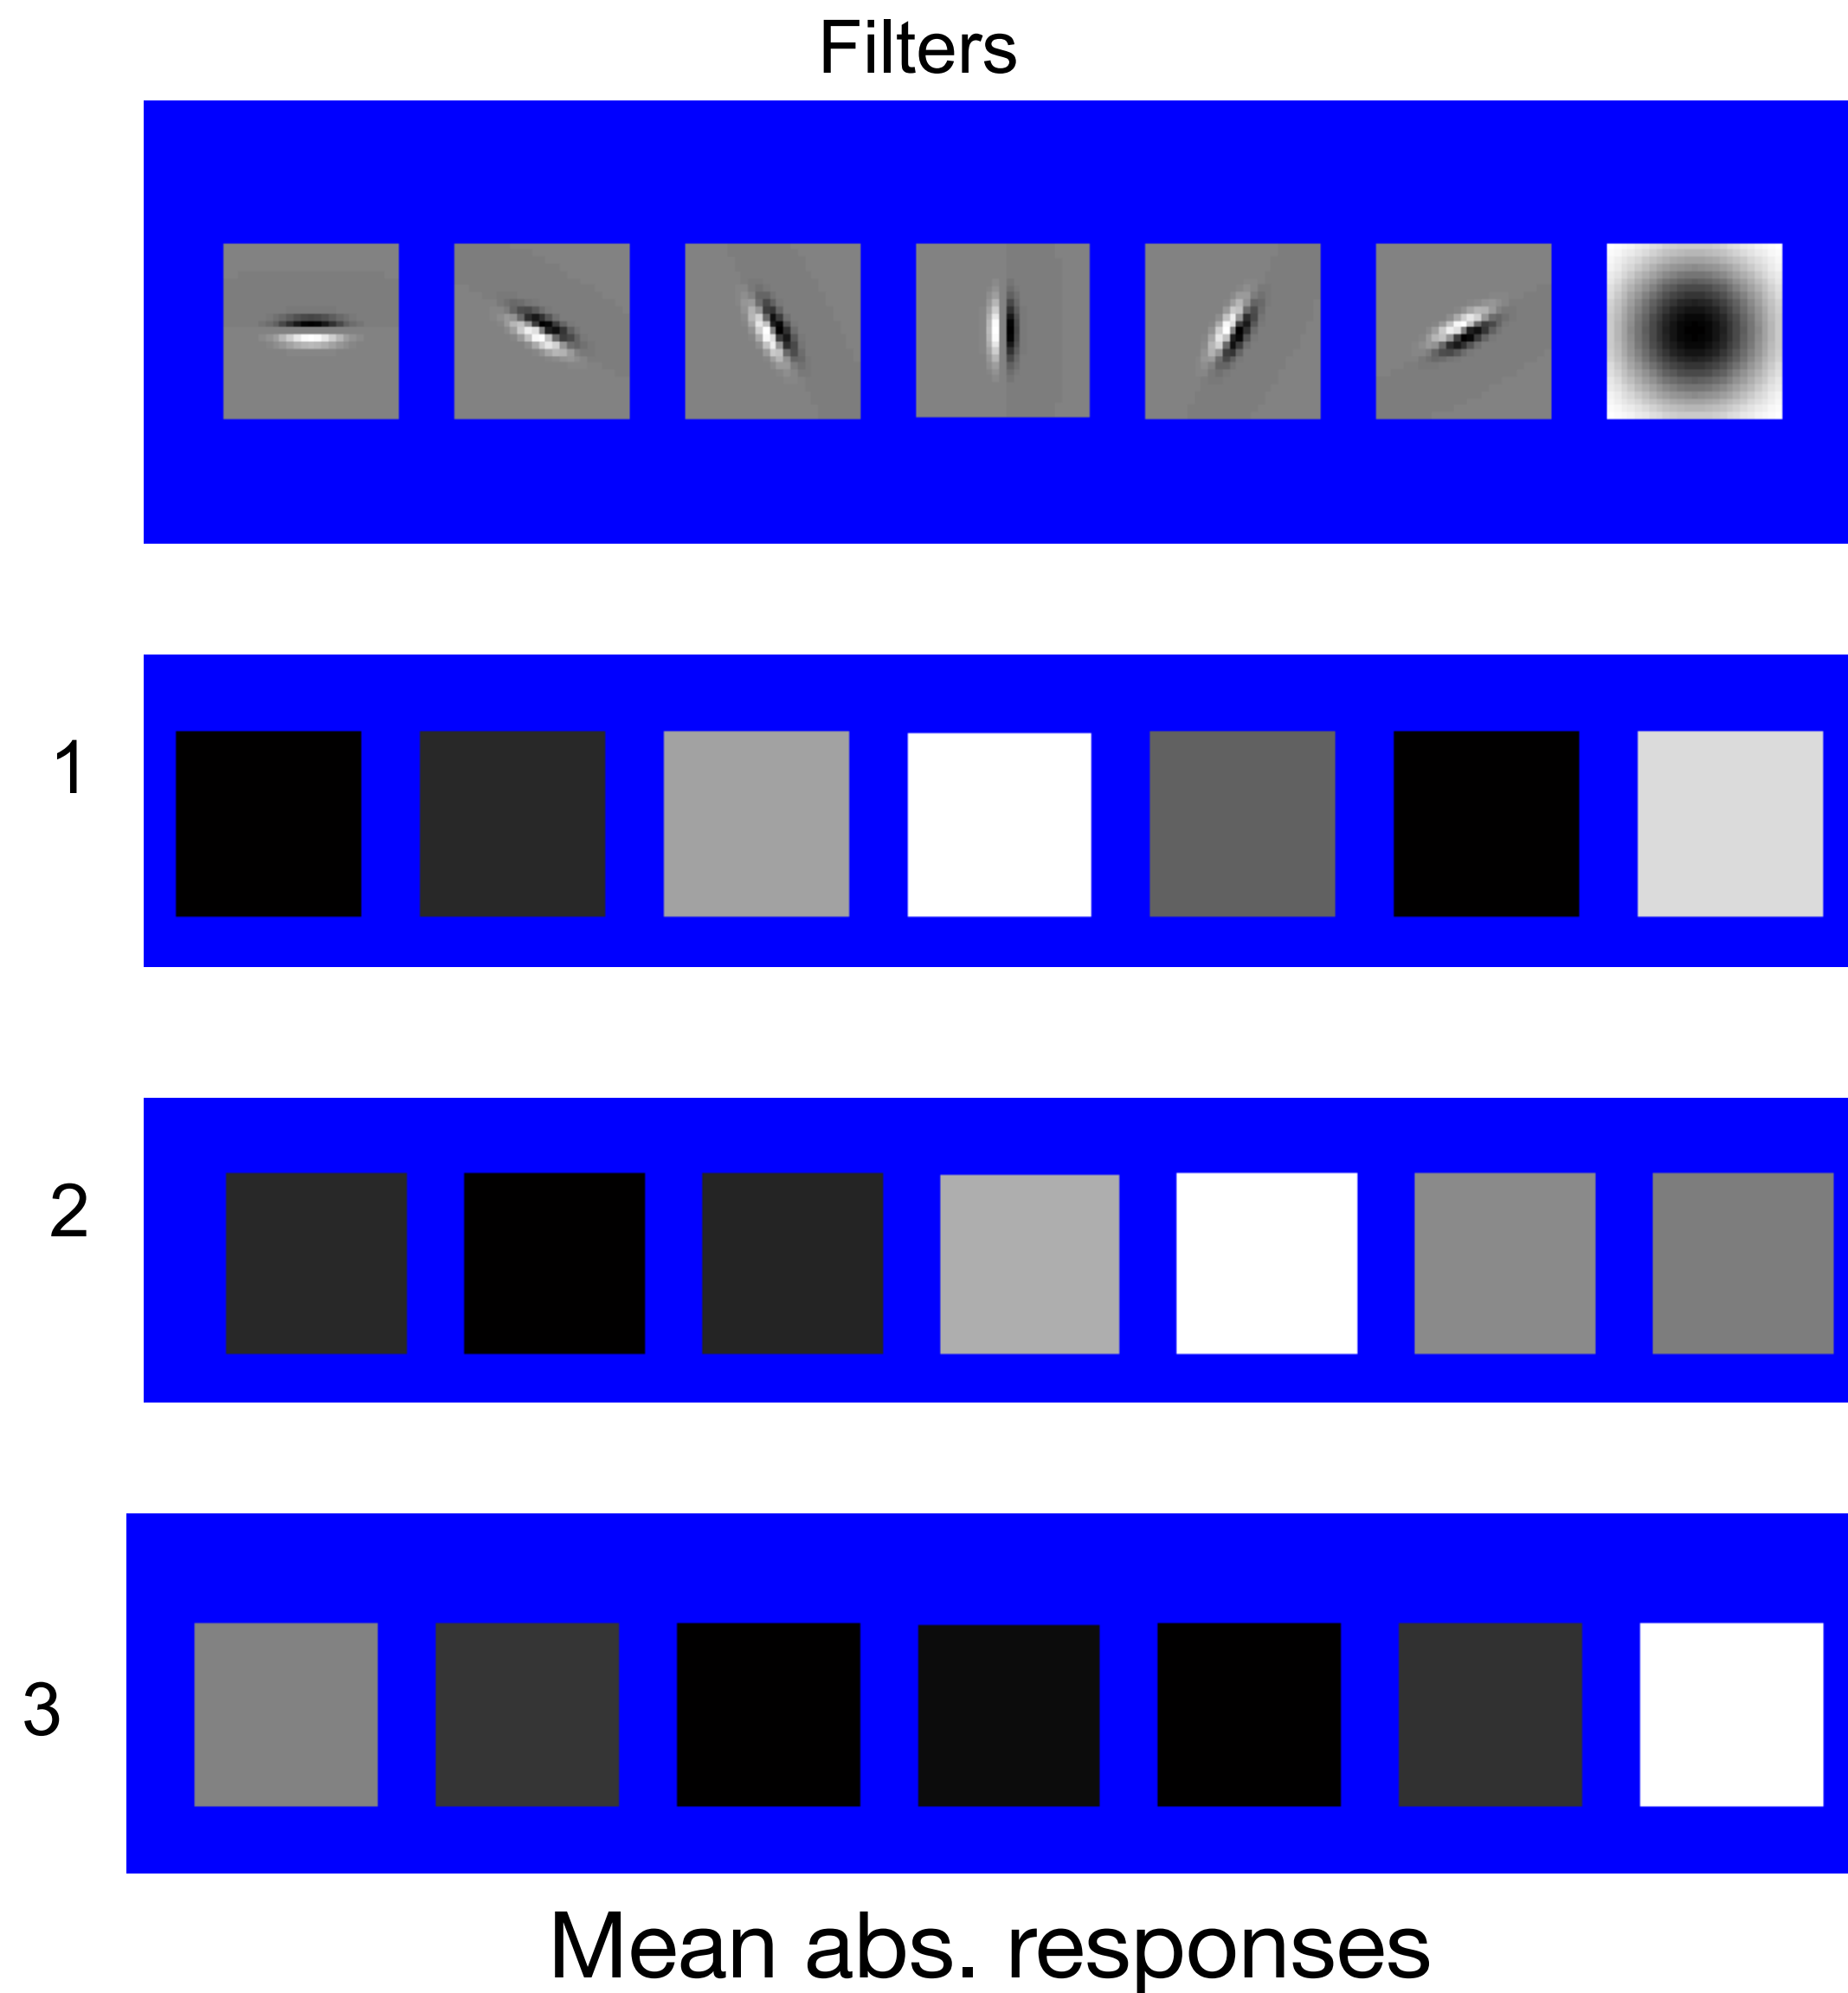
# How can we represent texture in natural images?

- Idea 1: Record simple statistics (e.g., mean, std.) of absolute filter responses





# Can you match the texture to the response?





# How can we represent texture in natural images?

- Generalize this to “orientation histogram”
- Idea 2: Marginal histograms of filter responses
  - One histogram per filter







# Steerable filter decomposition

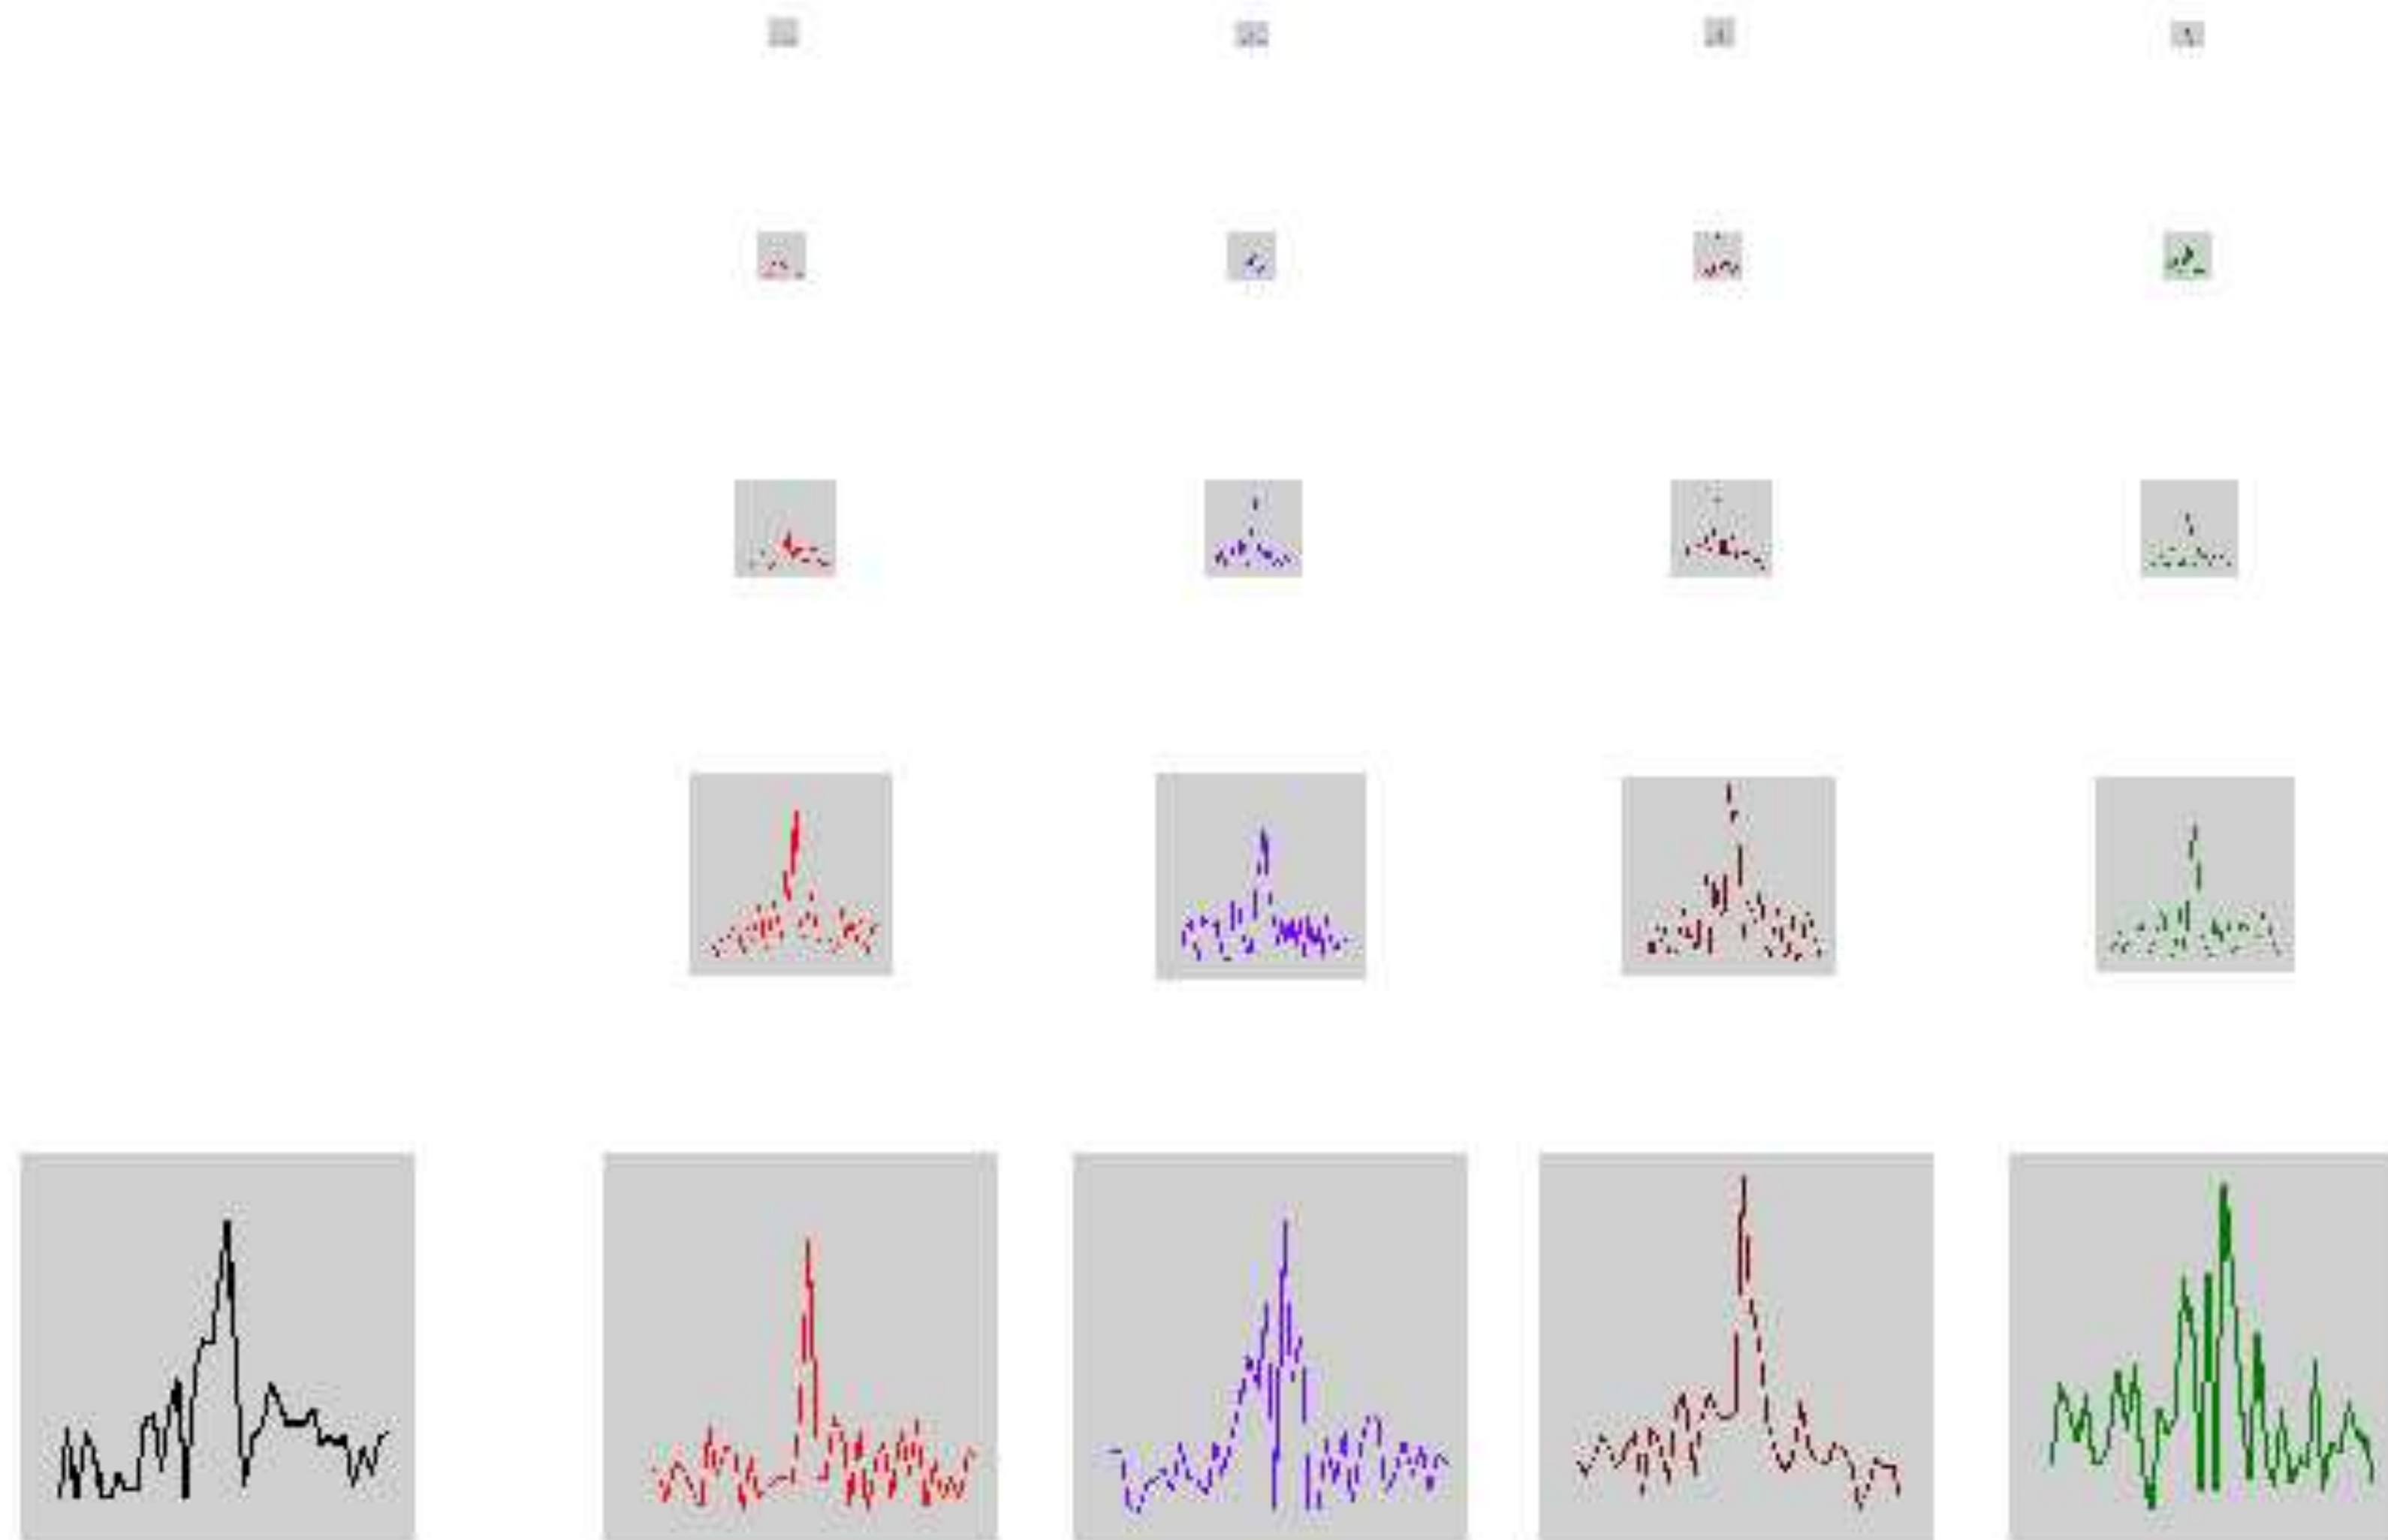
Filter bank



Input image



# Filter response histograms






# Texture synthesis


Start with a noise image as output.

## Main loop:

- Match pixel histogram of output image to input
- Decompose input/output images using a Steerable Pyramid
- Match subband histograms of input and output pyramids
- Reconstruct input and output images (collapse the pyramids)







Source: A. Efros










Figure 7: (Left pair) Inhomogeneous input texture produces blotchy synthetic texture. (Right pair) Homogeneous input.




Figure 8: Examples of failures: wood grain and red coral.



Figure 9: More failures: hay and marble.



