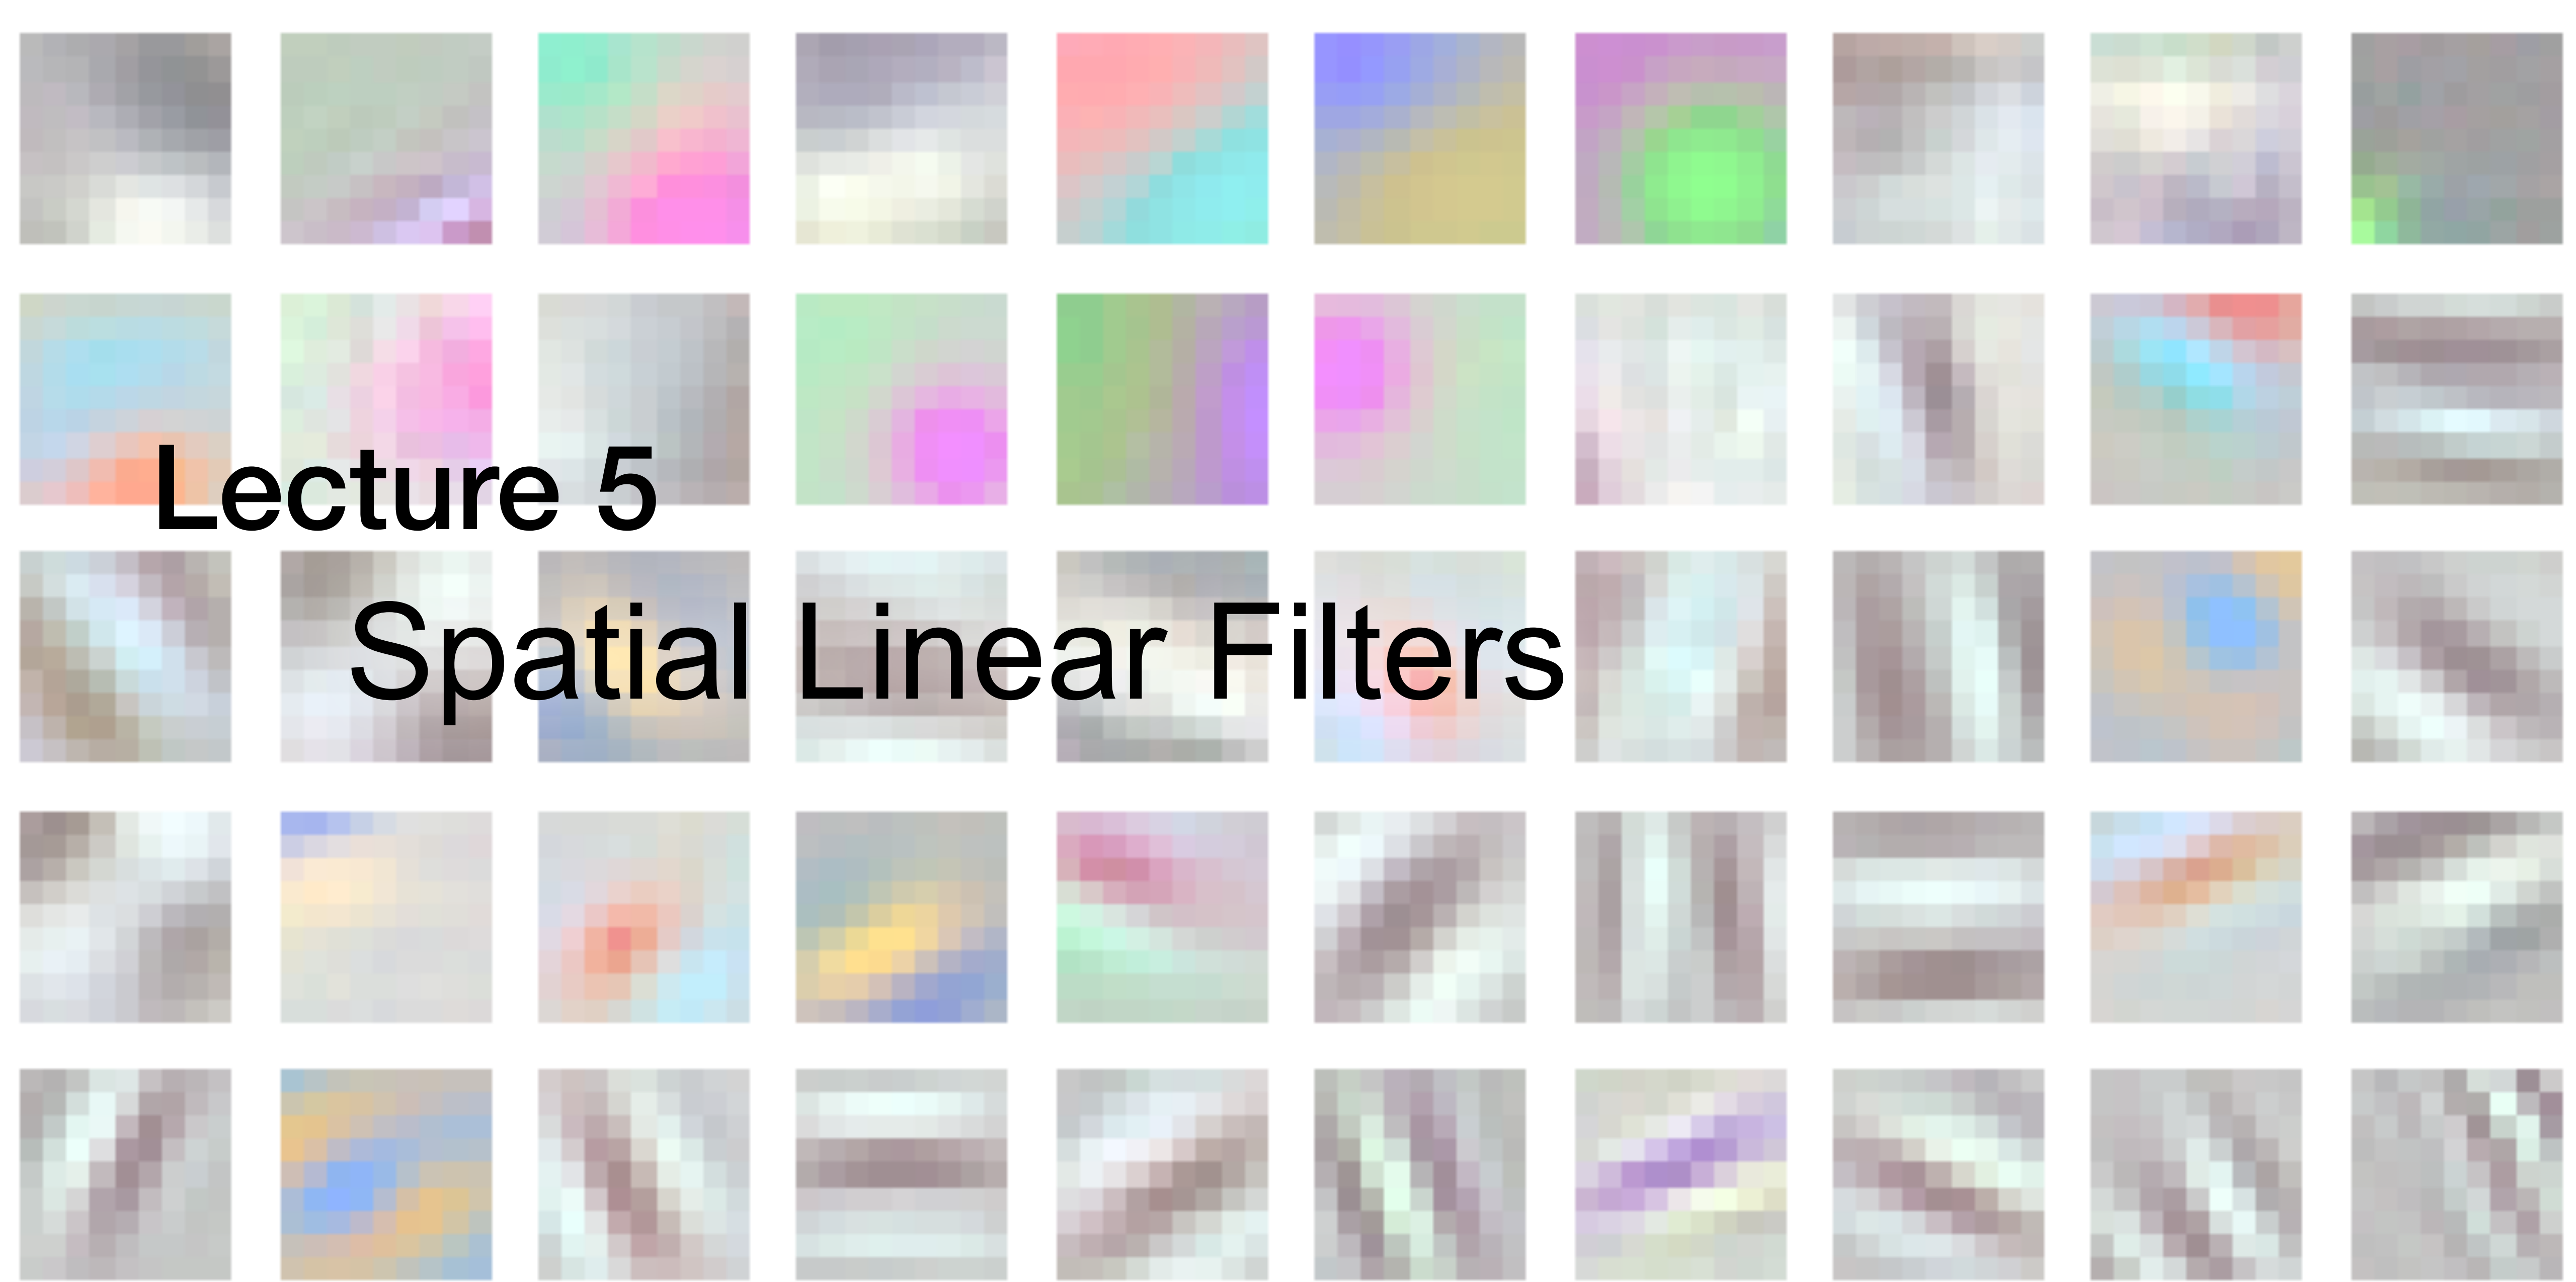


Lecture 5

Spatial Linear Filters



Today's lecture: Spatial Filtering

Problem set 1 due today. Problem set 2 out today (due Thursday of next week).

This week: Tuesday: spatial filtering. Thursday: temporal filtering

Next week: no class Tuesday, then on Thursday: image pyramids.

Today:

- Fourier transforms and signal processing, continued
- Fourier processing by human visual system
- Spatial digital filters

The Discrete Fourier transform

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

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

The inverse Discrete Fourier transform

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

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

The inverse of the 2D DFT is:

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

Visualizing the image Fourier transform

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

The values of $F[u, v]$ are complex.

Using the real and imaginary components:

$$F[u, v] = \text{Re} \{F[u, v]\} + j \text{Imag} \{F[u, v]\}$$

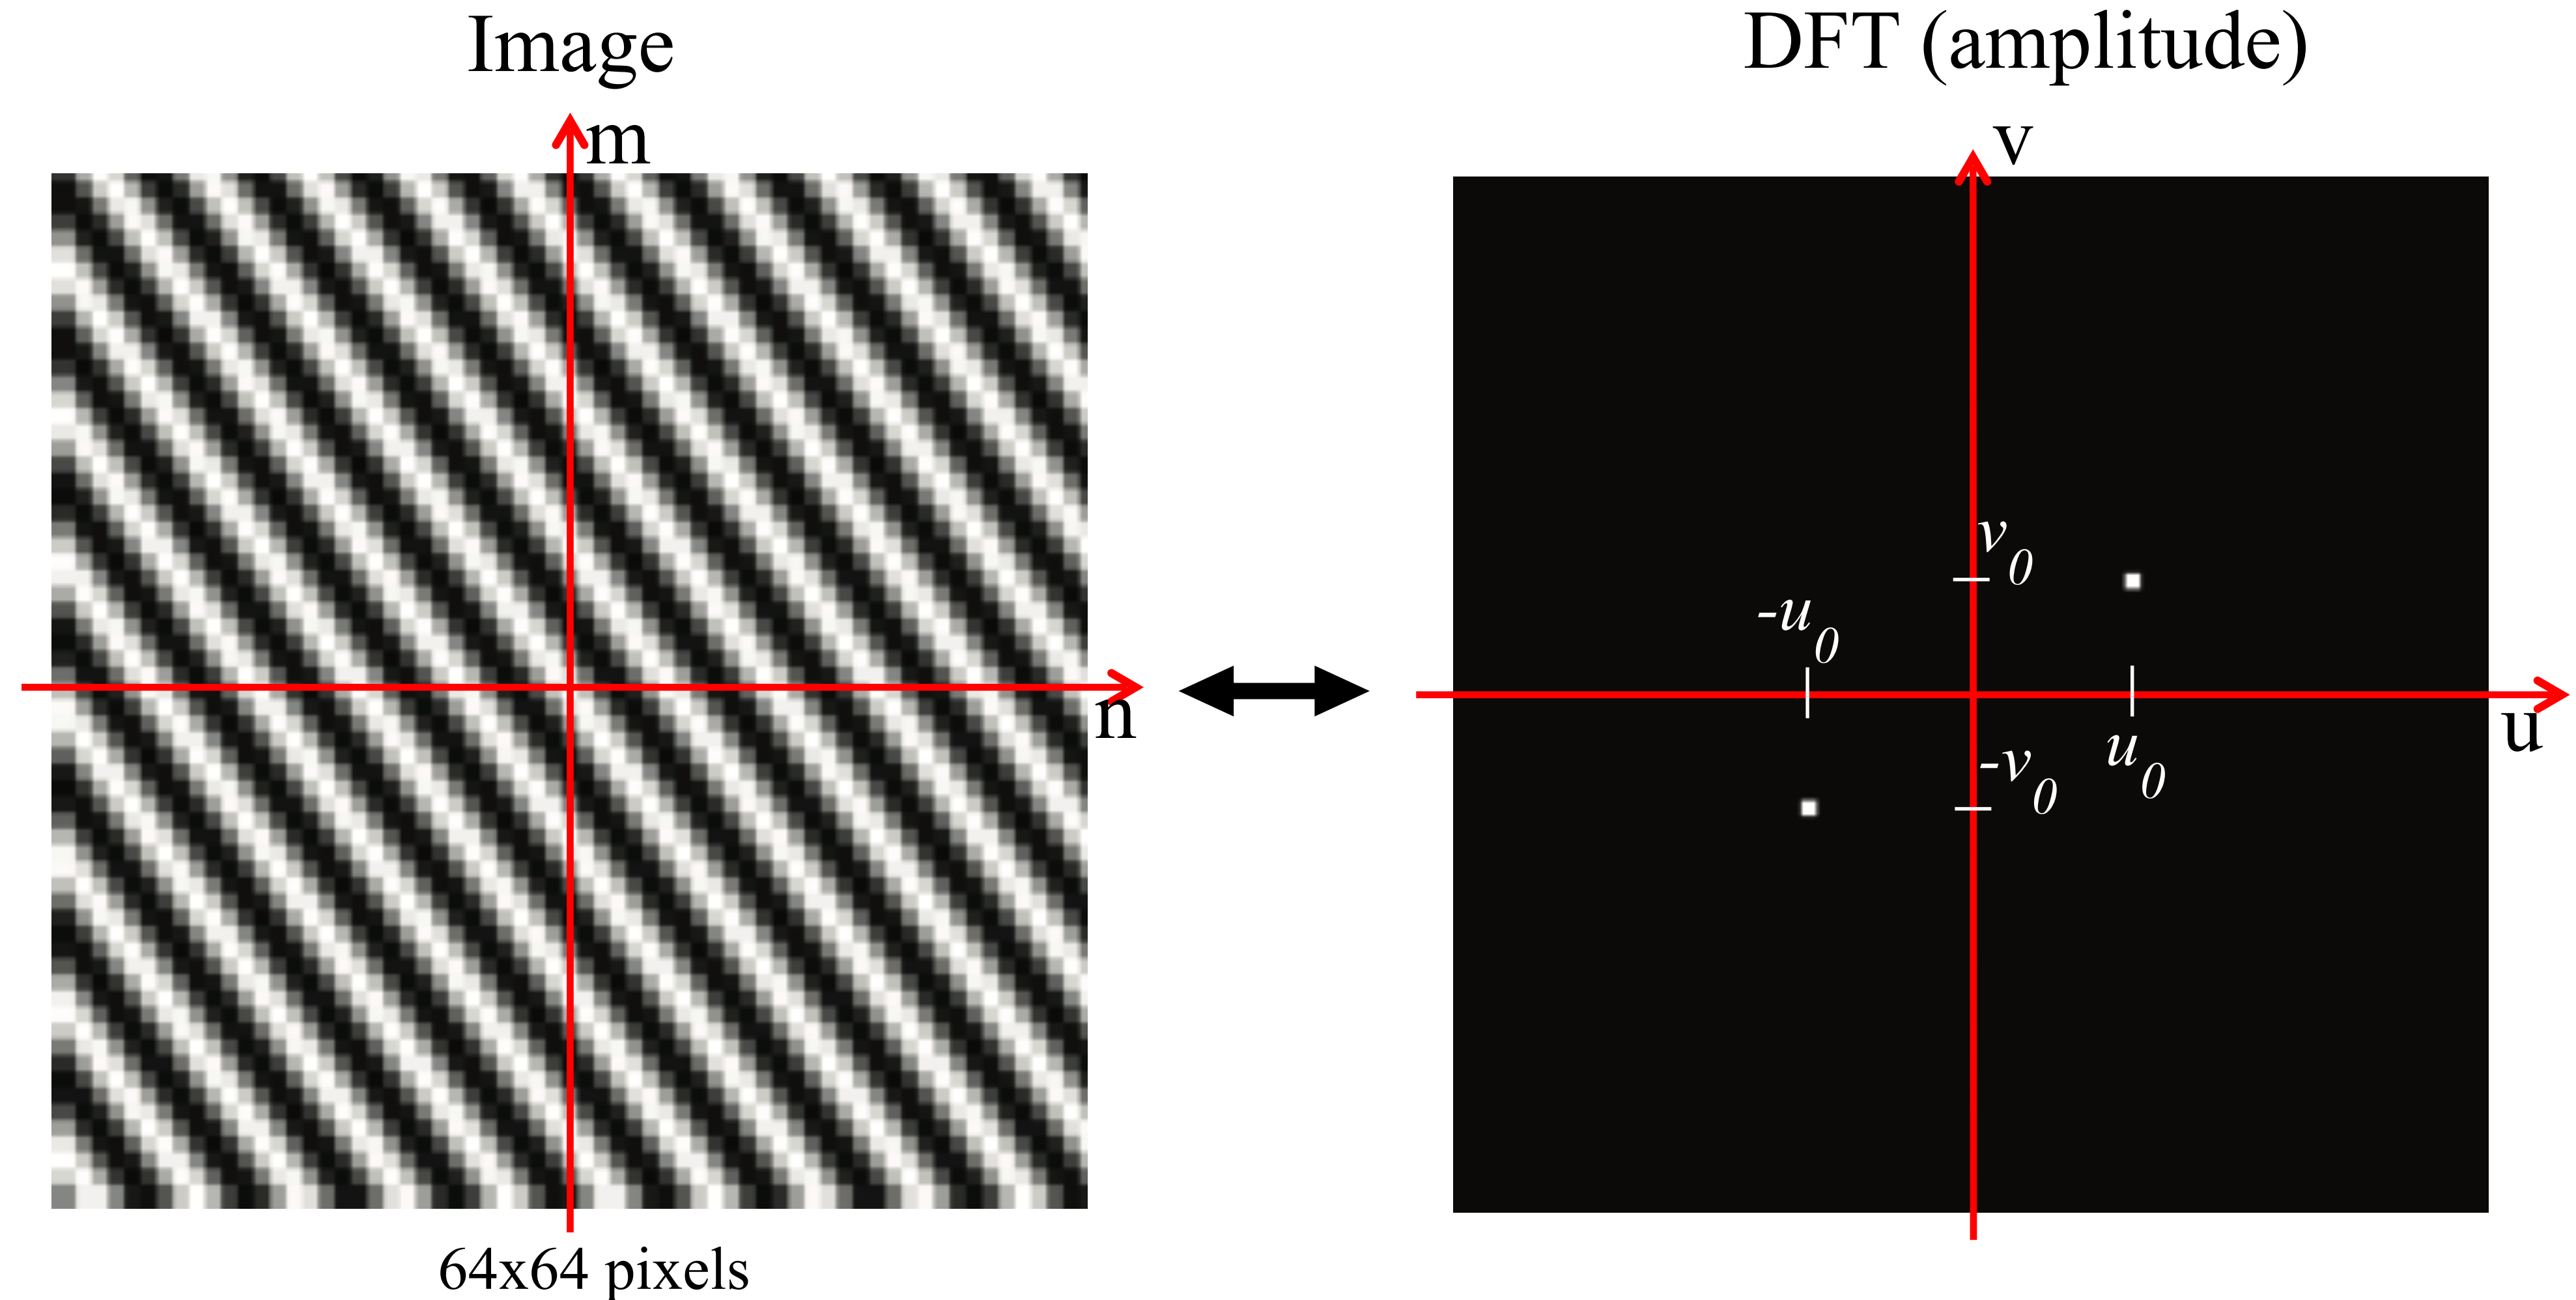
Or using a polar decomposition:

$$F[u, v] = A[u, v] \exp(j\theta[u, v])$$

Amplitude

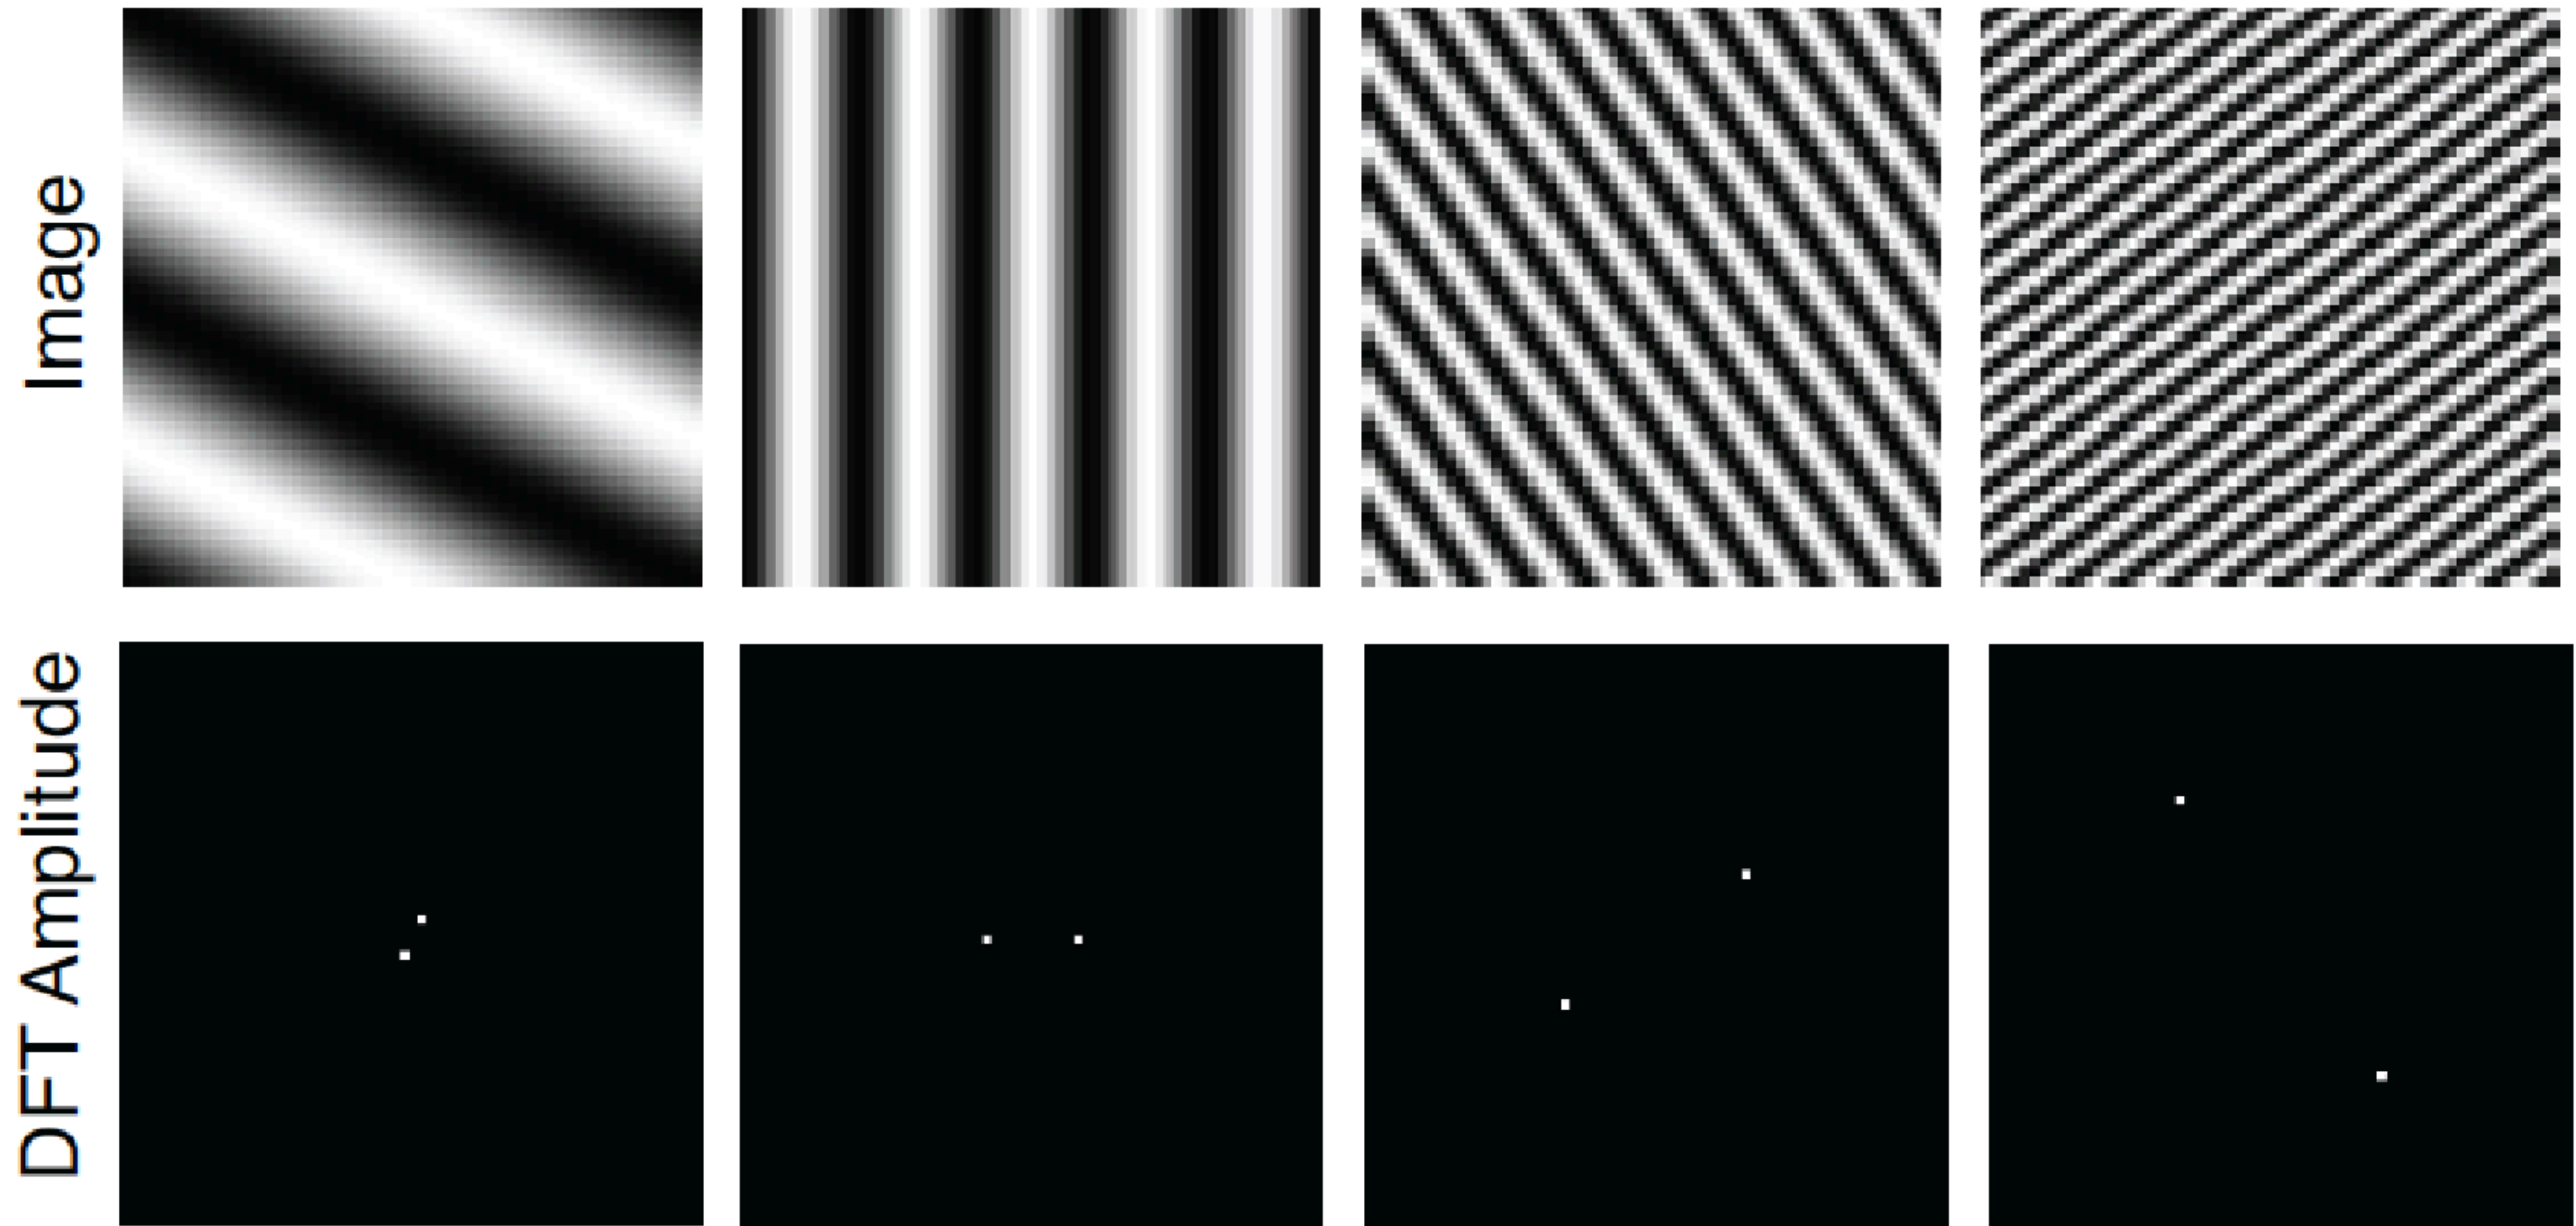
Phase

Simple Fourier transforms



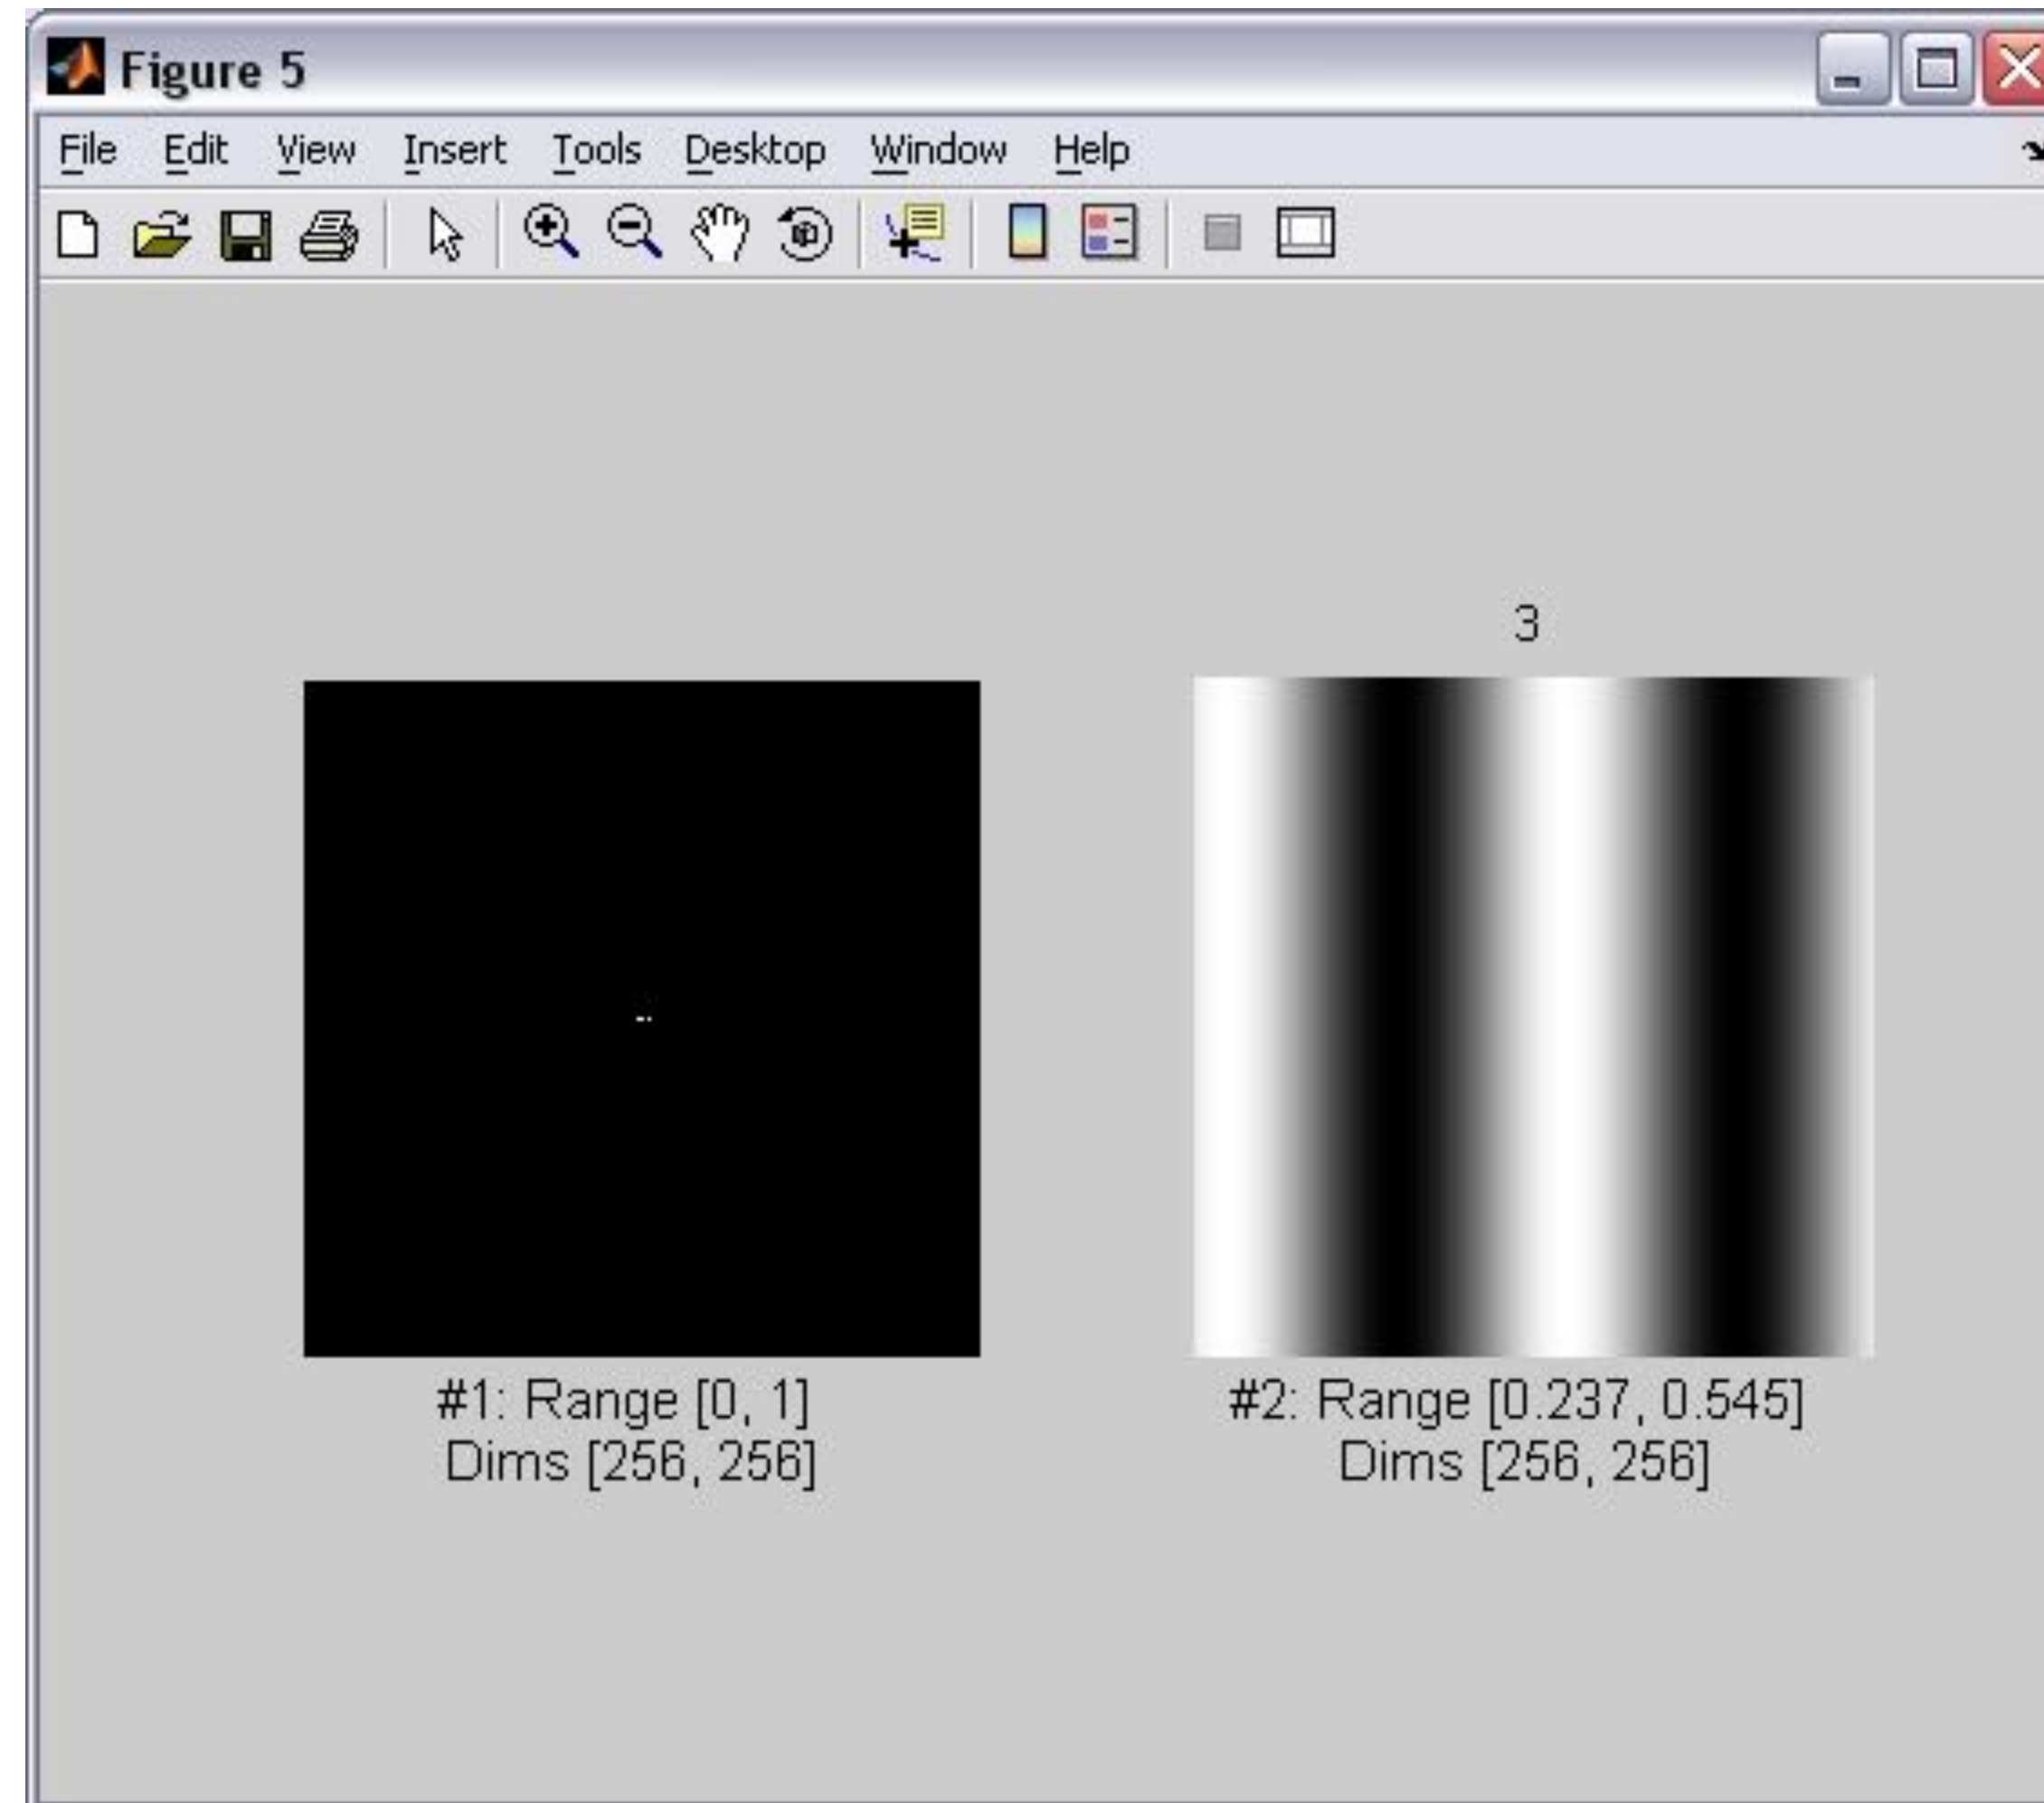
$$\cos \left(2\pi \left(\frac{u_0 n}{N} + \frac{v_0 m}{M} \right) \right) \longleftrightarrow \frac{1}{2} (\delta [u - u_0, v - v_0] + \delta [u + u_0, v + v_0])$$

Simple Fourier transforms



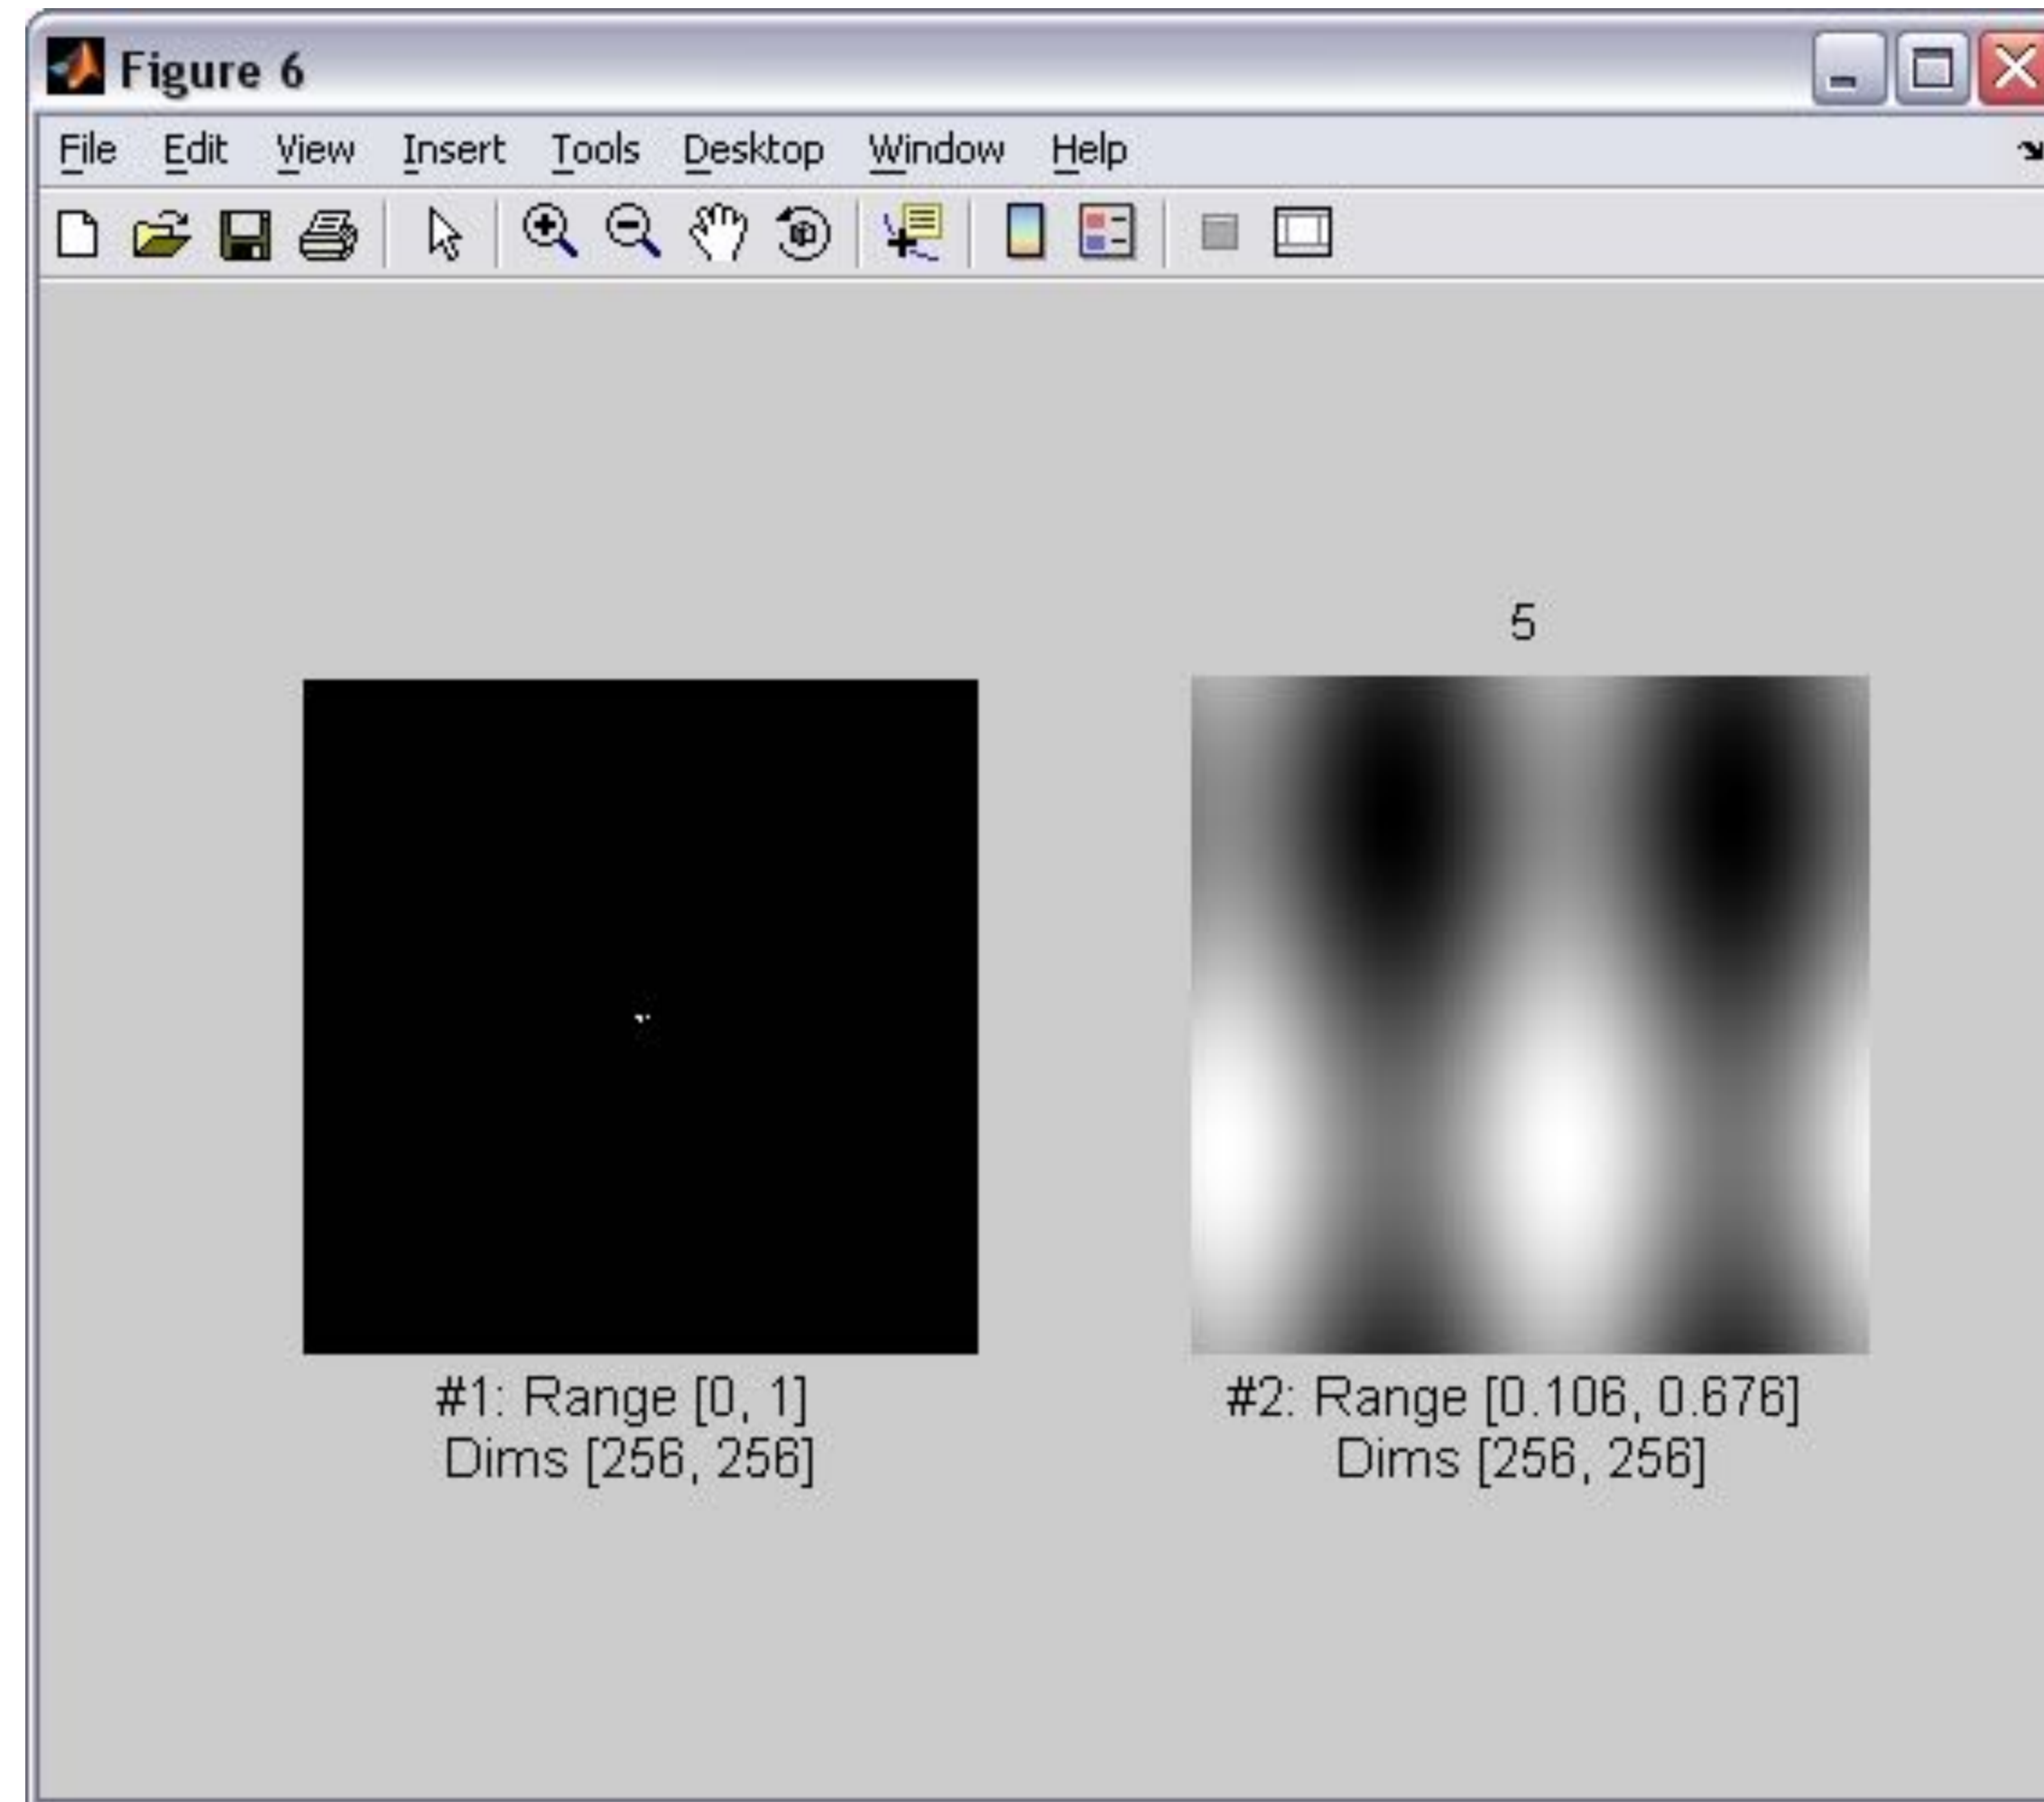
Images are 64x64 pixels. The wave is a cosine, therefore DFT phase is zero.

3

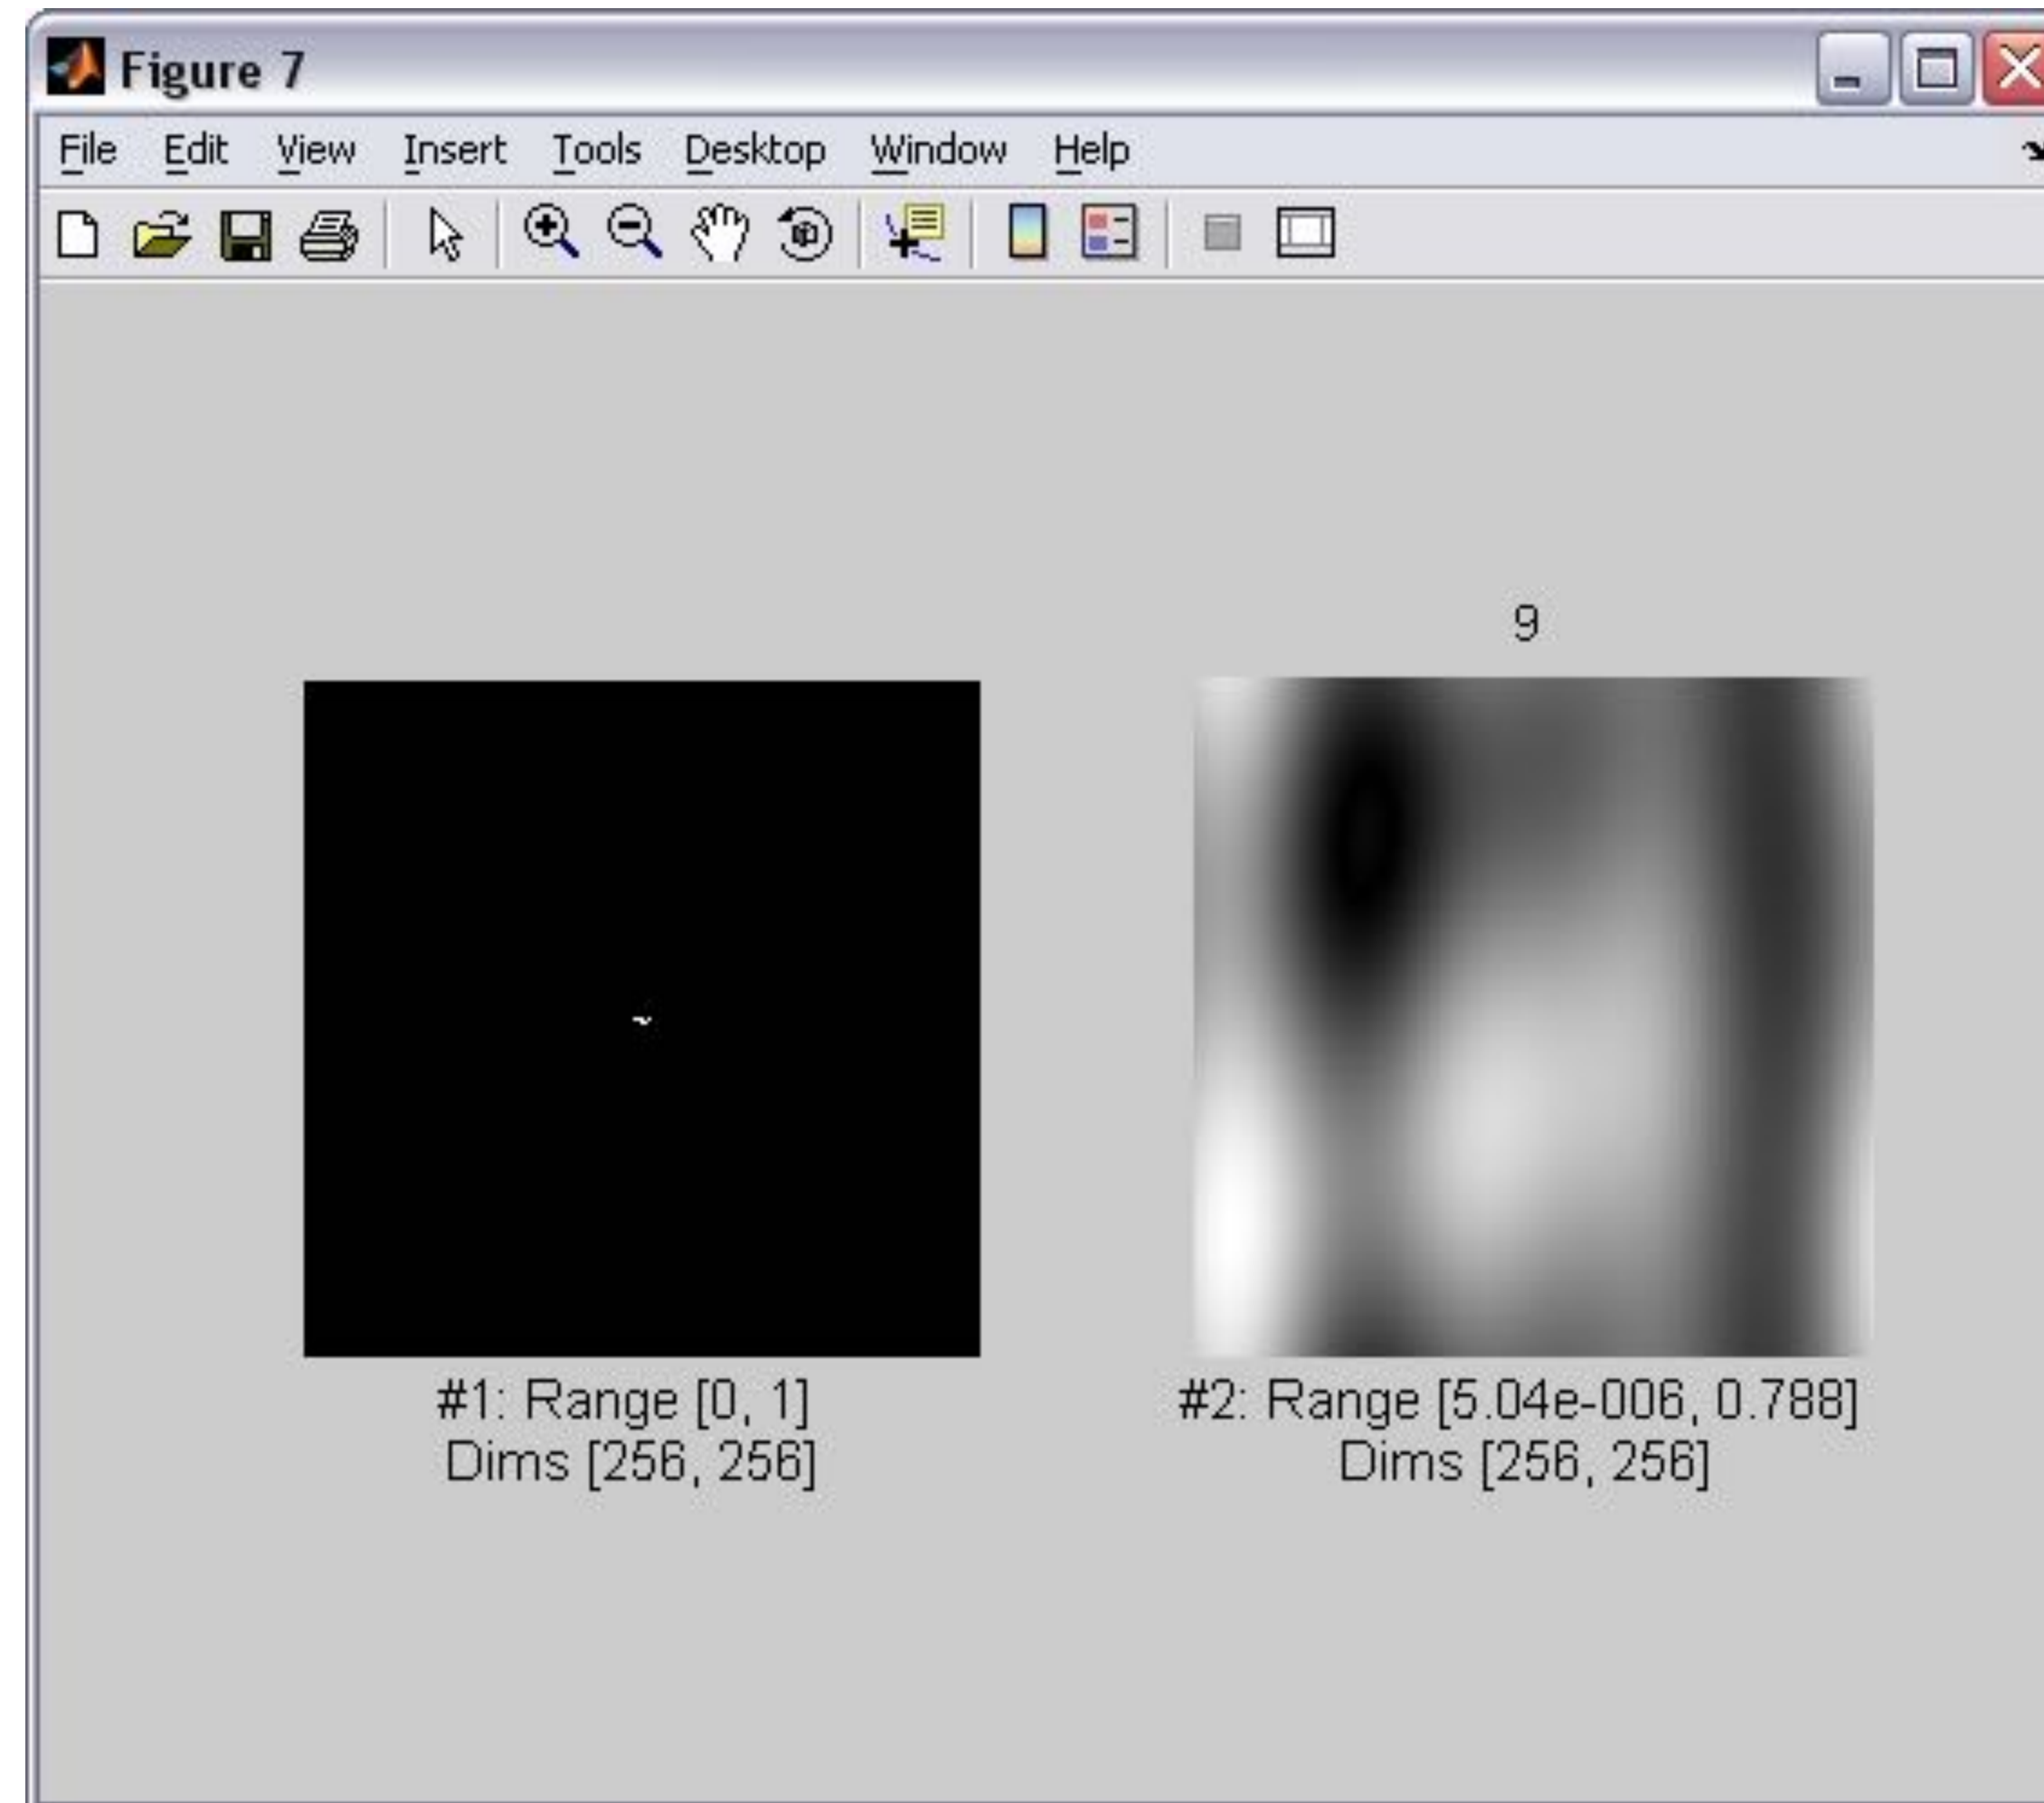


Now, an analogous sequence of images, but selecting Fourier components in descending order of magnitude.

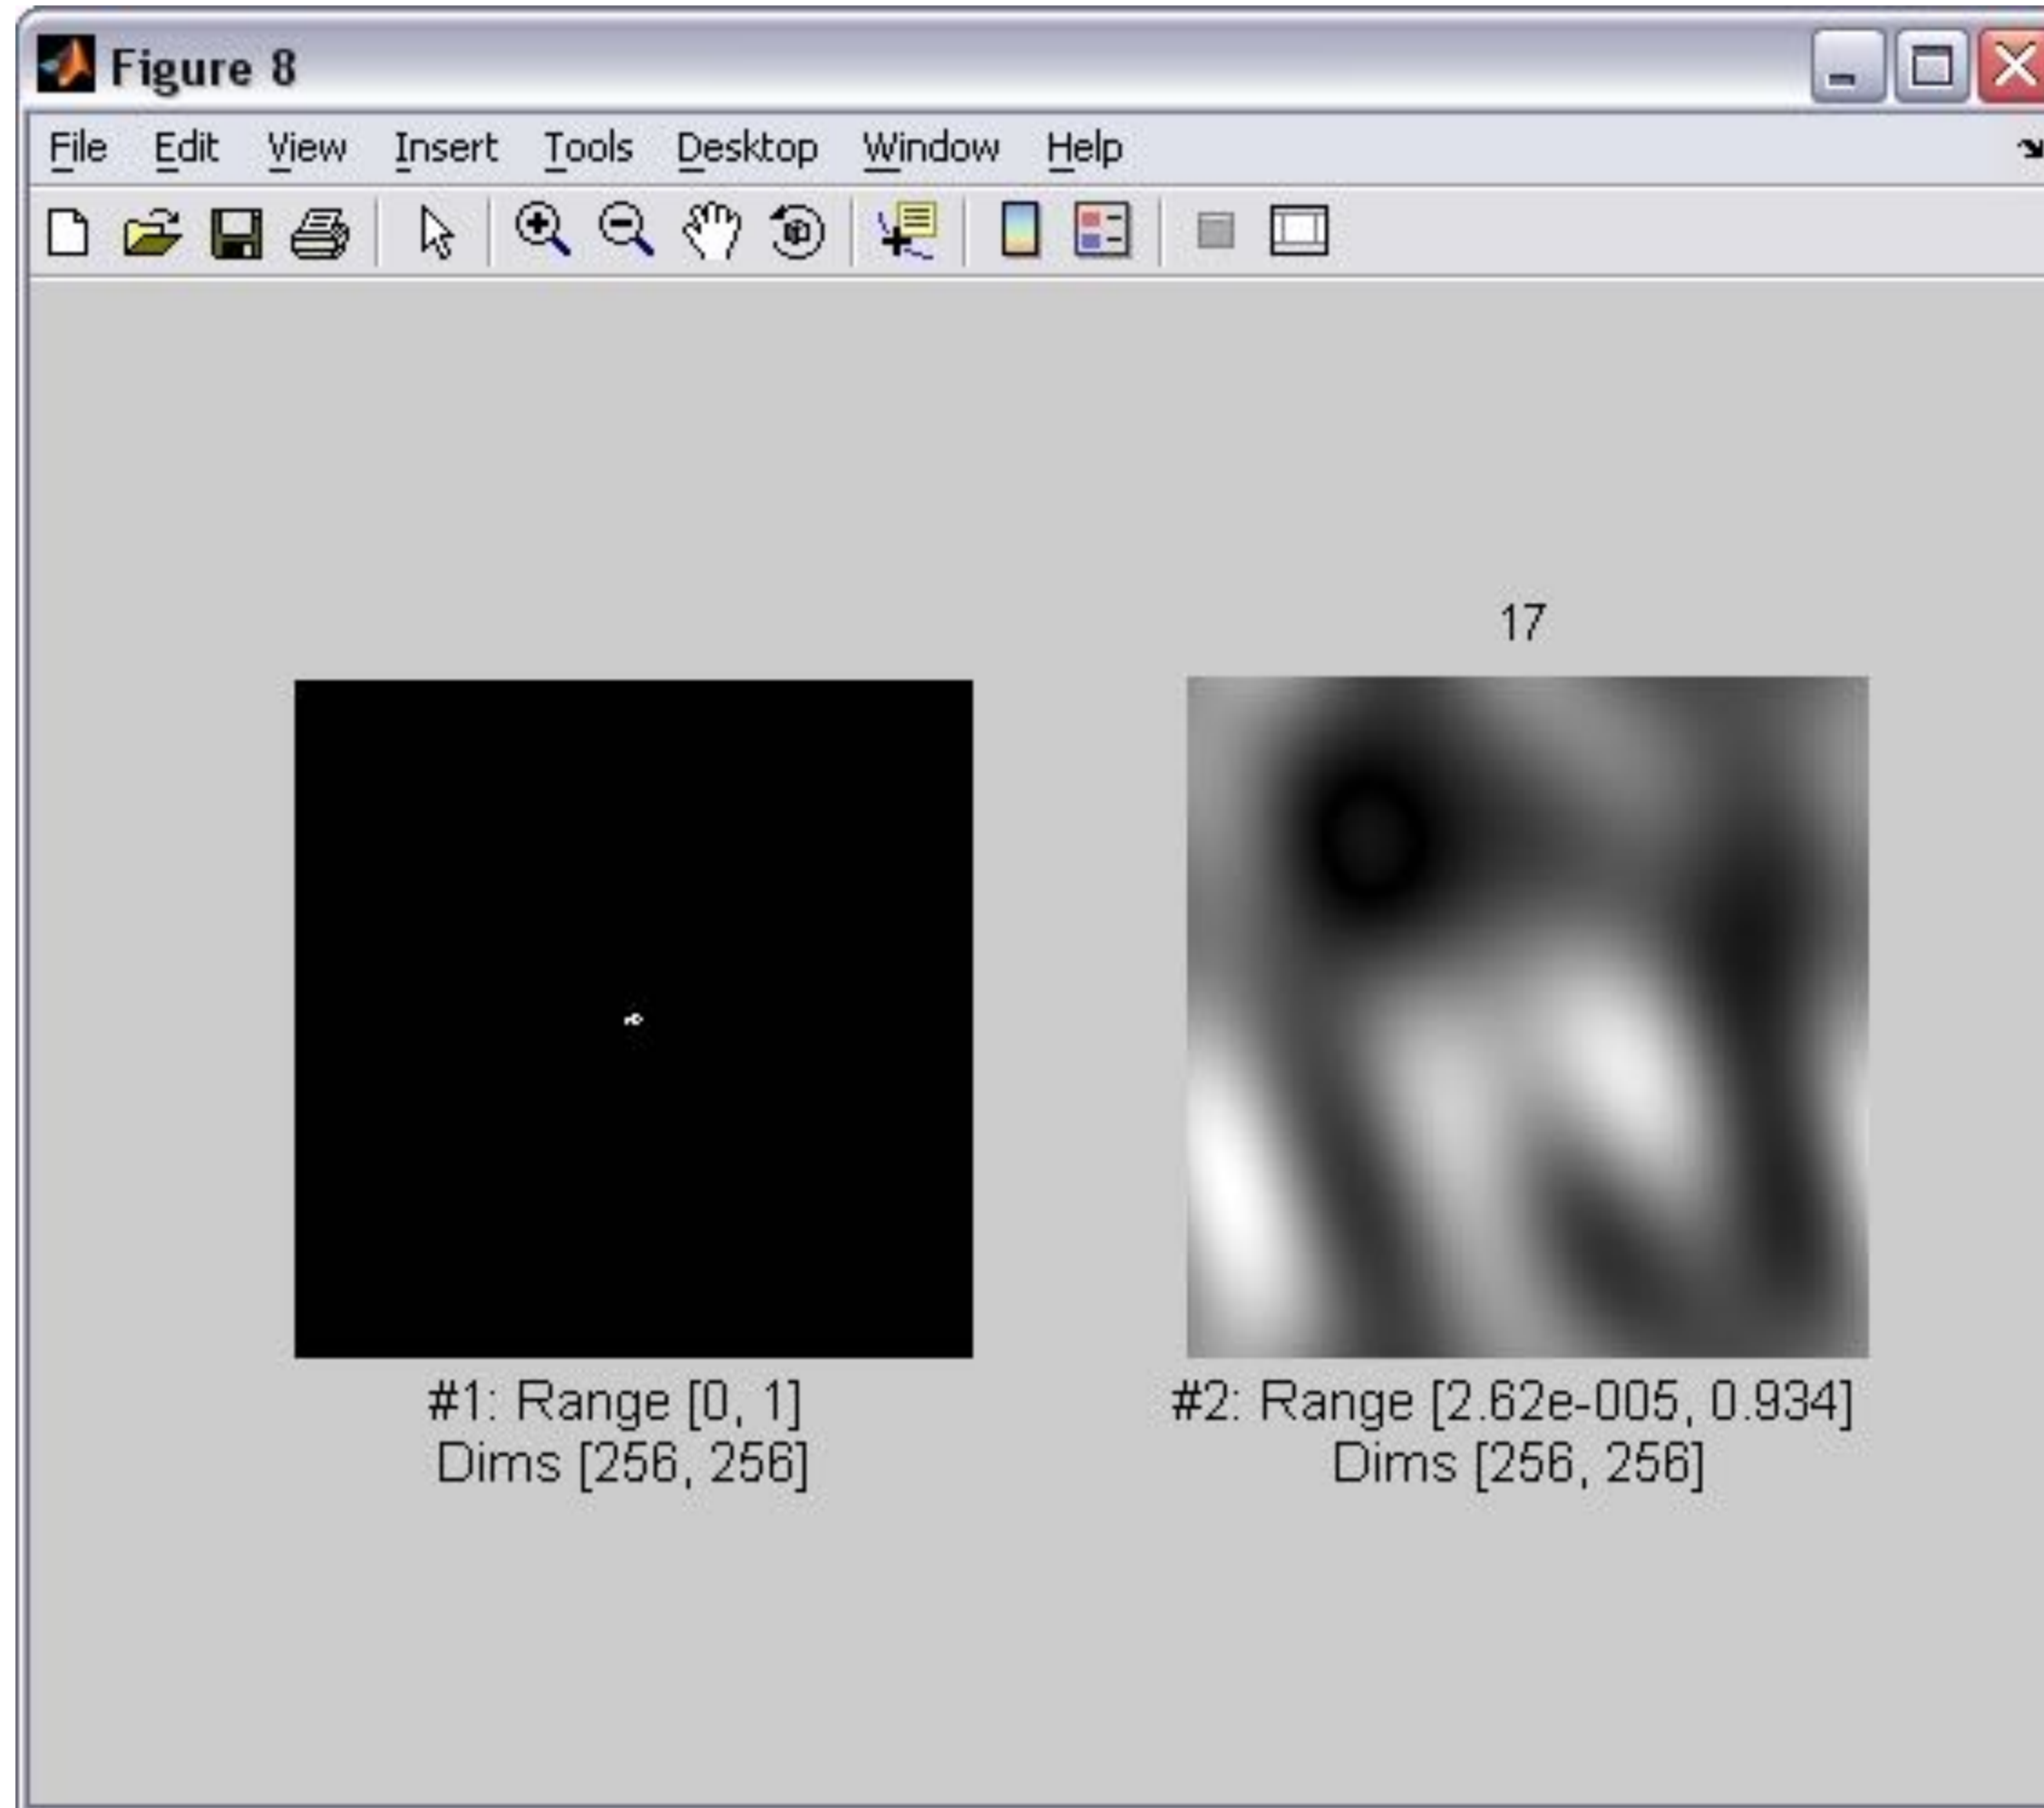
5



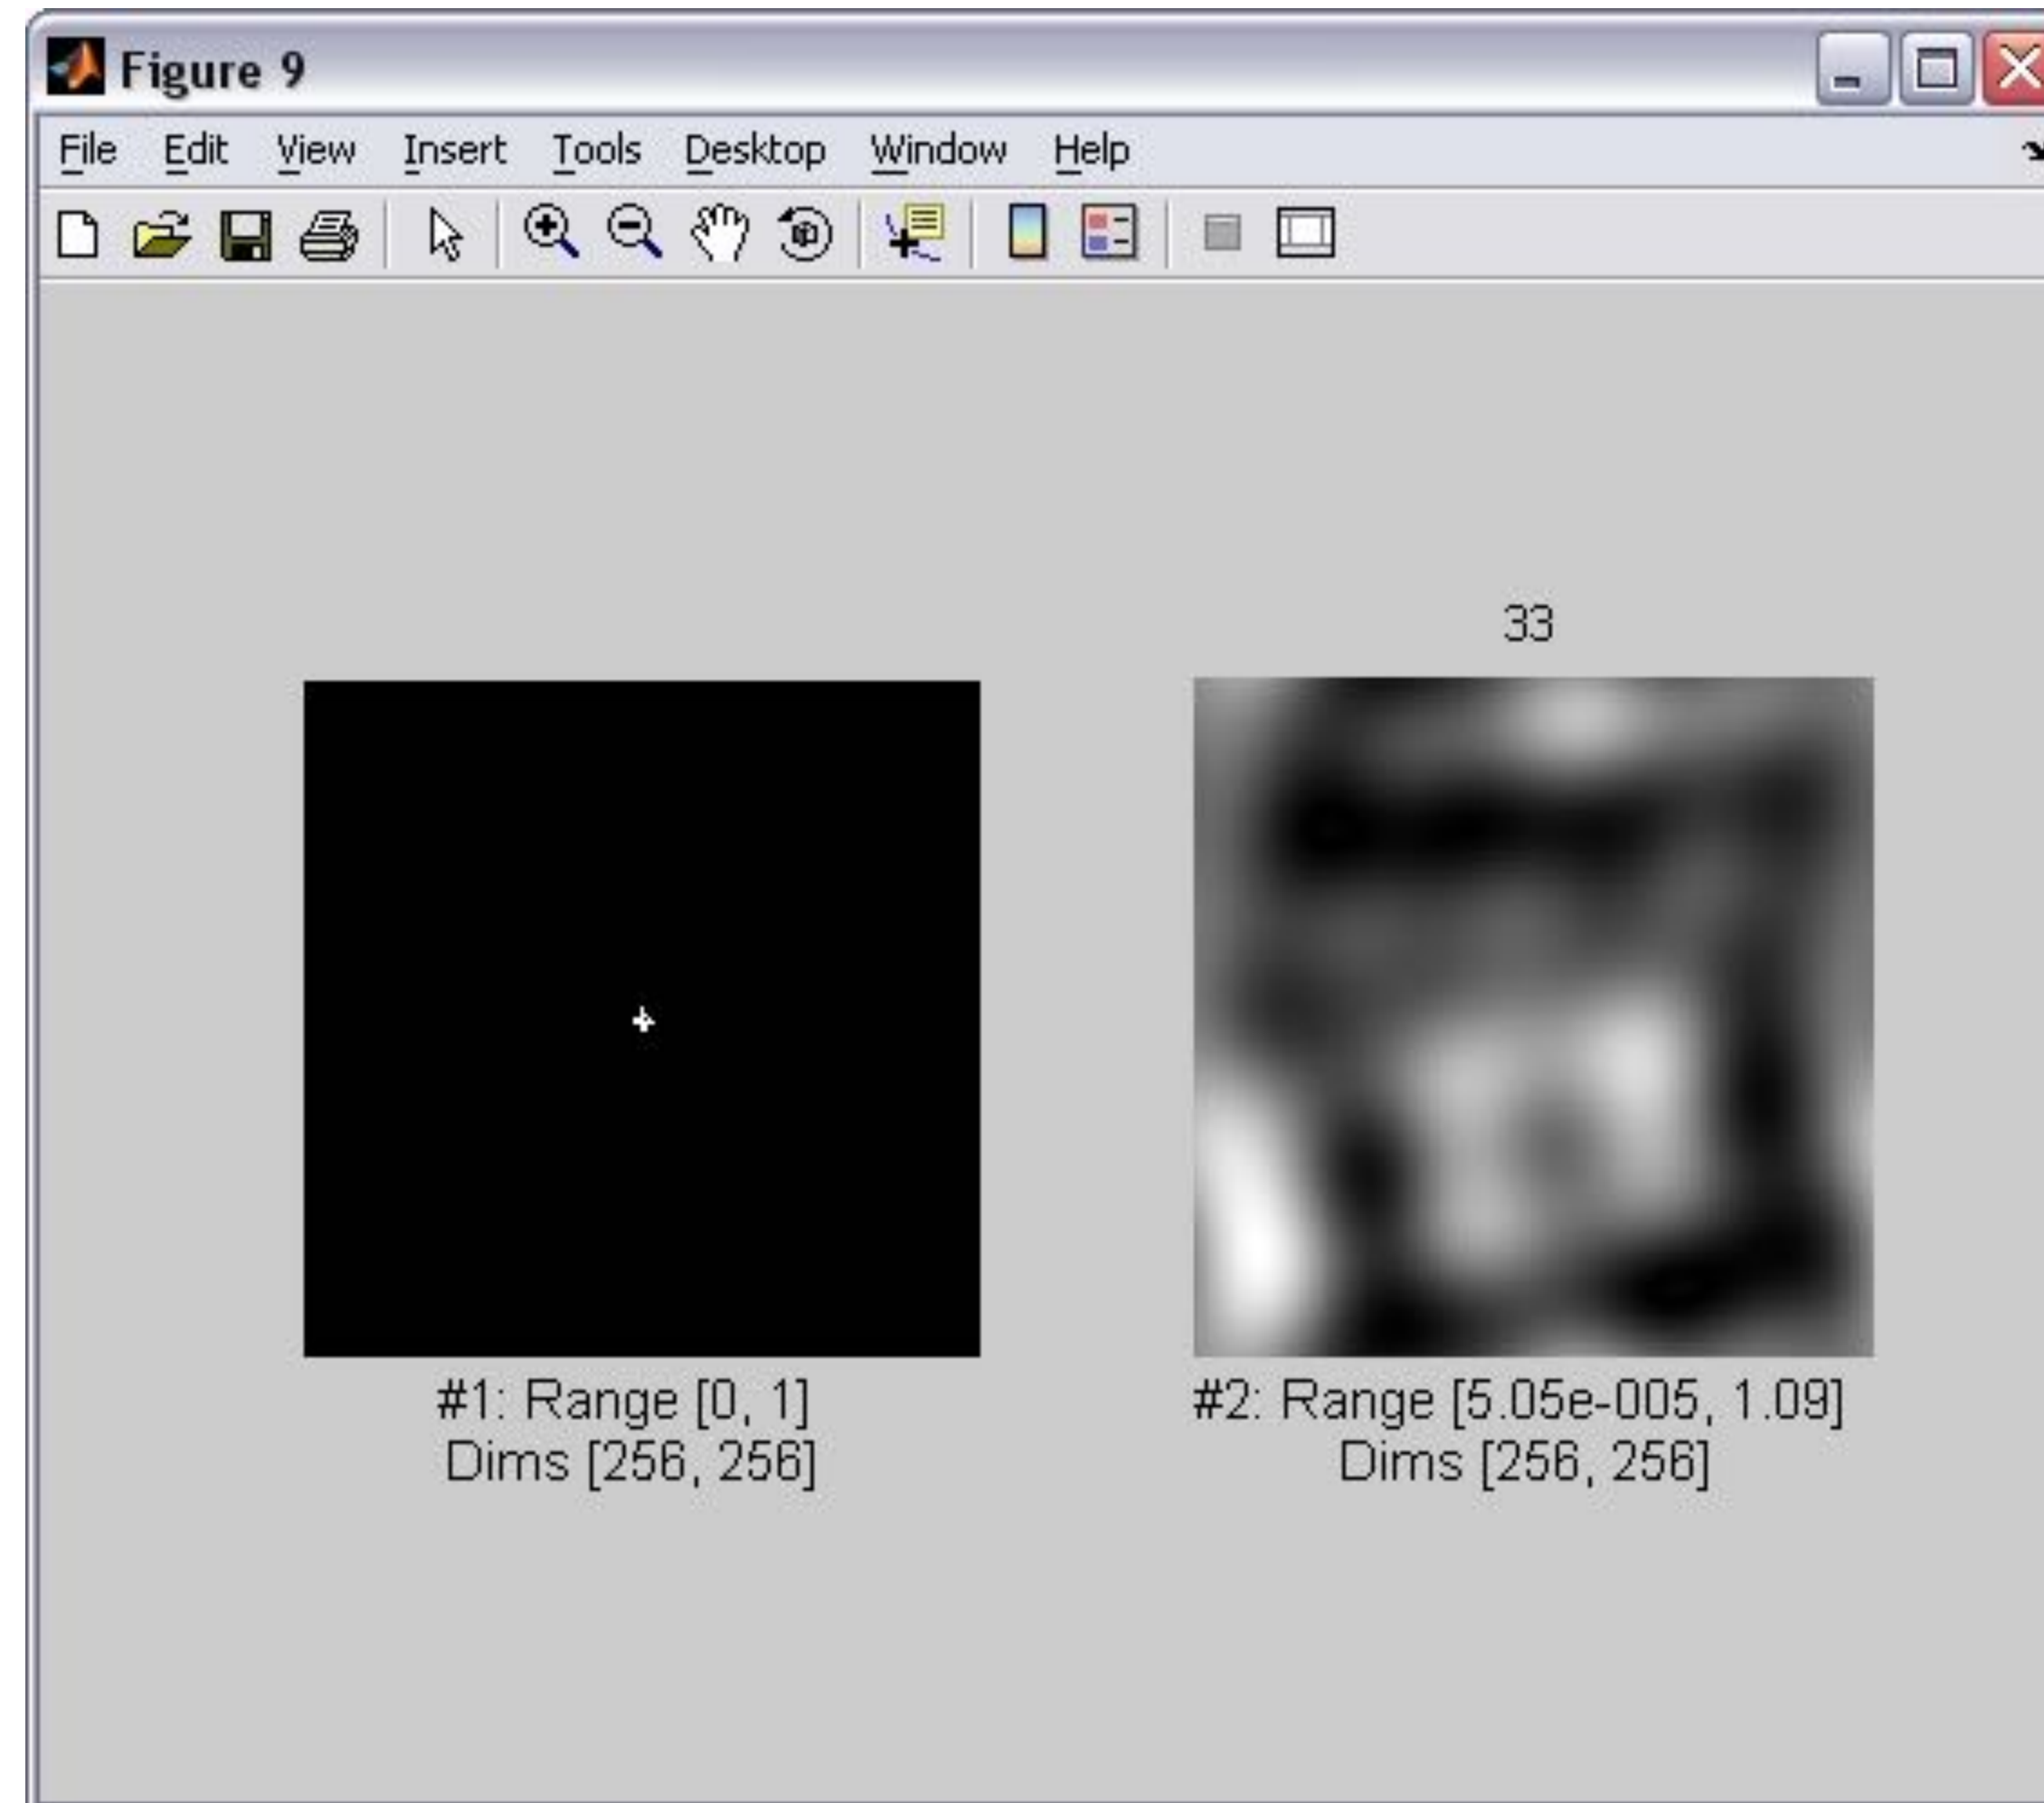
9



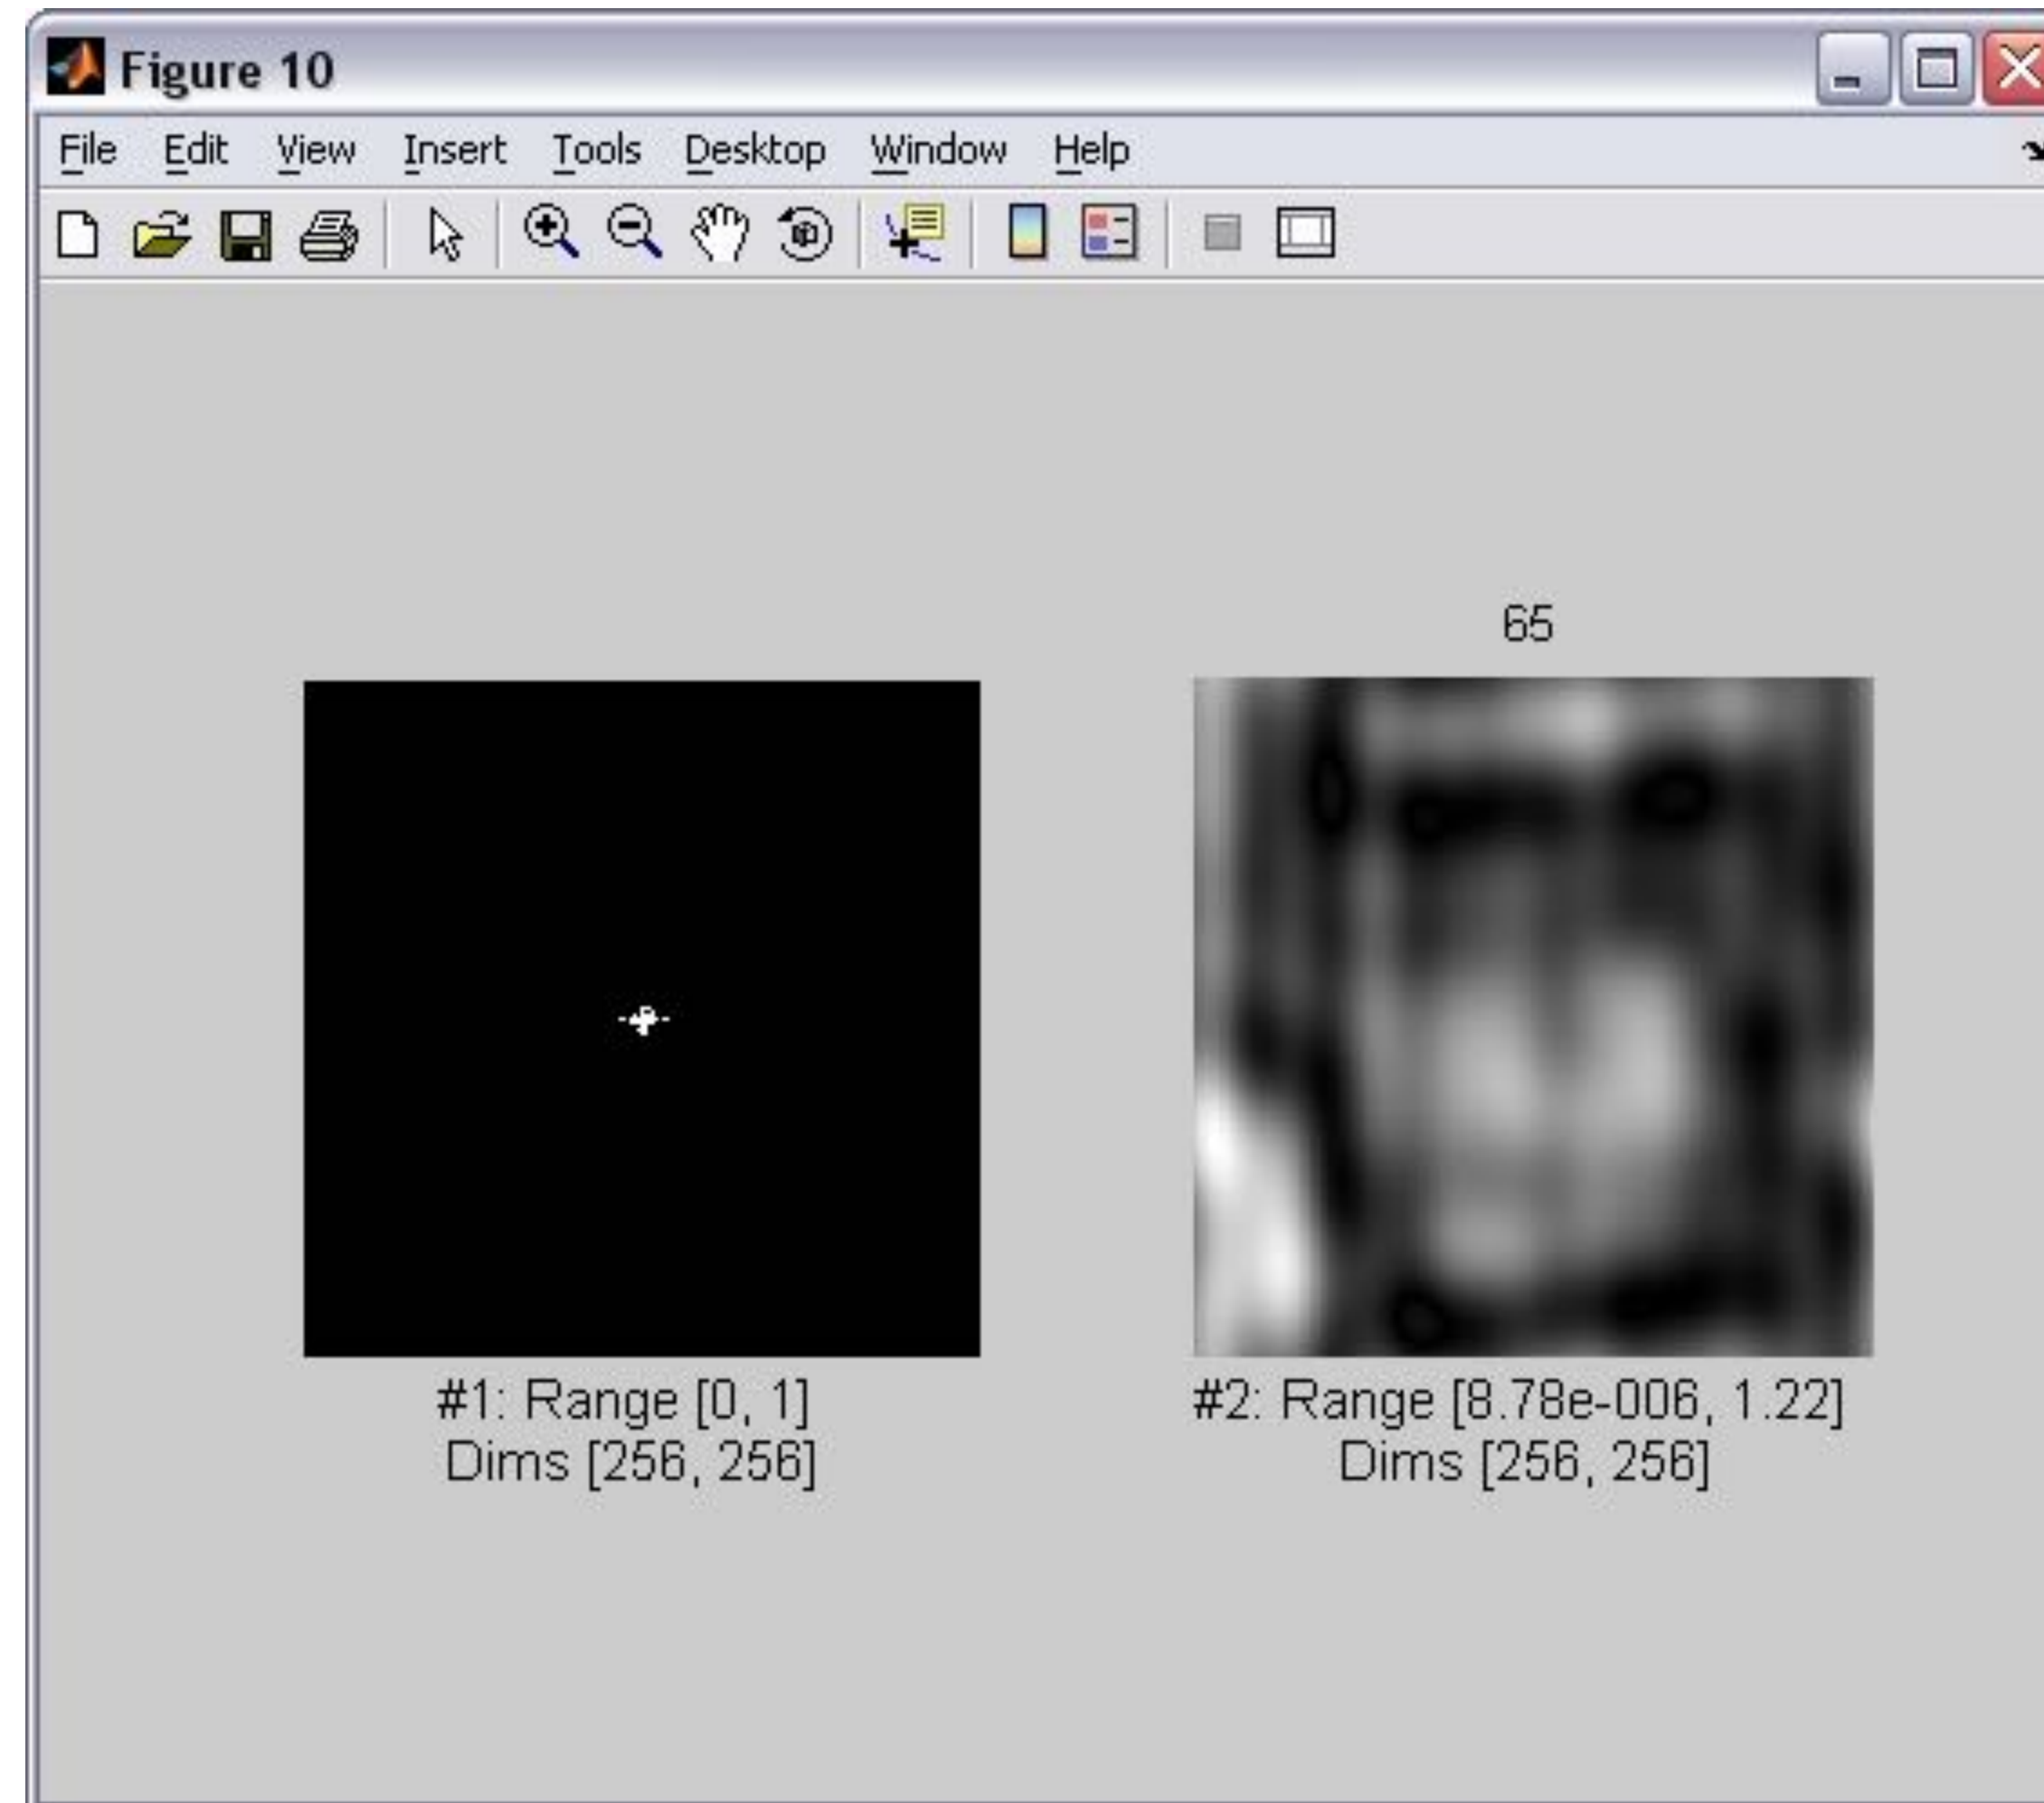
17



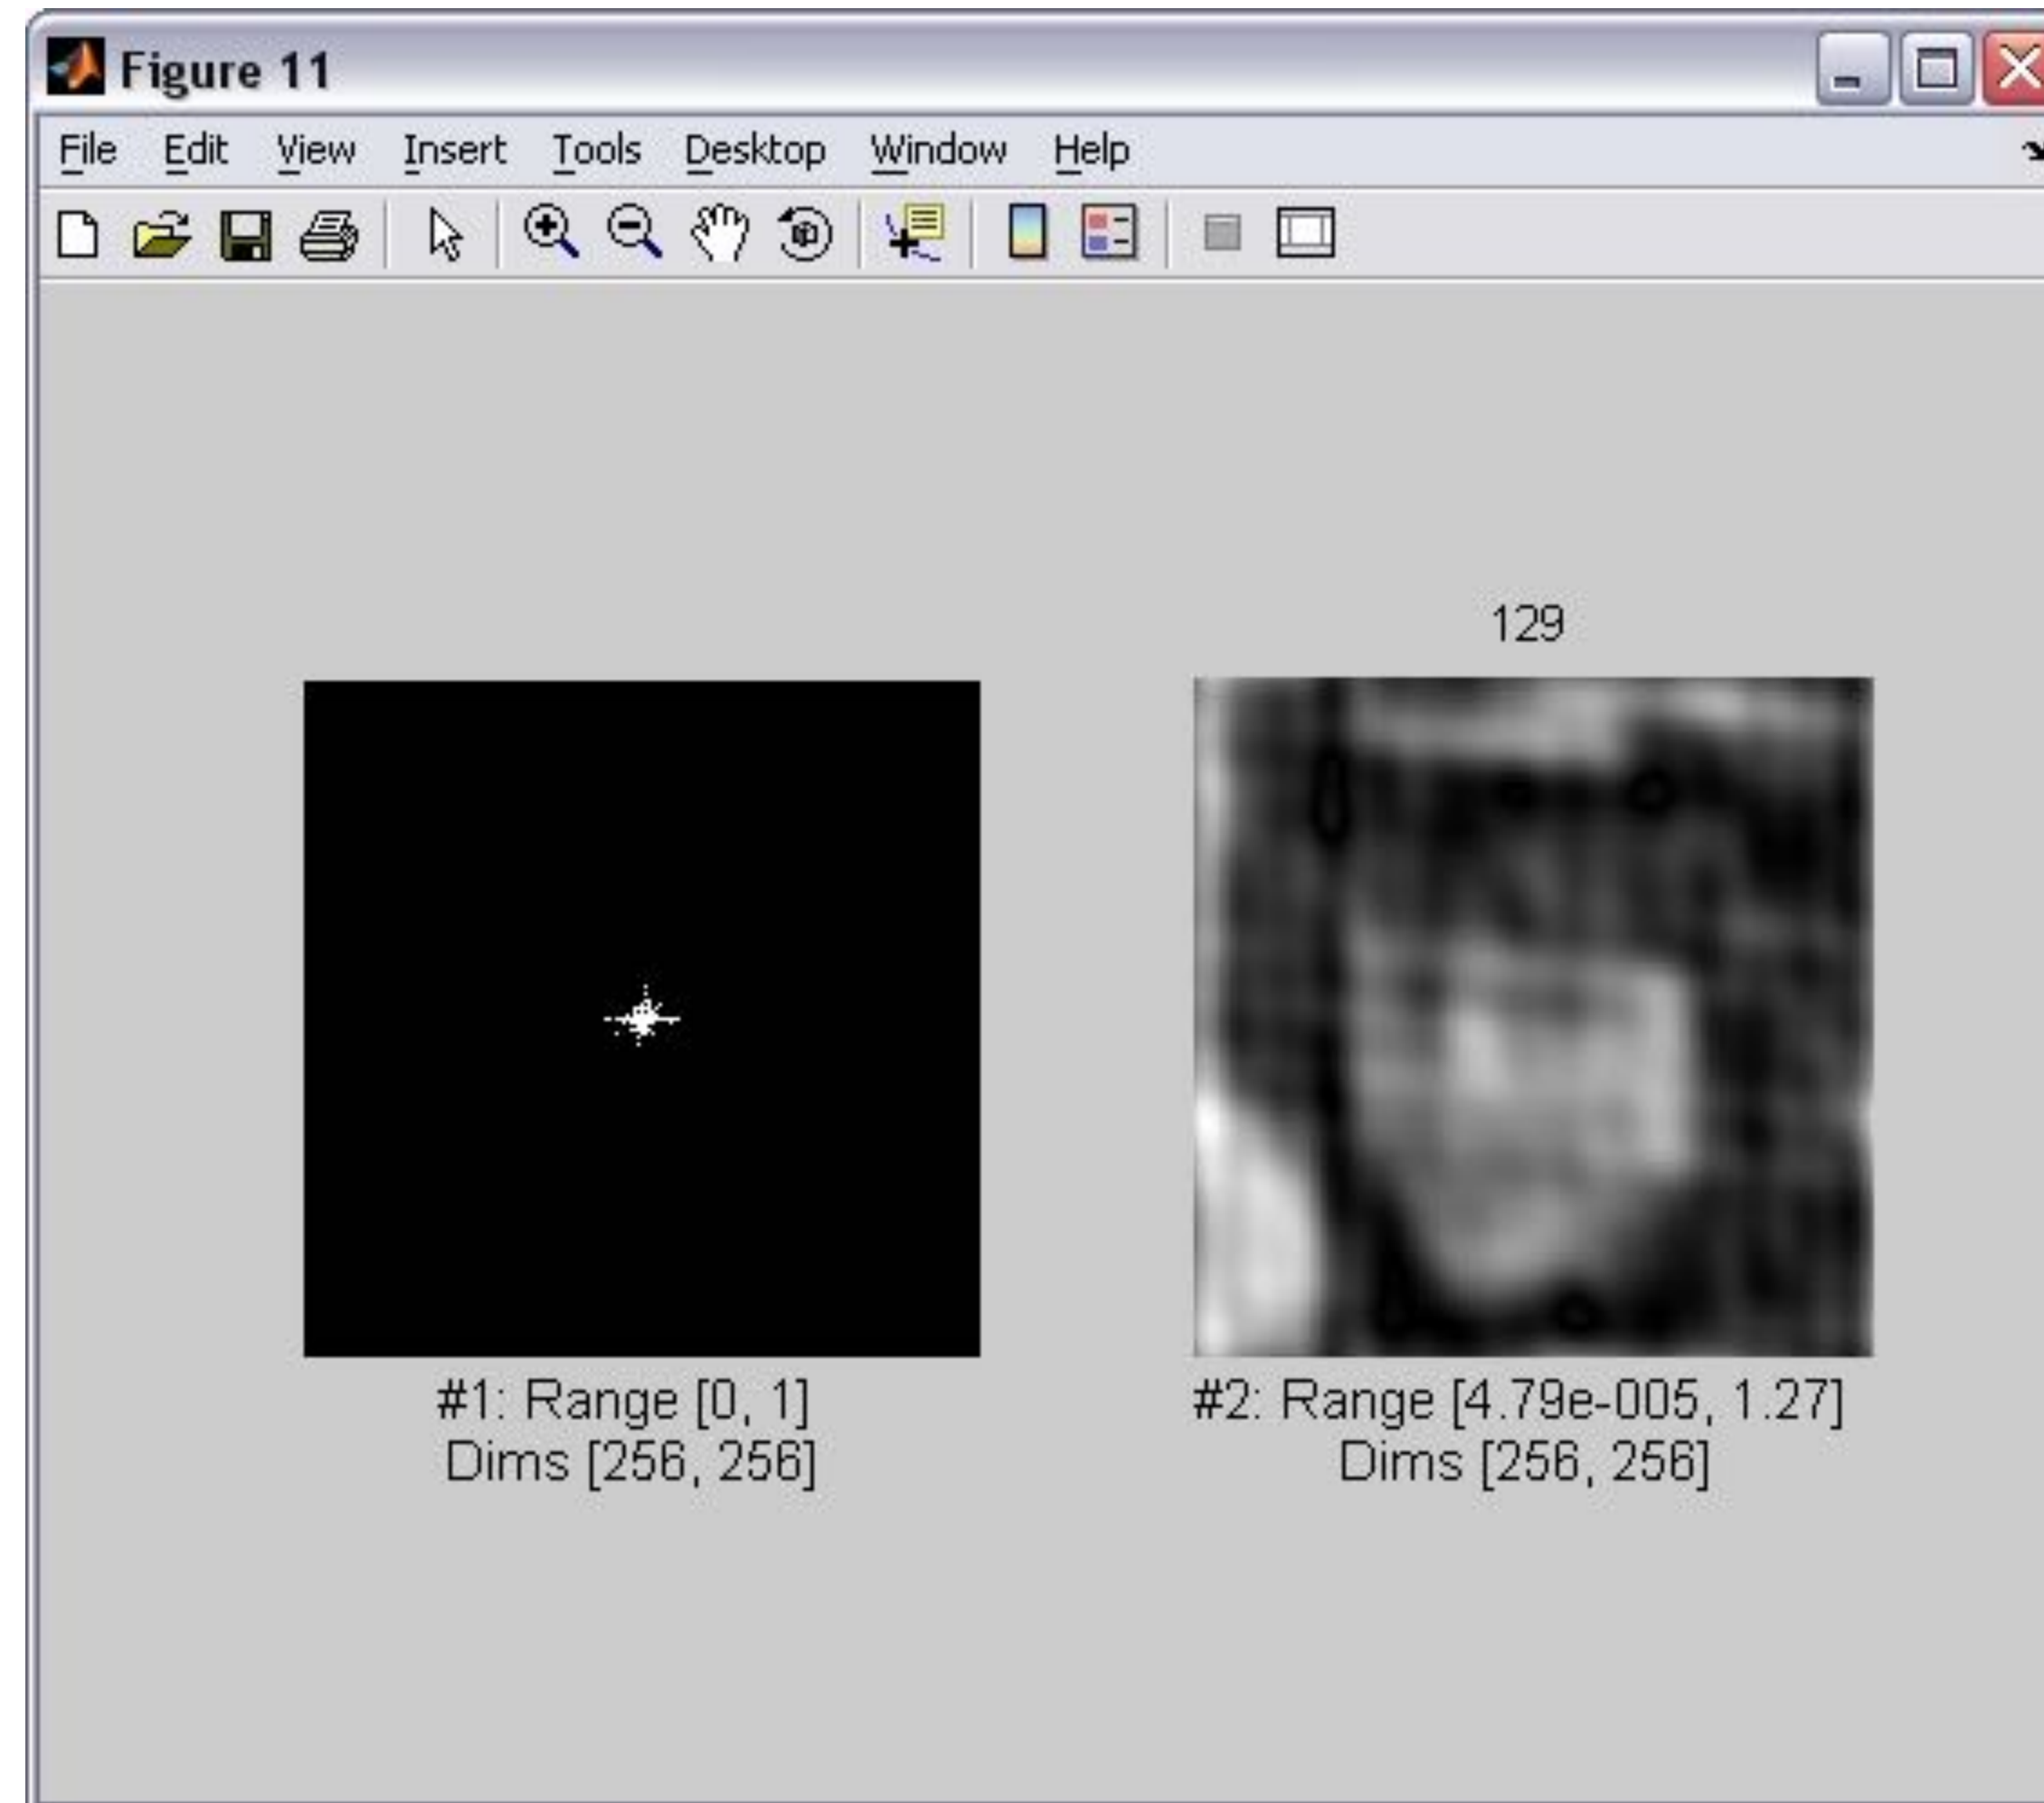
33



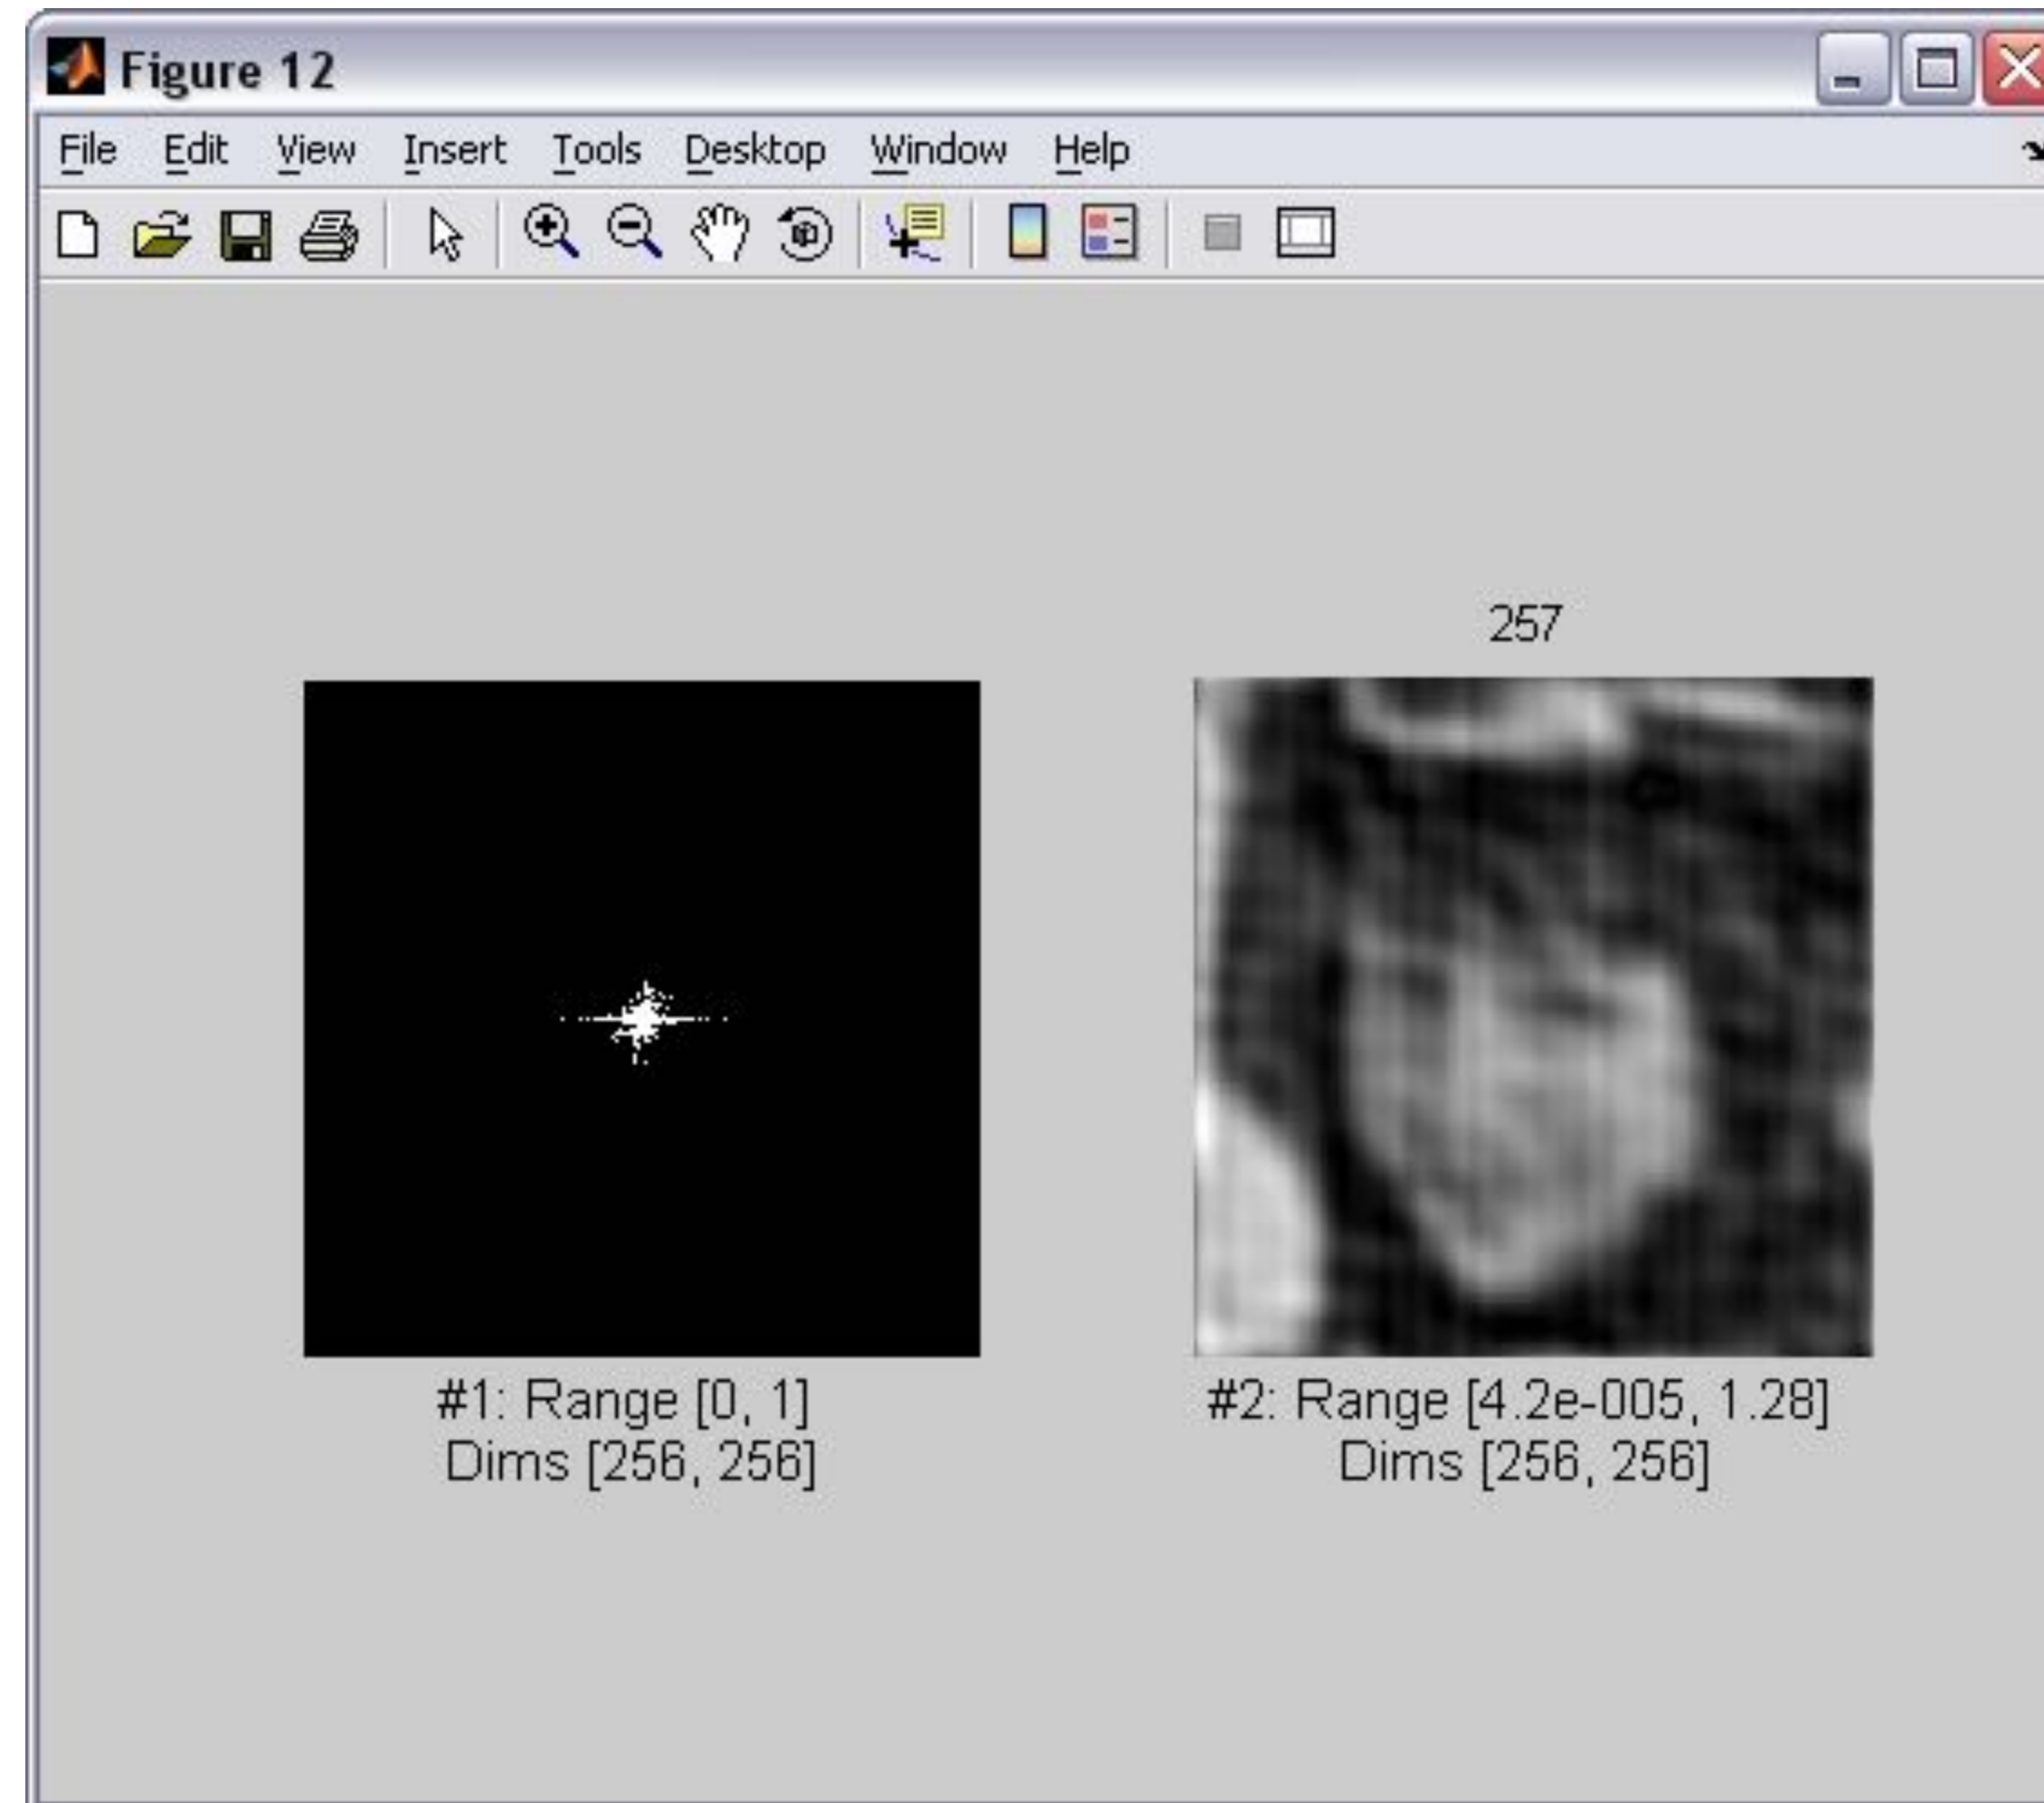
65



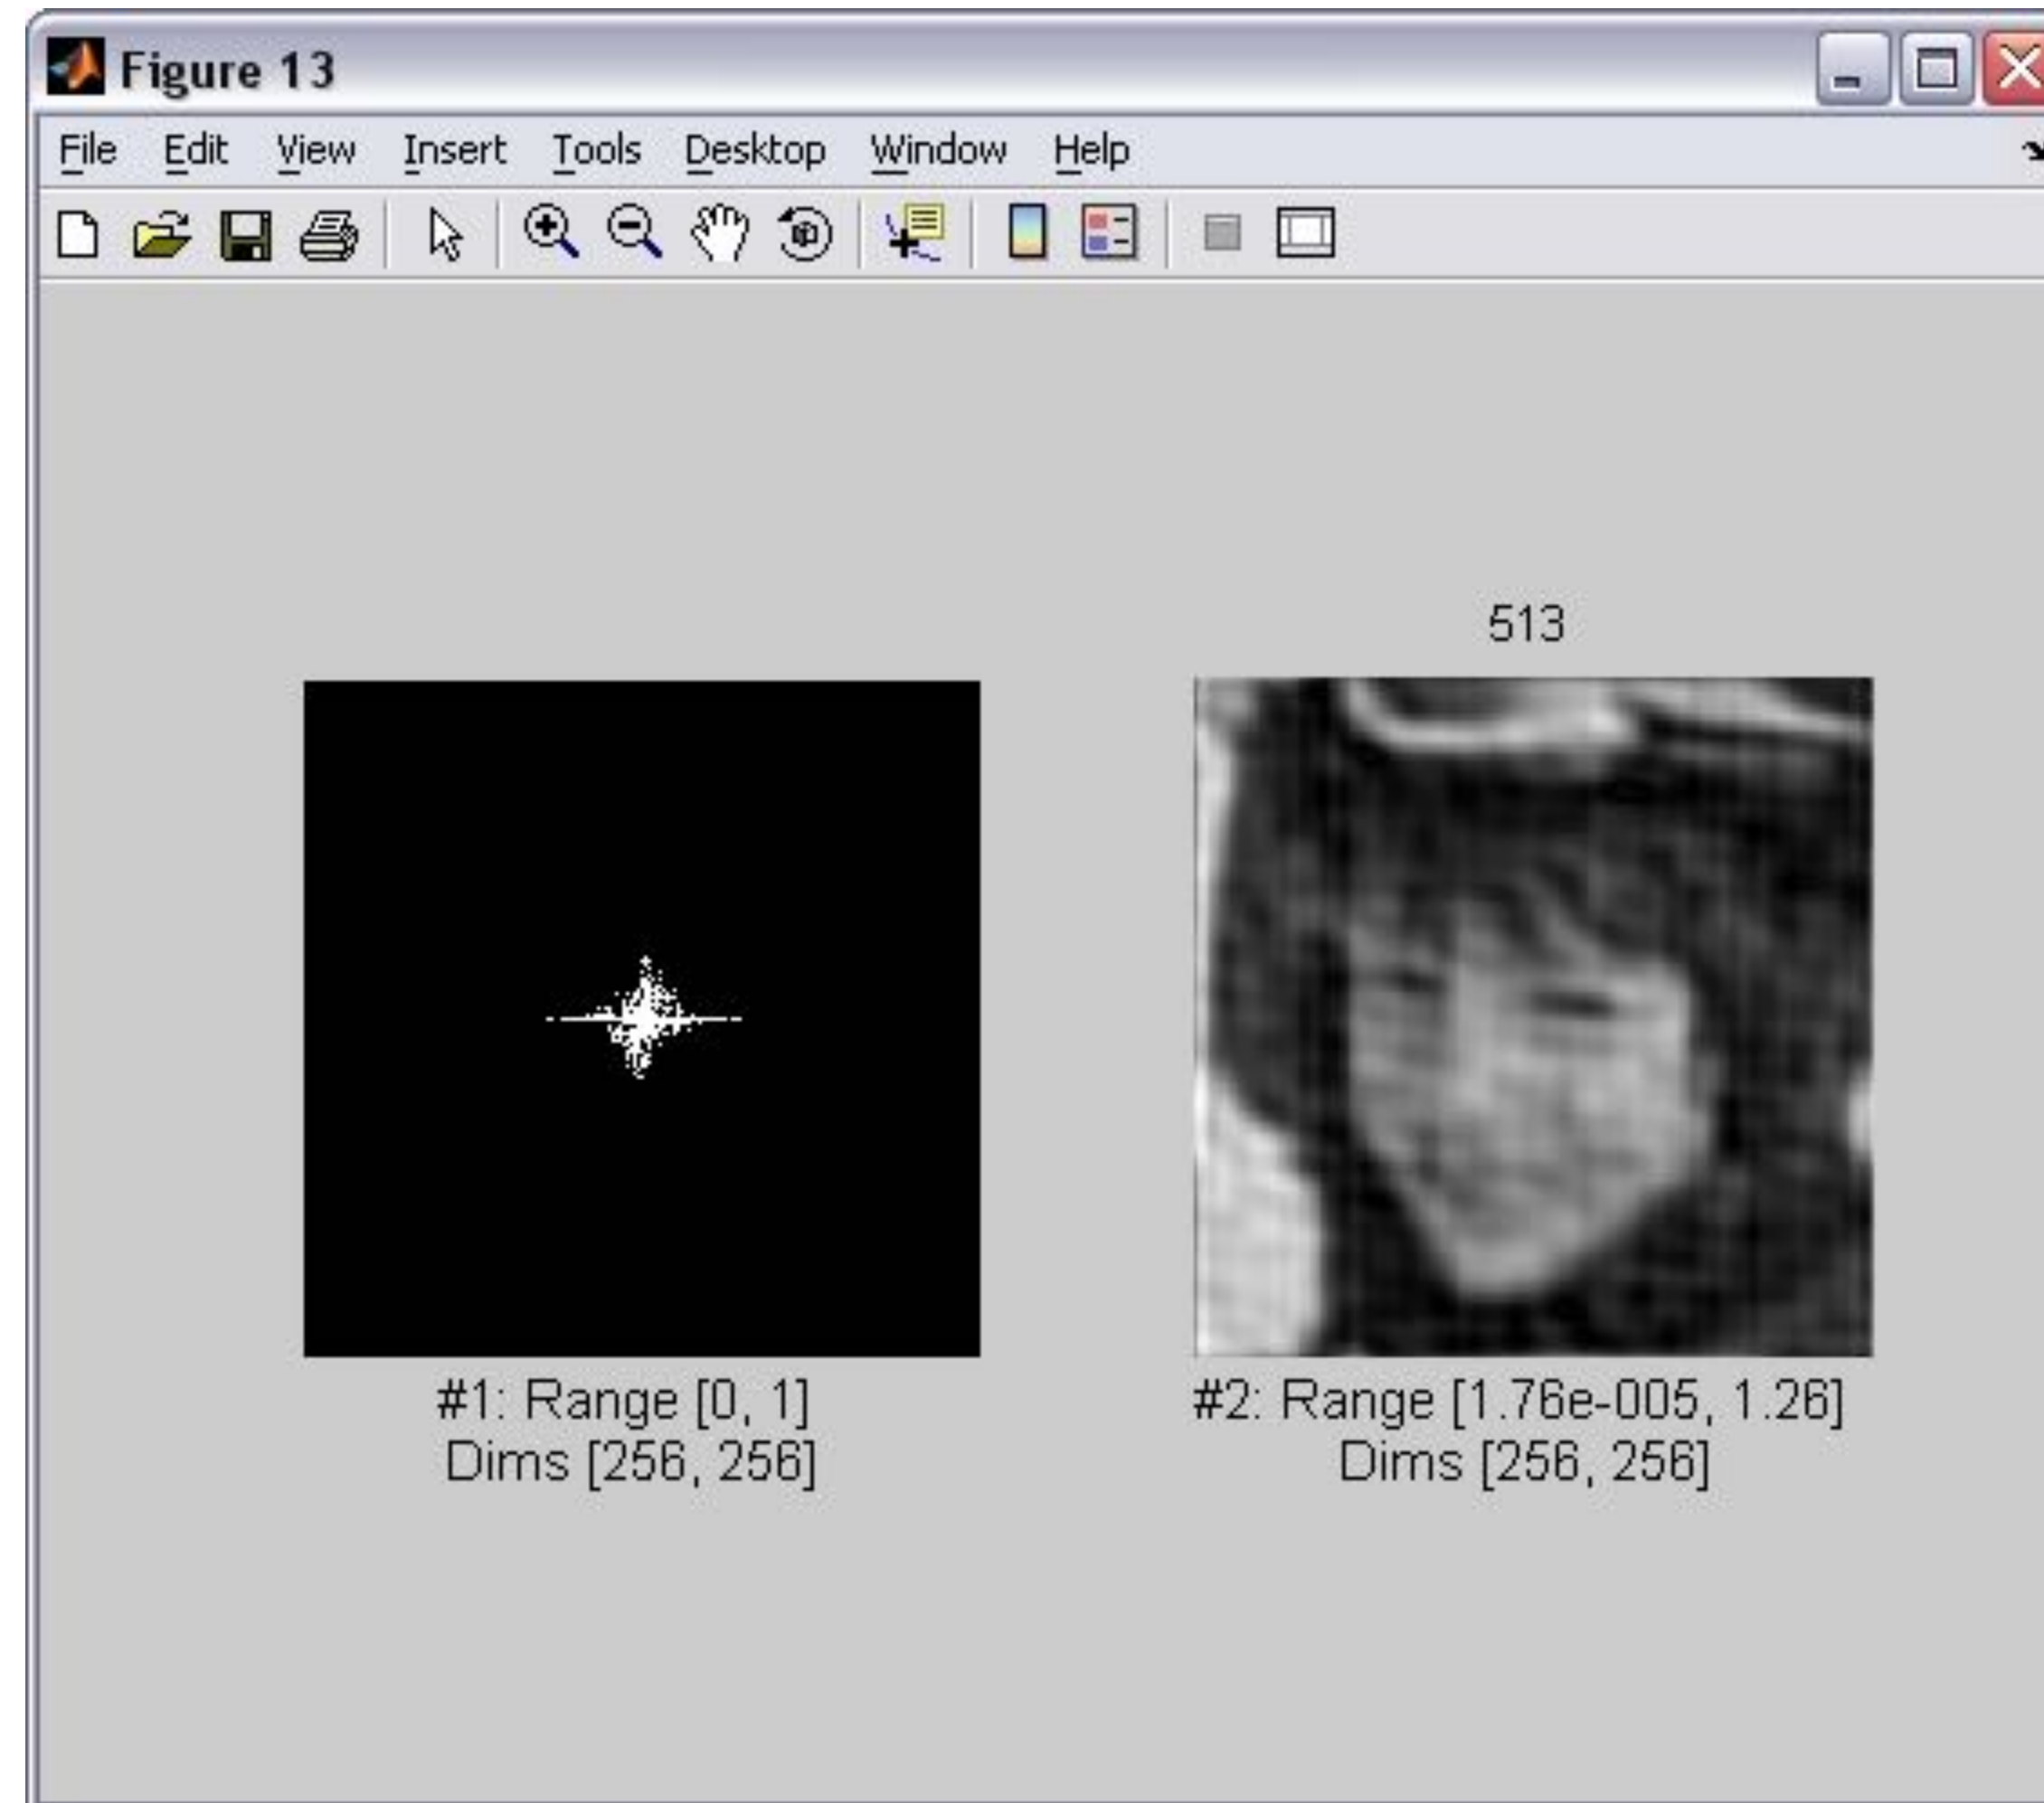
129



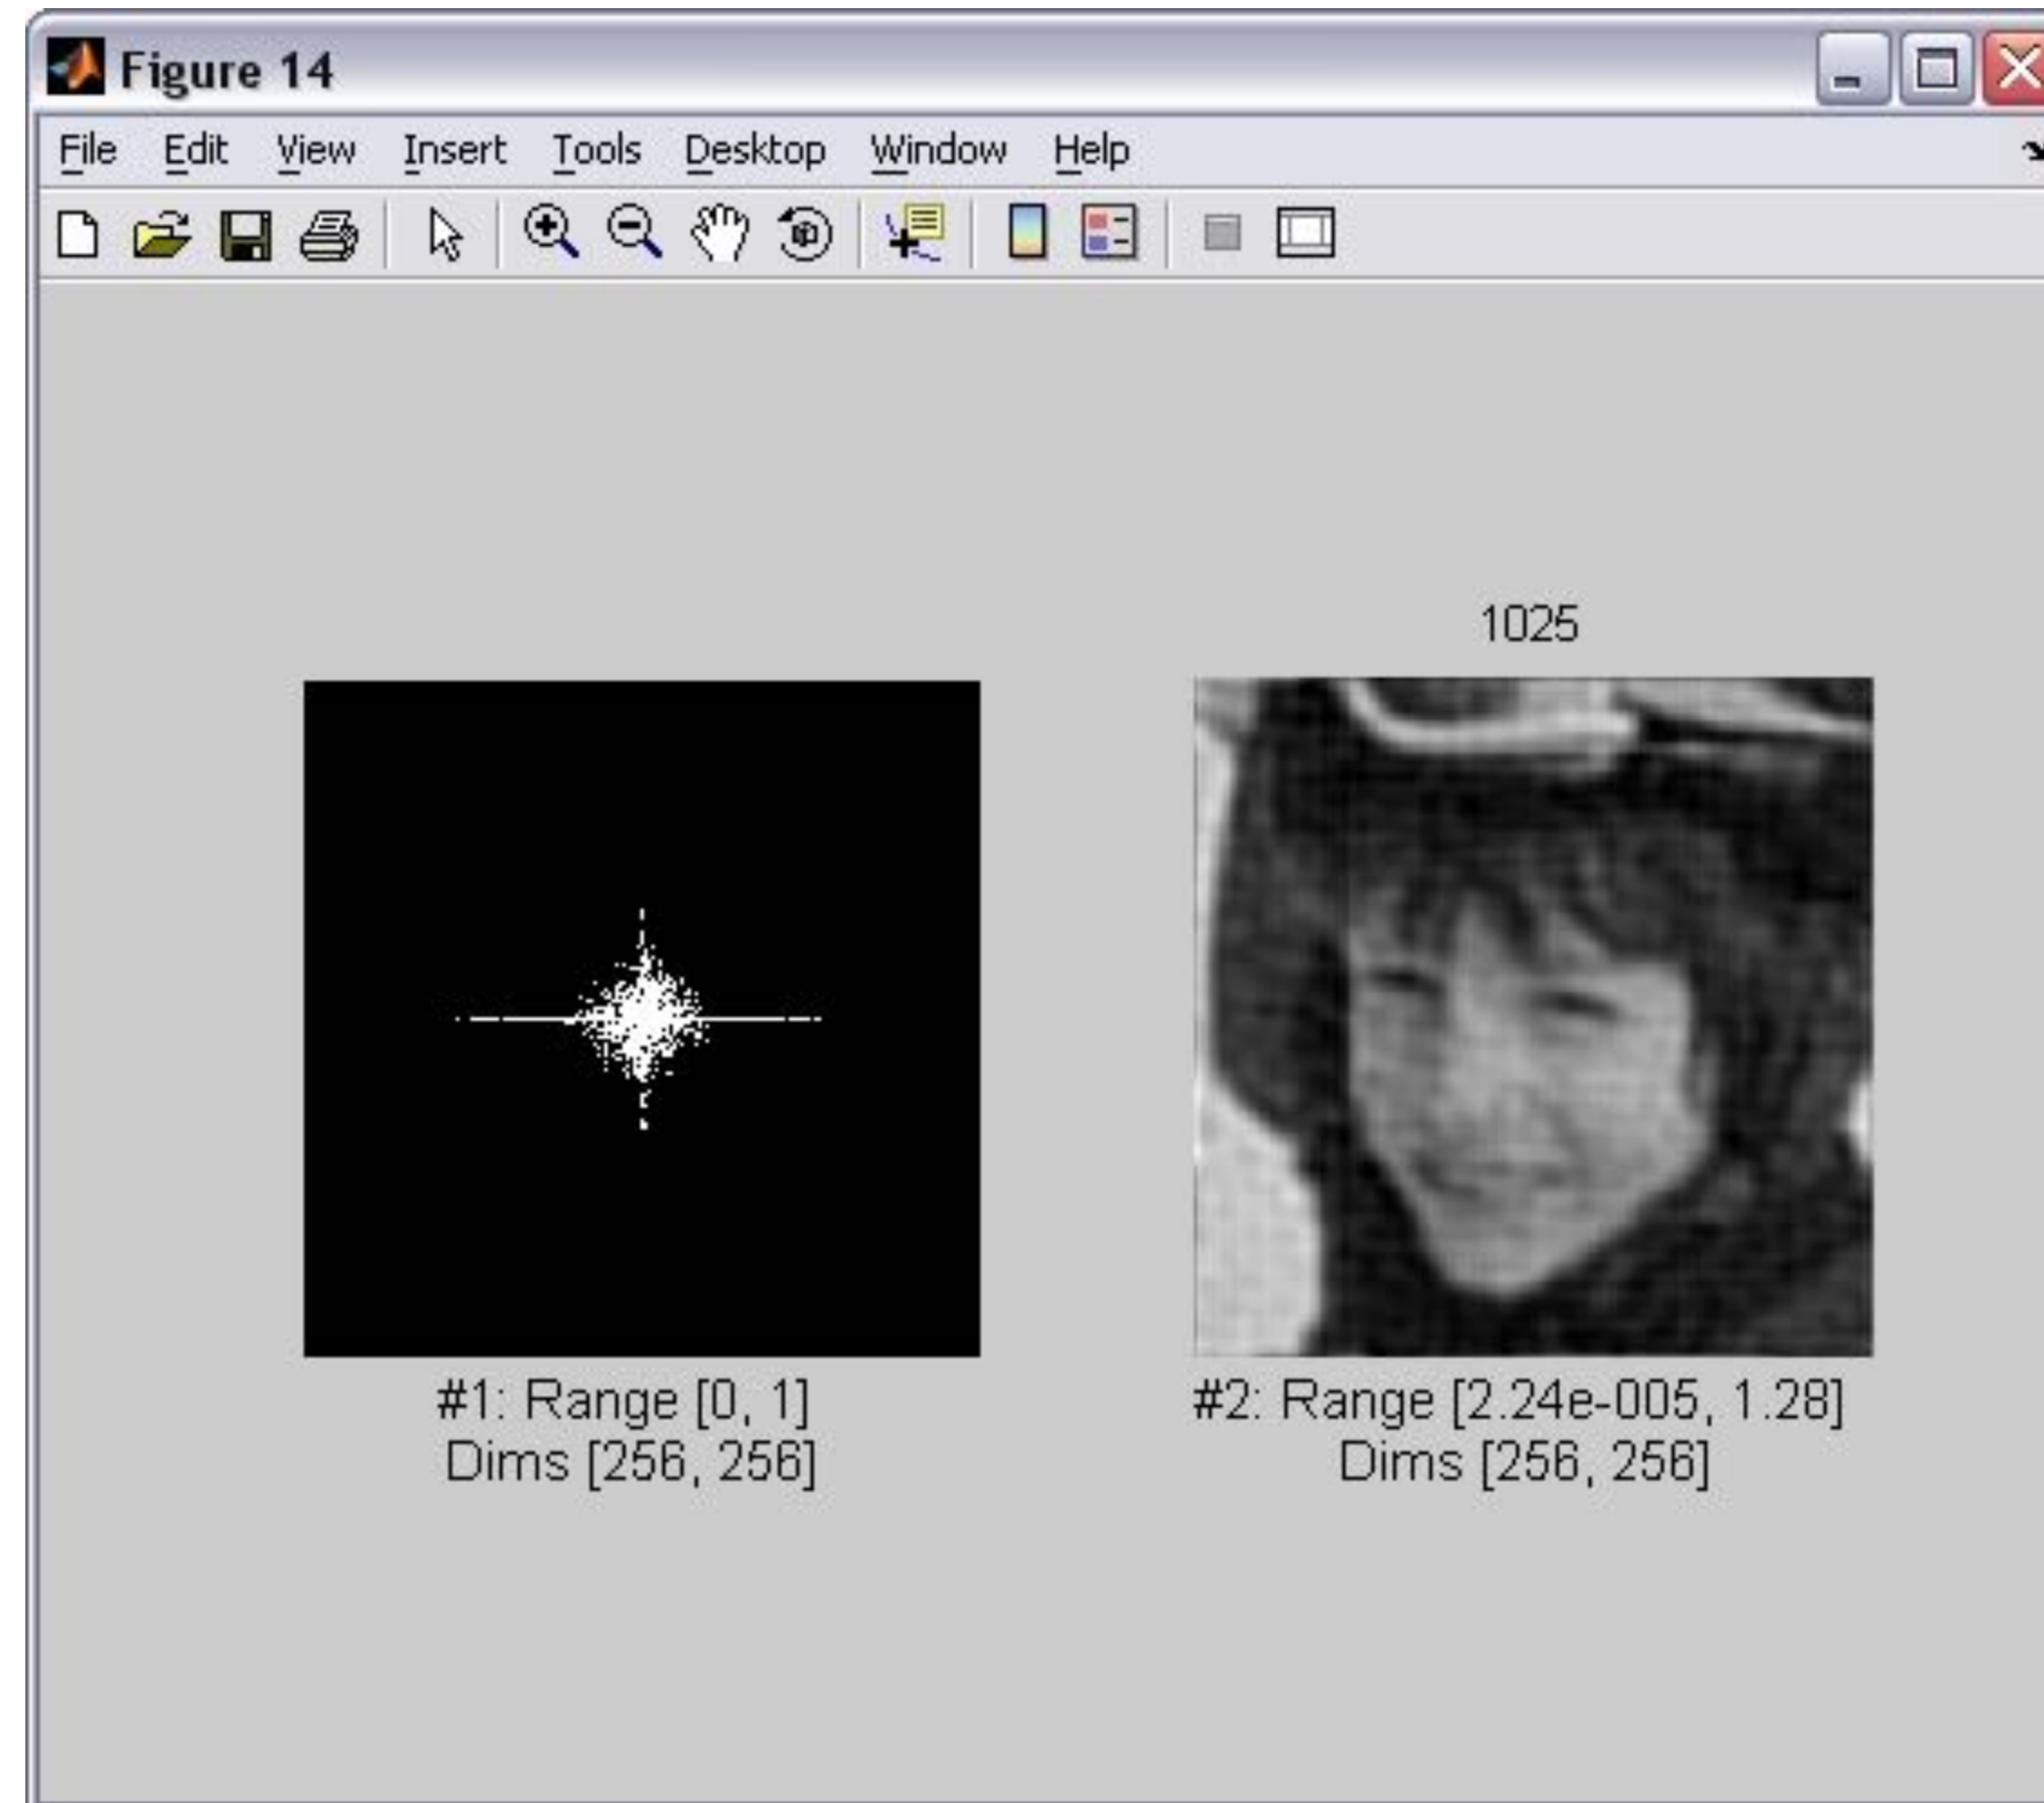
257



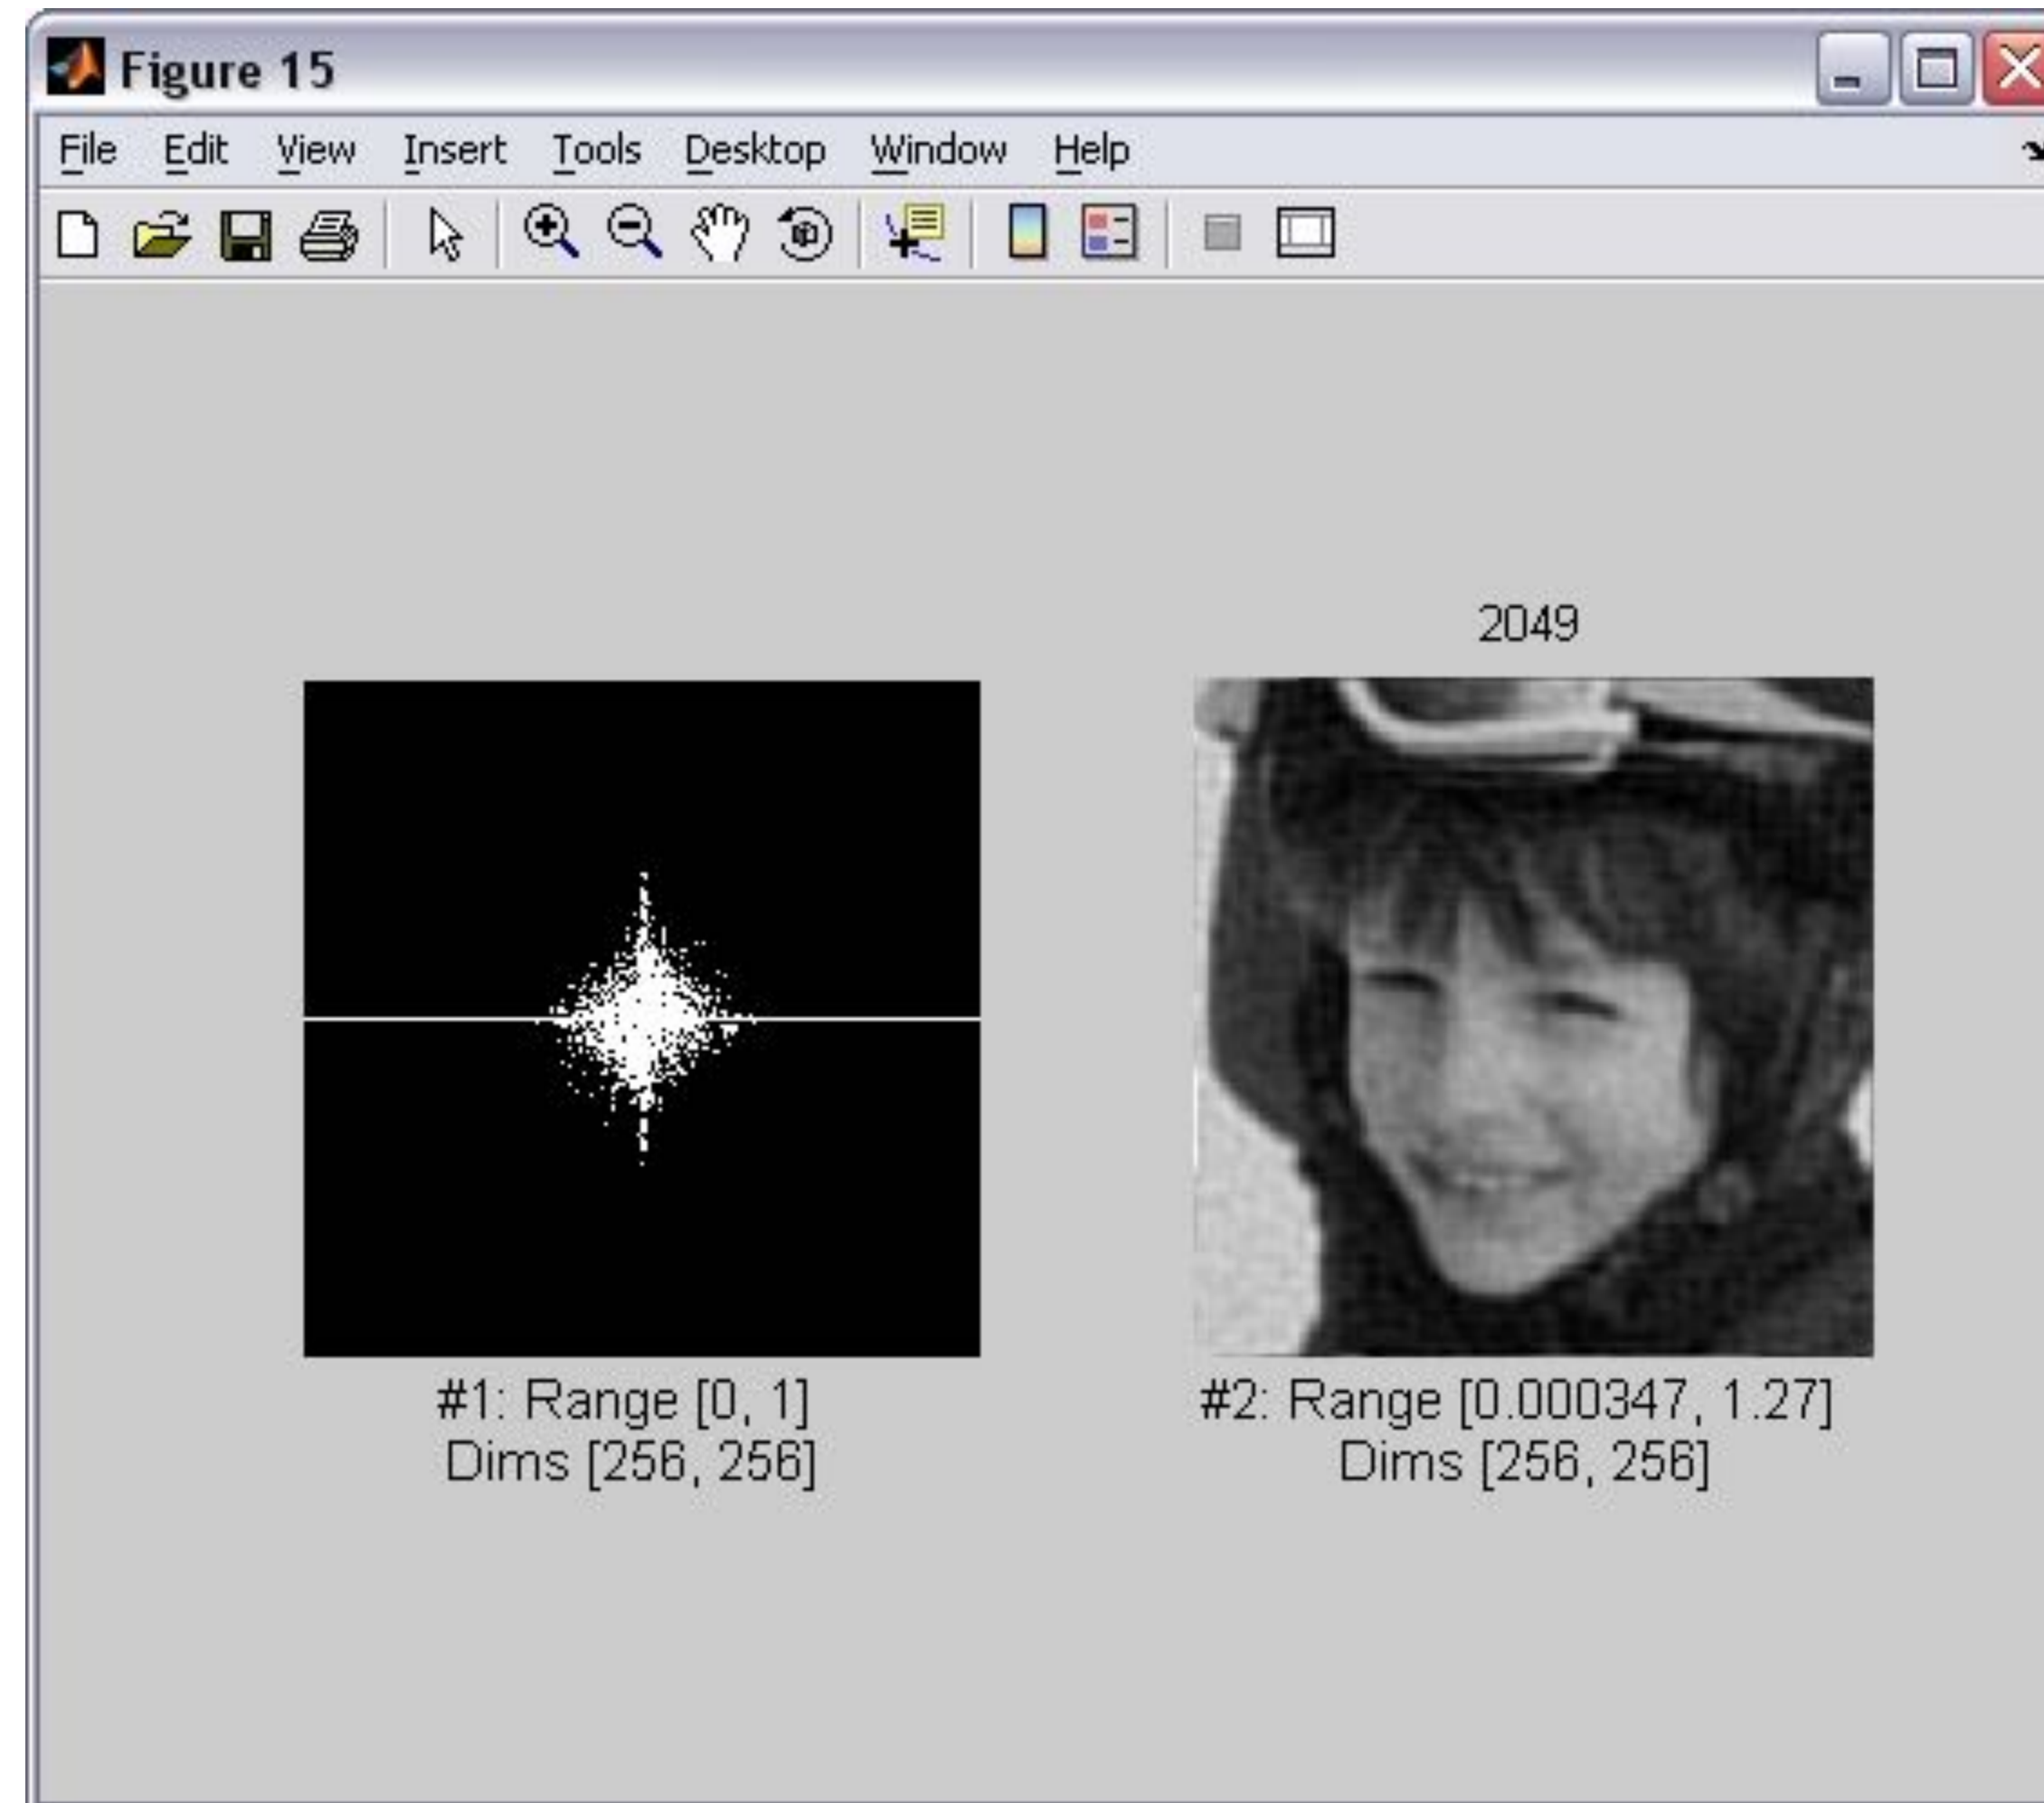
513



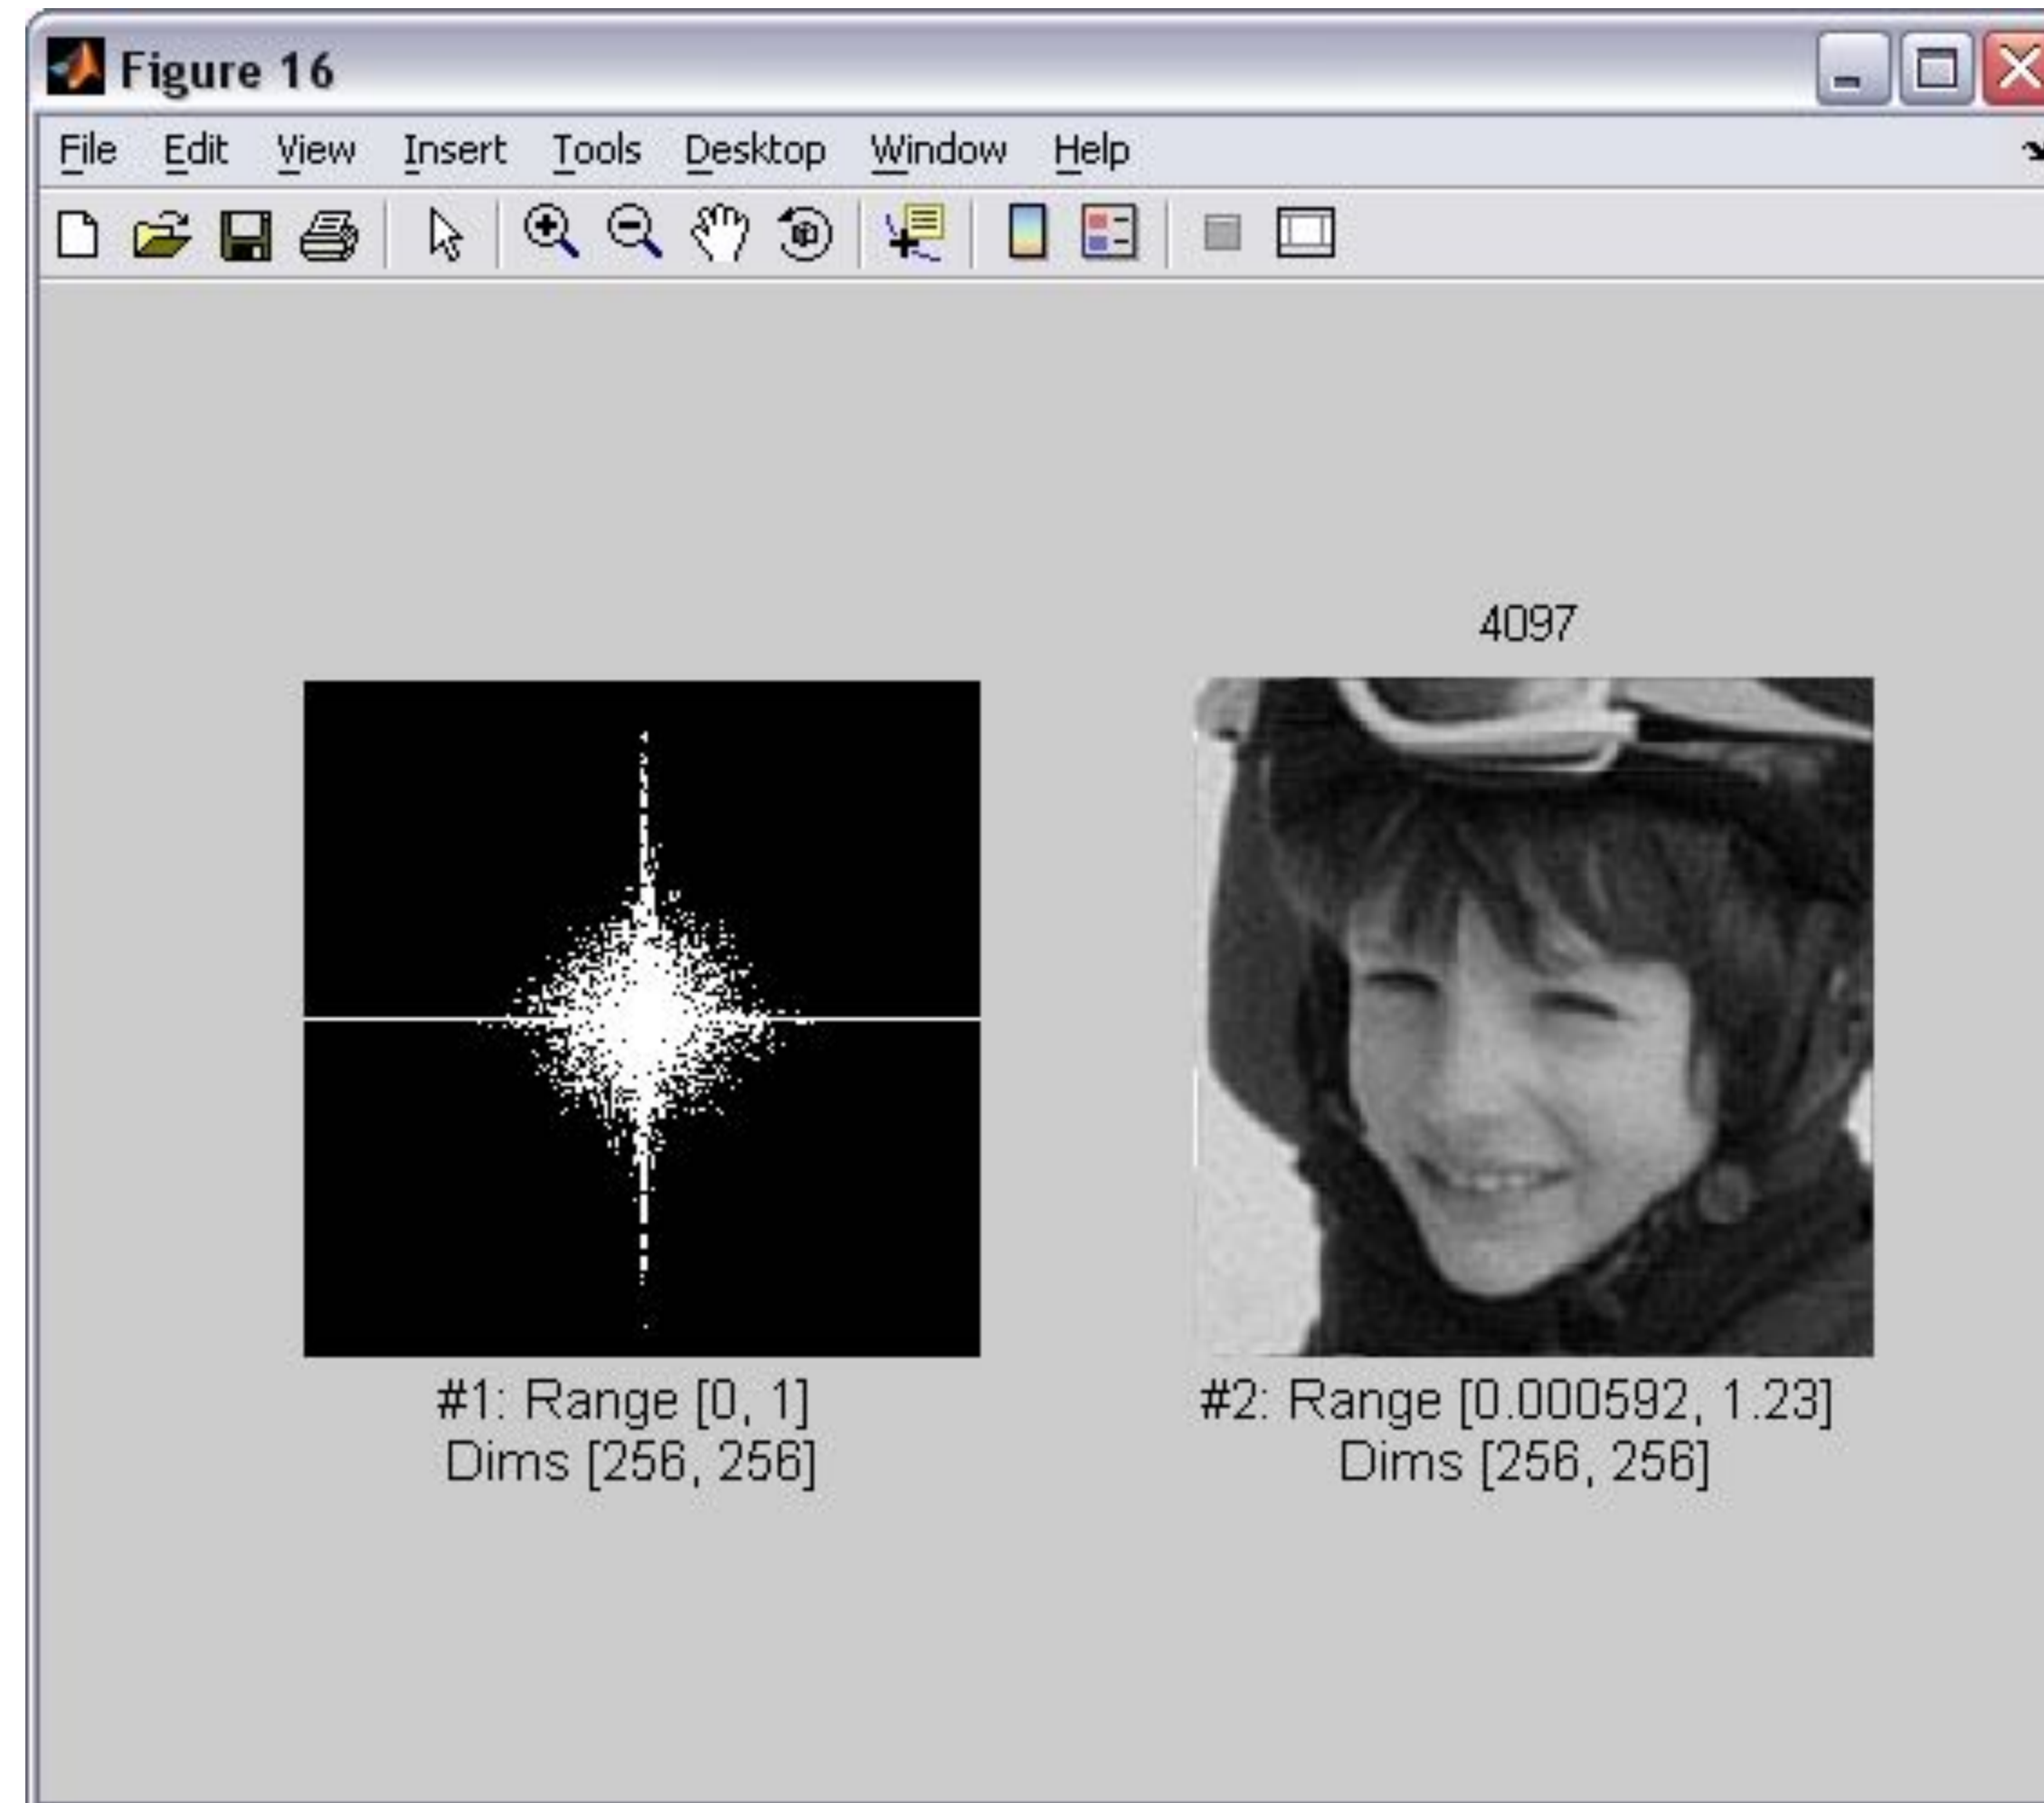
1025



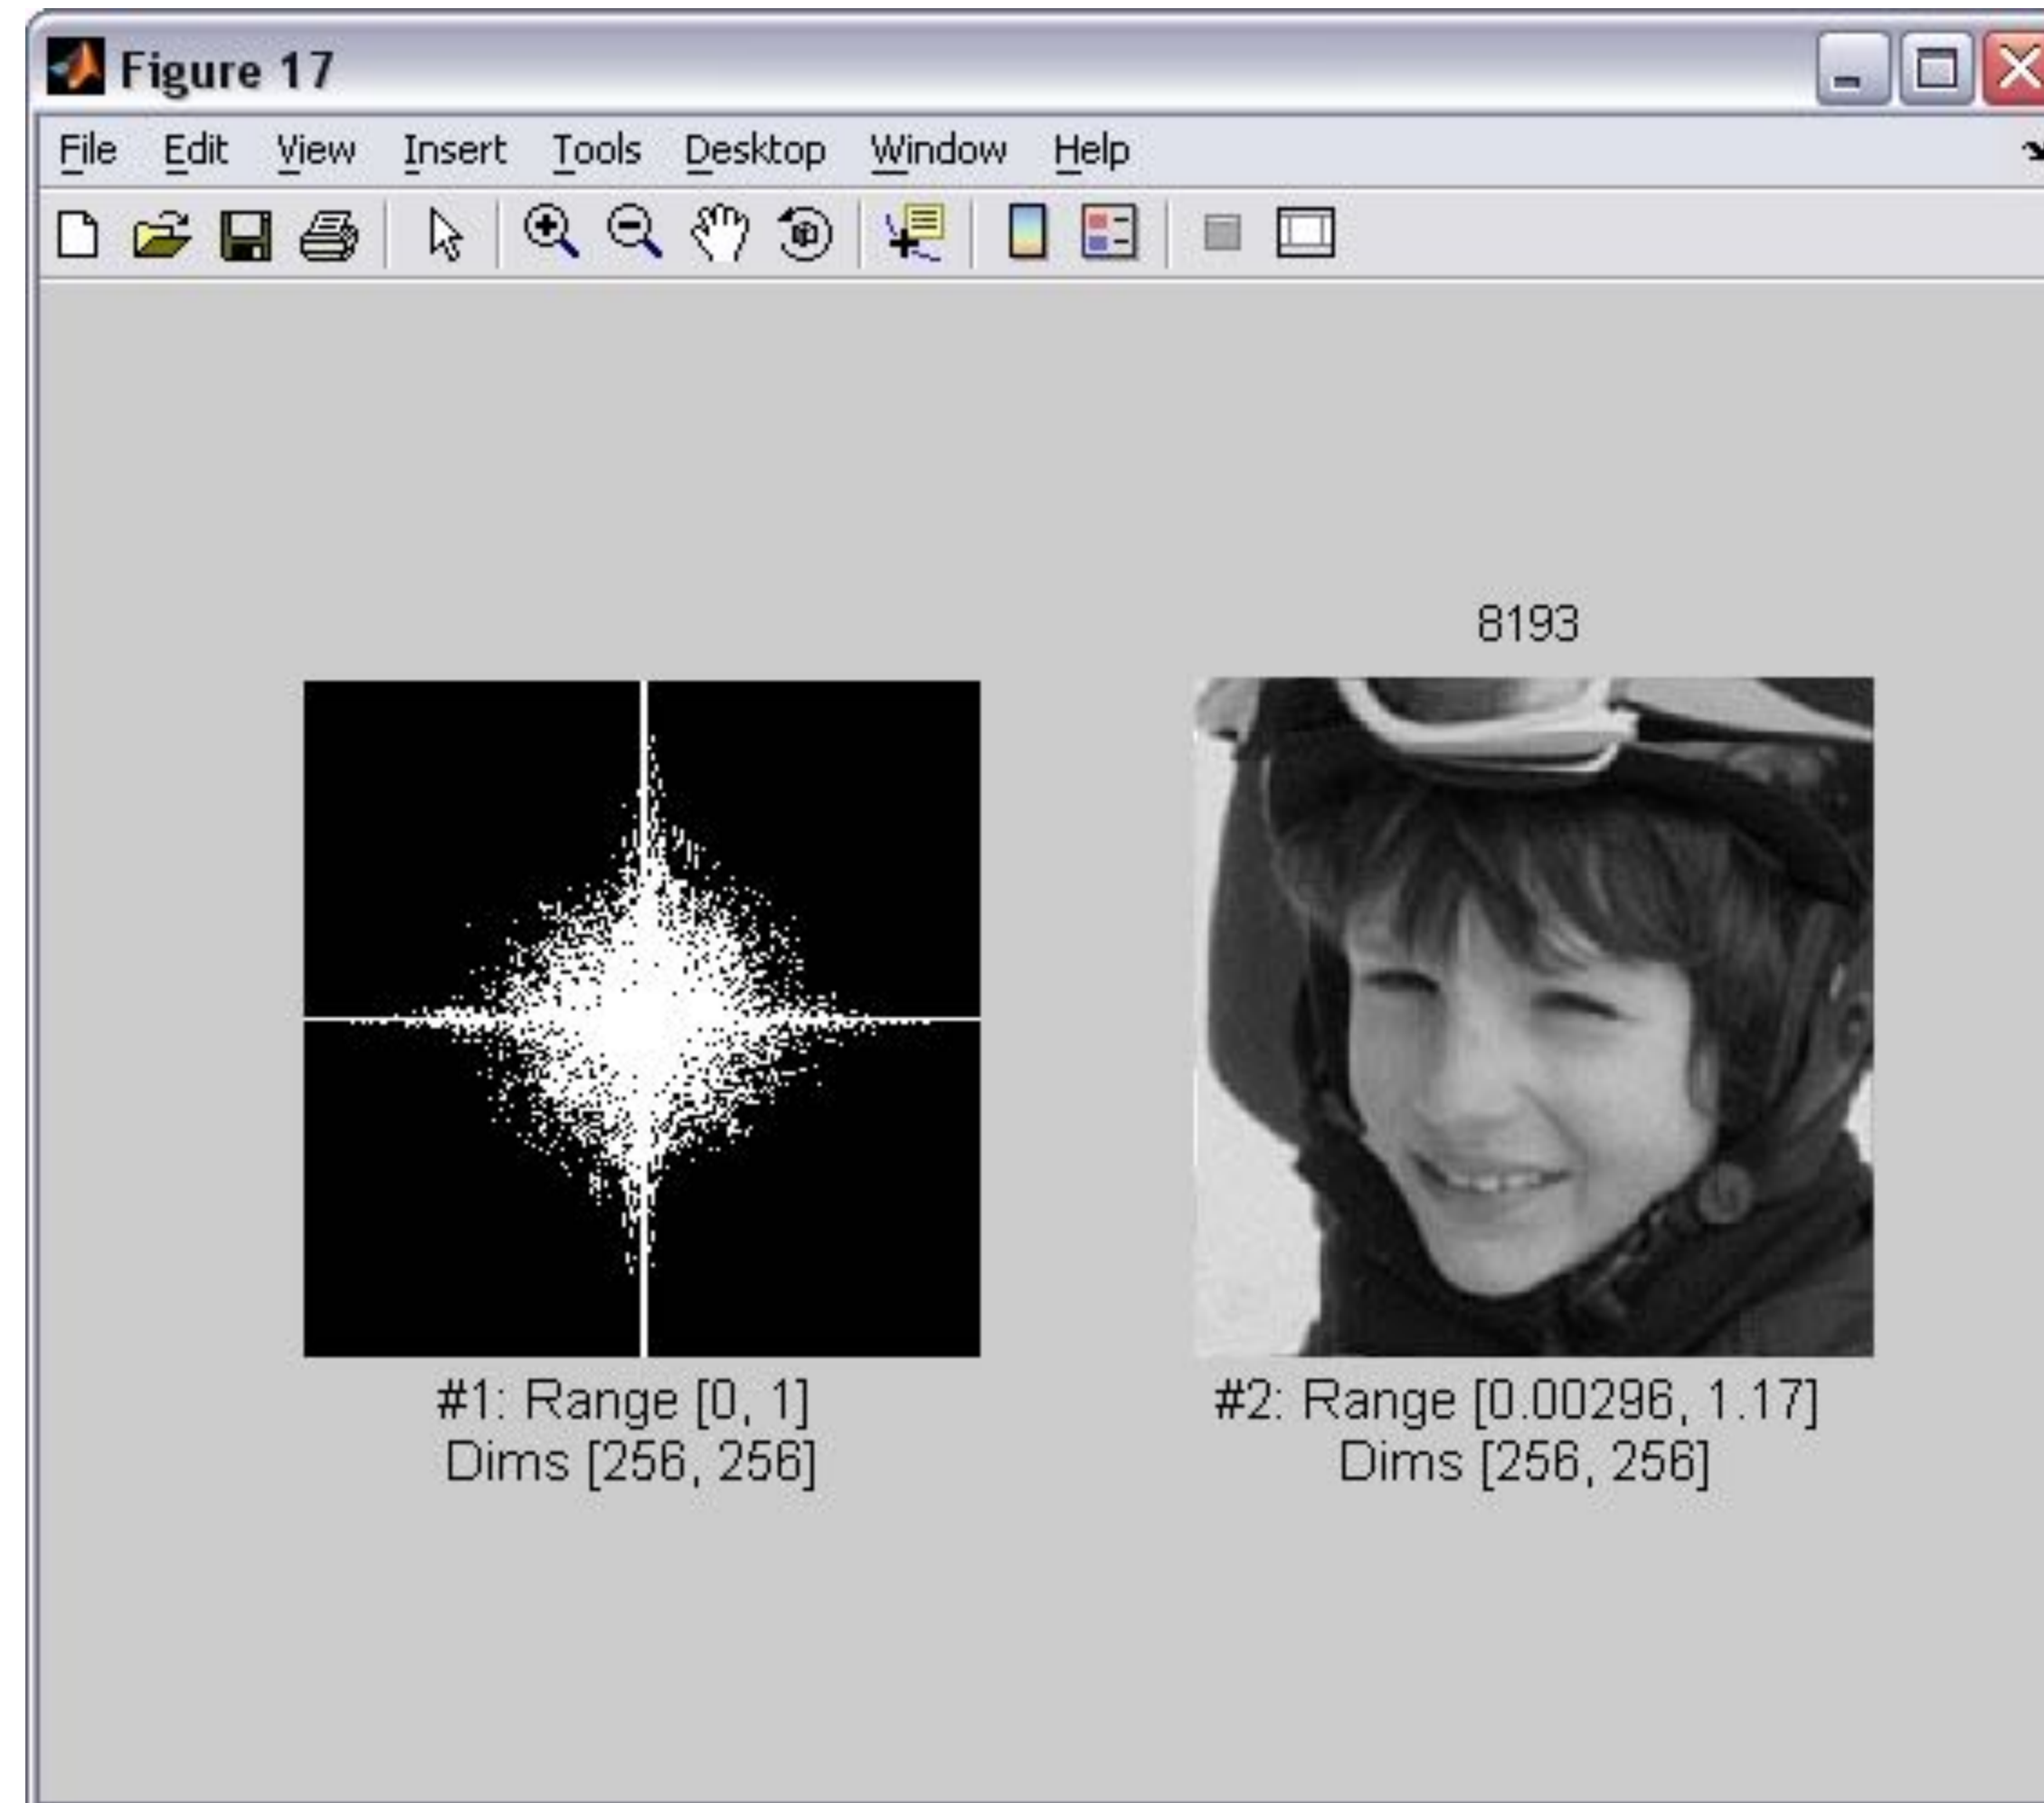
2049



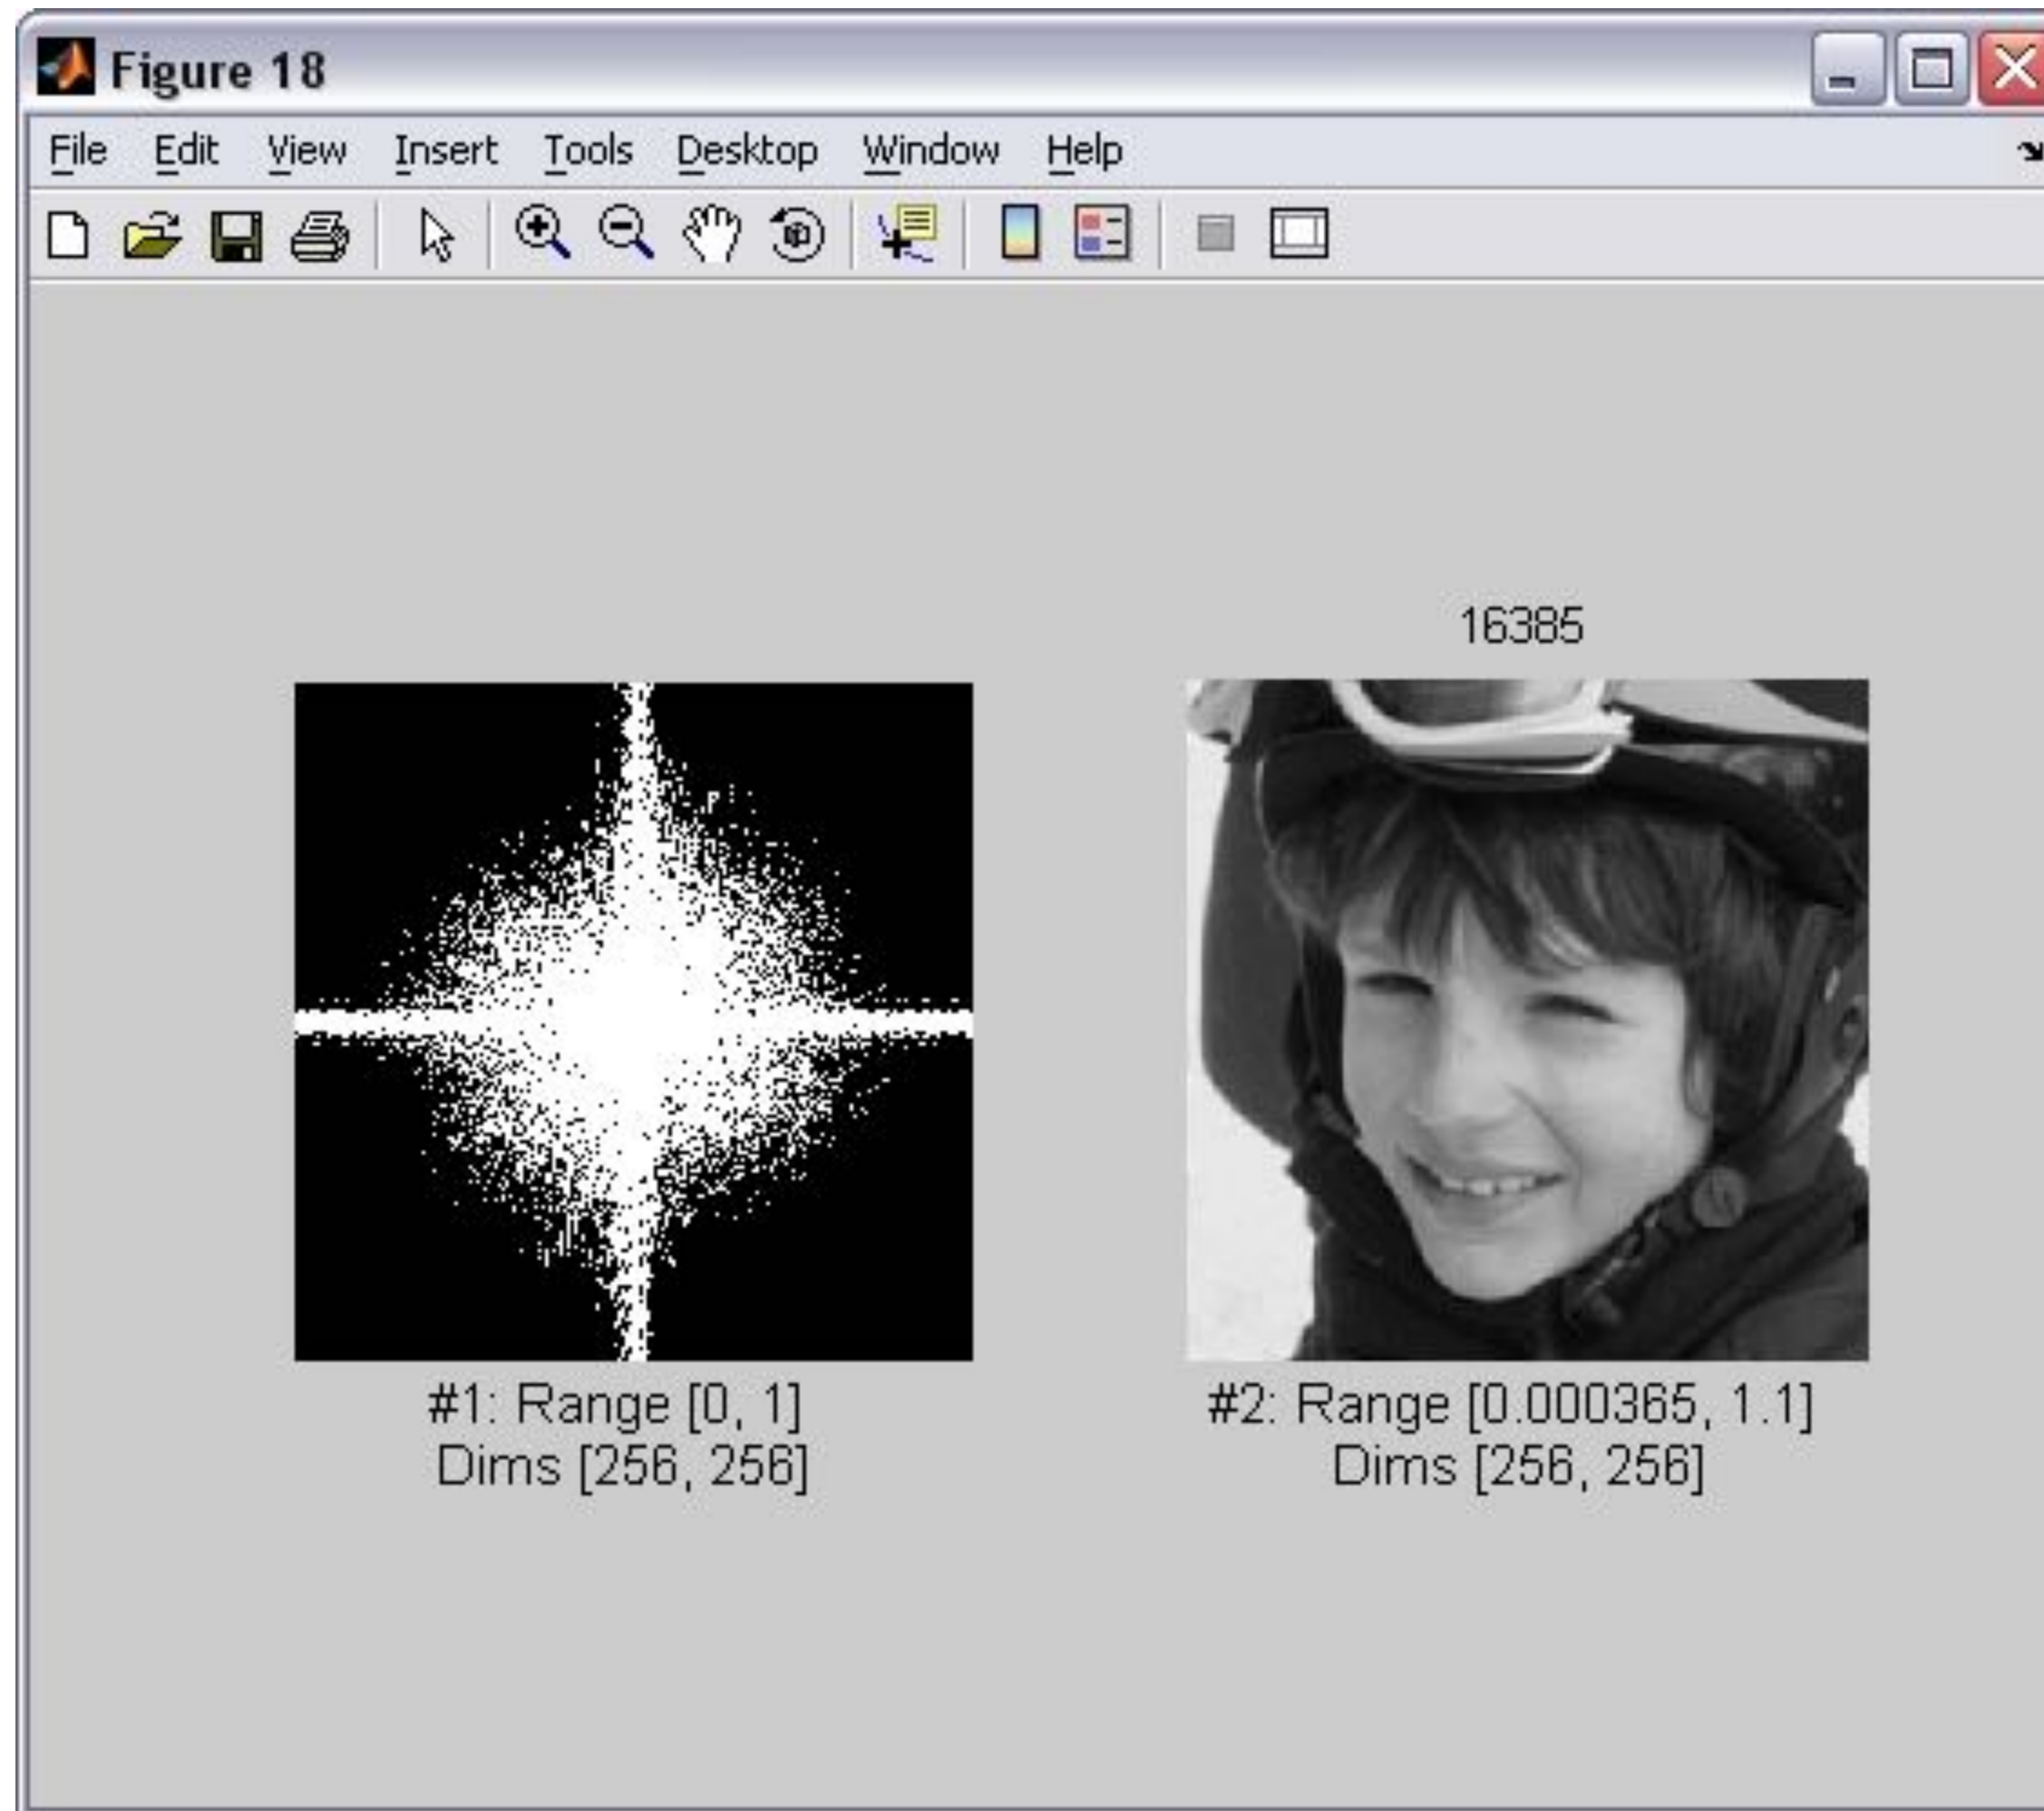
4097



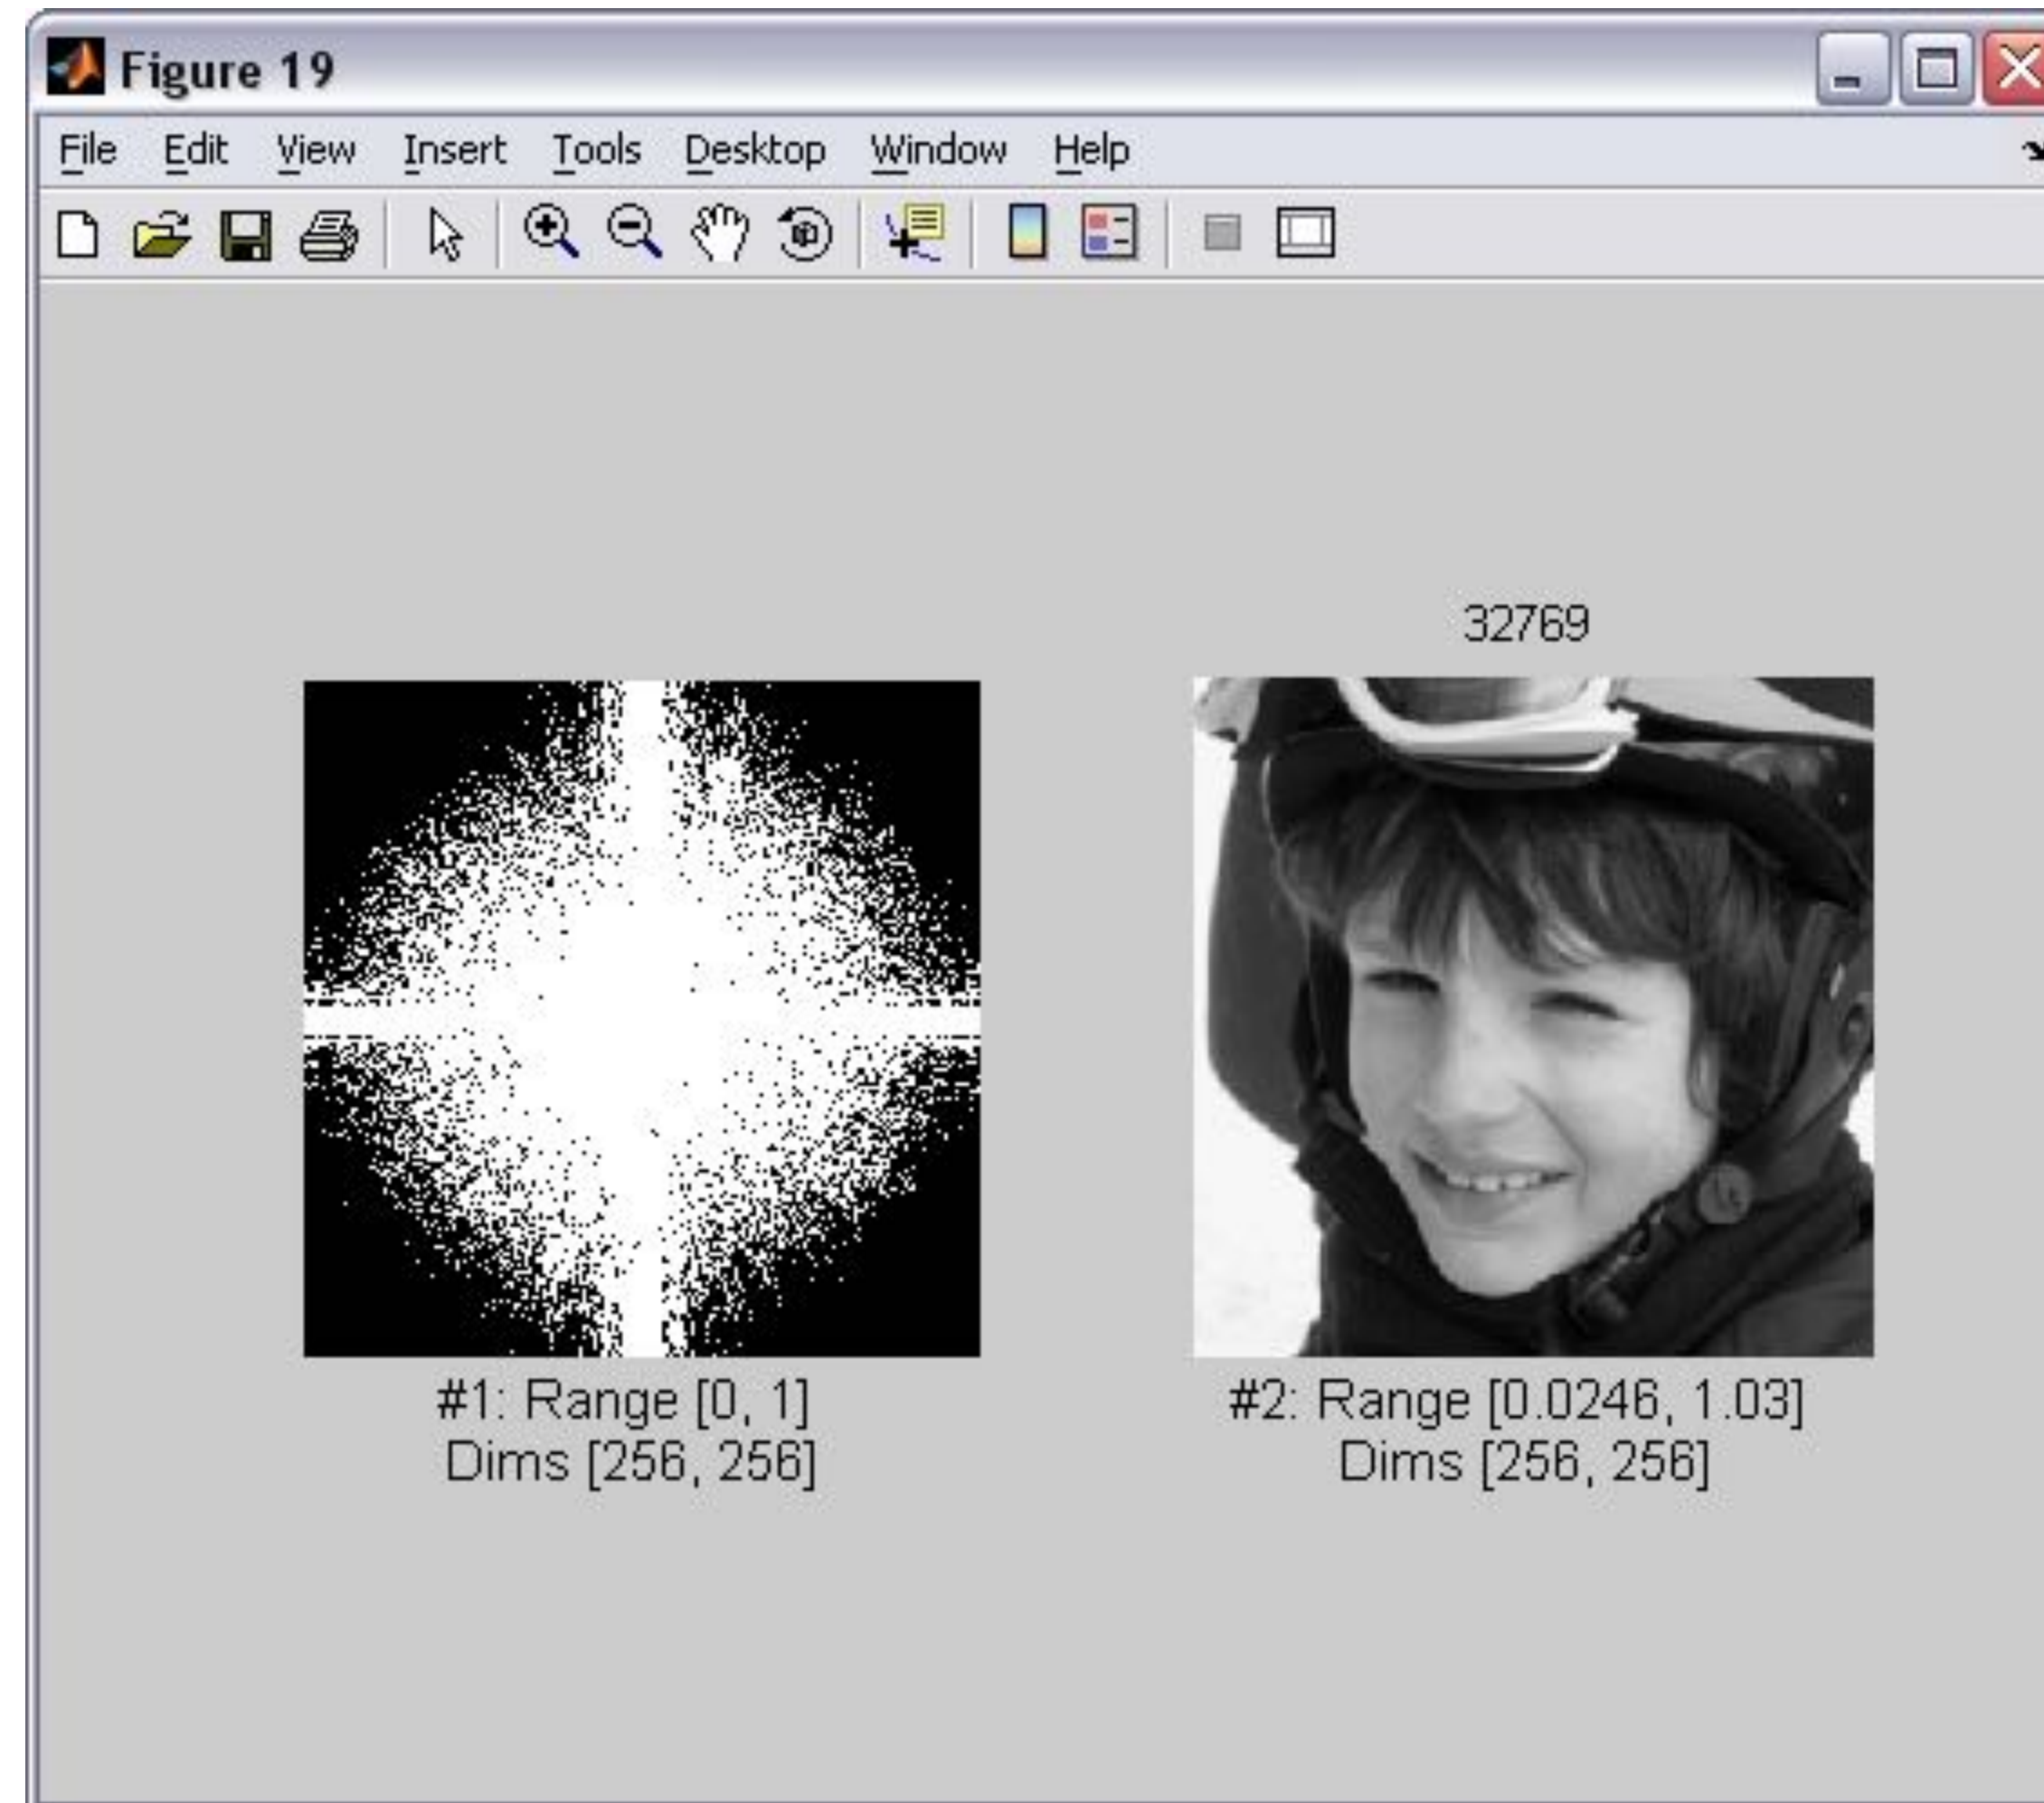
8193



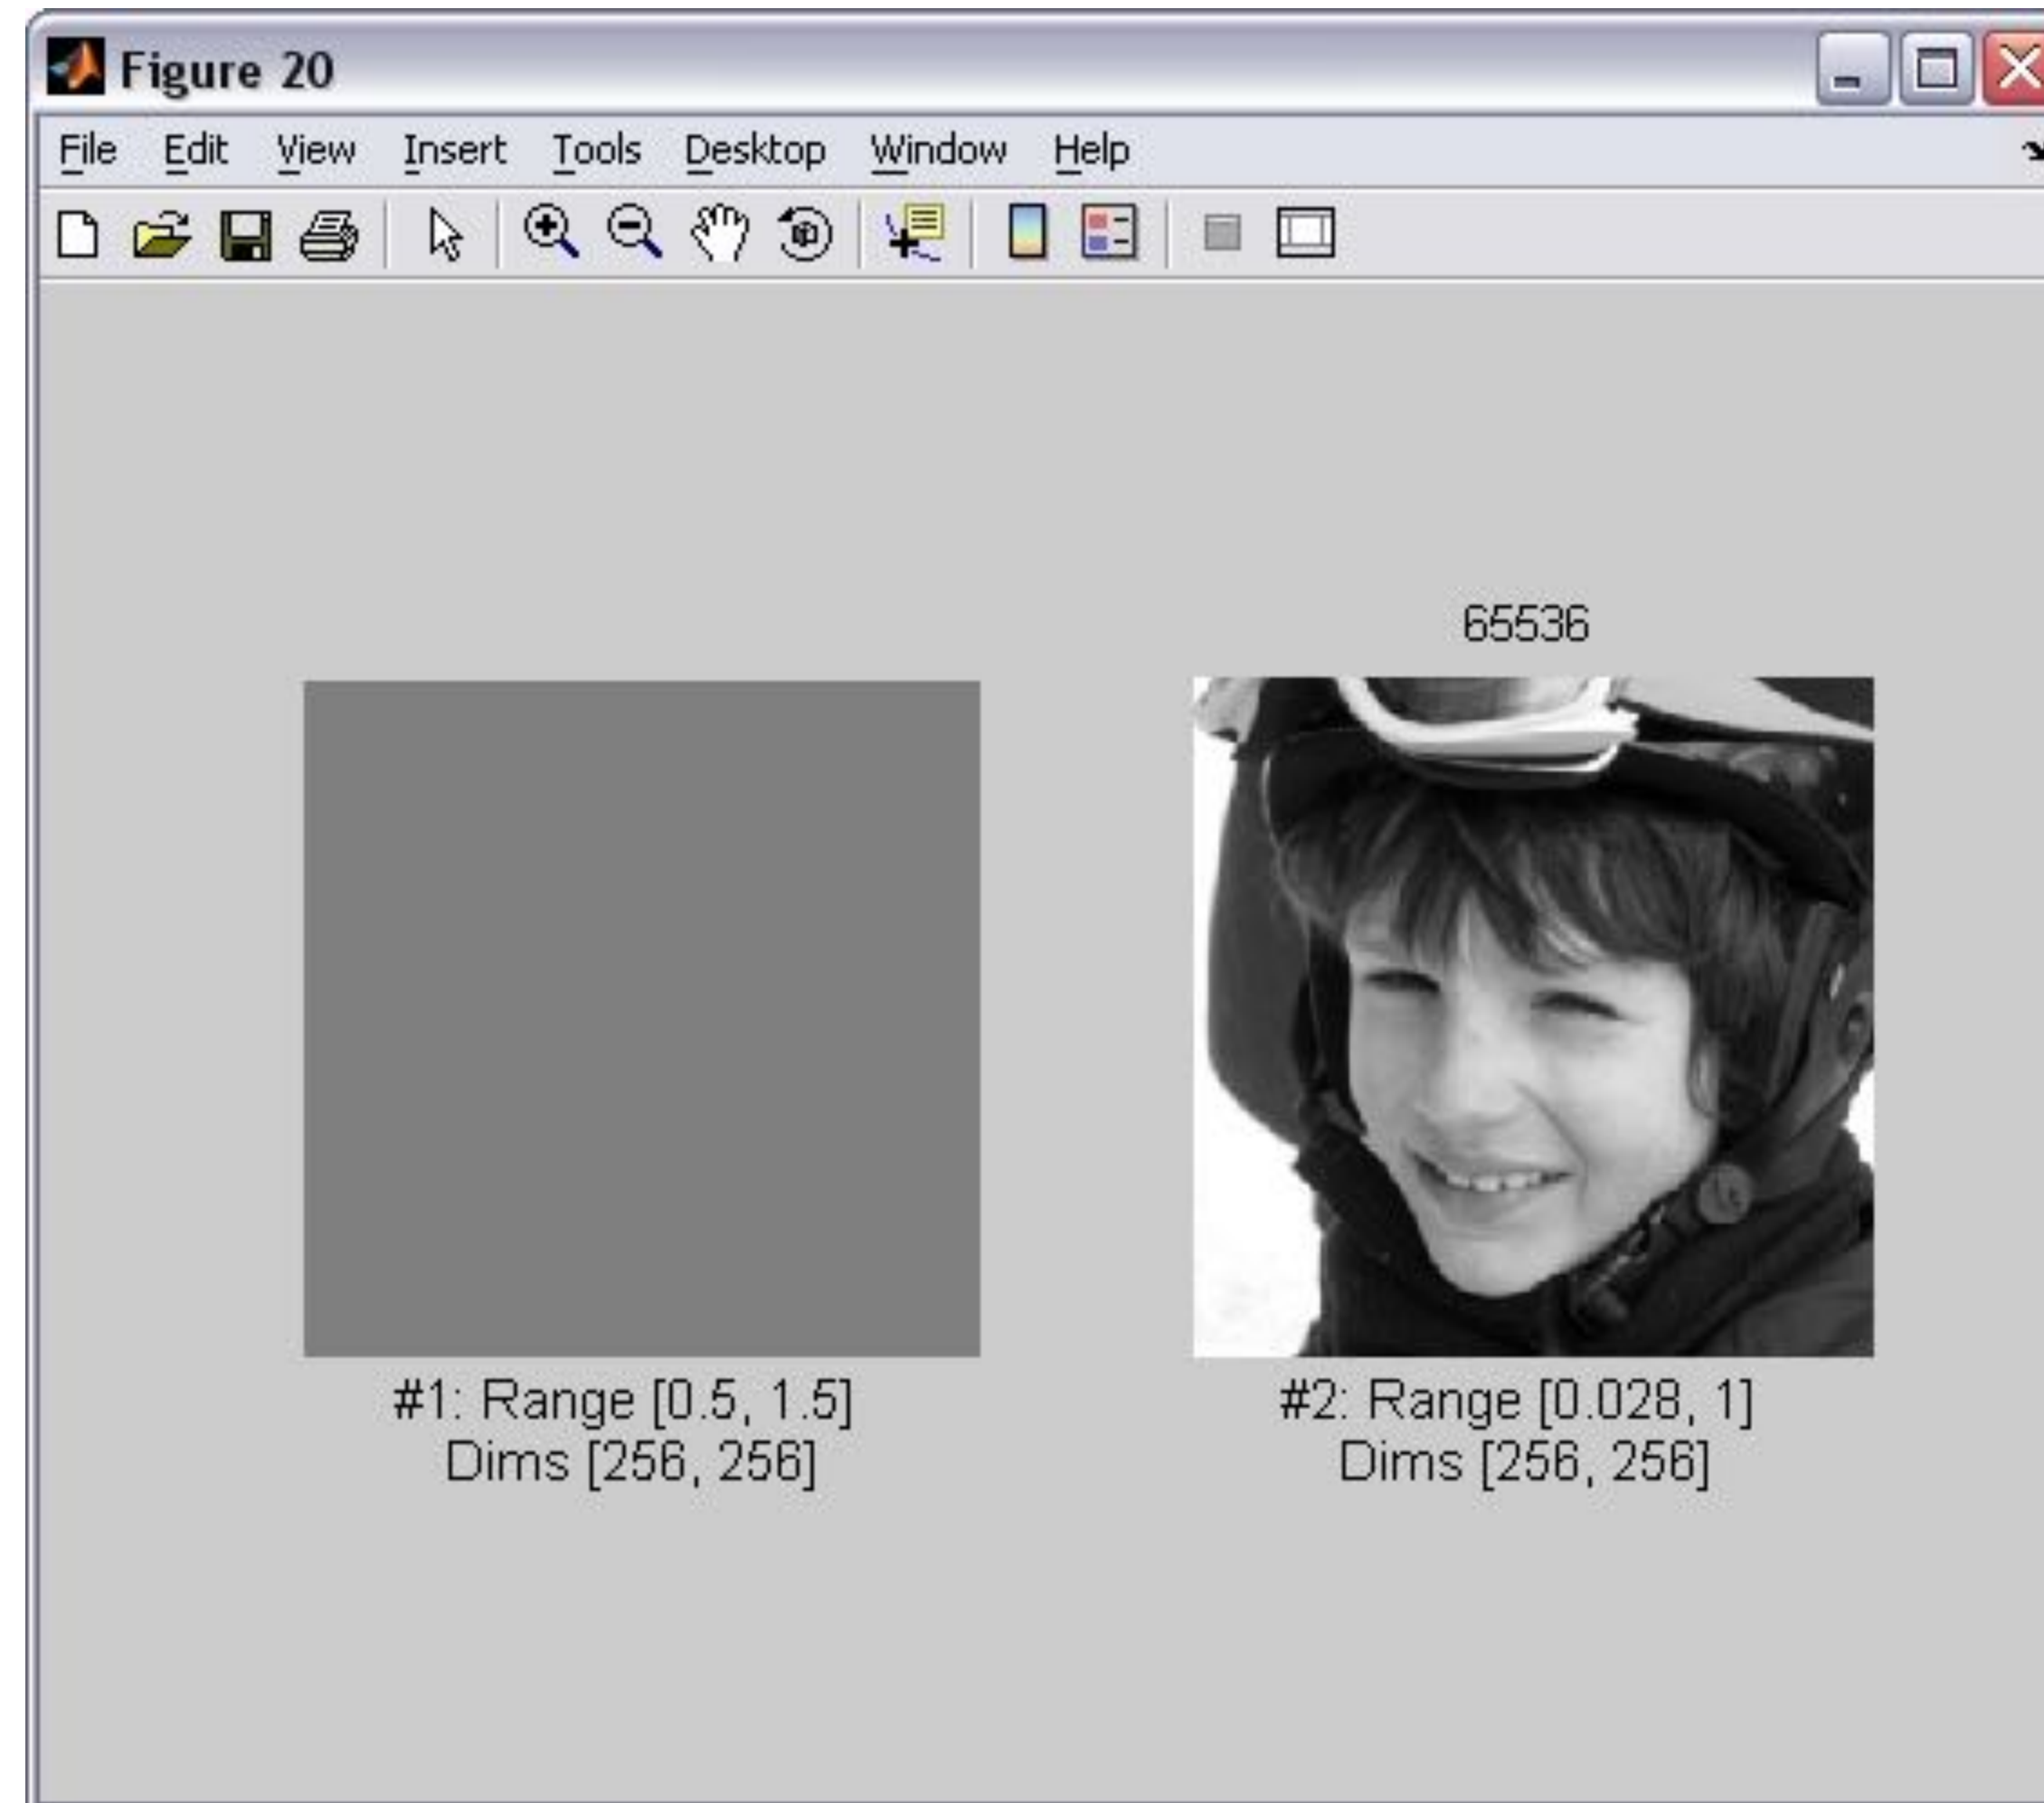
16385



32769

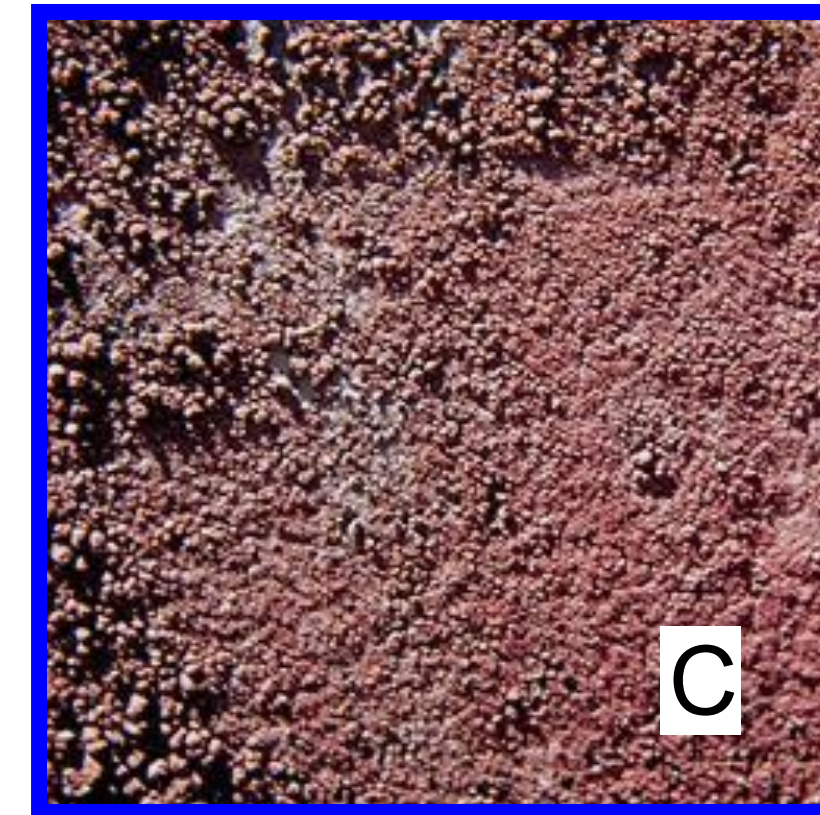
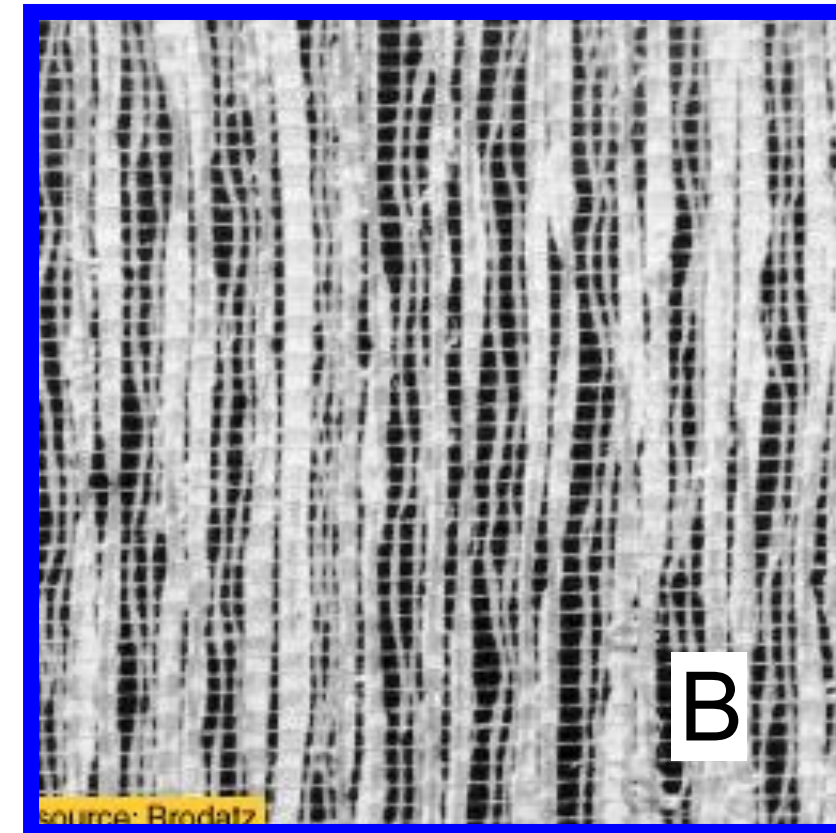


65536

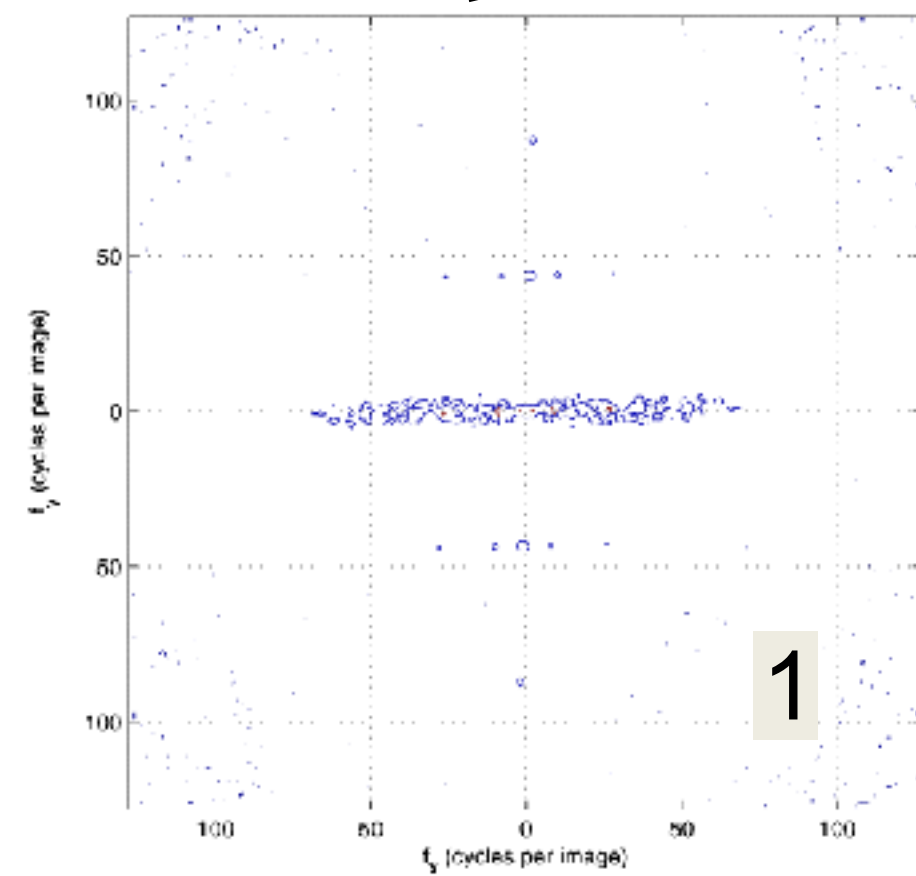


The DFT Game: find the right pairs

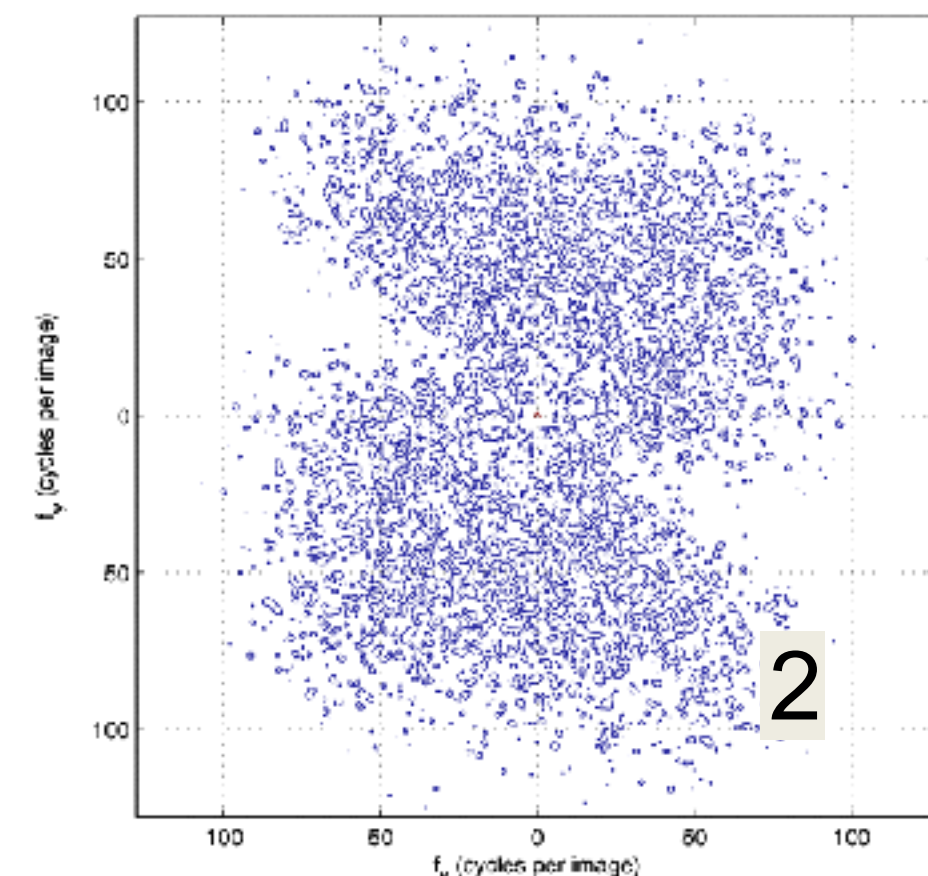
Images



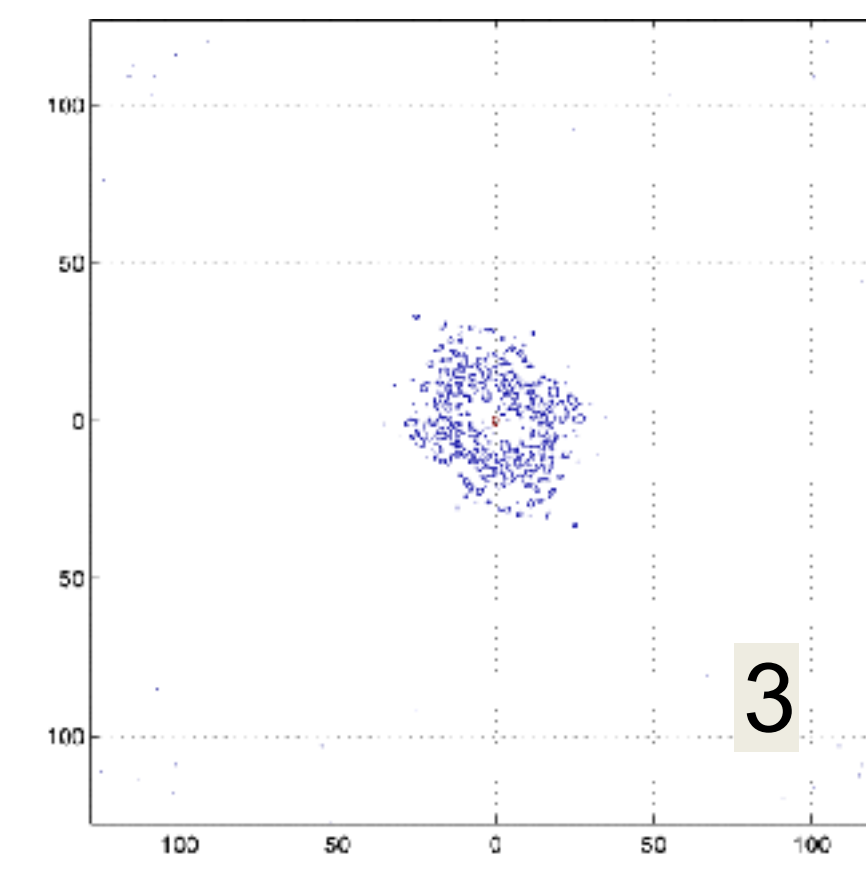
DFT
magnitude



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



a)



b)



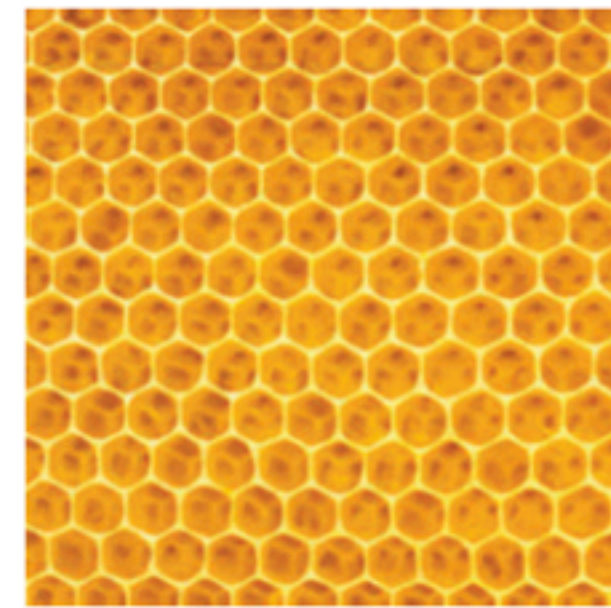
c)



d)



e)



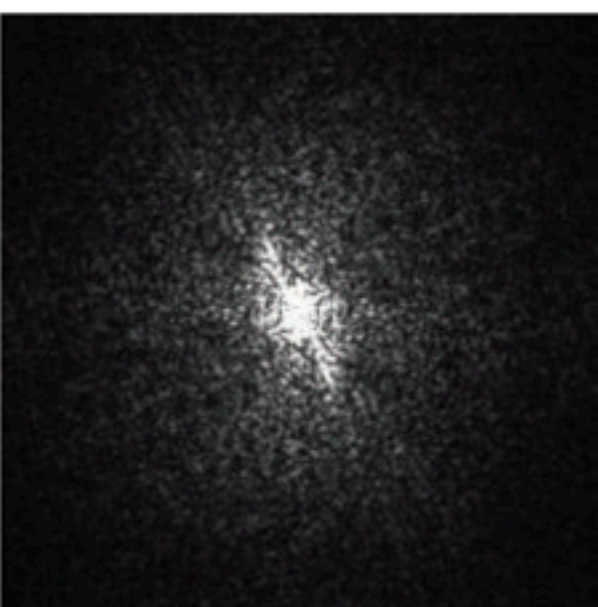
f)



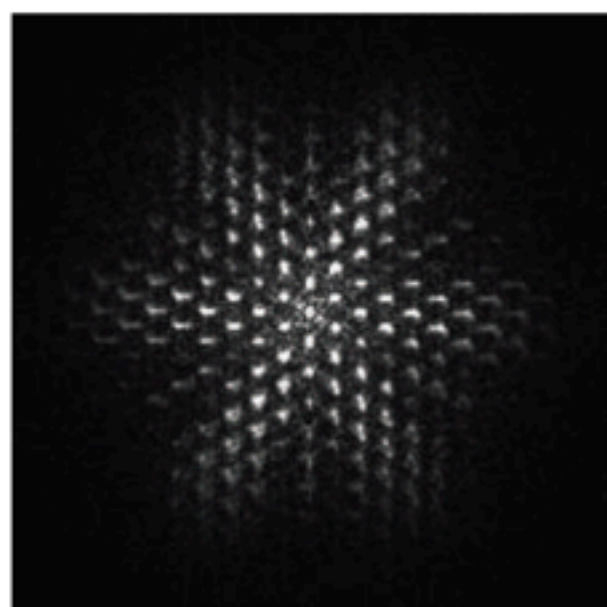
g)



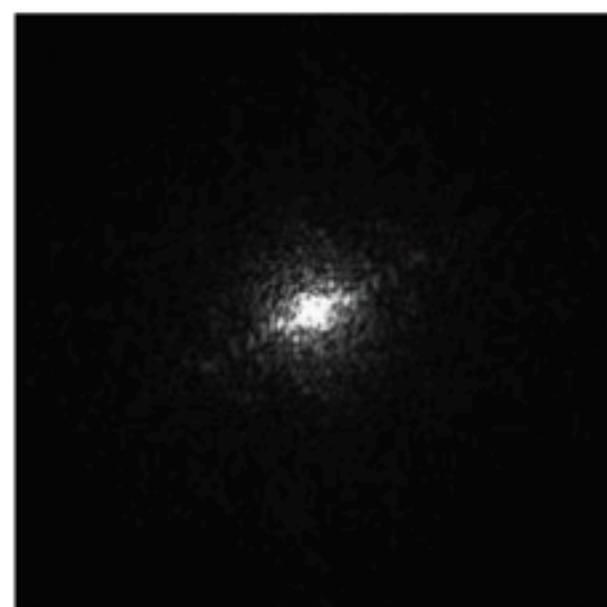
h)



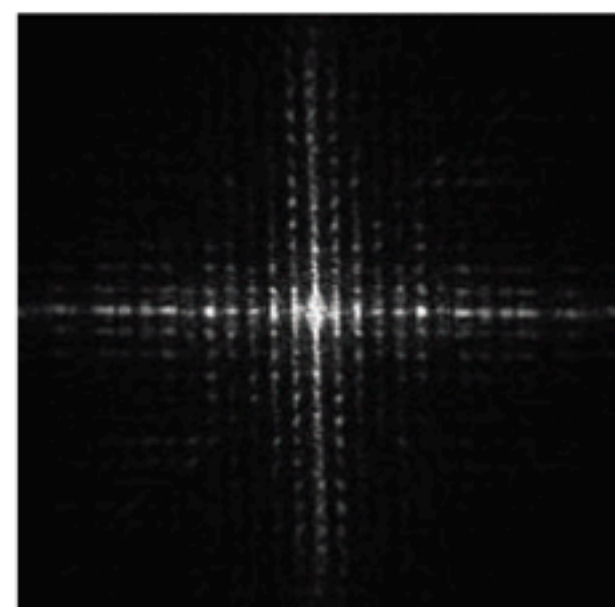
1)



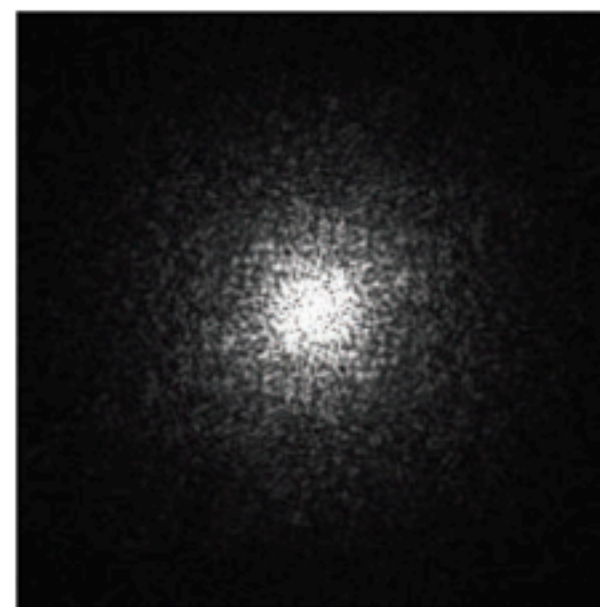
2)



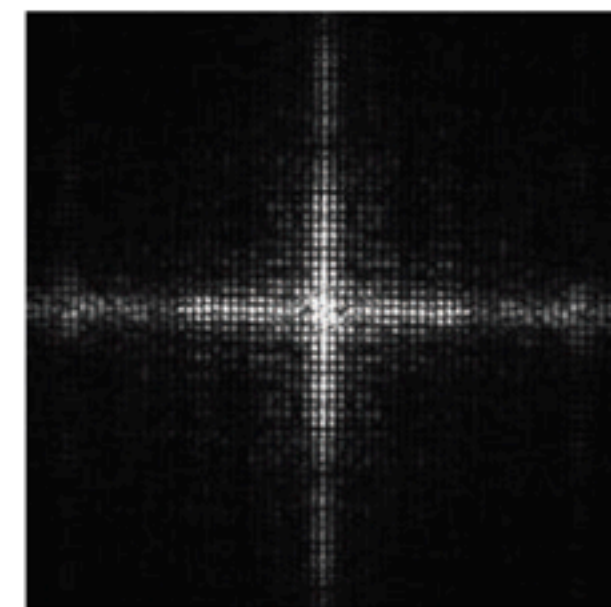
3)



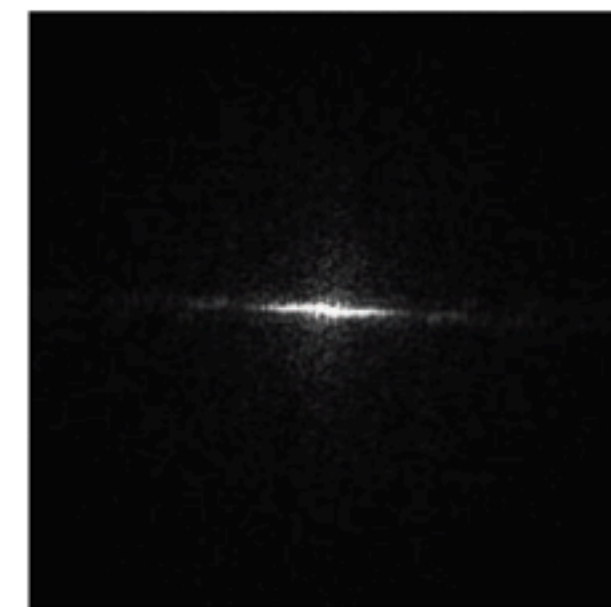
4)



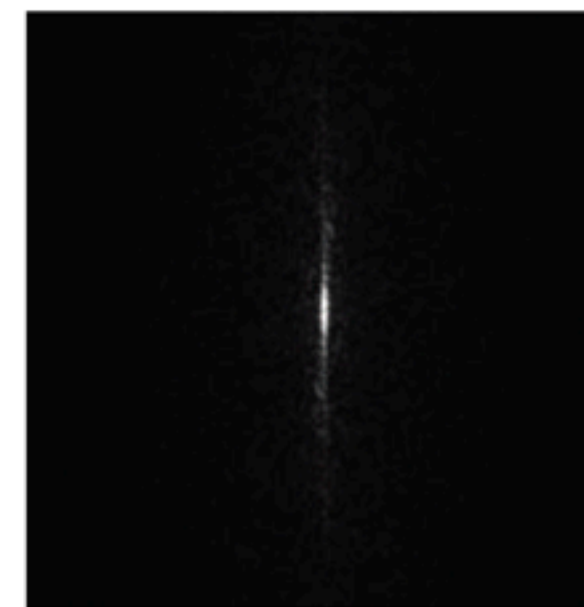
5)



6)



7)



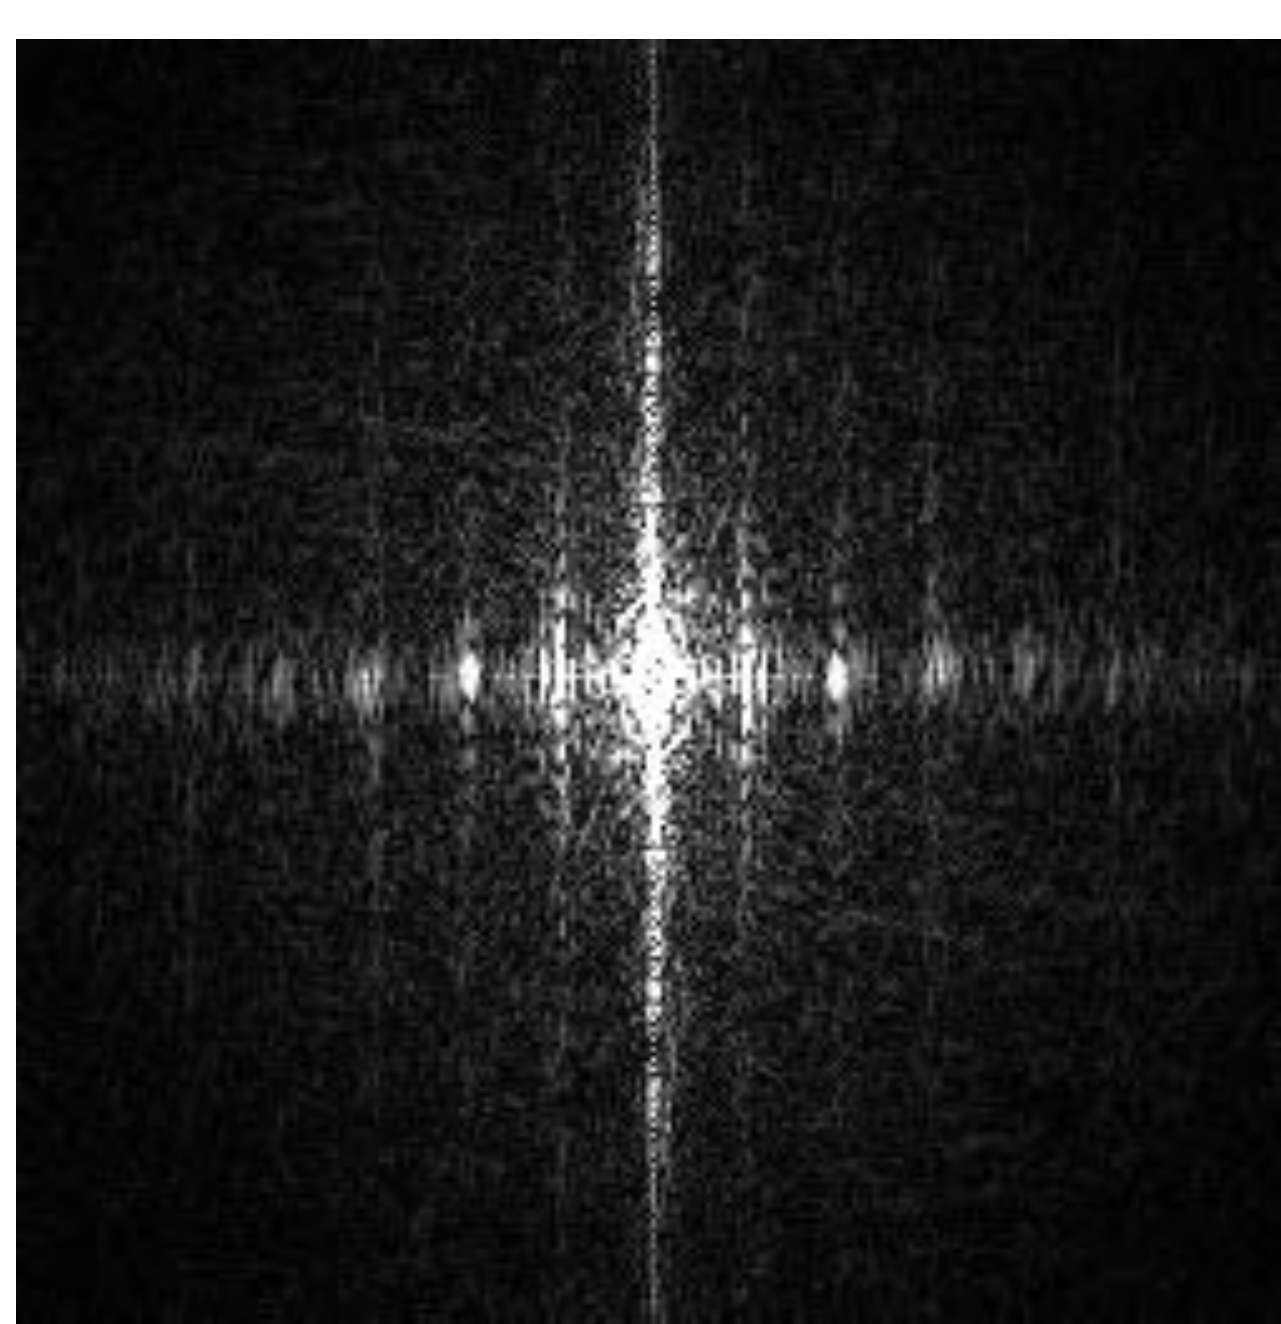
8)

(Solution in the class notes)

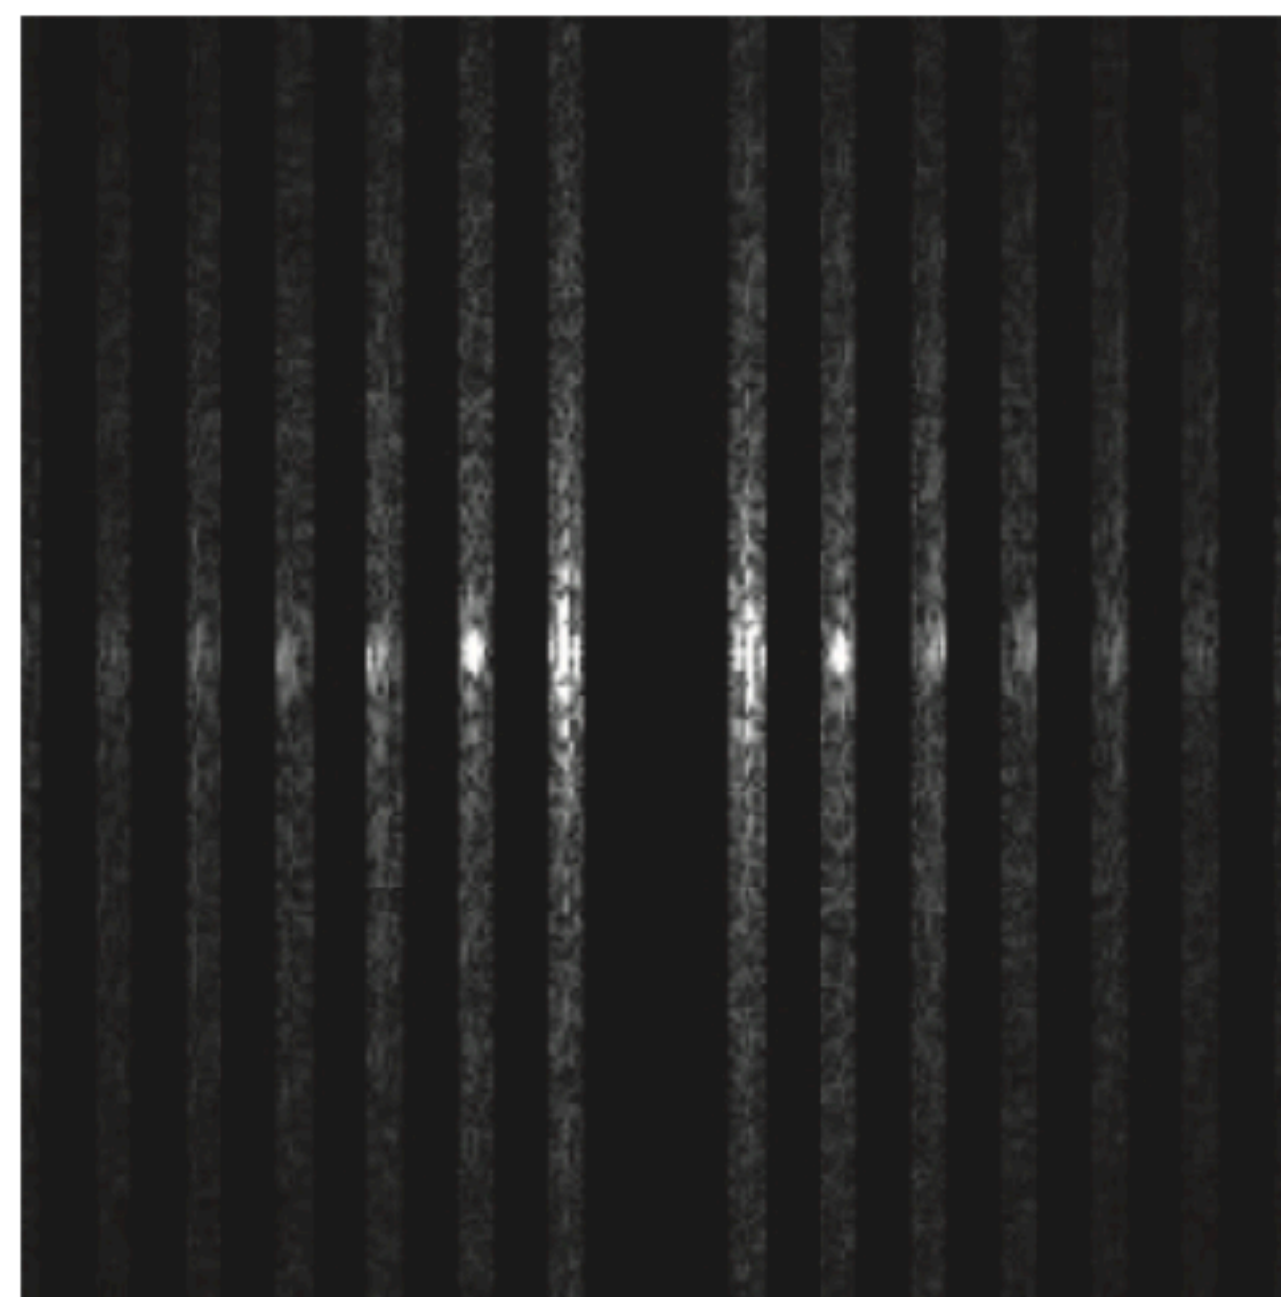


DFT
→

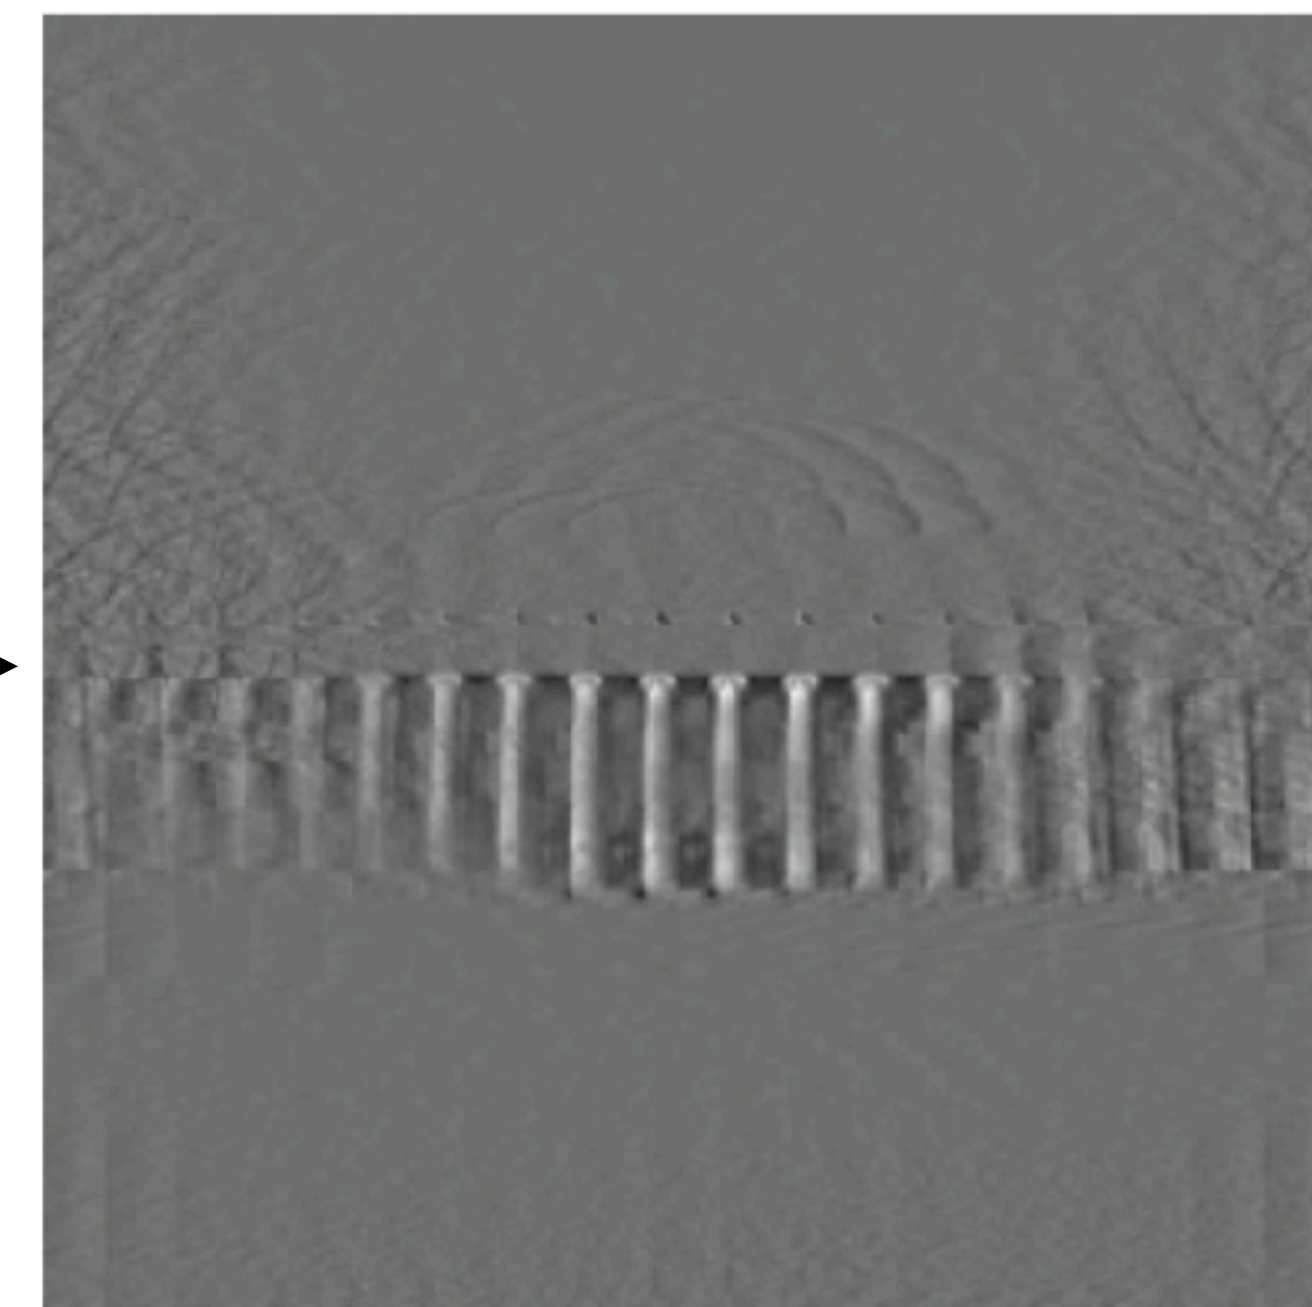
DFT squared
magnitude
(contrast
compressed
for visibility)



Retain Fourier
components
corresponding to
the repeating
columns

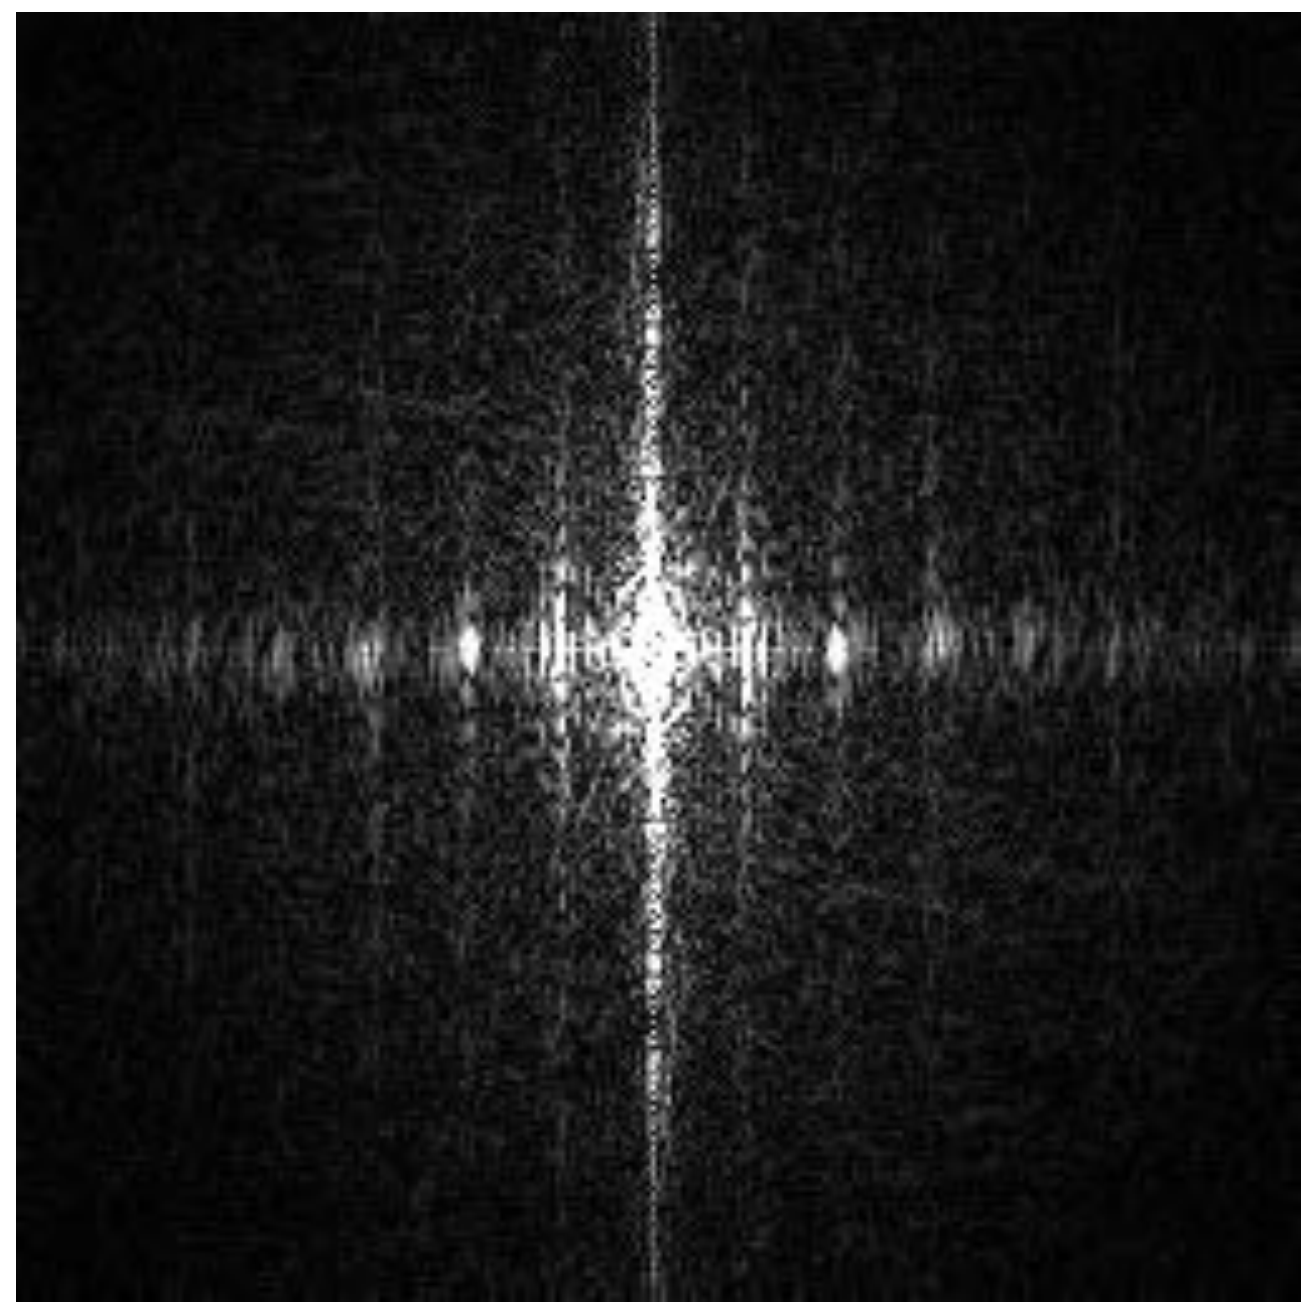


DFT⁻¹
→

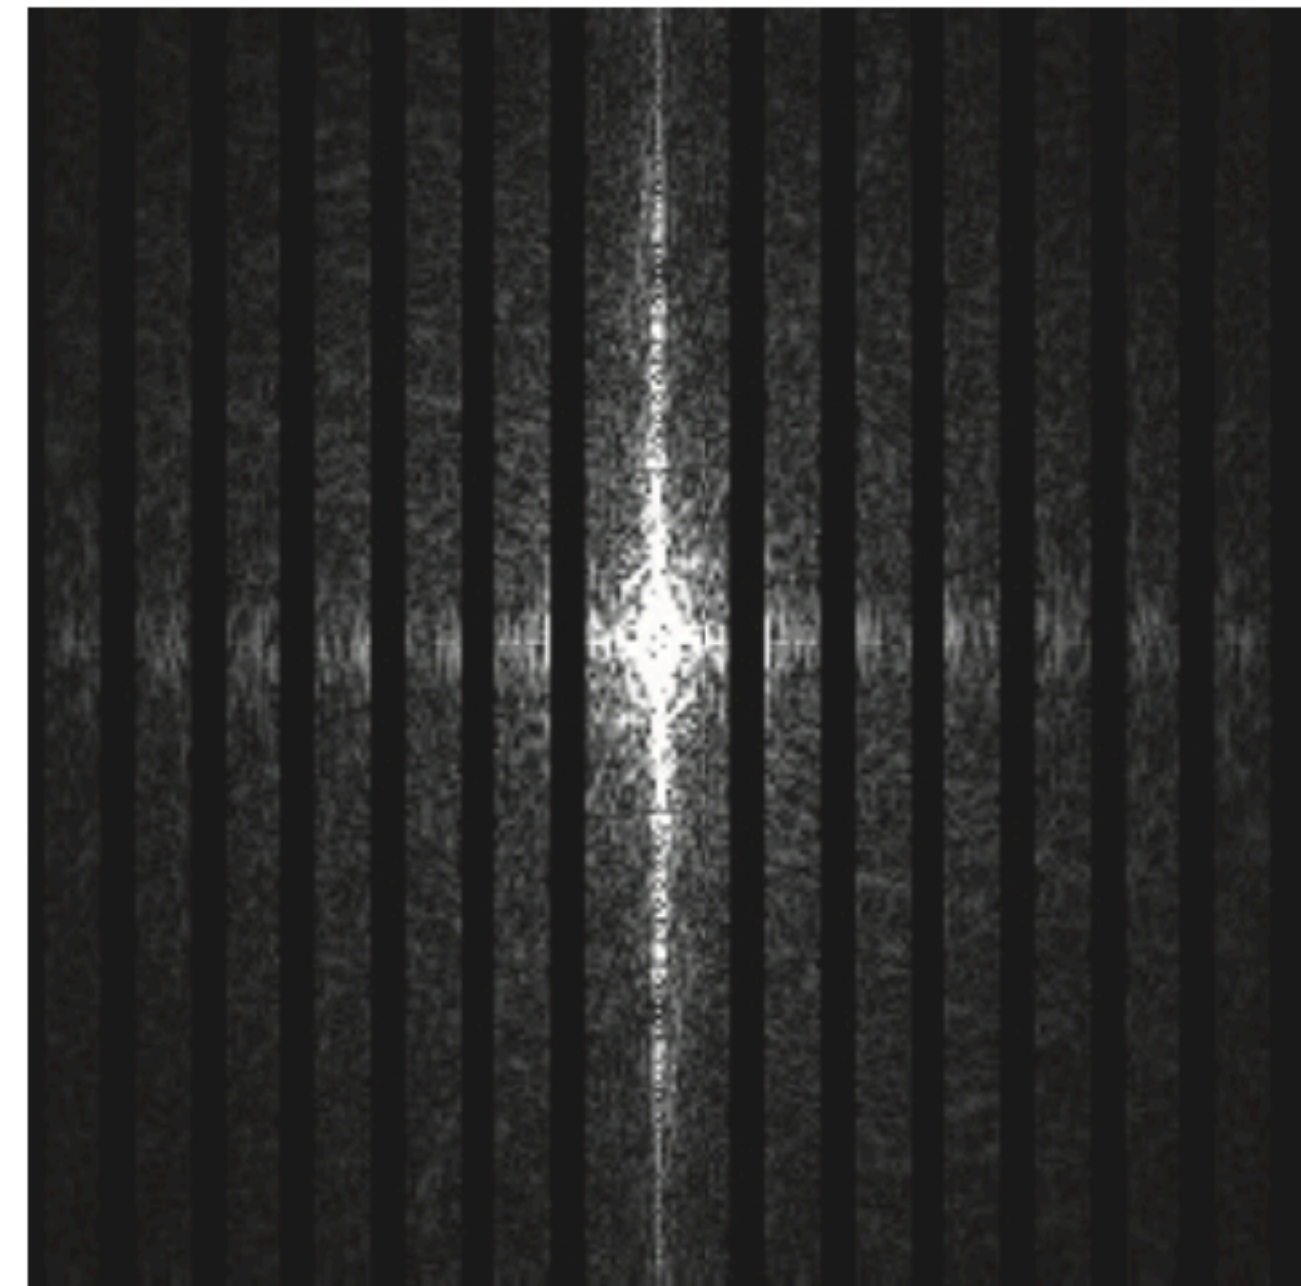




DFT
→



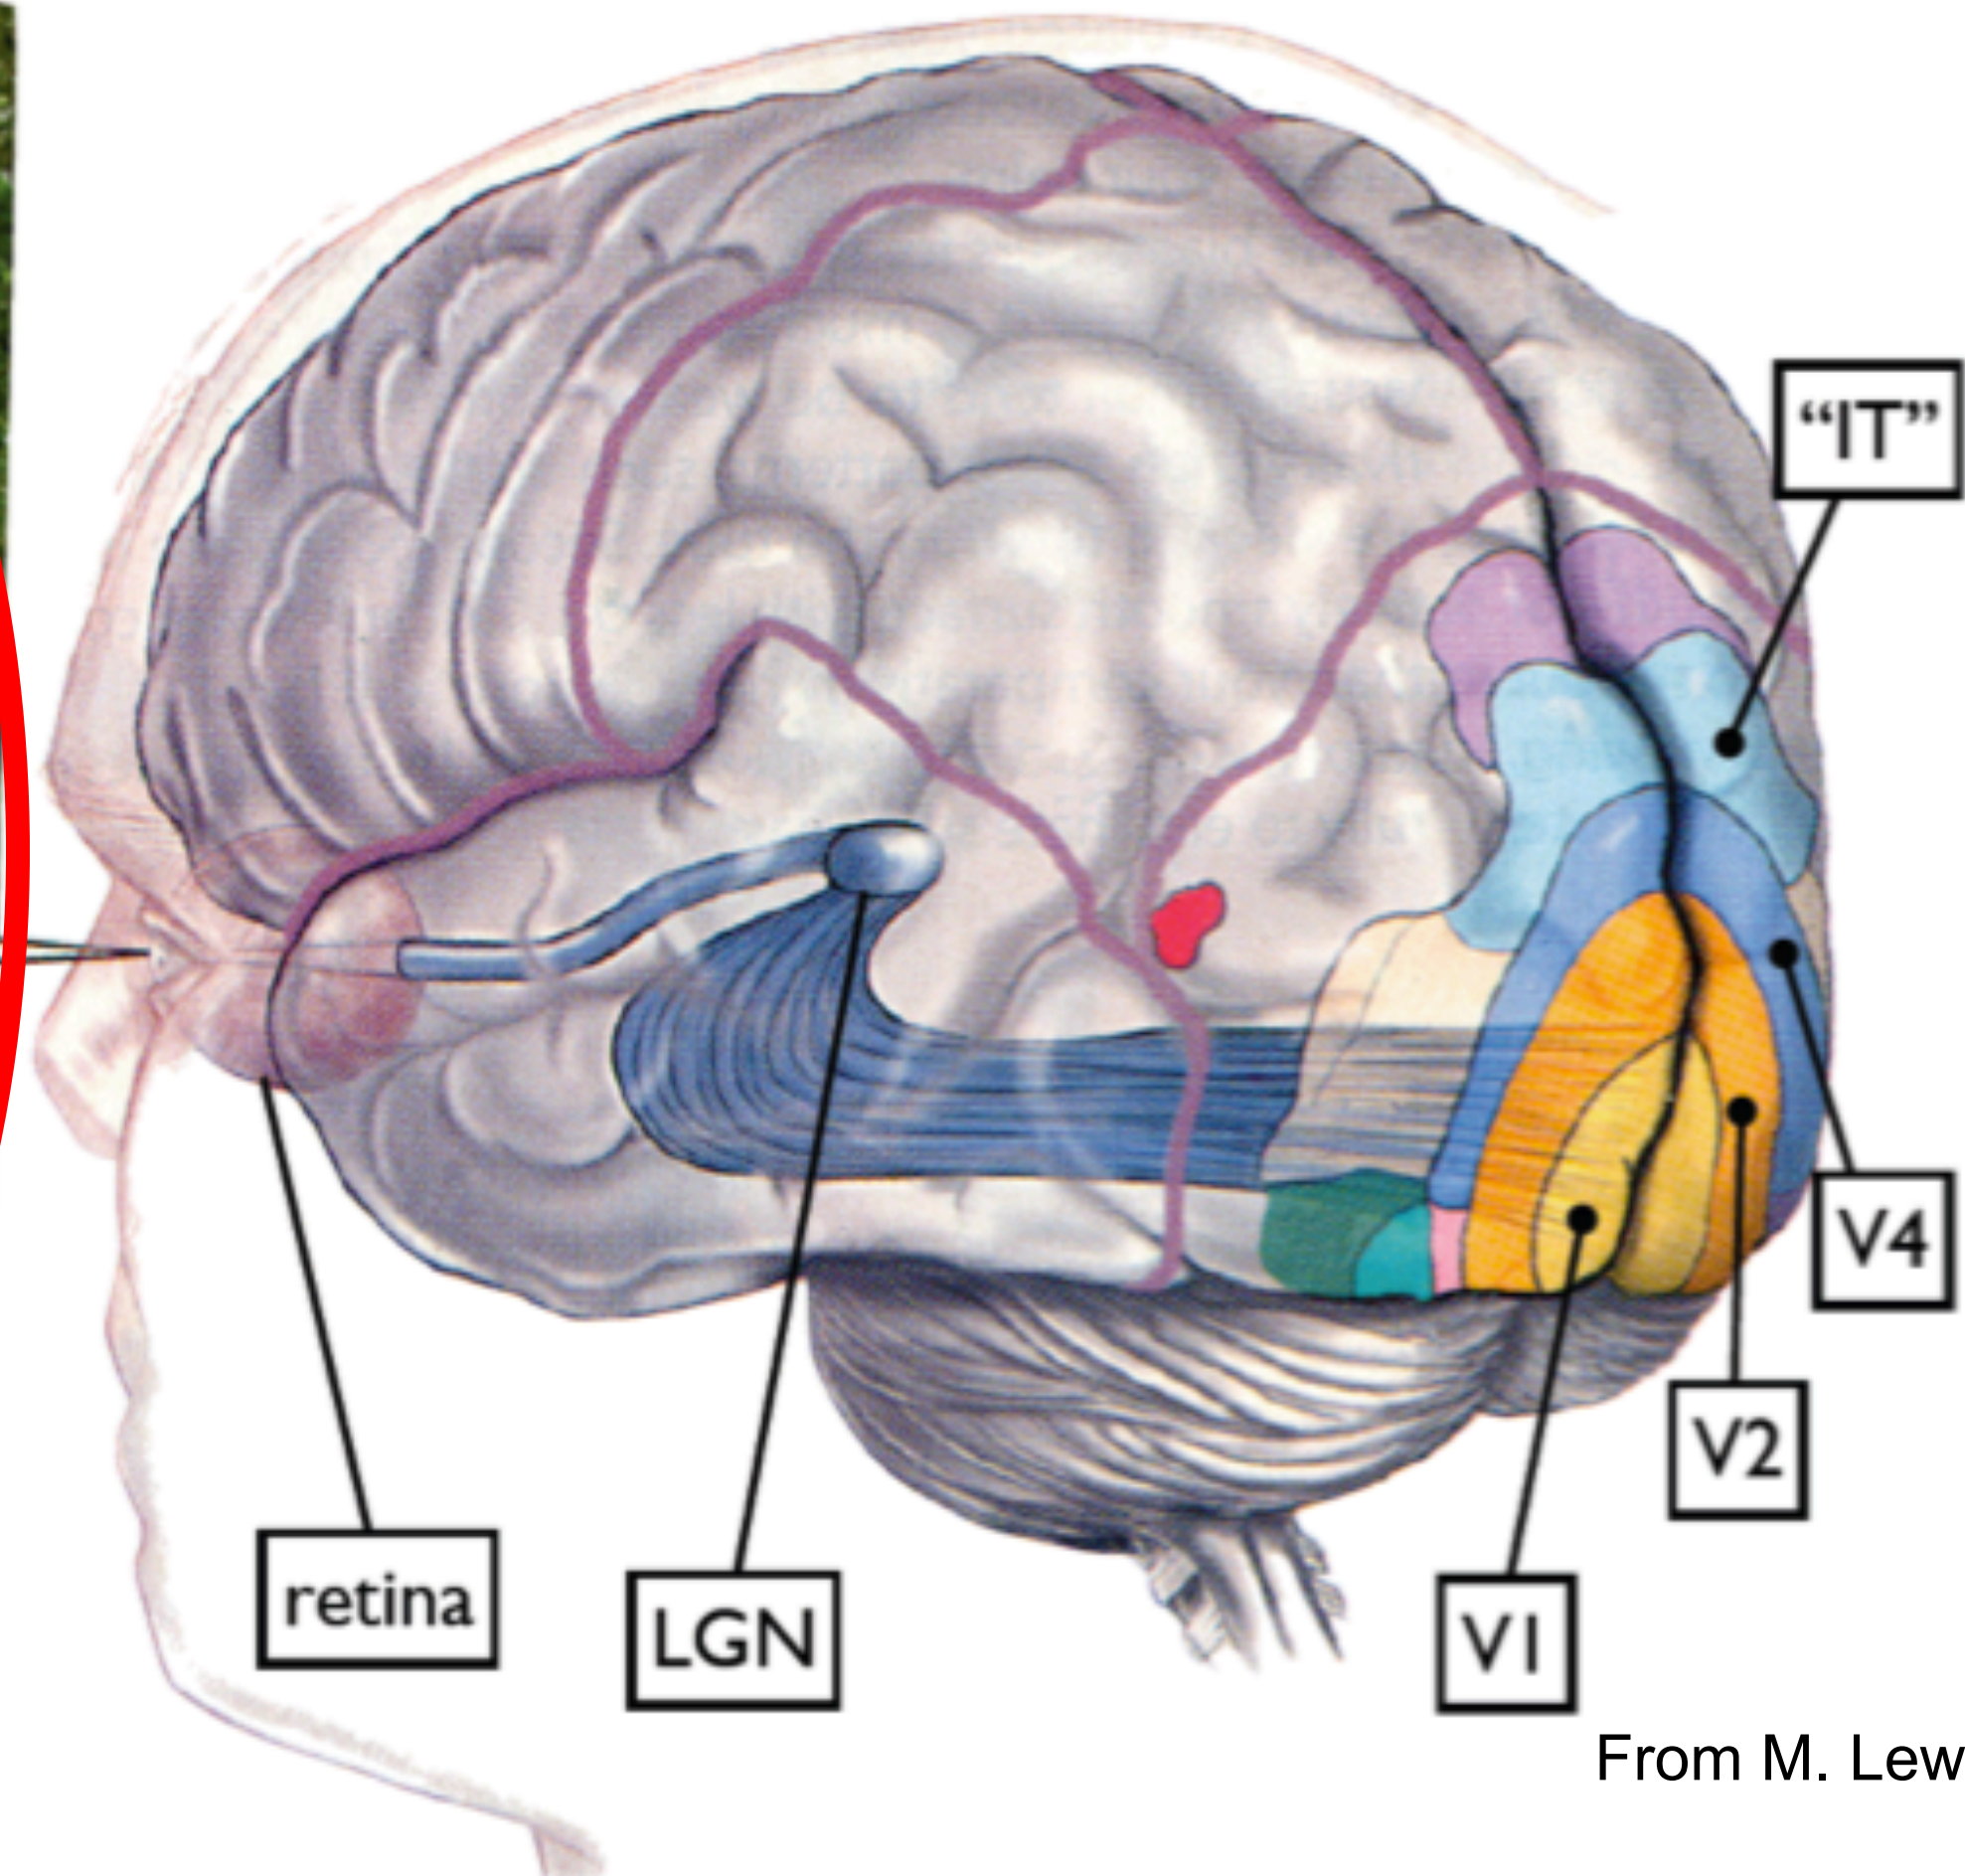
Remove Fourier components corresponding to the repeating columns



DFT⁻¹
→



Some visual areas...



From M. Lewicky

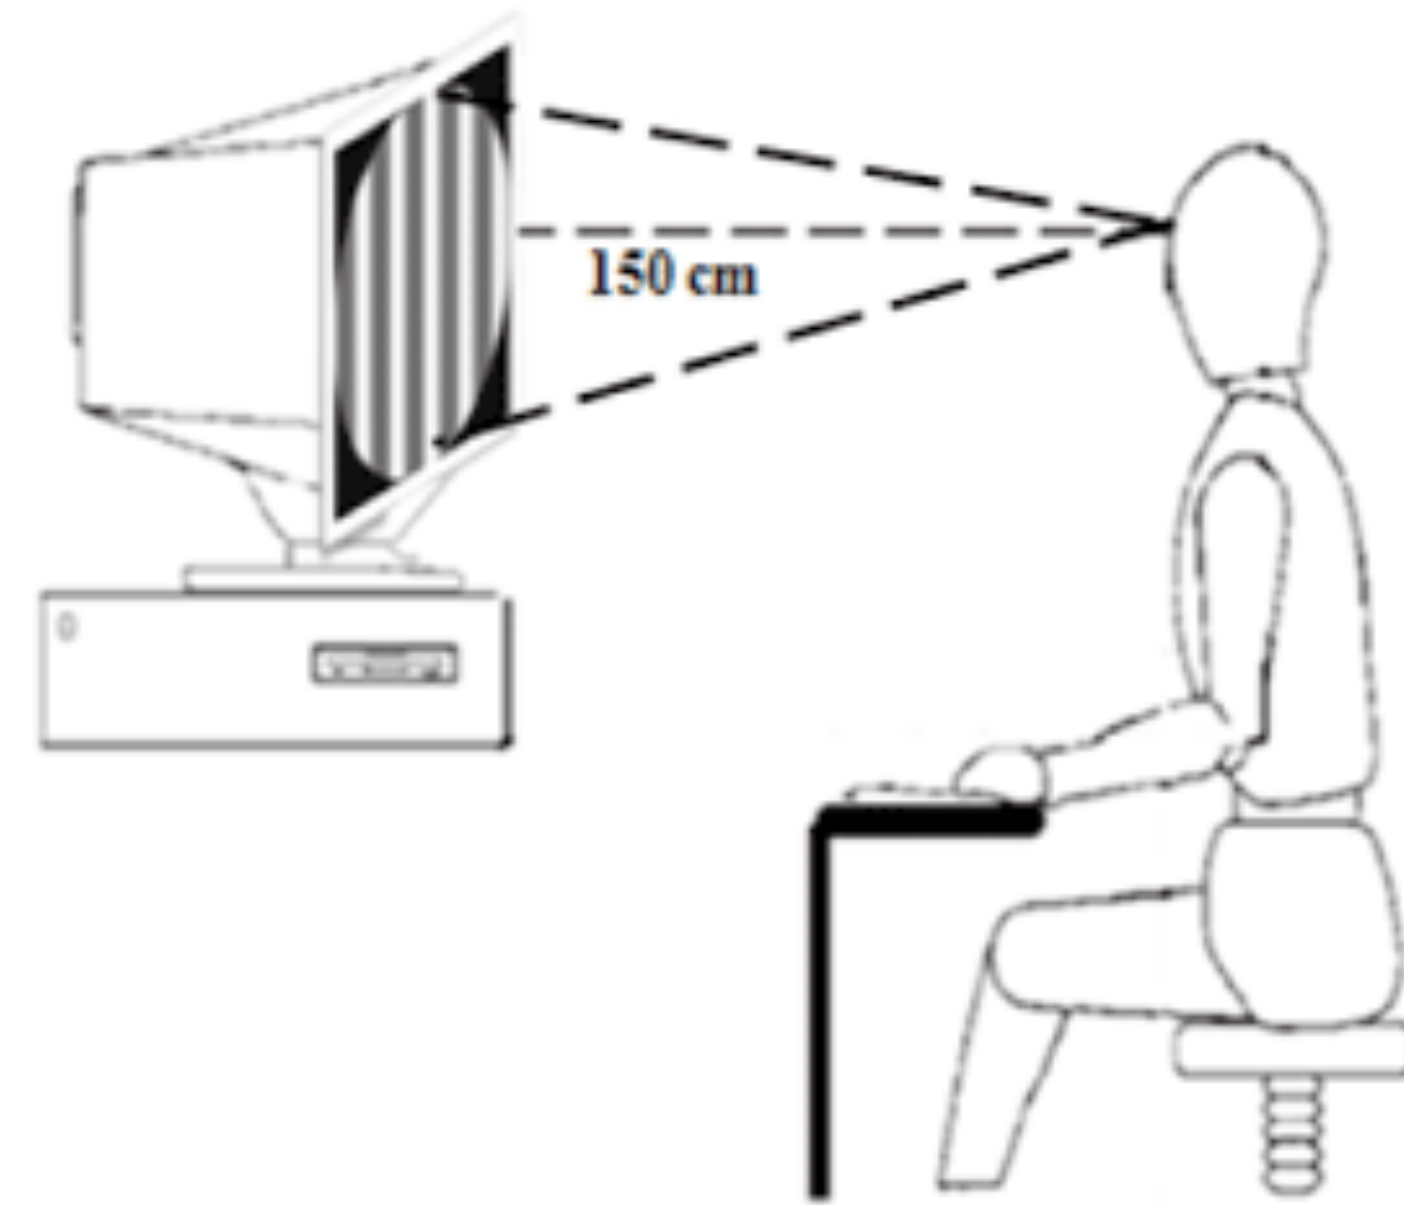
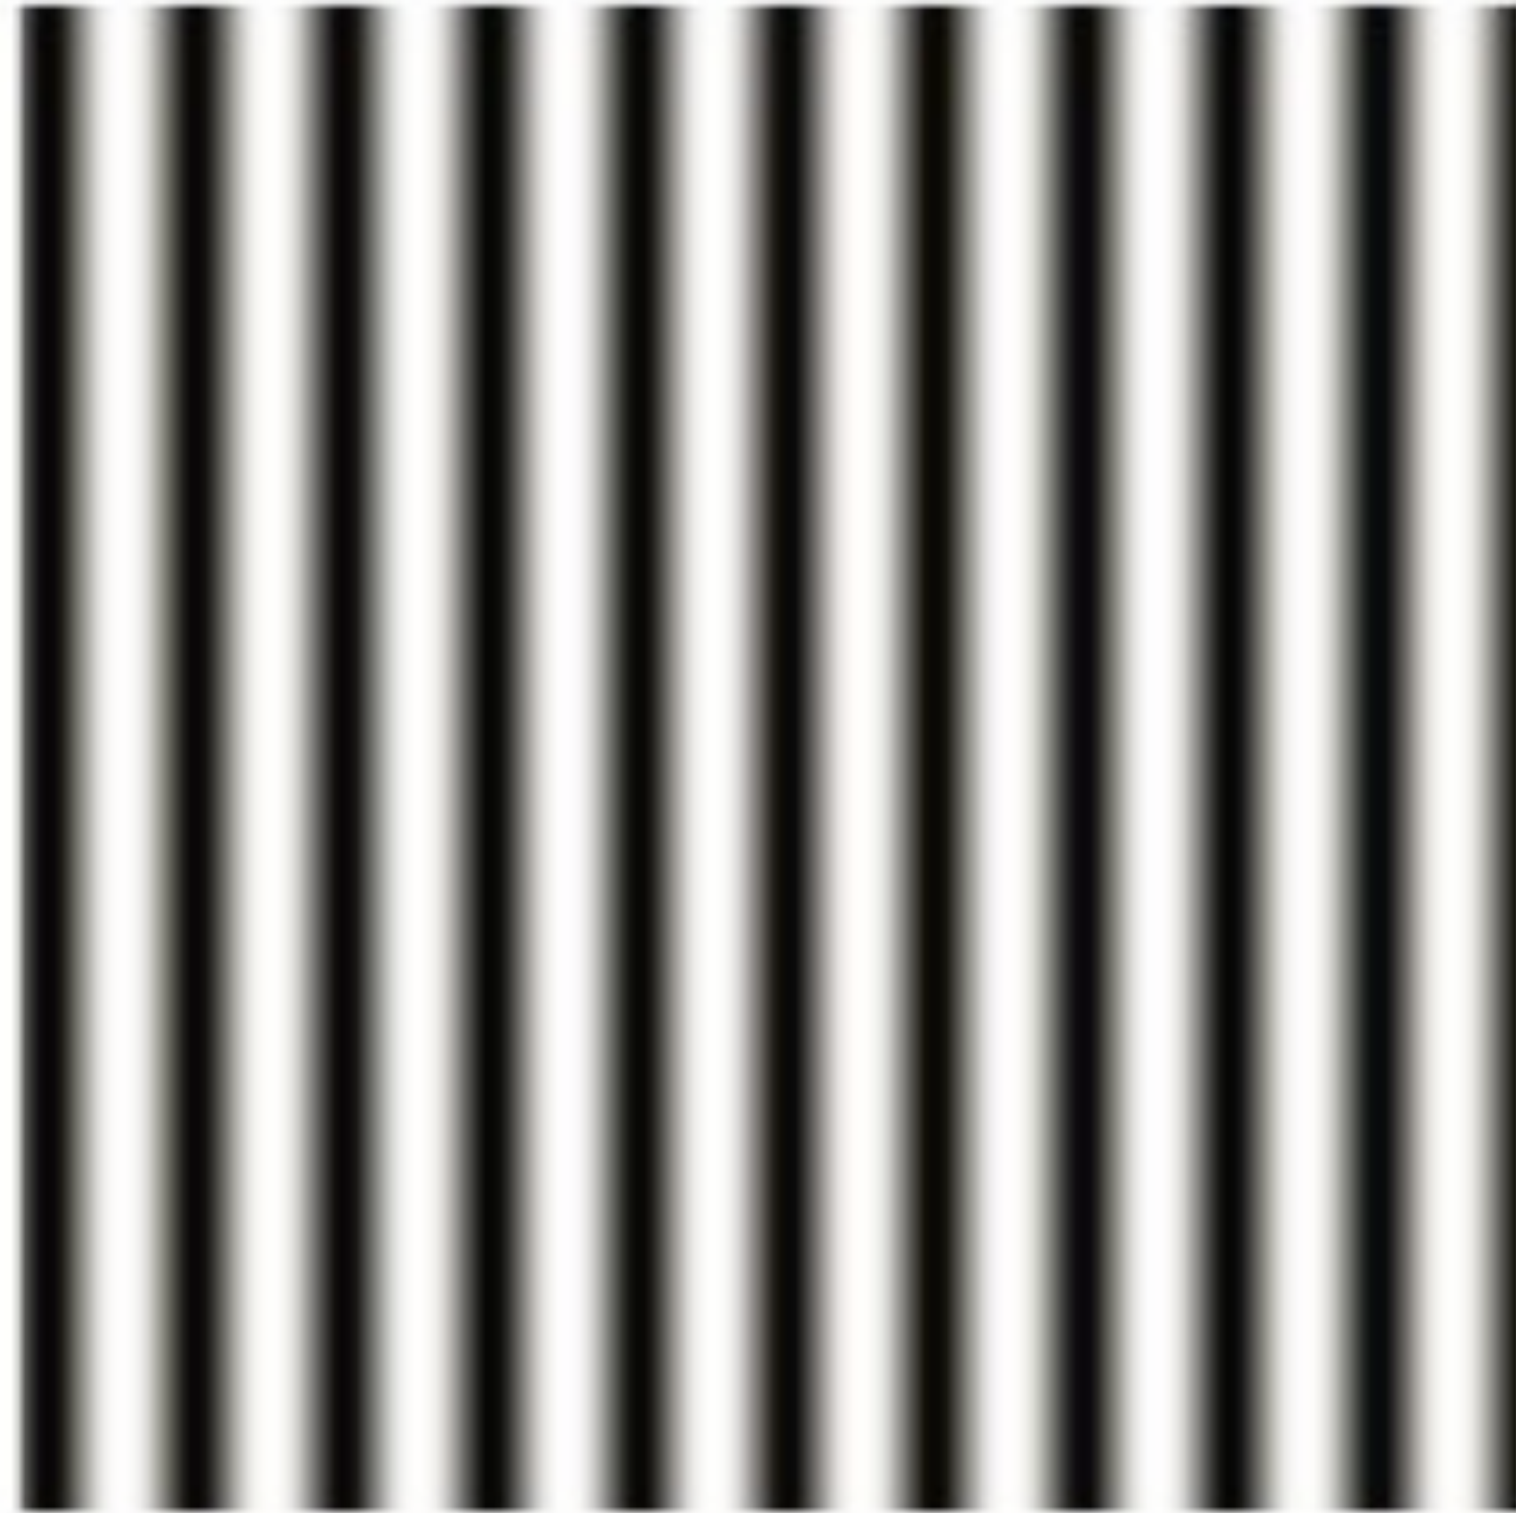
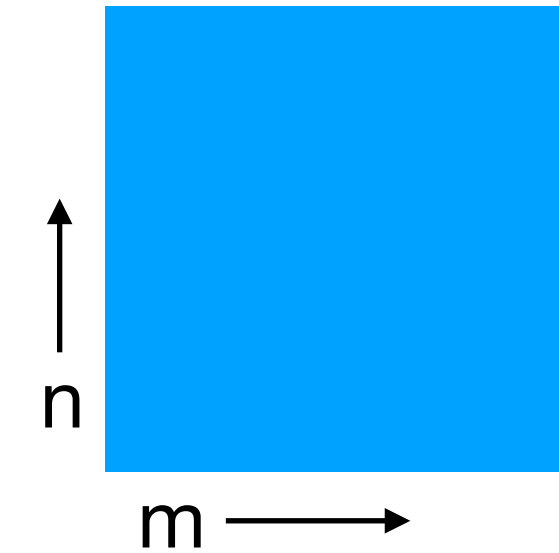


Figure 1. Stimulus presentation scheme. The stimuli were originally calibrated to be seen at a distance of 150 cm in a 19" display.

Campbell & Robson chart

Let's define the following image:

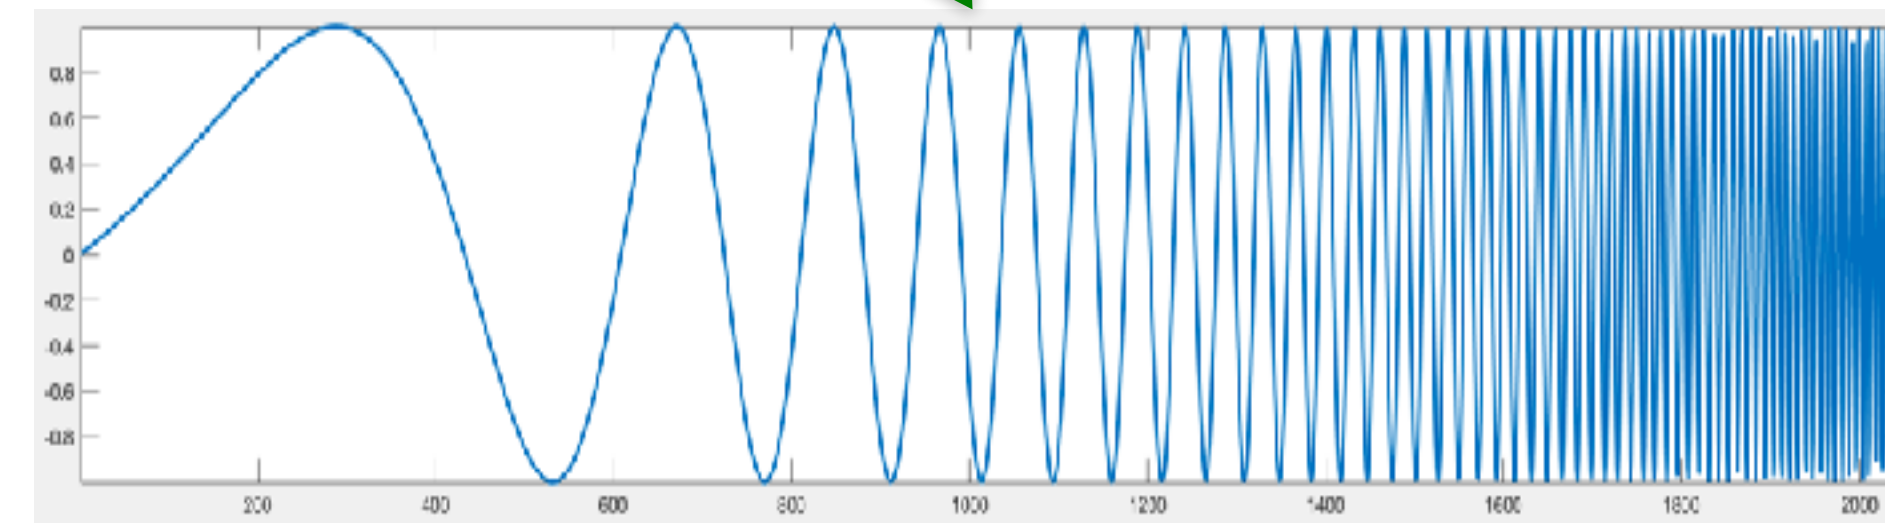
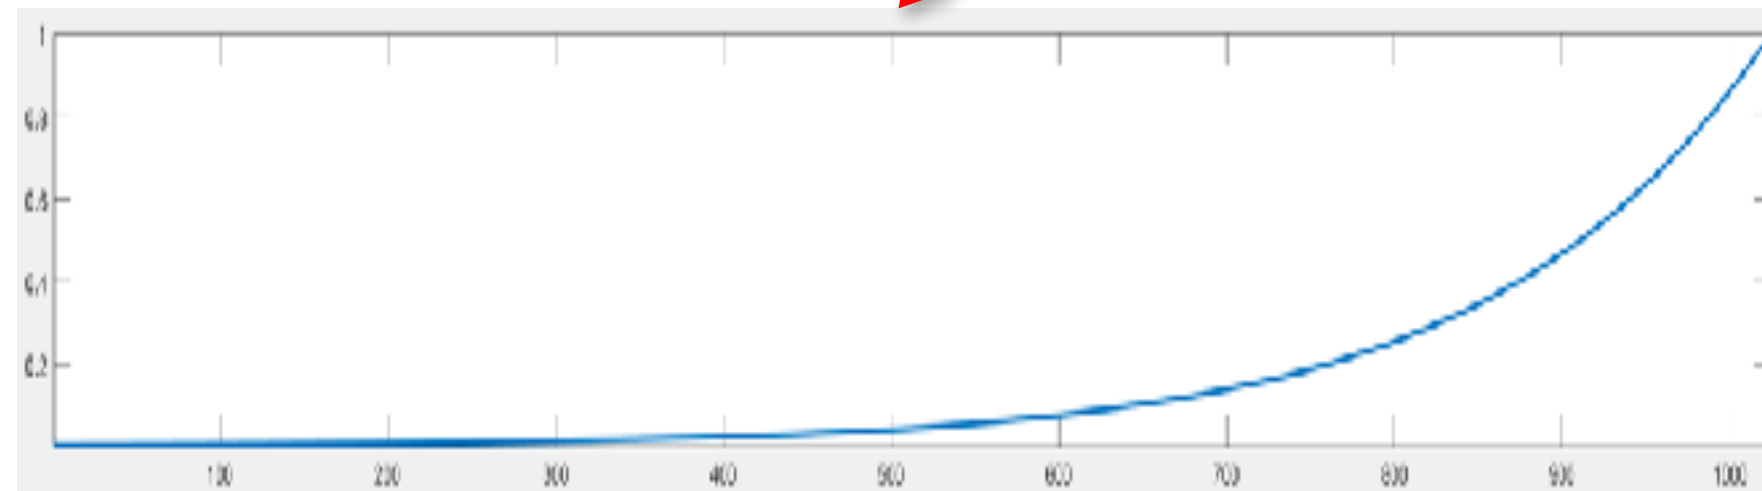
$$\mathbf{I}[n, m] = A[n] \sin(2\pi f[m] m/M)$$



With:

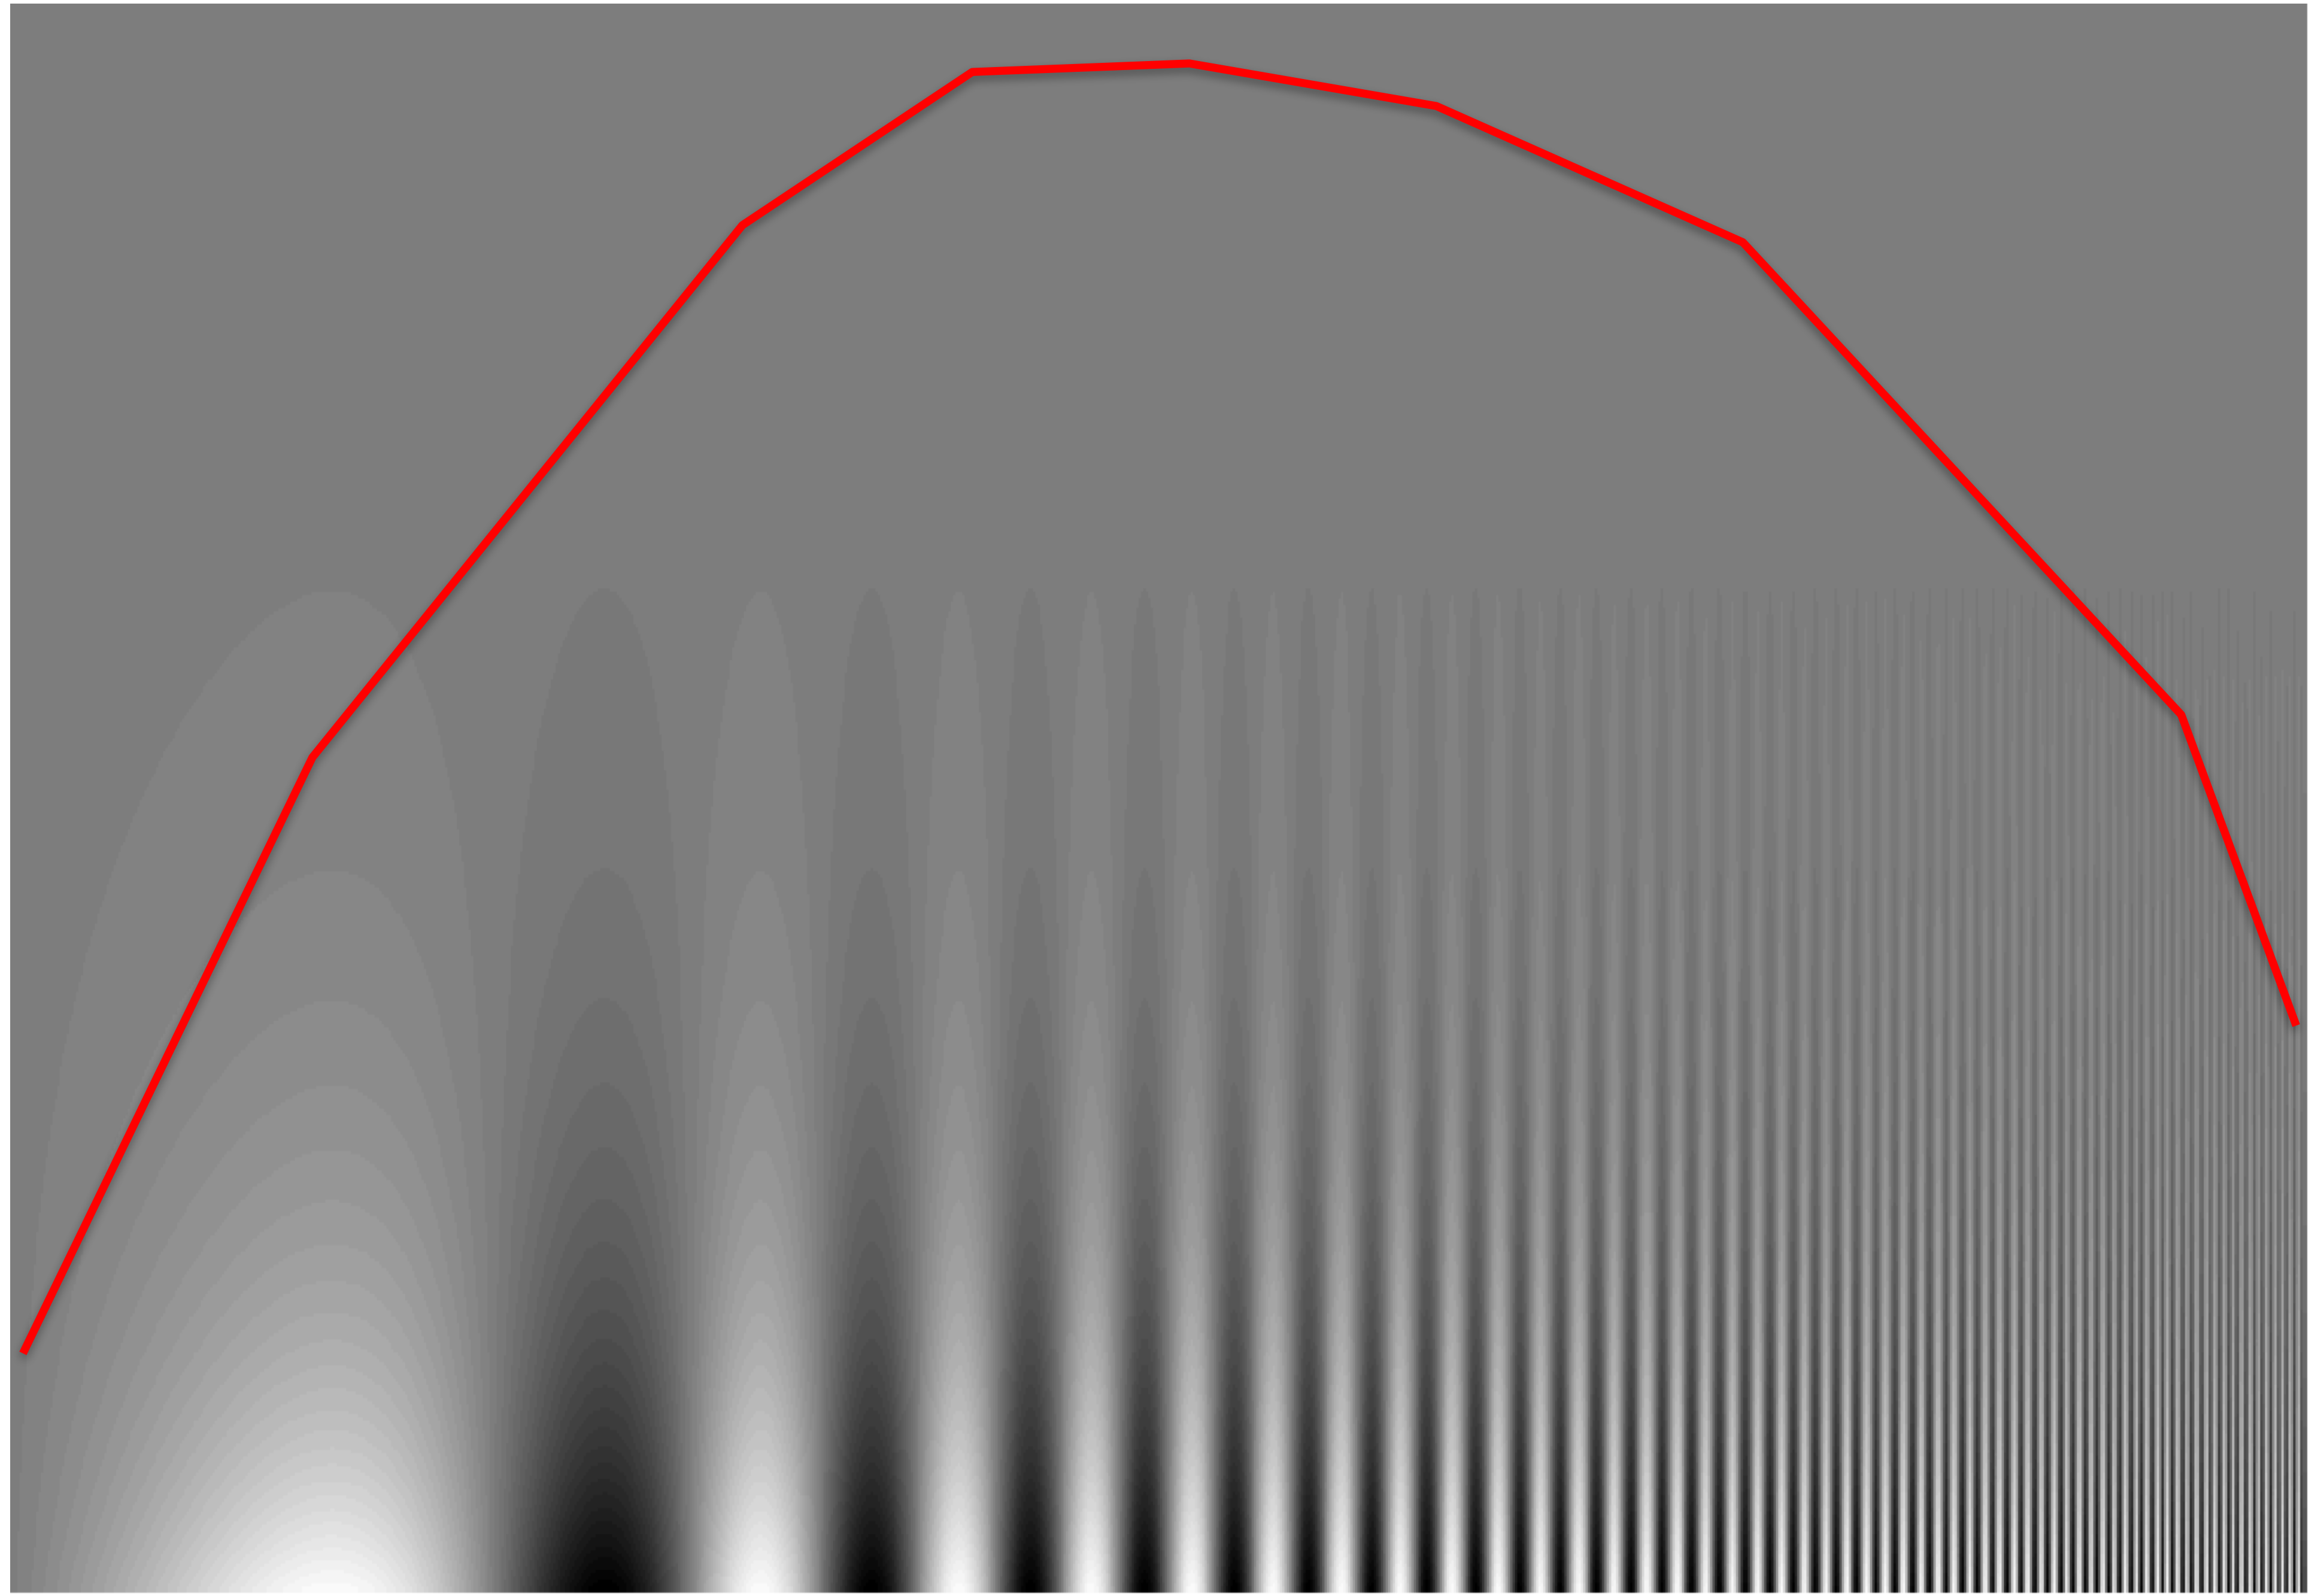
$$A[n] = A_{min} \left(\frac{A_{max}}{A_{min}} \right)^{n/N}$$

$$f[m] = f_{min} \left(\frac{f_{max}}{f_{min}} \right)^{m/M}$$



What do you think you should see when looking at this image?

$$\mathbf{I}[n, m] = A[n] \sin(2\pi f[m] m/M)$$

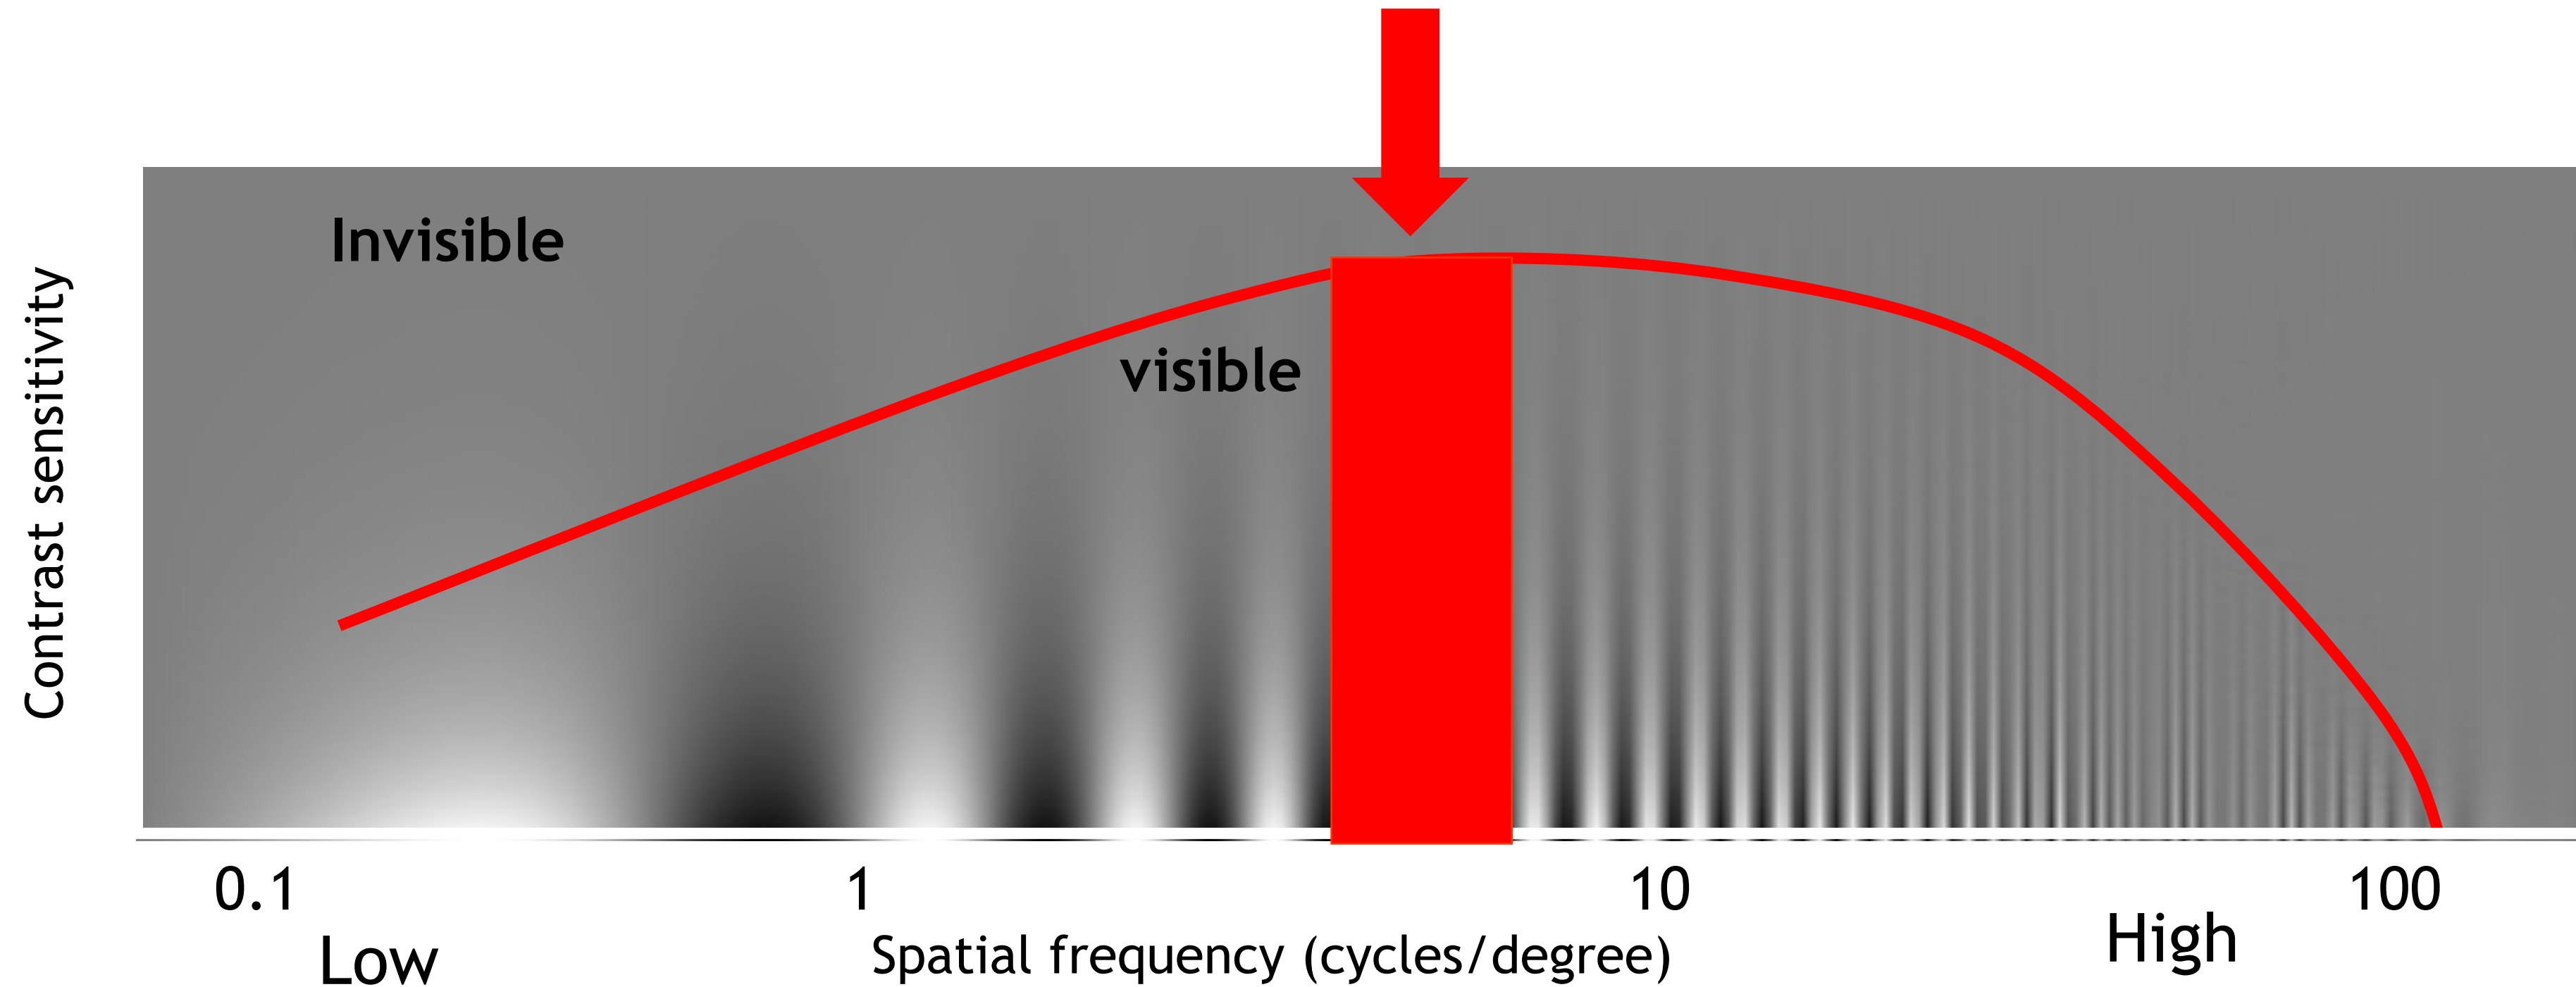


Contrast Sensitivity Function

Blackmore & Campbell (1969)

Maximum sensitivity

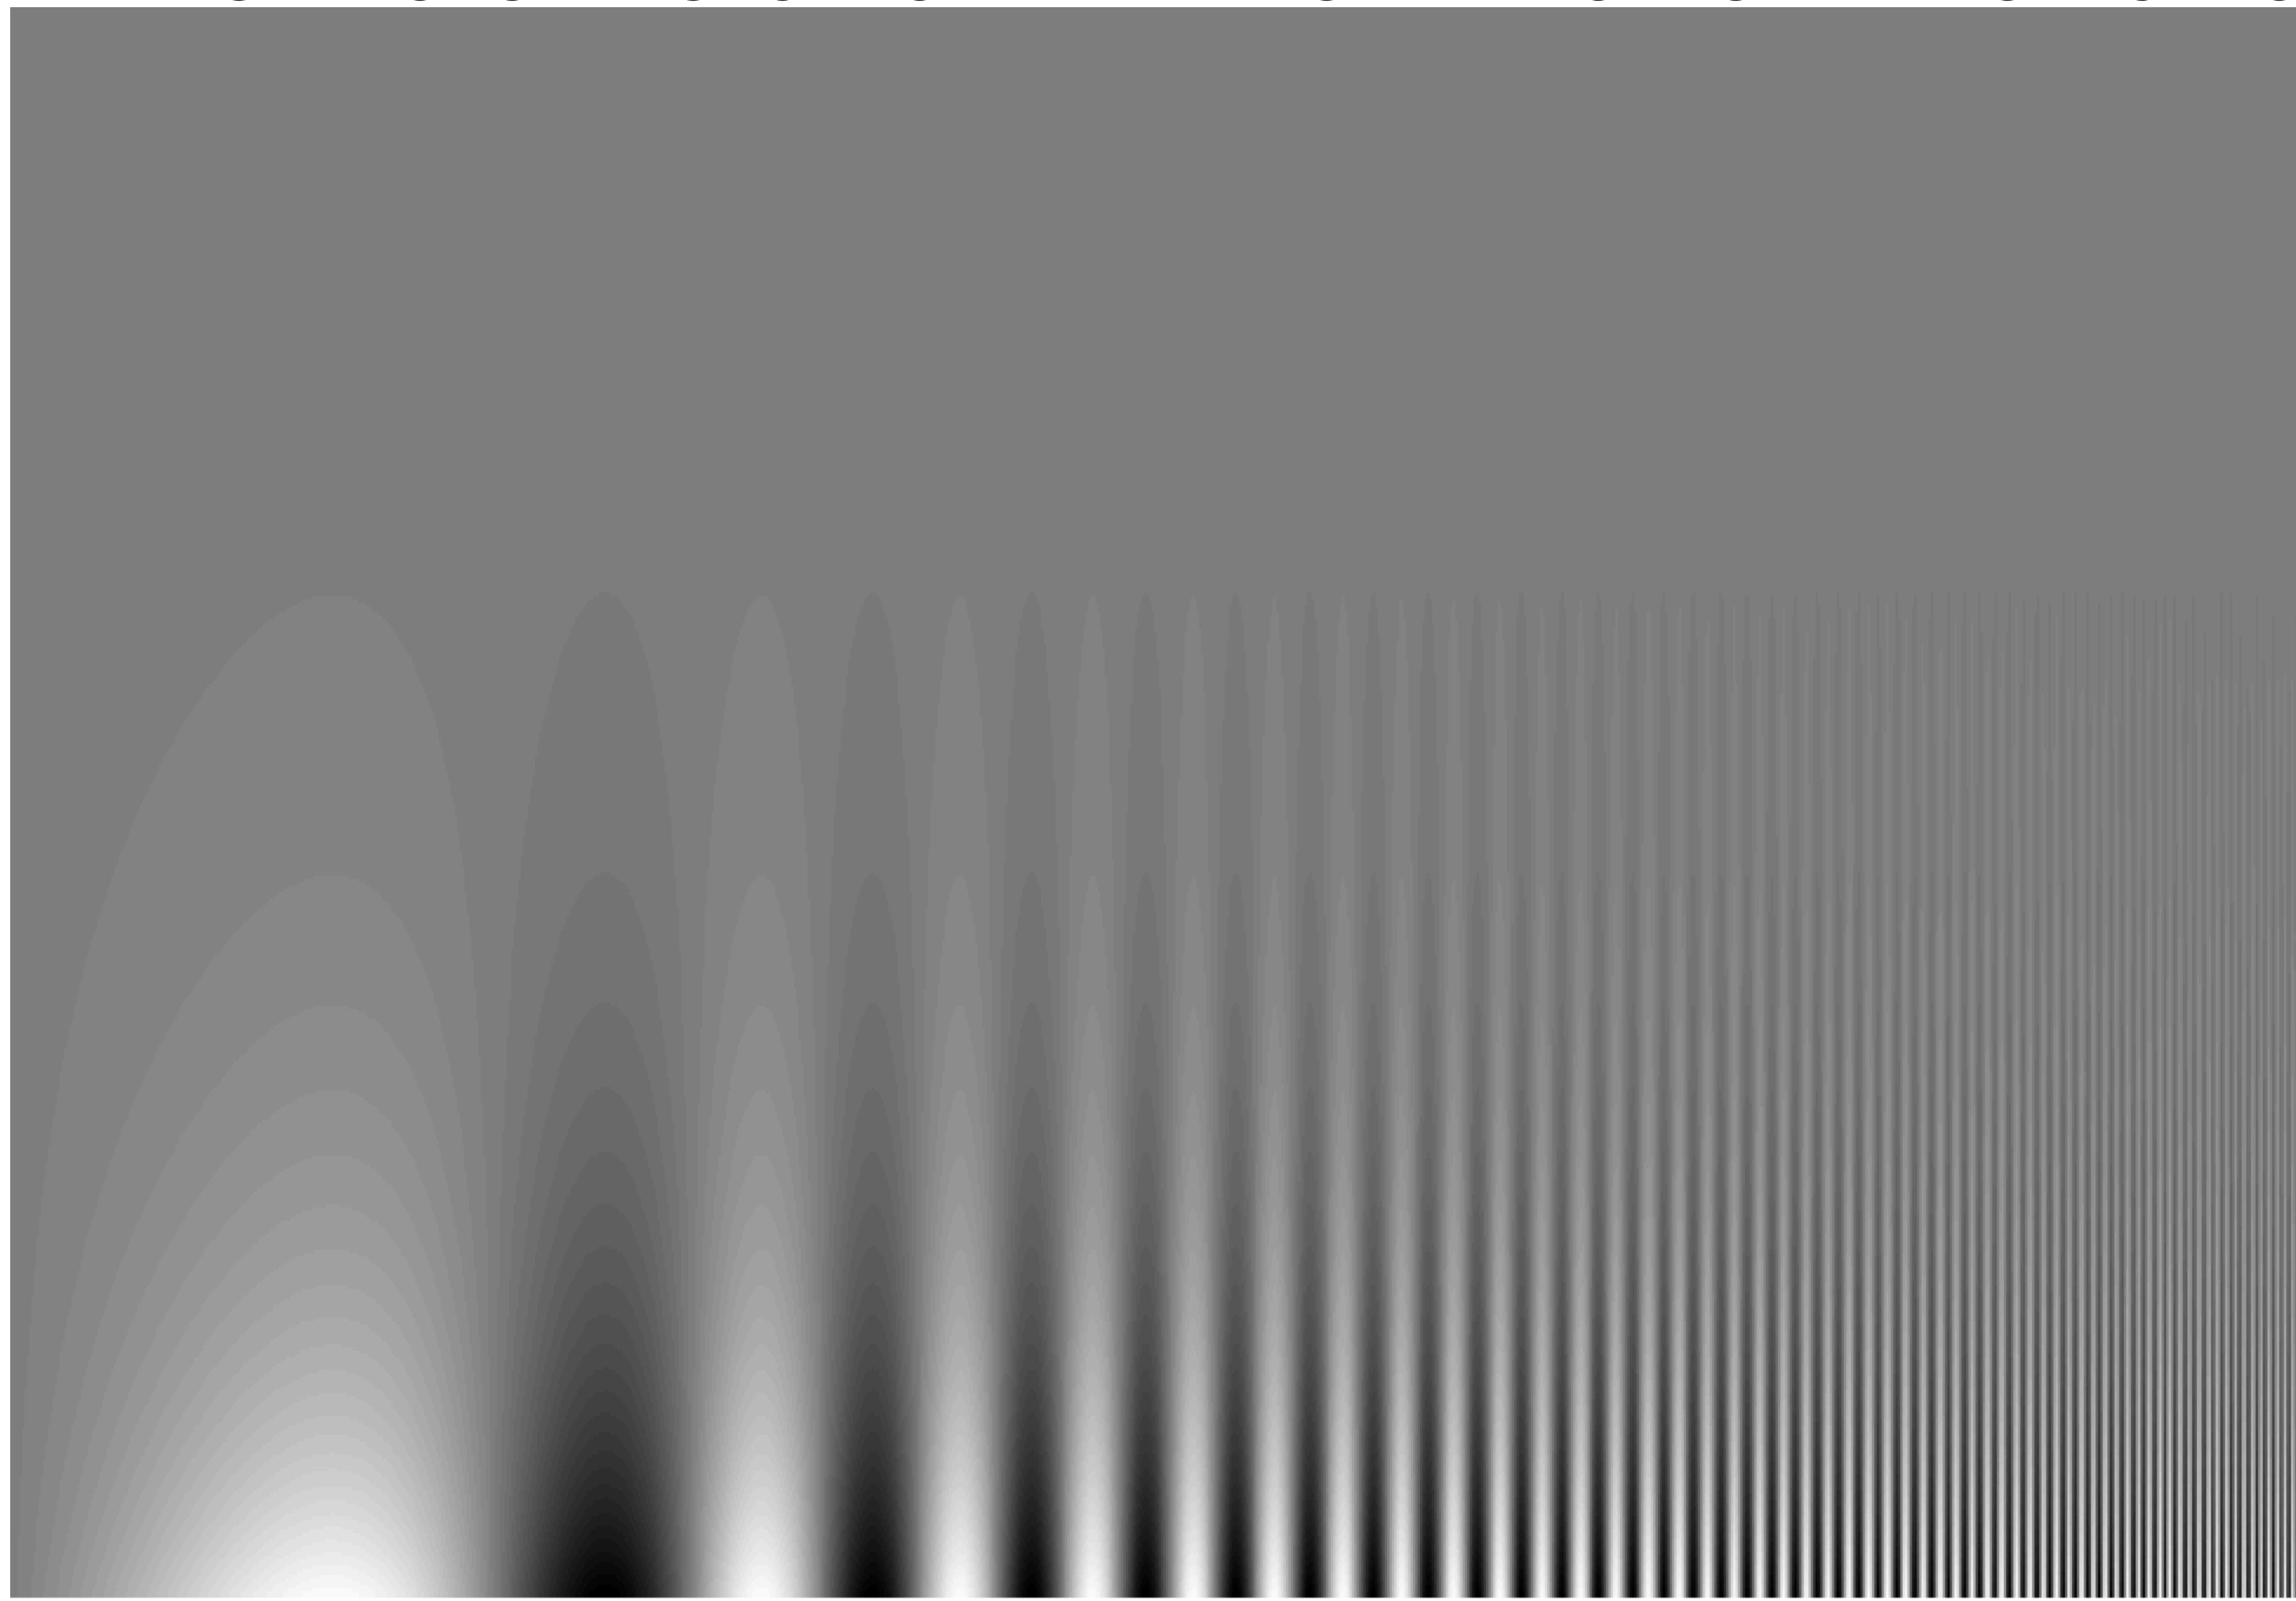
~ **6** cycles / degree of visual angle

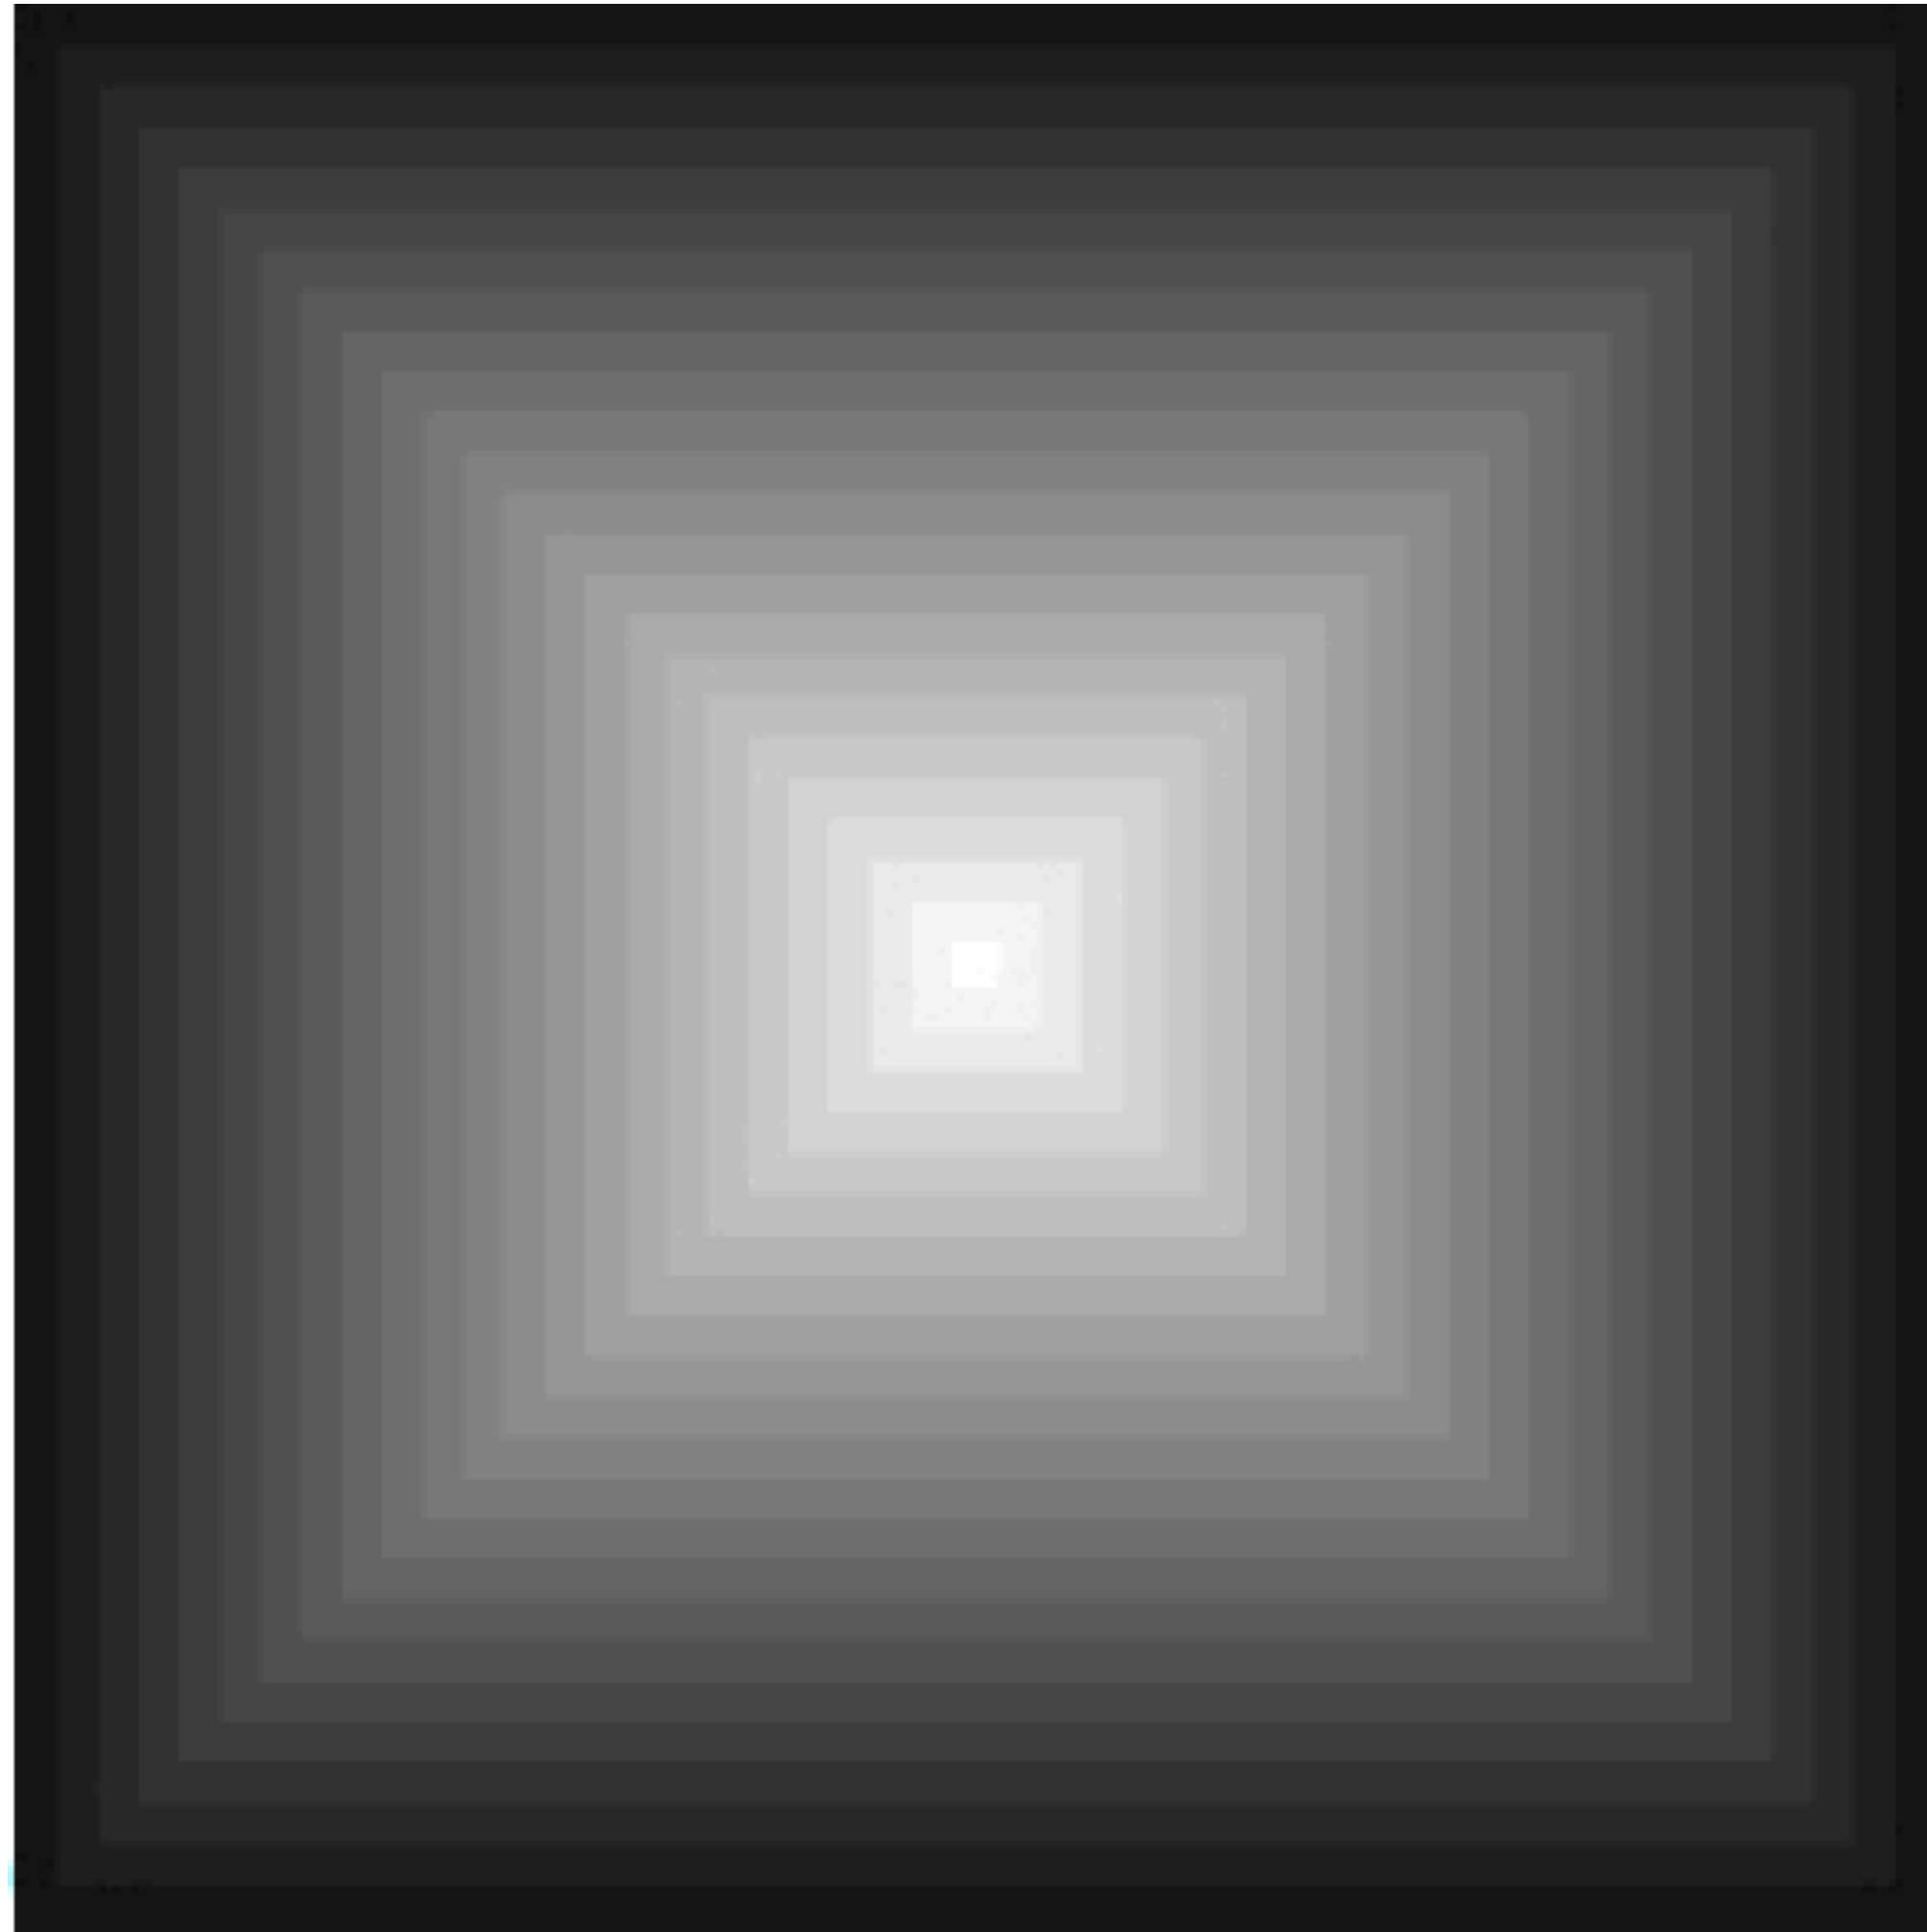


Things that are very close
and/or large are hard to see

Things far away
are hard to see

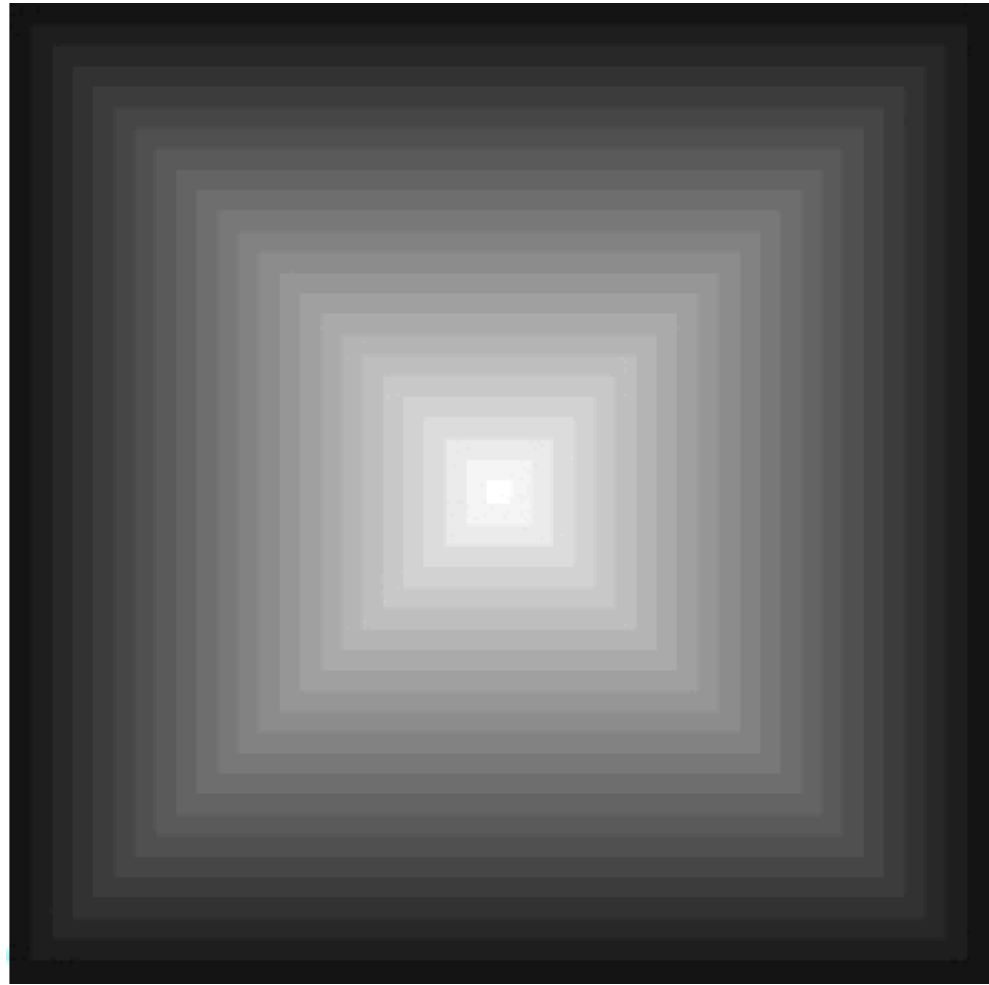
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20



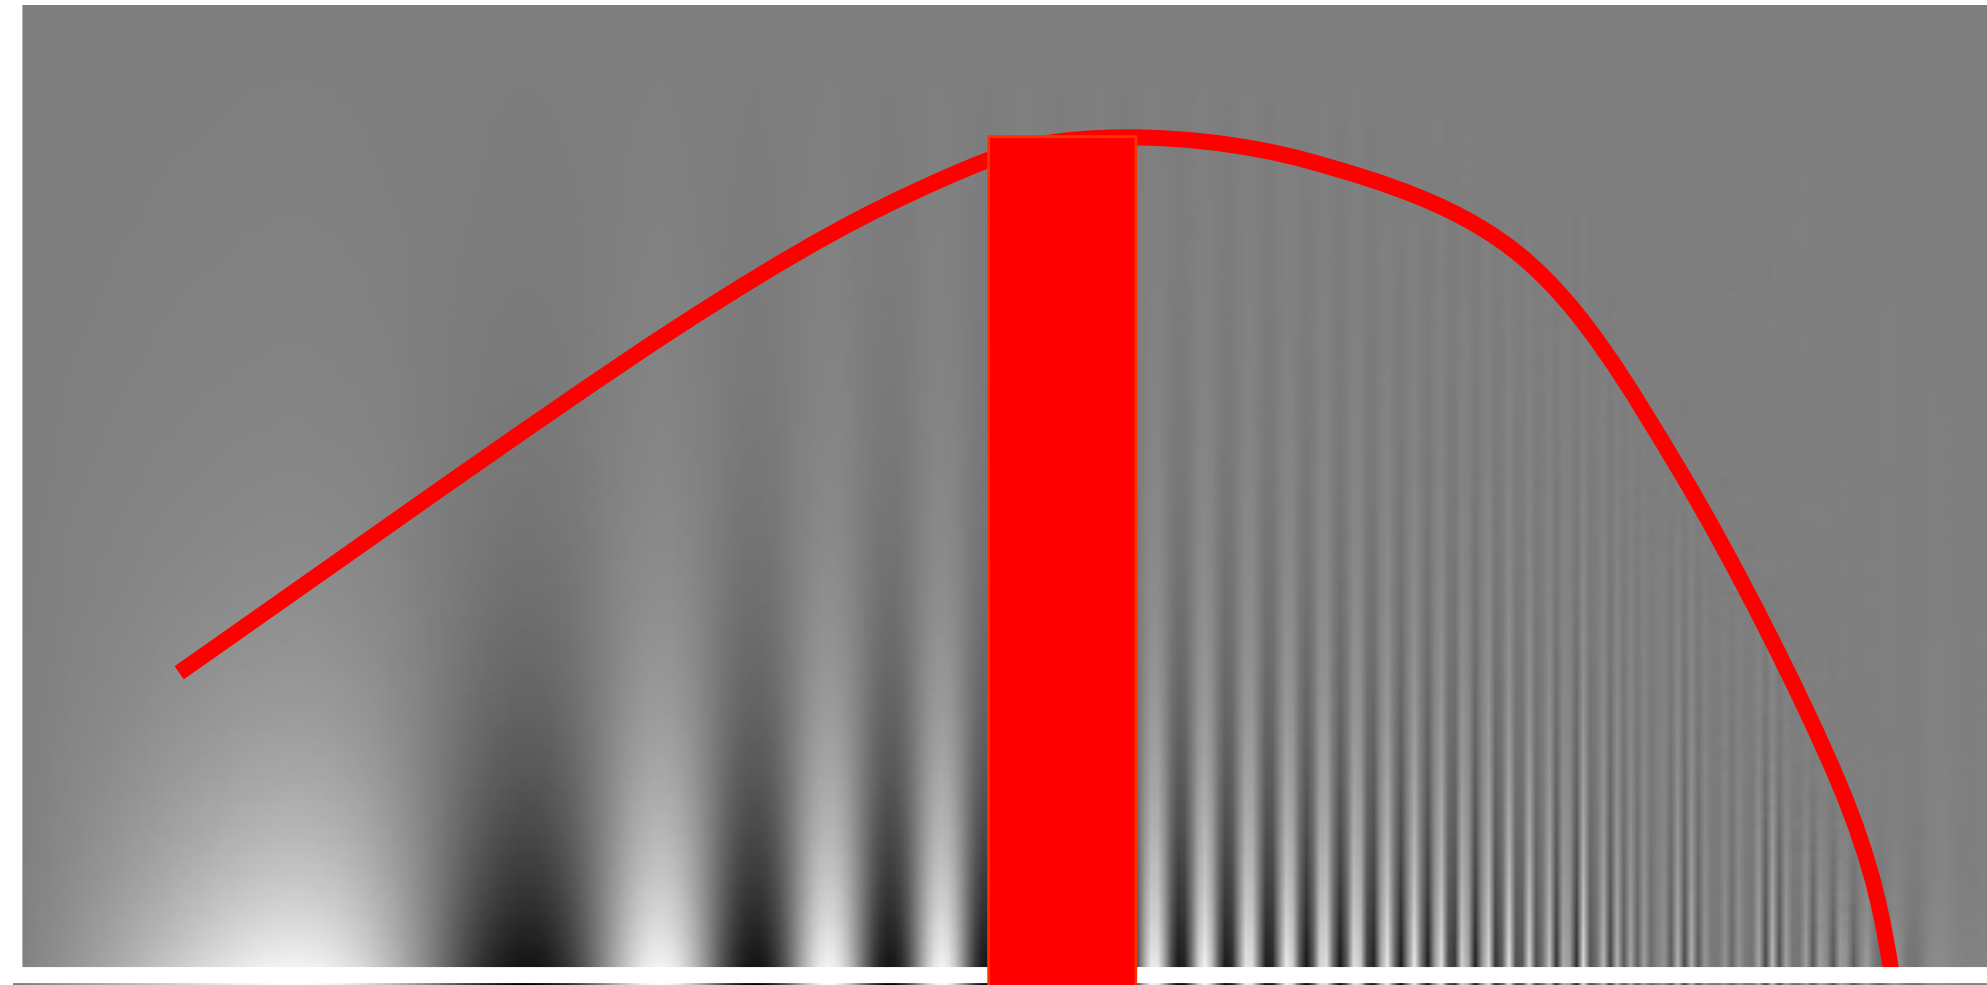


Vasarely visual illusion

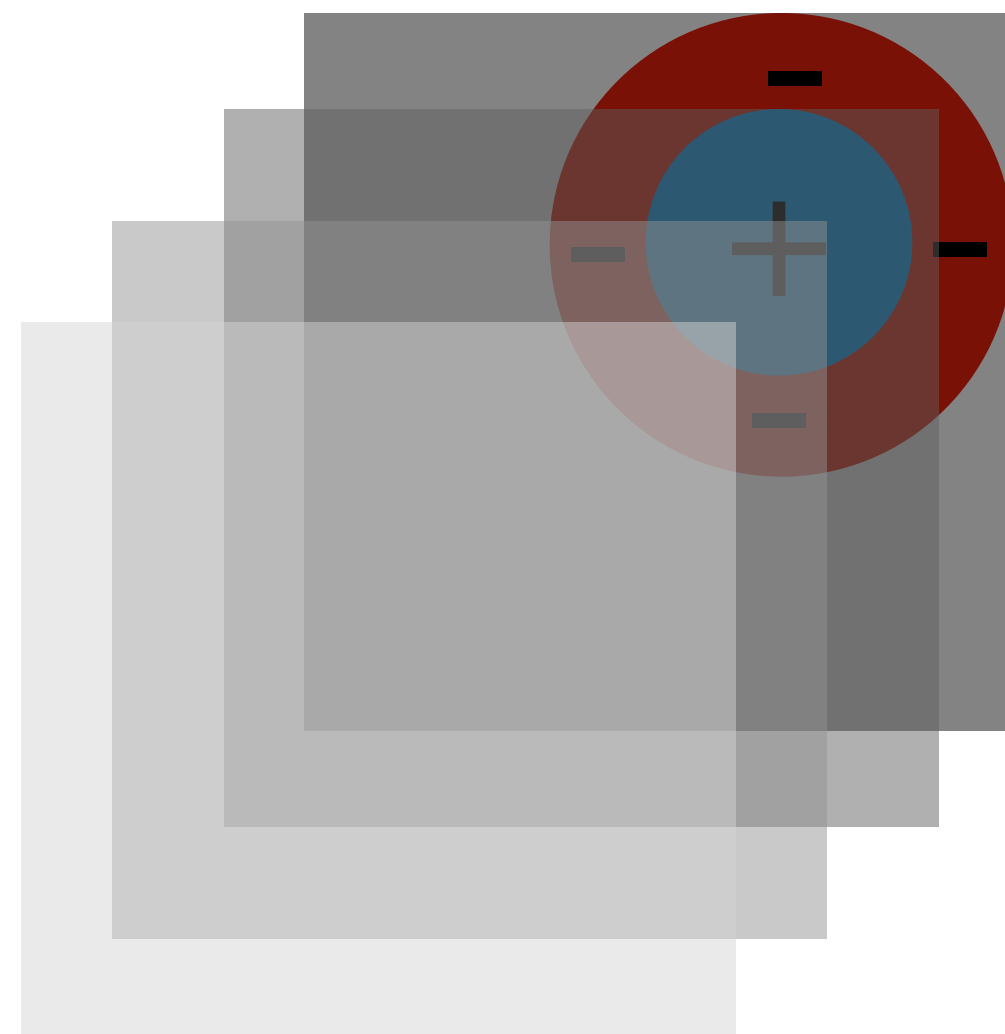
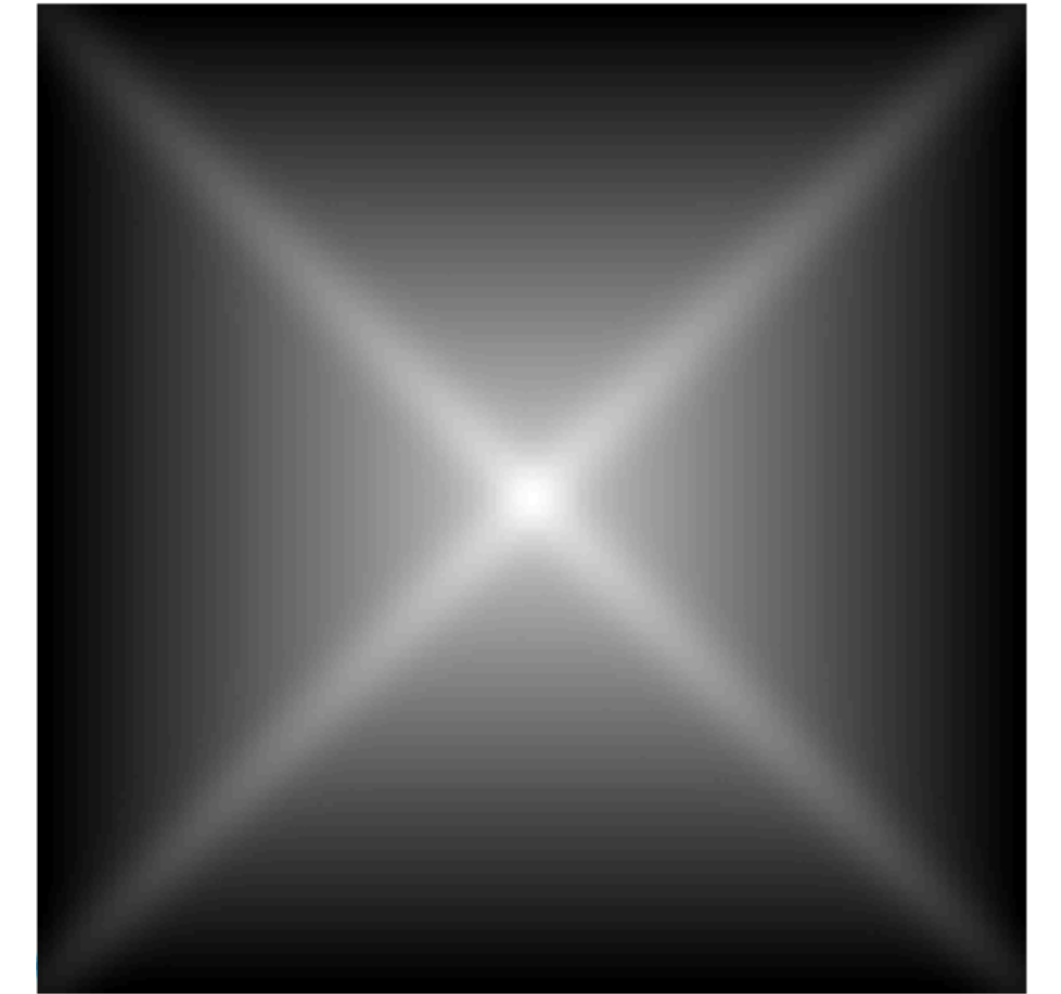
Input visual stimulus



Frequency filtering of human visual system

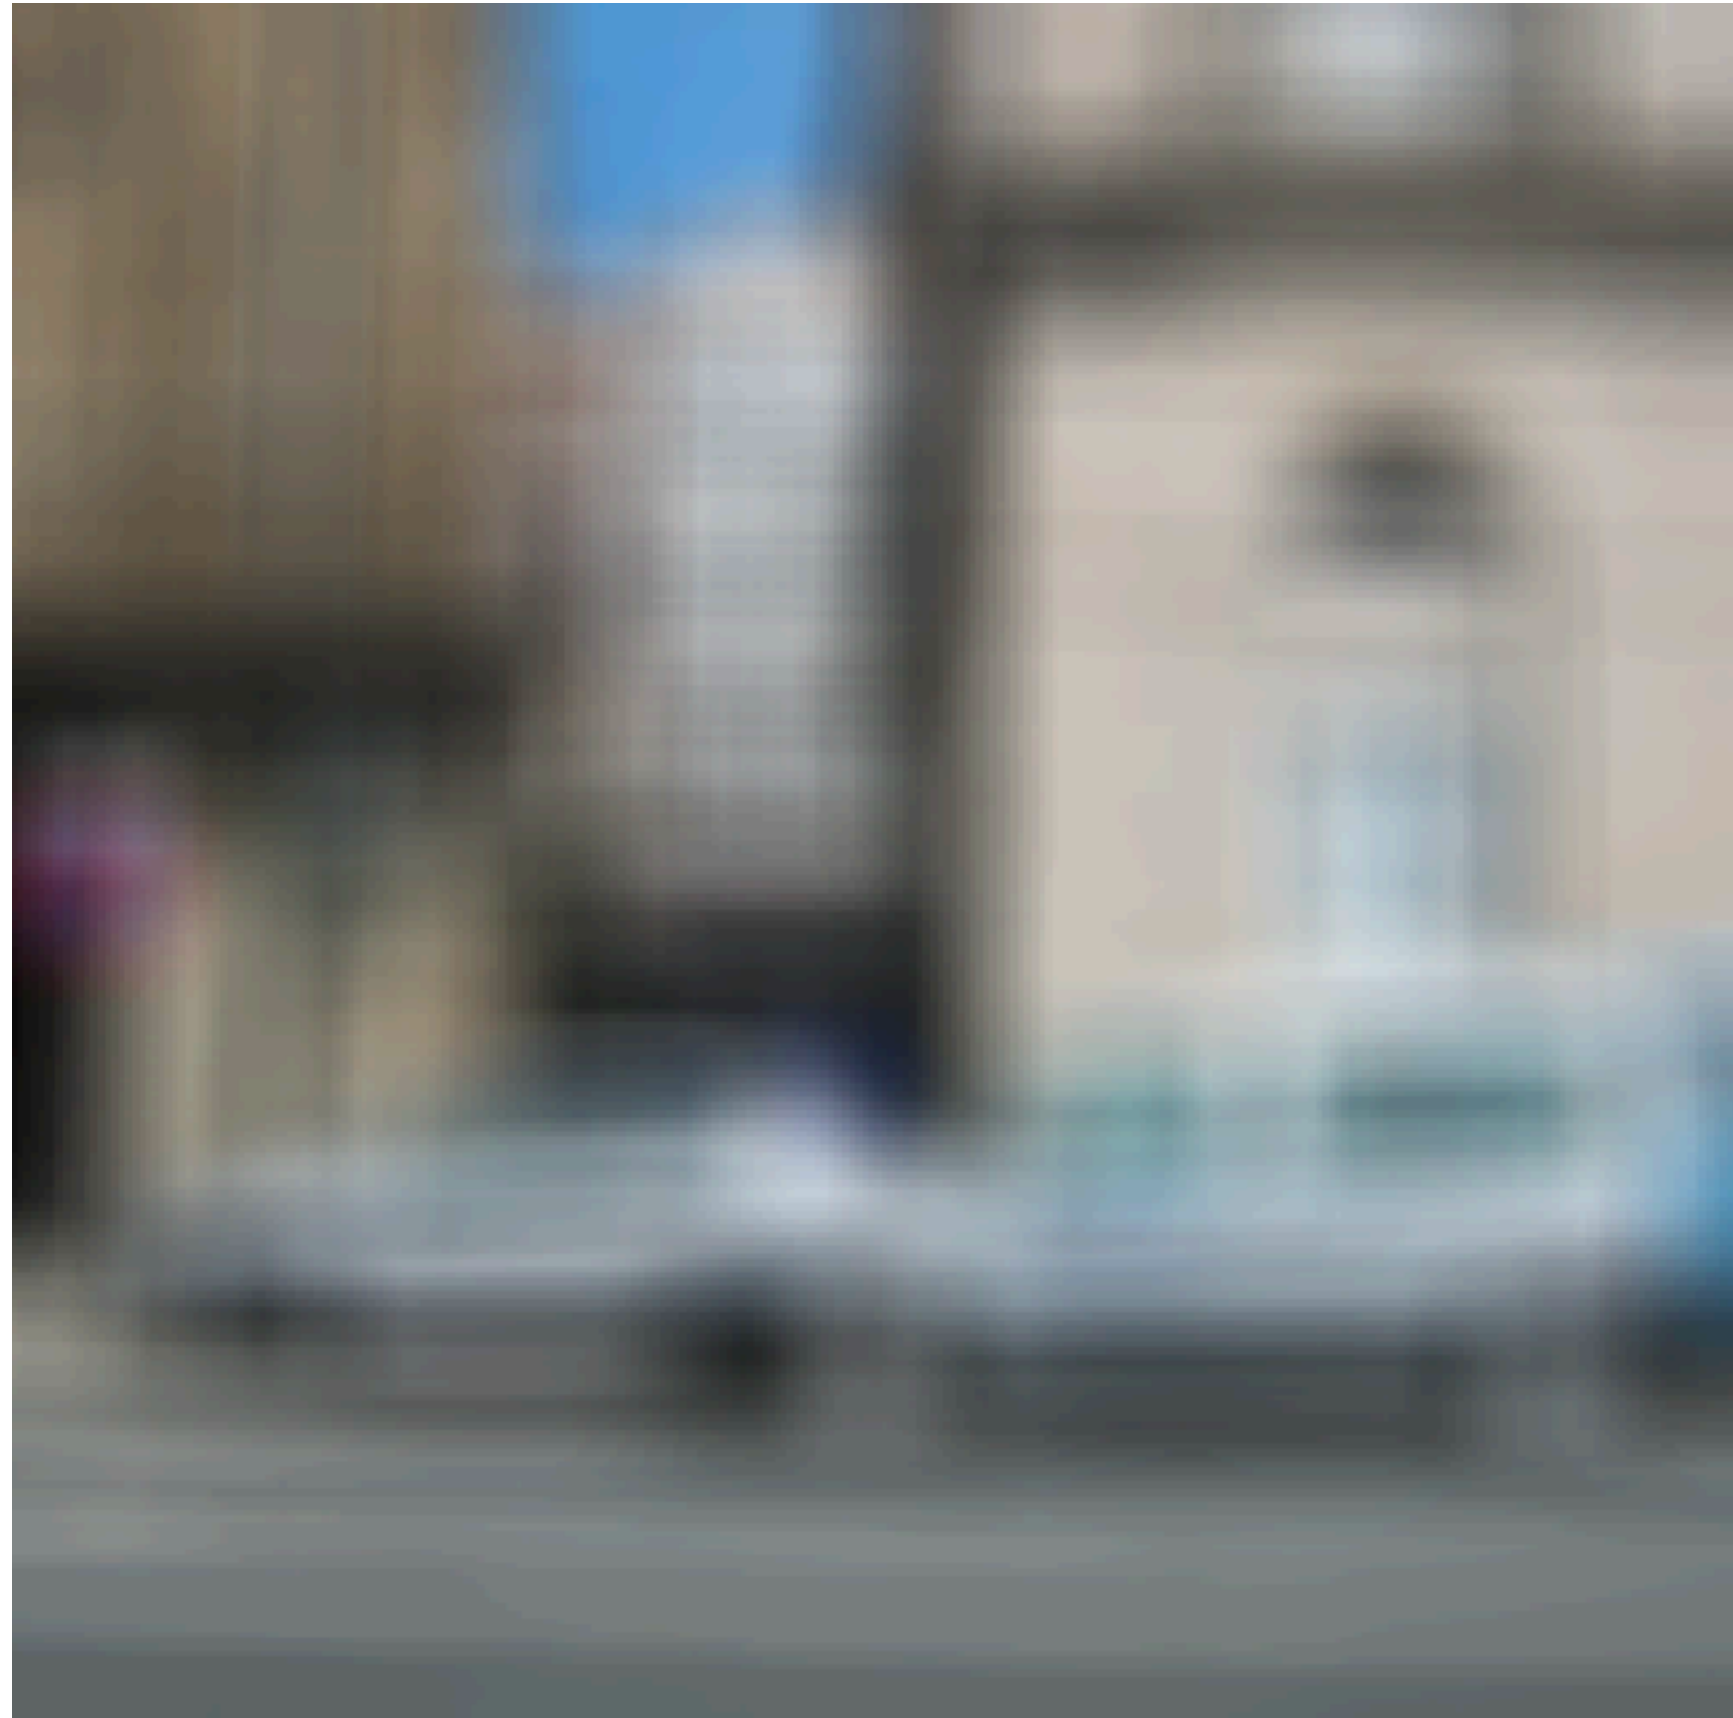


bandpass filtered output



Center-surround spatial filtering of human visual system is subtracting less positive intensity at the corners, giving a bright line there

Today: A collection of useful filters



Low-pass filters



High-pass filters

BLUR

Low pass-filters

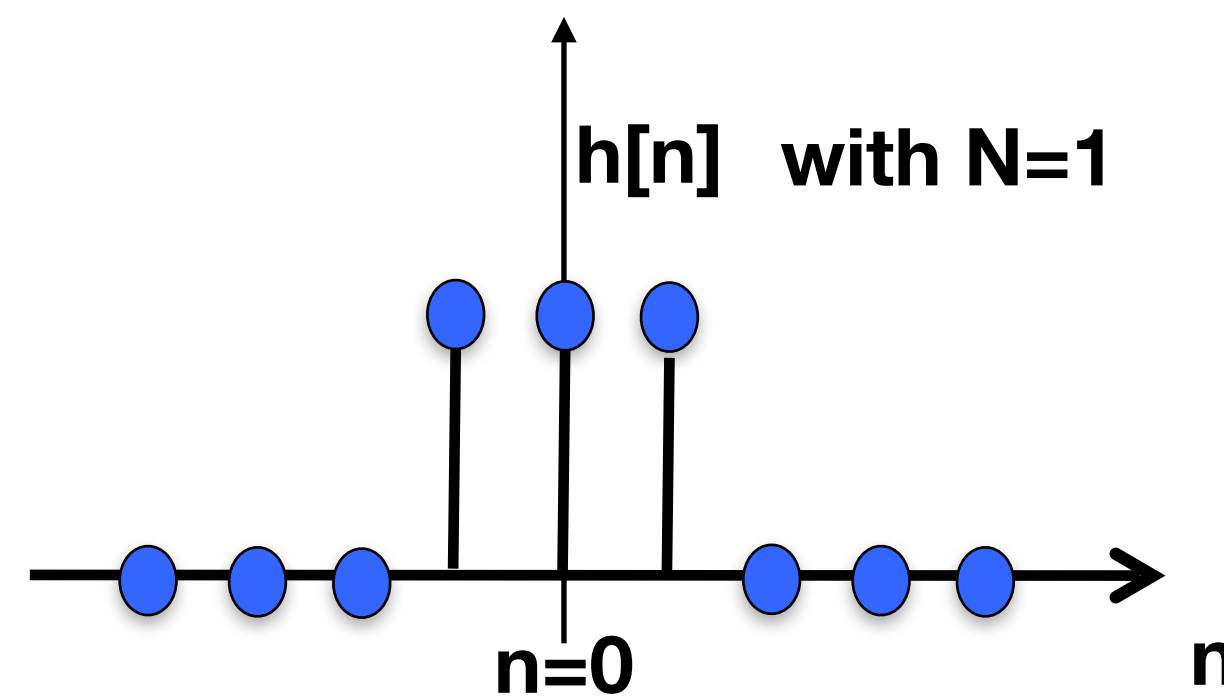
Box filter

$2N+1$

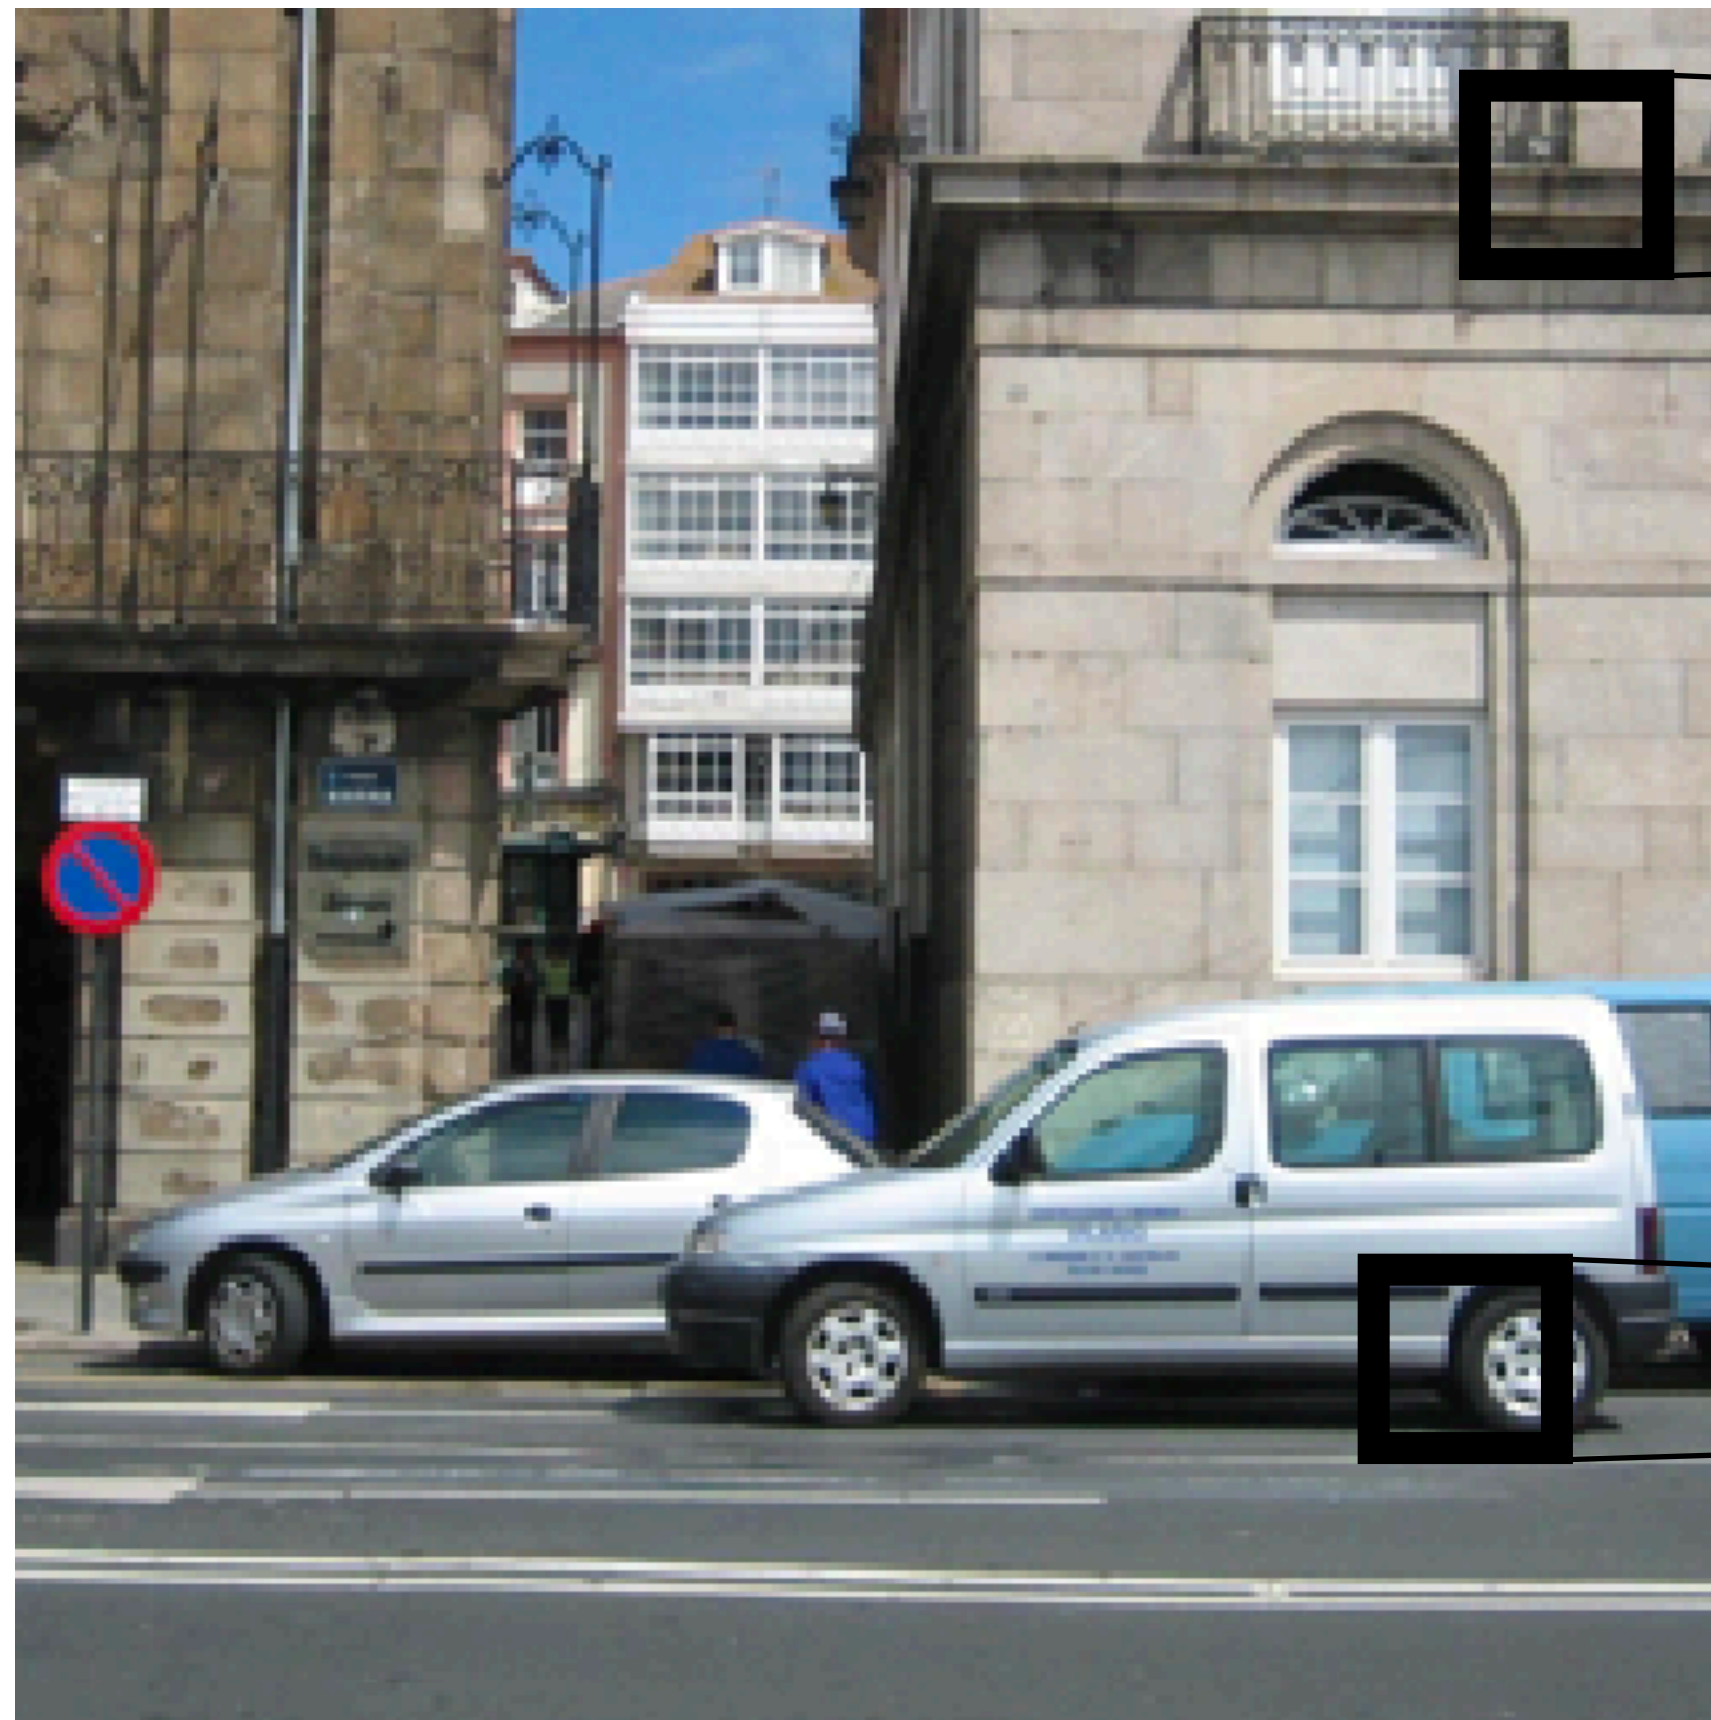
1	1	...	1
1	1		1
1	1		1
...			
1	1	1	1

$2M+1$

$$h_{N,M}[n,m] = \begin{cases} 1 & \text{if } -N \leq n \leq N \text{ and } -M \leq m \leq M \\ 0 & \text{otherwise} \end{cases}$$



Box filter

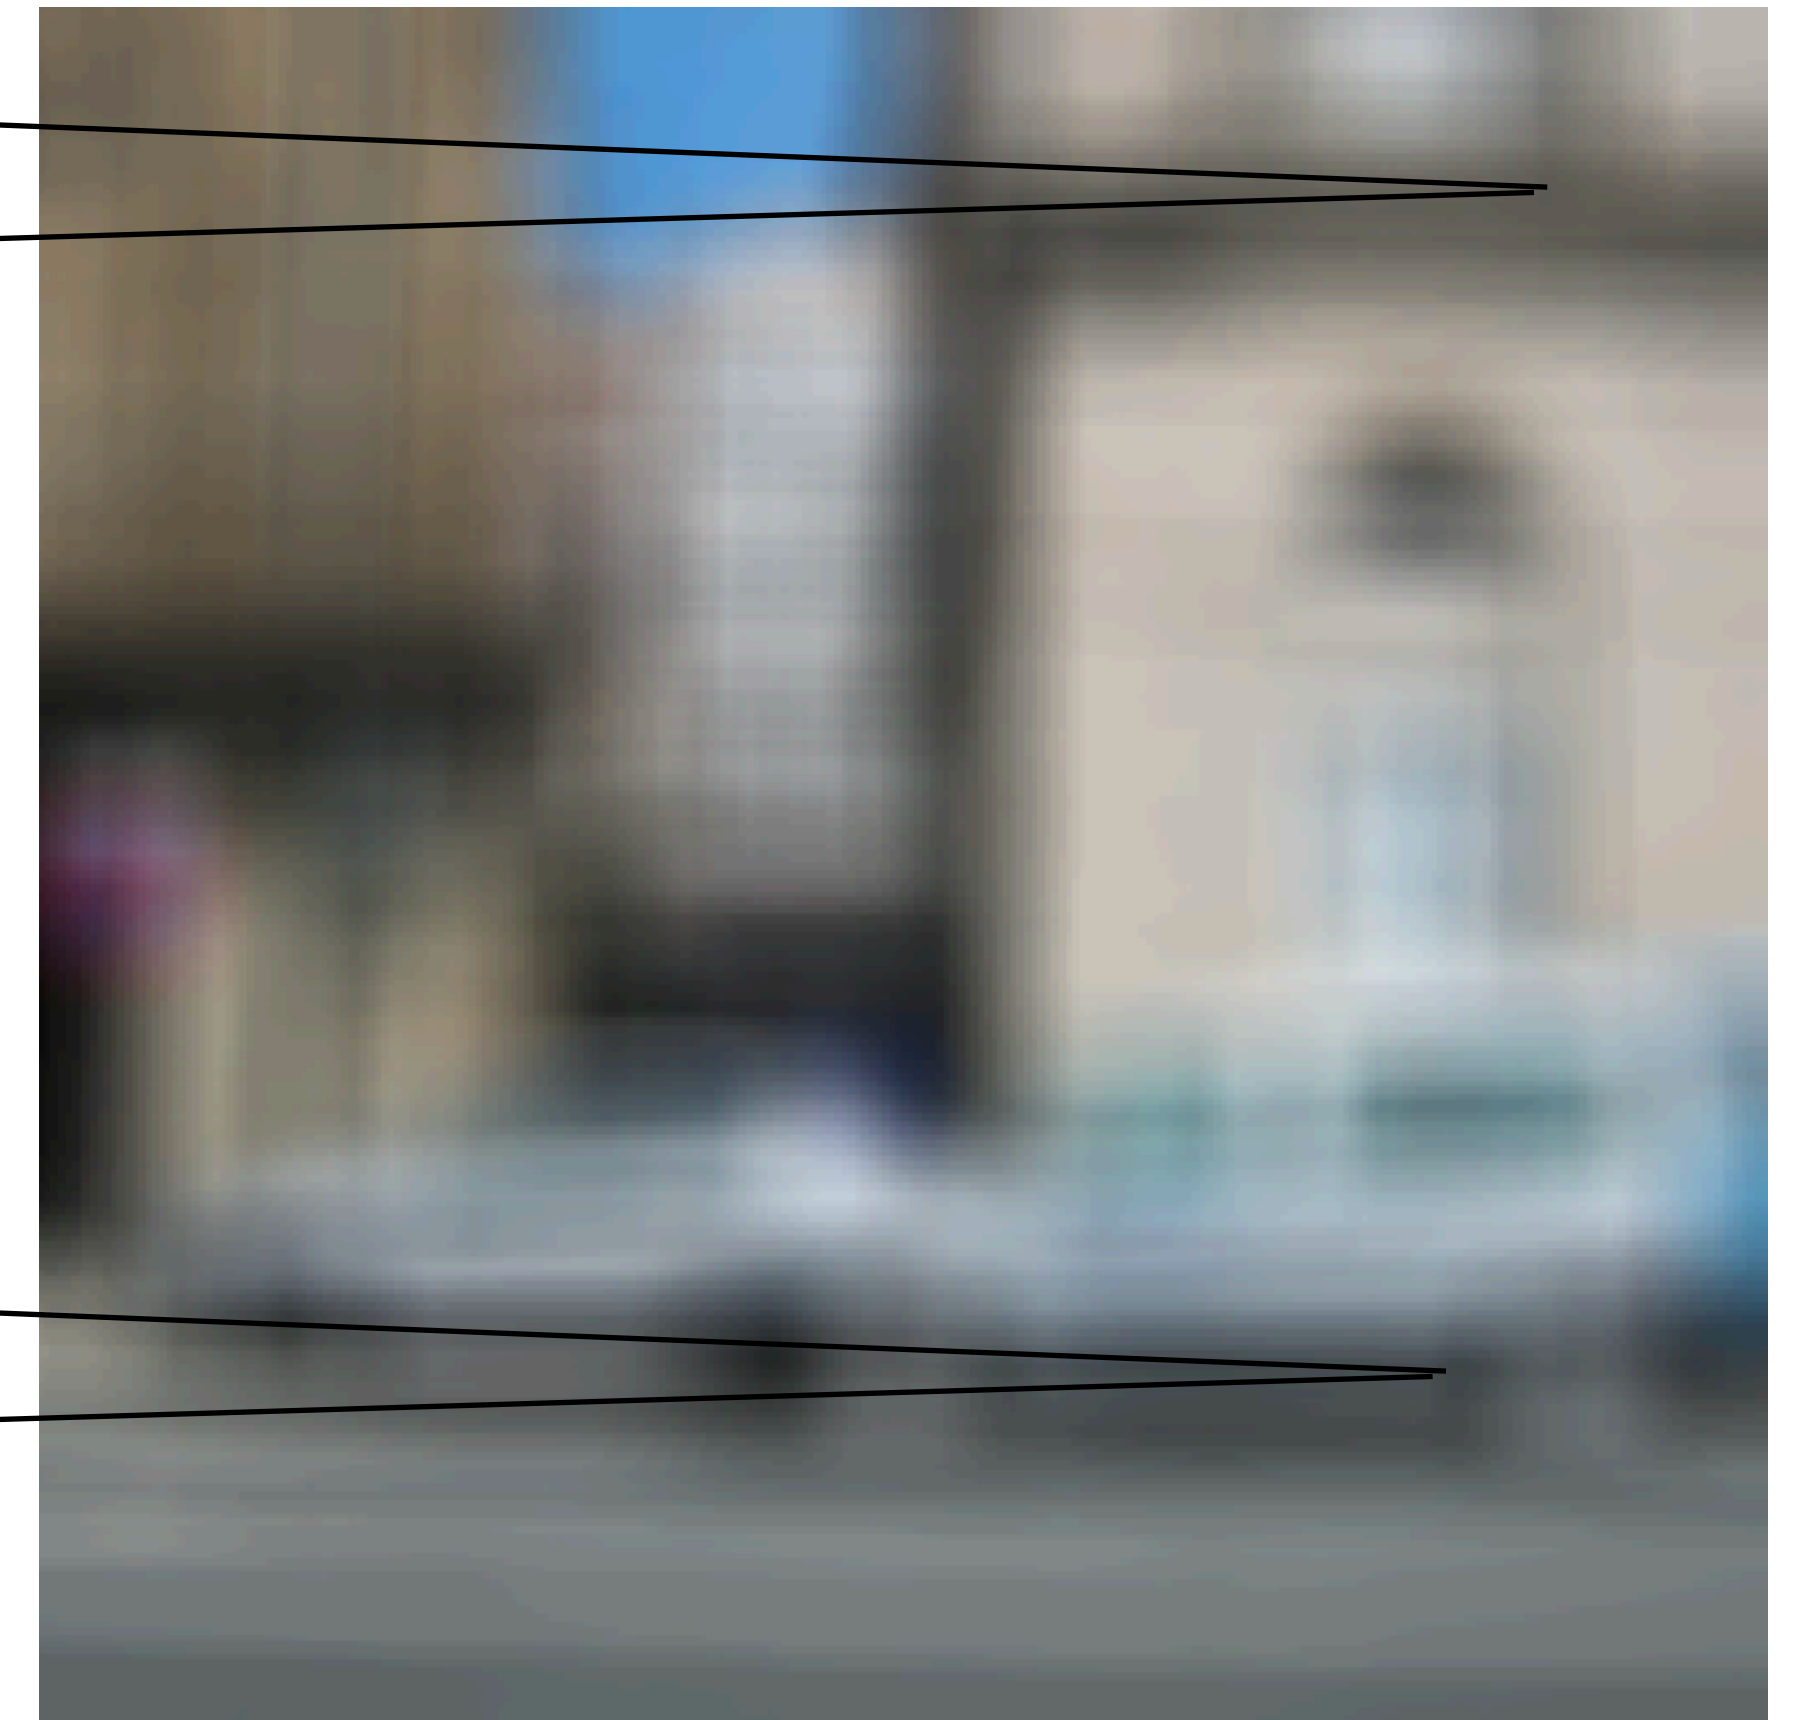


256X256

mean

$$\bigcirc \quad \frac{1}{21 \times 21} \quad \square =$$

mean



256X256

What does it do?

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

Box filter

The box filter is separable as it can be written as the convolution of two 1D kernels

$$h_{N,M}[n,m] = h_{N,0} \circ h_{0,M}$$

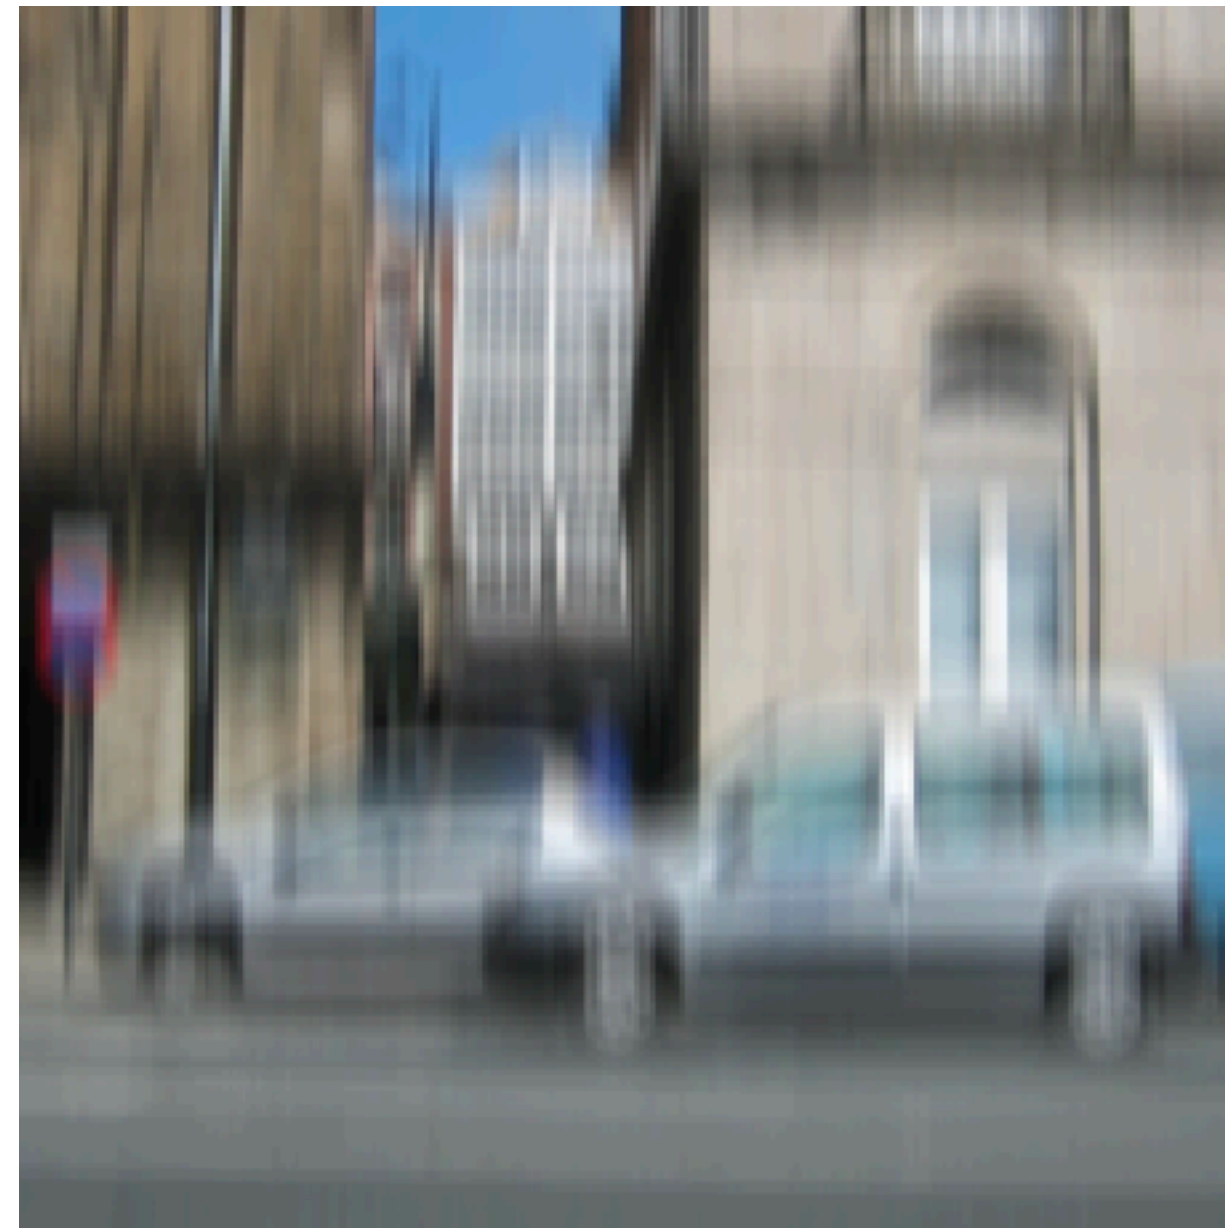
$$\begin{array}{c} 1 \\ 1 \\ 1 \end{array} \circ \begin{array}{cc} 1 & 1 \end{array} = \begin{array}{cc} 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{array}$$

Box filter

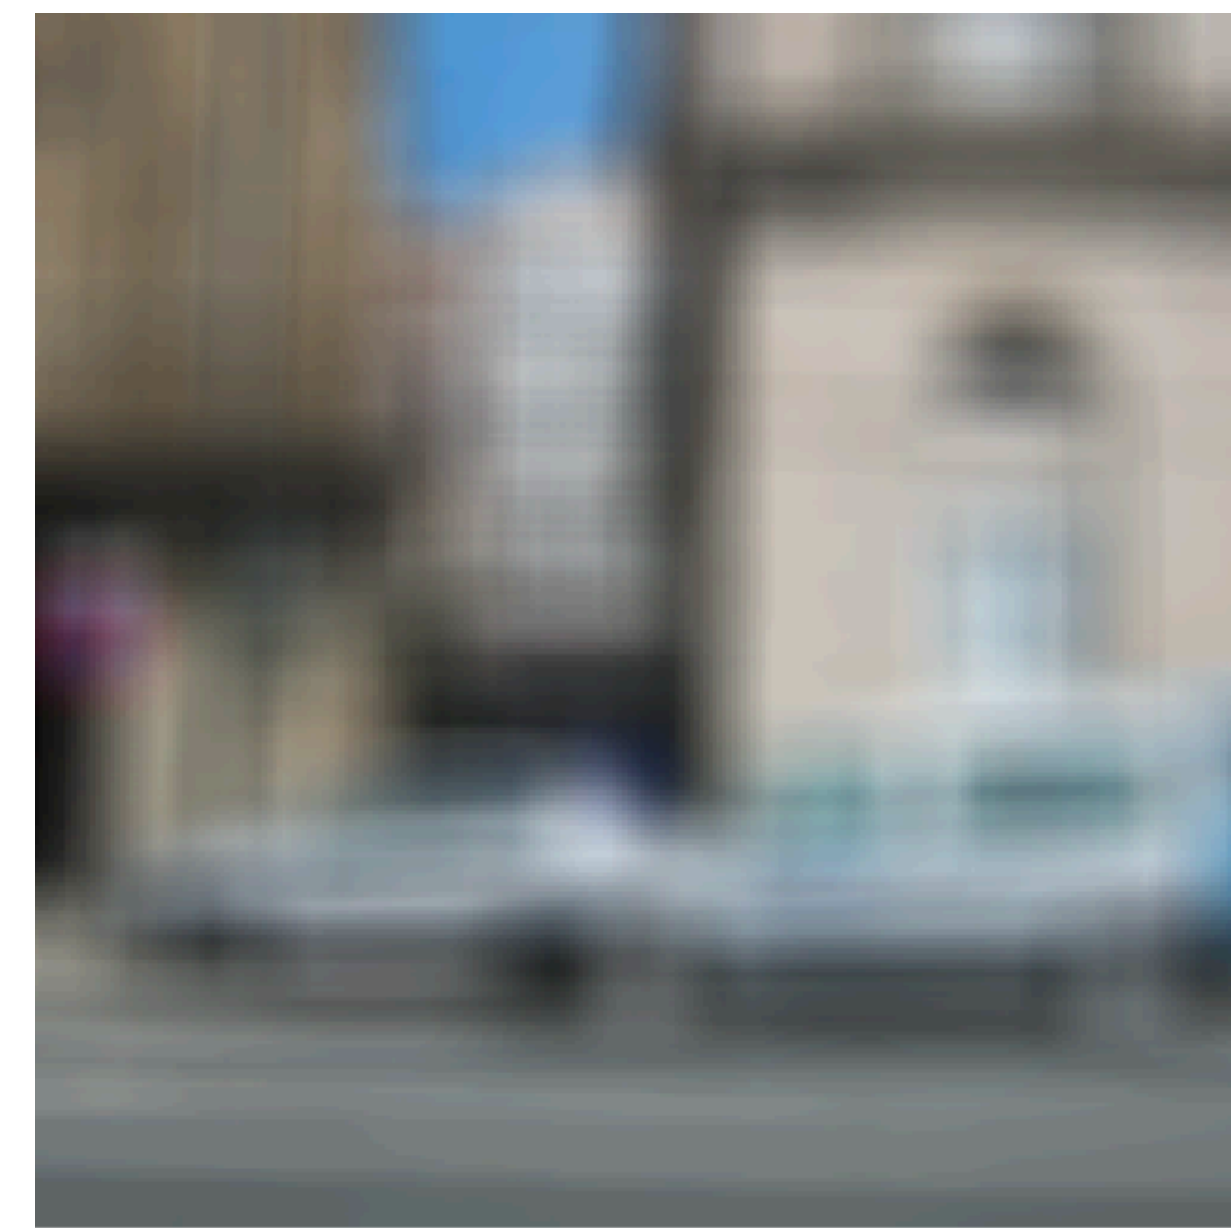


256X256

$$\bigcirc \frac{1}{21} \begin{array}{|c|} \hline \\ \hline \end{array} \rightarrow$$



$$\bigcirc \frac{1}{21} \begin{array}{|c|c|} \hline & \\ \hline \end{array} \rightarrow$$



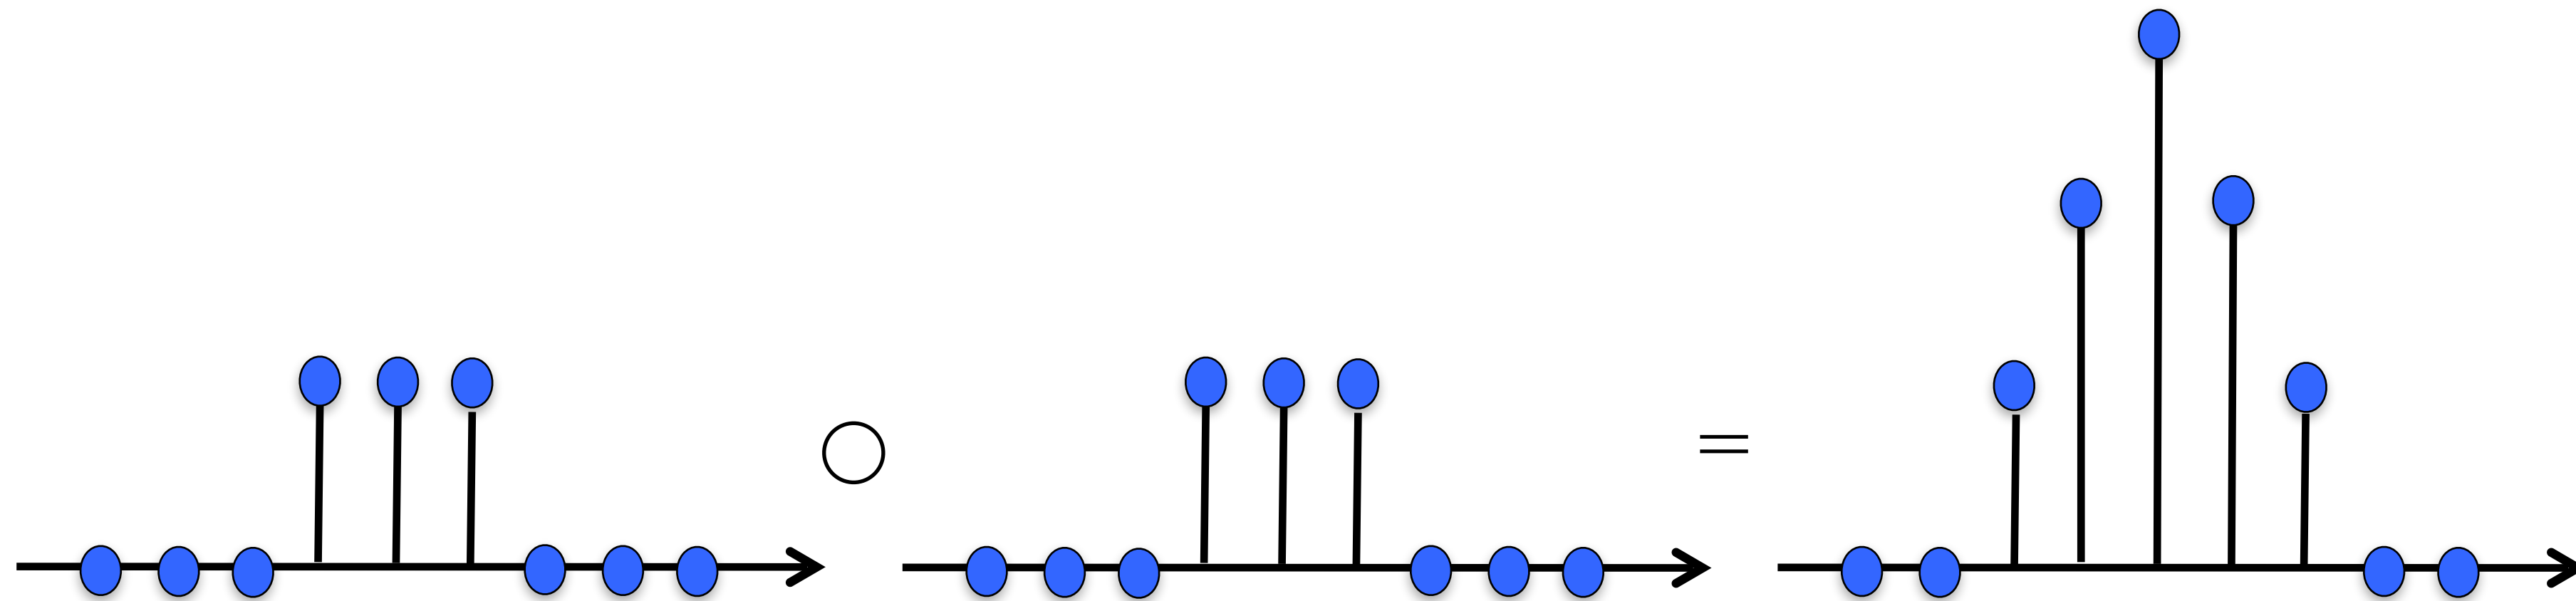
256X256

Requires $N+N$ sums, instead of $N*N$

Box filter

If you convolve two boxes:

$$1 \ 1 \ 1 \ \circlearrowleft \ 1 \ 1 \ 1 = 1 \ 2 \ 3 \ 2 \ 1$$

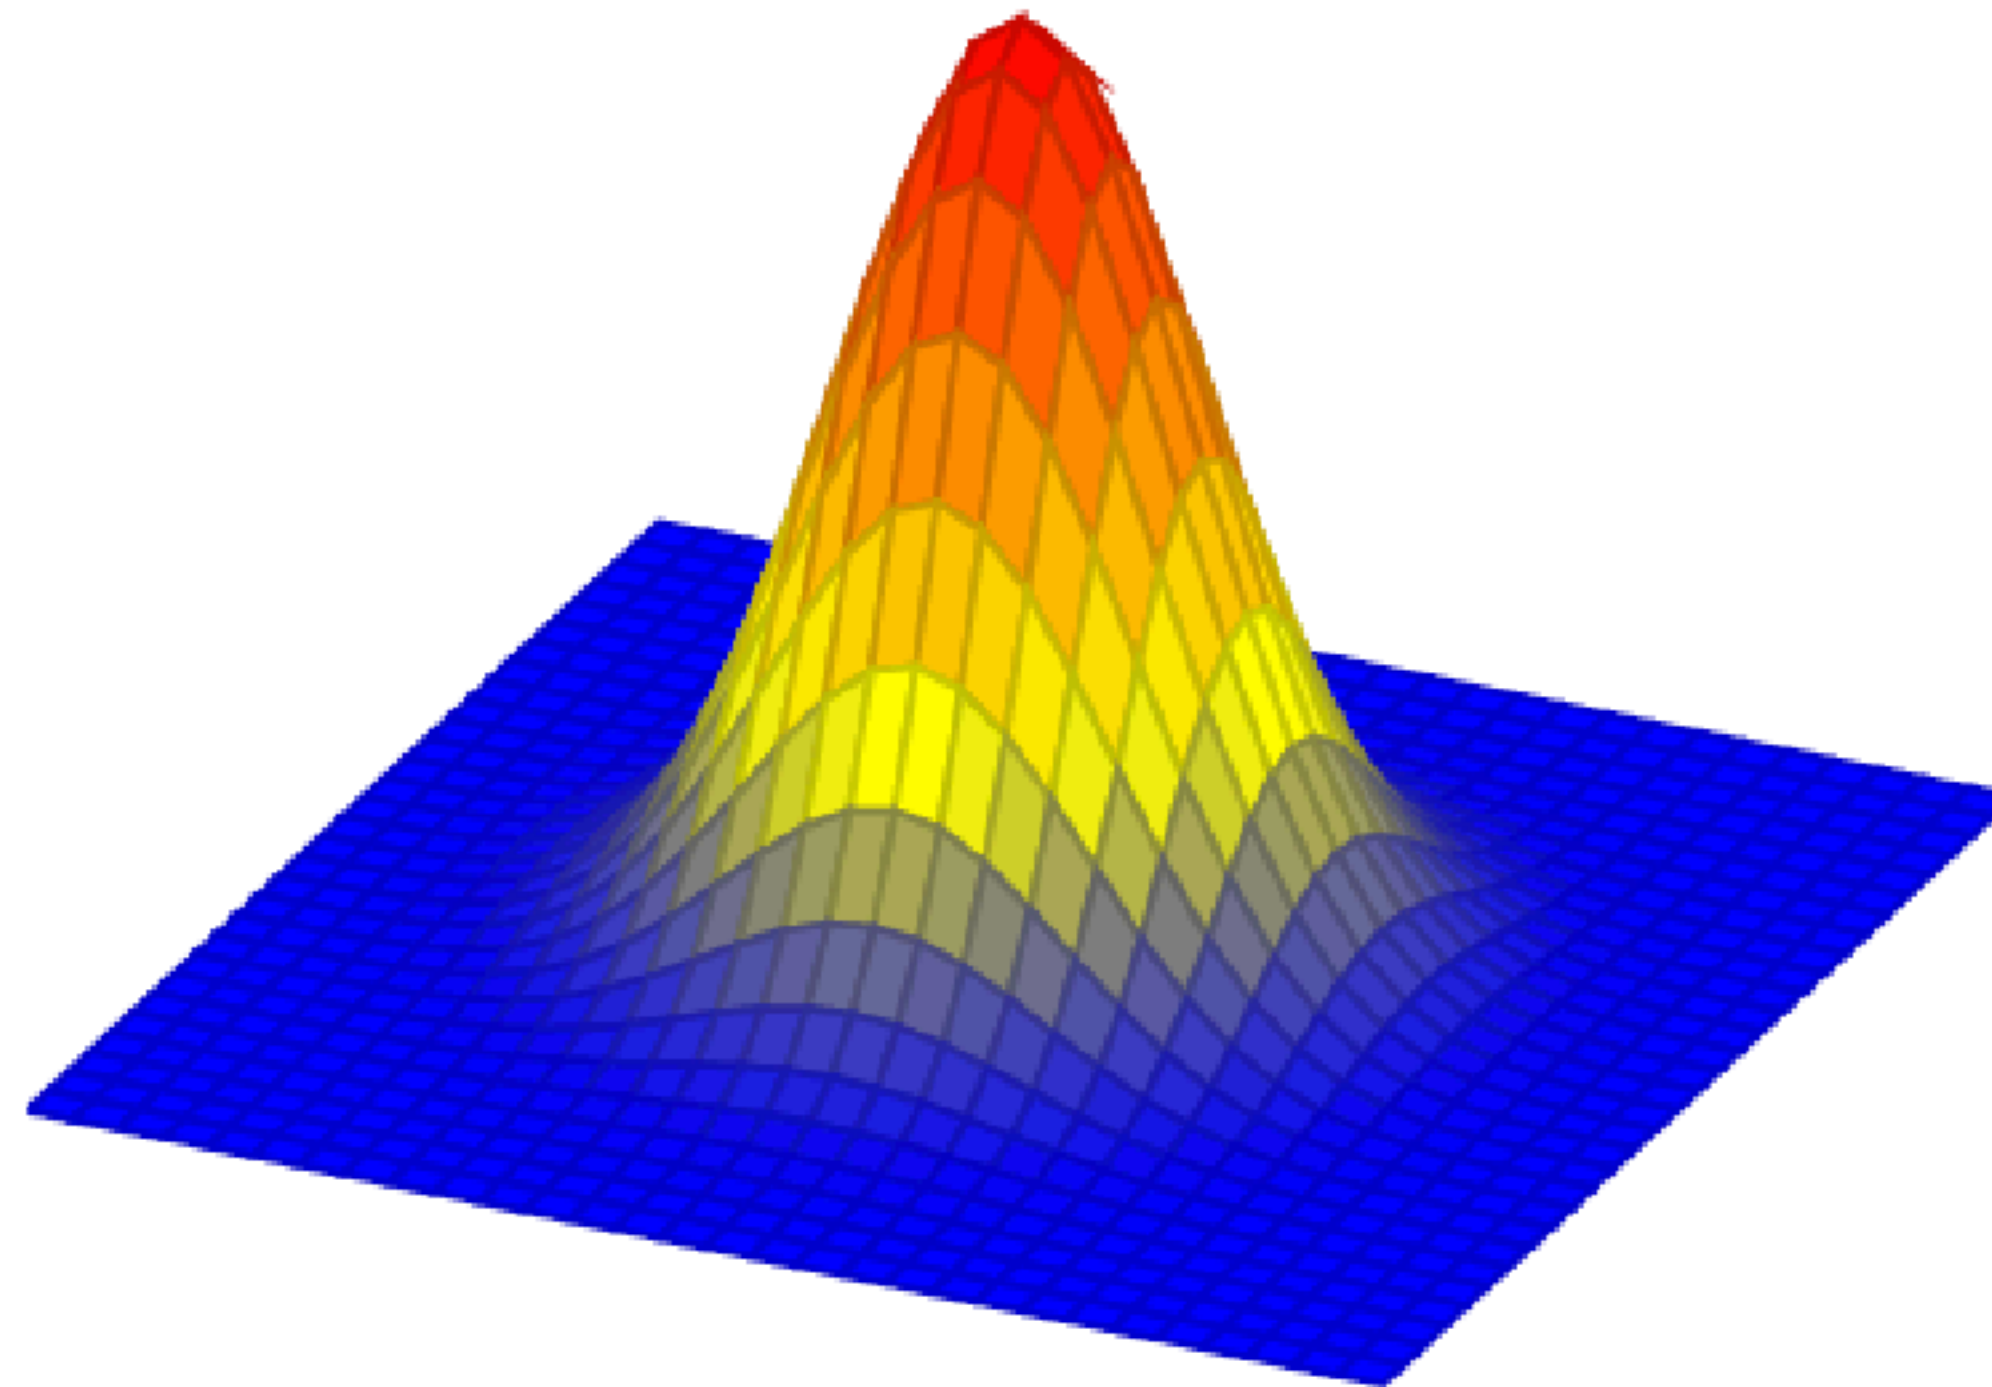


The convolution of two box filters is not another box filter.
It is a triangular filter.

Gaussian filter

In the continuous domain:

$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$



Gaussian filter

$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

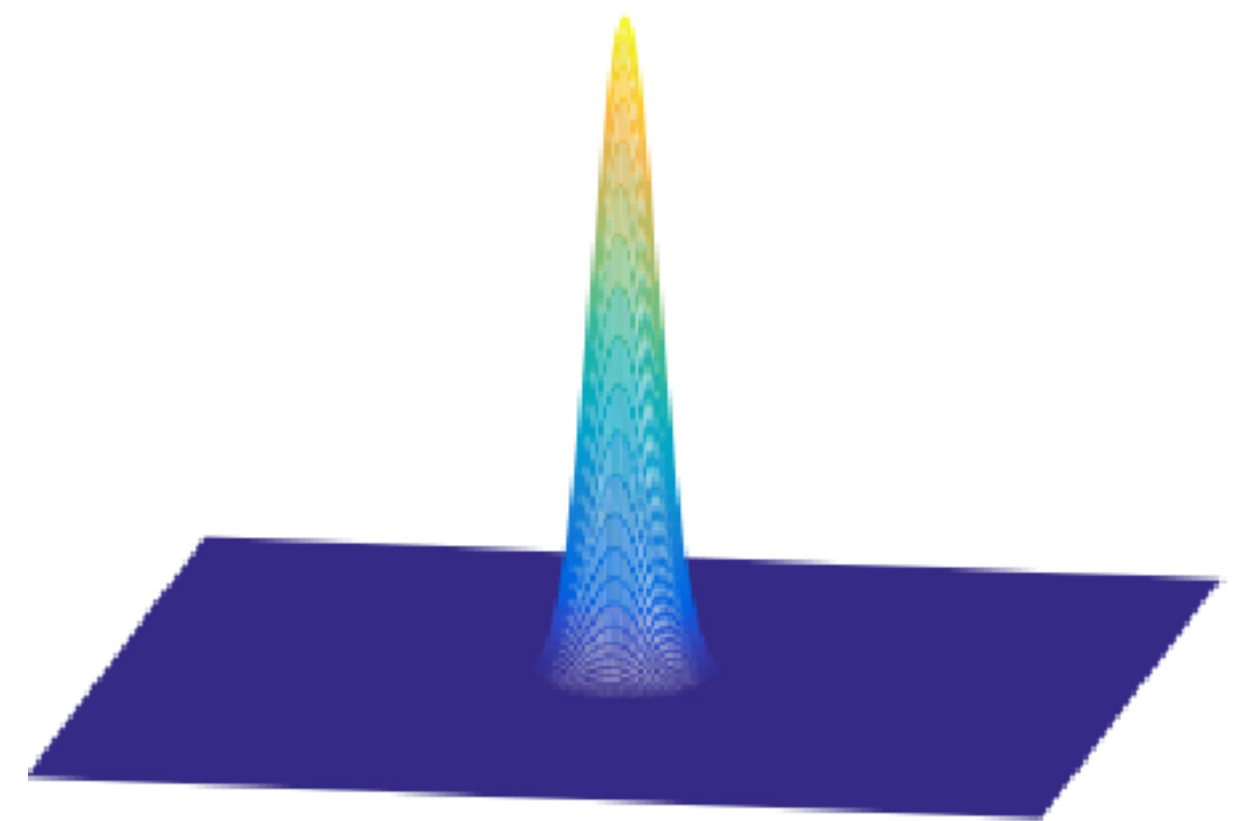
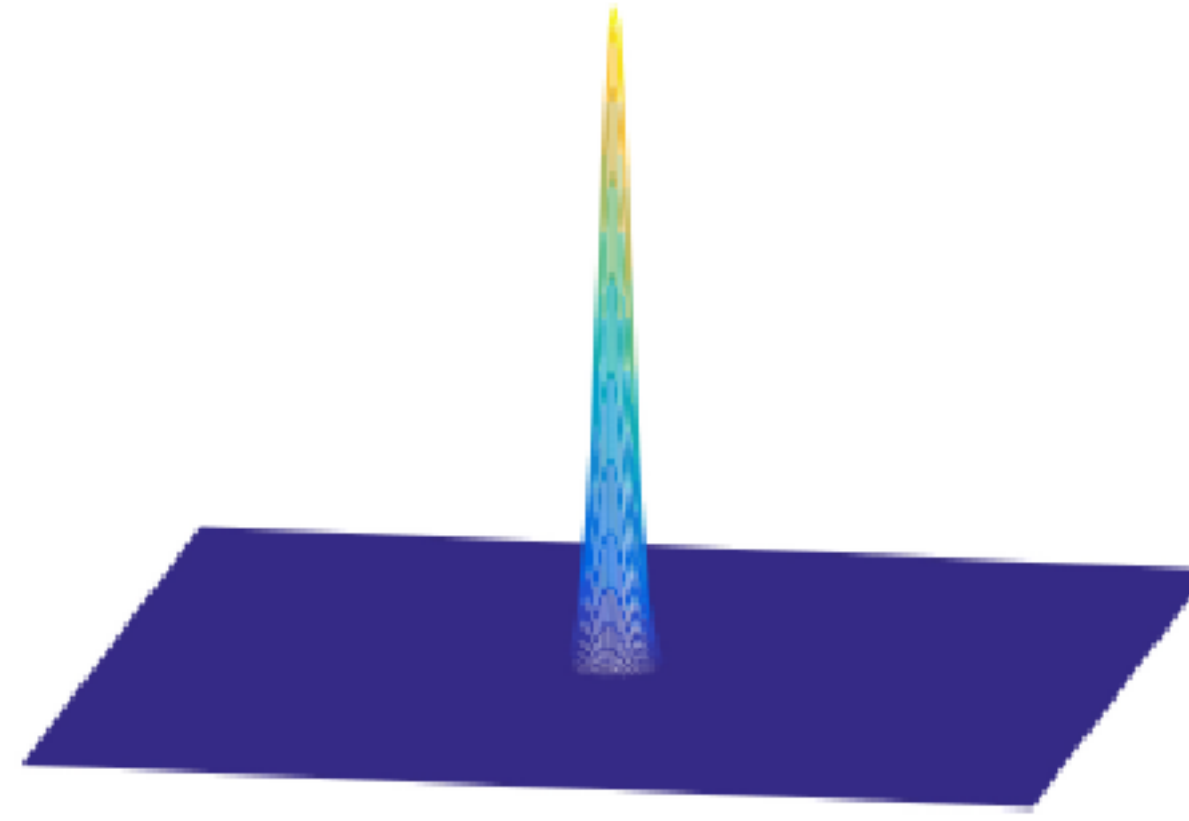
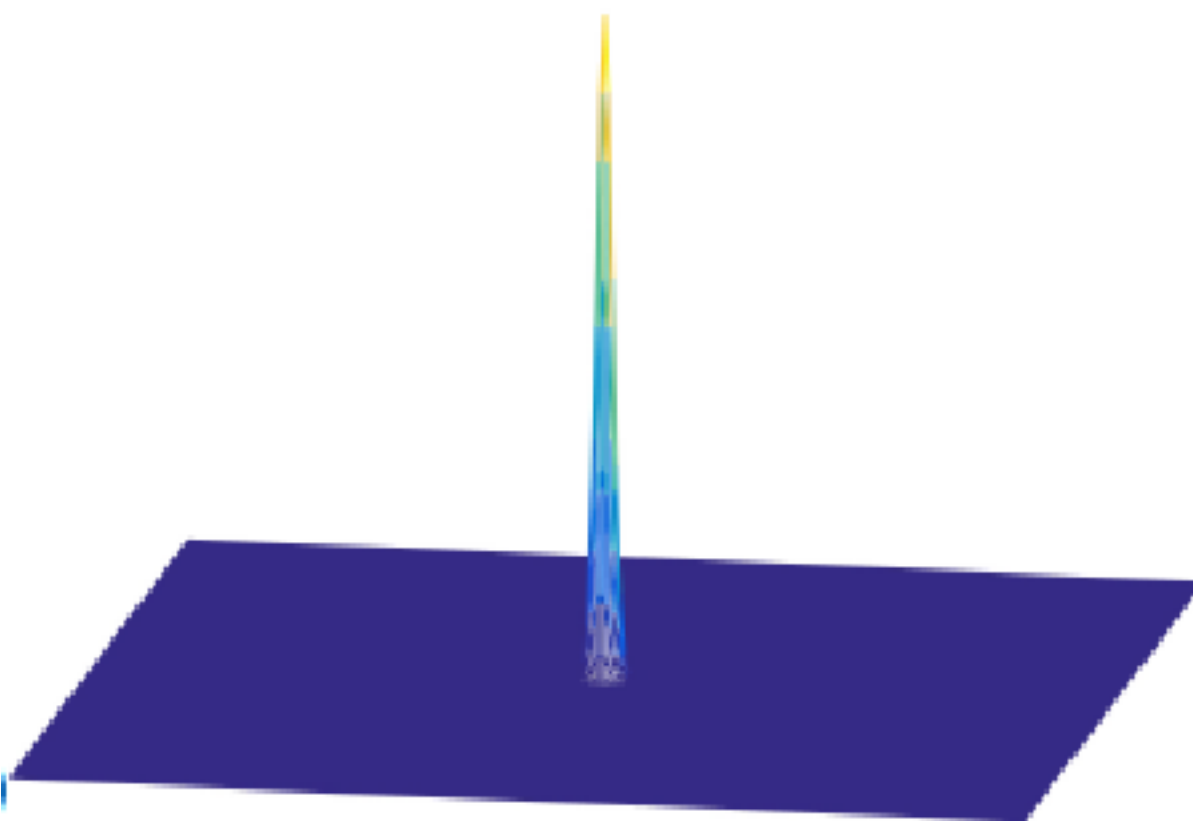
Discretization of the Gaussian:

At 3σ the amplitude of the Gaussian is around 1% of its central value

$$g[m, n; \sigma] = \exp\left(-\frac{m^2 + n^2}{2\sigma^2}\right)$$

Scale

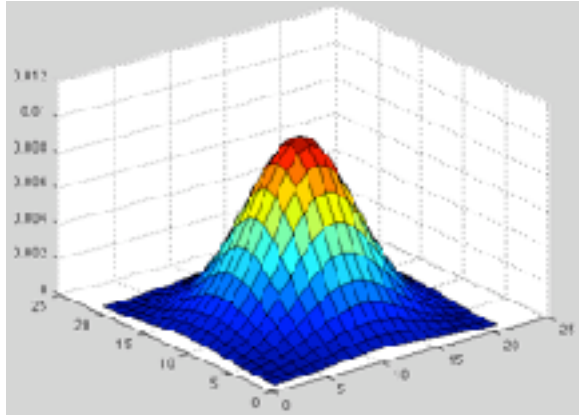
$$g[m, n; \sigma] = \exp\left(-\frac{m^2 + n^2}{2\sigma^2}\right)$$



Gaussian filter for low-pass filtering



Dali



Properties of the Gaussian filter

$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

- The n-dimensional Gaussian is the only completely circularly symmetric operator that is separable.
- The (continuous) Fourier transform of a Gaussian is another Gaussian

$$G(u, v; \sigma) = \exp\left(-2\pi^2(u^2 + v^2)\sigma^2\right)$$

Properties of the Gaussian filter

$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

- The convolution of two n-dimensional Gaussians is an n-dimensional Gaussian.

$$g(x, y; \sigma_1) \circ g(x, y; \sigma_2) = g(x, y; \sigma_3)$$

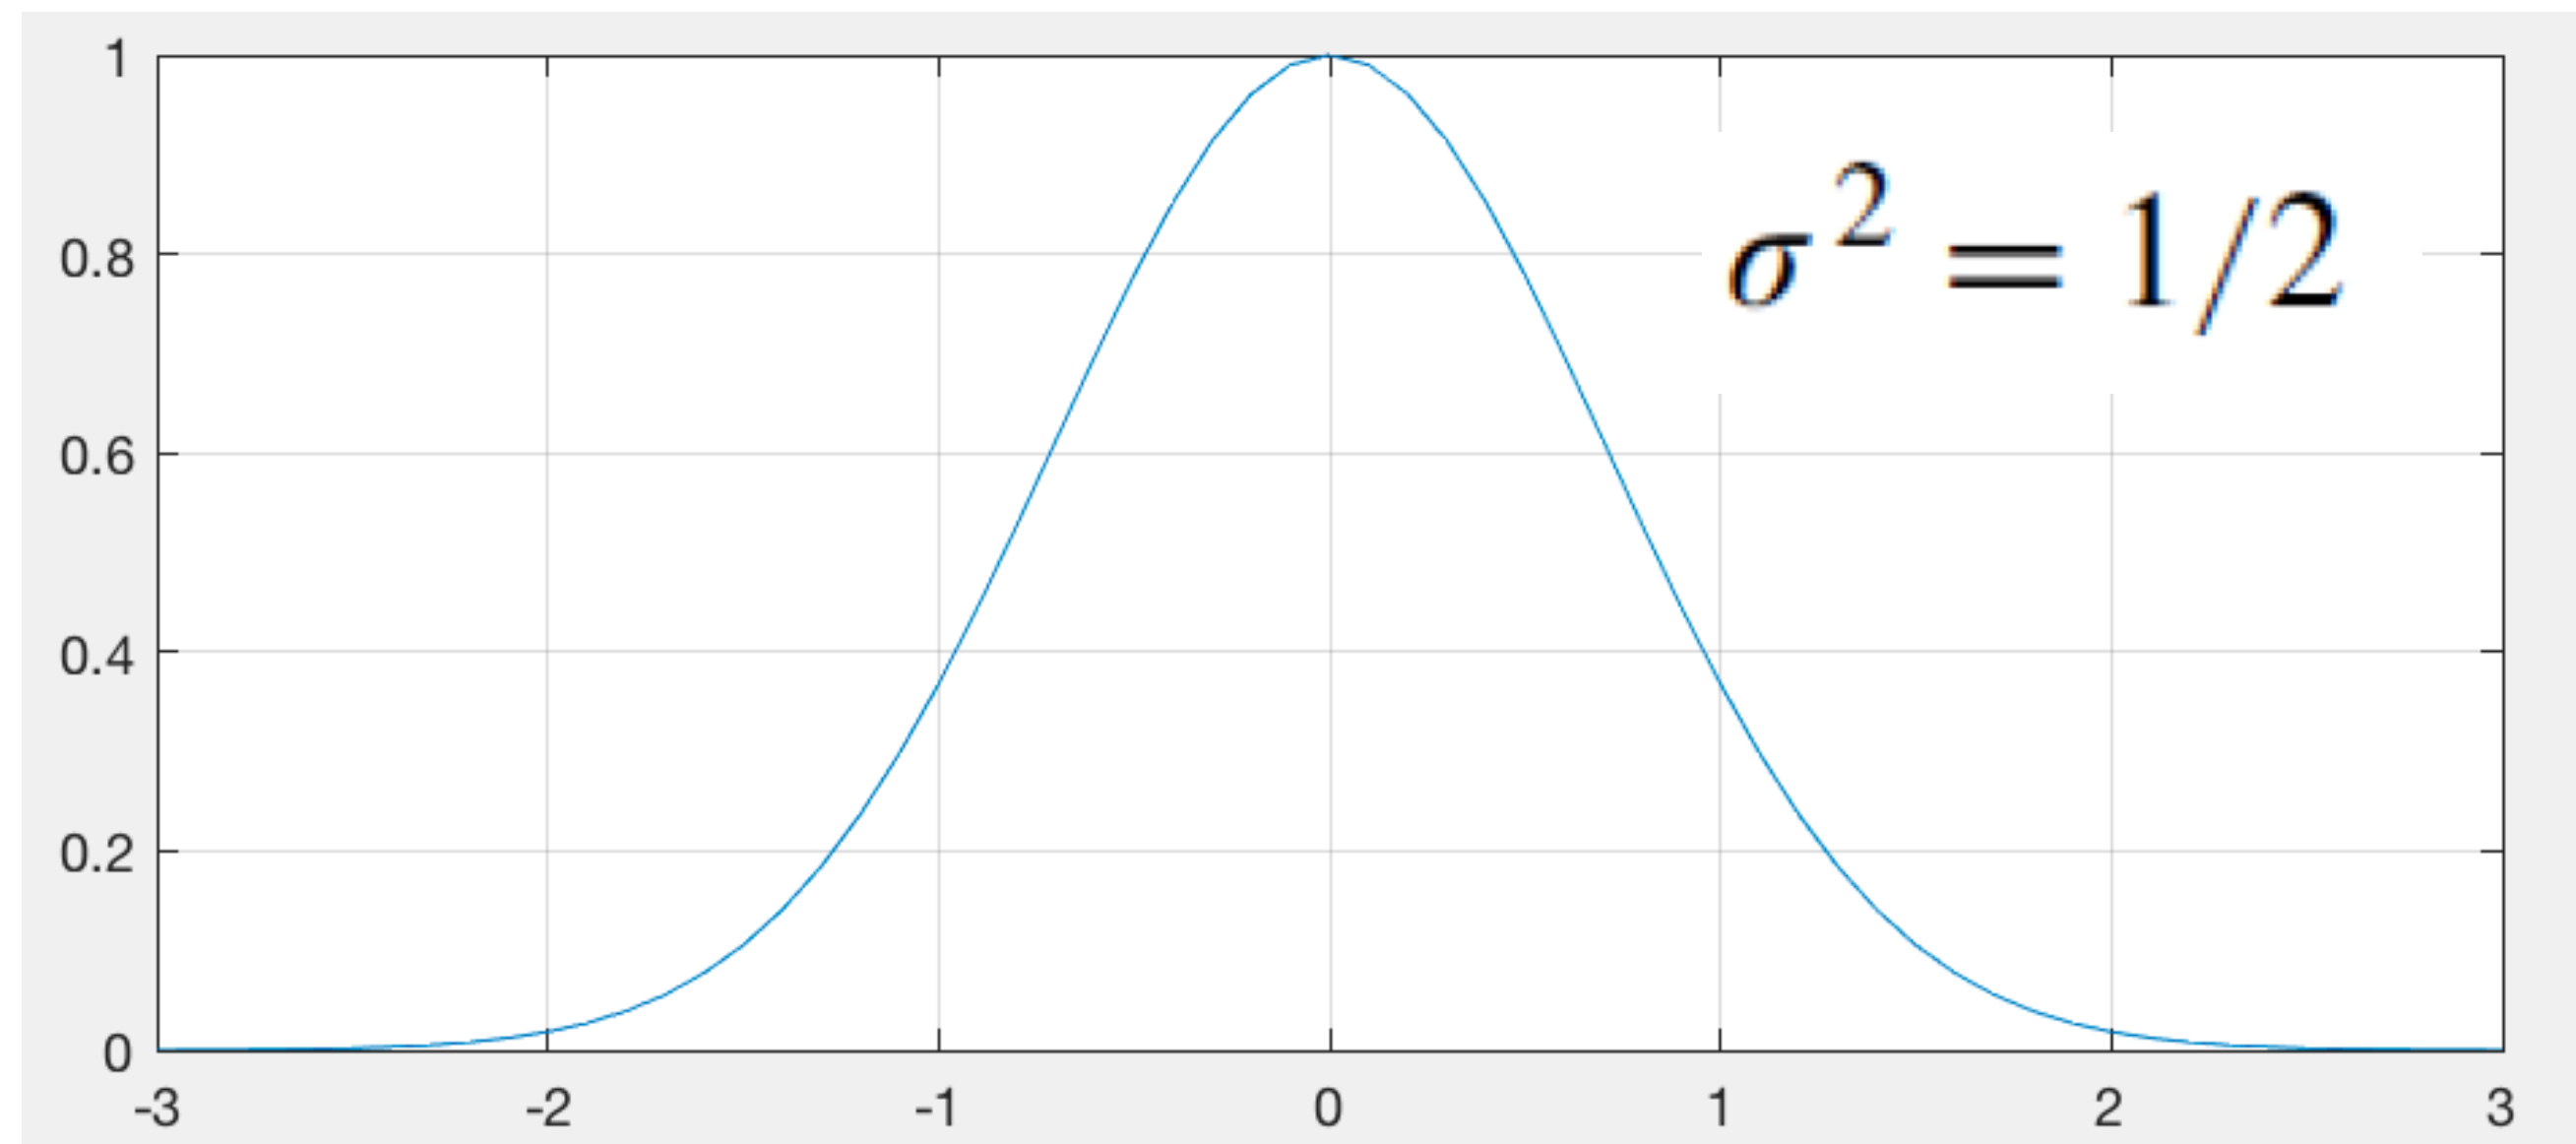
where the variance of the result is the sum

$$\sigma_3^2 = \sigma_1^2 + \sigma_2^2$$

(it is easy to prove this using the FT of the Gaussian)

Discretization of the Gaussian

There are very efficient approximations to the Gaussian filter for certain values of σ with nicer properties than when working with discretized gaussians.



$$g_5 [n] = [0.0183, 0.3679, 1.0000, 0.3679, 0.0183]$$

Binomial filter

Binomial coefficients provide a compact approximation of the gaussian coefficients using only integers.

The simplest blur filter (low pass) is

$$[1 \ 1]$$

Binomial filters in the family of filters obtained as successive convolutions of $[1 \ 1]$

Binomial filter

$$\mathbf{b}_1 = [1 \ 1]$$

$$\mathbf{b}_2 = [1 \ 1] \circ [1 \ 1] = [1 \ 2 \ 1]$$

$$\mathbf{b}_3 = [1 \ 1] \circ [1 \ 1] \circ [1 \ 1] = [1 \ 3 \ 3 \ 1]$$

Binomial filter

b_1					1		1													$\sigma_1^2 = 1/4$		
b_2					1		2		1												$\sigma_2^2 = 1/2$	
b_3					1		3		3		1										$\sigma_3^2 = 3/4$	
b_4					1		4		6		4		1								$\sigma_4^2 = 1$	
b_5					1		5		10		10		5		1						$\sigma_5^2 = 5/4$	
b_6					1		6		15		20		15		6		1				$\sigma_6^2 = 3/2$	
b_7					1		7		21		35		35		21		7		1		$\sigma_7^2 = 7/4$	
b_8					1		8		28		56		70		56		28		8		1	$\sigma_8^2 = 2$

Properties of binomial filters

- Sum of the values is 2^n
- The variance of b_n is $\sigma^2 = n/4$
- The convolution of two binomial filters is also a binomial filter

$$b_n \circ b_m = b_{n+m}$$

With a variance:

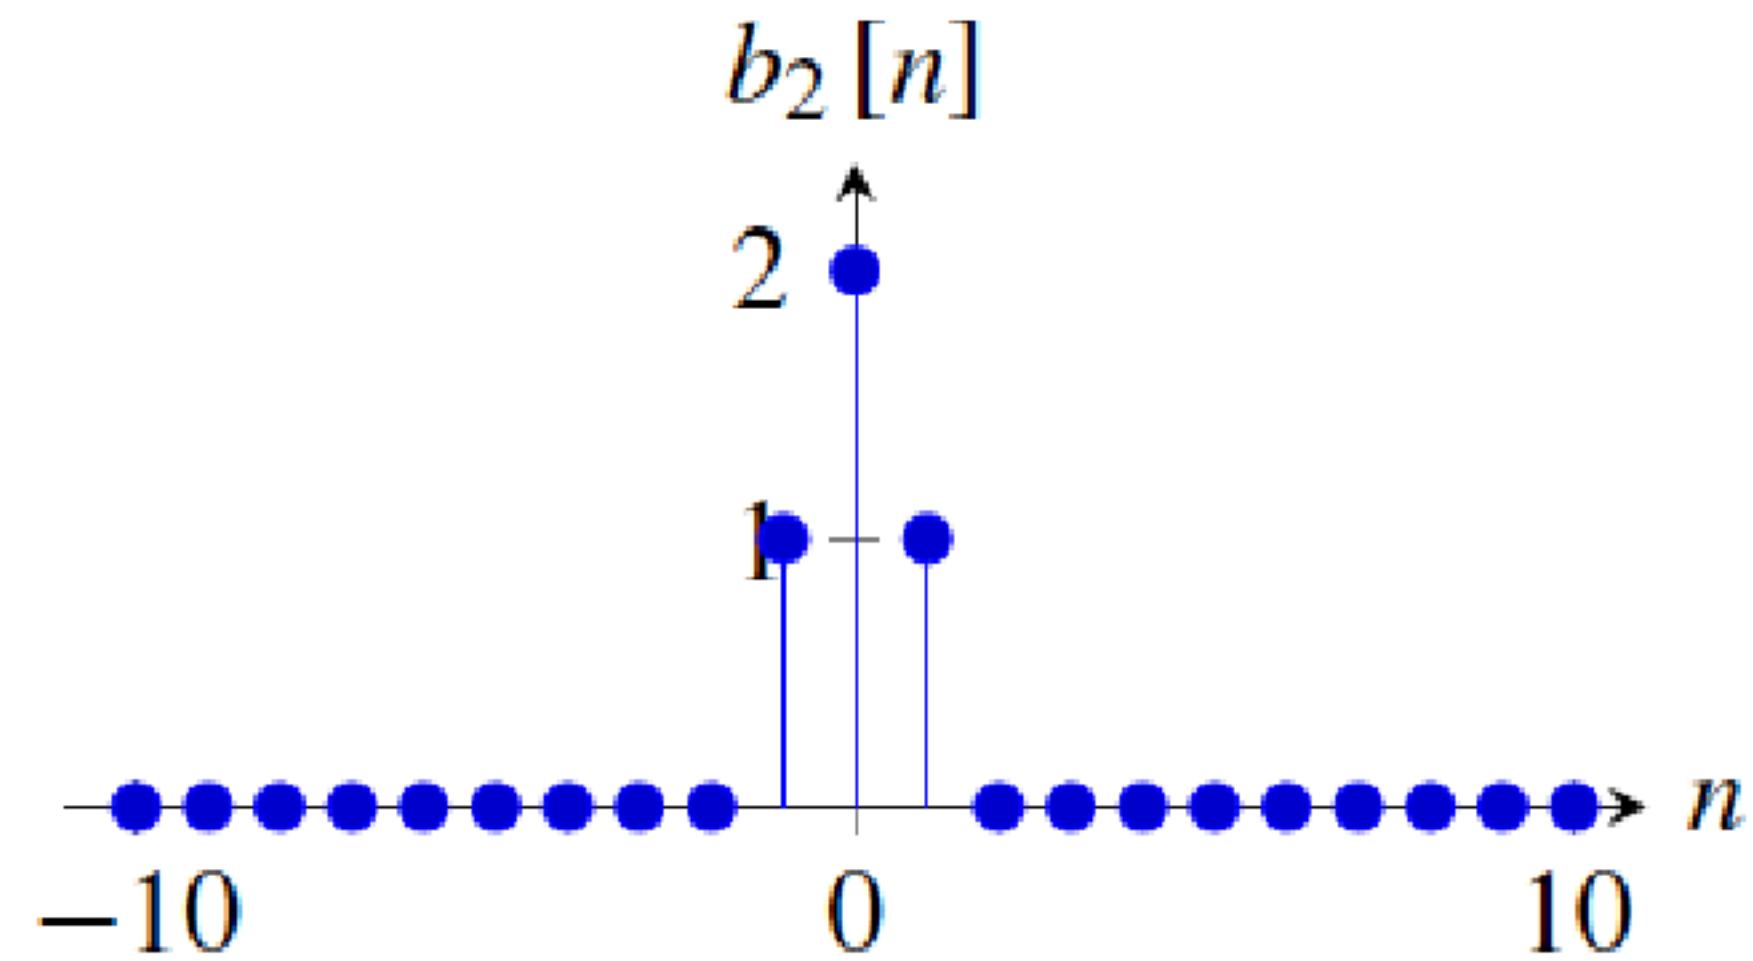
$$\sigma_n^2 + \sigma_m^2 = \sigma_{n+m}^2$$

These properties are analogous to the gaussian property in the continuous domain (but the binomial filter is different than a discretization of a gaussian)

B2[n]

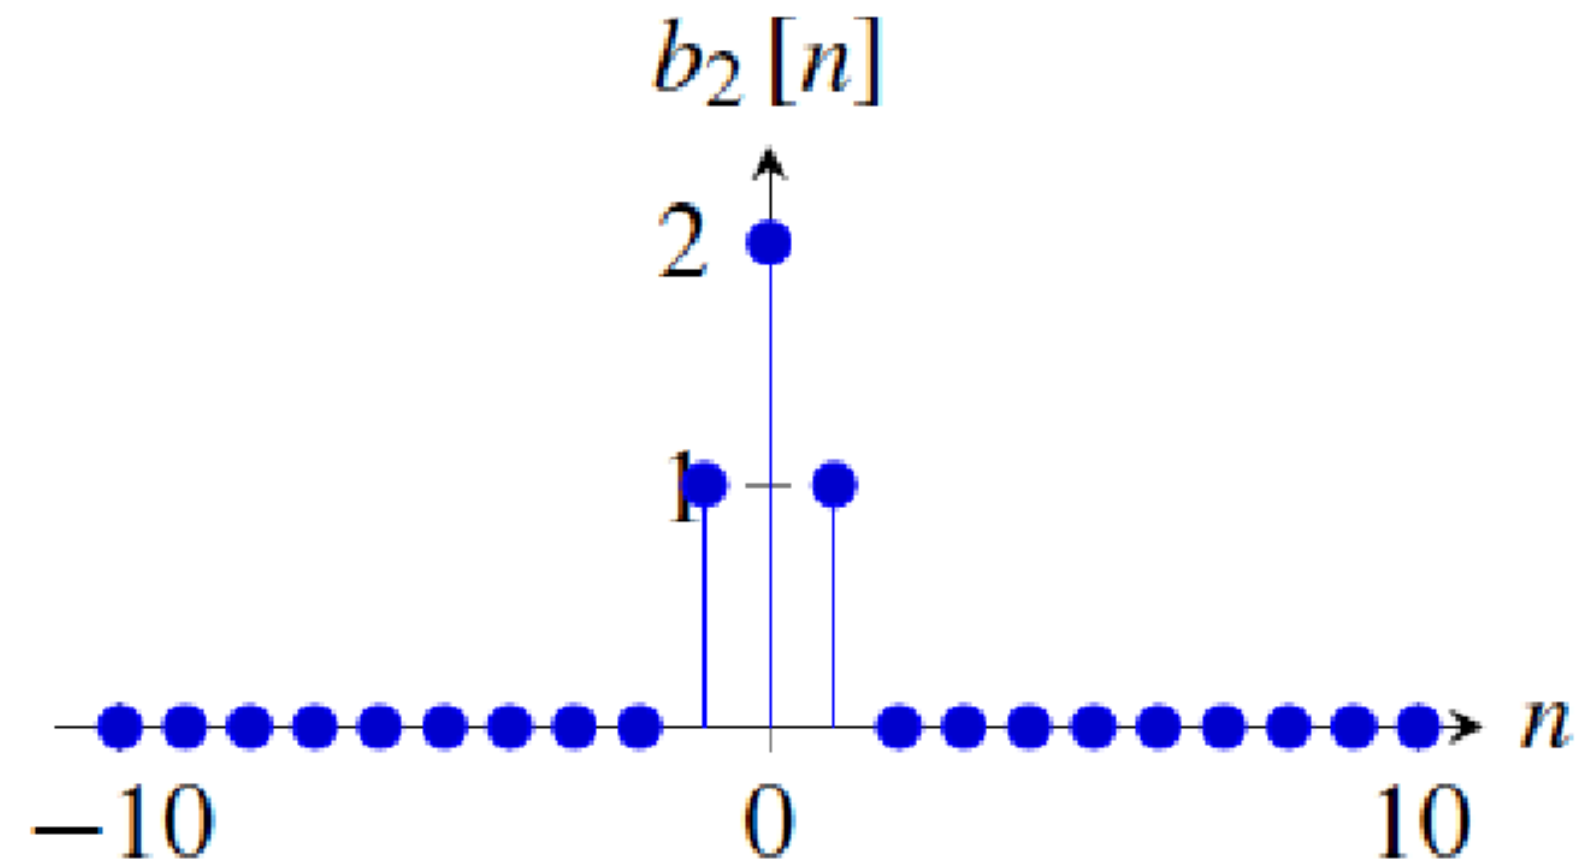
The simplest approximation to the Gaussian filter is the 3-tap kernel:

$$b_2 = [1, 2, 1]$$

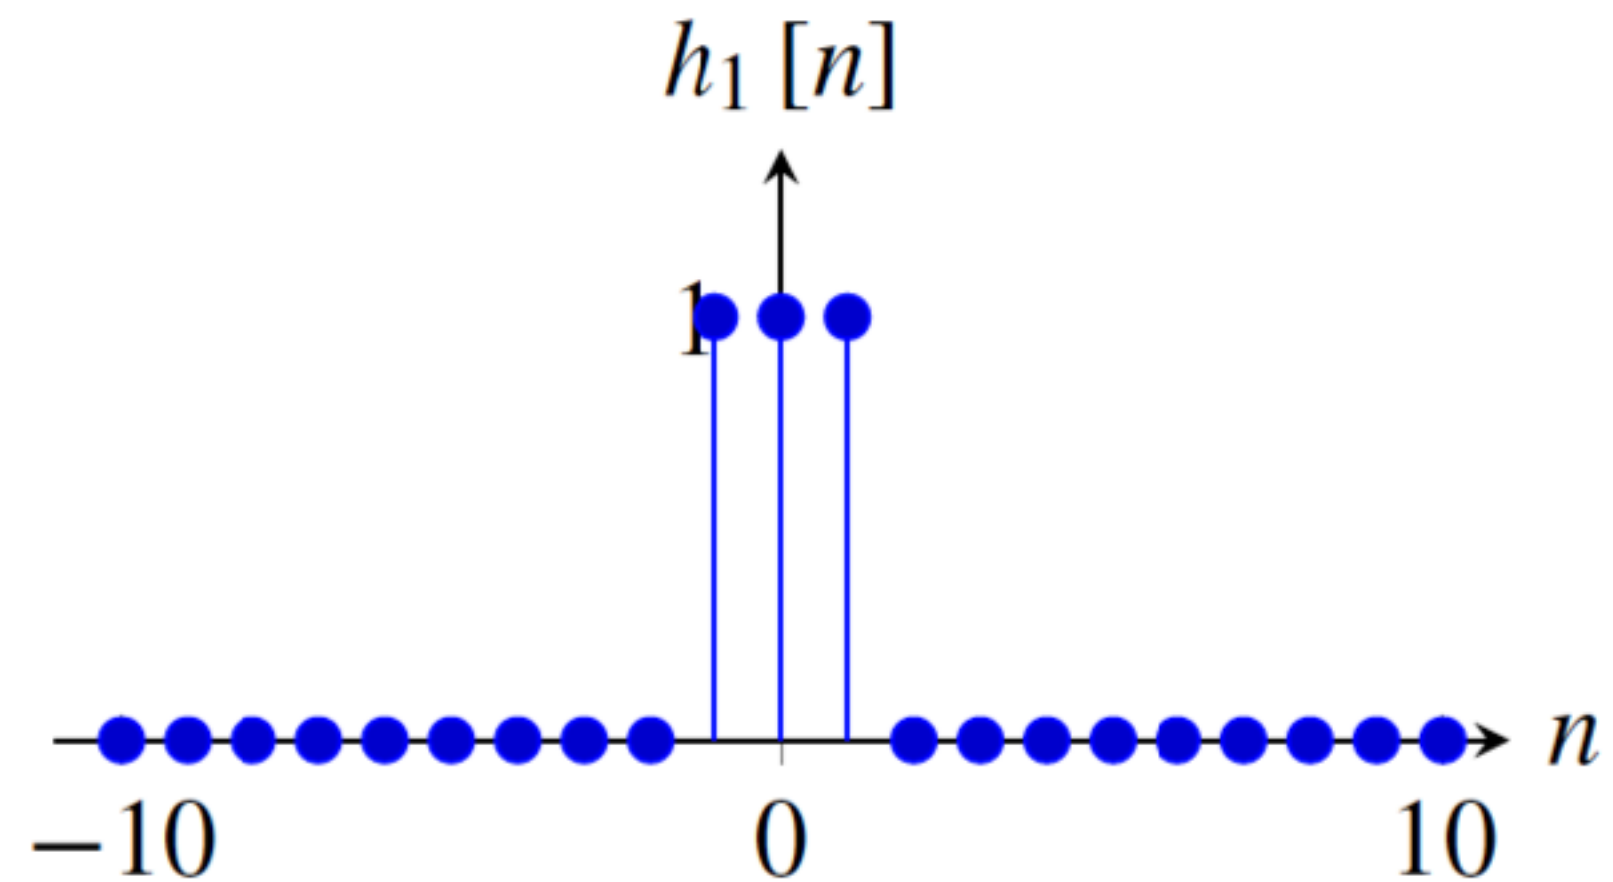


B2[n] versus the 3-tap box filter

[1 2 1]



[1 1 1]



Which one is better?

B2[n]

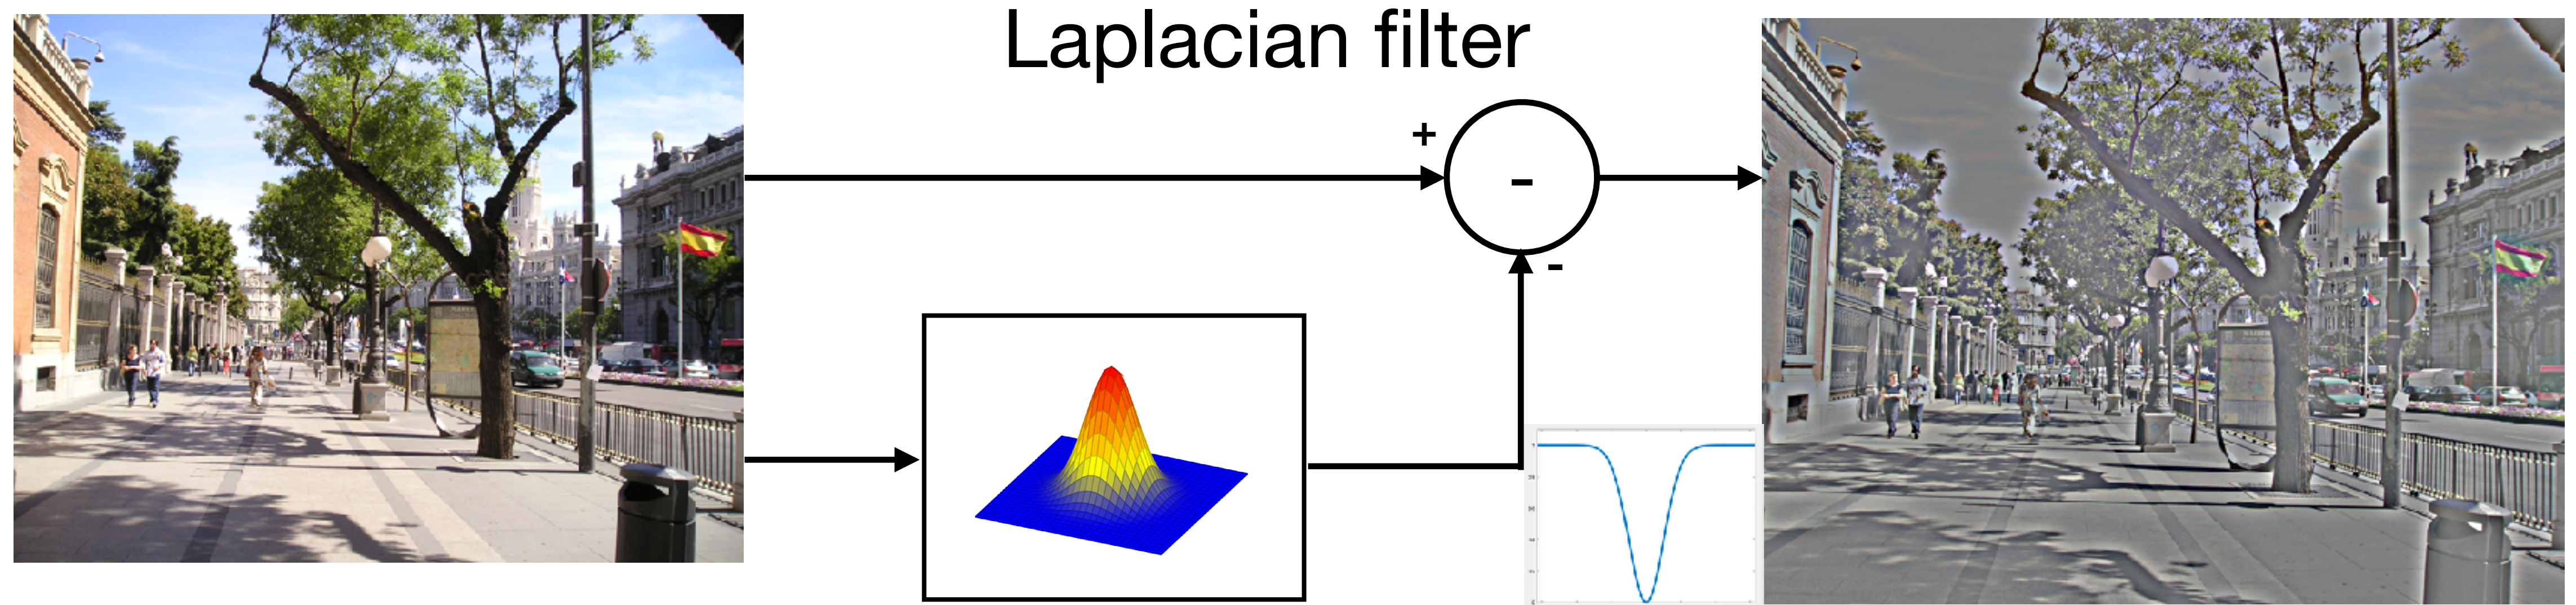
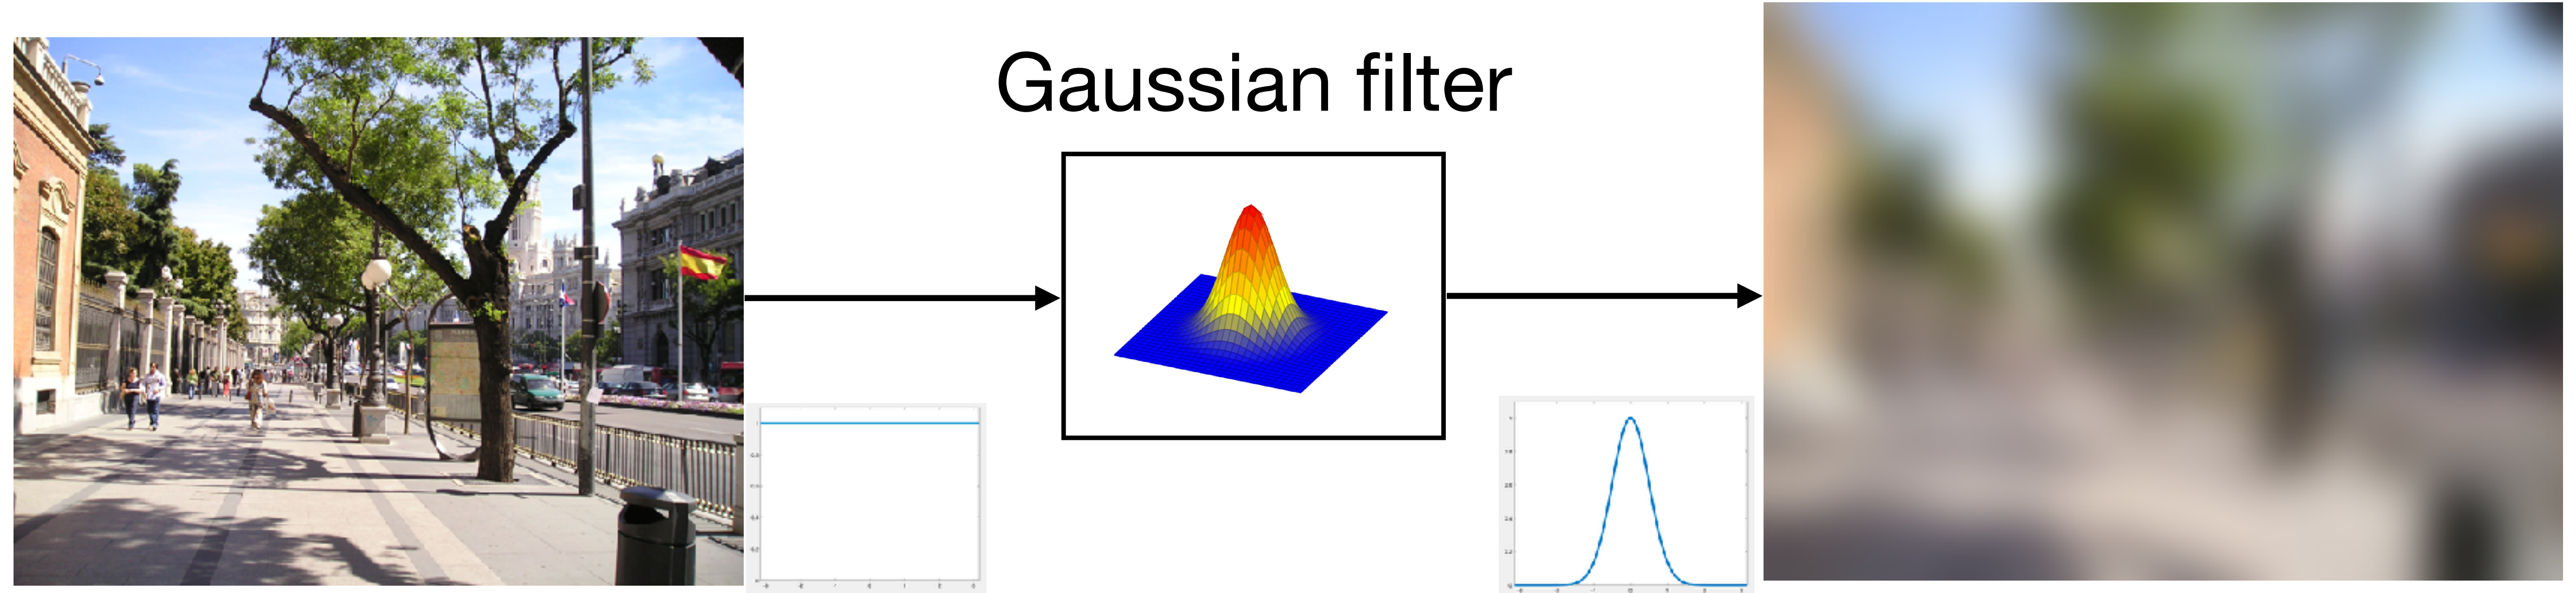
$$[1, 1, 1] \circ [\dots, 1, -1, 1, -1, 1, -1, \dots] = [\dots, -1, 1, -1, 1, -1, 1, \dots]$$

$$[1, 2, 1] \circ [\dots, 1, -1, 1, -1, 1, -1, \dots] = [\dots, 0, 0, 0, 0, 0, 0, \dots]$$

B2[n]

$$b_{2,2} = b_{2,0} \circ b_{0,2} = \begin{bmatrix} 1 & 2 & 1 \end{bmatrix} \circ \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

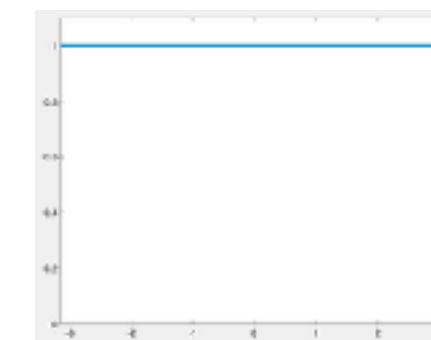
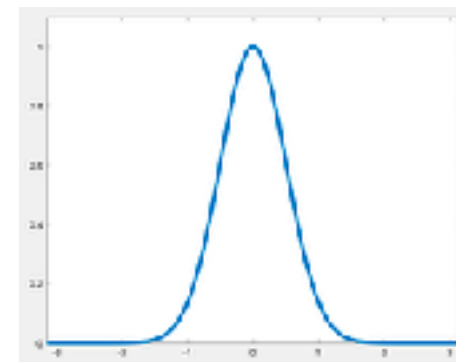
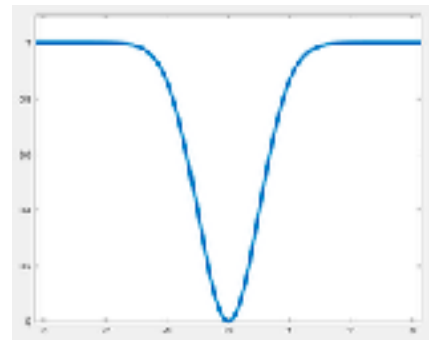
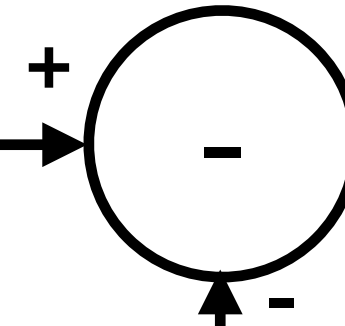
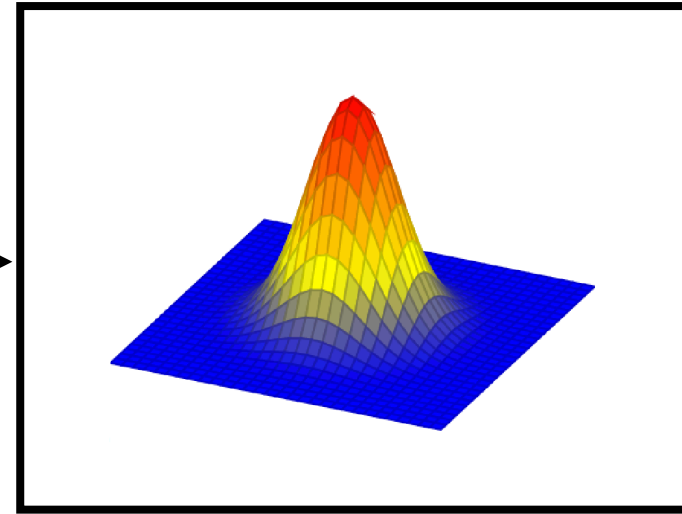
What about the opposite of blurring?



Laplacian filter



Gaussian filter



+

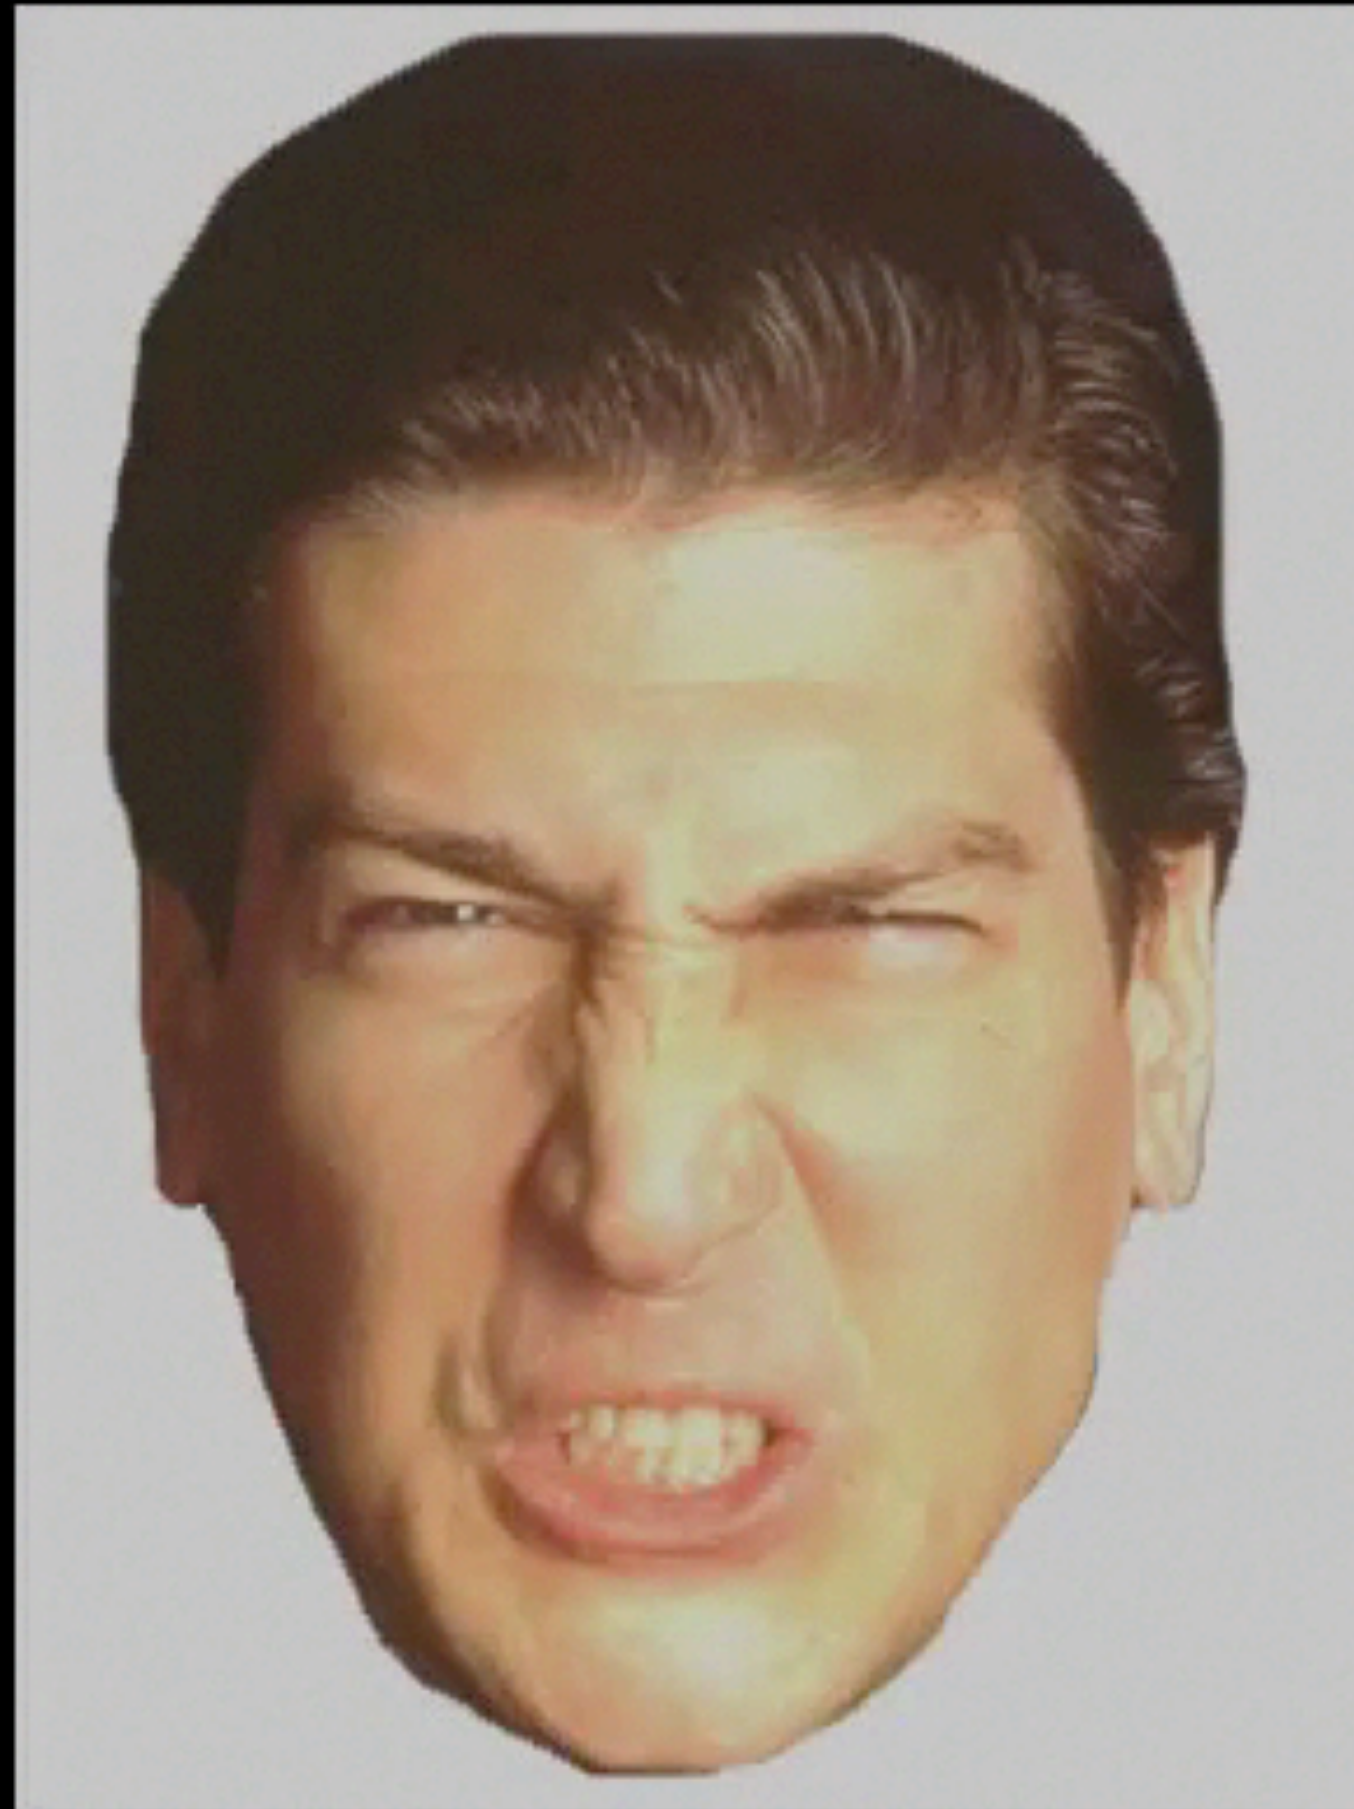


=



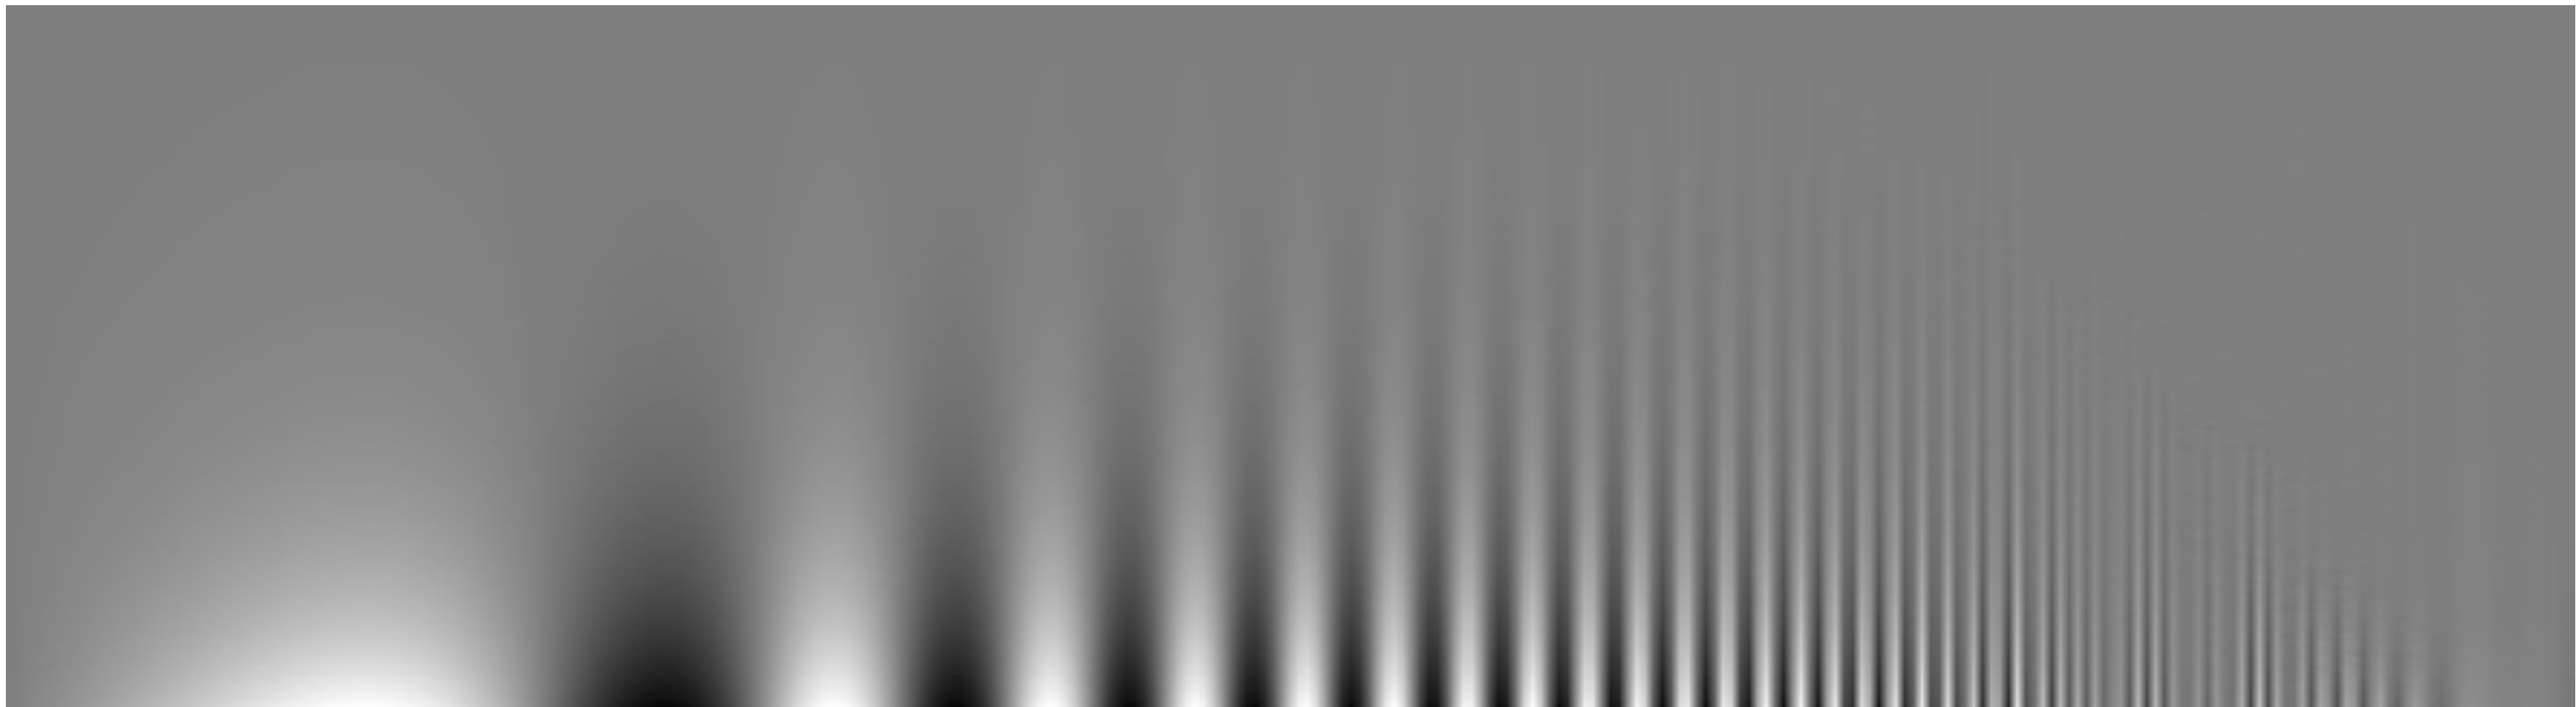
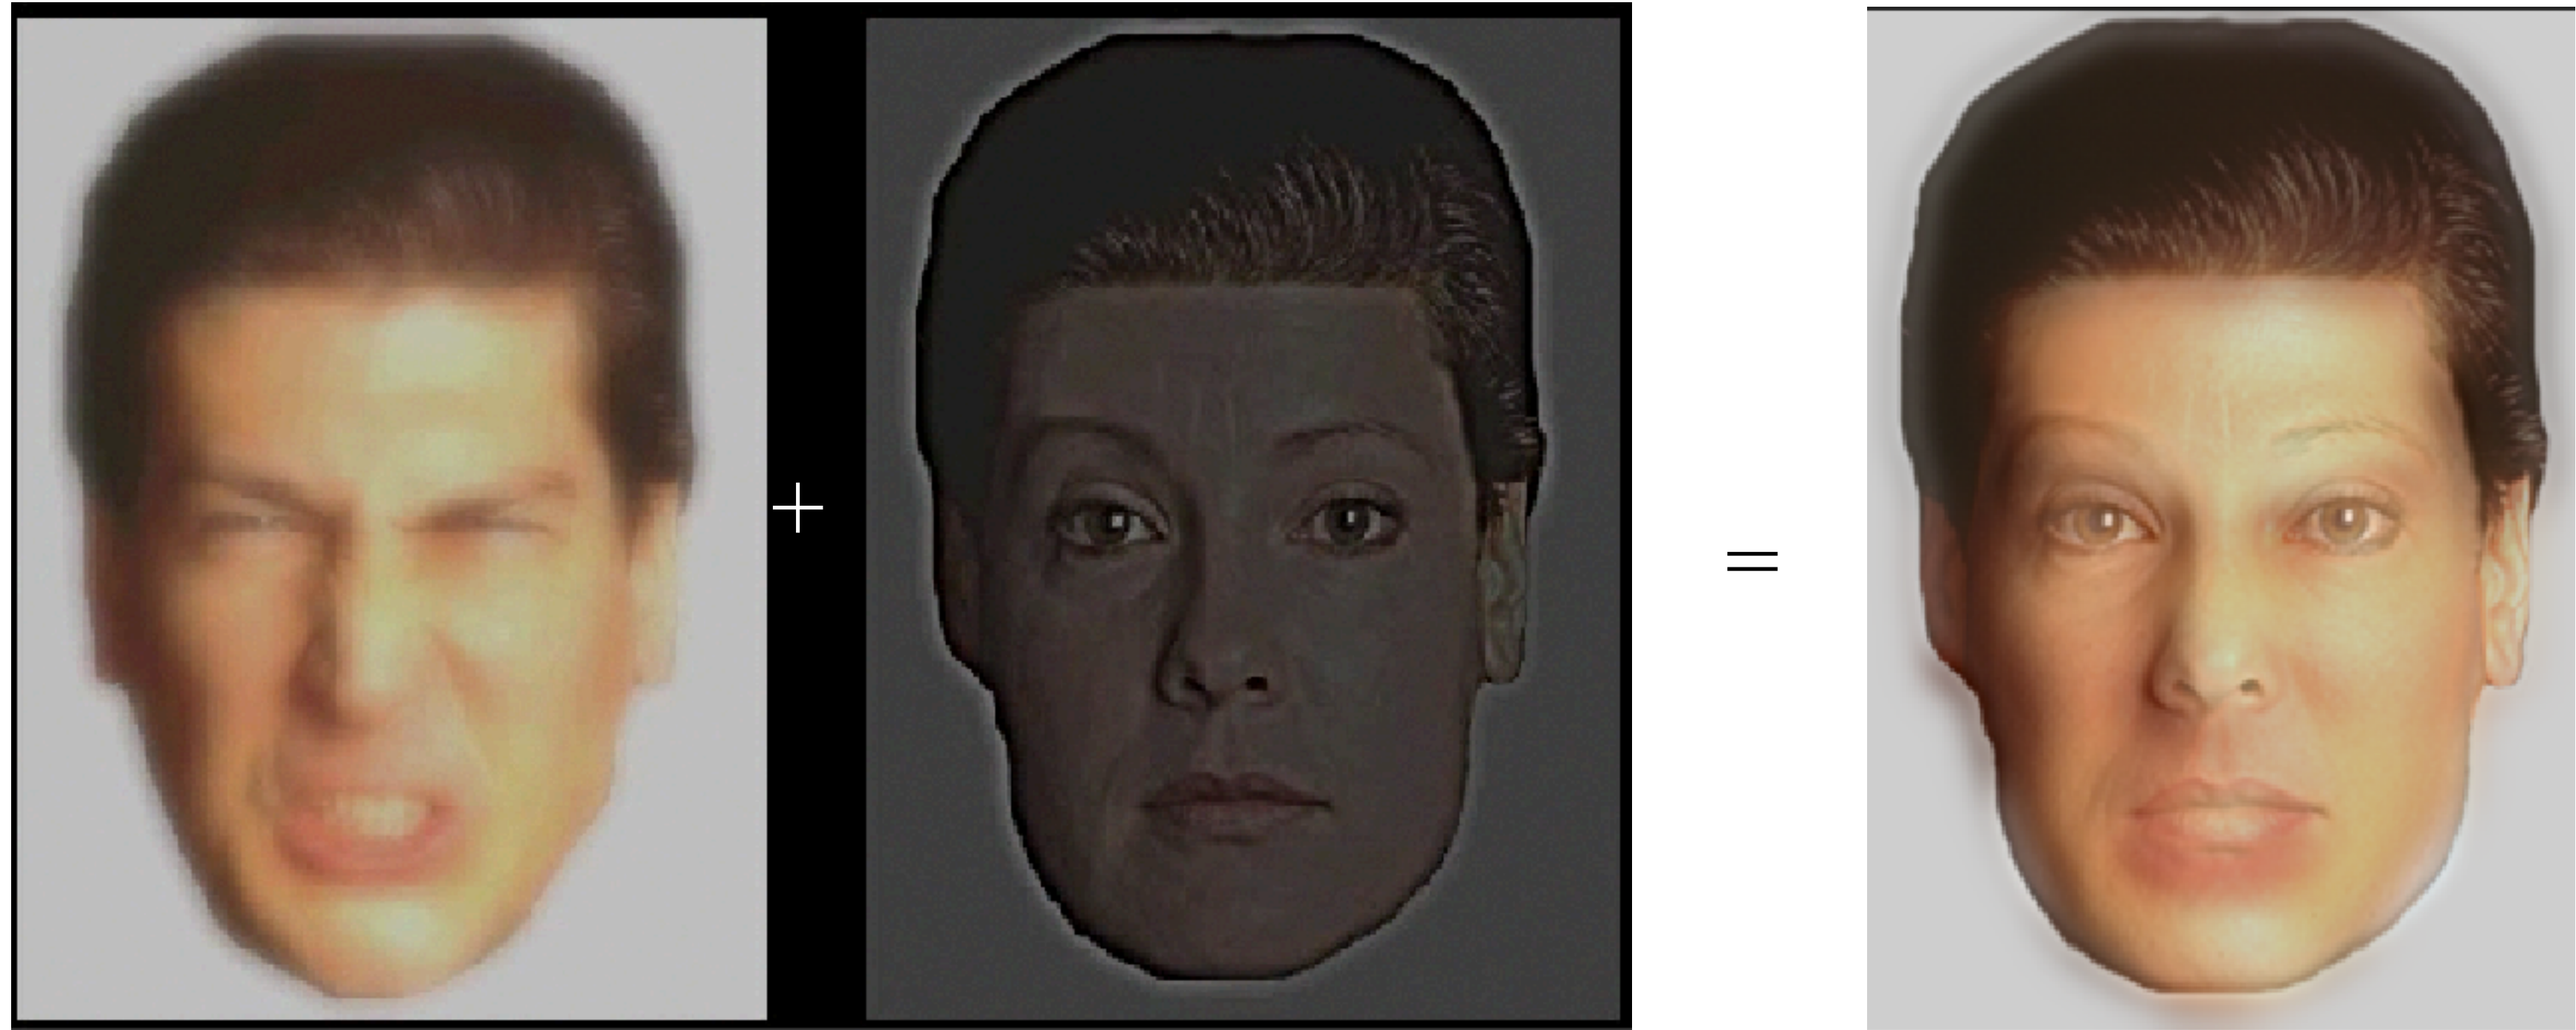
Hybrid Images

Oliva & Schyns

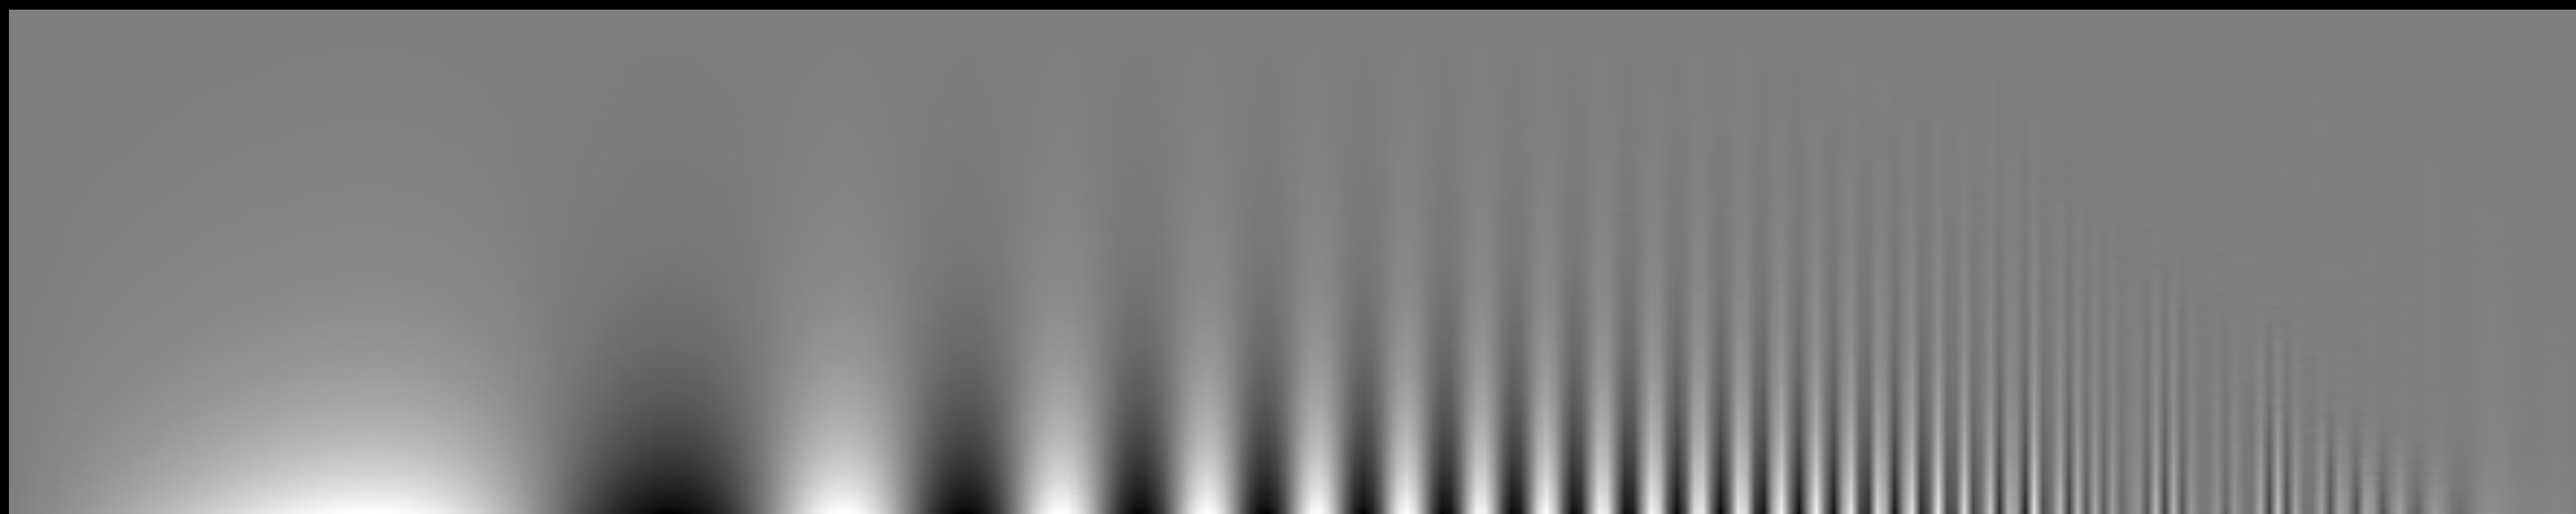


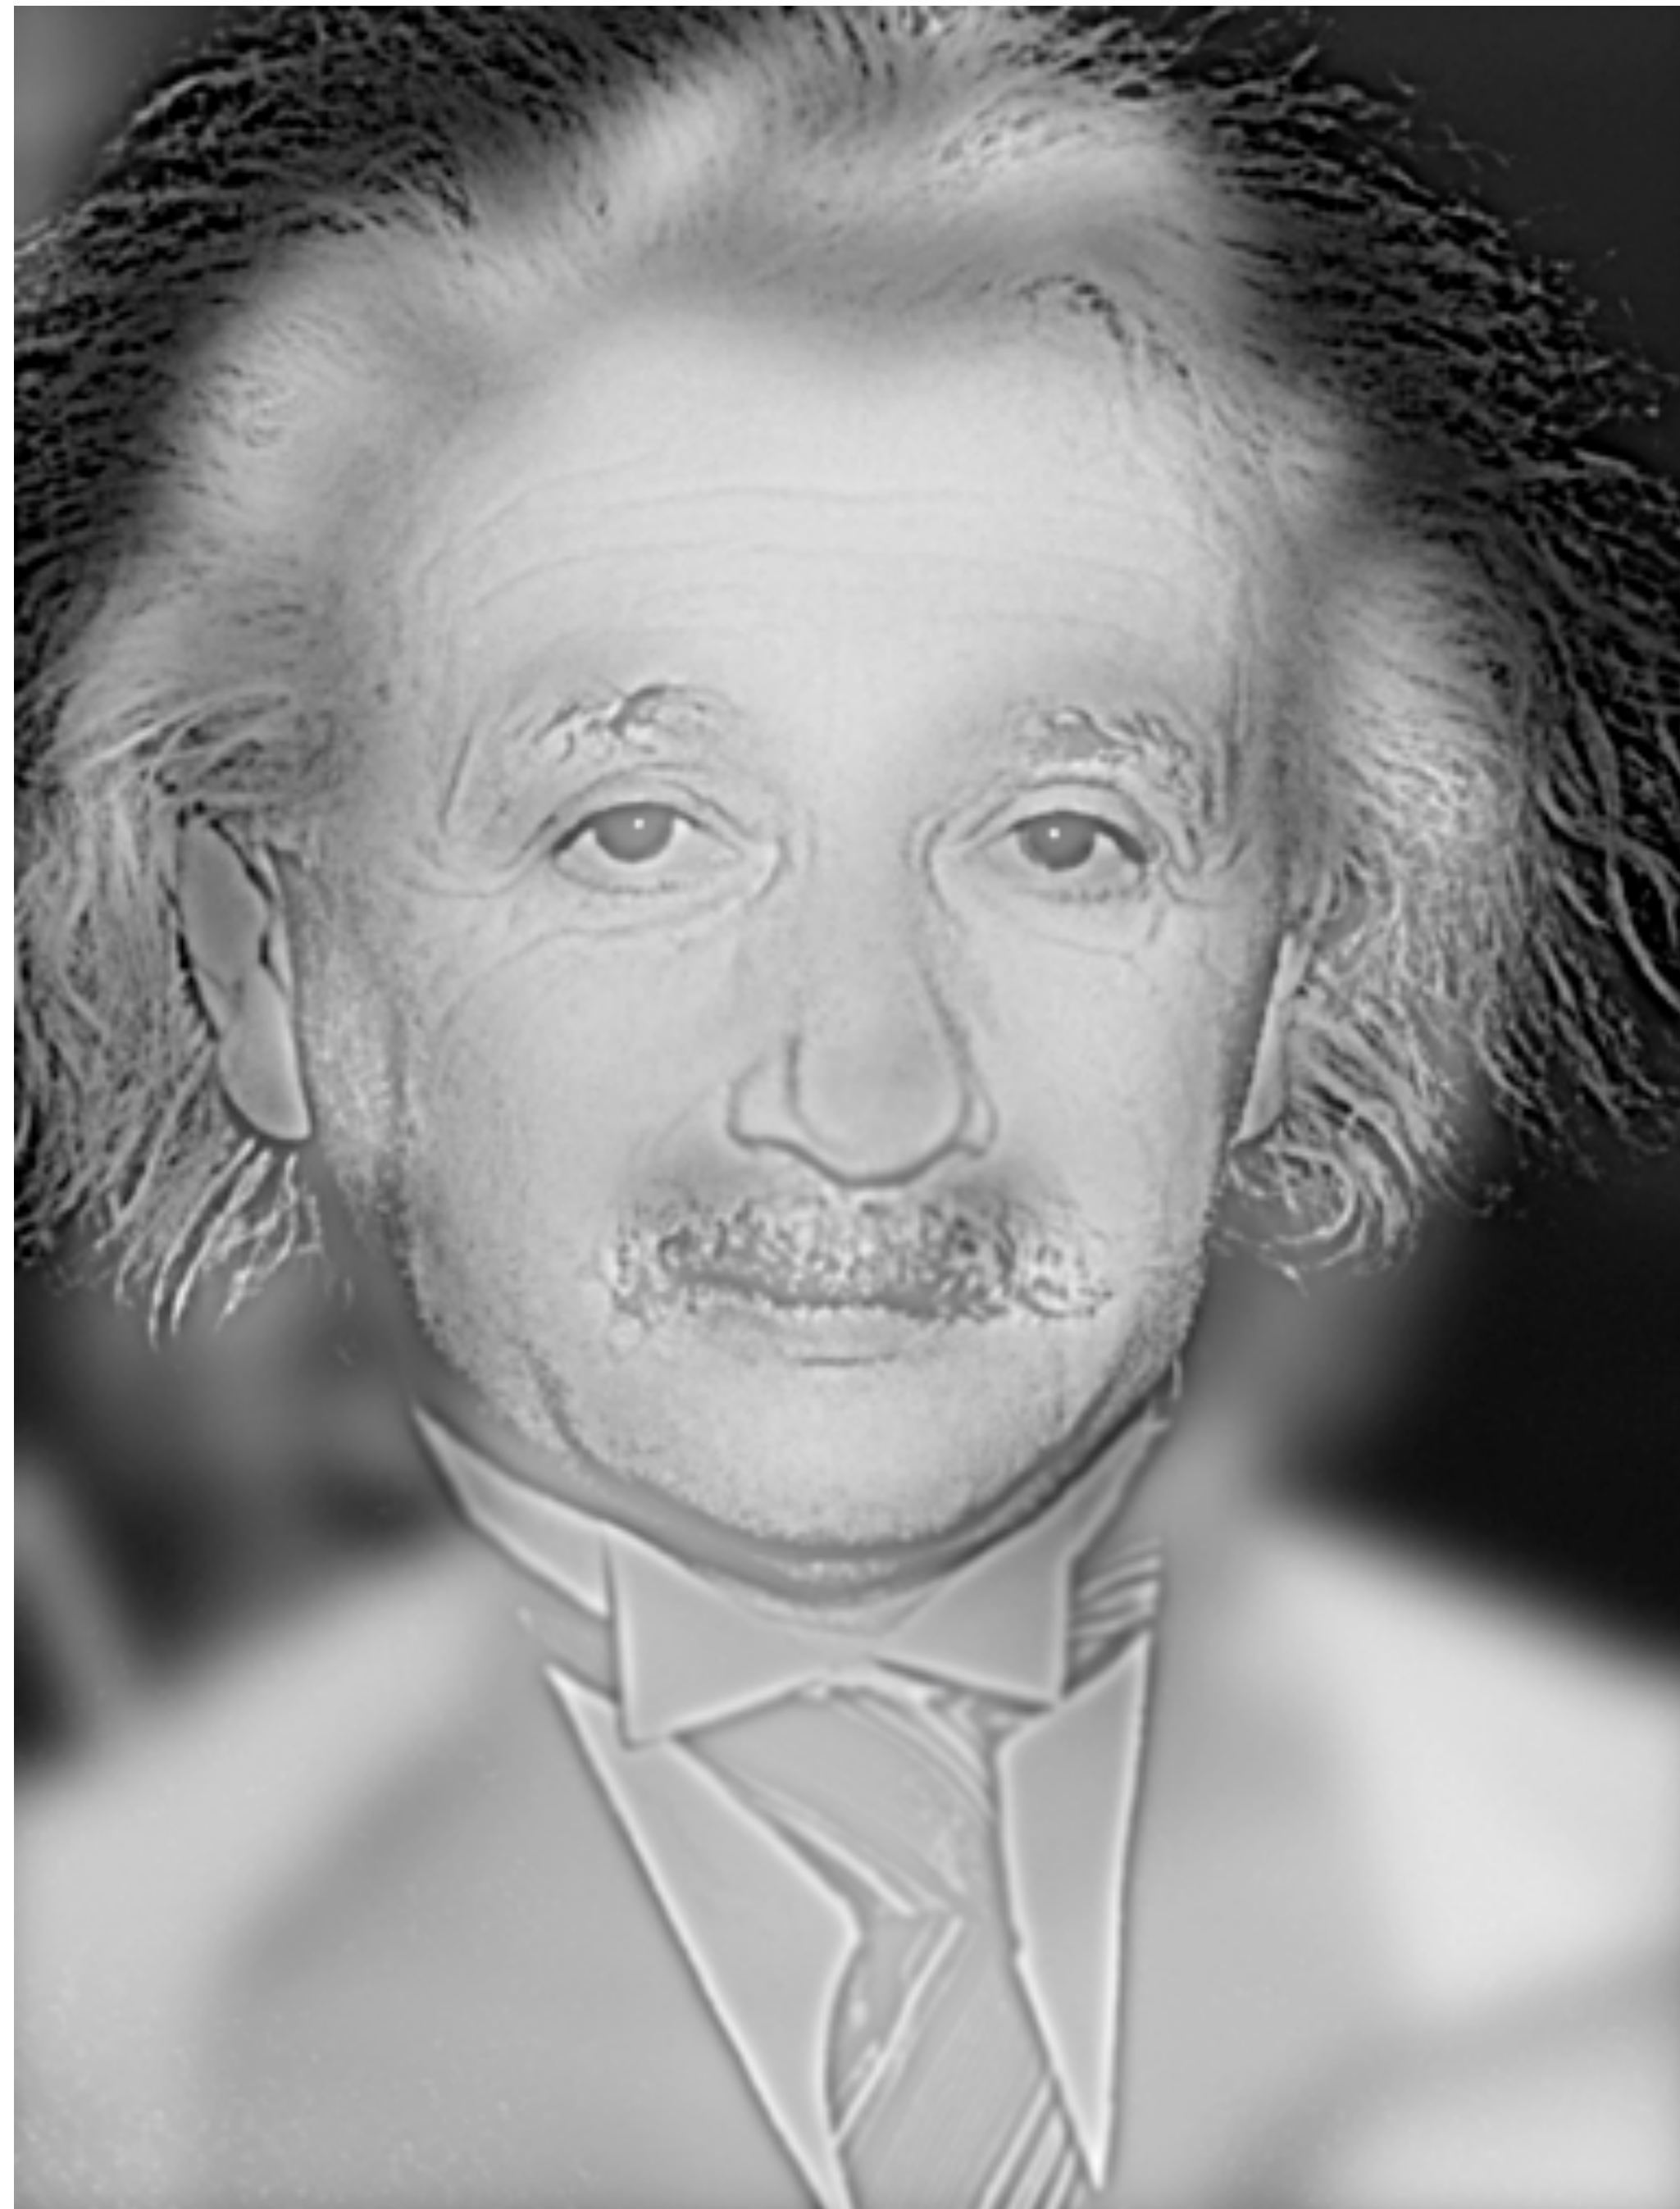
Hybrid Images

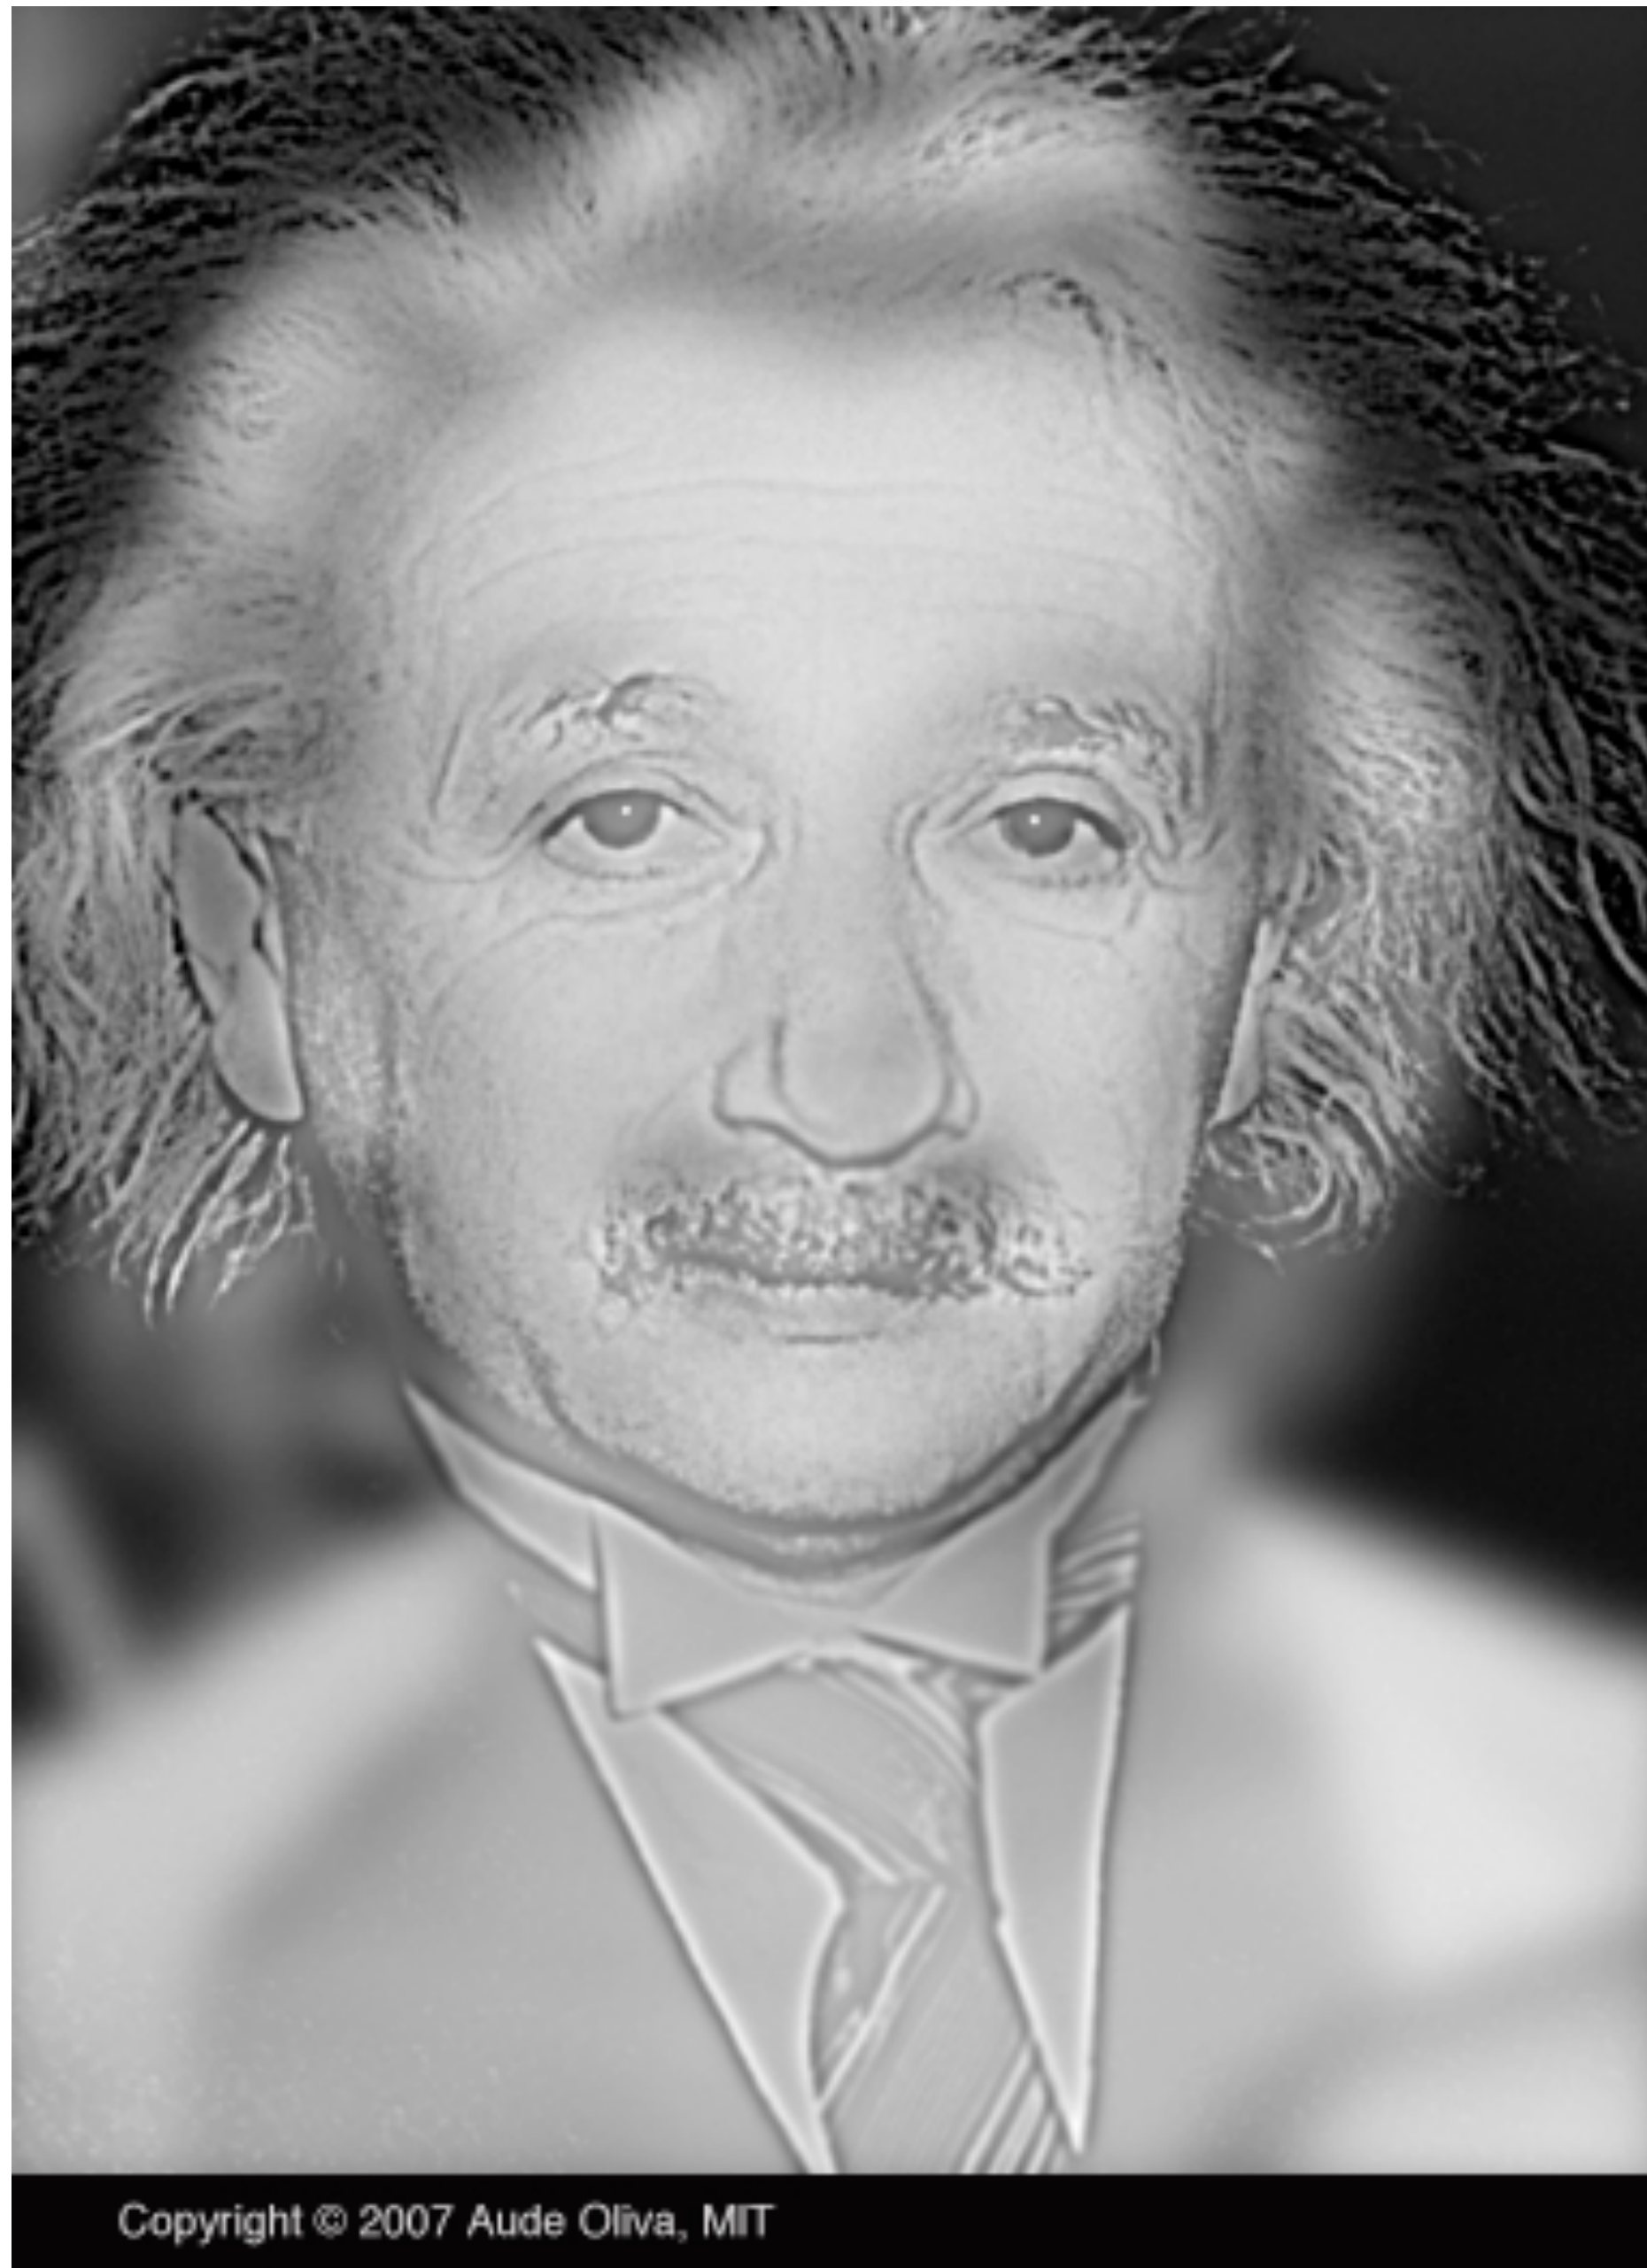




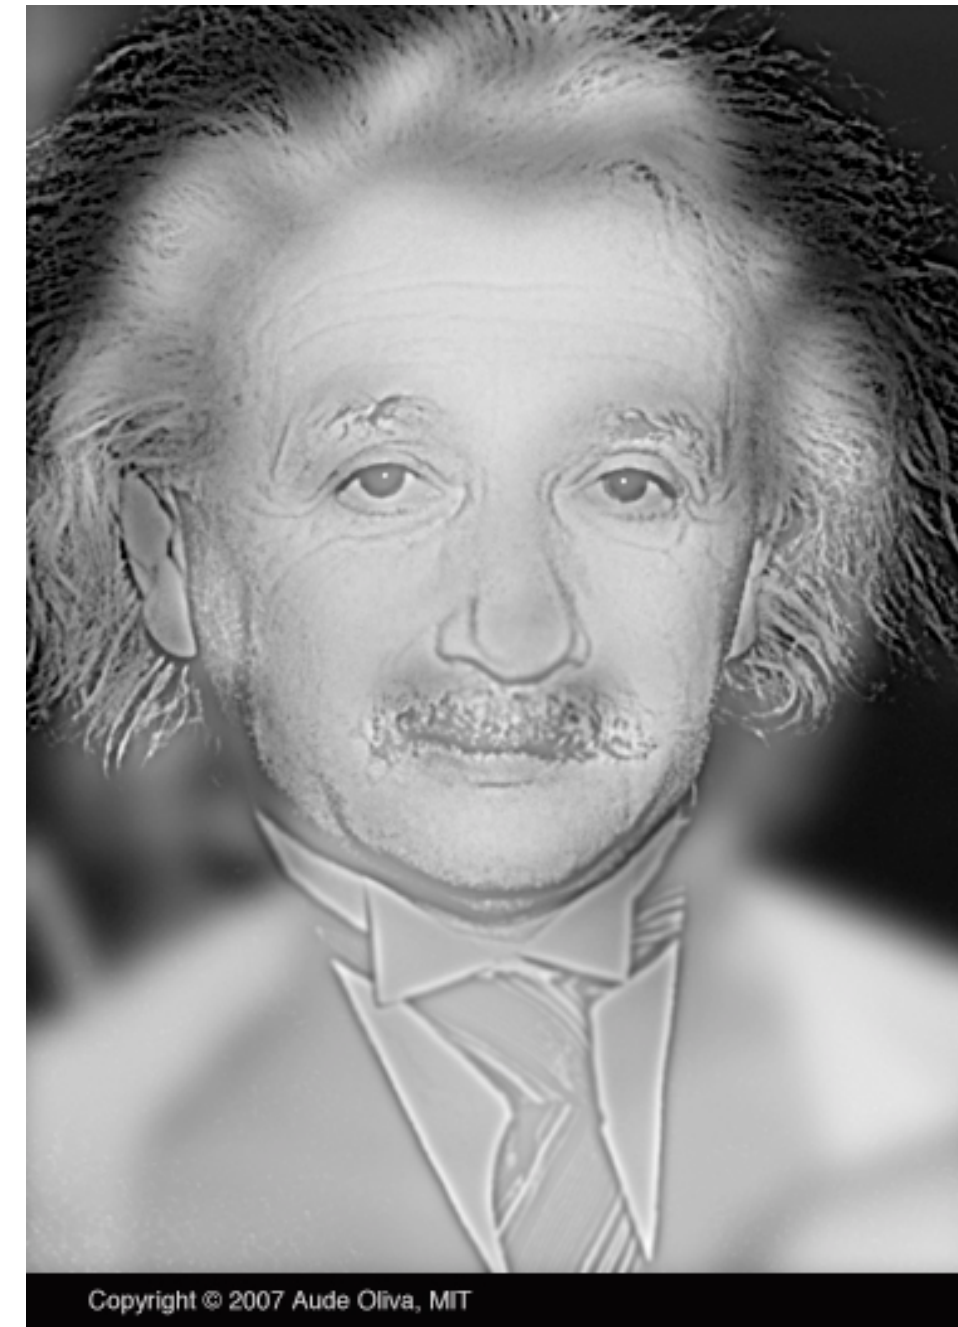
Hybrid Images



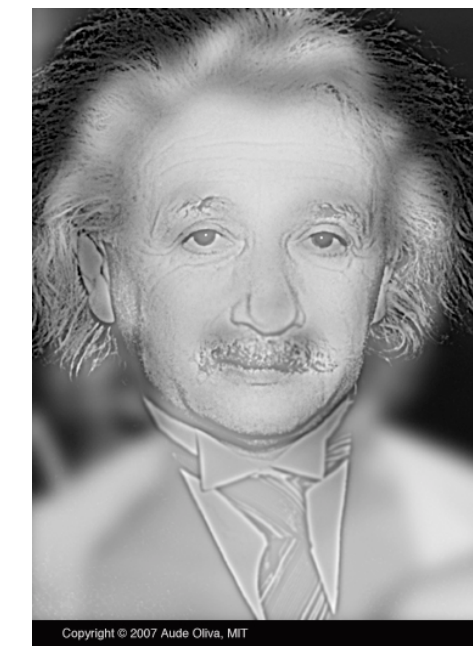




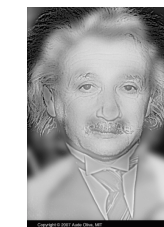
Copyright © 2007 Aude Oliva, MIT



Copyright © 2007 Aude Oliva, MIT



Copyright © 2007 Aude Oliva, MIT



DERIVATIVES

DERIVATIVES

High pass-filters

Finding edges in the image



Image gradient:

$$\nabla \mathbf{I} = \left(\frac{\partial \mathbf{I}}{\partial x}, \frac{\partial \mathbf{I}}{\partial y} \right)$$

Approximation image derivative:

$$\frac{\partial \mathbf{I}}{\partial x} \simeq \mathbf{I}(x, y) - \mathbf{I}(x - 1, y)$$

Edge strength

$$E(x, y) = |\nabla \mathbf{I}(x, y)|$$

Edge orientation:

$$\theta(x, y) = \angle \nabla \mathbf{I} = \arctan \frac{\partial \mathbf{I} / \partial y}{\partial \mathbf{I} / \partial x}$$

Edge normal:

$$\mathbf{n} = \frac{\nabla \mathbf{I}}{|\nabla \mathbf{I}|}$$

$$[-1 \ 1]$$

$$\frac{\partial I}{\partial x} \simeq I(x, y) - I(x - 1, y)$$



$g[m,n]$

\otimes

$[-1, 1]$

$=$

$h[m,n]$



$f[m,n]$

$$[-1 \ 1]^T$$

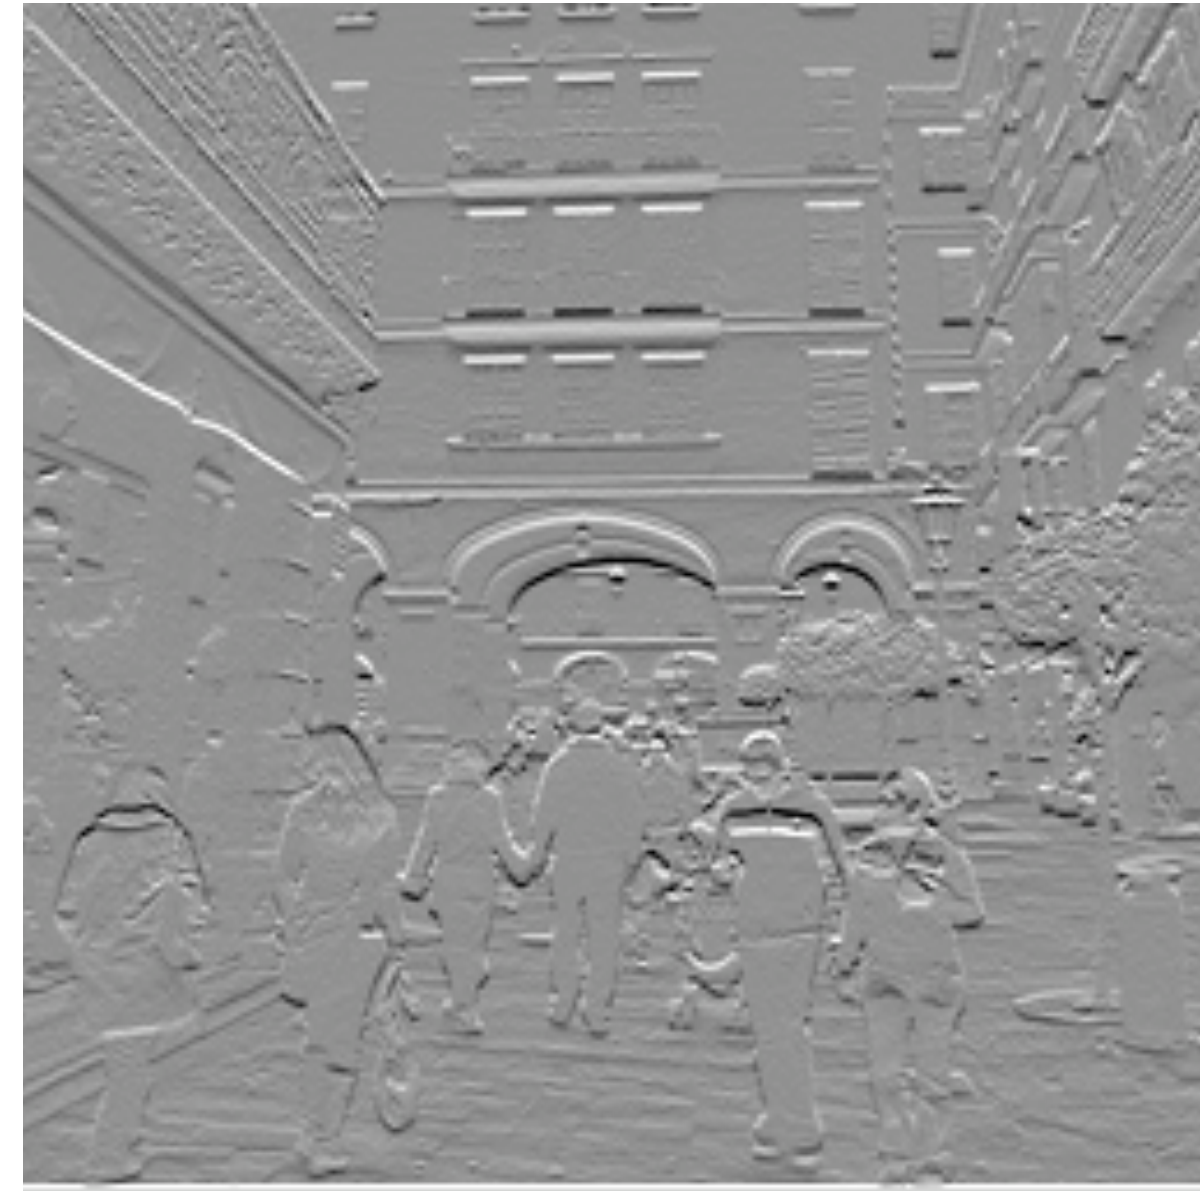


$g[m,n]$

\otimes

$$[-1, 1]^T =$$

$$h[m,n]$$



$f[m,n]$

Discrete derivatives

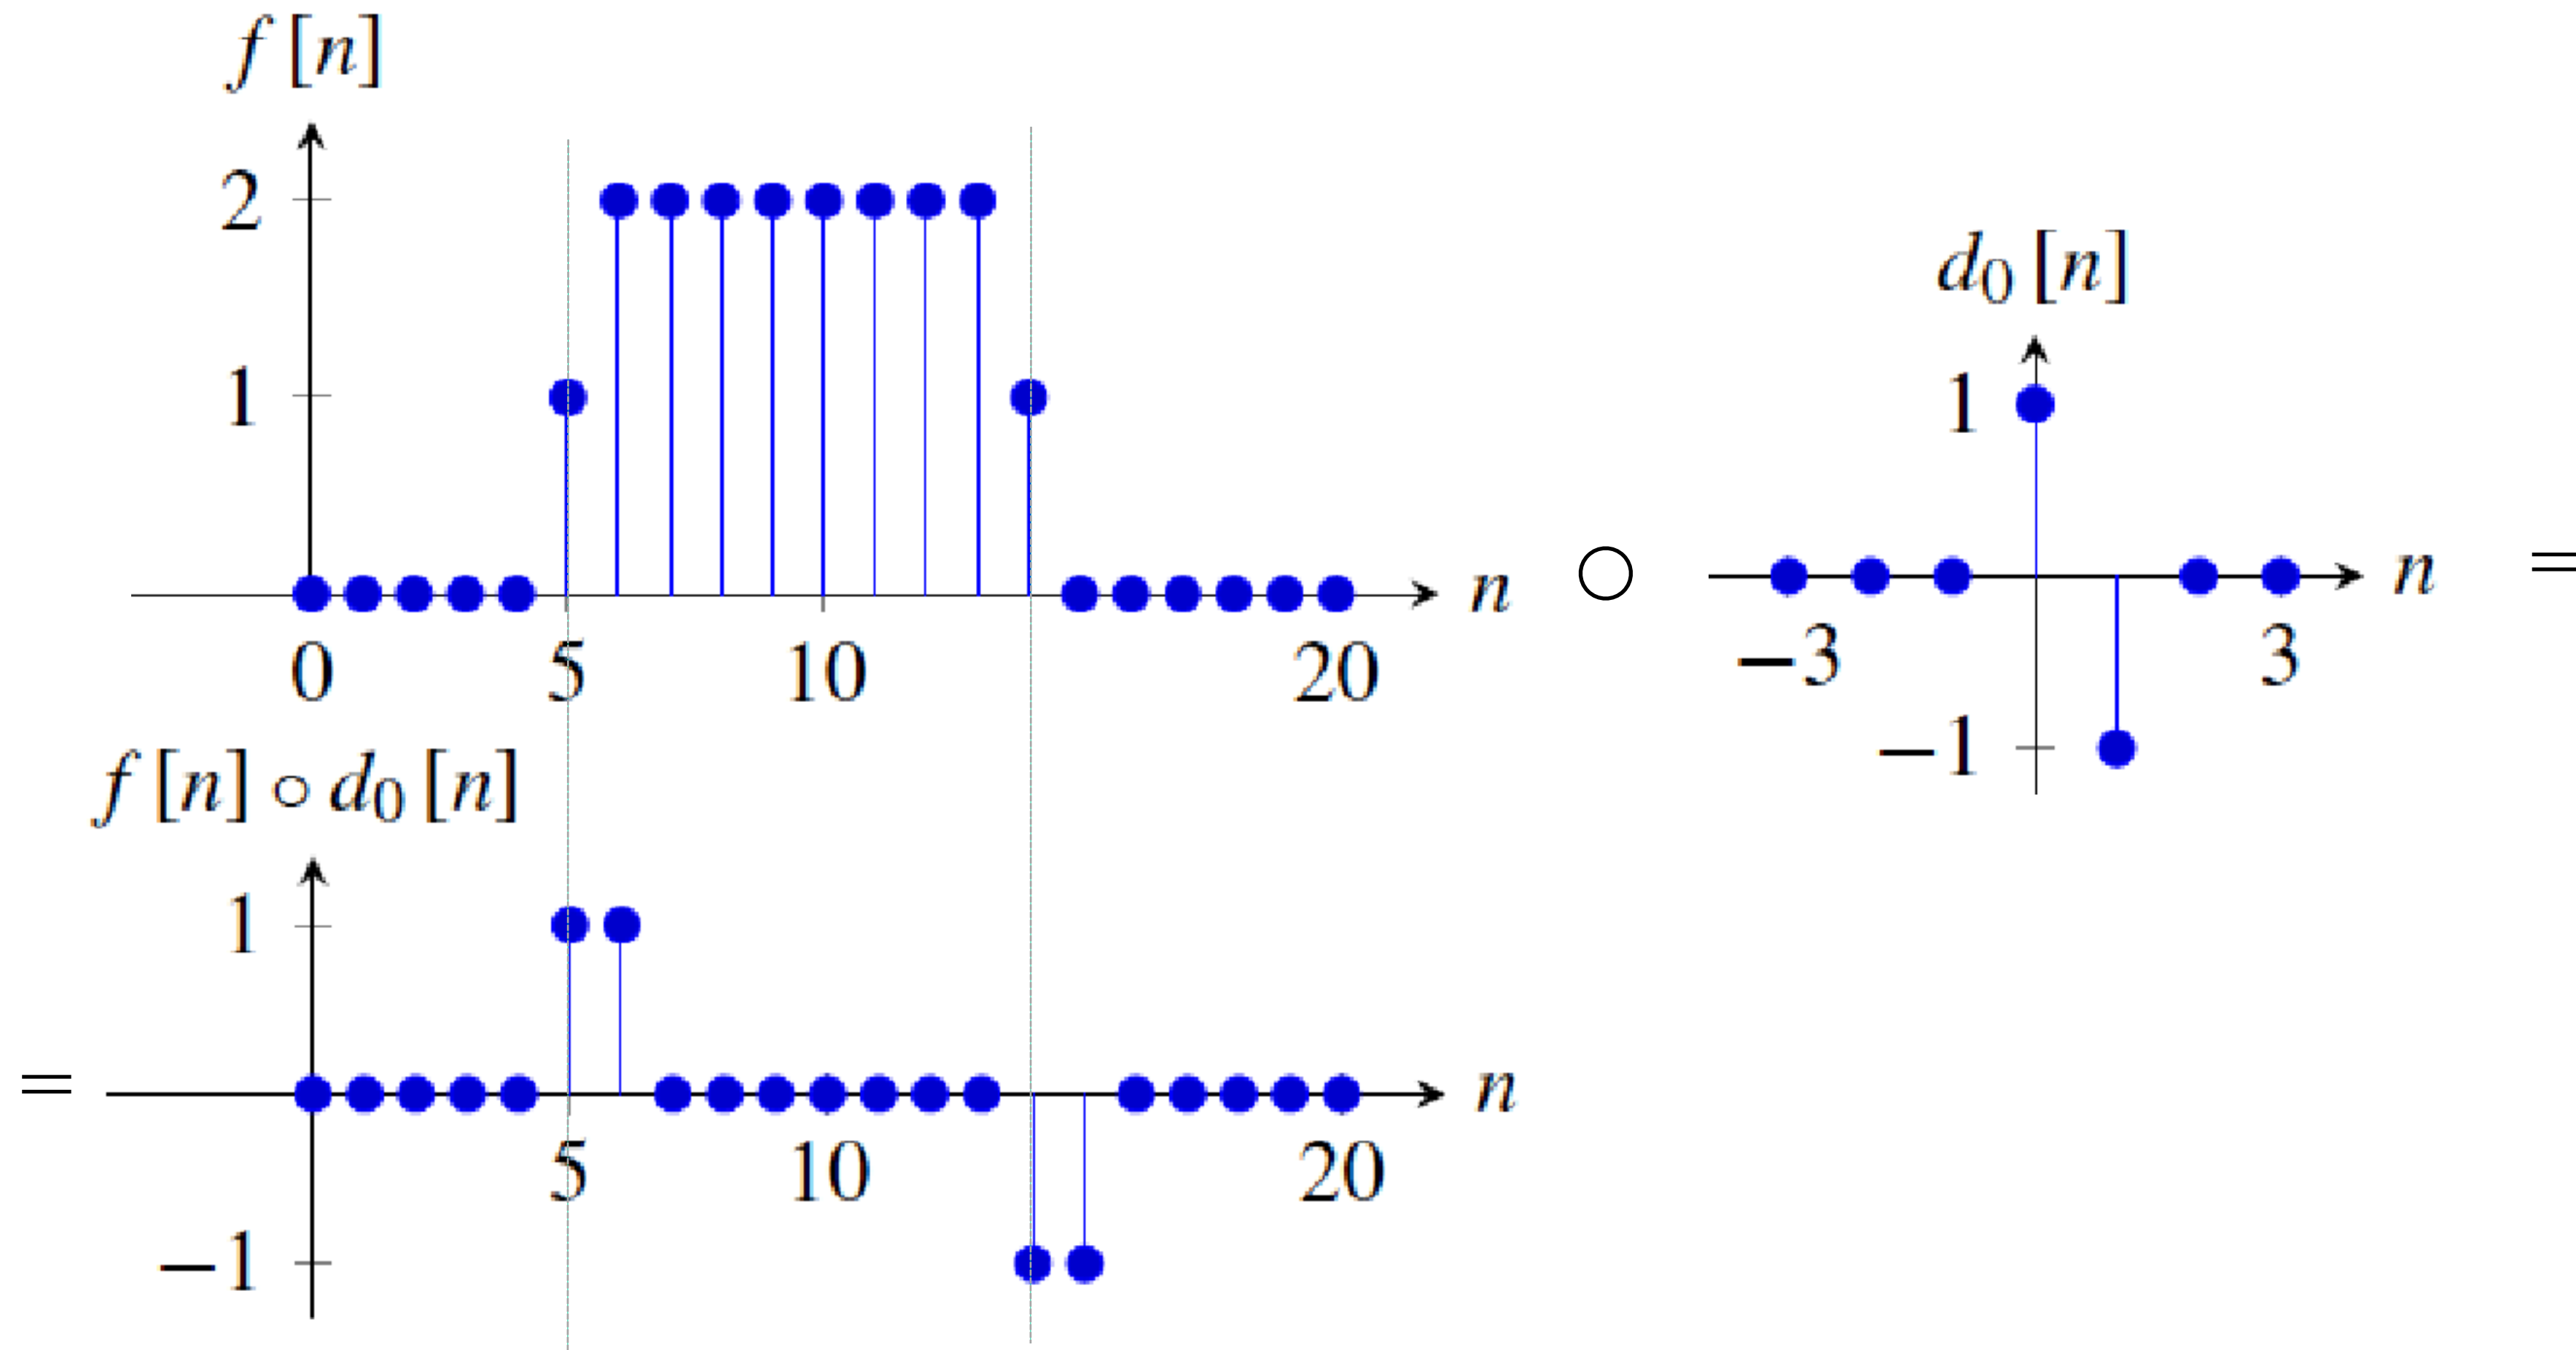
$$d_0 = [1, -1]$$

$$f \circ d_0 = f[n] - f[n - 1]$$

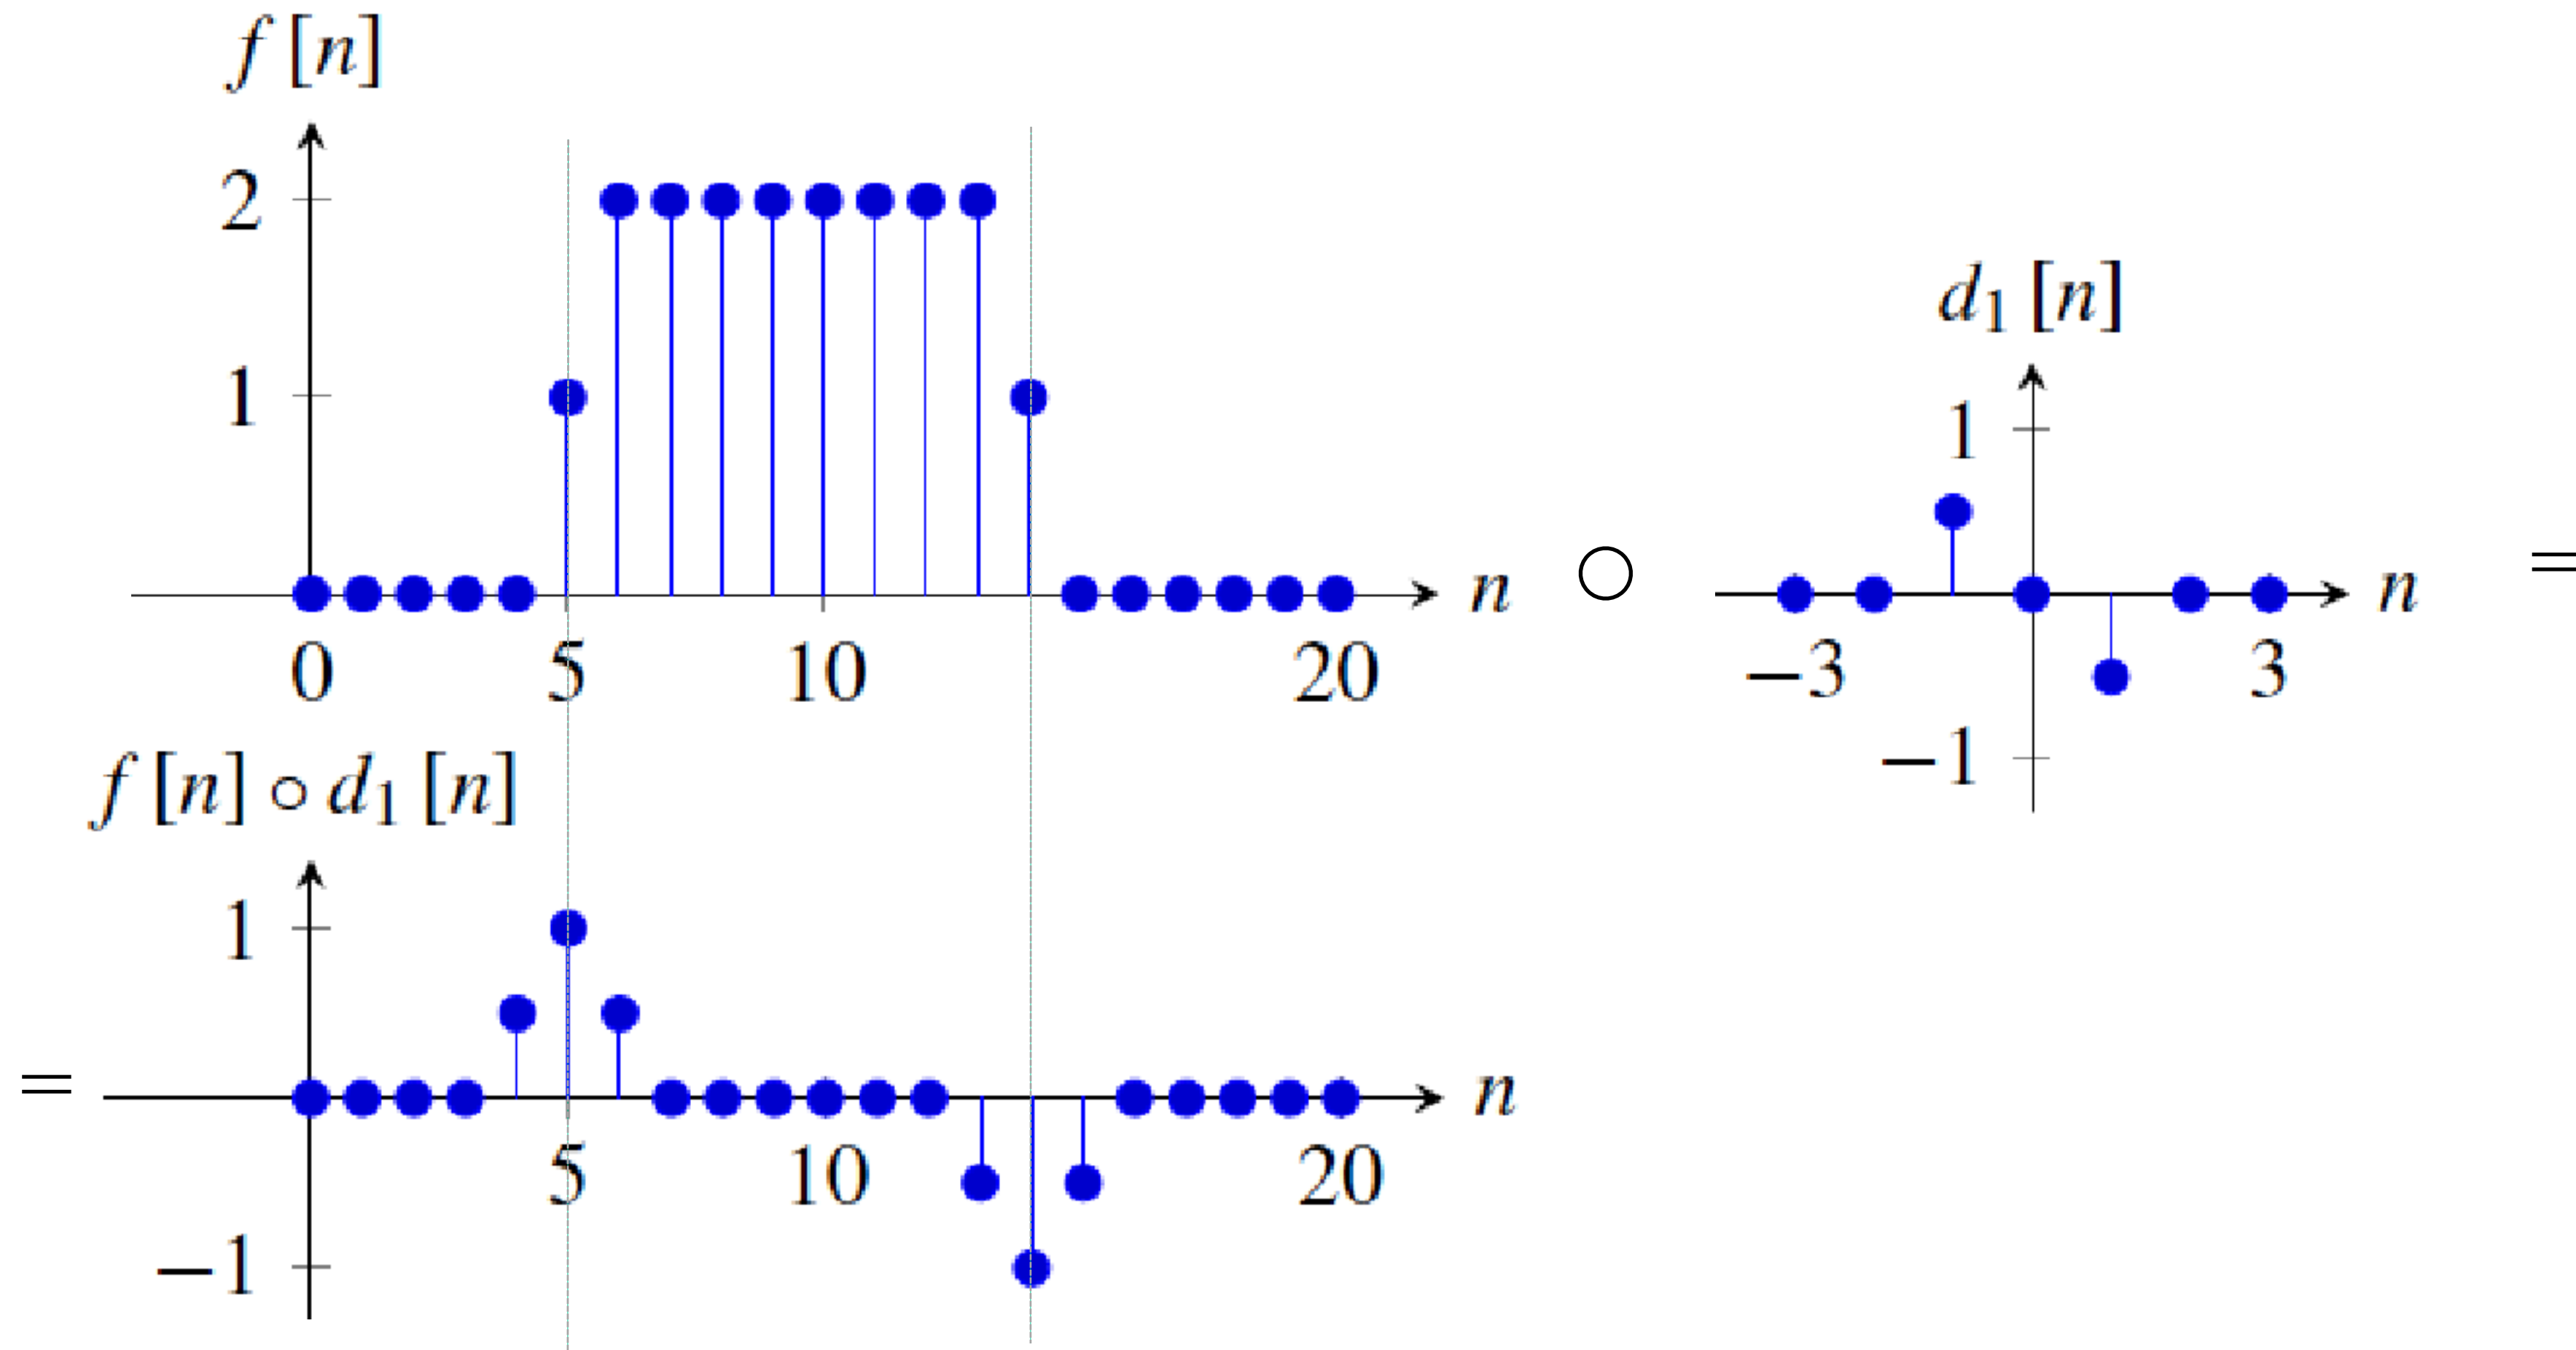
$$d_1 = [1, 0, -1] / 2$$

$$f \circ d_1 = \frac{f[n + 1] - f[n - 1]}{2}$$

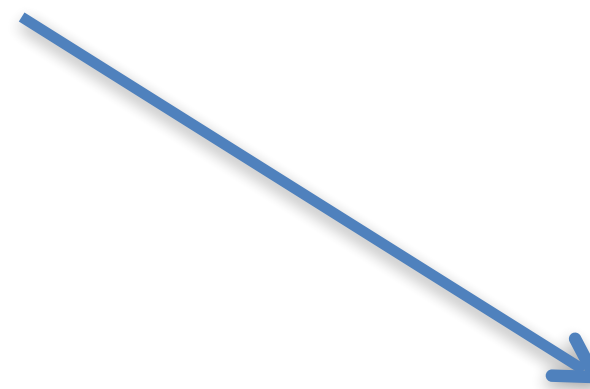
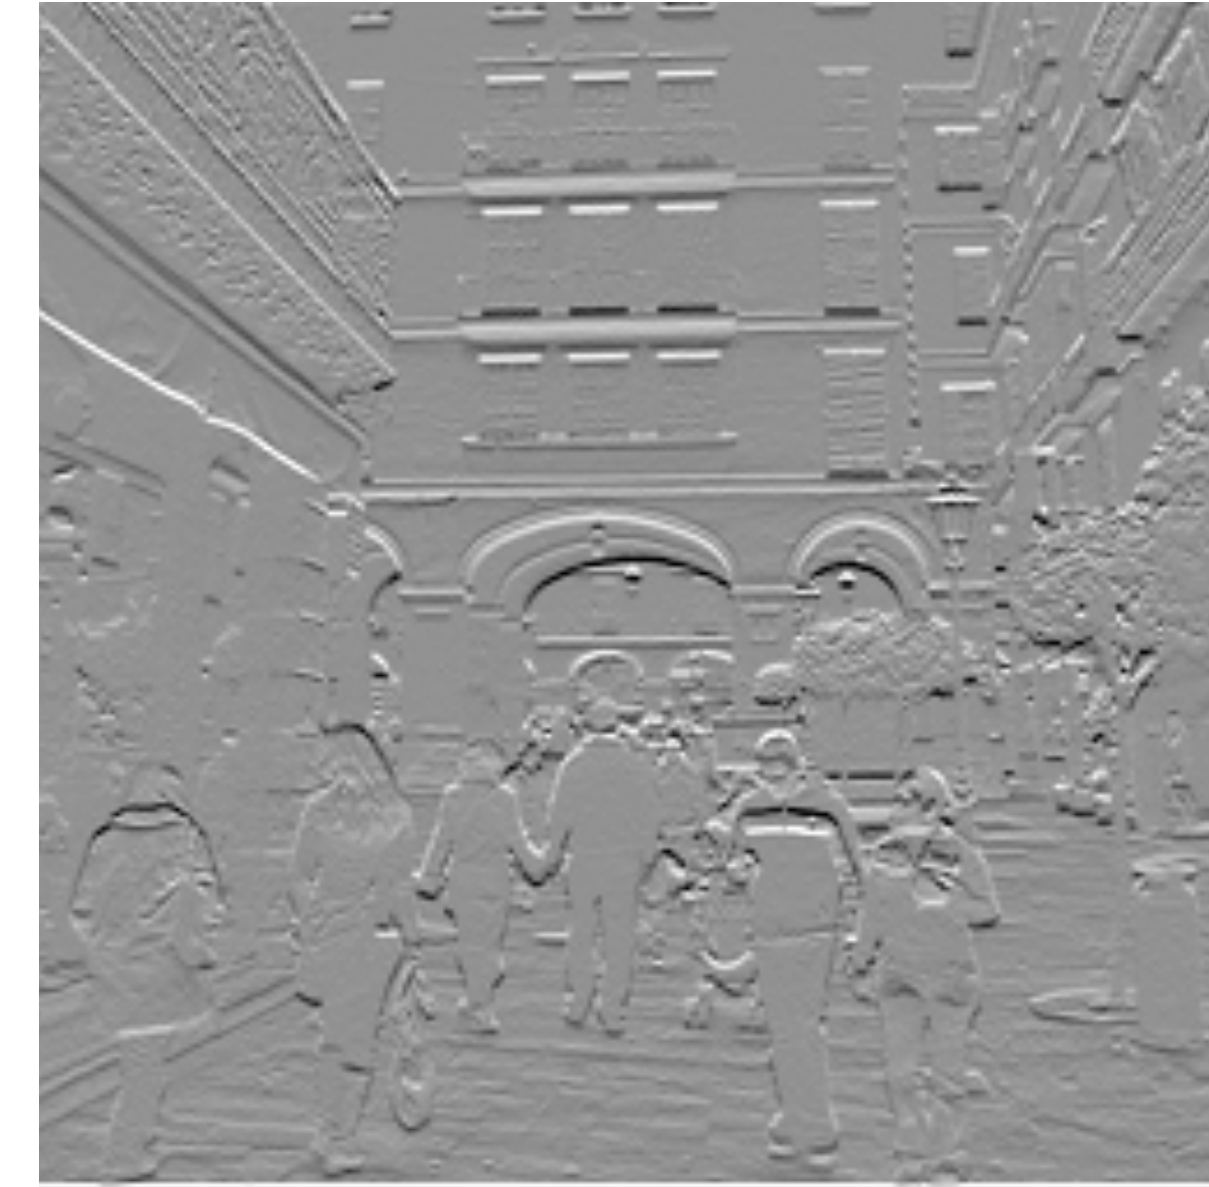
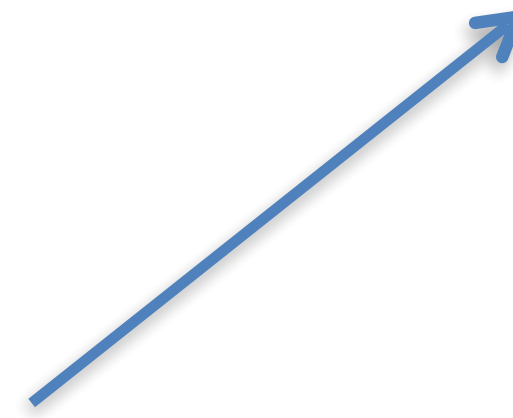
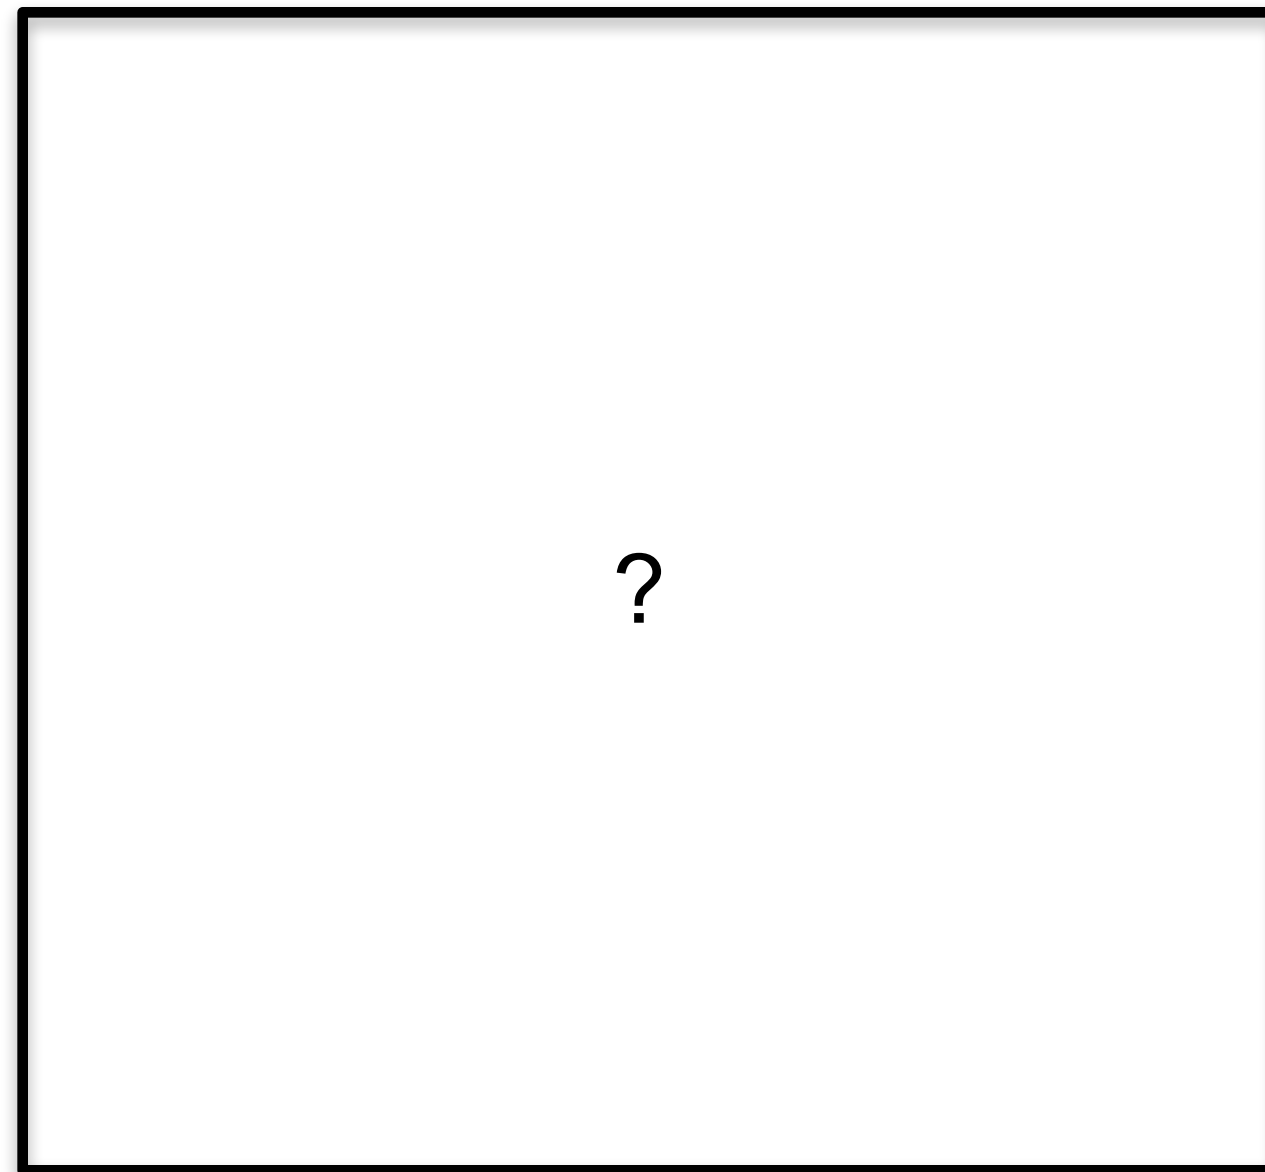
Discrete derivatives



Discrete derivatives



Can you go from the derivatives
back to the original image?



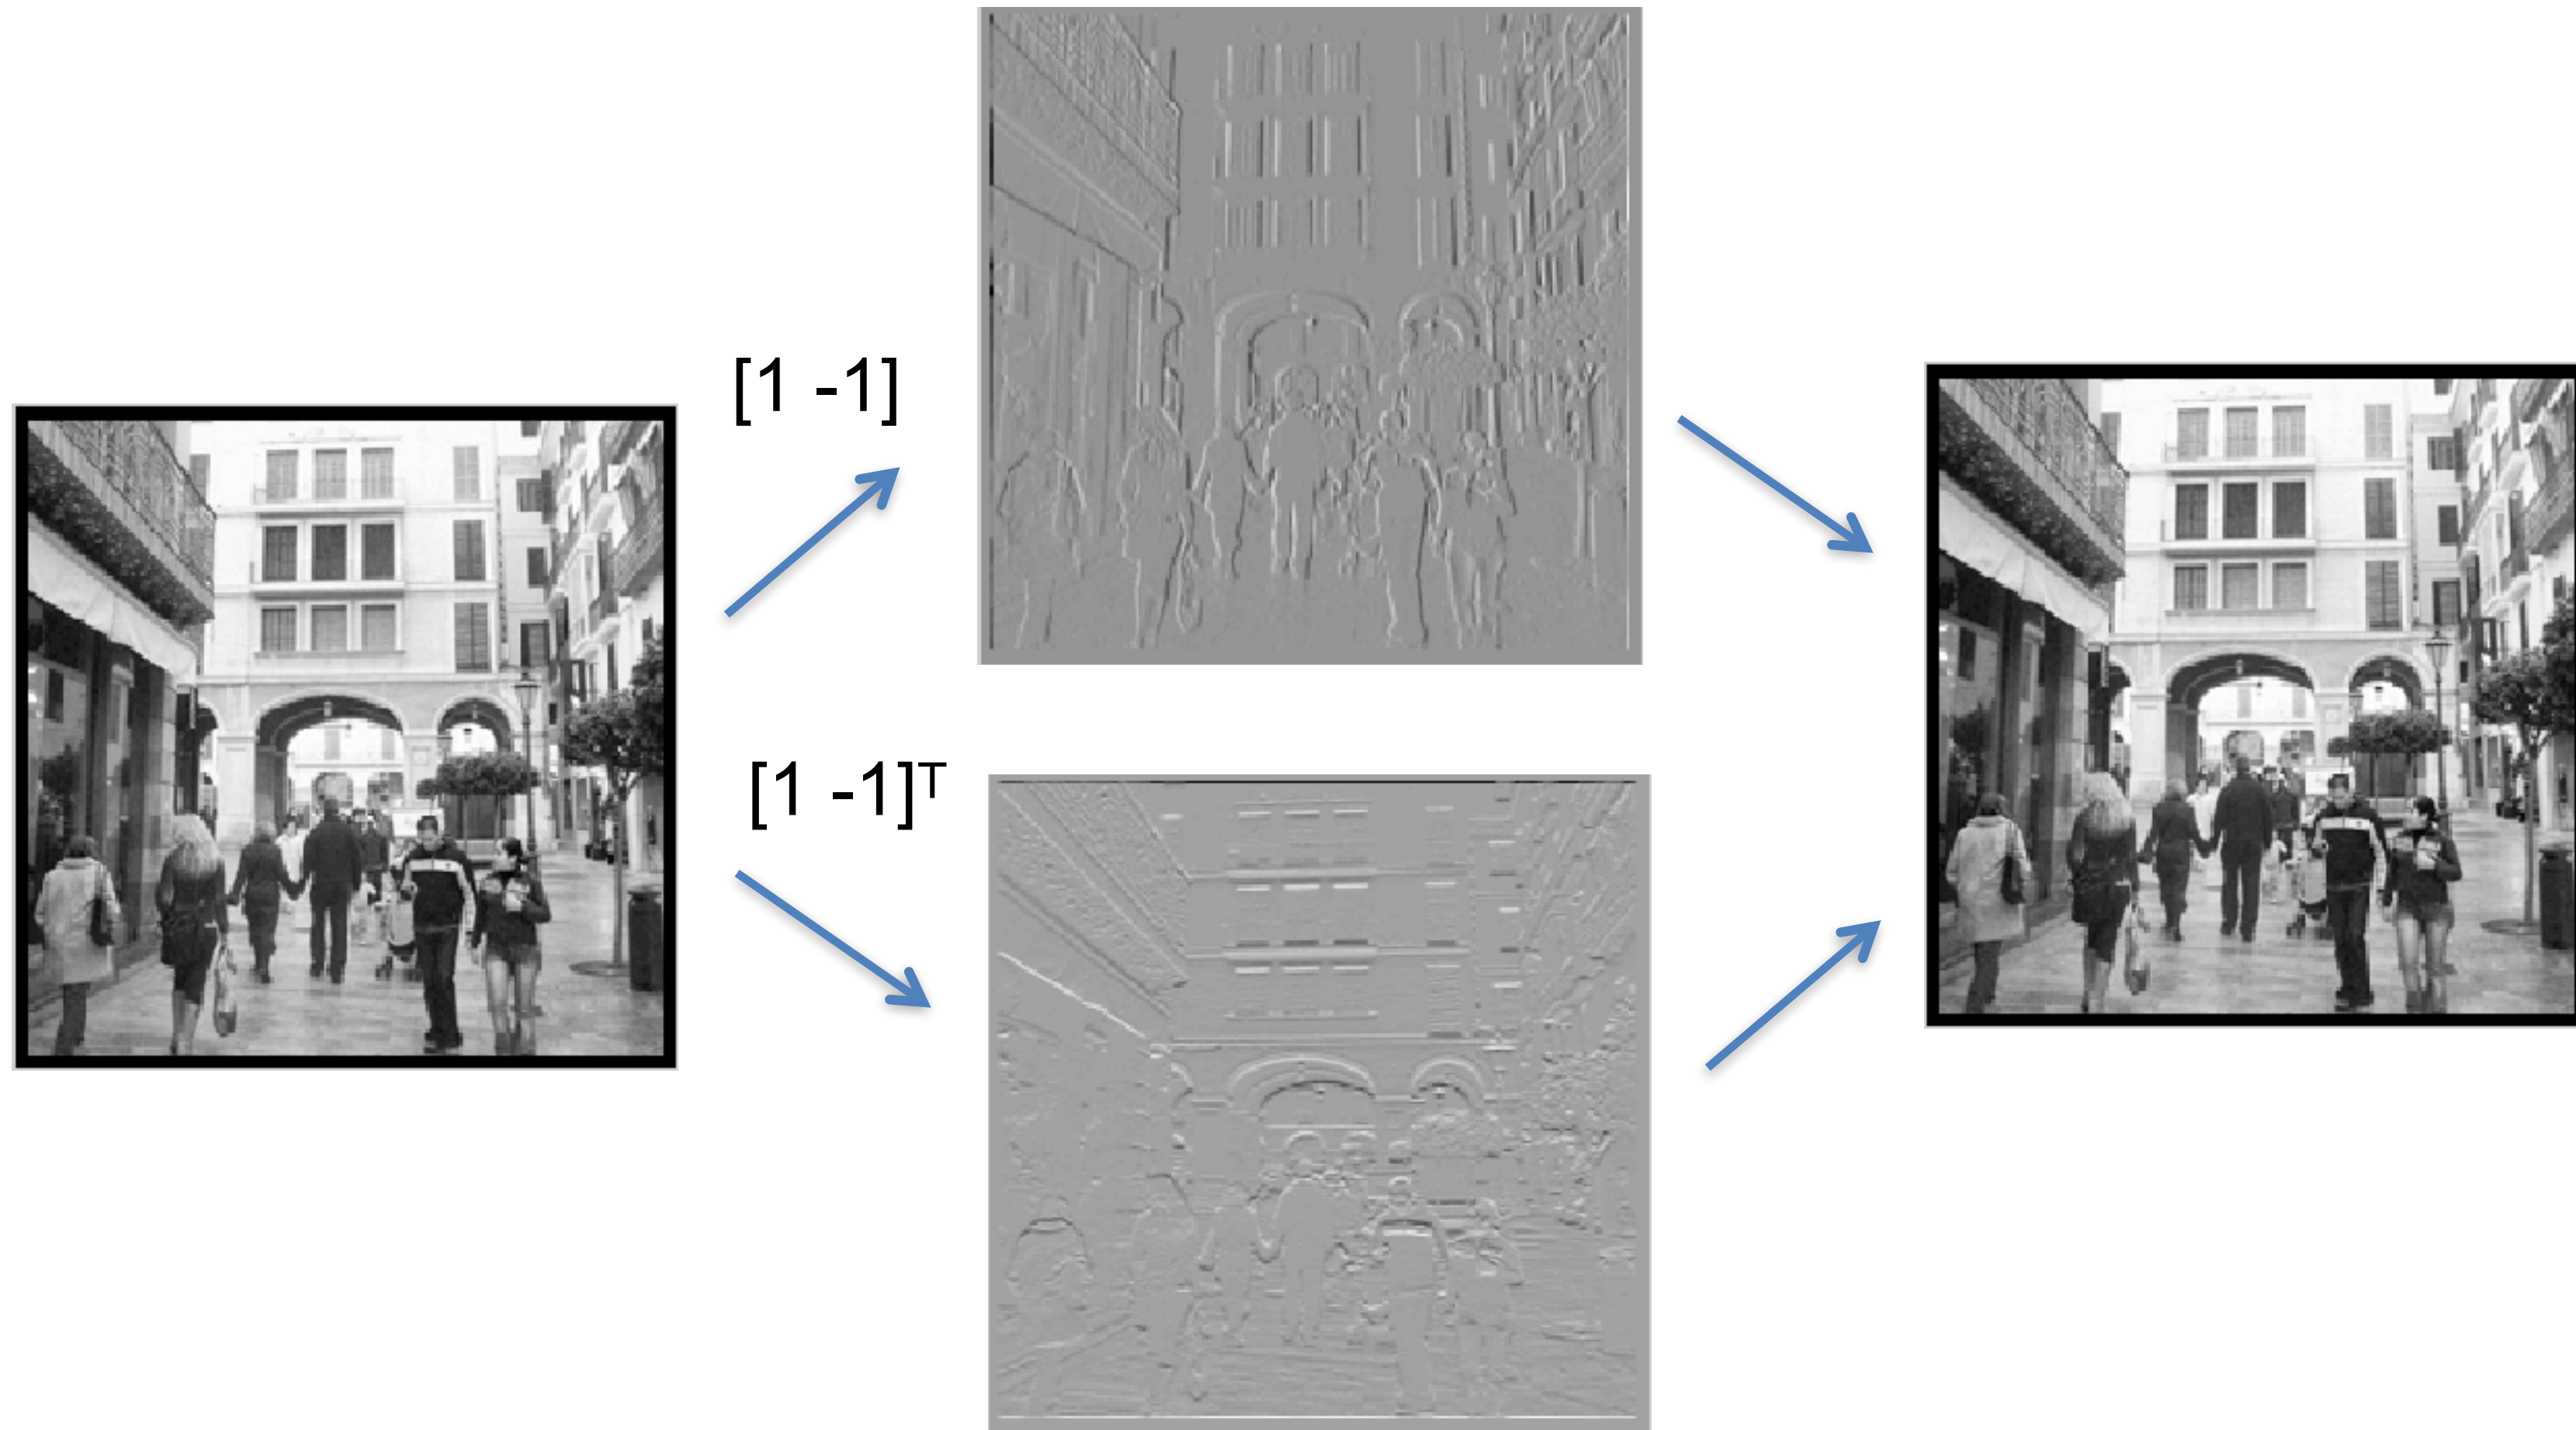
Reconstruction from 2D derivatives

In 2D, we have multiple derivatives (along n and m)

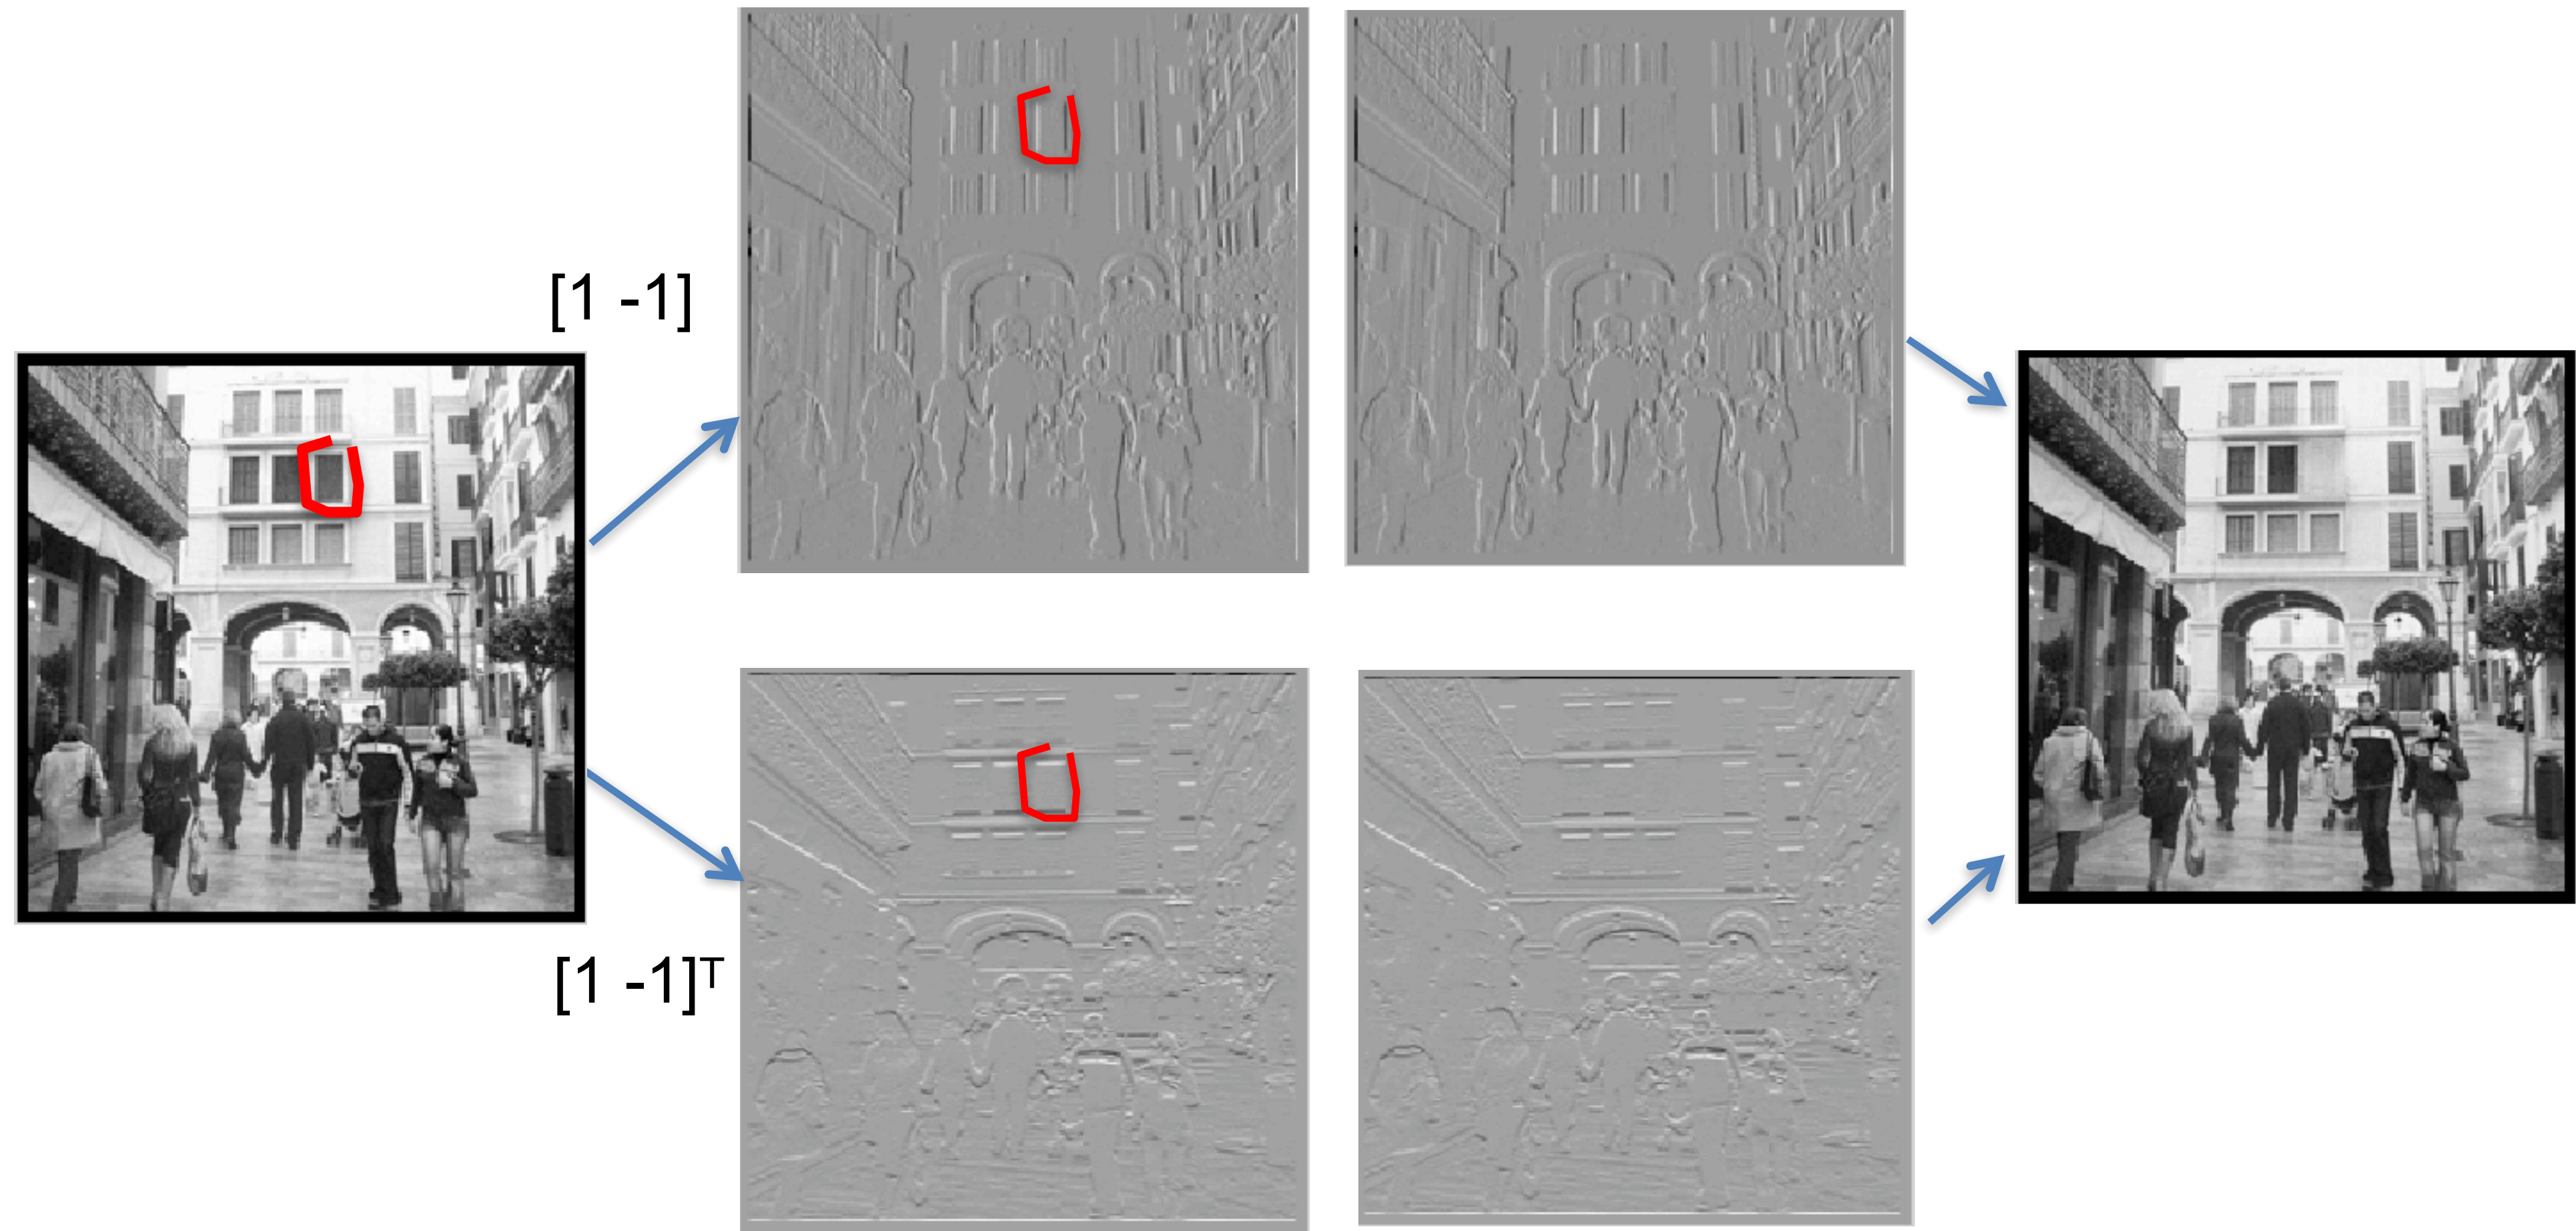
$$\begin{bmatrix} \text{red } c \\ \text{blue } c \end{bmatrix} = \begin{bmatrix} \text{red } [-1 \ 1] \\ \text{blue } [-1 \ 1]^T \end{bmatrix} \begin{bmatrix} \text{blue } c \end{bmatrix}$$

and we compute the pseudo-inverse of the full matrix.

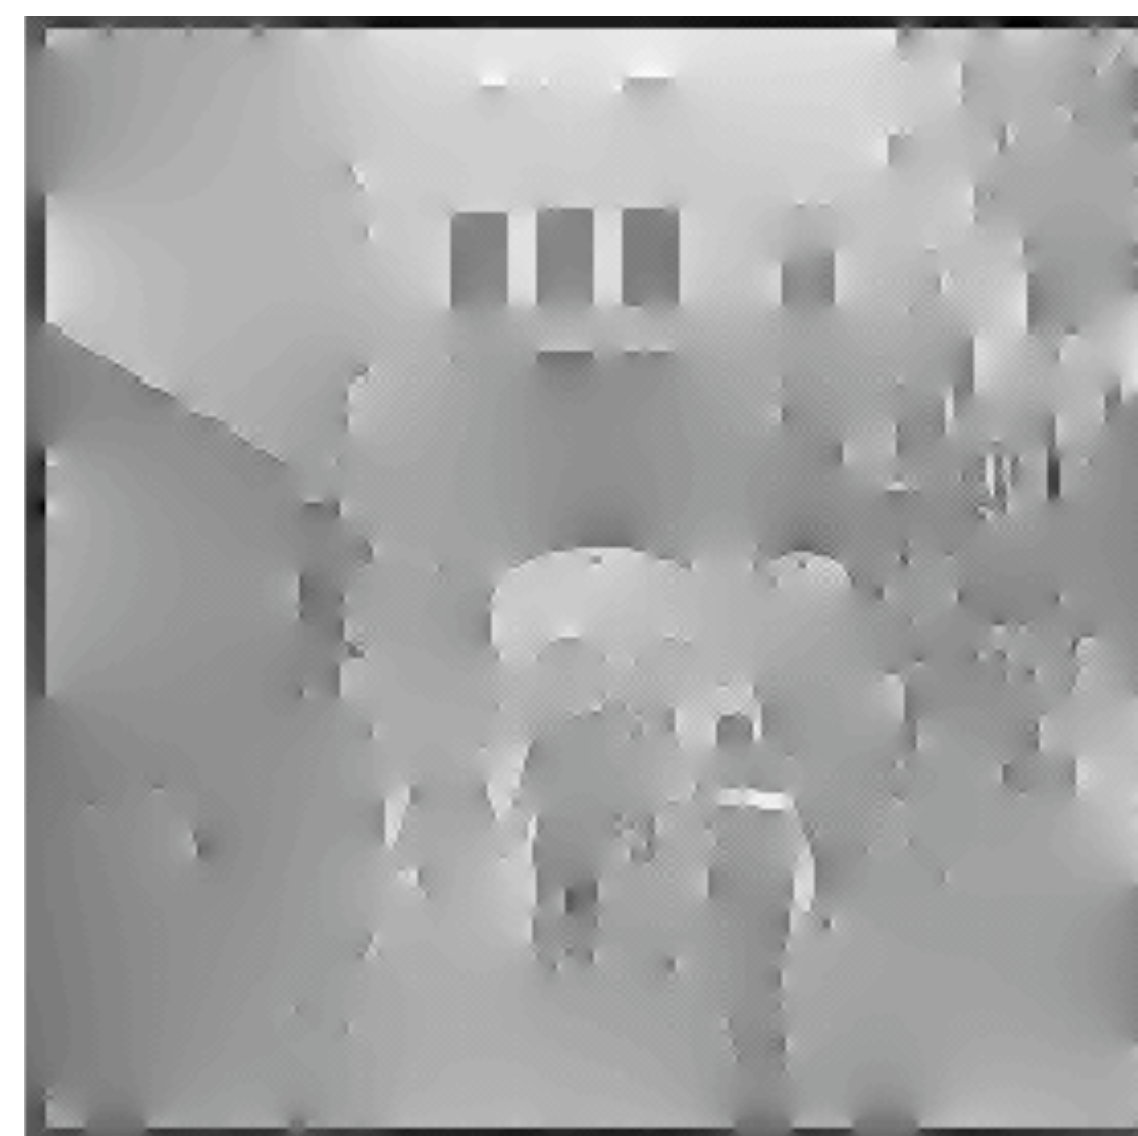
Reconstruction from 2D derivatives



Editing the edge image

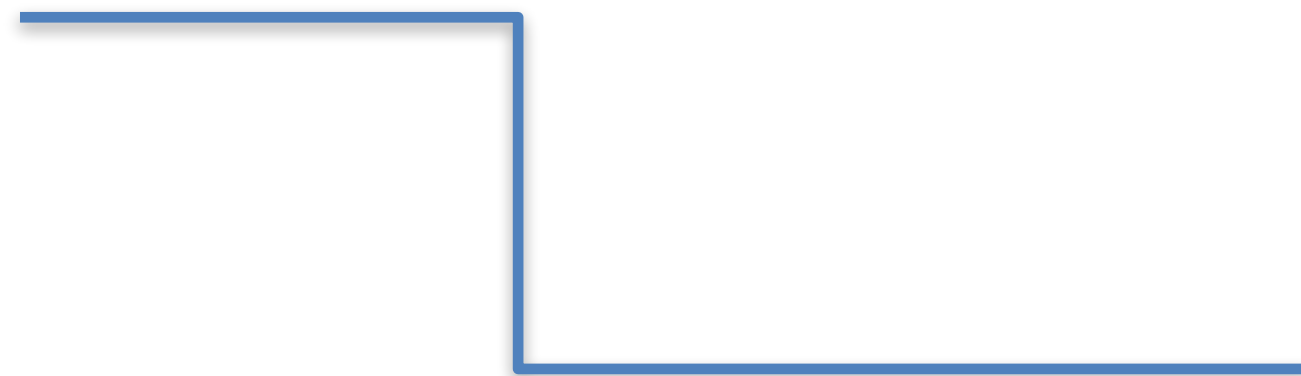


Thresholding edges



Issues with image derivatives

- Derivatives are sensitive to noise
- If we consider continuous image derivatives, they might not be defined in some regions (e.g., object boundaries, ...)



Derivatives

We want to compute the image derivative:

$$\frac{\partial f(x, y)}{\partial x}$$

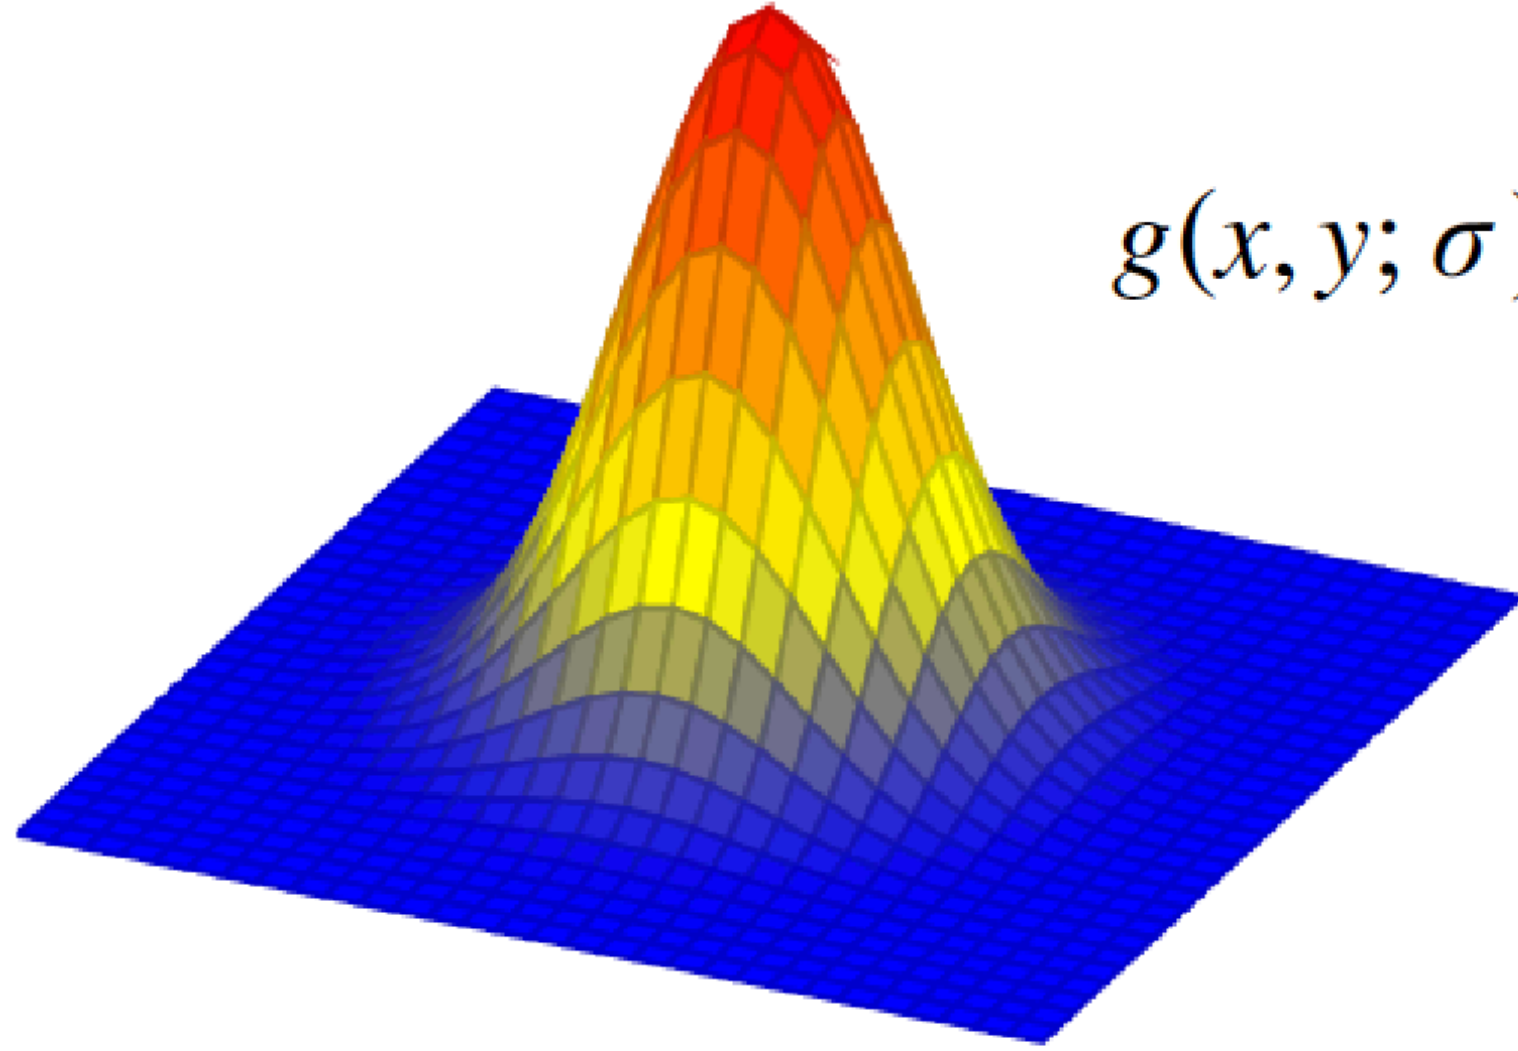
If there is noise, we might want to “smooth” it with a blurring filter

$$\frac{\partial f(x, y)}{\partial x} \circ g(x, y)$$

But derivatives and convolutions are linear and we can move them around:

$$\frac{\partial f(x, y)}{\partial x} \circ g(x, y) = f(x, y) \circ \frac{\partial g(x, y)}{\partial x}$$

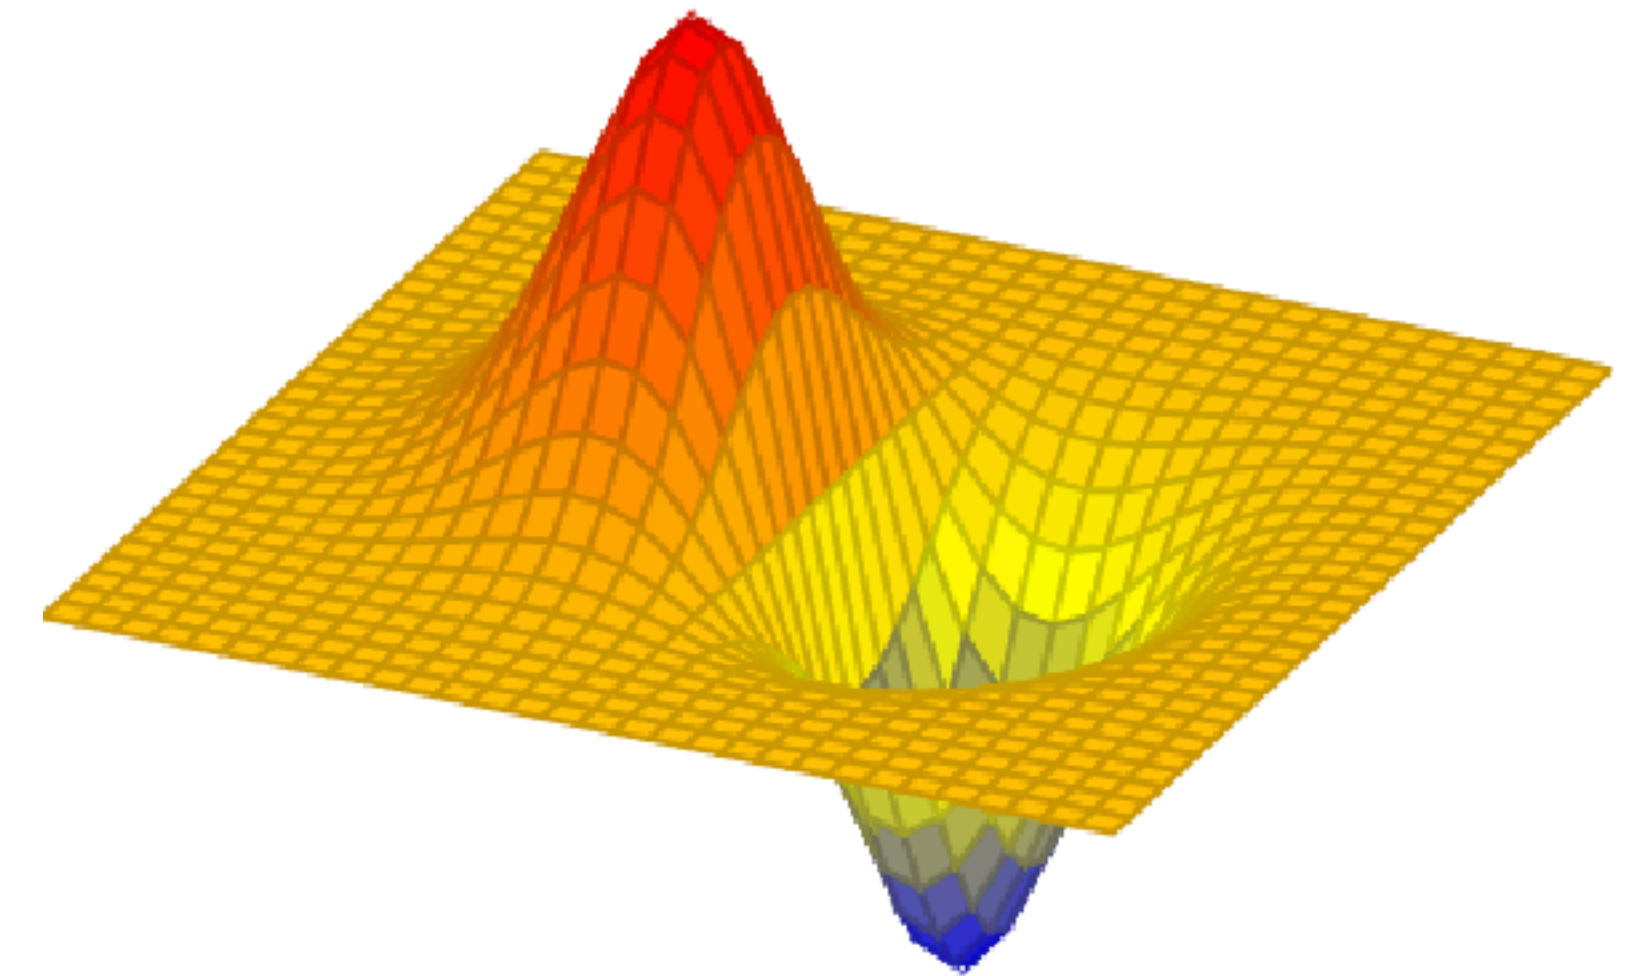
Gaussian derivatives



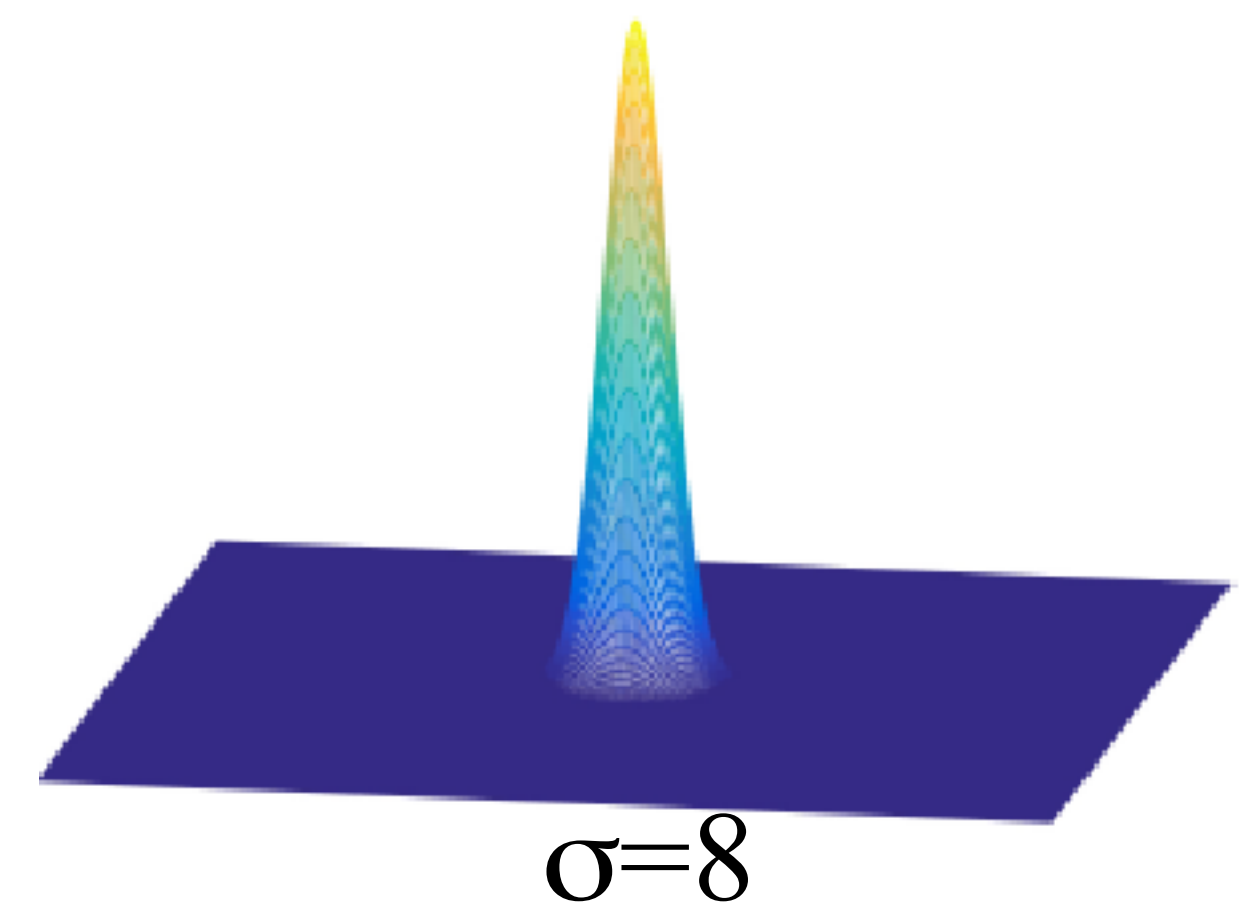
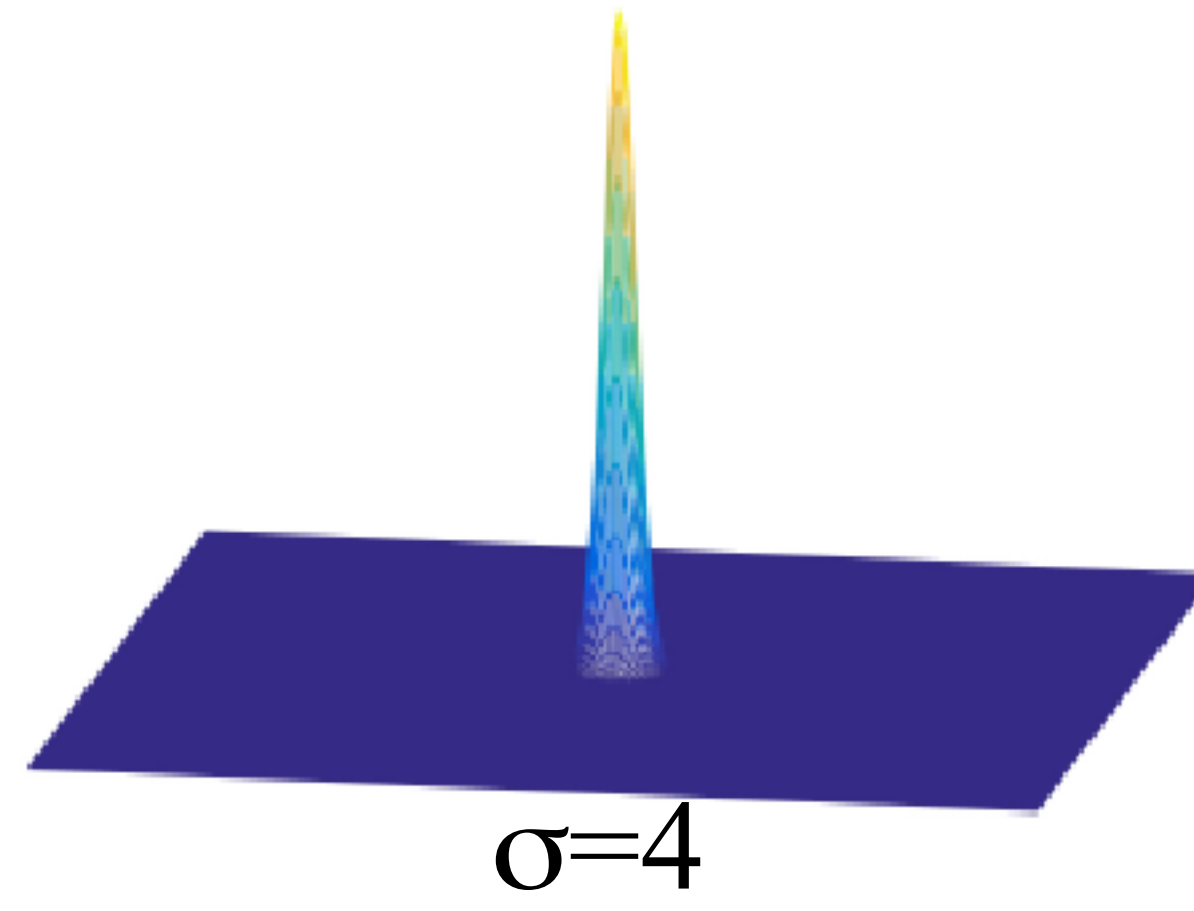
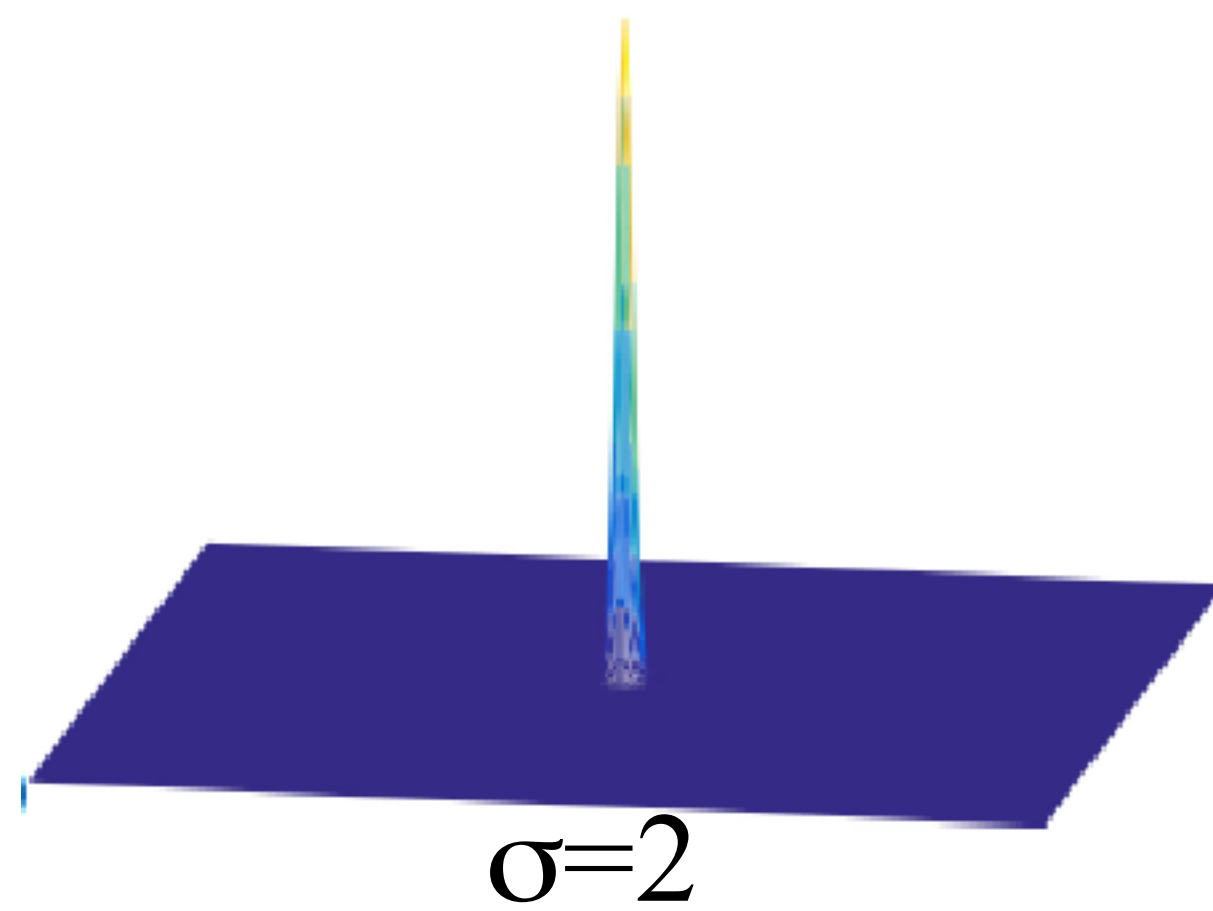
$$g(x, y; \sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

The continuous derivative is:

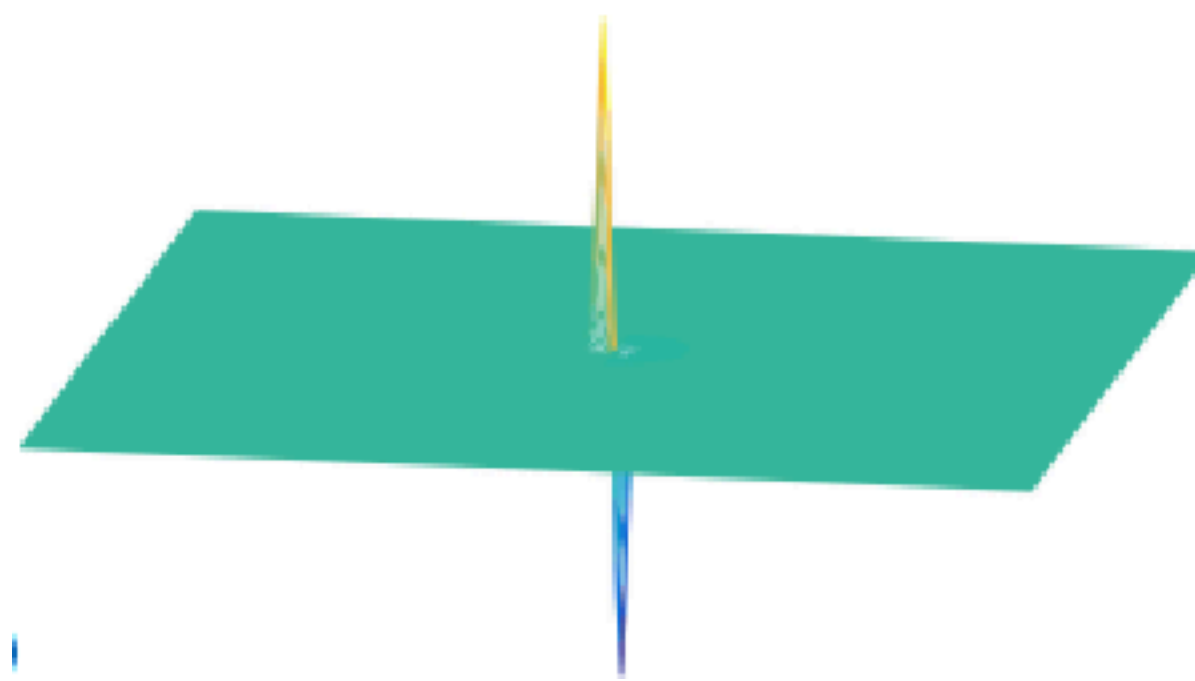
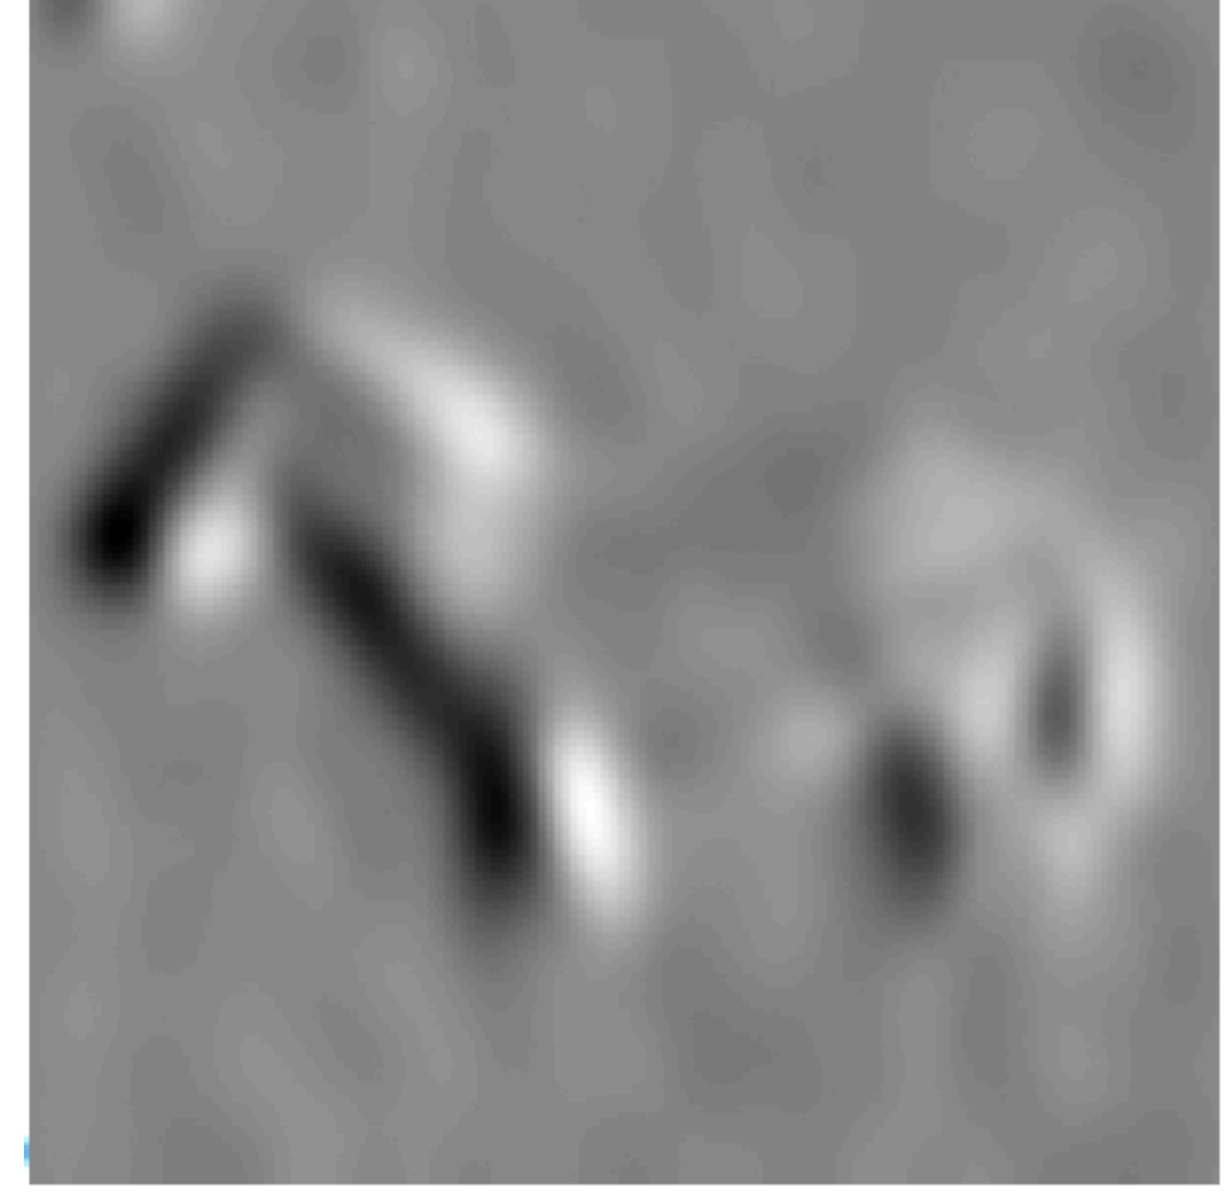
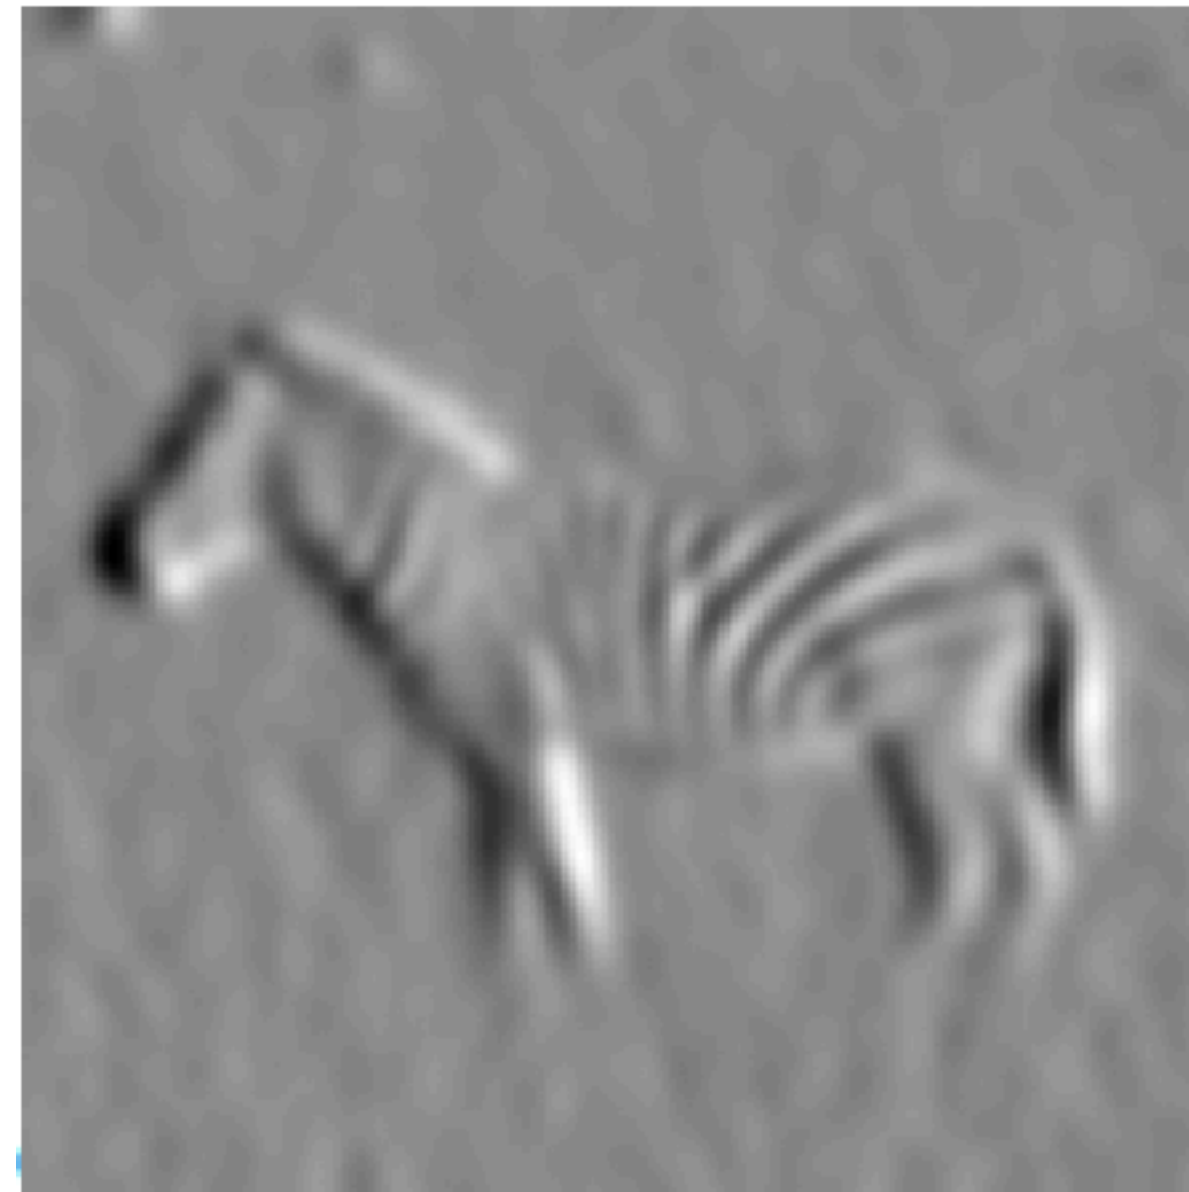
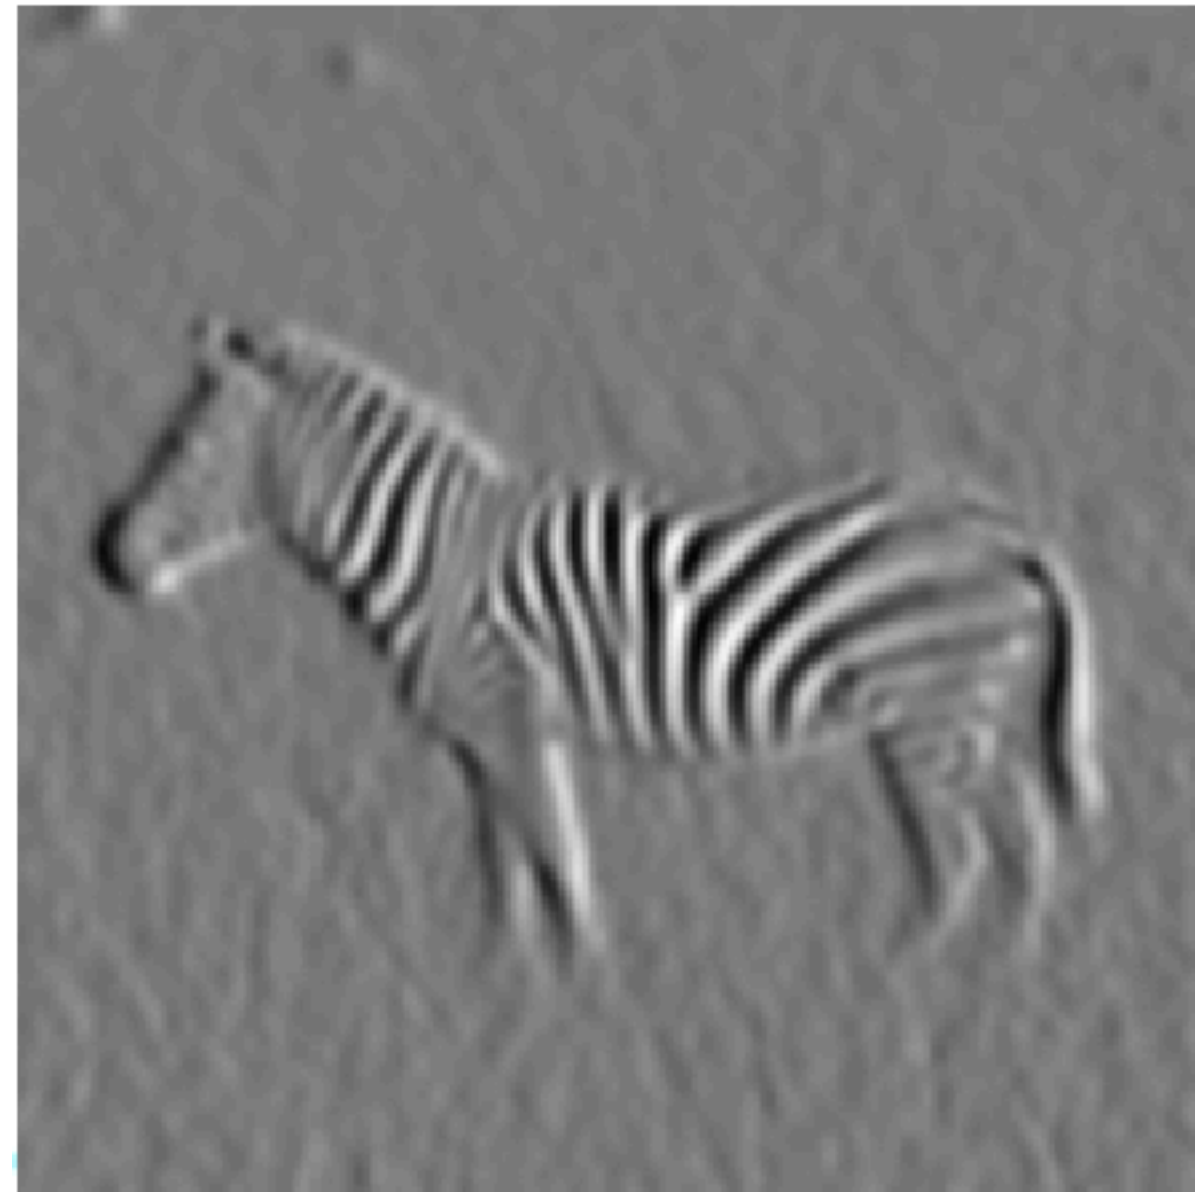
$$\begin{aligned} g_x(x, y; \sigma) &= \frac{\partial g(x, y; \sigma)}{\partial x} = \\ &= \frac{-x}{2\pi\sigma^4} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ &= \frac{-x}{\sigma^2} g(x, y; \sigma) \end{aligned}$$



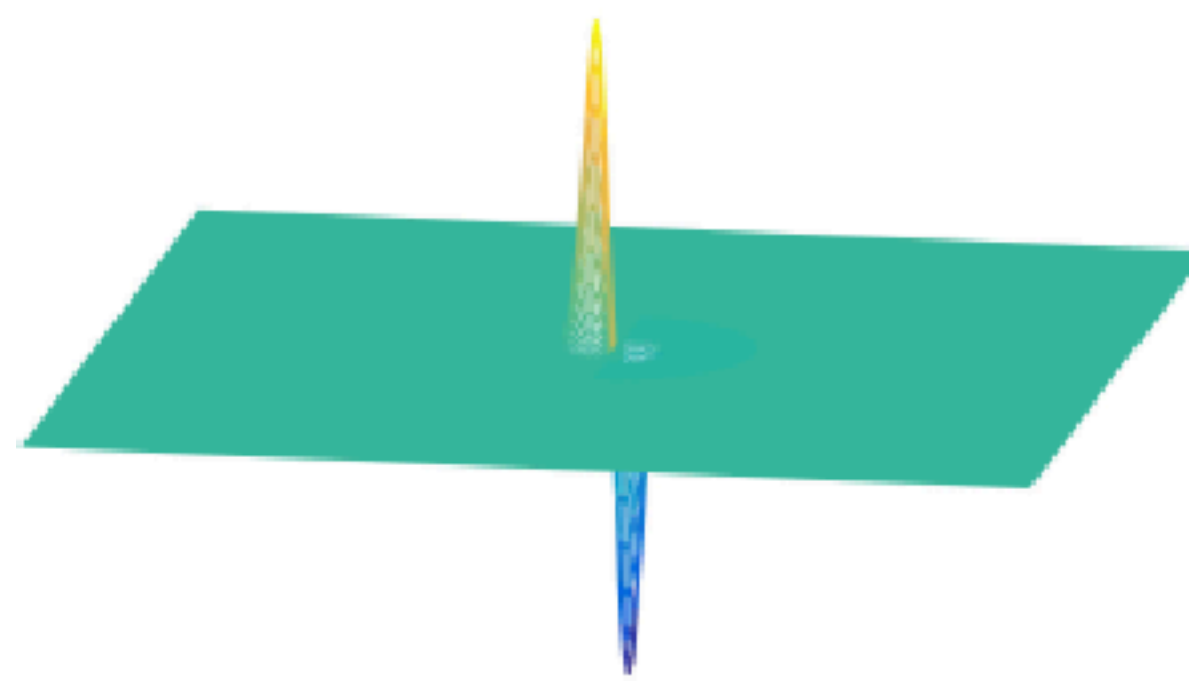
Gaussian Scale



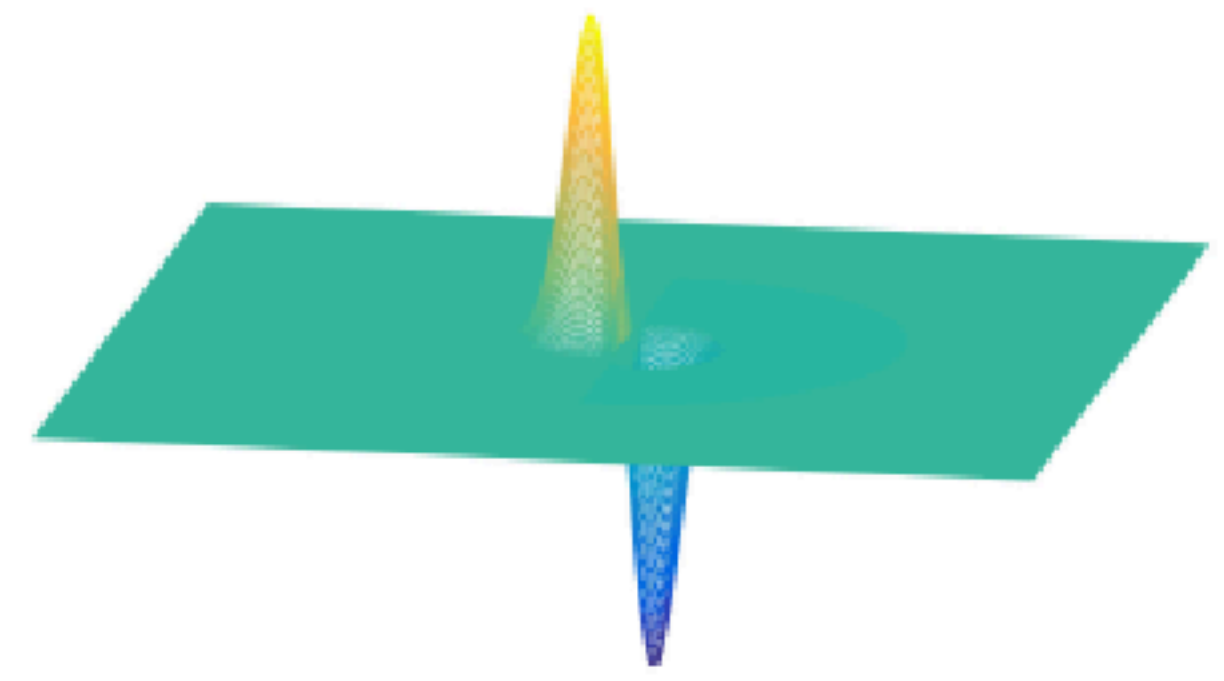
Derivatives of Gaussians: Scale



$\sigma=2$



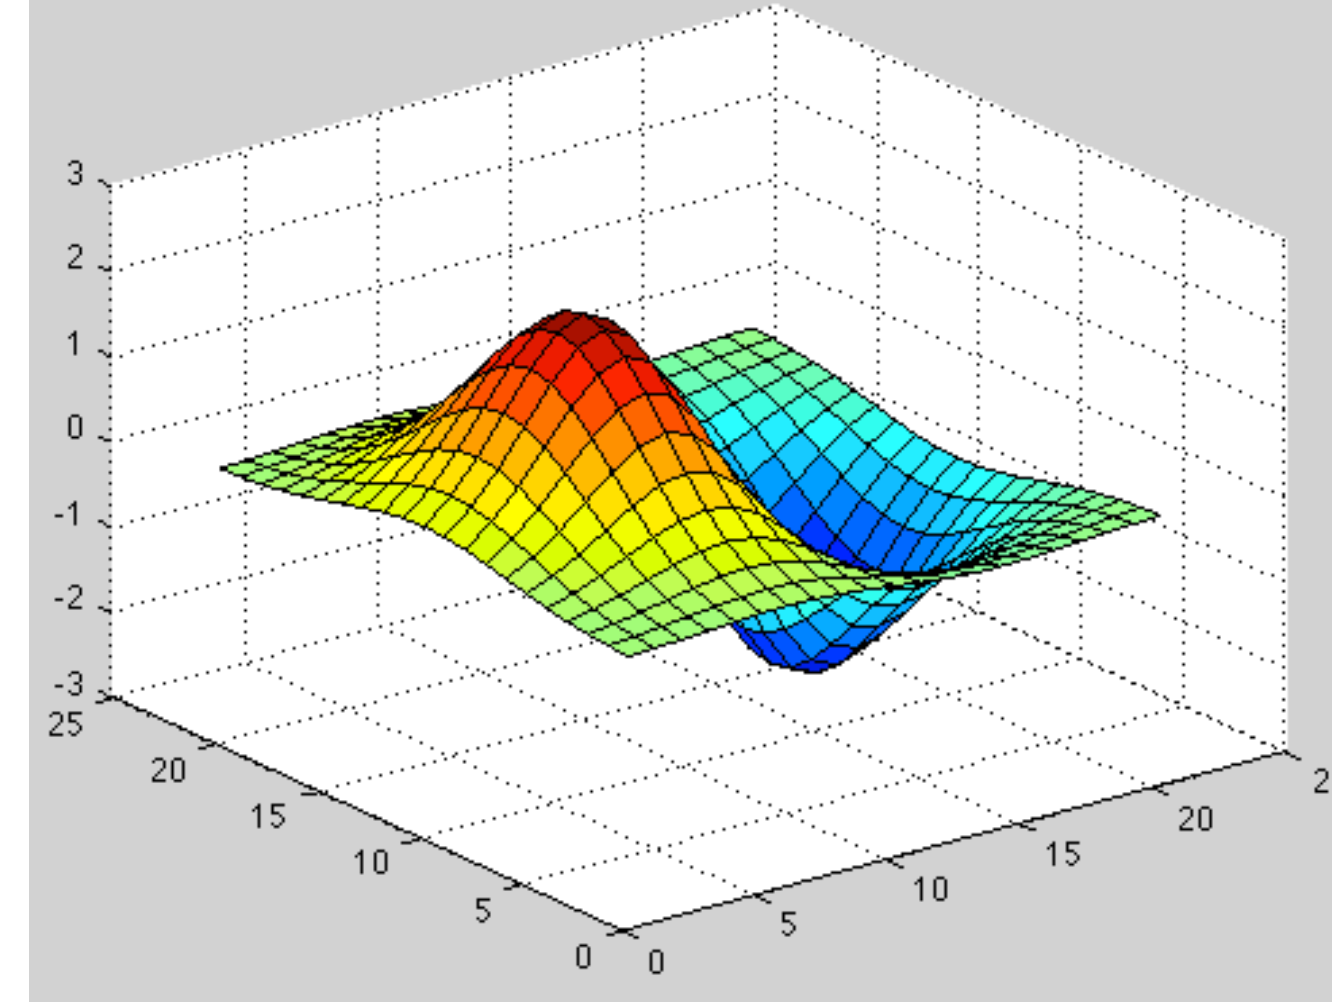
$\sigma=4$



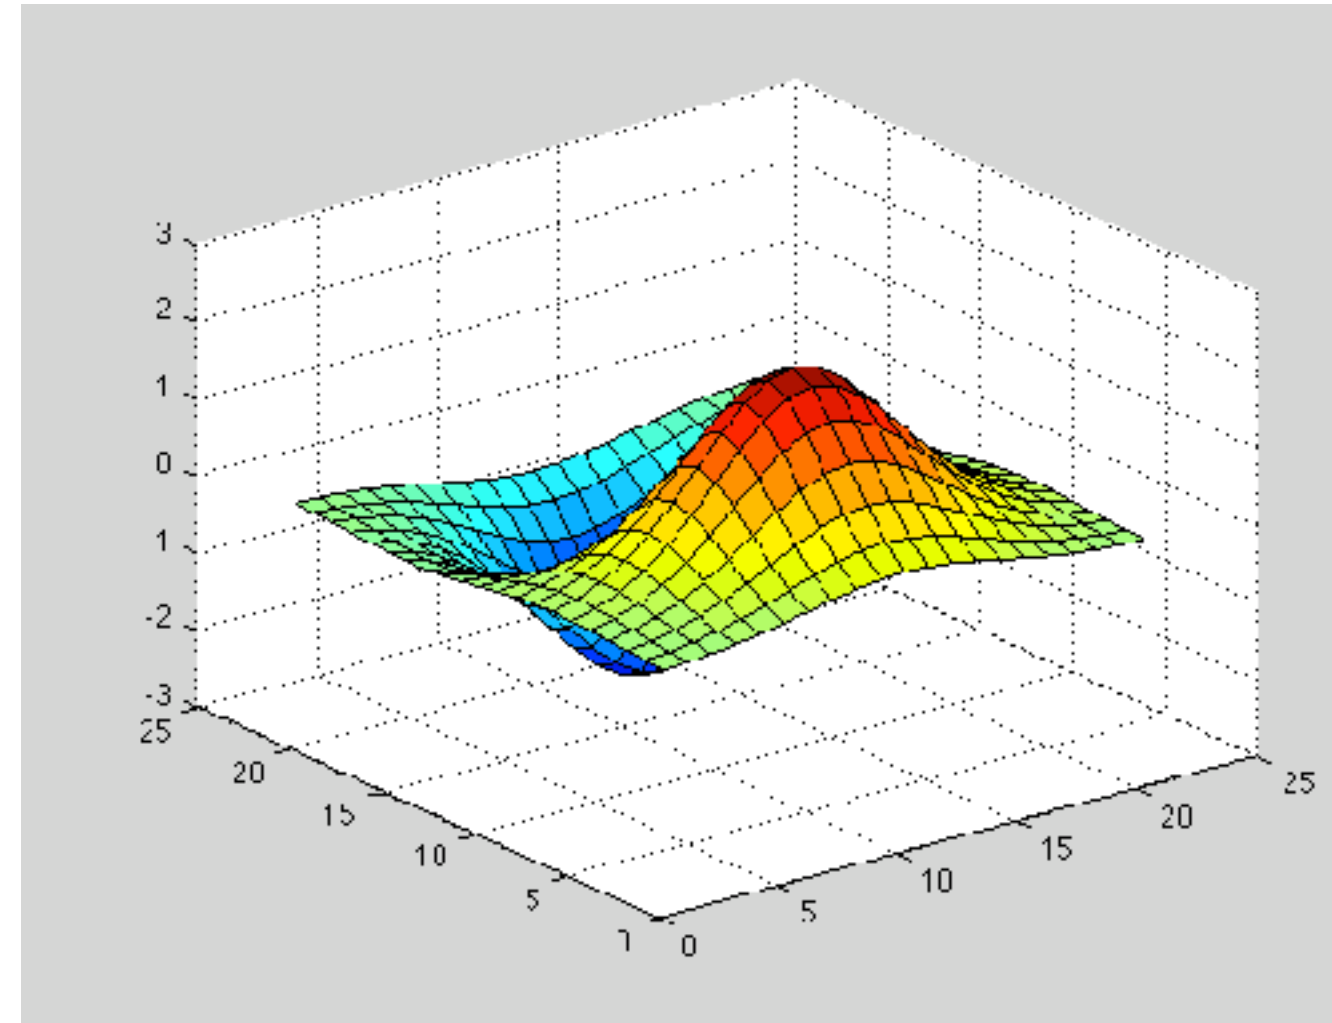
$\sigma=8$

Orientation

$$g_x(x,y) = \frac{\partial g(x,y)}{\partial x} = \frac{-x}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

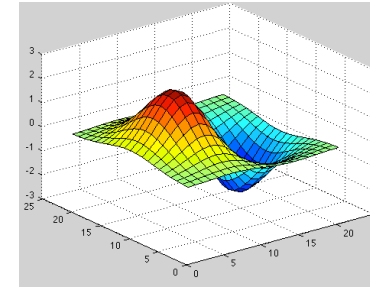


$$g_y(x,y) = \frac{\partial g(x,y)}{\partial y} = \frac{-y}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

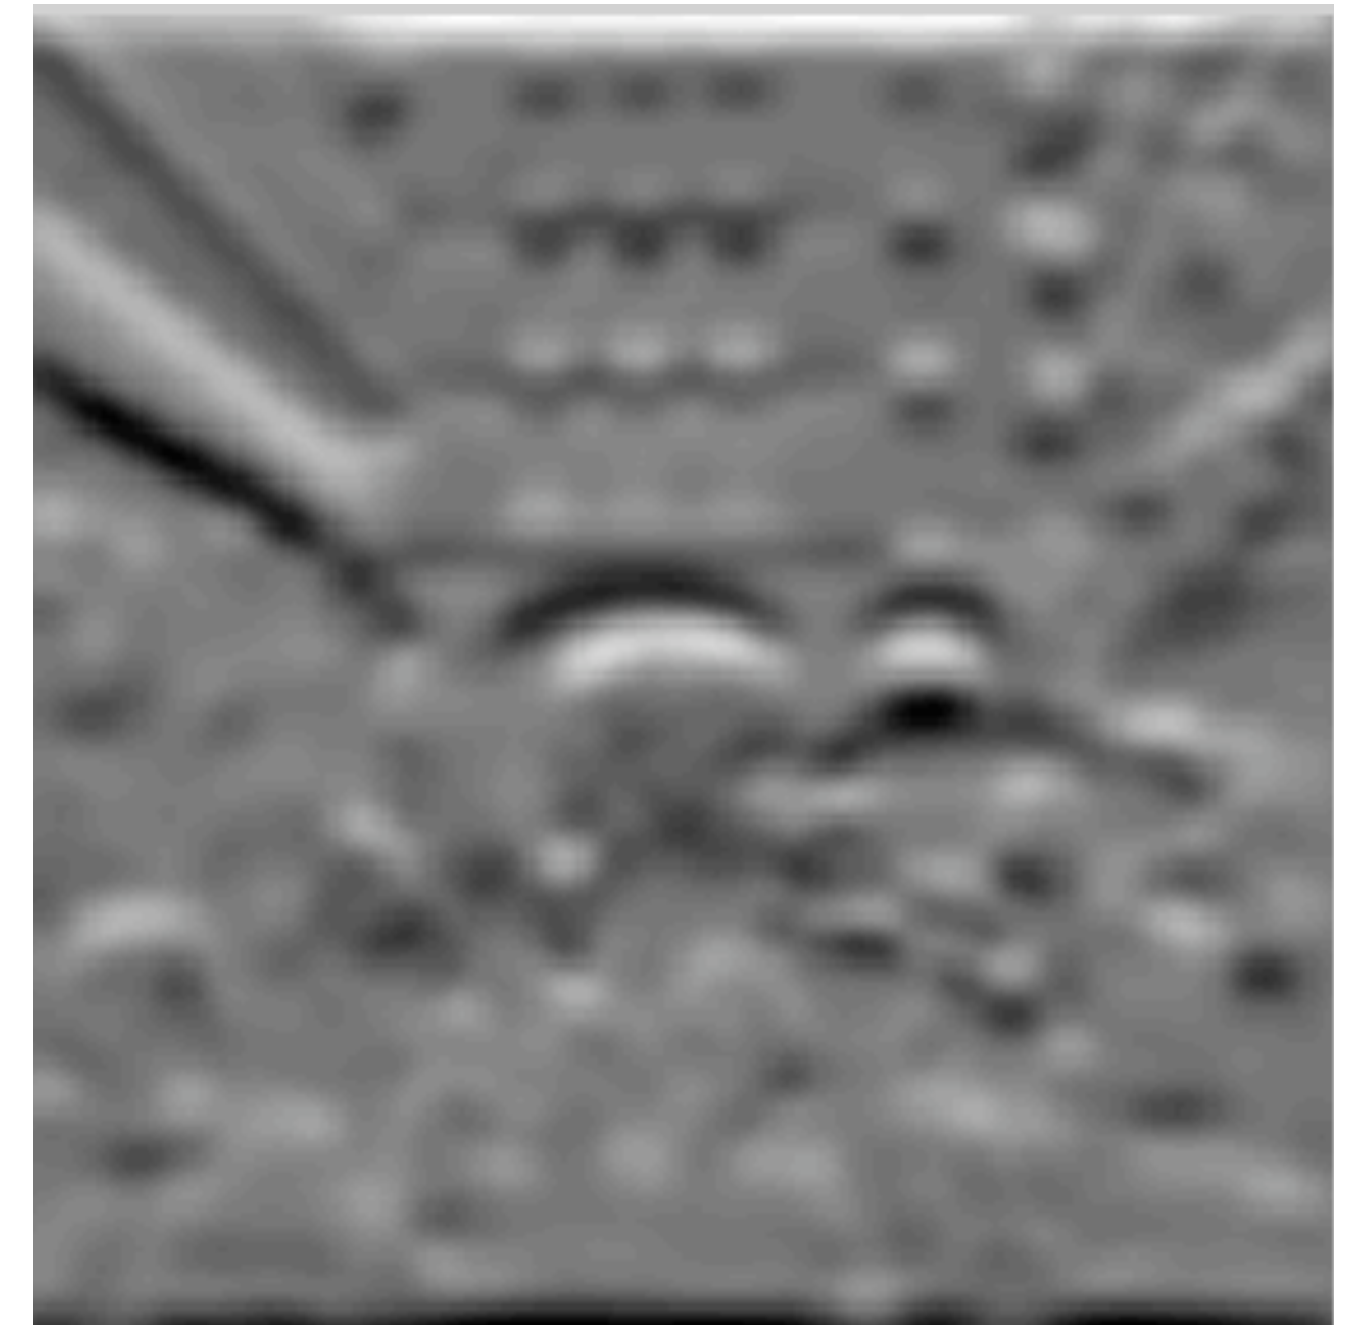
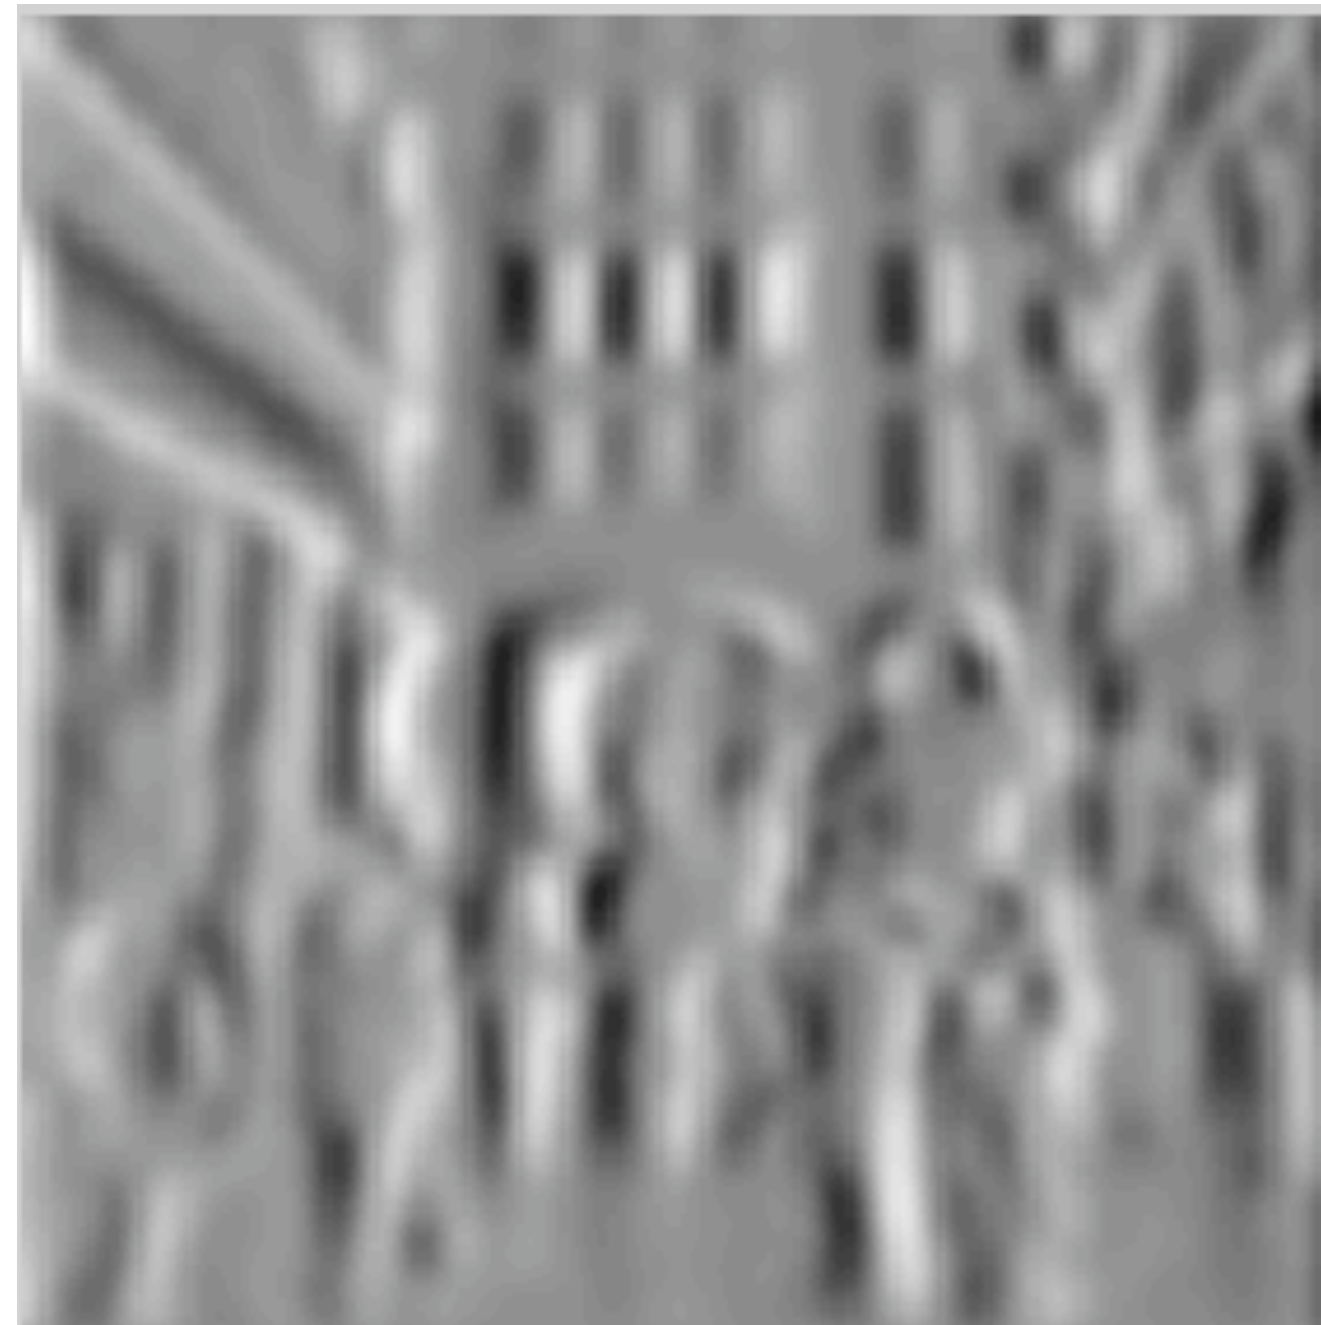
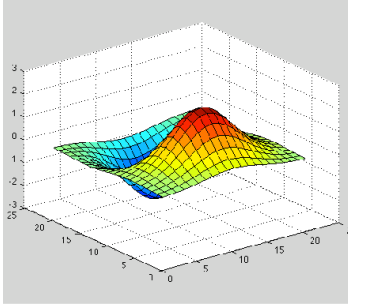


Orientation

$$g_x(x,y) = \frac{\partial g(x,y)}{\partial x} = \frac{-x}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

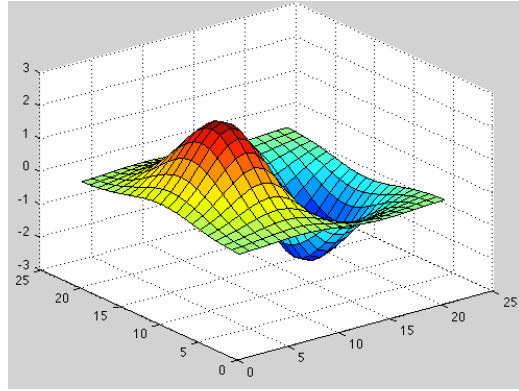


$$g_y(x,y) = \frac{\partial g(x,y)}{\partial y} = \frac{-y}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

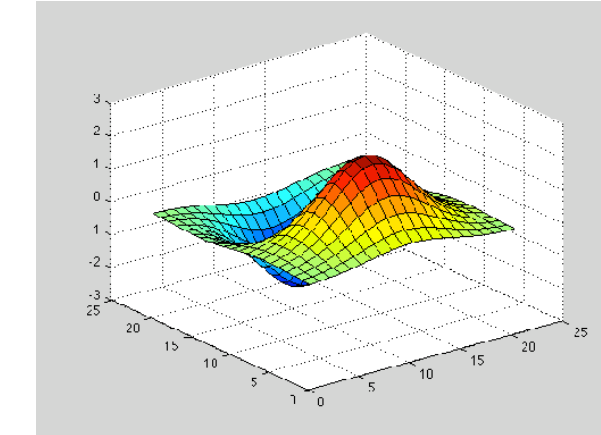


What about other orientations not axis aligned?

Orientation

$$g_x(x,y) = \frac{\partial g(x,y)}{\partial x} = \frac{-x}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$


$$g_y(x,y) = \frac{\partial g(x,y)}{\partial y} = \frac{-y}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$



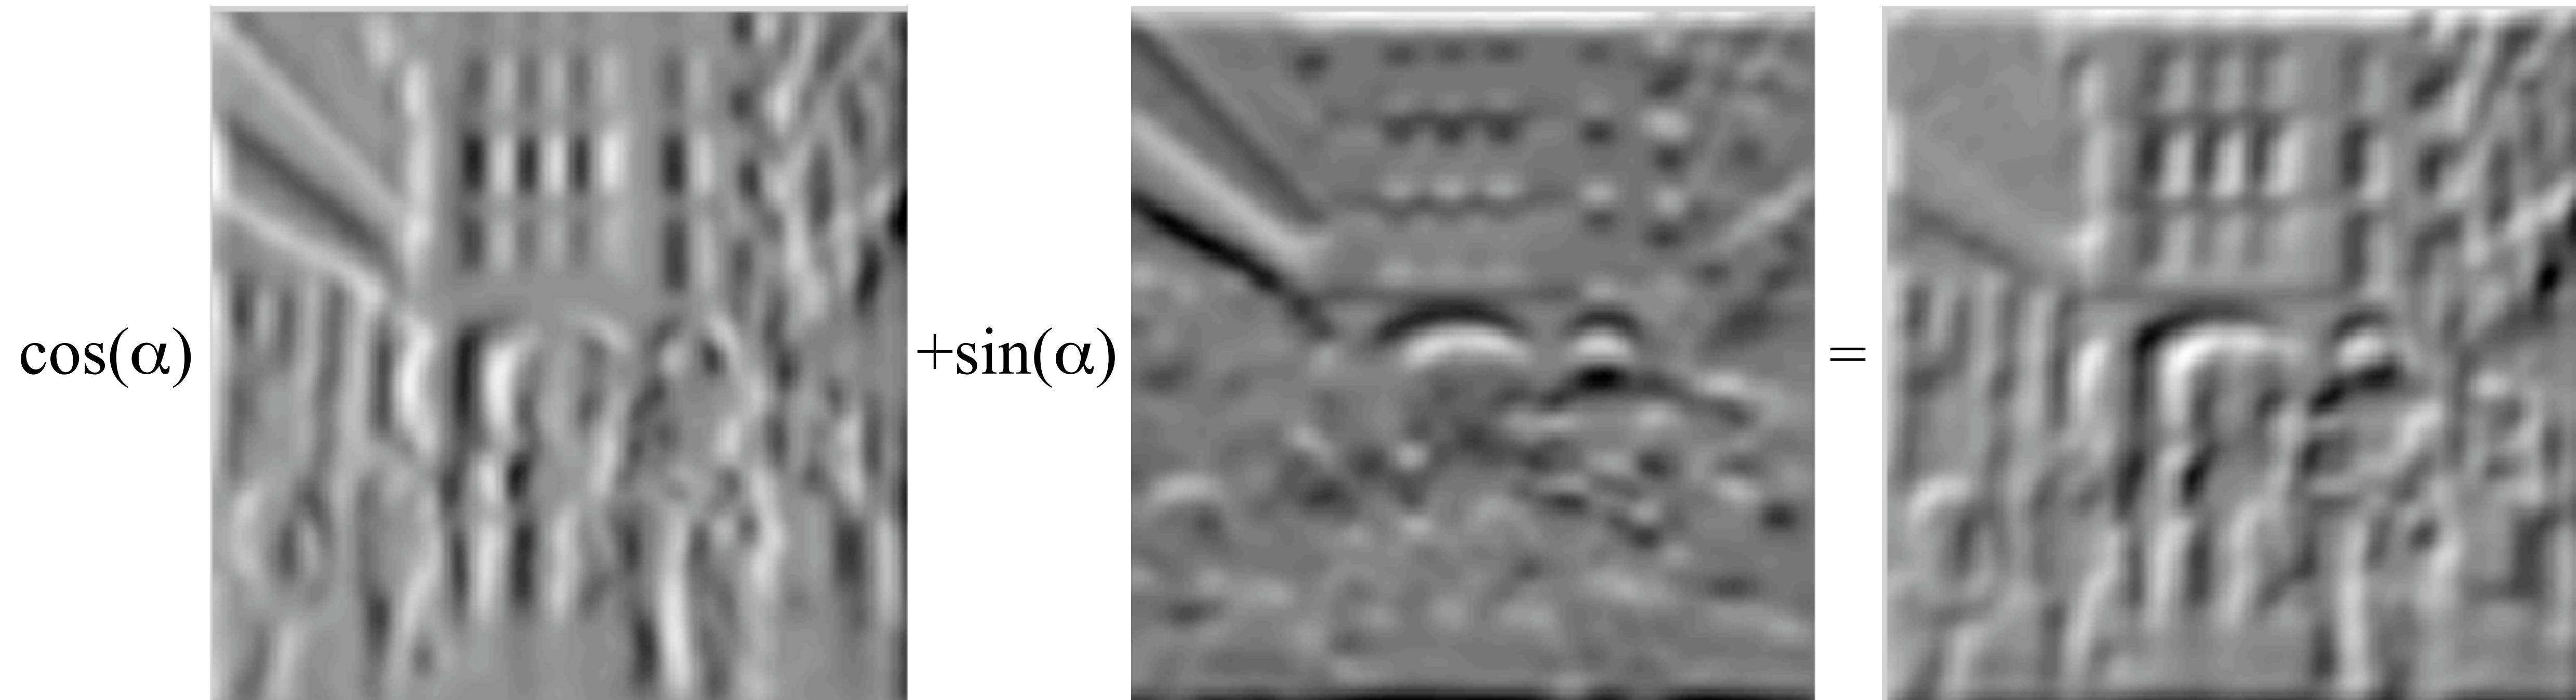
The smoothed directional gradient is a linear combination of two kernels

$$u^T \nabla g \otimes I = \left(\cos(\alpha) g_x(x,y) + \sin(\alpha) g_y(x,y) \right) \otimes I(x,y) =$$

Any orientation can be computed as a linear combination of two filtered images

$$= \cos(\alpha) g_x(x,y) \otimes I(x,y) + \sin(\alpha) g_y(x,y) \otimes I(x,y)$$

Orientation



Discretization 2D Gaussian derivatives

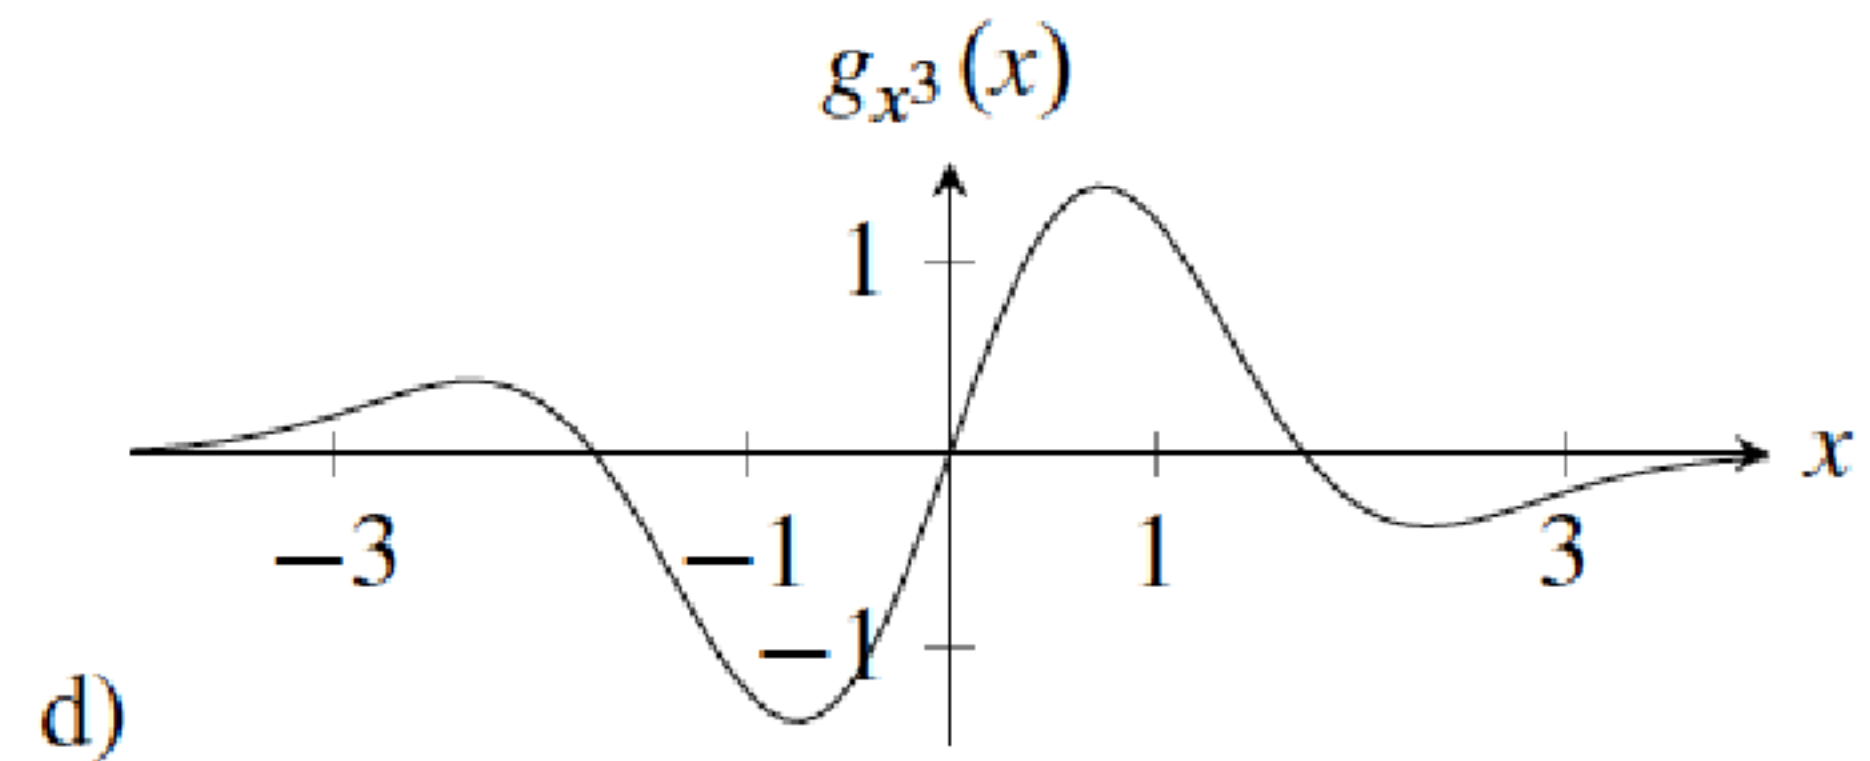
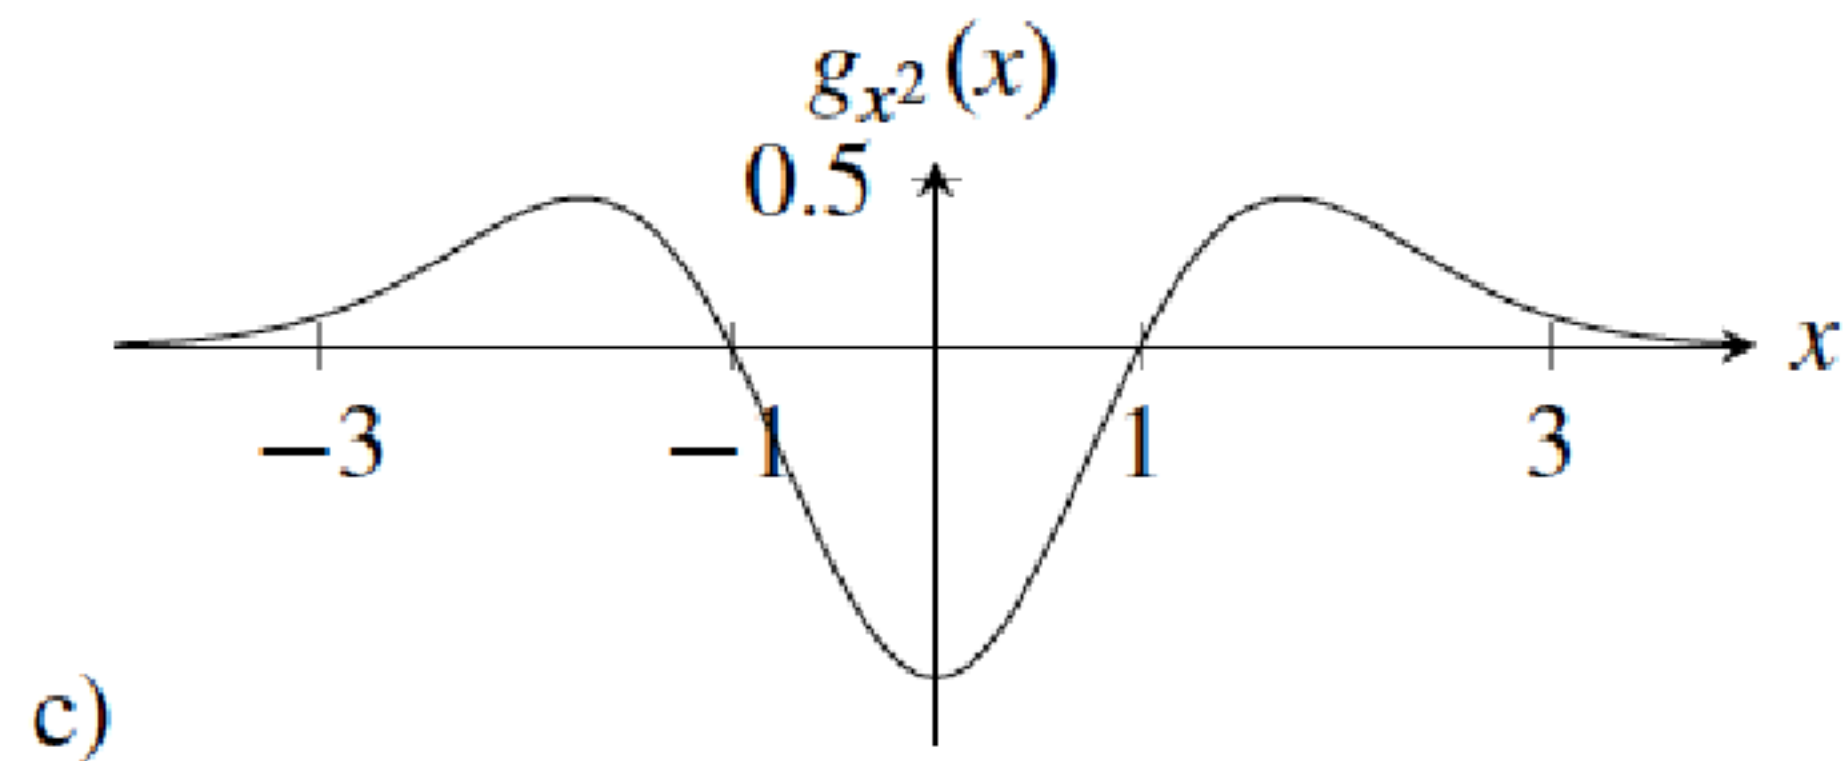
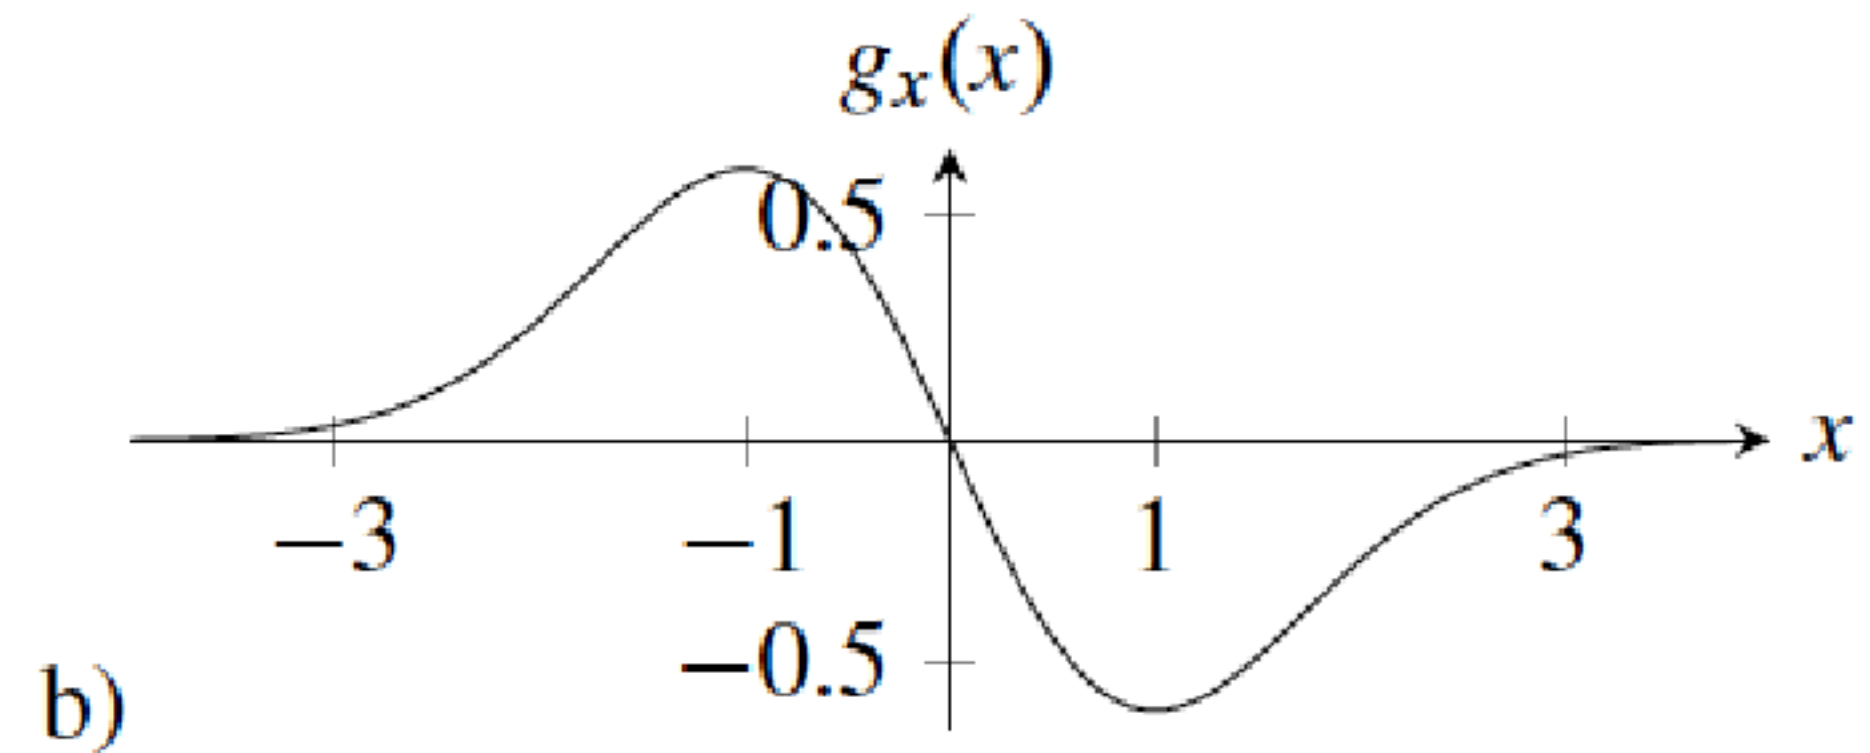
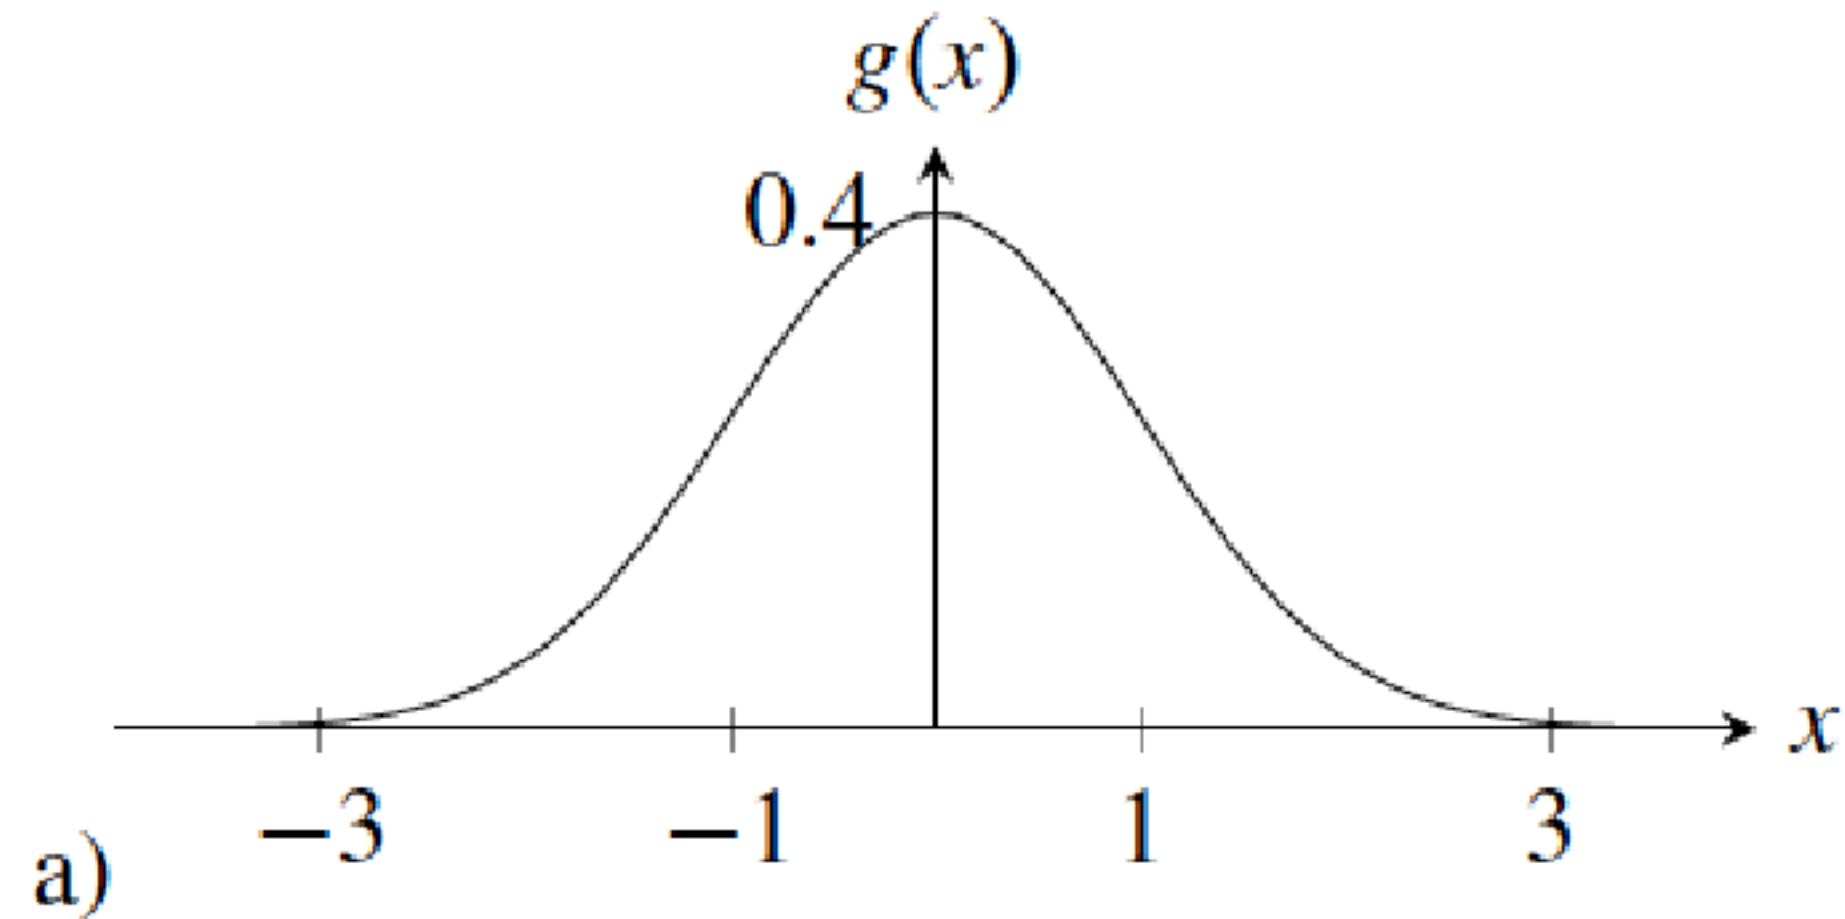
As Gaussians are separable, we can approximate two 1D derivatives and then convolve them.

One example is the Sobel-Feldman operator:

$$Sobel_x = \begin{bmatrix} 1 & 0 & -1 \end{bmatrix} \circ \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

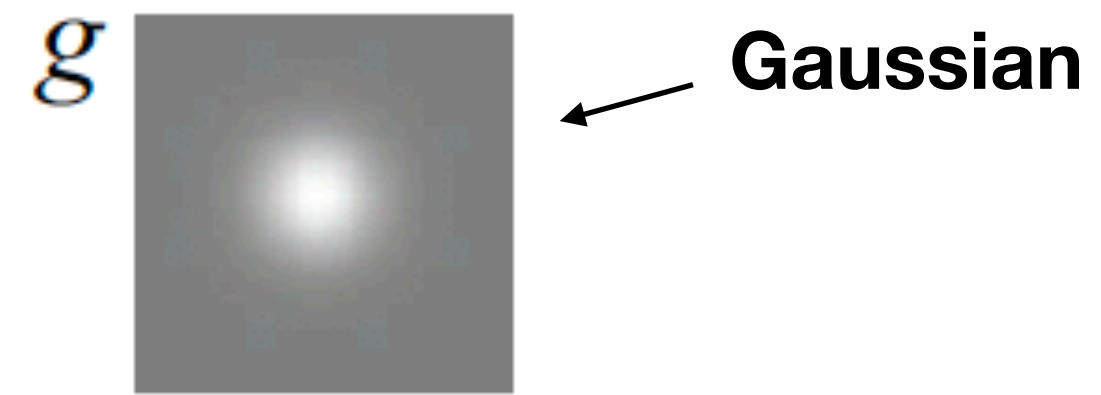
$$Sobel_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

n-th order Gaussian derivatives



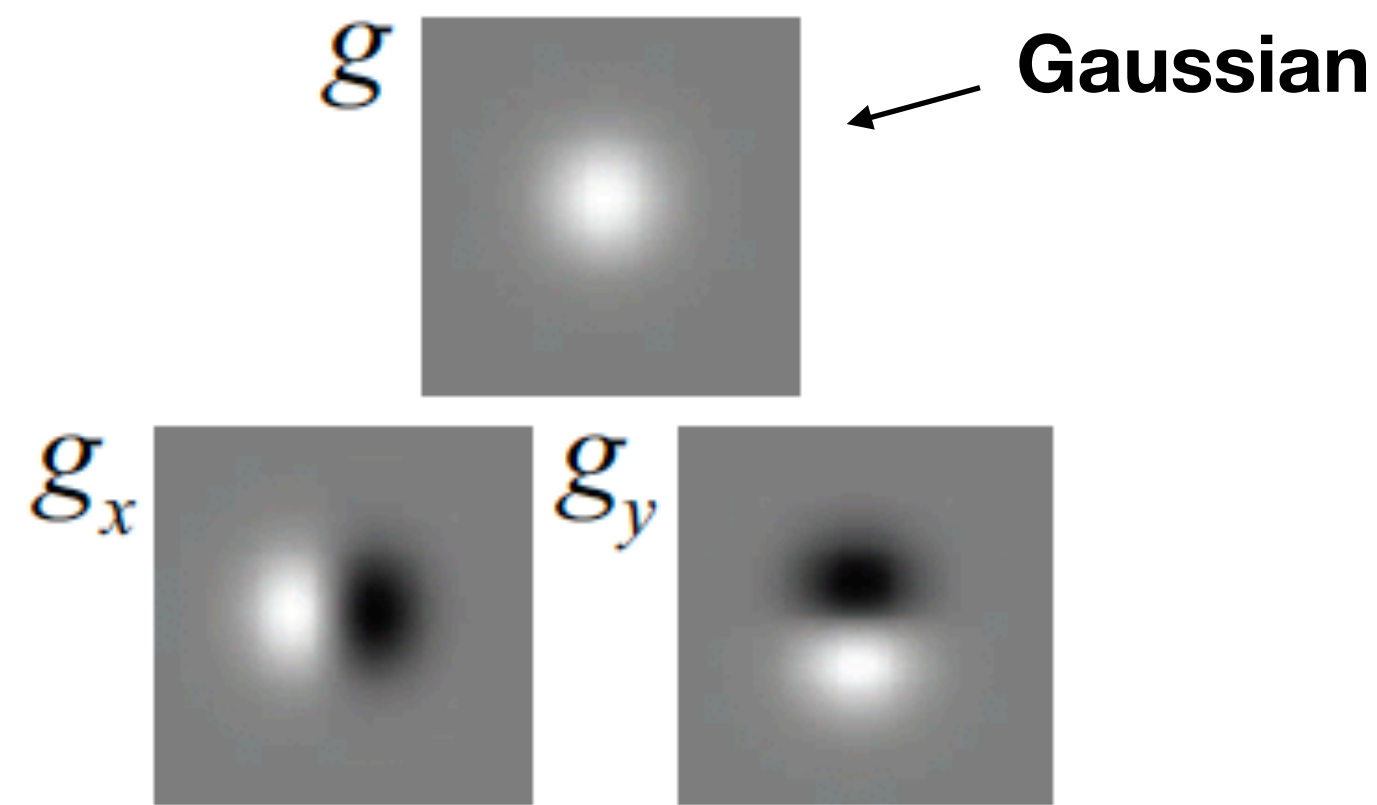
$$g_{x^n, y^m}(x, y; \sigma) = \frac{\partial^{n+m} g(x, y)}{\partial x^n \partial y^m} = \left(\frac{-1}{\sigma \sqrt{2}} \right)^{n+m} H_n \left(\frac{x}{\sigma \sqrt{2}} \right) H_m \left(\frac{y}{\sigma \sqrt{2}} \right) g(x, y; \sigma)$$

n-th order Gaussian derivatives



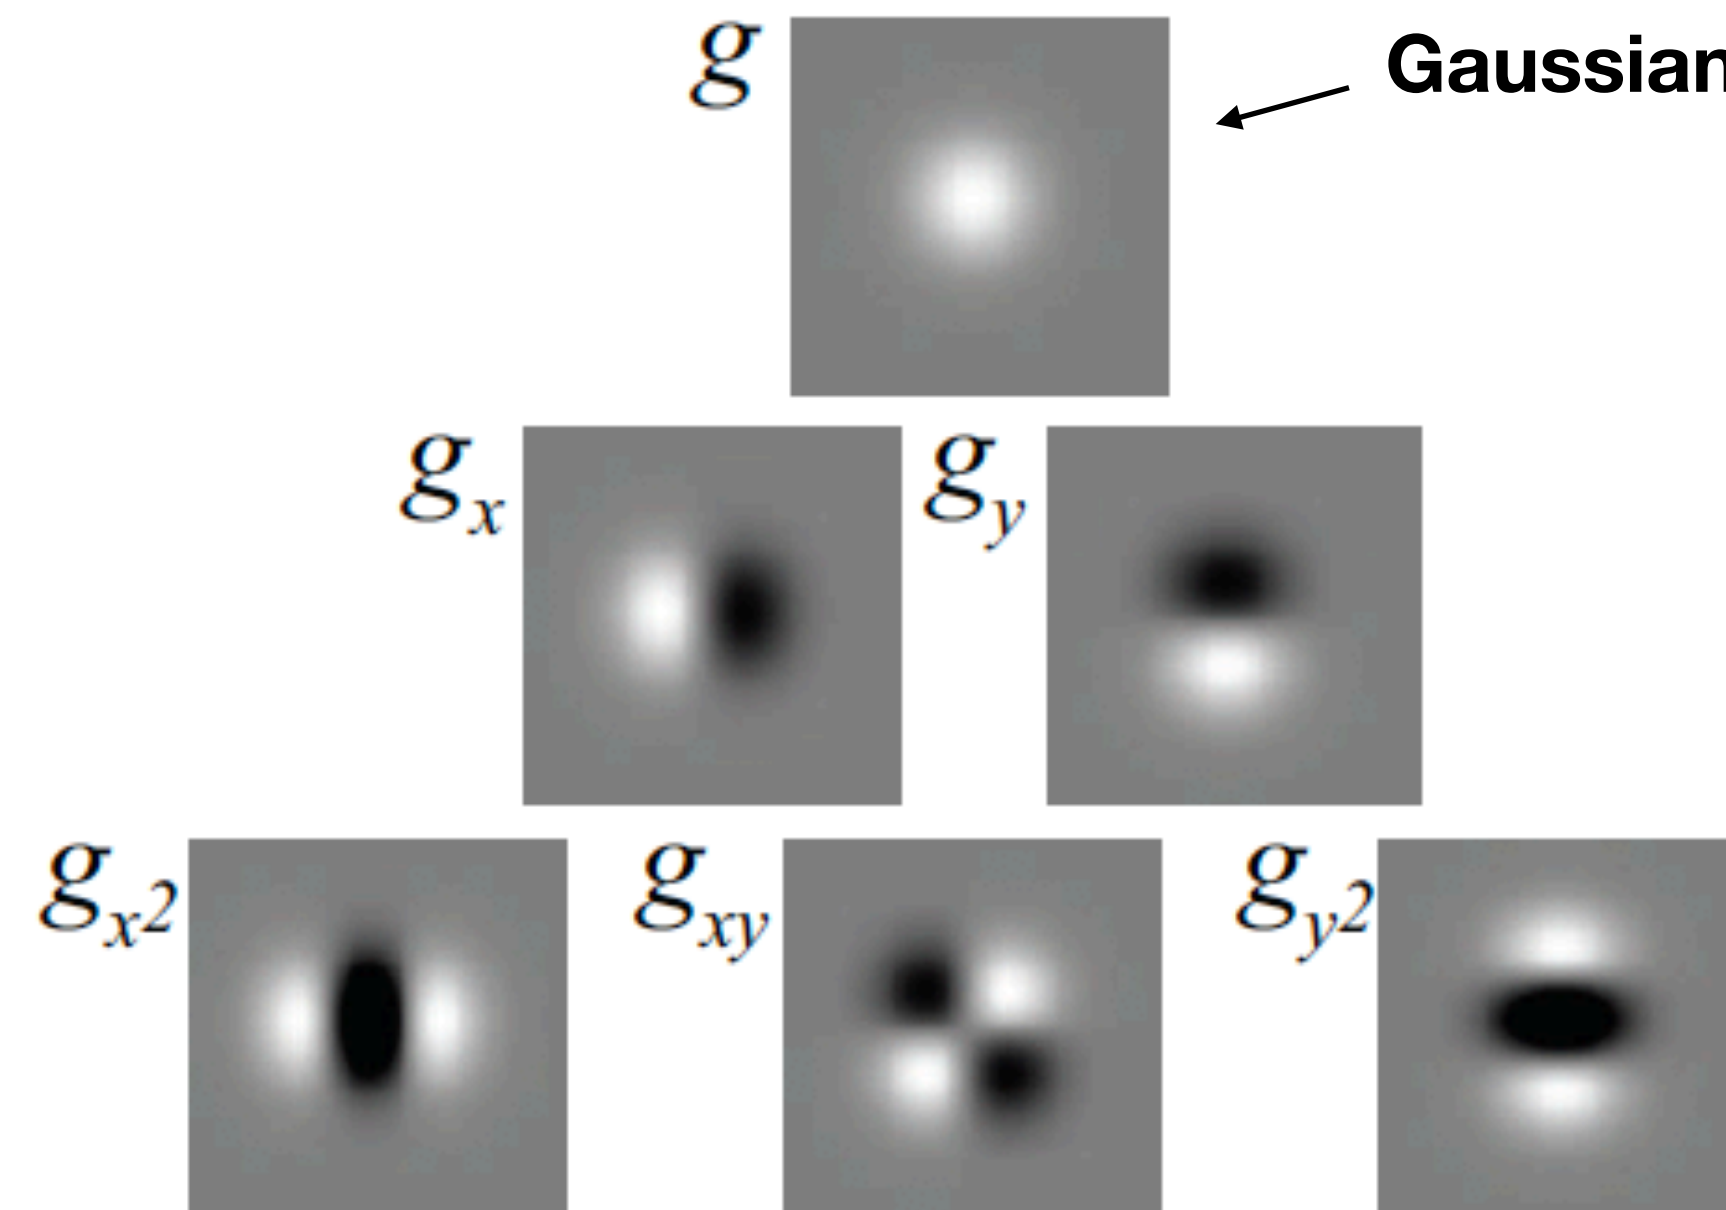
$$g_{x^n, y^m}(x, y; \sigma) = \frac{\partial^{n+m} g(x, y)}{\partial x^n \partial y^m} = \left(\frac{-1}{\sigma \sqrt{2}} \right)^{n+m} H_n \left(\frac{x}{\sigma \sqrt{2}} \right) H_m \left(\frac{y}{\sigma \sqrt{2}} \right) g(x, y; \sigma)$$

n-th order Gaussian derivatives



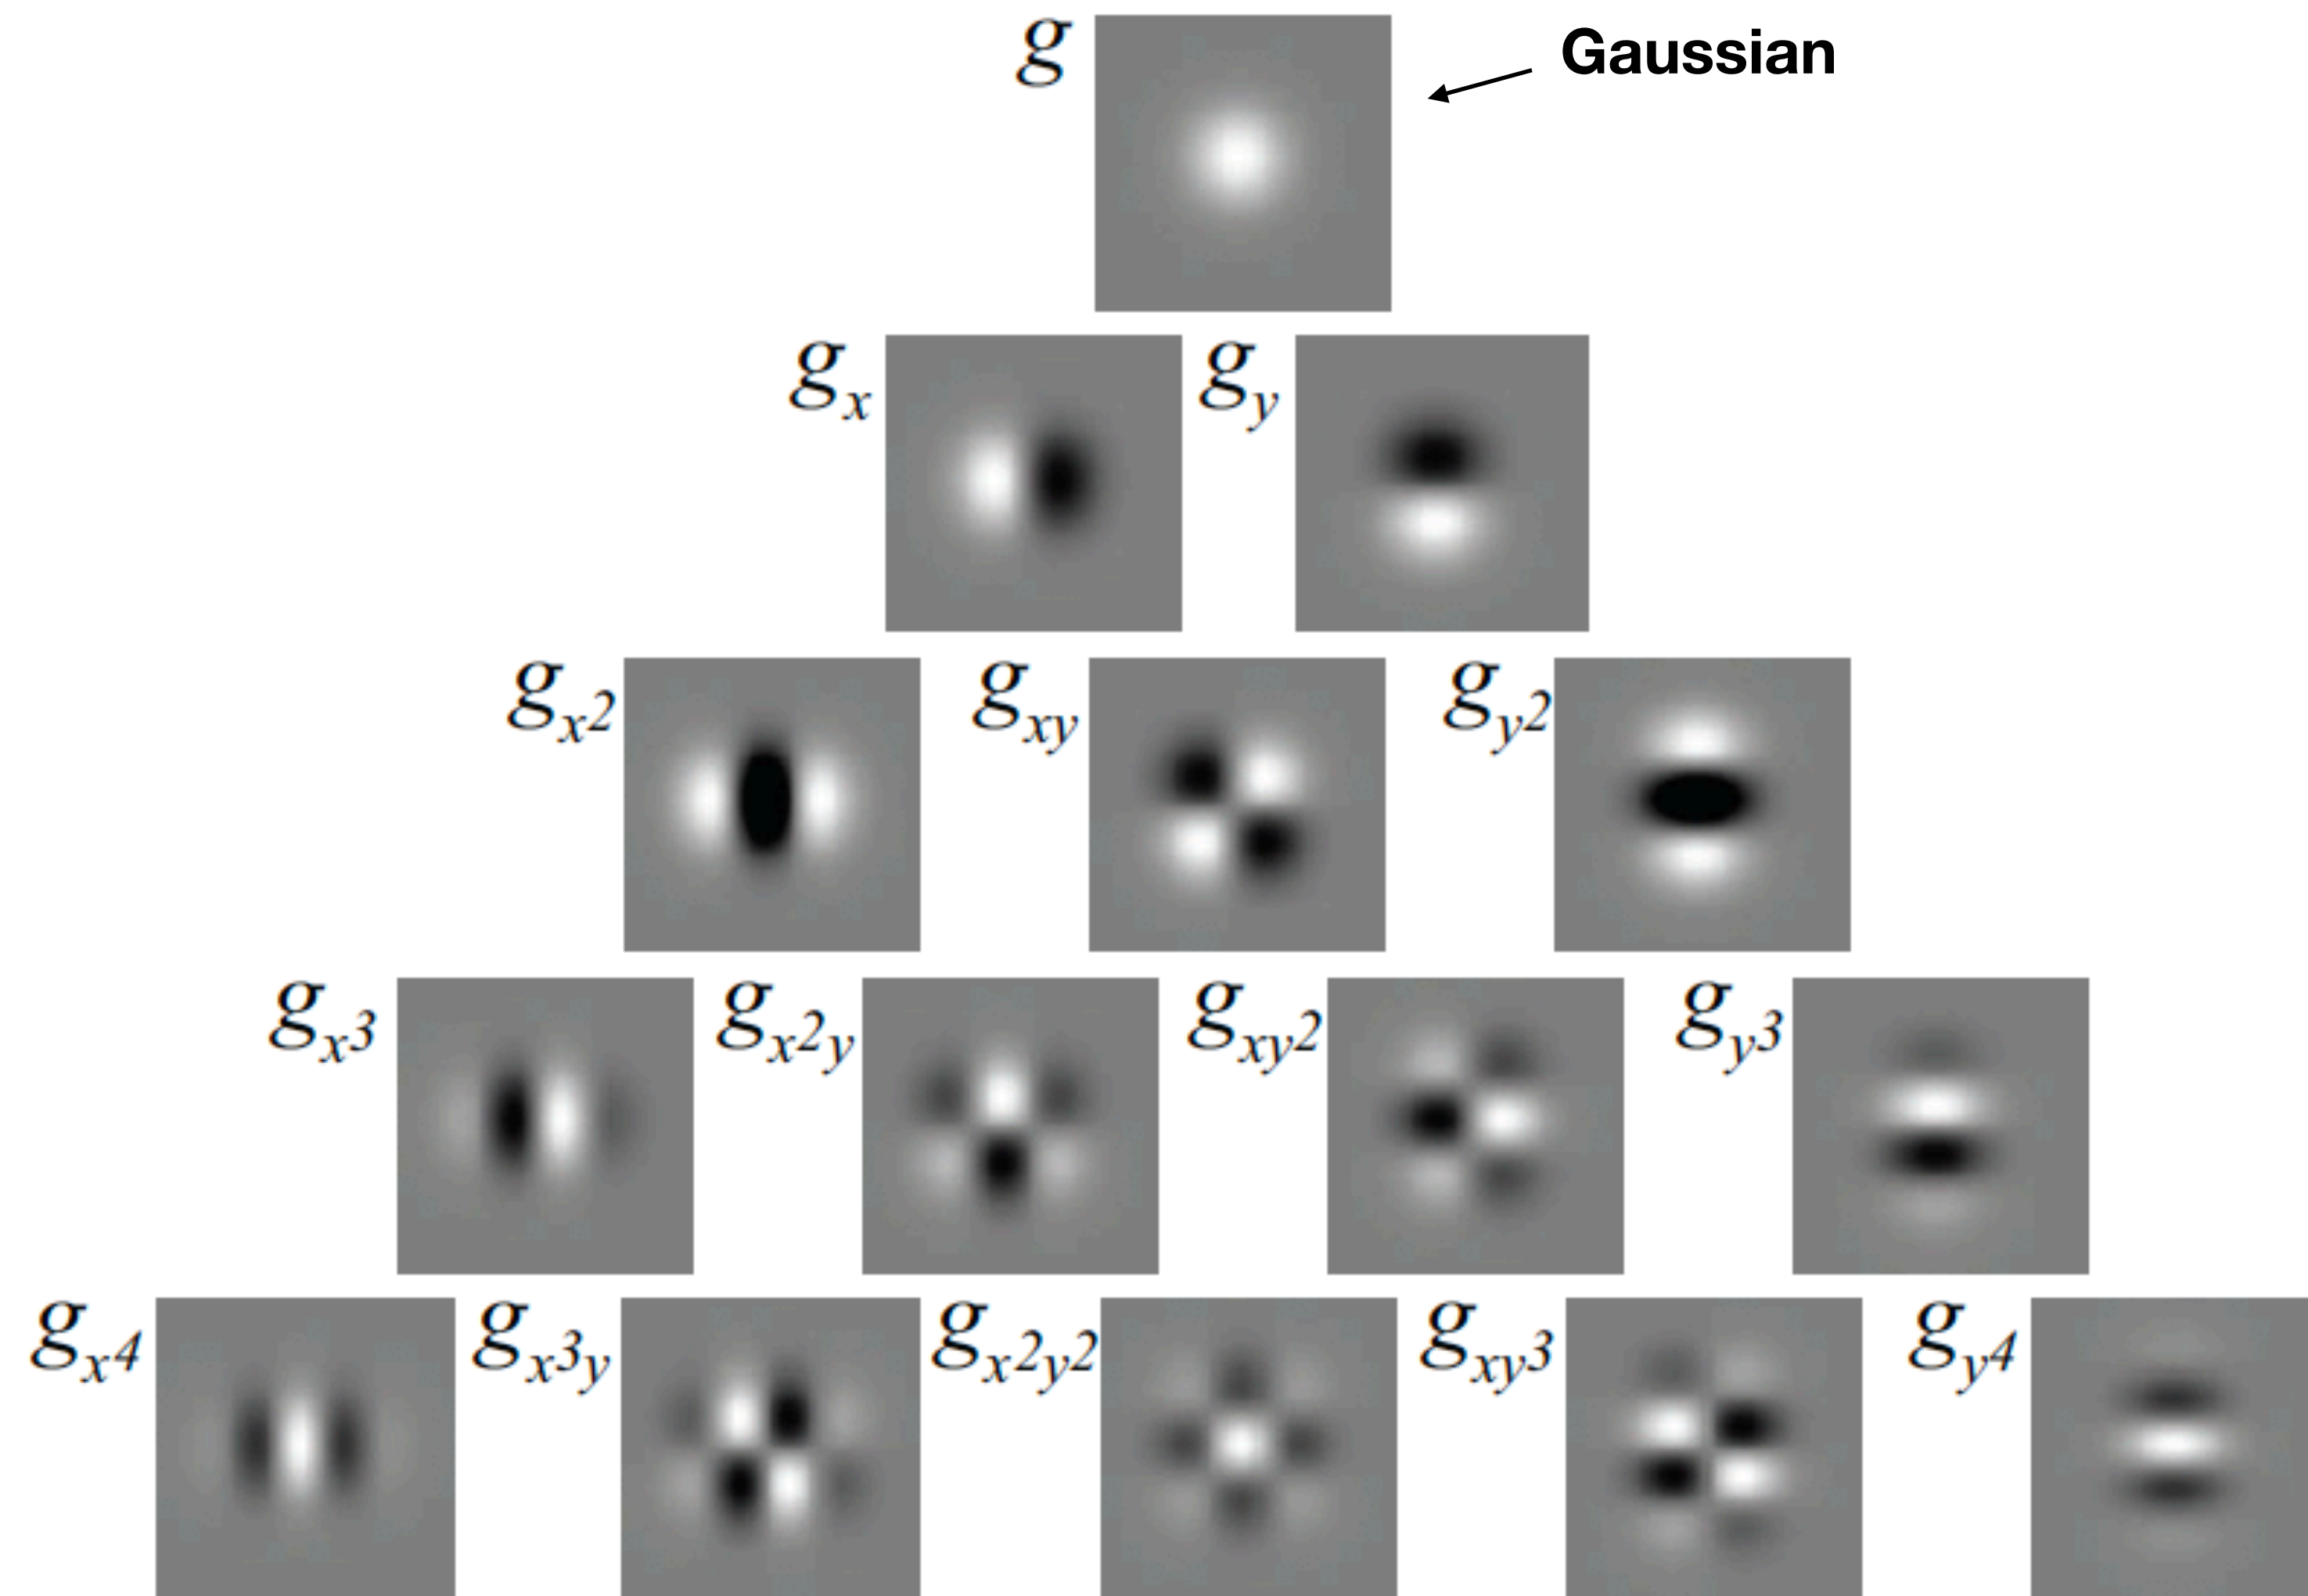
$$g_{x^n, y^m}(x, y; \sigma) = \frac{\partial^{n+m} g(x, y)}{\partial x^n \partial y^m} = \left(\frac{-1}{\sigma \sqrt{2}} \right)^{n+m} H_n \left(\frac{x}{\sigma \sqrt{2}} \right) H_m \left(\frac{y}{\sigma \sqrt{2}} \right) g(x, y; \sigma)$$

n-th order Gaussian derivatives



$$g_{x^n, y^m}(x, y; \sigma) = \frac{\partial^{n+m} g(x, y)}{\partial x^n \partial y^m} = \left(\frac{-1}{\sigma \sqrt{2}} \right)^{n+m} H_n \left(\frac{x}{\sigma \sqrt{2}} \right) H_m \left(\frac{y}{\sigma \sqrt{2}} \right) g(x, y; \sigma)$$

n-th order Gaussian derivatives



$$g_{x^n, y^m}(x, y; \sigma) = \frac{\partial^{n+m} g(x, y)}{\partial x^n \partial y^m} = \left(\frac{-1}{\sigma \sqrt{2}} \right)^{n+m} H_n \left(\frac{x}{\sigma \sqrt{2}} \right) H_m \left(\frac{y}{\sigma \sqrt{2}} \right) g(x, y; \sigma)$$

Image sharpening filter



Image sharpening filter

Subtract away the blurred components of the image:

$$\text{sharpening filter} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} - \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

This filter has an overall DC component of 1. It de-emphasizes the blur component of the image (low spatial frequencies).

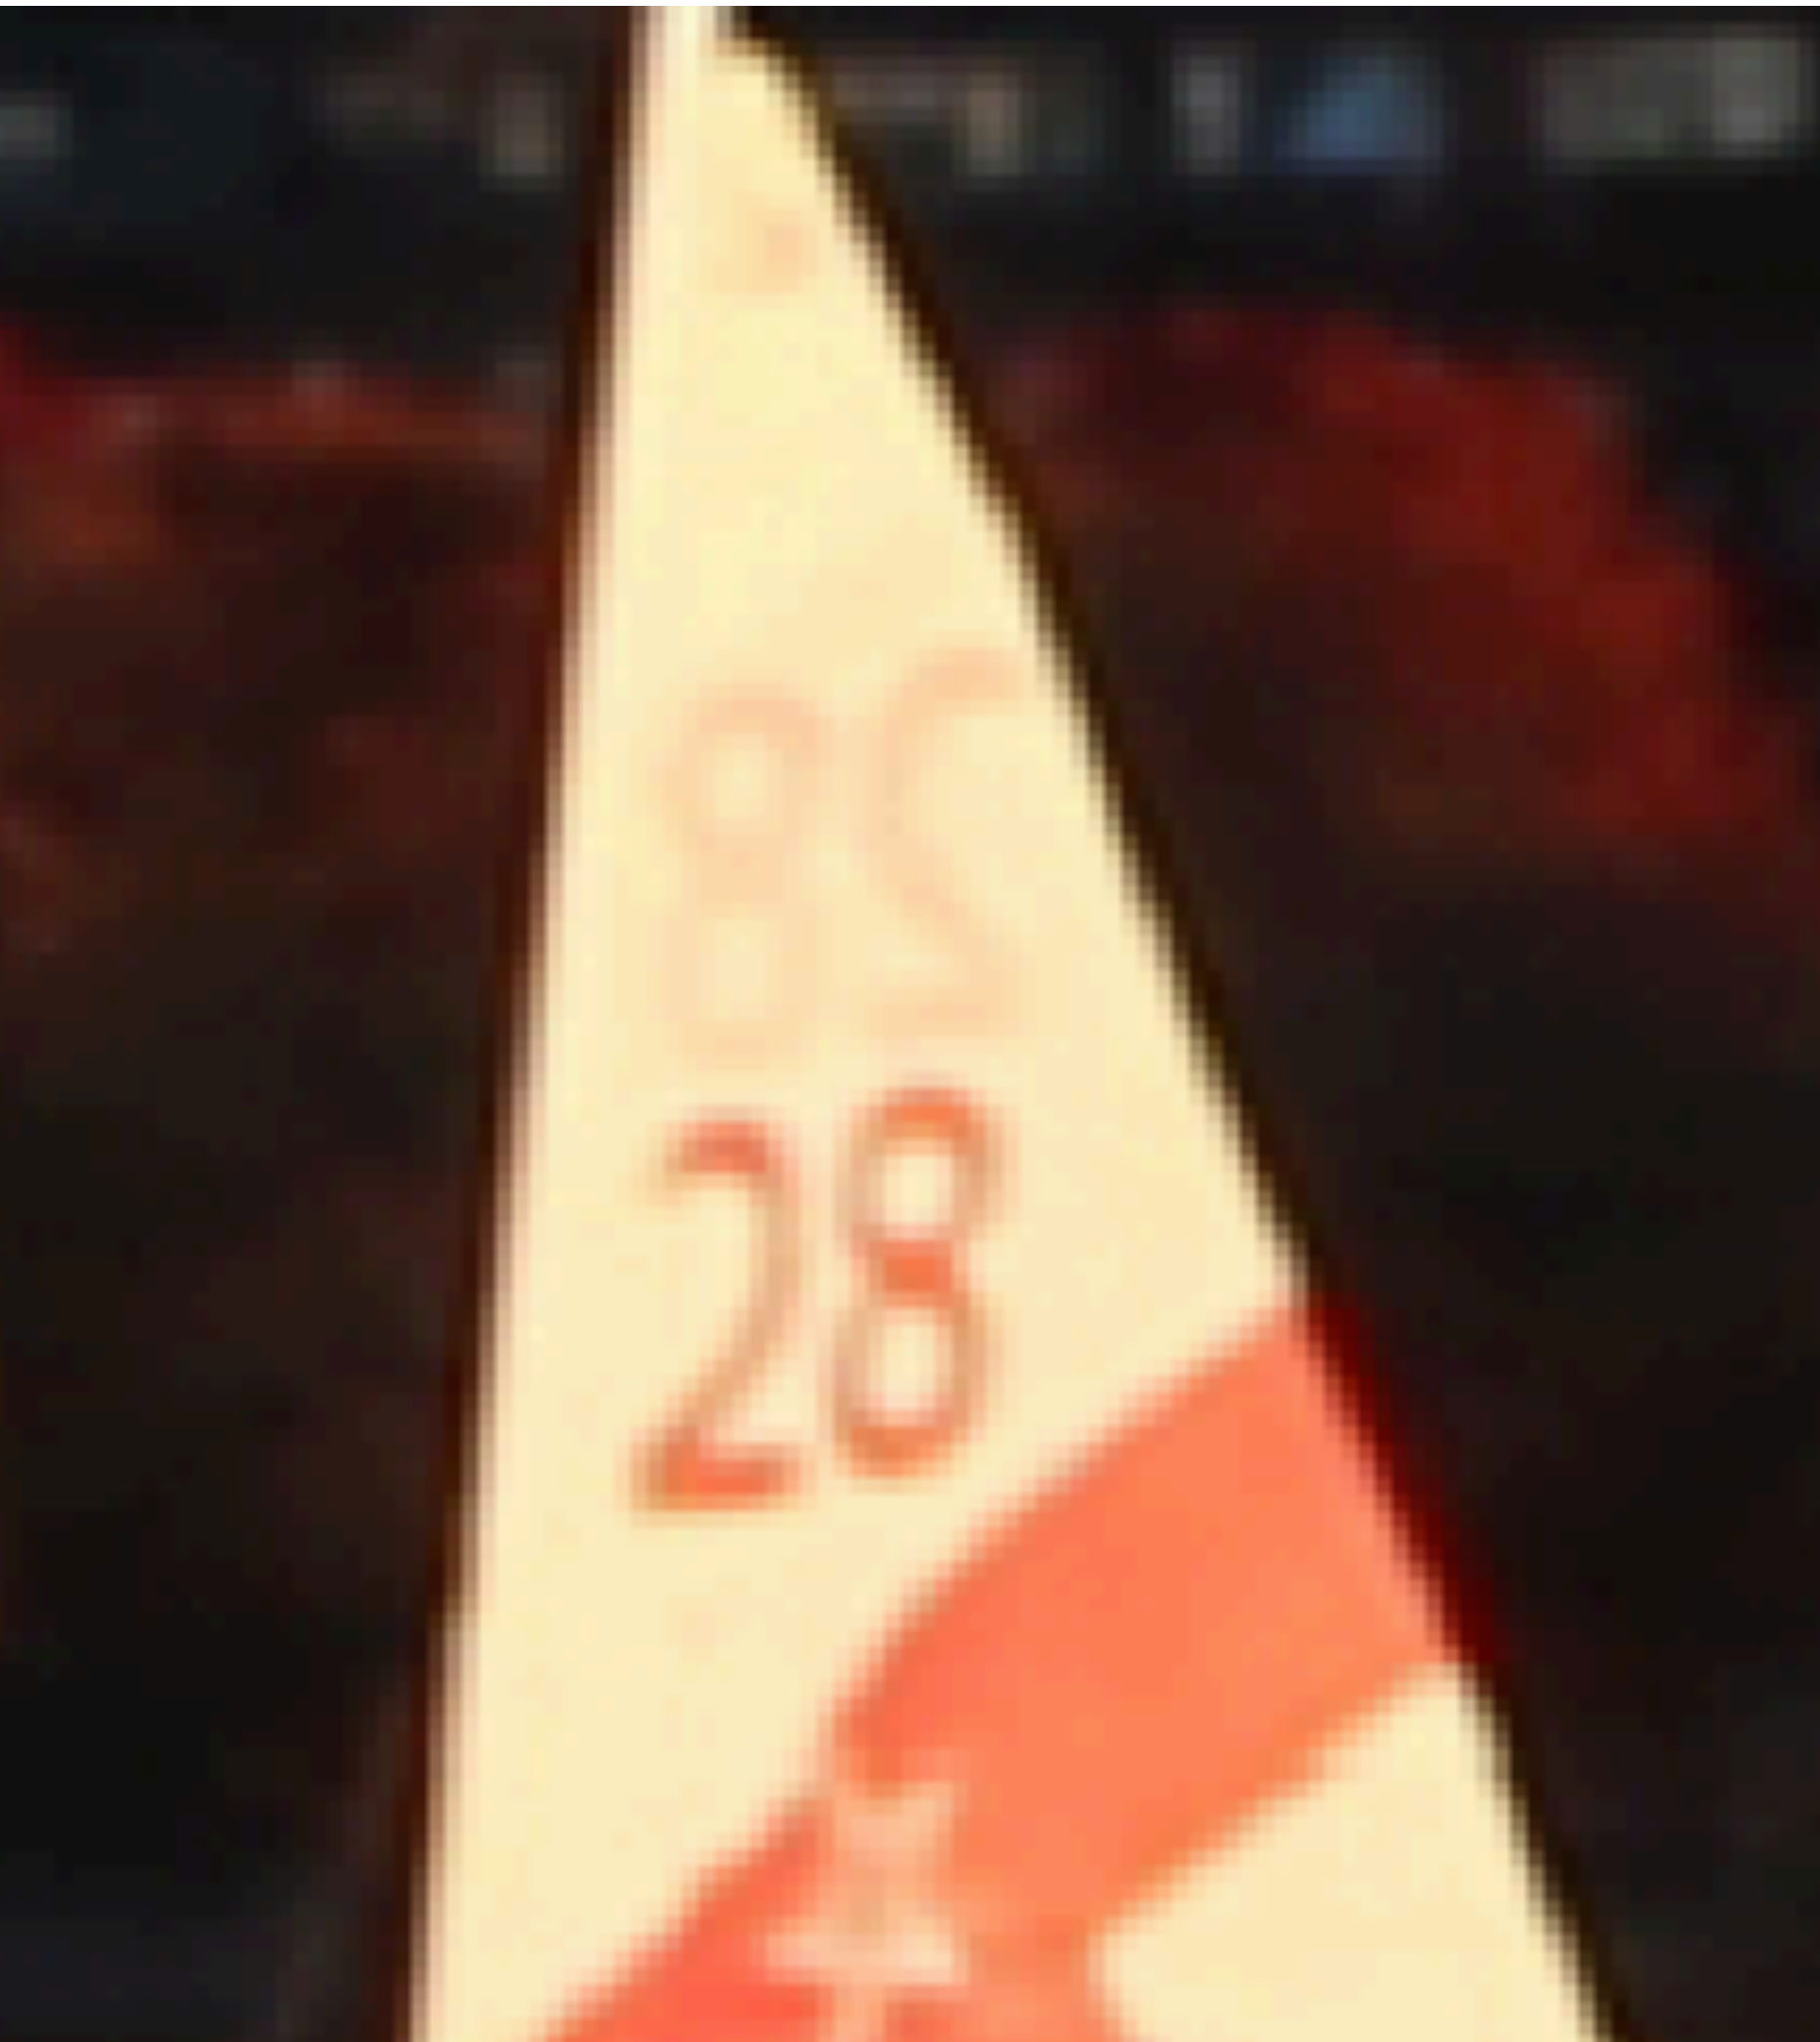
Input image



Sharpened



Input image



Sharpened



Input image

