


Lecture 3


Signal Processing

material for the course:
textbook class notes
lectures
lecture slides
lecture recordings

6.003 Signals and systems






Time continuous signal




Time discrete signal

Some visual areas...



From M. Lewicky

Signals and systems



One important class of systems is the set of linear systems.


A function f is linear if it satisfies:

$$f(\alpha \mathbf{x}) = \alpha f(\mathbf{x})$$

$$f(\mathbf{x} + \mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})$$

Linear system: $y = f(\mathbf{x})$


A linear function f can be written as a matrix multiplication:



n indexes rows,
 k indexes columns


Linear system: $y = f(\mathbf{x})$

A linear function f can be written as a matrix multiplication:



n indexes rows,
 k indexes columns


It can also be represented as a fully connected linear neural network



$h[n, k]$ Is the strength of the connection between $x[k]$ and $y[n]$


Linear system: $y = f(\mathbf{x})$

A linear function f can be written as a matrix multiplication:



n indexes rows,
 k indexes columns


It can also be represented as a fully connected linear neural network




$h[n, k]$ is the strength of the connection between $x[k]$ and $y[n]$


Images are turned into column vectors by concatenating all image columns

4x4 image



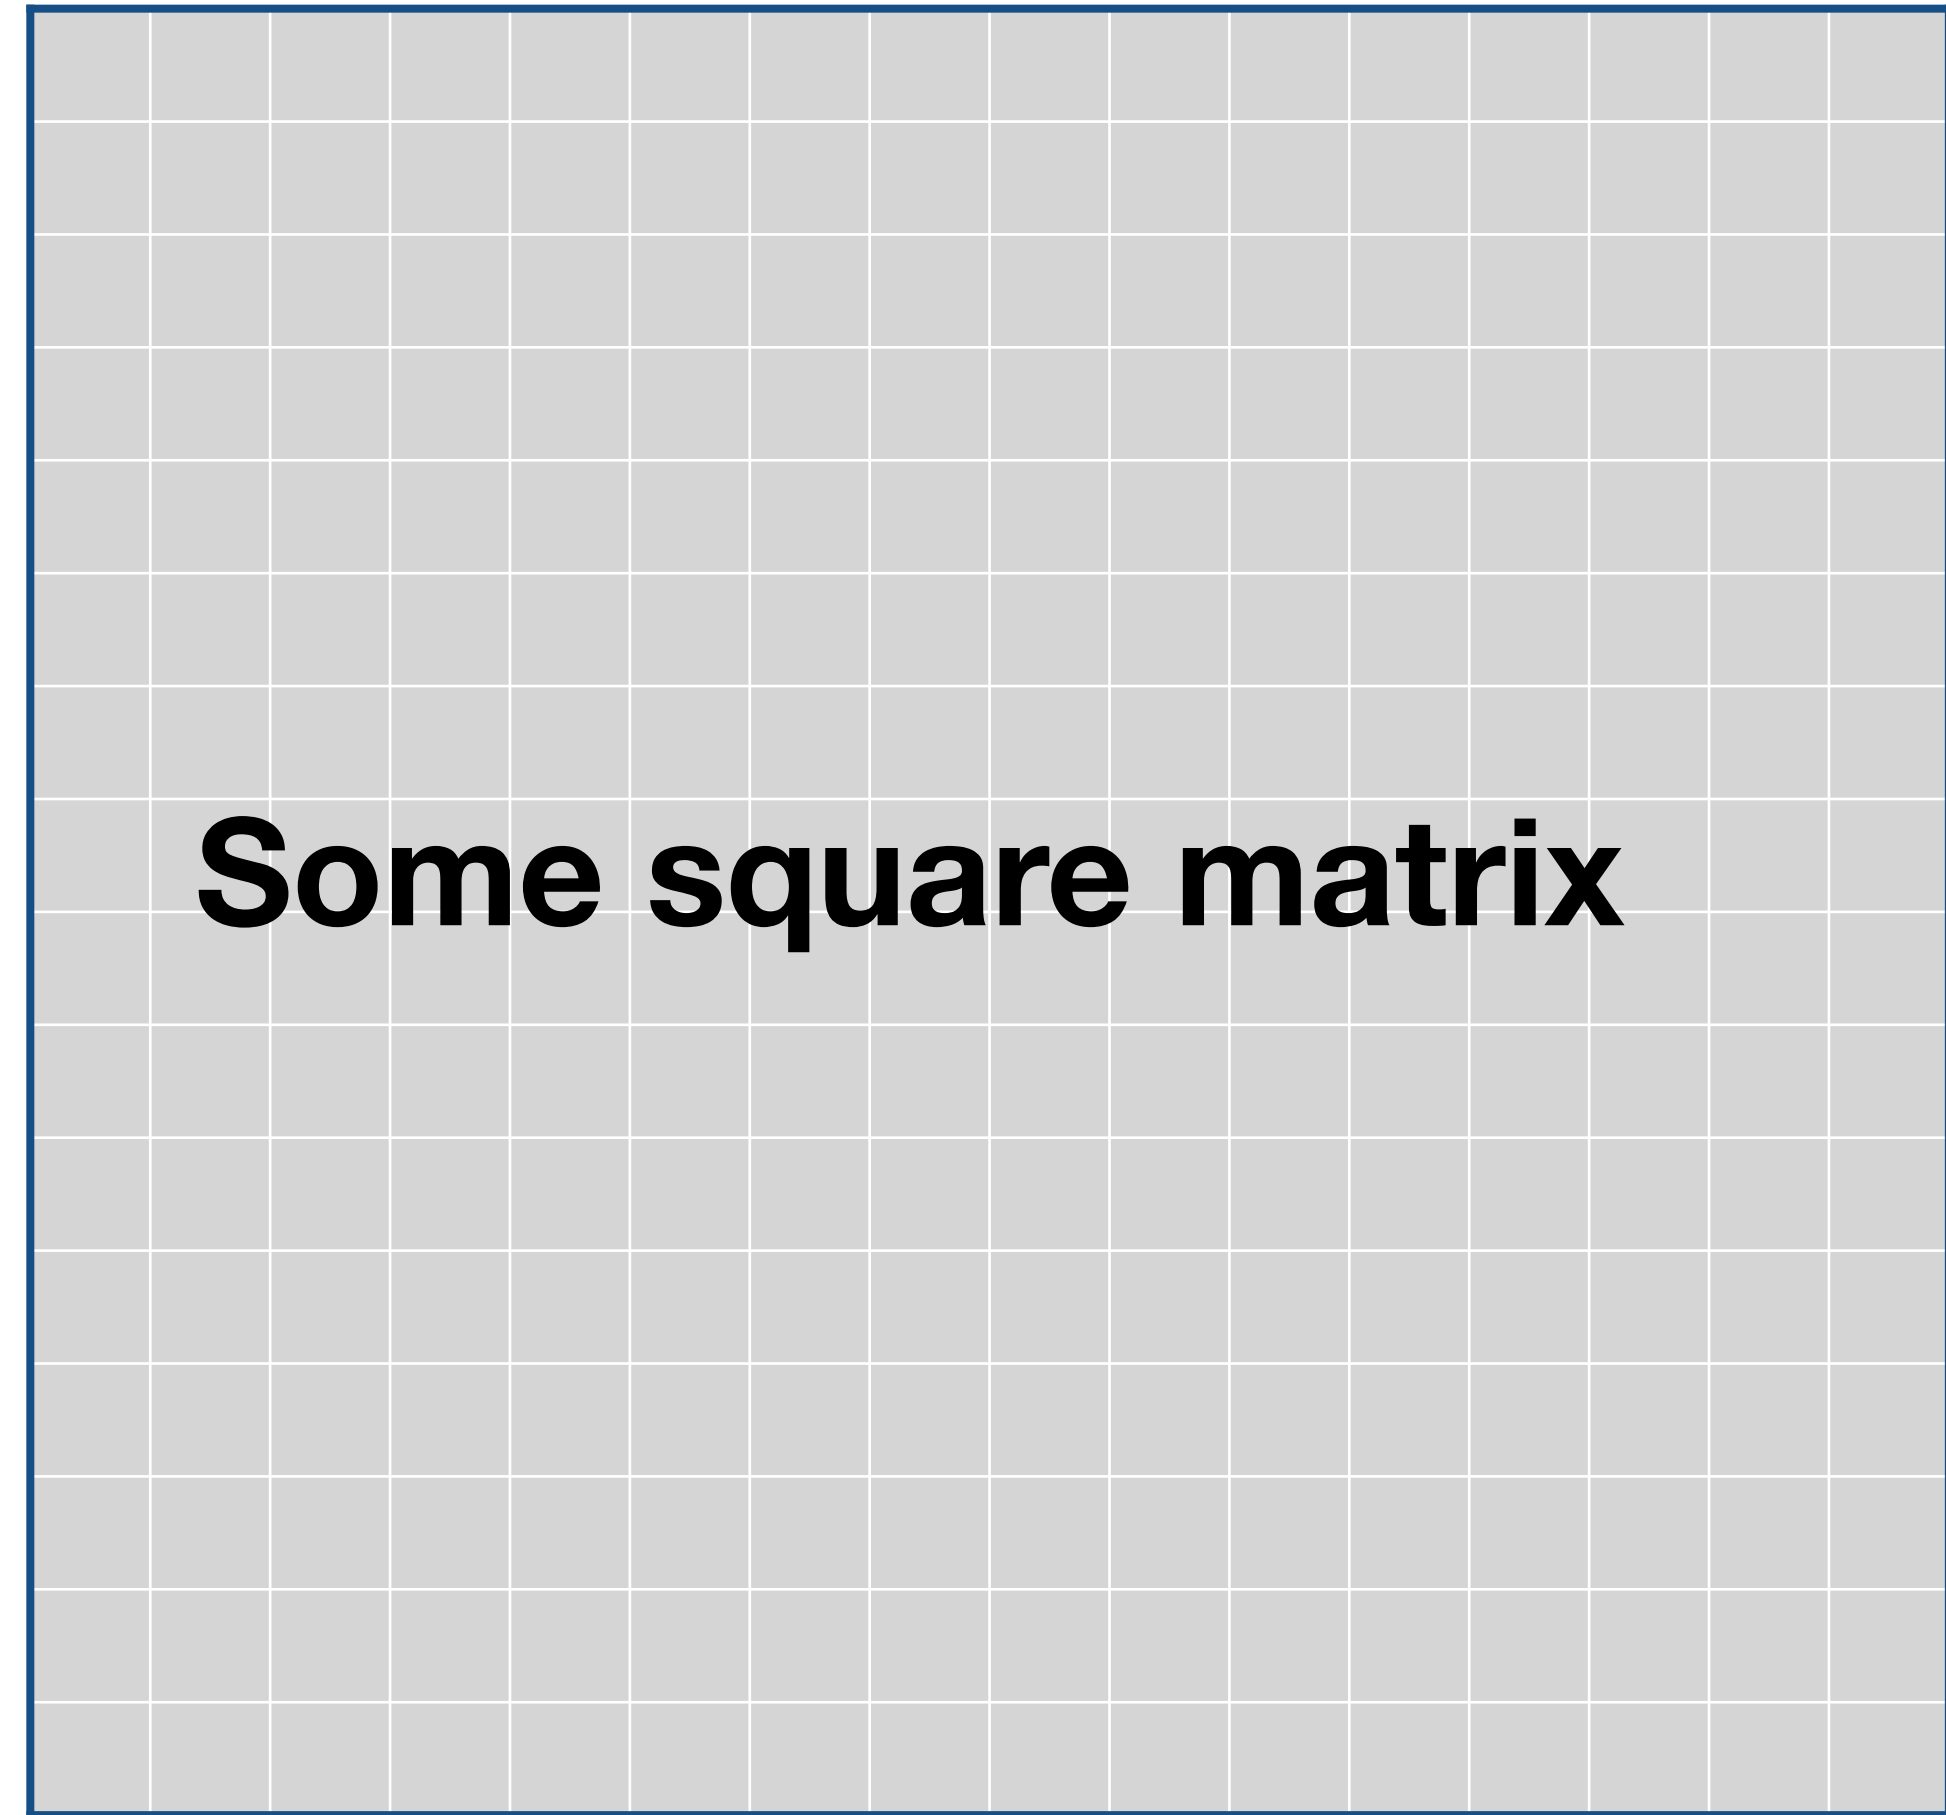


Column vector of length 16






=








4x4 image




Quiz: what operation is linear?

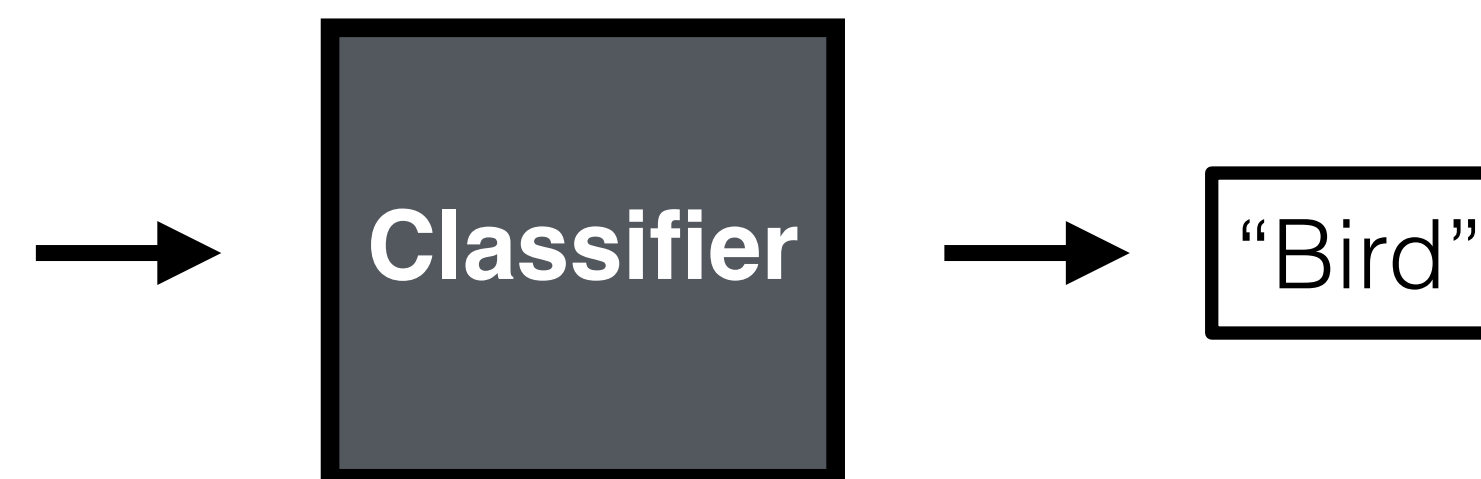

Quiz: what operation is linear?










We need translation invariance















Sky	Sky	Sky	Sky	Sky	Sky	Sky	Bird
Sky	Sky	Sky	Sky	Sky	Sky	Sky	Sky
Sky	Sky	Sky	Sky	Sky	Sky	Sky	Sky
Bird	Bird	Bird	Sky	Bird	Sky	Sky	Sky
Sky	Sky	Sky	Bird	Sky	Sky	Sky	Sky




Convolutional Neural Networks

- Neural network with specialized connectivity structure



LeCun et al. 1989

Convolution



The same weighting occurs within each window

Convolution: running example

x

0	0	0	0	0	0	0	0
0	0	0	0	0	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	0	1	1	1	0	0	0
0	0	1	1	1	0	0	0
0	0	0	0	0	0	0	0

Convolution: running example

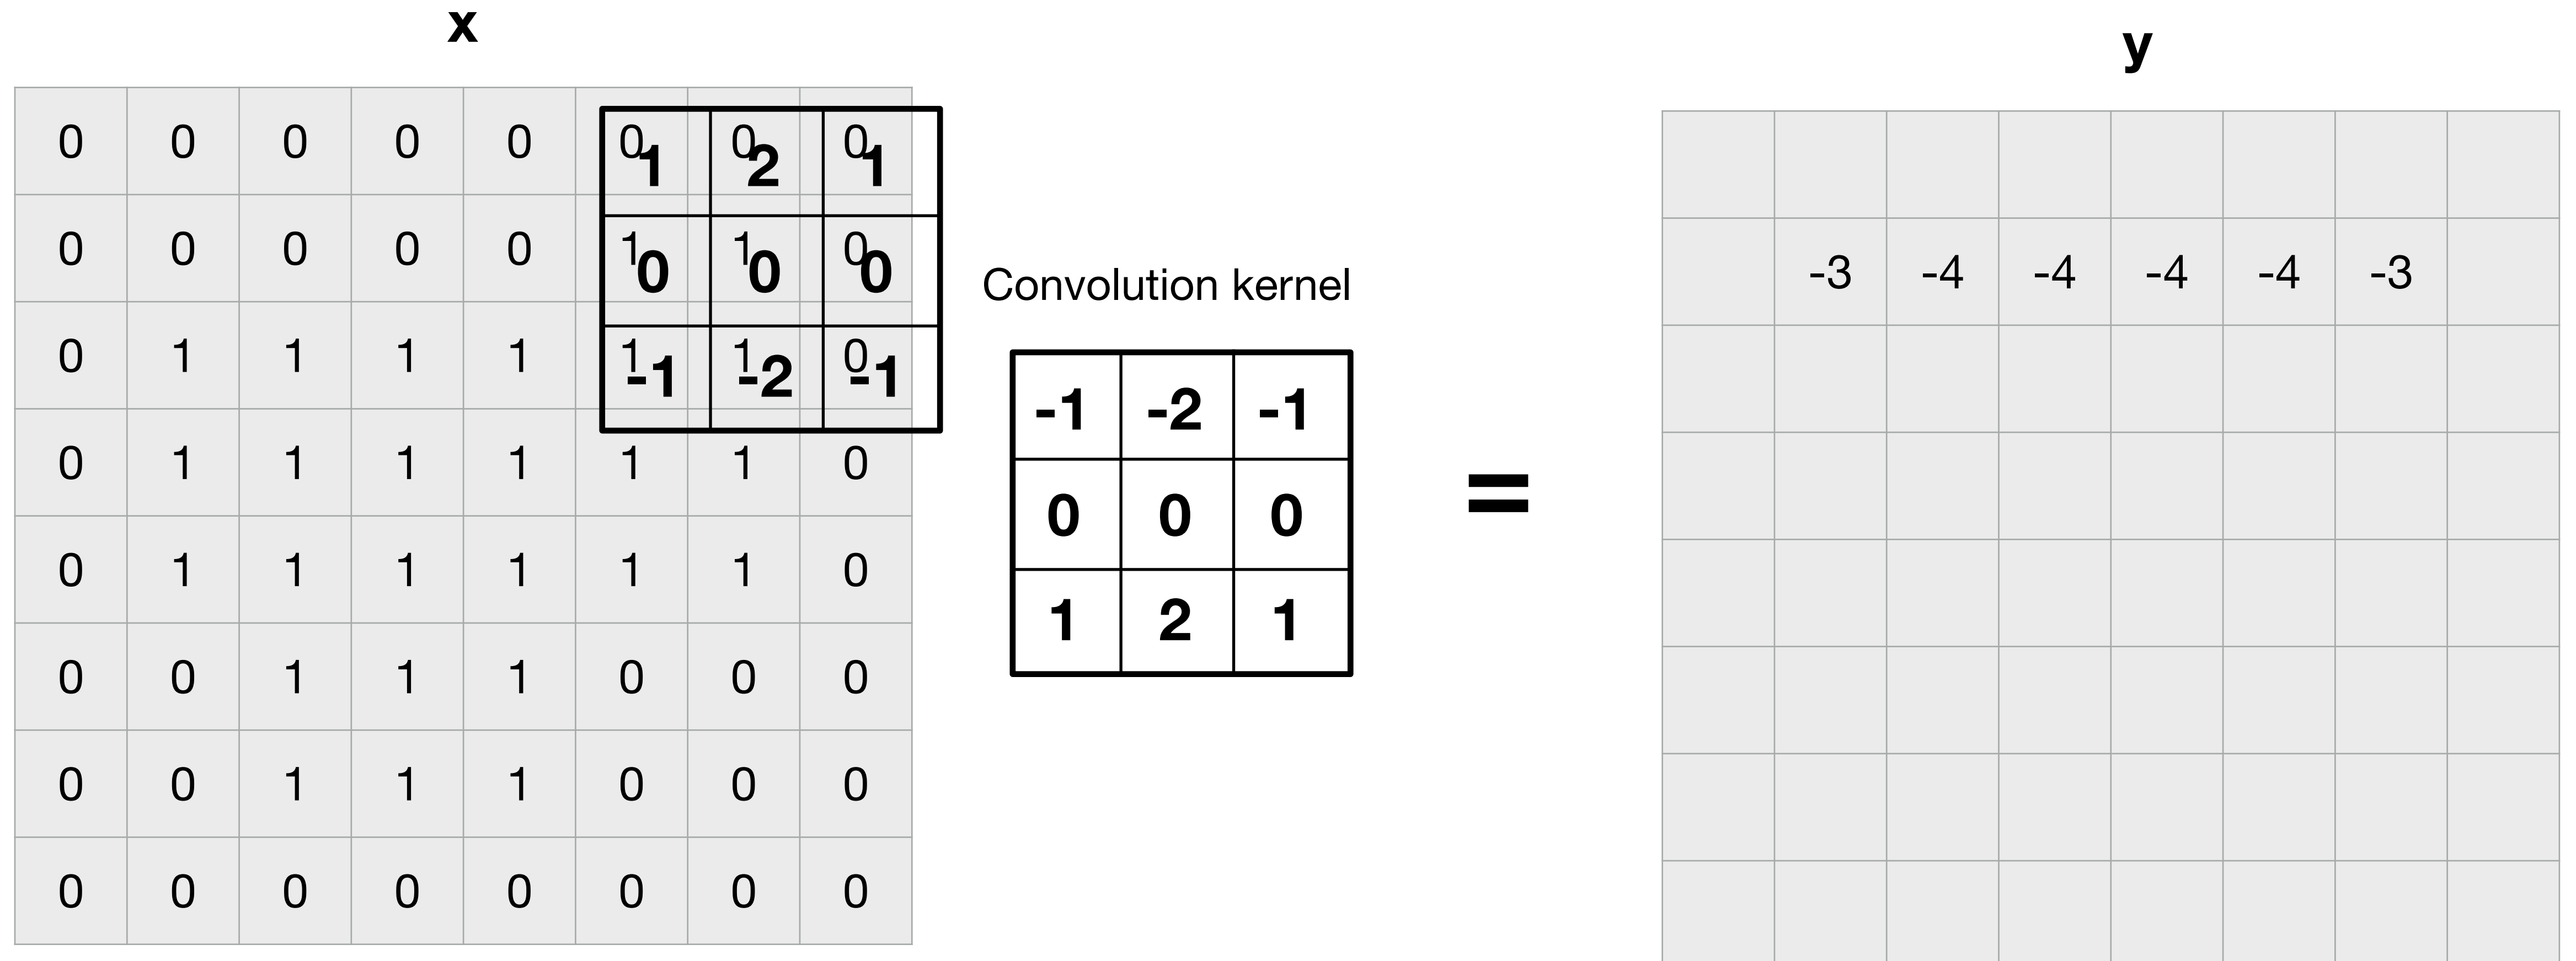
x

0	0	0	0	0	0	0	0
0	0	0	0	0	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	0	1	1	1	0	0	0
0	0	1	1	1	0	0	0
0	0	0	0	0	0	0	0

Convolution kernel

-1	-2	-1
0	0	0
1	2	1

Convolution: running example



Convolution: running example

x

0	0	0	0	0	0	0	0
0	0	0	0	0	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	0	1	1	1			
0	0	1	1	1			
0	0	0	0	0			

⁰ 1	⁰ 2	⁰ 1
⁰ 0	⁰ 0	⁰ 0
⁰ -1	⁰ -2	⁰ -1

Convolution kernel

-1	-2	-1
0	0	0
1	2	1

=

y

		-3	-4	-4	-4	-4	-3
		-3	-4	-4	-3	-1	0
		0	0	0	0	0	0
		2	1	0	1	3	3
		2	1	0	1	3	3
		1	3	4	3	1	0

Convolution: running example

x

0	0	0	0	0	0	0	0
0	0	0	0	0	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	1	1	1	1	1	1	0
0	0	1	1	1	0	0	0
0	0	1	1	1	0	0	0
0	0	0	0	0	0	0	0

Convolution kernel

-1	-2	-1
0	0	0
1	2	1

=

y

?							
	-3	-4	-4	-4	-4	-3	
	-3	-4	-4	-3	-1	0	
	0	0	0	0	0	0	
	2	1	0	1	3	3	
	2	1	0	1	3	3	
	1	3	4	3	1	0	

Convolution vs cross-correlation

Convolution $y[m, n] = x \circ h = \sum_{k, l=-N}^N x[m - k, n - l] h[k, l]$

Cross-correlation $y[m, n] = x * h = \sum_{k, l=-N}^N x[m + k, n + l] h[k, l]$

In the convolution, the kernel h is inverted left-right and up-down, while in the cross-correlation is not

Convolution

Cross-correlation

Kernel





-1	-2	-1
0	0	0
1	2	1

0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	
0	1	2	1	0	0	1	1	0
0	0	0	0	1	1	1	1	0
0	-1	-2	-1	1	1	1	1	0
0	1	1	1	1	1	1	0	


0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	
0	-1	-2	-1	0	0	1	1	0
0	0	0	0	1	1	1	1	0
0	1	2	1	1	1	1	1	0
0	1	1	1	1	1	1	0	

Quiz: what operation is the result of a convolution?


Quiz: what operation is the result of a convolution?



Examples

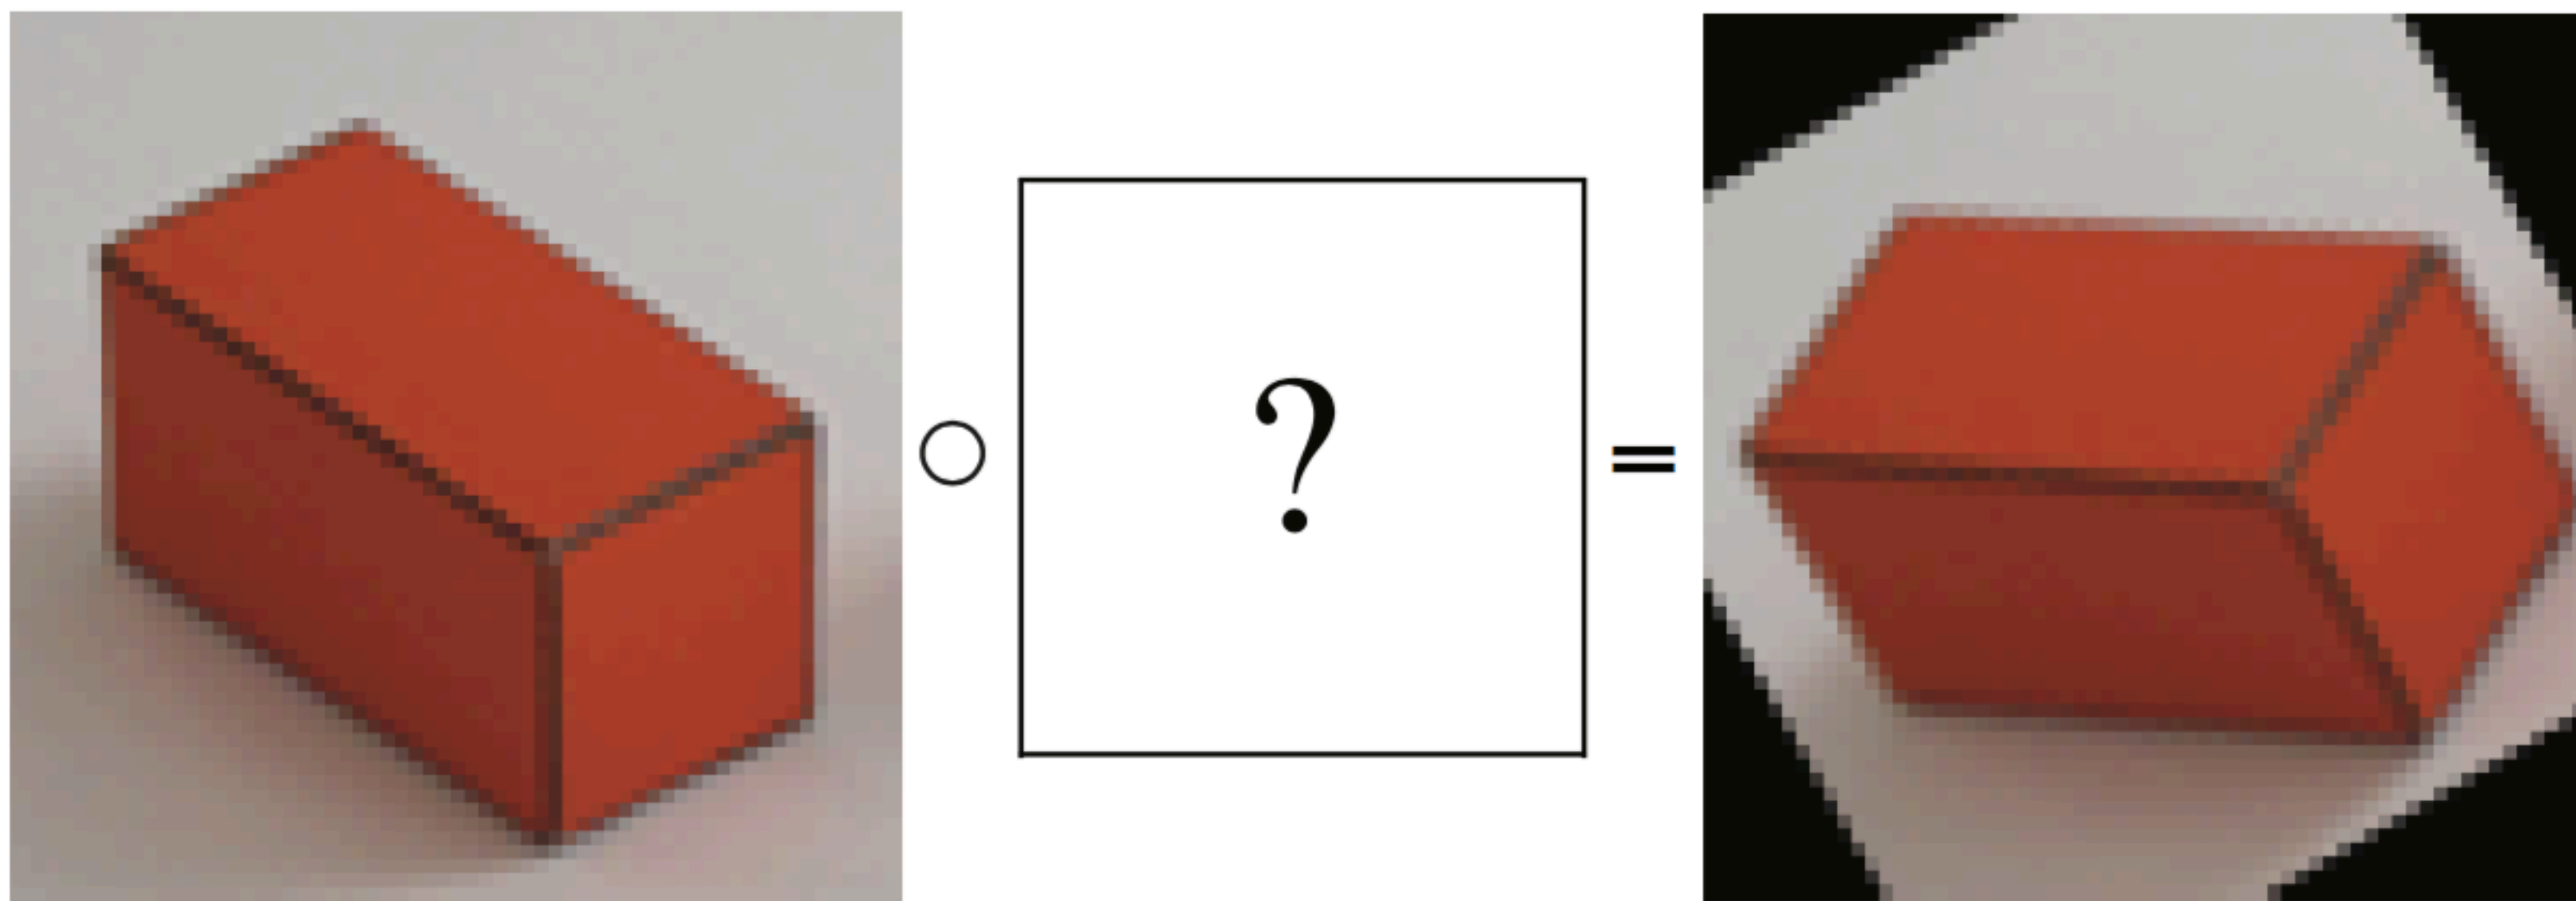


Defocus/blurring





Computes the local average over windows of size 5 x 5 pixels



Examples





Examples





Examples




Examples




Examples



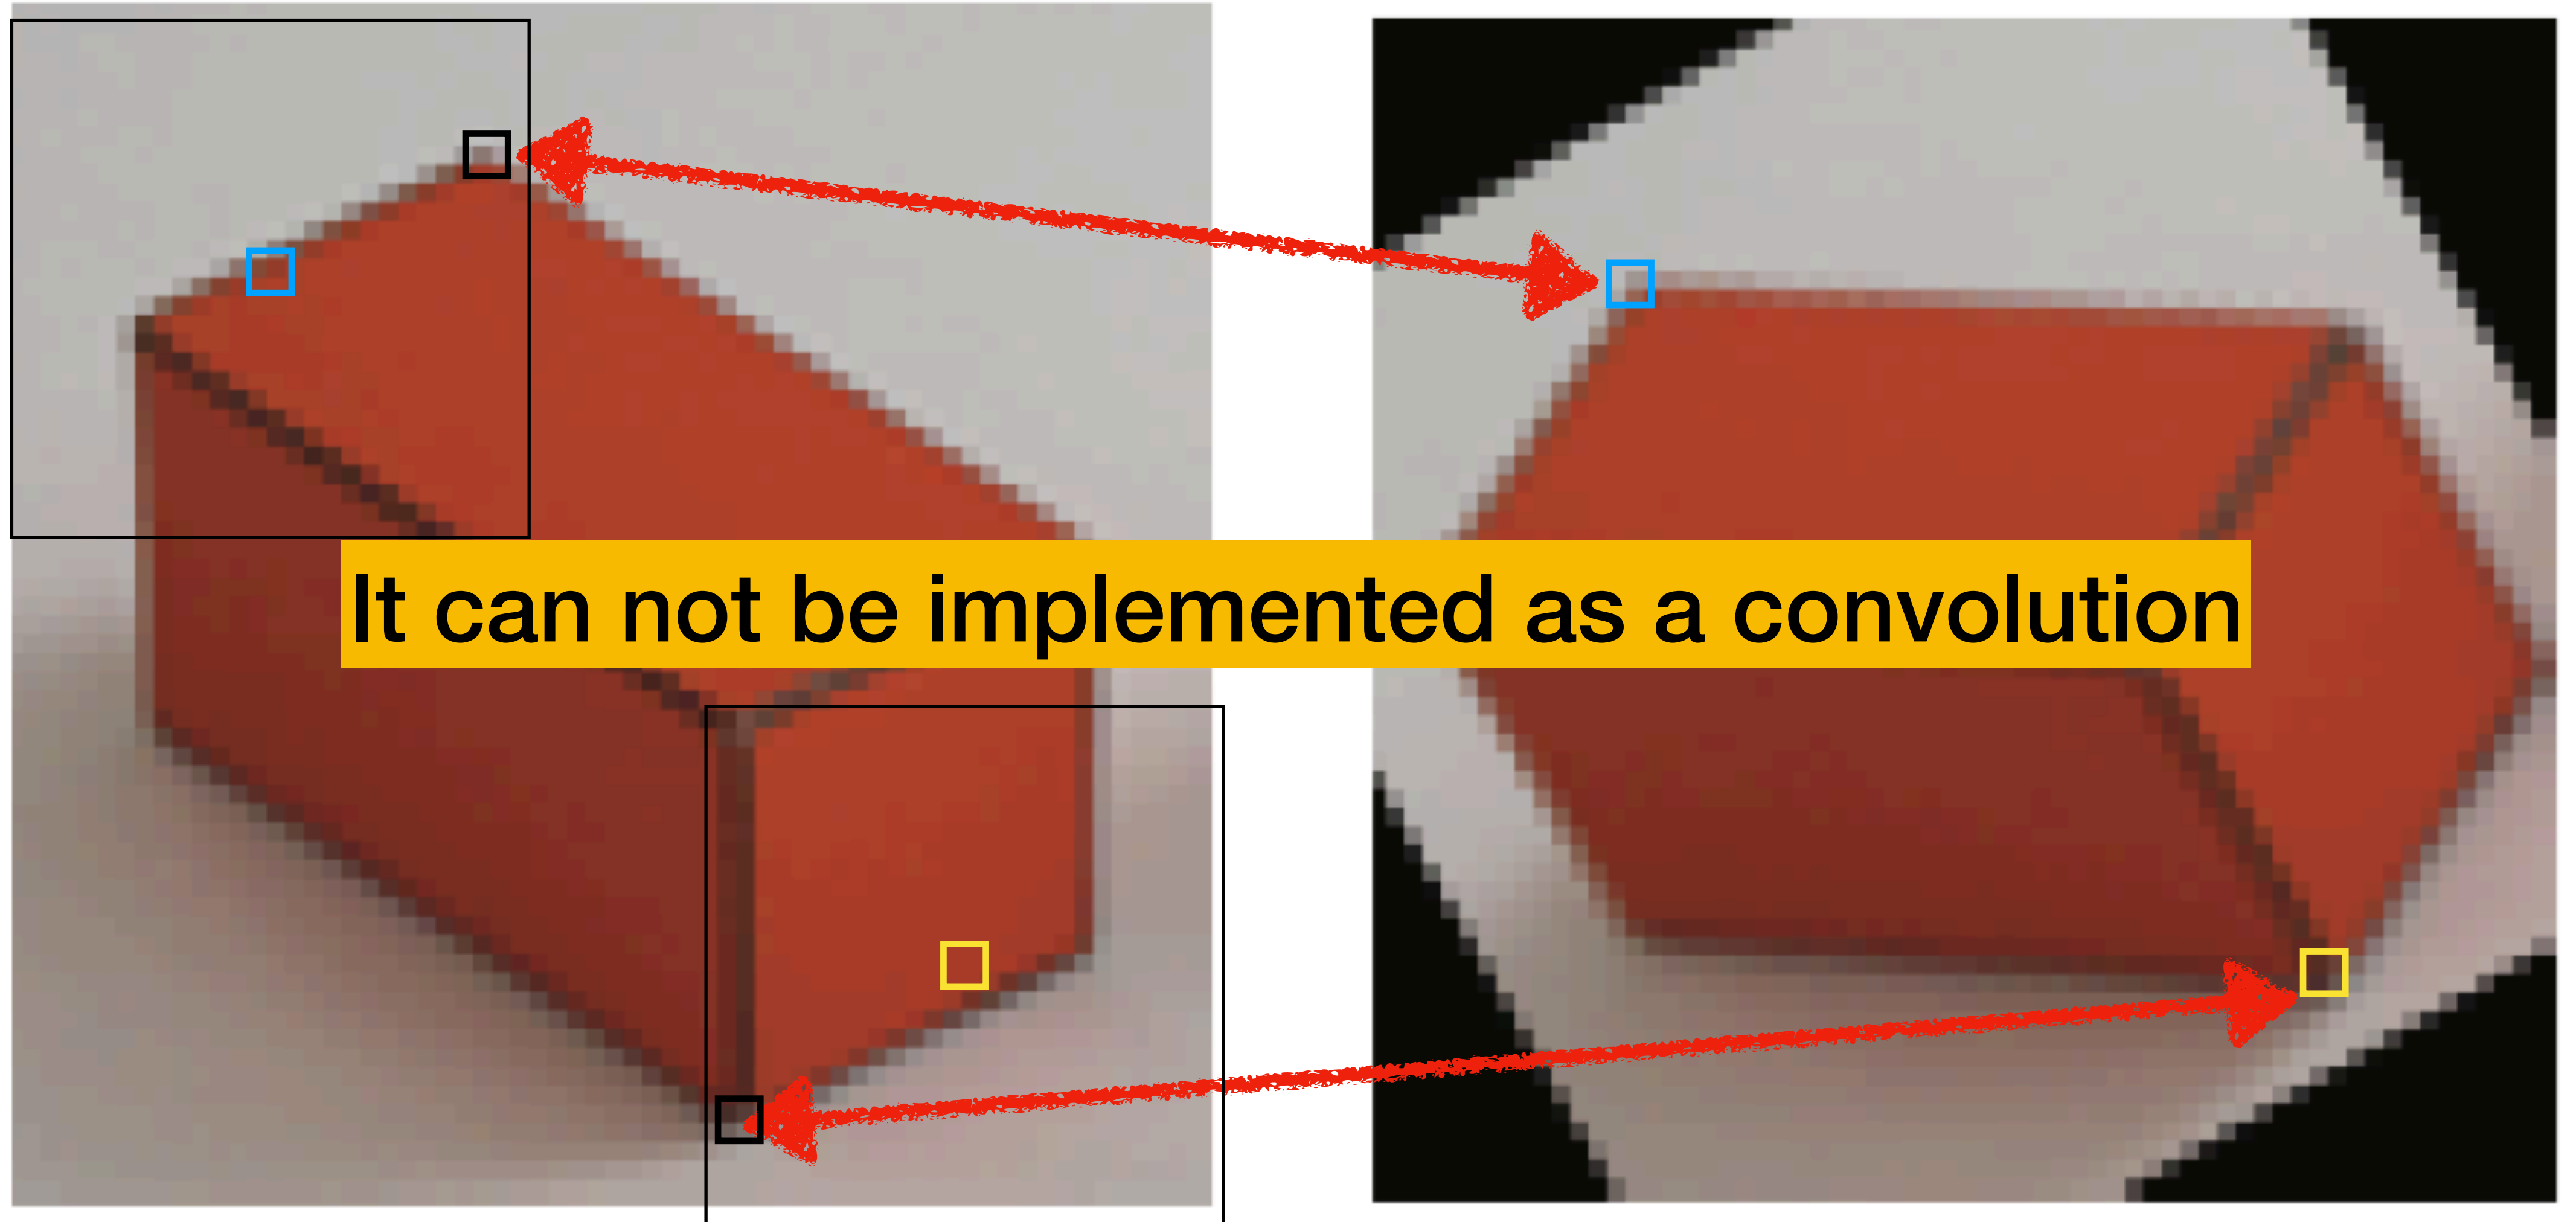
Examples



Examples



Examples



Convolution as matrix multiplication

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \circ \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 & 0 & 0 \end{bmatrix}$$

Convolution as matrix multiplication

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \circ \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 & 0 & 0 \end{bmatrix}$$

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 \\ 30 \\ 60 \\ 90 \\ 90 \\ 90 \\ 90 \\ 60 \\ 30 \end{bmatrix} = \begin{bmatrix} & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 90 \\ 90 \\ 90 \\ 90 \\ 90 \\ 0 \\ 0 \end{bmatrix}$$

Convolution as matrix multiplication

In the 1D case, it helps to make explicit the structure of the matrix:

A diagram illustrating 1D convolution as a dot product of two vectors. The first vector has 10 elements: three 0s, five 90s, and two 0s. The second vector has three 1/3s. The result vector has 10 elements: 0, 30, 60, 90, 90, 90, 60, 30, and two empty cells.

$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} & 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 & \end{bmatrix}$$

In the 1D case, it helps to make explicit the structure of the matrix:

A diagram illustrating 1D convolution as matrix multiplication. A vertical vector of values [0, 30, 60, 90, 90, 90, 90, 60, 30] is equal to a 9x10 grid matrix multiplied by a vertical vector of values [0, 0, 0, 90, 90, 90, 90, 90, 0, 0]. The first row of the matrix has 1/3 in the first three columns, and the rest are empty.

$$\begin{bmatrix} 0 \\ 30 \\ 60 \\ 90 \\ 90 \\ 90 \\ 90 \\ 60 \\ 30 \end{bmatrix} = \begin{bmatrix} 1/3 & 1/3 & 1/3 & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 90 \\ 90 \\ 90 \\ 90 \\ 90 \\ 0 \\ 0 \end{bmatrix}$$

Convolution as matrix multiplication

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \circ \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} & 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 & \end{bmatrix}$$

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 \\ 30 \\ 60 \\ 90 \\ 90 \\ 90 \\ 90 \\ 60 \\ 30 \end{bmatrix} = \begin{bmatrix} 1/3 & 1/3 & 1/3 & & & & & & & \\ & & 1/3 & 1/3 & 1/3 & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 90 \\ 90 \\ 90 \\ 90 \\ 90 \\ 0 \\ 0 \end{bmatrix}$$

Convolution as matrix multiplication

In the 1D case, it helps to make explicit the structure of the matrix:


$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \circ \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 \end{bmatrix}$$

In the 1D case, it helps to make explicit the structure of the matrix:

$$\begin{bmatrix} 0 \\ 30 \\ 60 \\ 90 \\ 90 \\ 90 \\ 90 \\ 60 \\ 30 \end{bmatrix} = \begin{bmatrix} 1/3 & 1/3 & 1/3 & & & & & & & \\ & 1/3 & 1/3 & 1/3 & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \\ & & & & & & & & & \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 90 \\ 90 \\ 90 \\ 90 \\ 90 \\ 0 \\ 0 \end{bmatrix}$$


Linear translation invariant system:

A LTI function f can be written as a matrix multiplication:



$h[n - k]$ n indexes rows,
 k indexes columns

It can also be represented as a convolutional layer of neural net:



$h[n - k]$ is the strength of the connection
between $x[k]$ and $y[n]$

$$y[n] = \sum_{k=-1}^1 h[k] x[n - k]$$

Rectangular filter



$g[m,n]$

\otimes



$h[m,n]$

=



$f[m,n]$

Rectangular filter



$g[m,n]$

\otimes



=

$h[m,n]$




$f[m,n]$

Rectangular filter



$g[m,n]$

\otimes

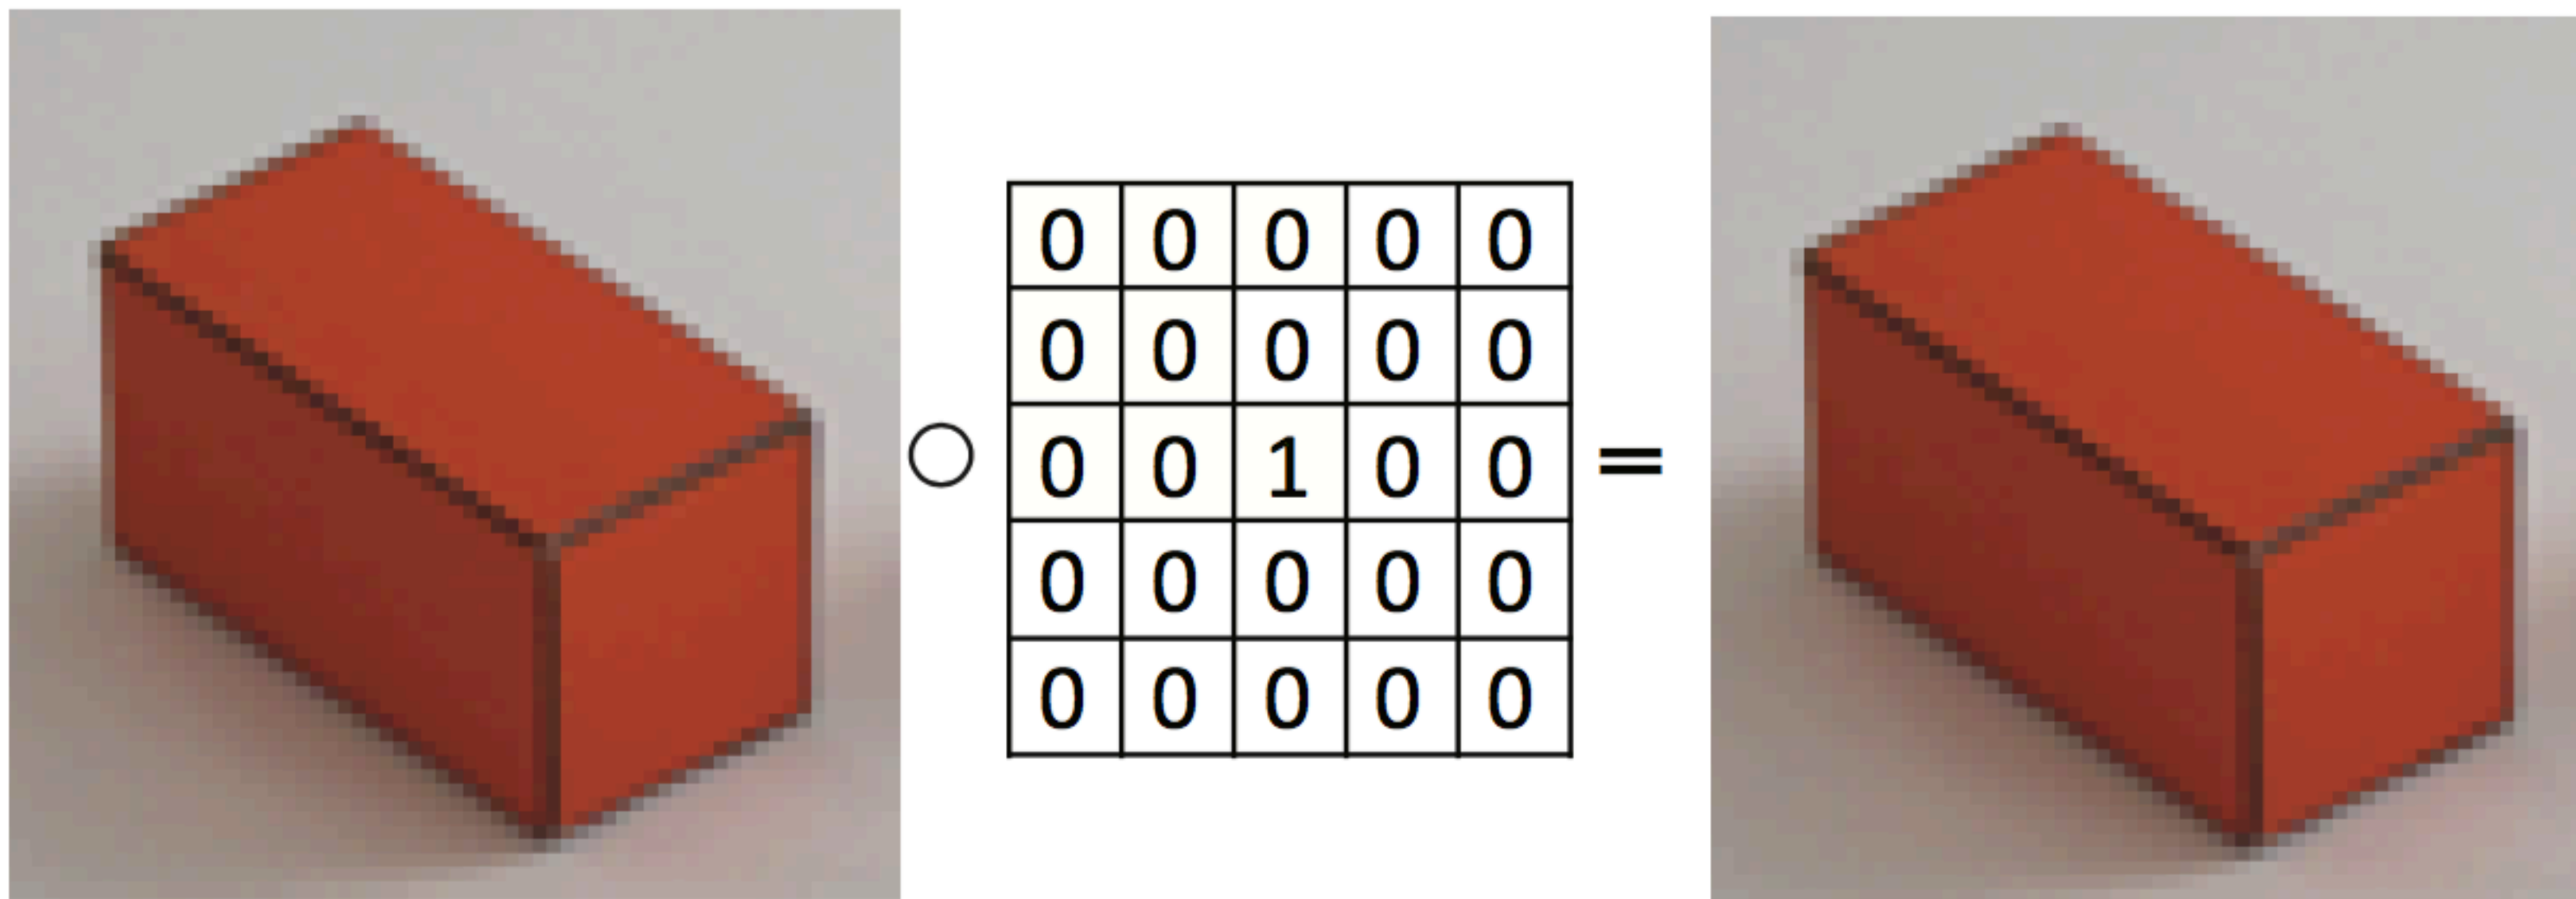


=




$f[m,n]$

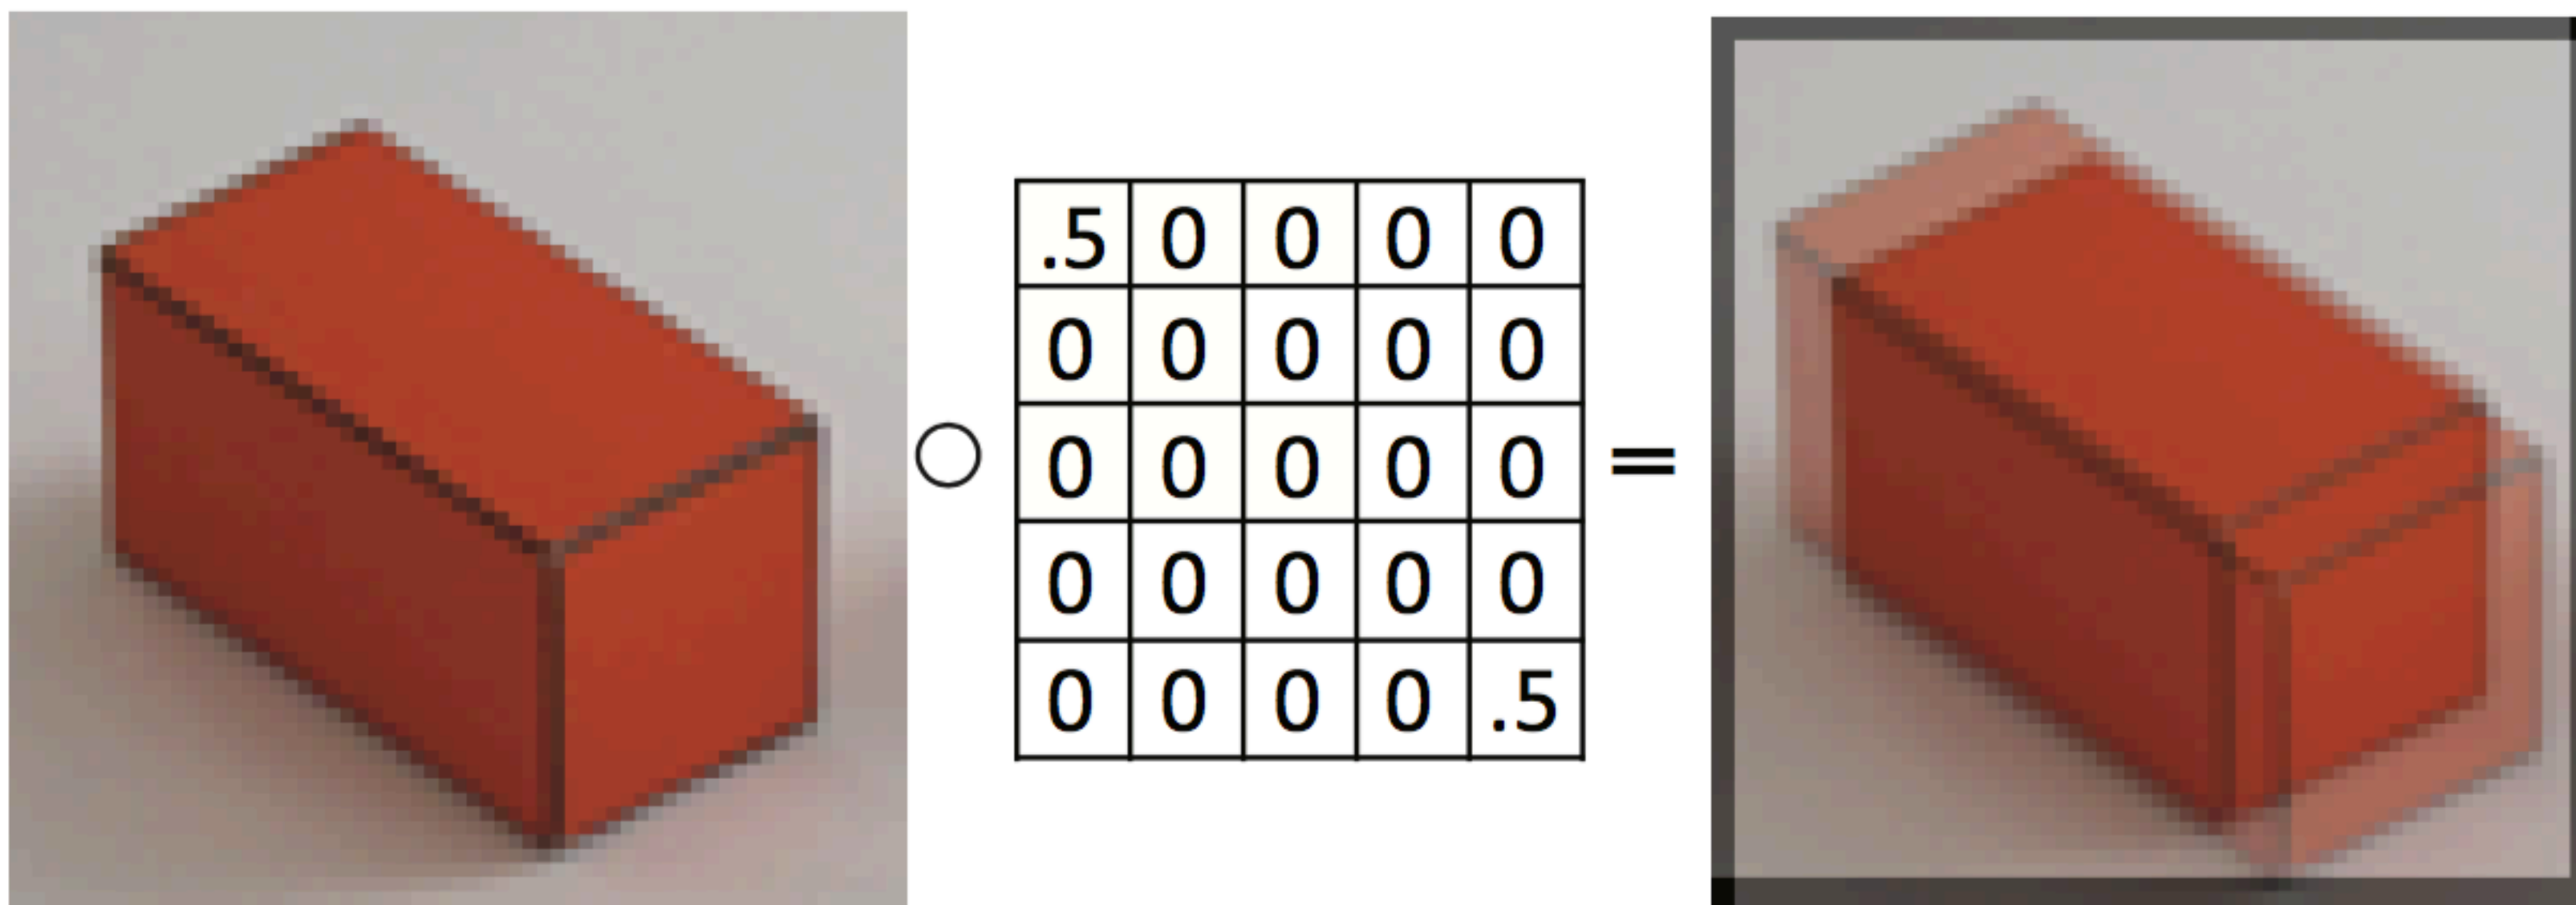
The identity



A shift



Examples


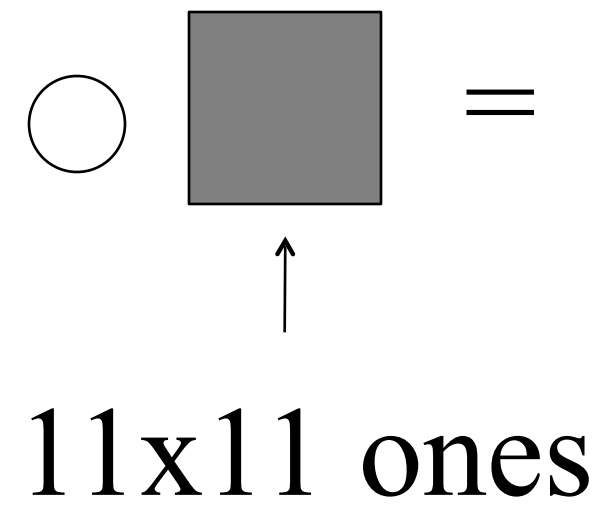



Handling boundaries



Handling boundaries

Zero padding



Handling boundaries

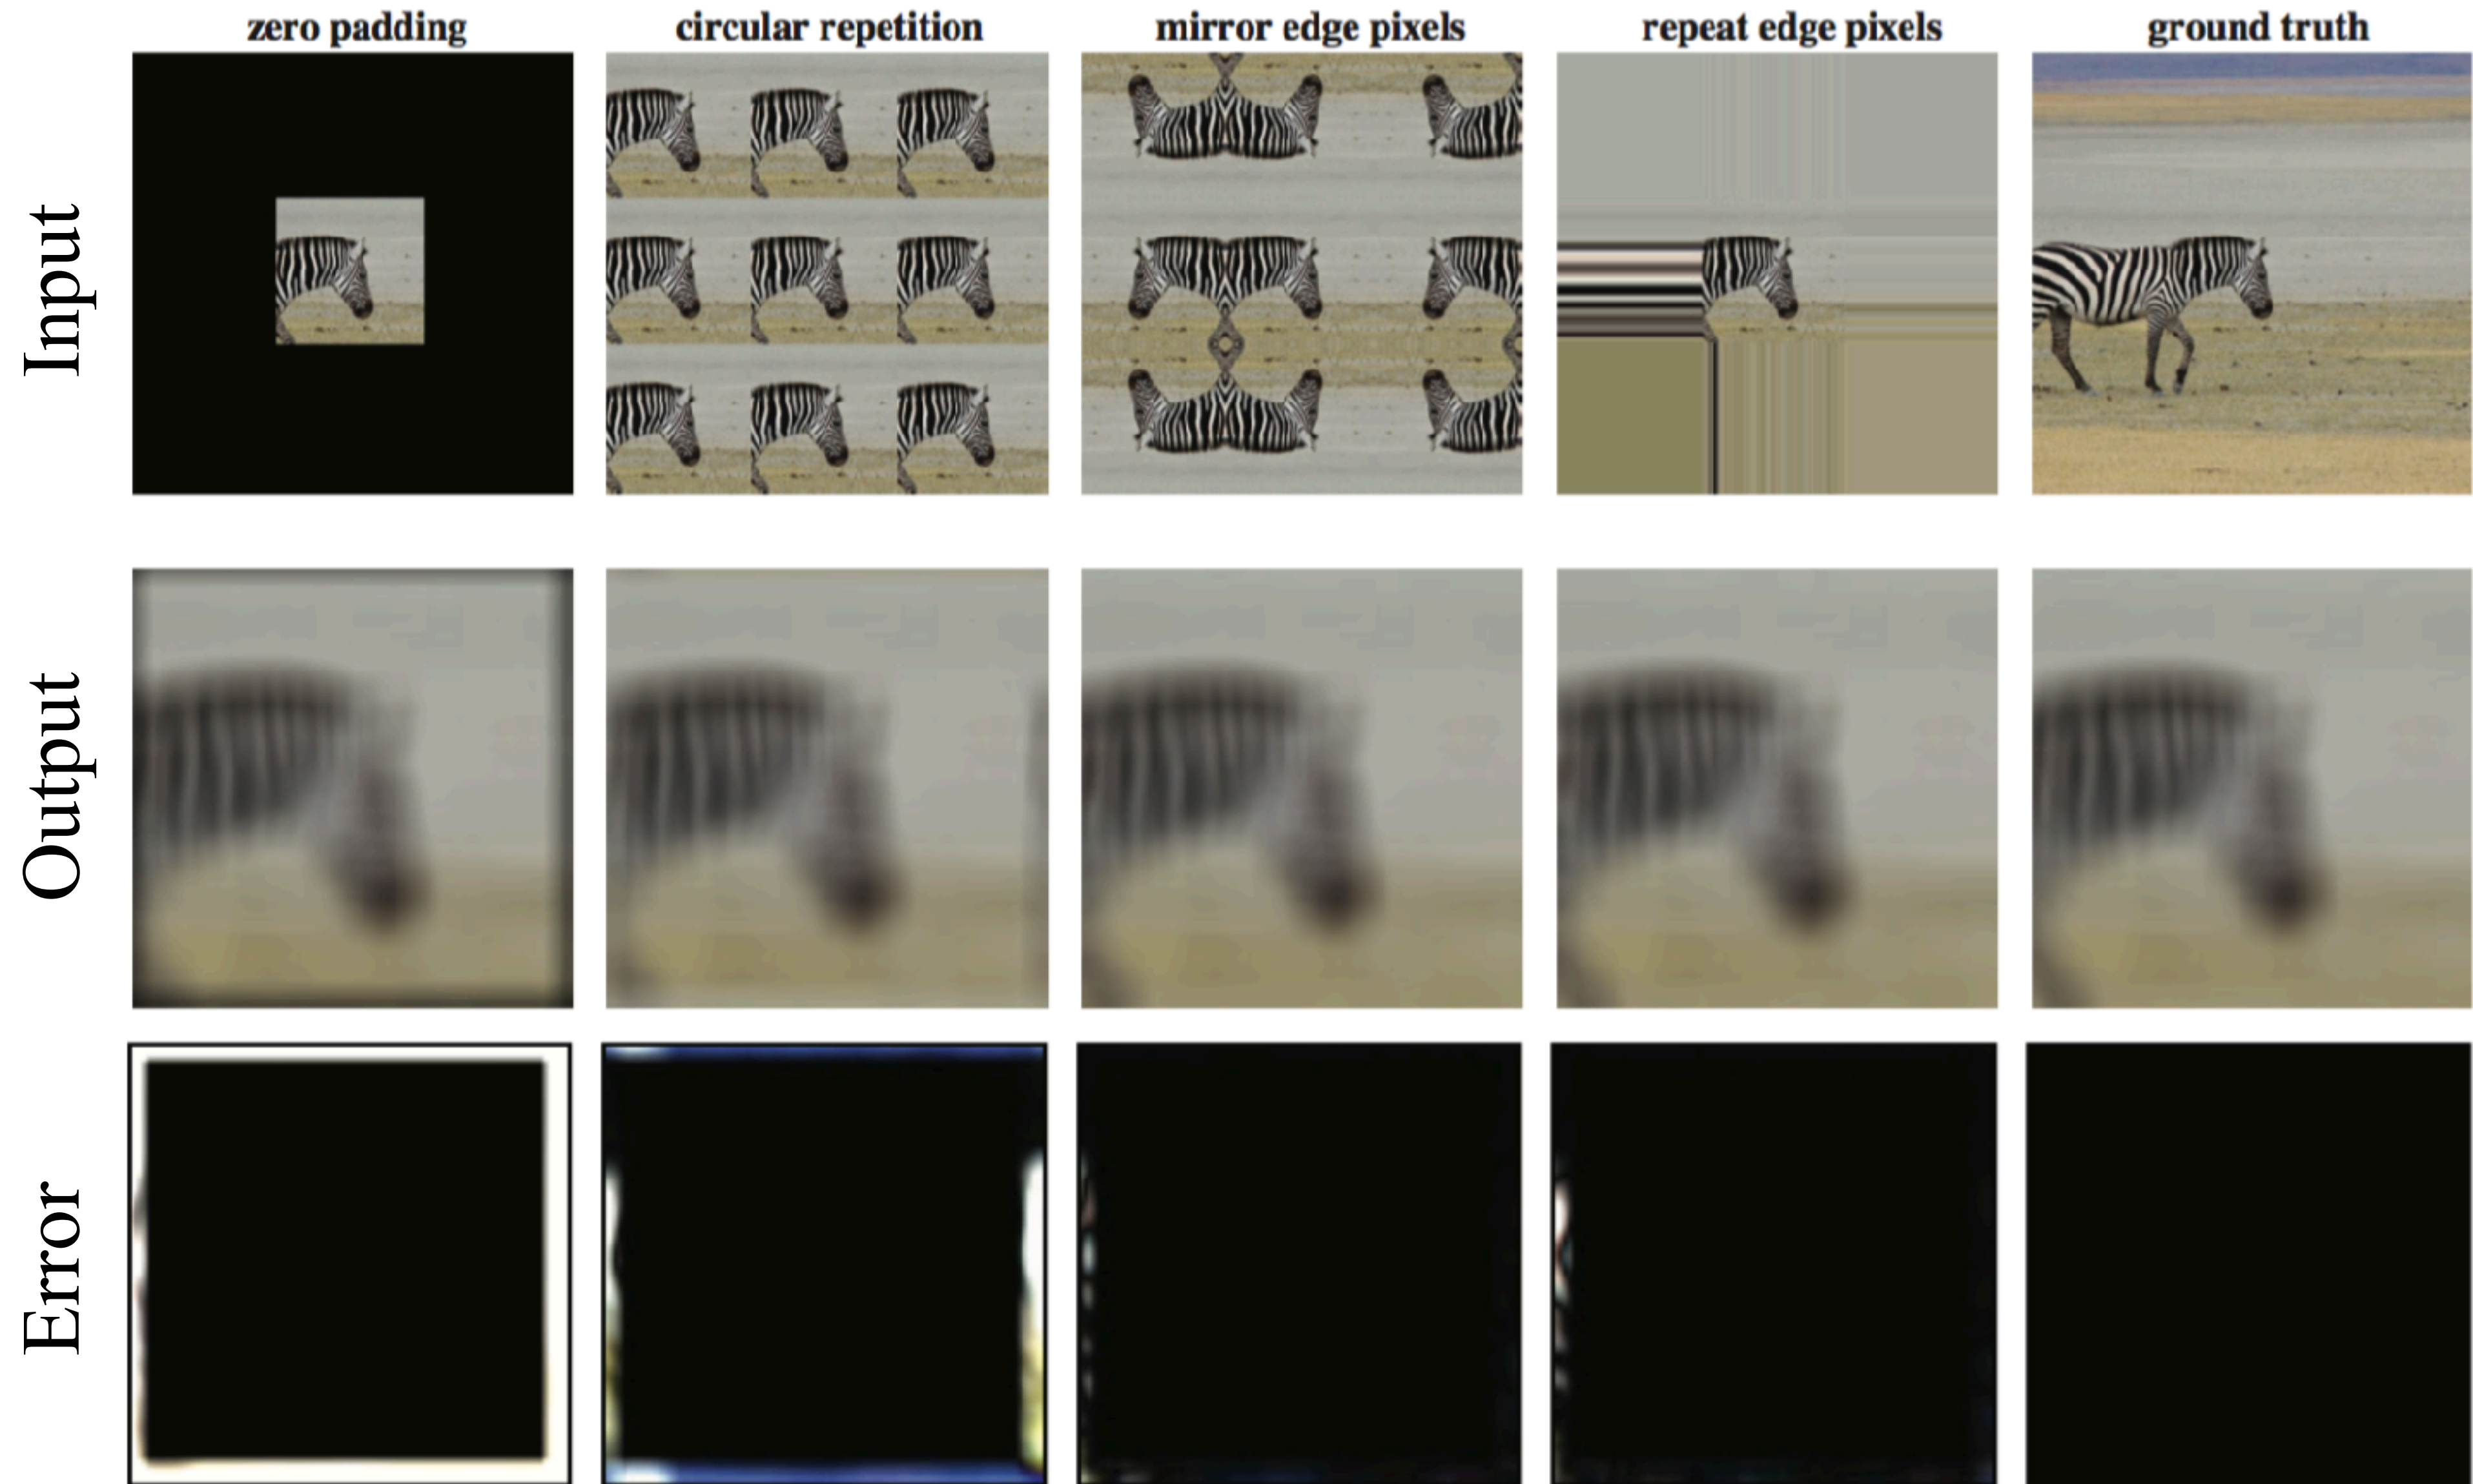

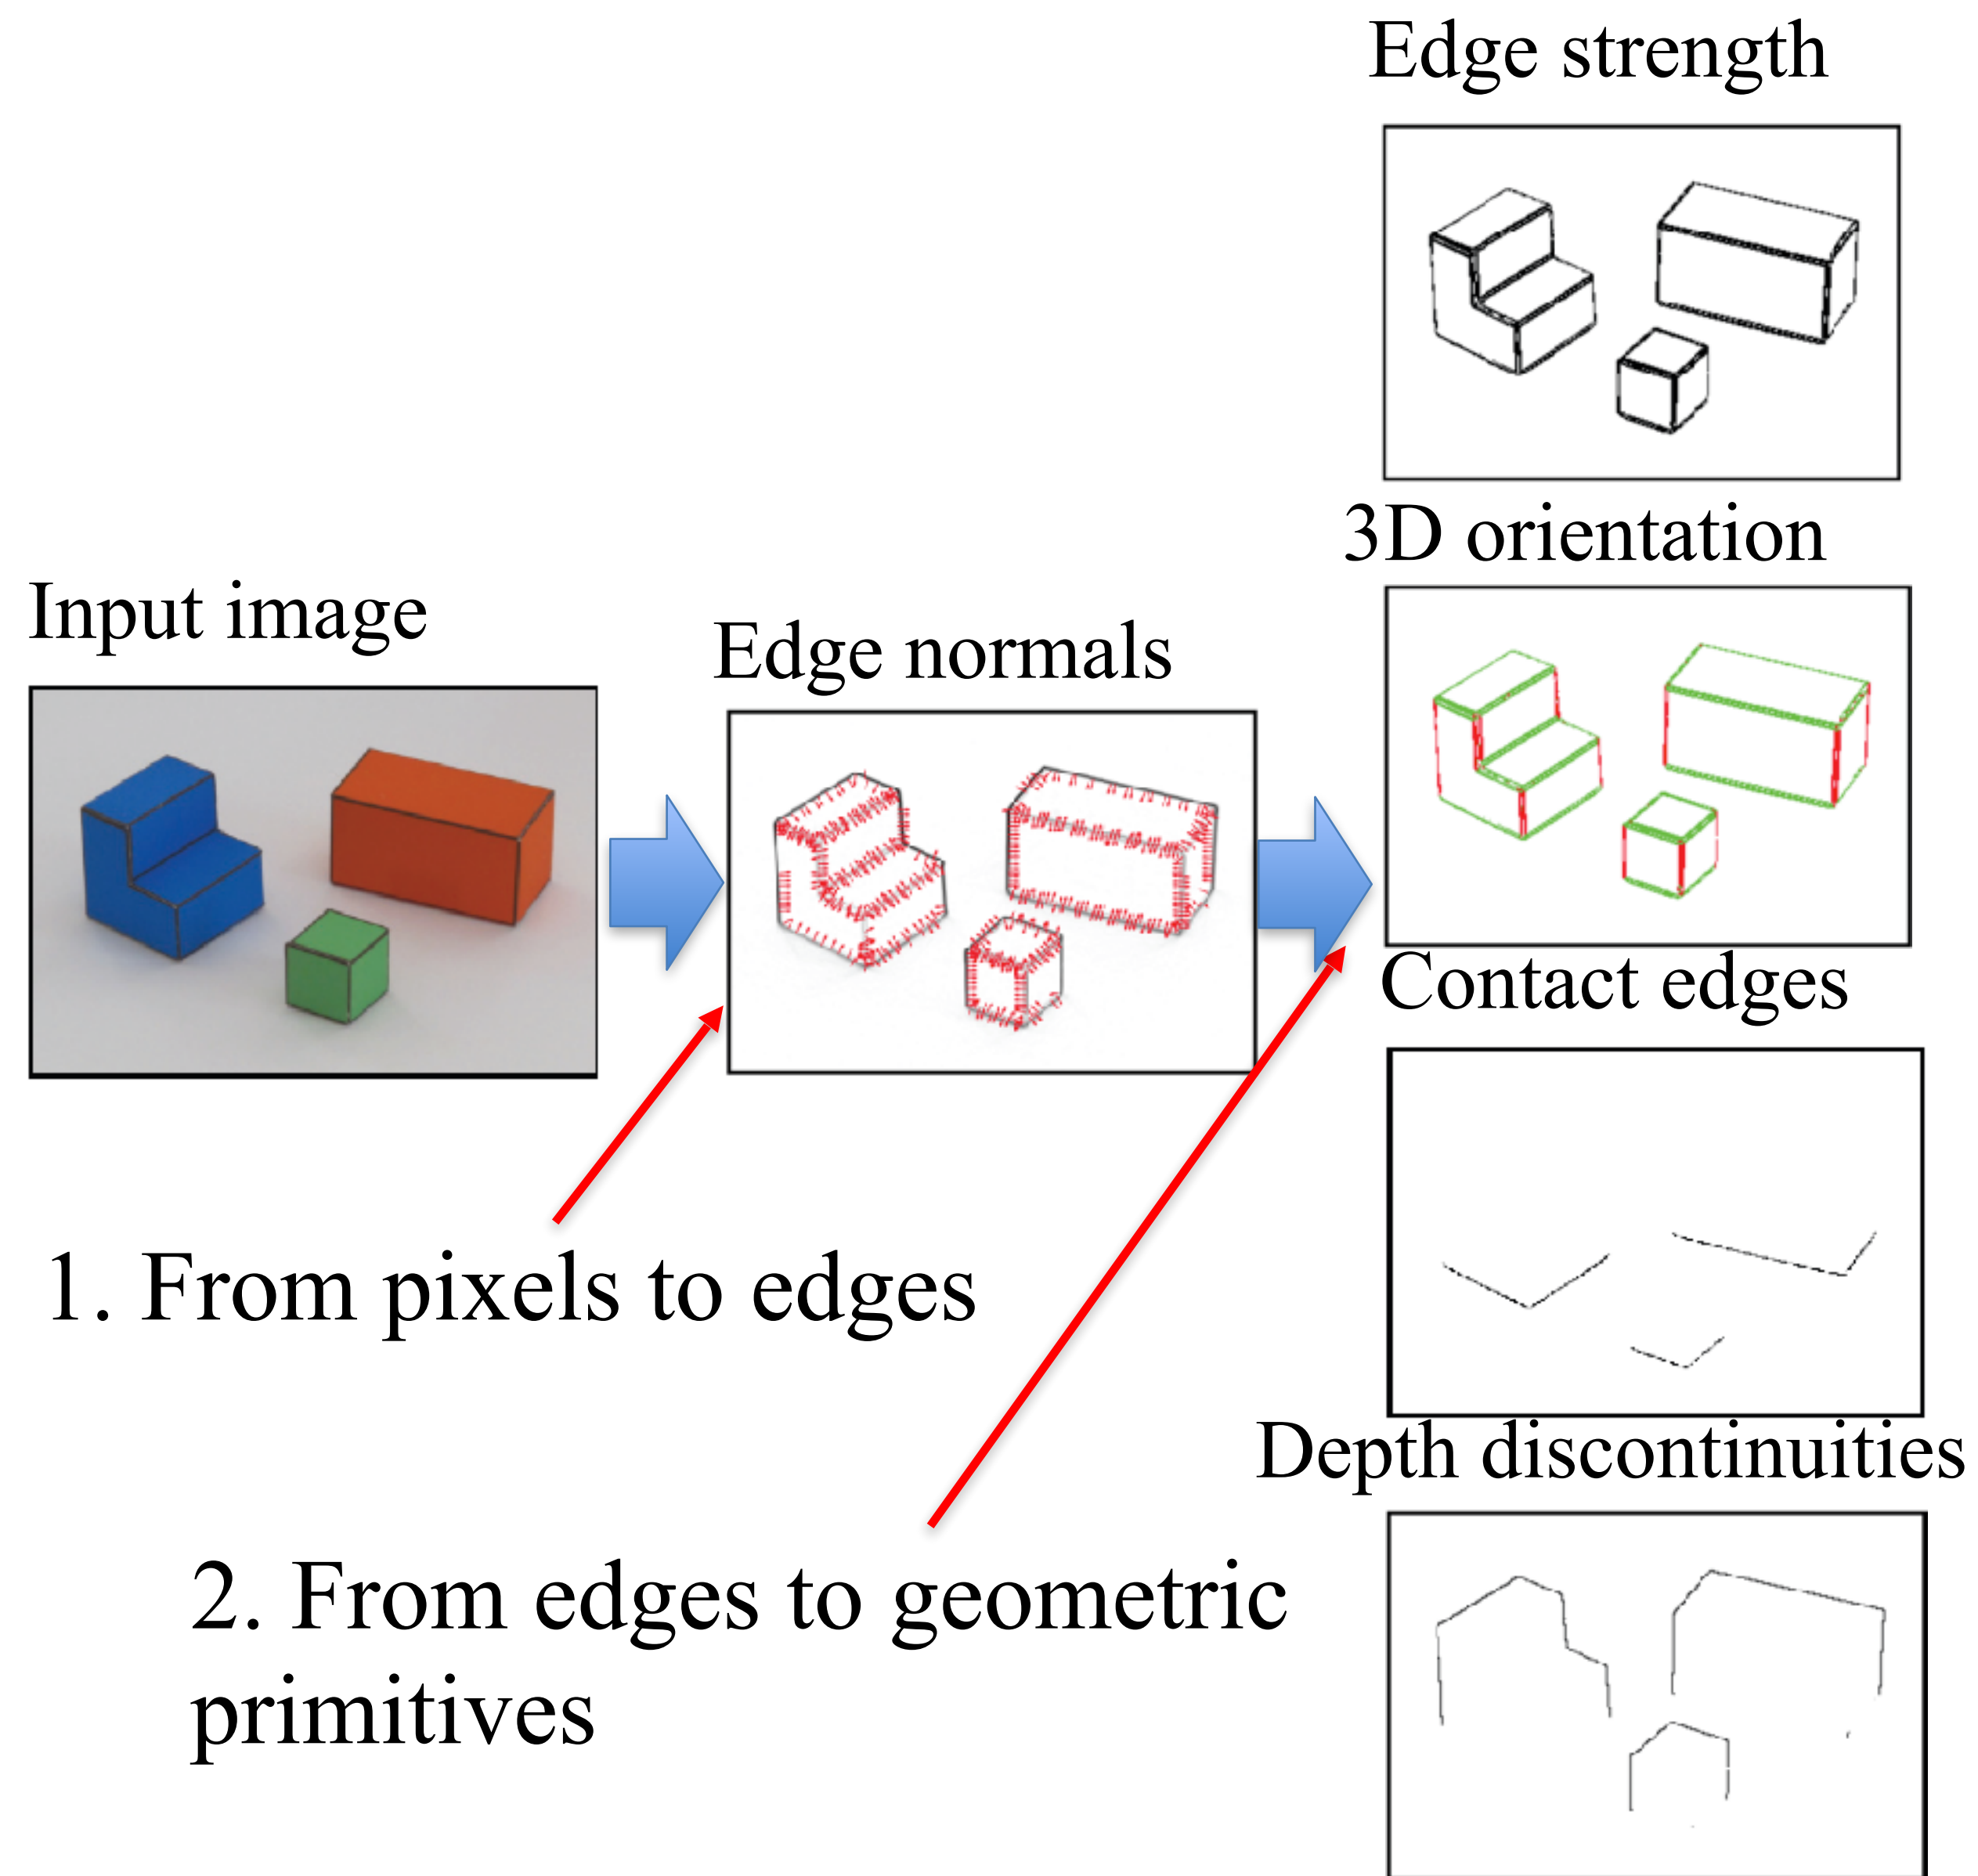
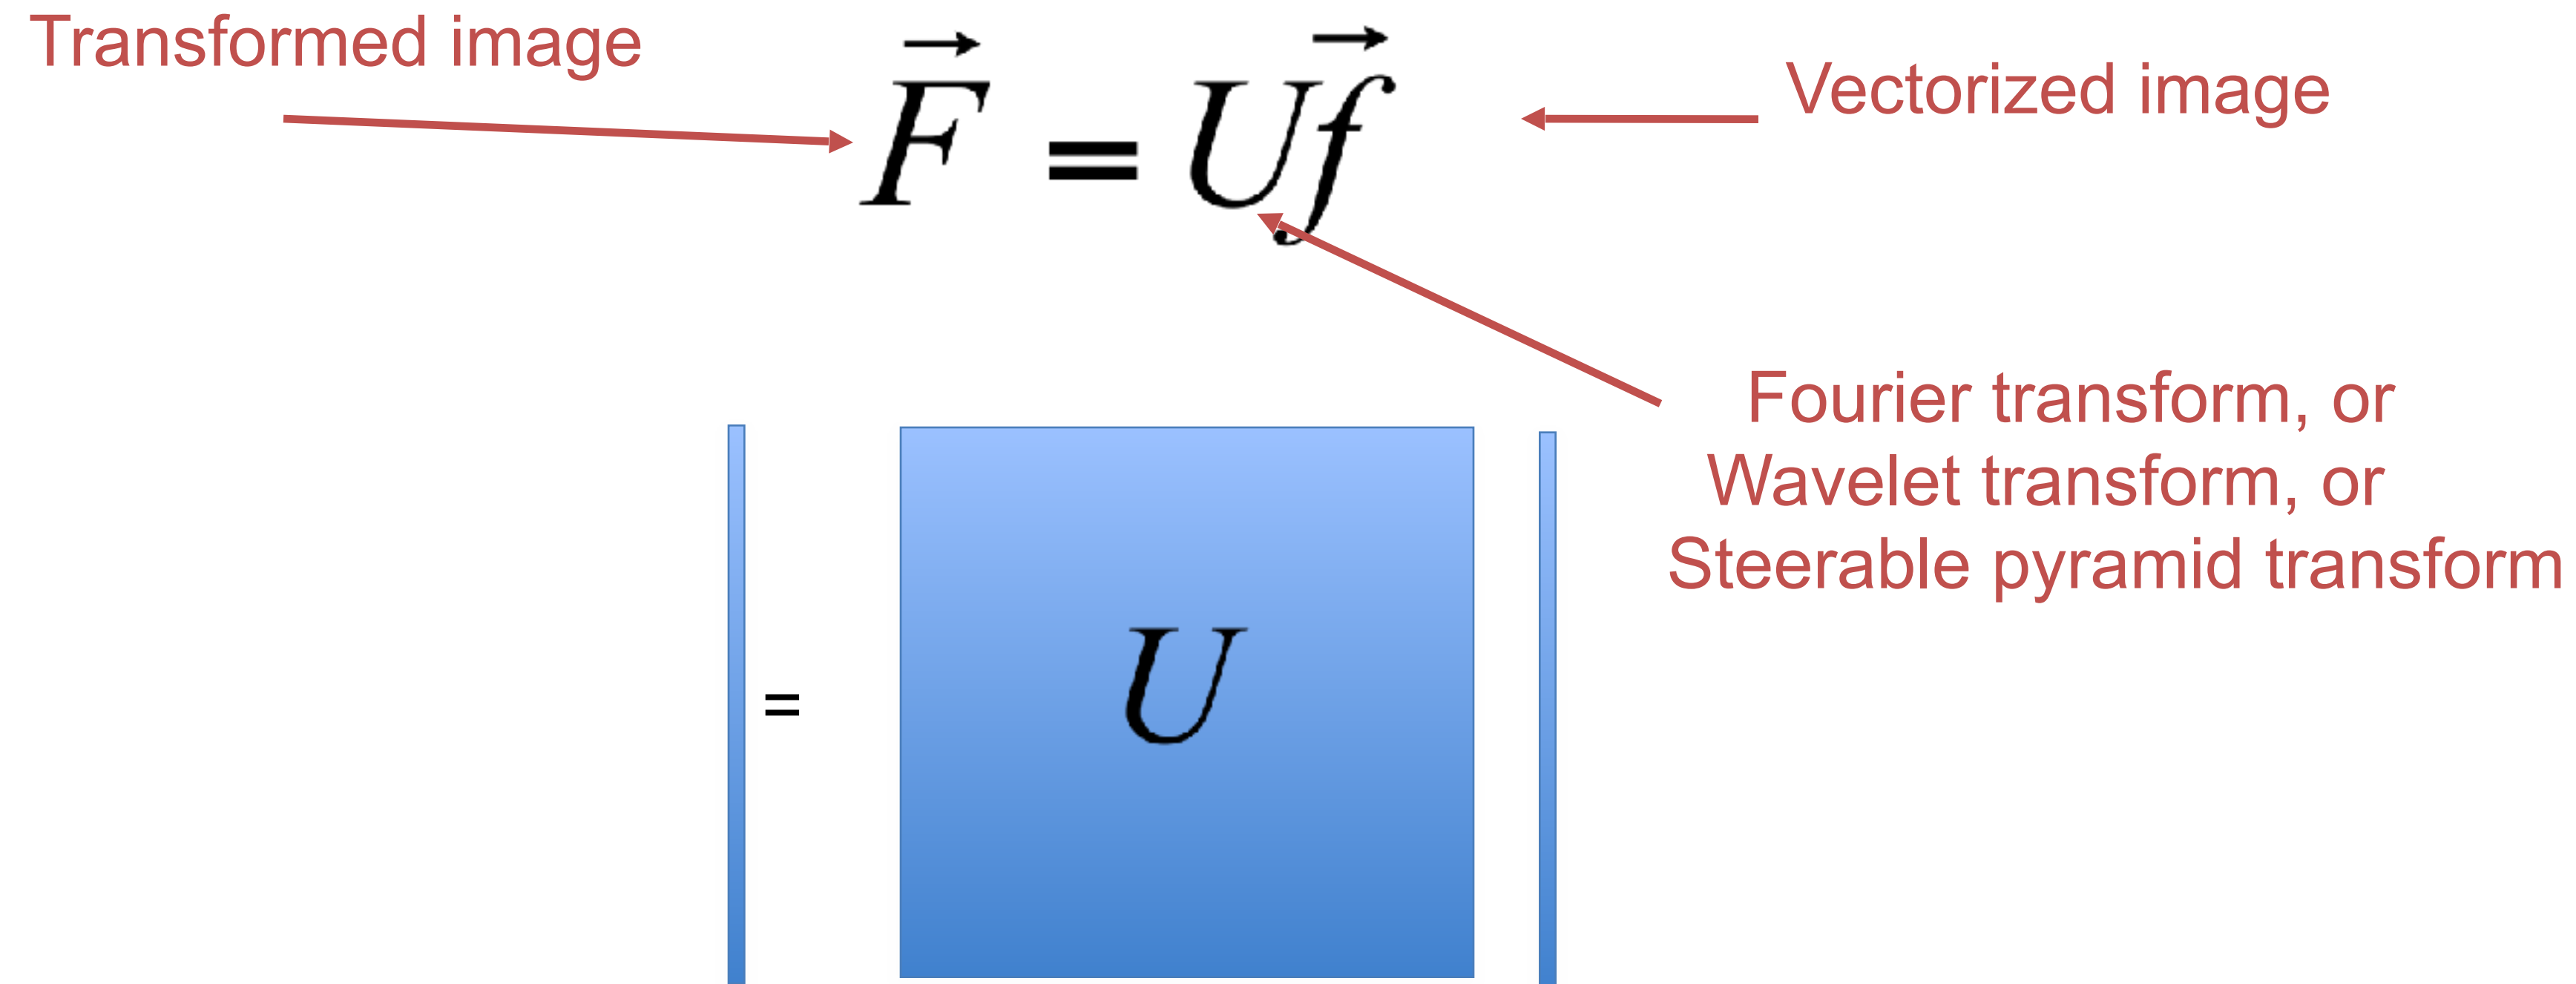


Image transformations



Linear image transformations


In analyzing images, it's often useful to make a change of basis.



Fourier series


$$f(t) = a_1 \sin(t) + a_2 \sin(2t) + a_3 \sin(3t) + \dots \quad \text{With } a_n = \frac{2}{\pi} \int_0^{\pi} f(t) \sin(nt) dt$$

One of Fourier's original examples of sine series is the expansion of the ramp signal:



$$\frac{1}{2}t = \sin(t) - \frac{1}{2} \sin(2t) + \frac{1}{3} \sin(3t) - \frac{1}{4} \sin(4t) + \dots$$

The result of this series approximates the ramp with increasing accuracy as we add more terms.



THÉORIE
ANALYTIQUE
DE LA CHALEUR,

PAR M. FOURIER.





CHEZ FIRMIN DIDOT, PÈRE ET FILS,
LIBRAIRES POUR LES MATHÉMATIQUES, L'ARCHITECTURE HYDRAULIQUE
ET LA MARINE, RUE JACOB, N° 24.

1822.

Fourier series as change of representation

$$\frac{1}{2}t = \sin(t) - \frac{1}{2}\sin(2t) + \frac{1}{3}\sin(3t) - \frac{1}{4}\sin(4t) + \dots$$



It is useful to think of the Fourier series of a signal as a **change of representation**. Instead of representing the signal by the sequence of values specified by the original function $f(t)$, the same function can be represented by the infinite sequence of coefficients (a_n).

The Discrete Fourier transform

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

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

The inverse of the DFT is:

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

The signal $f[n]$ is a weighted linear combination of complex exponentials with weights $F[u]$

The Discrete Fourier transform

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

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

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

$$F = \begin{matrix} \text{u} \downarrow \\ \begin{matrix} \text{n} \rightarrow \\ ? & ? & ? & ? & ? & ? & ? & ? & \dots & ? \end{matrix} \\ \exp\left(-2\pi j \frac{un}{N}\right) \end{matrix} f$$

NxN array

Lets visualize the transform coefficients


Visualizing the Fourier transform

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

$$\exp(\alpha j) = \cos(\alpha) + j \sin(\alpha)$$



$$\cos\left(2\pi \frac{un}{N}\right) - j \sin\left(2\pi \frac{un}{N}\right)$$

For:
 $u=1$
 $N=16$



Visualizing the transform coefficients

$$\exp\left(-2\pi j \frac{un}{N}\right) \quad \text{For } N=16$$





The inverse of the Discrete Fourier transform

Discrete Fourier Transform (DFT):

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

Its inverse:

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



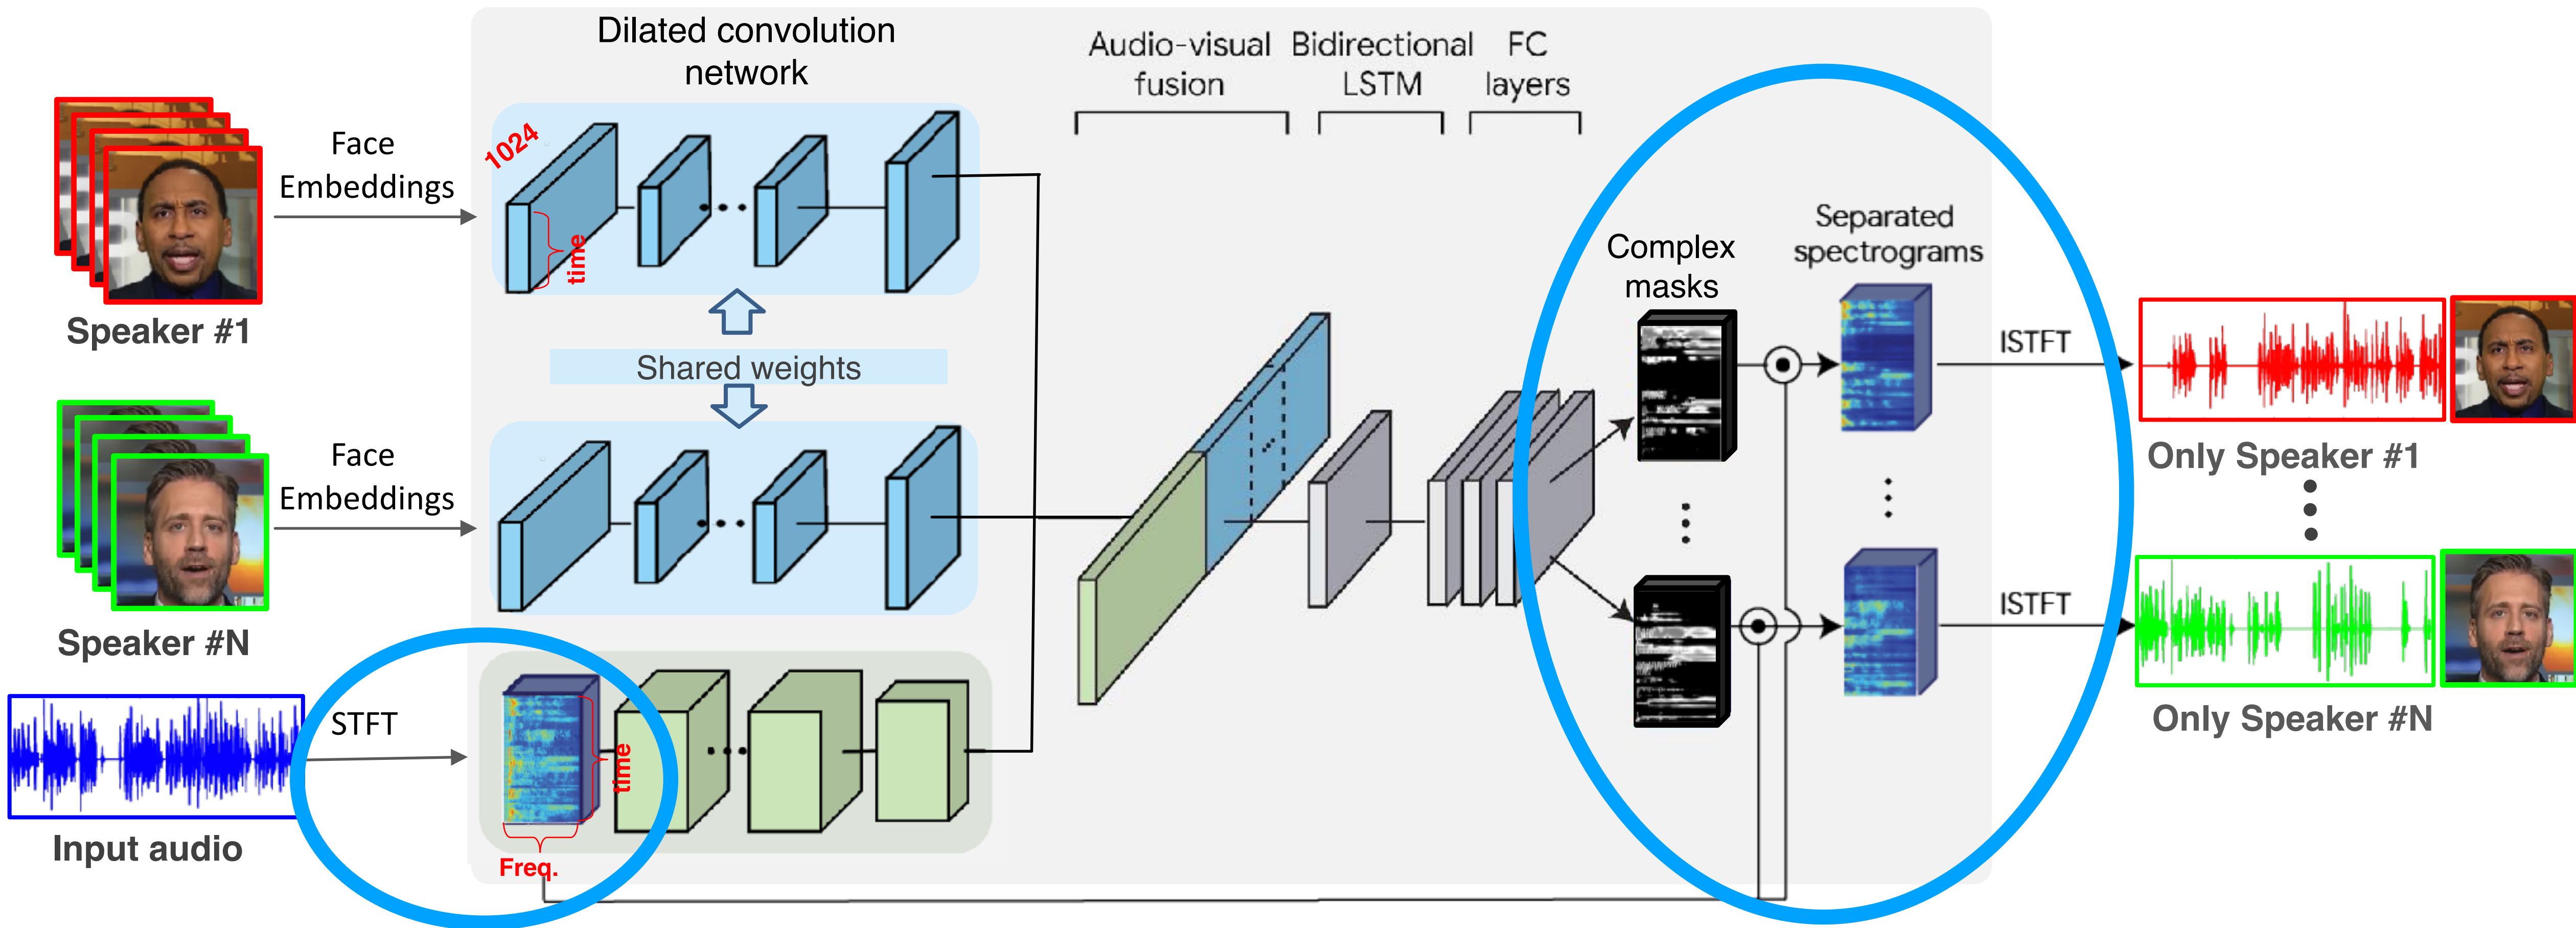
Google technology: Looking to Listen

uses Fourier transform representation for audio (overlapping short-time Fourier transforms of audio signal)

Input video



Teaser for neural network algorithms we'll discuss later in class:
"Looking to Listen" processes audio in the spectral domain to separate speakers



Looking to Listen processing result

Only the driver

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:

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


Visualizing the 2D DFT coefficients

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

For N=M=16



A remarkable property of Fourier transform



In the spatial domain, the output of f is the convolution:

$$f[m,n] = h \circ g = \sum_{k,l} h[m-k, n-l] g[k,l]$$

In the frequency domain:

$$F[u,v] = G[u,v] H[u,v]$$

Terminology:

Impulse response: $h[m,n]$

Transfer function: $H[u,v]$

Dual convolution property

The Fourier transform of the convolution is the product of Fourier transforms


$$f[m, n] = h \circ g$$



$$F[u, v] = G[u, v] H[u, v]$$

The Fourier transform of the product is the convolution of Fourier transforms

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



$$F[u, v] = \frac{1}{NM} G[u, v] \circ H[u, v]$$

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

Visualizing the image Fourier transform


$f[n, m]$




$F[u, v]$




magnitude



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.

Some important Fourier transforms




Some important Fourier transforms




Some important Fourier transforms and properties


Image



Magnitude DFT






Phase DFT




Translation

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




Some important Fourier transforms and properties


Image



Magnitude DFT




Phase DFT




Scale


Small image details produce content in high spatial frequencies



Magnitude DFT




Phase DFT




Some important Fourier transforms and properties


Image



Magnitude DFT






Phase DFT





Orientation

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




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)

Phase and Magnitude (even though we just finished matching images based on FT magnitudes, we should point out that phase is also very important)

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



Each color channel is processed in the same way.

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:

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


How does summing waves ends up giving back a picture?

2

2




#1: Range [0, 1]
Dims [256, 256]




#2: Range [0.000109, 0.0267]
Dims [256, 256]

6

6




#1: Range [0, 1]
Dims [256, 256]




#2: Range [1.89e-007, 0.226]
Dims [256, 256]

18

18



#1: Range [0, 1]
Dims [256, 256]




#2: Range [4.79e-007, 0.503]
Dims [256, 256]

50

50



#1: Range [0, 1]
Dims [256, 256]




#2: Range [8.5e-006, 1.7]
Dims [256, 256]

82

82




#1: Range [0, 1]
Dims [256, 256]




#2: Range [3.85e-007, 2.21]
Dims [256, 256]

136

136




#1: Range [0, 1]
Dims [256, 256]




#2: Range [8.25e-006, 3.48]
Dims [256, 256]

282

282




#1: Range [0, 1]
Dims [256, 256]



#2: Range [1.39e-005, 5.88]
Dims [256, 256]

538

538




#1: Range [0, 1]
Dims [256, 256]




#2: Range [6.17e-006, 8.4]
Dims [256, 256]

1088

1088




#1: Range [0, 1]
Dims [256, 256]




#2: Range [9.99e-005, 15]
Dims [256, 256]

2094

2094




#1: Range [0, 1]
Dims [256, 256]




#2: Range [8.7e-005, 19]
Dims [256, 256]

4052.

4052




#1: Range [0, 1]
Dims [256, 256]




#2: Range [0.000556, 37.7]
Dims [256, 256]

8056.

8056




#1: Range [0, 1]
Dims [256, 256]



#2: Range [0.00032, 64.5]
Dims [256, 256]

15366

15366




#1: Range [0, 1]
Dims [256, 256]




#2: Range [0.000231, 91.1]
Dims [256, 256]

28743

28743




#1: Range [0, 1]
Dims [256, 256]




#2: Range [0.00109, 146]
Dims [256, 256]

49190.

49190



#1: Range [0, 1]
Dims [256, 256]




#2: Range [0.00758, 294]
Dims [256, 256]

65536.

65536.




#1: Range [0.5, 1.5]
Dims [256, 256]



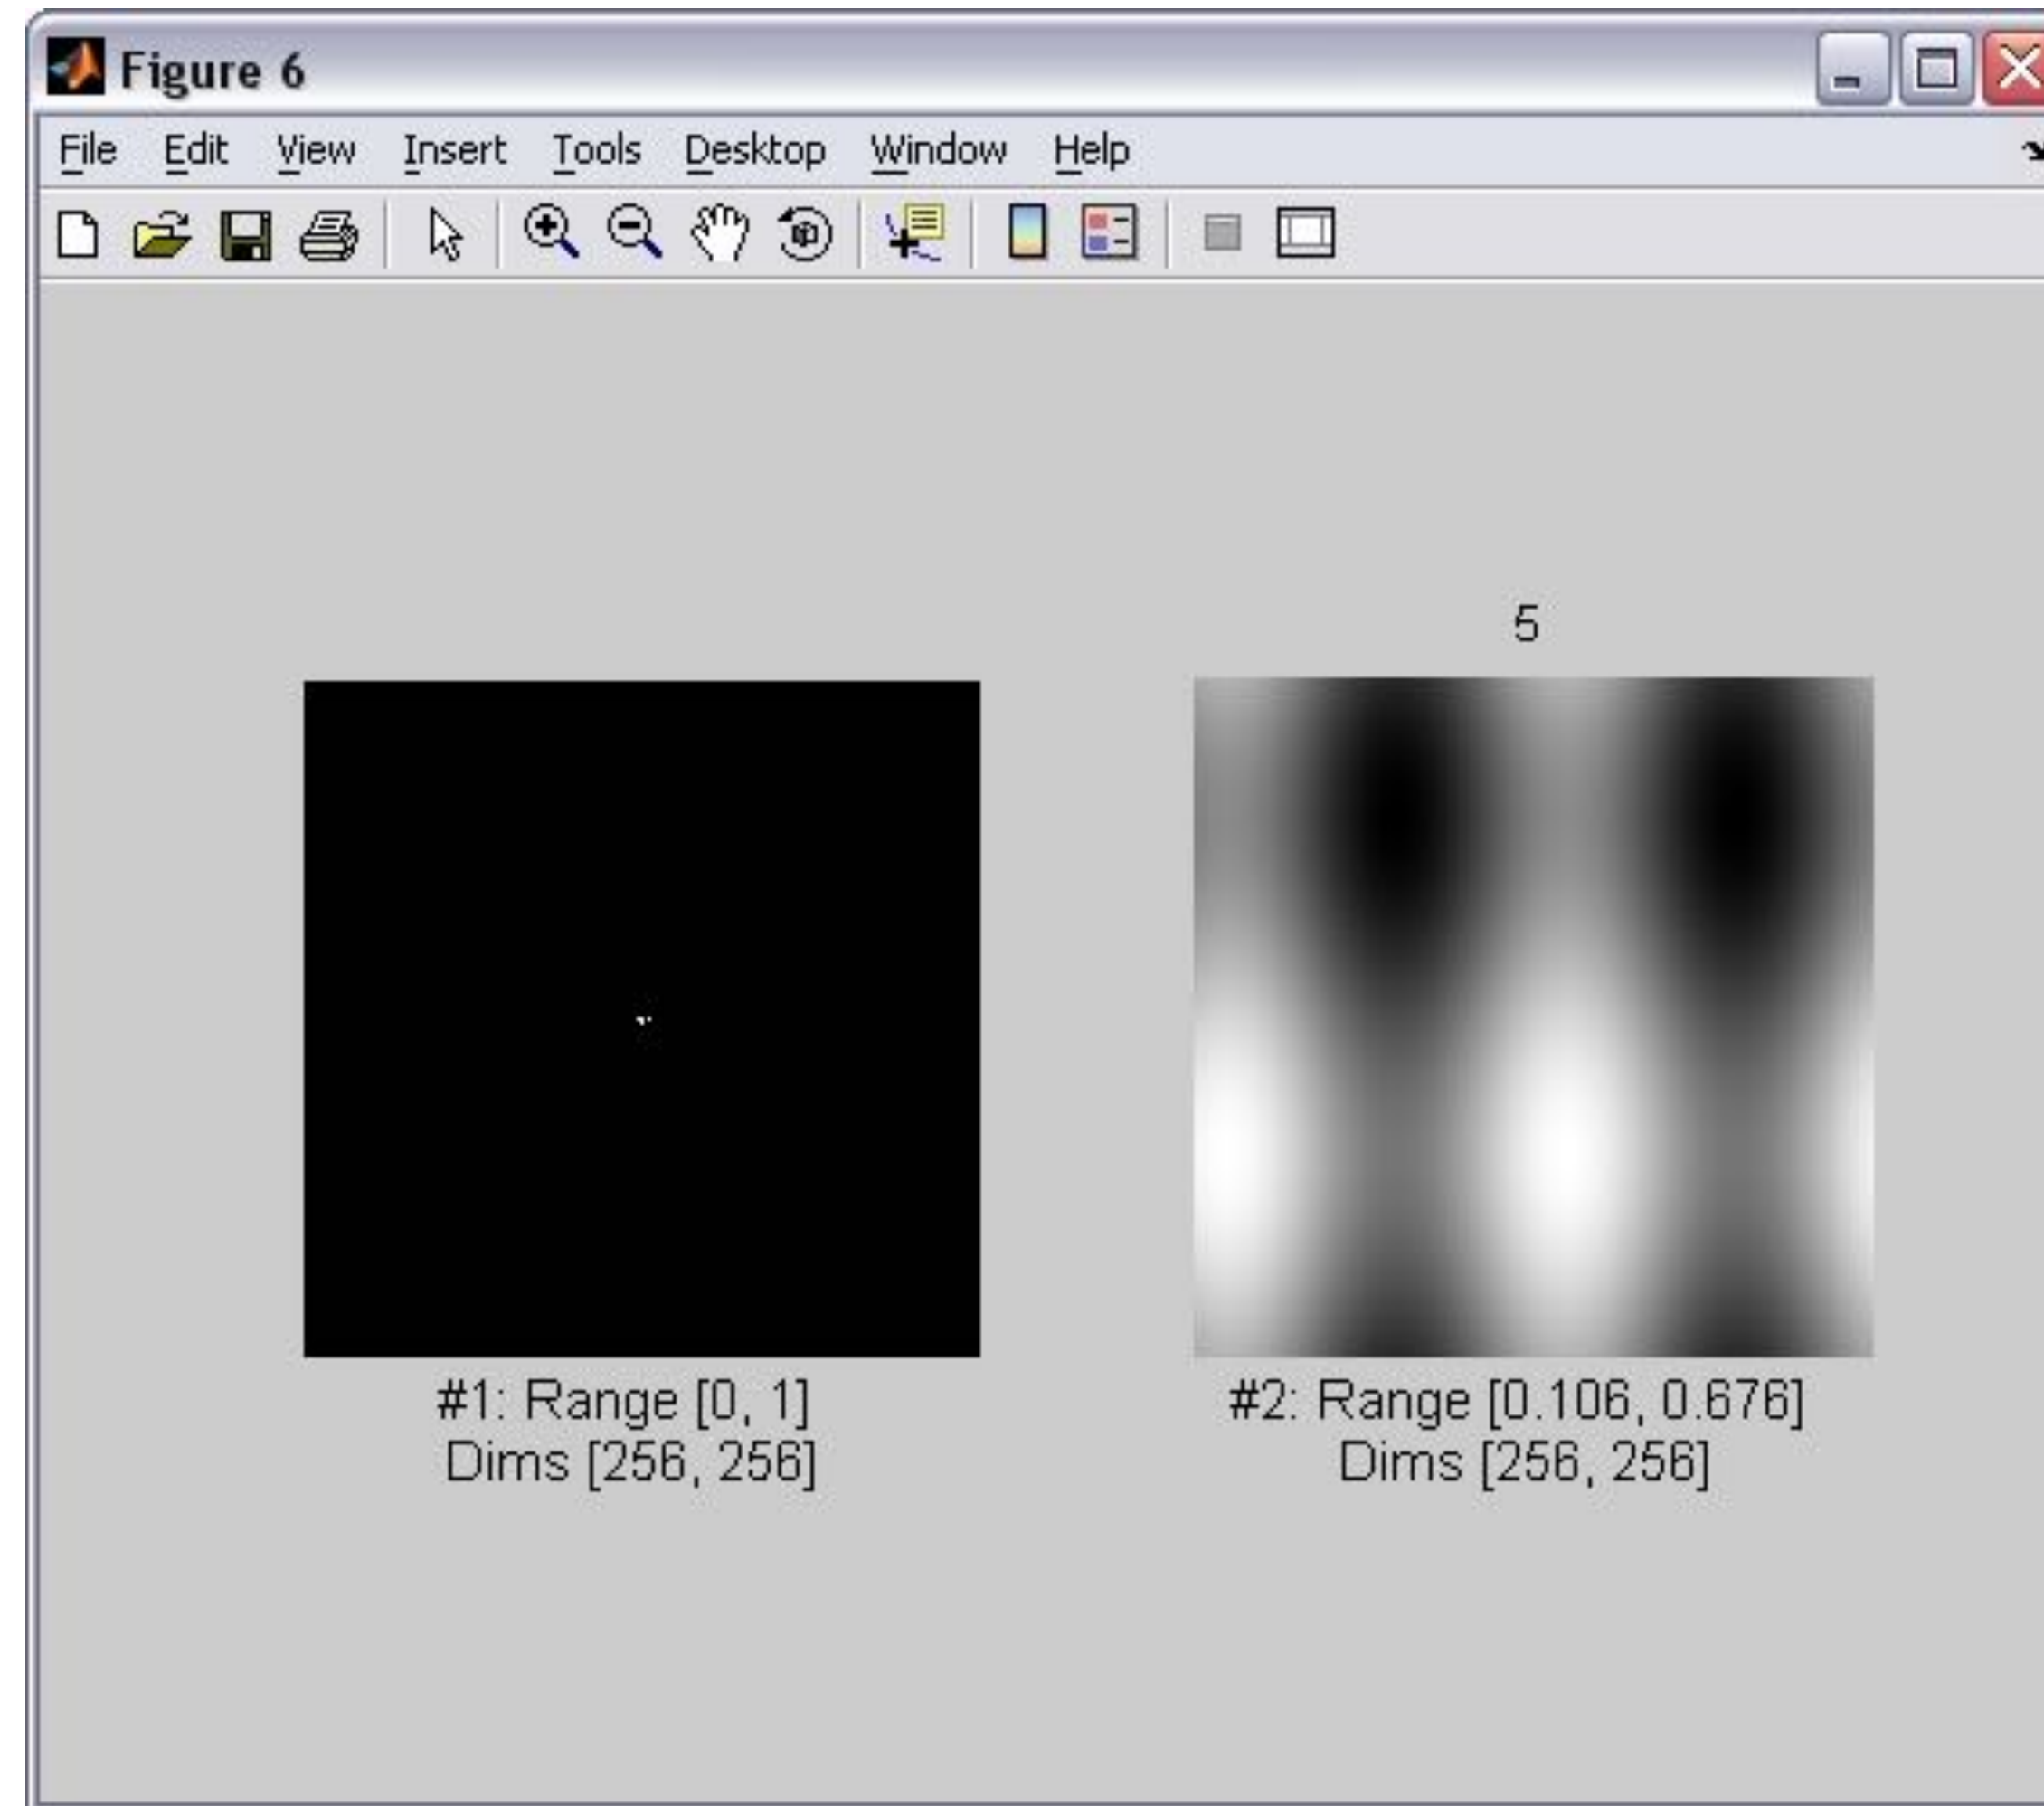
#2: Range [4.43e-015, 255]
Dims [256, 256]

3

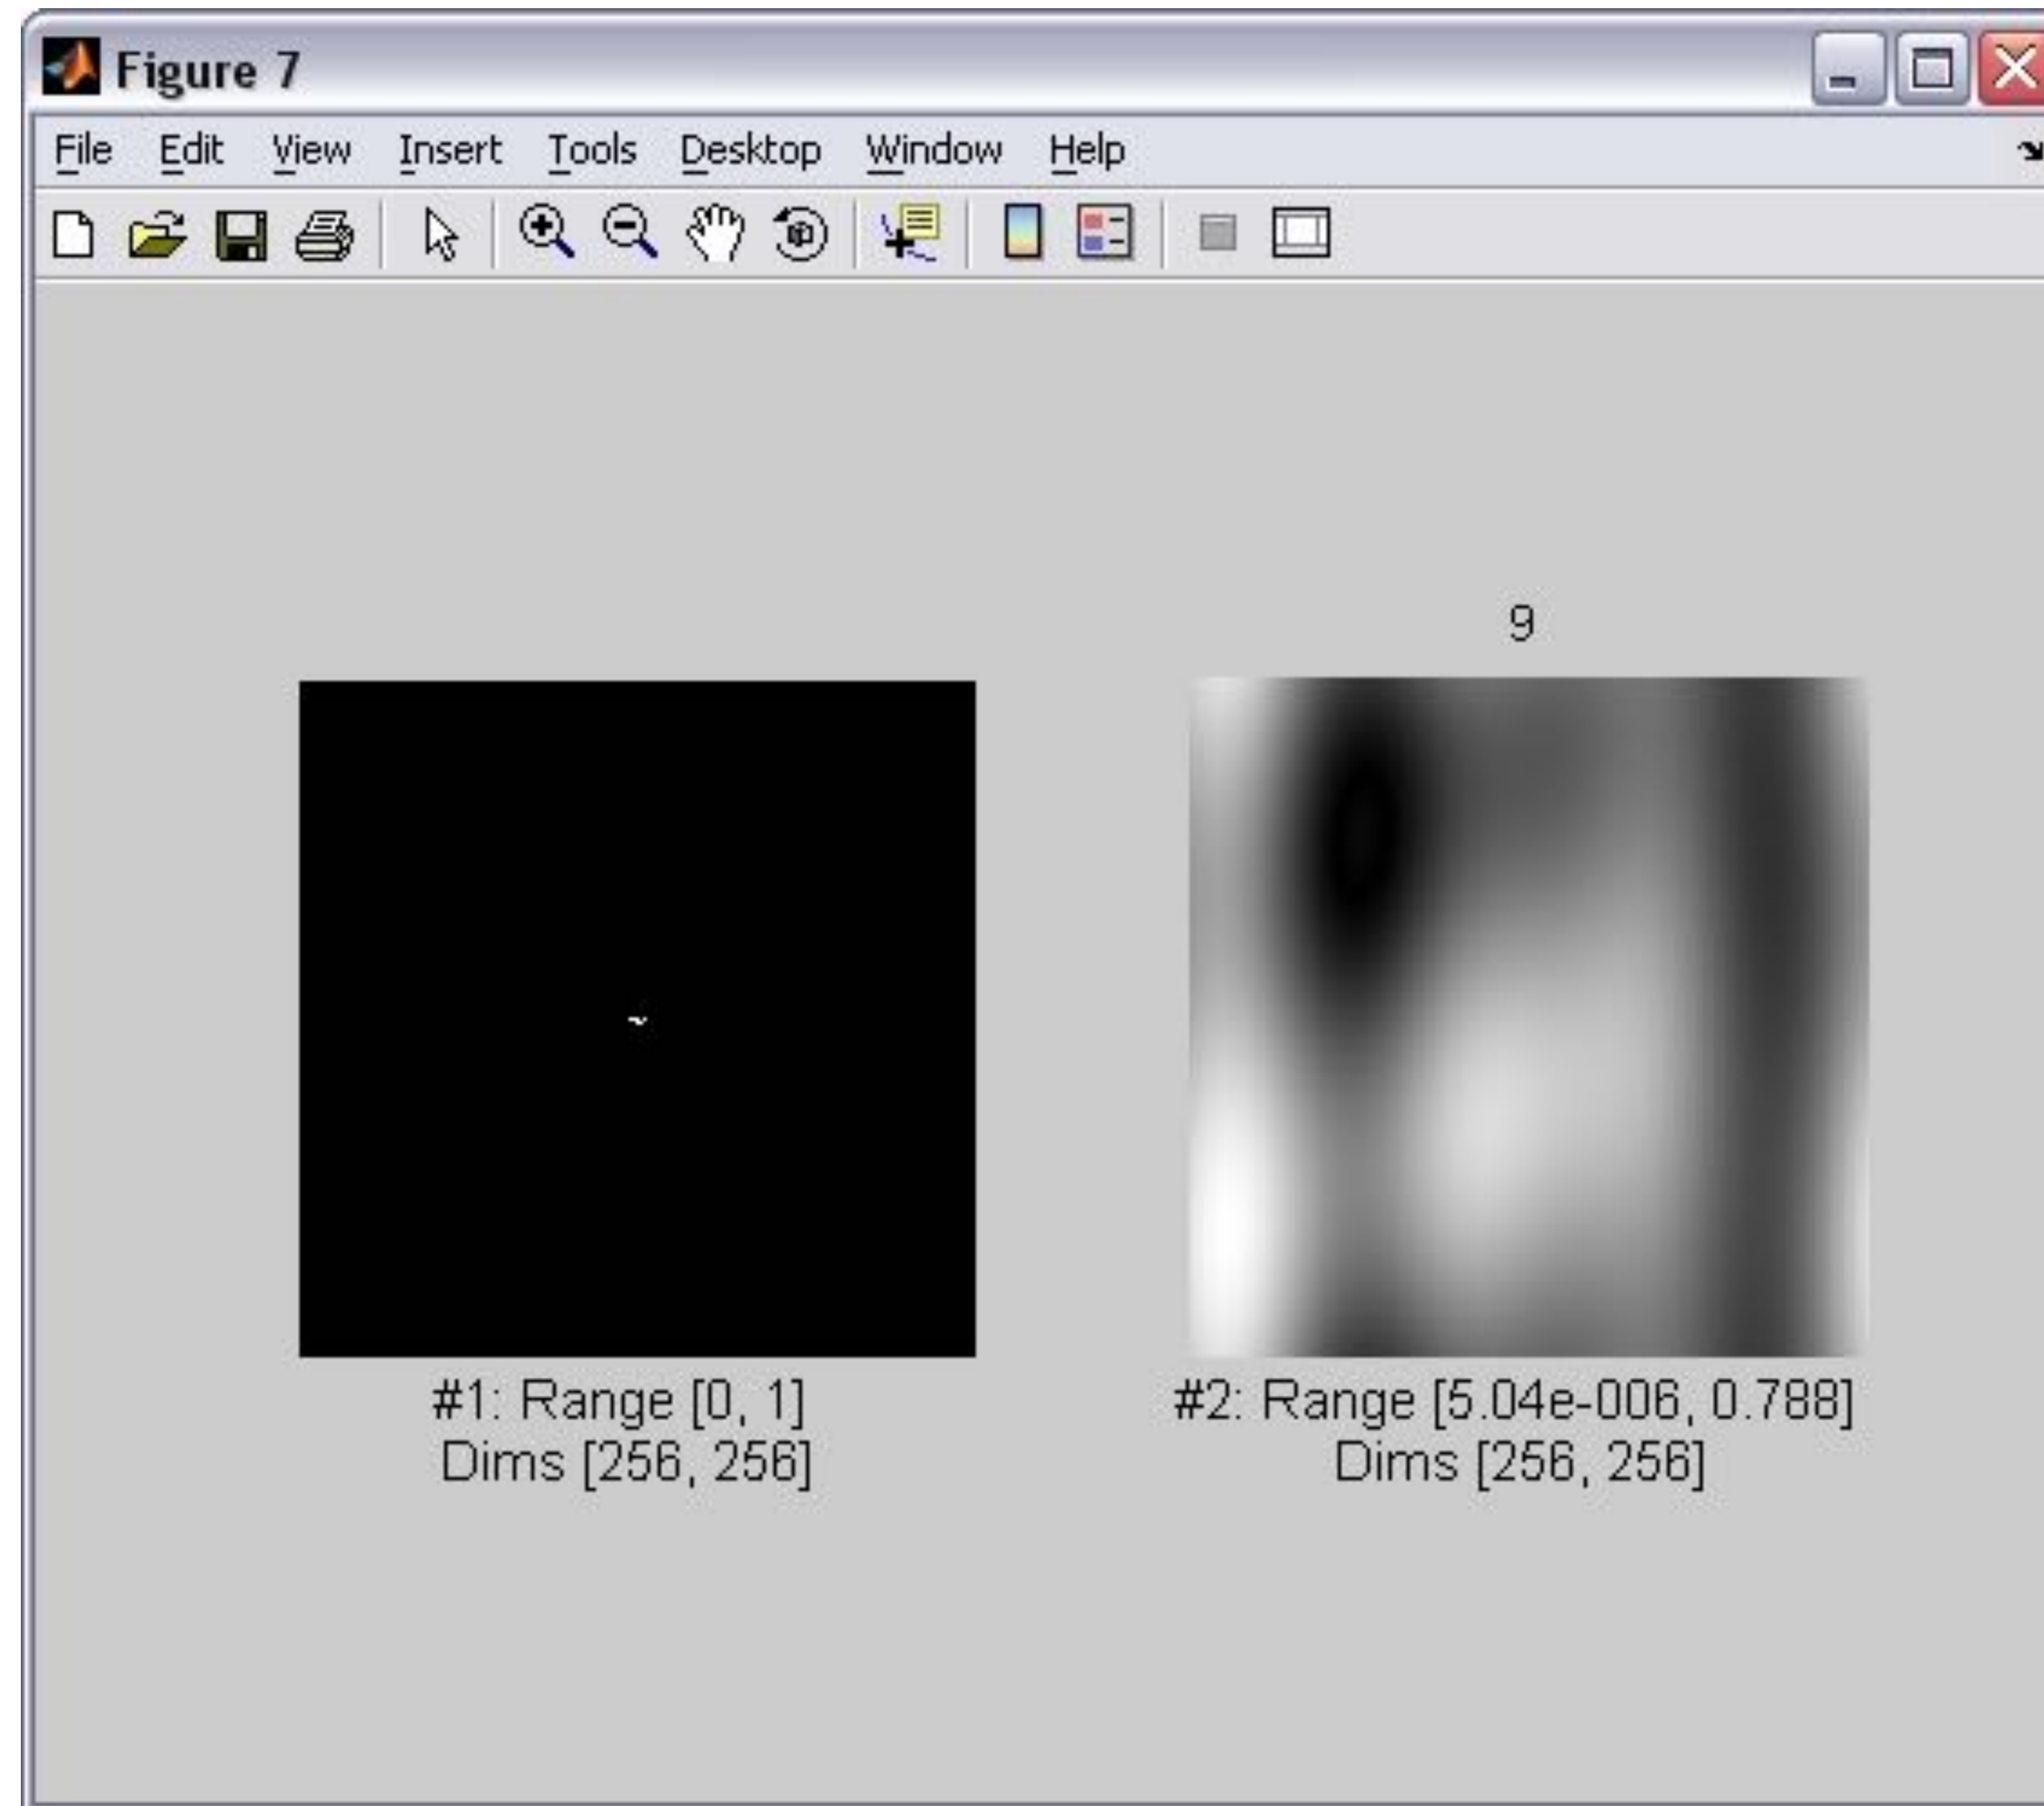


Now, an analogous sequence of images, but selecting Fourier components in descending order of magnitude.

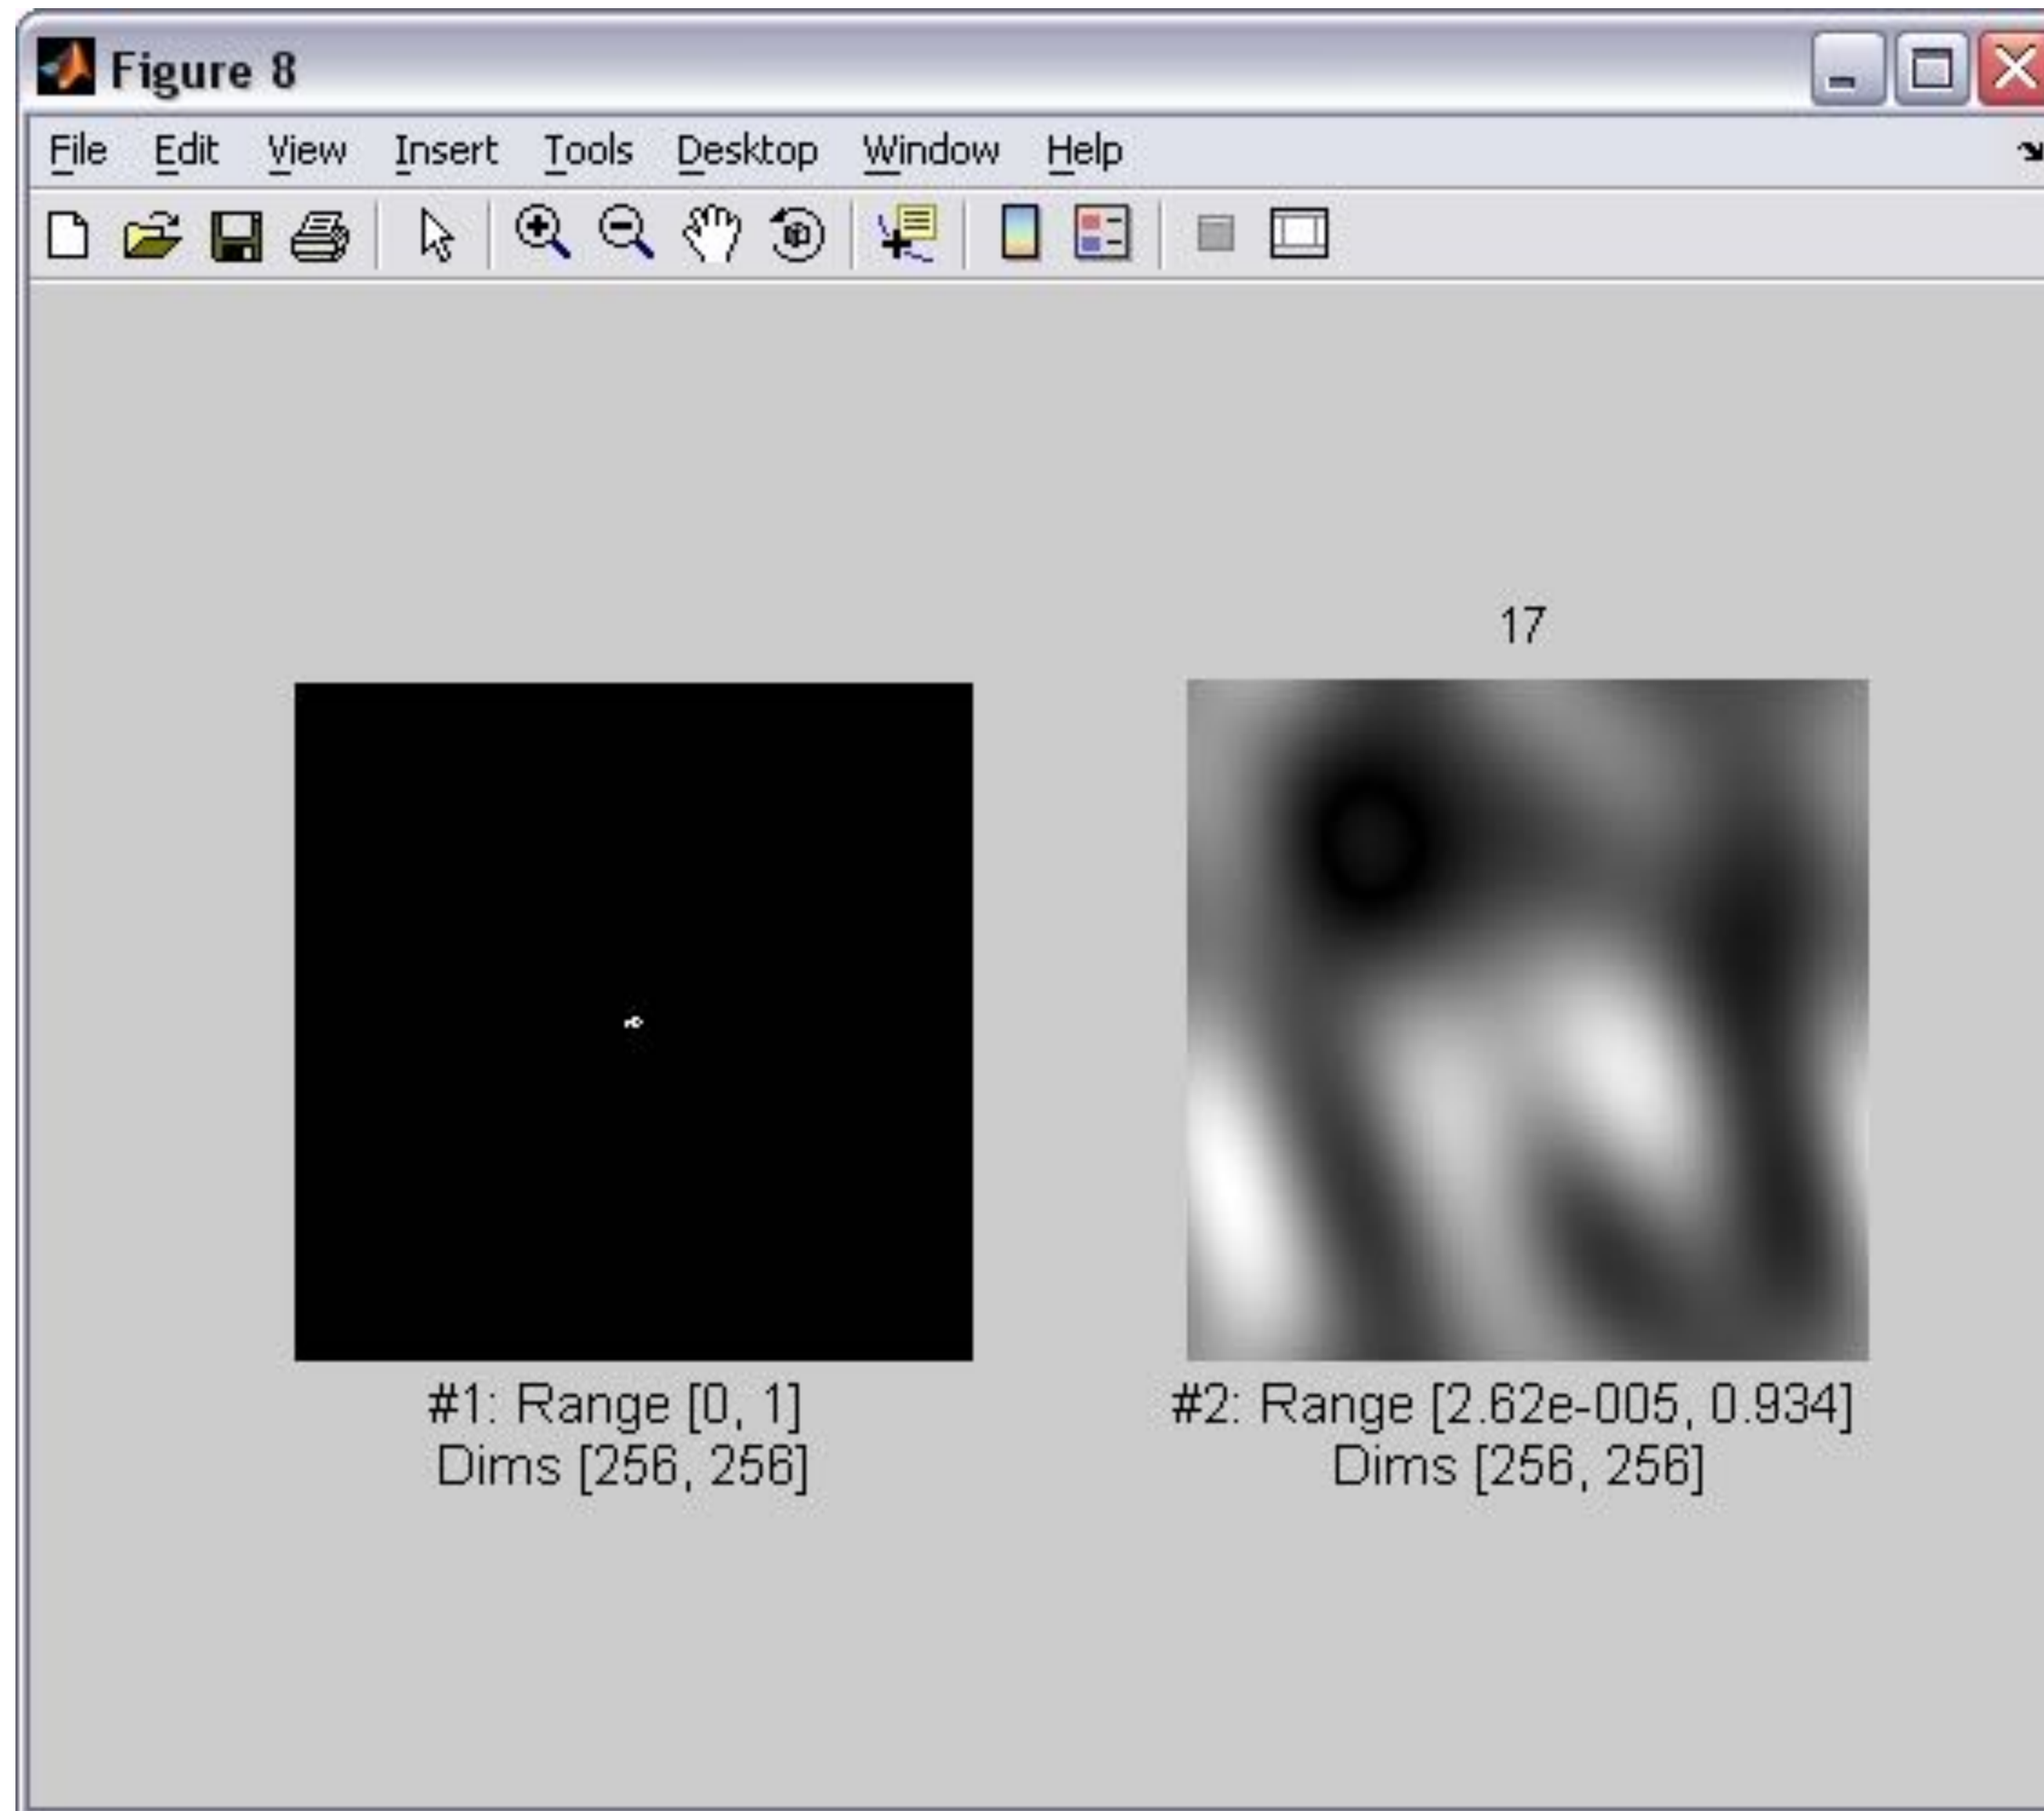
5




9



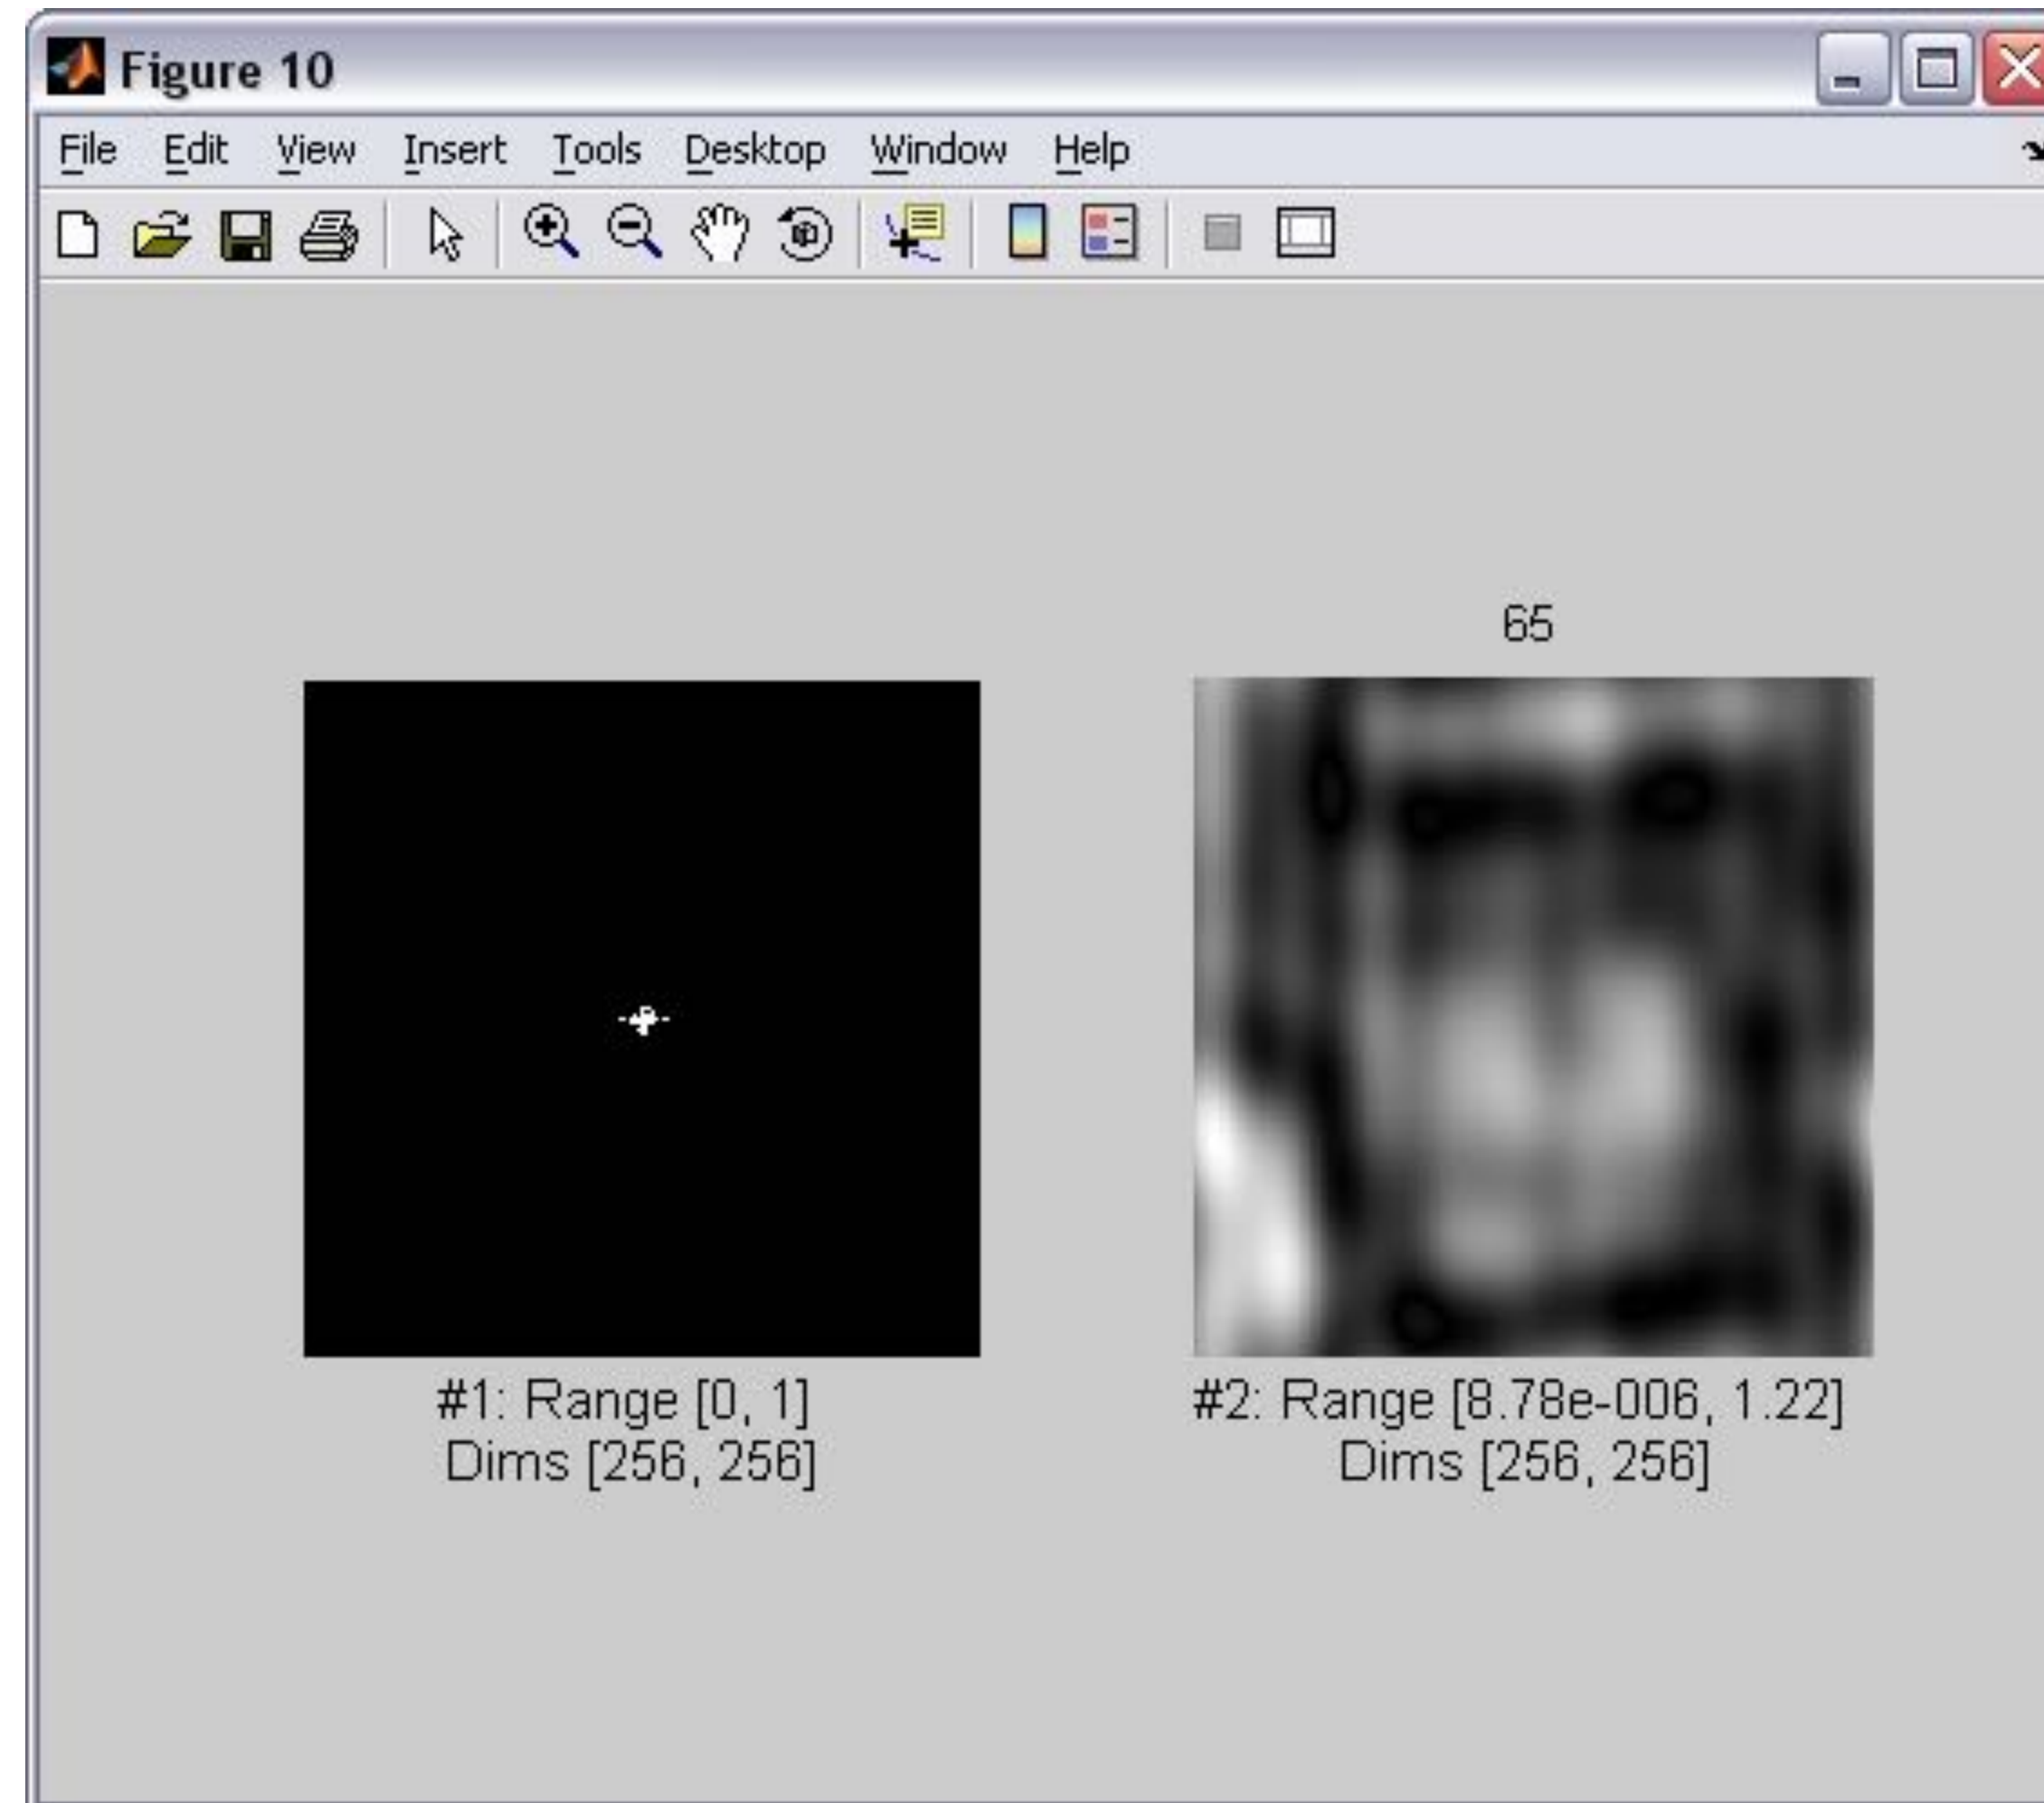
17



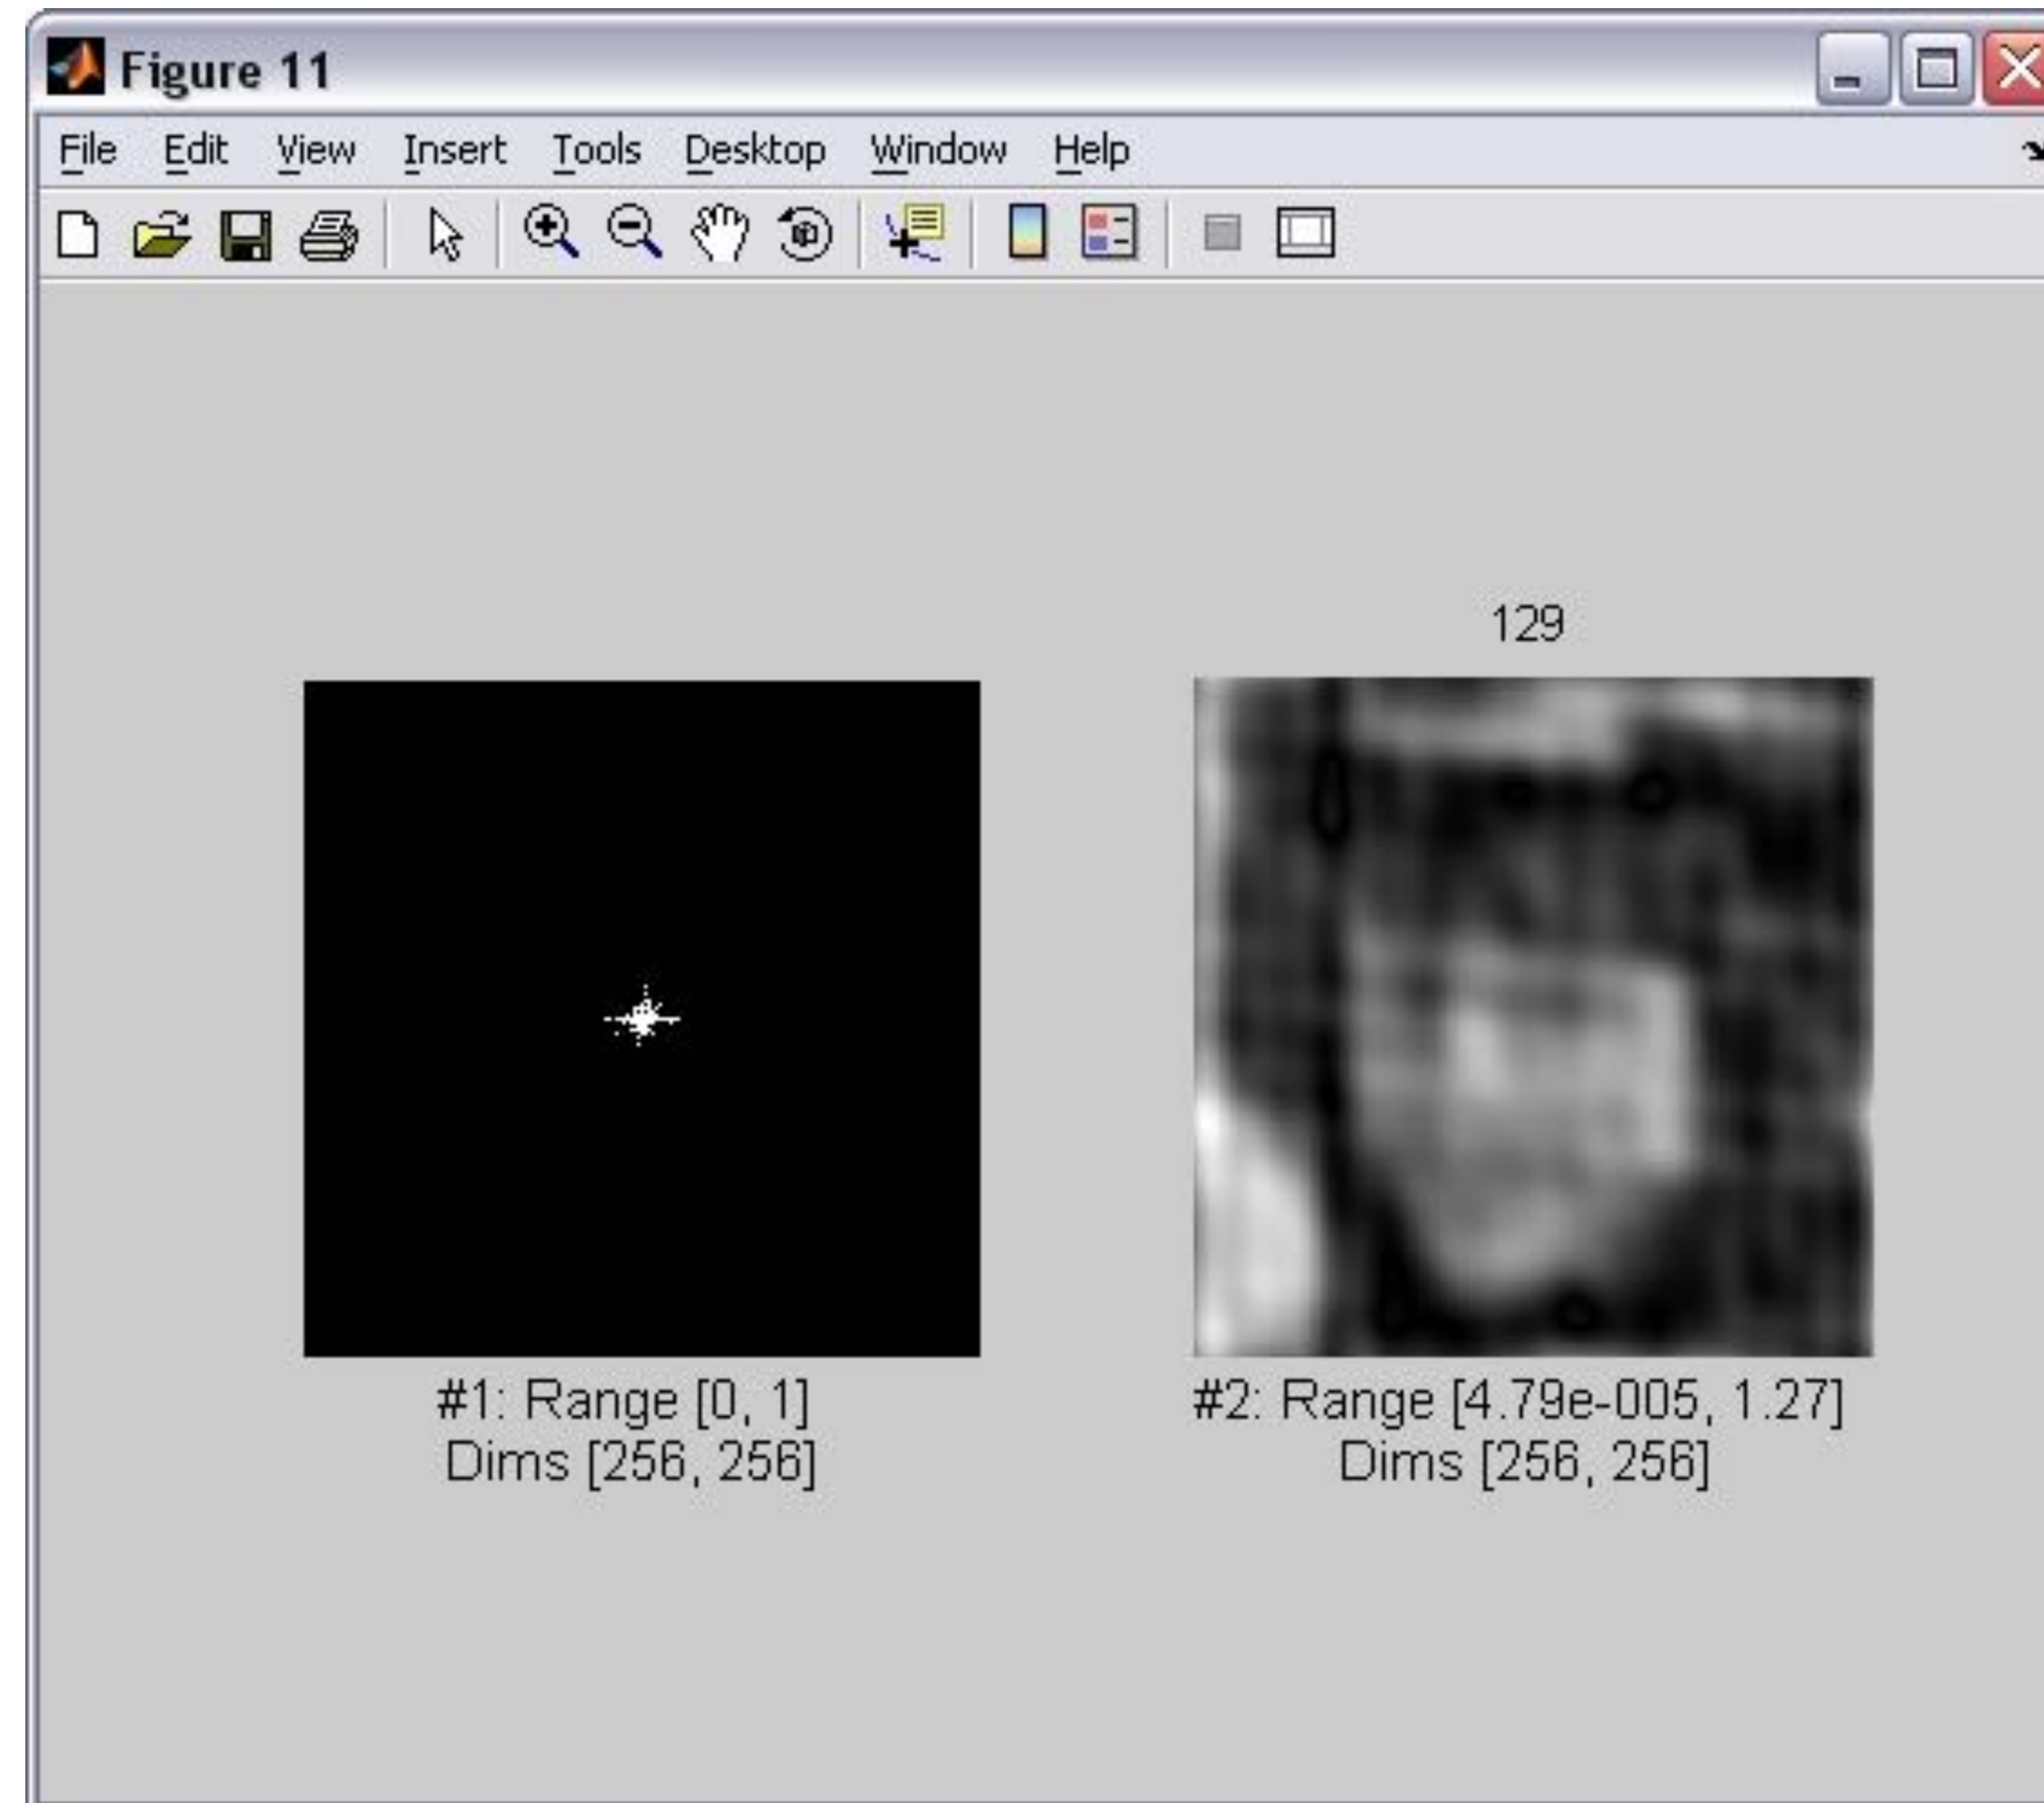
33



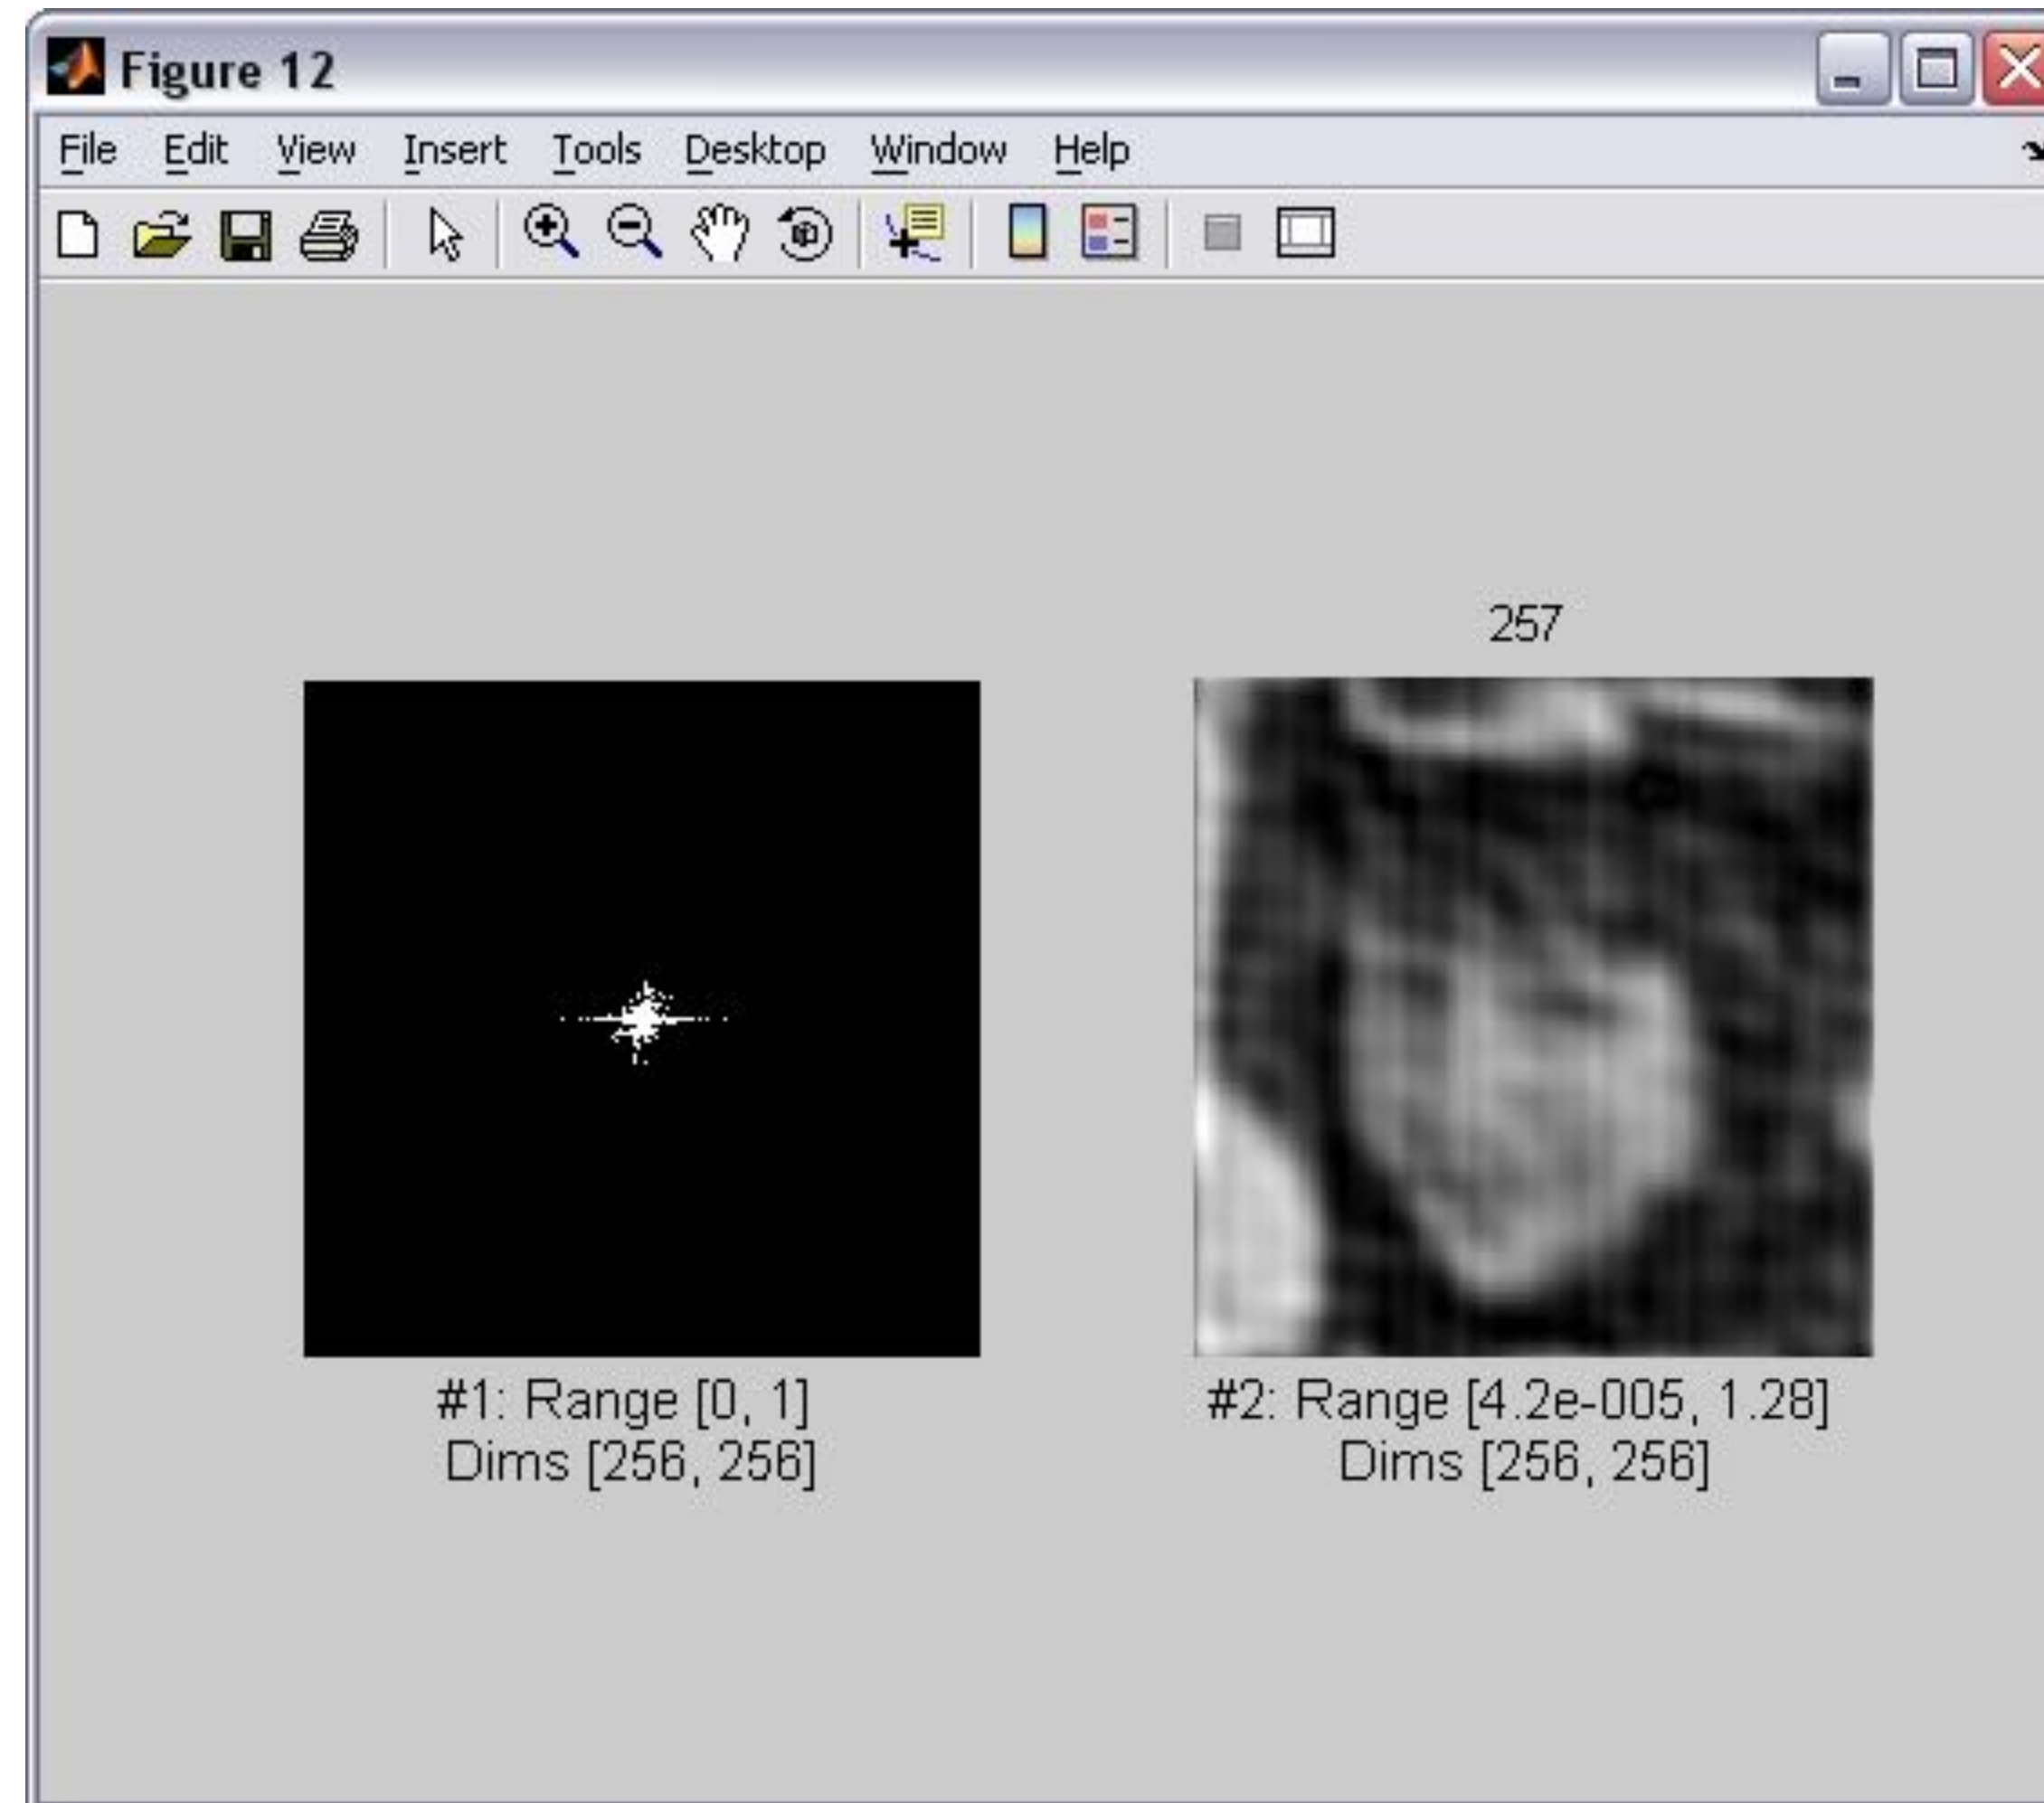
65




129




257




513




1025




2049



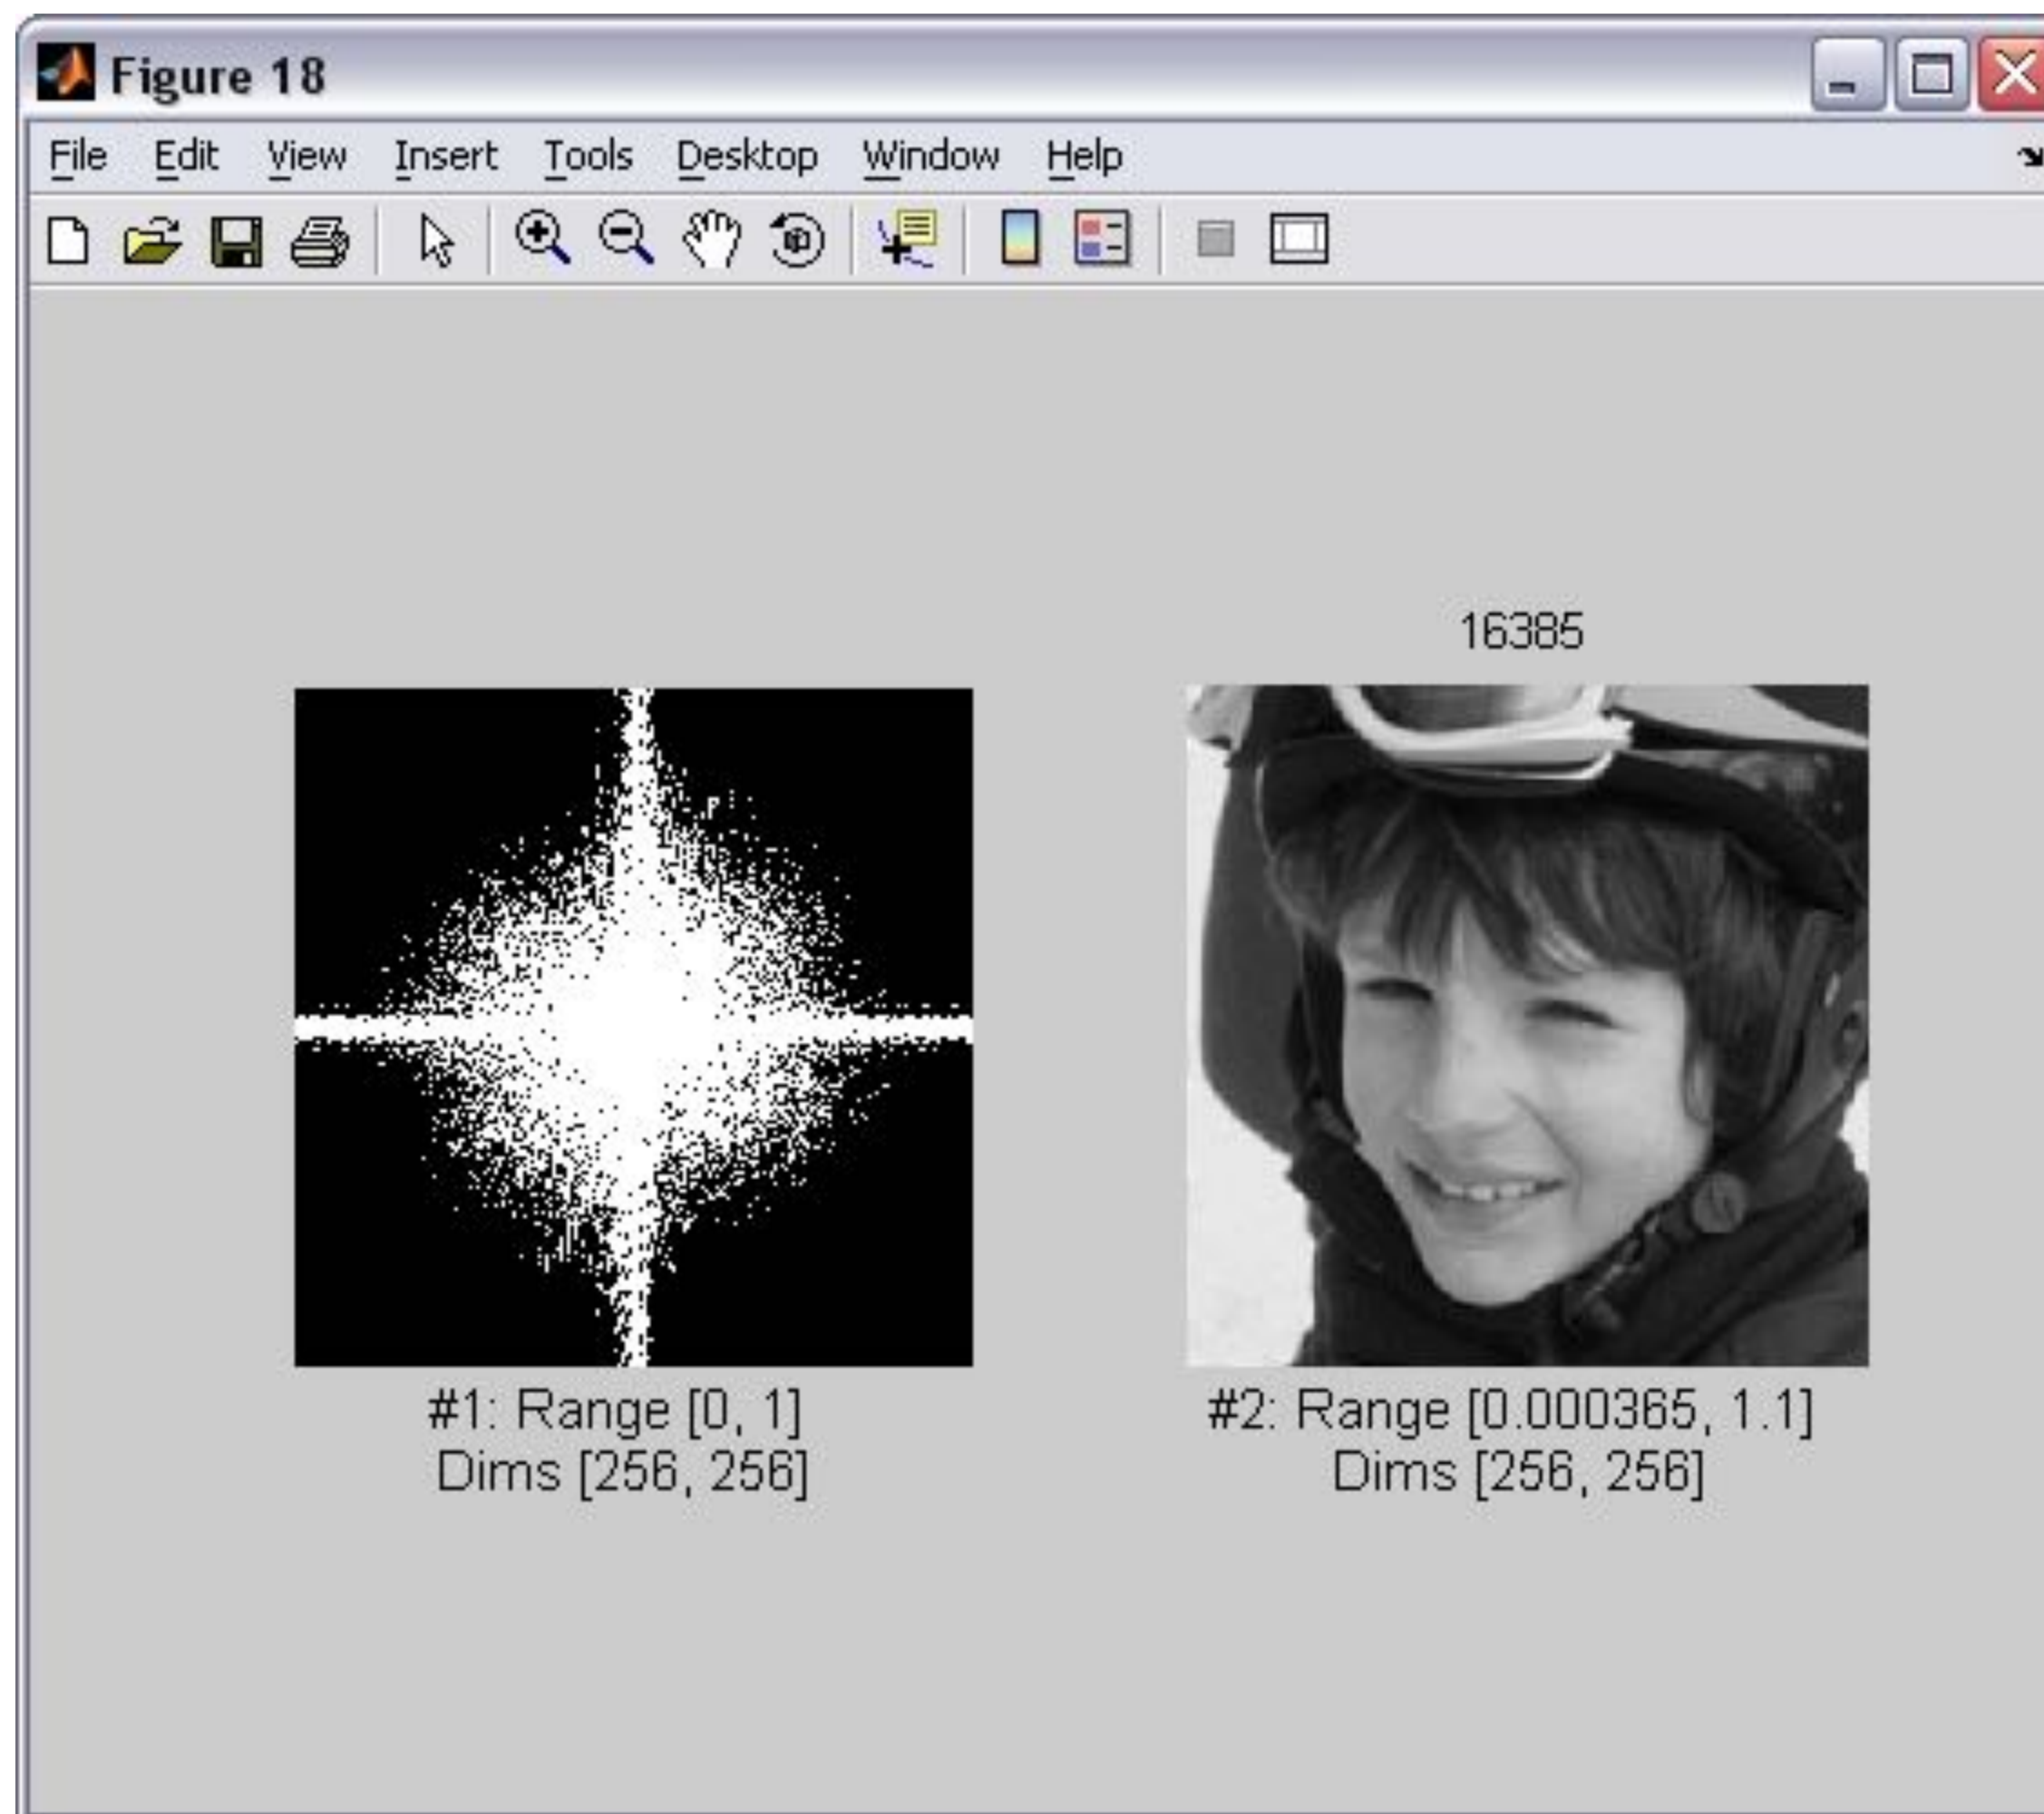
4097




8193




16385



32769




65536







DFT
→




DFT⁻¹
→





DFT
→



DFT⁻¹
→

