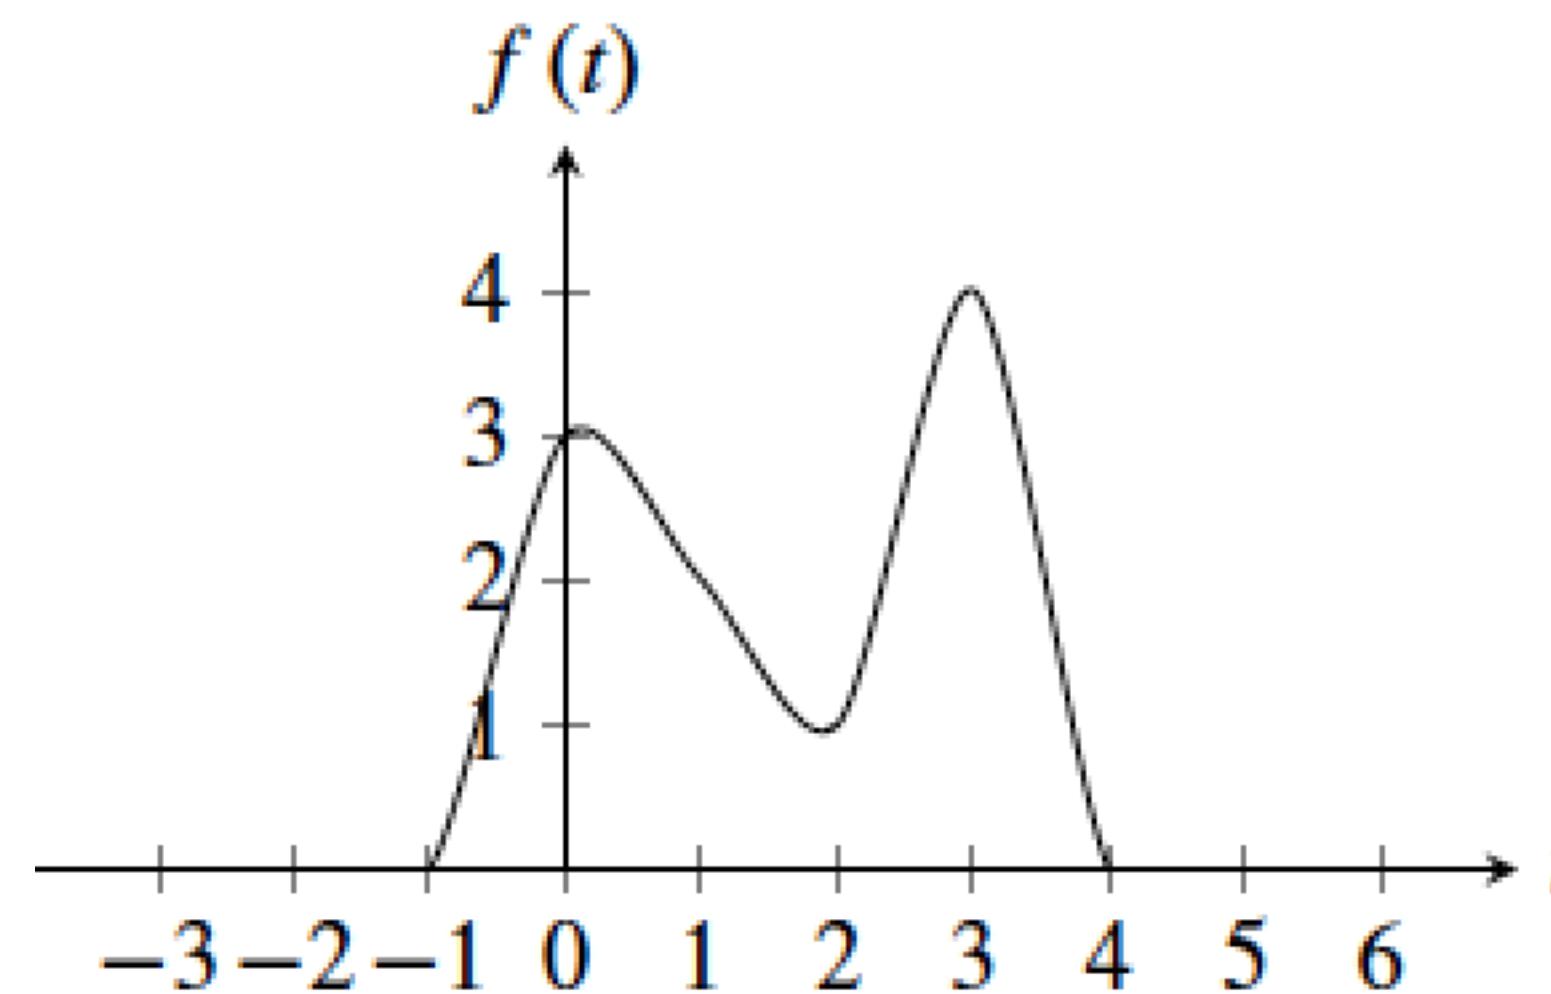


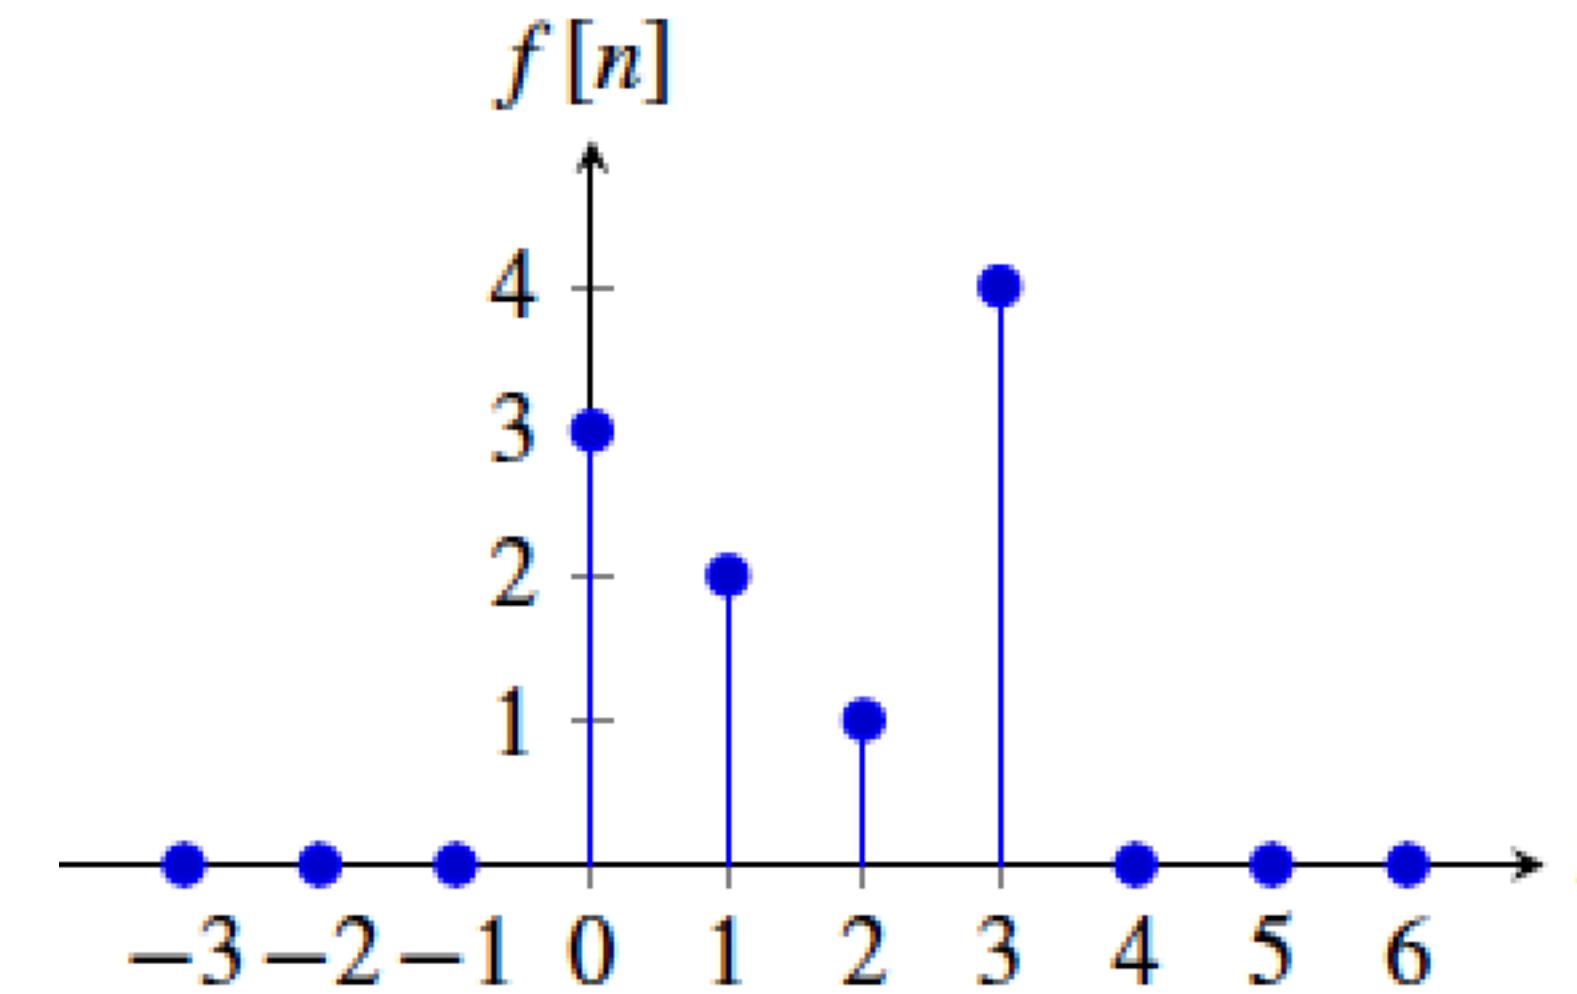
Lecture 4

Signal Processing

6.003 Signals and systems



Time continuous signal

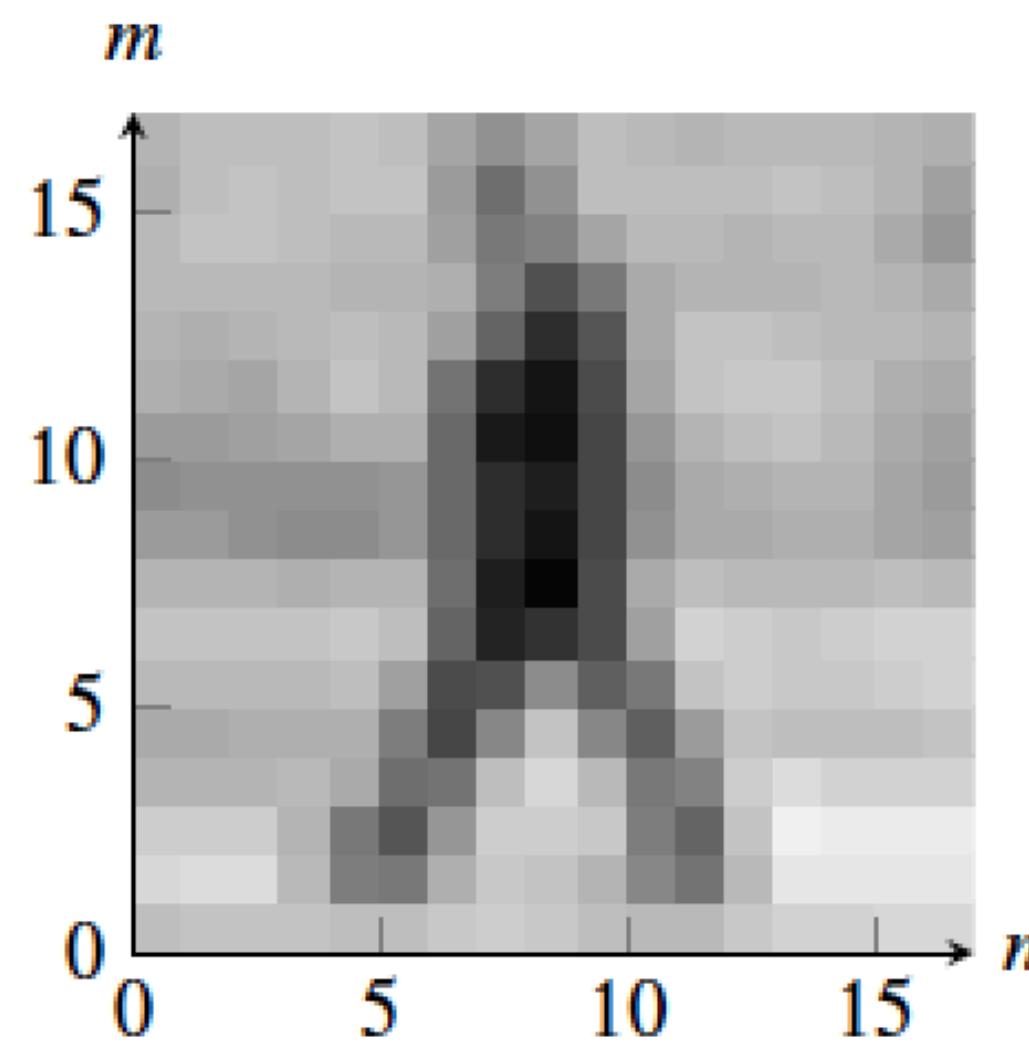


Time discrete signal

Review class notes!

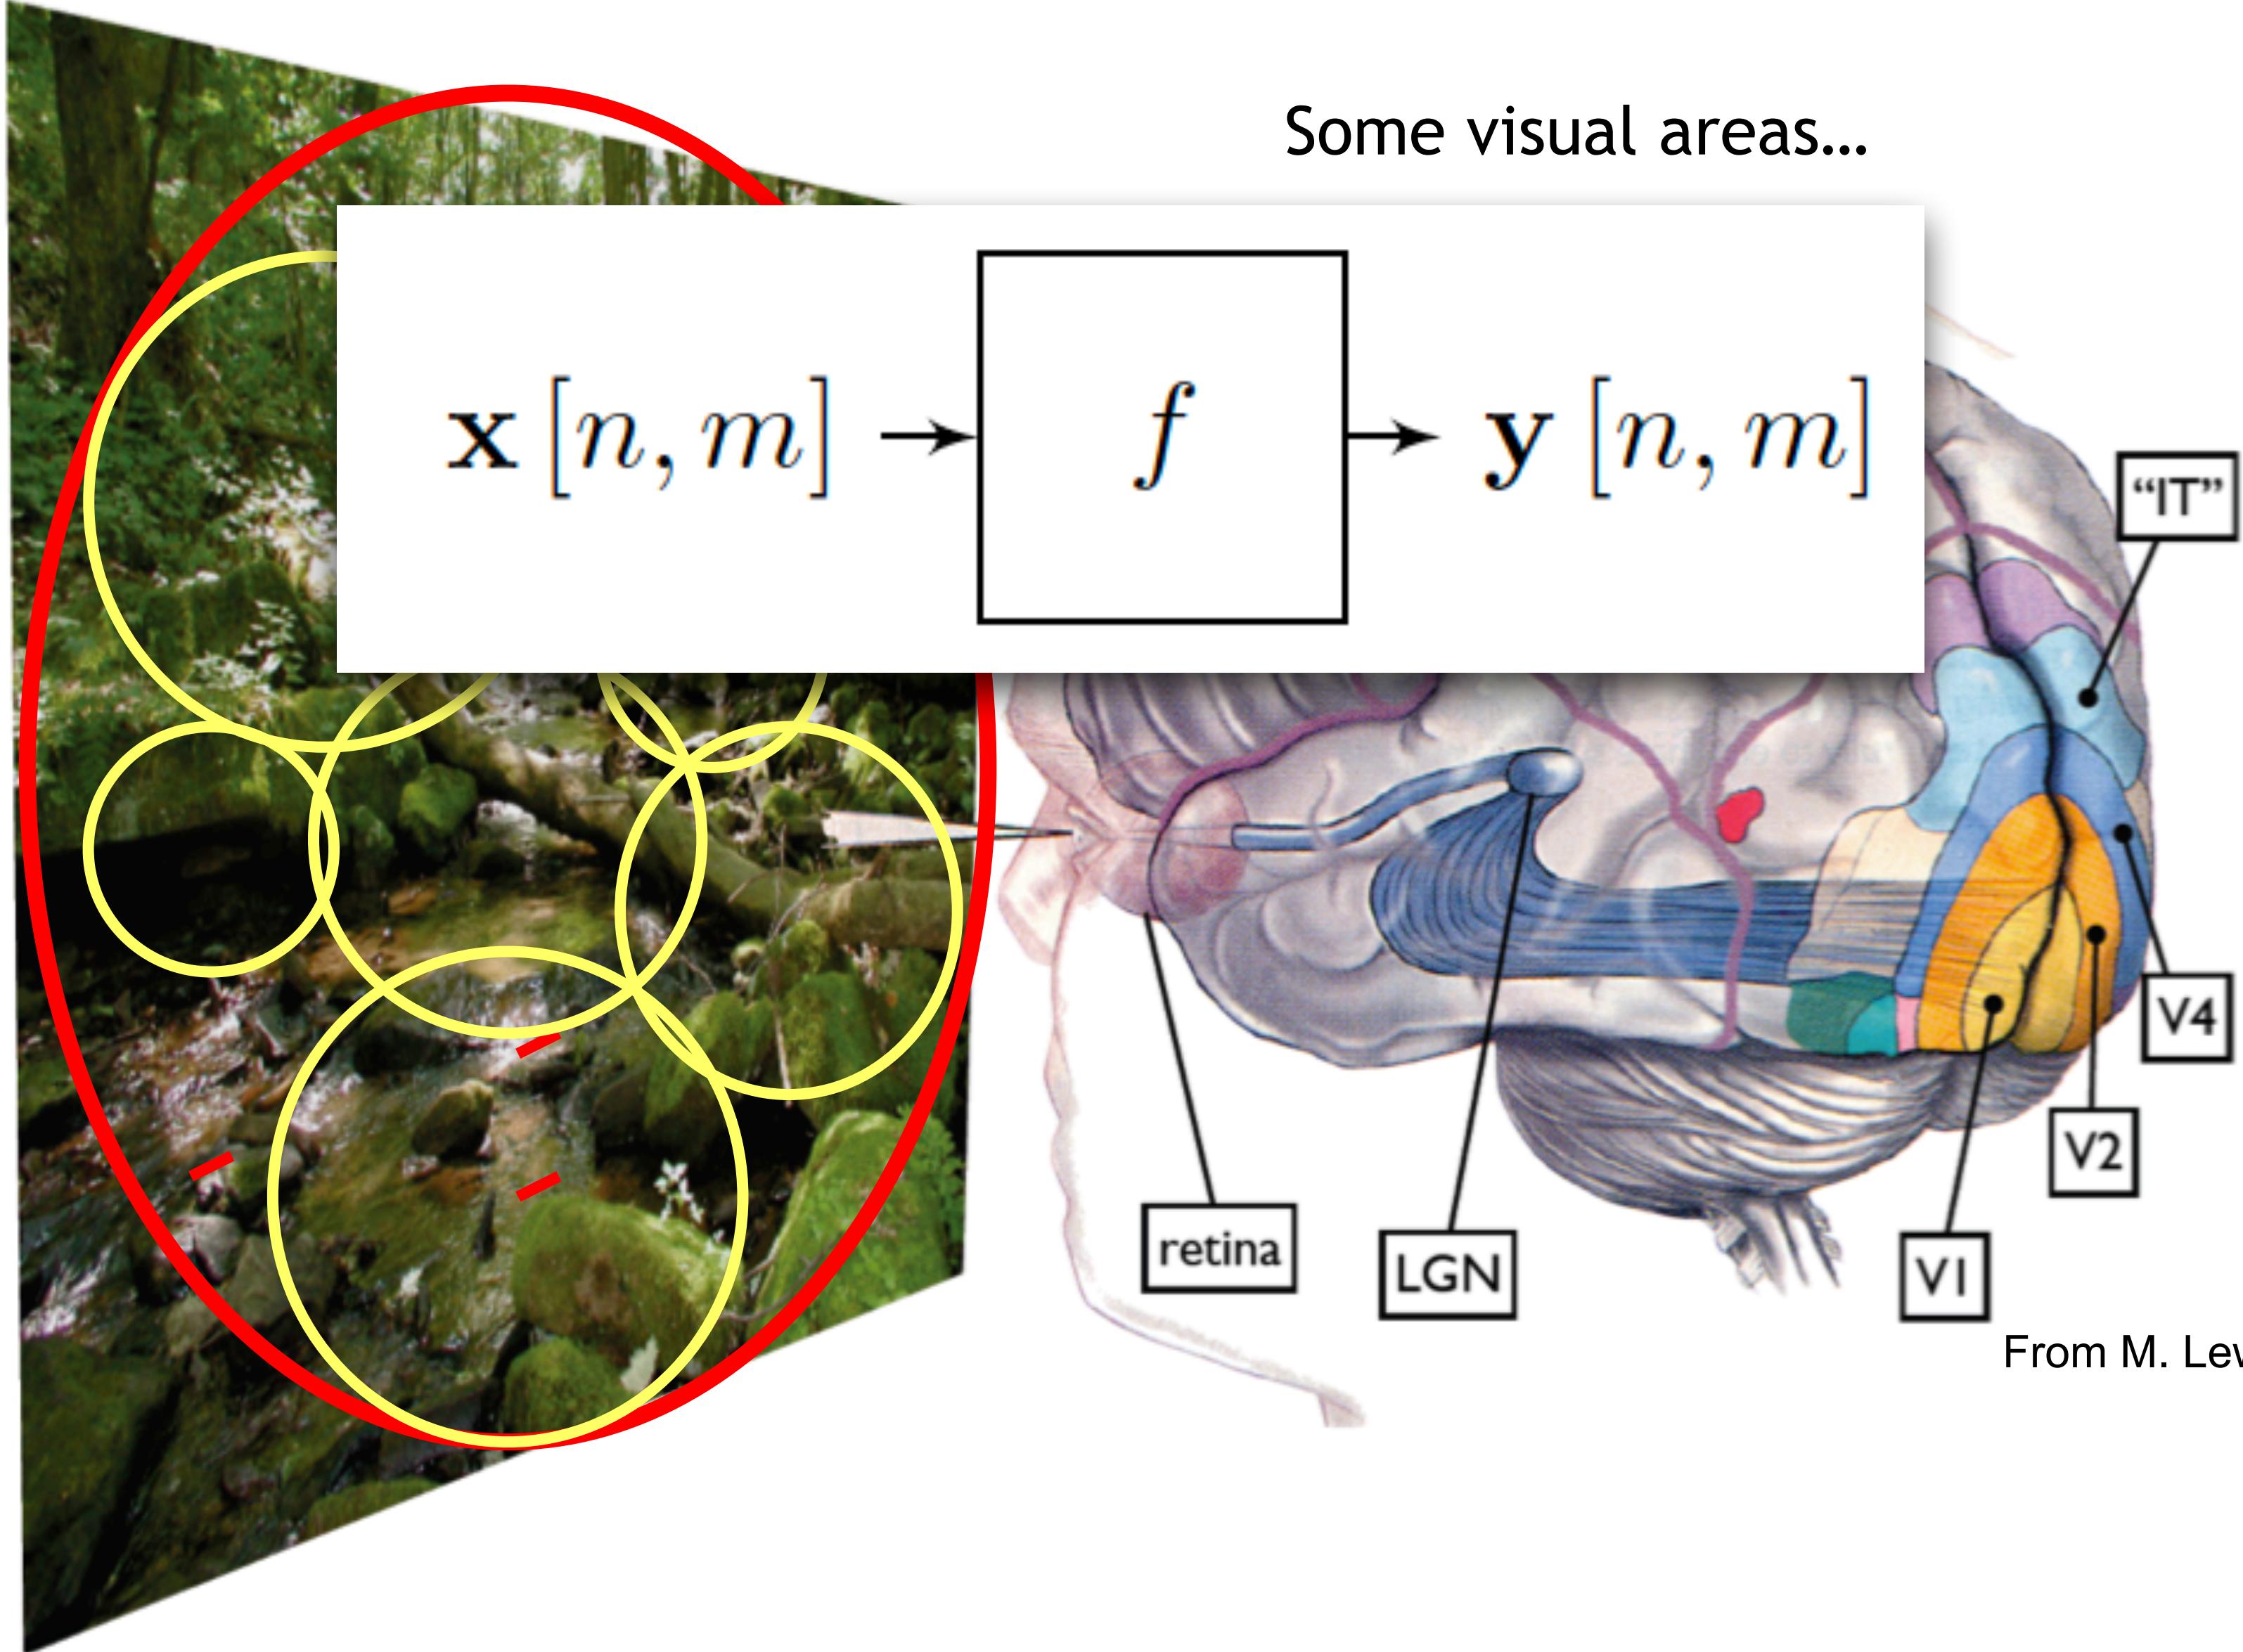
Remember, an image is just an array of numbers

What we see

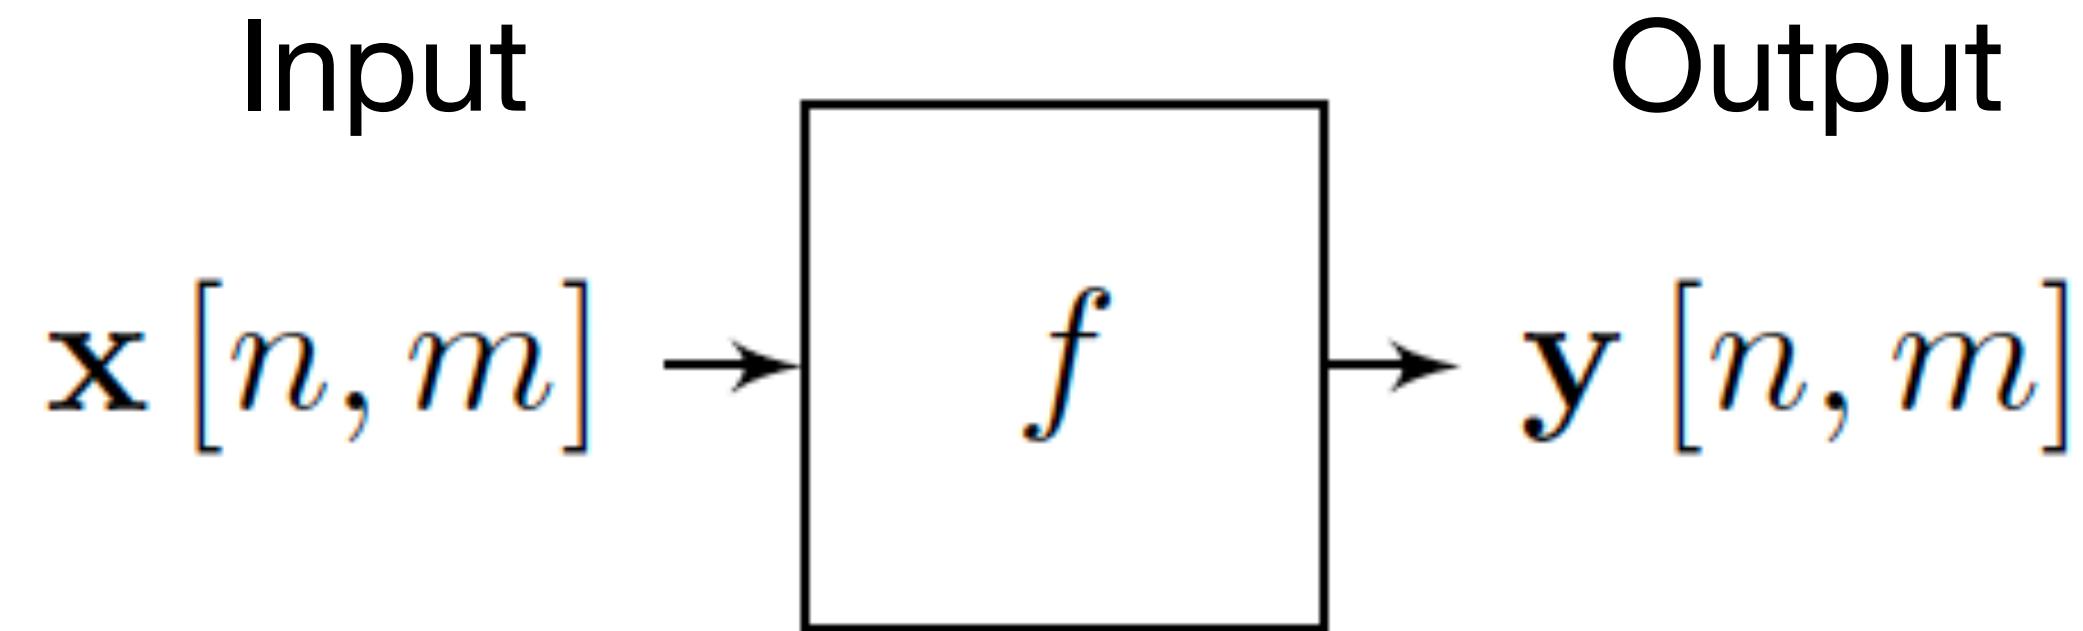


What the machine gets

$$\mathbf{I} = \begin{bmatrix} 160 & 175 & 171 & 168 & 168 & 172 & 164 & 158 & 167 & 173 & 167 & 163 & 162 & 164 & 160 & 159 & 163 & 162 \\ 149 & 164 & 172 & 175 & 178 & 179 & 176 & 118 & 97 & 168 & 175 & 171 & 169 & 175 & 176 & 177 & 165 & 152 \\ 161 & 166 & 182 & 171 & 170 & 177 & 175 & 116 & 109 & 169 & 177 & 173 & 168 & 175 & 175 & 159 & 153 & 123 \\ 171 & 174 & 177 & 175 & 167 & 161 & 157 & 138 & 103 & 112 & 157 & 164 & 159 & 160 & 165 & 169 & 148 & 144 \\ 163 & 163 & 162 & 165 & 167 & 164 & 178 & 167 & 77 & 55 & 134 & 170 & 167 & 162 & 164 & 175 & 168 & 160 \\ 173 & 164 & 158 & 165 & 180 & 180 & 150 & 89 & 61 & 34 & 137 & 186 & 186 & 182 & 175 & 165 & 160 & 164 \\ 152 & 155 & 146 & 147 & 169 & 180 & 163 & 51 & 24 & 32 & 119 & 163 & 175 & 182 & 181 & 162 & 148 & 153 \\ 134 & 135 & 147 & 149 & 150 & 147 & 148 & 62 & 36 & 46 & 114 & 157 & 163 & 167 & 169 & 163 & 146 & 147 \\ 135 & 132 & 131 & 125 & 115 & 129 & 132 & 74 & 54 & 41 & 104 & 156 & 152 & 156 & 164 & 156 & 141 & 144 \\ 151 & 155 & 151 & 145 & 144 & 149 & 143 & 71 & 31 & 29 & 129 & 164 & 157 & 155 & 159 & 158 & 156 & 148 \\ 172 & 174 & 178 & 177 & 177 & 181 & 174 & 54 & 21 & 29 & 136 & 190 & 180 & 179 & 176 & 184 & 187 & 182 \\ 177 & 178 & 176 & 173 & 174 & 180 & 150 & 27 & 101 & 94 & 74 & 189 & 188 & 186 & 183 & 186 & 188 & 187 \\ 160 & 160 & 163 & 163 & 161 & 167 & 100 & 45 & 169 & 166 & 59 & 136 & 184 & 176 & 175 & 177 & 185 & 186 \\ 147 & 150 & 153 & 155 & 160 & 155 & 56 & 111 & 182 & 180 & 104 & 84 & 168 & 172 & 171 & 164 & 168 & 167 \\ 184 & 182 & 178 & 175 & 179 & 133 & 86 & 191 & 201 & 204 & 191 & 79 & 172 & 220 & 217 & 205 & 209 & 200 \\ 184 & 187 & 192 & 182 & 124 & 32 & 109 & 168 & 171 & 167 & 163 & 51 & 105 & 203 & 209 & 203 & 210 & 205 \\ 191 & 198 & 203 & 197 & 175 & 149 & 169 & 189 & 190 & 173 & 160 & 145 & 156 & 202 & 199 & 201 & 205 & 202 \\ 153 & 149 & 153 & 155 & 173 & 182 & 179 & 177 & 182 & 177 & 182 & 185 & 179 & 177 & 167 & 176 & 182 & 180 \end{bmatrix}$$



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: $y = f(x)$

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

The diagram illustrates a convolutional layer mapping an input X to an output y . The input X is represented by a 4x4 grid of gray squares. The output y is represented by a 2x2 grid of white squares. A blue box highlights the element $h[n, k]$ at the top-left position of the output grid. To the left of the input grid, the equation $y =$ is shown, indicating the mapping from X to y .

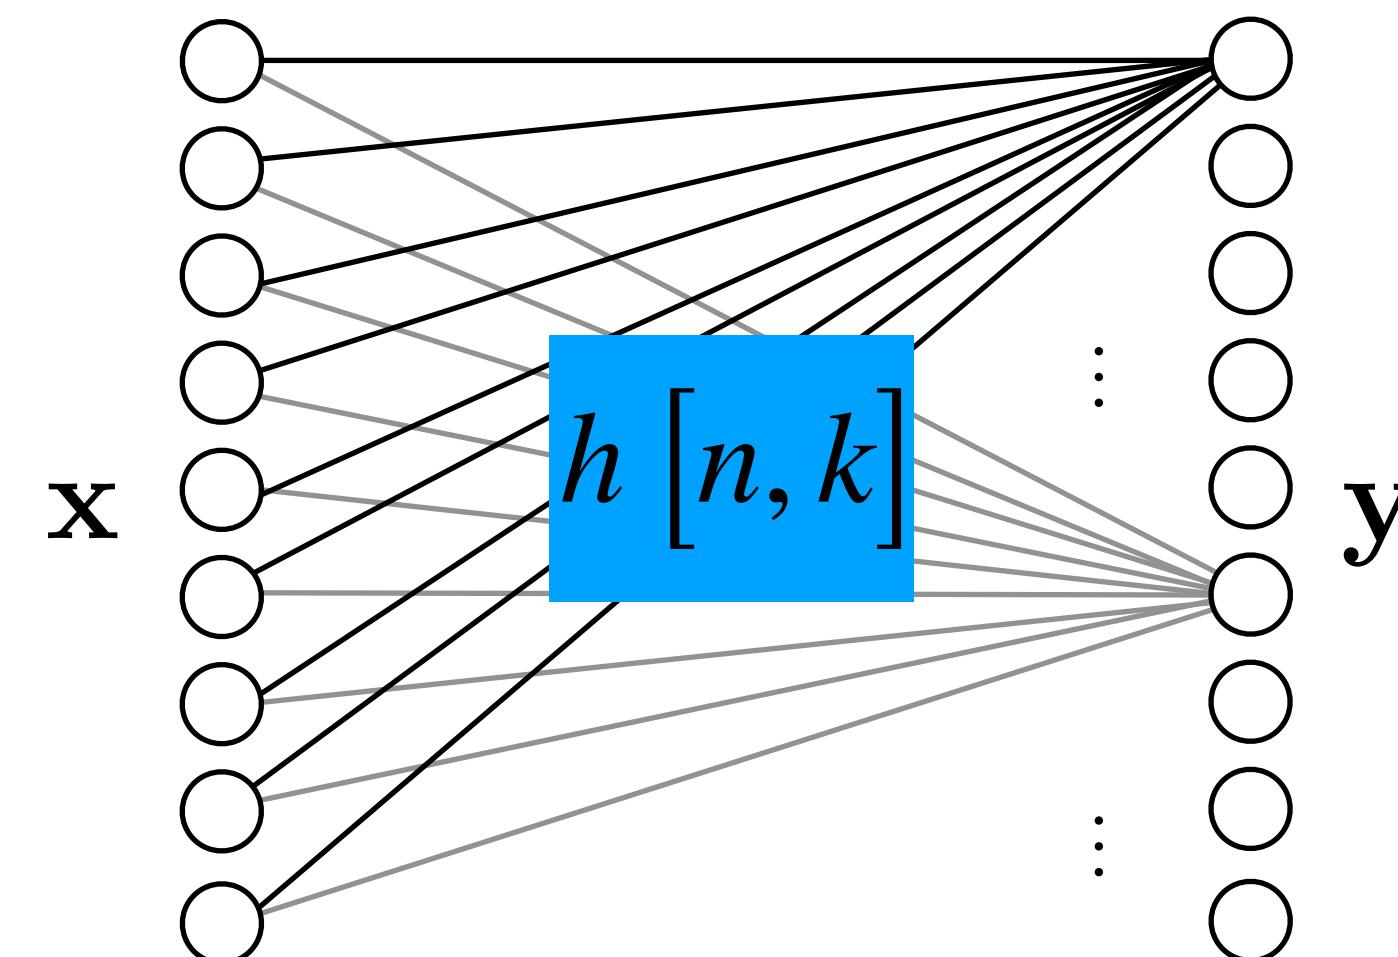
n indexes rows,
k indexes columns

Linear system: $y = f(x)$

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 $y[n]$

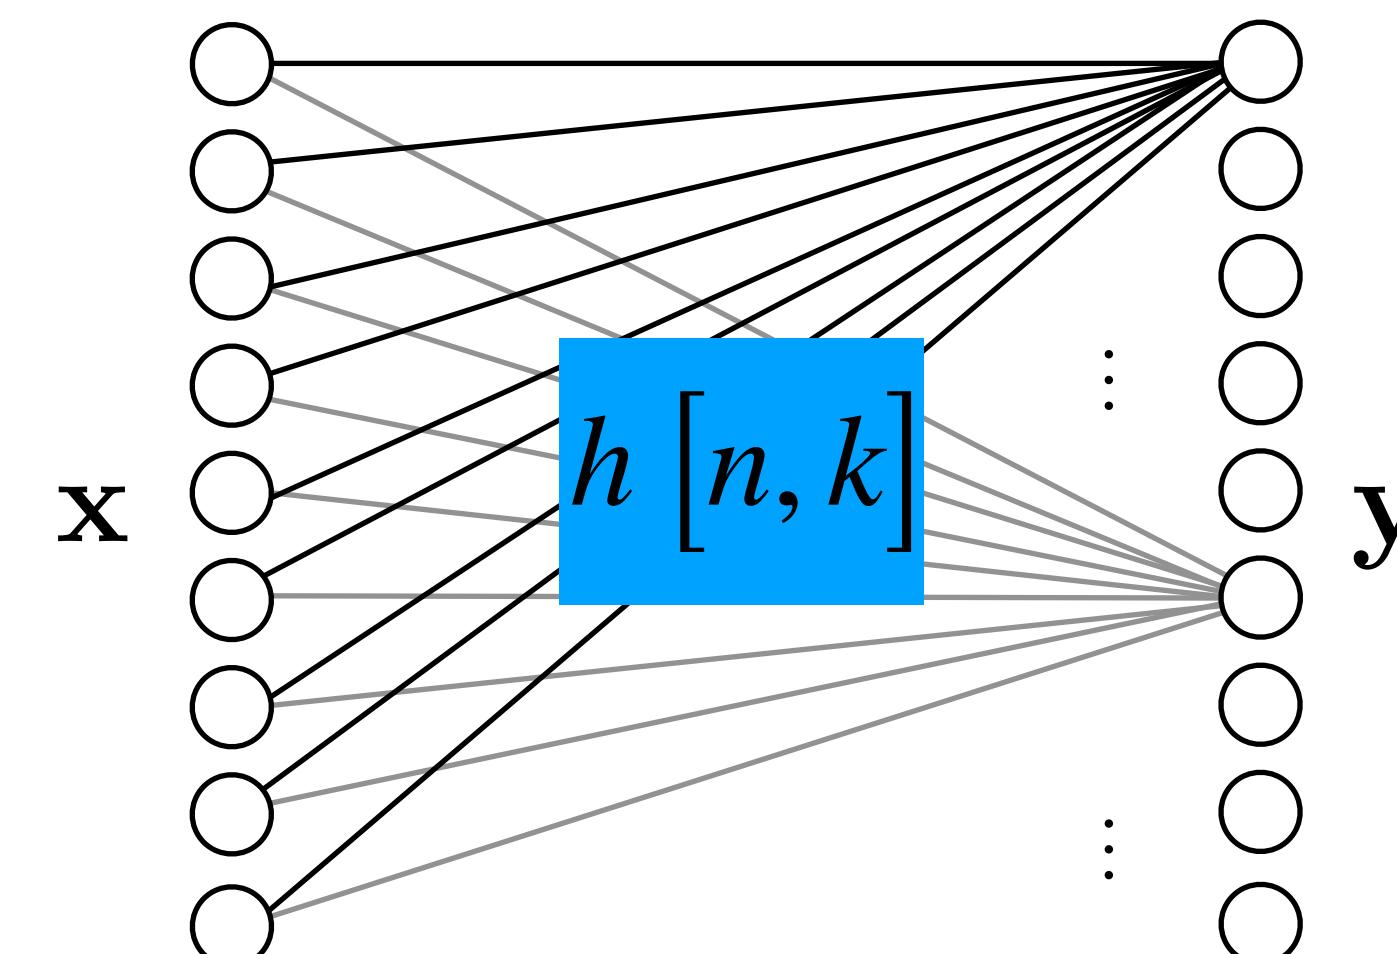
Linear system: $y = f(x)$

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

A 10x10 grid representing a matrix. The grid has alternating light gray and white squares. A blue box highlights the element at position $[n, k]$.

n indexes rows,
k indexes columns

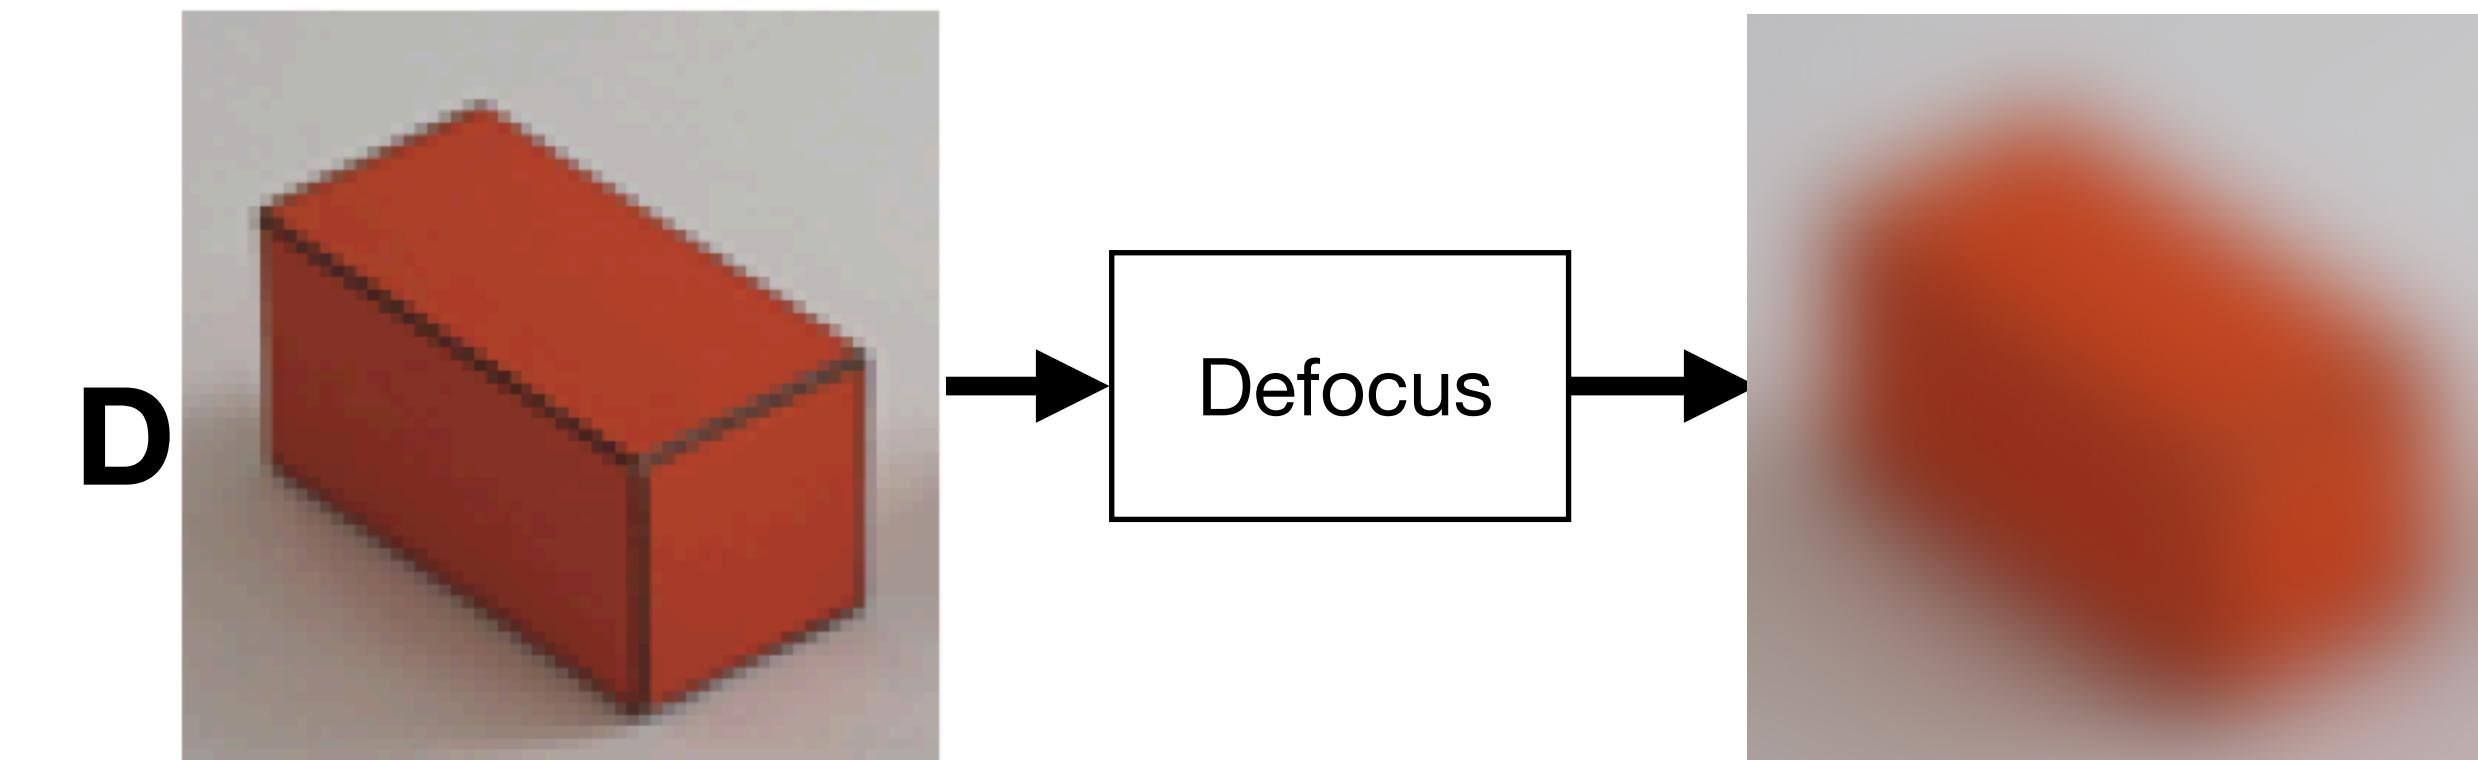
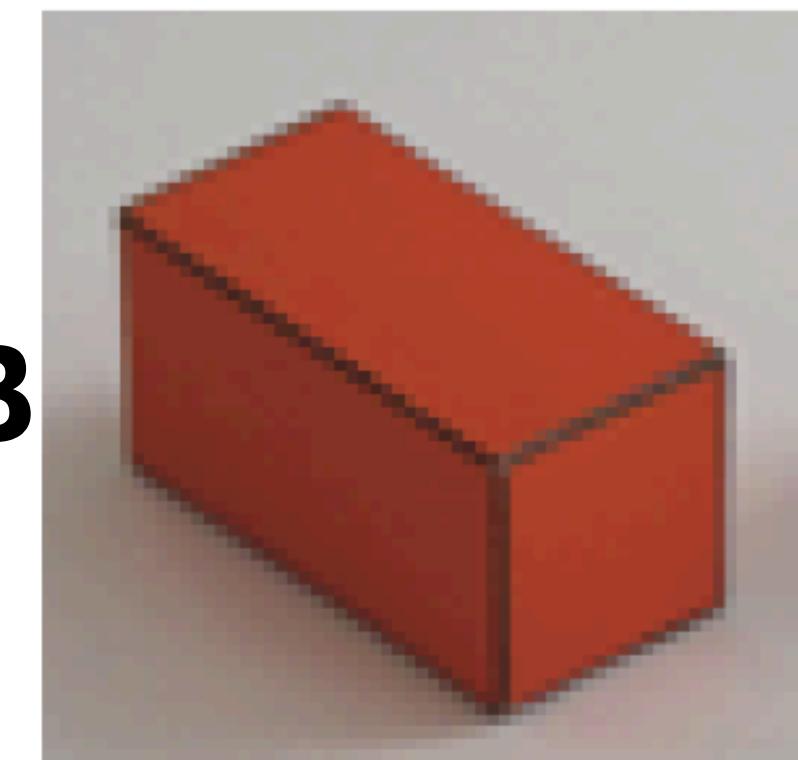
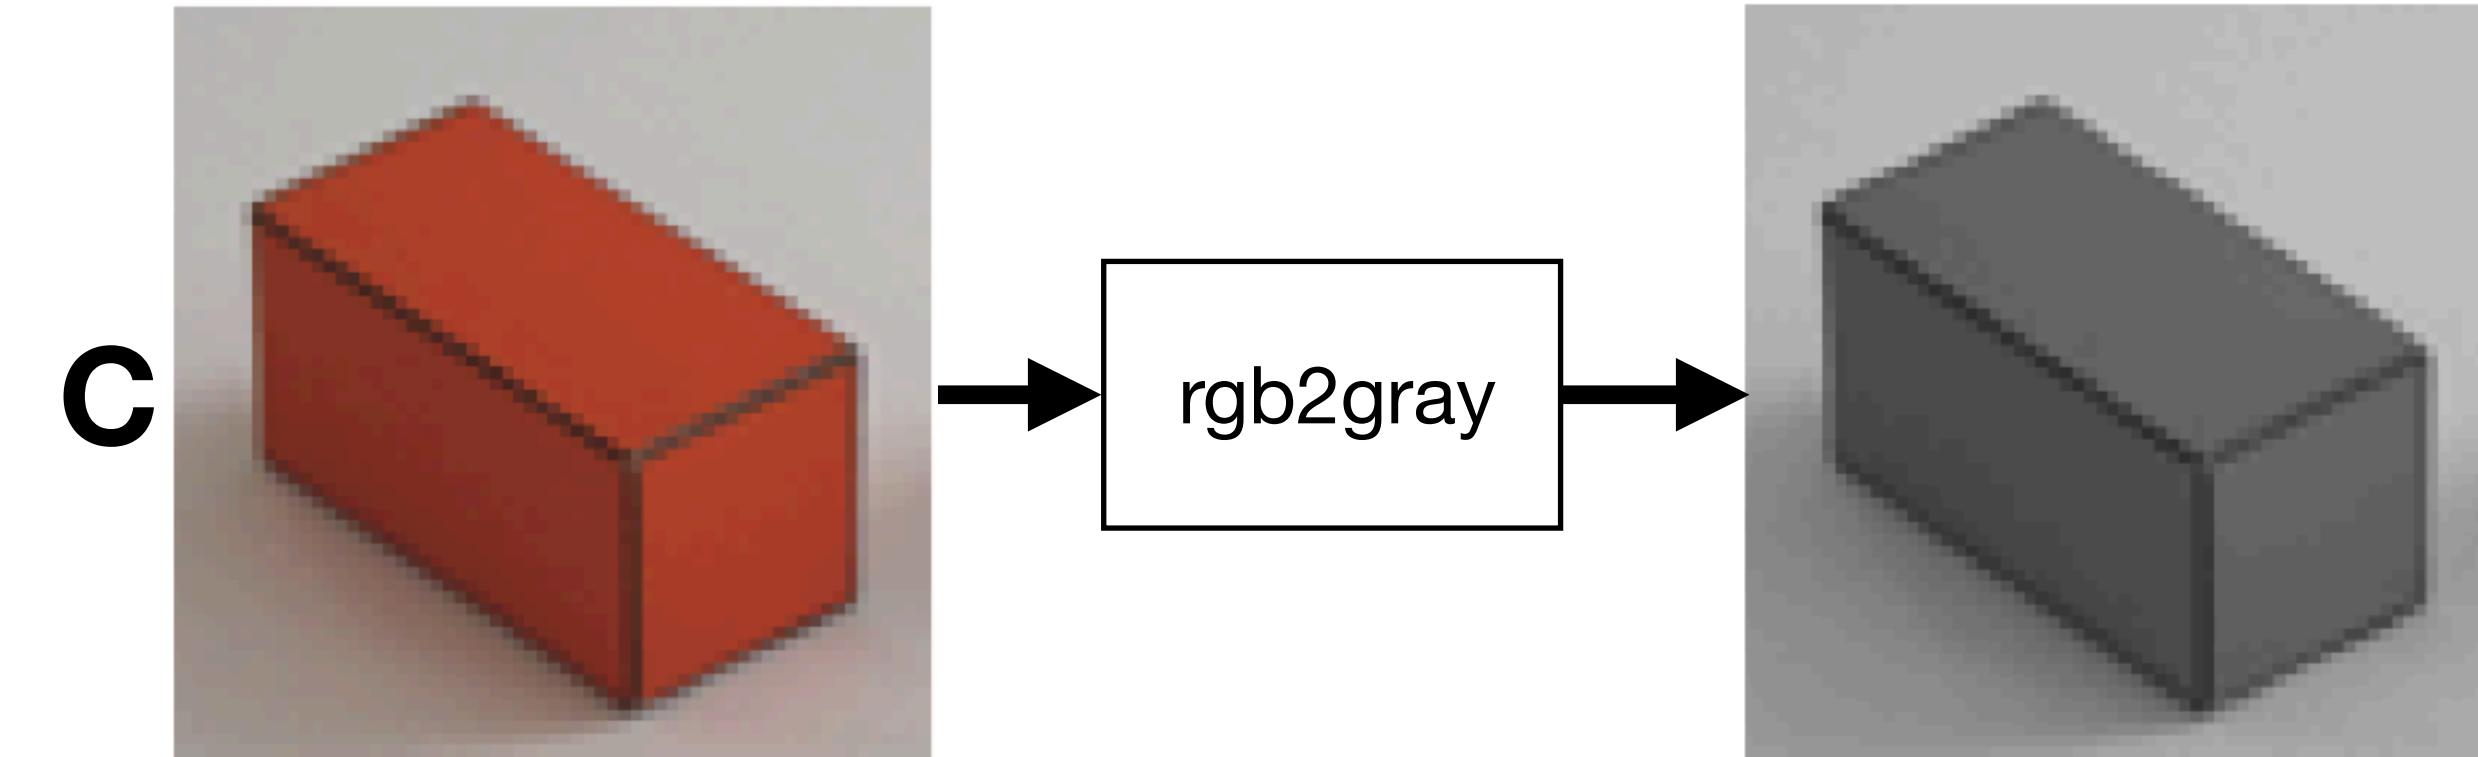
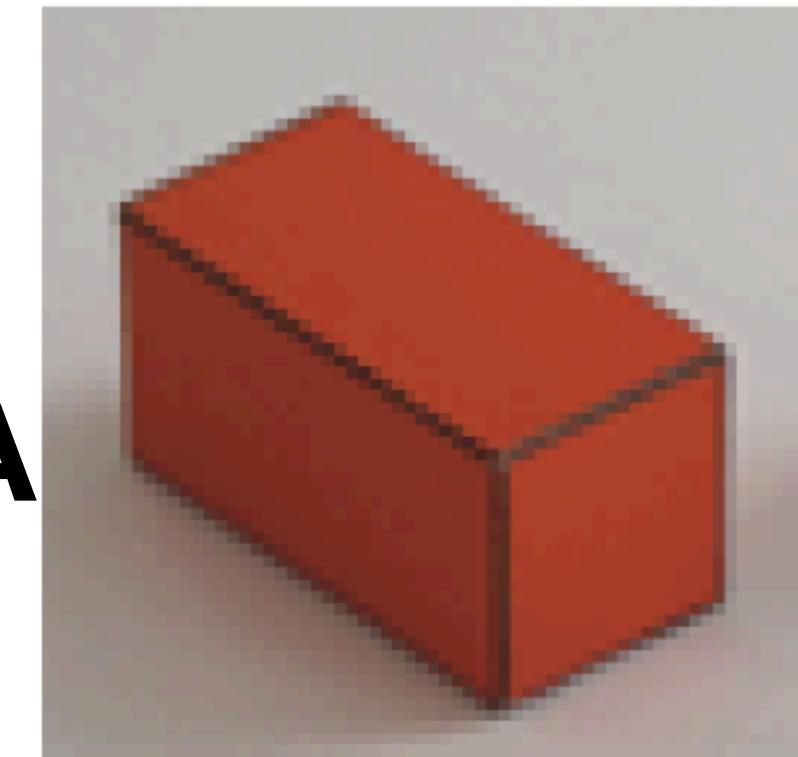
It can also be represented as a fully connected linear neural network

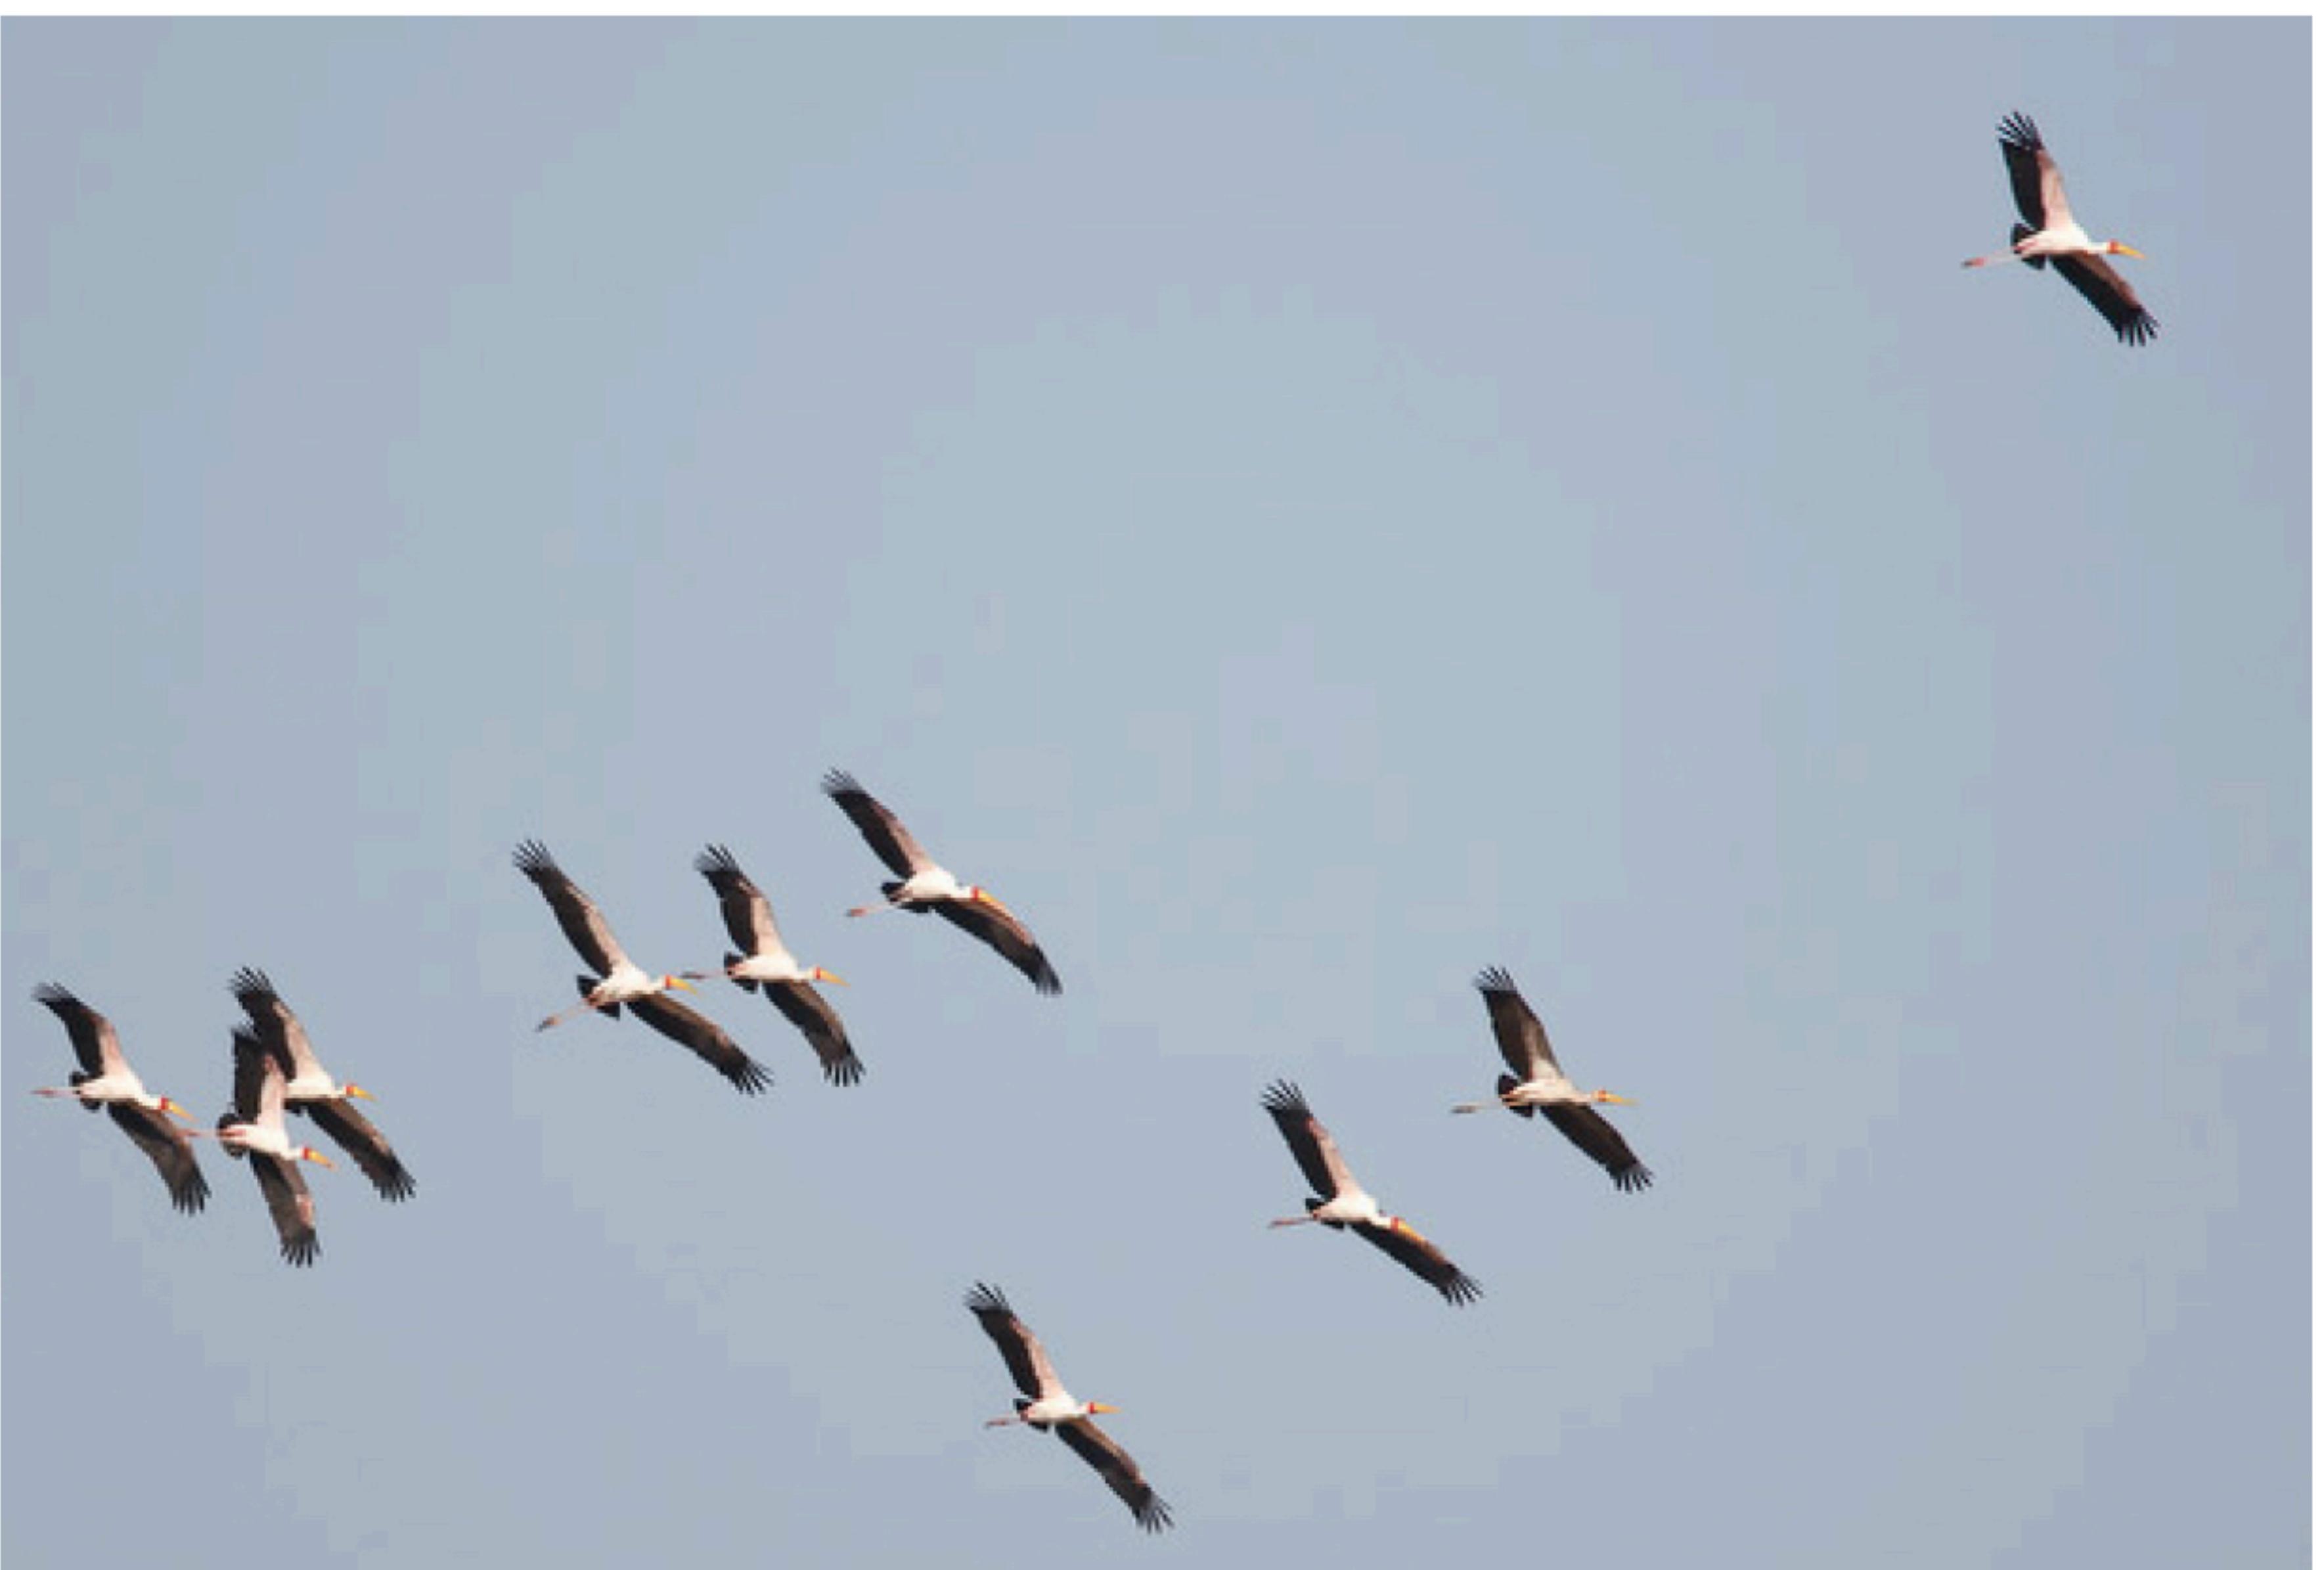


$h[n, k]$ Is the strength of the connection between $x[k]$ and $y[n]$

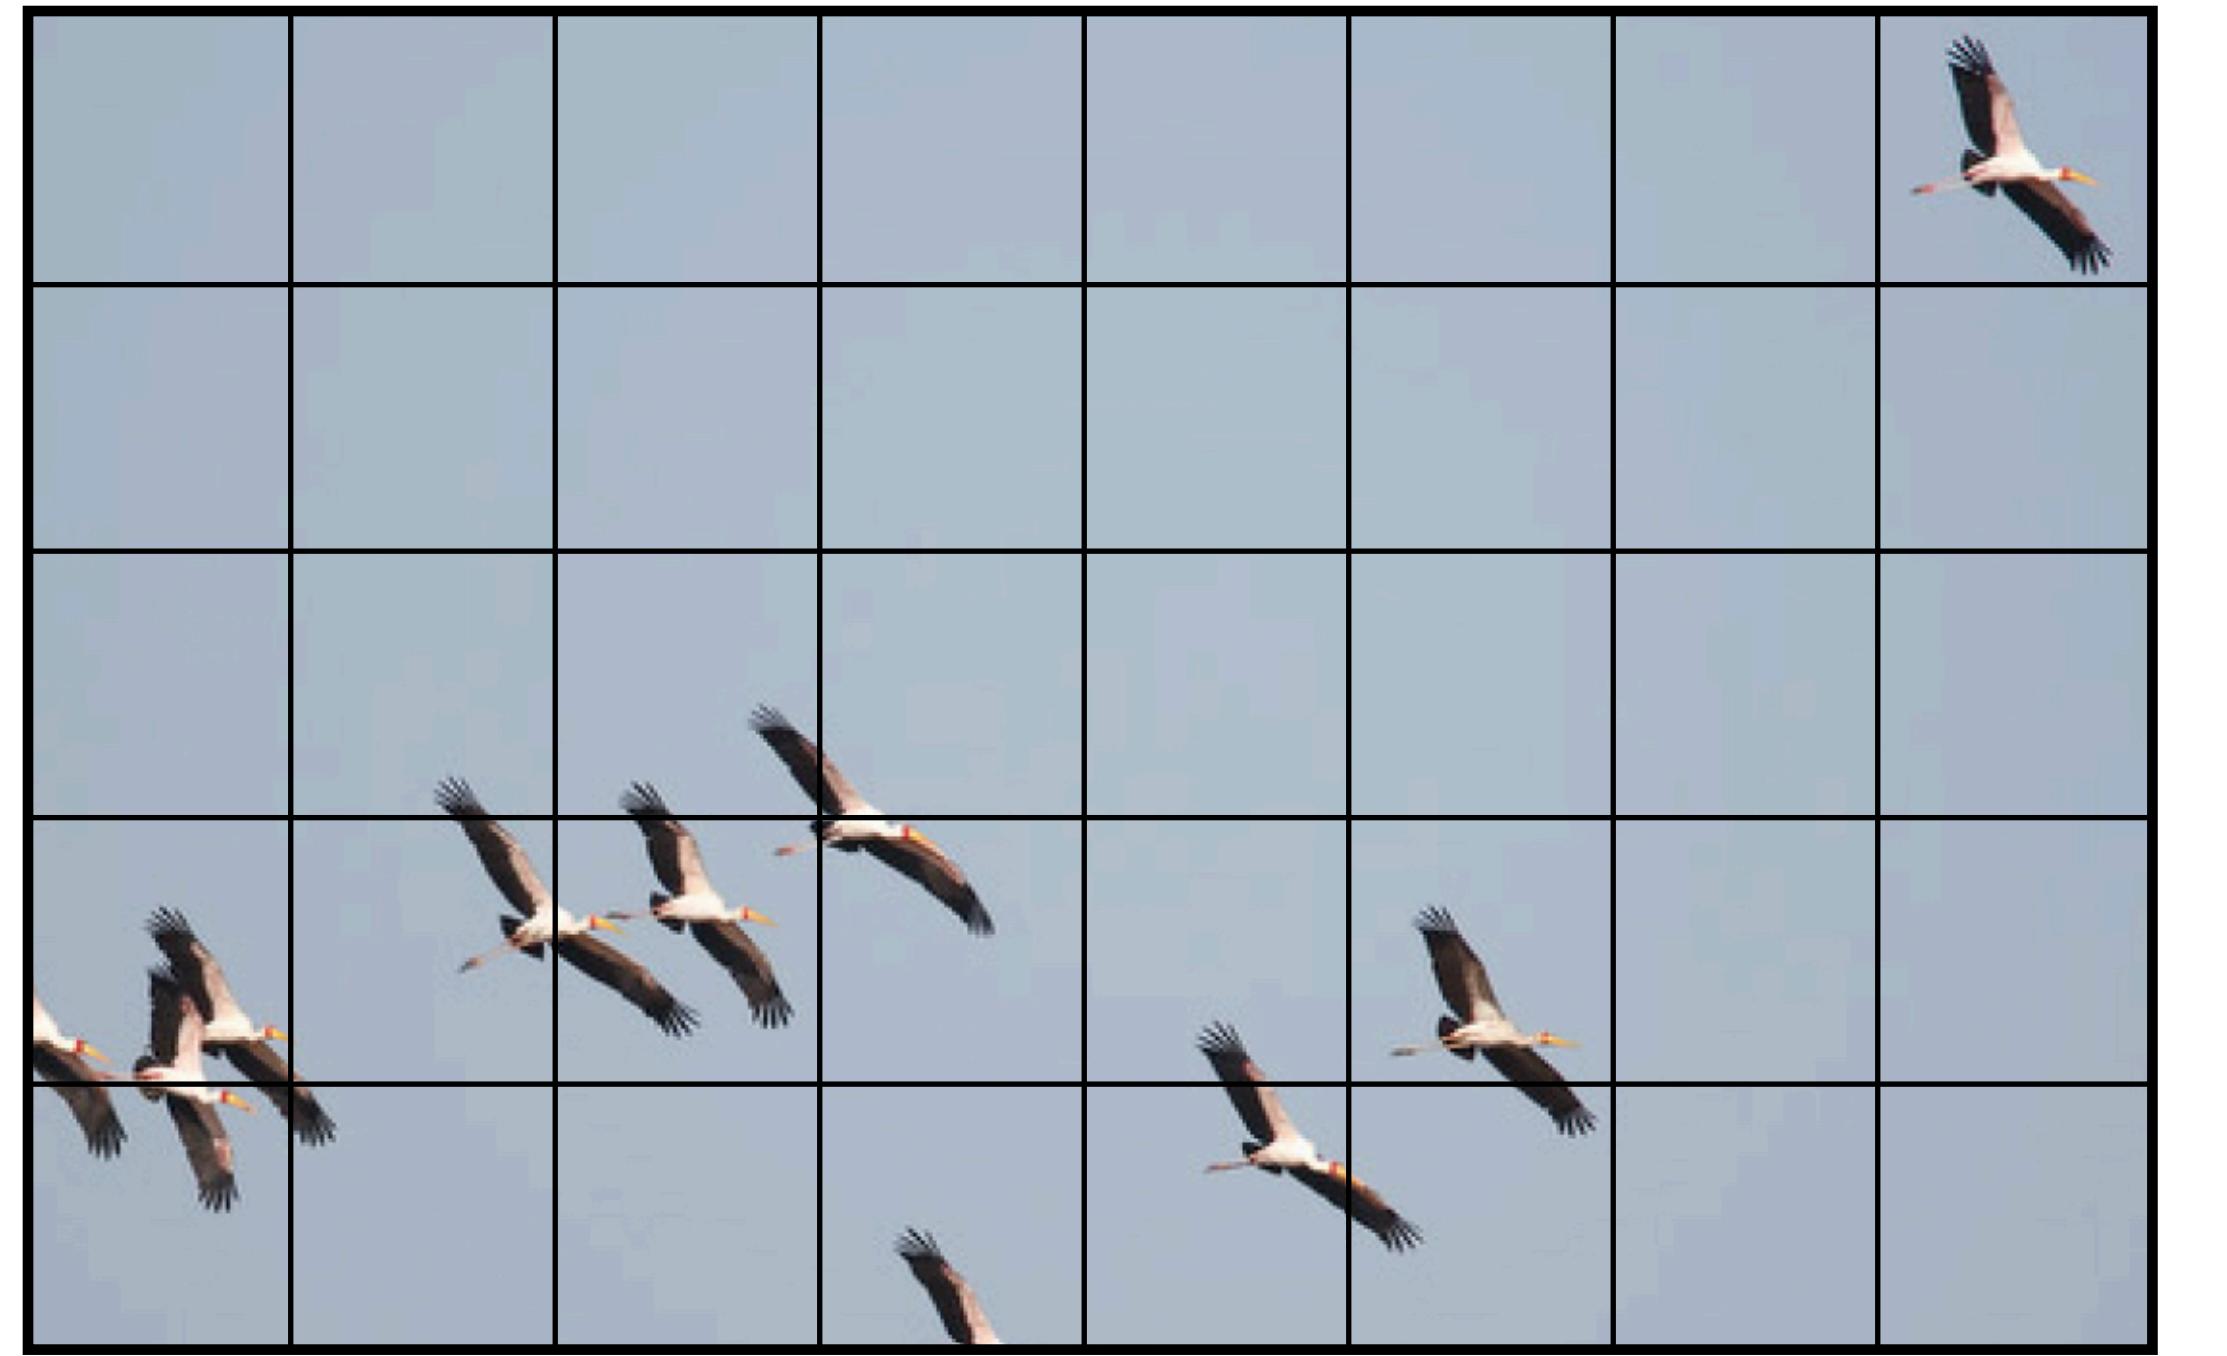
Quiz: what operation is linear?

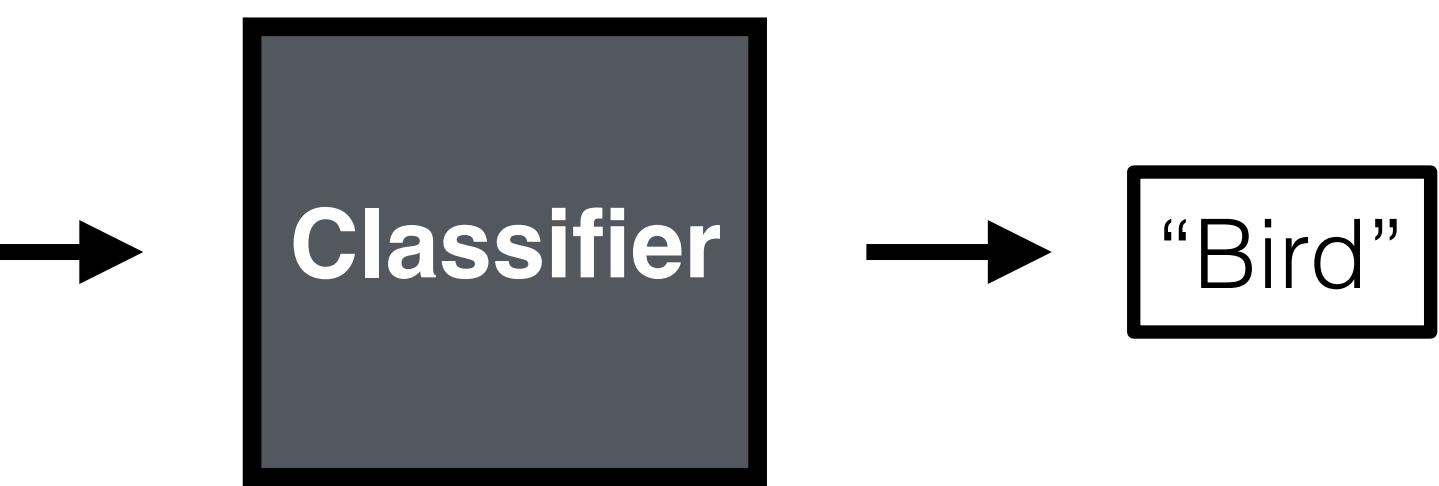
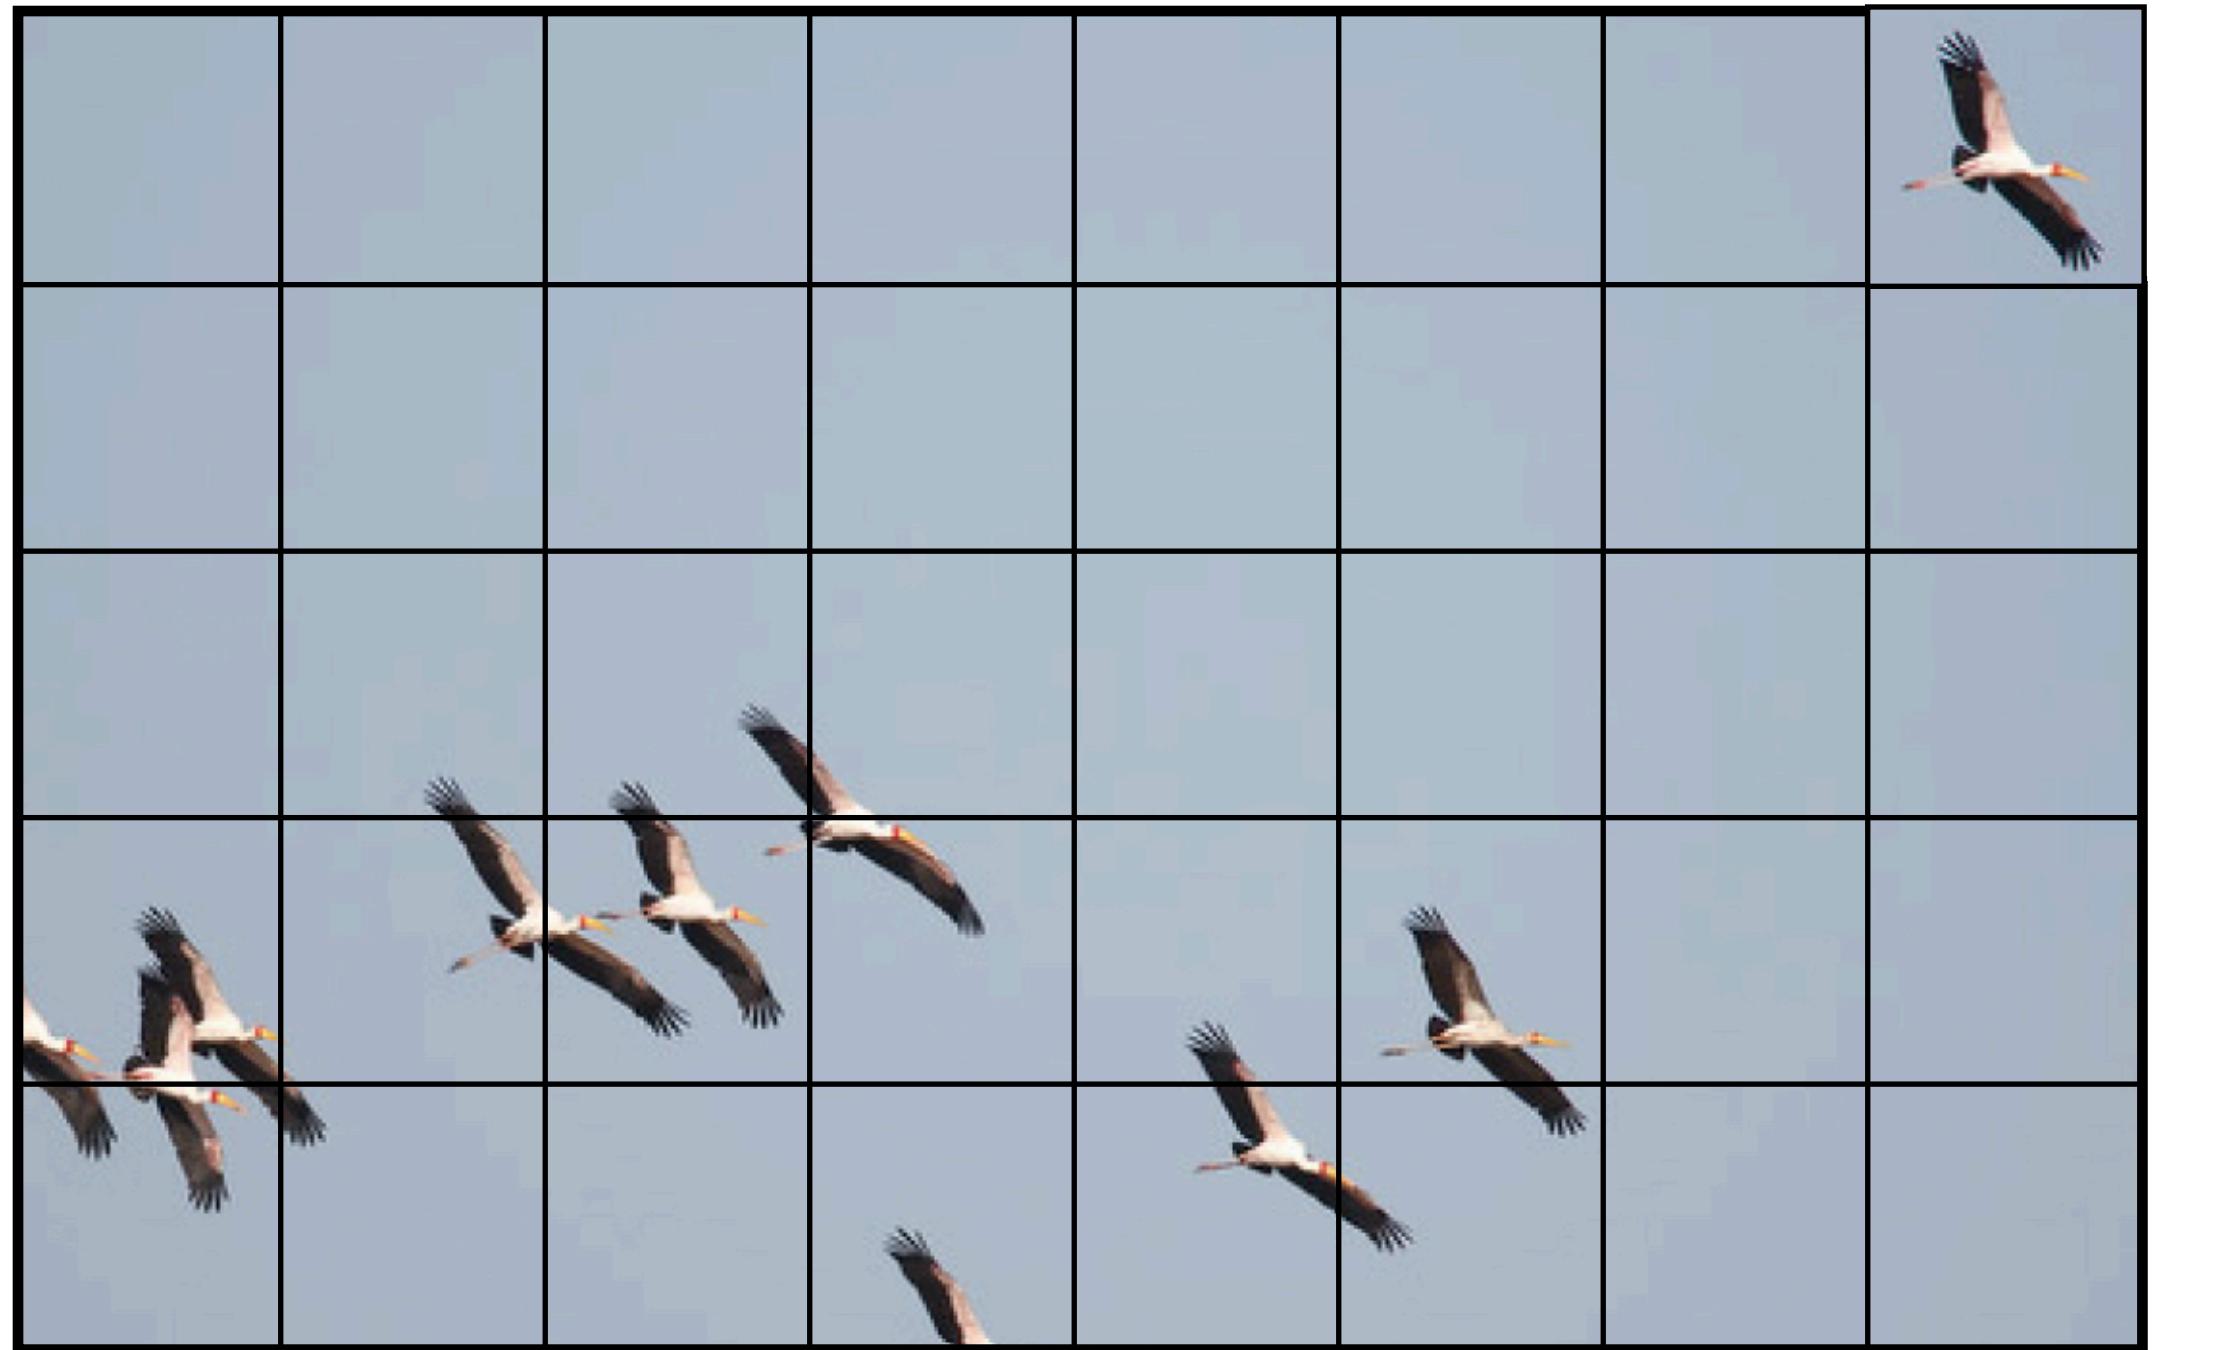
Quiz: what operation is linear?

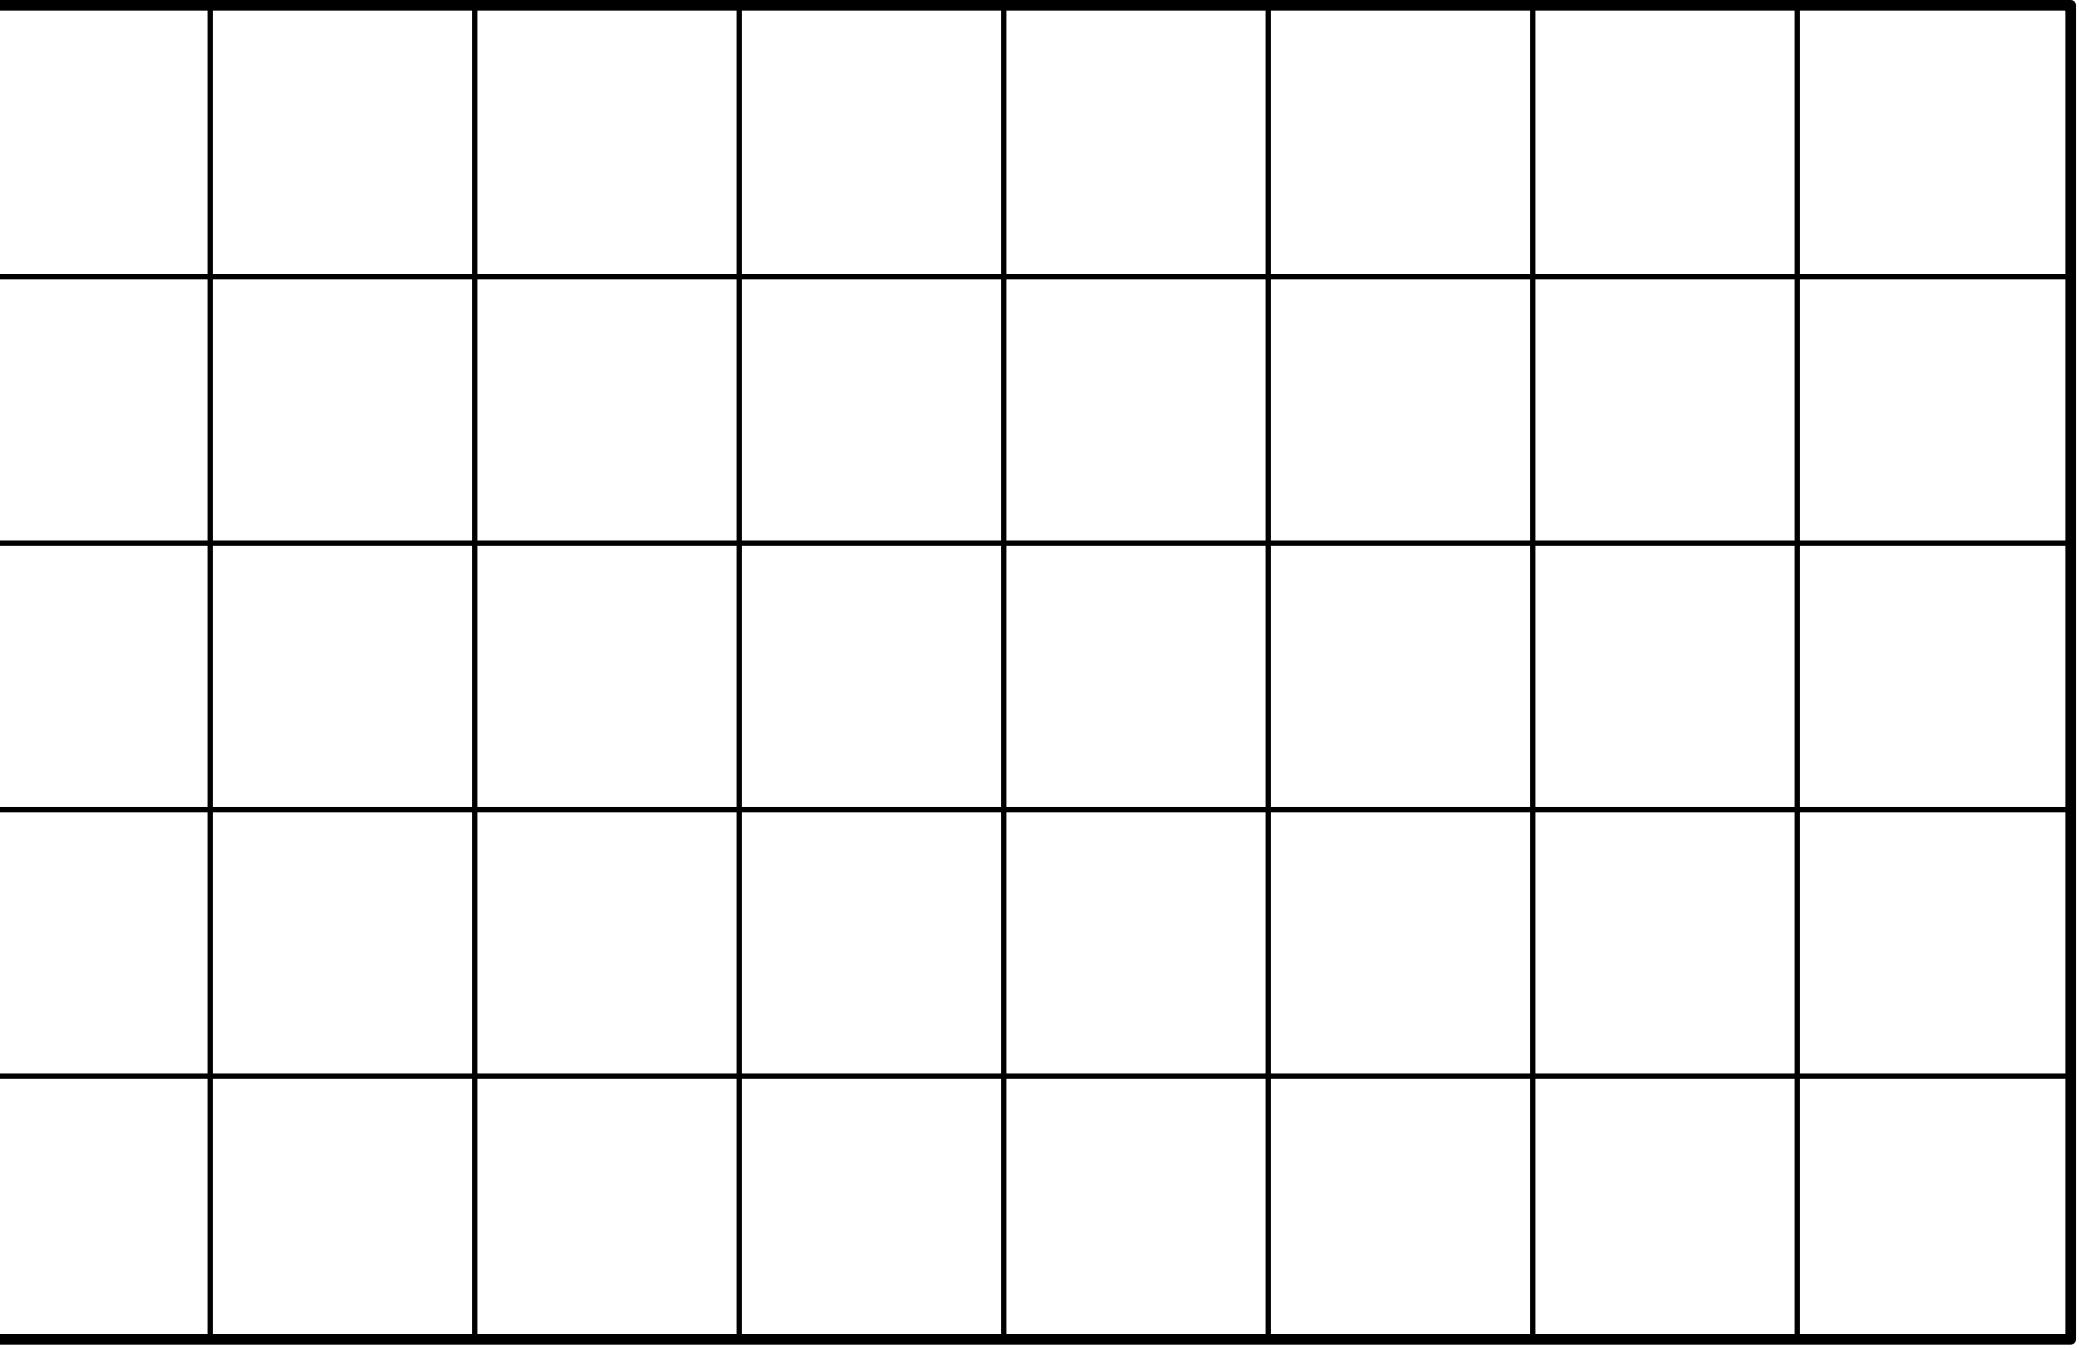
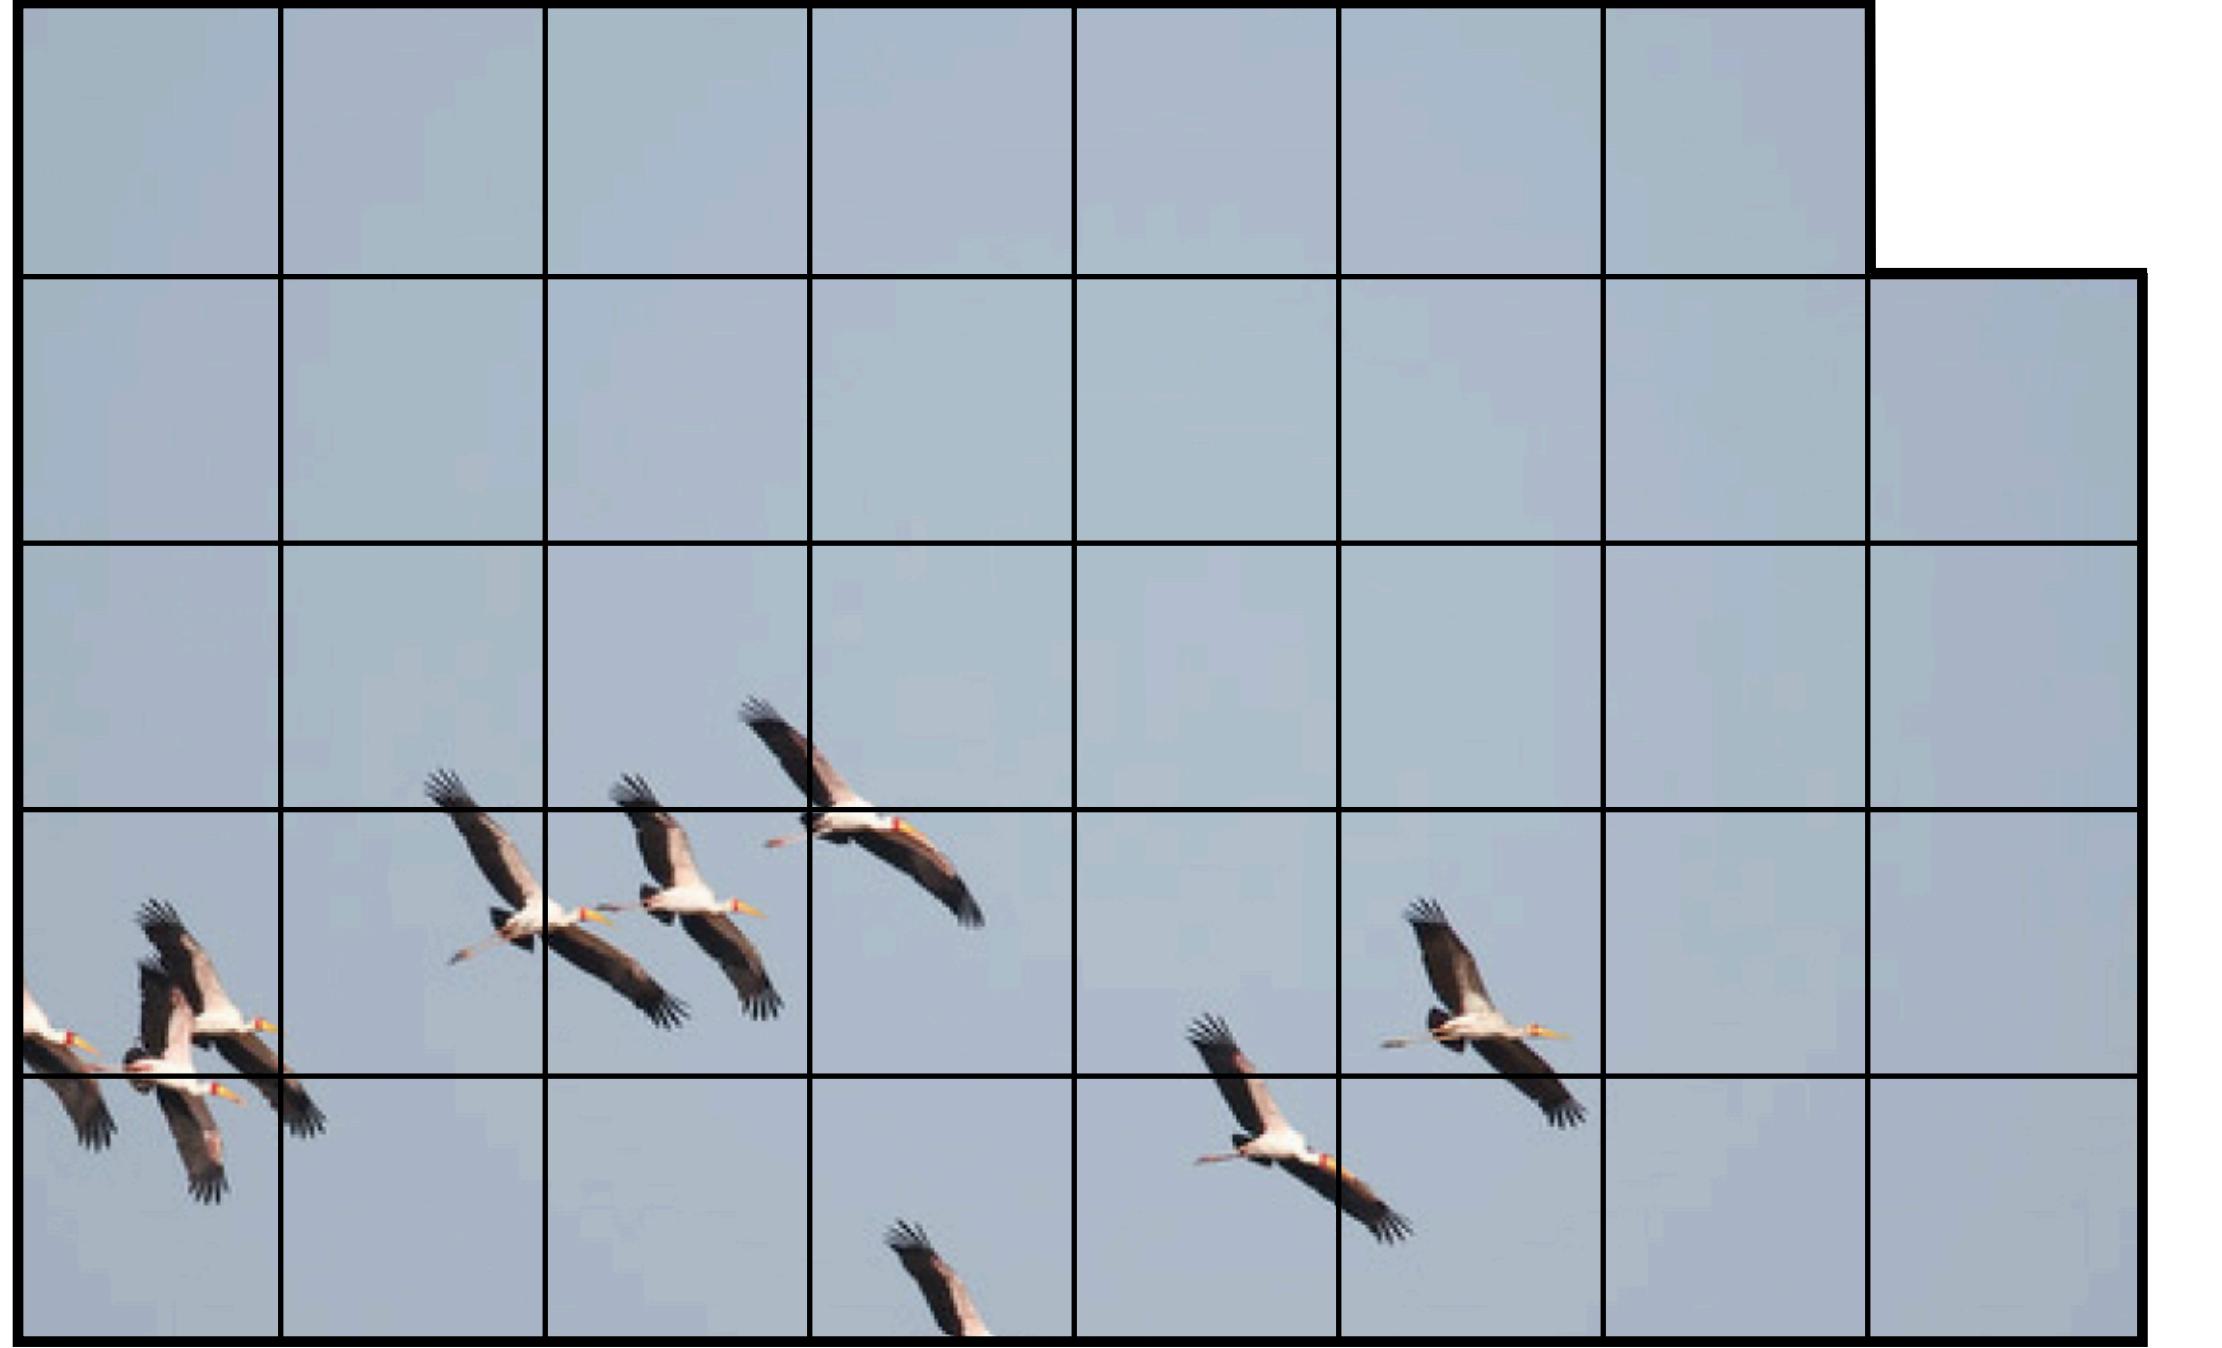


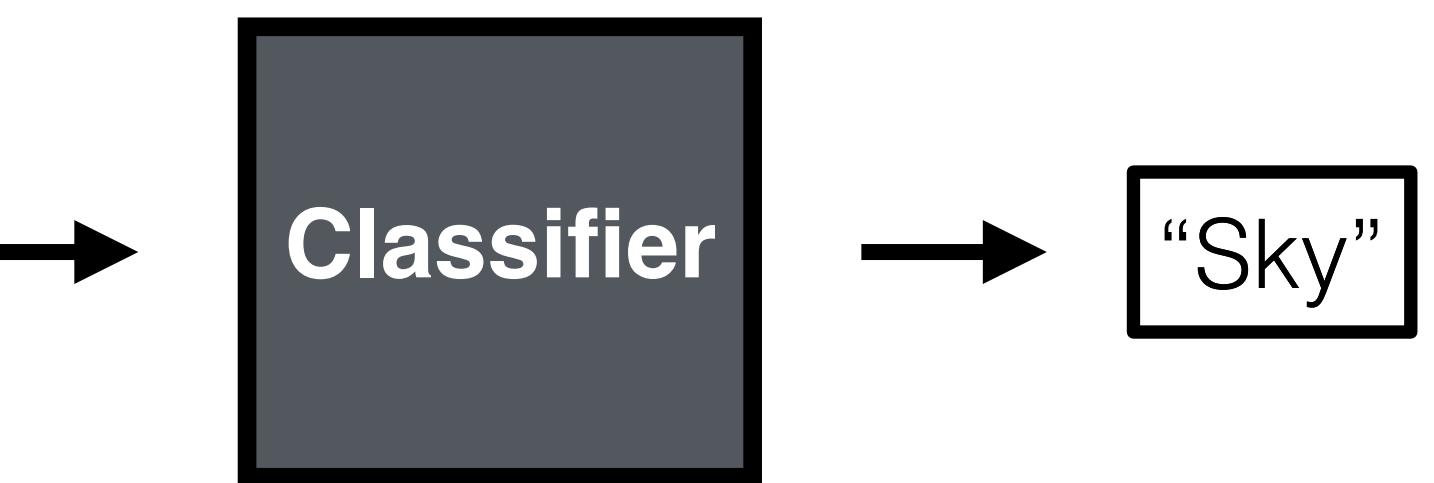
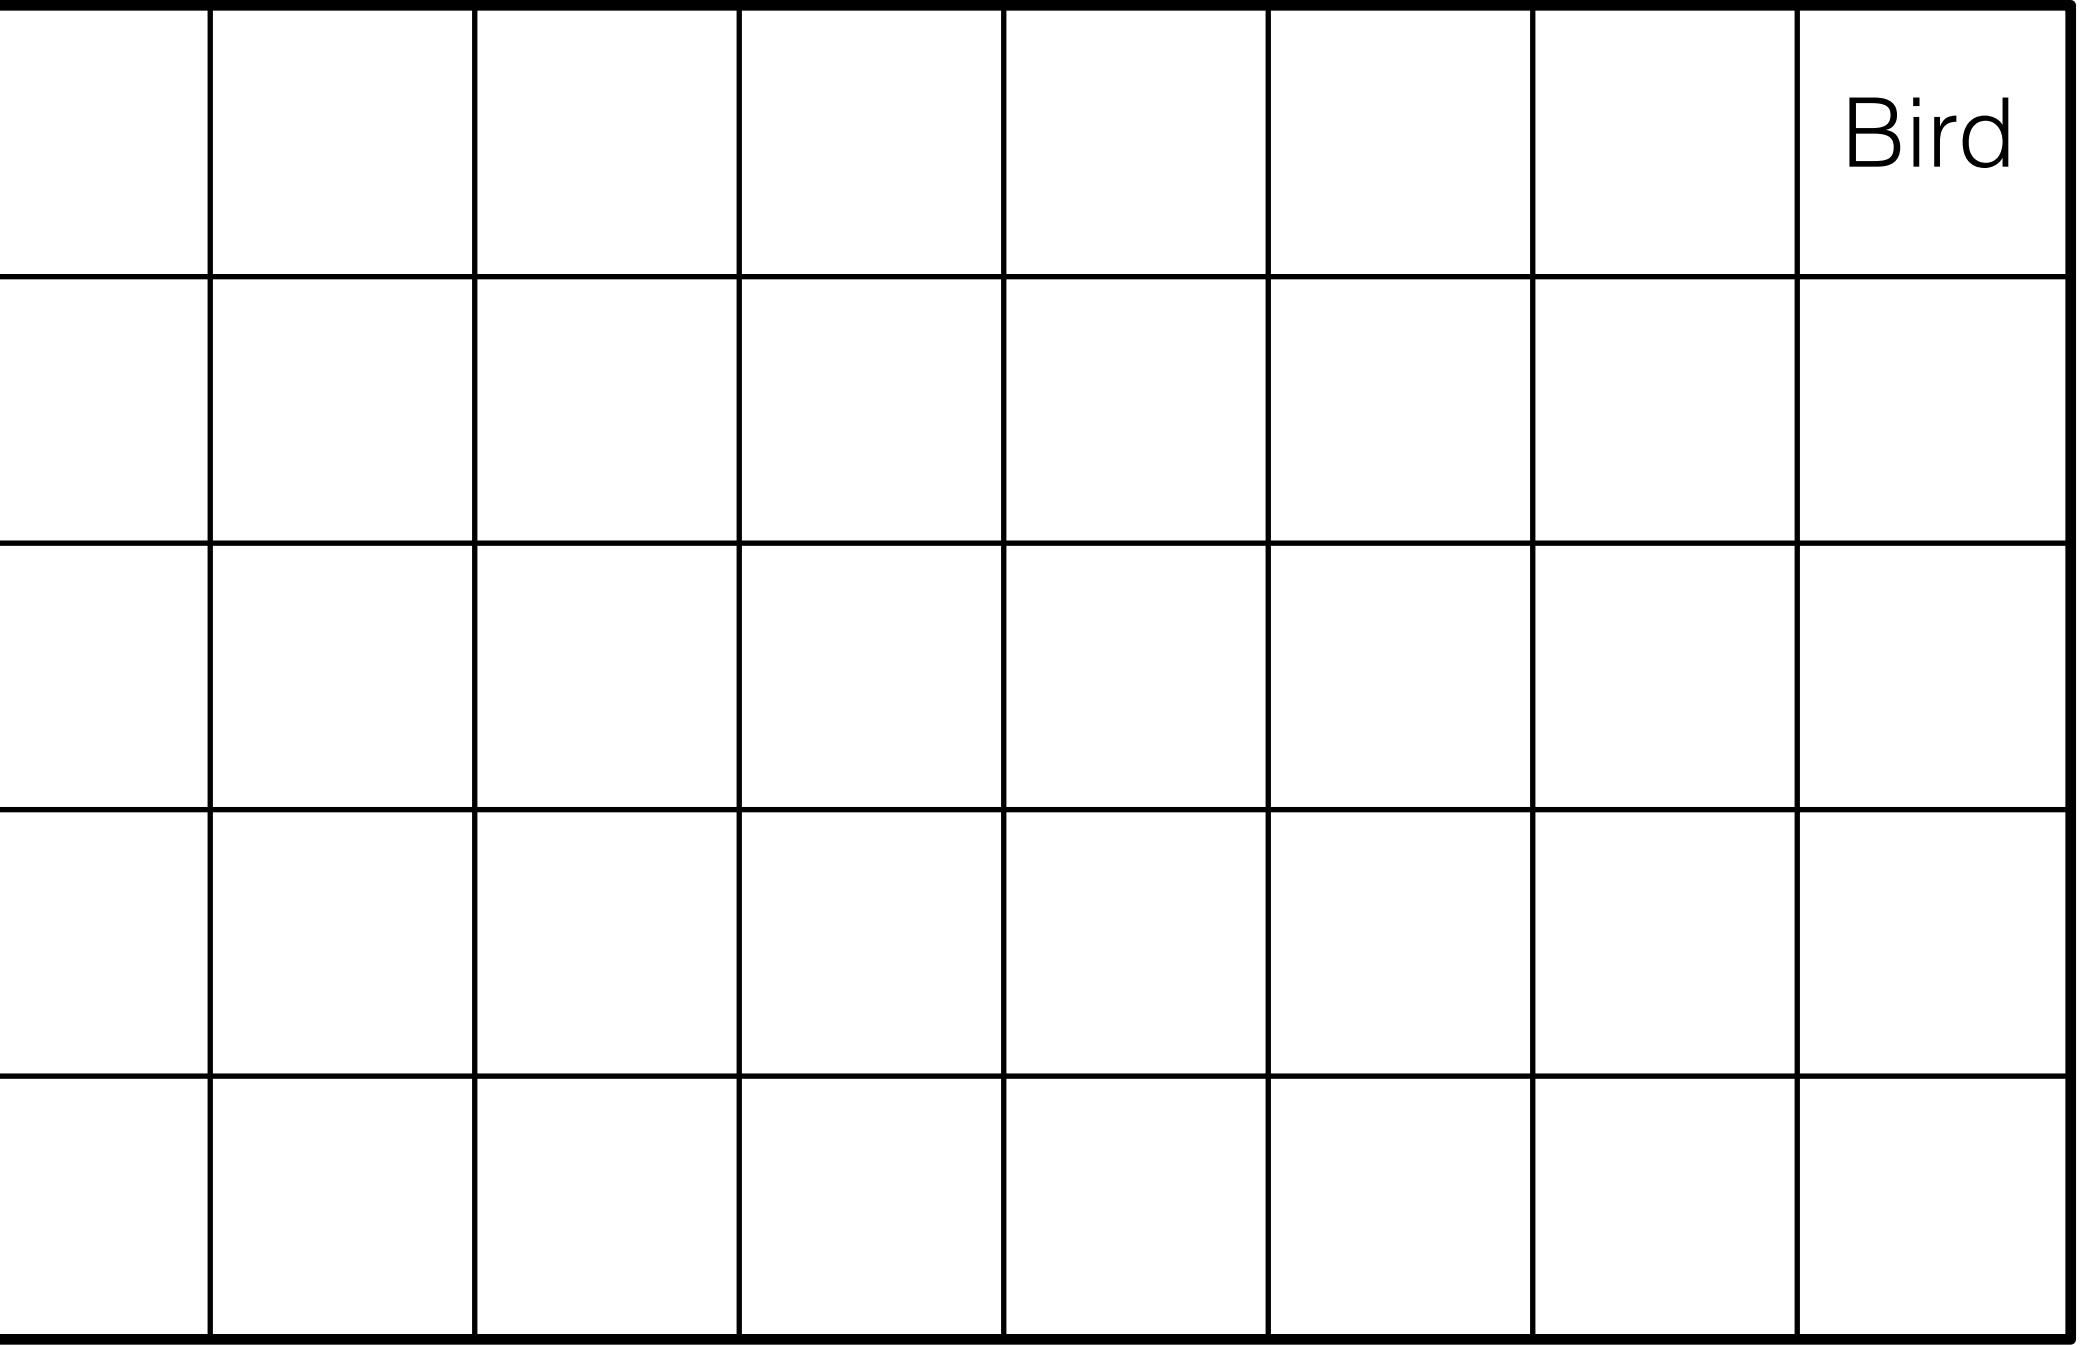
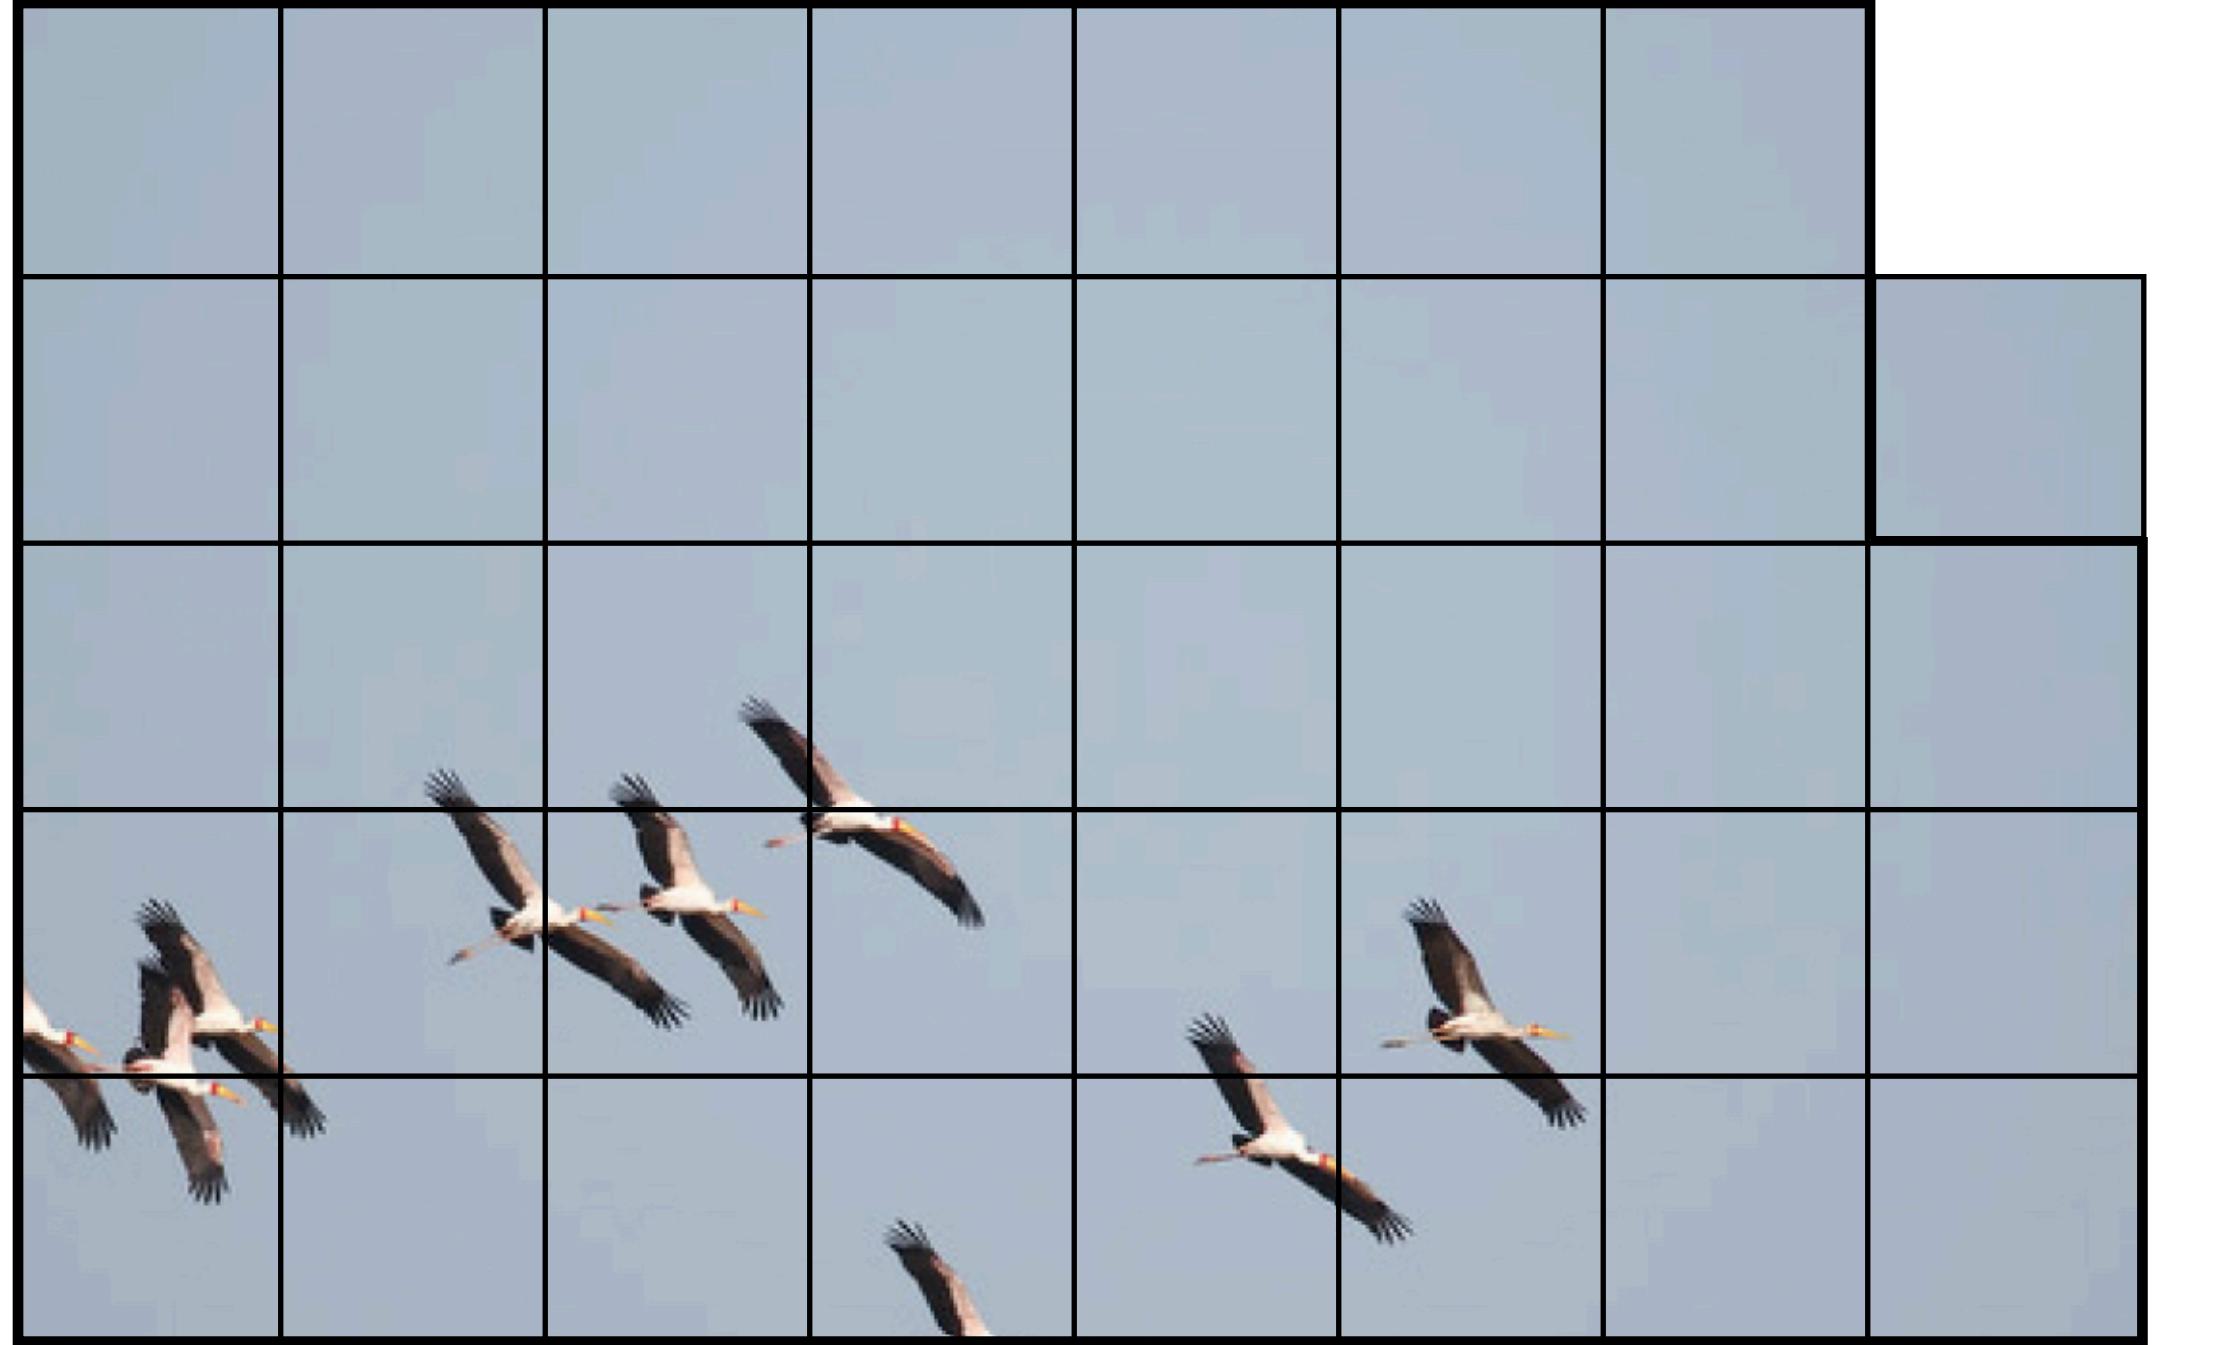


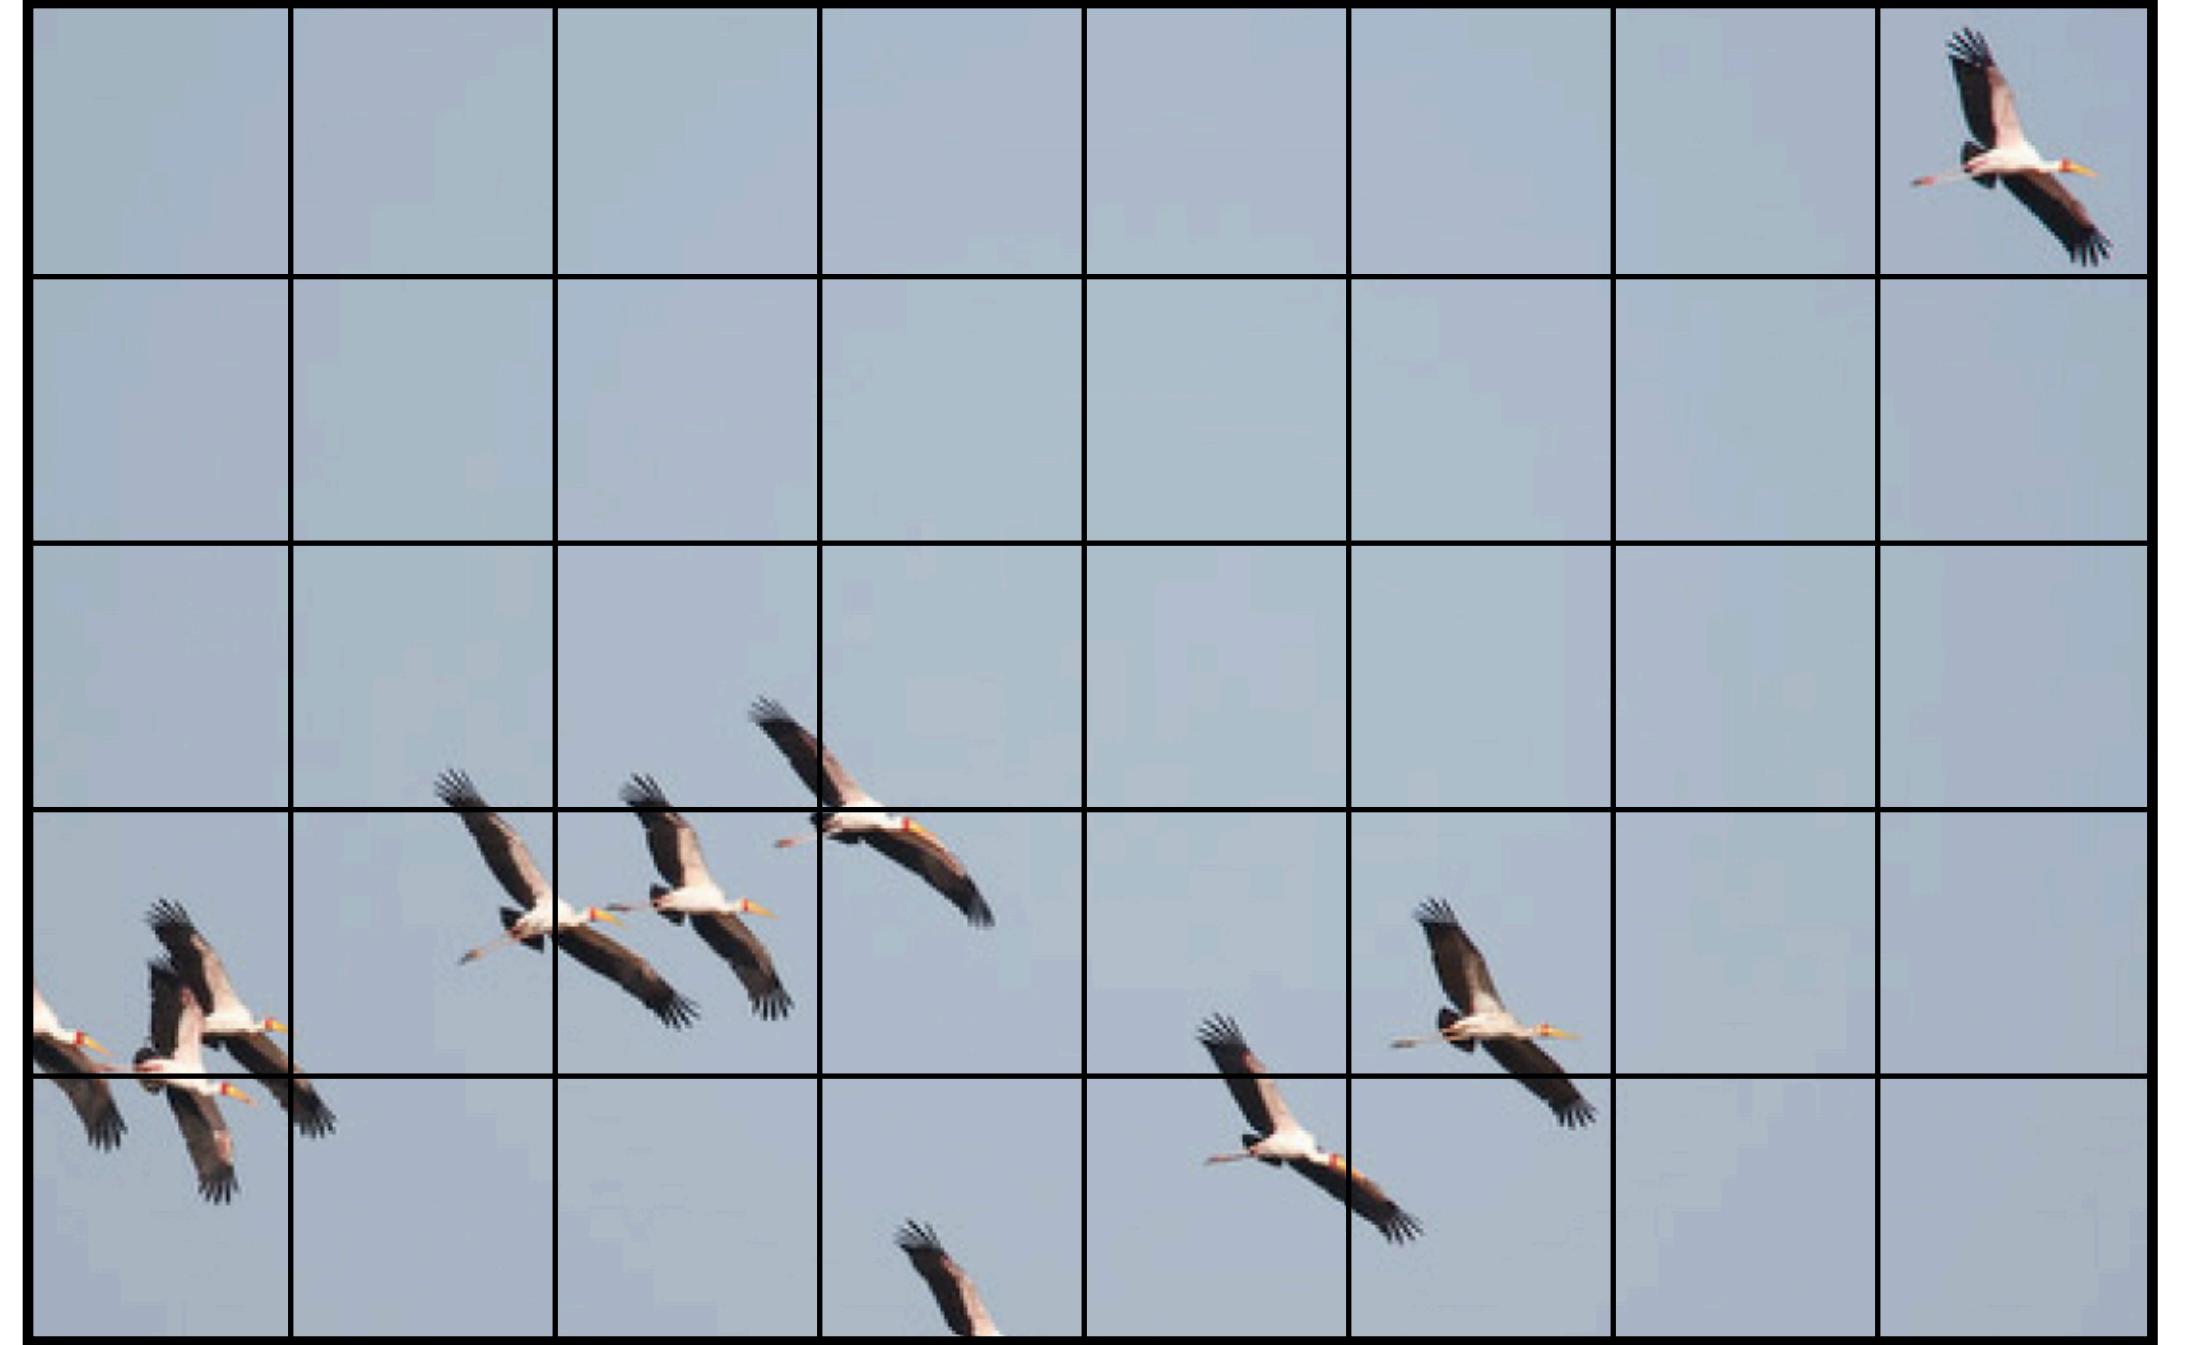
We need translation invariance



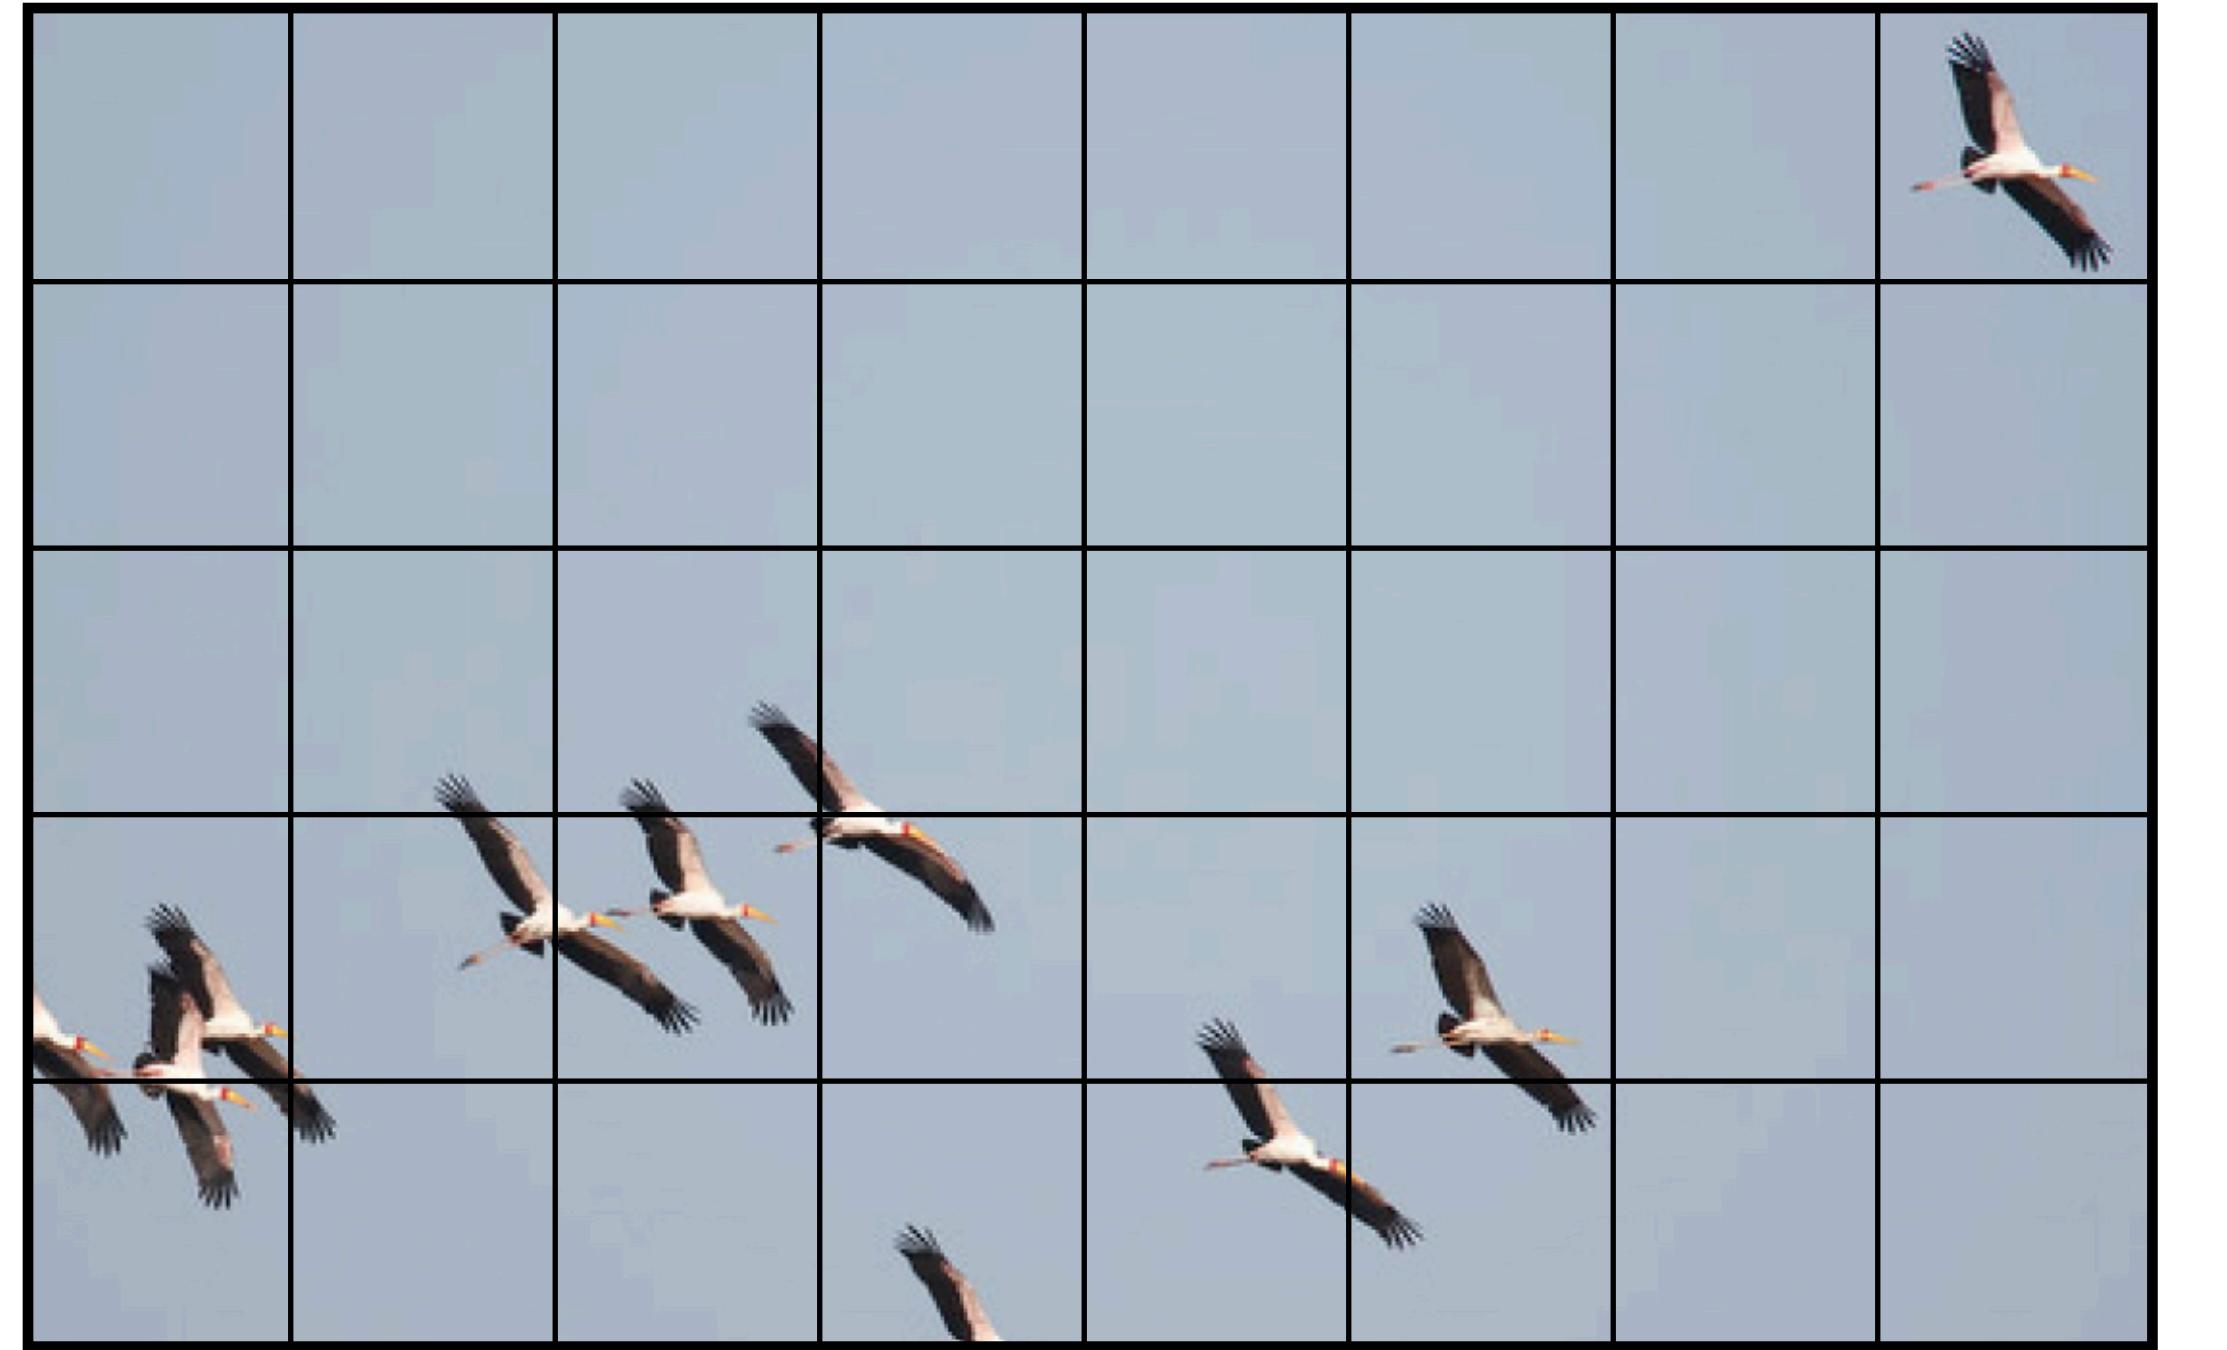




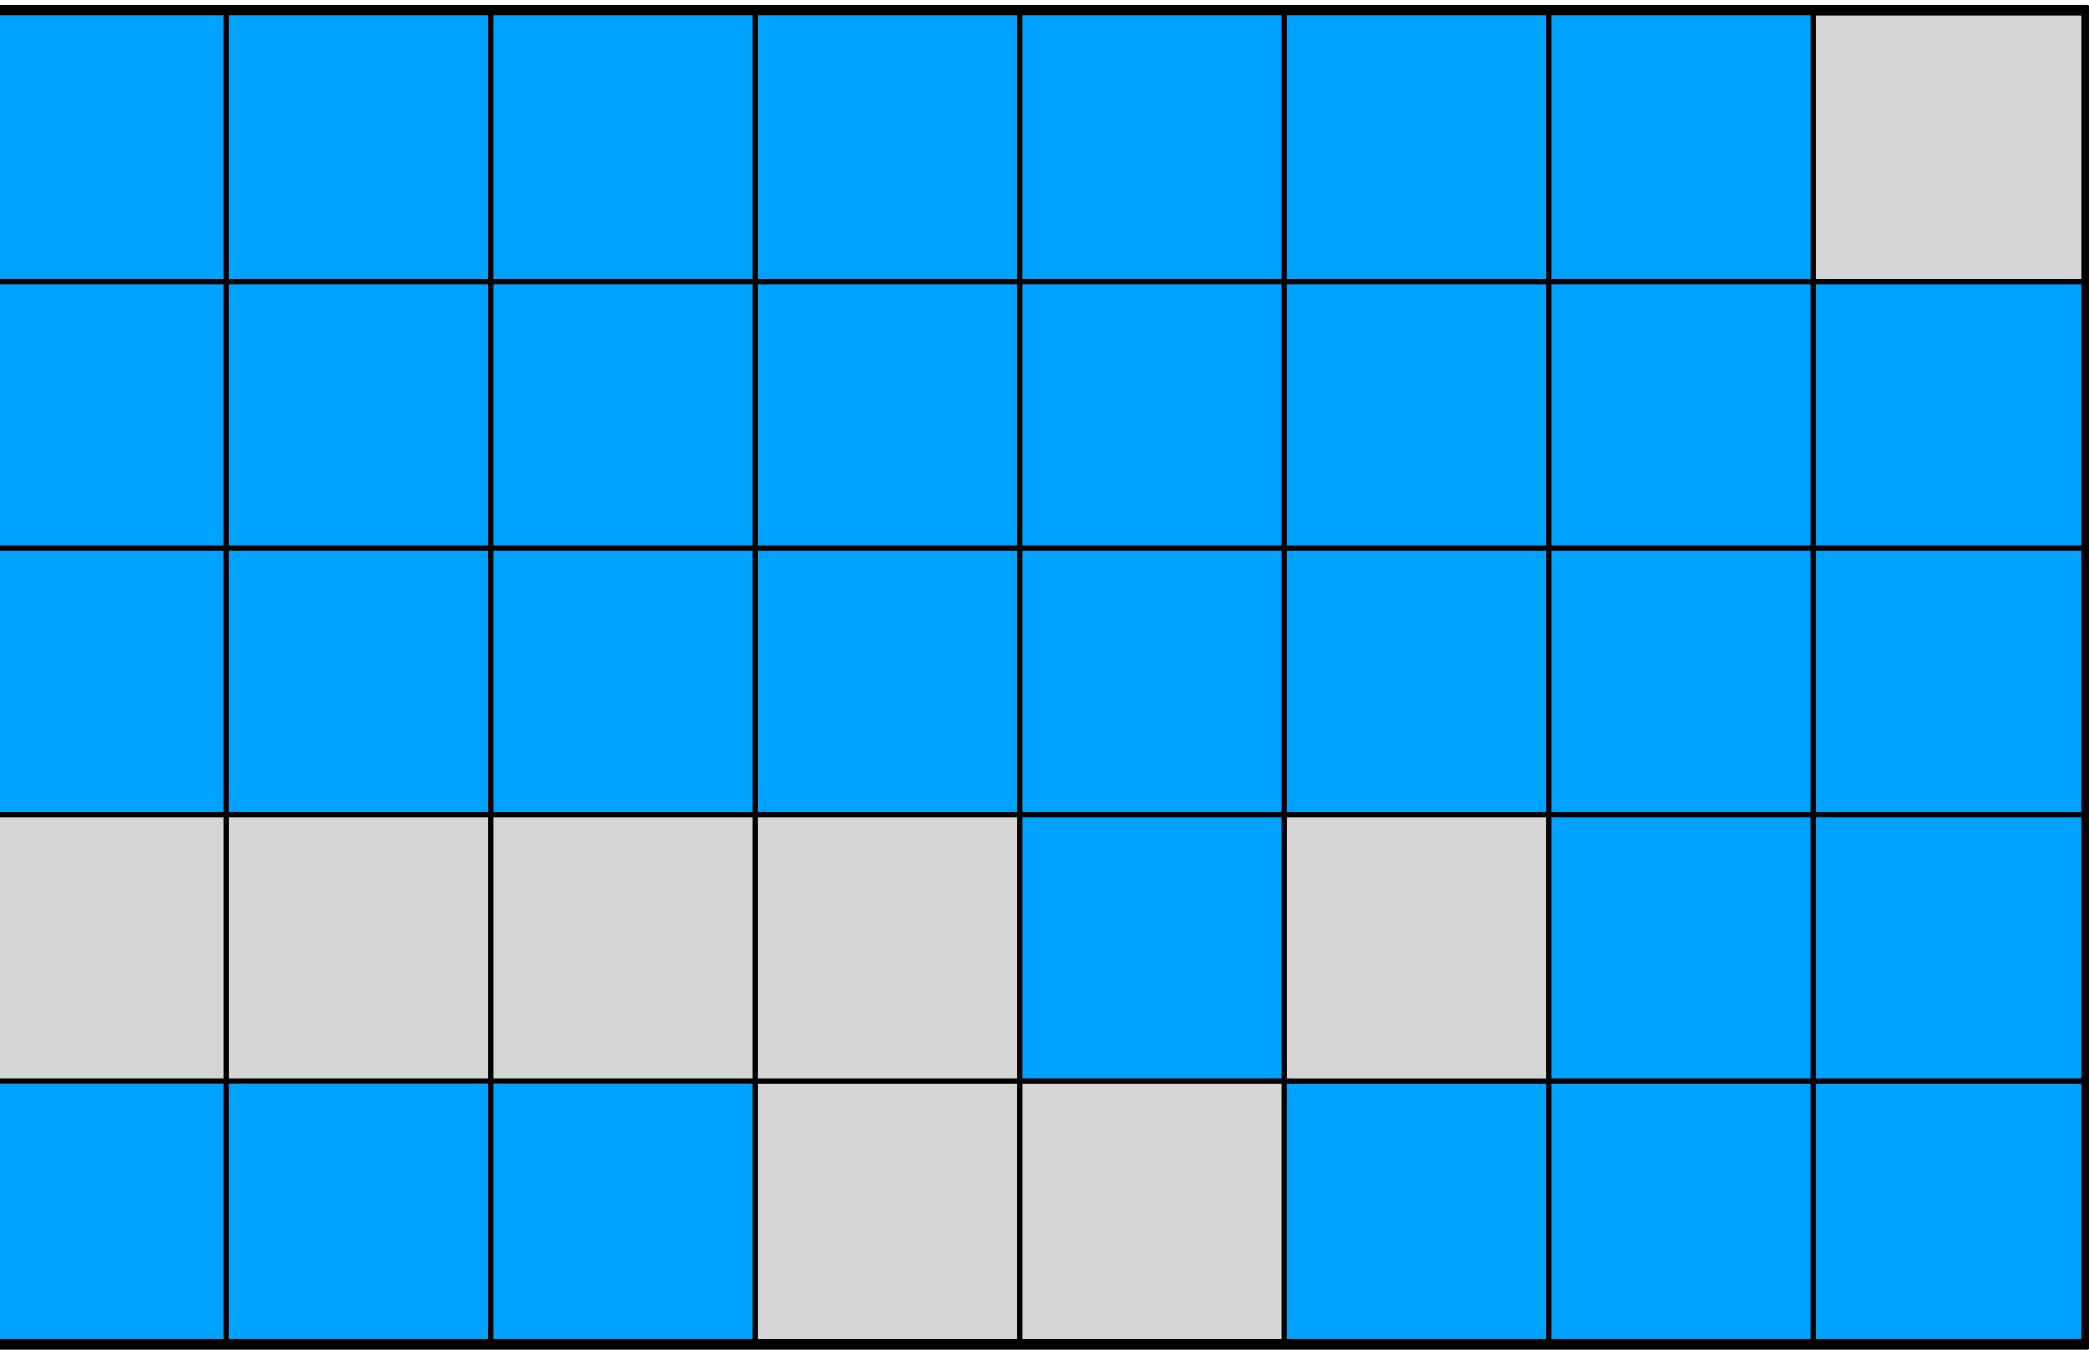




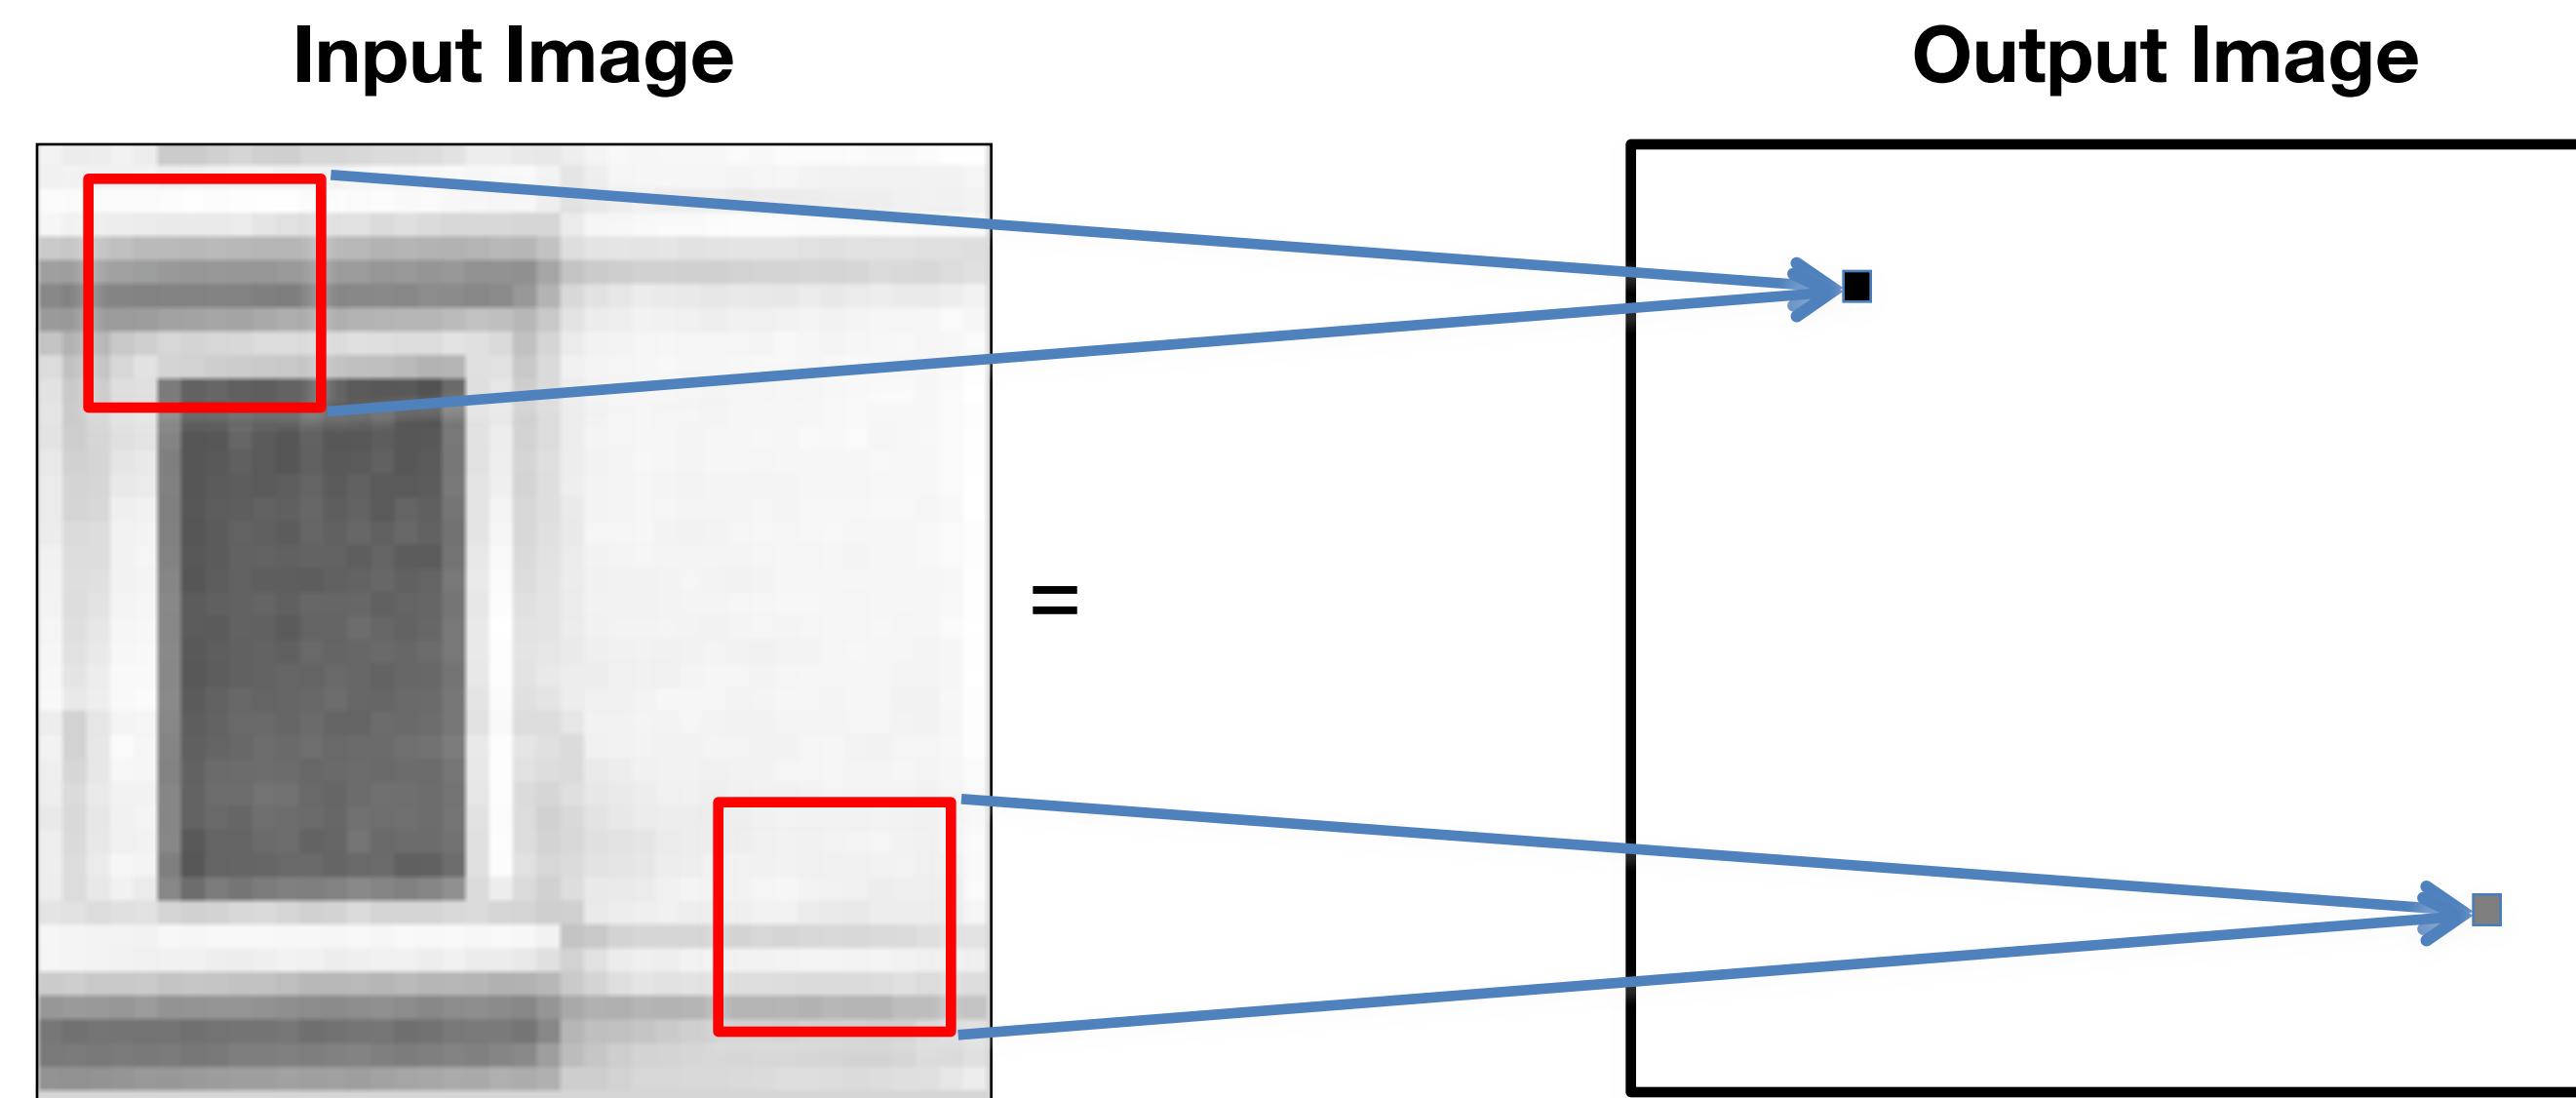
Sky	Sky	Sky	Sky	Sky	Sky	Sky	Bird
Sky	Sky	Sky	Sky	Sky	Sky	Sky	Sky
Sky	Sky	Sky	Sky	Sky	Sky	Sky	Sky
Bird	Bird	Bird	Sky	Bird	Sky	Sky	Sky
Sky	Sky	Sky	Bird	Sky	Sky	Sky	Sky



$$f$$



Convolution



The same weighting occurs within each window

Convolution: running example

Convolution: running example

x

0	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	1	0
0	1	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

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

1

y

Convolution: running example

X

Convolution kernel

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

1

y

Convolution: running example

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

Convolution kernel

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

=

y

-3

Convolution: running example

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

Convolution kernel

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

=

y

-3

Convolution: running example

X

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

10

y

Convolution: running example

X

Convolution kernel

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

10

y

Convolution: running example

X

Convolution kernel

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

10 of 10

y

Convolution: running example

Convolution kernel

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

10

Convolution: running example

Convolution kernel

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

10

Convolution: running example

X

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

10 of 10

y

Convolution: running example

x	0	0	0	0	0	0	0	0
0	0	0	0	0	0	1	1	0
0	1	1	1	1	1	1	1	0
0	1	1	1	1	1	1	1	0
0	1	1	1	1	1	1	1	0
0	0	1	1	1	1	01	02	01
0	0	1	1	1	1	00	00	00
0	0	0	0	0	0	-1	-2	-1

Convolution kernel

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

=

	-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

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

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

Convolution kernel

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

=

?								
	-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

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

x

Convolution kernel

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

=

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

y

Convolution vs cross-correlation

Convolution

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

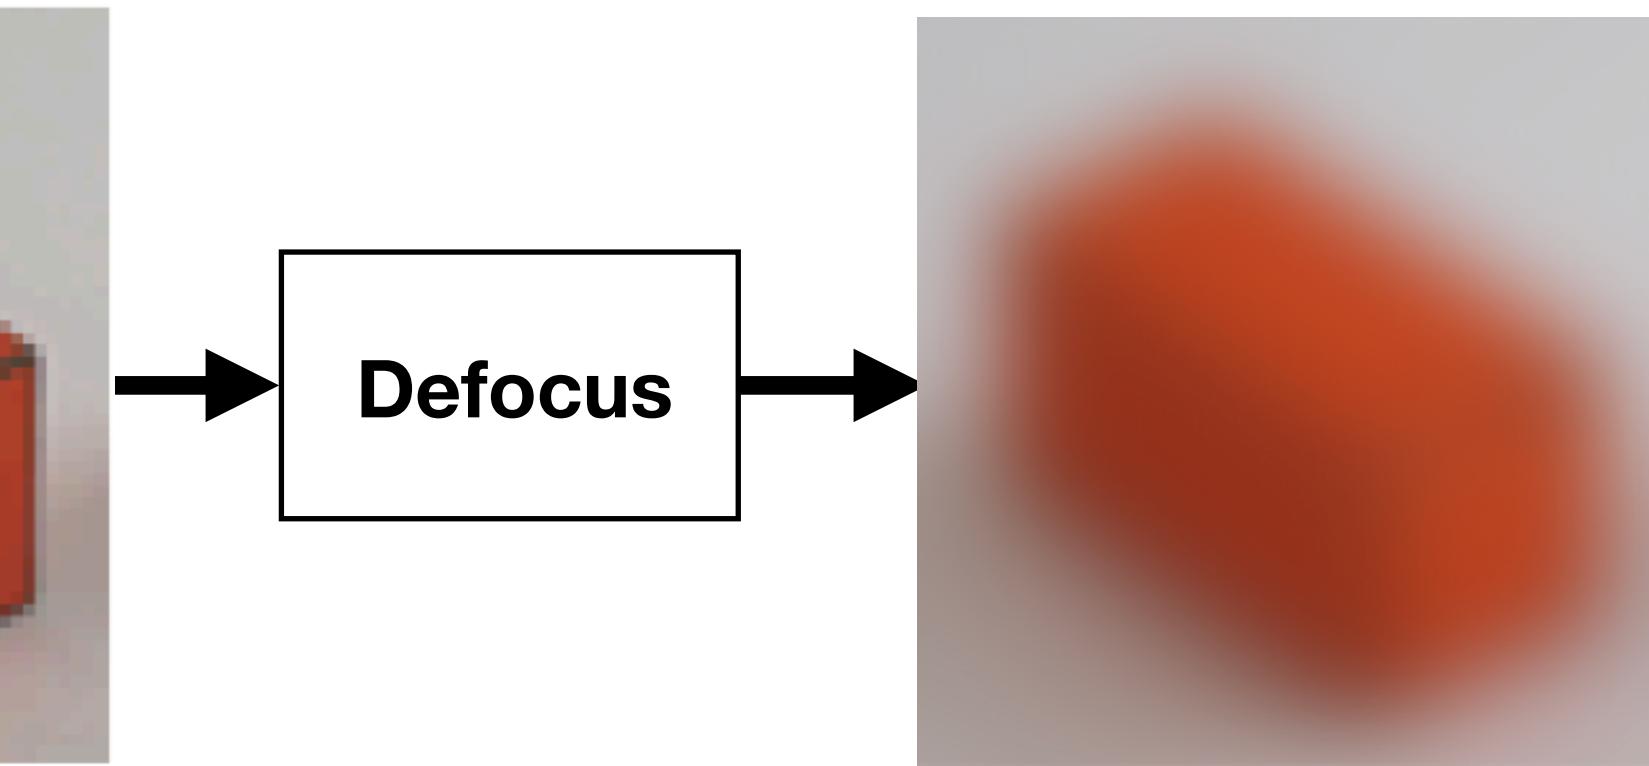
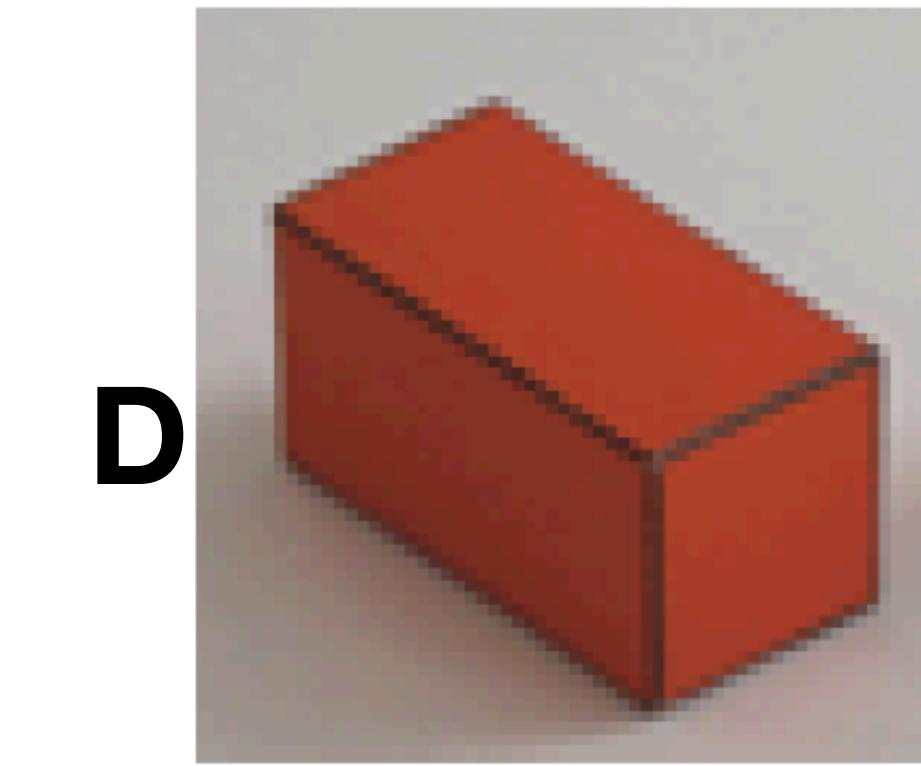
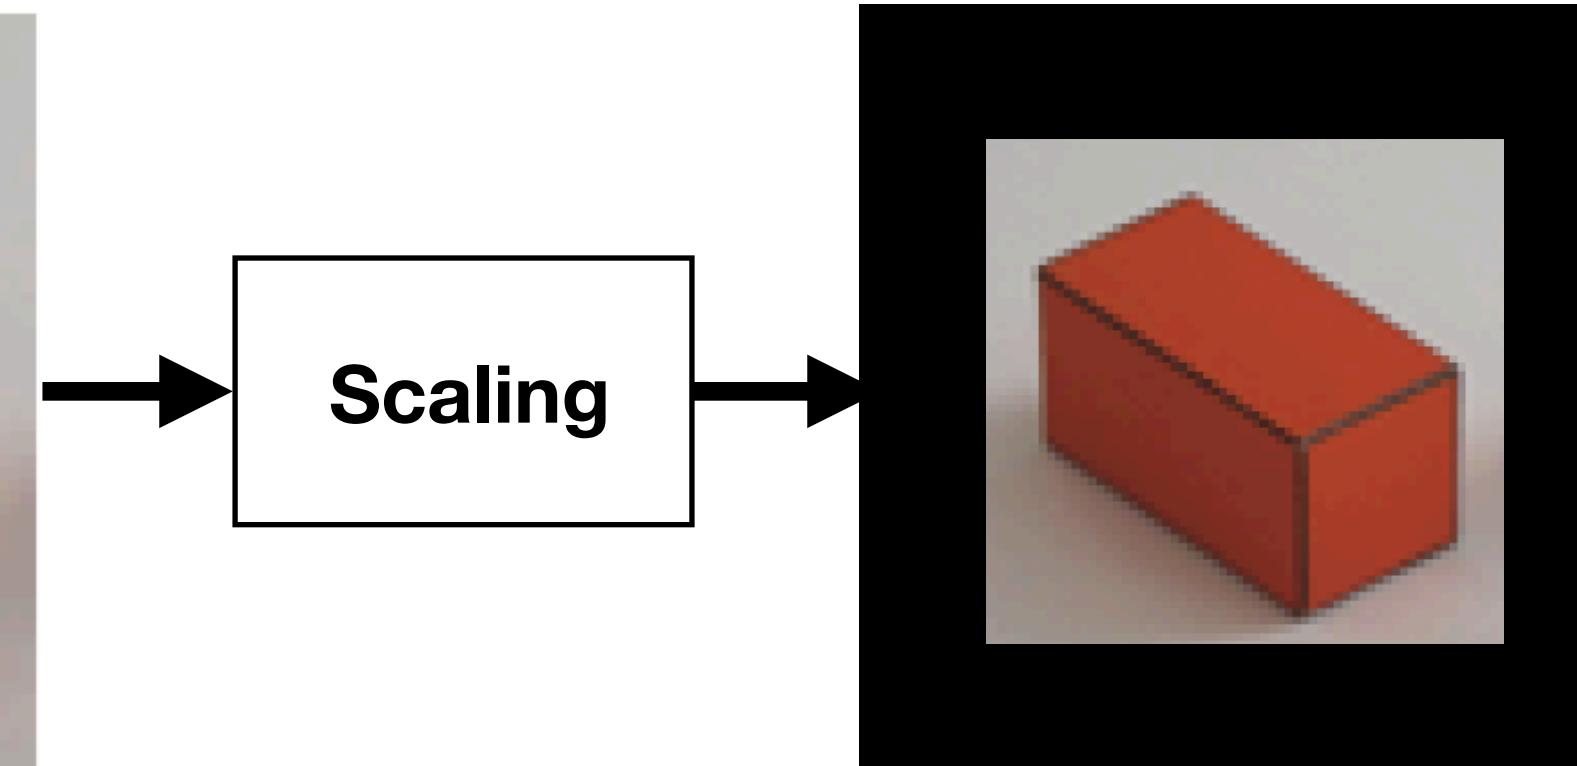
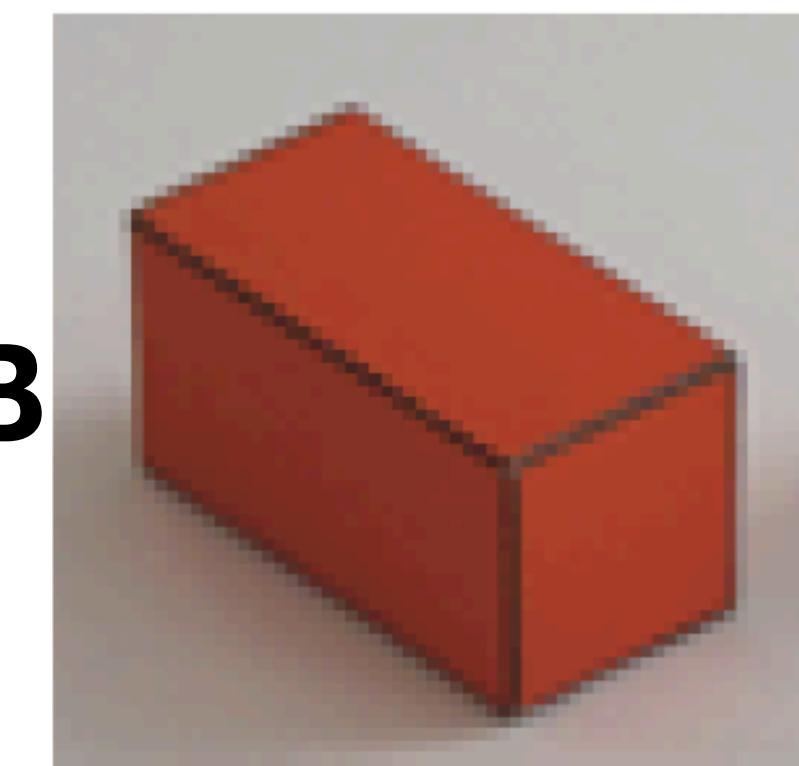
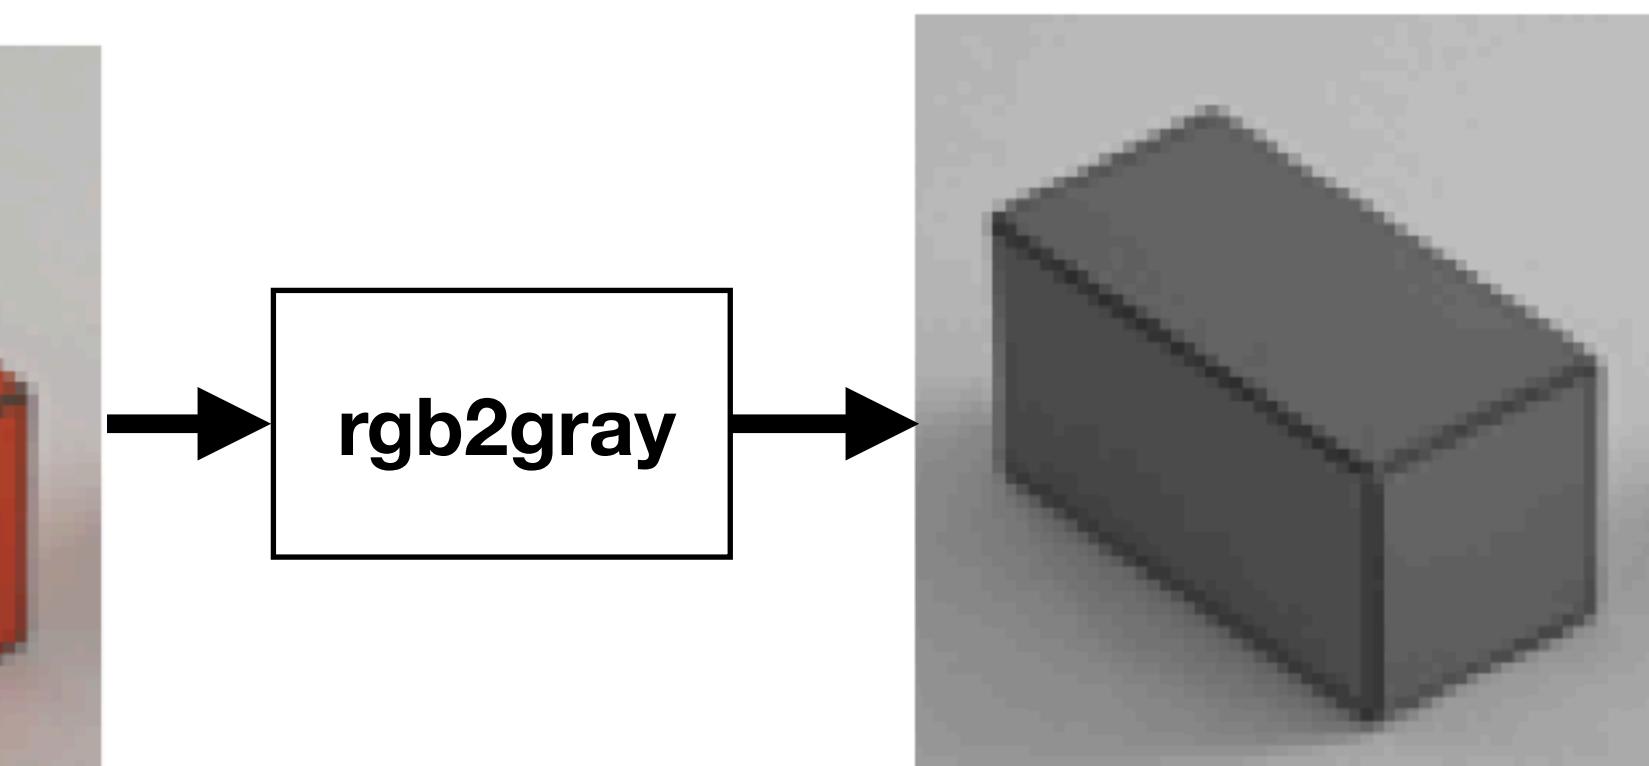
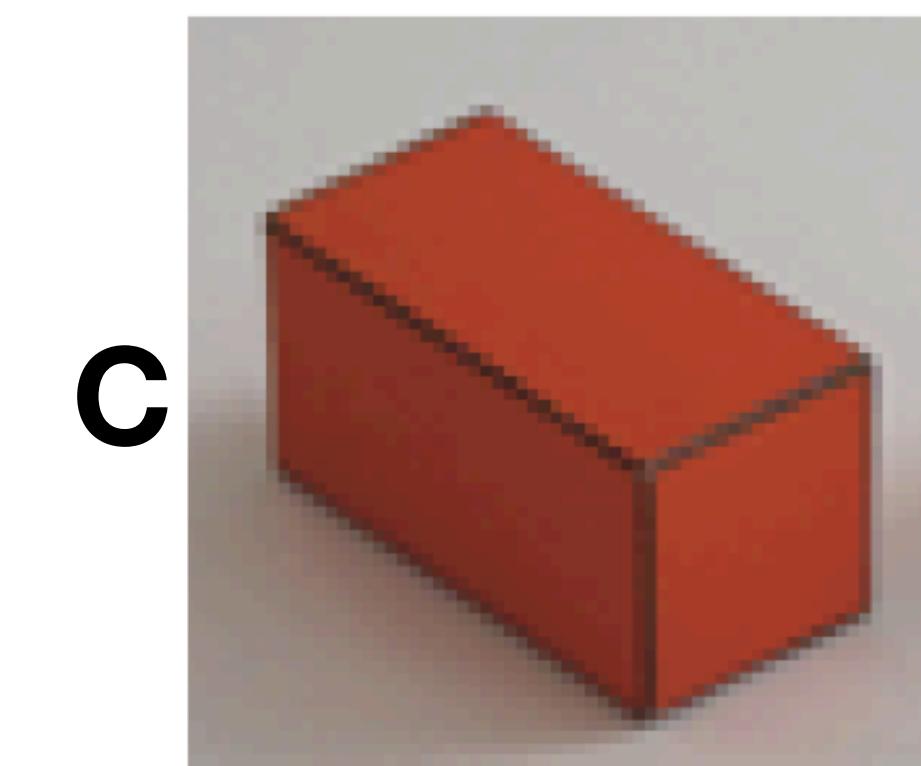
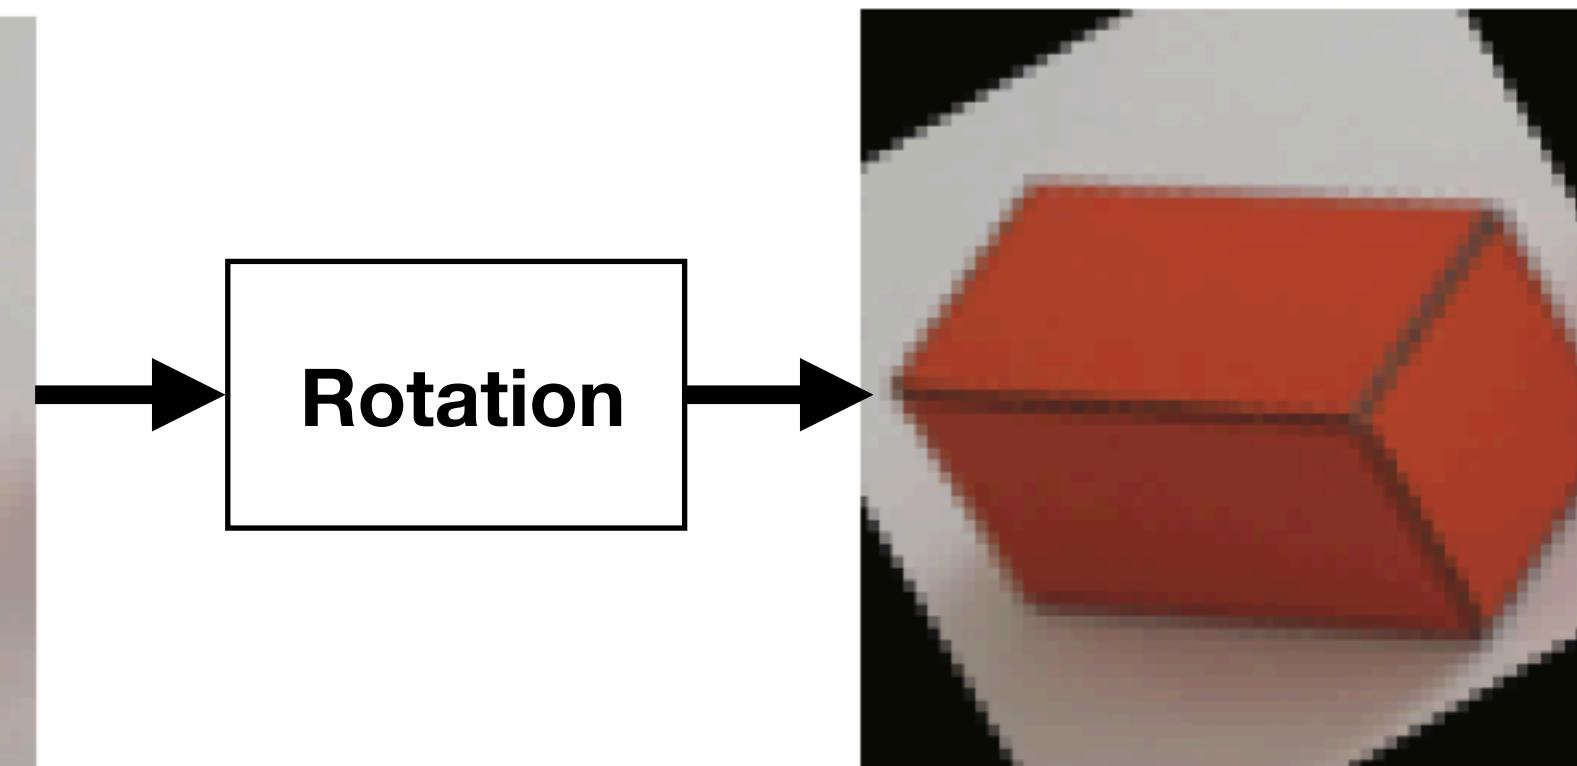
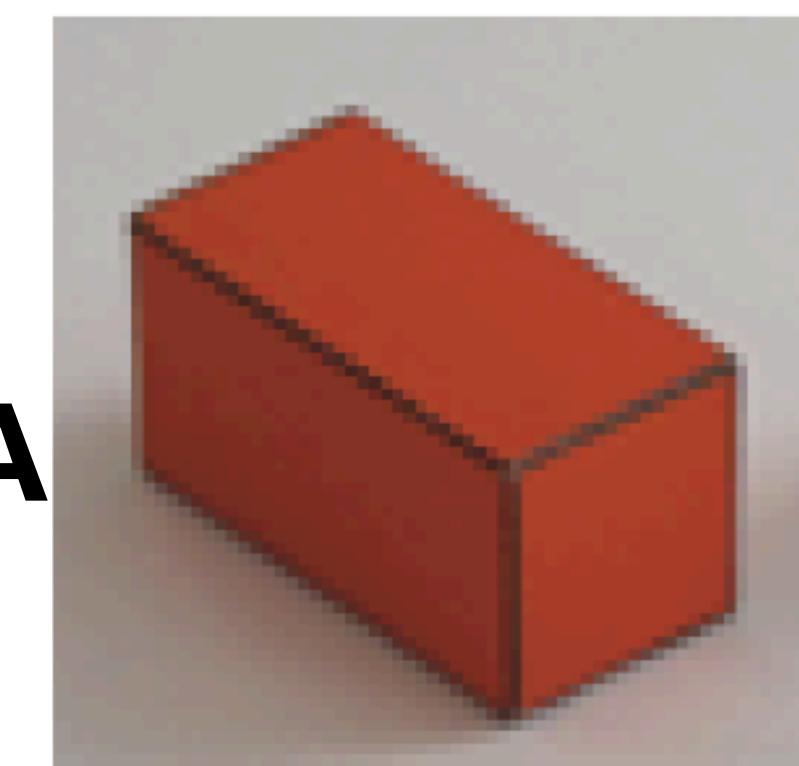
Cross-correlation

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

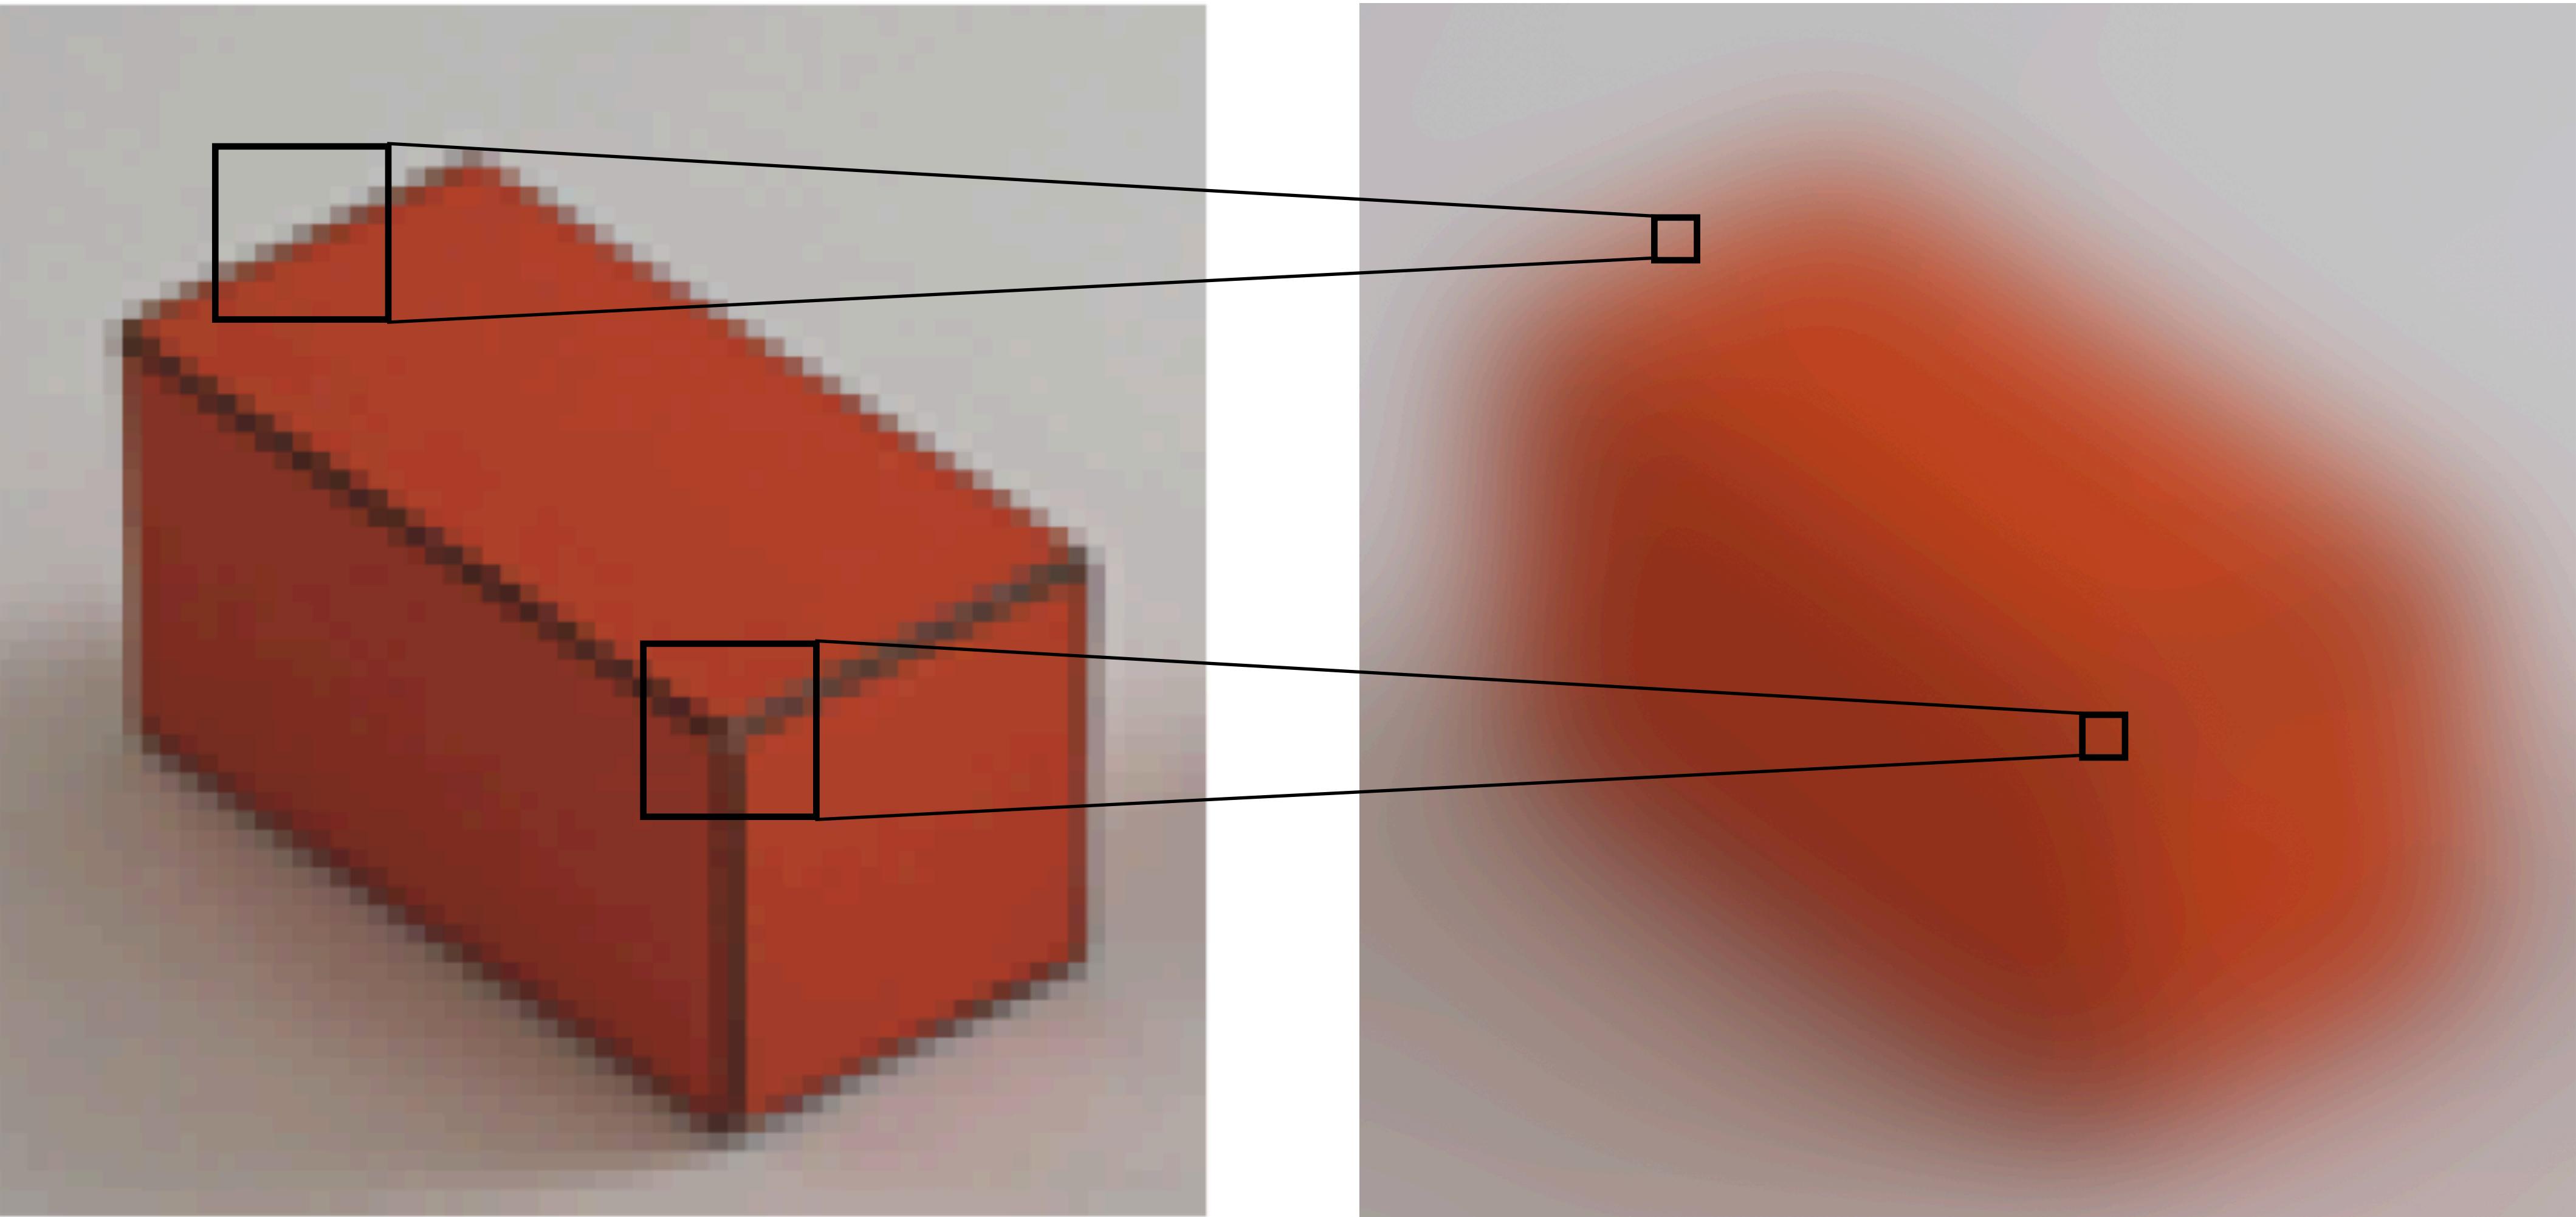
In the convolution, the kernel h is inverted left-right and up-down, while in the cross-correlation is not

Quiz: what operation is the result of a convolution?

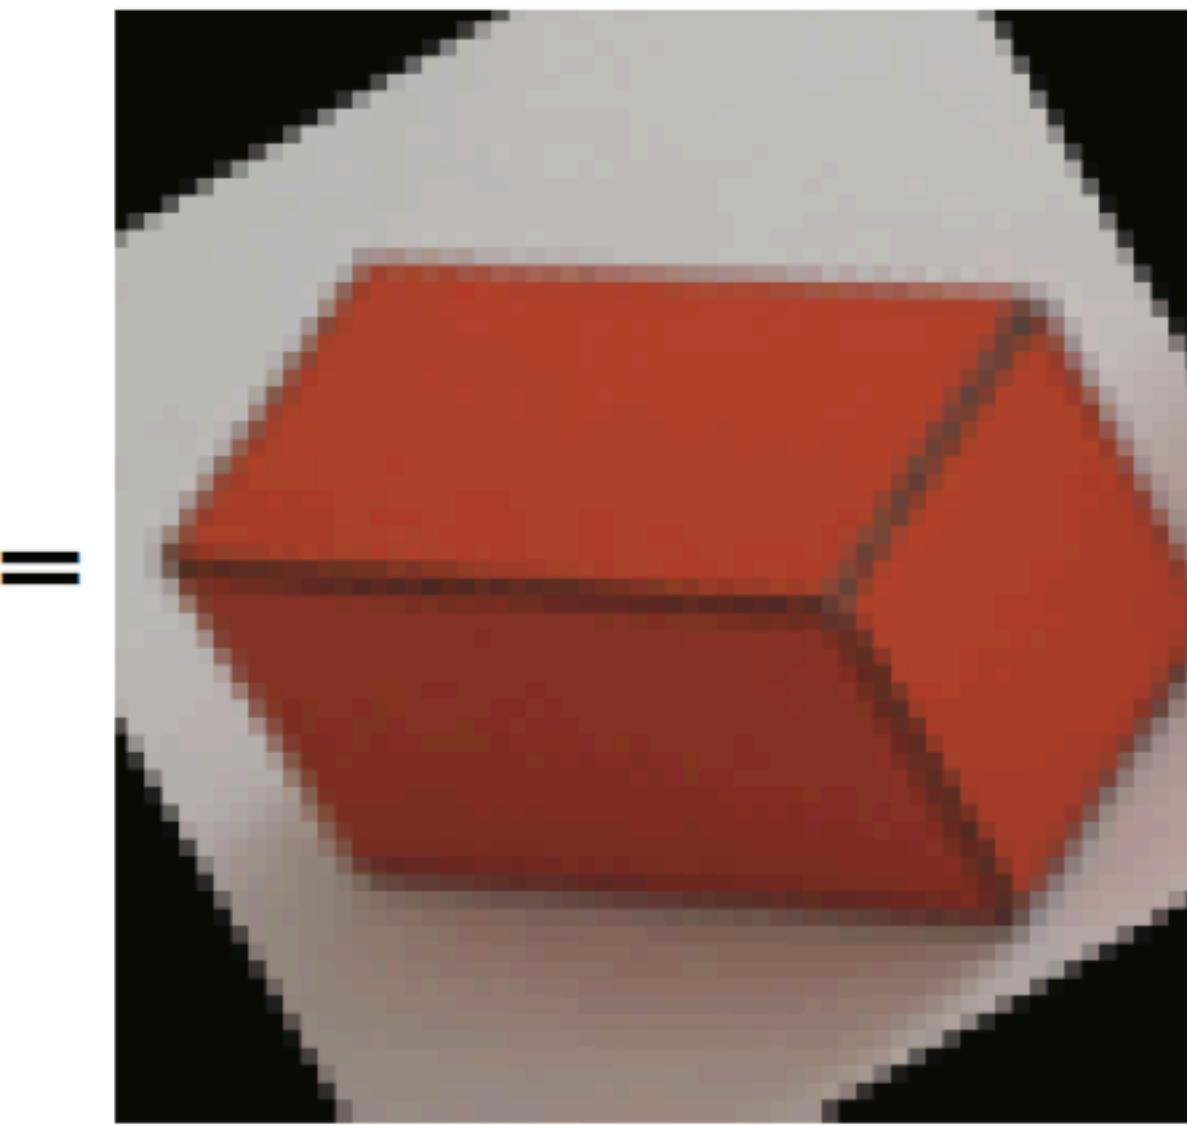
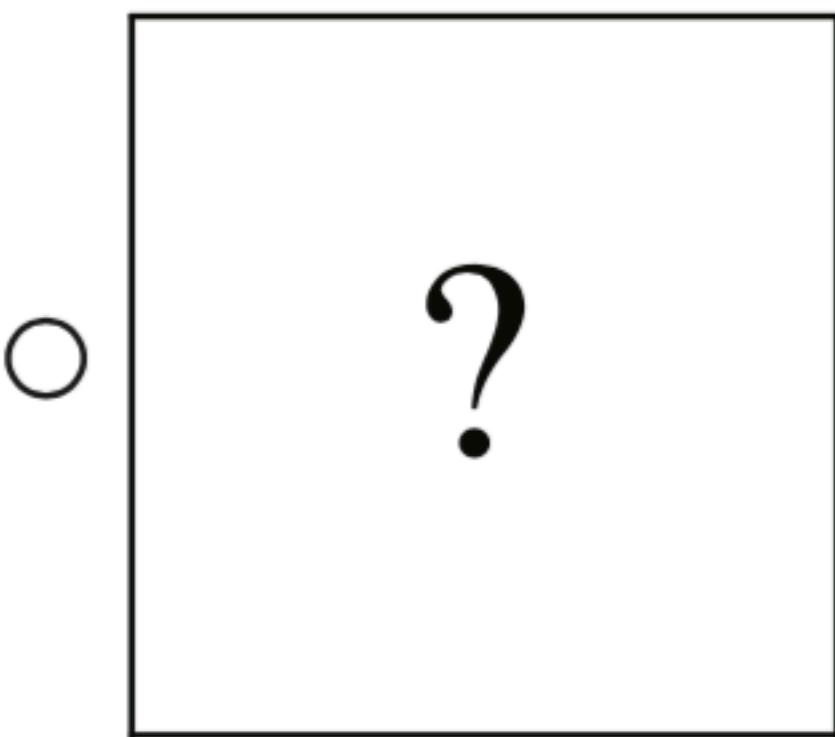
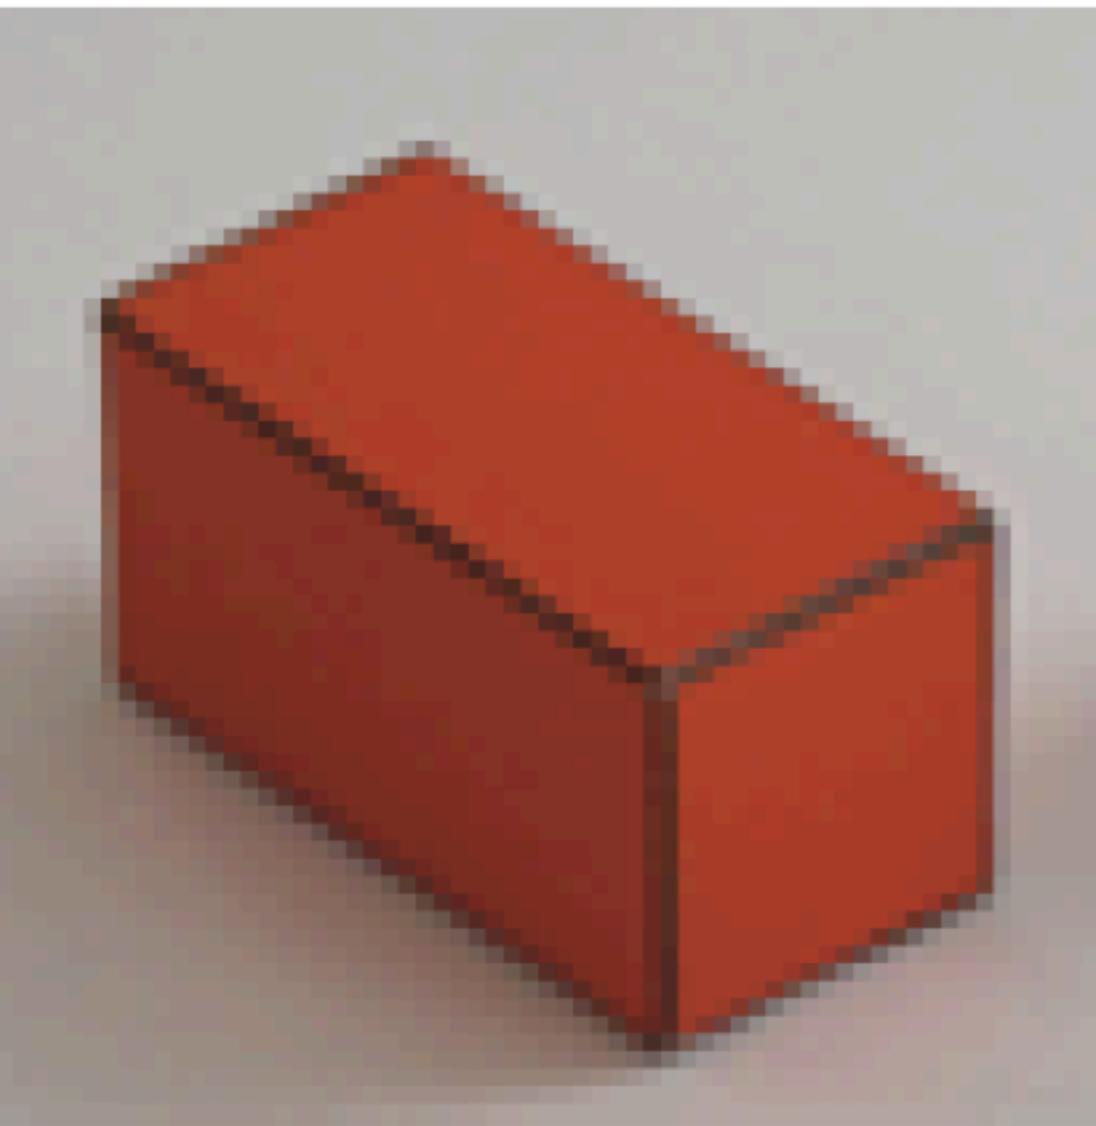
Quiz: what operation is the result of a convolution?



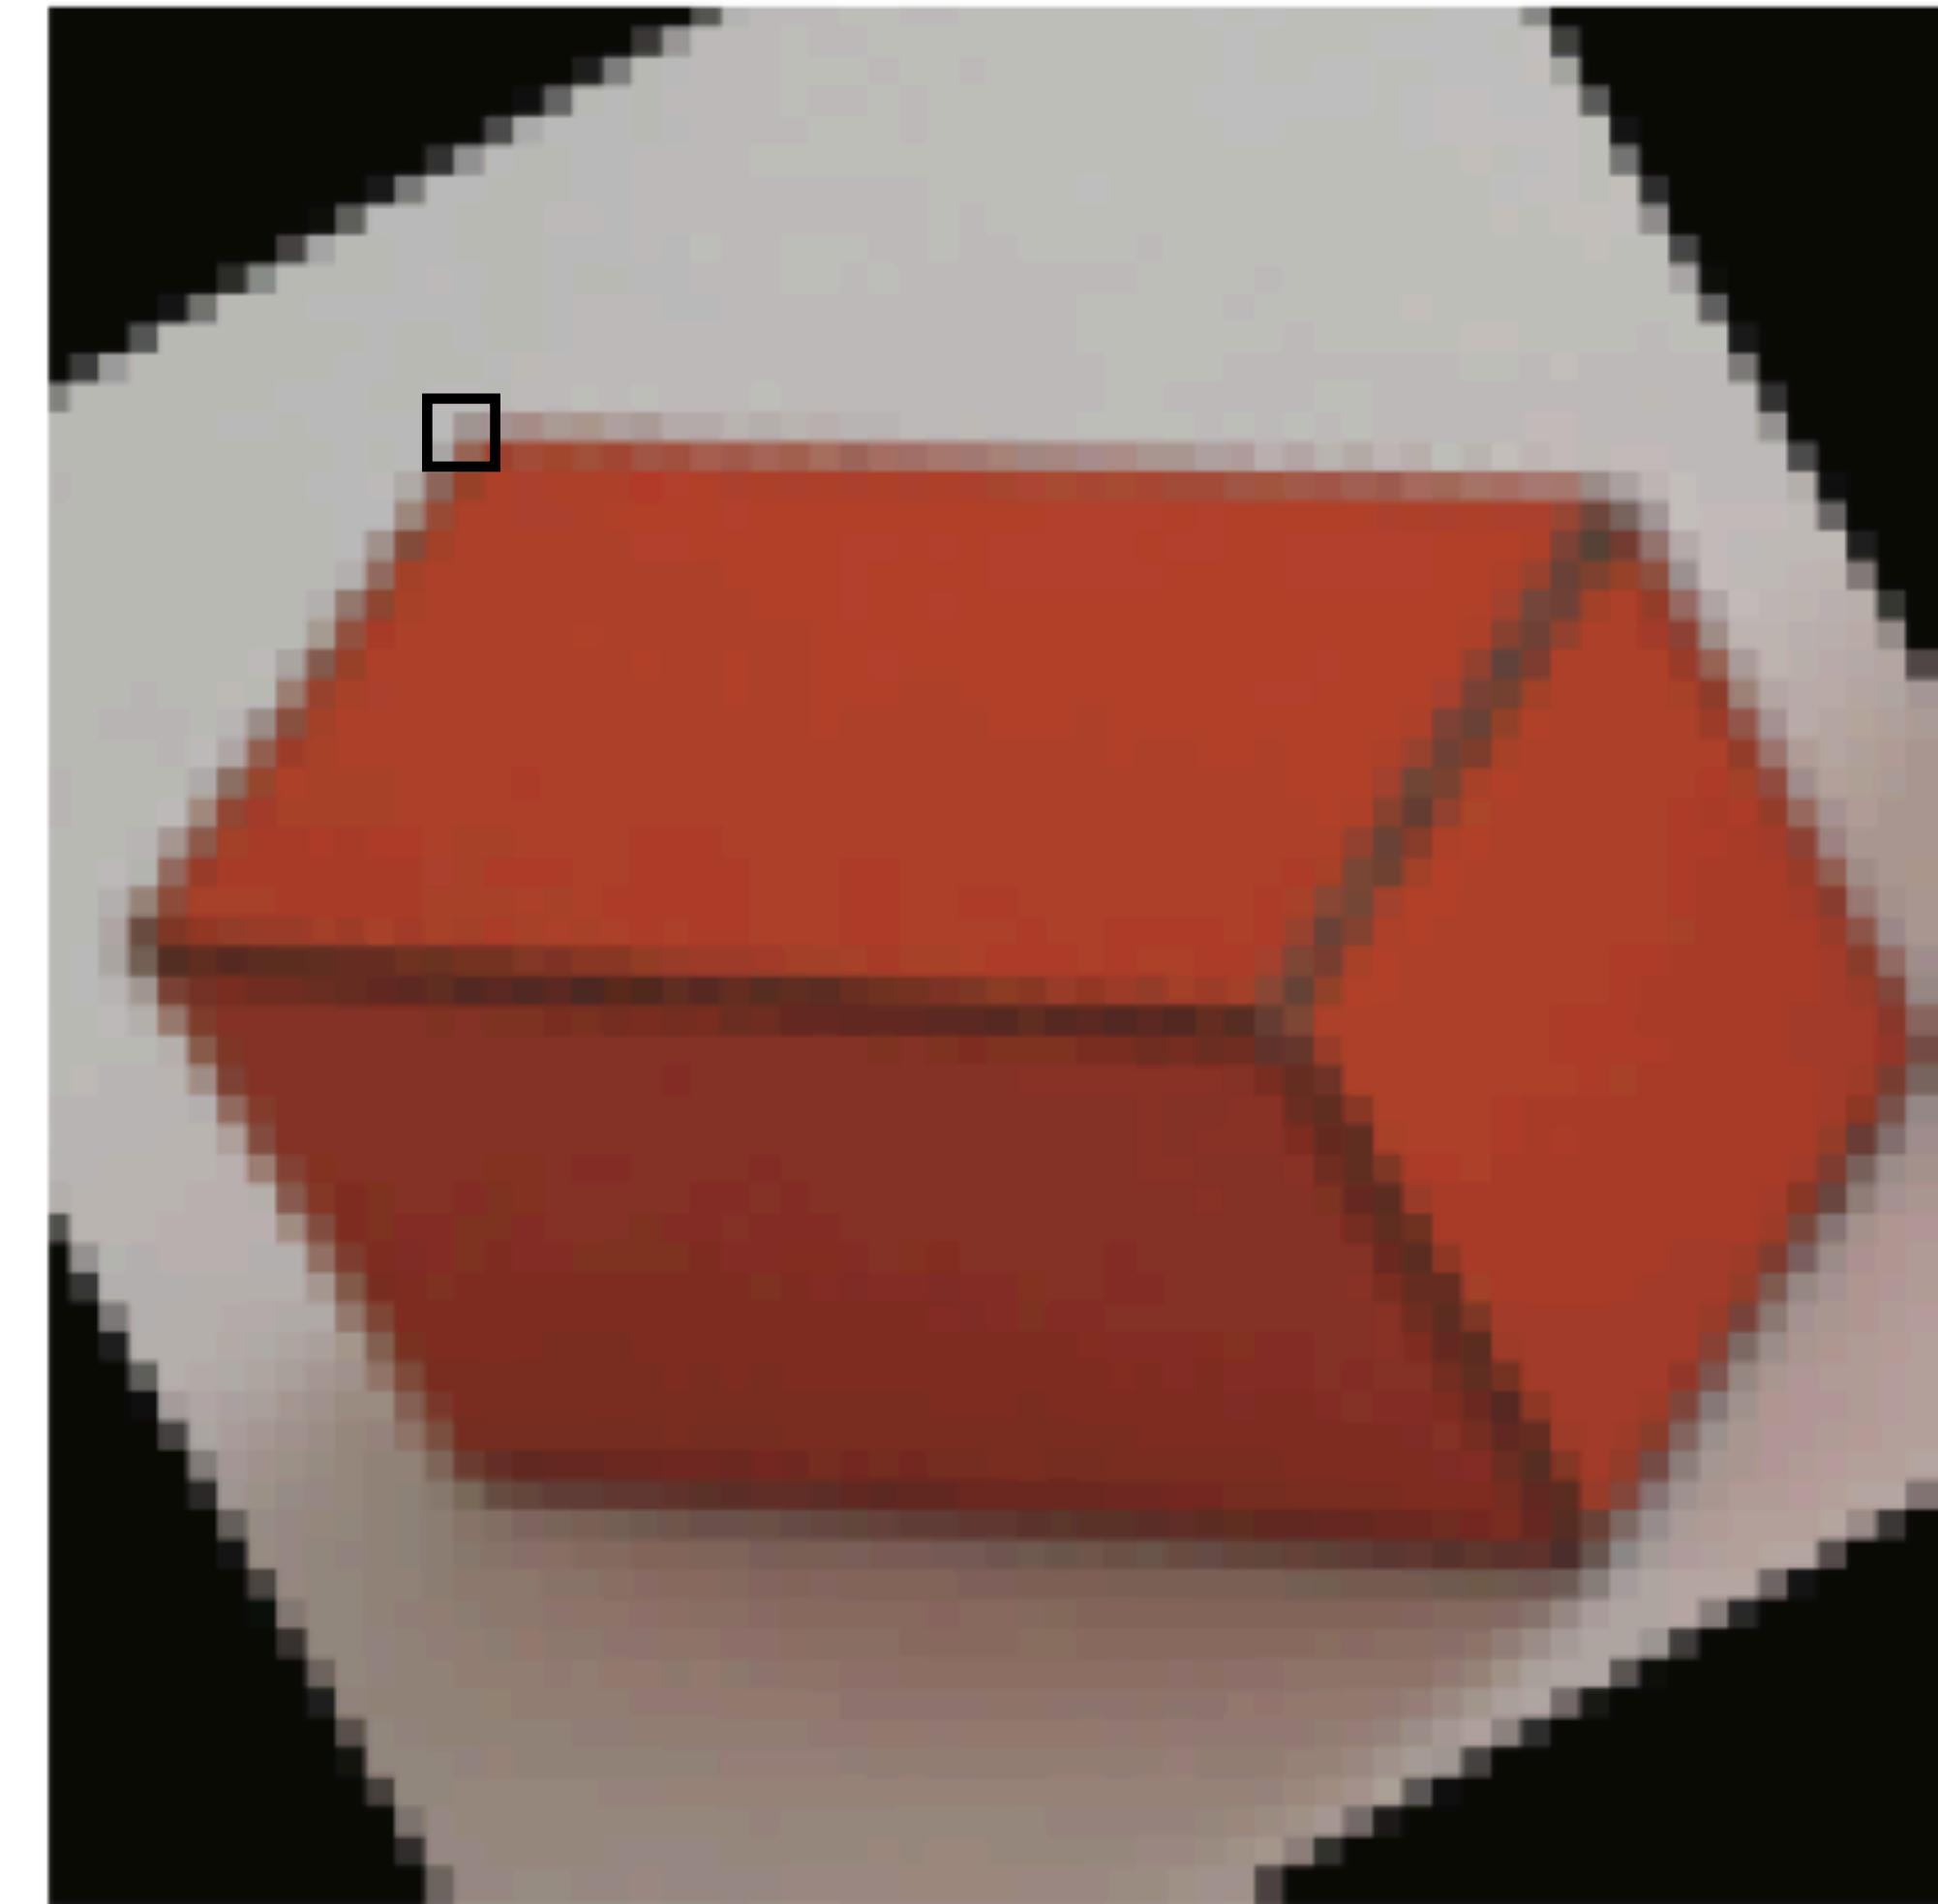
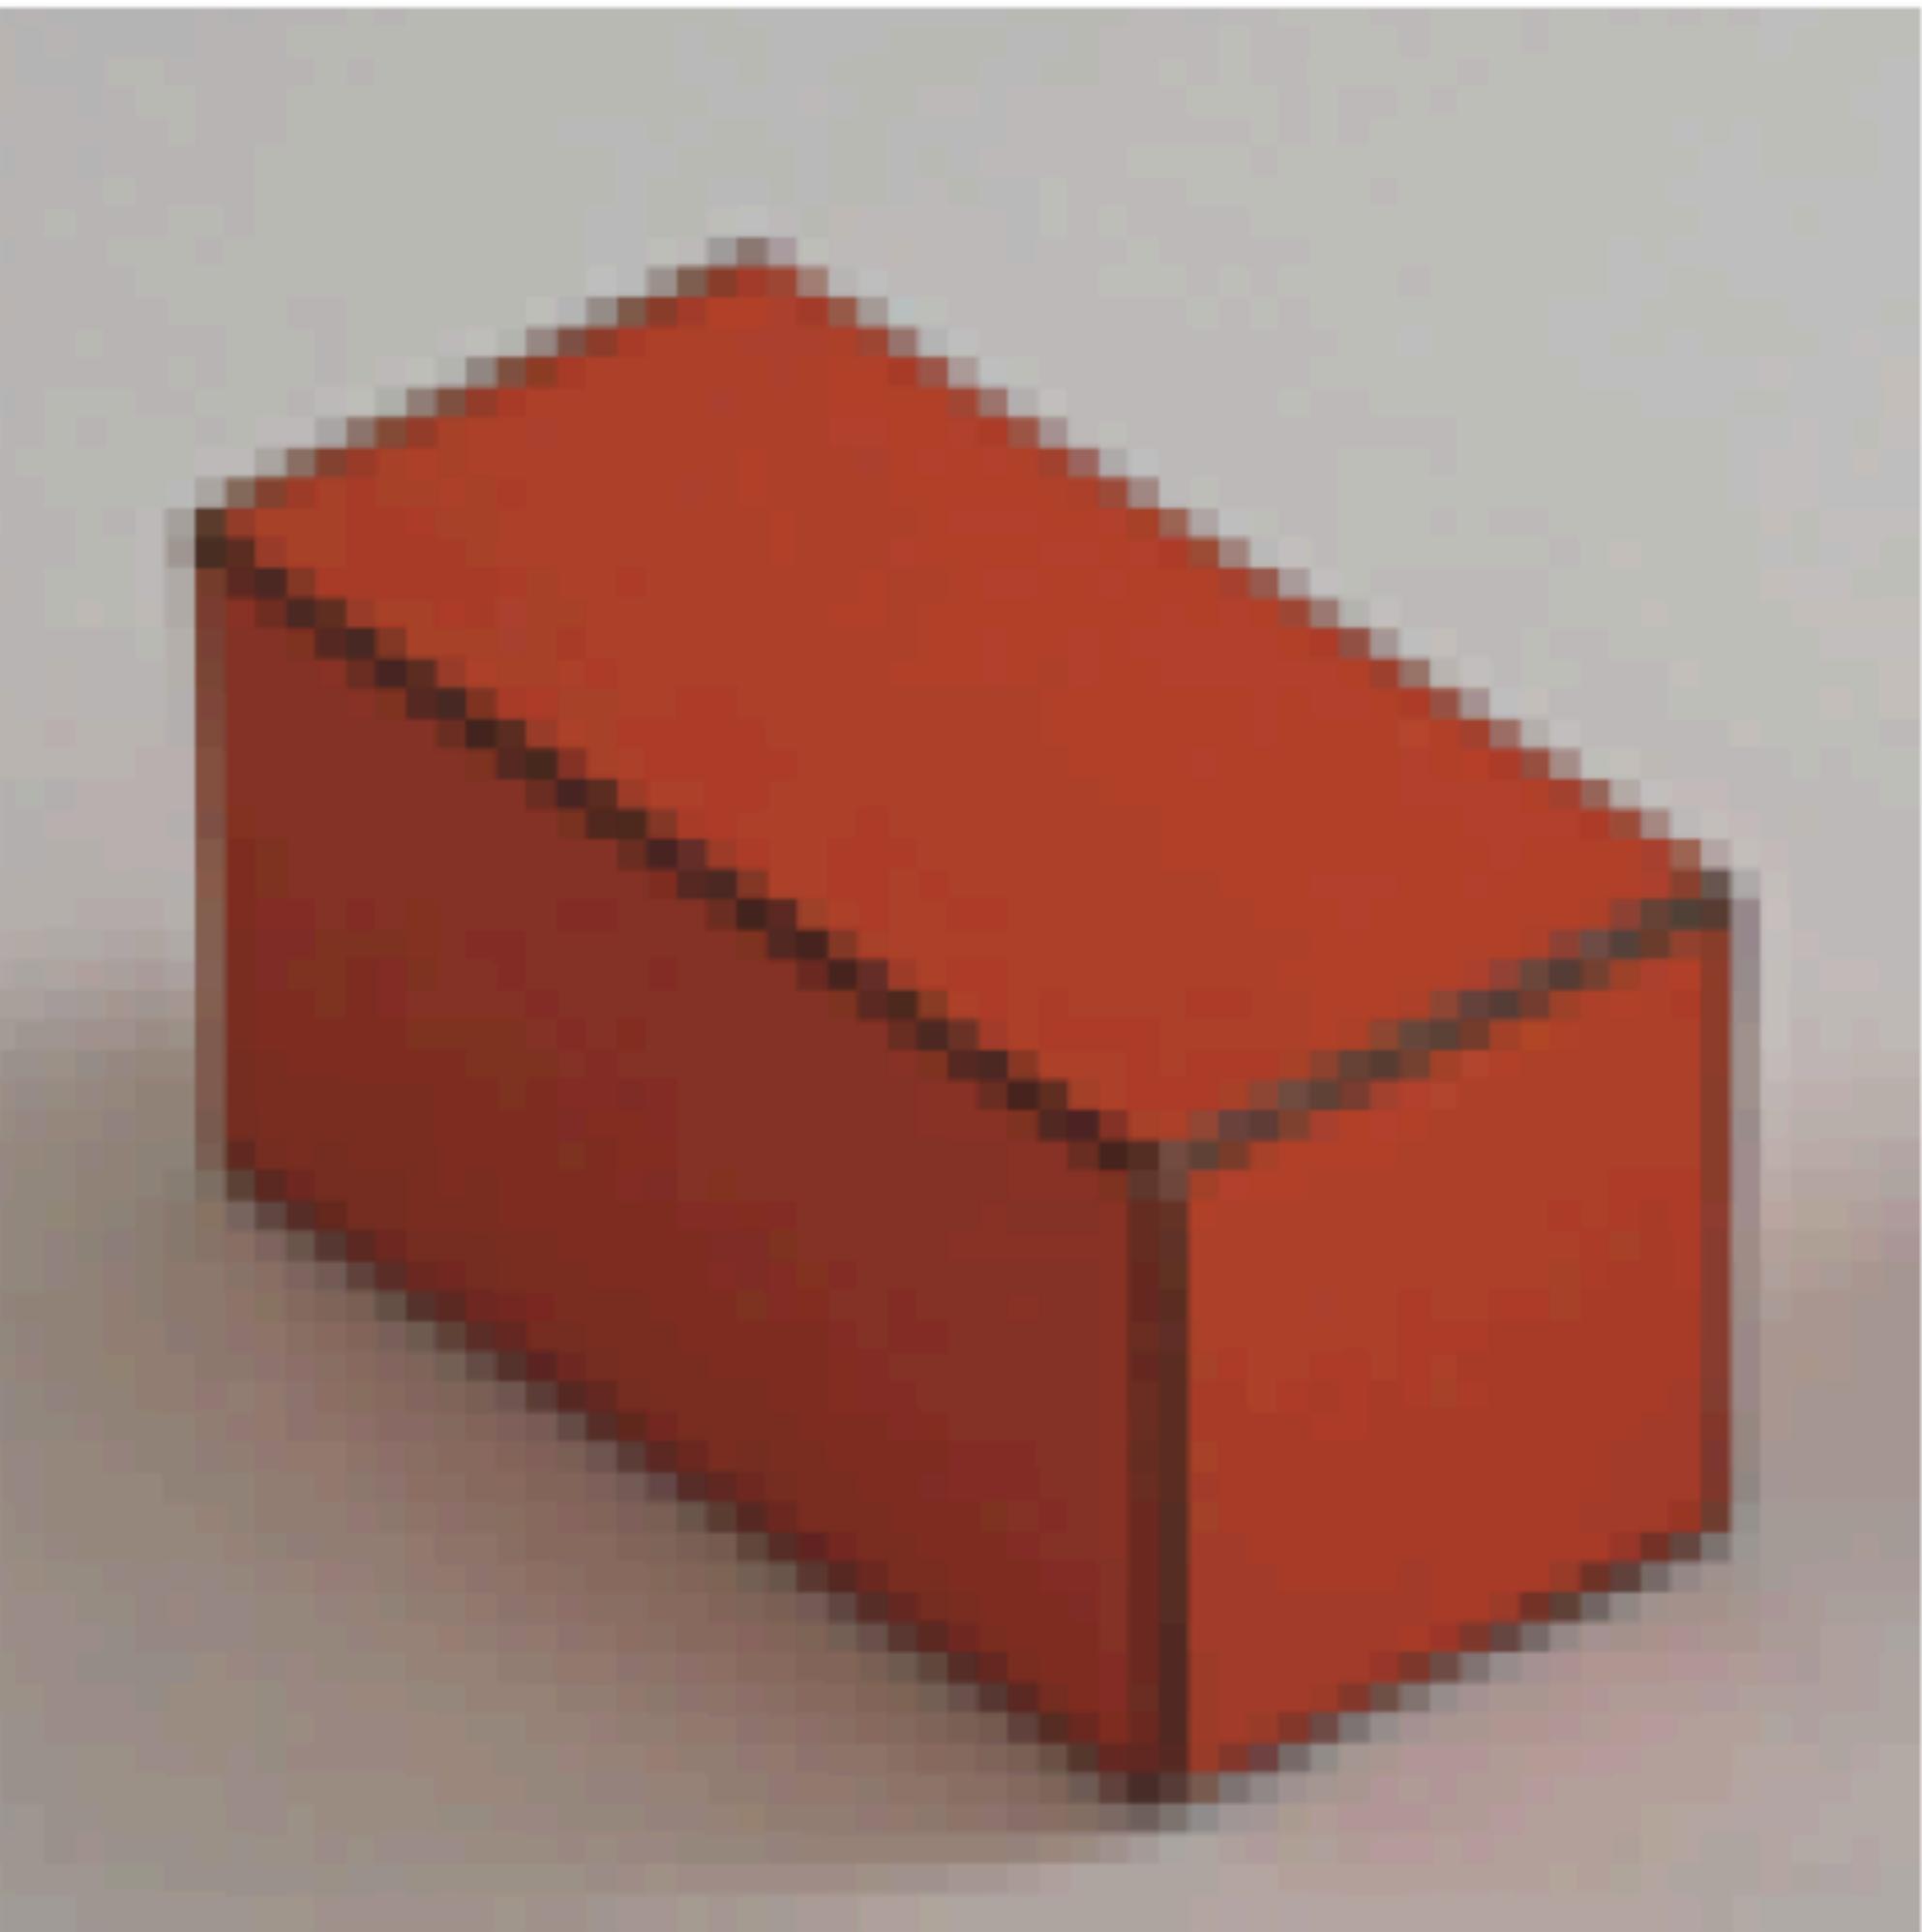
Examples



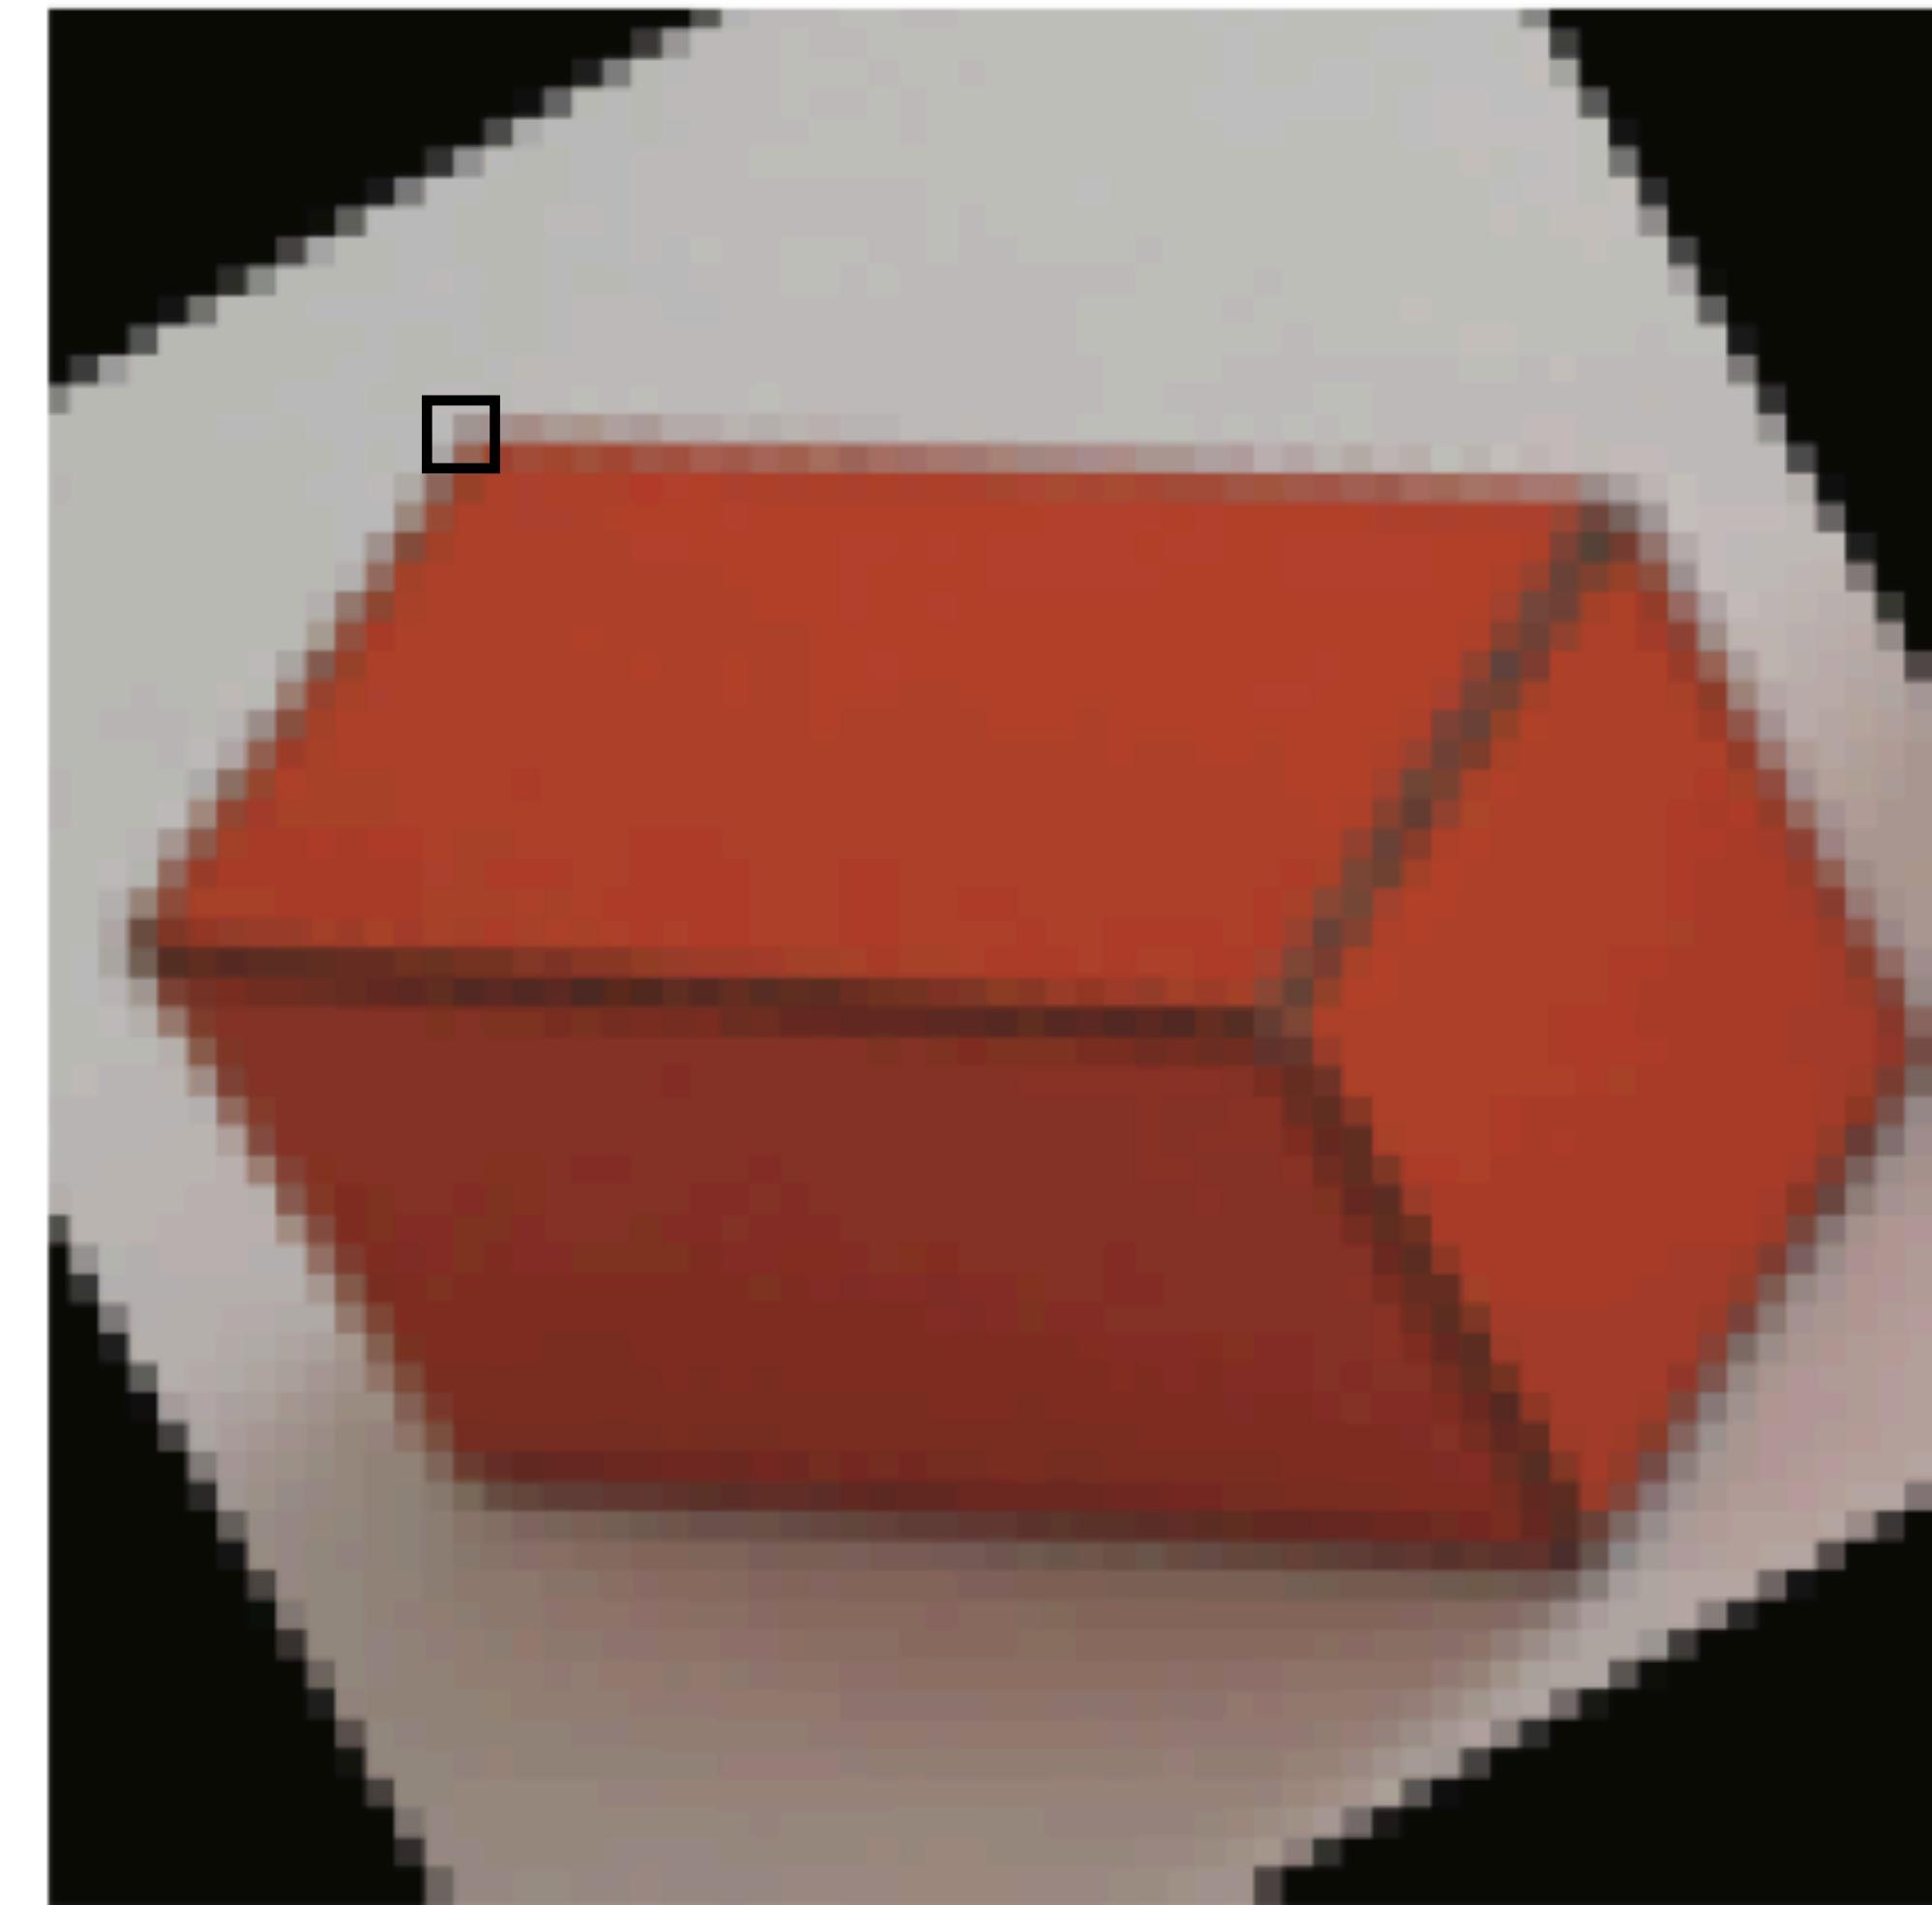
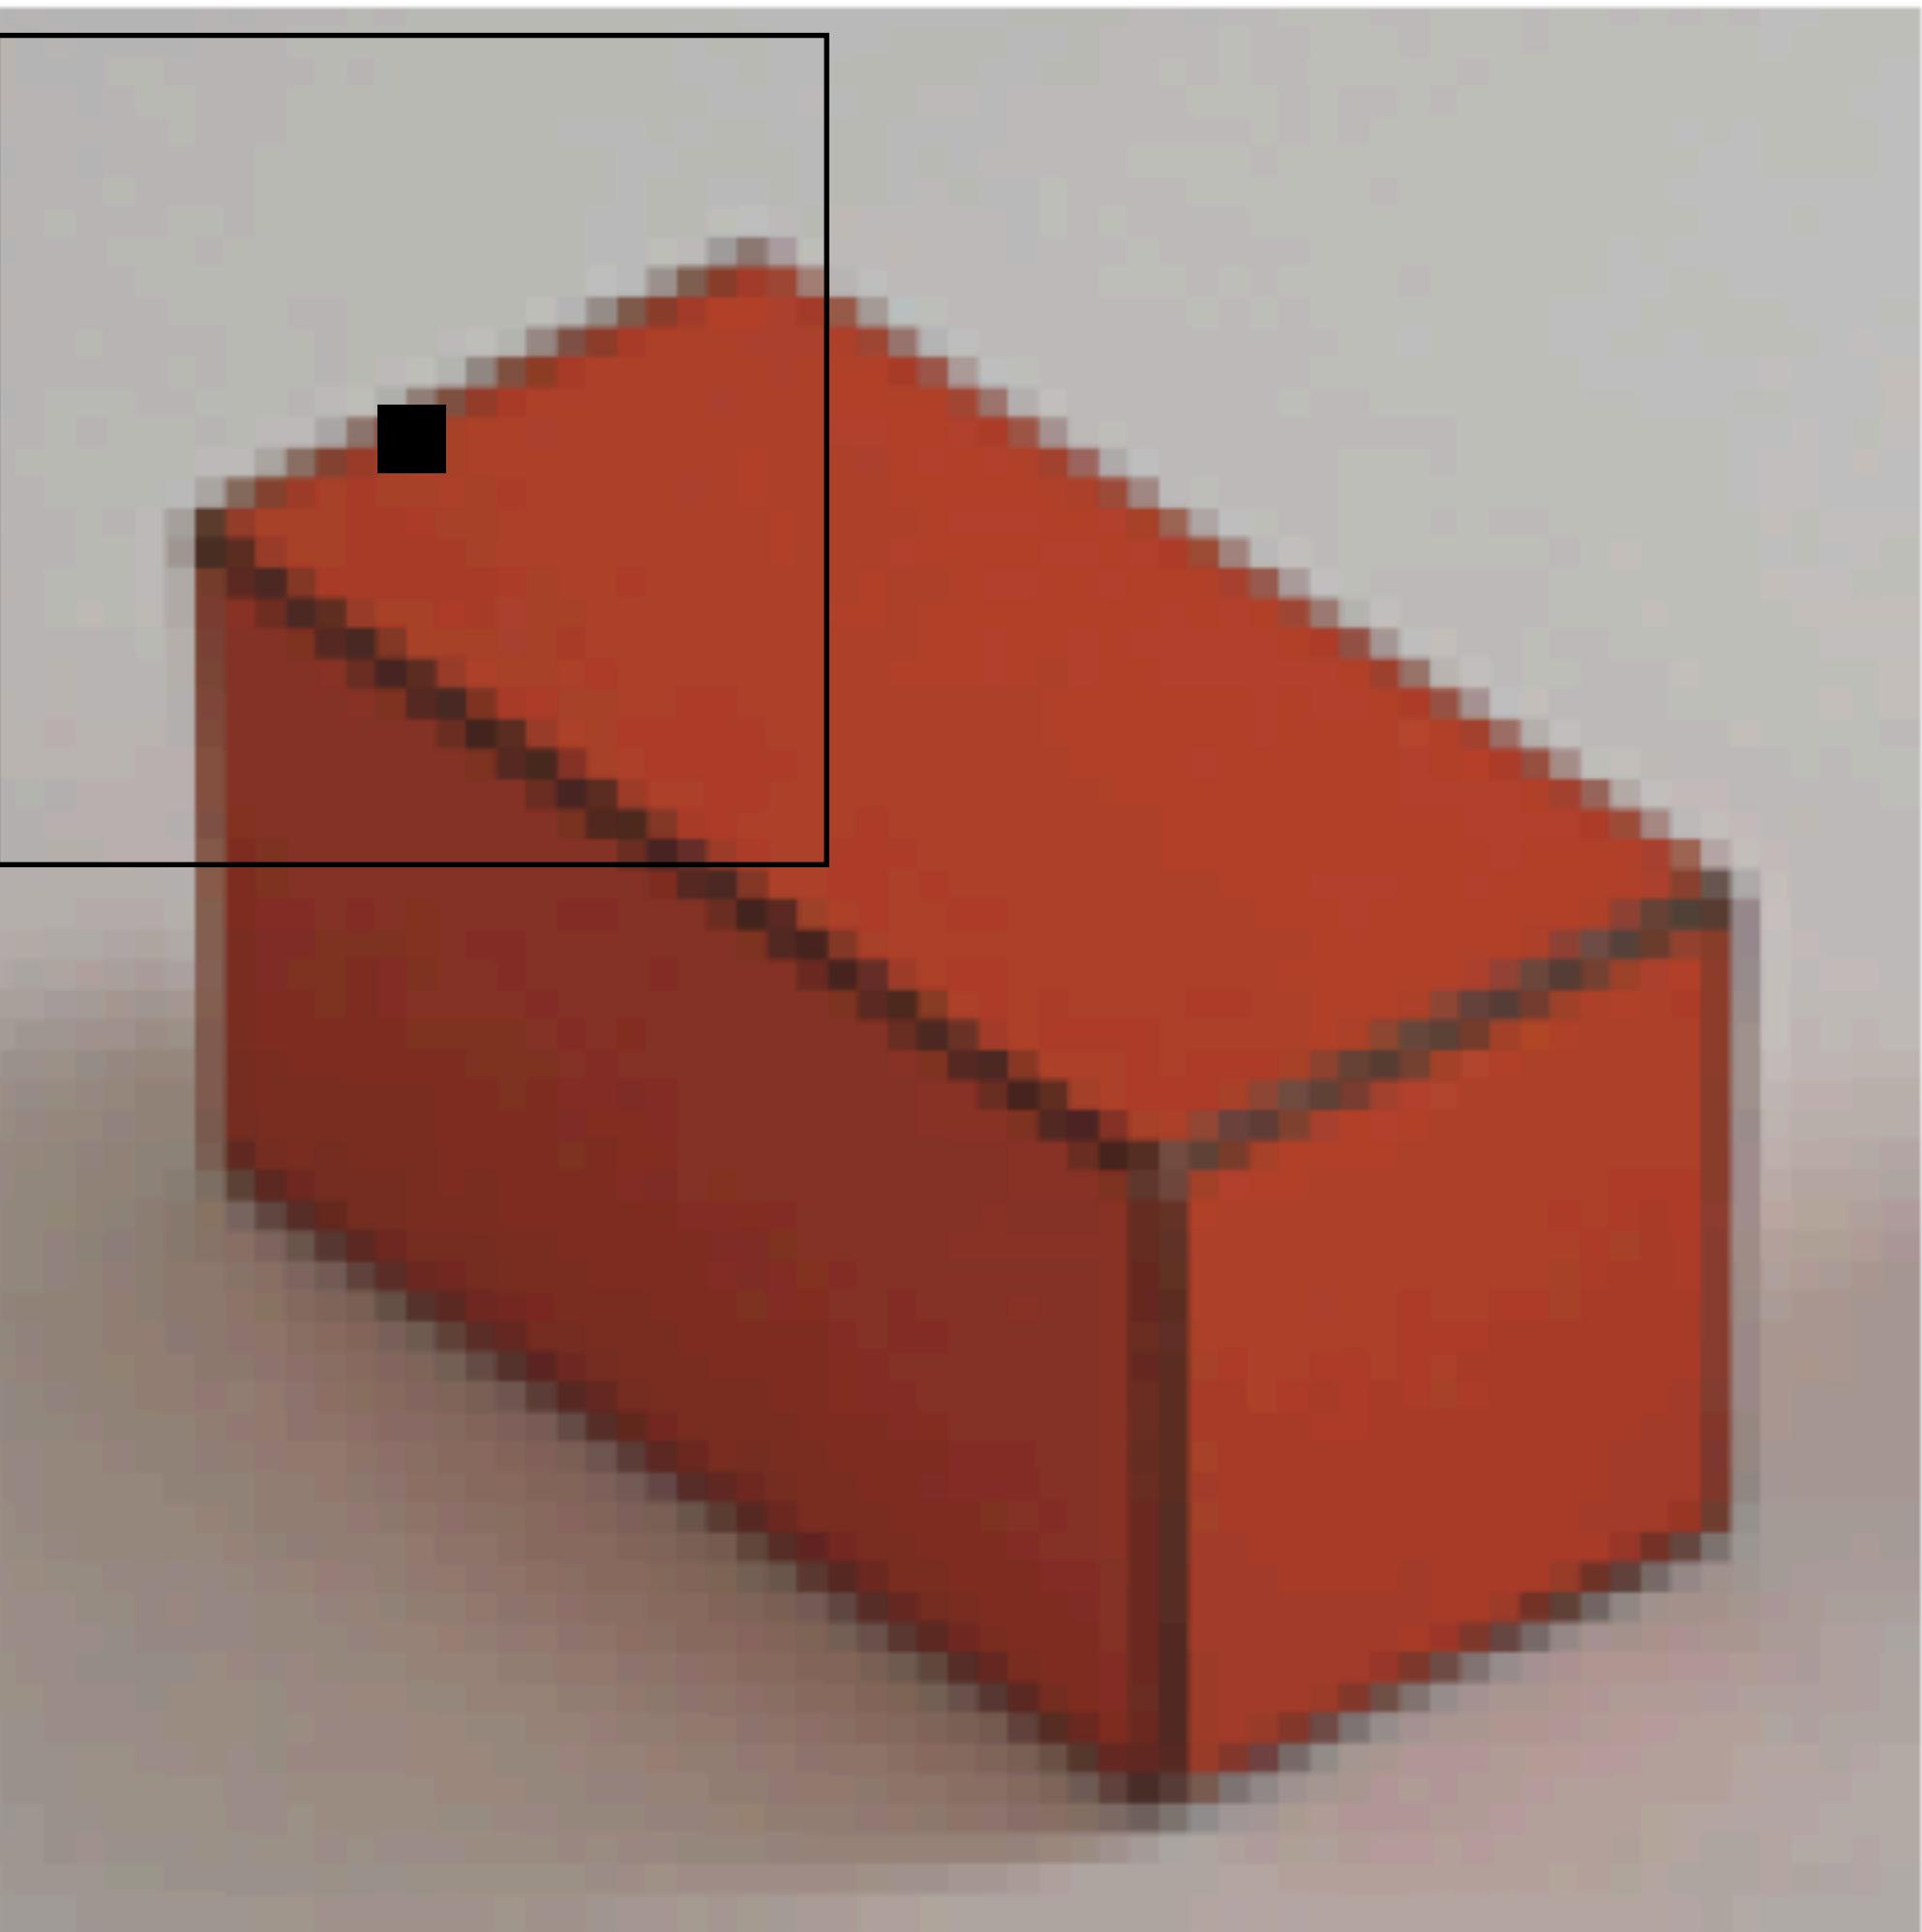
Examples



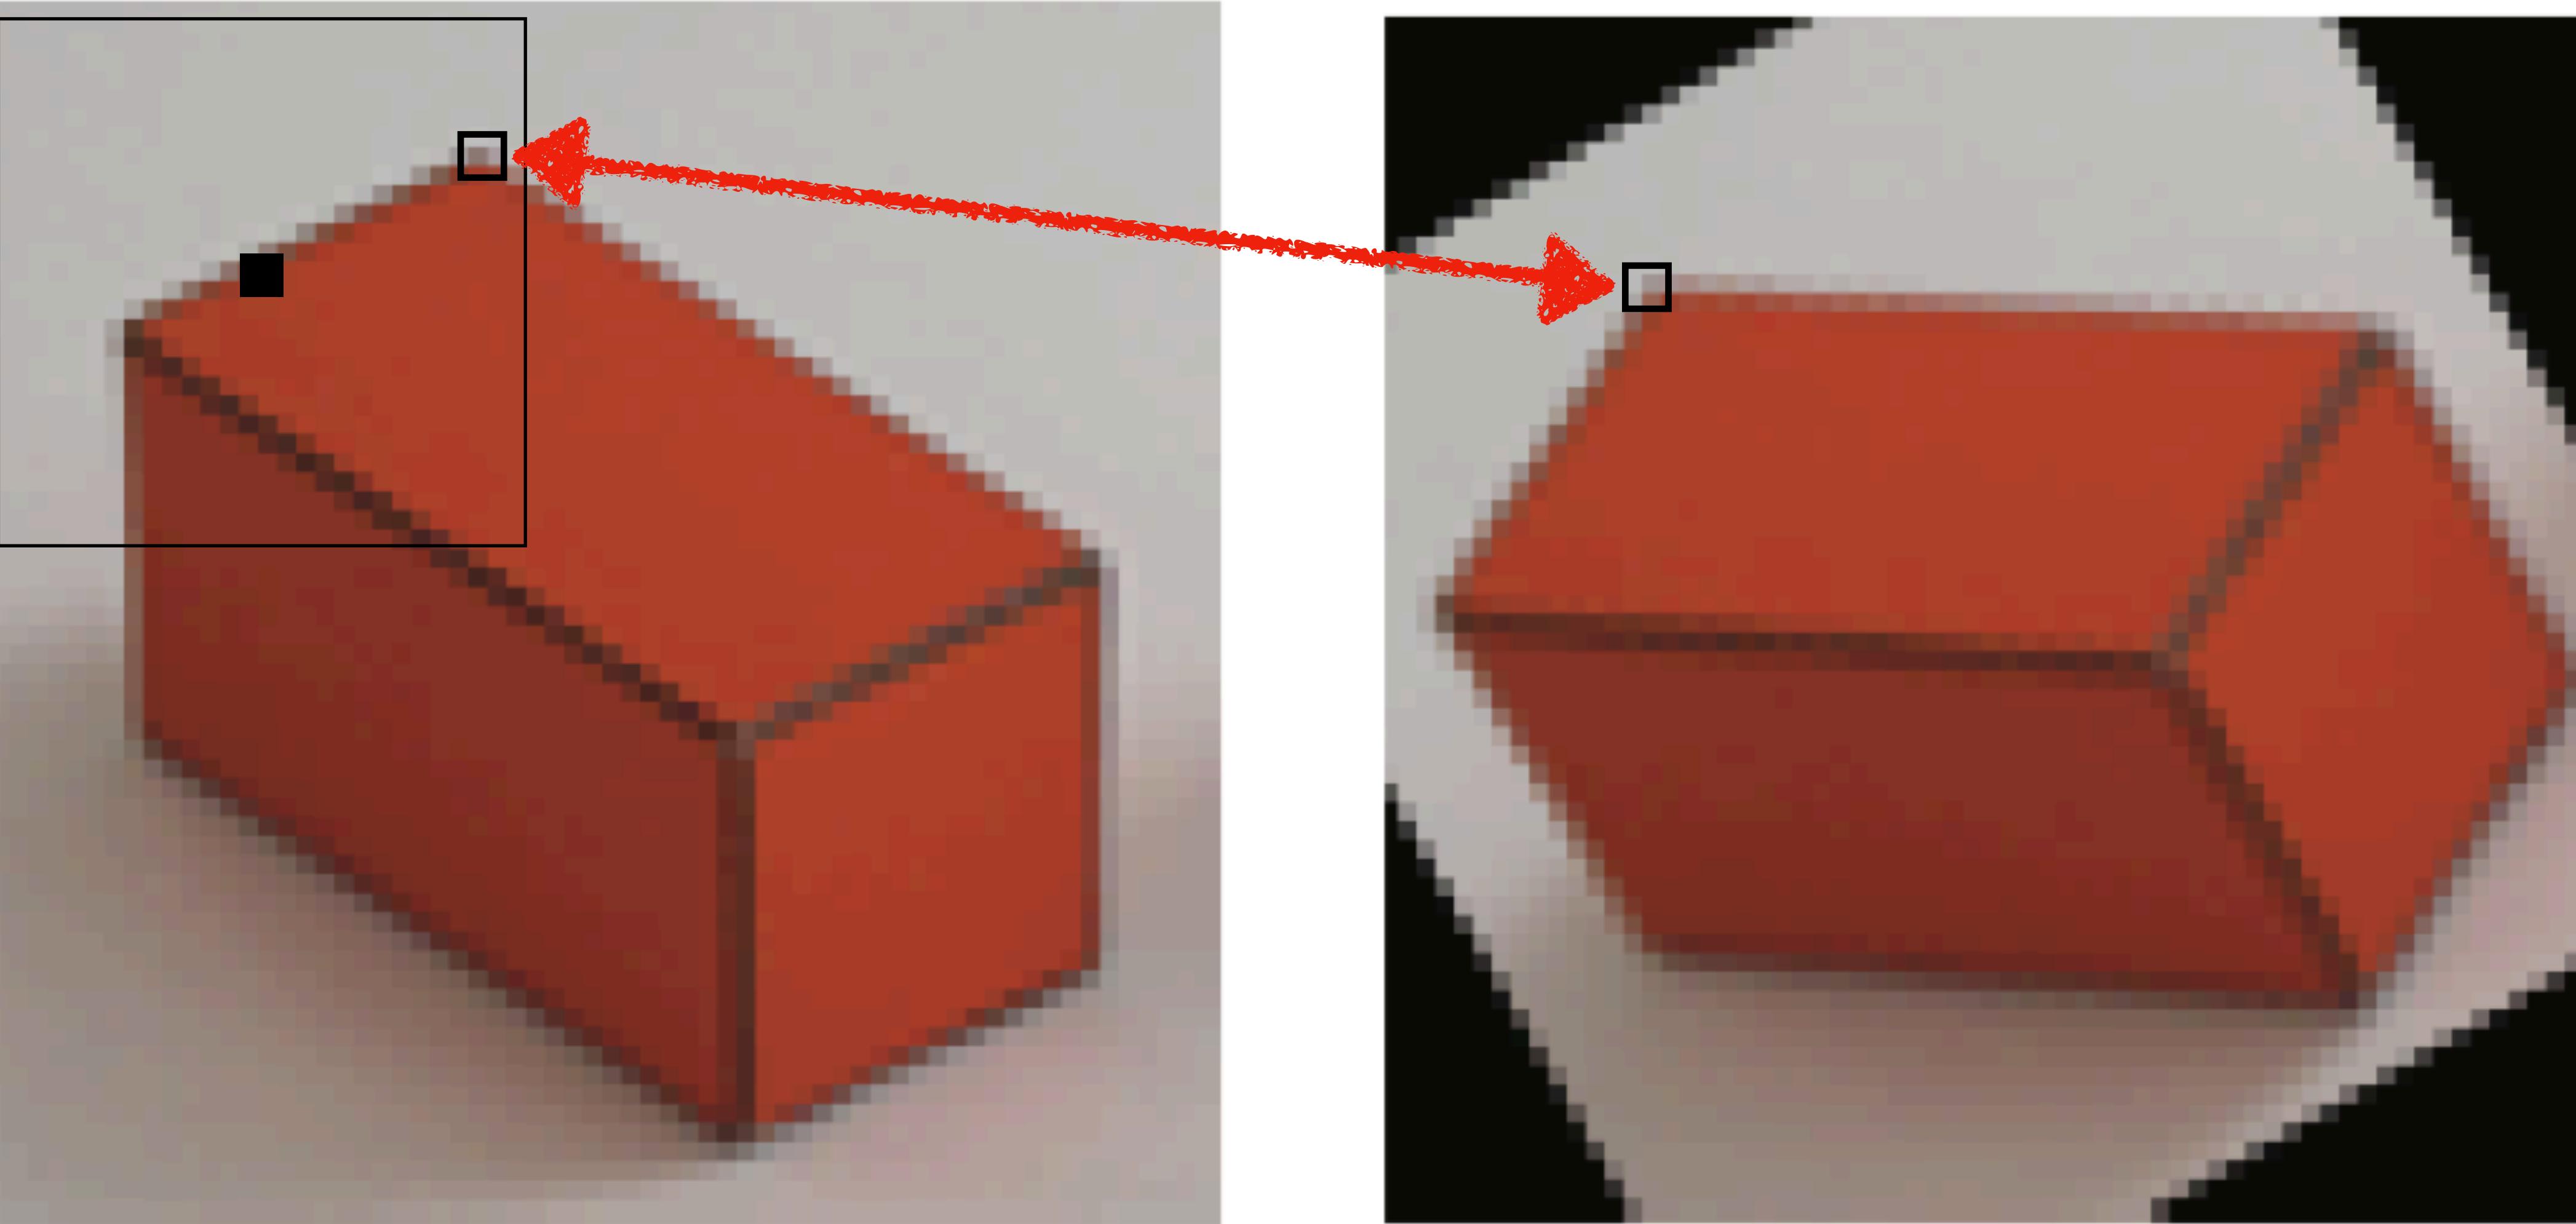
Examples



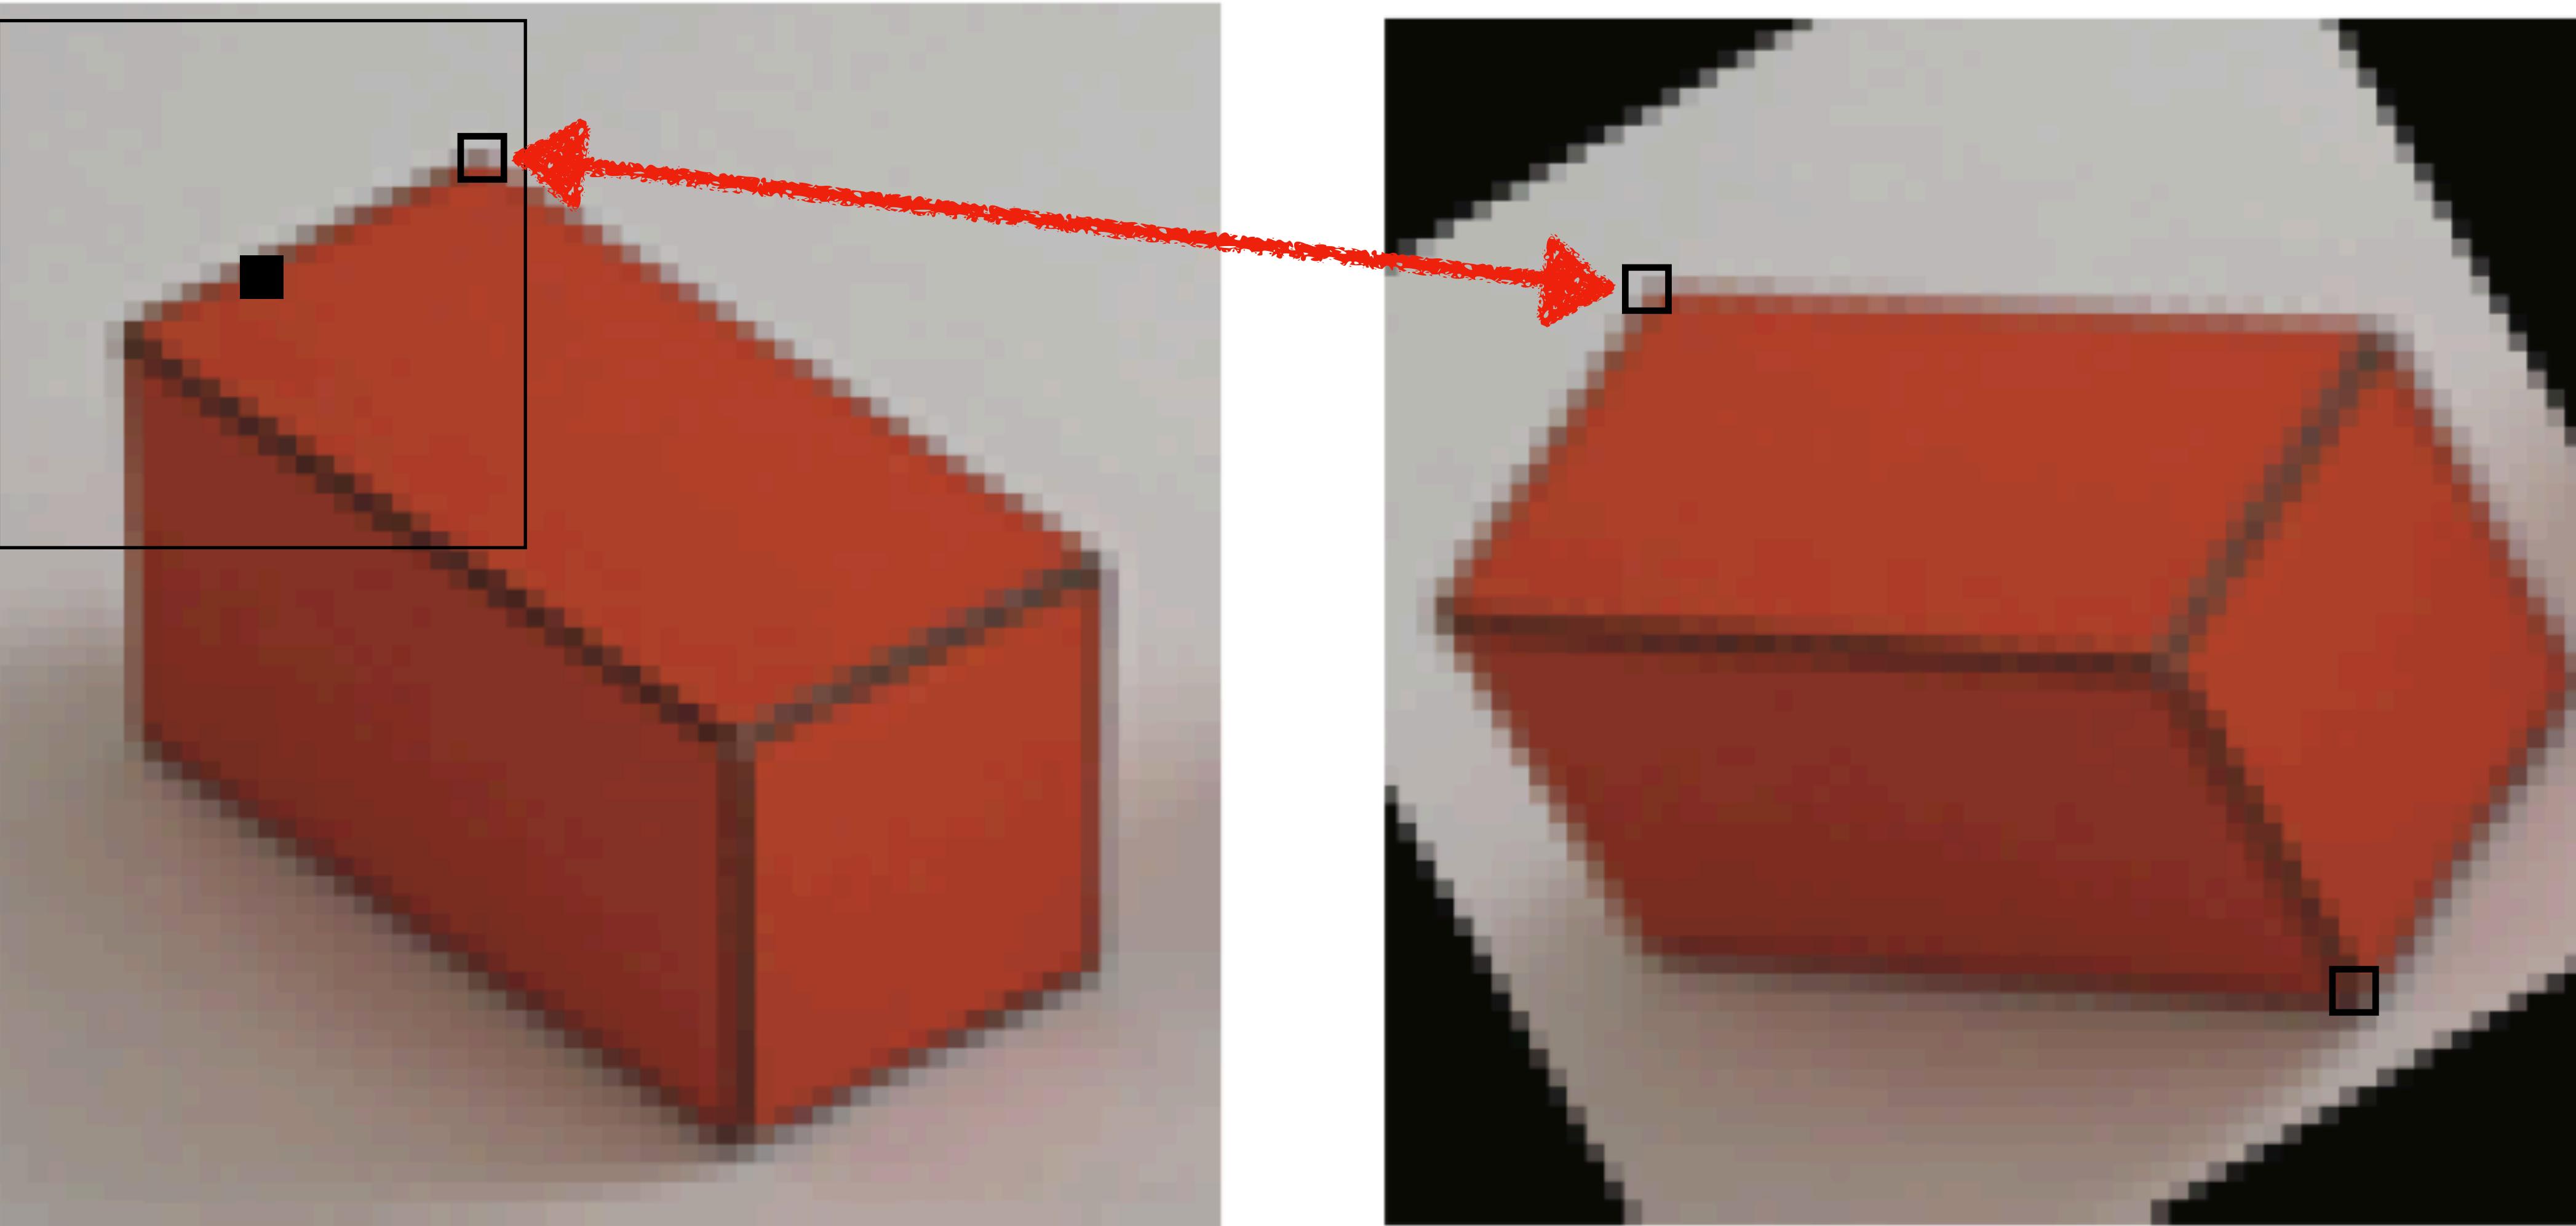
Examples



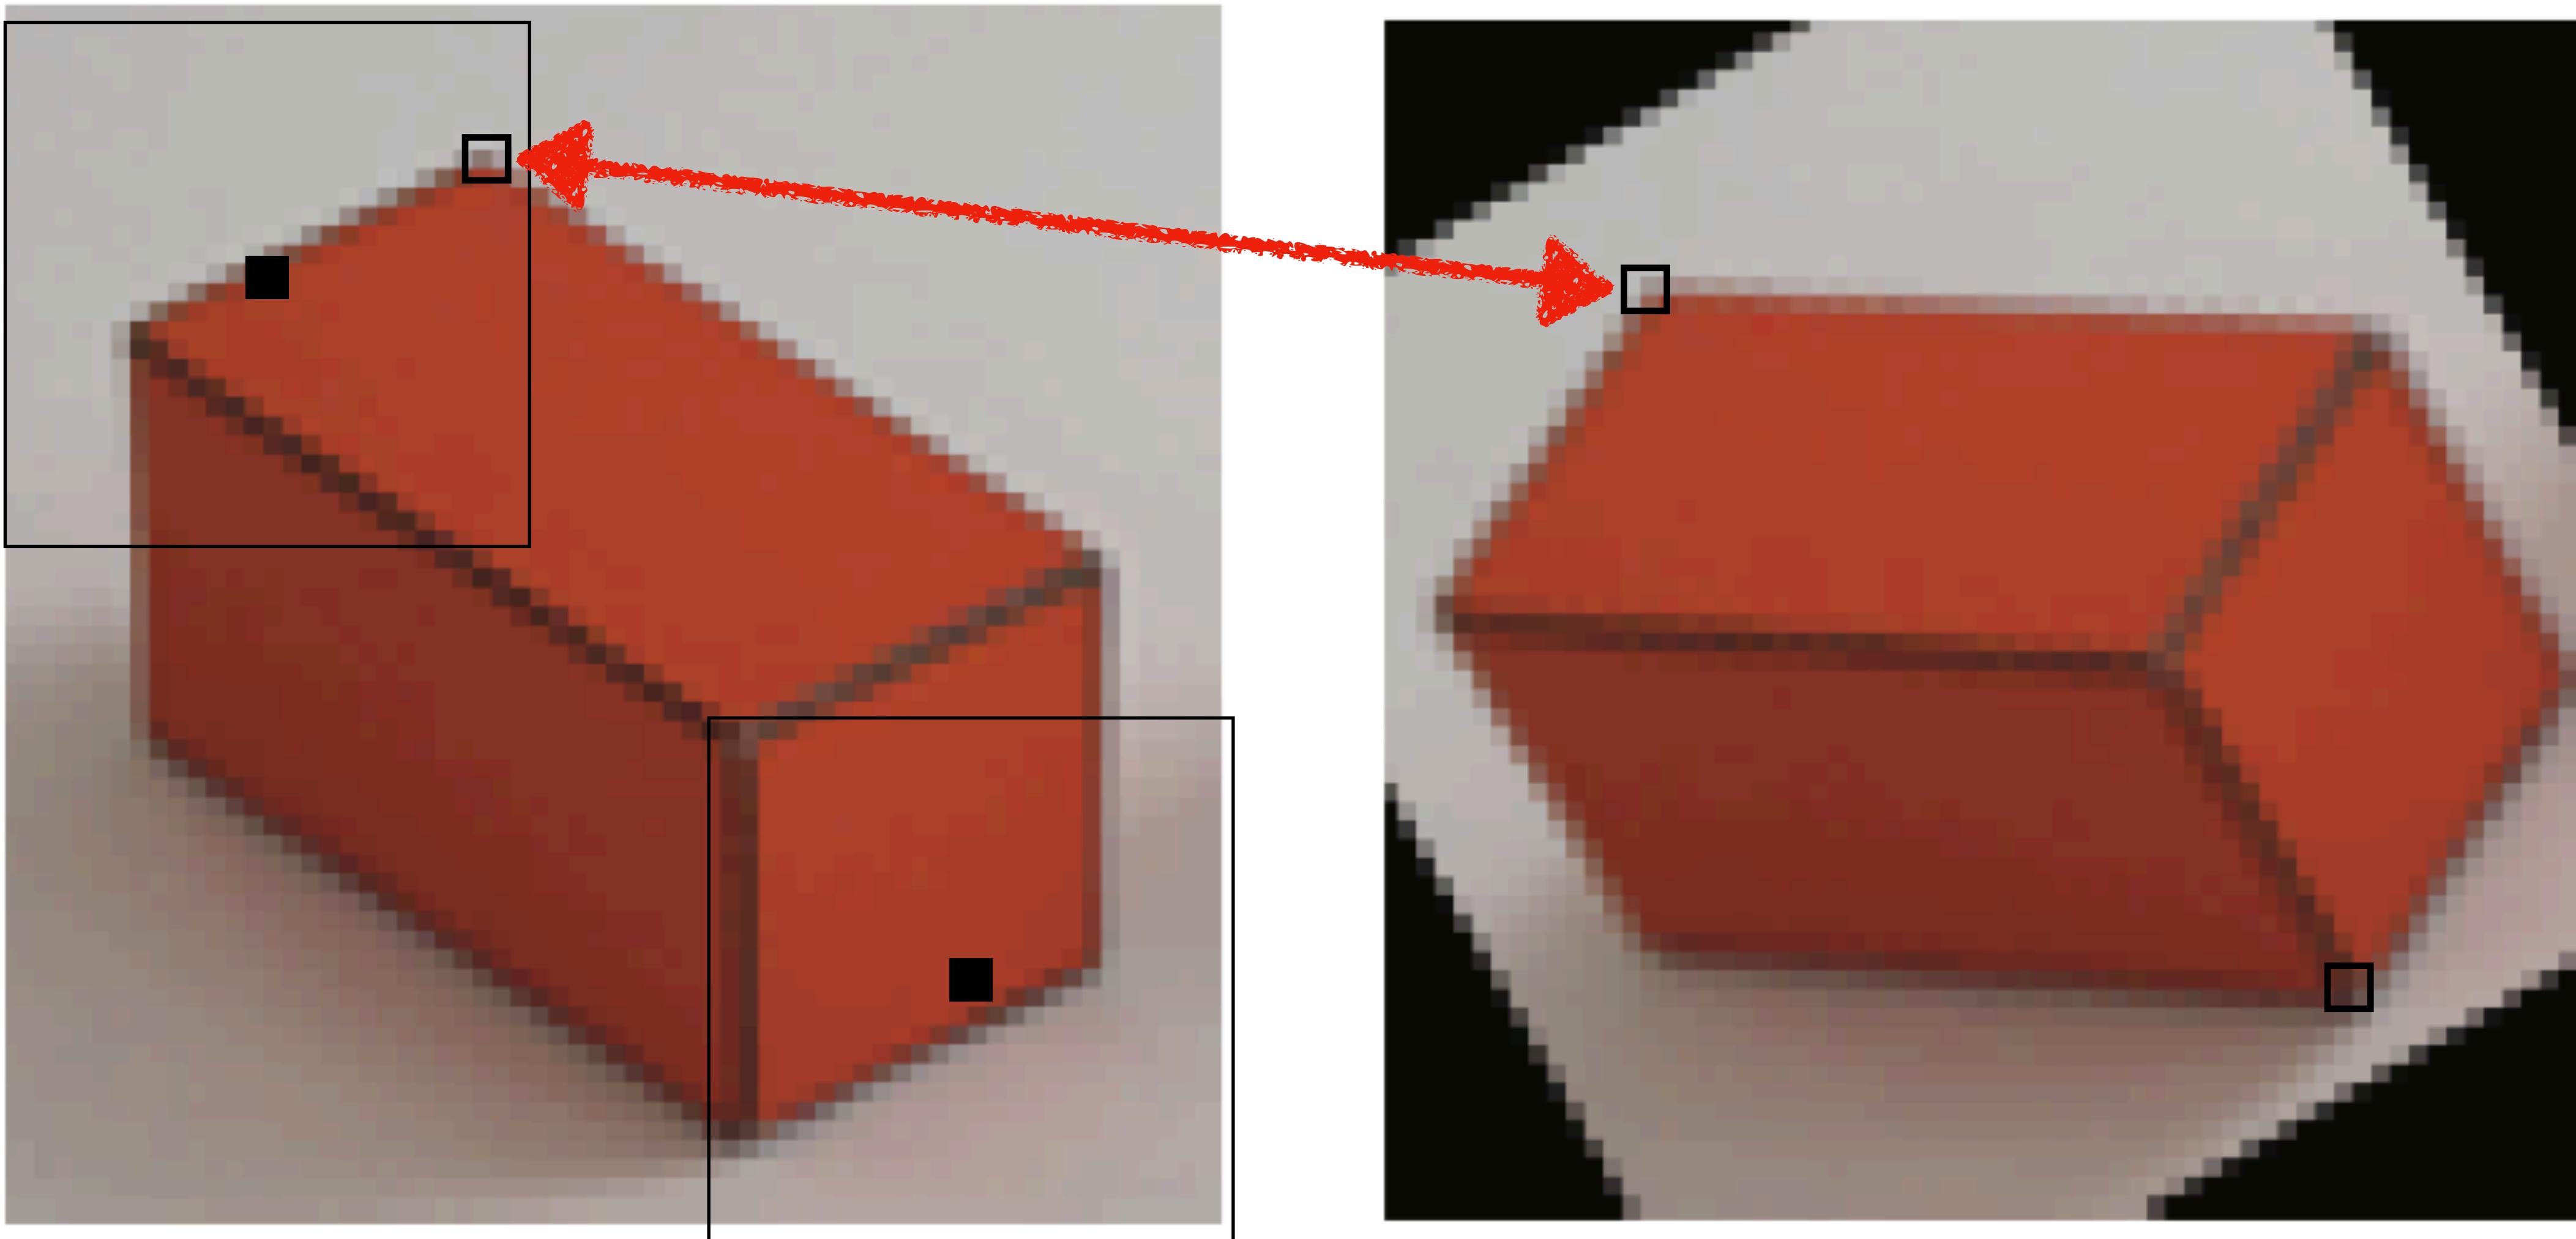
Examples



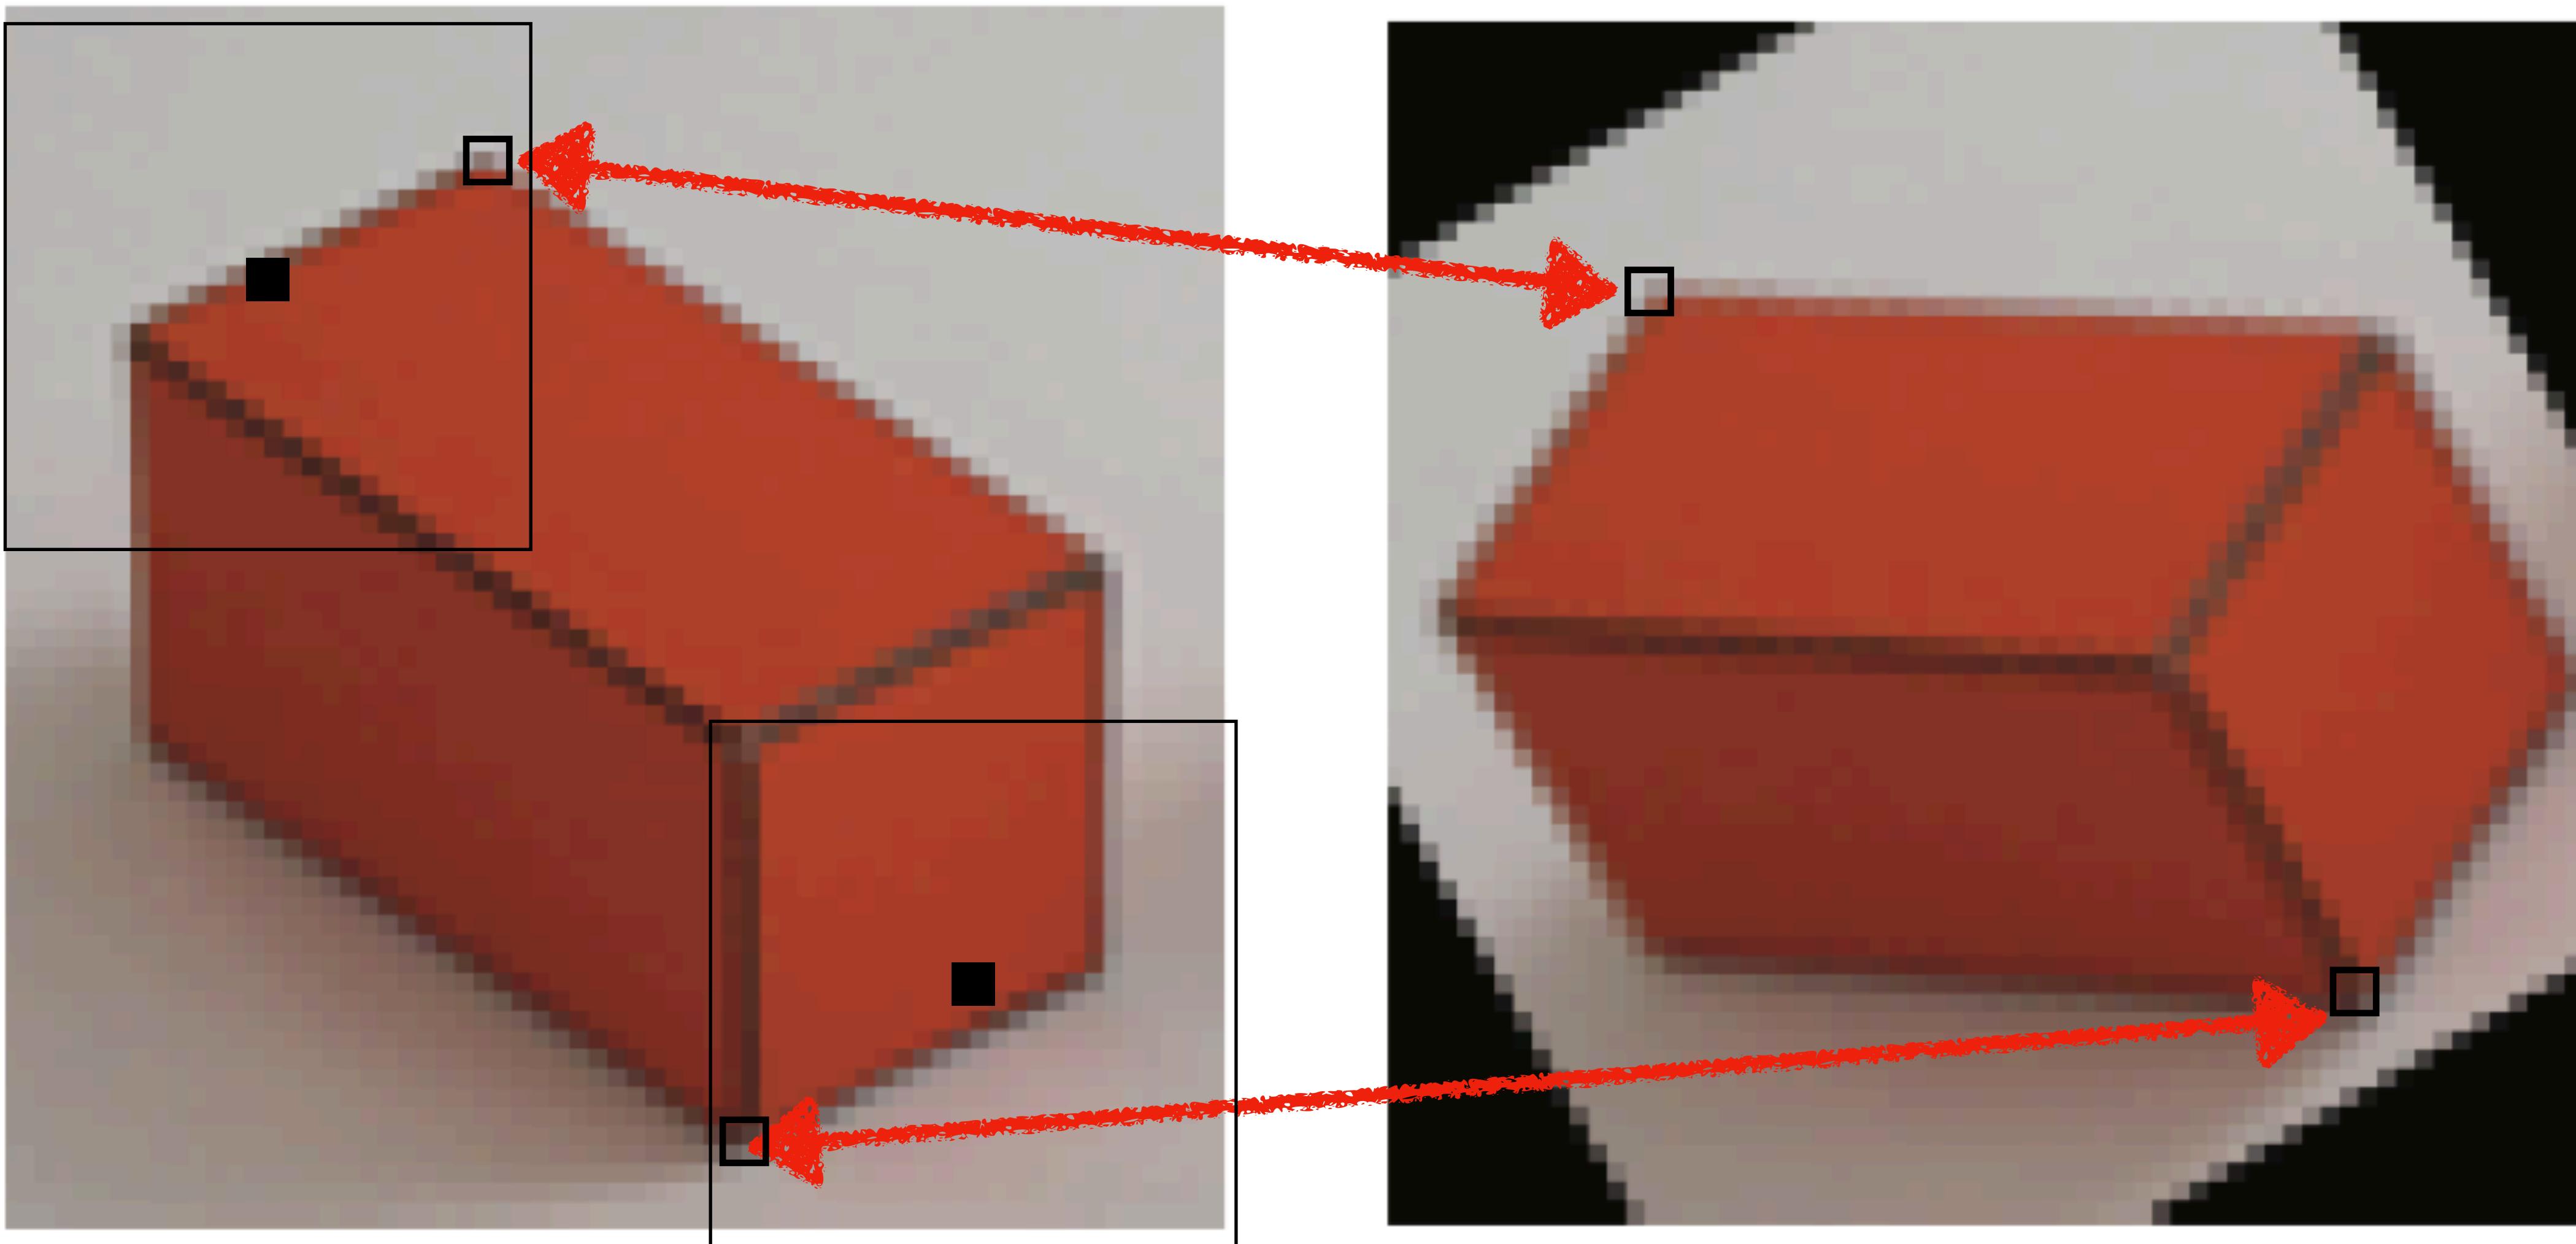
Examples



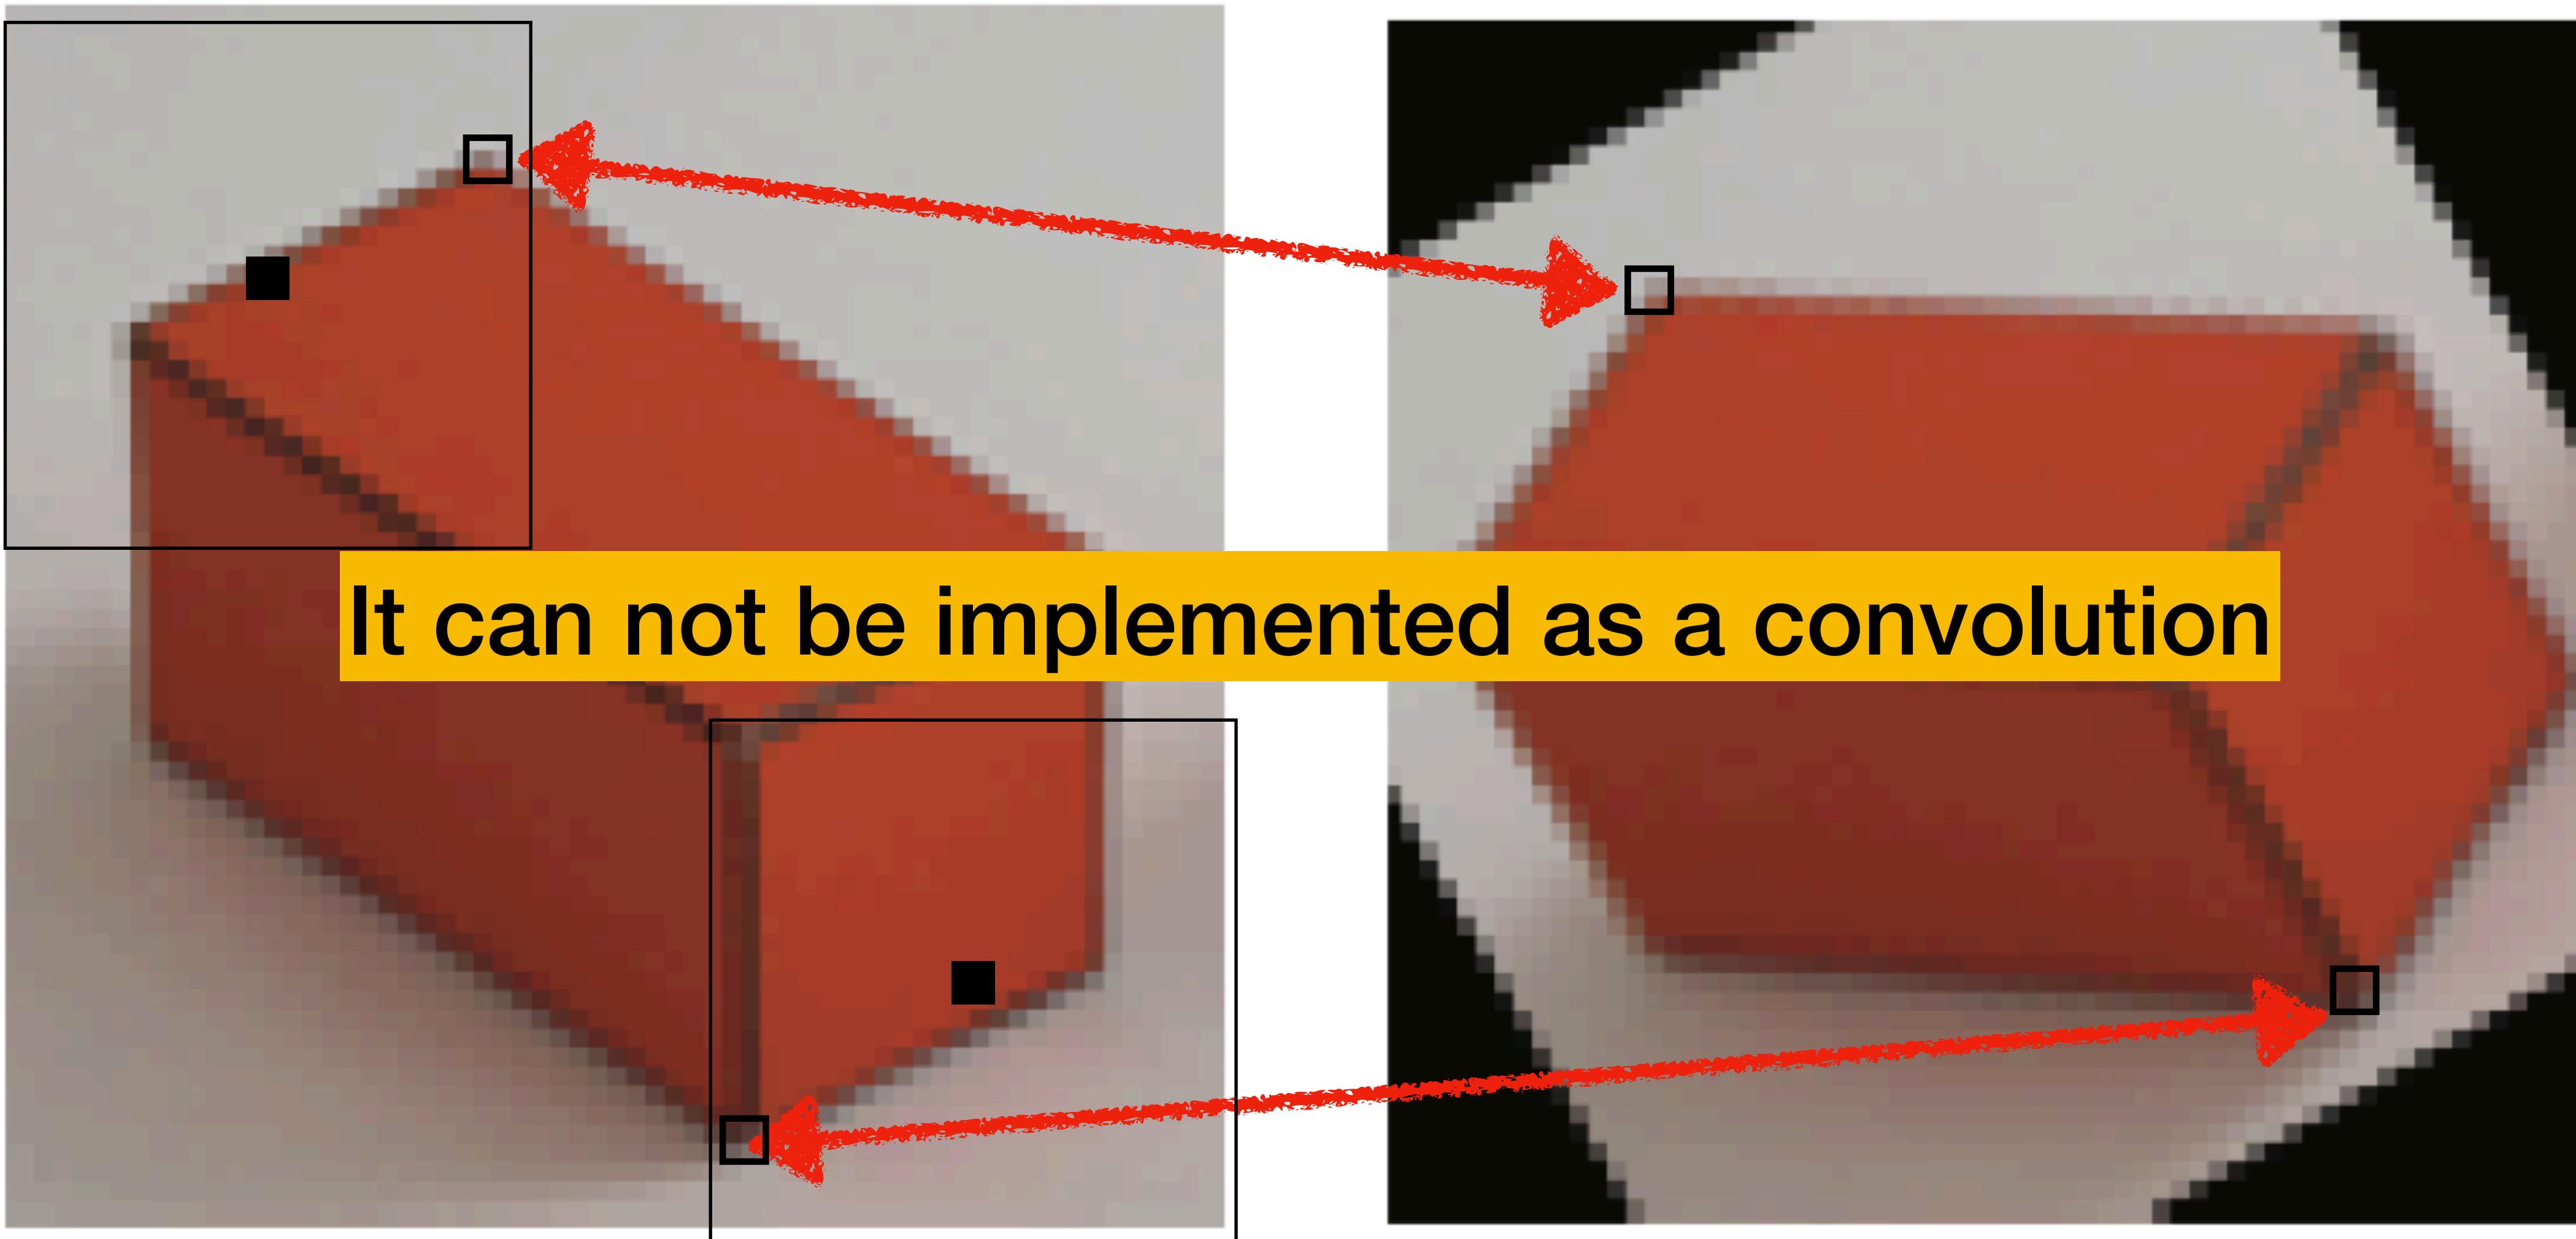
Examples



Examples



Examples



Convolution as matrix multiplication

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

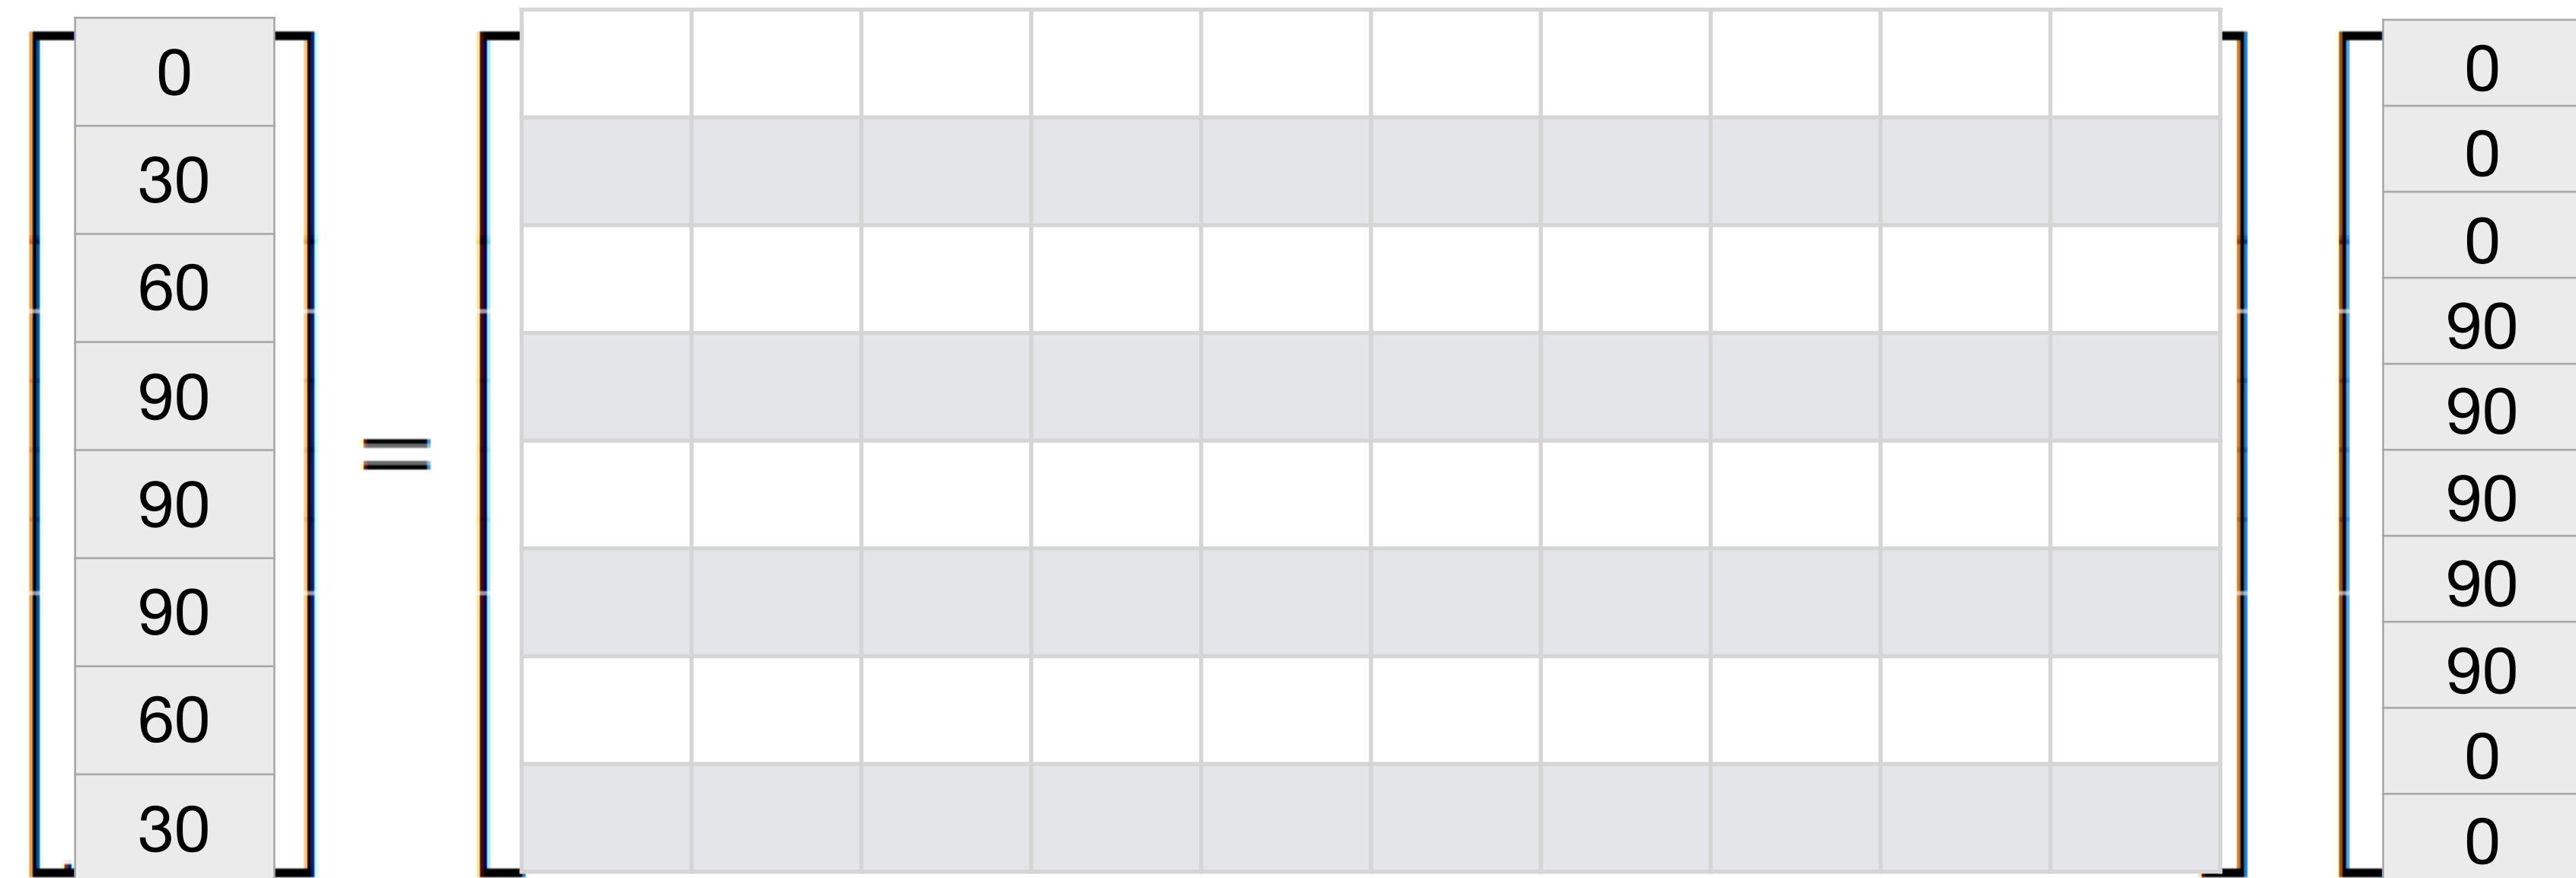
$$\begin{bmatrix} 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \end{bmatrix} \circ \begin{bmatrix} 1/3 & 1/3 & 1/3 \end{bmatrix} = \begin{bmatrix} \quad & 0 & 30 & 60 & 90 & 90 & 90 & 60 & 30 & \quad \end{bmatrix}$$

Convolution as matrix multiplication

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

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

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

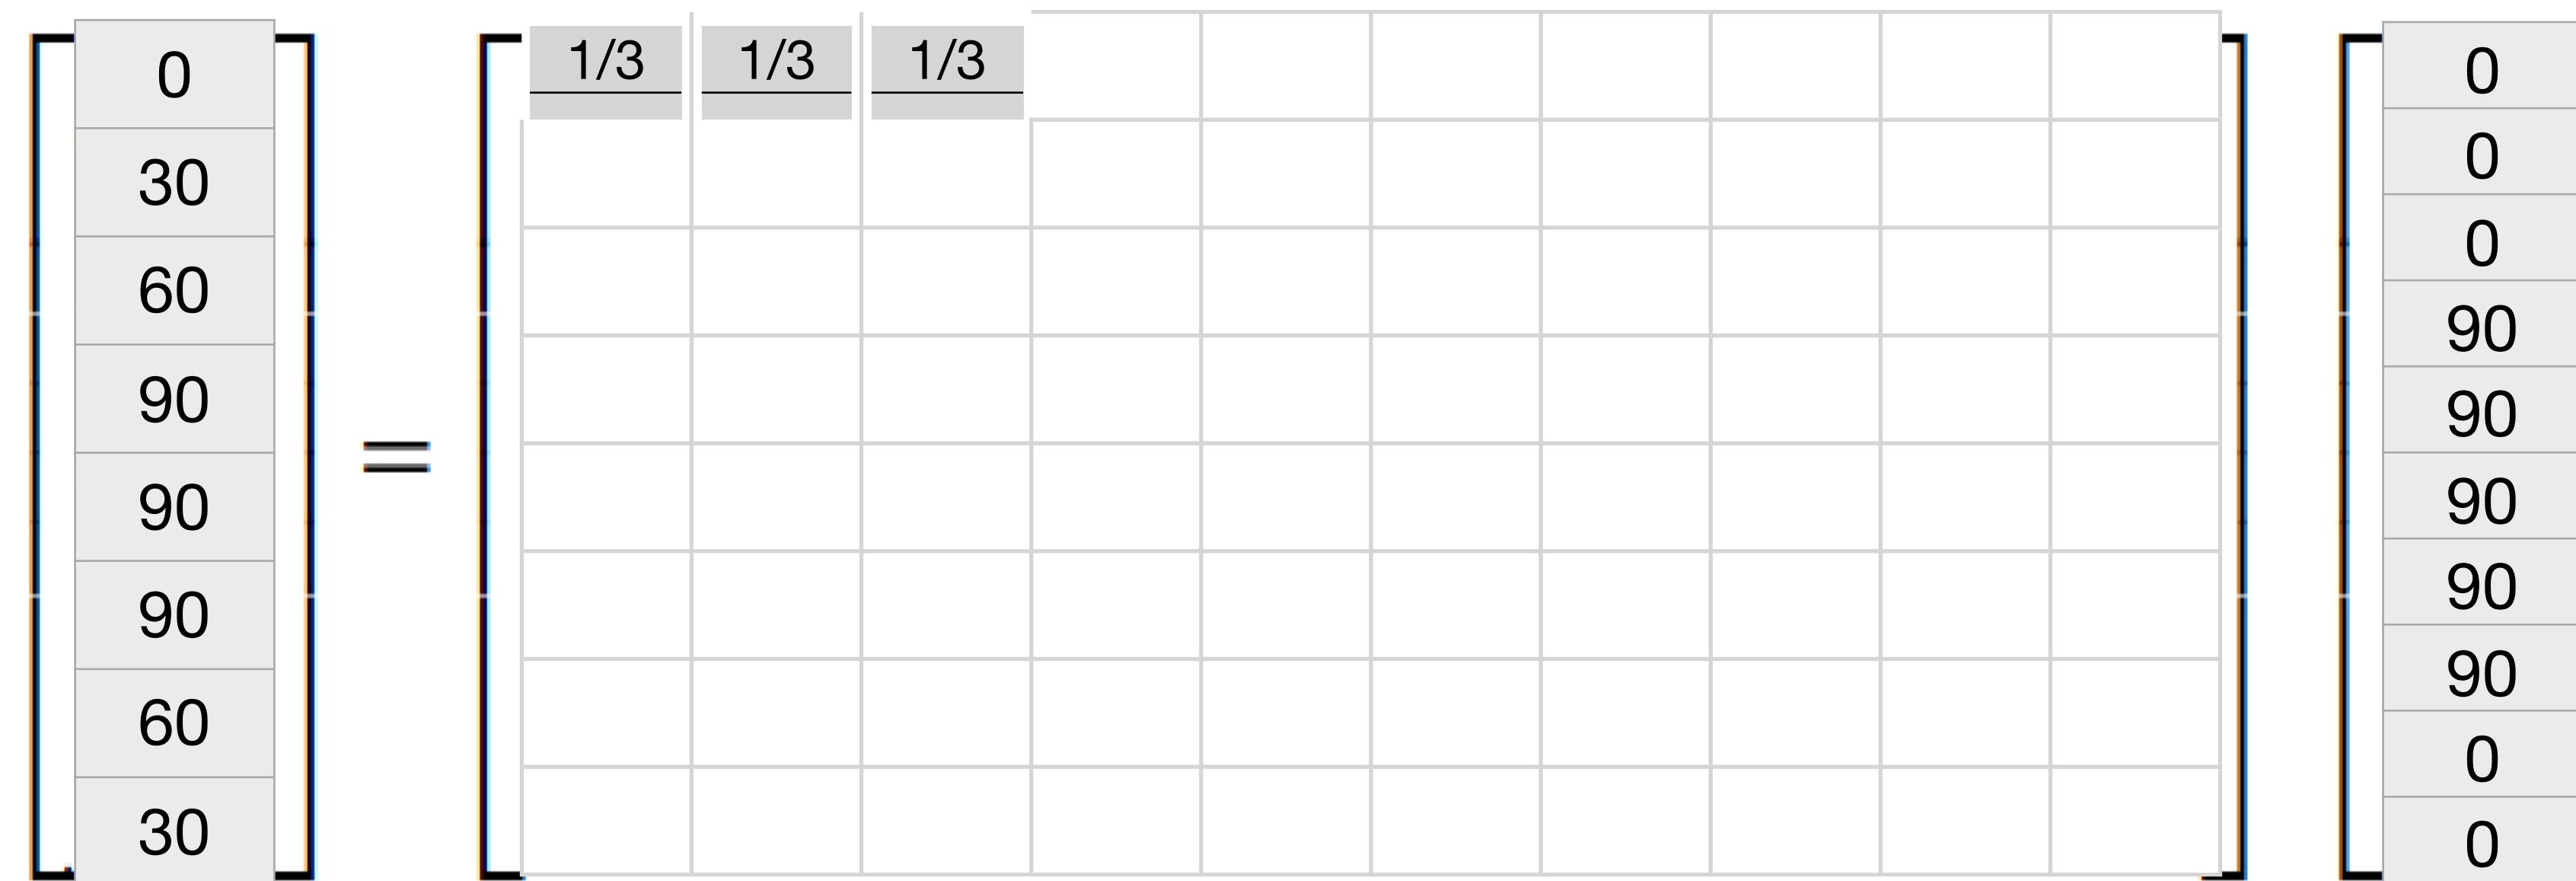


Convolution as matrix multiplication

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

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

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

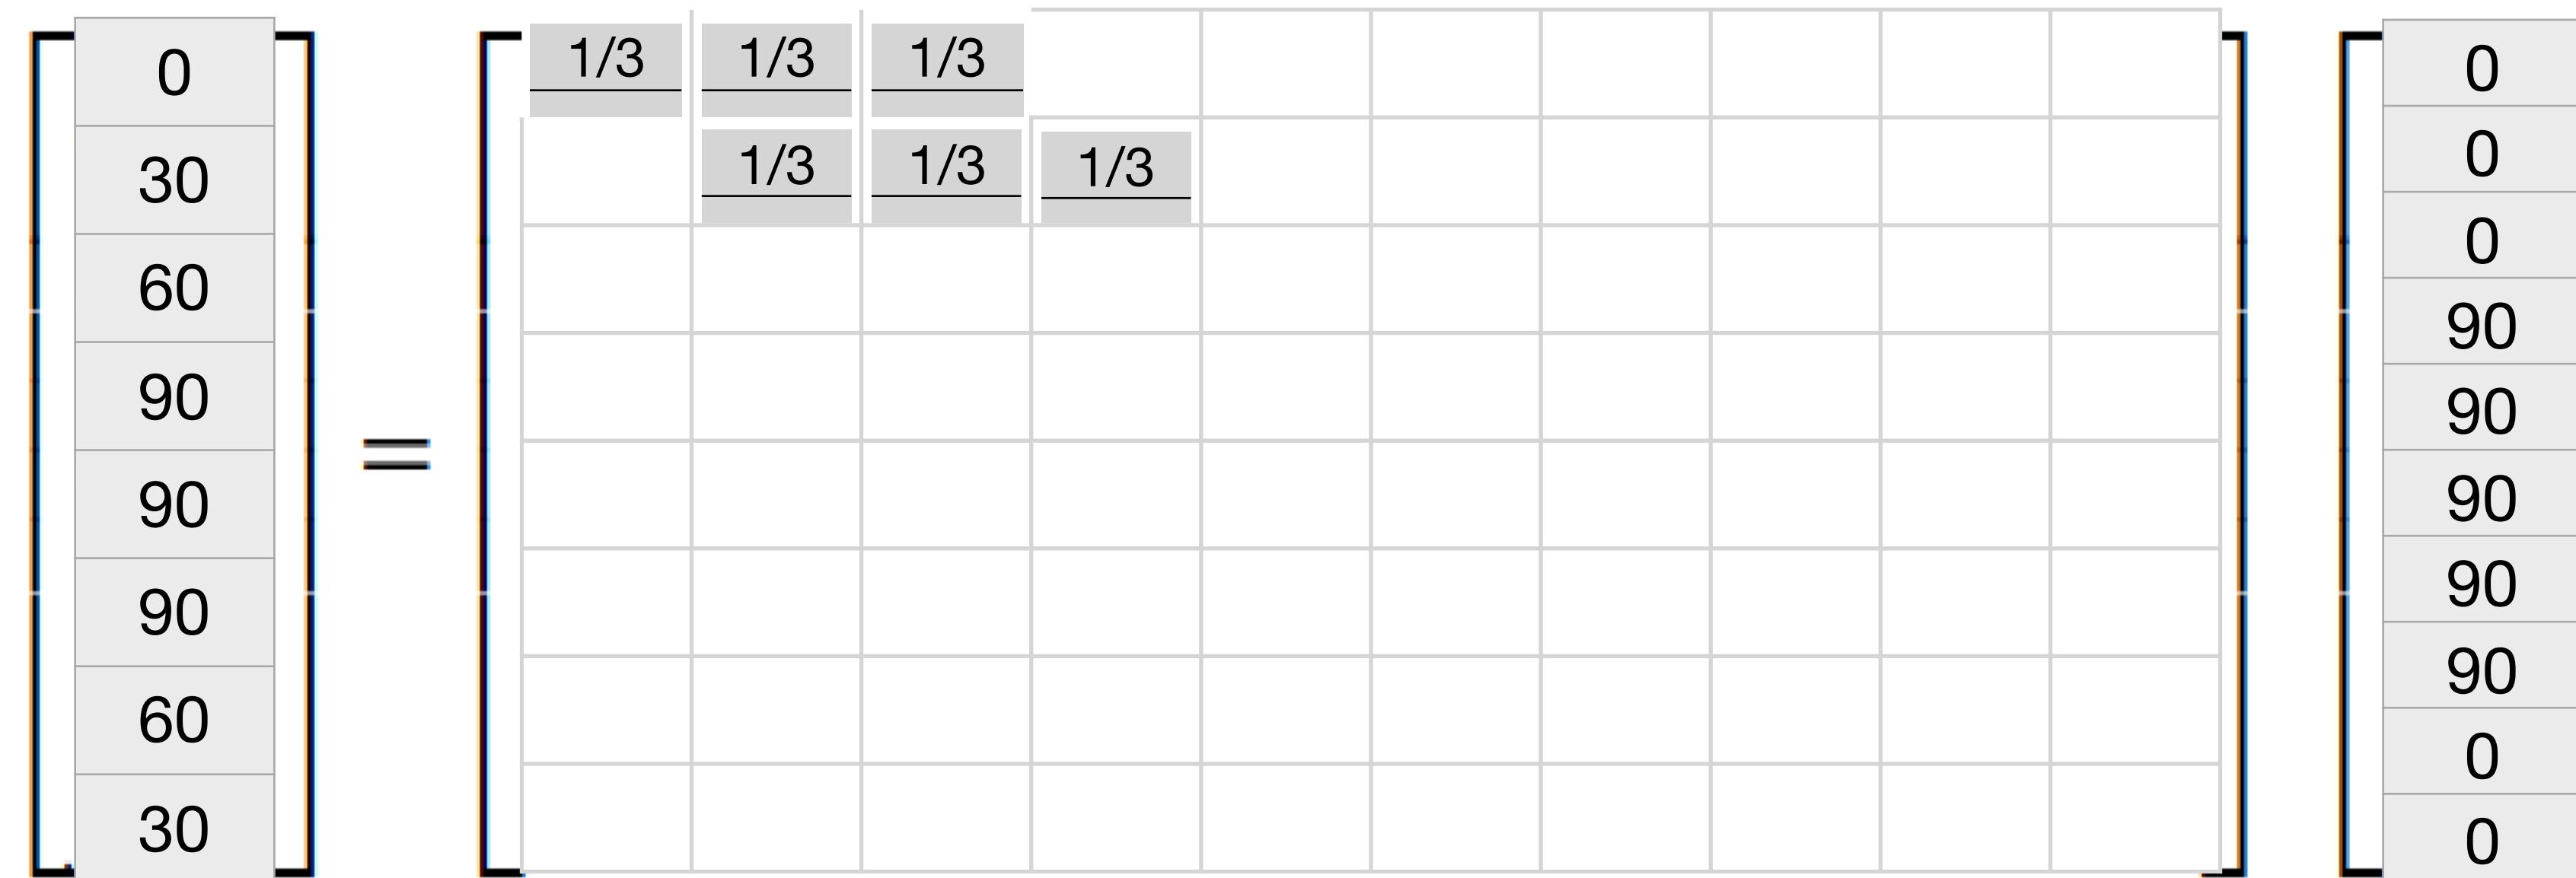


Convolution as matrix multiplication

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

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

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

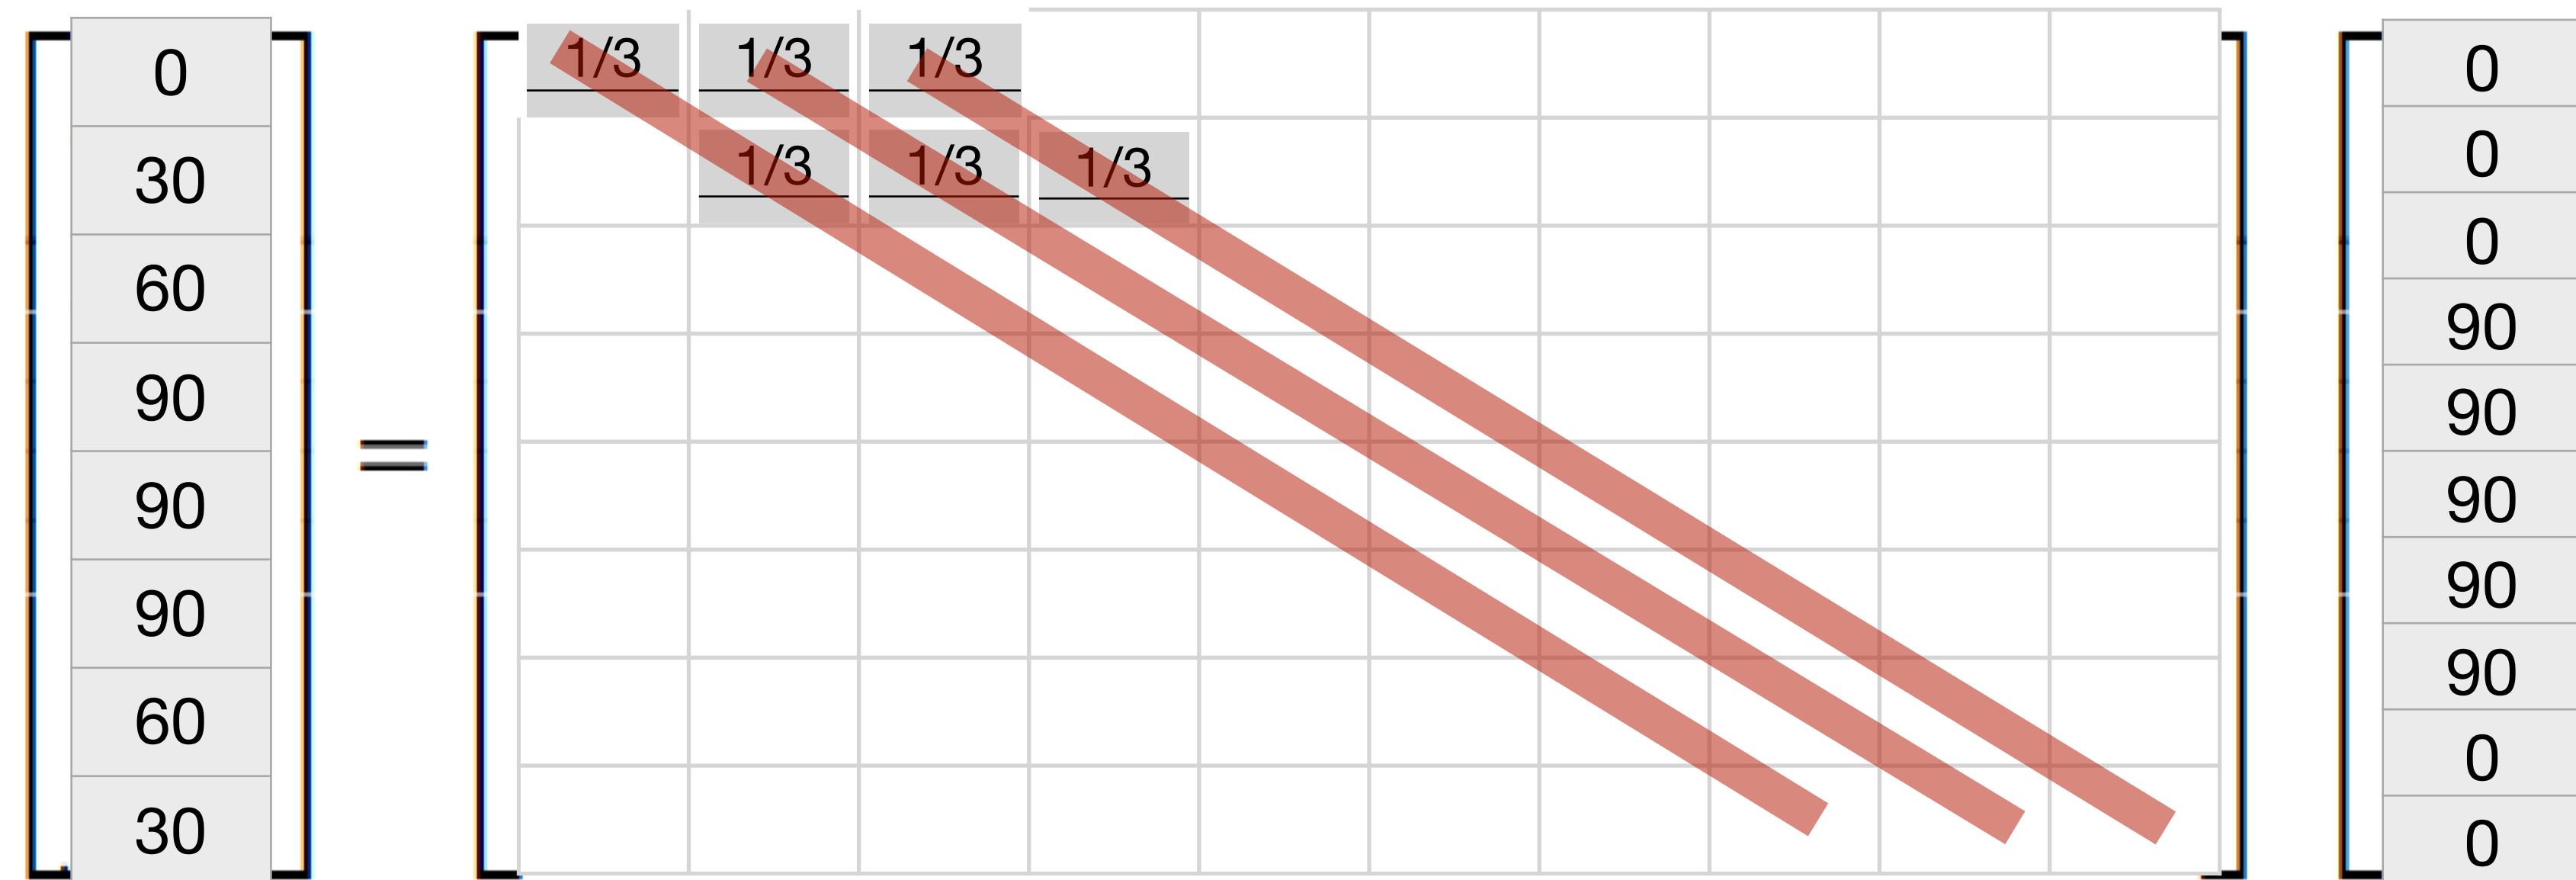


Convolution as matrix multiplication

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

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

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



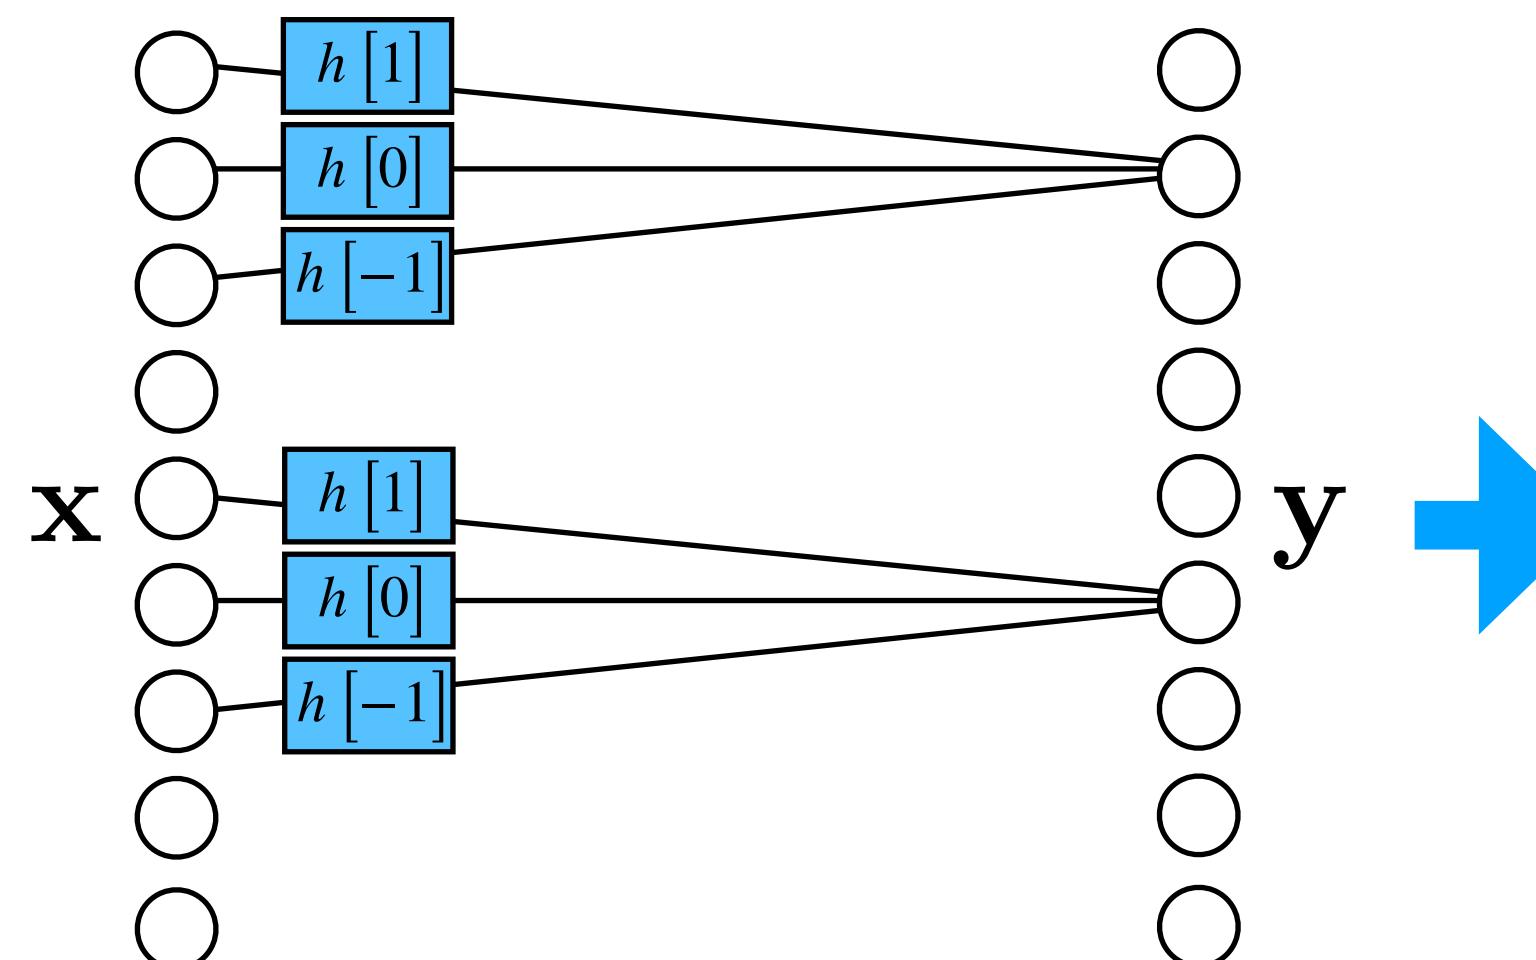
Linear translation invariant system:

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

A diagram illustrating a convolutional layer's receptive field. A 3x3 input grid is shown with red diagonal stripes representing the receptive fields of the central output unit. Three blue boxes labeled $h[1]$, $h[0]$, and $h[-1]$ indicate the receptive fields of the input units at the bottom-left, middle, and top-right respectively. The x and y axes are labeled.

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

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

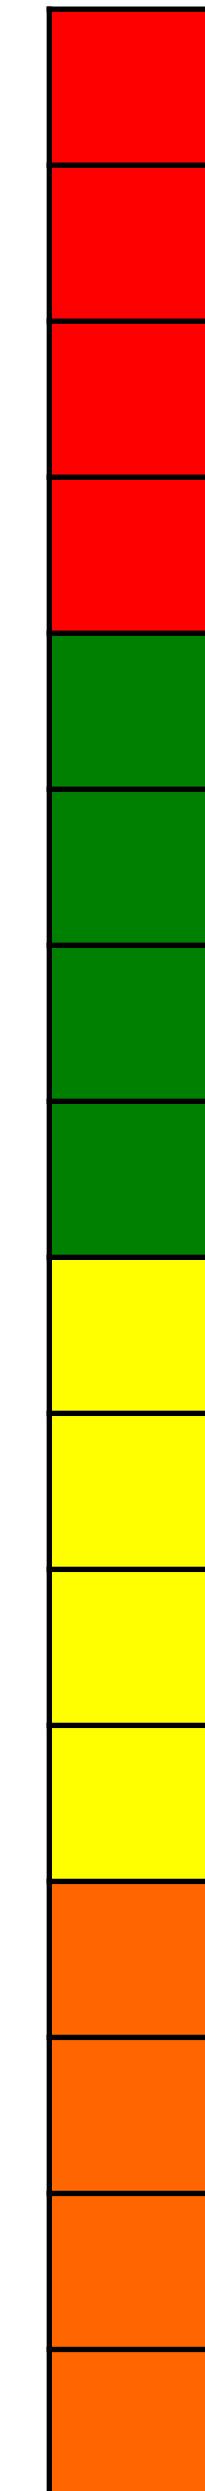
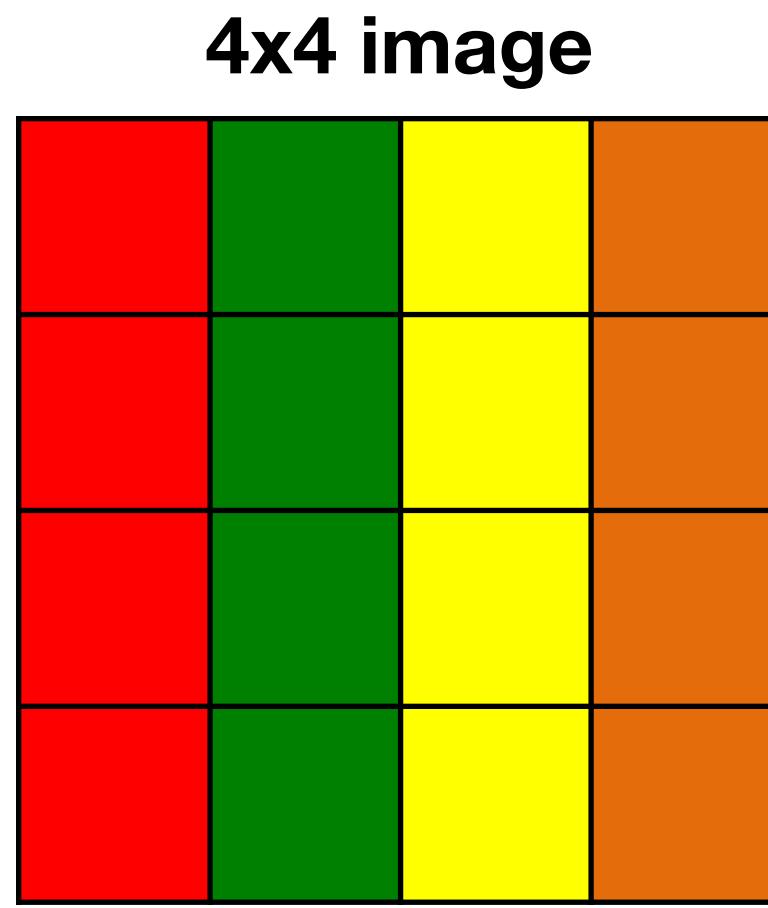


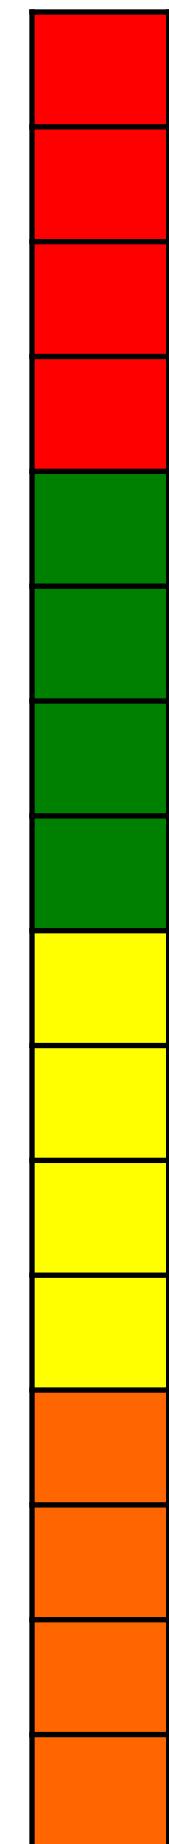
$$= \sum_{k=-1}^1 h[k] x[n-k]$$

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

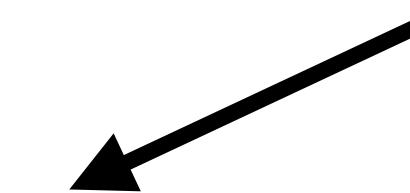
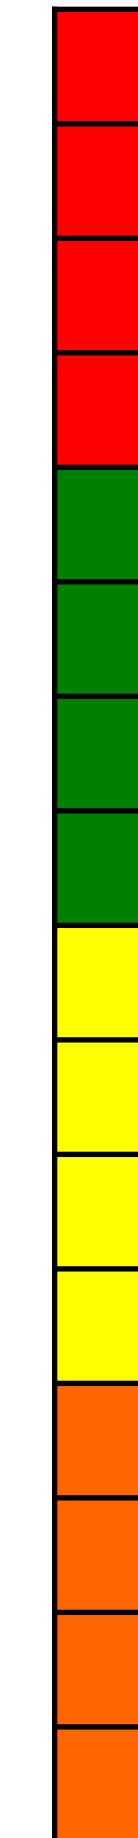
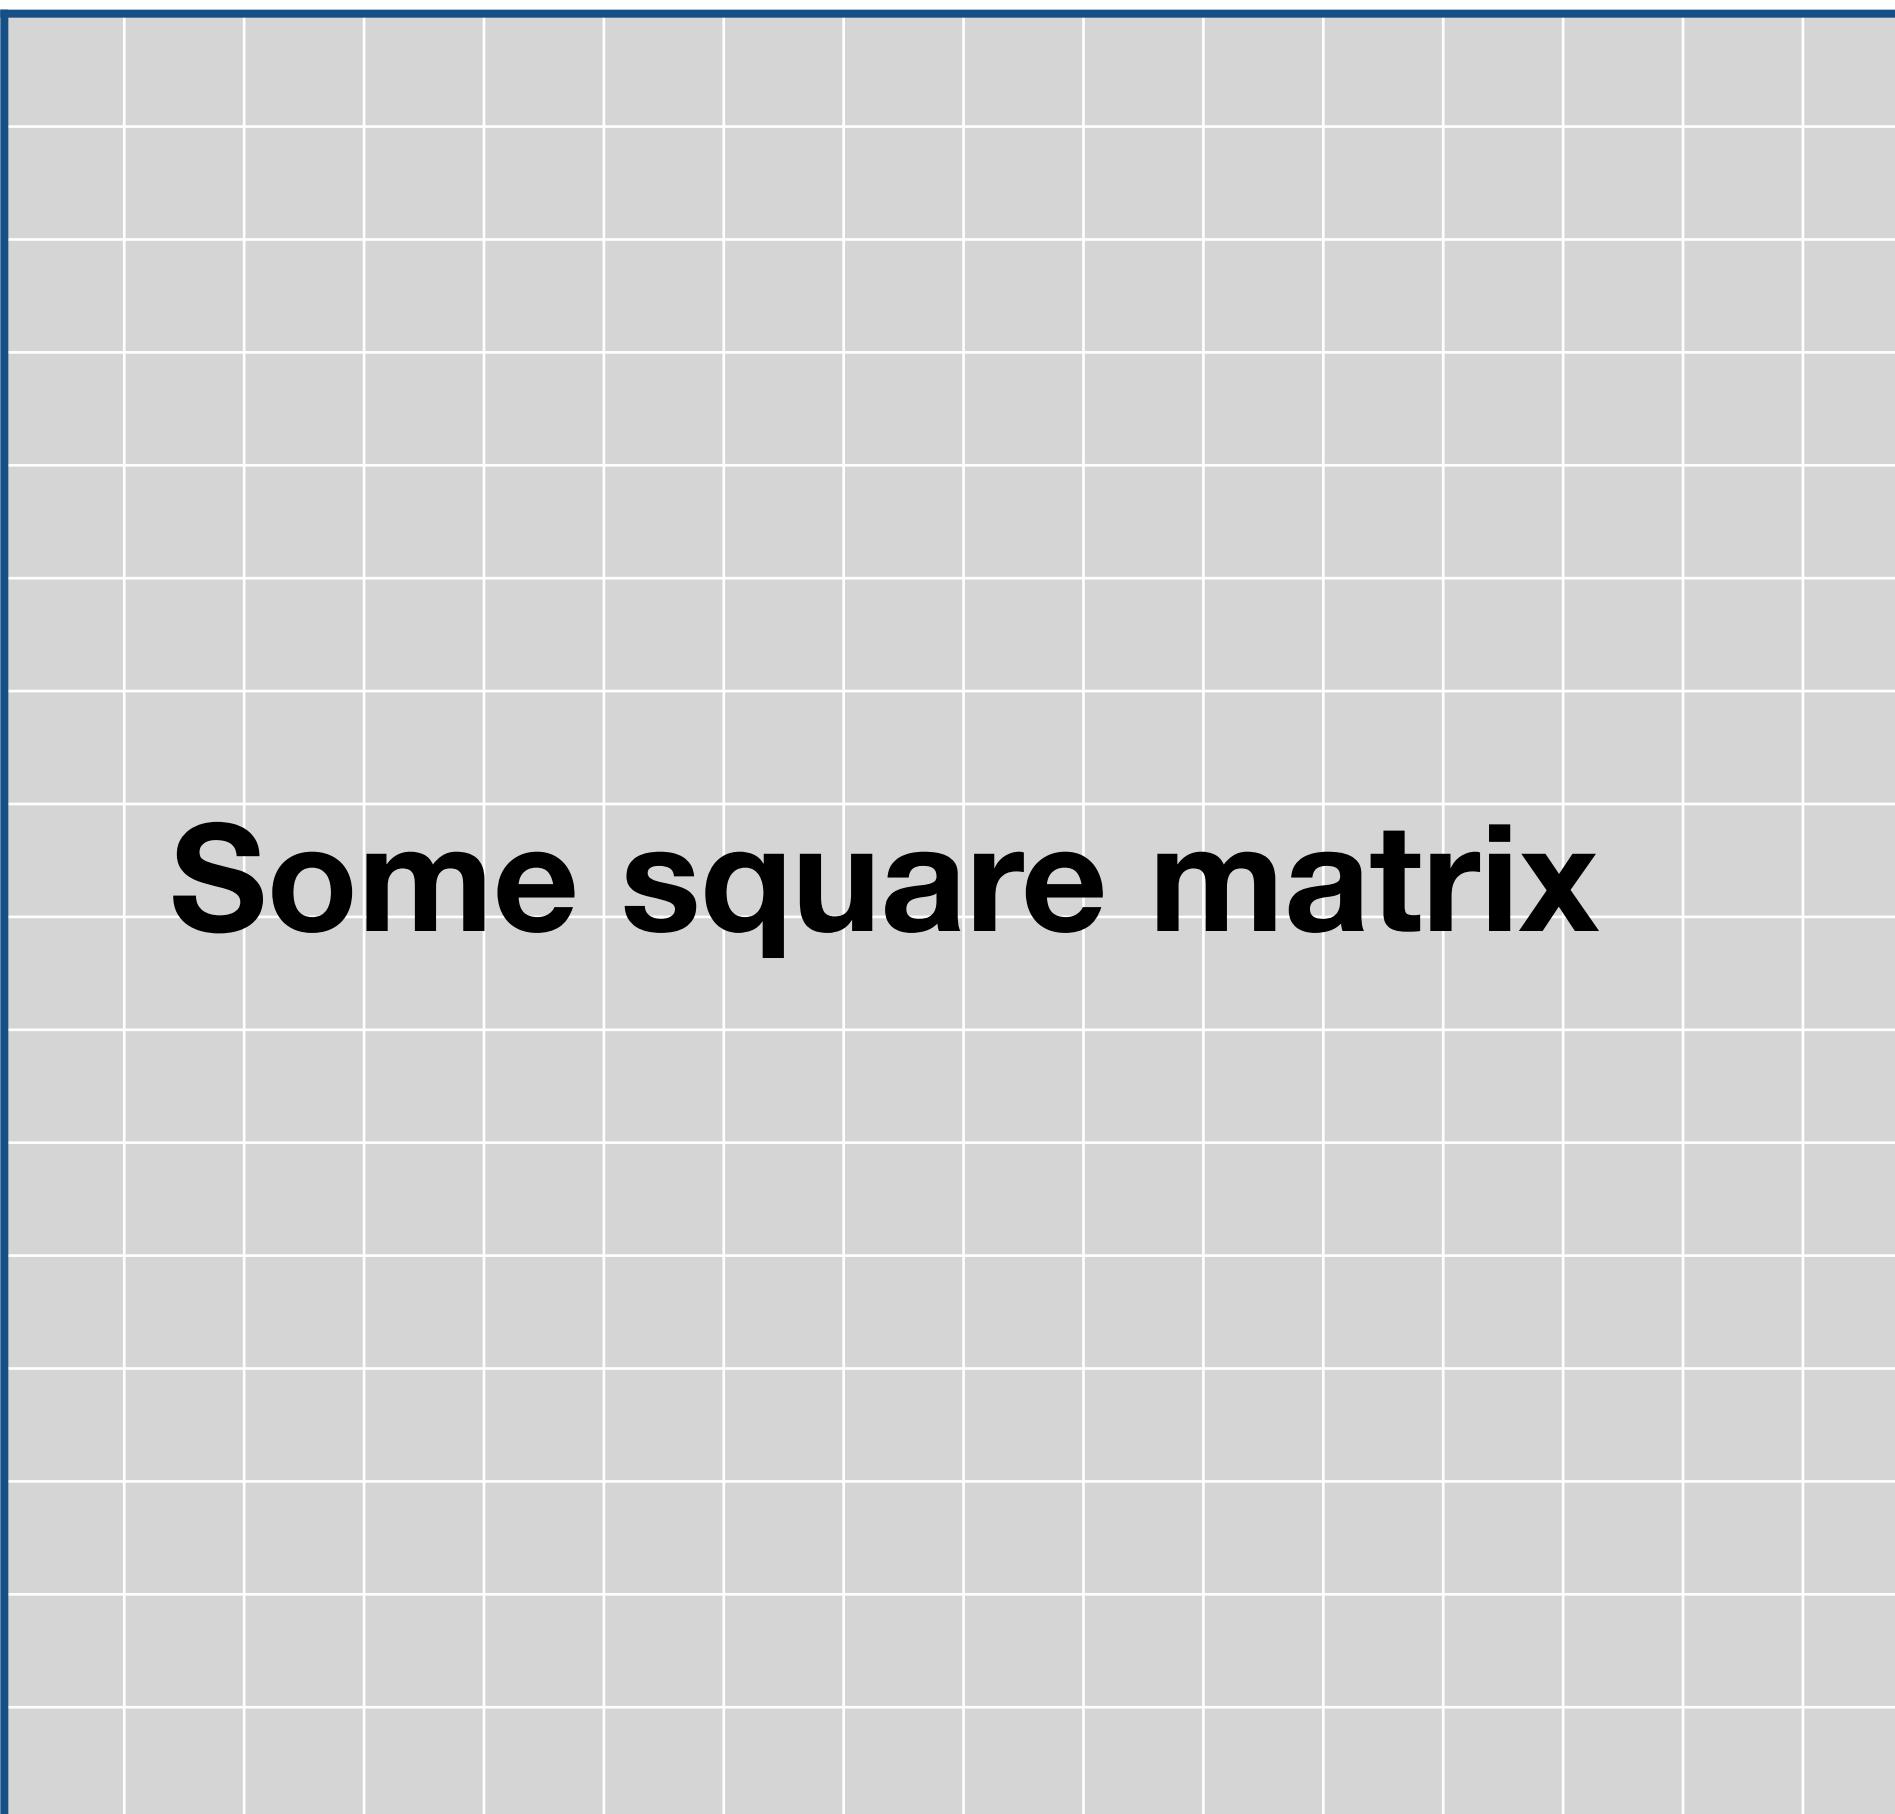
Images are turned into column vectors
by concatenating all image columns

Column vector of length 16

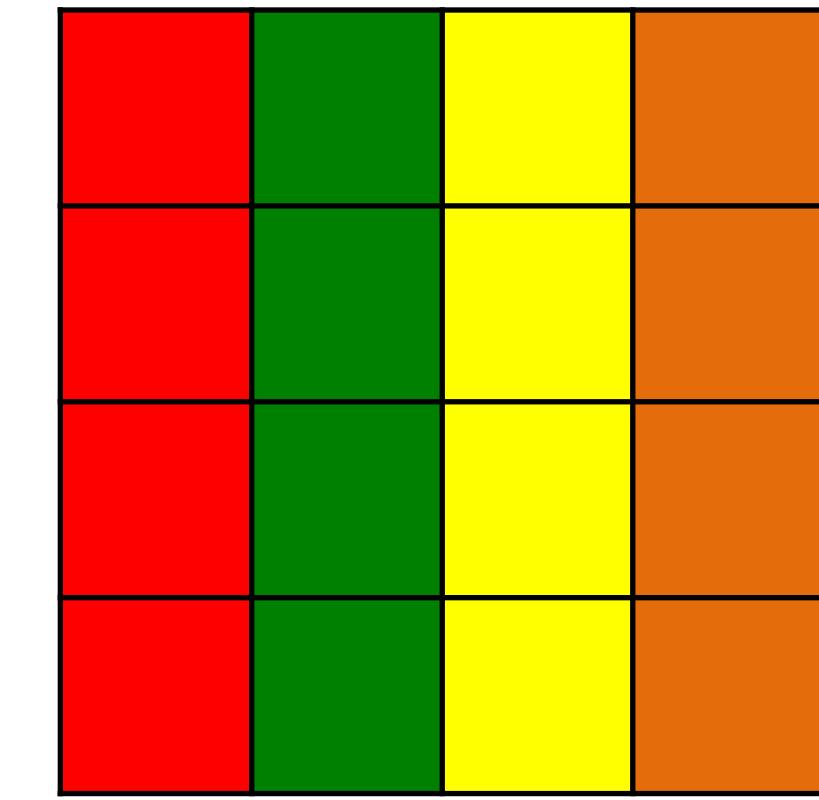




=

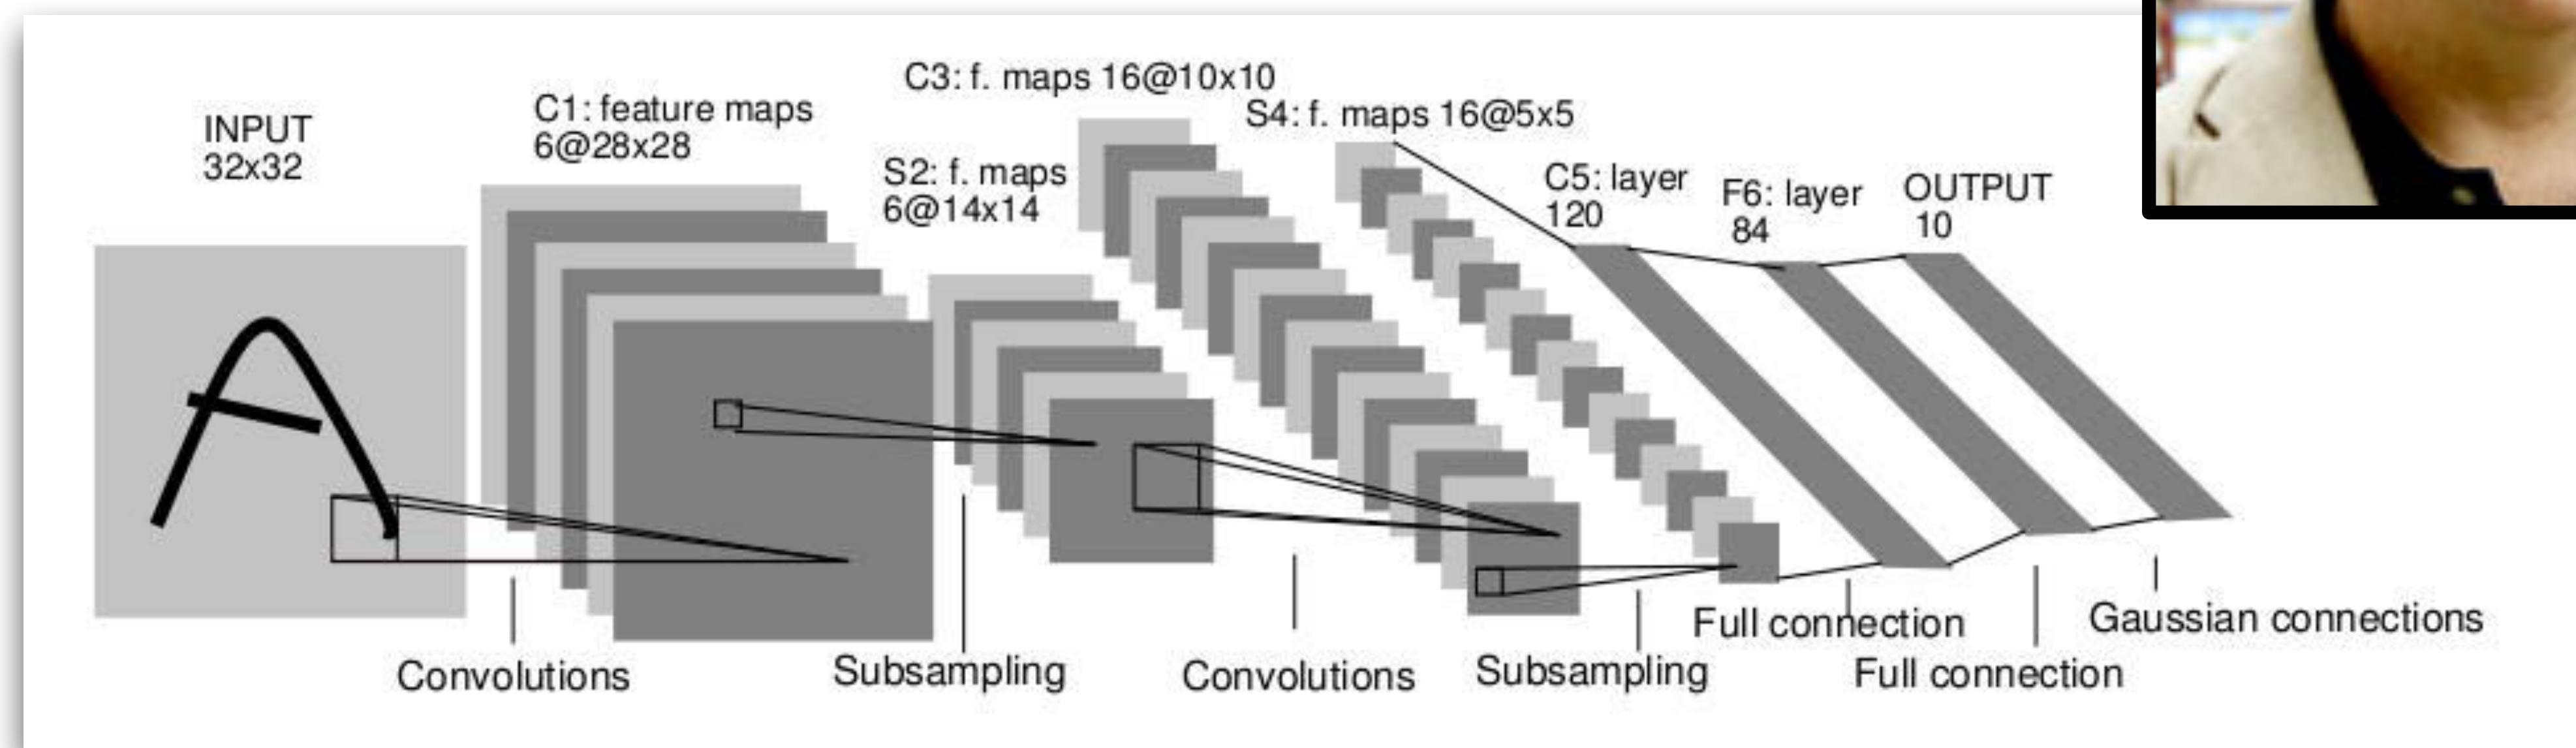


4x4 image



Convolutional Neural Networks

- Neural network with specialized connectivity structure

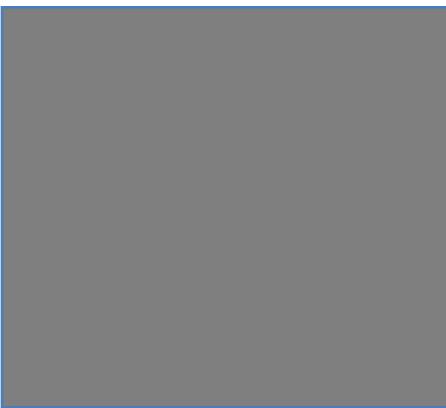


LeCun et al. 1989

Rectangular filter



$g[m,n]$



$h[m,n]$

=



$f[m,n]$

What does it do?

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

$$\frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

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

Input image

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

Convolution output

Rectangular filter



$g[m,n]$



$h[m,n]$



$f[m,n]$

Rectangular filter



$g[m,n]$

\otimes

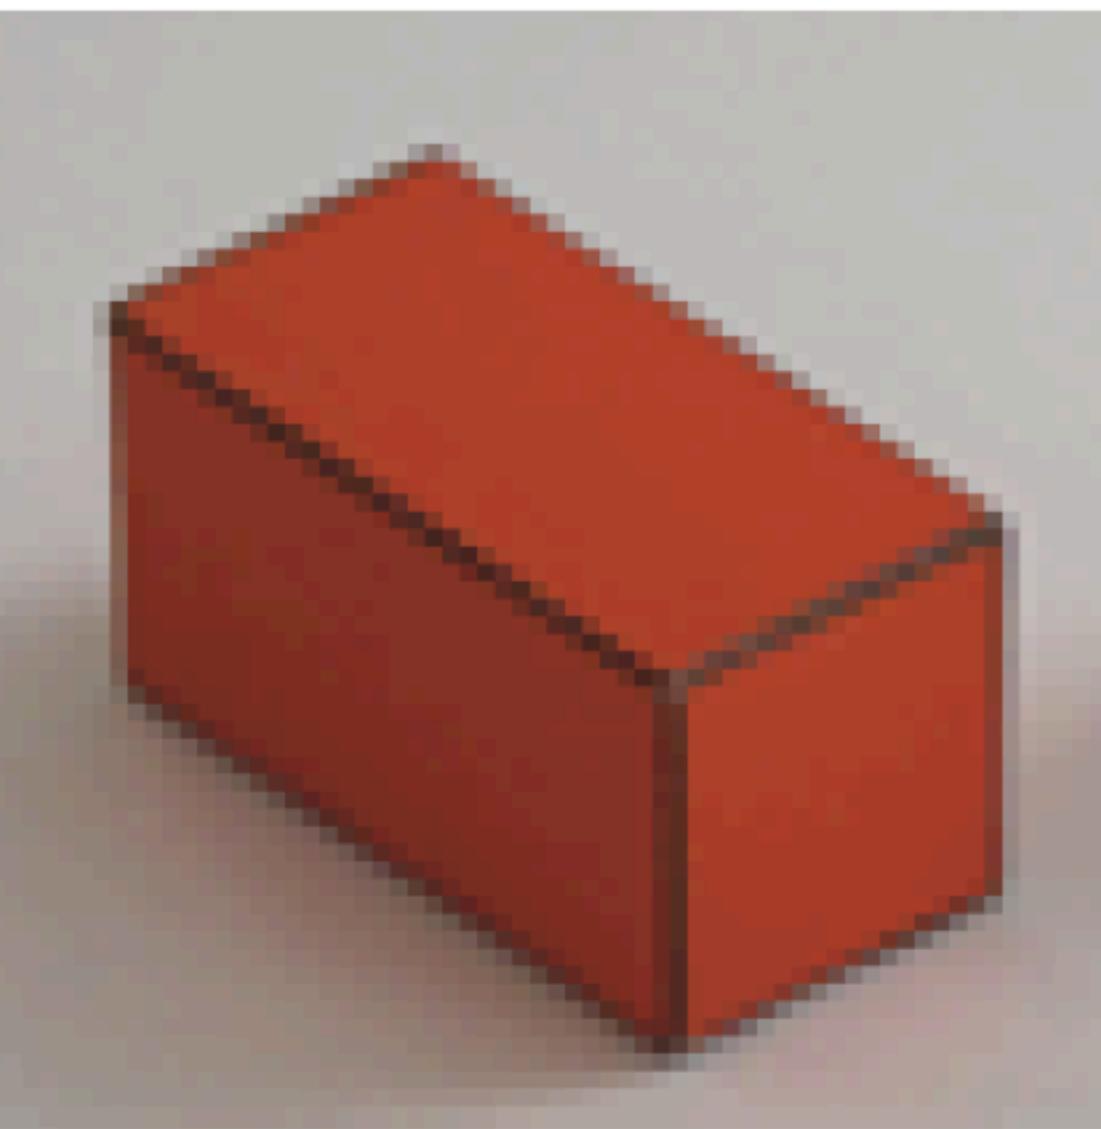
$h[m,n]$

=



$f[m,n]$

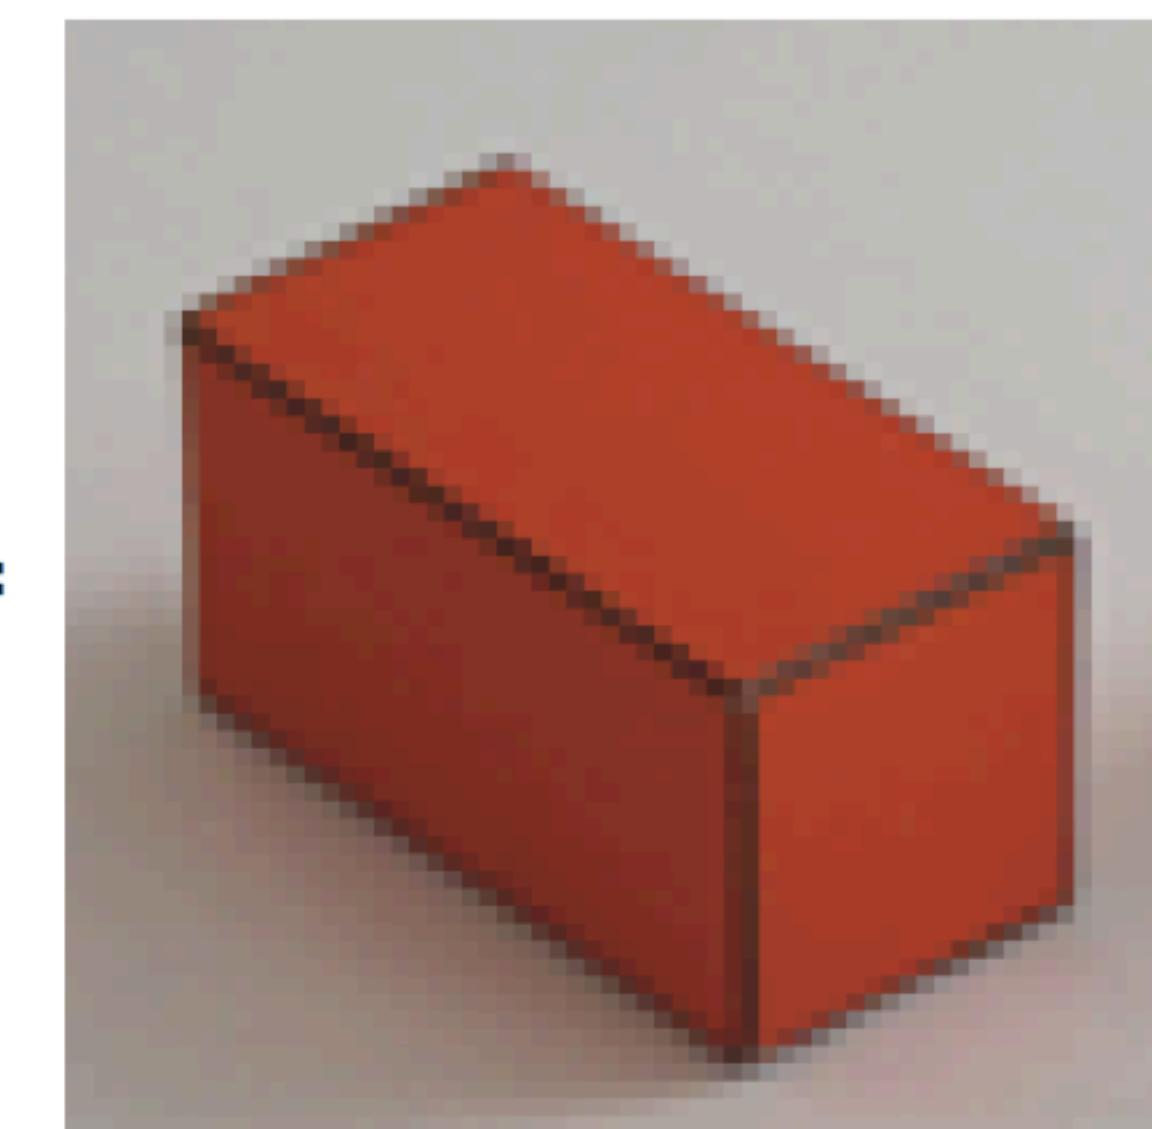
The identity



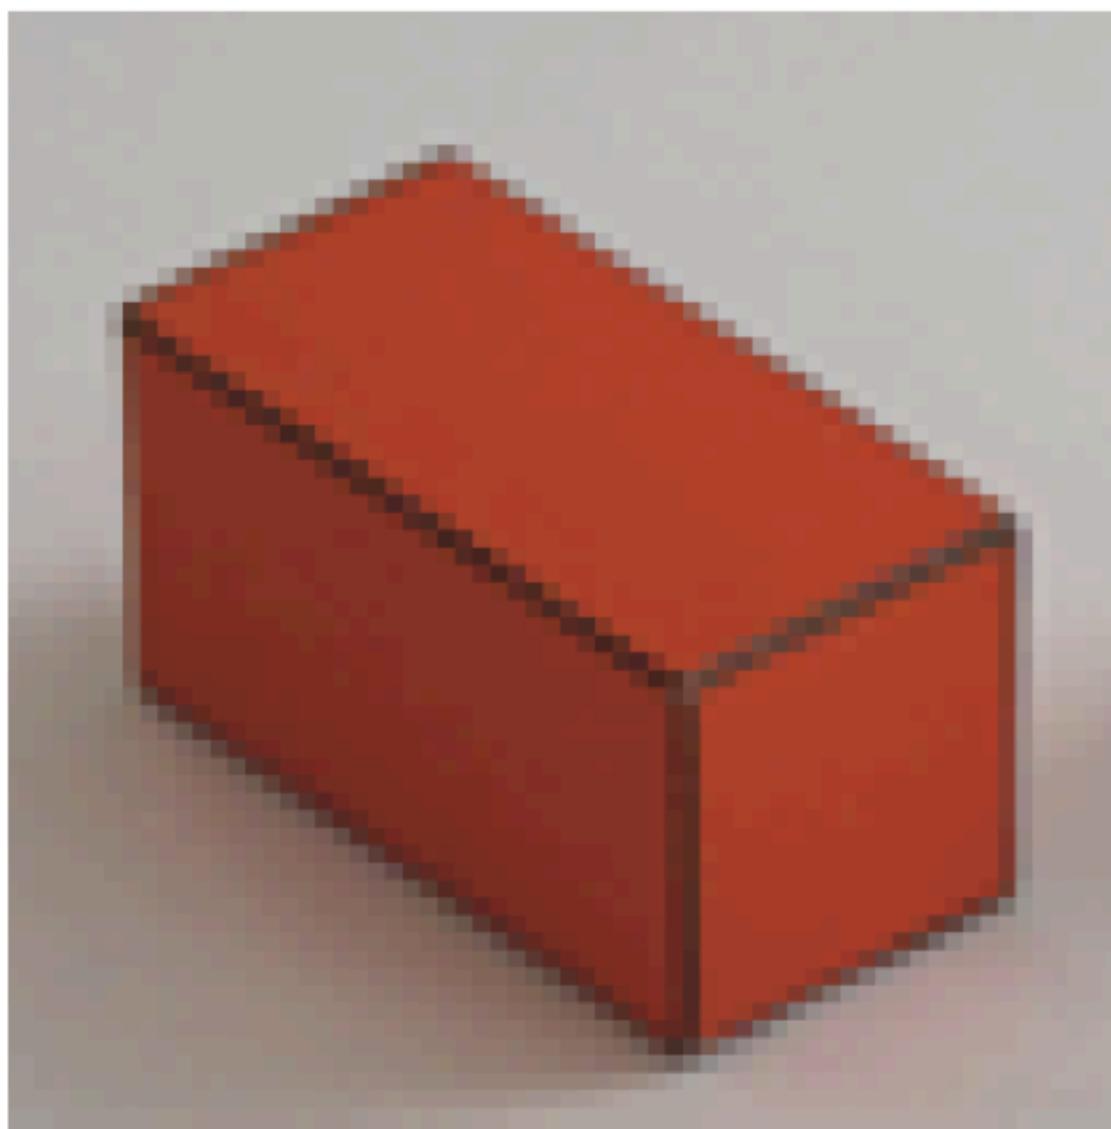
○

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

=



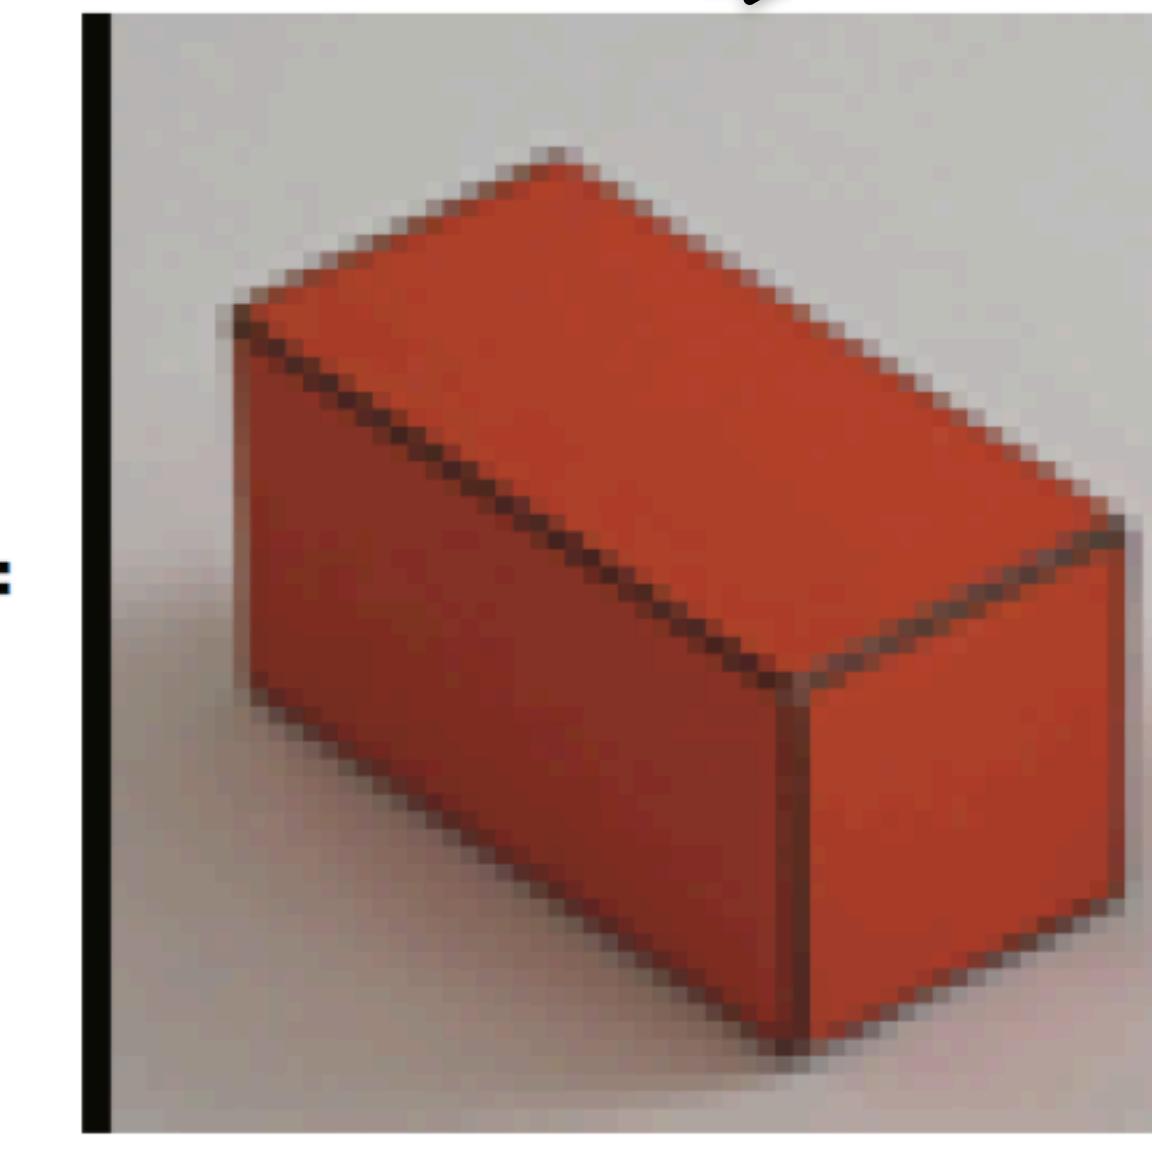
A shift



○

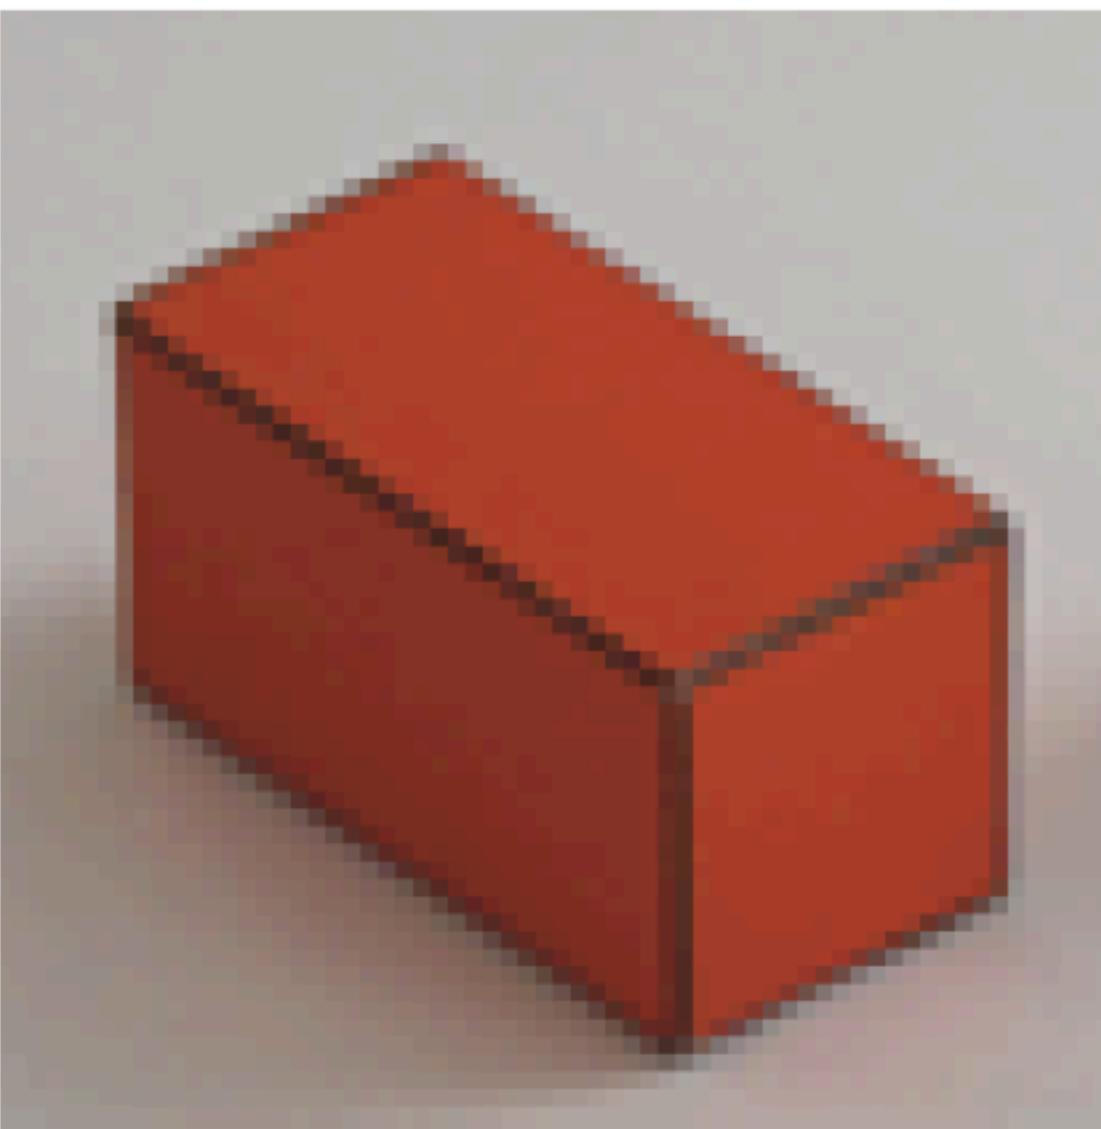
0	0	0	0	0
0	0	0	0	0
0	0	0	0	1
0	0	0	0	0
0	0	0	0	0

=

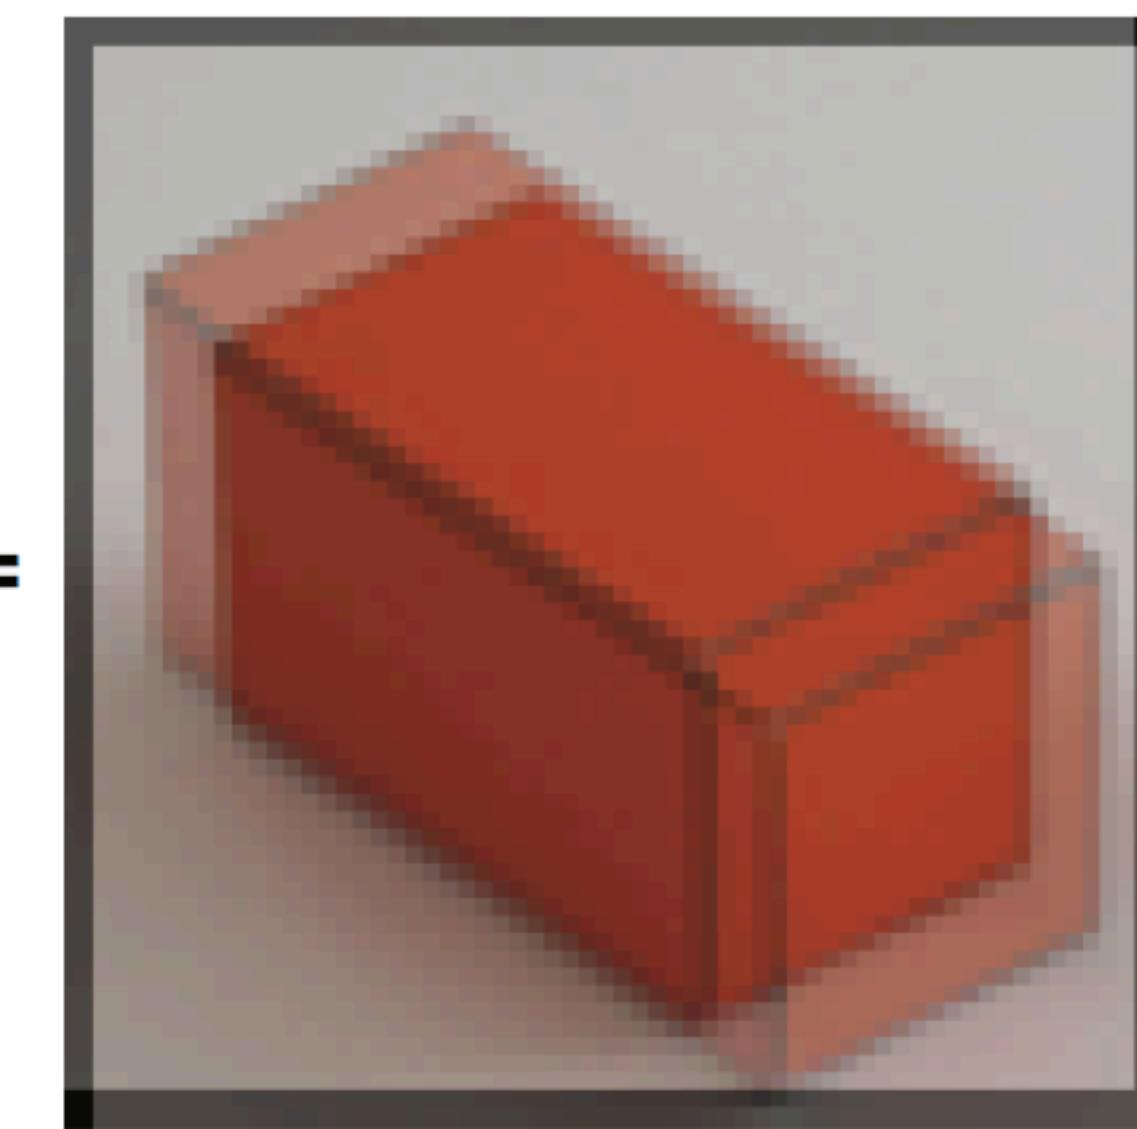


(using zero padding)

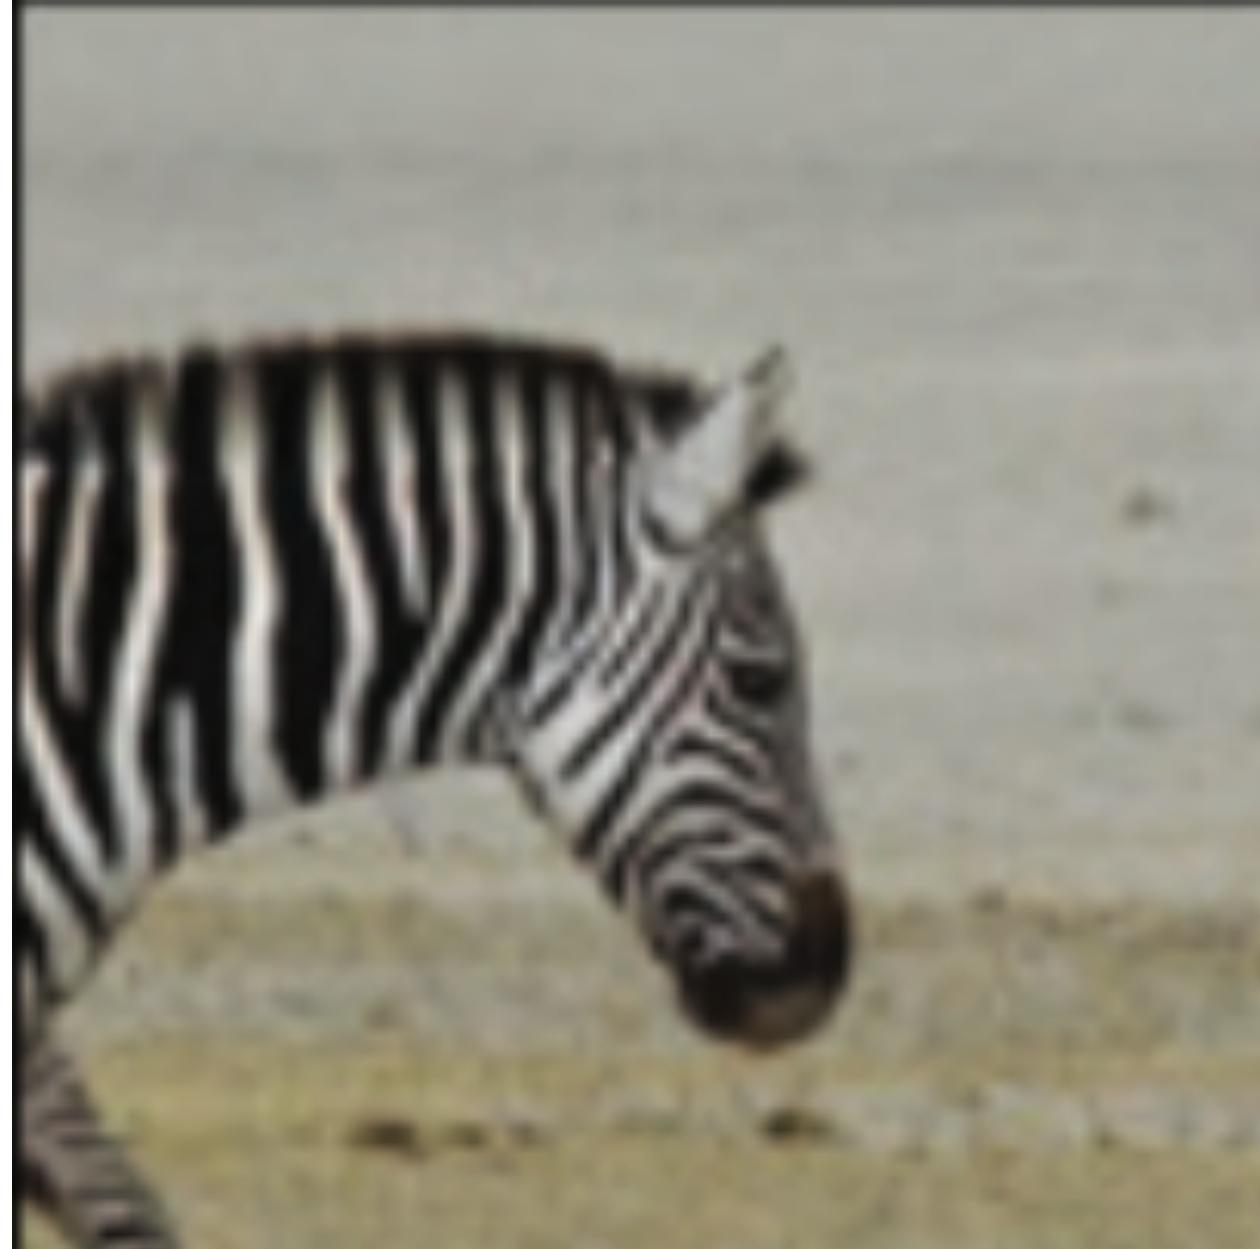
Examples



$$\odot \begin{array}{|c|c|c|c|c|} \hline .5 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & .5 \\ \hline \end{array} =$$

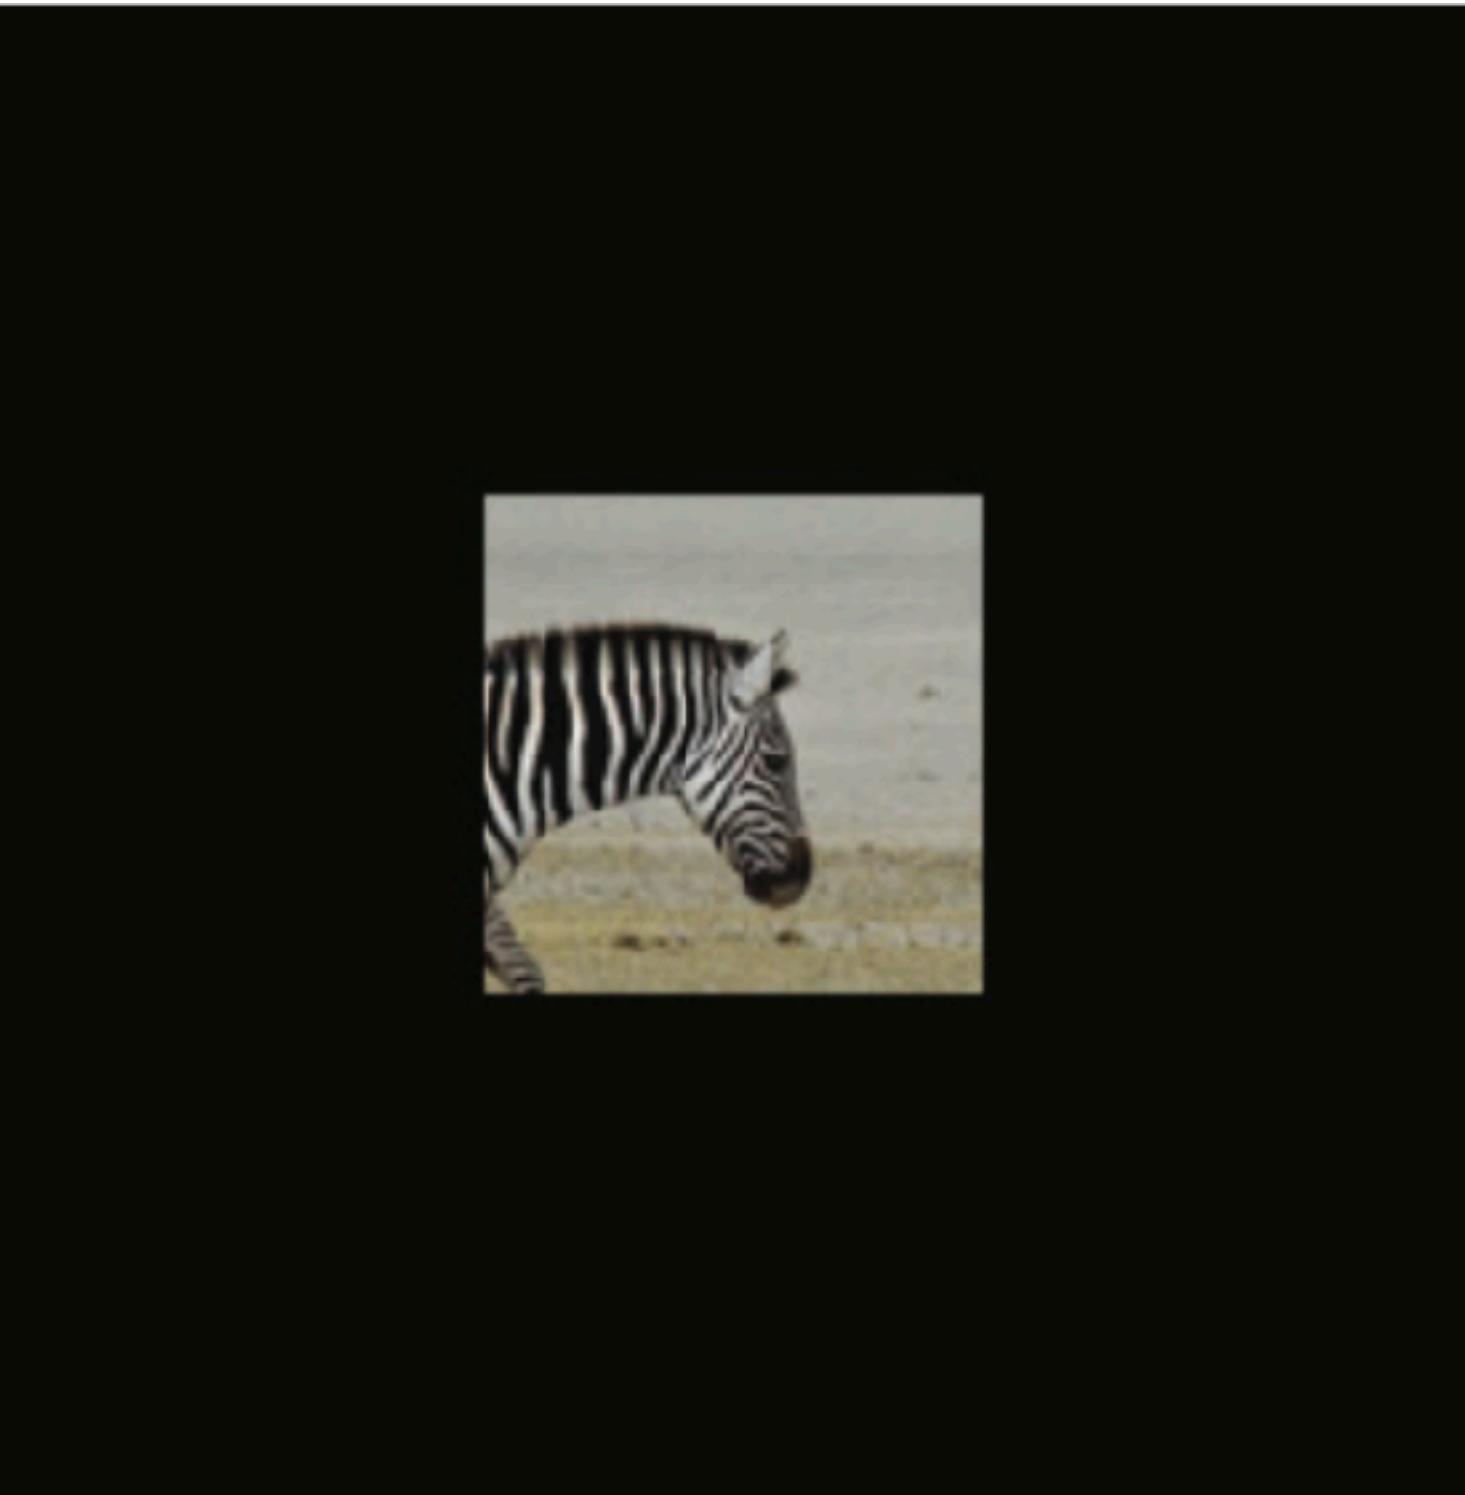


Handling boundaries

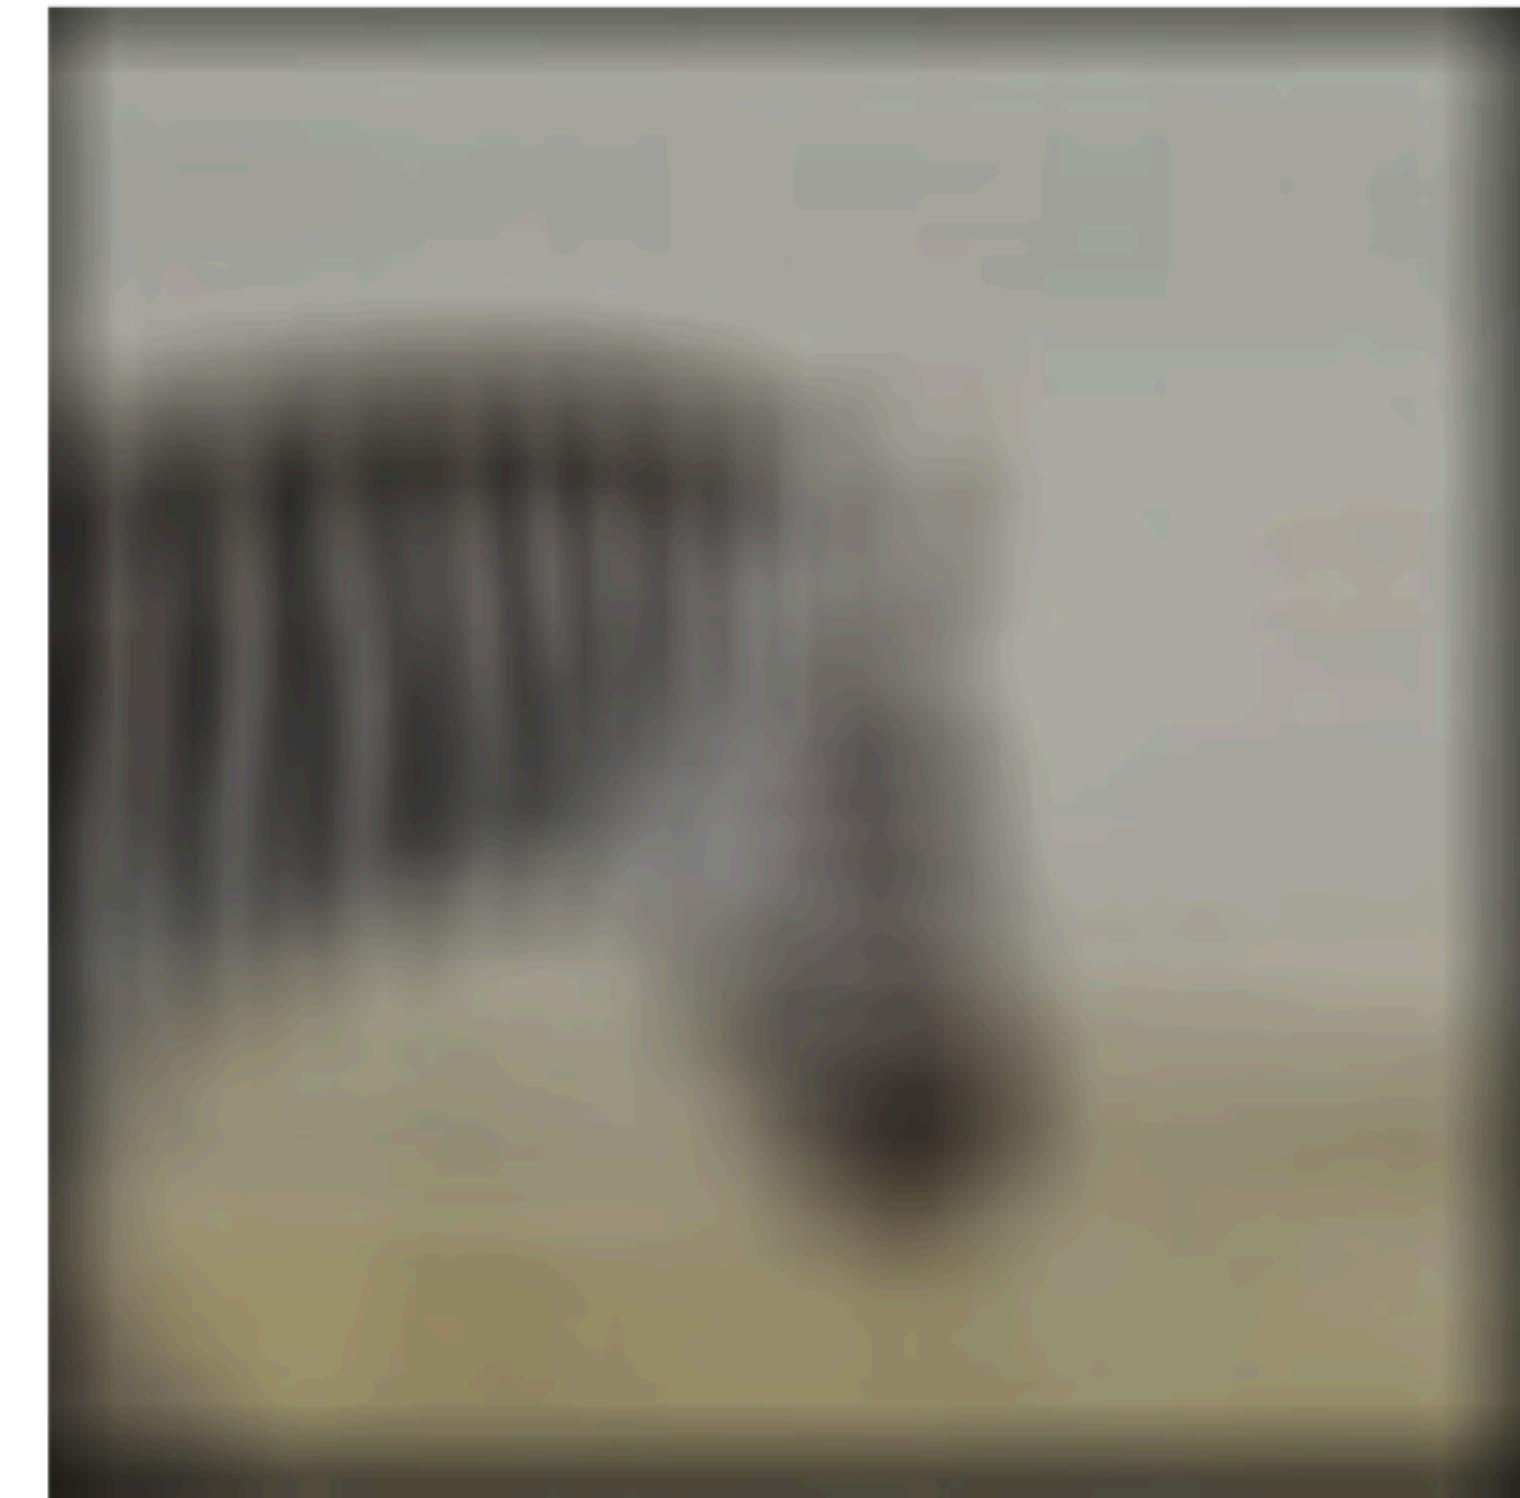


Handling boundaries

Zero padding



$$\bigcirc \quad \boxed{\textcolor{gray}{\square}} = \\ \uparrow \\ 11 \times 11 \text{ ones}$$



Handling boundaries

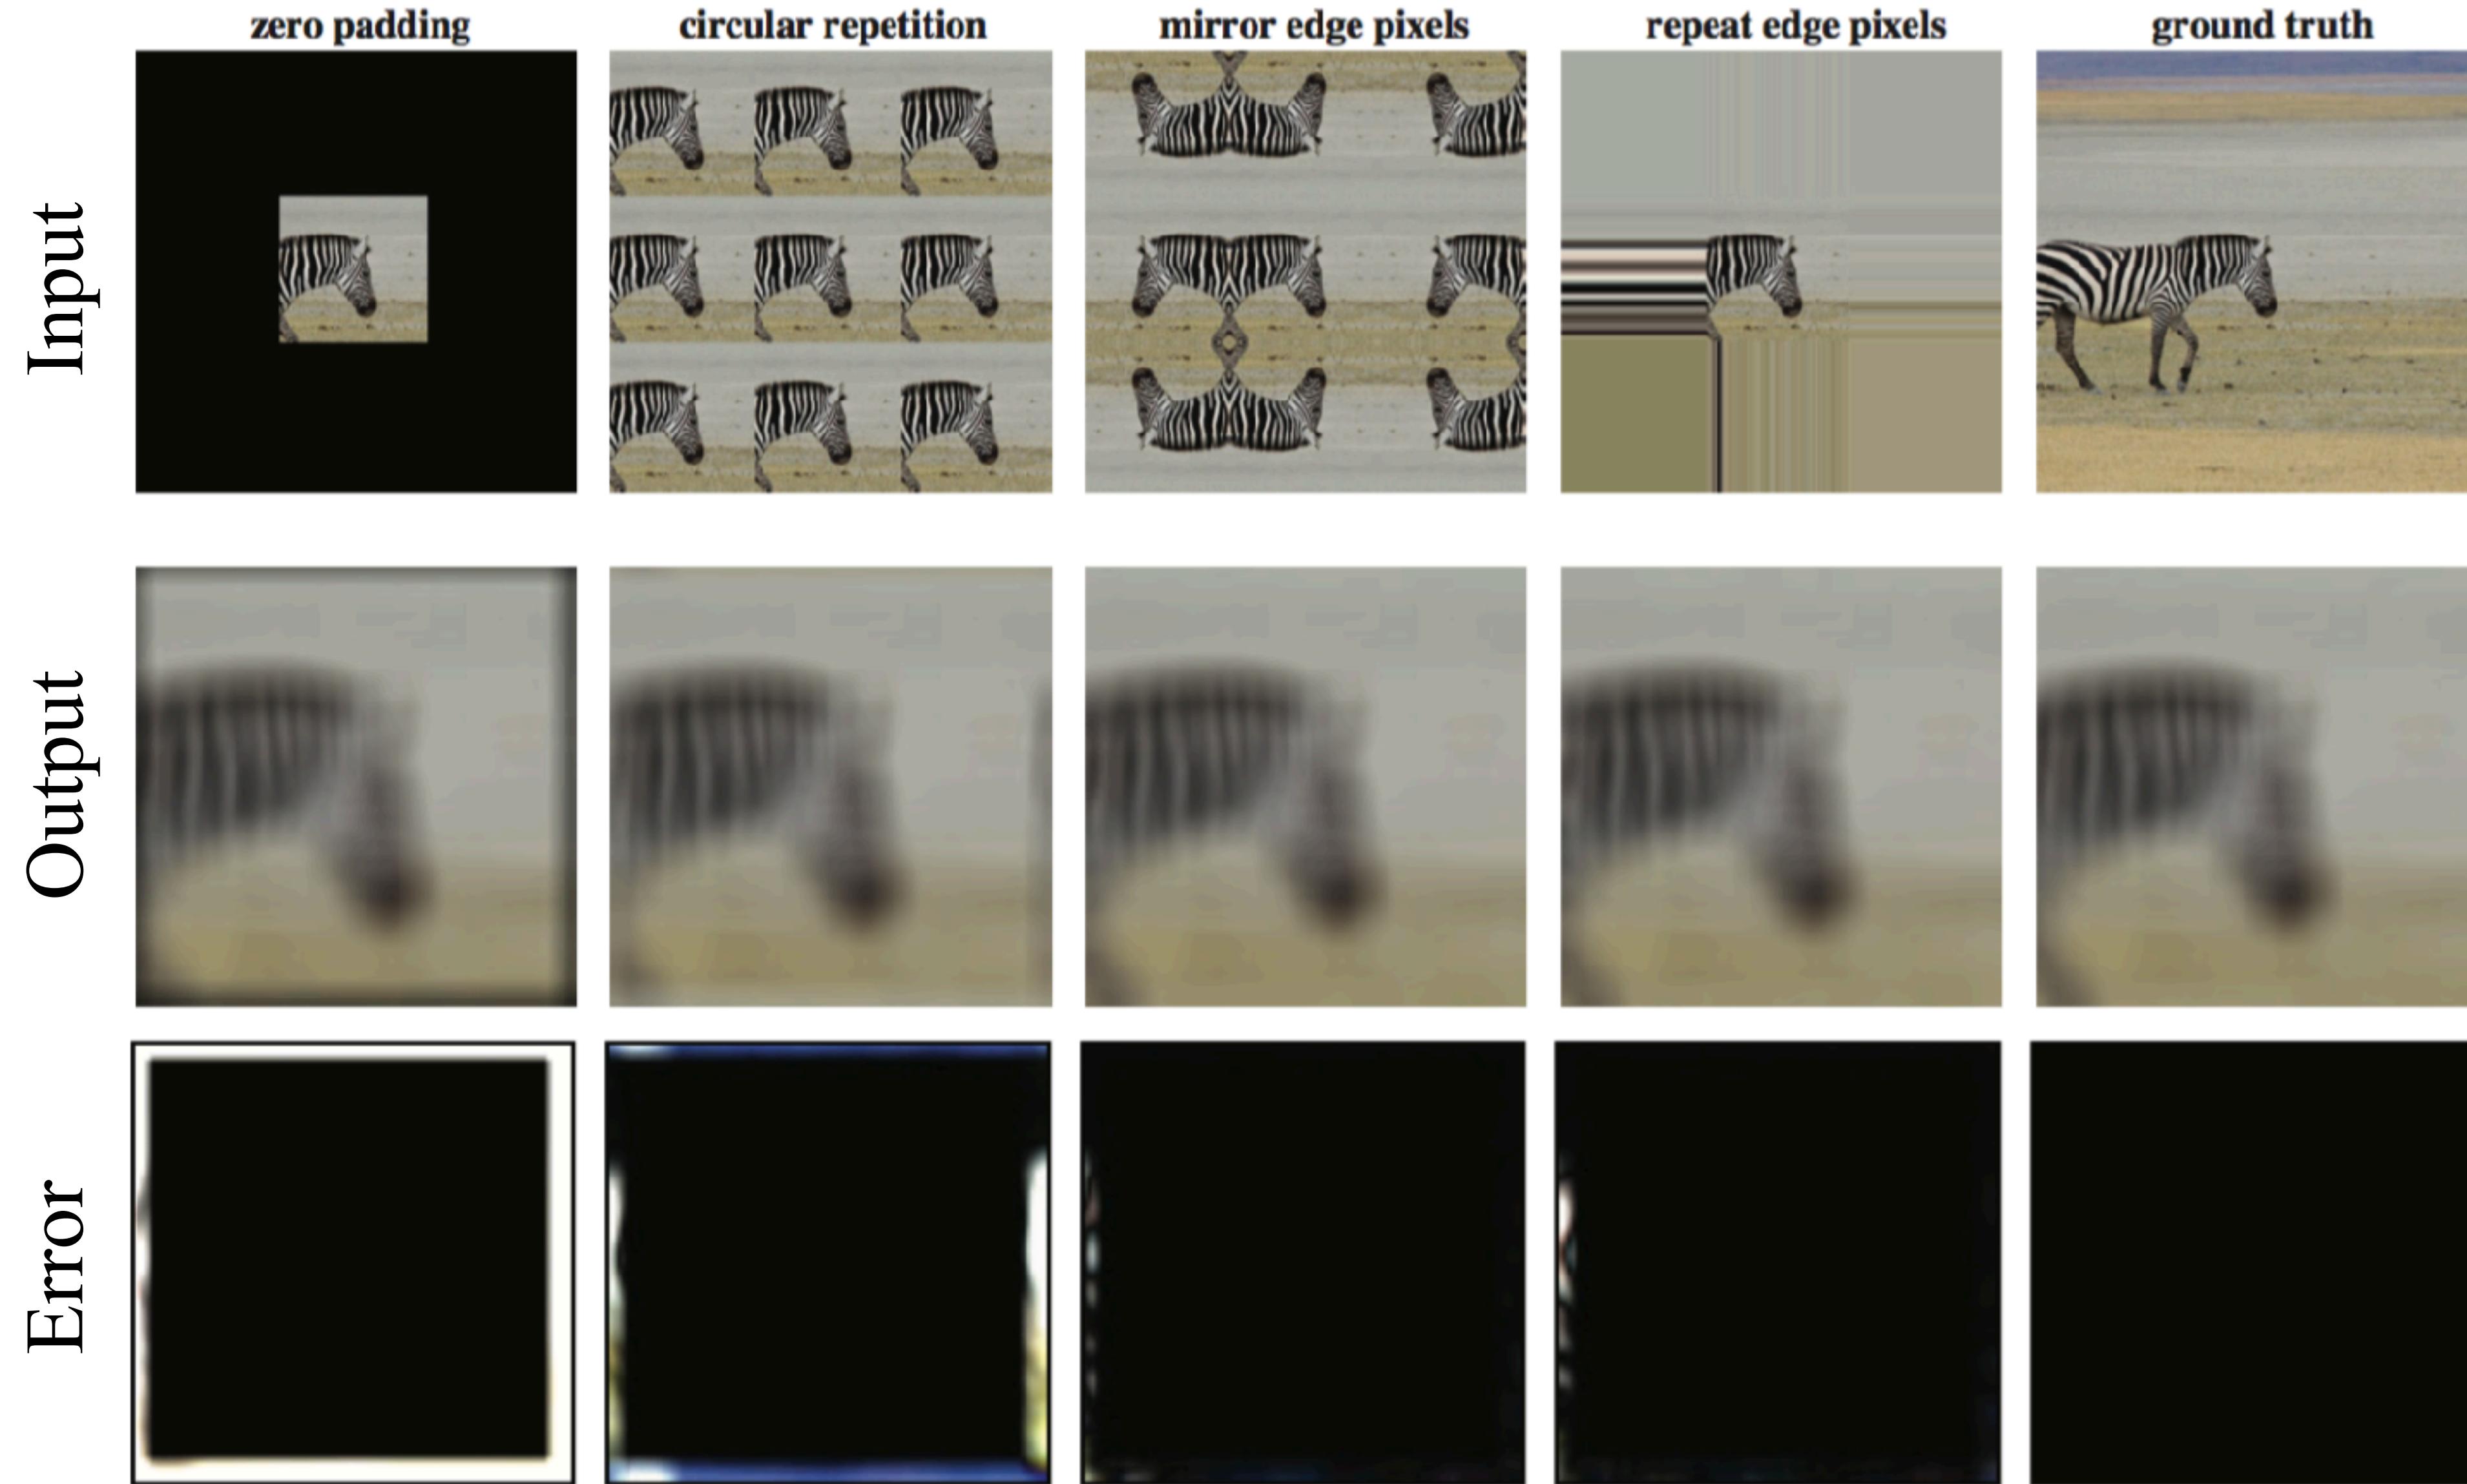
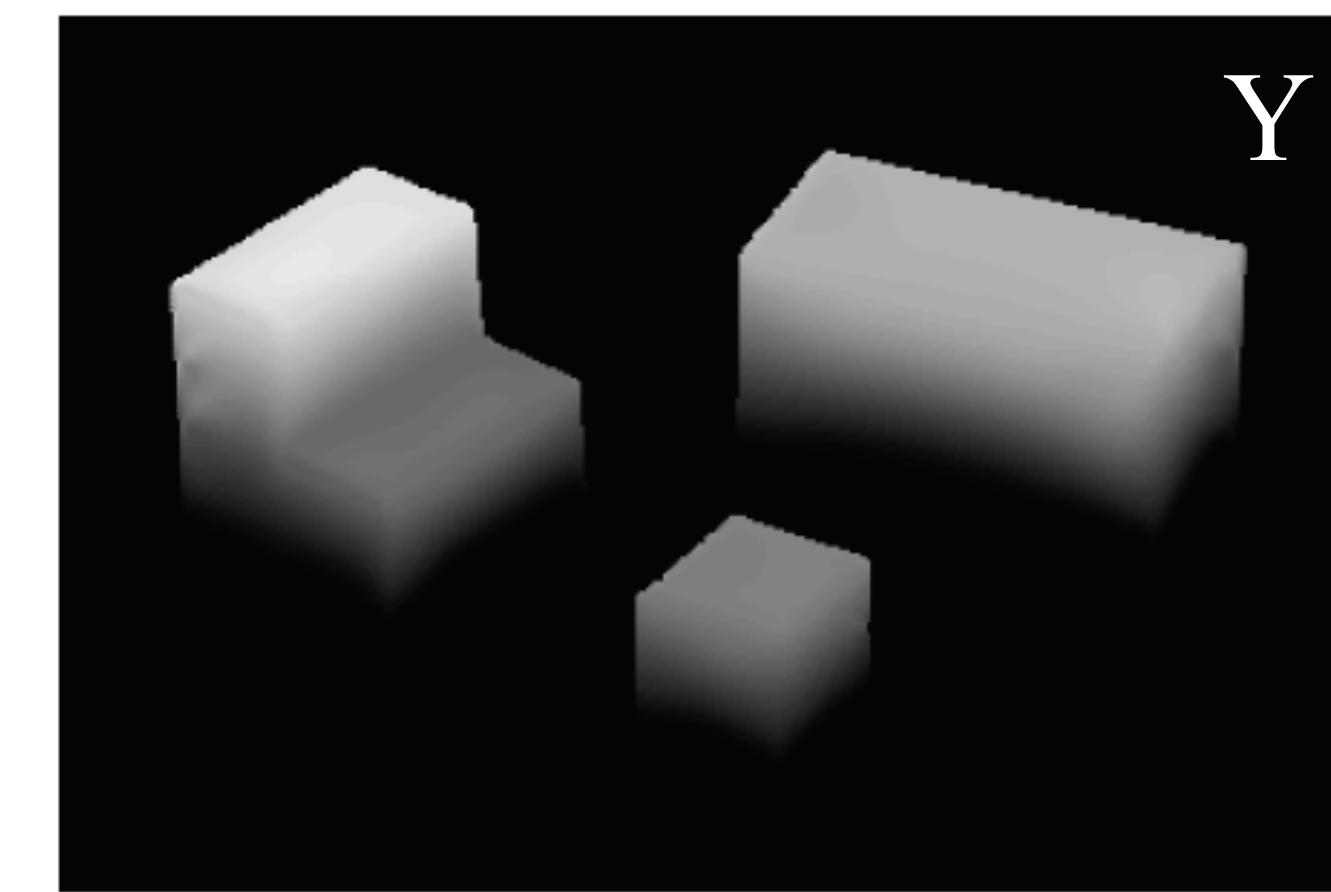
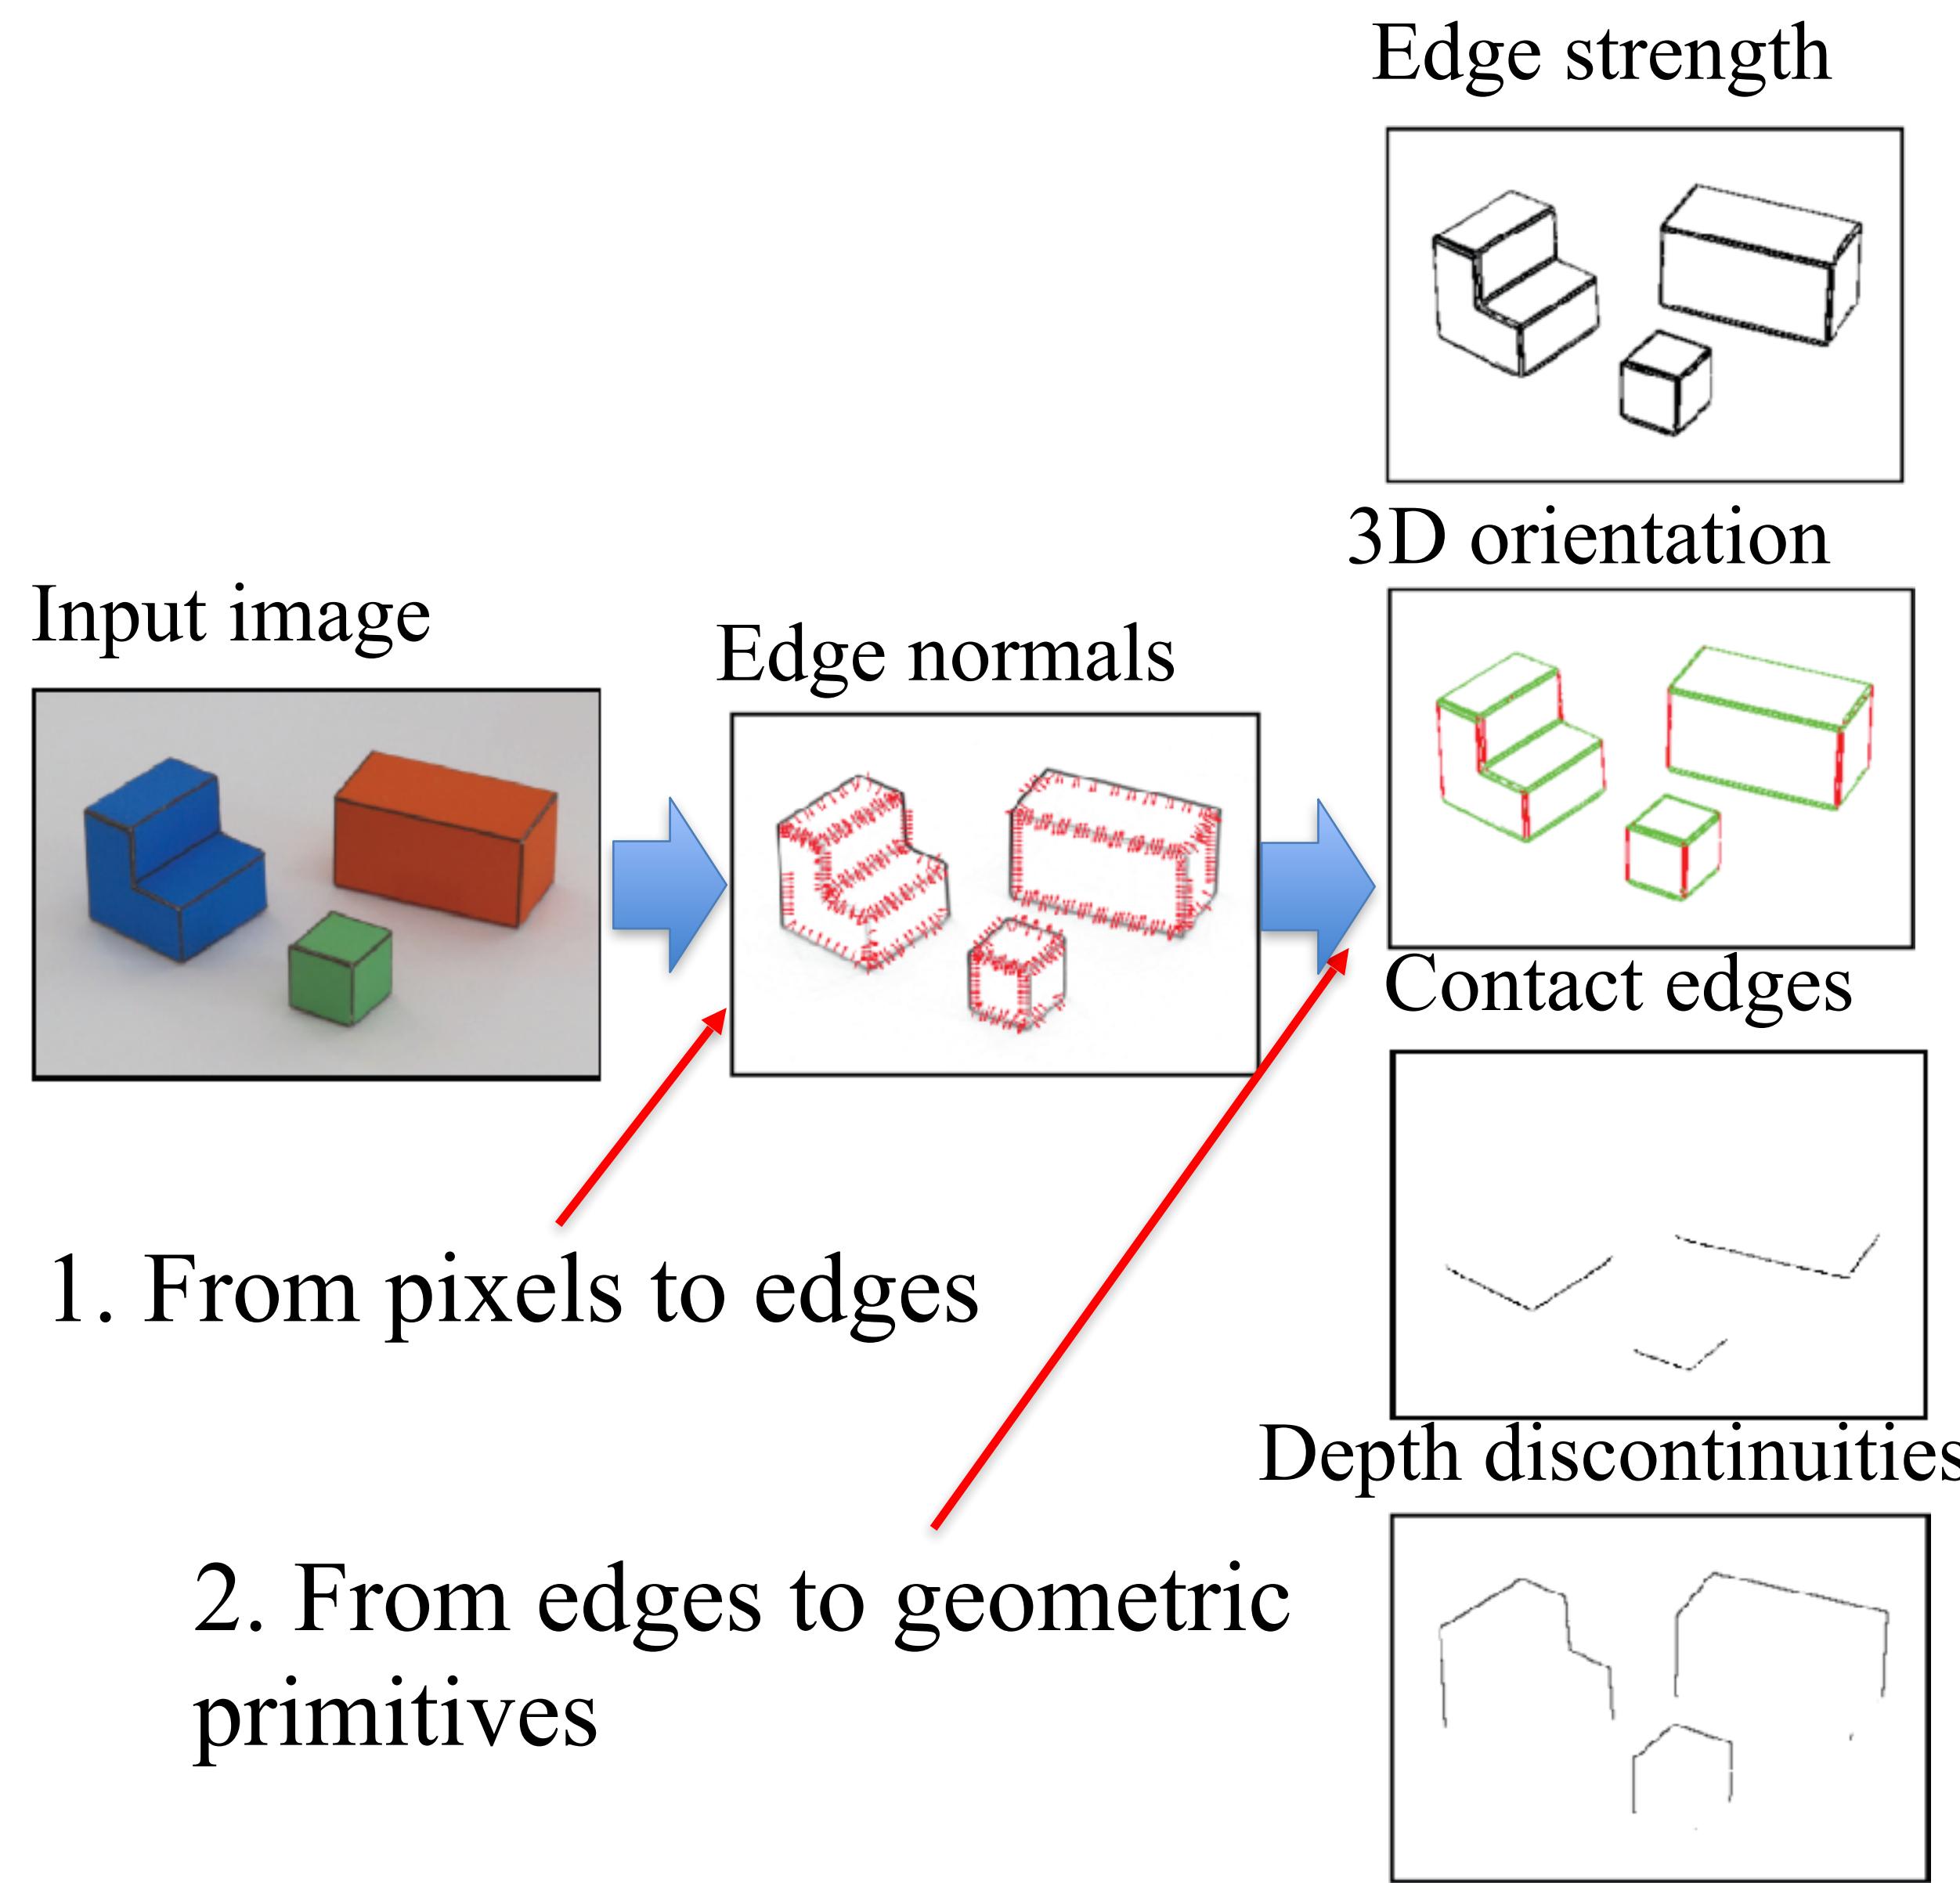


Image transformations

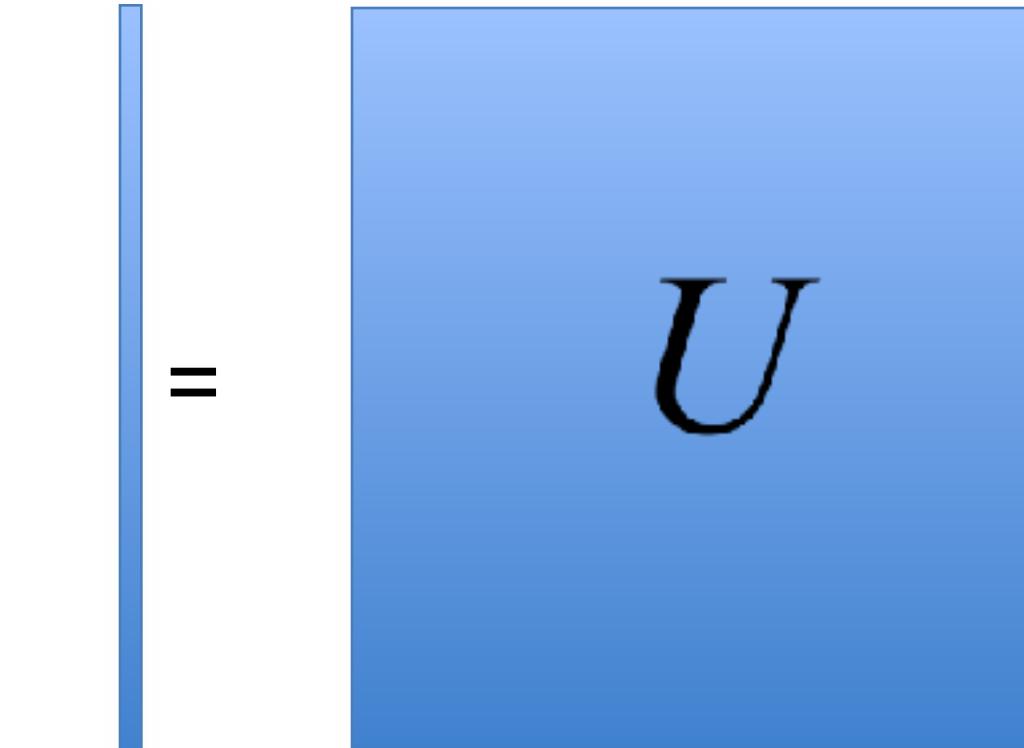


Linear image transformations

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

$$\vec{\tilde{F}} = \vec{U} \vec{f}$$

Transformed image → $\vec{\tilde{F}}$ = $\vec{U} \vec{f}$ ← Vectorized image

=  Fourier transform, or
Wavelet transform, or
Steerable pyramid transform

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:

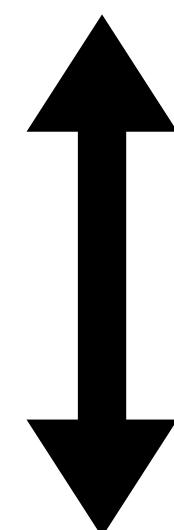
$$F = \begin{matrix} & \xrightarrow{n} \\ \downarrow u & \begin{matrix} ? & ? & ? & ? & ? & ? & ? & ? & \dots & ? \end{matrix} \\ \text{NxN array} \end{matrix} \quad f$$

Lets visualize the transform coefficients

Visualizing the Fourier transform

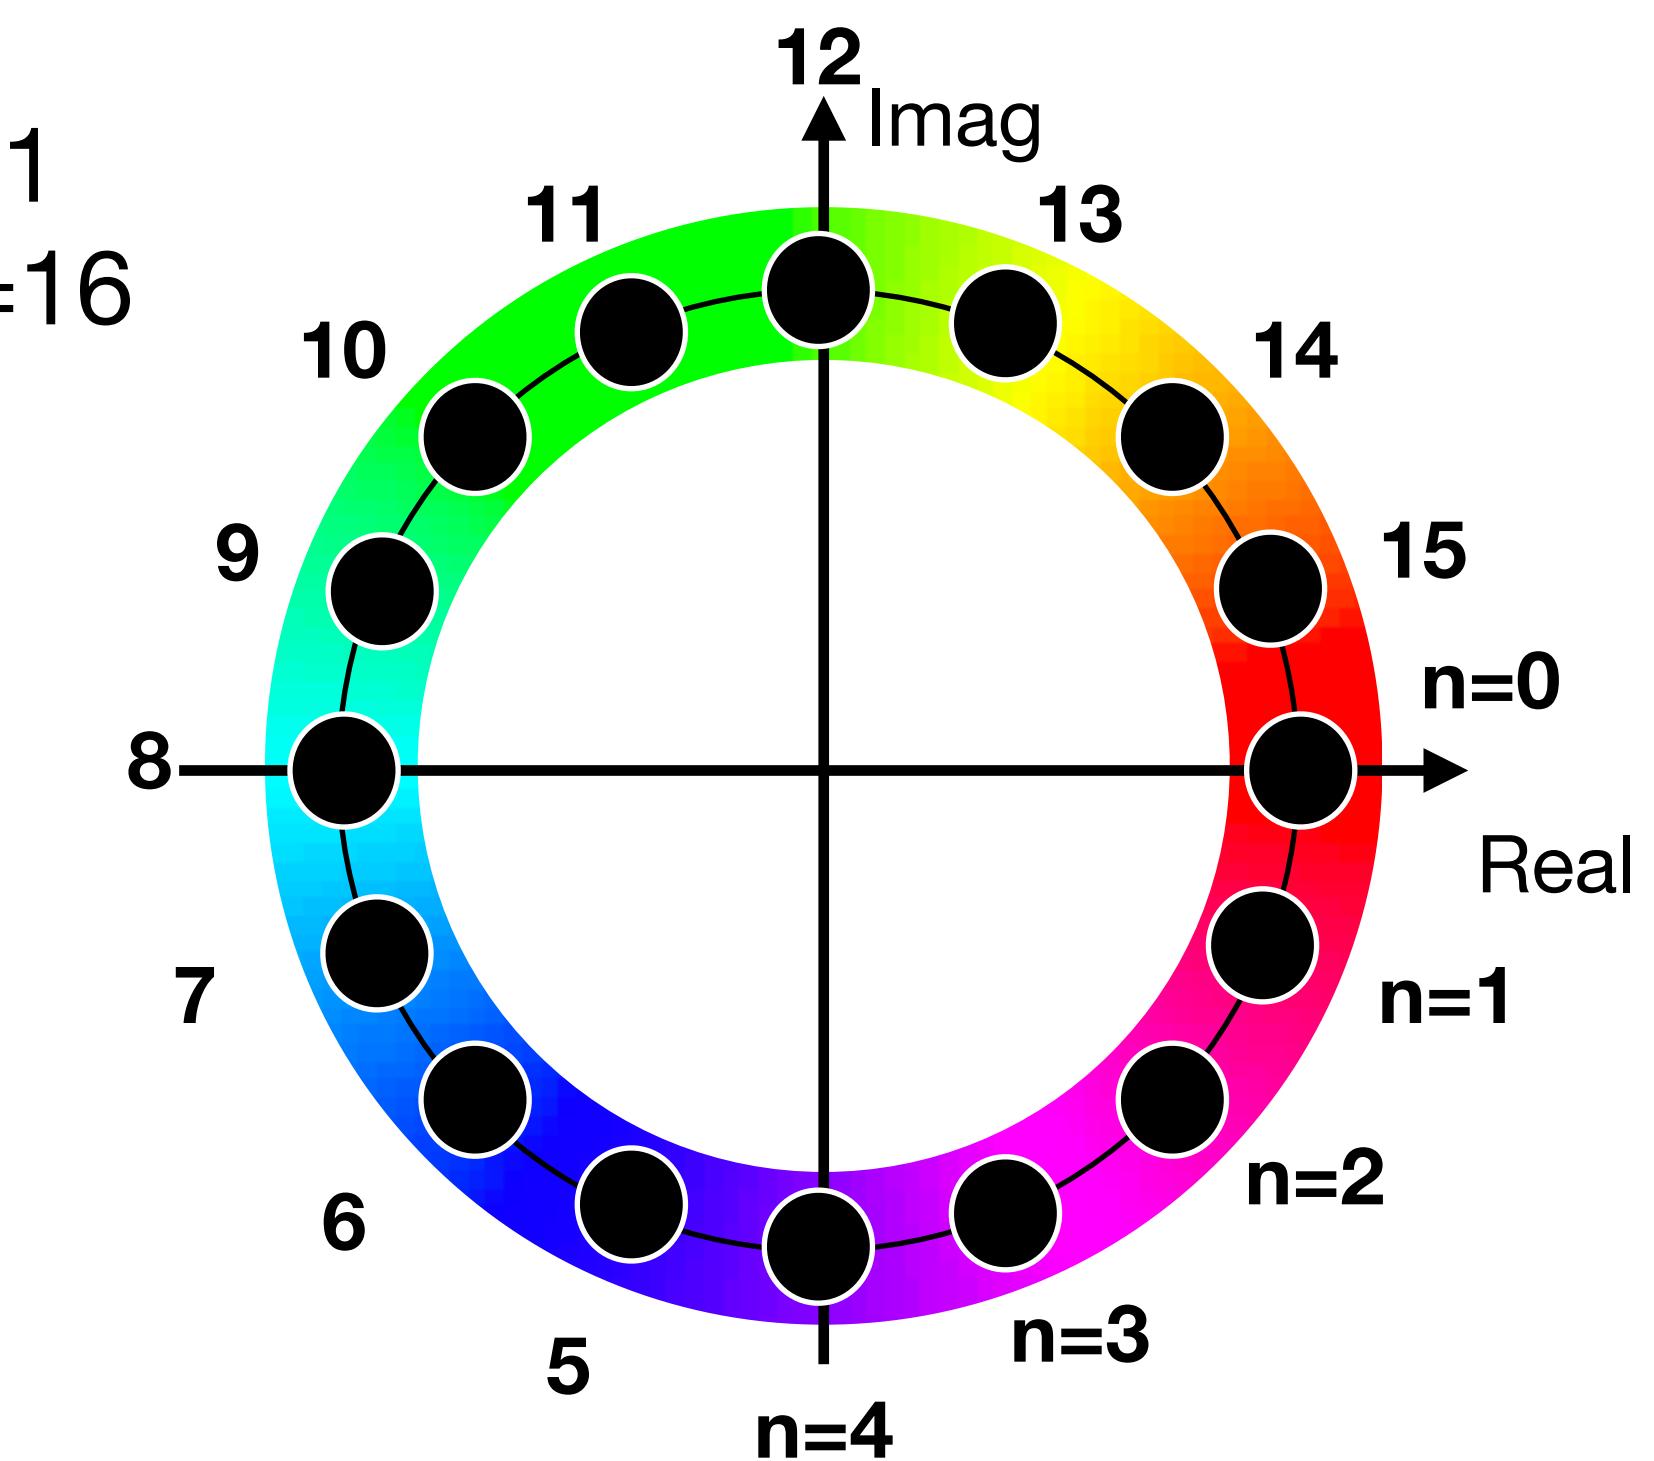
$$F[u] = \sum_{n=0}^{N-1} f[n] \exp\left(-2\pi j \frac{un}{N}\right)$$

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



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

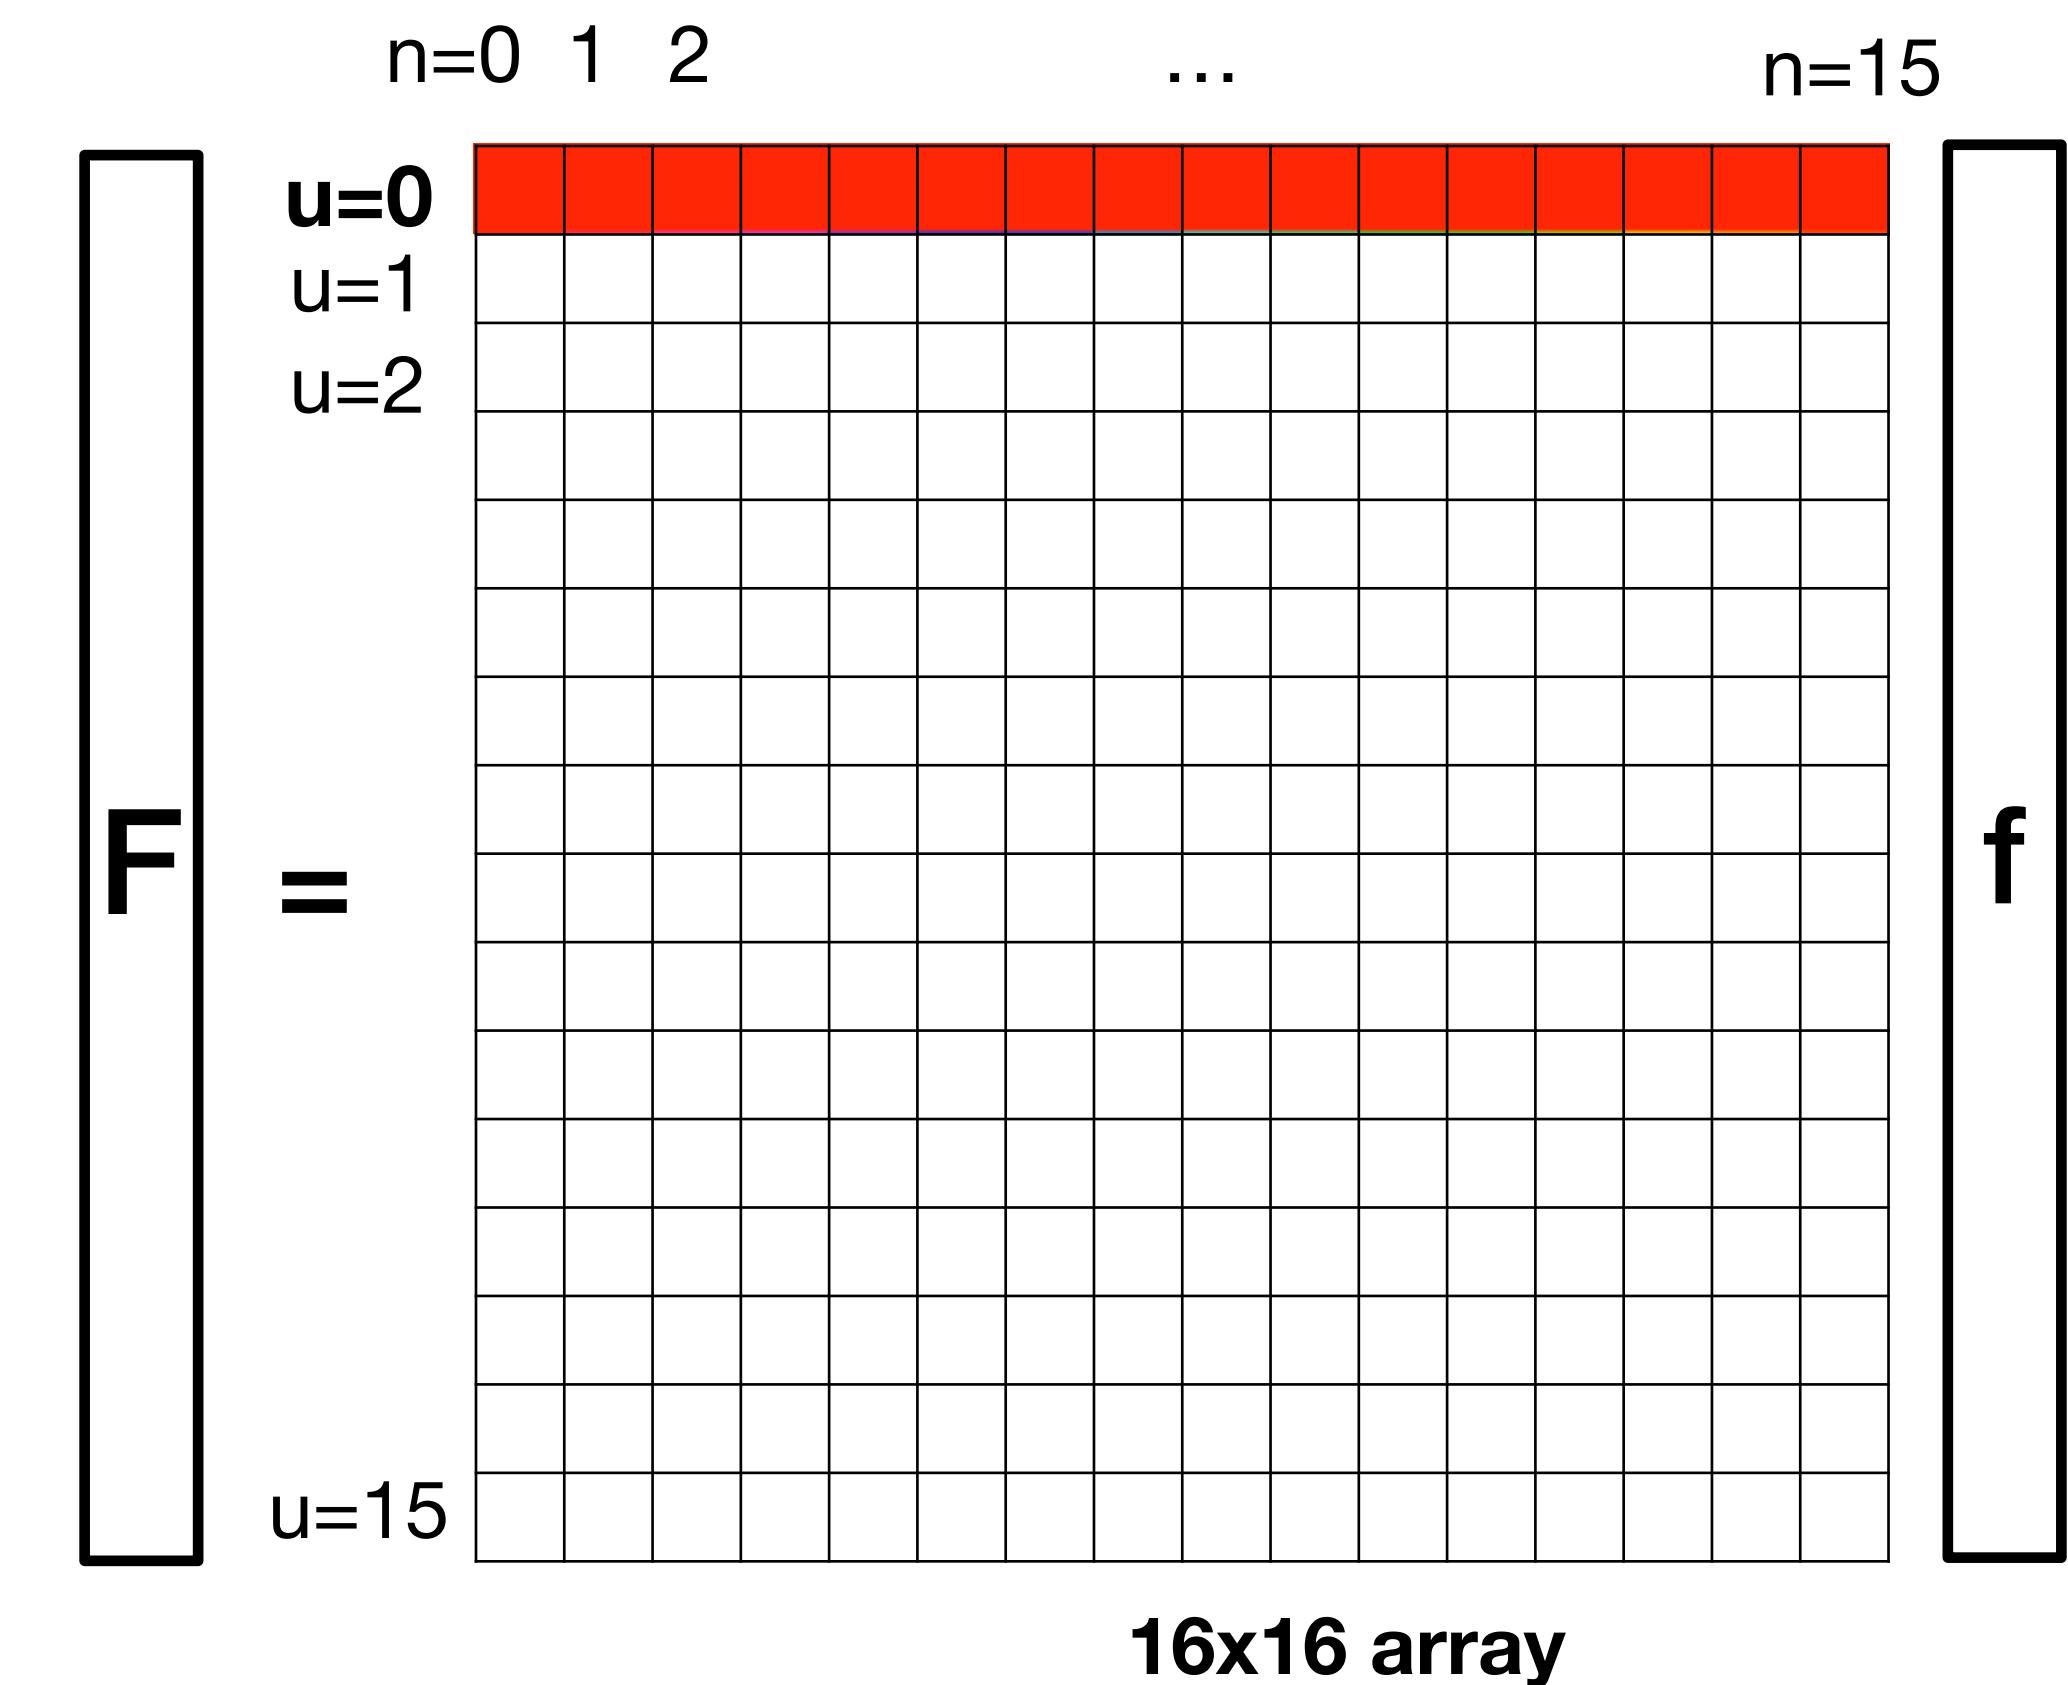
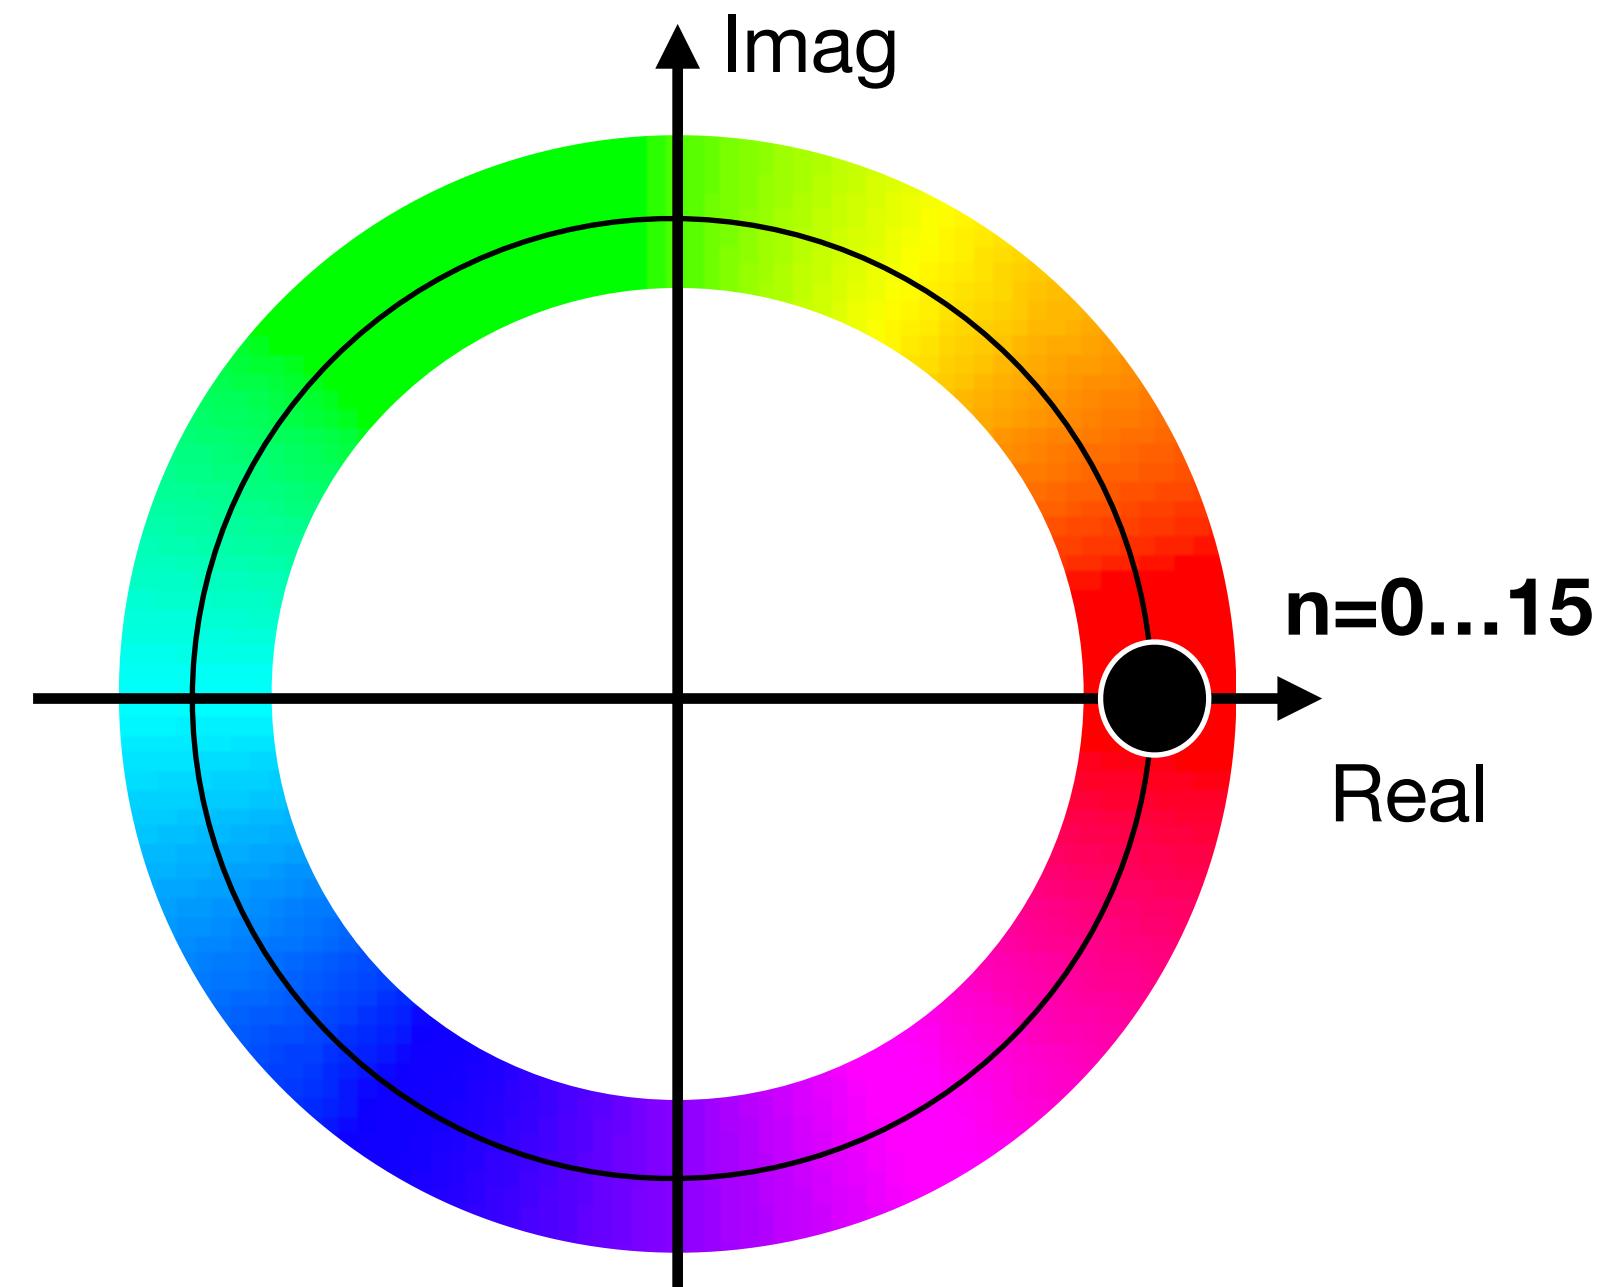
For:
 $u=1$
 $N=16$



Visualizing the transform coefficients

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

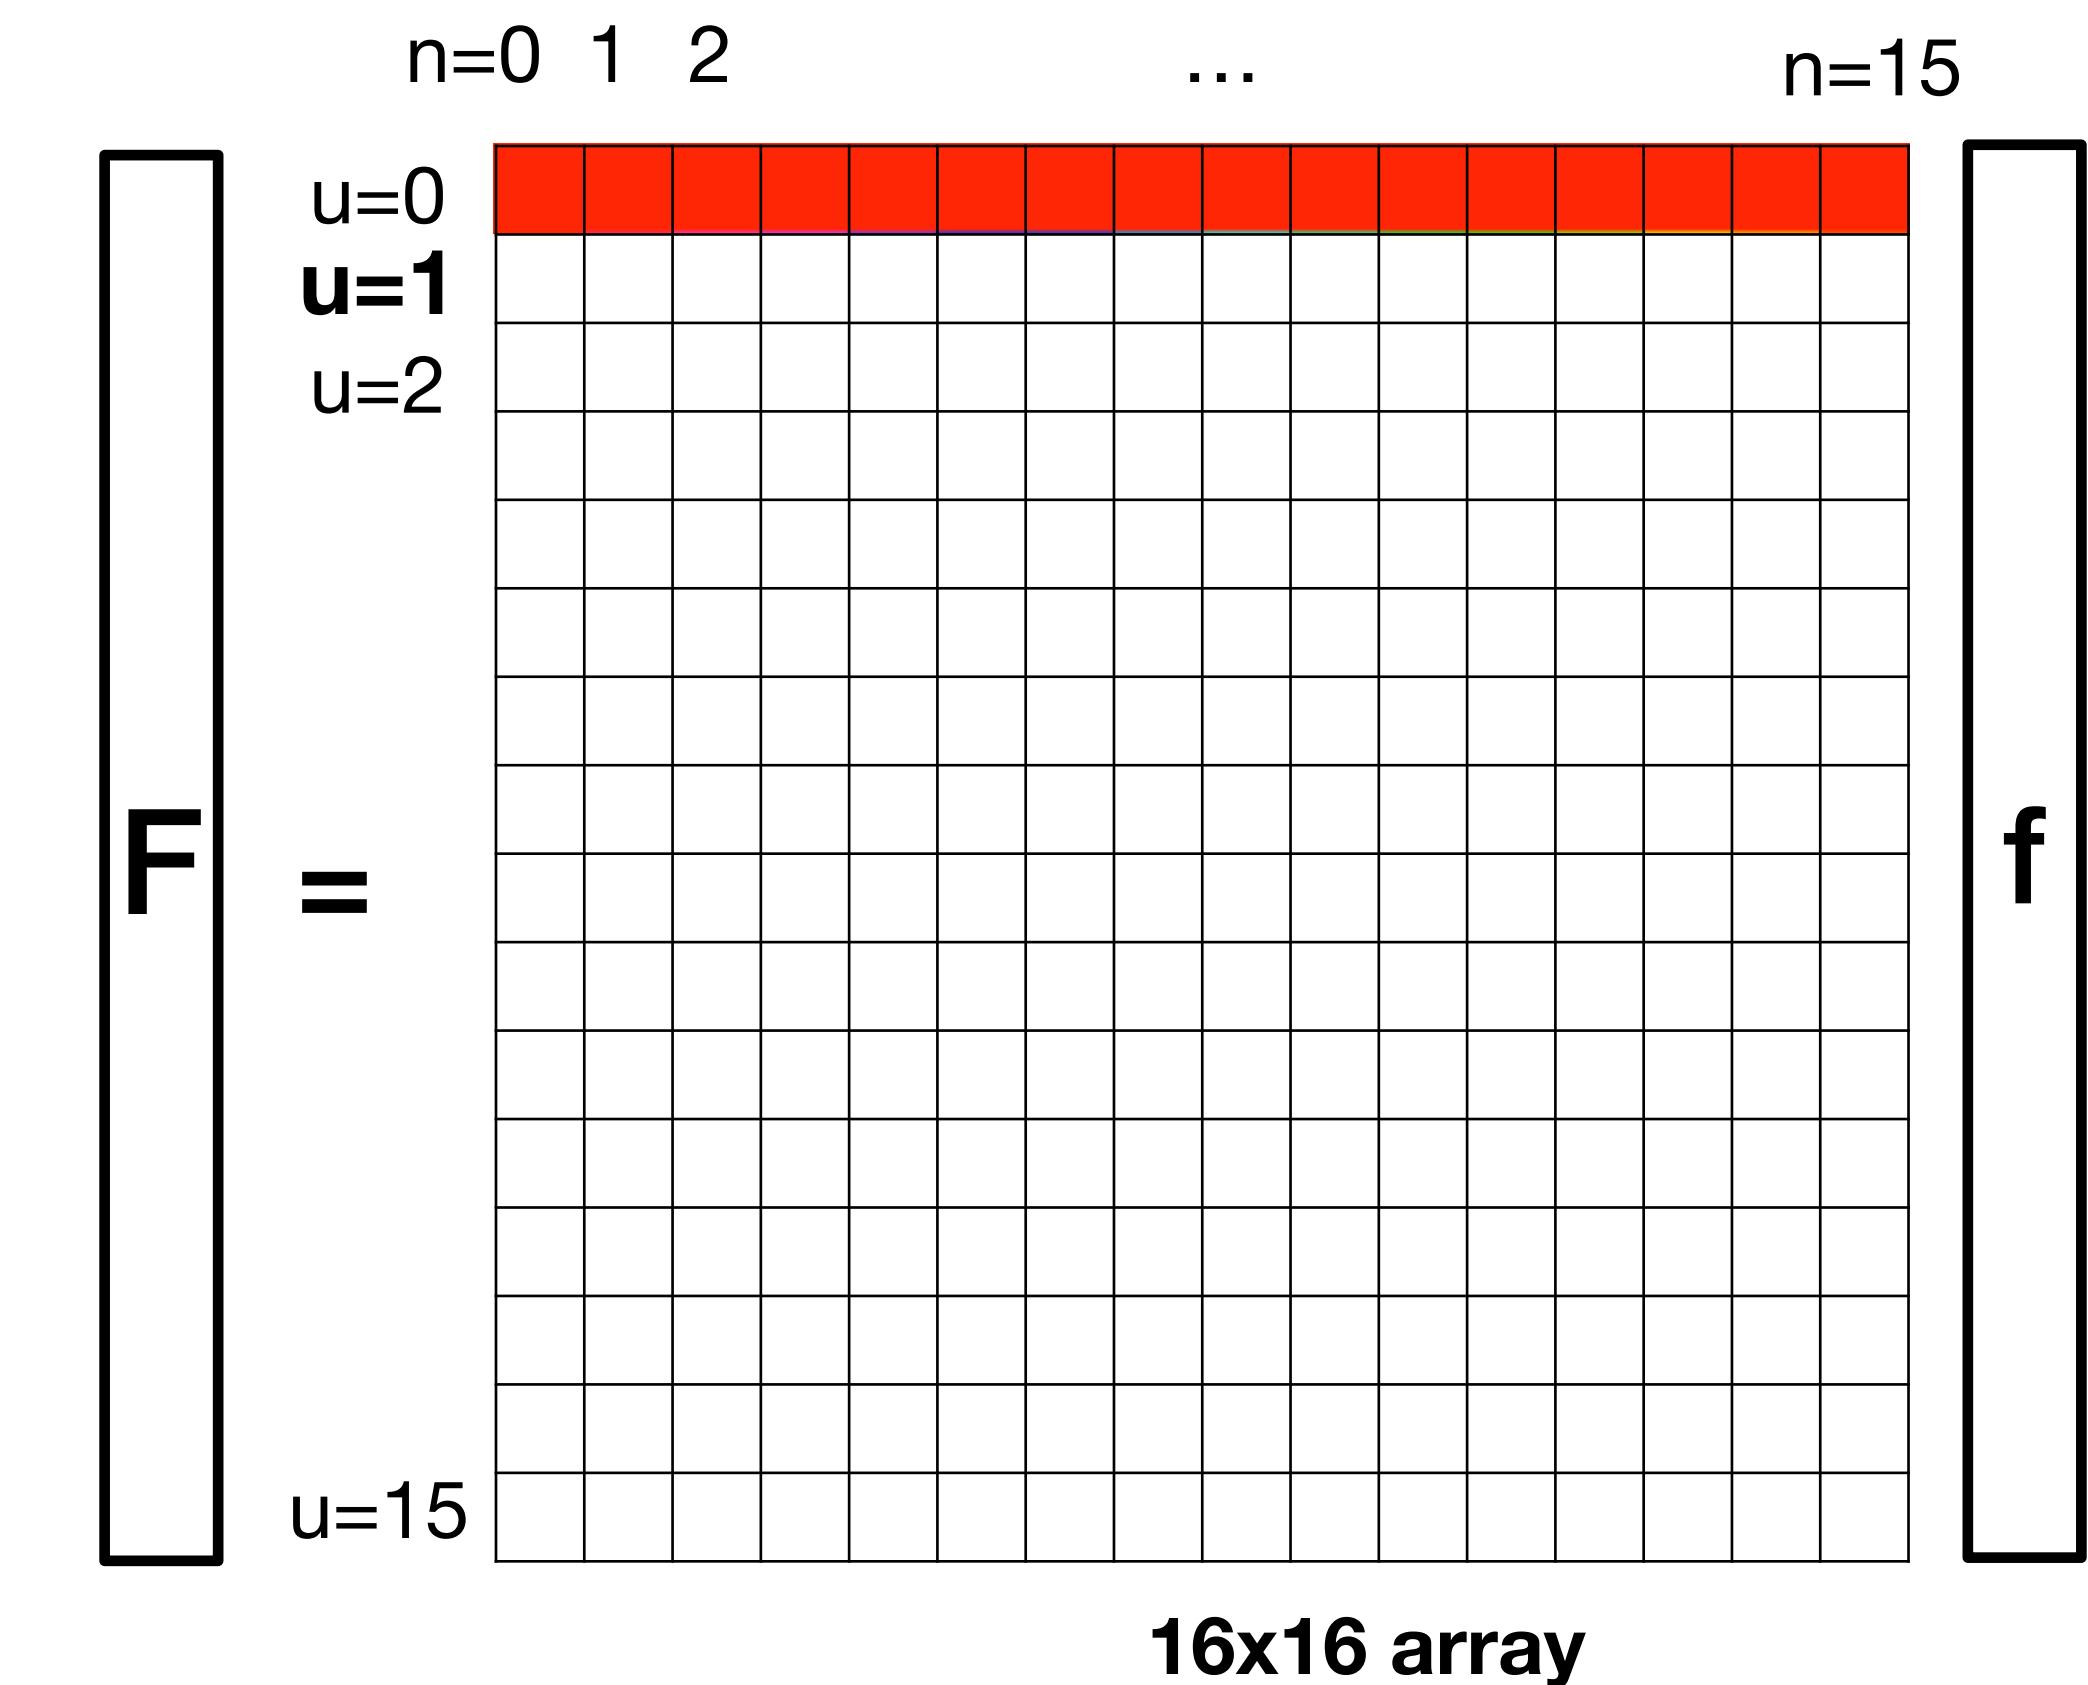
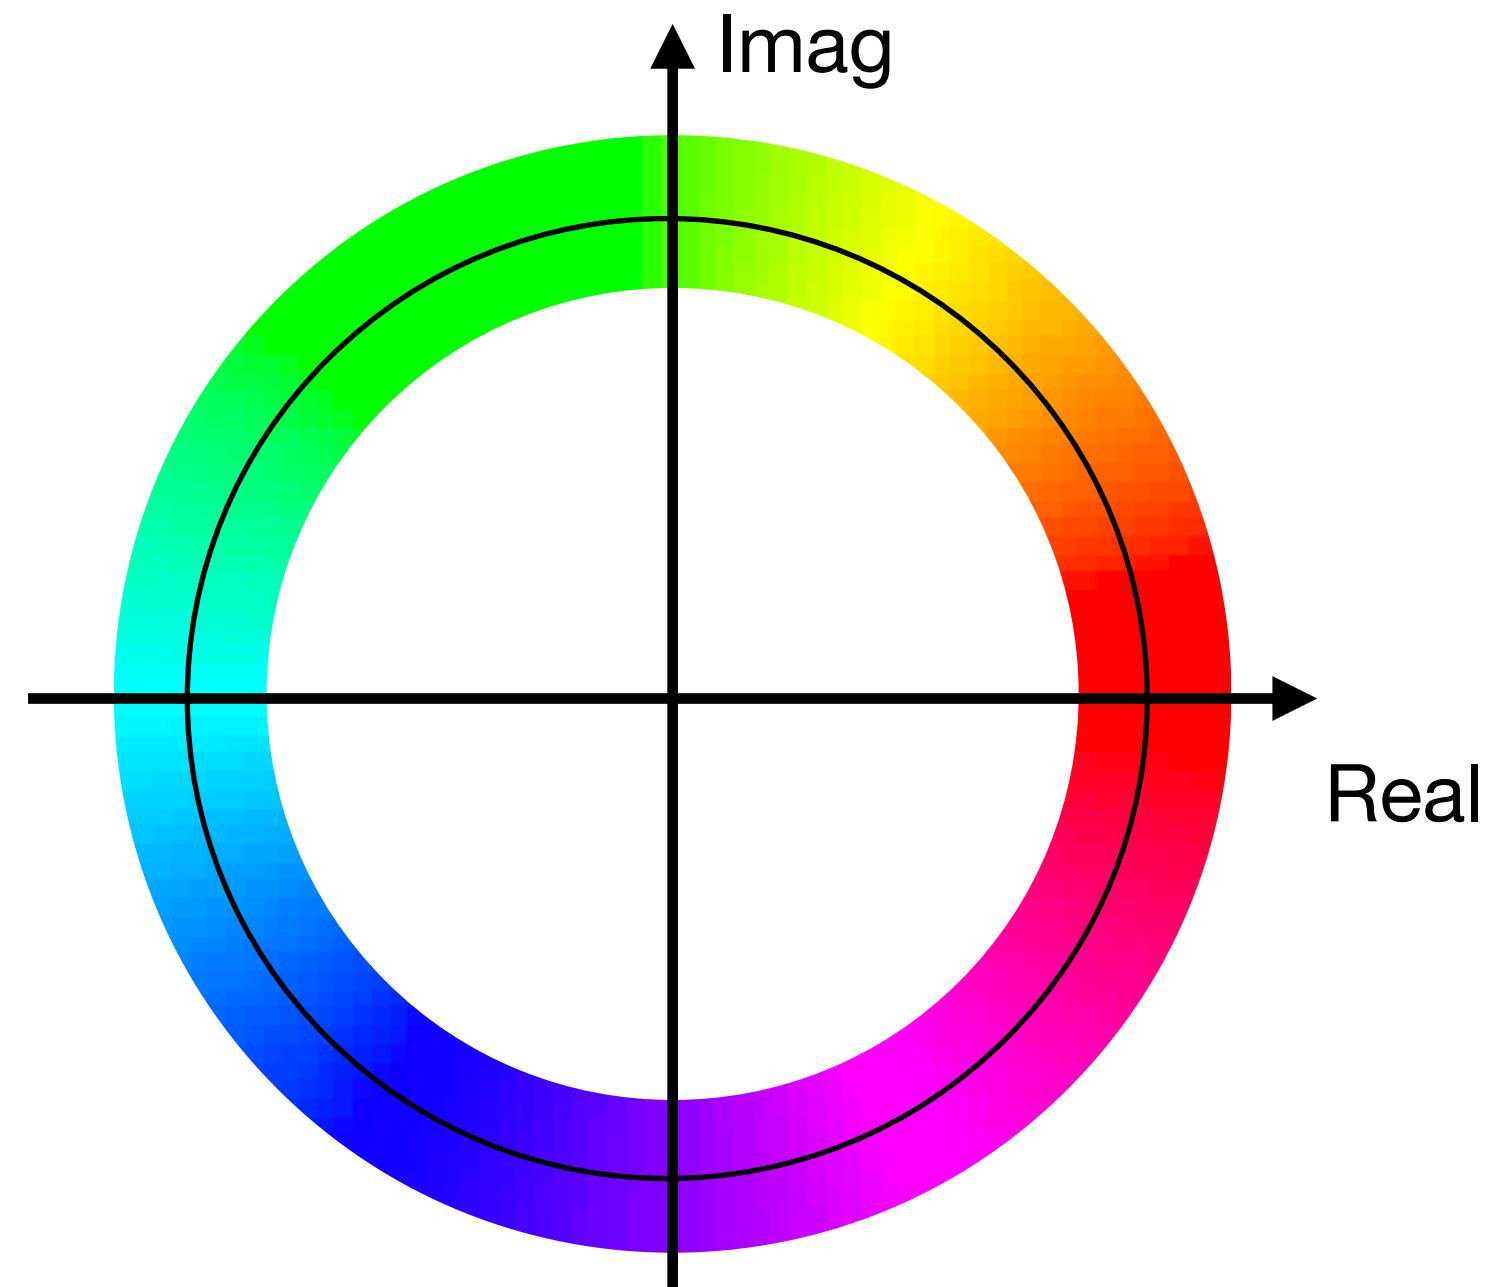
For N=16



Visualizing the transform coefficients

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

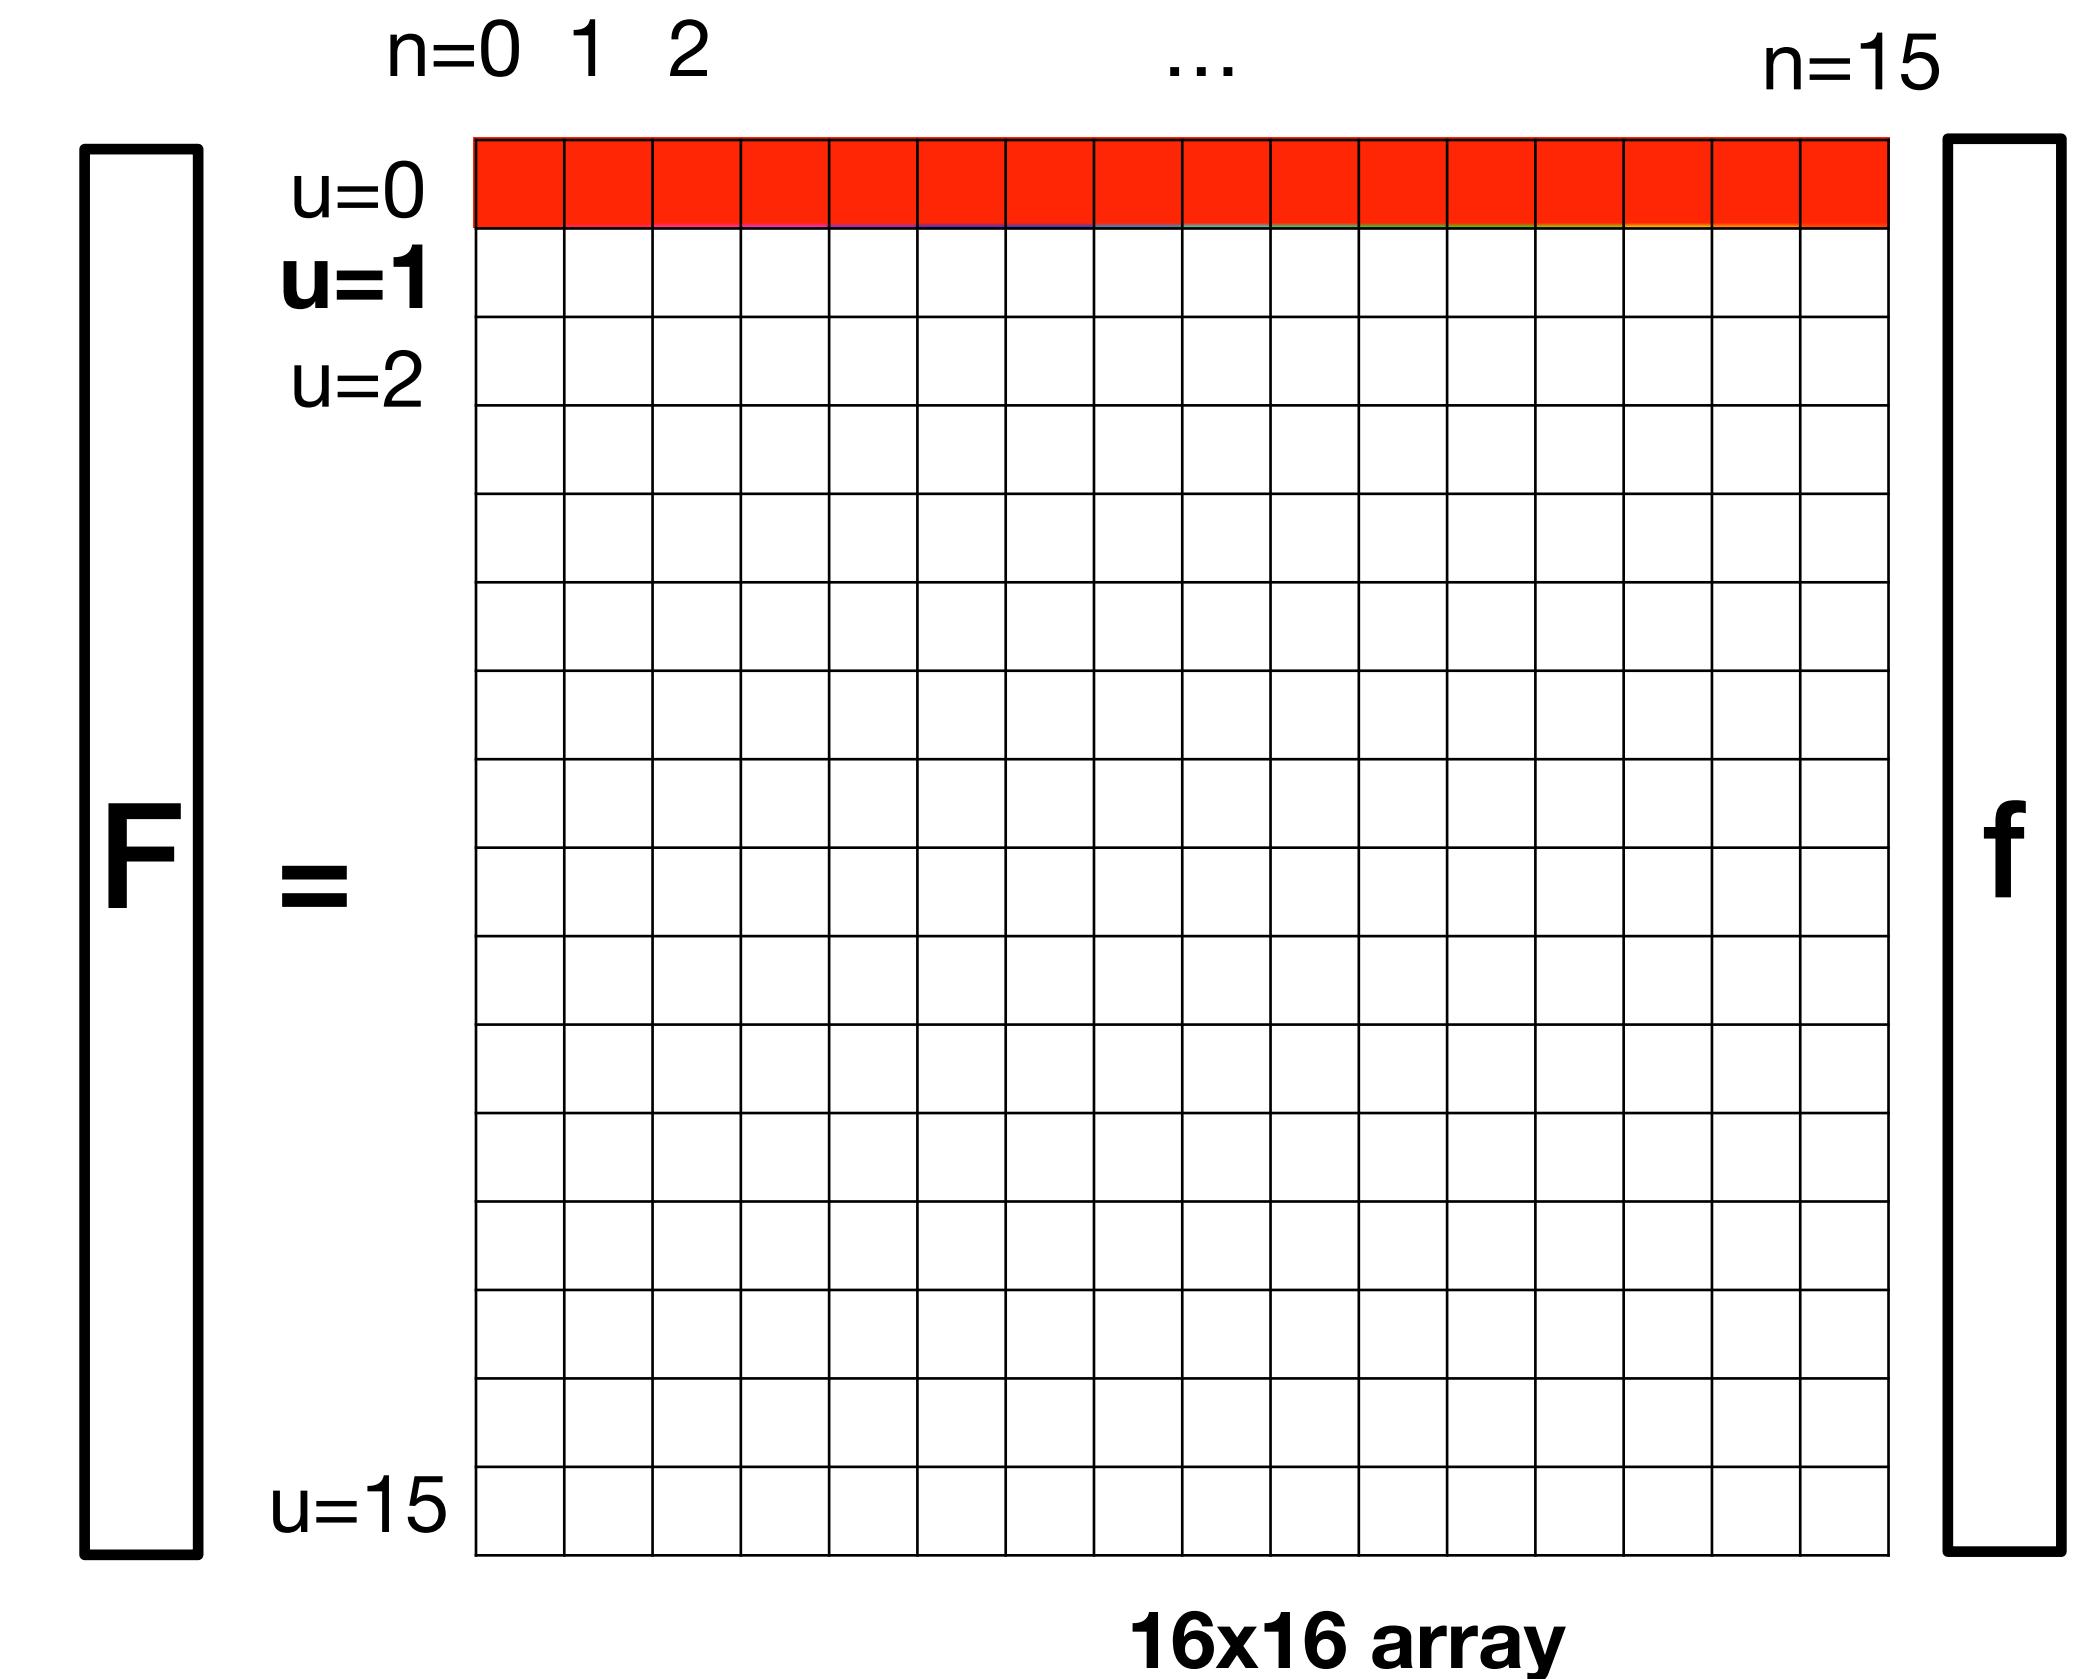
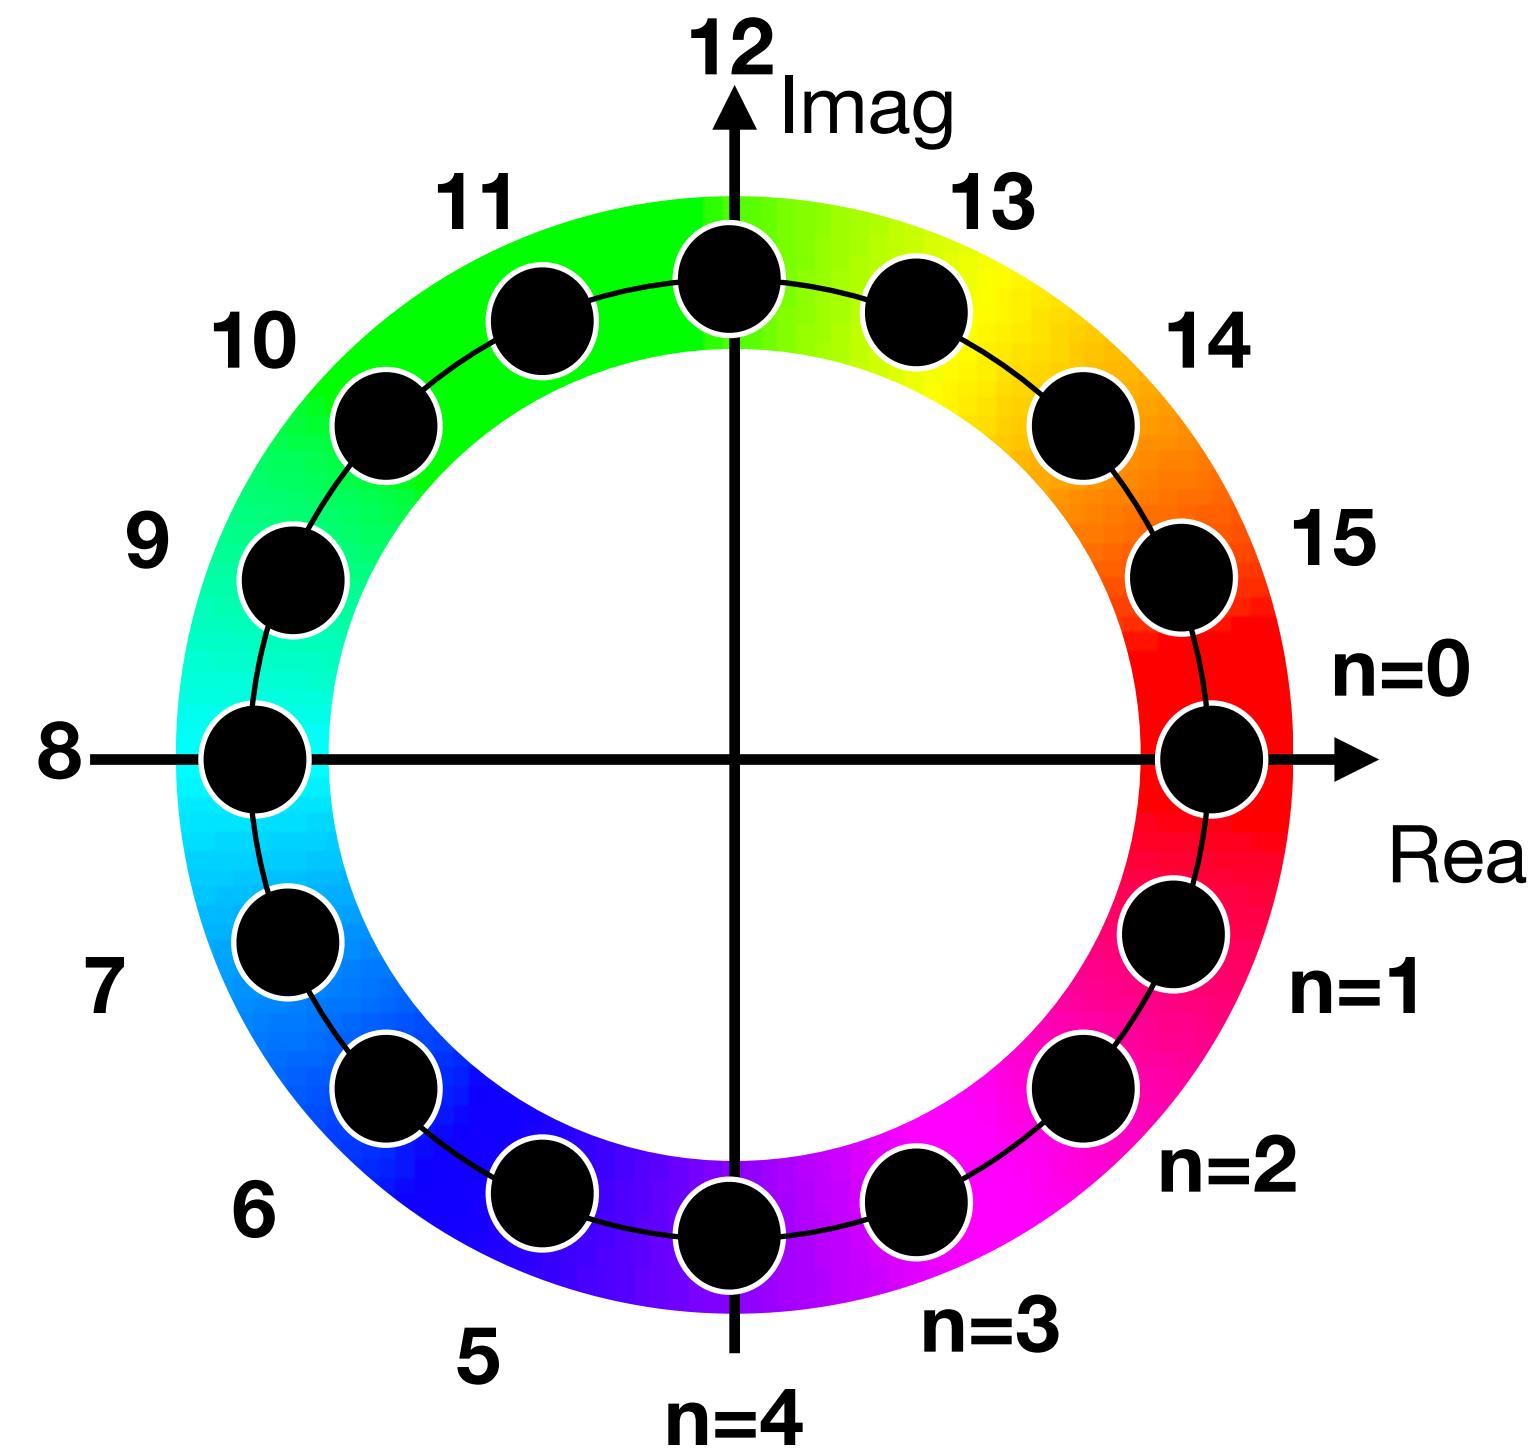
For N=16



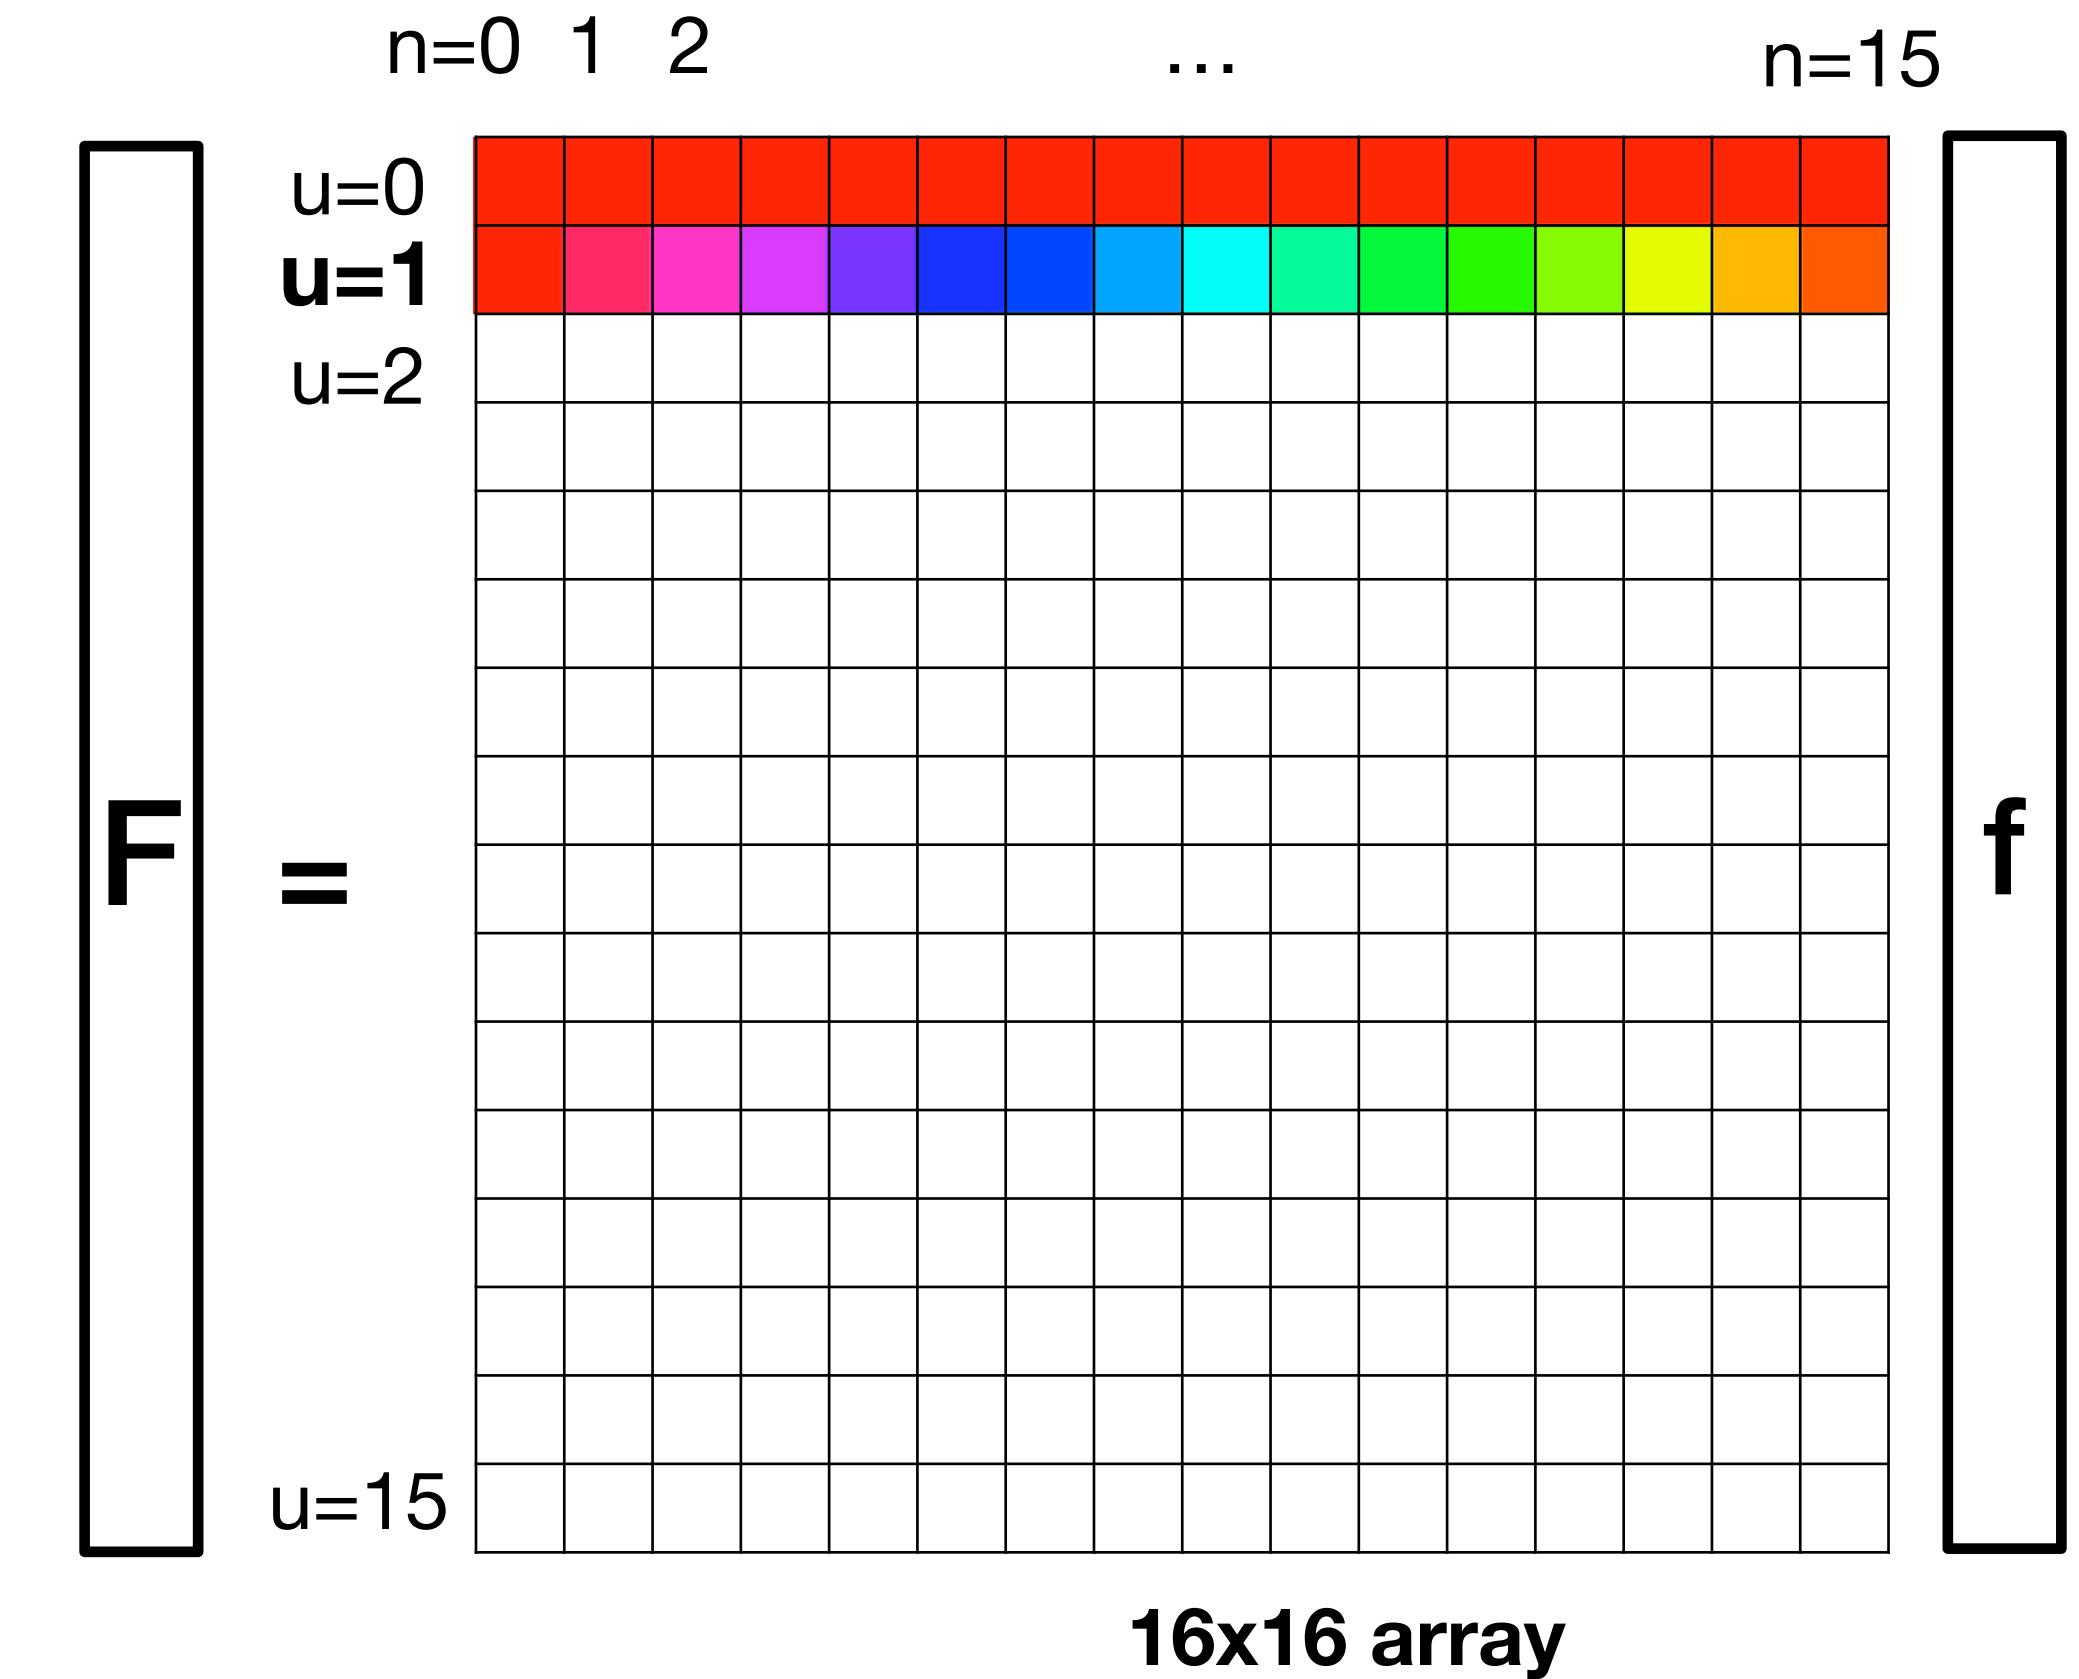
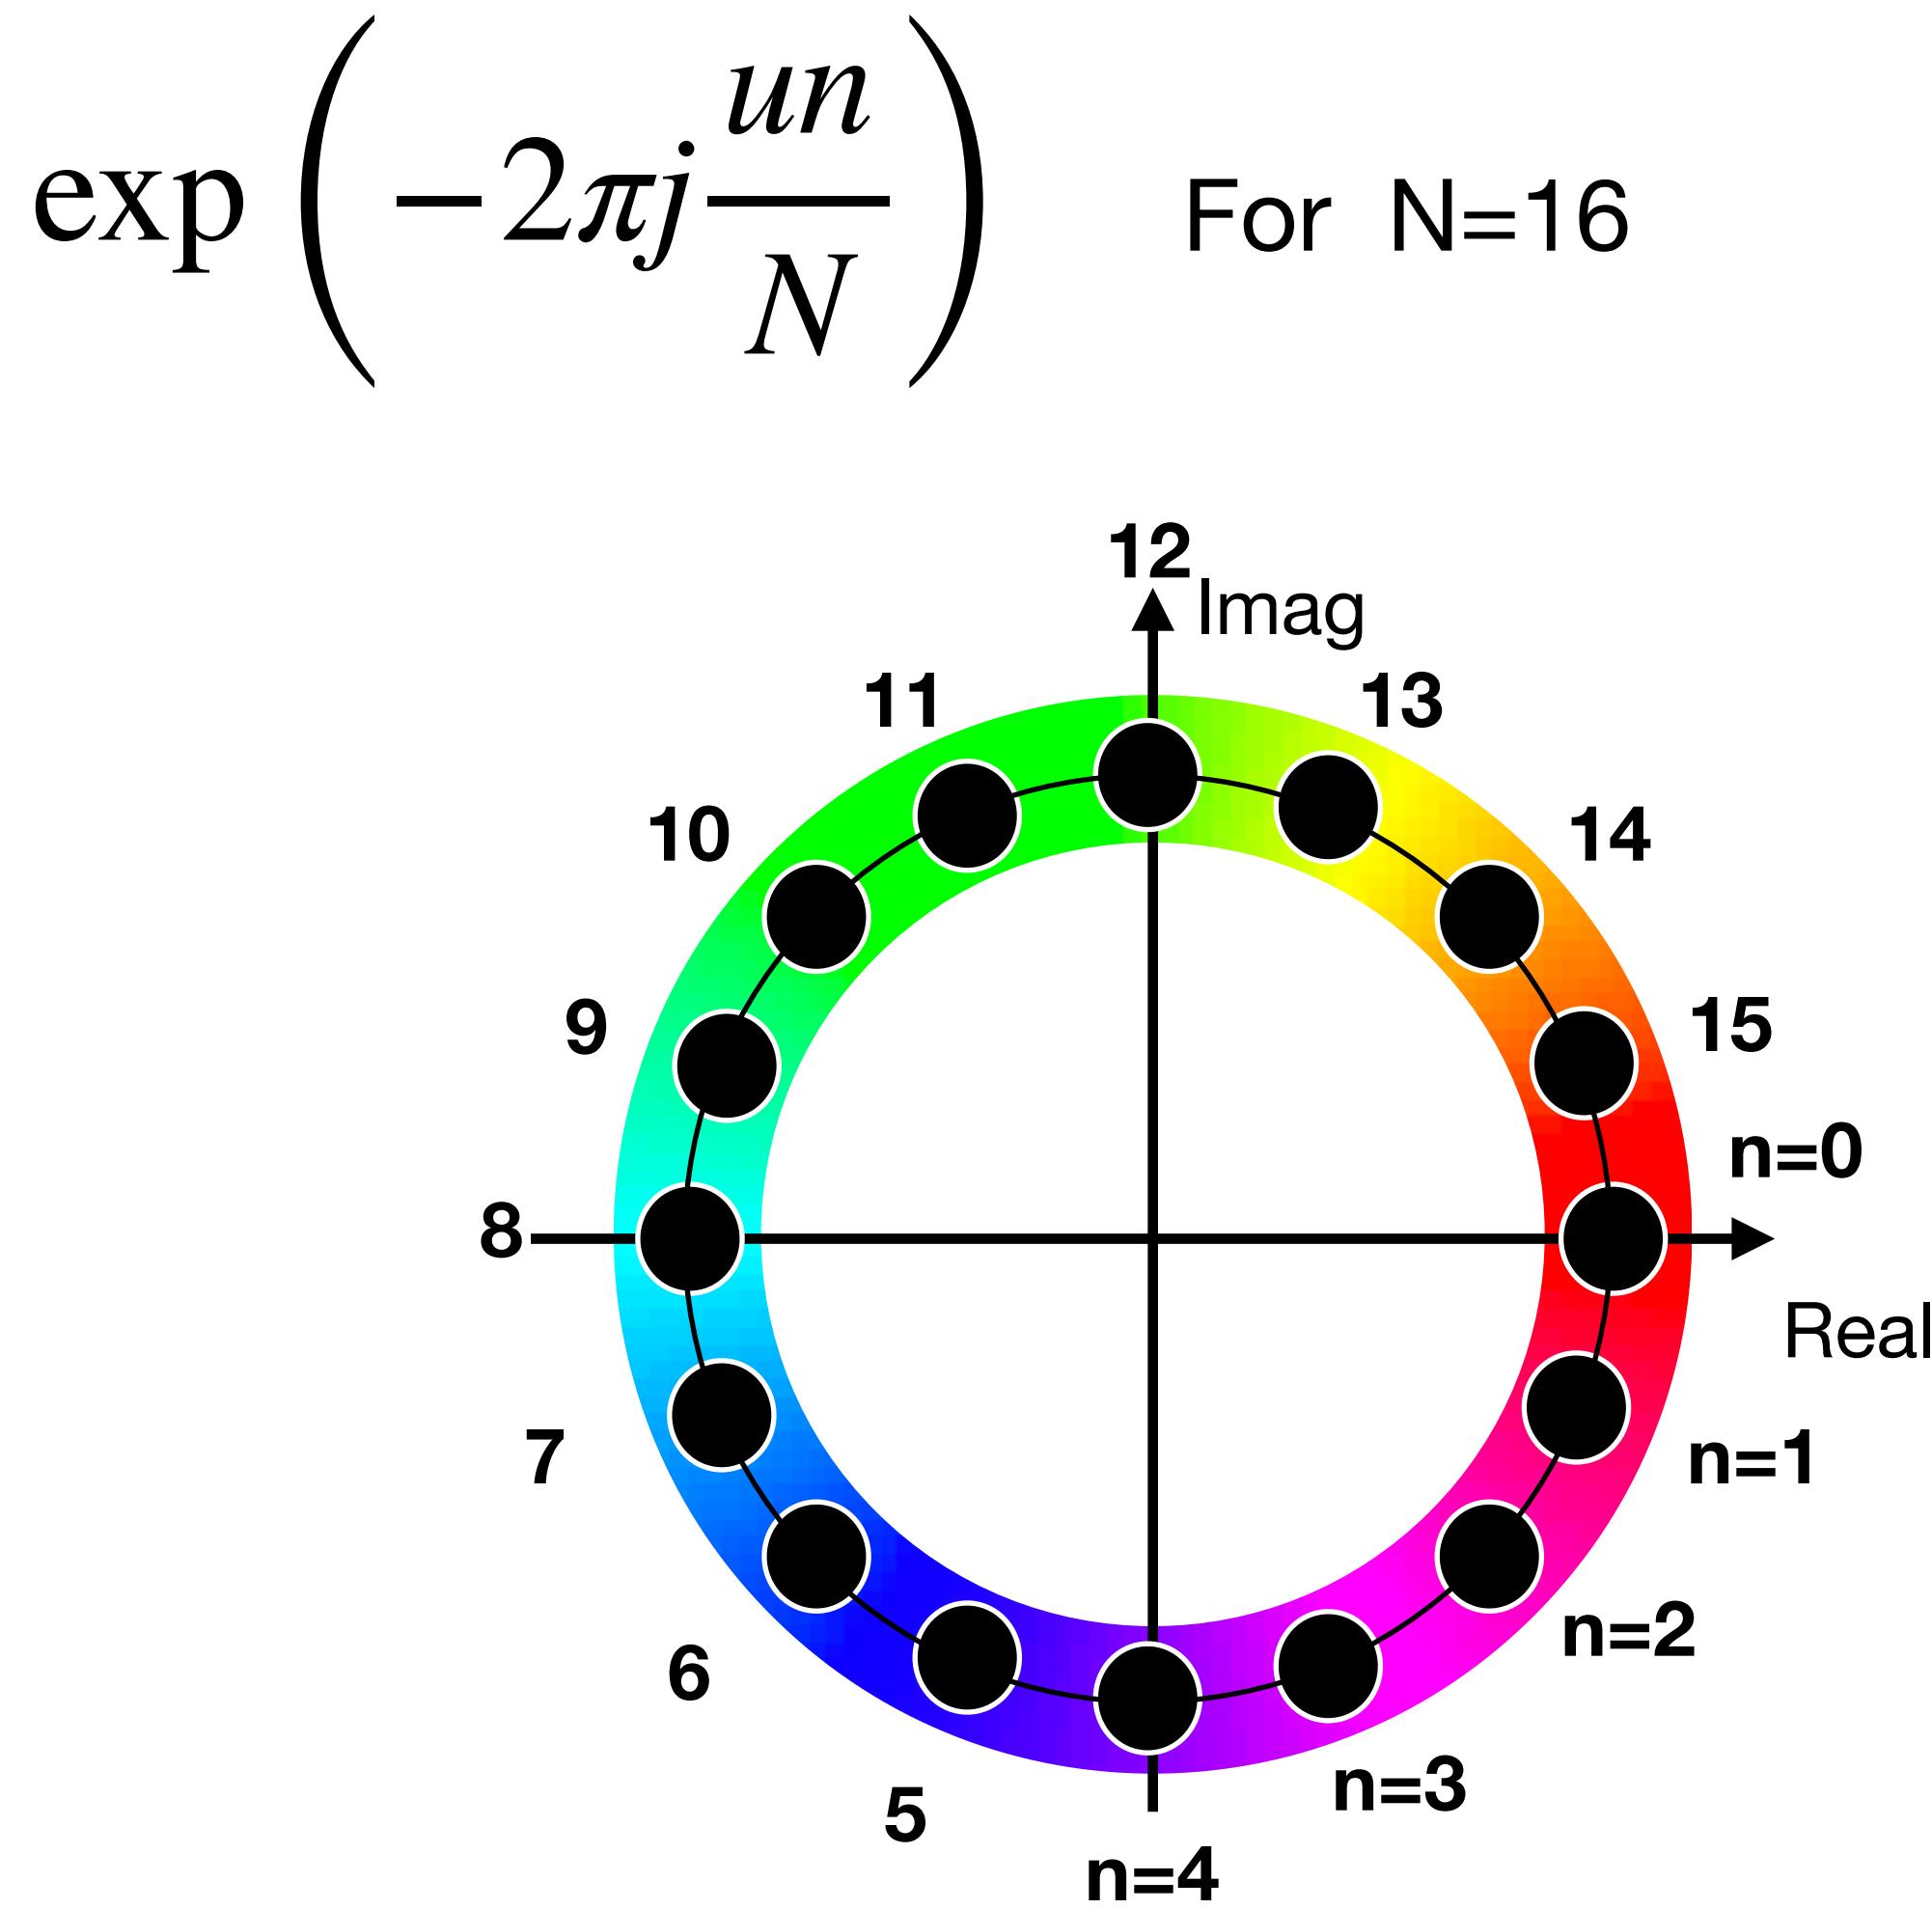
Visualizing the transform coefficients

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

For N=16



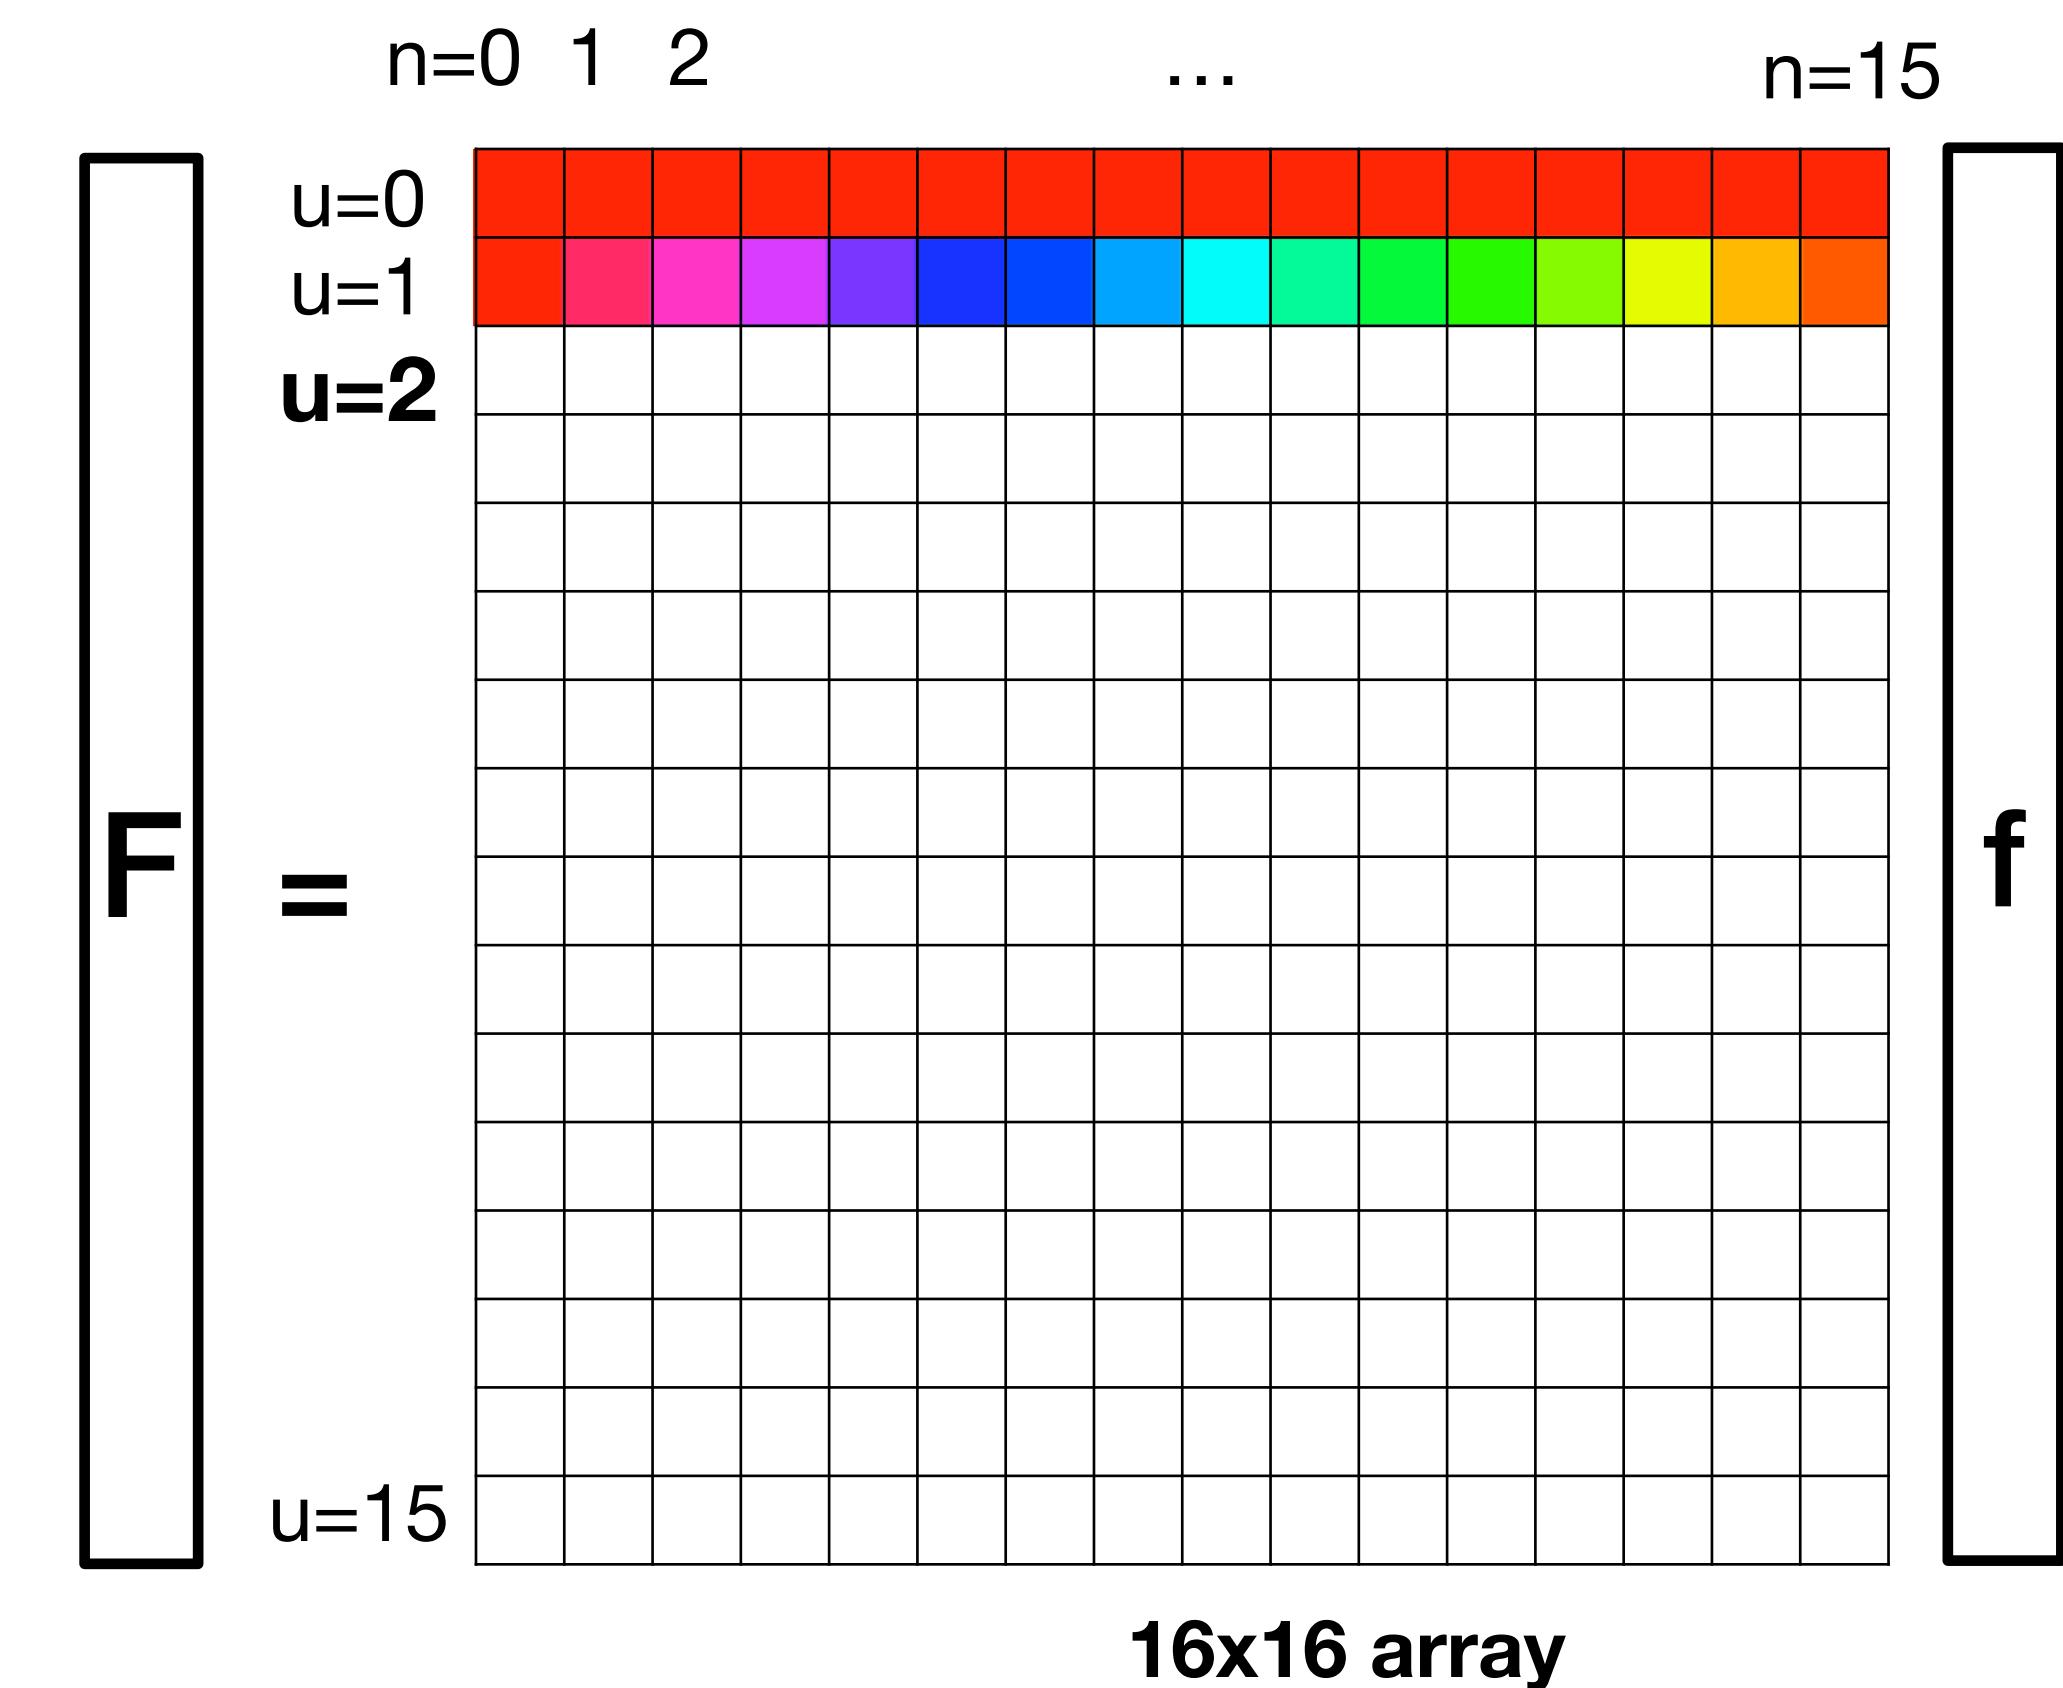
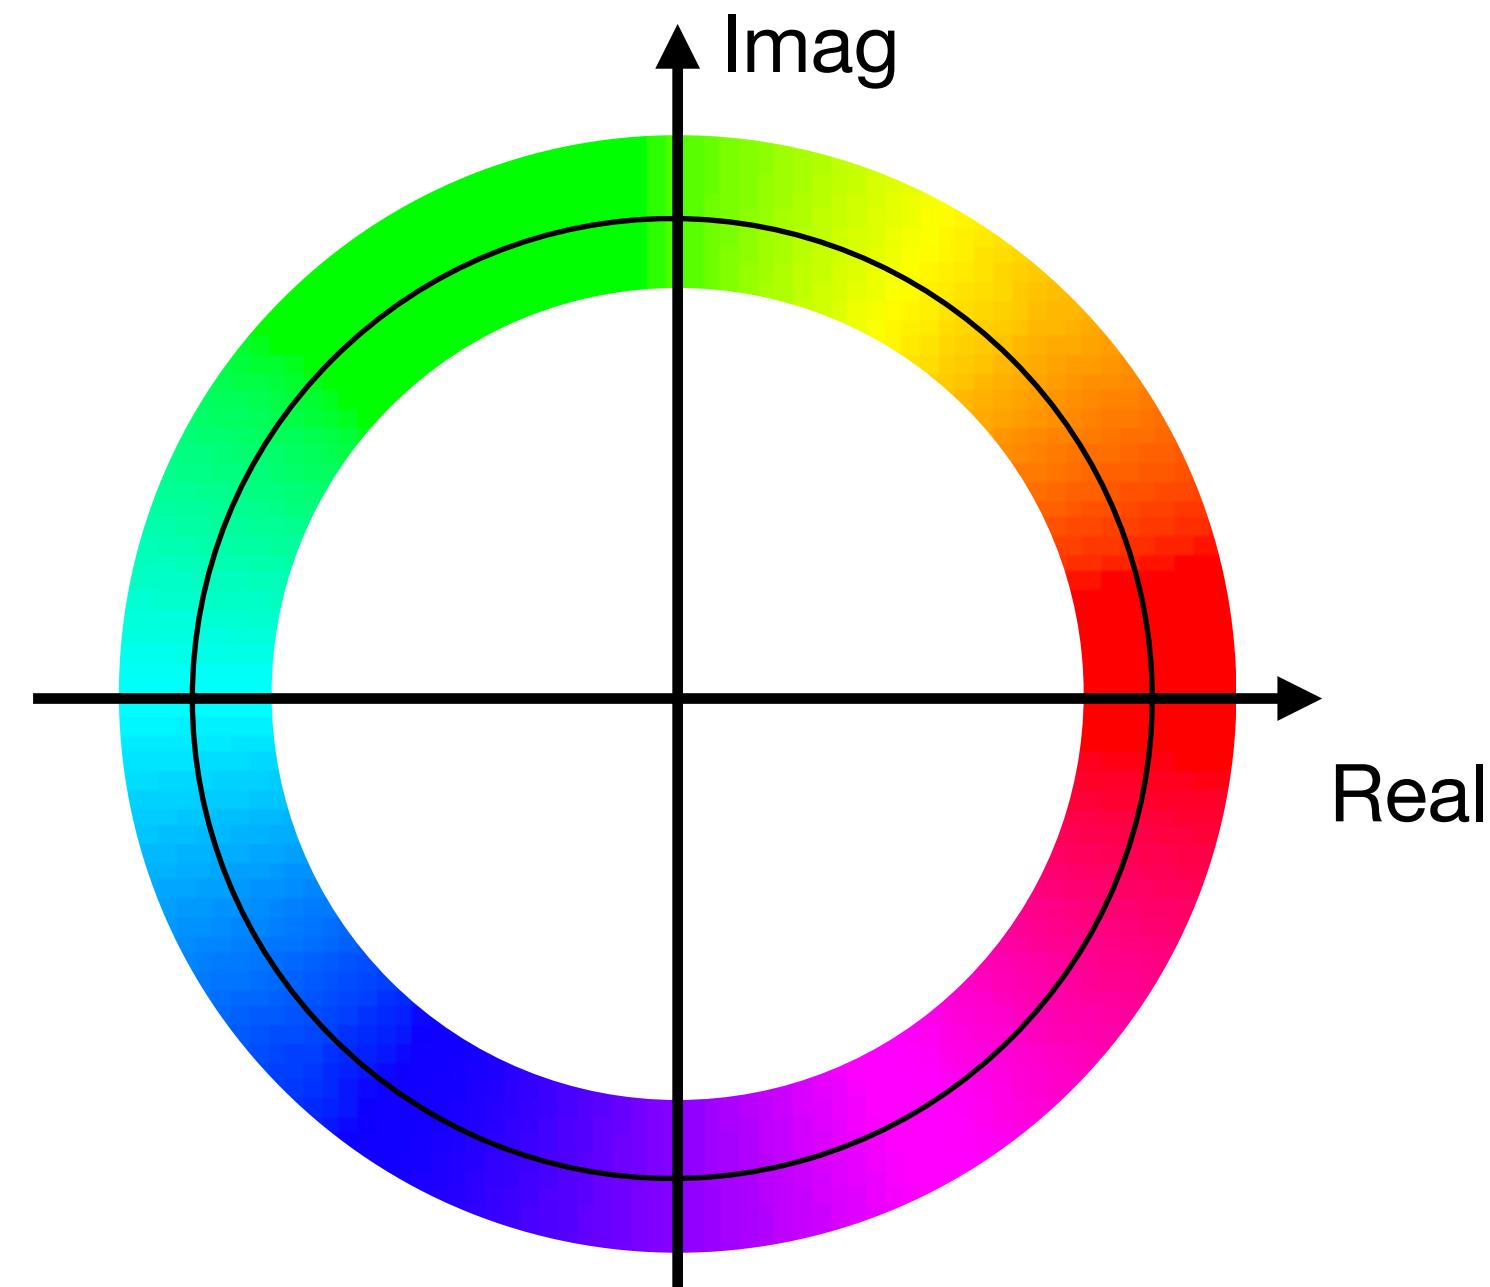
Visualizing the transform coefficients



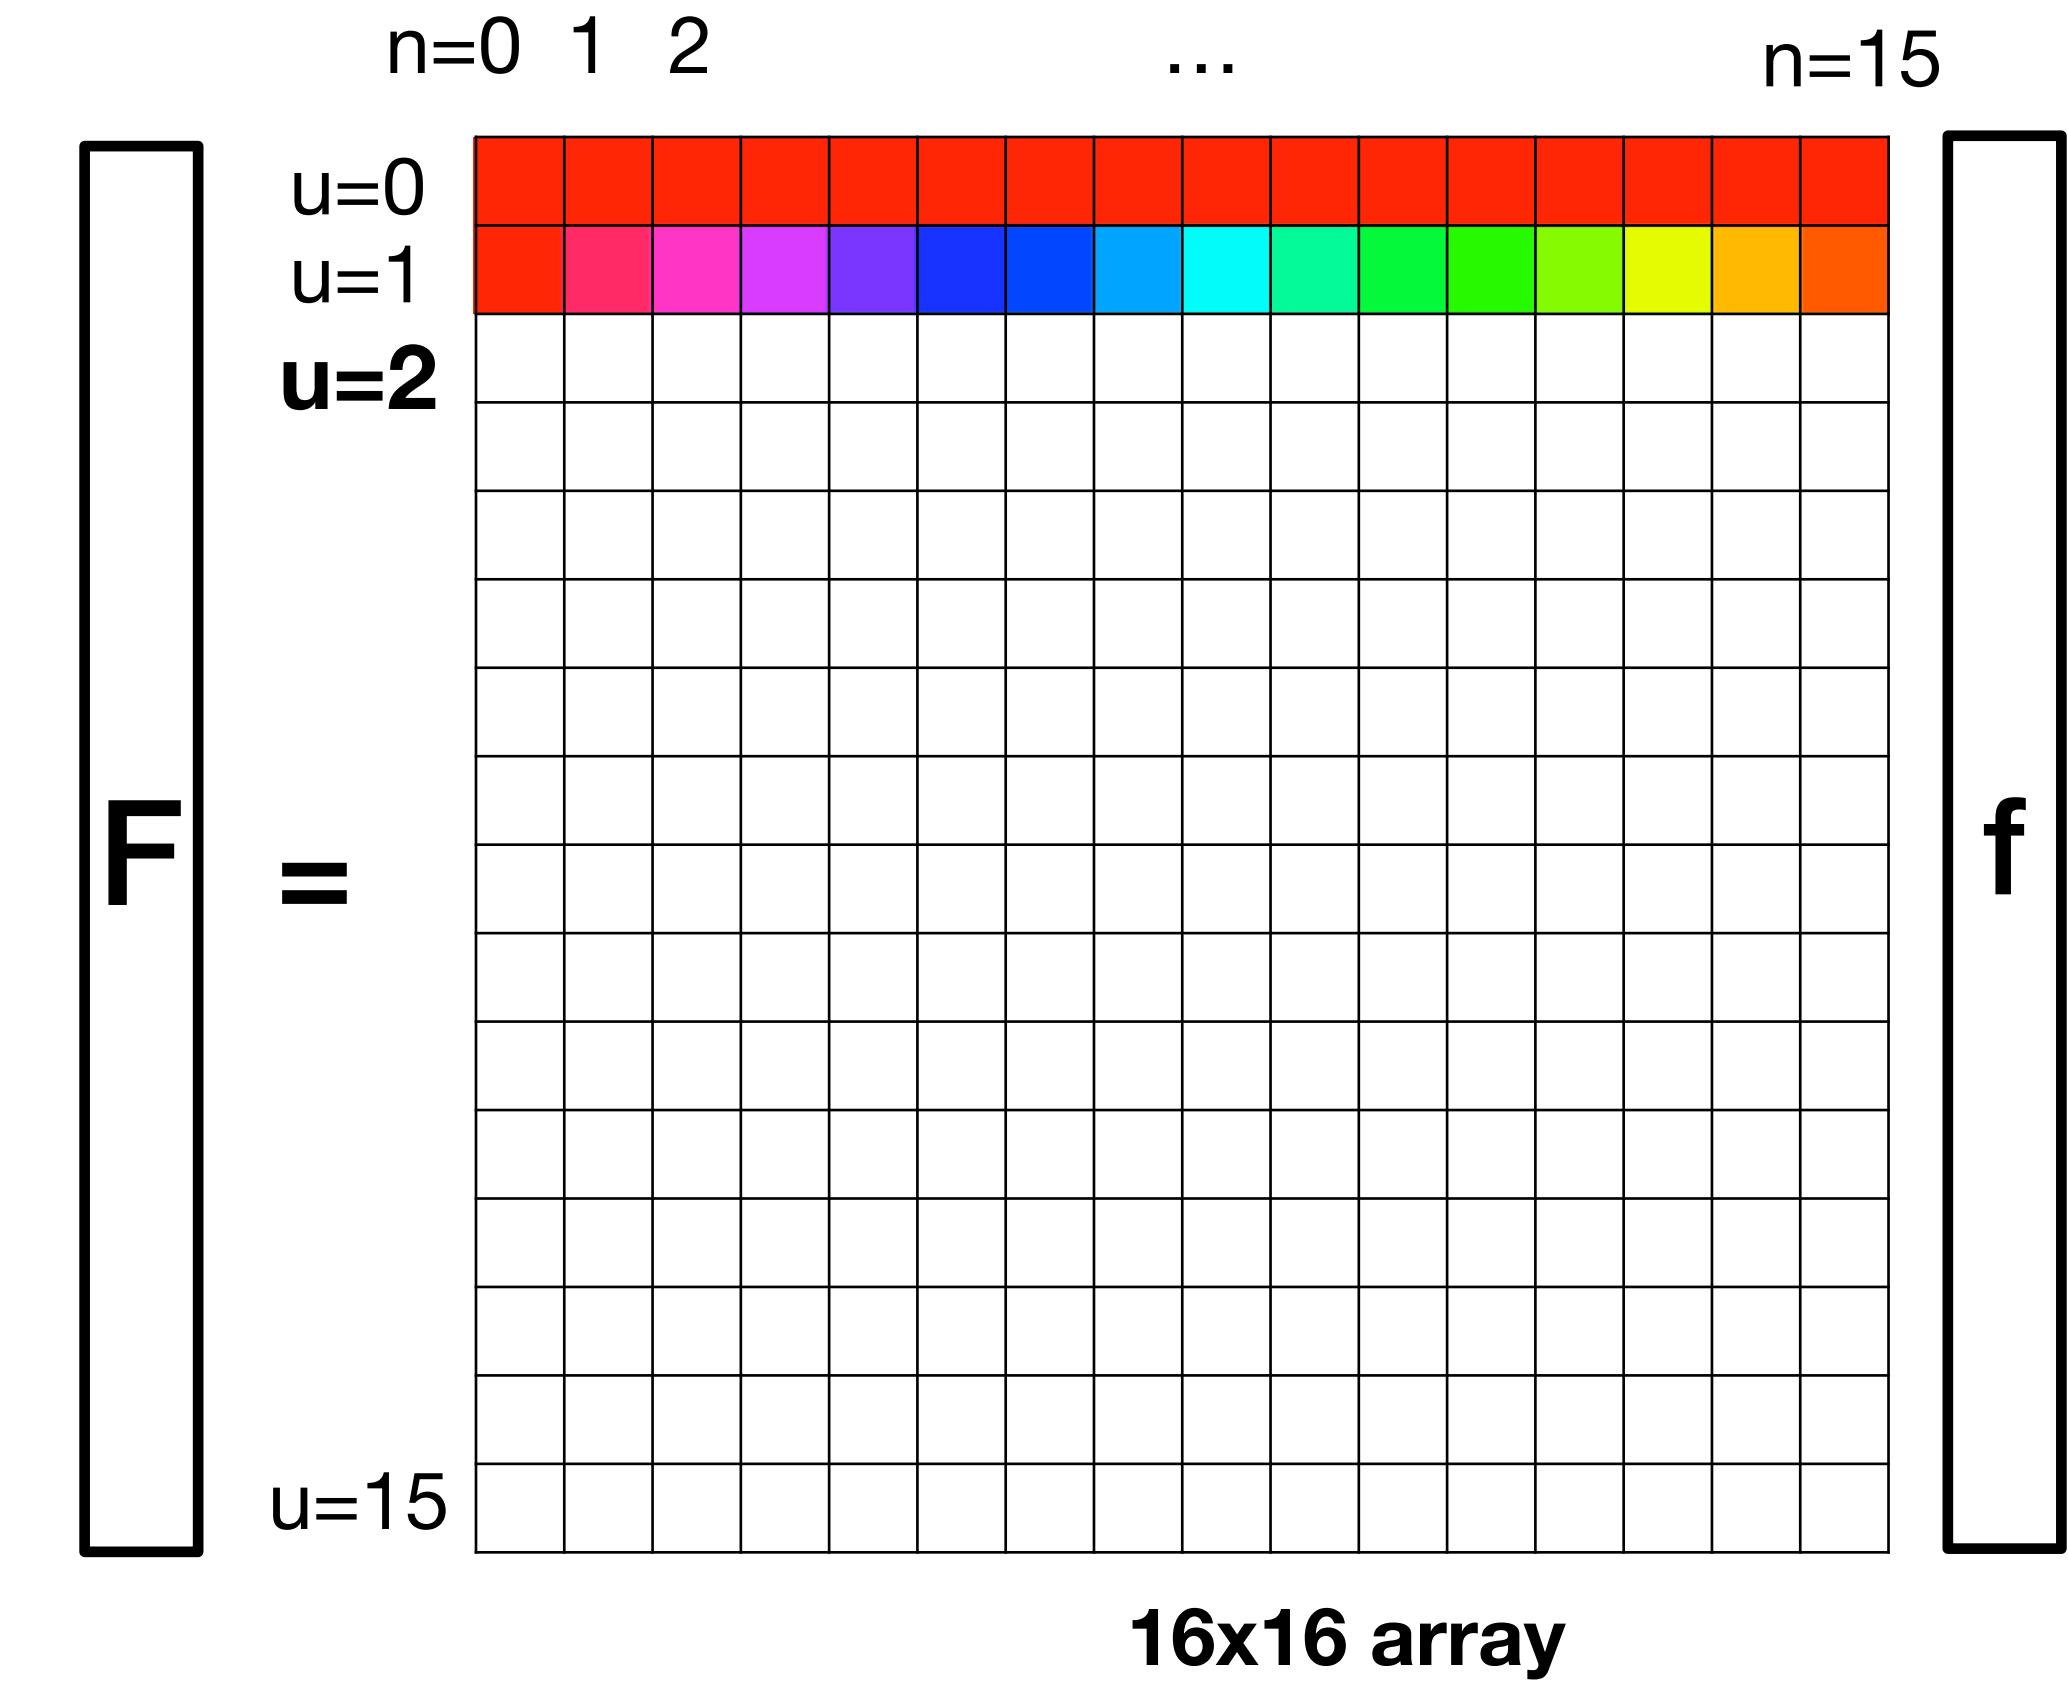
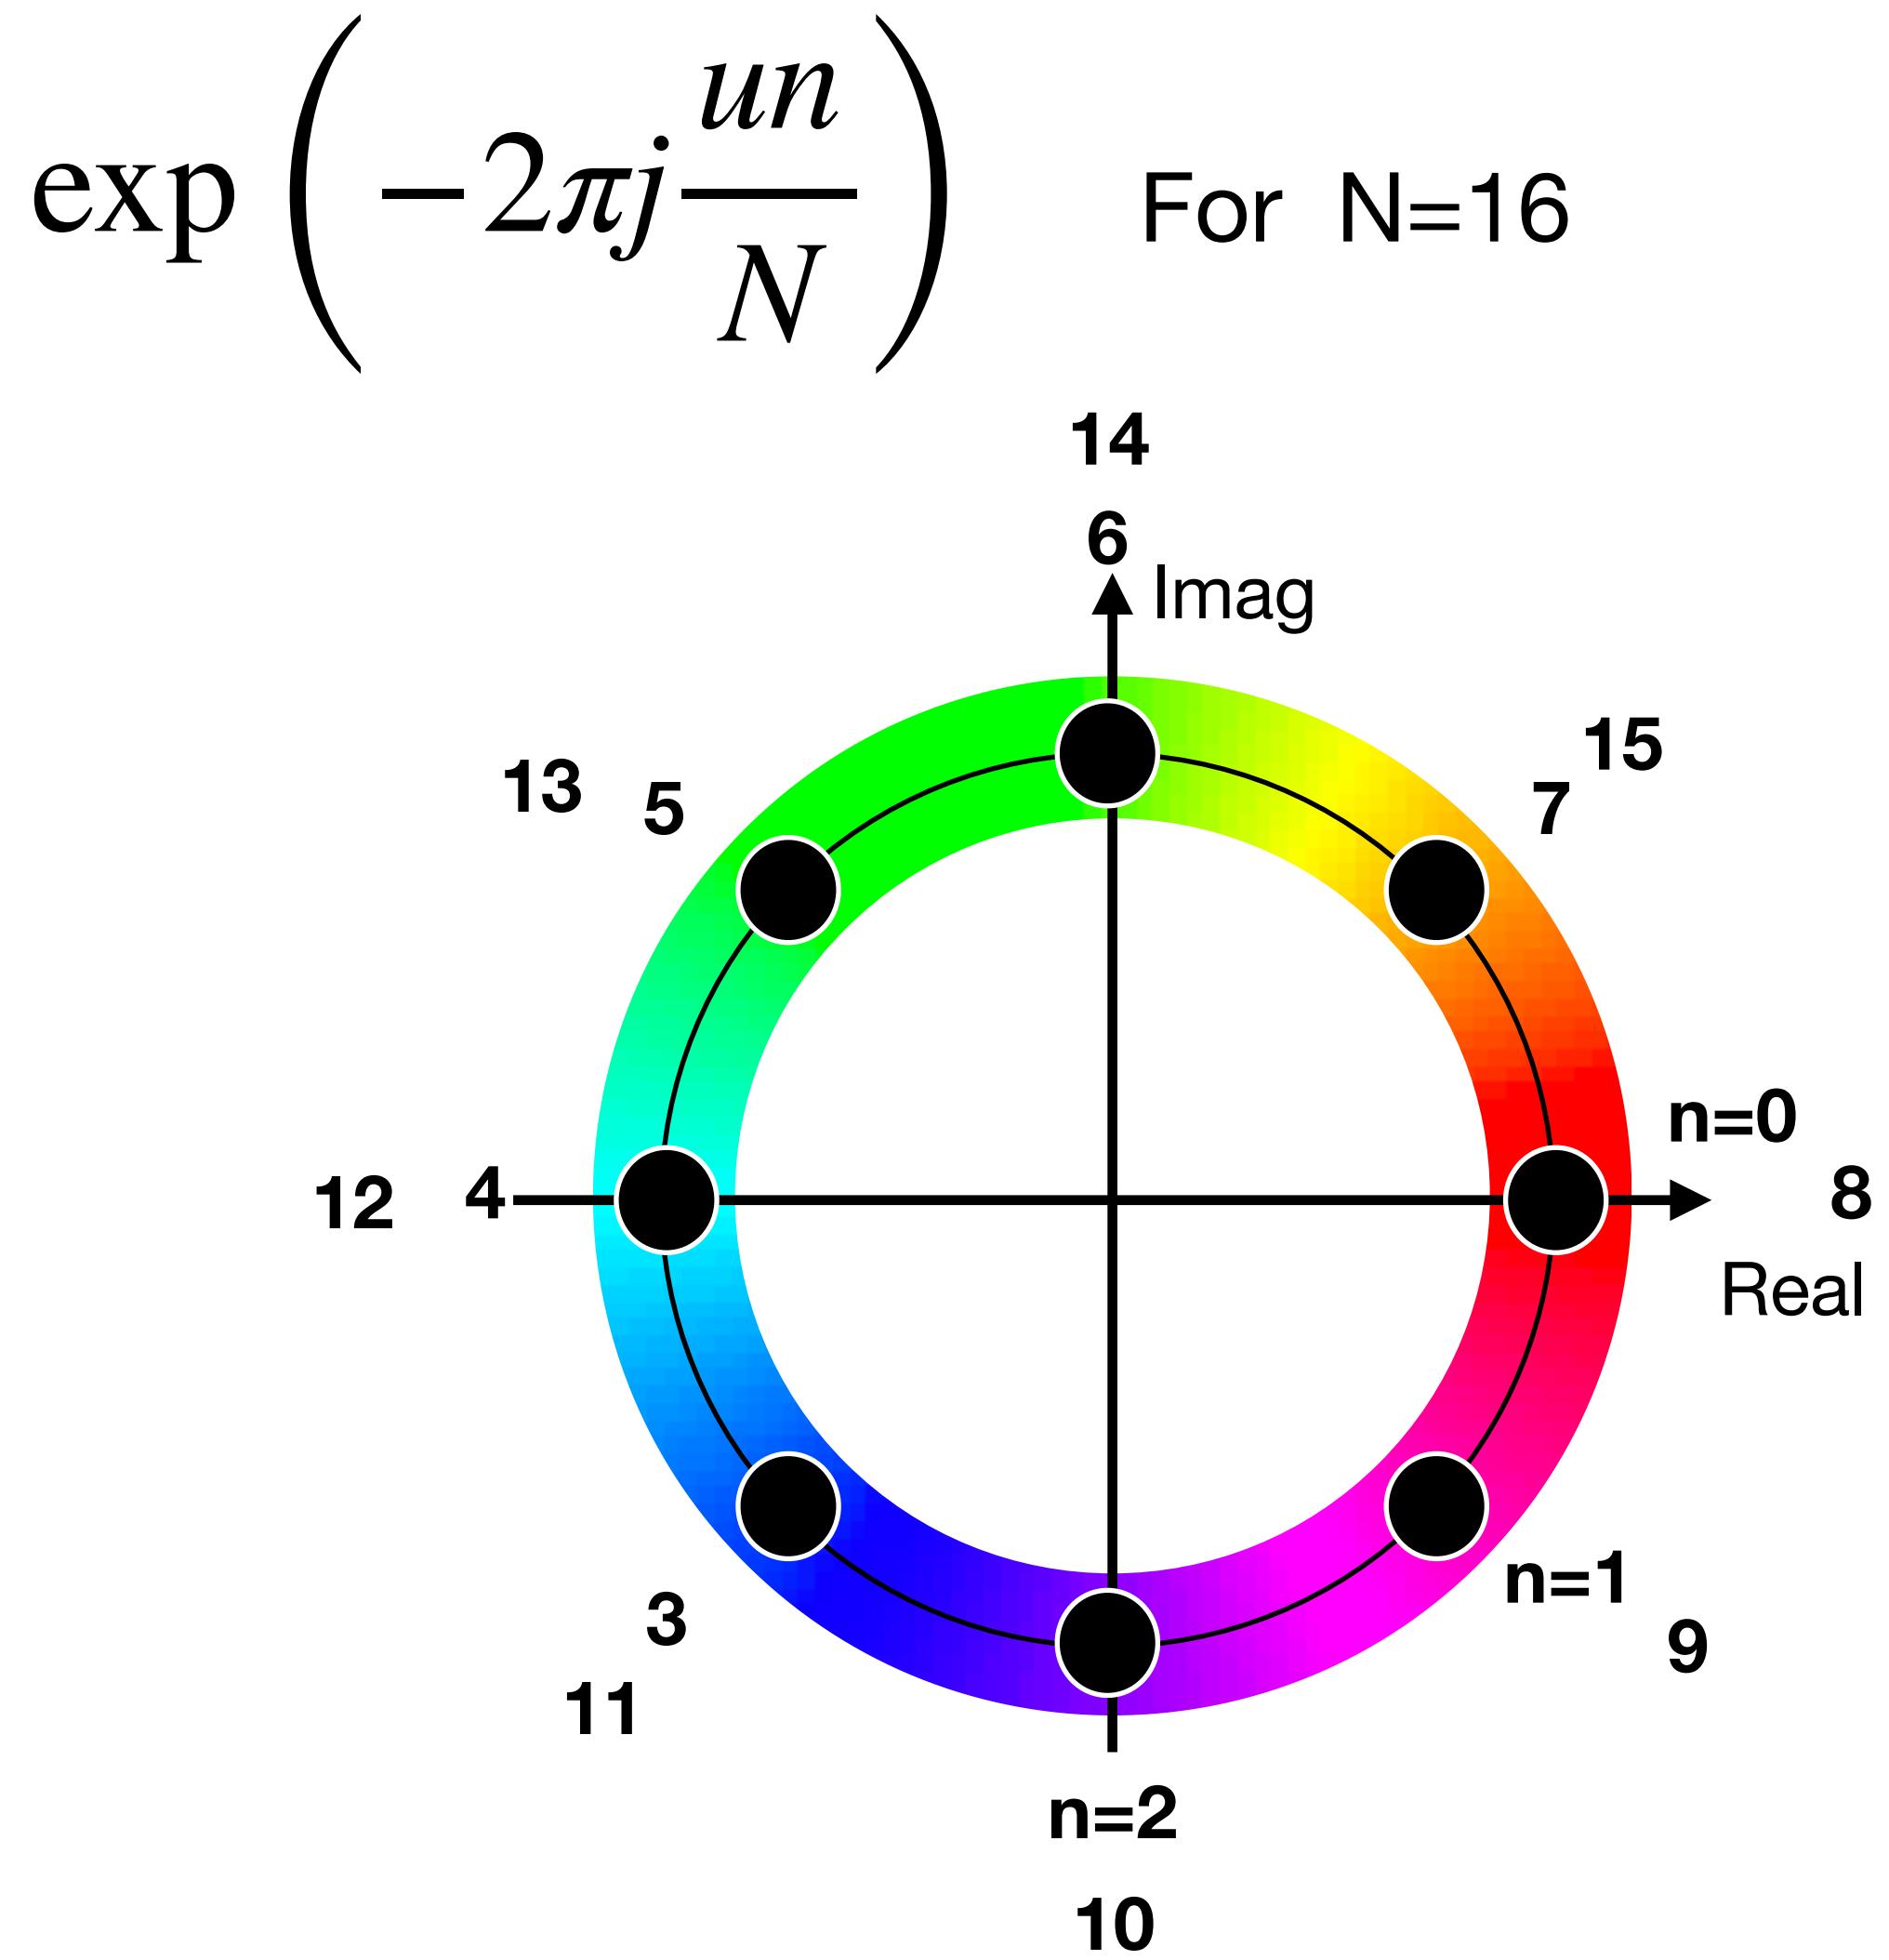
Visualizing the transform coefficients

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

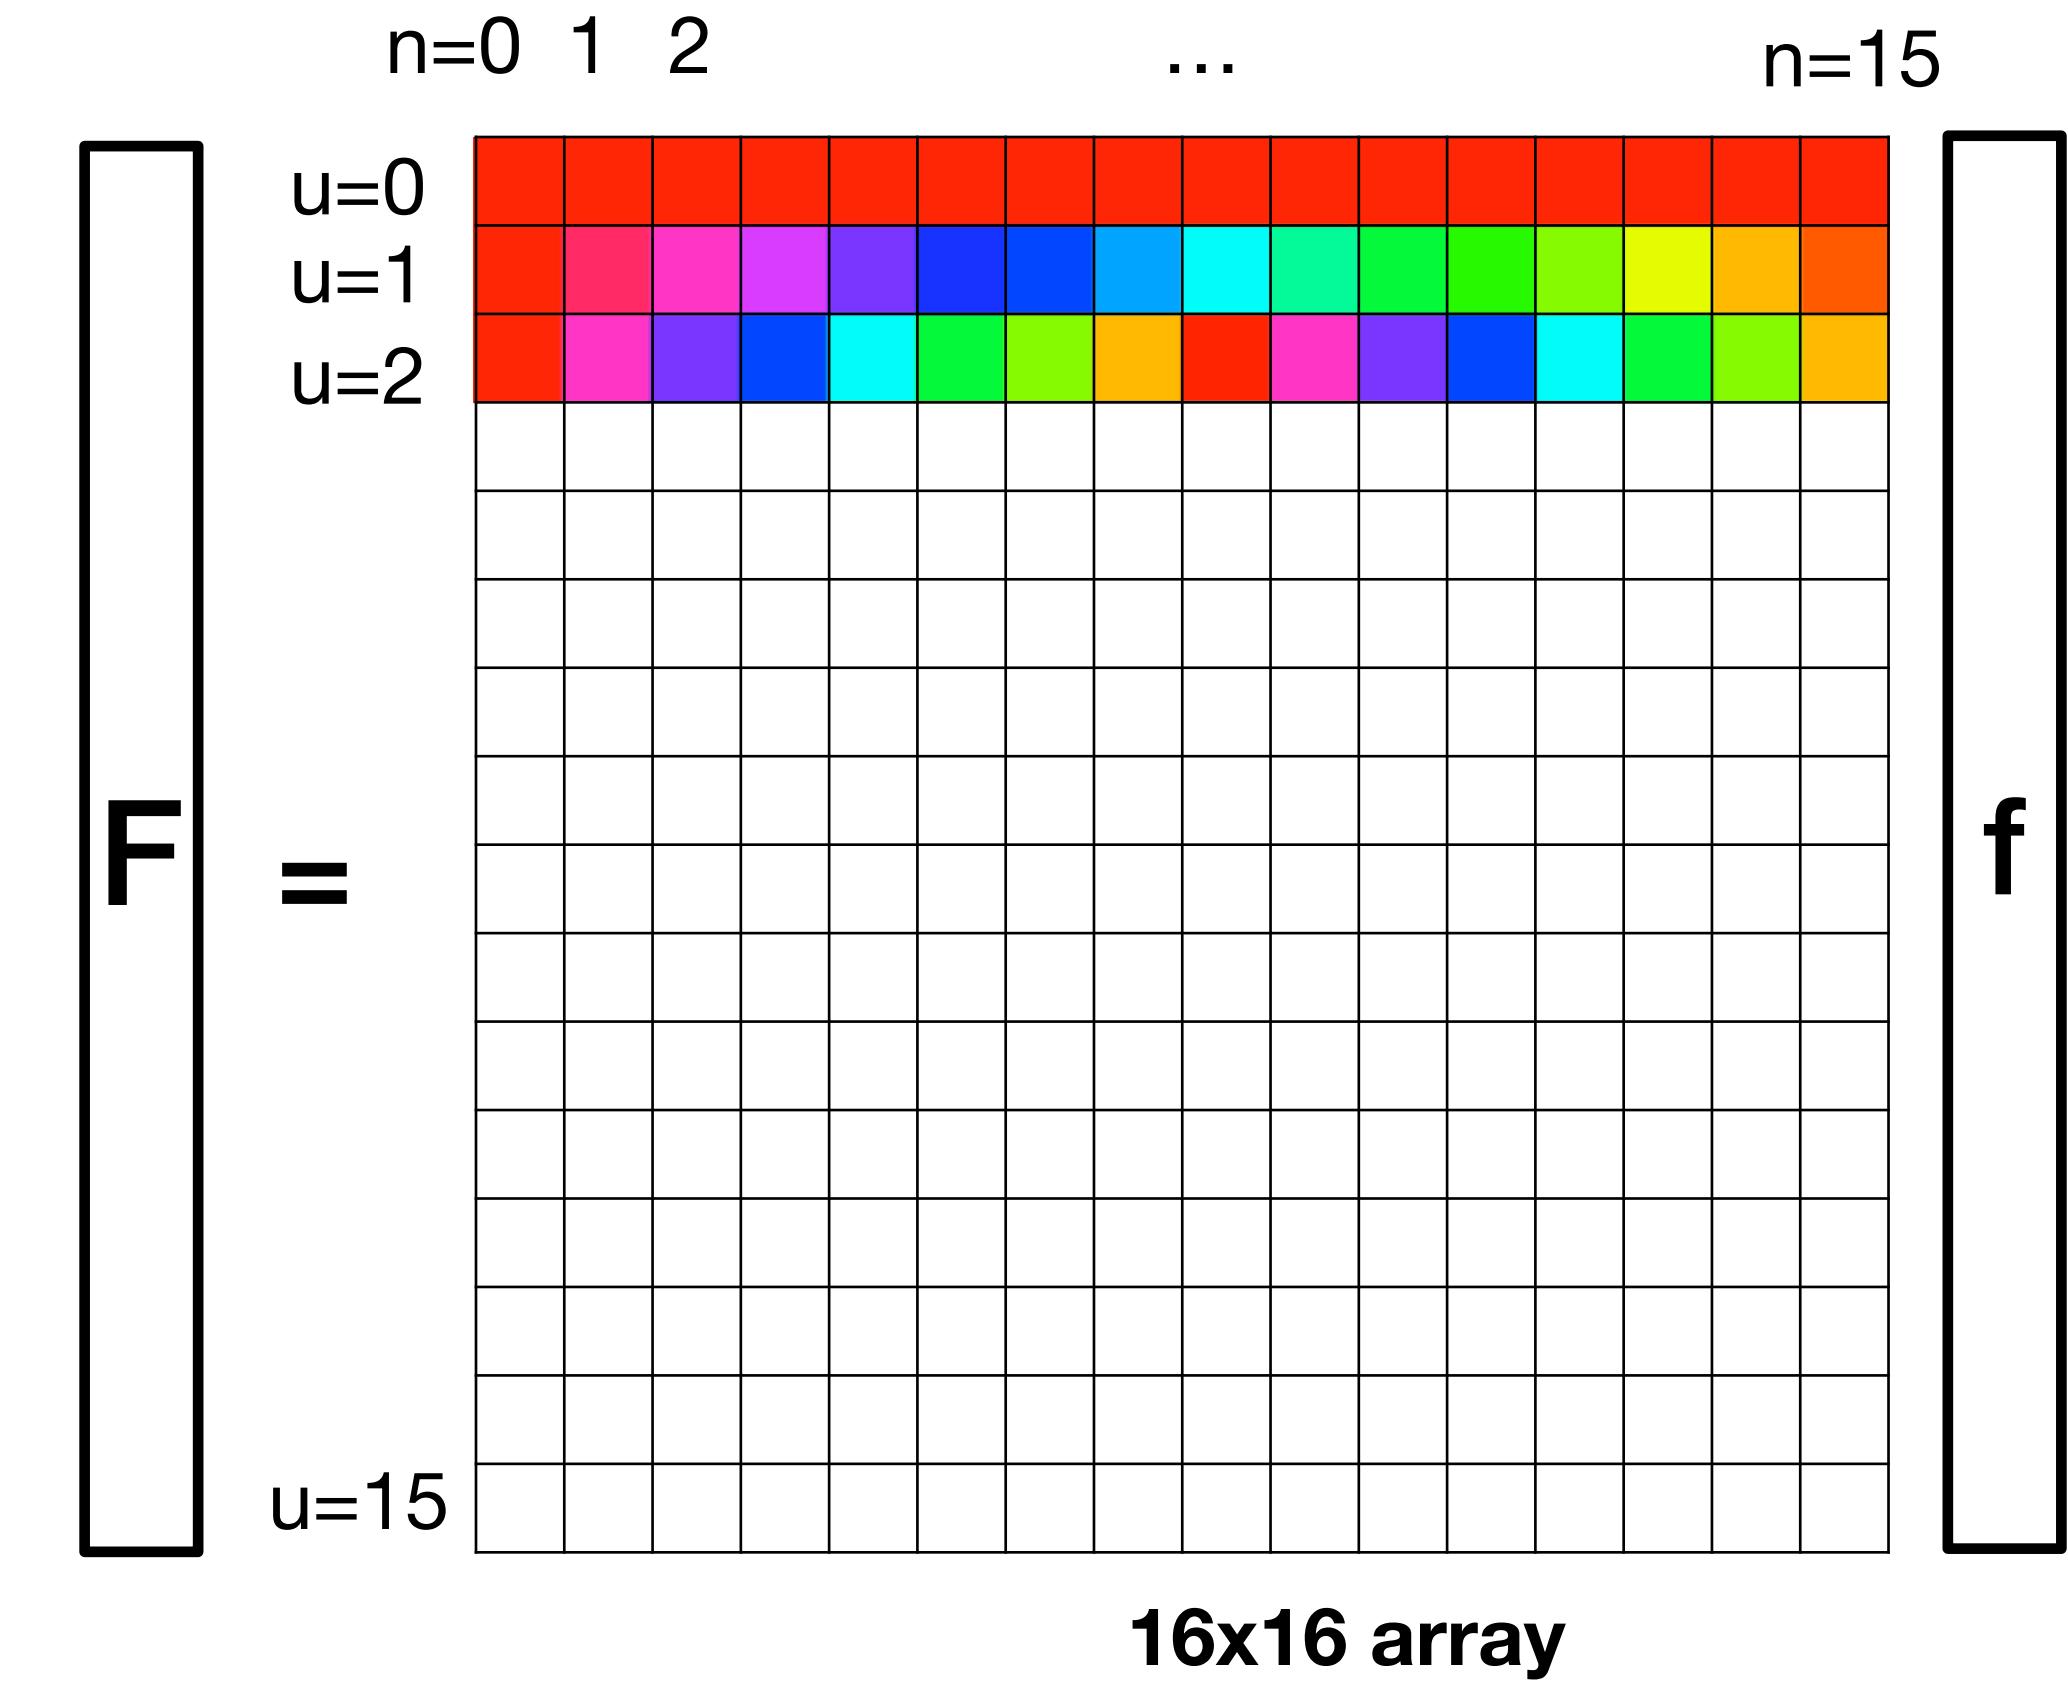
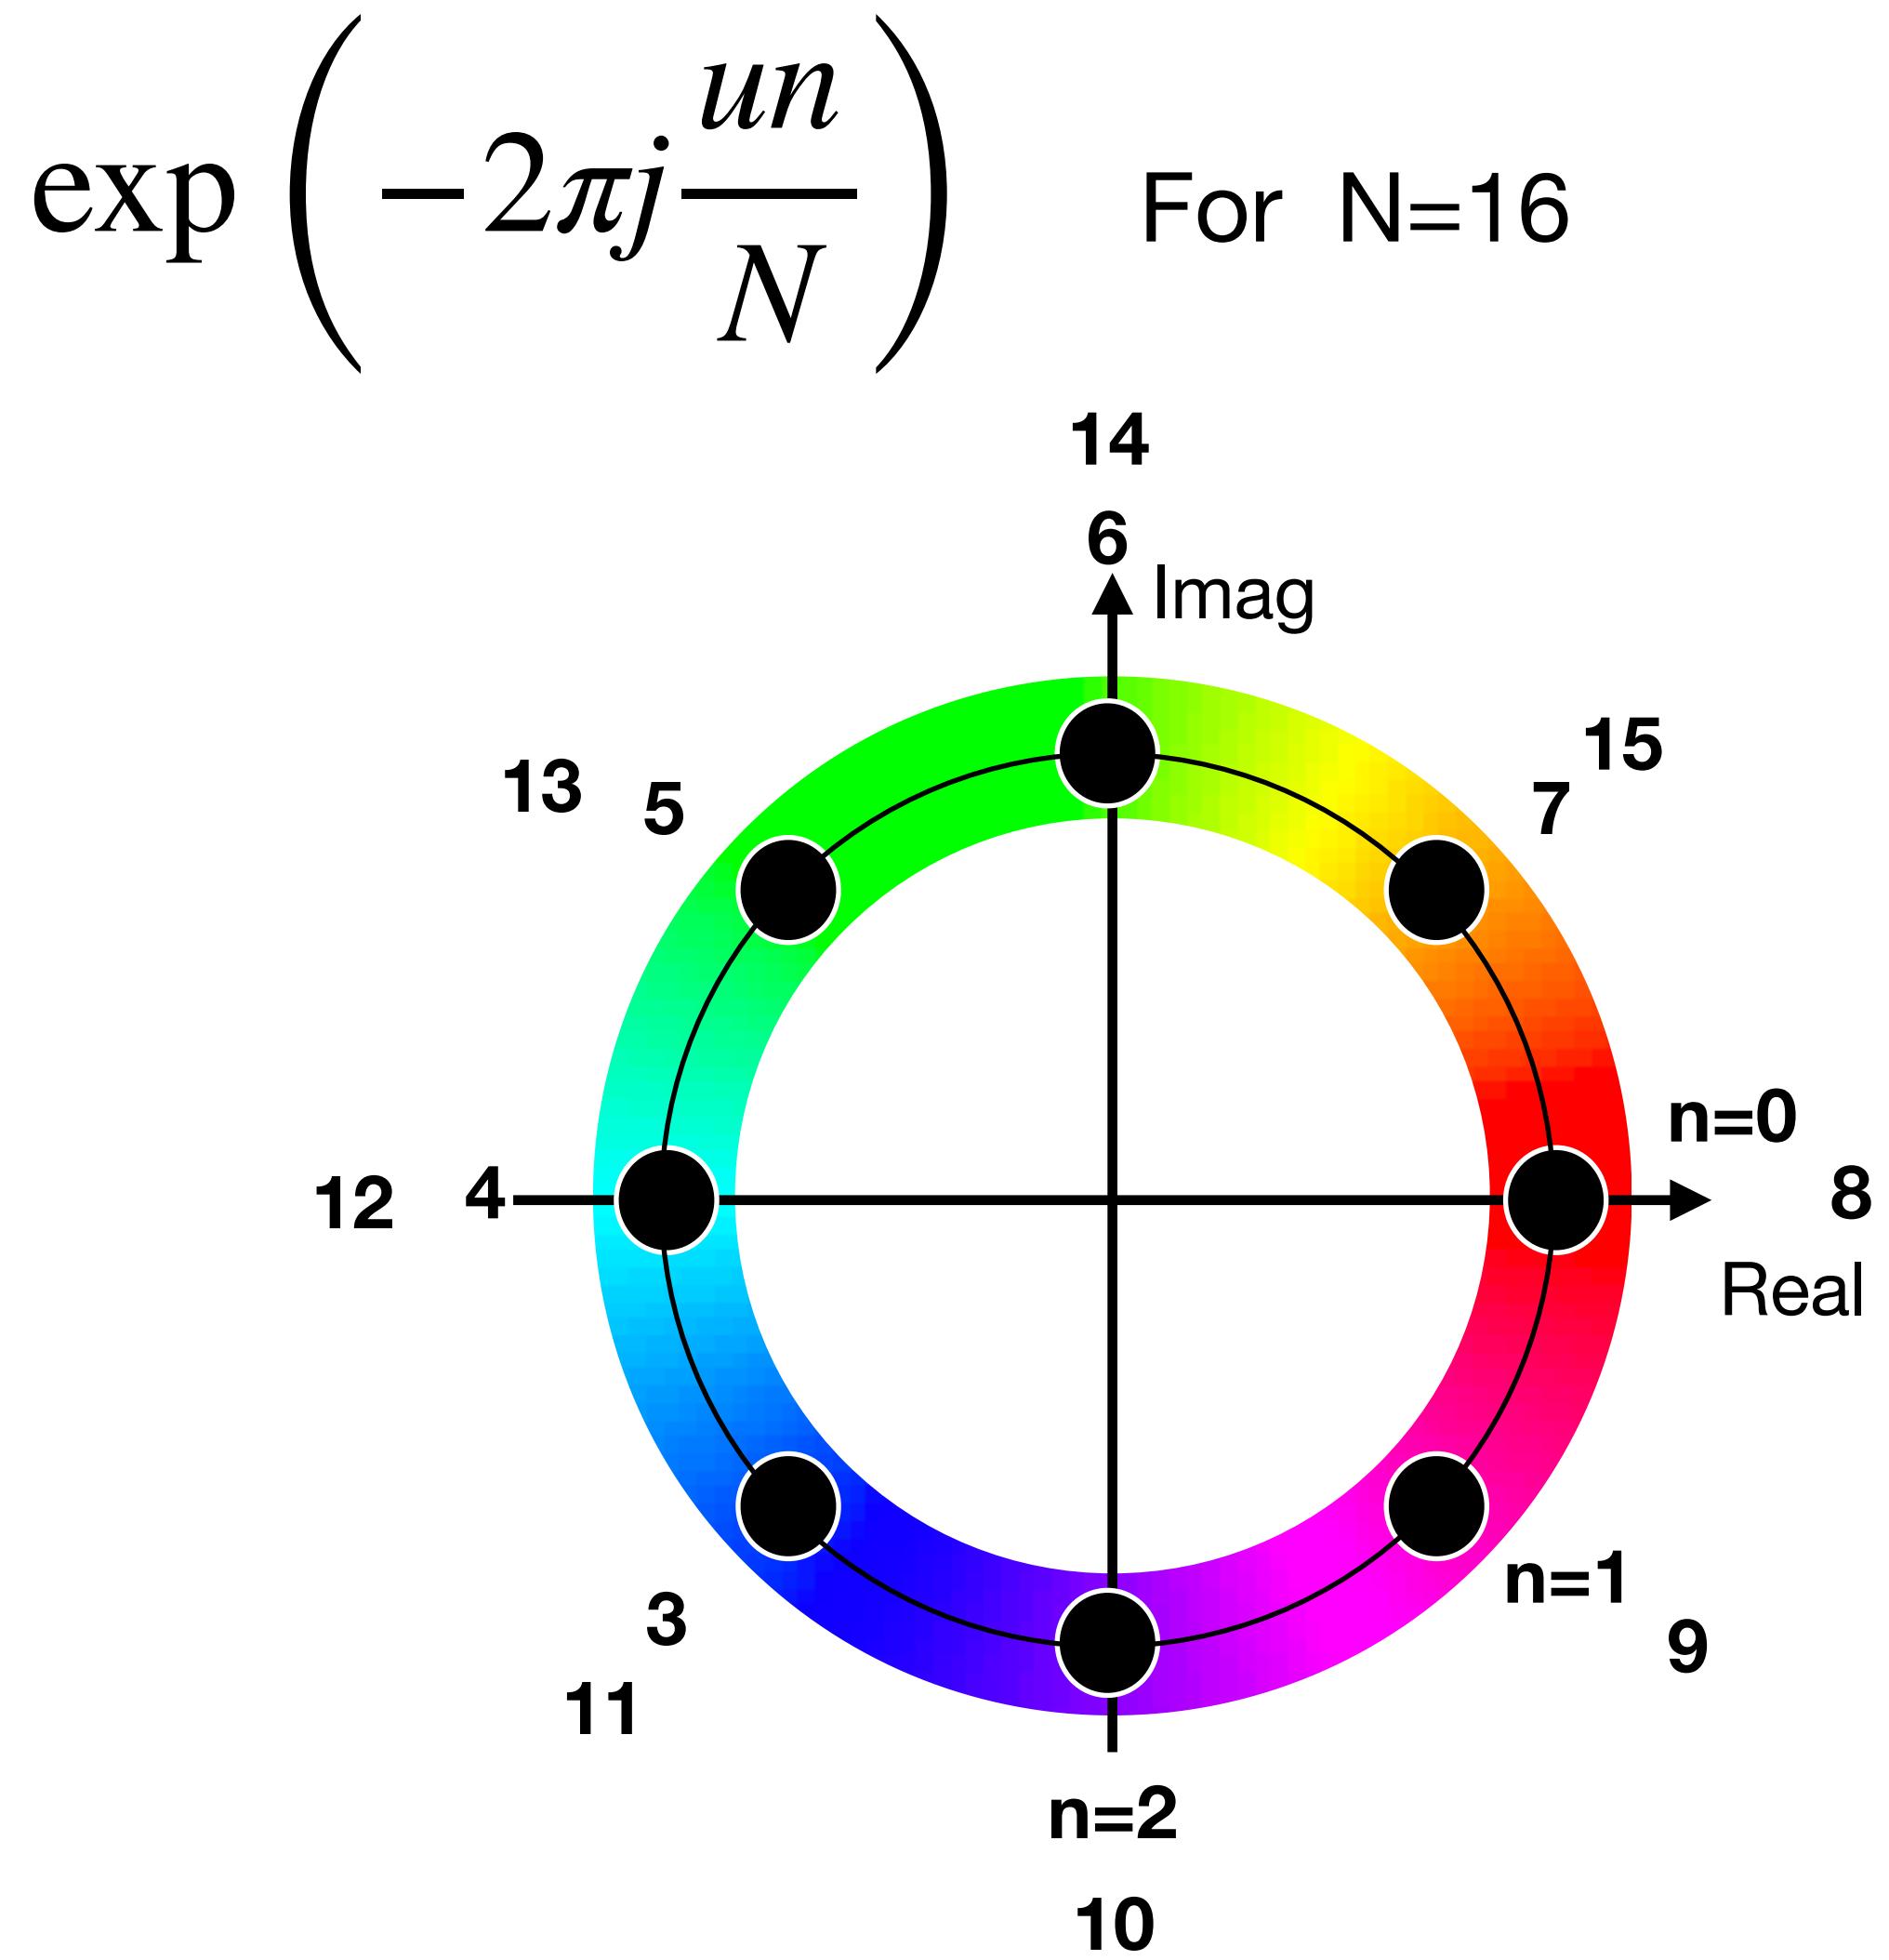
For N=16



Visualizing the transform coefficients



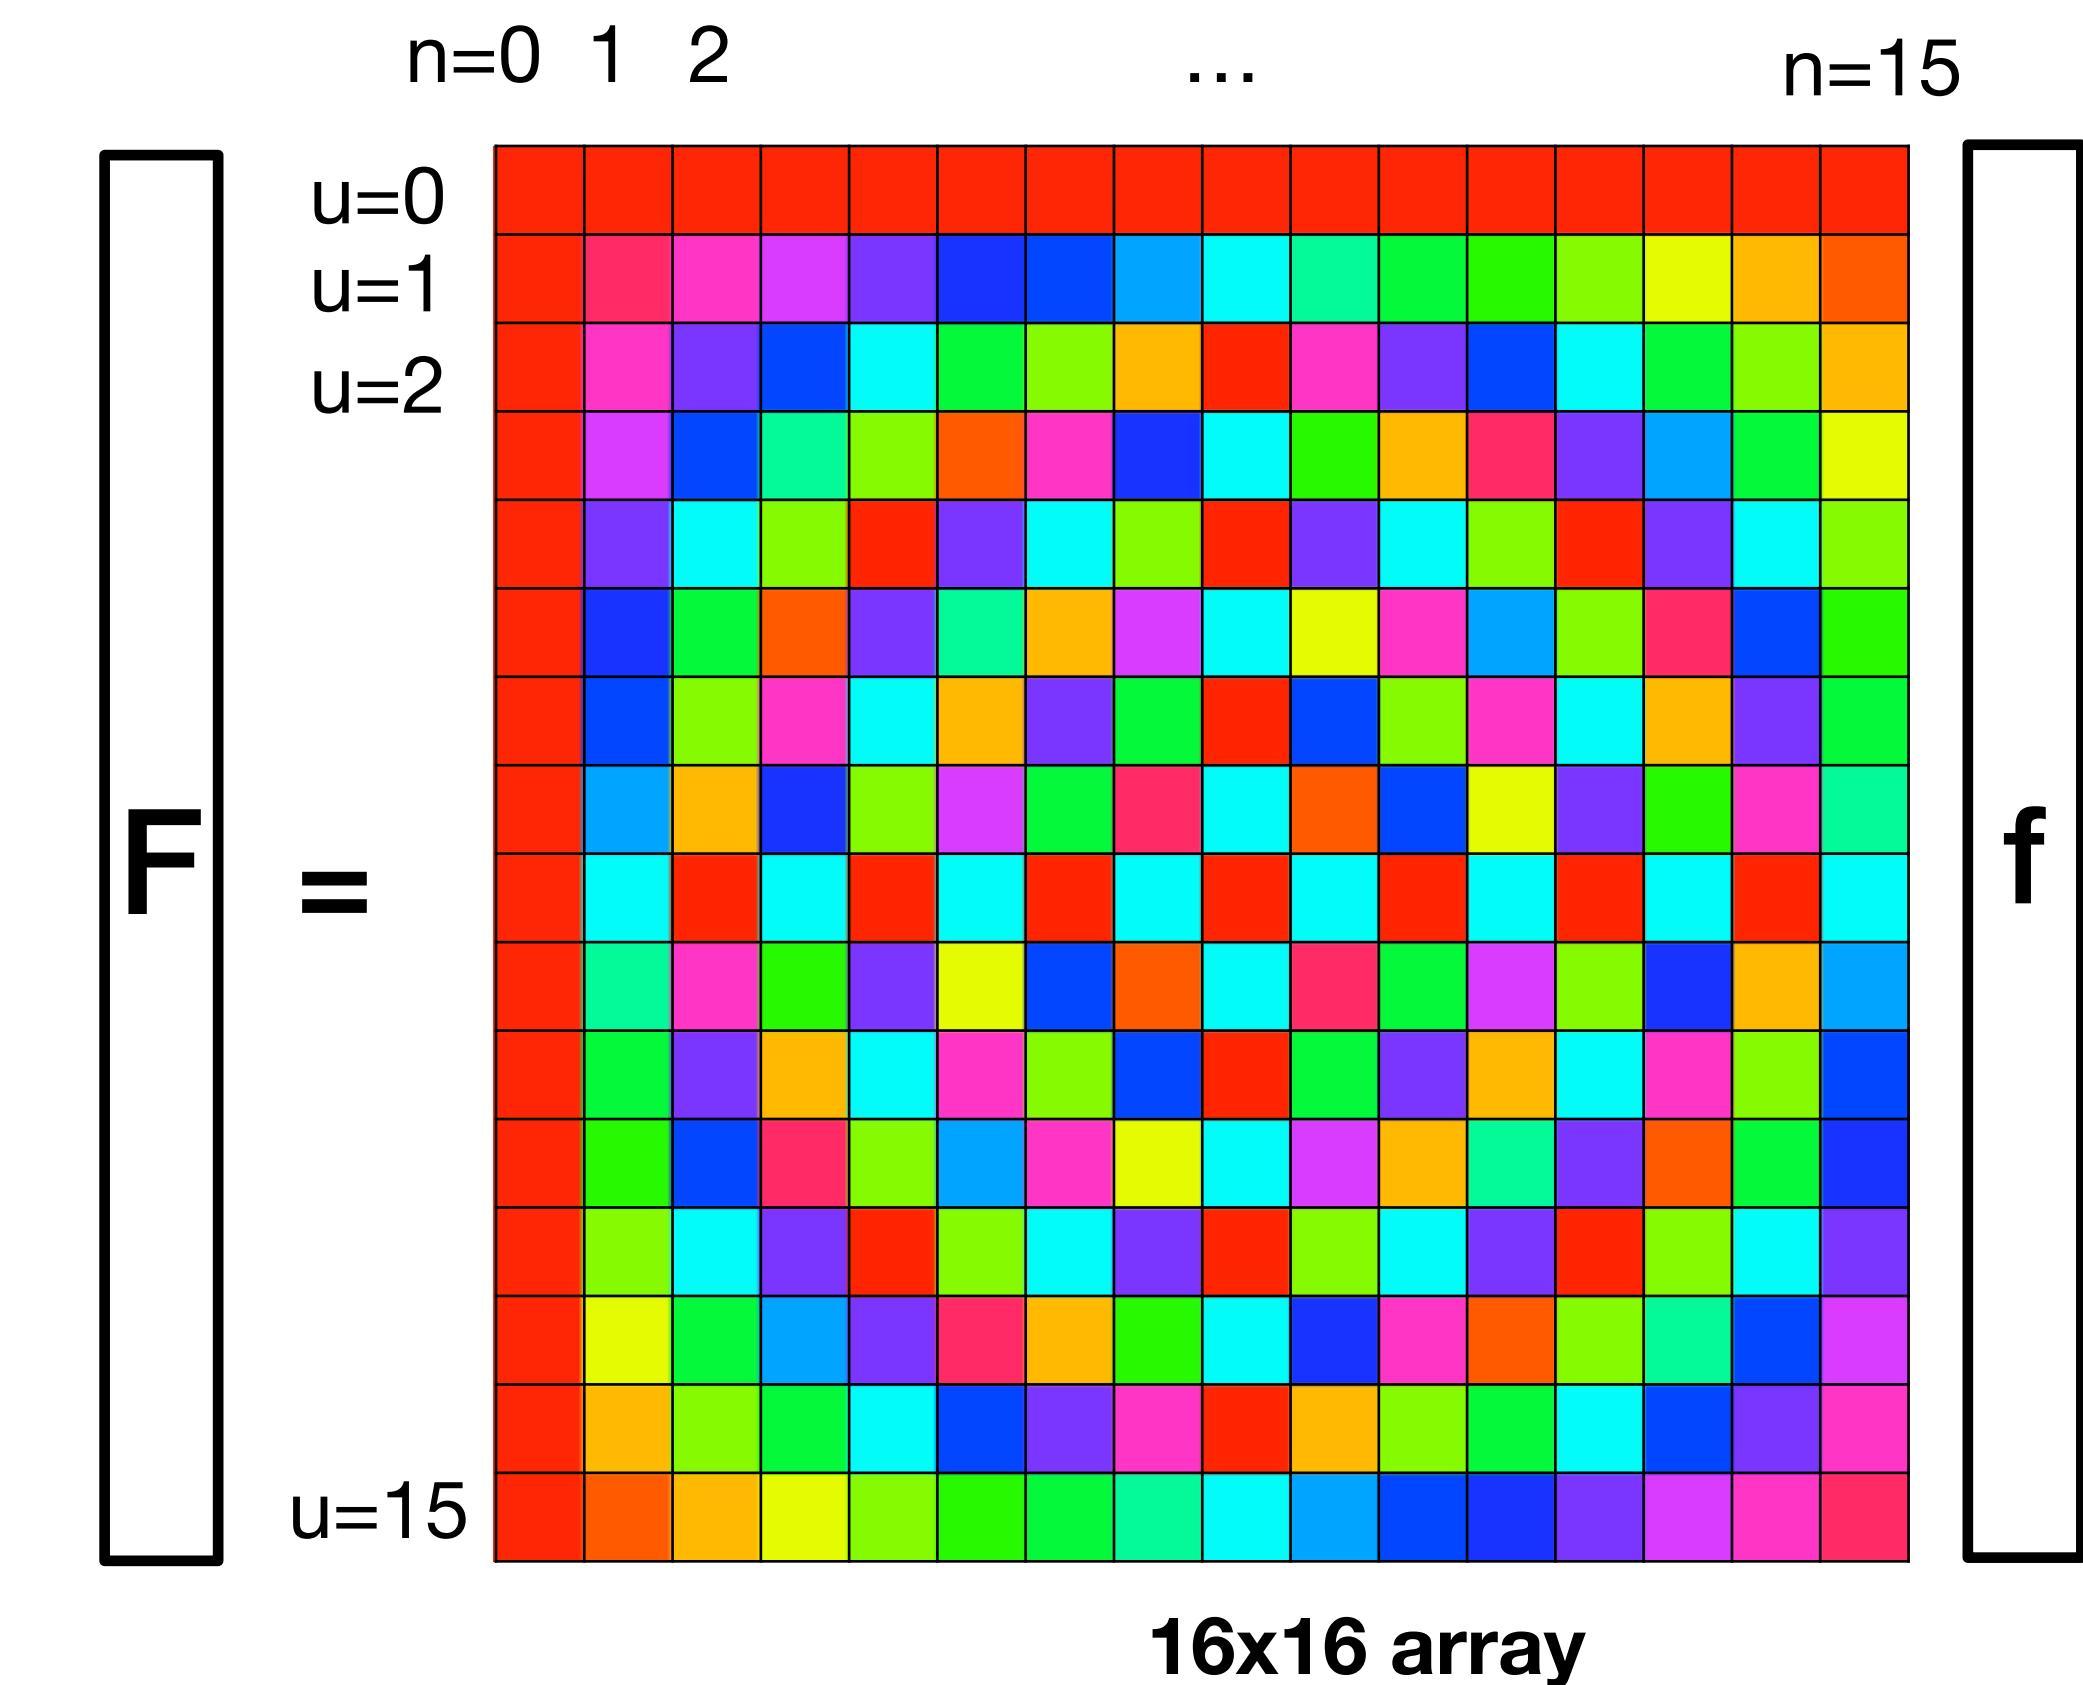
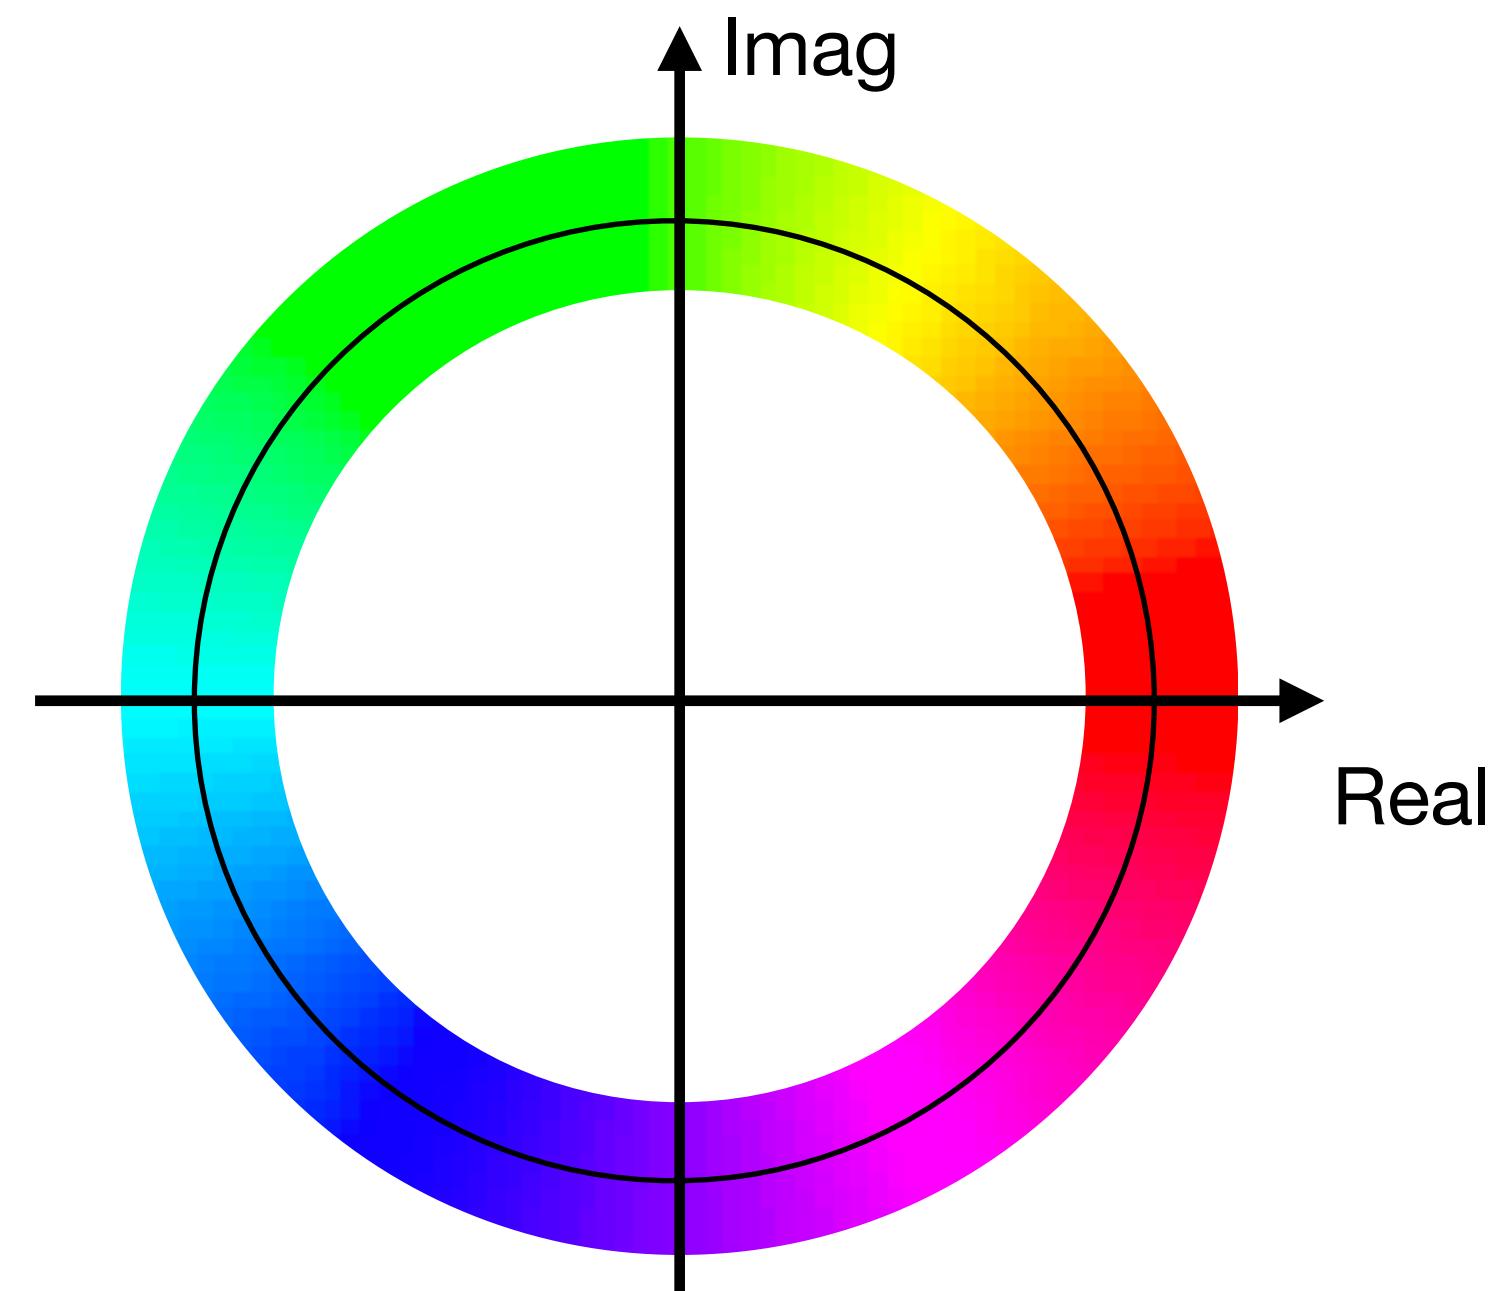
Visualizing the transform coefficients



Visualizing the transform coefficients

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

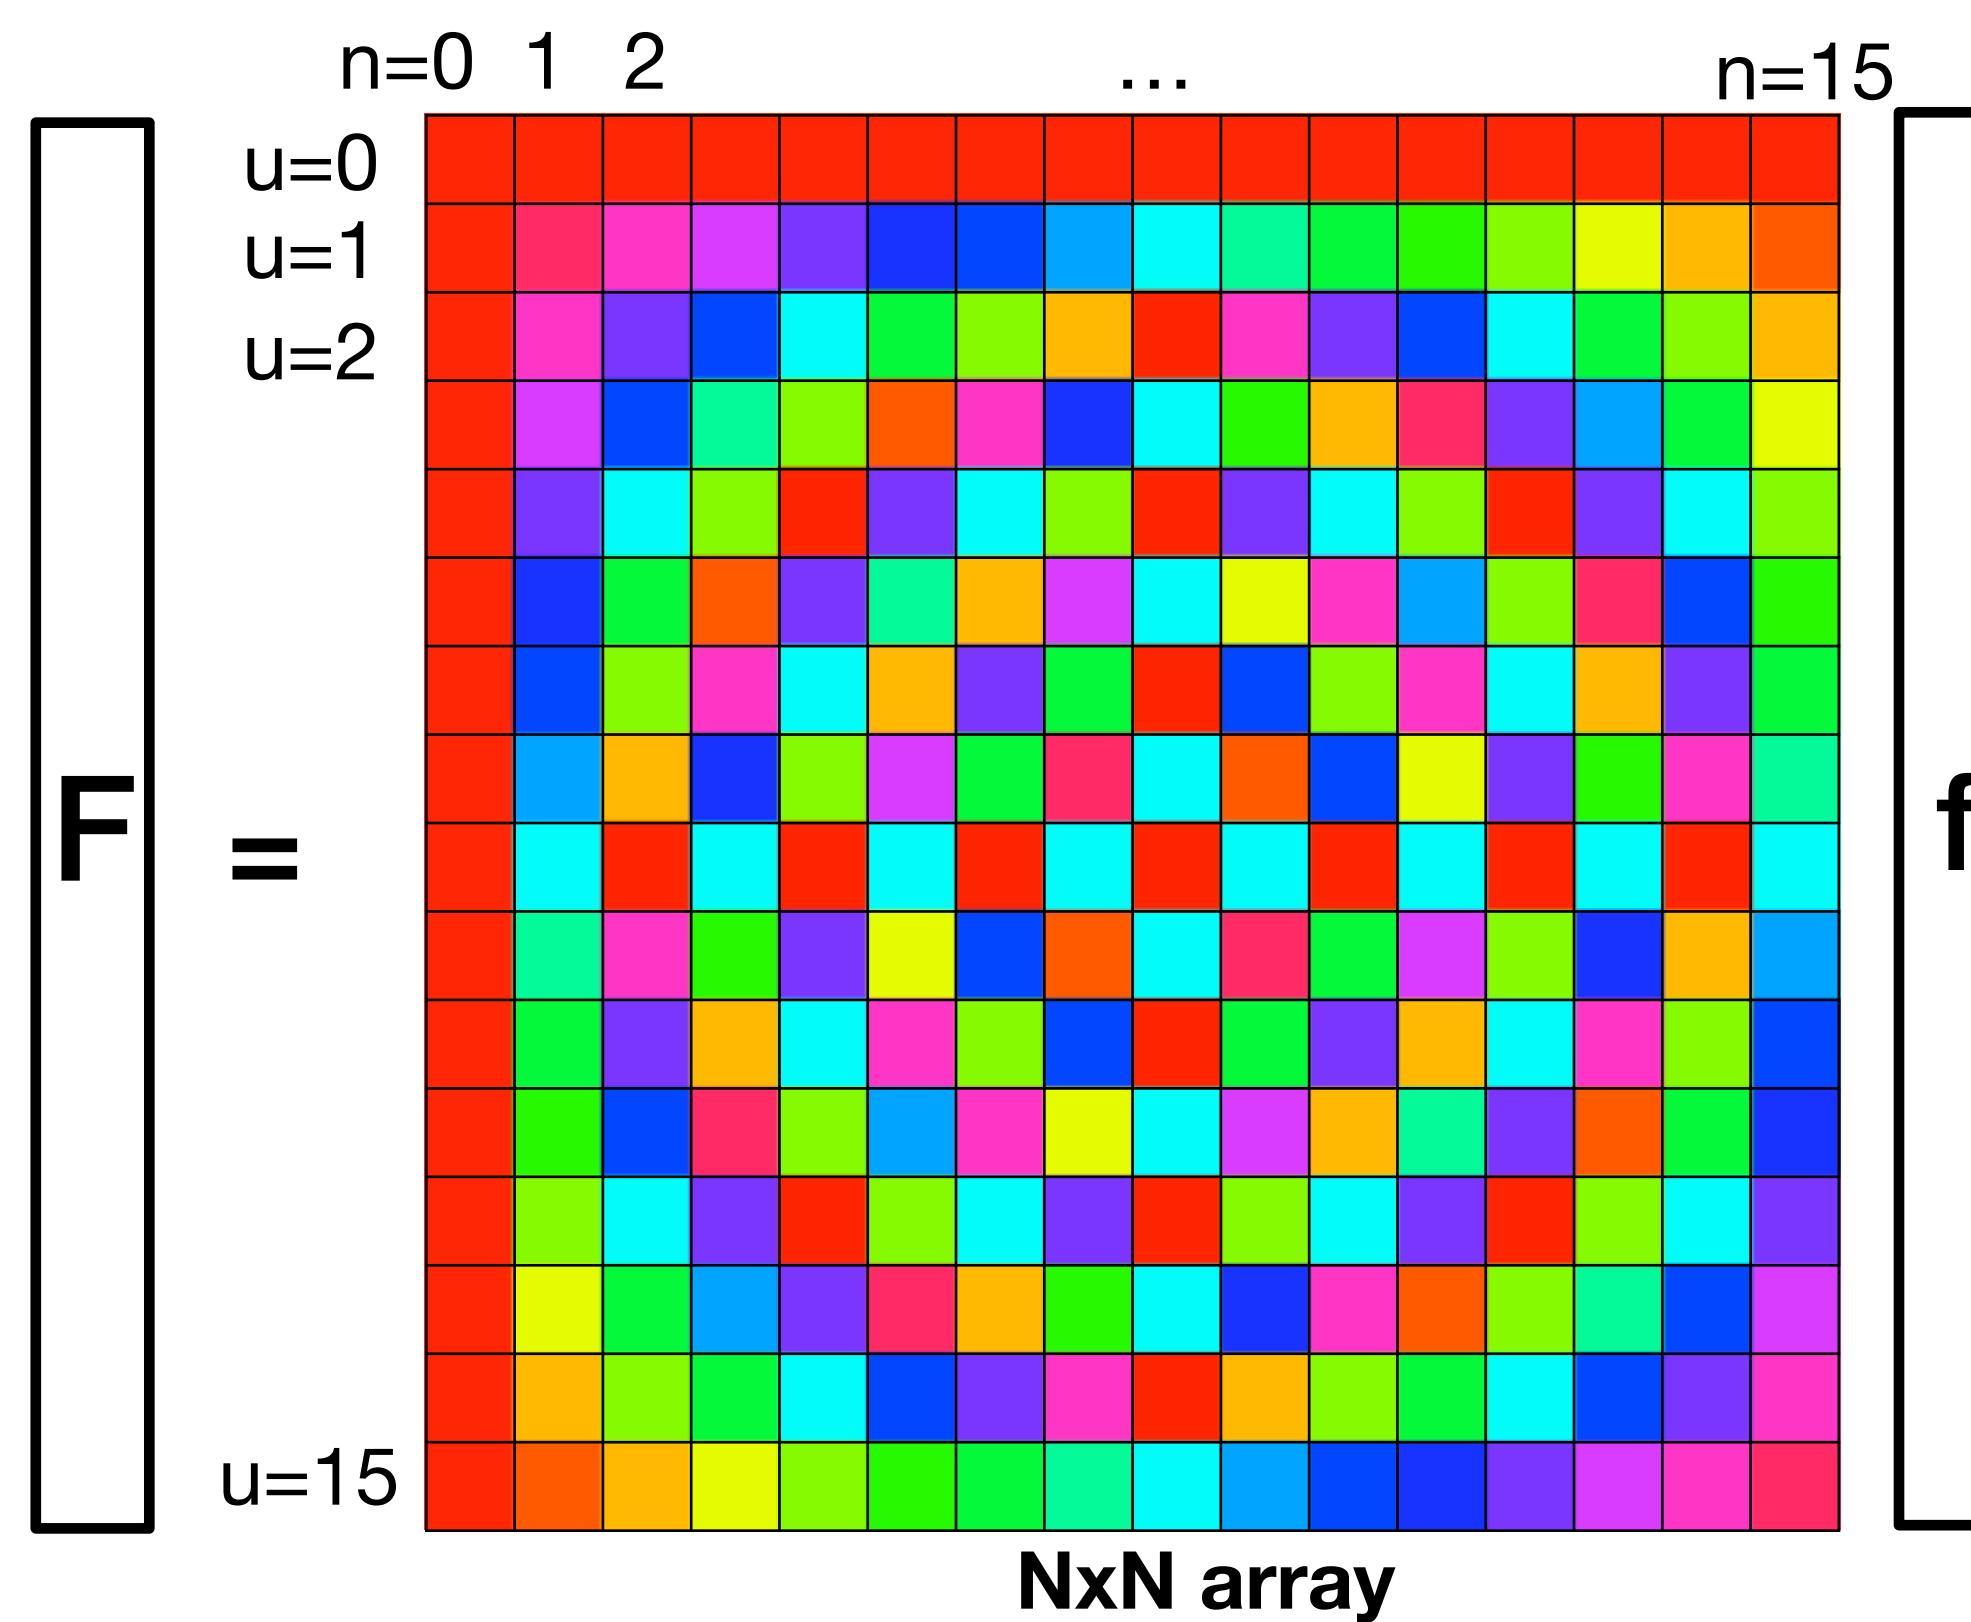
For N=16



The inverse of the Discrete Fourier transform

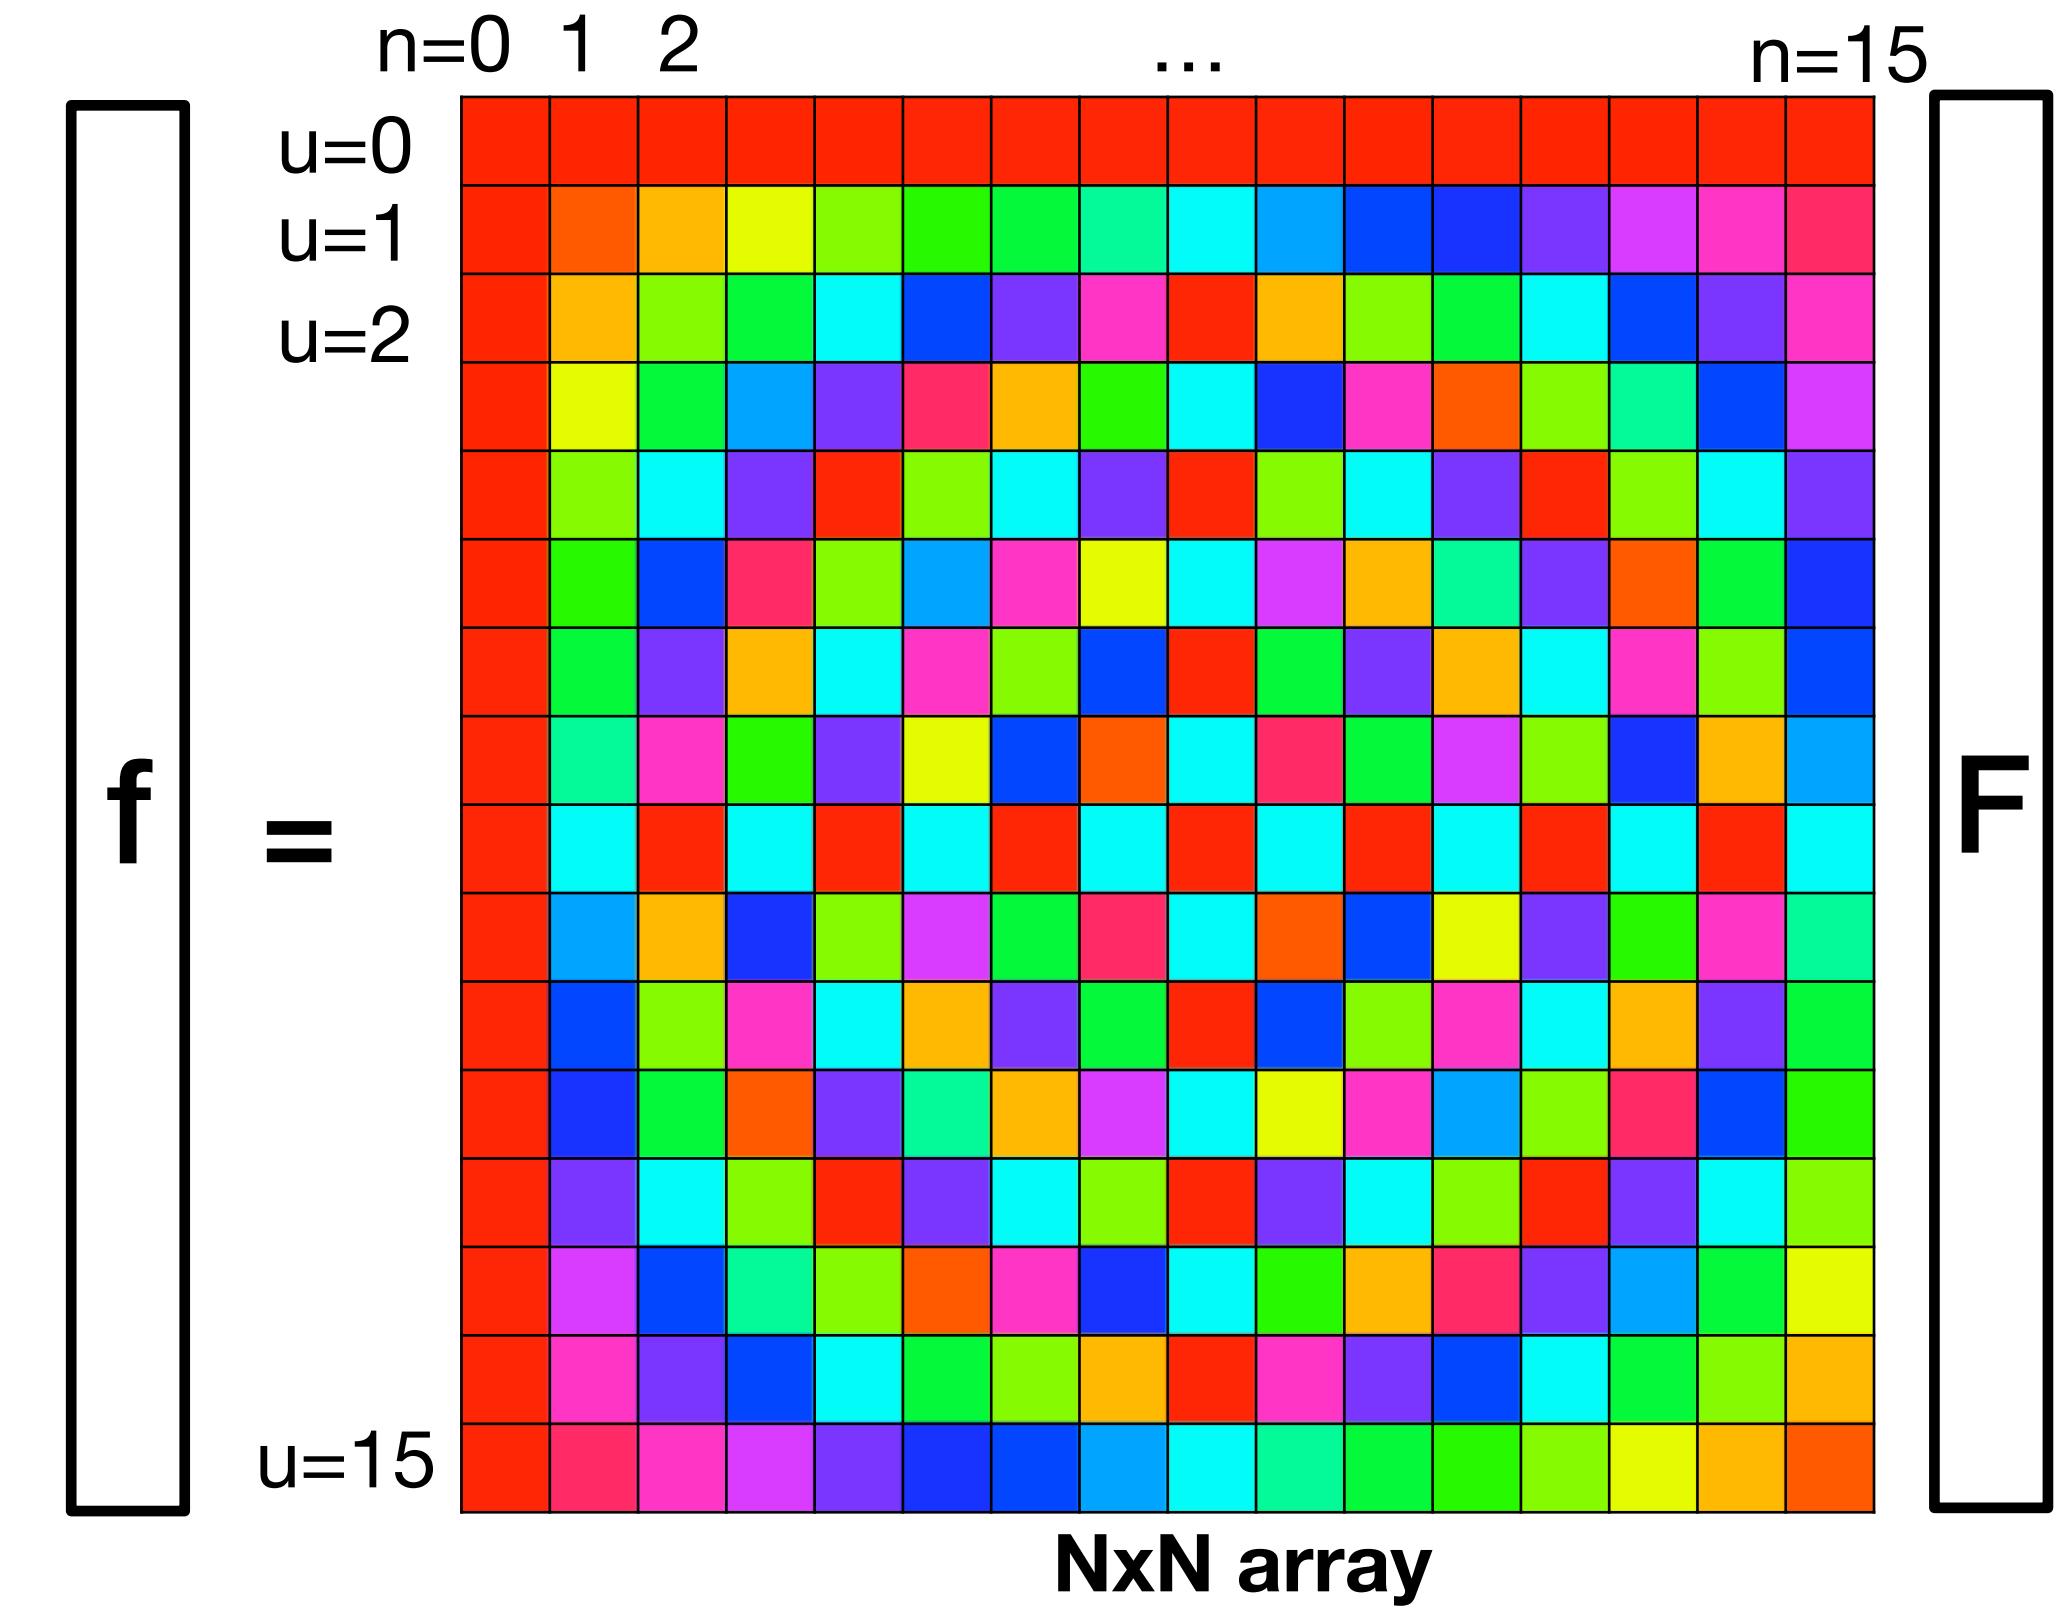
Discrete Fourier Transform (DFT):

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



Its inverse:

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



For images, the 2D DFT

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

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

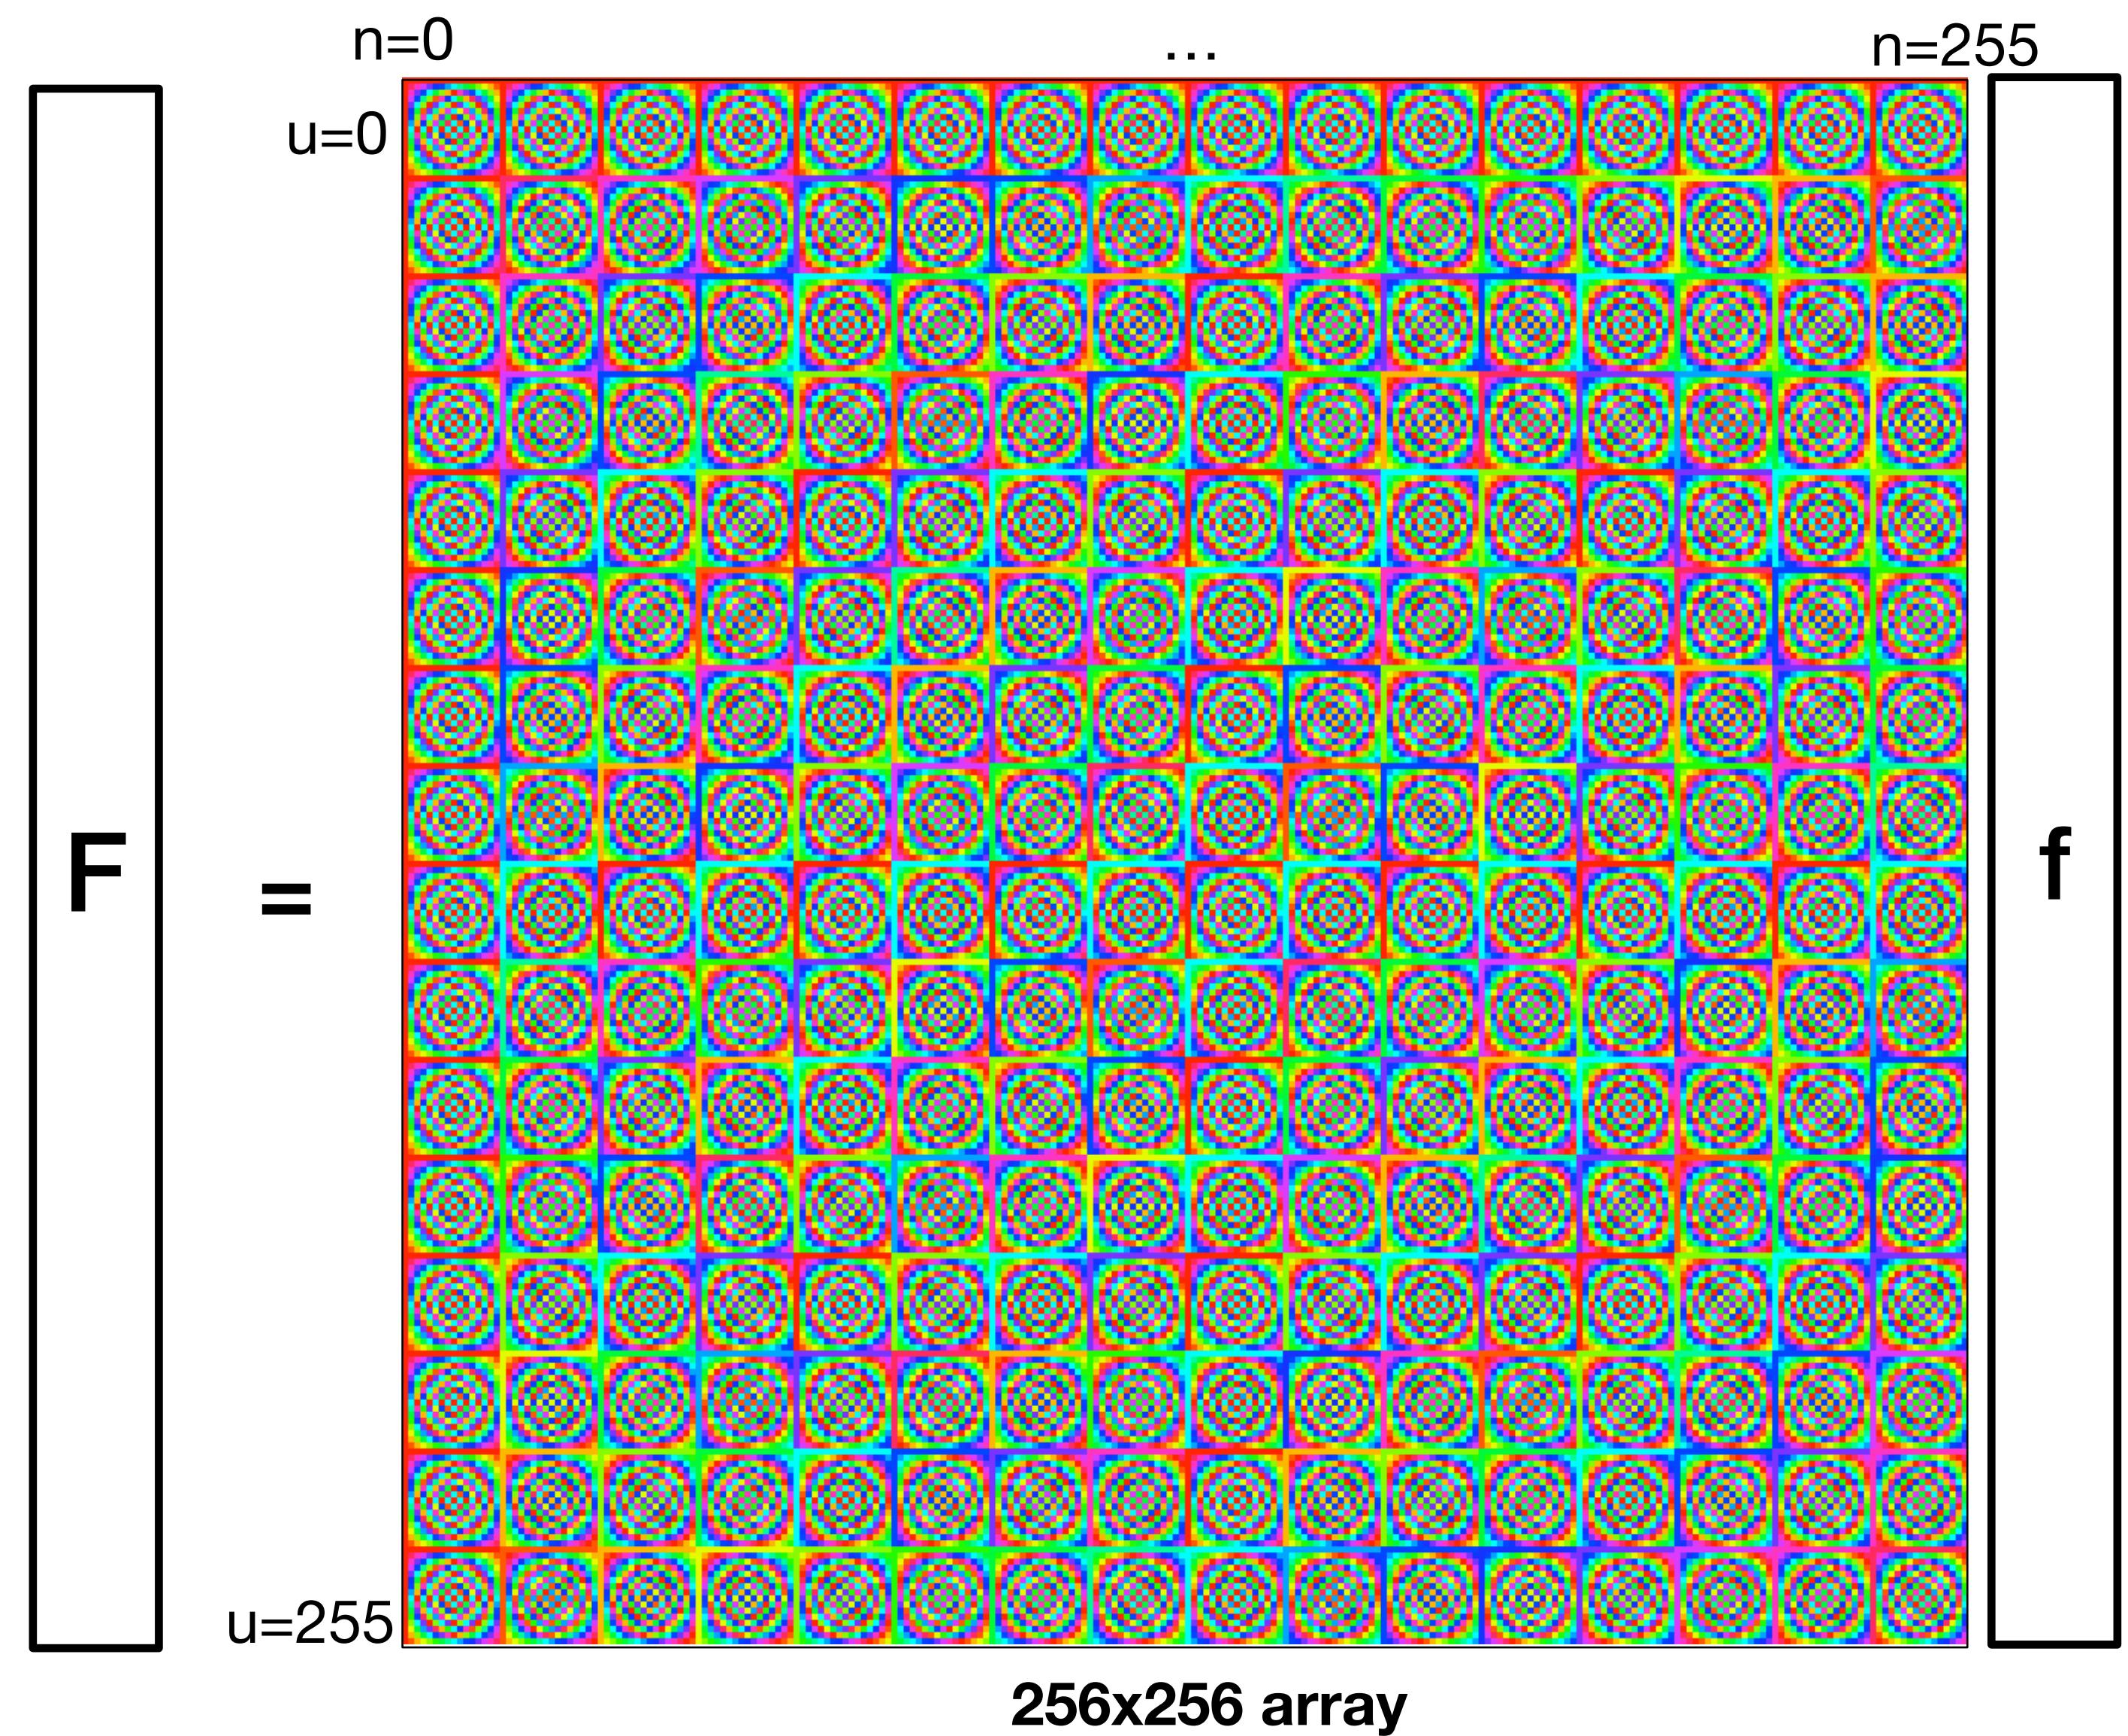
2D Discrete Fourier Transform (DFT) transforms an image $f[n,m]$ into $F[u,v]$ as:

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

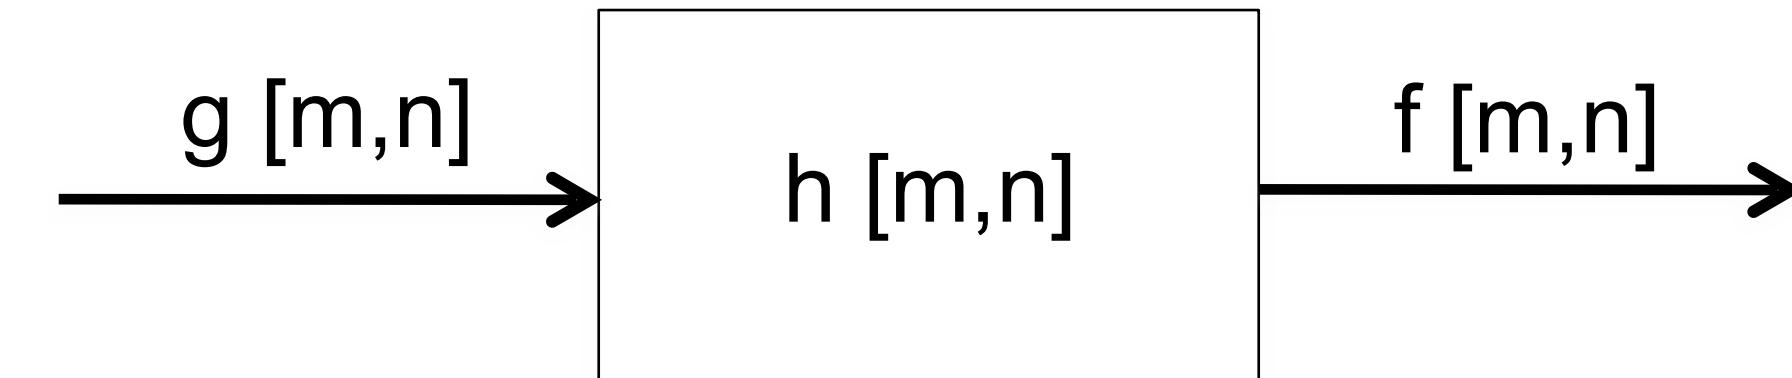
Visualizing the 2D DFT coefficients

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

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:

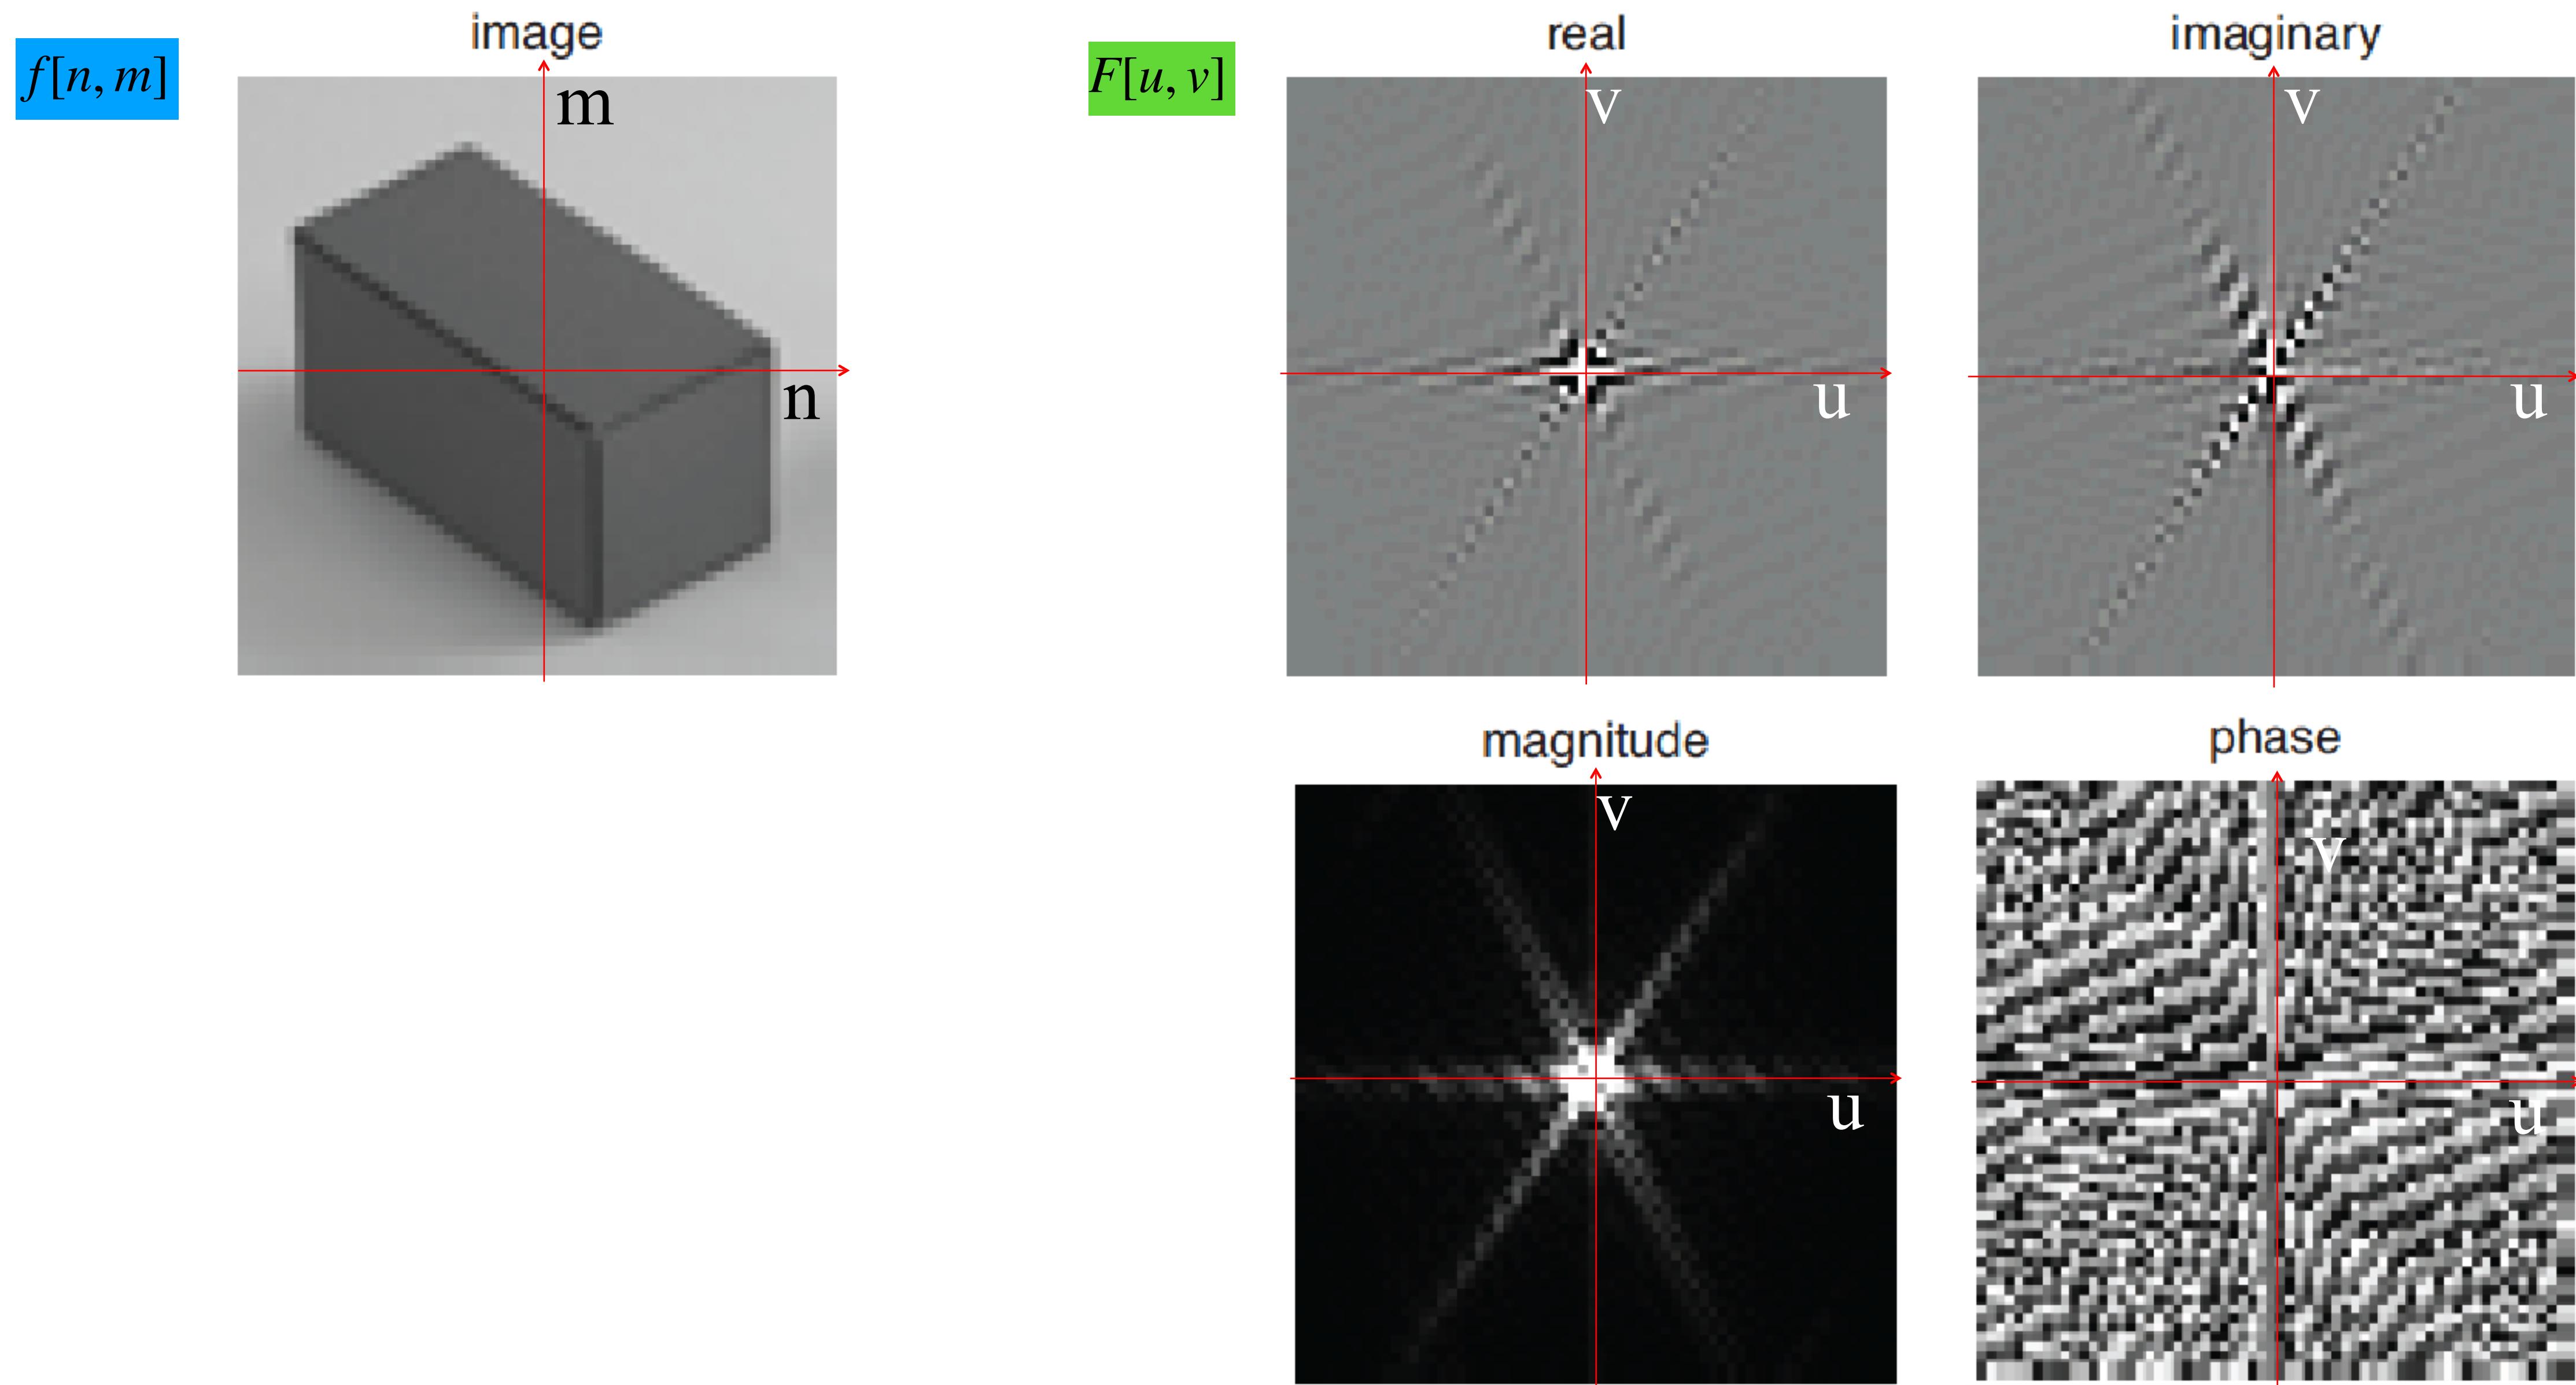
$$F[u, v] = Re\{F[u, v]\} + j Imag\{F[u, v]\}$$

Or using a polar decomposition:

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

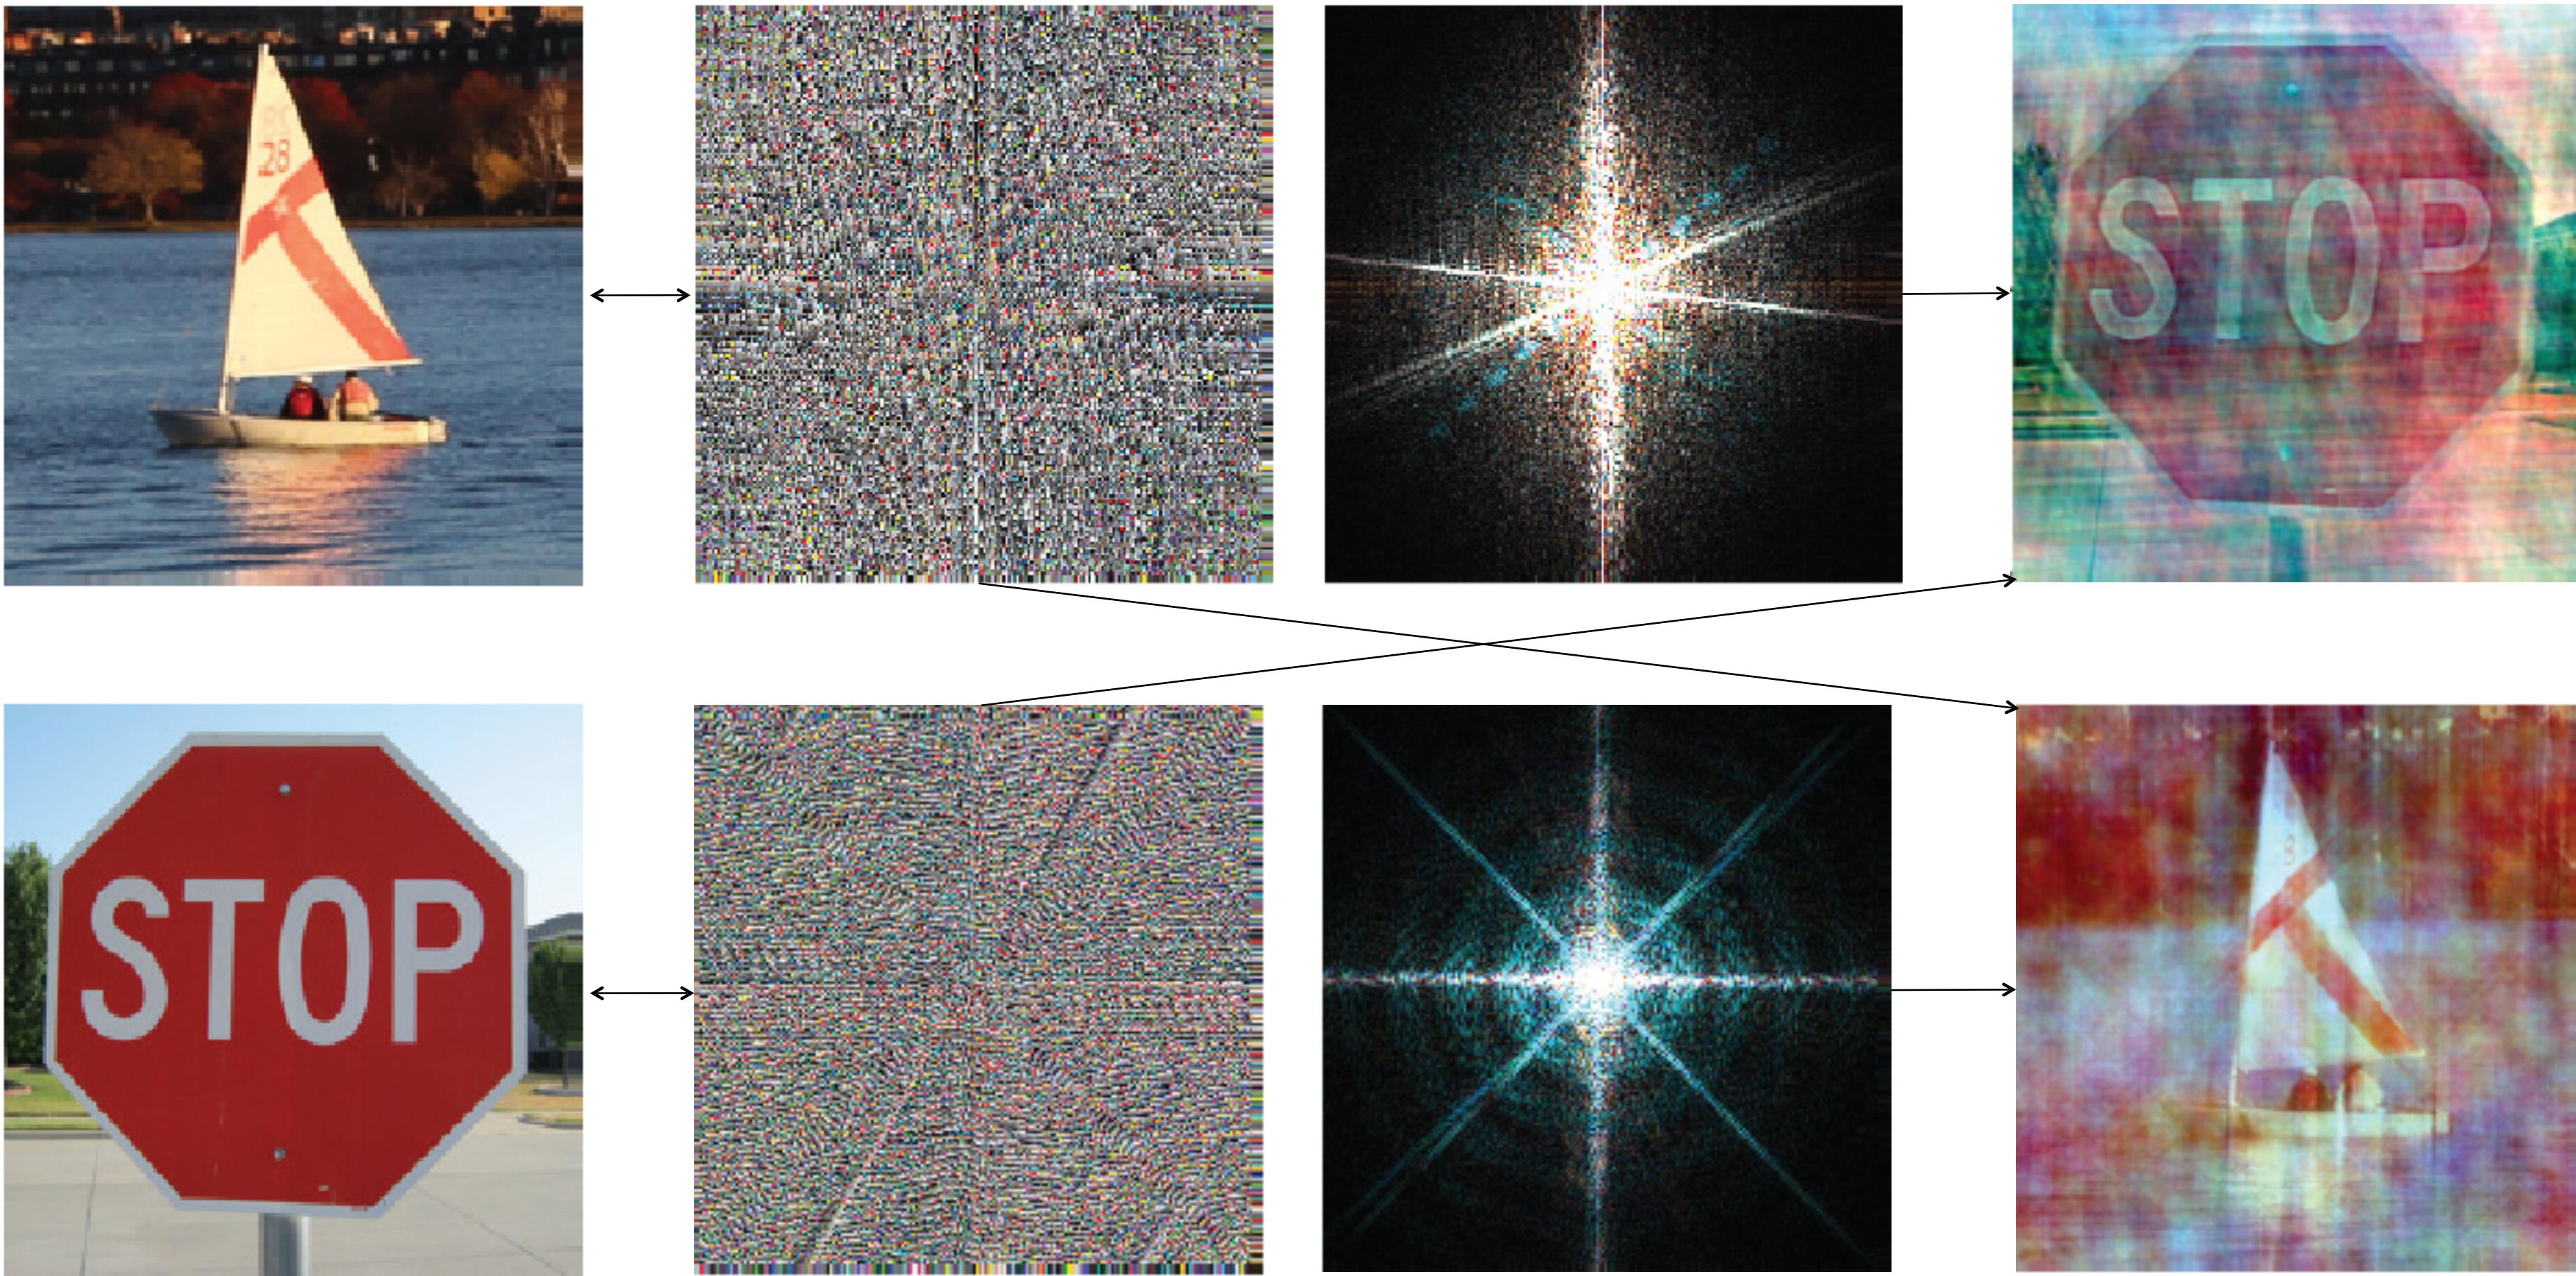
Amplitude Phase

Visualizing the image Fourier transform



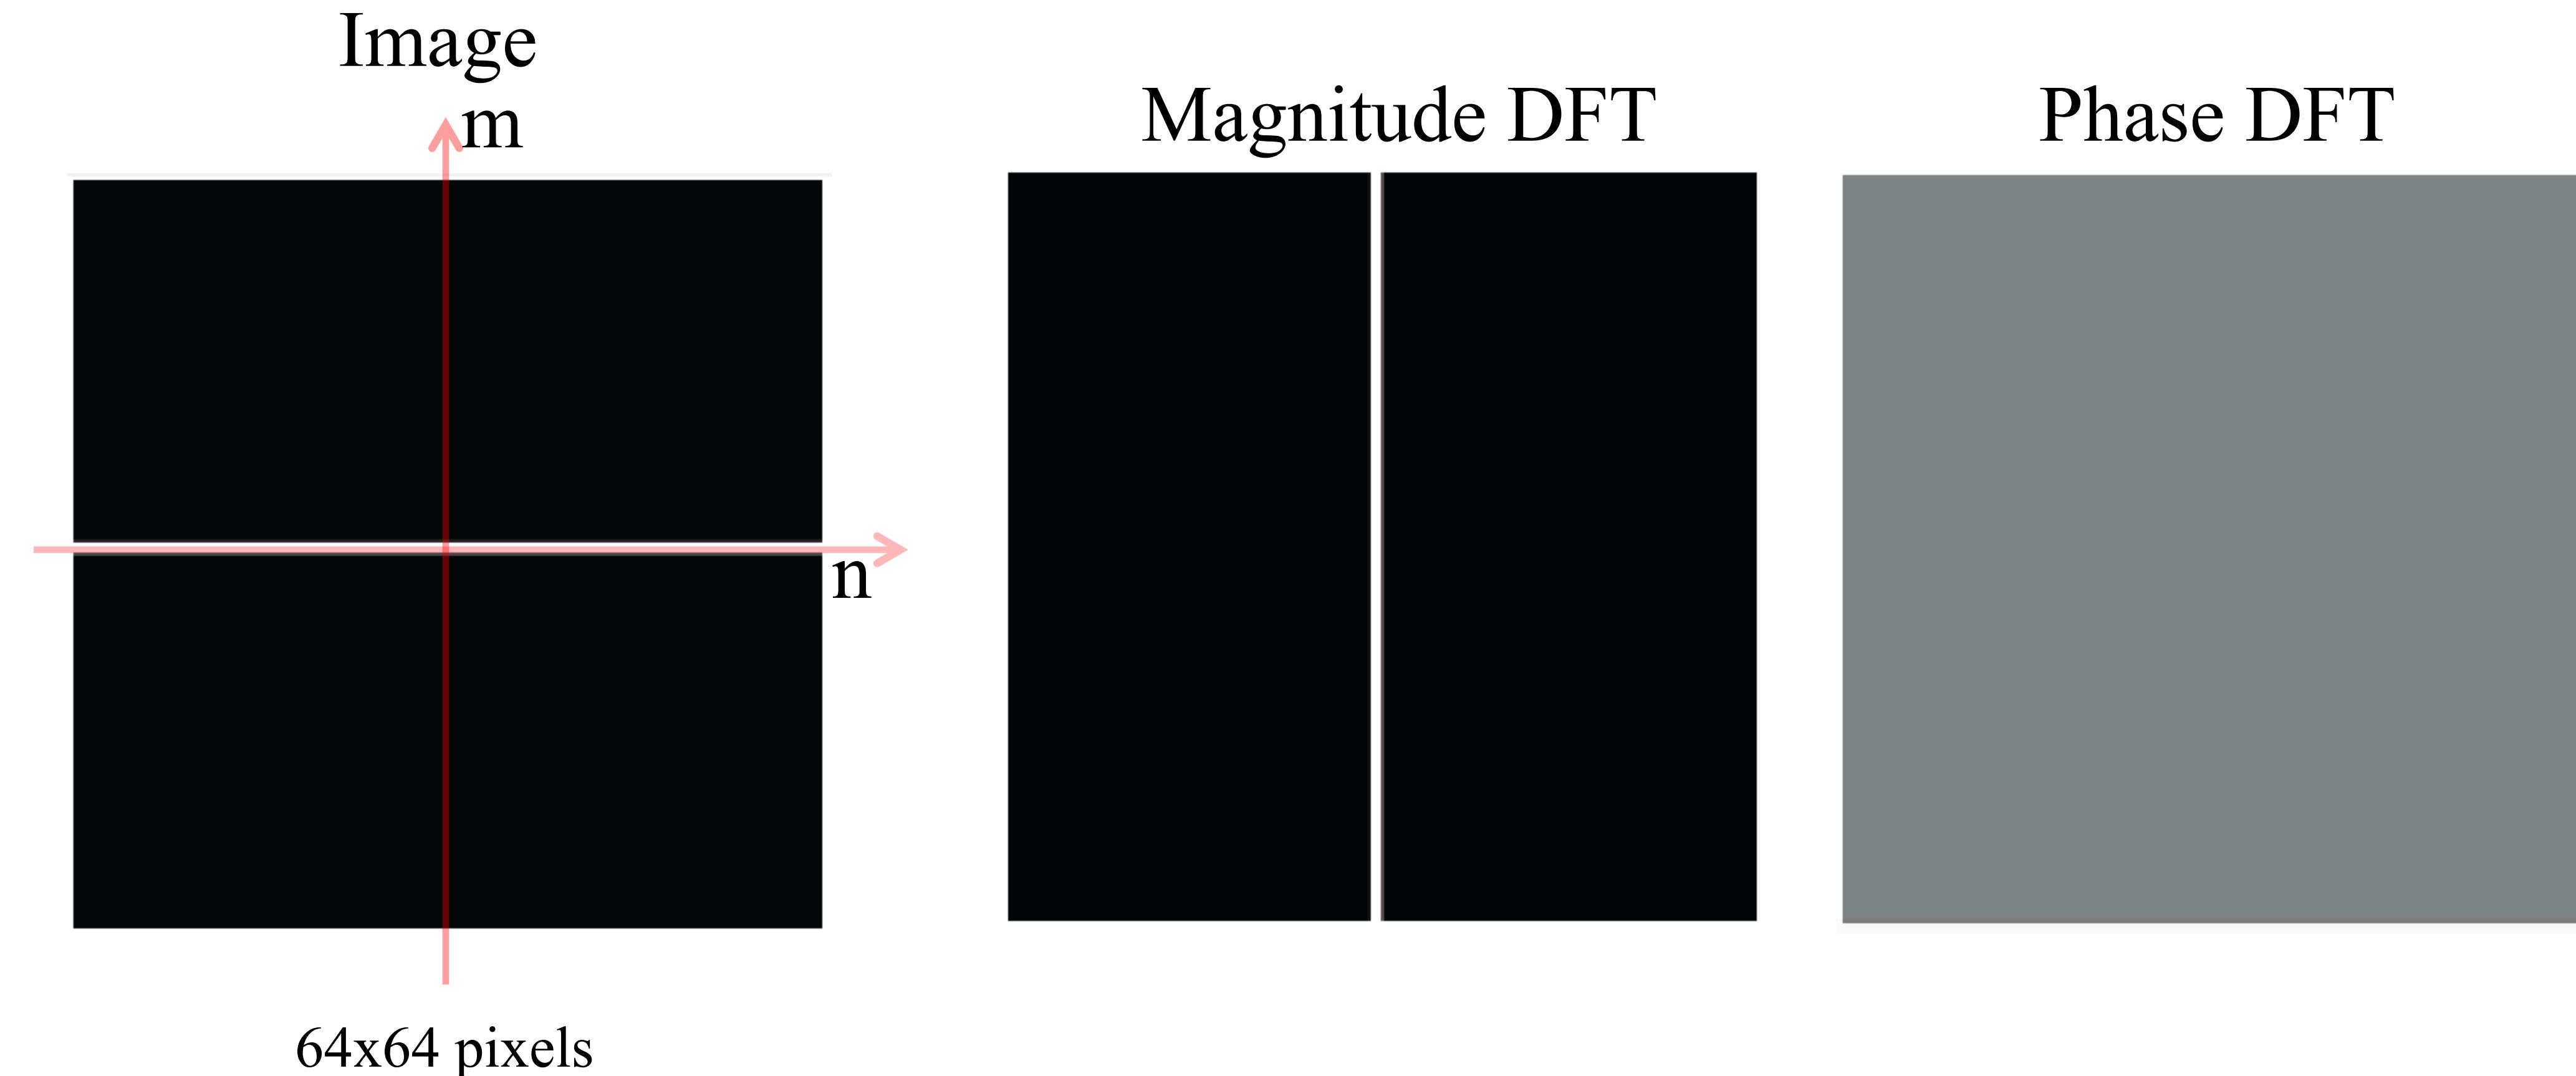
Phase and Magnitude

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



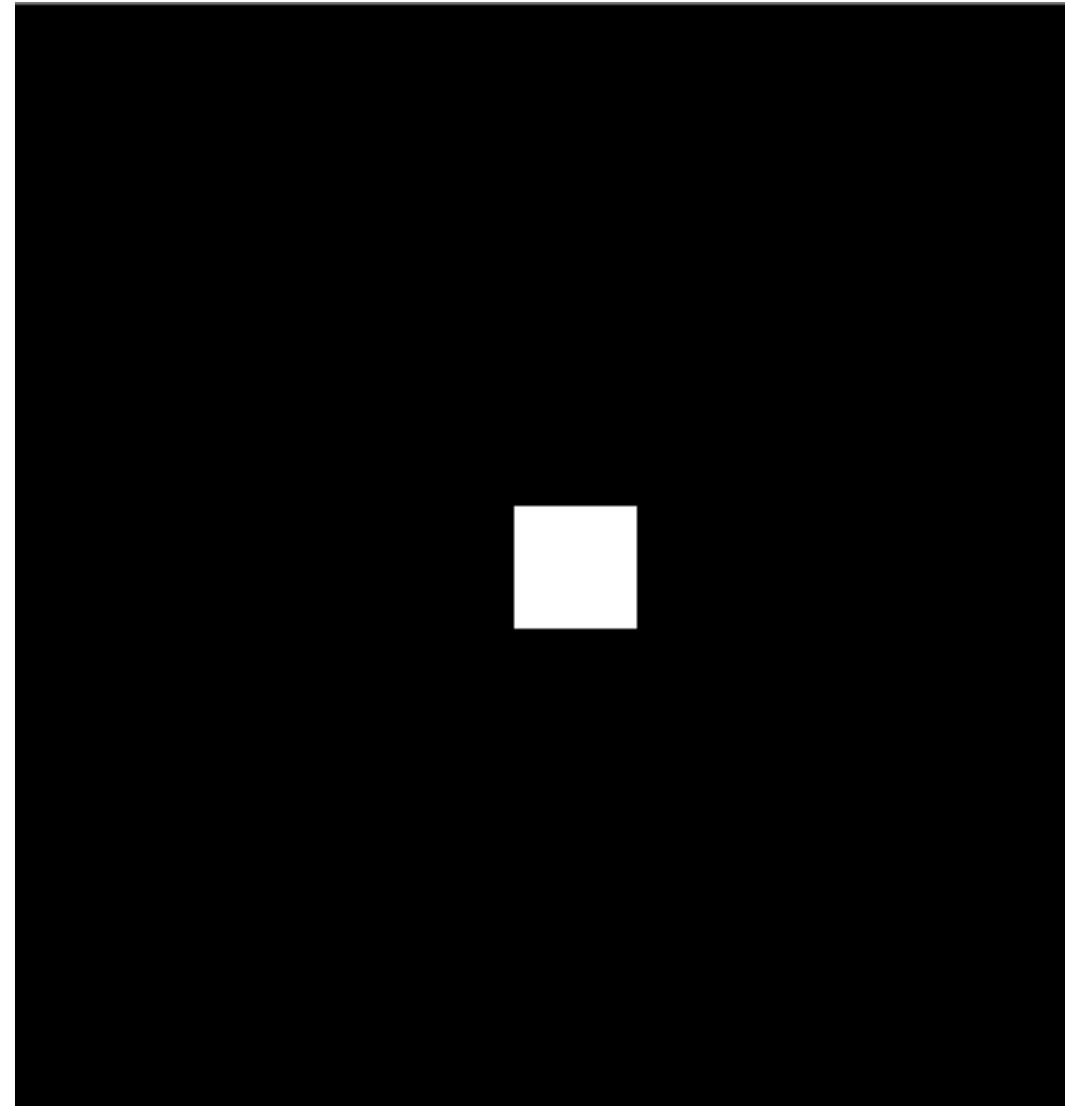
Each color channel is processed in the same way.

Some important Fourier transforms

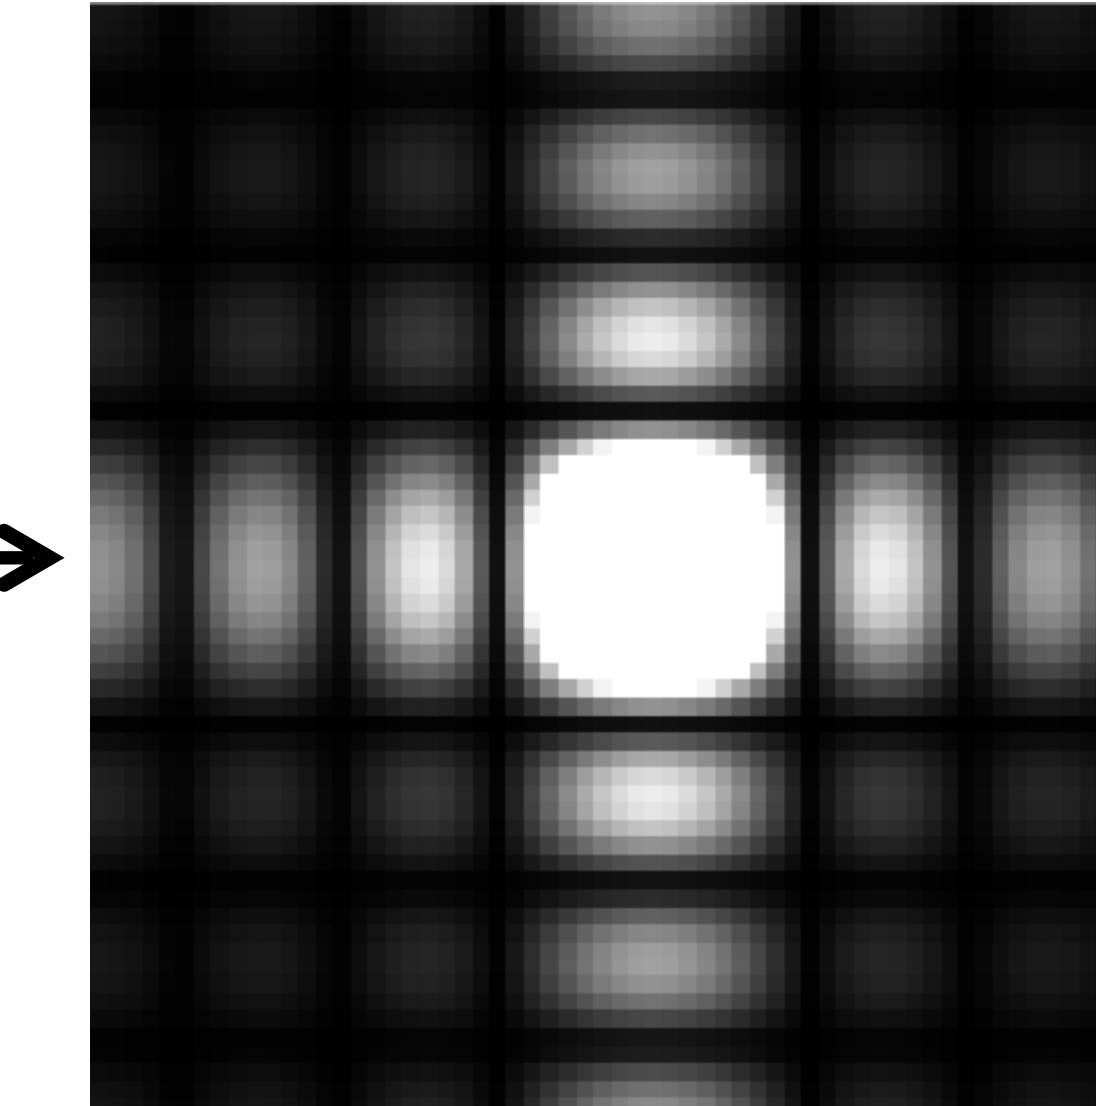


Some important Fourier transforms

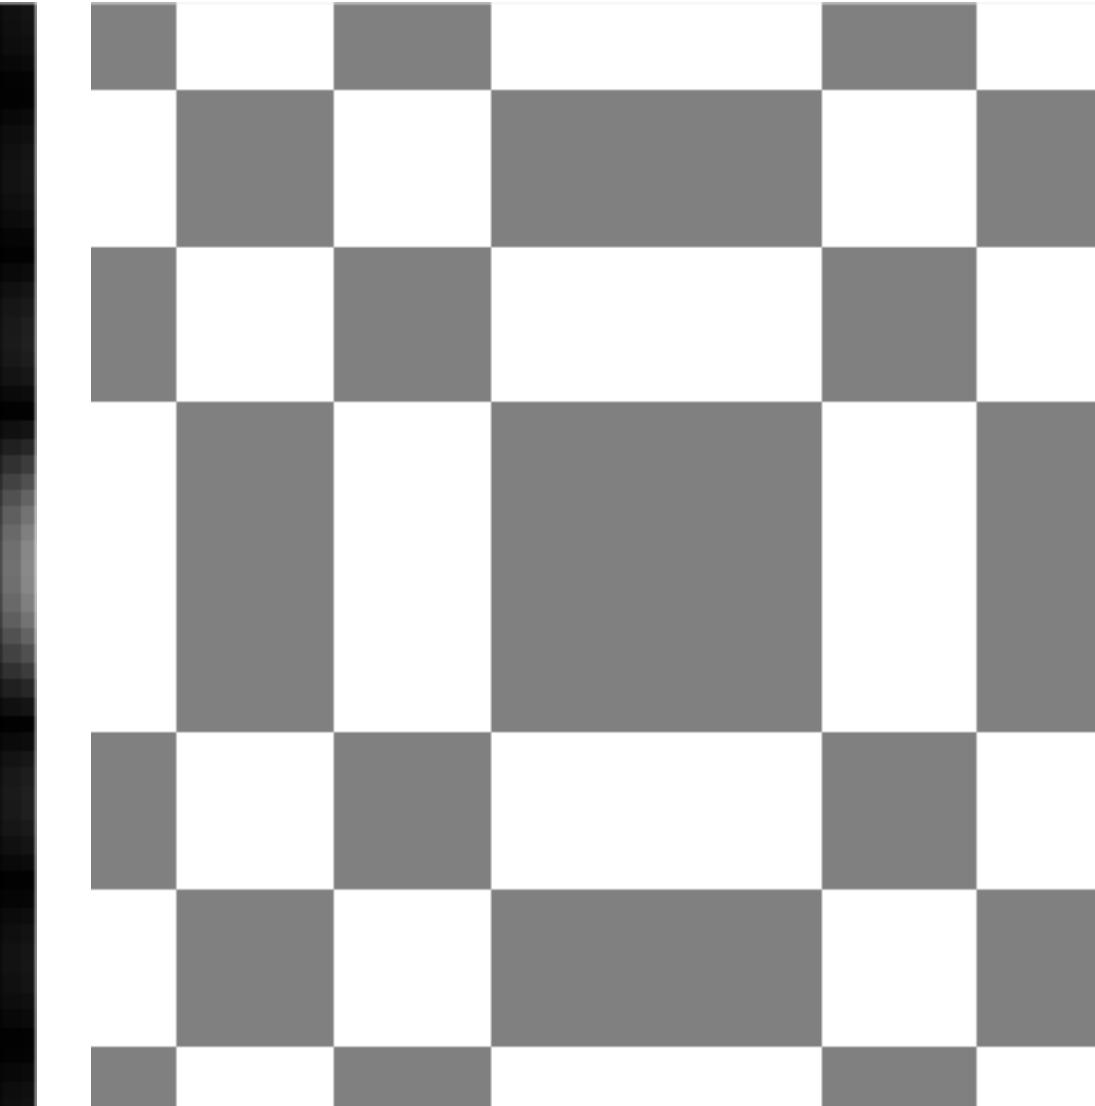
Image



Magnitude DFT

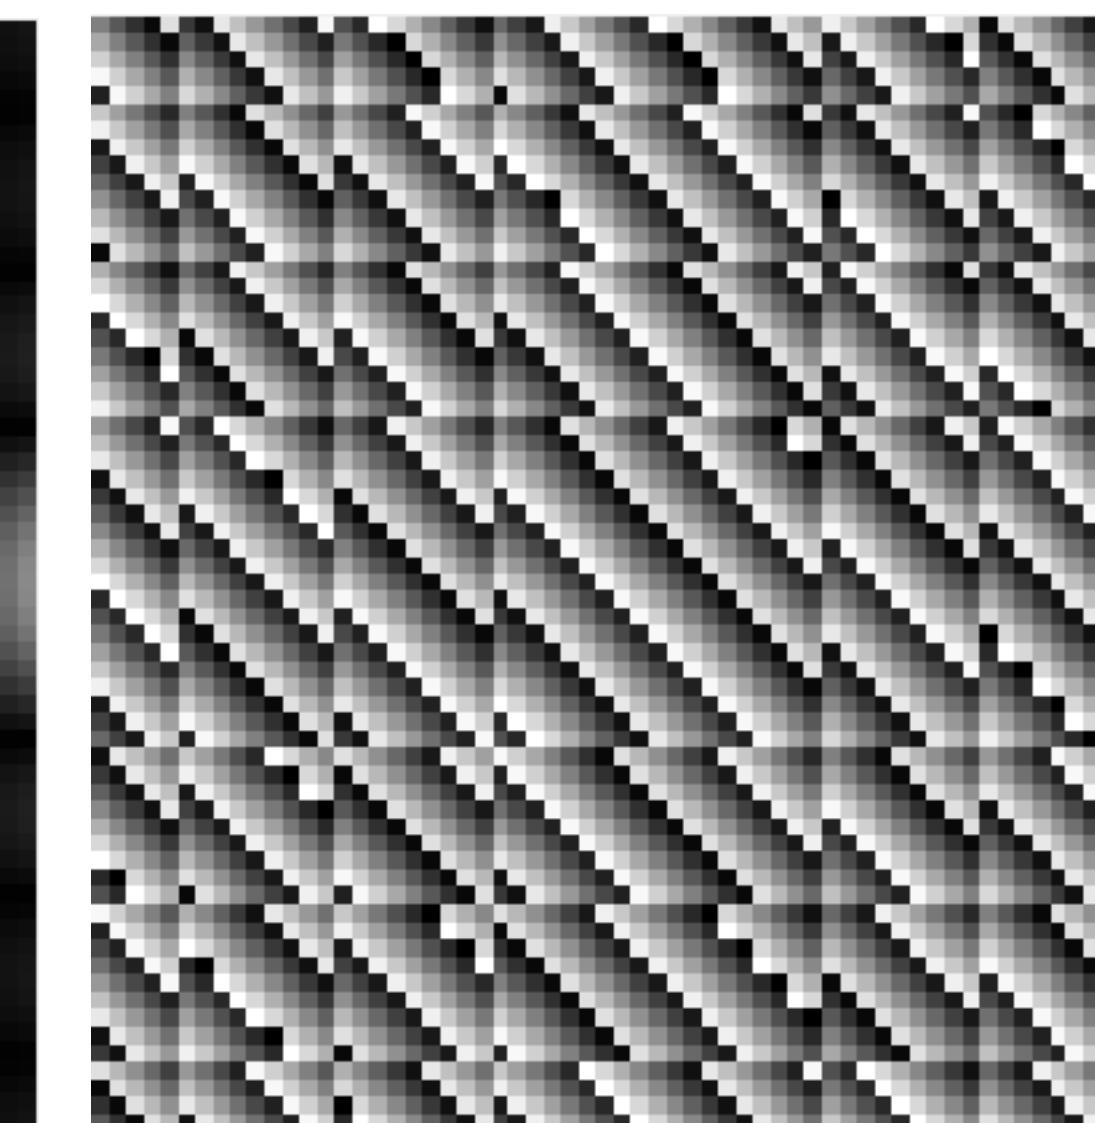
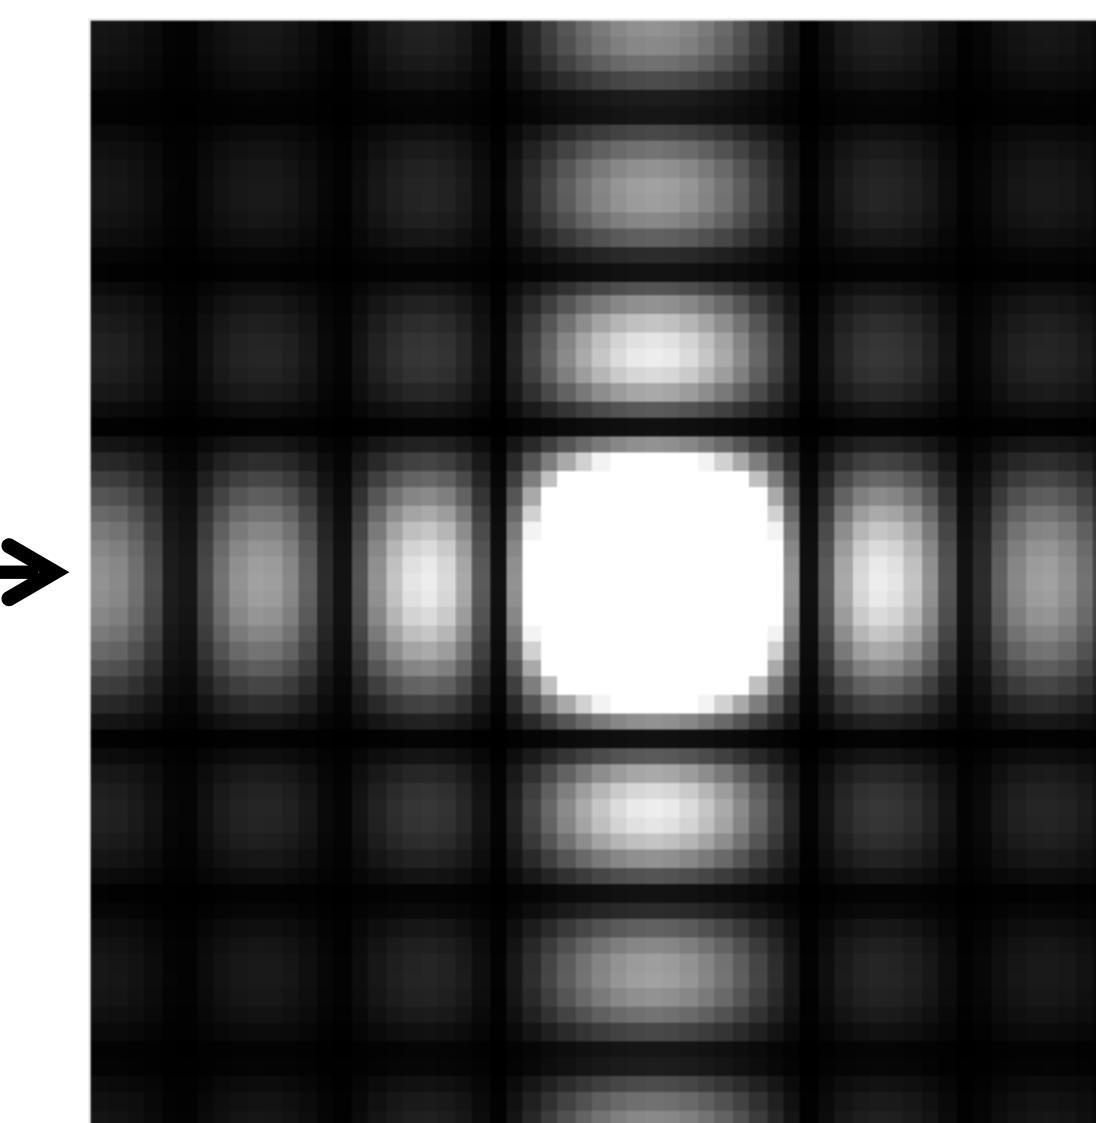
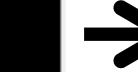
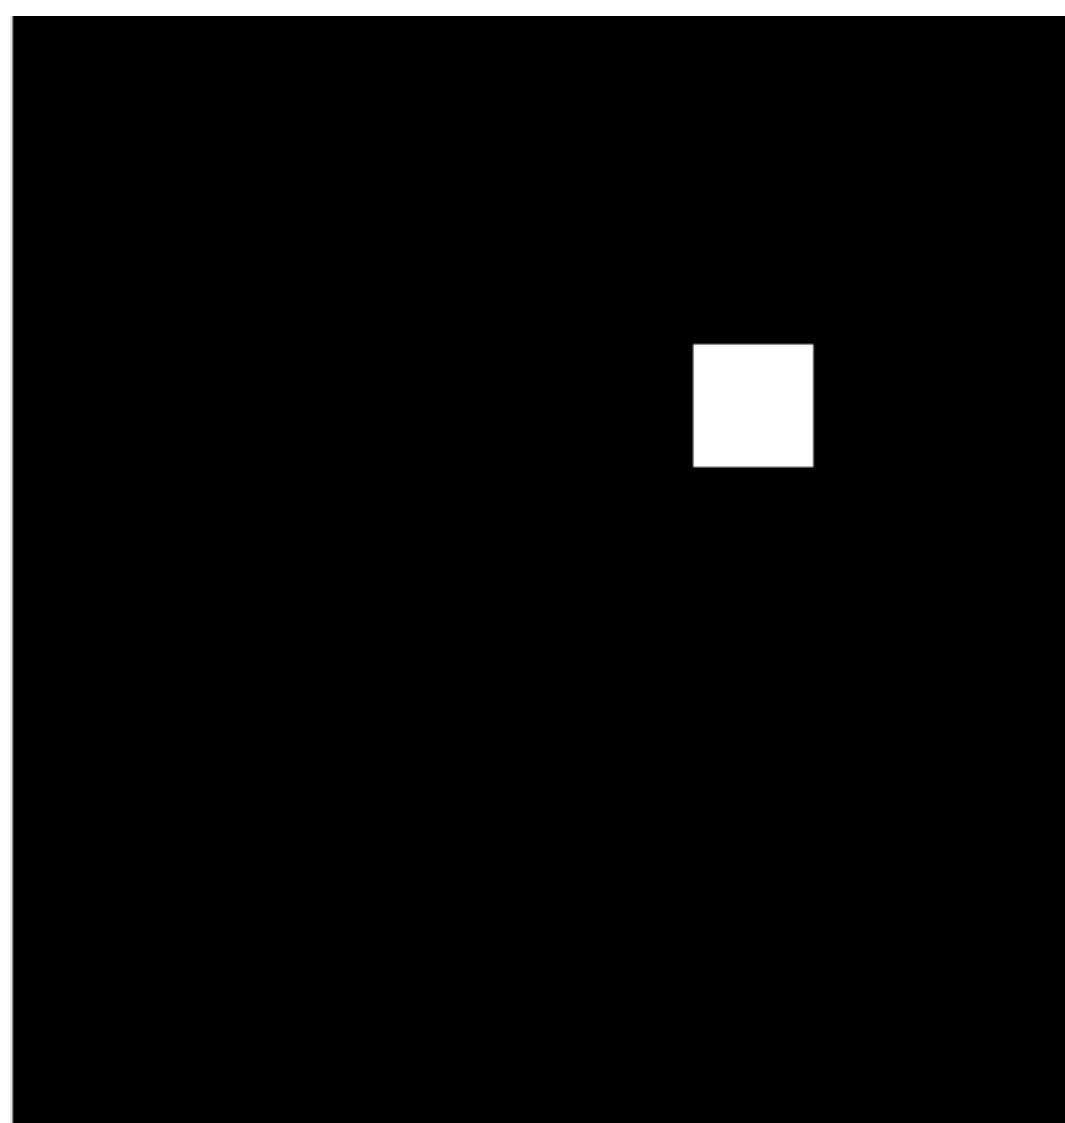


Phase DFT



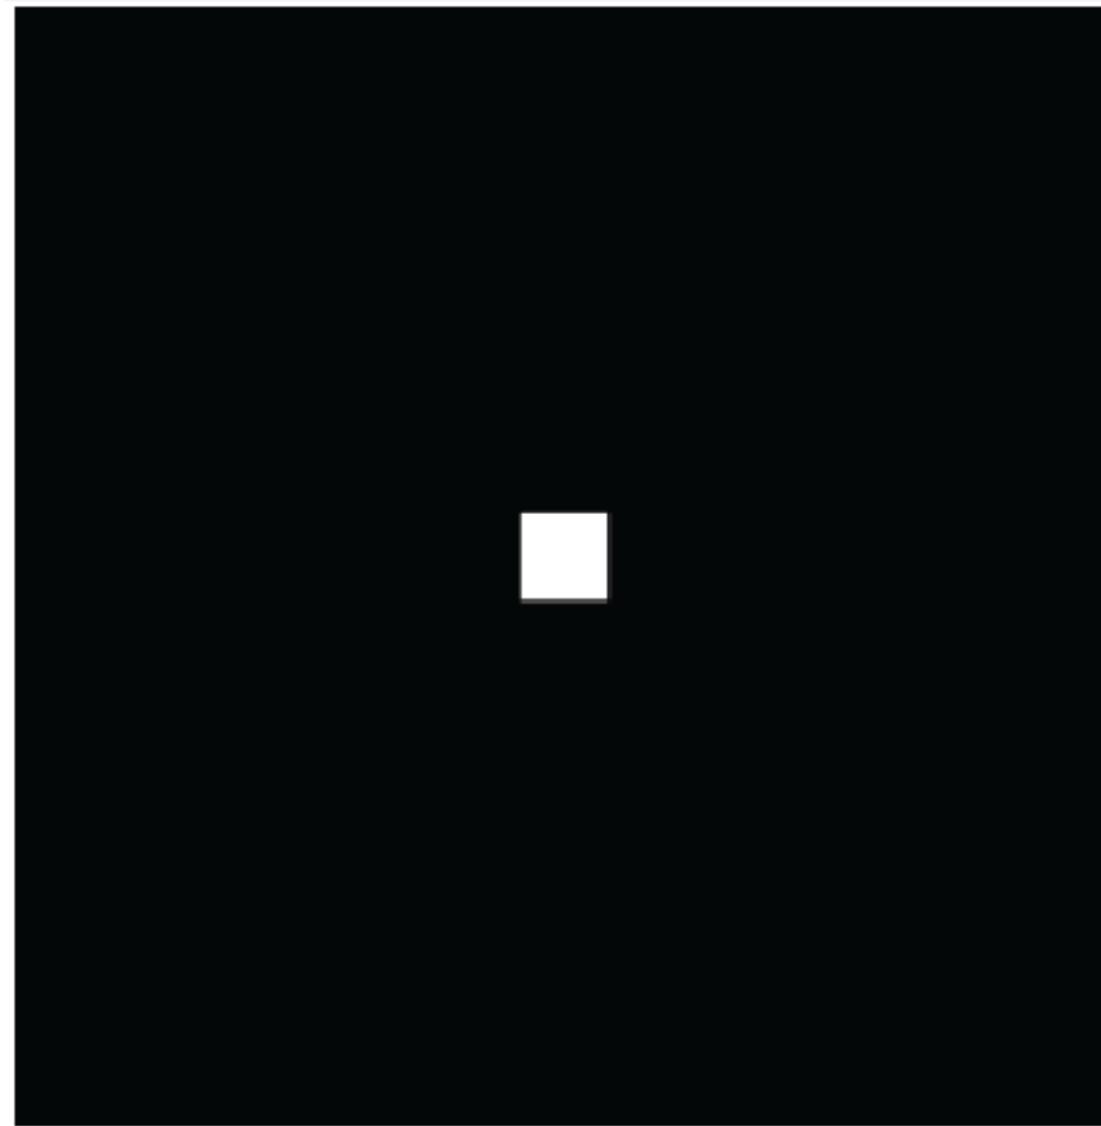
Translation

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

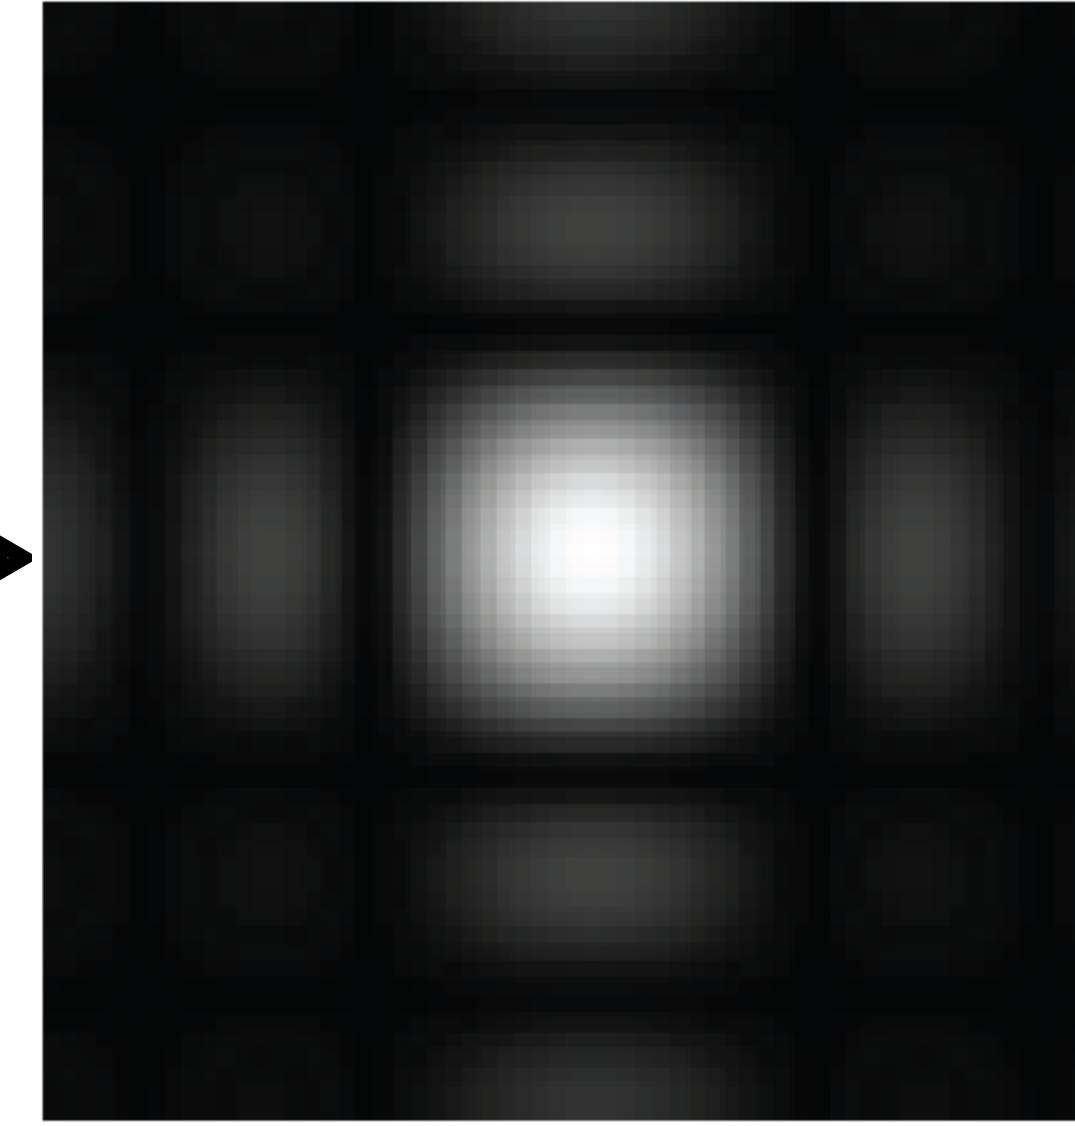


Some important Fourier transforms

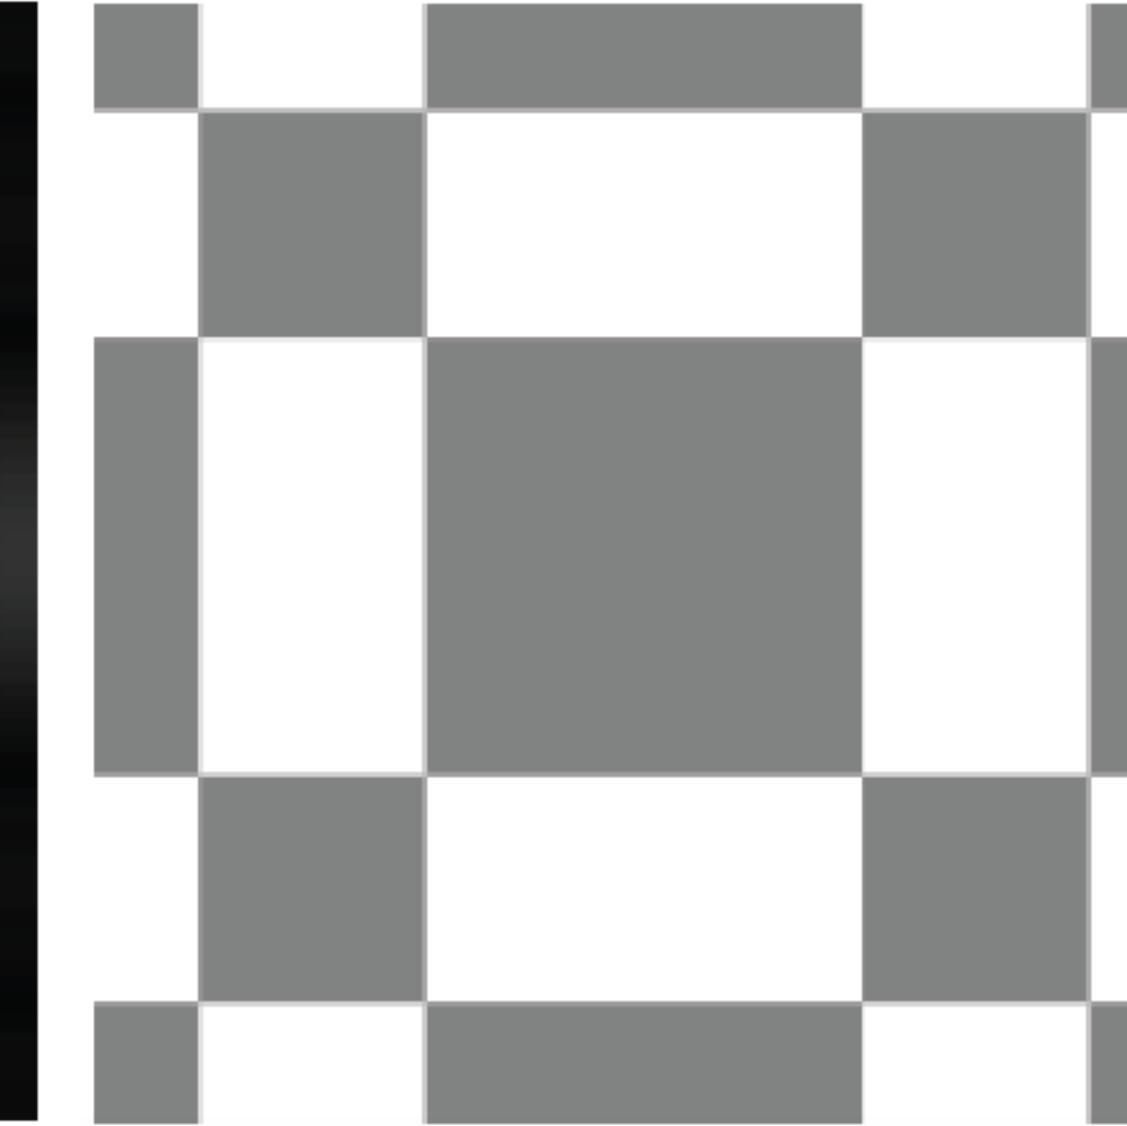
Image



Magnitude DFT

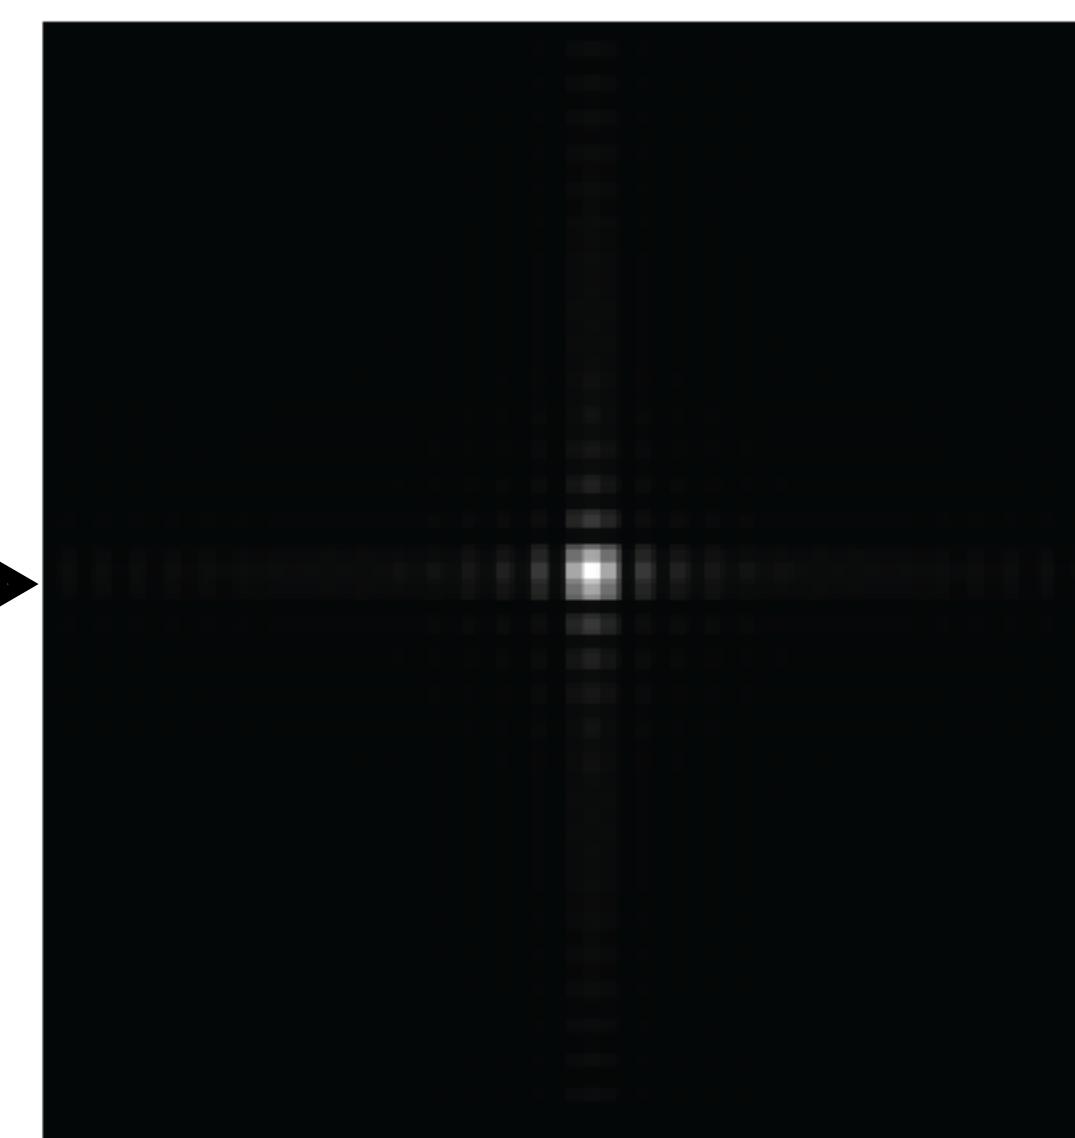
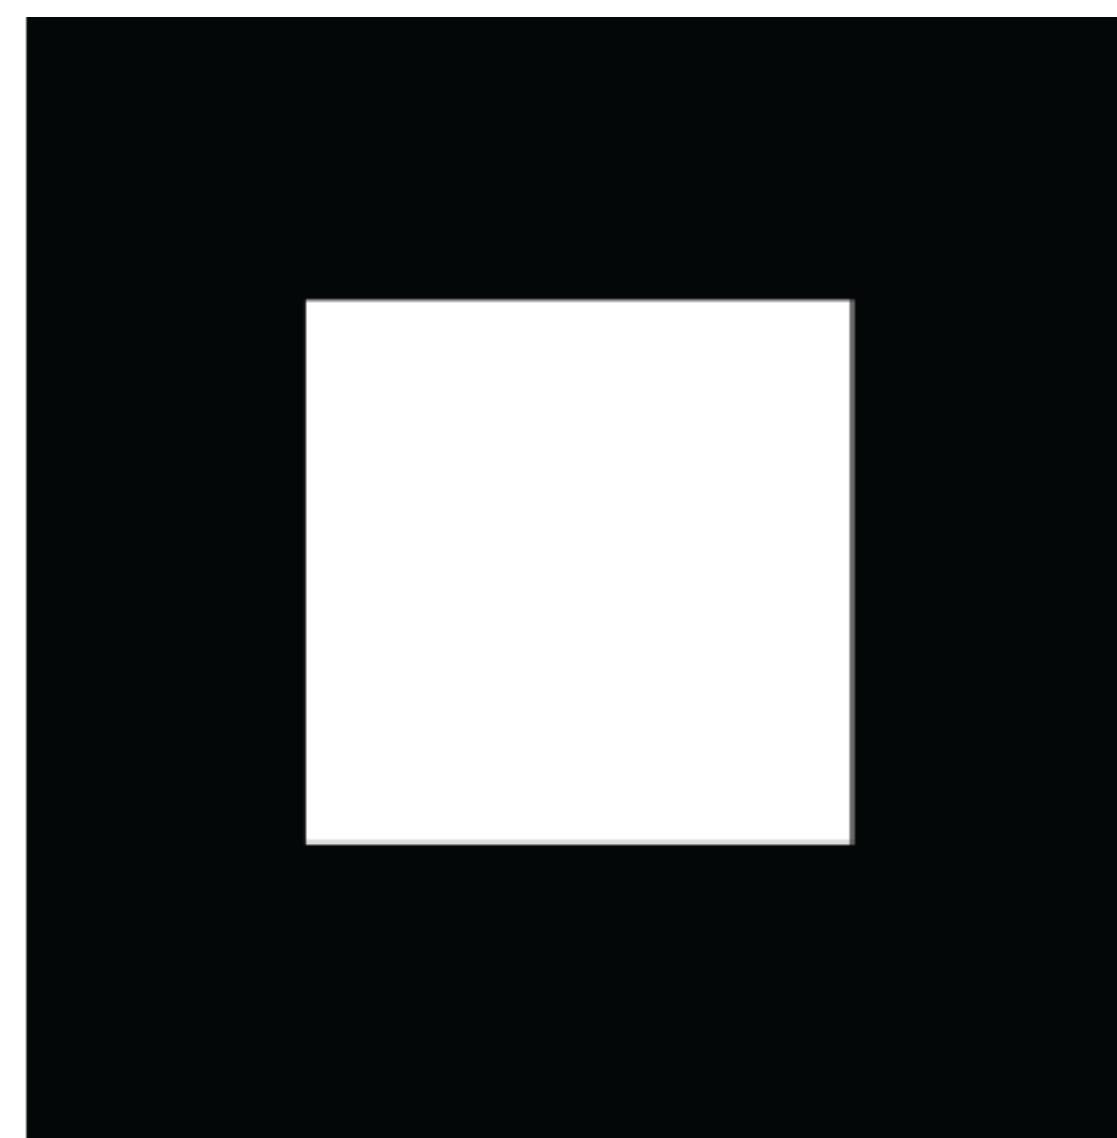


Phase DFT



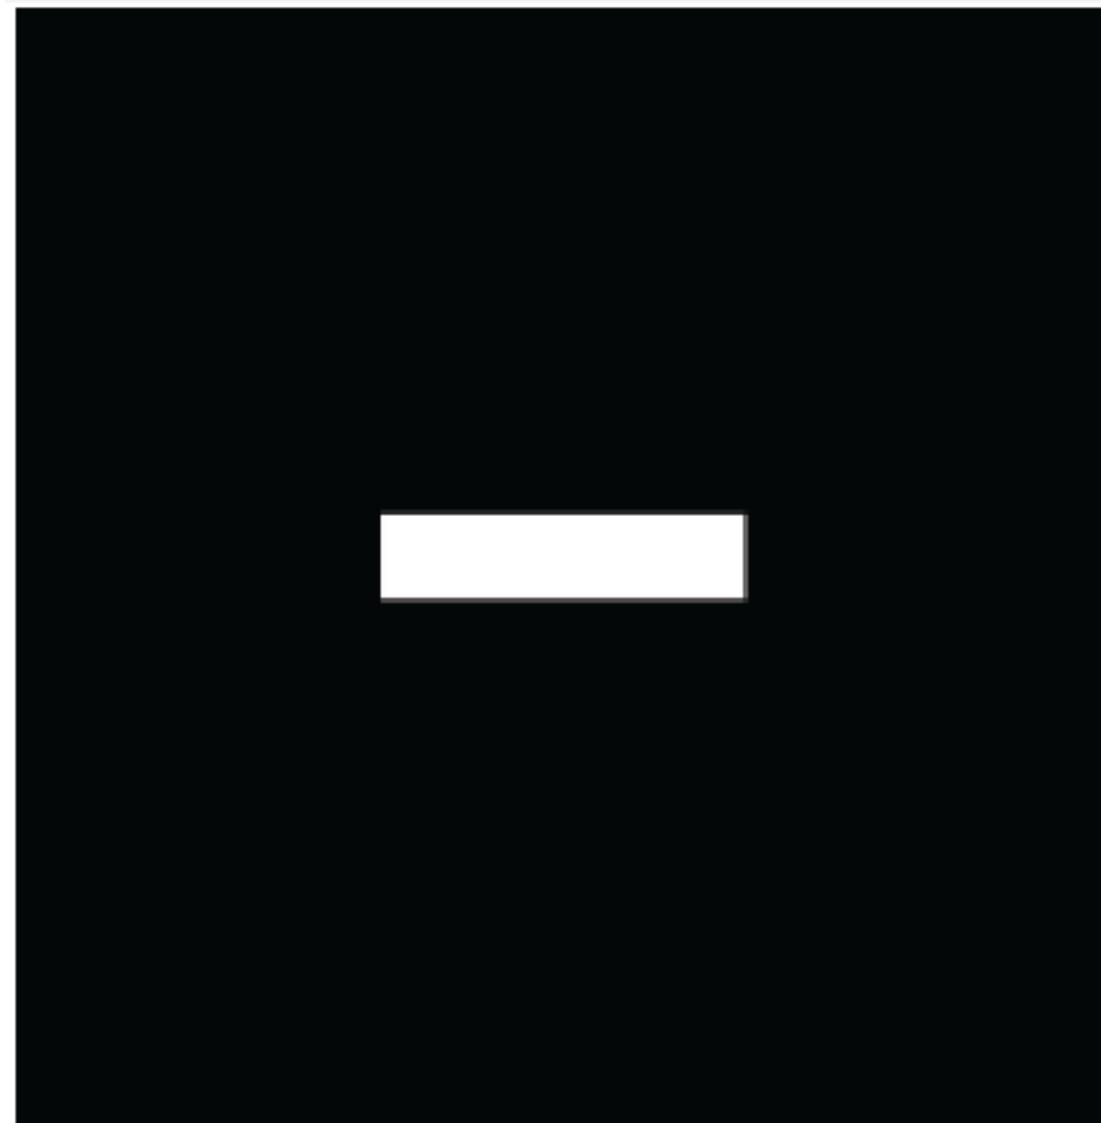
Scale

Small image details
produce content in high
spatial frequencies

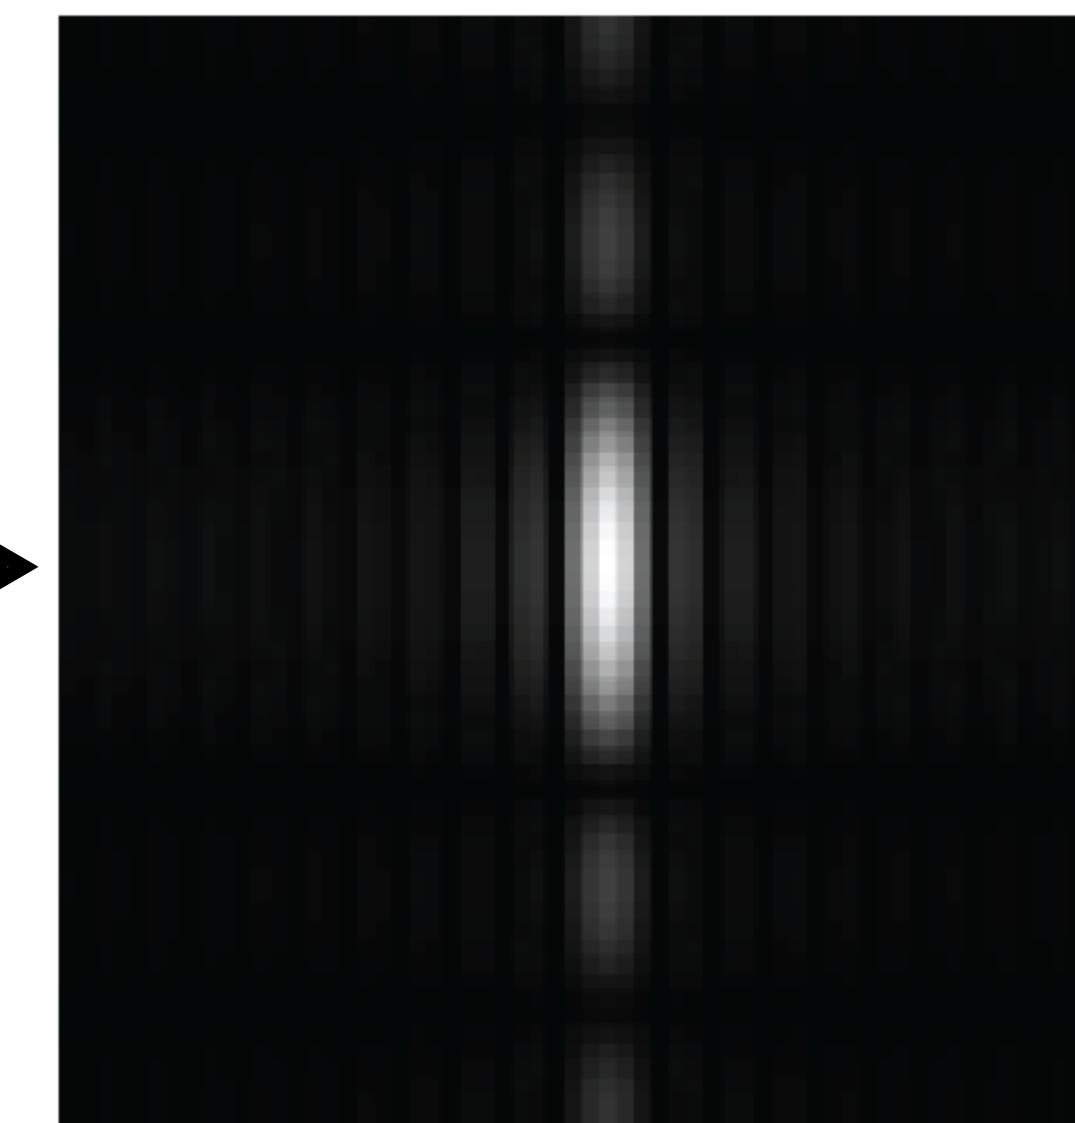


Some important Fourier transforms

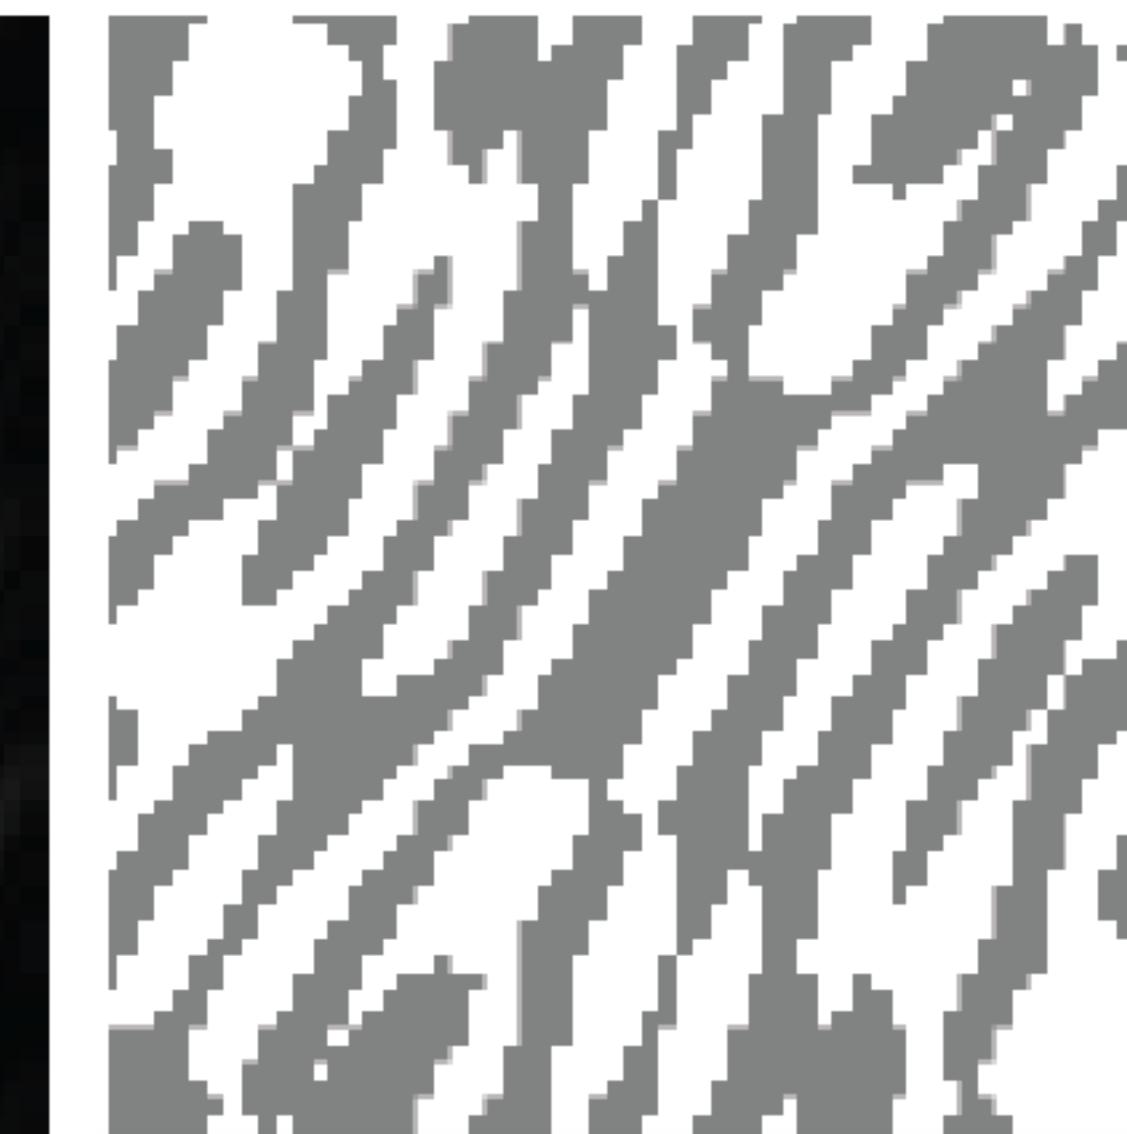
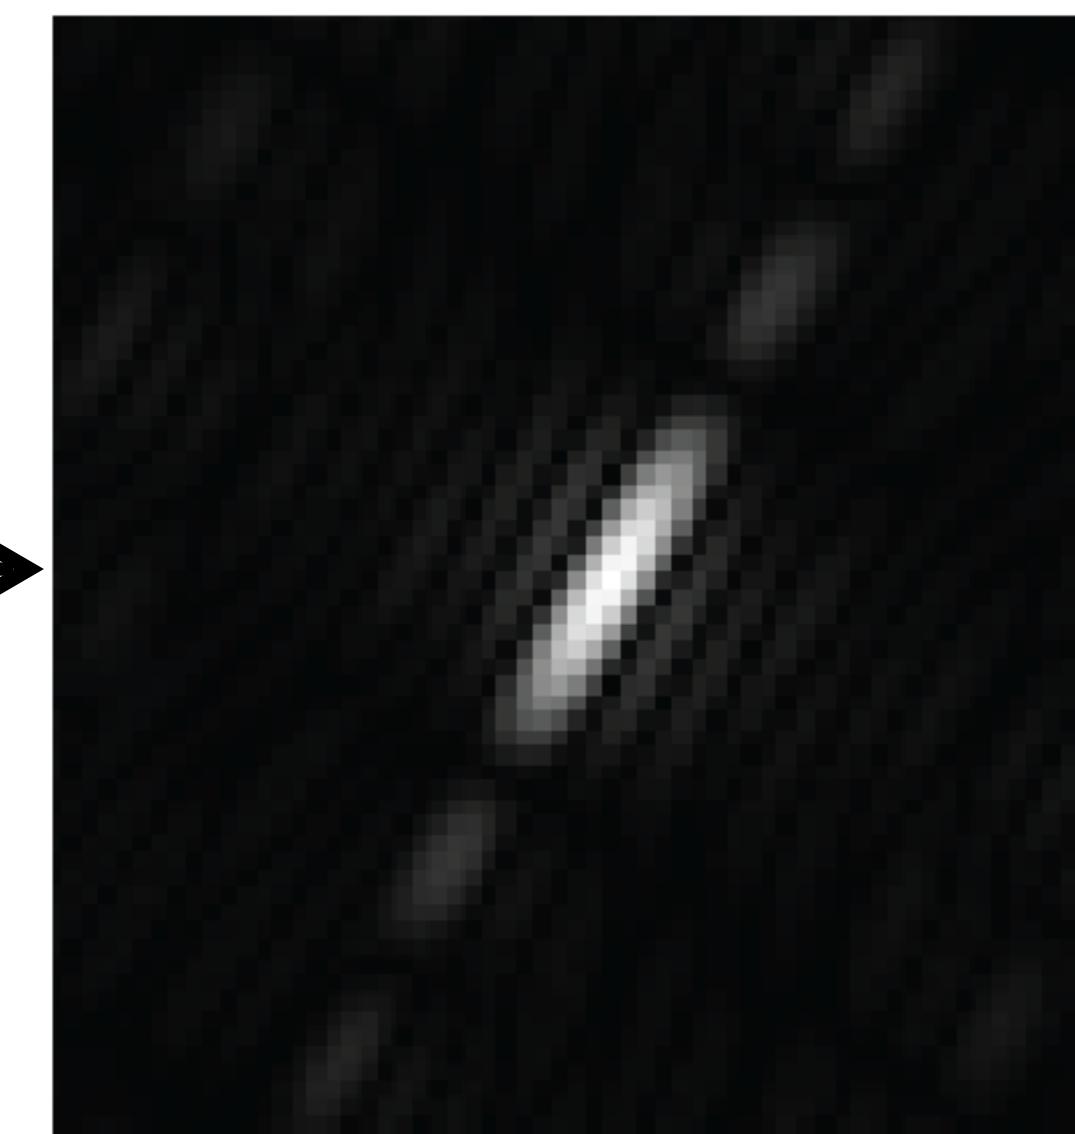
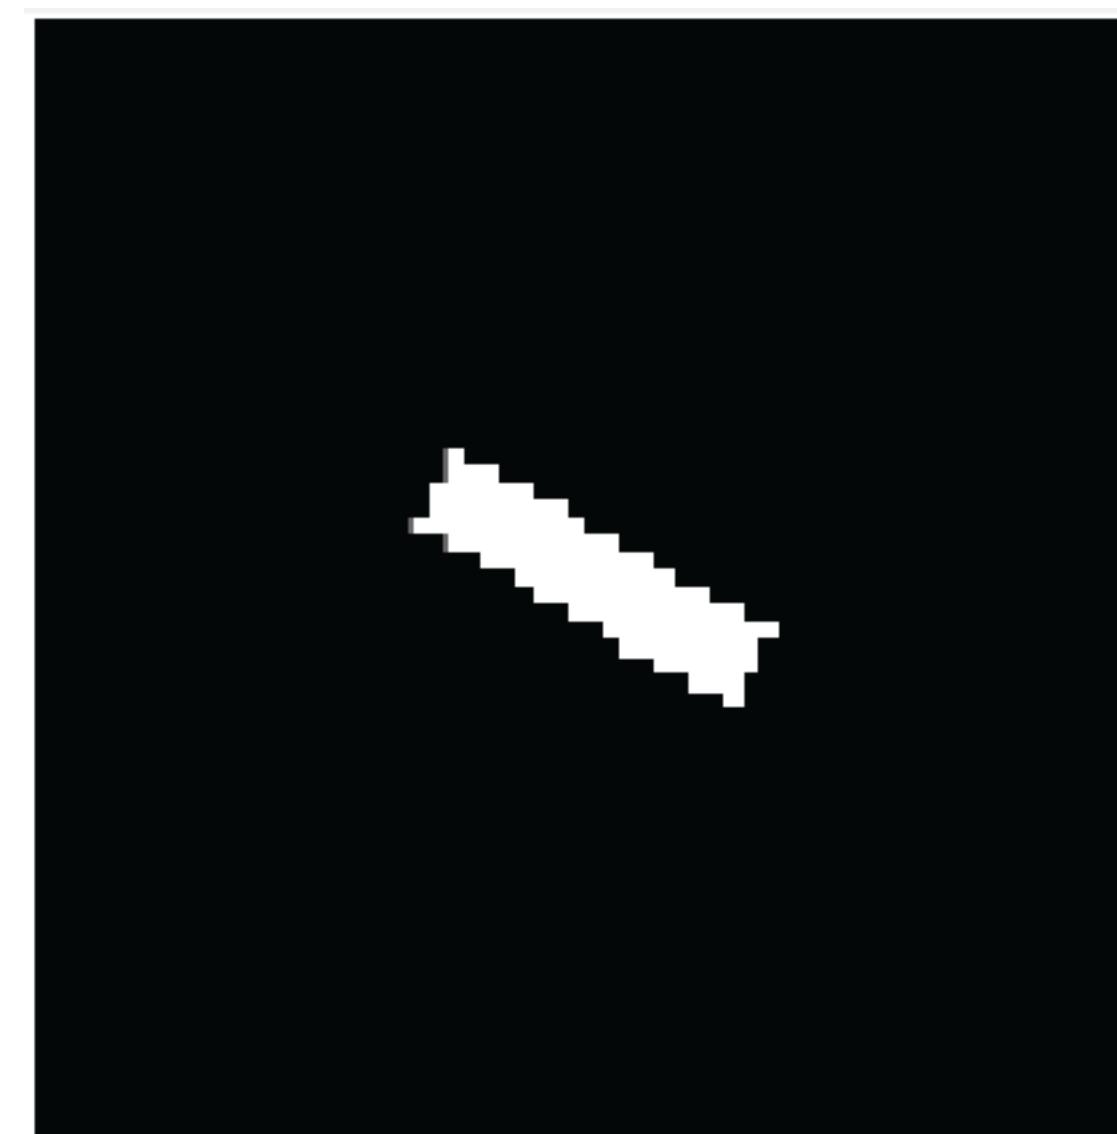
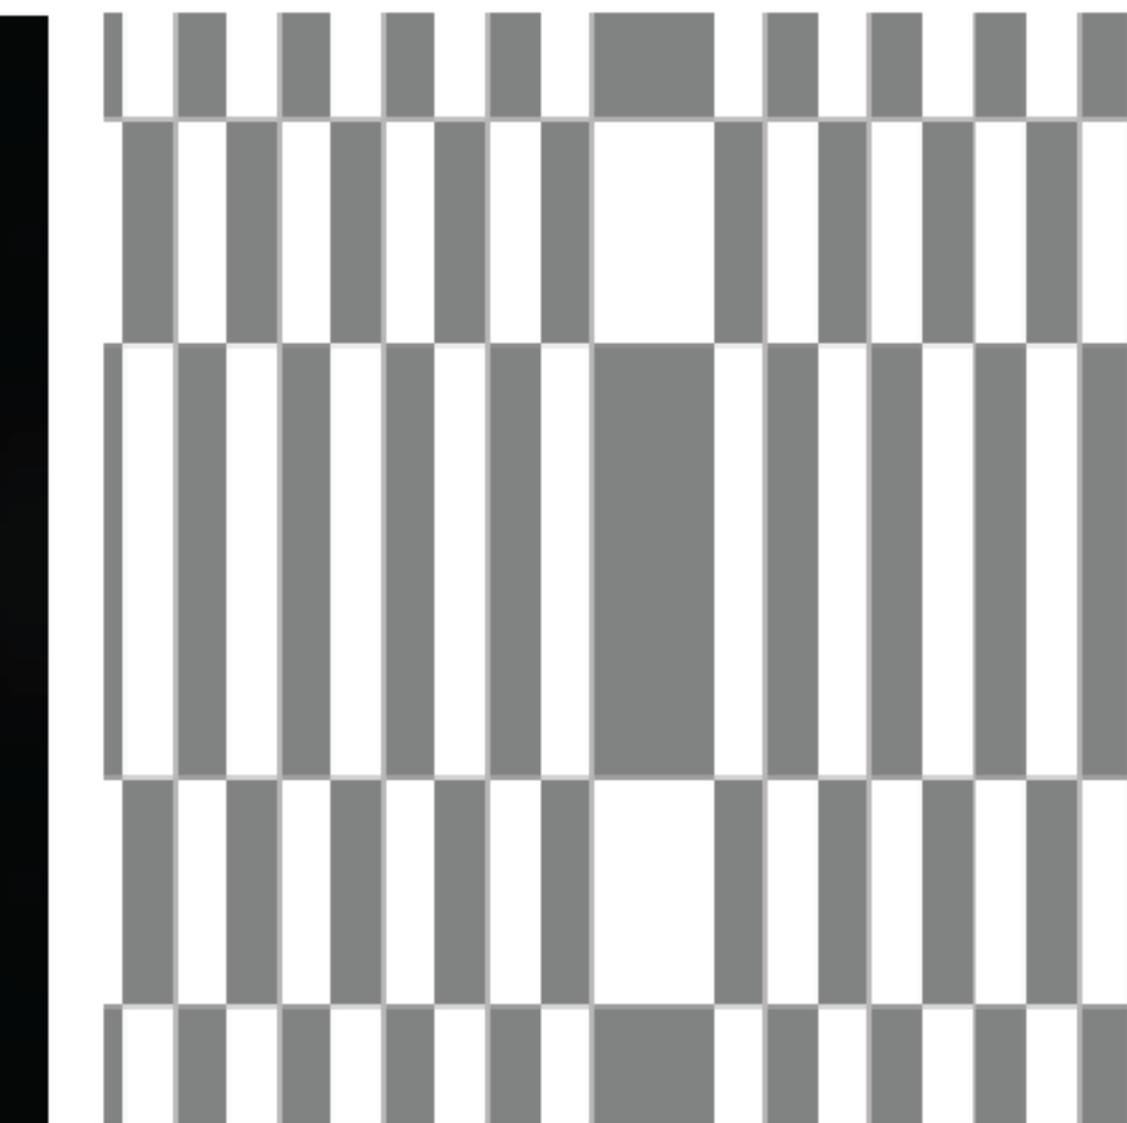
Image



Magnitude DFT



Phase DFT

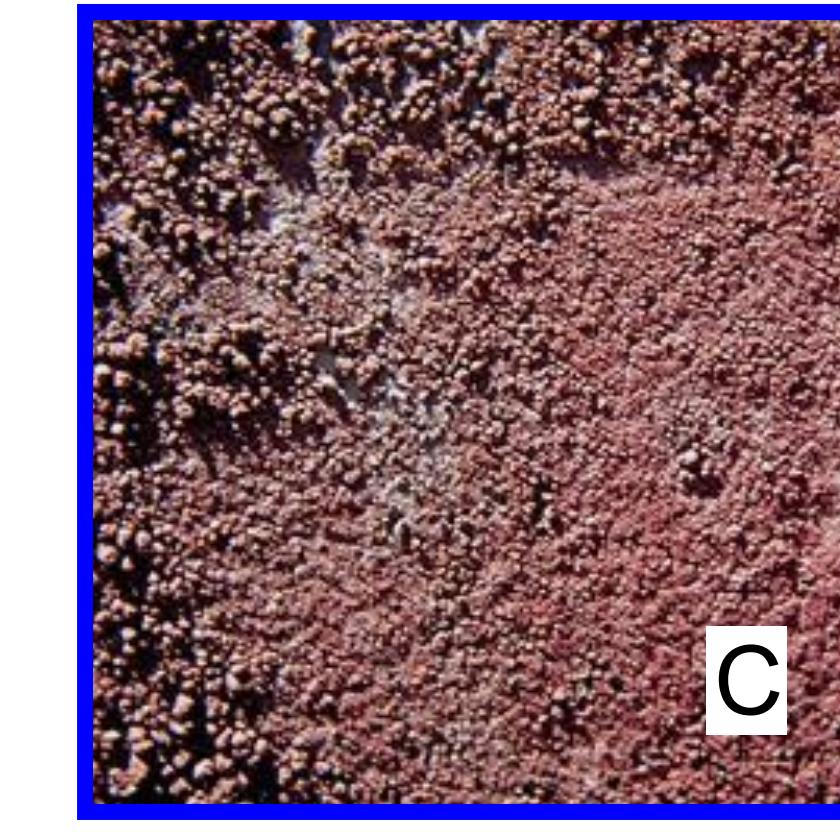
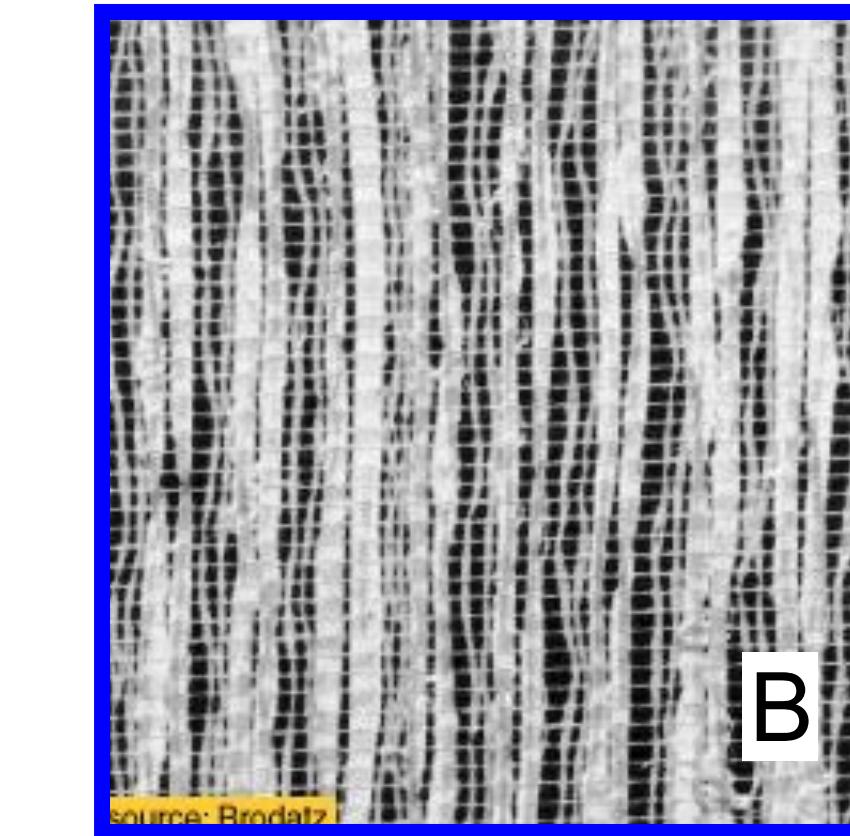
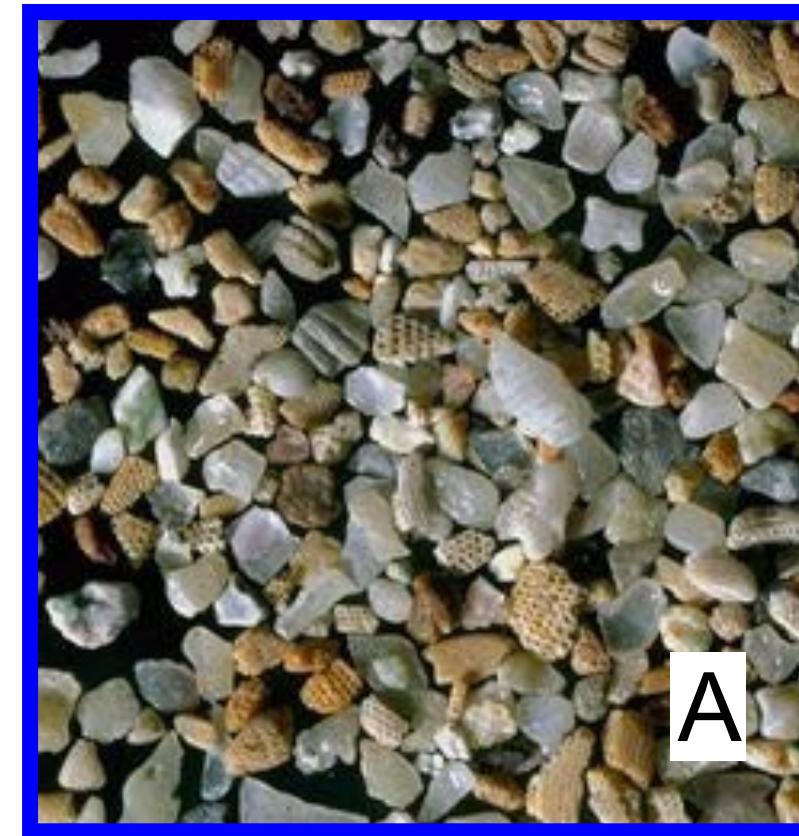


Orientation

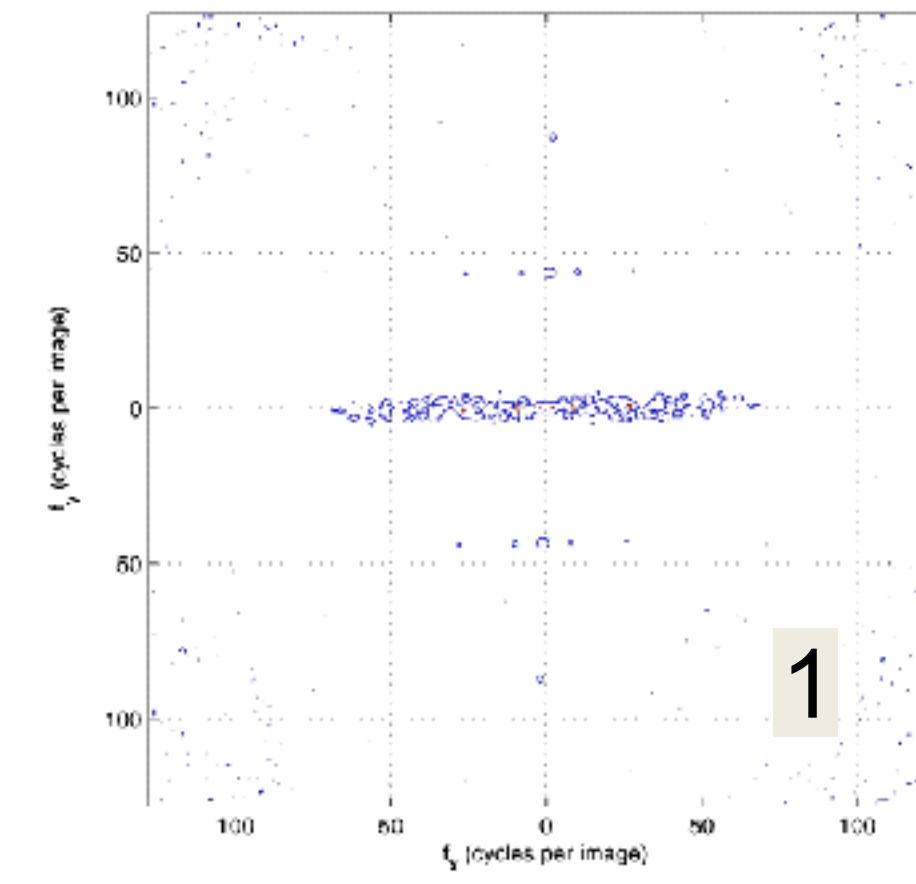
A line transforms to a line oriented perpendicularly to the first.

The DFT Game: find the right pairs

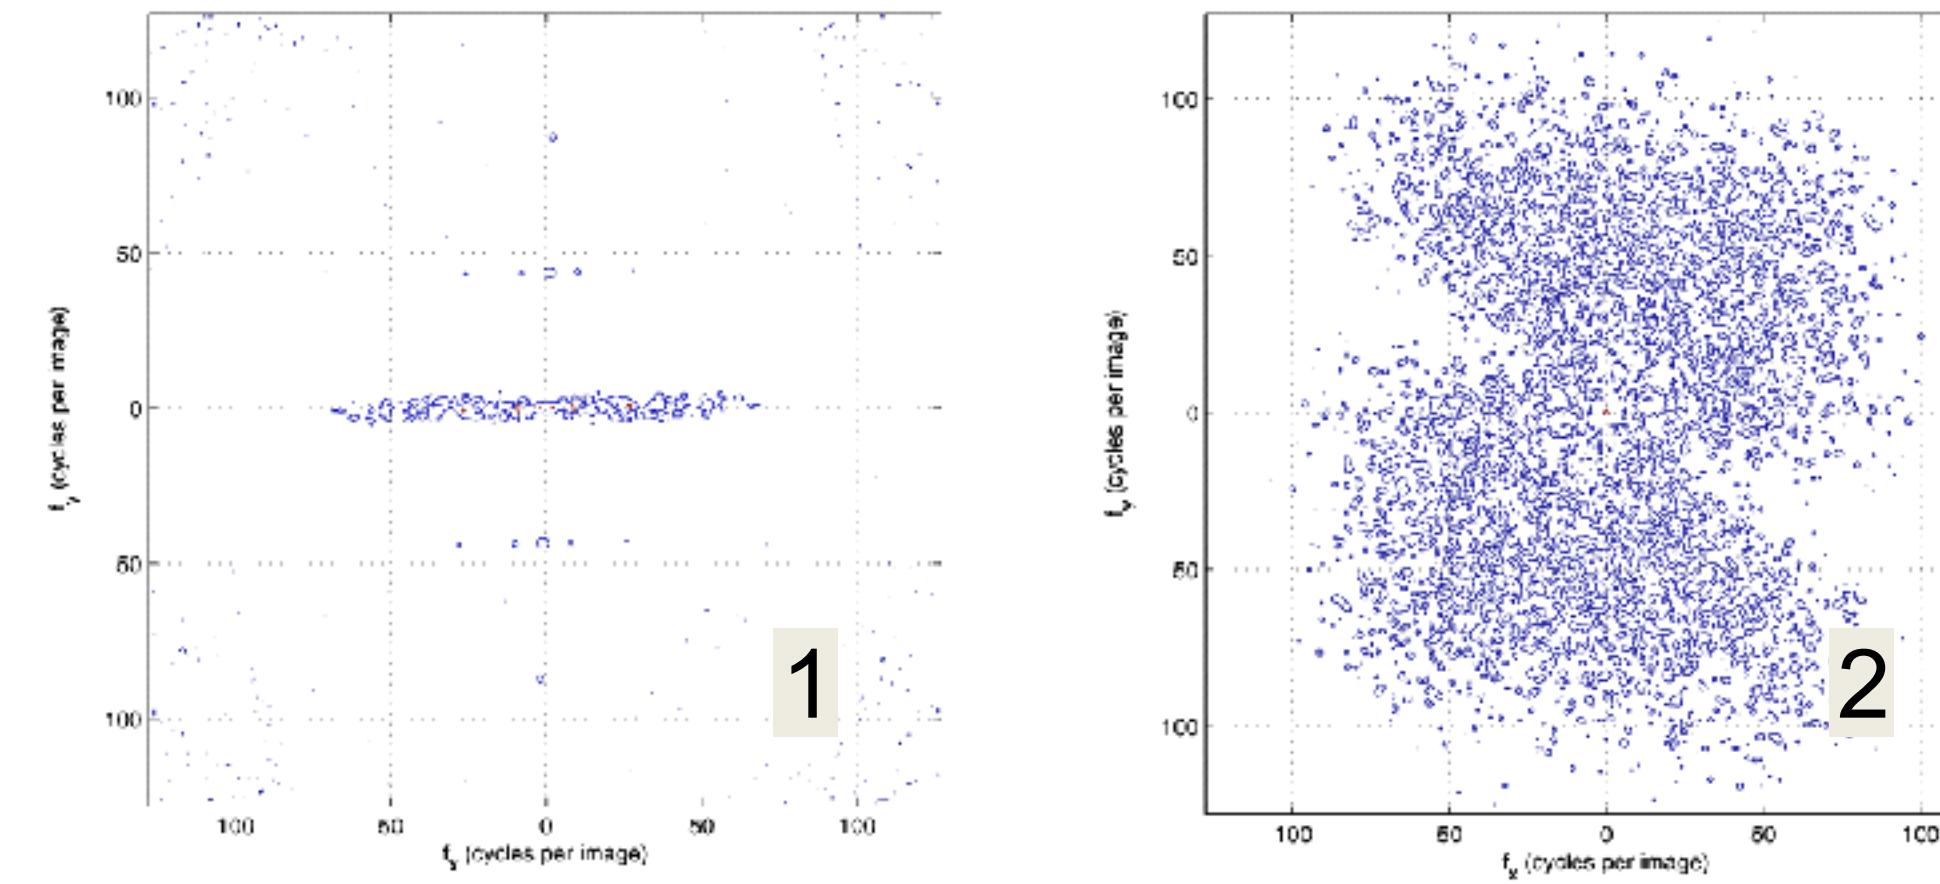
Images



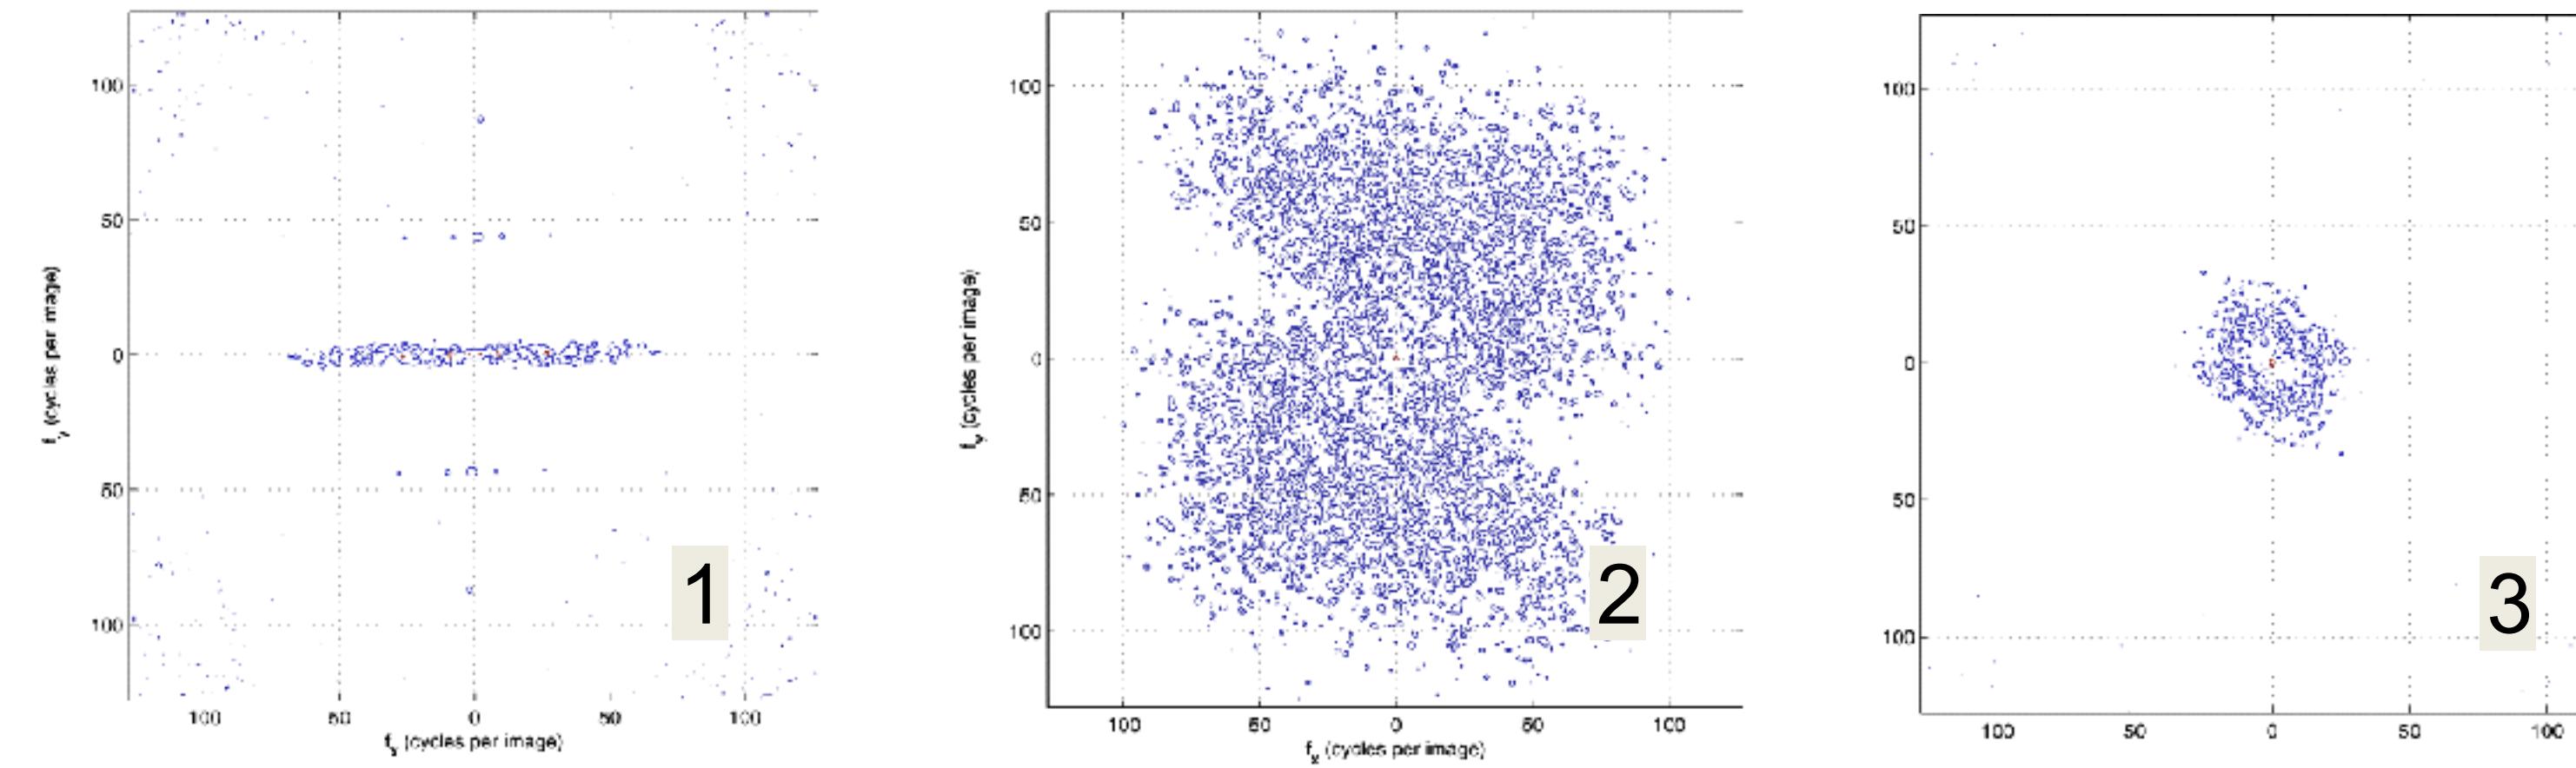
DFT
magnitude



f_x (cycles/image pixel size)

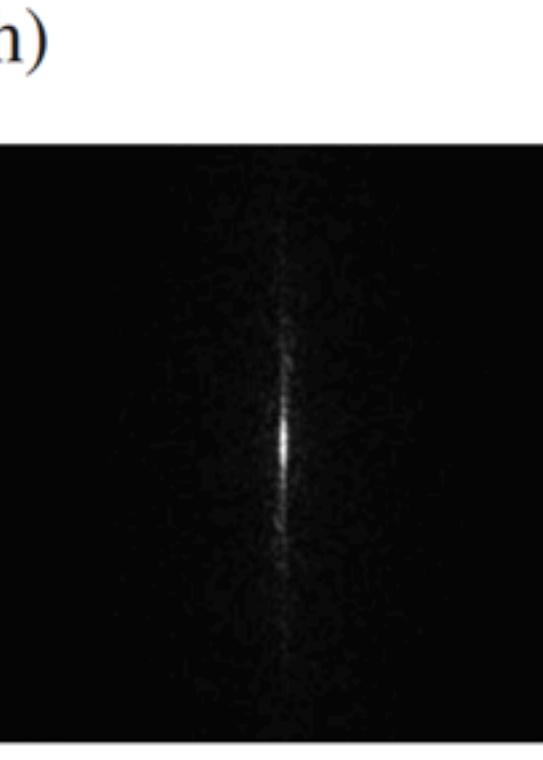
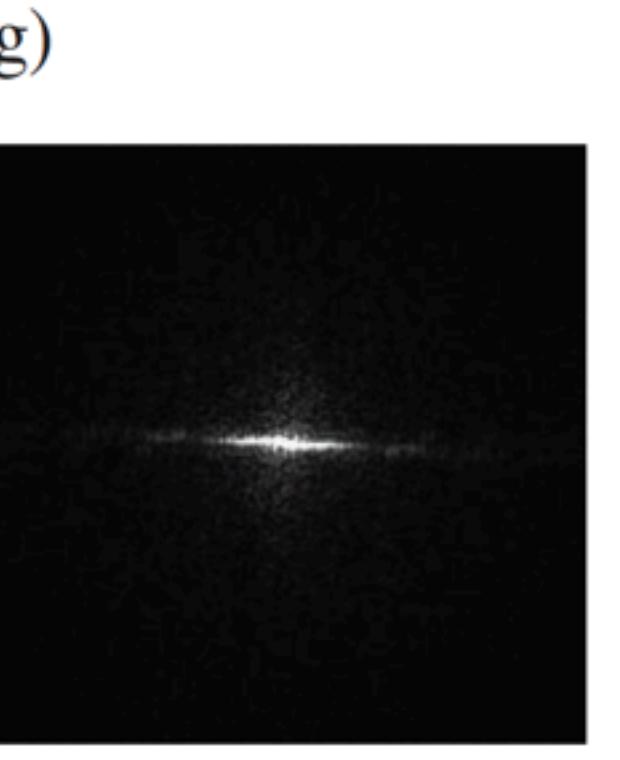
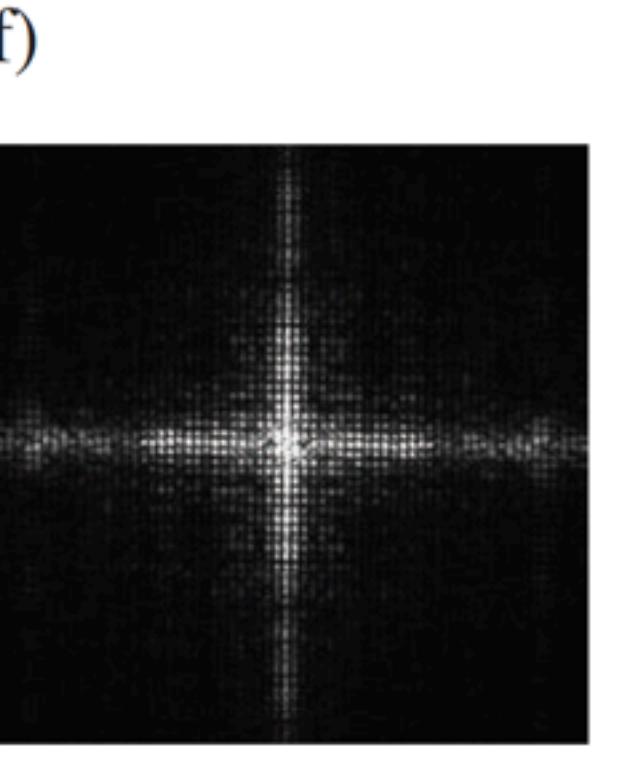
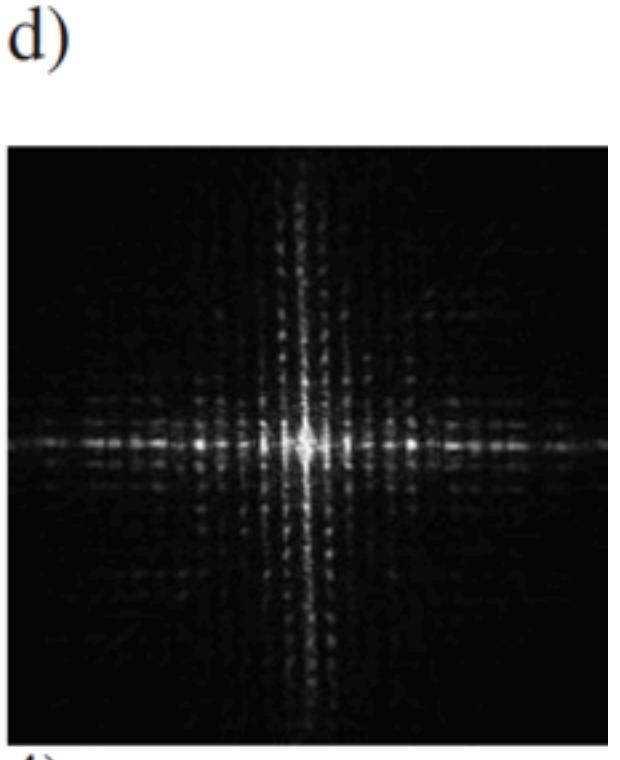
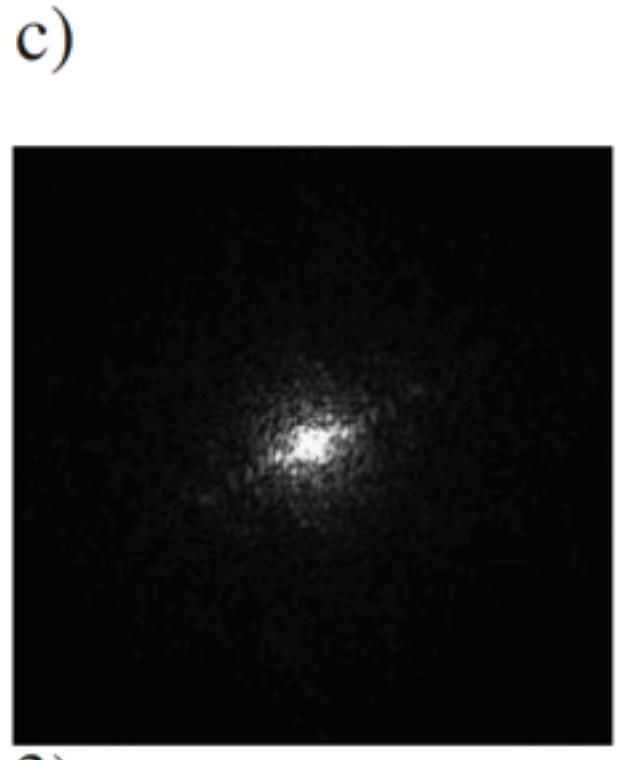
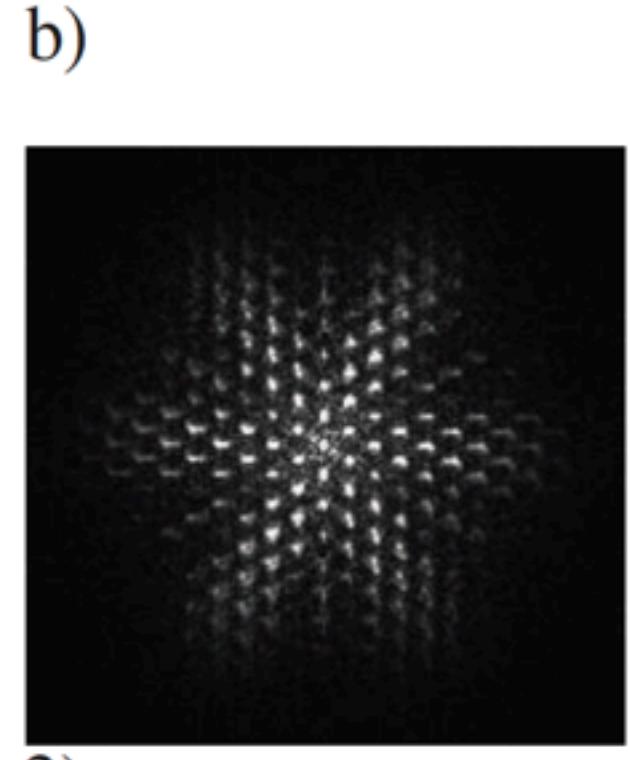
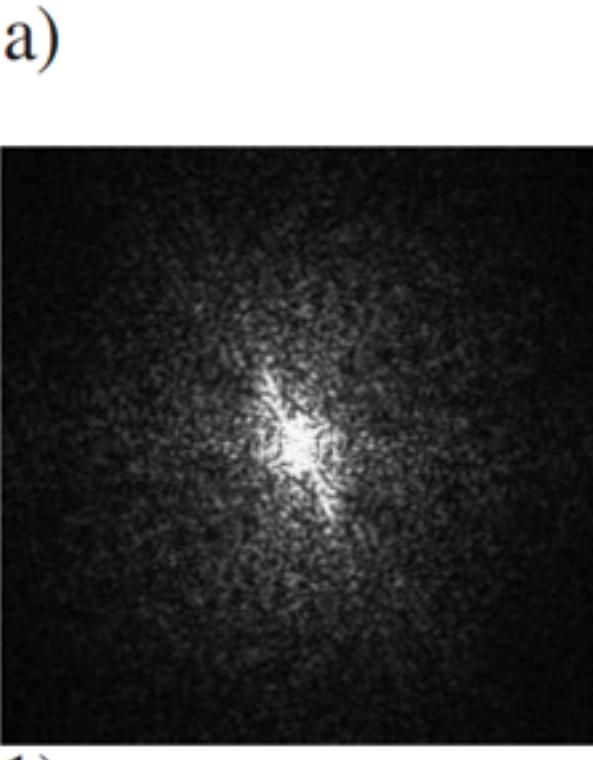
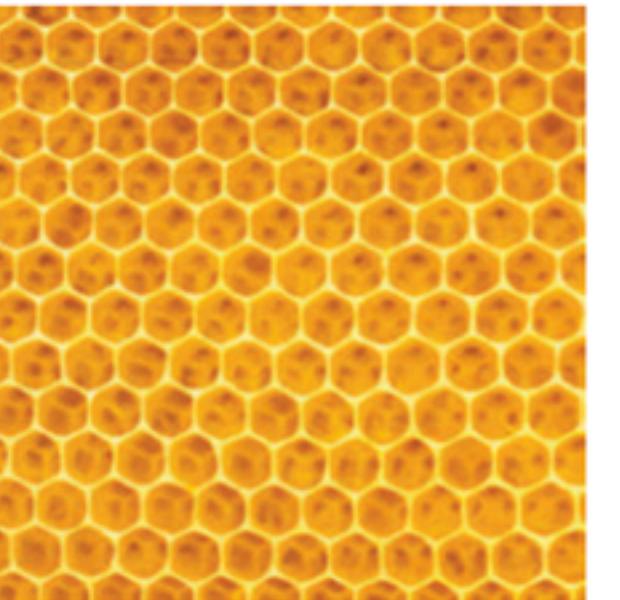


f_x (cycles/image pixel size)



f_x (cycles/image pixel size)

The DFT Game: find the right pairs



(Solution in the class notes)

The inverse Discrete Fourier transform

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

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

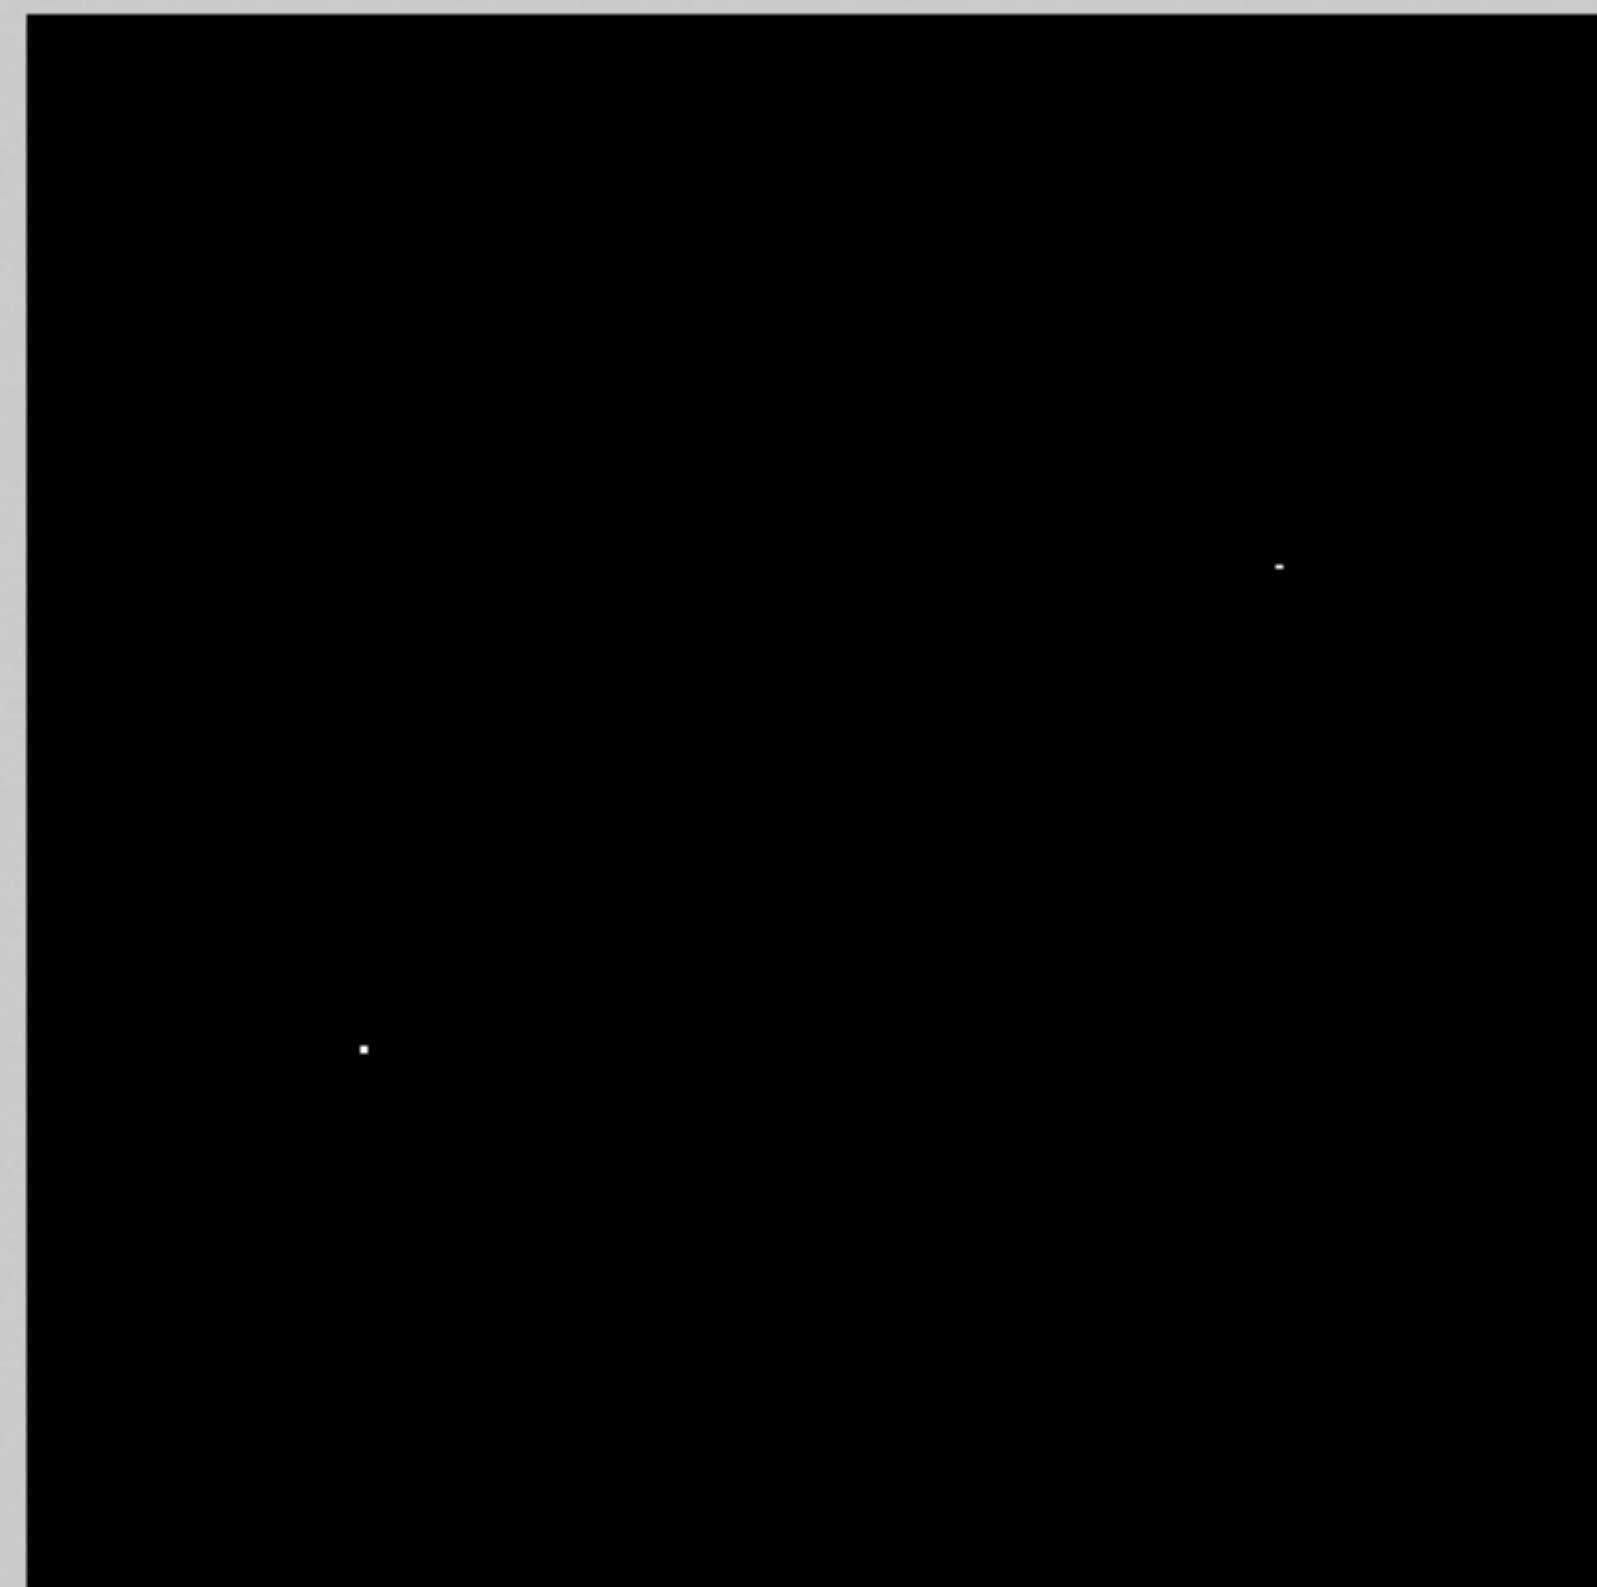
The inverse of the 2D DFT is:

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

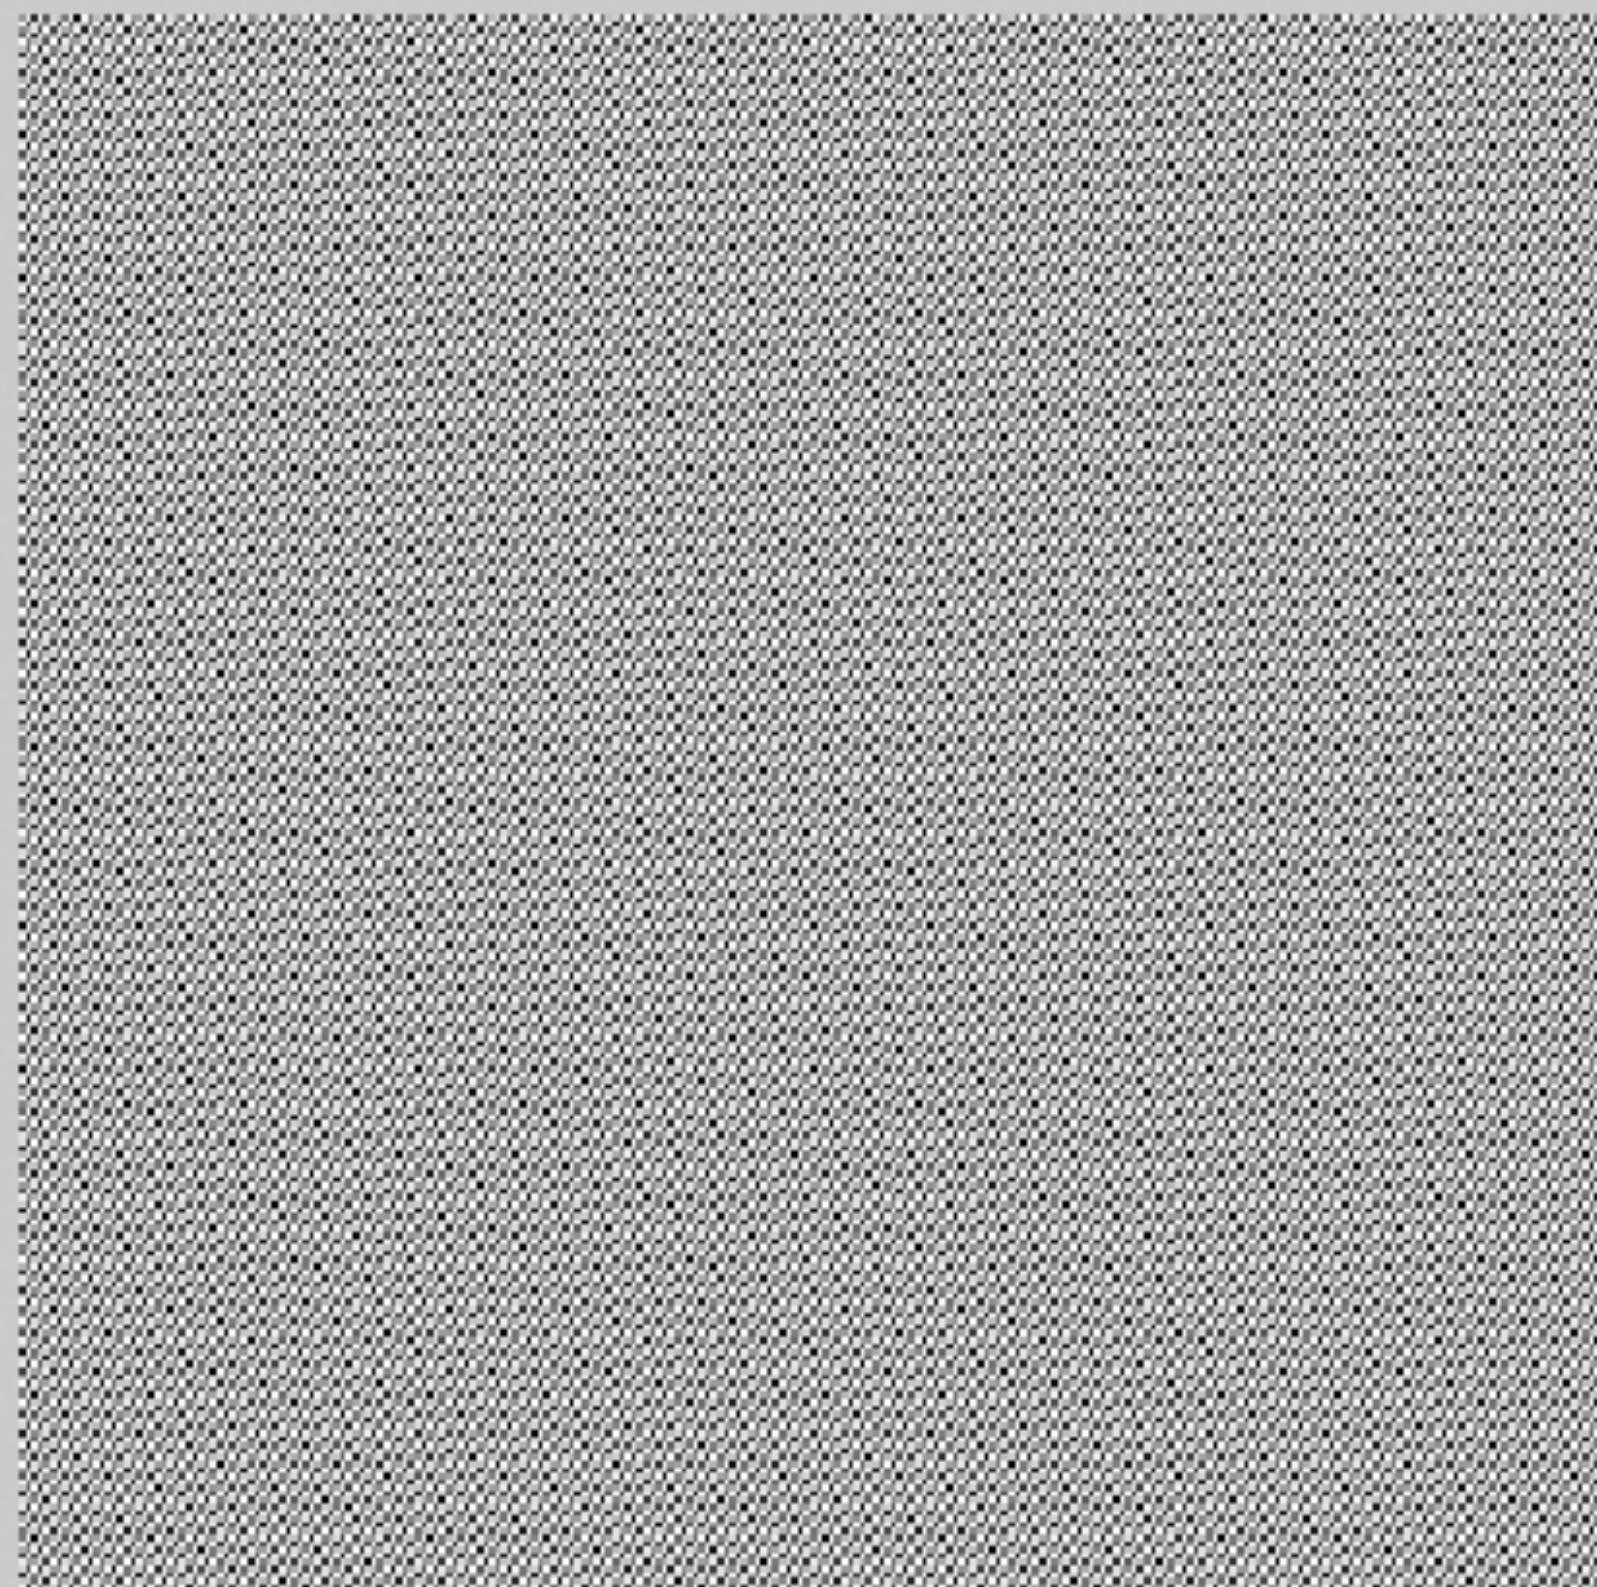
How does summing waves ends up giving back a picture?

2

2



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



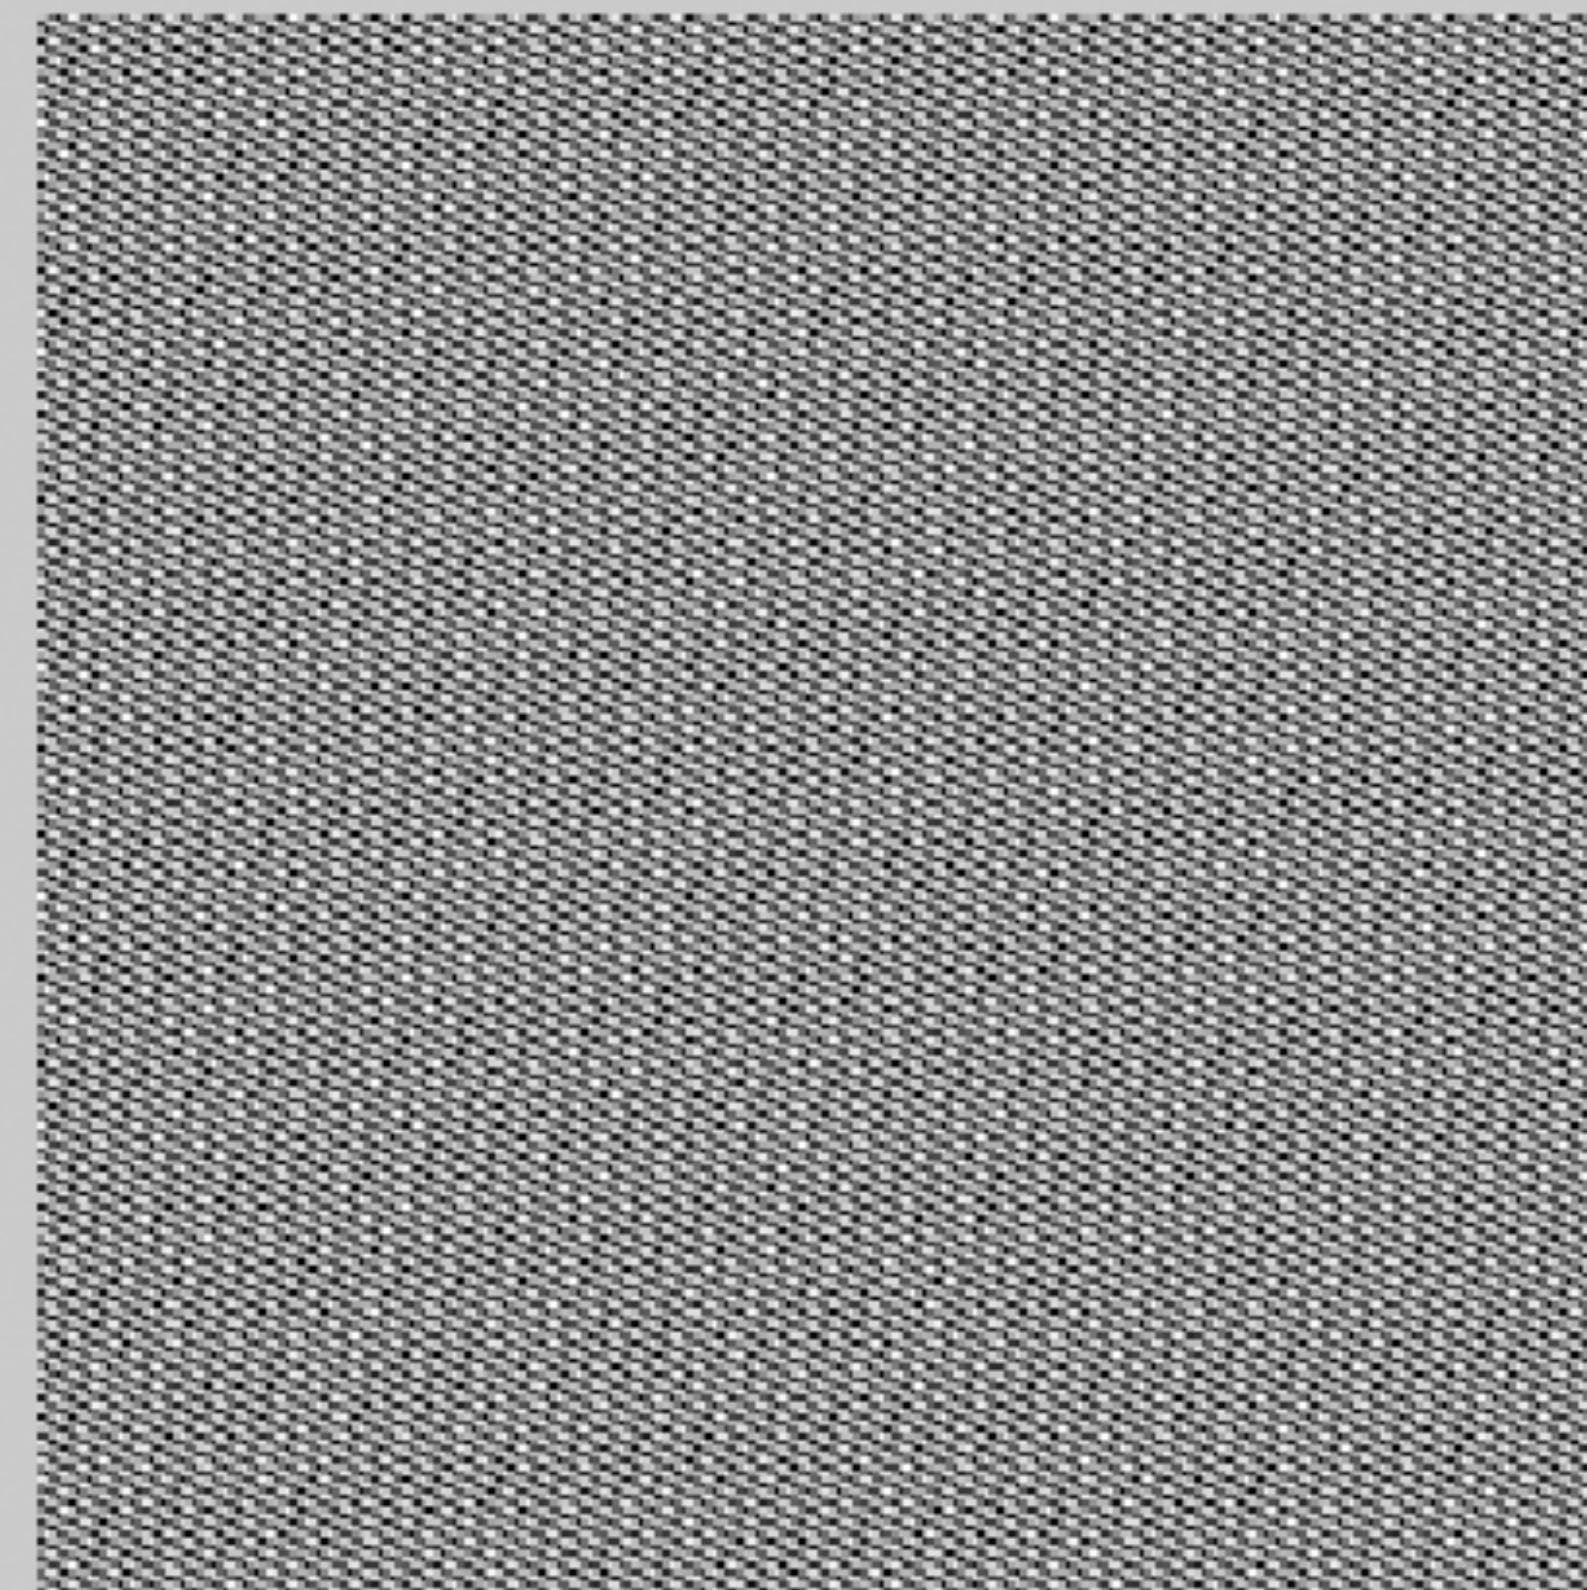
#2: Range [0.000109, 0.0267]
Dims [256, 256]

6

6

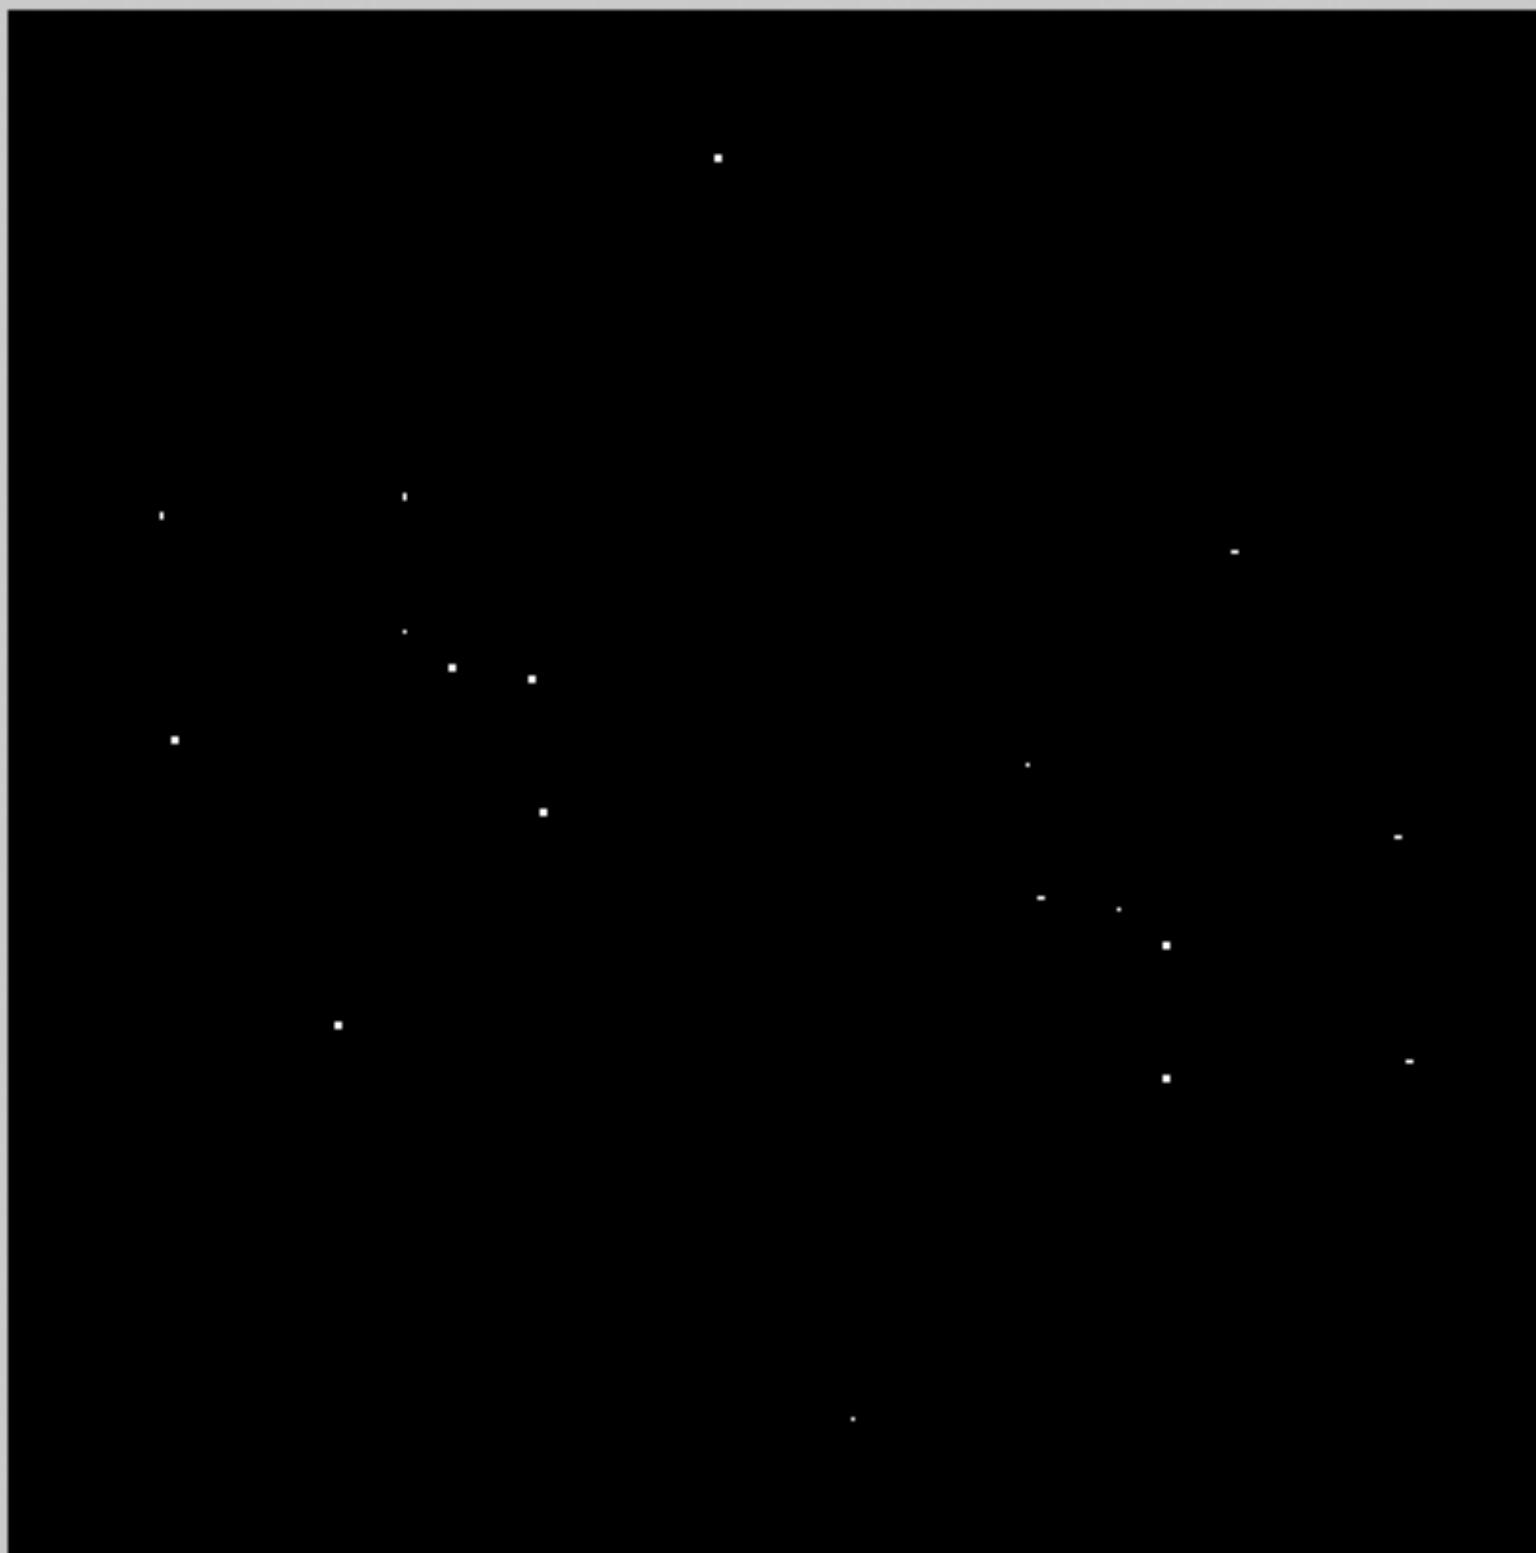


#1: Range [0, 1]
Dims [256, 256]

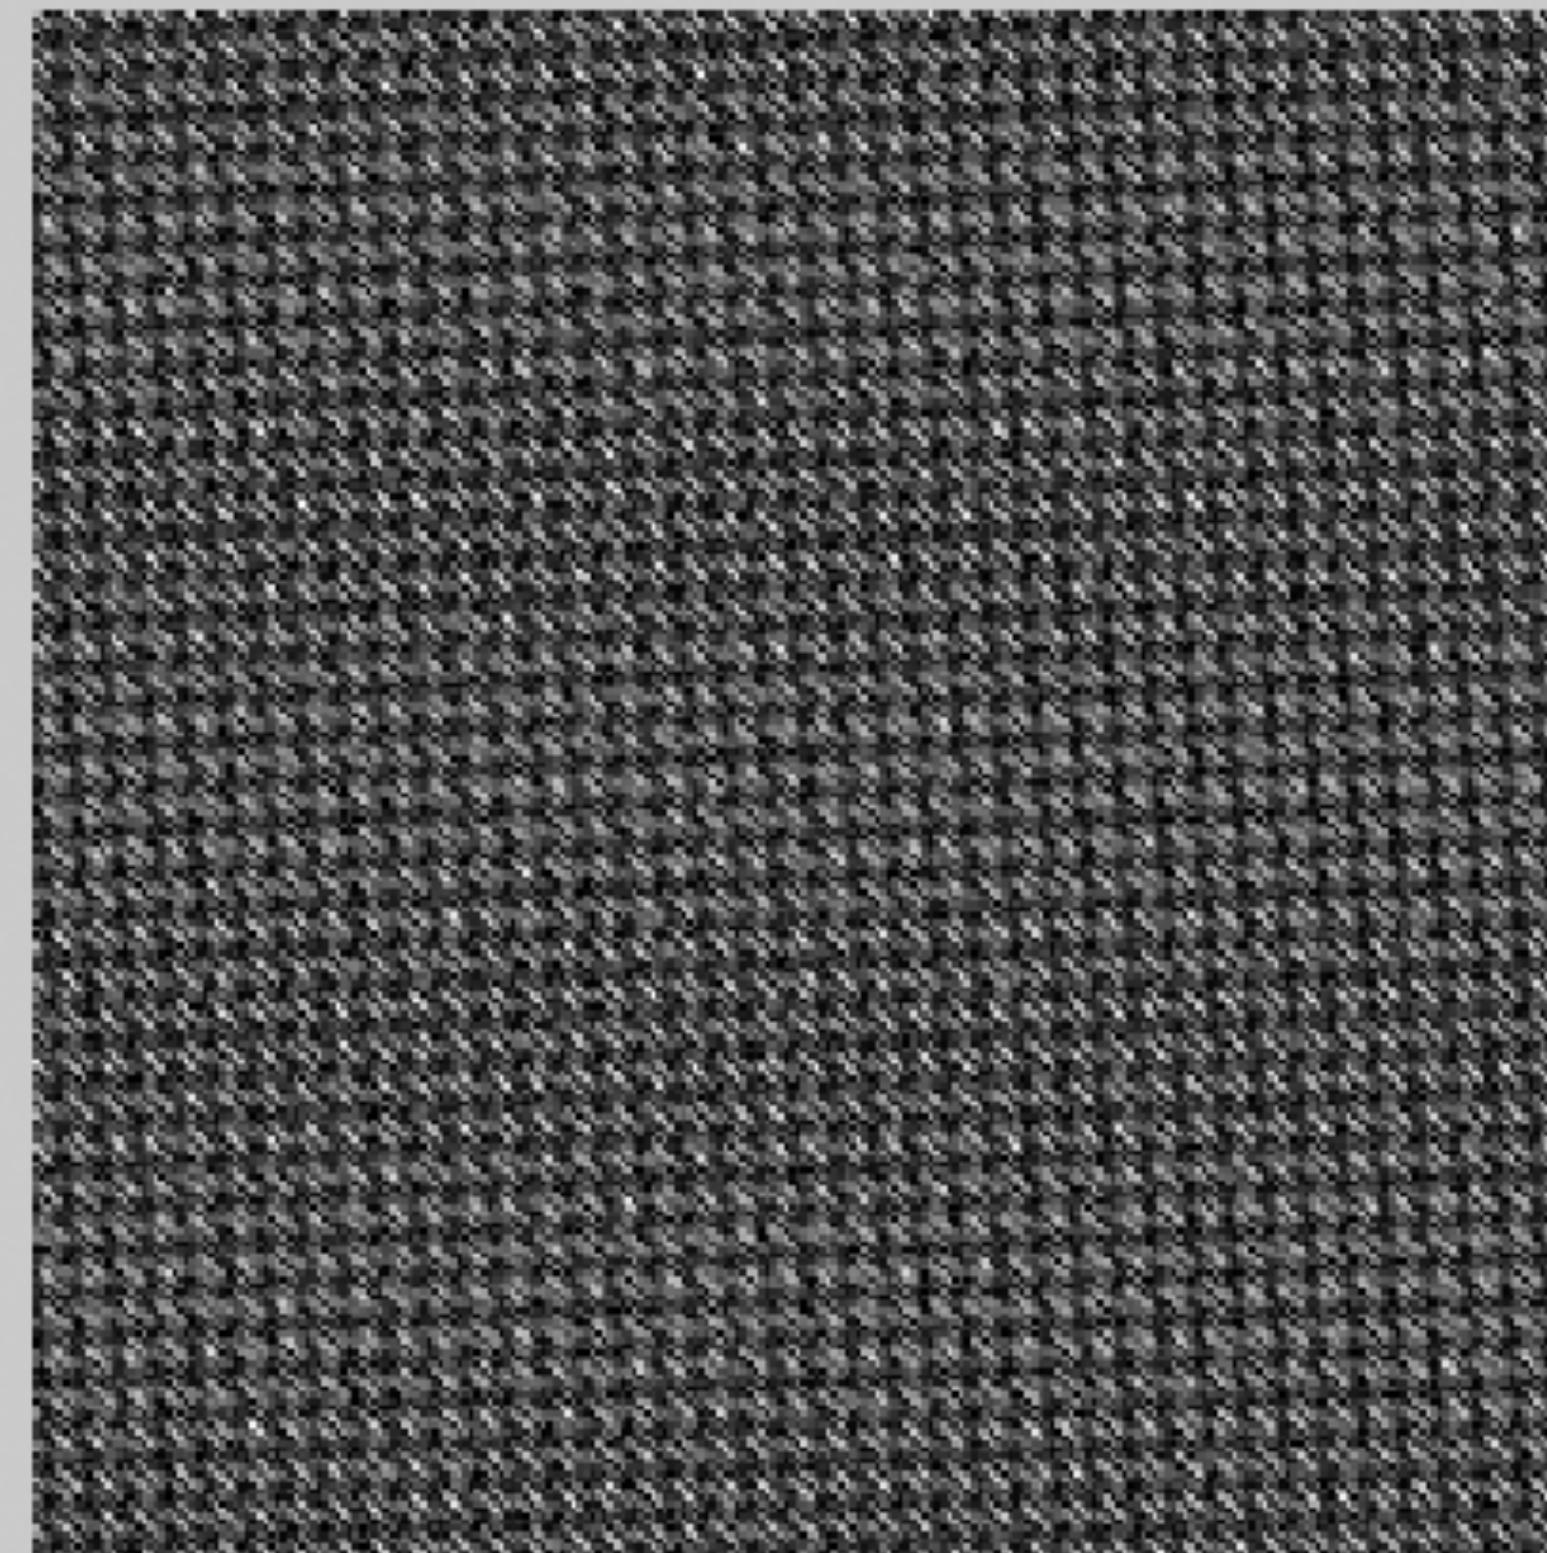


#2: Range [1.89e-007, 0.226]
Dims [256, 256]

18



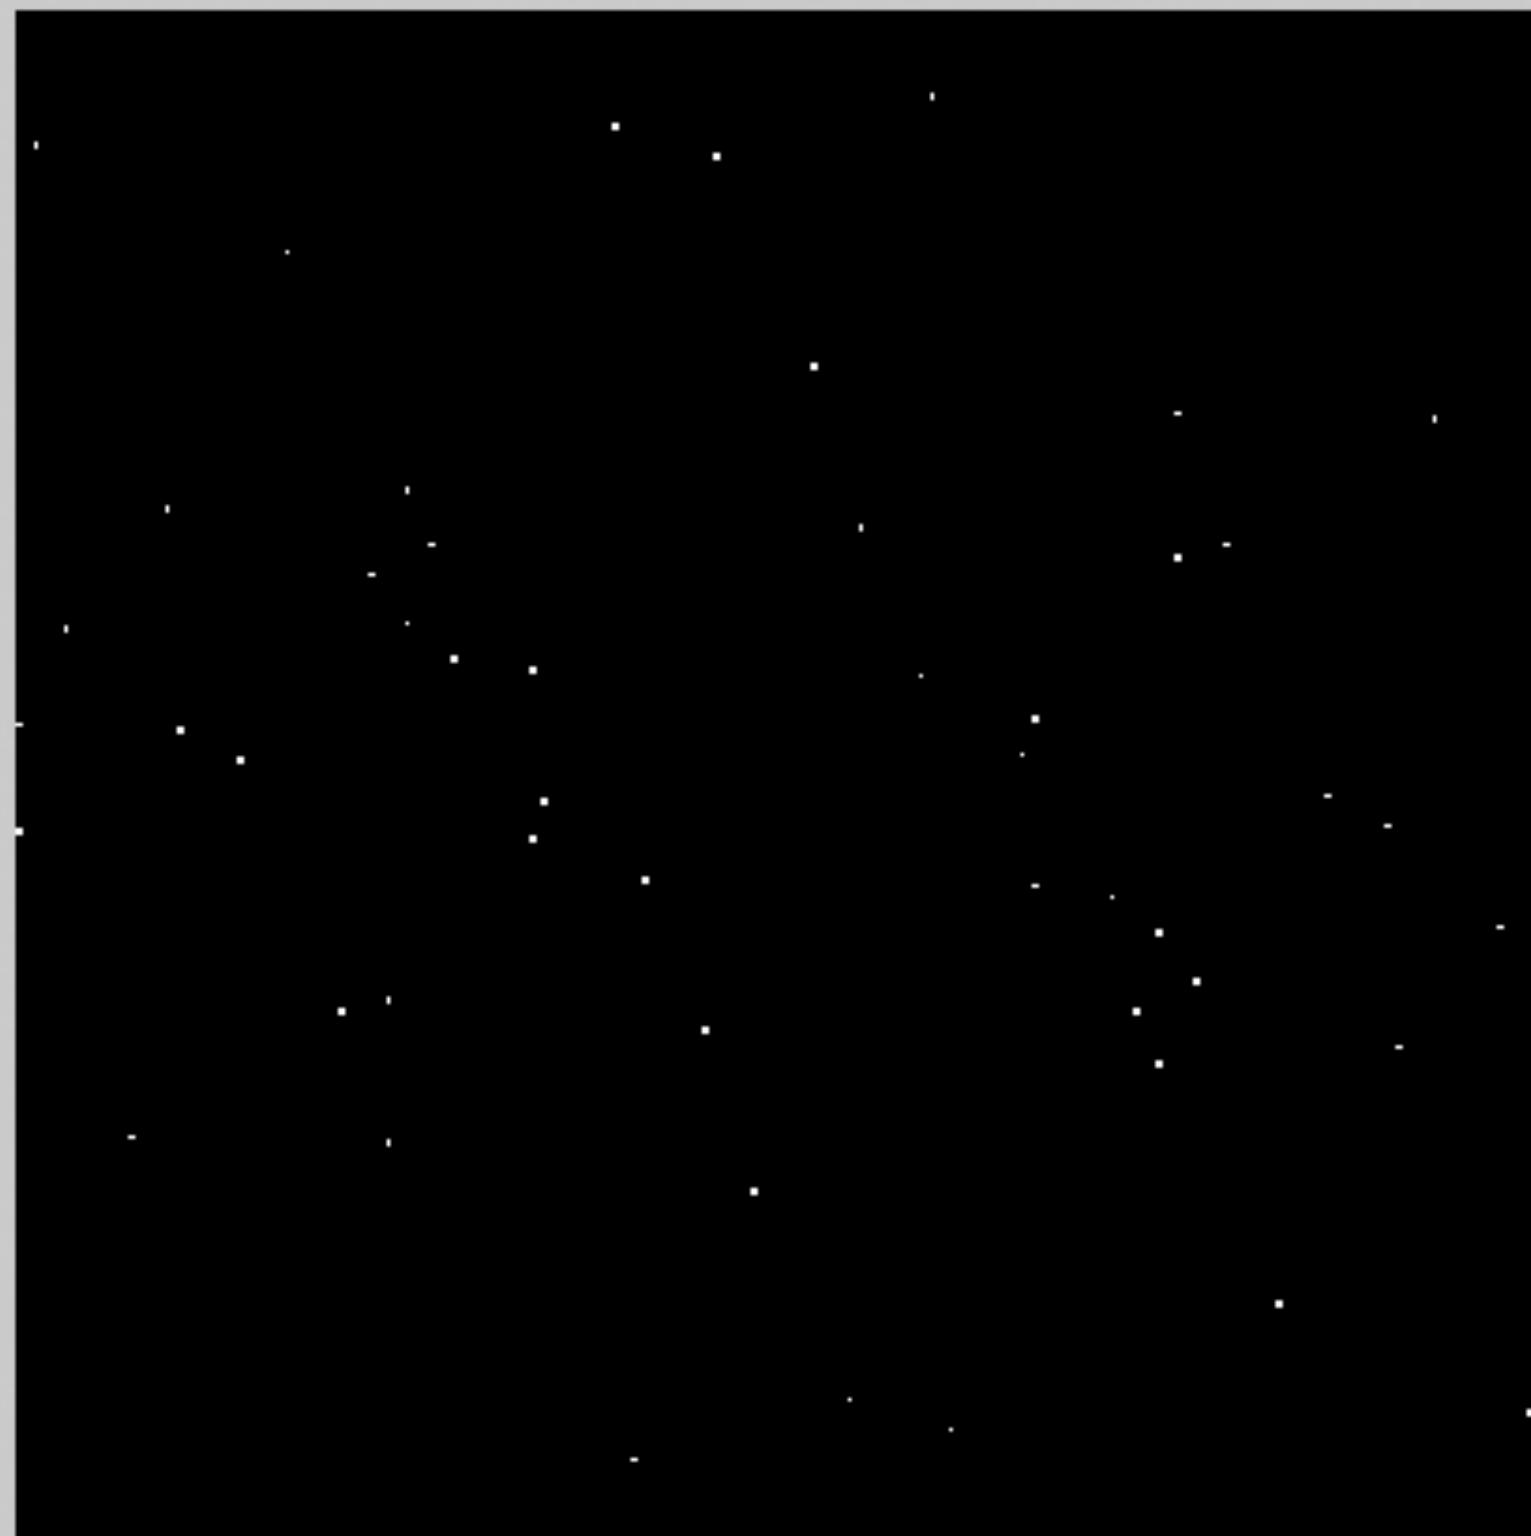
#1: Range [0, 1]
Dims [256, 256]



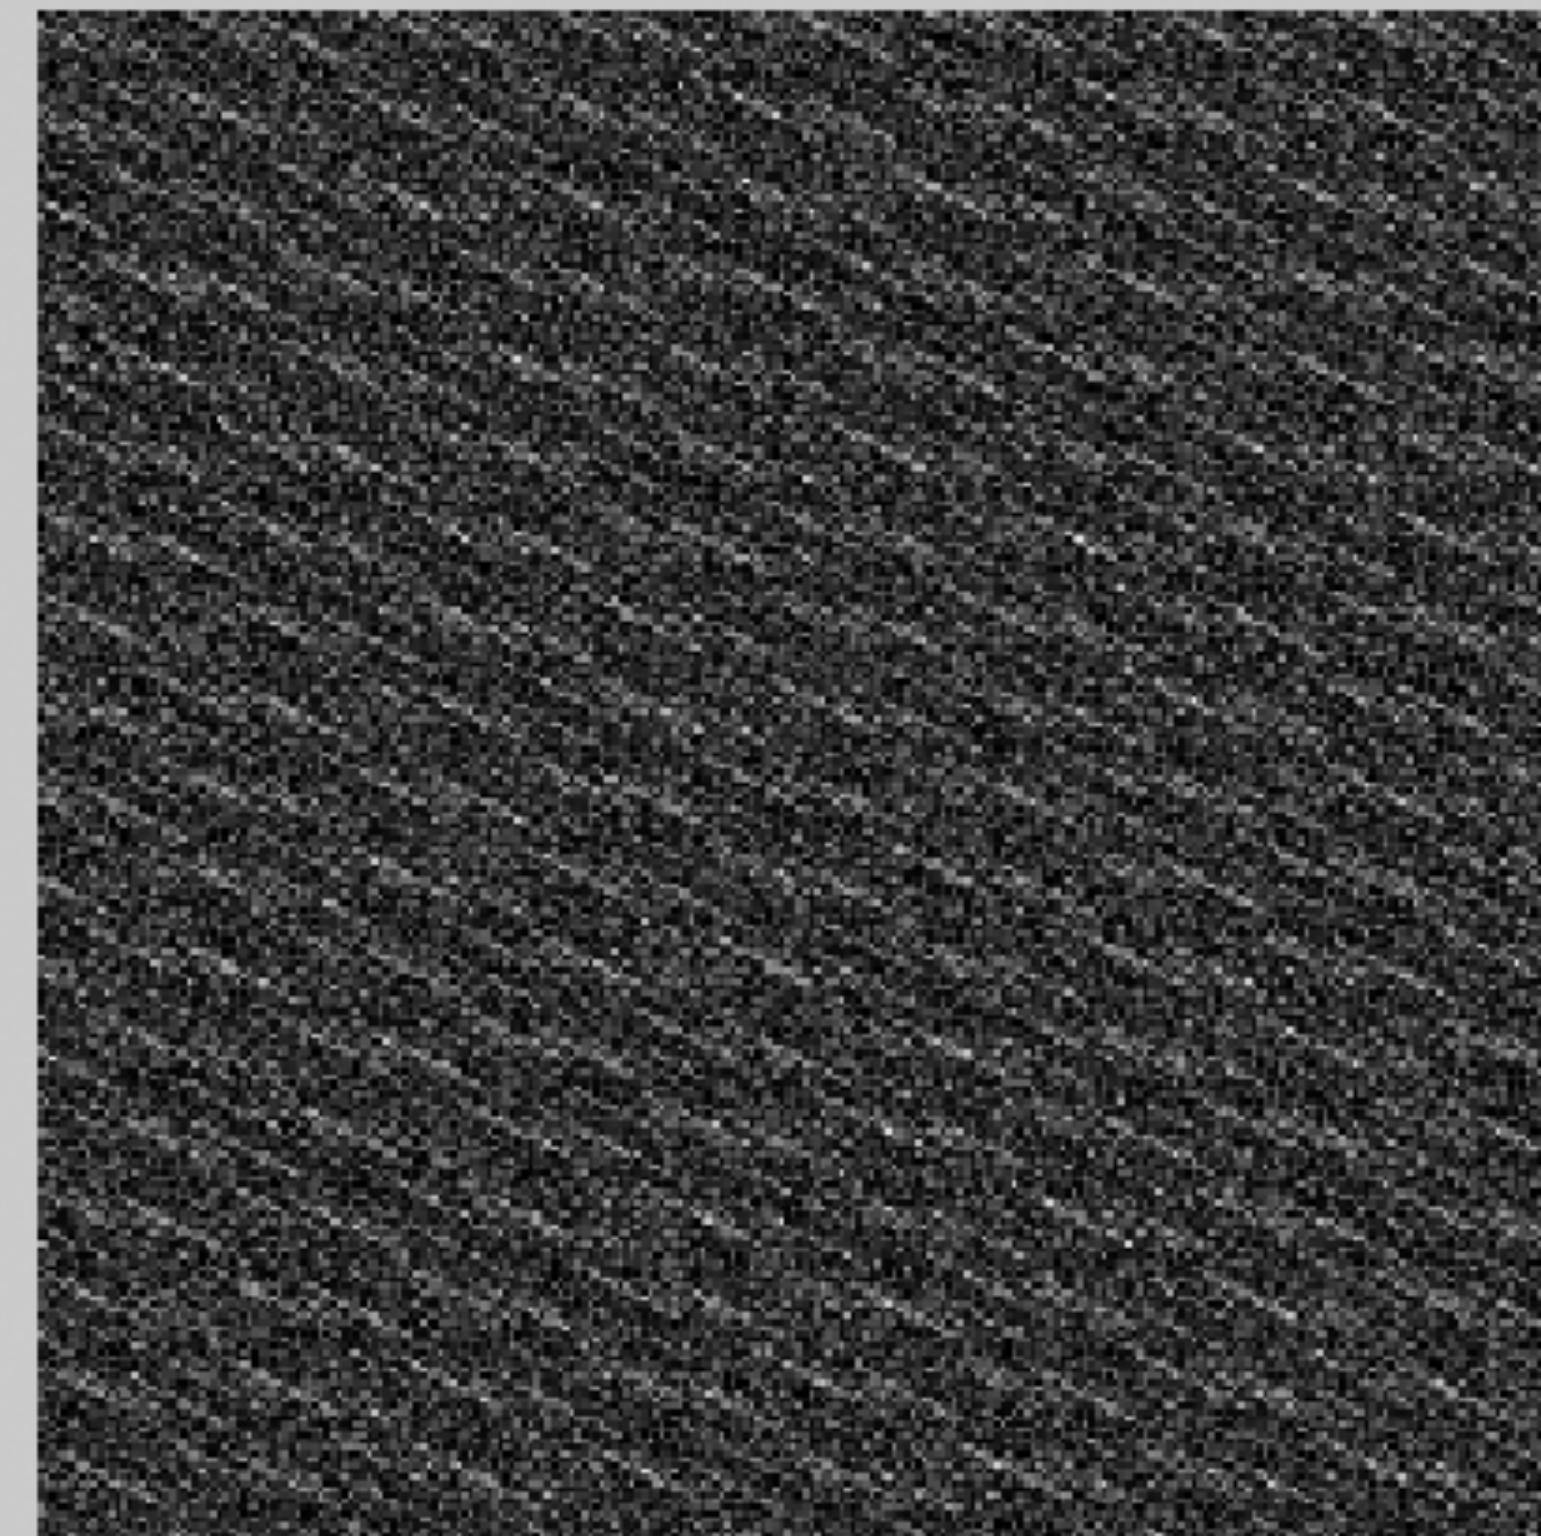
#2: Range [4.79e-007, 0.503]
Dims [256, 256]

50

50



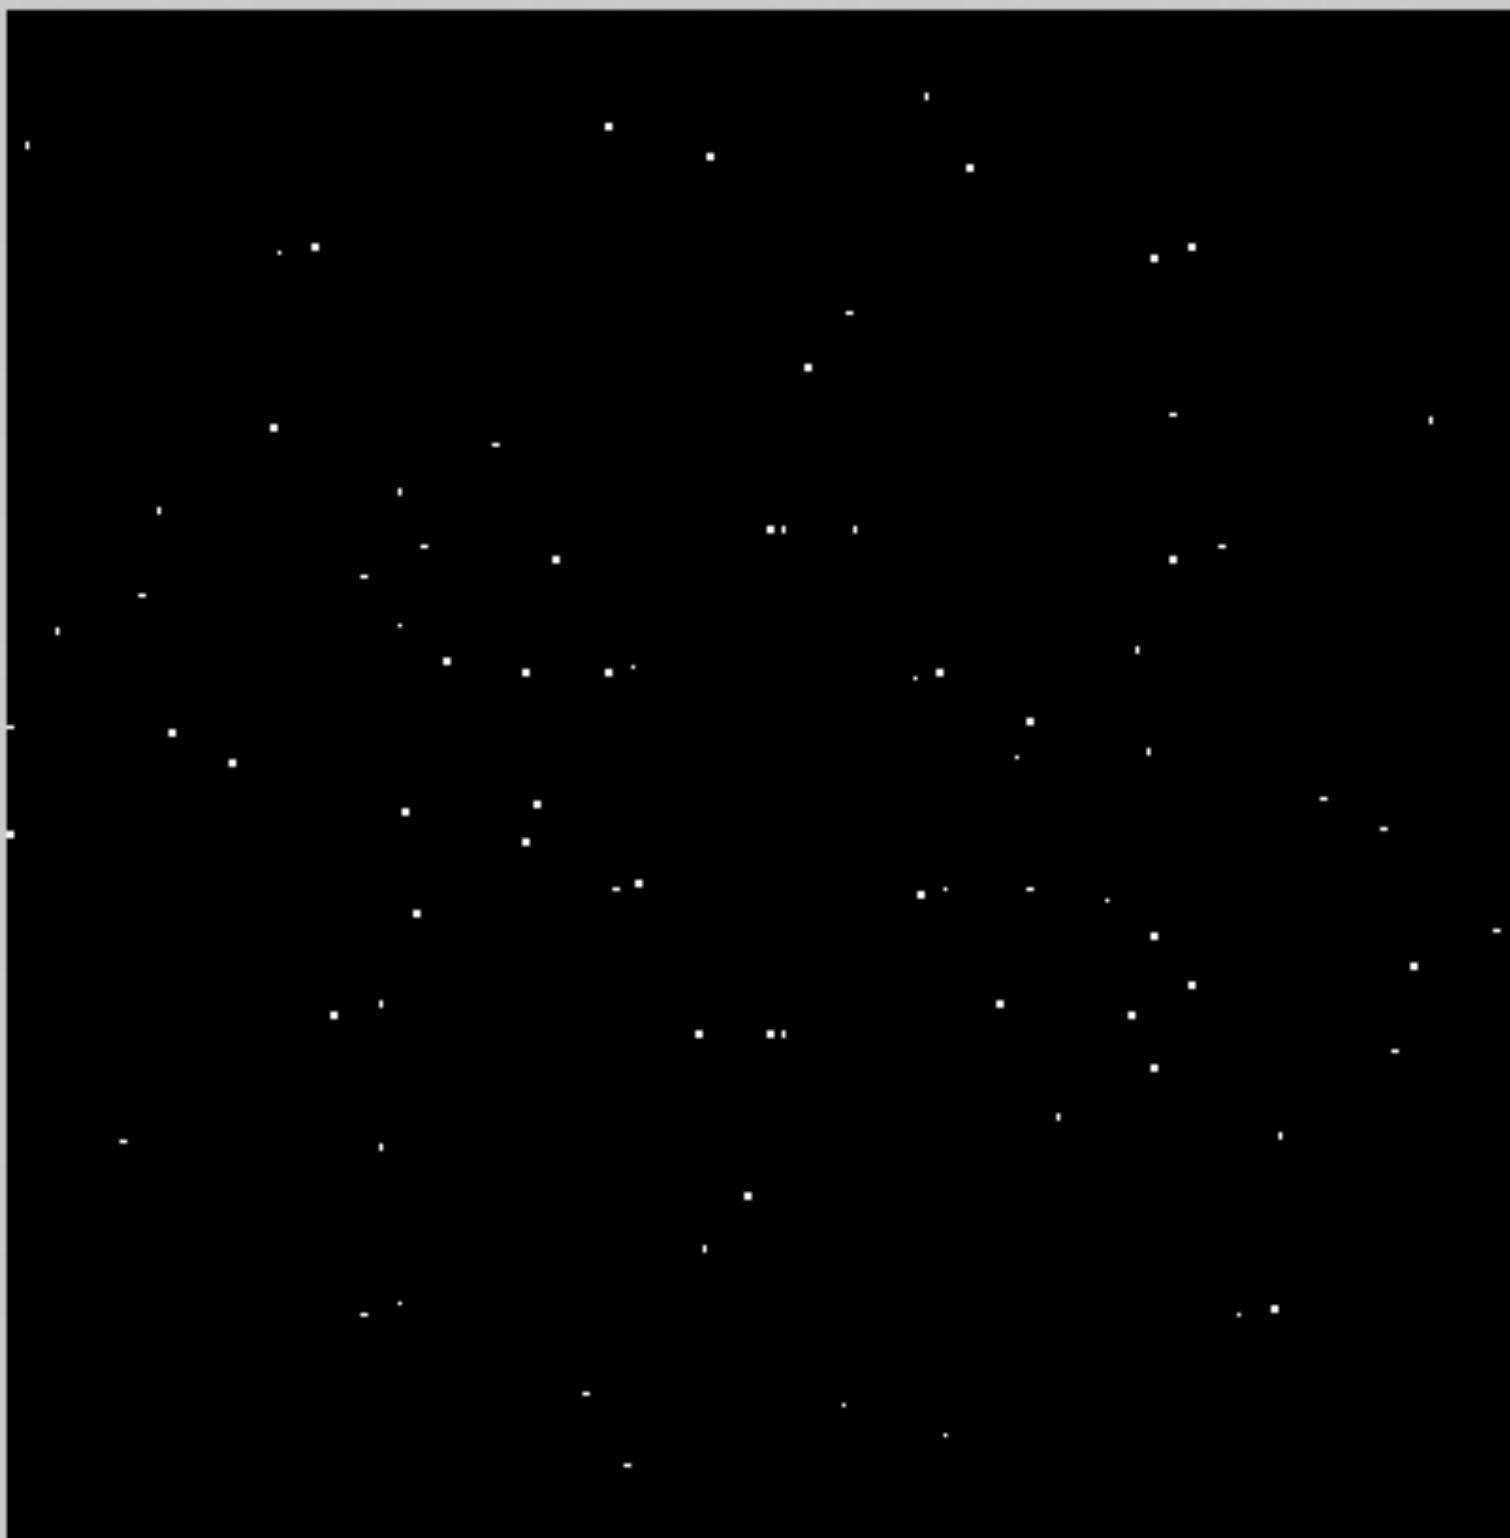
#1: Range [0, 1]
Dims [256, 256]



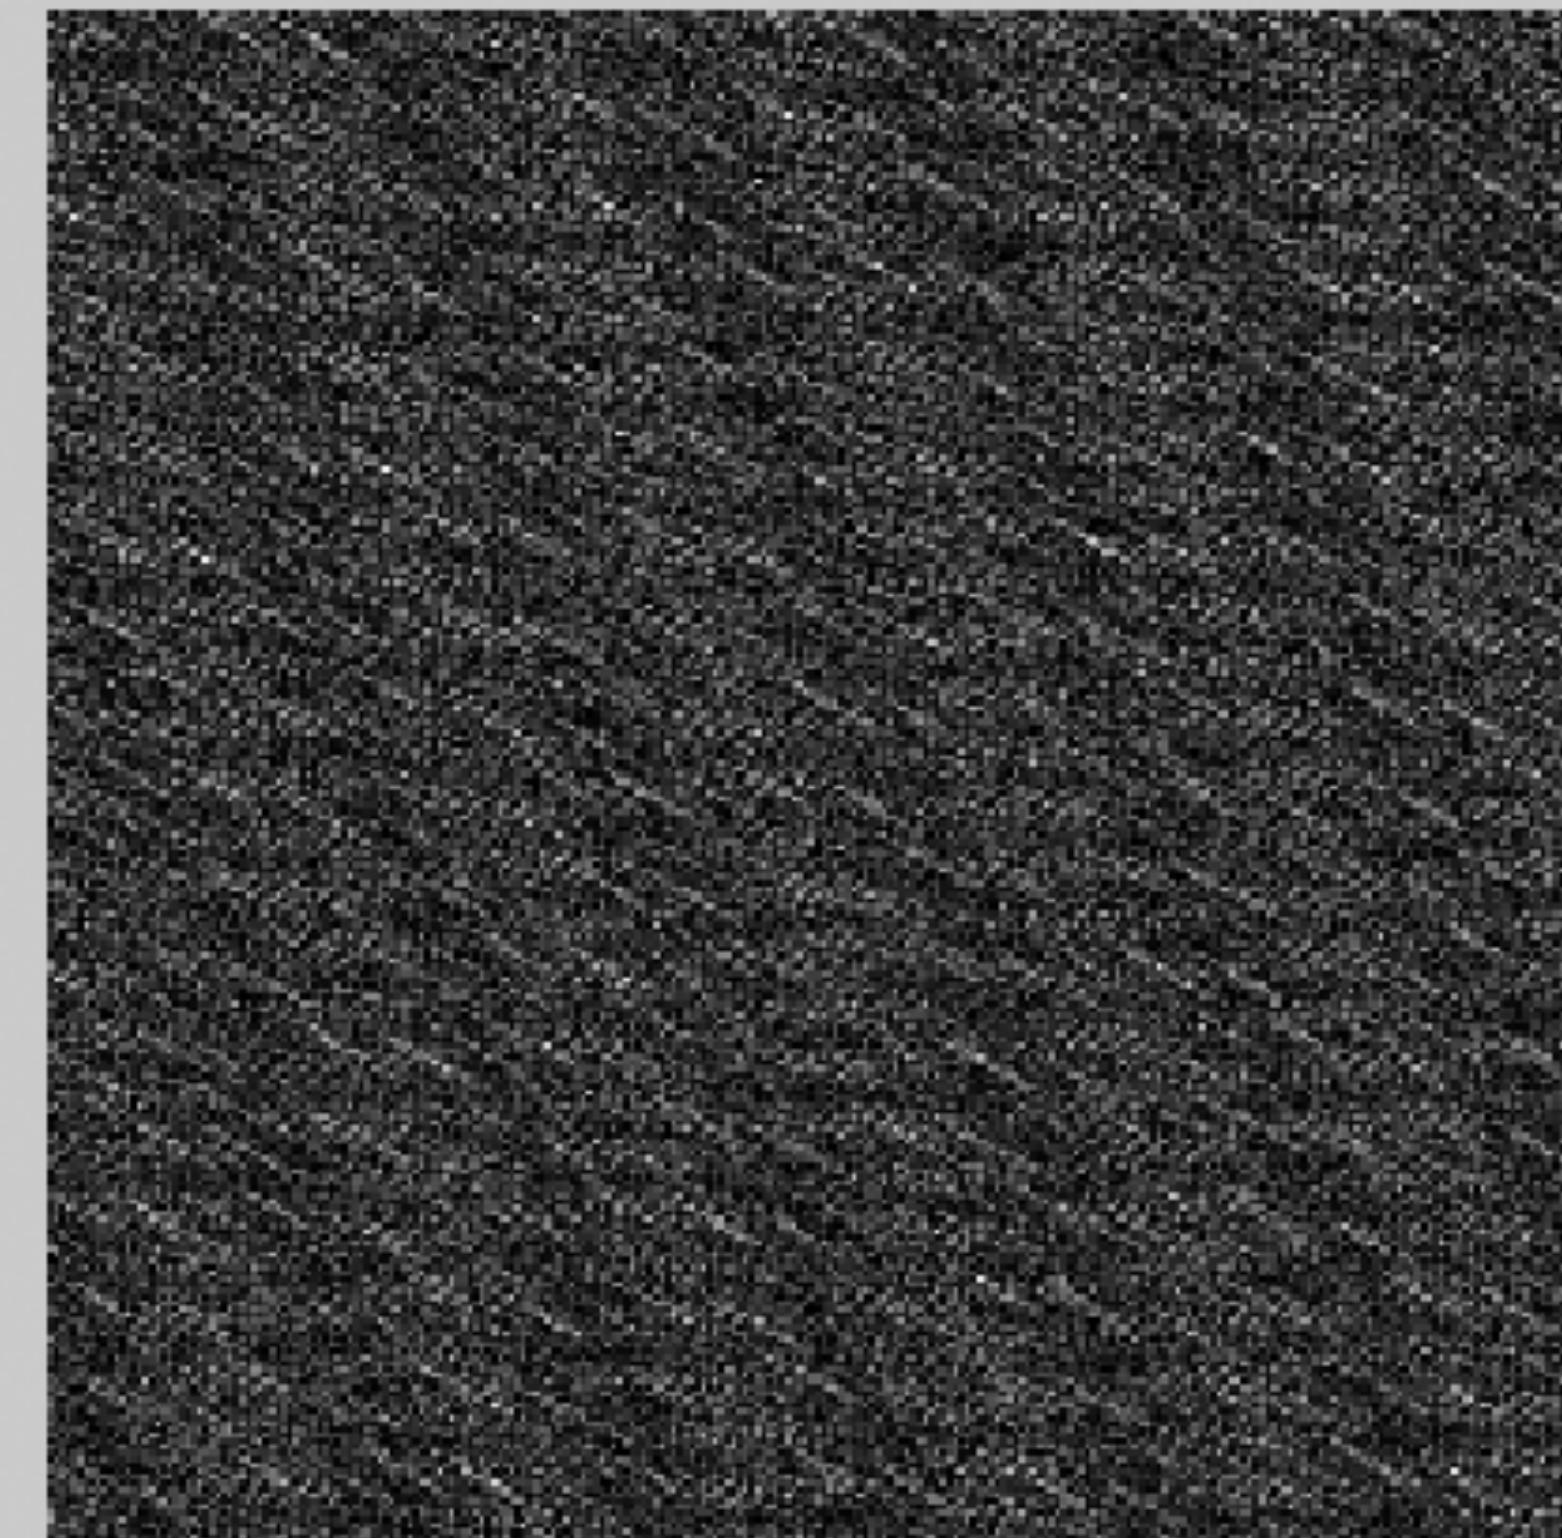
#2: Range [8.5e-006, 1.7]
Dims [256, 256]

82

82



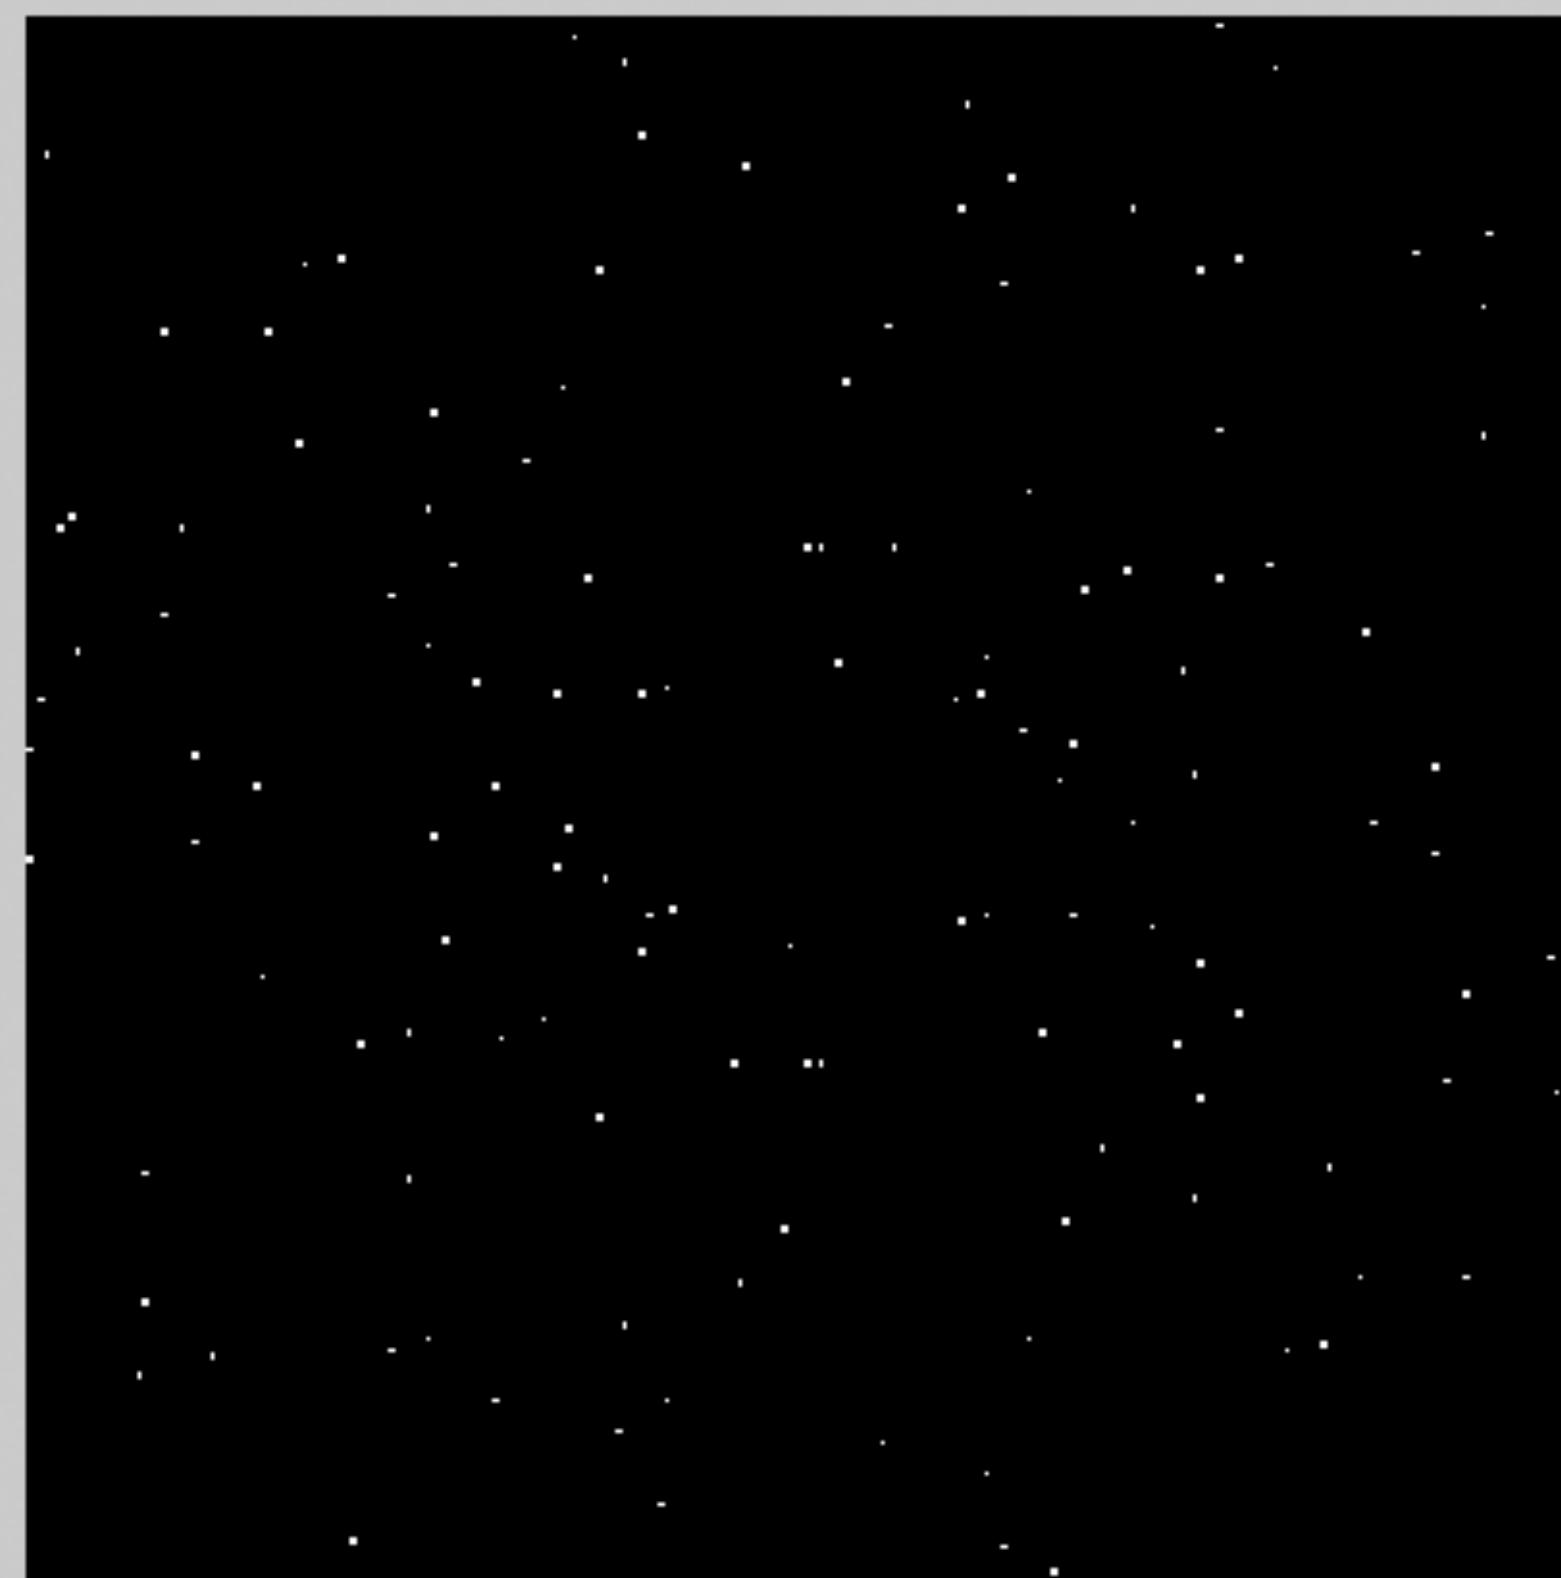
#1: Range [0, 1]
Dims [256, 256]



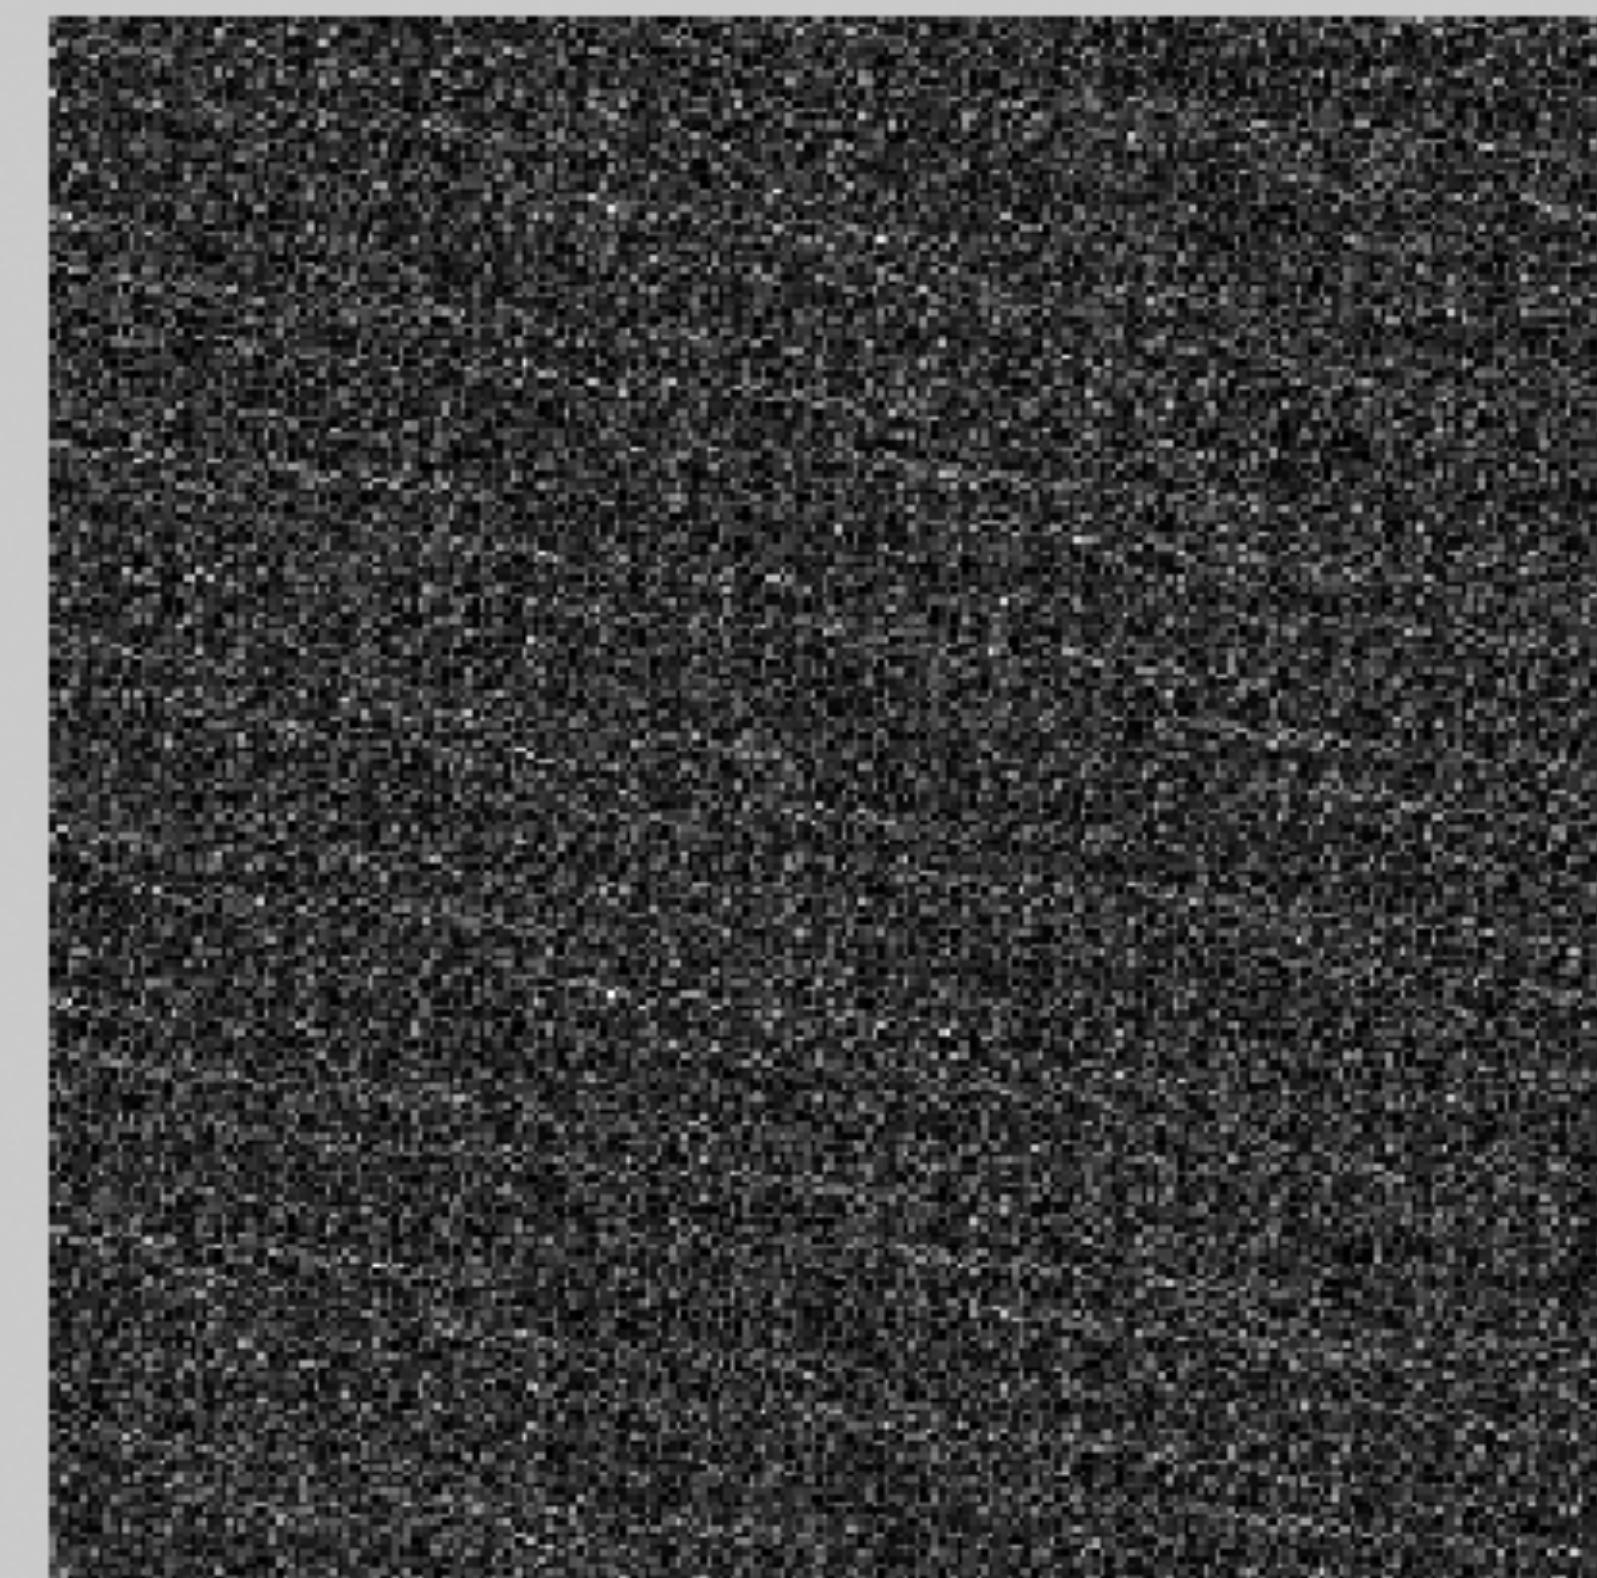
#2: Range [3.85e-007, 2.21]
Dims [256, 256]

136

136



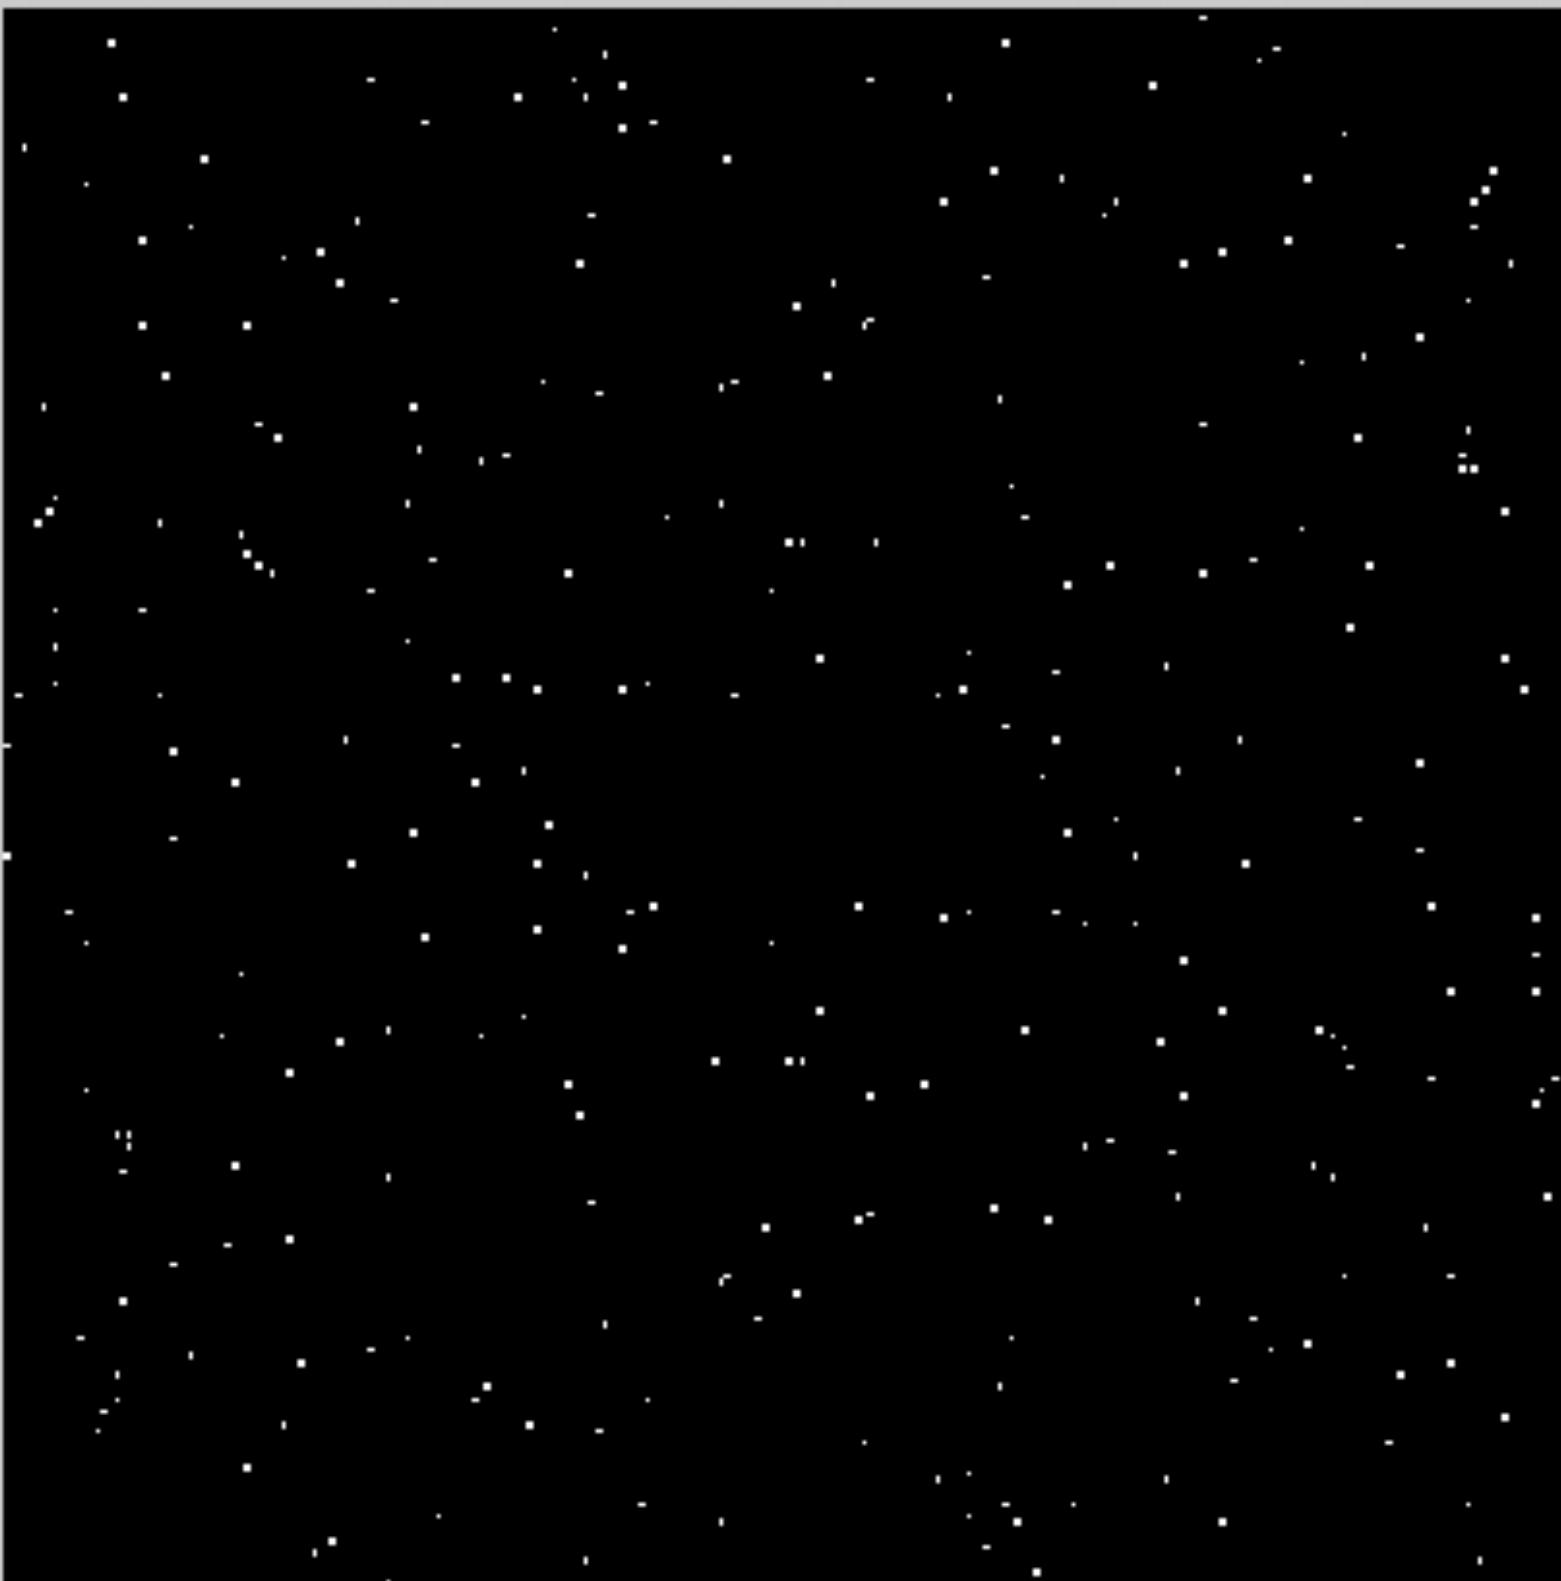
#1: Range [0, 1]
Dims [256, 256]



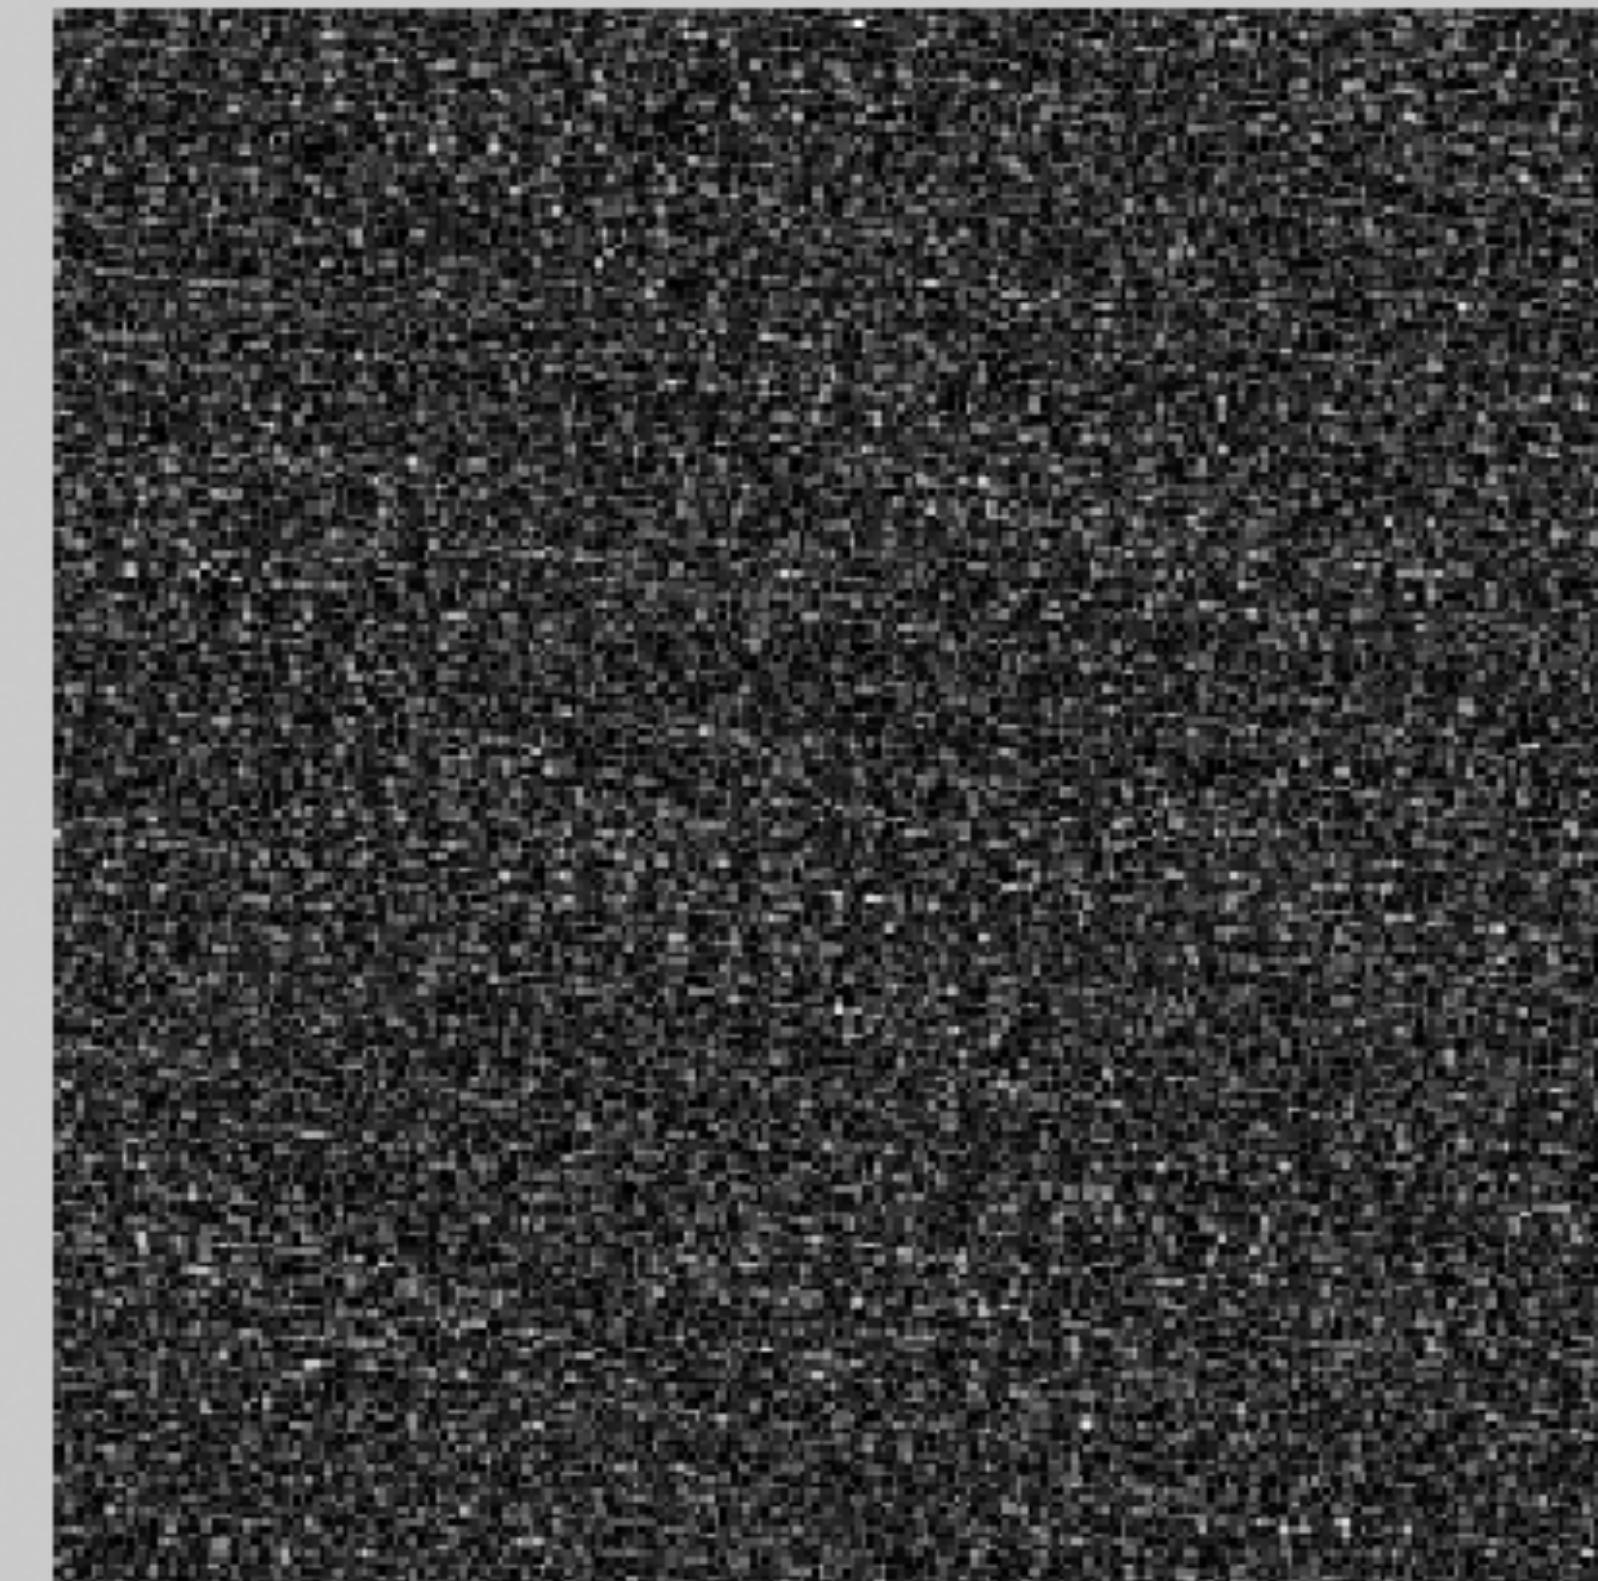
#2: Range [8.25e-006, 3.48]
Dims [256, 256]

282

282



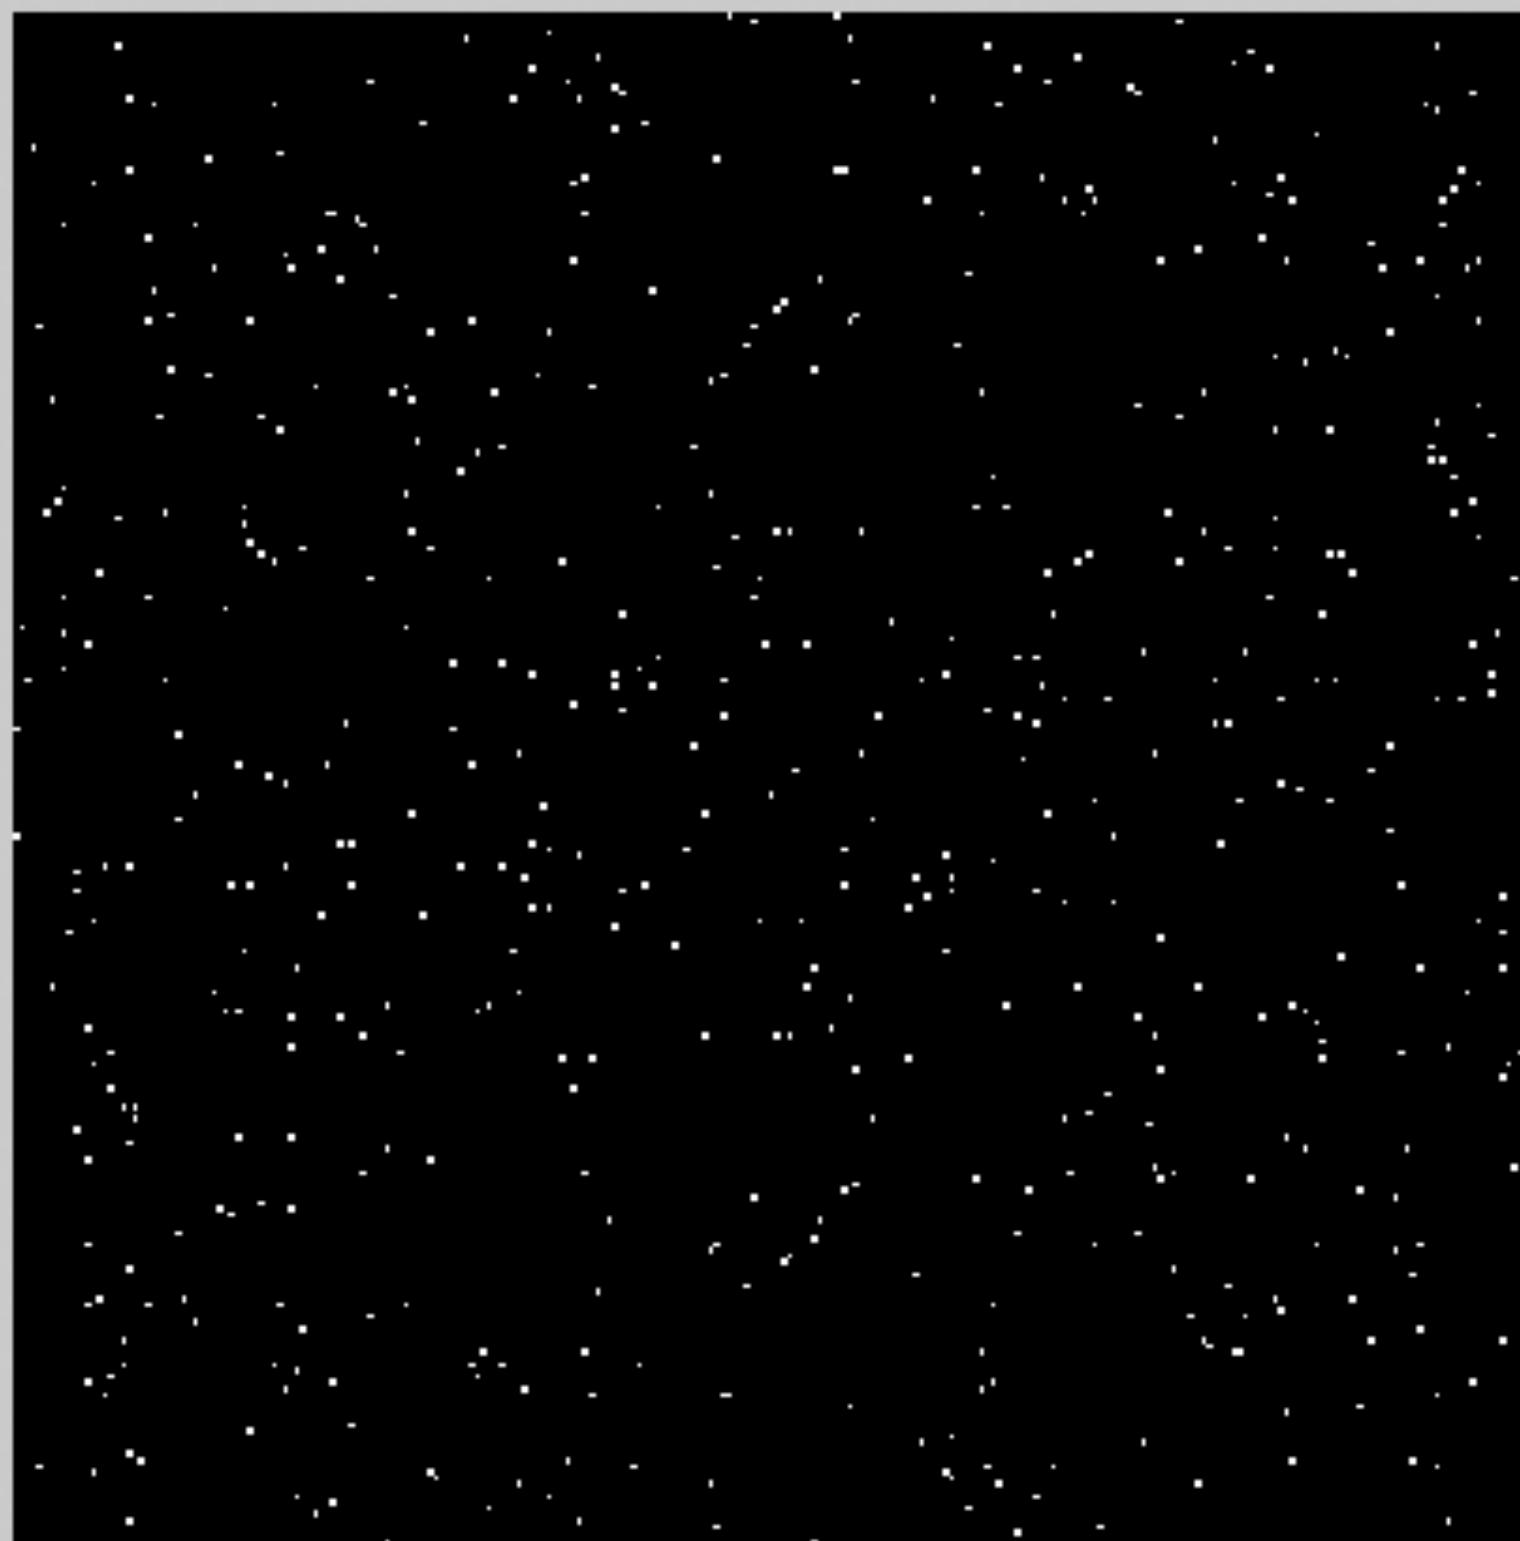
#1: Range [0, 1]
Dims [256, 256]



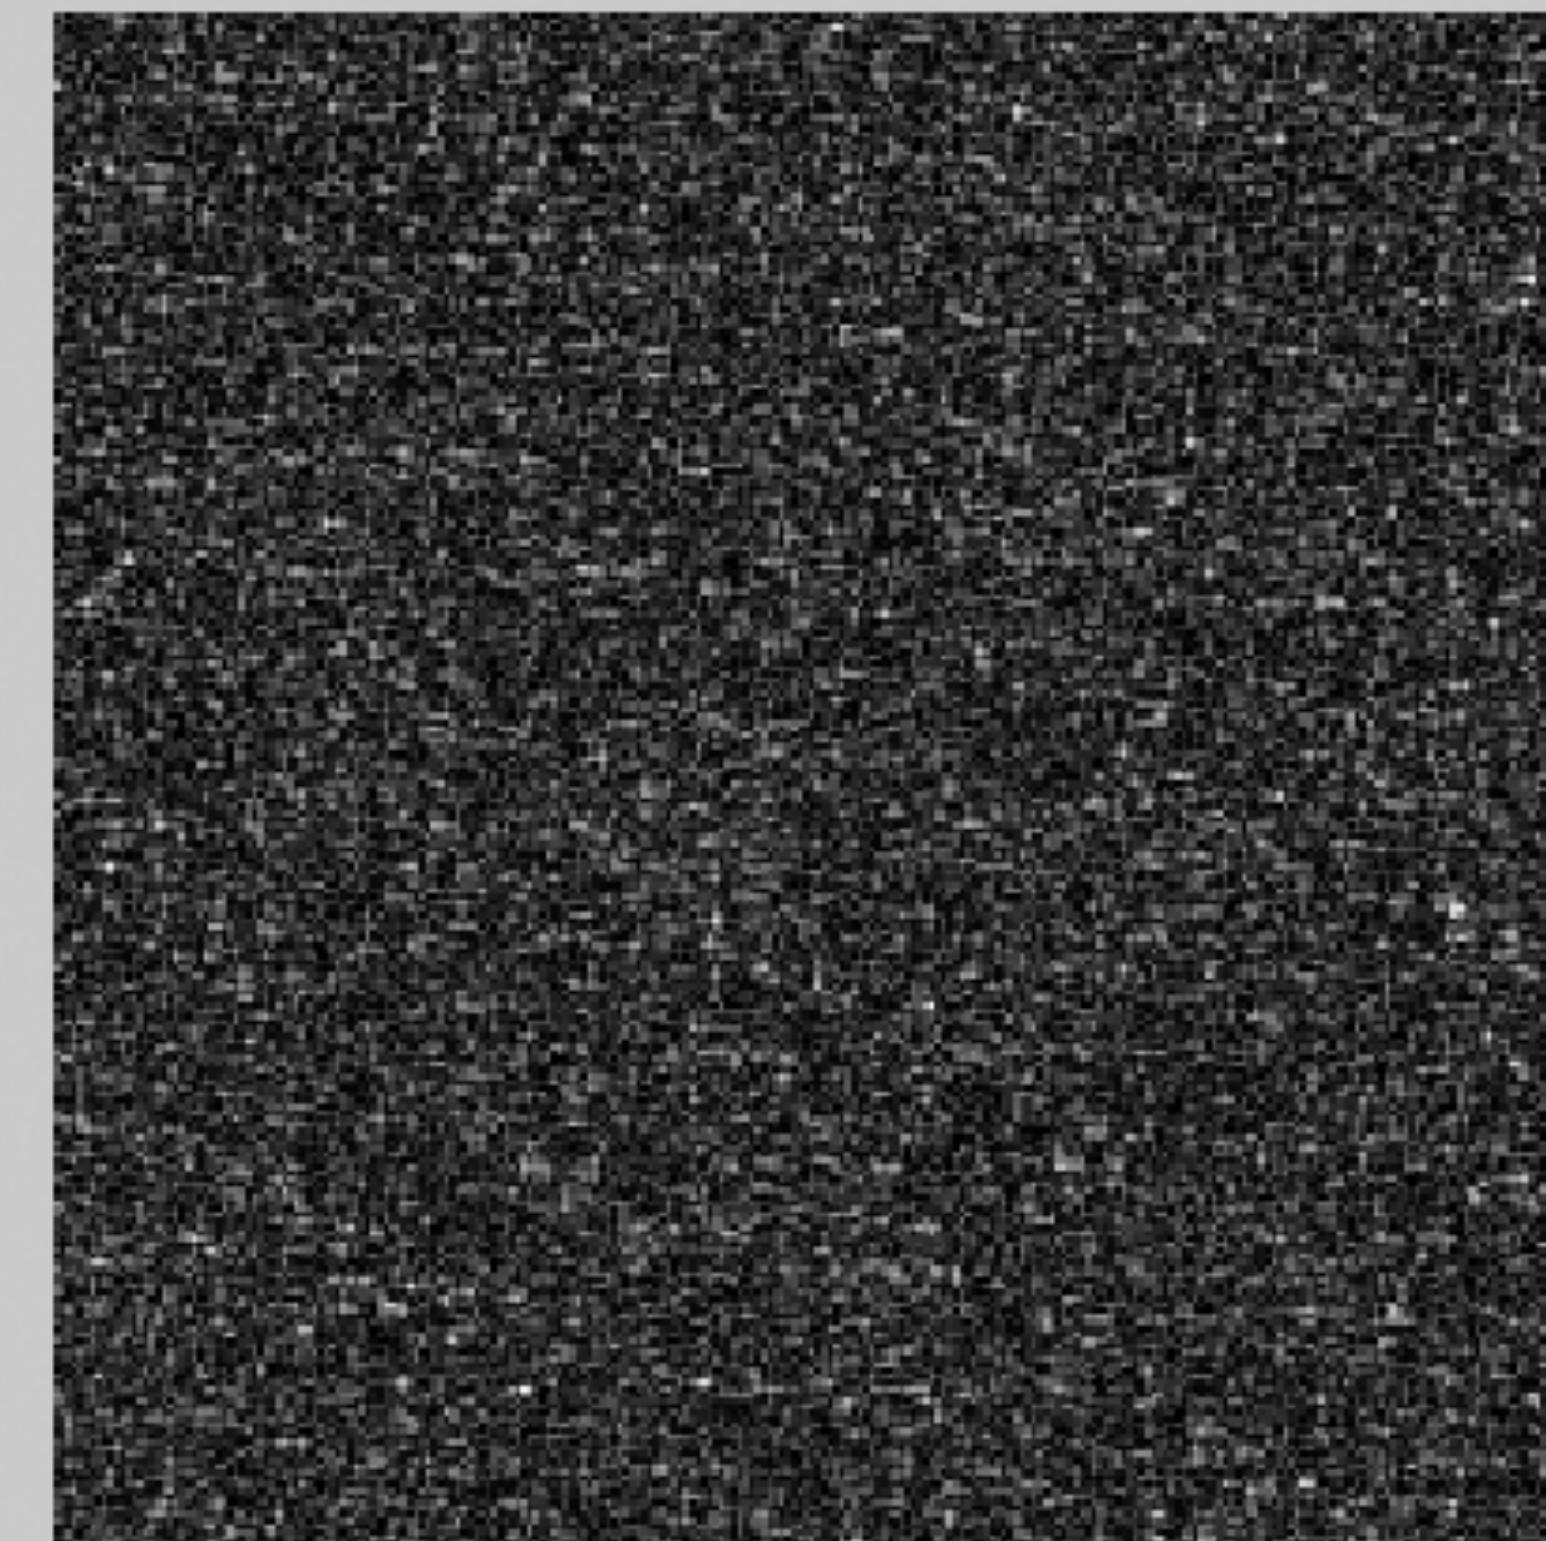
#2: Range [1.39e-005, 5.88]
Dims [256, 256]

538

538



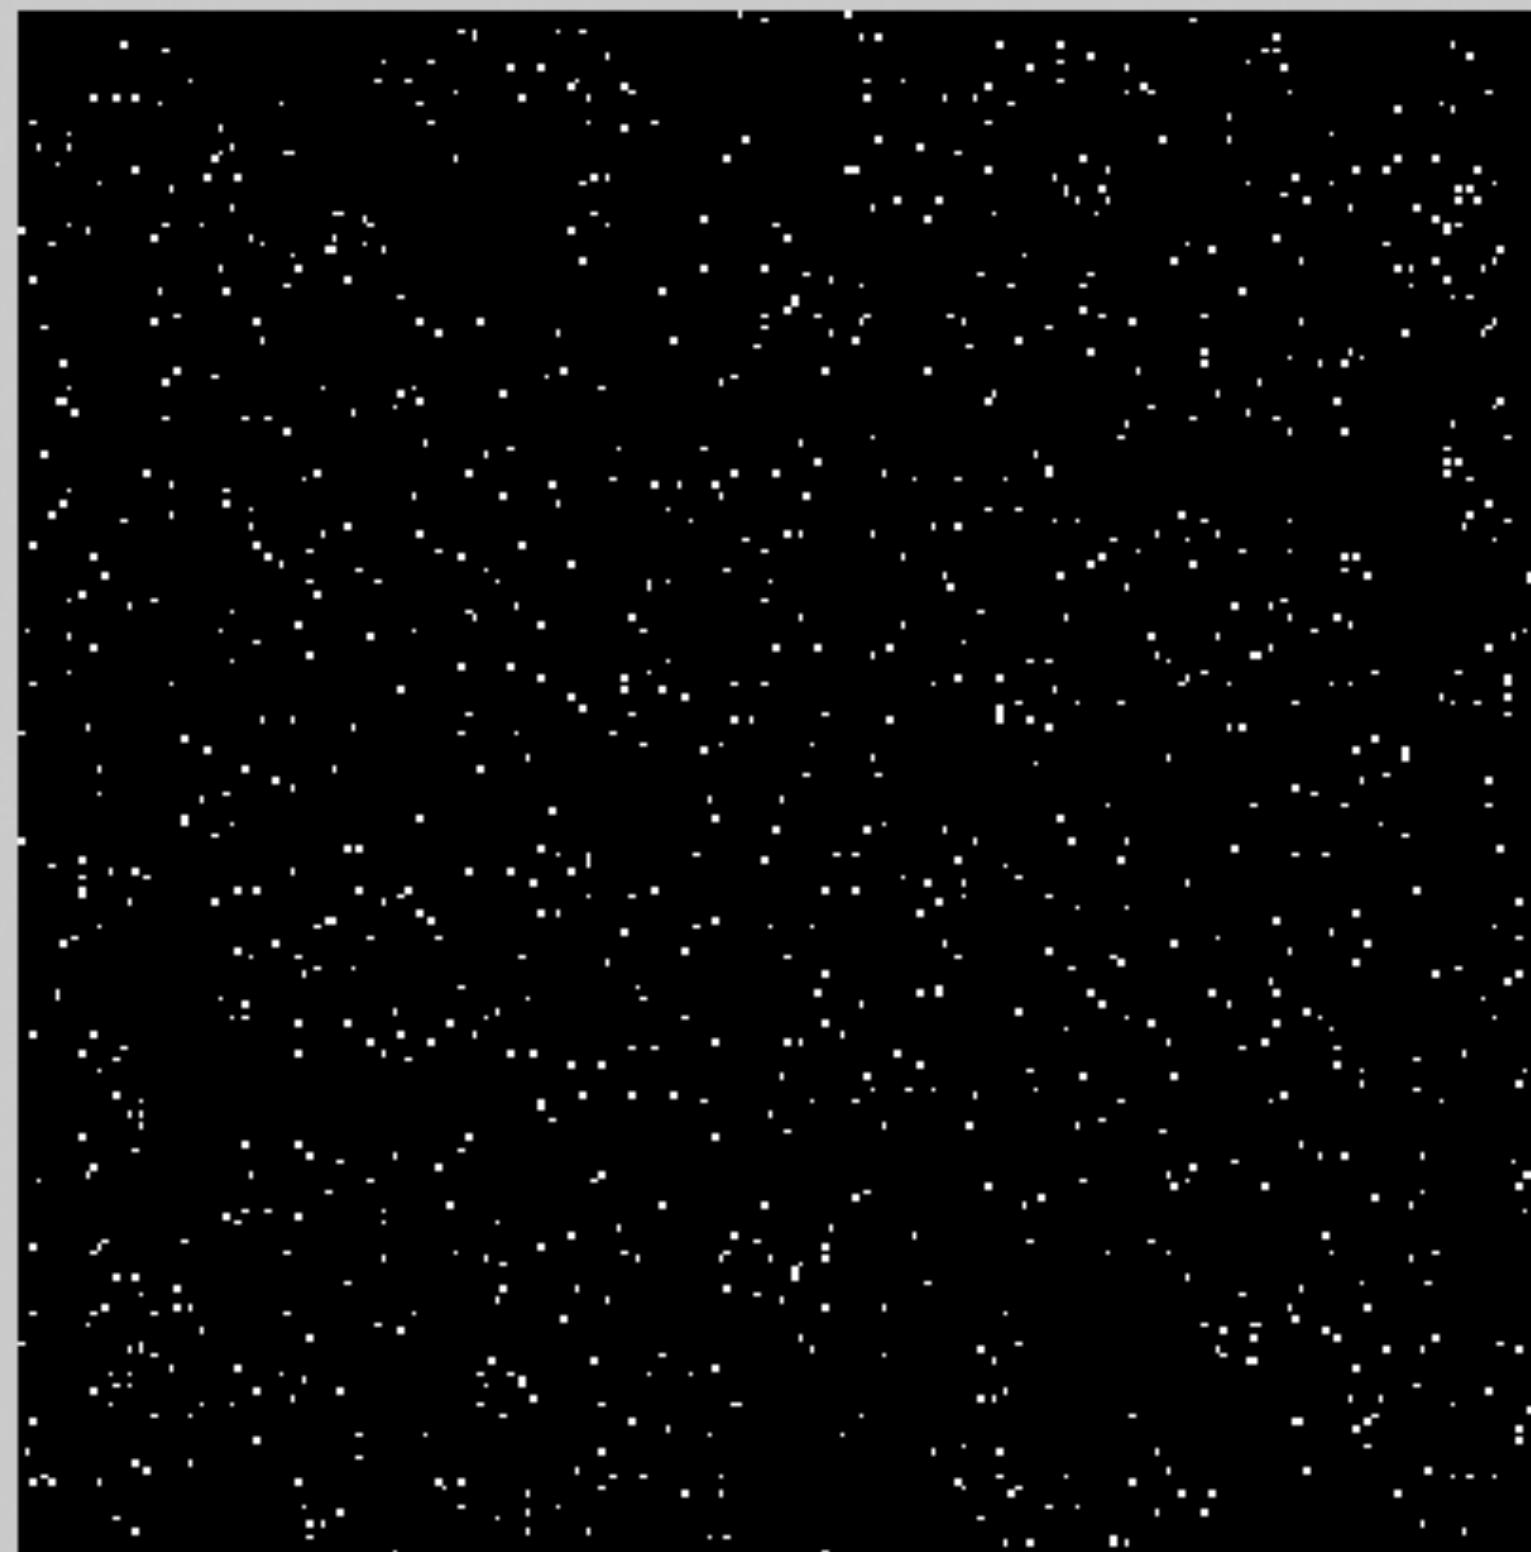
#1: Range [0, 1]
Dims [256, 256]



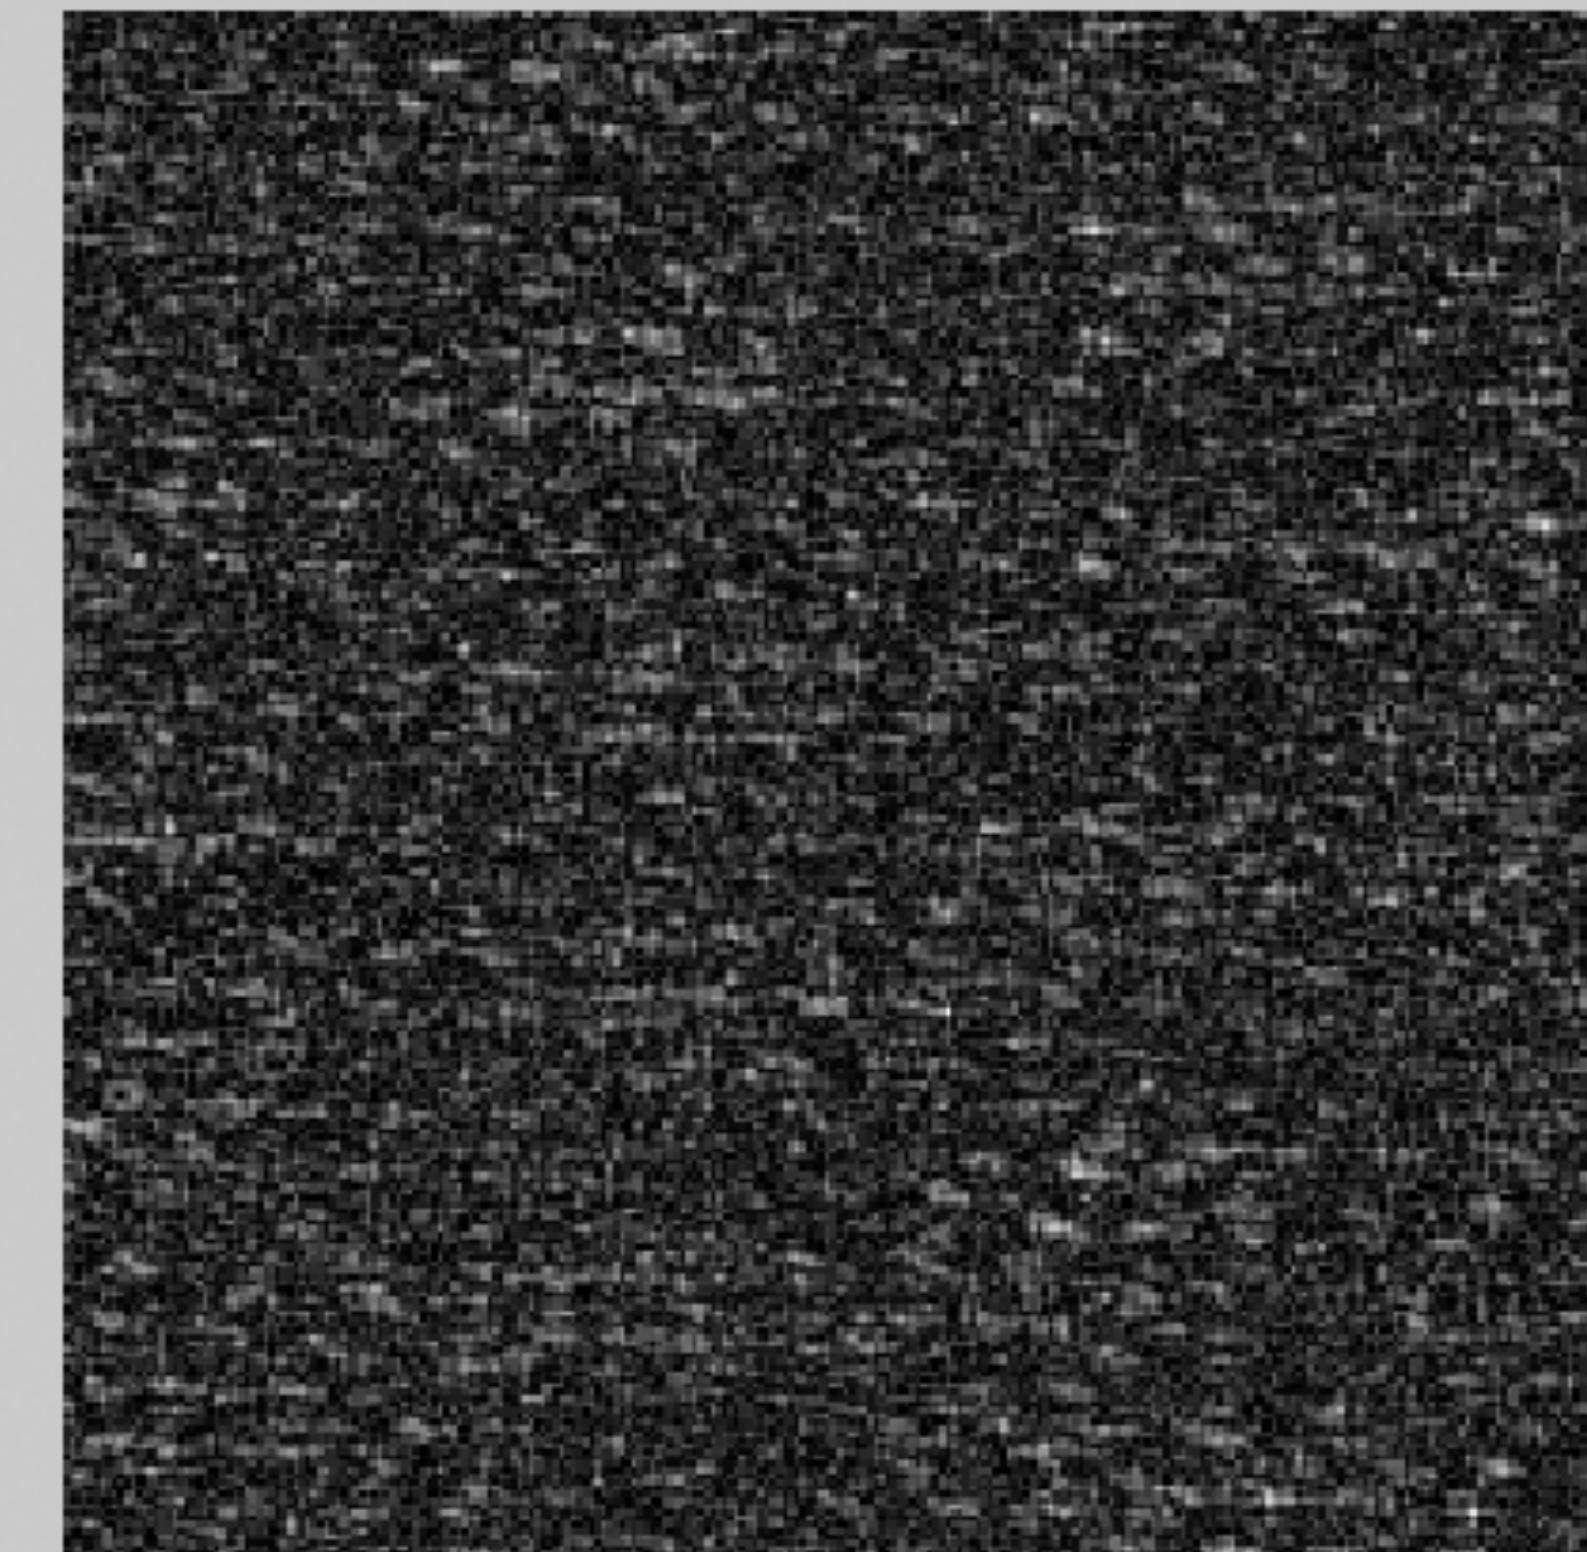
#2: Range [6.17e-006, 8.4]
Dims [256, 256]

1088

1088



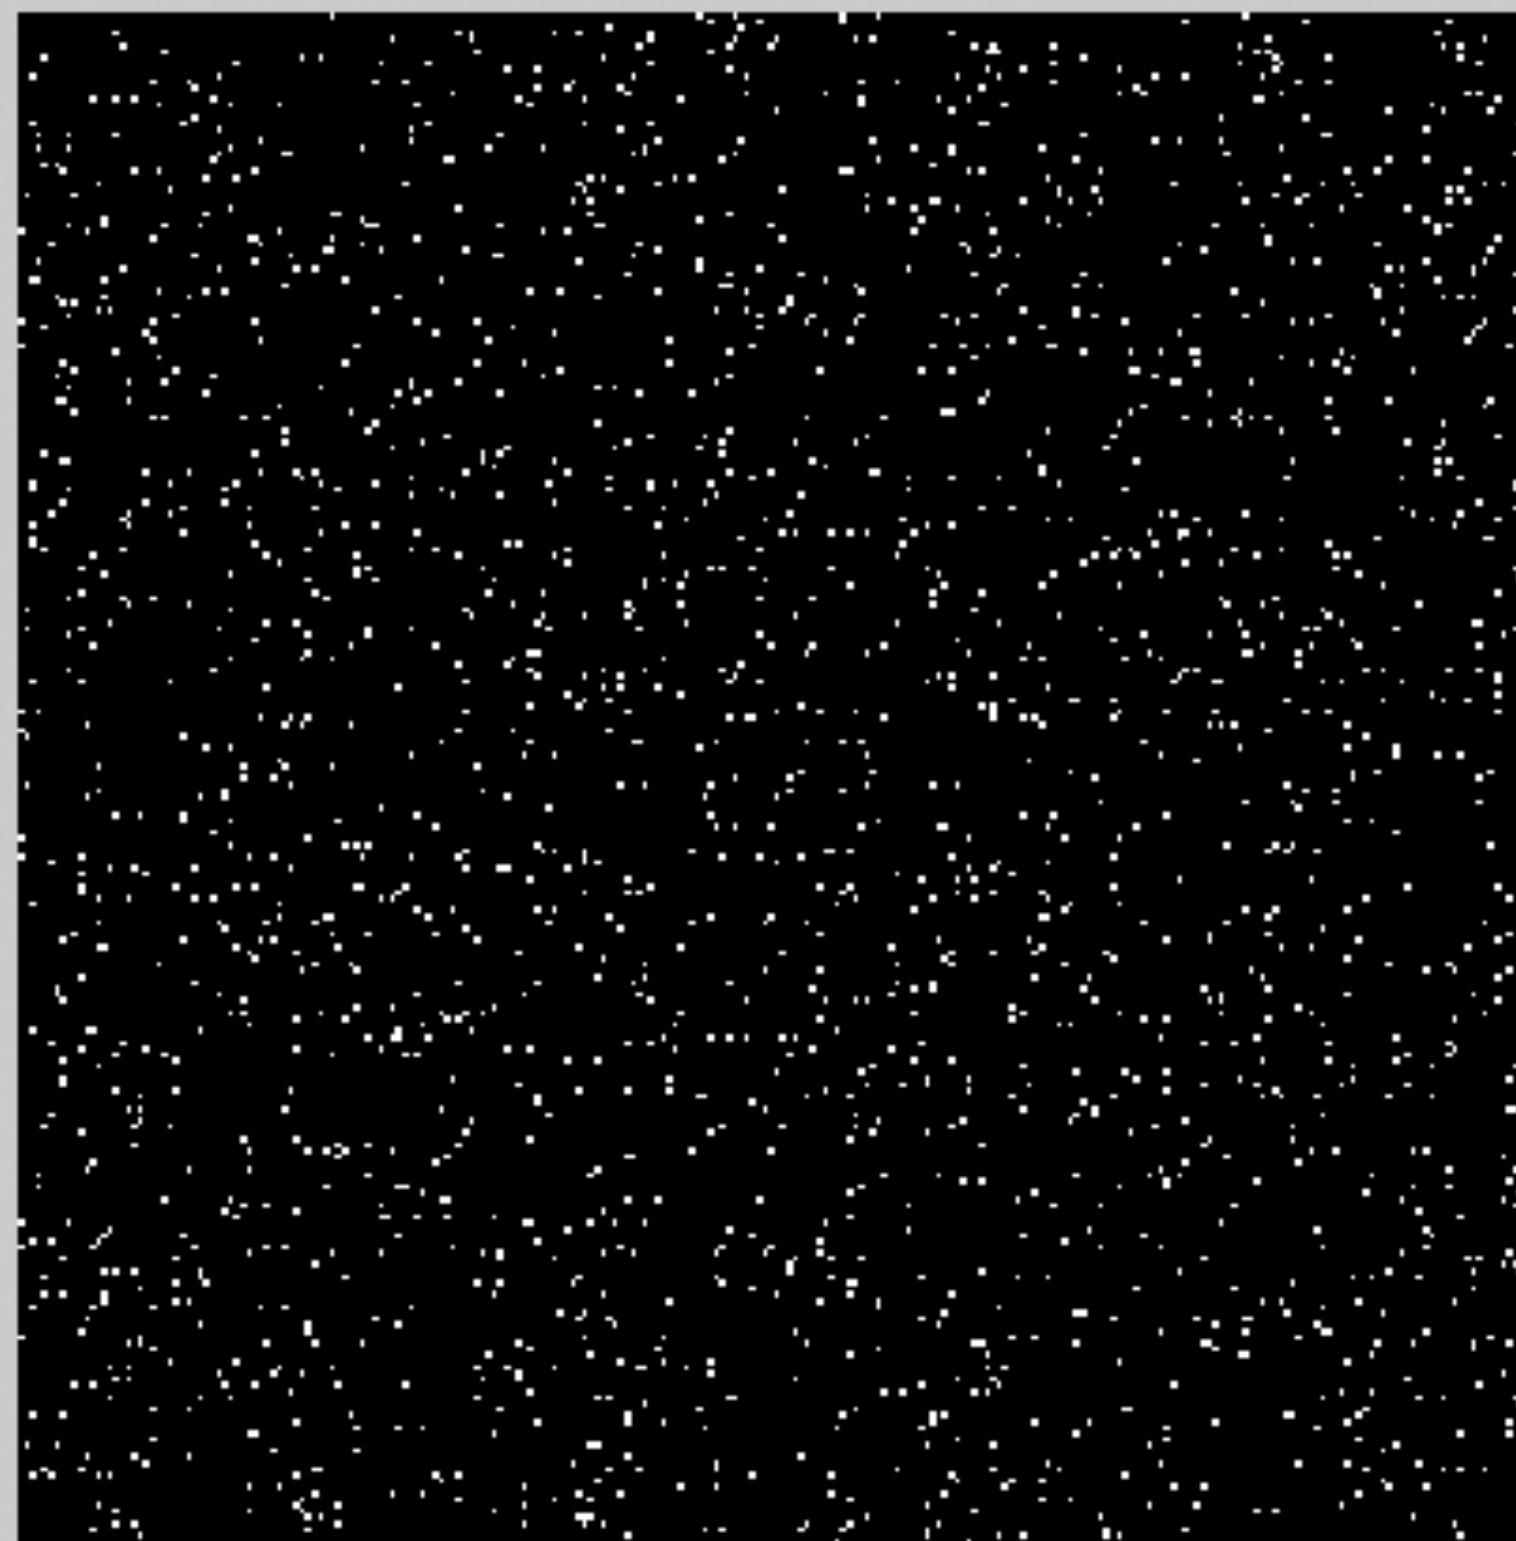
#1: Range [0, 1]
Dims [256, 256]



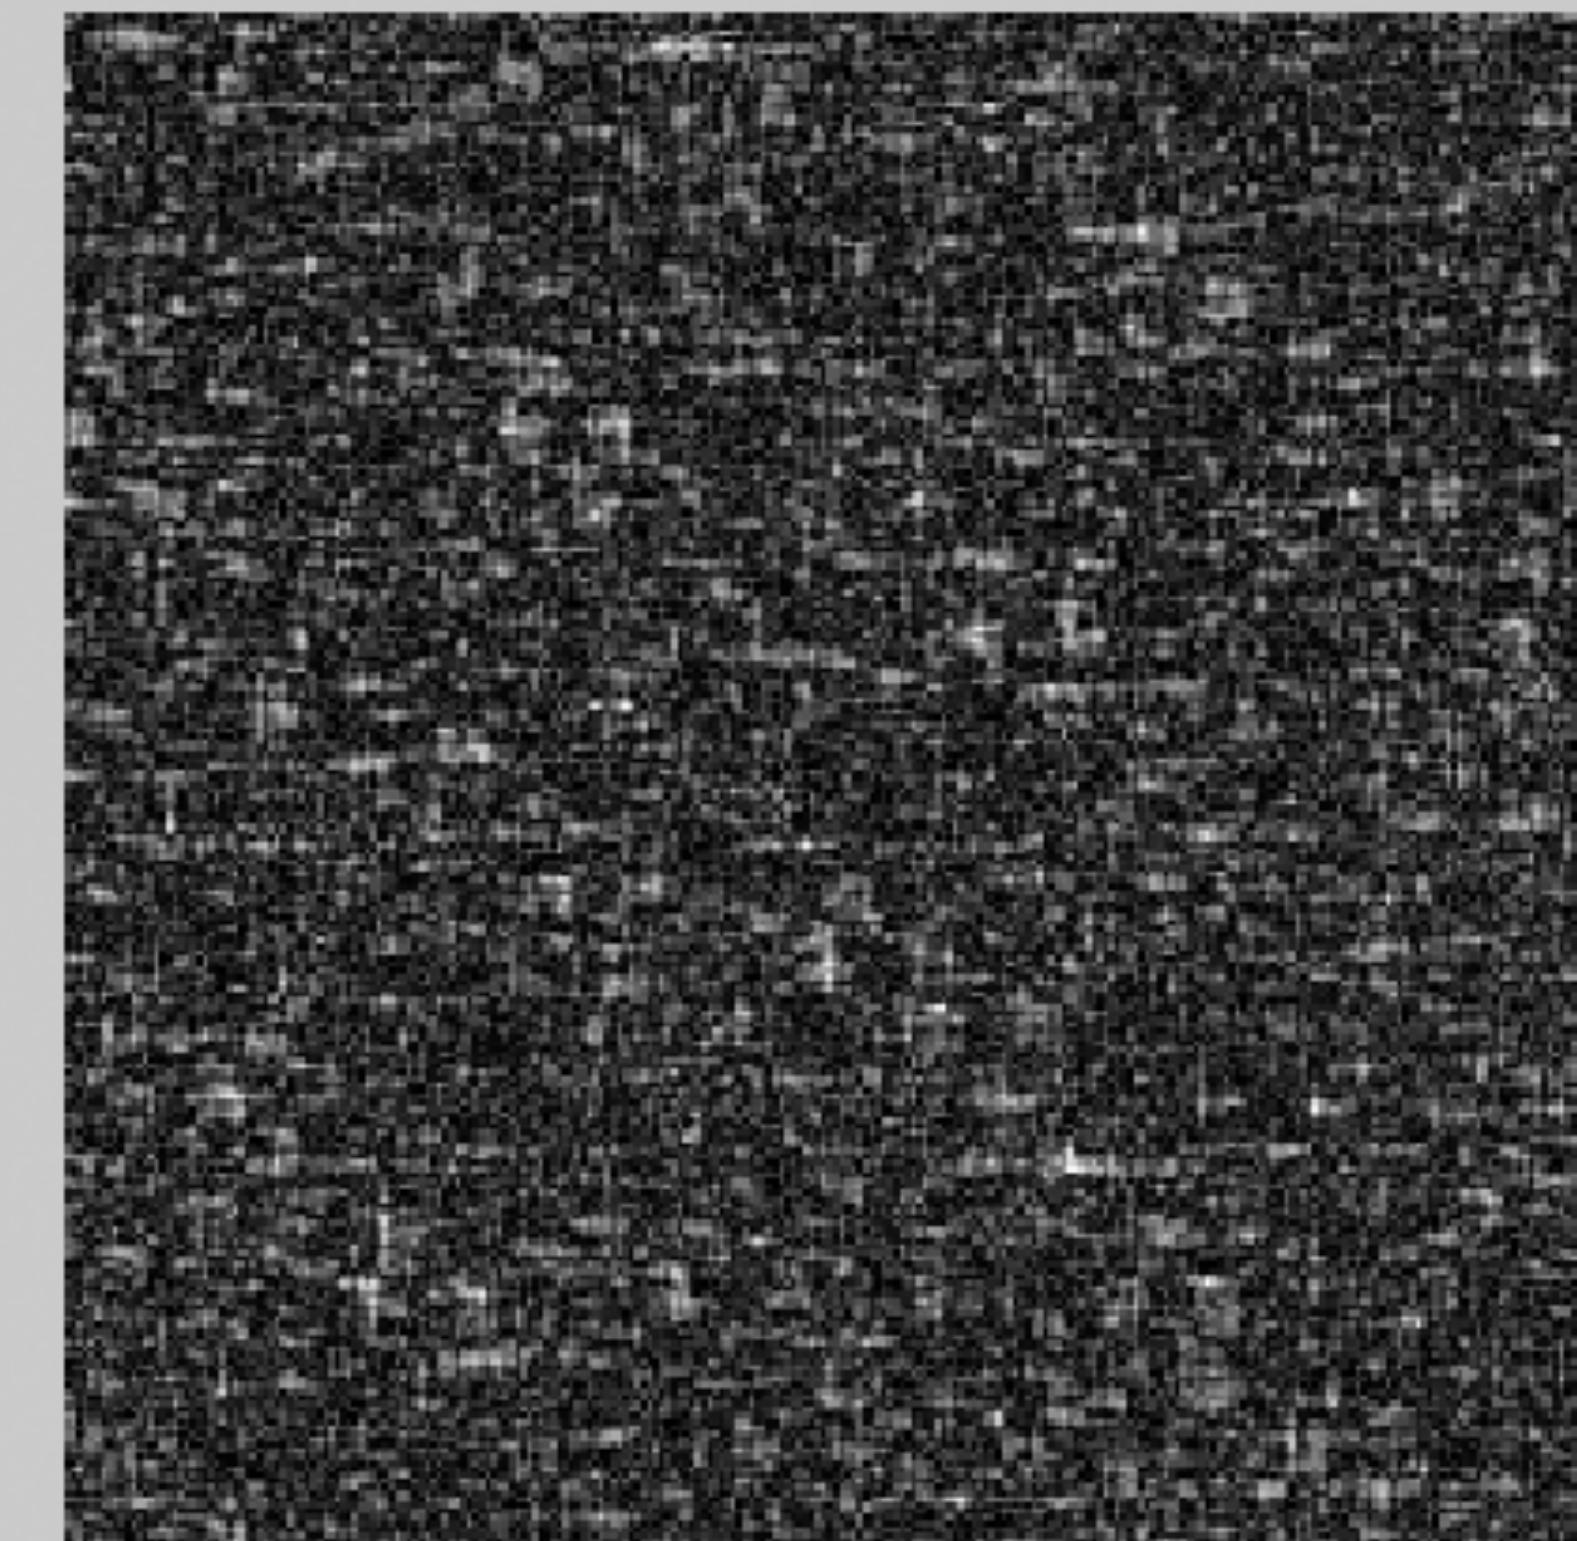
#2: Range [9.99e-005, 15]
Dims [256, 256]

2094

2094



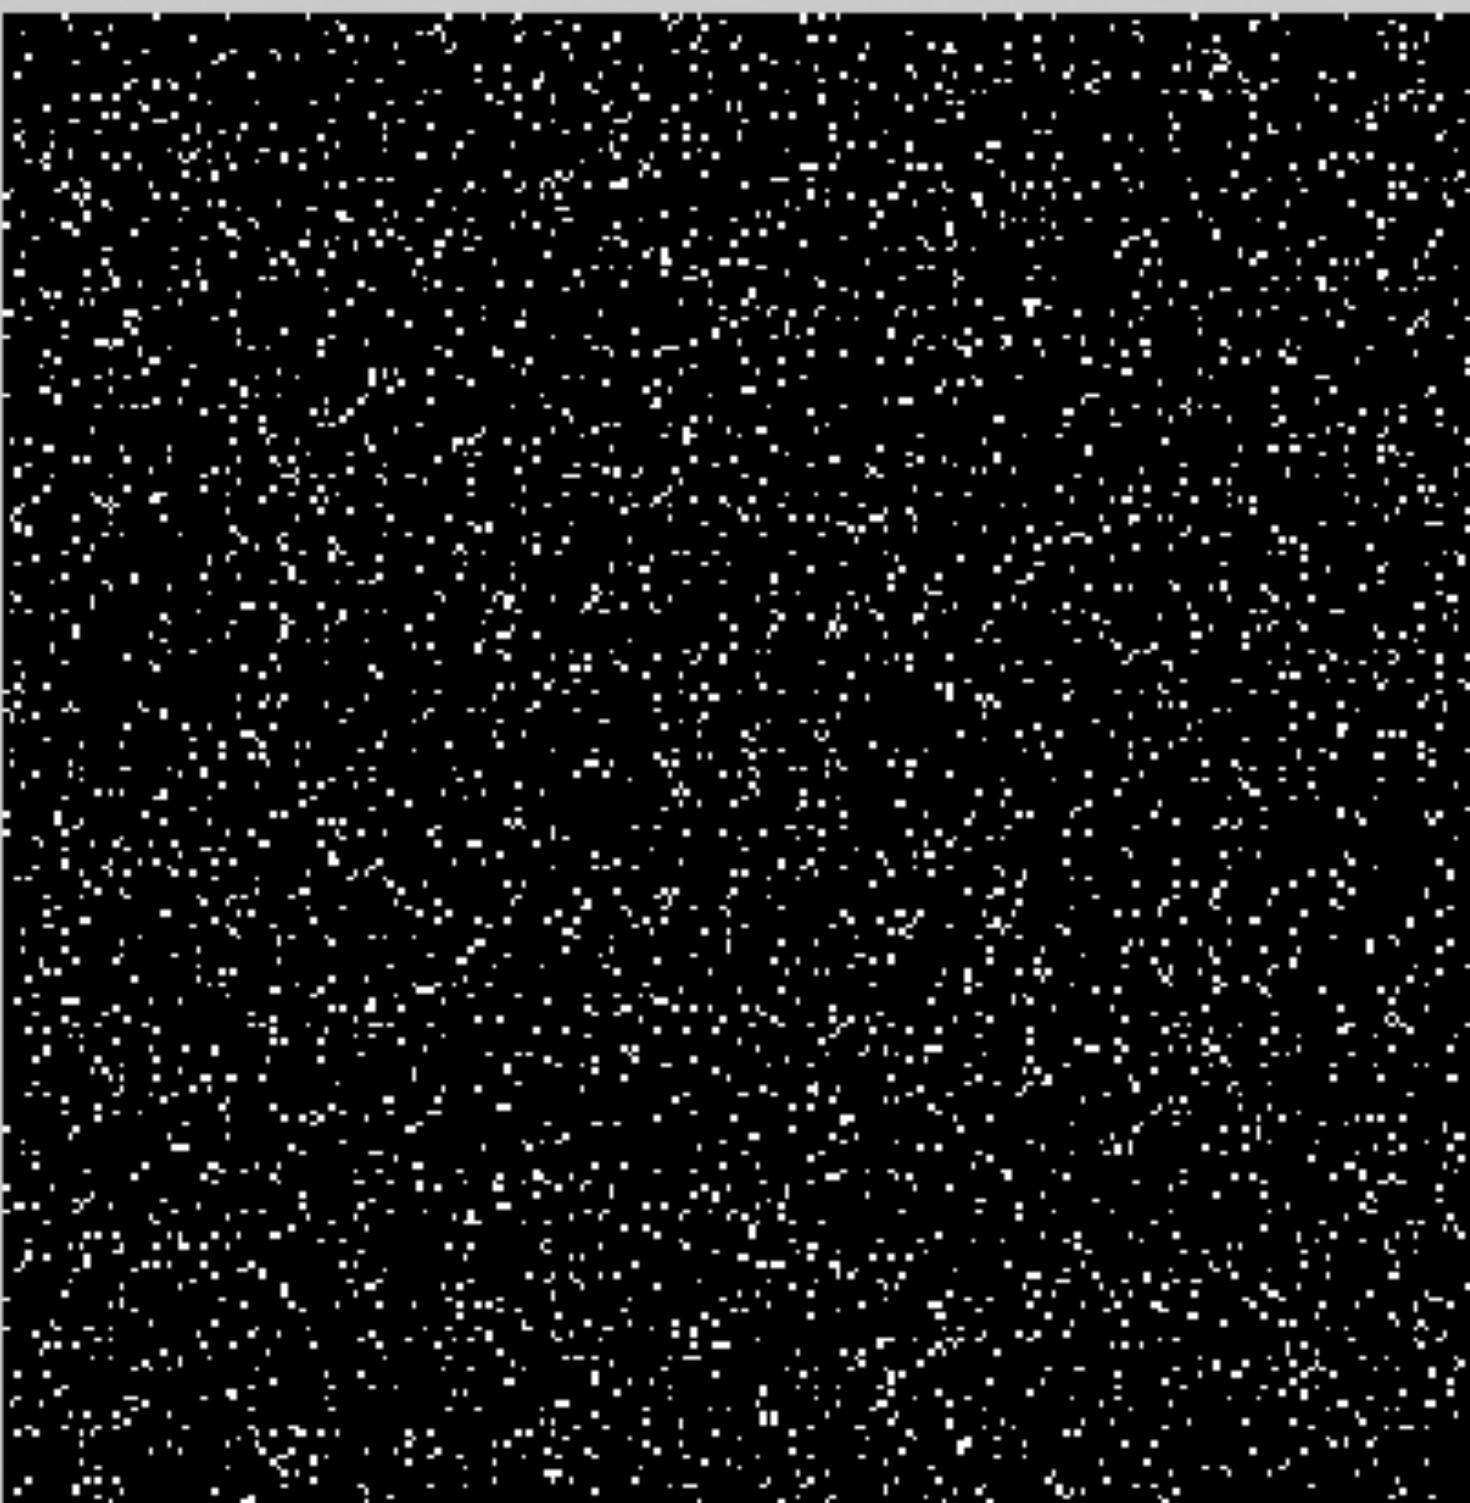
#1: Range [0, 1]
Dims [256, 256]



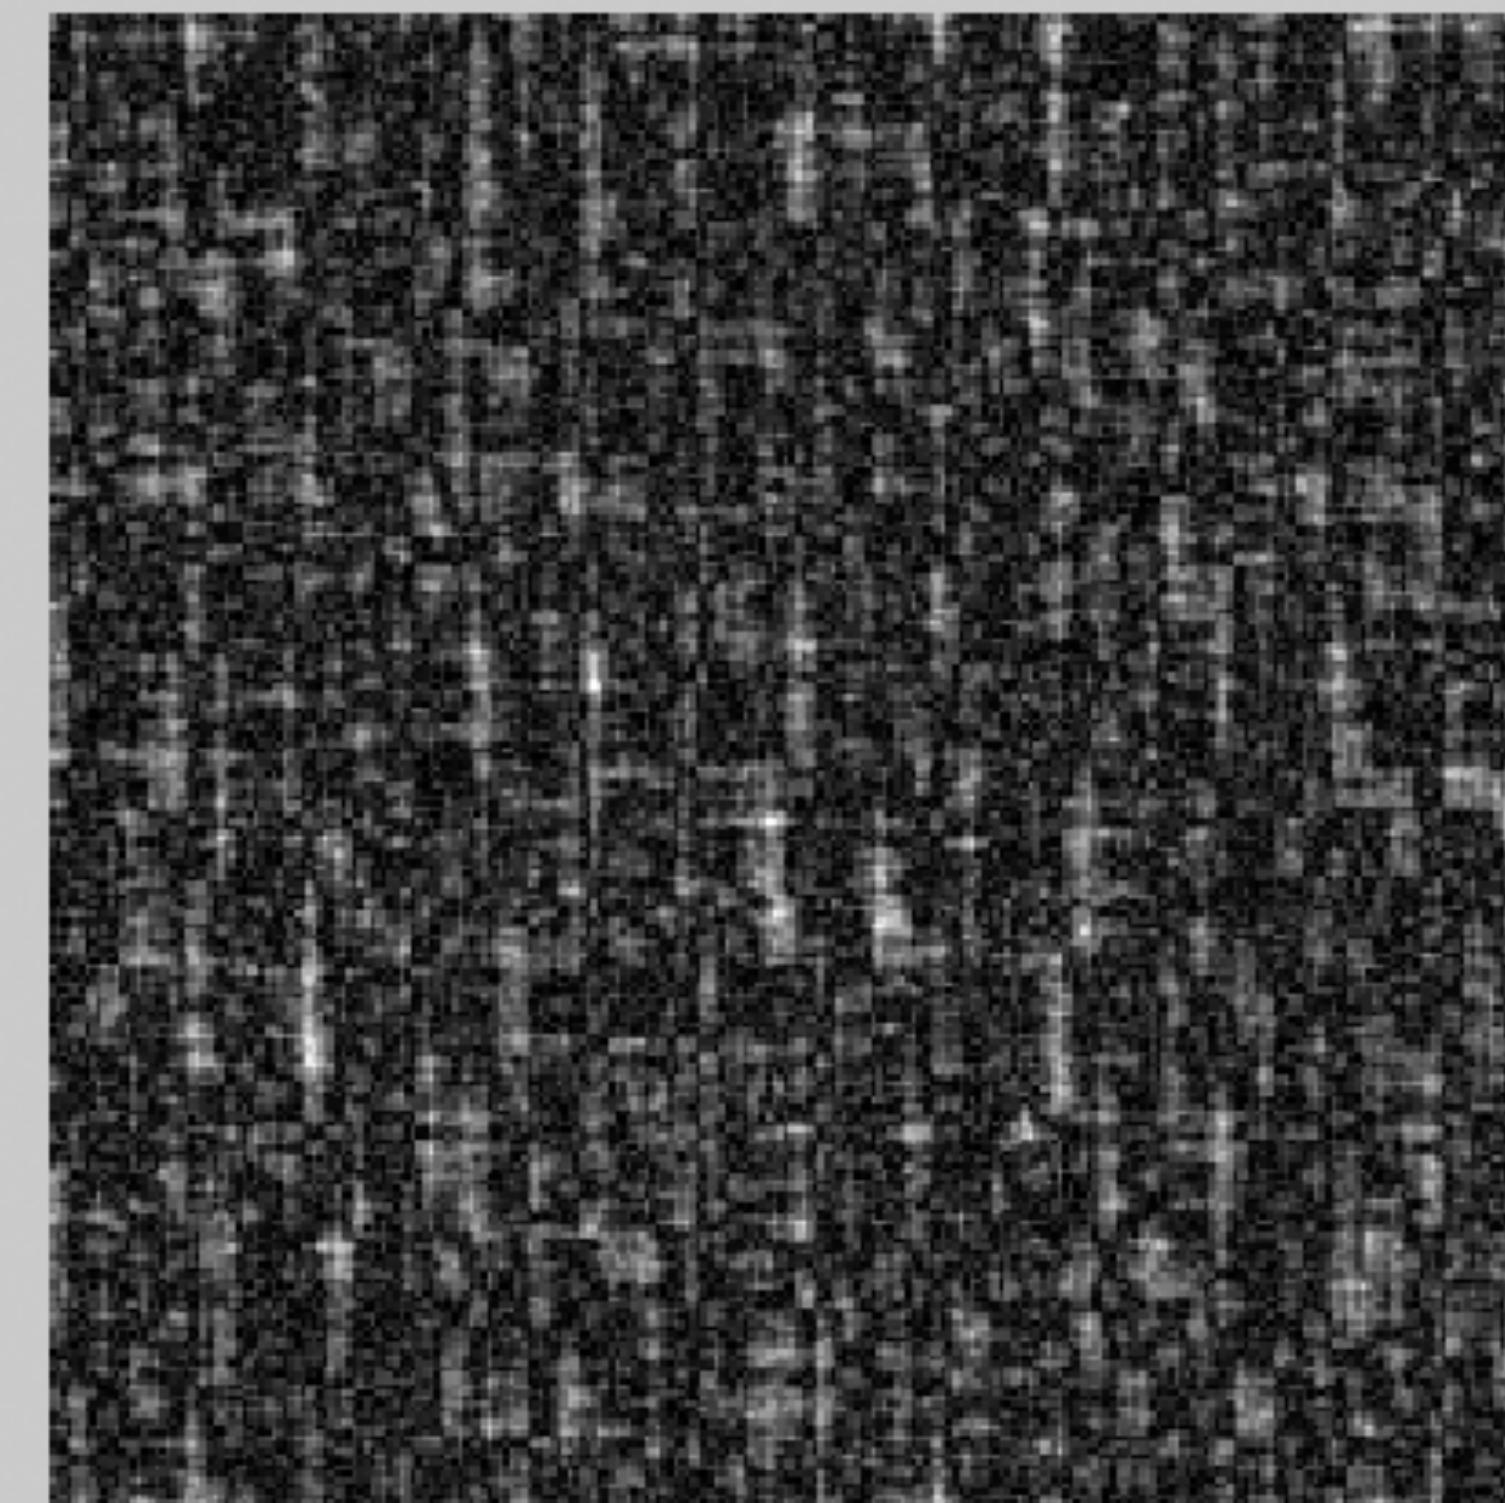
#2: Range [8.7e-005, 19]
Dims [256, 256]

4052.

4052



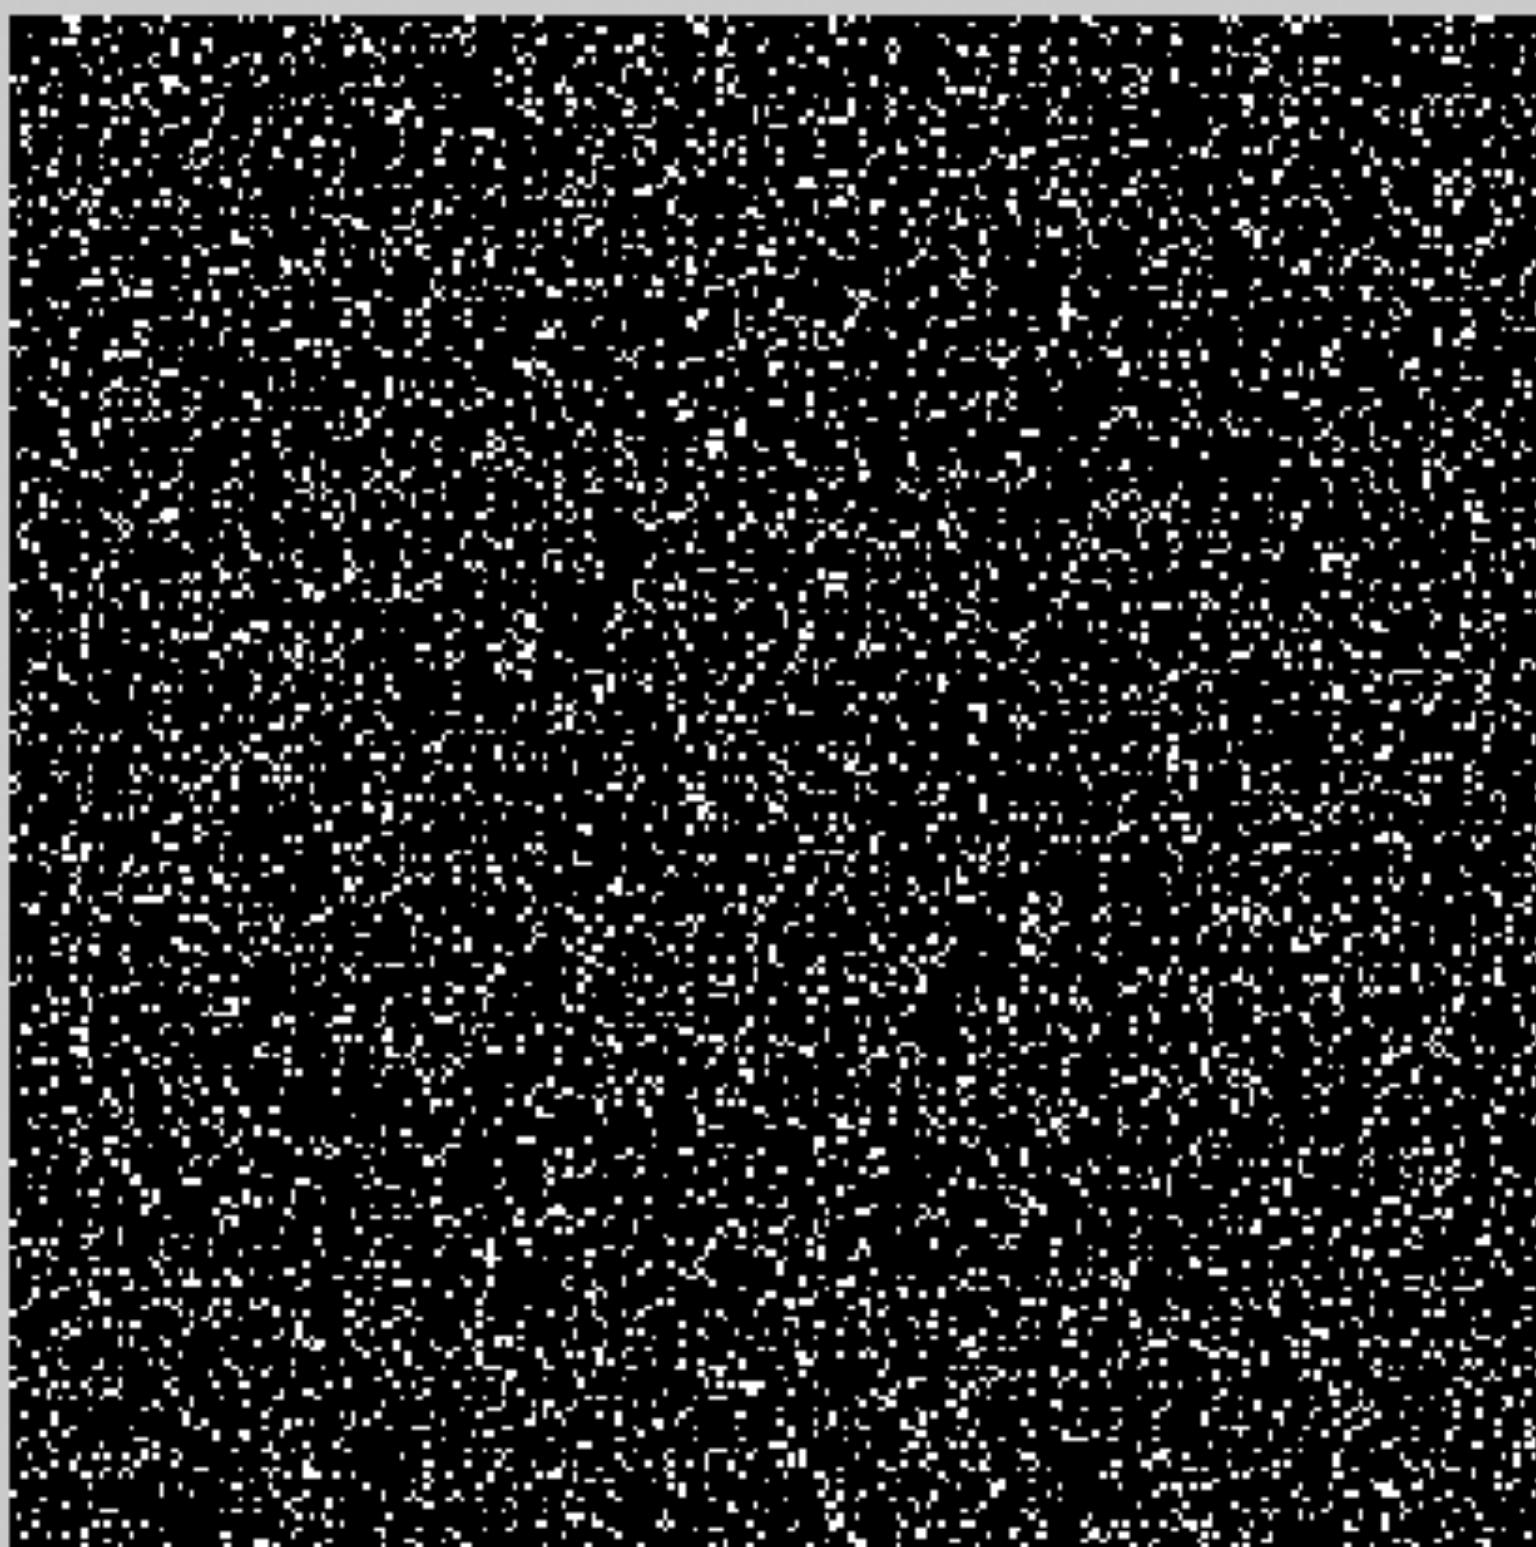
#1: Range [0, 1]
Dims [256, 256]



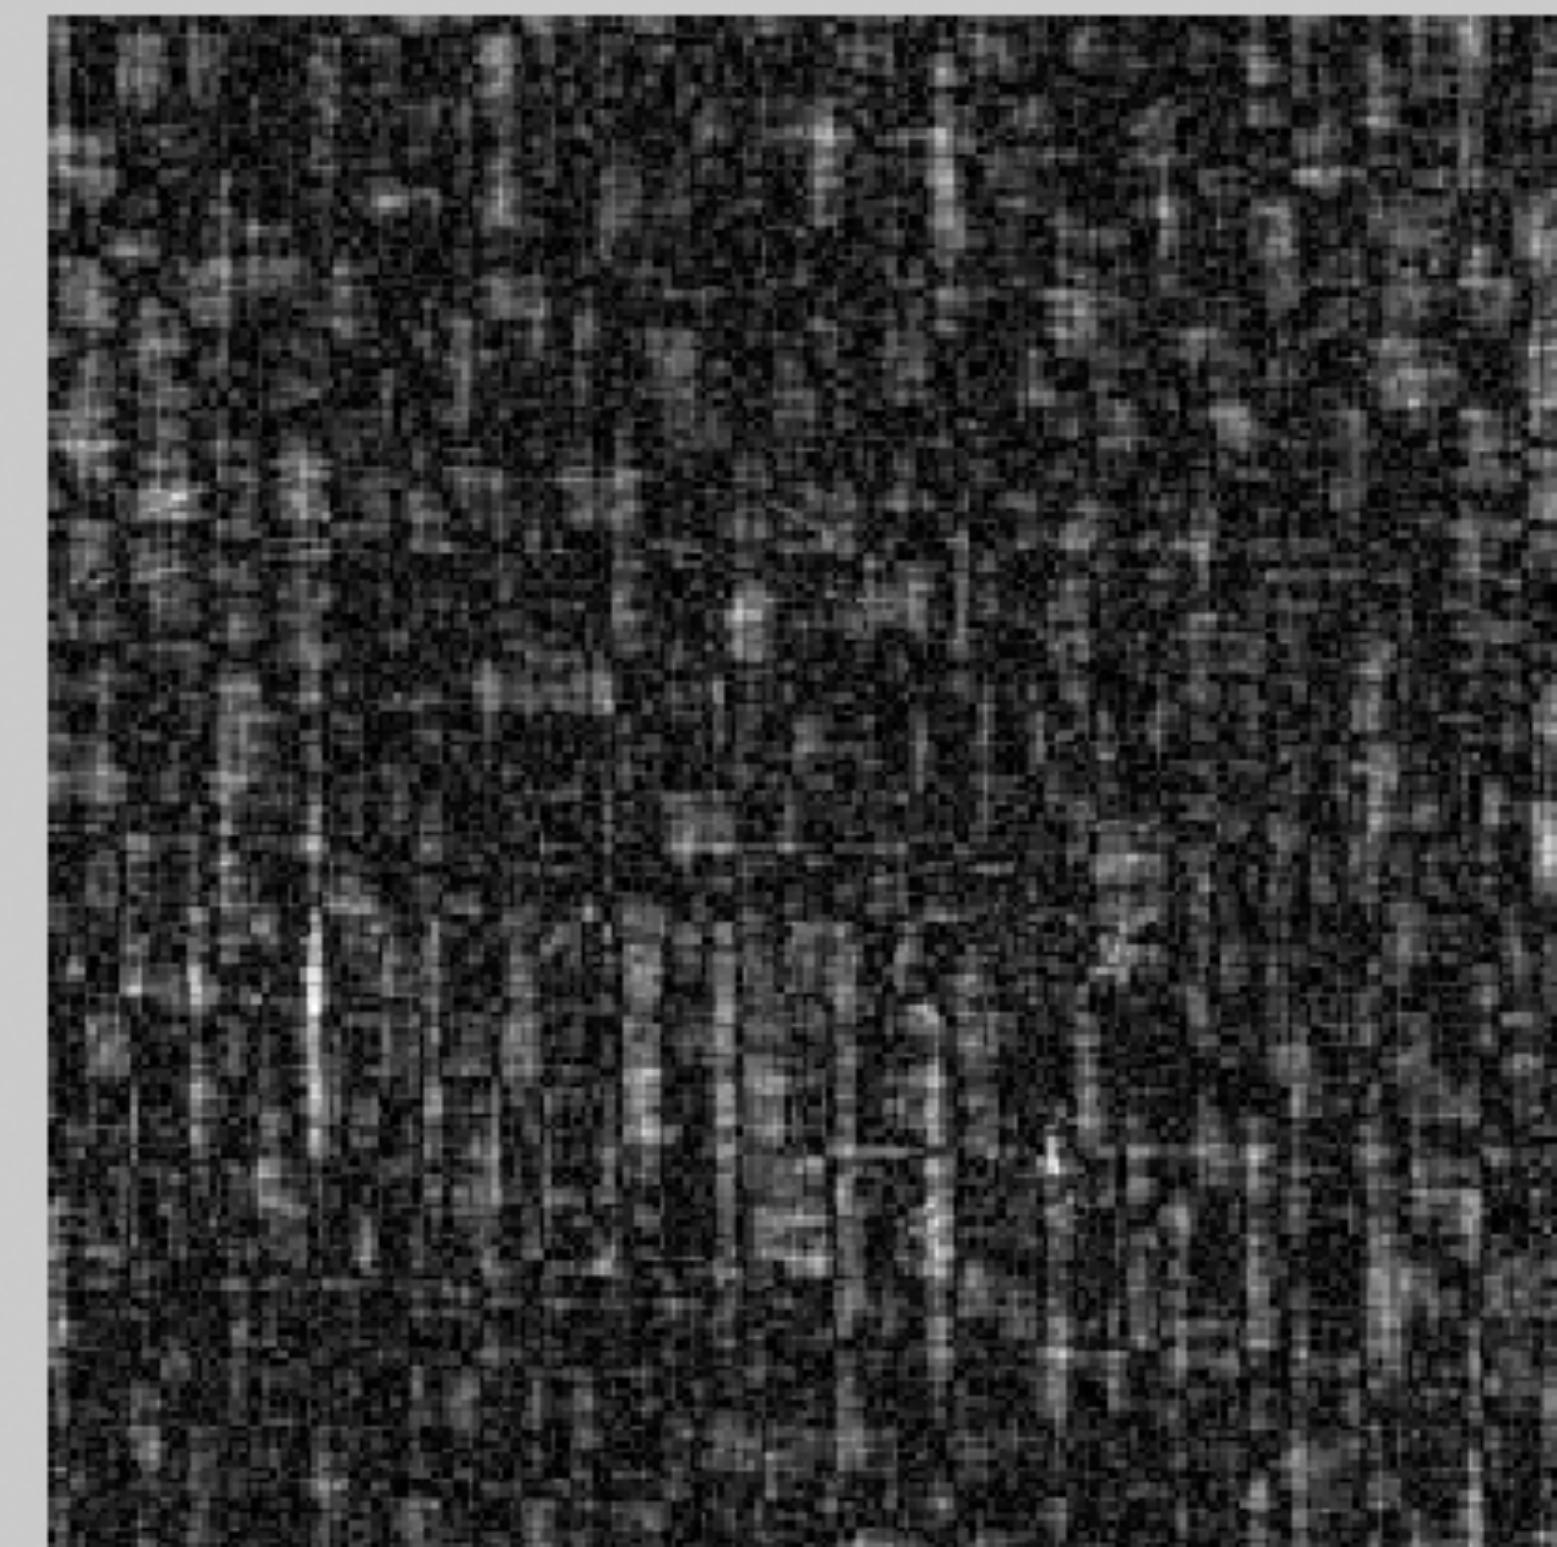
#2: Range [0.000556, 37.7]
Dims [256, 256]

8056.

8056



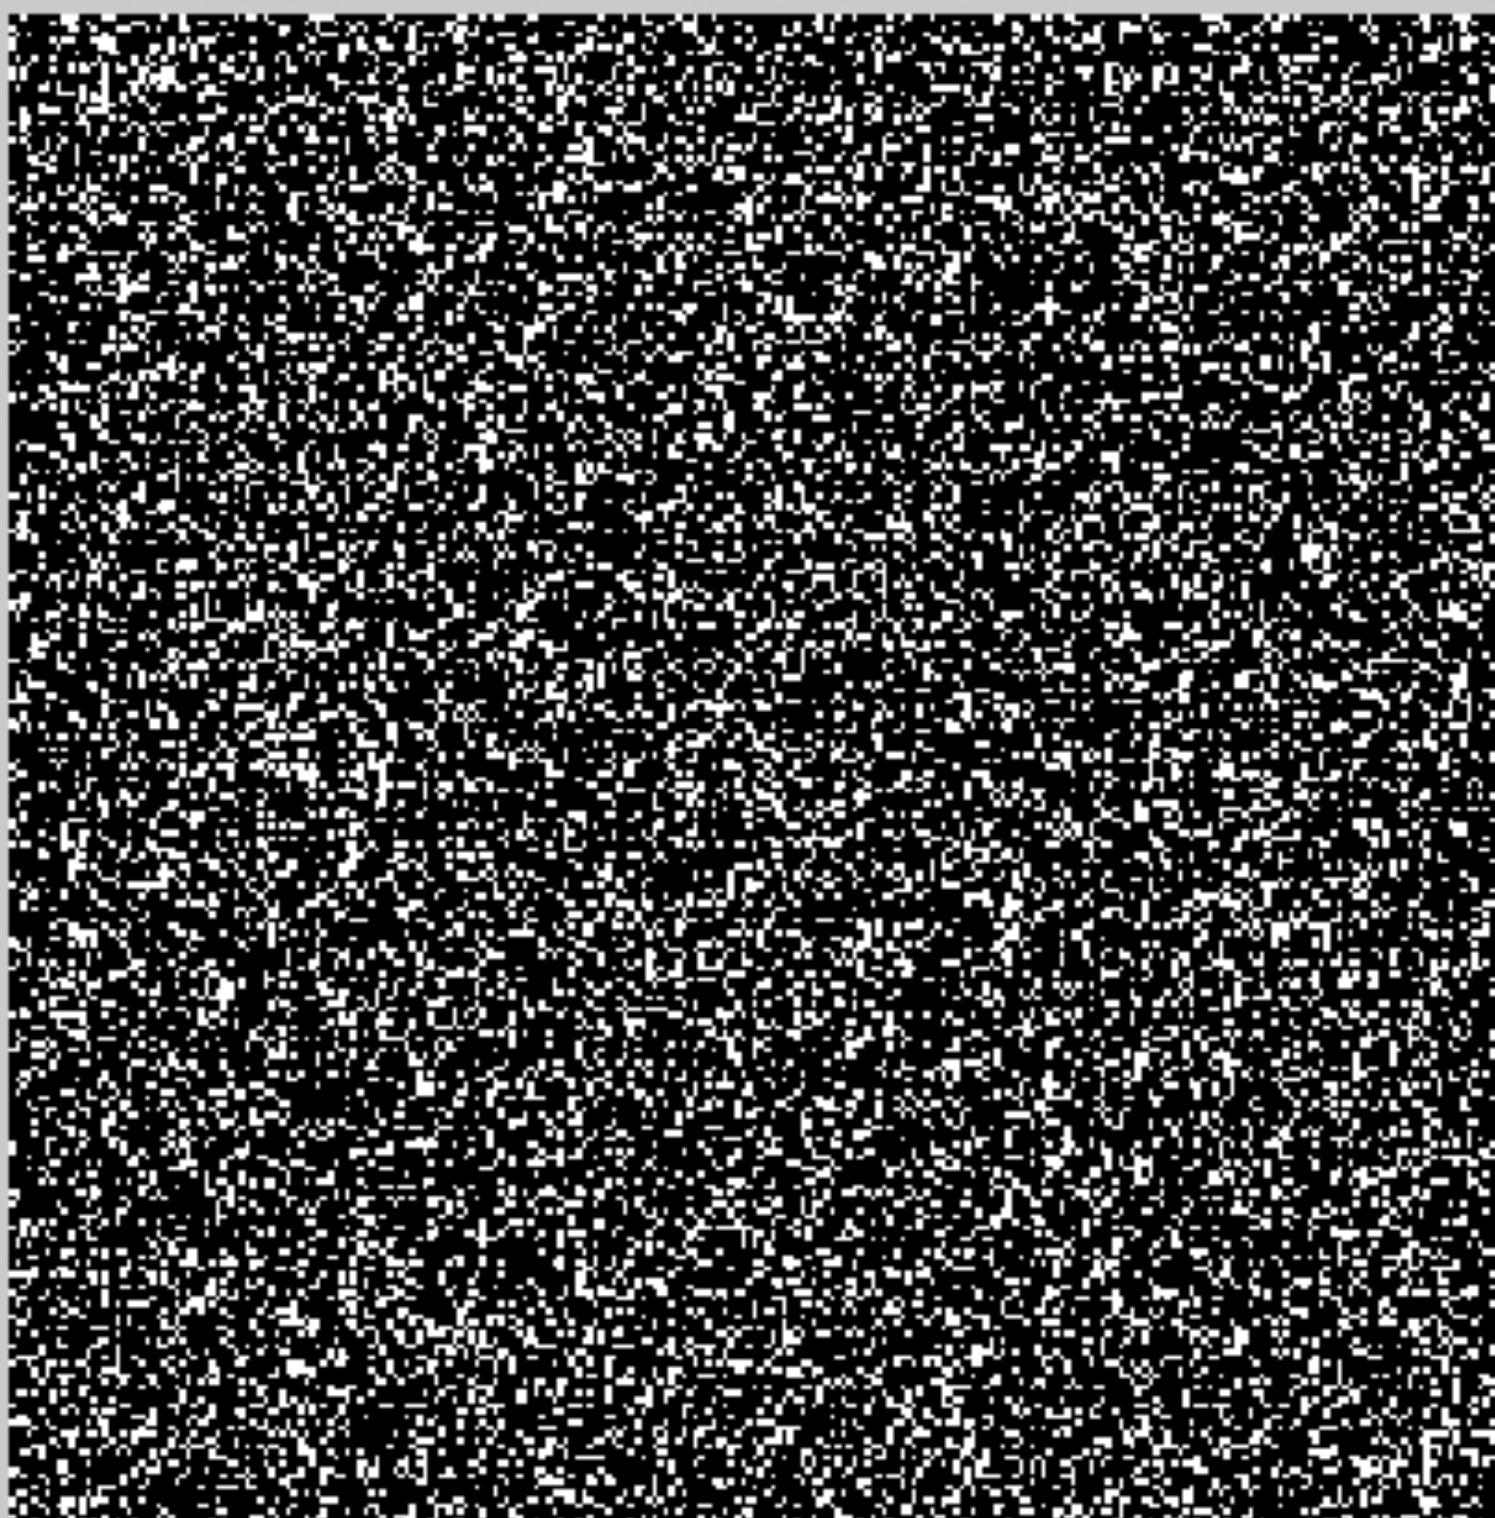
#1: Range [0, 1]
Dims [256, 256]



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

15366

15366



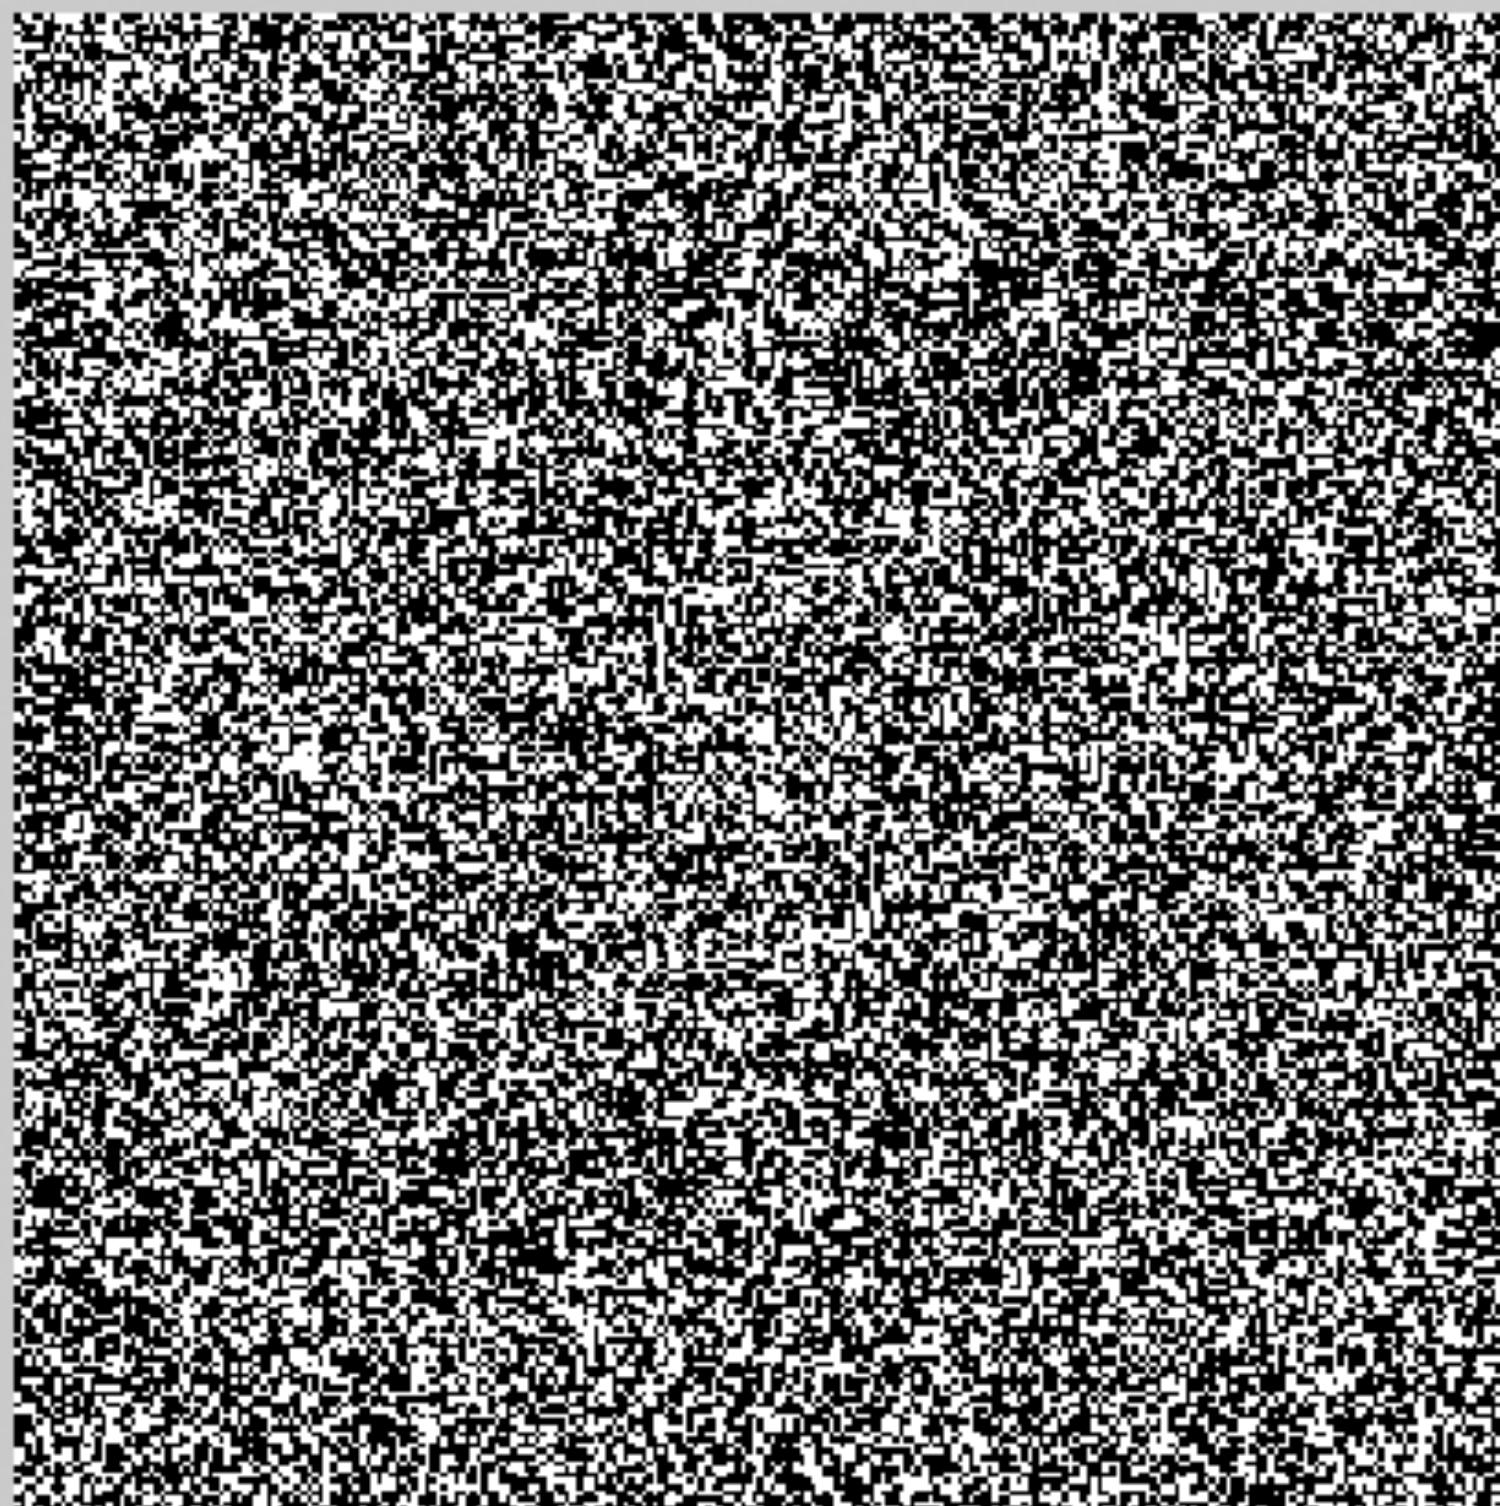
#1: Range [0, 1]
Dims [256, 256]



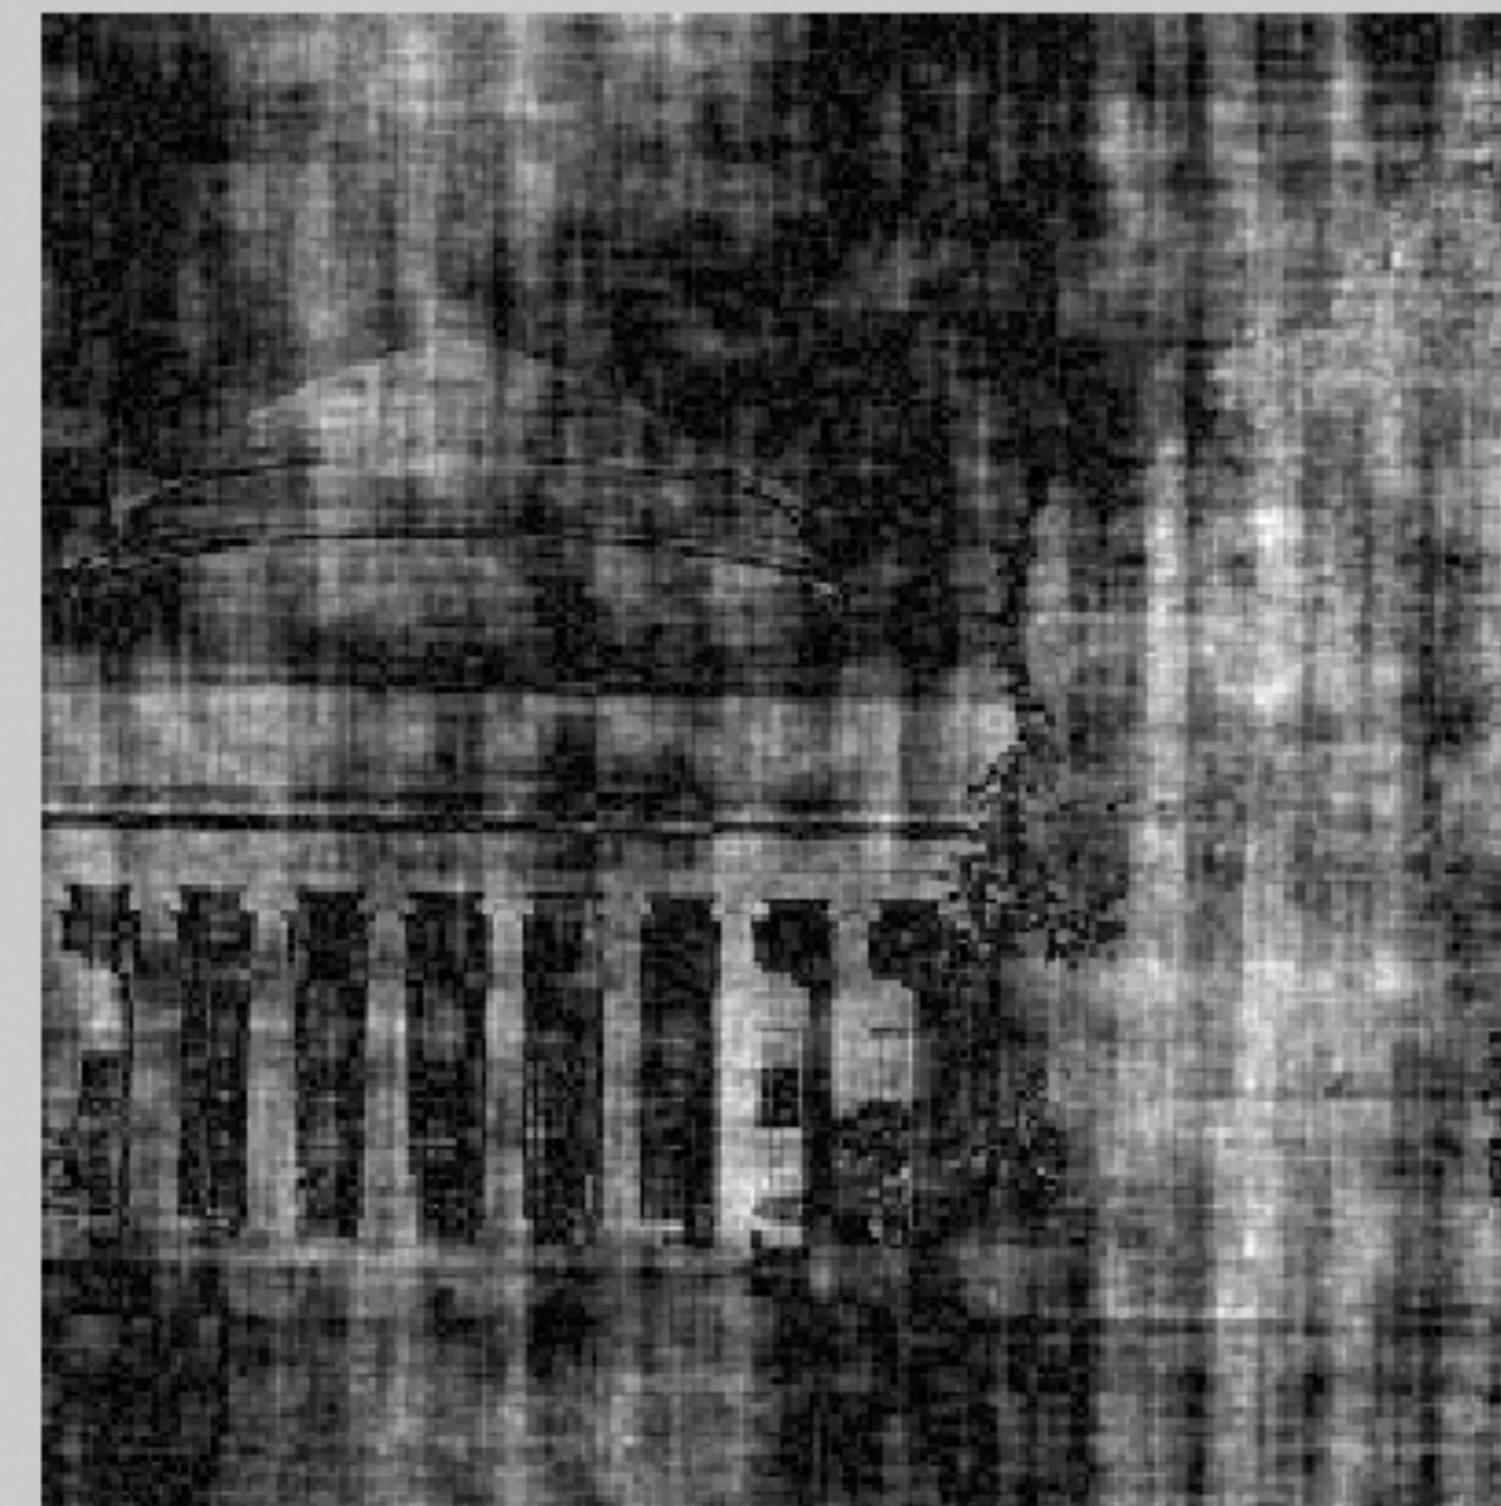
#2: Range [0.000231, 91.1]
Dims [256, 256]

28743

28743



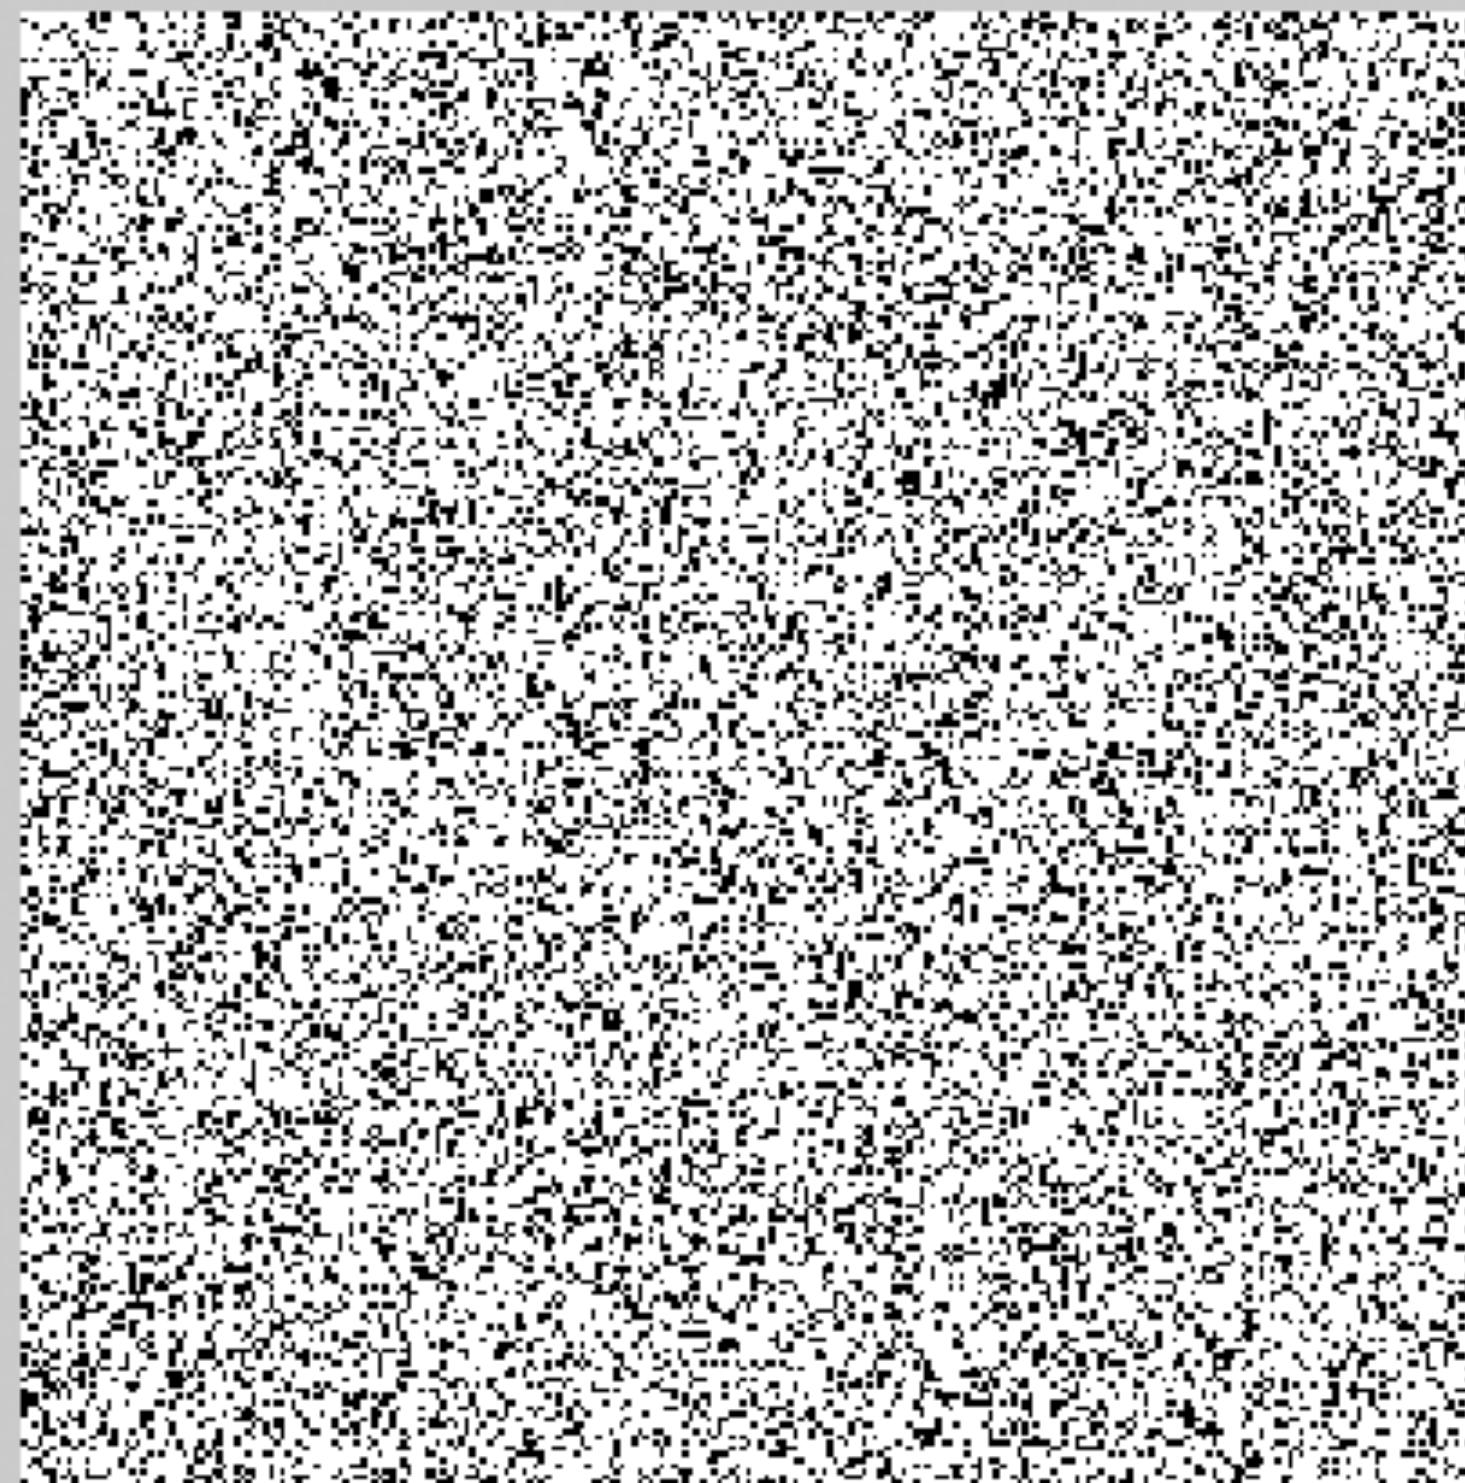
#1: Range [0, 1]
Dims [256, 256]



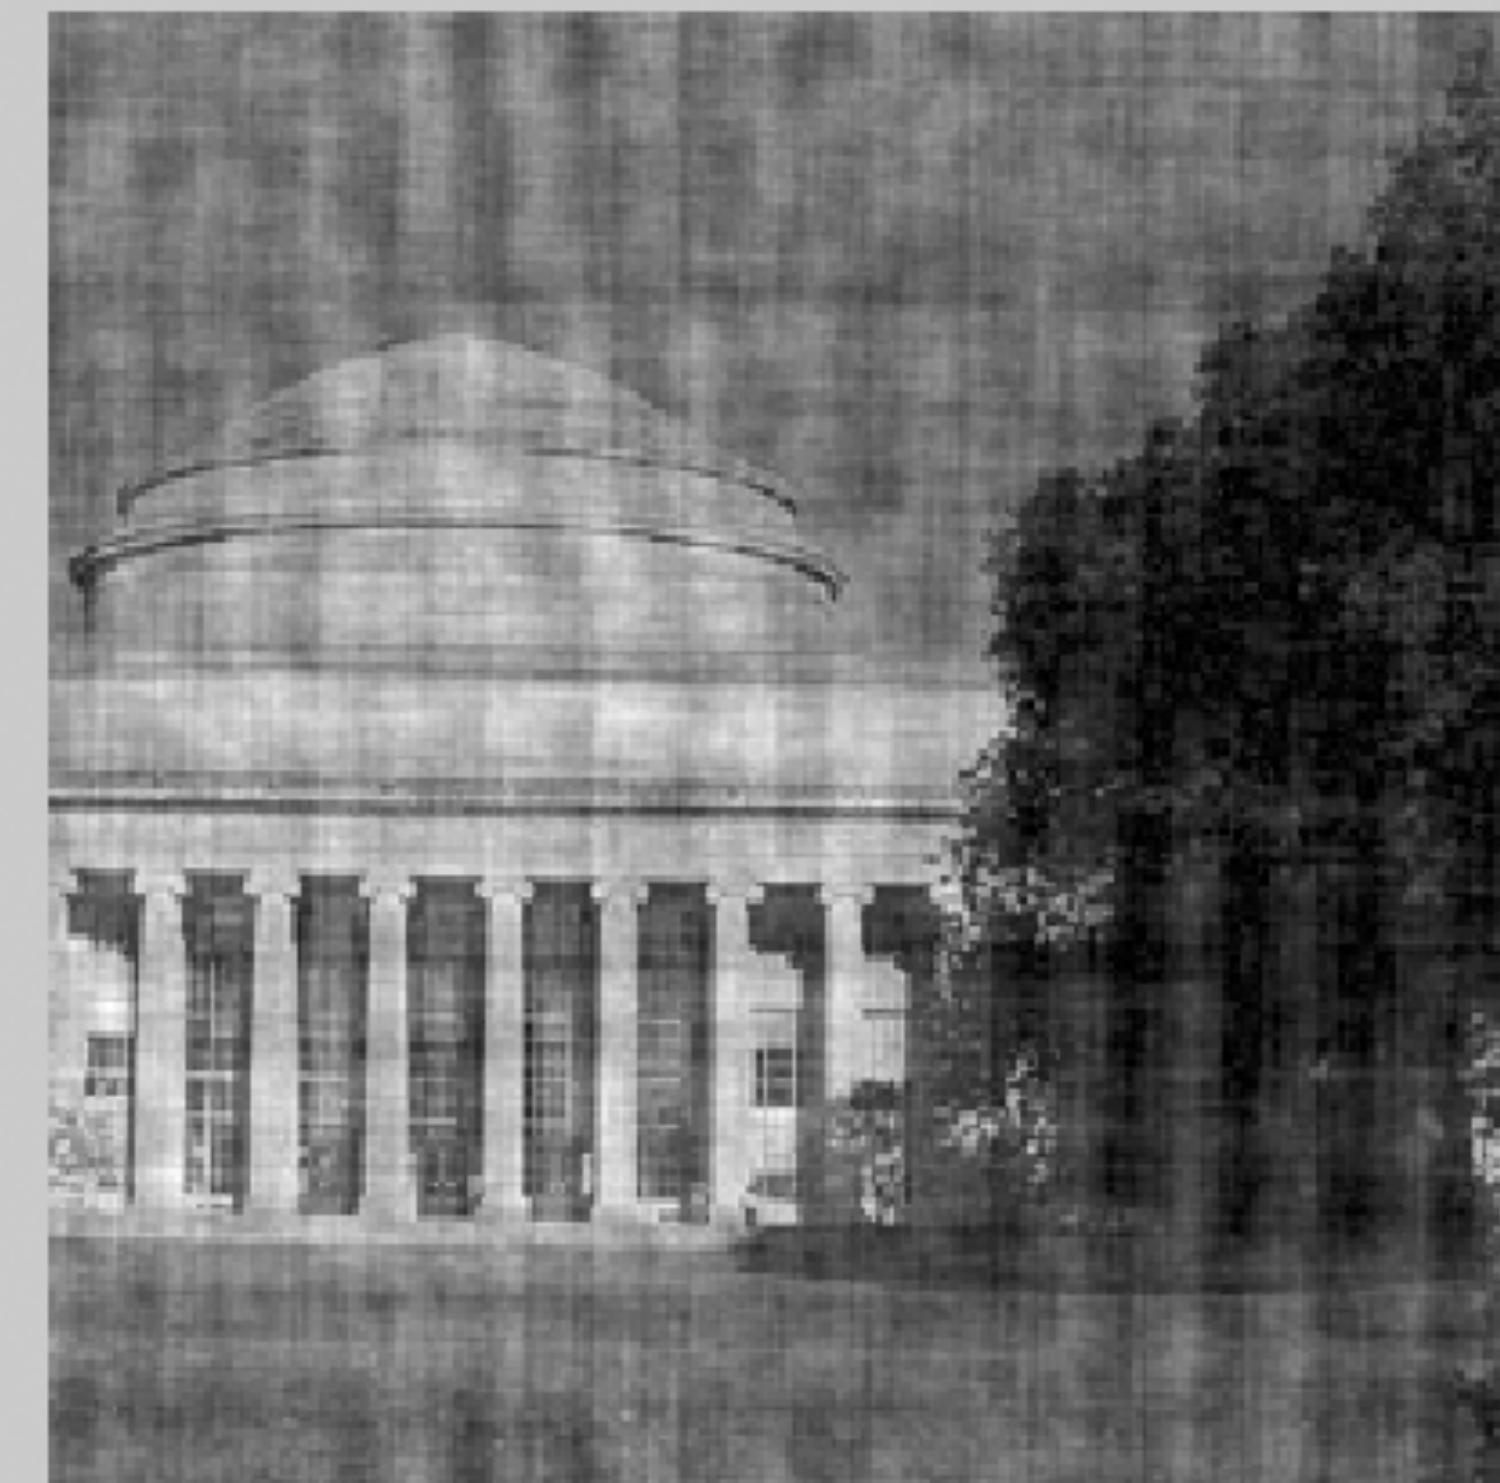
#2: Range [0.00109, 146]
Dims [256, 256]

49190.

49190



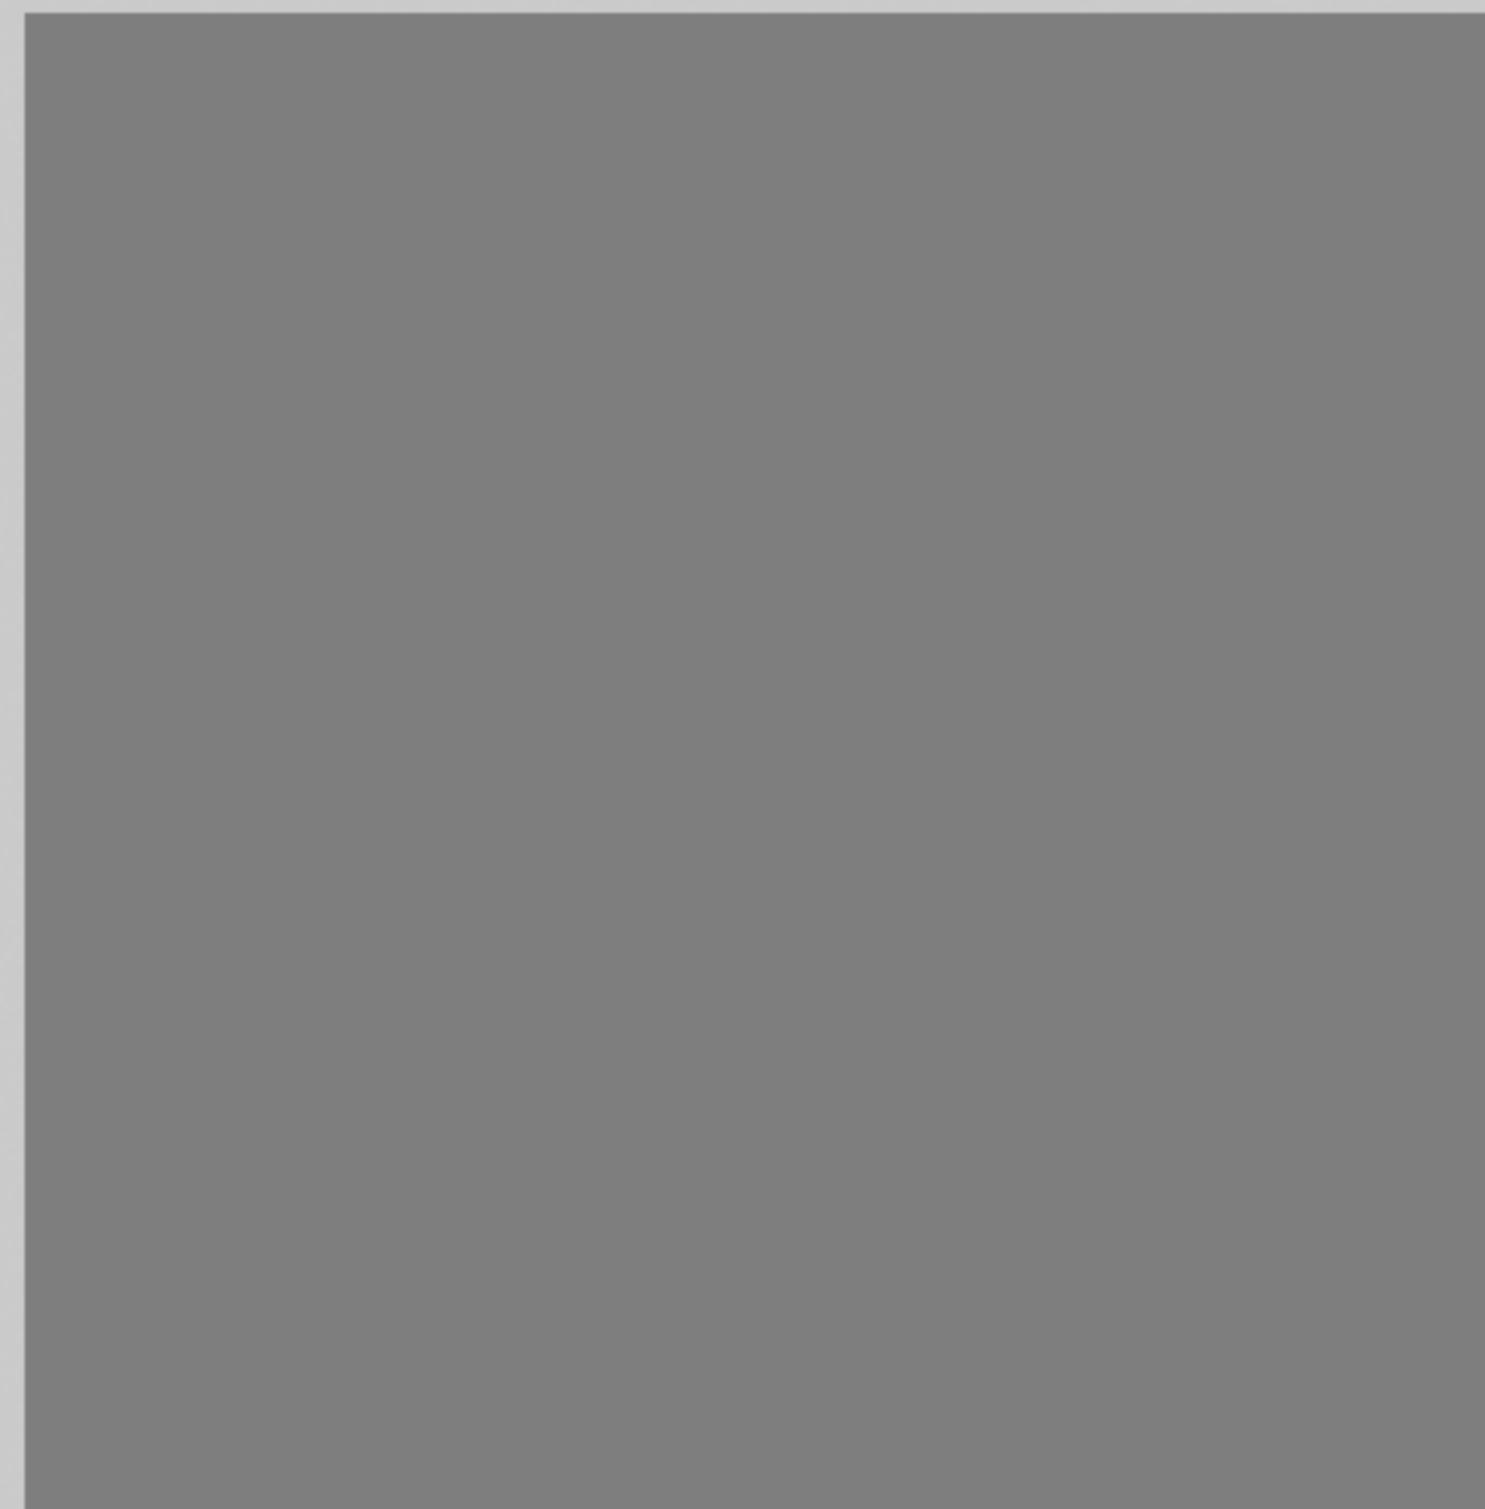
#1: Range [0, 1]
Dims [256, 256]



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

65536.

65536.

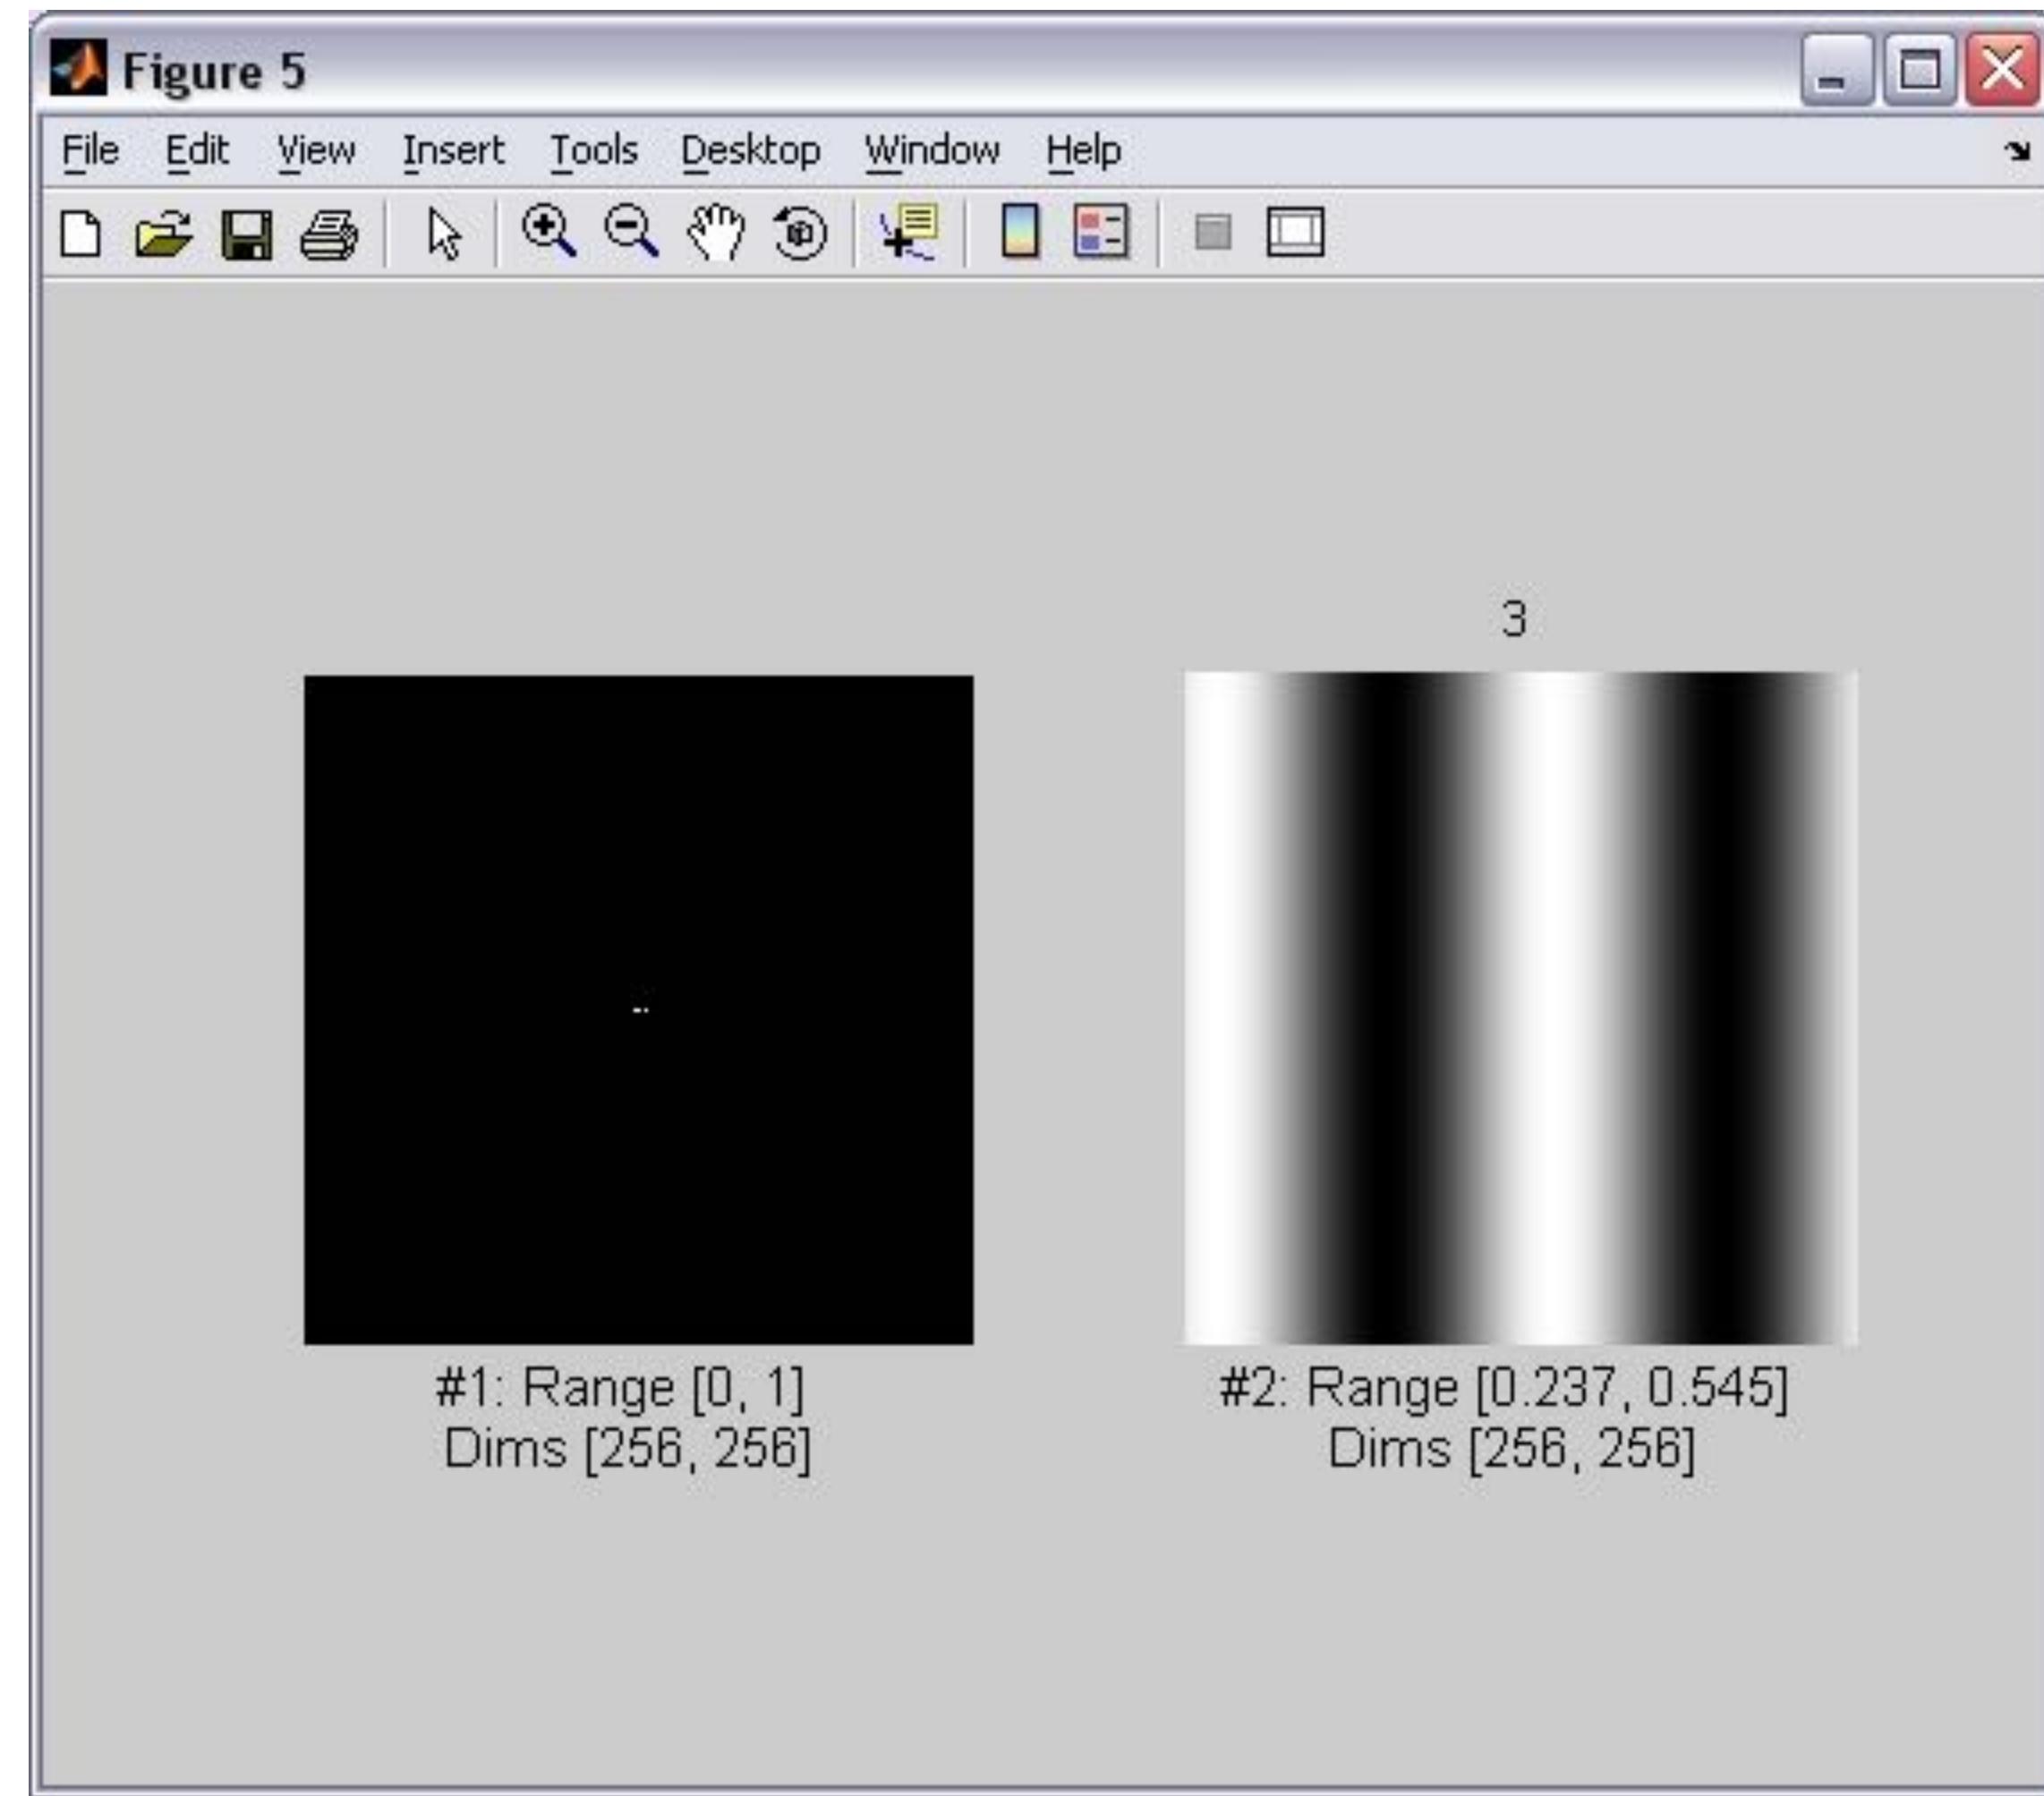


#1: Range [0.5, 1.5]
Dims [256, 256]



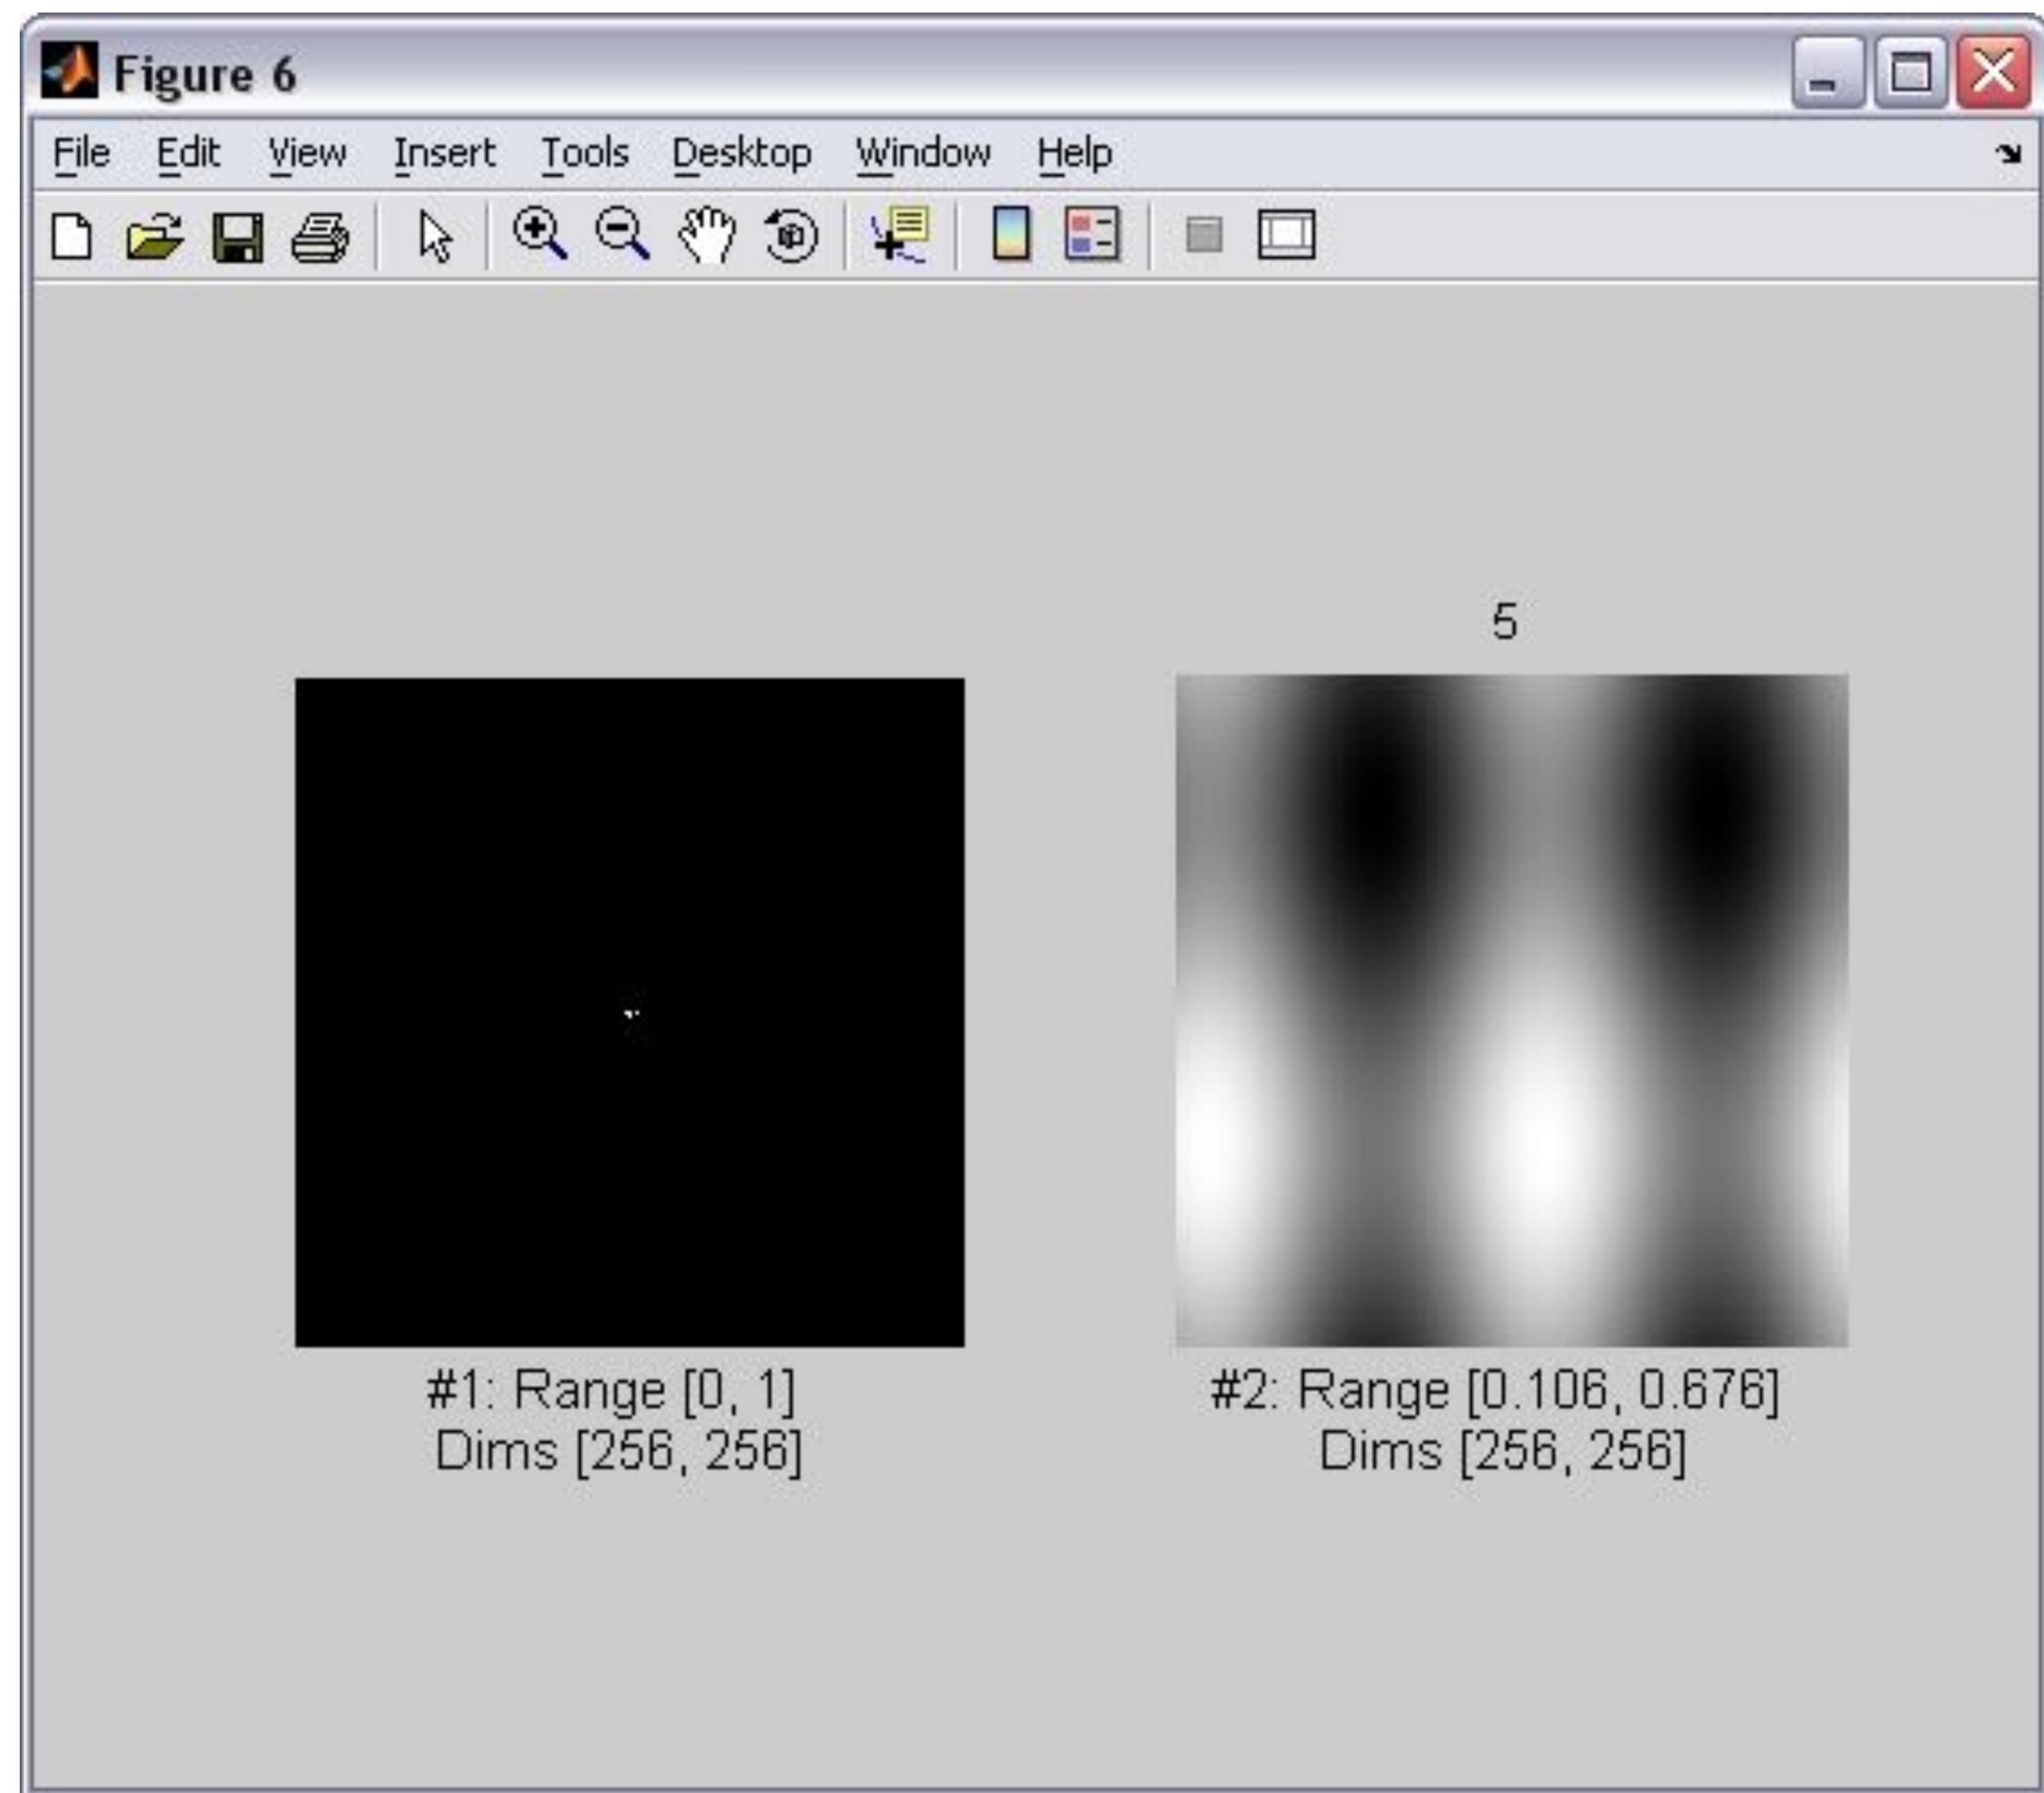
#2: Range [4.43e-015, 255]
Dims [256, 256]

3

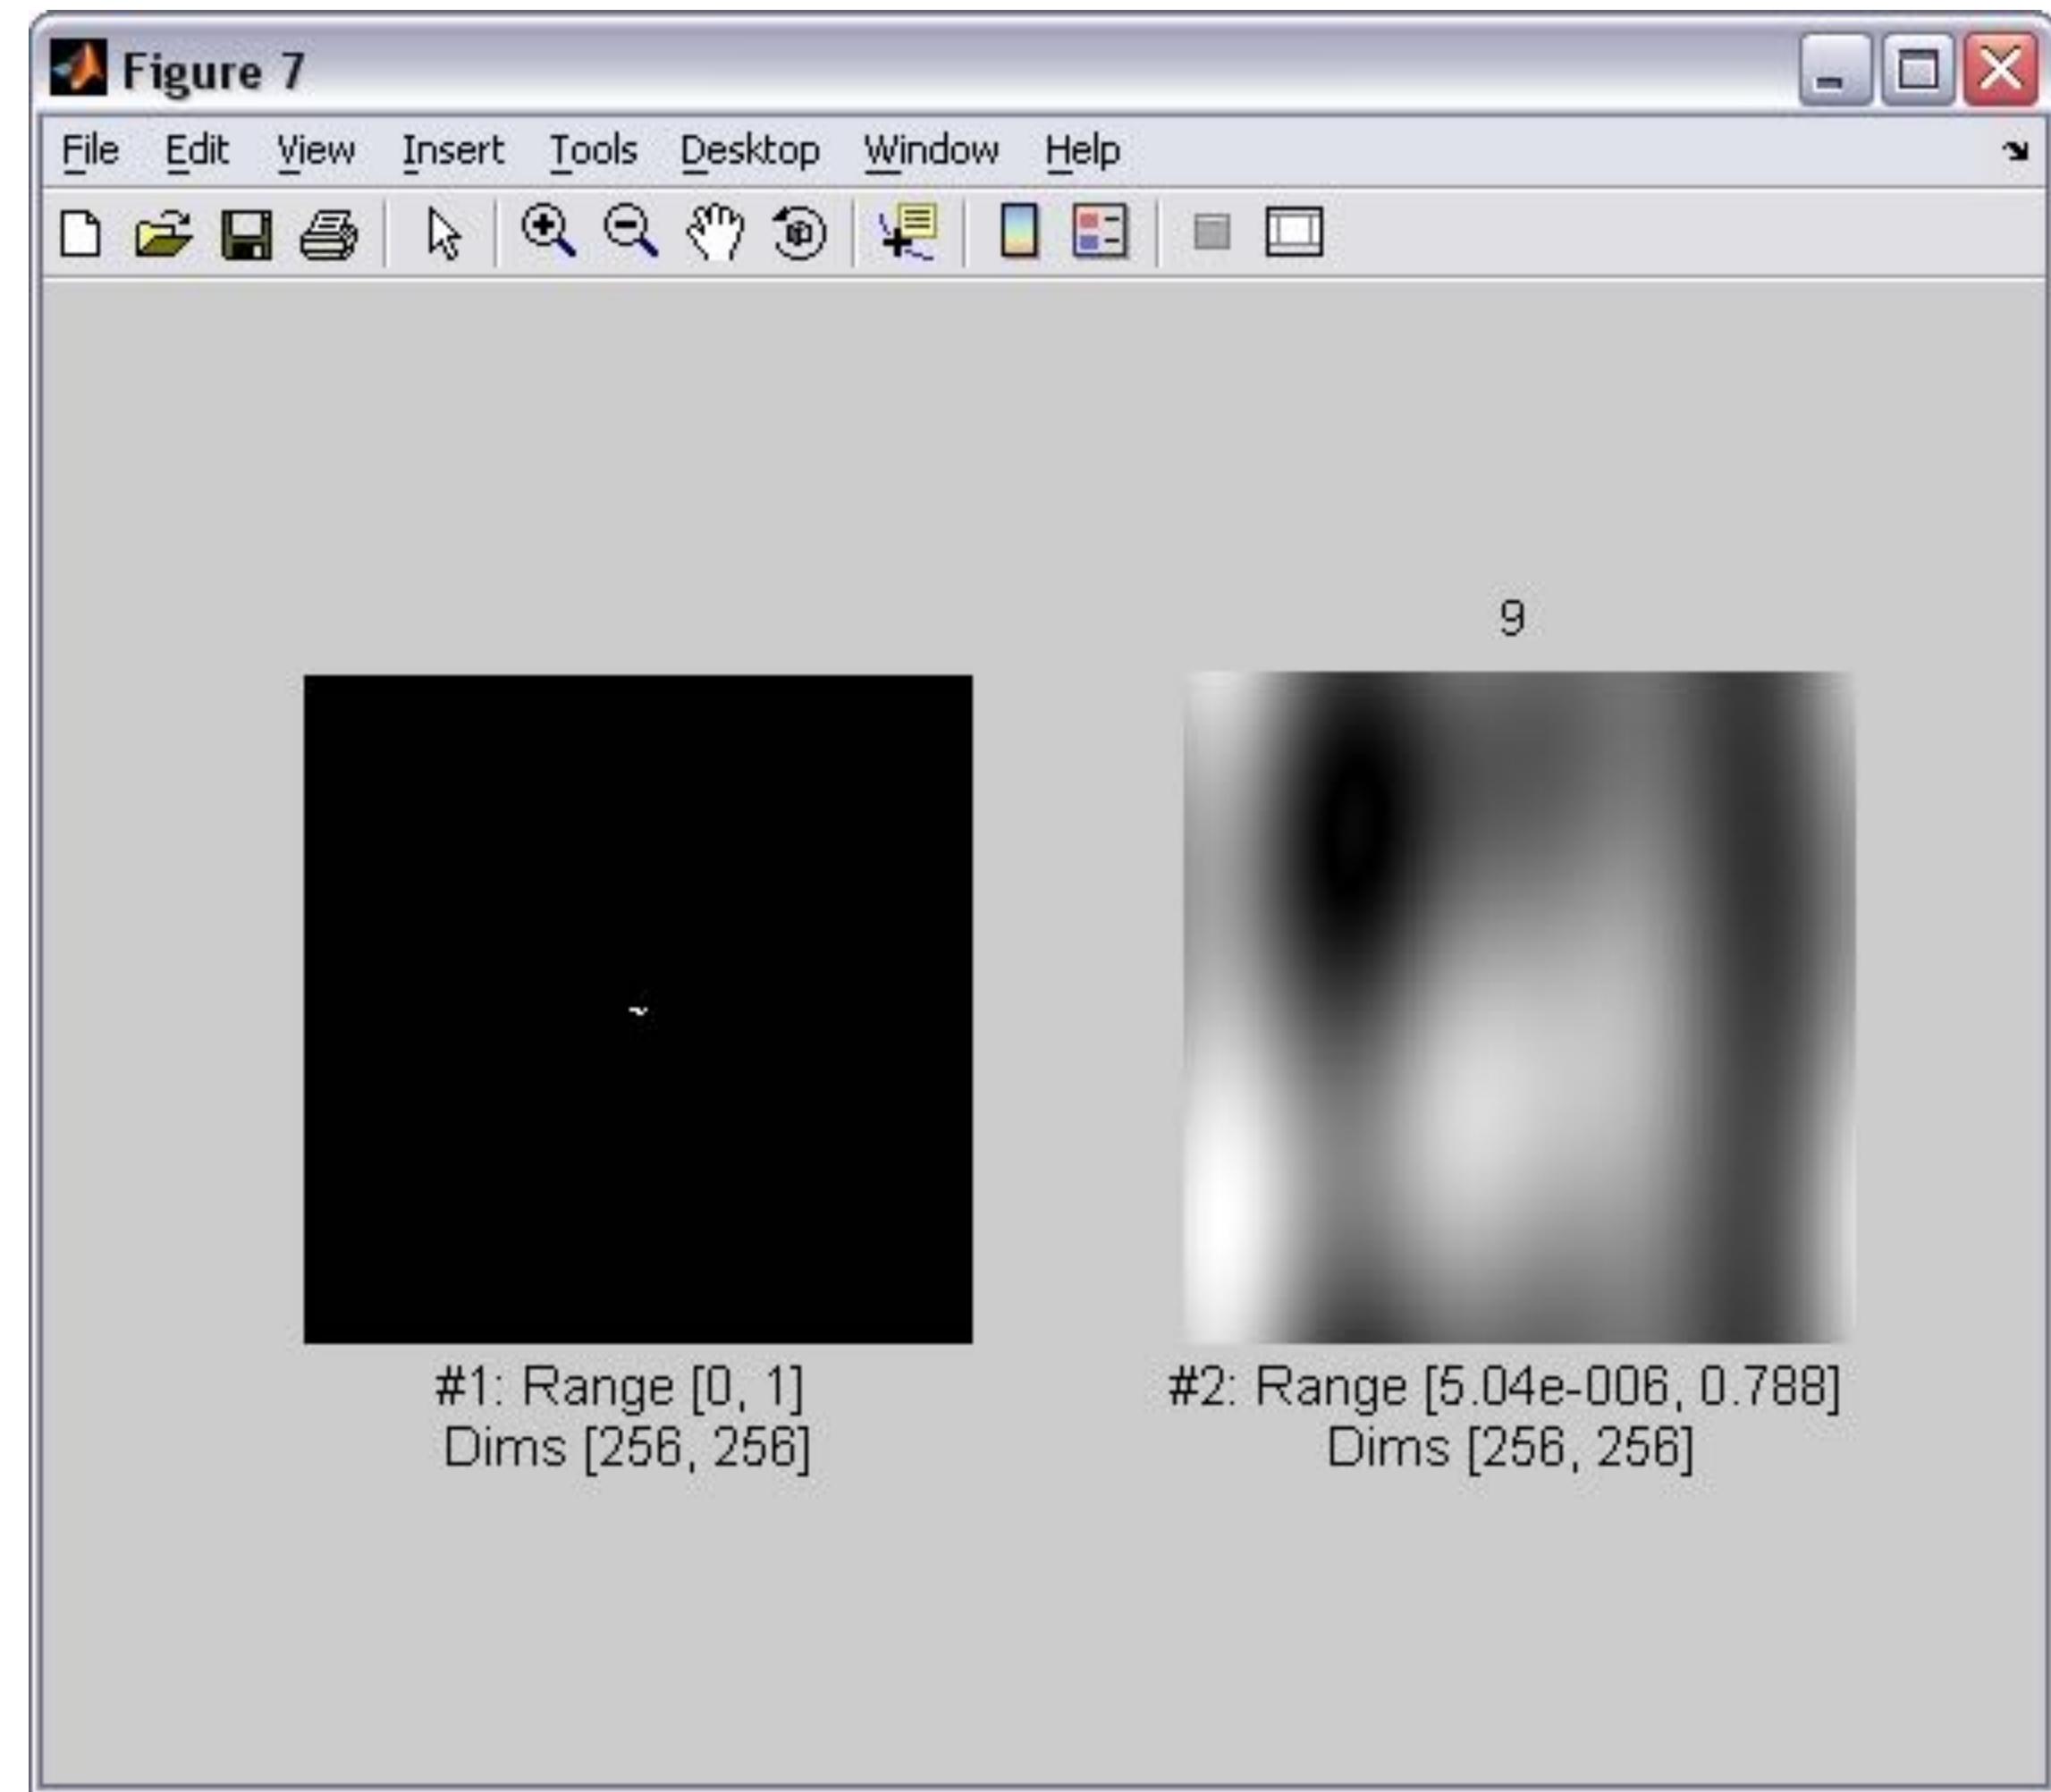


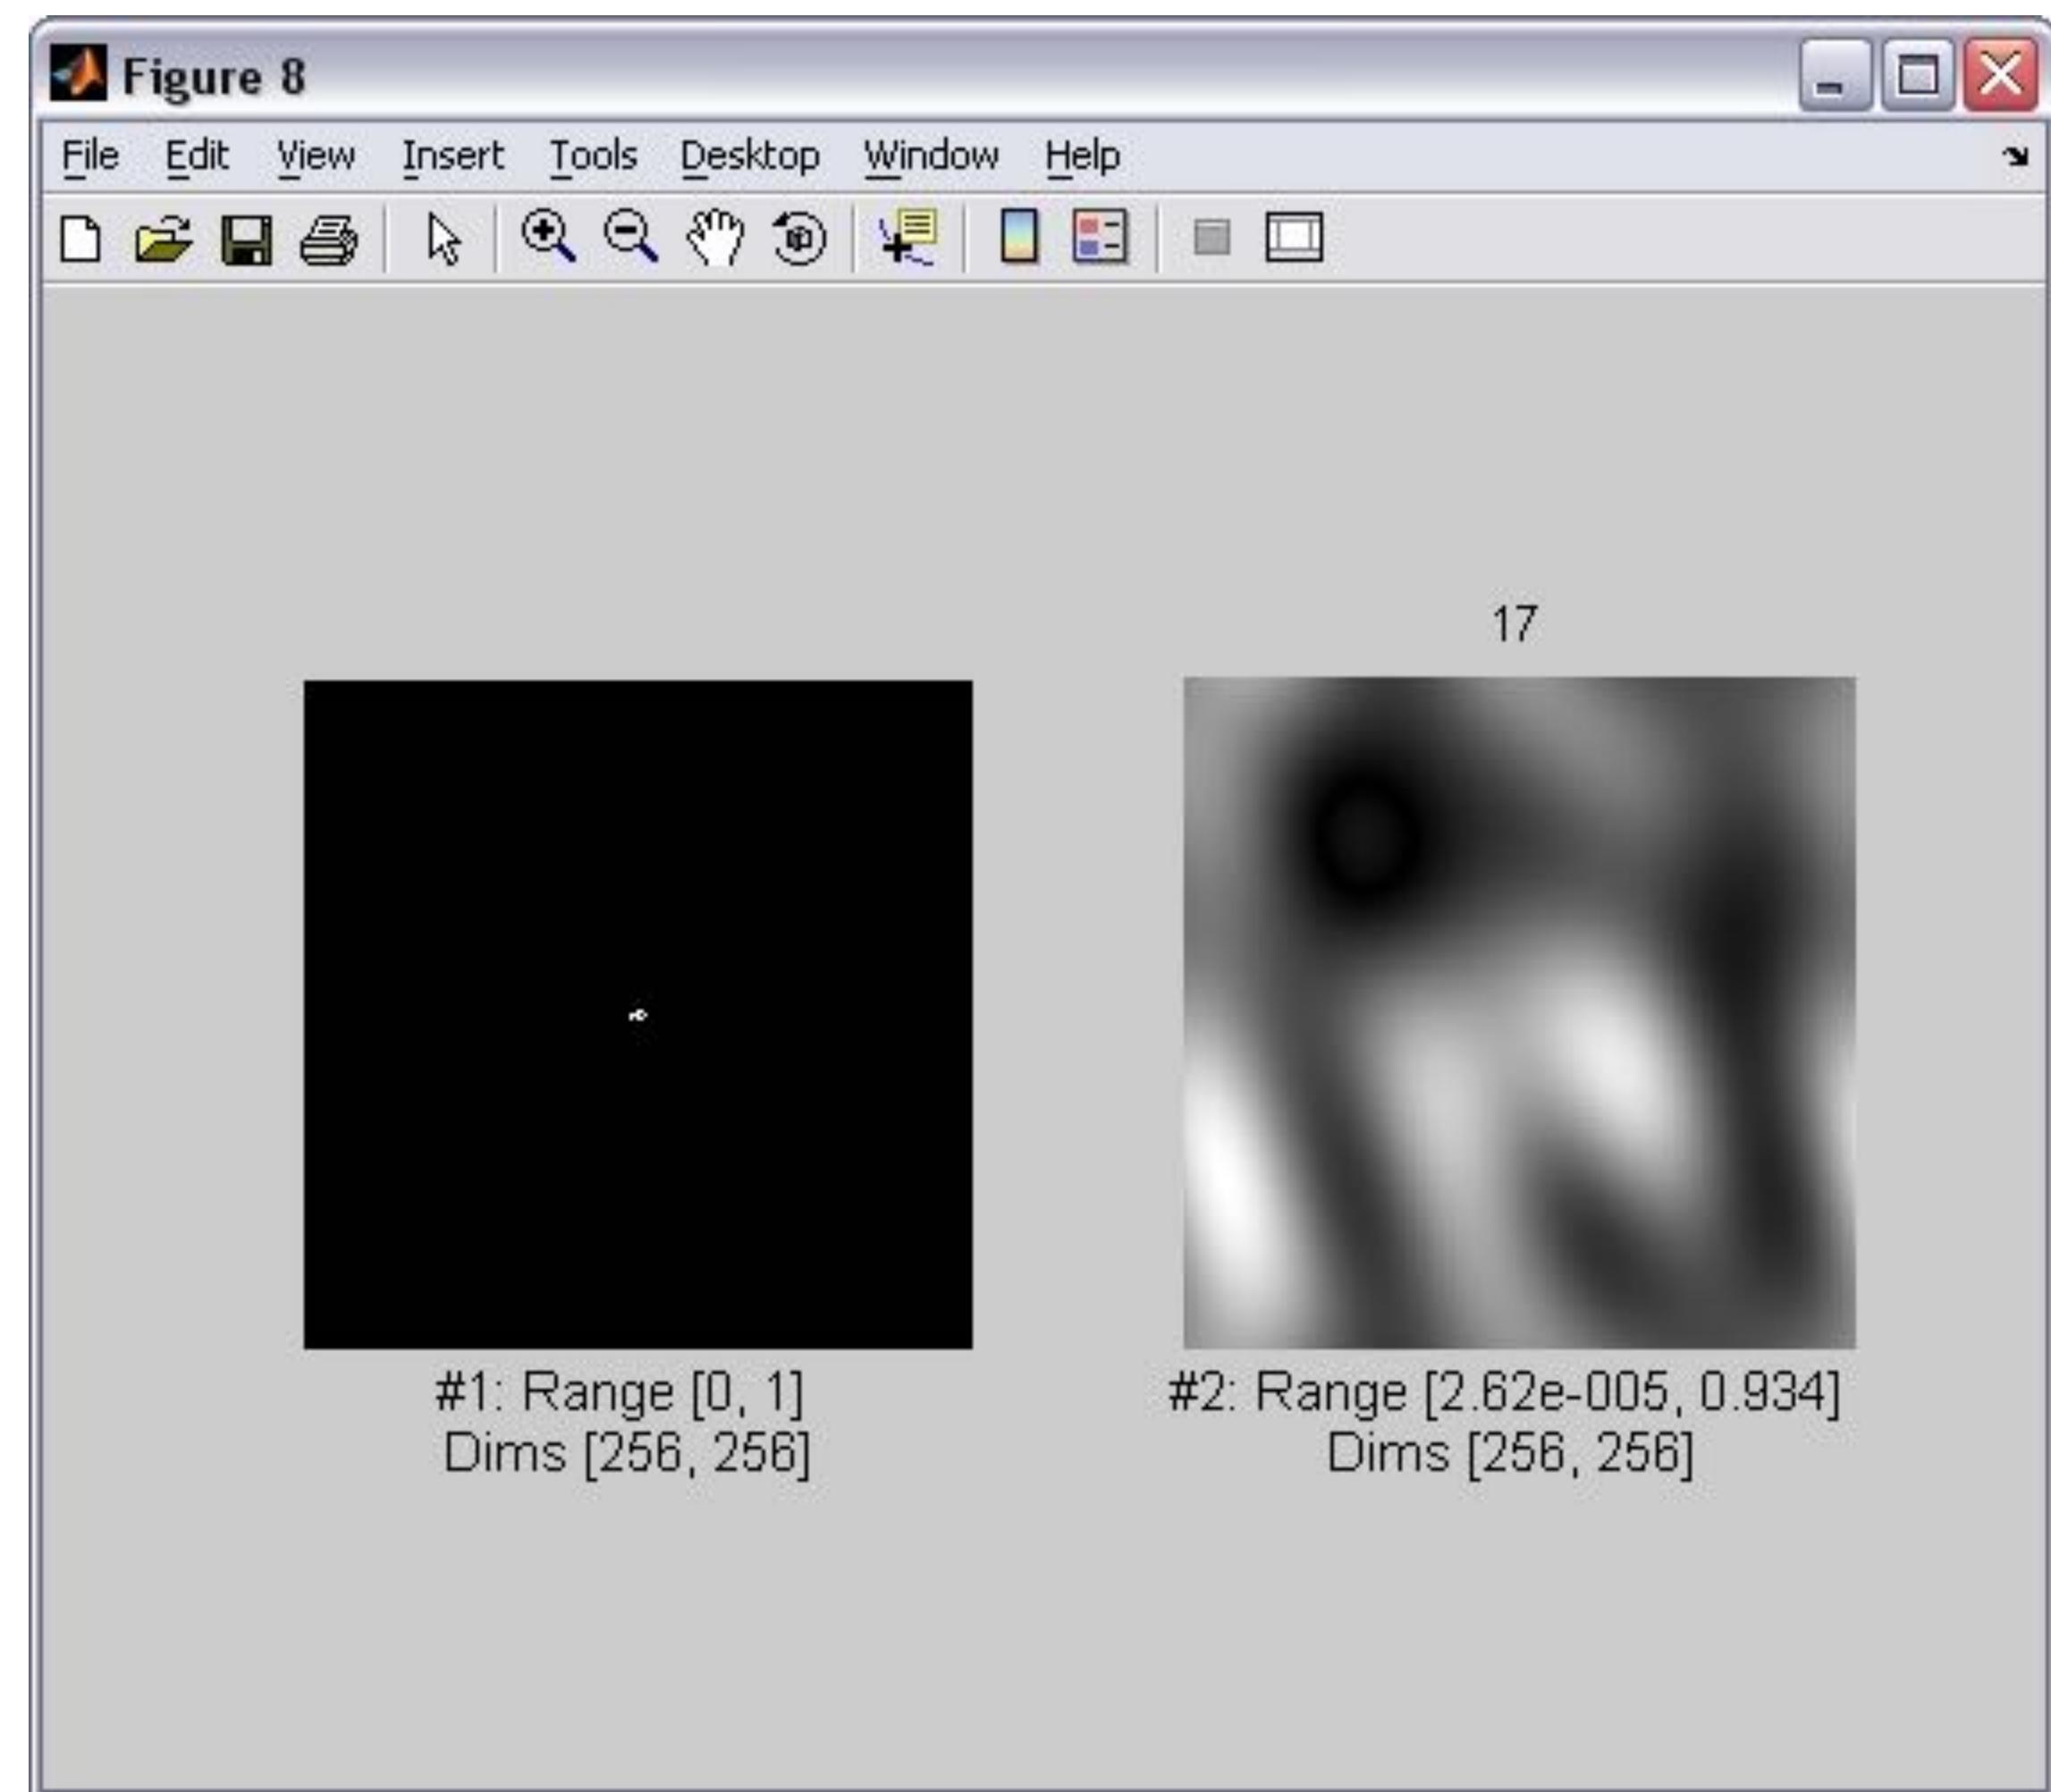
Now, an analogous sequence of images, but selecting Fourier components in descending order of magnitude.

5

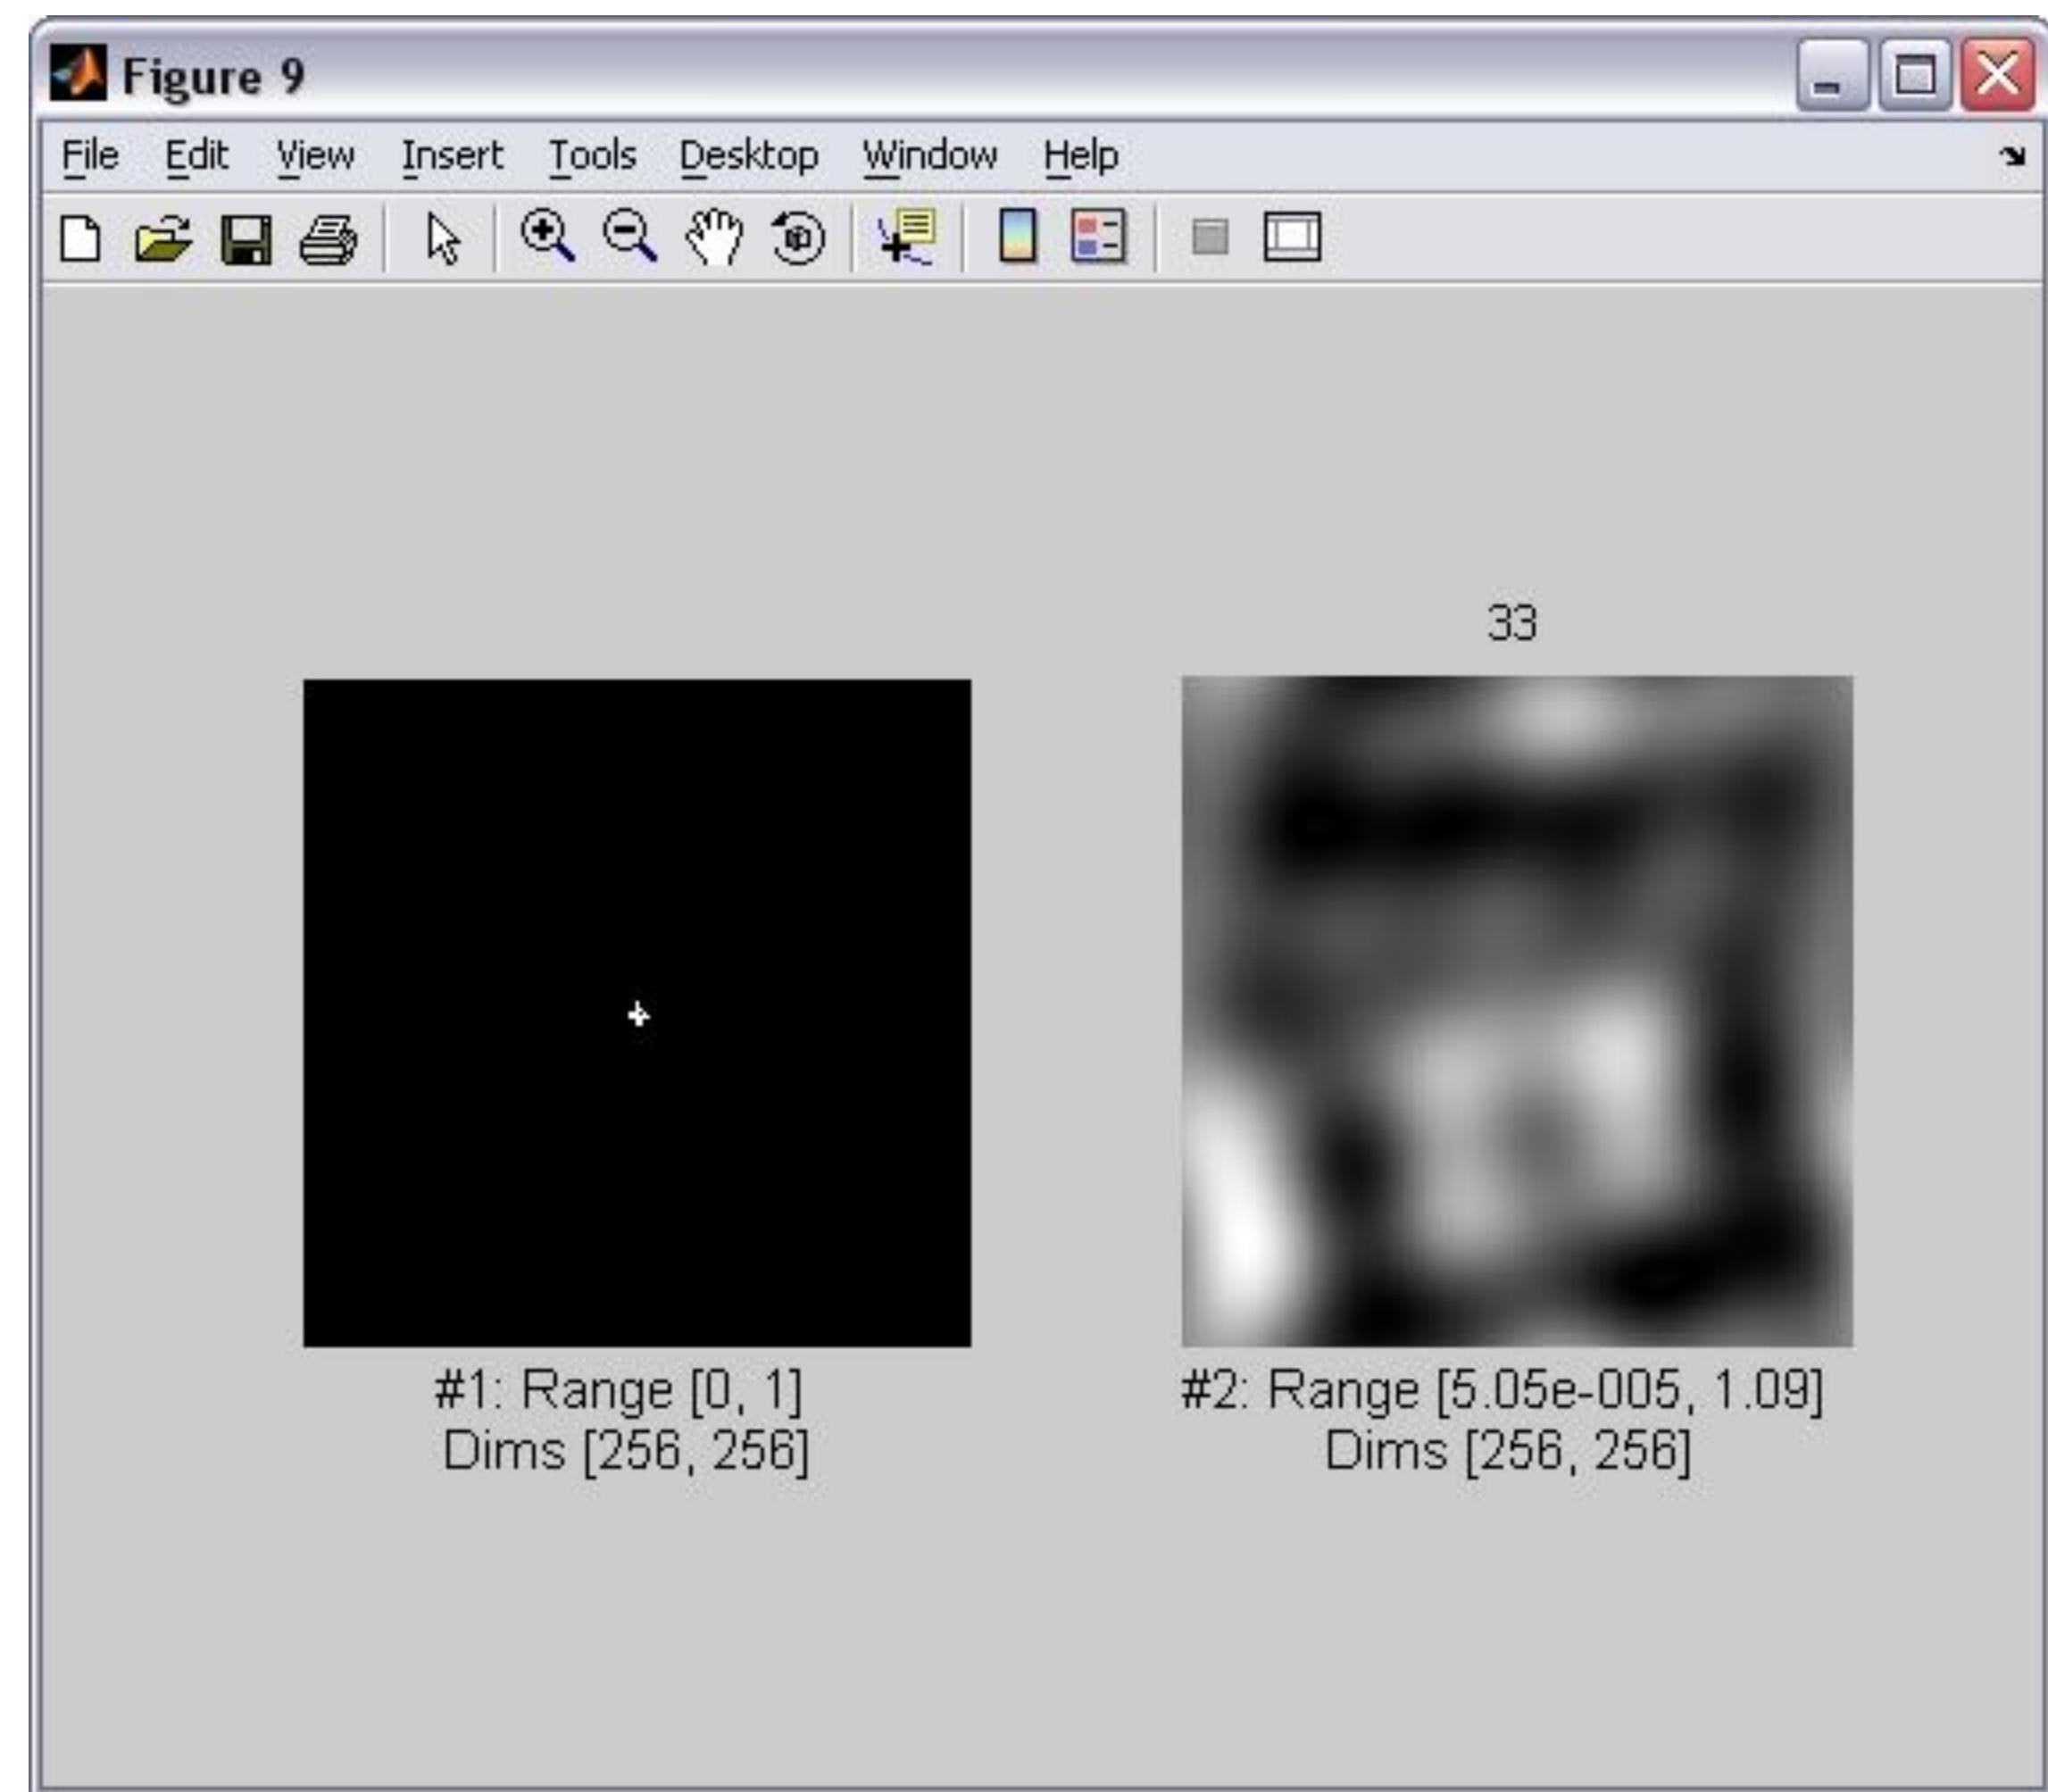


9

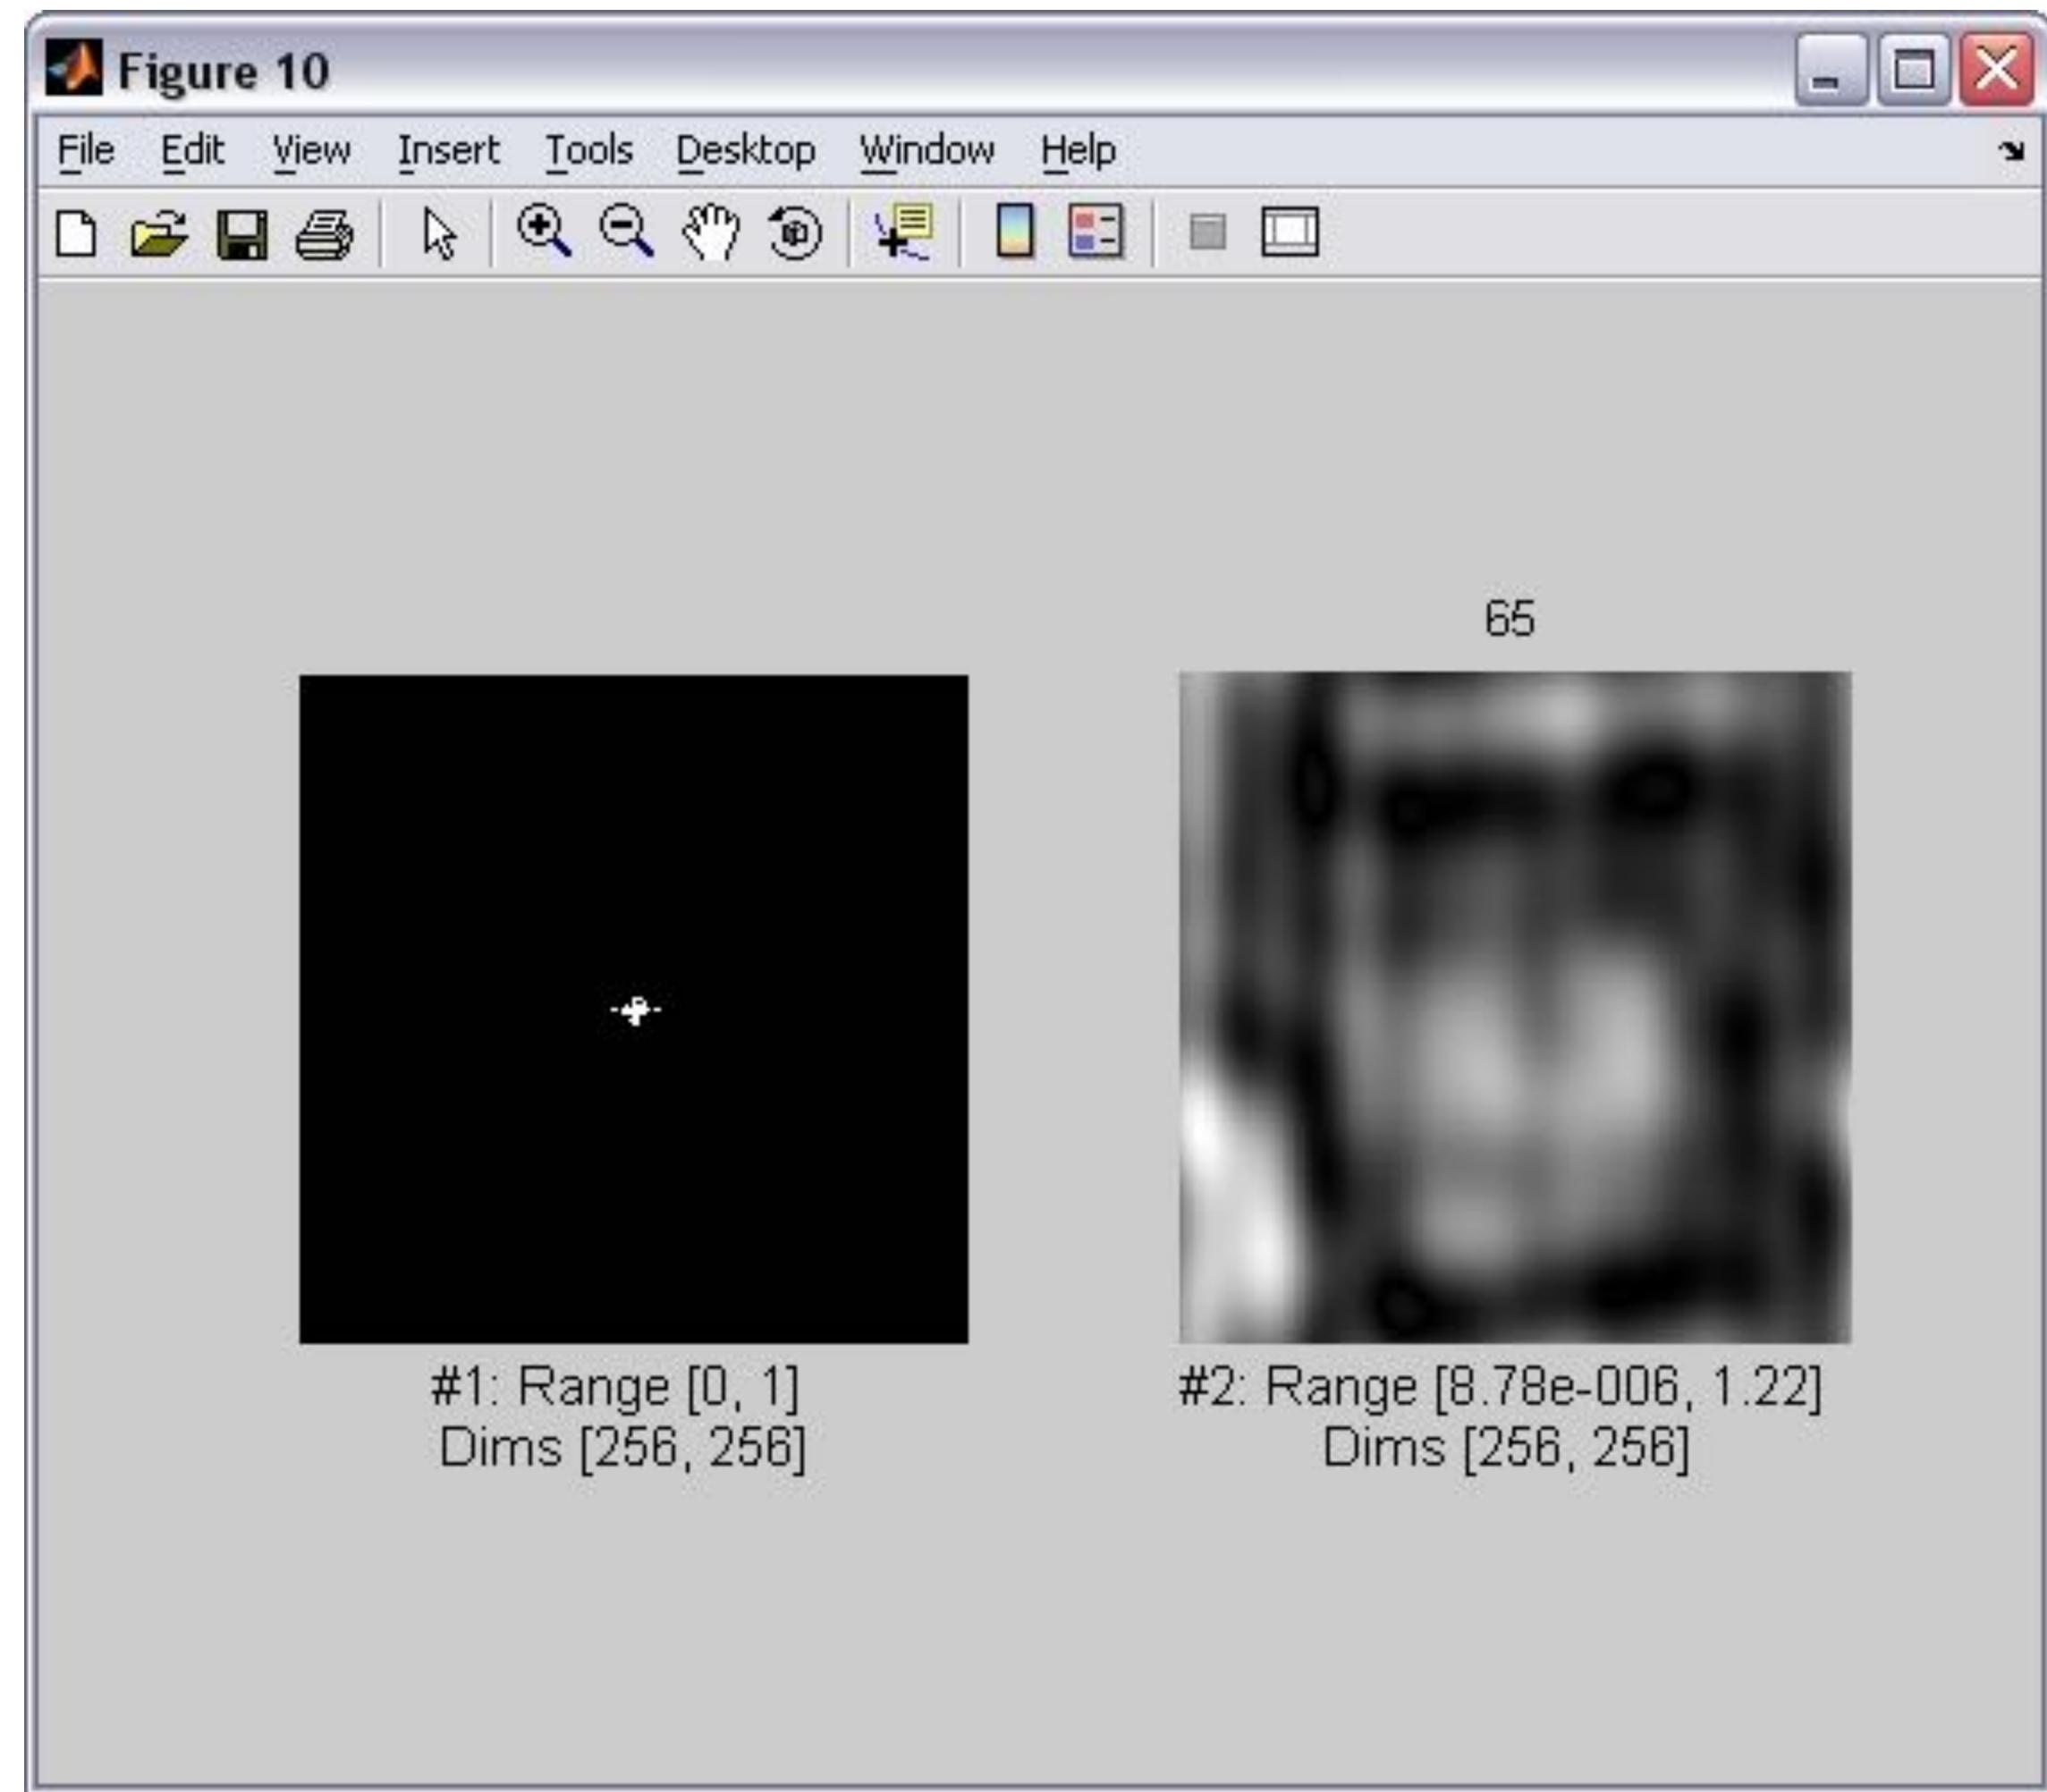




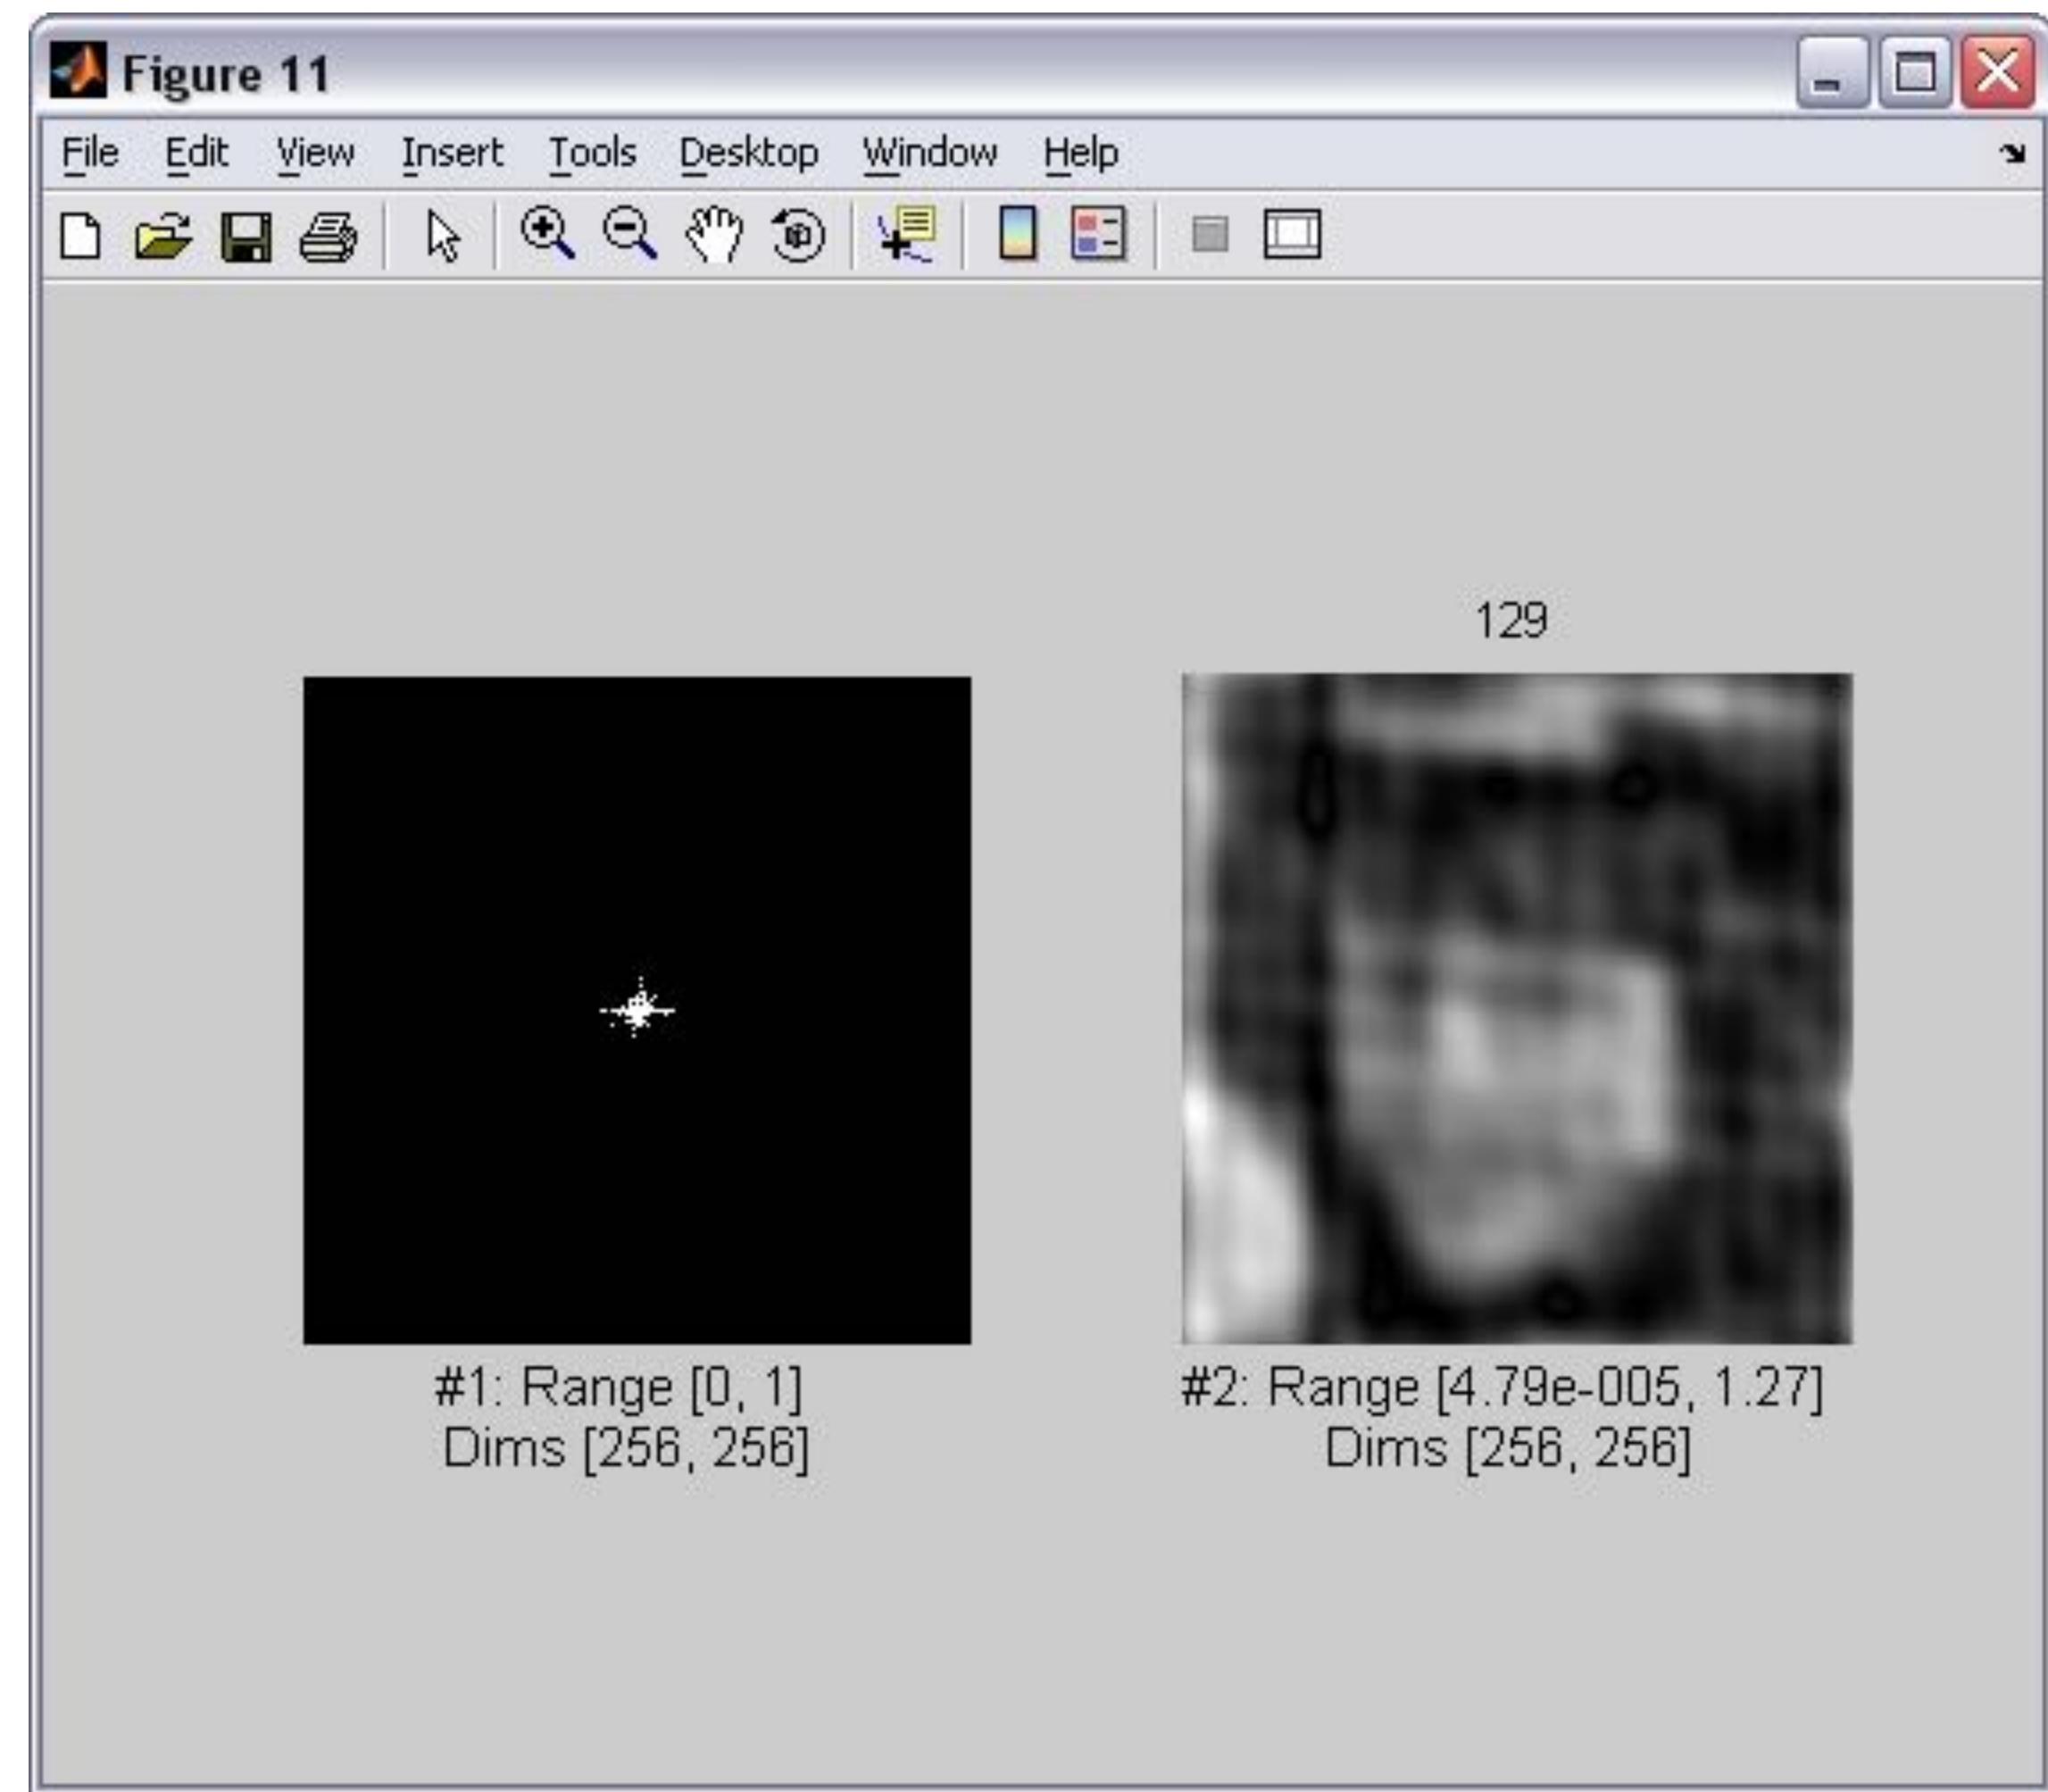
33



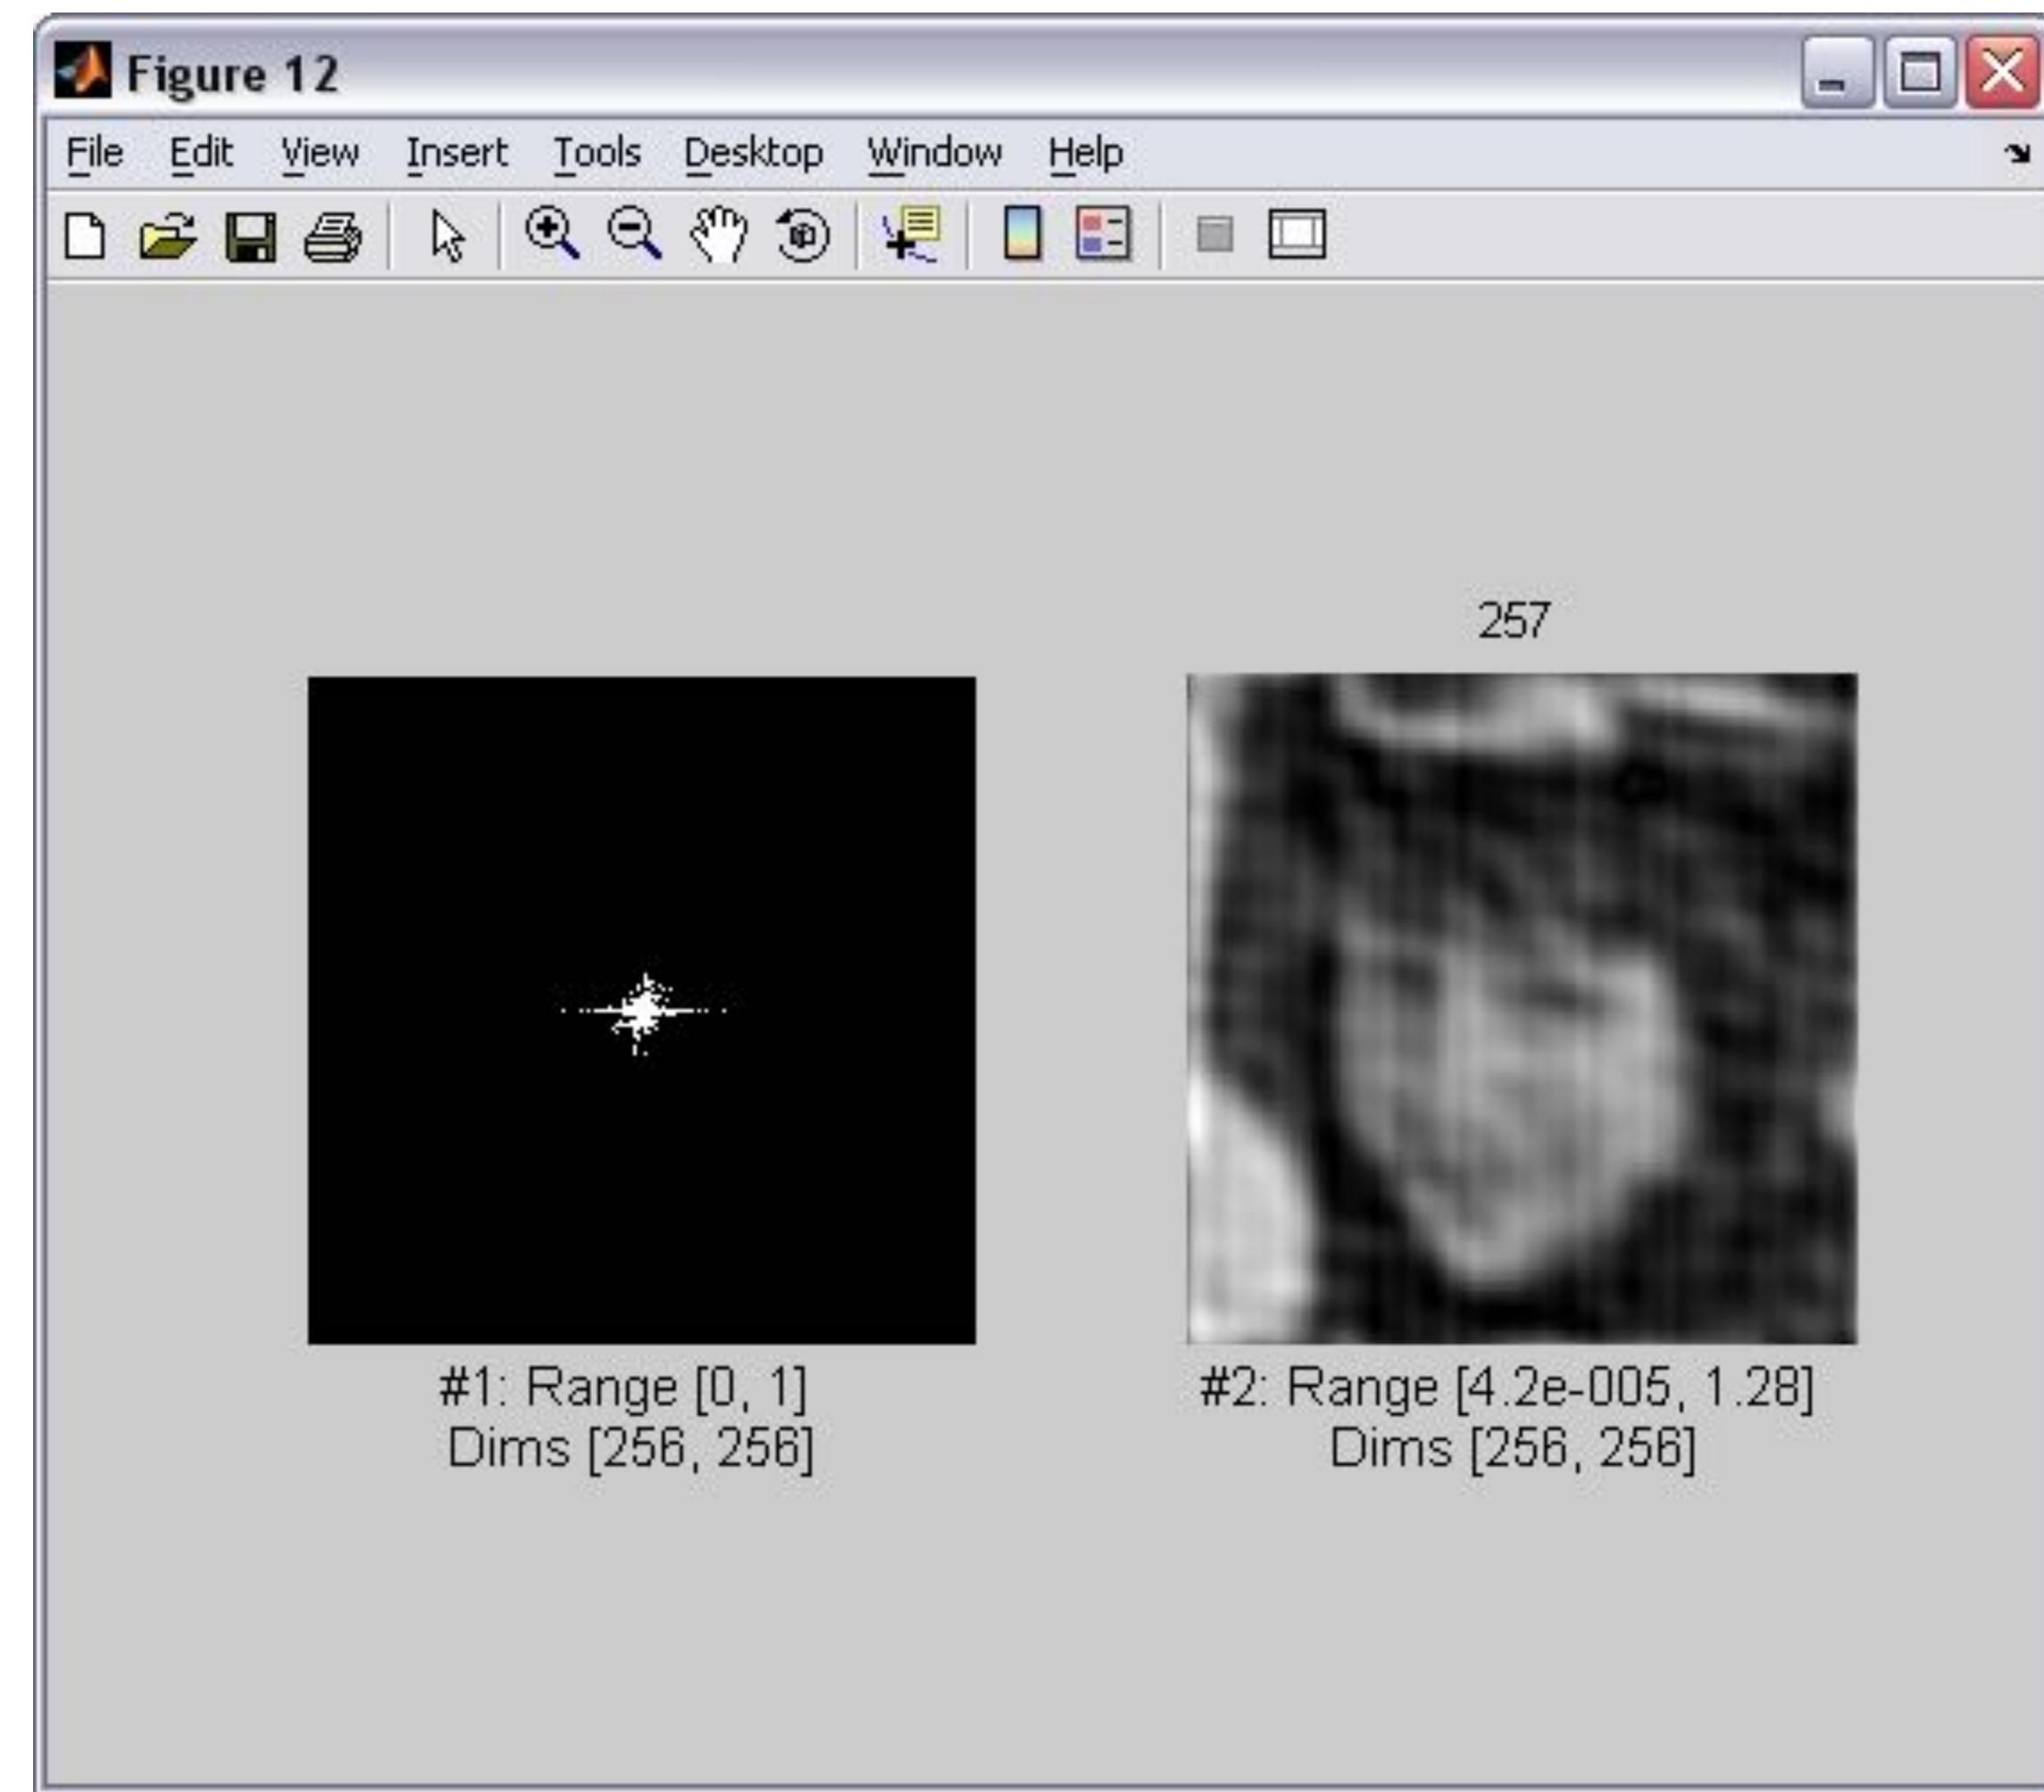
65



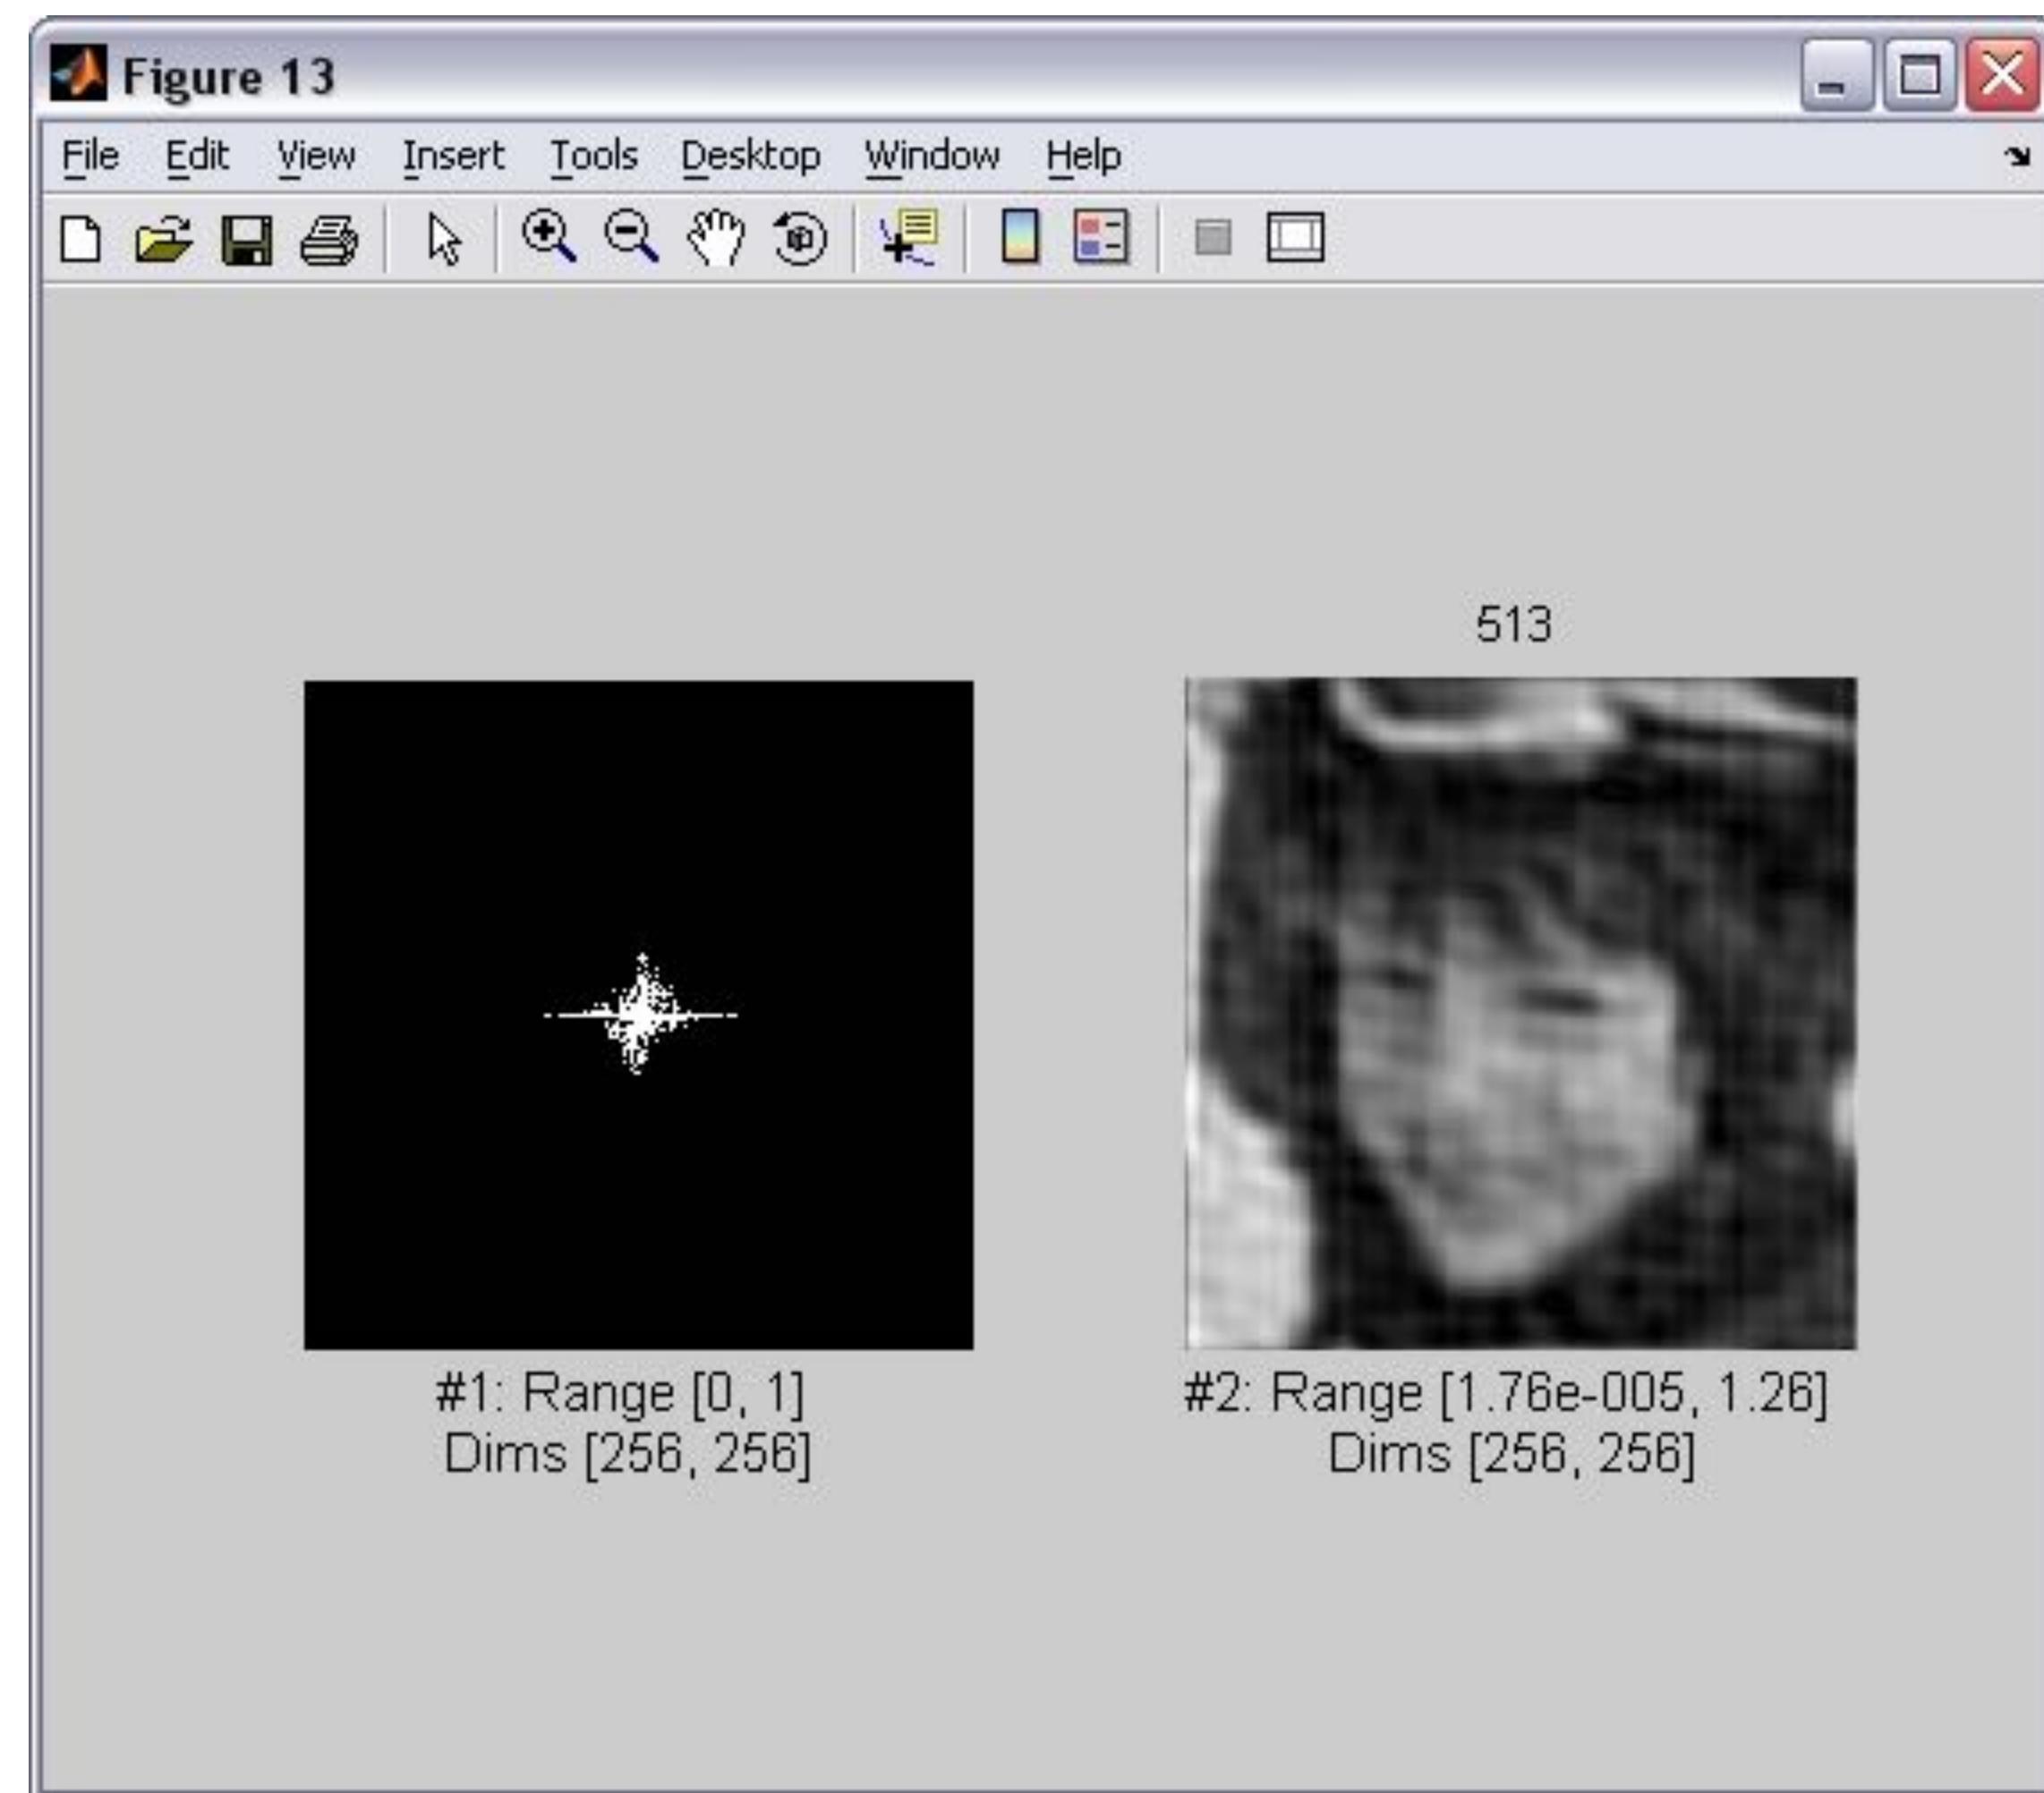
129



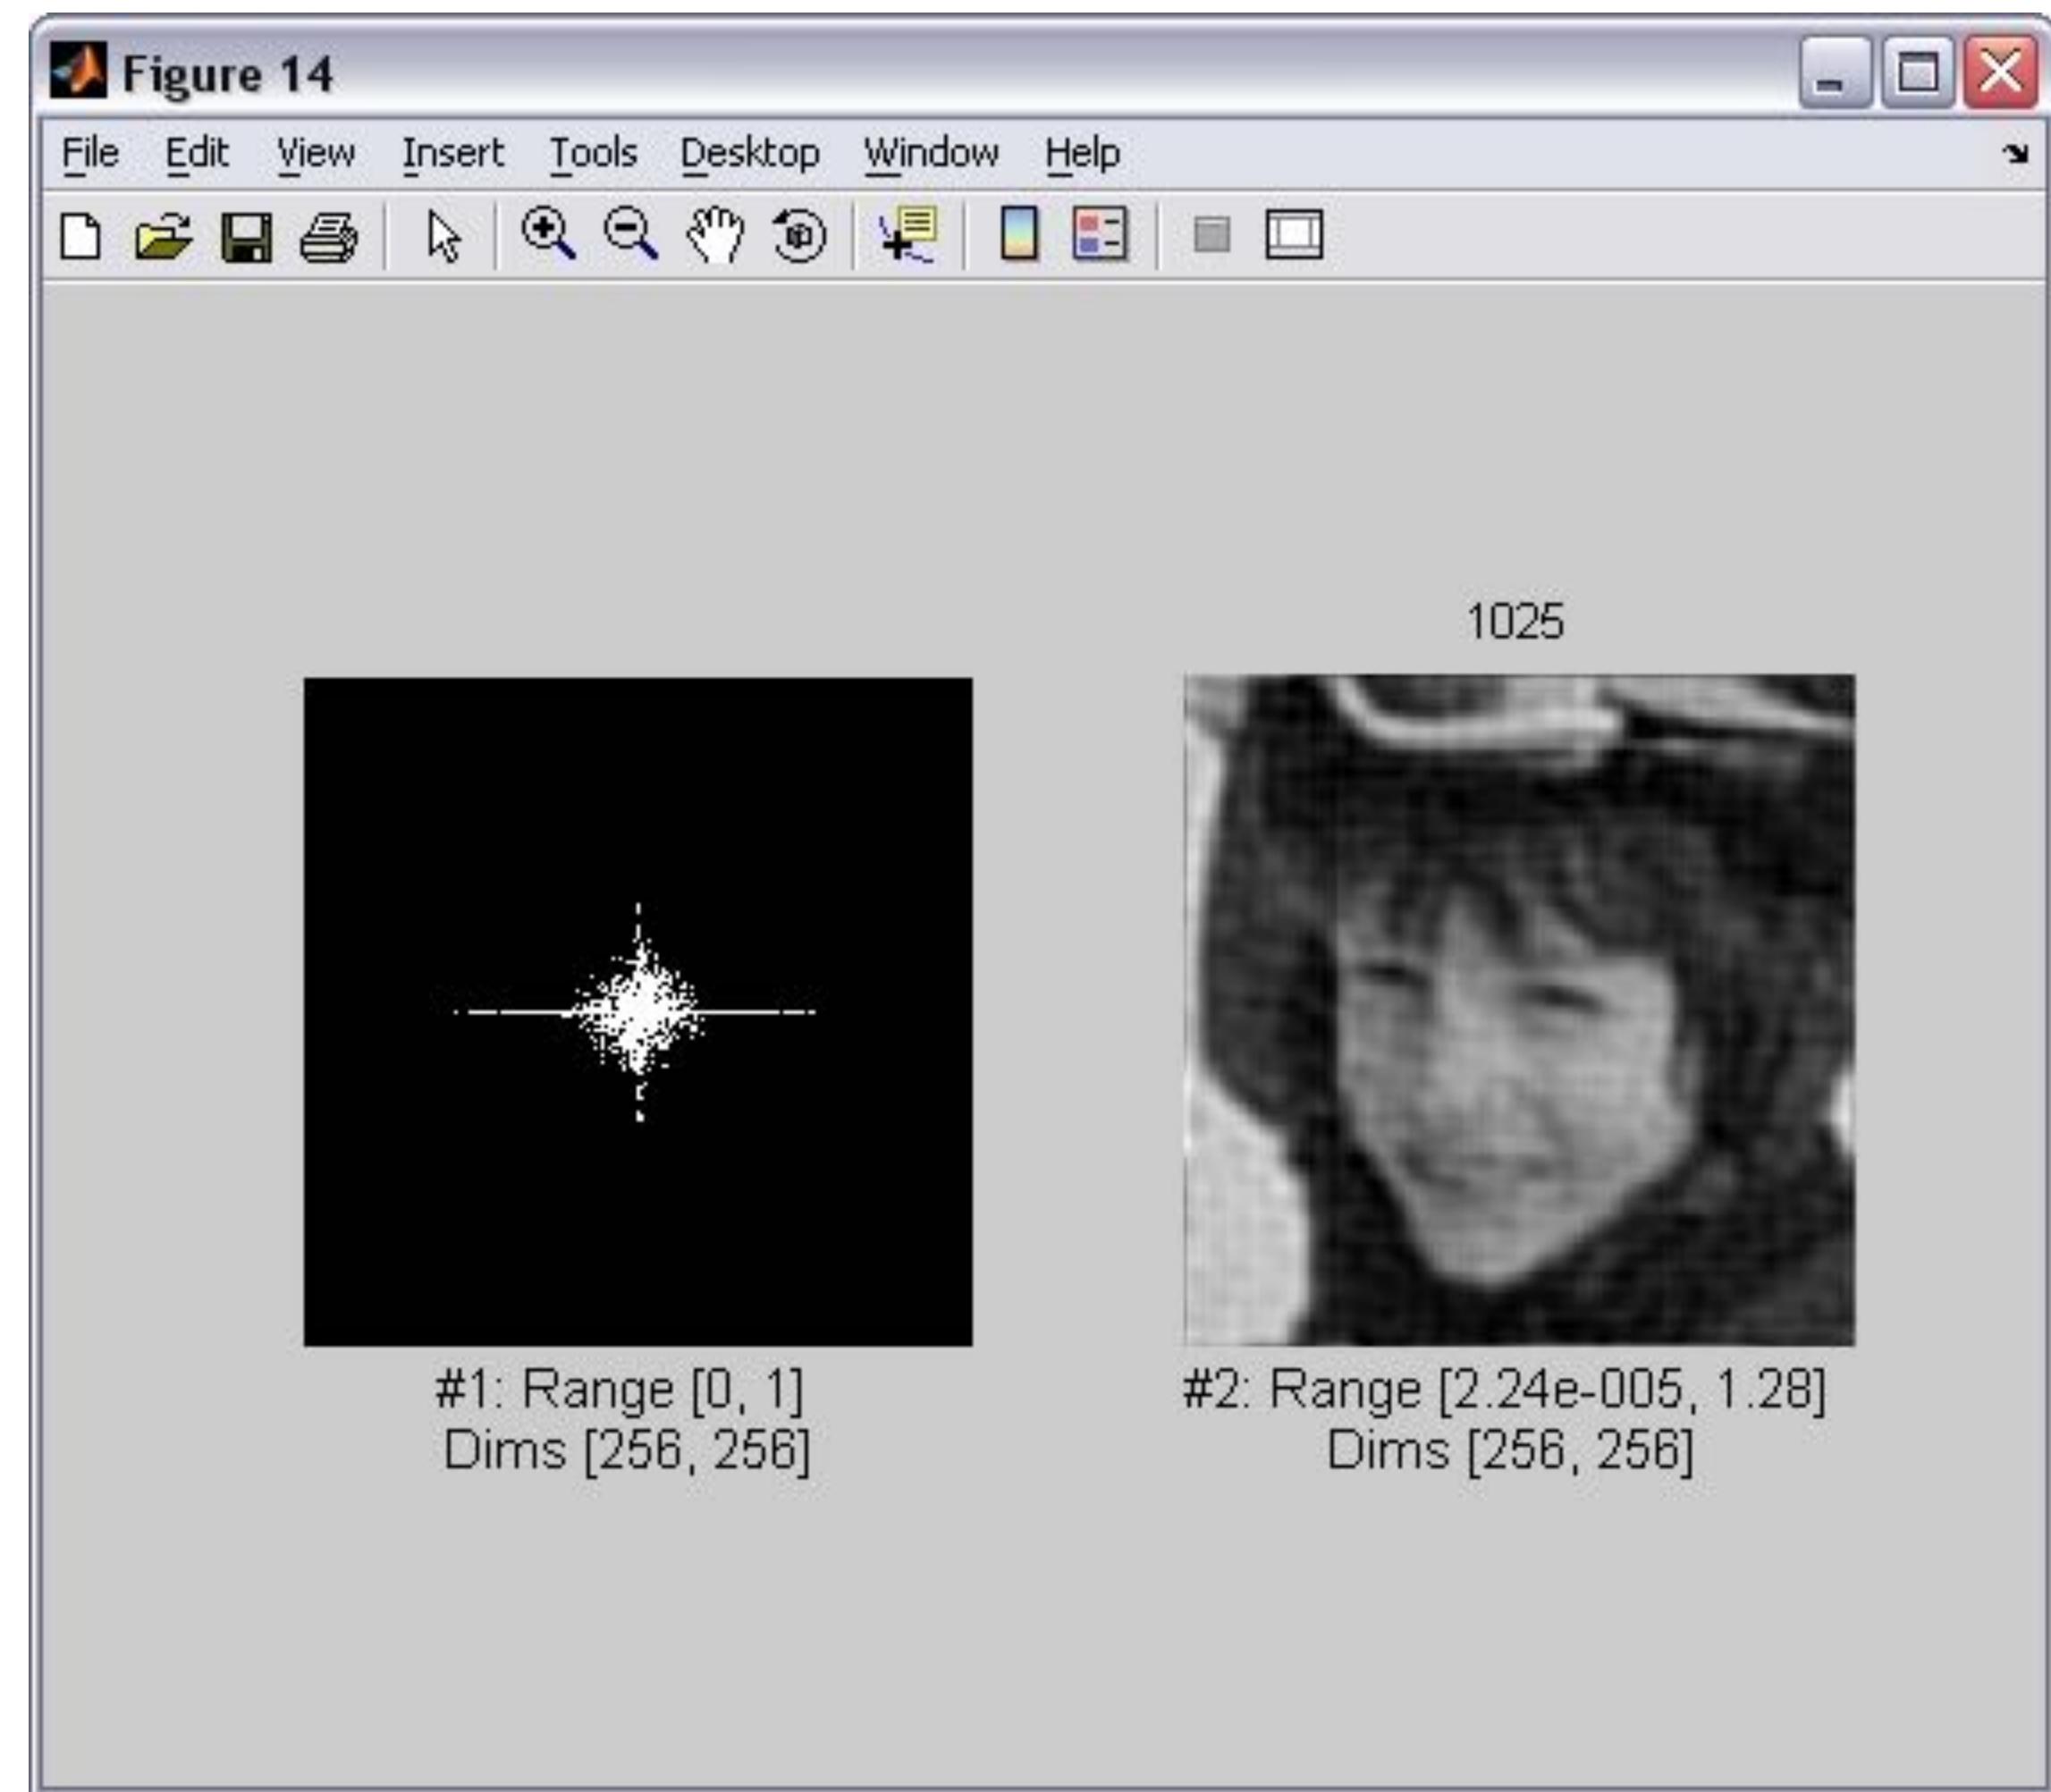
257



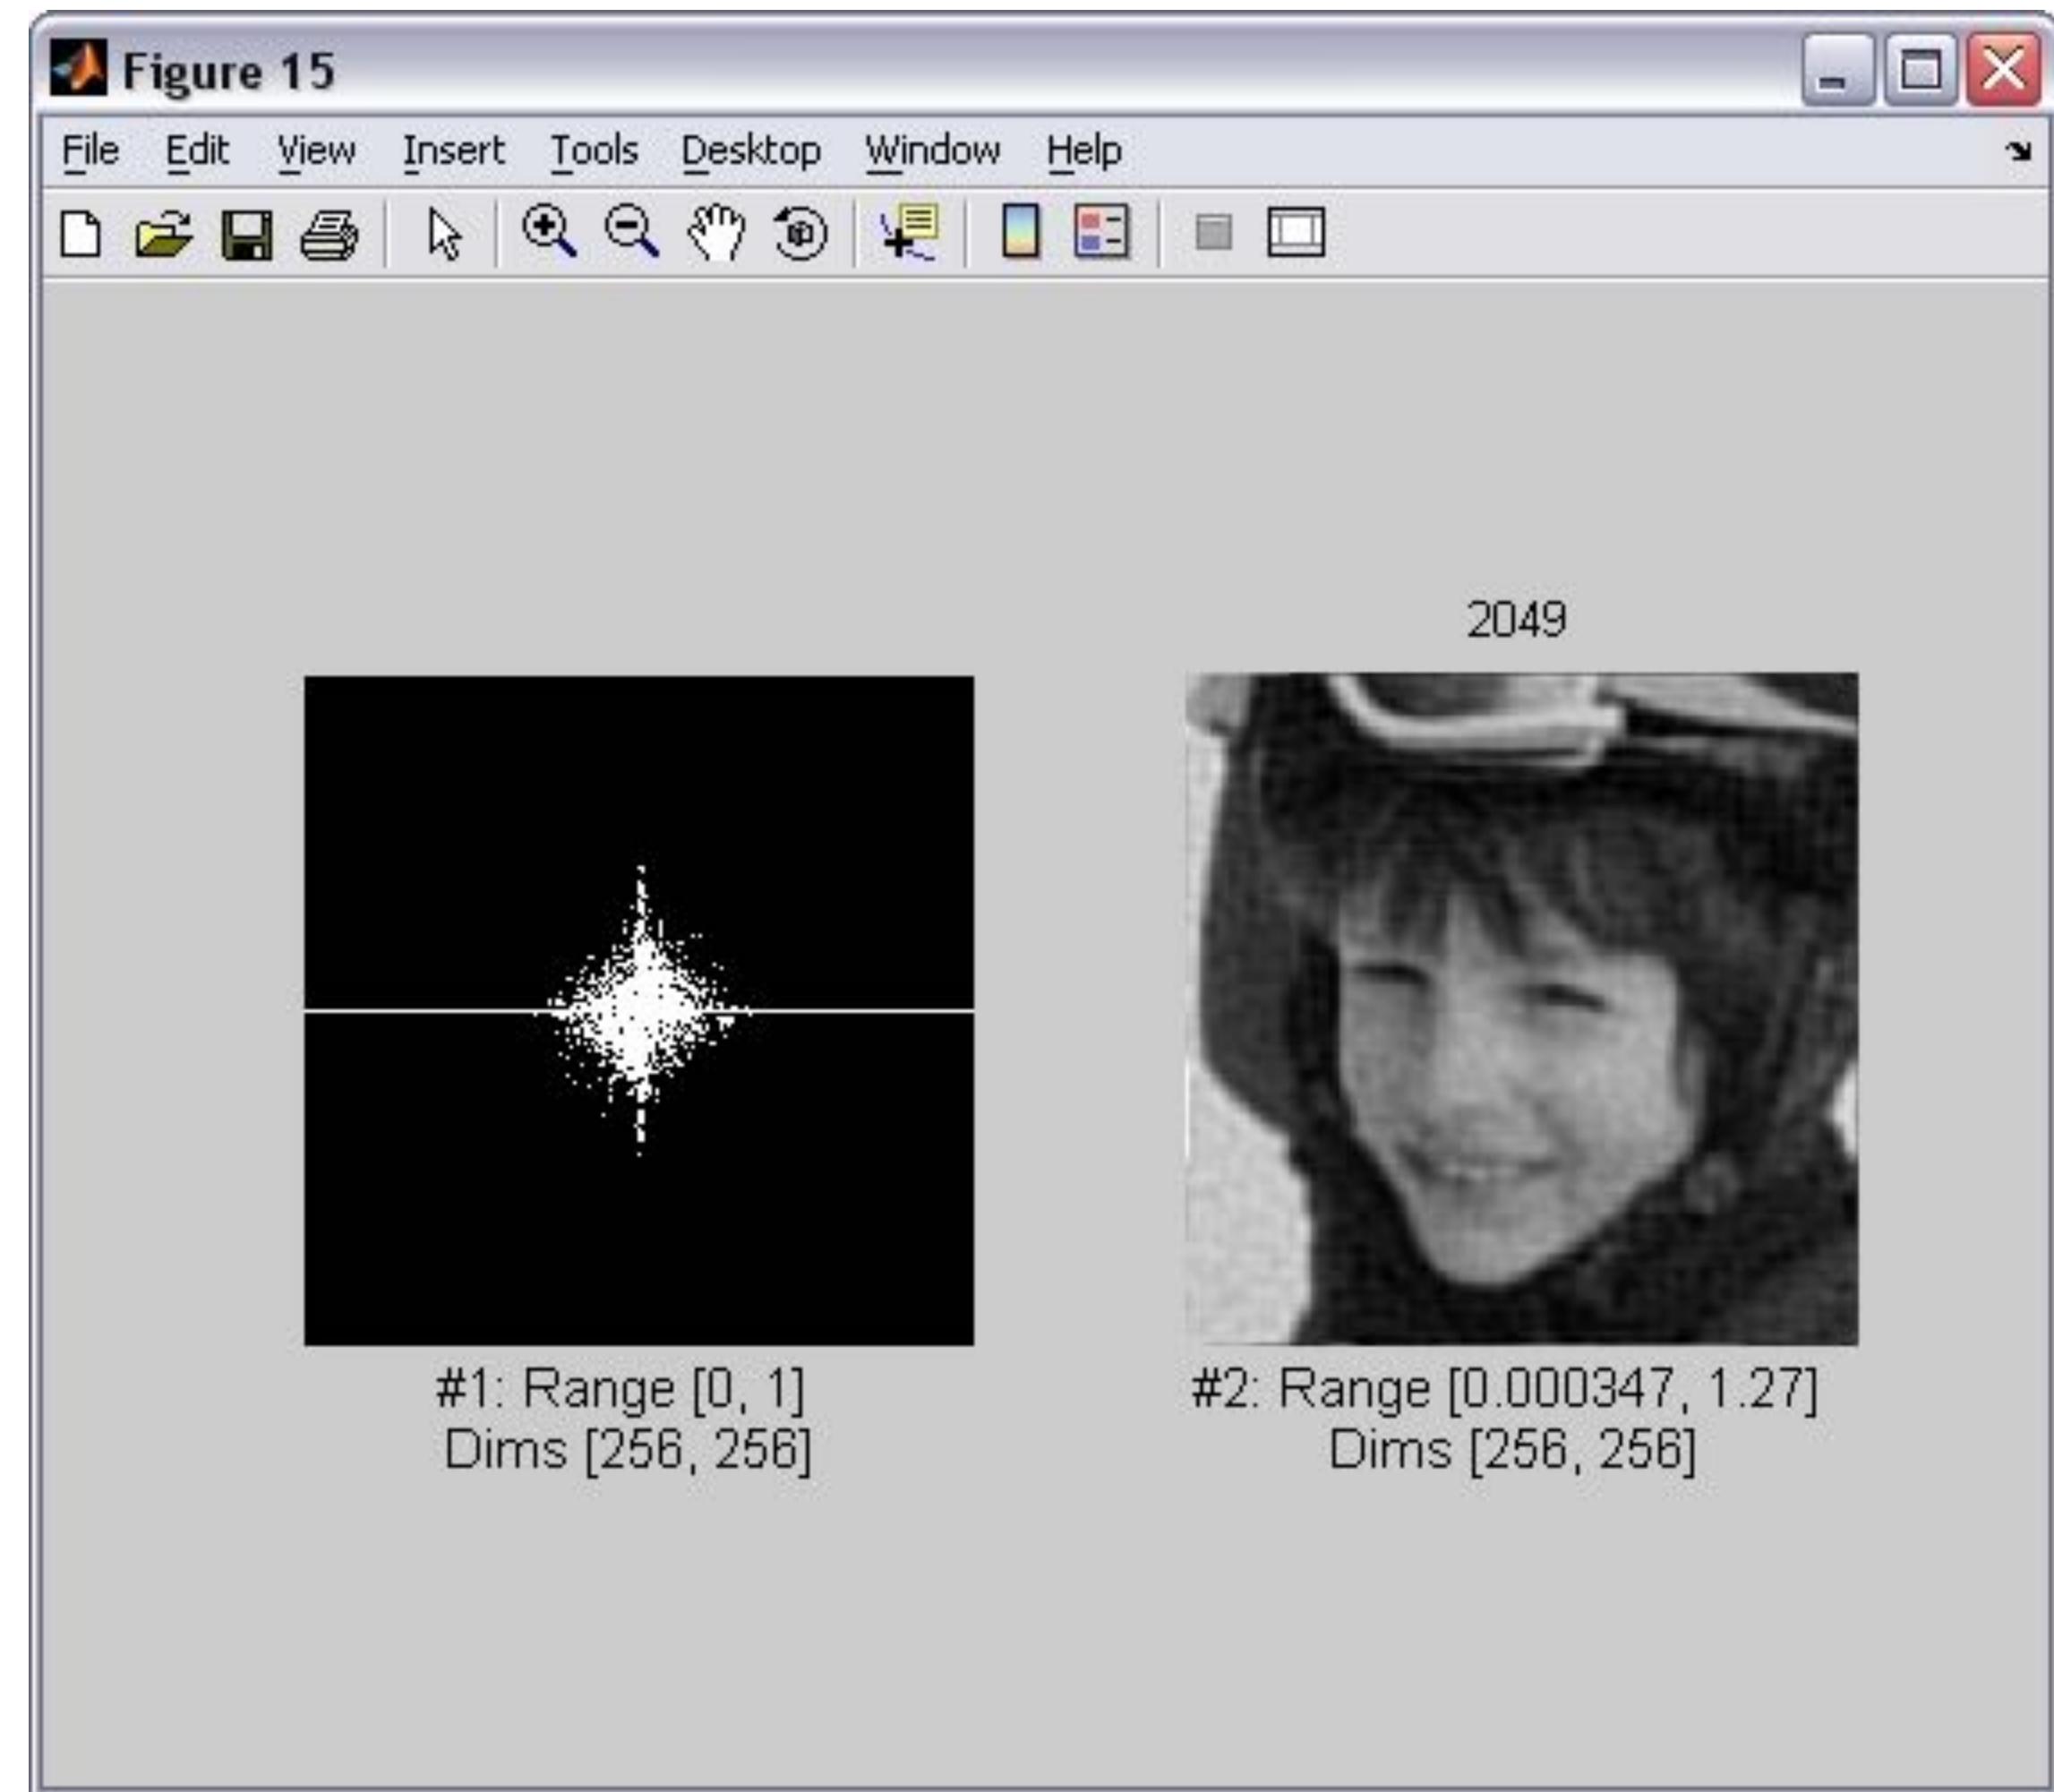
513



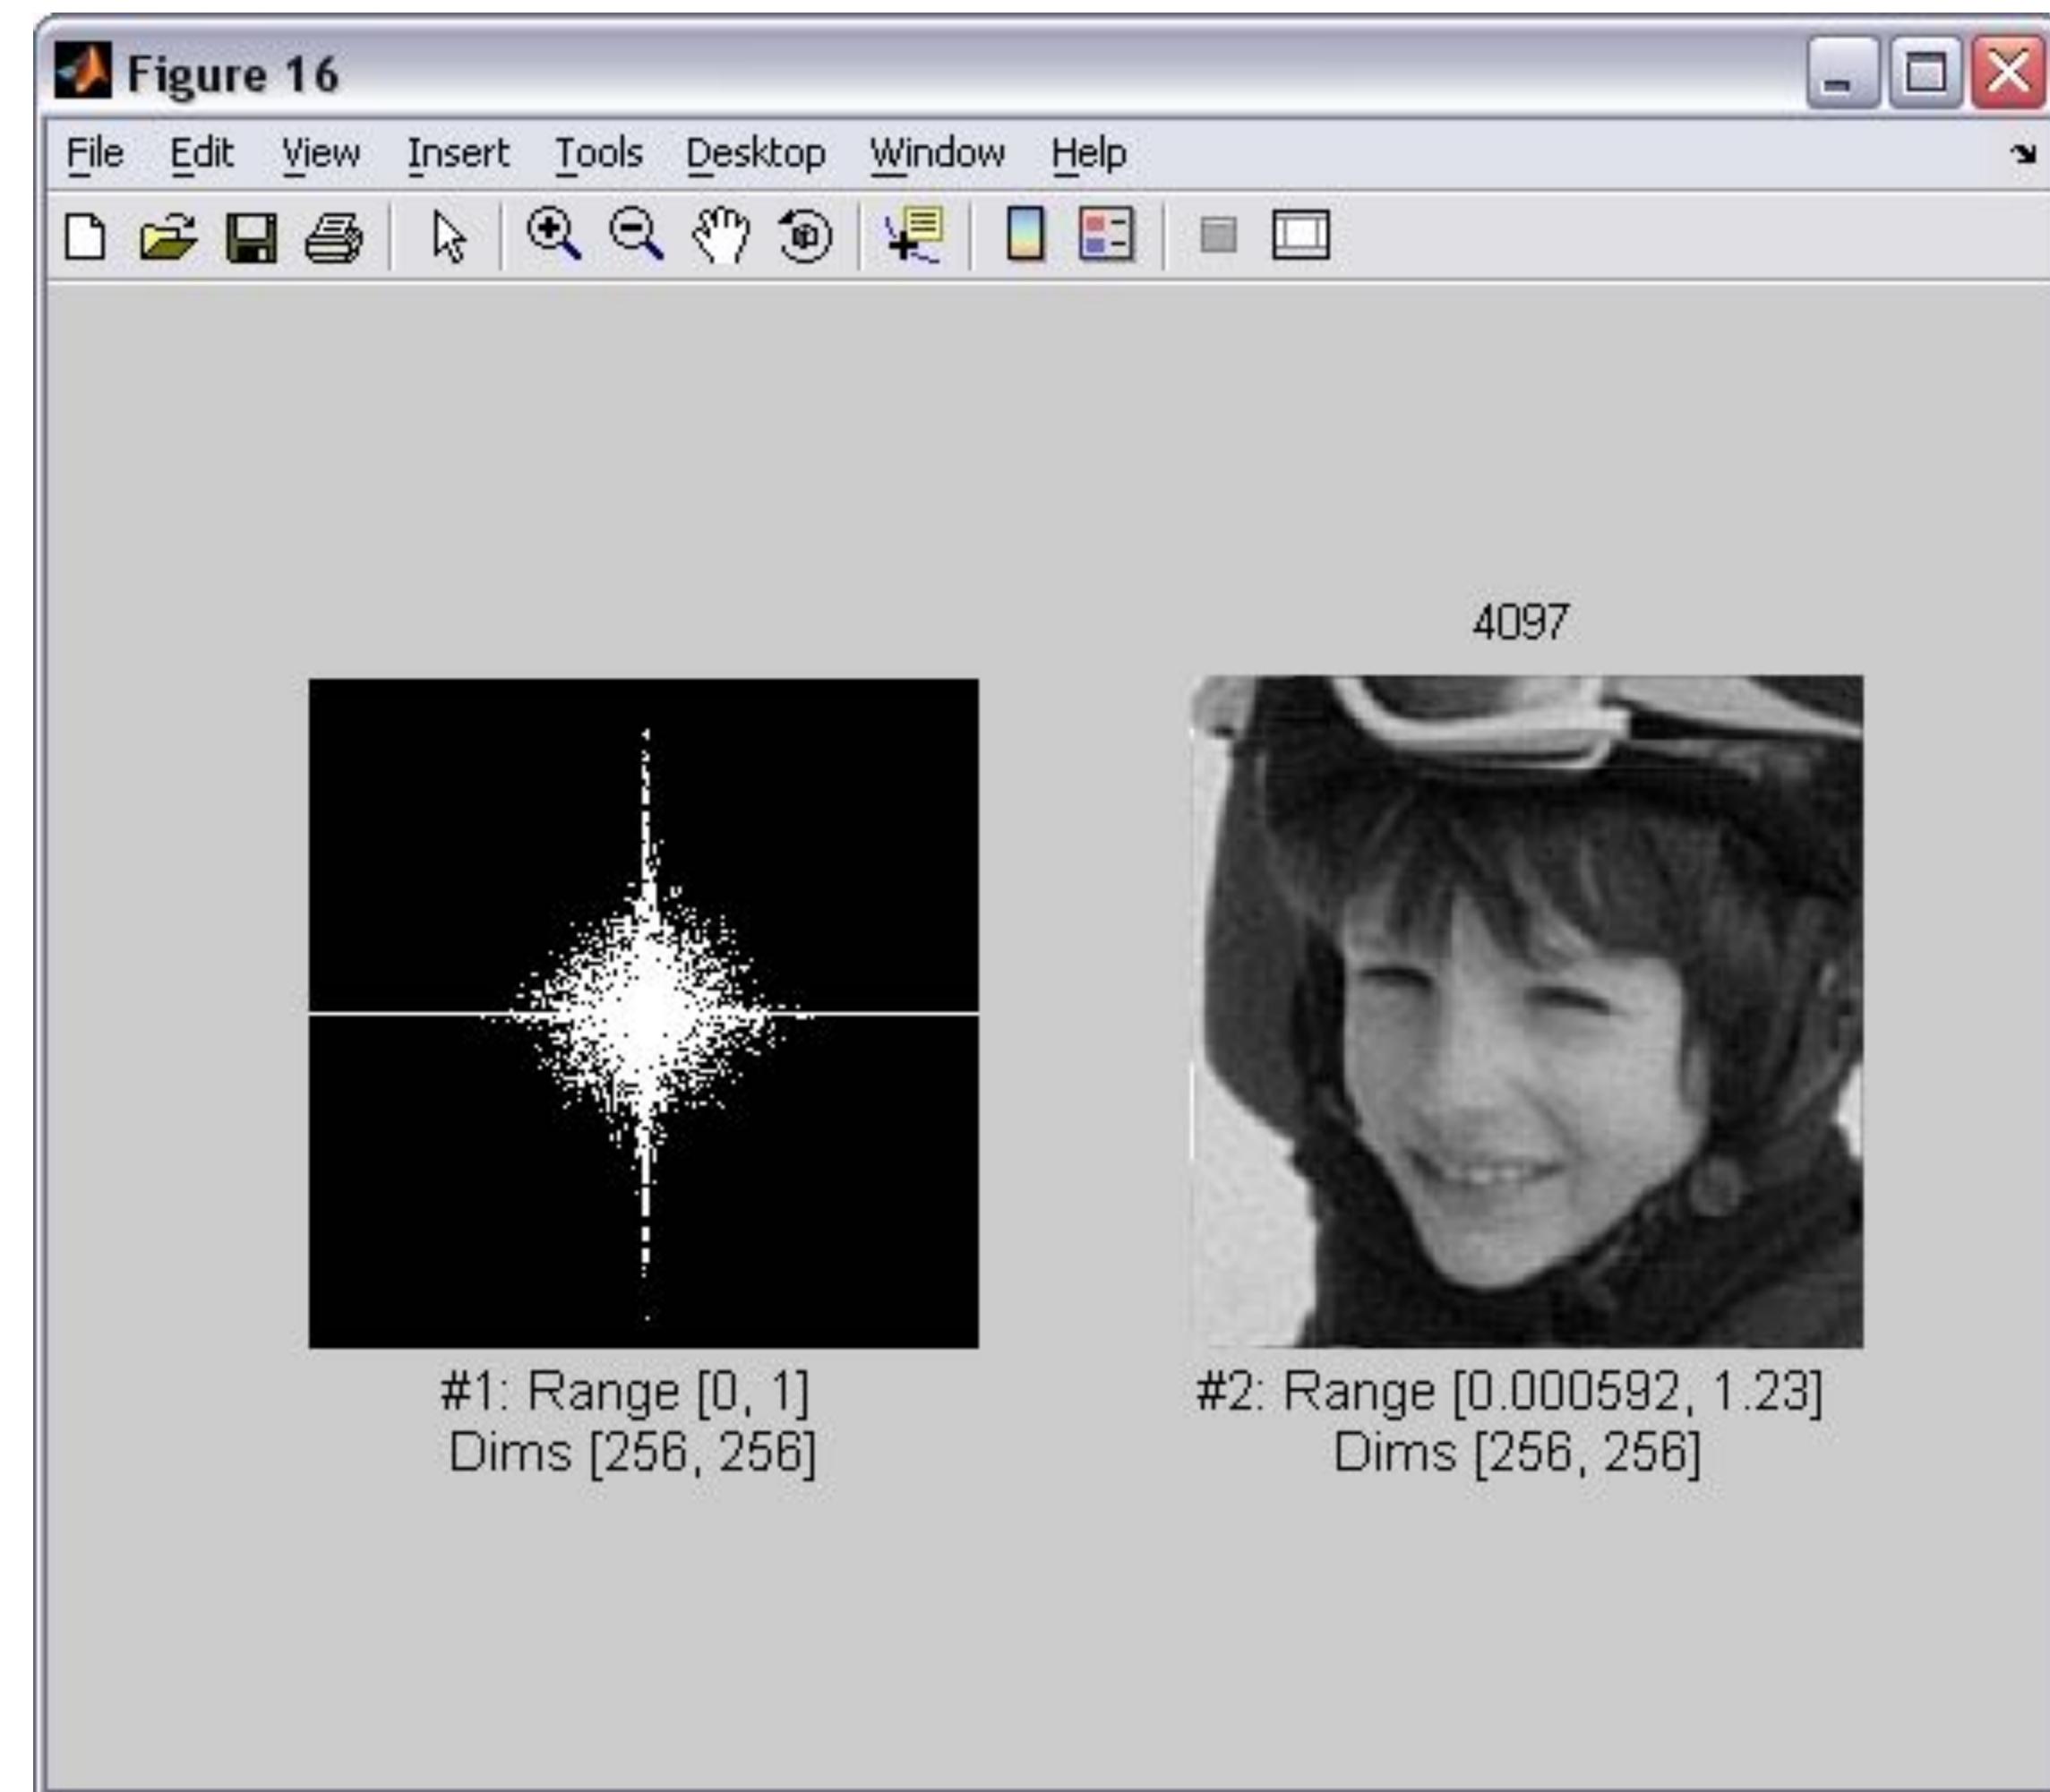
1025



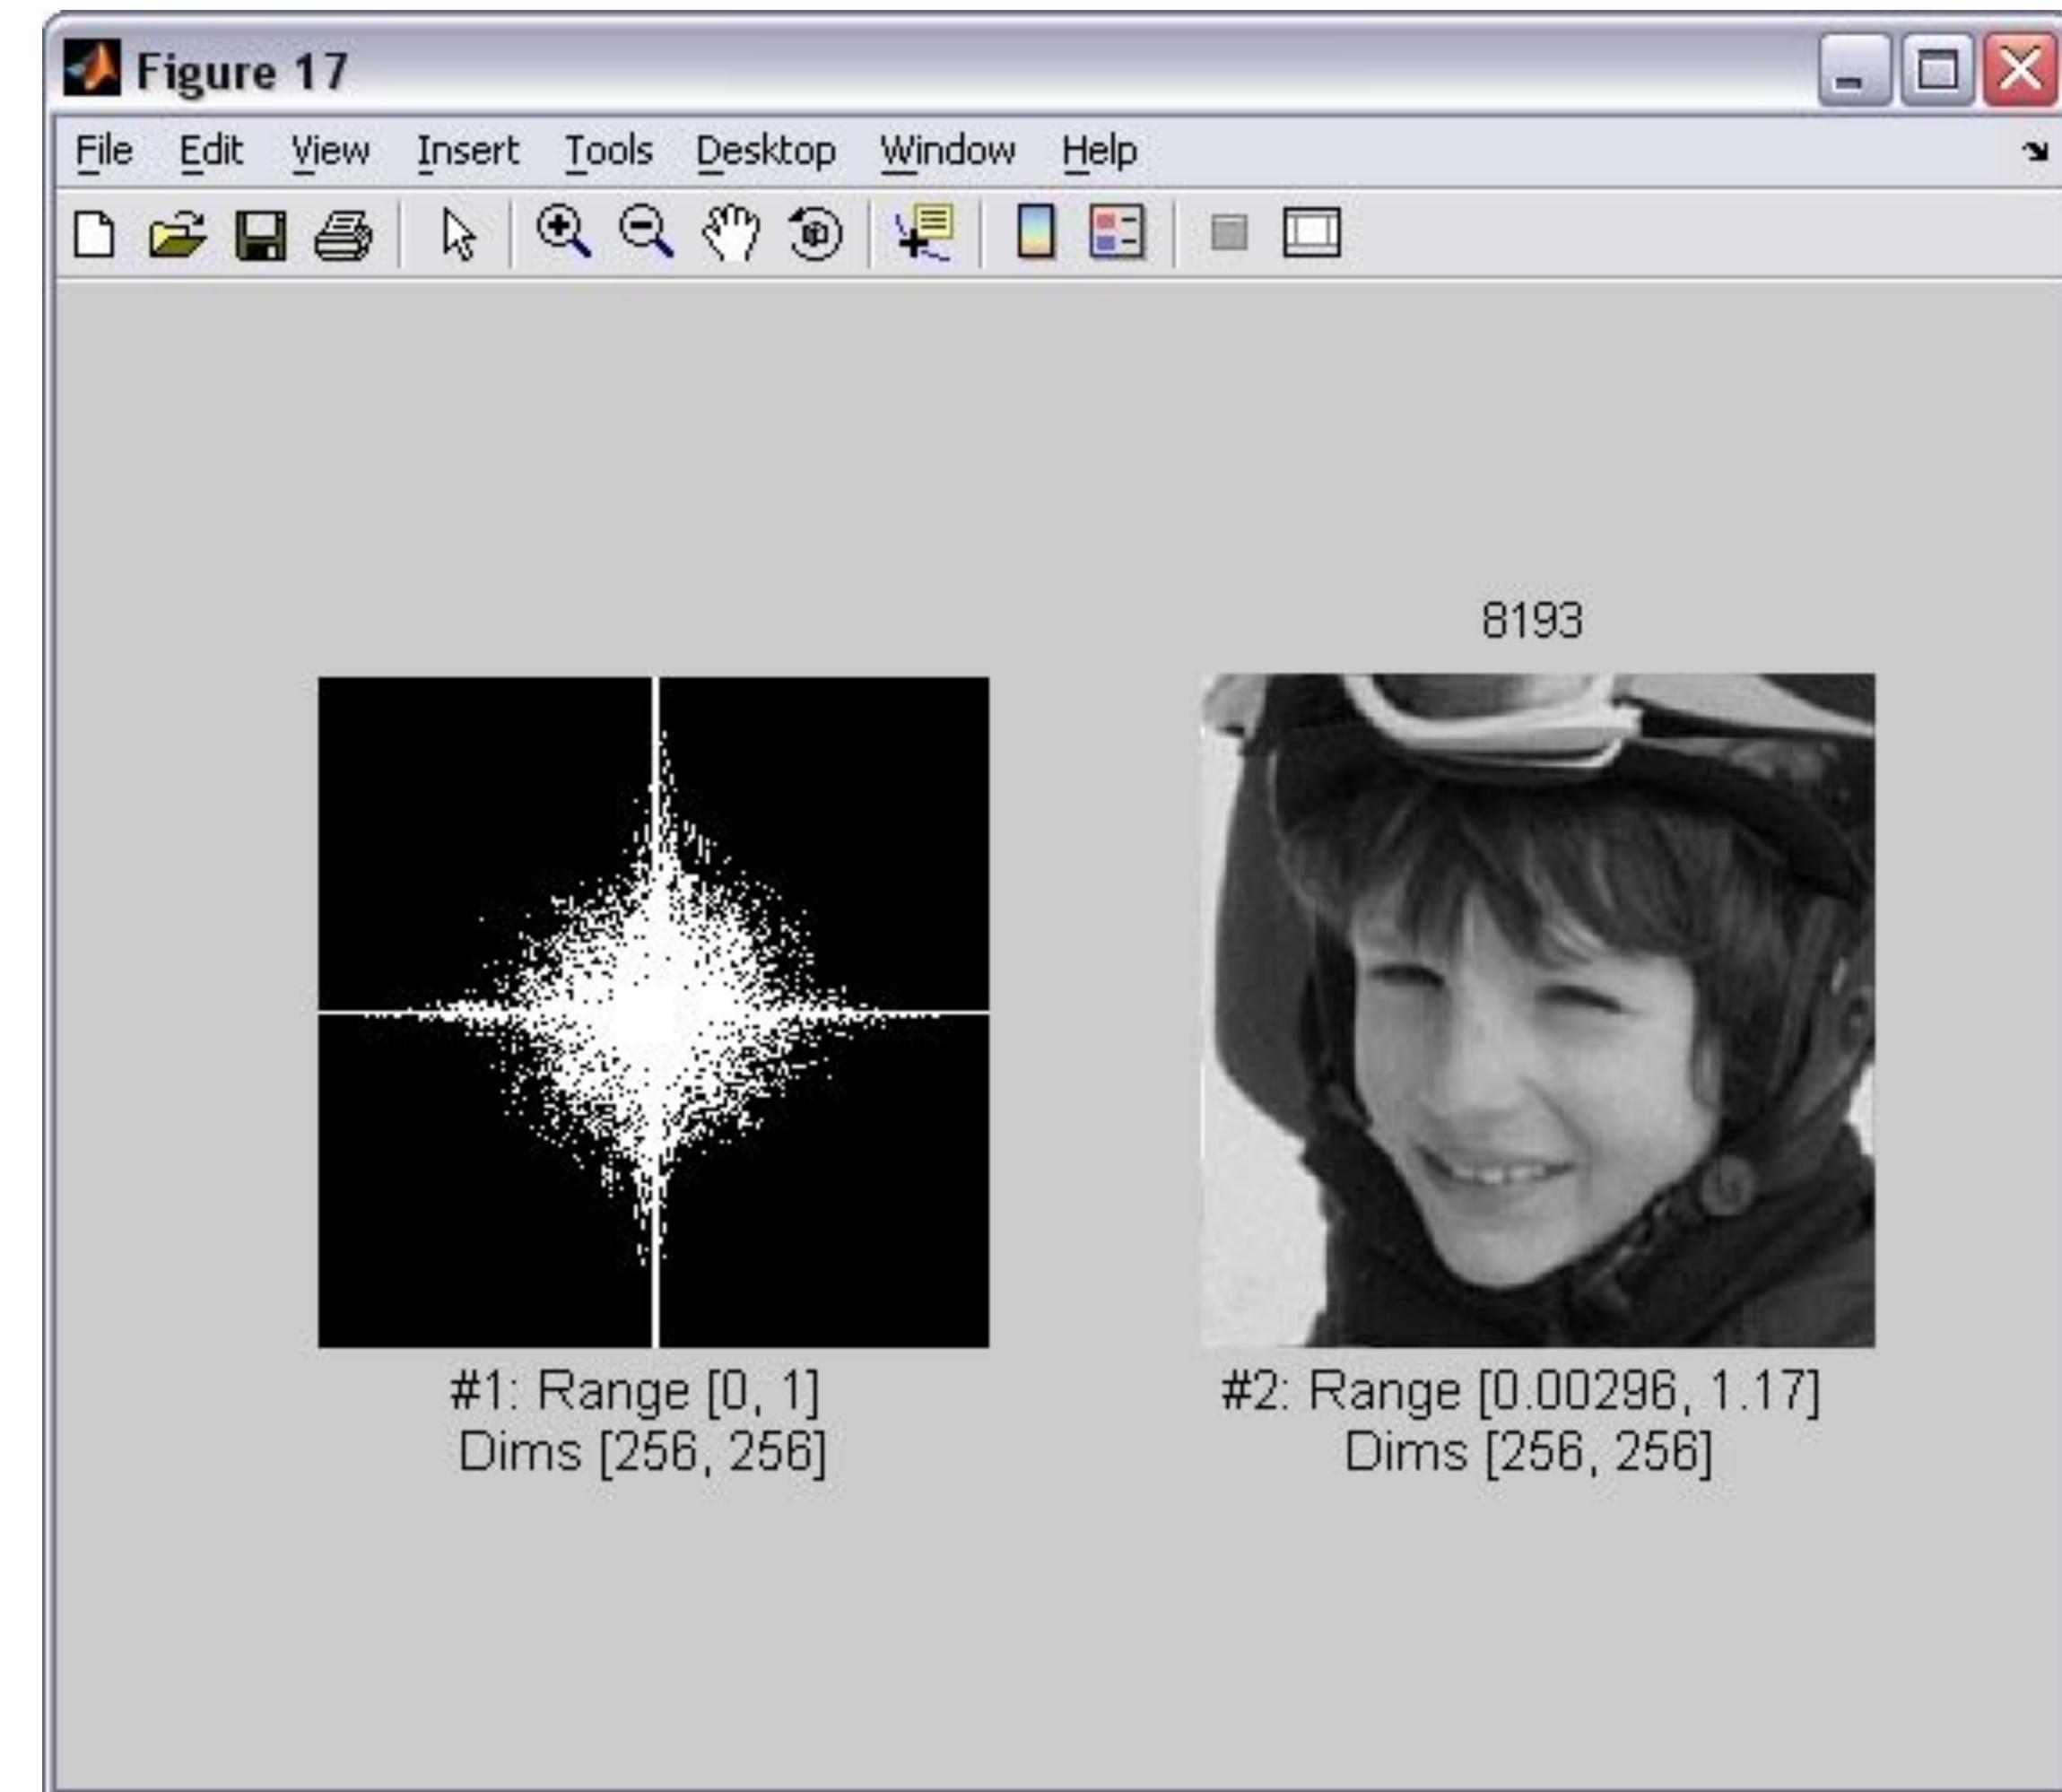
2049



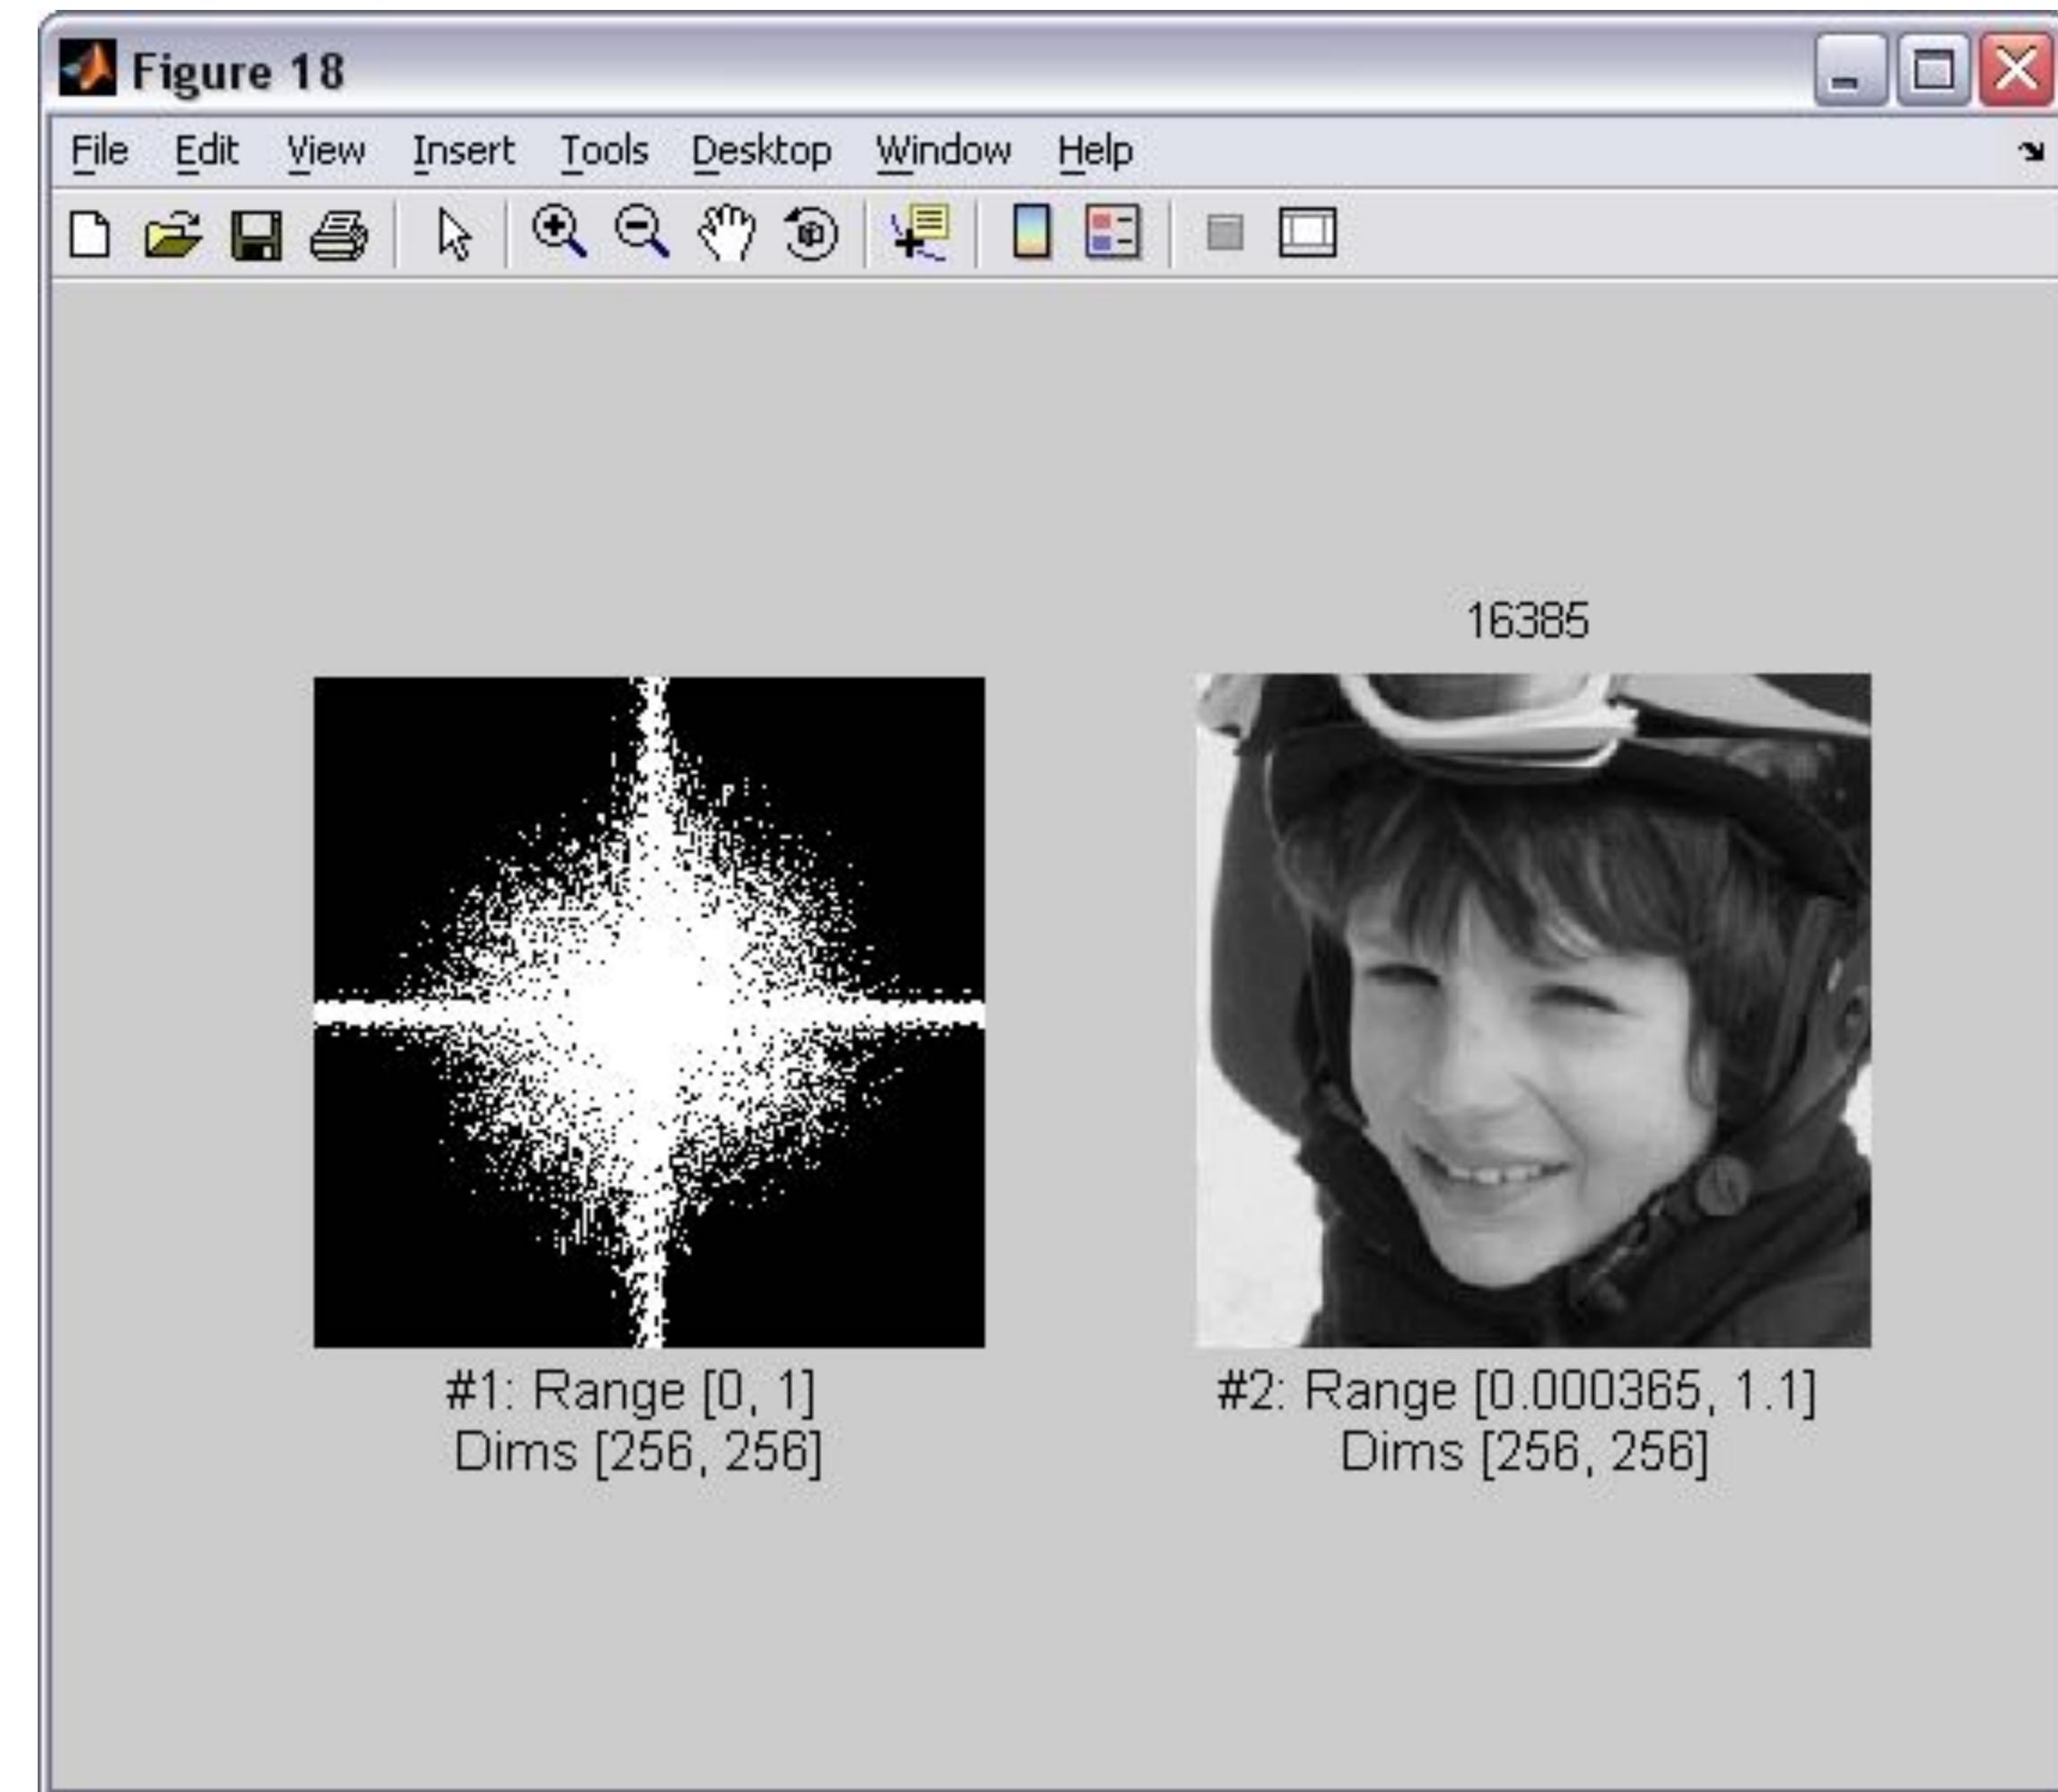
4097



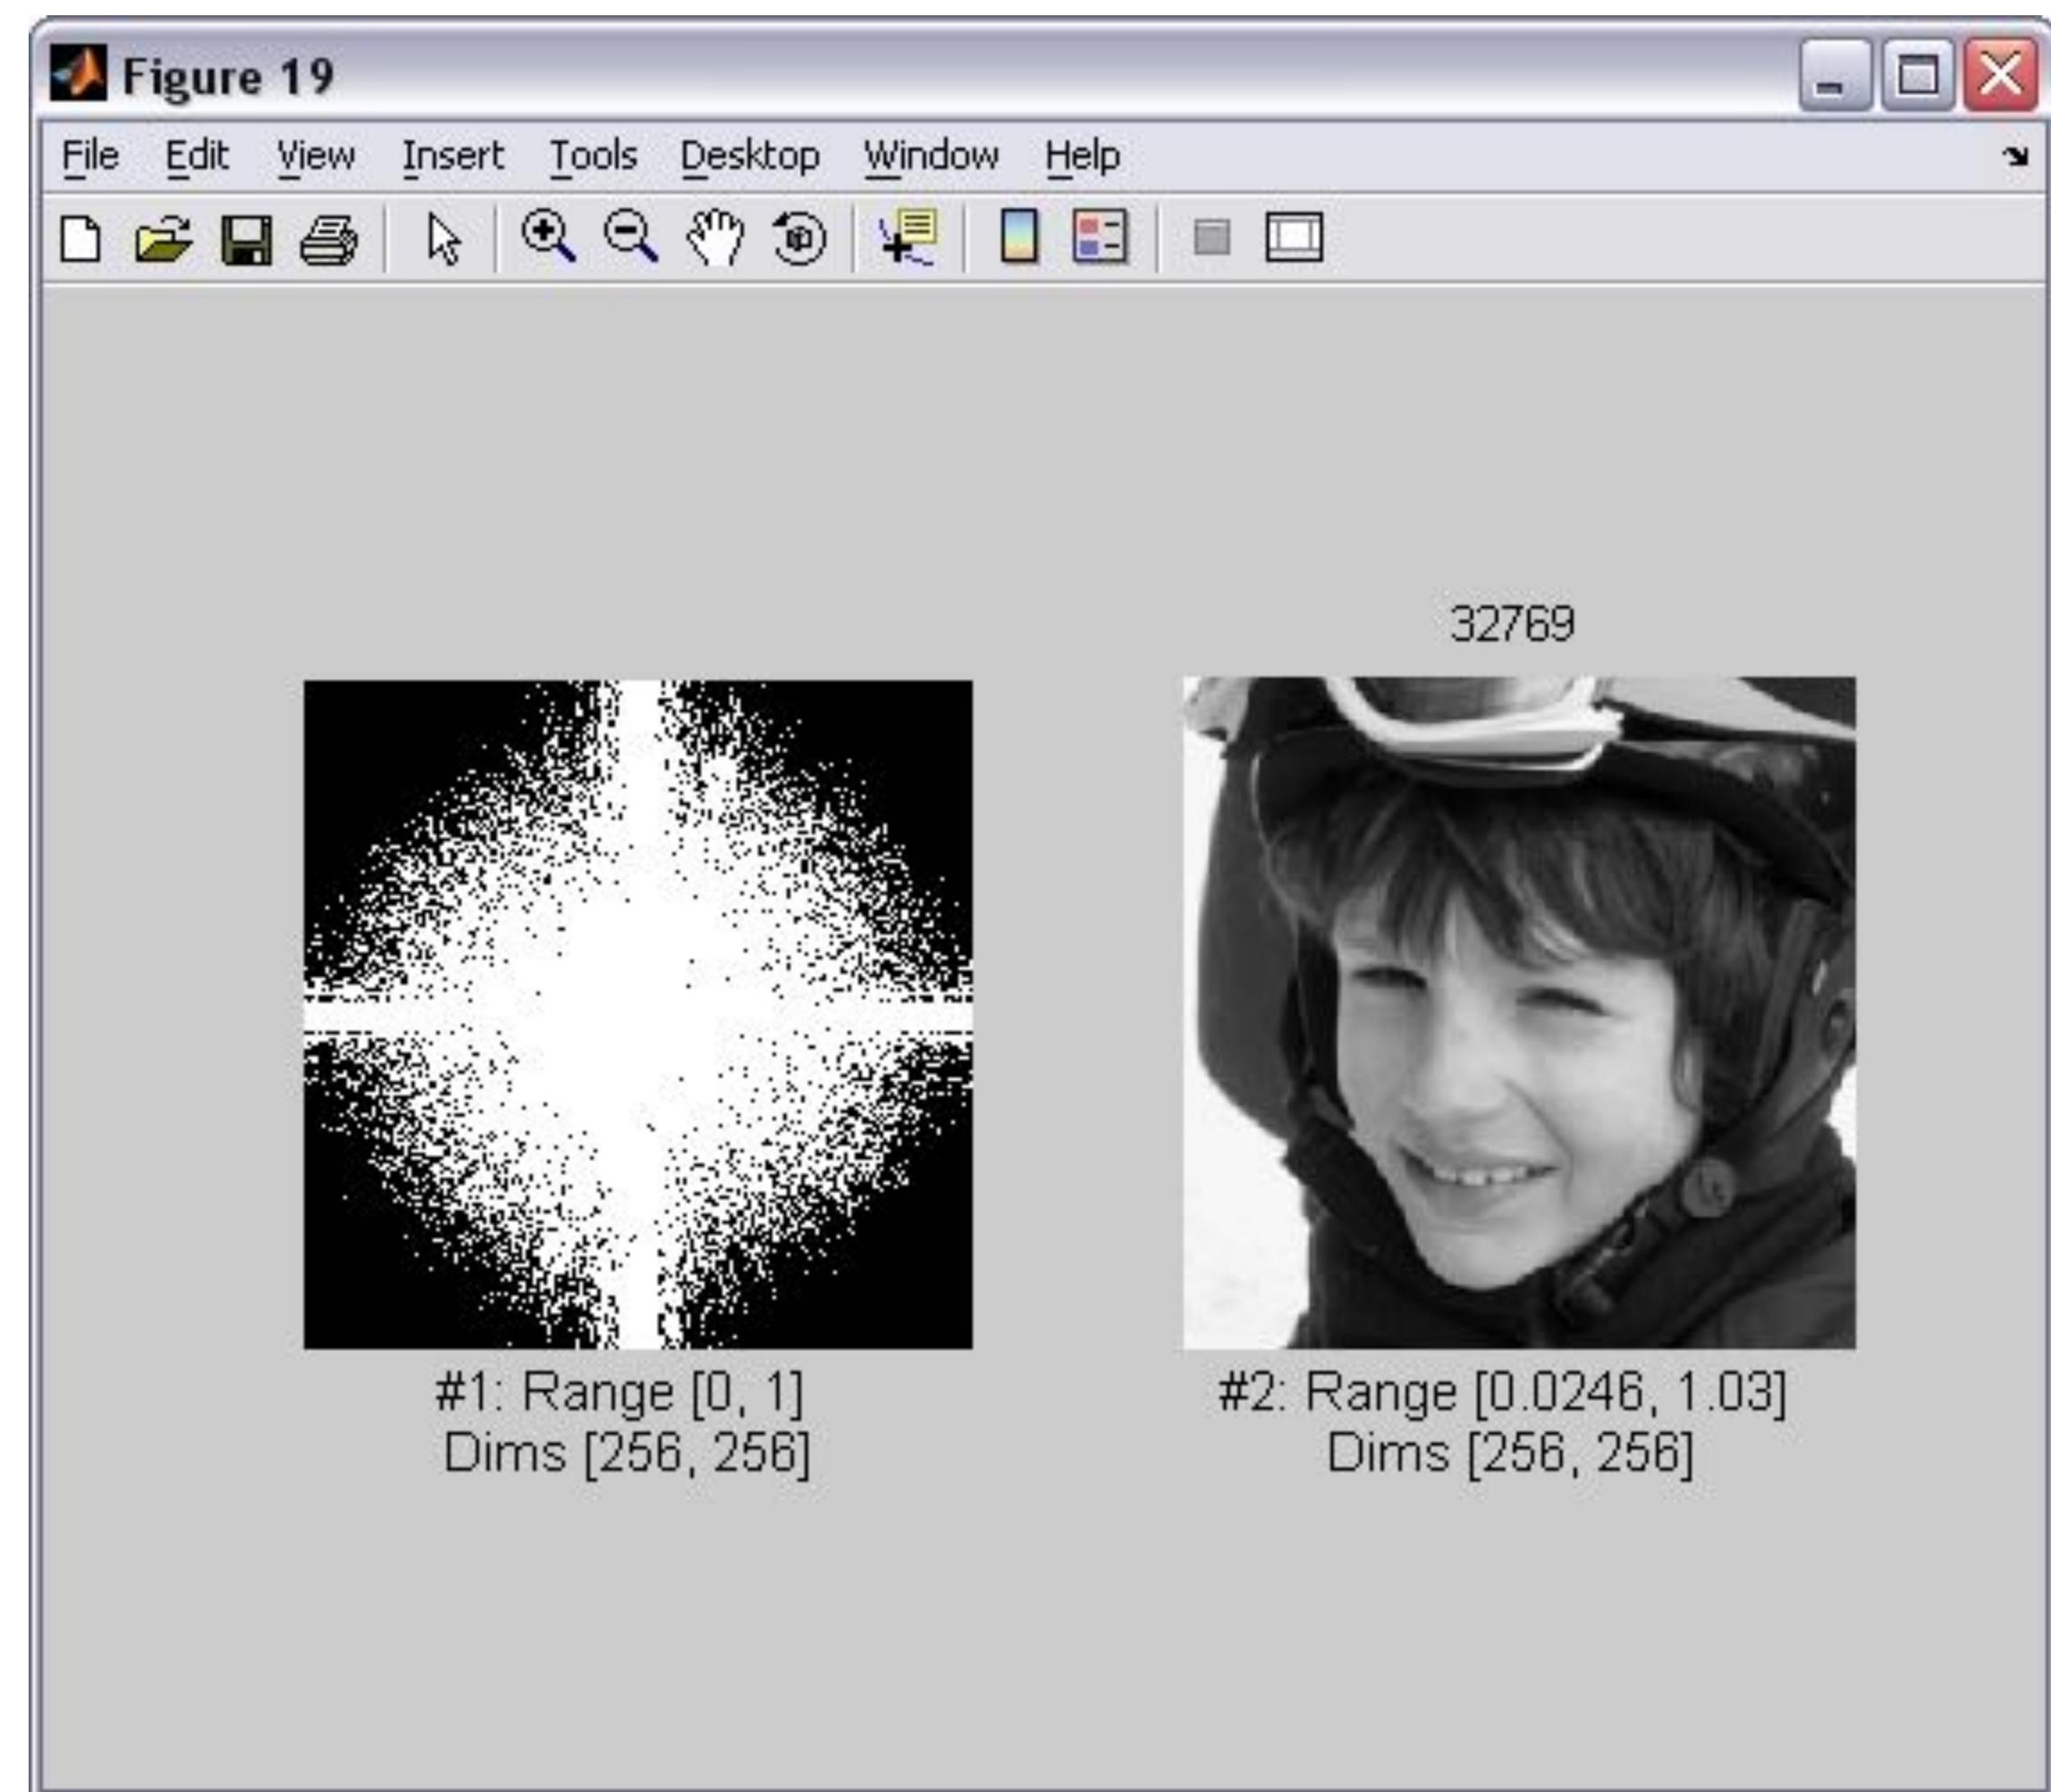
8193



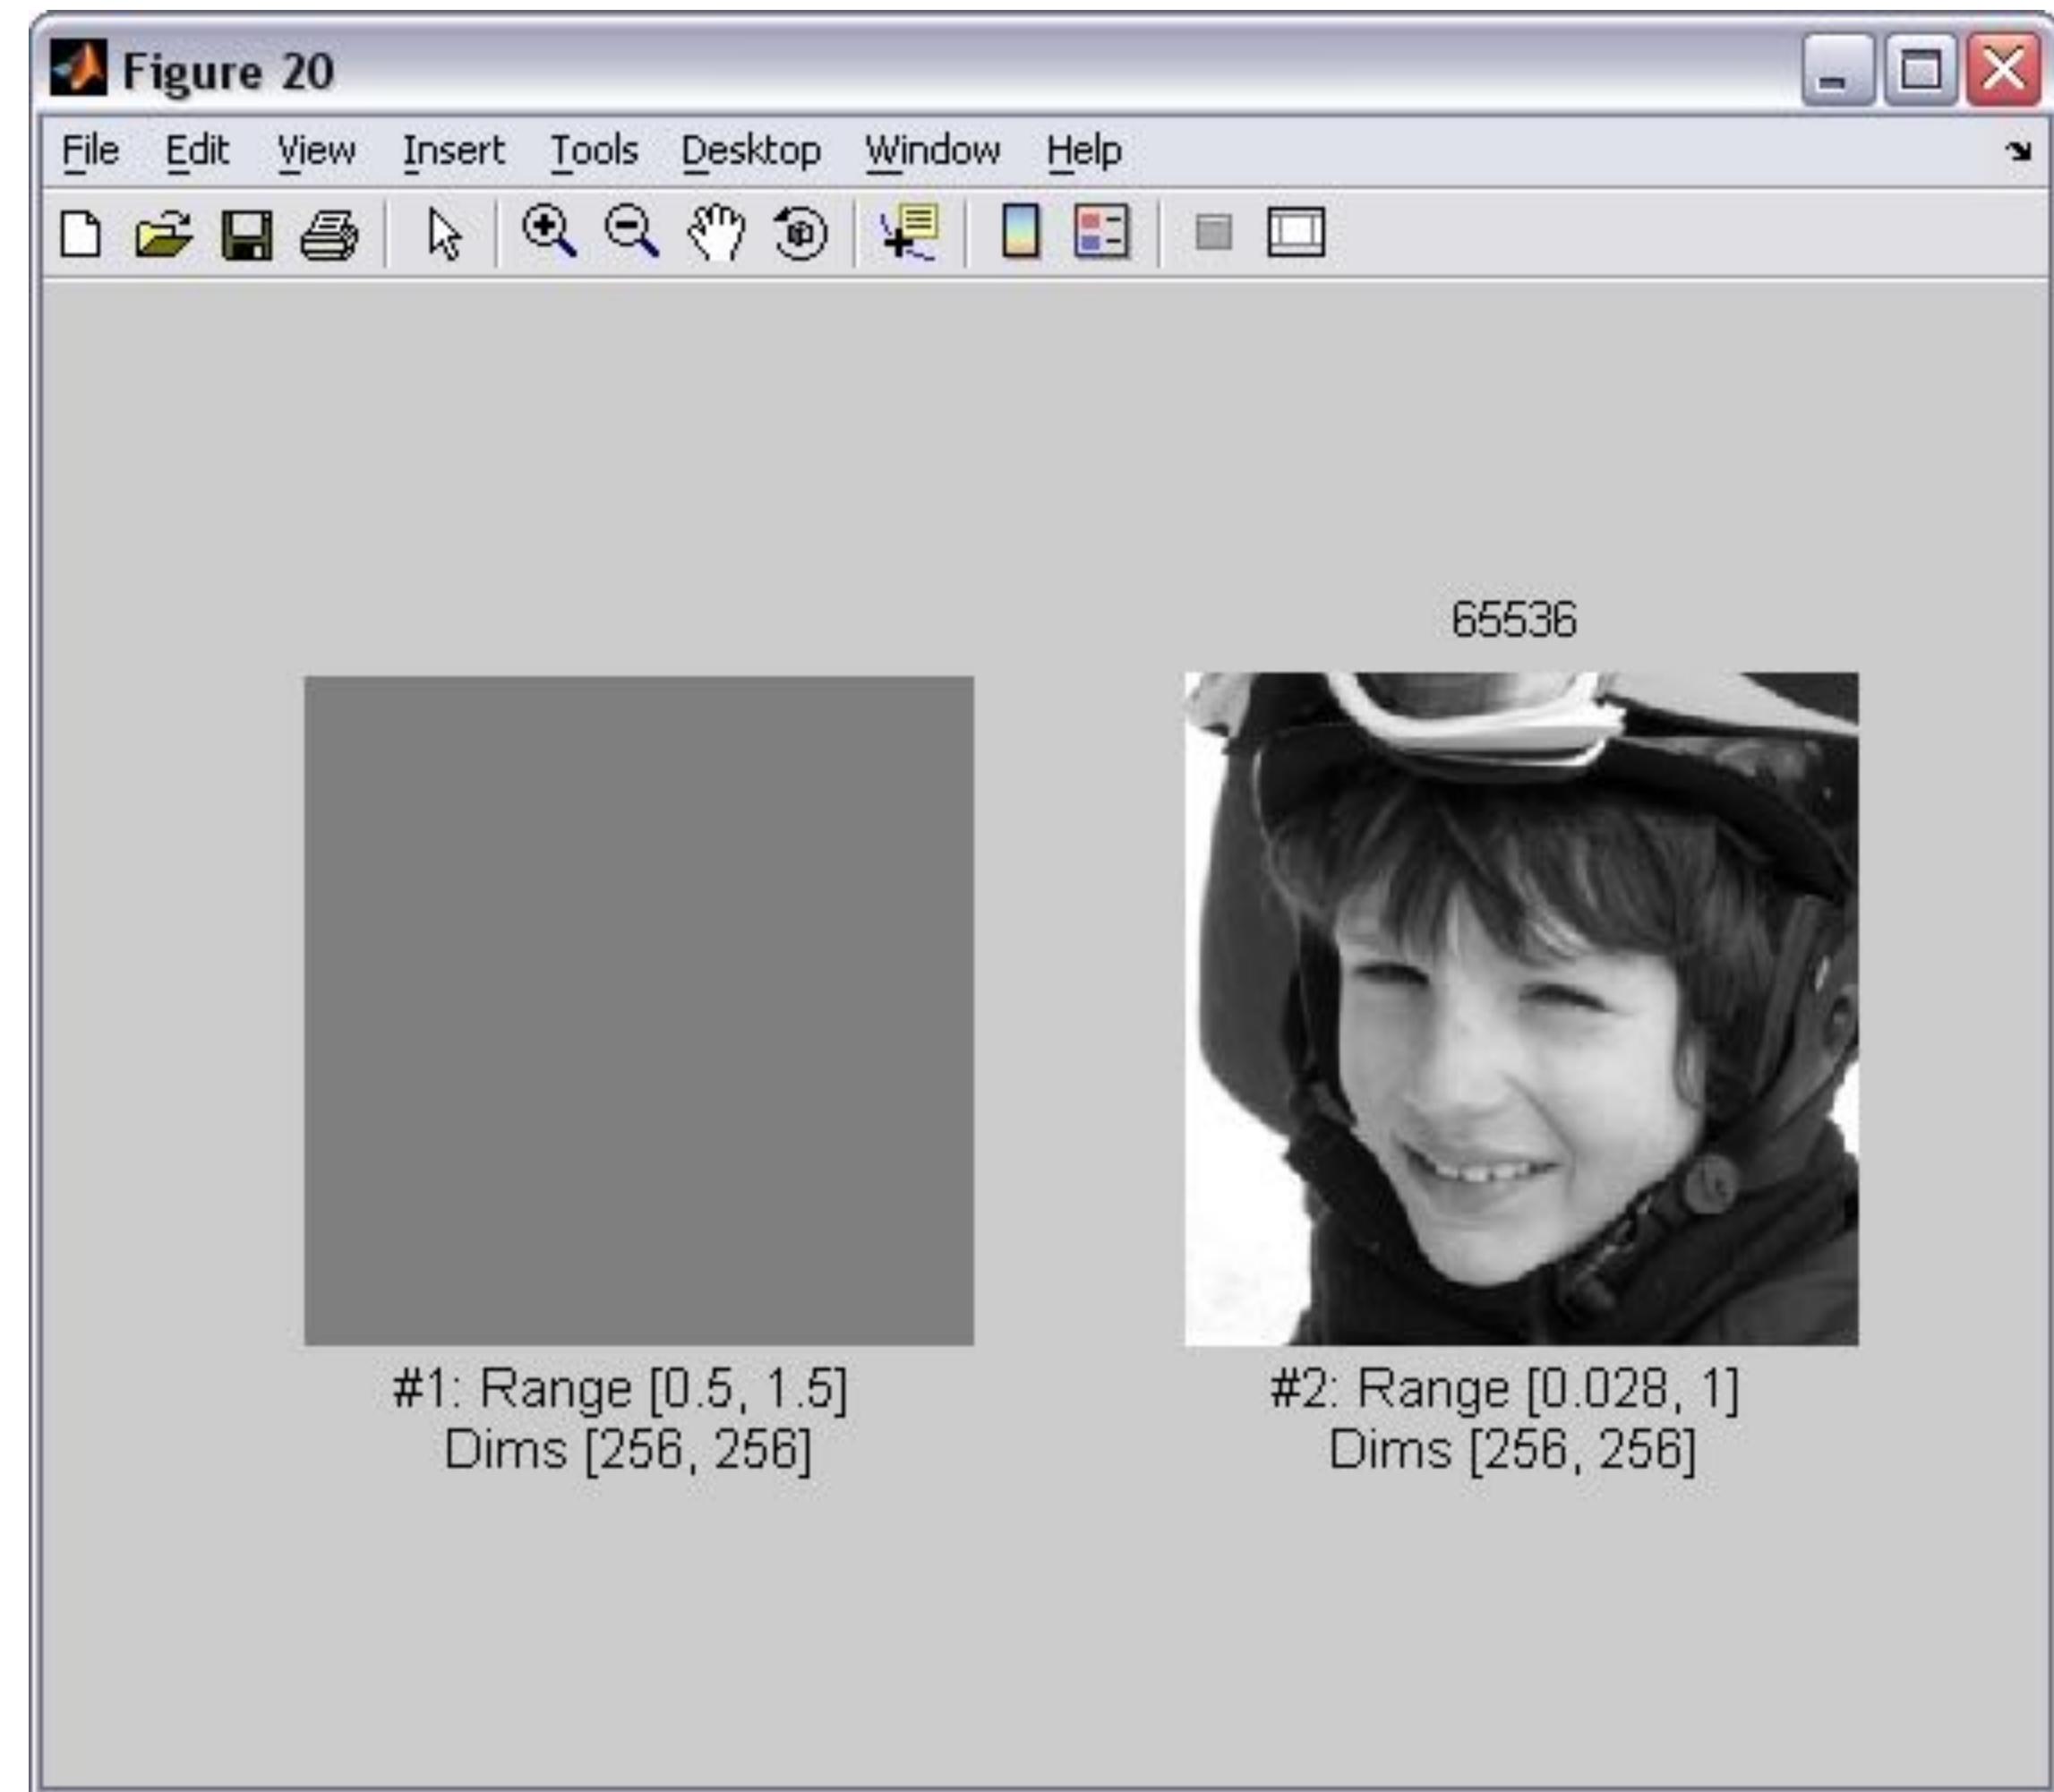
16385



32769

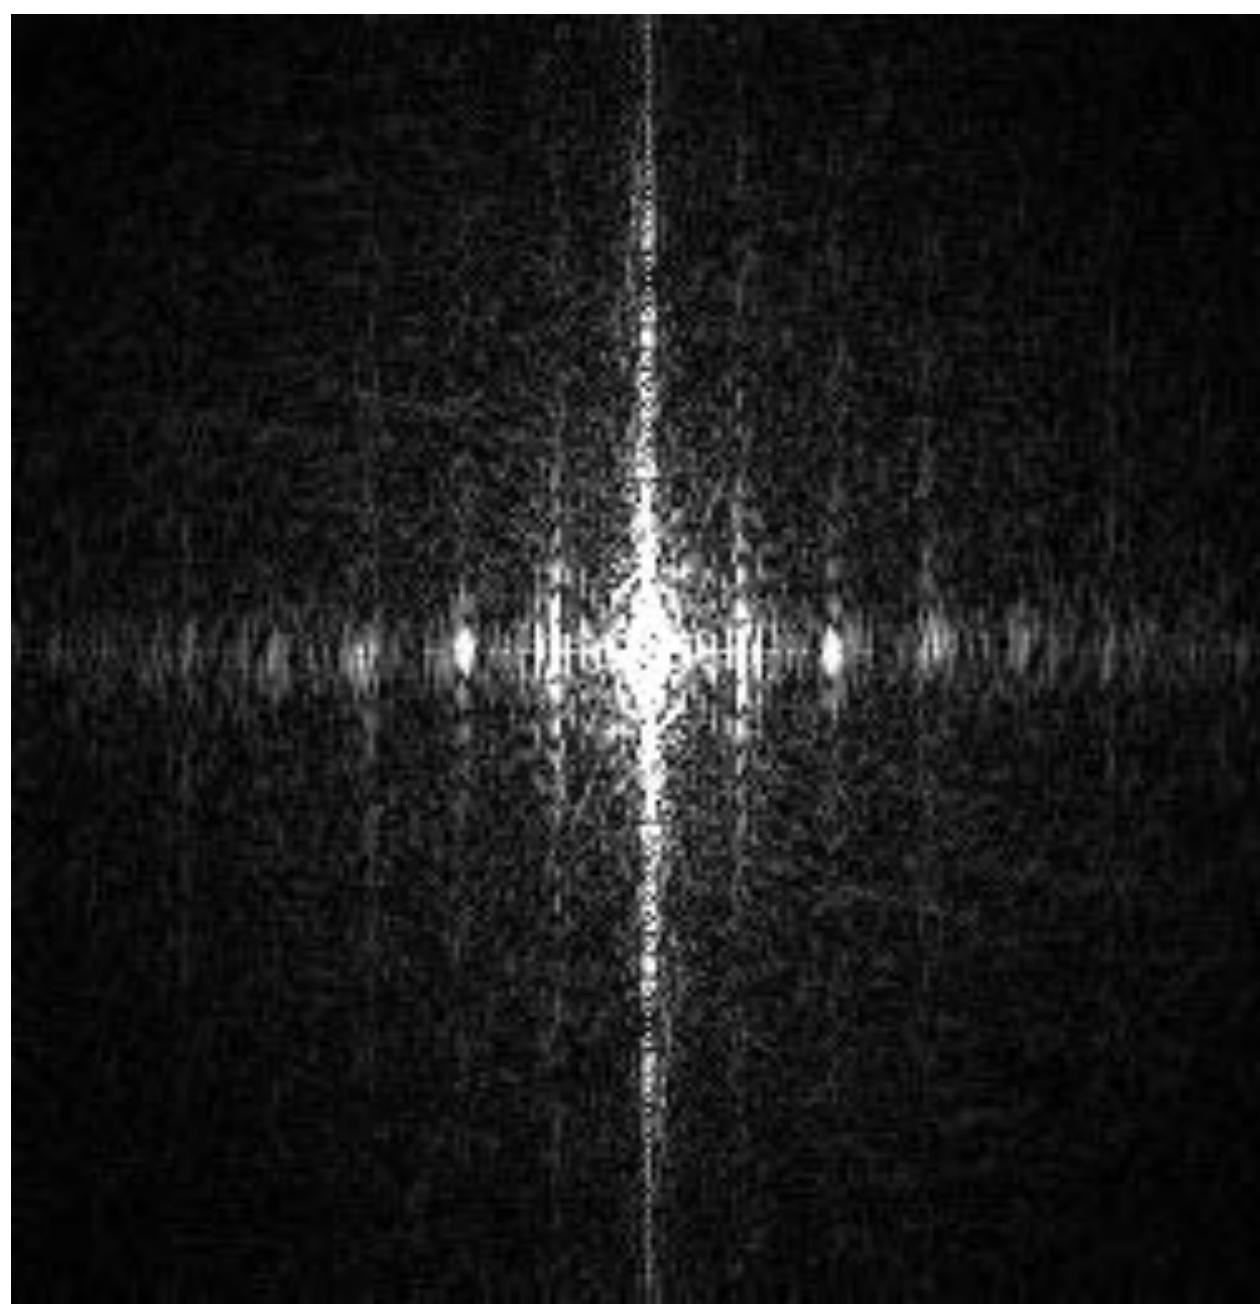


65536

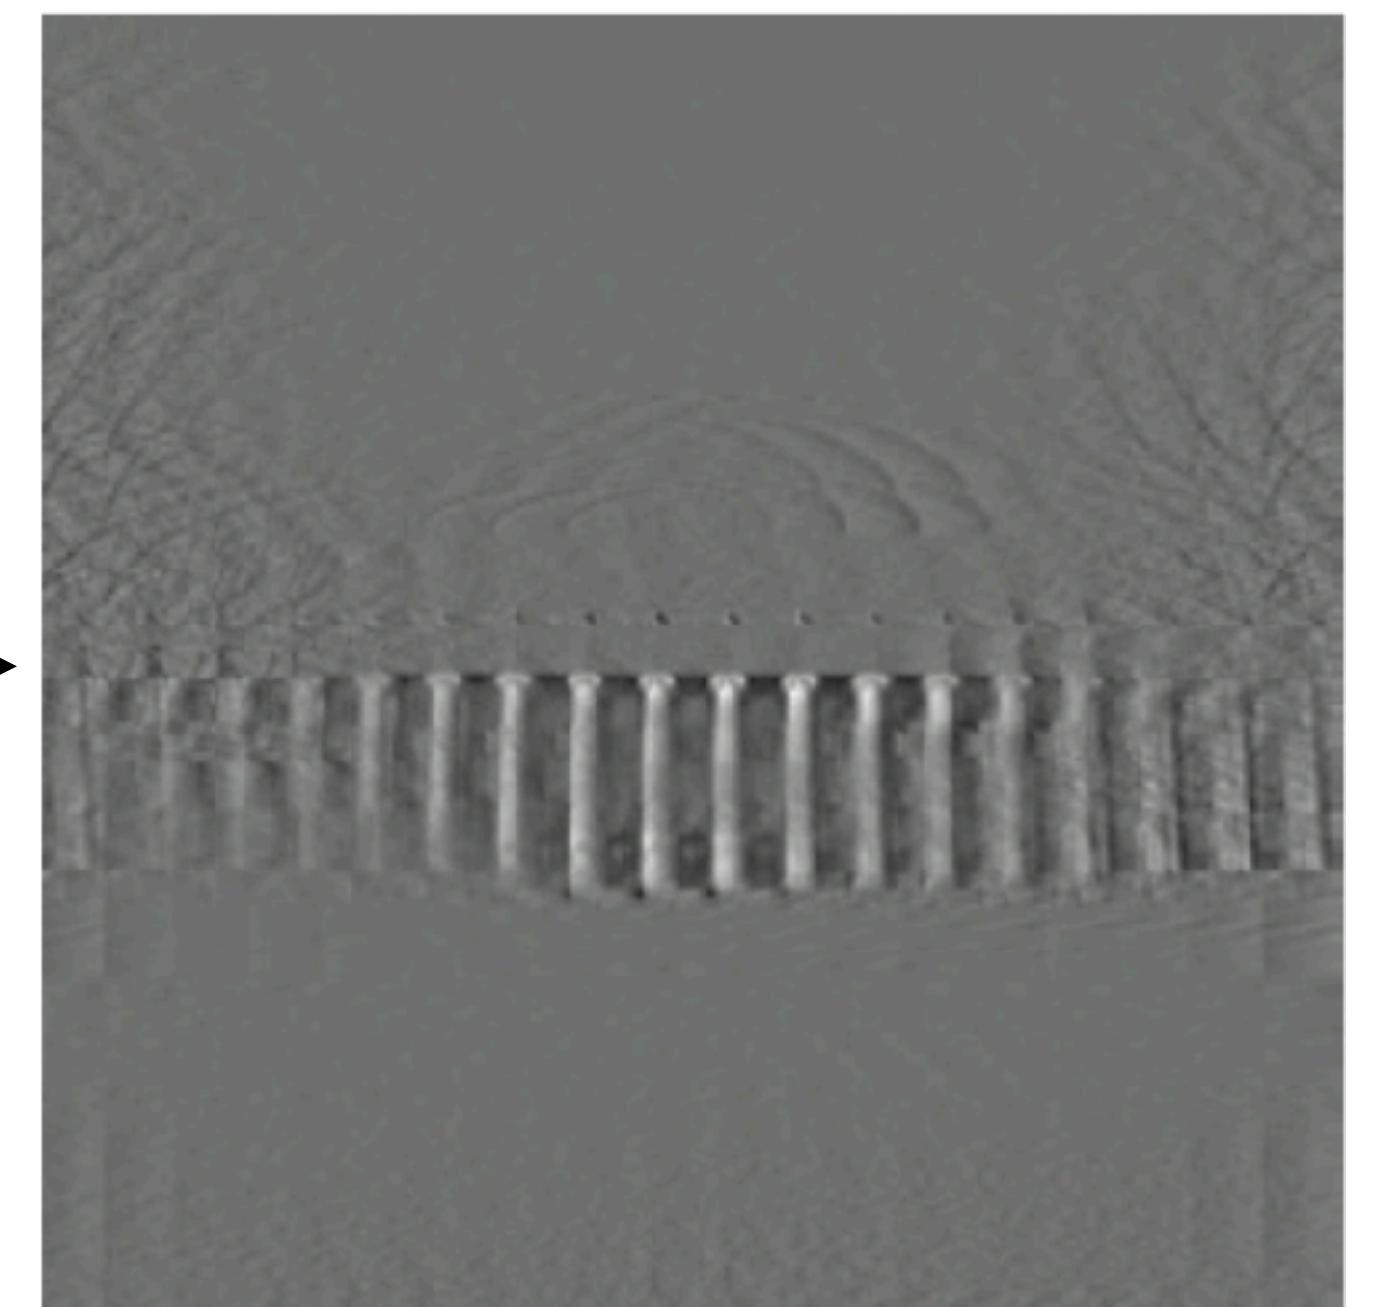
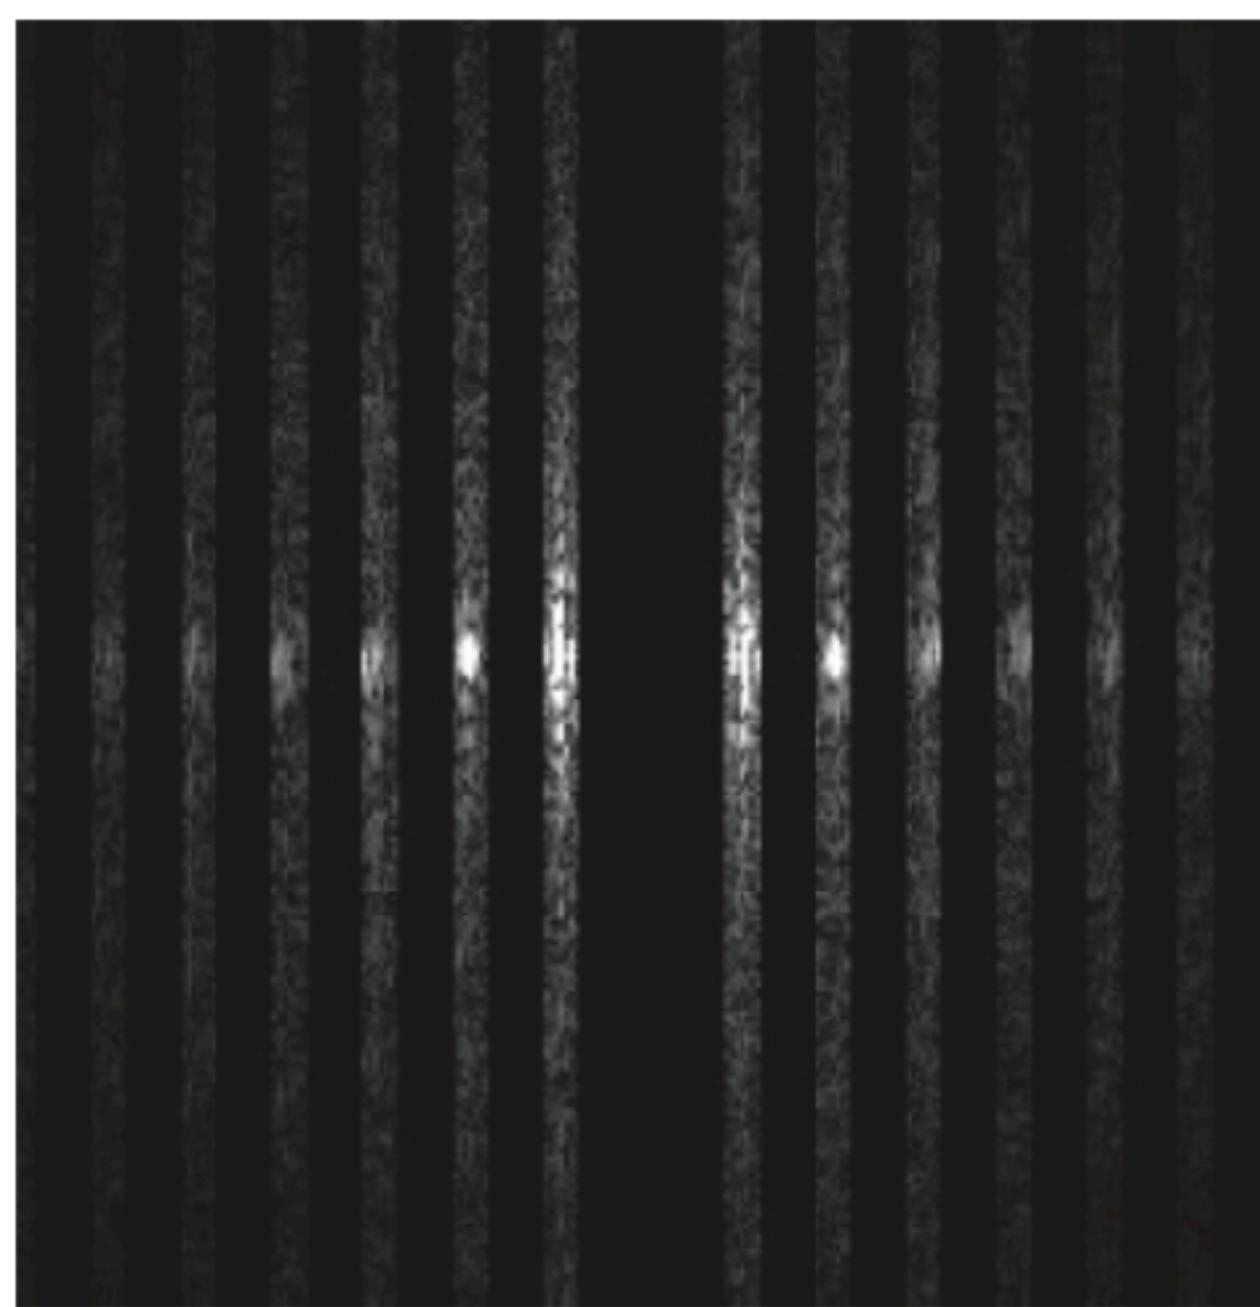




DFT
→

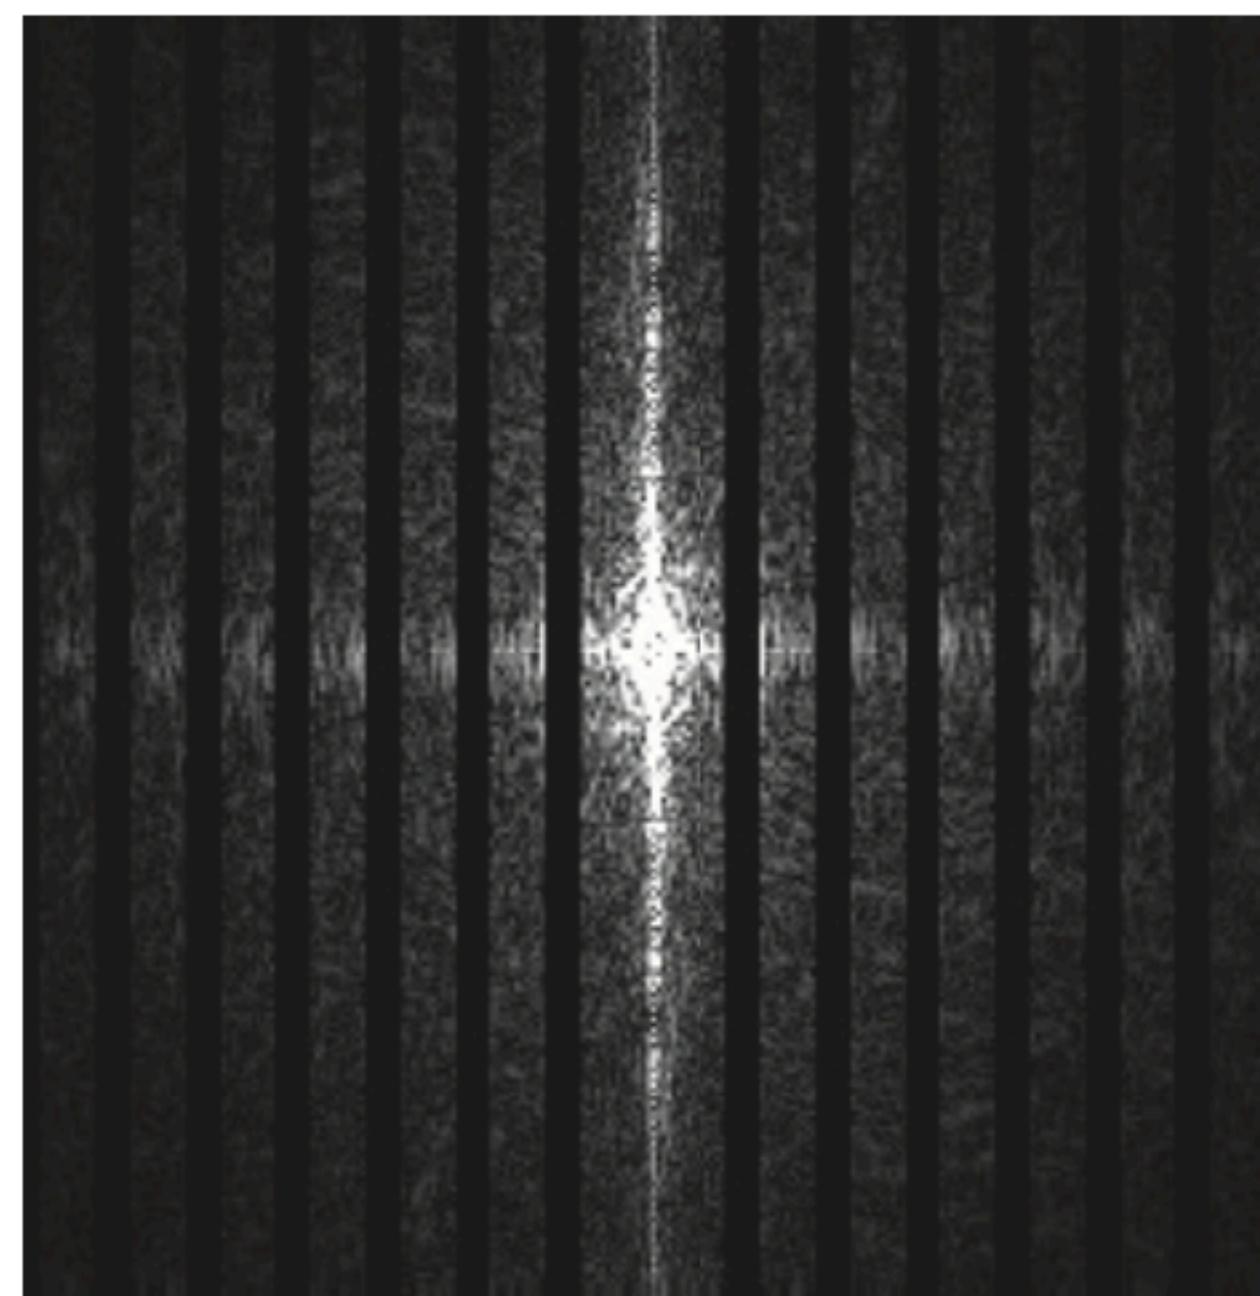
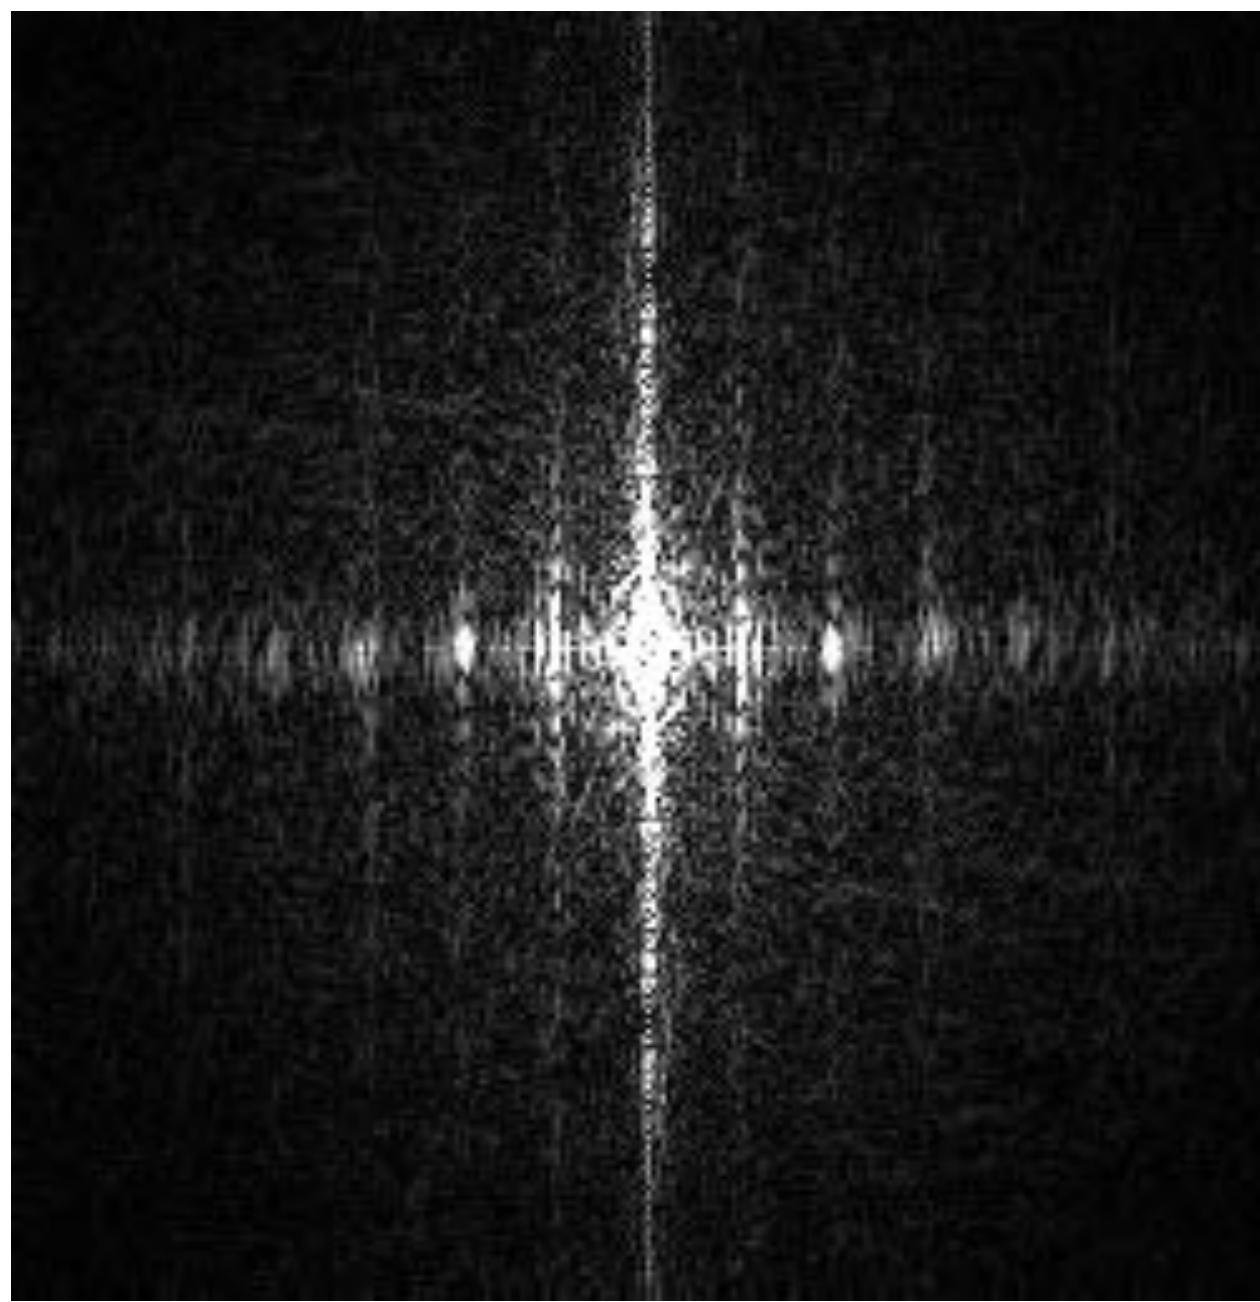


DFT⁻¹
→





DFT
→



DFT^{-1}
→

