6.869 Advances in Computer Vision Spring 2010

http://people.csail.mit.edu/torralba/courses/6.869/6.869. computervision.htm

Lecture 2
The structure of images

http://people.csail.mit.edu/torralba/courses/6.869/6.869. computervision.htm

6.869 Advances in Computer Vision

Spring 2010

Overview

Advanced topics in computer vision with a focus on the use of machine learning techniques and applications in graphics and human-computer interface. Topics include image representations, texture models, structure-from-motion algorithms, Bayesian techniques, object and scene recognition, tracking, shape modeling, and image databases. Applications may include face recognition, multimodal interaction, interactive systems, cinematic special effects, and photorealistic rendering. Covers topics complementary to 6.801/6.866; these subjects may be taken in sequence.

General information

Lecture: Mondays/Wednesdays 1:00-2:30pm
Room: 2-139 (where is this?)

Instructor: Antonio Torralba
E -mail: trrlb@mt.d (fill the missing vowels)
Office: D432

T.A.: Joseph Lim

Material:

- Textbook: new book by Rick Szeliski (not published yet, but a draft is available online)
- Textbook: Computer vision: a modern approach, by Forsyth and Ponce. Prentice Hall, 2002.
- The class will make use of MATLAB.

Grading:

- problem sets: 1/3
- 2 take-home exams: 1/3
- final project: 1/3

Announcements

First class, Wednesday Feb 3rd

Schedule

Lecture	Date	Topic	Slides	Readings	Assignments	Aditional material

The structure of ambient light

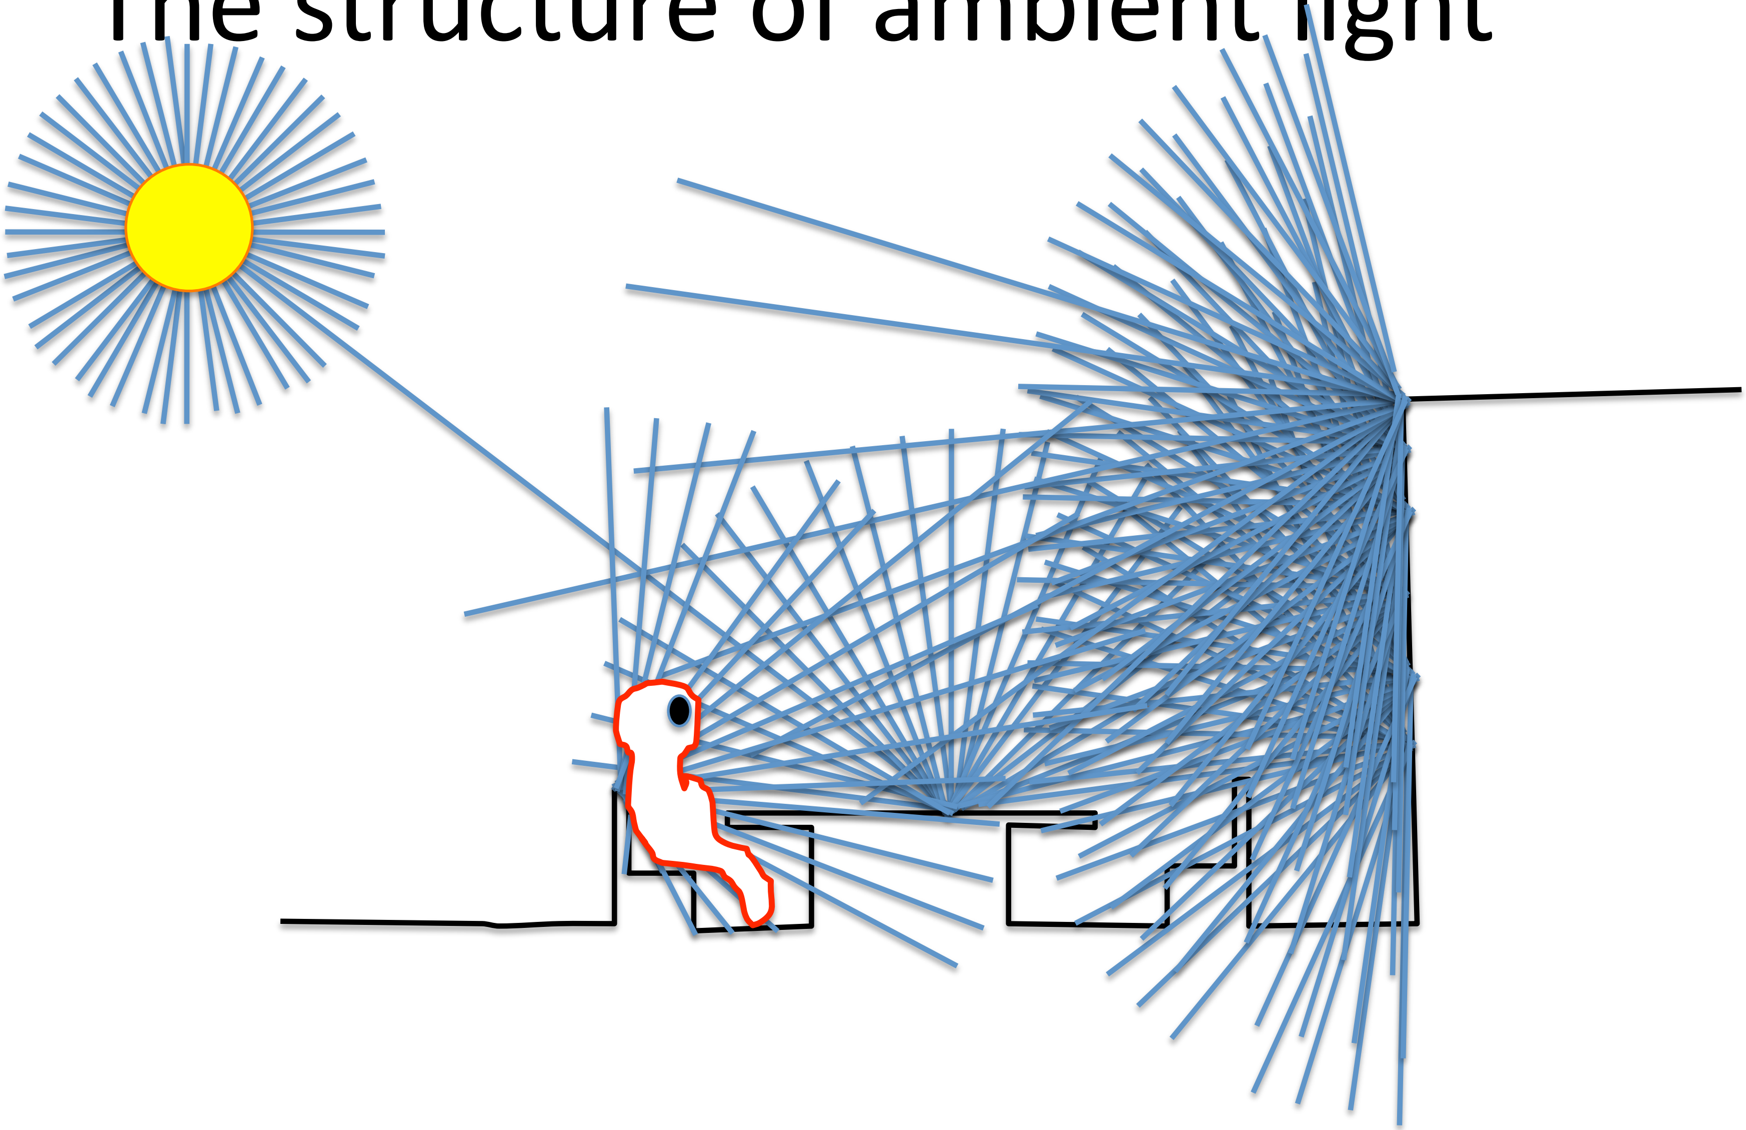

Edge like structures

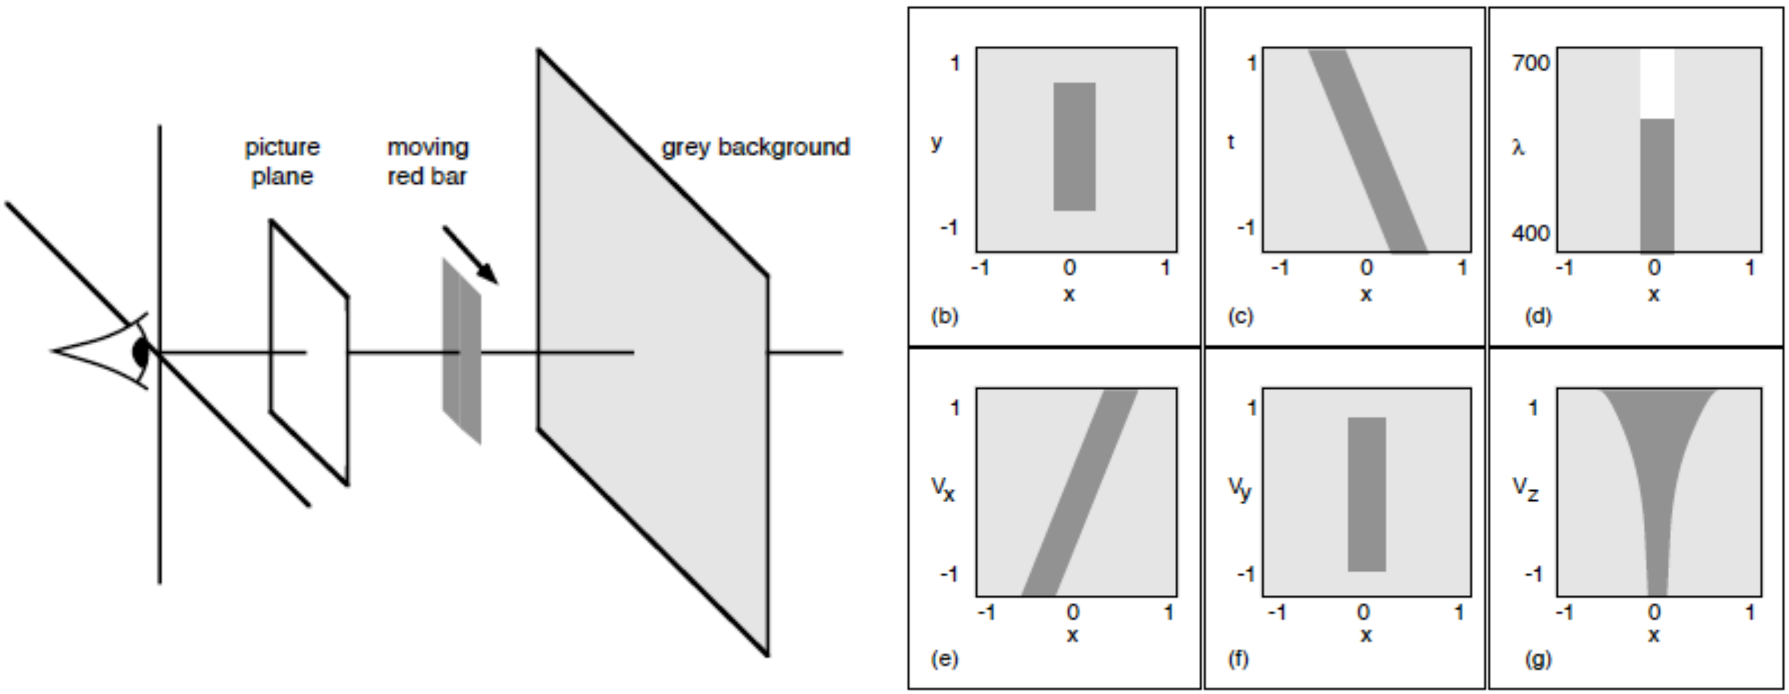

Default values:
$x = 0$
$y = 0$
$t = 0$
$\lambda = 550\text{nm}$
$V_x = 0$
$V_y = 0$
$V_z = 0$

- Proposition 1. The primary task of early vision is to deliver a small set of useful measurements about each observable location in the plenoptic function.

- Proposition 2. The elemental operations of early vision involve the measurement of local change along various directions within the plenoptic function.

What are we tuned to?

The visual system is tuned to process structures typically found in the world.

The visual system seems to be tuned to a set of images:

Demo inspired from D. Field

What is a natural image?

What is a natural image?

What is a natural image?

The visual system seems to be tuned to a set of images:

What is a natural image?

The visual system seems to be tuned to a set of images:

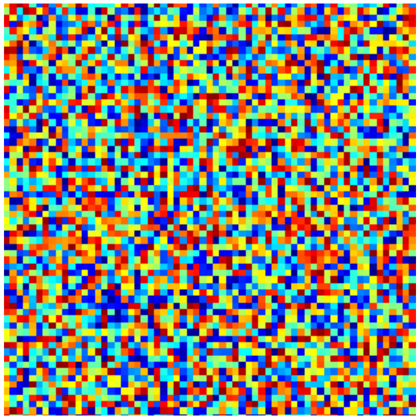

Did you saw this image?

6 images



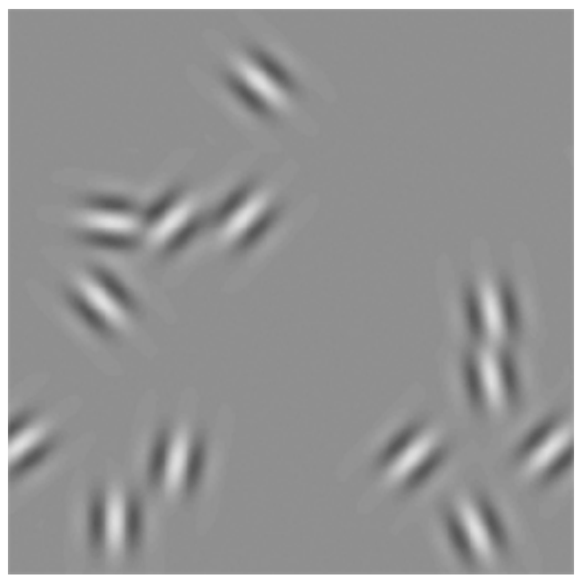


Not all these images are the result of sampling a real-world plenoptic function

• Proposition 1. The primary task of early vision is to deliver a small set of useful measurements about each observable location in the plenoptic function.

• Proposition 2. The elemental operations of early vision involve the measurement of local change along various directions within the plenoptic function.

• Goal: to transform the image into other representations (rather than pixel values) that makes scene information more explicit

What we think we see

What we really see

Cavanagh, Perception 95

What are “visual features”?

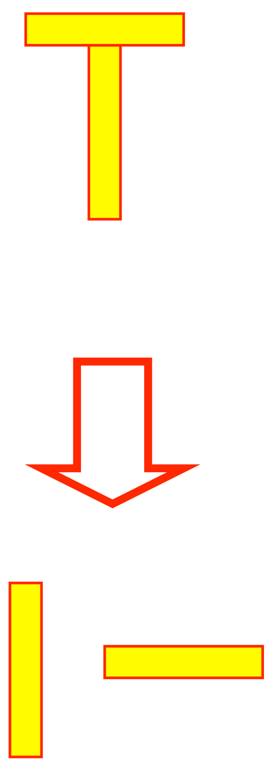

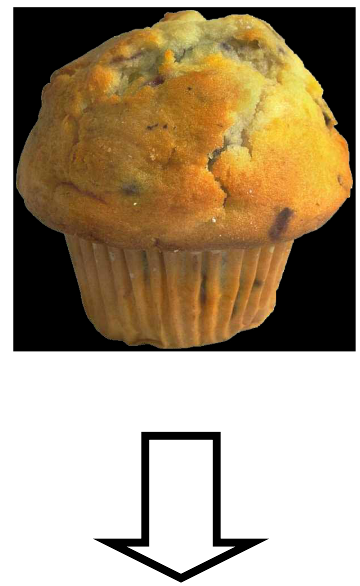

Shape, color, texture, etc

The image as a “surface”

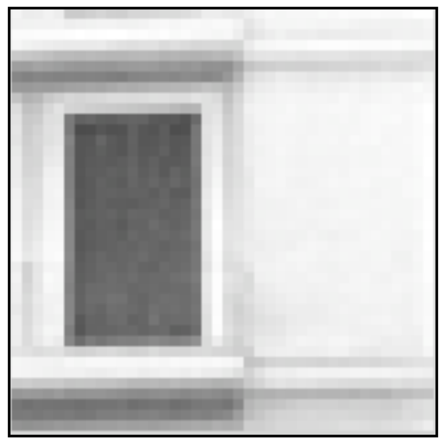

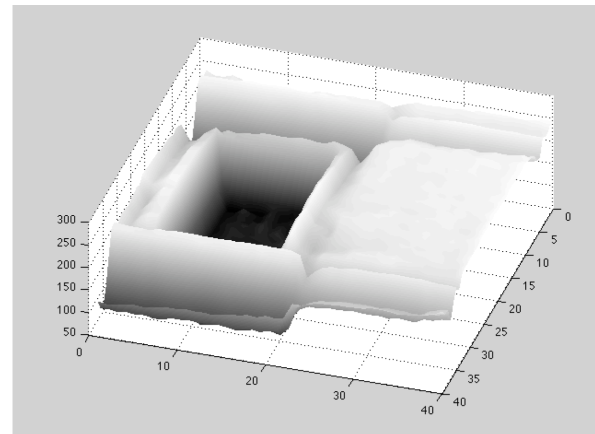

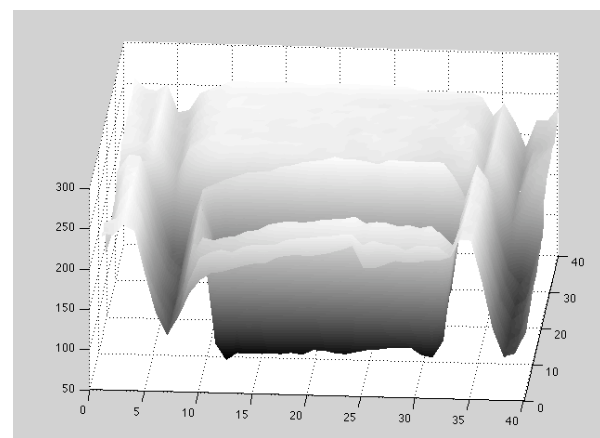

Filtering

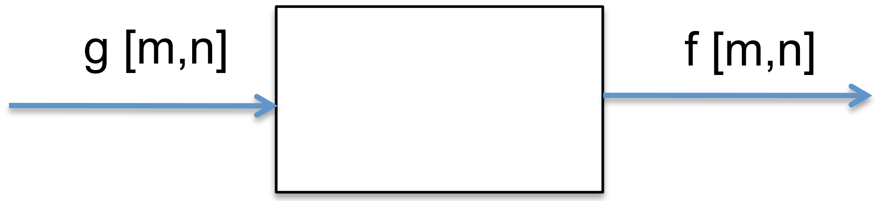

We want to remove unwanted sources of variation, and keep the information relevant for whatever task we need to solve

Linear filtering

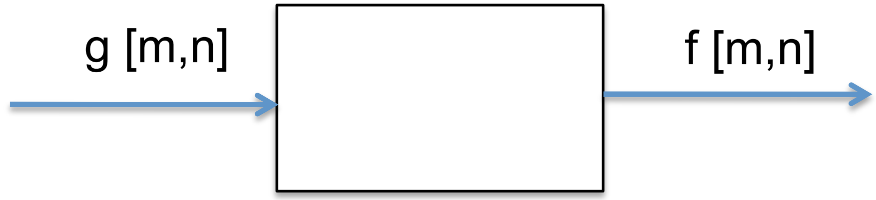

For a linear system, each output is a linear combination of all the input values:

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

In matrix form:

F = H G

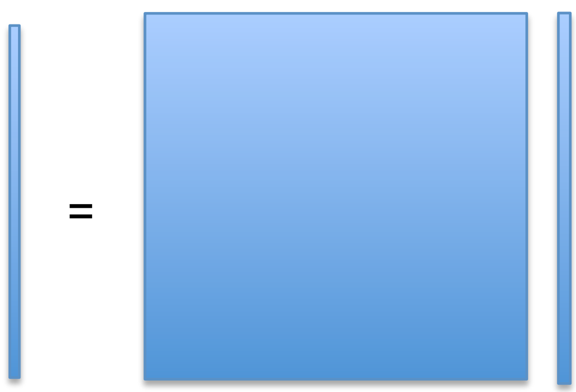

Linear filtering

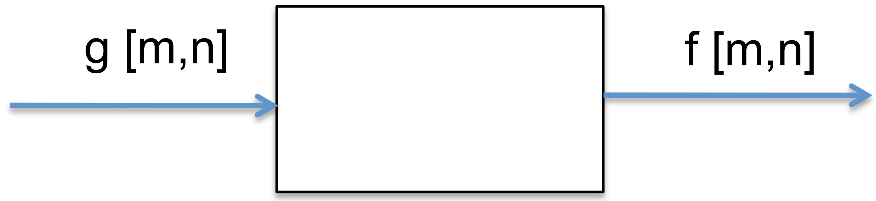

In vision, many times, we are interested in operations that are spatially invariant.
For a linear spatially invariant system:

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

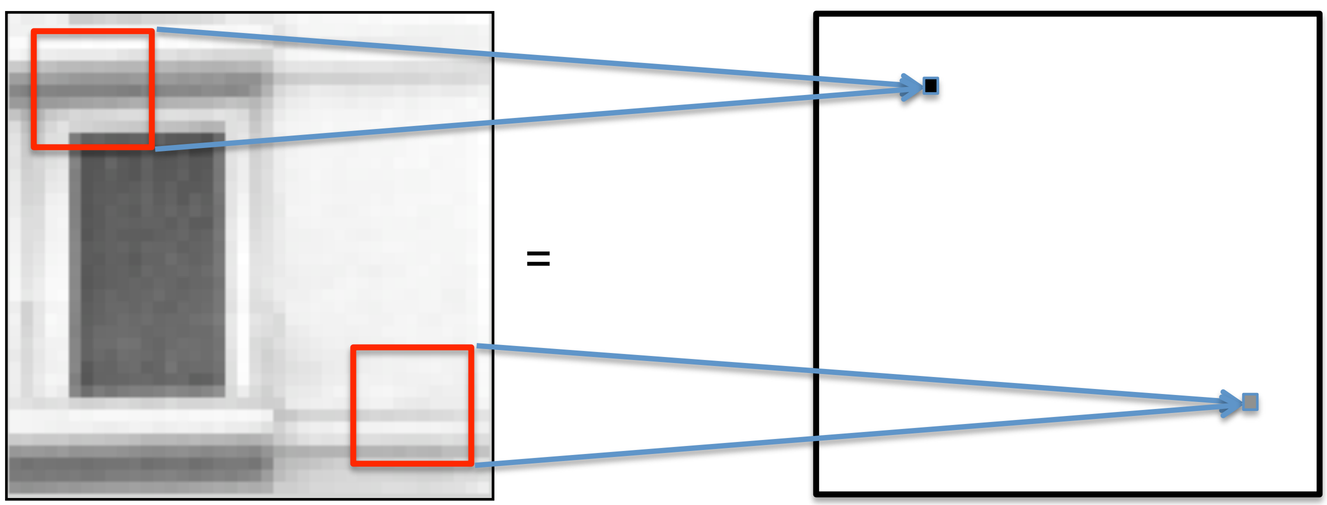

Linear filtering

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

Linear system:

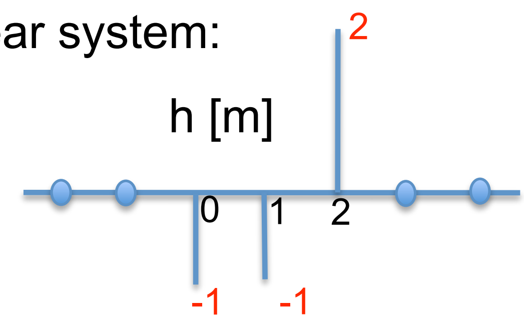

Input:

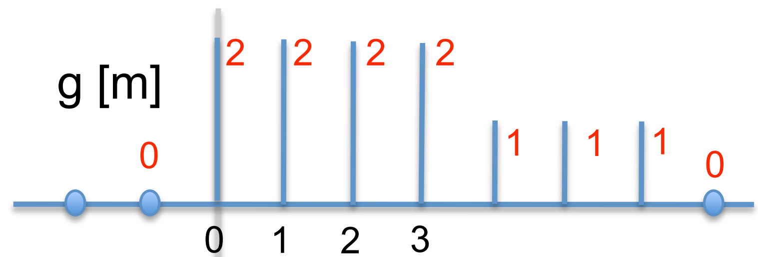

Output?

$$f[m=0] = \sum_{k} h[-k]\,g[k]$$

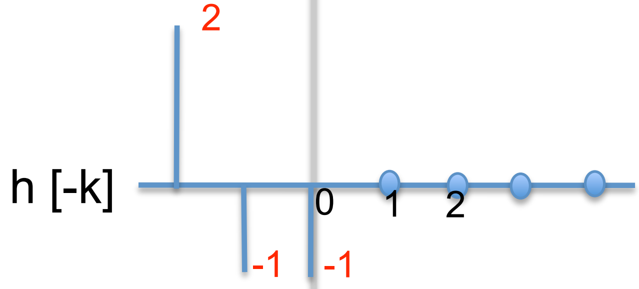

f [m=0]=-2

$$f[m=1] = \sum_{k} h[1-k]\,g[k]$$

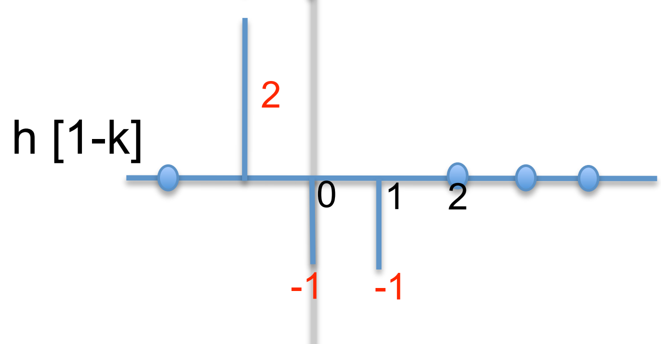

f [m=1]=-4

$$f[m=2] = \sum_{k} h[2-k]\,g[k]$$

f [m=2]=0

Linear filtering

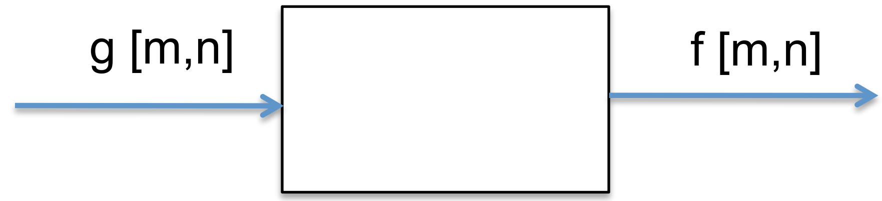

For a linear spatially invariant system

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

m=0 1 2 …

111	115	113	111	112	111	112	111
135	138	137	139	145	146	149	147
163	168	188	196	206	202	206	207
180	184	206	219	202	200	195	193
189	193	214	216	104	79	83	77
191	201	217	220	103	59	60	68
195	205	216	222	113	68	69	83
199	203	223	228	108	68	71	77

g[m,n]

$\otimes$

-1	2	-1
-1	2	-1
-1	2	-1

h[m,n]

=

?	?	?	?	?	?	?	?
?	-5	9	-9	21	-12	10	?
?	-29	18	24	4	-7	5	?
?	-50	40	142	-88	-34	10	?
?	-41	41	264	-175	-71	0	?
?	-24	37	349	-224	-120	-10	?
?	-23	33	360	-217	-134	-23	?
?	?	?	?	?	?	?	?

f[m,n]

Borders

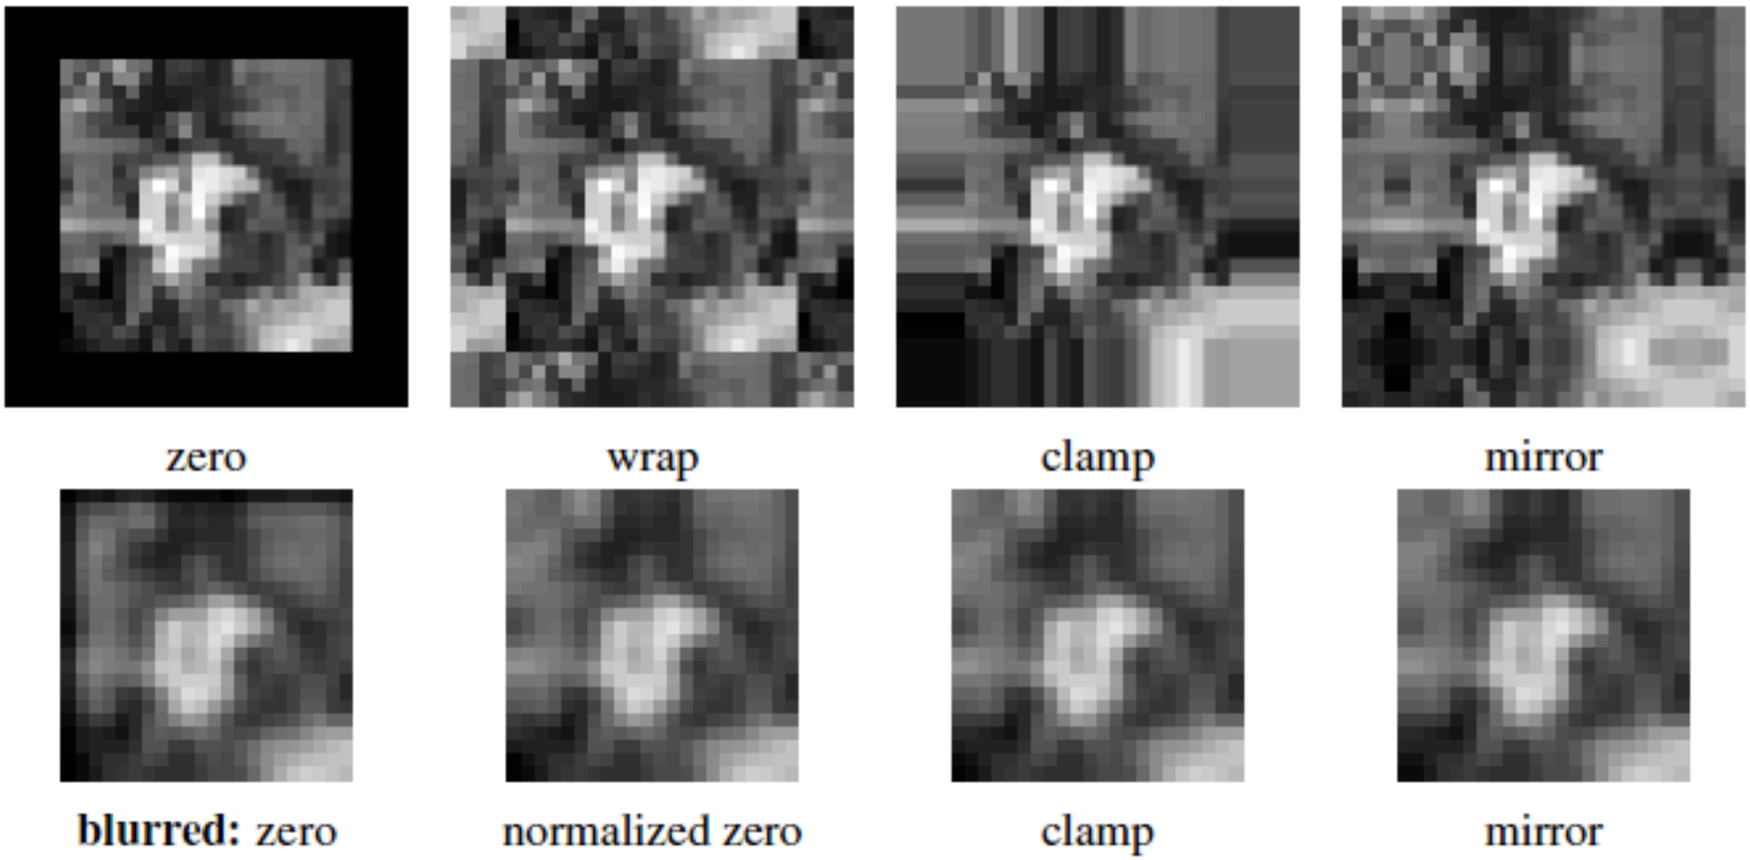

From Rick's book

A taxonomy of useful filters

- Impulse, Shifts,
- Blur
  - Rectangular blur (see artifacts)
  - Gaussian
  - Bilateral exponetial
  - Asymmetrical filter: motion blur
- Edges
  - [-1 1]
  - Derivative filter
  - Derivative of a gaussian
  - Oriented filters
  - Gabor filter
  - Quadrature filters: phase and magnitude. Iris code.
  - Elongated edges: filling gaps…

Impulse

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

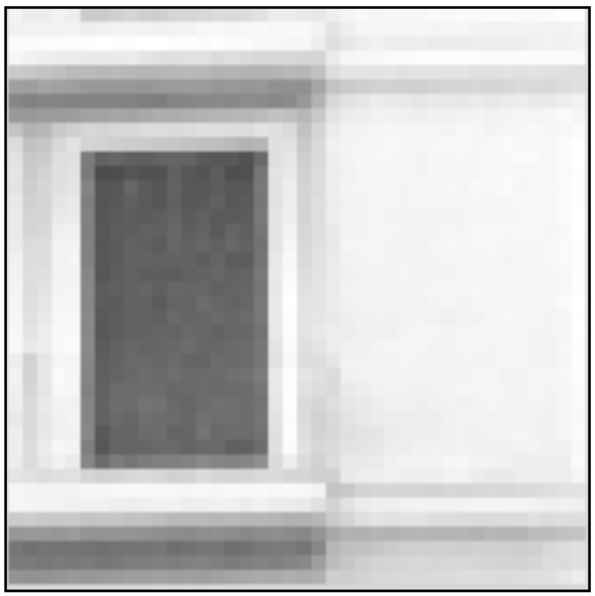

g[m,n]

$\otimes$

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

h[m,n]

=

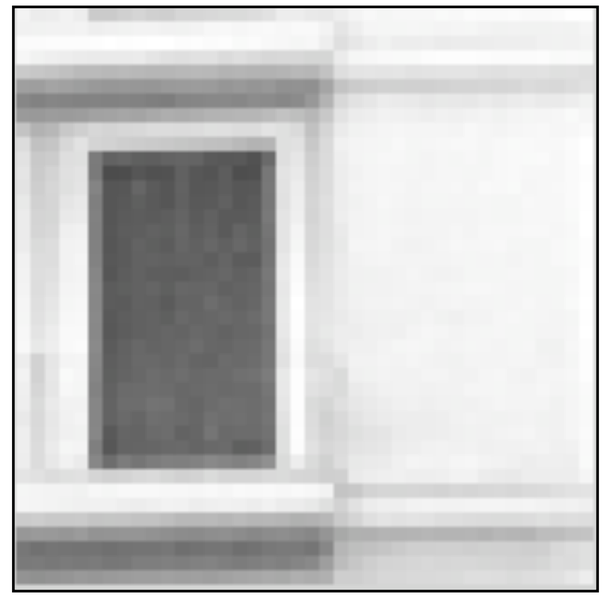

f[m,n]

Shifts

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

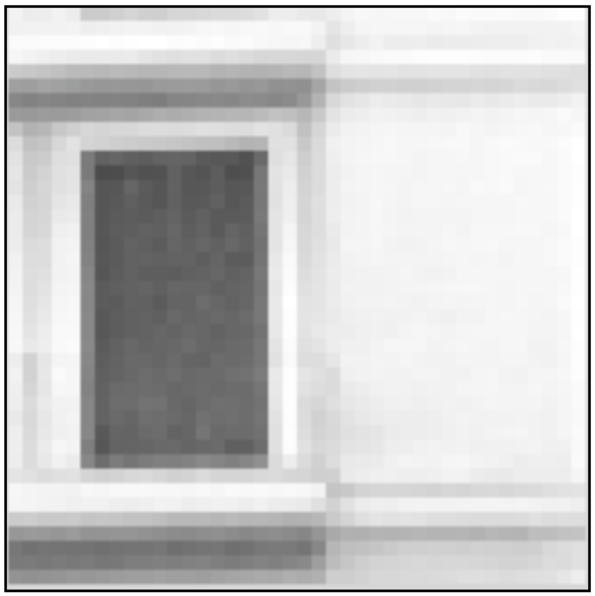

g[m,n]

$\otimes$

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

h[m,n]

=

2pixels

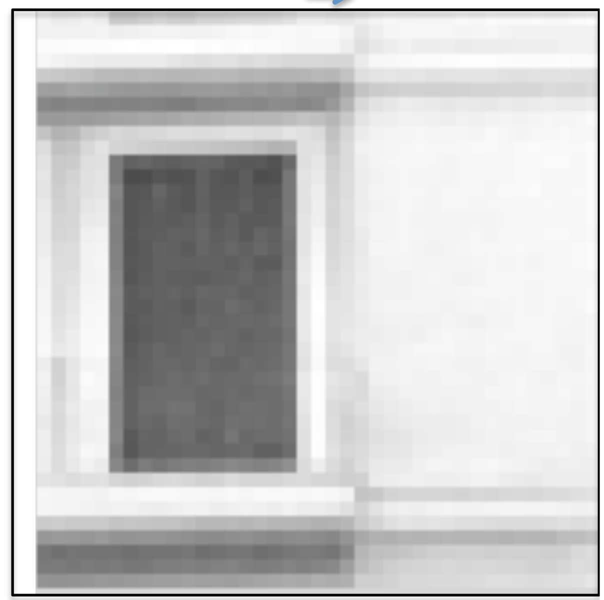

f[m,n]

Image rotation

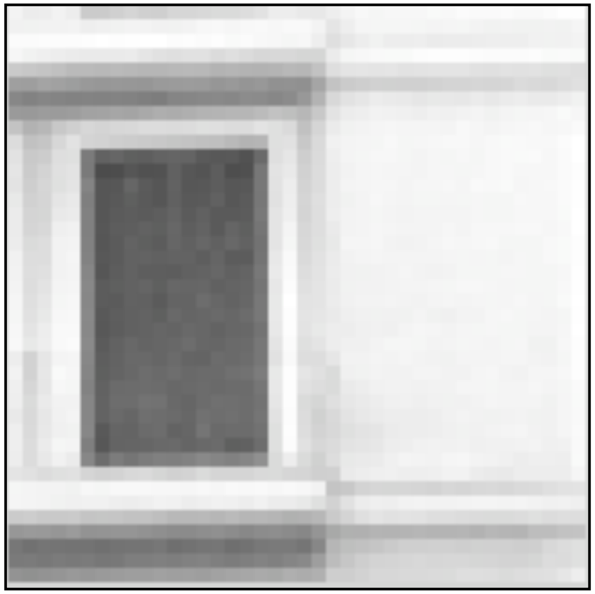

g[m,n]

$\otimes$? =

h[m,n]

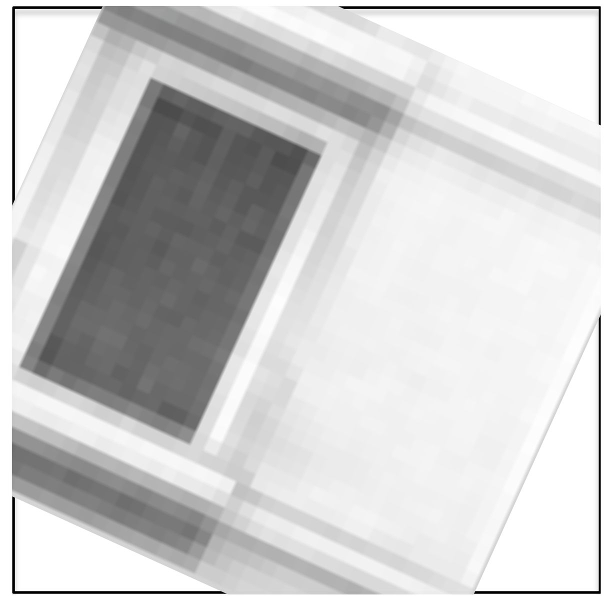

f[m,n]

It is linear, but not a spatially invariant operation. There is not convolution.

Rectangular filter

$\otimes$

h[m,n]

=

f[m,n]

g[m,n]

Rectangular filter

g[m,n]

$\otimes$

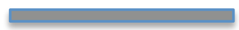

h[m,n]

=

f[m,n]

Rectangular filter

g[m,n]

$\otimes$

h[m,n]

=

f[m,n]

Integral image

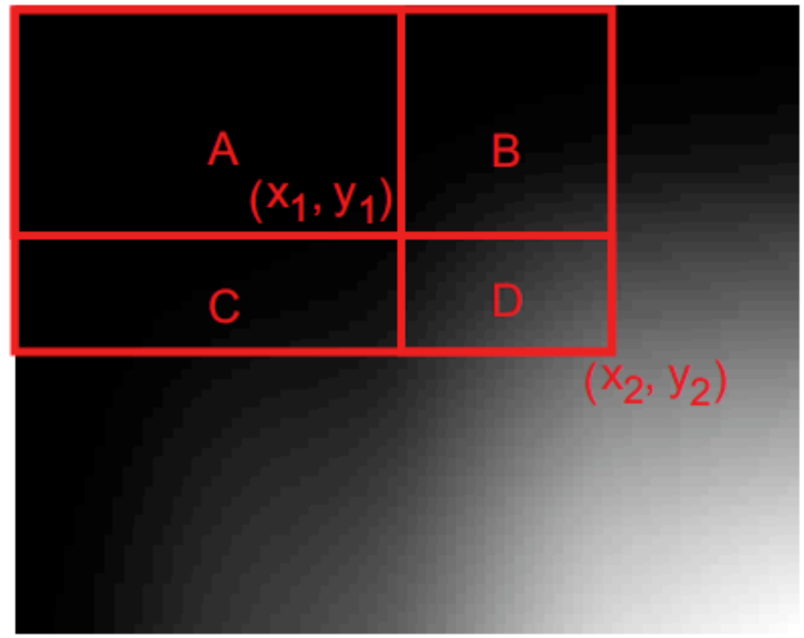

Sharpening

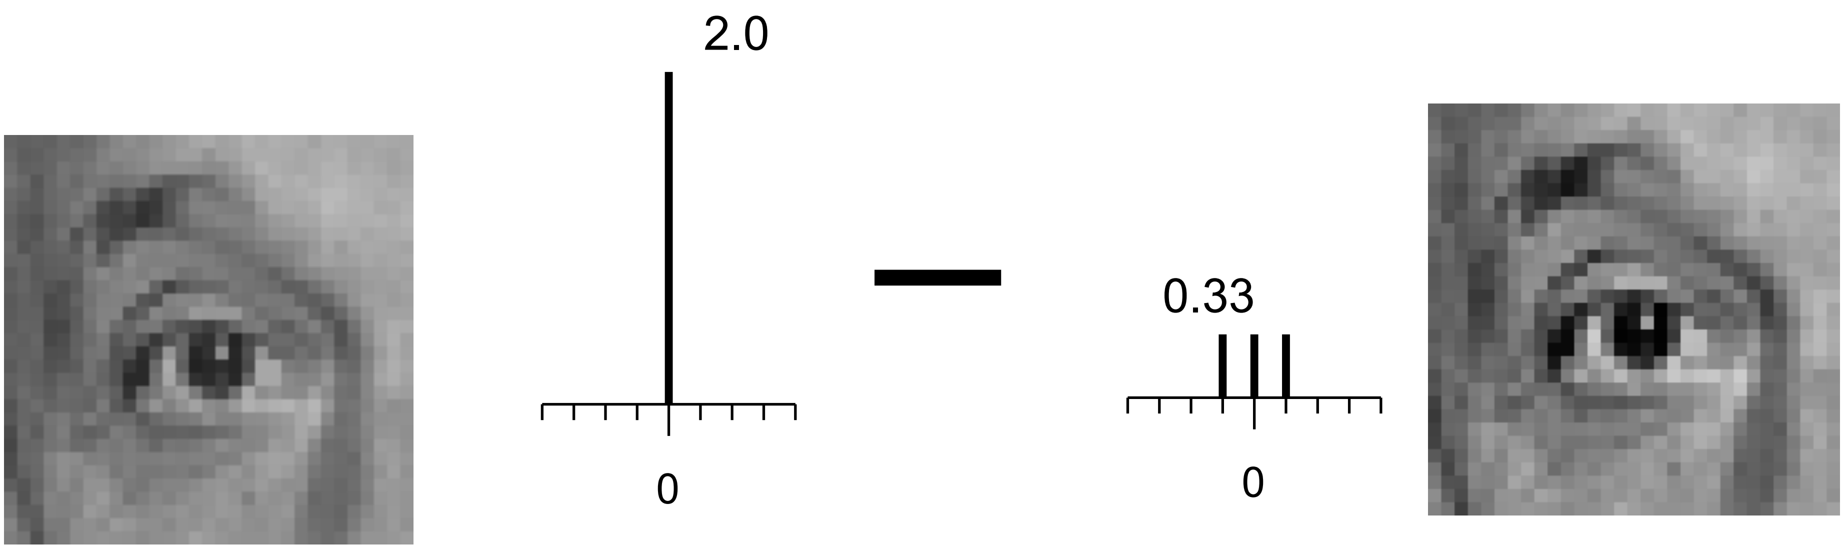

original

Sharpened
original

Sharpening example

filter

result

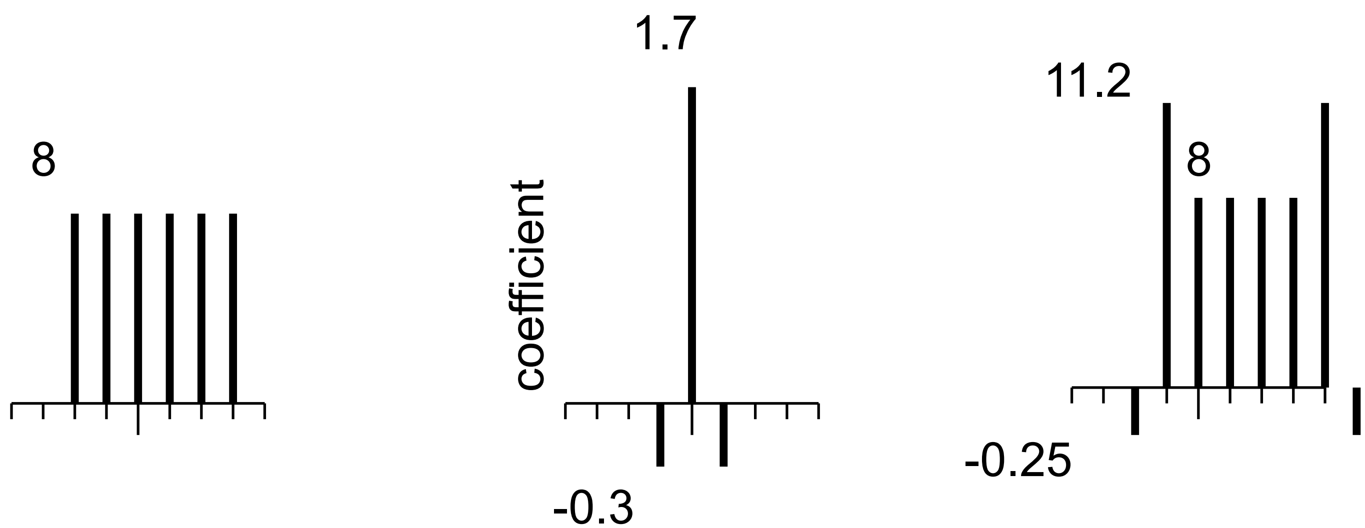

original

Sharpened
(differences are
accentuated; constant
areas are left untouched).

Sharpening

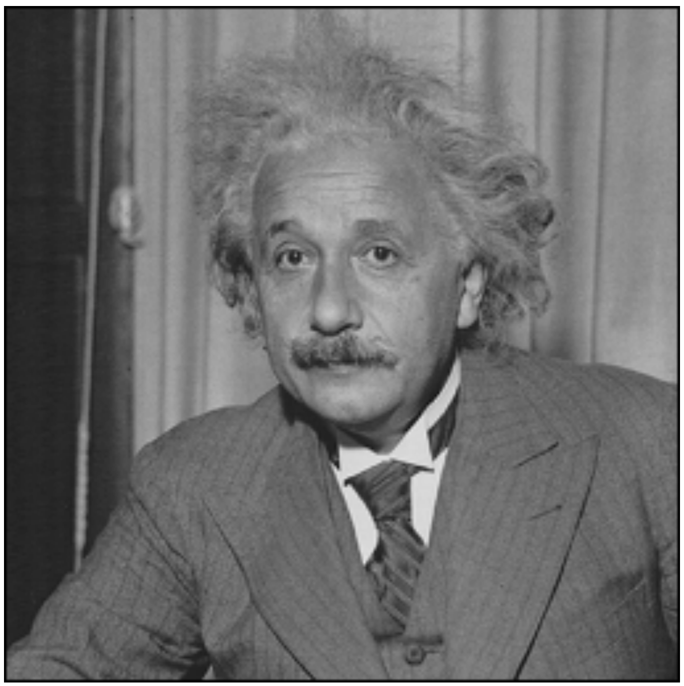

before

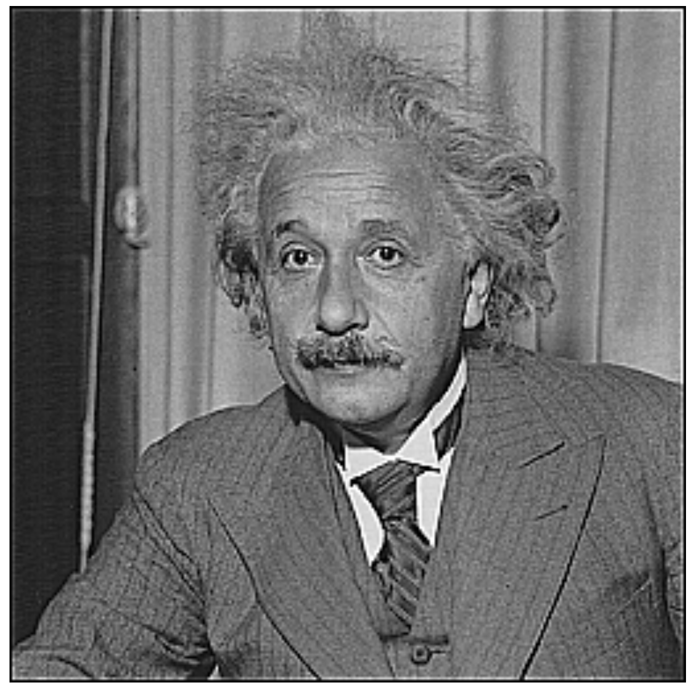

after

Gaussian filter

$$G(x, y; \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

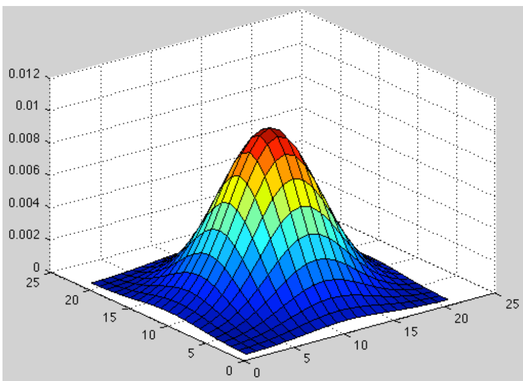

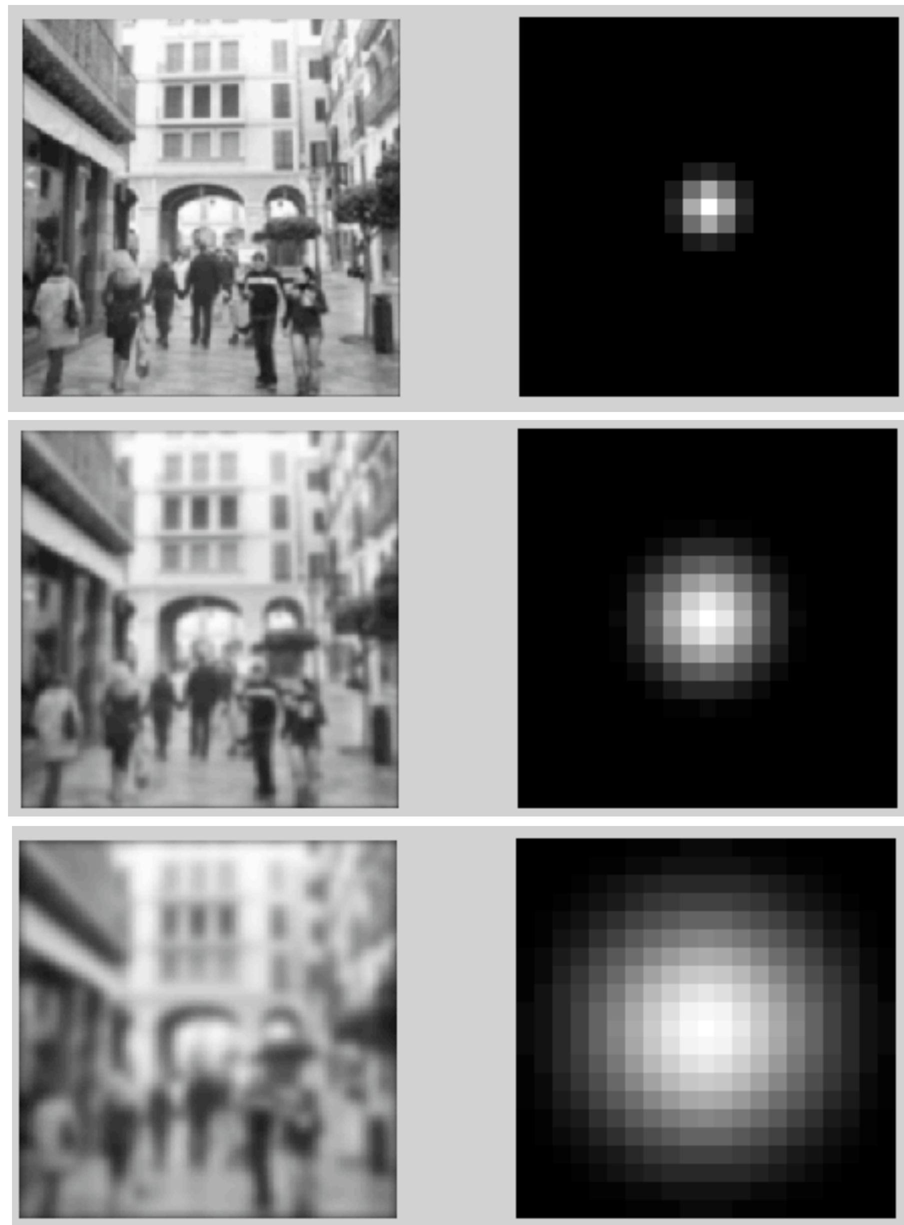

σ=1

σ=2

σ=4

Some desirable properties for a blur kernel

- Positivity: $h(m) >= 0$
- Symmetry: $h(m) = h(-m)$
- Unimodality: $h(m) >= h(m+1)$ for $m >= 0$
- Normalized: $\Sigma\, h(m) = 1$
- Equal contribution: $\Sigma\, h(2m) = \Sigma\, h(2m+1)$

Some kernels that verify this are:

[½ ½]

[¼ ½ ¼]

Global to Local Analysis

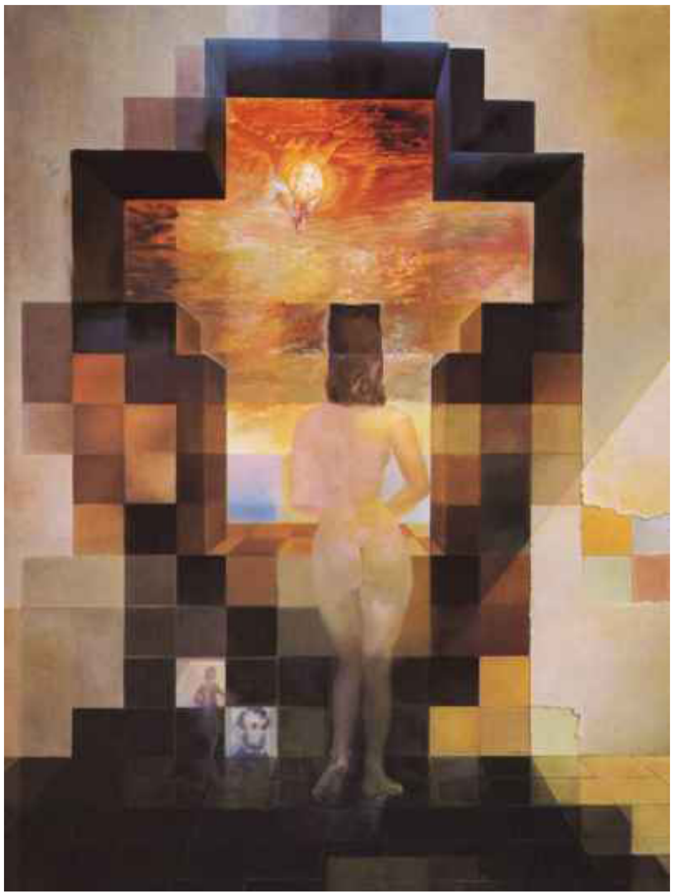

Dali

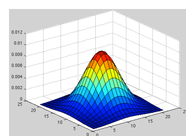

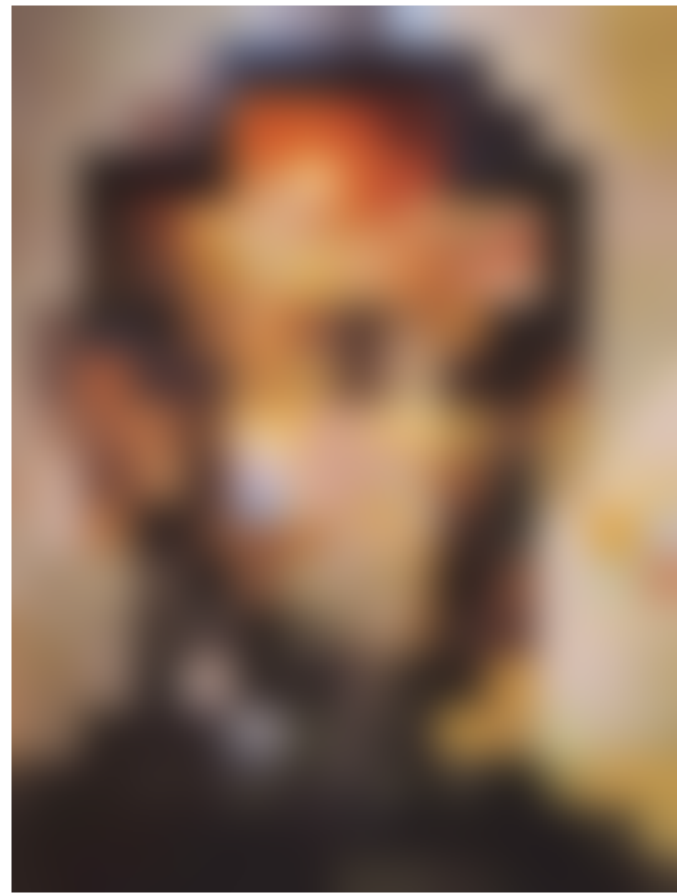

Linear image transformations

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

transformed image

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

Vectorized image

Fourier transform, or
Wavelet transform, or
Steerable pyramid transform

Self-inverting transforms

Same basis functions are used for the inverse transform

$$\vec{f} = U^{-1}\vec{F}$$

$$= U^{+}\vec{F}$$

U transpose and complex conjugate

An example of such a transform: the Discrete Fourier transform

Forward transform

$$F[m,n] = \sum_{k=0}^{M-1}\sum_{l=0}^{N-1} f[k,l]e^{-\pi i\left(\frac{km}{M}+\frac{\ln}{N}\right)}$$

Inverse transform

$$f[k,l] = \frac{1}{MN}\sum_{k=0}^{M-1}\sum_{l=0}^{N-1} F[m,n]e^{+\pi i\left(\frac{km}{M}+\frac{\ln}{N}\right)}$$

To get some sense of what basis elements look like, we plot a basis element --- or rather, its real part --- as a function of x,y for some fixed u, v. We get a function that is constant when (ux+vy) is constant. The magnitude of the vector (u, v) gives a frequency, and its direction gives an orientation. The function is a sinusoid with this frequency along the direction, and constant perpendicular to the direction.

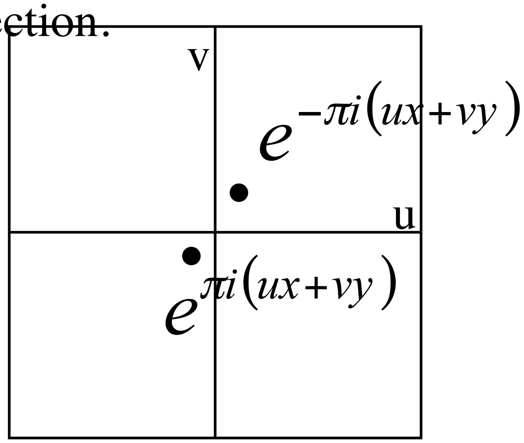

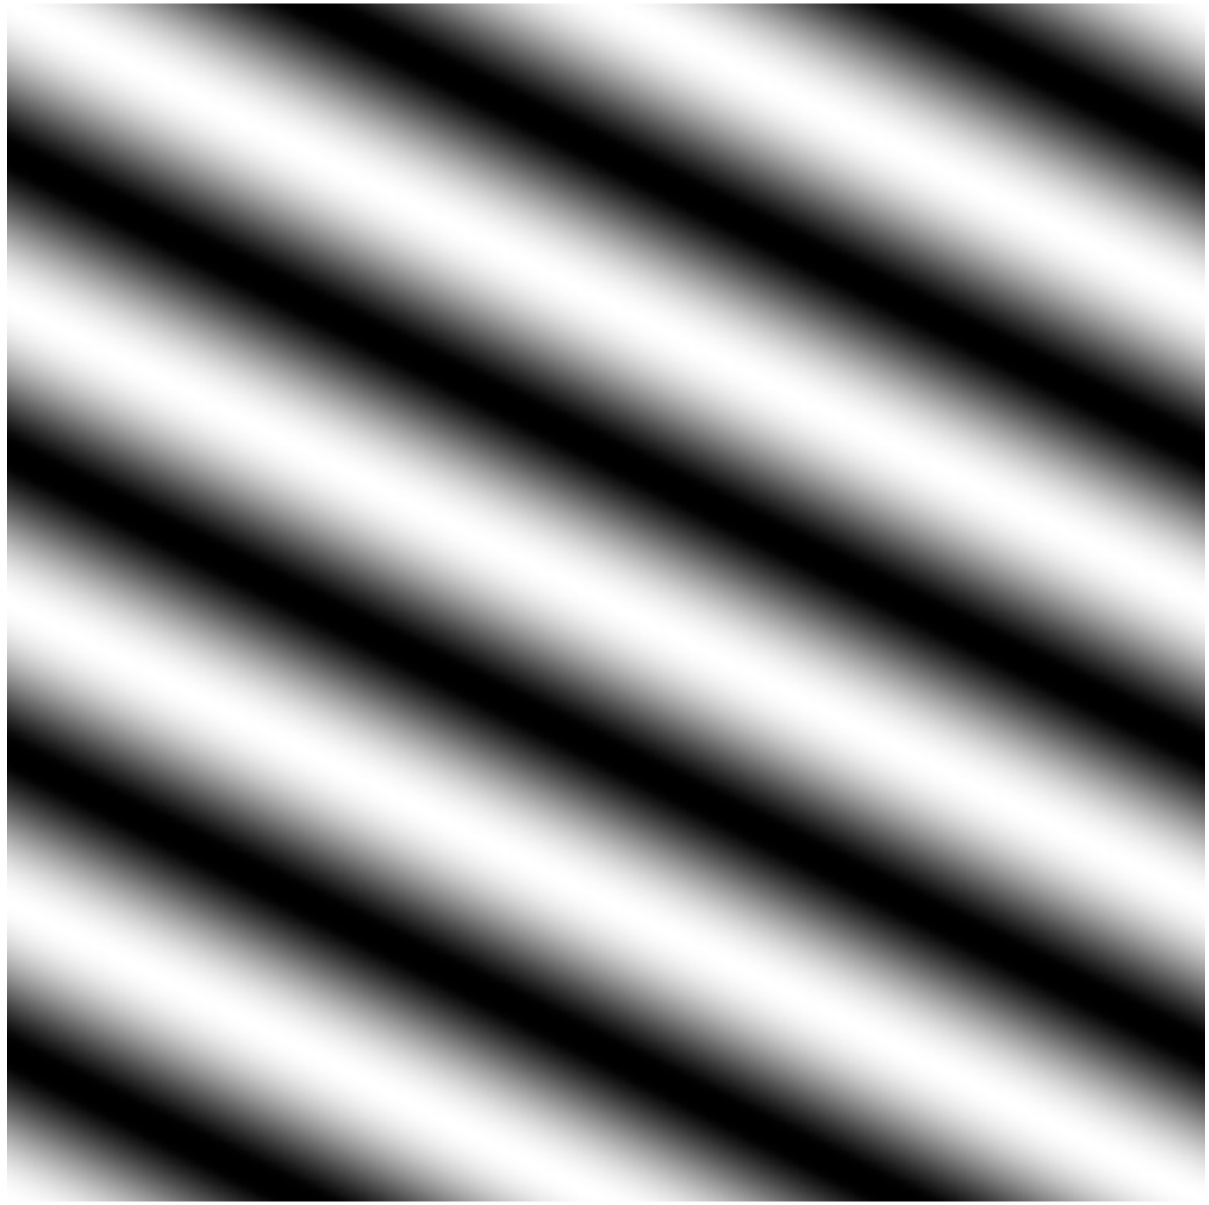

Here u and v are larger than in the previous slide.

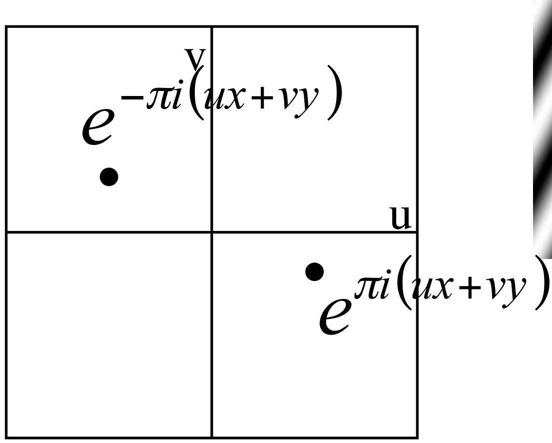

And larger still...

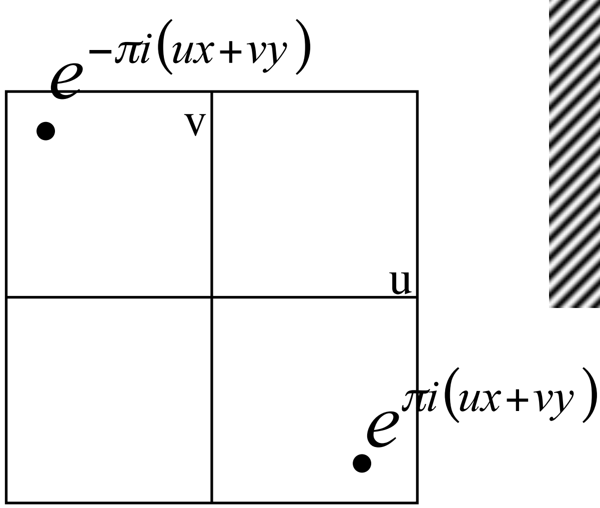

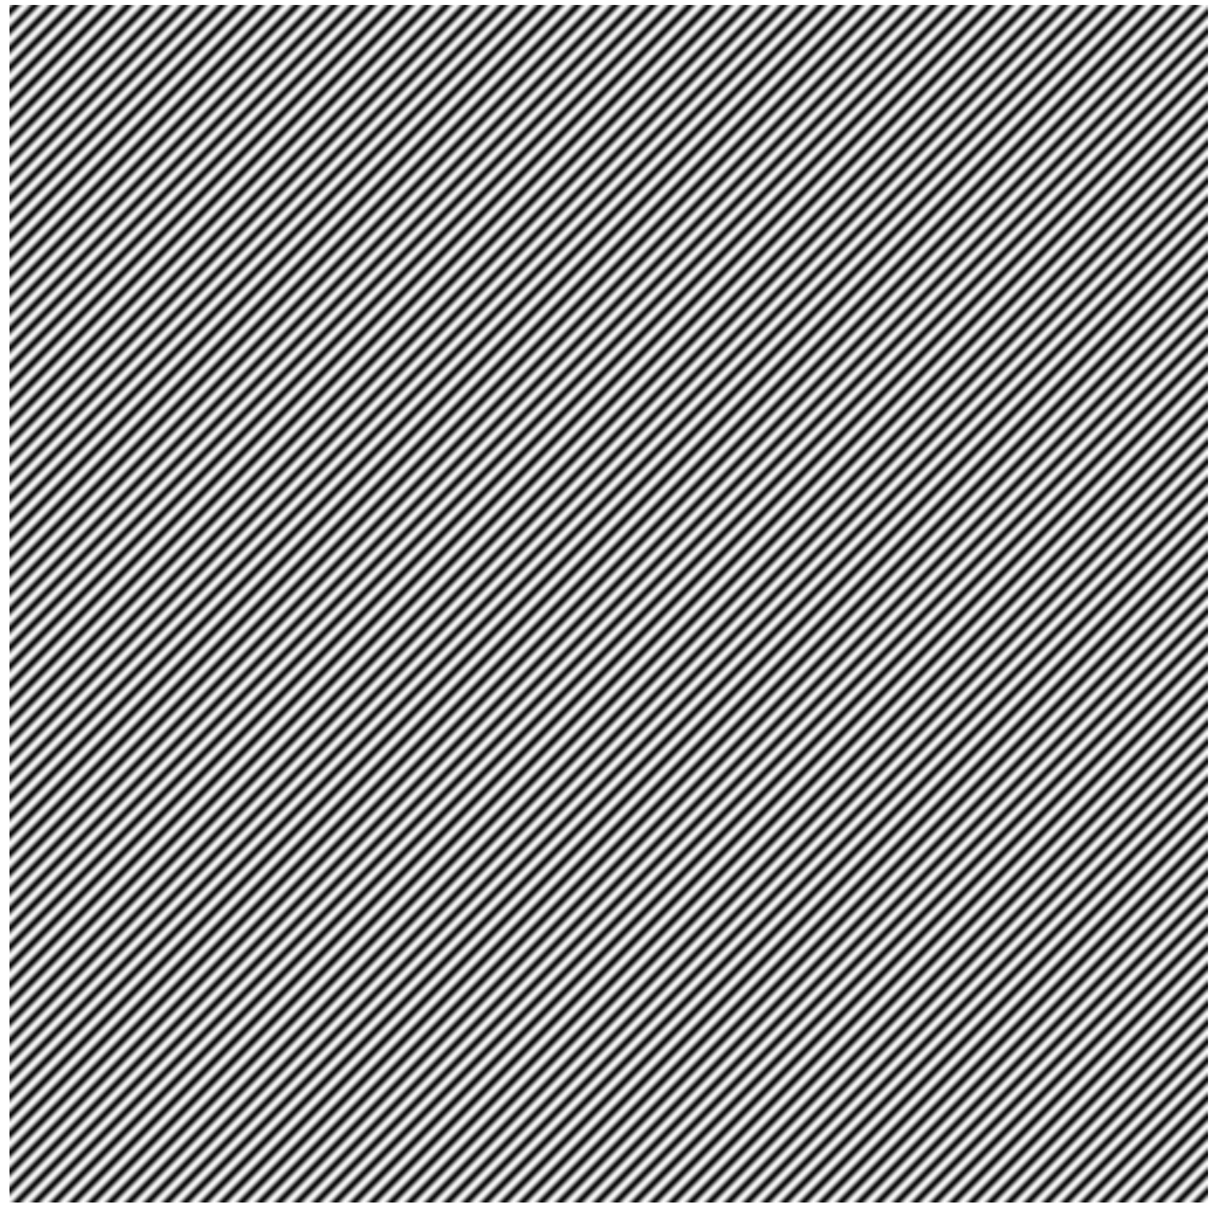

How to interpret a Fourier Spectrum

Vertical orientation

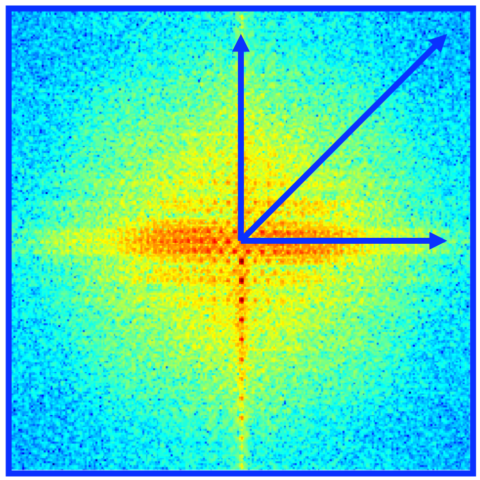

45 deg.

0

Horizontal orientation

0 f_{max}

fx in cycles/image

Low spatial frequencies

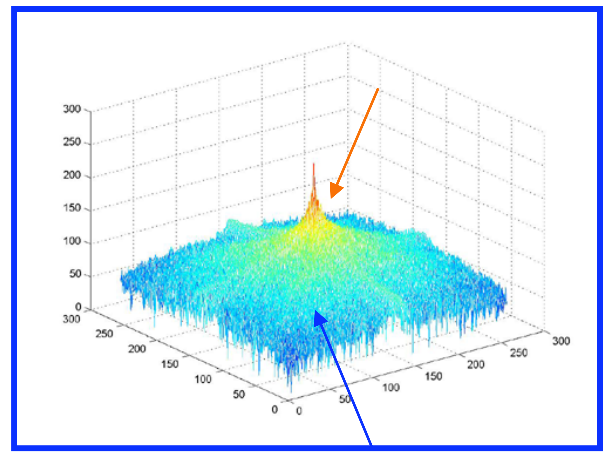

High spatial frequencies

Log power spectrum

Fourier Amplitude Spectrum

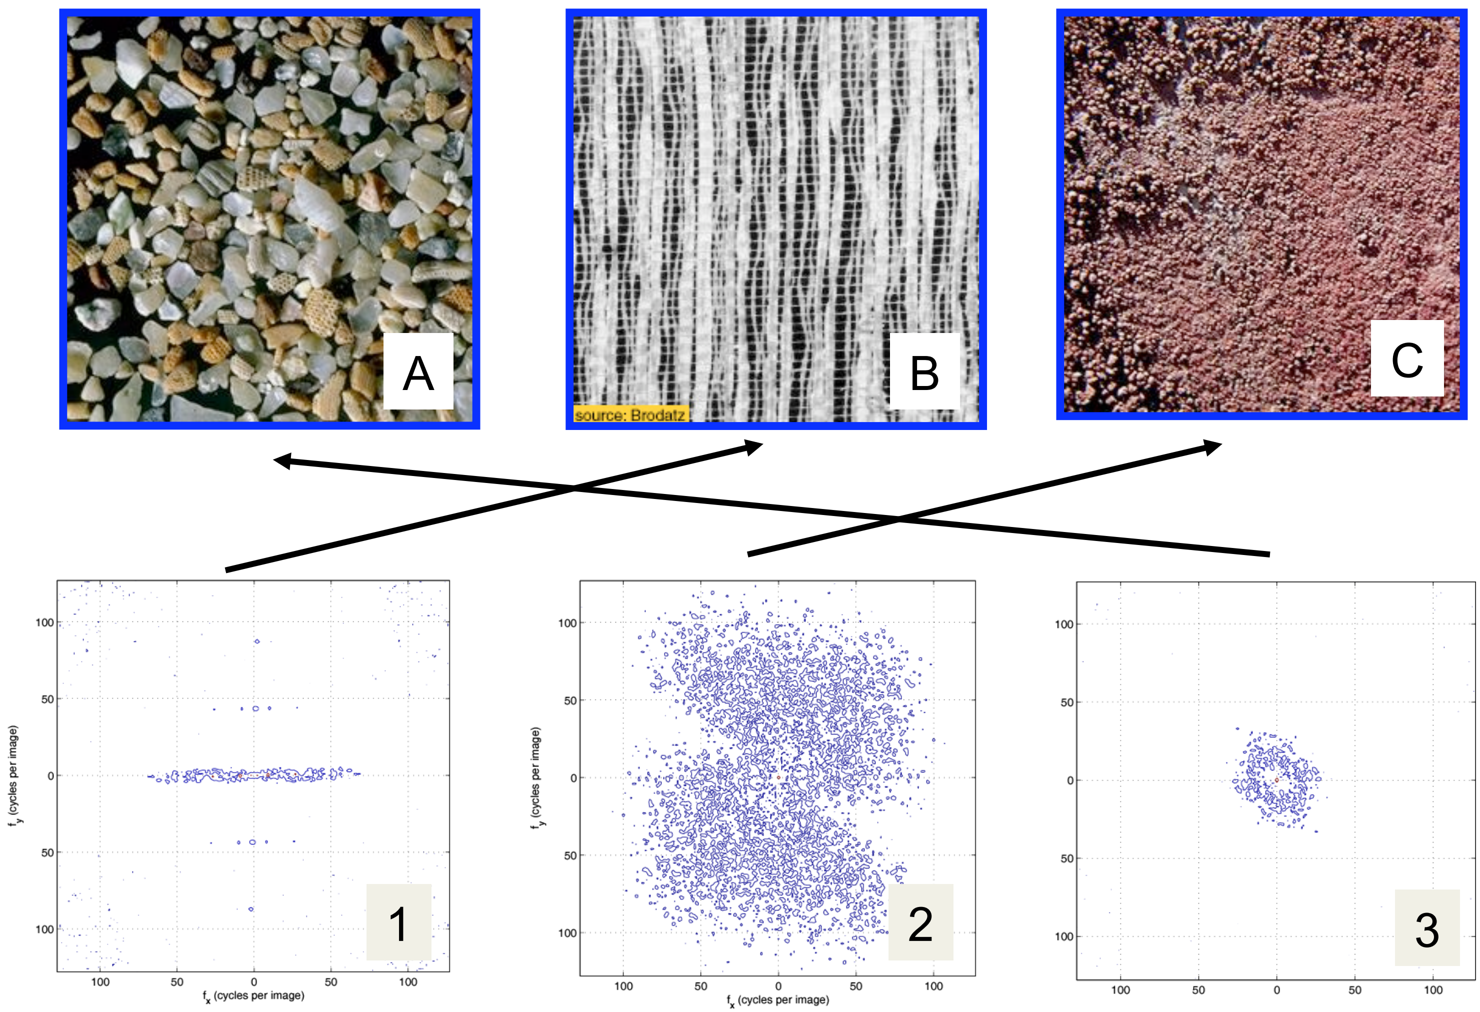

fx(cycles/image pixel size) fx(cycles/image pixel size) fx(cycles/image pixel size)

2

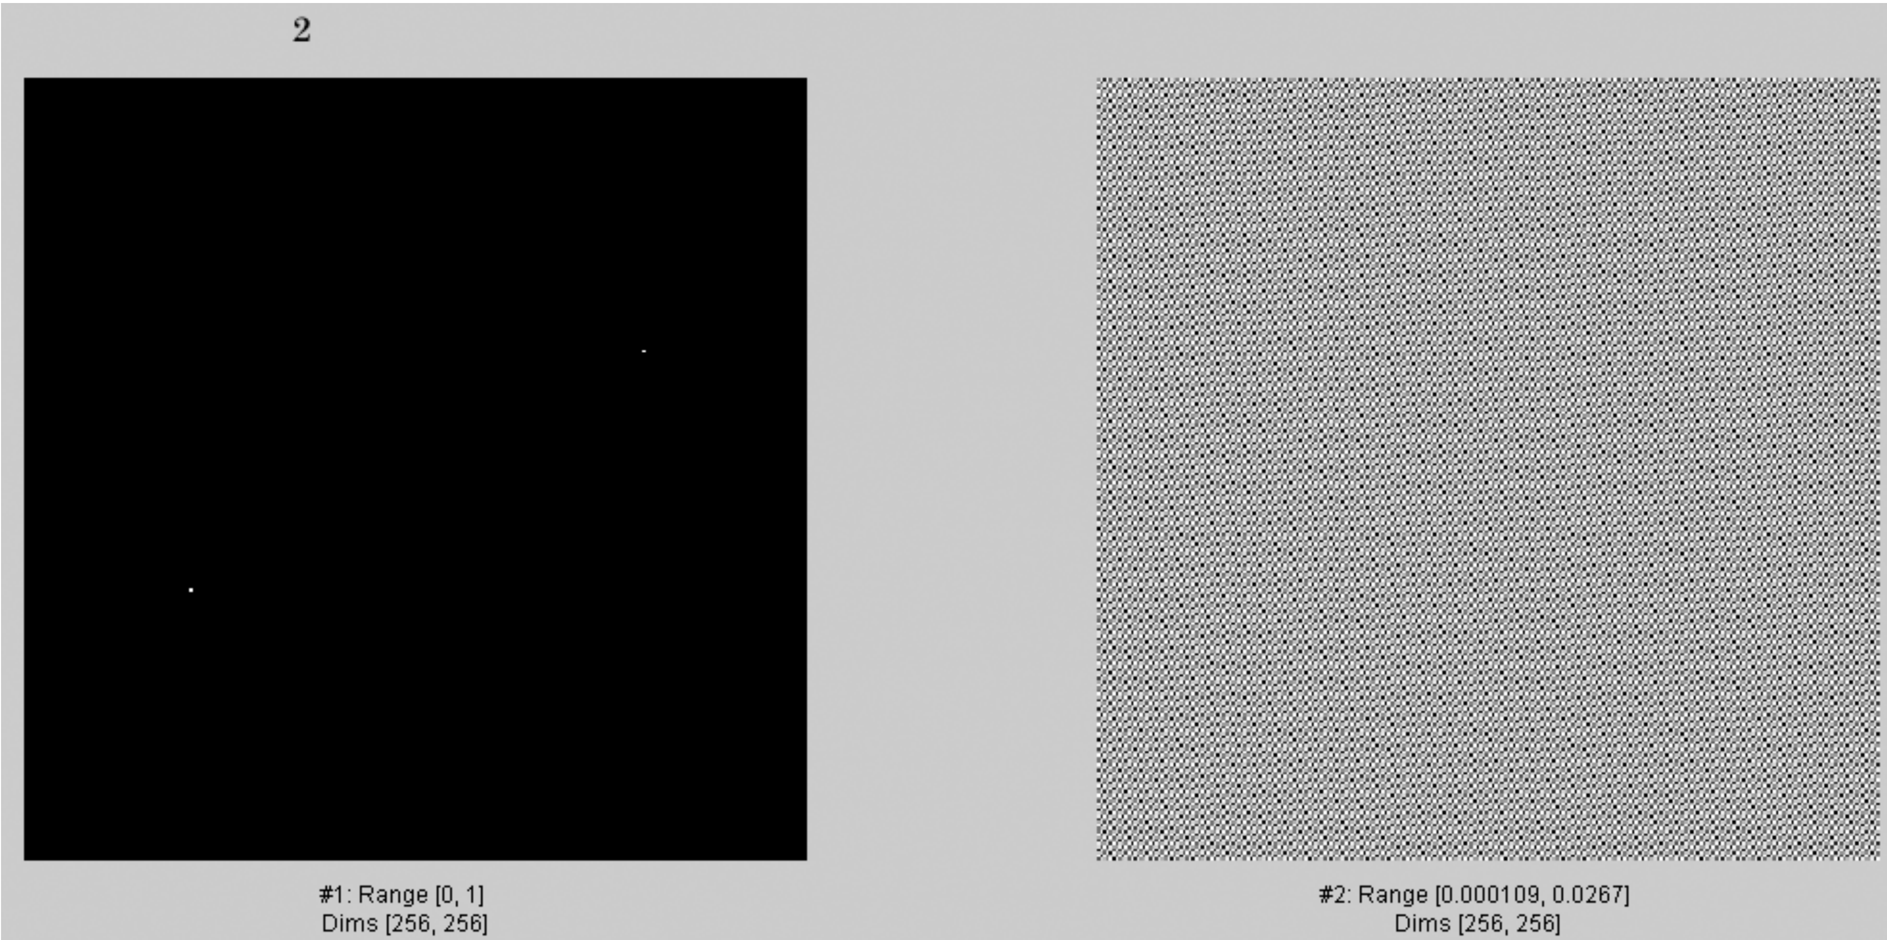

6

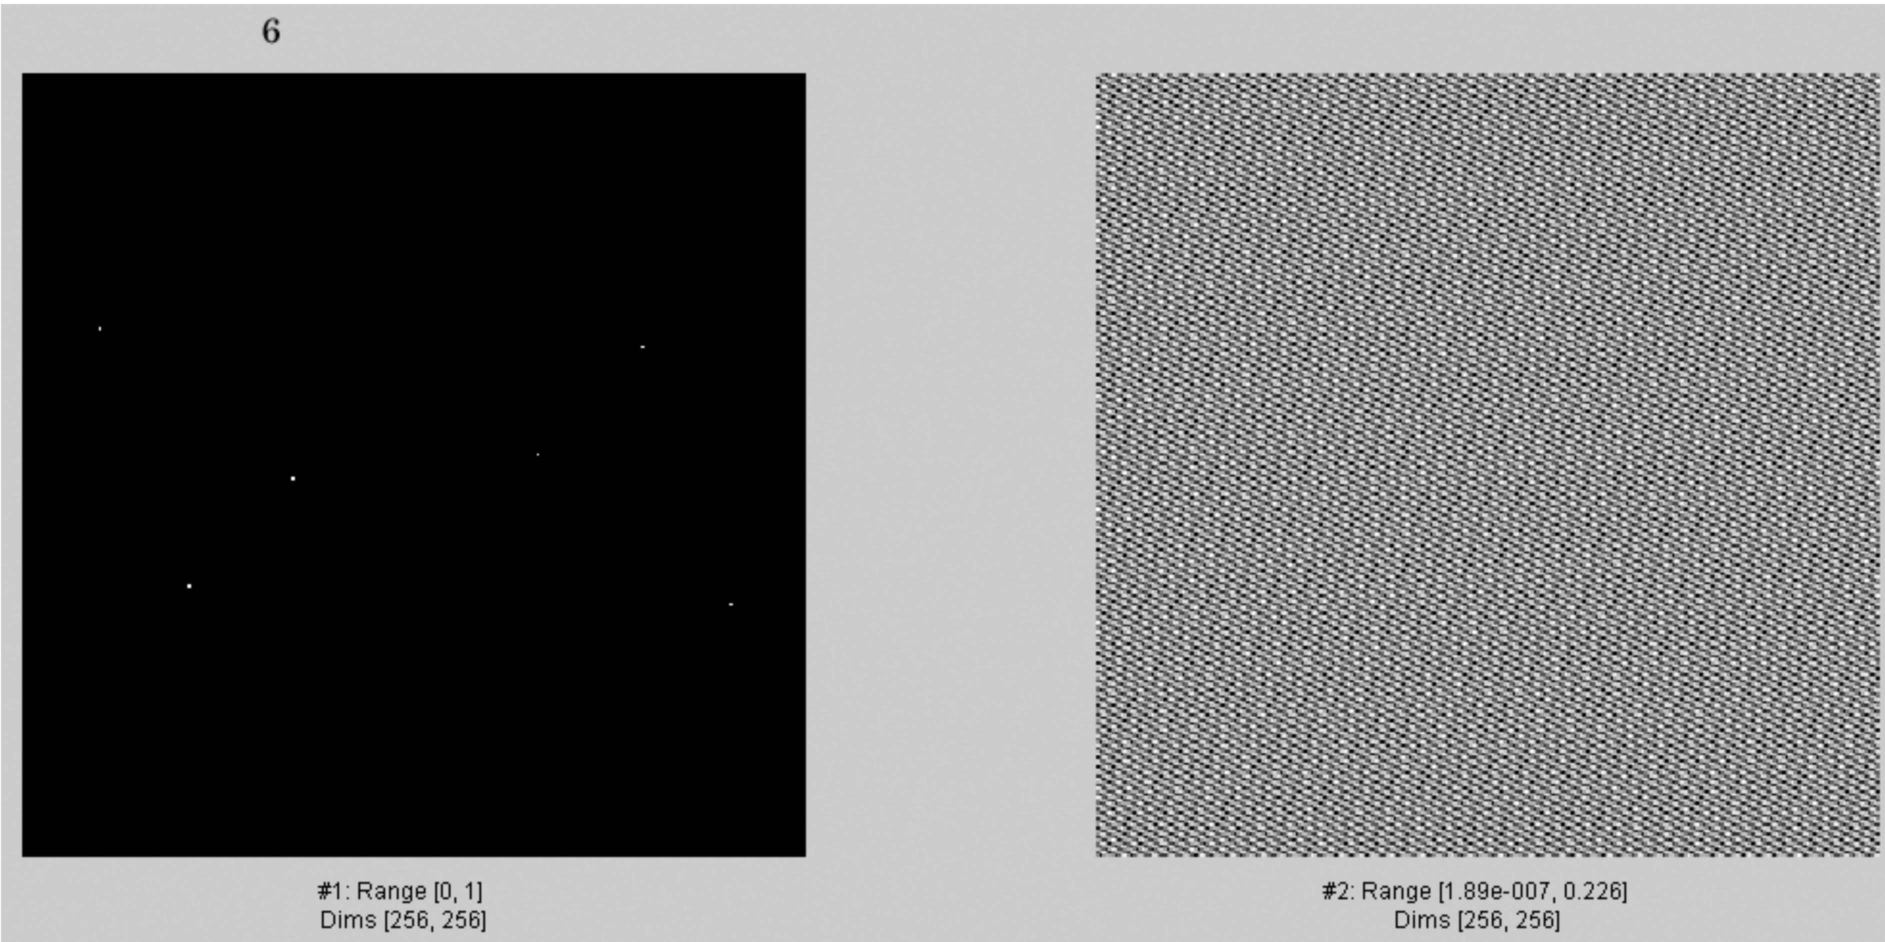

18

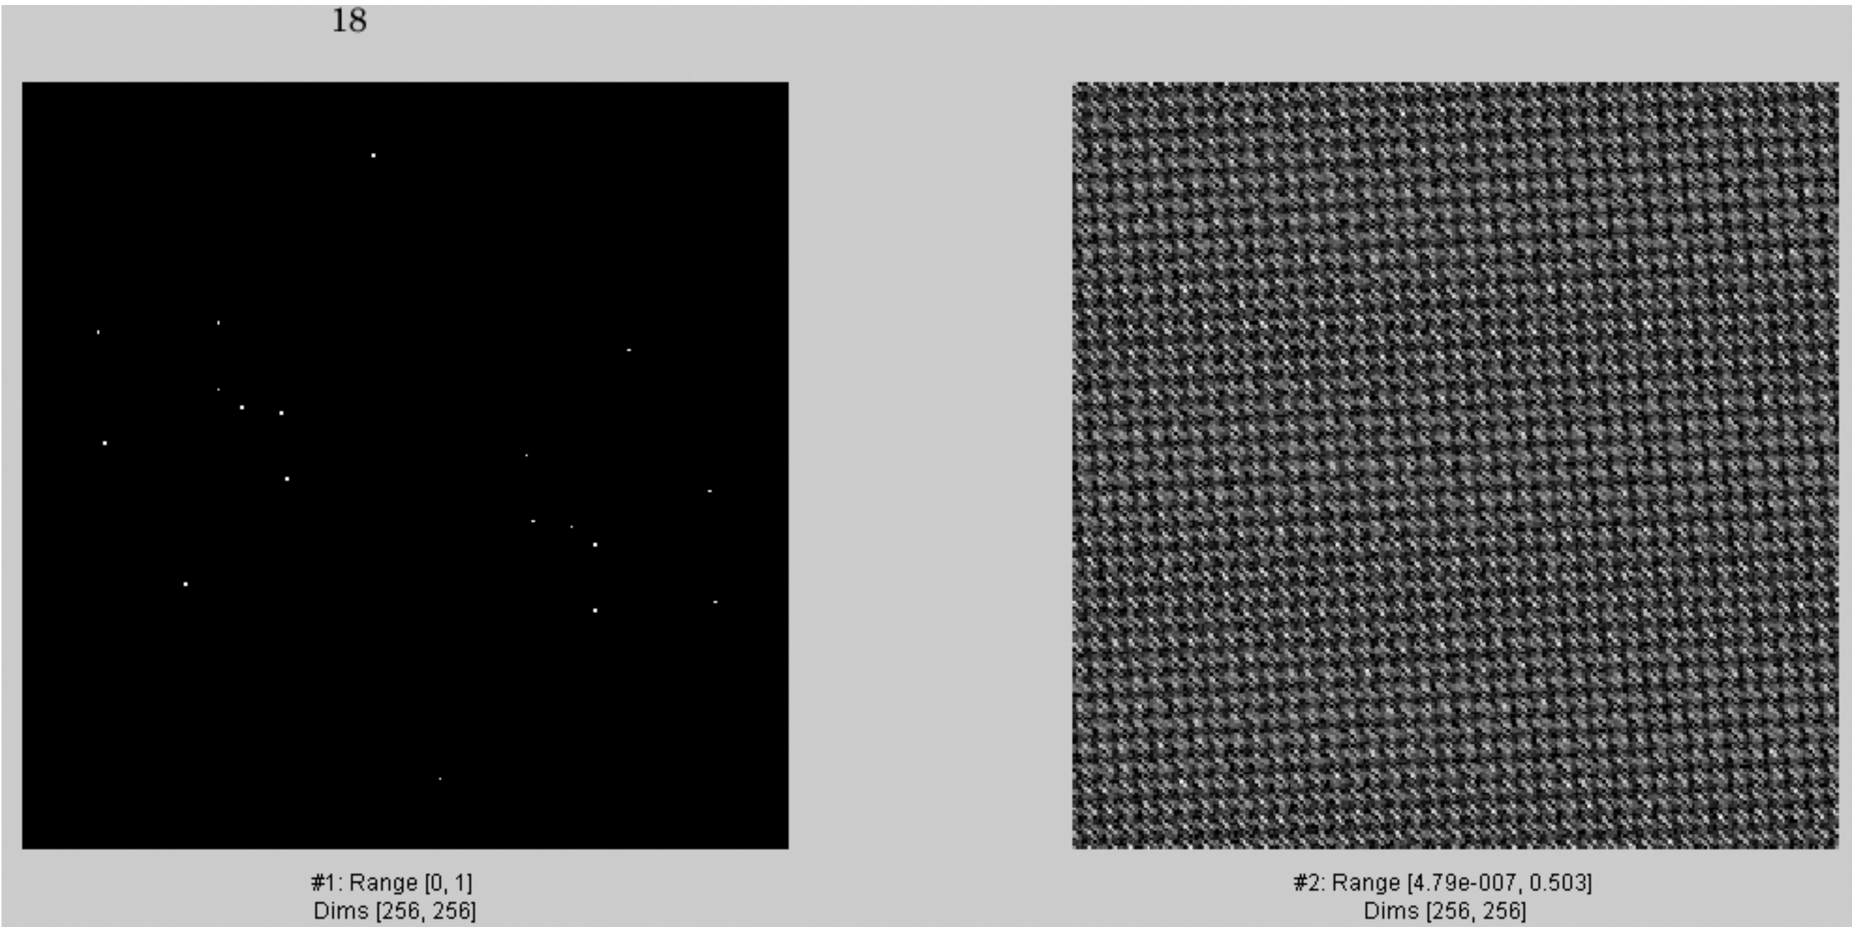

50

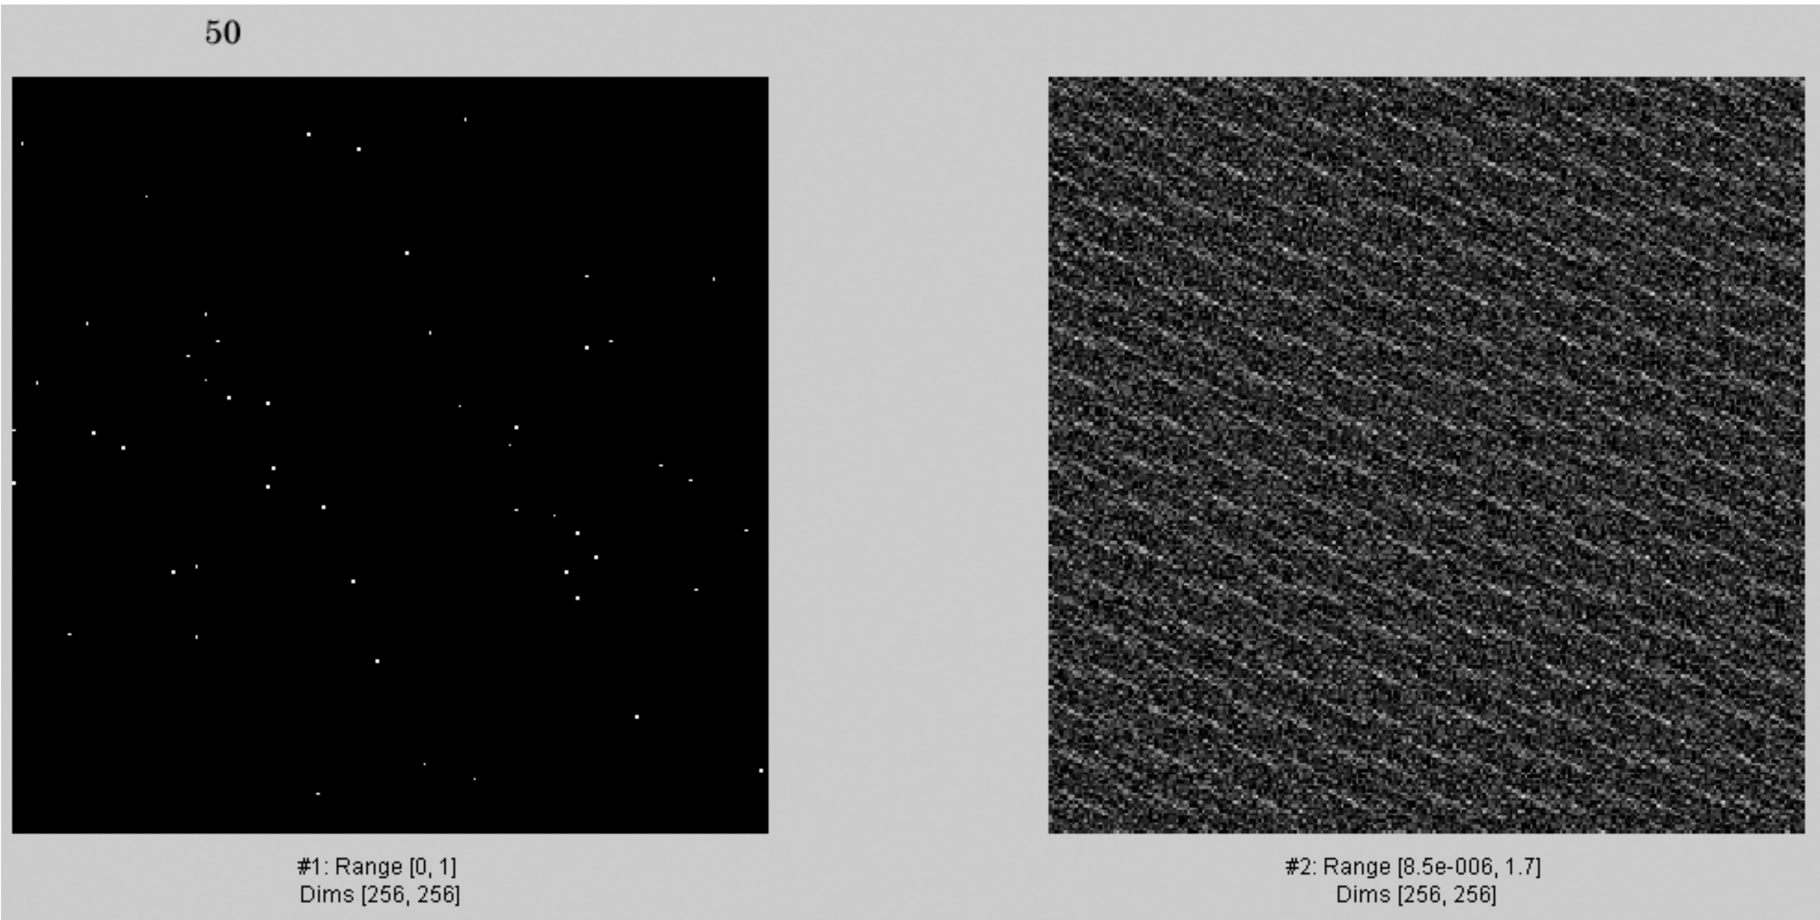

82

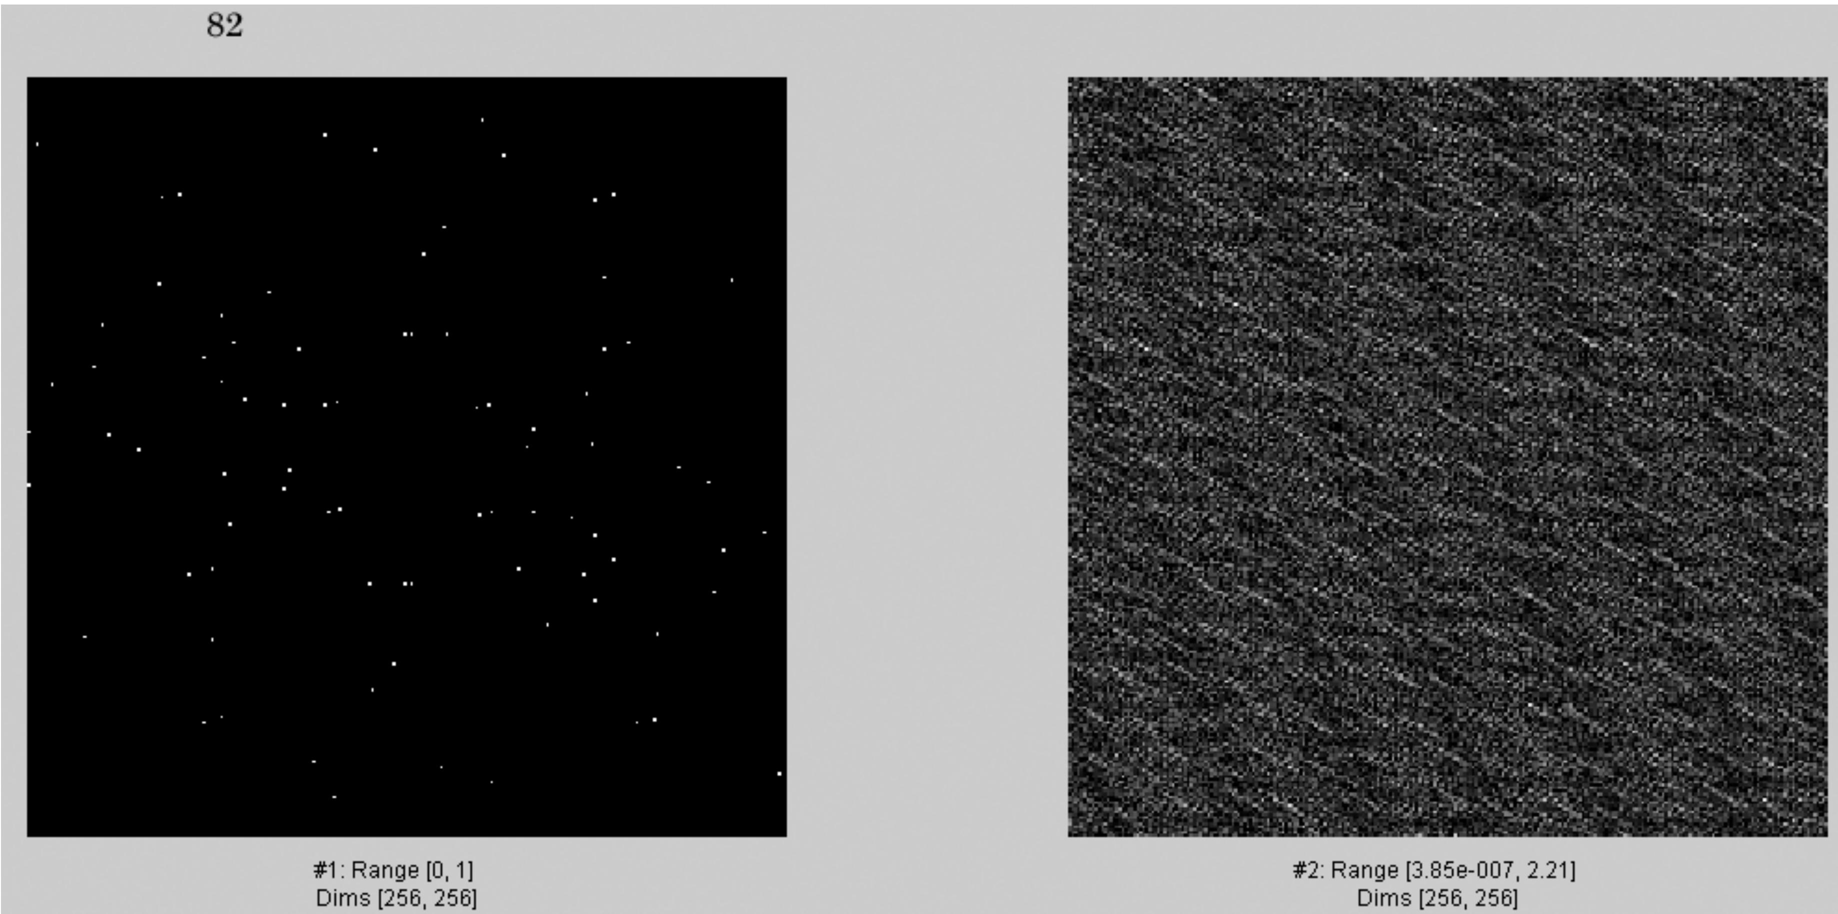

136

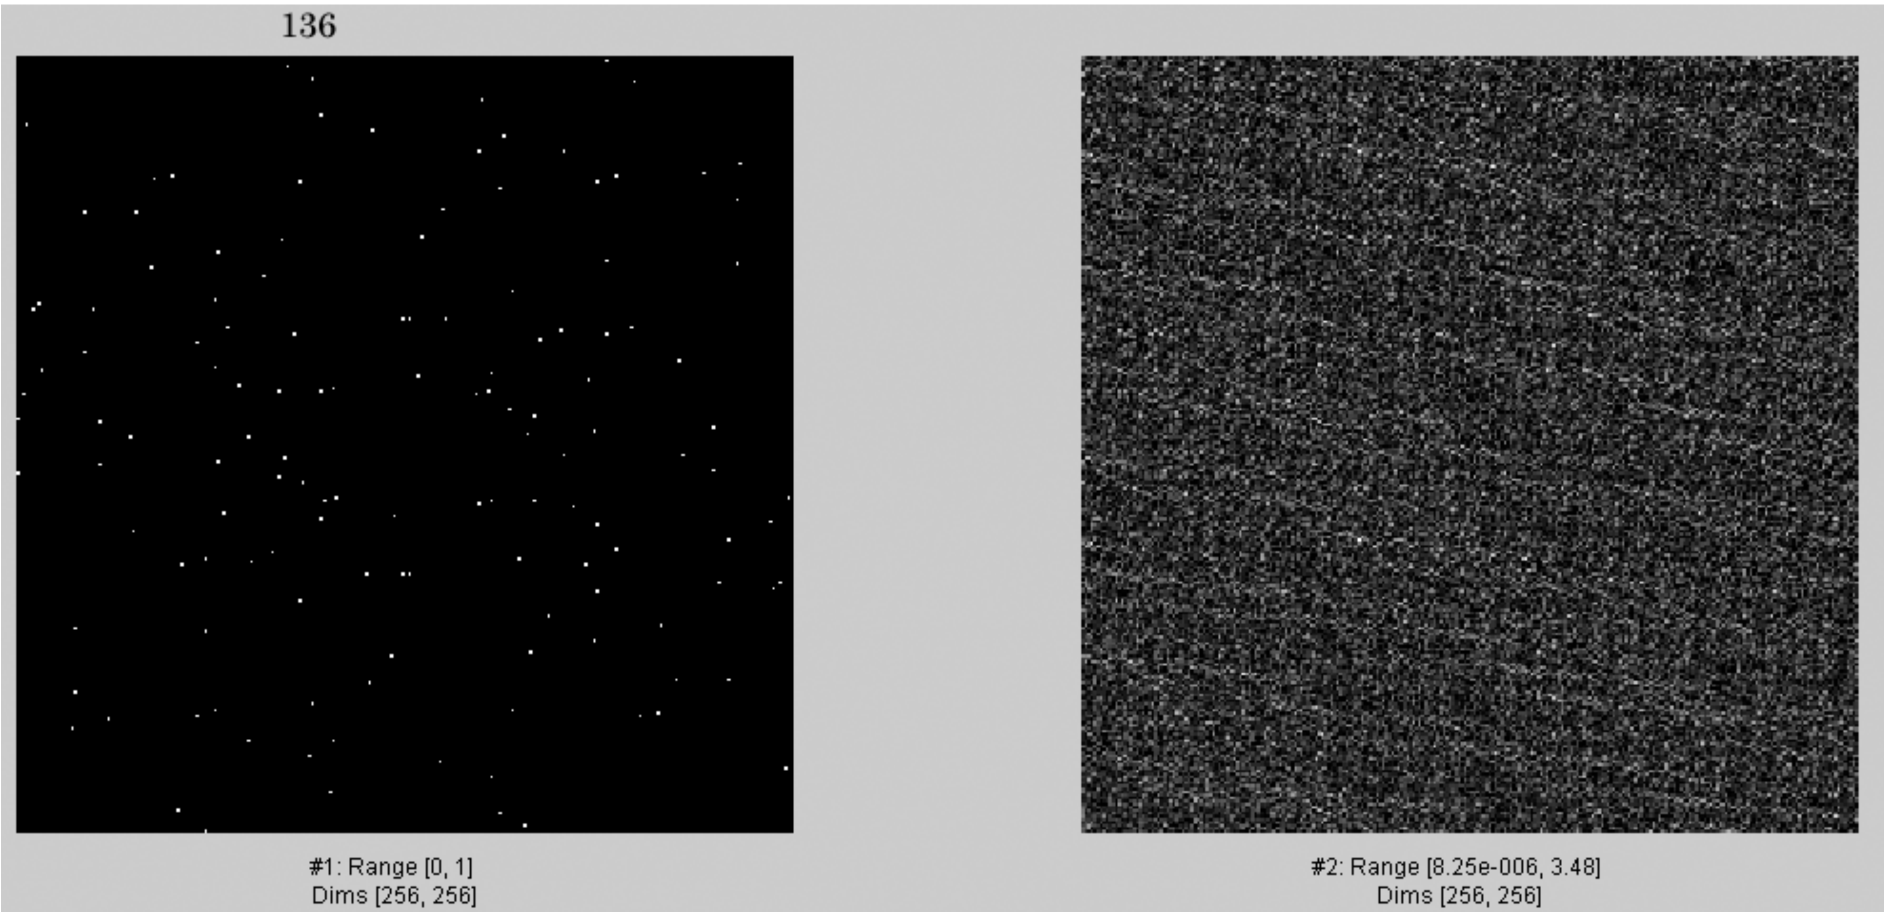

282

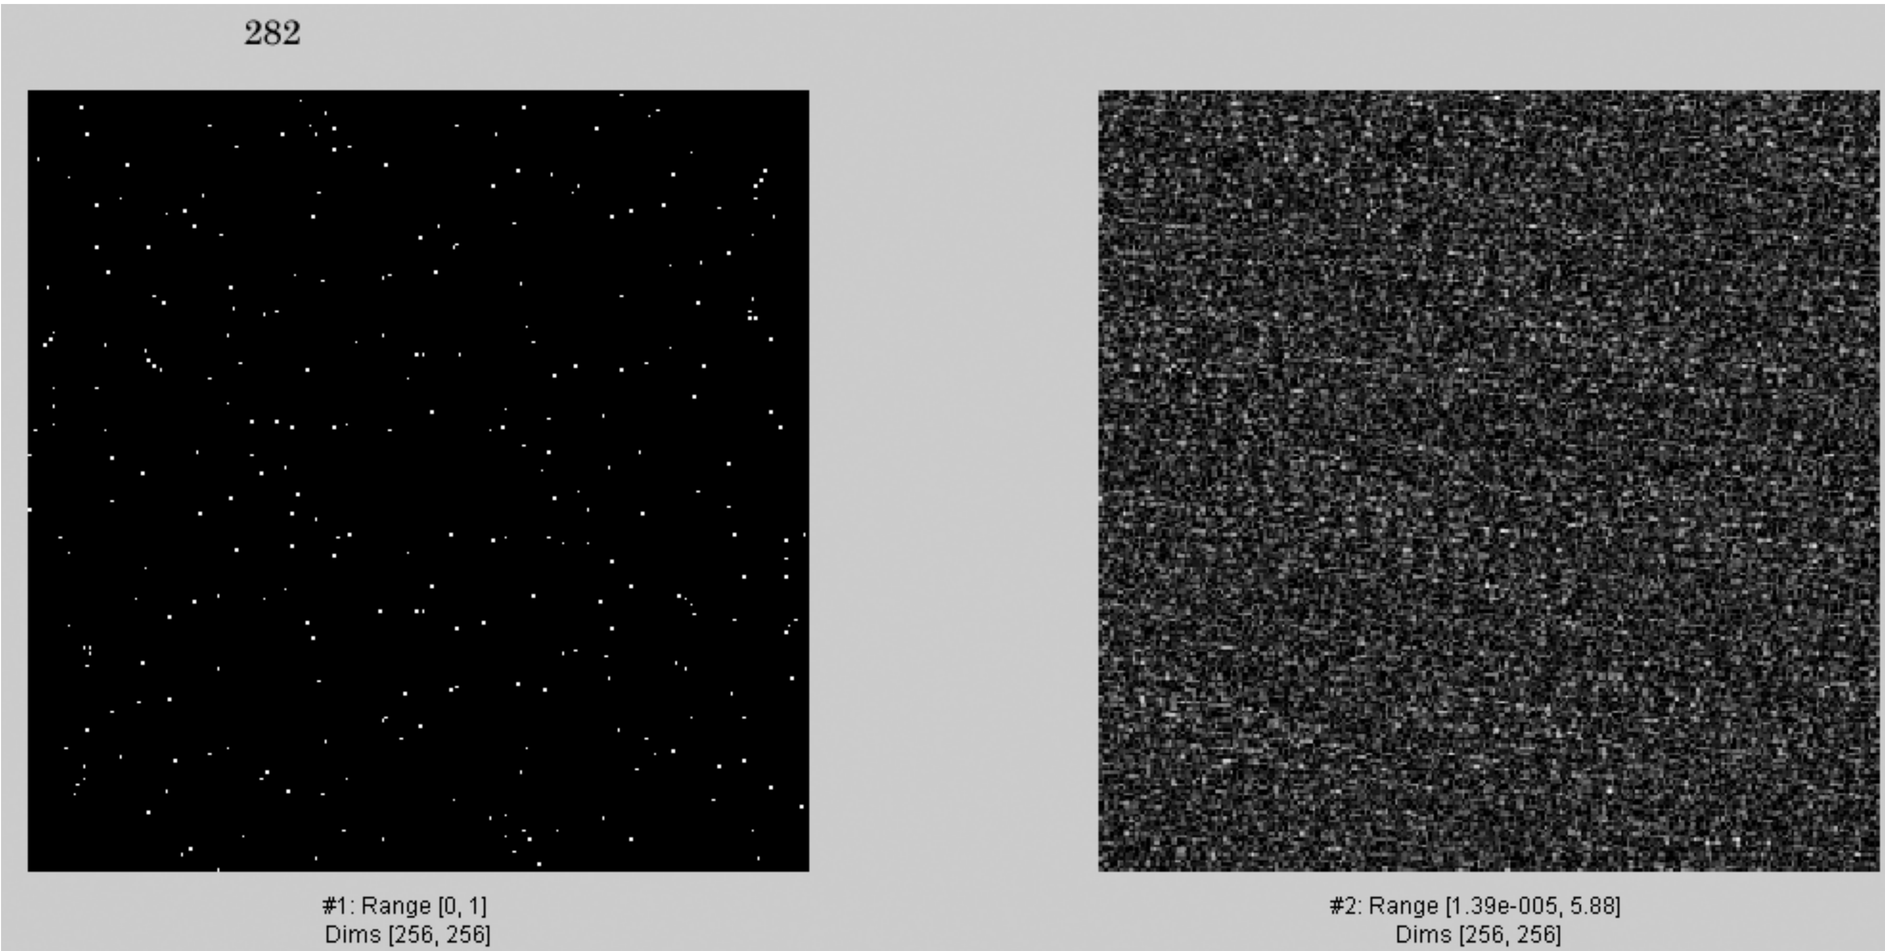

538

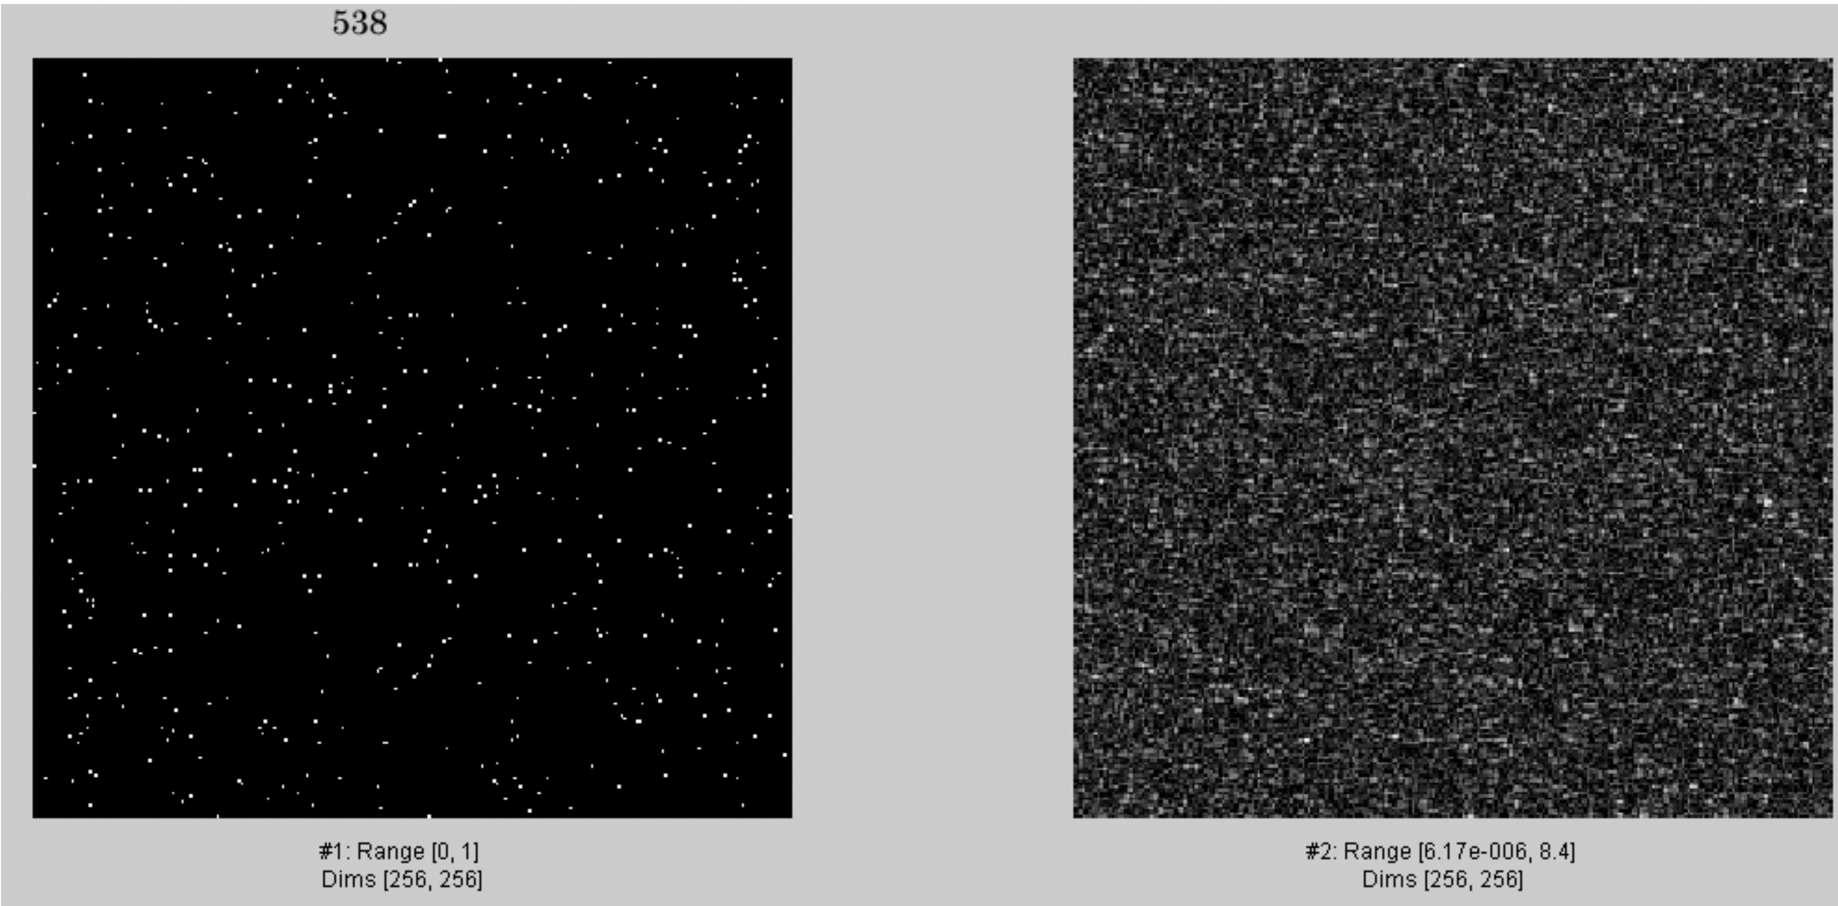

1088

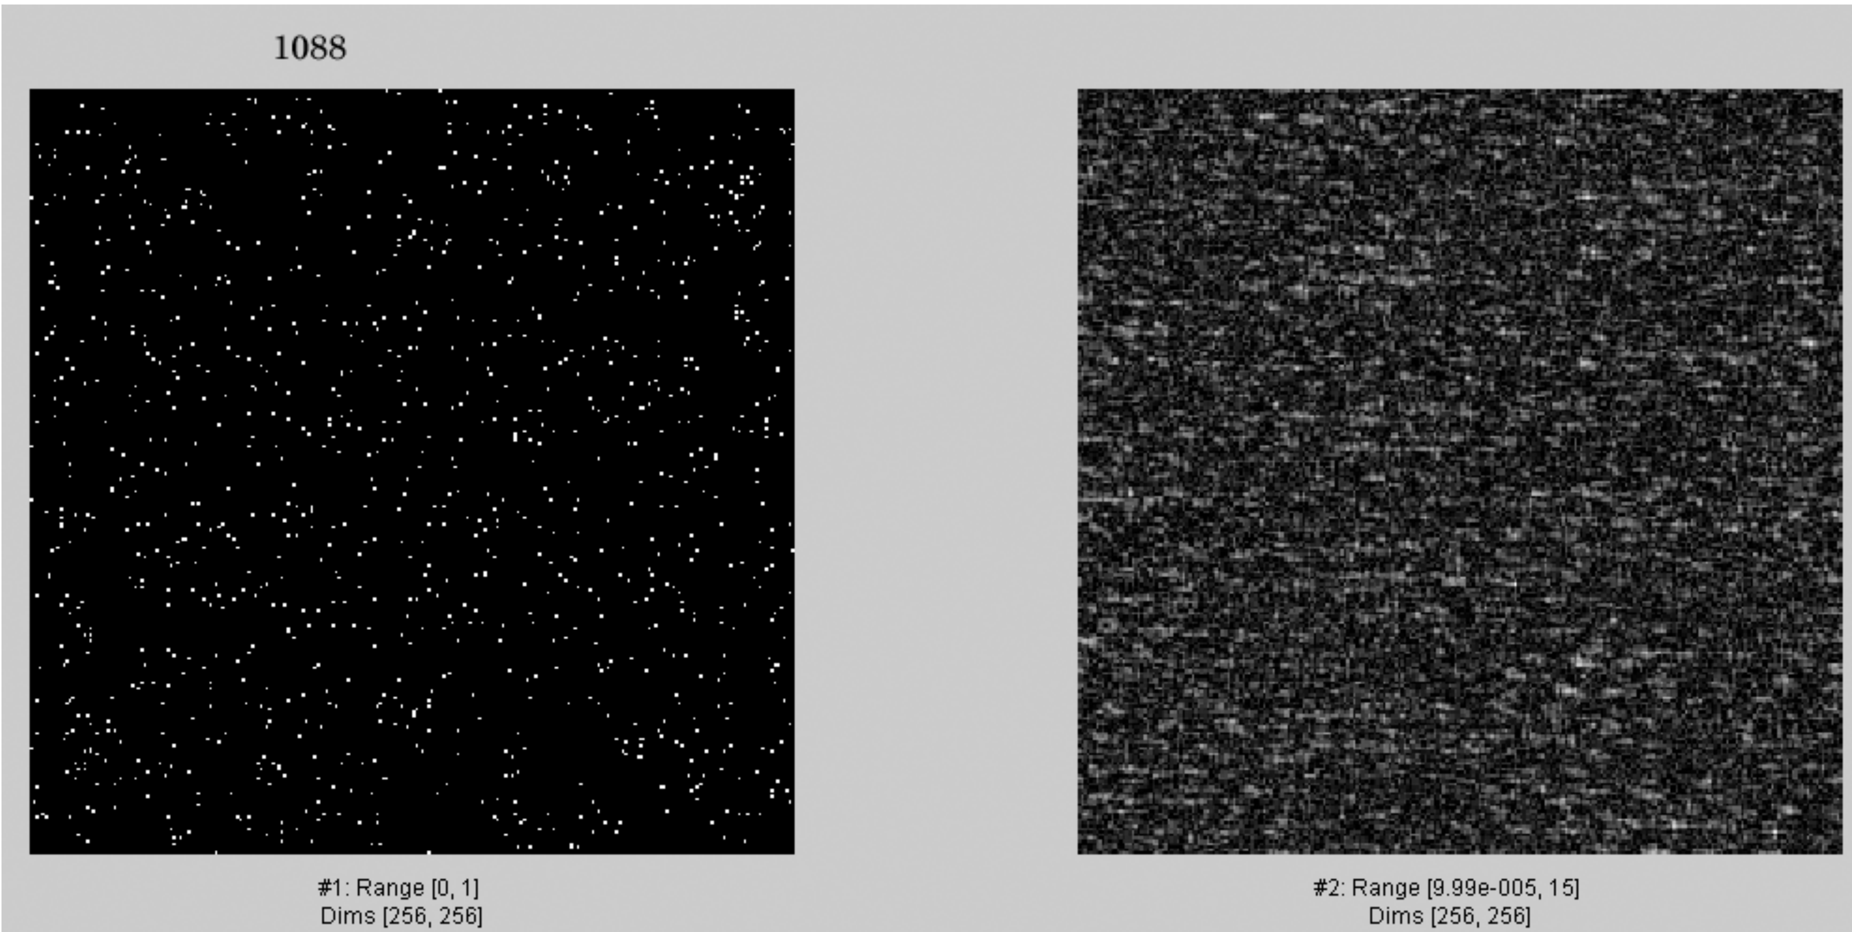

2094

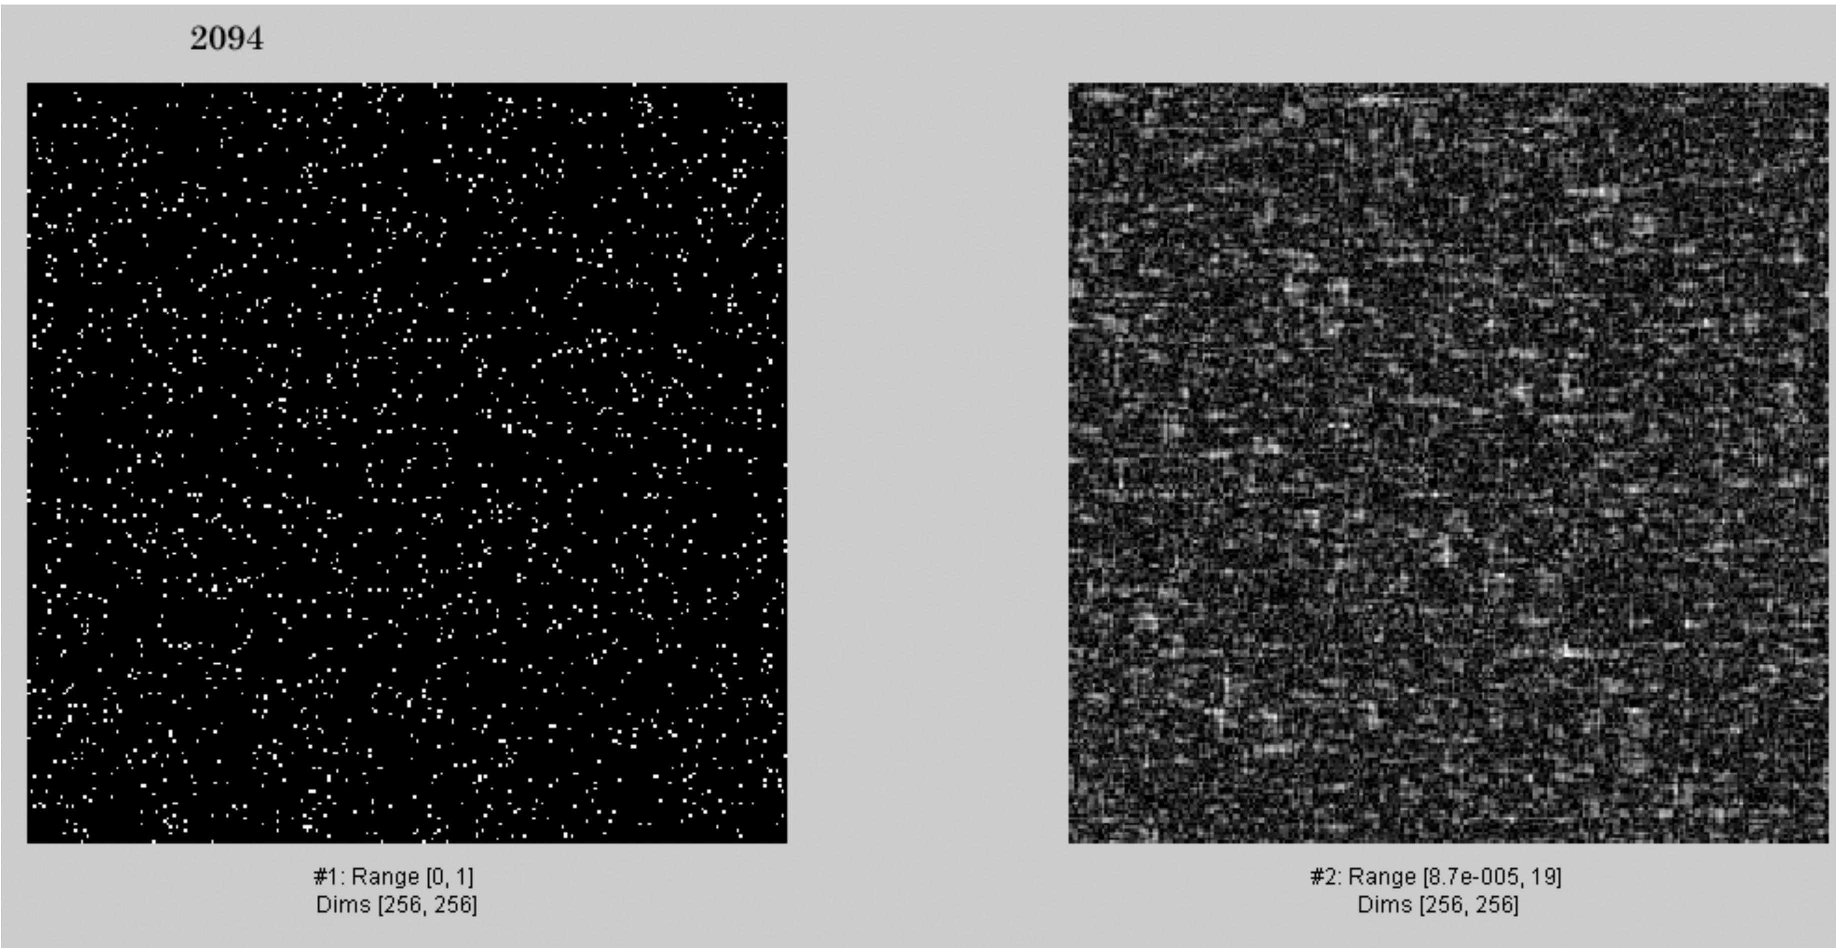

4052.

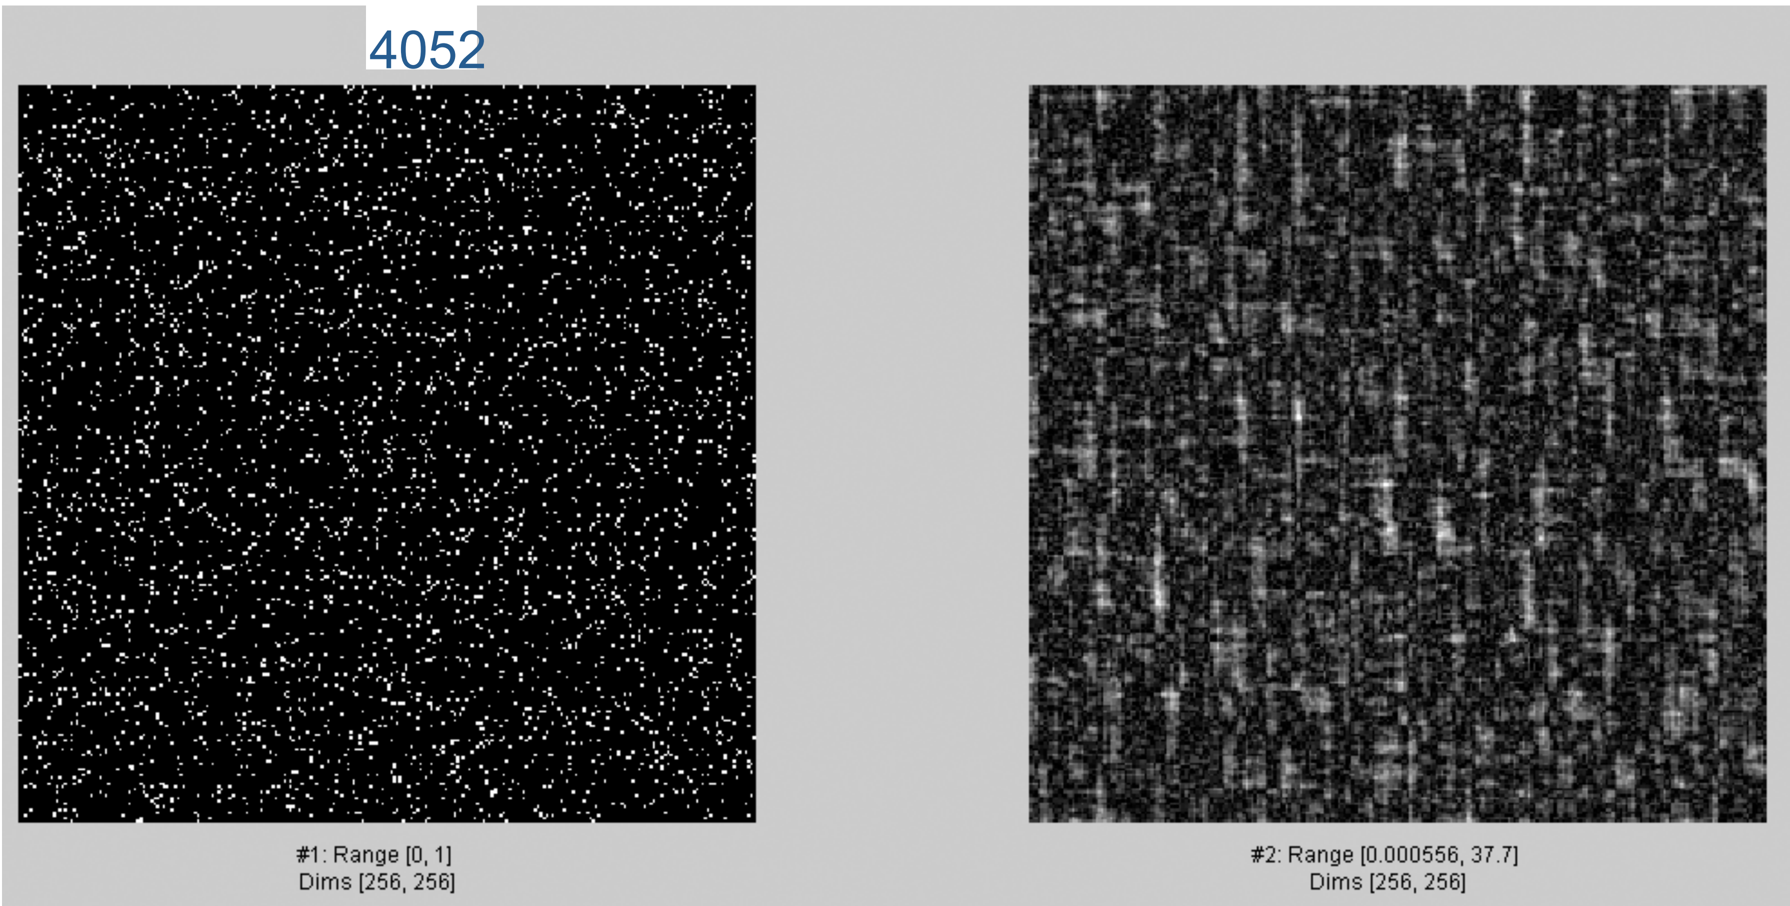

8056.

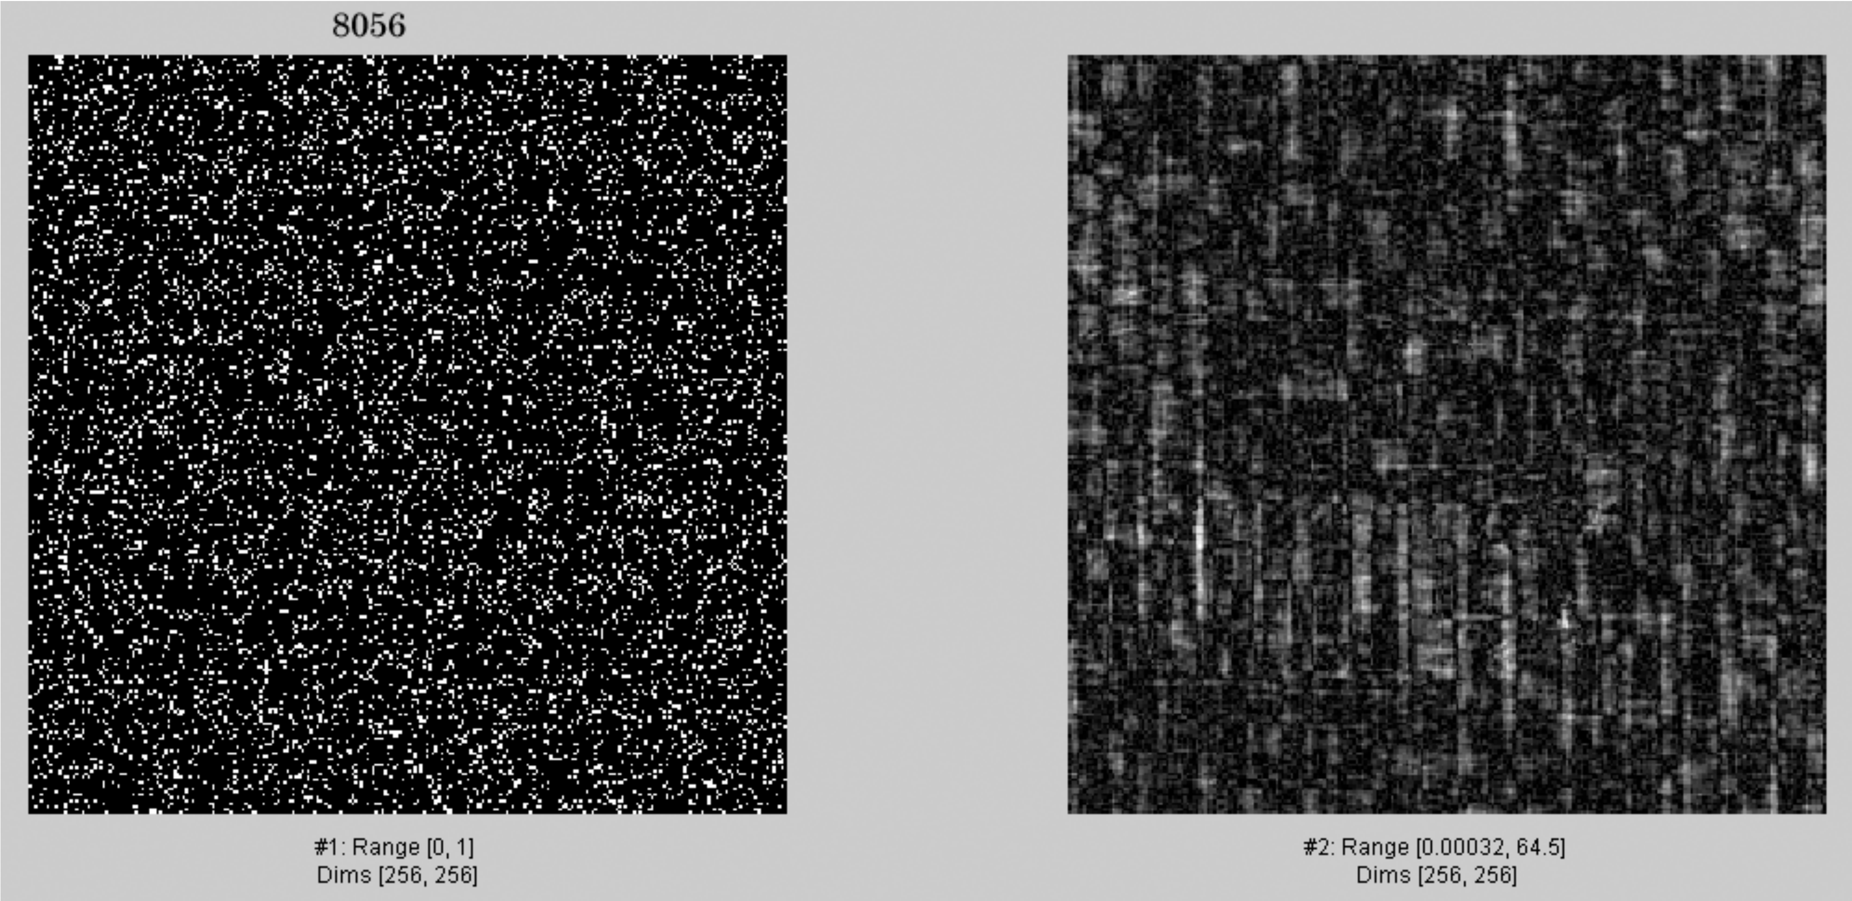

15366

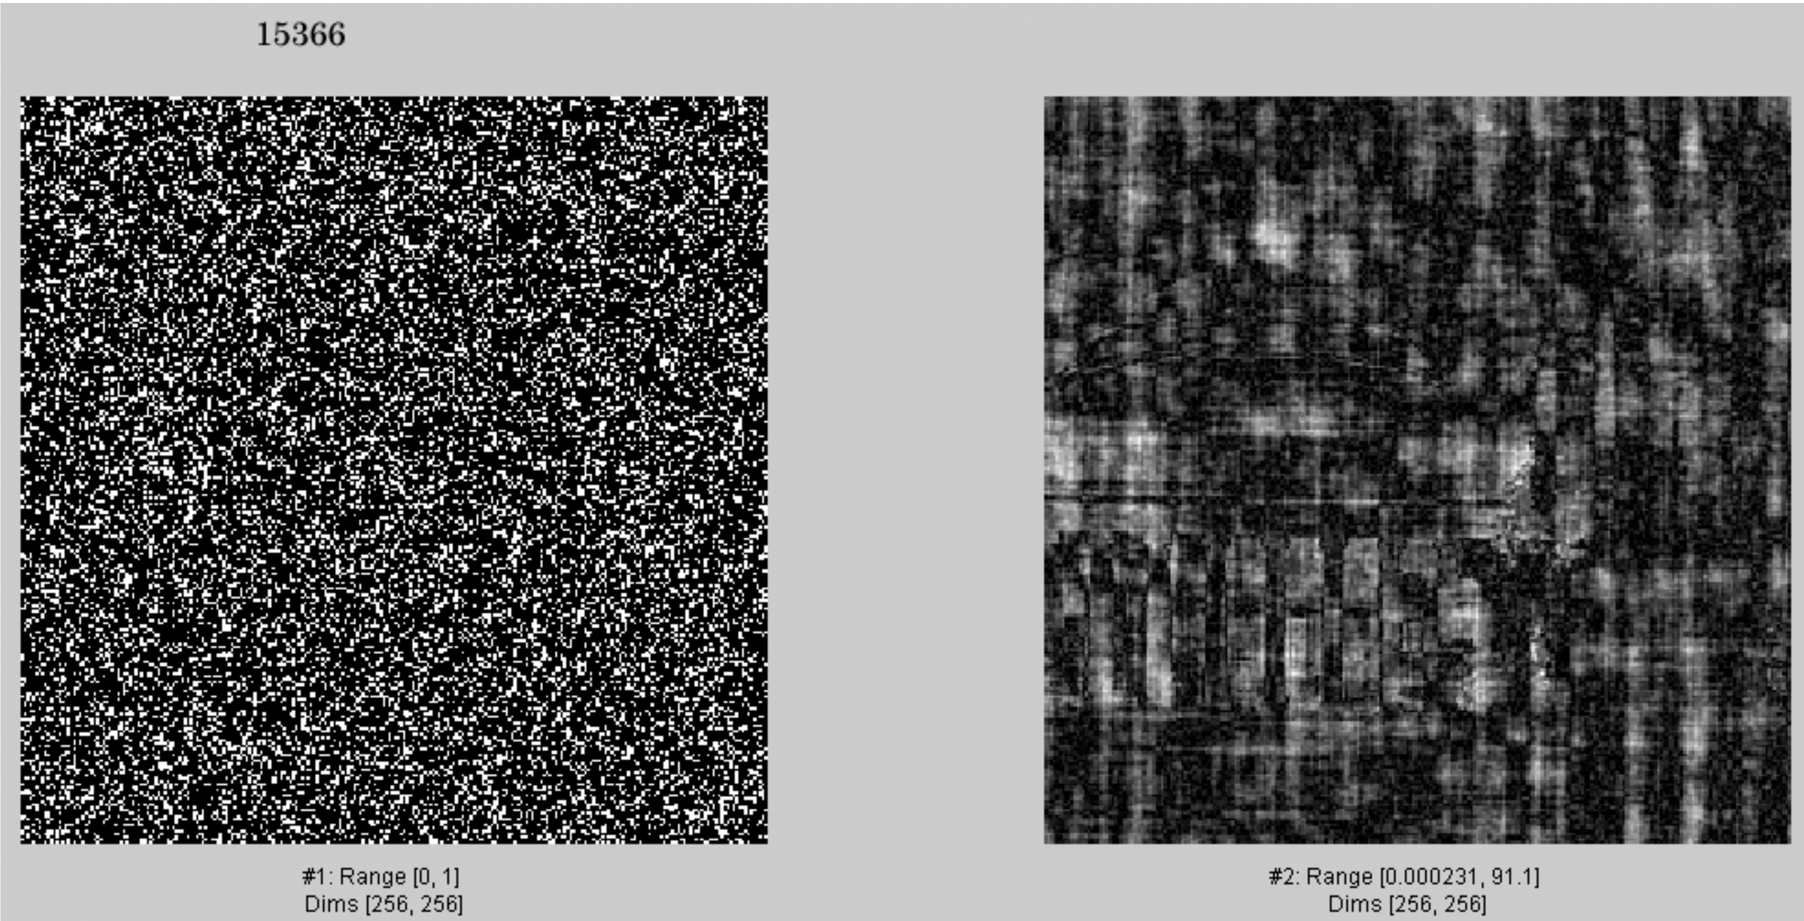

28743

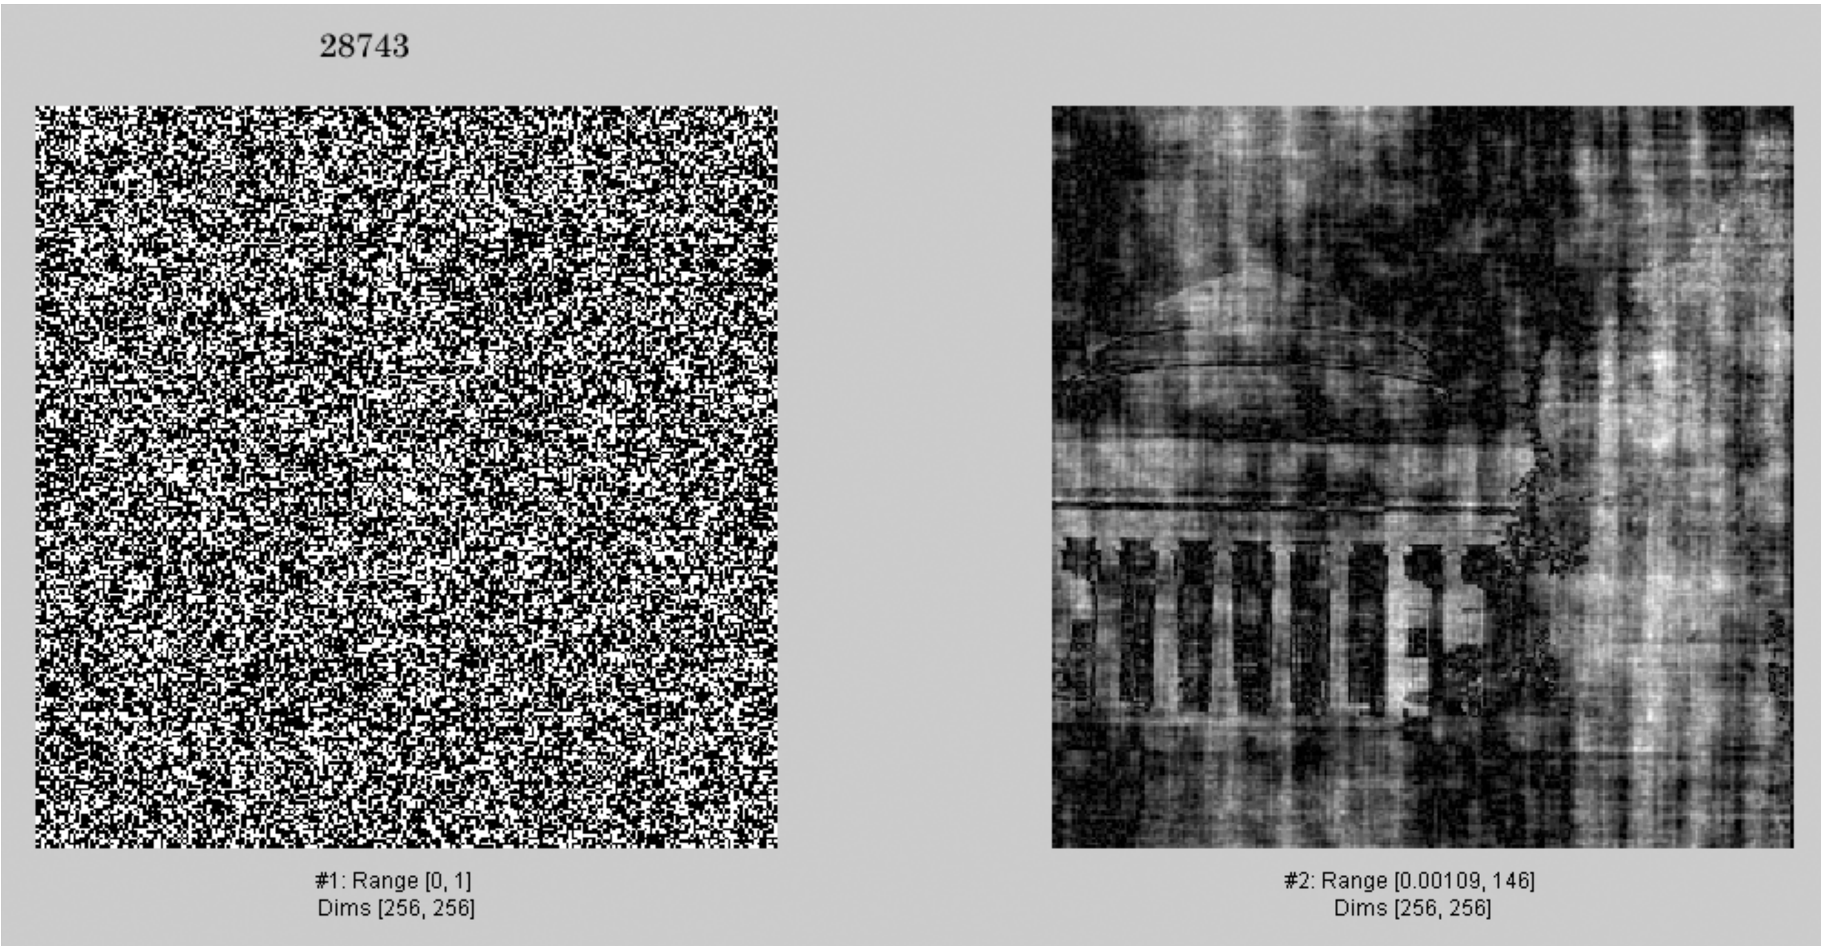

49190.

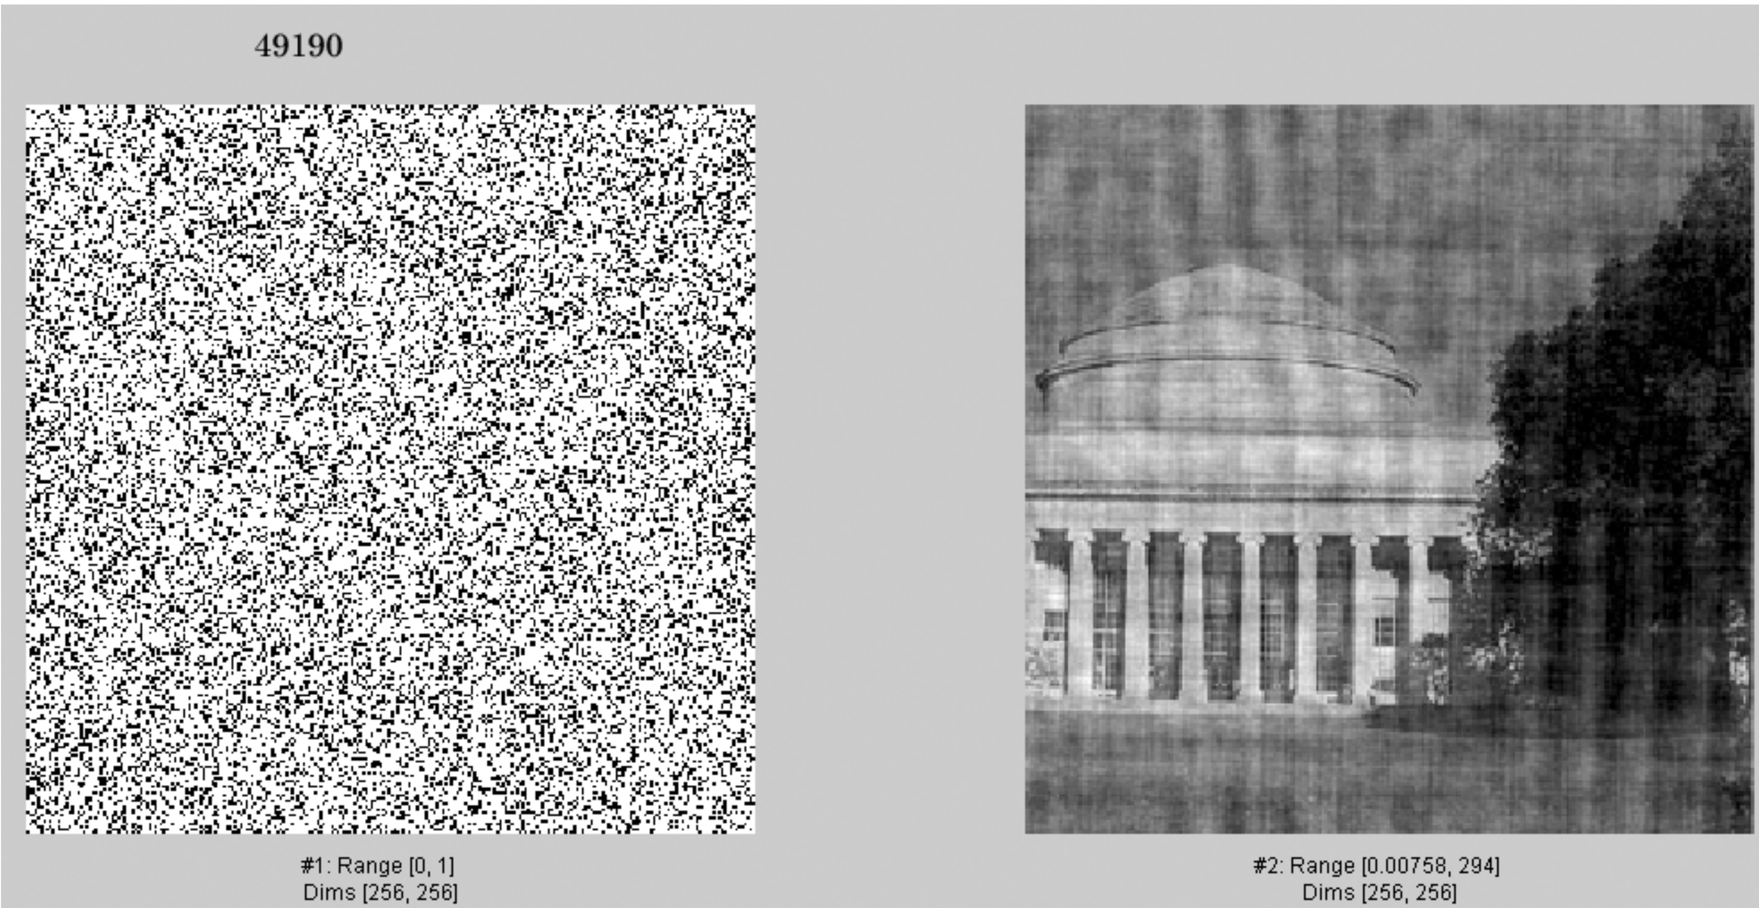

65536.

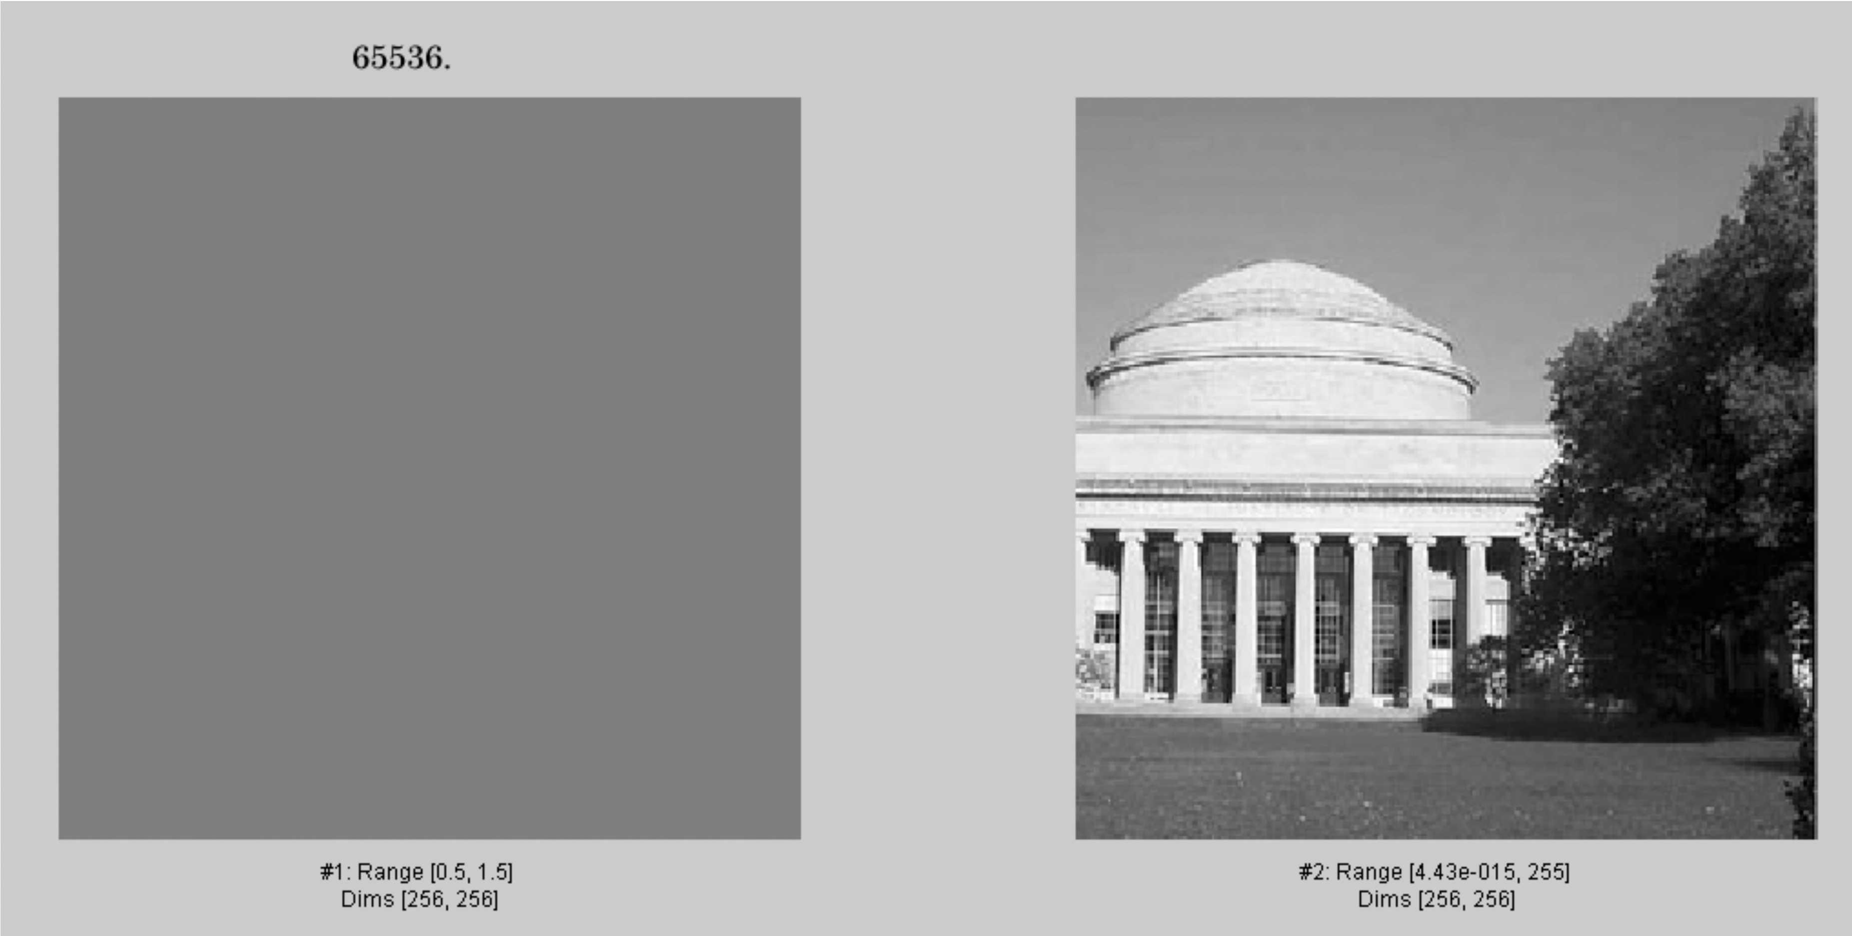
65536.
#1: Range [0.5, 1.5]
Dims [256, 256]
#2: Range [4.43e-015, 255]
Dims [256, 256]

3

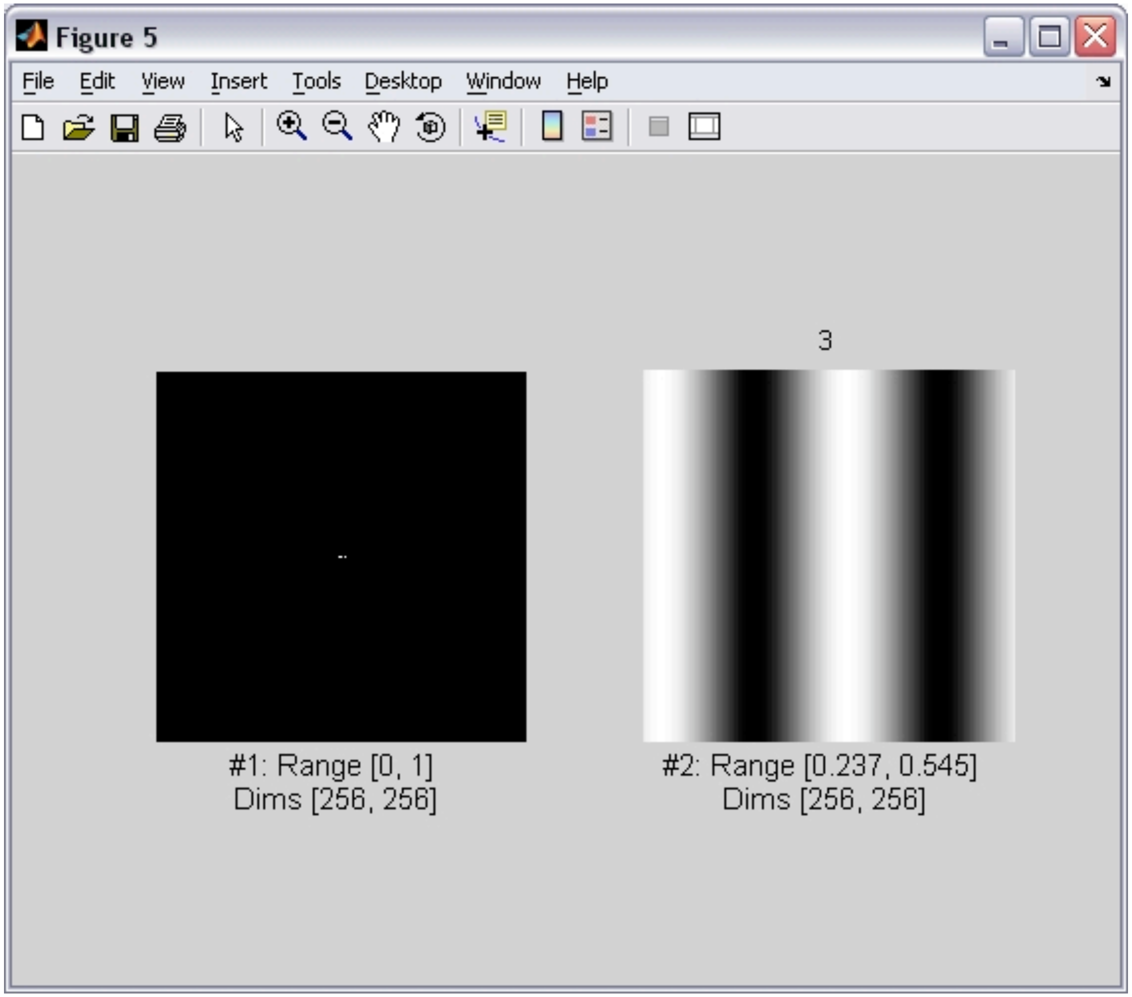

Now, an analogous sequence of images, but selecting Fourier components in descending order of magnitude.

5

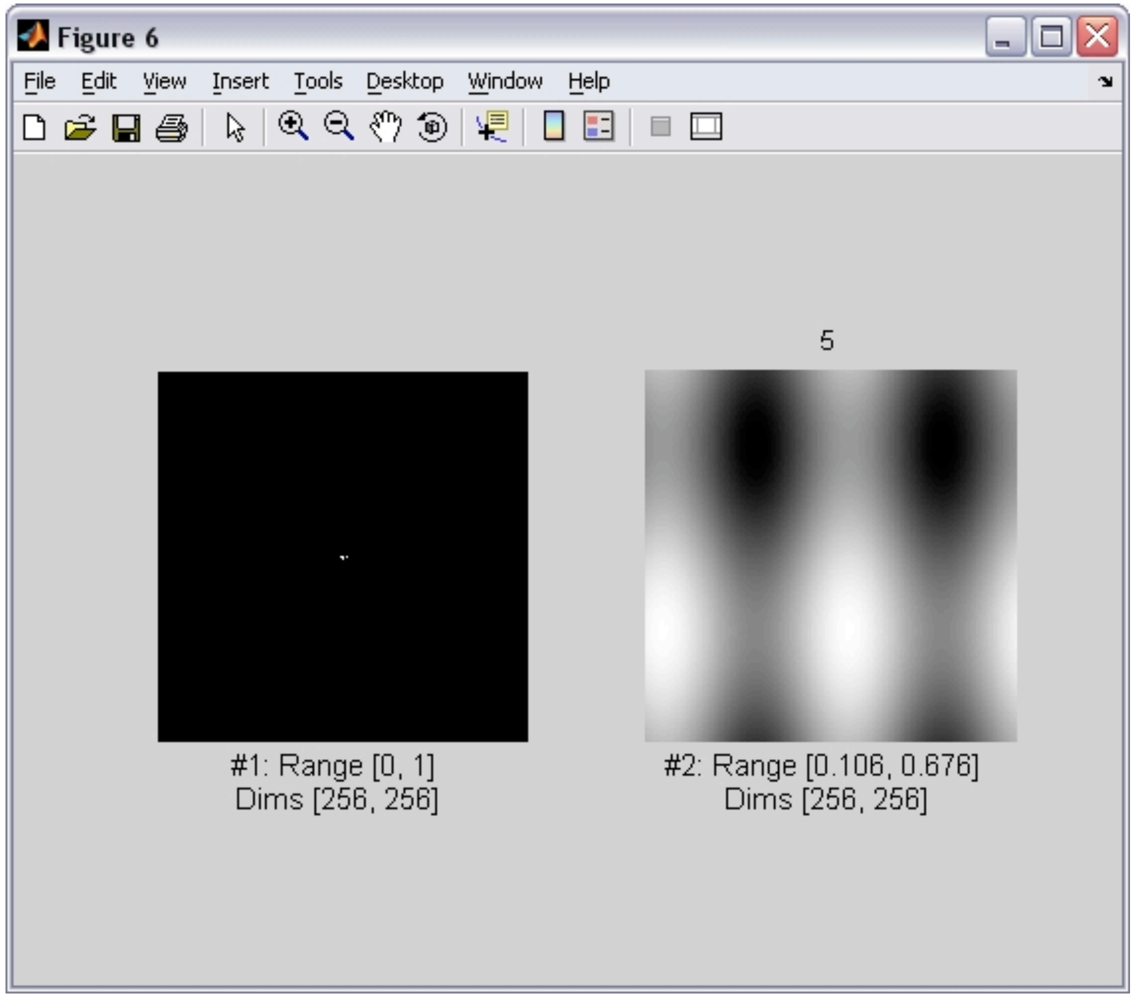
Figure 6
File Edit View Insert Tools Desktop Window Help
5
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.106, 0.676]
Dims [256, 256]

9

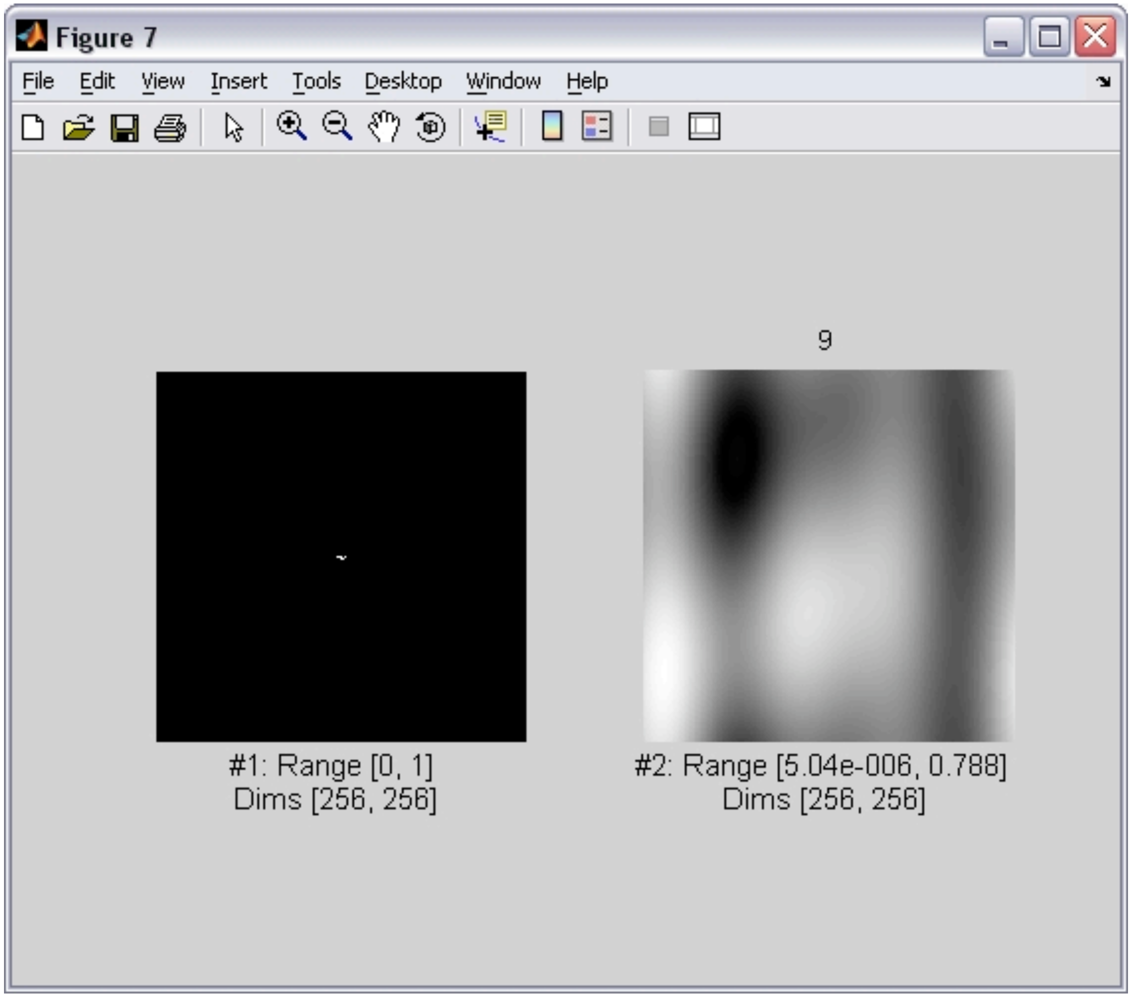
Figure 7
File Edit View Insert Tools Desktop Window Help
9
#1: Range [0, 1]
Dims [256, 256]
#2: Range [5.04e-006, 0.788]
Dims [256, 256]

17

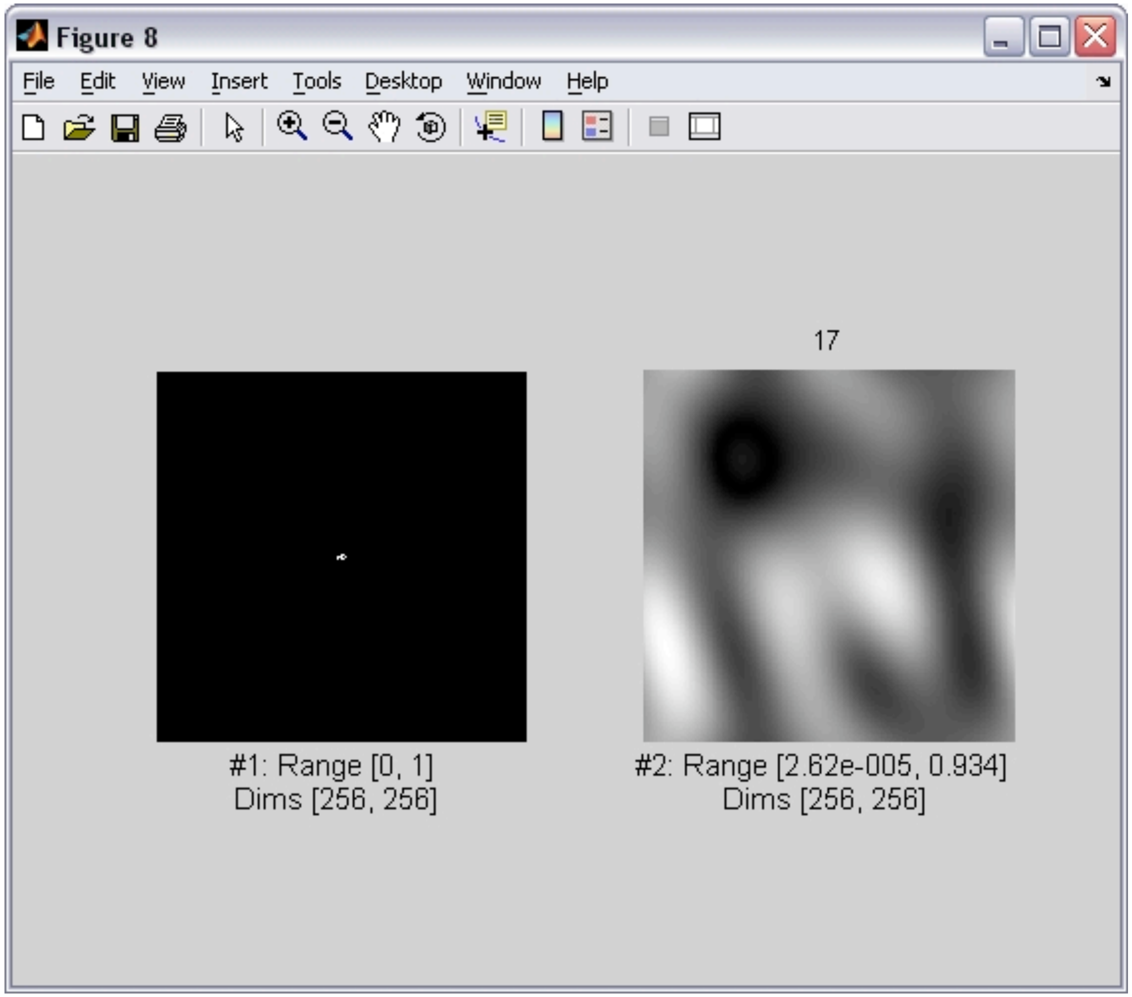
Figure 8
File
Edit
View
Insert
Tools
Desktop
Window
Help
17
#1: Range [0, 1]
Dims [256, 256]
#2: Range [2.62e-005, 0.934]
Dims [256, 256]

33

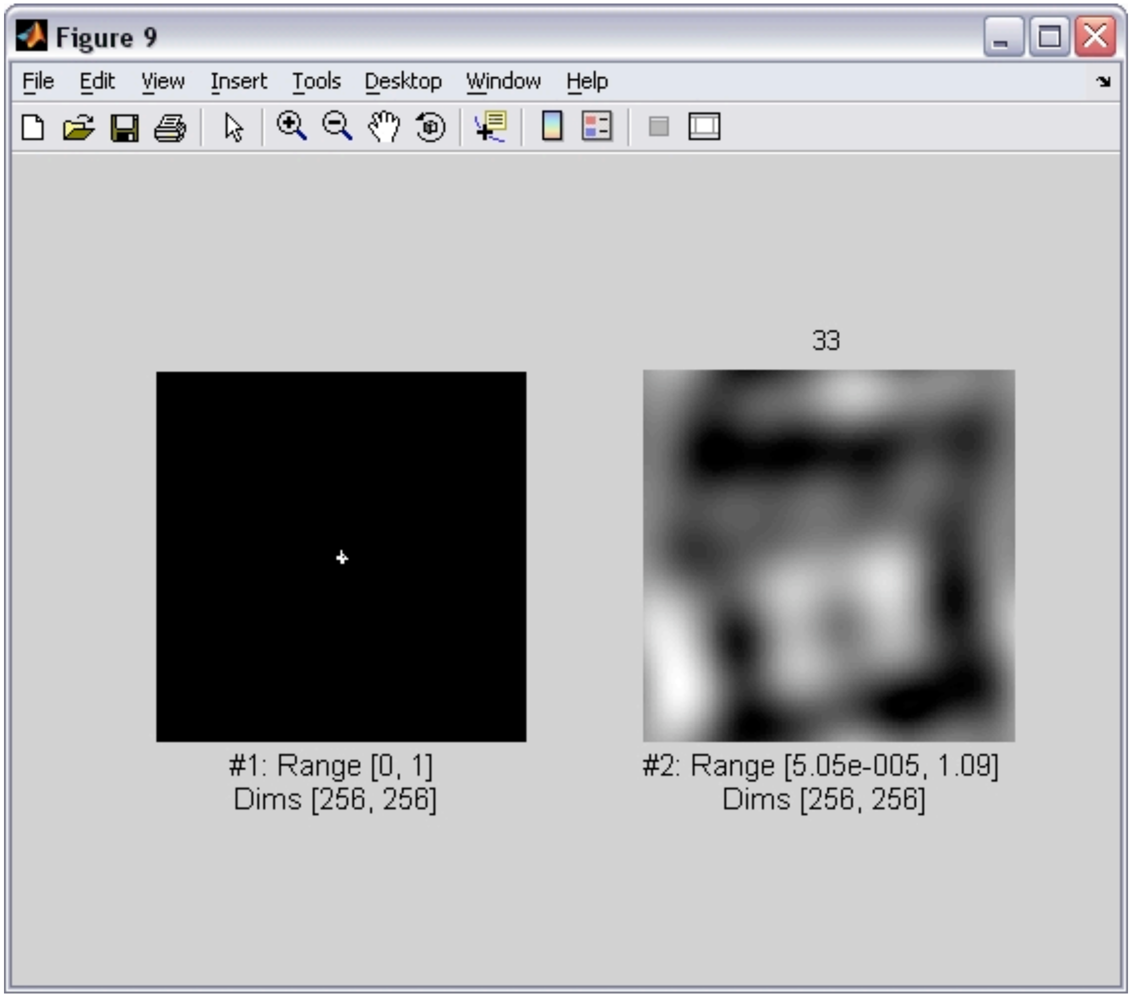
Figure 9
File Edit View Insert Tools Desktop Window Help
33
#1: Range [0, 1]
Dims [256, 256]
#2: Range [5.05e-005, 1.09]
Dims [256, 256]

65

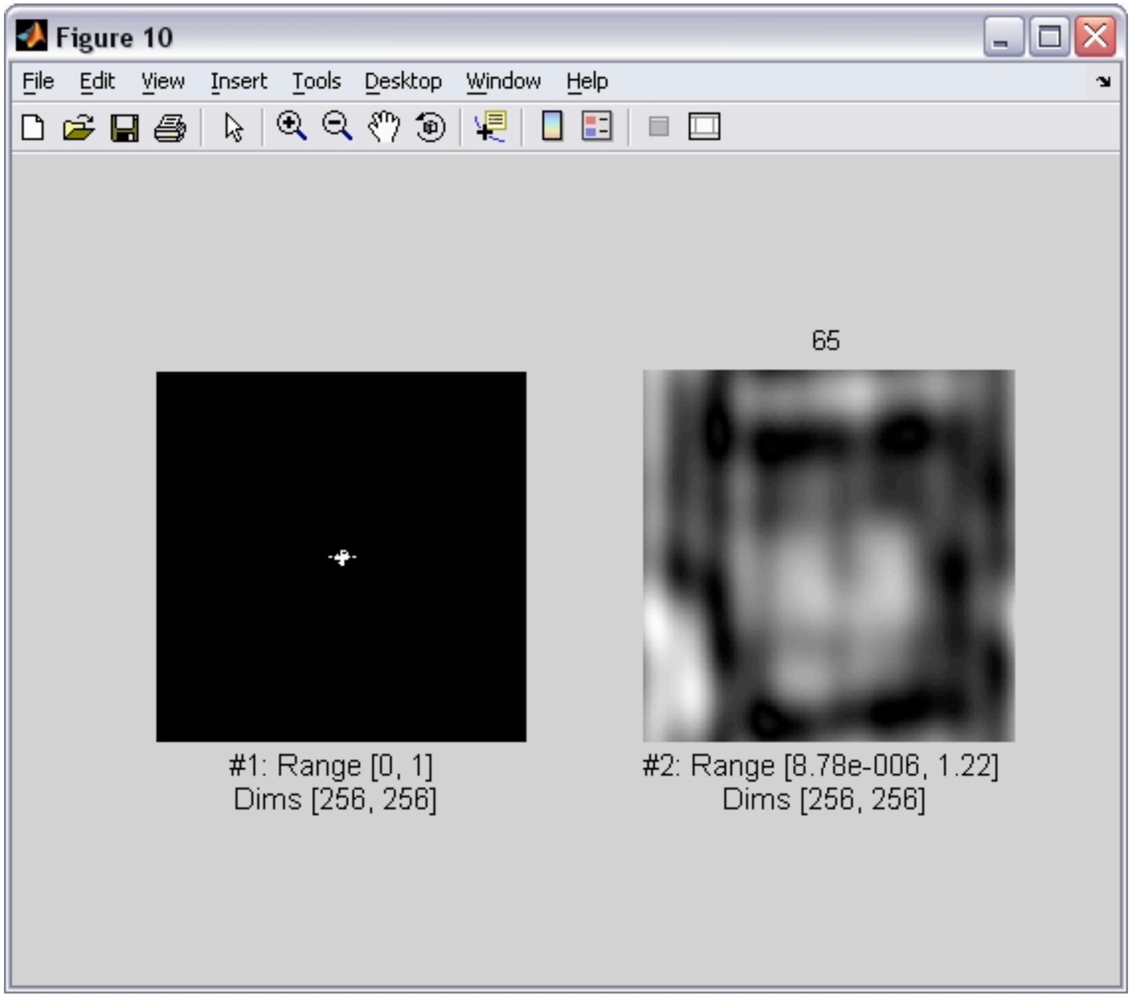
Figure 10
File
Edit
View
Insert
Tools
Desktop
Window
Help
65
#1: Range [0, 1]
Dims [256, 256]
#2: Range [8.78e-006, 1.22]
Dims [256, 256]

129

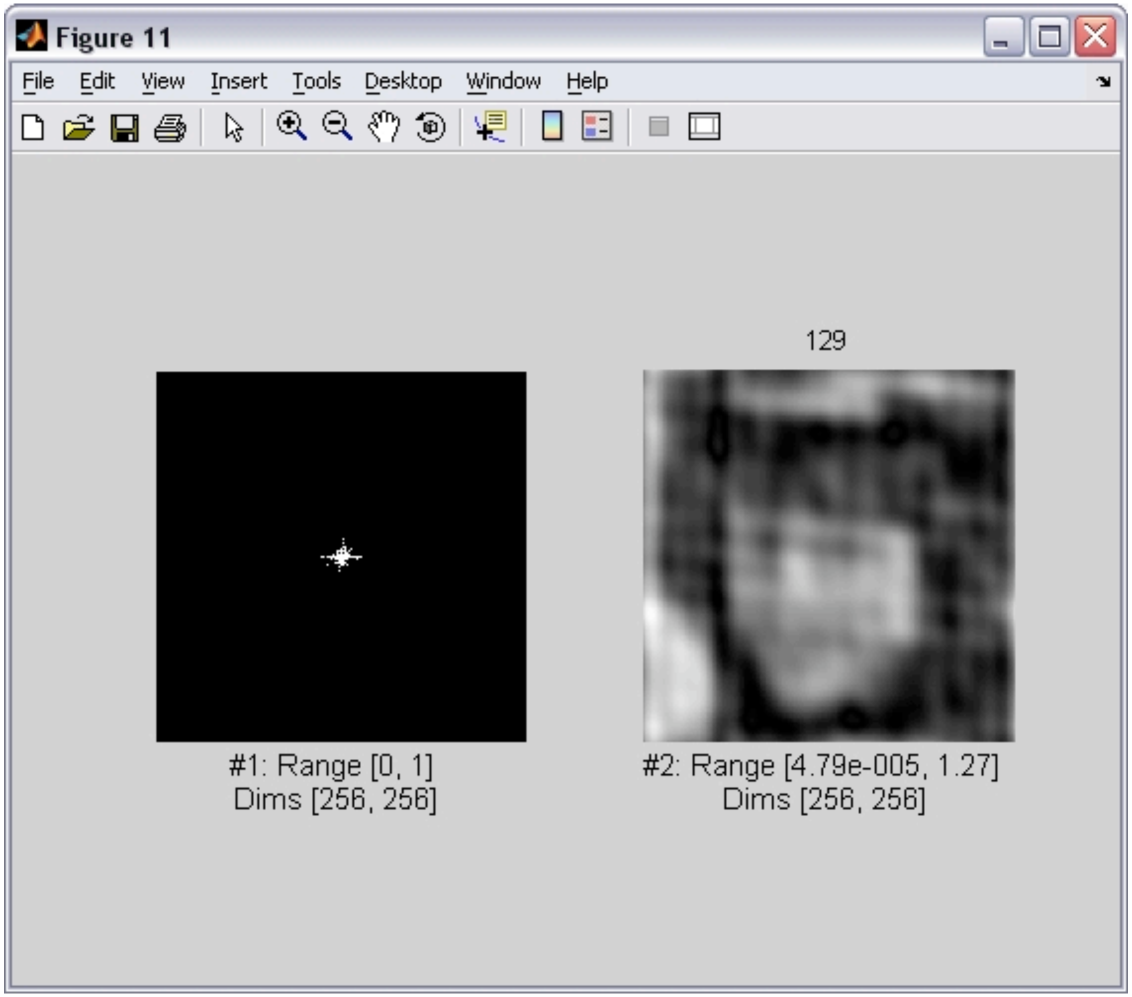
Figure 11
File Edit View Insert Tools Desktop Window Help
129
#1: Range [0, 1]
Dims [256, 256]
#2: Range [4.79e-005, 1.27]
Dims [256, 256]

257

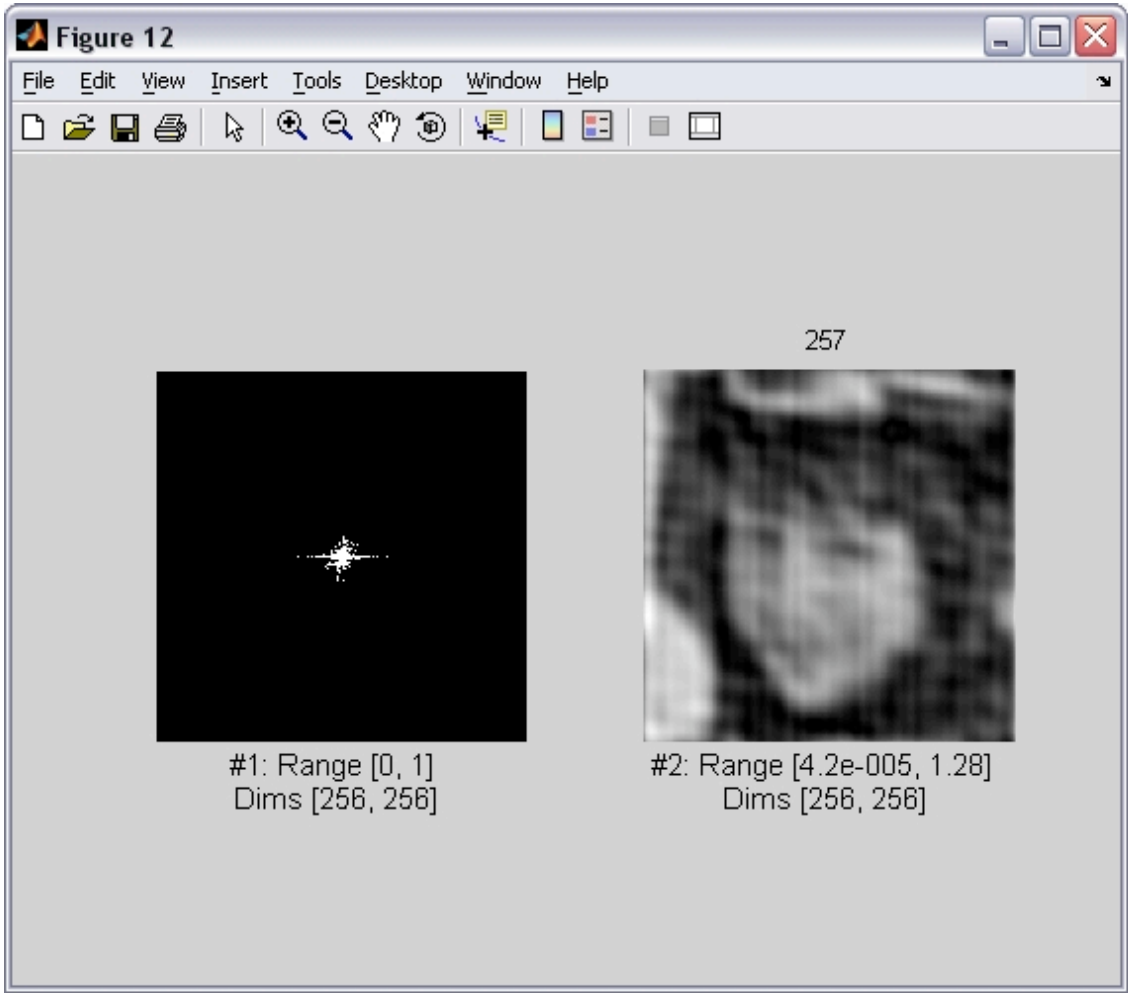
Figure 12
File Edit View Insert Tools Desktop Window Help
257
#1: Range [0, 1]
Dims [256, 256]
#2: Range [4.2e-005, 1.28]
Dims [256, 256]

513

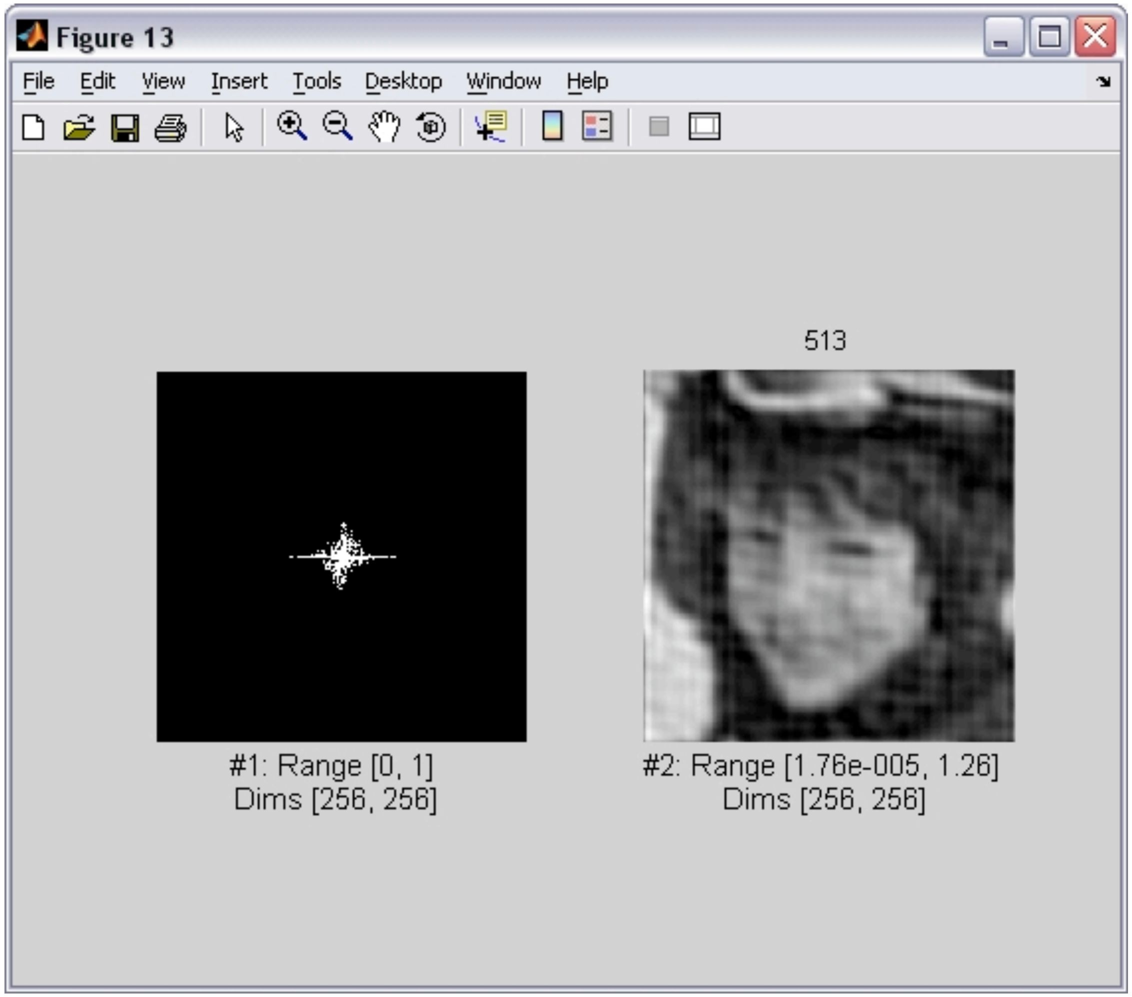
Figure 13
File Edit View Insert Tools Desktop Window Help
513
#1: Range [0, 1]
Dims [256, 256]
#2: Range [1.76e-005, 1.26]
Dims [256, 256]

1025

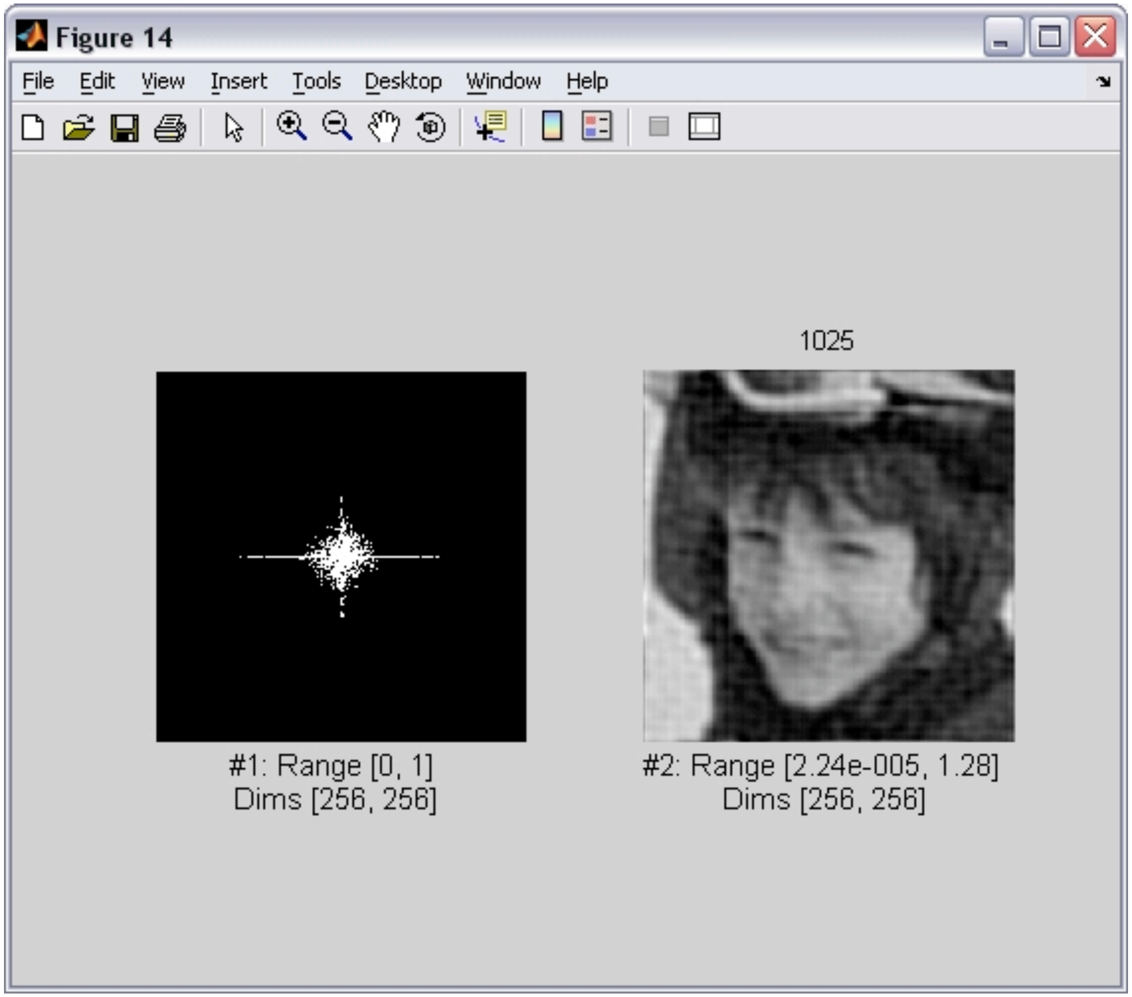
Figure 14
File Edit View Insert Tools Desktop Window Help
1025
#1: Range [0, 1]
Dims [256, 256]
#2: Range [2.24e-005, 1.28]
Dims [256, 256]

2049

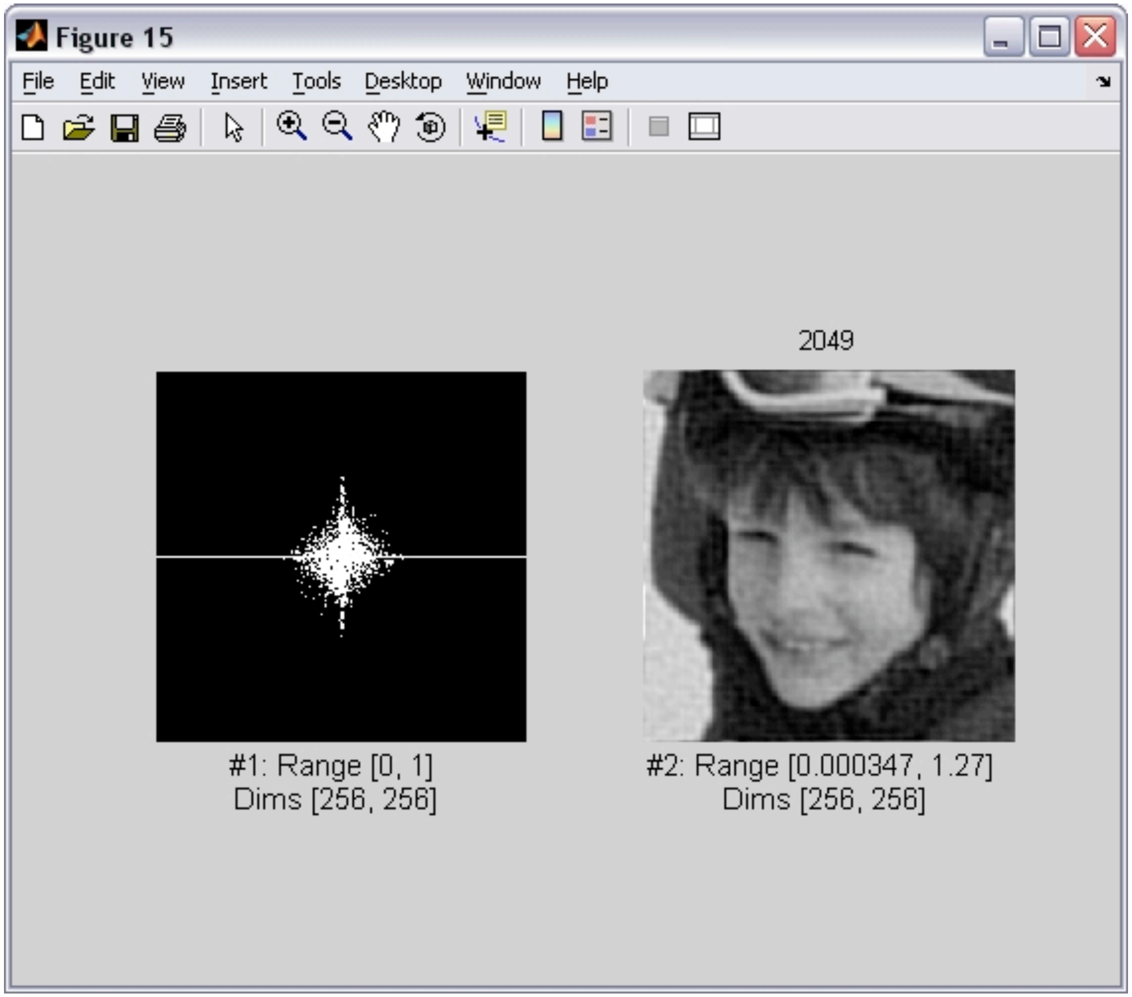
Figure 15
File
Edit
View
Insert
Tools
Desktop
Window
Help
2049
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.000347, 1.27]
Dims [256, 256]

4097

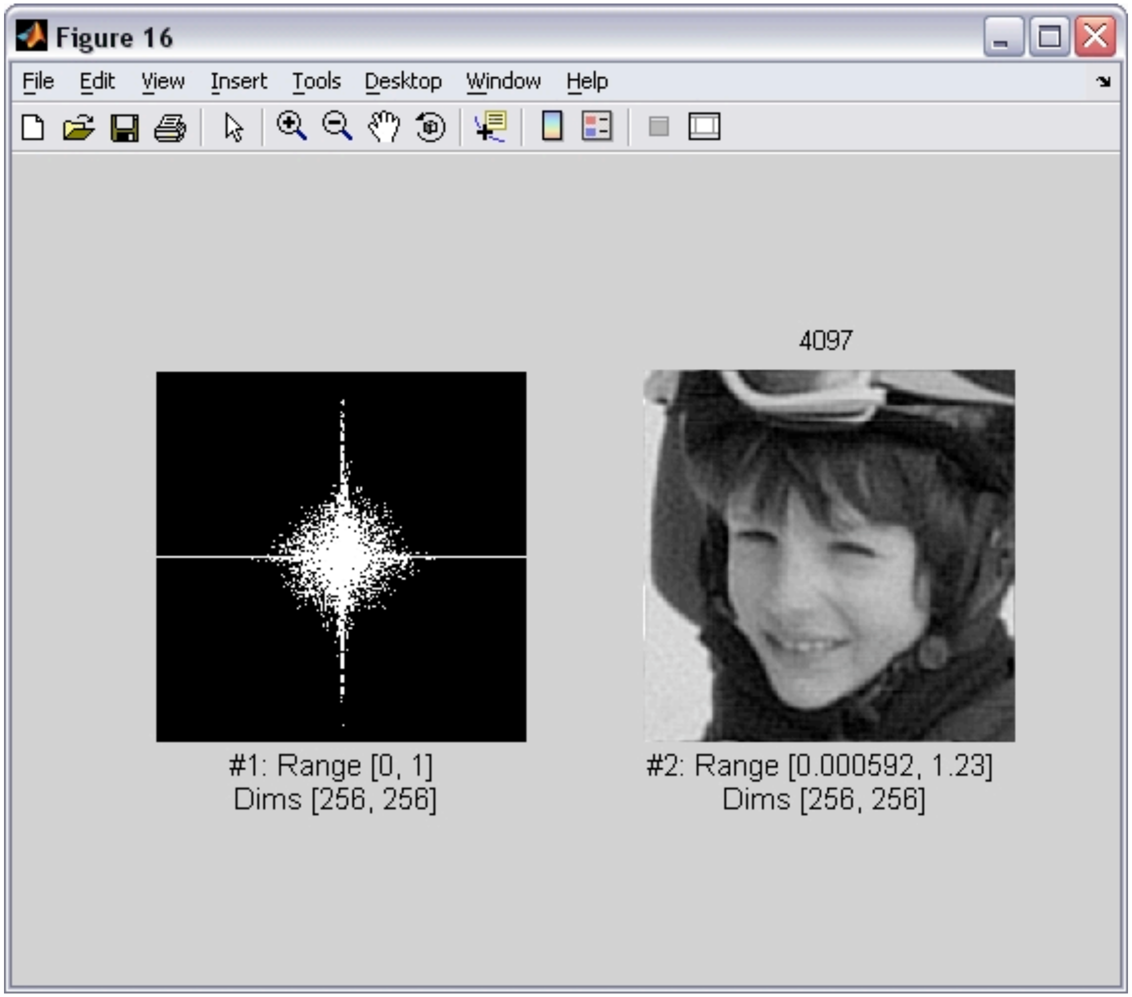
Figure 16
File
Edit
View
Insert
Tools
Desktop
Window
Help
4097
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.000592, 1.23]
Dims [256, 256]

8193

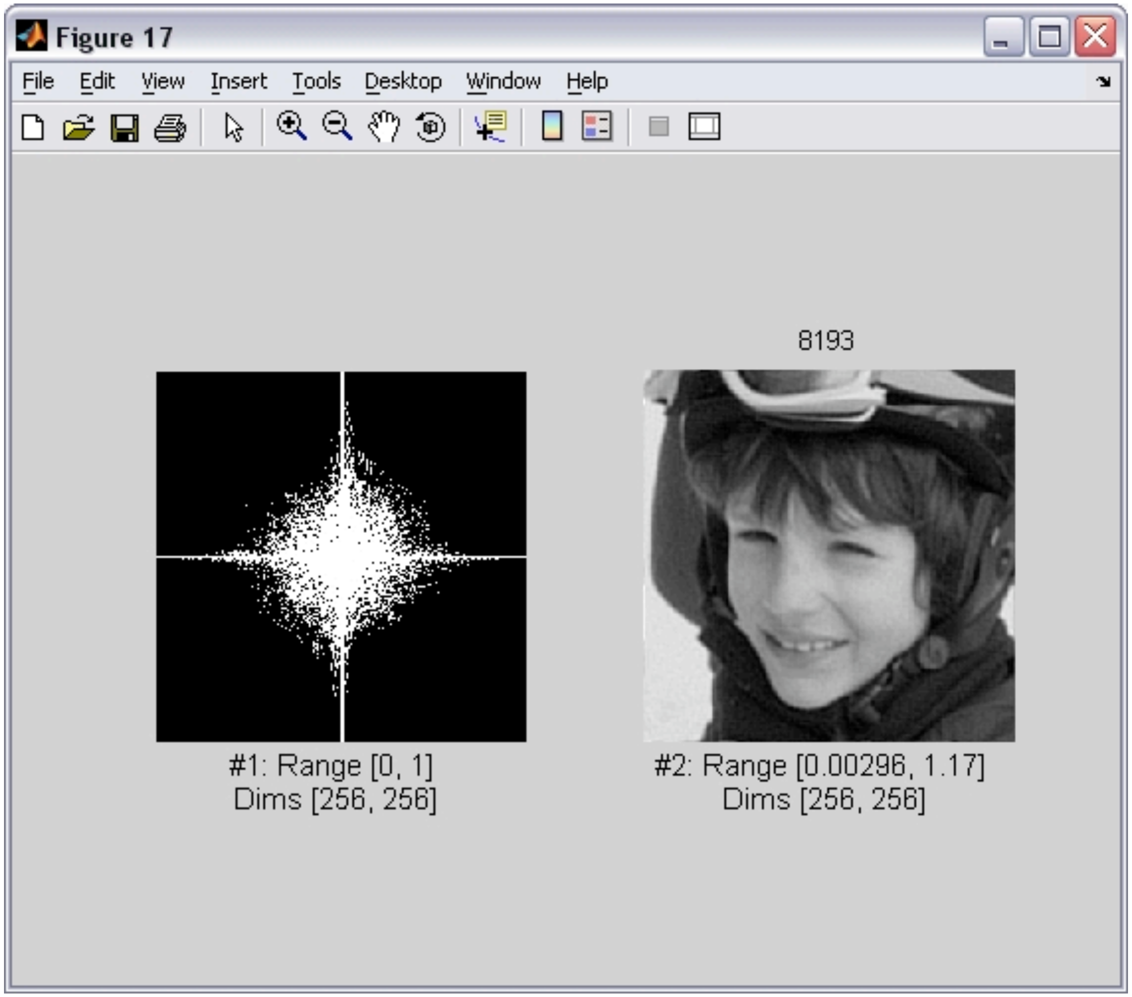
Figure 17
File Edit View Insert Tools Desktop Window Help
8193
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.00296, 1.17]
Dims [256, 256]

16385

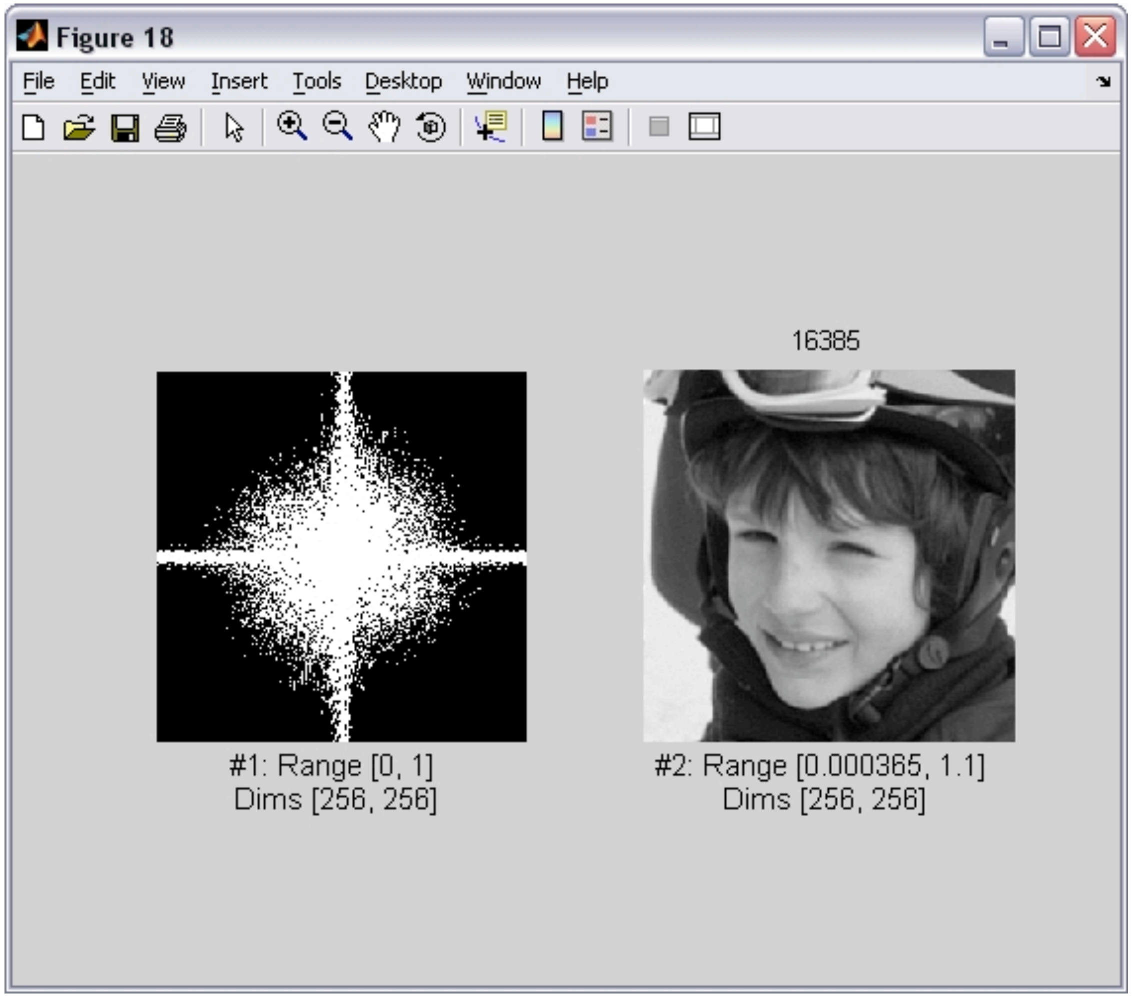
Figure 18
File
Edit
View
Insert
Tools
Desktop
Window
Help
16385
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.000365, 1.1]
Dims [256, 256]

32769

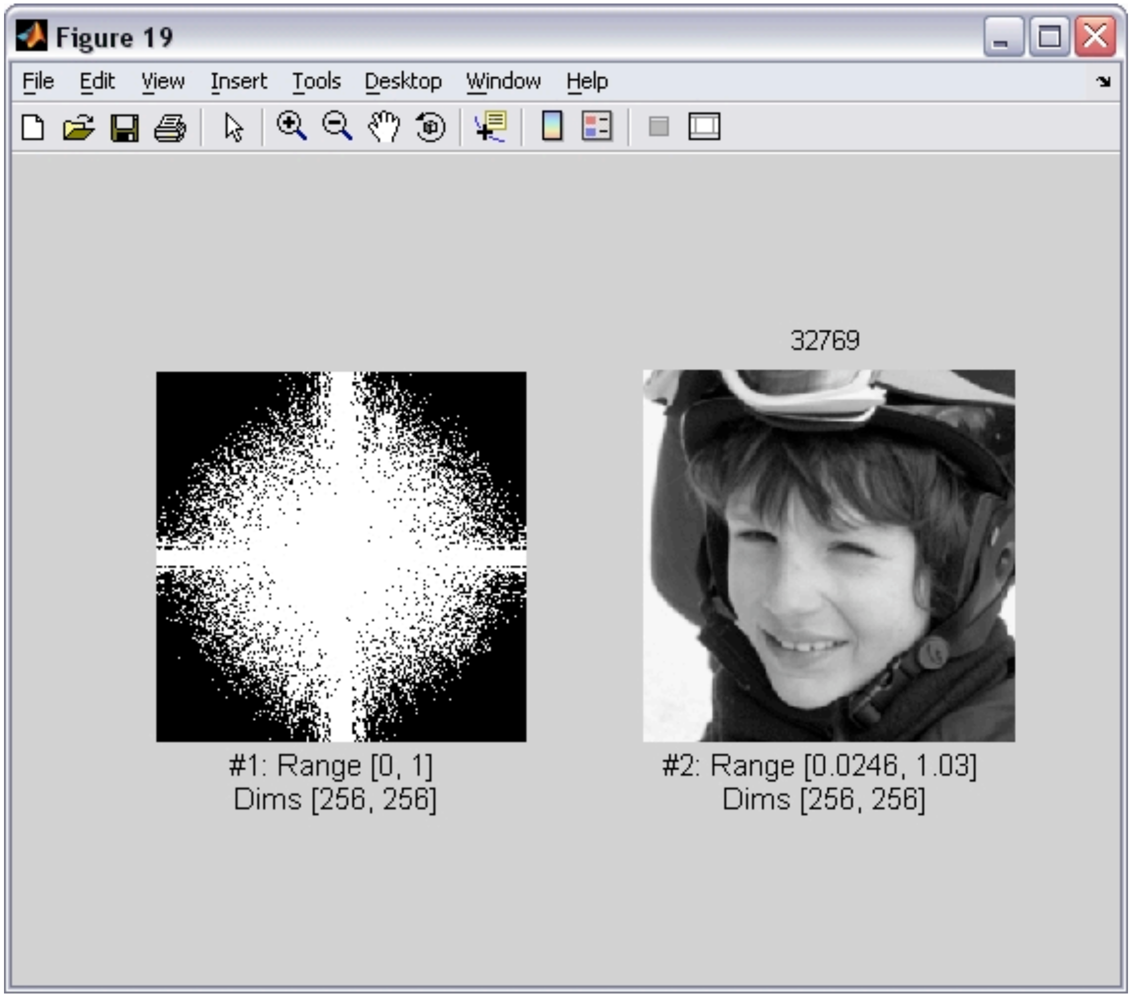
Figure 19
File Edit View Insert Tools Desktop Window Help
32769
#1: Range [0, 1]
Dims [256, 256]
#2: Range [0.0246, 1.03]
Dims [256, 256]

65536

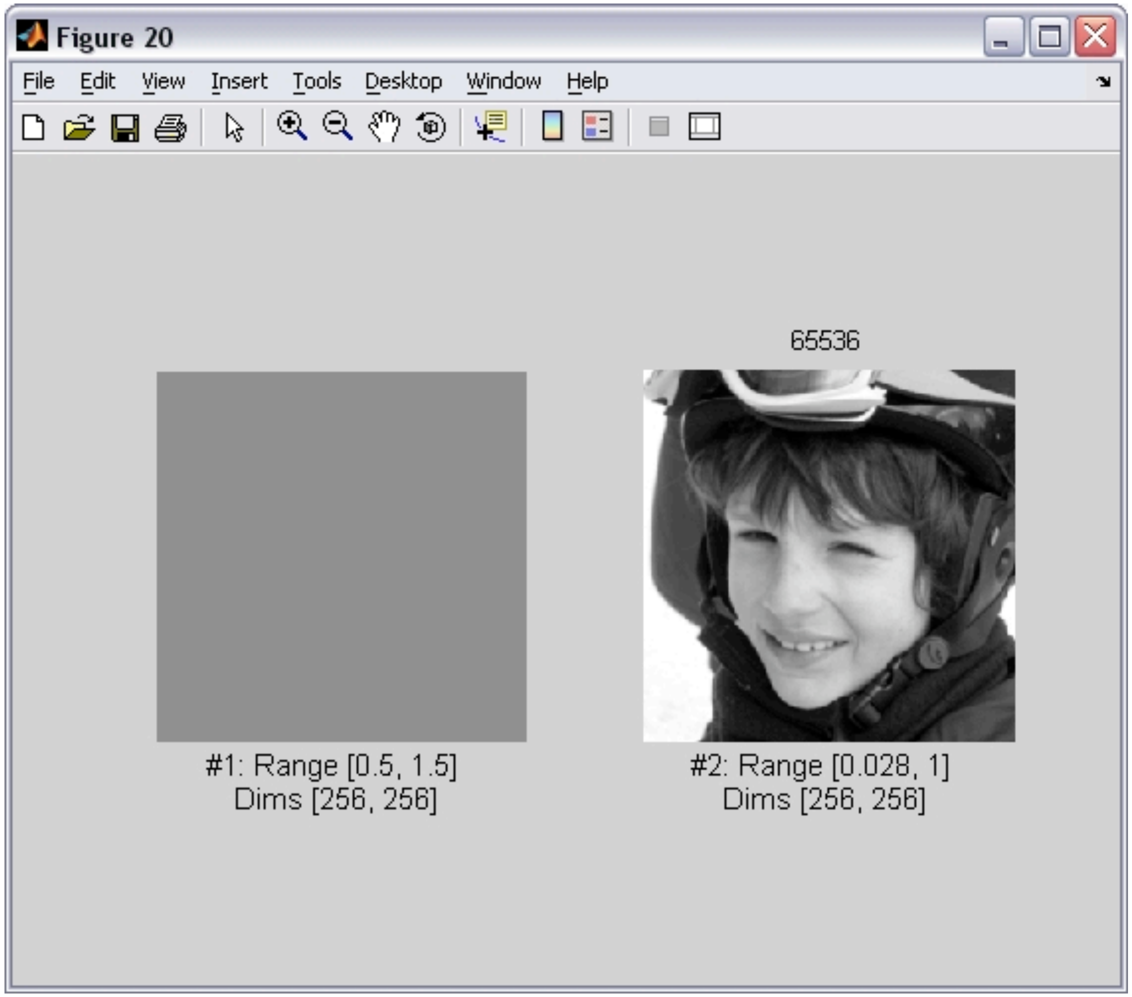
Figure 20
File
Edit
View
Insert
Tools
Desktop
Window
Help
65536
#1: Range [0.5, 1.5]
Dims [256, 256]
#2: Range [0.028, 1]
Dims [256, 256]

Fourier transform magnitude

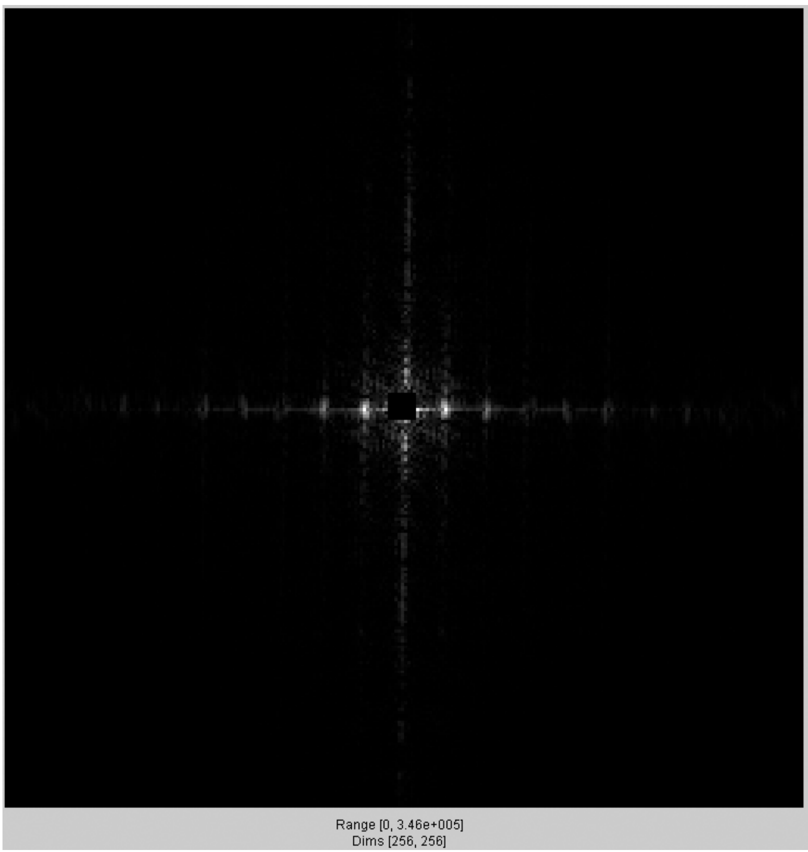

Masking out the fundamental and harmonics from periodic pillars

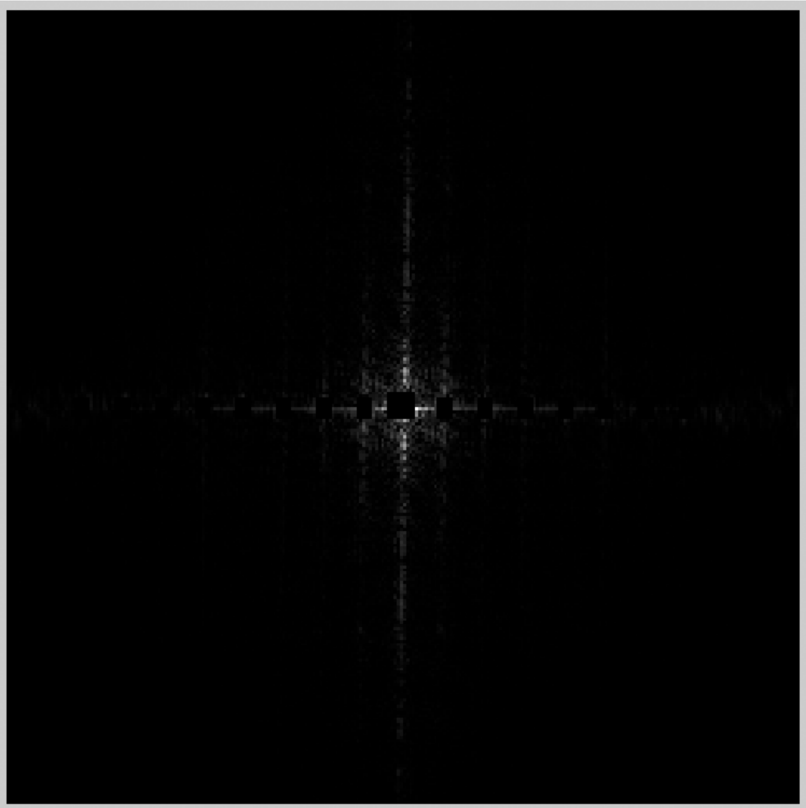

Range [0, 3.29e+005]
Dims [256, 256]

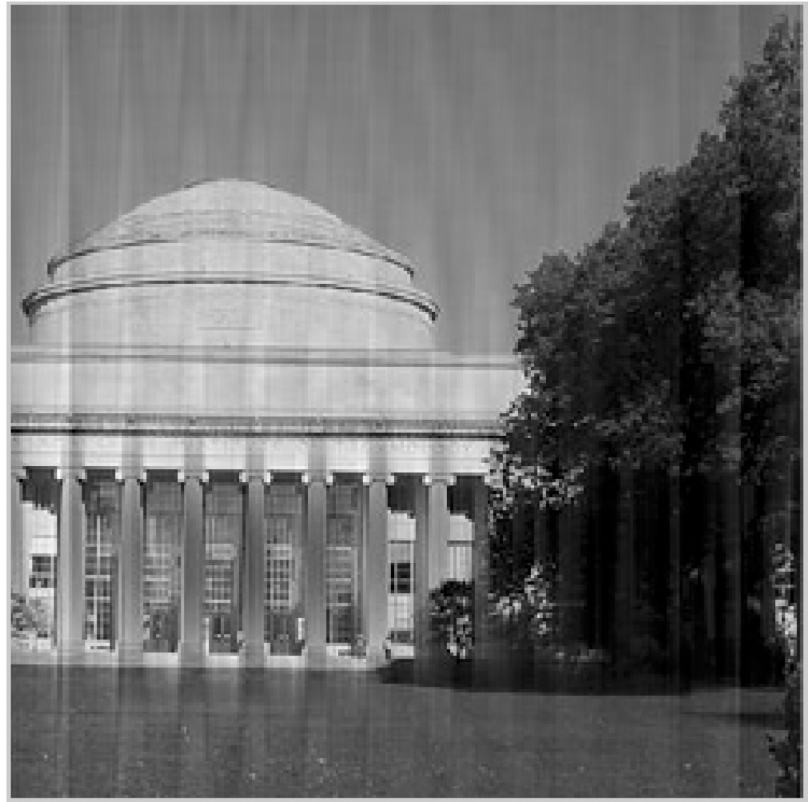

Range [0.000551, 297]
Dims [256, 256]

Fourier Transform

- Fourier transform of a real function is complex
  - difficult to plot, visualize
  - instead, we can think of the phase and magnitude of the transform
- Phase is the phase of the complex transform
- Magnitude is the magnitude of the complex transform

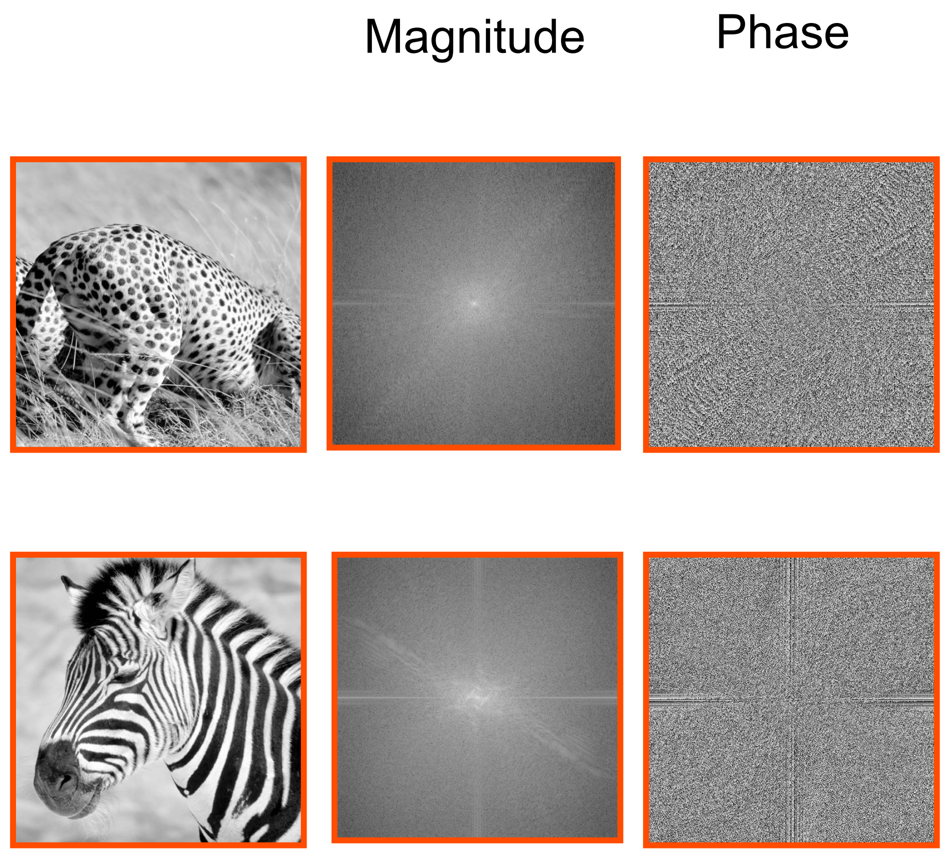

Phase and Magnitude

- Curious fact
  - all natural images have about the same magnitude transform
  - hence, phase seems to matter, but magnitude largely doesn't

- Demonstration
  - Take two pictures, swap the phase transforms, compute the inverse - what does the result look like?

This is the magnitude transform of the cheetah pic

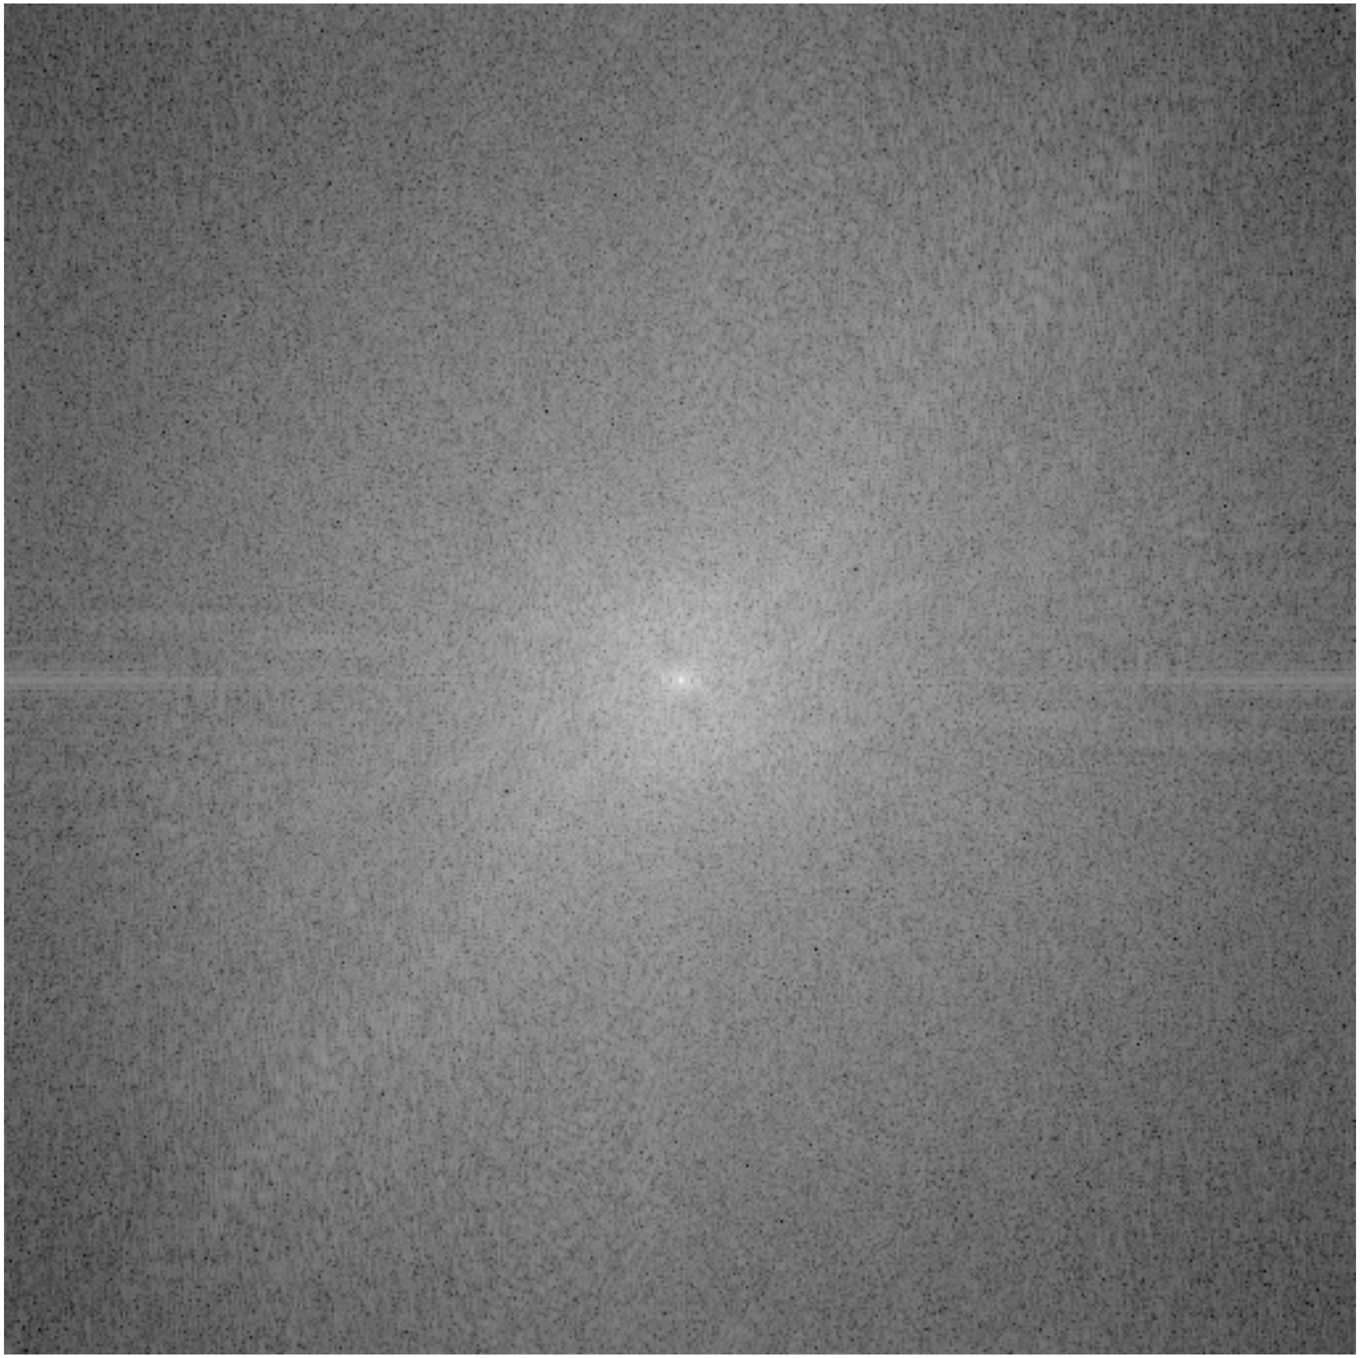

This is the phase transform of the cheetah pic

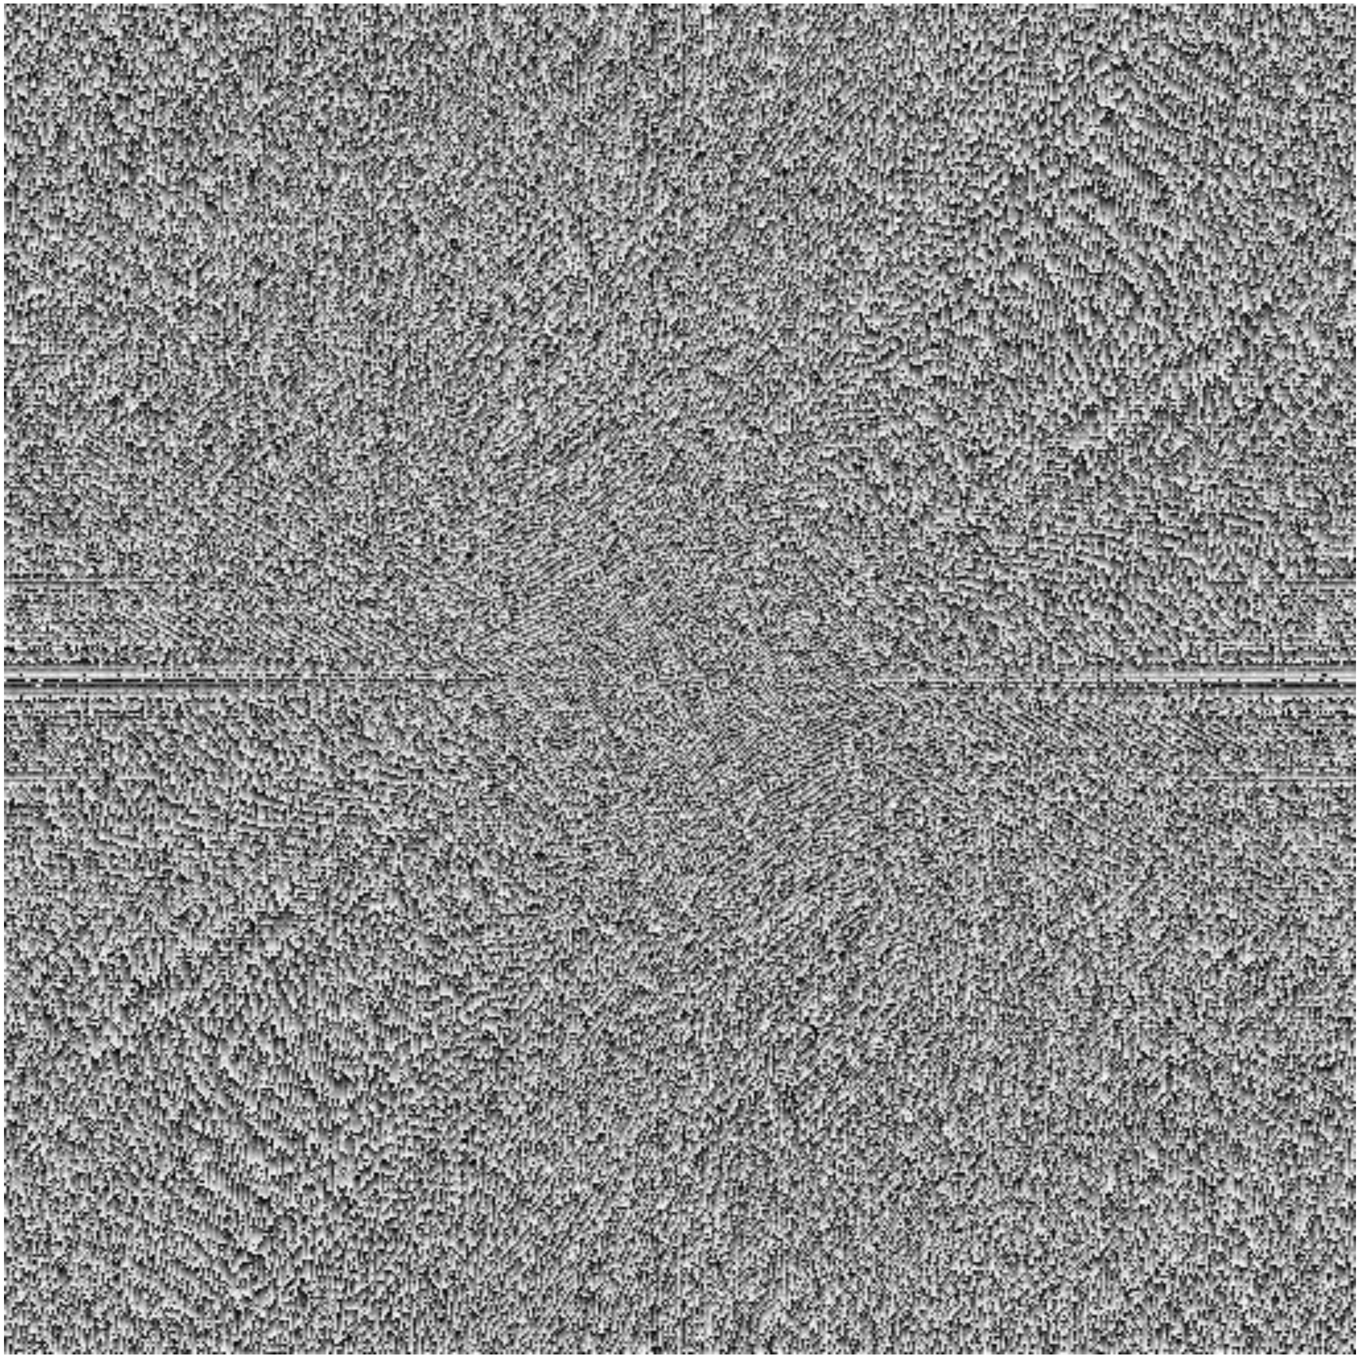

This is the
magnitude
transform
of the
zebra pic

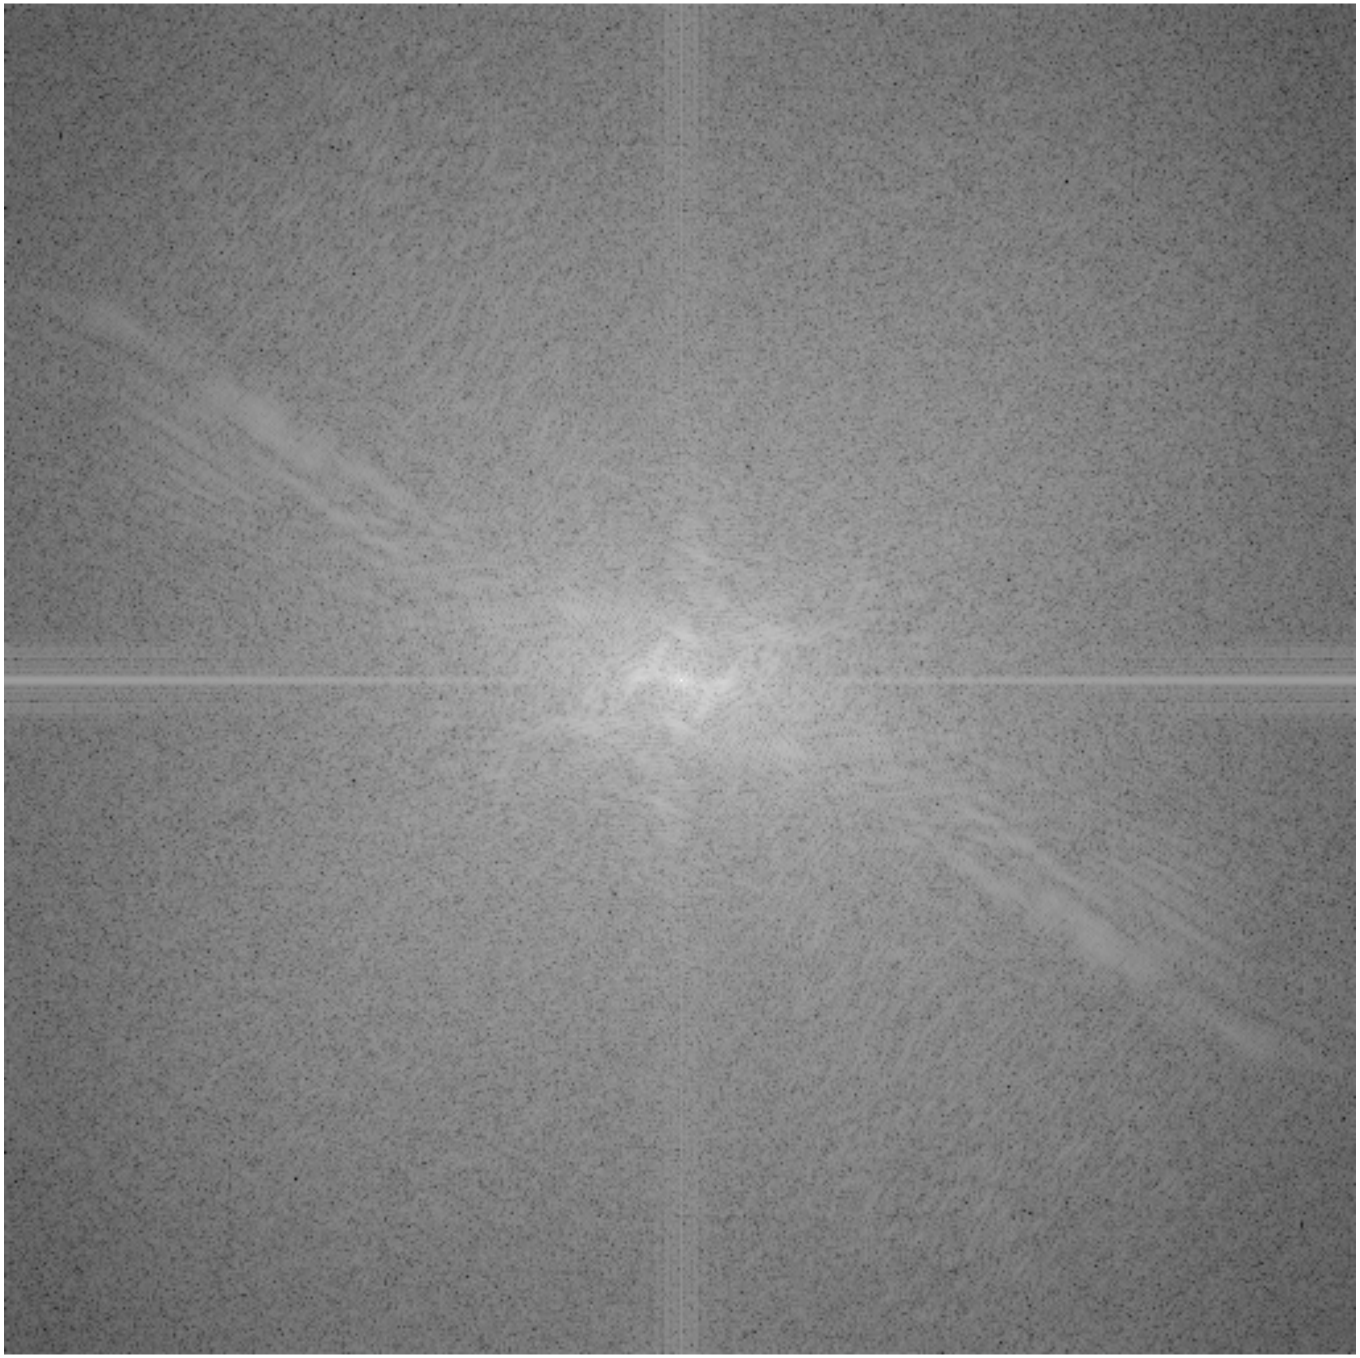

This is the
phase
transform
of the
zebra pic

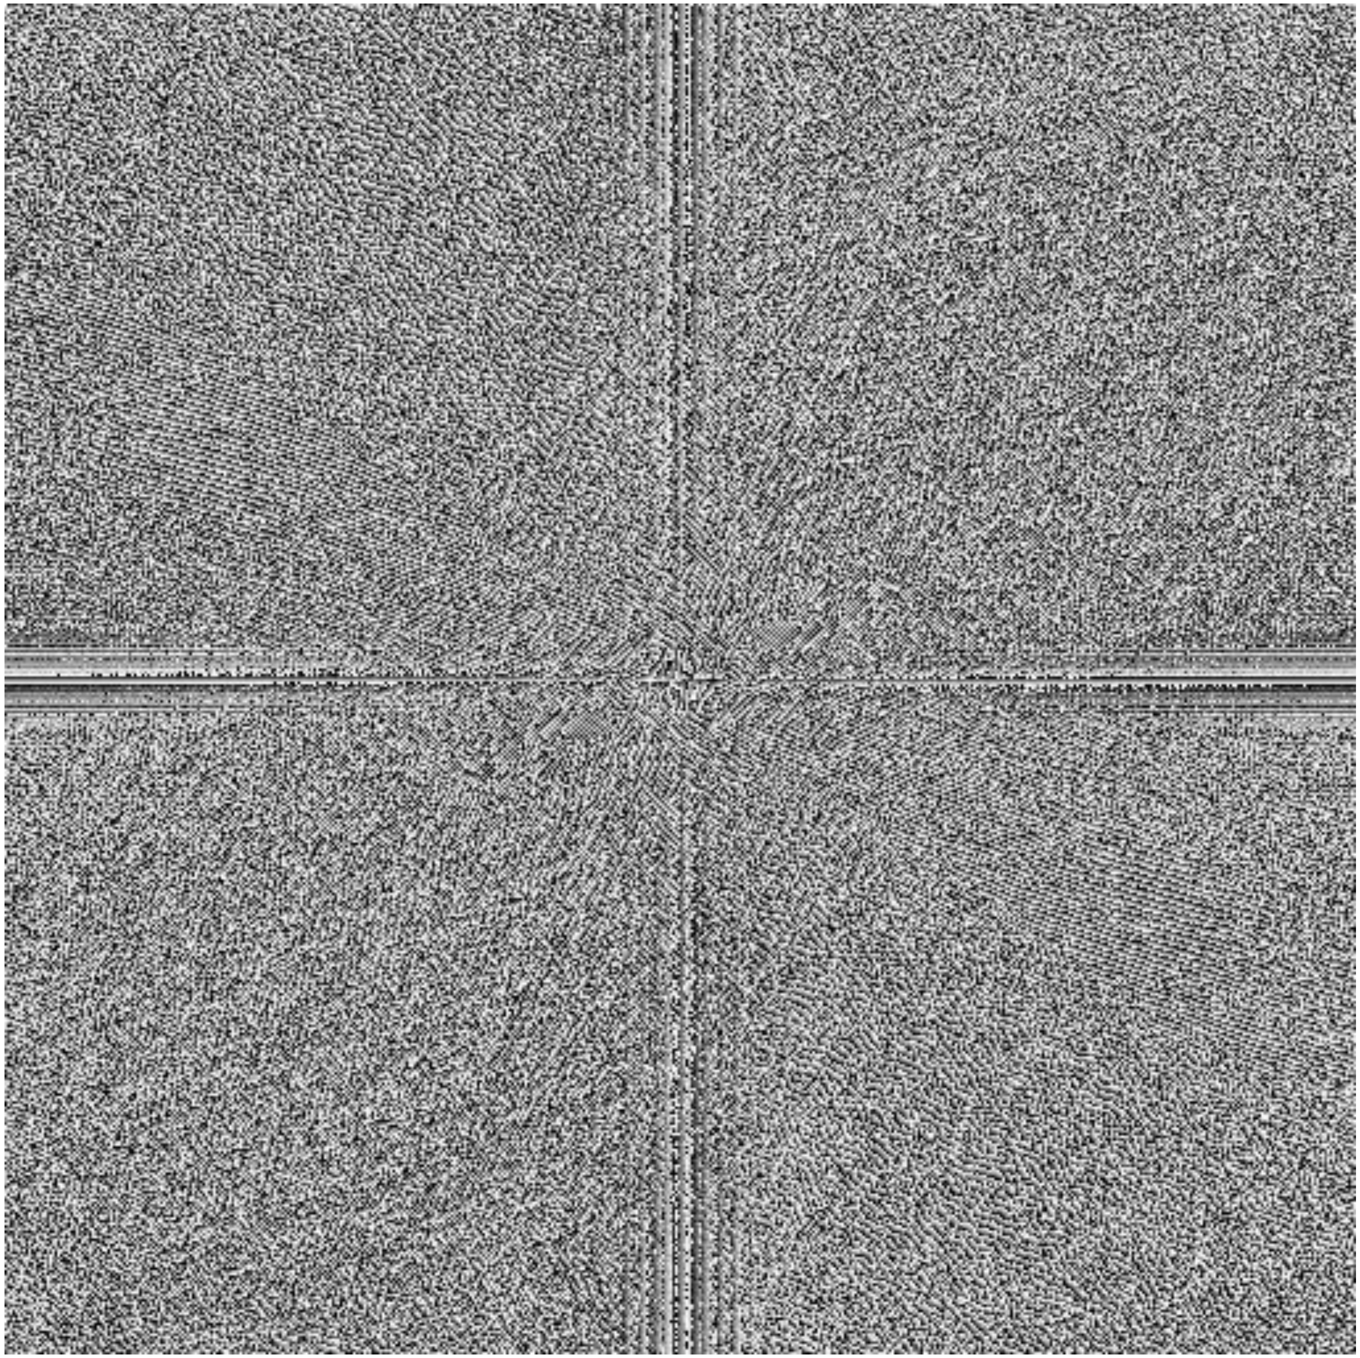

Reconstruction with zebra phase, cheetah magnitude

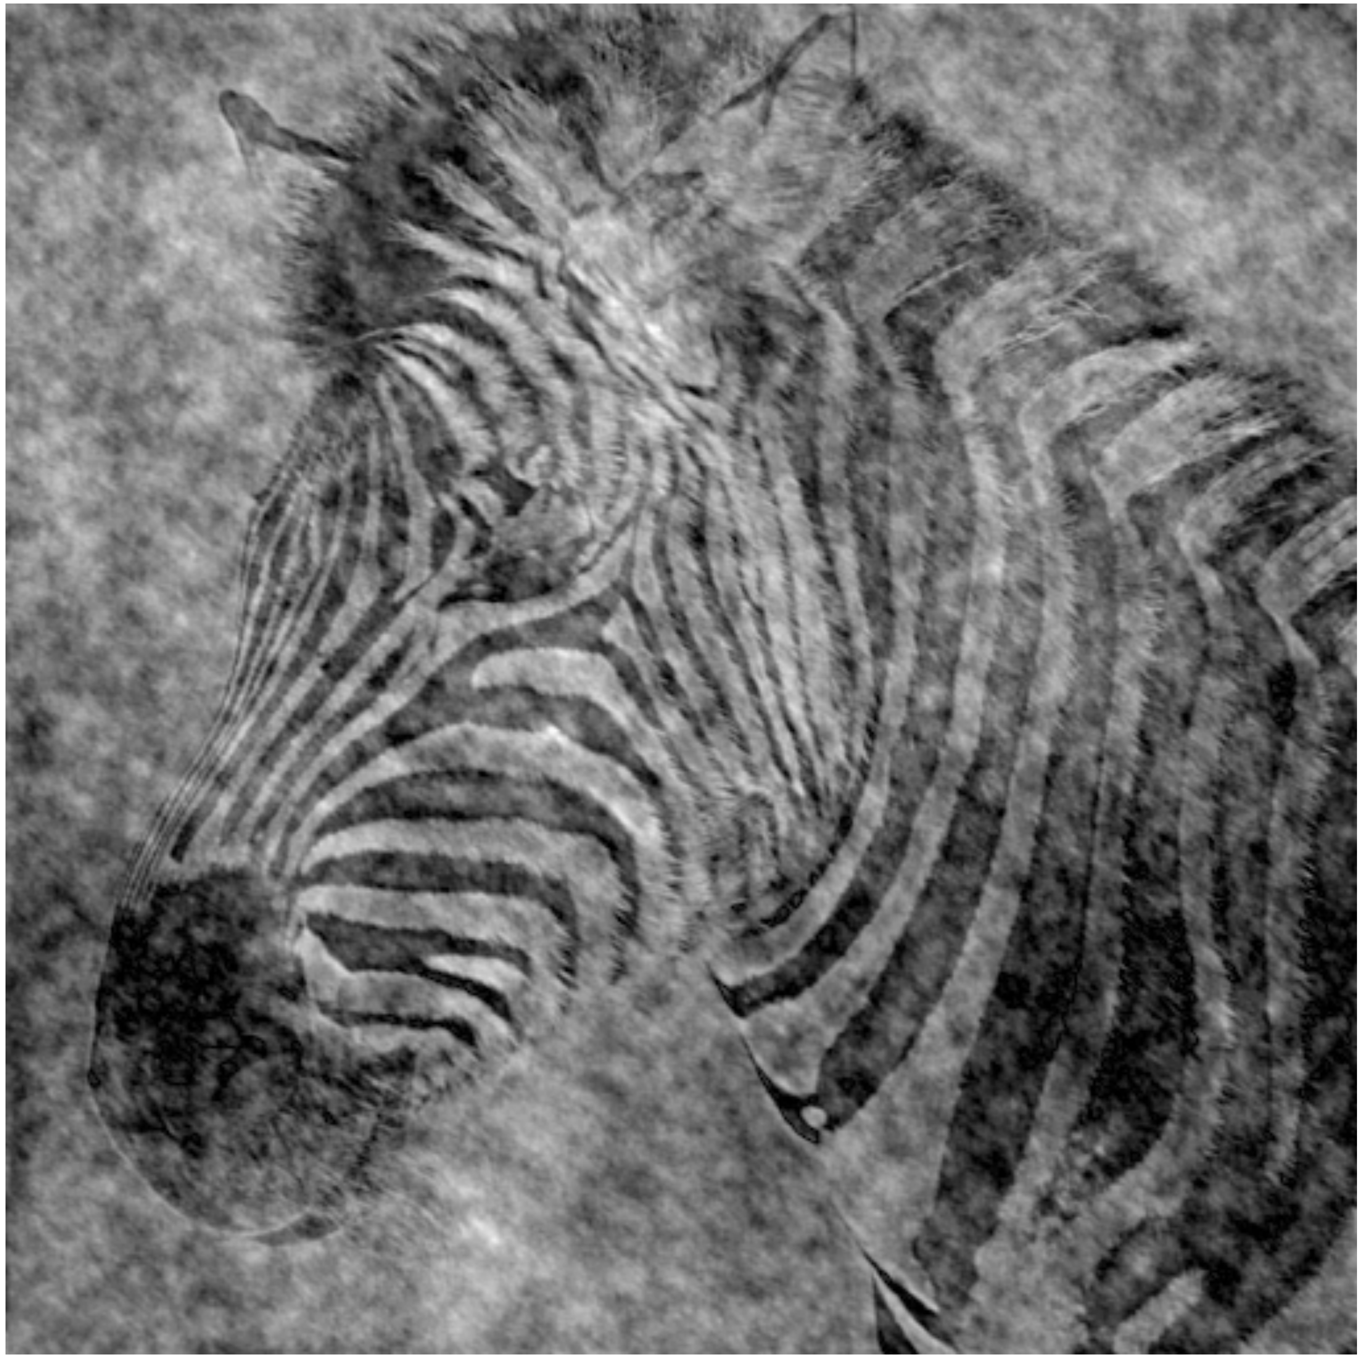

Reconstruction
with cheetah
phase, zebra
magnitude

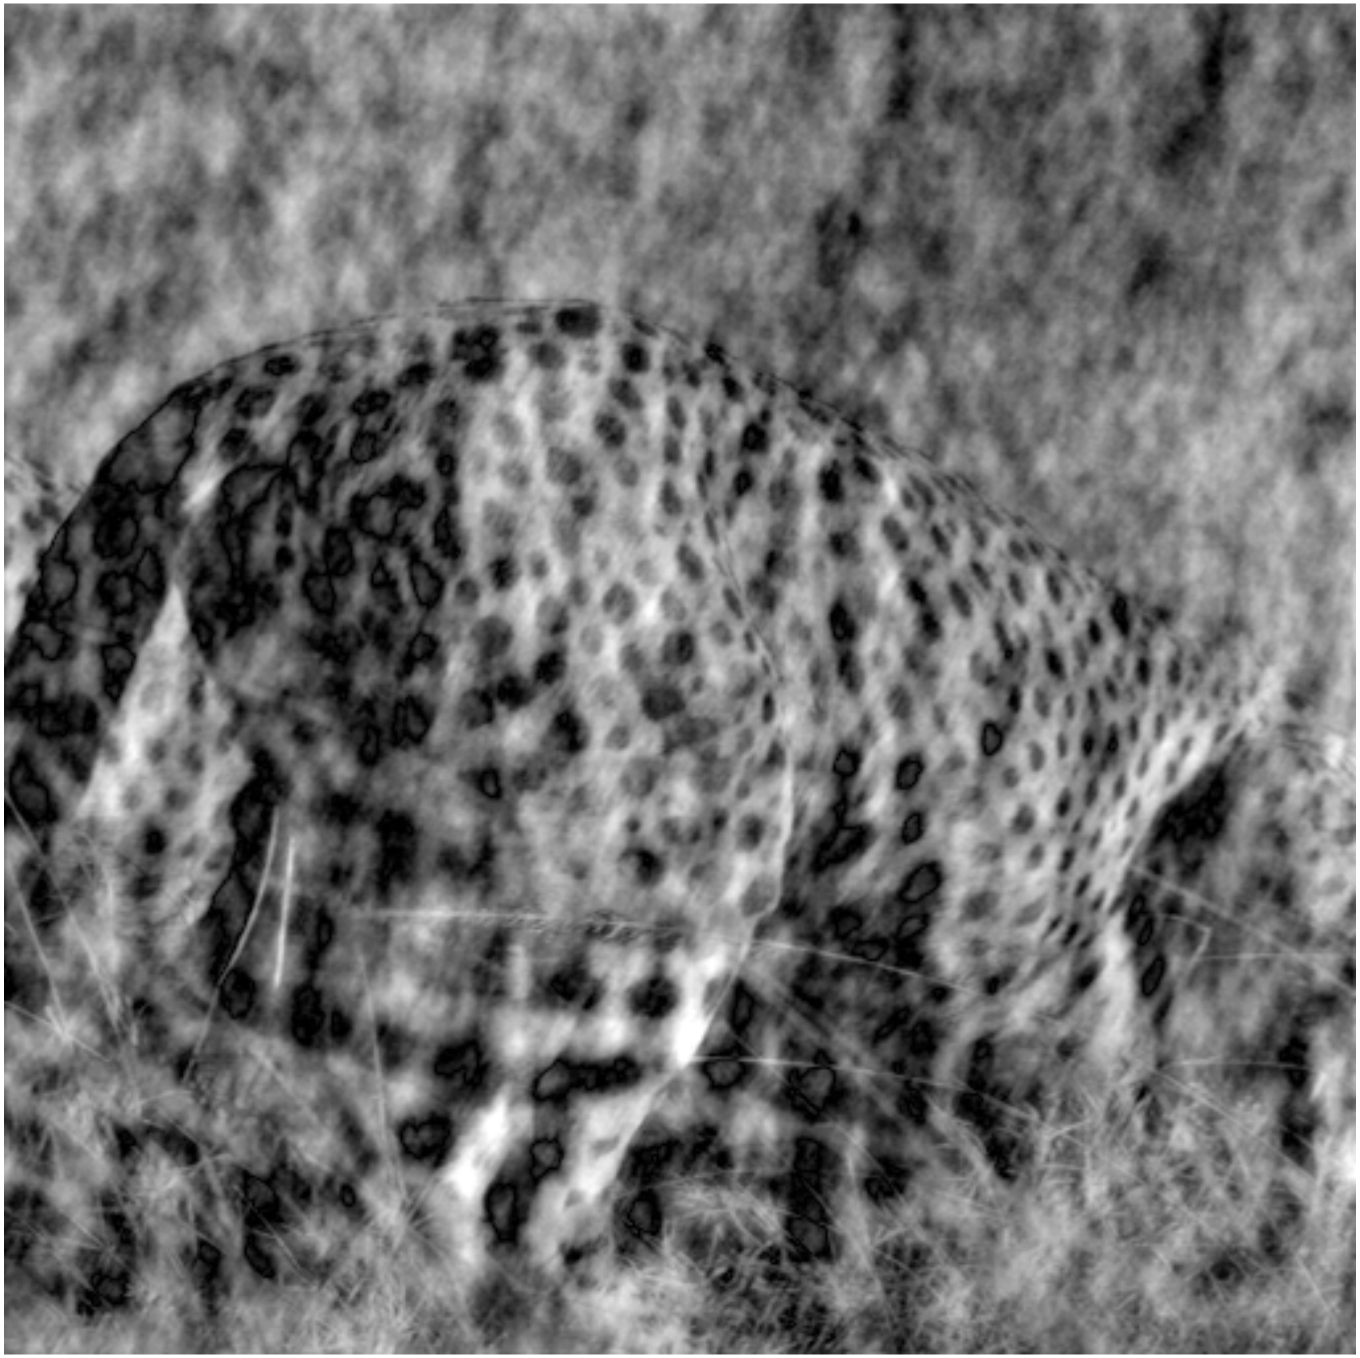

Phase and Magnitude

Image with cheetah phase
(and zebra magnitude)

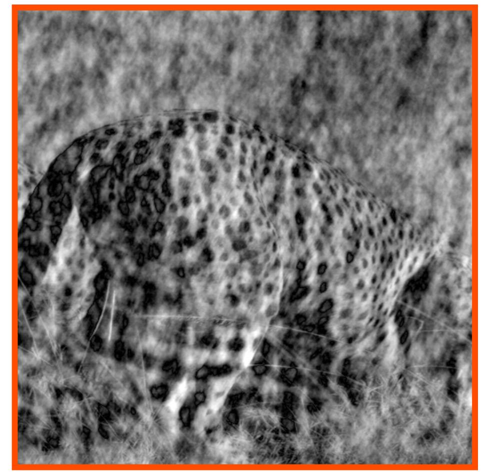

Image with zebra phase
(and cheetah magnitude)

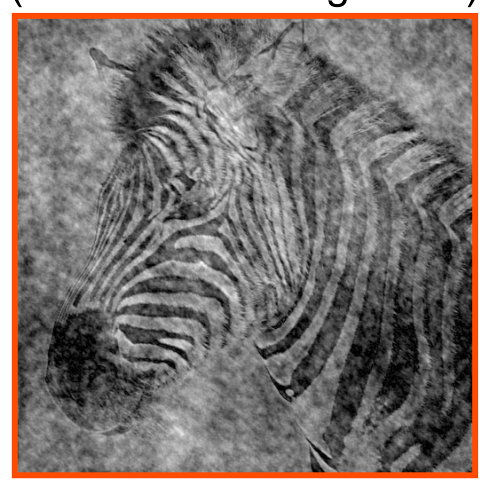

Randomizing the phase

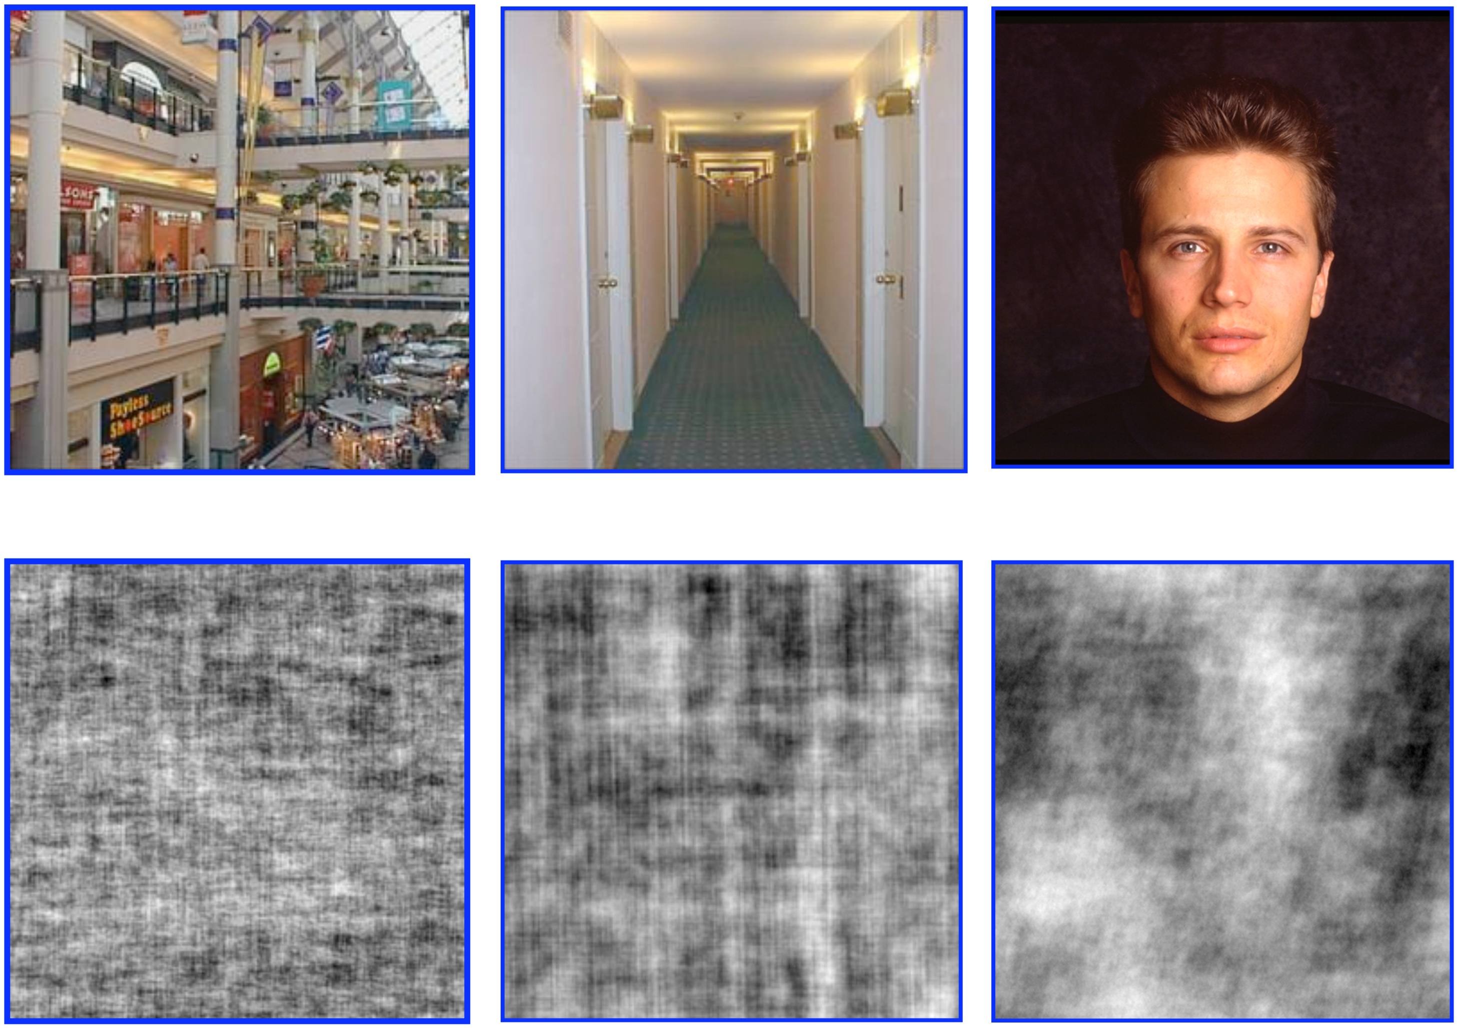

Fourier Characteristics of Natural Images

Power spectra
fall off as $1/f^2$

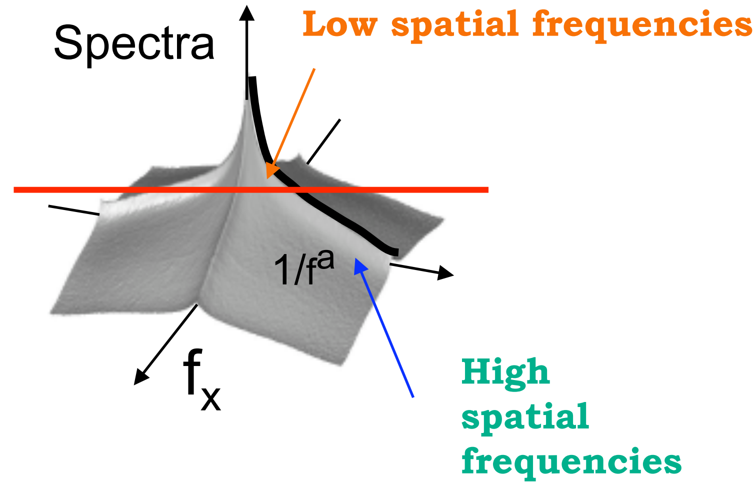

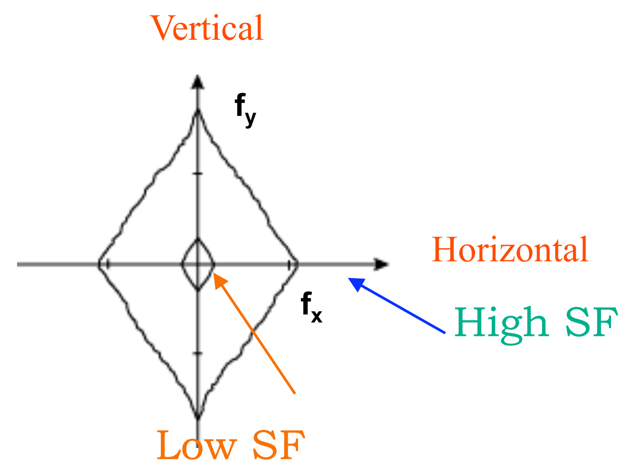

D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," J. Opt. Soc. Am. A **4**, 2379- (1987)

Torralba and Oliva, *Statistics of Natural Image Categories*. Network: Computation in Neural Systems 14 (2003) 391-412.

Power Spectrum of Images

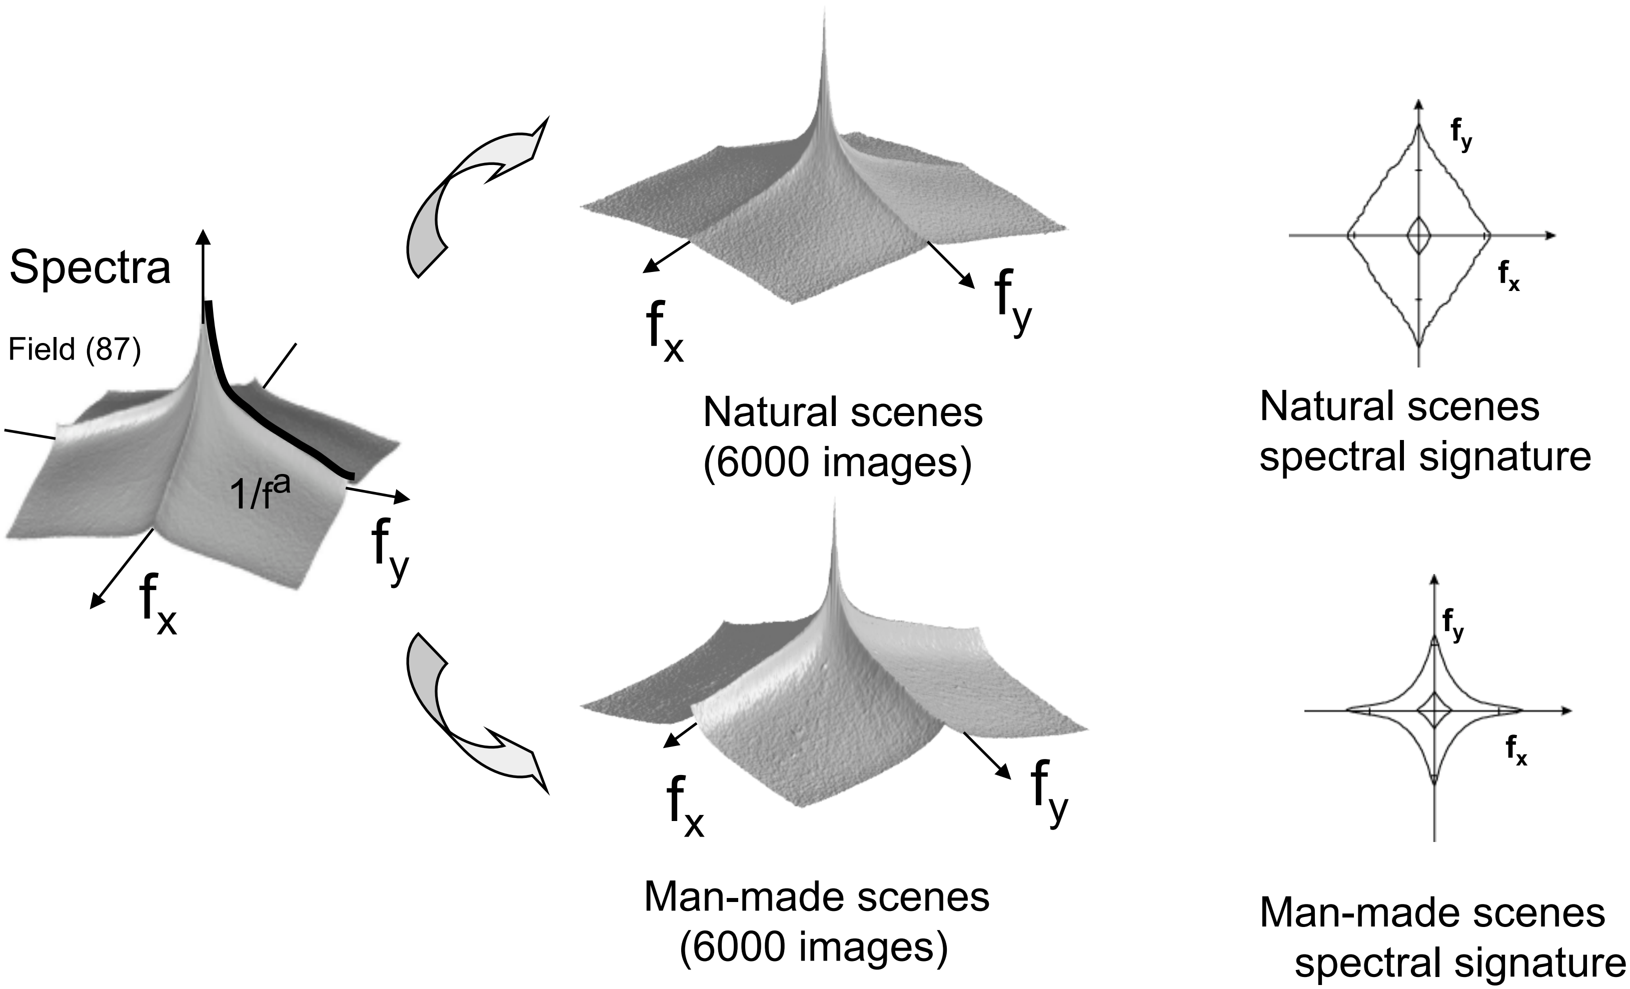

Torralba and Oliva, *Statistics of Natural Image Categories*. Network: Computation in Neural Systems 14 (2003) 391-412.

Edges

$[-1 \ 1]$

[-1 1]

g[m,n]

$\otimes$

[-1, 1]

h[m,n]

=

f[m,n]

$[-1\ 1]^T$

$\otimes$

$[-1, 1]^T$

h[m,n]

=

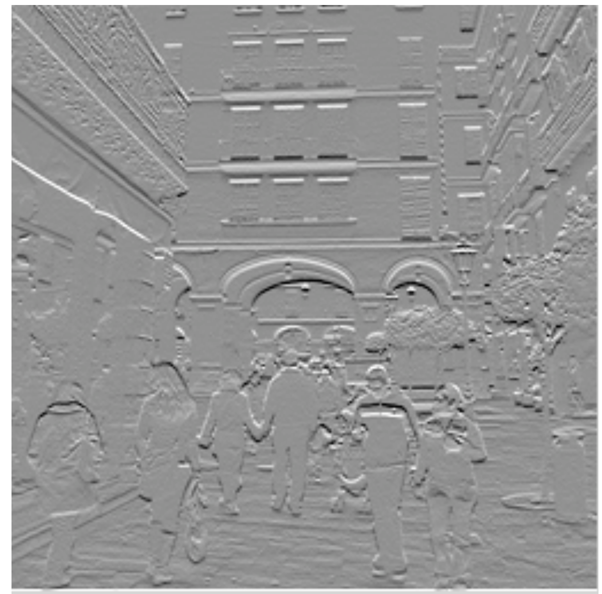

g[m,n]

f[m,n]

Observation: Sparse filter response

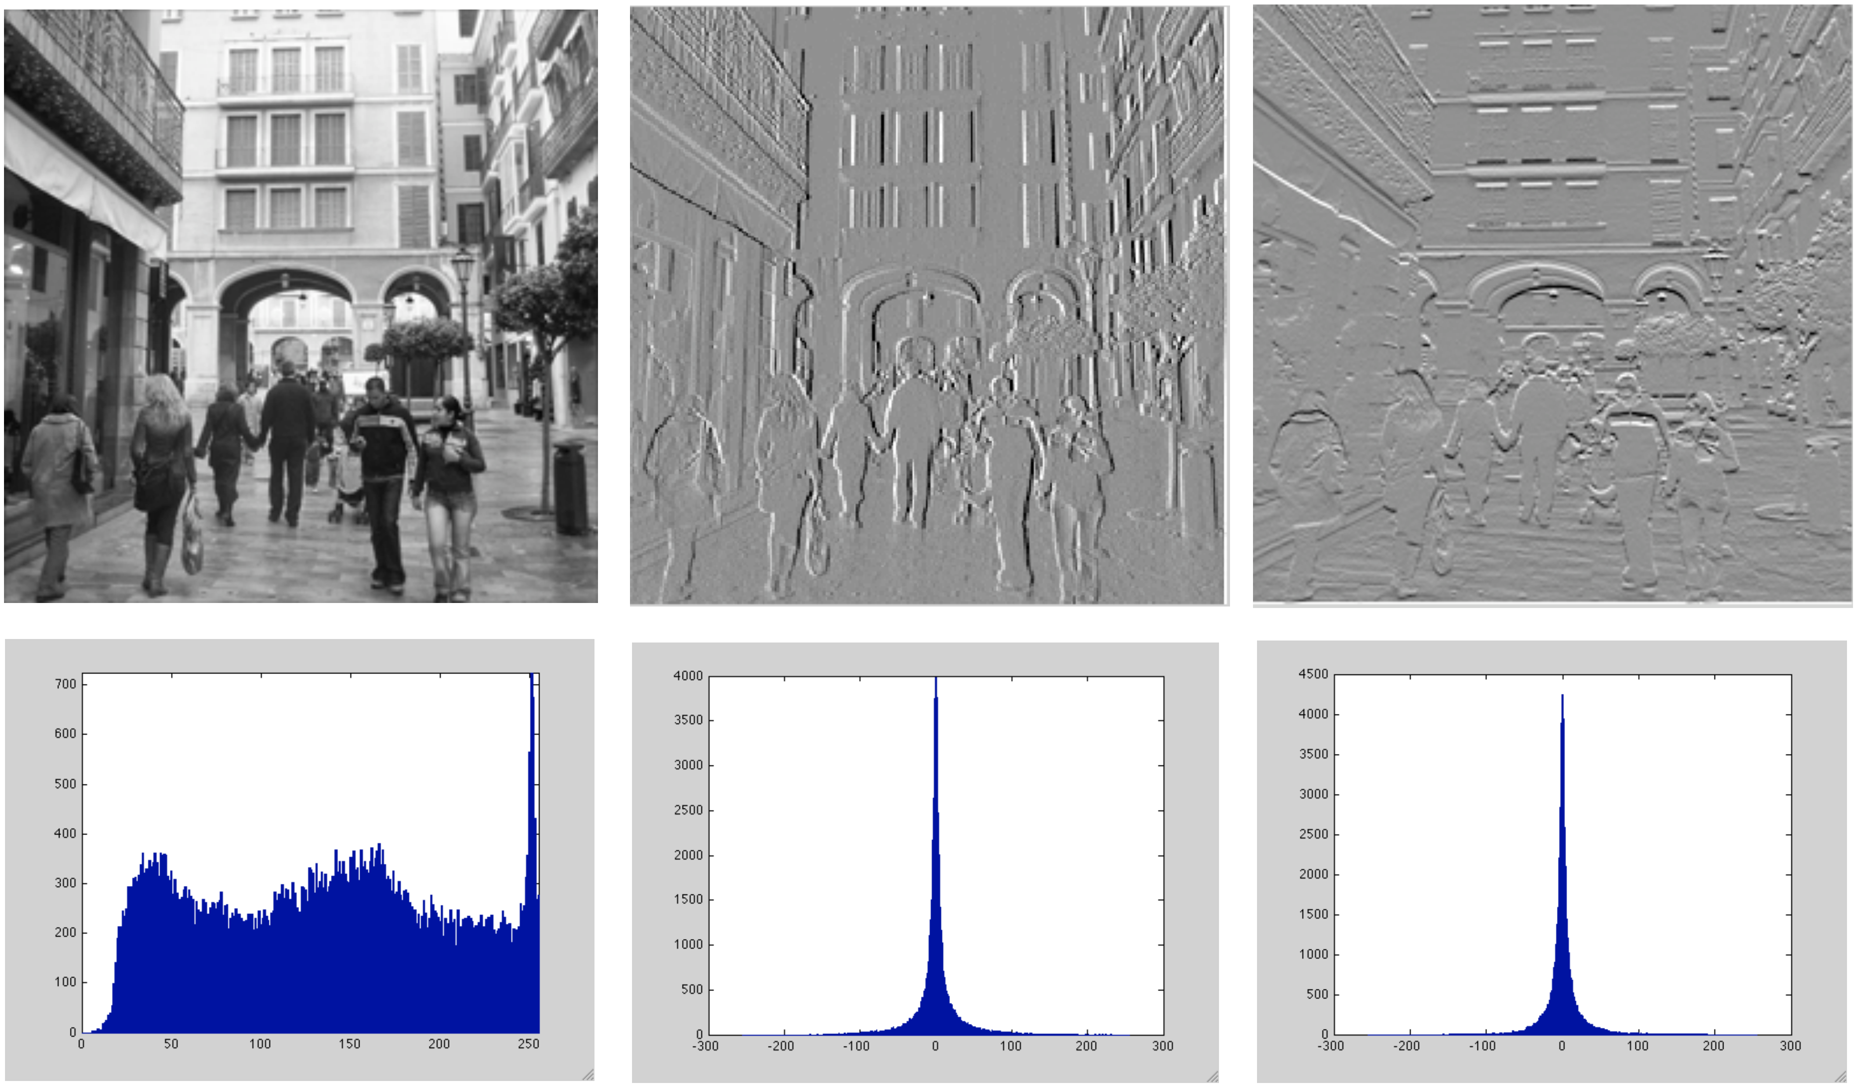

Back to the image

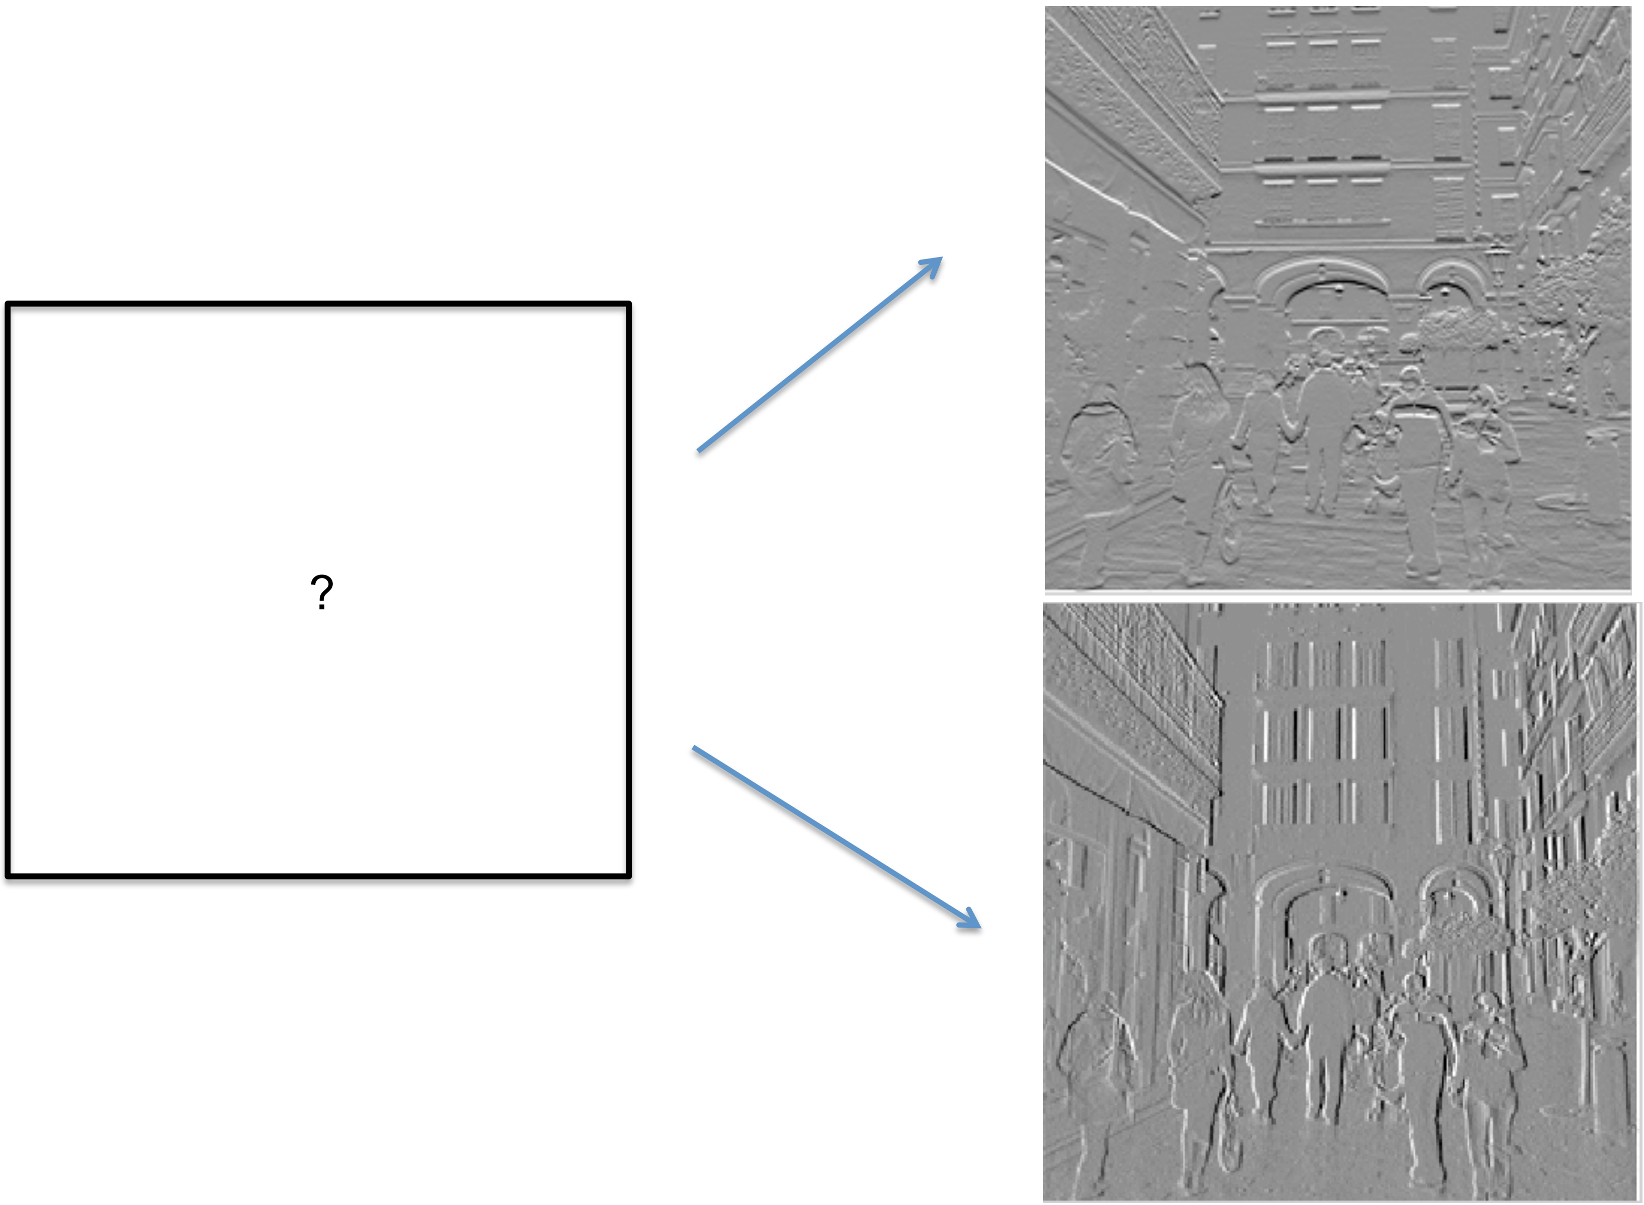

Reconstruction from derivatives

F = H G

$$
= \begin{bmatrix}
1 & -1 & & & & & & \\
 & 1 & -1 & & & & & \\
 & & 1 & -1 & & & & \\
 & & & 1 & -1 & & & \\
 & & & & 1 & -1 & & \\
 & & & & & 1 & -1 & \\
 & & & & & & 1 & -1 \\
 & & & & & & & 1
\end{bmatrix}
$$

If we have multiple filter outputs:

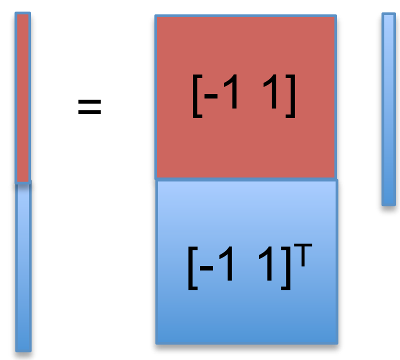

If the transformation H is not invertible, we can compute the pseudo-inverse:

$$\hat{G} = (H^T H)^{-1} H^T F$$

Reconstruction

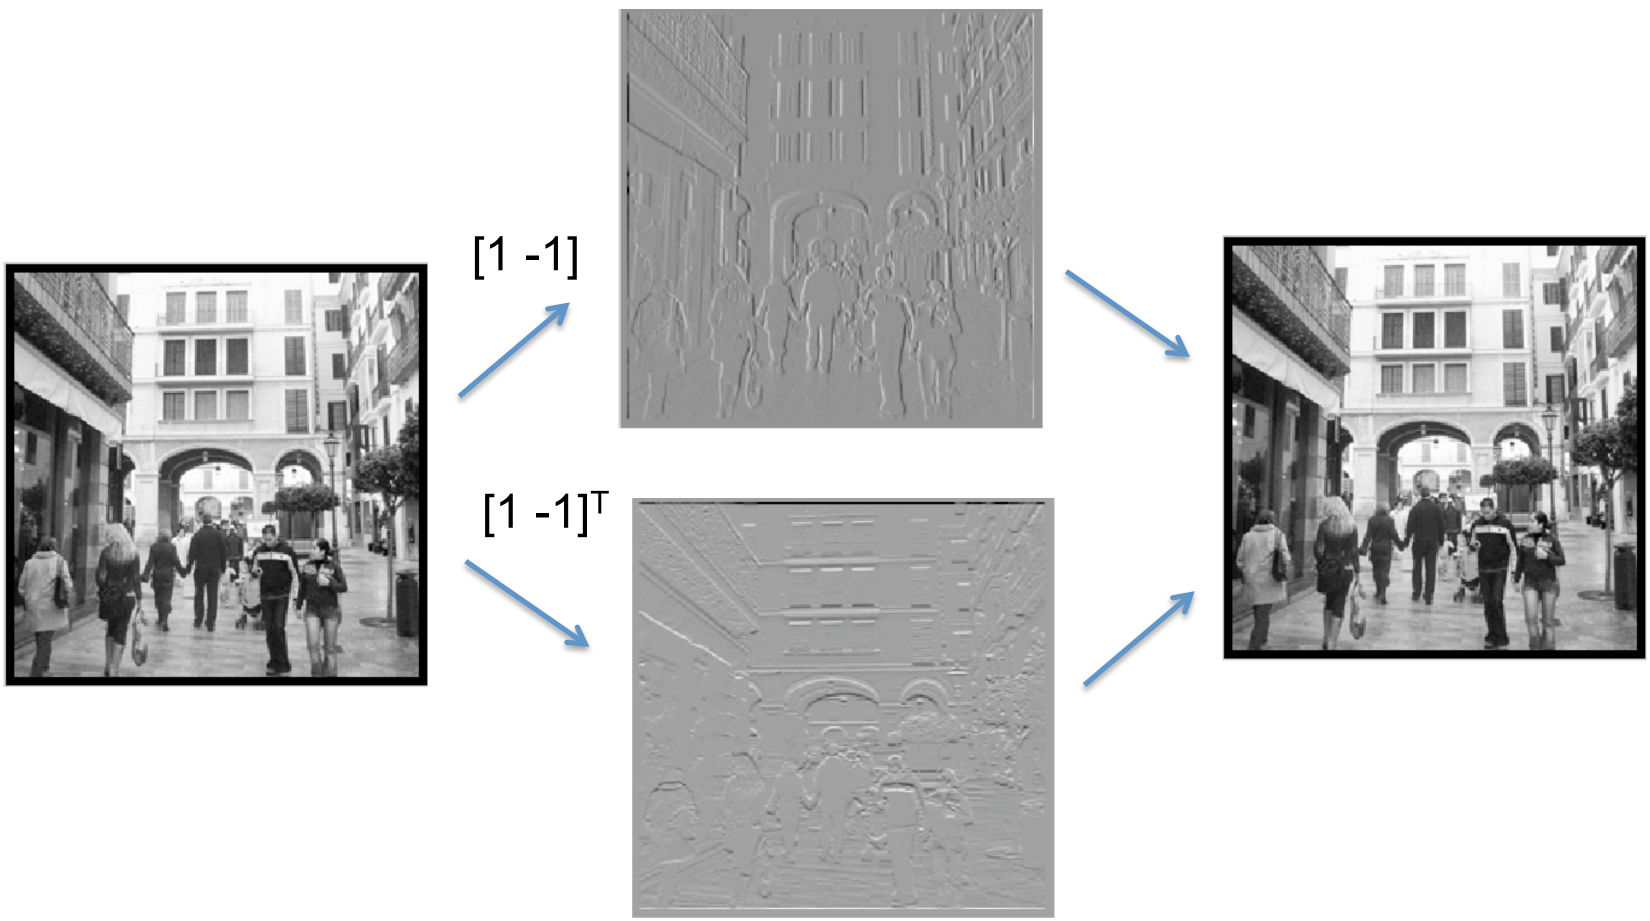

Editing the edge image

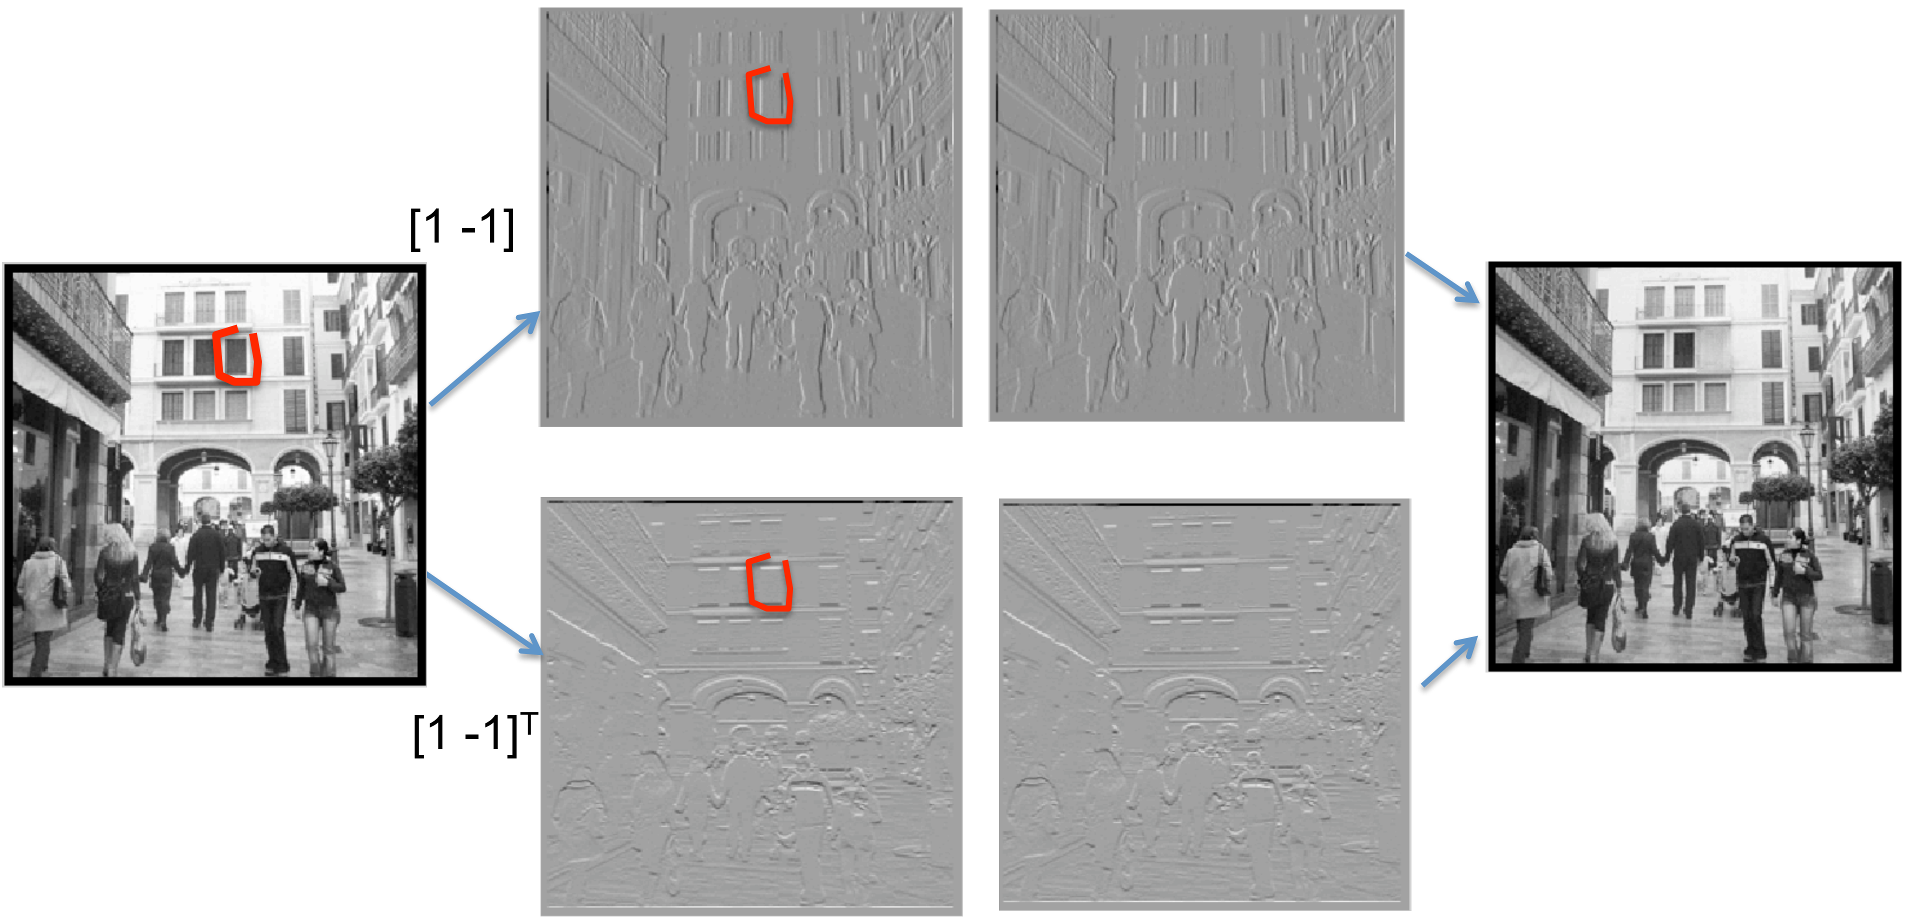

Thresholding edges

Intrinsic images

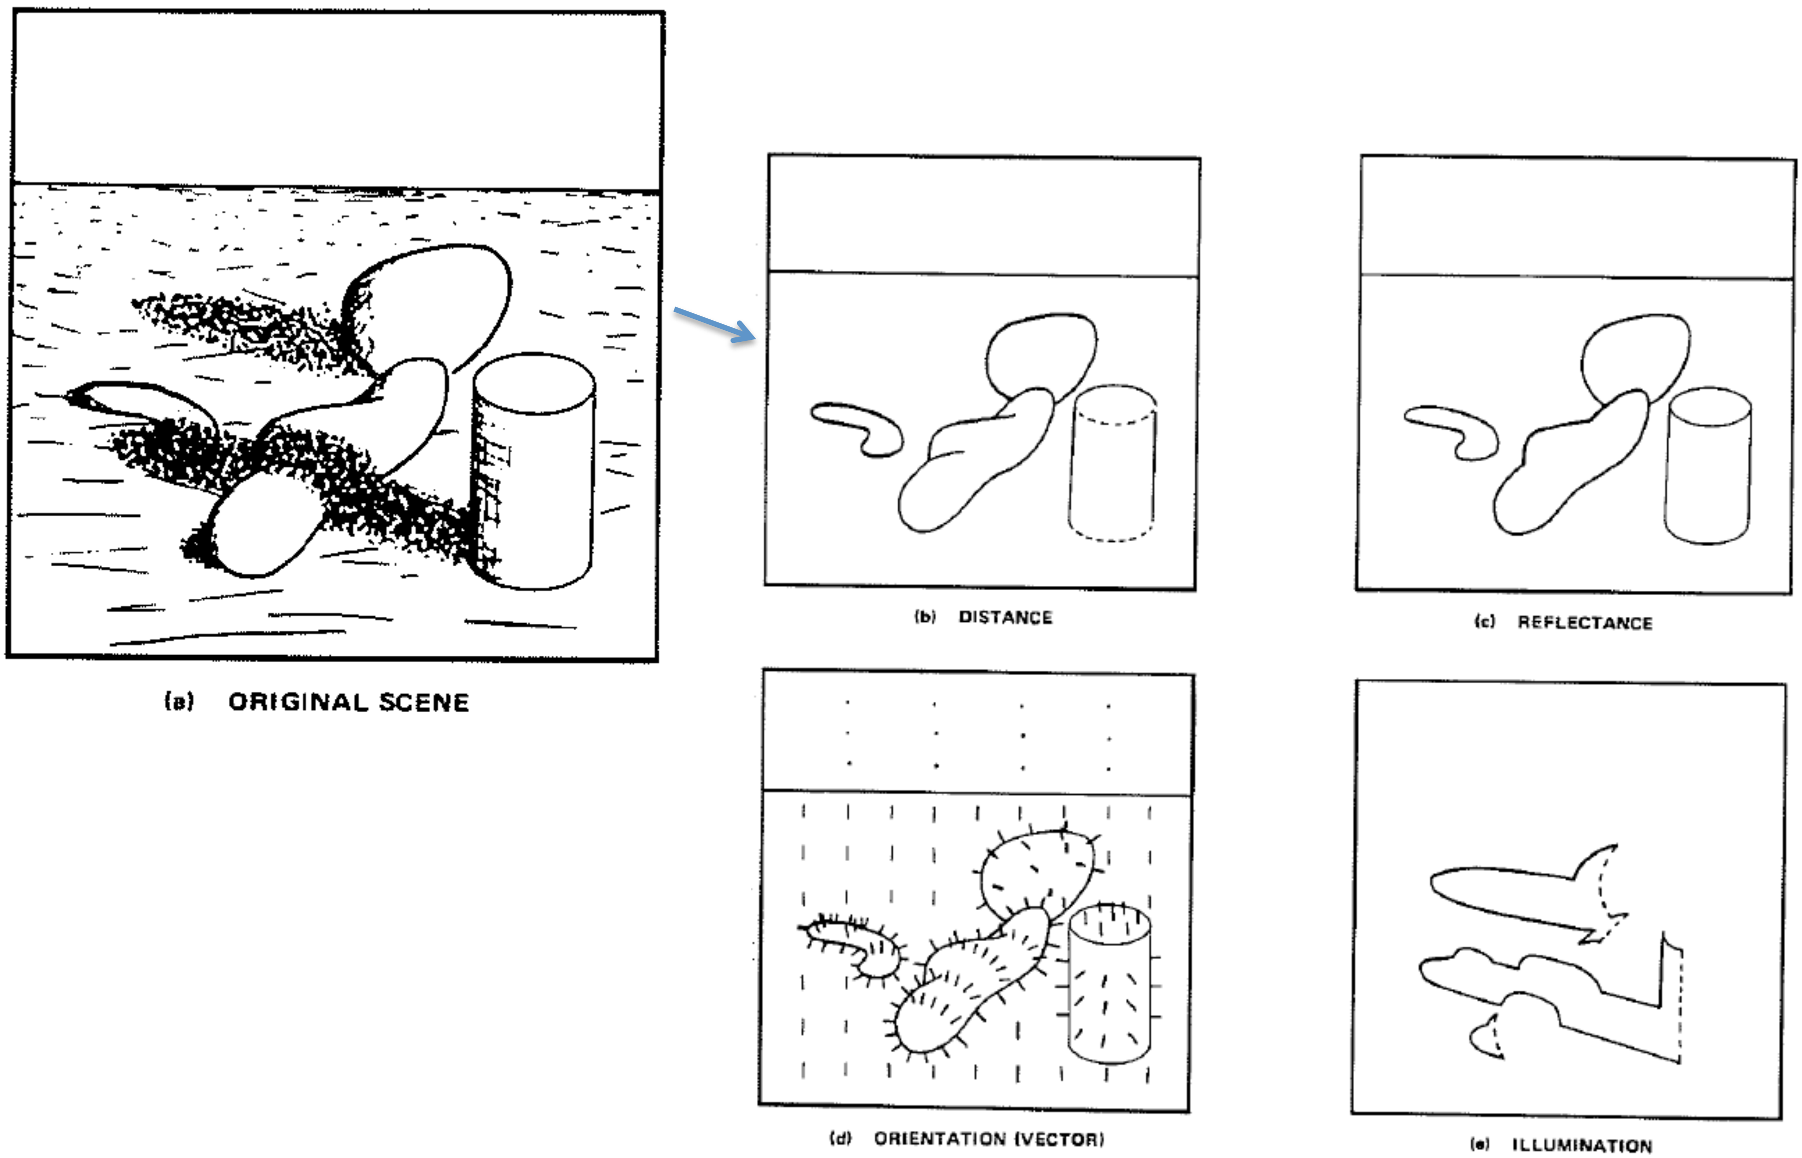

RECOVERING INTRINSIC SCENE CHARACTERISTICS FROM IMAGES

Technical Note 157

April 1978

By: Harry G. Barrow
J. Martin Tenenbaum
Artificial Intelligence Center

The research reported herein was supported by the National Science Foundation, under NSF Grant No. ENG76-01272.

To appear in *Computer Vision Systems*, A. Hanson and E. Riseman, eds., (Academic Press, New York, in press).

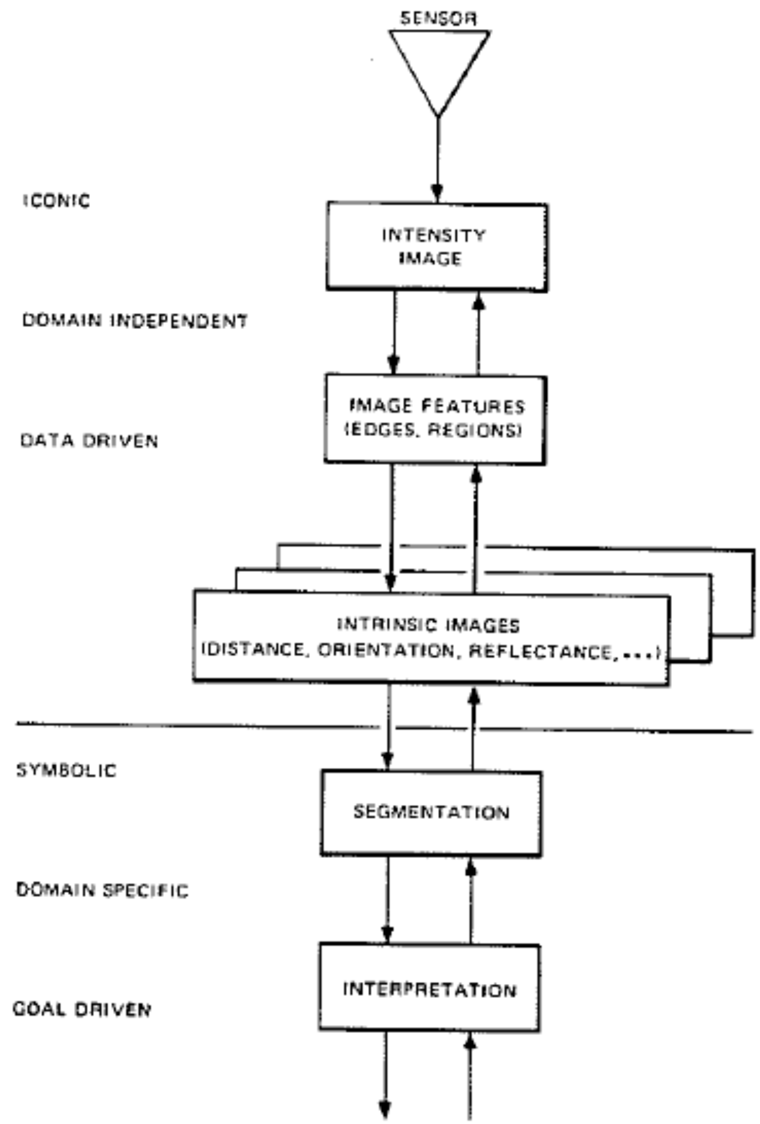

Figure 10 Organization of a visual system

Forming an Image

Illuminate the surface to get:

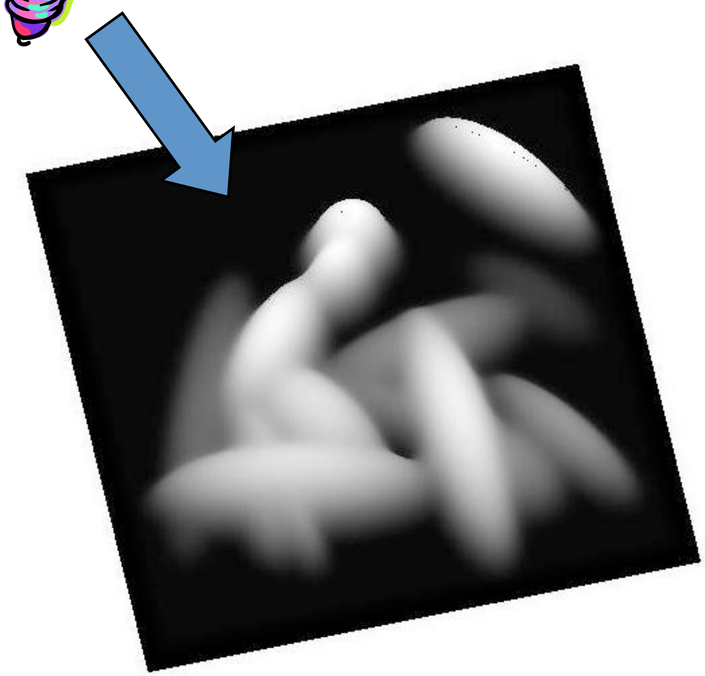

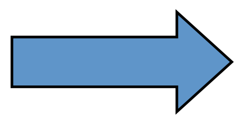

Surface (Height Map)

Shading Image

The shading image is the interaction of the shape of the surface and the illumination

Slide: Marshal Tappen

Painting the Surface

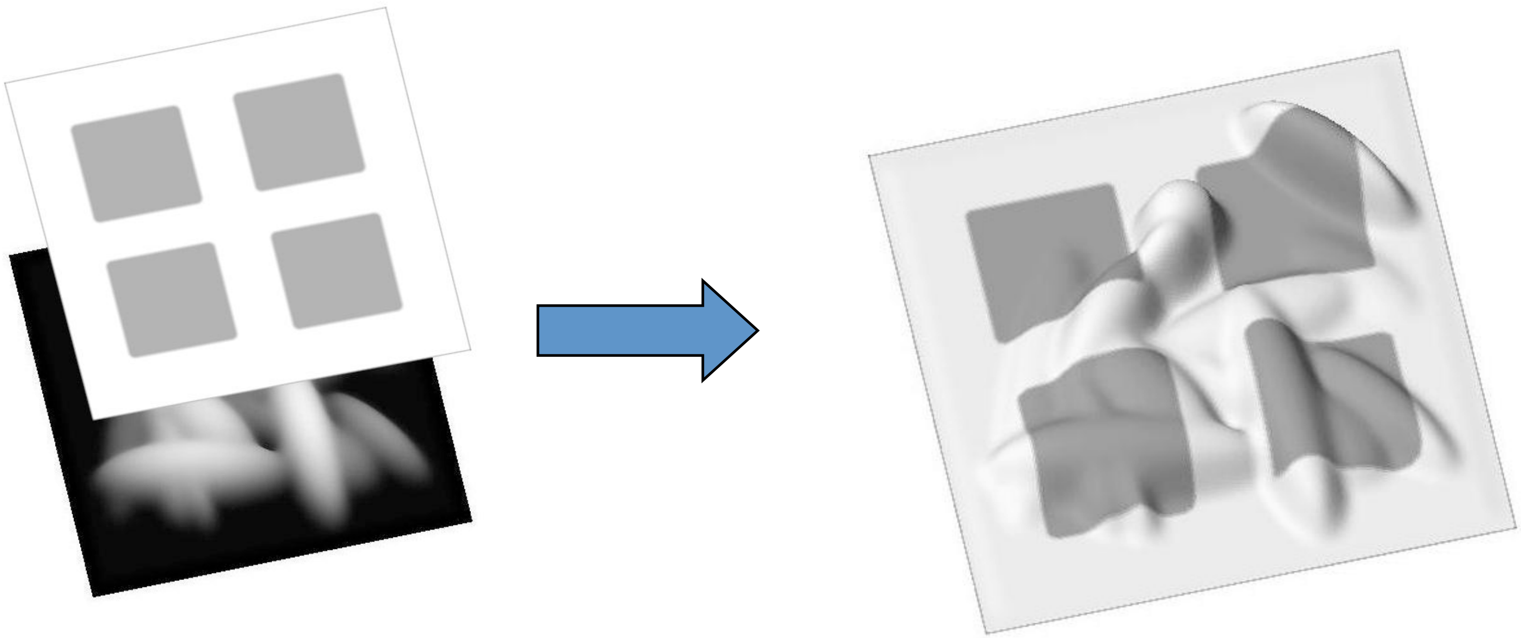

Scene Image

Add a reflectance pattern to the surface. Points inside the squares should reflect less light

Slide: Marshal Tappen

Goal

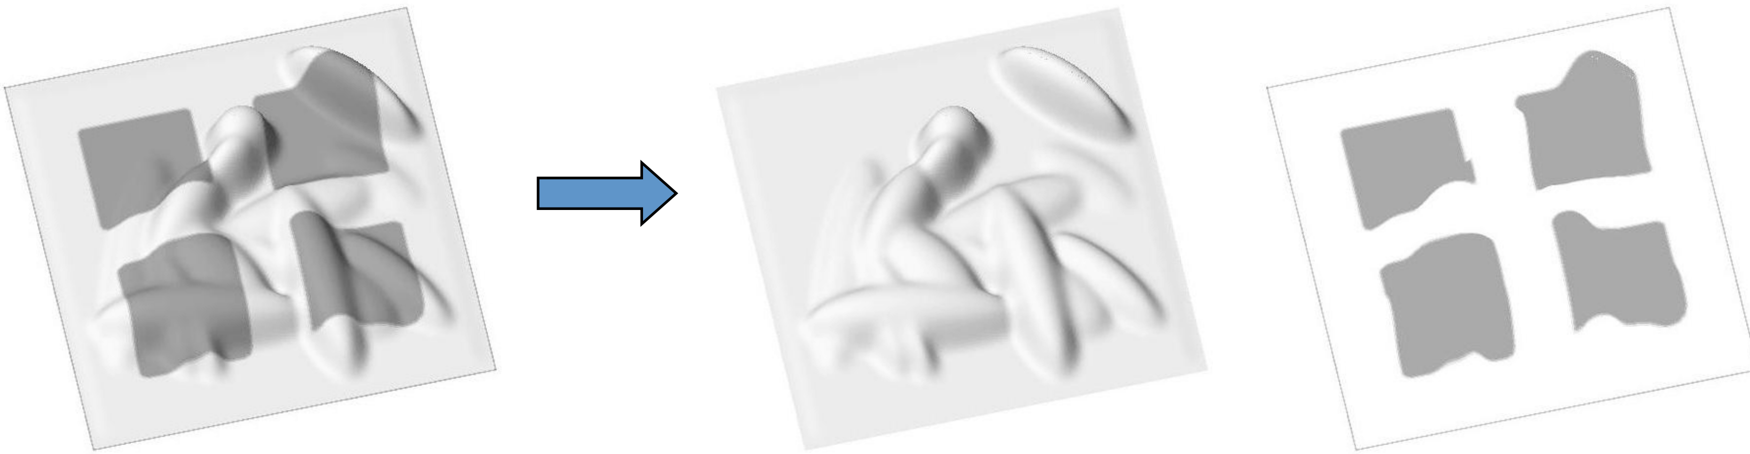

Image

Shading Image

Reflectance Image

Slide: Marshal Tappen

Retinex

E.H. Land, J.J. McCANN - Journal of the Optical society of America, 1971

Journal of the
OPTICAL SOCIETY
of AMERICA

VOLUME 61, NUMBER 1 JANUARY 1971

Lightness and Retinex Theory

EDWIN H. LAND* AND JOHN J. MCCANN
Polaroid Corporation, Cambridge, Massachusetts 02139
(Received 8 September 1970)

The reflectance tends to be constant across space except for abrupt changes at the transitions between objects or pigments. Thus a reflectance change shows itself as step edge in an image, while illuminance changes gradually over space. By this argument one can separate reflectance change from illuminance change by taking spatial derivatives: High derivatives are due to reflectance and low ones are due to illuminance.

Retinex

=

x

Again, we are trying to solve an ill-posed problem:

24 = ? x ?

From M. Tappen, PhD

Retinex

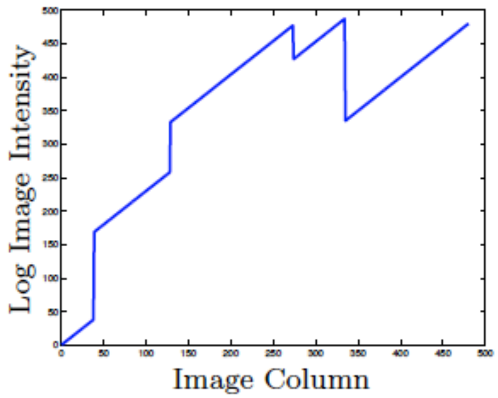

(a) One column from the observed image.

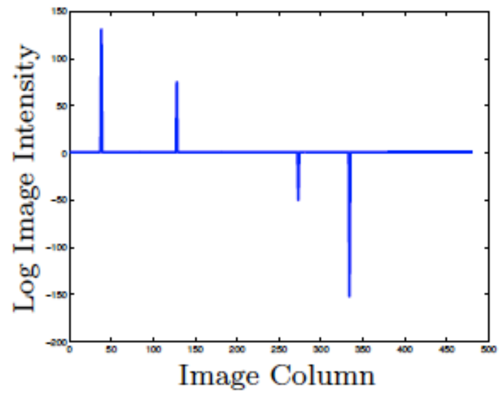

(b) The derivative of the plot from (a).

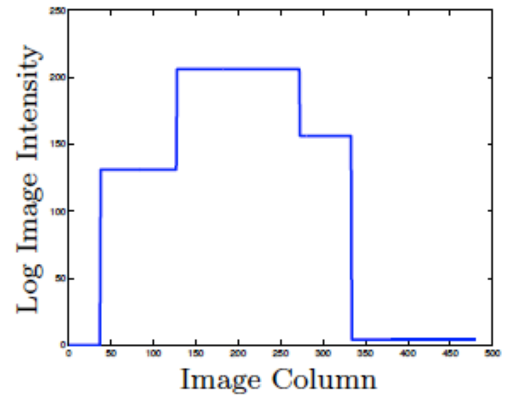

(c) The estimate of the log shading

From M. Tappen, PhD

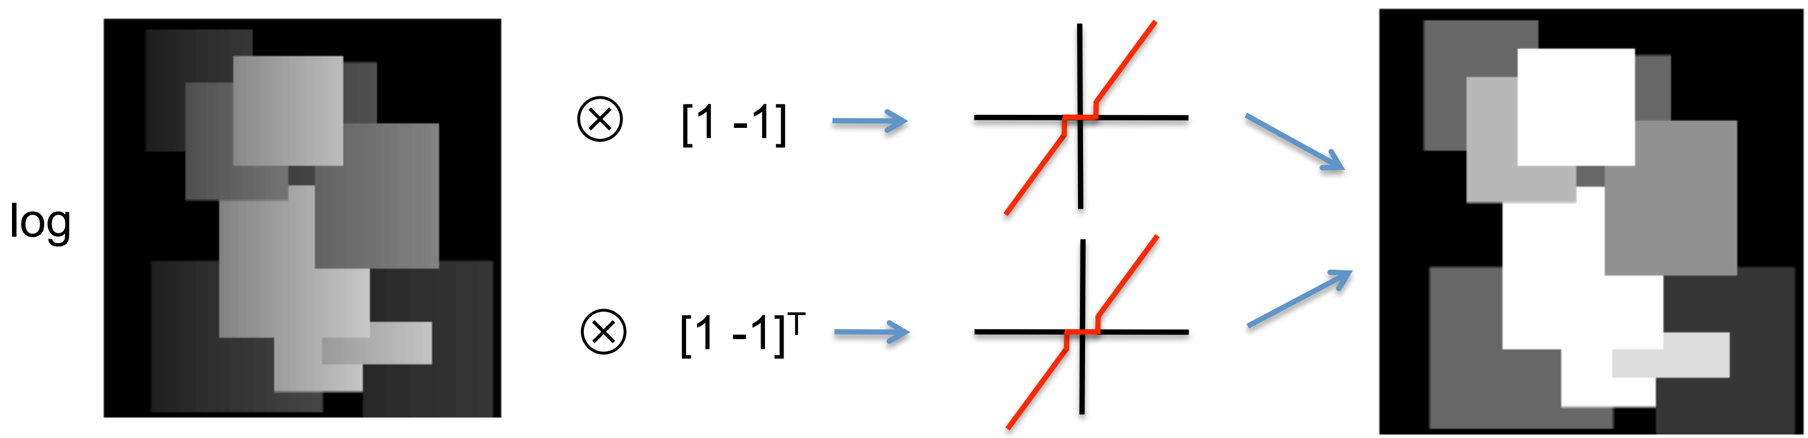

From M. Tappen, PhD

Craik-O'Brien-Cornsweet effect

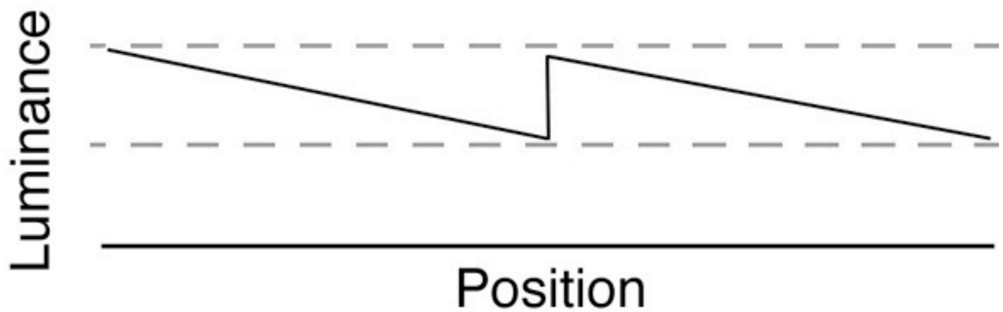

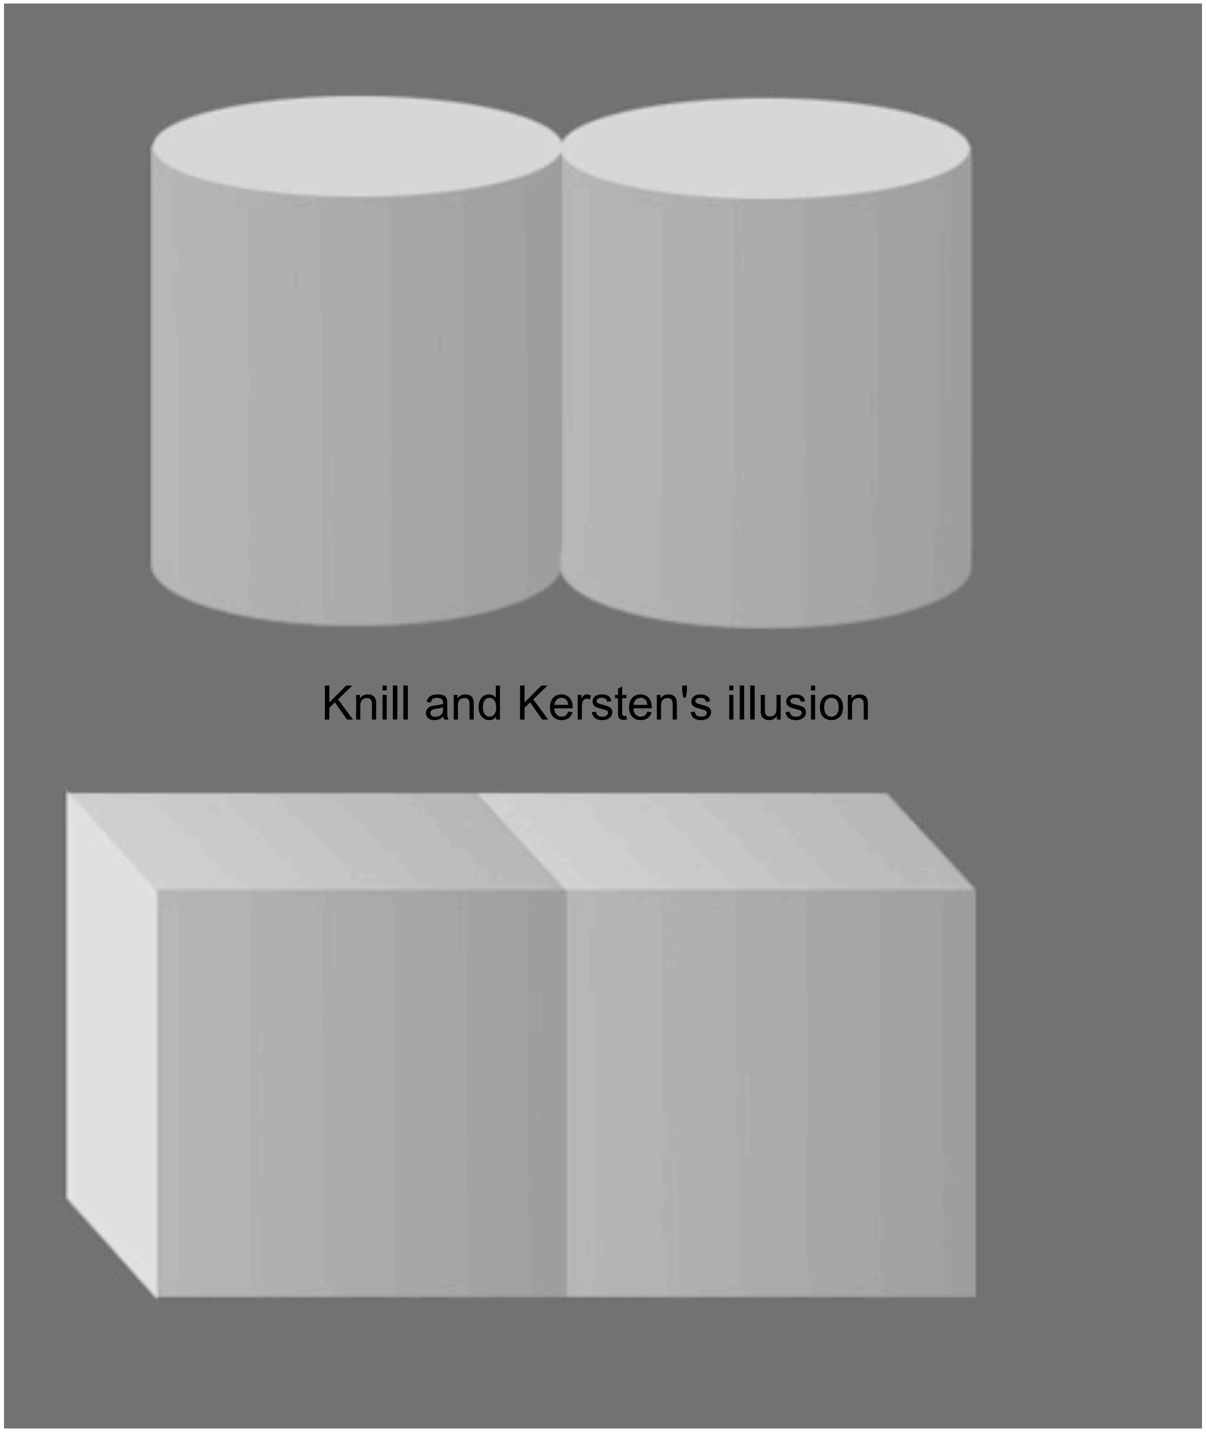

This illusion highlights the importance of scene interpretation.

The effect is gone

and it comes back when the gradient is not explained by the shape.

A more general approach

1. Compute the *x* and *y* image derivatives
2. Classify each derivative as being caused by *either* shading or a reflectance change
3. Set derivatives with the wrong label to zero.
4. Recover the intrinsic images by finding the least-squares solution of the derivatives.

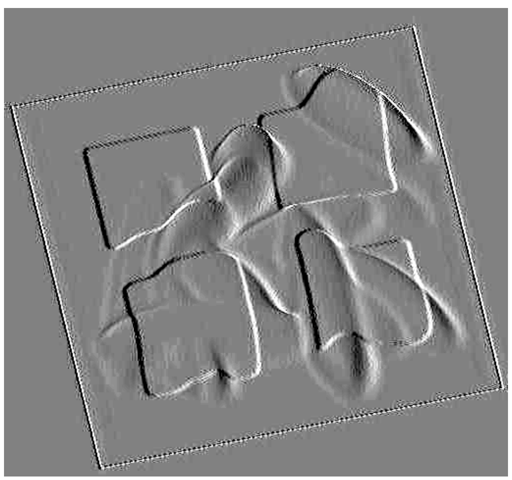

Original *x* derivative image

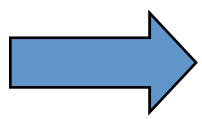

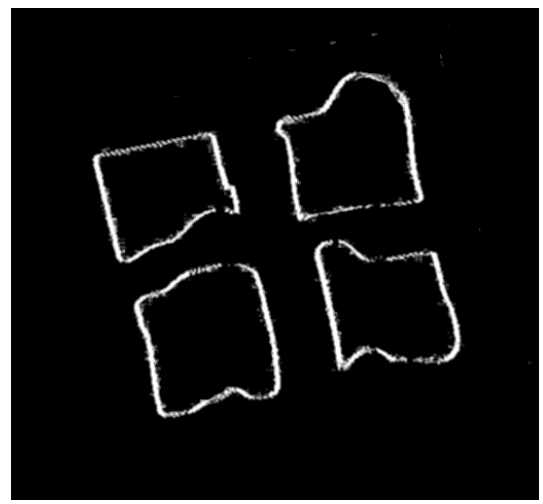

Classify each derivative
(White is reflectance)

Slide: Marshal Tappen

What edges are important?


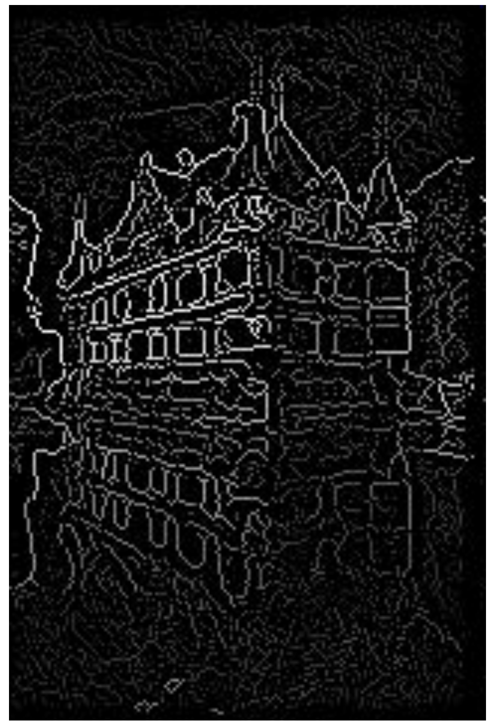
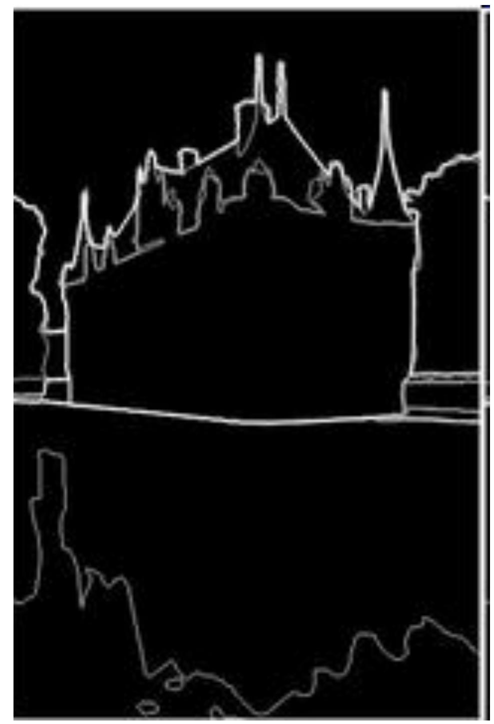

Differential Geometry Descriptors

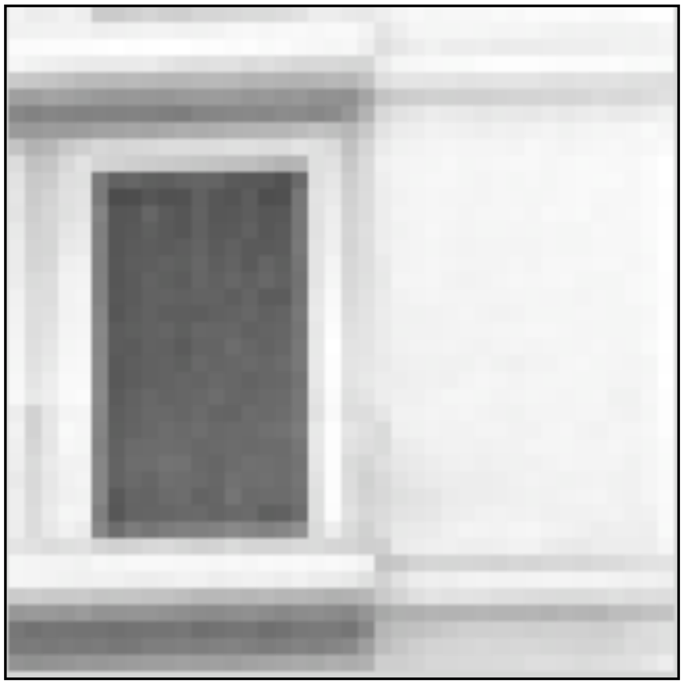

g(n,m)

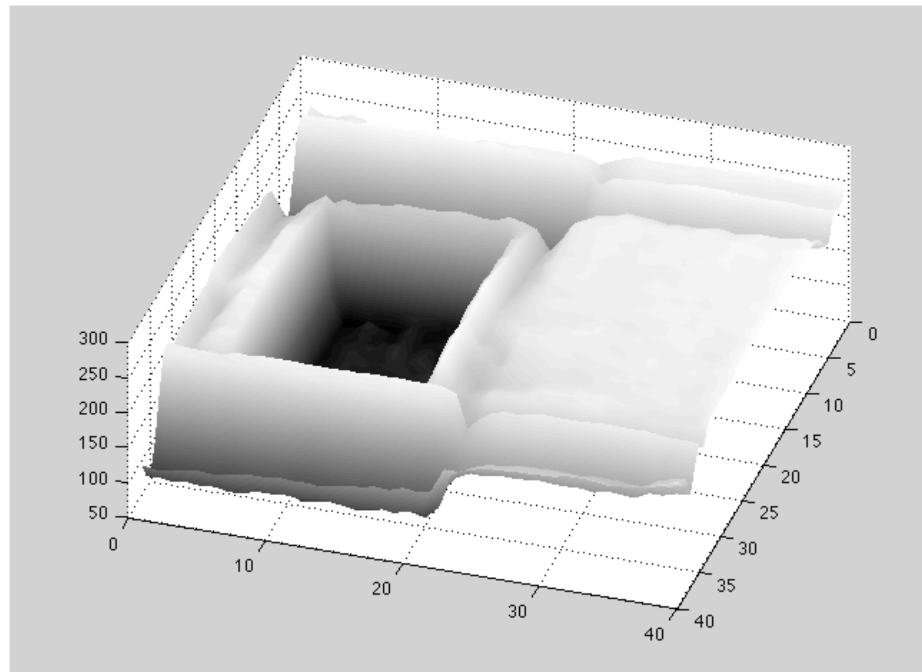

Differential Geometry Descriptors

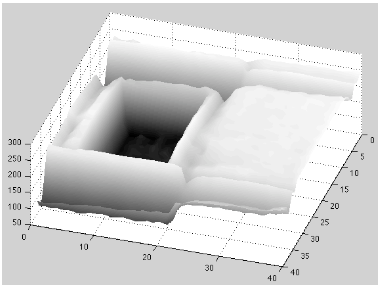

If we think of the image as a continuous function

g(x,y)

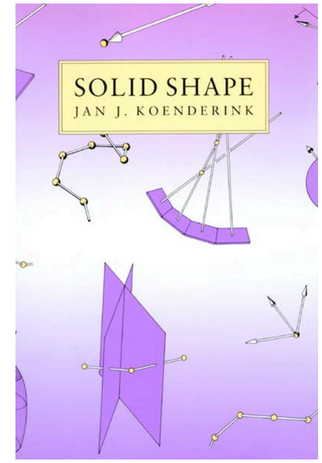

Image gradient:

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

Directional gradient:

|u|=1

$$u^T \nabla g = \cos(\alpha) \frac{\partial g(x,y)}{\partial x} + \sin(\alpha) \frac{\partial g(x,y)}{\partial y}$$

Laplacian:

$$\nabla^2 g = \frac{\partial^2 g(x,y)}{\partial x^2} + \frac{\partial^2 g(x,y)}{\partial y^2}$$

$$h(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

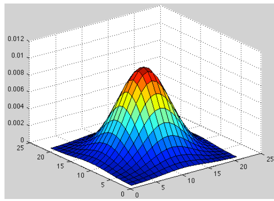

$$\frac{\partial g(x,y)}{\partial x} \otimes h(x,y) =$$

$$= \frac{\partial g(x,y) \otimes h(x,y)}{\partial x} =$$

$$= g(x,y) \otimes \frac{\partial h(x,y)}{\partial x}$$

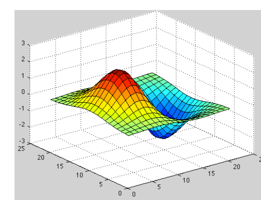

$$\frac{\partial h(x,y)}{\partial x} = \frac{-x}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

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

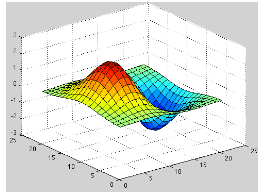

$$h_y(x,y) = \frac{\partial h(x,y)}{\partial x} = \frac{-x}{2\pi\sigma^4} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

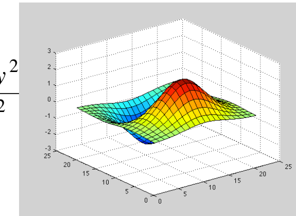

The smoothed directional gradient is a linear combination of two kernels

$$u^T \nabla g \otimes h = \left(\cos(\alpha) h_x(x,y) + \sin(\alpha) h_y(x,y)\right) \otimes g(x,y) =$$

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

$$= \cos(\alpha) h_x(x,y) \otimes g(x,y) + \sin(\alpha) h_y(x,y) \otimes g(x,y) =$$

= cos(α)

+sin(α)

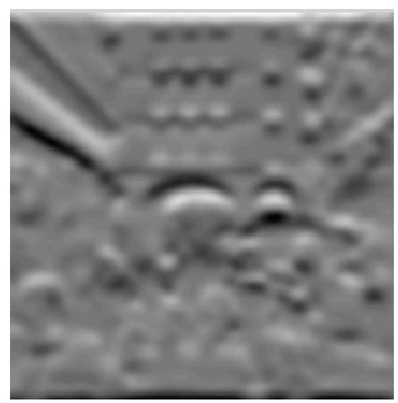

=

Steereability of gaussian derivatives, Freeman & Adelson 92

Laplacian

$$h(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

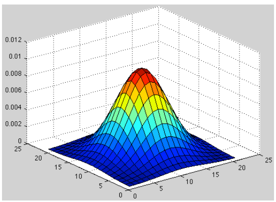

$$\nabla^2 g \otimes h = \left(\frac{\partial^2 g(x,y)}{\partial x^2} + \frac{\partial^2 g(x,y)}{\partial y^2} \right) \otimes h(x,y)$$

$$\nabla^2 g \otimes h = g \otimes \nabla^2 h$$

$$\nabla^2 h(x,y) = \left(\frac{x^2+y^2}{\sigma^4} - \frac{2}{\sigma^2} \right) h(x,y)$$

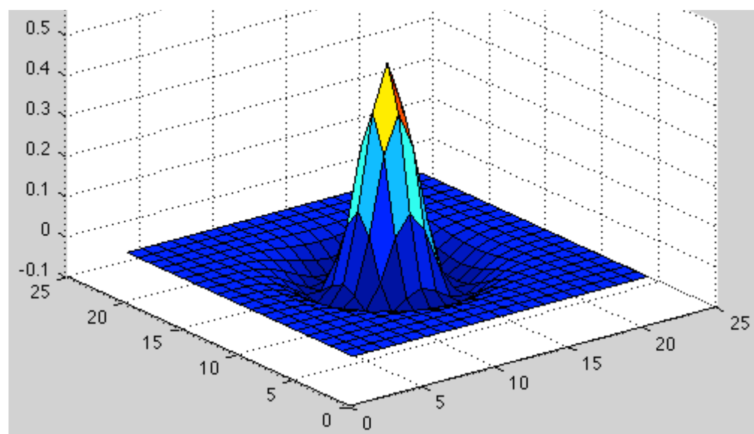

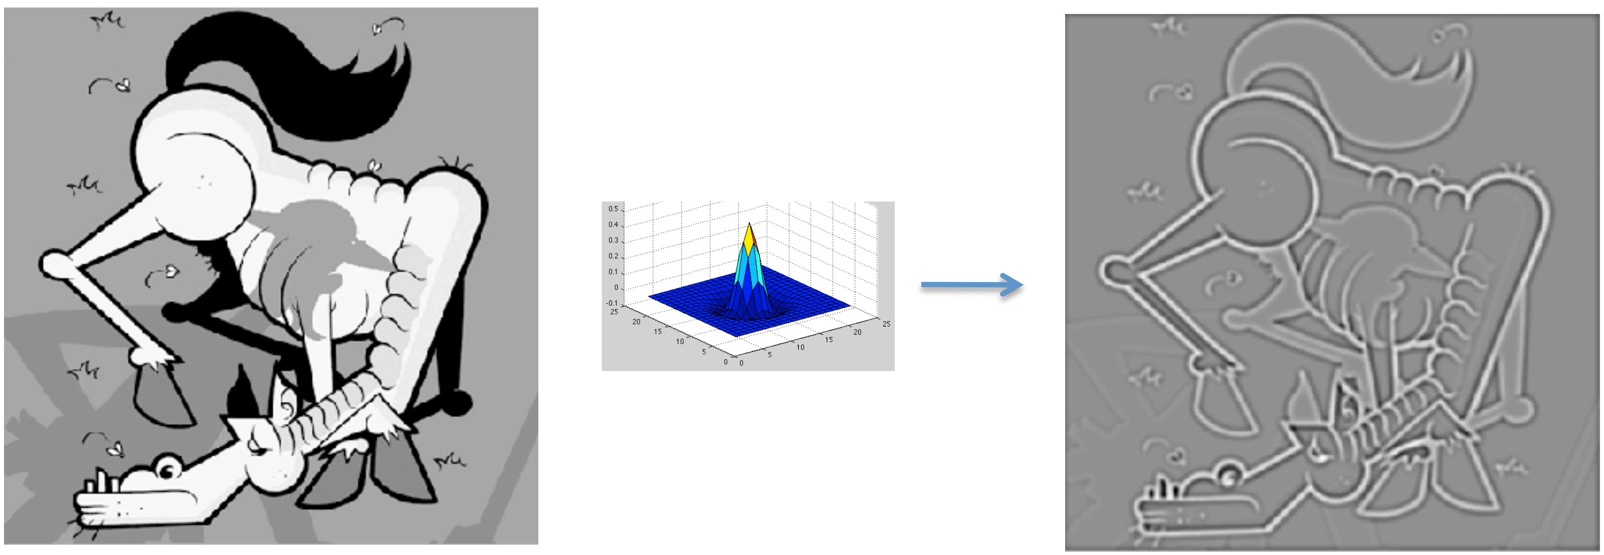

A "summary" of visual features

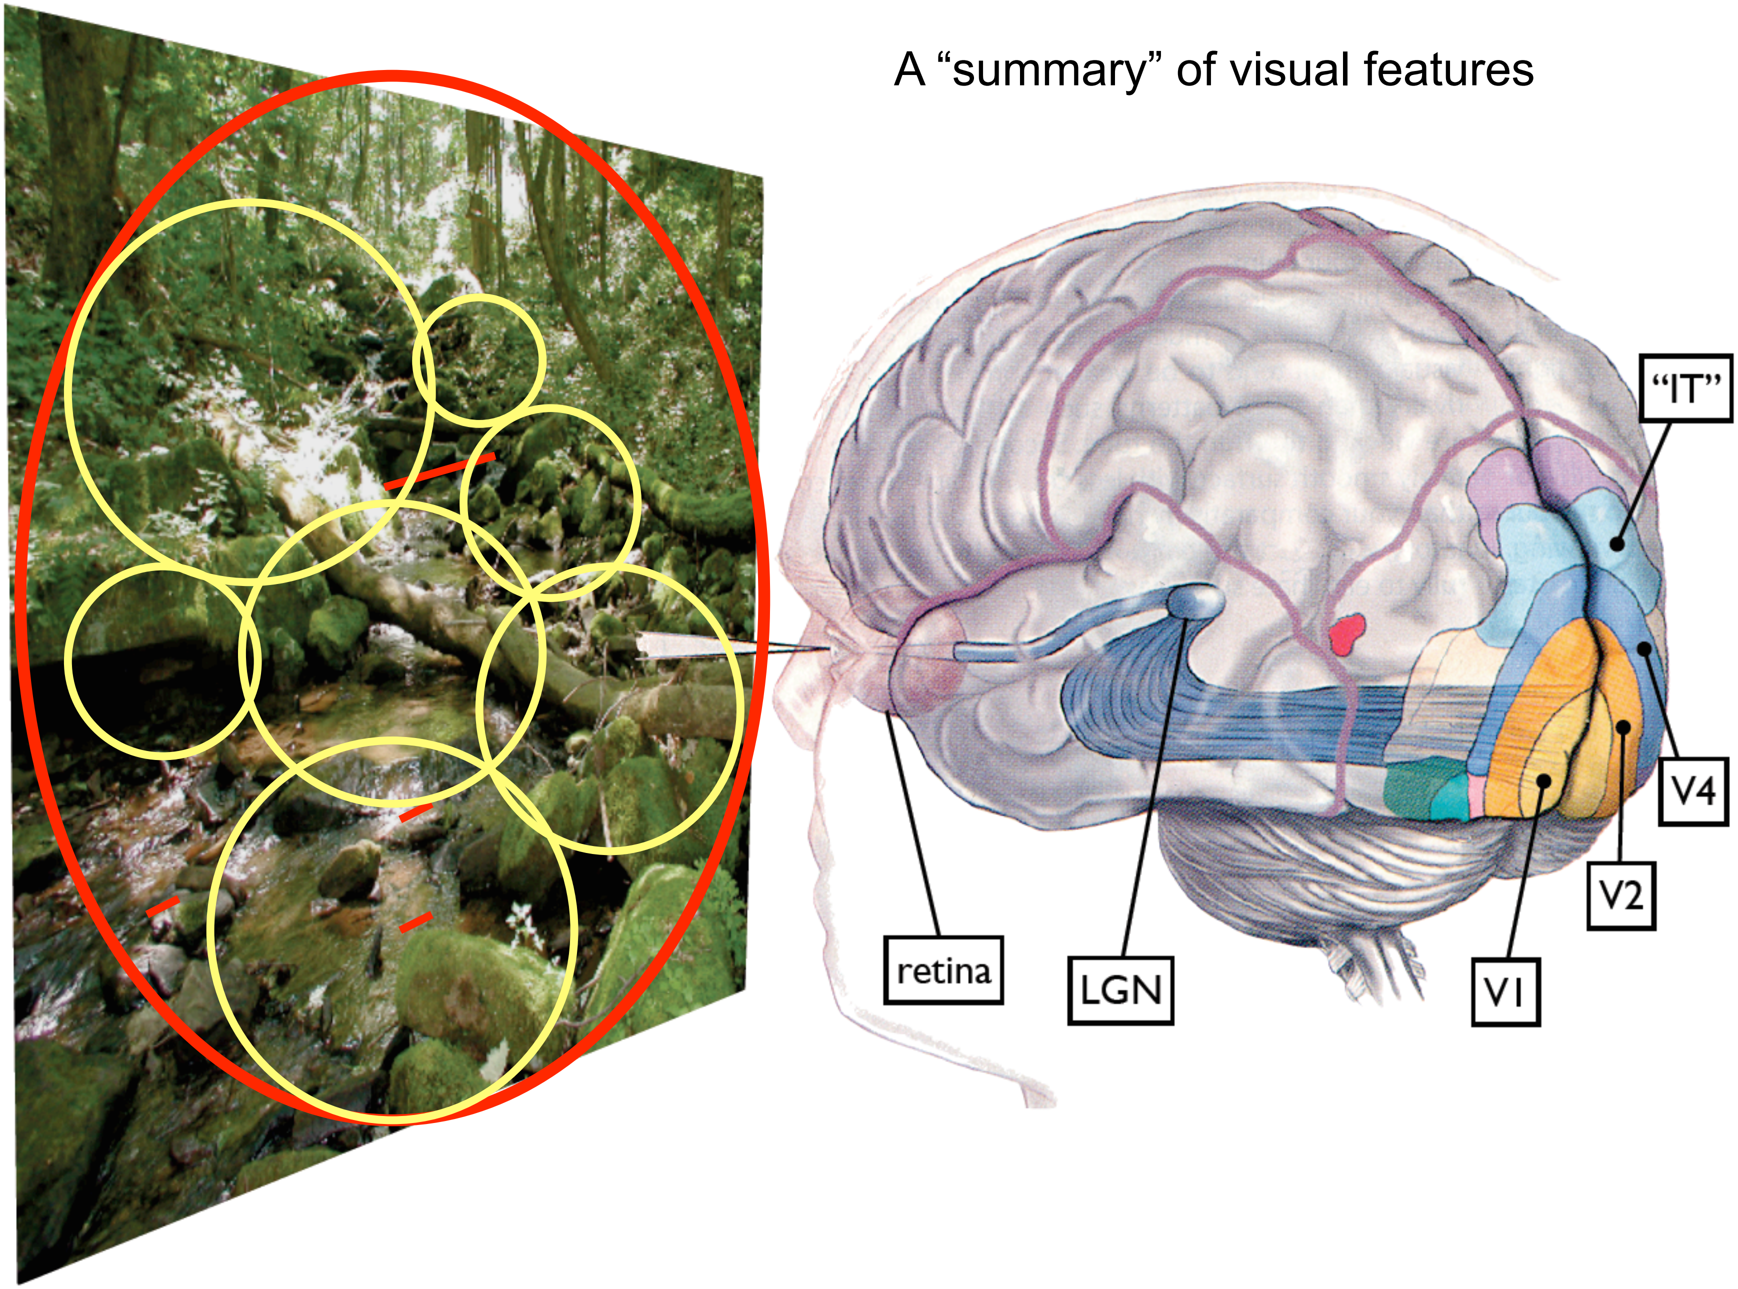

Comparing Human and Machine Perception

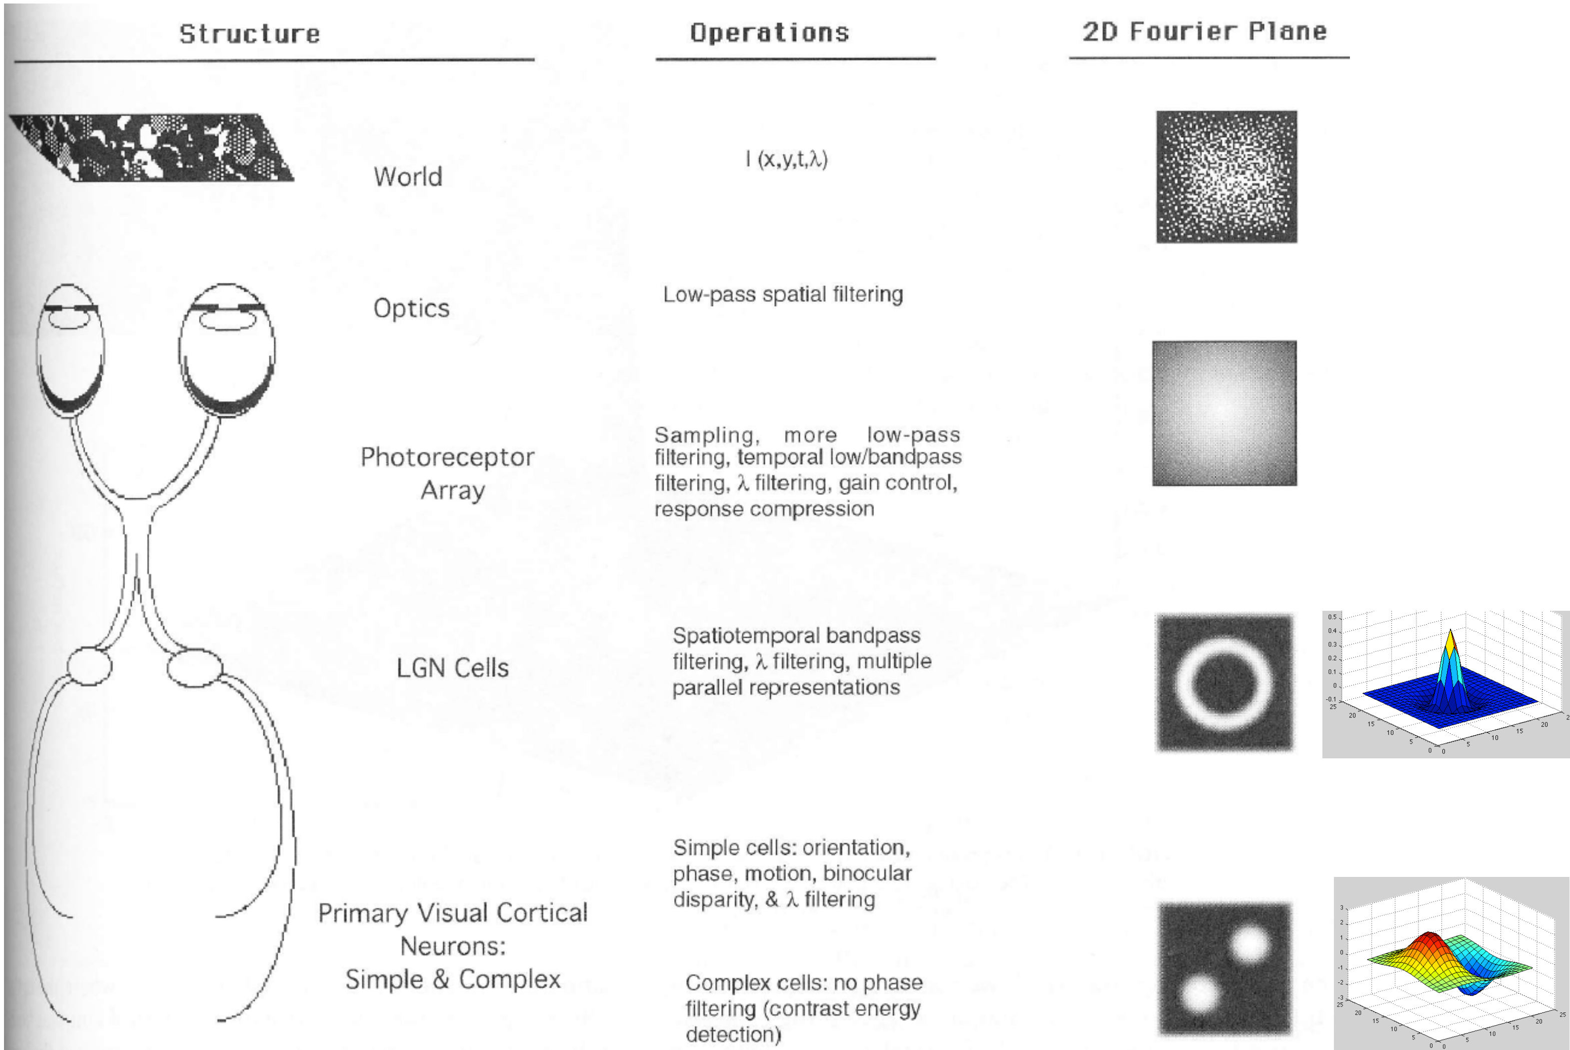

FIGURE 1 Schematic overview of the processing done by the early visual system. On the left, are some of the major structures to be discussed; in the middle, are some of the major operations done at the associated structure; in the right, are the 2-D Fourier representations of the world, retinal image, and sensitivities typical of a ganglion and cortical cell.

Contrast Sensitivity Function

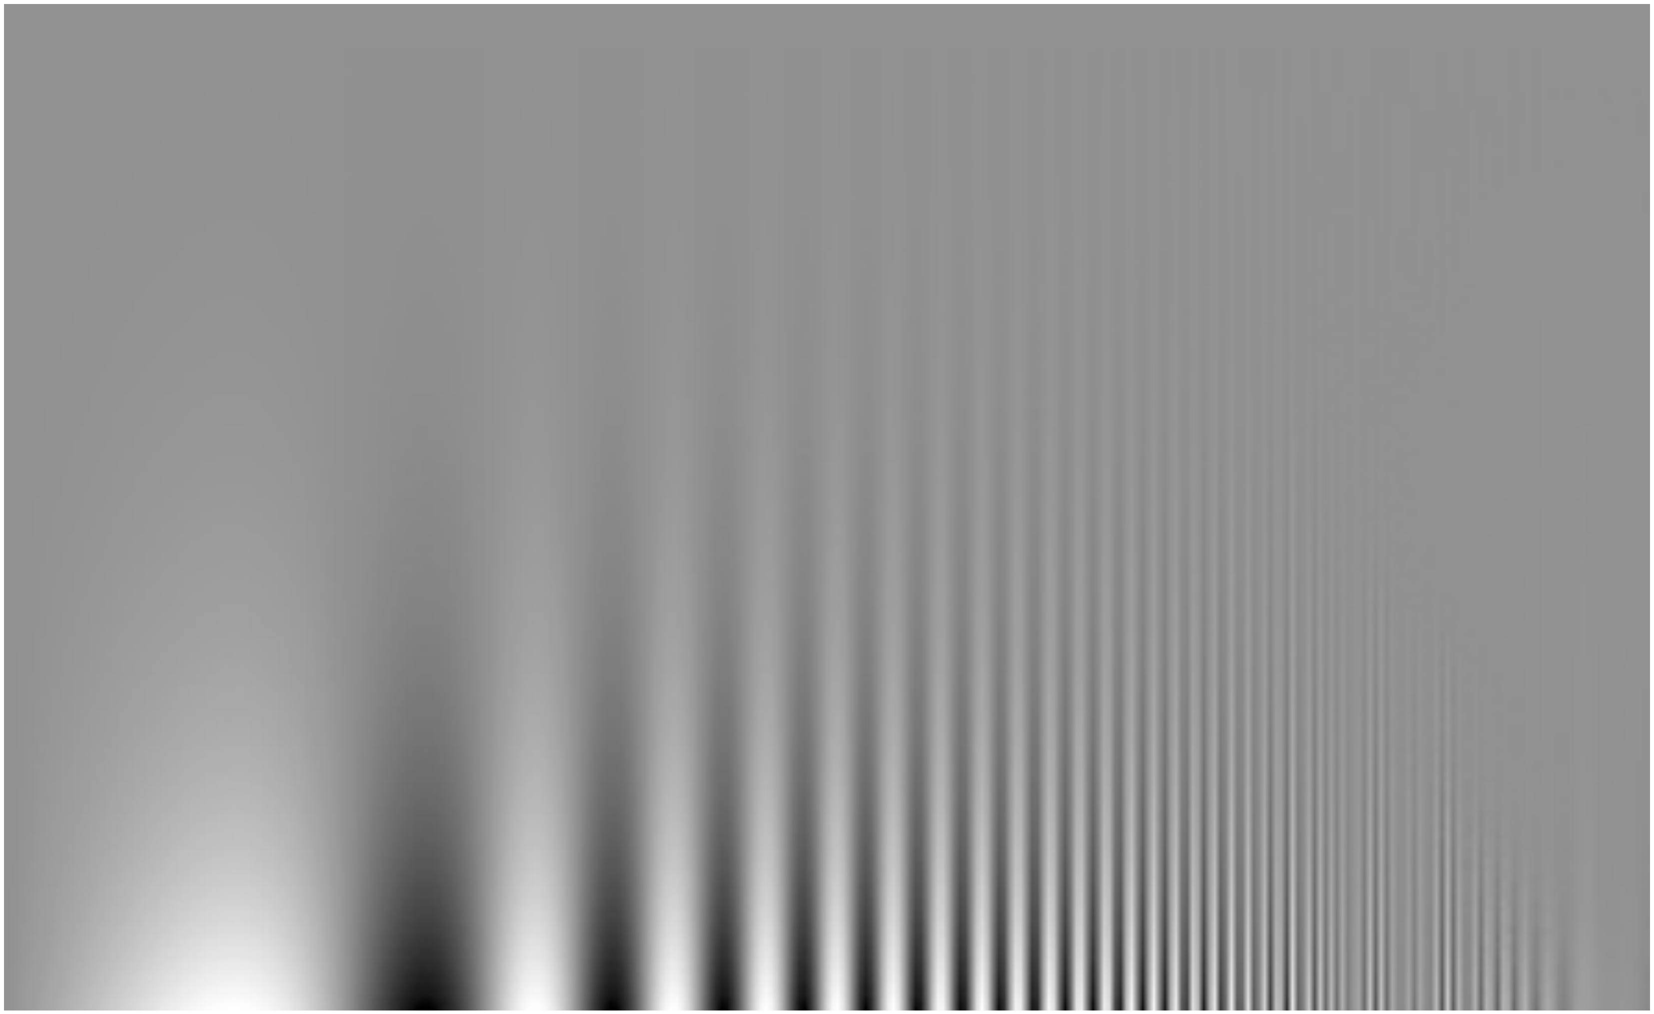

A demo of human contrast sensitivity as a function of spatial frequency. Frequency rises from left to right at a constant rate. Contrast drops from bottom to top at a constant rate. The bars are visible further up for middle frequencies, showing these are more salient to the human visual system.

Contrast Sensitivity Function

Blackmore & Campbell (1969)

Maximum sensitivity

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