# Fourier Transform and Frequency Domain

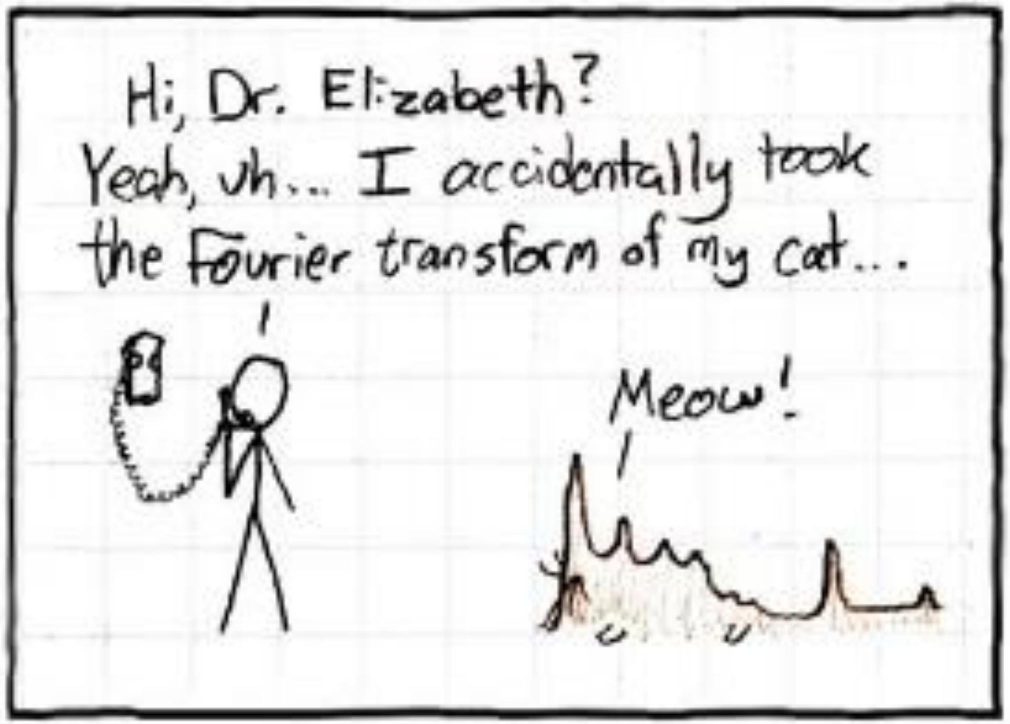


15-463, 15-663, 15-862
Computational Photography
Fall 2017, Lecture 6

http://graphics.cs.cmu.edu/courses/15-463

# Course announcements

- Last call for responses to Doodle about rescheduling the September 27th lecture!
  - Link available on Piazza.
  - Currently 17 responses. I’ll pick a date on Tuesday evening.

- Homework 1 is being graded.
  - Grades with comments will be uploaded on Canvas hopefully by Wednesday.
  - How was it?

- Homework 2 has been posted.
  - Much larger than homework 1.
  - Start early! Experiments take a long time to run.
  - How many have read/started/finished it?

# Overview of today's lecture

- Some history.
- Fourier series.
- Frequency domain.
- Fourier transform.
- Frequency-domain filtering.
- Revisiting sampling.

# Slide credits

Most of these slides were adapted from:

- Kris Kitani (15-463, Fall 2016).

Some slides were inspired or taken from:

- Fredo Durand (MIT).
- James Hays (Georgia Tech).

# Some history

# Who is this guy?

# What is he famous for?

Jean Baptiste Joseph Fourier
(1768-1830)

# What is he famous for?

Jean Baptiste Joseph Fourier
(1768-1830)

The Fourier series claim (1807):

*'Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.'*

# What is he famous for?

Jean Baptiste Joseph Fourier
(1768-1830)

The Fourier series claim (1807):

*'Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.'*

… and apparently also for the discovery of the greenhouse effect

# Is this claim true?

Jean Baptiste Joseph Fourier
(1768-1830)

The Fourier series claim (1807):

> *'Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.'*

# Is this claim true?

Jean Baptiste Joseph Fourier
(1768-1830)

The Fourier series claim (1807):

> *'Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.'*

Well, almost.

- The theorem requires additional conditions.
- Close enough to be named after him.
- Very surprising result at the time.

# Is this claim true?

Jean Baptiste Joseph Fourier
(1768-1830)

The Fourier series claim (1807):

> *‘Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.’*

Well, almost.

- The theorem requires additional conditions.
- Close enough to be named after him.
- Very surprising result at the time.

Malus

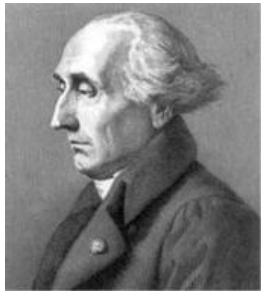

Lagrange

Legendre

Laplace

The committee examining his paper had expressed skepticism, in part due to not so rigorous proofs

# Amusing aside


Only known portrait of Adrien-Marie Legendre

1820 watercolor caricatures of French mathematicians Adrien-Marie Legendre (left) and Joseph Fourier (right) by French artist Julien-Leopold Boilly

For two hundred years, people were misidentifying this portrait as him

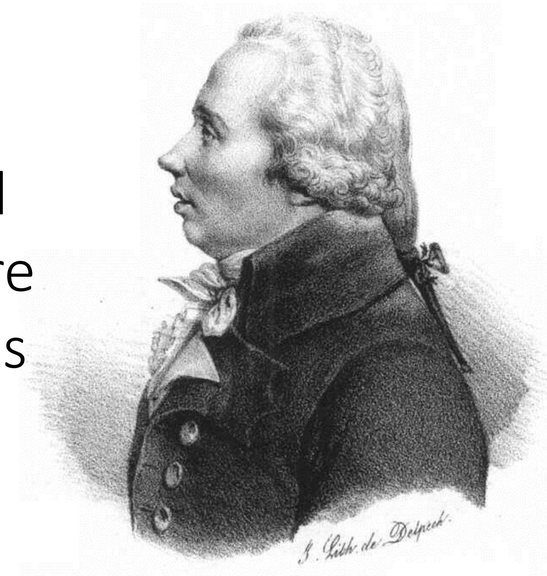

Louis Legendre (same last name, different person)

# Fourier series

# Basic building block

$$A \sin(\omega x + \phi)$$

Fourier's claim: Add enough of these to get <u>any</u> *periodic* signal you want!

# Basic building block

$$A \sin(\omega x + \phi)$$

amplitude

sinusoid

angular frequency

variable

phase

Fourier's claim: Add enough of these to get <u>any</u> *periodic* signal you want!

# Examples

How would you generate this function?

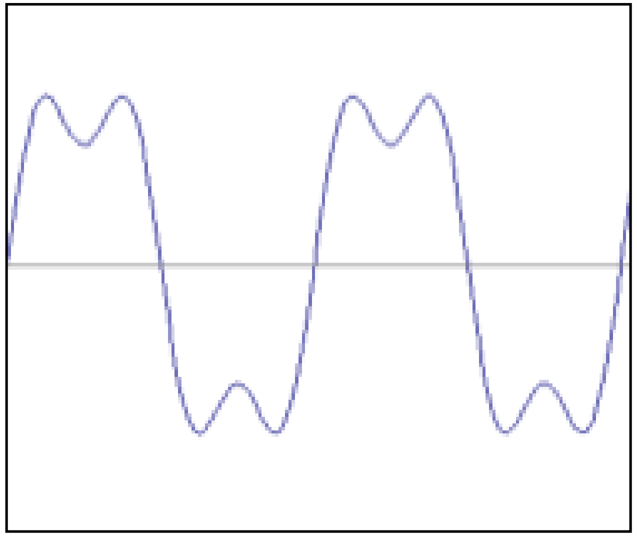 = ? + ?

# Examples

How would you generate this function?

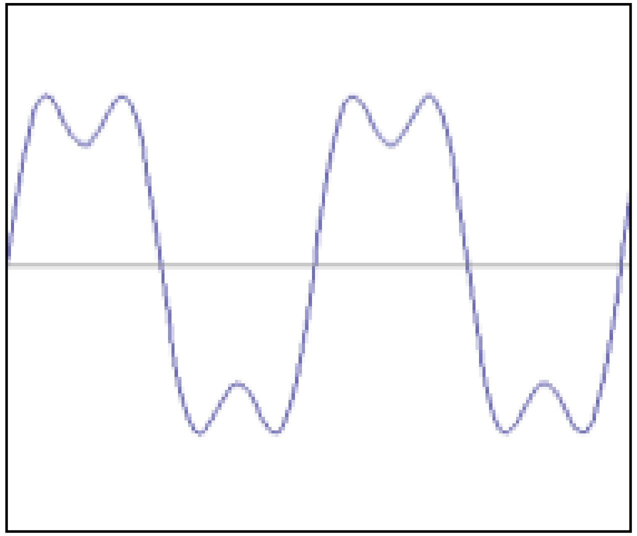

=

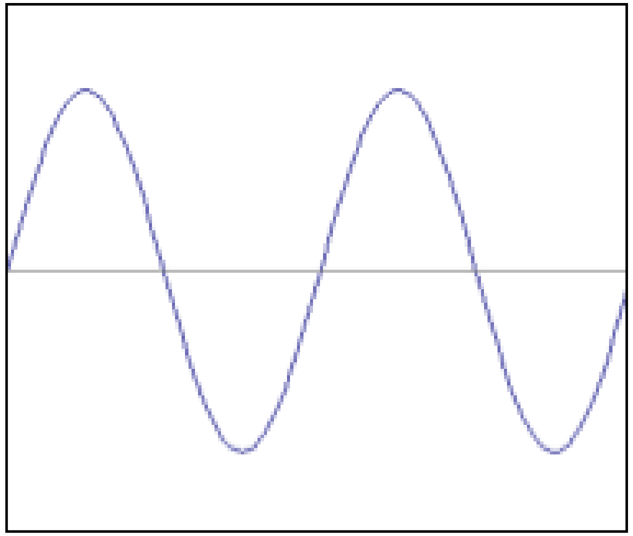

$\sin(2\pi x)$

+ ?

# Examples

How would you generate this function?

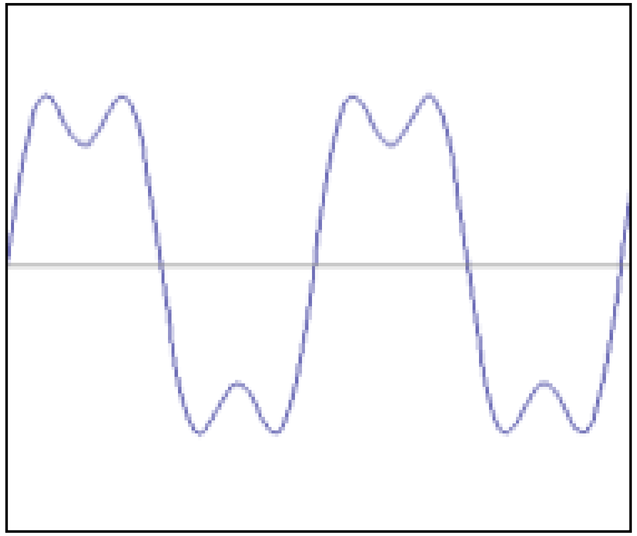

$$f(x) = \sin(2\pi x) + \frac{1}{3}\sin(2\pi 3x)$$

=

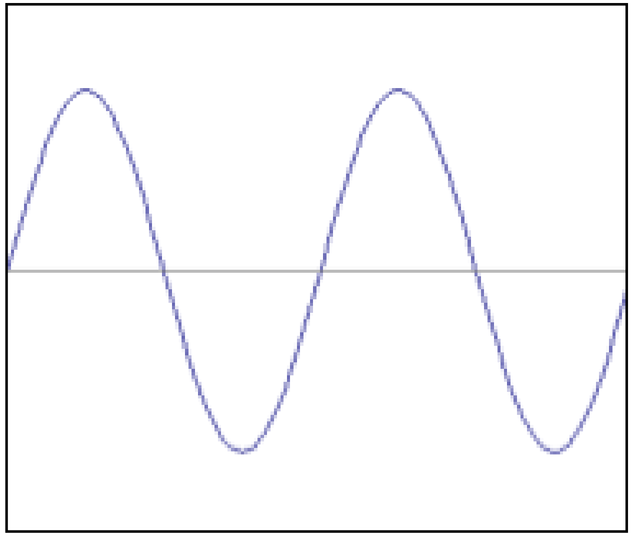

$$\sin(2\pi x)$$

+

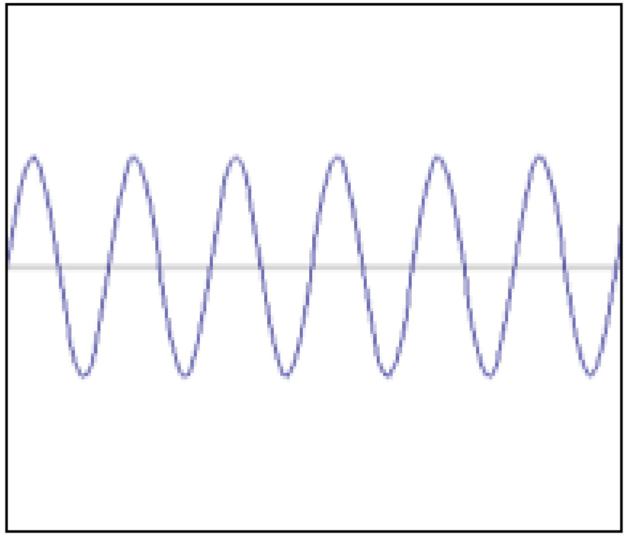

$$\frac{1}{3}\sin(2\pi 3x)$$

# Examples

How would you generate this function?

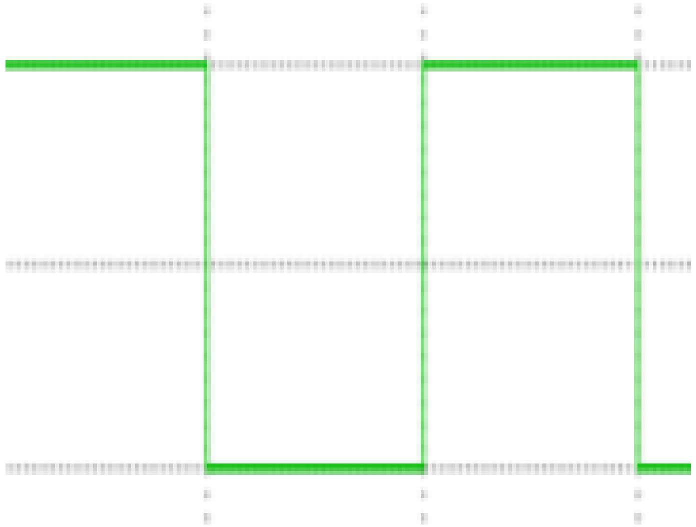

square wave

= ? + ?

# Examples

How would you generate this function?

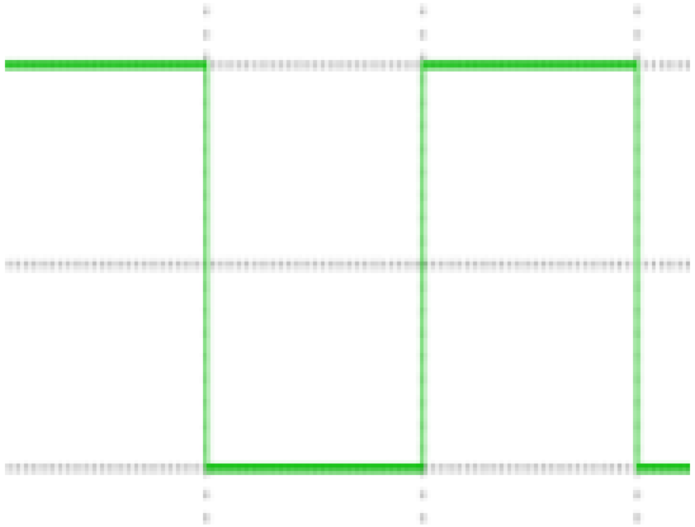

square wave

≈

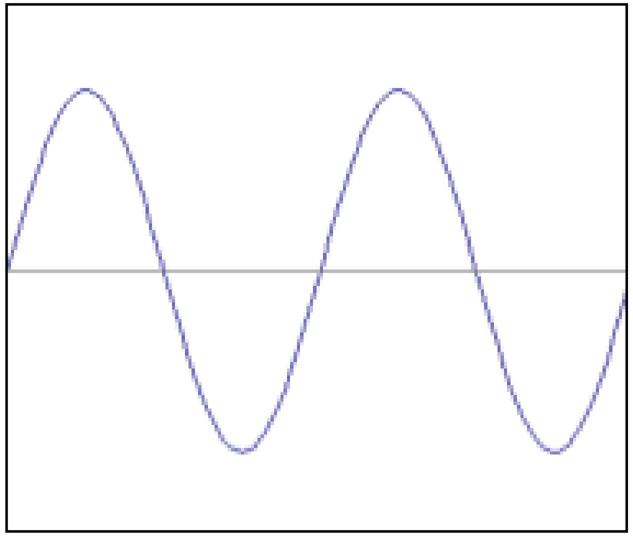

+

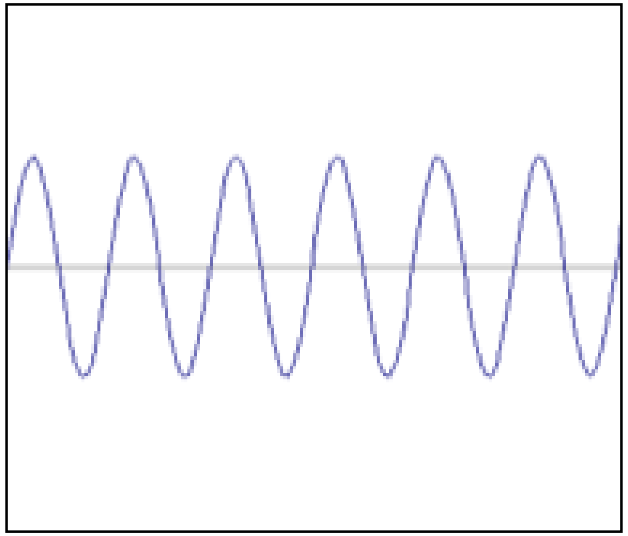

=

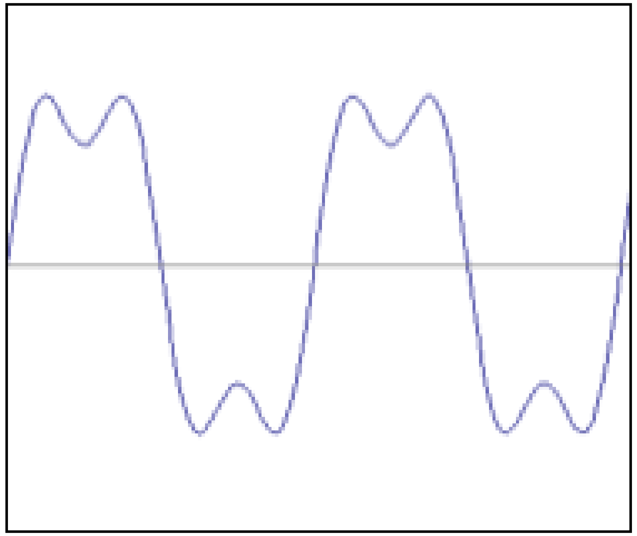

# Examples

How would you generate this function?

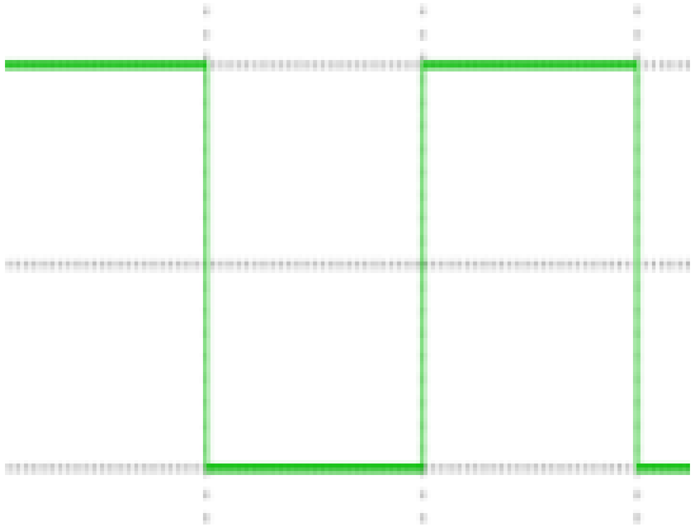

square wave

≈

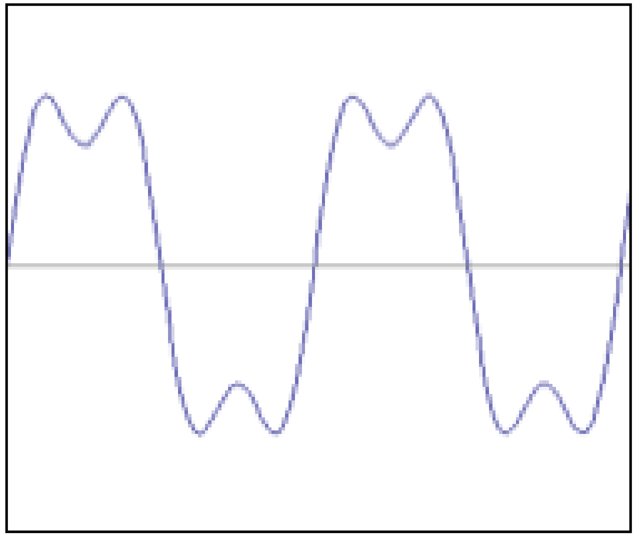

+

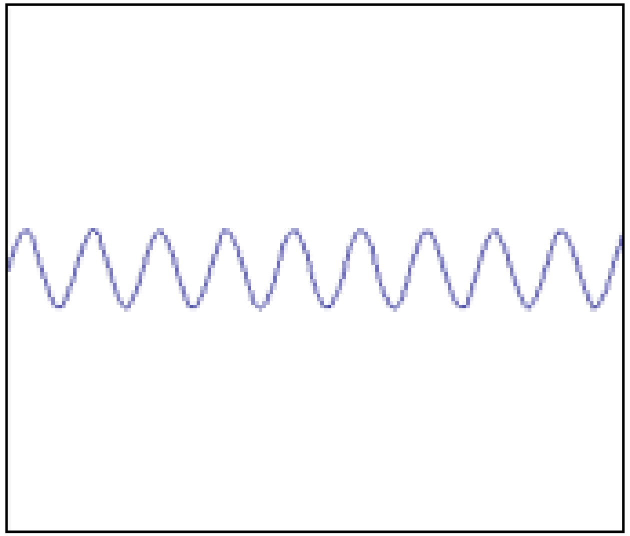

=

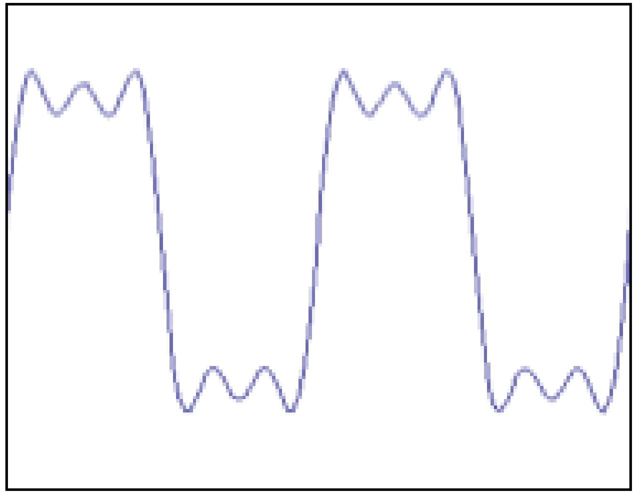

# Examples

How would you generate this function?

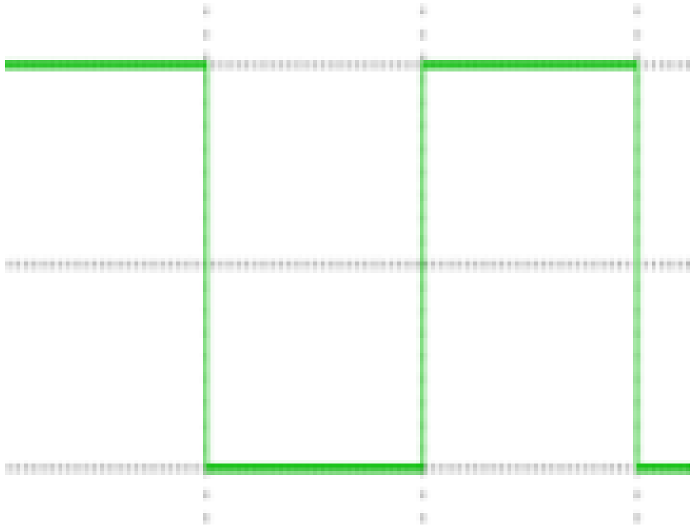

square wave

≈

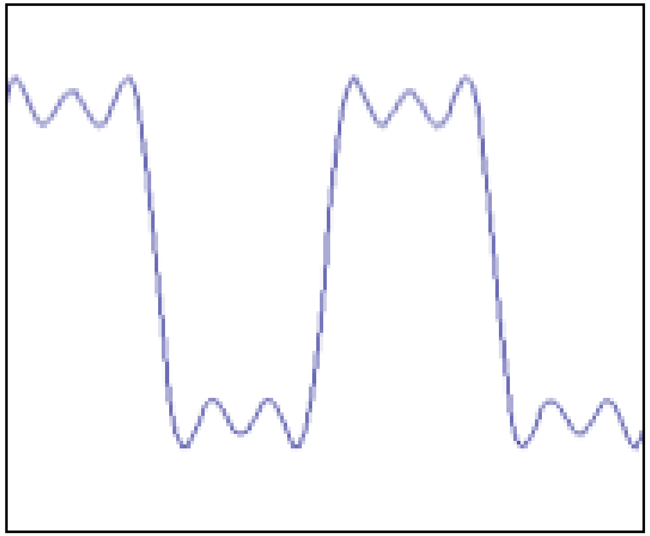

+

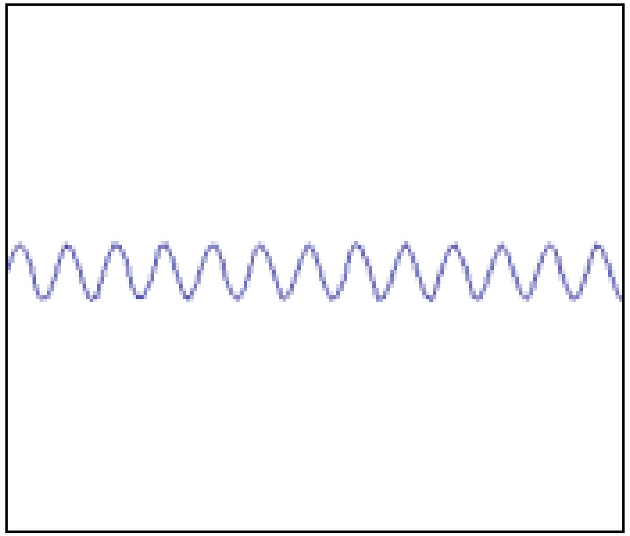

=

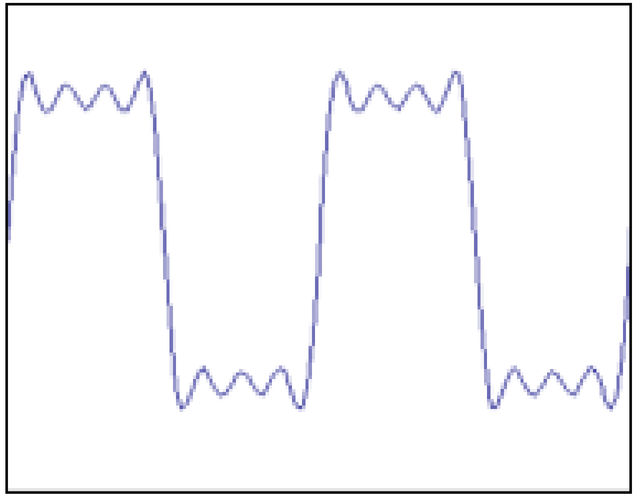

# Examples

How would you generate this function?

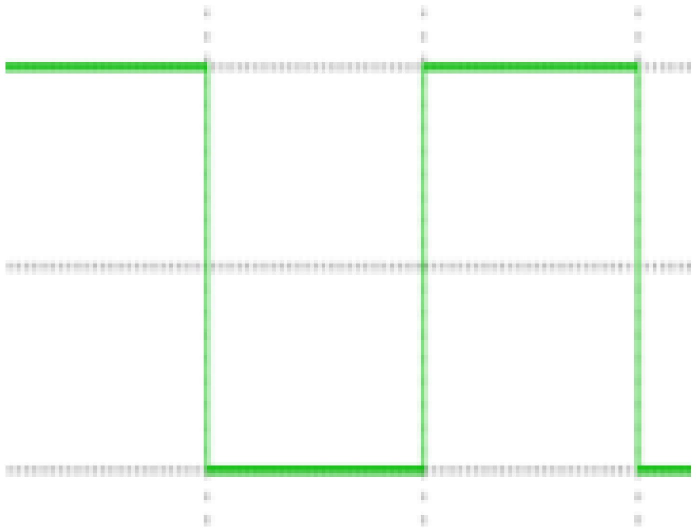

square wave

≈

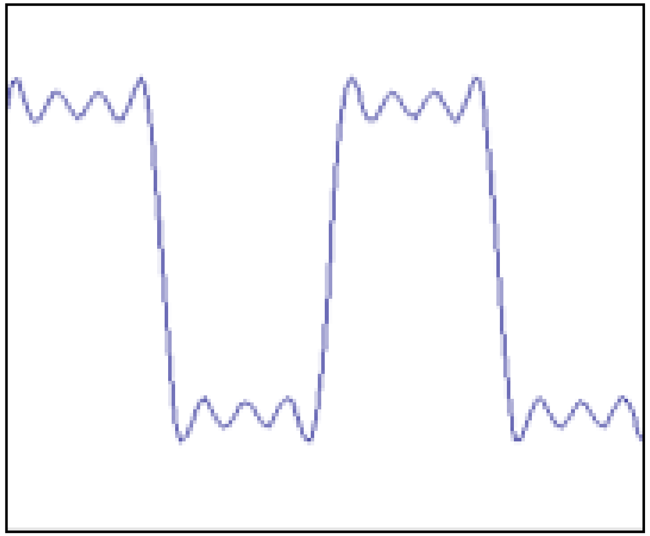

+

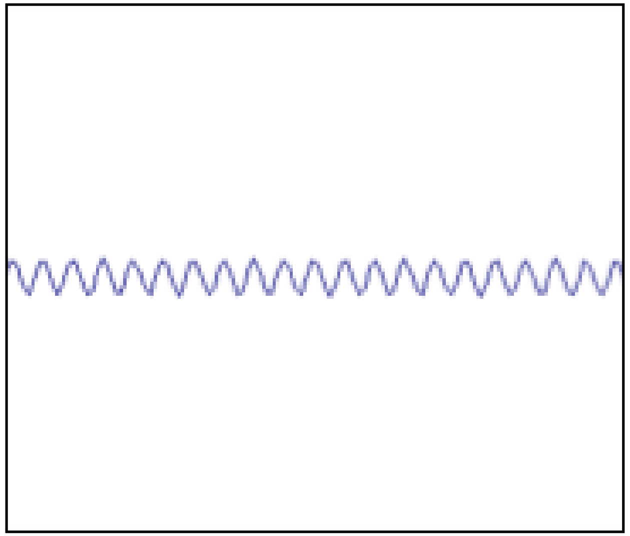

=

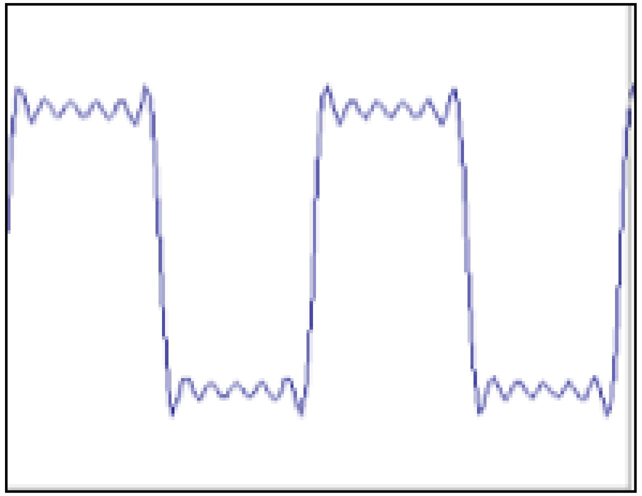

# Examples

How would you generate this function?

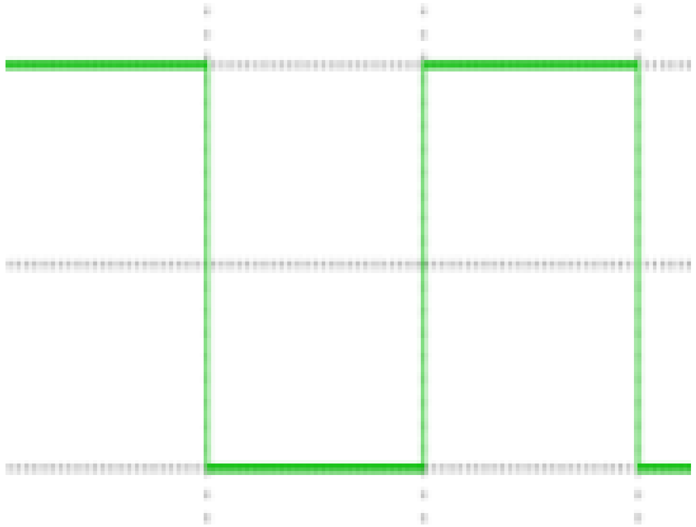

square wave

≈

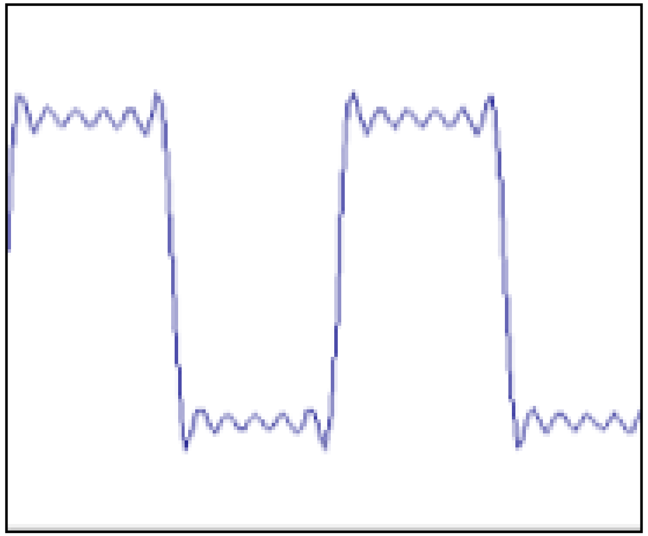

+

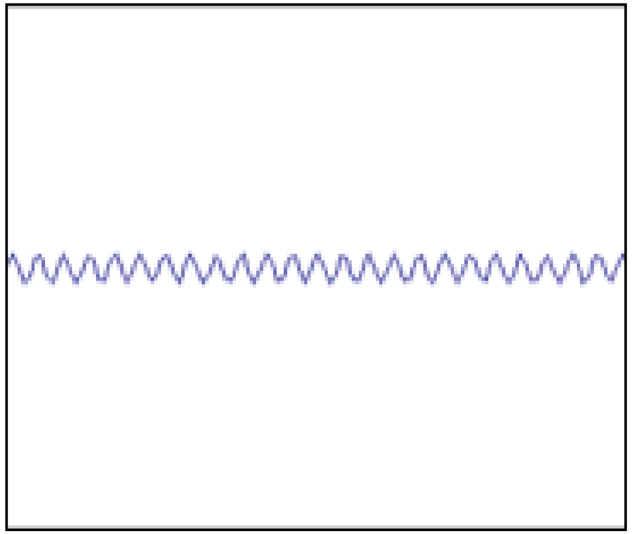

=

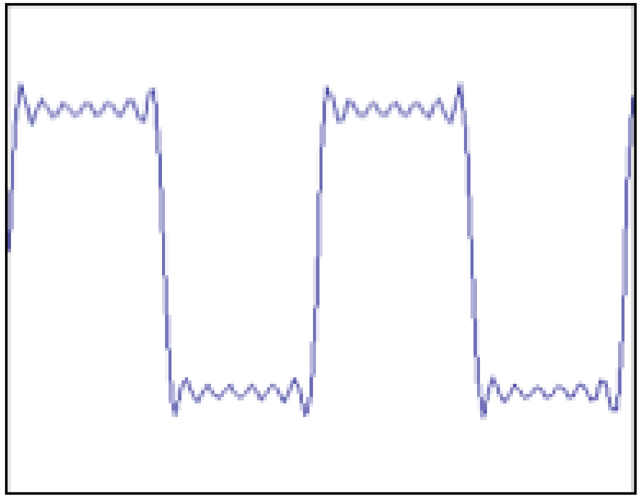

How would you express this mathematically?

# Examples

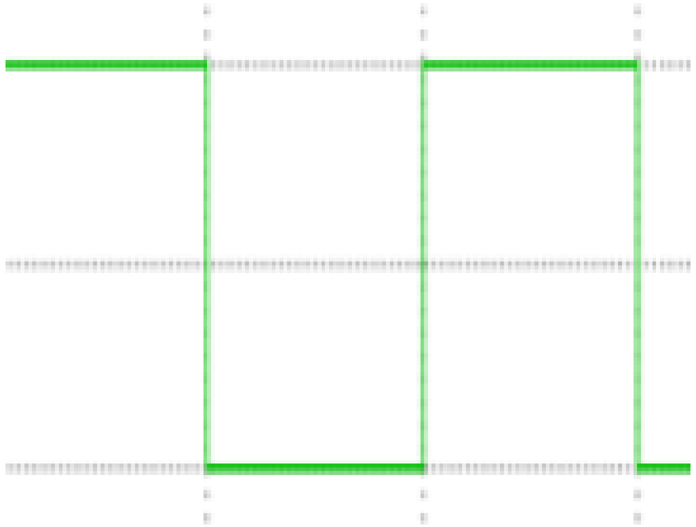

square wave

$$= A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kx)$$

infinite sum of sine waves

How would could you visualize this in the frequency domain?

# Examples

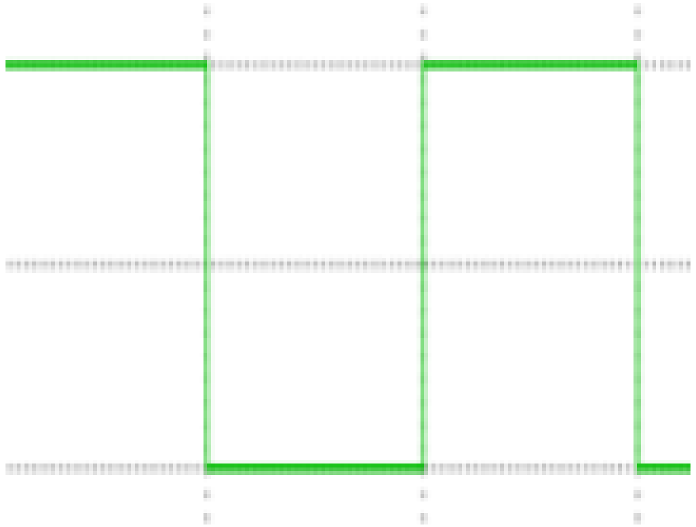

$$= \quad A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kx)$$

square wave

infinite sum of sine waves

magnitude

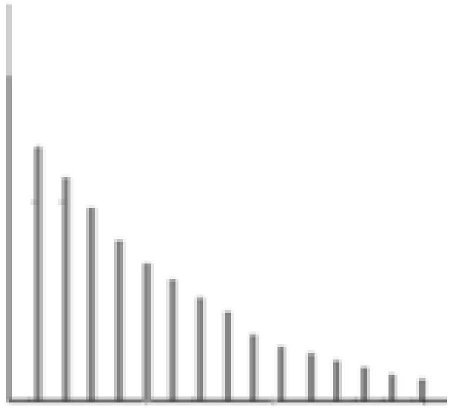

frequency

# Frequency domain

# Visualizing the frequency spectrum

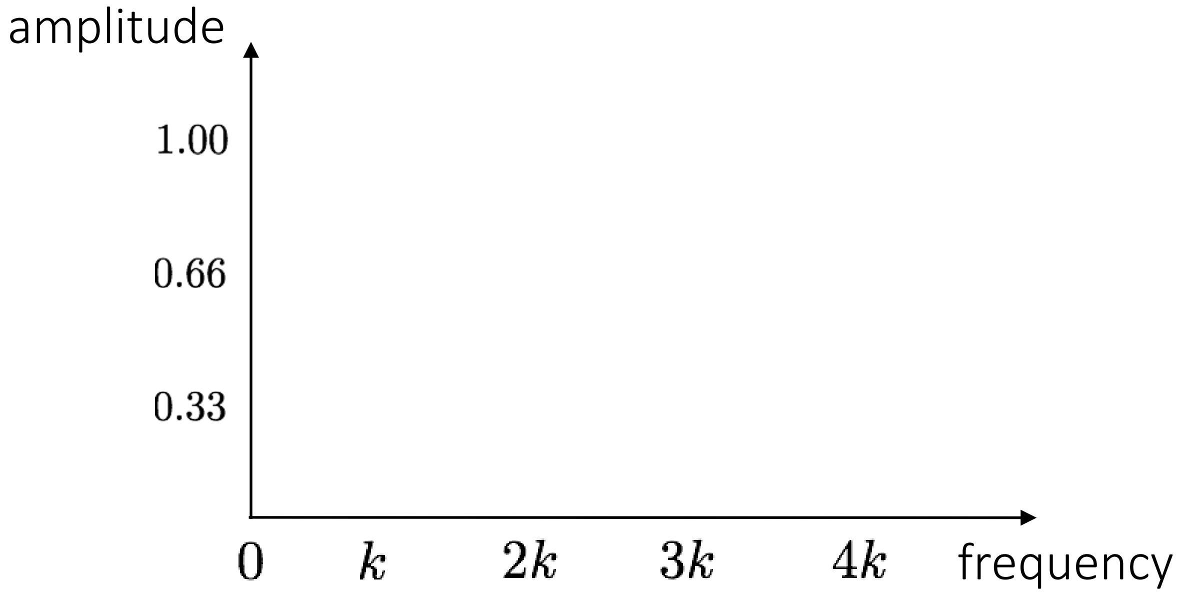

# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

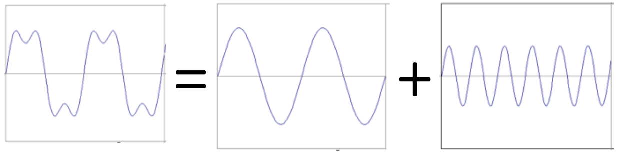

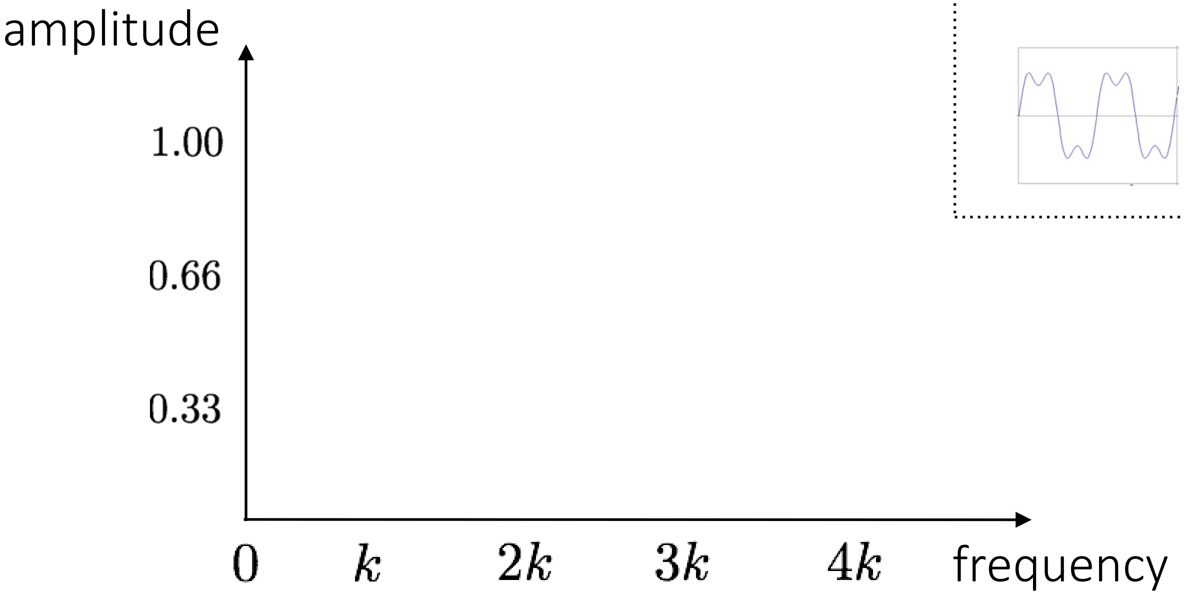

# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

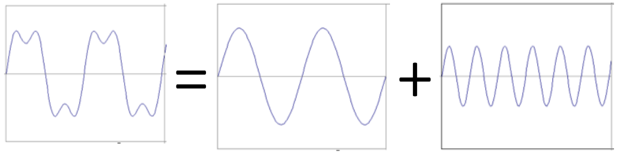


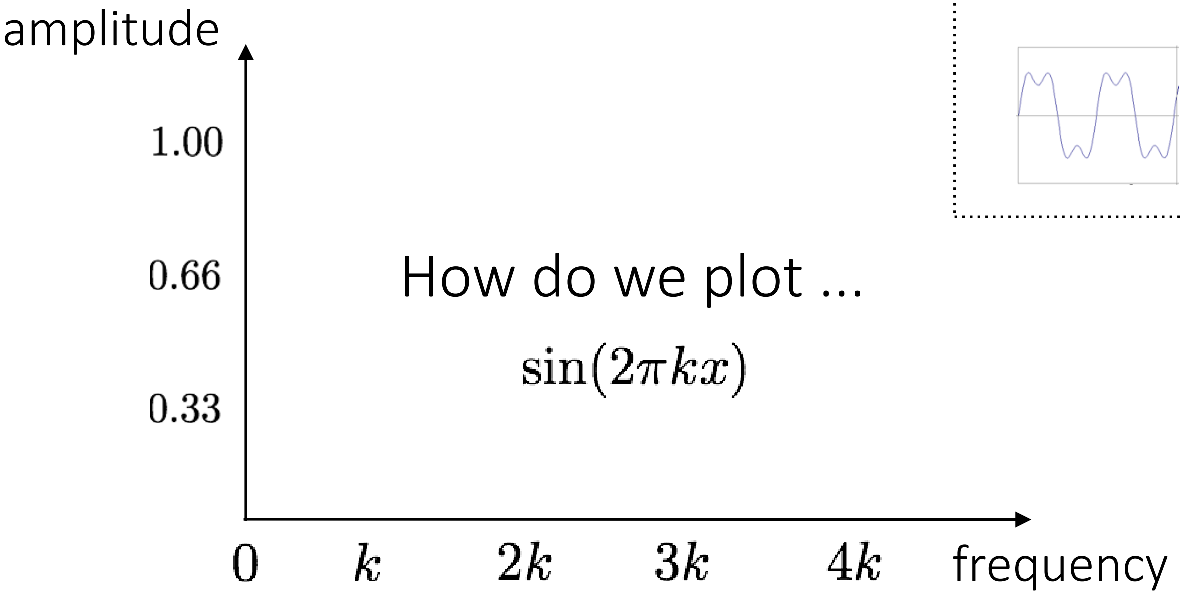


How do we plot ...

$$\sin(2\pi kx)$$

# Visualizing the frequency spectrum

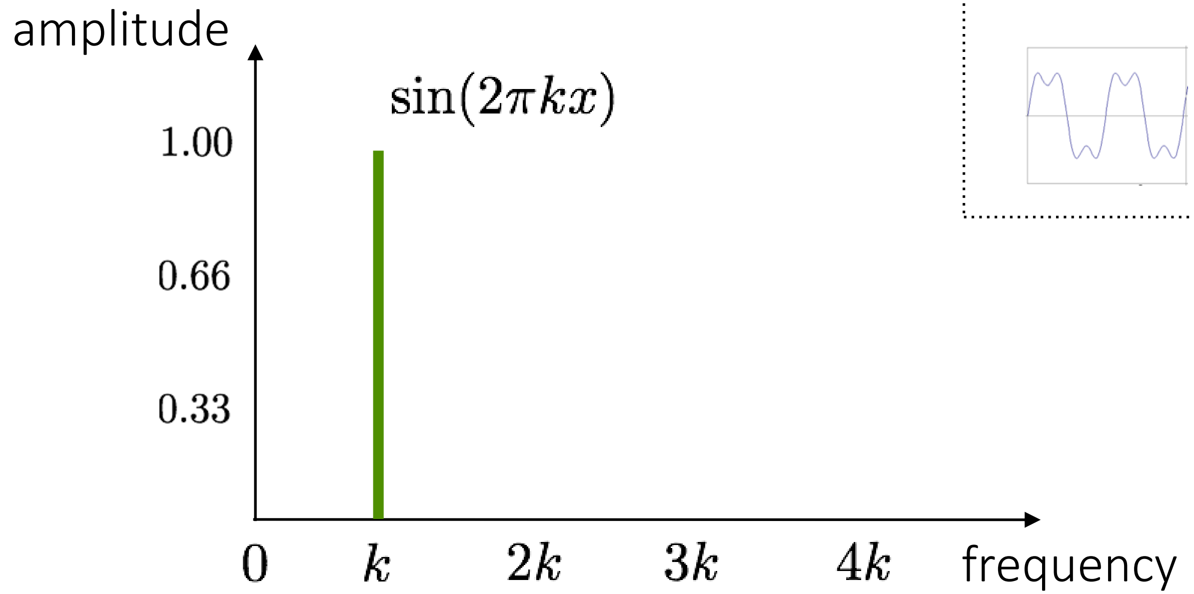


Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

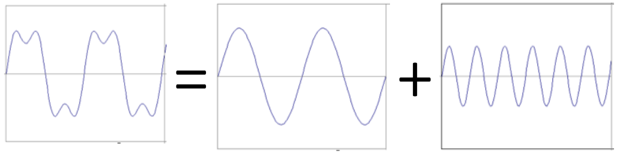

# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

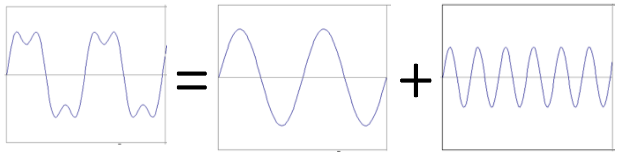


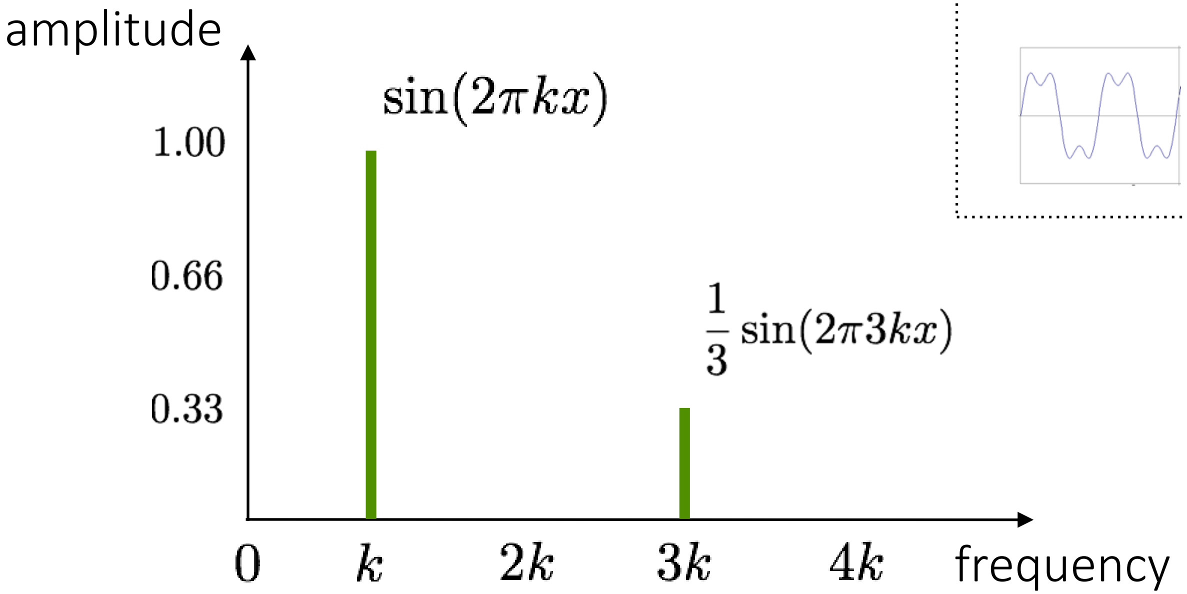

# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

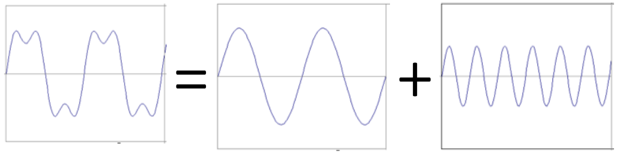

not visualizing the
symmetric negative part

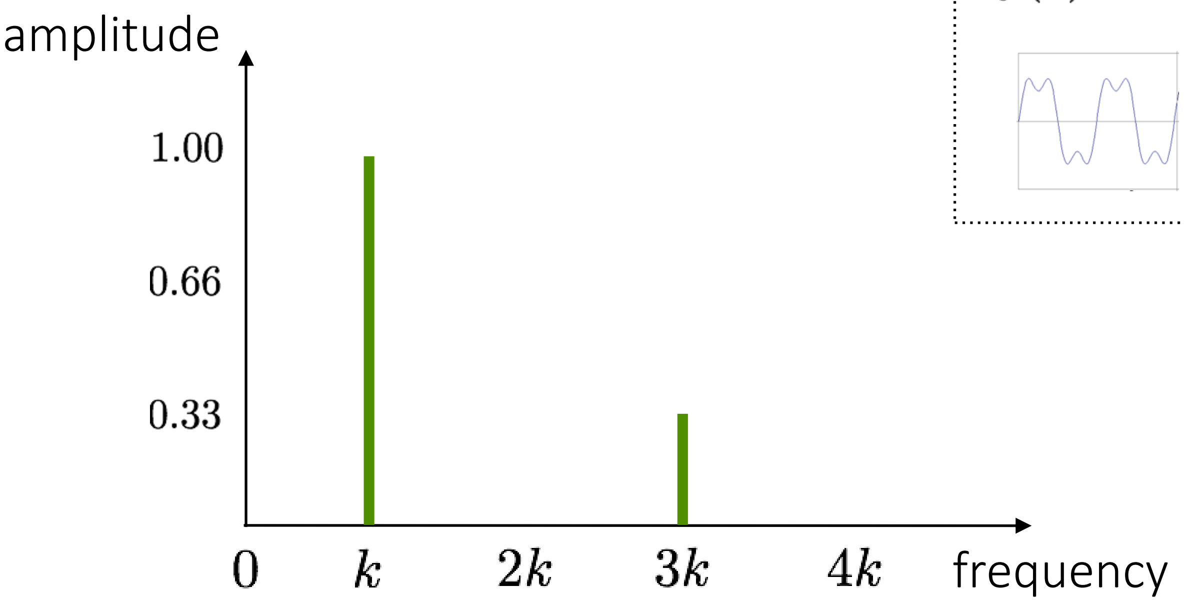


What is at zero
frequency?

Need to understand this to
understand the 2D version!

# Visualizing the frequency spectrum

Recall the temporal domain visualization

$$f(x) = \sin(2\pi kx) + \frac{1}{3}\sin(2\pi 3kx)$$

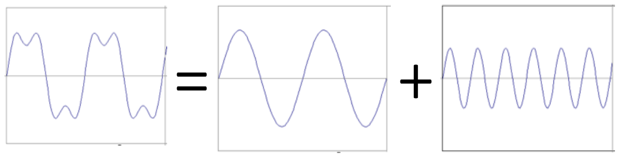

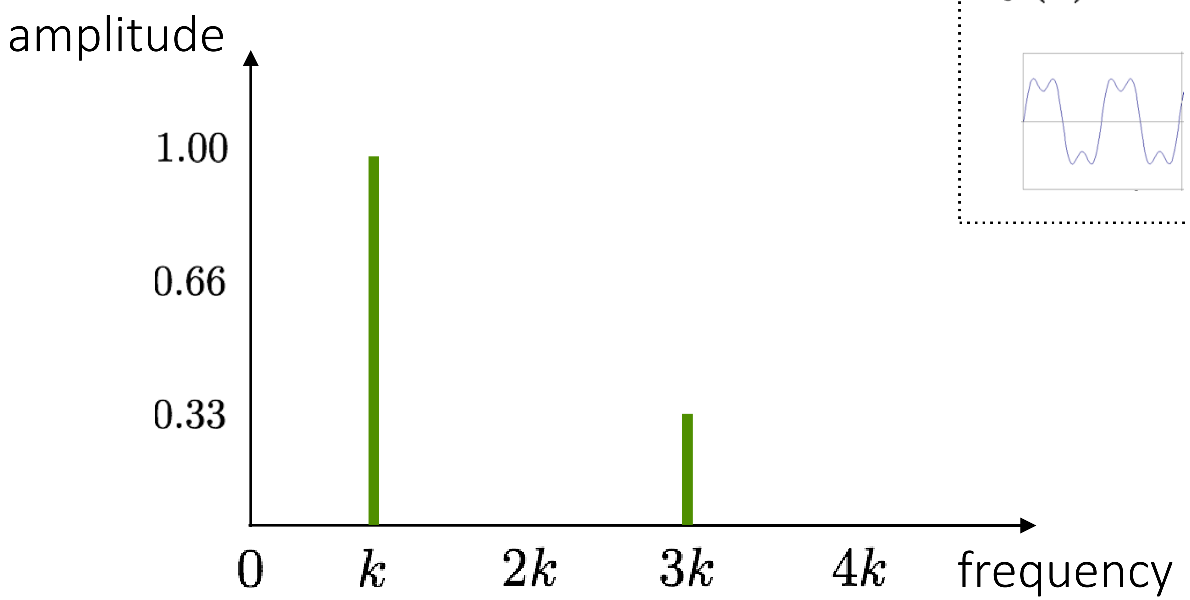


not visualizing the symmetric negative part

signal average (zero for a sine wave with no offset)

Need to understand this to understand the 2D version!

# Examples

Spatial domain visualization

Frequency domain visualization

1D

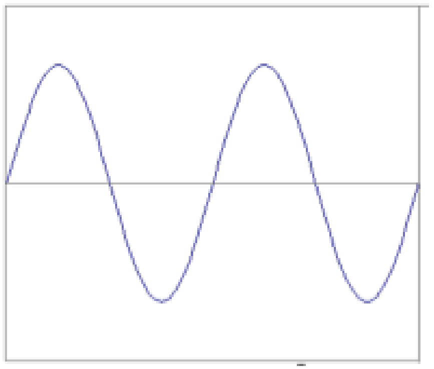

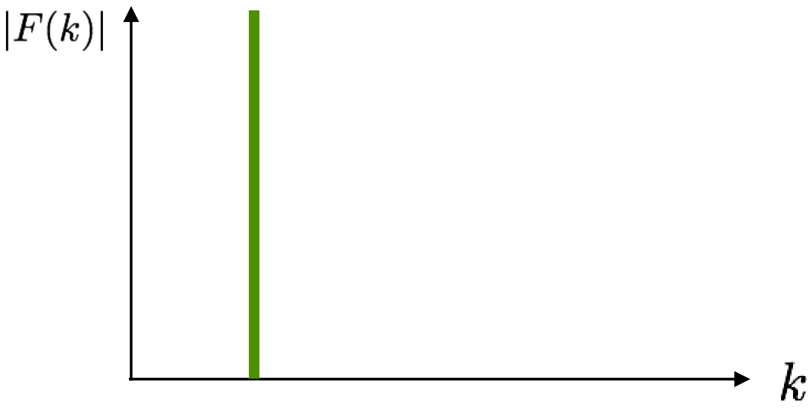


2D

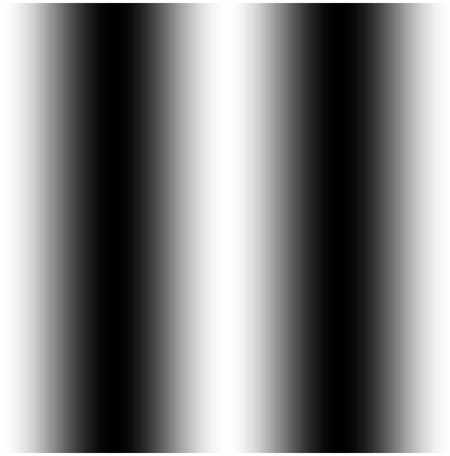

?

# Examples

Spatial domain visualization

Frequency domain visualization

1D

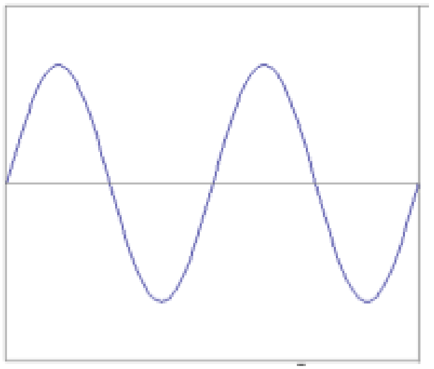

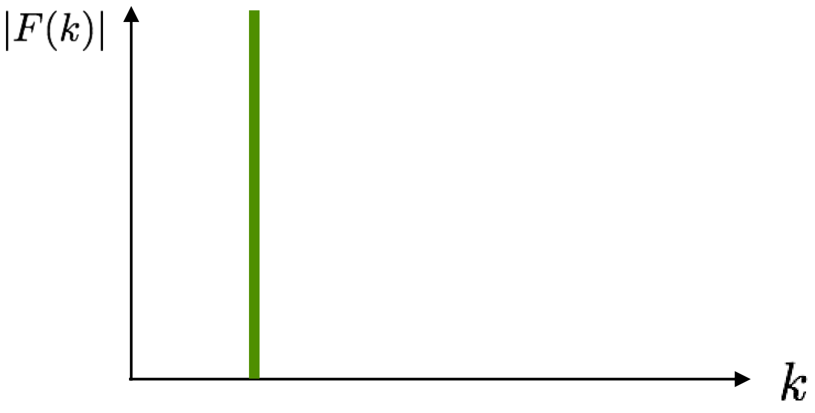


2D

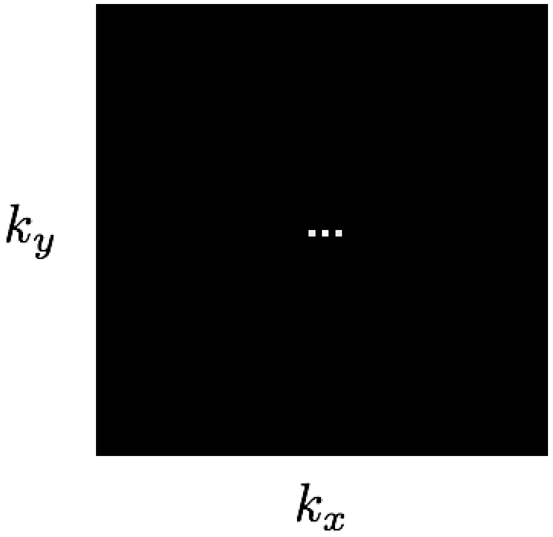


What do the three dots correspond to?

# Examples

Spatial domain visualization

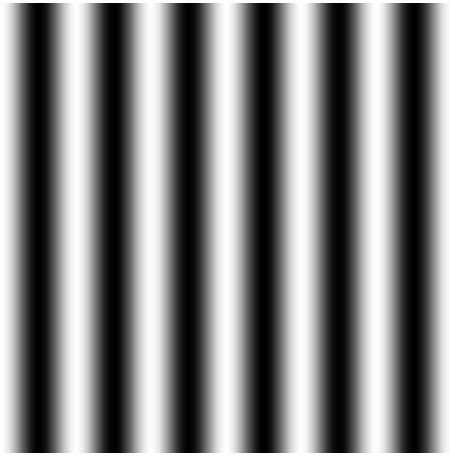

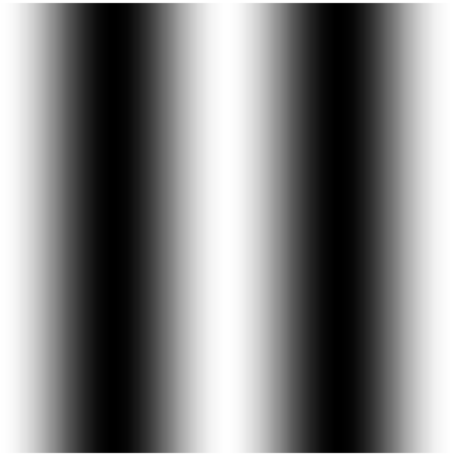

Frequency domain visualization

?

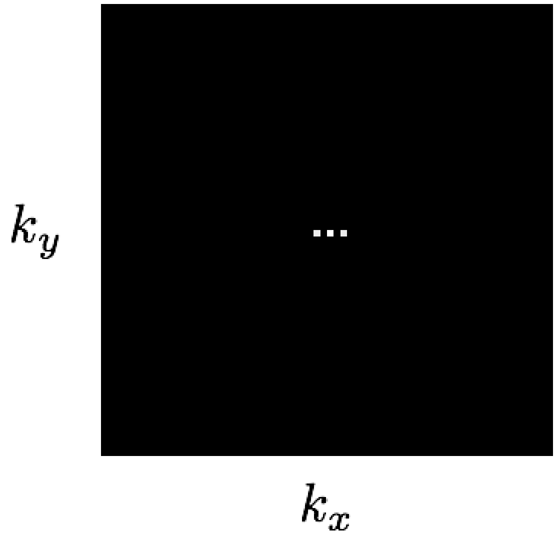

# Examples

Spatial domain visualization

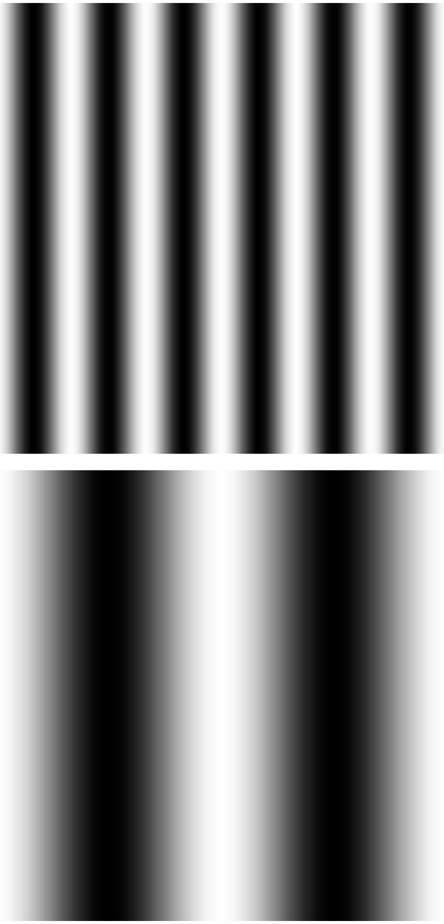

Frequency domain visualization

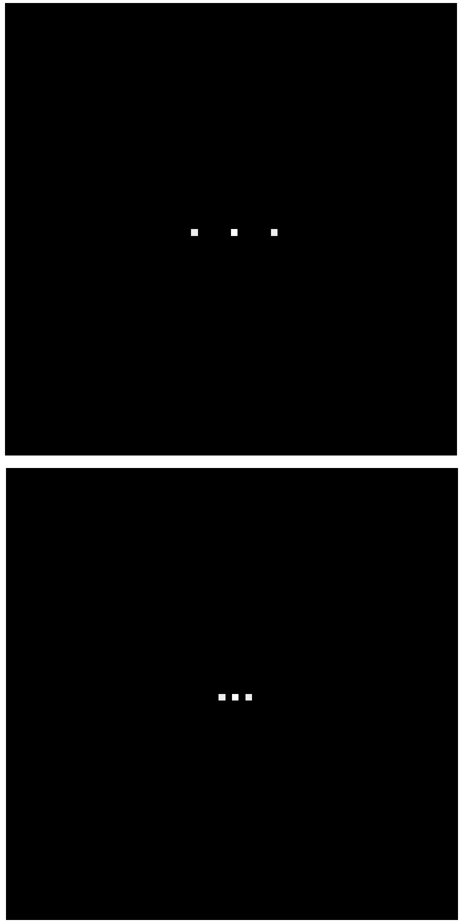

$k_y$

$k_x$

# Examples

How would you generate this image with sine waves?

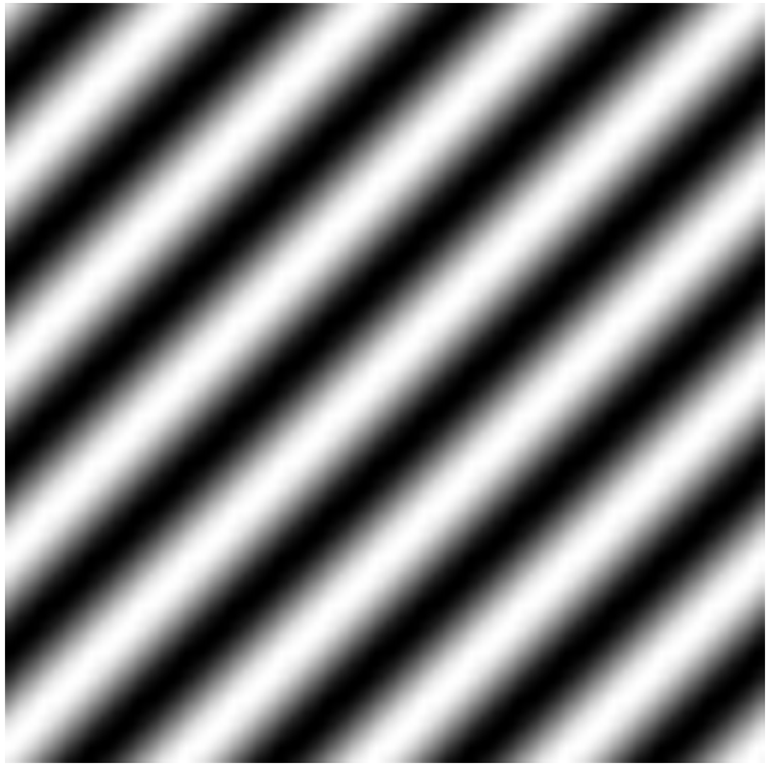

# Examples

How would you generate this image with sine waves?

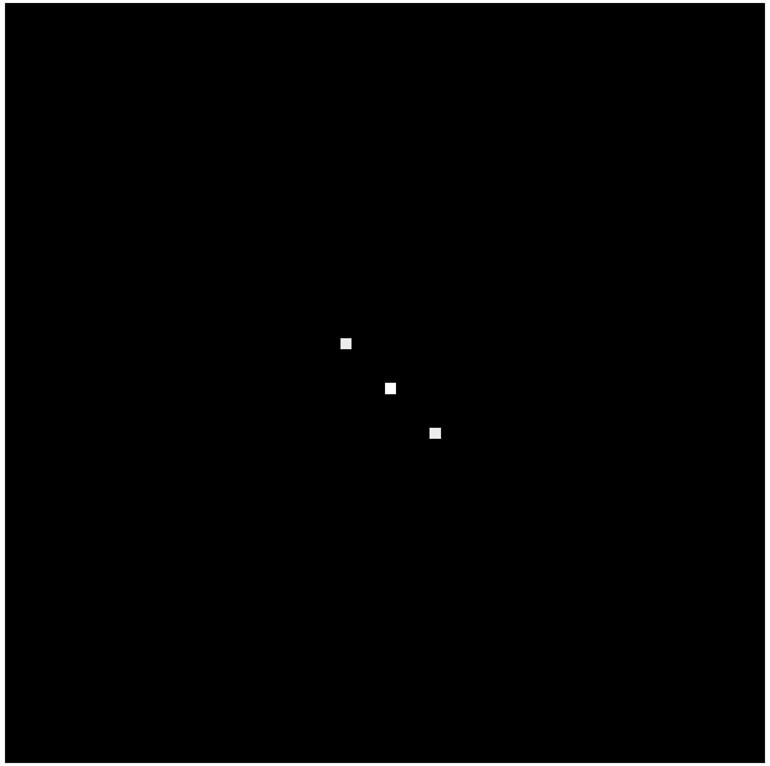

Has both an x and
y components

# Examples

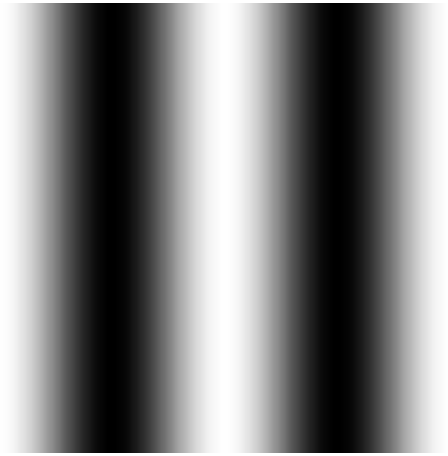 + 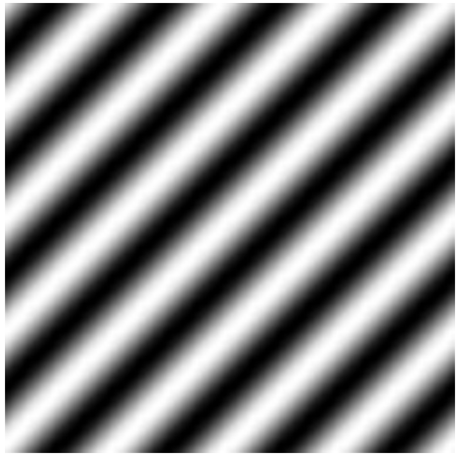 = ?

# Examples

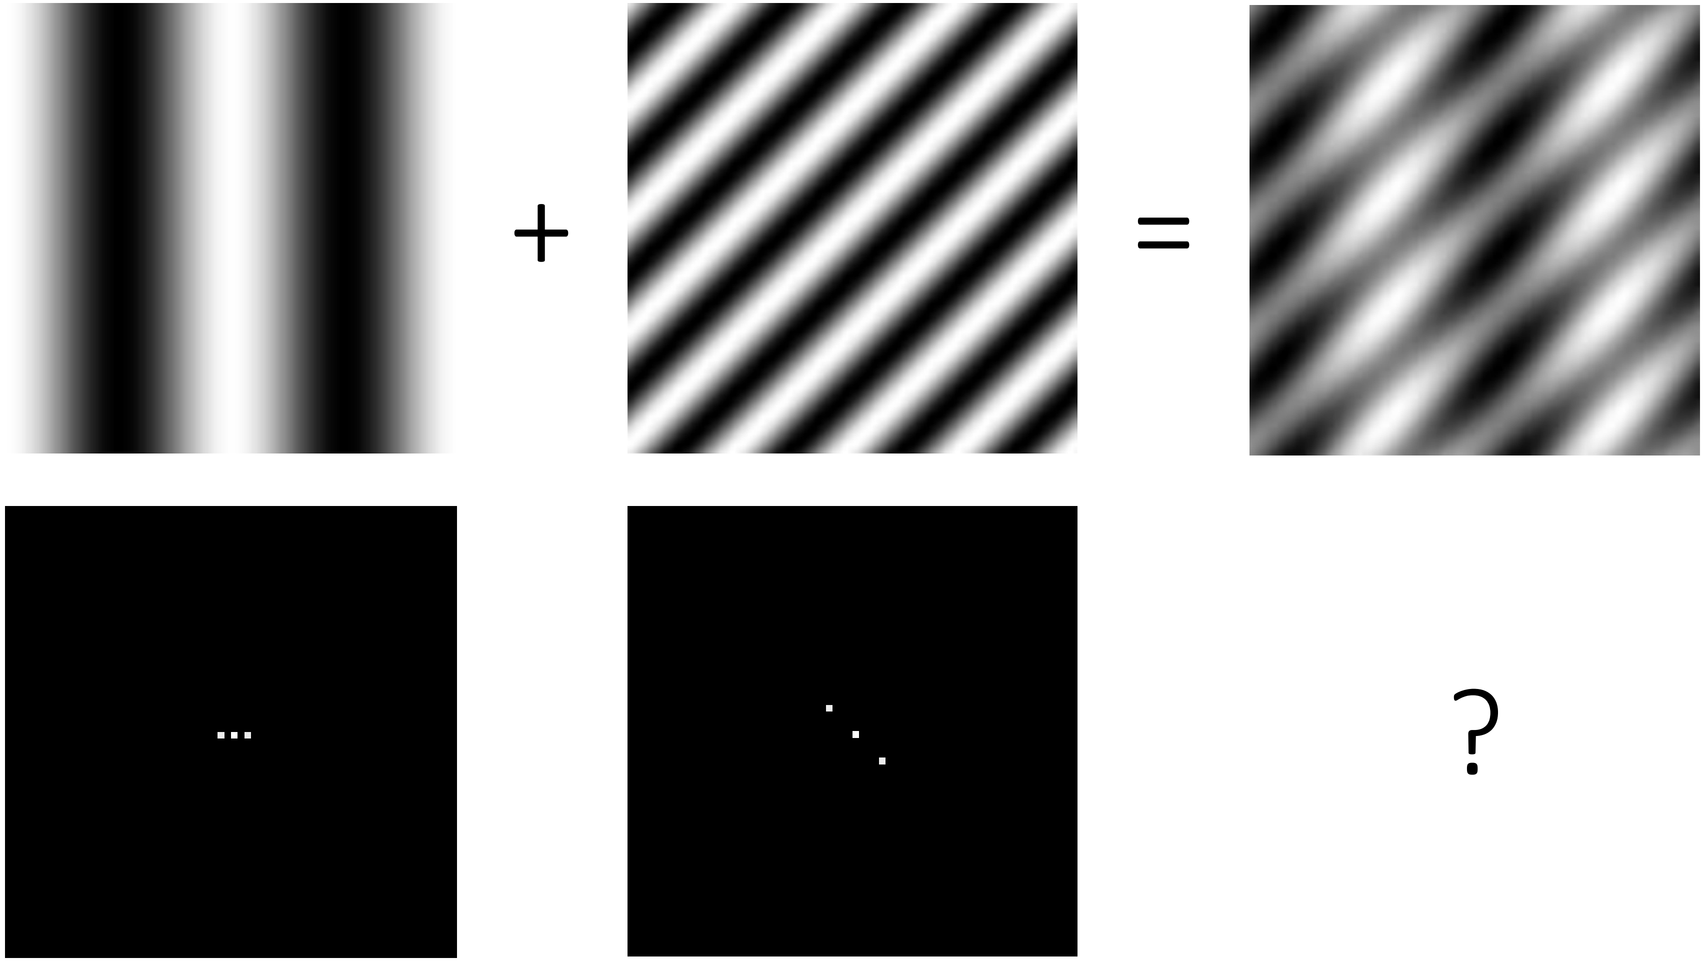

# Examples

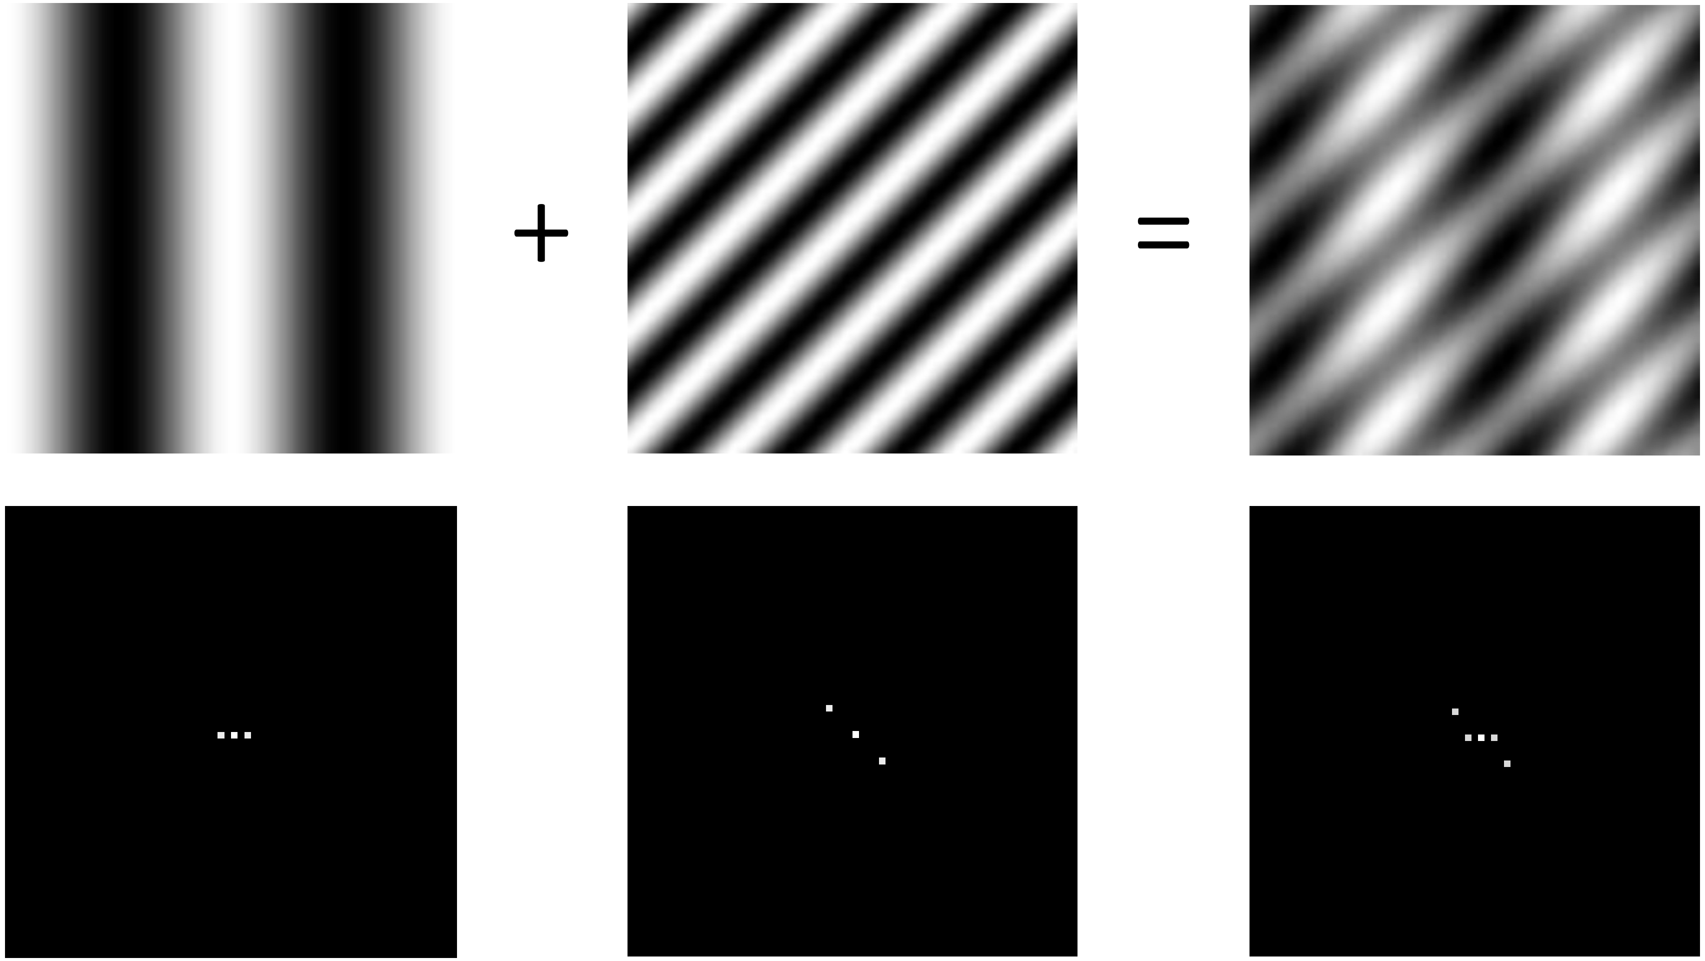

# Basic building block

$$A \sin(\omega x + \phi)$$

amplitude

sinusoid

angular frequency

variable

phase

What about non-periodic signals?

Fourier's claim: Add enough of these to get <u>any</u> *periodic* signal you want!

# Fourier transform

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$R + jI$$

what‘s this? what‘s this?

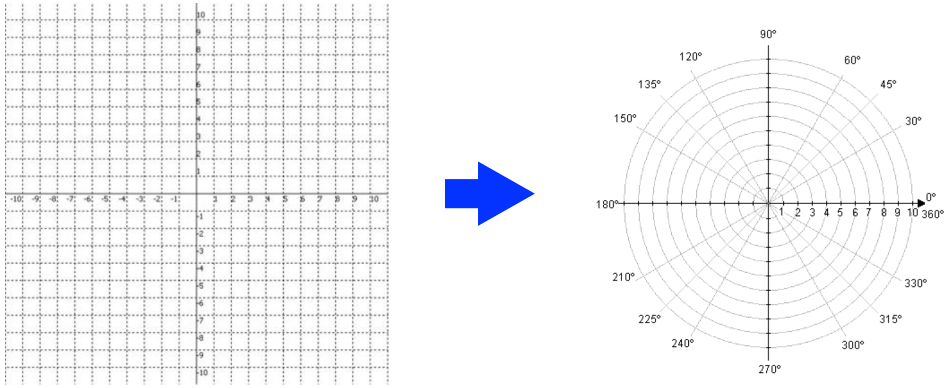

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$\underbrace{R}_{\text{real}} + \underbrace{jI}_{\text{imaginary}}$$

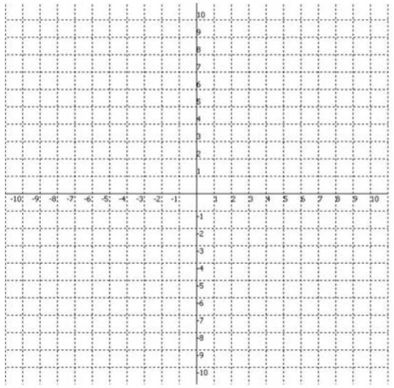

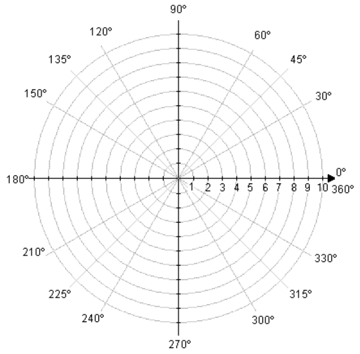

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$\underbrace{R}_{\text{real}} + \underbrace{jI}_{\text{imaginary}}$$

Alternative reparameterization:

polar coordinates

$$r(\cos\theta + j\sin\theta)$$

how do we compute these?

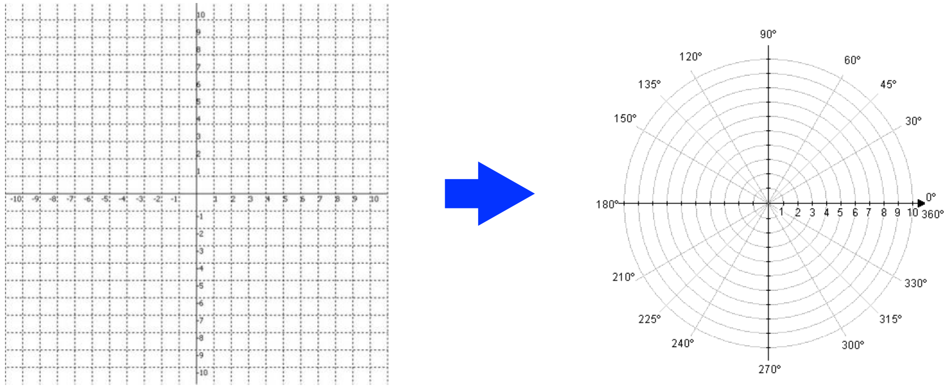


polar transform

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$R + jI$$

real imaginary

Alternative reparameterization:

polar coordinates

$$r(\cos\theta + j\sin\theta)$$

polar transform

$$\theta = \tan^{-1}\left(\frac{I}{R}\right) \quad r = \sqrt{R^2 + I^2}$$

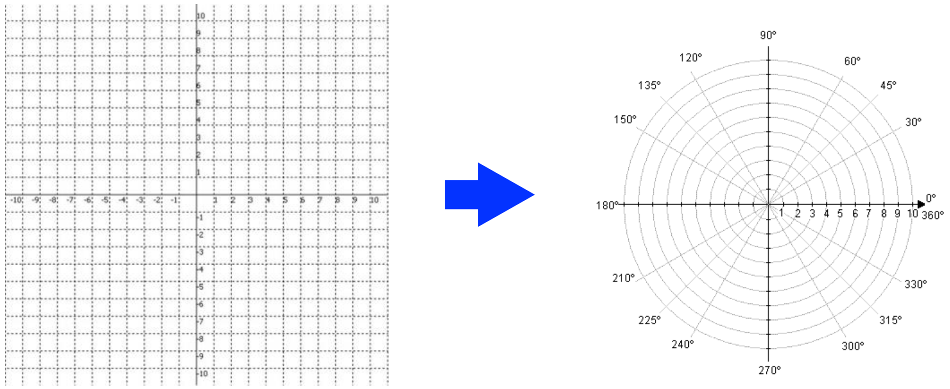


polar transform

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$R + jI$$

real imaginary

Alternative reparameterization:

polar coordinates

$$r(\cos\theta + j\sin\theta)$$

polar transform

$$\theta = \tan^{-1}\left(\frac{I}{R}\right) \quad r = \sqrt{R^2 + I^2}$$

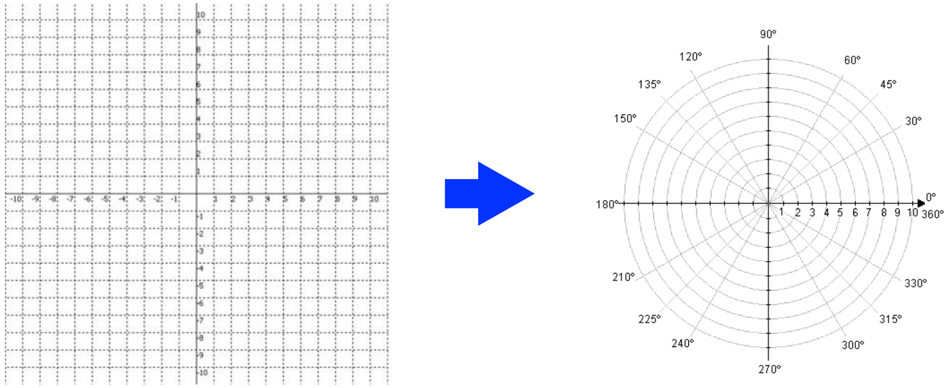


polar transform

How do you write these in exponential form?

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$R + jI$$

real imaginary

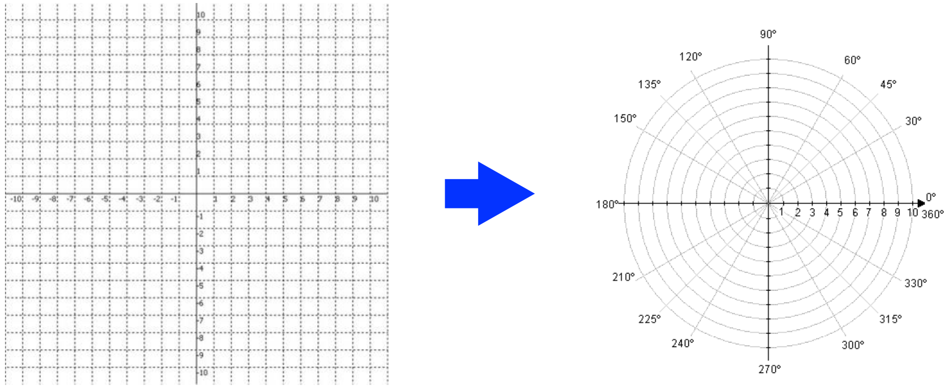

Alternative reparameterization:

polar coordinates

$$r(\cos\theta + j\sin\theta)$$

polar transform

$$\theta = \tan^{-1}\left(\frac{I}{R}\right) \quad r = \sqrt{R^2 + I^2}$$

or equivalently

$$re^{j\theta}$$

how did we get this?

exponential form

# Recalling some basics

Complex numbers have two parts:

rectangular coordinates

$$R + jI$$

real imaginary

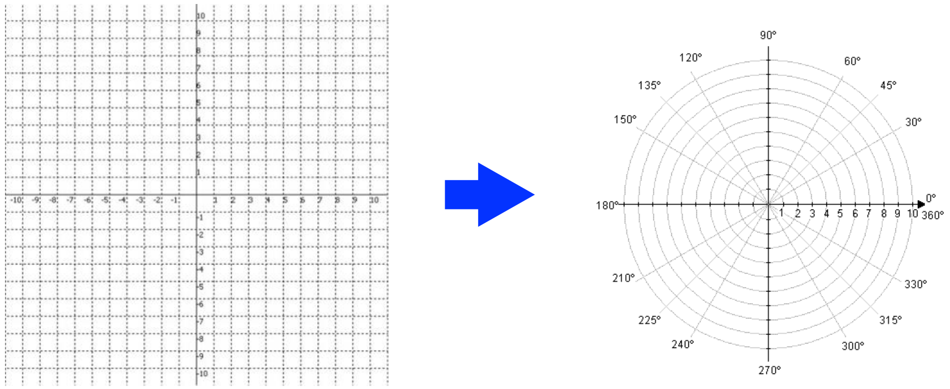


Alternative reparameterization:

polar coordinates

$$r(\cos\theta + j\sin\theta)$$

polar transform

$$\theta = \tan^{-1}\left(\frac{I}{R}\right) \quad r = \sqrt{R^2 + I^2}$$

or equivalently

$$re^{j\theta}$$

Euler's formula

$$e^{j\theta} = \cos\theta + j\sin\theta$$

exponential form

This will help us understand the Fourier transform equations

# Fourier transform

| | Fourier transform | inverse Fourier transform |
|---|---|---|
| continuous | $F(k) = \int_{\infty}^{-\infty} f(x) e^{-j2\pi kx} dx$ | $f(x) = \int_{\infty}^{-\infty} F(k) e^{j2\pi kx} dk$ |
| discrete | $F(k) = \frac{1}{N} \sum_{x=0}^{N-1} f(x) e^{-j2\pi kx/N}$ <br> $k = 0, 1, 2, \ldots, N-1$ | $f(x) = \sum_{k=0}^{N-1} F(k) e^{j2\pi kx/N}$ <br> $x = 0, 1, 2, \ldots, N-1$ |

Where is the connection to the 'summation of sine waves' idea?

# Fourier transform

| | Fourier transform | inverse Fourier transform |
|---|---|---|
| continuous | $F(k) = \int_{\infty}^{-\infty} f(x) e^{-j2\pi kx} dx$ | $f(x) = \int_{\infty}^{-\infty} F(k) e^{j2\pi kx} dk$ |
| discrete | $F(k) = \frac{1}{N} \sum_{x=0}^{N-1} f(x) e^{-j2\pi kx/N}$ <br> $k = 0, 1, 2, \ldots, N-1$ | $\boxed{f(x) = \sum_{k=0}^{N-1} F(k) e^{j2\pi kx/N}}$ <br> $x = 0, 1, 2, \ldots, N-1$ |

Where is the connection to the ‘summation of sine waves’ idea?

# Fourier transform

Where is the connection to the 'summation of sine waves' idea?

$$f(x) = \sum_{k=0}^{N-1} F(k) e^{j2\pi kx/N}$$

Euler's formula

$$e^{j\theta} = \cos\theta + j\sin\theta$$

sum over frequencies

$$f(x) = \sum_{k=0}^{N-1} F(k) \left\{ \cos(2\pi kx) + j\sin(2\pi kx) \right\}$$

scaling parameter

wave components

# Fourier transform pairs

spatial domain

frequency domain

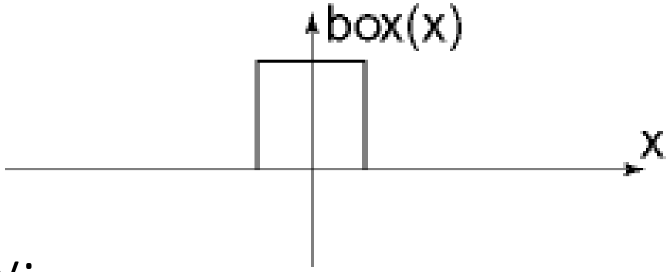


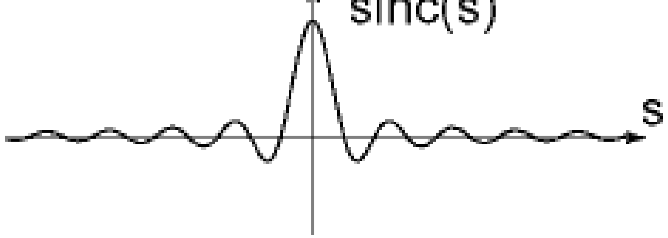


Note the symmetry: duality property of Fourier transform

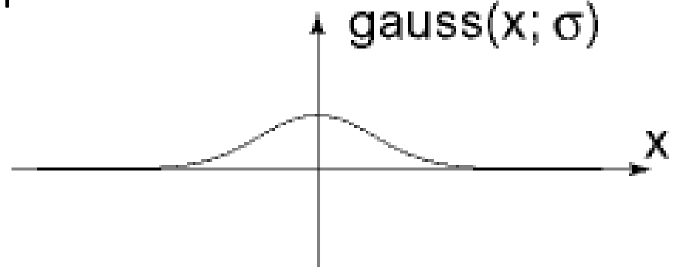


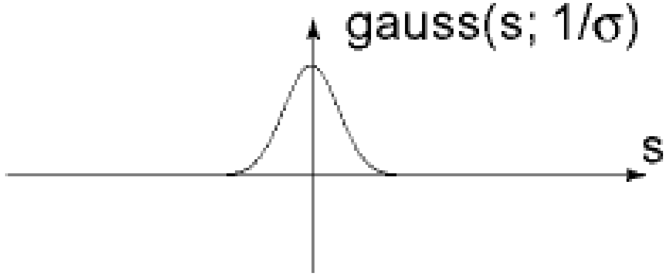


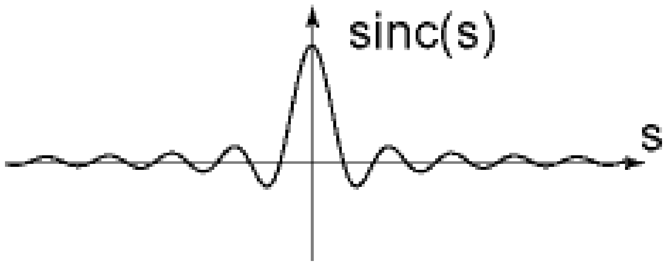


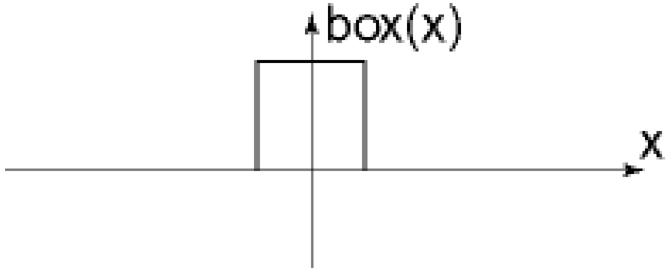

# Computing the discrete Fourier transform (DFT)

# Computing the discrete Fourier transform (DFT)

$$F(k) = \frac{1}{N}\sum_{x=0}^{N-1} f(x) e^{-j2\pi kx/N}$$ is just a matrix multiplication:

$$\boldsymbol{F} = \boldsymbol{W}\boldsymbol{f}$$

$$\begin{bmatrix} F(0) \\ F(1) \\ F(2) \\ F(3) \\ \vdots \\ F(N-1) \end{bmatrix} = \begin{bmatrix} W^0 & W^0 & W^0 & W^0 & \cdots & W^0 \\ W^0 & W^1 & W^2 & W^3 & \cdots & W^{N-1} \\ W^0 & W^2 & W^4 & W^6 & \cdots & W^{N-2} \\ W^0 & W^3 & W^6 & W^9 & \cdots & W^{N-3} \\ \vdots & & & & \ddots & \vdots \\ W^0 & W^{N-1} & W^{N-2} & W^{N-3} & \cdots & W^1 \end{bmatrix} \begin{bmatrix} f(0) \\ f(1) \\ f(2) \\ f(3) \\ \vdots \\ f(N-1) \end{bmatrix} \qquad W = e^{-j2\pi/N}$$

In practice this is implemented using the *fast Fourier transform* (FFT) algorithm.

# Another way to compute the Fourier transform

Use a lens!

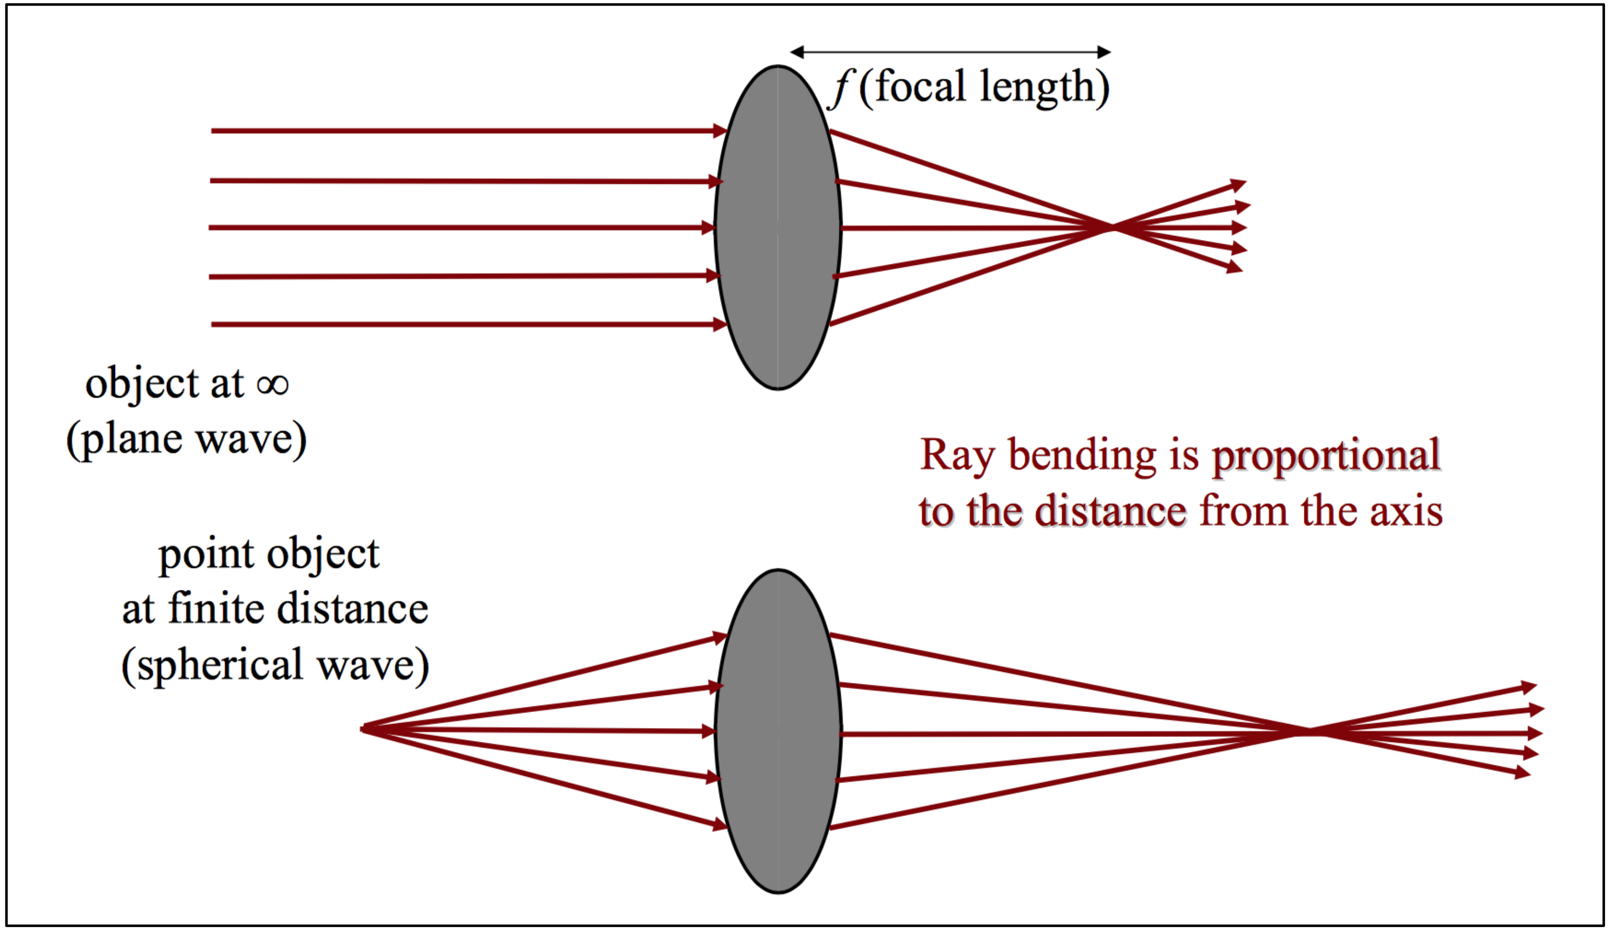


An ideal thin lens is an optical Fourier transform engine.

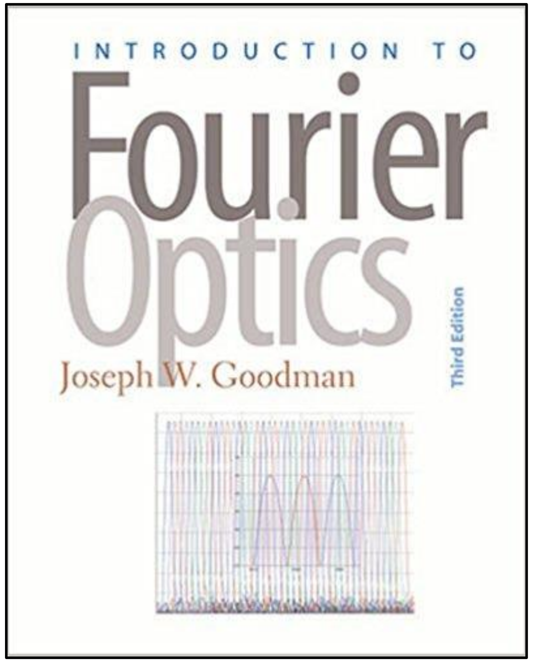

# Fourier transforms of natural images

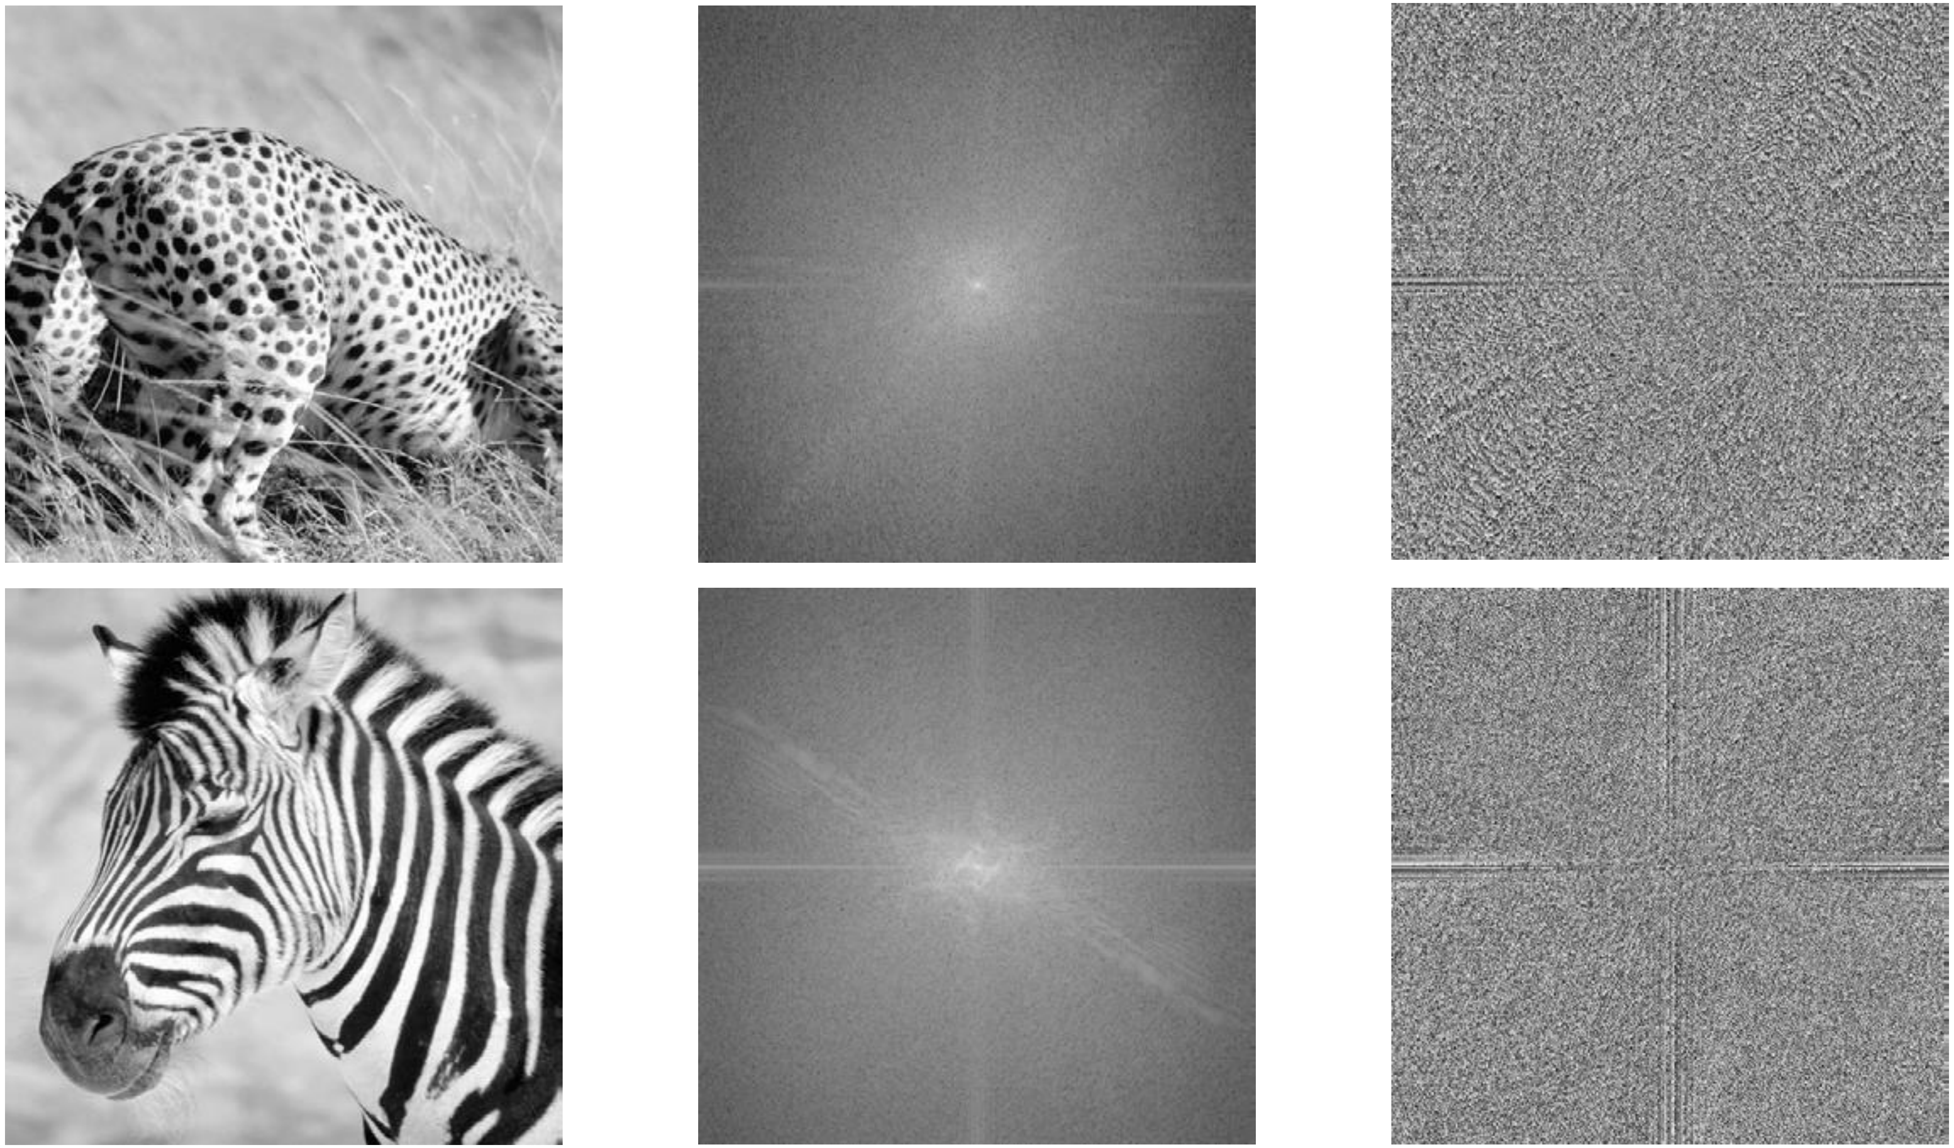

original amplitude phase

# Fourier transforms of natural images

Image phase matters!

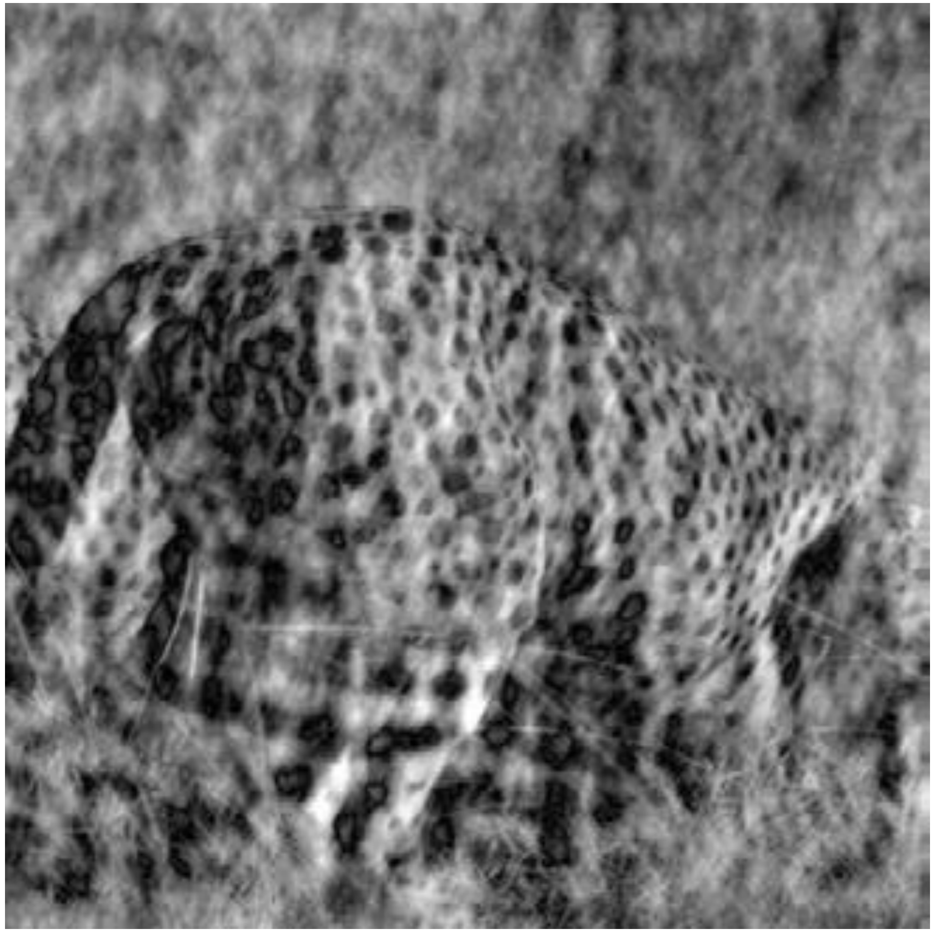

cheetah phase with zebra amplitude

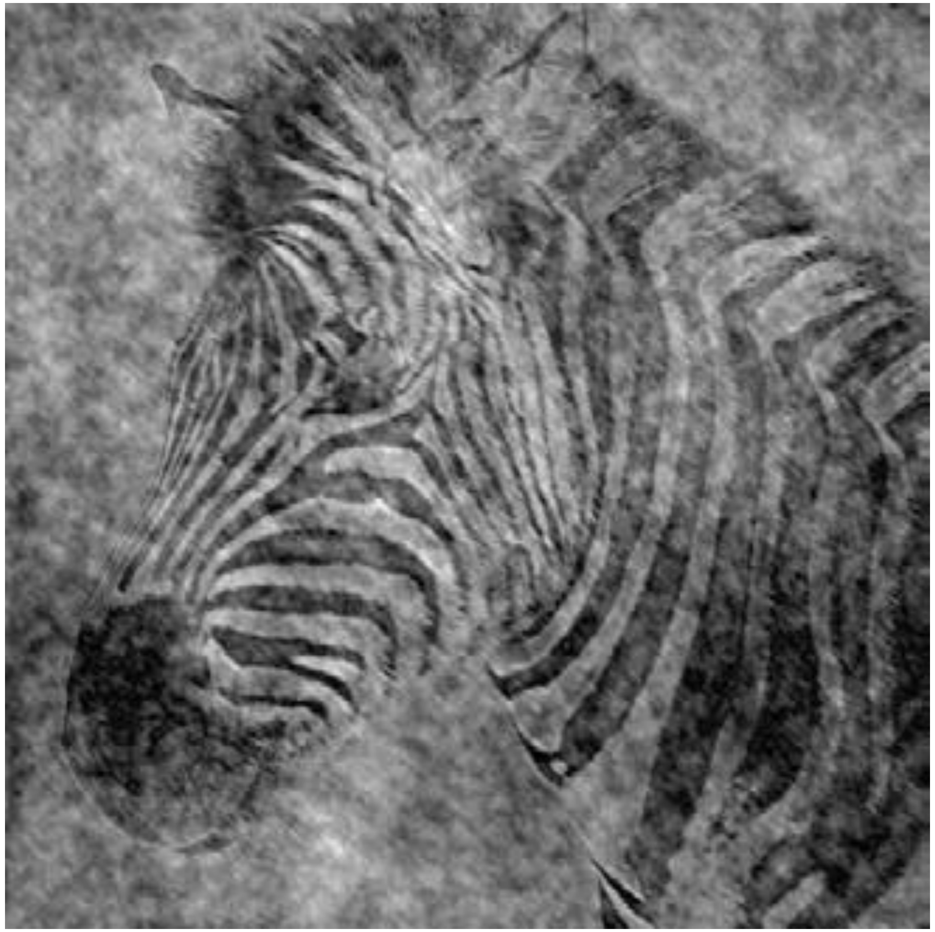

zebra phase with cheetah amplitude

# Frequency-domain filtering

# Why do we care about all this?

# The convolution theorem

The Fourier transform of the convolution of two functions is the product of their Fourier transforms:

$$\mathcal{F}\{g * h\} = \mathcal{F}\{g\}\mathcal{F}\{h\}$$

The inverse Fourier transform of the product of two Fourier transforms is the convolution of the two inverse Fourier transforms:

$$\mathcal{F}^{-1}\{gh\} = \mathcal{F}^{-1}\{g\} * \mathcal{F}^{-1}\{h\}$$

Convolution in spatial domain is equivalent to multiplication in frequency domain!

# What do we use convolution for?

# Convolution for 1D continuous signals

Definition of linear shift-invariant filtering as convolution:

$$(f * g)(x) = \int_{-\infty}^{\infty} f(y) g(x - y) dy$$

filtered signal

filter

input signal

Using the convolution theorem, we can interpret and implement all types of linear shift-invariant filtering as multiplication in frequency domain.

Why implement convolution in frequency domain?

# Frequency-domain filtering in Matlab

Filtering with `fft`:

```
im = double(imread('…'))/255;
im = rgb2gray(im); % "im" should be a gray-scale floating point image
[imh, imw] = size(im);

hs = 50; % filter half-size
fil = fspecial('gaussian', hs*2+1, 10);

fftsize = 1024; % should be order of 2 (for speed) and include padding
im_fft  = fft2(im,  fftsize, fftsize);                    % 1) fft im with
padding
fil_fft = fft2(fil, fftsize, fftsize);                    % 2) fft fil, pad to
same size as image
im_fil_fft = im_fft .* fil_fft;                           % 3) multiply fft
images
im_fil = ifft2(im_fil_fft);                               % 4) inverse fft2
im_fil = im_fil(1+hs:size(im,1)+hs, 1+hs:size(im, 2)+hs); % 5) remove padding
```

Displaying with `fft`:

```
figure(1), imagesc(log(abs(fftshift(im_fft)))), axis image, colormap jet
```

# Spatial domain filtering

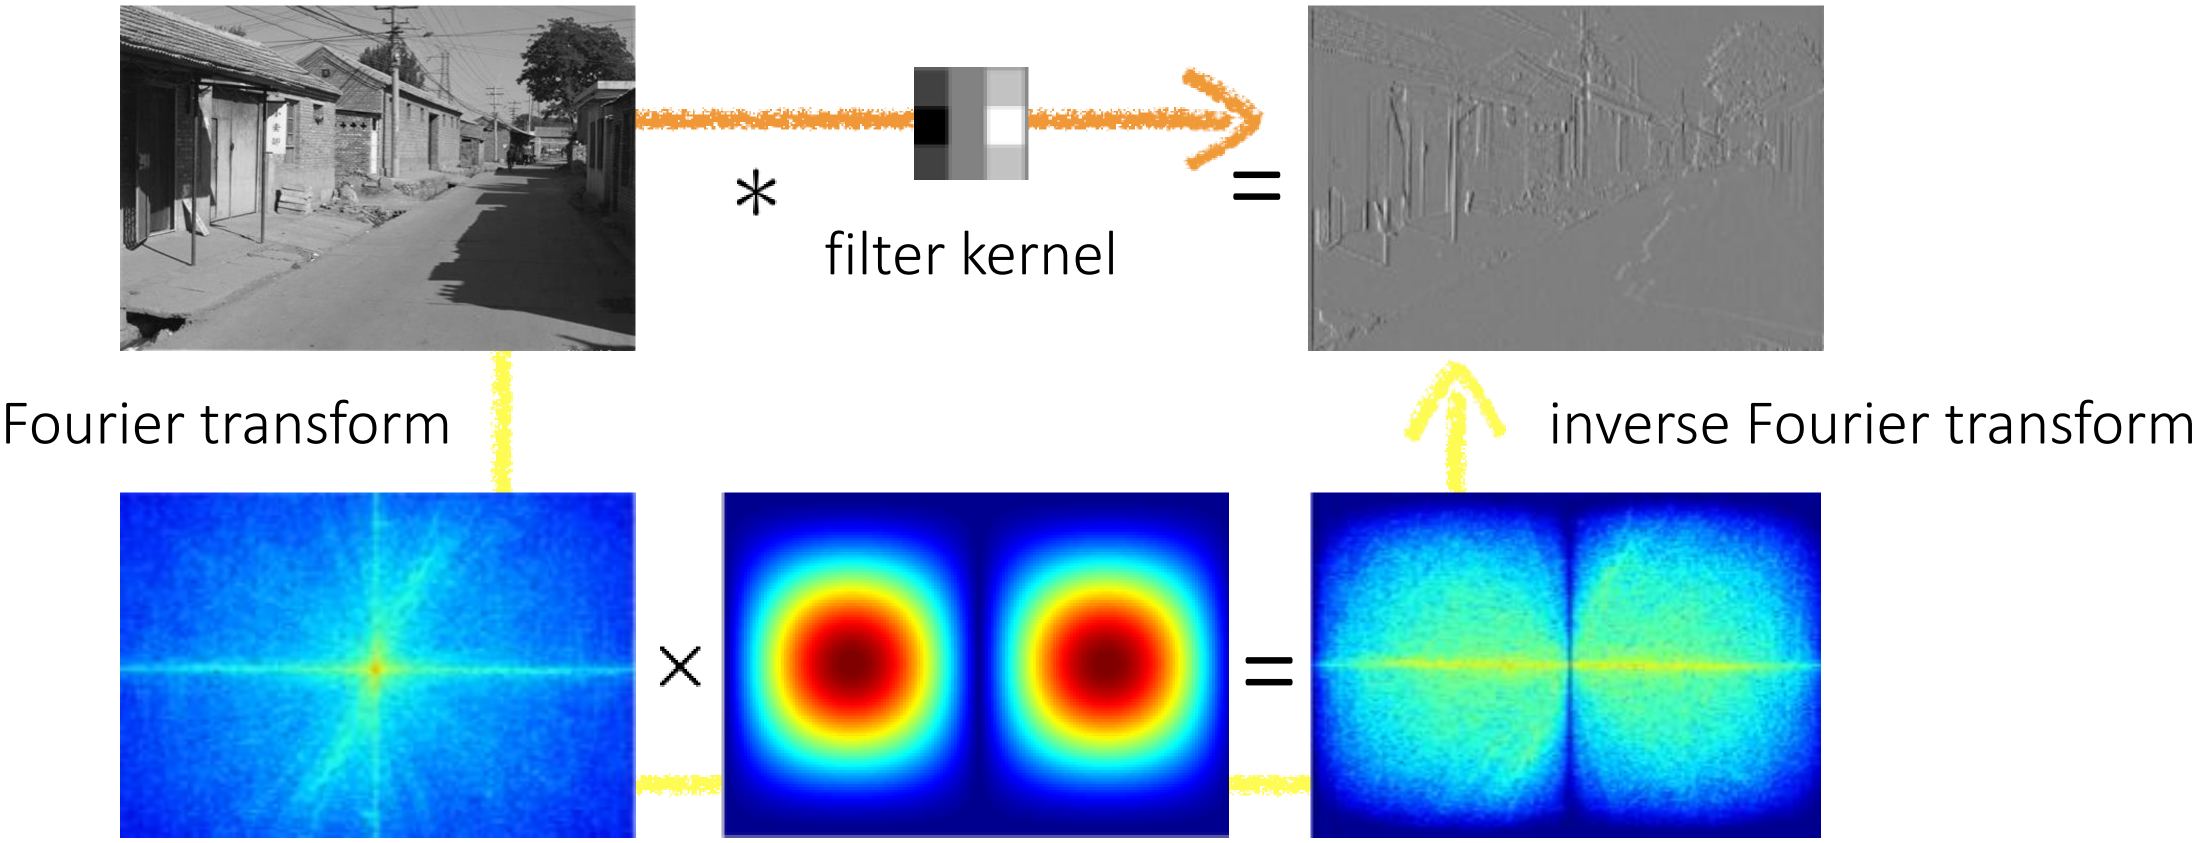


# Frequency domain filtering

# Revisiting blurring

Why does the Gaussian give a nice smooth image, but the square filter give edgy artifacts?

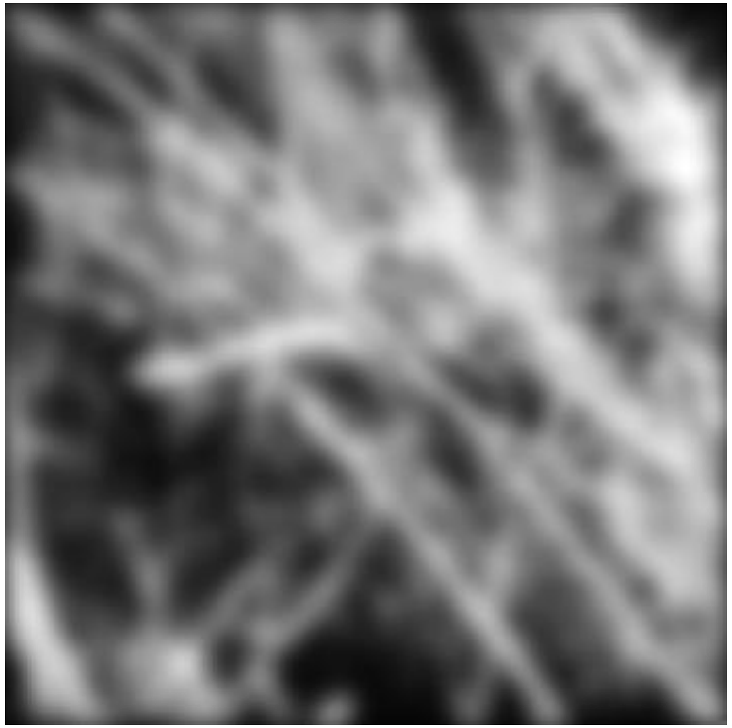

Gaussian filter

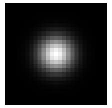

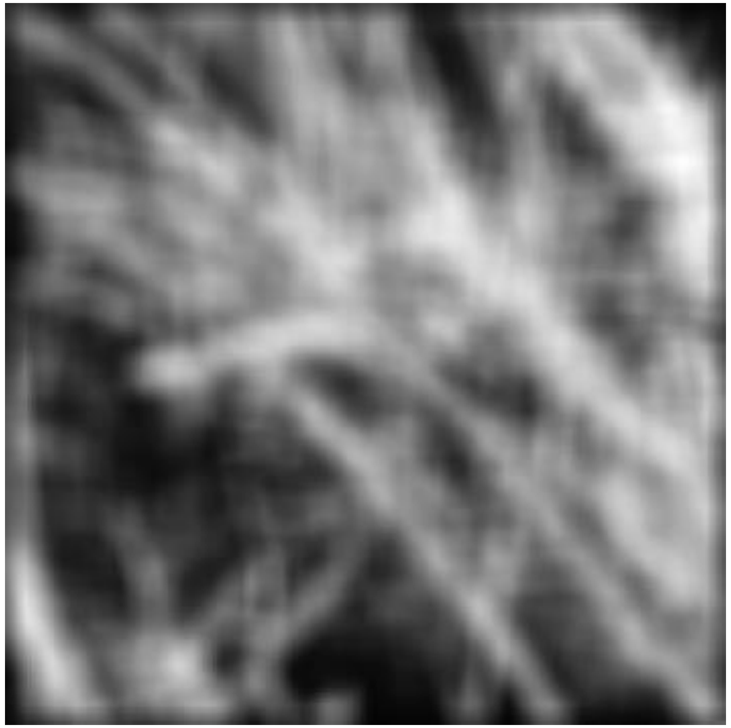

Box filter

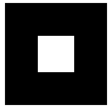

# Gaussian blur

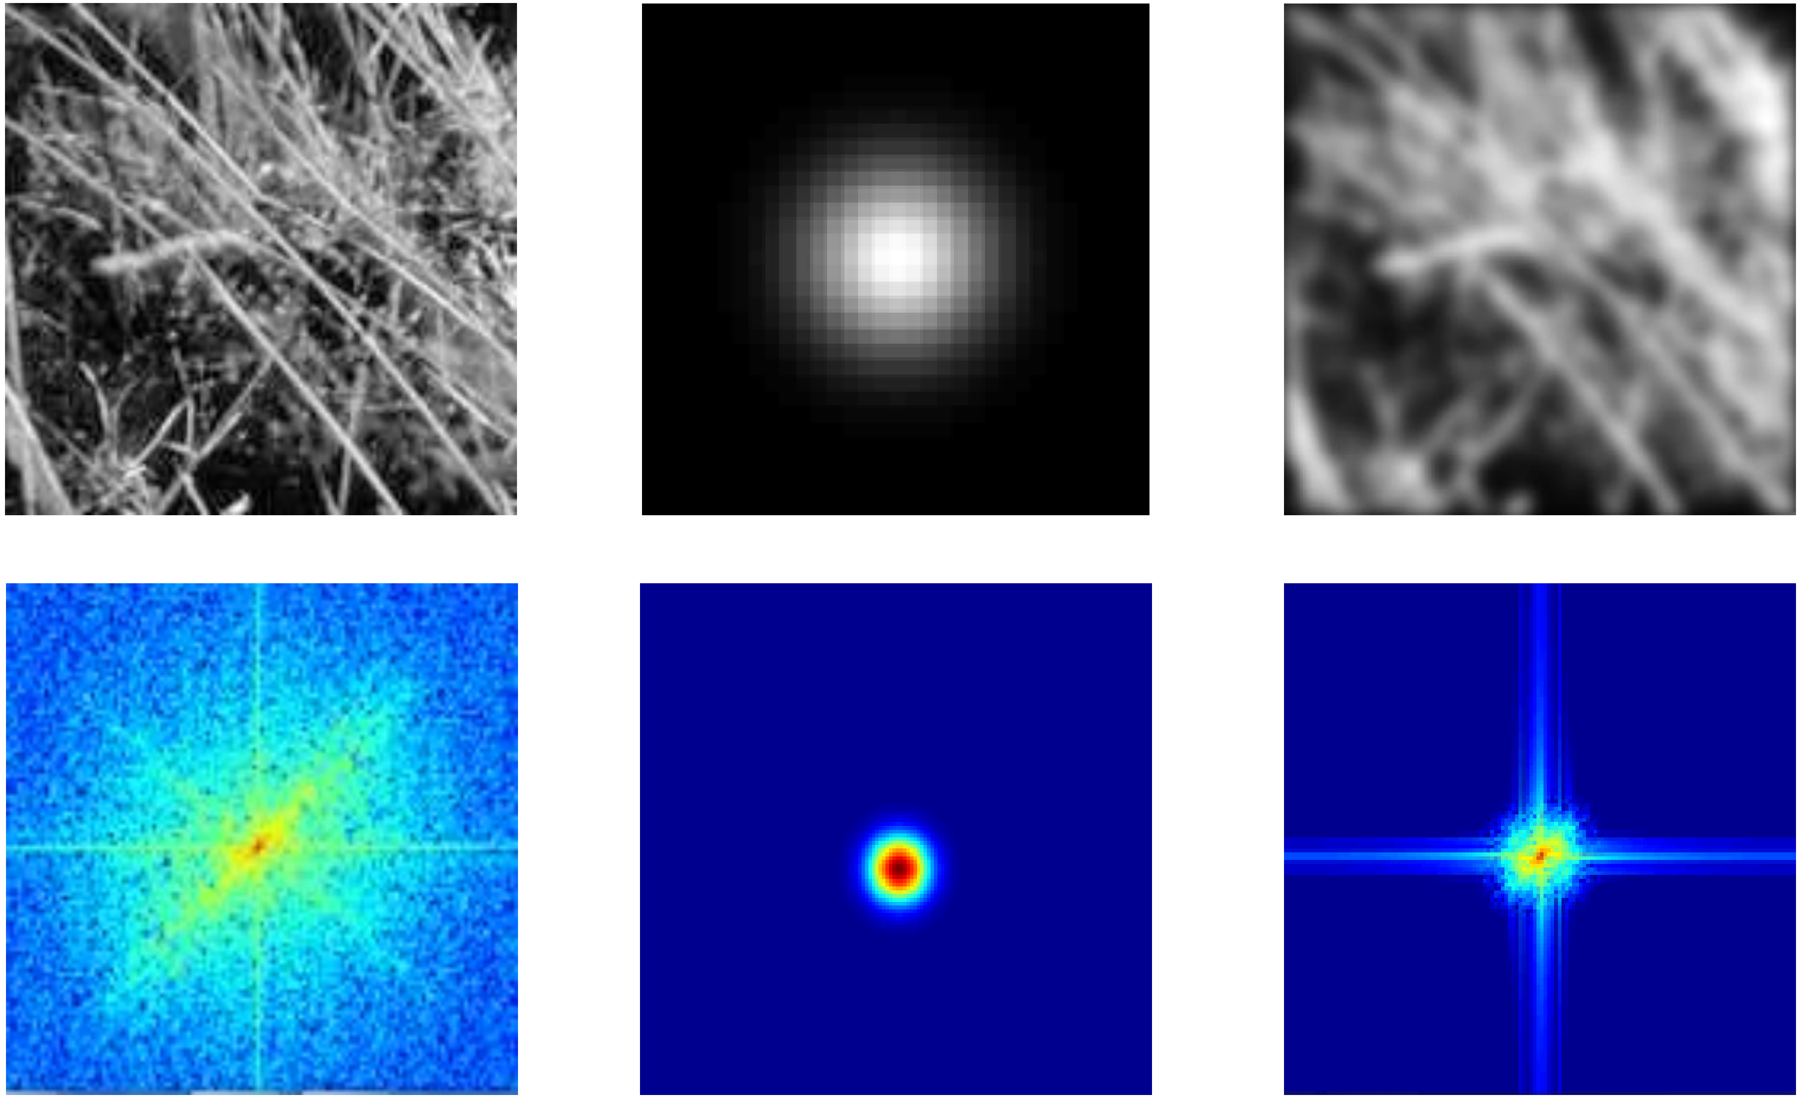

# Box blur

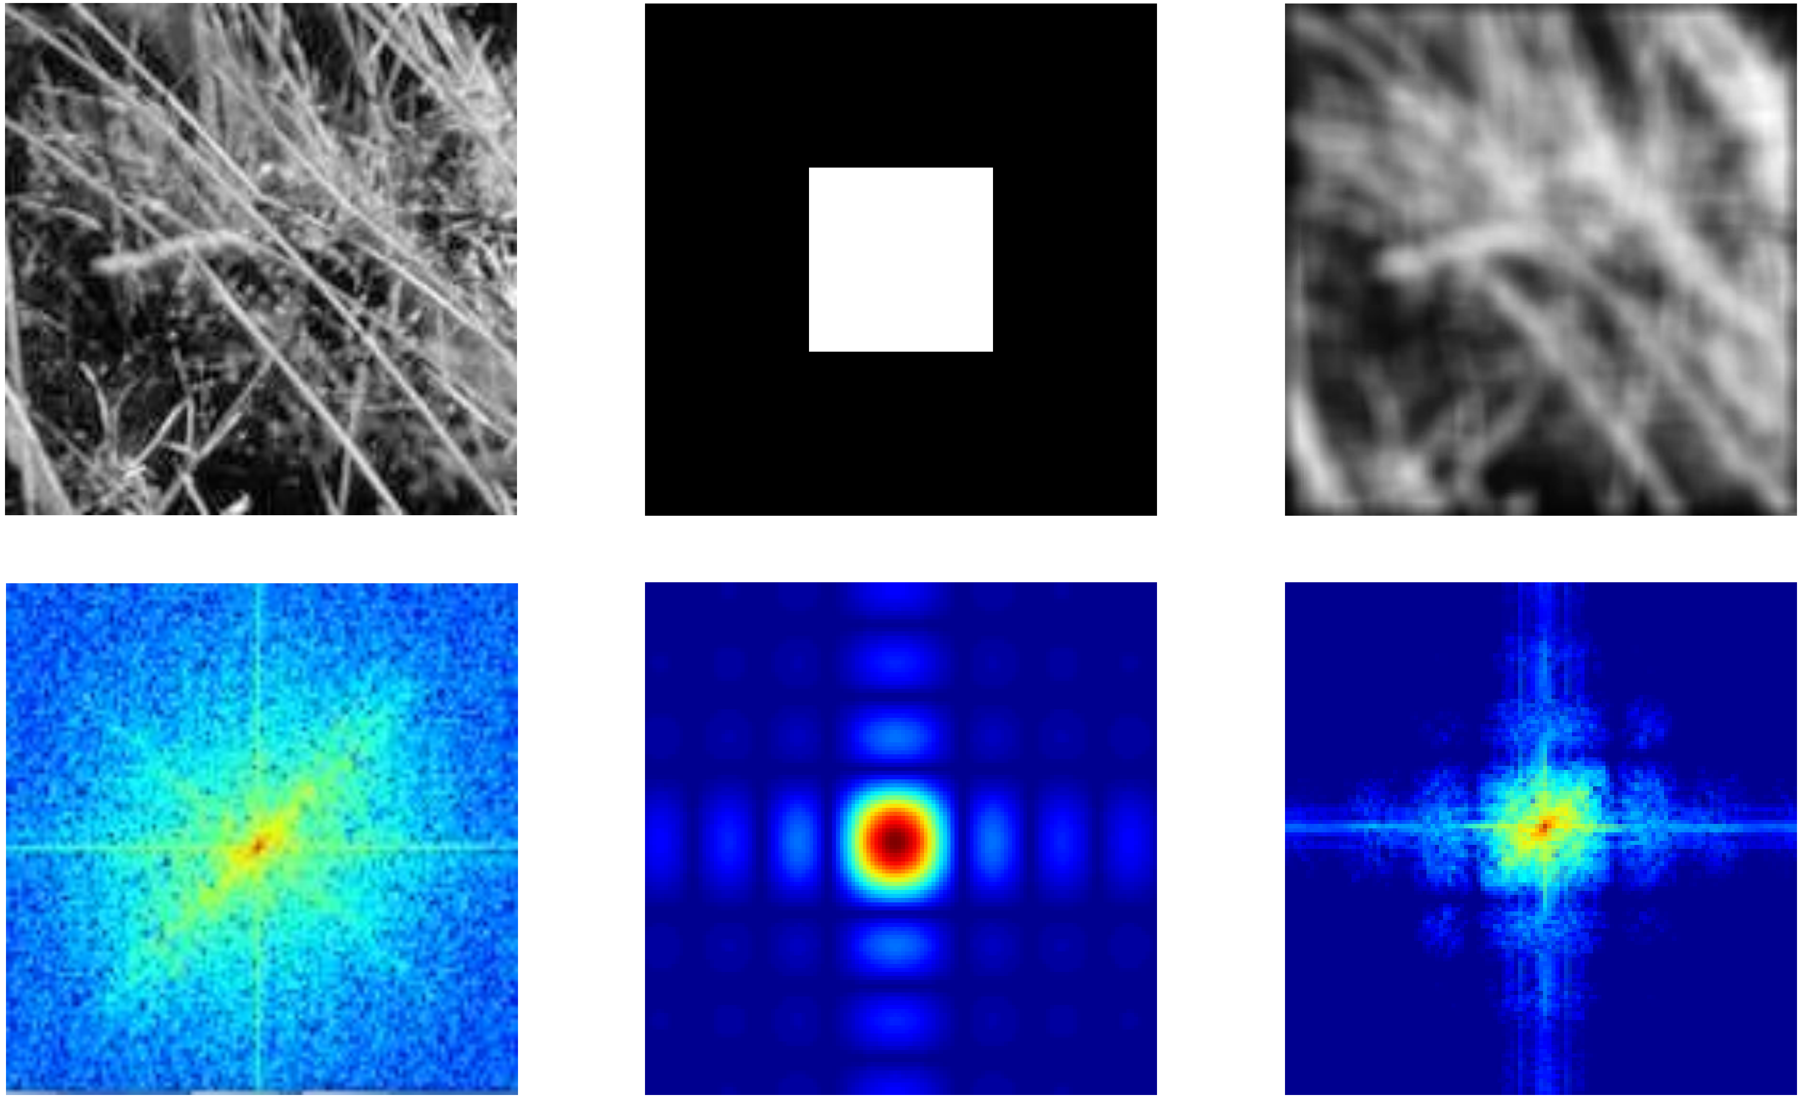

# A lens' kernel is its aperture

This is (one of the reasons) why we try to make lens apertures as circular as possible.

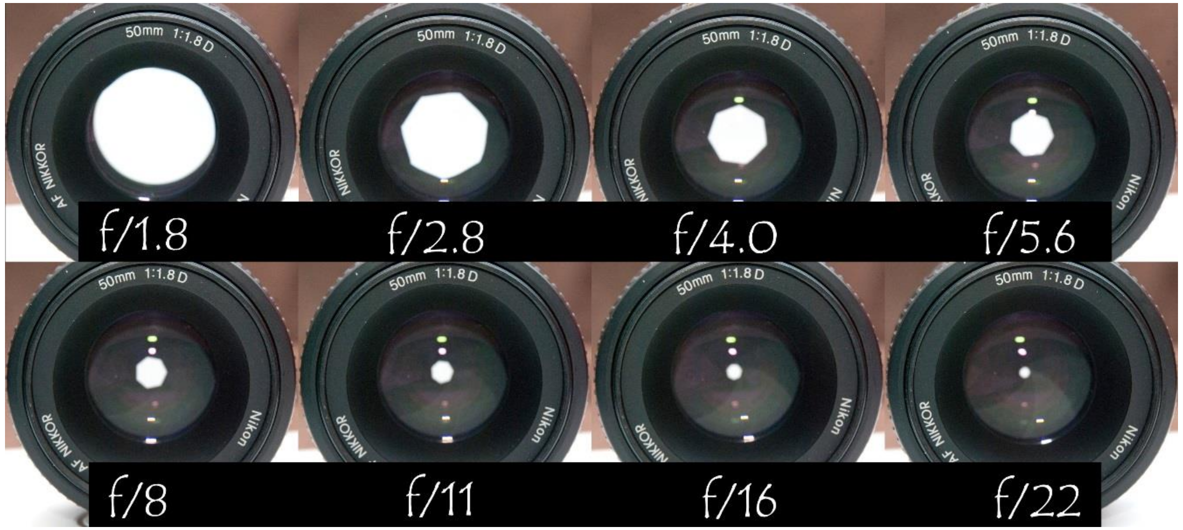


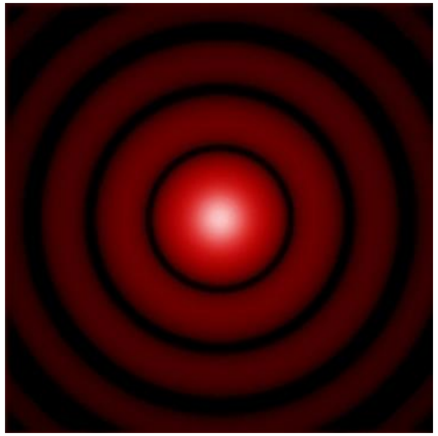

circular aperture (Airy disk)

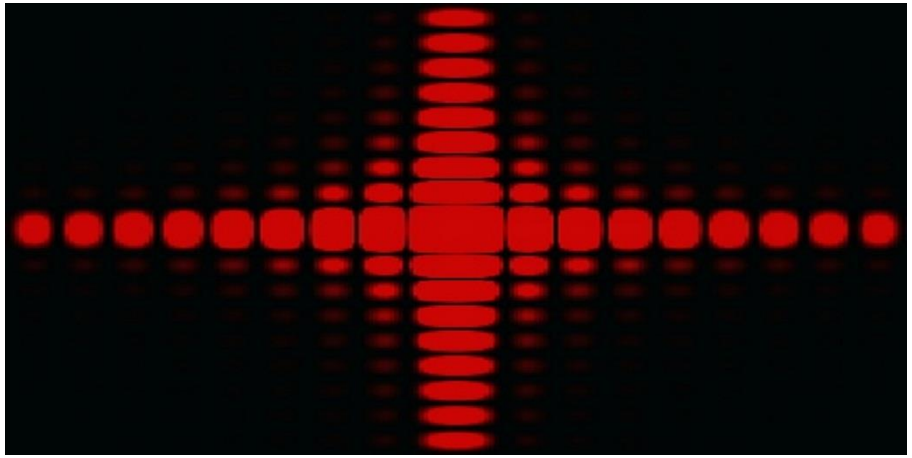

rectangular aperture

An ideal thin lens is an optical Fourier transform engine.

# More filtering examples

?

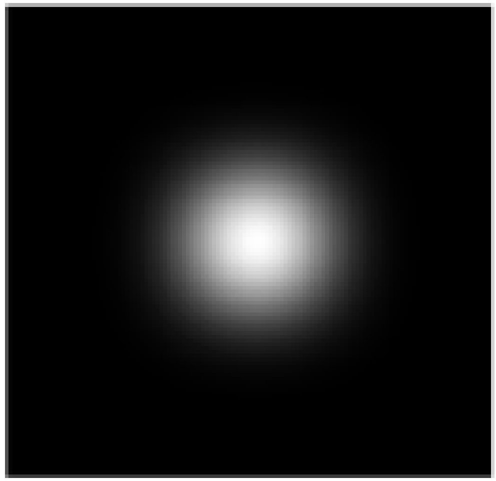

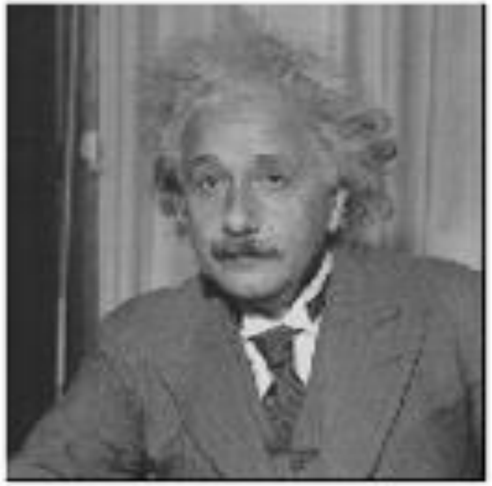

?

filters shown in frequency-domain

# More filtering examples

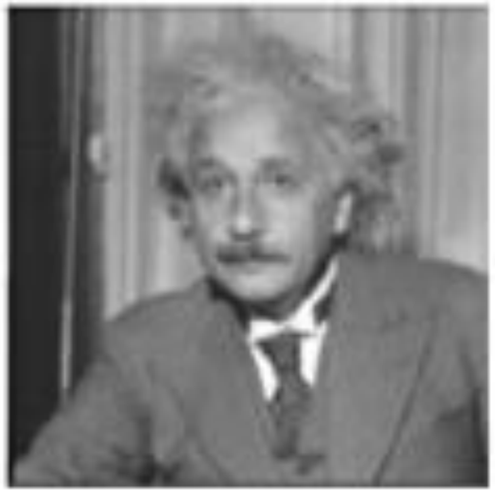

low-pass

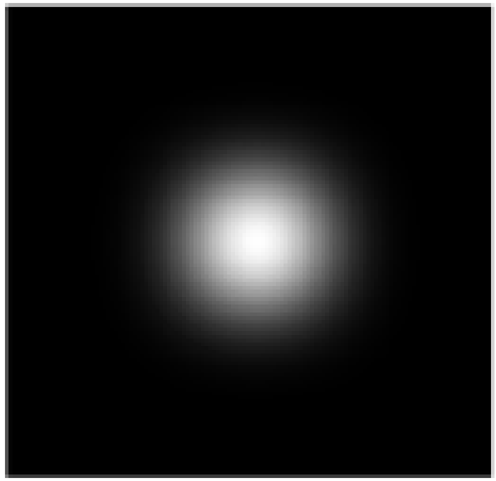

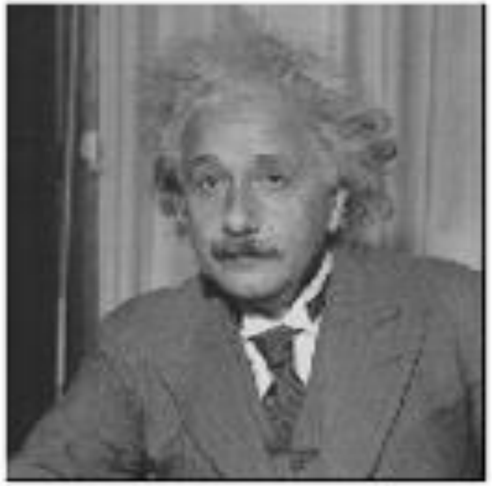

band-pass

filters shown in frequency-domain

# More filtering examples

?

high-pass

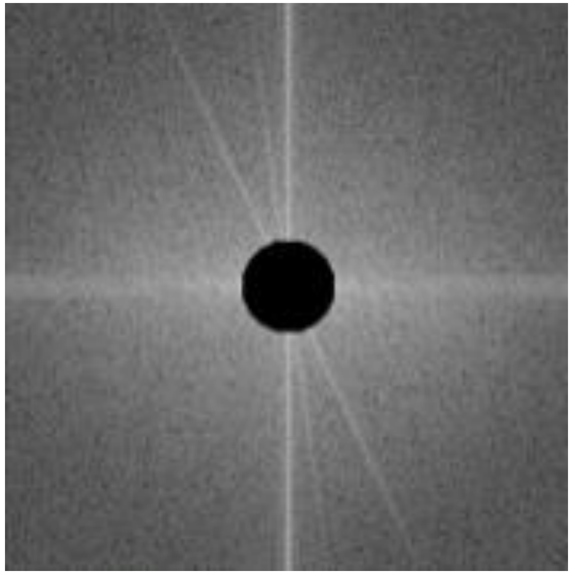

# More filtering examples

high-pass

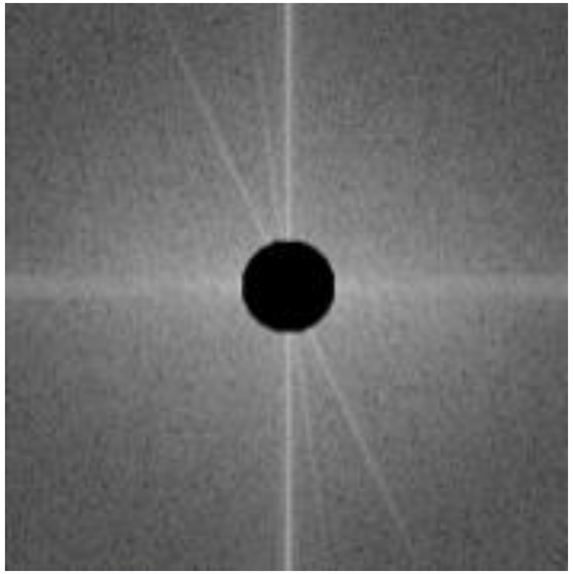

# More filtering examples

original image

low-pass filter

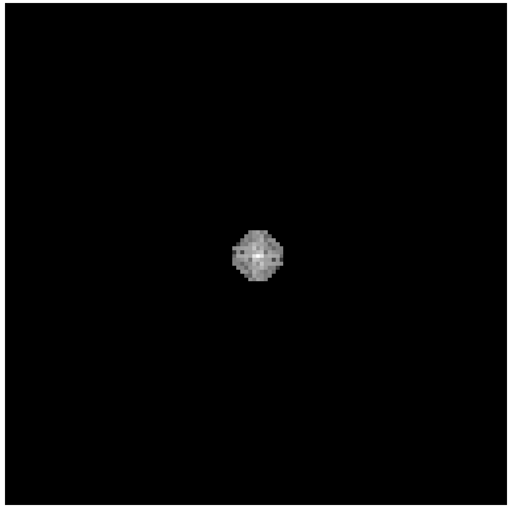

?

frequency magnitude

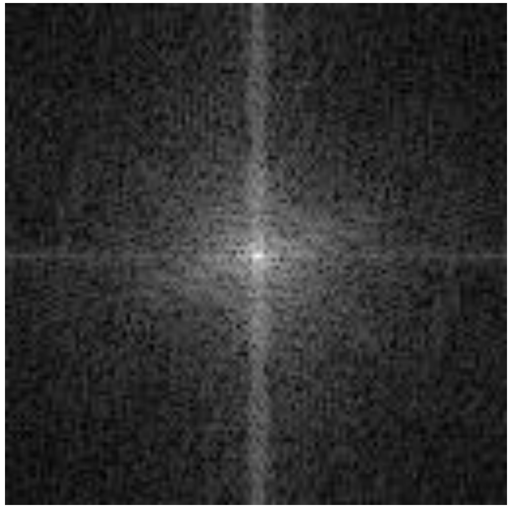

# More filtering examples

original image

low-pass filter

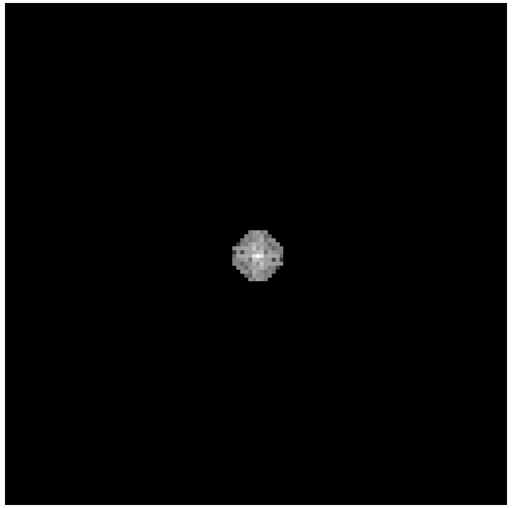

frequency magnitude

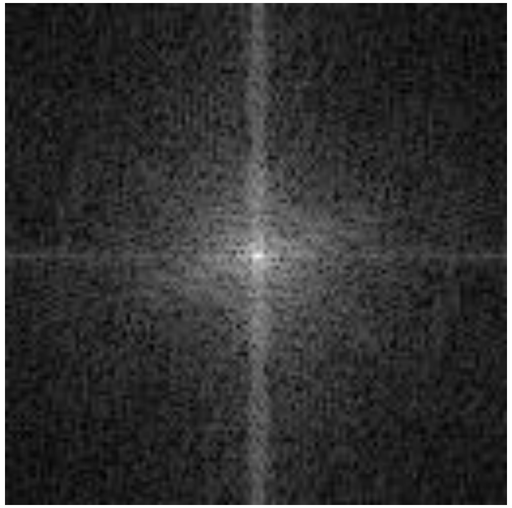

# More filtering examples

original image

high-pass filter

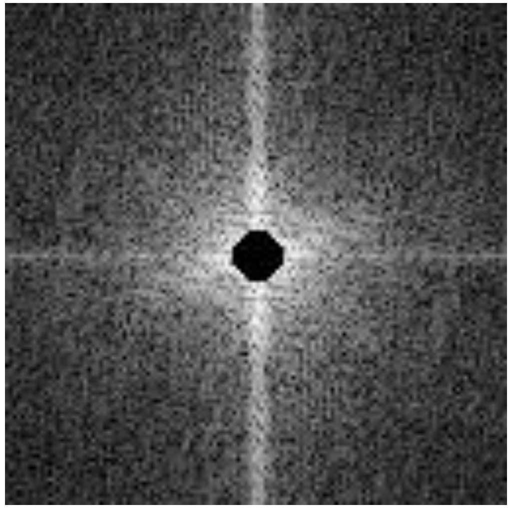

?

frequency magnitude

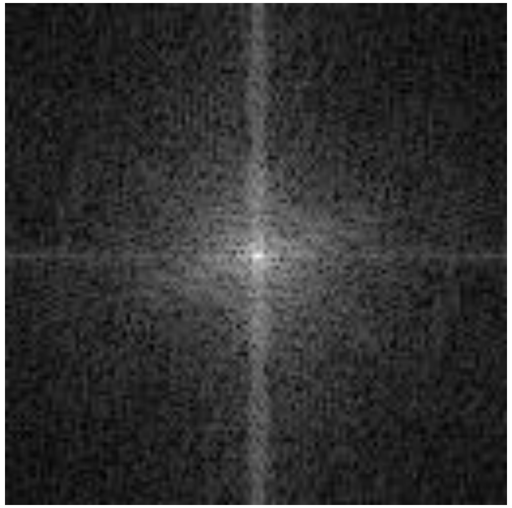

# More filtering examples

original image

high-pass filter

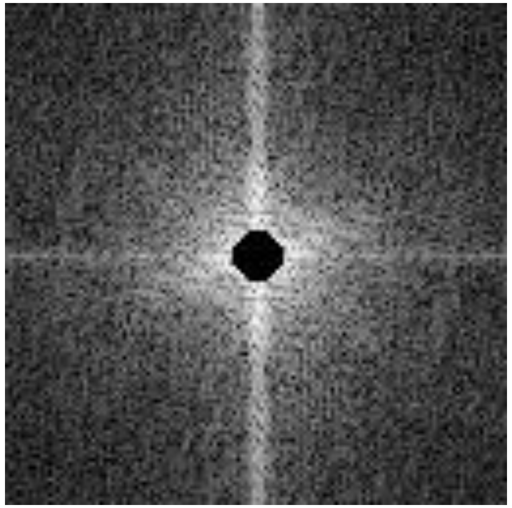

frequency magnitude

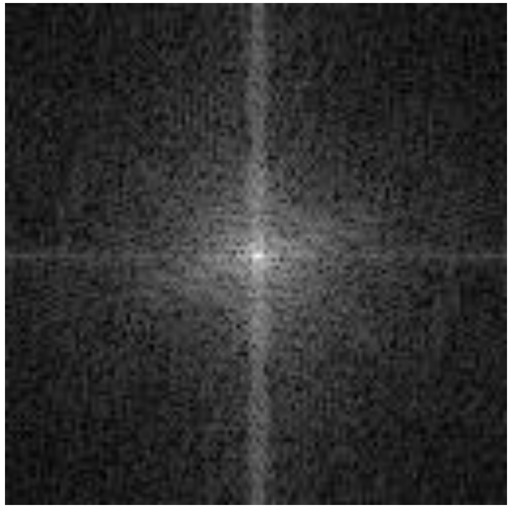

# More filtering examples

original image

band-pass filter

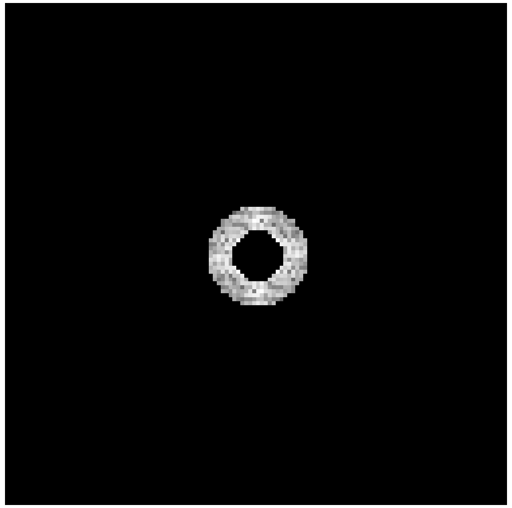

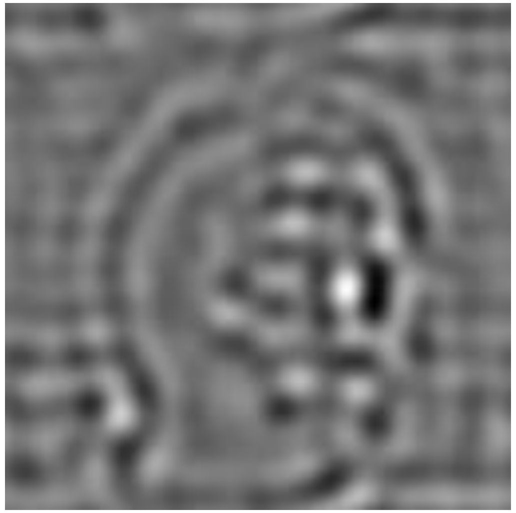

frequency magnitude

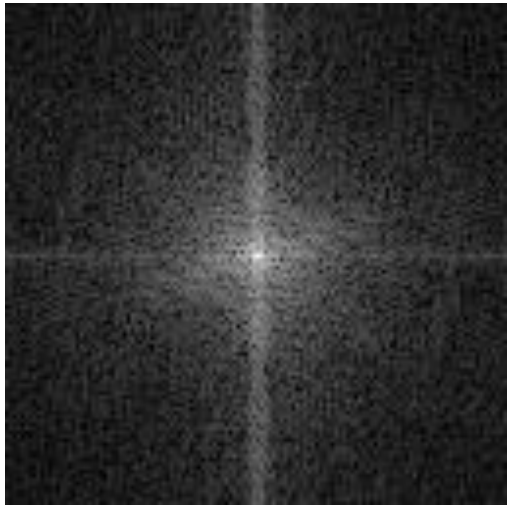

# More filtering examples

original image

band-pass filter

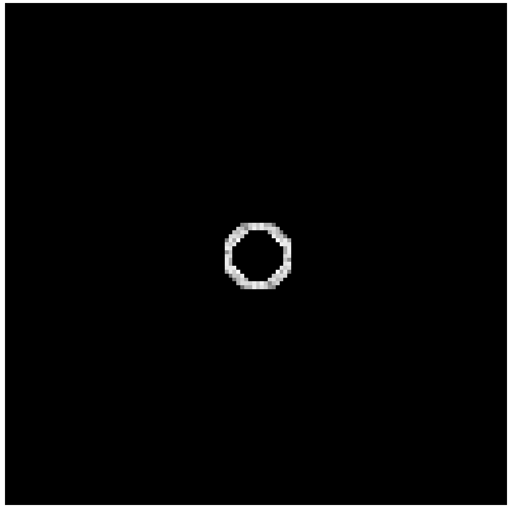

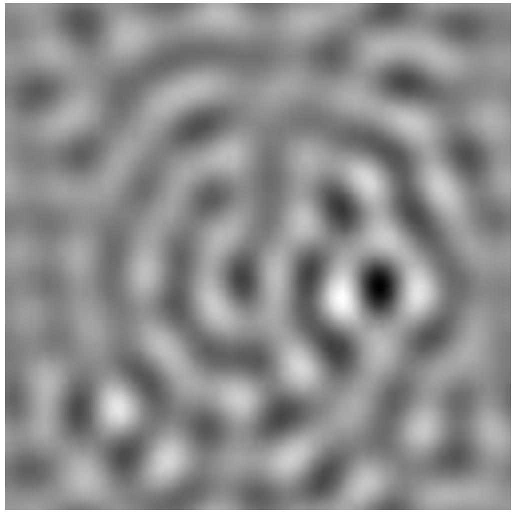

frequency magnitude

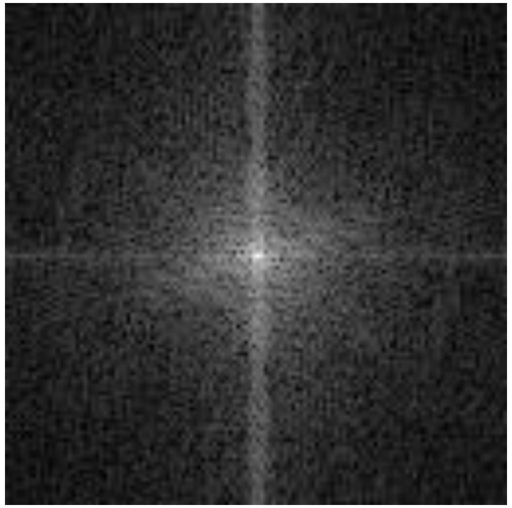

# More filtering examples

original image

band-pass filter

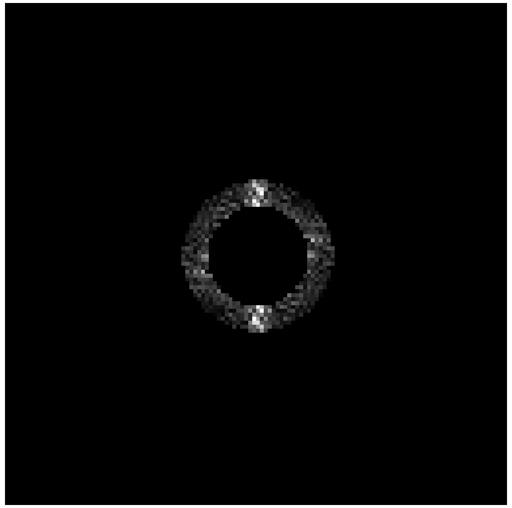

frequency magnitude

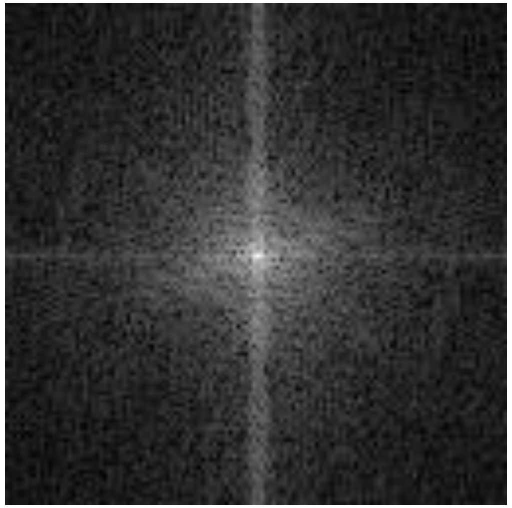

# More filtering examples

original image

band-pass filter

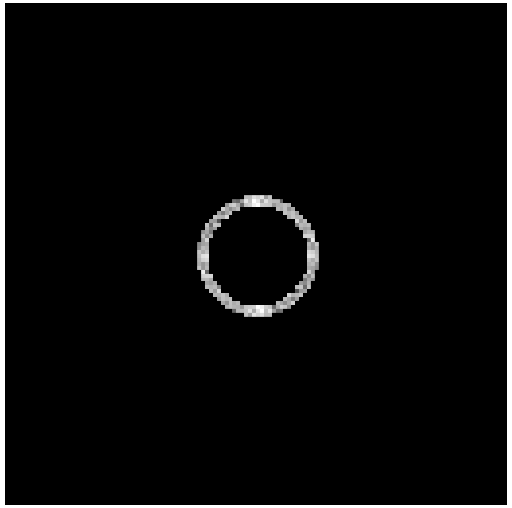

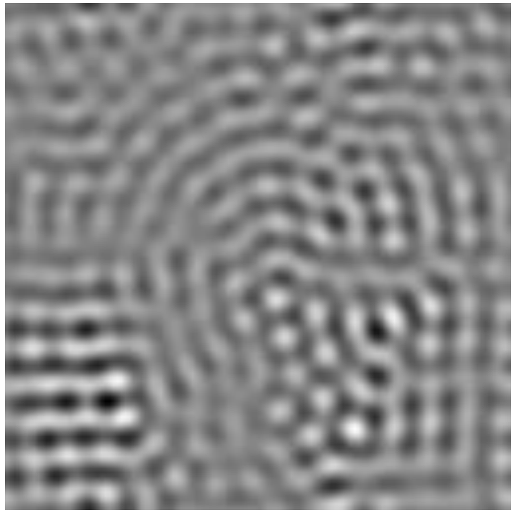

frequency magnitude

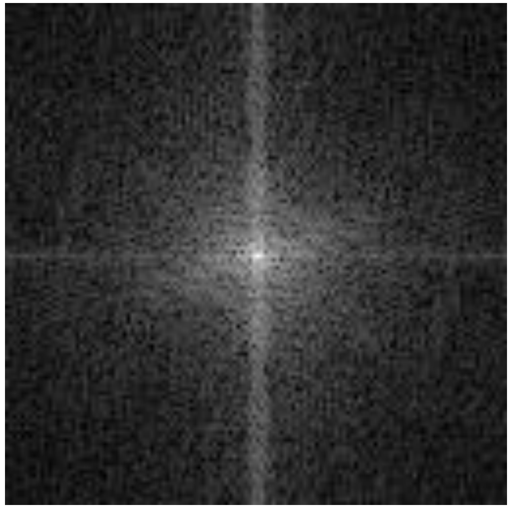

# Revisiting sampling

# The Nyquist-Shannon sampling theorem

A continuous signal can be perfectly reconstructed from its discrete version if sampling occurred with frequency:

$$f_s \geq 2 f_{\max}$$

This is called the Nyquist frequency

Equivalent reformulation: When downsampling, aliasing does not occur if samples are taken at the Nyquist frequency or higher.

# The Nyquist-Shannon sampling theorem

A continuous signal can be perfectly reconstructed from its discrete version if sampling occurred with frequency:

$$f_s \geq 2 f_{\max} \longleftarrow$$

This is called the Nyquist frequency

# Gaussian pyramid

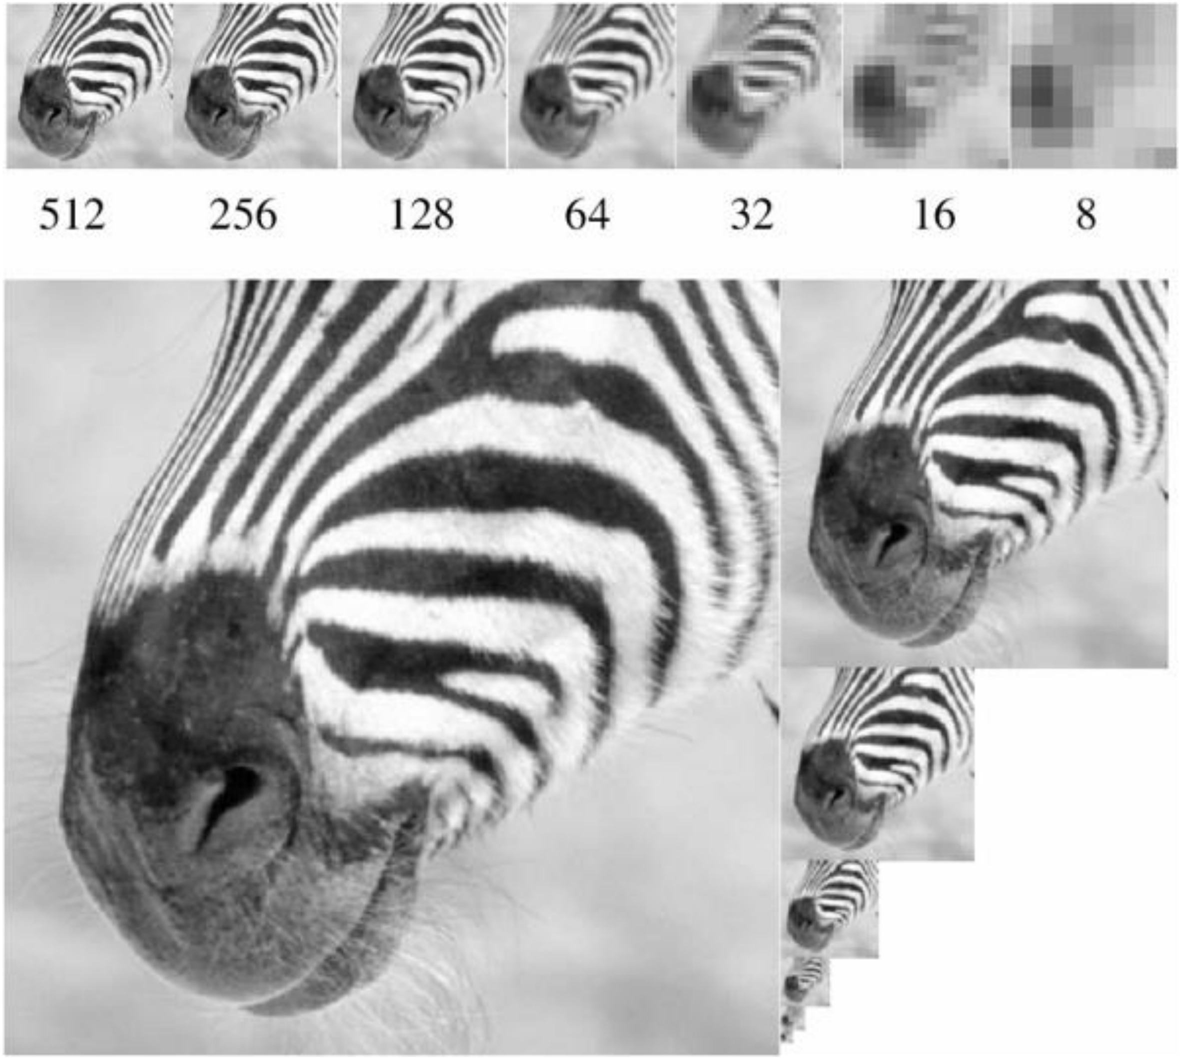


How does the Nyquist-Shannon theorem relate to the Gaussian pyramid?

# Gaussian pyramid

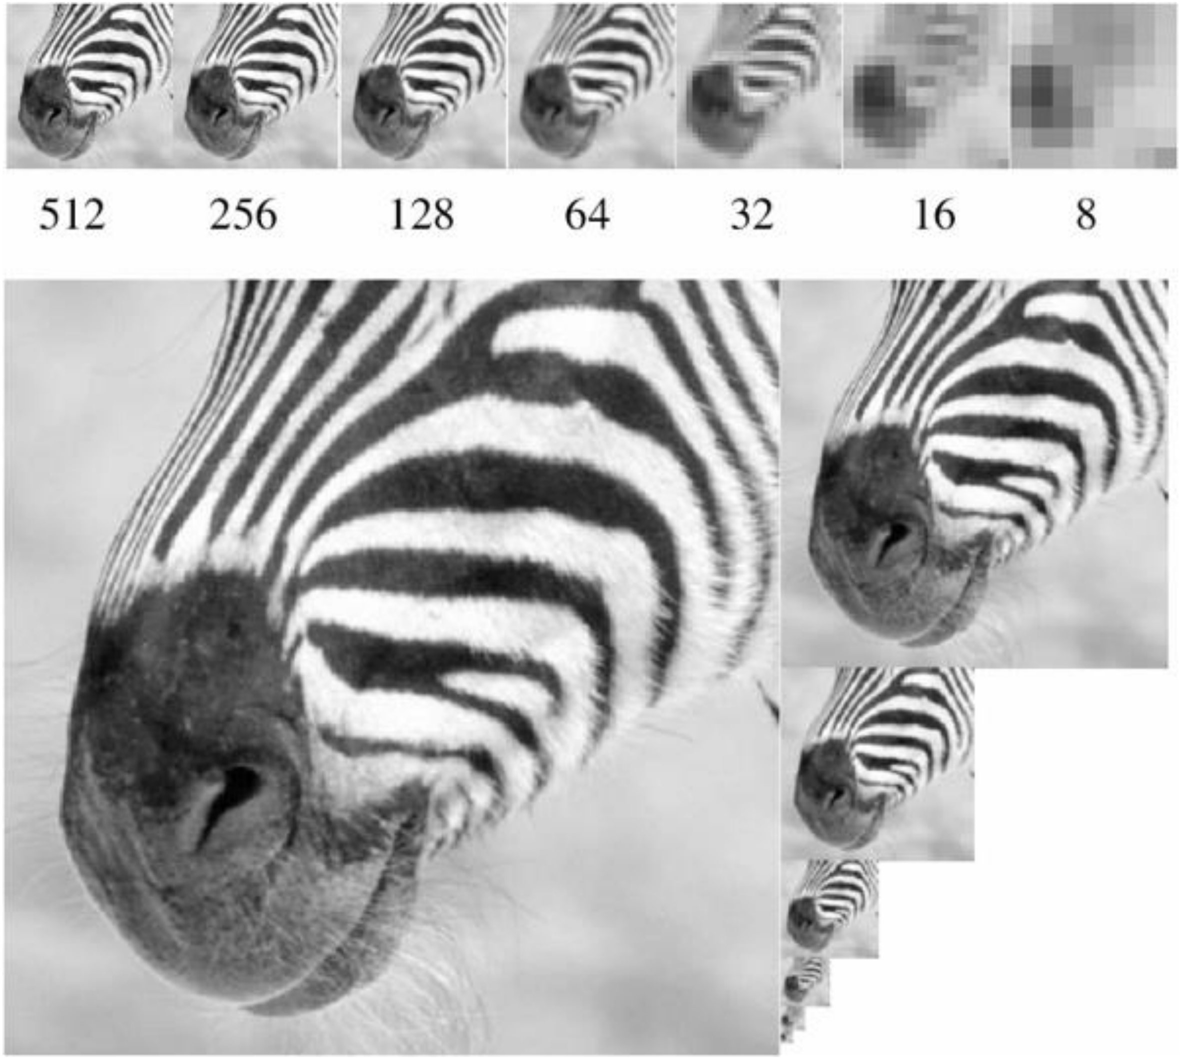


How does the Nyquist-Shannon theorem relate to the Gaussian pyramid?

- Gaussian blurring is low-pass filtering.
- By blurring we try to sufficiently decrease the Nyquist frequency to avoid aliasing.

How large should the Gauss blur we use be?

# Frequency-domain filtering in human vision

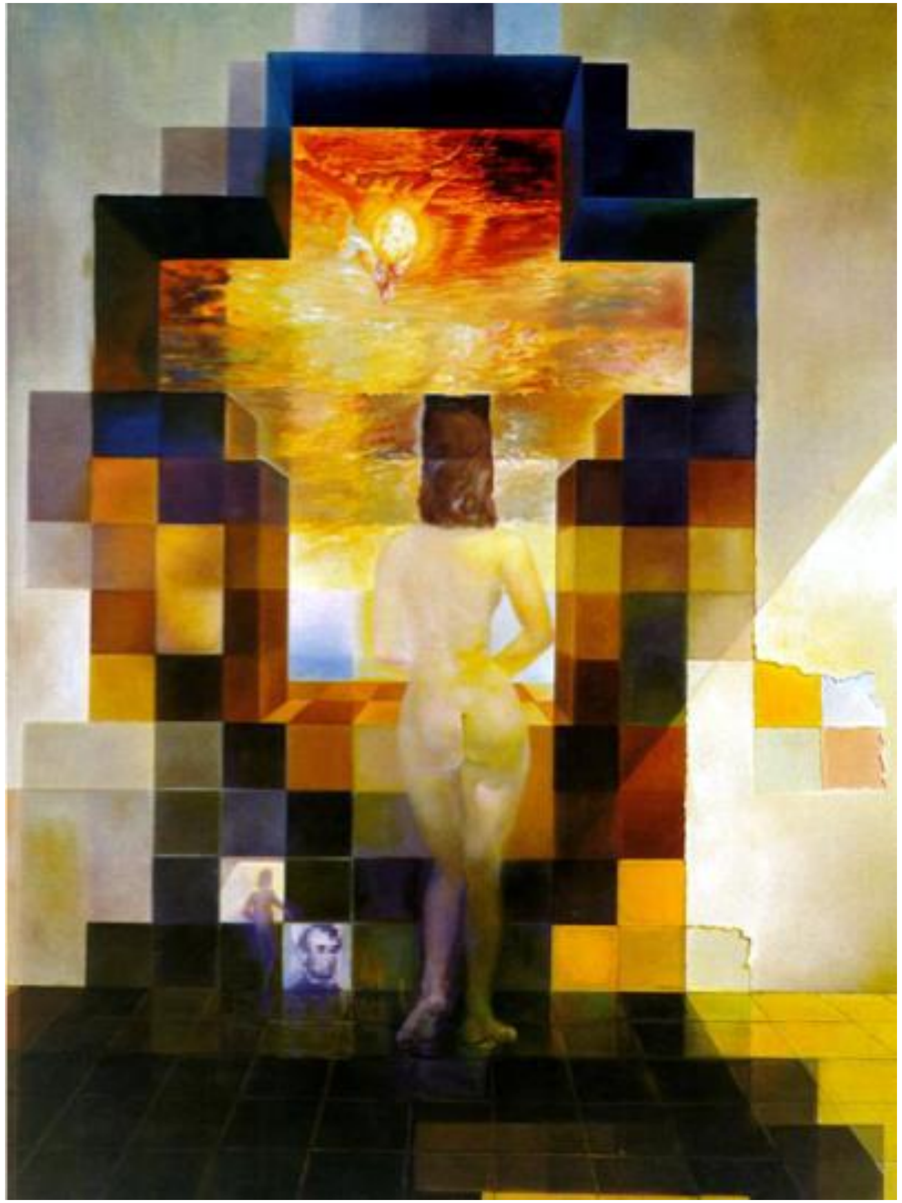

*Gala Contemplating the Mediterranean Sea Which at Twenty Meters Becomes the Portrait of Abraham Lincoln (Homage to Rothko)*

Salvador Dali, 1976

# Frequency-domain filtering in human vision

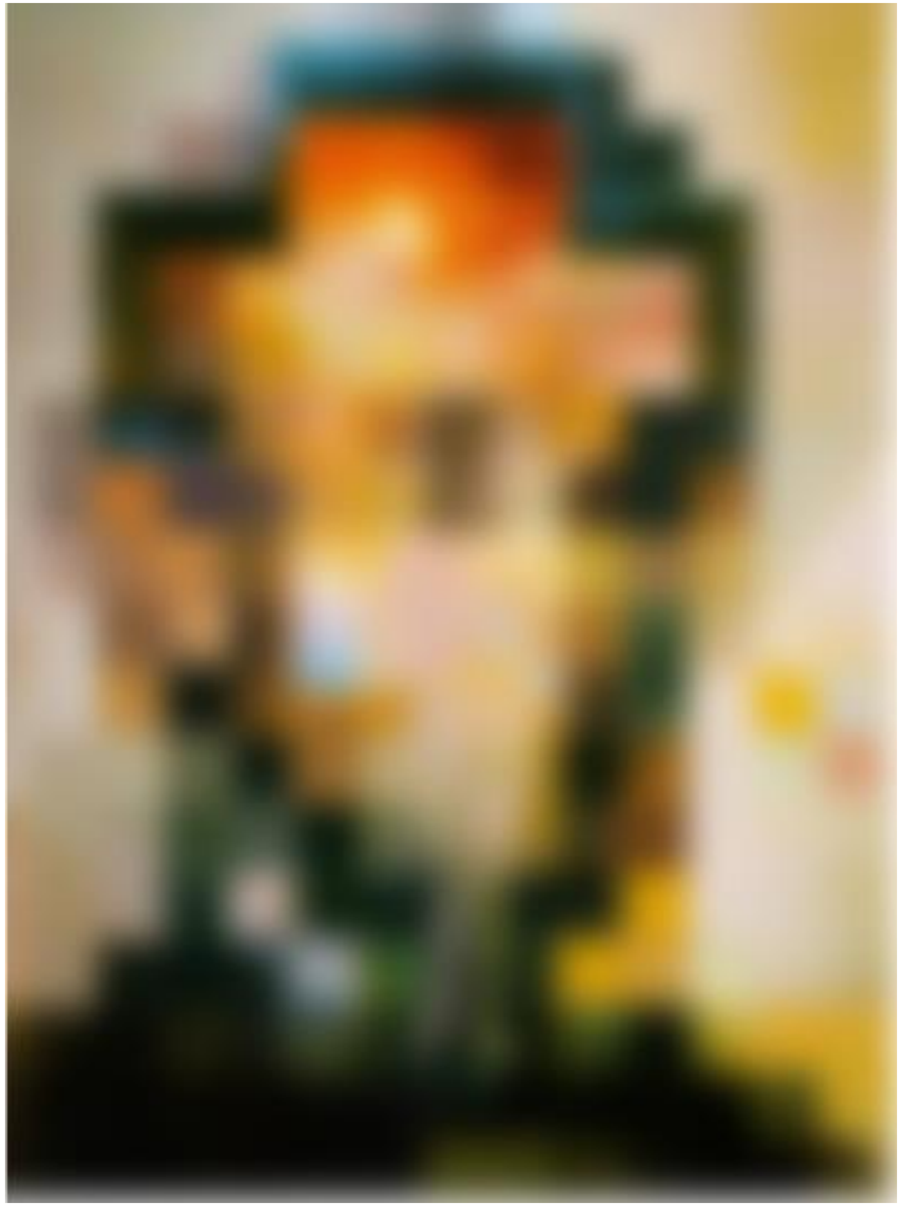

Low-pass filtered version

# Frequency-domain filtering in human vision

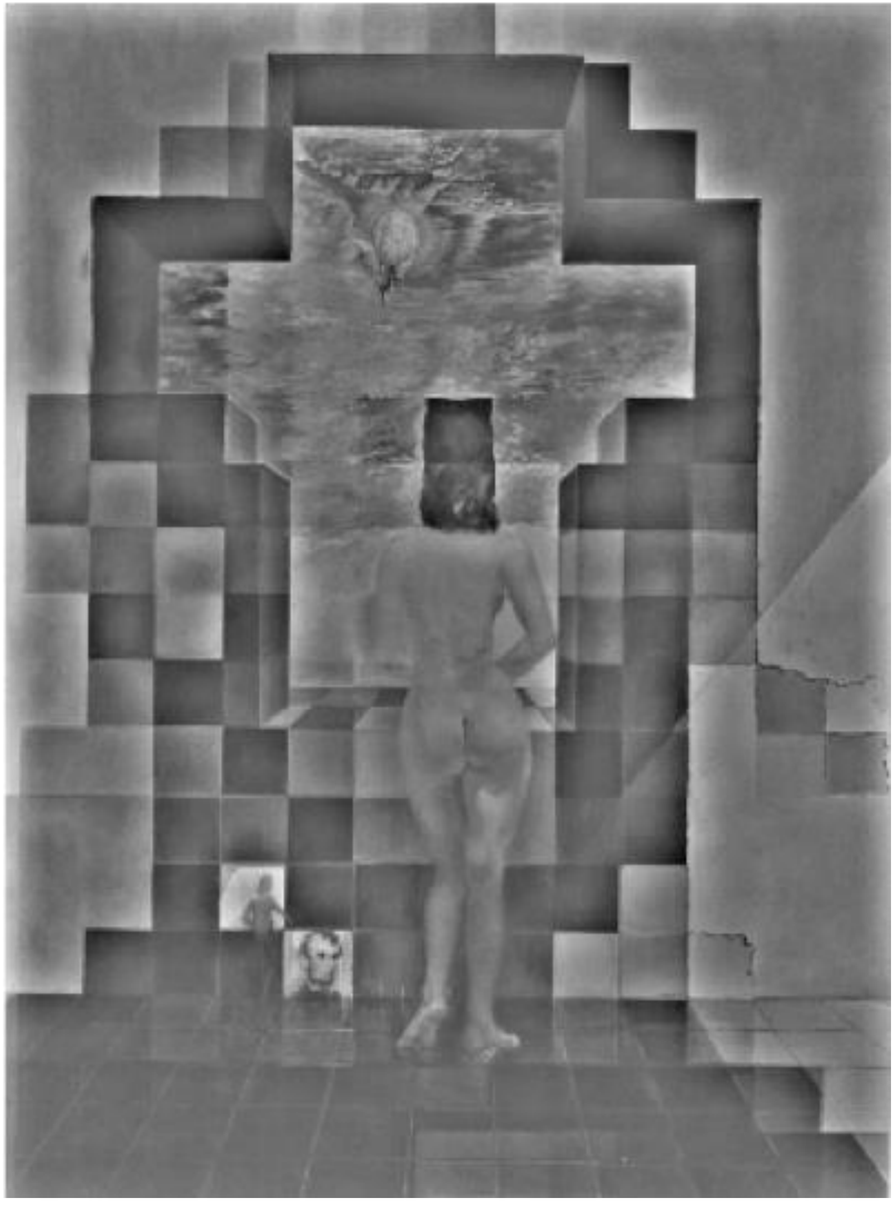

High-pass filtered version

# Variable frequency sensitivity

Experiment: Where do you see the stripes?

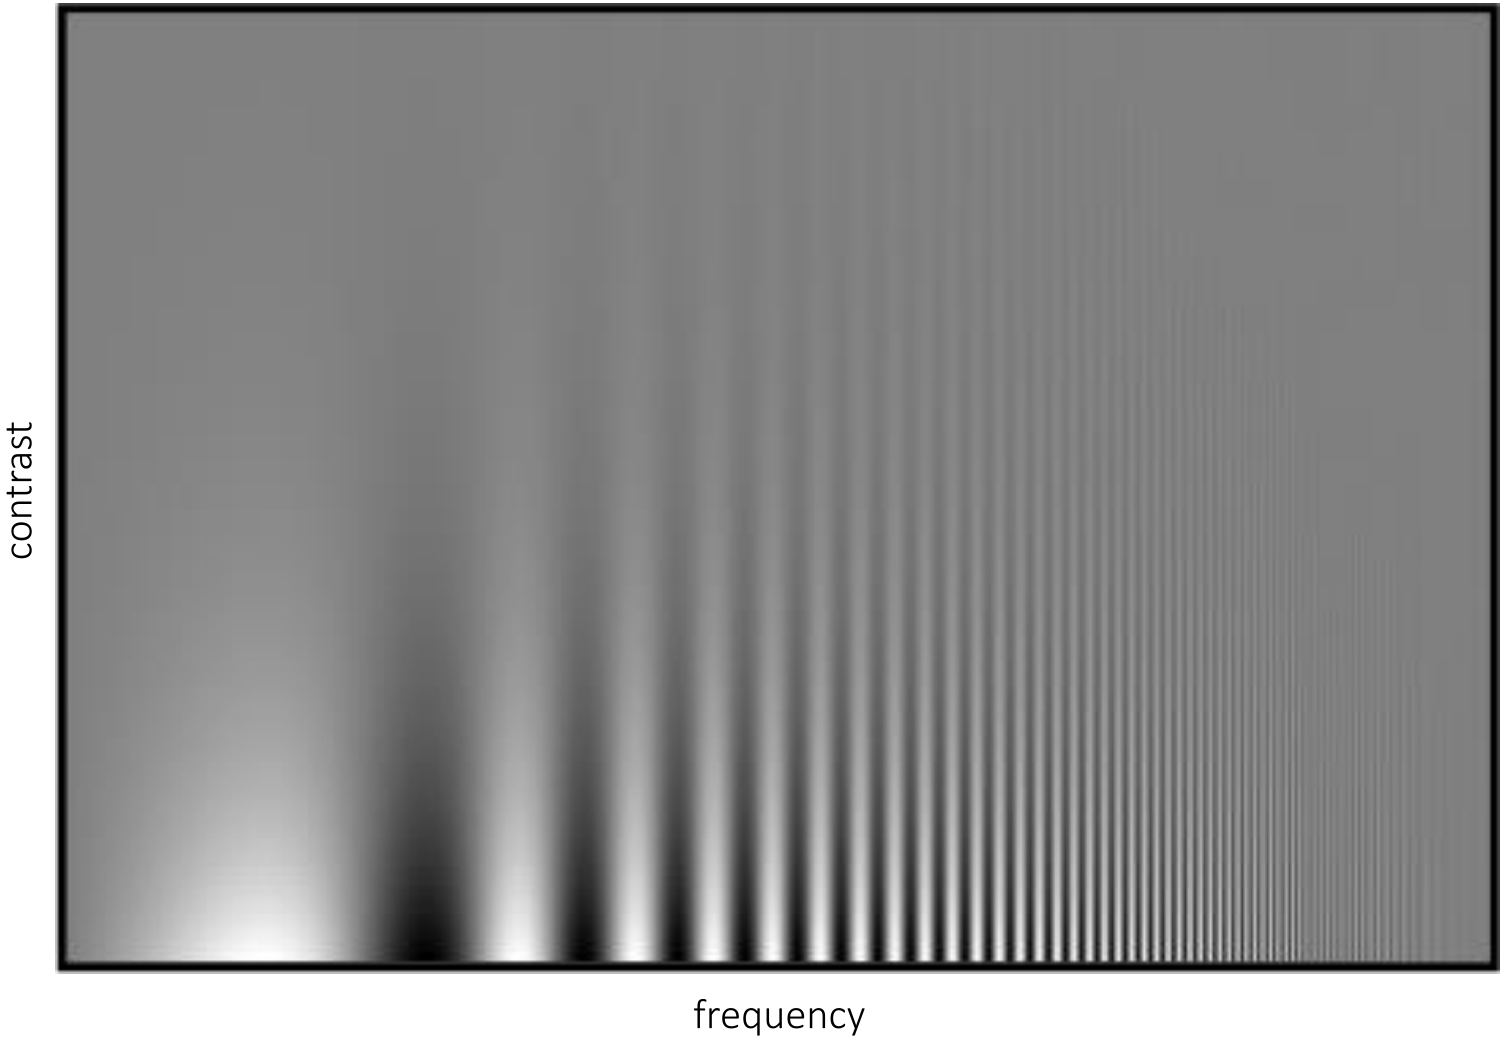

# Variable frequency sensitivity

Campbell-Robson contrast sensitivity curve

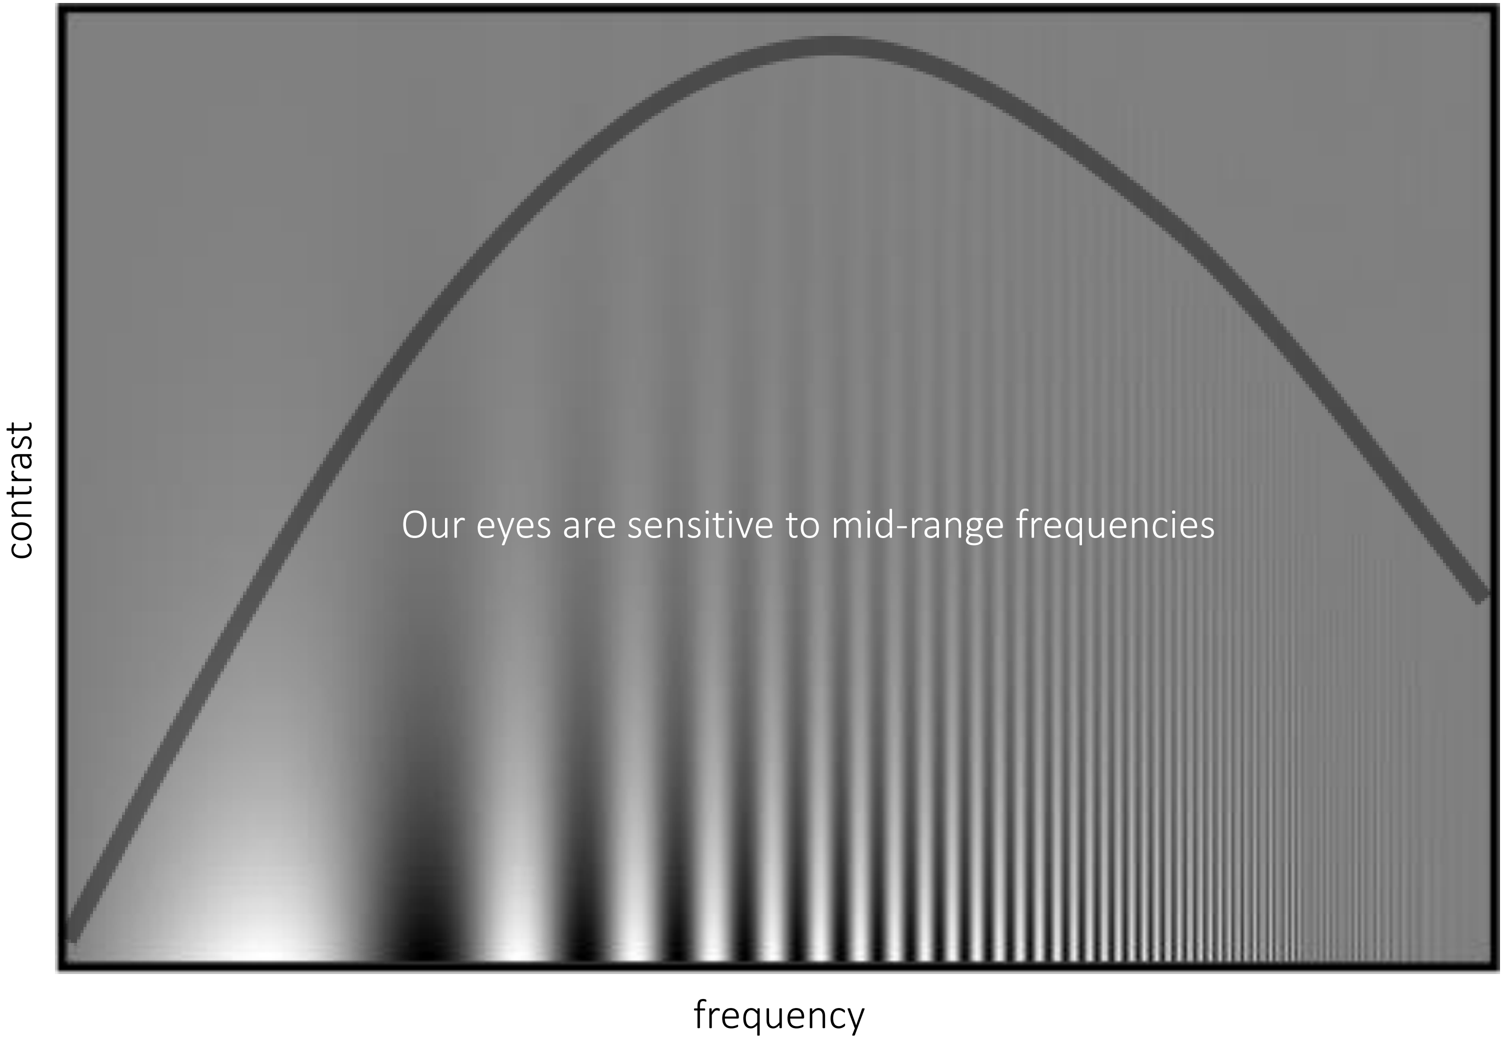


- Early processing in humans filters for various orientations and scales of frequency
- Perceptual cues in the mid frequencies dominate perception

# References

Basic reading:

- Szeliski textbook, Sections 3.4.

Additional reading:

- Goodman, “Introduction to Fourier Optics,” W.H.Freeman Publishing 2004.
    the standard reference on Fourier optics
- Hubel and Wiesel, “Receptive fields, binocular interaction and functional architecture in the cat's visual cortex,” The Journal of Physiology 1962
    a foundational paper describing information processing in the visual system, including the different types of filtering it performs; Hubel and Wiesel won the Nobel Prize in Medicine in 1981 for the discoveries described in this paper