Fourier Analysis

15-463: Rendering and Image Processing

Alexei Efros

Image Scaling

This image is too big to fit on the screen. How can we reduce it?

How to generate a half-sized version?

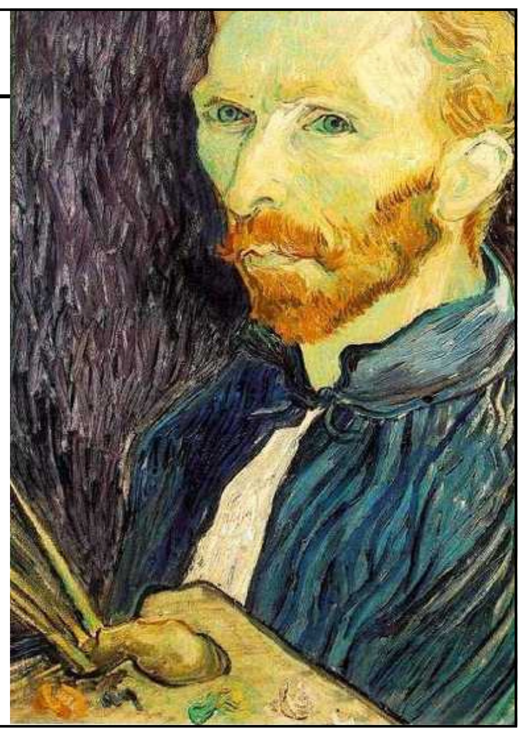

Image sub-sampling

1/4

1/8

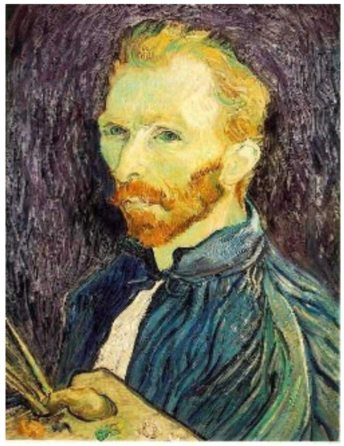

Throw away every other row and column to create a 1/2 size image
- called *image sub-sampling*

Slide by Steve Seitz

Image sub-sampling

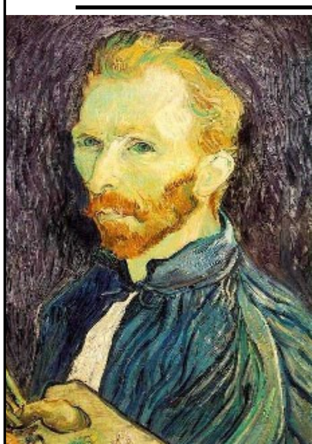

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Why does this look so crufty?

Slide by Steve Seitz

Even worse for synthetic images

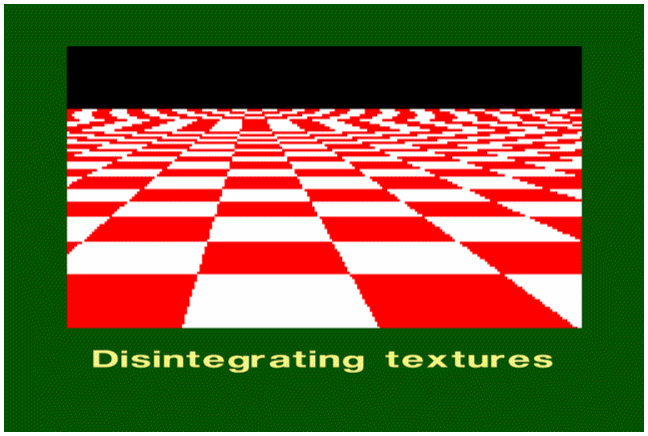

Slide by Steve Seitz

Really bad in video

Imagine a spoked wheel moving to the right (rotating clockwise). Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):

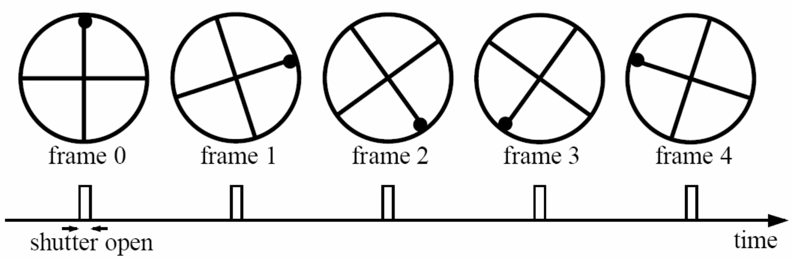

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Paul Heckbert

Alias: n., an assumed name

Input signal:

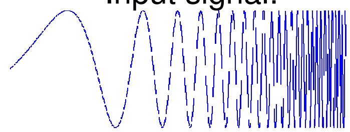

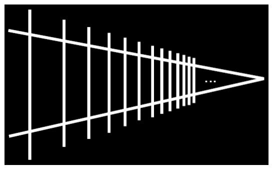

Picket fence receding
Into the distance will
produce aliasing…

Matlab output:

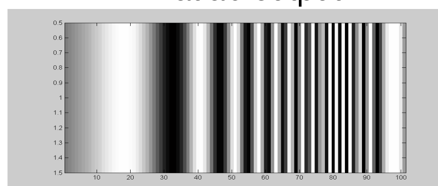

x = 0:.05:5; imagesc(sin((2.^x).*x))

WHY?

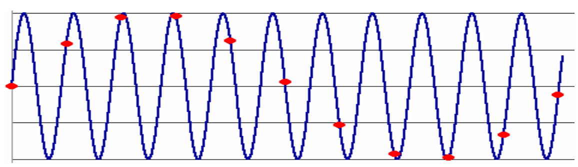

Aj-aj-aj:
Alias!

Not enough samples

Aliasing

- occurs when your sampling rate is not high enough to capture the amount of detail in your image
- Can give you the wrong signal/image—an *alias*

Where can it happen in graphics?

- During image synthesis:
  - **sampling** continous singal into discrete signal
  - e.g. ray tracing, line drawing, function plotting, etc.
- During image processing:
  - **resampling** discrete signal at a different rate
  - e.g. Image warping, zooming in, zooming out, etc.

To do sampling right, need to understand the structure of your signal/image

Enter Monsieur Fourier…

Jean Baptiste Fourier (1768-1830)

had crazy idea (1807):

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

Don't believe it?

- Neither did Lagrange, Laplace, Poisson and other big wigs
- Not translated into English until 1878!

But it's true!

- called Fourier Series

A sum of sines

Our building block:

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

Add enough of them to get any signal *f(x)* you want!

How many degrees of freedom?

What does each control?

Which one encodes the coarse vs. fine structure of the signal?

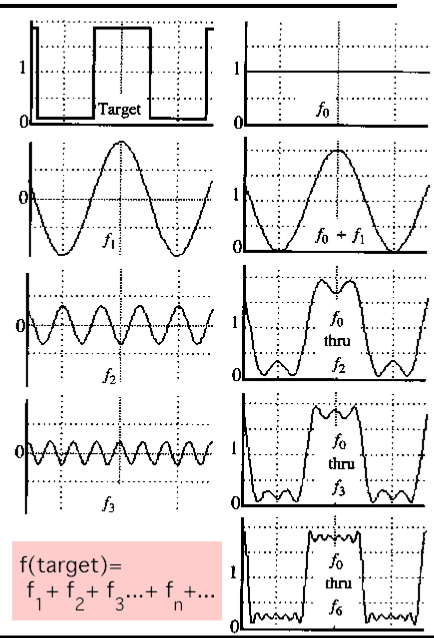

Fourier Transform

We want to understand the frequency ω of our signal. So, let's reparametrize the signal by ω instead of *x*:

f(x) → [Fourier Transform] → ***F(ω)***

For every ω from 0 to inf, ***F(ω)*** holds the amplitude *A* and phase ϕ of the corresponding sine $A\sin(\omega x+\phi)$

- How can *F* hold both? Complex number trick!

$$F(\omega)=R(\omega)+iI(\omega)$$

$$A=\pm\sqrt{R(\omega)^2+I(\omega)^2} \qquad \phi=\tan^{-1}\frac{I(\omega)}{R(\omega)}$$

We can always go back:

F(ω) → [Inverse Fourier Transform] → ***f(x)***

Frequency Spectra

Usually, amplitude is more interesting than phase:

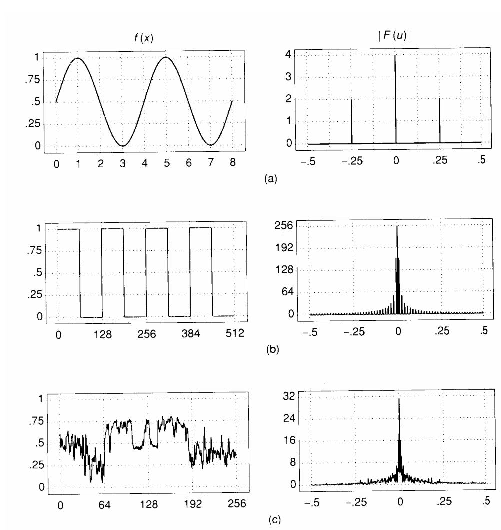

FT: Just a change of basis

M * $f(x) = F(\omega)$

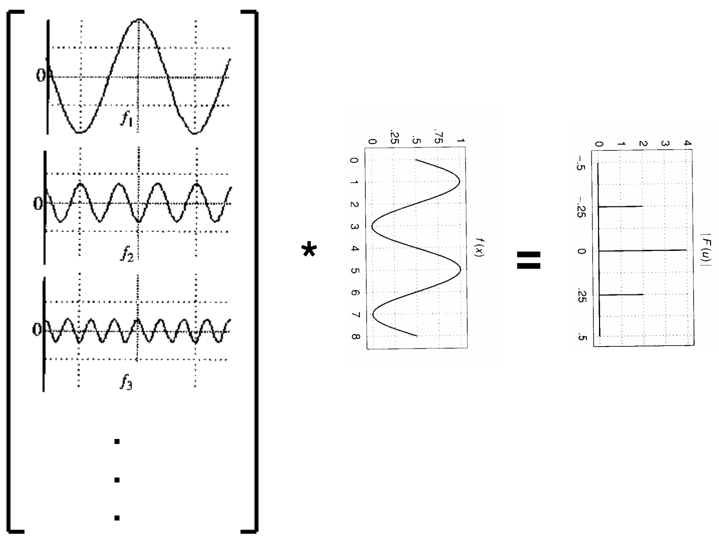

IFT: Just a change of basis

M^{-1} * $F(\omega) = f(x)$

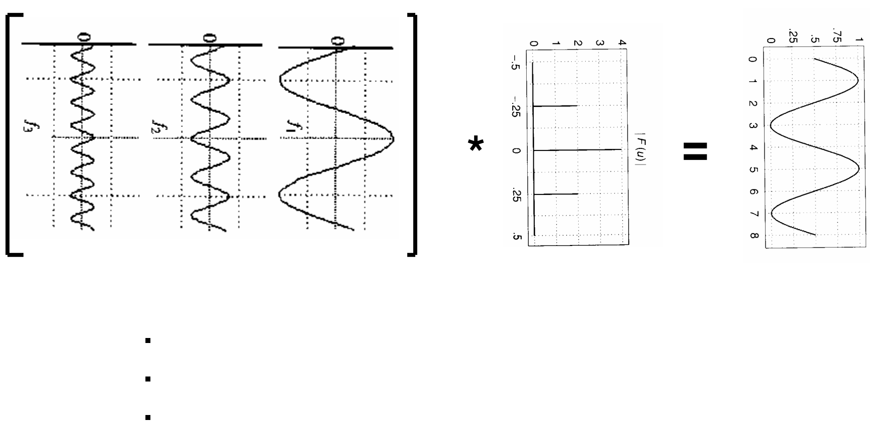

Finally: Scary Math

$$\text{Fourier Transform}: \quad F(\omega) = \int_{-\infty}^{+\infty} f(x) e^{-i\omega x} dx$$

$$\text{Inverse Fourier Transform}: \quad f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} F(\omega) e^{i\omega x} d\omega$$

Finally: Scary Math

$$\text{Fourier Transform}: \quad F(\omega) = \int_{-\infty}^{+\infty} f(x) e^{-i\omega x} dx$$

$$\text{Inverse Fourier Transform}: \quad f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} F(\omega) e^{i\omega x} d\omega$$

…not really scary: $e^{i\omega x} = \cos(\omega x) + i\sin(\omega x)$

is hiding our old friend: $A\sin(\omega x + \phi)$

phase can be encoded by sin/cos pair →

$$P\cos(x) + Q\sin(x) = A\sin(x+\phi)$$

$$A = \pm\sqrt{P^2 + Q^2} \qquad \phi = \tan^{-1}\left(\frac{P}{Q}\right)$$

So it's just our signal *f(x)* times sine at frequency ω

Extension to 2D

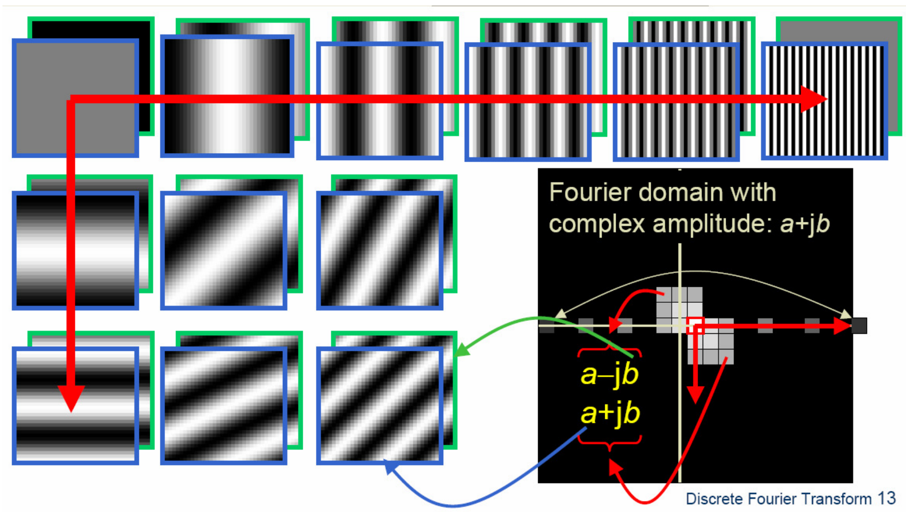

in Matlab, check out: imagesc(log(abs(fftshift(fft2(im)))));

This is the magnitude transform of the cheetah pic

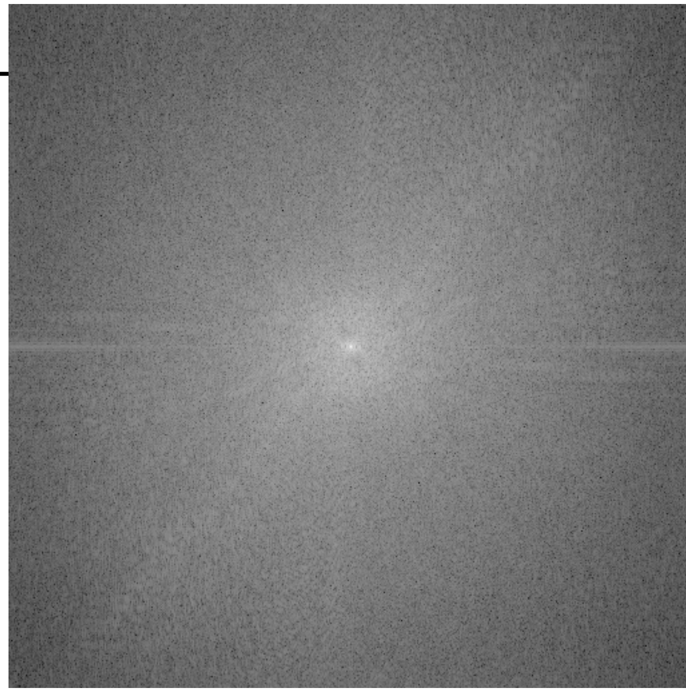

This is the phase transform of the cheetah pic

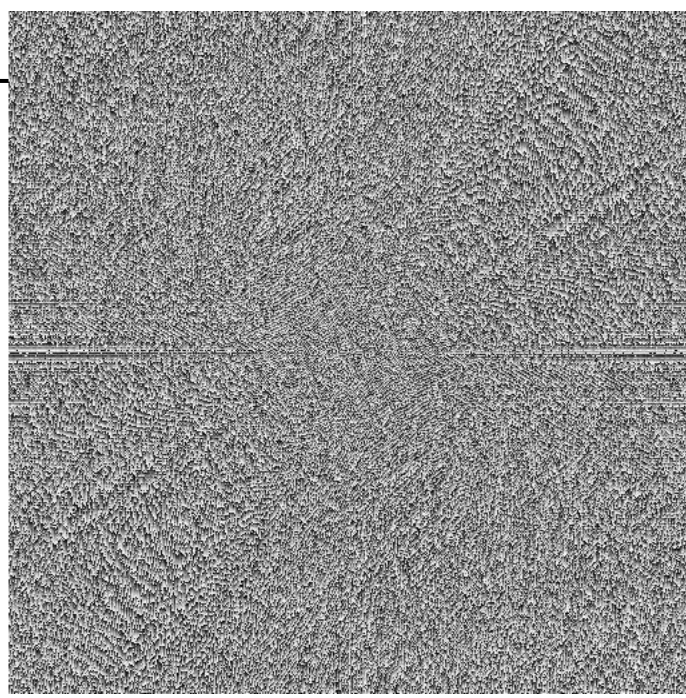

This is the magnitude transform of the zebra pic

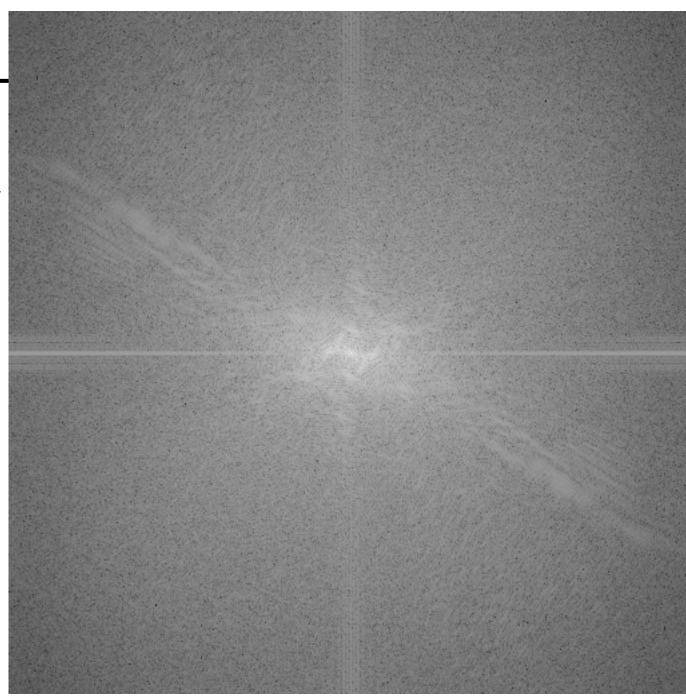

This is the phase transform of the zebra pic

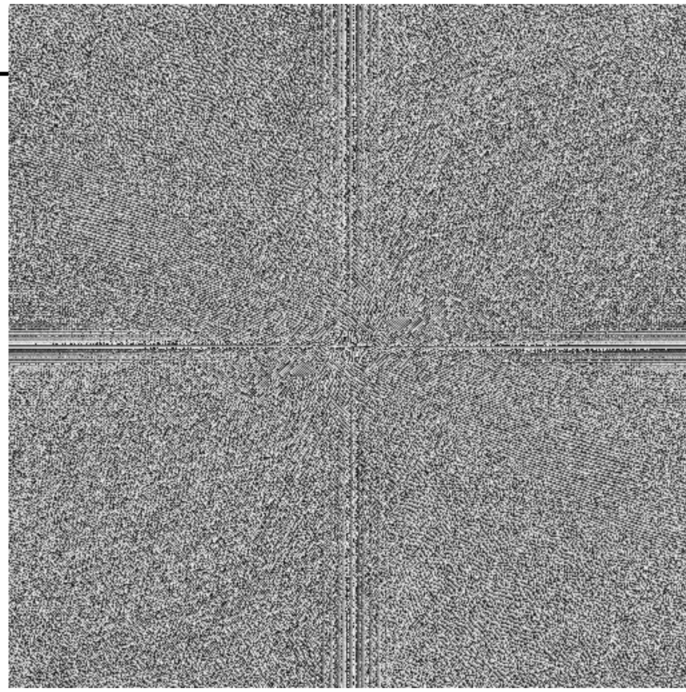

Curious things about FT on images

The magnitude spectra of all natural images quite similar

- Heavy on low-frequencies, falling off in high frequences
- Will any image be like that, or is it a property of the world we live in?

Most information in the image is carried in the phase, not the amplitude

- Seems to be a fact of life
- Not quite clear why

Reconstruction with zebra phase, cheetah magnitude

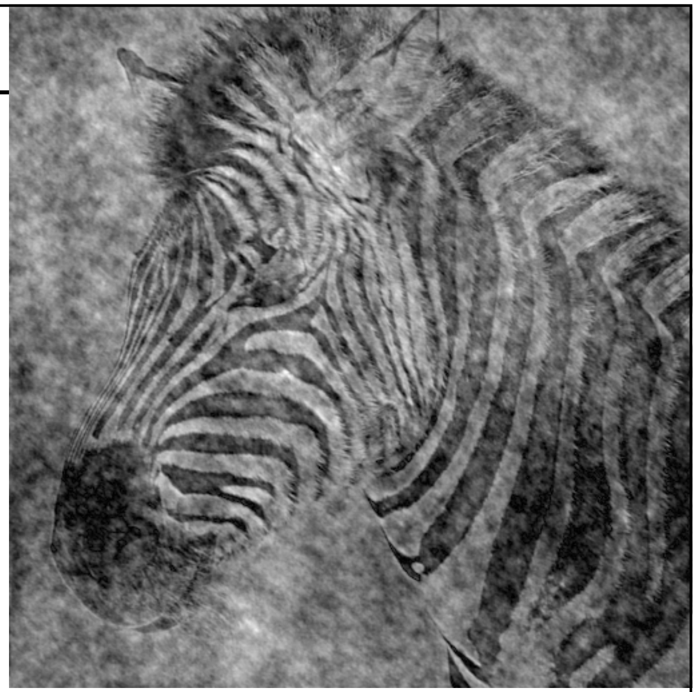

Reconstruction with cheetah phase, zebra magnitude

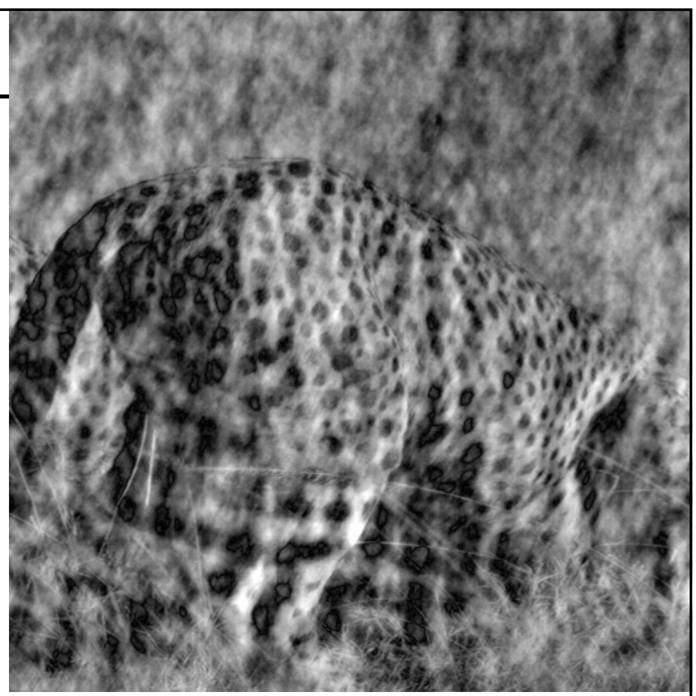

Various Fourier Transform Pairs

Important facts

- The Fourier transform is linear
- There is an inverse FT
- if you scale the function's argument, then the transform's argument scales the other way. This makes sense --- if you multiply a function's argument by a number that is larger than one, you are stretching the function, so that high frequencies go to low frequencies
- The FT of a Gaussian is a Gaussian.

The convolution theorem

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

Slide by David Forsyth

2D convolution theorem example

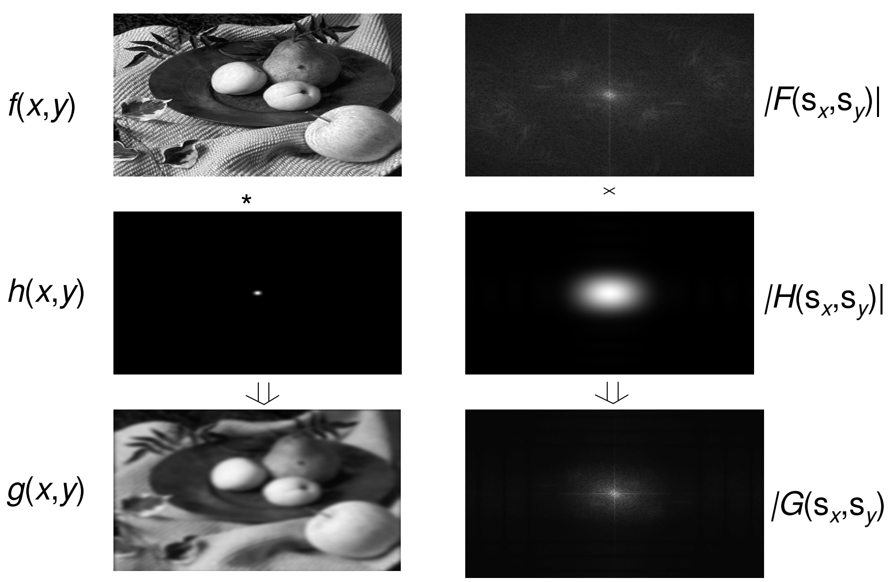

Slide by Steve Seitz

Low-pass, Band-pass, High-pass filters

low-pass:

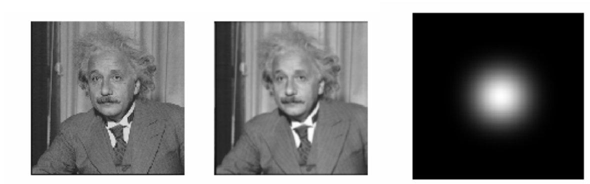

band-pass:

what's high-pass?