Lecture 2: 2D Fourier transforms and applications

B14 Image Analysis Michaelmas 2014 A. Zisserman

- Fourier transforms and spatial frequencies in 2D
  - Definition and meaning

- The Convolution Theorem
  - Applications to spatial filtering

- The Sampling Theorem and Aliasing

Much of this material is a straightforward generalization of the 1D Fourier analysis with which you are familiar.

Reminder: 1D Fourier Series

Spatial frequency analysis of a step edge

$$f(x) = \begin{cases} -1 & \text{if } x < 0 \\ 1 & \text{otherwise} \end{cases}$$

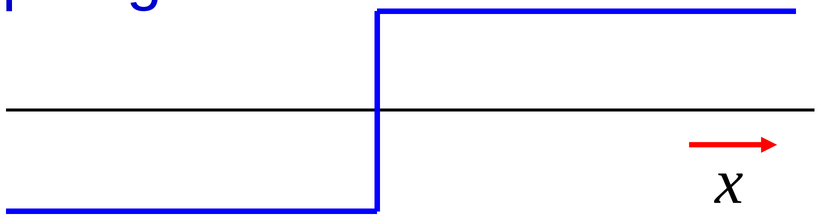

Fourier decomposition

Fourier Series

$$f(x) = \sum_n a_n \sin nx$$

$$a_n = \frac{1}{\pi}\int_{-\pi}^{\pi} f(x) \sin nx \, dx$$

$$= \frac{2}{\pi}\int_{0}^{\pi} \sin nx \, dx = \begin{cases} \frac{4}{n\pi} & \text{if n odd} \\ 0 & \text{otherwise} \end{cases}$$

$$f(x) = \sum_{n=1} \frac{4}{(2n-1)\pi} \sin(2n-1)x$$

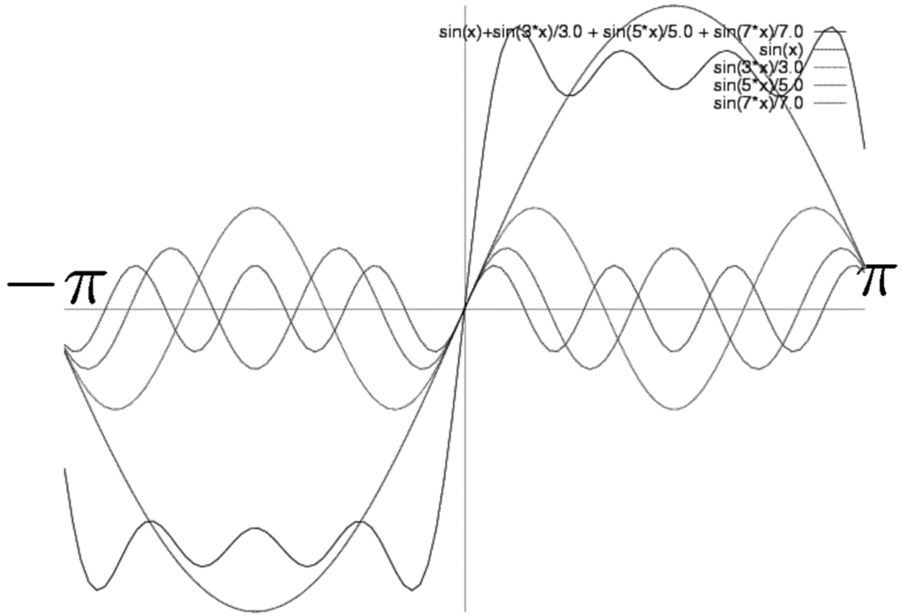

Fourier series reminder

Example

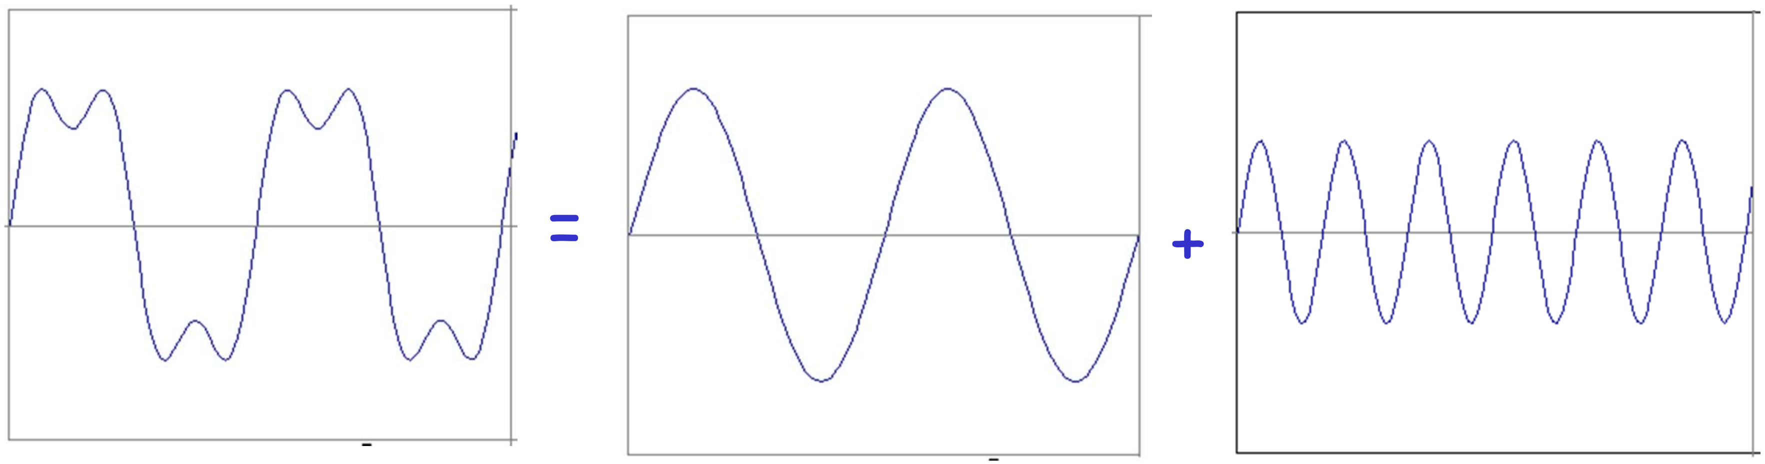

$$f(x) \;=\; \sin x + \frac{1}{3}\sin 3x + \ldots$$

Fourier series for a square wave

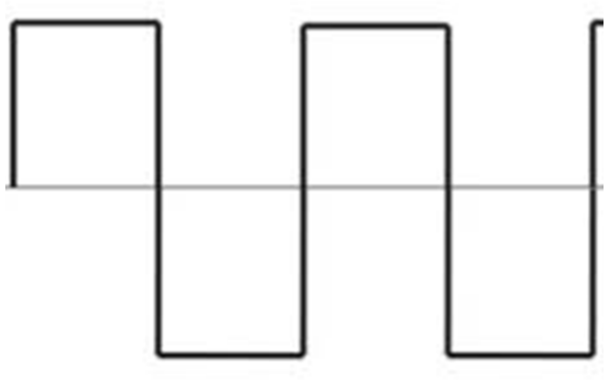

= $$f(x) = \sum_{n=1,3,5,\ldots} \frac{1}{n} \sin nx$$

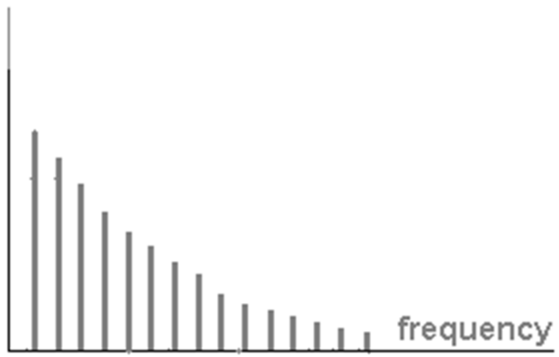

Fourier series: just a change of basis

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

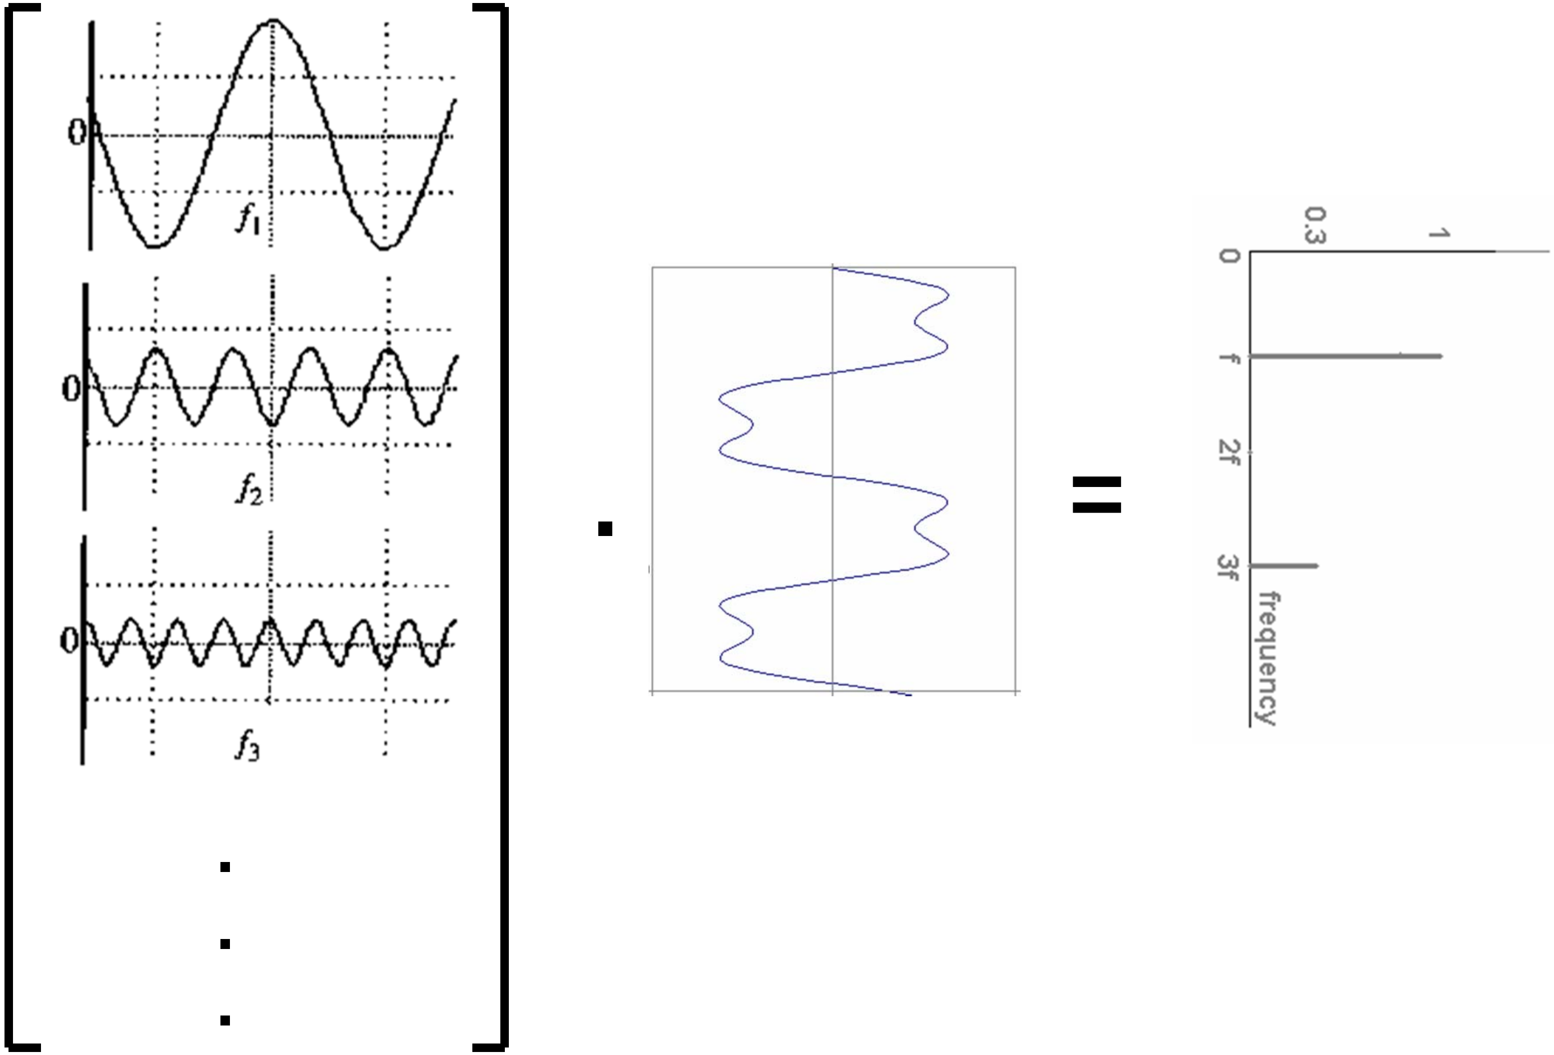

Inverse FT: Just a change of basis

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

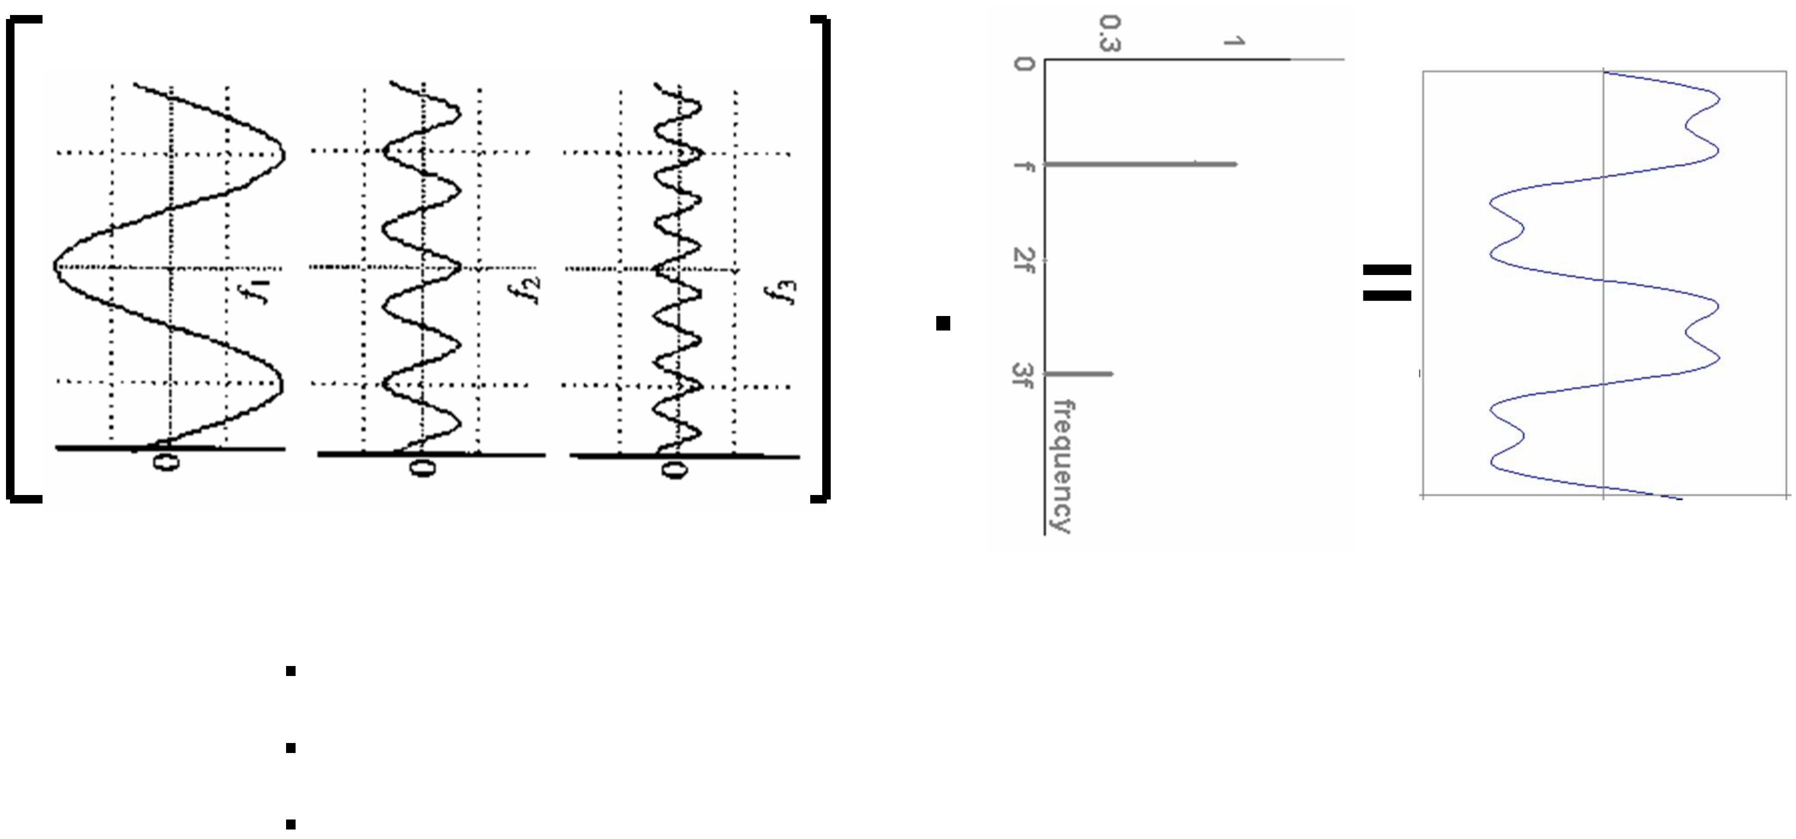

1D Fourier Transform

Reminder transform pair - definition

$$F(u) = \int_{-\infty}^{\infty} f(x) e^{-j2\pi ux}\, dx,$$
$$f(x) = \int_{-\infty}^{\infty} F(u) e^{j2\pi ux}\, du$$

Example

$$f(x) = \begin{cases} 1, & |x| < \frac{X}{2}, \\ 0, & |x| \geq \frac{X}{2}. \end{cases}$$

$$\begin{aligned} F(u) &= \int_{-\infty}^{\infty} f(x) e^{-j2\pi ux} dx \\ &= \int_{-X/2}^{X/2} e^{-j2\pi ux} dx \\ &= \frac{1}{-j2\pi u}[e^{-j2\pi uX/2} - e^{j2\pi uX/2}] \\ &= X\frac{\sin(\pi X u)}{(\pi X u)} = X\text{sinc}(\pi \text{X} \text{u}). \end{aligned}$$

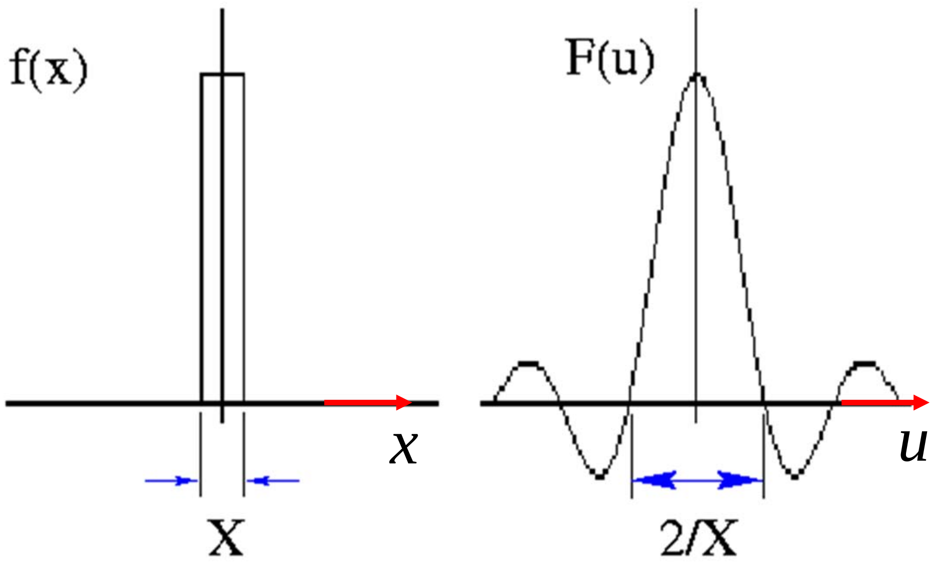

2D Fourier transforms

2D Fourier transform

Definition

$$F(u,v) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)e^{-j2\pi(ux+vy)}\,dx\,dy,$$
$$f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(u,v)e^{j2\pi(ux+vy)}\,du\,dv$$

where u and v are spatial frequencies.

Also will write FT pairs as $f(x,y) \Leftrightarrow F(u,v)$.

- $F(u,v)$ is complex in general,

$$F(u,v) = F_{\mathrm{R}}(u,v) + jF_{\mathrm{I}}(u,v)$$

- $|F(u,v)|$ is the magnitude spectrum
- $\arctan(F_{\mathrm{I}}(u,v)/F_{\mathrm{R}}(u,v))$ is the phase angle spectrum.
- Conjugacy: $f^*(x,y) \Leftrightarrow F(-u,-v)$
- Symmetry: $f(x,y)$ is even if $f(x,y) = f(-x,-y)$

Sinusoidal Waves

In 1D the Fourier transform is based on a decompostion into functions $e^{j2\pi ux} = \cos 2\pi ux + j \sin 2\pi ux$ which form an orthogonal basis. Similarly in 2D

$$e^{j2\pi(ux+vy)} = \cos 2\pi(ux+vy) + j \sin 2\pi(ux+vy)$$

The real and imaginary terms are sinusoids on the x, y plane. The maxima and minima of $\cos 2\pi(ux+vy)$ occur when

$$2\pi(ux+vy) = n\pi$$

write $ux+vy$ using vector notation with $\mathbf{u} = (u,v)^\top, \mathbf{x} = (x,y)^\top$ then

$$2\pi(ux+vy) = 2\pi\mathbf{u}.\mathbf{x} = n\pi$$

are sets of equally spaced parallel lines with normal $\mathbf{u}$ and wavelength $1/\sqrt{u^2+v^2}$.

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.

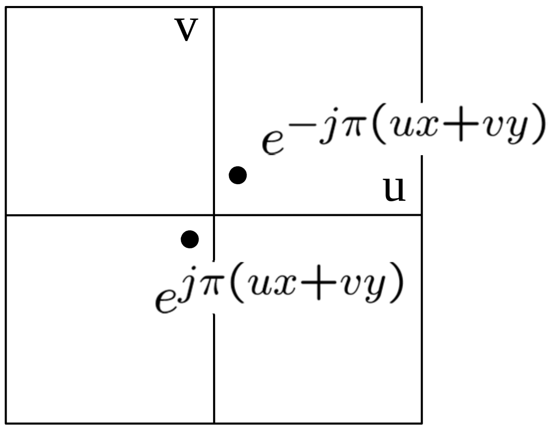

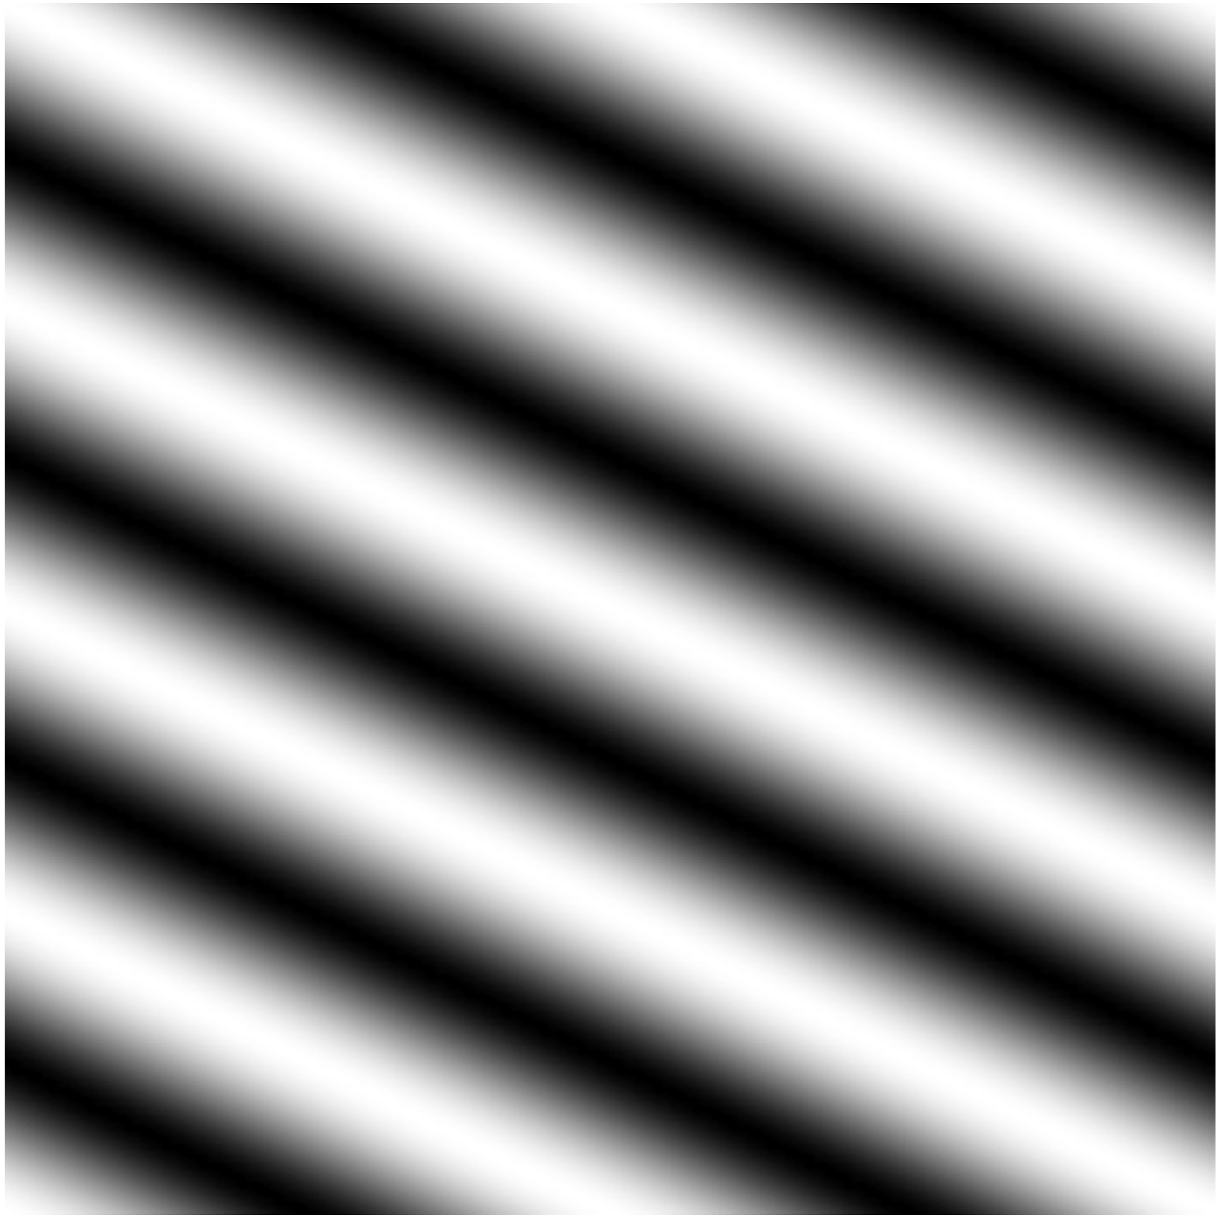

slide: B. Freeman

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

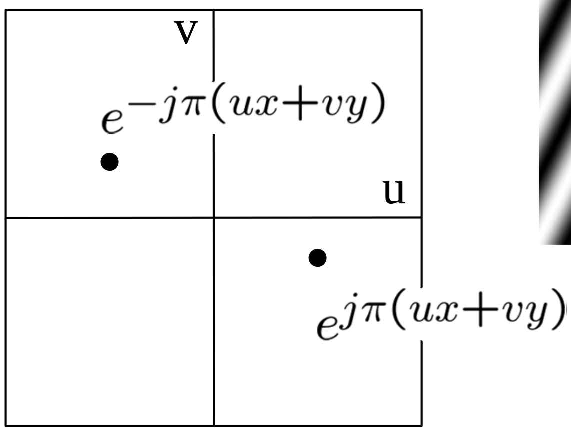

And larger still...

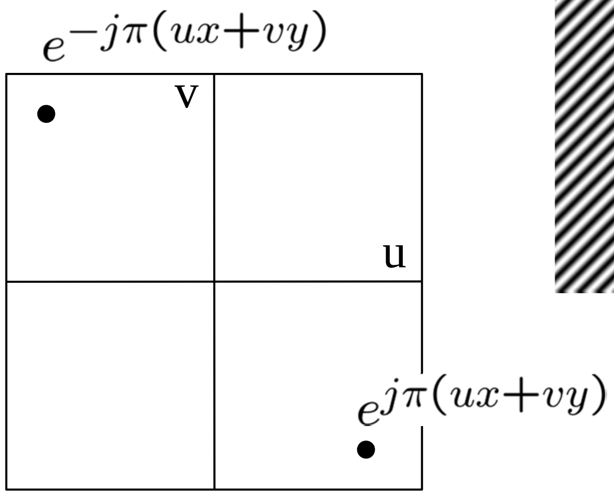

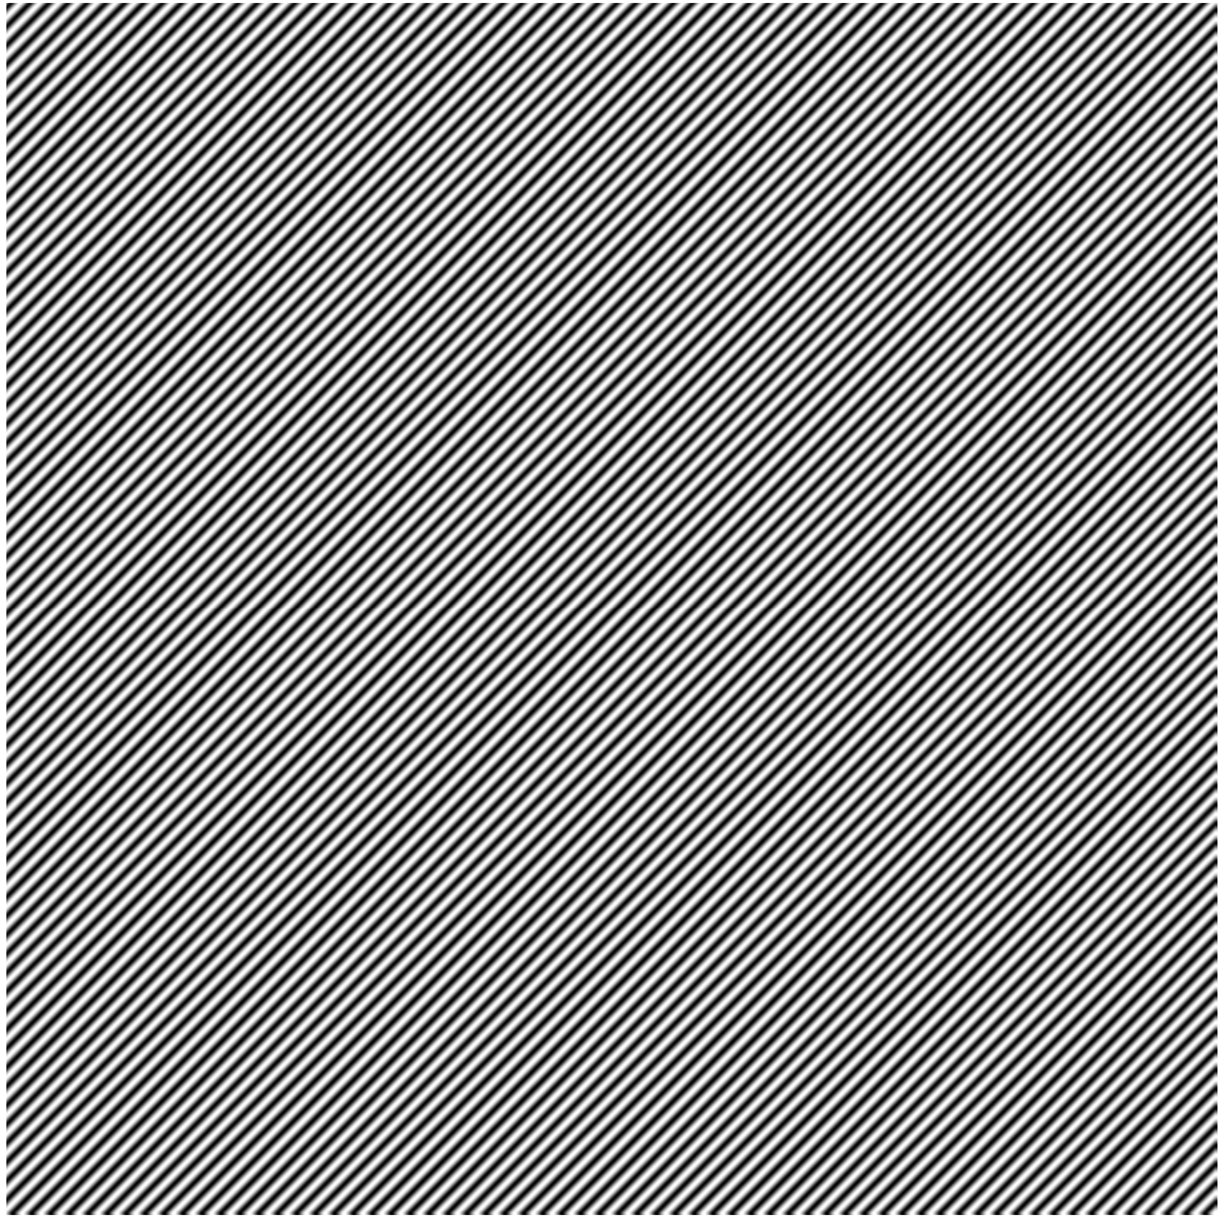

Some important Fourier Transform Pairs

FT pair example 1

rectangle centred at origin
with sides of length X and Y

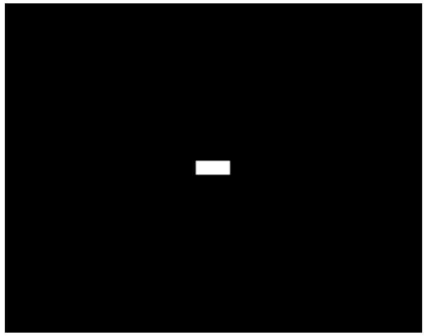

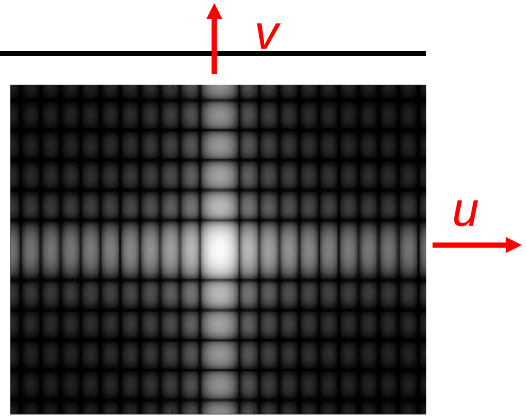

$f(x,y)$ $|F(u,v)|$

$$F(u,v) = \int\int f(x,y) e^{-j2\pi(ux+vy)} dx dy,$$

$$= \int_{-X/2}^{X/2} e^{-j2\pi ux} dx \int_{-Y/2}^{Y/2} e^{-j2\pi vy} dy,$$ separability

$$= \left[\frac{e^{-j2\pi ux}}{-j2\pi u}\right]_{-X/2}^{X/2} \left[\frac{e^{-j2\pi vy}}{-j2\pi v}\right]_{-Y/2}^{Y/2},$$

$$= \frac{1}{-j2\pi u}\left[e^{-juX} - e^{juX}\right] \frac{1}{-j2\pi v}\left[e^{-jvY} - e^{jvY}\right],$$

$$= XY \left[\frac{sin(\pi X u)}{\pi X u}\right]\left[\frac{sin(2\pi Y v)}{\pi Y v}\right]$$

$$= XY \mathrm{sinc}(\pi \mathrm{Xu}) \mathrm{sinc}(\pi \mathrm{Yv}).$$

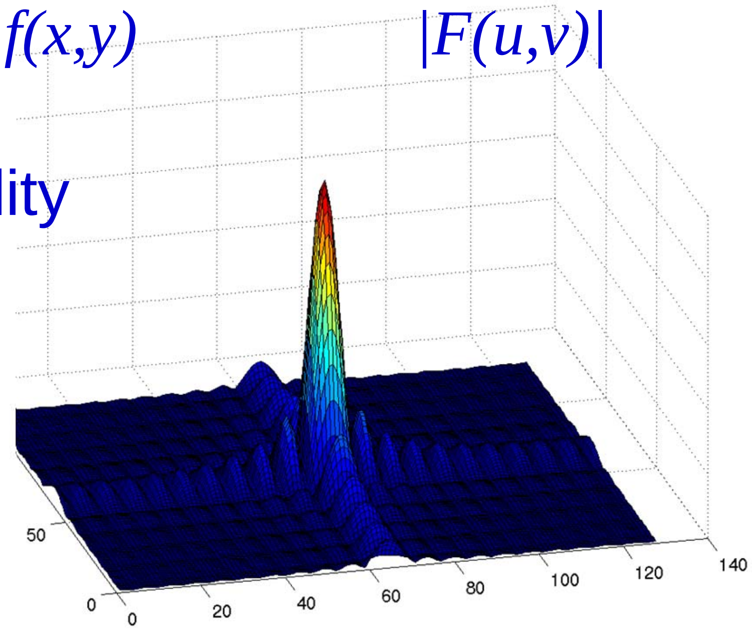

$|F(u,v)|$

FT pair example 2

Gaussian centred on origin

$$f(r) = \frac{1}{2\pi\sigma^2} e^{-r^2/2\sigma^2}$$

where $r^2 = x^2 + y^2$.

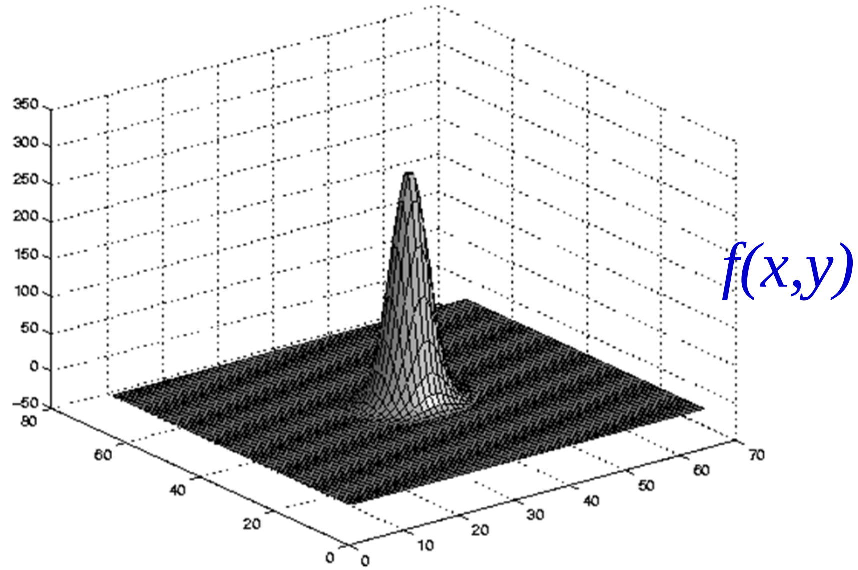

$$F(u,v) = F(\rho) = e^{-2\pi^2\rho^2\sigma^2}$$

where $\rho^2 = u^2 + v^2$.

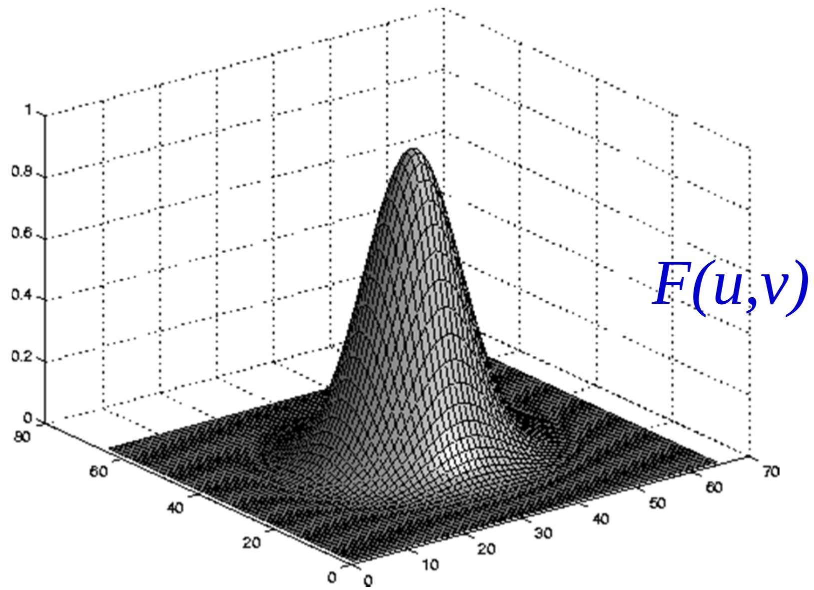

- FT of a Gaussian is a Gaussian
- Note inverse scale relation

FT pair example 3

Circular disk unit height and radius *a* centred on origin

$$f(x,y) = \begin{cases} 1, & |r| < a, \\ 0, & |r| \geq a. \end{cases}$$

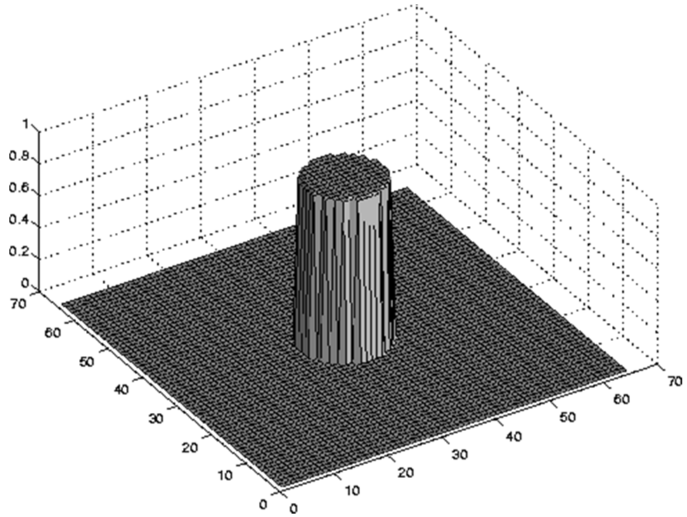

f(x,y)

$$F(u,v) = F(\rho) = aJ_1(\pi a \rho)/\rho$$

where $J_1(x)$ is a Bessel function.

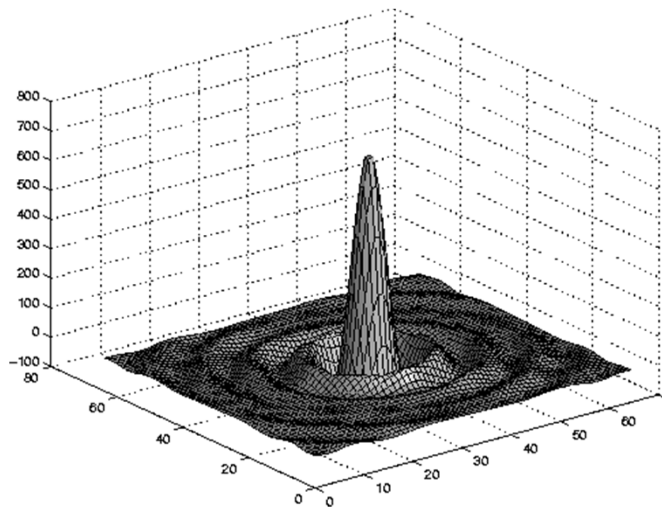

F(u,v)

- rotational symmetry
- a ‘2D’ version of a sinc

FT pairs example 4

$$f(x,y) = \delta(x,y) = \delta(x)\delta(y)$$

$$F(u,v) = \int\int \delta(x,y) e^{-j2\pi(ux+vy)}\, dx dy$$

$$= 1$$

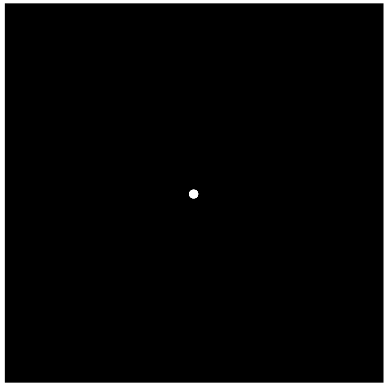

f(x,y)

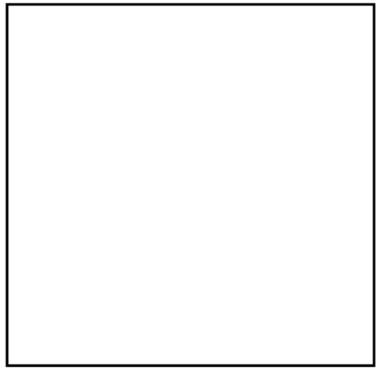

F(u,v)

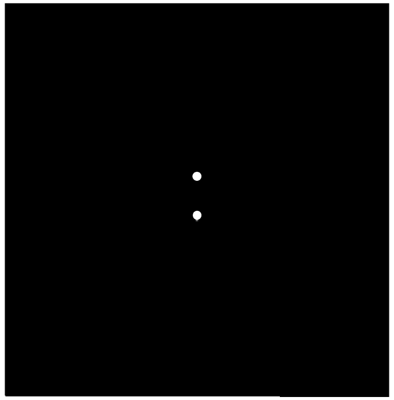

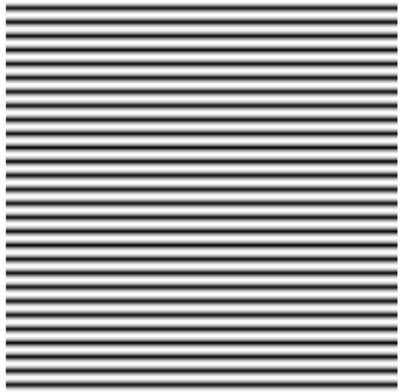

$$f(x,y) = \frac{1}{2}\left(\delta(x, y-a) + \delta(x, y+a)\right)$$

$$F(u,v) = \frac{1}{2}\int\int \left(\delta(x,y-a) + \delta(x,y+a)\right) e^{-j2\pi(ux+vy)}\, dx dy$$

$$= \frac{1}{2}\left(e^{-j2\pi av} + e^{j2\pi av}\right) = \cos 2\pi av$$

Summary

The spatial function $f(x, y)$

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) e^{j2\pi(ux+vy)} \, du \, dv$$

is decomposed into a weighted sum of 2D orthogonal basis functions in a similar manner to decomposing a vector onto a basis using scalar products.

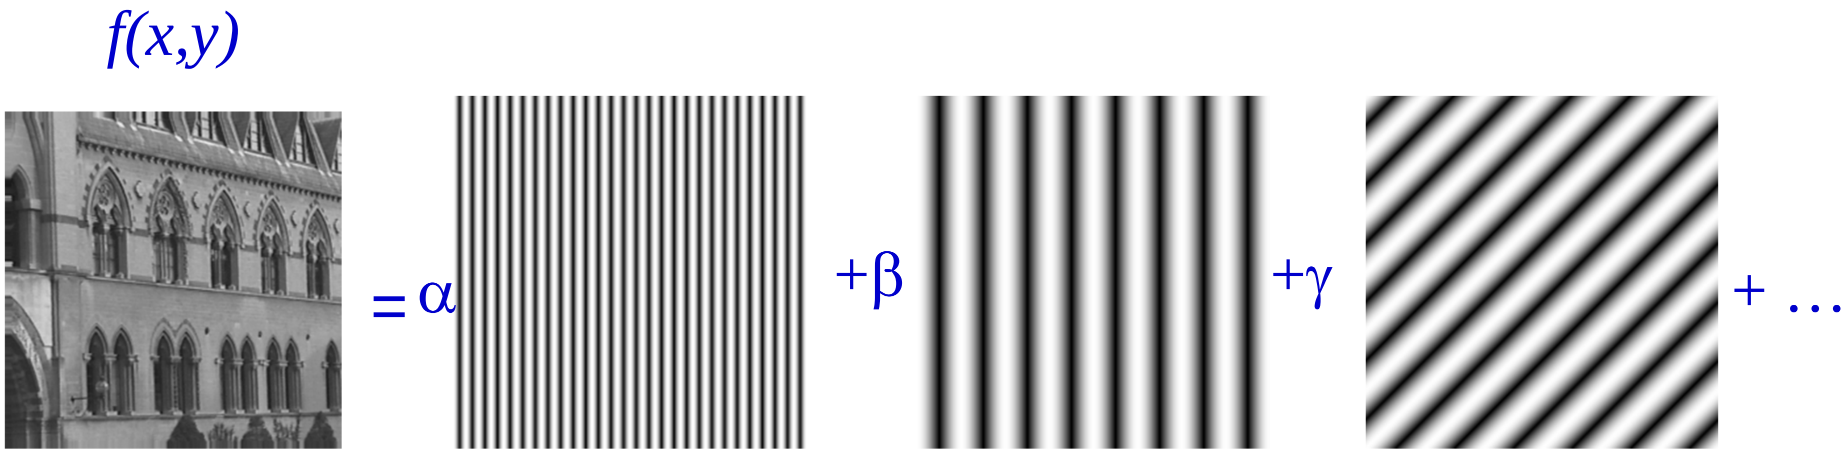

Example: action of filters on a real image

original | low pass | high pass

f(x,y)

|F(u,v)|

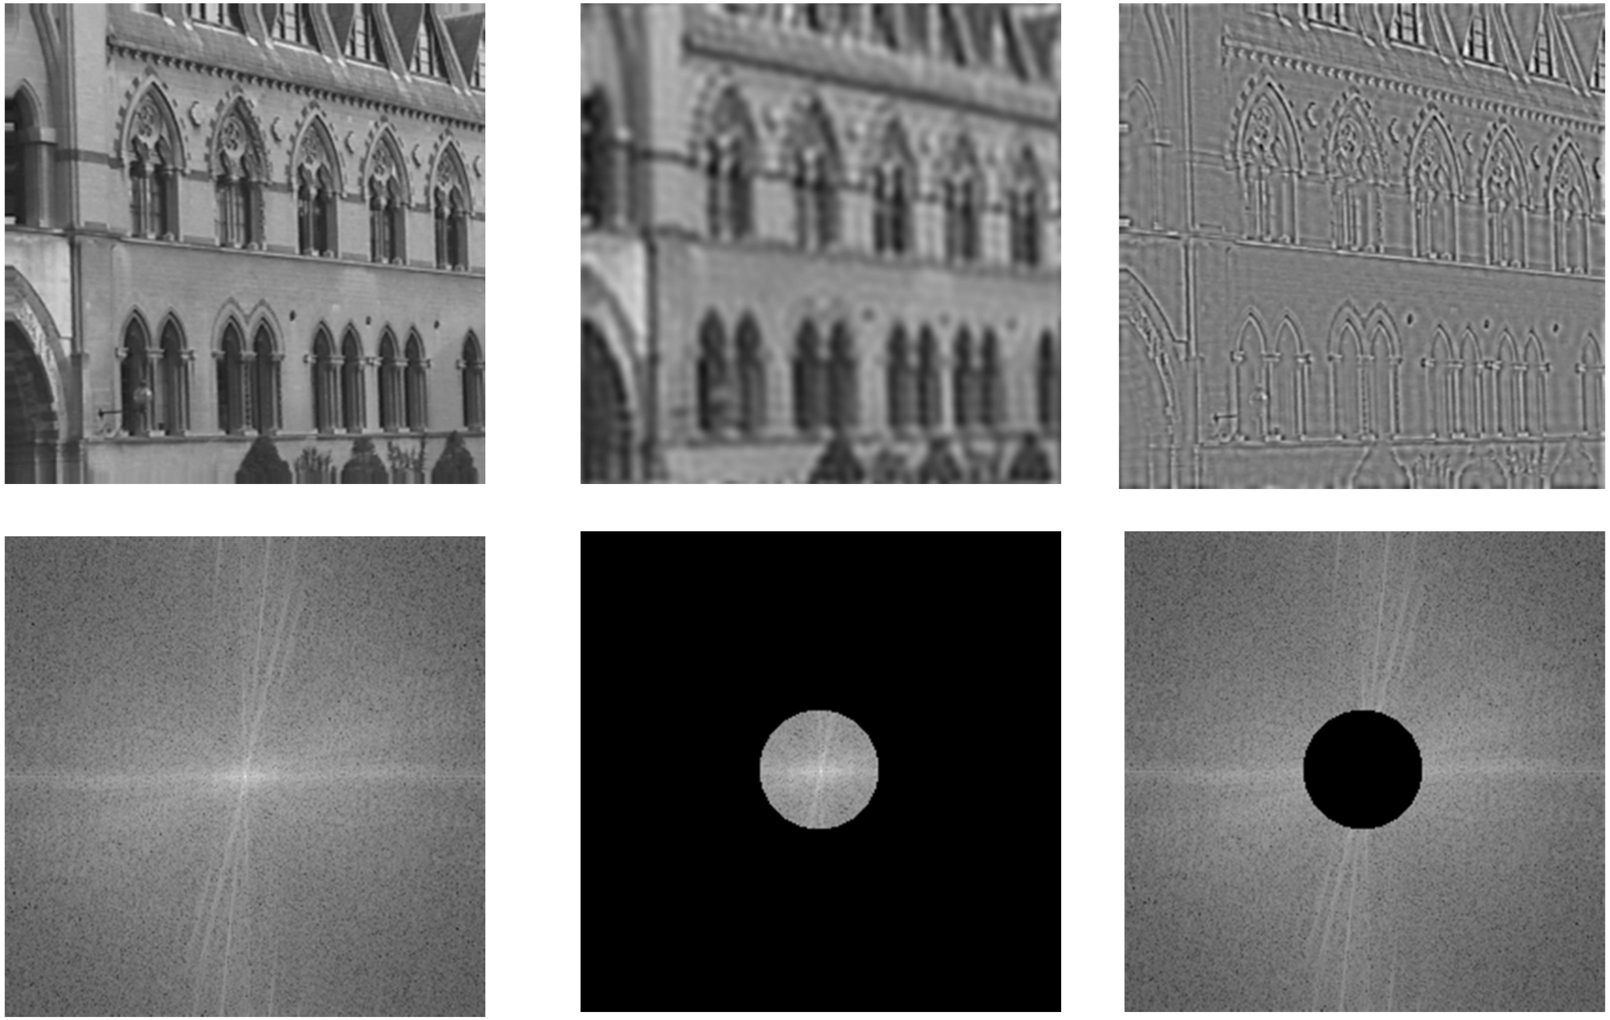

Example 2D Fourier transform

Image with periodic structure

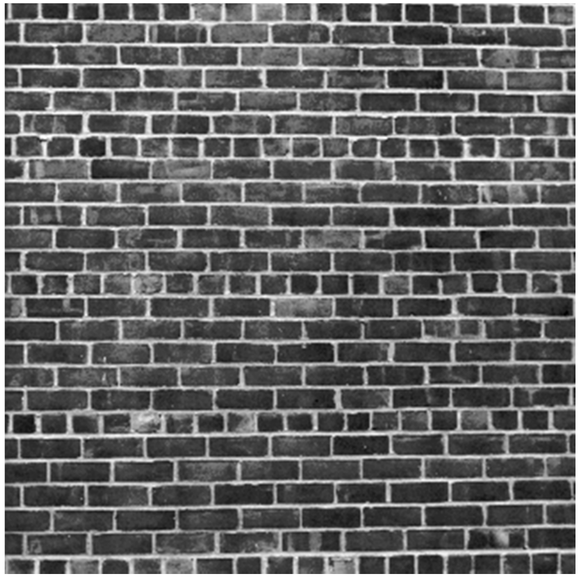

f(x,y)

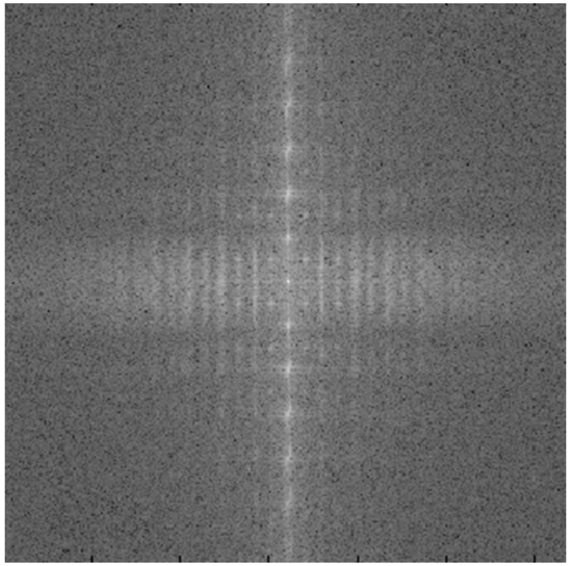

|F(u,v)|

FT has peaks at spatial frequencies of repeated texture

Example – Forensic application

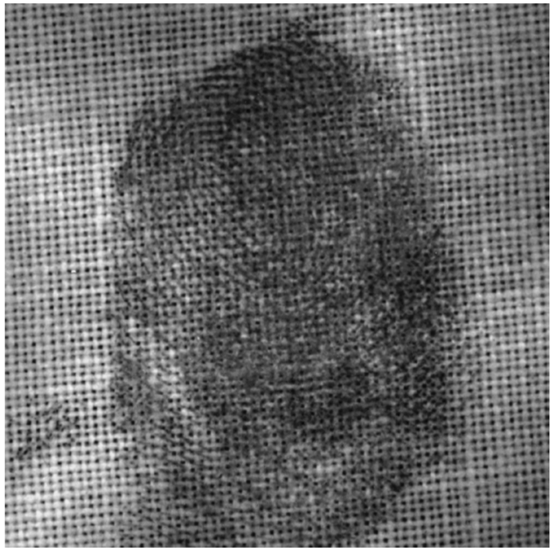

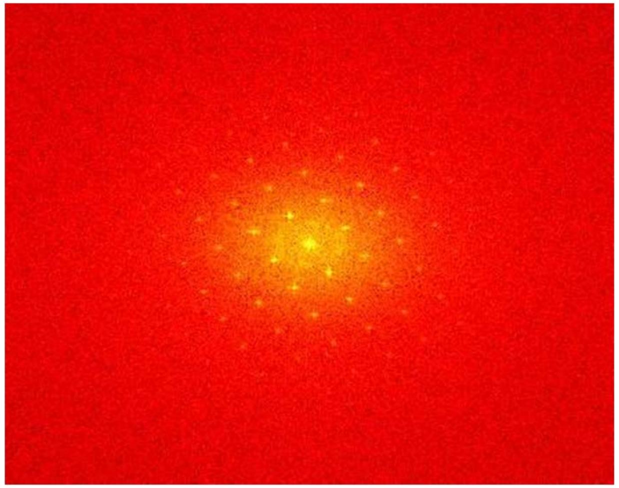

$|F(u,v)|$

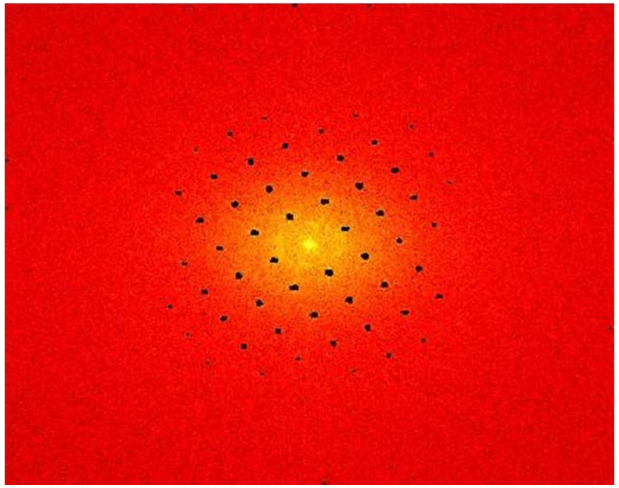

remove peaks

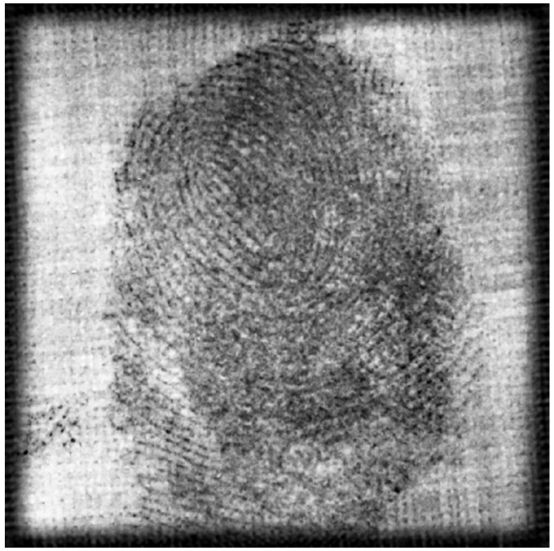

Periodic background removed

Example – Image processing

Lunar orbital image (1966)

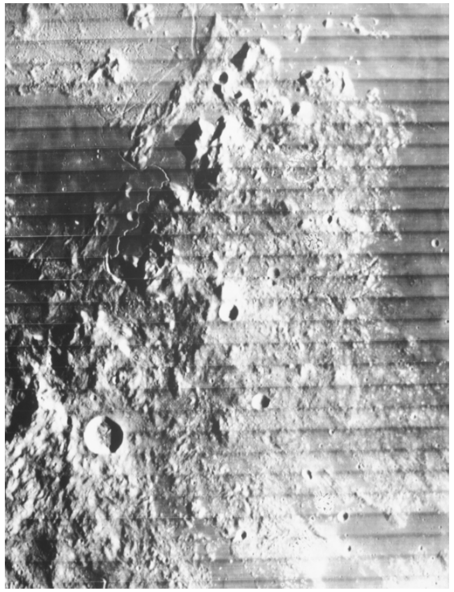

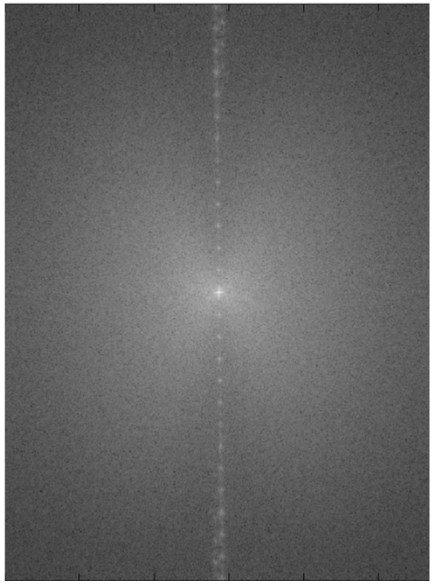

$|F(u,v)|$

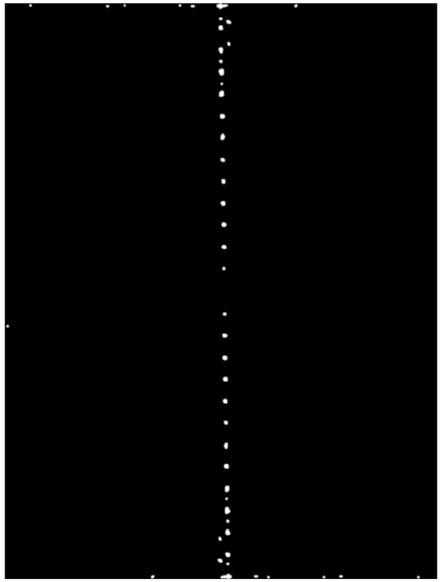

remove peaks

join lines removed

Magnitude vs Phase

f(x,y)

|F(u,v)|

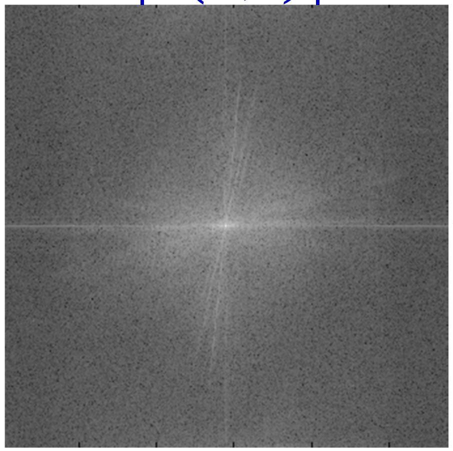

cross-section

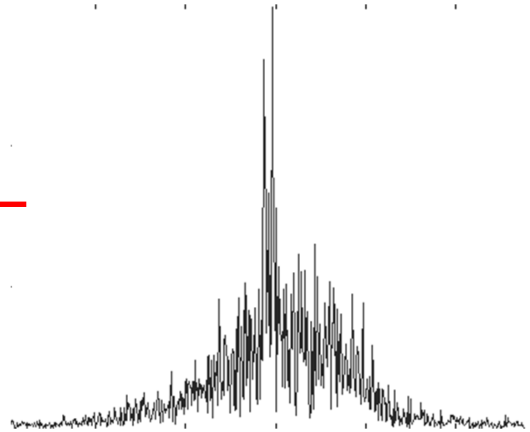

phase F(u,v)

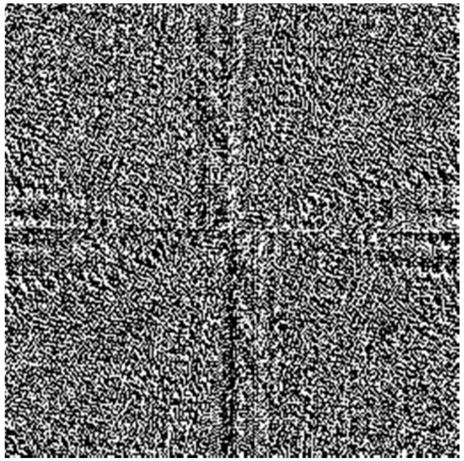

- |f(u,v)| generally decreases with higher spatial frequencies
- phase appears less informative

The importance of phase

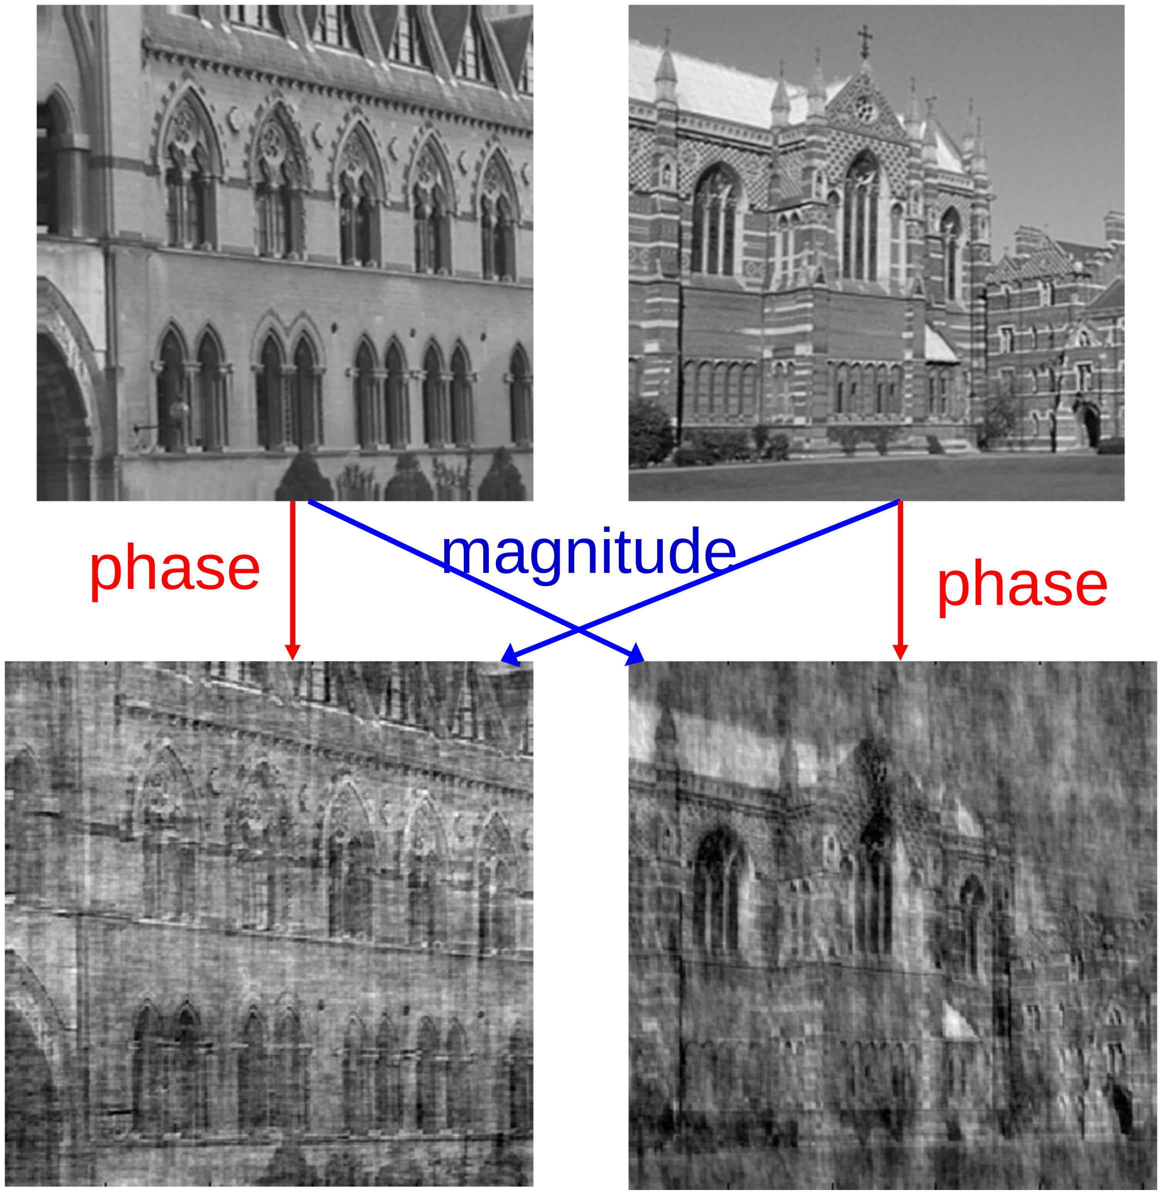

A second example

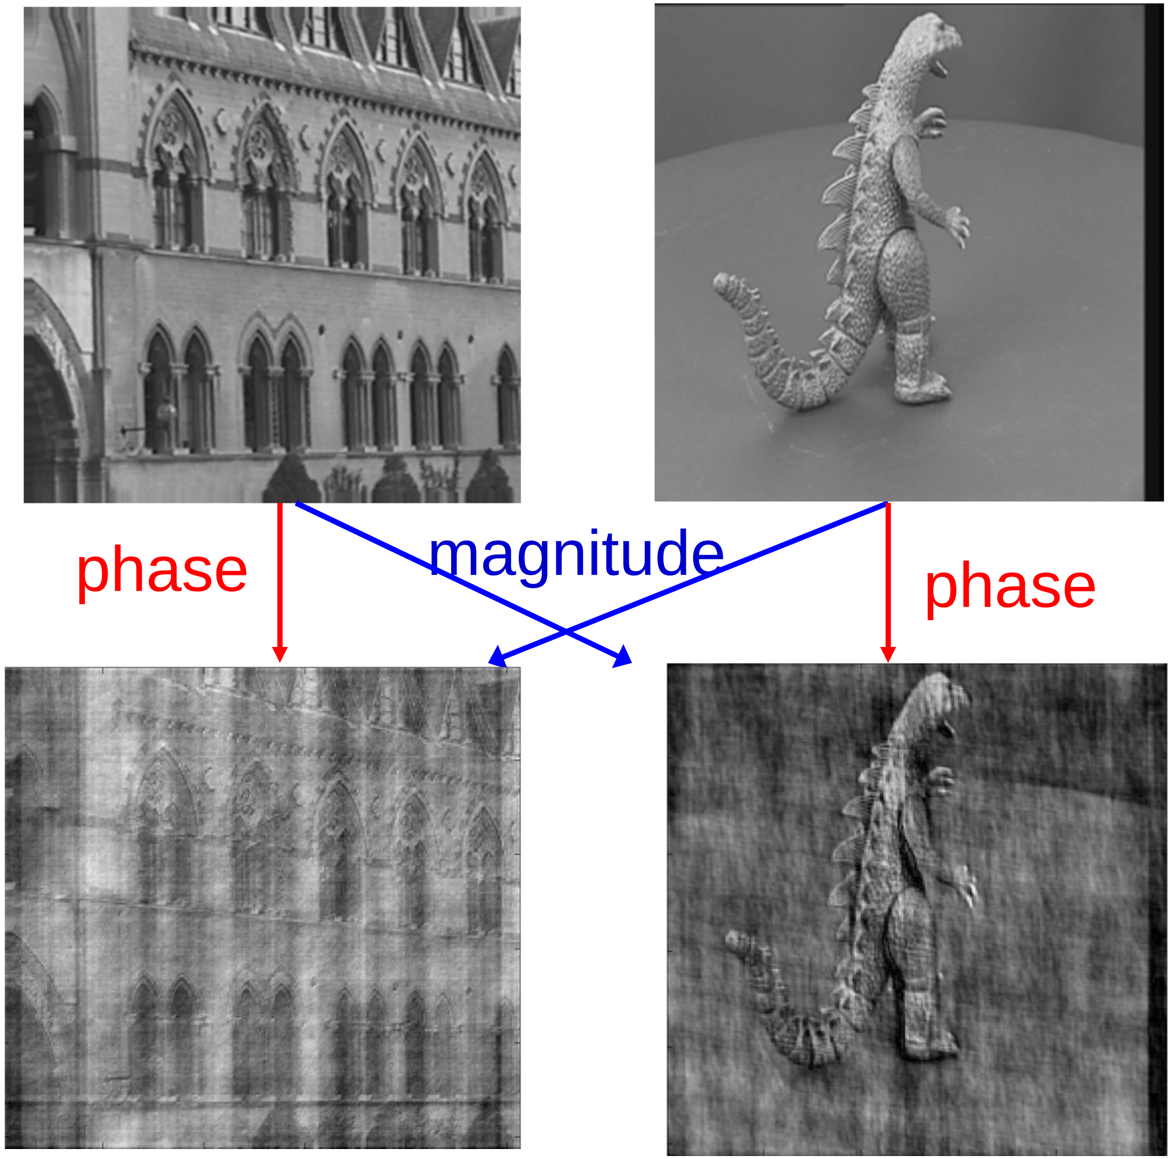

Transformations

As in the 1D case FTs have the following properties

- Linearity

$$\alpha f(x,y) + \beta g(x,y) \Leftrightarrow \alpha F(u,v) + \beta G(u,v).$$

- Similarity

$$f(ax,by) \Leftrightarrow \frac{1}{ab} F(\frac{u}{a}, \frac{v}{b}).$$

This applies, for example, when an image is scaled

- Shift

$$f(x-a, y-b) \Leftrightarrow e^{j2\pi(au+bv)} F(u,v)$$

This might apply, for example, if an object moved.

In 2D can also rotate, shear etc

Under an affine transformation: $\mathbf{x} \to \mathtt{A}\mathbf{x}$ $\mathbf{u} \to \mathtt{A}^{-\top}\mathbf{u}$

Example

How does $F(u,v)$ transform if $f(x,y)$ is rotated by 45 degrees?

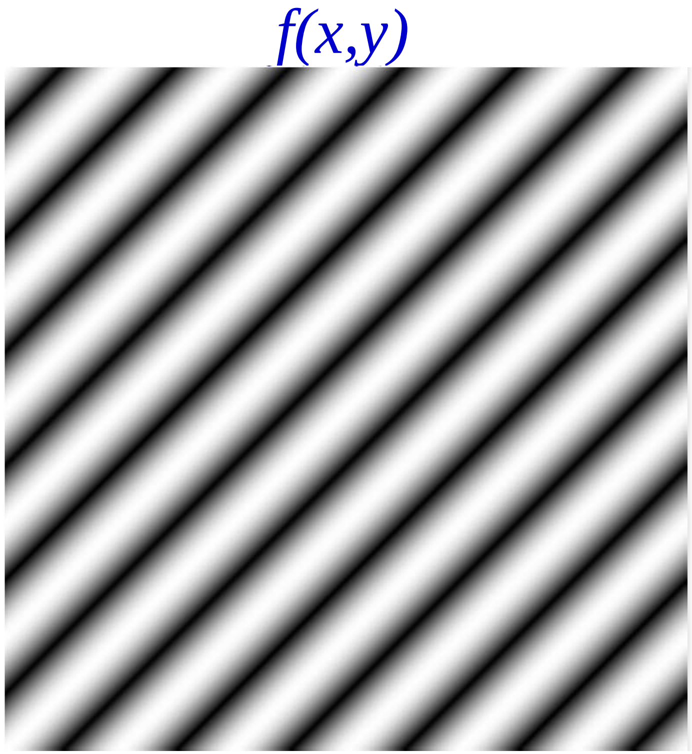

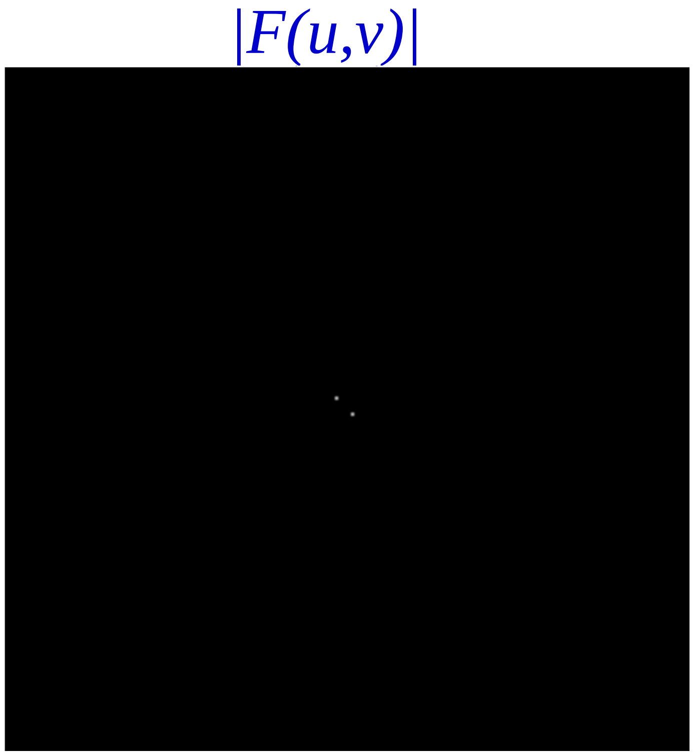

If $\mathtt{A} = \mathtt{R}$ then $\mathtt{A}^{-\top} = \mathtt{R}$.

i.e. FT undergoes the same rotation.

The convolution theorem

Filtering vs convolution in 1D

$$g(x) = \sum_i f(x+i)h(i)$$

filtering f(x) with h(x)

f(x)	100 | 200 | 100 | 200 | 90 | 80 | 80 | 100 | 100
h(x)	1/4 | 1/2 | 1/4
g(x)	| 150 | | | | | | |

molecule/template/kernel

$$g(x) = \int f(u)h(x-u)\,du$$

convolution of f(x) and h(x)

$$= \int f(x+u')h(-u')\,du'$$

after change of variable $u' = u - x$

$$= \sum_i f(x+i)h(-i)$$

- note negative sign (which is a reflection in x) in convolution
- *h(x)* is often symmetric (even/odd), and then (e.g. for even)

$$g(x) = \sum_i f(x+i)h(i)$$

Filtering vs convolution in 2D

convolution

$$g(x,y) = h(x,y) * f(x,y) = f(x,y) * h(x,y)$$
$$= \int\int f(u,v)h(x-u,y-v)\,du\,dv$$

filtering

image *f(x,y)*

filter / kernel *h(x,y)*

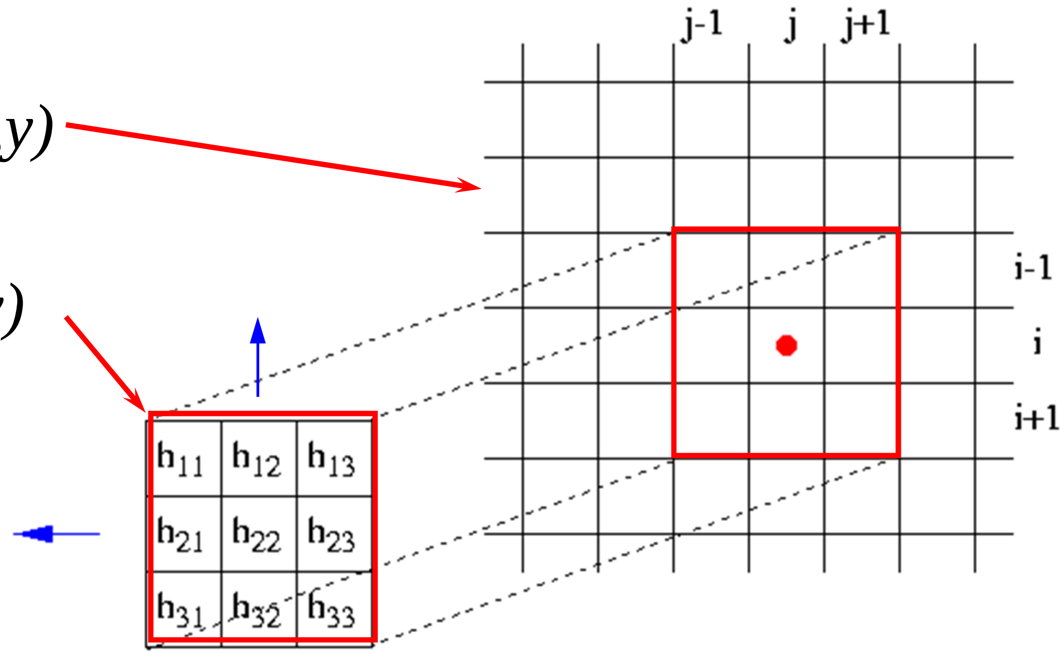

$$\begin{aligned} g(x,y) = \; & h_{11}\, f(i-1,j-1) + h_{12}\, f(i-1,j) + h_{13}\, f(i-1,j+1) + \\ & h_{21}\, f(i,j-1) + h_{22}\, f(i,j) + h_{23}\, f(i,j+1) + \\ & h_{31}\, f(i+1,j-1) + h_{32}\, f(i+1,j) + h_{33}\, f(i+1,j+1) \end{aligned}$$

for convolution, reflect filter in x and y axes

Convolution

- Convolution:
  - Flip the filter in both dimensions (bottom to top, right to left)

$$g[i,j] \;=\; \sum_{u=-k}^{k}\sum_{v=-k}^{k} h[u,v]\,f[i-u,j-v]$$

convolution with h

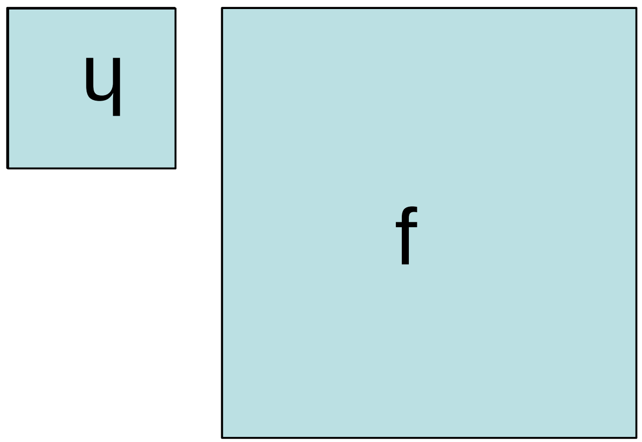

slide: K. Grauman

Filtering vs convolution in 2D in Matlab

2D filtering

f=image
h=filter

- `g=filter2(h,f);`

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

2D convolution

- `g=conv2(h,f);`

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

Convolution theorem

$$f(x,y) * h(x,y) \quad \Leftrightarrow \quad F(u,v)H(u,v)$$

Space convolution = frequency multiplication

In words: the Fourier transform of the convolution of two functions is the product of their individual Fourier transforms

Proof: exercise

Why is this so important?

Because linear filtering operations can be carried out by simple multiplications in the Fourier domain

The importance of the convolution theorem

It establishes the link between operations in the frequency domain and the action of linear spatial filters

Example smooth an image with a Gaussian spatial filter

*

Gaussian
scale=20 pixels

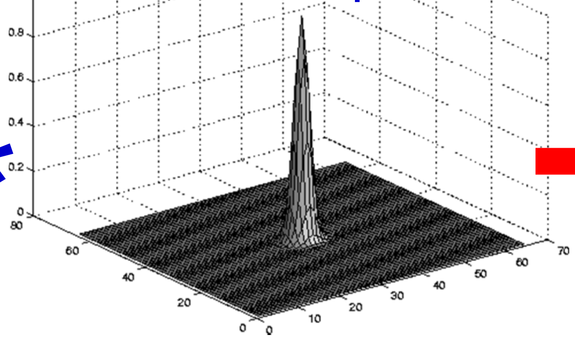

1. Compute FT of image and FT of Gaussian
2. Multiply FT's
3. Compute inverse FT of the result.

$$f(x, y) * g(x, y) \Leftrightarrow F(u, v)G(u, v)$$

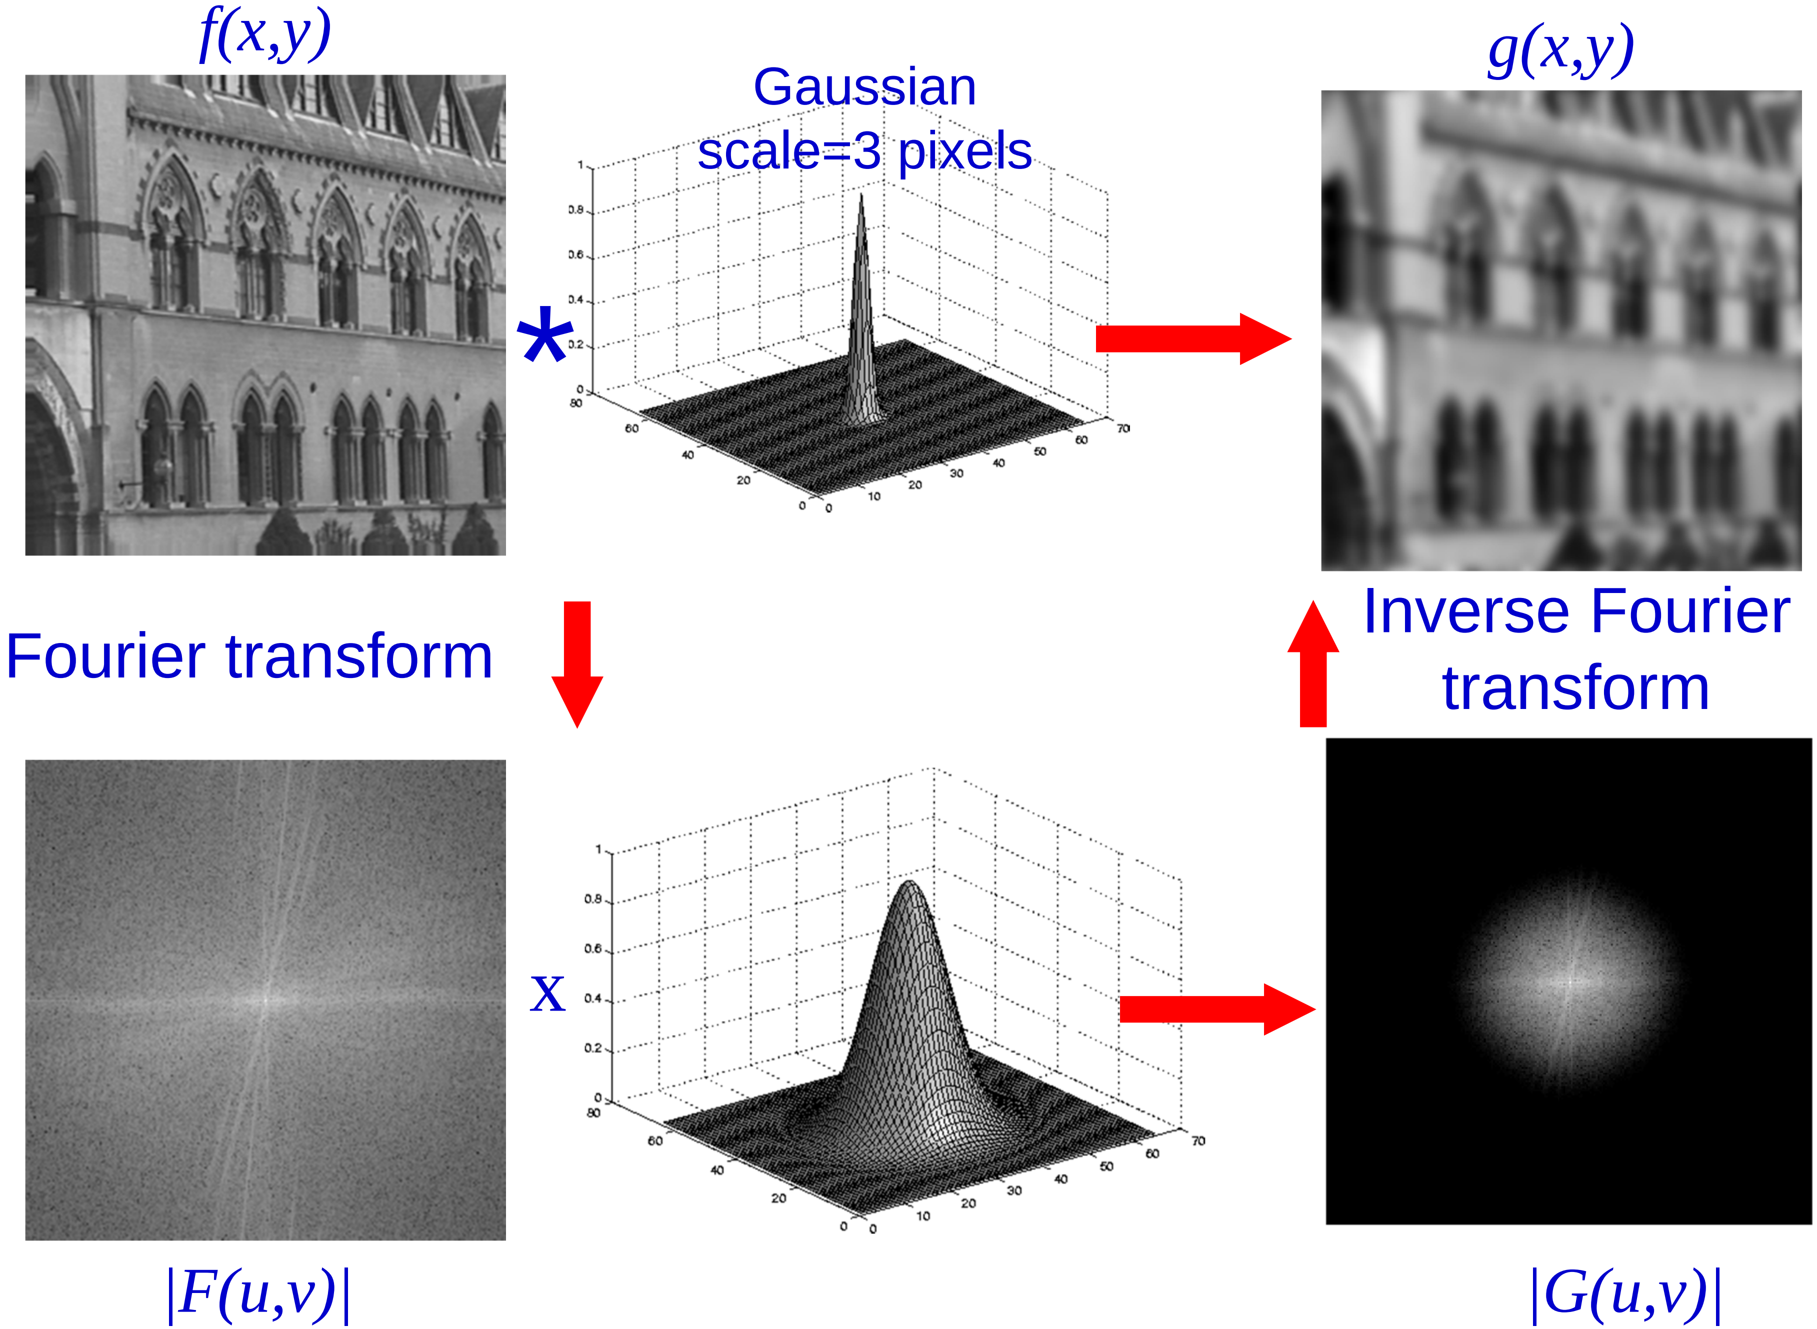

f(x,y)
*
Gaussian
scale=3 pixels
g(x,y)
Fourier transform
Inverse Fourier
transform
x
|F(u,v)|
|G(u,v)|

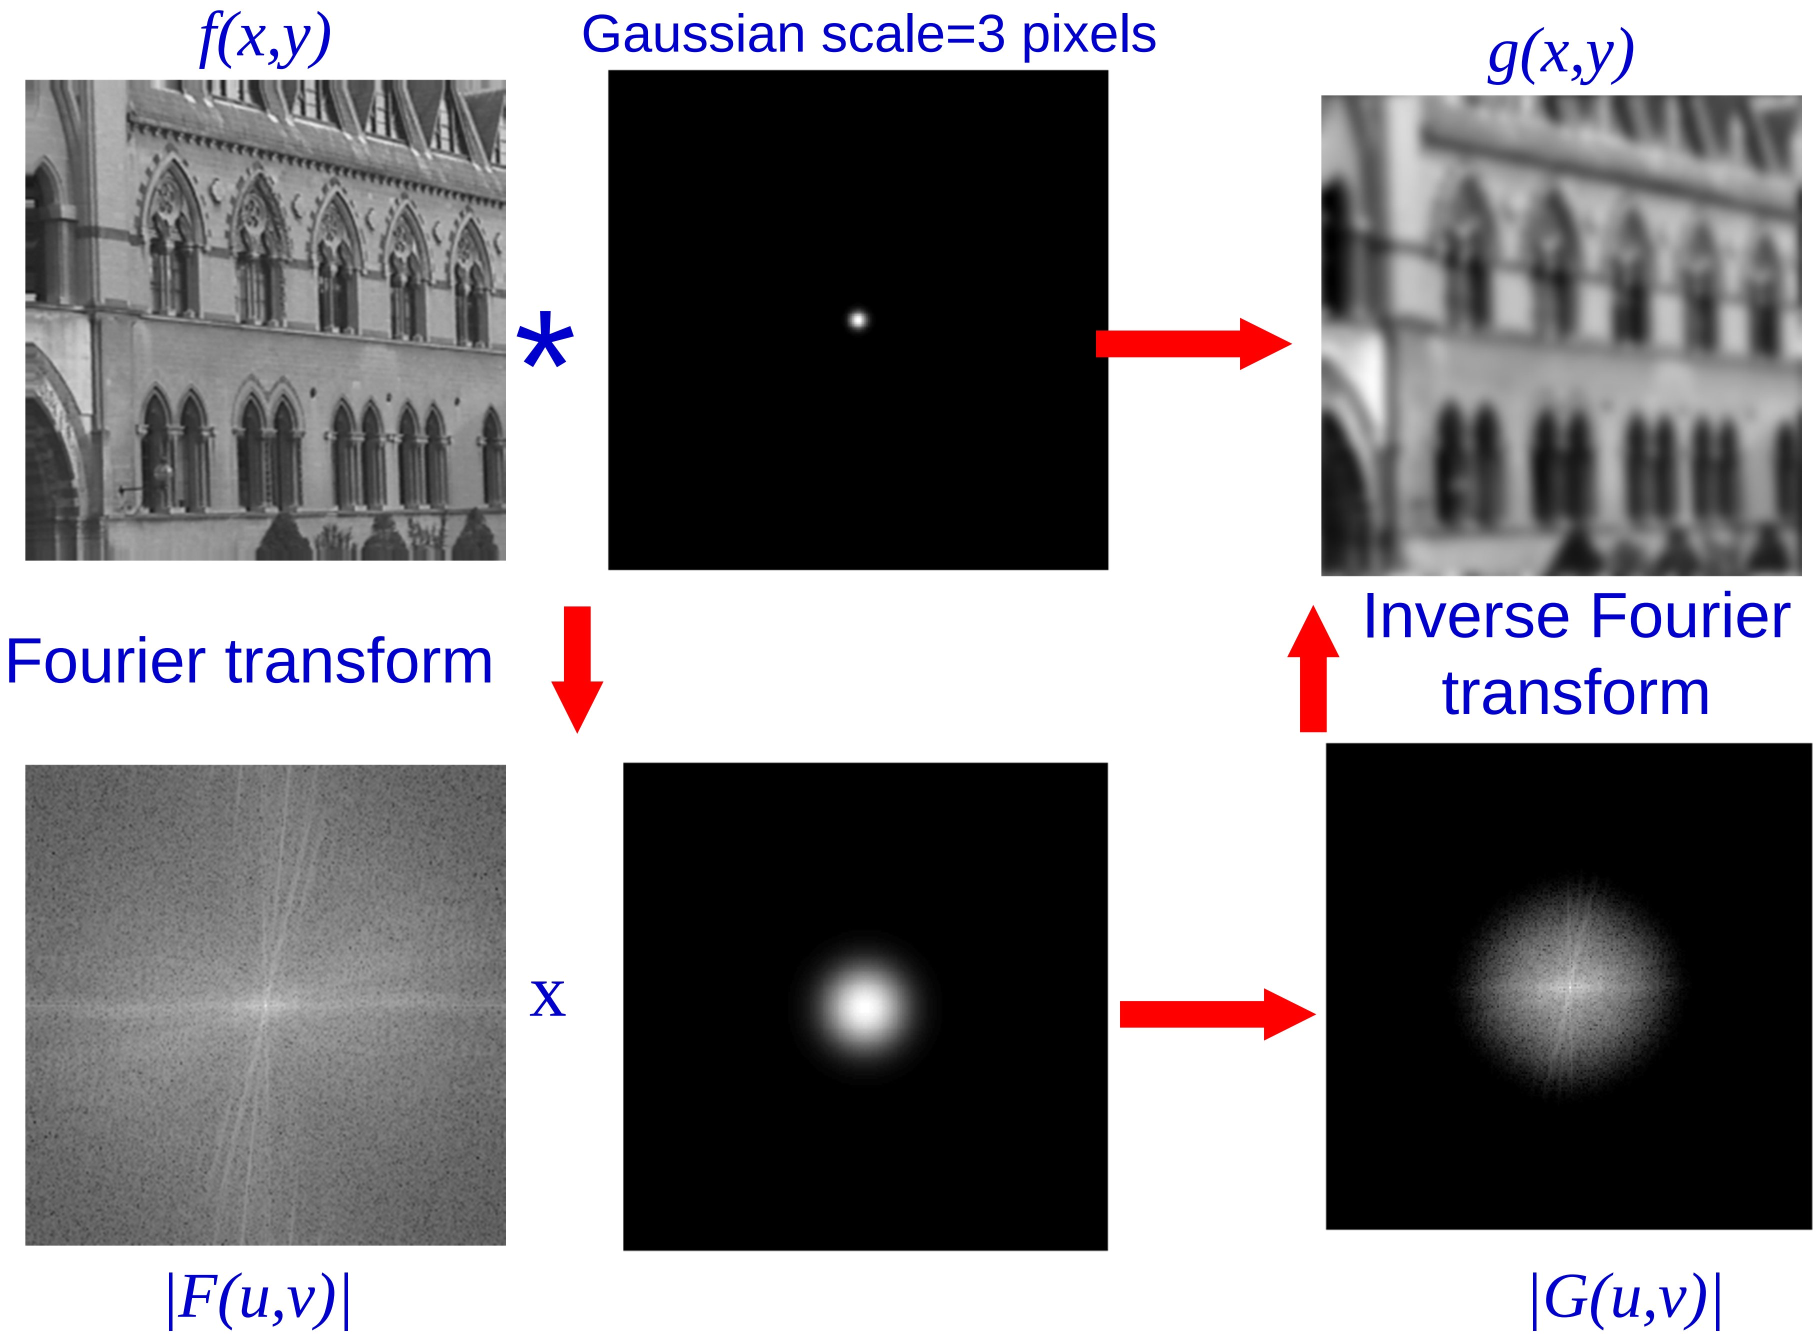

f(x,y)
Gaussian scale=3 pixels
g(x,y)
*
Fourier transform
Inverse Fourier transform
x
|F(u,v)|
|G(u,v)|

There are two equivalent ways of carrying out linear spatial filtering operations:

1. Spatial domain: convolution with a spatial operator
2. Frequency domain: multiply FT of signal and filter, and compute inverse FT of product

Why choose one over the other ?

- The filter may be simpler to specify or compute in one of the domains
- Computational cost

Exercise

What is the FT of …

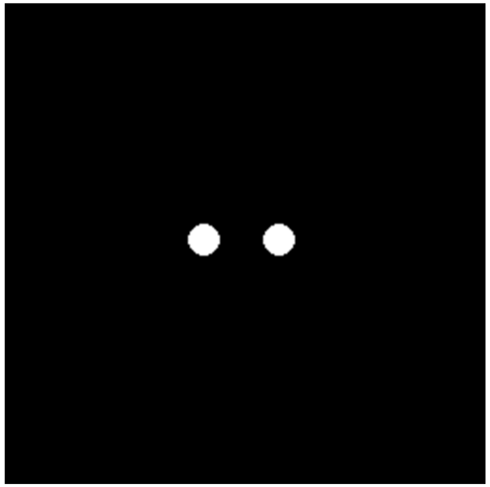

?

2 small disks

The sampling theorem

Discrete Images - Sampling

In 1D model the image as a set of point samples obtained my multiplying $f(x)$ by the comb function

$$comb(x) = \sum_{n=-\infty}^{\infty} \delta(x - nX)$$

an infinite set of delta functions spaced by X.

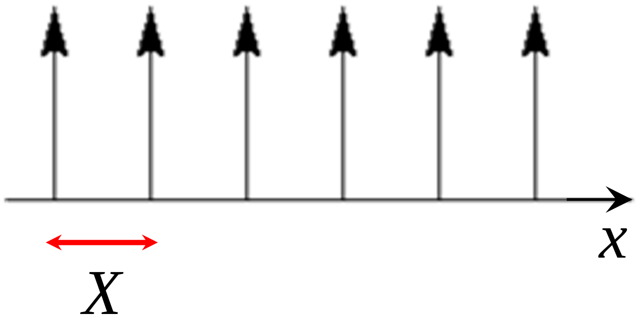

X

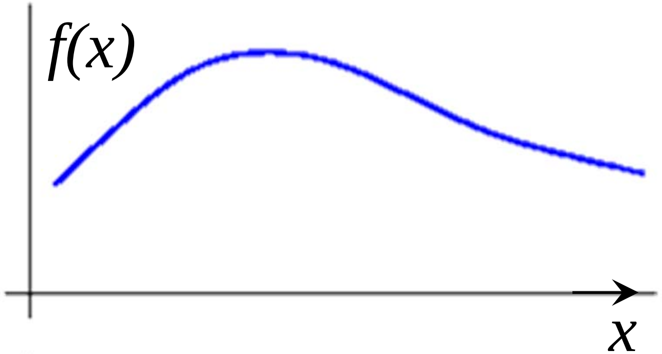

$$f_s(x) = \sum_{n=-\infty}^{\infty} \delta(x - nX) f(x)$$

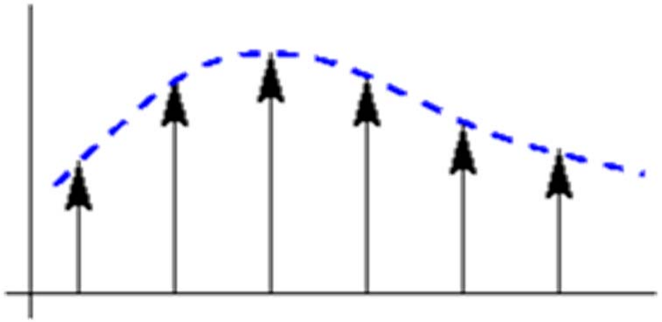

In 2D the equivalent of a comb is a bed-of-nails function

$$\sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \delta(x - nX)\delta(y - mY)$$

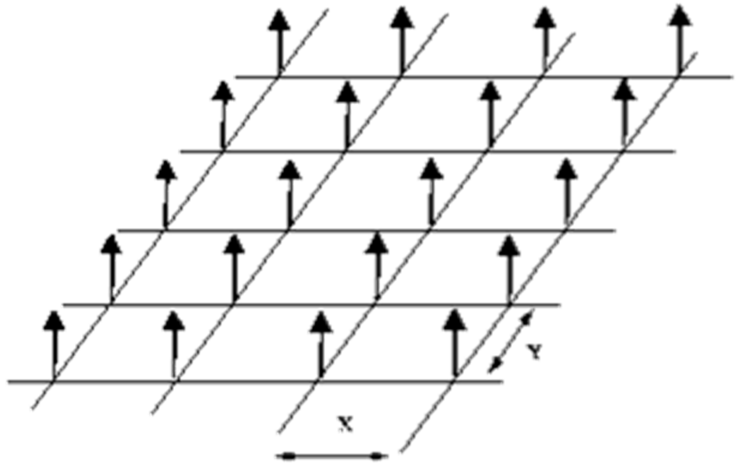

Fourier transform pairs

$$\sum_{n=-\infty}^{\infty} \delta(x - nX) \leftrightarrow \frac{1}{X} \sum_{n=-\infty}^{\infty} \delta(u - n/X)$$

$$\sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \delta(x - nX)\delta(y - mY) \leftrightarrow \frac{1}{XY} \sum_{n=-\infty}^{\infty} \delta(u - n/X) \sum_{m=-\infty}^{\infty} \delta(v - n/Y)$$

Sampling Theorem in 1D

spatial domain

frequency domain

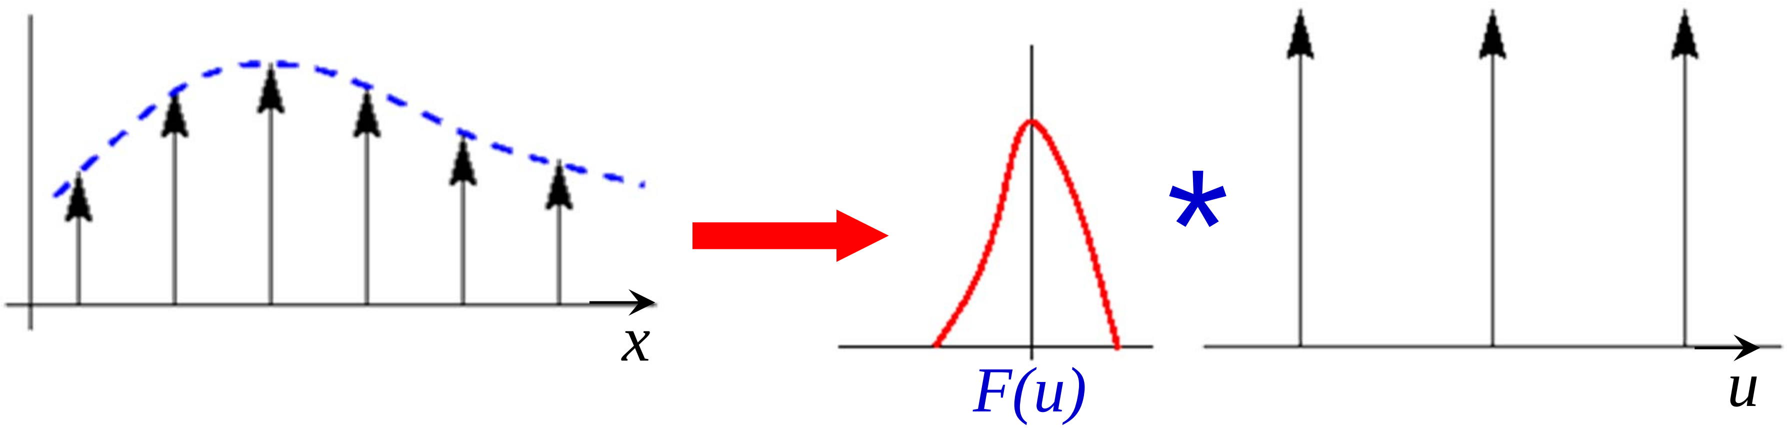

$$f_s(x) = \sum_{n=-\infty}^{\infty} \delta(x - nX) f(x)$$
$$= \sum_{n=-\infty}^{\infty} f(nX)\delta(x - nX)$$

$$F_s(u) = \frac{1}{X}\sum_{n=-\infty}^{\infty} \delta(u - n/X) * F(u) = \frac{1}{X}\sum_{n=-\infty}^{\infty} F(u - n/X)$$

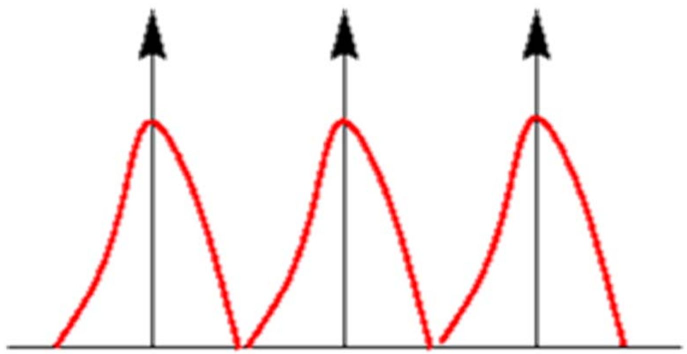

replicated copies of $F(u)$

Apply a box filter

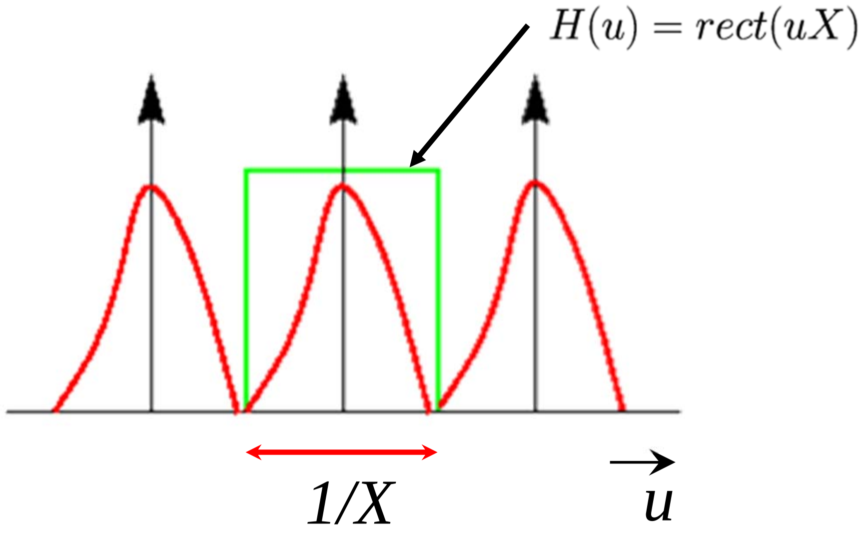

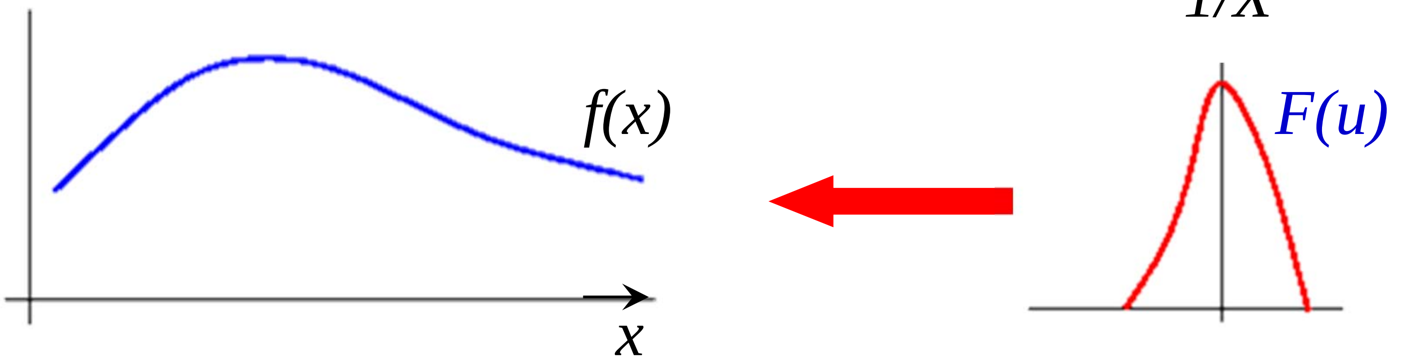

$$F(u) = F_s(u)H(u)$$

$$f(x) = \sum_{n=-\infty}^{\infty} f(nX)\delta(x - nX) * \text{sinc}\frac{\pi x}{X}$$

$$= \sum_{n=-\infty}^{\infty} f(nX)\text{sinc}\frac{\pi}{X}(x - nX)$$

The original continuous function *f(x)* is completely recovered from the samples provided the sampling frequency (*1/X*) exceeds twice the greatest frequency of the band-limited signal. (Nyquist sampling limit)

The Sampling Theorem and Aliasing

if sampling frequency is reduced …

spatial domain

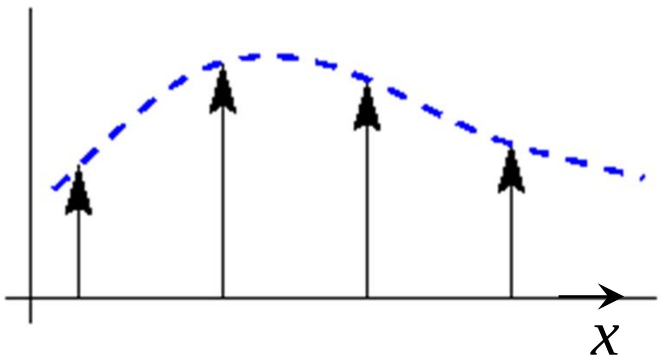

frequency domain

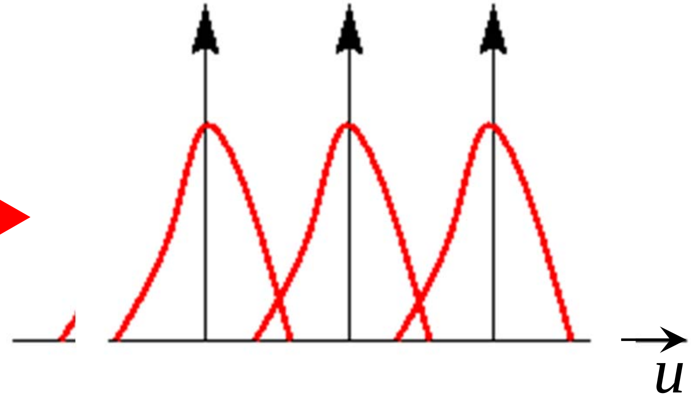

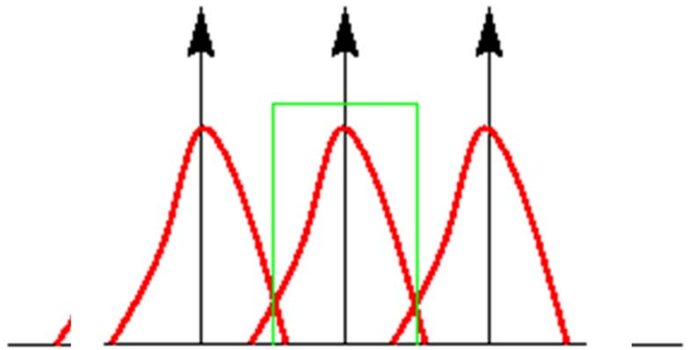

Frequencies above the Nyquist limit are ‘folded back’ corrupting the signal in the acceptable range.

The information in these frequencies is not correctly reconstructed.

Sampling Theorem in 2D

frequency domain

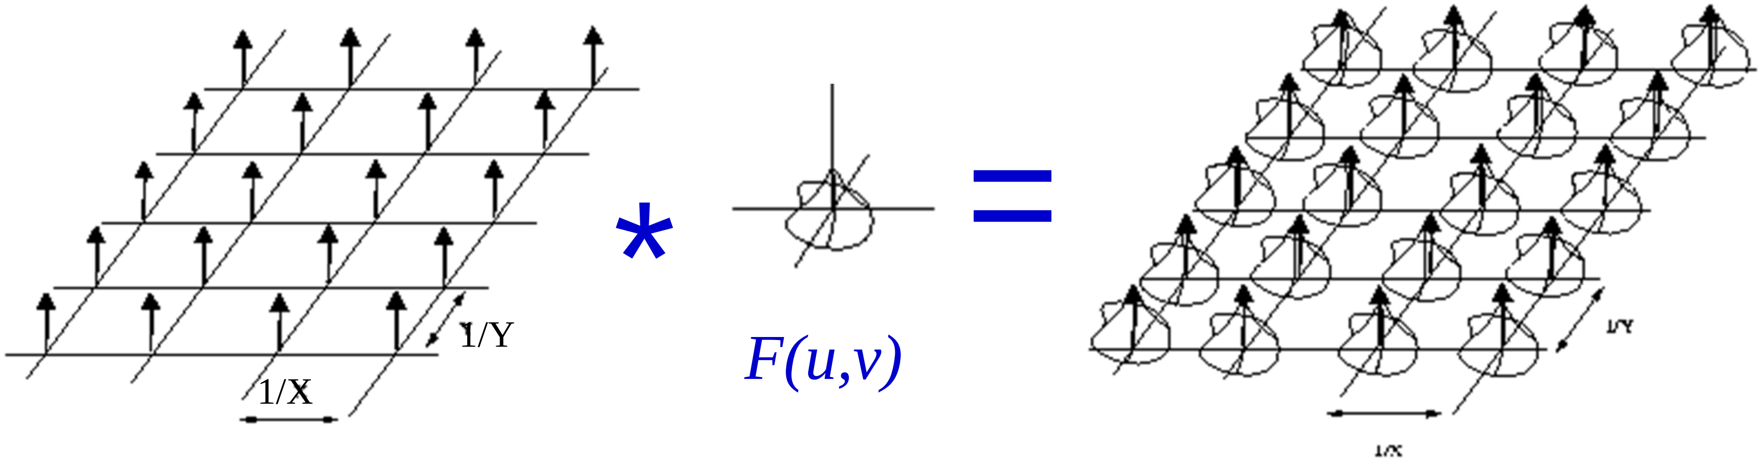

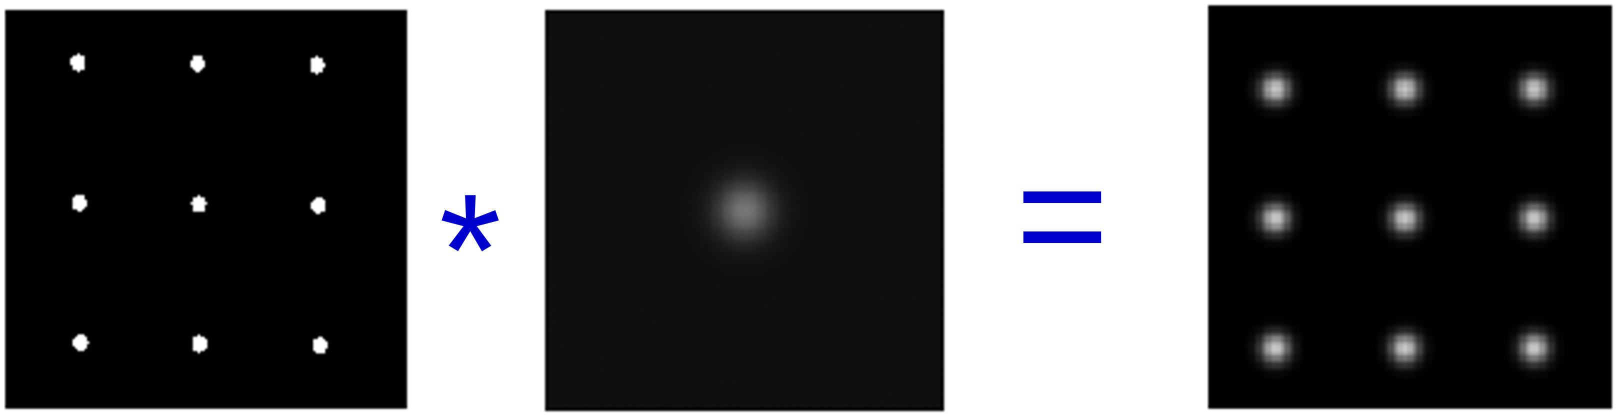

$H(u,v) = rect(uX)rect(vY)$

$$f(x,y) = \sum_{n=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} f(nX, mY)\,\mathrm{sinc}\frac{\pi}{X}(x - nX)\,\mathrm{sinc}\frac{\pi}{Y}(y - nY)$$

The sampling theorem in 2D

If the Fourier transform of a function $f(x,y)$ is zero for all frequencies beyond u_b and v_b,i.e. if the Fourier transform is *band-limited*, then the continuous function $f(x,y)$ can be completely reconstructed from its samples as long as the sampling distances w and h along the x and y directions are such that $w \leq \frac{1}{2u_b}$ and $h \leq \frac{1}{2v_b}$

Aliasing

Aliasing : 1D example

If the signal has frequencies above the Nyquist limit …

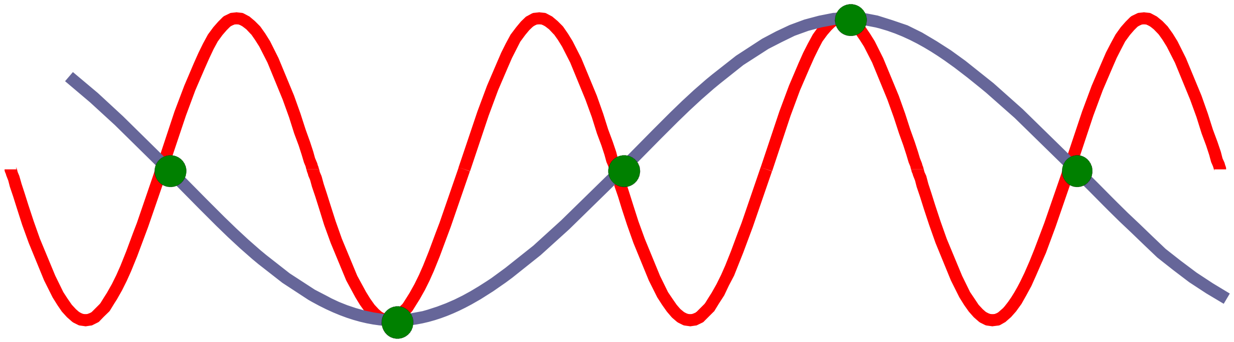

Insufficient samples to distinguish the high and low frequency

aliasing: signals “travelling in disguise” as other frequencies

Aliasing 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):

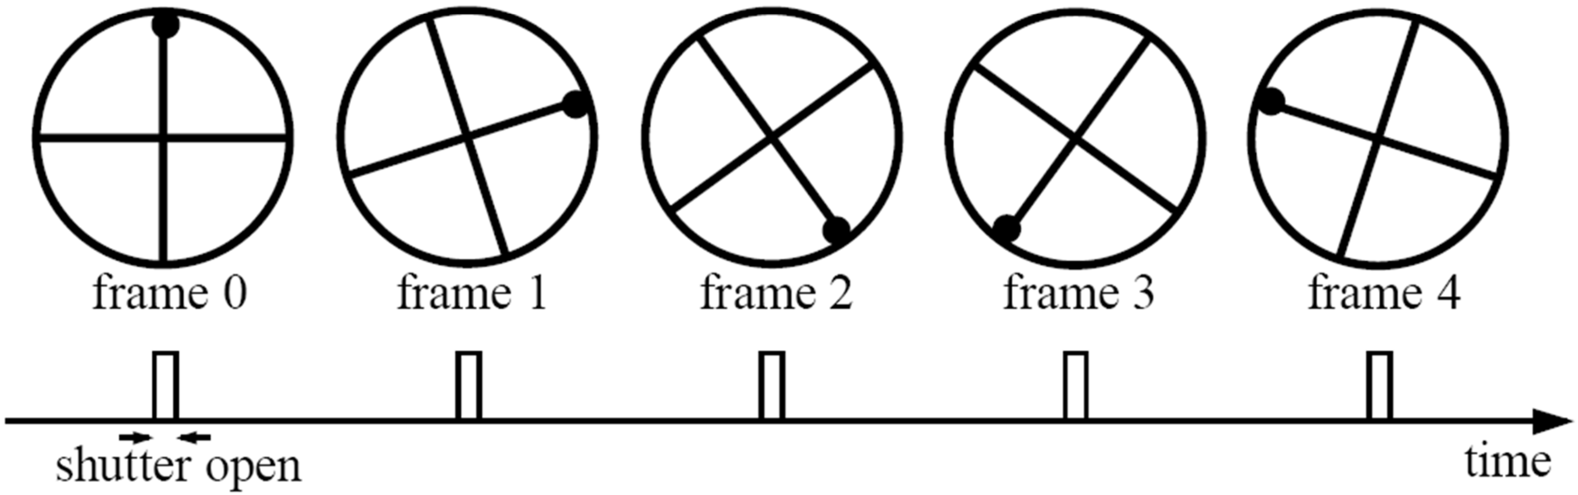

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Steve Seitz

Aliasing in 2D – under sampling example

original

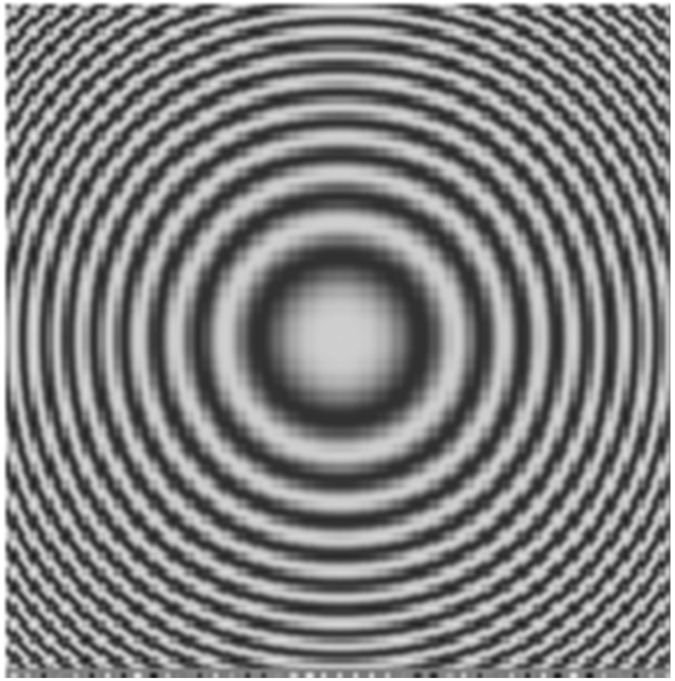

signal has frequencies
above Nyquist limit

reconstruction

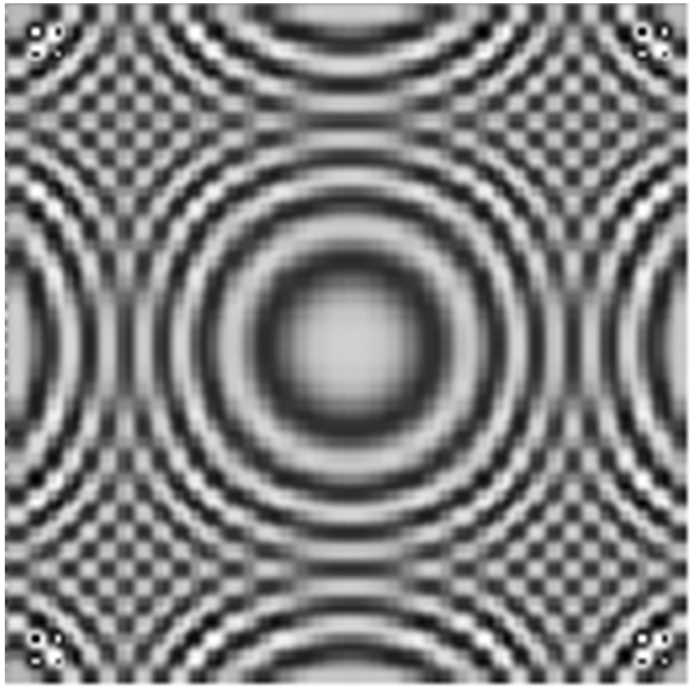

Aliasing in images

What’s happening?

Input signal:

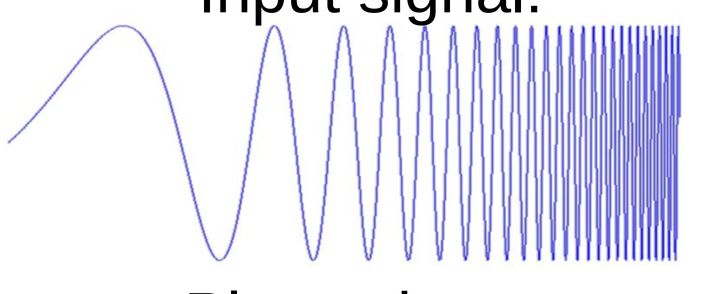

Plot as image:

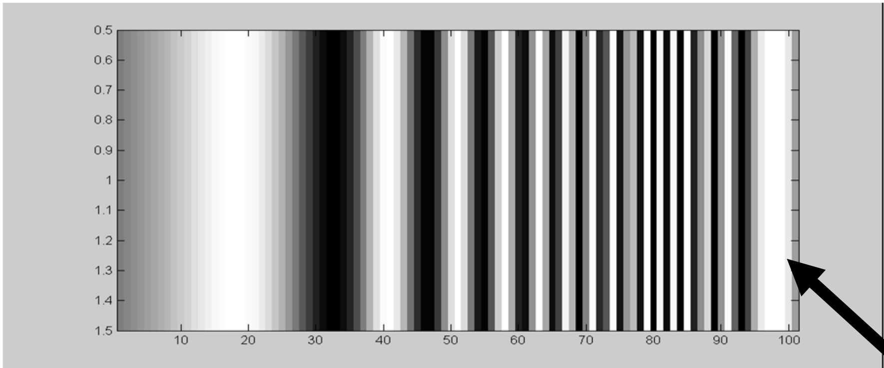

x = 0:.05:5; imagesc(sin((2.^x).*x))

Aliasing
Not enough samples

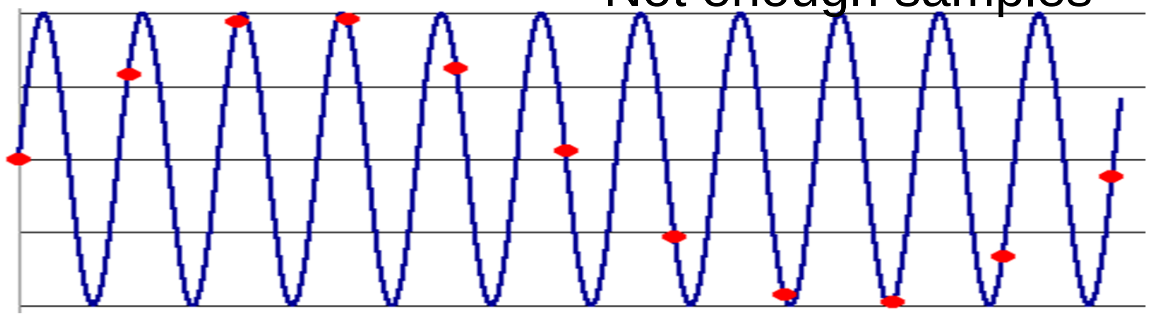

Anti-Aliasing

- Increase sampling frequency
  - e.g. in graphics rendering cast 4 rays per pixel

- Reduce maximum frequency to below Nyquist limit
  - e.g. low pass filter before sampling

Example

down sample by factor of 4

4 x zoom

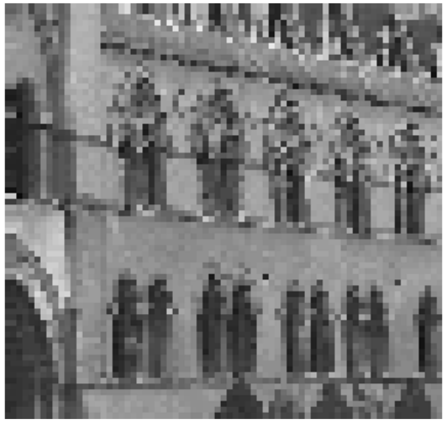

convolve with Gaussian

*

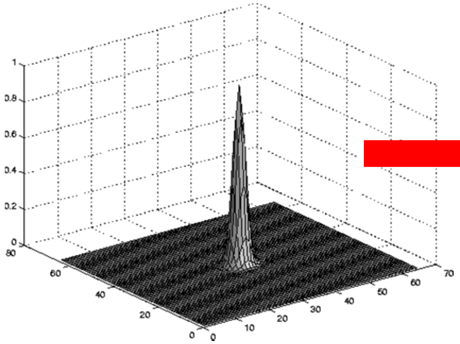

down sample by factor of 4

Hybrid Images

Frequency Domain and Perception

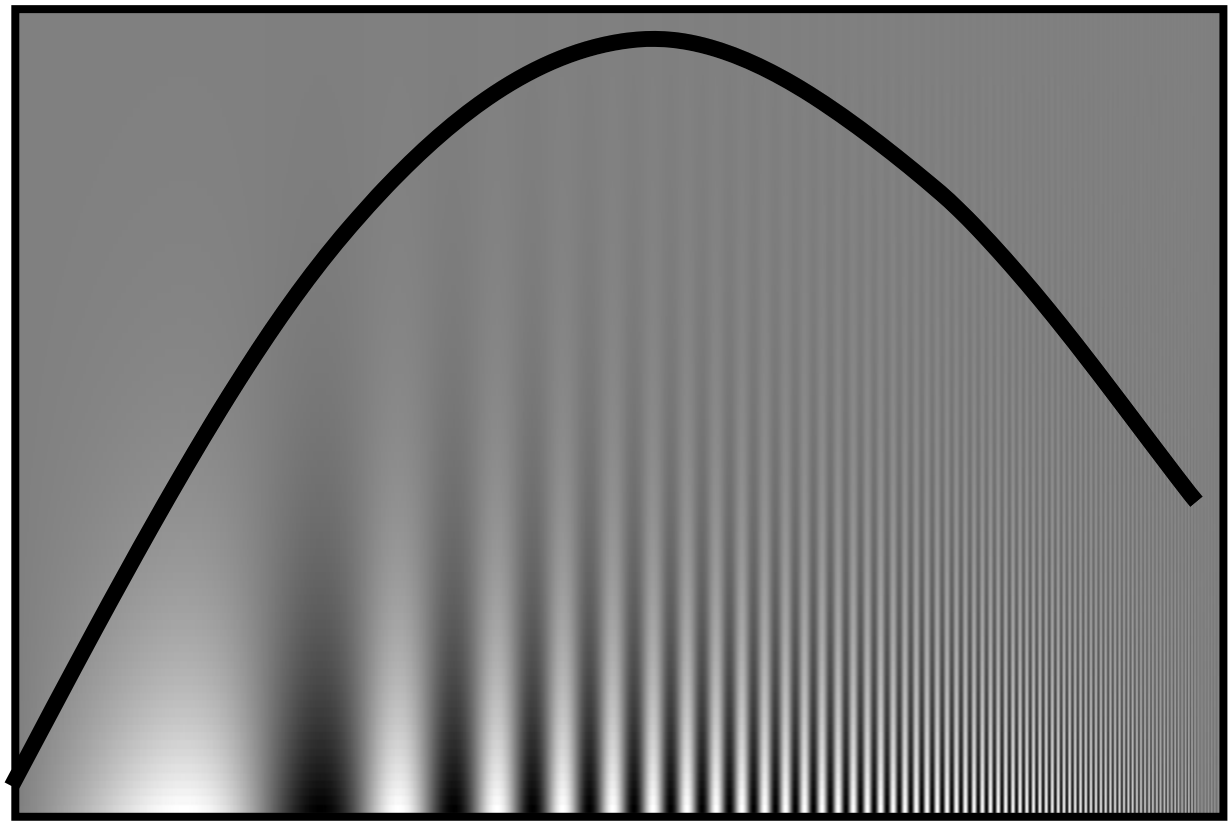

Campbell-Robson contrast sensitivity curve

slide: A. Efros

Perception of hybrid images

SIGGRAPH2006

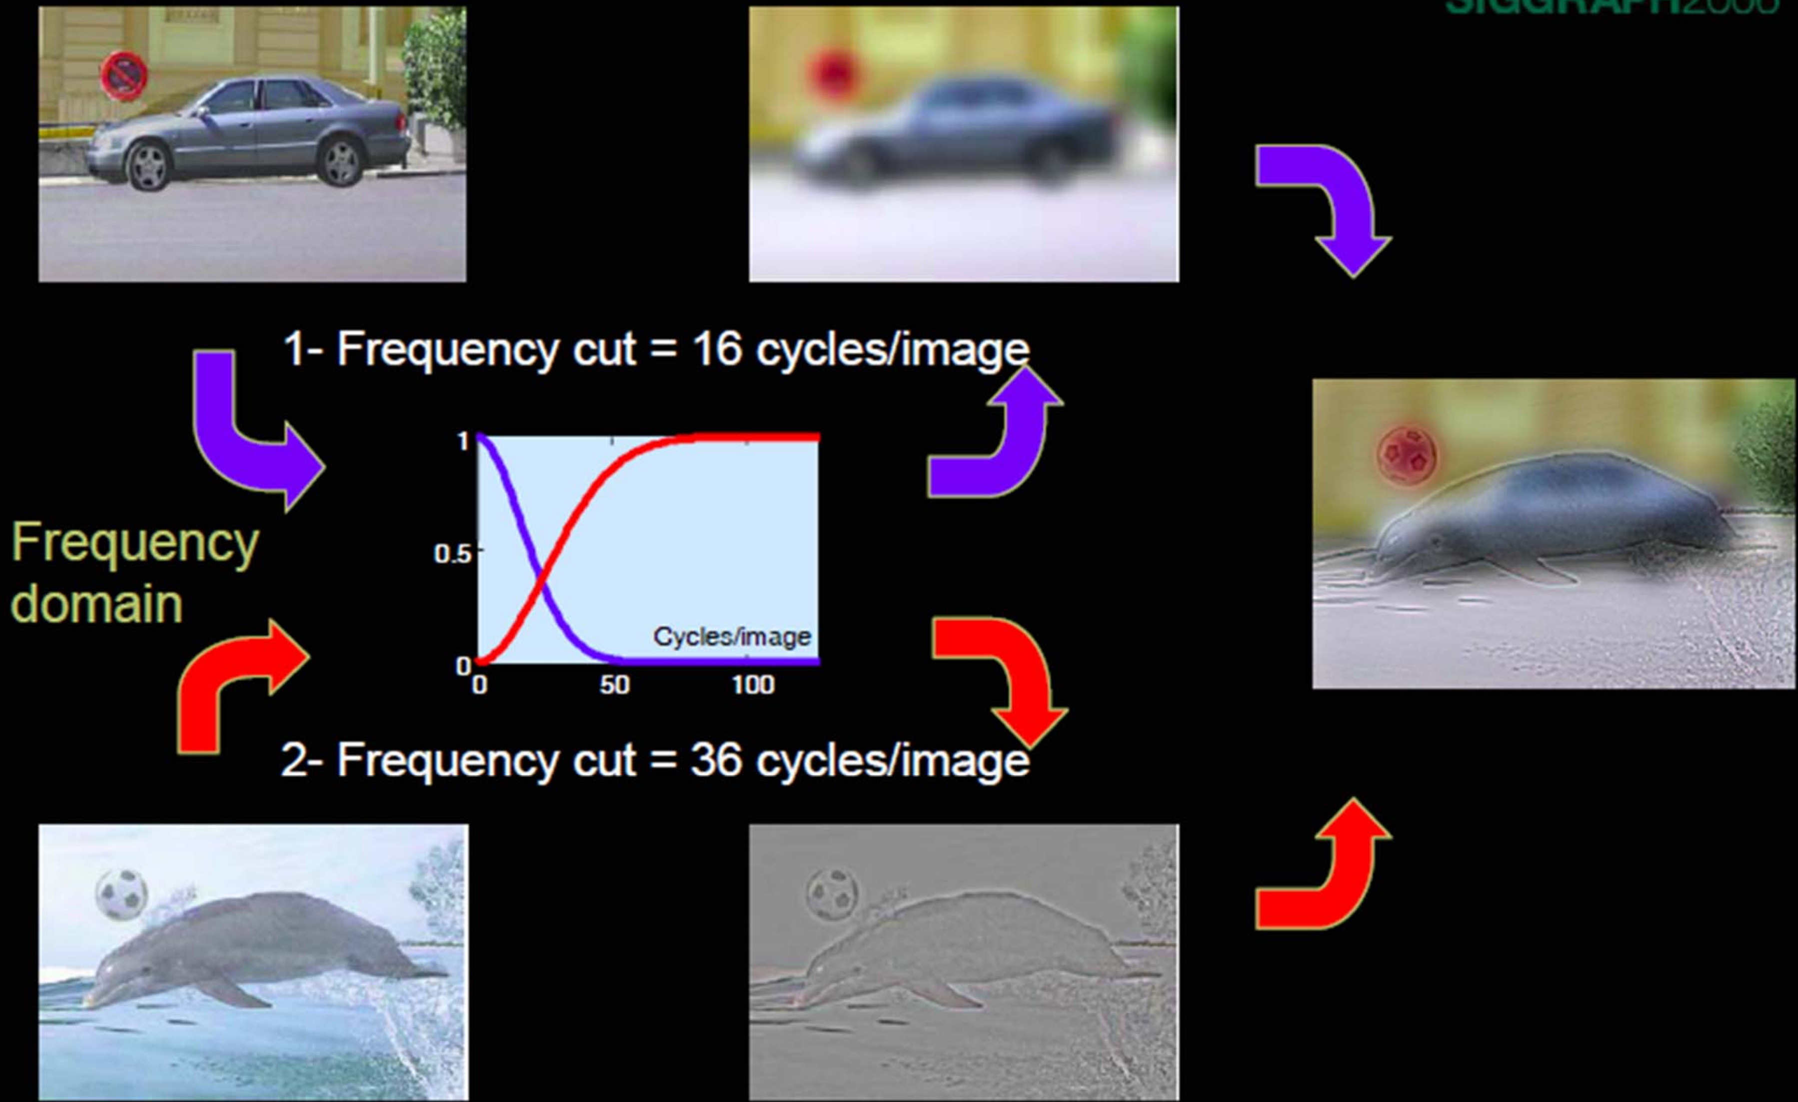

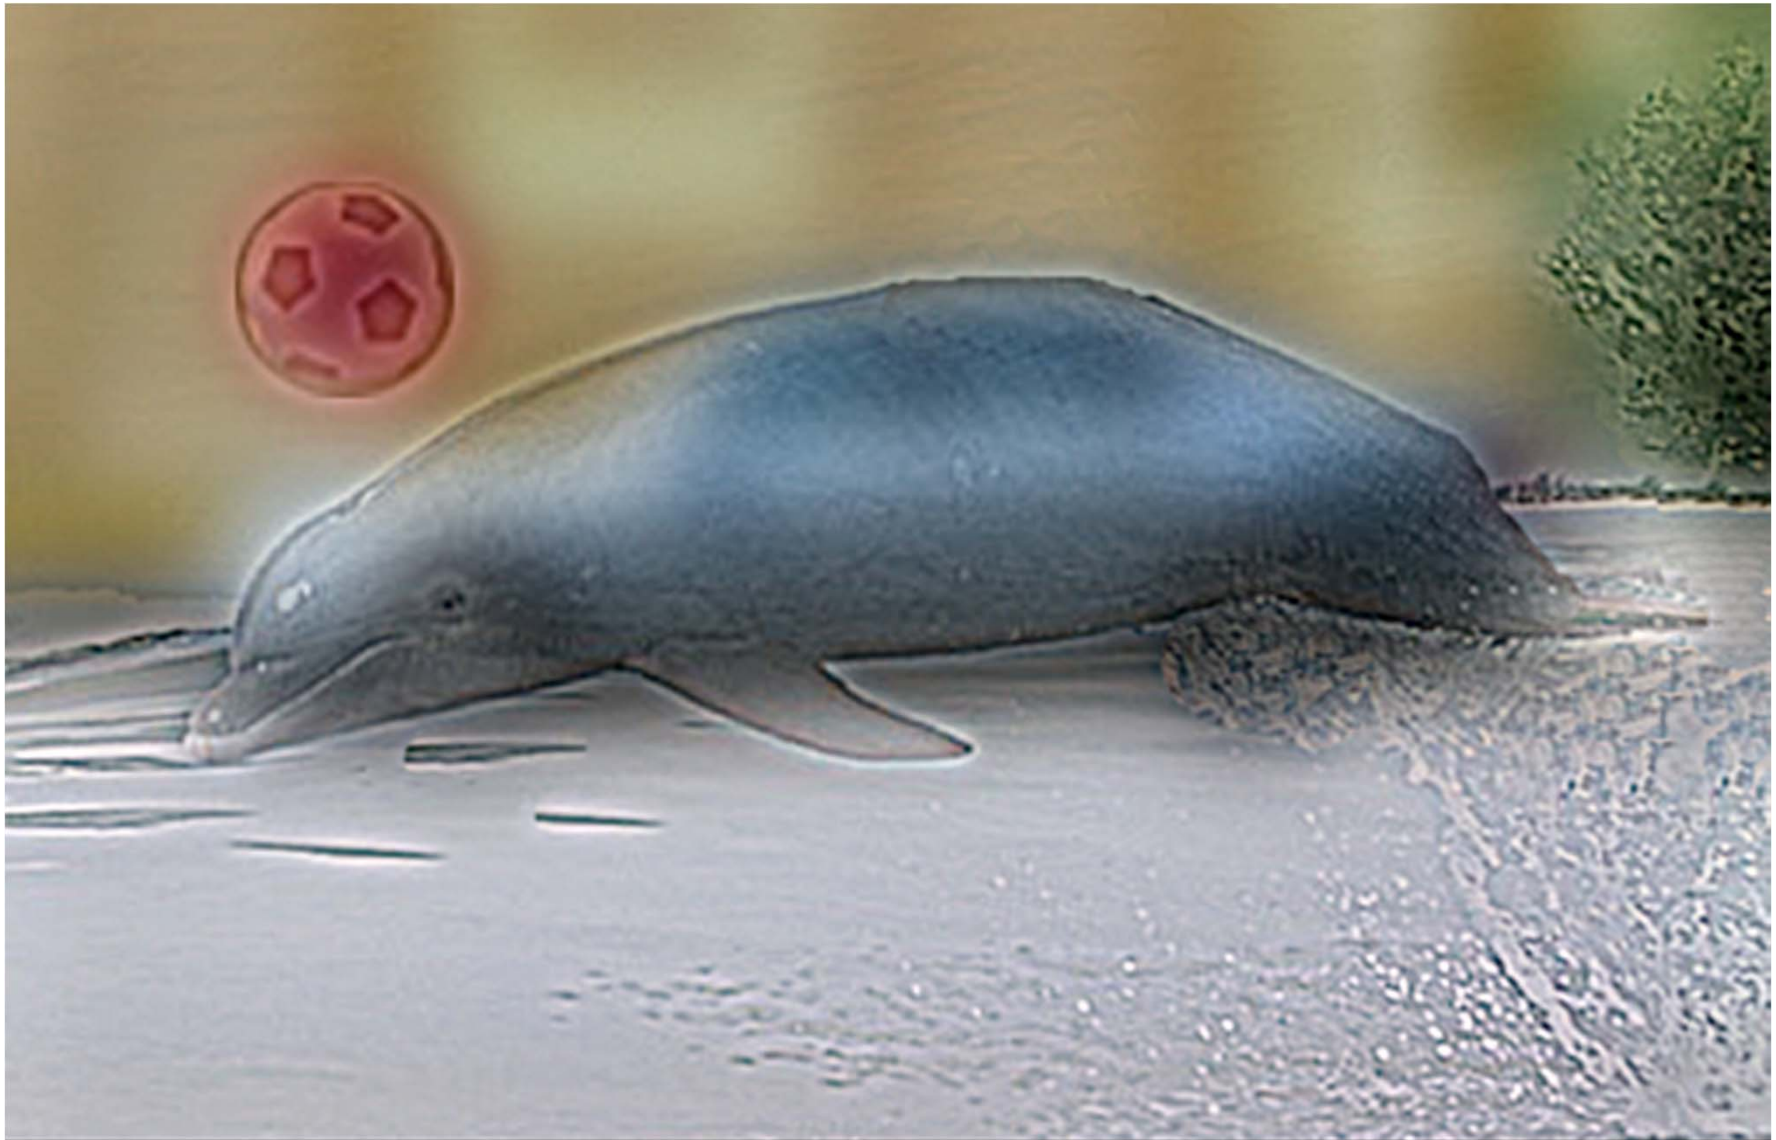

Changing expression

Sad ⟷ Surprised

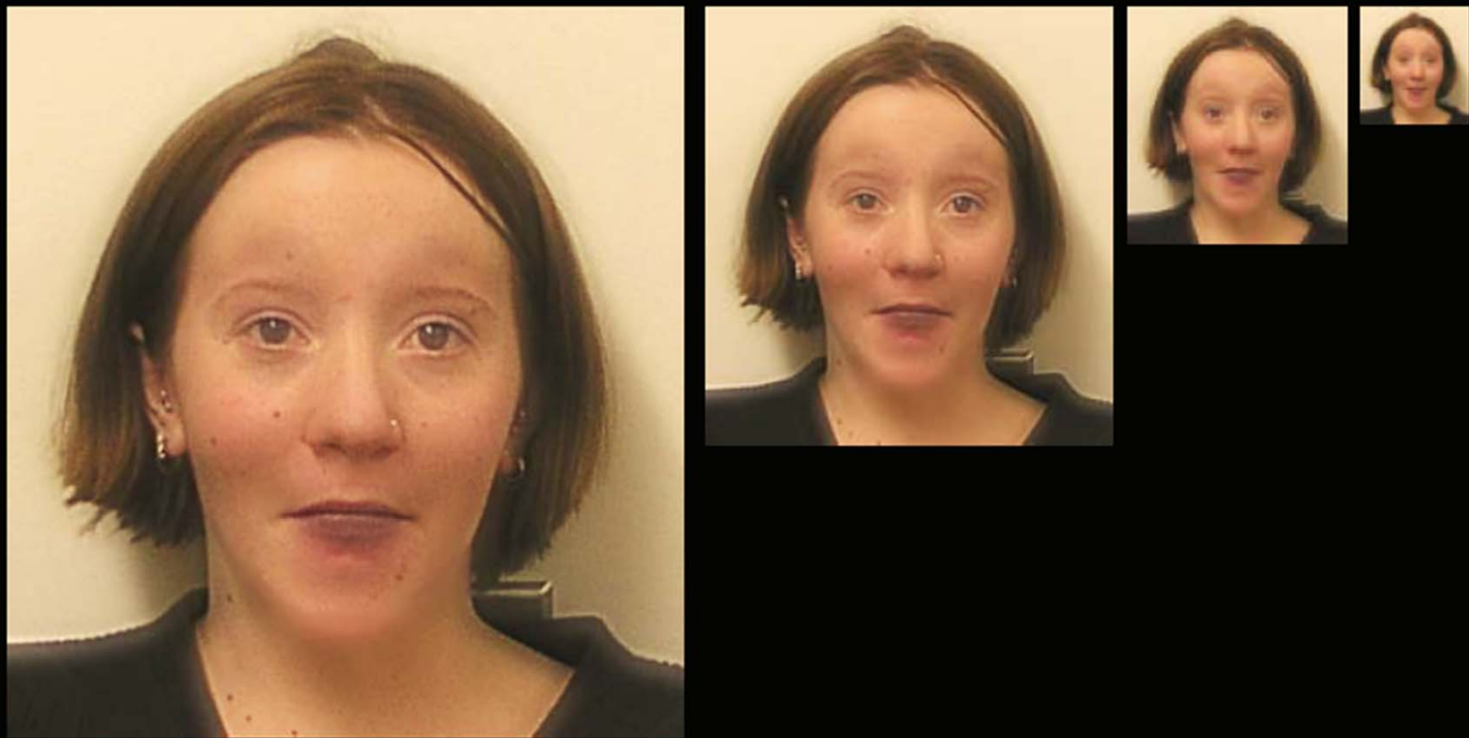

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006