

2. Fourier analysis and sampling theory

Reading

Required:

- ♦ Watt, Section 14.1

Recommended:

- ♦ Ron Bracewell, The Fourier Transform and Its Applications, McGraw-Hill.
- ♦ Don P. Mitchell and Arun N. Netravali, "Reconstruction Filters in Computer Computer Graphics," Computer Graphics, (Proceedings of SIGGRAPH 88). 22 (4), pp. 221-228, 1988.

What is an image?

We can think of an **image** as a function, f , from \mathbb{R}^2 to \mathbb{R} :

- $f(x, y)$ gives the intensity of a channel at position (x, y)
- Realistically, we expect the image only to be defined over a rectangle, with a finite range:
 - $f: [a,b] \times [c,d] \rightarrow [0,1]$



A color image is just three functions pasted together. We can write this as a “vector-valued” function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

We'll focus in grayscale (scalar-valued) images for now.

3

Images as functions



4


Digital images

In computer graphics, we usually create or operate on **digital (discrete)** images:

- ♦ **Sample** the space on a regular grid
- ♦ **Quantize** each sample (round to nearest integer)

If our samples are Δ apart, we can write this as:

$$f[i, j] = \text{Quantize}\{ f(i \Delta, j \Delta) \}$$



5

Motivation: filtering and resizing

What if we now want to:

- ♦ smooth an image?
- ♦ sharpen an image?
- ♦ enlarge an image?
- ♦ shrink an image?

Before we try these operations, it's helpful to think about images in a more mathematical way...

6

Fourier transforms


We can represent functions as a weighted sum of sines and cosines.

We can think of a function in two complementary ways:

- **Spatially** in the **spatial domain**
- **Spectrally** in the **frequency domain**

The **Fourier transform** and its inverse convert between these two domains:

$$\begin{array}{ccc} \text{Spatial domain} & \xrightarrow{F(s) = \int_{-\infty}^{\infty} f(x)e^{-i2\pi sx} dx} & \text{Frequency domain} \\ & \xleftarrow{f(x) = \int_{-\infty}^{\infty} F(s)e^{i2\pi sx} ds} & \end{array}$$



Fourier transforms (cont'd)

$$\begin{array}{ccc} \text{Spatial domain} & \xrightarrow{F(s) = \int_{-\infty}^{\infty} f(x)e^{-i2\pi sx} dx} & \text{Frequency domain} \\ & \xleftarrow{f(x) = \int_{-\infty}^{\infty} F(s)e^{i2\pi sx} ds} & \end{array}$$

Where do the sines and cosines come in?

$$a_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos(kx) dx$$

$$b_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(kx) dx$$

$$F(s) = A(s) + iB(s)$$

$$= |F(s)| e^{i2\pi\theta(s)}$$

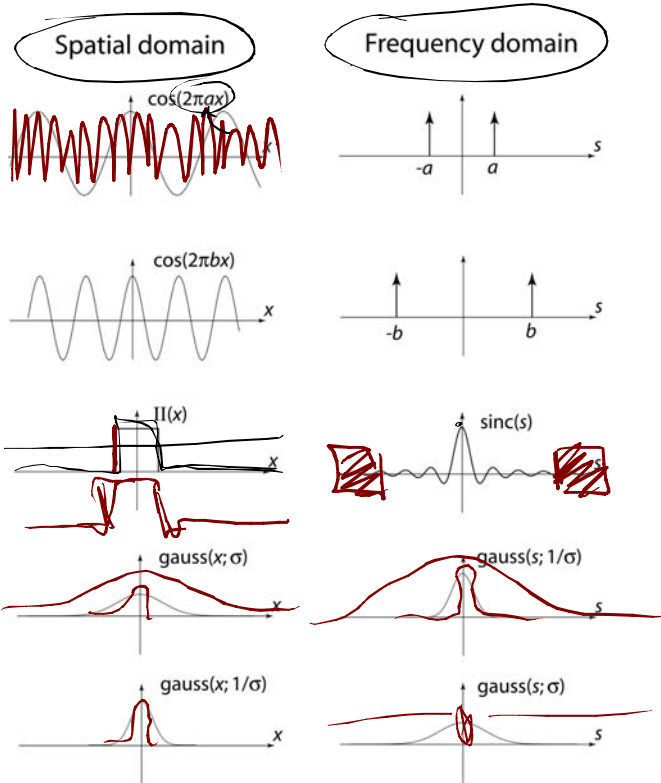
If $f(x)$ is symmetric, i.e., $f(x) = f(-x)$, then $F(s) \neq A(s)$. Why?

$k=4$

a_4

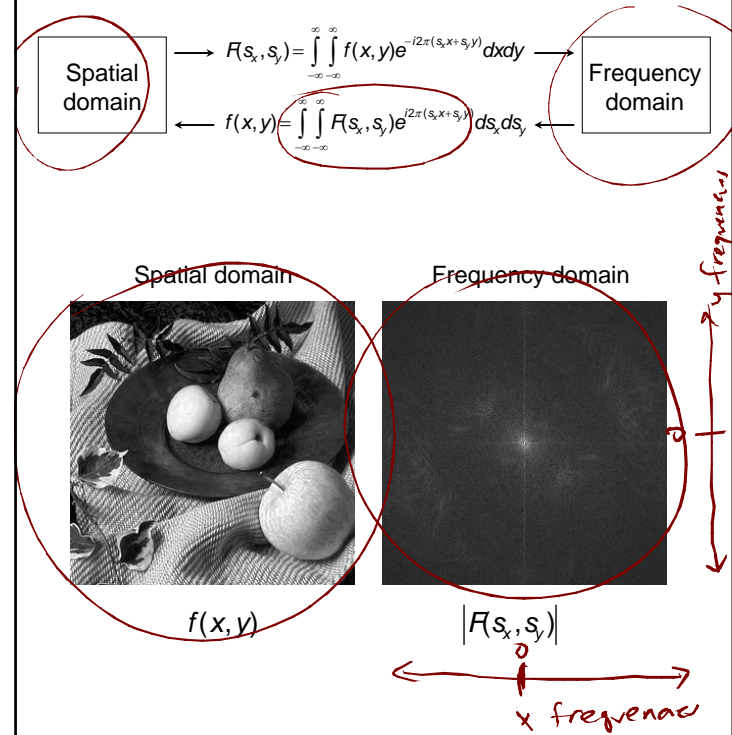
$\cos(4x)$

1D Fourier examples



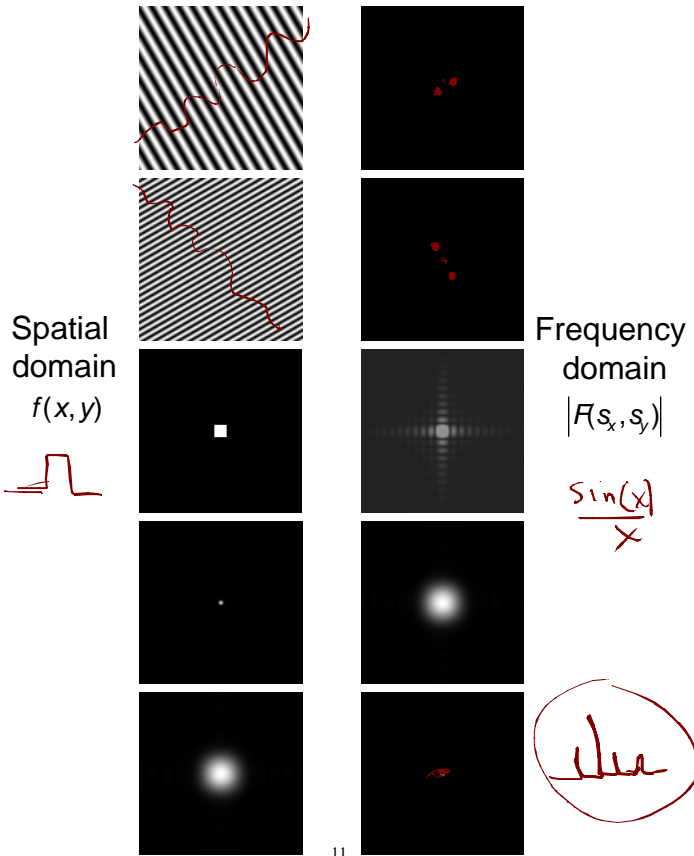
9

2D Fourier transform



10

2D Fourier examples



Convolution

One of the most common methods for filtering a function is called **convolution**.

In 1D, convolution is defined as:

$$g(x) = f(x) * h(x)$$

$$= \int_{-\infty}^{\infty} f(x') h(x - x') dx'$$

$$g(x) = \int_{-\infty}^{\infty} f(x') \tilde{h}(x' - x) dx'$$

where $\tilde{h}(x) \equiv h(-x)$

Convolution properties

Convolution exhibits a number of basic, but important properties.

Commutativity:

$$a(x) * b(x) = b(x) * a(x)$$

Associativity:

$$[a(x) * b(x)] * c(x) = a(x) * [b(x) * c(x)]$$

Linearity:

$$a(x) * [k \cdot b(x)] = k \cdot [a(x) * b(x)]$$

$$a(x) * (b(x) + c(x)) = a(x) * b(x) + a(x) * c(x)$$

13

Convolution in 2D


In two dimensions, convolution becomes:

$$\begin{aligned} g(x, y) &= f(x, y) * h(x, y) \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x', y') h(x - x', y - y') dx' dy' \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x', y') \tilde{h}(x' - x, y' - y) dx' dy' \end{aligned}$$

where $\tilde{h}(x, y) = h(-x, -y)$.



$f(x, y)$



$h(x, y)$




$g(x, y)$

14

Convolution theorems

Convolution theorem: Convolution in the *spatial* domain is equivalent to *multiplication* in the *frequency* domain.



Symmetric theorem: Convolution in the *frequency* domain is equivalent to *multiplication* in the *spatial* domain.


$$f \cdot h \longleftrightarrow F * H$$

15

1D convolution theorem example

Spatial domain

Frequency domain




16

Sifting and shifting

For sampling, the delta function has two important properties.


Sifting:

$$f(x)\delta(x-a) = f(a)\delta(x-a)$$



Shifting:

$$f(x)*\delta(x-a) = f(x-a)$$




19

The shah/comb function

A string of delta functions is the key to sampling. The resulting function is called the **shah** or **comb** function:

$$\text{III}(x) = \sum_{n=-\infty}^{\infty} \delta(x - nT)$$


which looks like:



Amazingly, the Fourier transform of the shah function takes the same form:

$$\text{III}(s) = \sum_{n=-\infty}^{\infty} \delta(s - ns_0)$$

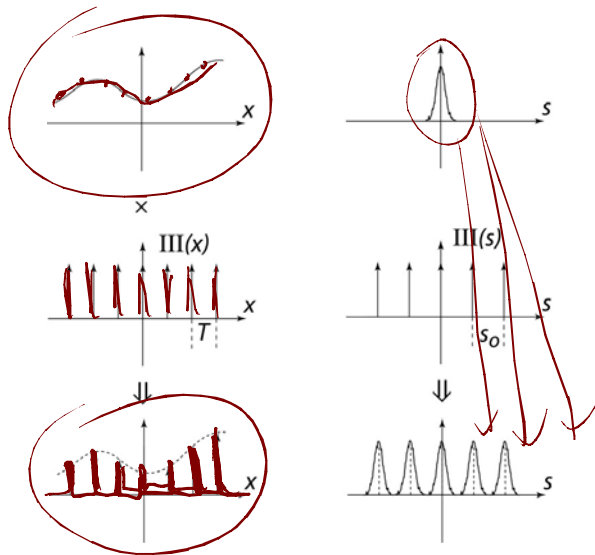
where $s_0 = 1/T$.



20

Sampling

Now, we can talk about sampling.

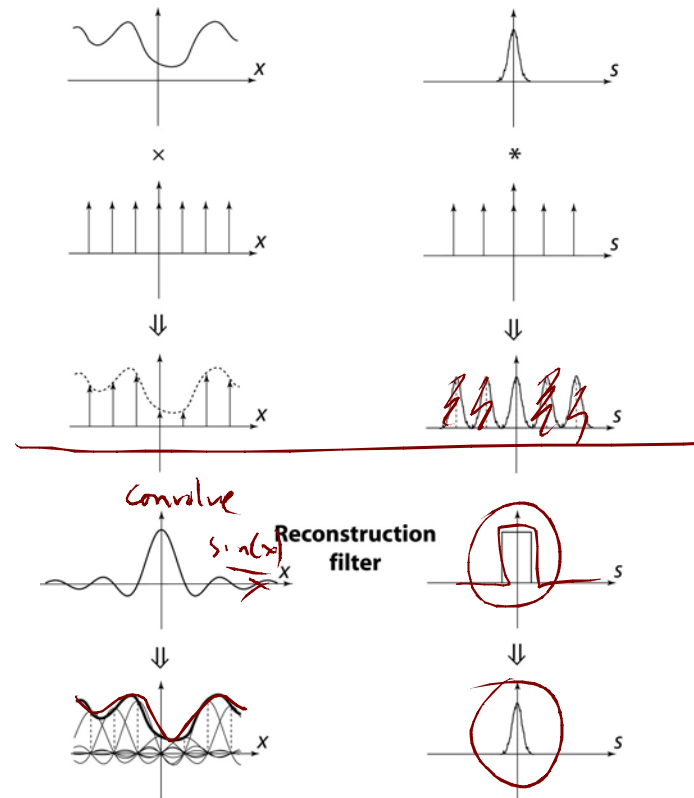


The Fourier spectrum gets *replicated* by spatial sampling!

How do we recover the signal?

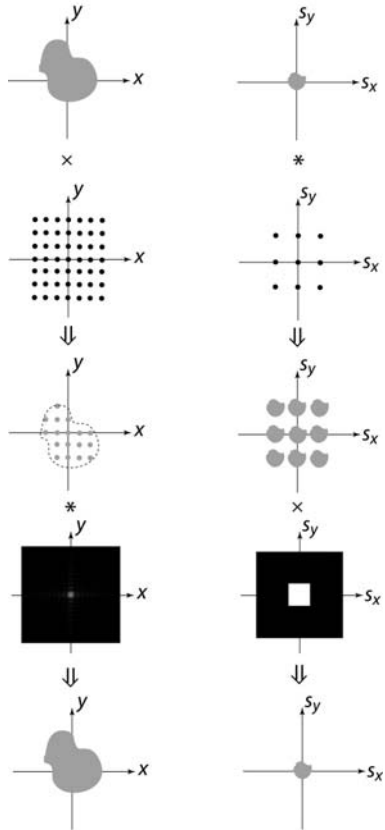
21

Sampling and reconstruction



22

Sampling and reconstruction in 2D



23

Sampling theorem

This result is known as the **Sampling Theorem** and is due to Claude Shannon who first discovered it in 1949:

A signal can be reconstructed from its samples without loss of information, if the original signal has no frequencies above $\frac{1}{2}$ the sampling frequency.

For a given **bandlimited** function, the minimum rate at which it must be sampled is the **Nyquist frequency**.


24

Reconstruction filters

The sinc filter, while “ideal”, has two drawbacks:

- ♦ It has large support (slow to compute)
- ♦ It introduces ringing in practice

We can choose from many other filters...



25


Cubic filters

Mitchell and Netravali (1988) experimented with cubic filters, reducing them all to the following form:

$$r(x) = \frac{1}{6} \begin{cases} (12 - 9B - 6C)|x|^3 + (-18 + 12B + 6C)|x|^2 + (6 - 2B) & |x| < 1 \\ ((-B - 6C)|x|^3 + (6B + 30C)|x|^2 + (-12B - 48C)|x| + (8B + 24C)) & 1 \leq |x| < 2 \\ 0 & \text{otherwise} \end{cases}$$

The choice of B or C trades off between being too blurry or having too much ringing. B=C=1/3 was their “visually best” choice.

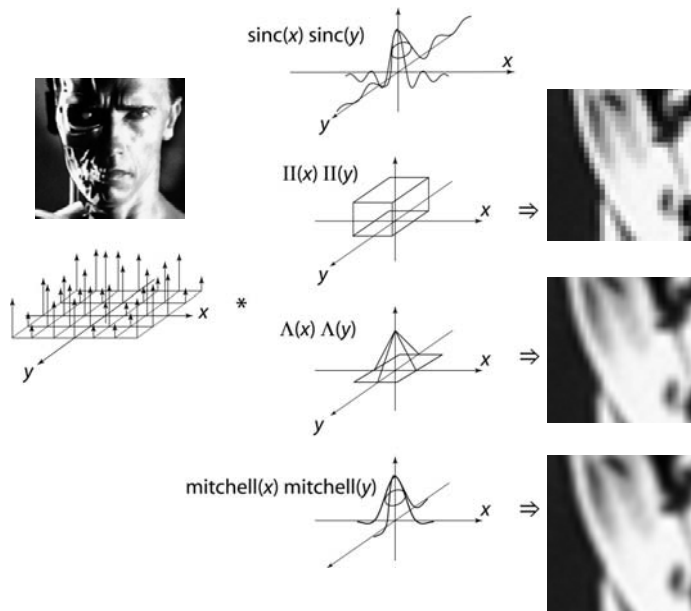
The resulting reconstruction filter is often called the “Mitchell filter.”



26

Reconstruction filters in 2D

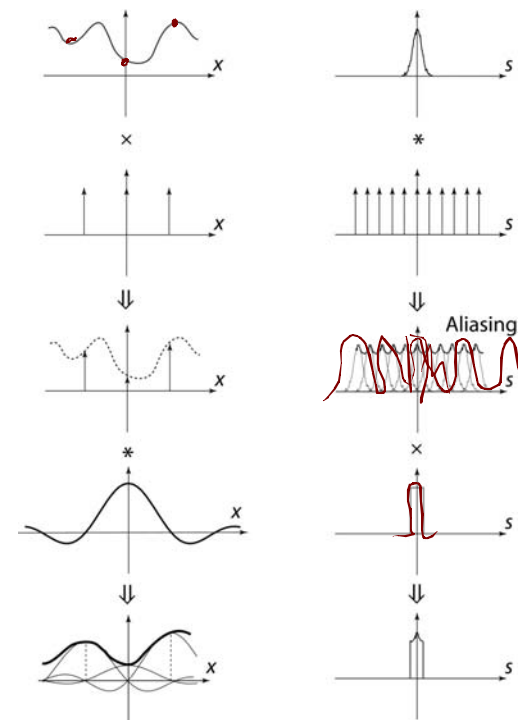
We can also perform reconstruction in 2D...



27

Aliasing


What if we go below the Nyquist frequency?



28

Anti-aliasing


Anti-aliasing is the process of *removing* the frequencies before they alias.



29

Anti-aliasing by analytic prefiltering

We can fill the "magic" box with analytic pre-filtering of the signal:




Why may this not generally be possible?

30

Filtered downsampling

Alternatively, we can sample the image at a higher rate, and then filter that signal:



We can now sample the signal at a lower rate. The whole process is called **filtered downsampling** or **supersampling and averaging down**.


31

Practical upsampling

When resampling a function (e.g., when resizing an image), you do not need to reconstruct the complete continuous function.

For zooming in on a function, you need only use a reconstruction filter and evaluate as needed for each new sample.

Here's an example using a cubic filter:




32

Practical upsampling

This can also be viewed as:

1. putting the reconstruction filter at the desired location
2. evaluating at the original sample positions
3. taking products with the sample values themselves
4. summing it up



Important: filter should always be normalized!


33

Practical downsampling

Downsampling is similar, but filter has larger support and smaller amplitude.

Operationally:

1. Choose filter in downsampled space.
2. Compute the downsampling rate, d , ratio of new sampling rate to old sampling rate
3. Stretch the filter by $1/d$ and scale it down by d
4. Follow upsampling procedure (previous slides) to compute new values



$$e^{i\theta} = \cos(\theta) + i\sin(\theta)$$

34

2D resampling

We've been looking at **separable** filters:

$$r_{2D}(x, y) = r_{1D}(x)r_{1D}(y)$$

$$\text{magnitude} = \sqrt{a^2 + b^2}$$
$$\text{phase} = \tan^{-1}\left(\frac{b}{a}\right)$$

magnitude
phase

How might you use this fact for efficient resampling in 2D?

