2. Sampling theory

Reading

Required:

Watt, Section 14.1

Recommended:

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.

Images as functions

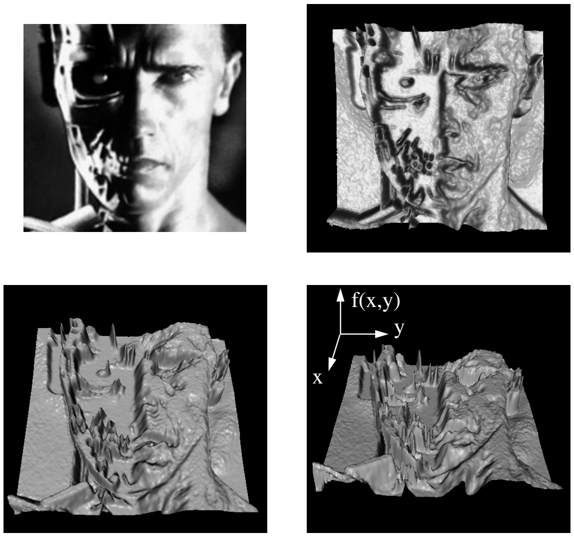

Digital images

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:

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

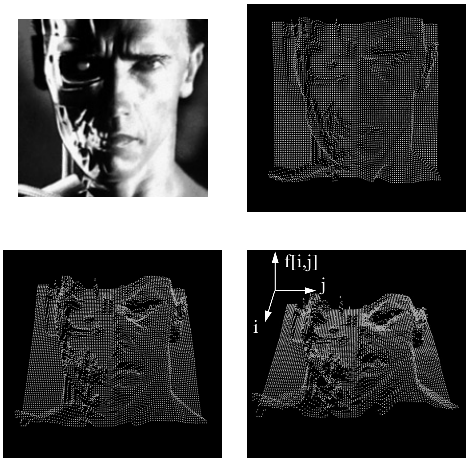

Motivation: filtering and resizing

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…

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:

Spatial domain $\longrightarrow$ Frequency domain:

$$F(s) = \int_{-\infty}^{\infty} f(x) e^{-i2\pi sx}\, dx$$

Frequency domain $\longrightarrow$ Spatial domain:

$$f(x) = \int_{-\infty}^{\infty} F(s) e^{i2\pi sx}\, ds$$

1D Fourier examples

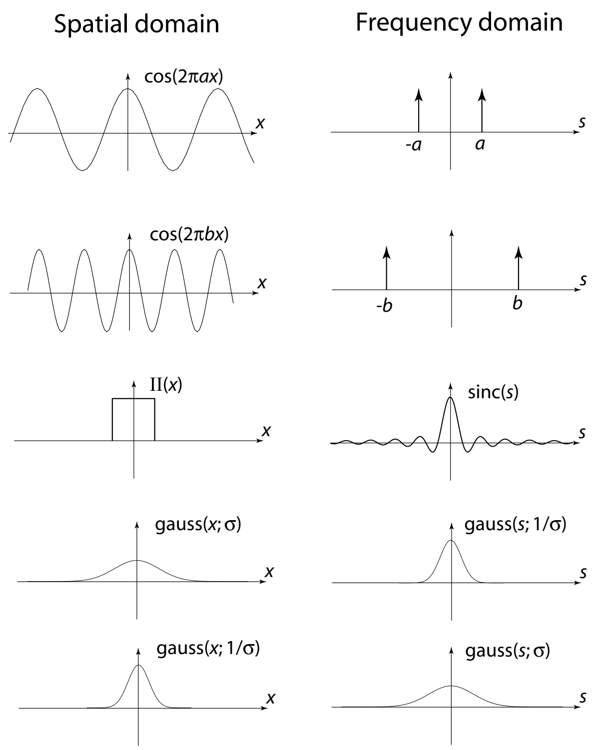

2D Fourier transform

Spatial domain → $F(s_x, s_y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y) e^{-i2\pi(s_x x + s_y y)} dx dy$ → Frequency domain

Spatial domain ← $f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(s_x, s_y) e^{i2\pi(s_x x + s_y y)} ds_x ds_y$ ← Frequency domain

Spatial domain

Frequency domain

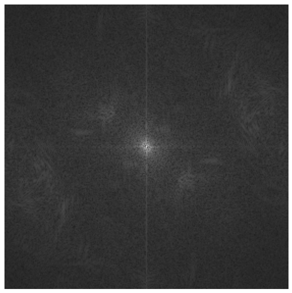

Convolution

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

In 1D, convolution is defined as:

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

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

Convolution in 2D

In two dimensions, convolution becomes:

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

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

$*$

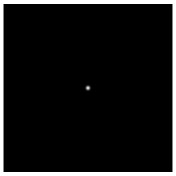

$=$

f(*x*,*y*) *h*(*x*,*y*) *g*(*x*,*y*)

Convolution theorems

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

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

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

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

1D convolution theorem example

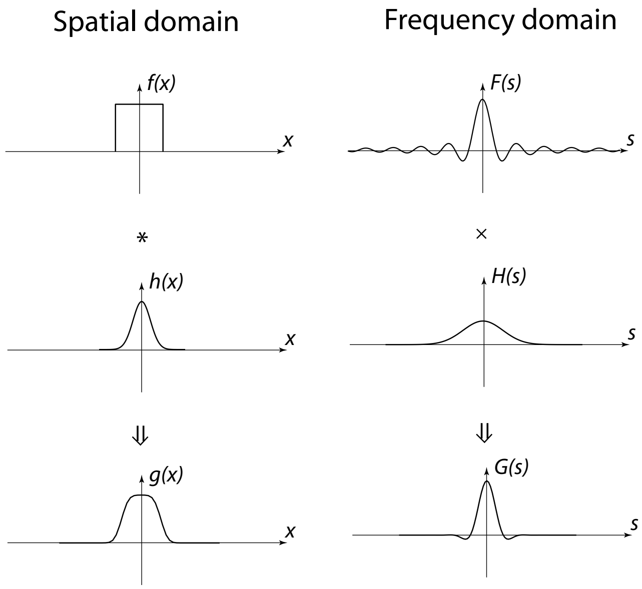

2D convolution theorem example

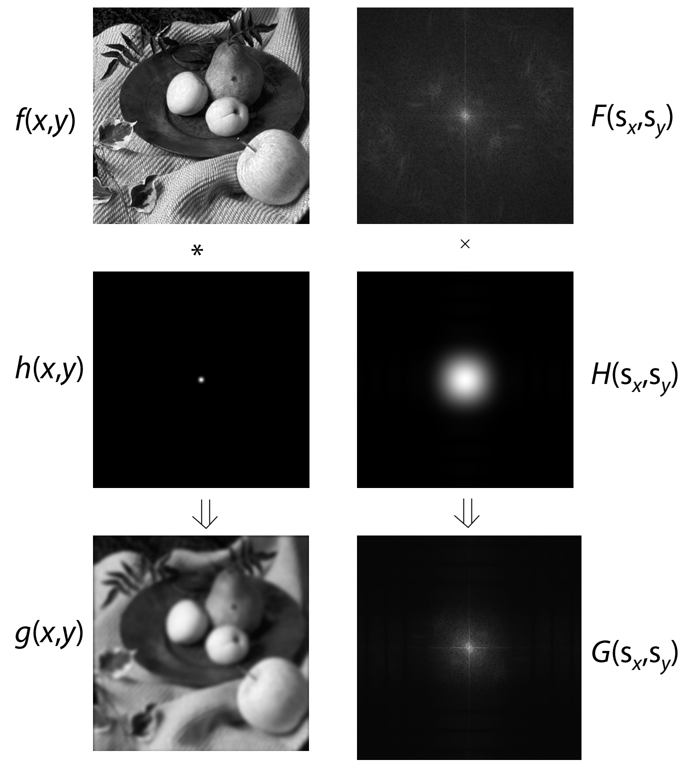

The delta function

The **Dirac delta function**, $\delta(x)$, is a handy tool for sampling theory.

It has zero width, infinite height, and unit area.

It is usually drawn as:

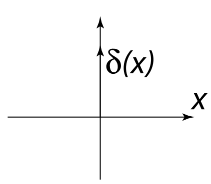

Sifting and shifting

For sampling, the delta function has two important properties.

Sifting:

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

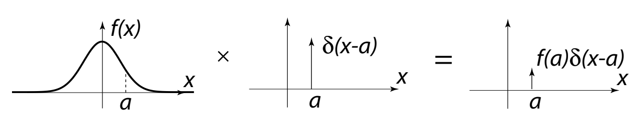

Shifting:

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

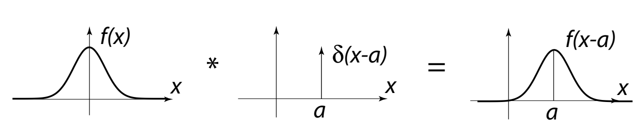

The shah/comb function

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

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

which looks like:

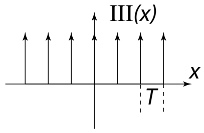

Amazingly, the Fourier transform of the shah function takes the same form:

$$\mathrm{III}(s) = \sum_{n=-\infty}^{\infty} \delta(s - ns_o)$$

where $s_o = 1/T$.

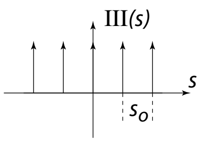

Sampling

Now, we can talk about sampling.

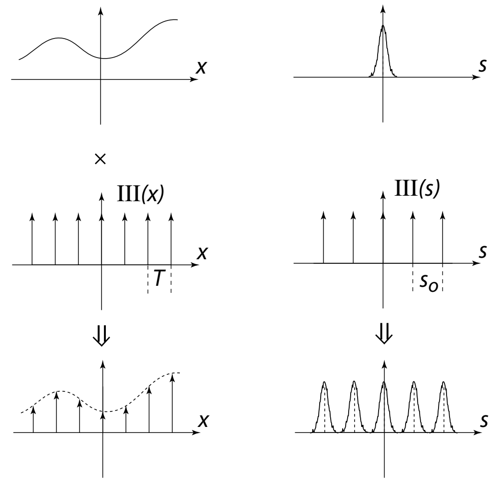

The Fourier spectrum gets *replicated* by spatial sampling!

How do we recover the signal?

Sampling and reconstruction

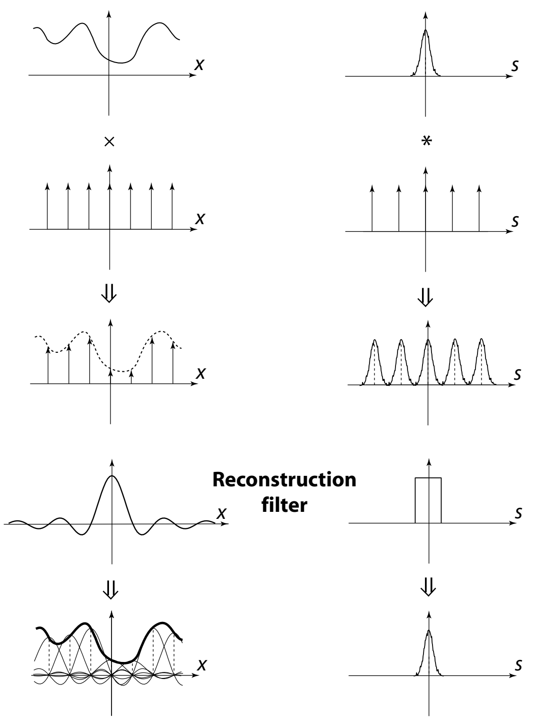

Sampling and reconstruction in 2D

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 ½ the sampling frequency.

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

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…

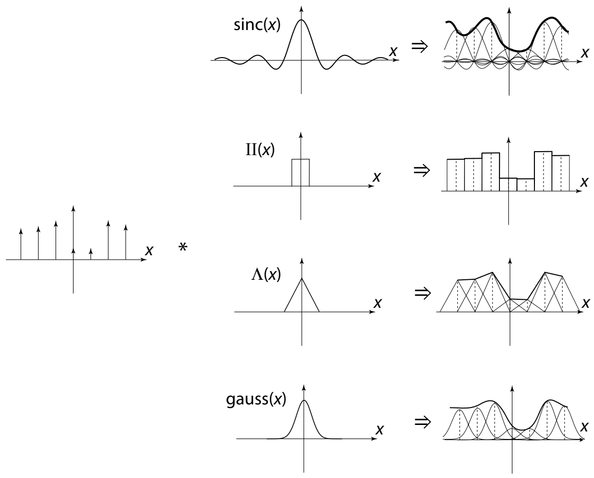

Reconstruction filters in 2D

We can also perform reconstruction in 2D…

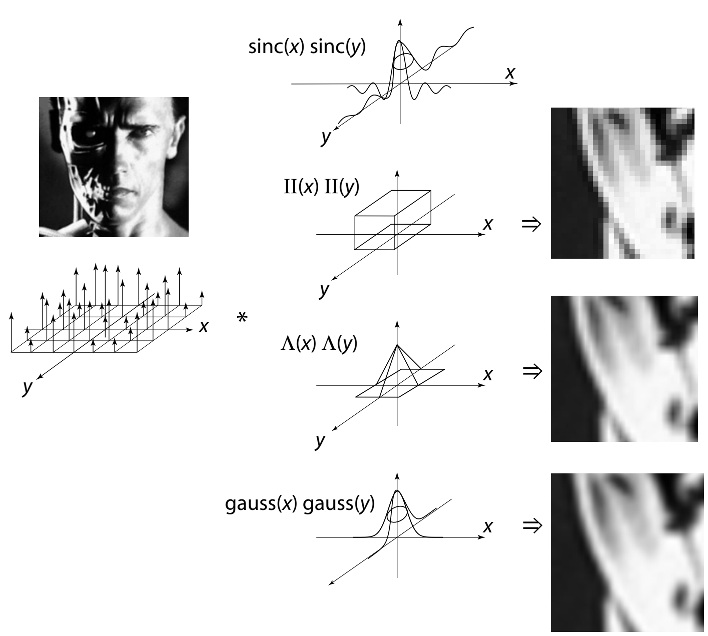

Aliasing

What if we go below the Nyquist frequency?

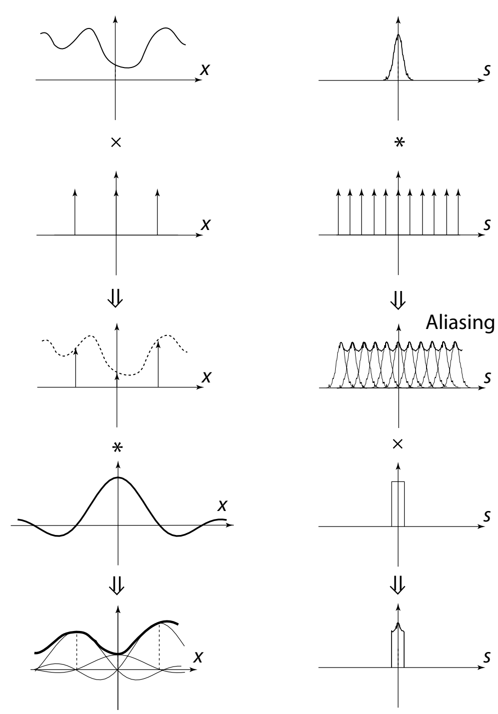

Anti-aliasing

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

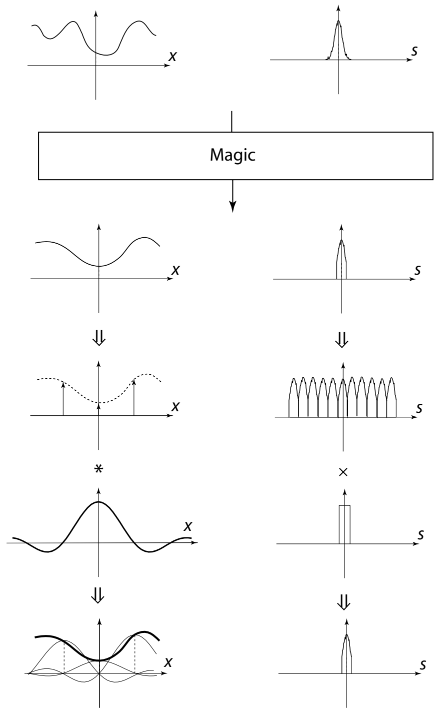

Anti-aliasing by prefiltering

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

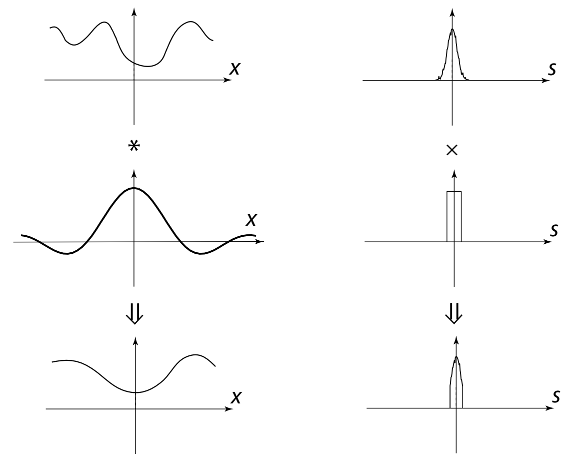

Why may this not generally be possible?

Filtered downsampling

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

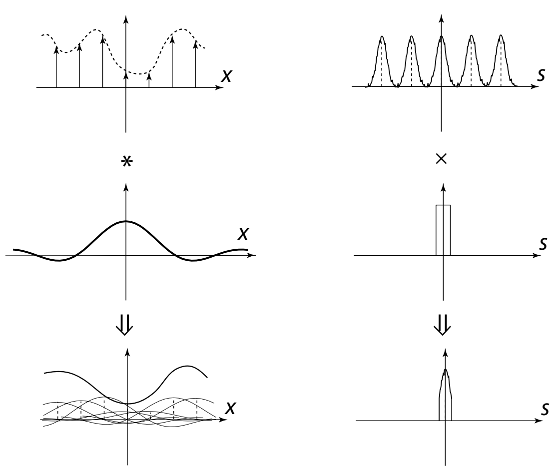

We can now sample the signal at a lower rate. The whole process is called filtered **downsampling** or **supersampling and averaging down**.