https://youtu.be/yr3ngmRuGU

Aliasing, Image Sampling and Reconstruction

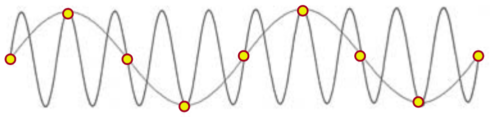

Many of the slides are taken from Thomas Funkhouser course slides and the rest from various sources over the web…

Recall: a pixel is a point…

- It is NOT a box, disc or teeny wee light
- It has no dimension
- It occupies no area
- It can have a coordinate
- *More* than a point, it is a *SAMPLE*

Image Sampling

- An image is a 2D rectilinear array of samples
  - Quantization due to limited intensity resolution
  - Sampling due to limited spatial and temporal resolution

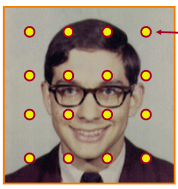

Pixels are infinitely small point samples

Imaging devices area sample.

- In video camera the CCD array is an area integral over a pixel.
- The eye: photoreceptors

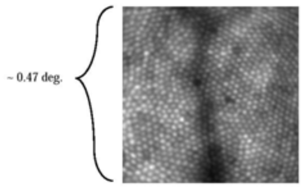

J. Liang, D. R. Williams and D. Miller, "Supernormal vision and high- resolution retinal imaging through adaptive optics," J. Opt. Soc. Am. A 14, 2884- 2892 (1997)

Aliasing

Good Old Aliasing

Reconstruction artefact

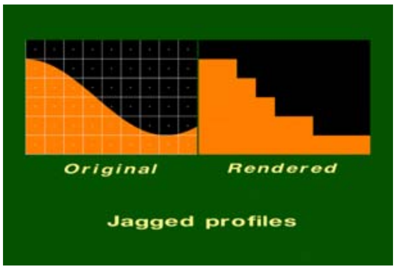

Slide © Rosalee Nerheim-Wolfe

Sampling and Reconstruction

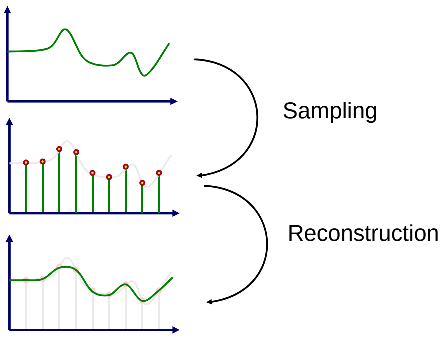

Sources of Error

- Intensity quantization
  - Not enough intensity resolution
- Spatial aliasing
  - Not enough spatial resolution
- Temporal aliasing
  - Not enough temporal resolution

Aliasing (in general)

- In general:
  - Artifacts due to under-sampling or poor reconstruction
- Specifically, in graphics:
  - Spatial aliasing
  - Temporal aliasing

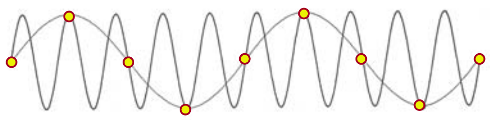

Under-sampling

Figure 14.17 FvDFH

Sampling & Aliasing

- Real world is continuous
- The computer world is discrete
- Mapping a continuous function to a discrete one is called sampling
- Mapping a continuous variable to a discrete one is called quantizaion
- To represent or render an image using a computer, we must both sample and quantize

Spatial Aliasing

- Artifacts due to limited spatial resolution

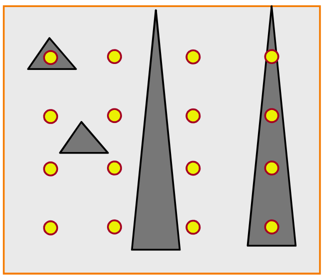

Can be a serious problem…

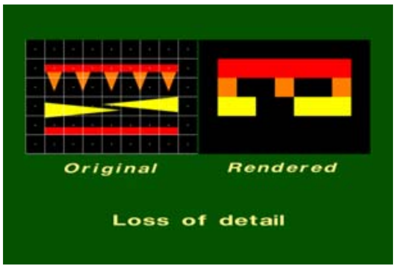

Slide © Rosalee Nerheim-Wolfe

Spatial Aliasing

- Artifacts due to limited spatial resolution

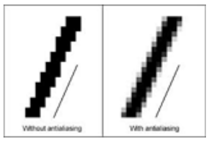

“Jaggies”

Aliasing

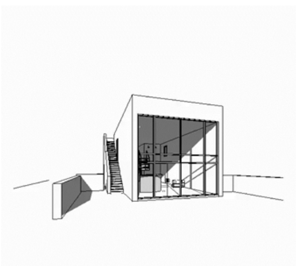

Anti-aliasing

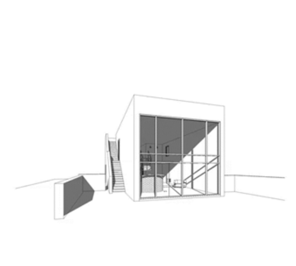

Spatial Aliasing

- Artifacts due to limited spatial resolution

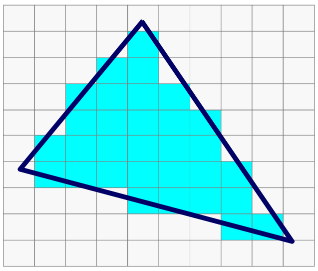

"Jaggies"

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

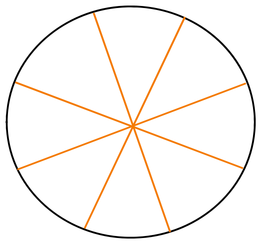

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

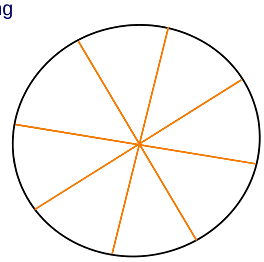

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

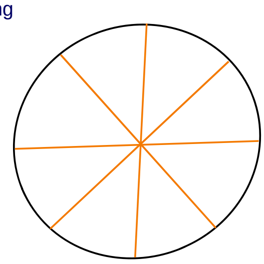

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

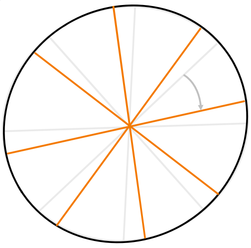

Temporal Aliasing

Staircasing or *Jaggies*

The raster *aliasing* effect – removal is called *antialiasing*

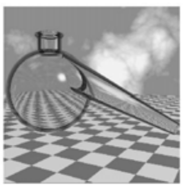

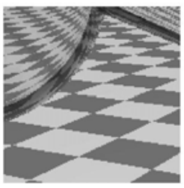

Images by Don Mitchell

…very serious problem!

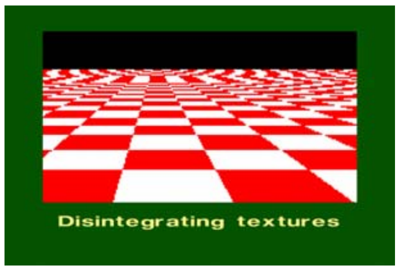

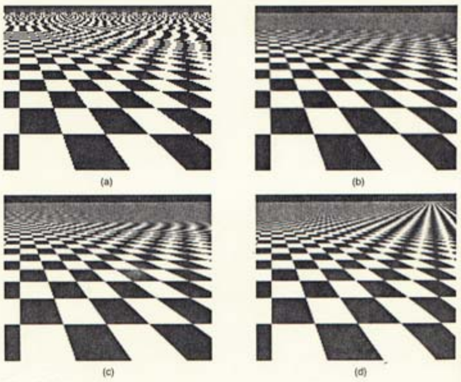

Fig. 17.12 (a) Point sampling of the source image. (b) MIP map filtering. (c) Summed area table filtering. (d) Elliptical weighted average using MIP maps and a Gaussian filter. (Courtesy of P. Heckbert.)

Blurring doesn't work well.

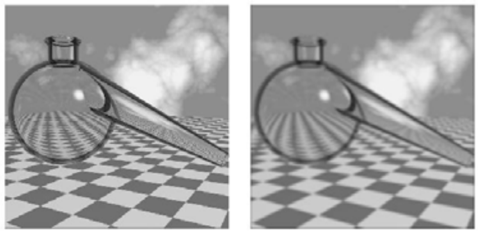

Removed the *jaggies*, but also all the detail ! → Reduction in resolution

Unweighted Area Sampling

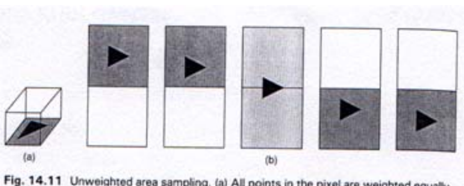

Fig. 14.11 Unweighted area sampling. (a) All points in the pixel are weighted equally. (b) Changes in computed intensities as an object moves between pixels.

Weighted Area Sampling

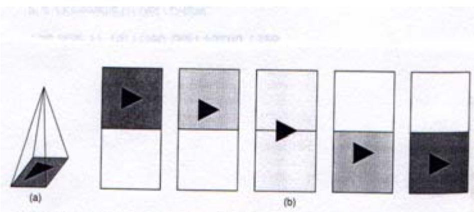

Fig. 14.12 Weighted area sampling. (a) Points in the pixel are weighted differently. (b) Changes in computed intensities as an object moves between pixels.

…with Overlap

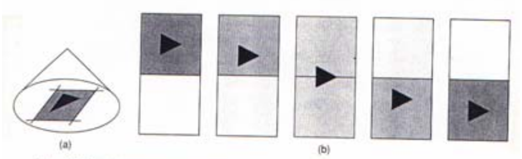

Fig. 14.13 Weighted area sampling with overlap. (a) Typical weighting function. (b) Changes in computed intensities as an object moves between pixels.

Sampling and Reconstruction

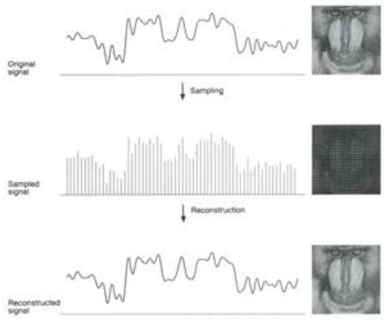

Figure 19.9 FvDFH

Antialiasing

- Sample at higher rate
  - Not always possible
  - Doesn't always solve problem
- Pre-filter to form bandlimited signal
  - Form bandlimited function (low-pass filter)
  - Trades aliasing for blurring

Must consider sampling theory!

How is antialiasing done?

- We need some mathematical tools to
  - analyse the situation.
  - find an optimum solution.
- Tools we will use :
  - Fourier transform.
  - Convolution theory.
  - Sampling theory.

We need to understand the behavior of the signal in frequency domain

Spectral Analysis / Fourier Transforms

- Spectral representation treats the function as a weighted sum of sines and cosines
- Every function has two representations
  - Spatial (time) domain - normal representation
  - Frequency domain - spectral representation
- The *Fourier transform* converts between the spatial and frequency domains.

Spectral Analysis / Fourier Transforms

Note the Euler formula : $e^{it} = \cos t + i \sin t$

Spatial domain → $F(\omega) = \int_{-\infty}^{\infty} f(x) e^{-i\omega x} dx$ → Frequency domain.

Frequency domain. → $f(x) = \frac{1}{2\Pi} \int_{-\infty}^{\infty} F(\omega) e^{i\omega x} d\omega$ → Spatial domain

- The *Fourier transform* converts between the spatial and frequency domain.
- Real and imaginary components.
- Forward and reverse transforms very similar.

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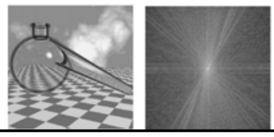 $F(\omega)$

Sampling Theory

- How many samples are required to represent a given signal without loss of information?
- What signals can be reconstructed without loss for a given sampling rate?

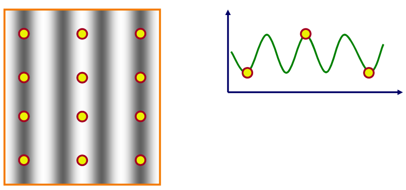

Spectral Analysis

- Spatial domain:
  - Function: f(x)
  - Filtering: convolution
- Frequency domain:
  - Function: F(u)
  - Filtering: multiplication

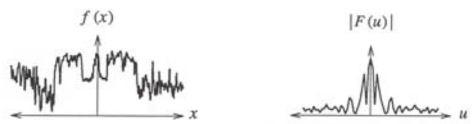

Any signal can be written as a sum of periodic functions.

Fourier Transform

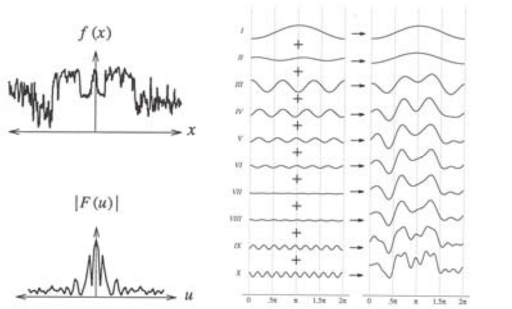

Figure 2.6 Wolberg

Fourier Transform

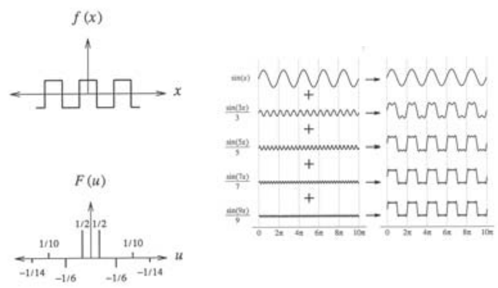

Figure 2.5 Wolberg

Sampling Theorem

- A signal can be reconstructed from its samples, if the original signal has no frequencies above 1/2 the sampling frequency - Shannon
- The minimum sampling rate for bandlimited function is called “Nyquist rate”

A signal is bandlimited if its highest frequency is bounded. The frequency is called the bandwidth.

Convolution

- Convolution of two functions (= filtering):

$$g(x) = f(x) \otimes h(x) = \int^{\infty} f(\lambda) h(x - \lambda) d\lambda$$

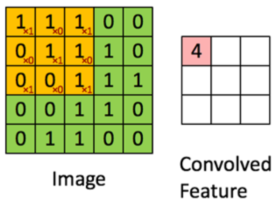

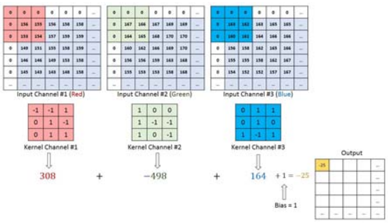

Convolution in frequency domain is same as multiplication in spatial domain, and vice-versa

$$f \otimes g = F \times G$$

$$F \otimes G = f \times g$$

Image Processing

- Image processing is a resampling problem

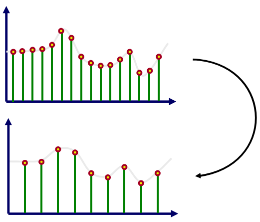

Resampling

Antialiasing in Image Processing

- General Strategy
  - Pre-filter transformed image via convolution with low-pass filter to form bandlimited signal
- Rationale
  - Prefer blurring over aliasing

Filtering in the frequency domain

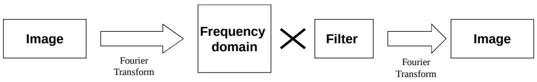

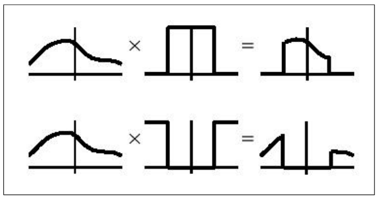

Lowpass filter

Highpass filter

Low-pass Filtering

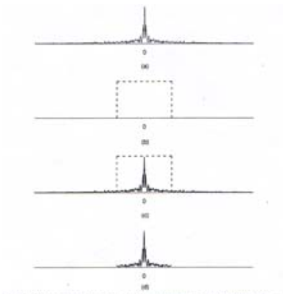

Fig. 14.21 Low-pass filtering in the frequency domain. (a) Original spectrum. (b) Low-pass filter. (c) Spectrum with filter. (d) Filtered spectrum. (Courtesy of George Wolberg, Columbia University.)

Low and High Pass Filtering.

- Low pass

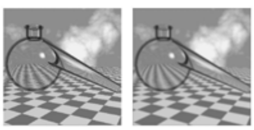

- High pass

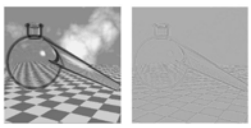

Low-pass Filtering

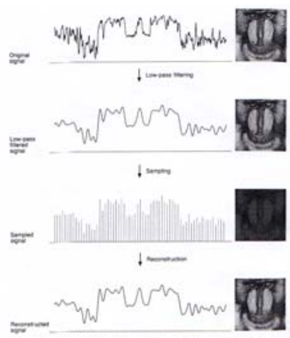

Fig. 14.20 The sampling pipeline with filtering. (Courtesy of George Wolberg, Columbia University.)

Low-pass Filtering

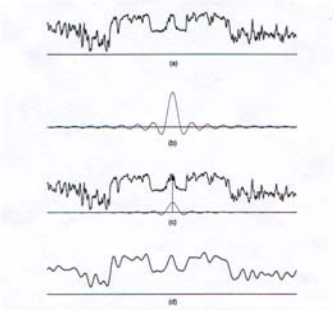

Fig. 14.23 Low-pass filtering in the spatial domain. (a) Original signal. (b) Sinc filter. (c) Signal with filter, with value of filtered signal shown as a black dot (•) at filter's origin. (d) Filtered signal. (Courtesy of George Wolberg, Columbia University.)

How can we represent sampling ?

Multiplication of the sample with a regular train of delta functions.

Sampling: Frequency domain.

The Fourier transform of regular comb of delta functions is a comb.
Spacing is inversely proportional

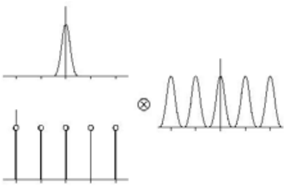

Sampling , the Comb function

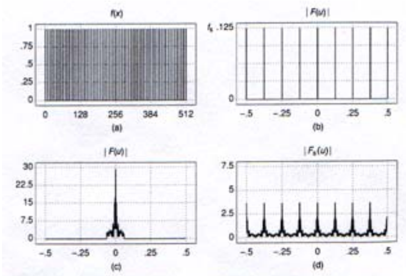

Fig. 14.26 (a) Comb function and (b) its Fourier transform. Convolving the comb's Fourier transform with (c) a signal's Fourier transform in the frequency domain yields (d) the replicated spectrum of the sampled signal. (Courtesy of George Wolberg, Columbia University.)

Reconstruction in frequency domain

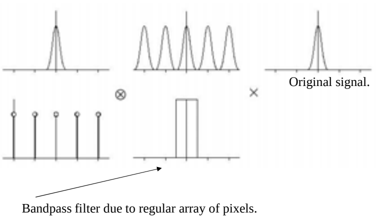

Undersampling leads to aliasing.

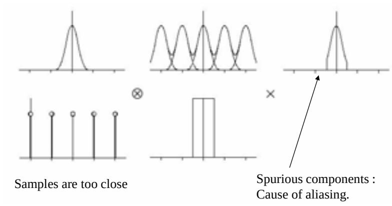

The Sampling Theorem.

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 1/2 the sampling frequency***

For a given bandlimited function, the rate at which it must be sampled is called the *Nyquist Frequency*

Sampling at the Nyquist Frequency

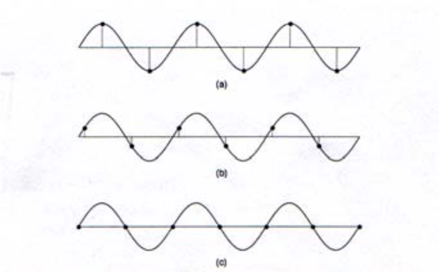

Fig. 14.16 Sampling at the Nyquist frequency (a) at peaks, (b) between peaks, (c) at zero crossings. (Courtesy of George Wolberg, Columbia University.)

Aliasing in the frequency domain

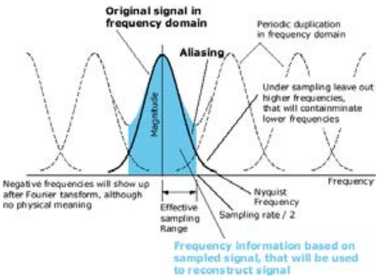

Removing aliasing is called *antialiasing*

How do we remove aliasing ?

- Perfect solution - prefilter with perfect bandpass filter.

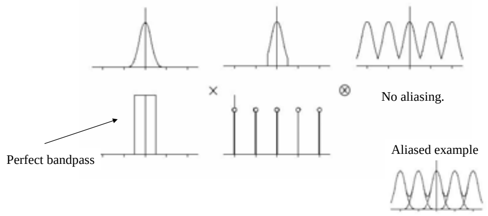

How do we remove aliasing ?

- Perfect solution - prefilter with perfect bandpass filter.
  - Difficult/Impossible to do in frequency domain.
- Convolve with sinc function in space domain
  - Optimal filter - better than area sampling.
  - Sinc function is infinite !!
  - Computationally expensive.

Adequate sampling

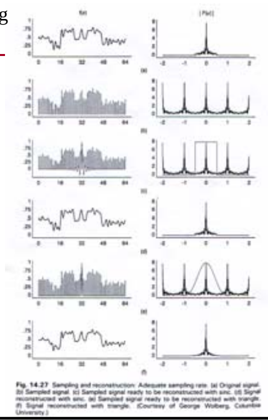

Inadequate sampling

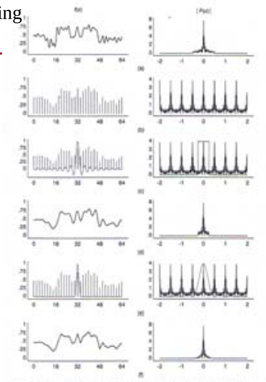

Fig. 14.28 Sampling and reconstruction: Inadequate sampling rate. (a) Original signal. (b) Sampled signal. (c) Sampled signal ready to be reconstructed with sinc. (d) Signal reconstructed with sinc. (e) Sampled signal ready to be reconstructed with triangle. (f) Signal reconstructed with triangle. (Courtesy of George Wolberg, Columbia University.)

Pre-filtering

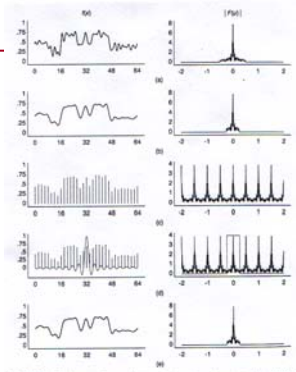

Fig. 14.29 Filtering, sampling, and reconstruction: Sampling rate is adequate after filtering. (a) Original signal. (b) Low-pass filtered signal. (c) Sampled signal. (d) Sampled signal ready to be reconstructed with sinc. (e) Signal reconstructed with sinc. (Courtesy of George Wolberg, Columbia University.)

The 'Sinc' function.

$$\int_{-\infty}^{\infty} square(x)e^{-i\omega x}dx = \int_{-1/2}^{1/2} e^{-i\omega x}dx$$

$$= \left.\frac{e^{-i\omega x}}{-i\omega x}\right|_{-\frac{1}{2}}^{\frac{1}{2}}$$

$$= \frac{e^{-\frac{1}{2}i\omega} - e^{\frac{1}{2}i\omega}}{\frac{1}{2}i\omega + \frac{1}{2}i\omega} = \frac{\sin\frac{1}{2}\omega}{\frac{1}{2}\omega}$$

$$e^{it} = \cos t + i \sin t$$

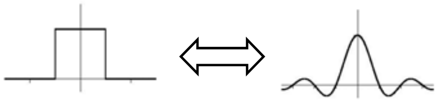

The Sinc Filter

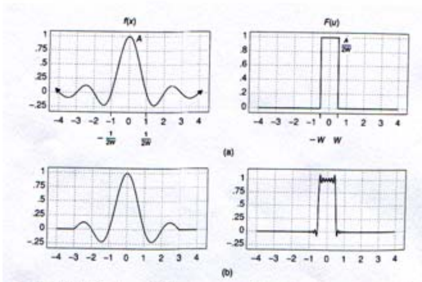

Fig. 14.24 (a) Sinc in spatial domain corresponds to pulse in frequency domain. (b) Truncated sinc in spatial domain corresponds to ringing pulse in frequency domain. (Courtesy of George Wolberg, Columbia University.)

Common Filters

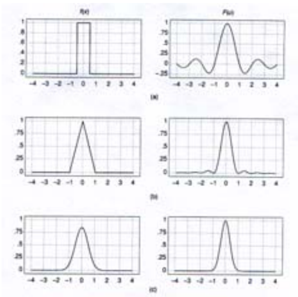

Fig. 14.25 Filters in spatial and frequency domains. (a) Pulse—sinc. (b) Triangle—sinc². (c) Gaussian—Gaussian. (Courtesy of George Wolberg, Columbia University.)

Sample-and-Hold

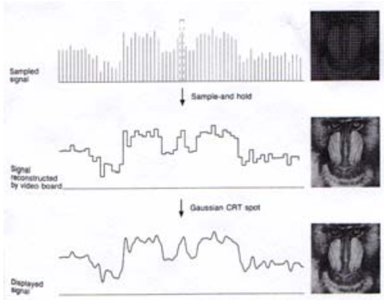

Fig. 14.30 Reconstruction by sample and hold and Gaussian CRT spot. (Courtesy of George Wolberg, Columbia University.)

Image Reconstruction

- Re-create continuous image from samples
  - Example: cathode ray tube

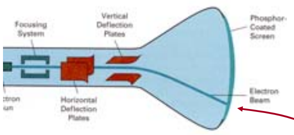

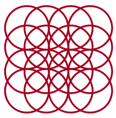

Image is reconstructed by displaying pixels with finite area (Gaussian)

End…

Adjusting Brightness

- Simply scale pixel components
  - Must clamp to range (e.g., 0 to 255)

Original

Brighter

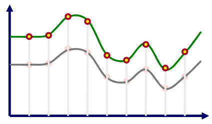

Adjusting Contrast

- Compute mean luminance L for all pixels
  - luminance = 0.30*r + 0.59*g + 0.11*b
- Scale deviation from L for each pixel component
  - Must clamp to range (e.g., 0 to 255)

Original

More Contrast

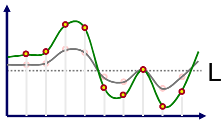

Image Processing

- Quantization
  - Uniform Quantization
  - Random dither
  - Ordered dither
  - Floyd-Steinberg dither
- Pixel operations
  - Add random noise
  - Add luminance
  - Add contrast
  - Add saturation
- Filtering
  - Blur
  - Detect edges
- Warping
  - Scale
  - Rotate
  - Warps
- Combining
  - Morphs
  - Composite

Adjust Blurriness

- Convolve with a filter whose entries sum to one
  - Each pixel becomes a weighted average of its neighbors

Original

Blur

$$\text{Filter} = \begin{bmatrix} \tfrac{1}{16} & \tfrac{2}{16} & \tfrac{1}{16} \\ \tfrac{2}{16} & \tfrac{4}{16} & \tfrac{2}{16} \\ \tfrac{1}{16} & \tfrac{2}{16} & \tfrac{1}{16} \end{bmatrix}$$

Edge Detection

- Convolve with a filter that finds differences between neighbor pixels

Original

Detect edges

$$\text{Filter} = \begin{bmatrix} -1 & -1 & -1 \\ -1 & +8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Image Processing

- Quantization
  - Uniform Quantization
  - Random dither
  - Ordered dither
  - Floyd-Steinberg dither
- Pixel operations
  - Add random noise
  - Add luminance
  - Add contrast
  - Add saturation
- Filtering
  - Blur
  - Detect edges
- Warping
  - Scale
  - Rotate
  - Warps
- Combining
  - Morphs
  - Composite

Scaling

- Resample with triangle or Gaussian filter

More on this next lecture!

Original

1/4X resolution

4X resolution

Summary

- Image processing is a resampling problem
  - Avoid aliasing
  - Use filtering

Triangle Filter

- Convolution with triangle filter

- Convolution with Gaussian filter

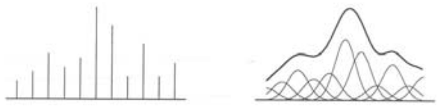

Figure 2.4 Wolberg