Image Sampling and Reconstruction

Thomas Funkhouser
Princeton University
C0S 426, Fall 2000

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

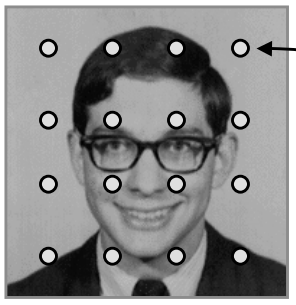

Pixels are infinitely small point samples

Image Reconstruction

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

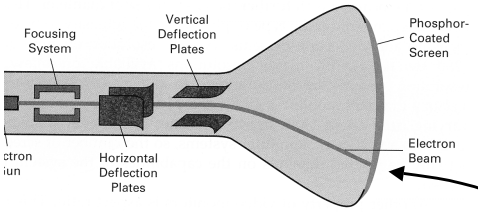

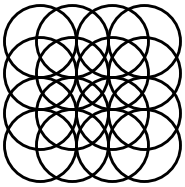

Image is reconstructed by displaying pixels with finite area (Gaussian)

Sampling and Reconstruction

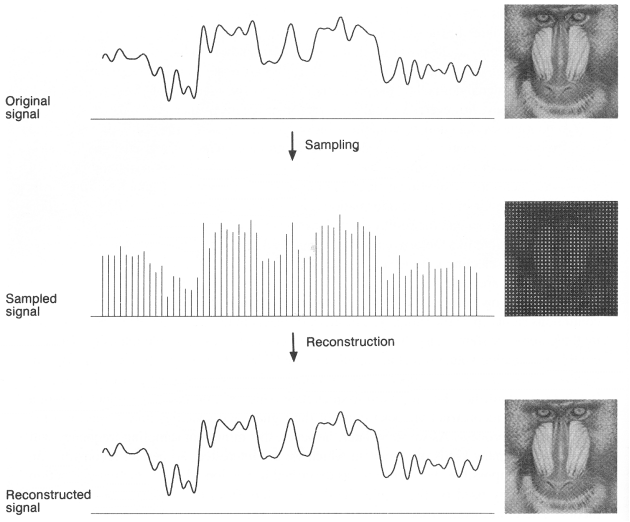

Figure 19.9 FvDFH

Sampling and Reconstruction

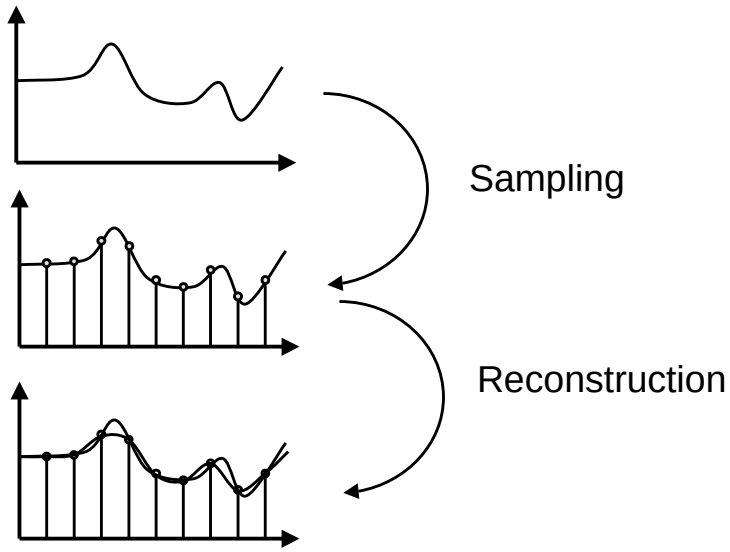

Sources of Error

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

$$E^2 = \sum_{(x,y)} \left(I(x,y) - P(x,y)\right)^2$$

Aliasing (in general)

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

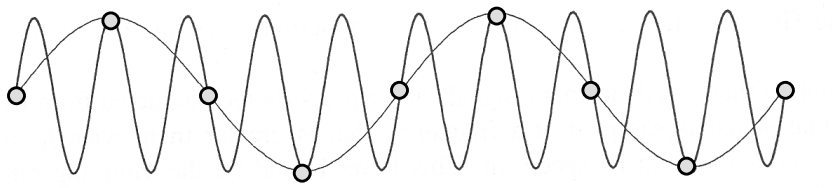

Under-sampling

Figure 14.17 FvDFH

Spatial Aliasing

- Artifacts due to limited spatial resolution

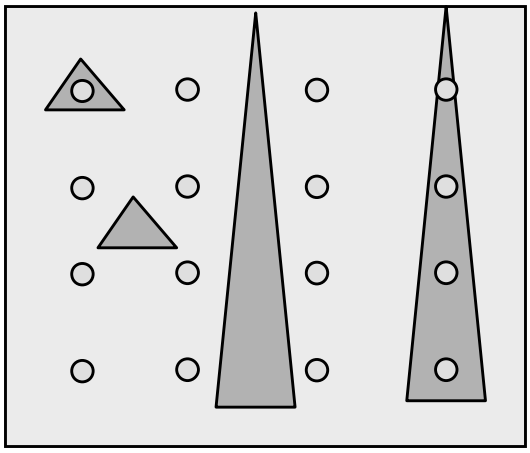

Spatial Aliasing

- Artifacts due to limited spatial resolution

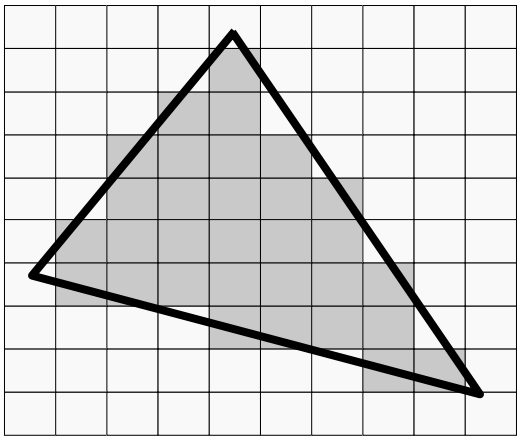

"Jaggies"

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

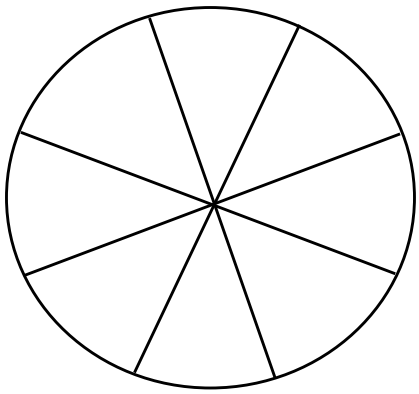

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

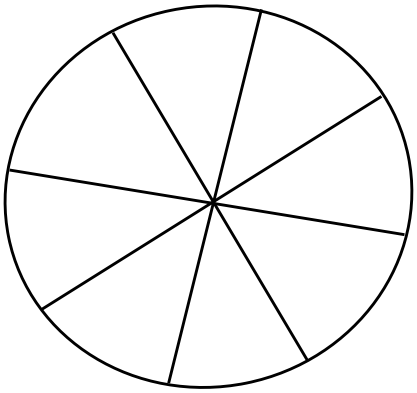

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

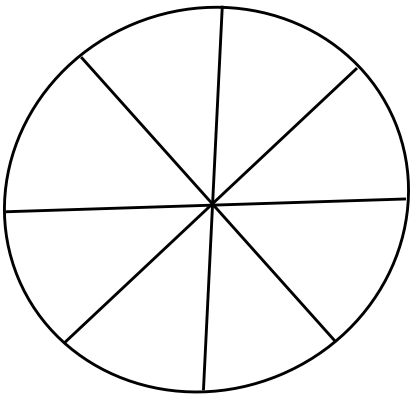

Temporal Aliasing

- Artifacts due to limited temporal resolution
  - Strobing
  - Flickering

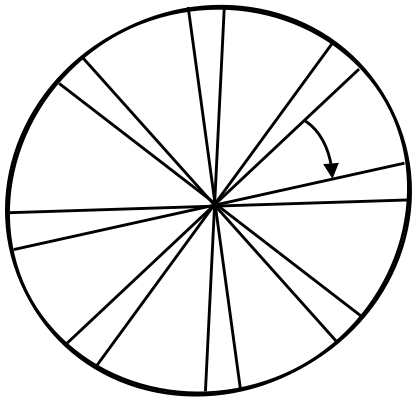

Antialiasing

- 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!

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?

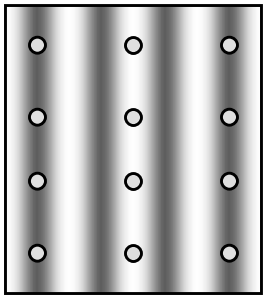

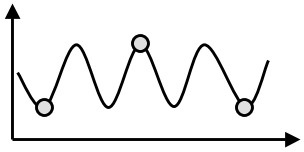

Spectral Analysis

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

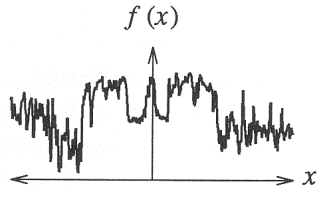

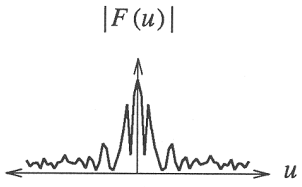

Any signal can be written as a sum of periodic functions.

Fourier Transform

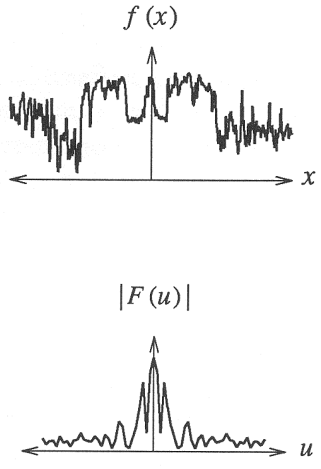

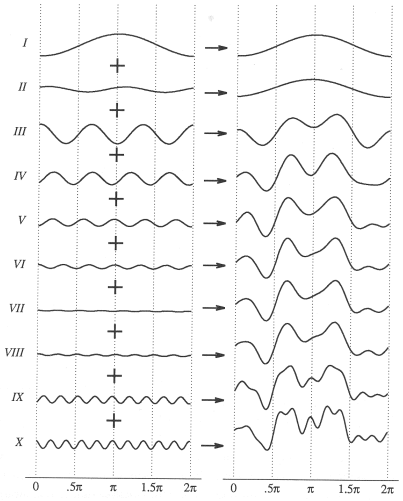

Figure 2.6 Wolberg

Fourier Transform

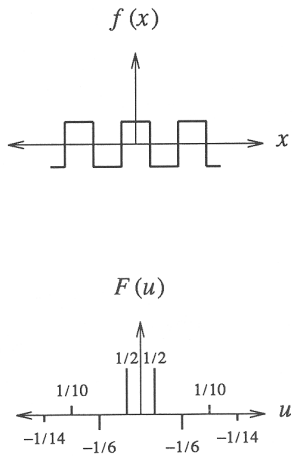

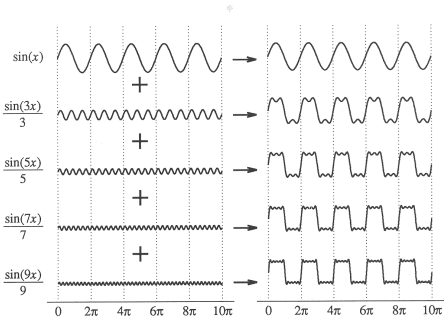

Figure 2.5 Wolberg

Fourier Transform

- Fourier transform:

$$F(u) = \int_{-\infty}^{\infty} f(x) e^{-i 2\pi x} dx$$

- Inverse Fourier transform:

$$f(x) = \int_{-\infty}^{\infty} F(u) e^{+i 2\pi u} du$$

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):

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

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

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

Image Processing

- Consider reducing the image resolution

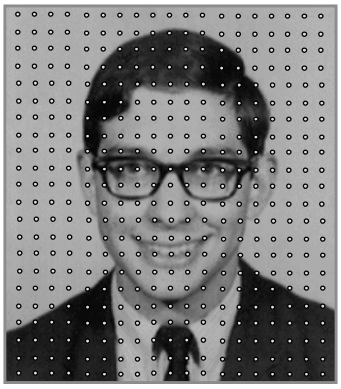

Original image

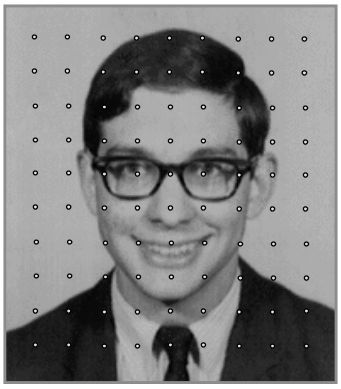

1/4 resolution

Image Processing

- Image processing is a resampling problem

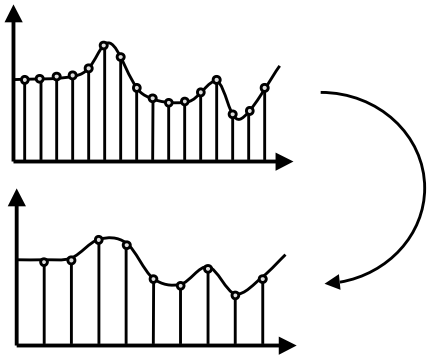

Thou shalt avoid aliasing!

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

Image Processing

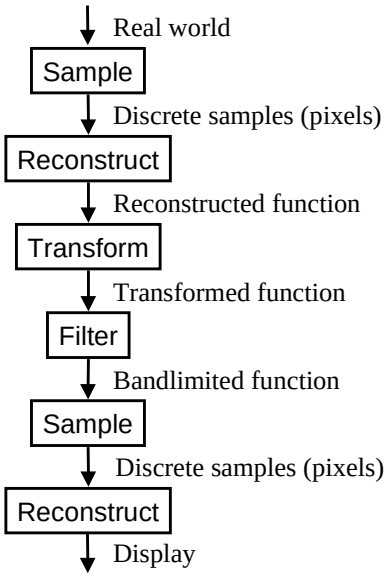

Image Processing

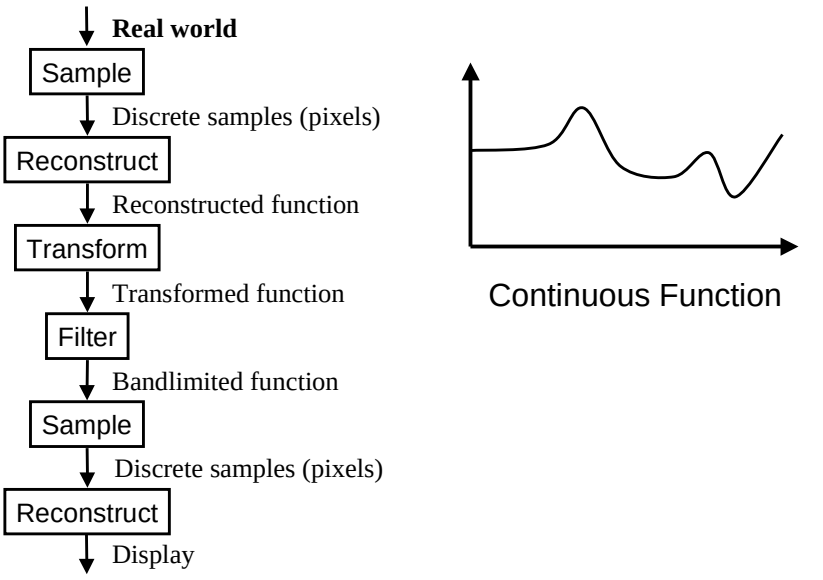

Image Processing

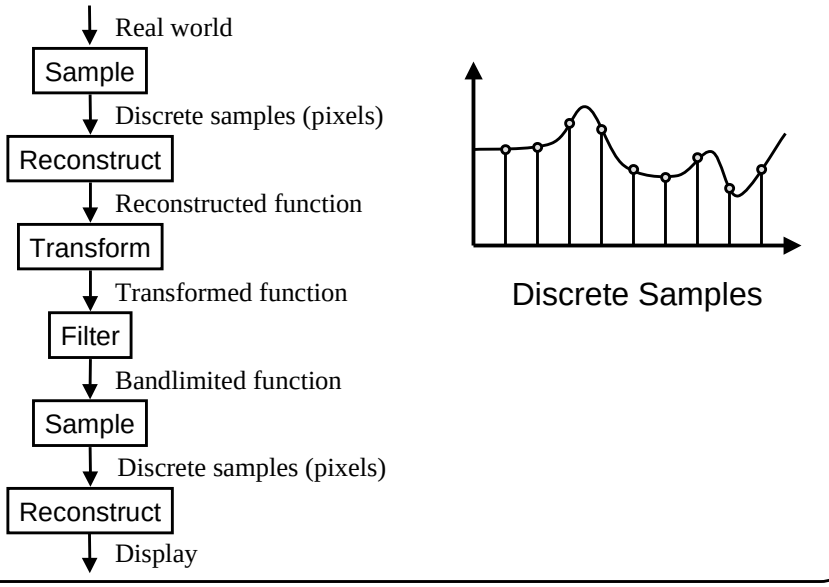

Image Processing

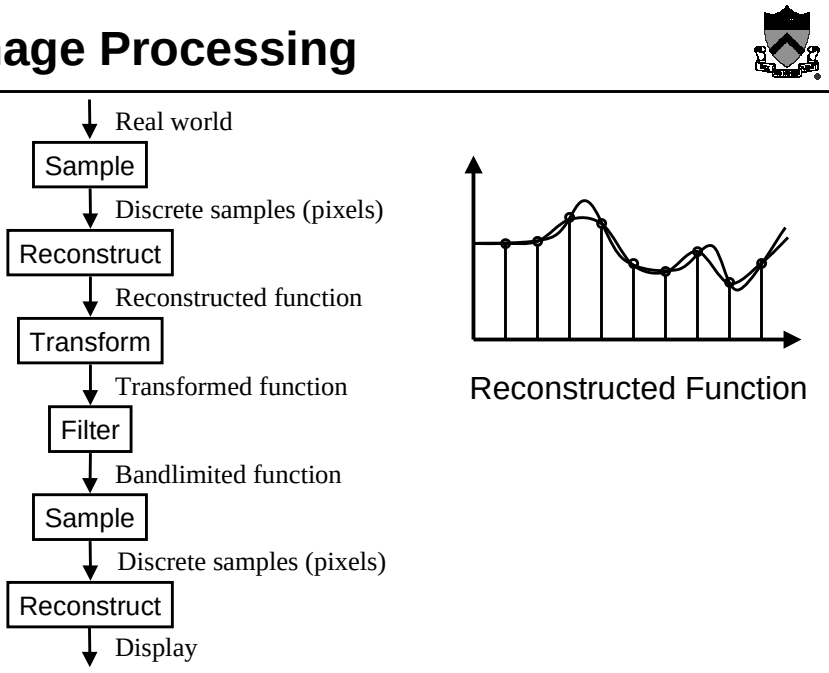

Reconstructed Function

Image Processing

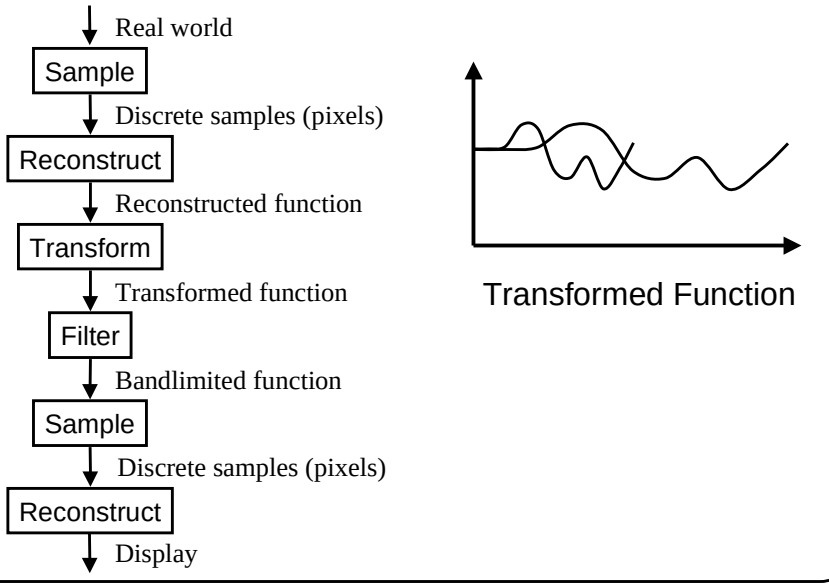

Transformed Function

Image Processing

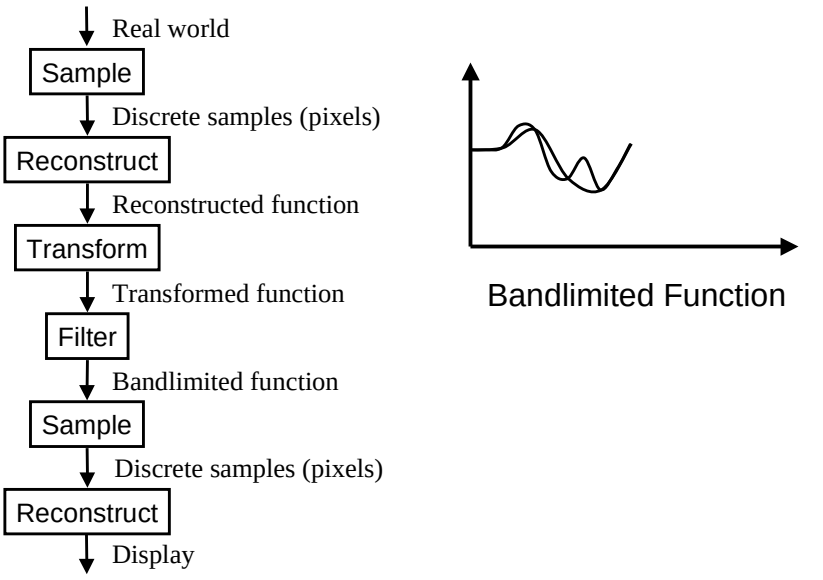

Image Processing

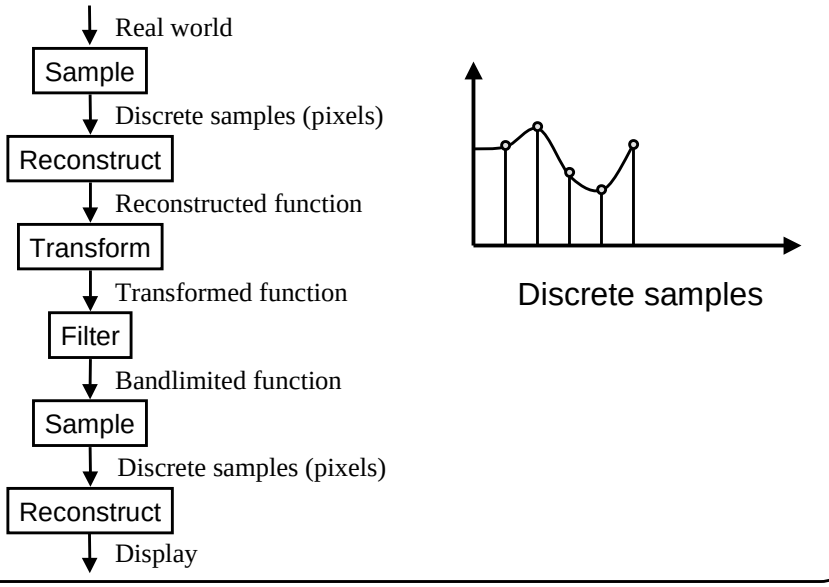

Image Processing

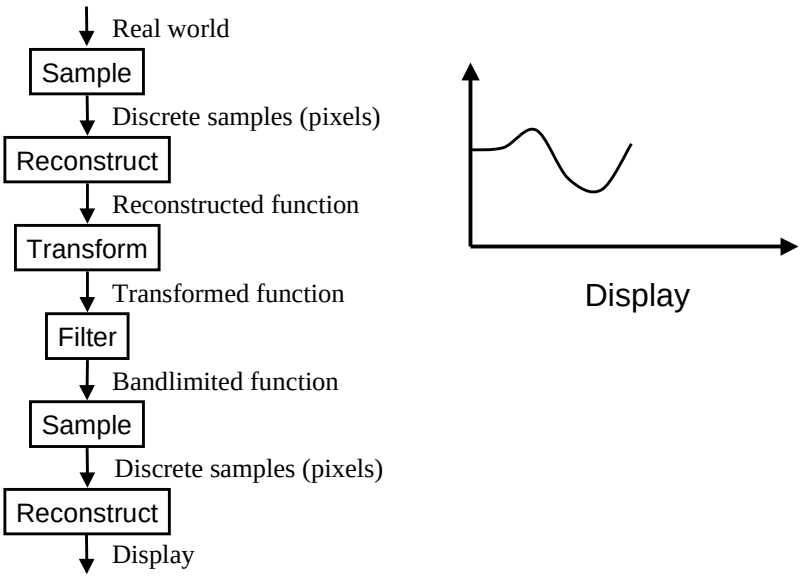

Ideal Low-Pass Filter

- Frequency domain

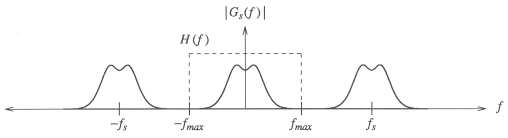

- Spatial domain

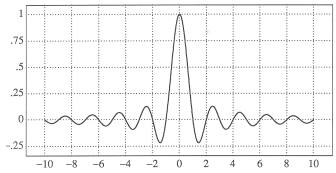

$$Sinc(x) = \frac{\sin \pi x}{\pi x}$$

Figure 4.5 Wolberg

Practical Image Processing

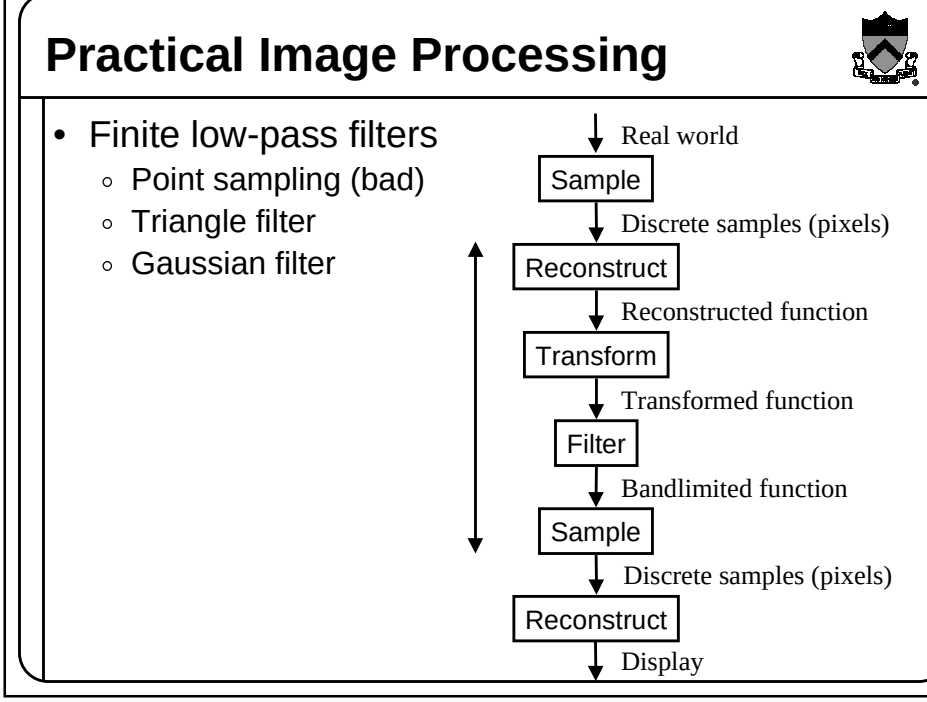

- Finite low-pass filters
  - Point sampling (bad)
  - Triangle filter
  - Gaussian filter

Triangle Filter

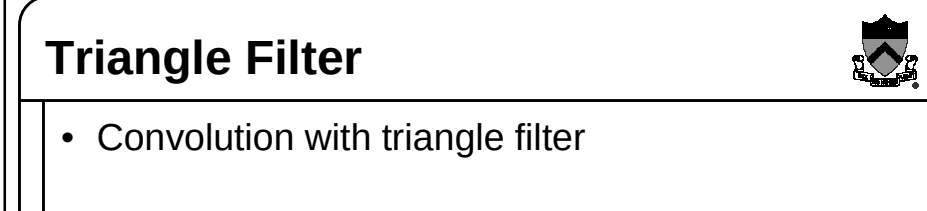

- Convolution with triangle filter

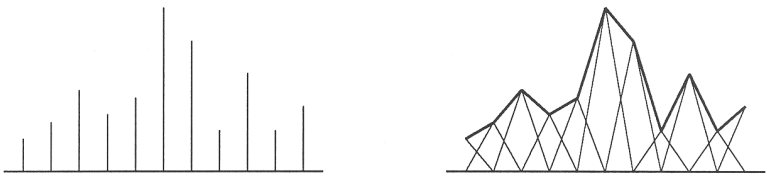

Input Output

Figure 2.4 Wolberg

Gaussian Filter

- Convolution with Gaussian filter

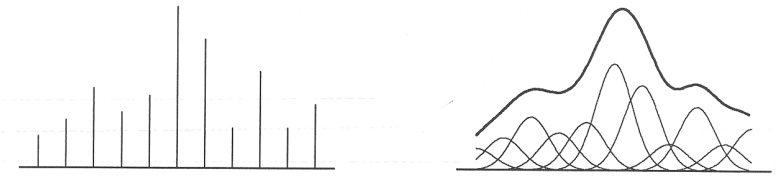

Input Output

Figure 2.4 Wolberg

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

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

Adjusting Brightness

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

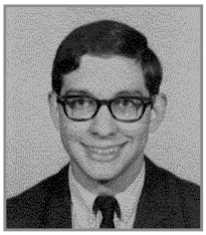

Original

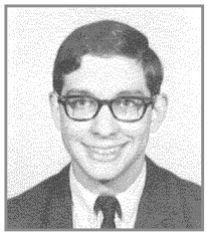

Brighter

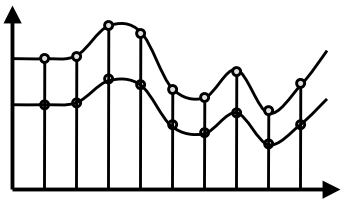

Adjusting Contrast

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

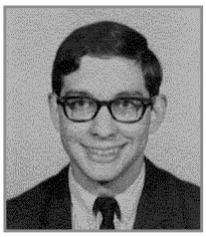

Original

More Contrast

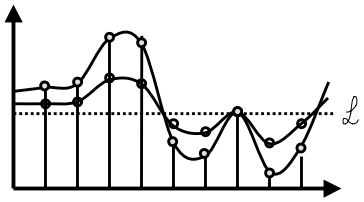

Image Processing

- 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

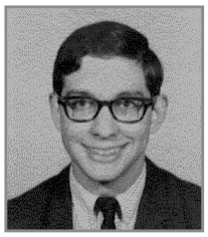

Original

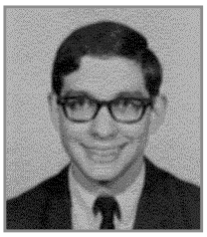

Blur

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

Edge Detection

- Convolve with a filter that finds differences between neighbor pixels

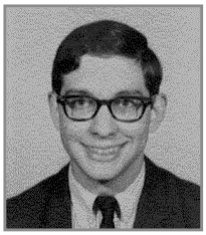

Original

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!

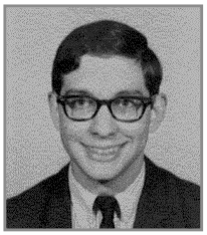

Original

1/4X resolution

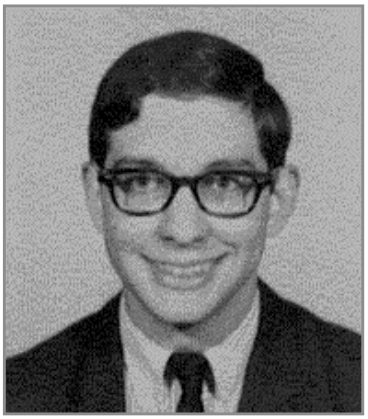

4X resolution

Summary

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