




Image Processing

COS 426




What is a Digital Image?

A digital image is a discrete array of samples representing a continuous 2D function



Continuous function



Discrete samples



Limitations on Digital Images

- Spatial discretization
- Quantized intensity
- Approximate color (RGB)
- (Temporally discretized frames for digital video)



Image Processing

- Changing intensity/color
 - Linear: scale, offset, etc.
 - Nonlinear: gamma, saturation, etc.
 - Add random noise
- Filtering over neighborhoods
 - Blur
 - Detect edges
 - Sharpen
 - Emboss
 - Median
- Moving image locations
 - Scale
 - Rotate
 - Warp
- Combining images
 - Composite
 - Morph



Digital Image Processing

- Changing intensity/color
 - Linear: scale, offset, etc.
 - Nonlinear: gamma, saturation, etc.
 - Add random noise
- Filtering over neighborhoods
 - Blur
 - Detect edges
 - Sharpen
 - Emboss
 - Median
- Moving image locations
 - Scale
 - Rotate
 - Warp
- Combining images
 - Composite
 - Morph
- Quantization
- Spatial / intensity tradeoff
 - Dithering



Adjusting Brightness


- Simply scale pixel components
 - Must clamp to range, e.g. [0..1] or [0..255]



Original



Brighter



Note: this is “contrast” on your monitor!
“Brightness” adjusts black level (offset)



Adjusting Contrast


- Compute mean luminance L for all pixels
 - $\text{luminance} = 0.30*r + 0.59*g + 0.11*b$
- Scale deviation from L for each pixel component
 - Must clamp to range (e.g., 0 to 1)



Original




More Contrast



Digression: Perception of Intensity

- Perception of intensity is nonlinear




Modeling Nonlinear Intensity Response

- Brightness (B) usually modeled as a logarithm or power law of intensity (I)

$$B = k \log I$$

$$B = I^{1/3}$$



- Exact curve varies with ambient light, adaptation of eye

Cameras

- Original cameras based on Vidicon obey power law for Voltage (V) vs. Intensity (I):

$$V = I^\gamma$$

$$\gamma \approx 0.45$$



Vidicon tube [wikipedia.org]

CRT Response

- Power law for Intensity (I) vs. applied voltage (V)

$$I = V^\gamma$$

$$\gamma \approx 2.5$$



CRT [wikipedia.org]

- Vidicon + CRT = almost linear!
- Other displays (e.g. LCDs) contain electronics to emulate this law

CCD Cameras

- Camera gamma codified in NTSC standard
- CCDs have linear response to incident light
- Electronics to apply required power law

- So, pictures from most cameras (including digital still cameras) will have $\gamma = 0.45$
 - sRGB standard: partly-linear, partly power-law curve well approximated by $\gamma = 1 / 2.2$



Digital Image Processing


- Changing intensity/color
 - Linear: scale, offset, etc.
 - Nonlinear: gamma, saturation, etc.
 - Add random noise
- Filtering over neighborhoods
 - Blur
 - Detect edges
 - Sharpen
 - Emboss
 - Median
- Moving image locations
 - Scale
 - Rotate
 - Warp
- Combining images
 - Composite
 - Morph
- Quantization
- Spatial / intensity tradeoff
 - Dithering



Basic Operation: Convolution



Output value is weighted sum of values in neighborhood of input image

- Pattern of weights is the “filter” or “kernel”






Convolution with a Triangle Filter






Input



Output





Convolution with a Triangle Filter






Convolution with a Triangle Filter

What if the filter runs off the end?



Input

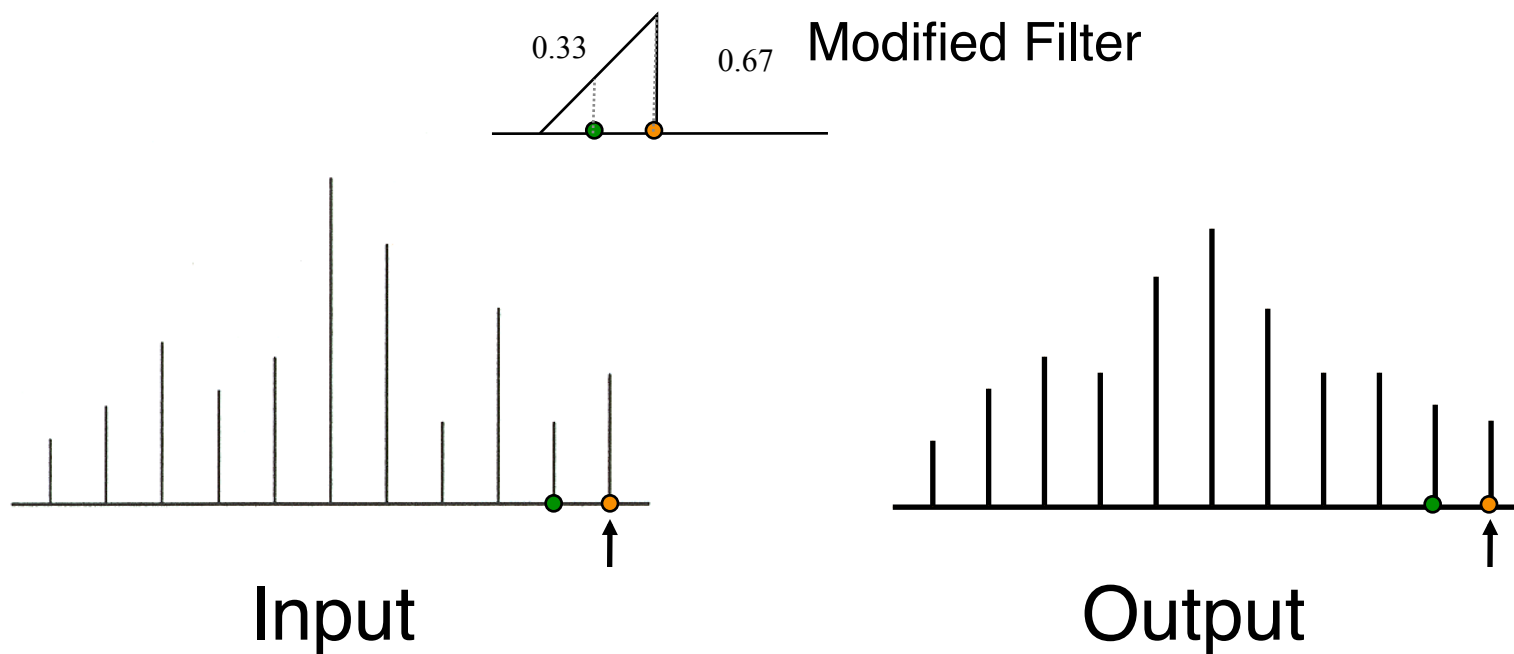


Output



Convolution with a Triangle Filter

Common option: normalize the filter





Convolution with a Gaussian Filter

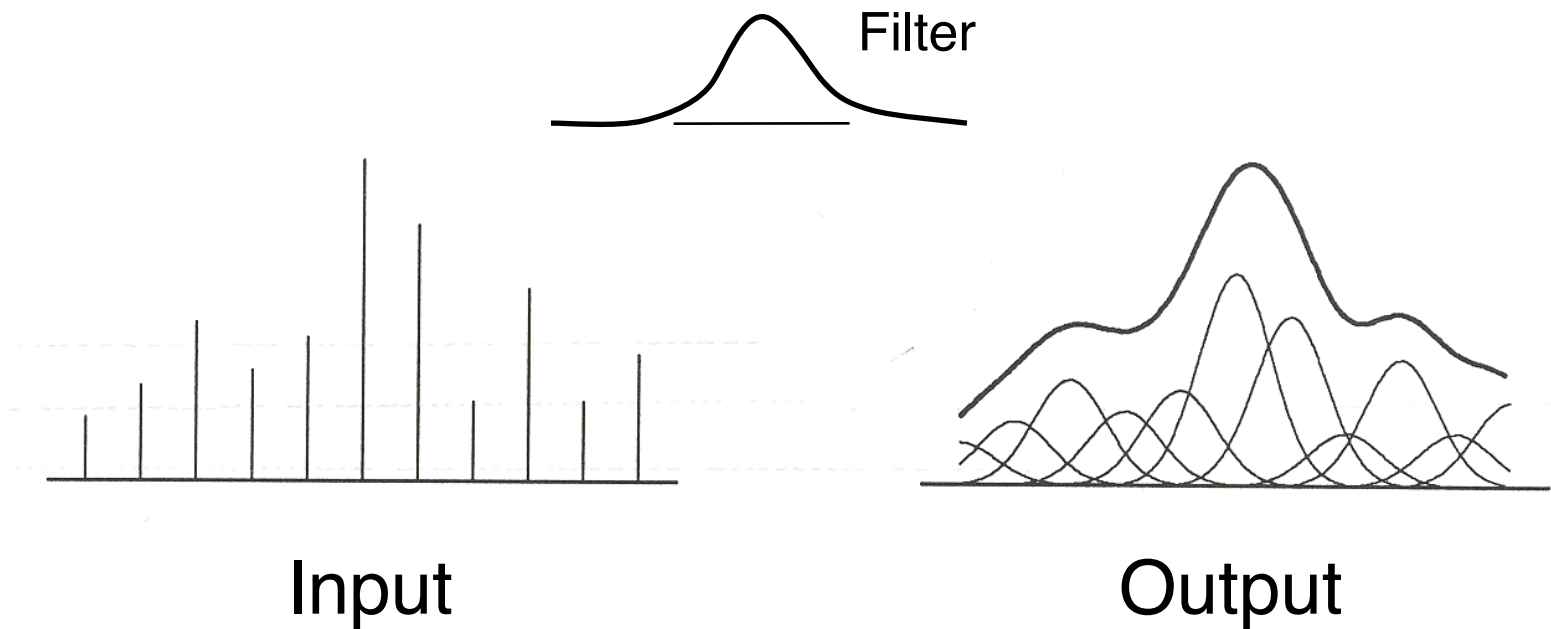



Figure 2.4 Wolberg



Linear Filtering

2D Convolution


- o Each output pixel is a linear combination of input pixels in neighborhood with weights prescribed by a filter



Input Image

$$\otimes \begin{array}{|c|} \hline \cdot \\ \cdot \\ \cdot \\ \hline \end{array} =$$

Filter




Output Image



Linear Filtering

2D Convolution


- o Each output pixel is a linear combination of input pixels in neighborhood with weights prescribed by a filter



Input Image



Filter




Output Image



Linear Filtering

2D Convolution


- o Each output pixel is a linear combination of input pixels in neighborhood with weights prescribed by a filter



Input Image



Filter




Output Image



Linear Filtering


2D Convolution

- o Each output pixel is a linear combination of input pixels in neighborhood with weights prescribed by a filter



Input Image

$$\otimes \begin{matrix} \square \\ \text{Filter} \end{matrix} =$$




Output Image



Linear Filtering

2D Convolution


- o Each output pixel is a linear combination of input pixels in neighborhood with weights prescribed by a filter



Input Image

$$\otimes \begin{array}{|c|} \hline \cdot \\ \hline \cdot \\ \hline \cdot \\ \hline \end{array} =$$

Filter



Output Image



Blur

Convolve with a filter whose entries sum to one

- o Each pixel becomes a weighted average of its neighbors



Original



Blur


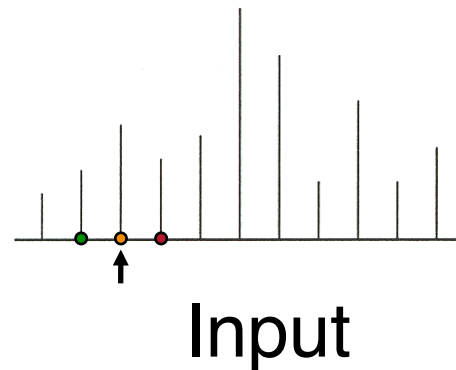



$$\text{Filter} = \begin{bmatrix} 1/16 & 2/16 & 1/16 \\ 2/16 & 4/16 & 2/16 \\ 1/16 & 2/16 & 1/16 \end{bmatrix}$$



Separable Filters

Remember our 1-D convolution...?



Separate X & Y dimensions:

- o Apply 1-D convolution across every row of image.
- o Then repeat for every column of the image.
- o What is the impact on performance?



Edge Detection

Convolve with a filter that finds differences between neighbor pixels



Original



Detect edges

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



Sharpen

Sum detected edges with original image



Original



Sharpened

$$\text{Filter} = \begin{bmatrix} -1 & -1 & -1 \\ -1 & +9 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$



Emboss

Convolve with a filter that highlights gradients in particular directions



Original



Embossed

$$\text{Filter} = \begin{bmatrix} -1 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$



Non-Linear Filtering


Each output pixel is a non-linear function of input pixels in neighborhood (filter depends on input)



Original



Paint



Stained Glass



Digital Image Processing


- Changing intensity/color
 - Linear: scale, offset, etc.
 - Nonlinear: gamma, saturation, etc.
 - Add random noise
- Filtering over neighborhoods
 - Blur
 - Detect edges
 - Sharpen
 - Emboss
 - Median
- Moving image locations
 - Scale
 - Rotate
 - Warp
- Combining images
 - Composite
 - Morph
- Quantization
- Spatial / intensity tradeoff
 - Dithering



Quantization

Reduce intensity resolution


- o Frame buffers have limited number of bits per pixel
- o Physical devices have limited dynamic range






Uniform Quantization

$P(x, y) = \text{round}(I(x, y))$
where $\text{round}()$ chooses nearest value that can be represented.



$I(x,y)$



$P(x,y)$
(2 bits per pixel)



Uniform Quantization

Images with decreasing bits per pixel:



8 bits



4 bits



2 bits



1 bit

Notice contouring.



Reducing Effects of Quantization

- Intensity resolution / spatial resolution tradeoff
- Dithering
 - Random dither
 - Ordered dither
 - Error diffusion dither
- Halftoning
 - Classical halftoning



Dithering

Distribute errors among pixels


- o Exploit spatial integration in our eye
- o Display greater range of perceptible intensities



Original
(8 bits)



Uniform
Quantization
(1 bit)





Floyd-Steinberg
Dither
(1 bit)



Random Dither

Randomize quantization errors

- o Errors appear as noise



$$P(x, y) = \text{round}(I(x, y) + \text{noise}(x,y))$$




Random Dither



Original
(8 bits)



Uniform
Quantization
(1 bit)



Random
Dither
(1 bit)



Ordered Dither

Pseudo-random quantization errors

o Matrix stores pattern of thresholds

$$i = x \bmod n$$

$$j = y \bmod n$$

$$e = I(x,y) - \text{trunc}(I(x,y))$$

$$\text{threshold} = (D(i,j)+1)/(n^2+1)$$


if ($e > \text{threshold}$)

$$P(x,y) = \text{ceil}(I(x,y))$$

else

$$P(x,y) = \text{floor}(I(x,y))$$

$$D_2 = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}$$





Ordered Dither

Bayer's ordered dither matrices

$$D_n = \begin{bmatrix} 4D_{n/2} + D_2(1,1)U_{n/2} & 4D_{n/2} + D_2(1,2)U_{n/2} \\ 4D_{n/2} + D_2(2,1)U_{n/2} & 4D_{n/2} + D_2(2,2)U_{n/2} \end{bmatrix}$$

$$D_2 = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}$$

$$D_4 = \begin{array}{cc|cc} 15 & 7 & 13 & 5 \\ 3 & 11 & 1 & 9 \\ \hline 12 & 4 & 14 & 6 \\ 0 & 8 & 2 & 10 \end{array}$$




Ordered Dither



Original
(8 bits)



Random
Dither
(1 bit)




Ordered
Dither
(1 bit)



Error Diffusion Dither

Spread quantization error over neighbor pixels

- o Error dispersed to pixels right and below
- o Floyd-Steinberg weights:



$$\frac{3}{16} + \frac{5}{16} + \frac{1}{16} + \frac{7}{16} = 1.0$$

Figure 14.42 from H&B




Error Diffusion Dither




Original
(8 bits)



Random
Dither
(1 bit)




Ordered
Dither
(1 bit)





Floyd-Steinberg
Dither
(1 bit)



Classical Halftoning



From *Town Topics*, Princeton






Classical Halftoning

Use ink dots of varying size to represent intensities

- o Area of dots proportional to intensity in image
- o Digital halftoning uses matrices (like ordered dither)



$I(x,y)$



$P(x,y)$

Summary



- Image filtering
 - Compute new values for image pixels based on function of old values
- Halftoning and dithering
 - Reduce visual artifacts due to quantization
 - Distribute errors among pixels
 - » Exploit spatial integration in our eye



Next Time...

- Changing intensity/color
 - Linear: scale, offset, etc.
 - Nonlinear: gamma, saturation, etc.
 - Add random noise
- Filtering over neighborhoods
 - Blur
 - Detect edges
 - Sharpen
 - Emboss
 - Median
- Moving image locations
 - Scale
 - Rotate
 - Warp
- Combining images
 - Composite
 - Morph
- Quantization
- Spatial / intensity tradeoff
 - Dithering