

# Introducing Assignment 1: Image Processing

---

COS 426: Computer Graphics (Spring 2019)

Jiaqi Su & Carlo Rosati

# GUI

COS426 Assignment 1

## Image Processing — Interactive Mode

Switch to: [Writeup](#)

Student Name <NetID>



- Push Image
- Batch Mode
- Animation
- MorphLines
- SetPixels
- ▾ Luminance
  - Brightness
  - Contrast
  - Gamma
  - Vignette
  - Histogram
- Color
- Filters
- Dithering
- Resampling
- Composite
- Misc
- Close Controls

- ▾ History
  - ▾ 1: Push Image
    - image name
    - Delete Below
  - ▾ 2: Brightness
    - brightness
    - Delete
- Close Controls

# GUI

---

- Useful functions

- Push Image
- Animation: generate gif animation using (min, step, max).
- MorphLines: specify line correspondences for morphing
- BatchMode: fix current parameter settings

# GUI

- Features to implement
  - SetPixels: set pixels to certain colors (A0)
  - Luminance: change pixel luminance
  - Color: remap pixel colors
  - Filter: convolution/box filter
  - Dithering: reduce visual artifacts due to quantization  $\approx$  cheat our eyes
  - Resampling: interpolate pixel colors
  - Composite: blending two images
  - Misc

# A few reminders...

---

- Don't try to exactly replicate example images.
- Choose parameters which give you best results.
- Have fun!

# Changing Contrast

- GIMP formula
  - $\text{value} = (\text{value} - 0.5) * (\tan ((\text{contrast} + 1) * \text{PI}/4)) + 0.5;$
- Notes:
  - When contrast=1,  $\tan(\text{PI}/2)$  is infinite. Using `Math.PI` can avoid this issue.
  - Do `pixel.clamp()` after computing the value.
  - Apply to each channel separately.




# Gamma correction

- $R = R^{\text{gamma}}$
- $G = G^{\text{gamma}}$
- $B = B^{\text{gamma}}$
- R,G,B are typically in  $[0, 1]$  (default in the code base)
- argument of `gammaFilter()` is  $\log(\text{gamma})$



# Vignette

- Pixels within innerR remain unchanged
- Pixels outside outerR are black
- Pixels between innerR and outerR should be multiplied with a value in [0, 1]:
  - Multiplier =  $1 - (R - \text{innerR}) / (\text{outerR} - \text{innerR})$
  - $R = \text{sqrt}(x^2 + y^2) / \text{halfdiag}$





# Histogram Equalization

Transform an image so that it has flat histogram of luminance values.



Before



After

# Histogram Matching

Transform an image so that it has same histogram of luminance values as reference image.




*reference image: town*




*reference image: flower*

# Histogram Equalization/Matching




pdf



cdf

# Histogram Equalization/Matching


- Image:  $x$
- Number of gray levels:  $L$
- $pdf(i) = \frac{n_i}{n}$   $n_i$  = number of pixels of the  $i$ -th gray level
- $cdf(j) = \sum_{i=0}^j pdf(i)$
- Target cdf:
  - Equalization:
    - $cdf_{ref}(i) = \frac{i}{L-1}$
  - Matching:
    - cdf of the reference image



(source:[http://paulbourke.net/texture\\_colour/equalisation/](http://paulbourke.net/texture_colour/equalisation/))

# Histogram Equalization/Matching

- Target cdf:
  - Equalization:
    - $cdf_{ref}(i) = \frac{i}{L-1}$
  - Matching:
    - cdf of the reference image
- Implementation
  - Equalization
    - $x' = cdf(x) * (L - 1) / (L - 1)$
  - Matching
    - $x' = \underset{i}{\operatorname{argmin}} |cdf(x) - cdf_{ref}(i)|$
    - Convert back to gray level:  $x' = \frac{x'}{L-1}$



# Saturation

- $\text{pixel} = \text{pixel} + (\text{pixel} - \text{gray}(\text{pixel})) * \text{ratio}$
- Do clamp()



# White balance

`whitebalance(image,  $rgb_w$ )`

$[L_w, M_w, S_w] = \text{rgb2lms}(rgb_w)$

for each pixel  $x$  in image

$[L, M, S] = \text{rgb2lms}(\text{image}(x))$

$L = L / L_w$

$M = M / M_w$


$S = S / S_w$

`image_out(x) = lms2rgb(L, M, S)`

- Hints:
  - Use `rgbToXyz()`, `xyzToLms()`, `lmsToXyz()`, `xyzToRgb()`
  - Do `clamp()`

# Convolution (Gaussian/Sharpen/Edge)

w1 w2 w3  
w4 w5 w6  
w7 w8 w9





# Convolution (Gaussian/Sharpen/Edge)

---

- Weights can be normalized depending on the application
- Edges? (not required)
  - Mirror boundary
  - Zero padding
  - Use part of the kernel only

# Gaussian filter


- Create a new image to work on
- Weights should be normalized, so that they sum to 1.
- Formula: 
$$G(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}$$
  - $x$  = distance to the center of the kernel
- Speed up:
  - First apply a 1D Gaussian kernel vertically and then a 1D Gaussian kernel horizontally

# Edge

- Kernel:

-1	-1	-1
-1	8	-1
-1	-1	-1

Inside boundary



At boundary


- Don't normalize weights
- Optional to invert the edge map for visualization:  
 $\text{pixel} = 1 - \text{pixel}$

# Sharpen

- Kernel:

-1	-1	-1
-1	<b>9</b>	-1
-1	-1	-1

Inside boundary



At boundary

- Don't normalize weights

# Edge Filter vs Sharpen Filter

<b>-1</b>	<b>-1</b>	<b>-1</b>
<b>-1</b>	<b>8</b>	<b>-1</b>
<b>-1</b>	<b>-1</b>	<b>-1</b>

Edge Filter

<b>-1</b>	<b>-1</b>	<b>-1</b>
<b>-1</b>	<b>9</b>	<b>-1</b>
<b>-1</b>	<b>-1</b>	<b>-1</b>

Sharpen Filter

$\text{Convolution}(\text{Image}, \text{Sharpen Filter}) = \text{Convolution}(\text{Image}, \text{Edge Filter}) + \text{Image}$

# Median

- Use a window (similar to convolution)
- Choose the median within the window
- Sorting: sort by RGB separately / sort by luminance



1



2



3



4



5

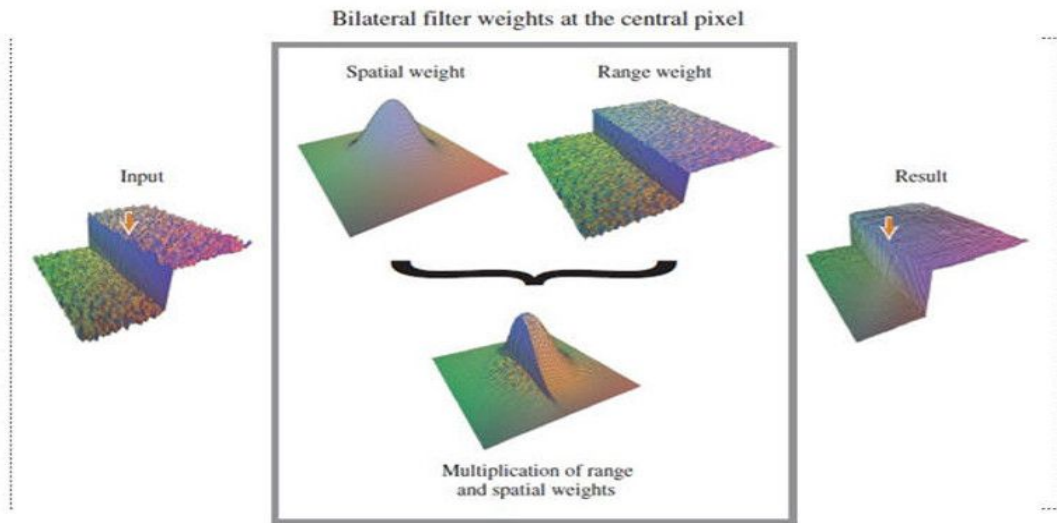
RGB Example

# Bilateral

- Combine Gaussian filtering in both spatial domain and color domain
- Weight formula of filter for pixel (i, j):

$$w(i, j, k, l) = e^{\left(-\frac{(i-k)^2 + (j-l)^2}{2\sigma_d^2} - \frac{\|I(i, j) - I(k, l)\|^2}{2\sigma_r^2}\right)}$$

- Similar color -> large weights, Different color -> smaller weights



# Quantization

- Quantize a pixel within  $[0, 1]$  using  $n$  bits
  - $\text{round}(p * (2^n - 1)) / (2^n - 1)$




$n=1$  example



# Random dithering



- Before quantization:
  - $p = p + (\text{random}() - 0.5)/(2^n - 1)$
  - $n$  is number of bits per channel

$n=1$  example



# Floyd-Steinberg error diffusion

- Loop over pixels line by line
  - Quantize pixel
  - Compute quantization error (the difference of the original pixel and the quantized pixel)
  - Spread quantization error over four unseen neighboring pixels with weights (see left figure below)
- Results look more natural



# Ordered dithering

## Pseudo code for n-bit case:

```
i = x mod m
j = y mod m
err = I(x, y) - floor_quantize(I(x, y))
threshold = (D(i, j) + 1) / (m^2 + 1)
if err > threshold
    P(x, y) = ceil_quantize(I(x, y))
else
    P(x, y) = floor_quantize(I(x, y))
```

- $\text{floor\_quantize}(p)$   
=  $\text{floor}(p * (2^{n-1})) / (2^{n-1})$
- $\text{ceil\_quantize}(p)$   
=  $\text{ceil}(p * (2^{n-1})) / (2^{n-1})$

$$m = 4, D = \begin{bmatrix} 15 & 7 & 13 & 5 \\ 3 & 11 & 1 & 9 \\ 12 & 4 & 14 & 6 \\ 0 & 8 & 2 & 10 \end{bmatrix}$$




n=1 example

# Resampling

- Bilinear interpolation

$$f(x, y) = \frac{1}{(x_2 - x_1)(y_2 - y_1)} (f(Q_{11})(x_2 - x)(y_2 - y) + f(Q_{21})(x - x_1)(y_2 - y) + f(Q_{12})(x_2 - x)(y - y_1) + f(Q_{22})(x - x_1)(y - y_1))$$

(from wikipedia)




# Resampling

- Gaussian interpolation

- Weights:


$$G(d, \sigma) = e^{-d^2 / (2\sigma^2)}$$

- Weights need to be normalized, so that sum up to 1



# Transformation (translate/scale/rotate/swirl)

- Inverse mapping



Inverse mapping guarantees that every pixel in the transformed image is filled!

# Transformation (translate/scale/rotate/swirl)

- To fill in a pixel in the target image, apply the inverse transform to the pixel location and look it up in the input image (with resampling technique) for pixel value.

- i.e. For translation of  $x' = x + tx$ ,  $y' = y + ty$ :

$$l'(x', y') = l(x' - tx, y' - ty)$$

- i.e. For scale of  $x' = x * sx$ ,  $y' = y * sy$ :

$$l'(x', y') = l(x' / sx, y' / sy)$$

# Composite

- $\text{output} = \text{alpha} * \text{foreground} + (1 - \text{alpha}) * \text{background}$
- alpha is the alpha channel foreground



backgroundImg



foregroundImg



foregroundImg(alpha channel)



Result



# Morph



- Basic concepts
  - transform the background image to the foreground image
  - $\alpha = 0$ : show background
  - $\alpha = 1$ : show foreground
  - $\alpha$  is the blending factor / timestamp
- General approach
  - specify correspondences (morphLines.html)
  - create an intermediate image with interpolated correspondences ( $\alpha$ )
  - warp the background image to the intermediate image
  - warp the foreground image to the intermediate image
  - blend using  $\alpha$

# Morph

```
GenerateAnimation(Image0, L0[...], Image1, L1[...])
begin
  foreach intermediate frame time t do
    for i = 1 to number of line pairs do
      L[i] = line t-th of the way from L0 [i] to L1 [i]
    end
    Warp0 = WarpImage(Image0, L0, L)
    Warp1 = WarpImage(Image1, L1, L)
    foreach pixel p in FinalImage do
      Result(p) = (1-t) Warp0 + t Warp1
    end
  end
end
```



# Warp Image

- $u = \frac{(X-P) \cdot (Q-P)}{\|Q-P\|^2}$
- $v = \frac{(X-P) \cdot \text{Perpendicular}(Q-P)}{\|Q-P\|}$  unit vector
- $X' = P' + u \cdot (Q' - P') + \frac{v \cdot \text{Perpendicular}(Q' - P')}{\|Q' - P'\|}$  unit vector
- *dist* = shortest distance from  $X$  to  $PQ$ 
  - $0 \leq u \leq 1$ :  $\text{dist} = |v|$
  - $u < 0$ :  $\text{dist} = \|X - P\|$
  - $u > 1$ :  $\text{dist} = \|X - Q\|$
- $\text{weight} = \left(\frac{\text{length}^p}{a + \text{dist}}\right)^b$ 
  - we use  $p = 0.5$ ,  $a = 0.01$ ,  $b = 2$



Contribution of this line segment  $PQ$  to the warping of  $X$ 's location

# Warp Image



Warped background or foreground (currently black)

Pixel source (background or foreground)

# Warp Image

For each pixel  $X$  in the destination

$$DSUM = (0,0)$$

$$weightsum = 0$$

For each line  $P_i Q_i$

calculate  $u, v$  based on  $P_i Q_i$

calculate  $X'_i$  based on  $u, v$  and  $P'_i Q'_i$

calculate displacement  $D_i = X'_i - X_i$  for this line

$dist$  = shortest distance from  $X$  to  $P_i Q_i$



$$weight = (length^p / (a + dist))^b$$

$$DSUM += D_i * weight$$


$$weightsum += weight$$

$$X' = X + DSUM / weightsum$$


$$destinationImage(X) = sourceImage(X')$$



# Interpolate Morph Lines




Background Image



Foreground Image

$$\text{current\_line}[i] = (1 - \alpha) * \text{background\_lines}[i] + \alpha * \text{foreground\_lines}[i]$$

# Blending



# Blending

$\alpha = 0.5$  (also the blending factor)



Background Image



Foreground Image



# Q&A

---