



COSC579: Image Processing Basics

Jeremy Bolton, PhD

Assistant Teaching Professor

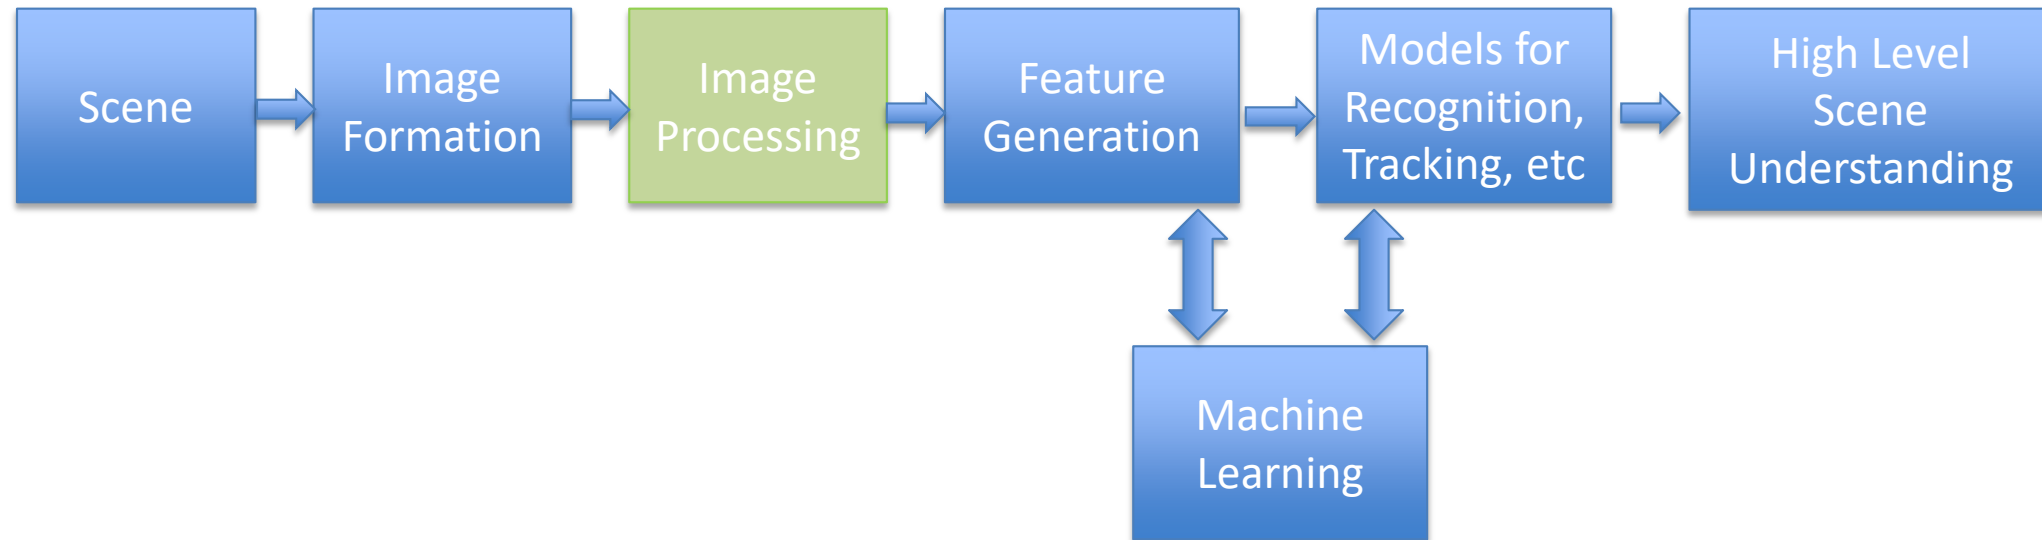
Outline

- I. Point Operators
 - I. Color Transforms
 - II. Normalization / Whitening
 - III. Histogram Normalization
- II. Linear Filters
 - I. Mathematical Formulation
 - I. Cross Correlation
 - II. Convolutions
 - II. Common Filters
 - I. Average
 - II. Gaussian
 - III. Laplacian

Reading

- Szeliski, Chapter 3.1-3.2
- Prince, Chapter 13.1.1 – 13.1.2



Our Progression Through Topics




What is an image?



What is an image?




Digital Image



The Eye

What is an image?

- A grid (matrix) of intensity values




=

255	255	255	255	255	255	255	255	255	255	255	255
255	255	255	255	255	255	255	255	255	255	255	255
255	255	255	20	0	255	255	255	255	255	255	255
255	255	255	75	75	75	255	255	255	255	255	255
255	255	75	95	95	75	255	255	255	255	255	255
255	255	96	127	145	175	255	255	255	255	255	255
255	255	127	145	175	175	175	255	255	255	255	255
255	255	127	145	200	200	175	175	95	255	255	255
255	255	127	145	200	200	175	175	95	47	255	255
255	255	127	145	145	175	127	127	95	47	255	255
255	255	74	127	127	127	95	95	95	47	255	255
255	255	255	74	74	74	74	74	74	255	255	255
255	255	255	255	255	255	255	255	255	255	255	255
255	255	255	255	255	255	255	255	255	255	255	255

(common to use one byte per value: 0 = black, 255 = white)

What is an image?

- We can think of a (grayscale) image as a **function**, f , from \mathbb{R}^2 to \mathbb{R} (or a 2D *signal*):
 - $f(x,y)$ gives the **intensity** at position (x,y)



- A **digital** image is a discrete (**sampled, quantized**) version of this function

Image transformations

- As with any function, we can apply operators to an image



$$g(x,y) = f(x,y) + 20$$




$$g(x,y) = f(-x,y)$$

- Operators
 - *Point Operators*
 - Histogram equalization
 - Color or intensity transformations
 - *Linear Filtering*: convolution and cross correlation

Pixel (intensity) transformation

- Transformation can be constant for each pixel

$$g(i,j) = a f(i,j) + b$$



$$g(i,j) = f(i,j) + 20$$

- Transform can vary from pixel-by-pixel

$$g(i,j) = a(i,j) f(i,j) + b(i,j)$$

Pixel Transforms based on image statistics

- Often an image is assumed to be a collection of observations of some random variable (a sample).
- Image statistics can provide a means for a data-based transformation (rather than some arbitrary scaling).



Histogram Equalization

- Problem: Image has undesirable histogram.
 - Distribution of intensity values.
- Solution: Map intensities to a desirable histogram using Histogram (PDF) and Cumulative Histogram (CDF)
- An often desired histogram is a Uniform distribution (all intensities are equal in frequency)

Histogram Equalization

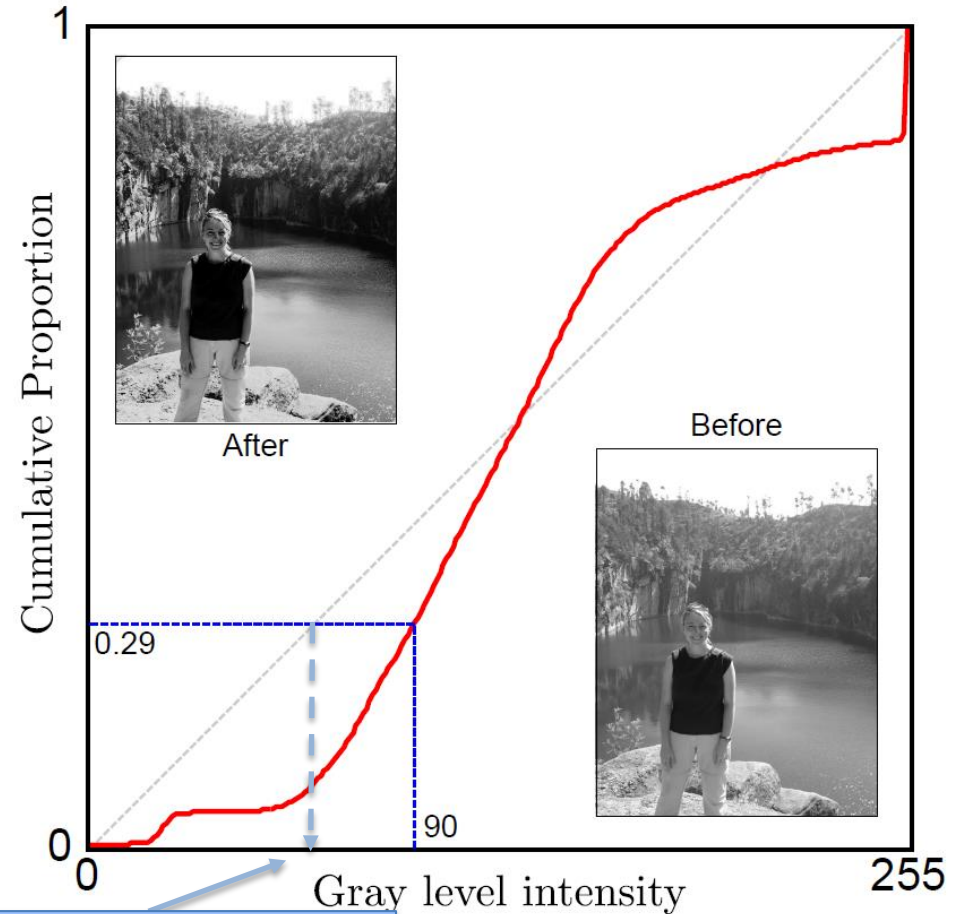
MATLAB

```
I = imread('tire.tif');  
J = histeq(I); %Enhance the contrast  
of an intensity image using histogram  
equalization  
imshowpair(I,J,'montage') %Display  
the original image and the adjusted  
image  
axis off  
figure;  
imhist(I,64)  
imhist(J,64)
```



Histogram Equalization

1. Compute Image Histogram (PDF)
2. Compute CDF
3. Compute CDF of desired histogram. (Assume Uniform)
4. Transform to desired CDF

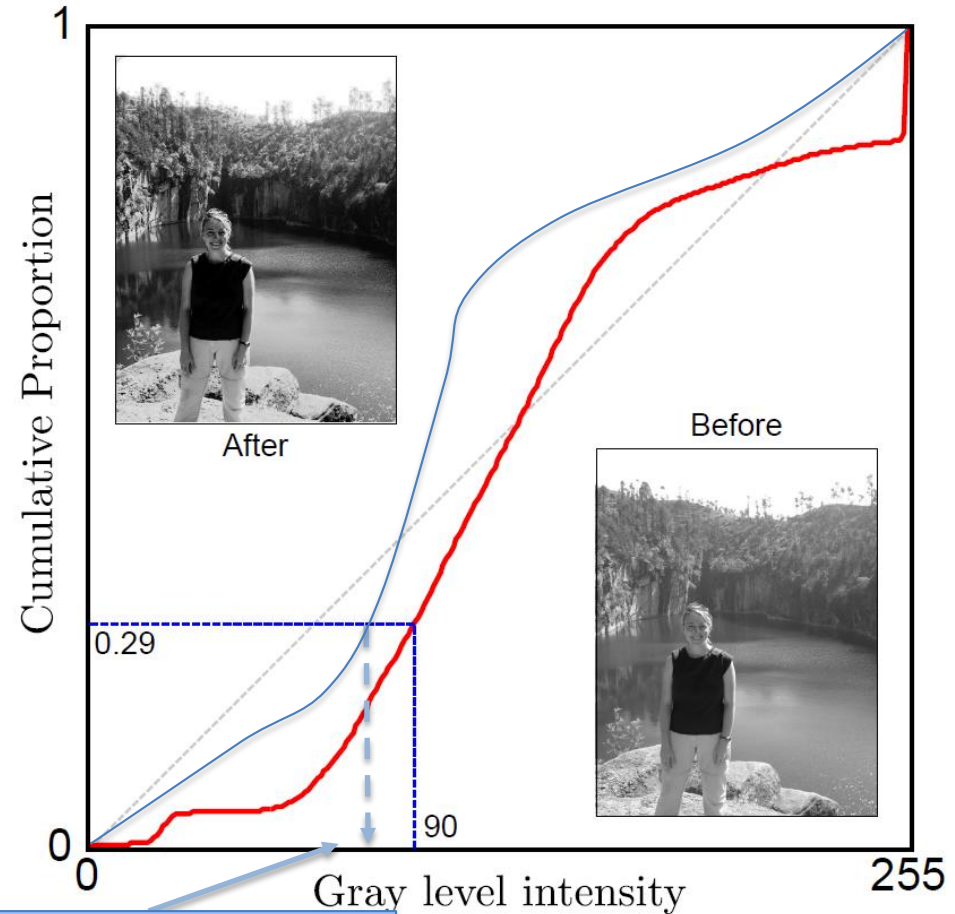


Map back to the original domain if you do not wish to change the range of intensities

Histogram Equalization

1. Compute Image Histogram (for all possible values k)
2. Compute CDF
3. Compute inverse of CDF of desired histogram.
 - Some CDFs are invertible (some are not)
 - If no closed form exists, then simply build a lookup table
4. Transform to desired CDF

$$f_{new}(x, y) = CDF_{desired}^{-1}(CDF(f(x, y)))$$



Map back to the original domain if you do not wish to change the range of intensities

Whitening

1. Compute Image Statistics

- Mean and standard deviation

$$\mu = \frac{\sum_{i=1}^I \sum_{j=1}^J p_{ij}}{IJ}$$
$$\sigma^2 = \frac{\sum_{i=1}^I \sum_{j=1}^J (p_{ij} - \mu)^2}{IJ}.$$

2. Apply transform to each pixel

- Use mean and standard deviation to “whiten” each pixel

$$x_{ij} = \frac{p_{ij} - \mu}{\sigma}.$$

Normalization Via Whitening

- Fix first and second moments to standard values
- Reduces contrast and constant additive luminance variations

Before



After

Normalization and Histogram Equalization

Make all of the moments the same by forcing the histogram of intensities to be the same



Before/ normalized/ Histogram Equalized

Linear Filter

- Basic Point Operators perform the same operation on each pixel
 - Not necessarily based on image data (local or global)
- Normalization / Whitening
 - Processing based on global (or local) image statistics
- Local Operators AKA (neighborhood operators)
 - Often implemented as **Linear Filters**
 - Local or neighborhood information is used in transformation

Image filtering

- Modify the pixels in an image based on some function of a local neighborhood of each pixel
 - Example: 3x3 neighborhood

10	5	3
4	5	1
1	1	7

Local image data

Some function




	7	

Modified image data

Linear filtering

- One simple version: linear filtering (cross-correlation, convolution)
 - Replace each pixel by a linear combination of its neighbors
- This linear combination is an binary operator which maps two functions to one function.
 - First operand is the input image
 - Second operand is called the “kernel” (or “mask”, “filter”)
 - The resulting function is the output image



Cross-correlation

Let F be the image, H be the kernel (of size $2k+1 \times 2k+1$), and G be the output image

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$

This is called a **cross-correlation** operation:

$$G = H \otimes F$$

Convolution


- Same as cross-correlation, except that the kernel is “flipped” (horizontally and vertically)

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i - u, j - v]$$


This is called a **convolution** operation: $G = H * F$

- Convolution is **commutative** and **associative**

Convolution



Mean filtering



Linear filters: examples



Original



0	0	0
0	1	0
0	0	0



Identical image

Linear filters: examples



Original



0	0	0
1	0	0
0	0	0



Shifted left
By 1 pixel

Linear filters: examples




Original



$\frac{1}{9}$

1	1	1
1	1	1
1	1	1



Blur (with a mean filter)

Linear filters: examples




Original

$$* \left(\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} \right) =$$




Sharpening filter
(accentuates edges)

Sharpening





before





after

Smoothing with box filter revisited







Gaussian Kernel




$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$


Gaussian filters



$\sigma = 1$ pixel



$\sigma = 5$ pixels




$\sigma = 10$ pixels




$\sigma = 30$ pixels

Sharpening




before




after

Sharpening revisited

- What does blurring take away?



Let's add it back:



Sharpen filter



Blurring (convolve with Gaussian)






Figure B.3 Image blurring. a) Original image. b) Result of convolving with a Gaussian filter (filter shown in bottom right of image). The image is slightly blurred. c-e) Convolving with a filter of increasing standard deviation causes the resulting image to be increasingly blurred.

Looking Ahead: Gradient Filters (Edge Detection)



Prewitt (vertical)

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$$




Prewitt (horizontal)

$$\begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix}$$




Laplacian


$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$



Laplacian of Gaussian

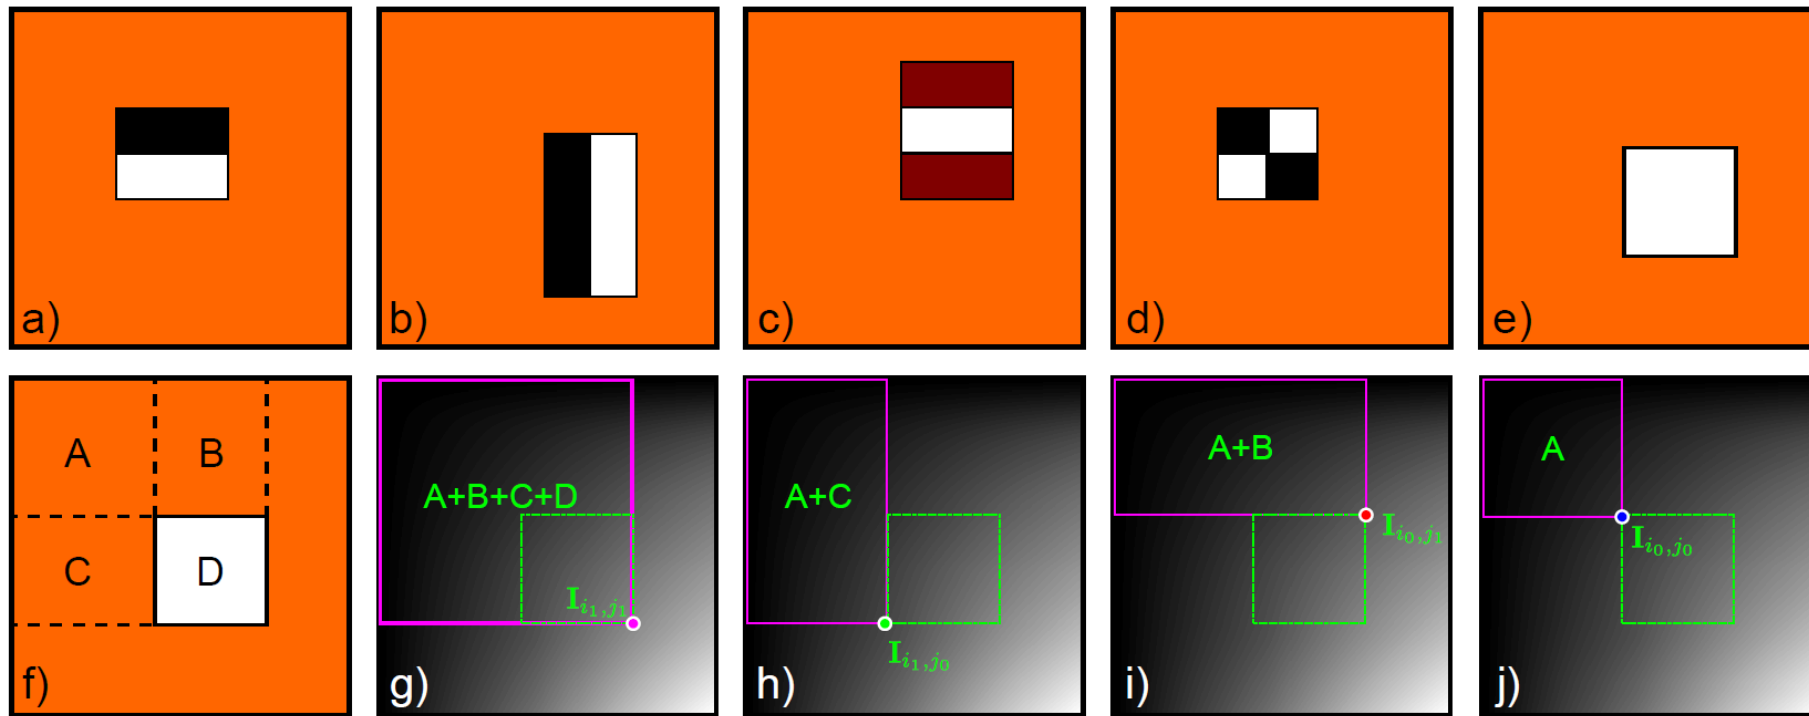


Difference of Gaussians





- Rule of thumb: big response when image matches filter

Haar Filters







Looking Ahead: Filtering for Texture (Textons)

- An attempt to characterize texture
- Replace each pixel with integer representing the texture 'type'



Computing Textons

Take a bank of filters and apply to lots of images



For new pixel, filter surrounding region with same bank, and assign to nearest cluster