


# Image Processing

# What is an image?


- We can think of an **image** as a function,  $f$ , from  $\mathbb{R}^2$  to  $\mathbb{R}$ :
  - $f(x, y)$  gives the **intensity** at position  $(x, y)$
  - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
    - $f: [a,b] \times [c,d] \rightarrow [0,1]$
- A color image is just three functions pasted together. We can write this as a “vector-valued” function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$



Image



Brightness values



$$I(x,y)$$



62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30

# What is a digital image?

- In computer vision we usually operate on **digital (discrete)** images:
  - **Sample** the 2D space on a regular grid
  - **Quantize** each sample (round to nearest integer)
- If our samples are  $\Delta$  apart, we can write this as:
- $f[i, j] = \text{Quantize}\{ f(i \Delta, j \Delta) \}$
- The image can now be represented as a matrix of integer values

$\xrightarrow{j}$

$i$   $\downarrow$

62	79	23	119	120	105	4	0
10	10	9	62	12	78	34	0
10	58	197	46	46	0	0	48
176	135	5	188	191	68	0	49
2	1	1	29	26	37	0	77
0	89	144	147	187	102	62	208
255	252	0	166	123	62	0	31
166	63	127	17	1	0	99	30



`im[1:10,1:10,0]`

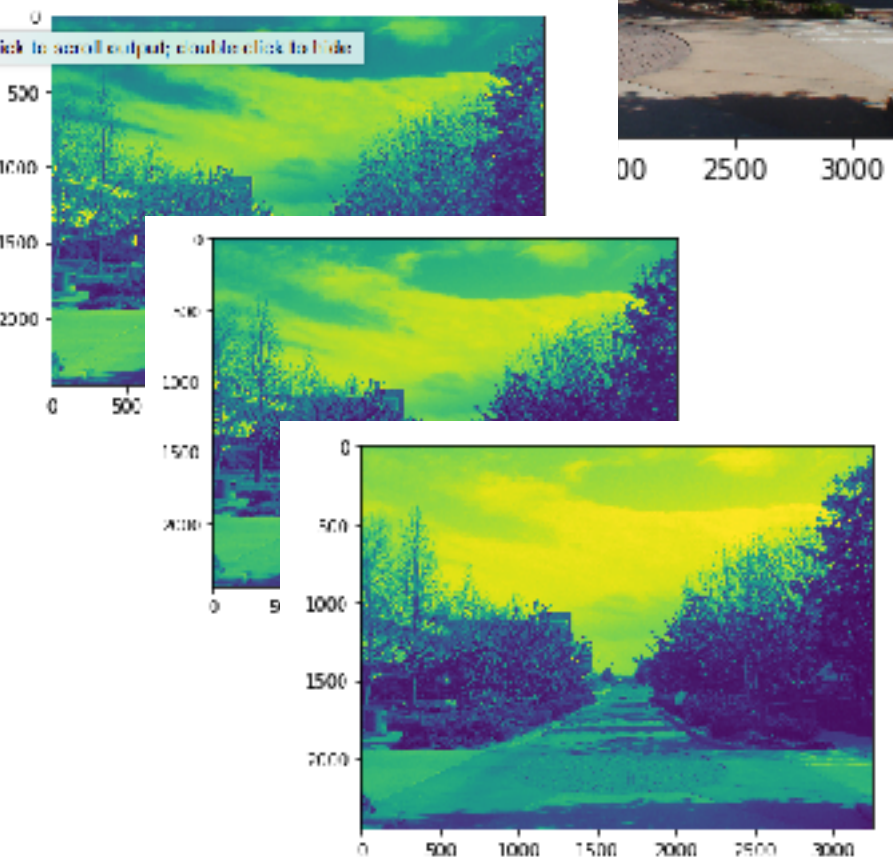
```
array([[88, 89, 90, 90, 91, 91, 92, 92, 91],
       [89, 90, 91, 91, 91, 92, 92, 92, 92],
       [90, 92, 92, 92, 91, 92, 92, 93, 93],
       [91, 93, 93, 92, 91, 92, 92, 93, 93],
       [91, 92, 92, 92, 92, 92, 93, 93, 94],
       [91, 91, 92, 92, 93, 93, 94, 95, 94],
       [92, 92, 92, 93, 94, 94, 96, 96, 96],
       [93, 93, 93, 93, 92, 92, 91, 94, 94],
       [94, 94, 94, 94, 93, 92, 92, 94, 94]], dtype=uint8)
```


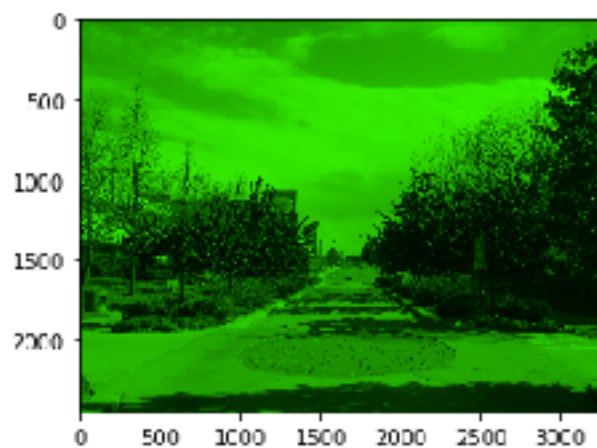
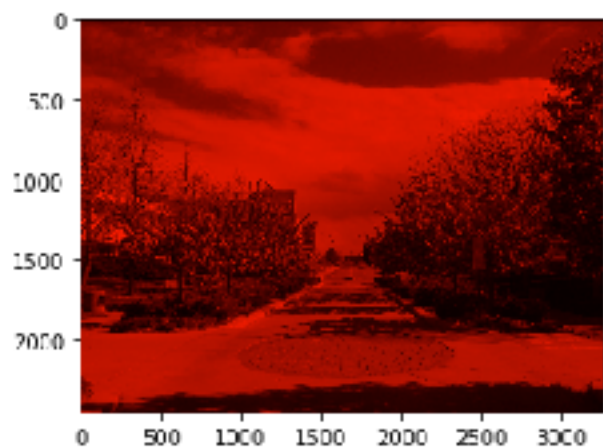

`im[1:10,1:10,1]`


```
array([[134, 135, 136, 136, 137, 137, 138, 138, 137],
       [136, 136, 138, 138, 137, 138, 138, 139, 139],
       [137, 138, 139, 138, 137, 138, 138, 139, 139],
       [137, 138, 138, 138, 138, 138, 139, 139, 139],
       [137, 137, 138, 138, 139, 139, 140, 140, 139],
       [138, 138, 139, 139, 140, 140, 141, 141, 141],
       [137, 137, 137, 137, 138, 139, 137, 140, 140],
       [138, 138, 138, 138, 139, 139, 138, 140, 140]], dtype=uint8)
```

`im[1:10,1:10,2]`

```
array([[194, 195, 196, 196, 197, 197, 198, 198, 196],
       [195, 196, 197, 197, 197, 197, 197, 197, 197],
       [195, 197, 197, 197, 196, 197, 197, 198, 198],
       [196, 198, 198, 197, 196, 197, 197, 198, 198],
       [196, 197, 197, 197, 197, 197, 198, 198, 198],
       [196, 196, 197, 197, 198, 198, 199, 199, 198],
       [197, 197, 197, 198, 199, 199, 200, 200, 200],
       [198, 198, 198, 198, 197, 197, 196, 199, 199],
       [199, 199, 199, 199, 198, 197, 197, 199, 199]], dtype=uint8)
```







```
array([[ 127.086,  128.086,  129.086,  129.086,  130.086,  130.086,
        [ 129.972,  129.972,  130.972,  130.972,  130.972,  130.972],
        [ 128.972,  130.972,  130.972,  130.972,  129.972,  130.972,
        [ 129.972,  131.972,  131.972,  130.972,  129.972,  130.972,
        [ 130.972,  131.972,  131.972],
        [ 129.972,  130.972,  130.972,  130.972,  130.972,  130.972,
        [ 129.972,  129.972,  130.972,  130.972,  131.972,  131.972,
        [ 130.972,  130.972,  130.972,  130.972,  131.972,  132.972,
        [ 130.972,  130.972,  130.972,  131.972,  132.972,  132.972,
        [ 134.271,  134.271,  134.271],
        [ 130.798,  130.798,  130.798,  130.798,  130.972,  130.972,
        [ 129.972,  132.972,  132.972],
        [ 131.798,  131.798,  131.798,  131.798,  131.972,  130.972,
        [ 130.972,  132.972,  132.972]])
```





# Human perception of Color

- Human eye is visually sensitive to electromagnetic spectrum between 400 - 700 nm.
- HSL - hue, saturation and luminance
- measured on a scale of 0-240
- Hue provides the color choice - arranged in a linear strip Red (240), violet (200), indigo(170), blue(128), green(85), yellow(42), orange(21)

# Human perception of Color

- Hue provides the color choice - arranged in a linear strip  
Red (240), violet (200), indigo(170), blue(128), green(85),  
yellow(42), orange(21)
- Saturation indicates pureness, or intensity or the amount  
of gray.
- luminance indicates lightness of the color. A value of 0  
indicates black, value of 240 indicates pure white.
- The relationship between grayscale reflectance of a  
surface and its RGB color equivalent is given by:

$$Y = 0.299 *R+ 0.587*G + 0.114*B$$

# Image Processing

- An **image processing** operation typically defines a new image  $g$  in terms of an existing image  $f$ .
- We can transform either the domain or the range of  $f$ .

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

- **Range transformation:**

What kinds of operations can this perform?

- Some operations preserve the range but **change the domain** of  $f$ :

$$g(x, y) = f(t_x(x, y), t_y(x, y))$$

What kinds of operations can this perform?

# Thresholding



One of the simplest operations we can perform on an image is *thresholding*.

For example, if we take the swan image and threshold it with a threshold of  $T$ , we make all pixels  $\geq T$  into 1, and all pixels  $< T$  into 0.


# Threshold $T=128$



# Examples



# A simple image and its histogram




# Definition of histogram

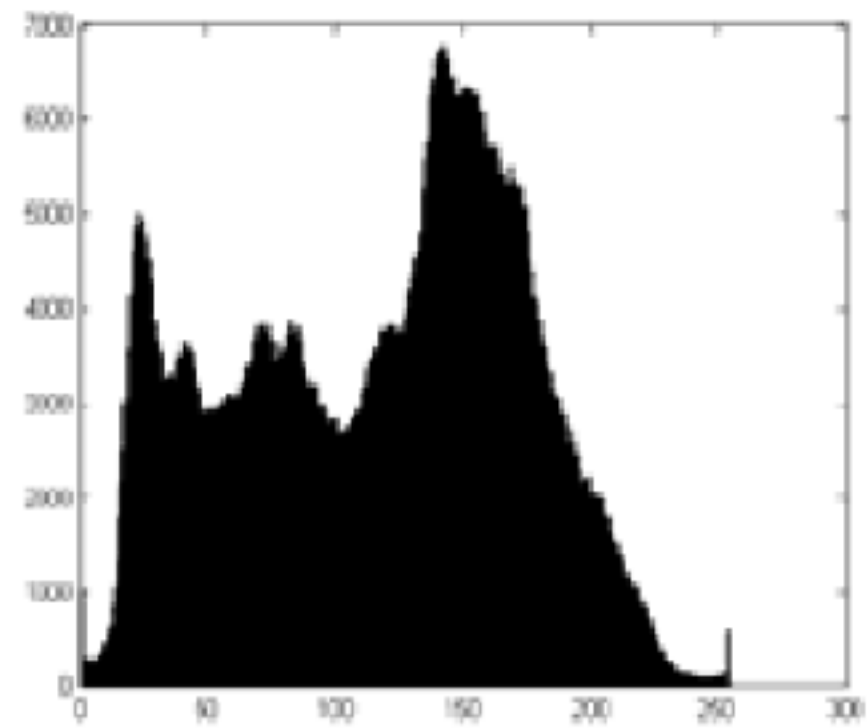
- To write this down, we might say that we have an image,  $I$ , in which the intensity at pixel with coordinates  $(x,y)$  is  $I(x,y)$ . We would write the histogram  $h$ , as  $h(i)$  indicating that intensity  $i$ , appears  $h(i)$  times in the image. If we let the expression  $(a=b)$  have the value 1 when  $a=b$ , and 0 otherwise, we can write for histogram  $h(i)$ :

$$h(i) = \sum_x \sum_y I(x,y) = i$$



# Examples





# Histograms allow image manipulation

- One reason to compute a histogram is that it allows us to manipulate an image by changing its histogram. We do this by creating a new image,  $J$ , in which:

$$J(x,y) = f(I(x,y))$$

- The trick is to choose an  $f$  that will generate a nice or useful image. Typically, we choose  $f$  to be monotonic. This means that: if  $u < v$  then  $f(u) < f(v)$ . Non-monotonic functions tend to make an image look truly different, while monotonic changes will be more subtle.

# Histogram Equalization

- The idea is to spread out the histogram so that it makes full use of the dynamic range of the image.
- For example, if an image is very dark, most of the intensities might lie in the range 0-50. By choosing  $f$  to spread out the intensity values, we can make fuller use of the available intensities, and make darker parts of an image easier to understand.
- If we choose  $f$  to make the histogram of the new image,  $J$ , as uniform as possible, we call this **histogram equalization**.

# How to do it

- **Cumulative Distribution Function (CDF).**

This encodes the fraction of pixels with an intensity that is equal to or less than a specific value.

If  $h$  is a histogram and  $C$  is a CDF, then  $h(i)$  indicates the number of pixels with intensity of  $i$ , while ,

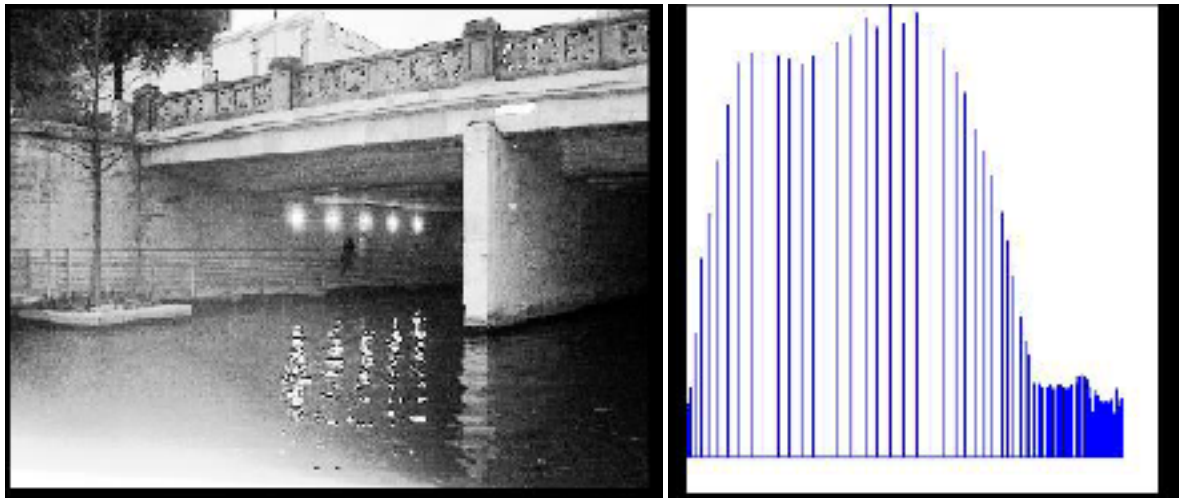

$$C(i) = \sum_{j \leq i} h(j) / N,$$

indicates the fraction of pixels with intensity less than or equal to  $i$ , assuming the image has  $N$  pixels.

# Actual construction

- $g$  is the histogram of  $J$ , and  $D$  its CDF.
- If there are  $k$  intensity levels in an image, then we want  $g=(N/k, N/k, \dots)$ , and  $D=(1/k, 2/k, 3/k, \dots)$ . This means that we want  $D(i) = i/k$ . Notice that  $C(i) = D(f(i))$ . That is, all the pixels in  $I$  that have an intensity less than or equal to  $i$  will have an intensity less than or equal to  $f(i)$  in  $J$  (since  $f$  is monotonic, if  $j < i$ ,  $f(j) < f(i)$ ).
- Putting these together, we have  $D(f(i))=f(i)/k=C(i)$ , so  $f(i) = kC(i)$ .

# Examples of histogram equalization



# 8 x 8 Image

52	55	61	59	70	61	76	61
62	59	55	104	94	85	59	71
63	65	66	113	144	104	63	72
64	70	70	126	154	109	71	69
67	73	68	106	122	88	68	68
68	79	60	79	77	66	58	75
69	85	64	58	55	61	65	83
70	87	69	68	65	73	78	90



# 8 x 8 Image

52	55	61	59	70	61	76	61
62	59	55	104	94	85	59	71
63	65	66	113	144	104	63	72
64	70	70	126	154	109	71	69
67	73	68	106	122	88	68	68
68	79	60					
69	85	64					
70	87	69					

Value	Count	Value	Count	Value	Count	Value	Count	Value	Count
52	1	64	2	72	1	85	2	113	1
55	3	65	3	73	2	87	1	122	1
58	2	66	2	75	1	88	1	126	1
59	3	67	1	76	1	90	1	144	1
60	1	68	5	77	1	94	1	154	1
61	4	69	3	78	1	104	2		
62	1	70	4	79	2	106	1		
63	2	71	2	83	1	109	1		

# 8 x 8 Image

Value	Count	Value	Count	Value	Count
52	1	64	2	72	1
55	3	65	3	73	2
58	2	66	2	75	1
59	3	67	1	76	1
60	1	68	5	77	1
61	4	69	3	78	1
62	1	70	4	79	2
63	2	71	2	83	1

v, Pixel Intensity	cdf(v)	h(v), Equalized v
52	1	0
55	4	12
58	6	20
59	9	32
60	10	36
61	14	53
62	15	57
63	17	65
64	19	73
65	22	85
66	24	93

$$h(v) = \text{round}\left(\frac{cdf(v) - cdf_{min}}{(M \times N) - cdf_{min}} \times (L - 1)\right)$$

$M$  – width

$M$  – height

$L$  – Number of gray levels

# Comparing histograms

- SSD

Let  $h$  and  $g$  be two histograms.

$$\|h - g\| = \sum_{i=1}^N (h(i) - g(i))^2$$

- Cosine

$$\cos(h, g) = \frac{\langle h, g \rangle}{\|h\| \|g\|}$$

# Treating histograms as probability distributions

- Chi-Squared

$$\chi^2(h_i, h_j) = \frac{1}{2} \sum_{m=1}^K \frac{[h_i(m) - h_j(m)]^2}{h_i(m) + h_j(m)}$$

- Smoothing probability distributions

- Use bigger buckets.
- Add a constant value (eg., 1) to every bucket.
- Gaussian smoothing