

CS 331

Computer Vision

09 / 28 / 2026

Instructor: Michael Eckmann

Today's Topics

- Questions / Comments?
- My answers to the in-class exercises from last time
- Histograms
- Distance measures to determine closeness of vectors


histograms

- A histogram of a greyscale (aka intensity) image is a function
 - whose domain is a set of grey values and
 - And whose value at a grey value is the number of pixels in the image with that grey value
- The domain could be the set of all grey values, where each grey value is a bin. Or it could be a smaller number of bins, where each contiguous set of grey values is a bin.
- e.g. for an 8 bit greyscale image, the range of values is 0 to 255
 - we could make 256 bins, where each bin is one grey value or
 - if we want 32 bins, then each bin will be associated with $256/32 = 8$ grey values. 0-7 in bin0, 8-15 in bin1, ..., 248-255 in bin31.
 - our histogram will be a function whose domain is the bins and whose value is the number of pixels in the image with any of the grey values corresponding to that bin

histograms

- A histogram then is a mapping from the domains (the bins) to the range (the number of pixels corresponding to a particular bin)

histograms



Michael Eckmann - Skidmore
College - CS 331 - Fall 2026

histograms

- Histograms can be used to describe the *global* intensity (or color) content of an image.
- Note well --- a histogram ignores where pixels are in the image
- very different images can have the exact same histogram
- A cumulative histogram is such that
 - The value in a bin is the sum of all the values in the bins $\leq$ that bin in the original histogram
 - e.g. if in the original histogram,
 - bin 0 has 25 pixels and bin 1 has 50 and bin2 has 35
 - The cumulative histogram would have
 - bin 0 has 25, bin1 has 75 ($= 25+50$) and bin2 has 110 ($=25+50+35$)

histograms

- Histograms have many uses in image processing.
- One use is for contrast enhancement (histogram equalization)
- In a nutshell histogram equalization examines the intensity histogram and makes the darkest pixels black and the brightest pixels white and spreads out the intermediate intensities in a controlled way.
- It creates a mapping of grey values of the original image to grey values in the resulting image so that the cumulative histogram of the resulting image will be as close to a straight line as possible.
- There are many details that I'm not interested in covering for this vision course that I do cover in image processing course.
- Let me show an example

histograms

- For color histograms, since we have multiple channels, we need a way to combine ranges of all the channels into a bin. Consider these 16 bins for a color image histogram:
- Bin0 those with R in 0-127, G in 0-63, B in 0-127
- Bin1 those with R in 0-127, G in 0-63, B in 128-255
- Bin2 those with R in 0-127, G in 64-127, B in 0-127
- Bin3 those with R in 0-127, G in 64-127, B in 128-255
- Bin4 those with R in 0-127, G in 128-191, B in 0-127
- Bin5 those with R in 0-127, G in 128-191, B in 128-255
- Bin6 those with R in 0-127, G in 192-255, B in 0-127
- Bin7 those with R in 0-127, G in 192-255, B in 128-255
- Bin8 those with R in 128-255, G in 0-63, B in 0-127
- Bin9 those with R in 128-255, G in 0-63, B in 128-255
- Bin10 those with R in 128-255, G in 64-127, B in 0-127
- Bin11 those with R in 128-255, G in 64-127, B in 128-255
- Bin12 those with R in 128-255, G in 128-191, B in 0-127
- Bin13 those with R in 128-255, G in 128-191, B in 128-255
- Bin14 those with R in 128-255, G in 192-255, B in 0-127
- Bin15 those with R in 128-255, G in 192-255, B in 128-255

Determining similarity

- Many problems in vision require us to compare how similar/dissimilar images are to each other.
- Trying to compare the raw image pixels to each other to determine how similar or dissimilar images are is not going to yield very good results.
- Instead we represent the images in some way. Then we compare how similar/dissimilar their *representations* are.

Determining similarity

- We can represent objects with those region properties (e.g. a vector of values of those properties) or as a histogram or in countless other ways or combinations of them.
- What we'll have is a vector of values as the representation of an image and we will need a way to compare the vectors to determine how close they are to each other (or how far apart they are).

Distance measures

- There are many ways to compute distance between two multidimensional vectors
- the L1 distance (aka Manhattan or Taxicab) is the sum of the absolute values of the differences of each dimension
- the L2 distance (aka Euclidean) is the square root of the sum of squares of the differences of each dimension
- These are easy to see in 2 dimensions.
- Say we have $p = (p_1, p_2)$ and $q = (q_1, q_2)$ then the L1 distance between them is: $\text{abs}(p_1 - q_1) + \text{abs}(p_2 - q_2)$
- The L2 distance between them is: $\text{sqrt}((p_1 - q_1)^2 + (p_2 - q_2)^2)$

Distance measures

- In general for n dimensions,
- if $p=(p_1, p_2, \dots, p_n)$ and $q = (q_1, q_2, \dots, q_n)$
- the L1 distance between them is:

$$\text{abs}(p_1-q_1) + \text{abs}(p_2-q_2) + \dots + \text{abs}(p_n-q_n)$$

- The L2 distance between them is:

$$\text{sqrt}((p_1-q_1)^2 + (p_2-q_2)^2 + \dots + (p_n-q_n)^2)$$

- These distance measures (as well as many others) could be used to compute the distance between any n-dimensional sets of values, not only histograms. For example the region measures (area, perimeter, etc.) can each be dimensions of a vector and then one can take the L1 or L2 distance between vectors to determine how similar or dissimilar the regions are.

Distance measures

- Let's do an example where the vectors have 5 values (5 dimensions)
- Example: $p = (5,3,2,6,9)$ and $q=(4,4,4,7,6)$ and $r = (10,6,4,5,0)$
- Let's see whether p is closer to q or r using L1
- Example calculation:
- L1 distance between p and $q = |5-4| + |3-4| + |2-4| + |6-7| + |9-6| = 1+1+2+1+3 = 8$
- L1 distance between p and $r = |5-10| + |3-6| + |2-4| + |6-5| + |9-0| = 5+3+2+1+9 = 20$
- So p is closer to q than r when considering L1 distance
- In other words we determined that p is more similar to q than it is to r
- Let's do the same with L2 distance and see how they compare