CS 331
Computer Vision

09 / 19 / 2024

Instructor: Michael Eckmann

Today's Topics

- Questions / Comments?
- Region properties
- Histograms
- Distance measures to determine closeness of vectors

histograms

- Reading on recent topics:
- section 3.3.3 Binary Image Processing, Distance Transforms and Connected Components
  - Pages 138 – top of 142
- section 3.1.4 Histogram Equalization which is the end of page 115 to the top of page 119

region properties of binary images

- Consider a region to be a set of connected ON pixels
- **area** of region is simply the count of the number of pixels in the region
- **centroid** = (row_c, col_c) where, row_c = mean row of all pixels in region and col_c = mean column of all pixels in region
- **perimeter** of a region without holes is the set of all pixels that are in the region AND have a neighbor outside of the region
  - neighbor can be defined as 4-connected or 8-connected
- **length of perimeter** can be computed by starting at a pixel and travelling around the whole perimeter and arriving back at the starting pixel
  - add 1 if the adjacent pixels are 4 connected
  - add sqrt(2) if the adjacent pixels are 8 connected but not 4 connected
  - Note: if there are n pixels in the perimeter, then there are n pairs of adjacent pixels, hence, n numbers added (each of which is either 1 or sqrt(2).

region properties of binary images

- “**Circularity**”
  - in some texts this is referred to as **compactness** and is defined as
  - C_1 = (length of perimeter)^2 / area
  - this is a dimensionless number and in the analog world is minimized by a disk (Ballard & Brown 1982)
  - however, for digital shapes, that measure is minimized when the shape is a diamond (when we use the length of perimeter definition as before)
  - a different circularity measure (Haralick 1974) which is similar for both digital and analog shapes and monotonically increases as the shape becomes more circular is
    - C_2 = mean of radial distances / standard dev. of radial distances
    - the radial distance of a pixel on the perimeter is the distance between that pixel's center and the centroid of the region

region properties of binary images

- **"Circularity"**
  - C_2 = mean of radial distances / standard dev. of radial distances
  - the radial distance of a pixel on the perimeter is the distance between that pixel's center and the centroid of the region
  - the mean and standard deviation are statistical measures
    - the mean is what we mostly think of as the average
    - the standard deviation is a measure of how spread out the group is in values
    - lower sd => more compact group, higher sd => more spread out
    - standard deviation is the square root of the variance
    - whatever units the original data is in, the variance is in those units squared, if you want a measure of the spread in the same units, then use standard deviation

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 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

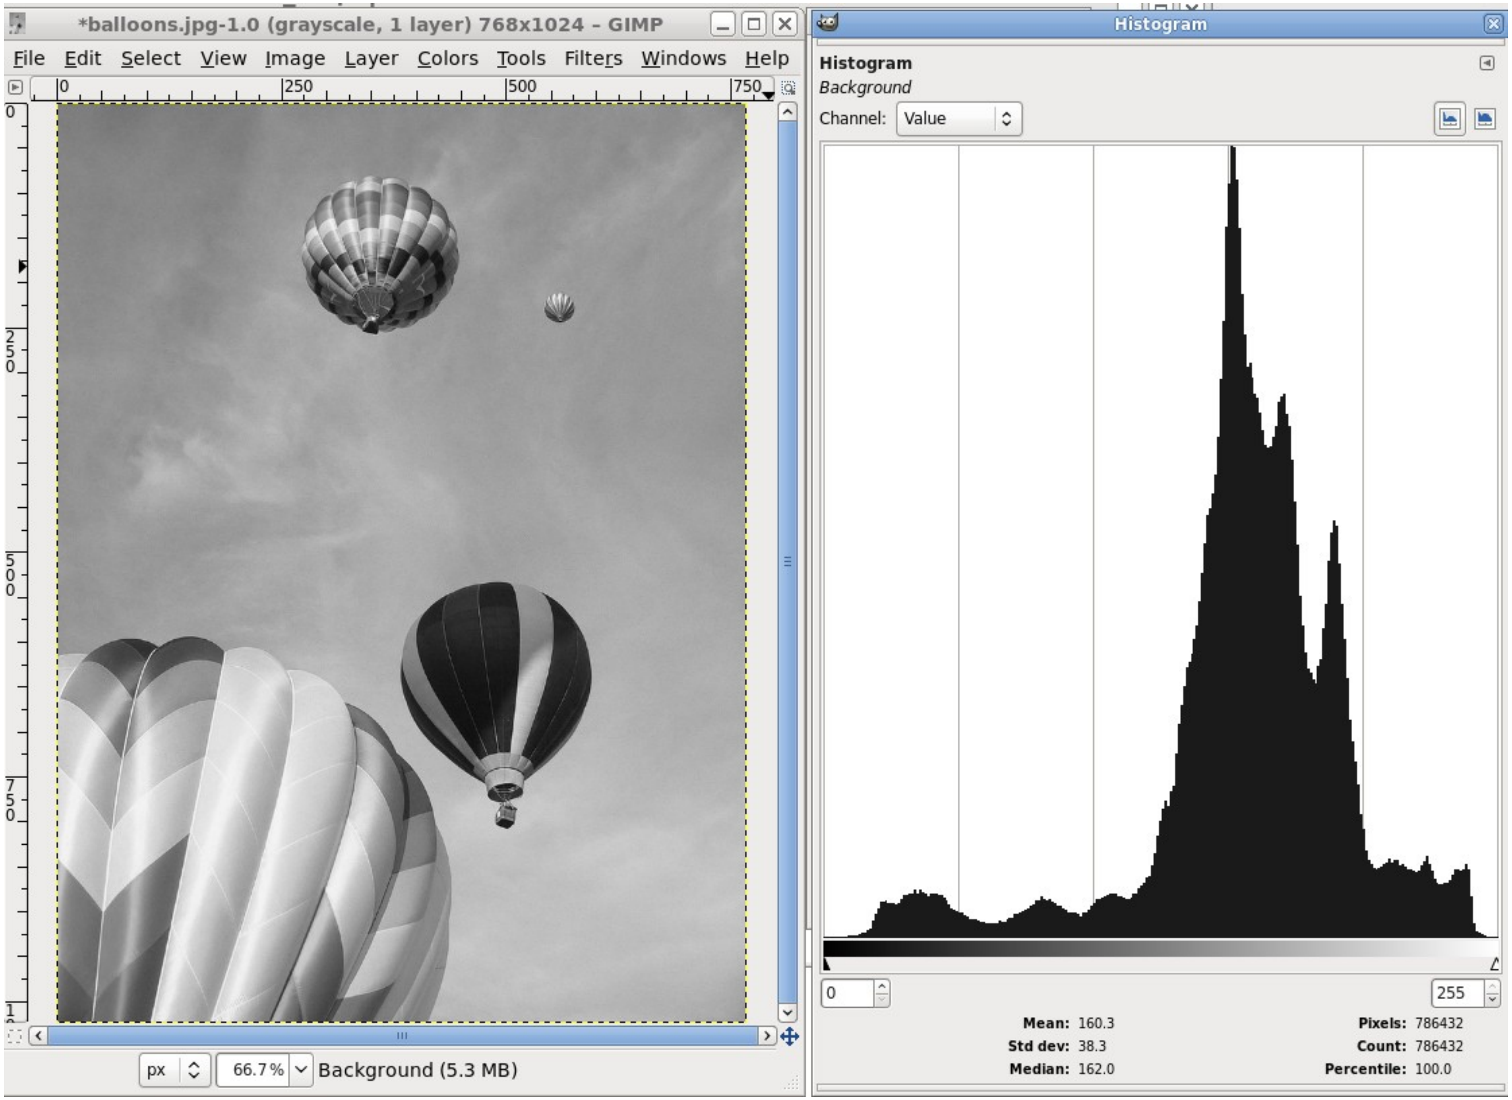

*balloons.jpg-1.0 (grayscale, 1 layer) 768x1024 – GIMP
File Edit Select View Image Layer Colors Tools Filters Windows Help
px 66.7% Background (5.3 MB)
Histogram
Histogram
Background
Channel: Value
0
255
Mean: 160.3
Std dev: 38.3
Median: 162.0
Pixels: 786432
Count: 786432
Percentile: 100.0

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 <= 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 culumlative 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 did cover in image processing course.
- Let me show an example

histograms

- For color histograms, since we have mutliple 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 or more accurately how similar/dissimilar their representations are.
- The hope is that the representation we choose captures enough of the essence of the images/objects that ones that should match should be determined to be similar and those that do not match should be determined to be not similar.
- 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 mutlidimensional 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 = (p1, p2) and q = (q1, q2) then the L1 distance between them is: abs(p1-q1) + abs(p2-q2)
- The L2 distance between them is: sqrt($(p1-q1)^2 + (p2-q2)^2$)

Distance measures

- In general for n dimensions,
- if p=(p1, p2, …, pn) and q = (q1, q2, …, qn)
- the L1 distance between them is:

abs(p1-q1) + abs(p2-q2) + … + abs(pn-qn)

- The L2 distance between them is:

sqrt($(p1-q1)^2 + (p2-q2)^2 + \ldots + (pn-qn)^2$)

- 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

histograms

- The p, q and r of last slide can each be histograms of multidimensions and the L1 and L2 distances can be used as a way to compute the distance between them
- The idea being that similar images will have similar histograms and the distance between their histograms will be small
- Whereas dissimilar images will have very dissimilar histograms and the distance between their histograms will be large
- But remember a histogram just counts up the number of pixels of certain ranges of grey values and disregards the position of the pixel values, so the histogram may not be the best representation of the image to compare for similarity between images
- If the images are of different size (that is different number of pixels) then we should normalize the histograms before computing distance
- Normalized histograms are those that have each bin value from the original histogram divided by the total number of pixels in the image. Therefore the sum of all the values of a normalized histogram will equal 1.

histograms

- Let's look back at these bins and then efficient operations to comput bin# based on R,G,B value.
- 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

Histograms

- How to determine bin # for an RGB (assumes each channel is divided into a power-of-2 range)

- Since R and B are each divided into 2 groups of values (which means divide by 256/2 = 128)
- G is divided into 4 groups of values (divide by 256/4 = 64)

- Rshifted = R >> 7 # same as R // 128 note: 2 to the 7th is 128
- Gshifted = G >> 6 # same as G // 64 note: 2 to the 6th is 64
- Bshifted = B >> 7

- Then to compute bin number:
- Rshifted << 3 + Gshifted << 1 + Bshifted
- Same as Rshifted * 8 (2 to the 3rd) + Gshifted * 2 + Bshifted

Histograms

- Notice that L1 and L2 distances are bin-to-bin comparison. Bins that are near each other vs. far from each other have the same affect on the distance.
- e.g. on the board
- Two pairs of histograms that have equal L1 distance, but with EMD (described on next slide), one pair has a lesser distance

Histograms

- Then let's consider that we have our image converted to HSV. We can create a histogram of Hues, where nearby bins have more similarity than bins that are further apart and instead of doing L1 or L2 which are bin-to-bin comparisons, consider a different distance measure that takes into account cells that are nearby like Earth Mover's Distance (EMD)
- Earth Mover's Distance
  - Think of the values in bins as amounts of dirt and the distance between two histograms is the least amount of work to transform one of the histograms to the other (least amount of work to move the dirt)
  - The amount of work to move dirt is equal to the amount of dirt being moved times the distance moved
  - Example on the board
  - Thoughts on difference between EMD and L1 or L2

In-class exercises

- Exercises on Morphology (dilation and erosion), connected components, computing information about a region, normalizing histograms and computing L1 distance