

SSP 100 - 10

How Computers See

09 / 22 / 2026

Instructor: Michael Eckmann

Today's Topics

- Questions? / Comments?
- Convolution
- In-class exercise
- Look at a different demo of Canny Edge detection
- How can we compare images for similarity?

Convolution

- Because the digital images are numeric we can do computation on them.
- Let's review the convolution operation.
- Technically we're doing cross-correlation operation which is similar to convolution.

In class exercise

- In class exercise on paper handout

Demo programs

- Let me show you a demo of the stages of canny edge detection so you can see the result of each of the steps below. (`canny_edge_stages_gui.py`)
- Also Canny Edge Detection which includes
 - first blurring the image to remove unwanted small details
 - edge detection
 - edge thinning and
 - edge following
 - to get 1 pixel thick edges that are connected into lines

Uses of edge detection

- Defect detection for manufacturing inspection
- Finding boundaries
 - of tumors in medical images
 - of lane lines for robot navigation and self driving
 - of features in satellite or drone imagery
 - and many others

Image Retrieval

- Now let's switch gears.
- Suppose you have a large number of images stored and you have a particular “query” image.
- Goal:
 - find the image or images that are most similar to the query image
 - Or better yet, output a sorted list of all the images stored that range from most similar to the query image down to least similar

Image Retrieval

- Would we compare individual pixel values to each other?
- Positives / negatives of this idea?

Histogram

- Let's consider first a 1 channel grayscale image, where each pixel is a value from 0 .. 255
- A 256 bin histogram of that image would be a structure of 256 cells, each of which holds the number of pixels for each of the values.
- Example on the board.

Histogram

- Instead of 256 bins we could make fewer bins like say 16 bins.
- Bin 0 holds a count of how many pixels in range 0..15
- Bin 1 holds a count of how many pixels in range 16..31
- Bin 2 holds a count of how many pixels in range 32..47
- ...
- Bin 14 holds a count of how many pixels in range 224..239
- Bin 15 holds a count of how many pixels in range 240..255

Histogram



- If we are considering a histogram of an RGB image, we could have over 16 million bins, 1 bin for each different RGB value.
- But just like gray value 26 and 27 look about the same brightness to our eyes, say $\text{RGB}=(120,30,20)$ looks about the same color to us as $\text{RGB}=(121,28,22)$.
- Example of 8 bins on the board (R, G, and B each divided into 2 groups (low values or high values))

Similarity

- To motivate how you might compare histograms for similarity, consider how you tell if a point in 2 dimensional space is closer to a second point or a third point.
- Let me draw on the board.
- Closer implies similarity, further implies dissimilarity.

Similarity / Distance

- Consider 3 dimensional space like the RGB cube.
- How might you tell how close two RGB colors (that is 2 points in 3d space) are?



Distance

- This works for any number of dimensional space (not only 2 or 3).
- For example, we can compute the distance between two 16 bin histograms in the same way.
- This distance is called the Euclidean distance, also known as L2 distance.
- Another distance measure is L1 (sometimes called taxicab or manhattan distance).