Image Processing

1

CS 554 – Computer Vision Pinar Duygulu Bilkent University Today

Image Formation

Point and Blob Processing

Binary Image Processing

Readings: – Gonzalez & Woods, Ch. 3

Slides are adapted from Alyosha Efros and Shapiro & Stockman

Imaging Process

- Light reaches surfaces in 3D.
- Surfaces reflect.
- Sensor element receives light energy.
- Intensity is important.
- Angles are important.
- Material is important.



Adapted from Shapiro and Stockman

Image Formation



FIGURE 2.15 An example of the digital image acquisition process. (a) Energy ("illumination") source. (b) An element of a scene. (c) Imaging system. (d) Projection of the scene onto the image plane. (e) Digitized image.

Sampling and Quantization



FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from A to B in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Sampling and Quantization



a b

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

What is an image?

We can think of an **image** as a function, f, from R² to 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:

 $-f:[a,b]\mathbf{x}[c,d] \rightarrow [0,1]$

A color image is just three functions pasted together. We can write this as a "vector-valued" function:

f(x,y) = [r(x,y) g(x,y) b(x,y)]

Images as functions









What is a digital image?

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 Δ apart, we can write this as:

f[i,j] =Quantize $\{f(i \Delta, j \Delta)\}$

The image can now be represented as a matrix of integer values

i

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

Point Processing

An image processing operation typically defines a new image g in terms of an existing image f.The simplest kind of range transformations are these independent of position x,y:

$$g = t(f)$$

This is called point processing.

Important: every pixel for himself – spatial information completely lost!

Basic Point Processing



Negative



a b FIGURE 3.4 (a) Original digital mammogram. (b) Negative image obtained using the negative transformation in Eq. (3.2-1). (Courtesy of G.E. Medical Systems.)

Image Enhancement

a b c d

FIGURE 3.9

(a) Aerial image. (b)–(d) Results of applying the transformation in Eq. (3.2-3) with c = 1 and $\gamma = 3.0, 4.0$, and 5.0, respectively. (Original image for this example courtesy of NASA.)



Contrast Streching



FIGURE 3.10 Contrast stretching. (a) Form of transformation function. (b) A low-contrast image. (c) Result of contrast stretching. (d) Result of thresholding. (Original image courtesy of Dr. Roger Heady, Research School of Biological Sciences, Australian National University, Canberra,

Image Histograms



CS554 Computer Vision ©



a b

FIGURE 3.15 Four basic image types: dark, light, low contrast, high contrast, and their corresponding histograms. (Original image courtesy of Dr. Roger Heady, Research School of Biological Sciences, Australian National University, Canberra, Australia.) Neighborhood Processing (filtering)

Q: What happens if I reshuffle all pixels within the image?





A: It's histogram won't change. No point processing will be affected...

Need spatial information to capture this.

Goal: Extract "Blobs"



What are "blobs"?

– Regions of an image that are somehow coherent Why?

- Object extraction, object removal, compositing, etc.
- …but are "blobs" objects?
- No, not in general

Blob's coherence

Simplest way to define blob coherence is as similarity in brightness or color:



The tools become blobs



The house, grass, and sky make different blobs

The meaning of a blob

Other interpretations of blobs are possible, depending on how you define the input image:

- Image can be a response of a particular detector
 - Color Detector
 - Face detector
 - Motion Detector
 - Edge Detector

100





 gx^2+gy^2

Why is this useful?



AIBO RoboSoccer (VelosoLab)

Ideal Segmentation



Result of Segmentation



Thresholding

Basic segmentation operation: mask(x,y) = 1 if im(x,y) > T mask(x,y) = 0 if im(x,y) < TT is threshold

User-defined
Or automati

Same as histogram partitioning:



a b c

FIGURE 10.28

(a) Original
image. (b) Image
histogram.
(c) Result of
global
thresholding with *T* midway
between the
maximum and
minimum gray
levels.



As Edge Detection





 gx^2+gy^2

 $gx^{2}+gy^{2} > T$

Sometimes works well...



...but more often not



Adaptive thresholding

Region growing



FIGURE 10.40 (a) Image showing defective welds. (b) Seed points. (c) Result of region growing. (d) Boundaries of segmented defective welds (in black). (Original image courtesy of X-TEK Systems, Ltd.).



- Start with initial set of pixels K
- Add to K any neighbors, if they are within similarity threshold
 Repeat until nothing changes

Is this same as global threshold? What can go wrong?

Color-Based Blob Segmentation

Automatic Histogram Partitioning

- Given image with N colors, choose K
- Each of the K colors defines a region
 - not necessarily contiguous
- Performed by computing color histogram, looking for modes



- This is what happens when you downsample image color range, for instance in Photoshop

Finding modes in a histogram



How Many Modes Are There?

· Easy to see, hard to compute

Mean-Shift (Comaniciu & Meer)



Iterative Mode Search

- 1. Initialize random seed, and fixed window
- 2. Calculate center of gravity of the window (the "mean")
- 3. Translate the search window to the mean
- 4. Repeat Step 2 until convergence

Mean-Shift (Comaniciu & Meer)





More Examples: http://www.caip.rutgers.edu/~comanici/segm_images.html

Issues:

Although often useful, all these approaches work only some of the time, and are considered rather "hacky".

Can't even handle our tiger:



Problem is that blobs != objects!

Binary images are handy in many cases (sprite extraction, compositing, etc).

Binary image processing is a well-studied field, based on set theory, called Mathematical Morphology

consists of a set of image analysis operations that are used to produce or process binary images, usually images of 0's and 1's.

0 represents the background1 represents the foreground

00010010001000 00011110001000 00010010001000

Application Areas

Document Analysis



Industrial Inspection



Medical Imaging



34

Separate objects from background and from one another

- Aggregate pixels for each object
- Compute features for each object

Example red blood cell image



- Many blood cells are separate objects
- Many touch to each other bad!
- Salt and pepper noise from thresholding
- How usable is this data?

- 63 separate objects detected
- Single cells have area about 50
- Noise spots
- Gobs of cells
More controlled images

- More uniform objects
- More uniform background
- Objects actually separated

- 15 objects detected
- Location known
- Area known
- 3 distinct clusters of 5 values of area; 85, 145, 293

Results on Coloring Pacmen



Thresholding



Binary images can be obtained from gray level images by thresholding

- Object region of interest has intensity distribution different from background
- Region pixel likely to be identified by intensity alone:
 - intensity > a
 - intensity < b
 - -a < intensity < b
- Works OK with flat-shaded scenes or engineered scenes.
- Does not work well with natural scenes.

CS554 Computer Vision © Pinar Duygulu

Use of histograms for threshold selection

- Cherry image with 3 regions
- Background is black
- Healthy cherry is bright
- Bruise is medium dark
- Histogram shows two cherry regions (black background has been removed)





How can we use a histogram to separate an image into 2 (or several) different regions?



Is there a single clear threshold? 2? 3?

Choosing Threshold

Detect peaks and valleys



Two distinct modes



Overlapped modes

CS554 Computer Vision © Pinar Duygulu

- Find the deepest valley between two modes of bimodal histogram
- Fit two or more Gaussian curves to the histogram
- Dynamic thresholding

Cleaning up thresholding results



- Delete object pixels on boundary to better separate parts.
- Fill small holes
- Delete tiny objects
- (last 2 are "salt-andpepper" noise)

Binary mathematical morphology consists of two basic operations

dilation and erosion

and several composite relations

closing and opening conditional dilation

. . .

Dilation expands the connected sets of 1s of a binary image.

It can be used for

1. growing features



2. filling holes and gaps





Erosion shrinks the connected sets of 1s of a binary image.



2. Removing bridges, branches and small protrusions



A structuring element is a shape mask used in the basic morphological operations.

They can be any shape and size that is digitally representable, and each has an origin.



The arguments to dilation and erosion are

a binary image B a structuring element S

dilate(B,S) takes binary image B, places the origin of structuring element S over each 1-pixel, and ORs the structuring element S into the output image at the corresponding position.



Dilation

a b c d e FIGURE 9.4 (a) Set *A*. (b) Square structuring element (dot is the center). (c) Dilation of A by B, shown shaded. (d) Elongated structuring element. (e) Dilation of A using this element.



Dilation

Historically, certain computer programs were written using only two digits rather than four to define the applicable year. Accordingly, the company's software may recognize a date using "00" as 1900 rather than the year 2000.



Historically, certain computer programs were written using only two digits rather than four to define the applicable year. Accordingly, the company's software may recognize a date using "00" as 1900 rather than the year 2000. a c b

joined.

FIGURE 9.5 (a) Sample text of poor resolution with broken characters (magnified view). (b) Structuring element. (c) Dilation of (a) by (b). Broken segments were

0	1	0
1	1	1
0	1	0

erode(B,S) takes a binary image B, places the origin of structuring element S over every pixel position, and ORs a binary 1 into that position of the output image only if every position of S (with a 1) covers a 1 in B.



Erosion



Erosion





Original image

Eroded image

Erosion







Eroded twice

CS554 Computer Vision © Pinar Duygulu

- Closing is the compound operation of dilation followed by erosion (with the same structuring element)
- Opening is the compound operation of erosion followed by dilation (with the same structuring element)

Opening and Closing

Opening : smoothes the contour of an object, breaks narrow isthmuses, and eliminates thin protrusions $A \circ B = (A \Theta B) \oplus B$

Closing : smooth sections of contours but, as opposed to opning, it generally fuses narrow breaks and long thin gulfs, eliminates small holes, and fills gaps in the contour

$$A \bullet B = (A \oplus B) \Theta B$$

Prove to yourself that they are not the same thing. Play around with bwmorph in Matlab.

1	1	1	1	1	1	1	
			1	1	1	1	
			1	1	1	1	
		1	1	1	1	1	
			1	1	1	1	
		1	1				

a) Binary image B

1	1	1
1	1	1
1	1	1

b) Structuring Element S

1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
	1	1	1	1	1	1	1
	1	1	1	1	1	1	1
	1	1	1	1	1	1	1
	1	1	1	1	1	1	1
	1	1	1	1			

c) Dilation $B \oplus S$

1	1	1	1	1	1	
	1	1	1	1	1	
	1	1	1	1	1	
	1	1	1	1	1	
	1	1	1	1	1	
	1	1				

e) Closing $B \bullet S$

1 1 1 1 1 1

d) Erosion $B \ominus S$

\square						
		1	1	1	1	
		1	1	1	1	
		1	1	1	1	
		1	1	1	1	
		1	1	1	1	

f) Opening $B \circ S$



В

0000100 0 0 $\mathbf{0}$ $\mathbf{0}$ \mathbf{O} \mathbf{O} \mathbf{O} \cap \cap \cap () \mathbf{O} \mathbf{O} \mathbf{O} 1 1 ()()1 1 1 1 ()()1 1 ()()()()()()1 1 1 1 1 1 1 1 1 1 1 1

S

1

1

111

1 1

Ω

1

0

1 1 1

1

 \mathbf{O}

1

0

59

Opening and Closing





THE TEST IMAGE OPENING: The original image eroded twice and dilated twice (opened). Most noise is removed

CLOSING: The original image dilated and then eroded. Most holes are filled.

CS554 (

Opening and Closing



Boundary Extraction

 β (*A*) = *A* - (*A* Θ *B*)



62

CS554 Computer Vision © Pinar Duygulu

Boundary Extraction



a b

FIGURE 9.14

(a) A simple binary image, with 1's represented in white. (b) Result of using Eq. (9.5-1) with the structuring element in Fig. 9.13(b).

Region Filling



 $X_k = (X_{k-1} \oplus B) \cap A^c \quad k = 1, 2, 3....$

Connected Component Labeling

Once you have a binary image, you can identify and then analyze each **connected set of pixels**.

The connected components operation takes in a binary image and produces a **labeled image** in which each pixel has the integer label of either the background (0) or a component.



Extraction of Connected Components



FIGURE 9.17 (a) Set *A* showing initial point *p* (all shaded points are valued 1, but are shown different from *p* to indicate that they have not yet been found by the algorithm). (b) Structuring element. (c) Result of first iterative step. (d) Result of second step. (e) Final result.

$$X_k = (X_{k-1} \oplus B) \cap A \quad k = 1, 2, 3, \dots$$

Labeling shown as pseudo-color





connected components of 1's from thresholded image





connected components of cluster labels

Connectivity

	Ν	
W	*	Ε
	\mathbf{S}	

NW	Ν	NE
W	*	Ε
SW	S	SE

4-neighborhood

8-neighborhood

Recursive labeling

1	1	0	1	1	1	0	1
1	1	0	1	0	1	0	1
1	1	1	1	0	0	0	1
0	0	0	0	0	0	0	1
1	1	1	1	0	1	0	1
0	0	0	1	0	1	0	1
1	1	0	1	0	0	0	1
1	1	0	1	0	1	1	1

a) binary image



b) connected components labeling





c) binary image and labeling, expanded for viewing

First Step : Run Length Encoding

Segment each image row into groups of similar pixels called *runs*

- Runs store a start and end point for each contiguous row of color







CS554 Computer Vision © Pinar Duygulu

Second Step : Merging Regions





Х

х

Final Results

Runs are merged into multi-row regions Image is now described as contiguous regions instead of just pixels



CS554 Computer Vision © Pinar Duygulu
Blob Properties

Now that we have nice, clean blobs, what can we do with them?

Compute Statistics:

- Area
- Perimeter
- Aspect ratio
- Center of mass
- best-fitting ellipse
- Average color
- Etc.

All this can be used to classify blobs and decide if they hold the objects we are interested in.





Area an Centroid

- We denote the set of pixels in a region by R.
- assuming square pixels: area:

$$A = \sum_{(r,c) \in R} - 1$$

centroid:

$$ar{r} = rac{1}{A} \quad \sum_{(r,c)\in R} r$$

 $ar{c} = rac{1}{A} \quad \sum_{(r,c)\in R} c$

- (\bar{r}, \bar{c}) is generally not a pair of integers.
- A precision of tenths of a pixel is often justifiable for the centroid.

Perimeter pixels and length

- Let perimeter P be the actual set of boundary pixels.
- P must be ordered in a sequence $P = \langle (r_o, c_o), \ldots, (r_{K-1}, c_{K-1}) \rangle$.
- Each pair of successive pixels in P are neighbors, including the first and last pixels.

perimeter length:

$$|P| = \#\{k | (r_{k+1}, c_{k+1}) \in N_4(r_k, c_k)\} + \sqrt{2}\#\{k | (r_{k+1}, c_{k+1}) \in N_8(r_k, c_k) - N_4(r_k, c_k)\}$$

where k + 1 is computed modulo K.

• Perimeter can vary significantly with object orientation.

 common measure of circularity of a region is length of the perimeter squared divided by area. circularity(1):

$$C_1 = \frac{|P|^2}{A}$$

• a second measure uses variation off of a circle circularity(2):

$$C_2 = \frac{\mu_R}{\sigma_R}$$

where μ_R and σ_R^2 are the mean and variance of the distance from the centroid of the shape to the boundary pixels (r_k, c_k) . mean radial distance:

$$\mu_R = \frac{1}{K} \sum_{k=0}^{K-1} \| (r_k, c_k) - (\bar{r}, \bar{c}) \|$$

variance of radial distance:

$$\sigma_R^2 = \frac{1}{K} \sum_{k=0}^{K-1} \left[\| (r_k, c_k) - (\bar{r}, \bar{c}) \| - \mu_R \right]^2$$



Bounding box



second-order row moment:

$$\mu_{rr} = rac{1}{A} \sum_{(r,c)\in R} (r-\bar{r})^2$$

second-order mixed moment:

$$\mu_{rc} = \frac{1}{A} \sum_{(r,c)\in R} (r-\bar{r})(c-\bar{c})$$

second-order column moment:

$$\mu_{cc} = \frac{1}{A} \sum_{(r,c)\in R} (c - \bar{c})^2$$

These are invariant to object location in the image.

Axis of least inertia

An axis which we could spin the pixels with least energy input The axis about which the pixels have least second moment



A region adjacency graph (RAG) is a graph in which each node represents a region of the image and an edge connects two nodes if the regions are adjacent.

0	0	0	0	0	0	0	0	0	0
0	1	1	1	1	1	0	2	2	0
0	1	-1	-1	-1	1	0	2	2	0
0	1	1	1	1	1	0	2	2	0
0	0	0	0	0	0	0	2	2	0
0	3	3	3	0	2	2	2	2	0
0	3	-2	3	0	2	-3	ŝ	2	0
0	3	-2	3	0	2	-3	-3	2	0
0	3	3	3	0	2	2	2	2	0
0	0	0	0	0	0	0	0	0	0

