

REGION & EDGE BASED SEGMENTATION



Fritz Albrechtsen 22.09.2010

We go through sections 10.1, 10.2.7 (briefly), 10.4, 10.5, 10.6.1

We cover the following segmentation approaches:

1. Edge-based segmentation
2. Region growing
3. Region split and merge
4. Watershed segmentation
5. Segmentation by motion

Assumed known:

1. Edge detection, point and line detection (10.2.1-10.2.6)
2. Segmentation by thresholding (10.3)
3. Basic mathematical morphology (9.1-9.4)

Need repetition of INF2310-stuff?

- Edge detection:
 - Read sections 10.2 and 3.6 in Gonzalez and Woods.
 - Look at spring 2009 lecture in INF2310 on "Filtering II" (in Norwegian).
- Thresholding:
 - Read section 10.3 in Gonzalez and Woods.
 - Go through spring 2010 lecture in INF2310 on "Segmentation" (in Norwegian)
- Morphology:
 - Read section 9.1-9.4 in Gonzalez and Woods.
 - Go through spring 2010 lecture in INF2310 on "Morphology" (in Norwegian)

What is segmentation

- Segmentation is the separation of one or more regions or objects in an image based on a discontinuity or a similarity criterion.
- A region in an image can be defined by its border (edge) or its interior, and the two representations are equal.
- If you know the interior, you can always define the border, and vice versa.
- Because of this, image segmentation approaches can typically be divided into two categories; **edge and region based methods.**

Visual segmentation is easy (?)



What if you wanted to group all the pixels corresponding to the car in this image into one group. Would that be easy?

What is a region?

- A region R of an image I is defined as a connected homogenous subset of the image with respect to some criterion such as gray level or texture (previous lecture)
- A segmentation of an image f is a partition of f into several homogeneous regions $R_i, i=1, \dots, m$
- An image f can be segmented into regions R_i such that:

$$1) f = \bigcup_{i=1}^m R_i$$

$$2) R_i \cap R_j = \emptyset, 1 \leq i, j \leq m \text{ and } i \neq j$$

$$3) P(R_i) = \text{true for all } i$$

$$4) P(R_i \cup R_j) = \text{false}, 1 \leq i, j \leq m \text{ and } i \neq j \text{ and}$$

R_i, R_j are adjacent

$P(R_i)$ is a logical predicate defined over all points in R_i . It must be true for all pixels inside the region and false for pixels in other regions.

Regions R_i and R_j are neighbors if their union forms a connected component.

Segmentation approaches

- **Pixel-based segmentation:**
each pixel is segmented based on gray-level values, no contextual information, only histogram.
 - Example: thresholding.
- **Region-based segmentation:**
considers gray-levels from neighboring pixels, either by including similar neighboring pixels (region growing), split-and-merge, or watershed segmentation.
- **Edge-based segmentation:**
Detects edge pixels and links them together to form contours.

Region vs. edge-based approaches



- Region based methods are robust because:
 - Regions cover more pixels than edges and thus you have more information available in order to characterize your region
 - When detecting a region you could for instance use texture which is not easy when dealing with edges
 - Region growing techniques are generally better in noisy images where edges are difficult to detect
- The edge based method can be preferable because:
 - Algorithms are usually less complex
 - Edges are important features in a image to separate regions
- The edge of a region can often be hard to find because of noise or occlusions
- Combination of results may often be a good idea

Edge-based segmentation

Two steps are needed:


1. Edge detection (to identify "edgels" - edge pixels)
 - (Gradient, Laplacian, LoG, Canny filtering)
2. Edge linking – linking adjacent "edgels" into edges
 - *Local Processing*
 - *magnitude* of the gradient
 - *direction* of the gradient vector
 - edges in a predefined neighborhood are linked if both magnitude and direction criteria is satisfied
 - *Global Processing via Hough Transform*

Edge detection

- A large group of methods.
- Operators detecting discontinuities
 - in gray level, color, texture, etc.
- Result from edge detection cannot be used directly.
- Post processing steps must follow to combine edges into edge chains to represent the region border
- The more prior information used in the segmentation process, the better the segmentation results can be obtained
- The most common problems of edge-based segmentation are:
 - edge presence in locations where there is no border
 - no edge presence where a real border exists

Why is a gradient operator not enough?

- Some images produce nice gradient maps.
- Most interesting/challenging images produce a very complicated gradient map
 - e.g. from a Sobel filter.
- **Only rarely will the gradient magnitude be zero, ... why?**
- **Distribution of gradient magnitude is often skewed, ... why?**
- **How do we determine gradient magnitude threshold, ... ?**




A challenging graylevel image




Russian ground to air missiles near Havana, Cuba.
Imaged by U2 aircraft 29.08.1962.

The gradient magnitude image




Magnitude of gradient image, min=0, max=1011

Edges from thresholding - I



Gradient magnitude thresholded at 250

Edges from thresholding - II




Gradient magnitude thresholded at 350

Improved edge detection

- What if we assume the following:
 - All gradient magnitudes above a strict threshold are assumed to belong to a bona fide edge.
 - All gradient magnitudes above an unstrict threshold and connected to a pixel resulting from the strict threshold are also assumed to belong to real edges.
- Hysteresis thresholding – Canny's edge detection (see INF 2310).


Edges from hysteresis thresholding



Result of hysteresis, are we really impressed?

Thinning of edges


- 1 Quantize the edge directions into four (or eight) directions.
- 2 For all nonzero gradient magnitude pixels, inspect the two neighboring pixels in the four (or eight) directions.
- 3 If the edge magnitude of any of these neighbors is higher than that under consideration, mark the pixel.
- 4 When all pixels have been scanned, delete or suppress the marked pixels.




Used iteratively in nonmaxima suppression.

Local edge linking

Left: Input image
Right: G_y



Left: G_x
Right: Result after edge linking



Hough-transform


- We have thresholded an edge magnitude image.
- Thus, we have n pixels that may partially describe the boundary of some objects.
- We wish to find sets of pixels that make up straight lines.
- Regard a point (x_i, y_i) on a straight line $y_i = ax_i + b$
 - There are many lines passing through the point (x_i, y_i) .
 - They all satisfy the equation for some set of parameters (a, b) .

- The equation can obviously be rewritten as follows:

$$b = -xa + y$$

- This is a line in the (a, b) plane parameterized by x and y .
- Another point (x, y) will give rise to another line in (a, b) space.

Hough transform – basic idea



Hough transform – basic idea

- Two points (x,y) and (z,k) define a line in the (x,y) plane.
- Two (x,y) points correspond to different lines in (a,b) .
- In (a,b) plane these lines will intersect in a point (a',b') .
- All points on the line defined by (x,y) and (z,k) in (x,y) space will parameterize lines that intersect in (a',b') in (a,b) space.
- Points that lie on a line will form a “cluster of crossings” in the (a,b) space.

Hough transform – algorithm

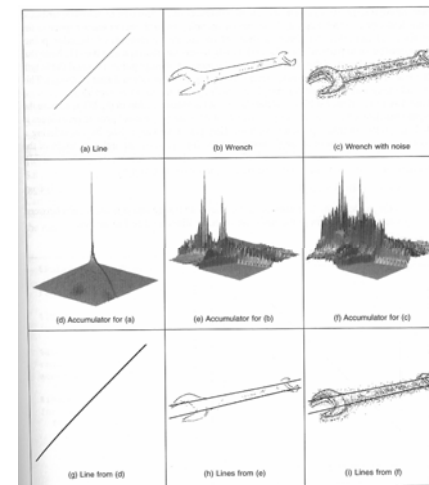
- Quantize the parameter space (a,b) into accumulator cells.
- Perform accumulation in (a,b) space:
 - For each point (x,y) with value 1 in the binary image, find the values of (a,b) in the range $[[a_{\min}, a_{\max}], [b_{\min}, b_{\max}]]$ defining the line corresponding to this point.
 - Increase the value of the accumulators along this line.
 - Then proceed with the next point in the image.
- Lines in (x,y) plane can be found as peaks in this accumulator space.

Images and accumulator space

Thresholded edge images

Visualizing the accumulator space
The height of the peak will be defined by the number of pixels in the line.

Thresholding the accumulator space and superimposing this onto the edge image



Note how noise affects accumulator. Even with noise, the largest peaks correspond to the major lines.

Hough transform – pros and cons

- Advantages:
 - Conceptually simple.
 - Easy implementation.
 - Handles missing and occluded data very gracefully.
 - Can be adapted to many types of forms, not just lines.
- Disadvantages:
 - Computationally complex for objects with many parameters.
 - Looks for only one single type of object.
 - Can be “fooled” by “apparent lines”.
 - The length and the position of a line segment cannot be determined.
 - Co-linear line segments cannot be separated.

Using HT for edge linking




FIGURE 10.34 (a) A 502×564 aerial image of an airport. (b) Edge image obtained using Canny's algorithm. (c) Hough parameter space (the boxes highlight the points associated with long vertical lines). (d) Lines in the image plane corresponding to the points highlighted by the boxes). (e) Lines superimposed on the original image.

Σ : Edge based segmentation

- Advantages
 - An approach similar to how humans segment images.
 - Works well in images with good contrast between object and background
- Disadvantages
 - Does not work well on images with smooth transitions and low contrast
 - Sensitive to noise
 - Robust edge linking is not trivial

Region growing (G&W:10.4.1)

- Grow regions by recursively including the neighboring pixels that are similar and connected to the seed pixel.
- Connectivity is needed not to connect pixels in different parts of the image.
- Similarity measures - examples:
 - Use difference in gray levels for regions with homogeneous gray levels
 - Use texture features for textured regions



Similarity: gray level difference



Cannot be segmented using gray level similarity.

Region growing

- Starts with a set of seeds (starting pixels)
 - Predefined seeds
 - All pixels as seeds
 - Randomly chosen seeds
 - Region growing steps (bottom-up method)
 - Find starting points
 - Include neighboring pixels with similar features (graylevel, texture, color) **A similarity measure must be selected.**
- Two variants:
1. Select seeds from the whole range of grey levels in the image. Grow regions until all pixels in image belong to a region.
 2. Select seed only from objects of interest (e.g. bright structures). Grow regions only as long as the similarity criterion is fulfilled.
- Problems:
 - Not trivial to find good starting points
 - Need good criteria for similarity

F3 22.09.10

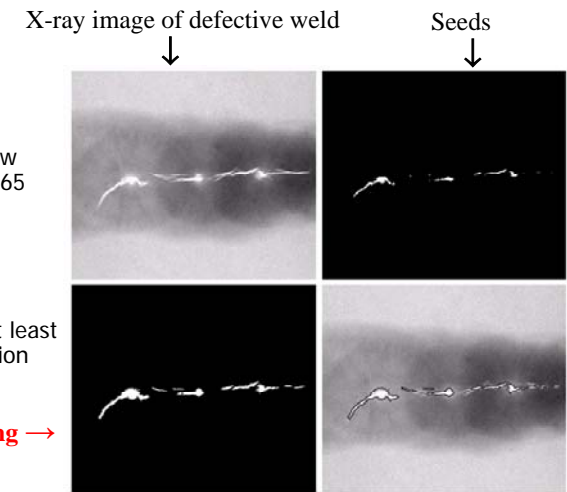
INF 4300

29

Region growing example

- Seeds: $f(x,y) = 255$
 - $P(R) = \text{TRUE}$ if
 - | seed gray level – new pixel gray level | < 65
- and
- New pixel must be 8-connected with at least one pixel in the region

Result of region growing →



F3 22.09.10

INF 4300

30

Similarity measures

- Graylevel or color?
- Graylevel and spatial properties, e.g., texture, shape, ...
- Intensity difference within a region (from pixel to seed or to region average so far)
- Within a value range (min, max)
- Distance between mean value of the regions (specially for region merging or splitting)
- Variance or standard deviation within a region.

F3 22.09.10

INF 4300

31

Region merging techniques

- Initialize by giving all the pixels a unique label
 - All pixels in the image are assigned to a region.
- The rest of the algorithm is as follows:
 - In some predefined order, examine the neighbor regions of all regions and decide if the predicate evaluates to true for all pairs of neighboring regions.
 - If the predicate evaluates to true for a given pair of neighboring regions then give these neighbors the same label.
 - The predicate is the similarity measure (can be defined based on e.g. region mean values or region min/max etc.).
 - Continue until no more mergings are possible.
 - Upon termination all region criteria will be satisfied.

F3 22.09.10

INF 4300

32

Region merging techniques

- The aim is to separate the apples from the background.
- This image poses more of a challenge than you might think.




Region merging techniques

- We run a standard region merging procedure where all pixels initially are given a unique label.
- If neighboring regions have mean values within 10 gray levels they are fused.
- Regions are considered neighbors in 8-connectivity.



Did Warhol merge Marilyn?



Region merging techniques

- Initialization is critical, results will in general depend on the initialization.
- The order in which the regions are treated will also influence the result:
 - The top image was flipped upside down before it was fed to the merging algorithm.
- Notice the differences!




Split and merge (G&W:10.4.2)

- Separate the image into regions based on a given similarity measure. Then merge regions based on the same or a different similarity measure.
- The method is also called "quad tree" division:
 - Set up some criteria for what is a uniform area e.g., mean, variance, texture, etc.
 - Start with the full image and **split** it in to 4 sub-images.
 - Check each sub-image. If not uniform, divide into 4 new sub-images
 - After each iteration, compare neighboring regions and **merge** if uniform according to the similarity measure.



(a) Whole Image


(b) First Split



(c) Second Split


(d) Merge

Split and merge example




Watershed segmentation (G&W:10.5)

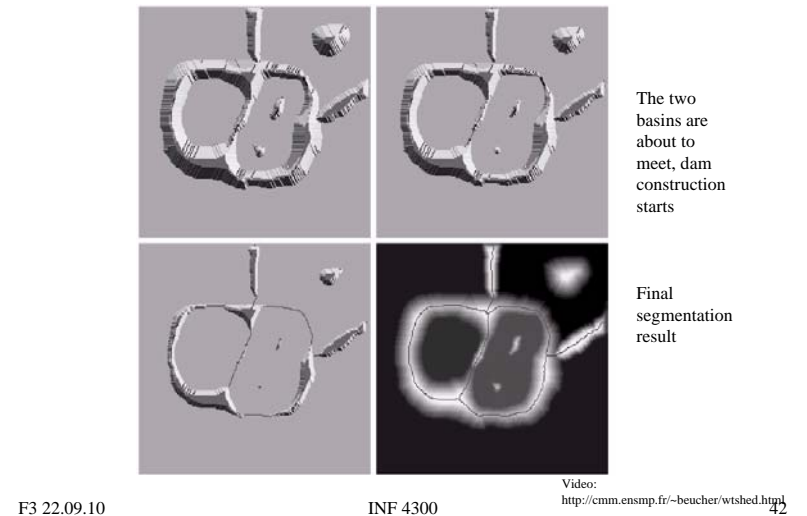
- Look at the image as a 3D topographic surface, (x,y,intensity), with both valleys and mountains.
- Assume that there is a hole at each minimum, and that the surface is immersed into a lake.
- The water will enter through the holes at the minima and flood the surface.
- To avoid two different basins to merge, a dam is built.
- Final step: the only thing visible would be the dams.
- The connected dam boundaries correspond to the watershed lines.



Watershed segmentation



Watershed segmentation



Watershed segmentation

- Can be used on images derived from:
 - The intensity image
 - Edge enhanced image
 - Distance transformed image
 - Thresholded image. From each foreground pixel, compute the distance to a background pixel.
 - Gradient of the image
- Most common: gradient image

Watershed algorithm

- Let $g(x,y)$ be the input image (often a gradient image).
- Let M_1, \dots, M_R be the coordinates of the regional minima.
- Let $C(M_i)$ be a set consisting of the coordinates of all points belonging to the catchment basin associated with the regional minimum M_i .
- Let $T[n]$ be the set of coordinates (s,t) where $g(s,t) < n$

$$T[n] = \{(s,t) \mid g(s,t) < n\}$$

This is the set of coordinates lying below the plane $g(x,y) = n$
 This is the candidate pixels for inclusion into the catchment basin, but we must take care that the pixels do not belong to a different catchment basin.


Watershed algorithm cont.

- The topography will be flooded with integer flood increments from $n=\min-1$ to $n=\max+1$.
- Let $C_n(M_i)$ be the set of coordinates of point in the catchment basin associated with M_i , flooded at stage n .
- This must be a connected component and can be expressed as $C_n(M_i) = C(M_i) \cap T[n]$ (only the portion of $T[n]$ associated with basin M_i)
- Let $C[n]$ be union of all flooded catchments at stage n :

$$C[n] = \bigcup_{i=1}^R C_n(M_i) \quad \text{and} \quad C[\max+1] = \bigcup_{i=1}^R C(M_i)$$

Dam construction

- Stage $n-1$: two basins forming separate connected components.
- To consider pixels for inclusion in basin k in the next step (after flooding), they must be part of $T[n]$, and also be part of the connected component q of $T[n]$ that $C_{n-1}[k]$ is included in.
- Use morphological dilation iteratively.
- Dilation of $C[n-1]$ is constrained to q .
- The dilation can not be performed on pixels that would cause two basins to be merged (form a single connected component)



Watershed algorithm cont.



- Initialization: let $C[\min+1] = T[\min+1]$
- Then recursively compute $C[n]$ from $C[n-1]$:
 - Let Q be the set of connected components in $T[n]$.
 - For each component q in Q , there are three possibilities:
 1. $q \cap C[n-1]$ is empty – **new minimum**
Combine q with $C[n-1]$ to form $C[n]$.
 2. $q \cap C[n-1]$ contains one connected component of $C[n-1]$
 q lies in the catchment basin of a regional minimum
Combine q with $C[n-1]$ to form $C[n]$.
 3. $q \cap C[n-1]$ contains more than one connected component of $C[n-1]$
 q is then on a ridge between catchment basins, and a dam must be built to prevent overflow.
Construct a one-pixel dam by dilating $q \cap C[n-1]$ with a 3x3 structuring element, and constrain the dilation to q .
- Constrain n to existing intensity values of $g(x,y)$ (obtain them from the histogram).

Watershed alg.

3. $q \cap C[n-1]$ contains more than one connected component of $C[n-1]$
 q is then on a ridge between catchment basins, and a dam must be built to prevent overflow.
Construct a one-pixel dam by dilating $q \cap C[n-1]$ with a 3x3 structuring element, and constrain the dilation to q .
- Constrain n to existing intensity values of $g(x,y)$ (obtain them from the histogram).

Watershed example

FIGURE 10.46
(a) Image of blobs. (b) Image gradient. (c) Watershed lines. (d) Watershed lines superimposed on original image. (Courtesy of Dr. S. Beucher, CMM/Ecole des Mines de Paris.)




F3 22.09.10

INF 4300

49

"Over-segmentation" or fragmentation




- Using the gradient image directly can cause over-segmentation because of noise and small irrelevant intensity changes
- Improved by smoothing the gradient image or using markers

F3 22.09.10

INF 4300

50

Solution: Watershed with markers



- A marker is an extended connected component in the image
- Can be found by intensity, size, shape, texture etc
- Internal markers are associated with the object (a region surrounded by bright point (of higher altitude))
- External markers are associated with the background (watershed lines)
- Segment each sub-region by some segmentation algorithm

F3 22.09.10

INF 4300

51

How to find markers

- Apply filtering to get a smoothed image
- Segment the smooth image to find the internal markers.
 - Look for a set of point surrounded by bright pixels.
 - How this segmentation should be done is not well defined.
 - Many methods can be used.
- Segment smooth image using watershed to find external markers, with the restriction that the internal markers are the only allowed regional minima. The resulting watershed lines are then used as external markers.
- We now know that each region inside an external marker consists of a single object and its background.
- Apply a segmentation algorithm (watershed, region growing, threshold etc.) only inside each watershed.


F3 22.09.10

INF 4300

52

Splitting objects by watershed

- Example: Splitting touching or overlapping objects.
 - Given graylevel (or color) image
 - Perform first stage segmentation
 - (edge detection, thresholding, classification,...)
 - Now you have a labeled image, e.g. foreground / background
 - Obtain distance transform image
 - From each foreground pixel, compute distance to background.
 - Use watershed algorithm on inverse of distance image.




F3 22.09.10

INF 4300

53


Watershed – advanced example



Original


Opening

Threshold, global




Distance transform
- Distance from a point in a region to the border of the region

F3 22.09.10



Watershed of inverse of distance transform

INF 4300




Second watershed to split cells

Matlab:
imopen
bwdist
imimposemin
imregionalmax
watershed

54

Watershed – example 3



Original


Opening

Threshold



Find internal and external markers from gradient image

F3 22.09.10



Watershed

INF 4300

Matlab:

- imopen
- imimposemin
- bwmorph
- watershed

55

Σ: Watershed

- Advantages
 - Gives connected components
 - A priori information can be implemented in the method using markers
- Disadvantages :
 - Often needs preprocessing to work well
 - Fragmentation or "over-segmentation" can be a problem

F3 22.09.10

INF 4300

56

Segmentation by motion

- In video surveillance one of the main tasks is to segment foreground objects from the background scene to detect moving objects.
- The background may have a lot of natural movements and changes
 - (moving trees, snow, sun, shadows etc)
- Motion detection may be restricted to given ROI
 - to avoid moving trees etc.

Background subtraction

- Thresholding intensity differences between frames:

$$d_{ij}(x, y) = \begin{cases} 1 & \text{if } |f(x, y, t_i) - f(x, y, t_j)| > T \\ 0 & \text{otherwise} \end{cases}$$

- Images must be of same size, co-registered, same lighting, ...
- Clean up small objects in d_{ij} , caused by noise.
- Possible to accumulate counts of differences per pixel (ADI's).
- Background subtraction:
 - Estimate the background image (per pixel) by computing the mean and variance of the n previous time frames at pixel (x,y) .
 - Subtract background image from current frame
 - Note that the background is updated by only considering the n previous frames. This eliminated slow trends.