



Image Segmentation

Brightness, Color,
Texture, Motion, Shape

Basic Concepts

Edge-based
Segmentation

Region-based
Segmentation

The Watershed
Transform

Summary

Chapter 5

Image Segmentation I: Basics

Variational Methods for Computer Vision

Winter 2014/15

Prof. Daniel Cremers
Chair for Computer Vision and Pattern Recognition
Department of Computer Science
Technische Universität München

Overview

- 1 Image Segmentation
- 2 Brightness, Color, Texture, Motion, Shape
- 3 Basic Concepts
- 4 Edge-based Segmentation
- 5 Region-based Segmentation
- 6 The Watershed Transform
- 7 Summary



Image Segmentation

Brightness, Color,
Texture, Motion, Shape

Basic Concepts

Edge-based
Segmentation

Region-based
Segmentation

The Watershed
Transform

Summary

Image Segmentation – A Difficult Problem



Original photograph by R. C. James



Image Segmentation

Brightness, Color,
Texture, Motion, Shape

Basic Concepts

Edge-based
Segmentation

Region-based
Segmentation

The Watershed
Transform

Summary



Image Segmentation

Brightness, Color,
Texture, Motion, Shape

Basic Concepts

Edge-based
Segmentation

Region-based
Segmentation

The Watershed
Transform

Summary

The goal of **image segmentation** (Bildsegmentierung) is to partition the image plane into “meaningful” components.


What is meaningful depends on the application. Typically one may want a segmentation where each region corresponds to a separate object or structure in the scene. In this sense, image segmentation is tightly coupled with figure-ground discrimination, **image interpretation** and semantic analysis.

Image segmentation is the **most studied problem** in image processing.

There exist many approaches. They typically differ in:

- **which local properties are considered** in the process (brightness, color, texture, motion,...).
- **how the partitioning is computed** (examples: region merging, region growing, watershed, graph cuts, level sets, convex relaxation techniques,...).

Segmentation: Brightness



Mumford, Shah '89, Chan, Vese, TIP '01



Segmentation: Color



Keuchel et al. PAMI '03



Nieuwenhuis, Cremers PAMI '13




Segmentation: Texture




Brox, Weickert, ECCV '04



Heiler, Schnörr, IJCV '05




Awate et al., ECCV '06



Awate et al., ECCV '06




Segmentation: Motion



Cremers, Soatto, *Motion Competition*, IJCV 2005.




Segmentation: Brightness and Shape



Cremers et al., ECCV '02

Schoenemann, Cremers PAMI '09

Semantic Multilabel Segmentation




Souiai et al., EMMCVPR '13

Semantic Multilabel Segmentation



Multi-label segmentation with length regularity



Multi-label segmentation with proportion priors

Nieuwenhuis, Strekalovskiy, Cremers, ICCV '13



Basic Concepts in Image Segmentation

Segmentation methods are generally based on two complementary concepts:

- **Edge-based methods** identify contours which approximate discontinuities of the color or texture. Whereas traditional methods heuristically group the output of edge detectors into connected curves, in recent years optimal boundaries are computed by energy minimization.
- **Region-based methods** identify regions in the image plane for which some criterion is more or less uniform (brightness, color, textures,...). Among these methods are the simple thresholding, region growing, region merging, but also a number of energy minimization methods.

Many segmentation methods exploit discontinuities or regional homogeneity. Yet, they are based on **different representations of the solution** (discrete vs. continuous, explicit vs. implicit) and on **different numerical solutions** (PDEs, maximum-flow algorithms, stochastic sampling,...).



Edge-based Segmentation in 8 Steps

Many strategies have been proposed to aggregate gradient based edge information into a coherent segmentation.

For example the following:

- 1 Identify edges by thresholding the gradient norm,
- 2 thin out regions (\rightarrow 1-dim. structures),
- 3 expand contour pieces (\rightarrow close gaps),
- 4 identify connected components,
- 5 eliminate smaller regions,
- 6 thin out regions (again),
- 7 introduce new boundary pixels (\rightarrow close gaps),
- 8 eliminate smaller regions.

W. A. Perkins, IEEE Trans. on PAMI 1980




Edge-based Segmentation: Laplace Zero Crossings

A strategy to automatically generate **closed curves** is to determine the zero crossings of the Laplacian of the image f :

$$\Delta(G_\sigma * f) = 0.$$

where G_σ denotes Gaussian smoothing of width σ .



Laplace zero-crossings for increasing σ (Author: D. Cremers)



Thresholding

The simplest method to compute segmentations is to threshold the input image $f : \Omega \rightarrow \mathbb{R}$:

$$g(x) = \begin{cases} 1, & f(x) > \theta \\ 0, & \text{else} \end{cases}$$

The threshold $\theta \in \mathbb{R}$ must be chosen appropriately.

For many images one can automatically determine appropriate thresholds by selecting **minima of the smoothed histogram of brightness values**.

Otsu proposed to select the threshold such that the **brightness variance of object and background are minimized**:

N. Otsu, *A Threshold Selection Method from Gray-Level Histograms*, IEEE Transactions on Systems, Man, and Cybernetics, vol. 9, no. 1, pp. 62-66, 1979.

Matlab: `level=graythresh(I);`



Thresholding



Image Segmentation

Brightness, Color,
Texture, Motion, Shape


Basic Concepts

Edge-based
Segmentation


Region-based
Segmentation

The Watershed
Transform


Summary




Input image




histogram




$\theta = 100$



$\theta = 150$



$\theta = 200$



Otsu: $\theta = 165.5$

Author: D. Cremers

Alternative Methods

There exists a number of adaptive thresholding techniques:

- Determine θ as the **minimum of the (smoothed) histogram**.
- **Otsu**: Determine θ by minimizing brightness variance.
- Set $\theta =$ **average brightness**.
- **Spatially adaptive**: $\theta =$ average brightness in the vicinity.
- **Clustering**: Determine mean intensity of inside (μ_1) and outside (μ_2):

$$\theta = \frac{\mu_1 + \mu_2}{2}$$

recompute segmentation and iterate.

- **Double thresholding**: For two threshold values $\theta_1 < \theta_2$ determine:
 - 1 all pixels i for which $f_i > \theta_2$ and
 - 2 all pixels j with $f_j > \theta_1$, which are connected with a pixel i (for which $f_i > \theta_2$).



Thresholding: Advantages and Drawbacks

Advantages:

- + Very fast to compute.
- + Adaptive variants allow to
 - + adapt to brightness distribution of the image,
 - + take into account brightness variance in the separated regions.

Drawbacks of thresholding:

- Thresholding methods neglect **spatial context**, for example the information that neighboring pixels are likely to be part of the same region. Instead all pixels are treated independently.
- Somewhat limited generality. A systematic (and mathematically more transparent) generalization of adaptive thresholding methods is given by **clustering methods**.



Segmentation by Color Clustering

Idea: Compute segmentations of an image by combining pixels of “similar” color in a single region.

Consider the colors of all pixels as **samples in the rgb color space \mathbb{R}^3** and apply a clustering algorithm.

There exist many possible clustering algorithms (see for example the Matlab help menue).

Among the best known methods is **k-means clustering**. This method determines a given number of k clusters by iteratively assigning data points to the nearest of k cluster centers and subsequently recomputing these cluster centers.

An extension of this amounts to fitting each cluster with a **multivariate Gaussian distribution** (i.e. an ellipsoid), thereby allowing adaptive stretching (in color space).



Segmentation with k-means Clustering



Input



Otsu threshold



2 clusters



5 clusters



5 cluster reconstr.



10 cluster reconstr.

Author: D. Cremers

The last two images are obtained by assigning each region the color of its cluster center. → color quantization



Region-based Segmentation

Region-based segmentation methods are related to thresholding and clustering methods. In addition, they consider (explicitly or implicitly) **spatial context**.

The central idea is as follows: Partition the image plane $\Omega \subset \mathbb{R}^2$ into n **pairwise disjoint regions** $\{\Omega_1, \dots, \Omega_n\}$ such that:

- (i) $\bigcup_{i=1}^n \Omega_i = \Omega$,
- (ii) $\Omega_i \cap \Omega_j = \emptyset$, if $i \neq j$.
- (iii) $P(\Omega_i) = \text{TRUE} \forall i$, $P(\Omega_i \cup \Omega_j) = \text{FALSE} \forall i \neq j$.

In addition, one typically assumes that neighboring pixels are preferably part of the same region (**spatial context**).

Condition (iii) states that region Ω_i should be **homogeneous with respect to some property P** . Any union of two regions, on the other hand, should not be.



Region Growing and Region Merging

Two rather old methods for computing region-based segmentations are region growing and region merging.

Region Growing (Adams & Bischof, PAMI '94):

- 1 Select a seed pixel.
- 2 Iteratively include neighboring pixels as long as their color is sufficiently similar.

Region Merging (Brice, Fennema 1970, Koepfler et al. '95):

- 1 Start with a partitioning for which each pixel is its own region.
- 2 Iteratively merge neighboring regions as long as they are sufficiently similar.


Advantages: interactive, always determines connected regions.

Drawbacks: Threshold values (“sufficiently similar”) needed, typically lack a systematic optimization criterion.



The Watershed Transform

An (edge-based) segmentation can be computed using the watershed transform (Wasserscheidentransformation).



(source: P. Soille 1998)

Intuition:

- Interpret the gradient of the smoothed image $|\nabla f_\sigma(x)| \equiv |\nabla G_\sigma * f|$ as a height profile.
- For each point let water drops go down-hill until they fill a basin (steepest descent).
- All points going to the same basin form a connected region (linear complexity).




The Watershed Transform




Input



$\sigma = 2.5$



$\sigma = 5$



$\sigma = 7$



$\sigma = 10$



$\sigma = 10$ Reconstr.

Author: D. Cremers


The last image is obtained by coloring respective regions with their average color value.




Watershed with Nonlinear Diffusion

One drawback of the watershed transform is that it typically leads to an **over-segmentation** (too many small regions).


On the other hand, a simple presmoothing delocalizes semantically important edge information. A better alternative is nonlinear presmoothing (for example with **Perona-Malik diffusion**).



Perona-Malik



watershed




segmentation

Author: D. Cremers



Watershed with Nonlinear Diffusion



Diffused



1104 regionen

Author: D. Cremers

Image Segmentation

Brightness, Color,
Texture, Motion, Shape

Basic Concepts

Edge-based
Segmentation


Region-based
Segmentation

The Watershed
Transform


Summary

Clustering with Nonlinear Diffusion

Nonlinear diffusion filtering can also be used as preprocessing in order to introduce spatial context into color clustering.



Diffused



10 cluster

Author: D. Cremers



Summary

- We have seen a number of edge-based and region-based segmentation methods (**Laplace zero-crossings, thresholding, clustering, region growing, region merging, watershed**).
- **Edge-based methods** exploit **brightness discontinuities** as a criterion of region boundaries. The integration of such edge information to coherent closed curves is an algorithmic challenge.
- **Region-based methods** exploit **color similarity** as a criterion for grouping pixels into coherent regions.
- The last methods (region growing, merging, watershed) also integrate some kind of **neighborhood information**, albeit in a somewhat heuristic manner.
- The above methods (except for clustering) **lack a mathematically transparent optimization criterion**. This is where optimization methods and statistical approaches enter the picture.

