

Lecture 23: Segmentation II: Optimisation Methods

Contents

1. Disadvantages of Previous Segmentation Methods
2. Mumford–Shah Cartoon Model
3. Theoretical Properties
4. Approximation Method
5. Algorithmic Realisation
6. Experiments
7. Active Contour Model of Chan and Vese
8. Extensions of the Chan–Vese Model

© 2001–2008 Joachim Weickert

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

Disadvantages of Previous Segmentation Methods

Disadvantages of Previous Segmentation Methods

- ◆ selection of multiple parameters (e.g. for double thresholding)
- ◆ sometimes less favourable results:
 - unsmooth region boundaries (e.g. for watersheds)
 - noise and too small structures create their own regions (e.g. for region merging)
- ◆ removal of less favourable results requires postprocessing steps (and hence new parameters)
- ◆ no tool for comparing the quality of different segmentations

Is there a systematic approach that avoids these shortcomings and requires only one single parameter (scale parameter)?

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

The Mumford–Shah Cartoon Model

- ◆ Let $f : \Omega \rightarrow \mathbb{R}$ be a continuous initial image.
- ◆ goal: segmentation (u, K) , where u is constant in each segment and K represents the set of segment boundaries (“cartoon”).
- ◆ simple variational ansatz: segmentation as minimiser of the energy functional

$$E(u, K) = \int_{\Omega \setminus K} (u - f)^2 dx dy + \lambda l(K),$$

where $l(K)$ measures the length of the segment boundaries and $\lambda > 0$ is a scale parameter.

- ◆ Term 1 controls the homogeneity within a segment, while term 2 optimises the length of the boundaries. Increasing λ makes it more attractive to have larger segments (at the expense of homogeneity).
- ◆ sufficiently complicated to represent basic features of a good segmentation; sufficiently simple to allow a thorough mathematical analysis

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

Theoretical Properties

Theoretical Properties

◆ Preservation of Mean

Term 1 implies that, within one segment, u has the same mean as f . Thus, given f , the segmentation is uniquely determined by the boundary set K , and we may write $E(K)$ instead of $E(u, K)$.

◆ Existence of a Global Minimiser, Regularity of the Boundaries

There exists (at least) one segmentation that minimises $E(K)$. The segment boundaries are either regular (at least C^1), or they are singular and of one of the following types:

- Three edges meet at an angle of 120 degrees.
- An edge meets the boundary of the image domain (at 90 degrees).

◆ Many Local Minimisers

The energy landscape may have numerous local minimisers. There are two options:

- very slow minimisation methods that guarantee to find the global minimum asymptotically (*simulated annealing*).
- more efficient approximations that usually find reasonable local minima but cannot guarantee to end up in a global minimum

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

An Approximation Method for Minimisation

In order to find good local minima of $E(u, K)$ in an efficient way, consider so-called 2-normal segmentations.

Definition (2-normal segmentation):

A segmentation K is called *2-normal*, if every segmentation K' , that results from merging two adjacent segments of K , has a larger energy:

$$E(K') > E(K).$$

Restricting ourselves to 2-normal segmentations allows to prove a number of useful properties.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

What are the Properties of 2-Normal Segmentations ?

◆ Limitation of the Number of Regions

$$n \leq c_0(\Omega) \frac{|\Omega| (\sup f - \inf f)^4}{\lambda^2}$$

where $|\Omega|$ denotes the area of the image domain Ω , and c_0 is a known constant that depends only on the geometry of Ω .

This shows:

- The larger the domain Ω , the more segments can be expected.
- High image contrast encourages many segments.
- Increasing λ encourages less segments.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

◆ Elimination of Small Regions

The area of each segment Ω_i , $i = 1, \dots, n$ is bounded from below:

$$|\Omega_i| \geq c_1(f, \lambda, \Omega).$$

◆ Elimination of Thin Regions

The length $l(\partial\Omega_i)$ of a segment boundary $\partial\Omega_i$ cannot become arbitrary large compared to the square root of its area:

$$l(\partial\Omega_i) \leq c_2(f, \Omega) |\Omega_i|^{1/2}$$

Since the segment area $|\Omega_i|$ is bounded by $|\Omega|$, there exists an upper limit for the length of a segment boundary.

◆ Smoothness of the Segment Boundaries

The segment boundaries are C^1 almost everywhere.

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

Algorithmic Realisation

Basic Idea:

- ◆ approximate global minimiser of cartoon model by local minimiser resulting from the restriction to 2-normal segmentations
- ◆ create 2-normal segmentations by region merging
- ◆ perform region merging by successively increasing the scale parameter λ

Useful Observation

- ◆ Merging two segments Ω_i, Ω_j with means u_i, u_j does not require to recompute the entire energy functional.
- ◆ One can show that the new energy is given by

$$E(K - \partial(\Omega_i, \Omega_j)) = E(K) + \frac{|\Omega_i| \cdot |\Omega_j|}{|\Omega_i| + |\Omega_j|} (u_i - u_j)^2 - \lambda l(\partial(\Omega_i, \Omega_j))$$

where $\partial(\Omega_i, \Omega_j)$ is the joint border between Ω_i and Ω_j .

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

The Region Merging Algorithm

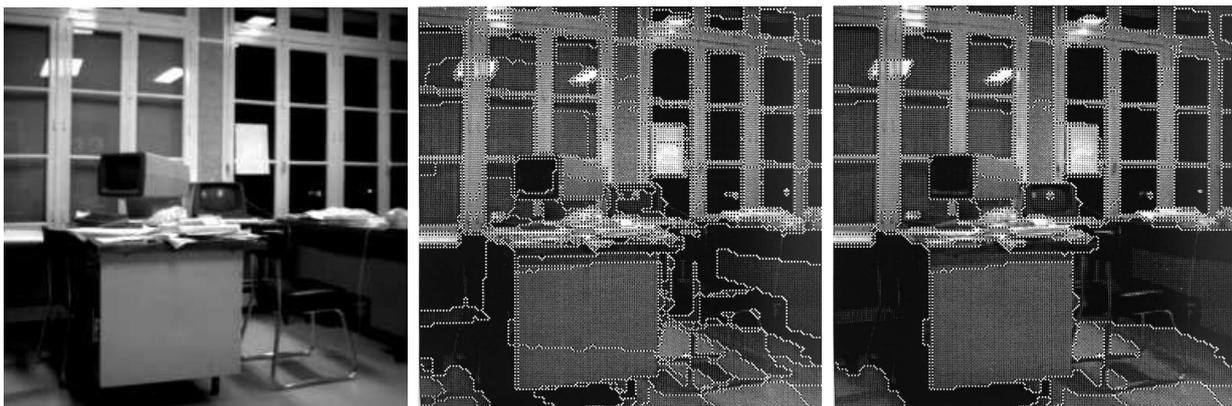
1. Initialisation: Each pixel is one segment.
2. For each pair of neighbouring pixels, one computes the λ value, for which a merging decreases the energy.
3. Merge the pair with the smallest λ value.
4. Compute the λ values for merging the new segment with its neighbouring segments.
5. Repeat the steps (2)–(4), until a desired number of regions or a specified value for λ is reached.

Important Properties

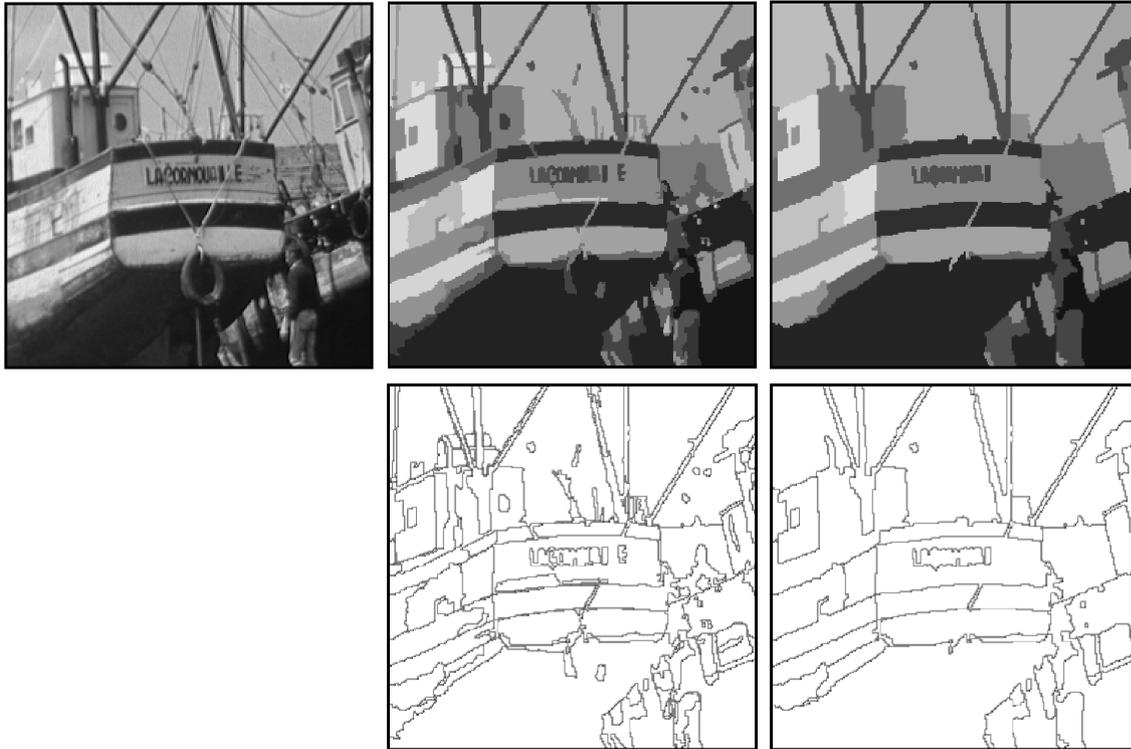
- ◆ *Fixed Point Property:*
For a piecewise constant image there exists some λ_0 such that it is a valid segmentation for all $\lambda < \lambda_0$.
- ◆ *Pyramid Property, Causality:*
For $\lambda_2 > \lambda_1$, the boundaries of a λ_2 -segmentation are a subset of the boundaries of a λ_1 -segmentation.

Experiments (1)

Experiments



(a) Left: Original image. (b) Middle: Segmented with $\lambda = 1024$. (c) Right: $\lambda = 4096$. Authors: G. Koepfler, J.-M. Morel, S. Solimini (1990).



(a) **Top left:** Original image. (b) **Top middle:** Segmented with $\lambda = 2022$. (c) **Top right:** Segmented with $\lambda = 6173$. (d) **Bottom middle:** Segmentation boundaries of (b). (e) **Bottom right:** Segmentation boundaries of (c). Author: G. Koepfler (1995).

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

The Active Contour Model of Chan and Vese (1)

The Active Contour Model of Chan and Vese

- ◆ Chan and Vese considered the cartoon model under the restriction that the image is segmented into two phases only.
- ◆ The energy functional then becomes a function of a contour C :

$$E(C) = \int_{\text{inside } C} (f(x,y) - u_1)^2 dx dy + \int_{\text{outside } C} (f(x,y) - u_2)^2 dx dy + \lambda l(C)$$

where u_1 and u_2 are the arithmetic means of $f(x,y)$ inside/outside C .

- ◆ The user specifies some initial contour C_0 close to the segmentation one wants to find.
- ◆ Minimising $E(C)$ drives the curve towards the segment boundaries (using so-called level set methods that even allow a contour to split; we omit the technical details).

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18



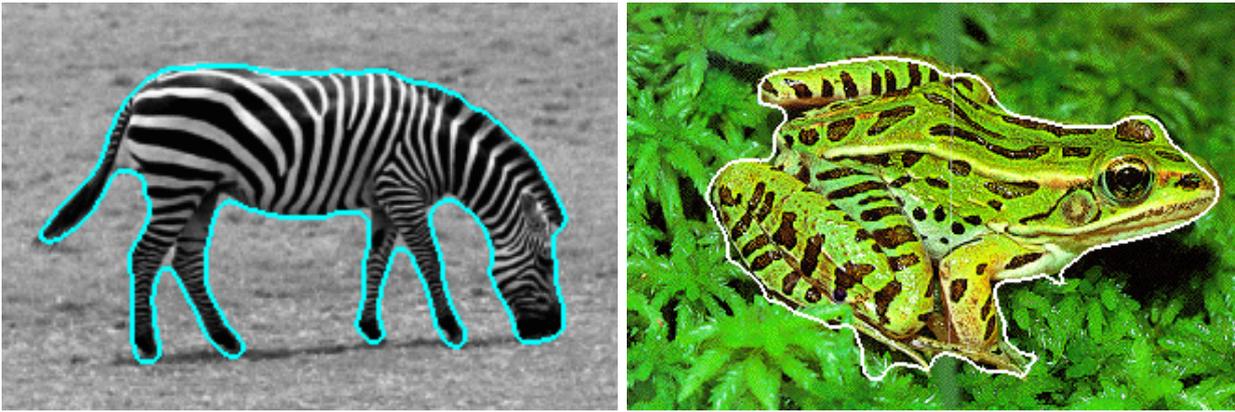
Top left to bottom right: Evolution of an active contour under the Chan-Vese model. The contour splits and creates a fairly good segmentation of Europe at night. Authors: T. Chan, L. Vese (2001).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

Extensions of the Chan–Vese Model

- ◆ recent modifications of the Chan–Vese model include
 - more sophisticated features (colour, texture, motion, ...)
 - additional statistical characterisations of a region (not only mean, but also standard deviation, ...)
 - multiphase models that create more than two types of segments
 - sometimes even a-priori knowledge using a statistical characterisation of the shapes to be expected (Lecture 30)
- ◆ yield state-of-the-art results in segmentation

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18



(a) **Left:** Segmented image using grey values and texture as features. (b) **Right:** Segmentation with colour and texture as feature channels. Author: T. Brox (2003).

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

Summary (1)

Summary

- ◆ The cartoon model of Mumford and Shah is an energy based segmentation method that rewards homogeneity within each region and a small length of segmentation boundaries.
- ◆ requires only a single parameter that is related to the scale
- ◆ relatively simple and mathematically well understood
- ◆ restriction to 2-normal segmentations automatically generates many desirable features that have to be engineered in an ad hoc way for most other methods: The segments are neither too many, too small, too thin, nor do they have too irregular boundaries.
- ◆ algorithmic realisation via region merging
- ◆ Active contour variants exist that allow interactive segmentation of high quality.

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

Summary (2)



Literature

- ◆ G. Koepfler: *Regionenwachstumsverfahren und Variationsmodelle*. Vortragsausarbeitung (Heidelberg, Nov. 1995). www.math-info.univ-paris5.fr/~gk/papers/heidelberg95.pdf (*well-readable overview*)
- ◆ J.-M. Morel, S. Solimini: Segmentation of images by variational methods: a constructive approach. *Revista Matematica de la Universidad Complutense de Madrid*, Vol. 1, No. 1–3, pp. 169–182, 1988. (*theoretical analysis of the cartoon model*)
- ◆ G. Koepfler, C. Lopez, J.-M. Morel: A multiscale algorithm for image segmentation by variational method. *SIAM Journal on Numerical Analysis*, Vol. 31, pp. 282–299, 1994. (*2-normal segmentation by region merging*)
- ◆ J.-M. Morel, S. Solimini: *Variational Methods in Image Segmentation*. Birkhäuser, Basel, 1994. (*an entire book on Mumford-Shah segmentation*)
- ◆ D. Mumford, J. Shah: Boundary detection by minimizing functionals, I. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (San Francisco, CA, June 1985), pp. 22–26, 1985. (*introduced the Mumford-Shah functional*)
- ◆ D. Mumford, J. Shah: Optimal approximation by piecewise smooth functions and associated variational problems. *Communications in Pure and Applied Mathematics*, Vol. 42, pp. 577–685, 1989. (*very detailed mathematical analysis*)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18

Summary (3)



- ◆ T. F. Chan, L. A. Vese: Active contours without edges. *IEEE Transactions on Image Processing*, Vol. 10, No. 2, pp. 266–277, Febr. 2001. (*Chan–Vese model*)
- ◆ T. Brox, M. Rousson, R. Deriche, J. Weickert: Unsupervised segmentation incorporating colour, texture, and motion. In N. Petkov, M. A. Westenberg: *Computer Analysis of Images and Patterns*. Springer LNCS Vol. 2756, pp. 353–360, Berlin, 2003. (*extensions of the Chan–Vese model*)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18