

Introducing Curvature into Globally Optimal Image Segmentation: Minimum Ratio Cycles on Product Graphs

Thomas Schoenemann and Daniel Cremers
Department of Computer Science
University of Bonn, Germany

Abstract

While the majority of competitive image segmentation methods are based on energy minimization, only few allow to efficiently determine globally optimal solutions. A graph-theoretic algorithm for finding globally optimal segmentations is given by the Minimum Ratio Cycles, first applied to segmentation in [8]. In this paper we show that the class of image segmentation problems solvable by Minimum Ratio Cycles is significantly larger than previously considered. In particular, they allow for the introduction of higher-order regularity of the region boundary.

The key idea is to introduce an extended graph representation, where each node of the graph represents an image pixel as well as the orientation of the incoming line segment. With each graph edge representing a pair of adjacent line segments, edge weights can depend on the curvature. This way arbitrary positive functions of curvature can be introduced into globally optimal segmentation by Minimum Ratio Cycles. In numerous experiments we demonstrate that compared to length-regularity the integration of curvature-regularity will drastically improve segmentation results.

Moreover, we show an interesting relation to the Snakes functional: Minimum Ratio Cycles provide a way to find one of the few cases where the Snakes functional has a meaningful global minimum.

1. Introduction and Related Work

Energy minimization is a popular means for image segmentation. While many approaches in the field penalize certain regularity properties of the curve, only a small percentage considers the curvature. This is surprising as curvature-based segmentation can be much closer to human perception than simple length regularization. Kaniza's experiments on illusory contours [9], for example, indicate that humans tend to perform line completion by finding closed curves of minimal curvature. Figure 1 shows segmentation results on test images which demonstrate that curvature-

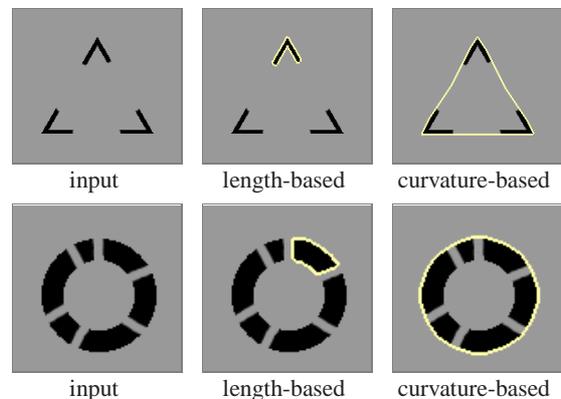


Figure 1. Effects of length-based and curvature-based regularization in image segmentation on artificial images. Results were generated using functionals (1) and (7), the latter with absolute curvature ($q = 1$). Note that curvature-regularity gives rise to fundamentally different segmentations and contour completion.

based regularization may have favorable properties with respect to line completion.

A large body of literature is dedicated to curvature-based regularity in areas such as image inpainting [12, 3] and perceptual grouping [14, 18]. Such approaches, however, are based on local optimization, either in a discrete setting or using partial differential equations and curve evolution. In the latter case, the resulting fourth-order terms in the Euler-Lagrange equations require highly sophisticated discretization schemes and small time steps to prevent numerical instabilities. In addition, these methods usually require appropriate initialization.

Approaches to image segmentation can be split into region-based functionals and edge-based ones. A prototype of a region-based approach is the functional of Mumford and Shah [13]. Its piecewise constant formulation can be seen as a spatially continuous generalization of the Ising model [7]. Further restricted to two phases with known mean intensities, the discrete version can be optimized globally via graph cuts [5]. We are aware of no other region-based approach where global optima are found efficiently.

Edge-based approaches are often line integrals along the

region boundary. In their seminal work, Kass et al. [10] proposed to integrate the negative squared image gradient plus the squared absolute of first and second derivative of the curve. While this effectively imposes length and curvature as regularity terms, the *global* minimum of this energy is typically not useful: for most weighting parameters it is given by an infinitesimally short curve or an infinitely long one. The choice of sensible parameters depends highly on the image.

Casselles et al. [2] proposed to drop the regularity terms and replace the negative squared absolute of the gradient by an edge indicator function $g(\cdot)$ which is a monotonically decreasing function of the image gradient. It is always strictly positive. Again the global minimum (with energy zero) is given by an infinitesimally short curve. Meaningful global optima can be determined when seed points are given for foreground and background [1].

Closely related to this are the normalized cuts of Shi and Malik [15]. Their approach works with very few seed pixels already. The functional is not minimized globally as the problem is NP-hard (see [15], Appendix A). Instead, a relaxed version is solved by spectral methods.

With their ratio regions, Cox et al. [4] proposed an edge-based functional where a useful global minimum can be found in polynomial time. In practice the complete search over all starting points takes prohibitively long. Integrating curvature appears very hard as the underlying graph theoretic approach is limited to planar graphs.

In their seminal work Jermyn and Ishikawa [8] proposed to consider the ratio of two line integrals. Many functionals of this form have a sensible global minimum. Moreover, thanks to a graph theoretic approach this minimum can be found efficiently. In subsequent works the approach was applied for the simultaneous segmentation of multiple images [16, 6] and to connecting pre-extracted line segments [17]. The latter considered curvature, but only for the pre-extracted segments.

To date the approach did not support the estimation of curvature integrals of entirely unknown curves *during* optimization. Exactly this will be developed in this paper. To the best of our knowledge this is the first time that a curvature-dependent regularization is imposed in globally optimal image segmentation.

2. Segmentation by Searching Region Boundaries

Given an image $I : \Omega \rightarrow \mathbb{R}$ we consider the task of segmenting the image plane Ω into a connected foreground region and background. Equivalently one can search for the region boundary, a closed curve $C : \mathbb{S}^1 \rightarrow \Omega$, where \mathbb{S}^1 is the unit sphere in the two-dimensional plane. To ease notation, throughout this paper the curve is assumed to be

parameterized uniformly. While we cannot exclude self-intersecting curves from the space of permissible solutions they did not arise in any of our experiments.

2.1. Energies for Region Boundaries

Finding the optimal region boundary can be stated as an energy minimization problem: each conceivable closed curve is assigned an energy, the curve with minimal energy is considered the optimal boundary.

A popular approach is to formulate the functional as an integral of costs along the curve. Globally minimizing such energies is usually an either trivial or ill-defined task: If the costs are positive everywhere, the minimum (with zero costs) is given by an infinitesimally short curve. If the costs are negative in a small region a minimum does not exist as the curve would grow infinitely in this region.

As a remedy, Jermyn and Ishikawa [8] proposed to minimize a ratio of two line integrals, a well-defined task for many functionals: Consider the task of minimizing

$$\min_C \frac{\int_{\mathbb{S}^1} \nabla I(C(s)) \cdot \vec{n}_C(s) ds}{|C|} \quad (1)$$

which will be referred to as *length ratio* in the following. Here $|C|$ denotes the length of the curve, “ \cdot ” the scalar product of two vectors and $\vec{n}_C(s)$ is the curve normal at point $C(s)$. When a curve is traversed in opposite orientation all normals are inverted. Hence all scalar products change sign and one obtains the same energy with opposite sign. This implies a negative minimum and that the absolute of (1) is effectively maximized. This task is usually well-defined: for a continuous image gradient the numerator terms integrated over a small closed curve will essentially cancel, giving rise to energy values close to zero. More generally Jermyn and Ishikawa considered energies of the form

$$\min_C \frac{\int_{\mathbb{S}^1} \vec{v}(C(s)) \cdot \vec{n}_C(s) ds + \int_{C_{in}} f(x) dx}{\int_{\mathbb{S}^1} g(C(s)) ds} \quad (2)$$

Here $\vec{v} : \Omega \rightarrow \mathbb{R}^2$ is an arbitrary vector field and $g : \Omega \rightarrow \mathbb{R}^+$ is strictly positive. The function $f : \Omega \rightarrow \mathbb{R}$ is integrated over the area enclosed by the curve.

2.2. Region Terms

In most applications region terms are not useful. For one thing the arising minimizations task are often ill-defined, e.g. there is no minimum for the ratio of area over boundary length. Moreover, the algorithm in [8] can only handle *oriented* areas, which leads to problems with self-intersecting curves as illustrated in figure 2. We will therefore not consider region terms in this paper.

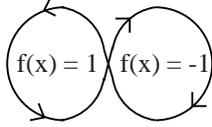


Figure 2. Region integrals are hard to handle for self-intersecting curves: When using the approach in [8] the integral of f is calculated as ± 2 although actually 0 is correct. This is as all areas are *oriented*. When this is not desired one has to make sure that the incorrect costs assigned to self-intersecting curves are always higher than the global optimum of the non self-intersecting ones.

3. Introducing Curvature-regularity

Minimizing functionals of form (2) is generally a well-defined task. However, in practice usually very small regions are found: the optimal curve needs strong numerator terms along the entire curve. If a region boundary passes through an area of low contrast, it will most likely not be assigned the optimal energy. Typically this results in very small curves enclosing a few pixels only.

In this paper we therefore propose to integrate the curvature of the curve in the ratio functional, thereby enabling meaningful segmentations even in the presence of low-contrast areas. Consider the functional

$$\min_C \frac{\int_{\mathbb{S}^1} \nabla I(C(s)) \cdot \vec{n}_C(s) ds}{\int_{\mathbb{S}^1} |\kappa_C(s)| ds} \quad (3)$$

where $\kappa_C(s)$ denotes the curvature of the curve C at $C(s)$. This minimization task is well-defined as a closed curve must have non-zero curvature somewhere, implying a strictly positive denominator. Now, for any straight line the numerator terms are accumulated along the line while the denominator term stays constant. A straight line segment passing a region of zero contrast will therefore not affect the total energy. This is demonstrated in Figure 1.

In this paper we consider a class of functionals which drastically generalizes functional (3): in the numerator and denominator integrals we consider functions depending on the location in the image $C(s)$, the tangent angle $\alpha_C(s)$ and the curvature $\kappa_C(s)$. That is, we solve minimization problems of form

$$\min_C \frac{\int_{\mathbb{S}^1} h(C(s), \alpha_C(s), \kappa_C(s)) ds}{\int_{\mathbb{S}^1} g(C(s), \alpha_C(s), \kappa_C(s)) ds} \quad (4)$$

The numerator function $h: \Omega \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ can be arbitrary, the denominator function $g: \Omega \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_0^+$ is non-negative. While $g(\cdot)$ may become 0 in some places, for any closed curve the integral of $g(\cdot)$ along the curve must be strictly positive.

4. Estimating Curvature via Product Graphs

To solve problem (4) it is discretized and reduced to the problem of searching cycles with minimal ratio in a graph.

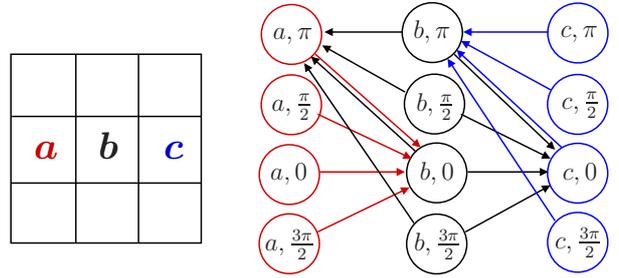


Figure 3. A sample image and a part of its associated graph (using a 4-neighborhood). A node in the graph encodes a pixel in the image together with the incoming direction. E.g. when going from pixel b to c the direction is 0 and one ends in the node $(c, 0)$. This construction allows for edge weights depending on curvature.

In this section we describe the graph, then give a brief description of the optimization algorithm in the next one.

4.1. Graph Construction

In the discretized problem the region boundary is composed of a finite set of pre-defined straight line segments. These are obtained by connecting each pixel to all pixels within a certain neighborhood of the pixel. The precise neighborhood relation is described in section 4.2.

This discrete problem can be expressed via a graph where each edge e is assigned a numerator weight $n(e)$ and a denominator weight $d(e)$ reflecting the numerator and denominator integrals in (4) *along the line segment* respectively. The curvature is considered fixed for the entire line segment. Minimizing the discretized ratio (4) is now equivalent to searching a cycle in the graph which minimizes

$$\min_C \frac{\sum_{e \in C} n(e)}{\sum_{e \in C} d(e)} \quad (5)$$

The key challenge is that for the estimation of curvature one needs to consider *pairs* of adjacent line segments, but the algorithm only supports weights defined for *single* edges. Hence edges cannot directly correspond to line segments as in [8]. To enable access to the previous line segment, it is stored in the nodes. More precisely its direction is stored as this suffices for the estimation of curvature. The node set is then

$$\mathcal{V} = \mathcal{P} \times \mathcal{A}$$

where \mathcal{P} is the set of pixels and \mathcal{A} is a finite set of directions of the incoming line segment, expressed as tangent angles between 0 and 2π . While this product space was used for several works involving curvature [14], to the best of our knowledge it is novel to path-based search.

Edges are constructed as follows: for a line segment from a pixel \vec{p} to a pixel \vec{q} in the image, there are several edges in the graph. All these edges end in the node representing the pixel \vec{q} as well as the tangent angle of the line

segment. They start in the various nodes corresponding to \vec{p} for all incoming directions of \vec{p} , as illustrated in Figure 3. Hence edges now correspond to *pairs* of line segments in the image: the current one connects \vec{p} and \vec{q} , the previous is given by the tangent angle stored in the root node of the edge.

The curvature of the pair of line segments corresponding to an edge from (\vec{p}, β) to (\vec{q}, α) is approximated by

$$\kappa(e = ((\vec{p}, \beta), (\vec{q}, \alpha))) = \frac{1}{l(\alpha) + l(\beta)} \cdot \begin{cases} \beta - \alpha & \text{if } |\beta - \alpha| \leq \pi \\ 2\pi - \beta + \alpha & \text{if } \beta - \alpha > \pi \\ -2\pi - \beta + \alpha & \text{if } \beta - \alpha < -\pi \end{cases}$$

where $l(\alpha)$ denotes the length of a line segment with tangent angle α . Notice that the edge weights $n(e)$ and $d(e)$ reflect the numerator and denominator integrals along the *current* line segment. The tangent angle of the previous line segment is only used to calculate the curvature. The numerator edge weights are computed as (abbreviating $\kappa(e)$ as κ)

$$n(e = ((\vec{p}, \beta), (\vec{q}, \alpha))) = |\vec{p} - \vec{q}| \cdot \left[\frac{1}{4}h(\vec{p}, \alpha, \kappa) + \frac{1}{2}h\left(\frac{\vec{p} + \vec{q}}{2}, \alpha, \kappa\right) + \frac{1}{4}h(\vec{q}, \alpha, \kappa) \right]$$

and – replacing $h(\cdot)$ by $g(\cdot)$ – likewise for the denominator weights $d(e)$. This discretization scheme balances two issues: Mathematically the tangent angle is not well-defined at the beginning and end of a line segment, so the center is emphasized. On the other hand the curve integrals should be sampled in intervals close to the pixel width. As some of our line segments skip several pixels, multiple samples along the edge are necessary.

4.2. Dealing with Large Graphs

To well approximate the continuum we need a sufficiently large set of directions \mathcal{A} . Directions are obtained by considering all line segments connecting pixels in a certain neighborhood¹. Due to the discretization on a uniform grid directions are not sampled in equidistant intervals. The size of the neighborhood determines the spatial scale at which curvature is approximated and with it the number of directions considered. We use a neighborhood of radius $\sqrt{13}$, corresponding to 32 directions.

With about $N \cdot |\mathcal{A}|$ nodes and $N \cdot |\mathcal{A}|^2$ edges for an image with N pixels, the resulting graphs are rather large. Already for an image of size 256×256 the available 2 GigaByte were exceeded. In such cases we do not store edges explicitly. Rather a node computes its edge list on demand. The same graph then fits into 140 MegaByte. The corresponding increase in run-time was about a factor of 4.

¹ We only consider those pixels where in the same direction there is no pixel with smaller distance.

5. Optimization of Ratio Functions

To minimize ratios of form (5) we use the variant in [8] of the Minimum Ratio Cycle algorithm [11]. It can be shown that a graph with same topology but with edge weights $w(e) = n(e) - \tau d(e)$ must possess a negative cycle for all ratios $\tau > \tau_{opt}$, with τ_{opt} the optimal ratio. Negative cycle detection can be done efficiently using distance calculations via the Moore-Bellman-Ford algorithm. The algorithm also returns a negative cycle if existent.

Starting with an upper bound on the ratio (e.g. the ratio of any cycle in the graph), the ratio is repeatedly corrected to the ratio of the found negative cycle until all negative cycles vanish. The last extracted cycle has optimal ratio. For integral edge weights $n(e)$ and $d(e)$ the algorithm terminates in a finite number of steps and has pseudo-polynomial complexity². If – as in all our experiments – the maximal absolute edge weight does not depend on the number of nodes, it has polynomial complexity. In practice the algorithm always terminated after less than 50 ratio adjustments. See [8, 11] for more details.

6. Minimum Ratio Cycles and Snakes

In [8] Jermyn and Ishikawa pointed out that the Minimum Ratio Cycle algorithm effectively minimizes a cost function consisting of the sum of two terms, where the second is weighted such that the optimal energy is exactly 0. With the extended class of functionals introduced in this paper, this allows us to shed a new light on the well-known Snakes functional [10]: Consider the task of minimizing

$$\tau_{opt} = \min_C \frac{\int_{\mathbb{S}^1} [\beta |C_s(s)|^2 + |\kappa_C(s)|^2] ds}{\int_{\mathbb{S}^1} |\nabla I(C(s))|^2 ds} \quad (6)$$

For this functional the proposed algorithm finds a τ so that

$$\min_C \int_{\mathbb{S}^1} [\beta |C_s(s)|^2 + |\kappa_C(s)|^2] ds - \tau \int_{\mathbb{S}^1} |\nabla I(C(s))|^2 ds = 0$$

This is essentially the Snakes energy³. Given a weighting factor for the relation between the two regularity terms, the algorithm hence finds a relative weight between the data term and the regularity terms such that the optimal Snakes energy becomes 0. This allows us to find one of the few parameter configurations for which the Snakes have a non-trivial global minimum.

7. Experiments

In several experiments we demonstrate that the proposed algorithm compares favorably to the length ratio (1) and the

² The algorithm in [11] is weakly polynomial but much slower in practice.

³ More precisely, the Snakes functional is obtained by replacing the squared curvature by the squared second derivative and dividing by τ .

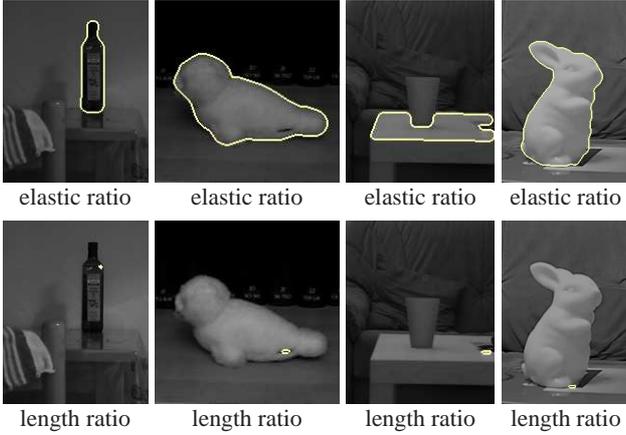


Figure 4. While length regularization typically does not yield useful regions, this is different for the elastic ratio.

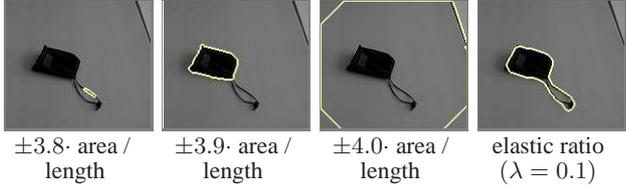


Figure 5. The length ratio enhanced by a balloon force is sensitive to weighting parameters.

Mumford-Shah functional. We use a functional of form

$$\frac{\int_{\mathbb{S}^1} \nabla I(C(s)) \cdot \vec{n}_C(s) ds}{\int_{\mathbb{S}^1} |\kappa_C(s)|^q ds + \lambda|C|} \quad (7)$$

For $q=2$ and positive λ the denominator integral is the well-known expression used in the *Elastica*. We will therefore refer to this combination as *elastic ratio*. Unless otherwise stated a length parameter of $\lambda=0.2$ was used.

7.1. Object Segmentation via the Elastic Ratio

Figure 4 shows that while the length-ratio typically finds small curves enclosing a few pixels only, curves with optimal elastic ratio often correspond to real objects in the image. As shown in Figure 5 enhancing the length ratio by a balloon force is usually not competitive compared to the elastic ratio: the resulting functional is extremely sensitive to the weighting factor for area.

Figures 6 and 7 demonstrate that segmentation with the elastic ratio is often closer to human perception than region-based approaches. Here we compare with a piecewise constant and a piecewise smooth Mumford-Shah, both with two phases. They are optimized locally via alternating minimization with graph cut segmentation. Results for different length weights are shown, for the piecewise smooth functional the smooth approximations are depicted.

For $\lambda = 0$ the denominator is minimized by convex shapes for absolute curvature and by circles for squared cur-



Figure 6. The elastic ratio often provides more meaningful (object) segmentations than the Mumford-Shah functionals (shown for different length weights).

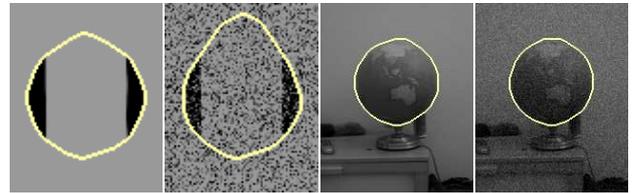


Figure 8. Regularization via squared curvature favors circle-like shapes.

vature. This effectively imposes a shape prior. The effects of absolute curvature are demonstrated in Figure 1 (page 1). For squared curvature Figure 8 indicates a certain noise-robustness.

7.2. Contrast-reversing Boundaries

The functionals 7 and 1 assign near-zero energies to contrast-reversing boundaries: if the region inside is darker than the outside in some places and lighter in others, the scalar products cancel as the gradient changes direction. To find contrast-reversing boundaries we therefore integrate the gradient absolute into a modified elastic ratio:

$$\frac{-\int_{\mathbb{S}^1} |\nabla I(C(s))|^p ds}{\int_{\mathbb{S}^1} |\kappa_C(s)|^q ds + \lambda|C|} \quad (8)$$

Figure 9 demonstrates the difference to the elastic ratio. For $p = q = 2$ minimizing (8) (and hence maximizing its absolute) is equivalent to minimizing the Snakes ratio (6): maximizing a ratio is equivalent to minimizing its reciprocal⁴.

Conclusion

In this paper we developed a framework which allows to integrate curvature into global shape optimization. The key idea is to introduce product graphs with edge weights corresponding to *pairs* of line segments in the image. This allows to represent discrete approximations of curvature in

⁴The Snakes integrate the *squared* absolute of the first derivative. However, for uniform parameterizations this is equivalent to considering the length of the curve.

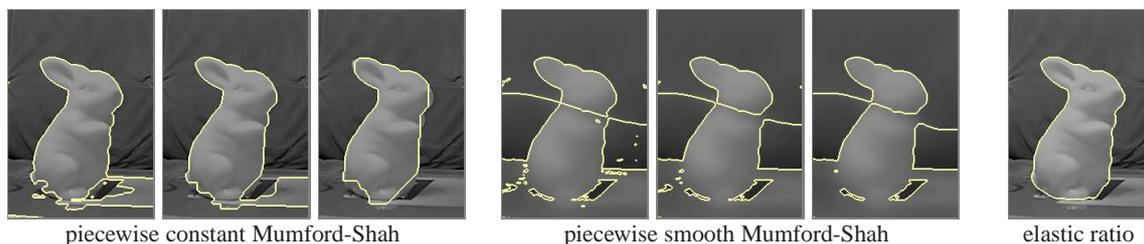


Figure 7. The elastic ratio is often more useful for object-based segmentations than region-based approaches (shown for different length weights).

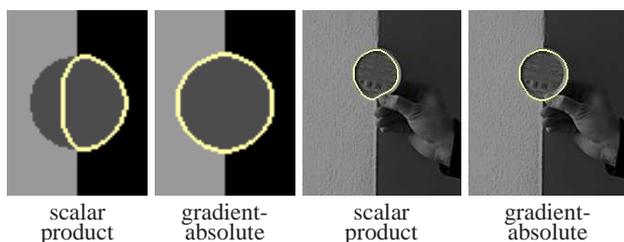


Figure 9. The gradient-absolute allows to find contrast-reversing boundaries ($p = 1, q = 2, \lambda = 0.1$).

the graph structure. Globally optimal segmentations are computed efficiently using Minimum Ratio Cycles.

Experiments on artificial and real images demonstrate that the proposed curvature regularization allows to generate contour completion in the sense that it finds segmenting boundaries which minimize the average curvature. Moreover, it provides more robust and meaningful segmentations on many real images than the previously considered length regularization. From a theoretical point of view, the proposed framework allows for the first time to study the effects of higher order regularity in a global shape optimization framework.

Acknowledgments

This work was supported by the German Research Foundation, grant #CR-250/1-1. We thank Mathias Hauptmann, Simon Masnou and Frank R. Schmidt for helpful discussions.

References

- [1] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *IEEE Int. Conf. on Comp. Vision*, volume 1, pages 26–33, Nice, France, 2003. 2
- [2] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int. J. of Comp. Vision*, 22(1):61–79, 1997. 2
- [3] T. F. Chan, S. H. Kang, and J. Shen. Euler’s elastica and curvature based inpaintings. *J. Appl. Math.*, 2:564–592, 2002. 1
- [4] I. J. Cox, S. B. Rao, and Y. Zhong. Ratio regions: a technique for image segmentation. In *IEEE Int. Conf. on Comp. Vision*, volume 2, pages 557–564, 1996. 2
- [5] D. M. Greig, B. T. Porteous, and A. H. Seheult. Exact maximum a posteriori estimation for binary images. *J. Roy. Statist. Soc., Ser. B.*, 51(2):271–279, 1989. 1
- [6] H. Ishikawa and I. H. Jermyn. Region extraction from multiple images. In *IEEE Int. Conf. on Comp. Vision*, Vancouver, Canada, 2001. 2
- [7] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 23:253–258, 1925. 1
- [8] I. H. Jermyn and H. Ishikawa. Globally optimal regions and boundaries as minimum ratio weight cycles. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 23(10):1075–1088, 2001. 1, 2, 3, 4
- [9] G. Kaniza. Contours without gradients or cognitive contours. *Italian Jour. Psych.*, 1:93–112, 1971. 1
- [10] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. of Comp. Vision*, 1(4):321–331, 1988. 1, 4
- [11] E. L. Lawler. Optimal cycles in doubly weighted linear graphs. In *Theory of Graphs: International Symposium*, pages 209–213, New York, USA, 1966. Gordon and Brear. 4
- [12] S. Masnou and J. M. Morel. Level-lines based disocclusion. In *Int. Conf. on Image Processing*, volume 3, pages 259–263, Chicago, USA, 1998. 1
- [13] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989. 1
- [14] P. Parent and S. W. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 11(8):823–839, 1989. 1, 3
- [15] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 22(8):888–905, 2000. 2
- [16] O. Veksler. Stereo matching by compact windows via minimum ratio cycle. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 24(12):1654–1660, 2002. 2
- [17] S. Wang, T. Kubota, J. Siskind, and J. Wang. Salient closed boundary extraction with ratio contour. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 27(4), 2005. 2
- [18] L. Williams and D. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation*, 9(4):837–858, 1997. 1