

The Elastic Ratio: Introducing Curvature into Ratio-based Image Segmentation

Thomas Schoenemann, Simon Masnou, and Daniel Cremers

Abstract—We present the first ratio-based image segmentation method which allows to impose curvature regularity of the region boundary. Our approach is a generalization of the ratio framework pioneered by Jernyn and Ishikawa so as to allow penalty functions that take into account the local curvature of the curve. The key idea is to cast the segmentation problem as one of finding cyclic paths of minimal ratio in a graph where each graph node represents a line segment.

Among ratios whose discrete counterparts can be globally minimized with our approach, we focus in particular on the elastic ratio

$$\frac{\int_0^{\mathcal{L}(C)} \nabla I(C(s)) \cdot (C'(s))^\perp ds}{\nu \mathcal{L}(C) + \int_0^{\mathcal{L}(C)} |\kappa_C(s)|^q ds},$$

that depends, given an image I , on the oriented boundary C of the segmented region candidate. Minimizing this ratio amounts to finding a curve, neither small nor too curvy, through which the brightness flux is maximal. We prove the existence of minimizers for this criterion among continuous curves with mild regularity assumptions. We also prove that the discrete minimizers provided by our graph-based algorithm converge, as the resolution increases, to continuous minimizers.

In contrast to most existing segmentation methods with computable and meaningful, i.e. non degenerate, global optima, the proposed approach is fully unsupervised in the sense that it does not require any kind of user input such as seed nodes. Numerical experiments demonstrate that curvature regularity allows to substantially improve the quality of segmentations.

Furthermore our results allow to draw conclusions about global optima of a parameterization-independent version of the Snakes functional: the proposed algorithm allows to determine parameter values where the functional has a meaningful solution and simultaneously provides the corresponding global solution.

I. INTRODUCTION

In this paper, we introduce a novel framework for image segmentation which allows to impose curvature regularity. We start with a brief overview of existing work on curvature regularity and on image segmentation.

A. Curvature in Vision and Image Processing

Curvature regularity plays an important role in many fields of computer vision and image processing - among them image segmentation, perceptual organization and inpainting. Inspired by the psychophysical studies of Kanizsa [42] which

Thomas Schoenemann is with the Department of Mathematical Sciences, Lund University, Sweden.

Simon Masnou is with the Institut Camille Jordan, Université Claude-Bernard Lyon 1, CNRS, France.

Daniel Cremers is with the Department of Computer Science, TU München, Germany.

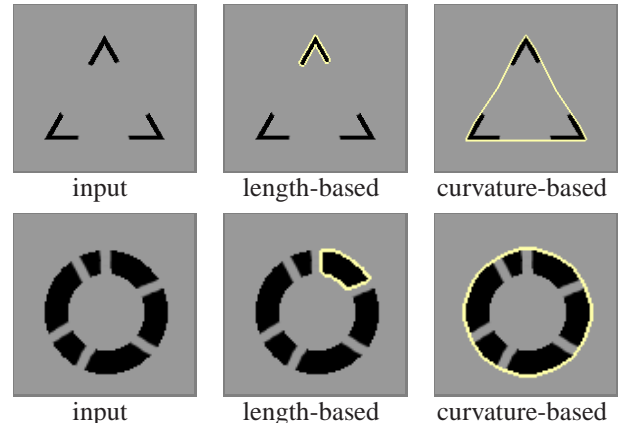


Fig. 1. Effects of length-based and curvature-based regularization in image segmentation on synthetic images. A comparison of the global minima of functionals (5) and (2) ($q = 1, \nu = 0$) shows that curvature regularity gives rise to fundamentally different segmentations. In particular it induces gap closing and contour completion.

indicated that it is a key component in human scene understanding, curvature was introduced in the context of shape completion through the celebrated Euler’s elastica energy $\int_0^{\mathcal{L}(C)} |\kappa_C(s)|^2 ds$, see [65], [38], [51] and the subsequent developments in [61], [44], [33], [19]. Other applications of curvature are image segmentation [43], [4], [3], [55], [14], inpainting [5], [17], [48], [30], [31], image smoothing and denoising (see [64] and references in [12]), image analysis [56], [27] or surface interpolation and smoothing [39], [63], [25], [7], [29].

A large body of this literature is devoted to optimizing curvature-dependent functionals in a continuous and/or discrete setting through a minimizing flow that merely yields *local* optimizers. Very few methods allow to compute *global* optimizers. An important exception is [4], [3] where line integrals depending on curvature are minimized by means of dynamic programming. Unfortunately, the run-time of this method is guaranteed to be quadratic in the number of image pixels. In practice this is too slow for reasonable image sizes. As a consequence the authors only consider curves in a band around an initial curve. In the context of image inpainting, another exception is [49], [48] where the specific functional of total absolute curvature can be optimized globally.

Further, recently combinatorial approaches to image segmentation have arising, and in the case of a 4-connectivity these can often be optimized globally [26], though without a polynomial time guarantee. The more general method [60] is practicable up to a 16-connectivity but is less likely to find global solutions.

B. Image Segmentation as an Optimization Problem

Over the past few decades numerous methods have been proposed for image segmentation. We focus on those that are based on minimizing a suitable energy functional. The corresponding functionals typically combine a data term with a regularity term (for a recent review see [21]). Both terms can incorporate either region or edge properties, yet the regularity term generally penalizes certain properties of the region boundaries.

a) Region-based Methods: In region-based image segmentation the intensity inside the region is assumed to be approximately constant [52], [18], to vary only slowly [52] or to be generated by a suitable probability model [69]. Recent methods allow to integrate flux into region-based formulations [41], [66], [45], [46], in particular for the segmentation of long and thin structures.

In the region-based framework, global optima are usually not available in polynomial time. A region-based functional with a very basic smoothness term could recently be optimized globally [24]. For length regularity and *given* region statistics, global optima can be computed [8], [16], [54]. For curvature regularity there are local evolution methods [28], [55] and, very recently, combinatorial methods [26], [60].

b) Edge-based Methods: Ever since the pioneering work of Kass, Witkin and Terzopoulos [43], edge-based methods [15], [68], [40], [67] have formed one of the major approaches to image segmentation.

Let $I : \Omega \rightarrow \mathbb{R}$ denote the image to be segmented, with $\Omega \subset \mathbb{R}^2$ the image domain. In their seminal work on the snakes model, Kass et al. [43] propose to minimize an energy computed along the region boundary described by a curve $C : [0, 1] \rightarrow \Omega$. The criterion to be minimized reads:

$$\alpha \int_0^1 |C'(t)|^2 dt + \beta \int_0^1 |C''(t)|^2 dt - \int_0^1 |\nabla I(C(t))|^2 dt$$

where α and β are positive constants. The first two terms of this functional measure the smoothness of the curve. Recall that the second derivative $C''(t)$ coincides with the local curvature for arc-length parameterized curves C . The third term aims at attracting the curve to places with large image gradients. It is a well known limitation of the model that for many choices of the weighting parameters, no meaningful global optima can be found: the only minimizers would be either points or curves with infinite length. In this paper we state a parameterization-invariant version of the functional and show that the proposed method allows to identify meaningful global optima of this functional.

Amini et al. [4] optimize a modified snakes functional using dynamic programming. To avoid meaningless solutions, they enforce the curve to be polygonal with an a priori known number of points, such that no two neighboring points are less than a given distance apart. Although their method provides polynomial time solutions, due to its quadratic run-time complexity it is only applied in a small band around an initial curve.

Caselles et al. [15] propose to minimize a line integral of a positive edge indicator function $g(|\nabla I|)$, where g decreases monotonically with increasing gradient strength. Although this

model is interesting for local optimization, it is easily seen that the global minimizers are meaningless: any degenerate curve reduced to a point in the image is a minimizer of this functional. Meaningful global optima can be found when seed nodes are given for foreground and background [9].

The normalized cuts of Shi and Malik [62] produce fair results without seed nodes. While the optimization problem is NP-hard, using relaxation techniques one can find a solution which is independent of initialization.

The ratio regions of Cox et al. [20] provide meaningful global optima together with a polynomial time complexity. However, the complete search over all starting points takes prohibitively long in practice. Moreover the method is limited to planar graphs.

Finally, Jermyn and Ishikawa propose in [41] a class of ratio functionals for image segmentation, including the ratio of flux over length, whose global minimizers can be efficiently found by iterative negative cycle detection in a suitable graph.

C. Contribution

In this paper we extend the class of globally optimizable ratio functionals considered in [41]: We show how to integrate curvature regularity of the region boundary. Specifically we propose an algorithm which computes in polynomial time closed **oriented** curves $C : [0, \mathcal{L}(C)] \rightarrow \Omega \subset \mathbb{R}^2$ of length $\mathcal{L}(C)$ which minimize (discrete versions of) functionals of the form

$$\frac{\int_0^{\mathcal{L}(C)} h(I, C(s), C'(s), \kappa_C(s)) ds}{\int_0^{\mathcal{L}(C)} g(C(s), C'(s), \kappa_C(s)) ds}, \quad (1)$$

where $h : C^2(\Omega, \mathbb{R}) \times \Omega \times \mathbb{S}^1 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is arbitrary (\mathbb{S}^1 denotes the circle and $C^2(\Omega, \mathbb{R})$ the space of twice continuously differentiable real-valued functions on Ω) and $g : \Omega \times \mathbb{S}^1 \times \mathbb{R}^2 \rightarrow \mathbb{R}^+$ such that the denominator is strictly positive for all closed curves C with strictly positive length. Since we assume arc-length parameterization, $C'(s)$ is a unit tangent vector and $\kappa_C(s) = C''(s)$ is the curvature.

A specific instance of (1) is the so-called *elastic ratio*:

$$\frac{\int_0^{\mathcal{L}(C)} \nabla I(C(s)) \cdot (C'(s))^\perp ds}{\nu \mathcal{L}(C) + \int_0^{\mathcal{L}(C)} |\kappa_C(s)|^q ds} =: \frac{Flux(C)}{\nu \mathcal{L}(C) + F_\kappa^q(C)}, \quad (2)$$

that allows to segment the input image $I : \Omega \rightarrow \mathbb{R}$ by determining a contour C which maximizes the flux of brightness along the normal $(C'(s))^\perp$ (denoting the tangent vector $C'(s)$ rotated by $+\frac{\pi}{2}$), while minimizing a weighted sum of boundary length and integrated local curvature. Here q and ν are positive real weighting factors and \cdot denotes the scalar product in \mathbb{R}^2 . Note that changing the orientation of the curve C in (2) does not modify the denominator but changes the orientation of $C'(s)$, therefore the orientation of $(C'(s))^\perp$, and, finally, the sign of the numerator. Therefore the minimum value is necessarily negative and the following equivalence holds

$$\min_C \frac{Flux(C)}{\nu \mathcal{L}(C) + F_\kappa^q(C)} \Leftrightarrow \max_C \frac{|Flux(C)|}{\nu \mathcal{L}(C) + F_\kappa^q(C)}.$$

This equivalence will be used in the proof of existence of continuous minimizers. We shall however keep using the "min" formulation because it is consistent with our graph approach based on Lawler's Minimum Ratio Cycle algorithm, and also to emphasize better the difference with the criterion of Jermyn and Ishikawa involving only length in the denominator [41], that will be discussed in section II-A.

Remark that, for constant images, all curves have null energy therefore all are solutions: there is no preferable segmentation, as it can be expected.

The key contribution of the curvature term in the denominator of (2) is to energetically favor curves along which angle variations are not too large. Figure 1 demonstrates this effect on synthetic images: In comparison to segmentations with pure length regularization, using Jermyn and Ishikawa's criterion, the proposed curvature regularity gives rise to contour completions that are reminiscent of those observed in Kanizsa's psychophysical studies [42].

Another instance of (1) that we consider in this paper is the *snakes ratio*:

$$\frac{-\int_0^{\mathcal{L}(C)} |\nabla I(C(s))|^p ds}{\nu \mathcal{L}(C) + F_\kappa^q(C)}, \quad (3)$$

with parameters $p, q, \nu \geq 0$. As we shall see later, the snakes ratio is closely related to a parameterization-invariant formulation of the famous snakes model [43]. As a consequence, our algorithm also allows to compute global minimizers of this reformulation and simultaneously provides a parameter set that gives rise to meaningful global optima.

The proposed framework exhibits the following properties, most of them being naturally inherited from the model of Jermyn and Ishikawa:

- We can find globally minimizing discrete curves for a bunch of functionals of the form (1) with a variety of data terms and regularity terms. In the numerator any functional dependence on location $C(s)$, local tangent $C'(s)$ and local curvature $\kappa_C(s)$ of the curve is allowed. In the denominator we require functions yielding positive integrals for all closed curves.
- The proposed method is fully unsupervised in the sense that it does not rely on any user input once the functional parameters and the discretization scale have been chosen. It arises from our experiments that a unique set of parameters can be used once for all. In contrast, most existing globally optimal segmentation methods either give rise to trivial minimizers (typically the empty set) or require some user seeds as additional input – see for example [8].
- Global optima of the discretized functional can be determined in polynomial time by means of ratio cycles in a graph where nodes correspond to line segments, i.e. each graph node is a pair of pixels. As a consequence of this global optimization, solutions do not depend on initialization or the choice of a particular numerical optimization scheme.

Further, we discuss the strengths and weaknesses of the proposed edge-based approach compared to recent works on

region-based combinatorial methods [60], [26] for curvature regularity which could be extended to handle ratios in a manner similar to [46], [53]. Region-based methods should be used when (a) one has to handle region-based data terms. (b) one wants to find more than one connected component or components with holes (e.g. when seed nodes are given). (c) one has run an edge-based method and got one of the rare cases where it produces self-intersections. In all other cases, the proposed method is preferable: it avoids a subdivision of pixels and hence has much less basic line segments to consider. Moreover, it does not have region variables at all. In summary, it is able to handle larger images and neighborhoods, the latter giving rise to a better approximation of continuous curvature. In particular, the neighborhood used in this work subsumes a 32-connectivity whereas [60] manages only 16 in practice and [26] only approximately handles 8 neighbors. Further, we show that it can be ported to the GPU where it is much faster. Let us finally mention that the problem of optimizing ratios defined on a graph has been recently investigated in [36], [46], [37]. The first paper adopts a relaxation approach for the NP-hard ratio optimization problem, while the two latter works transform, as we do, the ratio minimization into the minimization of an objective function, although using ratios different from ours.

A preliminary version of this work appeared in [59]. This extended version uses a different graph (each node represents a line segment instead of a pair of an image pixel and a direction), and addresses the issues of an efficient parallel implementation as well as theoretical considerations of the continuous optimization task: we address the question of existence of minimizers in the continuous domain (for the elastic ratio) and prove the convergence of discrete minimizers to continuous minimizers as the resolution increases.

The remainder of the manuscript is organized as follows. In Section II we review and discuss the ratio cycle framework of Jermyn and Ishikawa. In Section III, we generalize this framework so as to incorporate curvature regularity and we state the existence of solutions in the continuous setting and the consistency of the discretized formulation, the proofs of both claims being detailed in the appendix. In Section IV, we detail how a discretized version of the proposed functional can be represented by means of an appropriate graph structure. In Section V, we detail how optimal solutions can be computed in polynomial time by determining cycles in this graph. In Section VI, we analyze a parameterization-invariant version of the Snakes functional and show how our algorithm can be used to determine meaningful minimizers. In Section VII we provide numerical experiments which demonstrate that curvature regularity substantially improves the segmentation results. We end with a summary and a conclusion.

II. SEGMENTATION BY OPTIMIZING RATIOS

In [41] Jermyn and Ishikawa proposed to optimize a criterion that exhibits two nice properties:

- 1) Its discrete counterpart can be *globally* optimized in polynomial time. Interestingly, this property is quite unusual in the context of image segmentation where

for most fully unsupervised energy minimization approaches the global optima are either meaningless (typically the empty set) or not efficiently computable.

- 2) It does not suffer from a shrinking bias. Since the length of the curve is factored out, the formulation is essentially scale invariant: taking an image with a single binary object (or a smooth approximation of it), the energy remains constant for all rescalings of the object, see [41]. This is an important aspect in the context of image segmentation where, clearly, the same object may be observed at all sorts of scales. It is therefore preferable that the energy does not exhibit a bias towards smaller scales.

We will first state the contour-based formulation, and then turn to the region-based one.

A. Contour-based Problem Statement

In Jermyn and Ishikawa's approach, the image domain is segmented into two regions separated by an *oriented* curve C . In the most general setting, which has less invariance properties, this curve minimizes a ratio of the form

$$\frac{\int_0^{\mathcal{L}(C)} \vec{v}(C(s)) \cdot (C'(s))^\perp ds}{\int_0^{\mathcal{L}(C)} g(C(s)) ds}, \quad (4)$$

where the curve C is parameterized with arc-length, $\vec{v} : \Omega \rightarrow \mathbb{R}^2$ is a vector field, $(C'(s))^\perp$ is the unit normal to the curve that coincides with $C'(s)$ rotated by $+\frac{\pi}{2}$ and $g : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ is a positive weight function.

In particular, Jermyn and Ishikawa propose to minimize the average outward flux [66] that we call *length ratio* in this paper:

$$\frac{\int_0^{\mathcal{L}(C)} \nabla I(C(s)) \cdot (C'(s))^\perp ds}{\mathcal{L}(C)} = \frac{Flux(C)}{\mathcal{L}(C)}, \quad (5)$$

Like for the elastic ratio, due to the role of the curve orientation, minimizing (5) over oriented curves is equivalent to maximizing its absolute value.

It therefore amounts to finding curves aligned with strong image edges. The normalization by length avoids trivial optima (zero-length or infinite-length curves). Yet, in practice the found regions tend to be small and usually do not coincide with human perception. As a remedy, Jermyn and Ishikawa propose a balloon force. To understand how this works we must first look at the region-based interpretation.

B. Conversion to a Region-based Form

Jermyn and Ishikawa [41] also demonstrate that the edge criterion in (4) can be transformed into a region-based criterion by means of the Gauss-Green theorem. Denoting C a simple and smooth curve in \mathbb{R}^2 and C_{in} the connected region enclosed by C , the Gauss-Green Theorem states that for any smooth field \vec{v} on \mathbb{R}^2

$$\int_{C_{in}} \operatorname{div} \vec{v} dx = - \int_0^{\mathcal{L}(C)} \vec{v}(C(s)) \cdot \vec{n}_C(s) ds, \quad (6)$$

where $\vec{n}_C(s)$ is the *inner* unit normal to C at $C(s)$. Therefore, the minimization of (5) can equivalently be written as

$$\max_C \frac{|\int_{C_{in}} \Delta I(x) dx|}{\mathcal{L}(C)},$$

where $\Delta I = \operatorname{div}(\nabla I)$ is the Laplace operator. This expression gives some indications on the optimal regions and confirms what can be observed experimentally: the optimal regions are likely to contain high values of ΔI with constant sign, which often occurs for small regions in the vicinity of edges, i.e. in zones where ∇I changes a lot.

To include the balloon force weighted with β one only has to add $\pm\beta$ to $\Delta I(x)$ (see [41] for details). The arising problem can be written as

$$\max_C \frac{|Flux(C) \pm \int \beta dx|}{\mathcal{L}(C)}, \quad (7)$$

and will be termed *extended length ratio* in the following. It provides substantially better results in practice, although the choice of β is delicate for it influences the result significantly.

III. CURVATURE REGULARITY IN RATIO OPTIMIZATION

We propose a different remedy, instead of incorporating balloon forces as Jermyn and Ishikawa did, to the problem that minimizers of the length ratio are usually very small curves: in addition to the length of the curve, we will also penalize its curvature, and thus consider the minimization of (2).

Why should this particular energy help to avoid small curves more than the length ratio? The reason is that the curvature term discourages direction changes: the curvature along a line segment is zero. Hence the curvature term will not grow in parts where the curve goes straight, so that for long and not oscillating curves the denominator will be roughly similar to the length only, and the ratio will not differ much from the length ratio. In contrast, for small curves, the curvature term will be much larger than the length so the ratio will be much smaller, in absolute value, than the length ratio.

As the minimization of (2) is intensively studied in this paper, we now prove the existence of minimizers under some mild assumptions on the image I and on the length of curves.

A. Existence of minimizers of the elastic ratio for $q > 1$

Assuming that Ω is bounded, the existence of minimizers of the elastic ratio (2) whenever $q > 1$ is rather easy to prove among closed curves with length uniformly bounded by a constant – that can be chosen arbitrarily large – and that admits a uniform parameterization in the Sobolev space $W^{2,q}([0, 1], \bar{\Omega})$ [32]. Recall that this space contains all functions from $[0, 1]$ onto $\bar{\Omega}$ that are, together with their distributional first and second derivatives, L^q -integrable. Since $q > 1$, this space, endowed with the norm $\|C\|_{2,q} = (\|C\|^q + \|C'\|^q + \|C''\|^q)^{1/q}$, has the essential property that any bounded sequence admits a weakly converging subsequence, which is the key to find a minimizer. We shall not tackle the case $p = 1$, that corresponds to seeking parameterized curves with first derivative having

bounded variation, for it requires slightly more intricate arguments and is beyond the scope of this paper. Taking $q < 1$ yields a somewhat ill-posed problem since $\|\cdot\|_q$ does not satisfy the triangle inequality anymore, therefore is not a norm. In the sequel, we restrict to $q > 1$ for proving existence of continuous minimizers.

For simplicity, we assume that the image I is continuously differentiable on $\bar{\Omega}$. As mentioned before, due to the influence of curve orientation, minimizing (2) is equivalent to maximizing

$$\frac{|Flux(C)|}{\nu \mathcal{L}(C) + F_{\kappa}^q(C)}. \quad (8)$$

Using a uniform parameterization on $[0, 1]$ (indicated by a variable t instead of s) this ratio can be rewritten as

$$\frac{\left| \int_0^1 \nabla I(C(t)) \cdot C'(t)^\perp dt \right|}{\nu \mathcal{L}(C) + [\mathcal{L}(C)]^{1-2q} \int_0^1 |C''(t)|^q dt}, \quad (9)$$

where $C'(t)^\perp$ denotes the vector $C'(t)$ rotated by $\frac{\pi}{2}$. Let $A > 0$ and define

$$W_A^{2,q}([0, 1], \bar{\Omega}) = \{C \in W^{2,q}([0, 1], \bar{\Omega}), \mathcal{L}(C) \leq A\}.$$

We assume that I and A are such that there exists at least a simple closed curve in $W_A^{2,q}([0, 1], \bar{\Omega})$, $q > 1$, for which the numerator in (9) is non zero, otherwise the problem is trivial: for constant images, any curve is solution because all curves have null energy. The following theorem holds, whose proof is given in the appendix.

Theorem 1: Under the assumptions above, there exists a curve that minimizes the elastic ratio when $q > 1$ among all curves in $W_A^{2,q}([0, 1], \bar{\Omega})$ that are either simple or limit of simple curves. A minimizer also exists among the whole class of curves in $W_A^{2,q}([0, 1], \bar{\Omega})$.

B. The Class of Optimizable Functionals

We are not able to prove for the most general functional model (1) the existence of continuous global minimizers and their consistent discrete approximation with the proposed approach. Nevertheless, our algorithm can compute *discrete* global minimizers of *discrete* counterparts of (1), therefore global minimizers for a bunch of discrete functionals with a great variety of data and regularity terms.

In section IV, we will describe how ratios of the form (1) can be globally optimized after a suitable discretization. It should be noted that self-intersecting curves can occur. Yet, in our experiments we observed them only for very small length weights ν .

C. A Comment on Region-based Terms

Following Jermyn and Ishikawa [41], one might want to include regional terms into the algorithmic framework – such as log-likelihoods of color probabilities commonly used in region-based segmentation methods. Unfortunately, however, our algorithmic framework does not allow the introduction of arbitrary regional terms.




Fig. 2. An example of inconsistent normals for self-intersecting curves. While the ratio cycle algorithm does allow to estimate a local normal vector, there is no way of locally determining whether this normal vector corresponds to an *inner* or an *outer* normal. As a consequence, regional terms cannot be globally converted, by means of the Gauss-Green Theorem, to boundary terms written as flux in the direction of a normal defined continuously along the oriented curve (see text). Since self-intersecting curves cannot a priori be excluded, the proposed framework therefore does not allow a consistent integration of regional terms.

The reason is that, as mentioned above, we cannot exclude the emergence of self-intersecting curves – see Figure 2. Since the algorithm only deals with continuously oriented curves, in case of self-intersection the normal to the curve will actually be inner at some locations of the curve and outer at other locations. Therefore the conversion of the flux term into a regional term by means of the Gauss-Green theorem (as presented above) no longer applies. More specifically we cannot replace in equation (6) the inner normal to the domain's boundary by $(C'(s))^\perp$ for every s (or $-(C'(s))^\perp$ depending on the curve orientation). As a consequence, when including regional terms by means of the Gauss-Green theorem, for self-intersecting curves the respective cost no longer coincides with the actual regional energy.

Let us exemplify this on the specific example of Figure 2. While the regional integral of f over the area enclosed by the curve – corresponding to the left-hand-side of equation (6) – is $\int_{C_{in}} f dx = 2$ (assuming for simplicity that $|C_{in}| = 1$), the curve integral gives:

$$\int_0^{\mathcal{L}(C)} \vec{v}_f(C(s)) \cdot (C'(s))^\perp ds = 0, \quad (10)$$

with \vec{v} chosen such that $\text{div } \vec{v}_f = f$, using for instance the simple formula $\vec{v}_f(x, y) = \frac{1}{2} (\int_0^x f(t, y) dt, \int_0^y f(x, t) dt)$. To obtain (10), we have simply applied the Gauss-Green theorem separately to both loops with $(C'(s))^\perp$ representing the inner normal for one loop and the outer normal for the other. Thus for such self-intersecting curves, we cannot write

$$\int_{C_{in}} \text{div } \vec{v} dx = \pm \int_0^{\mathcal{L}(C)} \vec{v}(C(s)) \cdot (C'(s))^\perp ds.$$

A solution would be to choose *locally* either $(C'(s))^\perp$ or $-(C'(s))^\perp$ but there is no general rule to determine locally which one is the inner normal. Therefore boundary properties such as the flux in direction of $(C'(s))^\perp$ cannot be converted to respective region-based energies.

As a consequence of this limitation, we shall not consider regional terms in this paper, but rather restrict ourselves to edge-based expressions.

IV. THE OPTIMAL CURVE AS CYCLE IN A GRAPH

To globally optimize functionals of form (1) we discretize the space of all possible curves: a curve is now defined as a contiguous subset among finitely many line segments. Nevertheless, we have the continuous optimization task in mind, and the optimal discrete curve is viewed as a polygonal approximation of the optimal continuous one: its length and curvature estimates approximate the values of the continuous solution. Convergence is discussed in section IV-C.

To be able to use the path-based method described later on, we build a graph where each edge e is assigned a numerator edge weight $n(e)$ and a denominator edge weight $d(e)$. Each closed curve C in the discrete search space corresponds to some cycle Γ in the graph. In the end we will minimize the ratio problem

$$\min_{\Gamma} \frac{\sum_{e \in \Gamma} n(e)}{\sum_{e \in \Gamma} d(e)}. \quad (11)$$

The numerator and denominator sums should hence reflect the respective integrals in (1).

For the length ratio (5) Jermyn and Ishikawa [41] build a graph with one node for each image pixel and where edges represent line segments connecting pixels in an 8-neighborhood. The edge weights correspond to the respective integrals along the line segments.

For functionals of form (1) a more elaborate graph structure is needed: the optimization algorithm supports only edge weights depending on *single* edges. However, if edges directly correspond to line segments, one cannot approximate the curvature of the desired continuous curve: line segments always have zero curvature.

The graph in [59] has nodes corresponding to a pair of an image pixel and an incoming direction. In this work we take a slightly different approach, depicted in Figure 3, that allows a more consistent definition of discrete curvature: each node in the graph represents a line segment connecting two image pixels. More precisely the search space consists of all pairs of pixels that are spaced apart from one another not more than a certain distance R : if \mathcal{P} is the pixel set of the image, the node set $\mathcal{V} \subset \mathcal{P}^2$ of the graph is expressed as (with $|\cdot|$ the L_2 -norm)

$$\mathcal{V} = \{ (\vec{p}_1, \vec{p}_2) \mid 0 < |\vec{p}_1 - \vec{p}_2| \leq R \}.$$

Edges in the graph connect nodes sharing an image pixel. More precisely the edge set is

$$\mathcal{E} = \{ ((\vec{p}_1, \vec{p}_2), (\vec{p}_2, \vec{p}_3)) \in \mathcal{V}^2 \mid \vec{p}_1 \neq \vec{p}_3 \}.$$

An edge now represents an oriented polygonal curve consisting of two line segments, both of length $\leq R$. We now turn to the question of how to define suitable edge weights to approximate the continuous functional (1).

A. Estimating Curvature, Normals and Tangent Angles

To define the edge weights one first needs to calculate the tangent vector – or, equivalently, the tangent angle with respect to the x -axis – and the curvature of the corresponding part of




Fig. 3. **Top:** a grid, here representing a 4-connectivity, with a cycle (drawn in thick lines). Note that for every line segment two orientations are considered. **Bottom:** A part of the arising graph for curvature optimization. Each oriented line segment represents a node in the graph. Edges link consistently oriented line segments.

the curve. Recall that an edge represents a polygonal curve consisting of two adjacent line segments, say $\vec{p}_1\vec{p}_2$ and $\vec{p}_2\vec{p}_3$. For these two line segments the corresponding tangent angles with respect to the x -axis, $\alpha_{1,2}$ and $\alpha_{2,3}$, are computed using the function atan2 on the difference vector of the respective end points (recall that $\text{atan2}(y, x)$ is the principal value in $]-\pi, \pi]$ of the argument of $x + iy$, and that it is a routine function in C++). This difference vector also allows to calculate curve normals: to this end the vector is normalized and rotated by $+\frac{\pi}{2}$.

Estimating curvature is a more difficult issue. To allow optimal convergence properties, we follow the results of Bruckstein, Netravali, and Richardson [11]. Denoting $l_{1,2}$ and $l_{2,3}$ the lengths of the two line segments, the absolute curvature at point \vec{p}_2 is estimated as

$$|\kappa|(\vec{p}_1, \vec{p}_2, \vec{p}_3) = \frac{|\alpha_{1,2} - \alpha_{2,3}|_{\mathbb{S}^1}}{\frac{1}{2} \min(l_{1,2}, l_{2,3})}, \quad (12)$$

where the angle difference is taken on the circle \mathbb{S}^1 to correctly account for the jump over 2π . The sign of the curvature (if needed) is determined by whether the angle between the line segments exceeds 180 degrees or not. We discuss below convergence properties of this process.

B. Computing the Edge Weights

We first give edge weights for both ratios (2) and (3) considered in the experimental section. For the numerator we discretize each line segment via Bresenham's method [10]. Then, we evaluate the data term for each pixel using the above mentioned segment normal and sum the obtained values.

For the denominator, we evaluate length-based and curvature-based terms separately. The length of the curve is

Minimum Ratio Cycle Algorithm

Input: A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with two edge weights $n(e)$ and $d(e)$ for each edge.

Output: A cycle Γ minimizing the ratio $\sum_{e \in \Gamma} n(e) / \sum_{e \in \Gamma} d(e)$.

- 1) Find an upper bound λ on the optimal ratio λ_{opt}
- 2) Compute edge weights $w(e) = n(e) - \lambda d(e)$ for each edge $e \in \mathcal{E}$.
- 3) Call the Moore-Bellman-Ford algorithm (Fig. 5) for the graph \mathcal{G} and the edge weights w . If it returns a negative cycle, set λ to its ratio and go to 2). Otherwise output the last found cycle and **stop**.

Fig. 4. Ratio optimization after Lawler [47]. Shown is the linear search variant, also known as Dinkelbach's method [23].

readily calculated as the sum of all line segment lengths. For the curvature term, we evaluate the expression (12) and take the desired power of it. Finally, to get the integral of this term, the length of the segment needs to be included in the weights. Here again, we follow the results of [11], summing over all edges $(\vec{p}_1, \vec{p}_2, \vec{p}_3)$ the quantity

$$\frac{1}{2} \min(l_{1,2}, l_{2,3}) \cdot [|\kappa|(\vec{p}_1, \vec{p}_2, \vec{p}_3)]^q.$$

The general case (1) is more difficult to handle for it may contain terms like $\int_0^{\mathcal{L}(C)} I(C(s)) |\kappa_C(s)| ds$, where the integrand depends on position and curvature simultaneously. This makes it much harder to reflect the continuous functional in terms of discrete sums. Our solution is presently to calculate the Bresenham lines, then to calculate the values of $h(\cdot, \cdot, \cdot)$ and $g(\cdot, \cdot, \cdot)$ for each pixel, using the above given estimates for normals and curvature. These values are then summed to form the edge weights.

C. Consistency of the discrete elastic ratio minimization

We defer the question of how to compute the optimal polygonal curve to the next section and assume for the moment that it is solved. We should mention, however, that the optimization algorithms require the weights $n(e)$ and $d(e)$ to be subject to some regular quantization, i.e. they must be multiples of a certain $\epsilon > 0$. Termination is guaranteed for any such quantization, but the complexity depends on ϵ . Details are given in section V-D. The next theorem, whose proof is given in the appendix, states an important property of our approach when applied to the elastic ratio (2) with $q > 1$: it is consistent, i.e., continuous (global) minimizers are limits of discrete (global) minimizers.

Theorem 2: For every $q > 1$, any sequence of simple discrete minimizers of the discrete elastic ratio associated with increasing resolutions has a subsequence that converges to a continuous minimizer of the elastic ratio.

V. RATIO OPTIMIZATION OVER CYCLES IN A GRAPH

We now address the task of finding the cycle of optimal ratio in the described graph, i.e. how to solve the task

$$\min_{\Gamma} \frac{\sum_{e \in \Gamma} n(e)}{\sum_{e \in \Gamma} d(e)}$$

Moore-Bellman-Ford Algorithm

Input: A directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with (possibly negative) edge weights $w(e)$ for each edge. A root node $r \in \mathcal{V}$.

Output: A distance label $c(v)$ and a predecessor node $p(v)$ for every node $v \in \mathcal{V}$ in the graph. If the graph contains negative cycles such a cycle is returned.

- 1) Set $c(r) = 0$, $c(v) = \infty$ for $v \in \mathcal{V} \setminus \{r\}$. Mark $p(v)$ as invalid for all v .
- 2) Set `changes := false`
For all $v \in \mathcal{V}$: check all incoming edges $e = (\tilde{v}, v)$. **If** $c(\tilde{v}) + w(e) < c(v)$
 $c(v) = c(\tilde{v}) + w(e)$, $p(v) = \tilde{v}$
`changes := true`
- 3) **If** `changes = false` **stop**.
Otherwise check the predecessor entries p for cycles. If a cycle is found, return the cycle. Else go to 2).

Fig. 5. Distance calculation and negative cycle detection via the Moore-Bellman-Ford algorithm [34], [50], [6], [1].

over all cycles Γ , where $n(e)$ and $d(e)$ are subject to a regular ϵ -quantization with $d(e) \geq 0$, and such that the denominator sum is strictly positive for any cycle in the graph. To this end, we use a variant of the Minimum Ratio Cycle algorithm proposed by Lawler [47]: instead of binary search, we use linear search as proposed by Dinkelbach [23], which is much faster in practice.

The basic algorithm is shown in Figure 4. It is based on iterated negative cycle detection in a graph with single edge weights. Let λ be some ratio and define edge weights

$$w(e) = n(e) - \lambda d(e).$$

Now suppose the graph contains a negative cycle Γ with respect to the edge weights $w(e)$. By applying equivalence transformations one sees that any such cycle must be of better ratio than λ and vice versa:

$$\begin{aligned} \sum_{e \in \Gamma} w(e) < 0 &\Leftrightarrow \sum_{e \in \Gamma} [n(e) - \lambda d(e)] < 0 \\ &\Leftrightarrow \sum_{e \in \Gamma} n(e) < \lambda \sum_{e \in \Gamma} d(e) \Leftrightarrow \frac{\sum_{e \in \Gamma} n(e)}{\sum_{e \in \Gamma} d(e)} < \lambda. \end{aligned} \quad (13)$$

Notice that the third step is valid only because of the positivity of all conceivable denominator sums. This is the reason for the previously introduced restriction on the denominator.

The above equivalence transformation shows that the graph contains a negative cycle with respect to $w(e)$ if and only if the optimal ratio is lower than λ . If one is able to find negative cycles, this immediately gives rise to the algorithm in Figure 4: starting from some upper bound on the optimal ratio – 0 for all addressed problems – negative cycle detection and ratio adjustments are alternated. Every time a negative cycle is found, λ is set to its ratio. The last found cycle must be of optimal ratio.

Negative cycle detection is performed efficiently by the Moore-Bellman-Ford algorithm [34], [50], [6], [1] for distance calculations. The algorithm, depicted in Figure 5, is based on dynamic programming: starting from an initial distance labeling, the label of any node is reduced whenever the labels of its predecessors allow such an improvement. If the graph

does not have negative cycles, the algorithm terminates with the correct distance labeling. Otherwise, after a few steps the parent entries will permanently contain cycles. Regularly checking for cycles then allows to extract a negative cycle, which is necessary to update the ratio.

While the basic algorithm in Figure 4 must be carried out sequentially, the negative cycle detection in Figure 5 allows a lot of freedom for the implementation. We now discuss how to efficiently implement negative cycle detection, both in a sequential and in a parallel way. The key for efficiency lies in how to implement step 2) in Figure 5. Concerning the numerical implementation we noticed that both double precision and integer optimization lead to the global optimum. We use integer operations for both implementations.

A. Sequential Negative Cycle Detection

Efficient sequential implementations [1, page 140] make use of a queue for implementing step 2 in Figure 5. Nodes whose distance label *cannot* change in the present iteration (because none of their neighbors changed their label in the last one) will then not be visited. Every time the distance label of a node is changed, the node is added to the end of a queue. As long as there are nodes in the queue, the front one is removed and its neighbors are checked for possible distance improvements. While the worst case complexity remains the same, in practice significant speed-ups are obtained.

To optimize the run-time an explicit representation of the entire graph is suitable. However, its memory consumption is very high: only images up to size 256×256 can be processed with 2 GigaByte of memory. We therefore implemented a version where edges (and their weights) are computed on-the-fly. This solves the memory issues, but increases the run-time significantly.

B. Parallel Negative Cycle Detection

State-of-the-art graphics hardware allows highly parallel implementations of a certain class of algorithms. This class does not contain the queue-based implementation just described. However, in the form described in Figure 5, step 2) can be implemented in parallel. The method we use is to our knowledge unpublished, but can be shown to run in pseudo-polynomial time by a similar argument as [1, page 140]. It uses two buffers of distance labels, where the second is updated based on the first. Distances and parent pointers are stored in matrices, i.e. there are no node structures at all. The cycle check is done on the CPU every 25 iterations, its computational costs (including memory transfer between GPU and CPU) are negligible in practice.

C. Choosing the Root Node

For the Moore-Bellman-Ford algorithm for distance calculation (Fig. 5), a root node must be fixed. While the choice of this root node does not affect the optimality property of the ratio optimization process, it can have significant influence on the performance.

For the parallel implementation it is useful to add an extra root node and connect it to every node by an edge weighted

with 0. This amounts to initializing all distance labels with 0. After k iterations the distance label of any node contains the weight of the cheapest path of length k passing through it. While in theory one can still have $|\mathcal{V}|$ iterations until a negative cycle arises, in practice we expect a number of iterations in the order of the length of the most negative cycle in the graph.

This initialization could be used for the sequential implementation as well. However, we do not consider this sensible: first of all, the memory requirements are high since initially every node in the graph is added to the queue. Also one will have to visit every node in the graph at least once, which reduces the efficiency of the method in practice. For the first negative cycle detection, we choose a root node in the center of the image. In subsequent calls the root node is selected as one of the nodes in the last found cycle.

D. Complexity of the Method

The described graph to estimate curvature contains $\mathcal{O}(|\mathcal{P}| \cdot R^2)$ nodes. Since each node is connected with $\mathcal{O}(R^2)$ neighbors, there are $\mathcal{O}(|\mathcal{P}|R^4)$ edges. The Moore-Bellman-Ford algorithm is known to terminate in time $\mathcal{O}(nm)$ on a graph with n nodes and m edges. This gives us a worst case complexity of $\mathcal{O}(|\mathcal{P}|^2 R^6)$ for one negative cycle detection.

Finally, there is the issue of the number of distance calculations that need to be performed. Let $\epsilon > 0$ be the level of quantization, w_n be the maximum absolute numerator weight and w_d the maximal denominator weight, both before quantization. One can show [41] that the number of iterations is then $\mathcal{O}(m^3 w_d^2 w_n / \epsilon^3)$ in the worst case, with m the number of edges. In practice the number of iterations is less than 50 for $\epsilon = 10^{-3}$ and a radius $R = 3$.

In total this results in a run-time of $\mathcal{O}(|\mathcal{P}|^5 R^{18} w_d^2 w_n / \epsilon^3)$. While this seems very high, in practice we observe a linear dependence on the number of image pixels. On the GPU, even images of size 640×480 are processed in less than half an hour using a radius of 3.

VI. MINIMUM RATIO CYCLES AND SNAKES

In this section we show that the presented class of optimizable ratio functionals allows to draw conclusions about a parameterization-invariant version of the snakes model. In the original work of Kass et al. [43] the model was stated as

$$-\int_0^1 |\nabla I(C(s))|^2 dt + \alpha \int_0^1 |C'(t)|^2 dt + \beta \int_0^1 |C''(t)|^2 dt.$$

This expression is not invariant to the chosen parameterization of the curve. Kass et al. probably chose this formulation as it allows to remove numerical instabilities when dealing with explicit parameterizations of the curve. From a today's perspective one would want a parameterization-invariant formulation which might read like this:

$$-\int_0^{\mathcal{L}(C)} |\nabla I(C(s))|^2 ds + \lambda \nu \mathcal{L}(C) + \lambda F_{\kappa}^2(C). \quad (14)$$

Here we have modified the role of the weighting parameters: one is now given a relative weight $\nu > 0$ to balance the relation

of length and curvature regularity as well as a weight $\lambda > 0$ to influence both regularity terms at once.

This functional is closely related to the *snakes ratio* (3). When applying the Minimum Ratio Cycle algorithm to this ratio with $p = q = 2$ (having in mind the equivalences in (13)), one ends up computing a $\lambda_{\text{opt}} \leq 0$ and an optimal curve C_{opt} such that

$$\frac{-\int_0^{\mathcal{L}(C_{\text{opt}})} |\nabla I(C_{\text{opt}}(s))|^2 ds}{\nu \mathcal{L}(C_{\text{opt}}) + F_{\kappa}^2(C_{\text{opt}})} = \lambda_{\text{opt}},$$

thus

$$-\int_0^{\mathcal{L}(C_{\text{opt}})} |\nabla I(C_{\text{opt}}(s))|^2 ds + |\lambda_{\text{opt}}| \nu \mathcal{L}(C_{\text{opt}}) + |\lambda_{\text{opt}}| F_{\kappa}^2(C_{\text{opt}}) = 0,$$

and any other curve has larger energy with respect to the same parameter $|\lambda_{\text{opt}}|$. Hence, the snakes ratio provides valuable insights into the modified snakes model (14): given a relative weight ν between length and curvature regularity, minimizing the snakes ratio provides an absolute regularity weight $|\lambda_{\text{opt}}|$ for which the parameterization-invariant snakes model (14) has a meaningful optimum and the algorithm also provides the associated optimal curve. This means that now a model can be optimized globally for which previously only local solutions were available.

VII. EXPERIMENTS

On several images from different domains we demonstrate the performance of the proposed method. We focus on the elastic ratio with $q = 2$, but also discuss the snakes ratio (3). In particular we show that:

- The elastic ratio with squared curvature allows object segmentation for a large variety of domains. The length weight ν was adjusted experimentally on a variety of images. We found 0.15 to give reliable results that are stable for a fairly large range of values of ν around 0.15. Practically, having fixed $\nu = 0.15$ and $R = 3$, our algorithm turns to be fully unsupervised.
- Our fully unsupervised method is able to outperform region-based methods: it is less sensitive to shading effects which allows to find more precise boundaries.
- Our method is robust to noise, i.e. even for very noisy pictures it produces results comparable to those on noise-free pictures. Again, we stress that all results can be obtained with the same values of ν and R .

A. Length Ratio vs. Elastic Ratio

In Figure 6 we show a comparison of ratio functionals on images containing objects in front of cluttered background. In two cases the length ratio finds a meaningful object, i.e. when the entire object boundary has a high contrast. On these images the elastic ratio produces comparable results.

In the majority of cases, however, the length ratio tends to find small homogeneous regions. The figure shows that in many of these cases the elastic ratio is able to locate meaningful objects. This trend is confirmed by Figure 7, where we show some results on the Berkeley database.

Image		Run-time		
Name	Resolution	CPU-expl.	CPU-impl.	GPU
Seal	200 × 133	364s	812s	35s
Bunny	260 × 180	1567s	1593s	101s
Berkeley #3	321 × 481	N/A	8673s	1046s
Baseball #2	450 × 314	N/A	14810s	151s

TABLE I

Comparison of run-times for the different implementations. CPU run-time is shown with explicit storage of edges and with on-the-fly computation.

For this reason Jermyn and Ishikawa [41] proposed to integrate a suitably weighted balloon force. Figure 8¹ demonstrates that there are fairly large parameter ranges giving rise to almost identical segmentations. The functional therefore seems robust to the choice of the area weight. However, there is no parameter which works well for all the shown images.

Nonetheless the extended length ratio produces meaningful objects in several cases and we consider it somewhat complementary to the elastic ratio: each gives rise to segmentations that cannot be produced with the respective other. Whether one wants to favor objects with large area or with low curvature of the region boundary will generally depend on the application.

B. Efficiency on CPU and GPU

Due to the large search space, an efficient optimization of the elastic ratio is desirable. We implemented the algorithm both on the CPU and on the GPU as described in Section V.

The run-times for several images are given in Table I. For the smallest image the explicit graph uses roughly half the system memory. Here the explicit storage of edges is about twice as fast as the implicit one – both on the CPU. For the second image, both perform almost equally: here the entire system memory of 4 GB was needed for explicit storage.

The speed-up of the GPU version over the CPU one with implicit storage is between a factor of 8 and a factor of 100. The huge deviations are due to the different natures of the algorithms (queue-based vs. full parallel). In particular, these differences result in a different sequence of intermediate ratios.

All given run-times are quite high, lying above half a minute. However, we once again emphasize that the proposed method separates objects from the background in a *fully* unsupervised manner (i.e. does not rely on any user input, the parameters set being fixed once for all).

C. Robustness and Comparison to Region-based Approaches

For a comparison to region-based approaches we implemented variants of Mumford-Shah-like functionals [52]^{2,3}

$$E(u_1, u_2, \Omega_1) = \sum_{i=1,2} \int_{\Omega} \left(\frac{(I - u_i)^2}{\sigma_i^2} 1_{\Omega_i} + \lambda |\nabla u_i|^2 \right) dx + \nu |\partial \Omega_i|.$$

¹ We thank Greg Mori for sharing his data with us.

²Note that this functional is not identical with the original Mumford-Shah approach since the smoothness terms in the expression here are extended into the entire domain Ω .

³One may object that comparing explicit and implicit segmentation methods is unfair. We actually wanted to stress the stability issue.




Fig. 6. The elastic ratio provides more meaningful segmentations than the length ratio, in particular in the presence of partially low contrast.

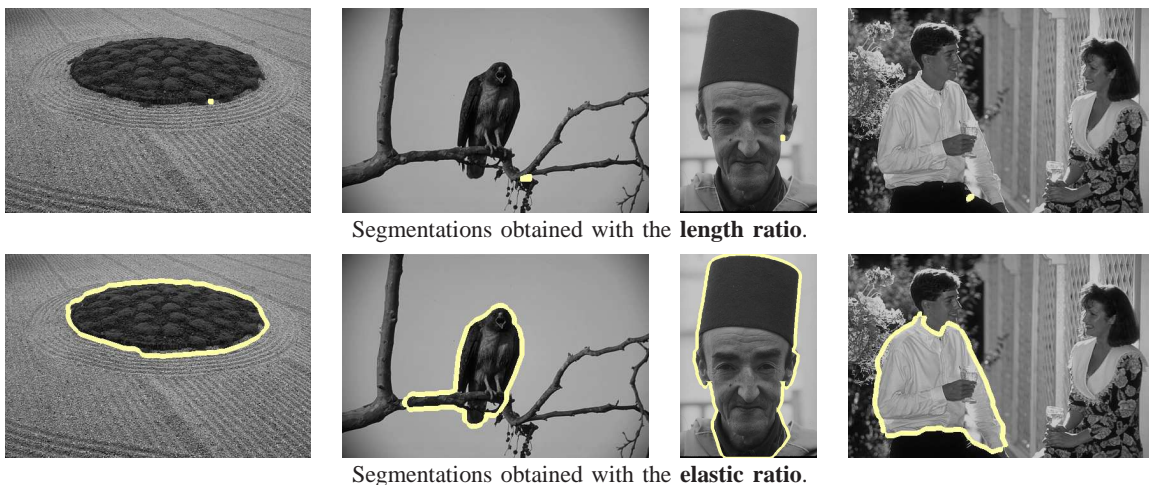


Fig. 7. The results on the Berkeley database confirm that the elastic ratio is better suited for object segmentation than the length ratio.

where a piecewise smooth approximation by two functions $u_1, u_2 : \Omega \rightarrow \mathbb{R}$ and a partition of the image plane Ω , expressed by characteristic functions 1_{Ω_i} , into two disjoint regions Ω_1 and Ω_2 is computed by alternating globally optimal updates for u_1, u_2 and Ω_1 in a manner similar to [35]. One could also use the recent method [57], [58], based on convex relaxation, and obtain an approximate solution.

In the comparative results of Figure 9, already for the noise-free image we could not find any length parameter where the above Mumford-Shah variants separate the object from the background. For the highly noisy images, despite the adaptive smoothness terms, numerous small regions arise. In contrast, the elastic ratio identifies the object almost perfectly *without* needing to adjust any parameter with respect to noise level.

The results presented in Figure 10 illustrate the robustness of the elastic ratio with respect to the length weight ν : the object is found for a fairly large range of ν . Up to a certain point, the contour becomes more complex with increasing ν . From this point on, the length term becomes dominant and the functional approaches the length ratio.

D. Results for the Snakes Ratio

Figure 11 presents results for a slightly modified snakes ratio: for robustness we use the gradient absolute instead

of its squared absolute. We recall from Section VI that all these results are global solutions of a suitably weighted parameterization-invariant reformulation of the snakes functional. More precisely, our algorithm always finds parameter sets for which the global minimizer has energy 0, therefore we cannot draw conclusions about the global minimizers of the parameterization-invariant snakes functional for other parameter sets. Still, we believe that other meaningful parameter sets do not lead to significantly better results.

When using the balancing weight $\nu = 0.15$ for length against curvature – which works well for the elastic ratio – the results are discouraging: in most cases the curve goes one way, turns around and goes almost exactly the same way back. We consider these solutions as valid in the original sense [43] since they do not self-intersect. While some line segments occur repeatedly with opposing directions of traversal, in the continuous solution space there will be a curve without repetitions and almost the same costs – at least if the image gradient is continuous.

When reducing the length weight ($\nu \rightarrow 0$) larger regions are found. These are often very close to convex and usually do not correspond to meaningful objects.




Fig. 8. Where the elastic ratio identifies body parts, the length ratio finds only small homogeneous regions. With a suitable area weight the extended length ratio can find meaningful regions.

VIII. CONCLUSION

We proposed an algorithmic framework which allows to impose curvature regularity in ratio optimization. The segmentation problem is cast as one of minimizing ratio functionals over cyclic paths in a graph where each node represents a line segment. Optimal cycles are determined in polynomial time using Lawler’s algorithm. While we cannot a priori exclude self-intersecting curves, we did not observe self intersections in the reported experiments except for very small length weights ν that are not meaningful. We proved the existence of minimizers of the elastic ratio in the continuous setting. Moreover, we proved that upon refinement of the discretization the solutions of the discretized energy converge – possibly taking a subsequence – to a minimizer of the continuous energy.

Numerous experimental comparisons demonstrate that curvature regularity allows to substantially improve the traditional

length-based regularity, it does not impose a shrinkage bias and mimics the notion of contour completion familiar from Kanizsa’s psychophysical experiments.

Lastly, we gave a parameterization-invariant reformulation of the snakes model and showed that a specific instance of our model allows to identify meaningful global solutions of this reformulation.

ACKNOWLEDGMENTS

This work was supported by the German Research Foundation, grant #CR-250/1-2. We thank Mathias Hauptmann and Frank R. Schmidt for helpful discussions. We also thank Greg Mori for sharing data with us.

APPENDIX

A. Existence of globally minimizing curves for the elastic ratio

This paragraph is devoted to proving Theorem 1. We will prove that there exists a curve that minimizes the elastic




Fig. 9. Noise Robustness and Comparison to Region-based Approaches. The elastic ratio extracts the object almost perfectly and is robust to noise, with no need to tune any parameter depending on the level of noise. In contrast, both the piecewise constant and the piecewise smooth Mumford-Shah functionals [52] fail to differentiate the object from the background (only local minimizers are shown, see text for details). For the piecewise smooth version, the smooth approximations are shown.




Fig. 10. Robustness to variations in the parameter settings of elastic ratio: For a fairly large parameter range a meaningful part of the image is found.

ratio, or equivalently that maximizes (9), among curves that are either simple or continuous limits of simple curves. The general result of existence of a minimizer among the whole class of curves in $W_A^{2,q}([0, 1], \bar{\Omega})$ can be proven analogously.

Let us turn to the maximization of (9). Under the assumptions of the theorem, the set of reachable energies is bounded thus admits a supremum. Take a maximizing sequence $(C_n)_{n \in \mathbb{N}}$, i.e. the energy of C_n tends to the supremum as n tends to infinity, of simple closed curves in $W_A^{2,q}([0, 1], \bar{\Omega})$ with uniform parameterization on $[0, 1]$. We will first prove that the sequence (C_n) is uniformly bounded in $W^{2,q}([0, 1], \bar{\Omega})$, which guarantees that there is a weakly converging subsequence. To show that a uniform bound exists, we prove that each term in the $W^{2,q}$ norm can be uniformly

controlled. With no loss of generality, we can assume that there exists $a_1 > 0$ such that, for every $n \in \mathbb{N}$

$$\begin{aligned} & \left| \int_0^1 \nabla I(C_n(t)) \cdot C_n'(t)^\perp dt \right| \\ & \geq a_1 (\nu \mathcal{L}(C_n) + [\mathcal{L}(C_n)]^{1-2q} \int_0^1 |C_n''(t)|^q dt). \end{aligned}$$

Due to the regularity of the image I , there exists also a_2 such that $\left| \int_0^1 \nabla I(C_n(t)) \cdot C_n'(t)^\perp dt \right| \leq a_2 \mathcal{L}(C_n) \leq a_2 A$, thus

$$\nu \mathcal{L}(C_n) + [\mathcal{L}(C_n)]^{1-2q} \int_0^1 |C_n''(t)|^q dt \leq \frac{a_2 A}{a_1}, \quad (15)$$

and therefore there exists a constant a_3 such that for every




Fig. 11. Experimental results using the snakes ratio demonstrate that the gradient norm is less powerful than the flux, since it does not account for edge orientation.

$n \in \mathbb{N}$

$$\int_0^1 |C_n''(t)|^q dt \leq a_3.$$

Observing that Ω is bounded and $|C_n'(t)| = \mathcal{L}(C_n) \leq A$ for every $t \in [0, 1]$ and every $n \in \mathbb{N}$, due to the assumption of uniform parameterization, we conclude that the sequence $(C_n)_{n \in \mathbb{N}}$ is uniformly bounded in $W^{2,q}([0, 1], \bar{\Omega})$, $q > 1$. Therefore (see for instance [32][Thm 1, p.144]), there exists a subsequence, still denoted as $(C_n)_{n \in \mathbb{N}}$, that converges weakly in $W^{2,q}([0, 1], \bar{\Omega})$ and strongly in $C^1([0, 1], \bar{\Omega})$ to a limit curve C . The strong convergence in $C^1([0, 1], \bar{\Omega})$ ensures that C is either simple or limit of simple curves with tangential self-contacts but no crossing. In addition, $\mathcal{L}(C_n) \rightarrow \mathcal{L}(C)$ (therefore $C \in W_A^{2,q}([0, 1], \bar{\Omega})$) and

$$\int_0^1 |C''(t)|^q dt \leq \liminf_{n \rightarrow \infty} \int_0^1 |C_n''(t)|^q dt. \quad (16)$$

Let us now check that the limit curve C has strictly positive length. From (15), we deduce that

$$\int_0^{\mathcal{L}(C_n)} |\kappa_{C_n}(s)|^q ds = [\mathcal{L}(C_n)]^{1-2q} \int_0^1 |C_n''(t)|^q dt \leq \frac{a_2 A}{a_1}.$$

Extending Fenchel's Theorem [13][Theorem 5.7.3] to $W^{2,q}$ curves by approximation, we know that for every $n \in \mathbb{N}$, $\int_0^{\mathcal{L}(C_n)} |\kappa_{C_n}(s)| ds \geq 2\pi$. By the Hölder inequality, it follows that

$$\mathcal{L}(C_n)^{q-1} \int_0^{\mathcal{L}(C_n)} |\kappa_{C_n}(s)|^q ds \geq (2\pi)^q,$$

thus $\mathcal{L}(C_n)^{q-1} \geq \frac{a_1(2\pi)^q}{a_2 A}$. Passing to the limit, we conclude that $\mathcal{L}(C) > 0$. Therefore, we can deduce from (16) that

$$\nu \mathcal{L}(C) + \frac{\int_0^1 |C''(t)|^q dt}{[\mathcal{L}(C)]^{2q-1}} \leq \liminf_{n \rightarrow \infty} \left(\nu \mathcal{L}(C_n) + \frac{\int_0^1 |C_n''(t)|^q dt}{[\mathcal{L}(C_n)]^{2q-1}} \right).$$

Besides, the continuity of ∇I and the pointwise convergence of $C_n(t)$ to $C(t)$, and $C_n'(t)$ to $C'(t)$ for every $t \in [0, 1]$ imply

that

$$\int_0^1 \nabla I(C_n(t)) \cdot C_n'(t)^\perp dt \rightarrow \int_0^1 \nabla I(C(t)) \cdot C'(t)^\perp dt,$$

and we finally get that

$$\begin{aligned} & \frac{\left| \int_0^1 \nabla I(C(t)) \cdot C'(t)^\perp dt \right|}{\nu \mathcal{L}(C) + [\mathcal{L}(C)]^{1-2q} \int_0^1 |C''(t)|^q dt} \\ & \geq \limsup_{n \in \mathbb{N}} \frac{\left| \int_0^1 \nabla I(C_n(t)) \cdot C_n'(t)^\perp dt \right|}{\nu \mathcal{L}(C_n) + [\mathcal{L}(C_n)]^{1-2q} \int_0^1 |C_n''(t)|^q dt}. \end{aligned}$$

The sequence (C_n) being maximizing, we conclude that C is a curve, limit of simple curves, that maximizes (8) in $W_A^{2,q}([0, 1], \bar{\Omega})$. The same proof can be used to establish the existence of maximizers among all $W_A^{2,q}$ curves, either simple or not, considering that Fenchel's Theorem also applies for nonsimple curves, see Remark 5 page 402 in [13].

B. Consistency of the discrete elastic ratio minimization

We prove now Theorem 2, i.e. that the limit, as the resolution increases, of a converging sequence of discrete simple minimizers of the discrete elastic ratio (i.e. $\sum n(e)/\sum d(e)$ with suitable weights computed as in section IV-B) is a minimizer of (2) in the continuous domain. Let us first recall that the usual way to study relations between discrete and continuous minimizers involves a particular notion of convergence for functionals, the Γ -convergence [22]. It has a particularly useful property: if a sequence of energy functionals F_n Γ -converges to a functional F and a sequence (x_n) of minimizers of F_n converges to x then x is a minimizer of F . In this framework the results of Bruckstein et al. in [11] are directly related to our problem. Bruckstein et al. consider the space of curves with finite length and finite total absolute curvature

endowed with the metric d defined by

$$d(C_1, C_2) = \inf_{\Psi: [0,1] \rightarrow [0,1]} \sup_{t \in [0,1]} |C_1(t) - C_2(\Psi(t))|,$$

with C_1, C_2 parameterized on $[0, 1]$ and Ψ in the class of all homeomorphisms from $[0, 1]$ to $[0, 1]$. Then they prove, using the discrete definition of curvature (12) and using d as convergence metric for sequence of curves, that the discrete counterpart of $\int_0^{\mathcal{L}(C)} |\kappa_C(s)|^q ds$ computed on polygons with n edges Γ -converges to $\int_0^{\mathcal{L}(C)} |\kappa_C(s)|^q ds$ as n tends to ∞ and the maximal length of polygon edges tends to zero. Now remark that the existence (in the continuous domain) of a curve maximizing (8) is equivalent – if (8) is not trivially zero – to the existence of a curve minimizing

$$\frac{\nu \mathcal{L}(C) + F_\kappa^q(C)}{|Flux(C)|} \quad (17)$$

in the class of $W^{2,q}$ curves with length uniformly greater than a suitable constant. If $\frac{1}{n}$ denotes the pixel size, let us define F_n as the functional that associates any polygon P_n defined on the grid with

$$F_n(P_n) = \sum_{e \in P_n} d(e),$$

where $d(e)$ is computed as in section IV-B and P_n is assumed to have a maximal edge length smaller than $\frac{\delta}{n}$ with δ a constant independent of P_n and n . According to the result by Bruckstein et al., F_n Γ -converges, as $q > 1$, to the functional

$$F(C) = \nu \mathcal{L}(C) + F_\kappa^q(C).$$

Besides, remark that the smoothness of I implies that its discrete gradient computed with finite differences uniformly converges to the continuous gradient ∇I . Take any sequence of simple polygons (P_n) with uniformly bounded length that converges for the metric d to a limit curve C . Let $\text{int}(P_n)$ and $\text{int}(C)$ denote the sets enclosed by P_n and C , respectively, and $\mathbb{1}_{\text{int}(P_n)}, \mathbb{1}_{\text{int}(C)}$ the associated characteristic functions. By the theory of functions of bounded variation [2], the derivatives $D\mathbb{1}_{\text{int}(P_n)}$ weakly- \star converge to $D\mathbb{1}_{\text{int}(C)}$ as $n \rightarrow \infty$. It follows from the Gauss-Green Theorem for BV functions [2] that

$$\sum_{e \in P_n} n(e) \rightarrow Flux(C),$$

and we deduce that the ratio $\frac{\sum_{e \in P_n} d(e)}{\sum_{e \in P_n} n(e)}$ Γ -converges to (17) as n tends to ∞ . Therefore, taking a sequence of simple discrete minimizers of this ratio, there exists a subsequence that converges to a minimizer of (17) in the continuous domain. Such minimizer being non degenerate according to our assumption that the length is uniformly bounded from below, we conclude that for any sequence of simple discrete minimizers of $\frac{\sum_{e \in P_n} n(e)}{\sum_{e \in P_n} d(e)}$, there exists a subsequence that converges to a continuous minimizer of (2) as $n \rightarrow \infty$. This achieves the proof of convergence.


REFERENCES

- [1] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford University Press, 2000.
- [3] A.A. Amini. *Using dynamic programming for solving variational problems in vision: Applications Involving Deformable Models for Contours and Surfaces*. PhD thesis, The University of Michigan, Ann Arbor, 1990.
- [4] A.A. Amini, T.E. Weymouth, and R.C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 12(9):855 – 867, September 1990.
- [5] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Trans. on Image Proc.*, 10(8):1200–1211, 2001.
- [6] R.E. Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87–90, 1958.
- [7] A.I. Bobenko and P. Schröder. Discrete Willmore Flow. In *Eurographics Symposium on Geometry Processing*, 2005.
- [8] Y. Boykov and M.-P. Jolly. Interactive organ segmentation using graph cuts. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pages 276–286, 2000.
- [9] 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.
- [10] J.E. Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems Journal*, 4(1):25–30, 1965.
- [11] A.M. Bruckstein, A.N. Netravali, and T.J. Richardson. Epi-convergence of discrete elastica. In *Applicable Analysis, Bob Carroll Special Issue*, volume 79, pages 137–171, 2001.
- [12] F. Cao. *Geometric Curve Evolution and Image Processing*, volume 1805 of *Lecture Notes in Mathematics*. Springer Verlag, 2003.
- [13] M.P. Do Carmo. *Differential Geometry of curves and surfaces*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1976.
- [14] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int. Jour. of Comp. Vision*, 22(1):61–79, 1997.
- [15] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Minimal surfaces based object segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 19(4):394–398, 1997.
- [16] A. Chambolle. Total variation minimization and a class of binary MRF models. In *Int. Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, pages 136–152, St. Augustine, Florida, 2005.
- [17] T.F. Chan, S.H. Kang, and J. Shen. Euler’s elastica and curvature based inpaintings. *SIAM Journal of Applied Mathematics*, 2:564–592, 2002.
- [18] T.F. Chan and L.A. Vese. Active contours without edges. *IEEE Trans. on Image Proc.*, 10(2):266–277, 2001.
- [19] G. Citti and A. Sarti. A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision*, pages 307–326, 2006.
- [20] 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.
- [21] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *Int. Jour. of Comp. Vision*, 72(2):195–215, April 2007.
- [22] G. Dal Maso. *An Introduction to Γ -Convergence*. Birkhäuser, Boston, 1993.
- [23] W. Dinkelbach. On nonlinear fractional programming. *Management Science*, 13:492–498, 1967.
- [24] X. Dou, W. Xiaodong, A. Wahle, and M. Sonka. Globally optimal surface segmentation using regional properties of segmented objects. In *IEEE Int. Conf. on Comp. Vision and Patt. Recog.*, Anchorage, Alaska, June 2008.
- [25] M. Droske and M. Rumpf. A level set formulation for Willmore flow. *Interfaces and Free Boundaries*, 6(3), 2004.
- [26] N.Y. El-Zehiry and L. Grady. Fast global optimization of curvature. In *IEEE Int. Conf. on Comp. Vision and Patt. Recog.*, San Francisco, California, June 2010.
- [27] J. Elder and S.W. Zucker. Computing contour closure. In *Europ. Conf. on Comp. Vision*, volume 1 of *LNCS*, pages 399–412, Cambridge, U.K., April 1996. Springer Verlag.
- [28] S. Esedoglu and R. March. Segmentation with depth but without detecting junctions. *Jour. Math. Imaging and Vision*, 18:7–15, 2003.

- [29] S. Esedoglu, S. Ruuth, and R. Tsai. Threshold dynamics for high order geometric motions. *Interfaces and Free Boundaries*, 10(3):263–282, 2008.
- [30] S. Esedoglu, S. Ruuth, and R.Y. Tsai. Threshold dynamics for shape reconstruction and disocclusion. In *Int. Conf. on Image Processing*, volume II, pages 502–505, Genova, Italy, 2005.
- [31] S. Esedoglu and J. Shen. Digital image inpainting by the Mumford-Shah-Euler image model. *European Journal of Applied Mathematics*, 13:353–370, 2002.
- [32] L.C. Evans and R.F. Gariepy. *Measure Theory and Fine Properties of Functions*. Studies in Advanced Math. C.R.C. Press, 1992.
- [33] C. Fantoni and W. Gerbino. Contour interpolation by vector field combination. *Journal of Vision*, 3:281–303, 2003.
- [34] L.R. Ford. Network flow theory. Paper P-923, The Rand Corporation, Santa Monica, 1956.
- [35] L. Grady and C. Alvino. The piecewise smooth Mumford-Shah functional on an arbitrary graph. *IEEE Trans. on Image Proc.*, 18(11):2547–2561, 2009.
- [36] L. Grady and E.L. Schwartz. Isoperimetric graph partitioning for image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 28:469–475, 2006.
- [37] D.S. Hochbaum. Polynomial time algorithms for ratio regions and a variant of normalized cut. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 32(5):889–898, 2010.
- [38] B.K.P. Horn. The curve of least energy. *ACM Transactions on Mathematical Software*, 9:441–460, 1983.
- [39] L. Hsu, R. Kusner, and J. Sullivan. Minimizing the squared mean curvature integral for surfaces in space forms. *Experimental Mathematics*, 1(3):191–207, 1992.
- [40] A. Jalba, M. Wilkinson, and J. Roerdink. CPM: A deformable model for shape recovery and segmentation based on charged particles. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 26(10):1320–1335, 2004.
- [41] 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.
- [42] G. Kanizsa. Contours without gradients or cognitive contours. *Italian Jour. Psych.*, 1:93–112, 1971.
- [43] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. Jour. of Comp. Vision*, 1(4):321–331, 1988.
- [44] B.B. Kimia, I. Frankel, and A.-M. Popescu. Euler spiral for shape completion. *Int. Jour. of Comp. Vision*, 54:159–182, 2003.
- [45] V. Kolmogorov and Y. Boykov. What metrics can be approximated by geo-cuts, or global optimization of length/area and flux. In *IEEE Int. Conf. on Comp. Vision*, Beijing, China, October 2005.
- [46] V. Kolmogorov, Y. Boykov, and C. Rother. Applications of parametric maxflow in computer vision. In *IEEE Int. Conf. on Comp. Vision*, Rio de Janeiro, Brazil, October 2007.
- [47] 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 Breach.
- [48] S. Masnou. Disocclusion: A variational approach using level lines. *IEEE Trans. on Image Proc.*, 11:68–76, 2002.
- [49] S. Masnou and J.M. Morel. Level-lines based disocclusion. In *Int. Conf. on Image Processing*, volume 3, pages 259–263, Chicago, USA, 1998.
- [50] E.F. Moore. The shortest path through a maze. In *International Symposium on the Theory of Switching*, pages 285–292. Harvard University Press, 1959.
- [51] D. Mumford. Elastica and computer vision. *Algebraic Geometry and its Applications*, pages 491–506, 1994.
- [52] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989.
- [53] F. Nicolls and P.H.S. Torr. Discrete minimum ratio curves and surfaces. In *IEEE Int. Conf. on Comp. Vision and Patt. Recog.*, San Francisco, California, June 2010.
- [54] M. Nikolova, S. Esedoglu, and T.F. Chan. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal of Applied Mathematics*, 66(5):1632–1648, 2006.
- [55] M. Nitzberg, D. Mumford, and T. Shiota. *Filtering, Segmentation and Depth*, volume 662 of *LNCS*. Springer Verlag, Berlin, 1993.
- [56] 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.
- [57] T. Pock, A. Chambolle, H. Bischof, and D. Cremers. A convex relaxation approach for computing minimal partitions. In *IEEE Int. Conf. on Comp. Vision and Patt. Recog.*, Miami, Florida, 2009.
- [58] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. An algorithm for minimizing the piecewise smooth Mumford-shah functional. In *IEEE Int. Conf. on Comp. Vision*, Kyoto, Japan, 2009.
- [59] T. Schoenemann and D. Cremers. Introducing curvature into globally optimal image segmentation: Minimum ratio cycles on product graphs. In *IEEE Int. Conf. on Comp. Vision*, Rio de Janeiro, Brazil, October 2007.
- [60] T. Schoenemann, F. Kahl, and D. Cremers. Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In *IEEE Int. Conf. on Comp. Vision*, Japan, 2009.
- [61] E. Sharon, A. Brandt, and R. Basri. Completion energies and scale. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 22(10):1117–1131, 2000.
- [62] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 22(8):888–905, 2000.
- [63] T. Tasdizen, R. Whitaker, P. Burchard, and S.J. Osher. Geometric Surface Processing via Normal Maps. In *ACM Transactions on Graphics*, 2003.
- [64] D. Tschumperlé and R. Deriche. Vector-valued image regularization with PDE's: A common framework for different applications. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 27(4), 2005.
- [65] S. Ullman. Filling-in the gaps: the shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25:1–6, 1976.
- [66] A. Vasilevskiy and K. Siddiqi. Flux-maximizing geometric flows. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 24(12):1565–1578, 2002.
- [67] X. Xie and M. Mirmehdi. MAC: Magnetostatic active contour model. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 30(4):632 – 646, 2008.
- [68] C. Xu and J. Prince. Generalized gradient vector flow external forces for active contours. *Signal Processing*, 71(2):131–139, 1998.
- [69] S.C. Zhu and A.L. Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multiband image segmentation. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 18(9):884–900, 1996.



Thomas Schoenemann received a BS (Vordiplom, 2002) and the MS (Diplom, 2005) in Computer Science from the Technical University RWTH Aachen, Germany. In 2009, he obtained his PhD in Computer Science from the University of Bonn, Germany. He is currently a postdoctoral researcher at the Centre for Mathematical Sciences, Lund Institute of Technology, Sweden. His research is focused on combinatorial optimization and applications to Computer Vision.



Simon Masnou obtained an Engineer's degree from Telecom Paris, France in 1992. He received the MS in Applied Mathematics from University Paris 9, France in 1993. From 1993 to 1995 he held a teaching position in Gabon. In 1998, he obtained a PhD degree in Mathematics from University Paris 9, France. In 1999, he was a Post-Doctoral Fellow at the Scuola Normale Superiore di Pisa, Italy. From 1999 to 2009 he was an Assistant Professor at University Paris 6, France. Since September 2009, he is a Professor at Institut Camille Jordan, University



Lyon 1, France. His research interests include applications of calculus of variations and geometric measure theory to image denoising and restoration.

Daniel Cremers received the BS in Mathematics (1994) and Physics (1994), and an MS (Diplom) in Theoretical Physics (1997) from the University of Heidelberg. In 2002 he obtained a PhD in Computer Science from the University of Mannheim, Germany. Subsequently he spent two years as a postdoctoral researcher at the University of California at Los Angeles and one year as a permanent researcher at Siemens Corporate Research in Princeton, NJ. From 2005 to 2009 he headed the Research Group for Computer Vision, Image Processing and Pattern Recognition at the University of Bonn, Germany. Since 2009, he heads the Computer Vision Group at the Department of Computer Science, Technical University of Munich, Germany. His research is focused on statistical and variational methods in computer vision. He received several awards, in particular the *Best Paper of the Year 2003* by the Pattern Recognition Society, the *Olympus Award 2004*, and the *UCLA Chancellor's Award for Postdoctoral Research 2005*.