

Curvature Regularity for Region-based Image Segmentation and Inpainting: A Linear Programming Relaxation

Thomas Schoenemann
Computer Science Department
Bonn University, Germany

Fredrik Kahl
Centre for Math. Sciences
Lund University, Sweden

Daniel Cremers
Computer Science Department
Bonn University, Germany

Abstract

We consider a class of region-based energies for image segmentation and inpainting which combine region integrals with curvature regularity of the region boundary. To minimize such energies, we formulate an integer linear program which jointly estimates regions and their boundaries. Curvature regularity is imposed by respective costs on pairs of adjacent boundary segments.

By solving the associated linear programming relaxation and thresholding the solution one obtains an approximate solution to the original integer problem. To our knowledge this is the first approach to impose curvature regularity in region-based formulations in a manner that is independent of initialization and allows to compute a bound on the optimal energy.

In a variety of experiments on segmentation and inpainting, we demonstrate the advantages of higher-order regularity. Moreover, we demonstrate that for most experiments the optimality gap is smaller than 2% of the global optimum. For many instances we are even able to compute the global optimum.

1. Introduction

Regularization is of central importance to image segmentation and inpainting [14, 2, 22, 4]. The introduction of higher-order regularizers in respective energy minimization approaches is known to give rise to substantial computational challenges. Some of the most powerful approaches to image segmentation are based on region integrals with regularity terms defined on the region boundaries [3, 15, 17, 5, 8, 9, 13]. While many such methods make use of length as a regularity term, only few use curvature regularity. This is in contrast to psychophysical experiments on contour completion [12] where curvature was identified as a vital part of human perception.

Length regularization has become an established paradigm because there exist many powerful algorithms for




Figure 1. In contrast to traditional length regularity, curvature regularity in image segmentation allows to preserve semantically relevant thin and elongated structures.



Figure 2. Curvature regularity strongly improves inpainting.

computing optimal solutions for length-regularized energies, either using discrete graph-theoretic approaches based on the min-cut/max-flow duality [11, 5] or using continuous PDE-based approaches using convex relaxation and thresholding theorems [16]. To date, region-based problems for segmentation using curvature regularity have been optimized using local optimization methods only (cf. [17, 9]). As a consequence, experimental results highly depend on the choice of initialization.

In this paper, we propose a relaxed version of region-based segmentation which can be solved optimally. The key idea is to cast the problem of region-segmentation with curvature regularity as an integer linear program (ILP). By solving its LP-relaxation and thresholding the solution we obtain a solution to the original integer problem and are able to evaluate a bound on its quality. In addition, we show that the method readily extends to the problem of inpainting.

Figure 1 demonstrates that the proposed method allows to segment objects in a way which preserves perceptually important thin and elongated parts. In this case the global

optimum was found since the relaxation was tight. Figure 2 demonstrates the superior performance of curvature regularity over length regularity in a corresponding inpainting experiment.

Existing Work on Curvature Regularity. For contour- or edge-based segmentation methods researchers have successfully developed algorithms to optimally impose curvature regularity using shortest path approaches [1] or ratio cycle formulations [19] on a graph representing the product space of image pixels and tangent angles [18]. In the region-based settings considered, curvature is usually handled by local evolution methods [7, 9, 17, 22]. The only exception we are aware of is the inpainting approach of Masnou and Morel [14] who can optimize the L_1 -norm of the curvature in the absence of data terms using dynamic programming.

In this paper we propose an LP-relaxation approach to minimize curvature in region-based settings. In contrast to [14] it allows to impose arbitrary functions of curvature and arbitrary data terms. The algorithmic formulation is based on the concepts of *cell complexes* and *surface continuation constraints* which have been pioneered by Sullivan [21] and Grady [10] in the context of 3D-surface completion.

2. An Alternative Approach to Length-Based Segmentation

In this section, we begin by casting the problem of length-based image segmentation as an integer linear program. In practice more efficient algorithms exist for this problem [11, 16]. However, the presented method is readily extended to include curvature regularity, see Section 3. In Section 4 it is shown how to apply the same methodology to inpainting.

Given an image $I : \Omega \rightarrow \mathbb{R}$, the task is to divide its domain $\Omega \subset \mathbb{R}^2$ into a foreground region R and a background region $\Omega \setminus R$ by minimizing the functional

$$E(R) = \int_R g_F(\mathbf{x}) d\mathbf{x} + \int_{\Omega \setminus R} g_B(\mathbf{x}) d\mathbf{x} + \nu |C| \quad (1)$$

with $g_F : \Omega \rightarrow \mathbb{R}$ and $g_B : \Omega \rightarrow \mathbb{R}$ arbitrary functions (depending on I), $\nu > 0$ a length weight and where $C = \partial R$ denotes the boundary of R with length $|C|$. The foreground region R can consist of several connected regions, each of which can have holes. By defining $g(\mathbf{x}) = g_F(\mathbf{x}) - g_B(\mathbf{x})$ one can consider the equivalent task of minimizing

$$E(R) = \int_R g(\mathbf{x}) d\mathbf{x} + \nu |C|. \quad (2)$$

In the proposed method the domain Ω is sub-divided into a set of N non-overlapping basic regions (a so-called *cell*




Figure 3. A foreground region and its boundary line. Note that the two have opposing senses of orientations.

complex). The optimal foreground region R is then composed out of a subset of these regions. This results in an optimization problem over all subsets of the basic regions.

In addition one needs access to the region *boundary* to evaluate the length-based regularity term. In the given setting this boundary is composed of the borders of the basic regions, called *edges*. We denote by M the total number of edges. The foreground region and the boundary are represented separately but estimated simultaneously. To make sure that both fit together, the following constraints are imposed:

Surface continuation constraints. If a basic region is part of the foreground region, along each of its edges the foreground either continues with another basic region or a boundary line.

That is, the foreground region may not terminate abruptly. Rather on its border a boundary line is fitted. We will now formalize this constraint in the form of an integer linear program. It involves an indicator variable $y_i \in \{0, 1\}$ for each basic region i . Here a value of 1 means that the region is part of the foreground, 0 that it belongs to the background.

For reasons that will become clear below, we consider *oriented* boundary lines as depicted in Figure 3. These are composed of basic line segments, where for each edge there are two basic line segments, one for each orientation. We obtain $2M$ line segments. Whether or not the line segment j is part of the solution is expressed by the variable $y_{N+j} \in \{0, 1\}$. We say that a region or a line segment is *active* if its indicator variable is 1. All variables are combined into a vector \mathbf{y} .

The energy function (2) can now be written as a scalar product $\mathbf{w}^\top \mathbf{y}$ of a cost vector $\mathbf{w} \in \mathbb{R}^{N+2M}$ and the variable vector \mathbf{y} . Here w_i reflects the integral of $g(\cdot)$ over the respective basic region and w_{N+j} the length (multiplied with ν) of the respective boundary segment.

It remains to encode the surface continuation constraints. A closer look reveals that there is one constraint for each edge. To formalize it we need the notion of *incidence* to an edge. This notion is visualized in Figure 4. To define




Figure 4. Line segments (red), regions and edges. Regions and edges are assigned an orientation. This forms the basis for defining incidence. For details see text.

it, each edge is arbitrarily assigned an orientation. Of the two line segments corresponding to the edge, the one with agreeing orientation is defined to be positive incident, the other one negative incident. All other line segments are not incident to the edge. The incidence of the line segment j to the edge k is reflected in the coefficient

$$a_{k,N+j} = \begin{cases} 1 & \text{if positive incident} \\ -1 & \text{if negative incident} \\ 0 & \text{otherwise.} \end{cases}$$

In addition there is also the incidence of basic regions to edges. To define it each basic region is assigned a sense of orientation (corresponding to the order in which the corner points are traversed, compare Figure 3). Now, the basic region i is positive incident to edge k if the edge is one of its borders and the two agree in orientation. If the edge is one of its borders but in the opposite orientation, we have negative incidence. Otherwise the two are not incident. This results in coefficients

$$a_{k,i} = \begin{cases} 1 & \text{if positive incident} \\ -1 & \text{if negative incident} \\ 0 & \text{otherwise.} \end{cases}$$

All coefficients are combined into a matrix $\mathbf{A} \in \{-1, 0, 1\}^{M \times (N+2M)}$. The surface continuation constraint for edge k simply says that for each edge there are as many positive incident as negative incident active elements:

$$\sum_i a_{k,i} y_i + \sum_j a_{k,N+j} y_{N+j} = 0.$$

Minimizing (2) can now be cast as solving the integer linear program

$$\begin{aligned} \min_{\mathbf{y}} \quad & \mathbf{w}^\top \mathbf{y} \\ \text{s.t.} \quad & \mathbf{A} \mathbf{y} = \mathbf{0} \\ & y_l \in \{0, 1\}, l = 1, \dots, N+2M. \end{aligned} \quad (3)$$

3. Generalization to Curvature Regularity

We now extend the above framework to include curvature regularity. Here we focus on functionals of the form

$$E(R) = \int_R g(\mathbf{x}) d\mathbf{x} + \nu|C| + \lambda \int_C |\kappa_C(\mathbf{x})|^p d\mathcal{H}^1(\mathbf{x}), \quad (4)$$

where compared to (2) the integral of some power $p > 0$ of the curvature κ_C of the region boundary C has been added. In fact, the proposed method can handle arbitrary dependences on position, tangent angle and curvature at points on the region boundary. One merely has to change the weight vector.

3.1. Curvature Cost via Pairs

To include dependences on curvature into the above described method, we consider *pairs* of adjacent line segments instead of single line segments. We only consider pairs corresponding to consistent orientations and denote \mathcal{F} the set of all such pairs.

We are now interested in which pairs of adjacent line segments occur in the boundary of the foreground region. This is again expressed by indicator variables, where for simplicity the indicator variable for the pair (j, j') is denoted $\bar{y}_{jj'}$. That is, we no longer give an order of all variables. Region indicators and indicators on pairs are combined into a vector $\bar{\mathbf{y}}$. This vector is paired with a weight vector $\bar{\mathbf{w}}$. While the entries \bar{w}_i for the region indicators are as above, the entries $\bar{w}_{jj'}$ now include a curvature cost for a line segment pair. If the lines have length l_j and $l_{j'}$, and between the lines the tangent angle changes by α , following Bruckstein et al. [6] the weight is given by

$$\bar{w}_{jj'} = \nu l_j + \lambda \min\{l_j, l_{j'}\} \left(\frac{\alpha}{\min\{l_j, l_{j'}\}} \right)^p.$$

Although the length of segment j' is not reflected in this weight, indeed the length of the boundary is correctly represented: the constraints stated in the next section ensure that each active boundary segment is present in two active pairs (once as the first and once as the second entry).

3.2. Deriving the Constraint Set

To ensure a consistent solution, two types of constraints are introduced. First of all, the surface continuation constraints are transferred to the new setting. To this end, the incidence of pairs of line segments to edges is defined. Our definition is based on the first of the two line segments. This is expressed by coefficients

$$b_{k,jj'} = \begin{cases} 1 & \text{if line } j \text{ is positive incident to edge } k \\ -1 & \text{if line } j \text{ is negative incident to edge } k \\ 0 & \text{otherwise.} \end{cases}$$

In other words $b_{k,jj'} = a_{k,N+j}$. When also defining $b_{k,i} = a_{k,i}$ for the region indicators, the surface continuation constraint for edge k reads

$$\sum_i b_{k,i} y_i + \sum_{(j,j') \in \mathcal{F}} b_{k,jj'} \bar{y}_{jj'} = 0.$$

In addition, there is another type of constraint that has to be imposed:

Boundary continuation constraints. If a pair of line segments is part of the boundary, there is another pair of line segments in the boundary that contains the second line segment as the first entry. Likewise, there is a pair that contains the first line segment as the second entry.

Here we have one constraint for every line segment. To formalize it we define coefficients

$$b_{M+l,jj'} = \begin{cases} 1 & \text{if } l = j \\ -1 & \text{if } l = j' \\ 0 & \text{otherwise} \end{cases}.$$

The above constraint set can now be expressed as equations

$$\sum_{(j,j') \in \mathcal{F}} b_{M+l,jj'} \bar{y}_{jj'} = 0.$$

In fact, these equations occur in standard shortest path methods.

When choosing a suitable order of the indicator variables, all coefficients can be combined into a matrix \mathbf{B} and we obtain the integer linear program

$$\begin{aligned} \min_{\bar{\mathbf{y}}} \quad & \bar{\mathbf{w}}^\top \bar{\mathbf{y}} \\ \text{s.t.} \quad & \mathbf{B} \bar{\mathbf{y}} = \mathbf{0} \\ & \bar{y}_l \in \{0, 1\}, \forall l. \end{aligned} \quad (5)$$

3.3. Optimization via LP-Relaxation

In general, solving integer linear programs is an NP-hard problem (e.g. [20, chapter 18.1]). We therefore proceed by solving the associated linear programming relaxation, where we use the dual simplex method as implemented in the Clp Package¹.

In some cases, solving the LP-relaxation results in integral solutions. Hence, in such cases the relaxation is tight and the globally optimal solution is obtained. In general, however, one finds fractional solutions that are not convex combinations of feasible integral solutions. From such a solution we obtain a (suboptimal) integral solution by thresholding the region variables and extracting the corresponding integral boundary variables.

¹The Clp package is available online at <http://www.coin-or.org/Clp>. It is based on the COIN library <http://www.coin-or.org>.

To evaluate the quality of the resulting solution, we compute the energy gap between the fractional and the integral solution. The cost of the optimal integral solution must be in between the two energies. Experimentally we have found the gap to be rather small, with a maximum of 2% relative.

4. Inpainting

The presented framework can be easily extended to the task of inpainting, i.e. to fill in damaged regions of a given image $I : \Omega \rightarrow \mathbb{R}$. We denote the minimal and maximal intensity level I_l and I_u respectively.

As discussed by Masnou and Morel [14], energies of the type

$$\int_{\Omega} |\nabla I(\mathbf{x})| \left| \operatorname{div} \left(\frac{\nabla I(\mathbf{x})}{|\nabla I(\mathbf{x})|} \right) \right|^p d\mathbf{x}$$

can be rewritten as integrals along level lines, i.e. as

$$\int_{I_l}^{I_u} \int_{\Gamma_t} |\kappa_{\Gamma_t}(\mathbf{x})|^p d\mathcal{H}^1(\mathbf{x}) dt,$$

where Γ_t denotes all level lines for level t .

For simply-connected inpainting domains, Masnou and Morel showed how to minimize the case of $p = 1$ globally via dynamic programming. The scheme presented in this work allows to handle arbitrary domains and arbitrary exponents p . In general the solutions will not be globally optimal: Ideally the boundary variables reflect level lines but for non-binary variables the corresponding solutions will not reflect level lines.

To adjust the scheme to inpainting, one only has to change the integrality constraints to

$$y_i \in [I_l, I_u], \quad i = 1, \dots, N$$

for the variables representing basic regions and to

$$\bar{y}_{jj'} \in [0, I_u]$$

for the variables representing boundary segments. In addition, one must fix the region variables on the border of the damaged region to the given intensities. Then, after setting the weights \bar{w}_i for the region variables to 0, the algorithm is run.

5. Experiments

We evaluate the proposed method on both image segmentation and inpainting tasks. For the accurate reflection of curvature it is important to have a sufficiently fine discretization of directions. To this end, pixels are subdivided into several basic regions as shown in Figure 5. The employed subdivisions correspond to 8- or 16-connectivities in graph cut frameworks. In all cases we use a curvature power of $p = 2$. The experiments were run on a Core 2 machine with 2.66 Ghz.




Figure 5. The chosen basic regions for a pixel and their equivalent connectivities for graph cuts.




Figure 6. Intensity-based segmentation with different curvature weights.

5.1. Image Segmentation

Our first experiment for image segmentation combines the data term of the piecewise constant Mumford-Shah functional [15] with a regularity term consisting of squared




Figure 8. Effect of the connectivity. Both results are globally optimal.

curvature:

$$\int_R (I(\mathbf{x}) - \mu_F)^2 d\mathbf{x} + \int_{\Omega \setminus R} (I(\mathbf{x}) - \mu_B)^2 d\mathbf{x} + \lambda \int_C |\kappa_C(\mathbf{x})|^2 d\mathcal{H}^1(\mathbf{x}).$$

Here we fix the mean values μ_F, μ_B based on the minimal and maximal intensities in the image. Figure 6 shows results on the well-known camera man image for different curvature weights. In this case the system to solve involves roughly 1 million variables and 300.000 constraints.

The results demonstrate that even with large weights one can segment long and thin objects. Moreover, the LP-relaxation is quite tight here: all shown results are within 2% relative² of the global optimum. They were generated using the 8-connectivity, where the run-times vary between 10 and 20 minutes.

In a second experiment the approach is extended to include seed nodes and the data terms are replaced by the negative logarithm of color histograms p_F and p_B computed for the user-labeled pixels. This time also a length term is included:

$$-\int_R \log(p_F(I(\mathbf{x}))) d\mathbf{x} - \int_{\Omega \setminus R} \log(p_B(I(\mathbf{x}))) d\mathbf{x} + \nu|C| + \lambda \int_C |\kappa_C(\mathbf{x})|^2 d\mathcal{H}^1(\mathbf{x}).$$

In Figure 7, basic regions corresponding to the 8-connectivity were used as the 16-connectivity gave comparable results.

However, Figure 8 shows that usually the 16-connectivity is preferable. For the latter there were roughly 500.000 basic regions. In Figure 9 it is shown that on these data curvature helps much to achieve connectedness. Moreover, here the global optima are found.

²This is w.r.t. the original energy as in (1) (but with curvature term), not the transformed one as in (2).




Figure 7. Comparison of length and curvature. We show results for the highest weights that do not disconnect the two front legs from the torso. The computed solutions are very close to the global optimum.




Figure 9. Curvature helps to establish connectedness. We show the visually best results for all approaches, using the 16-connectivity.

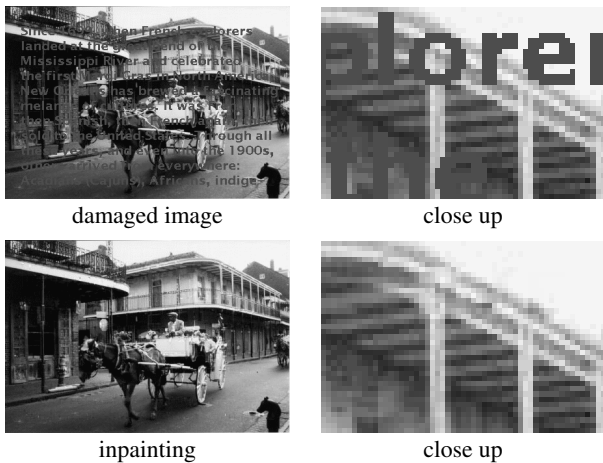


Figure 10. The proposed method is useful for inpainting.

The result in Figure 1 was also generated with the 16-connectivity. It was computed in 3.5 hours and is the global optimum of the energy. Note that curvature regularity allows to well preserve thin and elongated structures which are typically suppressed by algorithms which rely on length-based regularity.

5.2. Inpainting

We now turn to the problem of inpainting, where we present two experiments using the 8-connectivity. Figure 10 shows a result on the well-known New Orleans image. The close ups show that the proposed method works excellently in filling in the destroyed regions.

In Figure 11 we provide a comparison of length-based and curvature-based inpainting for a Japanese temple. Here it becomes evident that curvature terms are important as otherwise the regions are filled in with mostly homogeneous intensity values.

6. Conclusion

We proposed the first approach to minimize curvature regularity for region-based image segmentation that is independent of initialization. In addition, we showed that the method can be extended to the problem of inpainting. In both these applications, curvature plays an important role. To this end, the problem is formulated as an integer linear program and its LP-relaxation is solved. By thresholding the region variables an integral solution is obtained. The experiments indicate that the produced solutions are close to the global optimum.




Figure 11. Curvature is better suited for inpainting than length.

In numerous experiments on segmentation and inpainting we show that curvature regularity helps to substantially improve results with respect to algorithms which merely impose length regularity.

As we are representing both basic regions as well as pairwise line segments in our formulation, the sizes of the corresponding optimization problems become quite large even for moderate image resolutions. Nevertheless, in this work we have used a standard implementation of the dual simplex method for computing the solutions and obtained reasonable running times. To develop specialized solvers exploiting the particular problem structure at hand is an avenue of future research.

Acknowledgments We thank Simon Masnou for helpful comments on inpainting. We also thank Sara Vicente, Vladimir Kolmogorov and Carsten Rother for sharing data.

References

- [1] 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. **2**
- [2] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *ACM SIGGRAPH*, pages 417–424, New Orleans, Louisiana, July 2000. **1**
- [3] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987. **1**
- [4] F. Bornemann and T. März. Fast image inpainting based on coherence transport. *Jour. Math. Imaging and Vision*, 28(3):259–278, July 2007. **1**
- [5] Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in n-d images. In

- IEEE Int. Conf. on Comp. Vision*, volume 1, pages 105–112, Vancouver, Canada, 2001. **1**
- [6] A.M. Bruckstein, A.N. Netravali, and T.J. Richardson. Epicongvergence of discrete elastica. In *Applicable Analysis, Bob Carroll Special Issue*, volume 79, pages 137–171, 2001. **3**
- [7] T.F. Chan, S.H. Kang, and J. Shen. Euler’s elastica and curvature based inpainting. *SIAM Journal of Applied Mathematics*, 63:564–592, 2002. **2**
- [8] T.F. Chan and L.A. Vese. Active contours without edges. *IEEE Trans. on Image Proc.*, 10(2):266–277, 2001. **1**
- [9] S. Esedoglu and R. March. Segmentation with depth but without detecting junctions. *Jour. Math. Imaging and Vision*, 18:7–15, 2003. **1, 2**
- [10] L. Grady. Minimal surfaces extend shortest path segmentation methods to 3d. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 2009. To appear. <http://cns-web.bu.edu/lgrady/>. **2**
- [11] D.M. Greig, B.T. Porteous, and A.H. Seheult. Exact maximum *a posteriori* estimation for binary images. *Jour. Royal Stat. Soc., Series B*, 51(2):271–279, 1989. **1, 2**
- [12] G. Kanizsa. Contours without gradients or cognitive contours. *Italian Jour. Psych.*, 1:93–112, 1971. **1**
- [13] M. Klodt, T. Schoenemann, K. Kolev, M. Schikora, and D. Cremers. An experimental comparison of discrete and continuous shape optimization methods. In *Europ. Conf. on Comp. Vision*, Marseille, France, October 2008. **1**
- [14] 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, 2, 4**
- [15] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989. **1, 5**
- [16] 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. **1, 2**
- [17] M. Nitzberg, D. Mumford, and T. Shiota. *Filtering, Segmentation and Depth*, volume 662 of *LNCS*. Springer Verlag, Berlin, 1993. **1, 2**
- [18] 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. **2**
- [19] 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. **2**
- [20] A. Schrijver. *Theory of linear and integer programming*. Wiley-Interscience series in discrete mathematics. John Wiley and Sons, July 1994. **4**
- [21] J.M. Sullivan. Computing hypersurfaces which minimize surface energy plus bulk energy. *Motion by Mean Curvature and Related Topics*, pages 186–197, 1994. **2**
- [22] D. Tschumperlé. Fast anisotropic smoothing of multi-valued images using curvature-preserving PDE’s. *Int. Jour. of Comp. Vision*, 68(1):65–82, June 2006. **1, 2**