

Efficient Squared Curvature

Claudia Nieuwenhuis^{*,†} Eno Toeppe^{*,†} Lena Gorelick[†] Olga Veksler[†] Yuri Boykov[†]
^{*}UC Berkeley [†]University of Western Ontario

Abstract

Curvature has received increasing attention as an important alternative to length based regularization in computer vision. In contrast to length, it preserves elongated structures and fine details. Existing approaches are either inefficient, or have low angular resolution and yield results with strong block artifacts. We derive a new model for computing squared curvature based on integral geometry. The model counts responses of straight line triple cliques. The corresponding energy decomposes into submodular and supermodular pairwise potentials. We show that this energy can be efficiently minimized even for high angular resolutions using the trust region framework. Our results confirm that we obtain accurate and visually pleasing solutions without strong artifacts at reasonable runtimes.

1. Introduction

A number of vision tasks can be formulated as an energy minimization problem such as segmentation, 3D reconstruction, stereo and inpainting. The corresponding energies typically consist of a data affiliation term and a regularization term. The data affiliation term relates the solution to the image data, while the regularization term imposes some kind of prior knowledge on the result. Length-based regularizers give rise to sub-modular energies (Potts model) that can be globally and efficiently optimized and, therefore, are widely used. Their main disadvantage is an inherent shrinking bias, which tends to eliminate thin and elongated structures such as vessels or limbs, see Figure 1(b). In contrast, curvature-based regularizers preserve such fine details, but are more difficult to model, often resulting in non-submodular energies that are hard to optimize.

Previous approaches to curvature regularization are computationally intensive [14, 4, 8, 16], have strong discretization artifacts [14, 16, 6, 8, 15, 4] and are often restricted to specific angular resolutions [6] or grid complexes [14, 16, 8]. Curvature is often combined with length to compute Euler's elastica, which can alleviate some of the problems of

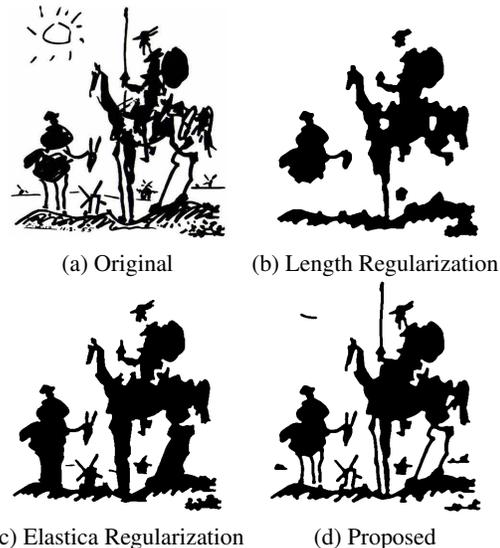


Figure 1. Segmentation results of Picasso's ink drawing using (b) length based regularization computed with the Potts model based on a 16 neighborhood using [2], (c) Euler's elastica combining squared curvature with length regularization, (d) the novel squared curvature regularizer.

pure curvature regularizers, see Figure 1(c). However, satisfactory solutions to the curvature regularization problem have not been found yet. In this paper we present a novel approach to modeling and efficiently optimizing squared curvature, which yields visually pleasing results without strong discretization artifacts, see Figure 1(d).

Our goal is to regularize the boundary of a binary labeling S based on an integral of squared curvature

$$K(S) = \int_{\partial S} \kappa^2 \cdot ds. \quad (1)$$

For example, a segmentation energy can combine this regularization term with a regional appearance term $\int_{int(S)} f(p) dp$ for some potential function f over pixels p .

We optimize a binary energy function with triple clique potentials, which can be reduced to an integer quadratic energy with both submodular and supermodular pairwise terms without adding auxiliary variables.

Such energies can be efficiently optimized with the LSA-

^{*}The first author was in part supported by the DAAD (German Academic Exchange Service)

TR method [7]. Our formalism is based on integral geometry, allows for high angular resolutions and yields excellent results compared to previous approaches.

Our contributions can be summarized as follows

- We propose a novel model for measuring squared curvature based on integral geometry and show how it relates to counting straight triple cliques.
- Our model can be formulated as a pairwise quadratic energy and optimized efficiently even for high angular resolutions.
- The proposed approach outperforms previous methods in terms of quality of the results as well as efficiency.

Related work: Most existing models for curvature regularization in vision are based on Bruckstein’s formula [14, 16, 8, 6] except for approaches based on learning [15, 11] or functional lifting [4]. The different formulations are optimized with methods from discrete or continuous optimization [10].

Olsson et al. [11] express curvature regularization as a multi-label optimization problem based on a set of local neighborhood patterns with learned weights and consistency constraints between neighboring patterns. This approach works well for small patterns but the runtime grows exponentially with the pattern size. Thus angular resolution is quite limited in practice. Shekhovtsov et al. [15] learn the costs for curvature in the form of “soft” patterns which serve as filters whose response is locally minimized. However, their MRF approach lacks accuracy due to missing consistency constraints.

Approaches based on the Bruckstein formula [5] express curvature as the exterior angle sum of an approximating polygon. Mostly this is done by formulating an optimization problem on a cell complex, which is a planar graph with fixed, regular structure. A segmentation is a consistent subset of faces and edges. Local “curvature” is then measured by the exterior angle between consecutive boundary edges.

The pioneering approach to image segmentation with curvature regularization on cell complexes was given by Schoenemann et al. [14]. They solve a linear program where each variable corresponds to a configuration of two edges with associated angle. Strandmark and Kahl [16] improve this framework by removing extraneous arcs and generalize it to 3D surfaces. Heber et al. [8] provide a formulation of the curvature model on cell complexes that can be optimized by approximating the envelope of the underlying non-convex functional.

The main drawback of the approaches formulated on cell complexes is their high runtime up to several hours, which is due to the large number of variables and consistency constraints. In addition, these methods suffer from a strong angular bias since they only allow for specific edge configurations. Although angular resolution can in theory be

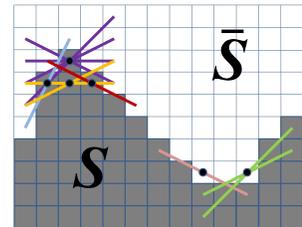
increased arbitrarily, solving the problem for high resolutions easily becomes infeasible. This is not the case for our approach since the number of triple cliques grows linearly with the angular resolution.

Another approach based on the Bruckstein formula which is related to cell complexes is El-Zehiry and Grady’s work [6]. They formulate their problem on a regular pixel grid, which is interpreted as a cell complex. Accordingly, their angular resolution is limited by 90 degrees, which leads to a coarse approximation of curvature. Extending this approach to a higher angular resolution is not possible due to inconsistencies with the Bruckstein formula.

Our method differs from the previous approaches in that we neither explicitly model our segmentation boundary as a polygon by means of cell complexes, nor do we measure angles between edges. Instead, we count the number of “active” straight line triple cliques and relate it to squared curvature based on integral geometry.

2. Integrating Squared Curvature

We propose a new discrete model for approximating the squared curvature integral $K(S)$ in (1) based on a certain class of triple cliques. Our combinatorial approach is presented below in the context of 2D segmentation and inpainting. It has straightforward extensions to segmentation problems in higher dimensions, but this is left for future work, see Section 5. Unlike previous discrete methods for curvature, we use a neighborhood system of symmetric ‘straight’ triple cliques. Our general intuition is illustrated on the right-hand side. Let the indicator variables x_p on pixels $p \in \Omega$ define the segment $S = \{p \mid x_p = 1\} \subset \Omega$ and consider a collection of triple cliques $(p - \Delta, p, p + \Delta)$ for each pixel p and its symmetric neighbors for fixed 2D shifts $\pm\Delta$ within a 5×5 neighborhood. The colors show triple cliques with binary configurations $(0, 1, 0)$ or $(1, 0, 1)$ at different pixels p . Such triple cliques ‘fire’ only on curved parts of the boundary of S . The number of responses increases for larger curvature. Configurations $(0, 1, 0)$ respond to positive and $(1, 0, 1)$ to negative curvature. These observations suggest that such triple cliques can measure curvature.



2.1. Notation: Variables, Cliques, Neighborhood

The proposed neighborhood structure is somewhat similar to the standard regular neighborhood of pairwise cliques (edges) commonly used for length-based regularization [3]. At each pixel p the neighborhood is bounded by a box centered at this pixel, see Figure 3. Interactions are imposed between pixel p and its m neighbors in a certain discrete

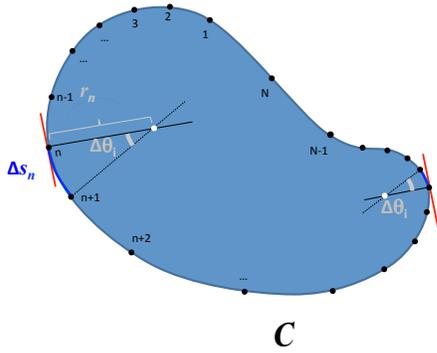


Figure 2. The partial sum approximation (3) of the integral $\int_C \kappa^2 ds$ uses a sequence of points on the contour C . Our approach is to select all contour points $\{n\}_1^N$ where the tangent is consistent with the orientation of some triple clique in a given neighborhood in Figure 3(a). For example, two points in this Figure are marked with a tangent of the same orientation as clique c_i in Figure 3(a). Each point n corresponds to a particular clique index $i(n)$.

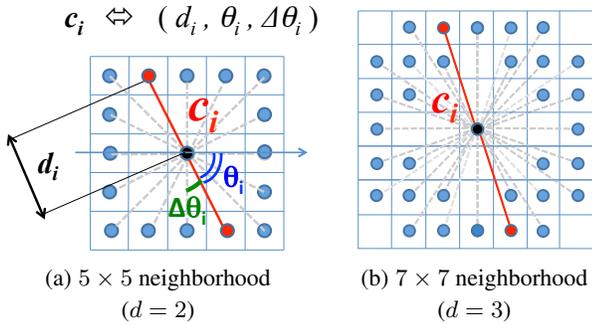


Figure 3. Neighborhood systems of size $(2d+1) \times (2d+1)$ for (a) $d = 2$ and (b) $d = 3$. To make all triple cliques approximately the same length, for each orientation $\theta_i \in \Theta$ we select the centered triple clique c_i whose length d_i is as close to d as possible. θ_i indicates the orientation of c_i with respect to the horizontal axis. $\Delta\theta_i$ denotes the angular rotation to the next clique.

set of directions $\Theta = \{\theta_i | i = 1, \dots, m\}$ limited by the grid locations within the box. Standard length-based regularization [3] corresponds to pairwise Potts interactions imposed over p and each neighbor $q = p \oplus i := p + \bar{d}_i$ where \bar{d}_i denotes a shift of given length d_i in direction θ_i . Our new curvature-based regularization model uses interaction potentials imposed over p and pairs of symmetric neighbors in directions \bar{d}_i and $-\bar{d}_i$ forming a straight triple clique

$$c_i(p) := (p \ominus i, p, p \oplus i) \quad (2)$$

centered at pixel p , see Figure 3.

We will use standard binary variables $x_p \in \{0, 1\}$ to denote the object/background label at pixel p . Vector $X = (x_p | p \in \Omega)$ will denote a configuration of binary labels of all pixels in Ω defining segment $S = \{p | x_p = 1\} \subset \Omega$. Vector $X_c = (x_{p \ominus i}, x_p, x_{p \oplus i})$ will denote a configuration

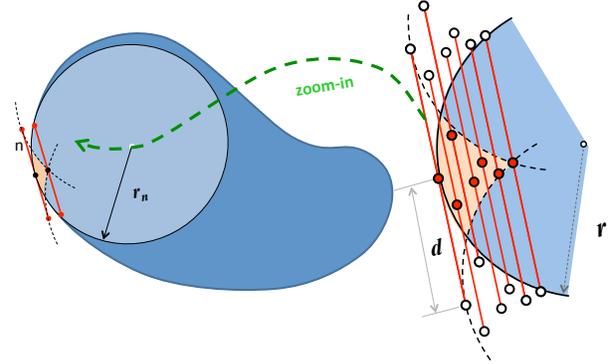


Figure 4. Examples of triple cliques of orientation θ_i that “fire” in a vicinity of point n on contour C in Figure 2. The brown region highlights points p where clique $c_i(p) = (p \ominus i, p, p \oplus i)$ has configuration $X_c = (0, 1, 0)$. The area of this brown region $A(\kappa, d)$ in (13), that is the number of such triple cliques, is determined by local curvature $\kappa_n = \frac{1}{r_n}$ and by clique size d_i .

of labels for pixels in clique $c_i(p)$.

Each triple clique c_i centered at pixel p is described by its orientation θ_i and distance d_i , see Figure 3(a). We will also use $\Delta\theta_i := \theta_{i+1} - \theta_i$ to denote the angular rotation to the next clique. Since our triple cliques include symmetric pairs of neighbors, the actual number of distinct triple cliques $c_i(p)$ centered at pixel p is half the number of its neighbors. To avoid confusion, in the rest of the paper m denotes the number of distinct triple cliques c_i at pixel p where $i \in \{1, \dots, m\}$ is an index of orientation in the set $\Theta = \{\theta_i | i = 1, \dots, m\}$ such that $\Theta \subset [0, \pi)$. This avoids identical cliques with orientations θ_i and $\pi + \theta_i$.

2.2. Partial Sum Approximation of (1)

Our discrete model for evaluating the integral of squared curvature (1) is based on a partial sum approximation. Without loss of generality, assume that the (continuous) segment S has genus zero and that its boundary ∂S is a closed contour C , as shown in Figure 2. Then, the integral in (1) can be approximated as

$$\int_C \kappa^2 \cdot ds \approx \sum_{n=1}^N \kappa_n^2 \cdot \Delta s_n \quad (3)$$

where $\{n\}_1^N$ is a finite sequence of points $n \in C$, Δs_n is a contour length between adjacent points, and κ_n is the curvature at point n . In general, such approximations converge for finer discretizations $\{n\}_1^N$ as $\max_n |\Delta s_n| \rightarrow 0$ if the contour C is sufficiently smooth.

It is common to select a sequence of approximating points $\{n\}_1^N$ at equal intervals. Our approach selects these points differently. We choose points $\{n\}_1^N$ where tangents coincide with orientations $\Theta = \{\theta_i | i = 1, \dots, m\}$ for a chosen neighborhood system, as illustrated in Figures 2 and 3.

Assume that curvature is nearly constant between adjacent points, then $\Delta s_n \approx r_n \cdot \Delta \theta_{i(n)} = \frac{\Delta \theta_{i(n)}}{|\kappa_n|}$ where r_n is the radius of the osculating circle at point n and $\theta_{i(n)} \in \Theta$ is the tangent orientation at point n . Thus, the partial sum (3) becomes

$$\sum_{n=1}^N \kappa_n^2 \cdot \Delta s_n \approx \sum_{n=1}^N |\kappa_n| \cdot \Delta \theta_{i(n)}. \quad (4)$$

Theorem 1. *Let contour point n have tangent orientation \bar{u} and osculating ball B of radius $r = \frac{1}{|\kappa|}$. Then, the set of all points $p \in B$ such that $\|p - n\| \leq r$ and $(p \pm d \cdot \bar{u}) \notin B$ for given distance $d < r$ has area $A(\kappa, d) = \frac{|\kappa| \cdot d^3}{4} + O(d^4)$, see brown region in Figure 4.*

This theorem is proved in the Appendix (Section 6). It allows to accurately estimate curvature κ_n at contour point n using triple cliques $c_i(p) = (p \ominus i, p, p \oplus i)$ of orientation $\theta_{i(n)}$ consistent with the tangent at point n . Assuming $\kappa_n \geq 0$, Theorem 1 implies that the number of triple cliques $c_i(p)$ in a vicinity of point n with configuration $(0, 1, 0)$ is

$$A(\kappa_n, d_i) \approx \frac{|\kappa_n| \cdot d_i^3}{4},$$

which is the brown area in Figure 4. In case $\kappa_n \leq 0$, the same number A corresponds to configurations $(1, 0, 1)$.

Assuming that triple cliques $c_i(p)$ of configurations $(0, 1, 0)$ or $(1, 0, 1)$ have penalties w_i depending on orientation i , the overall cost of all such cliques is

$$\sum_i \sum_p w_i \cdot \delta(X_{c_i(p)}) = \sum_i \sum_{n:i(n)=i} w_i \cdot A(\kappa_n, d_i)$$

where

$$\delta(X_c) := \begin{cases} 1 & \text{if } X_c = (1, 0, 1) \text{ or } (0, 1, 0) \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

is an indicator function. Choosing penalty

$$w_i = \frac{4 \cdot \Delta \theta_i}{d_i^3} \quad (6)$$

converts the right hand side in the equation above into

$$\sum_i \sum_{n:i(n)=i} \frac{4 \cdot \Delta \theta_i}{d_i^3} \cdot A(\kappa_n, d_i) = \sum_{n=1}^N |\kappa_n| \cdot \Delta \theta_{i(n)}.$$

In combination with (3) and (4) this demonstrates that

$$\int_C \kappa^2 \cdot ds \approx \sum_i \sum_p w_i \cdot \delta(X_{c_i(p)}) \quad (7)$$

which is our main technical result concluding this section.

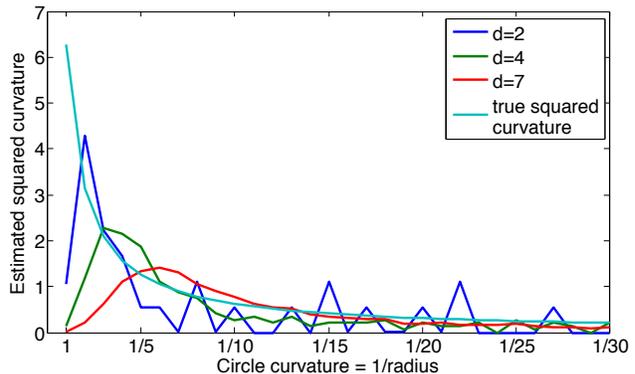


Figure 5. Accuracy of our squared curvature model. The proposed estimated squared curvature (vertical axis) is computed on a circle of decreasing total curvature (i.e. increasing radius, horizontal axis) based on increasing clique size d . To approximate large curvature small clique sizes d are required, whereas large clique sizes are more accurate to measure large curvature and yield less noisy measurements. For curvature values $\kappa > \frac{1}{d}$ we underestimate curvature, since in these cases the circle fits inside the neighborhood. This limits the estimated curvature energy by the number of pixels in the circle, i.e. its area, which grows quadratically.

2.3. Accuracy of Our Curvature Model

To demonstrate the accuracy of the proposed squared curvature model we generated a sequence of circles with decreasing total squared curvature (increasing radius r). We compared the computed curvature energy (7) obtained with different neighborhood sizes to the correct integral of squared curvature of the circle, $2\pi\kappa$. The plots in Figure 5 show that we indeed approximate squared curvature. We can also conclude that for smaller neighborhoods we more accurately measure larger curvature values but obtain more noisy results. In contrast, our results are less noisy for larger neighborhood masks but yield less accurate measurements for larger curvature. This is because our measurement is limited by the area of the circle, which is very small for small radius (large curvature) and grows quadratically. These findings are supported by Figure 6, which shows the contribution of each pixel to the measured curvature energy. To alleviate the issue of underestimating curvature for larger values of d the image can be scaled to subpixel accuracy according to the increase in neighborhood size or stronger regularization can be used, see Section 4.1.

3. Optimization

In (7) we derived our new squared curvature model, which can be used as a regularization energy

$$E_{curv}(X) = \sum_i \sum_p w_i \cdot \delta(x_p, x_{p \oplus i}, x_{p \ominus i}). \quad (8)$$

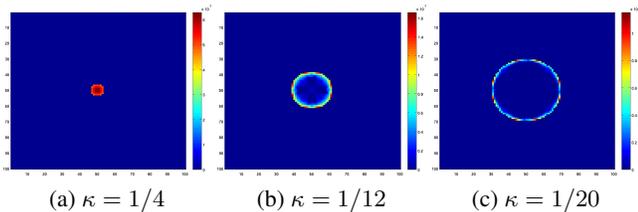


Figure 6. Response images showing contribution of each pixel to the curvature energy for circles of increasing radius (decreasing curvature), see Figure 13. The computations are based on a clique size $d = 9$. The images show that the brown area in Figure 4 is approximated with our approach. (a) also shows that large curvature is underestimated for large values of d since the response is limited by the area of the circle (compare Figure 5).

Let a, b, c denote the three pixels corresponding to a triple clique $c_i(p)$ given in (2). The function δ defined in (5) assigns each triple clique the value 1 if its configuration is (1,0,1) or (0,1,0). The corresponding triple clique energy in (8) naturally decomposes into a unary term, two pairwise submodular terms and a pairwise supermodular term without additional auxiliary variables

$$\begin{aligned} \delta(x_a, x_b, x_c) &= x_a(1 - x_b)x_c + (1 - x_a)x_b(1 - x_c) \\ &= x_b + x_ax_c - x_ax_b - x_bx_c. \end{aligned} \tag{9}$$

In general, energies of lower order are considered to be easier to optimize. This is an additional benefit of our integral geometry model that we get for ‘free’.

Our energy is nevertheless non-submodular and, therefore, cannot be optimized globally. We use a new method called Local Submodular Approximations with Trust Region (LSA-TR) proposed in [7]. The method is efficient and obtains state-of-the-art results that are very accurate on a wide range of applications.

We compare our results in terms of accuracy and runtime to other optimization approaches (TRW-S [9], LBP [12], QPBO-I [13] - an extension of QPBO [1]) in the experimental section 4.4. Throughout all experiments we use the same set of default parameters for LSA-TR.

4. Experiments

In this section we evaluate our squared curvature approach with respect to quality of the results and optimization and show a comparison to previous curvature regularization approaches.

We show results for binary segmentation and inpainting which minimize the following energy

$$E(X) = \sum_i D(x_i, I_i) + \lambda E_{curv}(X). \tag{10}$$

The data term $D(x_i, I_i)$ depends on the color of the i -th pixel in the image I . The weight λ balances the impact

of the regularizer with respect to the appearance term. We use a Gaussian of variance 0.4 for modeling the foreground and background data term in our experiments. The mean values are 0 and 0.6 for foreground and background respectively for the camera man examples and 0 and 1 for the Don Quixote image. For inpainting, we use the same energy but set the data term to 1 for foreground and background simultaneously for the region to be inpainted.

Note that we indicate the specific case of only four (horizontal and vertical) neighbors in a 3×3 neighborhood by “ 3×3 (4nh)” throughout the experimental section. If no nh-number is indicated we mean the full neighborhood.

4.1. Triple Clique Length d

We first show the effect of larger triple cliques in Figure 7. The larger the length of the triple cliques the more we underestimate large curvature values (see Figure 5) leading to noise in the segmentation. Yet, we would like to use large neighborhoods to attain a larger angular resolution of our clique orientations as shown in the next section. As a remedy we propose to scale the image according to triple clique length d in order to achieve subpixel accuracy. In this way we do not underestimate large curvature and avoid angular artifacts. Alternatively, one can increase regularization to reduce the noise but this might lead to less details in the segmentation. These points are illustrated in Figure 7.

4.2. Angular Resolution

We now show the effect of increasing angular resolution on the segmentation results. Figure 8 shows results for angular resolutions θ of 90 degrees: 3×3 (4nh), 45 degrees: 3×3 and 12.5 degrees: 7×7 . For small curvature weight the appearance is strong yielding comparable results for all angular resolutions. For very large curvature weight block structures become apparent. This is due to the fact that the algorithm is blind to curvature for points whose tangent direction is not in our clique set. Accordingly, for only 90 degree resolution horizontal and vertical edges become dominant in the result. For 45 degree resolution diagonal edges appear as well. For larger resolution we obtain smooth boundaries.

4.3. Inpainting Results

In this section we show results for the task of binary inpainting demonstrating that our approach is able to reconstruct shape with minimal curvature instead of length. We show two examples in Figure 9, for which we mask larger portions of the image for which we “occlude” its data: the data term in these regions is set to a constant value for both foreground and background. For the remaining part of the image we use a data term based on a Gaussian intensity model with fixed mean and variance for foreground and background. Figure 9 shows the results of our exper-

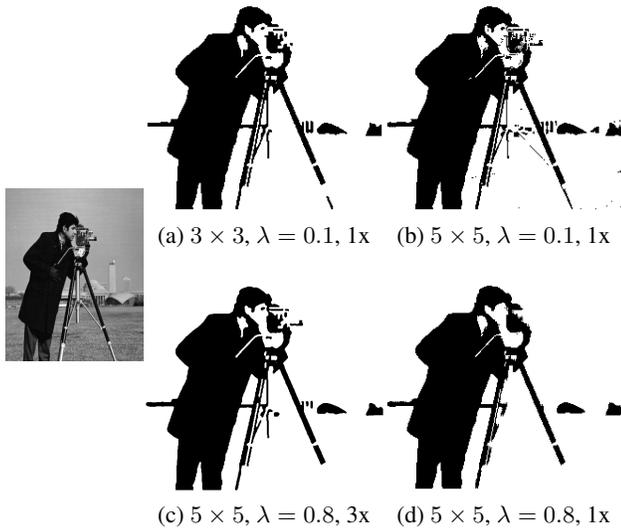


Figure 7. Segmentation results on the camera man image (left) showing the effect of using large triple cliques. (a) Result with novel squared curvature regularizer using a 3×3 neighborhood (8-connectivity) ($\lambda = 0.1$). (b) The same segmentation using a 5×5 neighborhood (16-connectivity). Note the noise which is due to underestimating large curvature locally (see text for details). (c) By scaling the image (and adapting λ accordingly), we avoid underestimation while taking advantage of a bigger angular resolution. (d) Similar effects but with less details can be achieved by increasing the regularizer weight instead of scaling the image.

iments and a comparison to length based inpainting. The results clearly show that we minimize the curvature of the object boundary, while length finds the shortest connection between boundary segments.

4.4. Optimization and Efficiency

We evaluated the performance and runtime of different optimization strategies for minimizing the energy in (10). In particular, we compared QPBO-I [13], TRW-S [9], Loopy Belief Propagation (LBP) [12] and LSA-TR [7]. For TRW-S we stopped computations after 50,000 steps without convergence of the algorithm. Running both TRW-S and LBP even for 500,000 steps did not improve the results. Figure 10 shows the energies we obtain for the respective methods on the vertical axis plotted against the regularization weight λ on the horizontal axis. Of all the tested algorithms, LSA-TR finds the lowest energy for our problem for almost all values λ . This is especially true for higher curvature weights, where QPBO-I and TRW-S compute trivial solutions of higher energy with almost all pixels labeled as background. In particular, QPBO is unable to label any pixels at all for $\lambda \geq 0.5$ and therefore QPBO-I cannot improve the result either. The high energies we get with LBP for $\lambda \geq 1$ correlates with the extremely noisy results the algorithm returns. Figure 11 shows a comparison in runtime re-

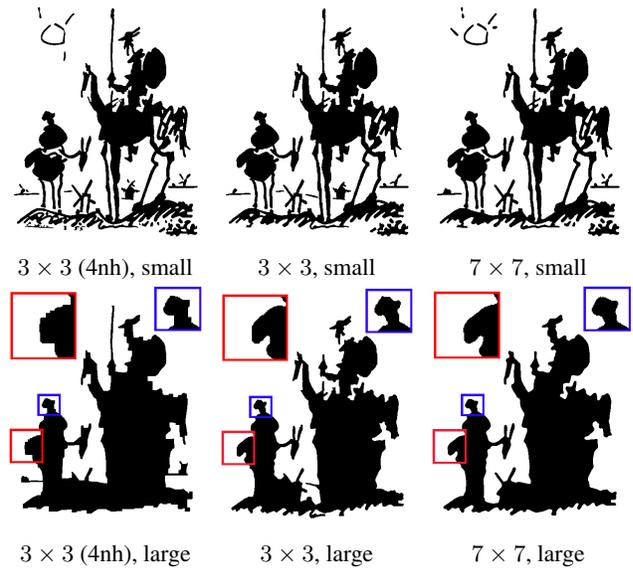


Figure 8. Segmentation results for Picasso's ink drawing for small and large curvature regularization weight and increasing neighborhood size. Note that the choice of the neighborhood size only becomes important for larger regularization weights. The images clearly show how 'blocky' structures disappear with larger neighborhood sizes.

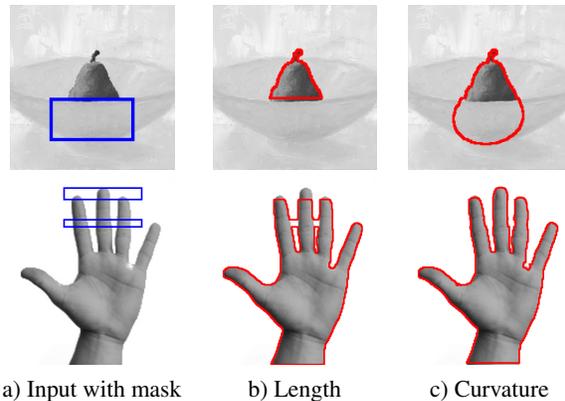


Figure 9. Results for binary inpainting: a) Original image with masked regions to be inpainted, b) Result of length based inpainting, c) Result of curvature based inpainting with our approach. The figure is best viewed in color.

vealing that TRW-S is least efficient of the four algorithms, whereas LSA-TR is the fastest optimization method. The proposed method also compares remarkably well to other approaches that compute squared curvature. For the results in Figure 12 the runtimes are as follows: Heber et al. [8] 1 to 5 minutes, Schoenemann et al. [14] 10 minutes to 3.5 hours, Bredies et al. [4] 5 minutes, Olsson et al. [11] between 0.9 seconds for 2×2 patches and 2.5 hours for 5×5 patches, and El-Zehiry and Grady's [6] 10 seconds per image. Strandmark and Kahl's approach [16] is only slightly

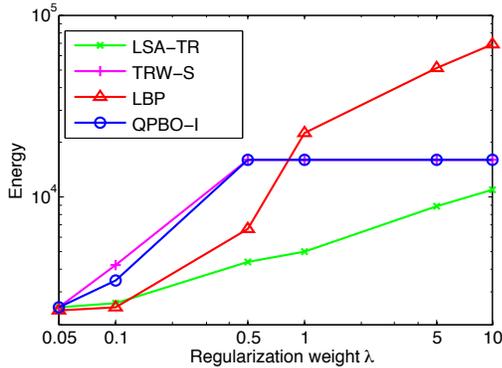


Figure 10. Comparison of optimized energies based on different regularization strengths λ for the following optimization algorithms: LSA-TR [7], TRW-S [9], LBP [12], QPBO-I [13]. LSA-TR clearly outperforms the other algorithms yielding reasonable results even for strong regularization, while the other approaches break down with trivial (TRWS, QPBO-I) or extremely noisy solutions (LBP).

λ	QPBO	TRW-S	LBP	LSA-TR (3x3)	LSA-TR (7x7)
0.05	0.7 (2%)	†	†	0.8	0.7
0.1	9.3 (30%)	†	†	1.4	1.6
0.5	15.2 (100%)	†	†	3.9	9.2
1	15.4 (100%)	†	†	10.6	48.2
5	15.3 (100%)	†	†	16.5	88
10	15.4 (100%)	†	†	7.9	110

Figure 11. Runtime comparison for different optimization methods based on a 3×3 neighborhood. For QPBO [1] the percentage of unlabeled pixels is indicated in parenthesis. For TRW-S and LBP we stopped computations after 50,000 steps (around 12.5 minutes) without convergence. LSA-TR is also listed for a 7×7 neighborhood with $\lambda \times 10$ to obtain comparable results.

faster than Schoenemann’s.

4.5. Comparison to Other Approaches

We compare our curvature regularizer to previous approaches in Figure 12 where we show results for small and large regularization for various methods. For all of these methods strong artifacts are evident, which become worse for stronger regularization. Heber et al. [8] compute elastica (i.e. length and squared curvature regularization), which preserves long structures, but also introduces block artifacts. El-Zehiry and Grady are limited by 90 degree resolution, i.e. their results are composed of blocks and contain large regions of unlabeled pixels QPBO [1] did not label. Schoenemann et al. [14] obtain block artifacts with increasing regularization. Strandmark and Kahl [16] do not preserve fine details and also suffer from some angular resolution artifacts. Bredies et al. [4] do not preserve elongated structures well for strong regularization and introduce block artifacts. Olsson et al. [11] introduce fine block artifacts. In

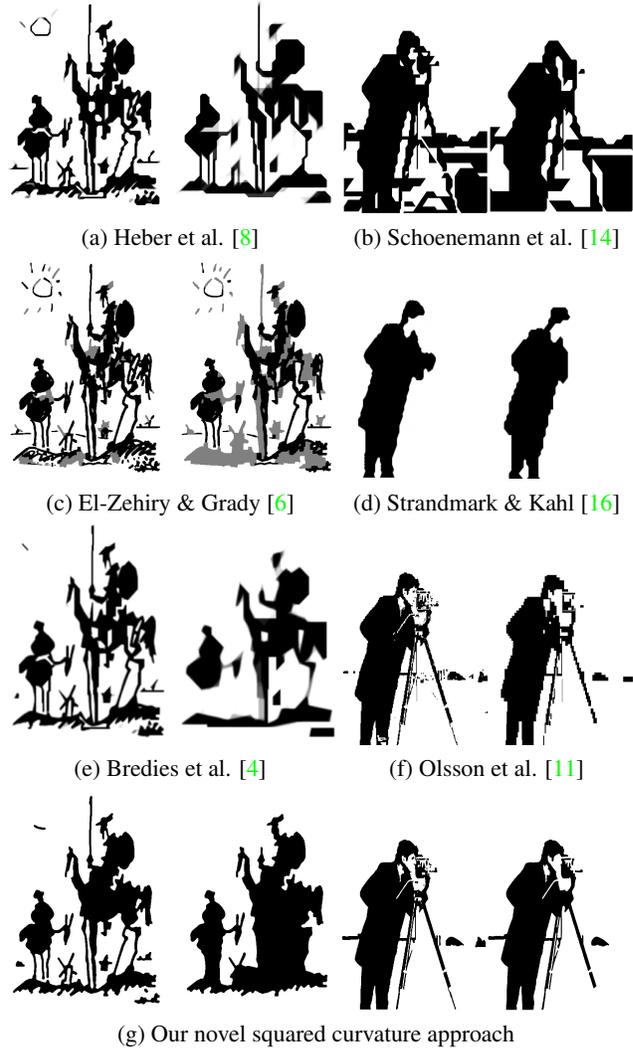


Figure 12. Comparison of our results for 7×7 neighborhood to previous curvature regularizers for smaller (left) and larger (right) curvature weight, see Figures 1 and 7 for original images. Due to lack of space we only show results on one image for each method. Grey pixels in (c) remain unassigned by QPBO - QPBO-I does not improve results. Grey pixels in (a) and (e) are due to non-binary solutions of the relaxed optimization problem.

contrast, our method (neighborhood size 7×7) preserves fine details in the segmentation and produces clean object boundaries for weak and strong regularization without requiring excessive runtimes.

5. Conclusion and Future Work

In this paper we proposed a novel approach to squared curvature computation and regularization. We gave an integral geometric derivation of our method which justifies usage of straight triple cliques for measuring squared curvature. Our triple cliques decompose into a set of submodular and supermodular pairwise cliques, which can be efficiently

optimized by LSA-TR. Our method outperforms previous methods in terms of quality and efficiency. Moreover, unlike previous methods our approach does not suffer from grid artifacts and works for high angular resolution.

Our approach naturally extends to higher dimensions for the regularization of 3D surfaces, which will be important e.g. for 3D reconstruction approaches. The mathematical model also directly extends to multi-label problems, but the optimization method [7] is restricted to binary problems so far. We will leave this for future work.

6. Appendix: Proof of Theorem 1

In this appendix we will prove Theorem 1. We will use an integral geometric argument to show that squared curvature is related to the brown area in Figure 4.

We first state that the partial area of the circle of curvature $\kappa = \frac{1}{r}$ above a half chord of length $d < r$ in Figure 13(a) is given by

$$\begin{aligned} G(\kappa, d) &:= \frac{1}{2\kappa^2}\gamma - \frac{d}{2\kappa} \cos \gamma \\ &= \frac{1}{2\kappa^2} \arcsin \kappa d - \frac{d}{2} \sqrt{\frac{1}{\kappa^2} - d^2}. \end{aligned} \quad (11)$$

Based on this expression we can now derive the area $A(\kappa, d)$ in Figure 13(b) with respect to a half chord of length d in a circle of radius r

$$\begin{aligned} A(\kappa, d) &:= 2 \left[G(d, \kappa) - \underbrace{2 \left(G\left(\frac{d}{2}, \kappa\right) - \frac{gd}{2}\right)}_{\text{area}B} \right] \\ &= -\frac{1}{\kappa^2} \arcsin \kappa d + 2 \frac{1}{\kappa^2} \arcsin \frac{\kappa d}{2} \\ &\quad - d \sqrt{\frac{1}{\kappa^2} - d^2} + d \sqrt{\frac{1}{\kappa^2} - \frac{d^2}{4}} \end{aligned} \quad (12)$$

Using Taylor approximation w.r.t. d we obtain

$$A(\kappa, d) \approx \frac{d^3 \kappa}{4}. \quad (13)$$

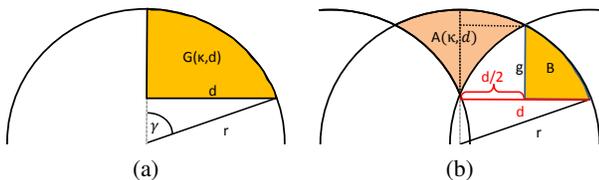


Figure 13. (a) Computing the partial area $G(\kappa, d)$ of a circle above a half chord of length d , (b) Computing the shaded area $A(\kappa, d)$

References

[1] E. Boros, P. Hammer, and X. Sun. Network flows and minimization of quadratic pseudo-boolean functions. *Technical report RRR 17-1991, RUTCOR*, 1991. 5, 7

[2] Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in n-d images. In *Int. Conf. Computer Vision*, 2001. 1

[3] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *Int. Conf. Computer Vision*, 2003. 2, 3

[4] K. Bredies, T. Pock, and B. Wirth. Convex relaxation of a class of vertex penalizing functionals. *Journal of Mathematical Imaging and Vision*, pages 1–25, 2012. 1, 2, 6, 7

[5] A. Bruckstein, A. Netravali, and T. Richardson. Epicongvergence of discrete elastica. *Applicable Analysis*, 79(1-2):137–171, 2001. 2

[6] N. El-Zehiry and L. Grady. Fast global optimization of curvature. In *Conf. Computer Vision and Pattern Recognition*, 2010. 1, 2, 6, 7

[7] L. Gorelick, O. Veksler, Y. Boykov, I. Ben Ayed, and A. DeLong. Local submodular approximations for binary pairwise energies. In *Conf. Computer Vision and Pattern Recognition*, 2014. 2, 5, 6, 7, 8

[8] S. Heber, R. Ranftl, and T. Pock. Approximate envelope minimization for curvature regularity. In *Proceedings of the 12th international conference on Computer Vision - Volume Part III*, European Conf. on Computer Vision, pages 283–292, 2012. 1, 2, 6, 7

[9] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1568–1583, 2006. 5, 6, 7

[10] C. Nieuwenhuis, E. Toeppe, and D. Cremers. A survey and comparison of discrete and continuous multilabel segmentation approaches. *Int. Journal of Computer Vision*, 2013. 2

[11] C. Olsson, J. Ulen, Y. Boykov, and V. Kolmogorov. Partial enumeration and curvature regularization. In *Int. Conf. Computer Vision*, 2013. 2, 6, 7

[12] J. Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In *National Conference on Artificial Intelligence*, pages 133–136, 1982. 5, 6, 7

[13] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer. Optimizing binary mrfs via extended roof duality. In *Conf. Computer Vision and Pattern Recognition*, pages 1–8, 2007. 5, 6, 7

[14] T. Schoenemann, F. Kahl, S. Masnou, and D. Cremers. A linear framework for region-based image segmentation and inpainting involving curvature penalization. *Int. Journal of Computer Vision*, 2012. 1, 2, 6, 7

[15] A. Shekhovtsov, P. Kohli, and C. Rother. Curvature prior for MRF-based segmentation and shape inpaint. In *DAGM*, 2012. 1, 2

[16] P. Strandmark and F. Kahl. Curvature regularization for curves and surfaces in a global optimization framework. In *EMMCVPR*, pages 205–218, 2011. 1, 2, 6, 7