

Introducing Total Curvature for Image Processing

Bastian Goldluecke and Daniel Cremers
TU Munich, Germany

Abstract

We introduce the novel continuous regularizer total curvature (TC) for images $u : \Omega \rightarrow \mathbb{R}$. It is defined as the Menger-Melnikov curvature of the Radon measure $|Du|$, which can be understood as a measure theoretic formulation of curvature mathematically related to mean curvature. The functional is not convex, therefore we define a convex relaxation which yields a close approximation. Similar to the total variation, the relaxation can be written as the support functional of a convex set, which means that there are stable and efficient minimization algorithms available when it is used as a regularizer in image processing problems. Our current implementation can handle general inverse problems, inpainting and segmentation. We demonstrate in experiments and comparisons how the regularizer performs in practice.

1. Introduction

Image processing problems are usually ill-posed due to inherent ambiguity, clutter and real-world noise. Therefore algorithms typically require some form of regularization or prior information to disambiguate the input data. While psychophysical studies on contour completion suggest that curvature plays a central role in human perception [12], length regularity is currently the dominant paradigm for denoising, deblurring, segmentation and geometric reconstruction. This is largely due to the fact that length regularity is mathematically well understood, and can be formulated with convex regularizers like the total variation for which a variety of fast and stable minimization algorithms exists.

Curvature, on the other hand, is a non-convex functional, and there is currently no known way to globally minimize it. The classical continuous mean curvature functional is non-convex and contains second-order derivatives, and thus its minimization is highly challenging. To date, most continuous approaches therefore employ local optimization methods only.

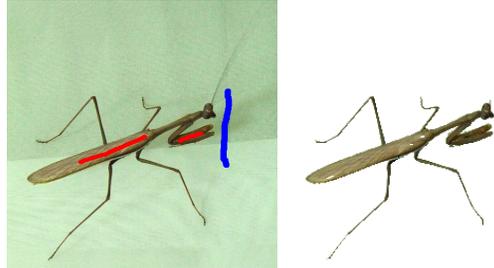


Figure 1: Total curvature regularity is ideal for segmentation of elongated structures. Image resolution 256×256 , smoothness parameter $\lambda = 0.002$.

Contribution.

In this paper, we introduce the concept of *total curvature* of a function, which is based on the Menger-Melnikov curvature of a measure. In contrast to the classical mean curvature functional, it is defined for discontinuous functions in the whole Hilbert space $\mathcal{L}^2(\Omega)$, and requires only first-order derivatives. Geometrically for characteristic (binary) functions, the total curvature gives a measure of the curvature of the jump set.

We introduce a one-parameter family of convex regularizers for total curvature, which has total variation as a limit case and is a very good approximation to total curvature for characteristic functions if the parameter is chosen to fit the discretization. We give the complete mathematical analysis necessary for minimization of total curvature regularized functionals. To this end we formulate the regularizers as support functionals of convex sets of dual vector fields – analogous to respective definitions of the total variation. As a consequence, we are able to employ well-known efficient and stable numerical schemes to solve image processing problems with total curvature regularity.

We provide detailed descriptions of the implementation and show a variety of experimental comparisons of total curvature with other continuous and discrete regularizers. These experiments demonstrate that total curvature is a suitable regularizer for a variety of image processing application such as image segmentation and inpainting.

Relation to previous work.

The proposed continuous approach to curvature regularity is most similar to the continuous length regularity approaches which have been analyzed extensively in previous work on region-based segmentation [6, 16, 18, 10]. In these works researchers minimize energies based on region integrals with the regularizer penalty proportional to the *length* of the region boundaries. Recent formulations are based on total variation [17, 21], allowing globally optimal solutions of the binary partitioning problem via relaxation and subsequent thresholding [17]. Total variation is also commonplace in more general image processing problems [8]. Due to its convexity and lower semi-continuity it lends itself to powerful minimization algorithms, which have been researched extensively in the last years [3, 4, 7, 8]. Since the proposed regularizer total curvature has similar mathematical properties, we can easily substitute it in the above methods and algorithms.

Continuous approaches to curvature regularity predominantly employ local evolution methods [9, 11, 18, 25], thereby being confined to sub-optimal local solutions. A notable exception is the approach [14] to inpainting, where the \mathcal{L}^1 -norm of the curvature is minimized globally using dynamic programming, provided there is no data term. Our approach differs from these works in that the original non-convex problem is relaxed to a closely related convex one, which can be solved in a globally optimal way. Discrete contour-based approaches to curvature regularity have been formulated using shortest path approaches [1] or ratio cycle formulations [23] on a graph representing the product space of image pixels and tangent angles [20].

A recent discrete region-based segmentation and inpainting method [24] is based on the concepts of cell complexes and surface continuation constraints. Their method is somewhat related to ours, since they also perform a relaxation of a binary non-convex problem to a convex one. Subsequent thresholding yields a solution within a known bound of the global optimum, typically below 0.1% in their experiments. However, discrete graph-based approaches often suffer from metrication errors [13], which make the results depend on the underlying grid. The continuous framework we present completely avoids these type of artifacts, as we will see in later comparisons.

2. Total Curvature

The Menger-Melnikov curvature [15] of degree $p > 0$ introduces the concept of the curvature of a measure.

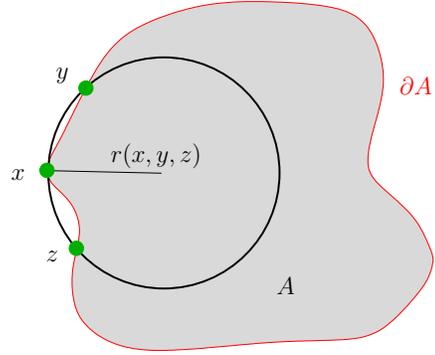


Figure 2: The Menger curvature weight is given by reciprocal of radius of circle through three points. When computing total curvature for the characteristic function of a set A , only triples (x, y, z) of points on the boundary of A contribute to the integral.

For a measure μ on a set $\Omega \subset \mathbb{R}^n$, it is defined as

$$c^p(\mu) := \int_{\Omega \times \Omega \times \Omega} c(x, y, z)^{2p} d\mu(x) d\mu(y) d\mu(z), \quad (1)$$

where $c(x, y, z) = r(x, y, z)^{-1}$, and $r(x, y, z)$ is the radius of the (unique) circle passing through x, y, z , see figure 2. In the case that the points are collinear, define $c(x, y, z) = 0$. Also, to avoid the issue of unbounded c , we set by common convention $c(x, y, z)$ to zero if x, y, z are closer together than a certain ϵ . Since after the eventual discretization this concern does not exist anymore, this detail is ignored in the following to not overly complicate the notation. Note that the definition makes sense in arbitrary dimension, since three points in n -dimensional space always lie in a two-dimensional plane. Also, only first-order derivatives are involved in the definition.

There are close ties of the Menger-Melnikov curvature to the mean curvature, which is not surprising since the mean curvature at each point of a curve in the plane is geometrically the inverse of the radius of an osculating (“kissing”) circle. Straight lines have zero curvature, while one can show [19] that if μ is the one-dimensional Hausdorff measure on a circle S_r with radius r , then $c^2(\mu) = 1/r$, in particular $c^2(1_{S_r}) = 1/r$. Figure 3 illustrates the similarity of the Menger-Melnikov curvature of degree p to the p th power of the mean curvature.

We now define the novel concept of *total curvature of a function* $u \in \mathcal{L}^2(\Omega)$ as the Menger-Melnikov curvature of the Radon measure $|Du|$.

Definition 2.1. Let $u \in \mathcal{L}^2(\Omega)$. Then its total curvature $TC_p(u)$ of degree p is defined as

$$TC_p(u) := c^p(|Du|). \quad (2)$$

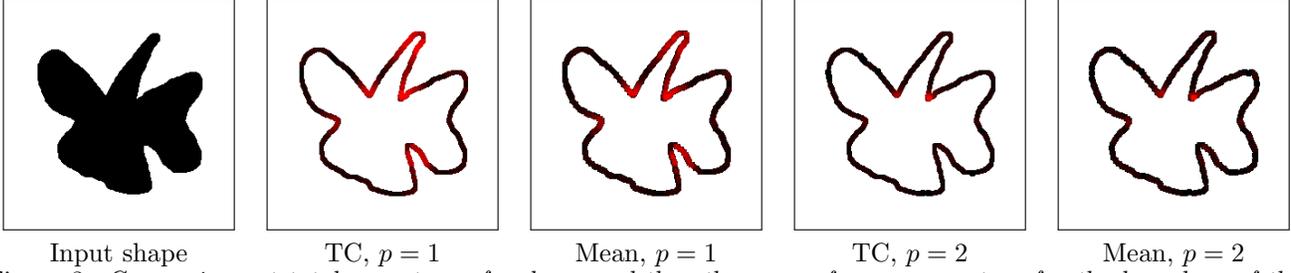


Figure 3: Comparison of total curvature of order p and the p th power of mean curvature for the boundary of the shape $A \subset \Omega$ on the left. Brighter shades of red mean higher curvature. We see that in each point of the boundary curve, both measures of curvature behave very similar. Note that the total curvature of order p in a point x is computed by keeping x fixed and integrating over y and z in (1), using $\mu = |D1_A|$. For illustration purposes, the boundary was dilated and values scaled to the same range.

The geometric insight is that for a characteristic function of a sufficiently regular (i.e. piecewise smooth boundary) set $A \subset \Omega$, the measure $|Du|$ is the $(n-1)$ -dimensional Hausdorff measure restricted to its boundary. Thus, in the integral in (1) only the points $x, y, z \in \partial A$ contribute to the curvature, i.e. the total curvature of a characteristic function 1_A is a measure of curvature for the boundary ∂A , see figure 2. This is analogous to the total variation, which equals the length of this boundary, $\text{TV}(1_A) = \mathcal{H}^1(\partial A)$. In the following, we suppress the dependency of the total curvature functional TC on p to not clutter notation. In experiments, we always use $p = 1$ if not otherwise noted.

3. Relaxation

Unfortunately, the functional TC is not convex on $\mathcal{L}^2(\Omega)$ and very hard to minimize if used as a regularizer. We therefore propose a one-parameter family TC_α of functionals which are natural relaxations of the *total curvature* and have in a sense both the total variation as well as the total curvature as limit cases. They have a number of desirable properties while preserving important characteristics of $\text{TC}(u)$. In this section, we only give the (primal) definition for differentiable u , which we extend to $\mathcal{L}^2(\Omega)$ via duality in section 4.

Definition 3.1. Let $\alpha > 0$ and $u \in \mathcal{C}^1(\Omega)$. Then the relaxation $\text{TC}_\alpha(u)$ for the total curvature of u is defined as

$$\text{TC}_\alpha(u) := \int_{\Omega \times \Omega \times \Omega} c(x, y, z)^{2p} \max(0, \dots \dots |\nabla u(x)|_2 + |\nabla u(y)|_2 + |\nabla u(z)|_2 - \alpha) \, d(x, y, z). \quad (3)$$

The relaxation TC_α is convex on $\mathcal{C}^1(\Omega)$, since concatenation with $\max(0, \cdot)$ preserves convexity. We exploit this in sections 4 and 5 in order to construct minimization algorithms for image processing problems.

Before we turn to that, we will discuss how TC_α is related to the curvature functional TC and the total variation.

Interpretation of the relaxation.

The relationship to the original total curvature functional for characteristic functions is the motivating factor in the definition of the relaxation (3). To understand it, we take a look at an indicator function 1_A in the discrete setting, when the gradient operator is given by standard forward differences. In this case, the value of $|\nabla 1_A|_2$ can be either 0, 1 or $\sqrt{2}$. Thus, if we set $\alpha = 2\sqrt{2}$, then this means that for indicator functions, the relaxation (3) is a sensible approximation to the exact curvature in the sense that in both cases, the integral gives a contribution if and only if each of x, y, z is on the boundary of the set A :

$$\begin{aligned} & |\nabla 1_A(x)|_2 + |\nabla 1_A(y)|_2 + |\nabla 1_A(z)|_2 > \alpha \\ \Leftrightarrow & |\nabla 1_A(x)|_2 \neq 0 \wedge |\nabla 1_A(y)|_2 \neq 0 \wedge |\nabla 1_A(z)|_2 \neq 0. \end{aligned} \quad (4)$$

For this reason, the case $\alpha = 2\sqrt{2}$ is the most relevant for us, and we use it throughout the experiments.

In the case $\alpha = 0$, the relaxation is equal to a weighted total variation with weight

$$\begin{aligned} g(w) = & \int_{\Omega^2} c^{2p}(w, y, z) \, d(y, z) + c^{2p}(x, w, z) \, d(x, z) + \dots \\ & \dots + c^{2p}(x, y, w) \, d(x, y), \end{aligned} \quad (5)$$

which reduces to a constant weight for $\Omega = \mathbb{R}^n$. Thus, cases of $0 < \alpha < 2\sqrt{2}$ lead to a regularization which is qualitatively somewhere between curvature and length regularity.

Experimental validation.

In order to verify how this theoretical interpretation works out in practice, we have experimentally compared the integrands for the exact total curvature, the relaxation for $\alpha = 2\sqrt{2}$ and $\alpha = \sqrt{2}$, as well as the total variation for reference. We computed the integrand

	$TC_{2\sqrt{2}}$	$TC_{\sqrt{2}}$	TC_0
TC	0.98	0.94	0.77
TV	0.70	0.87	0.89

Figure 4: Correlation coefficients for pointwise integrands of different regularizers. We clearly see that $TC_{2\sqrt{2}}$ is a very good approximation for TC. Data sets were computed from binary images obtained from a large image database.

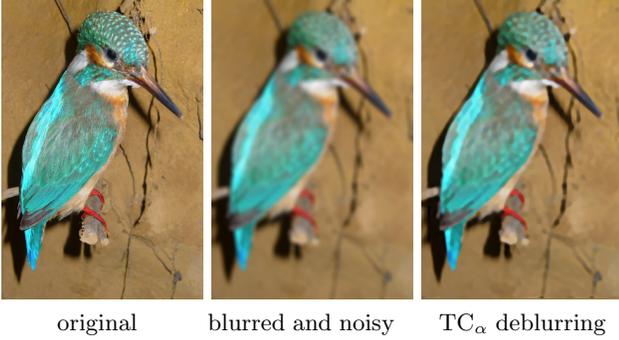


Figure 5: Deblurring with total curvature regularization. The blur kernel is Gaussian with a standard deviation of five pixels. Image resolution is 320×480 , deblurring was done for each color channel separately.

pointwise for each pixel in one hundred 256×256 binary images created by thresholding random images from the internet. Figure 4 shows the correlation coefficients for the resulting columns of data. We can see that indeed as expected, the relaxation $TC_{2\sqrt{2}}$ is a very close approximation of the exact total curvature TC (up to a constant linear transformation, which is unimportant if it is used as a regularizer), while the relaxation TC_0 is most closely correlated to TV.

4. Convex Analysis

We write $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ for the inner product on a vector space \mathcal{V} , and omit mentioning \mathcal{V} if it can be easily deduced from context. Usually, it will denote inner products on \mathcal{L}^2 Hilbert spaces. The following theorem computes the dual formulation of TC_α in terms of a support functional of a convex set, which is a foundation for efficient minimization algorithms.

Theorem 4.1. *Let $u \in C^1(\Omega)$. Then*

$$TC_\alpha(u) = \sup_{(p,q) \in C} \left\{ \langle u, K^*p \rangle_{\mathcal{L}^2(\Omega)} - \langle \alpha, q \rangle_{\mathcal{L}^2(\Omega^3)} \right\}, \quad (6)$$

with the convex set

$$C := \left\{ (p_x, p_y, p_z, q) \in C_c^1(\Omega^3, \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}) : \begin{aligned} &|p_i|_2 \leq q \text{ and } 0 \leq q \leq c^{2p}(\text{pointwise}) \end{aligned} \right\} \quad (7)$$

and the linear operator

$$K^*(p_x, p_y, p_z) := - \int_{\Omega^2} \operatorname{div}(p_x(\cdot, y, z)) d(y, z) - \int_{\Omega^2} \operatorname{div}(p_y(x, \cdot, z)) d(x, z) - \int_{\Omega^2} \operatorname{div}(p_z(x, y, \cdot)) d(x, y). \quad (8)$$

Proof. See additional material. \square

Note that K^* is the adjoint of the operator K given by

$$Ku(x, y, z) = [\nabla u(x) \quad \nabla u(y) \quad \nabla u(z)]^T. \quad (9)$$

As an immediate corollary, the above theorem extends definition 3.1 to all of $\mathcal{L}^2(\Omega)$ by establishing TC_α as the support functional of the convex set $[K^*, 1]C$, the notation meaning that the first component of the elements in C is transformed by K^* , while the second is unchanged.

Corollary 4.2. *Equation (6) extends definition 3.1 to all of $\mathcal{L}^2(\Omega)$. For all $u \in \mathcal{L}^2(\Omega)$,*

$$TC_\alpha(u) = \sigma_{[K^*, 1]C}(u, -\alpha), \quad (10)$$

where σ denotes the support functional. Thus, TC_α is a closed (or lower semi-continuous) functional on $\mathcal{L}^2(\Omega)$ in addition to being convex.

When total curvature is used as a regularizer in computer vision and image processing, then of special interest are problems of the form

$$\min_{u \in \mathcal{L}^2(\Omega)} \{TC_\alpha(u) + G(u)\}, \quad (11)$$

where G is a cost function on $\mathcal{L}^2(\Omega)$. In the next section, several archetypical image processing problems are given as examples. It is therefore a very interesting property of total curvature regularity that we are able to guarantee the existence of solutions to such problems if G has certain properties.

Theorem 4.3. *Let G be either*

- i. *convex, closed and coercive, or*
- ii. *weakly lower semi-continuous and coercive.*

Then problem (11) has a minimizer with bounded total curvature.

Proof. Property (i) implies (ii), see Theorem 3.3.3 in [2]. We have already shown that TC_α is closed on $\mathcal{L}^2(\Omega)$. Thus, in both cases the full functional is weakly lower semi-continuous and coercive. Since $\mathcal{L}^2(\Omega)$ is a reflexive space, this implies the existence of a minimizer according to the Weierstraß minimization theorem, 3.2.5 in [2]. Obviously, the minimizer must have bounded total curvature. \square

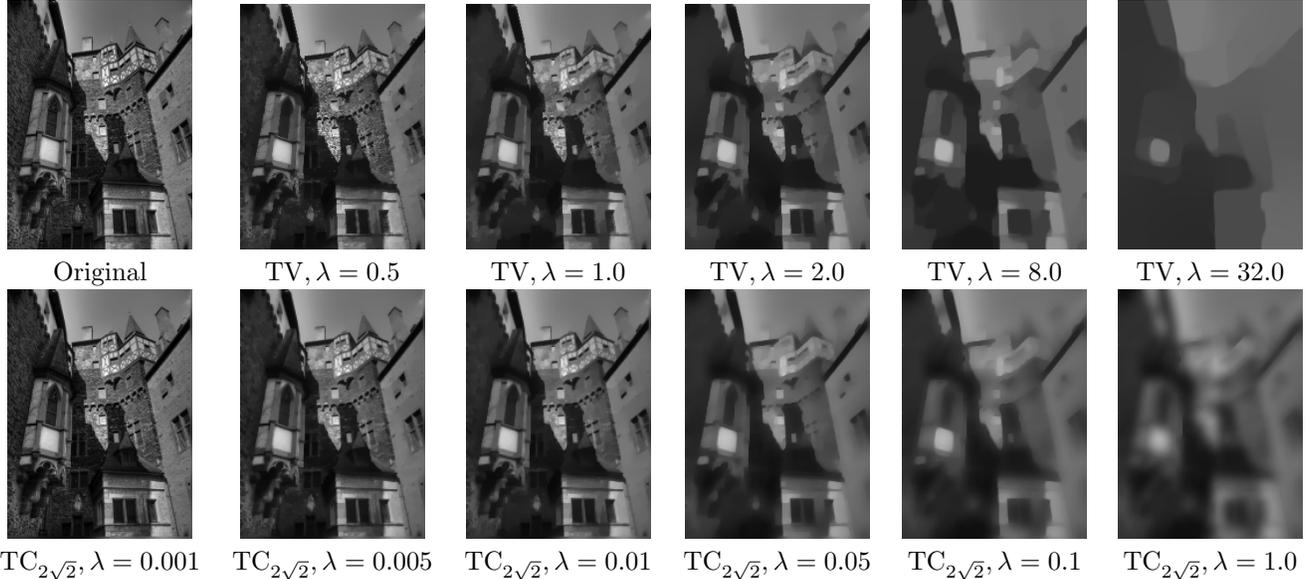


Figure 6: *Proximations of the relaxation of total curvature and total variation. The weights λ of both regularizers are unrelated. We can see that TC_α also allows discontinuous solutions and has similar edge-preserving properties than TV. However, at very high smoothing values, TC_α performs significantly different. In particular, it shows a saturation in regularization: the max function in the integrand has the effect of ignoring low-frequency discontinuities.*

A particularly important example for problem (11) is the TC_α - \mathcal{L}^2 problem (the analogous problem for the total variation is the celebrated Rudin-Osher-Fatemi (ROF) model [22]), which is to solve for given $f \in \mathcal{L}^2(\Omega)$ and $\lambda > 0$

$$\min_{u \in \mathcal{L}^2(\Omega)} \left\{ TC_\alpha(u) + \frac{1}{2\lambda} \|u - f\|_2^2 \right\}. \quad (12)$$

While it can also serve as a model for denoising, it is of large general interest because its minimizer is equal to the *proximation* of λTC_α , which can be interpreted as a subgradient descent step for TC_α with step size λ [3]. For this reason, it is at the heart of several algorithms designed to solve more general problems, in particular the fast iterative shrinkage and thresholding (FISTA) which we employ in our work [4]. Theorem 4.3 assures that the proximation for TC_α is well-defined, since the norm is obviously coercive on $\mathcal{L}^2(\Omega)$.

5. Image processing with TC_α regularity

In this section, we describe typical choices for the data term G in image processing applications, and detail algorithms how to solve the resulting problem with total curvature regularity.

Denoising.

The estimate for the original function u which maximizes the posterior probability under the assumption that we observe an image f distorted with pointwise

Gaussian noise of standard deviation λ is the minimizer of the proximation problem (12). We solve this problem with the Bermúdez-Moreno algorithm [5], which was recently reintroduced for image processing problems in [3], and earlier rediscovered in [7] as a solver for the ROF model with total variation regularization [22]. Specialized to problem (12), the algorithm is given in figure 10. It has a guaranteed convergence for a step size $\sigma < 2/(\lambda \|K^*\|^2)$. In figure 6, we demonstrate the behaviour of the proximation when the smoothing parameter is varied, and compare total curvature to total variation.

Inverse problems.

The general inverse problem is to find

$$\min_{u \in \mathcal{L}^2(\Omega)} \left\{ TC_\alpha(u) + \frac{1}{2\lambda} \|Au - f\|_2^2 \right\}, \quad (13)$$

where A is an operator on $\mathcal{L}^2(\Omega)$. It can be interpreted as a maximum a-posteriori estimate for the model $Au = f$, where the observation f is distorted by Gaussian noise with standard deviation λ . We solve it with the FISTA algorithm [4], which is suitable for general problems of the form (11) in the case of a data term G with L -Lipschitz continuous gradient. FISTA alternates a gradient descent in G with a subgradient descent in TC_α , which is given by the proximation. In both cases, the possible step size is $1/L$. An acceleration step guarantees a convergence rate of $O(1/N^2)$,

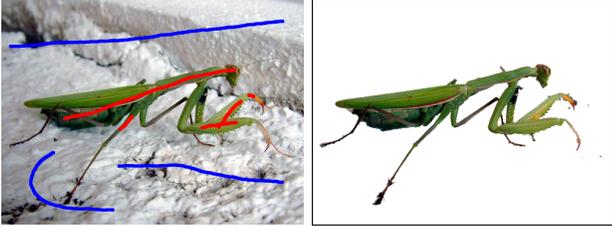


Figure 7: Segmentation with manually marked seed regions. Elongated structures are well preserved despite a relatively large amount of regularization, $\lambda = 0.005$. Image resolution is 512×384 .

where N is the number of iterations. The resulting algorithm for total curvature regularity is detailed in figure 10, where we use the Bermúdez-Moreno algorithm for the subgradient descent in TC_α . In the general inverse problem, $L = \|A^*A\|/\lambda$. A special case is deblurring, where $Au = b * u$ with a convolution kernel b , and $\|A\| = \|A^*\| = \|b\|_1$. An example for deblurring with curvature regularity can be found in figure 5. Note that in deblurring, the amount of regularization is usually very low, so that the choice of regularizer does not make a very noticeable impact on the result.

Inpainting.

A simple inpainting model is

$$\min_{u \in \mathcal{L}^2(\Omega)} \left\{ \text{TC}_\alpha(u) + \int_{\Omega \setminus M} \frac{\rho}{2} \|u - f\|_2^2 \right\}. \quad (14)$$

where $\rho > 0$ is constant and $M \subset \Omega$ is a mask denoting a damaged region in the image which is to be inpainted. If ρ is chosen large, the regularizer only reconstructs the damaged regions and leaves the rest of the image mostly unchanged. Since the data term has a Lipschitz continuous gradient with $L = \rho$, we can also solve it with the FISTA scheme, see figure 10. Inpainting with total curvature regularity and a comparison to total variation regularization can be found in figure 9. As reported before in [24], we see that curvature regularity is particularly well suited to inpainting and the result is visually significantly better.

Segmentation.

For segmentation with curvature regularity, we want to solve the binary problem

$$\min_{u \in \mathcal{L}^2(\Omega, \{0,1\})} \left\{ \text{TC}_\alpha(u) + \int_{\Omega} au \, dx \right\}. \quad (15)$$

The weight function $a \in \mathcal{L}^2(\Omega)$ is a cost function which denotes how expensive it is to assign a pixel to region 1. Note that it is negative if there is a preference for this assignment. A typical choice is $a = \log(P_b)/\log(P_f)$ if P_f, P_b are the probabilities for a pixel to belong to

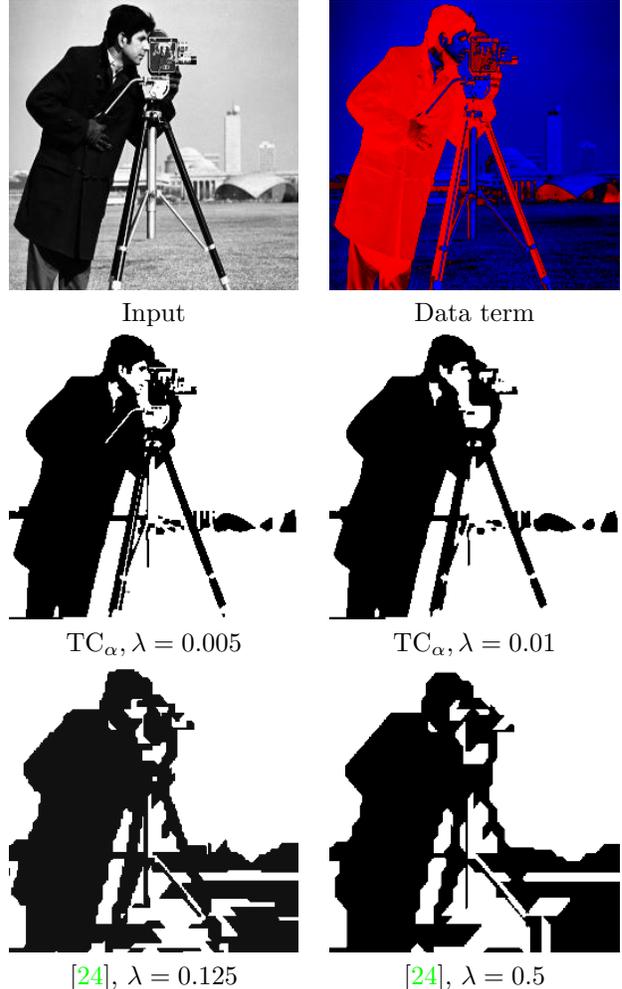


Figure 8: Comparison of segmentation results with Mumford-Shah data term on the cameraman image, with foreground intensity 1 and background intensity 0, respectively. Below are results taken from Schoenemann et al. [24] (with other average intensity levels). Discretization artifacts caused by metrication errors are clearly visible, which our continuous results do not suffer from.

foreground or background, respectively. The data term has again Lipschitz continuous derivative with $L > 0$ arbitrary, so we employ the FISTA scheme to solve the model, with the binary function space relaxed to $\mathcal{L}^2(\Omega, [0, 1])$, see figure 10.

In figures 1 and 7 we can see results with some parts of the image manually marked to build a histogram of the foreground and background color distribution. Elongated structures are segmented well although a relatively high amount of regularization was necessary. In figure 8 we compare segmentation results with different amounts of regularity to the results of Schoenemann et al. [24]. We could not reproduce their model

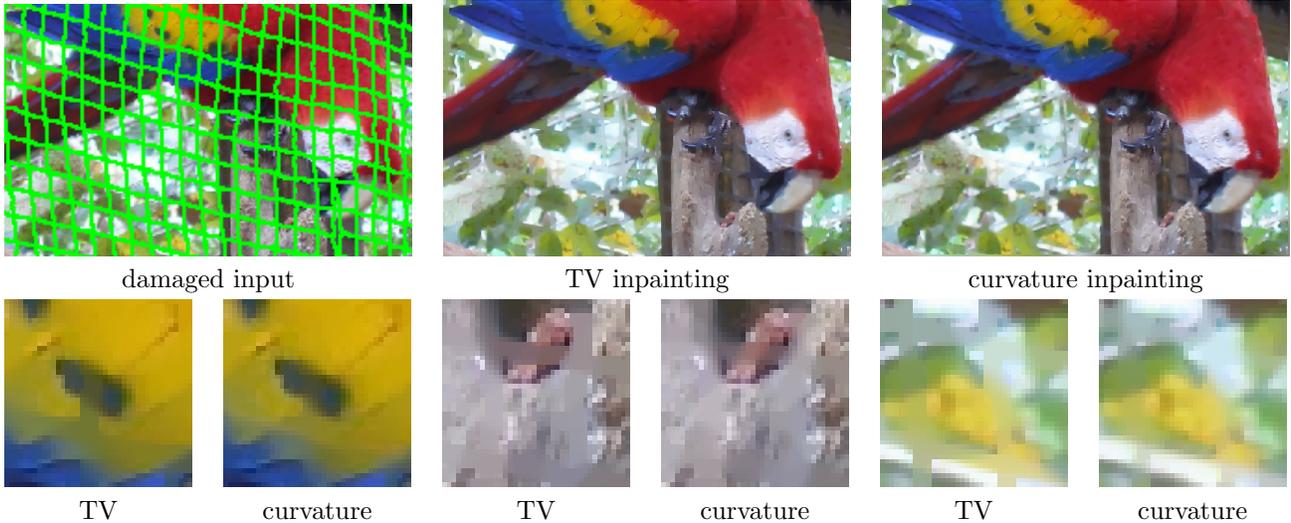


Figure 9: Uncaging a bird: *the mesh of the cage has been marked as a damaged regions and inpainted, with total variation regularity in the center and total curvature regularity on the right, respectively. The bottom row has closeups of the results for better comparison. While both regularizers are obviously only suitable for cartoon inpainting and are very similar at first glance, total curvature yields a visually more pleasing result, leaving less artifacts from the regular cage structure.*

parameters which is why the segmentations look quite different, however one can clearly see metrication errors in their results which our continuous method does not suffer from. Note that the results for TC_α are *not thresholded*, which shows that with a linear data term one gets almost binary solutions, so we are very close to the global optimum of the binary energy (15).

Implementation notes.

If we implement the full operator K^* in (8), then total curvature becomes extremely costly to minimize since computation time scales with the cube of the number of pixels. Instead, we note that the curvature weight c^{2p} decreases quickly if x, y, z are further apart, since the radius of the circle must be at least half as large as the smallest distance. For this reason, we restrict the domains of y and z to a window of $N \times N$ pixels around x to make the minimization feasible. Typically, $N = 5$ in our experiments. Note that the memory and computation time requirements are roughly N^4 times as high as for TV minimization. Since a correct and efficient implementation of the scheme is quite complex, we will offer sample CUDA code on our webpage ¹ after publication of the paper. Computation time for the examples in the paper was between 40 minutes and 3 hours on an nVidia Fermi class GPU.

6. Conclusion

We have introduced the novel regularizer *total curvature* and demonstrated its feasibility for image pro-

cessing applications. It is a measure of curvature based on the Menger-Melnikov curvature of a measure, and closely related to the mean curvature in that it behaves like the mean curvature integral over the boundary for the characteristic function of a set. We introduced a convex approximation which allows for efficient and stable minimization algorithms which compute the global minimizer of general inverse and inpainting problems with approximative curvature regularity. For segmentation, a relaxation yields solutions which are almost binary, i.e. very near the global optimum.

References

- [1] A. Amini, T. Weymouth, and R. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(9):855–867, 1990. 2
- [2] H. Attouch, G. Buttazzo, and G. Michaille. *Variational Analysis in Sobolev and BV Spaces*. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2006. 4
- [3] J.-F. Aujol. Some first-order algorithms for total variation based image restoration. *Journal of Mathematical Imaging and Vision*, 34(3):307–327, 2009. 2, 5
- [4] A. Beck and M. Teboulle. Fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2:183–202, 2009. 2, 5
- [5] A. Bermúdez and C. Moreno. Duality methods for solving variational inequalities. *Comp. and Maths. with Appls.*, 7:43–58, 1981. 5

¹<http://cvpr.in.tum.de>

<p>Initialize</p> $u_0 = 0, \mathbf{p}_0 = 0, q_0 = 0$ $\sigma < \frac{2}{\lambda \ K^*\ ^2} = \frac{1}{4 \cdot 3^2 \lambda N^4}$ <p>Iterate</p> $\tilde{\mathbf{p}}_{n+1} = \mathbf{p}_n + \sigma K u_n$ $\tilde{q}_{n+1} = q_n - \sigma \alpha$ $(\mathbf{p}_{n+1}, q_{n+1}) = \Pi_K(\tilde{\mathbf{p}}_{n+1}, \tilde{q}_{n+1})$ $u_{n+1} = f + \lambda \operatorname{div}(\mathbf{p}_{n+1})$ <p>Bermúdez-Moreno (BM) for proximation of TC_α</p>	<p>Initialize</p> $u_0 = 0, \bar{u}_0 = u_0, \mathbf{p}_0 = 0, q_0 = 0$ $L = L(\nabla G), t_0 = 1$ <p>Iterate</p> $g_n = \bar{u}_n - \frac{1}{L} \nabla G(u_n)$ <p>repeat k BM iterations with $\lambda = 1/L, f = g_n$</p> $t_{n+1} = (1 + \sqrt{1 + 4t_n^2})/2$ $\bar{u}_{n+1} = u_n + \frac{t_n - 1}{t_{n+1}}(u_{n+1} - u_n)$ <p>Beck-Teboulle (FISTA) for general data terms</p>	<p>General inverse problem</p> $\nabla G(u_n) = \frac{1}{\lambda} A^*(A\bar{u}_n - f)$ <p>Deblurring</p> $\nabla G(u_n) = \frac{1}{\lambda} \bar{b} * (b * u_n - f)$ <p>Inpainting</p> $\nabla G(u_n) = 1_{\Omega \setminus M} \rho(u_n - f)$ <p>Segmentation</p> $\nabla G(u_n) = a$ <p>Different gradients for FISTA</p>
--	--	--

Figure 10: Algorithms to solve image processing problems with TC_α regularization. The Bermúdez-Moreno algorithm is only suitable for the proximation, while FISTA can handle more general data terms, for example those detailed in the right box.

- [6] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987. 2
- [7] A. Chambolle. Total variation minimization and a class of binary MRF models. In *EMMCVPR*, pages 136–152, 2005. 2, 5
- [8] A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock. An introduction to total variation for image analysis. In *Theoretical Found. and Numerical Methods for Sparse Recovery*. De Gruyter, 2010. 2
- [9] T. Chan, S. Kang, and J. Shen. Eulers elastica and curvature based inpainting. *SIAM Journal of Applied Mathematics*, 63:564–592, 2002. 2
- [10] T. Chan and L. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001. 2
- [11] S. Esedoglu and R. March. Segmentation with depth but without detecting junctions. *Jour. Math. Imaging and Vision*, 18:7–15, 2003. 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 *ECCV*, 2008. 2
- [14] S. Masnou and J. Morel. Level-lines based disocclusion. In *Proceedings of the IEEE International Conference on Image Processing*, volume 3, pages 259–263, 1998. 2
- [15] M. Melnikov. Analytic capacity: a discrete approach and the curvature of measure. *Mat. Sb.*, 186(6):57–76, 1995. 2
- [16] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989. 2
- [17] M. Nikolova, S. Esedoglu, and T. Chan. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal of Applied Mathematics*, 66(5):1632–1648, 2006. 2
- [18] M. Nitzberg, D. Mumford, and T. Shiota. Filtering, segmentation and depth. In *LNCS*, volume 662. Springer, 1993. 2
- [19] H. Pajot. *Analytic Capacity, Rectifiability, Menger Curvature and the Cauchy Integral*. Number 1799 in Lecture notes in mathematics. Springer, 2002. 2
- [20] P. Parent and S. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(8):823–839, 1989. 2
- [21] T. Pock, A. Chambolle, H. Bischof, and D. Cremers. A convex relaxation approach for computing minimal partitions. In *CVPR*, 2009. 2
- [22] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, 1992. 5
- [23] T. Schoenemann and D. Cremers. Introducing curvature into globally optimal image segmentation: Minimum ratio cycles on product graphs. In *ICCV*, 2007. 2
- [24] T. Schoenemann, F. Kahl, and D. Cremers. Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In *ICCV*, pages 17–23, 2009. 2, 6
- [25] D. Tschumperlé. Fast anisotropic smoothing of multivalued images using curvature-preserving pdes. *International Journal of Computer Vision*, 68(1):65–82, 2006. 2

Introducing Total Curvature for Image Processing

Additional Material

Bastian Goldluecke and Daniel Cremers
TU Munich, Germany

1. Proof of theorem 4.1

Recall that the relaxation for the total curvature of u is for $\alpha > 0$ and $u \in \mathcal{C}^1(\Omega)$ defined as

$$\text{TC}_\alpha(u) := \int_{\Omega \times \Omega \times \Omega} c(x, y, z)^{2p} \max(0, |\nabla u(x)|_2 + |\nabla u(y)|_2 + |\nabla u(z)|_2 - \alpha) \, d(x, y, z). \quad (1)$$

The function $\max(0, \cdot)$ is convex and can be written as a support functional,

$$\max(0, v) = \sup_{w \in [0,1]} \{vw\}. \quad (2)$$

Using a standard density argument, we can thus rewrite (1) as

$$\begin{aligned} \text{TC}_\alpha(u) &= \int_{\Omega \times \Omega \times \Omega} c(x, y, z)^{2p} \max(0, |\nabla u(x)|_2 + |\nabla u(y)|_2 + |\nabla u(z)|_2 - \alpha) \, d(x, y, z) \\ &= \sup_{\tilde{q} \in \mathcal{C}_c^1(\Omega^3, [0,1])} \int_{\Omega \times \Omega \times \Omega} c(x, y, z)^{2p} \tilde{q}(x, y, z) (|\nabla u(x)|_2 + |\nabla u(y)|_2 + |\nabla u(z)|_2 - \alpha) \, d(x, y, z) \\ &= \sup_{q \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}), 0 \leq q \leq c^{2p}} \int_{\Omega \times \Omega \times \Omega} q(x, y, z) (|\nabla u(x)|_2 + |\nabla u(y)|_2 + |\nabla u(z)|_2 - \alpha) \, d(x, y, z). \end{aligned} \quad (3)$$

We are closer to what we would like to have, but still have to deal with the norms. For this, we write the norm also as a support functional,

$$|\mathbf{v}|_2 = \sup_{|\mathbf{w}|_2 \leq 1} \langle \mathbf{v}, \mathbf{w} \rangle.$$

We plug this into (3), and at first only look at the first term of the integral below the supremum:

$$\begin{aligned} &\int_{\Omega \times \Omega \times \Omega} q(x, y, z) |\nabla u(x)|_2 \, d(x, y, z) \\ &= \sup_{\tilde{\mathbf{p}}_x \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n), |\tilde{\mathbf{p}}_x|_2 \leq 1} \int_{\Omega \times \Omega \times \Omega} q(x, y, z) \tilde{\mathbf{p}}_x(x, y, z) \nabla u(x) \, d(x, y, z). \end{aligned} \quad (4)$$

We now use Gauss' integral theorem with respect to the integration variable x , rearrange and do a variable substitution:

$$\begin{aligned} &\sup_{\tilde{\mathbf{p}}_x \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n), |\tilde{\mathbf{p}}_x|_2 \leq 1} \int_{\Omega \times \Omega} \left[\int_{\Omega} q(x, y, z) \tilde{\mathbf{p}}_x(x, y, z) \nabla u(x) \, dx \right] \, d(y, z) \\ &= \sup_{\tilde{\mathbf{p}}_x \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n), |\tilde{\mathbf{p}}_x|_2 \leq 1} \int_{\Omega \times \Omega} \left[\int_{\Omega} \text{div}(q(\cdot, y, z) \tilde{\mathbf{p}}_x(\cdot, y, z))(x) u(x) \, dx \right] \, d(y, z) \\ &= \sup_{\mathbf{p}_x \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n), |\mathbf{p}_x|_2 \leq q} \int_{\Omega} u(x) \left[\int_{\Omega \times \Omega} \text{div}(\mathbf{p}_x(\cdot, y, z)) \, d(y, z) \right] (x) \, dx \end{aligned} \quad (5)$$

This is by definition of K^* nothing else than

$$\sup_{\mathbf{p}_x \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n), |\mathbf{p}_x|_2 \leq q} \int_{\Omega} u(x) K^* \mathbf{p}_x(x) dx = \langle u, K^* \mathbf{p}_x \rangle_{\mathcal{L}^2(\Omega)}. \quad (6)$$

Adapting the derivation slightly for the other terms of the sum, we arrive at the claim of the theorem.

2. Proof of corollary 4.2

The corollary is immediate, we just clarify notation and definitions. Recall that theorem 4.1 established that for $u \in \mathcal{C}^1(\Omega)$,

$$\text{TC}_{\alpha}(u) = \sup_{(\mathbf{p}, q) \in C} \left\{ \langle u, K^* \mathbf{p} \rangle_{\mathcal{L}^2(\Omega)} - \langle \alpha, q \rangle_{\mathcal{L}^2(\Omega^3)} \right\}, \quad (7)$$

with the convex set

$$C := \left\{ (p_x, p_y, p_z, q) \in \mathcal{C}_c^1(\Omega^3, \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}) : \right. \\ \left. |p_i|_2 \leq q \text{ and } 0 \leq q \leq c^{2p}(\text{pointwise}) \right\} \quad (8)$$

and the linear operator

$$K^*(p_x, p_y, p_z) := - \int_{\Omega^2} \text{div}(p_x(\cdot, y, z)) d(y, z) \\ - \int_{\Omega^2} \text{div}(p_y(x, \cdot, z)) d(x, z) \\ - \int_{\Omega^2} \text{div}(p_z(x, y, \cdot)) d(x, y). \quad (9)$$

We immediately see that

$$\text{TC}_{\alpha}(u) = \sup_{(\mathbf{p}, q) \in C} \left\{ \langle u, K^* \mathbf{p} \rangle_{\mathcal{L}^2(\Omega)} - \langle \alpha, q \rangle_{\mathcal{L}^2(\Omega^3)} \right\}, \\ = \sup_{(\tilde{\mathbf{p}}, q) \in [K^*, 1]C} \left\{ \langle u, \tilde{\mathbf{p}} \rangle_{\mathcal{L}^2(\Omega)} - \langle \alpha, q \rangle_{\mathcal{L}^2(\Omega^3)} \right\}, \\ = \sup_{(\tilde{\mathbf{p}}, q) \in [K^*, 1]C} \left\{ \langle [u, \alpha], [\tilde{\mathbf{p}}, q] \rangle_{\mathcal{L}^2(\Omega) \times \mathcal{L}^2(\Omega^3)} \right\} \quad (10)$$

This is by definition the support functional σ of the convex set $[K^*, 1]C$ evaluated at $[u, \alpha]$.