Approximate Envelope Minimization for Curvature Regularity

Stefan Heber, Rene Ranftl, and Thomas Pock

Institute for Computer Graphics and Vision, Graz University of Technology
Inffeldgasse 16, A-8010 Graz, Austria
{stefan.heber,ranftl,pock}@icg.tugraz.at
http://www.icg.tu-graz.ac.at

Abstract. We propose a method for minimizing a non-convex function, which can be split up into a sum of simple functions. The key idea of the method is the approximation of the convex envelopes of the simple functions, which leads to a convex approximation of the original function. A solution is obtained by minimizing this convex approximation. Cost functions, which fulfill such a splitting property are ubiquitous in computer vision, therefore we explain the method based on such a problem, namely the non-convex problem of binary image segmentation based on *Euler's Elastica*.

Keywords: Curvature, segmentation, convex conjugate, convex envelope.

1 Introduction

Many problems in computer vision such as image restoration, image segmentation, stereo and motion can be formulated as energy minimization problems. The energy functions to be minimized can be developed in different settings. In a MRF setting, the energy can be derived from the maximum a posteriori (MAP) formulation on a discrete graph with a node set representing the image pixels, and an edge set defining the pixel interactions to measure smoothness. In a variational setting, images are interpreted as continuous functions, and differential operators are used to measure smoothness of the functions. Although, both methods are very different in their theoretical background, they have in common that in the end one has to solve a numerical optimization problem. Hence, their success largely depends on the ability of the underlying numerical algorithm to find a solution close to the global optimizer.

Energies that can be solved globally are rare in computer vision. If the label set is binary and the pairwise terms are submodular, the energies can be minimized globally by computing a minimum cut [6,12] on the graph. The equivalent formulation in a variational setting is given by the total variation. See [3] for detailed relationships. Multi-label problems can not be solved globally in general. A remarkable exception is the case where the label set is ordered linearly and the pairwise terms are convex functions. It has been shown in [8] that this class of problems can be solved exactly by computing a minimum cut on an extended graph in higher dimensions. In a variational setting, it has been shown that the same class of energies can be minimized globally by solving a minimal surface problem in higher dimensions [16].

A. Fusiello et al. (Eds.): ECCV 2012 Ws/Demos, Part III, LNCS 7585, pp. 283–292, 2012.

Recently, higher order terms have attracted a lot of attention in the MRF community, because of their ability to model complex interactions between image pixels. Minimizing higher order terms is even harder than minimizing pairwise terms. One possibility is to apply reduction schemes such as the one presented in [9] to transform the higher order terms into pairwise terms and to subsequently make use of certain relaxation techniques such as QPBO [7,17,10] to solve the transformed energy. Although, the transformed energy is in general non-submodular, QPBO can still compute a partially optimal solution. Another possibility is to solve the higher order problem directly [10]. It is also possible to use linear programming (LP) relaxation techniques which can turn the problem into a tractable linear programming problem. This approach has been originally presented in [19] and has been extended and improved over the years in various ways *e.g.* [18,22,14].

In this paper, we propose a method for minimizing general non-convex energies, and we will discuss it based on the task of binary image segmentation with length and curvature regularization. Such curvature-based models have attained considerable attention in recent years, due to psychophysical experiments on contour completion [11]. These experiments pointed out, that curvature plays an important role in human perception. Thus, there is an increasing interest to incorporate curvature information as a prior to various imaging problems. Unfortunately, such curvature depending functionals are hard to minimize, due to their strong non-convexity.

In principal, any non-convex energy could be minimized by computing its convex envelope and by minimizing the convex envelope instead. However, computing the convex envelope is not tractable for most problems. Therefore, we propose to minimize an approximated convex envelope, which is computed by splitting the energy into combinatorially tractable parts, and then globally minimize the sum of the convex envelopes of those parts. Note, that the main idea of splitting the problem into appropriate subproblems is similar to the dual decomposition technique, which was used in [13] to address discrete MRF-based optimization problems in computer vision.

The paper is organized as follows. In Section 2 we will describe the task of binary image segmentation with length and curvature regularization, and we will show how to split up the according cost function into a set of simple functions. In Section 3 we will show how to minimize such type of functions. In Section 4 we will present some results and in the last section we will finally give a short conclusion.

2 Binary Image Segmentation with Curvature Regularity

The task of binary image segmentation is to divide the domain $\Omega \subset \mathbb{R}^2$ of a given image $I : \Omega \to \mathbb{R}$ into foreground R and background $\Omega \setminus R$ by minimizing the functional

$$f(R) = \lambda \int_R d(x)dx + \int_{\partial R} \left(\alpha + \beta \left|\kappa_{\partial R}(x)\right|^p\right) d\mathcal{H}^1(x) \,, \tag{1}$$

where $d(x) = d_\Omega(x) - d_{\Omega\setminus R}(x)$. d_Ω and $d_{\Omega\setminus R}$ are functions, that depend on the input image I, and $\kappa_{\partial R}$ is the curvature of the boundary of R. Thus, the first integral in (1) represents the data-term weighted by λ. The second integral is Mumford's elastica curve model [15], where α weights the boundary length and β the curvature.

In what follows we will discretize $f(R)$ and split it into a finite sum of local functions $f_i(R)$, $1 \leqslant i \leqslant m$. Therefore, we have to divide our image domain Ω into a set of non-overlapping basic regions, called a cell-complex. Next, we define m local neighborhoods N_i, that are sets of connected basic regions and satisfy $\bigcup_{i=1}^{m} N_i = \Omega$.

In order to discretize the energy $f(R)$ we choose an arbitrary planar graph $G = (V, E)$ covering Ω with a vertex-set V and an edge-set E. The face-set F of the graph G defines the set of basic regions. Hence, the discrete version of (1) can be defined as

$$f_D(R) = \lambda \sum_{k=1}^{|F|} R_k d_k + \sum_{j=1}^{|V|} \left(\alpha\, l_j(\pi_j(R)) + \beta\, \kappa_j(\pi_j(R))\right) \ , \tag{2}$$

where $R \in \{0,1\}^{|F|}$ and each element R_k indicates if the basic region F_k is part of the foreground or not. $d \in \mathbb{R}^{|F|}$ contains the data-term for each basic region, and l_j and κ_j are functions calculating the length and curvature costs at each vertex V_j. Note, that for calculating the regularization costs for V_j we only need those elements of R, that correspond to the adjacent regions of V_j. Therefore, we select the according elements of R with a linear operator π_j. Now, we define for each neighborhood N_i a function

$$\begin{aligned} f_i(R) = \lambda &\sum_{k=1}^{|F|} \chi_{N_i}(F_k)\ c_{ik}\ R_k\ d_k \\ &+ \sum_{j=1}^{|V|} \chi_{\hat{N}_i}(V_j)\ \hat{c}_{ij}\ \left(\alpha\, l_j(\pi_j(R)) + \beta\, \kappa_j(\pi_j(R))\right) \ , \end{aligned} \tag{3}$$

where $\hat{N}_i$ is the set of interior vertices of the neighborhood N_i, c_{ik} and $\hat{c}_{ij}$ are positive normalization coefficients satisfying

$$\sum_{i=1}^{m} \chi_{N_i}(F_k)\, c_{ik} = \sum_{i=1}^{m} \chi_{\hat{N}_i}(V_j)\, \hat{c}_{ij} = 1 \ . \tag{4}$$

As a result we can write (2) as a sum of simple functions $f_i(R)$.

$$f_D(R) = \sum_{i=1}^{m} f_i(R) \tag{5}$$

Now we explain how the functions l_j and κ_j calculate the regularization costs for the vertex V_j by an example. Consider the situation sketched in Fig. 1(a). Here, the gray regions belong to the foreground and the white regions belong to the background. Thus, the node V_1 belongs to the boundary of R with length cost $l_1(R_1, \ldots, R_5) = \frac{1}{2}(|e_1| + |e_4|)$, where $|e|$ is the length of the edge e. The division by two is necessary to get the correct length, when summing up over all vertices.

To measure the curvature we use a discrete formulation introduced by Bruckstein *et al*. [2]. Hence, in the case at hand, shown in Fig. 1(a), the curvature cost is

$$\kappa_1(R_1, \ldots, R_5) = \frac{\phi^p_{e_1,e_4}}{\min\{|e_1|, |e_4|\}^{p-1}} \ , \tag{6}$$

where ϕ_{e_1,e_4} is the angle between e_1 and e_4.

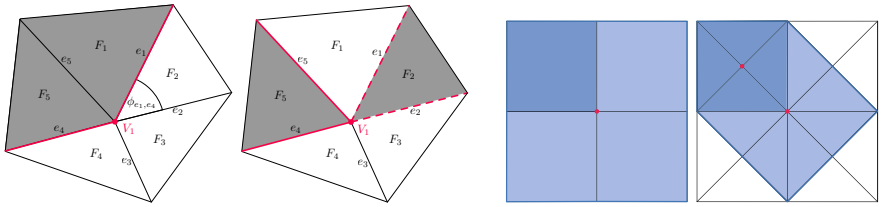

(a) (b) (c) 4-connectivity (d) 8-connectivity

Fig. 1. (a) and (b) show two foreground/background configurations for the adjacent regions of a vertex V_1. Such configurations are used to assign regularization costs to the vertex V_1. (a) shows a common boundary and (b) a situation where two boundary segments meet in a single vertex. (c) and (d) illustrate possible neighborhoods N_i (blue regions) and their corresponding interior vertices $\hat{N}_i$ (marked by purple dots), for 4- and 8-connectivity. These neighborhoods are defined for each pixel in the image, where a single pixel is indicated by the dark blue region.

The regularization costs are precomputed for certain configurations of R_k within a neighborhood N_i. Note that in the case where all adjacent regions of a vertex belong to either foreground or background the regularization cost vanishes. Configurations where different boundary segments meet at a single vertex (compare Fig. 1(b)) can be treated separately, *i.e.* we can calculate reasonable regularization cost, or we can also exclude such configurations by setting the costs to infinity. Thus, with our approach we can avoid problems with self-intersecting boundaries, present in the approach of Schoenemann *et al.* [20,21]. Furthermore, we can calculate the correct boundary cost for each configuration, including those, where boundaries meet at a single vertex. Such configurations are *e.g.* incorrectly handled in the 3-clique strategy presented by El-Zehiry and Grady [5].

Finally, Fig. 1(c) and 1(d) show some possibilities to define the cell-complex and the according neighborhoods.

3 Approximate Envelope Minimization

We are given the following minimization problem

$$\min_{x \in X} f(x), \quad \text{with} \quad f(x) = \sum_{i=1}^{m} \left(\tilde{f}_{A_i} \circ \pi_{A_i} \right)(x) \,, \tag{7}$$

where X is the n dimensional Euclidean space $\mathbb{R}^n$ and $f : X \to \mathbb{R}$ is a non-convex function. Moreover, the function $f(x)$ fulfills a certain splitting property, which allows us to rewrite it as the sum of functions $f_{A_i}(x) := (\tilde{f}_i \circ \pi_{A_i})(x)$. Here $\tilde{f}_{A_i} : X_{A_i} \to \mathbb{R}$ is a function defined on an appropriate subspace $X_{A_i} \subseteq X$. Each subspace X_{A_i} is implicitly defined according to an ordered index set $A_i = \{k_1, \ldots, k_{l_{A_i}}\}$ via the mapping $\pi_{A_i} : X \to X_{A_i}$,

$$\pi_{A_i}(x) := \left(x[k_1], \ldots, x[k_{l_{A_i}}] \right)^T \,, \tag{8}$$

where $x[t]$ is the t^{th} element of x. Furthermore, we define the following pseudo re-projection $\pi^+_{A_i} : X_{A_i} \to X$

$$\pi^+_{A_i}(x) := y \in X, \text{ with } y[s] = \begin{cases} x[r] & \text{if } s \in A_i \text{ and } r = |\{k \leqslant s \,|\, k \in A_i\}| \\ 0 & \text{else} , \end{cases} \tag{9}$$

where $|A|$ denotes the cardinality of the set A, and we assume that

$$X = \left\{ \pi^+_{A_i}(x_{A_i}) \mid x_{A_i} \in X_{A_i}, 1 \leqslant i \leqslant m \right\} . \tag{10}$$

It is essential for our approach that the domain of a single function $\tilde{f}_{A_i}$ is much smaller than the domain of the function f. Further, note that the convex conjugate f^* and bi-conjugate f^{**} are by definition convex functions since they are written through a point-wise supremum over affine functions. Thus, it is clear that the minimizer of the non-convex optimization problem (7) can be computed by minimizing the convex envelope f^{**} instead. Unfortunately, computing f^{**} is intractable in many situations and hence there is not much hope to find a closed form expression for f^{**}. However, there is hope that we can compute a weaker approximation to the *true* convex envelope by allowing for some simplifications. The key idea is presented in Proposition 1.

Proposition 1. *Given the definitions above, the following inequality holds*

$$f^{**}(x) = \left(\sum_{i=1}^{m} \tilde{f}_{A_i} \circ \pi_{A_i} \right)^{**} (x) \geqslant \sum_{i=1}^{m} \left(\tilde{f}^{**}_{A_i} \circ \pi_{A_i} \right) (x) =: \bar{f}(x) . \tag{11}$$

Proof. See supplementary material. □

Proposition 1 basically shows, that a weaker (less tight) envelope $\bar{f}(x)$ is obtained by computing the sum of the bi-conjugate functions instead of the bi-conjugate of the sum. On the one hand, the expected advantage of the former is that the functions $\tilde{f}^{**}_{A_i}$ might be much easier to compute than $(\sum_{i=1}^{m} f_{A_i})^{**}$. On the other hand, the quality of the relaxation depends on the splitting of f.

3.1 The Approximate Convex Envelope

As indicated by the task stated in Section 2, many problems in computer vision suggest a natural splitting of the energy, which is mainly motivated by the fact that energies are usually modeled by defining the interaction between certain elements in a local neighborhood. Thus each function $\tilde{f}_{A_i}$ in (7) can be defined to model the interaction in one specific neighborhood.

Moreover, based on a given splitting one can improve the relaxation by grouping certain functions f_{A_i} together to a new function f_{G_j}, *i.e.* summing up the according functions $\tilde{f}_{A_i}$. Note, that this grouping strategy improves the relaxation, but simultaneously increases the combinatorial complexity. In the following, we will work with an approximation to the true convex envelope given by

$$\hat{f}(x) = \sum_{j=1}^{\hat{m}} \left(\tilde{f}^{**}_{G_j} \circ \pi_{G_j} \right) (x) . \tag{12}$$

By using Proposition 1, one can easily verify, that the following inequalities hold for all $x \in X$.

$$f^{**}(x) \geqslant \hat{f}(x) \geqslant \bar{f}(x) \tag{13}$$

Hence, instead of minimizing f or f^{**}, we propose to minimize $\hat{f}$. The resulting minimization problem can be rewritten as the following saddle-point problem.

$$\min_{x \in X} \hat{f}(x) = \min_{x} \sum_{j=1}^{\hat{m}} \tilde{f}_{G_j}^{**}\left(\pi_{G_j}(x)\right) = \min_{x} \max_{y} \sum_{j=1}^{\hat{m}} \left\langle \pi_{G_j}(x), y_{G_j} \right\rangle - \tilde{f}_{G_j}^{*}(y_{G_j}) \ , \tag{14}$$

where $y = \left(y_{G_1}, \ldots, y_{G_{\hat{m}}}\right)^{\mathrm{T}}$ and $y_{G_j} \in X_{G_j}^{*}$. In order to solve (14), it is crucial to handle the $\tilde{f}_{G_j}^{*}$ functions, which can be done in various ways. For simplicity we will use a polyhedral approximation which is explained in Section 3.2.

3.2 Polyhedral Approximations

In many interesting cases we do not expect that we will find an explicit formula for the conjugate functions $\tilde{f}_{G_j}^{*}$. Hence, we need to find an implicit representation. The key comes through the following inequality, which is provided by the definition of the convex conjugate. Given a function g and its convex conjugate g^{*}, one has for any x, $y \in X$

$$g^{*}(y) \geqslant \langle x, y \rangle - g(x) \ . \tag{15}$$

By definition equality is reached by taking the supremum over x of the right hand side. Let us now interpret the above inequality in terms of the epigraphs of g and g^{*}. Let (x, s) be a point in the epigraph of g. Then one has for any (y, t) in the epigraph of g^{*}

$$t \geqslant \langle x, y \rangle - s \ . \tag{16}$$

In other words, we can refine a polyhedral approximation of g^{*} by successively sampling points from the epigraph of g, and by adding inequality constraints as denoted in (16). Clearly, the largest (and hence most successful) constraints will be generated by taking points on the graph of g, *i.e.* points $(x, g(x))$. The idea is now to obtain an approximation of g^{*} by computing a finite number of points $\{(x_i, g_i),\ 1 \leqslant i \leqslant k\}$ on the graph of g. As shown above, we can generate a set of constraints of the form

$$t \geqslant \langle x_i, y \rangle - g_i \quad \text{for} \quad 1 \leqslant i \leqslant k \ . \tag{17}$$

Of course, the more constraints we add, the better the approximation will be. Note, that in the case of discrete multi-labeling problems the domain X is already finite. Thus, by taking all points in X as sampling points, we obtain an exact representation.

3.3 Linear Programming Formulation

Now we propose a relaxed version of the saddle-point problem described in (14) by using polyhedral approximations according to (17). By approximating each $\tilde{f}_{G_j}^{*}(x)$ in (14) one obtains the following relaxation

$$\min_{x} \max_{y,\,t} \sum_{j=1}^{n} \left\langle \pi_{G_j}(x), y_{G_j} \right\rangle - t_{G_j} \tag{18}$$
$$\text{s.t. } t_{G_j} \geqslant \left\langle s^i_{G_j}, y_{G_j} \right\rangle - f_{G_j}(s^i_{G_j}), \text{ for } 1 \leqslant i \leqslant k_{G_j}\ ,\ 1 \leqslant j \leqslant n\ ,$$

where $t = (t_1, \dots t_n)^{\mathrm{T}}$, and $s^i_{G_j}$ denotes the i^{th} sampling position of the group G_j. By introducing Lagrange multipliers $\mu(s^i_{G_j})$ for each constraint we can identify the t_{G_j} variables as Lagrange multipliers for $\sum_{i=1}^{k_{G_j}} \mu(s^i_{G_j}) = 1$. Thus, the vector μ^j, which contains all elements of μ, that belong to the j^{th} group, has to be on the standard k_{G_j}-dimensional unit simplex $S_{k_{G_j}}$. In a final step, we identify the y_{G_j} variables as Lagrange multipliers for $\pi_{G_j}(x) = \sum_{i=1}^{k_{G_j}} s^i_{G_j}\, \mu(s^i_{G_j})$ and rewrite (18) as

$$\min_{x,\mu} \sum_{j=1}^{n} \sum_{i=1}^{k_j} f_{G_j}(s^i_{G_j})\, \mu(s^i_{G_j}) \tag{19}$$
$$\text{s.t. } \mu^j \in S_{k_{G_j}}\ , \quad \pi_{G_j}(x) = \sum_{i=1}^{k_{G_j}} s^i_{G_j}\, \mu(s^i_{G_j})\ , \quad 1 \leqslant j \leqslant n\ .$$

By taking a closer look at (19), one can see, that the solution x is obtained as a collection of convex combinations of sampling positions within the groups. Furthermore, due to the definition of $\pi_{G_j}(x)$ these convex combinations have to coincide with the solutions of neighboring groups for elements, where their domains intersect.

The LP in (19) shows an interesting connection to the Shlezinger relaxation [19], which had been generalized by Werner *et al*. [23]. The only differences are the so called marginalization constraints obtained by discretizing the label space, which has not been done in the proposed method.

3.4 Primal-Dual Algorithm

In this section we describe how to solve the saddle-point problem (14) via the primal-dual approach described by Chambolle and Pock [4]. Therefore, it will be convenient to introduce a linear assembling operator A, such that $Ay = \sum_{j=1}^{\hat{m}} \pi^+_{G_j}(y_{G_j})$. By using this operator, we can rewrite (14) as

$$\min_{x} \max_{y} \langle Ay, x \rangle - \sum_{j=1}^{\hat{m}} f^*_{G_j}(y_{G_j})\ . \tag{20}$$

A numerical algorithm to solve the above saddle-point problem is given by

$$\begin{cases} x^{k+1} = (x^k - \tau A y^k) \\ \bar{x}^{k+1} = x^{k+1} + \theta(x^{k+1} - x^k) \\ y^{k+1}_{G_j} = (\mathrm{id} + \sigma \partial f^*_{G_j})^{-1} \left(y^k_{G_j} + \sigma \left[A^* \bar{x}^{k+1} \right]_j \right) \quad \forall\, 1 \leqslant j \leqslant \hat{m}\ , \end{cases} \tag{21}$$

where τ and σ are chosen such that $\tau\sigma\|A\|^2 < 1$ and $\theta \in [0, 1]$.

Thus, the only interesting thing to show is the calculation of the resolvent operator $(I + \sigma \partial f^*_{G_j})^{-1}(z)$, which is defined as the solution of

$$\operatorname*{argmin}_w \frac{\|w - z\|^2}{2\sigma} + f^*_{G_j}(w) \; . \tag{22}$$

By using the polyhedral approximation (17), we can rewrite (22) as

$$\min_{w,t} \frac{\|w - z\|^2}{2\sigma} + t \tag{23}$$

$$\text{s.t. } \left\langle s^i_{G_j}, w \right\rangle - f_{G_j}(s^i_{G_j}) - t \leqslant 0 \; , \quad \forall\, 1 \leqslant i \leqslant k_{G_j} \; .$$

Introducing Lagrange multipliers $\mu = (\mu(s^1_{G_j}), \ldots, \mu(s^{k_{G_j}}_{G_j}))^{\mathrm{T}}$ for the inequality constraints, we arrive at

$$\min_{w,t} \max_{\mu \geqslant 0} \frac{\|w - z\|^2}{2\sigma} + t + \sum_{i=1}^{k_{G_j}} \left(\left\langle s^i_{G_j}, w \right\rangle - f_{G_j}(s^i_{G_j}) - t \right) \mu(s^i_{G_j}) \; . \tag{24}$$

Thus, minimizers in w are characterized by

$$w = z - \sigma \sum_{i=1}^{k_{G_j}} s^i_{G_j} \; \mu(s^i_{G_j}) \; . \tag{25}$$

By substituting this relation into the proximal map (24), we can get rid of w, and we see that t itself is just a Lagrange multiplier for $\sum_{i=1}^{k_{G_j}} \mu(s^i_{G_j}) = 1$, *i.e.* μ has to be in the standard k_{G_j}-dimensional unit simplex $S_{k_{G_j}}$. Hence we obtain

$$\min_{\mu \in S_{k_{G_j}}} \frac{\sigma}{2} \left\| \sum_{i=1}^{k_{G_j}} s^i_{G_j} \; \mu(s^i_{G_j}) - \frac{z}{\sigma} \right\|^2 + \sum_{i=1}^{k_{G_j}} f_{G_j}(s^i_{G_j}) \; \mu(s^i_{G_j}) \; , \tag{26}$$

which is a simplex constrained quadratic program, that can be solved *e.g.* with the FISTA algorithm presented by Beck and Teboulle [1].

Compared to the minimization strategy presented in Section 3.3, the quadratic program is defined point-wise, thus the μ variables need not to be saved for the entire set of neighborhoods, yielding a memory-efficient algorithm.

4 Results

In this section we present some experimental results for our example task of binary image segmentation based on Mumford's elastica curve model [15].

Here the data term in (2) is computed as $d = (I - \mu_f)^2 - (I - \mu_b)^2$, where I is the input image and $\mu_{f,b}$ are the mean values of the foreground and background. We set $\lambda = 1$ and evaluate the results for different α and β configurations (compare Fig. 2).

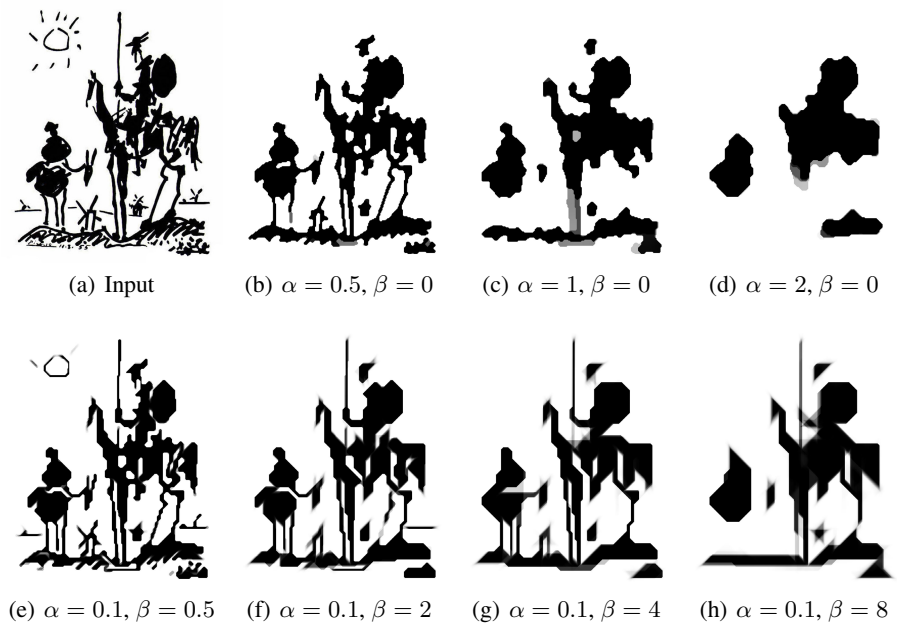

(a) Input (b) $\alpha = 0.5, \beta = 0$ (c) $\alpha = 1, \beta = 0$ (d) $\alpha = 2, \beta = 0$

(e) $\alpha = 0.1, \beta = 0.5$ (f) $\alpha = 0.1, \beta = 2$ (g) $\alpha = 0.1, \beta = 4$ (h) $\alpha = 0.1, \beta = 8$

Fig. 2. Segmentation results for the Don Quixote image by Pablo Picasso (256×216 px) using a cell-complex, that corresponds to a 8-connected graph in graph cut frameworks (compare Fig. 1(d)). We set $\lambda = 1$ and show results for different α and β values. The processing time is between 1 and 5 minutes, depending on the strength of the regularization.

As expected the segmentation results with curvature regularity ($\beta \neq 0$) favor certain directions, and lead to polygonal structures. Moreover, one can see that the curvature regularization tends to preserve elongated structures, whereas the length regularization tries to remove them.

5 Conclusion

In this paper we proposed a novel approach for minimizing non-convex functions, that fulfill a certain splitting property. As such functions are ubiquitous in computer vision, the proposed approach is applicable to a broad class of problems. Moreover, we demonstrated the general applicability of the approach on the task of binary image segmentation based on Mumford's elastica curve model [15].

References

1. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences 2(1), 183 (2009)
2. Bruckstein, A.M., Netravali, A.N., Richardson, T.J.: Epi-convergence of discrete elastica. Applicable Analysis 79, 137–171 (2001)

3. Chambolle, A.: Total Variation Minimization and a Class of Binary MRF Models. In: Rangarajan, A., Vemuri, B.C., Yuille, A.L. (eds.) EMMCVPR 2005. LNCS, vol. 3757, pp. 136–152. Springer, Heidelberg (2005)
4. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. Journal of Mathematical Imaging and Vision 40(1), 120–145 (2010)
5. El-Zehiry, N.Y., Grady, L.: Fast global optimization of curvature. In: CVPR, pp. 3257–3264 (2010)
6. Greig, D.M., Porteous, B.T., Seheult, A.H.: Exact maximum a posteriori estimation for binary images. Journal of the Royal Statistics Society 51(Series B), 271–279 (1989)
7. Hammer, P.L., Hansen, P., Simeone, B.: Roof duality, complementation and persistency in quadratic 0-1 optimization. Mathematical Programming 28(2), 121–155 (1984)
8. Ishikawa, H.: Exact optimization for markov random fields with convex priors. IEEE Transactions on Pattern Analysis and Machine Intelligence 25(10), 1333–1336 (2003)
9. Ishikawa, H.: Transformation of general binary MRF minimization to the first-order case. IEEE Transactions on Pattern Analysis and Machine Intelligence 33, 1234–1249 (2011)
10. Kahl, F., Strandmark, P.: Generalized roof duality for pseudo-boolean optimization. In: ICCV (2011)
11. Kanizsa, G.: Organization in Vision. Praeger, New York (1979)
12. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts. IEEE Transactions on Pattern Analysis and Machine Intelligence 26(2), 147–159 (2004)
13. Komodakis, N., Paragios, N., Tziritas, G.: MRF energy minimization and beyond via dual decomposition. IEEE Transactions on Pattern Analysis and Machine Intelligence 33, 531–552 (2011)
14. Komodakis, N., Tziritas, G.: A new framework for approximate labeling via graph cuts. In: ICCV, pp. 1018–1025 (2005)
15. Mumford, D.: Elastica and computer vision. In: Algebraic Geometry and Its Applications, pp. 491–506 (1994)
16. Pock, T., Cremers, D., Bischof, H., Chambolle, A.: Global solutions of variational models with convex regularization. SIAM Journal on Imaging Sciences 3(4), 1122–1145 (2010)
17. Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M.: Optimizing binary MRFs via extended roof duality. In: CVPR, pp. 1–8 (2007)
18. Schlesinger, D., Flach, B.: Transforming an arbitrary minsum problem into a binary one. Technical Report TUD-FI06-01, Dresden University of Technology (2006)
19. Schlesinger, M.: Syntactic analysis of two-dimensional visual signals in noisy conditions. Kibernetika 4, 113–130 (1976) (in Russian)
20. Schoenemann, T., Kahl, F., Cremers, D.: Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In: ICCV, September 29-October 2, pp. 17–23 (2009)
21. Schoenemann, T., Kahl, F., Masnou, S., Cremers, D.: A linear framework for region-based image segmentation and inpainting involving curvature penalization. International Journal of Computer Vision (to appear, 2012)
22. Werner, T.: A linear programming approach to max-sum problem: A review. IEEE Transactions on Pattern Analysis and Machine Intelligence 29(7), 1165–1179 (2007)
23. Werner, T.: Revisiting the linear programming relaxation approach to gibbs energy minimization and weighted constraint satisfaction. IEEE Trans. Pattern Anal. Mach. Intell. 32(8), 1474–1488 (2010)