

A Convex Framework for Image Segmentation with Moment Constraints

Maria Klodt and Daniel Cremers
TU Munich, Germany

{maria.klodt,daniel.cremers}@in.tum.de

Abstract

Convex relaxation techniques have become a popular approach to image segmentation as they allow to compute solutions independent of initialization to a variety of image segmentation problems. In this paper, we will show that shape priors in terms of moment constraints can be imposed within the convex optimization framework, since they give rise to convex constraints. In particular, the lower-order moments correspond to the overall volume, the centroid, and the variance or covariance of the shape and can be easily imposed in interactive segmentation methods. Respective constraints can be imposed as hard constraints or soft constraints. Quantitative segmentation studies on a variety of images demonstrate that the user can easily impose such constraints with a few mouse clicks, giving rise to substantial improvements of the resulting segmentation, and reducing the average segmentation error from 12% to 0.35%. GPU-based computation times of around 1 second allow for interactive segmentation.

1. Introduction

1.1. Shape Optimization in Computer Vision

Shape optimization is at the heart of several classical computer vision problems such as image segmentation and multi view reconstruction. Following a series of seminal papers [12, 1, 17], functional minimization has become the established paradigm to solve shape optimization problems such as image segmentation or 3D reconstruction. In the spatially discrete setting the study of the corresponding binary labeling problems goes back to the spin-glass models introduced in the 1920's [11]. Popular algorithms to solve the arising shape optimization problems include level set methods [18], graph cuts [9] or convex relaxation [4].

In this paper, we focus on a class of functionals of the form:

$$E(S) = \int_{\text{int}(S)} f(x) dx + \int_S g(x) dA, \quad (1)$$

where S denotes a hyper surface in \mathbb{R}^d , i.e. a set of closed

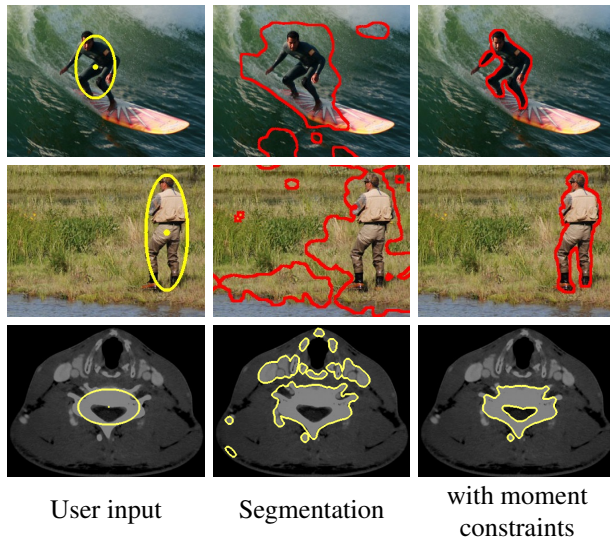


Figure 1. We propose a convex formulation for interactive image segmentation which allows to impose constraints on moments of arbitrary order. In particular, constraints on the lower order moments (area, centroid, covariance) are easily transmitted through mouse interaction (left). They allow to stabilize the segmentation process while preserving fine-scale details of the shape (right).

boundaries in the case of 2D image segmentation or a set of closed surfaces in the case of 3D segmentation and multi view reconstruction. The functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g : \mathbb{R}^d \rightarrow \mathbb{R}^+$ are application dependent. In a statistical framework for image segmentation, for example,

$$f(x) = \log p_{bg}(I(x)) - \log p_{ob}(I(x)),$$

may denote the log likelihood ratio for observing the color $I(x)$ at a point x given that x is part of the background or the object, respectively.

The second term in (1) corresponds to the area (for $d = 3$) or the boundary length (for $d = 2$), measured in a metric given by the function g . In the context of image segmentation, g may be a measure of the local edge strength – as in the geodesic active contours [3, 13] – which energetically favors segmentation boundaries along strong intensity gradients. In the context of multi view reconstruction,

$g(x)$ typically measures the photo-consistency among different views of the voxel x , where low values of g indicate a strong agreement from different cameras on the observed patch intensity.

1.2. Shape Priors for Image Segmentation

There has been much research on imposing prior shape knowledge into image segmentation. While it was shown that segmentation results can be substantially improved by imposing shape priors [10, 5, 8], existing approaches typically suffer from the following problems:

- Apart from a few exceptions such as [20] – computable solutions are only *locally* optimal. As a consequence, one typically needs appropriate initializations and solutions may be arbitrarily far from the globally minimal ones. Other notable exceptions were designed by Veksler and coworkers for specific scenarios such as “compact objects” [6] or “star-shaped objects” [21].
- Many shape priors have a rather fine granularity in the sense that they impose the object silhouette to be consistent with those silhouettes observed in a training set [5, 7]. The degree of abstraction is typically rather small. In particular, deviations of the observed shape from the training shapes are (elastically) suppressed by the shape prior. This is particularly undesirable in medical image segmentation where malformations of organs (that make it deviate from the training shapes of healthy organs) should be detected rather than ignored. It may therefore be of interest to merely impose some coarse-level shape information rather than imposing the exact form of the object.

An alternative approach that may provide a remedy for both of the above problems is to impose moment constraints. In particular, the lower-order moments allow to constrain the volume, the centroid and the size or covariance of objects without imposing any constraints on their local shape. A related idea of using *Legendre moments* (albeit in a local optimization scheme) was developed in [8].

In a convex formulation of multiple view 3D reconstruction, it was recently shown [14] that one can impose additional convex constraints which assure that the computed minimal surfaces are silhouette-consistent. Essentially this constraint can be seen as a *volume constraint*: The volume along any ray from the camera center must be at least 1 if that ray passes through the silhouette and zero otherwise. In the two-dimensional case, a related constraint was recently proposed as a bounding box prior for image segmentation [16].

1.3. Contribution

In this paper we show that one can impose an entire family of moment constraints in the framework of convex shape optimization, thereby generalizing from the zeroth order moment (volume) to higher order moments (centroid, scale, covariance, etc). In particular, all moment constraints - both soft or hard - correspond to convex constraints. As a consequence we can compute moment-constrained shapes which are independent of initialization and lie within a bound of the optimum.

The outline of the paper is as follows. In Section 2, we will briefly review a framework for convex relaxation and thresholding which allows to efficiently compute global minima of the above energies in a spatially continuous setting. In Section 3, we will then show that moment constraints can be imposed as convex constraints within the optimization framework. In Section 4 we show how the arising optimization problem can be minimized using efficient GPU-accelerated PDEs. In Section 5 we present experimental results and a quantitative evaluation showing that interactive segmentation results can be drastically improved using moment constraints.

2. Shape Optimization via Convex Relaxation

Functionals of the form (1) can be globally optimized in a spatially continuous setting by means of convex relaxation and thresholding [4]. To this end, one reverts to an implicit representation of the hyper surface S using an indicator function $u \in BV(\mathbb{R}^d; \{0, 1\})$ on the space of binary functions of bounded variation, where $u = 1$ and $u = 0$ denote the interior and exterior of S . The functional (1) defined on the space of surfaces S is therefore equivalent to the functional

$$E(u) = \int f(x) u(x) dx + \int g(x) |Du(x)|, \quad (2)$$

where the second term in (2) is the weighted total variation. Here Du denotes the distributional derivative which for differentiable functions u boils down to $Du(x) = \nabla u(x) dx$. By relaxing the binary constraint and allowing the function u to take on values in the interval between 0 and 1, the optimization problem becomes that of minimizing the convex functional (2) over the convex set $BV(\mathbb{R}^d; [0, 1])$. Global minimizers u^* of this relaxed problem can therefore efficiently be computed, for example by a simple gradient descent procedure.

The thresholding theorem [4] assures that thresholding the solution u^* of the relaxed problem preserves global optimality for the original binary labeling problem. We can therefore compute global minimizers for functional (2) in a spatially continuous setting as follows: Compute a global minimizer u^* of (2) on the convex set $BV(\mathbb{R}^d; [0, 1])$ and threshold the minimizer u^* at any value $\mu \in (0, 1)$.

3. Moment Constraints for Segmentation

In the following, we will successively constrain the moments of the segmentation and show how all of these constraints give rise to nested convex sets. To this end we will represent shapes in d dimensions as binary indicator functions $u \in BV(\Omega; \{0, 1\})$ of bounded variation on the domain $\Omega \subset \mathbb{R}^d$. We will denote the convex hull of this set by $\mathcal{B} = BV(\Omega; [0, 1])$.

3.1. Area Constraint

We can impose that the area of the shape u to be bounded by constants $c_1 \leq c_2$ by constraining u to lie in the set:

$$\mathcal{C}_0 = \left\{ u \in \mathcal{B} \mid c_1 \leq \int_{\Omega} u \, dx \leq c_2 \right\}.$$

Proposition 1. *For any constants $c_2 \geq c_1 \geq 0$, the set \mathcal{C}_0 is convex.*

Proof. Let $u_1, u_2 \in \mathcal{C}_0$ be two elements from this set. Then for any convex combination $u_{\alpha} = \alpha u_1 + (1 - \alpha)u_2, \alpha \in [0, 1]$ of these elements we have:

$$\int_{\Omega} u_{\alpha} \, dx = \alpha \int_{\Omega} u_1 \, dx + (1 - \alpha) \int_{\Omega} u_2 \, dx.$$

As a consequence we have $c_1 \leq \int_{\Omega} u_{\alpha} \, dx \leq c_2$ such that $u_{\alpha} \in \mathcal{C}_0$. \square

In practice, we can either impose an exact area by setting $c_1 = c_2$, or we can impose upper and lower bounds on the area. Alternatively, we can impose a soft area constraint by enhancing the functional (2) as follows:

$$E_{total}(u) = E(u) + \lambda \left(\int_{\Omega} u \, dx - c \right)^2, \quad (3)$$

which imposes a soft constraint with a weight $\lambda > 0$ favoring the area of the estimated shape to be near $c \geq 0$. Clearly, the functional (3) is also convex.

3.2. Centroid Constraint

Assume that someone gave us some bounds about the centroid (center of gravity) for the object we want to reconstruct. We can impose these bounds by constraining the solution u to the set \mathcal{C}_1 :

$$\mathcal{C}_1 = \left\{ u \in \mathcal{B} \mid \mu_1 \leq \frac{\int_{\Omega} x u \, dx}{\int_{\Omega} u \, dx} \leq \mu_2 \right\}, \quad (4)$$

where all inequalities are to be taken point wise and $\mu_1, \mu_2 \in \mathbb{R}^d$. This imposes the centroid to lie between the two constants $\mu_1 \leq \mu_2$. In particular, for $\mu_1 = \mu_2$, the centroid is fixed.

Proposition 2. *For any constants $\mu_2 \geq \mu_1 \geq 0$, the set \mathcal{C}_1 is convex.*

Proof. The inequality constraint in (4) is equivalent to

$$\mu_1 \int_{\Omega} u \, dx \leq \int_{\Omega} x u \, dx \leq \mu_2 \int_{\Omega} u \, dx,$$

which are clearly two linear inequality constraints. \square

Alternatively, we can impose the centroid as a soft constraint by minimizing the energy:

$$E_{total}(u) = E(u) + \lambda \left(\int_{\Omega} (\mu - x) u \, dx \right)^2,$$

which is also convex in u . Interestingly this soft constraint does not minimize the quadratic difference to the specified center μ ; the latter would not be convex. In contrast to the hard constraint, this soft constraint unfortunately exhibits a preference toward smaller shapes as it vanishes with decreasing object size:

$$\left(\int_{\Omega} (\mu - x) u \, dx \right)^2 = \left(\mu - \frac{\int_{\Omega} x u \, dx}{\int_{\Omega} u \, dx} \right)^2 b^2,$$

where $b = \int_{\Omega} u \, dx$ denotes the size of the object.

3.3. Imposing Higher-Order Moments

Clearly, the proposed concept can be generalized to moments of successively higher order, where we shall focus on so-called central moments (i.e. moments with respect to a specified centroid). In particular, the respective structures will generally be tensors of higher dimension. For example, one can impose the covariance structure by considering the following convex set:

$$\mathcal{C}_2 = \left\{ u \in \mathcal{B} \mid A_1 \leq \frac{\int_{\Omega} (x - \mu)(x - \mu)^{\top} u \, dx}{\int_{\Omega} u \, dx} \leq A_2 \right\}, \quad (5)$$

where the inequality constraint should be taken element wise. Here $\mu \in \mathbb{R}^d$ denotes the center and $A_1, A_2 \in \mathbb{R}^{d \times d}$ denote symmetric matrices such that $A_1 \leq A_2$ element wise. This constraint is particularly meaningful if one additionally constrains the centroid to be μ , i.e. considers the intersection of the set (5) with a set of the form (4).

Note that this allows, in particular, to constrain the scale σ of the object, because:

$$\sigma^2 = \frac{\int_{\Omega} (x - \mu)^2 u \, dx}{\int_{\Omega} u \, dx} = \text{tr} \frac{\int_{\Omega} (x - \mu)(x - \mu)^{\top} u \, dx}{\int_{\Omega} u \, dx}.$$

From the constraint in (5) it follows that:

$$\text{tr}(A_1) \leq \sigma^2 \leq \text{tr}(A_2).$$



Figure 2. Segmentation results with higher order moment constraints. By imposing constraints of increasing order (up to 12th order) more and more fine scale details of the shape are restored.

In general, we can impose constraints on moments of any order $k \in \mathbb{N}$:

$$\mathcal{C}_{i_1 \dots i_d} = \left\{ u \in \mathcal{B} \mid a_{i_1 \dots i_d} \leq \frac{\int_{\Omega} (x_1 - \mu_1)^{i_1} \dots (x_d - \mu_d)^{i_d} u \, dx}{\int u \, dx} \leq b_{i_1 \dots i_d} \right\},$$

where $i_1 + \dots + i_d = k$ and $a_{i_1 \dots i_d} \leq b_{i_1 \dots i_d}$ can be chosen arbitrarily to constrain the moment tensor of order k . Here x_i denotes the i -th component of x .

Proposition 3. *For all $i_1, \dots, i_d \in \mathbb{N}$ and for any constants $a_{i_1 \dots i_d} \leq b_{i_1 \dots i_d}$, the set $\mathcal{C}_{i_1 \dots i_d}$ is convex.*

Proof. The proof is analogous to that of proposition 1. \square

3.4. Imposing a Hierarchy of Shape Details

The above properties allow to impose various constraints on the shape associated with the indicator function u . Imposing more and more constraints of increasingly higher order leads to a smaller and smaller intersection of the associated convex sets as a feasible domain of the shape and a corresponding hierarchy of shape details being imposed in the segmentation. How much shape detail can one impose in this manner?

Proposition 4. *Similarity to any given shape can be imposed at arbitrary detail by imposing convex moment constraints of increasingly higher order.*

Proof. According to the uniqueness theorem of moments [19], the function u is uniquely defined by its moment sequence. \square

Figure 2 shows an example of segmentations with high order moment constraints: While the higher-order moments allow to recover fine-scale shape details, the shape improvements due to higher order constraints are fairly small. Furthermore imposing moments of higher order is not very practical: Firstly the user cannot estimate these moments visually. Secondly the user cannot transmit respective higher-order tensors through a simple mouse interaction. Instead, having the image data determine the shape’s fine scale structure turns out to be far more useful.

Therefore, for the application we consider – namely interactive image segmentation – we shall in the following

limit ourselves to imposing moments up to 2nd order (volume, center of mass, scale and covariance).

4. Reconstruction under Moment Constraints

Shape optimization and image segmentation with respective moment constraints can now be done by minimizing convex energies under respective convex constraints.

Let \mathcal{C} be a specific convex set containing knowledge about respective moments of the desired shape – given by an intersection of the above convex sets. Then we can compute segmentations by solving the convex optimization problem

$$\min_{u \in \mathcal{C}} E(u),$$

with $E(u)$ given in (2). In this paper we solved the Euler-Lagrange equations:

$$0 = \operatorname{div} \left(g \frac{\nabla u}{|\nabla u|} \right) - f$$

using the lagged diffusivity approach that was presented in [15]. We enforce constraints during the optimization by back-projecting the current segmentation onto the constraint set after every iteration using the algorithm of [2].

Unfortunately, the threshold theorem [4] guaranteeing optimality for the unconstrained binary labeling problem does not generalize to the constrained optimization problems considered here. Nevertheless, we can prove the following optimality bound.

Proposition 5. *Let $u^* = \arg \min_{u \in \mathcal{C}} E(u)$ be a minimizer of the relaxed problem and E_{opt} the (unknown) minimum of the corresponding binary problem. Then any thresholded version \hat{u} of the relaxed solution u^* is within a computable bound of the optimum E_{opt} .*

Proof. Since E_{opt} lies energetically in between the minimum of the relaxed problem and the energy of the thresholded version, we have:

$$E(\hat{u}) - E_{opt} \leq E(\hat{u}) - E(u^*).$$

\square

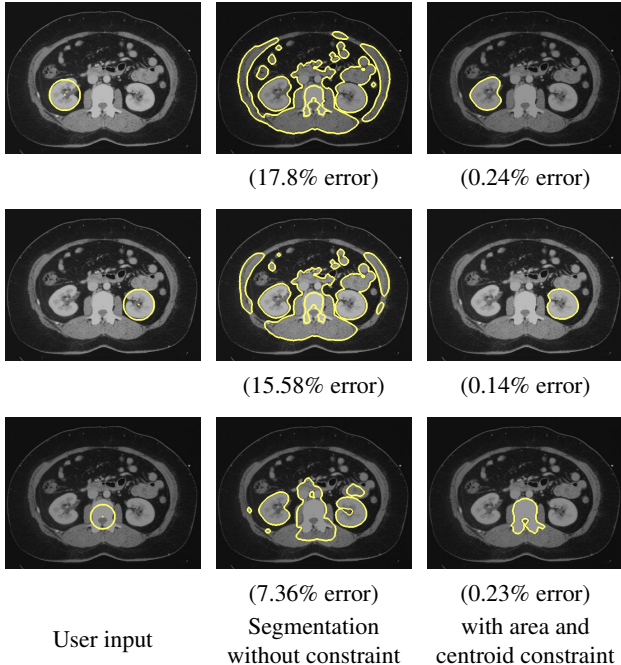


Figure 3. Segmentation of a CT image with kidneys and spine. The centroid and area constraints enable the user to specify the approximate location and size of the desired object that should be segmented. Imposing these moment constraints during optimization leads to drastic improvements in the segmentation.

5. Experimental Results

In this section we present a qualitative and quantitative evaluation of the proposed method on medical imagery and other real-world images and videos. For all experiments we use $g(x) = 1$ and $f(x) = \log p_{bg}(I(x))/p_{obj}(I(x))$ with input image $I : \Omega \rightarrow \mathbb{R}$. Respective moment constraints on centroid, area or covariance structure are easily imposed by simple mouse interactions. Solutions to the constrained convex optimization problems are computed on the fly. Typical run-times on the GPU are around 1 second for an image of the size 300×400 . The average energy bound is around 5% from the optimal solution.

5.1. Quantitative Evaluation on Medical Images

Centroid and Area Constraints Figure 3 shows a comparison of segmentation with and without a constraint on the area and centroid for a CT image of kidneys and spine: without constraints no location information is taken into account for the segmentation, resulting in a segmentation which includes many different regions. Enabling the area and centroid constraints leads to segmentations that prefer the center and the size of the circle that was clicked by the user. This leads to substantial improvements of the segmentations without affecting the fine-scale boundary estimation.

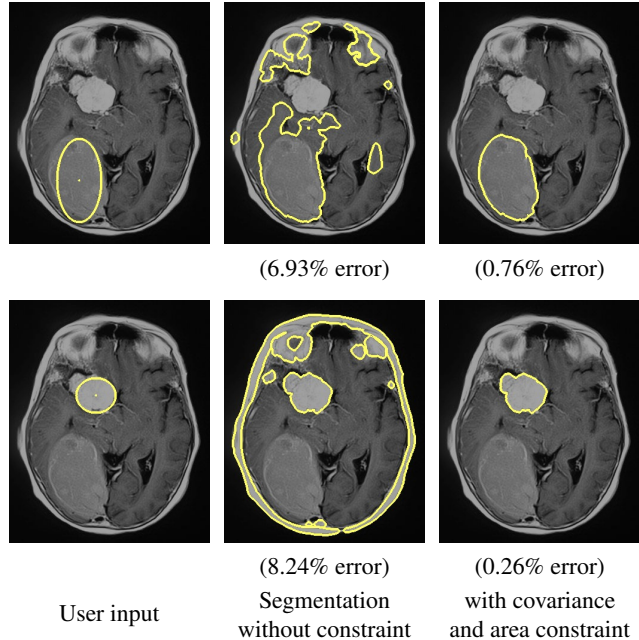


Figure 4. Tumor extraction in brain MR images using segmentation with and without constraints on covariance and area. While the algorithm does not require any local boundary information, constraining its second order moments by a simple user interaction suffices to generate the desired segmentation.

Higher Order Constraints More sophisticated structures can be specified with higher order moments. Since covariance matrices can be represented by ellipsoids, an intuitive user input is achieved by clicking an ellipse with the mouse. The axes of the ellipse define the entries of the corresponding covariance matrix, while the center and area of the ellipse define the centroid and area constraints. Figures 4 and 5 show segmentations with and without constraints resulting from user defined ellipses describing the approximation size, location and shape of the desired object.

5.1.1 Quantitative Performance Evaluation

Clearly, the user-specified moment constraints allow to visibly improve the segmentation. To quantify this improvement, table 1 shows average relative errors (i.e. the percentage of incorrectly labeled pixels per image) with standard deviations for an evaluation of the segmentation without constraint, with area constraint only, and with area and centroid and covariance constraint respectively. Some of the images that were used for the tests and their segmentations are shown in Figures 3, 4 and 5. The table shows that the use of these rather simple and easy to transmit constraints yield a reduction of incorrectly classified pixels by a factor of about 10.

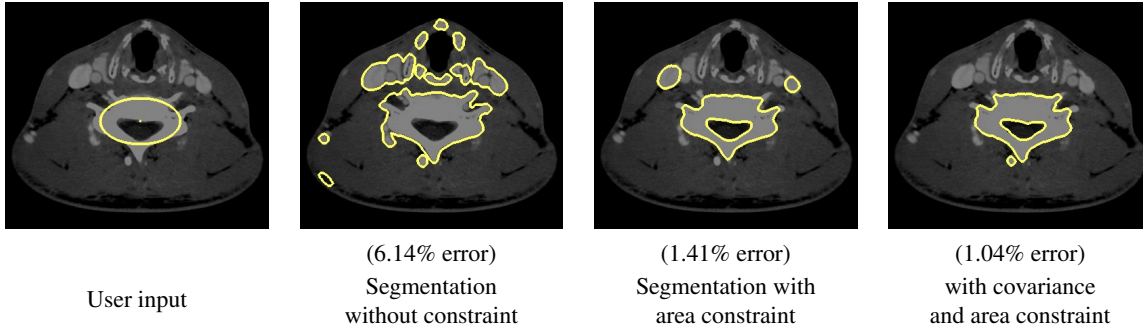


Figure 5. Segmentation without and with constraints for a CT image of the neck. The area constraint yields a segmentation which prefers the size of the ellipse that was clicked by the user, resulting in less incorrectly labeled pixels, compared to the segmentation without constraint. The covariance constraint also considers the dimensions of the ellipse which yields an even more accurate segmentation. Again, the convex constraints merely constrain respective moments of the solution leading to drastic improvements of the segmentation results.

	Avg. relative error
Segmentation without constraint	12.02 % \pm 0.89%
with area constraint	2.36 % \pm 0.11%
with centroid and area constraint	0.41 % \pm 0.05%
with centroid, area and covariance	0.35 % \pm 0.09%

Table 1. Average relative errors with standard deviations for segmentation without and with moment constraints.

5.2. Moment Constraints for Tracking

Figures 6 and 7 show how the proposed method can be applied to tracking objects in videos. As can be seen in Figure 7, the purely color-based segmentation does not suffice to correctly segment object from background in the case of non-unique color distributions.

We impose shape information by constraining the low order moments (area, centroid and covariance) throughout the entire image sequence. As can be seen in the first image of each sequence, the user initializes the method with two mouse clicks: an ellipse of the approximate size and location of the object is drawn on the first frame of the sequence. This is sufficient user input, since histograms and moment constraint parameters are derived from the ellipse: histograms for foreground and background are computed from the inside and outside of the ellipse, respectively, and the constraint parameters for area, centroid and covariance are derived from the ellipse’s area, center point and principal axes. The subsequent frames of the video use the histograms and moment constraints from the first frame, allowing a small deviation of the centroid from each frame to the next, which corresponds to a constraint on the maximum velocity. Since no previous learning of shapes is necessary, the approach naturally applies to arbitrary object shapes.

5.3. Moment Constraints for Real World Images

Figure 8 shows how moment constraints can improve segmentation of real world images. The purely color-based segmentations without moment constraints shown in the second column show that the color distributions of respective objects are not sufficiently different to discriminate the objects of interest. The third column of Figure 8 shows the segmentation results with constraints on area, centroid and covariance. All moment constraints are extracted from the user-specified ellipse. These constraints allow to quickly disambiguate the color information leading substantial improvements of the segmentation.

6. Conclusion

We proposed the use of moment constraints in a convex shape optimization framework. In particular, we showed that for an entire family of constraints on the area, the centroid, the covariance structure of the shape and respective higher-order moments, the feasible constraint sets are all convex. While, we cannot guarantee global optimality of the arising segmentations, all computed solutions are independent of initialization and within a known bound of the optimum. In both qualitative and quantitative experiments on interactive image segmentation, we demonstrated that respective moment constraints are easily imposed by the user and lead to drastic improvements of the segmentation results, reducing the average segmentation error from 12% to 0.35%. In contrast to existing works on shape priors in segmentation the use of low-order moment constraints does not require shape learning and is easily applied to arbitrary shapes since the recovery of fine scale shape details is not affected through the moment constraints. Efficient GPU-accelerated PDE solvers allow for computation times of about one second for images of size 300×400 , making this a practical tool for interactive image segmentation.



Figure 6. Moment constraints for object tracking. The user initializes the tracking by clicking an ellipse in the first frame, the moments of which constrain the segmentation in subsequent images. A small deviation of the centroid is allowed to track the moving object. Note that this approach is generic, as no shapes have to be previously learned.



Figure 7. Moment constraints for object tracking. The user initializes the tracking in form of an ellipse in the first frame (left), from which histograms and constraint parameters are derived. The first row shows results with moment constraints, where a deviation of the centroid is allowed from each frame to the next one to account for the object's motion. The second row shows results of histogram based tracking without constraints. This comparison shows that moment constraints can realize acceptable real-world object tracking with no previous learning of shapes.

References

- [1] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987. 1
- [2] J. P. Boyle and R. L. Dykstra. An method for finding projections onto the intersection of convex sets in Hilbert spaces. *Lecture Notes in Statistics*, 37:28–47, 1986. 4
- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. IEEE Intl. Conf. on Comp. Vis.*, pages 694–699, Boston, USA, 1995. 1
- [4] T. Chan, S. Esedoğlu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics*, 66(5):1632–1648, 2006. 1, 2, 4
- [5] D. Cremers, S. J. Osher, and S. Soatto. Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int. J. of Computer Vision*, 69(3):335–351, 2006. 2
- [6] P. Das, O. Veksler, V. Zavadsky, and Y. Boykov. Semiautomatic segmentation with compact shape prior. *Image and Vision Computing*, 27(1-2):206–219, 2008. 2
- [7] P. Etyngier, F. Segonne, and R. Keriven. Shape priors using manifold learning techniques. In *IEEE Int. Conf. on Computer Vision*, Rio de Janeiro, Oct 2007. 2
- [8] A. Foulonneau, P. Charbonnier, and F. Heitz. Affine-invariant geometric shape priors for region-based active contours. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 28(8):1352–1357, 2006. 2
- [9] D. M. Greig, B. T. Porteous, and A. H. Seheult. Exact maximum *a posteriori* estimation for binary images. *J. Roy. Statist. Soc., Ser. B.*, 51(2):271–279, 1989. 1
- [10] U. Grenander, Y. Chow, and D. M. Keenan. *Hands: A Pattern Theoretic Study of Biological Shapes*. Springer, New York, 1991. 2
- [11] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 23:253–258, 1925. 1
- [12] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. of Computer Vision*, 1(4):321–331, 1988. 1
- [13] S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. J. Yezzi. Gradient flows and geometric active contour models. In *IEEE Int. Conf. on Computer Vision*, pages 810–815, 1995. 1
- [14] K. Kolev and D. Cremers. Integration of multiview stereo and silhouettes via convex functionals on convex domains. In *European Conference on Computer Vision (ECCV)*, Marseille, France, October 2008. 2
- [15] K. Kolev, M. Klodt, T. Brox, and D. Cremers. Continuous global optimization in multiview 3d reconstruction. *International Journal of Computer Vision*, 2009. 4



(a) User input

(b) No Constraints

(c) Moment Constraints

Figure 8. Segmentation of real world images without and with moment constraints. (a) The user marks an ellipse at the approximate size and location of the object with two mouse clicks. Color histograms for foreground and background are determined from the inside and outside of the ellipse, respectively. (b) Segmentation results without constraints using only the histograms. (c) Segmentation results with constraints on area, centroid and covariance. Parameters of the constraints are derived from the ellipse marked by the user.

- [16] V. Lempitsky, P. Kohli, C. Rother, and T. Sharp. Image segmentation with a bounding box prior. In *IEEE Int. Conf. on Computer Vision*, Kyoto, Japan, 2009. 2
- [17] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42:577–685, 1989. 1
- [18] S. J. Osher and J. A. Sethian. Fronts propagation with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations. *J. of Comp. Phys.*, 79:12–49, 1988. 1
- [19] A. Papoulis and S. U. Pillai. *Probability, Random Variables,*

and Stochastic Processes. McGraw-Hill, New York, 4th edition edition, 2002. 4

- [20] T. Schoenemann and D. Cremers. A combinatorial solution for model-based image segmentation and real-time tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2009. 2
- [21] O. Veksler. Star shape prior for graph-cut image segmentation. In *Europ. Conf. on Computer Vision*, pages 454–467, 2008. 2