

Curvature Prior for MRF-Based Segmentation and Shape Inpainting^{*}

Alexander Shekhovtsov¹, Pushmeet Kohli², and Carsten Rother²

¹ Center for Machine Perception, Czech Technical University

² Microsoft Research Cambridge, UK

Abstract. Most image labeling problems such as segmentation and image reconstruction are fundamentally ill-posed and suffer from ambiguities and noise. Higher-order image priors encode high-level structural dependencies between pixels and are key to overcoming these problems. However, in general these priors lead to computationally intractable models. This paper addresses the problem of discovering compact representations of higher-order priors which allow efficient inference. We propose a framework for solving this problem that uses a recently proposed representation of higher-order functions which are encoded as lower envelopes of linear functions. Maximum a Posterior inference on our learned models reduces to minimizing a pairwise function of discrete variables. We show that our framework can learn a compact representation that approximates a low curvature shape prior and demonstrate its effectiveness in solving shape inpainting and image segmentation problems.

1 Introduction

A number of models encoding prior knowledge about scenes have been proposed in computer vision. The most popular ones have been in the form of a Markov Random Field (MRF). An important characteristic of an MRF is the factorization of the distribution into a product of factors. *Pairwise* MRFs can be written as a product of factors defined over two variables at a time. For discrete variables, this enables non-parametric representation of factors and the use of efficient optimization algorithms for approximate inference of the Maximum-a-Posteriori (MAP) solution. However, because of their restricted pairwise form, the model is not able to encode many types of powerful structural properties of images. Curvature is one such property which is known to be extremely helpful for inpainting (see figure 1), segmentation, and many other related problems.

Higher-order Priors. There has been a lot of research into priors based on high-level structural dependencies between pixels such as curvature. These priors can be represented in the probabilistic model using factors which may depend on more than two variables at a time. The largest number of variables in a factor is called the *order* of the model. Higher-order factors defined on discrete variables are computationally expensive to represent. In fact, the memory and time complexity for inferring the MAP solution with general inference algorithms grows exponentially with

^{*} A. Shekhovtsov was supported by EU project FP7-ICT-247870 NIFTi.

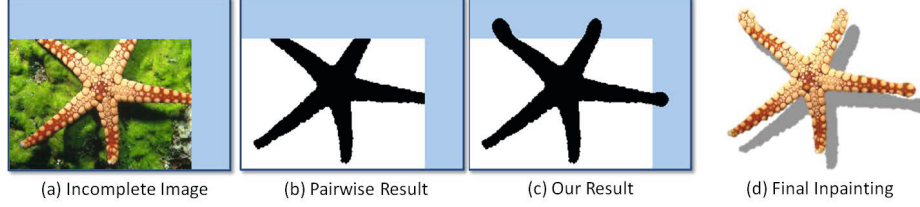


Fig. 1. (a) Input image (area for completion of starfish is shown in blue). (b) The starfish was interactively segmented from the image. Then the three arms of the starfish, which touch the image borders, were completed with an 8-connected pairwise MRF which encodes a standard length prior. (c) Completion of the shape with our higher-order curvature prior. (d) Finally, texture was added fully automatically using [2].

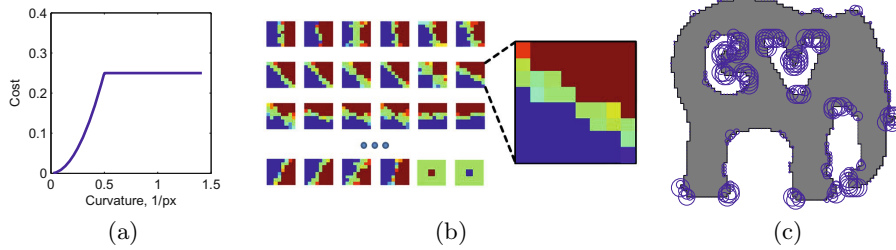


Fig. 2. (a) A given cost function for curvature. (b) “Soft” patterns whose lower envelope approximates the curvature cost of a binary labeling. A pattern fits well (has low cost) if all fore-/background pixels match to blue/red pattern weights, where green pixels can be assigned to both fore- and background. This lower envelope corresponds to a higher-order factor covering 8×8 pixels in every pixel-location in the model. The last two patterns are selected manually such that for interior and exterior pixels the value of the lower envelope is always 0. (c) An example demonstrating the curvature cost computed by our pattern-based approximation at different parts of the object boundary. Circle radius correspond to the assigned cost.

the order, and thus has limited the use of such models. The situation is a bit different for parametric models with continuous variables. Higher-order prior models such as Product of Experts (PoE) [11] or Field of Experts (FoE) [18] are differentiable in both parameters and hidden variables. These models thus enable inference using local gradient descent, and have led to impressive results for problems such as image restoration and optical flow. Recent research on discrete higher-order models has focused on identifying families of higher-order factors which allow efficient inference. The factors can be categorized into 3 broad categories: (a) *Reducable factors*, which allow MAP inference to be reduced to the problem of minimizing a pairwise energy function of discrete variables with the addition of some auxiliary variables [12–14, 17, 19], (b) *Message-enabled factors*, which allow efficient message computation and thus allow inference using message passing methods such as Belief Propagation (BP) and Tree Reweighted message passing (TRW) [10, 15, 25], and (c) *Constraint factors*, which impose global constraints that can be imposed efficiently in a relaxation framework [16, 26].

Pattern-based Representation. Pattern and lower-envelope based representations proposed in [12, 15, 19] can represent some families of *Reducable factors*. The higher-order potentials of [15, 19] are defined by enumerating important configurations (patterns) in a local window. The model of [19] additionally enables deviations from encoded patterns, by using linear weighting functions. The above models are generalized by the representation proposed in [12] which encodes higher-order functions as lower (or upper) envelopes of linear (modular) functions of the variables. The complexity of representing and performing inference depends on the number of linear functions (or patterns) used for representing the higher-order factor. A number of higher-order priors can be encoded using few linear functions (or patterns) and thus allow efficient inference. However, the use of a general higher-order prior would require exponential (in the order of the factor) number of linear functions (or patterns).

Our Contribution. This paper addresses the problem of discovering a compact representation of higher-order factors that encode a curvature prior for labelling problems. Given a set of training examples of labeling and their corresponding desired curvature-based costs, we find parameters of a linear-envelope representation that matches these costs. While the problem is difficult, we propose a simple yet effective algorithm for parameter learning. Figure 2 illustrates our discovered model. We applied the learned prior model to the problems of object segmentation and completion. The experimental results demonstrate that incorporation of this prior leads to much better results than those obtained using low-order (pairwise MRF) based models (see figure 1) and is comparable to other state-of-the-art curvature formulations.

Related work on Curvature. In this work we consider two closely related problems of shape inpainting and image segmentation with curvature prior. Given an image region with a lack of observations, the goal of *shape inpainting* is to complete the region from evidence outside of the region. This problem is related to inpainting of binary images which has been approached in the continuous setting with several curvature-related functionals [3, 6]¹. Image labeling with curvature regularization is an important topic of research, and both continuous and discrete formulations for the problem have been studied. Continuous formulations offer accurate models, however until recently, only local optimization methods were applied. For instance, [8] works with discretized Euler-Lagrange equations of the 4th order. Local optima found by such methods may be of poor quality and several methods solving convex relaxations have already been proposed [5, 9]. Discrete methods for image labeling with curvature regularization build on quantization and enumeration of boundary elements. Until recently, they were applied only in restricted scenarios where it is possible to reduce the problem to a search of the minimal path or minimum ratio cycle [23]. These cases enjoy global optimality, however they do not allow for arbitrary regional terms

¹ There is a vast literature on the general image inpainting problem, however these techniques, especially exemplar-based ones, do not extend to image segmentation problem, and are not relevant in the context of this paper.

or impose severe constraints on possible shapes. A series of recent works [20–22, 24] developed a global minimization method for the general discrete setting, where the regions, boundary segments and pairs of adjacent boundary segments can have arbitrary associated costs. The problem is formulated as an ILP and approached by a linear relaxation. It was shown that some image segmentation problems with (approximate) curvature regularization can be solved optimally in this model. For accurate approximation of the curvature, the discretization of the space must form a fine cell complex. Quantization of directions leads to visible artifacts of the segmentation (see section 4). Cell complexes with a finer quantization of directions and adaptive complexes are studied in [24]. A recent work [7] claims to give fast optimal solution for curvature regularization. However, their model is a crude approximation to the curvature functional. Its 4-neighborhood variant essentially penalizes the number of “corners” in the segmentation.

2 Higher-Order Model Representation and Optimization

We consider a set of pixels $\mathcal{V} = \{1 \dots N_X\} \times \{1 \dots N_Y\}$ and a binary set of labels $\mathcal{L} = \{0, 1\}$, where 1 means that a pixel belongs to the foreground (shape) and 0 to the background. Let $\mathbf{x}: \mathcal{V} \rightarrow \mathcal{L}$ be the labeling for all pixels with individual components denoted by x_v , $v \in \mathcal{V}$. Furthermore, let $V(h) \subset \mathcal{V}$ denote a square window of size $K \times K$ at location h , and \mathcal{U} is the set of all window locations. Windows are located densely in all pixels. More precisely, all possible $K \times K$ windows are considered which are fully inside the 2D-grid \mathcal{V} (see fig. 3 for illustration). Let $\mathbf{x}_{V(h)}: V(h) \rightarrow \mathcal{L}$ denote a restriction of labeling \mathbf{x} to the subset $V(h)$. We consider distribution of the form $p(\mathbf{x}) \propto \exp\{-E(\mathbf{x})\}$ with the following energy function:

$$E(\mathbf{x}) = \sum_{v \in \mathcal{V}} \theta_v(x_v) + \sum_{uv \in \mathcal{E}} \theta_{uv}(x_u, x_v) + \sum_{h \in \mathcal{U}} E_h(\mathbf{x}), \quad (1)$$

where notation uv stands for ordered pair (u, v) , $\theta_v: \mathcal{L} \rightarrow \mathbb{R}$ and $\theta_{uv}: \mathcal{L}^2 \rightarrow \mathbb{R}$ are unary and pairwise terms, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is a set of pairwise terms and E_h are higher-order terms. We consider the higher order terms E_h of the following form (equivalent to [19])

$$E_h(\mathbf{x}) = \min_{y \in P} \left(\langle \mathbf{w}_y, \mathbf{x}_{V(h)} \rangle + c_y \right). \quad (2)$$

This term is the minimum (lower envelope) of several modular functions of $\mathbf{x}_{V(h)}$ ². We refer to individual linear functions $\langle \mathbf{w}_y, \mathbf{x}_{V(h)} \rangle + c_y$ as “soft” patterns. Here $\mathbf{w}_y \in \mathbb{R}^{K^2}$ is a weight vector and $c_y \in \mathbb{R}$ is a constant term for the pattern. Vector \mathbf{w}_y is of the same size as the labeling patch $\mathbf{x}_{V(h)}$ and it can be visualized as an image (see fig. 2(b)). The variable $y \in P$ is called a pattern-switching variable. It is a discrete variable from the set $P = \{0, \dots, N_P\}$. We let the pattern which corresponds to $y = 0$ have the associated weights $\mathbf{w}_0 = \mathbf{0}$.

² This model has some similarities with a mixture model, as discussed in [1].

This pattern assigns a constant value c_0 to all labellings $\mathbf{x}_{V(h)}$ and it ensures that $E_h(\mathbf{x}) \leq c_0$ for all \mathbf{x} . (See [1] for detailed relation of this model to [19] and [15]). In our model it is used to represents the maximal cost f^{\max} of the curvature cost function. The minimization problem of energy (1) expresses as

$$\min_{\mathbf{x}} \left[E_0(\mathbf{x}) + \sum_{h \in \mathcal{U}} \min_{y \in P} \left(\langle \mathbf{w}_y, \mathbf{x}_{V(h)} \rangle + c_y \right) \right], \quad (3)$$

where unary and pairwise terms are collected into E_0 . The problem can also be written as a minimization of a pairwise energy

$$\min_{\substack{\mathbf{x} \in \mathcal{L}^{\mathcal{V}} \\ \mathbf{y} \in P^{\mathcal{U}}}} \left[E_0(\mathbf{x}) + \sum_h \langle \mathbf{w}_{y_h}, \mathbf{x}_{V(h)} \rangle + c_{y_h} \right], \quad (4)$$

where $\mathbf{y} : \mathcal{U} \rightarrow P$ is the concatenated vector of all pattern switching variables³. Clearly, problem (4) is a minimization of a pairwise energy function of discrete variables \mathbf{x}, \mathbf{y} . The problem is NP-hard in general, however, a number of approximate MAP inference techniques for pairwise energies can be used such as Block-ICM, TRW, BP, or Linear programming based relaxations. Here we report results obtained by the memory-efficient adaptation of TRW-S with post-processing by block-ICM (see details in [1]).

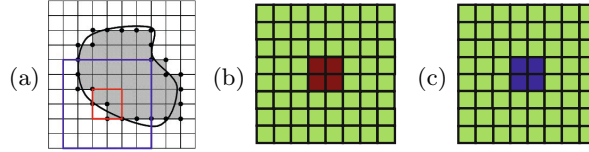


Fig. 3. (a) Continuous shape and its discretization. Filled, small circles show boundary locations. The large blue window illustrates $V(h)$ at location h . (b,c) Fore- and background patterns, which are active at none-boundary locations, with costs: green $w_{y,v} = 0$, red: $w_{y,v} = +B$, blue: $w_{y,v} = -B$, constant c_y is $-4B$ and $+4B$ respectively.

3 Learning a Curvature Cost Model

Suppose we are given a shape $S \subset \mathbb{R}^2$ such that we can calculate the curvature κ at every point of the boundary, ∂S . Let $f(\kappa) \geq 0$ be a curvature cost function, which defines a desired penalty on curvature, in this paper we consider $f(\kappa) = \min(\kappa^2, f^{\max})$. Let the total cost of the shape be $\int_{\partial S} f(\kappa) dl$. Our goal is to approximate this integral by the sum $\sum_{h \in \mathcal{U}} E_h(\mathbf{x})$, where functions E_h operate over a discretized representation of the shape, \mathbf{x} , and are of the form (2) with

³ We refer to components of \mathbf{y} by y_h , while y usually denotes an independent bound variable.

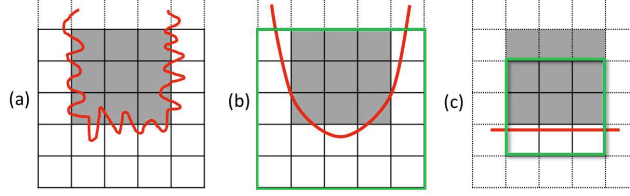


Fig. 4. Problem definition and motivation of large-sized windows. Examples above show discrete labelings on a pixel grid with a corresponding red continuous curve. Note, there are infinitely many continuous curves which give rise to the same discrete labelling - two examples are given in (a) and (b). The red curve in (b) is probably the one with lowest curvature given the discrete labelling. Our goal is to find an energy function which maps every discrete labelling to the corresponding cost of the continuous curvature with *lowest* curvature. (c) makes the important point that larger sized windows have inherently a better chance of predicting well the curvature at the center of the window. In (c) the green window is of size 3x3, while in (b) it is of size 5x5. The underlying discrete labelling is identical in both cases and the red curve is the optimal (lowest curvature) continuous curve given the window. The crucial point is that the curvature of the continuous curve, at the center of the window, is very different in (b) and (c). Note, this problem is to some extent mitigated by the fact that the total cost of segmentation is the sum of costs along the boundary.

weights \mathbf{w}, \mathbf{c} . Here \mathbf{w} and \mathbf{c} denote the concatenated vectors of all weights \mathbf{w}_y and \mathbf{c}_y , respectively. The learning problem is to determine the pattern weights \mathbf{w}, \mathbf{c} such that the approximation is most accurate. Since the mapping of continuous to discrete curves is a many-to-one mapping, we further formalize our exact goal in figure 4. In the figure we also motivate the important aspect that larger windows are potentially superior.

We first restrict the sum $\sum_{h \in \mathcal{U}} E_h(\mathbf{x})$ to take into account only boundary locations. We call h a *boundary location* for shape \mathbf{x} if the 2×2 window at h contains some pixels which are labeled foreground as well as some pixels which are labeled background, as illustrated in fig. 3. We constrain all soft patterns to be non-negative ($\langle w_y, x \rangle + c_y \geq 0$) and introduce two special patterns (fig. 3b,c), which have cost 0 for locations where the 2×2 window at location h contains only background or foreground pixels. These patterns make $E_h(\mathbf{x})$ vanish over all non-boundary locations, therefore such locations do not contribute to the sum. The learning task is now to determine $E_h(\mathbf{x})$, such that at each boundary location the true cost $f(\kappa)$ is approximated. In this way the discrete sum corresponds to the desired integral if we were to neglect the fact that the number of boundary locations does only approximate the true length of the boundary.

Point-wise learning procedure. Let us assume that in a local $K \times K$ window, shapes of low curvature can be well-approximated by simple quadratic curves⁴. The idea is to take many examples of such shapes and fit $E_h(\mathbf{x})$ to

⁴ Note, based on our definition in fig. 4 we select curves which are likely to be the ones of lowest curvature (among all curves) for the corresponding discrete labelling.

approximate their cost. We consider many quadratic shapes $(S^i)_{i=1}^N$ in the window $K \times K$ and derive their corresponding discretization on the pixel grid $(\mathbf{x}^i)_{i=1}^N$. Each continuous shape has an associated curvature cost $f^i = f(\kappa^i)$ at the central boundary location. We formulate the learning problem as minimization of the average approximation error:

$$\arg \min_{\mathbf{w}, \mathbf{c}} \sum_i |E_h(\mathbf{x}^i) - f^i|, \quad \text{s.t.} \quad \begin{cases} \mathbf{w}_0 = 0, c_0 = f^{\max}; \\ E_h(x) \geq 0 \quad \forall x, \end{cases} \quad (5)$$

where the first constraint represents the special implicit pattern ($w_0, c_0 = f^{\max}$), which ensures that $E_h(\mathbf{x}) \leq f^{\max}$. The second constraint makes sure that cost is non-negative. It is important for the following reason: the formulation of the approximation problem does not explicitly take into account “negative samples”, *i.e.* labellings which do not originate from smooth curves, and which must have high cost in the model. However, requiring that all possible negative samples in a $K \times K$ window have high cost would make the problem too constrained. The introduced non-negativity constraint is tractable and not too restrictive. This problem appears difficult, since $E_h(\mathbf{x}^i)$ is itself a concave function in the parameters \mathbf{w}, \mathbf{c} . We approach (5) by a k-means like procedure:

Algorithm 1. Iterative Factor Discovery

```

1 repeat /* iteration */
2   for  $i = 1 \dots N$  do
3      $y^i = \arg \min_y [\langle \mathbf{w}_y, \mathbf{x}^i \rangle + c_y]$ ; /* find matching patterns */
4   for  $y \in 1 \dots N_P$  do /* refit patterns */
5      $(\mathbf{w}_y, c_y) = \arg \min_{\xi} \sum_{i|y^i=y} |\langle \mathbf{w}_y, \mathbf{x}^i \rangle + c_y - f^i|$ , s.t.  $\begin{cases} \xi_v \leq w_{y,v}; \\ \xi_v \leq 0; \\ \sum_v \xi_v + c_y \geq 0. \end{cases}$ 
6 until convergence or maximum iterations;
```

The refitting step 5 is a linear optimization which can be solved exactly. The constraint in step 5 is an equivalent representation of the constraint $\langle \mathbf{w}_y, \mathbf{x} \rangle + c_y \geq 0 \forall x$, imposed by (5).

4 Experiments

We applied Algorithm 1 to learn a prior model with 96 patterns of size 8×8 pixels from 10000 randomly, synthetically generated smooth curves and their discretization. We initialize the curvature potential in the learning process by clustering the 10000 patches in 32 groups based on the orientation of the boundary at the patch center. Then each orientation group is further subdivided into 3 bins based on the curvature. To measure the accuracy of our curvature potential approximation, we sampled large shapes for which the true curvature cost can be

computed and then compared it with our approximated cost which is obtained by summing the response of our curvature cost potential along the boundary. Further details of the learning procedure are given in [1].

Shape Inpainting. We now demonstrate the learned prior model on the problem of shape inpainting. The goal is to reconstruct the full shape, while only some parts of the shape are visible to the algorithm. This is a useful test to inspect our shape prior. Let $F \subset \mathcal{V}$ be the set of pixels restricted to foreground (shape) and $B \subset \mathcal{V}$ pixels restricted to background. The unary terms of (1), $\theta_v(x_v)$, are set to ∞ if label x_v contradicts with the constraints and 0 otherwise. This ensures that the correct segmentation is inferred in the region $F \cup B$.

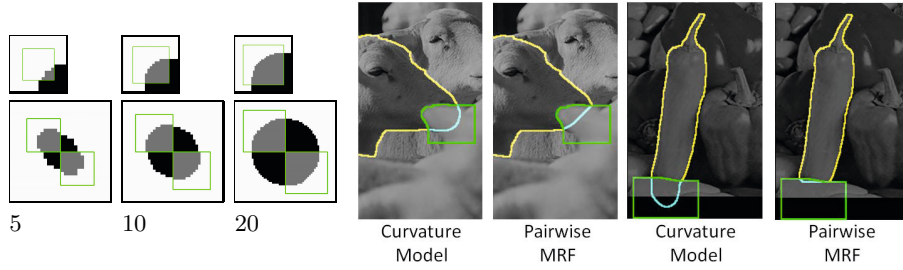


Fig. 5. Left. Inpainting of a corner and a circle. The green boxes show the area to be inpainted, where the size in pixels of the length of the green boxes is below the images. Pixels in gray show the estimated solution. Note, the boundary conditions are different: right-angle boundary condition (top) and circle boundary condition (bottom). **Right.** Two example for automatic shape completions of an occluded object. In both cases the left result is with a pure curvature prior and the right result with a pure length prior (8-connected). Note, the yellow curve (and a part of the green curve) indicate the original user-defined segmentation. Then the user defines the green area. Inside the green area, the method automatically finds the shape completion (blue curve).

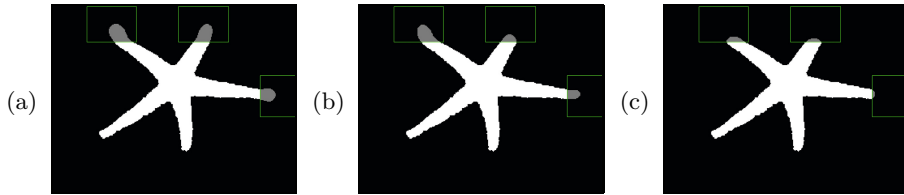


Fig. 6. Combining length and curvature for inpainting: (a) pure curvature, (b) curvature and length, (c) curvature and length (with high weight).

In the unknown region $\mathcal{V} \setminus (F \cup B)$ all unaries are exactly 0. Fig.5(Left) shows examples of inpainting of corners and circles of varying size. Fig.5(Right) demonstrates inpainting with real-world shapes and compares against a naive length regularization. It can be seen that the higher-order model which encodes curvature produces shape completions with smooth boundaries. An example of combining curvature prior with length prior is shown in fig. 6.

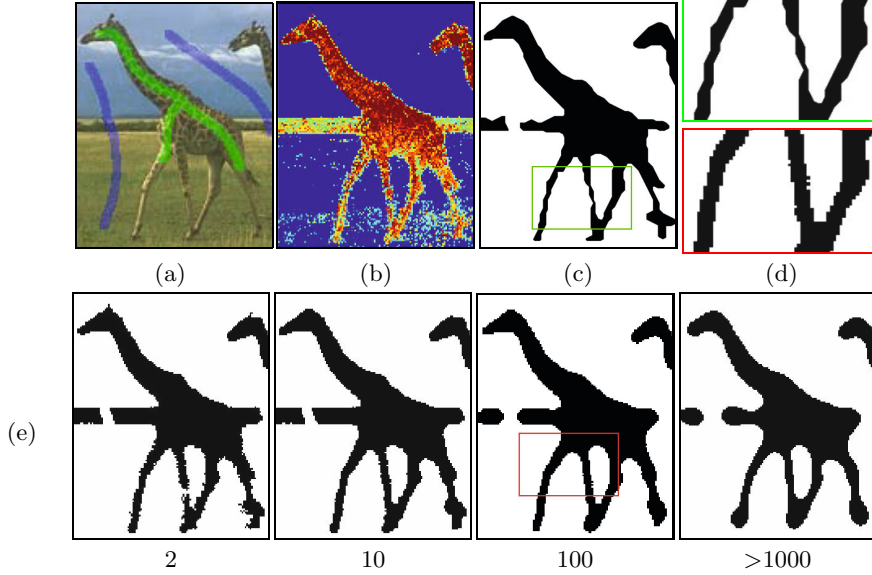


Fig. 7. Image segmentation. (a) Image with foreground (green) and background (blue) seeds; (b) Color based unary potential costs (red for foreground, and blue for background). (c) Result from [20] (d) Zoom-in from [20] (top) and our result (e,100). (e) Our model with various strength for curvature prior.

Image Segmentation. We use a simple model for the task of interactive fore- and background segmentation, as in [4]. Based on the user brush strokes (fig. 7(a)) we compute likelihoods using a Gaussian mixture model (GMM) with 10 components. The difference of the unaries $\theta_v(1) - \theta_v(0)$ correspond to the negative log-likelihood ratio of foreground and background. Results for our curvature model for various strengths of the prior are shown in fig.7(e). Increasing the strength of the prior above some limit (1000) has almost no effect on the smoothness of the solution, because each local 8×8 window is already maximally smooth according to the model. Note, that segmentation of this instance with length regularization cannot segment the legs of the giraffe correctly for arbitrary regularization strength (see [1]). Our result is visually superior to [20], see fig. 7(d), despite the fact that we use a grid with much coarser resolution than a fine cell-complex used in [20]. Further results and a detailed comparison to [20] is in [1].

5 Conclusions and Discussion

This paper has shown how to compute compact representations of higher-order priors which enable the use of standard algorithms for MAP inference. We presented results on the problem of learning a curvature-based shape prior for image inpainting and segmentation. Our higher-order shape prior operates on a large set of pixels and is therefore robust to discretization artifacts. In the future, it

would be interesting to extend the approach to incorporate other types of local shape properties, not necessarily defined by an analytic function but for instance by exemplars. Such a generalization would likely require a more general learning technique.

References

1. Shekhovtsov, A., Kohli, P., Rother, C.: Curvature prior for MRF-based segmentation and shape inpainting. Tech. rep., research Report CTU-CMP-2011-11, Czech Technical University (2011)
2. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* 28(3) (2009)
3. Bertozzi, A., Esedoglu, S., Gillette, A.: Inpainting of binary images using the Cahn-Hilliard equation. *IP* 16 (2007)
4. Boykov, Y., Jolly, P.: Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images. In: *ICCV* (2001)
5. Bredies, K., Pock, T., Wirth, B.: Convex relaxation of a class of vertex penalizing functionals (preprint)
6. Chan, T.F., Shen, J.: Non-texture inpainting by curvature-driven diffusions (CDD). *JVCIR* 12 (2001)
7. El-Zehiry, N.Y., Grady, L.: Fast global optimization of curvature. In: *CVPR* (2010)
8. Esedoglu, S., March, R.: Segmentation with depth but without detecting junctions. *JMIV* 18 (2003)
9. Goldluecke, B., Cremers, D.: Introducing total curvature for image processing. In: *ICCV* (2011)
10. Gupta, R., Diwan, A.A., Sarawagi, S.: Efficient inference with cardinality-based clique potentials. In: *ICML* (2007)
11. Hinton, G.E.: Products of experts. In: *ICANN* (1999)
12. Kohli, P., Kumar, M.: Energy minimization for linear envelope MRFs. In: *CVPR* (2010)
13. Kohli, P., Ladicky, L., Torr, P.: Robust higher order potentials for enforcing label consistency. *IJCV* 82 (2009)
14. Kolmogorov, V., Zabini, R.: What energy functions can be minimized via graph cuts? *PAMI* 26 (2004)
15. Komodakis, N., Paragios, N.: Beyond pairwise energies: Efficient optimization for higher-order MRFs. In: *CVPR* (2009)
16. Nowozin, S., Lampert, C.: Global connectivity potentials for random field models. In: *CVPR* (2009)
17. Ramalingam, S., Kohli, P., Alahari, K., Torr, P.: Exact inference in multi-label CRFs with higher order cliques. In: *CVPR* (2008)
18. Roth, S., Black, M.: Fields of experts. *IJCV* 82 (2009)
19. Rother, C., Kohli, P., Feng, W., Jia, J.: Minimizing sparse higher order energy functions of discrete variables. In: *CVPR* (2009)
20. Schoenemann, T., Kahl, F., Cremers, D.: Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In: *ICCV* (2009)
21. Schoenemann, T., Kahl, F., Masnou, S., Cremers, D.: A linear framework for region-based image segmentation and inpainting involving curvature penalization. *CoRR* abs/1102.3830 (2011)

22. Schoenemann, T., Kuang, Y., Kahl, F.: Curvature Regularity for Multi-label Problems - Standard and Customized Linear Programming. In: Boykov, Y., Kahl, F., Lempitsky, V., Schmidt, F.R. (eds.) EMMCVPR 2011. LNCS, vol. 6819, pp. 163–176. Springer, Heidelberg (2011)
23. Schoenemann, T., Masnou, S., Cremers, D.: The elastic ratio: Introducing curvature into ratio-based image segmentation. *IEEE Transactions on Image Processing* 20(9), 2565–2581 (2011)
24. Strandmark, P., Kahl, F.: Curvature Regularization for Curves and Surfaces in a Global Optimization Framework. In: Boykov, Y., Kahl, F., Lempitsky, V., Schmidt, F.R. (eds.) EMMCVPR 2011. LNCS, vol. 6819, pp. 205–218. Springer, Heidelberg (2011)
25. Tarlow, D., Zemel, R., Frey, B.: Flexible priors for exemplar-based clustering. In: *UAI* (2008)
26. Vicente, S., Kolmogorov, V., Rother, C.: Joint optimization of segmentation and appearance models. In: *ICCV* (2009)