Double Expansion for Multilabel Segmentation with Shape Prior

Lena Gorelick[1]
www.csd.uwo.ca/~ygorelic/
Zhengqin Li[2]
sites.google.com/a/eng.ucsd.edu/zhengqinli/
Olga Veksler[1]
www.csd.uwo.ca/~olga/

[1] University of Western Ontario
London ON, Canada
[2] University of California San Diego
San-Diego, CA, USA

Abstract

We propose a new class of energies for segmentation of multiple foreground objects with a common shape prior. Our energy involves infinity constraints. For such energies standard expansion algorithm has no optimality guarantees and in practice gets stuck in bad local minima. Therefore, we develop a new move making algorithm, we call *double expansion*. In contrast to expansion, the new move allows each pixel to choose a label from a *pair* of new labels or keep the old label. This results in an algorithm with optimality guarantees and robust performance in practice. We experiment with several types of shape prior such as star-shape, compactness and a novel symmetry prior, and empirically demonstrate the advantage of the double expansion.

1 Introduction

Image segmentation is one of the fundamental problems in computer vision. It is commonly addressed using energy minimization and regularization is often used to restrict the search space of possible solutions. Length [1, 9, 14], curvature [16] and shape regularization have been used to rule out solutions inconsistent with the assumed prior. Some notable shape priors are connectivity [17, 24], star-shape [11, 22], convexity [10, 20], part-based [8, 15, 25]. Most are traditionally used in the context of binary segmentation, but recently shape prior has also been incorporated in multi-object segmentation framework [6, 12, 13, 18].

We propose a new class of pairwise multilabel energies for simultaneous segmentation of multiple foreground objects with shape priors. We assume that objects do not occlude each other, since applying a shape prior to an occluded part is not well-defined. Using pairwise potentials only, we cannot distinguish between overlapping and touching pairs of objects. Our model therefore prohibits object touching as well and enforces at least one-pixel-wide background strip around each object. In practice, when objects do touch in the image, they will be separated by a thin background strip without significantly affecting the segmentation.

Prohibiting occlusion involves infinity constraints. For such energies, standard expansion algorithm has no optimality guarantees and can get stuck in bad local minima. We propose *double expansion*, a new move making algorithm. In contrast to α-expansion where each pixel chooses between α and the old label, double α-expansion allows each pixel to choose between three labels: α, background and the old label. The addition of the background label allows simultaneously expanding on α and correcting the shapes of other objects. Expanding on α without this correction may violate the shape of other objects yielding expansion infeasible. Expansion is a special case of double expansion, and, therefore, double expansion is guaranteed to be at least as effective as expansion. Moreover, double expansion has theoretical optimality guarantees and avoids poor local minima in practice. In [15, 23] there are moves that are similar to double expansion. However, in [23] moves are not submodular, with no local minimum guarantee. Moves in [15] can be applied only to a subset of pixels.

Also related to our work are the methods in [12, 13, 18]. In [12] they formulate a multi-object segmentation energy with a "hedgehog" shape prior, but optimize with standard expansion. In [13] they propose more general path moves, but have no optimality guarantees.

We experiment with the star shape [11, 22] and compact shape [6] priors and propose a novel symmetry shape prior. Symmetry is an important cue in human visual system for detecting objects [21]. For symmetry prior, we assume that objects have approximate bilateral reflection symmetry. We empirically compare double expansion to the standard expansion for all three shape priors. Our results demonstrate that double expansion avoids poor local minima and is very robust with respect to expansion label order.

2 Energy

We now describe the multilabel segmentation energy. Let $\mathcal{P} \subset \Omega$ be a set of image pixels and $\mathcal{L} = \{0, 1, \ldots, m\}$ be the set of labels, with 0 denoting the background and the rest denoting the foreground objects. Let $x_p \in \mathcal{L}$ be a multilabel variable corresponding to pixel p, and $\mathbf{x} = (x_p | p \in \mathcal{P})$ be a vector of all variables. Then the energy is defined as follows:

$$f(\mathbf{x}) = f^a(\mathbf{x}) + f^r(\mathbf{x}) + f^s(\mathbf{x}), \tag{1}$$

where $f^a(\mathbf{x})$ is the appearance term, $f^r(\mathbf{x})$ is the boundary regularization term and $f^s(\mathbf{x})$ is the shape prior term. The appearance term is a unary term given by

$$f^a(\mathbf{x}) = \sum_{p \in \mathcal{P}} u_p(x_p) = \sum_{p \in \mathcal{P}} -\log Pr(c_p | \theta^{x_p}), \tag{2}$$

where c_p is the color of pixel p and $\theta^l, l \in \mathcal{L}$ is the appearance model for label l.

The boundary regularization term $f^r(\mathbf{x})$ enforces background pixels around each foreground label. Let $\mathcal{N}$ be a standard 4- or 8- neighborhood system of ordered pairs, then

$$\begin{aligned} f^r(\mathbf{x}) &= \sum_{(p,q) \in \mathcal{N}} v_{pq}(x_p, x_q), \quad \text{where} \\ v_{pq}(x_p, x_q) &= \begin{cases} 0, & \text{if } x_p = x_q \\ w_{pq}, & \text{if } (x_p = 0 \vee x_q = 0) \wedge (x_p \neq x_q) \\ \infty & \text{otherwise.} \end{cases} \end{aligned} \tag{3}$$

Here $w_{pq} > 0$ denotes either edge sensitive or uniform length term.

We now turn to shape constraints. Let $\mathcal{S}^i$ be a set of ordered pixel pairs that are involved in shape constraints for object i. Then the shape prior $f^{s,i}(\mathbf{x})$ for object i is given by

$$f^{s,i}(\mathbf{x}) = \sum_{(p,q) \in \mathcal{S}_i} h^i_{pq}(x_p, x_q) \tag{4}$$

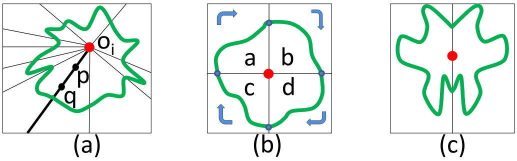

Figure 1: Illustration of (a) star, (b) compact and (c) symmetry shape priors for an object.

and $f^s(\mathbf{x})$ integrates these constraints for all objects $f^s(\mathbf{x}) = \sum_i f^{s,i}$. Depending on the shape prior, set $\mathcal{S}^i$ is determined either with respect to given object centers (for compact and star shape priors) or axes of symmetry. Potentials h^i_{pq} assign a penalty to configurations that violate the shape prior constraints, see details in Sec. 2.1,2.2,2.3.

2.1 Star Shape Prior

The star shape prior in [22] assumes that centers o_i for each object $0 < i \leq m$ are given. Fig. 1(a) illustrates the star constraints for object i with red pixel denoting the center o_i. Consider a pixel q inside object i, and the line segment connecting pixel q to center o_i. If object i is a star, any pixel p on this line segment must belong to object i. One star shape is outlined in green. To enforce the star prior, it is sufficient to place the constraints only between immediately adjacent pixel pairs along all lines emanating from the center o_i [22].

Let $\mathcal{S}_i$ be a set of all immediately adjacent ordered pixel pairs (p,q) along all discrete lines emanating from center o_i. We assume that pixel p is closer to the center than pixel q, that is $\|o_i - p\| \leq \|o_i - q\|$. The star-shape prior $f^{s,i}(\mathbf{x})$ for object i is then given by

$$f^{s,i}(\mathbf{x}) = \sum_{(p,q)\in\mathcal{S}^i} h^i_{pq}(x_p,x_q), \quad \text{where} \quad h^i_{pq}(x_p,x_q) = \begin{cases} K & \text{if } x_p \neq i \wedge x_q = i, \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

The larger is K, the stricter is the star shape prior. In practice, for each object $0 < i \leq m$, we consider all discrete lines connecting o_i and any pixel q on the image border.

2.2 Compact Prior

The compact shape prior in [4] assumes that an object can be partitioned into four quadrants around a given center. Within each quadrant the object contour is either monotonically decreasing or increasing in the direction allowed for the quadrant, see Fig. 1(b). Let $l \in \{a,b,c,d\}$ be a quadrant of object i with respect to o_i. Consider quadrant $l = a$ and denote by $\mathcal{S}^i_a$ the set of all adjacent ordered pixel pairs (p,q) in this quadrant. We use standard 4- or 8-neighborhood system and assume that q is immediately above pixel p or to the left of pixel p. The compact shape prior $f^{s,i}_l(\mathbf{x})$ for quadrant $l = a$ penalizes switching from label i to any other label in the direction from left to right and top to bottom:

$$f^{s,i}_a(\mathbf{x}) = \sum_{(p,q)\in\mathcal{S}^i_a} h^i_{pq}(x_p,x_q), \quad \text{where} \quad h^i_{pq}(x_p,x_q) = \begin{cases} K & \text{if } x_q = i \wedge x_p \neq i, \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

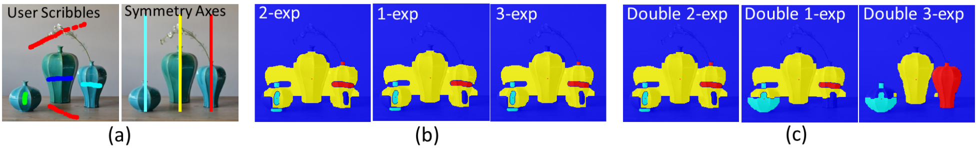

Figure 2: Expansion vs. double expansion: (a) input image with scribbles and symmetry axes, (b) expansion gets stuck in bad local minimum, (c) double expansion escapes it.

Similar pairwise potentials are defined for the other three quadrants with respect to their allowed orientations, see Fig. 1(b). The compact shape prior for object i is given by

$$f^{s,i} = \sum_{l \in \{a,b,c,d\}} f_l^{s,i}(\mathbf{x}).$$

2.3 Symmetry Prior

The symmetry prior assumes that an object is symmetric with respect to a given axis. For simplicity, here we assume that objects are vertically symmetric around their centers. Let $\mathcal{S}^i$ be a set of all pairs of symmetric pixels with respect of a given vertical axis of object i. The symmetry prior for object i penalizes pairs of pixels that do not agree on label i, that is

$$f^{s,i}(\mathbf{x}) = \sum_{(p,q)\in\mathcal{S}^i} h^i_{pq}(x_p,x_q), \quad \text{where} \quad h^i_{pq}(x_p,x_q) = \begin{cases} K & \text{if } x_p \neq x_q \wedge x_p = i, \\ K & \text{if } x_p \neq x_q \wedge x_q = i, \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

3 Optimization

Optimization with Expansion: First we review move-making algorithms [3]. Let $\mathbf{x}$ be a labeling and let $M(\mathbf{x})$ a set of labelings or *moves* that includes $\mathbf{x}$. We say $\hat{\mathbf{x}}$ is a local minimum w.r.t. $M(\mathbf{x})$ if $\forall \mathbf{x}' \in M(\mathbf{x}), f(\mathbf{x}') \geq f(\hat{\mathbf{x}})$. In an α-expansion move, each pixel p can stay with its current label x_p or switch to α. That is, a move $\mathbf{x}'$ is called an α-expansion if $x'_p = x_p$ whenever $x'_p \neq \alpha$. Let $M^\alpha(\mathbf{x})$ be the set of all α-expansions for a fixed α. Expansion algorithm iterates over $\alpha \in \mathcal{L}$ and finds an optimal α-expansion move for each α till convergence to a local minimum w.r.t expansion moves $\cup_\alpha M^\alpha(\mathbf{x})$.

Starting from a finite energy solution, we show that expansion moves are submodular for energy in Eq. 1. According to [2], it is enough to show that each pairwise potentials $g(x_p,x_q)$ in the energy satisfies $g(\beta,\gamma)+g(\alpha,\alpha) \leq g(\alpha,\gamma)+g(\beta,\alpha), \forall \alpha,\beta,\gamma \in \mathcal{L}$. First, it holds that $g(\alpha,\alpha) = 0, \forall \mathcal{L}$ for all pairwise potentials in (1). Second, since current solution must have a finite cost, we only need to show that the following inequality holds for $g(\beta,\gamma) \neq \infty$:

$$g(\beta,\gamma) \leq g(\alpha,\gamma)+g(\beta,\alpha), \;\; \forall \alpha,\beta,\gamma \in \mathcal{L}. \tag{8}$$

Now, consider boundary regularization potentials in (3). Suppose $v_{pq}(\beta,\gamma) \neq 0$. Then, $v_{pq}(\beta,\gamma) = w_{pq}$ and $\beta \neq \gamma$. Therefore, either $\alpha \neq \gamma$ or $\alpha \neq \beta$ or both. This implies, that either $v_{pq}(\alpha,\gamma) \geq w_{pq}$ or $v_{pq}(\beta,\alpha) \geq w_{pq}$ or both. Therefore (8) holds.

Consider star shape potentials in (5). If $h_{pq}(\beta,\gamma) = K$, then $\beta \neq i$ and $\gamma = i$. If $\alpha = i$ then $h_{pq}(\beta,\alpha) = K$. If $\alpha \neq i$, then $h_{pq}(\alpha,\gamma) = K$. Inequality (8) holds in both cases. Expansion submodularity of compact (6) and symmetry (7) shape priors is shown in a similar way.

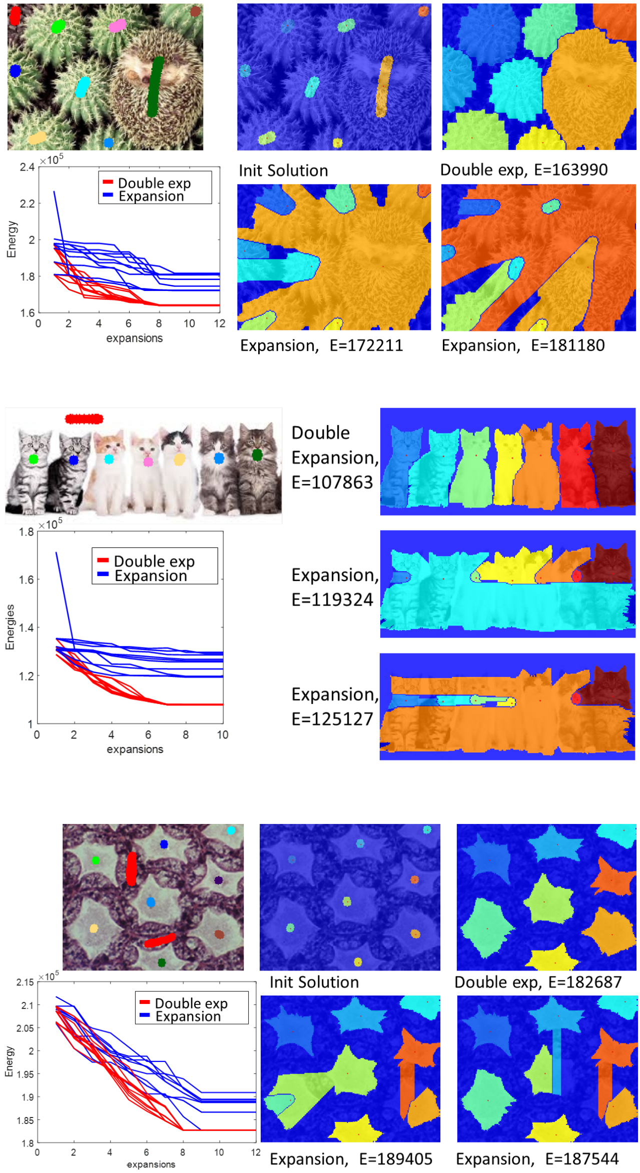

Figure 3: Comparison of double α-expansion with expansion for multilabel starshape prior. Ten runs with random label order were performed. In contrast to expansion, double expansion is robust to label order and always converges to a lower energy solution (top-right). For expansion we show two randomly chosen solutions (bottom-right) at convergence, illustrating convergence to poor local minima.

Optimization with Double Expansion: To motivate double expansion, we first analyze drawbacks of standard expansion. Consider the image in Fig. 2 (a) with three symmetrical objects sharing the same appearance model. User scribbles are in (a-left) and the symmetry axes estimated from them are shown in (a-right). We show in (b) solutions for three consecutive runs of expansion. Initially, everything is assigned to background except the foreground scribbles. After expansion on label 2 (yellow), most of the foreground pixels are erroneously assigned to label 2 due to appearance. Note, user scribbles are hard constrained and, therefore, preserve their label. Subsequent expansion on label 1 (light blue) is unable to grab any pixels already assigned to the yellow object, because symmetry of the yellow object would be violated. Further expansion on label 3 (red) runs into the same problem.

To fix the expansion problem, we need a move that in addition to expanding on a new object label, can “erase" areas erroneously assigned to other object labels. For example the blue object would be able to expand if the right part of the yellow object was simultaneously “erased”. Erasing means replacing erroneous object pixels with the background label. This move requires expanding on two labels at once: the new object label α and the background. We develop such *double expansion* move.

Given label $\alpha \in \mathcal{L}$, in a double α-expansion move, each pixel chooses one of the three options: keep its current label, switch to α or to 0. That is a move $\mathbf{x}'$ is a double α-expansion if $x'_p = x_p$ whenever $x'_p \notin \{a, 0\}$. Let $M_d^\alpha(\mathbf{x})$ be the set of all double α-expansions for a fixed α. Double expansion algorithm iterates over $\alpha \in \mathcal{L}$ and finds an optimal double α-expansion move for each α until convergence to a local minimum w.r.t. $\cup_\alpha M_d^\alpha(\mathbf{x})$.

Consider Fig. 2 (c). Double expansion on label 2 results in an identical solution to standard expansion. Next, using the same label order for expansion, double expansion on label 1 (light-blue) grabs a big portion of the yellow object. This can be done, because background label can fix the symmetry of the yellow object while expanding on the light-blue object. Expansion on the label 3 (red) reduces the area erroneously assigned to the yellow object even more. With double expansion, new objects can compete for the foreground pixels while accounting for the shape prior of the label that is expanding and the shape prior of the other labels. The result at convergence is in Fig. 5.

Double expansion is a more powerful move compared to expansion, that is $M^\alpha(\mathbf{x}) \subset M_d^\alpha(\mathbf{x})$. Starting from the same labeling, a single move of double α-expansion is guaranteed to be at least as good as a single move of α-expansion. While we cannot guarantee lower energy of a sequence of double expansion moves, we observe this behavior in practice.

Finding the optimal double α-expansion is formulated as minimization of ternary (3-label) energy. Let $\mathbf{x}$ be the current labeling. Let $\mathcal{L}_3 = \{0, 1, 2\}$. We introduce a variable $y_p \in \mathcal{L}_3$ for each p, and store them in $\mathbf{y} = (y_p \,|\, p \in \mathcal{P})$. There is a one-to-one correspondence $\mathbf{e}$ between all double α-expansion moves from $\mathbf{x}$ in $M_d^\alpha(\mathbf{x})$ and all labellings $\mathbf{y}$. For each $\mathbf{y}$, let us denote the corresponding double α-expansion as $\mathbf{e}(\mathbf{y})$[1]. Then

$$\mathbf{e}(\mathbf{y})_p \quad = \quad \begin{cases} \alpha & \text{if } y_p = 0 \\ 0 & \text{if } y_p = 1 \\ x_p & \text{if } y_p = 2 \end{cases} \tag{9}$$

We now define ternary energy $t(\mathbf{y})$. Let $t_p(y_p) = u_p(\mathbf{e}(\mathbf{y})_p)$, for any $(p,q) \in \mathcal{N}$, let $t^r_{pq}(y_p, y_q) = v_{pq}(\mathbf{e}(\mathbf{y})_p, \mathbf{e}(\mathbf{y})_q)$, and for any $(p,q) \in \mathcal{S}$, let $t^s_{pq}(y_p, y_q) = h_{pq}(\mathbf{e}(\mathbf{y})_p, \mathbf{e}(\mathbf{y})_q)$. Energy $t(\mathbf{y})$ is the sum of all unary and pairwise terms. It is straightforward to check that $t(\mathbf{y}) = f(\mathbf{e}(\mathbf{y}))$. Therefore $\mathbf{y}^*$ that minimizes $t(\mathbf{y})$ corresponds to an optimal double

[1]Note that $\mathbf{e}(\mathbf{y})$ also depends on $\mathbf{x}$, which we omit for less cumbersome notation.

α-expansion. In the supplementary material we show that $t(\mathbf{y})$ is multilabel submodular and thus can be optimized with the construction from [19]. Moreover, double expansion algorithm has optimality guarantees for energies of the form $f(\mathbf{x}) = f^a(\mathbf{x}) + f^r(\mathbf{x})$. The following theorem is proved in the supplementary materials.

Theorem 3.1 *Let $\hat{\mathcal{L}} = \{1, ..., m\}$, i.e. the set of labels excluding the background label* 0*. Let $\hat{\boldsymbol{x}}$ be a local minimum of the energy above with respect to double expansion moves and let $\boldsymbol{x}^*$ be a globally optimal solution. Then $f(\hat{\boldsymbol{x}}) \leq |\hat{\mathcal{L}}^*| f(\boldsymbol{x}^*)$, where $|\hat{\mathcal{L}}^*|$ is the number of labels from $\hat{\mathcal{L}}$ that appear in the optimal solution $\boldsymbol{x}^*$.*

4 Applications

Below we apply our double α-expansion algorithm to multilabel segmentation with compact [4], star [22] and vertical symmetry shape prior for each label except the background. These problems are formulated as a pairwise multilabel energy as we explained in Sec. 2.

Fig. 3-5 show examples of images with multiple foreground objects. In Fig. 3 we show segmentation results with star shape prior applied to each foreground object, whereas Fig. 4-5 show similar results with compactness and symmetry shape prior respectively. For each image, the input scribbles were used to extract both, the appearance models for foreground and background as well as the object centers/axes required for the shape prior.

All experiments were initialized with user scribbles and repeated ten times, randomizing the order of labels for double expansion. For comparison, we repeated the same experiments with the standard expansion [3]. All experiments were run till convergence, which always occurred after one pass over the labels. Our double expansion optimization is very robust to the order of expansion, converging to the same local minimum. In contrast, standard expansion is sensitive to the order of expansion and in almost all cases converges to a local minimum with a much higher energy. For each image we show the result obtained with our double expansion and two randomly chosen results obtained with standard expansion.

5 Conclusions

We propose a new class of energies for simultaneous segmentation of multiple foreground objects with a common shape prior and a new double expansion algorithm to optimize such energies. Compared to expansion, the double expansion algorithm converges to lower energy solutions, has optimality guarantees and is robust to order of expansion in practice. We experiment with several types of shape prior and empirically demonstrate the advantage of the double expansion. In the future we plan to incorporate more challenging, non-submodular and higher order shape priors such as convexity and connectivity. Moreover, we plan to extend the class of the energies to include a label cost term [7].

References

[1] Andrew Blake, Pushmeet Kohli, and Carsten Rother. *Markov Random Fields for Vision and Image Processing*. MIT Press, 2011.

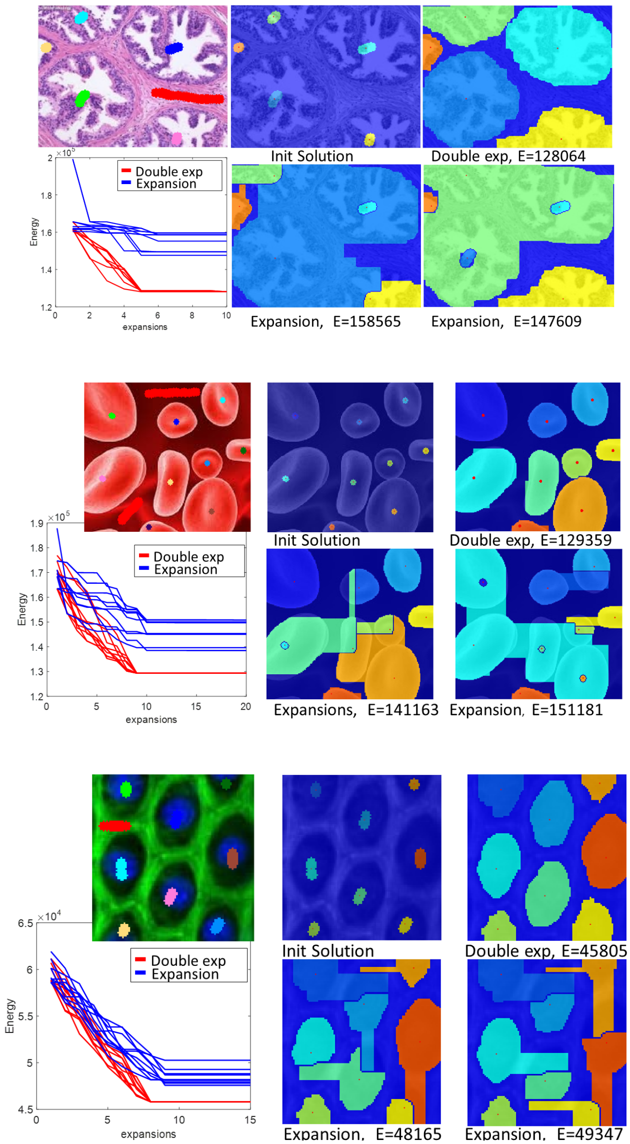

Figure 4: Comparison of double α-expansion with expansion for multilabel shape compactness. Ten runs with random label order were performed. In contrast to expansion, double expansion is robust to label order and always converges to a lower energy solution (top-right). For expansion we show two randomly chosen solutions (bottom-right) at convergence, illustrating convergence to poor local minima.

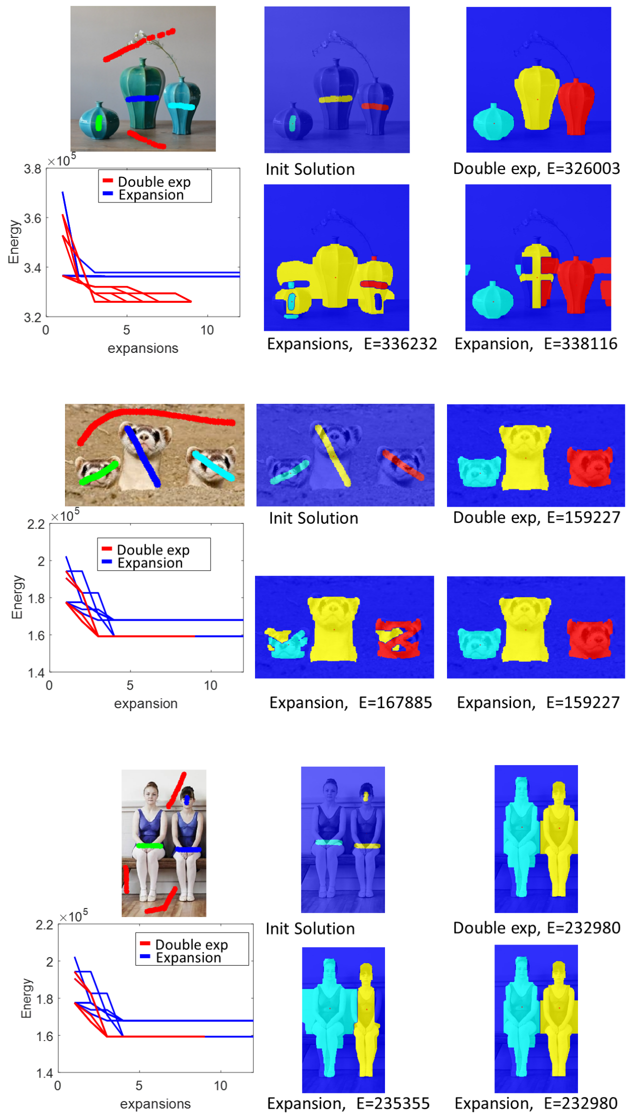

Figure 5: Comparison of double α-expansion with expansion for multilabel symmetry prior. Ten runs with random label order were performed. In contrast to expansion, double expansion is robust to label order and always converges to a lower energy solution (top-right). For expansion we show two randomly chosen solutions (bottom-right) at convergence, illustrating convergence to poor local minima.

[2] Yuri Boykov and Olga Veksler. Graph cuts in vision and graphics: Theories and applications. In Nikos Paragios, Yunmei Chen, and Olivier D. Faugeras, editors, *Handbook of Mathematical Models in Computer Vision*, pages 79–96. Springer, 2006. ISBN 978-0-387-26371-7.

[3] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23: 1222–1239, November 2001.

[4] Piali Das, Olga Veksler, Vyacheslav Zavadsky, and Yuri Boykov. Semiautomatic segmentation with compact shape prior. *Image Vision Computing*, 27(1-2):206–219, 2009.

[5] Piali Das, Olga Veksler, Vyacheslav Zavadsky, and Yuri Boykov. Semiautomatic segmentation with compact shape prior. *Image Vision Comput.*, 27(1-2):206–219, 2009.

[6] Andrew Delong and Yuri Boykov. Globally optimal segmentation of multi-region objects. In *In ICCV*, pages 1–8, 2009.

[7] Andrew Delong, Anton Osokin, Hossam N. Isack, and Yuri Boykov. Fast Approximate Energy Minimization with Label Costs. *International Journal of Computer Vision*, 96 (1):1–27, 2012.

[8] P. Felzenszwalb and O. Veksler. Tiered scene labeling with dynamic programming. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.

[9] Stuart Geman and Christine Graffigne. Markov random field image models and their applications to computer vision. In *Proceedings of the International Congress of Mathematicians*, 1986.

[10] Lena Gorelick, Olga Veksler, Yuri Boykov, and Claudia Nieuwenhuis. Convexity shape prior for segmentation. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, pages 675–690, 2014.

[11] Varun Gulshan, Carsten Rother, Antonio Criminisi, Andrew Blake, and Andrew Zisserman. Geodesic star convexity for interactive image segmentation. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pages 3129–3136, 2010.

[12] Hossam Isack, Olga Veksler, Milan Sonka, and Yuri Boykov. Hedgehog shape priors for multi-object segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[13] Hossam Isack, Olga Veksler, Ipek Oguz, Milan Sonka, and Yuri Boykov. Efficient optimization for hierarchically-structured interacting segments (HINTS). In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[14] Vladimir Kolmogorov and Yuri Boykov. What metrics can be approximated by geo-cuts, or global optimization of length/area and flux. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 - Volume 01*, pages 564–571, Washington, DC, USA, 2005.

[15] Xiaoqing Liu, Olga Veksler, and Jagath Samarabandu. Order-preserving moves for graph-cut-based optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(7):1182–1196, 2010.

[16] Claudia Nieuwenhuis, Eno Toeppe, Lena Gorelick, Olga Veksler, and Yuri Boykov. Efficient Squared Curvature. In *IEEE conf. on Comp. Vision and Pattern Recognition (CVPR)*, 2014.

[17] Sebastian Nowozin and Christoph H. Lampert. Global interactions in random field models: A potential function ensuring connectedness. *SIAM J. Imaging Sciences*, 3(4): 1048–1074, 2010.

[18] Loïc Alain Royer, David L. Richmond, Carsten Rother, Bjoern Andres, and Dagmar Kainmueller. Convexity shape constraints for image segmentation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June, 2016*, pages 402–410, 2016.

[19] Dmitrij Schlesinger. Exact solution of permuted submodular minsum problems. In *Energy Minimization Methods in Computer Vision and Pattern Recognition, 6th International Conference, EMMCVPR 2007, Ezhou, China, August 27-29, 2007, Proceedings*, pages 28–38, 2007.

[20] E. Strekalovskiy and D. Cremers. Generalized ordering constraints for multilabel optimization. In *International Conference on Computer Vision (ICCV)*, 2011.

[21] Christopher W. Tyler. *Human symmetry perception and its computational analysis*. Psychology, 1996.

[22] Olga Veksler. Star shape prior for graph-cut image segmentation. In *European Conference on Computer Vision (ECCV)*, 2008.

[23] Olga Veksler. Multi-label moves for mrfs with truncated convex priors. *International Journal of Computer Vision*, 98(1):1–14, 2012.

[24] Sara Vicente, Vladimir Kolmogorov, and Carsten Rother. Graph cut based image segmentation with connectivity priors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[25] John Winn and Jamie Shotton. The layout consistent random field for recognizing and segmenting partially occluded objects. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 37–44, 2006.