

Iterative Graph Cuts for Image Segmentation with a Nonlinear Statistical Shape Prior

Joshua C. Chang · Tom Chou

Published online: 4 May 2013
© Springer Science+Business Media New York 2013

Abstract Shape-based regularization has proven to be a useful method for delineating objects within noisy images where one has prior knowledge of the shape of the targeted object. When a collection of possible shapes is available, the specification of a shape prior using kernel density estimation is a natural technique. Unfortunately, energy functionals arising from kernel density estimation are of a form that makes them impossible to directly minimize using efficient optimization algorithms such as graph cuts. Our main contribution is to show how one may recast the energy functional into a form that is minimizable iteratively and efficiently using graph cuts.

Keywords Image segmentation · MM · Graph cuts · Energy minimization · Statistical shape prior · Kernel density estimation

1 Introduction

Graph cuts provides an ingenious technique for image segmentation that relies on transforming the problem of energy minimization into the problem of determining the maximum flow or minimum cut on an edge-weighted graph. By using

graph cuts, segmentations can be found efficiently in low-order polynomial time. When images are noisy, segmentation performed using only image data produces poor results. In these cases, regularization is needed.

Typically, when one is analyzing images, one has expectations for what he or she is looking for—people, birds, cars, cows, bacteria, whatever the object may be. In these cases it is natural to try to restrict segmentation to results that match the shape of the desired object. Within any class of objects there is often variability in shape. Such variability can arise even when dealing with shape-invariant objects, for example by changes in pose or perspective. Probabilistic expression of shape knowledge is a natural way of capturing this variability.

1.1 Related Prior Work

In Chang et al. [5], we presented a technique for tracking an object whose boundary motion is modeled using the level-set method. Such objects do not retain any particular a-priori shape. We remarked that one could adapt our regularization method to the segmentation of objects that do retain their shapes. The purpose of this article is to explicitly present a method for performing such a task.

Our method is grounded in the theory of graph cuts-based image segmentation with shape-based regularization, where segmentation is performed using a-priori shape knowledge. The prior literature on this subject is vast. Briefly, we mention a few influential articles. Slabaugh and Unal [29] developed a graphcuts-based method for segmentation assuming that the desired object is an ellipse. Allowing for more flexibility, Freedman and Zhang [14] provided a method to segment any particular shape based on the signed distance embedding. Some methods have concentrated on looking at shapes with certain characteristics such as compactness or

J.C. Chang (✉)
Mathematical Biosciences Institute, The Ohio State University,
Jennings Hall, 3rd Floor, 1735 Neil Avenue, Columbus,
OH 43210, USA
e-mail: chang.1166@mbi.osu.edu

T. Chou
UCLA Biomathematics and Mathematics, BOX 951766,
Room 5303 Life Sciences, Los Angeles, CA 90095-1766, USA
e-mail: tomchou@ucla.edu

local convexity [11, 31]. Other methods have treated the task of space specification as a statistical problem.

Statistical shape priors have taken many forms in the literature [10, 23, 35]. Two of the most common forms of shape representation are kernel-PCA and kernel-based level-set embeddings. Both methods learn a probability distribution for the target shape based on a set of training shapes or templates. They differ in how they use the training images. Kernel PCA-based methods project the training templates into principal-component space and define probability distributions over the resulting vector spaces. Level-set-based methods define a distance between implicitly embedded shapes and define a probability distribution with respect to the distance.

One major disadvantage of the PCA-based approach is that the resulting probability distribution for the shape is Gaussian. Given a set of training templates, this approach will tend to favor shapes that are more similar to the space-averaged shape, thereby potentially making certain training shapes improbable. This problem also exists for exponential-family shape priors where the log-prior is a linear combination of the training templates. Combinations of shapes may not be valid shapes [9], and under this approach certain training templates may become improbable as well. Instead, we opt for a kernel-density based approach.

Kernel density estimation (KDE) is a non-parametric method for the estimation of probability distributions. Its main advantage is its flexibility and its ability to capture multiple modes. Our method here, and that given in Chang et al. [5], can be described as a graph-cuts adaptation of the method of Cremers et al. [8]. In their paper, they treated shape-prior specification as a density estimation problem and used KDE to derive a prior from an ensemble of training shapes. The resulting energy functionals, being nonlinear, require some intervention to use in a graph-cuts framework.

In this manuscript, we make the main contribution of demonstrating the use of majorization-minimization (MM) for the iterative relaxation of a nonlinear energy functional using graph-cuts. The MM algorithm is a generalization of the well-known expectation-maximization (EM) algorithm. It is often useful for linearizing nonlinear objective functions. We use MM to find a surrogate minimization problem that is solved using graph-cuts. In doing so, we expose a mathematical link between the energy functional of Cremers et al. [8], and the linear energy functionals commonly used for graph-cuts with statistical priors [10, 23, 35]. The subsequent use of graph cuts results in a computationally more efficient approach relative to that of using level set approaches [15].




Fig. 1 Embedding of shapes using signed Euclidean distance functions. A shape (pictured is a leaf) is embedded as a signed Euclidean distance function where values inside are the positive distance from the boundary, and values outside are negative distances from a boundary. The boundary of the outlined leaf is the 0-level set of the function

2 Mathematical Method

2.1 Shapes

Signed-distance functions provide a handy tool to represent shapes mathematically. Shapes or regions $\Omega \subset \mathbb{R}^d$ can be implicitly represented by a function $\phi_\Omega(\mathbf{s}) : \mathbb{R}^d \rightarrow \mathbb{R}$, which for every $\mathbf{s} \in \mathbb{R}^d$ is the signed Euclidean distance from s to the boundary of Ω (Fig. 1). In this paper we take the convention that $\phi_\Omega(\mathbf{s})$ is positive if $\mathbf{s} \in \Omega$, and negative if $\mathbf{s} \in \mathbb{R}^d \setminus \Omega$.

In many applications, since the shape of an object is only approximately known, it is advantageous to represent the knowledge of the shape probabilistically. We use a kernel density estimate of the distribution of possible shapes embedded as a collection of individual discretized level sets. Let us first denote χ_Ω the characteristic function for a region Ω ,

$$\chi_\Omega(\mathbf{s}) = \begin{cases} 1, & \mathbf{s} \in \Omega; \\ 0, & \mathbf{s} \notin \Omega. \end{cases} \tag{1}$$

Then, for two shapes Ω and Λ , embedded as signed distance functions ϕ_Ω and ϕ_Λ , we introduce an (asymmetric) pairwise shape energy

$$U_{\text{shape}}(\Omega, \Lambda) = \underbrace{\int_{\mathbb{R}^d} (\chi_\Omega(\mathbf{s}) - \chi_\Lambda(\mathbf{s}'))^2 |\phi_\Lambda(\mathbf{s}')|^\lambda ds}_{\text{mass mismatch}} + \underbrace{\oint_{\partial\Omega} |\phi_\Lambda(\mathbf{z}')|^\lambda dz'}_{\text{boundary mismatch}}, \tag{2}$$

where $\lambda \geq 0$. This expression is a generalization of other shape energies seen in the literature, where commonly

$\lambda = 0, 1, 2$ [9, 10, 16, 21, 22, 24, 34, 35]. Although symmetric shape energies are desirable [7], for general use in our optimization scheme we require an asymmetric energy that does not depend on the Euclidean-distance shape embedding of the evolving segmentation. If one uses $\lambda = 0$, then it is possible to make this energy symmetric with small modifications (see Sect. 4).

To make this energy robust to changes in scale, orientation, and location, a transformation $\mathbf{s}' = \mathbb{T}(\mathbf{s})$ that maps points \mathbf{s} in the reference frame of Ω to points \mathbf{s}' in the reference frame of Λ is needed. In Appendix A, we provide a method for finding rigid transformations of the form

$$\mathbb{T}(\mathbf{s}) = \alpha \mathbf{R}(\omega)(\mathbf{s} - \mathbf{c})$$

where α is a scaling factor, $\mathbf{R}(\omega)$ is a rotation matrix, and \mathbf{c} is a translation vector. These three parameters are chosen to minimize Eq. (2) (see Appendix A for details on calculating these parameters). This method worked well for our examples; however, other shape alignment schemes may be used [1, 18, 20, 30, 33].

After transformation, we can use the transformation-minimized energy in Eq. (2) to assemble reference shapes $\{\Omega_j\}_{j=1}^J$ into a probability distribution by using a kernel of the form

$$K(\Omega, \Lambda) = \sqrt{\frac{\beta}{2\pi}} e^{-\beta U_{\text{shape}}(\Omega, \Lambda)}. \tag{3}$$

We can then represent a probability distribution over shapes, $p(\Omega)$ according to:

$$p(\Omega) \propto \sum_{j=1}^J w_j K(\Omega, \Omega_j), \tag{4}$$

where w_j are weighting coefficients that sum to one. This representation of the prior is a kernel density estimate [28] (KDE) of the distribution of shapes. We set β , the reciprocal temperature, to be the scale-normalized kernel width [6, 28]:

$$\beta = \left[\sum_{j=1}^J \frac{w_j}{s_j^\lambda} \min_{k \neq j} U_{\text{shape}}(\Omega_k, \Omega_j) \right]^{-1}, \tag{5}$$

where s_j is the inertial scale of a shape Ω_j as defined in Jiang and Tomas [18]. The division by s_j^λ makes β scale-invariant, so that it need only be computed once. Note that this expression of the shape prior is potentially multi-modal.

2.2 Generative Image Modeling

The task is to identify an object Ω in an image $I : S \subset \mathbb{R}^d \rightarrow \mathbb{R}^q$, where q refers to the dimension of the color space used. We model an image probabilistically with distributions

of intensity values conditional on a parameter vector θ , and region:

$$p(I(\mathbf{s})|\theta) = \begin{cases} p_\Omega(I|\theta) & \mathbf{s} \in \Omega \text{ (foreground);} \\ p_\Delta(I|\theta) & \mathbf{s} \in \Delta = S \setminus \Omega \text{ (background).} \end{cases} \tag{6}$$

Information about θ and Ω can be incorporated as prior probability distributions $p(\theta|\Omega)$ and $p(\Omega)$. With these prior distributions, we can write the joint posterior distribution $p(\Omega, \theta|I) \propto p(I|\Omega, \theta)p(\theta|\Omega)p(\Omega) = e^{-U(\Omega, \theta)}$, where $U(\Omega, \theta)$ is an energy (not to be confused with U_{shape}).

We wish to infer the segmentation Ω by maximizing the posterior probability relative to Ω . To this end, we will maximize the logarithm of the posterior, or equivalently, minimize the energy

$$\begin{aligned} U(\Omega, \theta) &= -\log \left[\overbrace{p(I|\Omega, \theta)}^{\text{likelihood}} \overbrace{p(\theta|\Omega)p(\Omega)}^{\text{prior}} \right] \\ &= -\sum_{\mathbf{s} \in \Omega} \log p_\Omega(I(\mathbf{s})|\theta) - \sum_{\mathbf{s} \in \Delta} \log p_\Delta(I(\mathbf{s})|\theta) \\ &\quad - \log \underbrace{\sum_{j=1}^J w_j K(\Omega, \Omega_j)}_{\text{Eqs. (3), (4)}} \\ &\quad - \log p(\theta|\Omega) + \text{const.} \end{aligned} \tag{7}$$

2.3 Majorization-Minimization

The energy functional in Eq. (7) resembles the energy functional given by Cremers et al. [8], where level-set based gradient descent is used for relaxation. Unfortunately, optimization performed in this manner is slow, since only small deformations of an evolving contour occur in each iteration.

We take an iterative two-step approach to minimizing this energy. Given the segmentation estimate $\Omega^{(n)}$ in the n -th step, we minimize the energy with respect to θ [25].

Given $\theta^{(n)}$, we find the optimal $\Omega^{(n+1)}$ by iteratively minimizing

$$\begin{aligned} Q(\Omega|\Omega^{(n)}) &= -\sum_{\mathbf{s} \in \Omega} \log p_\Omega(I(\mathbf{s})|\theta) - \sum_{\mathbf{s} \in \Delta} \log p_\Delta(I(\mathbf{s})|\theta) \\ &\quad + \underbrace{\frac{\beta}{2} \sum_{j=1}^J c_j^{(n)} U_{\text{shape}}(\Omega, \Omega_j)}_{\text{majorization}}, \end{aligned} \tag{8}$$

where

$$c_j^{(n)} = \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)}. \tag{9}$$

In Eq. (8), we have replaced logarithm-sum term of Eq. (7) with a linearized majorization term, by exploiting the concavity of the logarithm function. The majorization term defines a surrogate minimization problem that is more easily solved—in our case using graph cuts—albeit at the price of iteration (Algorithm 1). Such an algorithm is commonly known as a majorization-minimization or MM algorithm [17] (see Appendix B). This approach is generally possible for other energies that contain components of the form $f(J[\Omega])$, where the function $f : \mathbb{R} \rightarrow \mathbb{R}$ is convex, and J is a functional that is dependent on Ω .

2.4 Graph Cuts for Surrogate Energy Relaxation

Here, we describe minimization of the surrogate energy in Eq. (8) using graph cuts. Graph cut methods have their grounding in combinatorial optimization theory, and are concerned with finding the minimum-cost cut in an undirected graph. A cut is a partition of a connected graph into

Algorithm 1 MM-Graphcut algorithm

Obtain initial guess of image intensity model parameters $\theta^{(0)}$.
 Obtain initial guess of the segmentation $\Omega^{(0)}$, by segmentation without a shape prior.
 Align and rescale the shape prior templates to $\Omega^{(0)}$.
 $n = 1$
while $U_{\text{shape}}(\Omega^{(n)}, \Omega^{(n-1)}) > \text{tol}$ **do**
 Recalculate the shape weights $c_j^{(n)}$.
 Reweight the edges and perform max flow to estimate $\Omega^{(n+1)}$
 Re-estimate the image intensity parameters $\theta^{(n+1)}$ (as needed)
 Re-align the shape prior templates to the guess $\Omega^{(n)}$ (as needed)
 $n++$
end while

two disconnected sets. The cost of a cut is the sum of the edge weights along a cut, and a max-flow min-cut algorithm finds the cut with the lowest cost. To use graph cuts for image segmentation, we must express our energy function in terms of edge-weights on a graph. Following Boykov and Jolly [2], we begin by expressing the energy given in Eq. (8) as a function of the vertices V and edges \mathcal{E} of a graph $\mathcal{G} = (V, \mathcal{E})$:

$$U(\mathcal{G}) = \sum_{s \in V} U_V(s) + \sum_{(s, \mathbf{u}) \in \mathcal{E}} U_{\mathcal{E}}(s, \mathbf{u}).$$

In the graph that we construct, each pixel is assigned a node, and edges exist between nodes representing neighboring pixels (Fig. 2). The neighbor edges are known as neighbor-links (n-links). Two special nodes called the source (foreground) and sink (background) are added, along with edges connecting these nodes to each pixel node (Fig. 2). These edges are known as terminal-links (t-links).

We want to infer an unknown two-coloring on the nodes of the graph that represents inclusion of a node s into either the foreground set Ω , or the background set Δ . This inference involves splitting the graph into two parts (Fig. 2), where a pixel’s connection to the source represents inclusion into the foreground set.

Discretizing the shape divergence (Eq. (2)) over an eight-connected neighbor graph yields

$$U_{\text{shape}}(\Omega, \Delta) = \sum_s \overbrace{|\chi_{\Omega}(s) - \chi_{\Delta}(s')|}^{\text{indicator of label mismatch}} |\phi_{\Delta}(s')|^{\lambda} + \sum_{(s, \mathbf{u}) \in \mathcal{N}} \overbrace{|\chi_{\Omega}(s)(1 - \chi_{\Omega}(\mathbf{u})) + \chi_{\Omega}(\mathbf{u})(1 - \chi_{\Omega}(s))|}^{\text{indicator of } s, \mathbf{u} \text{ lying across boundary of } \Omega} \times \frac{\pi}{8\|s - \mathbf{u}\|} \left| \phi_{\Delta} \left(\frac{s' + \mathbf{u}'}{2} \right) \right|^{\lambda}, \tag{10}$$




Fig. 2 (Left) Embedding of segmentation into a graph. In the graph-cuts framework, the problem of segmentation is transformed into the problem of finding the minimum cut in an edge-weighted graph. The energy to minimize is rewritten as a sum of atomic terms and a sum of

neighbor-interaction terms. The edges are weighted such that the cut is the energy of the corresponding segmentation. (Right) Resulting segmentation. The cutting of the graph yields two disconnected node sets where Ω is the foreground and Δ is the background

where \mathbf{N} refers to the set of tuples of nodes that are neighboring in space. The level set function ϕ_Δ^λ is taken to be the Euclidean distance embedding of Δ , transformed to the center, scale, and orientation of Ω . As mentioned before, we use an asymmetric representation of the shape divergence since the signed-distance function is a property of a global segmentation and is not obtainable from local labeling.

Using this expression, one sees that

$$\begin{aligned}
 U_V(\mathbf{s}) &= -\chi_\Omega(\mathbf{s}) \log p_\Omega(I(\mathbf{s})|\boldsymbol{\theta}) \\
 &\quad - (1 - \chi_\Omega(\mathbf{s})) \log p_\Delta(I(\mathbf{s})|\boldsymbol{\theta}) \\
 &\quad + \frac{\beta}{2} \sum_{j=1}^J c_j^{(n)} |\phi_{\Omega_j}(\mathbf{s}')|^\lambda |\chi_\Omega(\mathbf{s}) - \chi_{\Omega_j}(\mathbf{s}')| \quad (11)
 \end{aligned}$$

and

$$\begin{aligned}
 U_{\mathcal{E}}(\mathbf{s}, \mathbf{u}) &= \frac{\pi\beta}{16\|\mathbf{s} - \mathbf{u}\|} \left[\sum_{j=1}^J c_j^{(n)} \left| \phi_{\Omega_j} \left(\frac{\mathbf{s}' + \mathbf{u}'}{2} \right) \right|^\lambda \right. \\
 &\quad \left. \times |\chi_\Omega(\mathbf{s})(1 - \chi_\Omega(\mathbf{u})) + \chi_\Omega(\mathbf{u})(1 - \chi_\Omega(\mathbf{s}))| \right]. \quad (12)
 \end{aligned}$$

The neighbor energy given in Eq. (12) obeys the *submodularity property*, that is,

$$\begin{aligned}
 U_{\mathcal{E}}(\mathbf{s} \in \Omega, \mathbf{u} \notin \Omega) + U_{\mathcal{E}}(\mathbf{s} \notin \Omega, \mathbf{u} \in \Omega) &= \frac{\pi\beta}{8\|\mathbf{s} - \mathbf{u}\|} \sum_{j=1}^J c_j^{(n)} \left| \phi_{\Omega_j} \left(\frac{\mathbf{s}' + \mathbf{u}'}{2} \right) \right|^\lambda \\
 &\geq U_{\mathcal{E}}(\mathbf{s} \in \Omega, \mathbf{u} \in \Omega) + U_{\mathcal{E}}(\mathbf{s} \notin \Omega, \mathbf{u} \notin \Omega) \\
 &= 0.
 \end{aligned}$$

As a result, the energy (Eq. (8)) is minimizable with graph-cuts in polynomial time [3, 13].

To embed $U_V(\mathbf{s})$ into the graph, we set the weights of the t-links between each pixel and the source to the following

$$\begin{aligned}
 w(\mathbf{s}, \text{source}) &= -\log p_\Delta(I(\mathbf{s})|\boldsymbol{\theta}) \\
 &\quad + \frac{\beta}{2} \sum_{j=1}^J c_j^{(n)} \chi_{\Omega_j}(\mathbf{s}') |\phi_{\Omega_j}(\mathbf{s}')|^\lambda, \quad (13)
 \end{aligned}$$

and the weights of the t-links between each pixel and the sink to the following

$$w(\mathbf{s}, \text{sink}) = -\log p_\Omega(I(\mathbf{s})|\boldsymbol{\theta})$$

$$+ \frac{\beta}{2} \sum_{j=1}^J c_j^{(n)} (1 - \chi_{\Omega_j}(\mathbf{s}')) |\phi_{\Omega_j}(\mathbf{s}')|^\lambda. \quad (14)$$

The cutting of the edge from an \mathbf{s} to the source implies that $\mathbf{s} \in \Delta$, so it adds to the cost of the cut by the contribution of \mathbf{s} into the total energy as if $\mathbf{s} \in \Delta$. In other words, these weights can be interpreted as a pixel’s strength of belonging to each region.

To embed $U_{\mathcal{E}}$ into the graph, we set the n-links between pairwise neighboring pixels \mathbf{s} and \mathbf{u} to

$$\frac{\pi\beta}{16\|\mathbf{s} - \mathbf{u}\|} \sum_{j=1}^J c_j^{(n)} \left| \phi_{\Omega_j} \left(\frac{\mathbf{s}' + \mathbf{u}'}{2} \right) \right|^\lambda. \quad (15)$$

Our surrogate energy is now minimized by finding the minimum cut of the graph. For details on how to perform this optimization, we refer the reader to Boykov and Kolmogorov [3].

3 Results

We have implemented our method in Java using the Fiji application programming interface [27] and tested it on the regularization of luminosity-based segmentation of a variety of objects in images (Fig. 3). For comparison, we have obtained results of segmentation using a shape-free global length penalty [12]. In these examples, we have used Laplacian statistics for $p(I|\boldsymbol{\theta})$ and uninformative priors for $\boldsymbol{\theta}$ (by setting $p(\boldsymbol{\theta}) = 1$). The parameter λ was set to 2, and the inverse-temperature β was set to τ^2 . The algorithm was terminated when the labeling became stationary.

The first images we processed were low-contrast noisy images of a van. A van (Dodge B2500) was driven across a bumpy road and videotaped as it entered and exited the camcorder’s field of view. Snapshots of the recording were randomly extracted, showing the van in different regions an of the image, under different orientations and different scale factors. We manually constructed a shape prior for the van by hand-drawing five templates of a van in many poses. These templates of the van were aligned using the method detailed in Appendix A. Shape-free segmentation of the van captured the dorsal aspect of the van, but was inaccurate in delineating its ventral features. Using the shape prior, we obtained an admissible result that is recognizable as a van. Although a van maintains a rigid structure, its shape in images can still be considered probabilistic due to uncertainty in pose.

We also processed pictures of leaves from the Caltech 101 image database.¹ In the given examples, background objects interfere with the segmentation. Without a

¹http://www.vision.caltech.edu/Image_Datasets/leaves/.




Fig. 3 Shape-regularized segmentation. From left to right: *original image*, *length-penalized segmentation*, *formation of a statistical shape prior using a collection of templates*, *shape-regularized segmentation*. Segmentation of the original image of a van without shape information misses the wheels. A collection of templates (shown aligned and

summed) taken together define a probabilistic shape prior. Such representation of shapes is useful since it can account for intrinsic variation in a class of shapes, as well as variations that result from changing pose. In these templates, silhouettes of a van are taken in different poses




Fig. 4 Segmentation using a hybrid shape-prior. Using a combination of three and five-tip leaves, a “hybrid” shape prior is constructed. Although the average of the shape templates is a leaf with seven tips, the

algorithm relaxed the energy into segmentations of leaves with the correct number of tips. This feature is a result of the kernel-density energy functional being multimodal

shape prior, segmentations based on pixel intensities include these background distractors. To construct templates, we hand-traced six representative shapes for each of the two types of leaves. The regularized segmentations that we obtained were able to control for the background distractors.

In Fig. 4, we segmented both leaves using a prior composed of both three-tip and five-tip templates. Although the mean shape of the hybrid-shape prior is a leaf with seven tips, the algorithm relaxed the energy into segmentations of leaves with the correct number of tips.

4 Discussion

We have presented a method for regularization of segmentation using a shape prior that is learned through kernel-density-estimation. Using a majorization-minimization trick, we showed how one can minimize the resulting energy, that includes a non-linear kernel-density term, using graph-cuts.

The advantages of our method are of two types. First, our method uses kernel density estimation. Although there have been previous studies that include statistical shape priors in graph-cuts segmentation [6, 35], these methods have used

unimodal shape prior distributions where the input shape templates no longer retain their a-priori weights. As an example, we successfully segmented two types of leaves with very different shapes (Fig. 4). Segmentation of the same images using a hybrid shape-prior with the linear energy of previous graphcuts studies results in a collapse, as both types of shapes are improbable under the prior, and the mean shape (a leaf with seven tips) is improbable under the likelihood. Our shape-prior is a multi-modal kernel density estimator like the one used by Cremers et al. [8], thereby enjoying many of the benefits outlined in Cremers et al. [8].

Second, our method retains the speed advantage of other graphcut methods. Speed can be a significant limitation in applications involving “big data.” It is also an issue in tracking applications, where many successive images need to be processed. Using an MM algorithm, we find surrogate energy functionals that can be minimized quickly using graph cuts. The surrogate energy functional (Eq. (8)) contains an iteratively re-weighted shape regularization term, where shapes that are the most-similar the current segmentation estimate are the most influential. In our examples, convergence of segmentation defined as stationarity typically takes place in few (<10) iterations.

Closer examination of Eq. (8) is revealing. Essentially, our algorithm takes advantage of the mathematical link between the relaxation of the nonlinear kernel-shape-prior energy functional of Cremers et al. [8] and iterative relaxation of the linear shape-prior energy functional found in many studies on incorporating shape priors into graph cuts [6, 34, 35]. As a result, our method has a computational complexity that is similar to that of these graph-cuts methods, which also rely on iterative template-alignment, while retaining many of the advantages of kernel-density estimation.

4.1 Limitations

As in other iterative schemes for image segmentation, our optimization method is local. Consequently, the iteration can become trapped in local minima, as in the level-set method of Cremers et al. [8]. It is important to note, however, that the notion of locality is different between this method and in the level-set method. In the level-set method, a boundary is propagated to steady state, where the boundary undergoes only small deformations between iterations.

In this method, the boundary is able to jump between iterations [5]. The locality of our approach refers to locality in the influence of the templates—the weights—as the segmentation iterates. Each iteration involves the optimization of an energy (Eq. (8)) that is linear in U_{shape} . While the method we provide for minimization of this energy is local, because of the locality of the template-alignment procedure, other non-local techniques for minimizing the surrogate energy [21]

may be used. In this manner, one could trade-off speed for globality.

To prevent poorly-aligned templates from dominating over the likelihood, particularly in early iterations, one may treat β as an annealment parameter, as in Vu and Manjunath [34].

The main modeling deficiency of our method compared to the method of Cremers et al. [8] is the use of an asymmetric shape divergence. The advantages of using a symmetric shape distance have been documented in Cremers and Soatto [7]. If a symmetric distance is desired, one may work with $\lambda = 0$, thereby making our distance function the Hamming distance combined with a length-penalty.

4.2 Computational Considerations

Each iteration of the MM-algorithm is solved quickly using graph-cuts. Many efficient algorithms for this optimization exist, yielding excellent performance on modern hardware. There are many articles discussing the performance advantages of graph-cuts [4, 15]. For even more speed, there are GPU-accelerated approaches [32], and flow-reuse approaches [19].

The minimization of the surrogate energy is not the bottleneck of our method. To unleash the full potential of this shape regularization method, one needs to optimize the processing steps that are in the periphery the graph cut optimization—the inference of the image intensity model, and the alignment of the templates.

As far as we know, other methods of shape-regularized segmentation also require these processing steps. Fortunately, these operations can be easily parallelized. Furthermore, they need not be repeated in every step of the algorithm. The image intensity parameters need only change if a significant proportion of the labels change. Similarly, the templates only need re-alignment if the segmentation has significantly changed. To improve performance, one may calculate the transformation needed for a single template, and use that transformation for the entire set of templates. One may also use a fast moment-based method [18, 26] to provide an initial state for shape alignment.

4.3 Future Directions

The MM algorithm is useful for linearizing objective functionals. In this paper, we have presented it as a technique for linearizing a nonlinear shape prior, yielding a linearized surrogate energy functional. It may prove to be a useful tool for the graph-cut community for the relaxation other useful nonlinear functionals. Finally, the MM and EM algorithms may have use in inference of more-robust likelihood models. In our opinion, it is an avenue worth investigating.

Acknowledgements This material is based upon work supported by the National Science Foundation under Agreement No. 0635561. J.C. and T.C. also acknowledge support from the National Science Foundation through grants DMS-1032131 and DMS-1021818, and from the Army Research Office through grant 58386MA.

Appendix A: Shape Alignment

In order to align the shape templates, we choose the affine transformation $\mathbb{T}(\mathbf{s}) = \alpha \mathbf{R}(\omega)(\mathbf{s} - \mathbf{c})$ that minimizes the shape energy

$$U_{\text{shape}}(\mathbb{T}(\mathbf{c}, \omega, \alpha)) = \underbrace{\int_{\mathbb{R}^d} (\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbb{T}(\mathbf{s})))^2 |\phi_{\Lambda}(\mathbb{T}(\mathbf{s}))|^\lambda ds}_{\text{mass mismatch}} + \underbrace{\int_{\mathbb{R}^d} \delta(\phi_{\Omega}(\mathbf{s})) |\phi_{\Lambda}(\mathbb{T}(\mathbf{s}))|^\lambda ds}_{\text{boundary mismatch}}$$

In this expression, we have rewritten the contour integral on $\partial\Omega$ in terms of an integral over the zero-level set of ϕ_{Ω} . Here, we develop a local Newton-Raphson algorithm for finding the optimal transform. Let us denote the $2 \cdot d$ column vector of transformation parameters $\varphi = [\alpha \ \mathbf{c} \ \omega]$. We estimate φ using the iterative updates

$$\varphi_{n+1} = \varphi_n - [\mathbf{H}_{\varphi}(\varphi_n)]^{-1} \nabla_{\varphi} U_{\text{shape}}(\varphi_n)$$

where

$$\nabla_{\varphi} E = \begin{bmatrix} \frac{\partial}{\partial \alpha} & \nabla_{\mathbf{c}}^T & \nabla_{\omega}^T \end{bmatrix}^T U_{\text{shape}}$$

and

$$\mathbf{H}_{\varphi} = \begin{bmatrix} \frac{\partial^2}{\partial \alpha^2} & \nabla_{\mathbf{c}}^T \frac{\partial}{\partial \alpha} & \nabla_{\omega}^T \frac{\partial}{\partial \alpha} \\ \nabla_{\mathbf{c}} \frac{\partial}{\partial \alpha} & \nabla_{\mathbf{c}} \nabla_{\mathbf{c}}^T & \nabla_{\omega}^T \nabla_{\mathbf{c}} \\ \nabla_{\omega} \frac{\partial}{\partial \alpha} & \nabla_{\omega} \nabla_{\mathbf{c}}^T & \nabla_{\omega} \nabla_{\omega}^T \end{bmatrix} U_{\text{shape}}$$

To populate the matrix, we need to calculate the gradients with respect to α , \mathbf{c} , and ω . For $\lambda > 0$, the first-order derivatives take the form

$$\nabla_{\varphi_i} U_{\text{shape}} = \lambda \int_{\mathbb{R}^d} [(\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}'))^2 + \delta(\phi_{\Omega}(\mathbf{s}))] \times |\phi_{\Lambda}(\mathbf{s}')|^{\lambda-1} \text{sgn}(\phi_{\Lambda}(\mathbf{s}')) \underbrace{\frac{\partial \mathbb{T}}{\partial \varphi_i}}_{p \times d} \underbrace{\nabla \phi_{\Lambda}(\mathbf{s}')}_{d \times 1} ds$$

The sign function comes about from differentiation of the absolute value function in the distributional sense.

For $\lambda \neq 1$, the second-order derivatives constitute tensors that take the form

$$\begin{aligned} \nabla_{\varphi_j}^T \nabla_{\varphi_i} U_{\text{shape}} &= \lambda(\lambda - 1) \int_{\mathbb{R}^d} [(\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}'))^2 + \delta(\phi_{\Omega}(\mathbf{s}))] \\ &\times |\phi_{\Lambda}(\mathbf{s}')|^{\lambda-2} \frac{\partial \mathbb{T}}{\partial \varphi_j} \nabla \phi_{\Lambda}(\mathbf{s}') \nabla^T \phi_{\Lambda}(\mathbf{s}') \left(\frac{\partial \mathbb{T}}{\partial \varphi_i} \right)^T ds \\ &+ 2\lambda \int_{\mathbb{R}^d} \delta(\phi_{\Lambda}(\mathbf{s}')) [(\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}'))^2 + \delta(\phi_{\Omega}(\mathbf{s}))] \\ &\times |\phi_{\Lambda}(\mathbf{s}')|^{\lambda-1} \frac{\partial \mathbb{T}}{\partial \varphi_i} \nabla \phi_{\Lambda}(\mathbf{s}') \nabla^T \phi_{\Lambda}(\mathbf{s}') \left(\frac{\partial \mathbb{T}}{\partial \varphi_j} \right)^T ds \\ &+ \lambda \int_{\mathbb{R}^d} [(\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}'))^2 + \delta(\phi_{\Omega}(\mathbf{s}))] |\phi_{\Lambda}(\mathbf{s}')|^{\lambda-1} \\ &\times \text{sgn} \phi_{\Lambda}(\mathbf{s}') \frac{\partial \mathbb{T}}{\partial \varphi_i} \nabla \nabla \phi_{\Lambda}(\mathbf{s}') \left(\frac{\partial \mathbb{T}}{\partial \varphi_j} \right)^T ds \\ &+ \lambda \int_{\mathbb{R}^d} [(\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}'))^2 + \delta(\phi_{\Omega}(\mathbf{s}))] \\ &\times |\phi_{\Lambda}(\mathbf{s}')|^{\lambda-1} \text{sgn} \phi_{\Lambda}(\mathbf{s}') \underbrace{\frac{\partial^2 \mathbb{T}}{\partial \varphi_j^T \partial \varphi_i}}_{p \times p \times d} \nabla \phi_{\Lambda}(\mathbf{s}') ds. \end{aligned}$$

If $\lambda = 1$, the first term is zero and the following term is added

$$\begin{aligned} &-2\lambda \int_{\mathbb{R}^d} \delta(\phi_{\Lambda}(\mathbf{s}')) (\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}')) \text{sgn}(\phi_{\Lambda}(\mathbf{s}')) \\ &\times \frac{\partial \mathbb{T}}{\partial \varphi_j} \nabla \phi_{\Lambda}(\mathbf{s}') \left(\frac{\partial \mathbb{T}}{\partial \varphi_i} \nabla \phi_{\Lambda}(\mathbf{s}') \right)^T ds. \end{aligned}$$

The contour integrals in these expressions can be calculated using a regularized version of the Dirac delta function such as

$$\delta_{\epsilon}(x) = \frac{1}{\pi} \frac{\epsilon}{\epsilon^2 + x^2} \quad \epsilon \rightarrow 0^+.$$

Similarly, the characteristic function can be interpreted as the Heaviside function acting on the signed-distance shape embedding, which can be approximated using an approximation of the Heaviside function such as

$$H_{\epsilon}(x) = \frac{1}{2} + \frac{1}{\pi} \arctan \frac{x}{\epsilon} \quad \epsilon \rightarrow 0^+.$$

The nonzero transformation derivatives are as follows

$$\begin{aligned} \frac{\partial \mathbb{T}}{\partial \mathbf{c}} &= -\alpha \mathbf{R} & \frac{\partial \mathbb{T}}{\partial \alpha} &= \mathbf{R}(\mathbf{s} - \mathbf{c}) & \frac{\partial \mathbb{T}}{\partial \omega} &= \alpha \frac{\partial \mathbf{R}}{\partial \omega}(\mathbf{s} - \mathbf{c}) \\ \frac{\partial^2 \mathbb{T}}{\partial \omega \partial \mathbf{c}} &= -\alpha \frac{\partial \mathbf{R}}{\partial \omega} & \frac{\partial^2 \mathbb{T}}{\partial \alpha \partial \omega} &= \frac{\partial \mathbf{R}}{\partial \omega}(\mathbf{s} - \mathbf{c}) \end{aligned}$$




Fig. 5 Inputted shape templates for a van. These templates are aligned using the method detailed in Appendix A in order to generate the aligned templates of Fig. 3

$$\frac{\partial^2 \mathbb{T}}{\partial \omega \partial \omega} = \alpha \frac{\partial^2 \mathbf{R}}{\partial \omega \partial \omega} (\mathbf{s} - \mathbf{c}) \quad \frac{\partial^2 \mathbb{T}}{\partial \alpha \partial \mathbf{c}} = -\mathbf{R}.$$

The results of this alignment method are shown in the third column Fig. 3. As an example, the raw input templates for the vans in this figure are shown in Fig. 5.

Appendix B: MM Algorithm for Iterative Graph Cuts

The shape term $\log \sum w_j K(\Omega, \Omega_j)$ can make minimization of the energy difficult, since its formulation involves a sum within a logarithm, making the energy functional nonlinear with respect to the labeling of pixels (into background and foreground) in the image.

To linearize the shape contribution, we will derive a surrogate function with separated terms. A function $f(x|x_k)$, with fixed x_k , is said to majorize a function $g(x)$ at x_k if the following holds [17]

$$g(x) \leq f(x|x_k) \tag{16}$$

$$f(x_k) = g(x_k|x_k). \tag{17}$$

We wish to perform iterative inference by finding a sequence of segmentations $\Omega^{(n+1)} = \arg \min_{\Omega} Q(\Omega|\Omega^{(n)})$, where $Q(\Omega|\Omega^{(n)})$ majorizes Eq. (7). By the descent property of the MM algorithm [17], this sequence converges to a local minimum.

For any convex function $f(x)$, the following holds [17]

$$f\left(\sum_i \alpha_i t_i\right) \leq \sum_i \alpha_i f(t_i).$$

Noting that $-\log(\cdot)$ is convex, we have

$$\begin{aligned} & \underbrace{-\log \sum_{j=1}^J w_j K(\Omega, \Omega_j)}_{\text{shape kernel density}} \\ & \leq -\sum_{j=1}^J \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)} \end{aligned}$$

$$\begin{aligned} & \times \log \left[\frac{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)}{w_j K(\Omega^{(n)}, \Omega_j)} w_j K(\Omega, \Omega_j) \right] \\ & = -\sum_{j=1}^J \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)} \log K(\Omega, \Omega_j) + \text{const} \\ & = \underbrace{\frac{\beta}{2} \sum_{j=1}^J \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)} U_{\text{shape}}(\Omega, \Omega_j)}_{\text{separated log shape kernel density}} + \text{const}, \tag{18} \end{aligned}$$

verifying that the inequality condition (Eq. (16)) holds. In the case that $\Omega = \Omega^{(n)}$, we have

$$\begin{aligned} & -\sum_{j=1}^J \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)} \\ & \times \log \left[\frac{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)}{w_j K(\Omega^{(n)}, \Omega_j)} w_j K(\Omega^{(n)}, \Omega_j) \right] \\ & = -\log \sum_{j=1}^J w_j K(\Omega^{(n)}, \Omega_j), \end{aligned}$$

verifying that the equality condition (Eq. (17)) holds. So the two majorizing conditions (16), (17) are met, and we can minimize our original energy by iteratively minimizing

$$\begin{aligned} & Q(\Omega|\Omega^{(n)}) \\ & = -\sum_{\mathbf{s} \in \Omega} \log p_{\Omega}(I(\mathbf{s})|\boldsymbol{\theta}) \\ & \quad - \sum_{\mathbf{s} \in \Delta} \log p_{\Delta}(I(\mathbf{s})|\boldsymbol{\theta}) - \log p(\boldsymbol{\theta}|\Omega) \\ & \quad + \frac{\beta}{2} \sum_{j=1}^J \frac{w_j K(\Omega^{(n)}, \Omega_j)}{\sum_{k=1}^J w_k K(\Omega^{(n)}, \Omega_k)} U_{\text{shape}}(\Omega, \Omega_j). \tag{19} \end{aligned}$$

Because the distance function can be written as a sum over vertices and edges of a graph, so can Eq. (19). As a result, it is possible to minimize Eq. (19) within the graph cuts framework.

References

1. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(4), 509–522 (2002)
2. Boykov, Y., Jolly, M.: Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In: *Proceedings. Eighth IEEE International Conference on Computer Vision, 2001. ICCV 2001*, vol. 1, pp. 105–112. IEEE Press, New York (2001)
3. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(9), 1124–1137 (2004)

4. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.* **23**(11), 1222–1239 (2001)
5. Chang, J., Chou, T., Brennan, K.: Tracking monotonically advancing boundaries in image sequences using graph cuts and recursive kernel shape priors. *IEEE Trans. Med. Imaging* **13**(5), 1008–1020 (2012)
6. Cremers, D., Grady, L.: Statistical priors for efficient combinatorial optimization via graph cuts. In: *Computer Vision—ECCV 2006*, pp. 263–274 (2006)
7. Cremers, D., Soatto, S.: A pseudo-distance for shape priors in level set segmentation. In: *2nd IEEE Workshop on Variational, Geometric and Level Set Methods in Computer Vision*, pp. 169–176 (2003). Citeseer
8. Cremers, D., Osher, S., Soatto, S.: Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int. J. Comput. Vis.* **69**(3), 335–351 (2006)
9. Cremers, D., Schmidt, F., Barthel, F.: Shape priors in variational image segmentation: convexity, Lipschitz continuity and globally optimal solutions. In: *IEEE Conference on Computer Vision and Pattern Recognition 2008. CVPR 2008*, pp. 1–6. IEEE Press, New York (2008)
10. Dambreville, S., Rathi, Y., Tannenbaum, A.: A framework for image segmentation using shape models and kernel space shape priors. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(8), 1385–1399 (2008)
11. Das, P., Veksler, O., Zavadsky, V., Boykov, Y.: Semiautomatic segmentation with compact shape prior. *Image Vis. Comput.* **27**(1–2), 206–219 (2009)
12. El Zehiry, N., Xu, S., Sahoo, P., Elmaghraby, A.: Graph cut optimization for the Mumford–Shah model. In: *The Seventh IASTED International Conference on Visualization, Imaging and Image Processing*, pp. 182–187. ACTA Press, Calgary (2007)
13. Freedman, D., Drineas, P.: Energy minimization via graph cuts: settling what is possible. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005. CVPR 2005*, vol. 2, pp. 939–946. IEEE Press, New York (2005)
14. Freedman, D., Zhang, T.: Interactive graph cut based segmentation with shape priors. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2005*, vol. 1, pp. 755–762. IEEE Press, New York (2005)
15. Grady, L., Alvino, C.: The piecewise smooth Mumford–Shah functional on an arbitrary graph. *IEEE Trans. Image Process.* **18**(11), 2547–2561 (2009)
16. Heimann, T., Meinzer, H.: Statistical shape models for 3D medical image segmentation: a review. *Med. Image Anal.* **13**(4), 543–563 (2009)
17. Hunter, D., Lange, K.: A tutorial on MM algorithms. *Am. Stat.* **58**(1), 30–37 (2004)
18. Jiang, T., Tomasi, C.: Robust shape normalization based on implicit representations. In: *19th International Conference on Pattern Recognition, 2008. ICPR 2008*, pp. 1–4. IEEE Press, New York (2008)
19. Kohli, P., Torr, P.H.: Dynamic graph cuts for efficient inference in Markov random fields. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(12), 2079–2088 (2007)
20. Lee, S., Park, J.: Nonlinear shape normalization methods for the recognition of large-set handwritten characters. *Pattern Recognit.* **27**(7), 895–902 (1994)
21. Lempitsky, V., Blake, A., Rother, C.: Image segmentation by branch-and-mincut. In: *Computer Vision—ECCV 2008*, pp. 15–29 (2008)
22. Lempitsky, V., Blake, A., Rother, C.: Branch-and-mincut: global optimization for image segmentation with high-level priors. *J. Math. Imaging Vis.* **44**(3), 315–329 (2012)
23. Malcolm, J., Rathi, Y., Tannenbaum, A.: Graph cut segmentation with nonlinear shape priors. In: *IEEE International Conference on Image Processing, 2007. ICIP 2007*, vol. 4. IEEE Press, New York (2007)
24. Malcolm, J., Rathi, Y., Tannenbaum, A.: Multi-object tracking through clutter using graph cuts. In: *IEEE 11th International Conference on Computer Vision, 2007. ICCV 2007*, pp. 1–5. IEEE Press, New York (2007)
25. O’Hagan, A., Forster, J., Kendall, M.: *Bayesian Inference*. Arnold, Sevenoaks (2004)
26. Pei, S.C., Lin, C.N.: Image normalization for pattern recognition. *Image Vis. Comput.* **13**(10), 711–723 (1995)
27. Schindelin, J.: Fiji is just ImageJ—batteries included. In: *Proceedings of the ImageJ User and Developer Conference*, Luxembourg (2008)
28. Silverman, B.: *Density Estimation for Statistics and Data Analysis* vol. 26. Chapman & Hall, London (1986)
29. Slabaugh, G., Unal, G.: Graph cuts segmentation using an elliptical shape prior. In: *IEEE International Conference on Image Processing. ICIP 2005*, vol. 2. IEEE Press, New York (2005)
30. Tabbone, S., Wendling, L.: Binary shape normalization using the radon transform. In: *Discrete Geometry for Computer Imagery*, pp. 184–193. Springer, Berlin (2003)
31. Veksler, O.: Star shape prior for graph-cut image segmentation. In: *Computer Vision—ECCV 2008*, pp. 454–467 (2008)
32. Vineet, V., Narayanan, P.: Cuda cuts: fast graph cuts on the GPU. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPRW’08*, pp. 1–8. IEEE Press, New York (2008)
33. Viola, P., Wells, W. III: Alignment by maximization of mutual information. *Int. J. Comput. Vis.* **24**(2), 137–154 (1997)
34. Vu, N., Manjunath, B.: Shape prior segmentation of multiple objects with graph cuts. In: *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008*, pp. 1–8. IEEE Press, New York (2008)
35. Zhu-Jacquot, J., Zabih, R.: Graph cuts segmentation with statistical shape priors for medical images. In: *Third International IEEE Conference on Signal-Image Technologies and Internet-Based System, SITIS’07*, pp. 631–635. IEEE Press, New York (2007)



Joshua C. Chang has been an NSF postdoctoral fellow at the Mathematical Biosciences institute in Columbus Ohio since 2012. He received his B.S. (2004) and Ph.D. (2012) degrees from UCLA in Mathematics and Biomathematics, respectively, and is trained in applied mathematics and mathematical biology. His research involves applications of stochastic processes and statistical mechanics to problems in mathematical biology and inverse problems.



Tom Chou has been a professor in the departments of Biomathematics and Mathematics at UCLA since 2000. He received his S.B. (1989) and Ph.D. (1995) degrees from MIT and Harvard, respectively, and is trained in condensed matter and statistical physics. His research involves applications of stochastic processes and statistical mechanics to problems in molecular biophysics, cellular biology, and populations.