Published in final edited form as:

J Math Imaging Vis. 2014 March 1; 49(1): 87–97. doi:10.1007/s10851-013-0440-9.

Iterative graph cuts for image segmentation with a nonlinear statistical shape prior

Joshua C. Chang and

Mathematical Biosciences Institute, The Ohio State University, Jennings Hall, 3rd Floor, 1735 Neil Avenue, Columbus, Ohio 43210

Tom Chou

UCLA Biomathematics and Mathematics, BOX 951766, Room 5303 Life Sciences, Los Angeles, CA 90095-1766

Joshua C. Chang: chang.1166@mbi.osu.edu; Tom Chou: tomchou@ucla.edu

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 compactedness or 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].

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 $\mathbf{s}$ to the boundary of Ω . In this paper we take the convention that $\phi_\Omega(\mathbf{s})$ is positive if $\mathbf{s} \in \Omega$, and negative if $\mathbf{s} \in \mathbb{R}^d/\Omega$.

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 ,

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

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

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

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 Discussion).

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

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

where α is a scaling factor, $\mathbf{R}(\boldsymbol{\omega})$ 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 references shapes $\{\Omega_j\}_{j=1}^J$ into a probability distribution by using a kernel of the form

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

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

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

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]:

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

where s_j is the inertial scale of a shape Ω_j as defined in Jiang and Tomasi [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 multimodal.

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:

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

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

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

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

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

where

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

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. 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 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 $\mathcal{G} = (V, \mathcal{E})$:

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

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

$$\begin{aligned}
U_{\text{shape}}(\Omega, \Lambda) = & \sum_{\mathbf{s}} \overbrace{|\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}')|}^{\text{indicator of label mismatch}} |\varphi_{\Lambda}(\mathbf{s}')|^{\lambda} \\
& + \sum_{(\mathbf{s}, \mathbf{u}) \in \mathbf{N}} \overbrace{[(\chi_{\Omega}(\mathbf{s})(1 - \chi_{\Omega}(\mathbf{u})) + \chi_{\Omega}(\mathbf{u})(1 - \chi_{\Omega}(\mathbf{s}))]}^{\text{indicator of } \mathbf{s}, \mathbf{u} \text{ lying across boundary of } \Omega} \times \frac{\pi}{8 \|\mathbf{s} - \mathbf{u}\|} \left| \varphi_{\Lambda} \left(\frac{\mathbf{s}' + \mathbf{u}'}{2} \right) \right|^{\lambda}, \quad (10)
\end{aligned}$$

where $\mathbf{N}$ refers to the set of tuples of nodes that are neighboring in space. The level set function $\varphi_{\Lambda}^{\lambda}$ is taken to be the Euclidean distance embedding of Λ , 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

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

and

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

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

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

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

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

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

$$w(\mathbf{s}, \text{sink}) = -\log p_{\Omega}(I(\mathbf{s})|\boldsymbol{\theta}) + \frac{\beta}{2} \sum_{j=1}^J c_j^{(n)} (1 - \chi_{\Omega_j}(\mathbf{s}')) |\varphi_{\Omega_j}(\mathbf{s}')|^{\lambda}. \quad (14)$$

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_{ε} into the graph, we set the n-links between pairwise neighboring pixels $\mathbf{s}$ and $\mathbf{u}$ to

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

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 $p(I|\theta)$ and uninformative priors for θ (by setting $p(\theta) = 1$). 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 database¹. In the given examples, background objects interfere with the segmentation. Without a 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

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

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.

Acknowledgments

JC acknowledges support by grant number T32GM008185 from the National Institute of General Medical Sciences. JC and TC also acknowledge support from the the National Science Foundation through grants DMS-1032131 and DMS-1021818, and from the Army Research Office through grant 58386MA.

References

1. Belongie S, Malik J, Puzicha J. Shape matching and object recognition using shape contexts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2002; 24(4):509–522.
2. Boykov, Y.; Jolly, M. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. *Computer Vision, 2001; ICCV 2001. Proceedings. Eighth IEEE International Conference on; IEEE*. 2001. p. 105-112.
3. Boykov Y, Kolmogorov V. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2004; 26(9):1124–1137.
4. Boykov Y, Veksler O, Zabih R. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2001; 23(11):1222–1239.
5. Chang J, Chou T, Brennan K. Tracking monotonically advancing boundaries in image sequences using graph cuts and recursive kernel shape priors. *Medical Imaging, IEEE Transactions on*. 2012; 13(5):1008–1020.

6. Cremers, D.; Grady, L. Computer Vision–ECCV 2006. 2006. Statistical priors for efficient combinatorial optimization via graph cuts; p. 263-274.
7. Cremers, D.; Soatto, S. A pseudo-distance for shape priors in level set segmentation. 2nd IEEE workshop on variational, geometric and level set methods in computer vision; Citeseer. 2003. p. 169-176.
8. Cremers D, Osher S, Soatto S. Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *International journal of Computer Vision*. 2006; 69(3):335–351.
9. Cremers, D.; Schmidt, F.; Barthel, F. Shape priors in variational image segmentation: Convexity, lipschitz continuity and globally optimal solutions. *Computer Vision and Pattern Recognition, 2008; CVPR 2008. IEEE Conference on; IEEE*. 2008. p. 1-6.
10. Dambreville, S.; Rath, Y.; Tannenbaum, A. *IEEE transactions on pattern analysis and machine intelligence*. 2008. A framework for image segmentation using shape models and kernel space shape priors; p. 1385-1399.
11. Das P, Veksler O, Zavadsky V, Boykov Y. Semi-automatic segmentation with compact shape prior. *Image and Vision Computing*. 2009; 27(1–2):206–219.
12. El Zehiry, N.; Xu, S.; Sahoo, P.; Elmaghraby, A. The Seventh IASTED International Conference on Visualization, Imaging and Image Processing. ACTA Press; 2007. Graph cut optimization for the Mumford-Shah model; p. 182-187.
13. Freedman, D.; Drineas, P. Energy minimization via graph cuts: Settling what is possible. *Computer Vision and Pattern Recognition, 2005; CVPR 2005. IEEE Computer Society Conference on; IEEE*. 2005. p. 939-946.
14. Freedman, D.; Zhang, T. Interactive graph cut based segmentation with shape priors. *Computer Vision and Pattern Recognition, 2005; CVPR 2005. IEEE Computer Society Conference on; IEEE*. 2005. p. 755-762.
15. Grady L, Alvino C. The piecewise smooth mumford–shah functional on an arbitrary graph. *Image Processing, IEEE Transactions on*. 2009; 18(11):2547–2561.
16. Heimann T, Meinzer H. Statistical shape models for 3d medical image segmentation: A review. *Medical image analysis*. 2009; 13(4):543–563. [PubMed: 19525140]
17. Hunter D, Lange K. A tutorial on MM algorithms. *The American Statistician*. 2004; 58(1):30–37.
18. Jiang, T.; Tomasi, C. Robust shape normalization based on implicit representations. *Pattern Recognition, 2008; ICPR 2008. 19th International Conference on; IEEE*. 2008. p. 1-4.
19. Kohli P, Torr PH. Dynamic graph cuts for efficient inference in markov random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*. 2007; 29(12):2079–2088.
20. Lee S, Park J. Nonlinear shape normalization methods for the recognition of large-set handwritten characters. *Pattern Recognition*. 1994; 27(7):895–902.
21. Lempitsky, V.; Blake, A.; Rother, C. *Computer Vision–ECCV 2008*. 2008. Image segmentation by branch-and-mincut; p. 15-29.
22. Lempitsky V, Blake A, Rother C. Branch-and-mincut: Global optimization for image segmentation with high-level priors. *Journal of Mathematical Imaging and Vision*. 2012:1–15.
23. Malcolm, J.; Rath, Y.; Tannenbaum, A. Graph cut segmentation with nonlinear shape priors. *Image Processing, 2007; ICIP 2007. IEEE International Conference on; IEEE*. 2007.
24. Malcolm, J.; Rath, Y.; Tannenbaum, A. Multi-object tracking through clutter using graph cuts. *Computer Vision, 2007; ICCV 2007. IEEE 11th International Conference on; IEEE*. 2007. p. 1-5.
25. O'Hagan, A.; Forster, J.; Kendall, M. Bayesian inference. Arnold; 2004.
26. Pei SC, Lin CN. Image normalization for pattern recognition. *Image and Vision Computing*. 1995; 13(10):711–723.
27. Schindelin, J. Fiji is just ImageJ – batteries included. *Proceedings of the ImageJ User and Developer Conference; Luxembourg*. 2008.
28. Silverman, B. Density estimation for statistics and data analysis. Vol. 26. Chapman & Hall/CRC; 1986.
29. Slabaugh, G.; Unal, G. Graph cuts segmentation using an elliptical shape prior. *Image Processing, 2005; ICIP 2005. IEEE International Conference on; IEEE*. 2005.

30. Tabbone, S.; Wendling, L. Discrete Geometry for Computer Imagery. Springer; 2003. Binary shape normalization using the radon transform; p. 184-193.
31. Veksler, O. Computer Vision–ECCV 2008. 2008. Star shape prior for graph-cut image segmentation; p. 454-467.
32. Vineet, V.; Narayanan, P. Cuda cuts: Fast graph cuts on the gpu. Computer Vision and Pattern Recognition Workshops, 2008; CVPRW'08. IEEE Computer Society Conference on; IEEE. 2008. p. 1-8.
33. Viola P, Wells W III. Alignment by maximization of mutual information. International journal of computer vision. 1997; 24(2):137–154.
34. Vu, N.; Manjunath, B. Shape prior segmentation of multiple objects with graph cuts. Computer Vision and Pattern Recognition, 2008; CVPR 2008. IEEE Conference on; IEEE. 2008. p. 1-8.
35. Zhu-Jacquot, J.; Zabih, R. Graph cuts segmentation with statistical shape priors for medical images. Signal-Image Technologies and Internet-Based System, 2007; SITIS'07. Third International IEEE Conference on; IEEE. 2007. p. 631-635.

Biographies

Dr. 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.

Dr. Josh 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.

A Shape alignment

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

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

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 $\boldsymbol{\phi} = [\alpha \quad \mathbf{c} \quad \boldsymbol{\omega}]$. We estimate $\boldsymbol{\phi}$ using the iterative updates

$$\boldsymbol{\phi}_{n+1} = \boldsymbol{\phi}_n - [\mathbf{H}_{\boldsymbol{\phi}}(\boldsymbol{\phi}_n)]^{-1} \nabla_{\boldsymbol{\phi}} U_{\text{shape}}(\boldsymbol{\phi}_n)$$

where

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

and

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

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

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

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

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

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

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

$$-2\lambda \int_d \delta(\varphi_{\Lambda}(\mathbf{s}')) (\chi_{\Omega}(\mathbf{s}) - \chi_{\Lambda}(\mathbf{s}')) \text{sgn}(\varphi_{\Lambda}(\mathbf{s}')) \times \frac{\partial \mathbb{T}}{\partial \phi_j} \nabla \varphi_{\Lambda}(\mathbf{s}') \left(\frac{\partial \mathbb{T}}{\partial \phi_i} \nabla \varphi_{\Lambda}(\mathbf{s}') \right)^T d\mathbf{s}.$$

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

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

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

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

The nonzero transformation derivatives are as follows

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

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.

B MM algorithm for iterative graph cuts

The shape term $\log \sum_j w_j K(\Omega, \Omega_j)$ 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]

$$g(x) \leq f(x|x_k) \quad (16)$$

$$f(x_k) = g(x_k|x_k). \quad (17)$$

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

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

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

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

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

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

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.

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 leaf is the 0-level set of the function (red outline).

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.

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 kernel-density shape prior model is able to settle into a shape with the proper configuration tips for both example images. This feature is a result of the kernel-density energy functional being multimodal.

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.

Algorithm 1**MM-Graphcut algorithm**

Obtain initial guess of image intensity model parameters $\theta^{(0)}$.

Obtain initial guess of the segmentation $\mathcal{I}^{(0)}$, by segmentation without a shape prior.

Align and rescale the shape prior templates to $\mathcal{I}^{(0)}$.

$n = 1$

while $U_{\text{shape}}(\mathcal{I}(n), \mathcal{I}^{(n-1)}) > \text{tol}$ **do**

Recalculate the shape weights $c_j^{(n)}$.

Reweight the edges and perform max flow to estimate $\mathcal{I}^{(n+1)}$

Re-estimate the image intensity parameters $\theta^{(n+1)}$ (as needed)

Re-align the shape prior templates to the guess $\mathcal{I}^{(n)}$ (as needed)

$n++$

end while