

Bounded Labeling Function for Global Segmentation of Multi-Part Objects with Geometric Constraints

Masoud S. Nosrati, Shawn Andrews, Ghassan Hamarneh
 Medical Image Analysis Lab, Simon Fraser University, BC, Canada
 {smn6, sda56, hamarneh}@sfu.ca

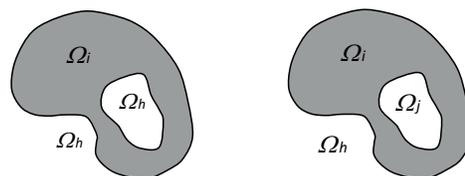
Abstract

The inclusion of shape and appearance priors have proven useful for obtaining more accurate and plausible segmentations, especially for complex objects with multiple parts. In this paper, we augment the popular Mumford-Shah model to incorporate two important geometrical constraints, termed *containment* and *detachment*, between different regions with a specified minimum distance between their boundaries. Our method is able to handle multiple instances of multi-part objects defined by these geometrical constraints using a single labeling function while maintaining global optimality. We demonstrate the utility and advantages of these two constraints and show that the proposed convex continuous method is superior to other state-of-the-art methods, including its discrete counterpart, in terms of memory usage, and metrication errors.

1. Introduction

The piecewise constant Mumford-Shah (MS) model [18] is one of the most popular models in image segmentation. In recent years, many efforts have been made to improve this model in terms of *optimizability*, by simplifying the objective function and formulating it as a convex energy functional [2, 7, 20], and *fidelity*, by making the objective function more faithful to the underlying segmentation tasks. In image segmentation literature, improving objective function fidelity has taken several forms: (i) adding new energy terms, e.g. edge, region, and, shape prior terms [1]; (ii) extending binary segmentation methods to multi-label segmentation [8, 26]; and (iii) incorporating spatial relationships between labels, objects, or object parts [14]. In many image labeling tasks, some geometric relationships are known beforehand, and incorporating this information into a segmentation algorithm improves results.

In this paper, we focus on segmentation tasks where two regions must be separated by a third. Specifically, we focus primarily on the geometric constraint *containment*,



(a) Standard labeling function setting (b) Our setting

Figure 1: The inside vs. outside ambiguity in (a) is resolved by our *containment* constraint in (b).

where one region separates a second region from the background (cf. Fig. 1). Other geometric constraints can also be enforced using the same framework, such as *detachment*, where the background separates two regions. This paper addresses the problem of multi-region segmentation with these two important geometrical constraints, *containment* and *detachment* with a minimum distance (or *thickness*) between regions' boundaries, in a continuous framework while maintaining global optimality. We choose these two geometrical constraints due to their intuitive definitions, descriptive power, and ability to help properly segment regions with weak intensity/color appearance models. Using a continuous framework provides several advantages over discrete methods: **1)** no metrication error; **2)** less memory usage; **3)** efficient parallelizability, and **4)** allowance for sub-pixel resolution.

1.1. Previous works

Improving segmentation via encoding spatial relations between multiple target objects is not new. For example, some methods encode spatial relationships via relative inter-object distances [17] while other works have incorporated fuzzy spatial relationships [9], but those methods do not handle stricter geometric constraints such as containment. Another popular approach is to perform atlas-based segmentation, which has been particularly useful in medical image analysis applications since an atlas encodes the spatial relationships between multiple anatomical structures, or

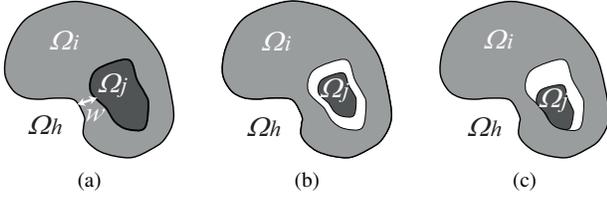


Figure 2: Containment vs. similar configurations (h : background). According to (3), “object i contains object j ” in (a) with $Th(\Omega_h, \Omega_i, \Omega_j) \geq w$, but the relationship between i and j in (b) and (c) is *not* containment.

organs [10]. However, those methods have been designed to encode constraints on a *single* instance of an object in an image, not multiple instances, spatially-recurring throughout the image.

A few recent works have focused on *tiered segmentation* to encode adjacency relationships [12, 21]. Strelakovsky *et al.* [21] proposed a generalized label ordering constraint which can enforce many complex geometric constraints while maintaining convexity. This method requires that the constraint term obey the triangle inequality, a requirement that was later relaxed by introducing a convex relaxation method for non-metric priors [22]. Both [21] and [22] are designed to penalize transitions between adjacent labels, but for meaningful containment and detachment constraints, a minimum *spatial* distance between non-adjacent labels is required. Other recent works have focused on incorporating *topological* constraints into a segmentation objective function. In the continuous domain, some methods incorporate the containment constraint into their segmentation framework by simultaneously evolving a coupled surfaces [19, 23, 25, 28]. However, these works are limited to objects with two surfaces and also are limited to segmenting a single instance of each object in an image. In the discrete domain, Li *et al.* [16] proposed a method to segment nested objects, but their method is limited to star-shaped objects. Delong and Boykov [11] and Ulén *et al.* [24] proposed segmentation methods that encode geometric constraints (including containment) between distinct regions into a graph cut framework. Our work can be viewed as a continuous analogue to these works, providing several advantages, as noted earlier and as will demonstrated in Section 4.

We represent our segmentation using a single continuous labeling function, assigning each region to an interval of label values. We perform segmentation via energy minimization, and we ensure a globally optimal solution using a “functional lifting” technique, similar to what is used by Pock *et al.* [20], to convexify our data energy term by reformulating the problem in a higher dimensional space. This approach exhibits several important qualities, such as the ability to deal with topological changes (e.g. spatially recurring instances of an object), and extendibility to higher

dimensional images.

We introduce the containment and detachment constraints in Section 2. We show how to encode these two constraints in a continuous segmentation framework and show how our formulation can be convexified by functional lifting in Section 2.1. Section 3 explains how the energy is optimized. Different examples as well as comparisons with other popular state-of-the-art methods are given in Section 4, followed by our conclusions in Section 5.

2. Methodology

In this section we explicitly define *containment* and *detachment* and show how we encode them in a MS based model while maintaining global optimality.

We first consider a containment constraint in a 3 region segmentation. We divide the image domain, $\Omega \subset \mathbb{R}^2$, into three non-overlapping parts (Fig. 1(b)): the outside or background region Ω_h , the outer region Ω_i , and the contained region Ω_j , where $\Omega = \Omega_h \cup \Omega_i \cup \Omega_j$. In many binary segmentation applications that use relaxed labeling functions, label values below $1/2$ correspond to background and values above $1/2$ correspond to foreground. We extend this definition as follows. Given a label set $\Gamma = [0, 1]$, we define our labeling function $u : \Omega \rightarrow \Gamma$, such that

$$\begin{aligned} 0 \leq u(\mathbf{x}) < 1/3 &\iff \mathbf{x} \in \Omega_h \\ 1/3 \leq u(\mathbf{x}) < 2/3 &\iff \mathbf{x} \in \Omega_i \\ 2/3 \leq u(\mathbf{x}) < 1 &\iff \mathbf{x} \in \Omega_j. \end{aligned} \quad (1)$$

To precisely define containment, we introduce a function that measures the *thickness* of the outer region Ω_i :

$$Th(\Omega_h, \Omega_i, \Omega_j) = \min_{x_1 \in \Omega_j} \min_{x_2 \in \Omega_h} \|x_1 - x_2\|. \quad (2)$$

We define containment for 3 regions as:

Definition 1 (Containment). We say object i contains object j with thickness w if and only if

$$Th(\Omega_h, \Omega_i, \Omega_j) \geq w. \quad (3)$$

We note that $\Omega = \Omega_h \cup \Omega_i \cup \Omega_j$ is assumed here.

An example is shown in Fig. 2(a) where the light gray object, i , contains the dark gray object, j , with a minimum thickness of w . The related configurations between i and j seen in Fig. 2(b) and (c) are *not* containment based on our definition in (3). However, (b) can be seen as containment in a 4 region segmentation: i contains the interior white region, and the interior white region contains j .

Given an input image¹ $I : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, for objects i and j and the background h , let μ_i , μ_j and μ_h be constant approximations of the regional intensities and define

¹Our method can be extended to vector valued images, e.g. color images or tensor fields, by modifying the data terms, as in [6], [27]

$g_k(\mathbf{x}) = |I(\mathbf{x}) - \mu_k|^2$ for $k = \{h, i, j\}$. To segment I such that i contains j we solve the following energy minimization problem:

$$\begin{aligned} & \arg \min_{u \in D} E(u, \mathbf{g}) \\ & = \arg \min_{u \in D} \int_{\Omega} |\nabla_{\mathbf{x}} u(\mathbf{x})| + \rho(\mathbf{x}, u(\mathbf{x}), \mathbf{g}) d\mathbf{x}, \quad (4) \\ & D = \left\{ u \mid u(\mathbf{x}) = 0 \text{ for } \mathbf{x} \in \partial\Omega \text{ and} \right. \\ & \quad \left. Th(\Omega_h, \Omega_i, \Omega_j) \geq w \right\}. \end{aligned}$$

Here $\nabla_{\mathbf{x}}$ is the gradient in x and y directions, $\mathbf{g} = (g_h, g_i, g_j)$, and $\rho(\mathbf{x}, u(\mathbf{x}), \mathbf{g}) : \Omega \rightarrow \mathbb{R}^+$ is a non-negative data term that encourages u to satisfy (1), e.g.

$$\rho(\mathbf{x}, u(\mathbf{x}), \mathbf{g}) = \begin{cases} g_h(\mathbf{x}) & \text{if } 0 \leq u(\mathbf{x}) < 1/3 \\ g_i(\mathbf{x}) & \text{if } 1/3 \leq u(\mathbf{x}) < 2/3 \\ g_j(\mathbf{x}) & \text{if } 2/3 \leq u(\mathbf{x}) < 1 \end{cases}. \quad (5)$$

For convenience, we often let ρ be a function of \mathbf{g} implicitly, and write $\rho(\mathbf{x}, u(\mathbf{x}))$.

Constraining u to D , $E(u, \mathbf{g})$ ensures that object j and object h have no shared boundaries, resulting in j being contained in i . In other words, the segmentation corresponding to u cannot abruptly change from object j to object h , and thus the value of u cannot change from $u \geq 2/3$ to $u \leq 1/3$ in a distance less than w .

This leads us to the fact that the constraint $Th(\Omega_h, \Omega_i, \Omega_j) \geq w$ can be replaced by the more convenient constraint $|\nabla_{\mathbf{x}} u| \leq \frac{1}{3w}$, which limits the rate that u can change spatially. This lets us rewrite D as

$$D = \left\{ u \mid u(\mathbf{x}) = 0 \text{ for } \mathbf{x} \in \partial\Omega, |\nabla_{\mathbf{x}} u| \leq \frac{1}{3w} \right\}, \quad (6)$$

To better understand the equivalence of these constraints, let's consider the example shown in Fig. 3. Here, a black and white image is segmented into three regions, with μ_h corresponding to black pixels, μ_i corresponding to (non-existent) gray pixels, and μ_j corresponding to white pixels. Fig. 3(b) illustrates the labeling function u corresponding to the segmentation in Fig. 3(a), with no thickness constraint. Here, u becomes discontinuous (unbounded $|\nabla_{\mathbf{x}} u|$) in order to avoid assigning any pixels to the exterior object i . By enforcing the restriction $|\nabla_{\mathbf{x}} u| \leq \frac{1}{3w}$, u is not able to jump from $< 1/3$ (background) to $\geq 2/3$ (white object) in less than distance w (Fig. 3(d)). By restricting u , we force the white object to be contained by an intermediate region of thickness of w (cf. Fig. 3(c)). We note that if w is large enough, the energy increase from hallucinating the intermediate region in Fig. 3(c) will become greater than the energy increase from not segmenting the inner white region, and the result will be $u = 0$ across the image.

Using a similar formulation to containment, we can incorporate a constraint ensuring two regions are detached.

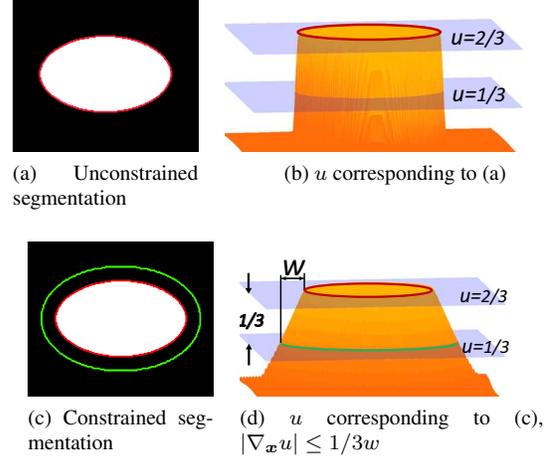


Figure 3: Constrained vs. unconstrained labeling function u . In (a), the 3-region labeling function u , as defined in (1), is used to segment the white object from the black background, while the intermediate region has zero thickness. In (b), we see the u corresponding to (a). Without a thickness constraint, u is allowed to become discontinuous, skipping over the interval corresponding to the intermediate region, i.e. Ω_i in (1). In (c), we see how the segmentation changes when the thickness constraint (3) is enforced, with the intermediate region being hallucinated around the white object. In (d), we see the u corresponding to (c). By bounding the rate of change of u , a band of thickness w must be assigned to the intermediate region.

Definition 3 (Detachment). Object i and object j are detached with thickness w if and only if

$$Th(\Omega_i, \Omega_h, \Omega_j) \geq w. \quad (7)$$

In other words, object j does not share a boundary with object i , thereby ensuring a separation between i and j by enforcing the labeling function u to pass through the background, h , as it travels from i to j . To encode detachment, we simply swap g_h and g_i in (4).

Now, we note that the data term (5) is not convex, making standard gradient descent based optimization schemes prone to local minima. In the next section, we discuss how to convexify (4) and thus find a globally minimizing segmentation.

2.1. Function convexification

We use a “functional lifting” technique similar to the one proposed by Pock *et al.* [20] (motivated by Ishikawa’s work in the discrete Markov random field setting [15]) to transfer our energy functional to a higher dimensional space, where it becomes convex. The objective is to solve the following minimization problem

$$\arg \min_{u \in D} \left\{ \int_{\Omega} |\nabla_{\mathbf{x}} u(\mathbf{x})| d\mathbf{x} + \int_{\Omega} \rho(\mathbf{x}, u(\mathbf{x})) d\mathbf{x} \right\}. \quad (8)$$

The first term is a convex total variation (TV) term, but the second term $\rho(\mathbf{x}, u(\mathbf{x}))$ can be non-convex. To lift the original energy to a higher dimensional space, we represent u in terms of its *super level set*, $\varphi : \Omega \times \Gamma \rightarrow \{0, 1\}$ by

$$\varphi(\mathbf{x}, \gamma) = \mathbf{1}_{\{u \geq \gamma\}}(\mathbf{x}) = \begin{cases} 1 & \text{if } u(\mathbf{x}) \geq \gamma \\ 0 & \text{otherwise} \end{cases}. \quad (9)$$

u can be recovered from φ via the layer-cake formula

$$u(\mathbf{x}) = \int_{\Gamma} \varphi(\mathbf{x}, \gamma) d\gamma. \quad (10)$$

The TV term in (8) can be re-written with respect to φ using the generalized co-area formula [13]:

$$\int_{\Omega} |\nabla_{\mathbf{x}} u(\mathbf{x})| d\mathbf{x} = \int_{\Omega} \int_{\Gamma} |\nabla_{\mathbf{x}} \varphi(\mathbf{x}, \gamma)| d\gamma d\mathbf{x}. \quad (11)$$

By observing that $\delta(u(\mathbf{x}) - \gamma) \equiv |\partial_{\gamma} \varphi(\mathbf{x}, \gamma)|^2$, where $\delta(\cdot)$ is the Dirac delta function, the data term in (8) can then be re-written as [20]:

$$\begin{aligned} \rho(\mathbf{x}, u(\mathbf{x})) &= \int_{\Gamma} \rho(\mathbf{x}, \gamma) \delta(u(\mathbf{x}) - \gamma) d\gamma \\ &= \int_{\Gamma} \rho(\mathbf{x}, \gamma) |\partial_{\gamma} \varphi(\mathbf{x}, \gamma)| d\gamma. \end{aligned} \quad (12)$$

Now, using (11) and (12), the equivalent form of (8) is

$$\arg \min_{\{\varphi | u \in D\}} \int_{\Omega} \int_{\Gamma} |\nabla_{\mathbf{x}} \varphi(\mathbf{x}, \gamma)| + \rho(\mathbf{x}, \gamma) |\partial_{\gamma} \varphi(\mathbf{x}, \gamma)| d\gamma d\mathbf{x}. \quad (13)$$

Note that the non-convex function ρ does not depend on φ any more and (13) is convex in φ . As the last stage of convexification, φ , (9), is relaxed so it varies continuously between zero and one, i.e. $\varphi \in \Omega \times \Gamma \rightarrow [0, 1]$. To recover u , we threshold φ and apply (10).

Now, it is not immediately clear what form the constraint $\{\varphi | u \in D\}$ will take. If we ignore the thickness constraint from D , we could use $\varphi \in \mathcal{D}_1$ where

$$\mathcal{D}_1 = \{\varphi | \varphi(\mathbf{x}, 0) = 1, \varphi(\mathbf{x}, 1) = 0, \partial_{\gamma} \varphi \leq 0\}. \quad (14)$$

It is evident from (9) that every φ constructed from the super level sets of some u are in \mathcal{D}_1 .

We now present a theorem describing how to enforce the thickness constraint in φ . Let $\nabla_3 \varphi$ be the gradient of φ in all components, i.e. $\nabla_3 \varphi = (\partial_x \varphi, \partial_y \varphi, \partial_{\gamma} \varphi)$.

Theorem 1. *If*

$$\frac{|\nabla_{\mathbf{x}} \varphi|}{|\partial_{\gamma} \varphi|} \leq \frac{1}{3w}, \quad (15)$$

²From (9), it is observed that the derivative of φ with respect to γ is zero everywhere except where φ changes, i.e. $u(\mathbf{x}) = \gamma$.

then constraint $|\nabla_{\mathbf{x}} u| \leq \frac{1}{3w}$ is satisfied by any u constructed by thresholding φ at some value and applying (10).

Proof: The gradient of $\nabla_3 \varphi$ at point (x, y, γ) is perpendicular to its level set surface at that point. That is, if we let L_{ν} to be the ν -level surface of φ , let v_1 be tangent to L_{ν} at (x, y, γ) , and let $v_2 = \nabla_3 \varphi(x, y, \gamma)$, then $\langle v_1, v_2 \rangle = 0$.

Based on (9), u is a level set of φ no matter where φ is thresholded. The standard technique for finding a vector tangent to the surface defined by u is to choose a unit vector in x and y and set the γ component equal to the rate of change of u in the chosen direction. Specifically, this means

$$v = \left(\frac{\partial_x u}{|\nabla_{\mathbf{x}} u|}, \frac{\partial_y u}{|\nabla_{\mathbf{x}} u|}, |\nabla_{\mathbf{x}} u| \right) \quad (16)$$

is tangent to the surface u , and thus is orthogonal to $\nabla_3 \varphi$ at any point $(x, y, u(x, y))$. Thus

$$0 = \langle v, \nabla_3 \varphi \rangle \quad (17)$$

$$= \frac{\langle \nabla_{\mathbf{x}} u, \nabla_{\mathbf{x}} \varphi \rangle}{|\nabla_{\mathbf{x}} u|} + \partial_{\gamma} \varphi |\nabla_{\mathbf{x}} u| \quad (18)$$

$$|\nabla_{\mathbf{x}} u| = -\frac{\langle \nabla_{\mathbf{x}} u, \nabla_{\mathbf{x}} \varphi \rangle}{|\nabla_{\mathbf{x}} u|} \cdot \frac{1}{\partial_{\gamma} \varphi} \quad (19)$$

$$|\nabla_{\mathbf{x}} u| \leq \frac{|\nabla_{\mathbf{x}} \varphi|}{|\partial_{\gamma} \varphi|} \leq \frac{1}{3w}, \quad (20)$$

where the last step uses $\frac{\langle a, b \rangle}{|a|} \leq |b|$. \square

The objective now is to solve:

$$\arg \min_{\varphi \in \mathcal{D}_2} \int_{\Omega} \int_{\Gamma} |\nabla_{\mathbf{x}} \varphi| + \rho |\partial_{\gamma} \varphi| d\gamma d\mathbf{x}, \quad (21)$$

$$\mathcal{D}_2 = \left\{ \varphi \in \mathcal{D}_1 \mid |\nabla_{\mathbf{x}} \varphi| \leq \frac{|\partial_{\gamma} \varphi|}{3w} \right\}. \quad (22)$$

3. Optimization

Due to the discontinuity in the Euler-Lagrange equation of (21), we use a primal-dual algorithm, [4, 5], to obtain the global solution. Defining the dual variable $\mathbf{p} = (p_1, p_2, p_3)^T$, we can write the total variation part of (21) as

$$|\nabla_{\mathbf{x}} \varphi(\mathbf{x}, \gamma)| = \max_{|\mathbf{p}'| \leq 1} \langle \mathbf{p}'(\mathbf{x}, \gamma), \nabla_{\mathbf{x}} \varphi(\mathbf{x}, \gamma) \rangle, \quad (23)$$

where $\mathbf{p}' = (p_1, p_2)$. In (21), recalling that $-1 \leq \partial_{\gamma} \varphi \leq 0$, it can be easily shown

$$\rho(\mathbf{x}, \gamma) |\partial_{\gamma} \varphi(\mathbf{x}, \gamma)| = \max_{-p_3 \leq \rho} p_3(\mathbf{x}, \gamma) \partial_{\gamma} \varphi(\mathbf{x}, \gamma). \quad (24)$$

Using these dual variables, the optimization problem of (21) becomes a min-max problem

$$\arg \min_{\varphi \in \mathcal{D}_2} \left\{ \int_{\Omega} \int_{\Gamma} \left(\max_{|\mathbf{p}'| \leq 1} \langle \mathbf{p}'(\mathbf{x}, \gamma), \nabla_{\mathbf{x}} \phi(\mathbf{x}, \gamma) \rangle + \max_{-p_3 \leq \rho} p_3(\mathbf{x}, \gamma) \partial_{\gamma} \varphi(\mathbf{x}, \gamma) \right) d\gamma d\mathbf{x} \right\} \quad (25)$$

$$= \arg \min_{\varphi \in \mathcal{D}_2} \left\{ \int_{\Omega} \int_{\Gamma} \max_{\mathbf{p} \in \mathcal{C}} \langle \nabla_3 \varphi, \mathbf{p} \rangle \right\}, \quad (26)$$

Combining the constraints in (23) and (24) gives [20]

$$\mathcal{C} = \mathcal{C}_1 = \left\{ \mathbf{p} \in \mathbb{R}^3 \mid \sqrt{p_1^2 + p_2^2} \leq 1, p_3 \geq -\rho \right\}, \quad (27)$$

a cylinder with radius one, open on one end (cf. Fig. 4(a)). Recalling that $\partial_{\gamma} \varphi \leq 0$, p_3 that maximizes (26) is always as negative as possible, i.e. $p_3 = -\rho$. Thus for $\mathcal{C} = \mathcal{C}_1$, the maximizing \mathbf{p} is always on the circle at the base of the cylinder, highlighted in red in Fig. 4(a).

Unfortunately, while minimizing φ over \mathcal{D}_1 can be done efficiently, the gradient magnitude constraint \mathcal{D}_2 imposes (in order to enforce the thickness constraint) makes the minimization of (26) for φ difficult.

We will show that by moving the burden of enforcing the thickness constraint to the dual space, the optimization problem becomes much easier. Specifically, we will constrain φ to \mathcal{D}_1 , and introduce a new space for the dual variables, \mathcal{C}_2 , satisfying the following properties:

1. \mathcal{C}_2 is convex;
2. if φ satisfies (15) then

$$\max_{\mathbf{p} \in \mathcal{C}_2} \langle \nabla_3 \varphi, \mathbf{p} \rangle = \max_{\mathbf{p} \in \mathcal{C}_1} \langle \nabla_3 \varphi, \mathbf{p} \rangle; \quad (28)$$

3. if φ does not satisfy (15) then (26) becomes arbitrarily large and that choice of φ will be disallowed:

$$\max_{\mathbf{p} \in \mathcal{C}_2} \langle \nabla_3 \varphi, \mathbf{p} \rangle = \infty. \quad (29)$$

Note that since (15) must be enforced at each spatial location, conditions 2 and 3 must also hold at each spatial location, thus the integrals are dropped from (26).

Theorem 2. *A set that satisfies the three above conditions is the truncated cone defined by*

$$\mathcal{C}_2 = \left\{ \mathbf{p} \mid \sqrt{p_1^2 + p_2^2} \leq 3w(p_3 + \rho) + 1, p_3 \geq -\rho \right\}. \quad (30)$$

Proof. As a truncated cone, \mathcal{C}_2 is convex, so condition 1 is satisfied. To show 2 and 3 are satisfied, we will determine, for a given $\nabla_3 \varphi$, the value of

$$\max_{\mathbf{p} \in \mathcal{C}_2} \langle \nabla_3 \varphi, \mathbf{p} \rangle. \quad (31)$$

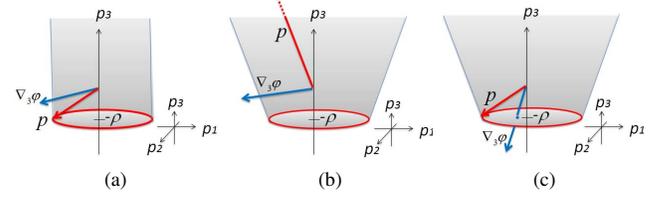


Figure 4: Valid sets for dual variables \mathbf{p} . (a) Set \mathcal{C}_1 [20] (without any geometrical constraint). (b-c) Set \mathcal{C}_2 ; the truncated cone that impose our constraint (15). (b) When (15) is not satisfied \mathbf{p} becomes ∞ . (c) When (15) is satisfied we obtain the same solution as (a).

We first note that the circle at the base of the truncated cone is the same as the circle at the base of the cylinder defined by \mathcal{C}_1 . Thus, if the \mathbf{p} maximizing (31) lies on this circle (the red set in Fig. 4(c)), then (28) is satisfied. We now simplify this problem by reducing the possible forms for $\nabla_3 \varphi$ and \mathbf{p} .

Since the condition (15) and \mathcal{C}_2 are rotationally symmetric with respect to the first two components, we can assume the second component of $\nabla_3 \varphi$ is 0, i.e. rotating $\nabla_3 \varphi$ does not change (31). Further, if $\nabla_3 \varphi$ is scaled by some constant, the maximizing \mathbf{p} is not changed. Thus, we let $\nabla_3 \varphi = (\ell, 0, a)$, where $\ell = |\nabla_{\mathbf{x}} \varphi|$ is held constant and $\partial_{\gamma} \varphi = a \leq 0$ is the only free parameter.

The maximizing \mathbf{p} is always on the boundary of \mathcal{C}_2 , since scaling \mathbf{p} until it is on the boundary increases (31). Further, the first two components of the maximizing \mathbf{p} must align with the first two components of $\nabla_3 \varphi$, so $p_2 = 0$.

If \mathbf{p} is on the circle at the base of the cone, it would take the form $\mathbf{p} = (b, 0, -\rho)$, where $0 \leq b \leq 1$ is a free variable, and (31) becomes

$$\max_{\mathbf{p}} \langle \mathbf{p}, \nabla_3 \varphi \rangle = \max_b (b\ell - a\rho), \quad (32)$$

which is maximized for $b = 1$, i.e. a \mathbf{p} on the edge of the circle (the red set in Fig. 4(c)), so no maximizing \mathbf{p} is on the inner part of the circle.

If \mathbf{p} is on the surface of the cone, it would take the form $\mathbf{p} = (3w(b + \rho) + 1, 0, b)$ (from (30)), where $b \geq -\rho$ is a free variable. This gives

$$\max_{\mathbf{p}} \langle \mathbf{p}, \nabla_3 \varphi \rangle = \max_b (\ell(3wb + 3w\rho + 1) + ab) \quad (33)$$

$$= \max_b (b(3w\ell + a) + 3w\rho\ell + \ell). \quad (34)$$

If $(3w\ell + a) \leq 0$, (34) is maximized by minimizing b , i.e. $b = -\rho$. This corresponds to \mathbf{p} on the edge of the circle (the red set in Fig. 4(c)), which implies (28) is satisfied. If $(3w\ell + a) > 0$, (34) is maximized by $b \rightarrow \infty$, and (31) gets arbitrarily large, satisfying (29) (cf. Fig. 4(b)).

To complete the proof, we note that $(3w\ell + a) \leq 0$ is

equivalent to the thickness constraint (15):

$$\begin{aligned} (3wl + a) &\leq 0 \\ 3w|\nabla_{\mathbf{x}}\varphi| &\leq -\nabla_{\gamma}\varphi \end{aligned} \quad (35)$$

$$|\nabla_{\mathbf{x}}\varphi| \leq \frac{|\nabla_{\gamma}\varphi|}{3w}. \quad (36)$$

□

To find the optimal solution for $\varphi \in \mathcal{D}_1$ and $\mathbf{p} \in \mathcal{C}_2$, we perform the following primal and dual steps [3]:

Primal step Find the minimum φ for a fixed \mathbf{p} :

$$\varphi^{k+1} = \arg \min_{\varphi \in \mathcal{D}_1} \int_{\Omega \times \Gamma} \langle \nabla_3 \varphi, \mathbf{p}^k \rangle + \frac{(\varphi - \varphi^k)^2}{2t_{\varphi}}. \quad (37)$$

Dual step Find the maximum \mathbf{p} for a fixed φ :

$$\mathbf{p}^{k+1} = \arg \max_{\mathbf{p} \in \mathcal{C}_2} \int_{\Omega \times \Gamma} \langle \nabla_3 \varphi^{k+1}, \mathbf{p} \rangle + \frac{(\mathbf{p} - \mathbf{p}^k)^2}{2t_p}. \quad (38)$$

t_{φ} and t_p are the step sizes of the primal and dual update equations (here we used $t_{\varphi} = 0.01$ and $t_p = 5$). The solutions of (37) and (38) are derived from the Euler-Lagrange equations and projecting the obtained solutions to their valid sets (cf. Algorithm 1 in Appendix A).

After finding the global solution for the relaxed optimization problem, the labeling function u is recovered by thresholding φ at 0.5 and applying (10). It can be proven that thresholded minimizers of the relaxed problem are the minimizers of the binary problem (13) (cf. Appendix B).

We note that our framework can be extended to multiple nested regions by dividing Γ into more than 3 intervals in (1). The thickness constraint between consecutive regions can be set by adjusting w and the interval length for each region in (1).

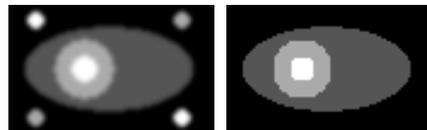
4. Experiments and discussion

In this section, we demonstrate advantages of our work over popular state-of-the-art segmentation methods and compare our framework with the analogous discrete work of Delong and Boykov (DB) [11].

4.1. Synthetic data

In our first experiment, we compare our method with DB in terms of memory usage and metrication error on a simple synthetic example.

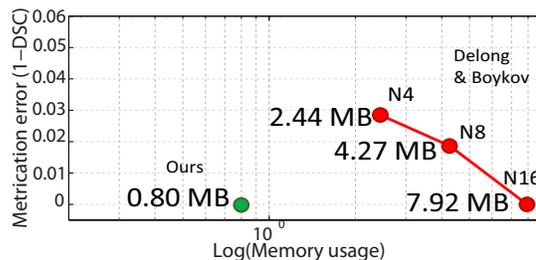
Metrication error is defined as the artifacts which appear in graph-based segmentation methods due to penalizing region boundaries only across axis aligned edges. In Fig. 5, the goal is to segment the three-region object from the background. Fig. 5(b-d) resulted from DB’s method for 4, 8 and 16 graph connectivity. Note the metrication artifacts in Fig. 5(b-c). Increasing the graph connectivity reduces metrication error, but also increases memory usage.



(a) Original image (b) GC: 4-connected



(c) GC: 8-connected (d) GC: 16-connected (e) Our result



(f) Metrication error vs. memory usage

Figure 5: Synthetic three-region object+background segmentation. (b-d) DB graph cuts based method [11] with different connectivities. (e) Our segmentation results. (f) Metrication error vs. memory usage: red curve: GC-based method; green circle: our method.

Memory usage of our method and the graph-based methods is seen in Fig. 5(f). The red curve in Fig. 5(f) illustrates the metrication error (1 - Dice similarity coefficient(DSC)) vs. memory usage of [11] for 4, 8 and 16 connectivity, while the green circle represents our method. Here, removing the metrication error in the graph-based method requires 16 connectivity, even for these smooth objects. This requires ~ 10 times more memory than our method (0.80 vs. 7.92 MB).

4.2. Microscopy/Histology cell segmentation

We applied our method to a set of 20 different histology and microscopy images consist of multiple instances of multi-region cells. In these experiments, we show how containment and detachment with thickness constraints are useful for cell segmentation and compare our method with DB [11] in terms of memory usage and metrication error.

Fig. 6 shows a typical microscopy image with multiple cells, where nuclei are typically contained inside a cell membrane. Fig. 6(b) and (c) show the segmentation results with a thickness of $w = 2$ pixels for DB (with 4-connectivity) and our method, respectively. Metrication error can be clearly seen in Fig. 6(b). Our method also requires less memory (7.91 MB vs. 33.90 MB).

Changing the thickness of the containing region allows

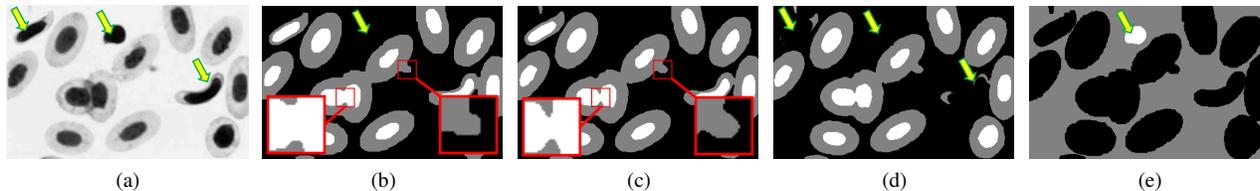


Figure 6: Cell segmentation in a microscopy imagery. (a) Original image, 250×395 pixels. Arrows show abnormal cells. (b) Result of [11], 33.90 MB. (c) Our result (thickness= 2 pixels, 7.91 MB). (d) Our result to segment only normal looking (elliptical) cells (thickness=10 pixels, 7.91 MB). Note that [11] needs ~ 313.41 MB extra memory (347.31 MB in total) to impose thickness of 10 pixels while the memory usage of our method is independent of thickness constraint. (e) Segmenting isolated nuclei by imposing *detachment* constraint (7).

Table 1: DSC and memory usage comparison (20 images).

Method	DSC	Memory (MB)
No containment [20]	0.6478 ± 0.06	8.9 ± 7.2
DB [11]	0.9065 ± 0.08	180.0 ± 204.7
Ours	0.9158 ± 0.07	8.9 ± 7.2

us to control which objects are segmented. In Fig. 6(a), to exclude the abnormal cells (arrows) from the segmentation, we increase the thickness of the outer region (membrane) from 2 to 10 pixels, resulting in Fig. 6(d). In DB, increasing the thickness requires more edges be added to the underlying graph, and increasing the thickness from 2 to 10 pixels requires an extra ~ 313.41 MB memory, an almost 10-fold increase. Thickness can be increased in our method by simply changing the value of w .

Fig. 6 (e) demonstrates the usage of a detachment constraint (cf. (7)), identifying nuclei that are not surrounded by a cell membrane.

Fig. 8 shows results for 7 other images, comparing segmentations generated using a continuous method without a containment constraint [20], DB with 4 connectivity, and our method. These results illustrate the importance of a containment constraint (first column) and also show the effects of metrication error (second column). Table 1 summarizes the mean accuracy and memory usage of the 3 methods across all 20 images, and Fig. 7 compares the memory usage vs. image size of our method and DB across all 20 images. We note that some of the images (e.g. the bottom two rows in Fig. 8) have 4 regions segmented, and Fig. 7 shows that the memory usage in graph-based methods tends to increase proportionally more than our continuous method in these cases.

On average, our method converges after 200 iterations for a 256×256 image. Using non-optimized MATLAB code on a standard 2.3 GHz CPU with 6GB RAM, the graph cuts-based method [11] tends to run 2-3 times faster than our method but with more memory usage.

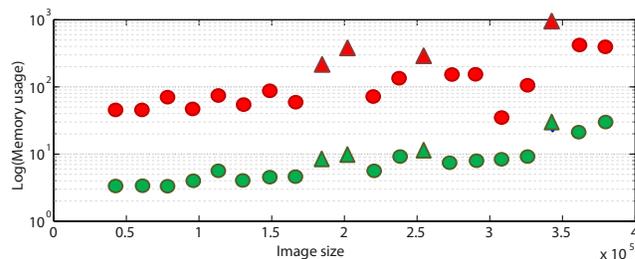


Figure 7: Memory efficiency: DB (in red) vs. ours (in green). \circ : 3-region segmentation; Δ : 4-region segmentation. Memory usage ratio (DB/ours) for 3 regions: 14.63 ± 4.52 and for 4 regions: 32.40 ± 8.35 .

5. Conclusion

We introduced a variational framework to augment the conventional MS model for segmenting multi-region objects. We proposed a labeling function that allows us to enforce useful geometric constraints such as containment and detachment. By using this framework, a user can easily apply high level intuitive geometric constraints to improve segmentation results without the need for a deep understanding of how the method works. Our method compares favorably with analogous graph cuts-based methods in terms of metrication error and memory efficiency.

References

- [1] X. Bresson, P. Vandergheynst, and J. Thiran. A variational model for object segmentation using boundary information and shape prior driven by the Mumford-Shah functional. *IJCV*, 68:145–162, 2006.
- [2] E. S. Brown, T. F. Chan, and X. Bresson. Completely convex formulation of the Chan-Vese image segmentation model. *IJCV*, 98(1):103–121, 2012.
- [3] A. Chambolle. An algorithm for total variation minimization and applications. *JMIV*, 20(1):89–97, 2004.
- [4] A. Chambolle. Total variation minimization and a class of binary MRF models. In *EMMCVPR*, pages 136–152, 2005.
- [5] T. Chan, G. Golub, and P. Mulet. A nonlinear primal-dual method for total variation-based image restoration. *SIAM J. Sci. Comput.*, 20(6):1964–1977, 1999.

- [6] T. Chan, B. Sandberg, and L. Vese. Active contours without edges for vector-valued images. *J. Visual Commun. Image Represent.*, 11(2):130–141, 2000.
- [7] T. F. Chan, S. Esedoglu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM J. Appl. Math.*, 66(5):1632–1648, 2006.
- [8] G. Chung and L. A. Vese. Image segmentation using a multilayer level-set approach. *Comput. Visual. Sci.*, 12(6):267–285, 2009.
- [9] O. Colliot, O. Camara, and I. Bloch. Integration of fuzzy spatial relations in deformable models—application to brain MRI segmentation. *Pattern Recognit.*, 39(8):1401–1414, 2006.
- [10] T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *IEEE TPAMI*, pages 484–498, 1998.
- [11] A. Delong and Y. Boykov. Globally optimal segmentation of multi-region objects. In *ICCV*, pages 285–292, 2009.
- [12] P. F. Felzenszwalb and O. Veksler. Tiered scene labeling with dynamic programming. In *CVPR*, pages 3097–3104, 2010.
- [13] W. Fleming and R. Rishel. An integral formula for total gradient variation. *Archiv der Mathematik*, 11(1):218–222, 1960.
- [14] S. Gould, J. Rodgers, D. Cohen, G. Elidan, and D. Koller. Multi-class segmentation with relative location prior. *IJCV*, 80(3):300–316, 2008.
- [15] H. Ishikawa. Exact optimization for Markov random fields with convex priors. *IEEE TPAMI*, 25(10):1333–1336, 2003.
- [16] K. Li, X. Wu, D. Chen, and M. Sonka. Optimal surface segmentation in volumetric images—a graph-theoretic approach. *IEEE TPAMI*, 28(1):119–134, 2006.
- [17] A. Litvin and W. C. Karl. Coupled shape distribution-based segmentation of multiple objects. In *IPMI*, pages 345–356, 2005.
- [18] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Commun. Pure Appl. Math.*, 42(5):577–685, 1989.
- [19] N. Paragios. A variational approach for the segmentation of the left ventricle in cardiac image analysis. *IJCV*, 50(3):345–362, 2002.
- [20] T. Pock, T. Schoenemann, G. Graber, H. Bischof, and D. Cremers. A convex formulation of continuous multi-label problems. In *ECCV*, pages 792–805. 2008.
- [21] E. Strelakovsky and D. Cremers. Generalized ordering constraints for multilabel optimization. pages 2619–2626, 2011.
- [22] E. Strelakovsky, C. Nieuwenhuis, and D. Cremers. Non-metric priors for continuous multilabel optimization. In *ECCV*, pages 208–221. 2012.
- [23] E. Ukwatta, J. Yuan, M. Rajchl, and A. Fenster. Efficient global optimization based 3D carotid AB-LIB MRI segmentation by simultaneously evolving coupled surfaces. In *MIC-CAI*, pages 377–384. 2012.
- [24] J. Ulén, P. Strandmark, and F. Kahl. An efficient optimization framework for multi-region segmentation based on Lagrangian duality. *IEEE TMI*, 2012.
- [25] A. Vazquez-Reina, E. Miller, and H. Pfister. Multiphase geometric couplings for the segmentation of neural processes. In *CVPR*, pages 2020–2027, 2009.

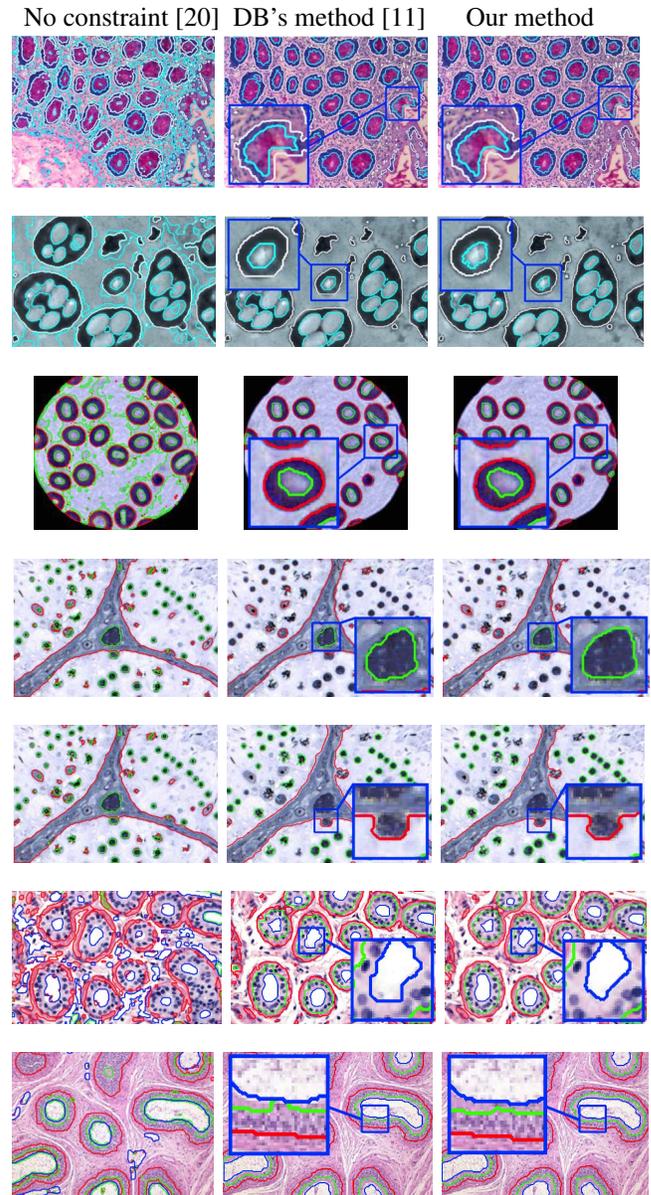


Figure 8: Incorporating geometrical constraint into the segmentation of histology and microscopy images.

- [26] L. Vese and T. Chan. A multiphase level set framework for image segmentation using the Mumford and Shah model. *IJCV*, 50(3):271–293, 2002.
- [27] Z. Wang and B. Vemuri. Tensor field segmentation using region based active contour model. In *ECCV*, pages 304–315. 2004.
- [28] X. Wu, X. Dou, A. Wahle, and M. Sonka. Region detection by minimizing intraclass variance with geometric constraints, global optimality, and efficient approximation. *IEEE TMI*, 30(3):814–827, 2011.