

Hausdorff Distance Constraint for Multi-surface Segmentation

Frank R. Schmidt¹ and Yuri Boykov²

¹ Université Paris Est, France

² University of Western Ontario, Canada

Abstract. It is well known that multi-surface segmentation can be cast as a multi-labeling problem. Different segments may belong to the same semantic object which may impose various inter-segment constraints [1]. In medical applications, there are a lot of scenarios where upper bounds on the Hausdorff distances between subsequent surfaces are known. We show that incorporating these priors into multi-surface segmentation is potentially NP-hard. To cope with this problem we develop a submodular-supermodular procedure that converges to a locally optimal solution well-approximating the problem. While we cannot guarantee global optimality, only feasible solutions are considered during the optimization process. Empirically, we get useful solutions for many challenging medical applications including MRI and ultrasound images.

1 Introduction

This paper addresses the problem of joint segmentation of multiple interacting surfaces. Inter-dependence between segmentation boundaries is relevant when extracting complex multi-part objects, e.g. organs in biomedical imaging. This problem can be cast as a multi-labeling problem [2,1]. While Potts model [3] would be the most straightforward approach, the intensity appearance models for each label/segment are not discriminative enough and additional geometric constraints for the labels/segments are necessary. Geometric constraints in a form of different shape priors [4,5,6,7] are widely used in binary segmentation, but their applicability or usefulness for multi-surface segmentation is limited.

Li *et al.* [2] proposed a nested surface approach that models geometric interaction by defining distance constraints on subsequent surfaces S_ℓ . In their approach, every surface is represented in polar coordinates with respect to a center point (cf. Fig. 1, left) or a center line in 3D. Therefore, they can only represent star shaped segmentations (cf. Fig. 1, middle). While this limitation does not have a strong affect on segmentations for near-circular objects like the left ventricle of a heart, it becomes problematic for complex objects like the brain. Another problem of using the polar coordinates is the need to re-sample and transform the image data usually given in Cartesian coordinates. As a result, energy and optimization in [2] is sensitive to the pre-selected center point/line.

The topological restrictions in [2] were addressed by Delong and Boykov [1]. They can incorporate the minimal distance constraint between the surfaces without relying on a center point/line. They keep the original Cartesian (grid-based)




Fig. 1. *Left:* Li *et al.* [2] can segment star shaped shapes. *Middle:* Shapes that are not star-shaped cannot be represented. *Right:* Delong and Boykov [1] are not limited in their segmentation, but they cannot enforce the Hausdorff distance (cf. Figure 2).

representation of the image (cf. Fig. 1, right) and enforce a tube of a certain distance around the enclosed surface (cf. Fig. 2, middle). We refer to this constraint as *tubular distance constraint*. A constraint on the maximal distance independent of the center point/line leads to the *Hausdorff distance constraint* (cf. Fig. 2, right). [1] were unable to incorporate this Hausdorff distance prior into their global optimization framework.

The *Hausdorff distance prior* is fundamentally different from the *tubular distance prior*. It constitutes a potentially NP-hard optimization problem because the involved energy is a combination of submodular and supermodular terms. We will show that this problem can be approximated efficiently using the submodular-supermodular procedure [8]. This technique was inspired by the concave-convex procedure for continuous functions [9]. Our algorithm changes the underlying graph in a way different from earlier submodular-supermodular techniques [10,11]. For example, besides changing the edge capacities, we also replace certain edges modifying the graph’s connectivity. This operation maintains the value of the previous solution, but the corresponding cut is no longer the global minimum and it can be recomputed. Our connectivity modifications are based on the signed distance maps of the current segmentation, which can be efficiently computed in linear time [12,13]. Typically, our method converges within 5-10 iterations.

This paper is organized as follows. We will present the overall energy function in Section 2 and in Section 3 we will show how to enforce the Hausdorff distance prior. In Section 4, we will present some medical examples. In Section 5 we will provide a conclusion of this work and discuss future work.




Fig. 2. *Left:* Without any distance constraint, an arbitrary multilabel segmentation is possible which implicitly defines separating surfaces S_ℓ . *Middle:* Incorporating a tubular distance t between subsequent labels is polynomial. *Right:* Incorporating an upper bound H on the Hausdorff distance makes the problem potentially NP-hard.



Fig. 3. Label Representation: Labels between 0 and k (here: $k = 3$) can be represented by k binary variables [15]. All these binary variables fulfill the logical conditions $x_\ell \Rightarrow x_{\ell-1}$ (left table). This relationship is illustrated in the four images for $u(x)$ (left image) and x_0, x_1 and x_2 (from middle to right).

2 Multi-surface Segmentation

In the following, we want to address the problem of multi-surface segmentation in the context of medical imaging. We therefore assume that a volumetric image data¹ $I : \Omega \rightarrow \mathbb{R}$ is given that assigns to every point x of the image domain $\Omega \subset \mathbb{R}^3$ a color value $I(x) \in \mathbb{R}$. The goal is now to find a labeling $u : \Omega \rightarrow \mathcal{L}$ that assigns to every point a label of the finite label space $\mathcal{L} = \{0, \dots, k\}$ in a way that the following energy function is minimized

$$E(u) = \underbrace{\sum_{\ell=0}^k \int_{u^{-1}(\ell)} D_\ell(x) \, dx}_{\text{data term}} + \underbrace{\sum_{\ell=1}^k \int_{S_\ell} V_\ell(s) \, ds}_{\text{smoothness term}}. \tag{1}$$

Here, the data terms D_ℓ describe the likelihood that a given point x should be labeled ℓ and the smoothness terms V_ℓ encode the likelihood that a point s should separate the label $\ell - 1$ from label ℓ . S_ℓ describes the separating boundary between the region of label ℓ and $\ell - 1$. It is popular to choose D_ℓ to be the negative log-likelihood of a precomputed GMM for the label ℓ and to use a GAC-like function for V_ℓ [14]. In particular, V_ℓ may depend on ℓ if we know beforehand that certain boundaries tend to be stronger than other boundaries. We will discuss this in detail for the experiments presented in Section 4.

In E , we assume that the area of label $\ell - 1$ is completely surrounded by the area of label ℓ . Otherwise, the second sum of (1) would contain $O(k^2)$ summands instead of k . This nested labeling setup is essential for the energy minimization approaches that we discuss in this paper. It was shown [15] that the global minimum of the spatially discretized energy E can be found by solving a single graph cut problem. While [15] did not enforce the nestedness of the labeling, this was circumvented in [1] by introducing a minimal distance constraint (cf. Section 2.1). In [15], each pixel x is represented by k binary variables x_0, \dots, x_{k-1} where $u(x) = \ell$ is encoded by the k variables $x_i = [i < \ell]$. Figure 3 illustrates

¹ While we assume gray-scaled volumetric images, the presented framework can easily be extended to more general appearance models and image domains.




Fig. 4. *Left:* By using a tubular distance prior of $t = 2$, we enforce the constraint that every pixel within distance t from any green one (labeled 2) must have label $\ell > 0$ (blue or green). *Right:* We enforce this by ∞ -edges from nodes x_1 (green layer) to x_0 (blue layer). The dashed edges are ∞ -edges enforcing $x_1 \Rightarrow x_0$ (cf. (2)).

this for $k = 3$. The data term for any pixel x can then be represented as the submodular function

$$D(x) = D_0(x)\bar{x}_0 + \sum_{\ell=1}^{k-1} D_\ell(x) \cdot \bar{x}_\ell x_{\ell-1} + D_k(x)x_{k-1} + \sum_{\ell=1}^{k-1} \infty \cdot \bar{x}_{\ell-1} x_\ell \quad (2)$$

While the ∞ -terms are not relevant in order to evaluate $D(x)$, a minimization of these terms for general binary x_i leads to values for x_i that guarantee a labeling $u(x)$ according to the table in Figure 3. This is because a minimum will enforce the logical condition $x_\ell \Rightarrow x_{\ell-1}$. As a result, the vector (x_0, \dots, x_{k-1}) starts with entries $x_\ell = 1$ followed by entries $x_\ell = 0$. Minimizing D leads therefore to a labeling $u(x) = \sum_{i < k} x_i$.

Since also the smoothness term of (1) can be represented as a submodular function on the binary variables, we obtain an overall submodular function that we like to minimize. As a result, the submodular function associated with E can be minimized using a graph cut approach [16,17]. We denote the involved graph as $\mathcal{G} = (\mathcal{X} \times \mathcal{L}^*, \mathcal{E}, c)$ where \mathcal{X} denotes the pixel set, \mathcal{L}^* the label set without the background label '0' and $c : \mathcal{E} \rightarrow \mathbb{R}_+$ the capacity function.

2.1 Tubular Distance Prior

For medical applications, we often know the minimal distance of two subsequent surfaces. In other words, we know a lower bound on the *tubular distance*²

$$\text{dist}_T(S_{\ell+1}, S_\ell) = \min_{s_1 \in S_{\ell+1}} \min_{s_2 \in S_\ell} \|s_1 - s_2\|.$$

Let us assume now that in our given application the minimal distances $t_1, \dots, t_{k-1} > 0$ are given. Then, solving the problem:

$$\min_{u: \mathcal{X} \rightarrow \mathcal{L}} E(u) \quad , \text{ s.t. } \text{dist}_T(S_{\ell+1}, S_\ell) \geq t_\ell \text{ for } \ell < k \quad (3)$$

² Note that dist_T is not a metric, but a pseudosemimetric.

can be addressed by adding ∞ -terms to the overall energy. The goal is to end up with an energy of ∞ as soon as the labeling is infeasible. In order to do so, let us denote as

$$\mathcal{N}_{t_\ell}(x) = \{y \in \mathcal{X} \mid \|x - y\| < t_\ell\}$$

the neighborhood of x that consists of all points y that are within a ball of radius $t_\ell > 0$. The logical constraint that we want to enforce for the global solution of (3) is that $u(x) = \ell$ implies $u(y) \geq \ell - 1$ for all $y \in \mathcal{N}_{t_\ell}(x)$. In [1] it was proposed to model this by considering the following energy (cf. Figure 4)

$$E_t(u) = E(u) + \sum_{\ell \in \mathcal{L}} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{N}_{t_\ell}(x)} \infty \cdot \bar{y}_{\ell-1} x_\ell. \tag{4}$$

In the following, we refer to the graph associated with the energy E_t as $\mathcal{G}_t = (\mathcal{X} \times \mathcal{L}^*, \mathcal{E}_t, c_t)$. This graph is a super-graph of the graph \mathcal{G} of Section 2. Besides the edges of \mathcal{G} , \mathcal{G}_t consists of ∞ -edges described in (4).

In the next section, we will show how a Hausdorff distance prior can also be added to this framework by adding more logical constraints. As long as they are of the form $\infty \cdot \bar{x}y$ they can always be modeled by introducing ∞ -edges into the existing graph. As it will turn out, finding the correct edges is potentially NP-hard and we will use a submodular-supermodular procedure in order to solve this problem.

3 Hausdorff Distance Prior

In the last section, we revisited how the global optimum of a classical multi-labeling problem can be solved if the surfaces S_ℓ that describe the boundaries of each label are enclosed by one another. In addition, we mentioned how to incorporate a tubular distance prior. It turns out that these tools are very useful for medical images. Nonetheless, the outer boundaries of an object might be rather weak. As a consequence, it is possible that the boundary of the outer labels is leaking and will only be constrained by the smoothness term of (1) (cf. Figure 7). This behavior cannot be explained anatomically, because every outer muscle or outer membrane can only expand until a certain distance from the enclosing organ is reached. This means that the Hausdorff distance between subsequent surfaces

$$\text{dist}_{\text{HD}}(S_\ell, S_{\ell-1}) = \max_{s_1 \in S_{\ell-1}} \min_{s_2 \in S_\ell} \|s_1 - s_2\| \tag{5}$$

is bounded from above. It turns out that incorporating the Hausdorff distance constraint into the multi-surface segmentation leads to an energy that is neither completely submodular nor completely supermodular. Interestingly, the supermodular part of this penalty vanishes in a certain restricted 1D-situation. This is essential to the approximation approach that we propose. In Section 3.1, we will show that the Hausdorff distance prior involves a combination of a sub- and




Fig. 5. *Left:* By using a Hausdorff distance prior of $H = 2$, we enforce the constraint that at least one pixel within a distance of H from any blue pixel (labeled 1) must have label $\ell > 1$ (green). *Right:* Therefore, we must choose the **correct** ∞ -edge from the nodes x_0 (blue layer) to x_1 (green layer). This makes the problem potentially NP-hard.

a supermodular energy part and in Section 3.2 we will show that the supermodular energy part vanishes for a specific 1D problem. In Section 3.3, we will use this observations in order to approximate the overall energy. To this end, we have to redefine the involved graph in every iteration. While the graph structure changes, the involved energy is always an upper envelope of the supermodular function and the energy of the global minima is non-increasing. Therefore, our approach is a variant of the submodular-supermodular procedure [8] that will always converge towards a local minimum. In Section 4 we will present these minima for some real-world examples.

3.1 Sub- and Supermodularity of the General Prior

In the following, we assume that for a pair of subsequent surfaces $(S_\ell, S_{\ell+1})$, both a lower bound $t_\ell > 0$ for the tubular distance and an upper bound $H_\ell > 0$ for the Hausdorff distance is given. Then, we are interested in solving the following energy problem:

$$\min_{u: \mathcal{X} \rightarrow \mathcal{L}} E(u) \quad , \text{ s.t. } \text{dist}_T(S_{\ell+1}, S_\ell) \geq t_\ell \text{ for } \ell < k$$

$$\text{dist}_{\text{HD}}(S_{\ell+1}, S_\ell) \leq H_\ell \text{ for } \ell < k \tag{6}$$

While transforming the tubular distance constraint into a graph cut problem can be easily done, it is much more difficult for the Hausdorff distance prior H_ℓ . The logical constraint that we want to enforce is:

$$u(x) = \ell \quad \Rightarrow \quad \exists y \in \mathcal{N}_{H_\ell}(x) : u(y) = \ell + 1 \tag{7}$$

Similar to the tubular distance prior, we can translate this into the following penalty term (cf. Figure 5)

$$p_\ell(x) = \infty \cdot x_{\ell-1} \cdot \prod_{y \in \mathcal{N}_{H_\ell}(x)} \bar{y}_\ell. \tag{8}$$

Theorem 1. *The implication of (7) is true iff $p_\ell(x) = 0$.*

Table 1. Hausdorff Distance Penalty: Because $p(1, \cdot, \cdot)$ is supermodular and $p(\cdot, 0, \cdot)$ is submodular, p cannot be minimized via graph cuts [16,17].

$x_{\ell-1}$	y_ℓ^0	y_ℓ^1	$p(x_{\ell-1}, y_\ell^0, y_\ell^1)$	$p(1, y_\ell^0, y_\ell^1)$	$p(x_{\ell-1}, 0, y_\ell^1)$
0	0	0	0	-	0
0	0	1	0	-	0
0	1	0	0	-	-
0	1	1	0	-	-
1	0	0	∞	∞	∞
1	0	1	0	0	0
1	1	0	0	0	-
1	1	1	0	0	-

Proof. Following the definitions of $x_{\ell-1}$ and y_ℓ , the expression $p_\ell(x) = 0$ is equivalent to

$$u(x) \geq \ell \quad \Rightarrow \quad \exists y \in \mathcal{N}_{H_\ell}(x) : u(y) \geq \ell + 1. \tag{9}$$

Since the two expressions (7) and (9) are both obviously true for $u(x) < \ell$, we have to consider the following two remaining cases

Case 1: $u(x) = \ell$

In that case, (7) implies (9). If on the other hand (7) is false, all $y \in \mathcal{N}_{H_\ell}(x)$ are not labeled $\ell + 1$. Because of the nested labeling constraint, these labels have to be all less than $\ell + 1$ which proves that also (9) is false.

Case 2: $u(x) > \ell$

In that case, (7) is always true. Therefore, we have to proof that the RHS of (9) is also always true. Since $x \in \mathcal{N}_{H_\ell}(x)$ and $u(x) \geq \ell + 1$, the RHS of (9) is in fact also always true.

□

The penalty term (8) can be used in order to enforce the Hausdorff distance constraints by considering the following overall energy:

$$E_{t,H}(u) = E_t(u) + \sum_{\ell \in \mathcal{L}} \sum_{x \in \mathcal{X}} p_\ell(x) \tag{10}$$

From Table 1 it is clear that $p_\ell(x)$ is not submodular as soon as the neighborhood $\mathcal{N}_{H_\ell}(x)$ consists of more then one pixel. This means that the energy cannot be minimized via graph cut [17, Theorem 6.1]. Therefore, we want to apply the submodular-supermodular procedure. The central point of this procedure is to find a submodular upper envelope $E^{(n)}$ to the energy $E_{t,H}$ that touches $E_{t,H}$ at the global minimizer of $E^{(n-1)}$. In Section 3.3 we will create energies that have exactly this property. The creation of these energies depends on an important property for certain $1D$ problems that we will discuss in Section 3.2.




Fig. 6. *Left:* Examples for a segmentation with Hausdorff distance prior $H_1 = 2$ if the inner segment is connected and contains y . *Right:* This can be enforced by ∞ -edges that follow the negative gradient of the signed distance map with respect to y .

3.2 Pure Submodularity of a 1D Prior

In Section 3.1 we showed that incorporating the general Hausdorff distance prior into multi-surface segmentation is potentially NP-hard. In particular, we showed that if there are at least two neighboring pixels involved for which the Hausdorff distance constraint might be violated, the problem becomes also potentially NP-hard. This means that even for a 1D image segmentation, the Hausdorff distance constraint cannot be solved globally optimal. The whole situation changes, if the constraint has only to be checked with respect to one other 1D-pixel. In that case it results in adding one additional edge per pixel into the graph \mathcal{G}_t . One of these situations is depicted in Figure 6. It assumes that we have a ternary labeling problem, *i.e.*, $\mathcal{L} = \{0, 1, 2\}$ and that the following holds:

- The segment with the label $\ell = 2$ is connected, *i.e.*, it is an interval.
- A pre-defined pixel y lies inside of this interval.

If we want to enforce a maximal Hausdorff distance of H_1 , we know that for every pixel $z \geq y + H_1$ or $z \leq y - H_1$, we only have to enforce the Hausdorff distance constraint with respect to $z - H_1$ resp. $z + H_1$. These constraints can be enforced by adding ∞ -edges into the graph associated with E_t . These edges point into the negative gradient's direction of the distance map with respect to the point y . The pixels in the interval $]y - H_1, y + H_1[$ cannot violate any constraint because the distance between these points and y is smaller than H_1 (Figure 6, left). Nonetheless, we can add edges from these nodes to y in order to ensure that y is chosen as a point of the interval (dashed edges in Figure 6).

While the application of this specific 1D image segmentation is rather limited, there is an obvious extension to arbitrary dimensions, namely the problem of multi-labeled image segmentation with a circular core of label $\ell = k$. A variation of this idea was used in [2], where the authors used star-shaped cores. While it is also easy to find the core of star-shaped areas by just following the line from any pixel towards the star center, this path is not the shortest path to the star's boundary. That is the reason why [2] does not optimize the Hausdorff distance. Instead it considers a maximal distance constraint for a polar representation of shapes. Besides the theoretical drawback of not minimizing the Hausdorff distance, there is also a practical drawback. Namely that only star-shaped surfaces can be detected. For more complex objects like the bifurcated

carotid artery or the airways of the lung, this strong shape constraint does not hold any longer and different approaches have to be used. In [2] this was addressed by detecting the core of the true segmentation in a pre-processing step that does not use the same global energy. While this approach works for applications where a coarse segmentation can be easily detected, this method becomes more difficult for applications like brain segmentation where detecting this core segment is as difficult as finding the right segmentation.

In this work, we want to advocate that such a coarse segmentation is in fact not necessary. To see this, consider again the above mentioned 1D-problem. Instead of the chosen ∞ -edges, we can also choose edges that follow the negative gradient's direction of the signed distance map with respect to the true segment $u^{-1}(2)$. If we do not know this segment, we can iterate between segmentation and adjusting the ∞ -edges associated to the Hausdorff distance. We will use this idea in Section 3.3 in order to develop an iterative method for the general segmentation problem. In order to do this, we re-define the signed distance transform of every label after each segmentation and iterate this process. We will show how this approach provides us with a submodular-supermodular procedure that will converge towards a local minimum of $E_{t,H}$.

3.3 Submodular-Supermodular Procedure

In Section 3.2 we explained how the challenging Hausdorff distance constraint can be incorporated into the multi-surface segmentation if the interior of every surface is circular and the center point is known. For this extremely limited case, we explained how the signed distance map with respect to the interior of $S_{\ell-1}$ can be used to define constraints for the label ℓ . These observations are the core elements of our submodular-supermodular procedure described in Algorithm 1. In every iteration the signed distance map is re-computed with respect to previously computed segments. In order to show that this algorithm is a *submodular-supermodular procedure*, we have to prove the following theorem:

Theorem 2. *Algorithm 1 is a submodular-supermodular procedure, i.e., it fulfills the following properties of the energies $E^{(n)}$ and their global minimum $u^{(n)}$:*

1. $E^{(n)}$ is submodular.
2. $E^{(n)}(u) \geq E_{t,H}(u)$ for all labelings $u : \mathcal{X} \rightarrow \mathcal{L}$.
3. $E^{(n+1)}(u^{(n)}) = E_{t,H}(u^{(n)})$.

Proof. The three properties follow from the described construction as follows:

1. Since the graph $\mathcal{G}^{(n)}$ can be used to minimize $E^{(n)}$, $E^{(n)}$ has to be submodular [16,17].
2. If $E^{(n)}$ or $E_{t,H}$ are finite, their energy is exactly E . Therefore, it is enough to study the cases when either of these two energies are infinite because some of the incorporated constraints are violated. If $E^{(n)}(u)$ is infinite, the inequality $E^{(n)}(u) \geq E_{t,H}(u)$ is naturally true. If on the other hand $E_{t,H}(u) = \infty$, there exists a point $x \in \mathcal{X}$ that violates the Hausdorff distance constraint H_ℓ

Algorithm 1. Submodular-Supermodular Procedure

Input: Multi-surface graph $\mathcal{G} = (\mathcal{X} \times \mathcal{L}, \mathcal{E}, c)$ associated to E ,
 tubular distance constraints t_ℓ and Hausdorff distance constraints H_ℓ .
Output: Labeling $u : \mathcal{X} \rightarrow \mathcal{L}$ that locally minimizes $E_{t,H}$.

- 1: Set $n = 0$.
- 2: Create $\mathcal{G}^{(n)} = \mathcal{G}_t$ according to Section 2.1.
- 3: Find global minimum $u^{(n)}$ of $E^{(n)}$ associated to $\mathcal{G}^{(n)}$.
- 4: **repeat**
- 5: **for** $\ell = 2, \dots, k$ **do**
- 6: Compute signed distance transforms $dt_\ell^{(n+1)}(x)$
 for the sets $\mathcal{X}_\ell = \{y \in \mathcal{X} | u^{(n)}(y) \geq \ell\}$.
- 7: **end for**
- 8: $n := n + 1$
- 9: Create graph $\mathcal{G}^{(n)} = \mathcal{G}_t$ according to Section 2.1
- 10: **for** $\ell = 1, \dots, k - 1$ **do**
- 11: **for** $x \in \mathcal{X}$ **do**
- 12: Let $v := H_\ell \cdot \nabla dt_{\ell+1}^{(n)}(x)$.
- 13: Find maximal $t \in [0, 1]$ such that for $y := x - t \cdot v$ the following holds
 $dt_\ell^{(n)}(y) = dt_\ell^{(n)}(x) - H_\ell \cdot t$.
- 14: Add ∞ -edge to $\mathcal{G}^{(n)}$ from x_ℓ to $y_{\ell+1}$
- 15: **end for**
- 16: **end for**
- 17: Find global minimum $u^{(n)}$ of $E^{(n)}$ associated to $\mathcal{G}^{(n)}$.
- 18: **until** $u^{(n)} = u^{(n-1)}$
- 19: **return** $u^{(n)}$

with $u(x) = \ell - 1$. This means that none of the points of $\{y \in \mathcal{X} | \|x - y\| \leq H_\ell\}$ are labeled ℓ . Since these points include also the point y computed in Step 13, one ∞ -edge of $\mathcal{G}^{(n)}$ has to be cut and thus, $E^{(n)}(u) = \infty$.

3. Since $E^{(n+1)}$ selects for every Hausdorff distance the active constraint with respect to $u^{(n)}$, $E^{(n+1)}(u^{(n)})$ is finite and the equality holds.

□

Note that Step 13 in Algorithm 1 follows the negative gradient of the signed distance map until either the Hausdorff distance H_ℓ or the medial axis is reached. This is an analogous case to the dashed edges in Fig. 6. While these edges were optional for the 1D case, they are important for Algorithm 1 to become a submodular-supermodular procedure. If we would not use these edges, the proof of Property 2 in Theorem 2 would otherwise not be possible.

All solutions $u^{(n)}$ computed by Algorithm 1 fulfill the Hausdorff distance constraints, but it is possible that the constraints enforced in Step 14 are too tight. Therefore, a re-iteration is necessary. While it is easy to prove that the algorithm will terminate, it would rely on the fact that the set of all possible label settings is finite. This is a very bad upper bound for the running time. In practice, we never experienced more than 10 iterations which makes this algorithm very attractive for practical applications.




Fig. 7. To segment the left ventricle of a heart (left), we use conservative seeds (2nd f.l.). The two right images show the segmentation w.r.t. a pure tubular constraint (2nd f.r.) and a combination of tubular and Hausdorff distance constraint (right).

4 Experiments

Geometrically constrained multi-surface segmentation is particularly useful in medical applications where segments must be consistent with the anatomy. We use grayscale medical images that should be partitioned into k nested segments or layers. Each segment is represented by one label. To guide the segmentation,




Fig. 8. To segment the brain (1st row, left), we use conservative seeds (2nd row, left). The 2nd column shows the result with mere tubular distance constraints and the 3rd column also incorporates the Hausdorff distance constraint. The 1st and 2nd row shows the segmentation results with 3 and 4 labels resp.




Fig. 9. To segment the carotid artery (left column), we use a 3-label image segmentation problem with tubular distance and Hausdorff distance constraint. The segmentation (middle column) defines a 3D image of the two segmented surfaces (right column).

we also employ some small set of “seeds” from a user. The seeds provide hard constraints and train GMM appearance model for every segment/label. We use the standard EM-style re-estimation of GMM after each convergent segmentation of our algorithm 1. In our examples the image edges are much stronger for the

inner segment than for the outer one. Therefore, the edge sensitivity for the core segment V_k is scaled up by a factor of 5 in all the experiments.

In our first experiment we segmented the left ventricle of a heart in an MR image. Very conservative seeds (cf. Figure 7) were mainly used to provide reasonable color models for each label. While the segmentation w.r.t. to the tubular constraint $t_1 = 2$ is locally satisfactory for the left ventricle, the label for the surrounding muscle strongly leaks into the background. This is mainly because the lower part of this muscle does not have strong image edges. Adding the Hausdorff distance constraint $H_1 = 10$ gives a much more accurate delineation of the ventricle. The wrong labeling in the upper right corner of the image also disappears. Note that we receive the same result for a range of t_1 from 1 to 3 and for a range for H_1 from 8 to 12.

In the second experiment, we apply the same technique to an MR image of the brain (cf. Figure 8). Both tubular constraints were set to $t_1 = t_2 = 3$ and the Hausdorff distance constraints were set to the tight value of $H_1 = 4$ for the outer segment and unconstrained $H_2 = \infty$ for the inner segment. Note that the Hausdorff distance constraint gave much better results with even the 3-label segmentation (top row). In particular, the fine structure of the gray matter (blue label) was detected. Also, the skull is automatically merged into the background label. These results are not possible without the Hausdorff distance constraint. As expected, the 4-label case (bottom row) gives semantically more reasonable results than the 3-label case.

In a last experiment, we segmented challenging 3D ultrasound data of the carotid artery. We used a tubular distance constraint of $t_1 = 2$ and a Hausdorff distance constraint of $H_1 = 10$. Even though the data is very noisy, we obtain a good segmentation that accurately follows the bifurcation (cf. Figure 9).

5 Conclusion and Future Work

We showed how the Hausdorff distance constraint can be incorporated into multi-surface segmentation without topological limitations. We provided a submodular-supermodular procedure that can find an approximative solution. We also showed that these solutions provide us with reasonable segmentation for challenging 2D and 3D data sets. One major advantage with respect to prior work is that we do not rely on a center point/line and, thus, are not restricted to star shaped segmentations. There are many possible extensions of our work that we would like to address in the future.

First of all, it makes sense to use more flexible model estimators than GMM. Secondly, the memory consumption of the tubular distance constraint is very extensive. Similarly to the Hausdorff distance constraint, one could try to add only those edges that are really necessary. Thirdly, we like to explore other distances reflecting the elasticity of the involved tissues and muscles.

Acknowledgements. This work was partially supported by Natural Sciences and Engineering Research Council of Canada (NSERC, Discovery program),

Canada Fund for Innovations (CFI, project 10318), the National Institute of Health (NIH, 5R01EB004640-07) and the Agence Nationale de la Recherche (ANR-2010-BLAN-0205-01). We like to thank Robarts Research Institute for providing medical images.

References

1. Delong, A., Boykov, Y.: Globally optimal segmentation of multi-region objects. In: IEEE ICCV, Kyoto, Japan, pp. 285–292 (2009)
2. Li, K., Wu, X., Chen, D.Z., Sonka, M.: Optimal surface segmentation in volumetric images—a graph-theoretic approach. IEEE PAMI 28, 119–134 (2006)
3. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE PAMI 23, 1222–1239 (2001)
4. Freedman, D., Zhang, T.: Interactive graph cut based segmentation with shape priors. In: IEEE Int. Conf. on CVPR, pp. 755–762 (2005)
5. Cremers, D.: Dynamical statistical shape priors for level set based tracking. IEEE PAMI 28, 1262–1273 (2006)
6. Veksler, O.: Star Shape Prior for Graph-Cut Image Segmentation. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part III. LNCS, vol. 5304, pp. 454–467. Springer, Heidelberg (2008)
7. Lempitsky, V., Blake, A., Rother, C.: Image Segmentation by Branch-and-Mincut. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part IV. LNCS, vol. 5305, pp. 15–29. Springer, Heidelberg (2008)
8. Narasimhan, M., Bilmes, J.A.: A submodular-supermodular procedure with applications to discriminative structure learning. In: UAI, pp. 404–412 (2005)
9. Yuille, A.L., Rangarajan, A.: The concave-convex procedure. Neural Computation 15, 915–936 (2003)
10. Rother, C., Minka, T., Blake, A., Kolmogorov, V.: Cosegmentation of image pairs by histogram matching - incorporating a global constraint into mrfs. In: IEEE Int. Conf. on CVPR, pp. 993–1000 (2006)
11. Ladicky, L., Russell, C., Kohli, P., Torr, P.H.S.: Graph Cut Based Inference with Co-occurrence Statistics. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part V. LNCS, vol. 6315, pp. 239–253. Springer, Heidelberg (2010)
12. Karzanov, A.V.: Quick algorithm for determining the distance from the points of the given subset of an integer lattice to the points of its complement. Cybernetics and System Analysis, 177–181 (1992); translation by Julia Komissarchik
13. Felzenszwalb, P., Huttenlocher, D.: Distance transforms of sampled functions. Computing and Information Science TR2004-1963, Cornell (2004)
14. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. Int. J. of Computer Vision 22, 61–79 (1997)
15. Ishikawa, H.: Exact optimization for markov random fields with convex priors. IEEE PAMI 25, 1333–1336 (2003)
16. Boros, E., Hammer, P.L.: Pseudo-boolean optimization. Discrete Applied Mathematics 123, 155–225 (2002)
17. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? IEEE PAMI 24, 657–673 (2004)