Computerized Medical Imaging and Graphics 69 (2018) 96–111

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

journal homepage: www.elsevier.com/locate/compmedimag

Optimal multi-object segmentation with novel gradient vector flow based shape priors

Junjie Bai[a], Abhay Shah[a], Xiaodong Wu[a,b,*]

[a] Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA 52242, USA
[b] Department of Radiation Oncology, The University of Iowa, Iowa City, IA 52242, USA

ARTICLE INFO

Article history:
Received 3 January 2018
Received in revised form 18 June 2018
Accepted 21 August 2018

Keywords:
Shape priors
Gradient vector flows
Multi-object segmentation
Segmentation

ABSTRACT

Shape priors have been widely utilized in medical image segmentation to improve segmentation accuracy and robustness. A major way to encode such a prior shape model is to use a mesh representation, which is prone to causing self-intersection or mesh folding. Those problems require complex and expensive algorithms to mitigate. In this paper, we propose a novel shape prior directly embedded in the voxel grid space, based on gradient vector flows of a pre-segmentation. The flexible and powerful prior shape representation is ready to be extended to simultaneously segmenting multiple interacting objects with minimum separation distance constraint. The segmentation problem of multiple interacting objects with shape priors is formulated as a Markov Random Field problem, which seeks to optimize the label assignment (objects or background) for each voxel while keeping the label consistency between the neighboring voxels. The optimization problem can be efficiently solved with a single minimum *s*-*t* cut in an appropriately constructed graph.

The proposed algorithm has been validated on two multi-object segmentation applications: the brain tissue segmentation in MRI images and the bladder/prostate segmentation in CT images. Both sets of experiments showed superior or competitive performance of the proposed method to the compared state-of-the-art methods.

1. Introduction

Shape priors have been widely used in medical image segmentation due to the similar intensity profiles and weak boundaries between the target objects and background. For example, the popular deformable model (Jalba et al., 2004; Weese et al., 2001; Martin et al., 2010; McInerney and Terzopoulos, 1996) and the LOGISMOS (layered optimal graph image segmentation of multiple objects and surfaces) (Wu and Chen, 2002; Oguz and Sonka, 2014; Yin et al., 2010; Song et al., 2010, 2009) methods use a mesh in the physical space to ensure that the resulting segmentation is roughly aligned to the initial model or the pre-segmentation. Recently, the star-shaped prior has been successfully encoded in the voxel grid space, which can be incorporated into the graph-based image segmentation methods with efficient global optimizers (Gulshan et al., 2010; Veksler, 2008; Bai et al., 2014).

1.1. Shape priors with a mesh-based representation

Deformable model first initializes a mesh representing the target object close to the desired location and orientation (Dornheim et al., 2007; Jalba et al., 2004; Weese et al., 2001). The boundary and regional information are then used to evolve the mesh vertices along their normal directions to adhere to the image content. An elastic force is also often used to limit the amount of warp to enforce the shape prior (Dornheim et al., 2007; Weese et al., 2001). Such evolution is repeated multiple times until convergence.

LOGISMOS has been widely used for multiple surface segmentation with the capability of enforcing mutual surface interactions, while still achieving globally optimal solutions (Wu and Chen, 2002; Li et al., 2006; Yin et al., 2010; Song et al., 2009, 2010). The method first builds a mesh based on an initial pre-segmentation or a shape model of the target object. The image is then resampled usually along the normal direction at each mesh vertex to form a graph node column. The graph nodes on each column represent the set of all possible target surface locations sampled at the normal direction at each mesh vertex. Various edges are then added among the nodes to ensure the smoothness of the surfaces, and the mutually interactive relationships between them. The globally optimal

* Corresponding author at: Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA 52242, USA.
E-mail address: xiaodong-wu@uiowa.edu (X. Wu).

https://doi.org/10.1016/j.compmedimag.2018.08.004
0895-6111/

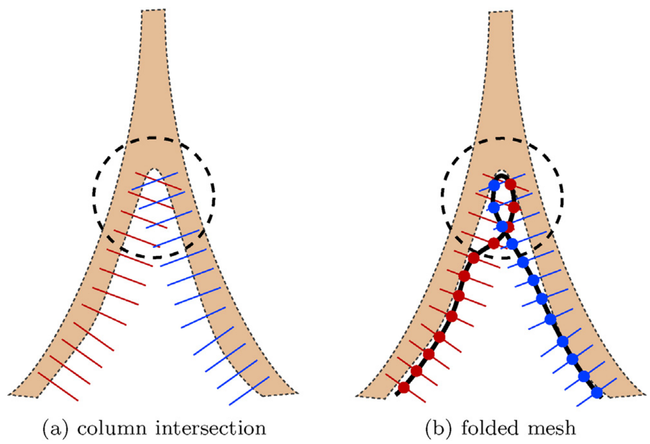

Fig. 1. Mesh folding at the bifurcation locations. (a) shows part of LOGISMOS graph columns built based on a pre-segmentation (light brown region). The red columns and blue columns are intersecting in the circled area. This may lead to folded mesh as shown in (b) in the solution. Although deformed model do not explicitly build such columns, the mesh folding may still happen during mesh evolving. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

surface locations are then computed by solving a single minimum *s-t* cut in the constructed graph.

Although widely used in medical image segmentation, the mesh-based shape prior representation has several shortcomings in practice. Firstly, the mesh may be "self-intersected" or "folded" in the output solution, causing undesired topological changes from the prior shape. The problem is especially noticeable at the high curvature locations such as concavities and bifurcations (see Fig. 1). To avoid this problem, the deformable method need to run expensive self-intersection detection and re-meshing algorithms in each evolution step (Park et al., 2001; Pons and Boissonnat, 2007; Wang et al., 2008). For LOGISMOS, various methods have been proposed to prevent the interference between graph node columns. The electrical field (Yin et al., 2009), flow lines (Petersen et al., 2014) and gradient vector flow (Oguz and Sonka, 2014) are used to build curved graph columns instead of the traditional straight-line columns along the normal directions. Veni et al. (2013, 2013) devised a flexible 3D lattice particle system to compute non-intersecting graph column trajectories. The graph columns for the same layer are generated with repulsive spring-like forces, while the columns for the interactions between different layers are computed with attractive spring-like forces. The whole spring potential system, if well balanced, could avoid mesh folding. However, solving the spring system is nontrivial. A local gradient descent method with asynchronous updates was used by the authors. The omnidirectional displacement method (Kainmueller et al., 2013) has also been proposed to distribute graph nodes uniformly in a sphere centered at each mesh vertex. However, the resulting MRF problem can no longer be solvable by the minimum *s-t* cut algorithm, and the high computational complexity restricts its wide application.

Secondly, it is difficult to enforce the interactions between the shape priors for multiple objects. For deformable model method, an expensive mesh-collision detection algorithm has to be run at the end of each evolving iteration to prevent two close objects from intersecting each other (Klinder et al., 2009; Teschner et al., 2005). For LOGISMOS-based methods, nesting surfaces must share the same base mesh in order to enforce the surface interaction constraints (Oguz and Sonka, 2014). For example, the same base mesh is used for segmenting both inner and outer walls of left atrium with soft smoothing penalties (Veni et al., 2013a,b). When multiple meshes for different shape priors are used, a workaround must be proposed to register multiple meshes into the same physical space. Song et al. (2010, 2009) proposed a heuristic method to merge the graph columns from the mesh for the prostate shape prior and the mesh for the bladder shape prior into a common "interactive region" in order to enforce the exclusion relationship of the two objects. To make use of multiple *different* shape priors for interactive objects efficiently, we should embed the shape priors directly in the voxel space instead of separate mesh spaces.

1.2. Shape priors defined directly in the voxel grid space

Another group of methods embed the shape prior directly in the voxel grid space, fundamentally eliminating the drawbacks of the mesh representation. Veksler (2008) introduced a flexible star-shaped prior which requires every point in the target object to be visible to a predefined *star center* through a straight line totally within the object. The prior is enforced by adding into the graph cut method (Boykov and Funka-Lea, 2006) a set of edges pointing towards the star center; each edge is set a weight of $+\infty$. Those edge mimic the Euclidean rays originating from the star center. However, the discretized rays forming with those edges computed by heuristics may not behave like their Euclidean counterparts. Bai et al. (2014) utilized the well-behaved *consistent digital rays* (Christ et al., 2012; Chun et al., 2009) to systematically build a set of edges so that the formed discretized rays have a tight Hausdorff distance bound to their Euclidean counterparts. Gulshan et al. (2010) extended Veksler's star-shaped prior by using geodesic paths instead of Euclidean rays to define a star shaped object. The geodesic paths are defined by combining the Euclidean distance and a penalty term for crossing edges along the path. Including gradient penalty term allows the geodesic paths to bend around image gradients to reduce the number of star centers needed to cover a complex shaped object. However, it also makes the star center locations tricky to find as Fig. 13 shows an intuitive placement of star centers leading to failed segmentation (discussed more in Section 4).

If more detailed template or pre-segmentation is available, then the distance transform of the aligned template can be incorporated into the pairwise term of the graph-cut framework to softly penalize the difference between segmentation and the template (Vu and Manjunath, 2008; Aslan et al., 2010). This method, however, requires a balance between the soft shape prior term and the other image information terms, which may be hard to tune.

1.3. Overview of contributions

A gradient vector flow (GVF) (Xu and Prince, 1998) defines a diffusion of the gradient vectors from a grayscale or binary edge map. When computed from an edge map, it defines a vector roughly pointing towards the closest strong edge at each voxel. In this paper, we show that if computed from a volumetric pre-segmentation, instead of an edge map, the GVF encodes descriptive shape information from the pre-segmentation. We thus propose a novel GVF-based shape prior which can be *directly* embedded into the voxel grid space. This shape prior representation avoids the cumbersome self-intersection detection problem in the mesh representation. It enables a simple incorporation of the inclusion/exclusion relationship in segmenting multiple interacting objects with shape priors. The segmentation problem of multiple interacting objects with shape priors is formulated as a Markov Random Field problem, which seeks to optimize the label assignment (objects or background) for each voxel while keeping the label consistency between the neighboring voxels. The optimization problem can be efficiently solved with a single minimum *s-t* cut in an appropriately constructed graph.

The main contributions of this paper are summarized, as follows.

- A novel GVF shape prior that can effectively capture global shape information is proposed.

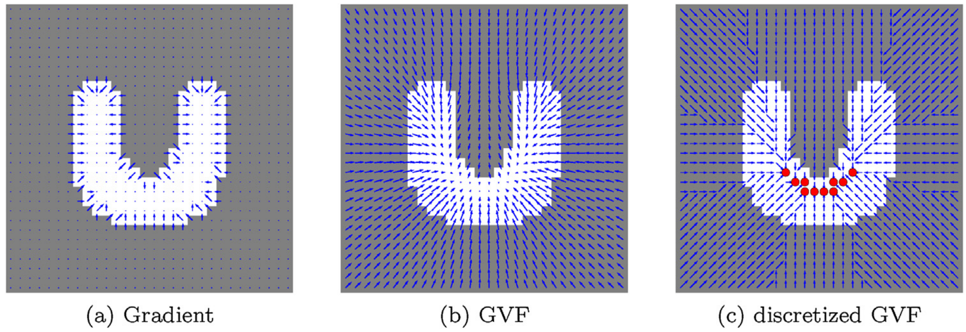

Fig. 2. The use of GVF to encode the shape prior information in the image. (a) shows that the raw gradient of the pre-segmentation is restricted to a narrow band around the boundary. GVF can propagate it smoothly over the whole image, as in (b). A discretized GVF in (c) generates a GVF path for every non-core voxels connecting to one of the core voxels (red circles). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

- The GVF-based shape prior can be directly embedded into the image grid space, facilitating the incorporation of the interaction of multiple objects.
- The GVF-based shape prior representation avoids the cumbersome self-intersection detection problem in the mesh representation.
- The segmentation is formulated as a Markov random field (MRF) optimization problem encoded within it the GVF-based shape prior.
- The resulting optimization problem for optimal multi-object segmentation can be efficiently solved by the minimum *s-t* cut algorithm, yielding a globally optimal solution.

2. Methods

We first introduce the gradient vector flows and the GVF-based shape priors, then we show how to incorporate such a shape prior into an MRF formulation and how to segment multiple interacting objects simultaneously.

2.1. Gradient vector flows

Suppose P is the set of image voxels in a 3D grid. The binary segmentation assigns every voxel $p \in P$ a label $f_p \in L$, where $L = \{0, 1\}$ is the set of available labels. Assume that a pre-segmentation $\hat{S} = \{p|\hat{f}_p = 1, p \in P\}$ is available, then the gradients of the pre-segmentation $\nabla\hat{S}$ are those nonzeroes around the pre-segmentation boundary (Fig. 2(a)). In the narrow band with $\nabla\hat{S} \neq 0$, the gradient vectors encode the shape information of the pre-segmentation in the sense that they are the most direct directions pointing towards the inside of the pre-segmentation.

In order to enforce the shape prior in a larger range, we propose to use the GVF method (Xu and Prince, 1998) to propagate the gradient vectors. The GVF takes the raw gradient vector field $\nabla\hat{S}$ as input, and computes a new vector field $\mathbf{h}$ to minimize the following energy

$$E^{GVF}(\mathbf{h}) = \int\int\int_P \|\nabla\hat{S}\|^2\|\mathbf{h} - \nabla\hat{S}\|^2 + \mu\|\nabla\mathbf{h}\|^2\, dx\, dy\, dz. \tag{1}$$

The first term is the product of the gradient magnitude and the squared difference of the output vector field and the input gradient. The second term is the sum of the squares of the partial derivatives of the output GVF vectors. More specifically, the first term, which is dominant when the input gradient $\nabla\hat{S}$ is strong, encourages the output vector field $\mathbf{h}$ to well align with the input gradient. The second term, which is dominant when the input gradient is small, ensures the output vector field to change smoothly in the regions far away from the region with strong gradients as the input. The constant μ balances the fidelity to the original input gradient field $\nabla\hat{S}$, and the smooth extrapolation of the original gradients to remote regions.

Suppose we set μ to be zero, then the output GVF field $\mathbf{h}$ would be exactly the same as the input field $\nabla\hat{S}$. In another word, $\mathbf{h}$ would only be nonzero inside a narrow band around pre-segmentation boundary, and drops immediately to zeros once outside this narrow band. By setting μ to a nonzero value, the second term would gradually smooth this sharp drop in the output $\mathbf{h}$, effectively propagating the gradient information farther beyond the narrow band. The larger μ is, the longer the gradient propagation range is. If we set μ to a very large value, then it can effectively propagate the gradient throughout the whole image. Fig. 2(b) shows such an example where a GVF vector field propagates the gradient information (Fig. 2(a)) to all voxels of the image. However, setting μ to a larger value leads to longer time for computing the GVF field. Thus, one may only need to set μ to a value that covers the target object boundary search range.

2.2. GVF-based shape priors in the voxel grid space

In Fig. 2(b), we can see that the GVF magnitude is very small at the "core" of pre-segmentation due to the competition of boundary gradients from multiple directions. Formally we define the object *"core"* C as the set of voxels inside the pre-segmentation with their GVF magnitude smaller than a small threshold θ_C, i.e., $C = \{p \in P|p \in \hat{S}, \|\mathbf{h}_p\| < \theta_C\}$. The GVF directions of those core voxels are considered not reliable due to the conflicting gradient information of edges from very different directions. Fig. 2(c) shows an example in which the red circled voxels are the core voxels for the pre-segmentation (white mask).

To label each voxel in the voxel grid space, we discretize the GVF field so that the GVF vector at each voxel points towards exactly one of its neighbor voxels. More concretely, the GVF vector field $\mathbf{h}$ is transformed into the discretized GVF vector field $\mathbf{h}^D$. For every pixel p, $\mathbf{h}^D_p = \vec{pq}$ where $q = \mathrm{argmax}_{q \in N_p} \mathbf{h}_p \cdot \vec{pq}/(\|\mathbf{h}_p\|\|\vec{pq}\|)$. N_p is the set of neighboring voxels of voxel p. In another word, $\mathbf{h}^D_p$ is the vector from p to its neighbor voxel with the smallest angle error when approximating GVF vector $\mathbf{h}_p$. See Fig. 2(c) for an example of the discretized GVF field.

For any non-core voxel $p \notin C$, the *GVF path* of p is defined as $GP(p) = p, q_0, q_1, \ldots, q_N$, such that $\vec{pq_0}, \vec{q_0q_1}, \ldots, \vec{q_{N-1}q_N} \in \mathbf{h}^D$, and $q_N \in C$. All GVF paths form a GVF tree or forest. The *GVF-based shape prior* is defined as: if one voxel p is part of the object, then all voxels along the GVF path $GP(p)$ are also part of the object. In another word, the boundary surface of the target object intersects with any GVF path at most once. In this way, we maintain the global structure of the segmented object to align with the shape of the pre-segmentation, which essentially means that the target object can be

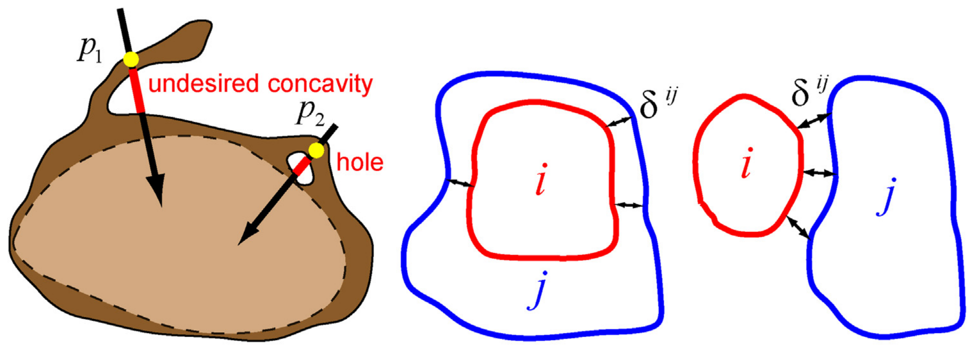

(a) GVF-based shape prior (b) inclusion (c) exclusion

Fig. 3. Illustrating GVF-based shape prior and the inclusion/exclusion multi-object interaction priors. (a) Enforcing the GVF shape prior can prevent a dramatic shape deviation from pre-segmentation. The light brown region indicates the pre-segmentation of the target object. The two thick black arrows represent the GVF paths from the voxels p_1 and p_2, respectively, with respect to the pre-segmentation. With the GVF shape prior, if p_1 is on the final segmentation (the dark and light brown regions), then all voxels along the thick red line segment should be part of the target object, avoiding the undesired concavity in the figure. Similarly, as p_2 is on the target object, enforcing the GVF shape prior should be able to remove the small hole in the final segmentation. (b) shows object i is included inside object j with a minimum separation distance of δ^{ij} between the boundary of two objects. (c) shows object i and j are mutually excluded from each other, with a minimum separation distance of δ^{ij} in between. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

obtained by "deforming the pre-segmentation with the constraint of the GVF tree (forest).

Fig. 3(a) shows an example where the GVF shape prior helps maintain the shape of pre-segmentation. Suppose the smaller light brown region is pre-segmentation and the larger dark brown region is the final segmentation. If the GVF paths connecting the foreground voxels p_1, p_2 to the object core of pre-segmentation are indicated by the thick black arrows, then the thick red line segments should not be labeled as background by enforcing the GVF shape prior, thus avoiding the dramatic shape deviation from pre-segmentation.

2.3. Penalty functions for the GVF-based shape prior

For a single voxel p, the GVF-based shape prior can be enforced by the following penalty function between p and all $q \in GP(p)$.

$$\phi(f_p, f_q) = \infty \cdot [f_p = 1, f_q = 0], \tag{2}$$

where $[\cdot]$ is the indicator function which returns 1 when the enclosed condition is true, and 0 otherwise. The penalty function $\phi(f_p, f_q)$ returns ∞ when p belongs to the object, while q is in the background. Otherwise, the penalty function returns 0. In a minimization problem, this ensures that any solution violating the GVF-based shape prior is not valid.

Due to the possible large number of voxels along the GVF path, enforcing $\phi(f_p, f_q)$ between p and all $q \in GP(p)$ naively is computationally expensive, especially considering that this process has to be repeated for every voxel $p \in P$. Fortunately, it can be shown that the GVF-based shape prior (i.e., the monotonicity of the boundary surface of the target object with respect to all the GVF paths) can be sufficiently enforced by using only $\phi(f_p, f_{q_0}), \phi(f_{q_0}, f_{q_1}), \ldots, \phi(f_{q_{N-1}}, f_{q_N})$ where $q_0, q_1, \ldots, q_N = GP(p)$ (Wu and Chen, 2002; Veksler, 2008). Thus, the penalty function only need to be enforced between p and its nearest neighbor indicated by the discretized GVF vector $\mathbf{h}_p^D$. This leads to the observation that to enforce the GVF-based shape prior for all voxels $p \in P$, we only need to iterate through the discretized GVF vector field $\mathbf{h}^D$ ($\left|\mathbf{h}^D\right| = \left|P - C\right|$), i.e.,

$$\sum_{\vec{pq} \in \mathbf{h}^D} \phi(f_p, f_q) = \sum_{\vec{pq} \in \mathbf{h}^D} \infty \cdot [f_p = 1, f_q = 0]. \tag{3}$$

2.4. MRF formulation

Suppose N_P is the neighborhood system in image P, then the overall MRF energy with the GVF-based shape prior for segmentation is

$$E(\mathbf{f}_P) = \sum_{p \in P} D_p(f_p) + \sum_{(p,q) \in N_P} \theta \cdot V_{pq}(f_p, f_q) + \sum_{\vec{pq} \in \mathbf{h}^D} \phi(f_p, f_q). \tag{4}$$

$D_p(f_p)$ is a data term describing the appearance information of voxel p. The more likely it belongs to the foreground, the smaller $D_p(f_p = 1)$ is, while the larger $D_p(f_p = 0)$ is. $V_{pq}(f_p, f_q)$ is a pairwise smoothness term based on the boundary/edge information between neighboring voxels p and q. If $f_p = f_q$, then $V_{pq}(f_p, f_q) = 0$; otherwise, $V_{pq}(f_p, f_q)$ is a nonnegative number that is inversely proportional to the likelihood of an edge between the neighboring voxels p and q. In general, the smoothness term encourages a smooth boundary between the object and the background. The more likely an edge lies in between the neighboring voxels, the smaller the smoothness penalty is. The smoothness term coefficient θ balances the data term and the smoothness term. The larger θ is, the more emphasis is put on the boundary/edge information instead of the appearance information. The third term in Eq. (4) enforces the GVF-based shape prior as in Eq. (3).

2.5. Solving MRF by minimum *s*-excess problem

Now we show how to encode the MRF formulation with the GVF-based shape prior as a *minimum s-excess problem*, which can be exactly solved by computing a minimum *s*-*t* cut (Wu and Chen, 2002; Hochbaum, 2001). The minimum *s*-excess problem in a directed graph $G = (V, E)$ is defined as a partition of the graph vertices into two disjoint sets according to the sum of weights on specific vertices and edges. More specifically, suppose every vertex $v \in V$ carries an arbitrary weight $w(v)$ and every directed edge $e \in E$ carries a nonnegative weight $c(e) \geq 0$, then the minimum *s*-excess problem seeks a vertex subset $H \subseteq V$ such that the cost $\gamma(H)$ is minimized, with $\gamma(H) = \sum_{v \in H} w(v) + \sum_{(u,v) \in Eu \in H, v \in \bar{H}} c(u, v)$, where $\bar{H} = V - H$. In another word, the *s*-excess cost of a vertex set (also called *source set*) H is defined as the sum of vertex weights for all vertices inside the set, and the edge weights for all edges from the set to the complementary of the set.

We construct a graph such that when the minimum *s*-excess problem is solved, the voxels corresponding to the vertices in the source set H will be labeled as foreground. Each voxel p in the image

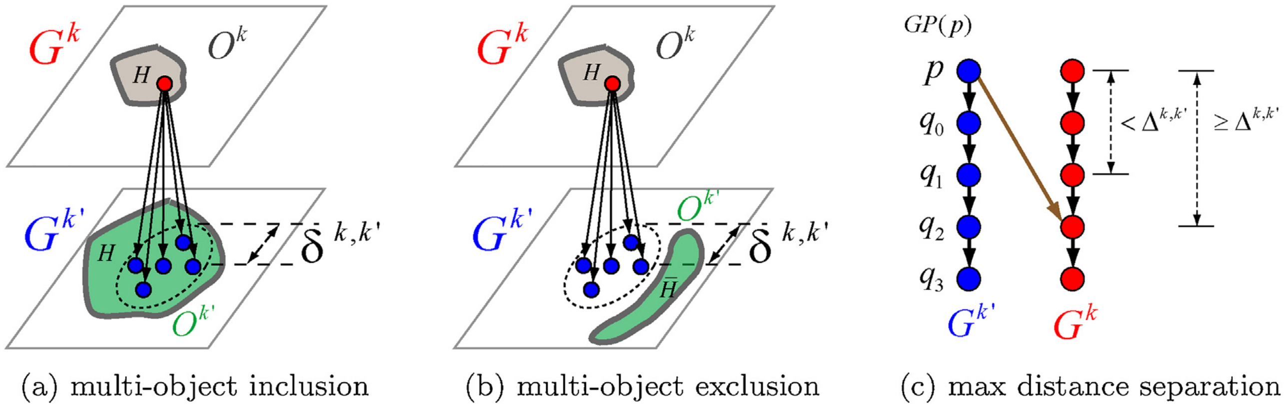

Fig. 4. Illustration of how to enforce the minimum surface separation constraints for the inclusion and the exclusion interaction relations, and how to enforce the maximum surface distance constraints when both objects share the same pre-segmentation. (a) and (b) show that additional edges within a "cone" of radius $\delta^{k,k'}$ need to be added to enforce the minimum distance between the nested and the excluded objects, respectively. (c) shows the added edges between subgraphs for enforcing the maximum distance between two object surfaces sharing the same pre-segmentation and thus the same GVF paths.

defines exactly one vertex v_p in the graph. The three terms in Eq. (4) are correctly encoded in the graph by the following means:

- The data term is encoded by assigning vertex weight $w(v_p) = D_p(1) - D_p(0)$, i.e., the data term difference when being labeled as foreground from background. In this way, the sum of the vertex weights in a source set H is $\sum_{v_p \in H} w(v_p) = \sum_{v_p \in H}(D_p(1) - D_p(0)) = \sum_{v_p \in H} D_p(1) + \sum_{v_p \in \bar{H}} D_p(0) - \sum_{p \in P} D_p(0)$. Note the first two terms are exactly the data term in Eq. (4) and the third term is a constant.
- The smoothness term is encoded by adding edges $(v_p, v_q), \forall(p, q) \in N_P$ and assigning edge weights $c(v_p, v_q) = V_{pq}(1, 0)$. Thus, the sum of the "excess" edge weights (i.e., edges from source set to outside) will be $\sum_{v_p \in H, v_q \in \bar{H}} c(v_p, v_q) = \sum_{v_p \in H, v_q \in \bar{H}} V_{pq}(1, 0) = \sum_{(p,q) \in N_P} V_{pq}(f_p, f_q)$, because all neighboring voxel pairs have the same label, except the ones with corresponding vertices $v_p \in H$ and $v_q \in \bar{H}$.
- The GVF-based shape prior term is enforced by adding edges $(v_p, v_q), \vec{pq} \in \mathbf{h}^D$ with an infinite edge weight. Similar to the edges for the smoothness term, it is easy to show that these edges correctly encodes the GVF-based shape prior term in Eq. (3). Any minimal s-excess solution would avoid including these edges as "excess" edges due to the associated infinite edges weight, thus enforcing the GVF-based shape prior.

Computing the minimum s-excess problem in the constructed graph will return a source set H, which corresponds to voxels segmented as foreground. All other voxels are labeled as background.

2.6. Simultaneous multi-objects segmentation

Medical image applications often require *multiple* objects to be segmented simultaneously, e.g., prostate and bladder (Song et al., 2009, 2010), white matter and gray matter in brain (Oguz and Sonka, 2014). It is advantageous to incorporate the interaction prior information between multiple objects when possible. Fig. 3(b) and (c) show examples of inclusion and exclusion interactions between two objects i and j, with a minimum distance δ^{ij} between them.

To simultaneously segment K objects ($K > 0$), we need to determine K binary variables $f_p^1, f_p^2, \ldots, f_p^K \in \{0, 1\}$ for each voxel p, with $\sum_{k=1}^{K} f_p^k \leq 1$. Here, $f_p^k = 1$ indicates voxel p is part of the k-th object, and vice versa. The energy function $E(\mathbf{f}_P)$ is then defined as $\sum_{k=1}^{K} E(\mathbf{f}_P^k)$. $E(\mathbf{f}_P^k)$ is the energy function for segmenting object O^k, defined as same as in Eq. (4).

We build K subgraphs $G^k = (V^k, E^k)$, $k = 1, \ldots, K$ to represent the K binary variables for each voxel. In every subgraph G^k, a vertex v_p^k is created for every voxel p. The idea is that within every subgraph, the graph is built in a way similar to Section 2.5 to encode the region term, the boundary term and the GVF shape prior. Since our proposed GVF shape prior is embedded directly in the voxel space, we can utilize the method proposed by Delong and Boykov (2009) to add additional edges between two subgraphs corresponding to two interacting objects to enforce the minimum separation distance constraints.

Now we discuss inclusion and exclusion scenarios separately since they differ slightly in graph construction. Note although we define interaction relationship only between two objects at a time, we can segment multiple interacting objects simultaneously since one object can interact with multiple other objects.

Inclusion/nesting relationship: Suppose object O^k is inside object $O^{k'}$ with a minimum distance of $\delta^{k,k'}$ between their boundaries. Without loss of generality, we assume $v_p^k \in H \Leftrightarrow f_p^k = 1$, i.e., the vertex for pixel p in subgraph G^k belonging to the source set in the solution to the s-excess problem corresponds to labeling pixel p as part of the k-th object O^k (as presented in Section 2.5). For both subgraphs G^k and $G^{k'}$, the vertex weights for the region term, the edge weights for the boundary term, and the edge weights for the GVF shape prior are assigned exactly the same as presented in Section 2.5.

To enforce the minimum separation constraints, we add edges $(v_p^k, v_{p'}^{k'})$ with an $+\infty$ weight for all voxels p' such that $\|\vec{pp'}\| \leq \delta^{k,k'}$. As shown in Fig. 4(a), these edges are from the vertex for voxel p in subgraph G^k to all vertices within the minimum distance from the voxel p in subgraph $G^{k'}$. The rationale is that if vertex v_p^k is in the source set, then vertex $v_{p'}^{k'}$ must also be in the source set. In another word, if voxel p is part of the k-th object, then all voxels p' that are "close" to it (closer than $\delta^{k,k'}$) must also be part of the k'-th object when we assume $v_p^{k'} \in H \Leftrightarrow f_p^{k'} = 1$. This ensures that the boundary of the k-th object is always inside the k'-th object with a minimum distance of $\delta^{k,k'}$ between them.

Exclusion relationship: Suppose object O^k does not overlap with object $O^{k'}$ with a minimum distance of $\delta^{k,k'}$ in between. Without loss of generality, we assume $v_p^k \in H \Leftrightarrow f_p^k = 1$, i.e., the vertex for pixel p in subgraph G^k belonging to the source set in the solution to the s-excess problem corresponds to labeling pixel p as part of the k-th object O^k. For subgraph G^k, the vertex weights for the region term, the edges and edge weights for the boundary term and GVF shape priors, are the same as presented in Section 2.5. For subgraph $G^{k'}$, the vertex weights for the region term and the edges and edge weights for the boundary term are also the same as presented in Section 2.5. The edge weights between the two subgraphs for the minimum separation distance constraints are also assigned in the same way as for the inclusion relationship.

However, we do need to flip the meaning of vertices in subgraph $G^{k'}$ such that $v_p^{k'} \in H \Leftrightarrow f_p^{k'} = 0$. Correspondingly, the edge directions for the GVF shape priors in subgraph $G^{k'}$ need to be reversed.

First, we analyze why flipping the meaning of $v_p^{k'}$ and adding edges $(v_p^k, v_p^{k'})$, $\|\vec{pp'}\| \leq \delta^{k,k'}$ with $+\infty$ weights correctly enforces the minimum distance separation constraints. Referring to Fig. 4(b), if voxel p is part of the k-th object, then vertex v_p^k is in the source set in the solution to the s-excess problem. The edges with an infinite weight ensure that vertex $v_{p'}^{k'}$ with $\|\vec{pp'}\| \leq \delta^{k,k'}$ also belongs to the source set. As we flip the meaning of vertices in subgraph $G^{k'}$, that is, a vertex in the source set corresponds to the voxel being labeled as *not* part of the k'-th object, any voxel that is within $\delta^{k,k'}$ from object O^k are not part of $O^{k'}$. This enforces the exclusion relation between the two objects with a minimum separation distance.

Now that we flip the mapping from vertices $v_{p'}^{k'}$ to labeling variables $f_{p'}^{k'}$, we have to reverse the directions of the edges for the GVF shape priors in subgraph $G^{k'}$ in order to encode the same GVF shape prior energy term in Eq. (3).

It is also possible to enforce the *maximum* surface distance between pairs of nested surfaces when they share the same pre-segmentation. In this case, the two surfaces share the common GVF paths which serve as graph columns in the grid space. This common column structure enables the maximum surface distance constraint to be enforced using techniques used in the LOGISMOS method (Bai et al., 2014).

More specifically, suppose object O^k is inside $O^{k'}$ and the two object boundaries are at most $\Delta^{k,k'}$ distance apart. Assume $GP(p) = q_0, q_1, \ldots, q_N$ is the shared GVF path connecting voxel p to the object core for both objects O^k and $O^{k'}$, we add an edge $(v_p^{k'}, v_{q_i}^k)$ such that $q_i \in GP(p)$ with $\|\vec{pq_i}\| \geq \Delta^{k,k'}$ and $\|\vec{pq_{i-1}}\| < \Delta^{k,k'}$, i.e., q_i is the first voxel along the GVF path that is more than $\Delta^{k,k'}$ away from voxel p. This ensures that if voxel p is in object $O^{k'}$, then all voxels that are at least $\Delta^{k,k'}$ away along the GVF path must also be in object O^k. In another word, the boundary distance between the two objects are never greater than the maximum separation distance constraints. Fig. 4(c) shows how edges are added to enforce the maximum distance constraints while nested surfaces share the same pre-segmentation.

3. Method validation

We validated the proposed method on two applications: the brain tissue segmentation in MRI T1 images, and the prostate and bladder segmentation in CT images. The brain tissue segmentation aims to segment two *nested* surfaces which separates three different types of tissues (white matter, gray matter, and cerebrospinal fluid), while the second application aims to segment two *mutually exclusive* objects.

3.1. Experiment settings

For both applications, we first obtained an initial segmentation, i.e., pre-segmentation. The proposed GVF shape prior was defined based on this pre-segmentation. The final segmentation output accuracy was assessed by two metrics: Dice similarity coefficient (DSC) and average symmetric surface distance (ASSD). The *Dice similarity coefficient* is used to measure how well two volumes overlap with each other. Assume A and B are two volumes, then the DSC between the two volumes is defined as $2|A \cap B|/(|A| + |B|)$, which ranges between 0 and 1. The larger the DSC is, the better the two volumes are aligned, with 1 indicating a perfect overlapping. The *average symmetric surface distance* (ASSD) is used to measure how close two segmented surfaces are. Assume given segmentation volume A, boundary surface S_A is defined to be the set of voxels in A that has at least one neighboring background voxel. Let $d(x, S_A)$ denote the shortest distance between a point x and any point on the boundary surface S_A of segmentation A. The ASSD between two segmentations A and B is defined as $ASSD = (\sum_{a \in S_A} d(a, S_B) + \sum_{b \in S_B} d(b, S_A))/(|S_A| + |S_B|)$. It measures the average distance from any point on a contour/surface to the other contour/surface. The ASSD metric has a range of $[0, +\infty]$. The smaller ASSD is, the better the two segmented surfaces/contours agree with each other. If $ASSD = 0$, then the two segmentations are identical.

3.2. Brain segmentation: data and compared methods

For the brain tissue segmentation, 18 T1-weighted scans of normal subjects from the Internet Brain Segmentation Repository (IBSR) were used (commonly know as "IBSR 18" in the literature) (Wels et al., 2011; Valverde et al., 2015). The image size is $256 \times 128 \times 256$ with the voxel spacing range from $0.837 \times 1.5 \times 0.837\,\mathrm{mm}^3$ to $1 \times 1.5 \times 1\,\mathrm{mm}^3$. Each scan comes with a brain mask marking voxels inside skull. Each voxel inside the brain mask is labeled by expert as one of three labels: white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF). Using this public dataset allows us to directly compare the proposed method to various state-of-art methods.

Valverde et al. (2015) validated 10 brain segmentation algorithms on the IBSR dataset, aiming to evaluate a wide set of available techniques and tools. We compare our method to the results reported. These ten algorithms are summarized below: FAST (Zhang et al., 2001) uses expectation maximization with K-means initialization to optimize an MRF model; SPM5 (Ashburner and Friston, 2005) and SPM8 (Ashburner et al., 2012) are two versions of the SPM toolbox based on the iterative Gaussian mixture model, atlas registration and bias correction techniques; GAMIXTURE (Tohka et al., 2007) estimates a Gaussian mixture model by a genetic algorithm; ANN (Tian and Linan, 2007) implements a self organizing map for clustering image data; FCM (Pham, 2001) uses a fuzzy c-means clustering algorithm; KNN (de Boer et al., 2009) performs k-nearest neighbor searches based on an automated registration of prior probability atlases; SVPASEG (Tohka et al., 2010) applies an iterative conditional modes (ICM) method with an initialization based on a genetic algorithm to optimize an MRF model; FANTASM (Pham, 2001) extends FCM by adding a spatial term in the objective function; and PVC (Shattuck et al., 2001) builds a maximum-a-posteriori (MAP) model with the ICM optimization. We refer readers to Table 1 in Valverde et al. (2015) for software sources and versions used in reporting these numbers.

In addition to the above 10 methods, we also included the results reported by the following state-of-art segmentation methods: (i) the discriminative model-constrained EM approach, which combines supervised discriminative modeling and unsupervised statistical expectation maximization (EM) segmentation into an integrated Baysian framework (Wels et al., 2011); (ii) the adaptive Markov modeling based on mutual information (Awate et al., 2006); (iii) the hidden Markov chain model (Bricq et al., 2008) which unifies partial volume effect, bias field correction, and a probabilistic atlas; and (iv) the prior knowledge driven multiscale segmentation (Akselrod-Ballin et al., 2007) which embeds an atlas prior in a multi-scale pyramid.

We also compared our method to Atropos (Avants et al., 2011), the segmentation tool distributed with the state-of-art brain MRI registration software Advanced Normalization Tools (ANTs) (Rajchl et al., 2016; Ou et al., 2014; Klein et al., 2009). It solves n-class segmentation formulated as a Baysian optimization problem by an expectation maximization (EM) method with Markov random field (MRF) modeling. We used the software downloaded from the author's website of the software (ANTs), with the same pre-segmentation used in our proposed method (Section 3.3.2) as the initial labeling for EM iterations, The MRF modeling smoothing parameter was set to 0.2 with every voxel as an MRF variable after a grid search for optimal parameters.

3.3. *Brain segmentation: workflow*

3.3.1. *Preprocessing*

Each image was first cropped by the brain mask to only include voxels inside the skull, reoriented to have the standard orientation ('RAI' orientation), and resampled to have unit voxel spacing. We used 2-fold cross validation for this experiment. The 18 images were randomly split into two disjoint set, with one of them being the training set and the other as the test set. Then the roles of training/testing set were reversed for both sets. To account for the inter-scan intensity profile difference, we first computed a mean brain histogram of the training set. Then we performed histogram matching to match every test image to the histogram.

3.3.2. *Pre-segmentation*

The BRAINSABC software[1] (Young Kim and Johnson, 2013) was used to classify every voxel within the brain mask into one of the three tissue classes: WM, GM, and CSF. BRAINSABC first performs a deformable registration to align the input image to an atlas, then conducts an atlas-based tissue classification using an expectation-maximization approach (Van Leemput et al., 1999). The largest connected component consisting of the white matter voxels is used as the pre-segmentation of the inner surface, i.e., the WM-GM boundary. The largest connected component consisting of the white matter *and* the gray matter voxels was used as the pre-segmentation of the outer surface, i.e., the GM-CSF boundary. The GVF-based shape priors were defined with respect to these pre-segmentations. To simplify notation, we would call the WM-GM boundary as the WM surface, and the GM-CSF boundary as the GM surface. The WM surface was set to be included in the GM surface with a minimum distance of 1 mm.

3.3.3. *Energy term design*

The random forest method using a simple four-dimensional feature vector was applied to generate voxel-wise probability maps for all three types of tissues. The four features were the *normalized* x, y, z coordinates (normalized to [0,1] range) and the intensity of each voxel. Using the techniques in Section 2.6, two variables $f_p^{\mathrm{WM}}, f_p^{\mathrm{WM/GM}} \in \{0, 1\}$ were introduced for each voxel p. If $f_p^{WM} = 1$, then p is WM. If $f_p^{WM/GM} = 1$, then p is WM *or* GM. The corresponding data terms $D(\cdot)$ in Eq. (4) for the two labels are defined in Eqs. (5) and (6). Fig. 6(b) shows an example WM region term $D^{WM}(f^{WM} = 1)$ computed from input image Fig. 6(a).

$$\begin{aligned} D_p^{\mathrm{WM}}(f_p^{\mathrm{WM}} = 1) &\quad \propto -\log Pr(\mathrm{WM}|p), \\ D_p^{\mathrm{WM/GM}}(f_p^{\mathrm{WM/GM}} = 1) &\quad \propto -\log Pr(\mathrm{WM/GM}|p). \end{aligned} \tag{5}$$

$$D_p^{l}(f_p^l = 0) \propto -\log(1 - Pr(l|p)),\ l \in \{\mathrm{WM}, \mathrm{WM/GM}\}. \tag{6}$$

Note that the two surfaces WM-GM and GM-CSF are highlighted at different intensity ranges. It makes sense to adjust the contrast of images when computing the boundary terms $V(\cdot)$ in Eq. (4). More specifically, the intensity is transformed by a sigmoid function controlled by parameters α^l and β^l in Eq. (7), which enhances the contrast roughly within the range $[\beta^l - 3\alpha^l, \beta^l + 3\alpha^l]$. Here $l \in \{\mathrm{WM}, \mathrm{GM/WM}\}$. These parameters are set to $(\alpha^{\mathrm{WM}}, \beta^{\mathrm{WM}}) = (10, 180)$ and $(\alpha^{\mathrm{GM/WM}}, \beta^{\mathrm{GM/WM}}) = (30, 120)$ by visually checking the training set images. Since the histograms of all images are matched to the reference mean histogram of the training datasets, the contrast adjustment parameters above should be robust to different scans. Fig. 6(d) shows the result of contrast adjustment before computing the WM-GM boundary terms.

$$\bar{I}_p^l \propto \frac{1}{1 + \exp\left(-\frac{I_p - \beta^l}{\alpha^l}\right)}. \tag{7)(8)(9}$$

The gradient between neighboring voxels p and q is then used as the boundary term (Eq. (8)). The more intensity difference of the neighboring voxels has, the smaller the boundary penalty is. On the other hand, if the neighboring voxels have very similar intensities, then the boundary penalty will be large, strongly encouraging the neighboring voxels to share the same labeling. The parameter σ in Eq. (8) was set to 0.1 by visually checking boundary terms generated on the training set scans for a good contrast between the edge and non-edge locations. Fig. 6(e) and (f) shows example boundary terms for the vertical and horizontal edges computed from the contrast-adjusted image (Fig. 6(d)). The smoothness term coefficient *theta* in Eq. (4) was set as 1.0. All parameters were fixed with the same values for all test images.

3.4. *Brain segmentation: results*

One example segmentation from different views is shown in Fig. 5. The four columns show the original image, the pre-segmentation, the manual contour and the segmentation of our GVF-based shape prior method, respectively. We can see the segmentation aligns well with the manual contour. Fig. 7 shows the 3D rendering of WM-GM and GM-CSF surfaces from three subjects. The proposed method is shown to agree well with the manual contour in the 3D space as well. In Fig. 6(i) and (l), we also show the WM segmentation by the proposed method with the region term, the GVF shape prior, and the boundary terms shown in Fig. 6(b), (c), (e) and (f). Compared to the pre-segmentation (Fig. 6(g) and (j)), the segmentation by the proposed method aligns better to the reference manual contour (Fig. 6(h) and (k)).

Quantitatively we compared the Dice coefficient (DSC) of WM and GM regions with various published methods (Section 3.2). Table 1 lists the DSC means and standard deviations of various state-of-art methods and the proposed method. Our method achieved the best accuracy on GM, with a large margin compared to most of the other compared methods. The WM accuracy was also very competitive. In fact, if we visualize the results in the column chart (Fig. 8), it was obvious that the proposed method resulted in accurate segmentation of both WM and GM. In contrast, most of the other compared methods performed well on WM, but relatively poorly on GM. The surface distance error metric (ASSD) was not reported since many published methods did not report them.

For each brain MRI scan, it took 1 h 31 min for BRAINSABC to do the pre-segmentation. Given the pre-segmentation, it only took, on average, 1 min 40 s to compute the GVF-based shape prior, the energy terms, and the final segmentation. In total, it took an average of 1 h 32 min and 40 s to process one scan.

3.5. *Brain segmentation: sensitivity to pre-segmentation*

Since the shape prior is defined based on a pre-segmentation, it is natural to see how sensitive the method is with respect to the pre-segmentation. To study the sensitivity, we artificially warped

[1] Available online: https://github.com/BRAINSia/BRAINSTools

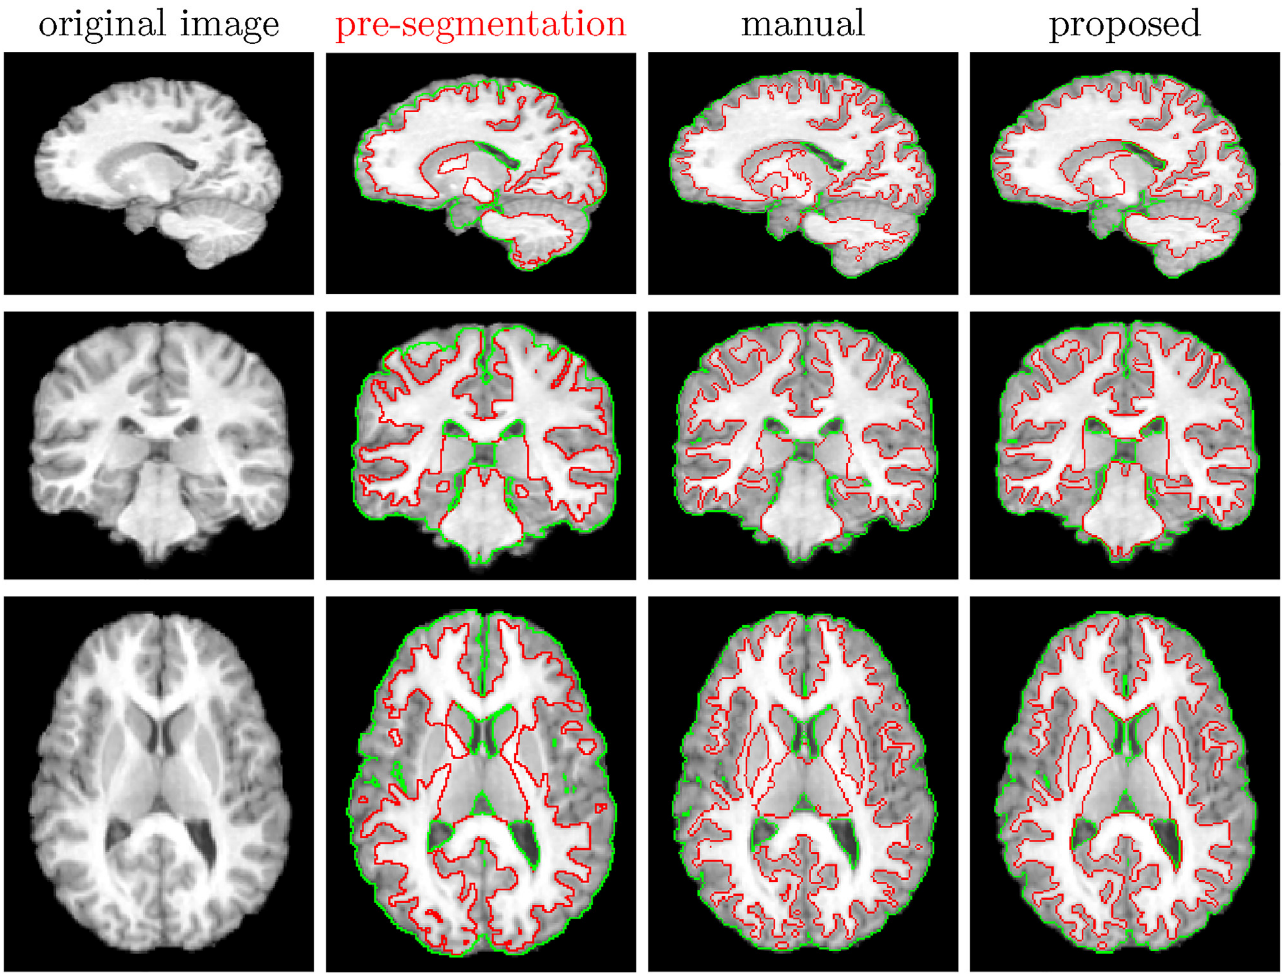

Fig. 5. Example segmentation using the proposed GVF-based shape priors. The red contour is the WM-GM boundary. The green contour is the GM-CSF boundary. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
The DSCs achieved by the proposed method compared to those by the other state-of-art methods on WM and GM in brain tissue segmentation. Bold items are the best accuracy for each tissue type among all methods. Our method resulted in the most accurate GM segmentation, and was very competitive on the WM segmentation.

Method	WM	GM
FAST (Zhang et al., 2001)	0.89 ± 0.02	0.74 ± 0.04
SPM5 (Ashburner and Friston, 2005)	0.86 ± 0.02	0.68 ± 0.07
SPM8 (Ashburner et al., 2012)	0.88 ± 0.01	0.81 ± 0.02
GAMIXTURE (Tohka et al., 2007)	0.87 ± 0.02	0.78 ± 0.08
ANN (Tian and Linan, 2007)	0.87 ± 0.03	0.70 ± 0.07
FCM (Pham, 2001)	0.88 ± 0.03	0.70 ± 0.06
KNN (de Boer et al., 2009)	0.86 ± 0.03	0.79 ± 0.03
SVPASEG (Tohka et al., 2010)	0.86 ± 0.02	0.81 ± 0.03
FANTASM (Pham, 2001)	0.88 ± 0.03	0.71 ± 0.06
PVC (Shattuck et al., 2001)	0.83 ± 0.07	0.70 ± 0.08
Awate et al. (2006)	$\mathbf{0.89 \pm 0.02}$	0.81 ± 0.04
Akselrod-Ballin et al. (2007)	0.87	0.86
Bricq et al. (2008)	0.87 ± 0.02	0.80 ± 0.06
Wels et al. (2011)	0.87 ± 0.05	0.83 ± 0.12
Atropos (Avants et al., 2011)	0.86 ± 0.03	0.82 ± 0.03
Proposed	0.87 ± 0.03	$\mathbf{0.88 \pm 0.02}$

the image on a $20 \times 20 \times 20$ grid using B-Spline. At each B-Spline grid point, a random Gaussian noise deformation force with a mean of zero and a standard deviation of σ^{ptb} along each dimension was applied. Then the proposed GVF-based shape prior segmentation algorithm was run using a pre-segmentation perturbed with various perturbation magnitudes σ^{ptb}. The DSC and ASSD between

Table 2
Sensitivity of proposed method with respect to deformation in pre-segmentation. The first column is the perturbation magnitude parameter σ^{ptb} value. The surface error (ASSD) values are reported in unit mm.

ptb mag	Pre-seg DSC		Final DSC		Pre-seg ASSD		Final ASSD	
	WM	GM	WM	GM	WM	GM	WM	GM
0	0.794	0.775	0.866	0.878	1.01	0.93	0.59	0.67
2	0.772	0.754	0.865	0.877	1.08	0.99	0.59	0.68
3	0.754	0.736	0.864	0.877	1.15	1.03	0.60	0.68
5	0.712	0.696	0.862	0.876	1.32	1.13	0.62	0.71
10	0.621	0.608	0.852	0.868	1.68	1.36	0.69	0.79
15	0.559	0.545	0.841	0.857	1.93	1.52	0.75	0.88
20	0.512	0.498	0.832	0.848	2.15	1.68	0.81	0.97

the perturbed pre-segmentation and the manual expert contour are measured to quantify the deformation magnitude (reported as 'pre-seg DSC/ASSD' in Table 2). The DSC and ASSD between the output of the proposed algorithm and the manual expert contours are reported to validate how sensitive the proposed method is to the pre-segmentation (reported as 'final DSC/ASSD' in Table 2).

As the perturbation magnitude parameter σ^{ptb} increases from 0 to 5, the output segmentation of the proposed method barely changed at all, even though the DSC between the deformed pre-segmentation and manual contour dropped by about 0.04 for both WM and GM. As the perturbation magnitude increased even larger, the accuracy of the proposed method started to decrease more significantly. But even at the largest perturbation magnitude

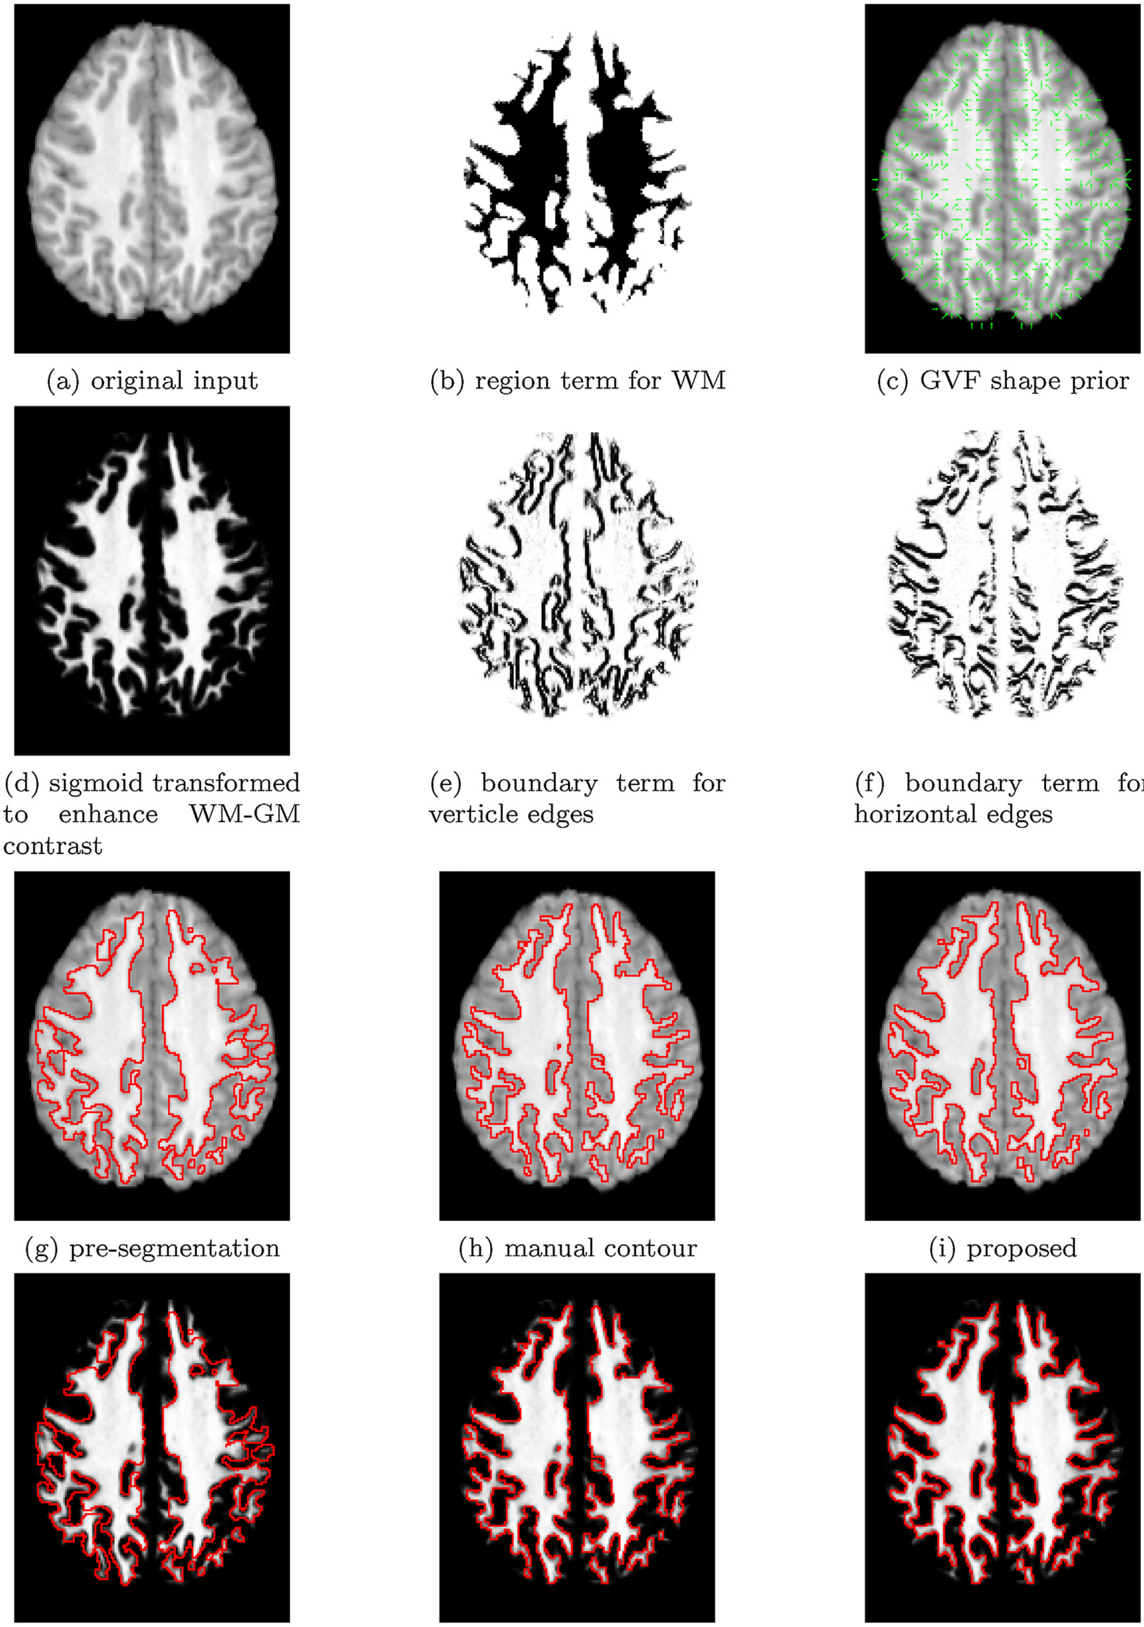

(a) original input (b) region term for WM (c) GVF shape prior

(d) sigmoid transformed to enhance WM-GM contrast (e) boundary term for verticle edges (f) boundary term for horizontal edges

(g) pre-segmentation (h) manual contour (i) proposed

(j) pre-segmentation (k) manual contour (l) proposed

Fig. 6. Example cost images and the GVF shape prior for segmenting WM ((b), (c), (e), (f)), and the corresponding pre-segmentation ((g), (j)), manual contour ((h), (k)), and final segmentation by the proposed method ((i), (l)), overlaid on original input image (a) and the sigmoid transformed image (d) to highlight the edge between WM and GM. To improve visibility in (c), the GVF shape prior at one voxel is shown within every 5-by-5 voxel region.

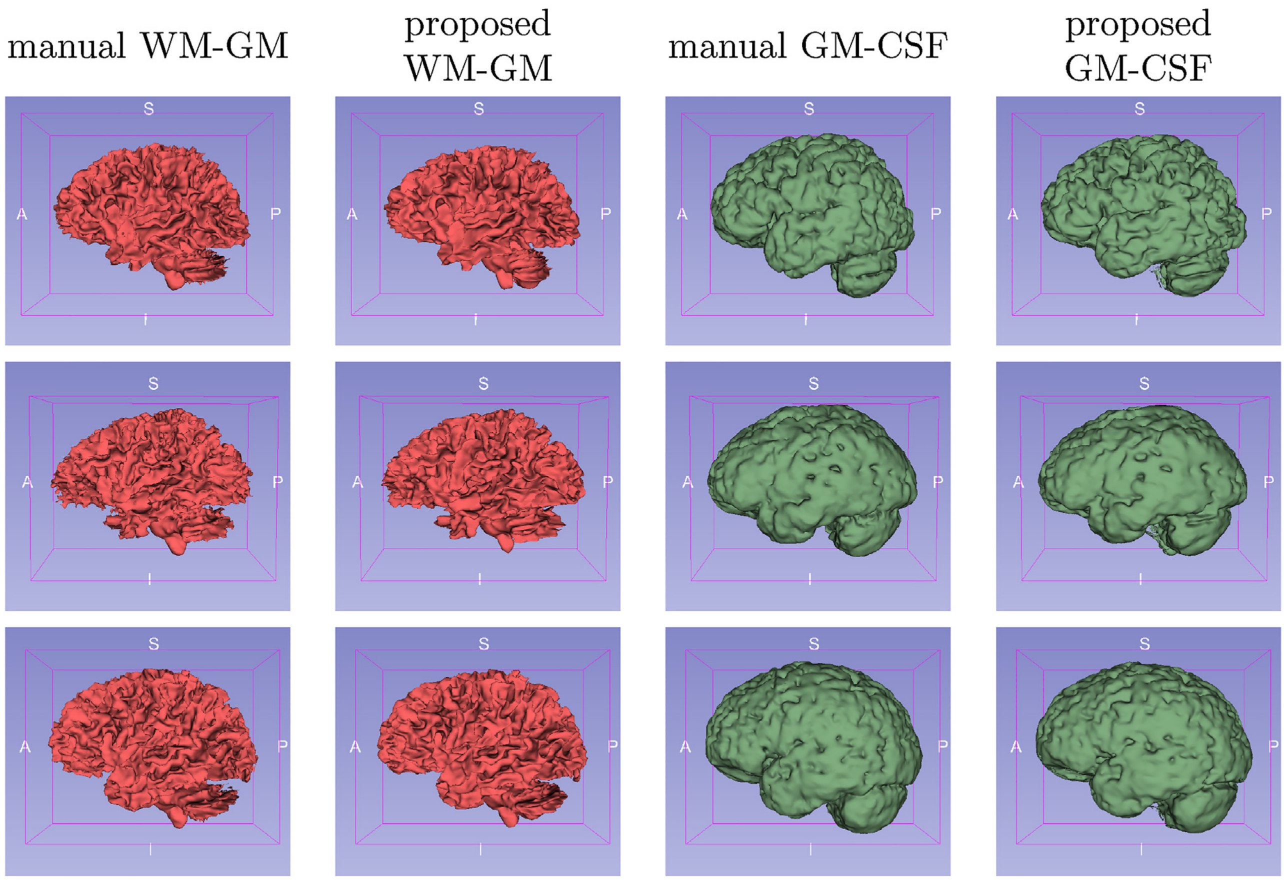

Fig. 7. 3D renderings of WM-GM and GM-CSF surfaces of manual contouring and the proposed GVF-based shape prior segmentation, from 3 different subjects.

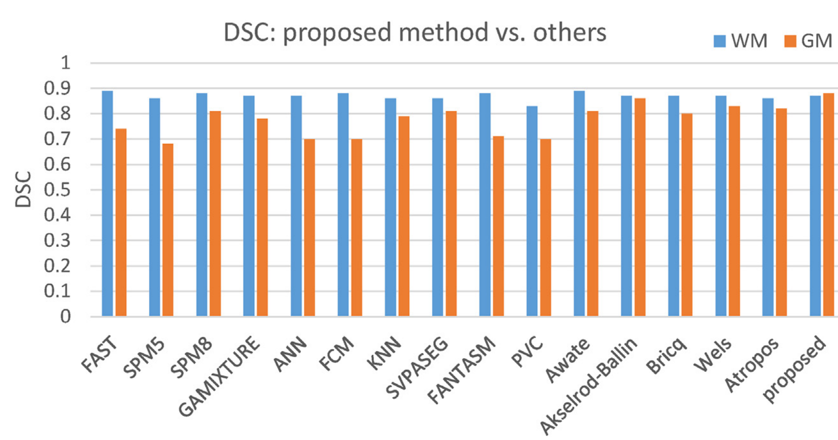

Fig. 8. Brain tissue segmentation shown with the DSC bar graph for the proposed method compared to those for the other state-of-art methods, generated from statistics reported in Table 1. The proposed method achieved accurate segmentation on both types of tissue, while most of the compared methods performed well on WM, but relatively poor on GM.

($\sigma^{ptb} = 20$), the DSC only droped around 0.03 for both WM and GM, while the deformed pre-segmentation only had 0.512 and 0.498, respectively, in DSC. In terms of surface distance error ASSD, while the pre-segmentation was even perturbed by 1.14 mm on WM and 0.75 mm on GM, the output ASSD's only drop by 0.22 and 0.30 mm, respectively.

Fig. 9 demonstrates the flexibility of the proposed GVF-based shape prior. Although the pre-segmentation (the white mask in Fig. 9(c)) is not very accurate compared to the manual delineation of white matter (the red contour), the segmentation using GVF-based shape prior (the white mask Fig. 9(d)) is still able to take advantage of the shape information, and achieve good boundary alignment with the manual contour. The GVF-based shape prior computed from the pre-segmentation is drawn as small blue arrows in Fig. 9(c) and (d) (you may zoom in to see them clearly).

3.6. Bladder/prostate segmentation: data and compared methods

For the prostate and bladder segmentation, 21 volumetric CT images from different patients with prostate cancer were used. The image spacing ranges from $0.98 \times 0.98 \times 3.00\,\text{mm}^3$ to $1.60 \times 1.60 \times 3.00\,\text{mm}^3$. Expert manual contours of bladder and prostate were available for each scan. Eight scans randomly selected were used as the training set. The other 13 scans were

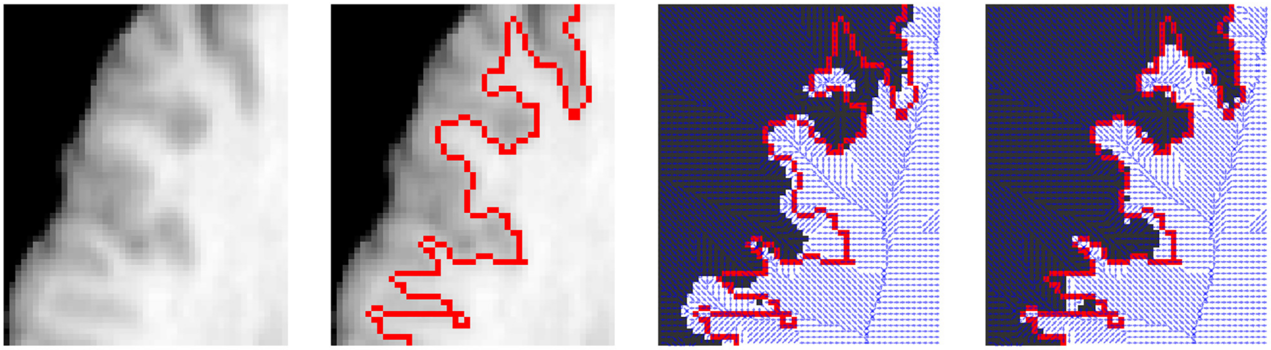

Fig. 9. The GVF-based shape prior is flexible. Even though the pre-segmentation boundary (white mask in (c)) is not quite accurate around the white matter boundary compared to the manual delineation (red contours), the segmentation using the GVF-based shape prior (white mask in (d)) is still able to take advantage of the shape information, and achieve accurate segmentation with respect to the manual contour. Small blue arrows in (c) and (d) are the representation of the GVF-based shape prior computed from the pre-segmentation (zoom in to see more details). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

used as the testing set. All parameters were tuned based on the training set. All reported results are from the testing set.

Our proposed method was compared against the mesh-based approach to enforce the shape priors (Song et al., 2010). It iteratively evolves the bladder and prostate *a prior* shape meshes using the LOGISMOS method. In each iteration, the bladder and prostate shape meshes are first evolved independently, and then deformed jointly by defining a "mutually interacting region" to ensure that the two meshes do not intersect. A soft shape prior is also incorporated in the prostate mesh to encourage the shape conforming to a trained model.

3.7. Bladder/prostate segmentation: workflow

3.7.1. Pre-segmentation

We used the same pre-segmentation used by Song et al. (2010). The prostate shape is relatively consistent across scans, thus a point distribution model built from training images was aligned to each test scan. The Procrustes analysis method was employed to get the prostate pre-segmentation. Due to the large shape variations in bladder, a geodesic active contour method was used to generate the pre-segmentation for bladder. The prostate and bladder were set to be two exclusive objects with a minimum separation distance of 0 mm.

3.7.2. Energy term design

Similar to the brain tissue segmentation, a random forest model with a four-dimensional feature vector (normalized $x/y/z$ coordinates and intensity) was trained to generate probability maps of bladder and prostate. Due to the poor soft-tissue contrast in CT scans, a Chan–Vese model (Chan and Vese, 2001) was also trained to more accurately model the intensity difference between prostate and bladder. The two probability maps generated by the random forest model and the Chan–Vase model were added together with an equal weight as the data term ($D_p(\cdot)$ in Eq. (4)). The smoothness term design was the same as that in brain tissue segmentation (Section 3.3).

3.8. Bladder/prostate segmentation: results

The segmentation accuracy of the proposed method is presented in Table 3. A 2-tailed paired t-test was perform to show that our method was significantly better than Song et al.'s method (Song et al., 2010) on bladder ($p < 0.05$) with both DSC and ASSD metrics. For the prostate, both methods were not statistically different, while our method showed a better DSC on average but slightly worse ASSD values.

Table 3
Bladder/prostate segmentation accuracy compared to mesh-based method (Song et al., 2010).

Metric	Object	Mesh based method (Song et al., 2010)	Proposed
DSC	Bladder	0.933 ± 0.015	**0.941** ± **0.017**
	Prostate	0.824 ± 0.036	**0.835** ± **0.015**
ASSD (mm)	Bladder	0.89 ± 0.26	**0.79** ± **0.28**
	Prostate	**1.48** ± **0.47**	1.64 ± 0.28

The bold values indicate the best results can be achieved among compared methods.

Table 4
Dice coefficient achieved by the methods with v.s. without the GVF shape prior.

Application	Object	Without GVF	With GVF
Brain	WM	0.871 ± 0.033	**0.875** ± **0.030**
	GM	0.880 ± 0.019	**0.884** ± **0.014**
Bladder/prostate	Bladder	0.933 ± 0.020	**0.941** ± **0.017**
	Prostate	0.815 ± 0.027	**0.835** ± **0.015**

The bold values indicate the best results can be achieved among compared methods.

Fig. 10 shows example segmentation results of the proposed method and the mesh-based LOGISMOS approach. The LOGISMOS method tends to produce spiky surfaces compared to the smooth appearance of the manual contour. The spiky error at the top of bladder in the mesh-based segmentation (third column in Fig. 10) might caused by the high risk of mesh folding, which has to be accounted for by special procedures discussed in Section 1.1. In contrast, the proposed segmentation method with the GVF-based shape priors does not need to handle the mesh self-intersection/folding problem, yielding accurate and smooth segmentation of bladder and prostate.

3.9. Effectiveness of the GVF shape priors

Table 4 shows the segmentation accuracy comparison between the methods with and without using the GVF-based shape priors on both brain segmentation and bladder/prostate segmentation. For both applications, the segmentation accuracy was improved by incorporating the GVF shape priors. The paired t-test for each object in each application demonstrated that the improvement was

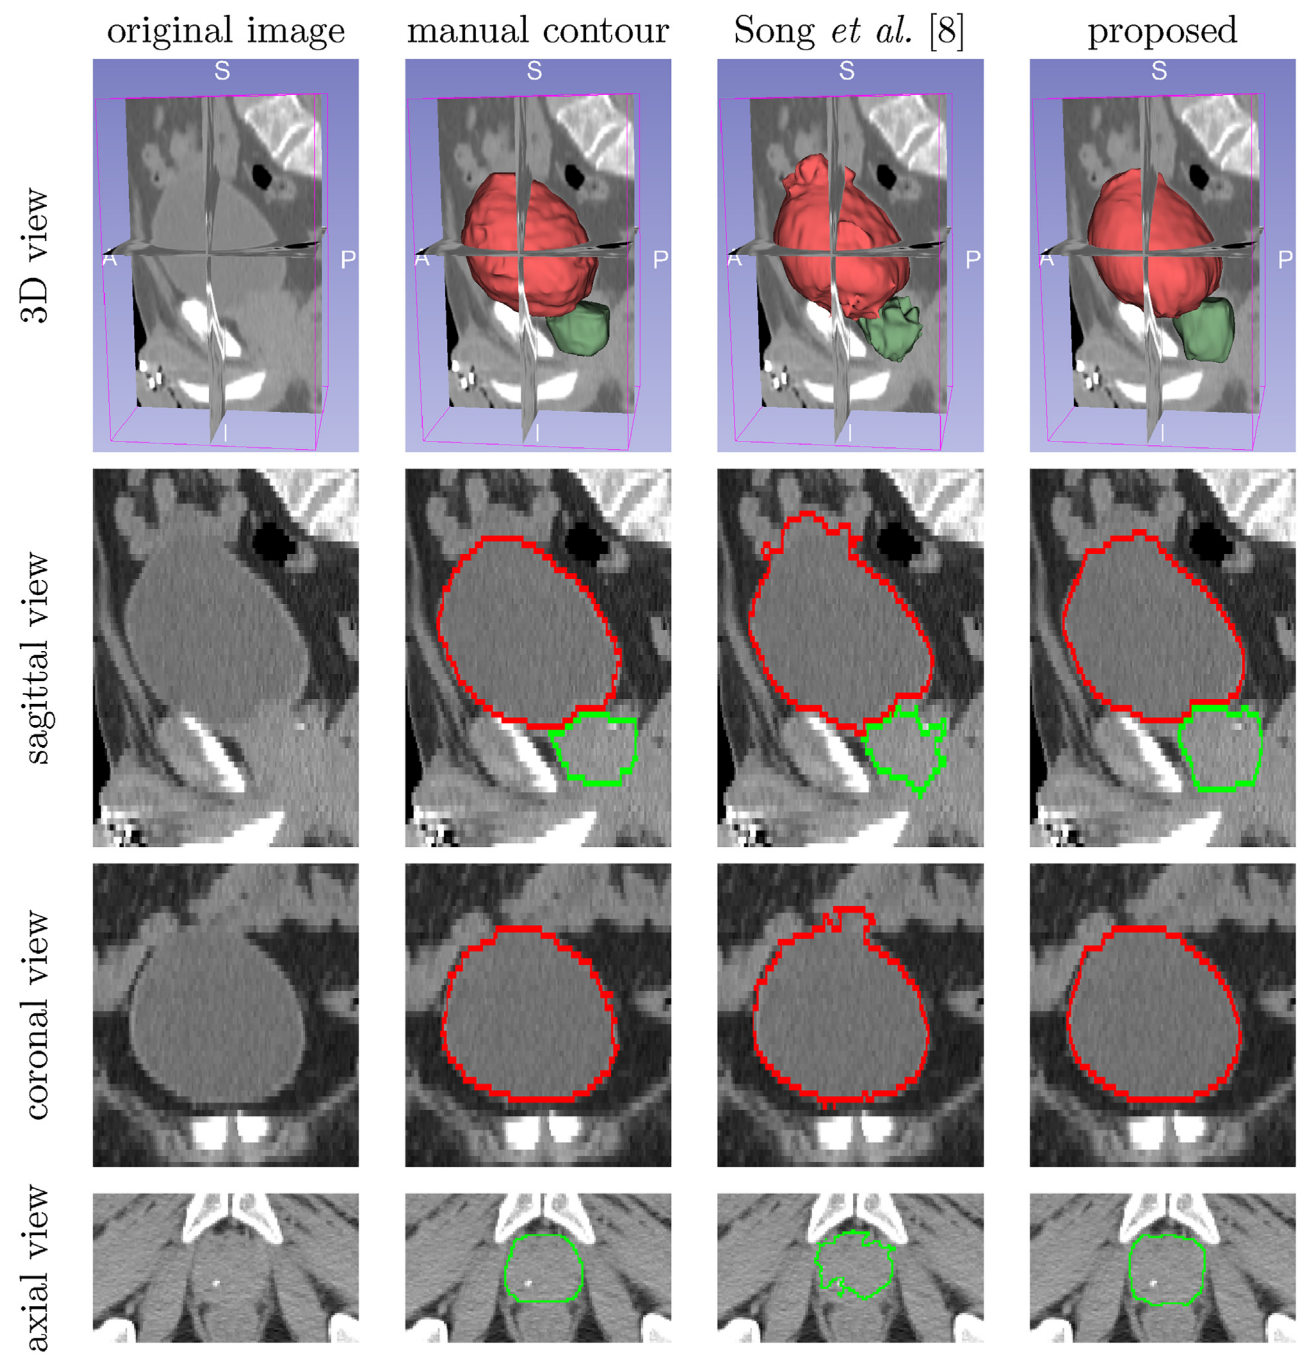

Fig. 10. An illustrative example of segmentations by the proposed method and the mesh-based LOGISMOS method (Song et al., 2010). The red and green contours are for bladder and prostate, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

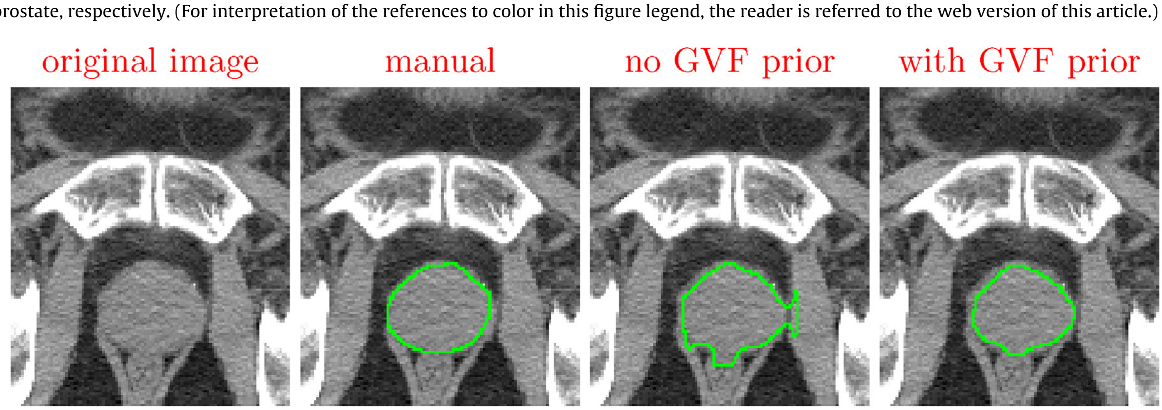

Fig. 11. The GVF-based shape prior enables to avoid leakage to the other adjacent tissues for prostate segmentation.

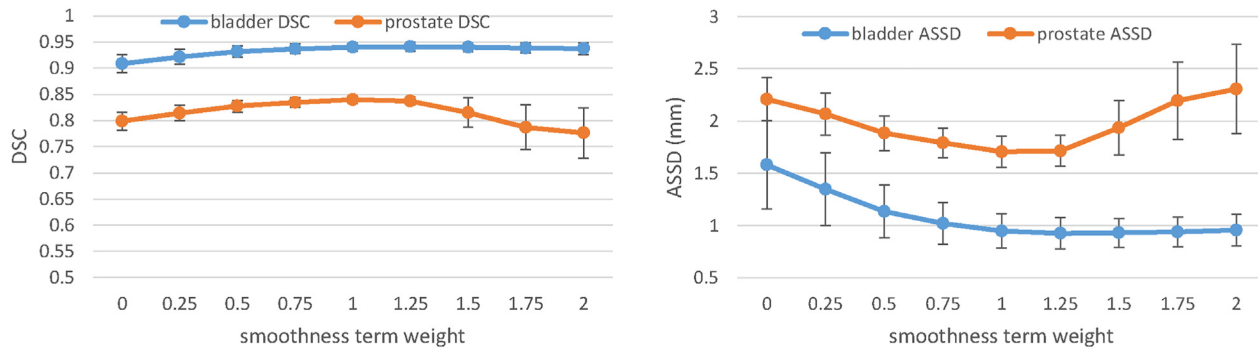

(a) DSC changes with the varying smoothness term coefficient θ values (b) ASSD changes with the varying smoothness term coefficient θ values

Fig. 12. Accuracy of the proposed method on bladder and prostate segmentation when varying the smoothness term coefficient θ values in Eq. (4). The error bars show the 95% confidence intervals for the mean DSC (ASSD) reported.

statically significant ($p < 0.05$). By incorporating the GVF shape priors, the method produced segmentation with less variation. Fig. 11 shows an example of prostate segmentation where the proposed GVF-based shape prior could avoid the leakage to other tissues by regulating the shape of target object.

3.10. Sensitivity to the varying parameters

To explore the performance variation of the proposed method with regard to the varying parameter values, the smoothness term coefficient θ in Eq. (4) was changed from 0.0 to 2.0 with a step size of 0.25 varied in the bladder/prostate segmentation application. All other parameters were fixed. The corresponding mean DSC and ASSD, and the corresponding 95% confidence interval for each θ value were recorded in Fig. 12. We can see that as we varied the smoothness term coefficient around 1, the segmentation accuracy first plateaus around the reported accuracy, and then decreases mildly. This shows the importance of utilizing both the appearance information from the data term and the boundary information from the smoothness term.

4. Discussion

The novelty of the proposed method compared to the star shape and the geodesic star methods: Mathematically, our method shares the same formulation as the star shape method (Veksler, 2008) and the geodesic star method (Gulshan et al., 2010) to enforce the corresponding shape priors. However, our method is more general and flexible, especially for medical image segmentation.

The star-shape prior (Veksler, 2008) is not flexible enough to handle the complex shapes presented in medical image analysis. The digital rays used to enforce the star-shape prior is not well analogous their corresponding Euclidean rays, with image regions uncovered by the digital rays. This results in no shape control for the corresponding part of the target object. In order to cover a complex shape, a number of star centers need to be put *accurately* inside the object, which can be especially cumbersome for 3D image segmentation. More problematically, the resulting optimization problem (given more than 3 star centers) is computationally intractable (Gibson et al., 2014).

Although the geodesic star method (Gulshan et al., 2010) reduces the number of star centers and alleviates the need for accurate star center locations, it still requires a very careful placement of the star centers for segmenting an object with a complex shape. The geodesic star method uses a geodesic distance which combines Euclidean distance with image gradient information. With the user input star centers, the method partitions the image domain into

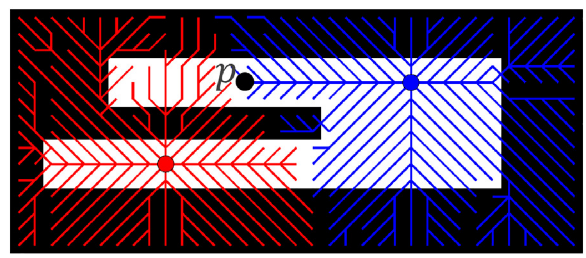

Fig. 13. Intuitive star center locations of the geodesic star method (Gulshan et al., 2010) may lead to undesirable result.

Voronoi cells, each including exactly one star center, such that each voxel in a Voronoi cell is closer to the center of the cell than any other star centers based on its geodesic distances. The assumption here is that the part of the target object in each Voronoi cell is a geodesic star with respect to the star center of the cell. This assumption requires a careful placement of the star centers. In addition, due to the contribution of gradient information in computing geodesic distance, it is sometime counter intuitive to place the star centers for users. Fig. 13 shows an example geodesic forest on a synthetic image, which is generated by putting 0.9 weight on gradient avoidance and 0.1 weight on Euclidean distance, as suggested in Gulshan et al. (2010). The white region is the target object and the red and blue dots are two star centers intuitively placed by the user. The image is partitioned into two Voronoi cells, the red one and the blue one, based on the geodesic distance. The pixel p is correctly covered by the blue geodesic tree due to the use of the gradient information in its geodesic distance, even though it is closer to the red star center in Euclidean distance. However, the upper left part of the object beyond pixel p is still covered by the red geodesic tree. If we enforce the star shape prior according to that geodesic forest, we would either not be able to get the top left part of the object, or wrongly include the gap in the red Voronoi cell between the upper and lower half of the object as part of the target object. In summary, if the geodesic star centers are not well-placed, the geodesic star method may fail the segmentation.

In comparison, our method does not require user to specify seed points (star centers). The pre-segmentation can potentially be obtained by a fully automated method, which saves a large amount of human labor. It is flexible since the pre-segmentation can have complex shapes. The GVF shape prior depends on pre-segmentation shapes instead of raw gradients in original image, and thus is more robust to weak and noisy gradients in medical images. The requirement for pre-segment is also relatively loose: the pre-segmentation can also be either over-segmented or under-segmented or both at

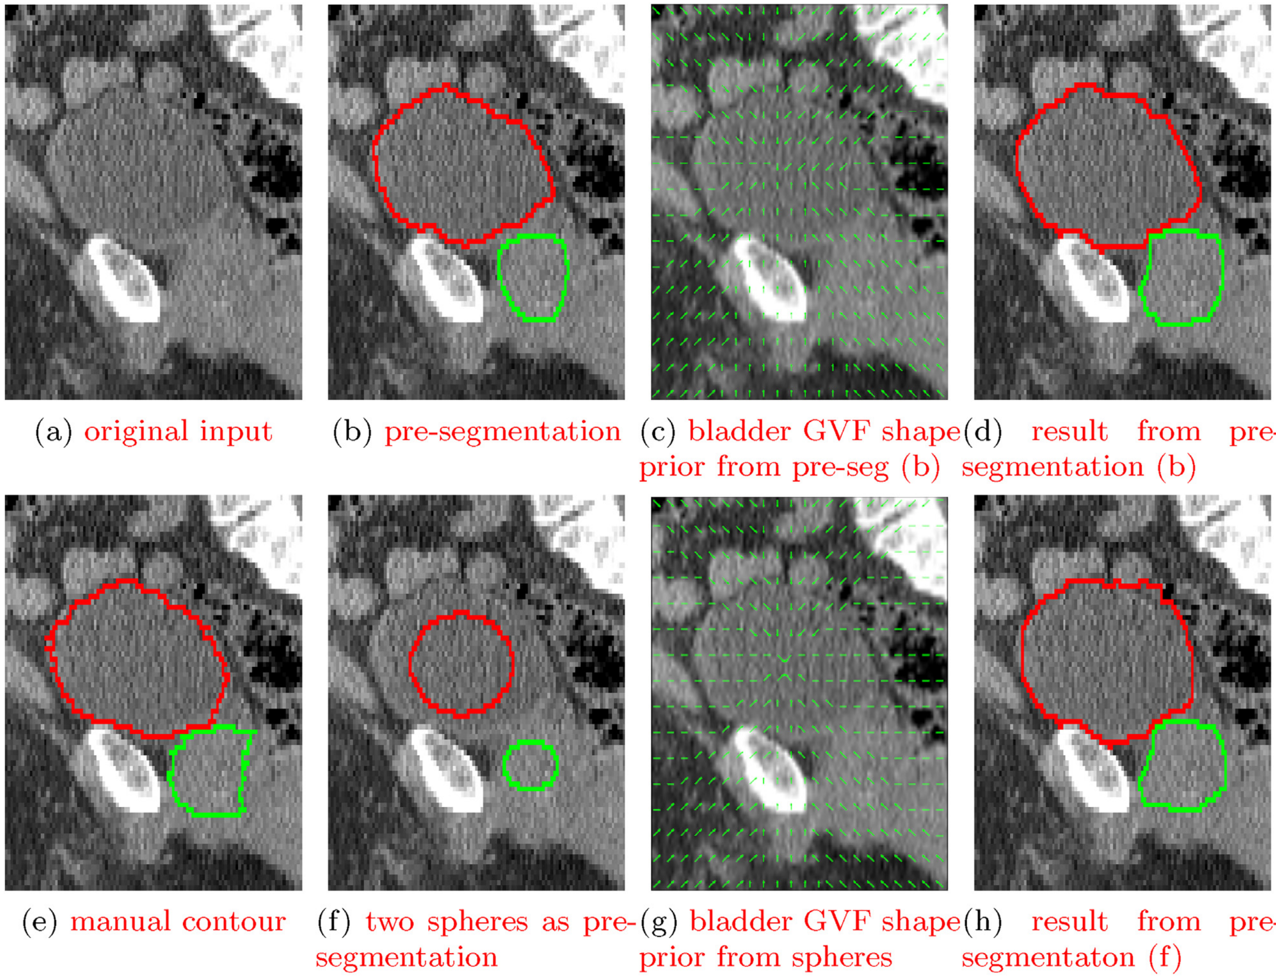

Fig. 14. Results obtained by rough pre-segmentation consisting of two under-segmented spheres. Note although the two spheres in (f) are under-segmented by a relatively large margin, the resulting GVF shape prior in (g) still sufficiently captures the bladder shape compared to the GVF shape prior in (c) computed from a relatively accurate pre-segmentation in (b). The final segmentation (h) is also close to the segmentation (d) obtained from relatively accurate pre-segmentation (b). To improve the visibility in (c) and (g), the GVFs are shown with every 5-by-5 voxel region.

different regions. Thus, our proposed GVF shape prior is more suitable to medical image analysis due to the less human labor, flexible shape representation, and robustness to noisy images.

Inaccurate pre-segmentation: Sometimes the pre-segmentation may be relatively rough. However, the flexibility of the GVF shape prior may still be able to capture useful shape information and regularize the final segmentation. To demonstrate that, we used two spheres that were under-segmented by a large margin as pre-segmentations for bladder and prostate (Fig. 14(f)). The resulting segmentation has Dice coefficients of 0.924 for bladder and 0.852 for prostate (Fig. 14(h)). In comparison, the result from the original relatively accurate pre-segmentation has Dice coefficients of 0.936 for bladder and 0.862 for prostate (Fig. 14(d)). The illustrative examples in Fig. 14 show that even a well under-segmented pre-segmentation can still provide sufficient shape information for the target object using the proposed GVF shape prior. The resulting segmentation can still be of high quality.

The fidelity to the GVF-based shape prior: In our proposed method, the deviation to the shape prior can be penalized by either an infinite or a finite weight, thus achieving the control of the fidelity to the specified shape prior in different degrees.

For our conducted experiments, we set the deviation penalty to an infinite weight due to the flexibility of the GVF-based shape prior. For example, it allows a large range of changes in local surface orientations. The local surface displacement is also well captured by the GVF shape prior. With our proposed method, we do scan-wise pre-segmentation to hopefully capture the topology of the target object correctly, and then use the flexible GVF shape prior to recover from moderate amount of pre-segmentation errors. The tolerance of our method to the errors introduced in the shape prior has been demonstrated in Section 3.5 by perturbing pre-segmentation. The perturbation to the pre-segmentation basically introduces errors in the shape prior. Fig. 9c and d visually shows how our method can recover from those errors even by using an infinite penalty.

If necessary, we can make use of a finite penalty to enforce the GVF shape prior to allow some GVF paths to trade off the prior information for what the data is telling. However, using a finite penalty introduces one extra parameter to tune, thus one may prefer to use an infinite penalty when the pre-segmentation is not too bad and rely on the flexibility of the GVF shape priors to refine it.

Comparison to topology correction algorithm: The topology correction algorithm (Bazin and Pham, 2007) proposed by Bazin and Pham uses the fast marching technique to propagate the topology of a given template to a binary or probabilistic segmentation. For example, the initial template can be as simple as a 2D donut ring for any single 2D object with one hole in it (the hole in initial template must be aligned with the hole in the object). Their algorithm strictly enforces the classic definition of *global topology*. In contrast, the proposed GVF shape prior enforces strong conformity to *local shape*. For example, the undesired concavity p_1

shown in Fig. 3(a) violates the GVF shape prior. But it would not be corrected by Bazin and Pham's method (Bazin and Pham, 2007). Our proposed method shows better local shape resemblance to pre-segmentation besides the global topology. The proposed GVF shape prior enables multiple interactive object segmentation while maintaining the inclusion/exclusion relationship with minimum separation distance. Such interactions are at least nontrivial to be enforced with Bazin and Pham's method.

Comparison to deep learning methods: Deep learning methods, notably Convolutional Neural Networks (CNN), have been applied widely on medical image segmentation and achieved state-of-art accuracy on multiple applications (Ronneberger et al., 2015; Pereira et al., 2015; Lyksborg et al., 2015; Kamnitsas et al., 2017; Dou et al., 2016). Informative low level features and high level representation models are automatically learned in those multi-layer neural networks. However, most of the networks are applied on 2D images (Ronneberger et al., 2015; Pereira et al., 2015) due to the huge memory footprint of 3D CNN's. Lyksborg et al. (2015) used 2D slices from 3 orthogonal views for segmentation to make use of 3D context information, which is effectively a 2.5D method. Methods using 3D CNN usually have to be applied on small patches of images (Kamnitsas et al., 2017; Dou et al., 2016). To get a complete 3D segmentation mask, images are divided into multiple small patches and segmented independently. The results are then pieced together. This may introduce segmentation inconsistencies around the patch boundary. Moreover, processing patches individually ignores the global location information within the whole image.

In comparison, the proposed method can efficiently process a complete 3D scan, fully utilizing the global spatial and shape information. One future direction is to explore the informative features extracted by deep learning and optimize the segmentation on the whole image using the proposed method.

5. Conclusion

We propose a novel GVF-based shape prior which can be directly embedded in the voxel grid space, fundamentally avoiding the mesh folding problem. Its flexibility and power enables to enforce the geometric interactions in multi-object segmentation in an efficient and straightforward fashion. The proposed segmentation method has been validated on the applications of brain tissue segmentation (two nesting/inclusive objects), and of the bladder/prostate segmentation (two exclusive objects). The experiment results demonstrated either superior or competitive segmentation accuracy compared to the other state-of-the-art methods.

Conflict of interest statement

The authors do not have any conflict of interest.

Acknowledgments

This work was supported in part by the National Science Foundation (NSF) under Grants CCF-1318996 and CCF-1733742, and in part by the National Institutes of Health (NIH) under Grants R01-EB004640 and R01-EB020665.

References

Advanced Normalization Tools (ANTs). URL http://stnava.github.io/ANTs/.
Akselrod-Ballin, A., Galun, M., Gomori, J.M., Brandt, A., Basri, R., 2007. Prior knowledge driven multiscale segmentation of brain MRI. Medical Image Computing and Computer-Assisted Intervention: MICCAI. International Conference on Medical Image Computing and Computer-Assisted Intervention, vol. 10(Pt. 2), 118–126.
Ashburner, J., Friston, K.J., 2005. Unified segmentation. NeuroImage 26 (3), 839–851.
Ashburner, J., Barnes, G., Chen, C., Daunizeau, J., Flandin, G., Friston, K., Gitelman, D., Kiebel, S., Kilner, J., Litvak, V., et al., 2012. SPM8 Manual. Functional Imaging Laboratory, Institute of Neurology.
Aslan, M., Ali, A., Chen, D., 2010. 3D vertebrae segmentation using graph cuts with shape prior constraints. International Conference on Image Processing (ICIP), 2193–2196.
Avants, B.B., Tustison, N.J., Wu, J., Cook, P.A., Gee, J.C., 2011. An open source multivariate framework for n-tissue segmentation with evaluation on public data. Neuroinformatics 9 (4), 381–400.
Awate, S.P., Tasdizen, T., Foster, N., Whitaker, R.T., 2006. Adaptive Markov modeling for mutual-information-based, unsupervised MRI brain-tissue classification. Med. Image Anal. 10 (5), 726–739.
Bai, J., Miri, M.S., Liu, Y., Saha, P., Garvin, M., Wu, X., 2014. Graph-based optimal multi-surface segmentation with a star-shaped prior: application to the segmentation of the optic disc and cup. In: 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI), IEEE, pp. 525–528, http://dx.doi.org/10.1109/ISBI.2014.6867924.
Bazin, P.-L., Pham, D.L., 2007. Topology correction of segmented medical images using a fast marching algorithm. Comput. Methods Programs Biomed. 88 (2), 182–190.
Boykov, Y., Funka-Lea, G., 2006. Graph cuts and efficient *n-d* image segmentation. Int. J. Comput. Vis. 70 (2), 109–131.
Bricq, S., Collet, C., Armspach, J.P., 2008. Unifying framework for multimodal brain MRI segmentation based on Hidden Markov Chains. Med. Image Anal. 12 (6), 639–652.
Chan, T.F., Vese, L.A., 2001. Active contours without edges. IEEE Trans. Image Process. 10 (2), 266–277.
Christ, T., Pálvölgyi, D., Stojaković, M., 2012. Consistent digital line segments. Discrete Comput. Geomet. 47 (4), 691–710.
Chun, J., Korman, M., Nöllenburg, M., Tokuyama, T., 2009. Consistent digital rays. Discrete Comput. Geomet. 42 (3), 359–378.
de Boer, R., Vrooman, H.A., van der Lijn, F., Vernooij, M.W., Ikram, M.A., van der Lugt, A., Breteler, M.M.B., Niessen, W.J., 2009. White matter lesion extension to automatic brain tissue segmentation on MRI. NeuroImage 45 (4), 1151–1161.
Delong, A., Boykov, Y., 2009. Globally optimal segmentation of multi-region objects. IEEE International Conference on Computer Vision, 285–292.
Dornheim, J., Seim, H., Preim, B., Hertel, I., Strauss, G., 2007. Segmentation of Neck Lymph Nodes in CT Datasets With Stable 3D Mass-Spring Models: Segmentation of Neck Lymph Nodes, vol. 14., pp. 1389–1399, http://dx.doi.org/10.1016/j.acra.2007.09.001.
Dou, Q., Chen, H., Jin, Y., Yu, L., Qin, J., Heng, P.-A., 2016. 3D deeply supervised network for automatic liver segmentation from CT volumes. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 149–157.
Gibson, M., Varadarajan, K., Wu, X., 2014. Computing regions decomposable into m stars. In: European Symposium on Algorithms. Springer Berlin Heidelberg, pp. 480–491.
Gulshan, V., Rother, C., Criminisi, A., Blake, A., Zisserman, A., 2010. Geodesic star convexity for interactive image segmentation. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, pp. 3129–3136.
Hochbaum, D.S., 2001. An efficient algorithm for image segmentation, Markov random fields and related problems. J. ACM 48 (4), 686–701.
Internet Brain Segmentation Repository. URL http://www.nitrc.org/projects/ibsr.
Jalba, A., Wilkinson, M., Roerdink, J., 2004. CPM: a deformable model for shape recovery and segmentation based on charged particles. IEEE Trans. Pattern Anal. Mach. Intell. 26 (10), 1320–1335.
Kainmueller, D., Lamecker, H., Heller, M.O., Weber, B., Hege, H.-C., Zachow, S., 2013. Omnidirectional displacements for deformable surfaces. Med. Image Anal. 17 (4), 429–441.
Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Med. Image Anal. 36, 61–78.
Klein, A., Andersson, J., Ardekani, B.A., Ashburner, J., Avants, B., Chiang, M.-C., Christensen, G.E., Collins, D.L., Gee, J., Hellier, P., Song, J.H., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R.P., Mann, J.J., Parsey, R.V., 2009. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. NeuroImage 46 (3), 786–802.
Klinder, T., Ostermann, J., Ehm, M., Franz, A., Kneser, R., Lorenz, C., 2009. Automated model-based vertebra detection, identification, and segmentation in CT images. Med. Image Anal. 13 (3), 471–482.
Li, K., Wu, X., Chen, D.Z., Sonka, M., 2006. Optimal surface segmentation in volumetric images – a graph-theoretic approach. IEEE Trans. Pattern Anal. Mach. Intell. 28 (1), 119–134.
Lyksborg, M., Puonti, O., Agn, M., Larsen, R., 2015. An ensemble of 2D convolutional neural networks for tumor segmentation. In: Scandinavian Conference on Image Analysis. Springer, pp. 201–211.
Martin, S., Troccaz, J., Daanen, V., 2010. Automated segmentation of the prostate in 3D MR images using a probabilistic atlas and a spatially constrained deformable model. Med. Phys. 37 (4), 1579.
McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Oguz, I., Sonka, M., 2014. LOGISMOS-B: layered optimal graph image segmentation of multiple objects and surfaces for the brain. IEEE Trans. Med. Imaging 33 (6), 1220–1235.

Ou, Y., Akbari, H., Bilello, M., Da, X., Davatzikos, C., 2014. Comparative evaluation of registration algorithms in different brain databases with varying difficulty: results and insights. IEEE Trans. Med. Imaging 33 (10), 2039–2065.

Park, J.Y., McInerney, T., Terzopoulos, D., Kim, M.H., 2001. A non-self-intersecting adaptive deformable surface for complex boundary extraction from volumetric images. Comput. Graph. (Pergamon) 25 (3), 421–440.

Pereira, S., Pinto, A., Alves, V., Silva, C.A., 2015. Deep convolutional neural networks for the segmentation of gliomas in multi-sequence MRI. In: International Workshop on Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries. Springer, pp. 131–143.

Petersen, J., Nielsen, M., Lo, P., Nordenmark, L.H., Pedersen, J.H., Wille, M.M.W., Dirksen, A., de Bruijne, M., 2014. Optimal surface segmentation using flow lines to quantify airway abnormalities in chronic obstructive pulmonary disease. Med. Image Anal. 18 (3), 531–541.

Pham, D.L., 2001a. Robust fuzzy segmentation of magnetic resonance images. IEEE Symposium on Computer-Based Medical Systems, 127–131.

Pham, D.L., 2001b. Spatial models for fuzzy clustering. Comput. Vis. Image Underst. 84 (2), 285–297.

Pons, J.P., Boissonnat, J.D., 2007. Delaunay deformable models: Topology-adaptive meshes based on the restricted delaunay triangulation. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1–8.

Rajchl, M., Baxter, J.S.H., Qiu, W., Khan, A.R., Fenster, A., Peters, T.M., Yuan, J., 2016. Fast deformable image registration with non-smooth dual optimization. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 465–472.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 234–241.

Shattuck, D.W., Sandor-Leahy, S.R., Schaper, K.a., Rottenberg, D.a., Leahy, R.M., 2001. Magnetic resonance image tissue classification using a partial volume model. NeuroImage 13 (5), 856–876.

Song, Q., Wu, X., Liu, Y., Smith, M., Buatti, J., Sonka, M., 2009. Optimal graph search segmentation using arc-weighted graph for simultaneous surface detection of bladder and prostate. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (Eds.), Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 827–835.

Song, Q., Liu, Y., Saha, P., Song, X.W.Q., Liu, Y., Liu, Y., Saha, P., Sonka, M., 2010. Graph search with appearance and shape information for 3-D prostate and bladder segmentation. Medical Image Computing and Computer-Assisted Intervention, vol. 13 (Pt. 3), 172–180.

Teschner, M., Kimmerle, S., Heidelberger, B., Zachmann, G., Raghupathi, L., Fuhrmann, A., Cani, M.P., Faure, F., Magnenat-Thalmann, N., Strasser, W., Volino, P., 2005. Collision detection for deformable objects. Comput. Graph. Forum 24 (1), 61–81.

Tian, D., Linan, F., 2007. A brain MR images segmentation method based on SOM neural network. The 1st Engineering, 2007. ICBBE 2007, 686–689.

Tohka, J., Krestyannikov, E., Dinov, I.D., Graham, A.M., Shattuck, D.W., Ruotsalainen, U., Toga, A.W., 2007. Genetic algorithms for finite mixture model based tissue classification in brain MRI. IEEE Trans. Med. Imaging 26 (5), 696–711.

Tohka, J., Dinov, I.D., Shattuck, D.W., Toga, A.W., 2010. Brain MRI tissue classification based on local Markov random fields. Magn. Reson. Imaging 28 (4), 557–573.

Valverde, S., Oliver, A., Cabezas, M., Roura, E., Llad, X., 2015. Comparison of 10 brain tissue segmentation methods using revisited IBSR annotations. J. Magn. Reson. Imaging 41 (1), 93–101.

Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based bias field correction of MR images of the brain. IEEE Trans. Med. Imaging 18 (10), 885–896.

Veksler, O., 2008. Star shape prior for graph-cut image segmentation. European Conference on Computer Vision (ECCV).

Veni, G., Fu, Z., Awate, S.P., Whitaker, R.T., 2013a. Proper ordered meshing of complex shapes and optimal graph cuts applied to atrial-wall segmentation from DE-MRI. Proceedings. IEEE International Symposium on Biomedical Imaging, vol. 2013, 1296–1299.

Veni, G., Fu, Z., Awate, S.P., Whitaker, R.T., 2013b. Bayesian segmentation of atrium wall using globally-optimal graph cuts on 3D meshes. Information Processing in Medical Imaging, vol. 23, 656–667.

Vu, N., Manjunath, B., 2008. Shape prior segmentation of multiple objects with graph cuts. 2008 IEEE Conference on Computer Vision and Pattern Recognition, 1–8.

Wang, X., Metaxas, D., Chen, T., Axel, L., 2008. Meshless deformable models for LV motion analysis. In: 26th IEEE Conference on Computer Vision and Pattern Recognition. CVPR, pp. 1–8.

Weese, J., Kaus, M., Lorenz, C., 2001. Shape constrained deformable models for 3D medical image segmentation. In: Information Processing in Medical Imaging.

Wels, M., Zheng, Y., Huber, M., Hornegger, J., Comaniciu, D., 2011. A discriminative model-constrained EM approach to 3D MRI brain tissue classification and intensity non-uniformity correction. Phys. Med. Biol. 56 (11), 3269–3300.

Wu, X., Chen, D., 2002. Optimal net surface problems with applications. Autom. Lang. Program., 1029–1042.

Xu, C., Prince, J., 1998. Snakes, shapes, and gradient vector flow. IEEE Trans. Image Process. 7 (3), 359–369.

Yin, Y., Song, Q., Sonka, M., 2009. Electric field theory motivated graph construction for optimal medical image segmentation. Graph-Based Representations in Pattern Recognition, 334–342.

Yin, Y., Zhang, X., Williams, R., Wu, X., Anderson, D., Sonka, M., 2010. LOGISMOS – Layered Optimal Graph Image Segmentation of Multiple Objects and Surfaces: cartilage segmentation in the knee joint. IEEE Trans. Med. Imaging 29 (12), 2023–2037.

Young Kim, E., Johnson, H.J., 2013. Robust multi-site MR data processing: iterative optimization of bias correction tissue, classification, and registration. Front. Neuroinf. 7 (November), 29.

Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images through a hidden Markov random field model and the expectation–maximization algorithm. IEEE Trans. Med. Imaging 20 (1), 45–57.