Shape complexes: the intersection of label orderings and star convexity constraints in continuous max-flow medical image segmentation

John S. H. Baxter
Jiro Inoue
Maria Drangova
Terry M. Peters

Shape complexes: the intersection of label orderings and star convexity constraints in continuous max-flow medical image segmentation

John S. H. Baxter,^{a,b,*} Jiro Inoue,^a Maria Drangova,^{a,b} and Terry M. Peters^{a,b}

^aWestern University, Robarts Research Institute, 1151 Richmond Street N., London, Ontario N6A 5B7, Canada

^bWestern University, Biomedical Engineering Graduate Program, 1151 Richmond Street N., London, Ontario N6A 5B7, Canada

Abstract. Optimization-based segmentation approaches deriving from discrete graph-cuts and continuous max-flow have become increasingly nuanced, allowing for topological and geometric constraints on the resulting segmentation while retaining global optimality. However, these two considerations, topological and geometric, have yet to be combined in a unified manner. The concept of “shape complexes,” which combine geodesic star convexity with extendable continuous max-flow solvers, is presented. These shape complexes allow more complicated shapes to be created through the use of multiple labels and super-labels, with geodesic star convexity governed by a topological ordering. These problems can be optimized using extendable continuous max-flow solvers. Previous approaches required computationally expensive coordinate system warping, which are ill-defined and ambiguous in the general case. These shape complexes are demonstrated in a set of synthetic images as well as vessel segmentation in ultrasound, valve segmentation in ultrasound, and atrial wall segmentation from contrast-enhanced CT. Shape complexes represent an extendable tool alongside other continuous max-flow methods that may be suitable for a wide range of medical image segmentation problems. © 2016 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.3.4.044005](https://doi.org/10.1117/1.JMI.3.4.044005)]

Keywords: continuous max-flow segmentation; star convexity constraint; optimization-based segmentation; convex optimization; ASETS library.

Paper 16199PR received Sep. 2, 2016; accepted for publication Nov. 28, 2016; published online Dec. 20, 2016.

1 Introduction

Encoding shape information and constraining the shape of possible segmented objects has long been considered fundamental to incorporating anatomical knowledge in segmentation. Active shape models^{1,2} and the general family of statistical shape models³ use shape information to constrain or guide the evolution of segmentation contours to adhere to a predefined point distribution model for either boundary or skeleton points. Multilevel statistical shape models⁴ embed hierarchical label orderings as a method of simplifying or sparsifying this information for multicompartiment objects or multiple objects in a single scene. Level-set based shape methods do not require an intermediate point representation but constrain shape information directly on the labeling function.⁵ In atlas-registration based segmentation approaches, the shape information is implicitly encoded in the atlas and constraints on the registration algorithm.⁶ Atlas combination methods such as shape-based averaging⁷ attempt to preserve shape information when multiple atlases are used for segmentation. An important disadvantage common to all of these approaches is that the shape information, such as the predefined point distribution model or the shape atlas, is composed of a training set of segmented images. This training set must contain the necessary variability to capture pathology, but limit excessive variability that would degrade performance.

More recently, image segmentation based on Markov random field (MRF) modeling has received a great deal of interest,

especially following the development of efficient graph-cut based solvers⁸ that guarantee optimality for MRFs with the property of submodularity.⁹ Shape information from training images has been incorporated into graph-cuts^{10,11} using an iterative approach similar to that of its level-set predecessors.⁵ These methods guarantee local optimality, but no longer have the global optimality guarantee from traditional graph-cuts.

However, graph-cut techniques are a sufficiently flexible paradigm to encode shape information that is independent of training data. Of particular note are star-shape constraints or “simple star convexity constraints”¹² in which every point in an object must be connected to a single vantage point using a linear path. (Examples of simple star convex objects are shown in Fig. 1) These can be extended to “geodesic star convexity constraints,”¹³ in which the path no longer has to be strictly linear but follows a predefined geodesic. Egger et al. have developed template-based frameworks, which perform a discrete analog of a coordinate system transformation. These can be used to ensure a particular predefined two-dimensional (2-D)^{14,15} or three-dimensional (3-D)^{16,17} shape or extrapolate a shape from user interaction.¹⁸ As these methods use minimal or no a priori training information, often relying solely on the identification of the vantage point, they are better suited for problems where a sufficient body of training segmentations cannot be collected. In addition, they retain the guarantee of global optimality for foreground–background segmentation problems.

Similarly, continuous max-flow segmentation, a continuous-space analog to the discrete “graph space” of traditional MRFs,

*Address all correspondence to: John S. H. Baxter, E-mail: jbaxter@robarts.ca

Fig. 1 Simple star convex objects with vantage points indicated with an “X.”

has gained traction with the development of highly parallelizable duality-based optimization approaches.¹⁹ Simple and geodesic star convexities in this continuous space have also been investigated. Yuan et al.²⁰ used an additional multiplier to allow for unconstrained flow along a predefined geodesic, allowing a single star convex label to represent the prostate. Ukwatta et al.²¹ developed a max-flow propagated level-sets framework using a coordinate system warping approach to ensure star convexity for both the blood vessel as a whole and the lumen. (This coordinate system warping can be seen as a continuous analog to the discrete graph template approach of Egger et al.^{14,18}) Ukwatta’s approach is conceptually important in that it used a combination of geodesic star convex labelings to ensure a more complex ring-shaped vessel wall. However, using coordinate system warping to reformulate shape constraints into topological equivalents introduces a series of problems including computational expense and possible ill-definition and ambiguity when performed on multiple or branching vessels.

The aim of this paper is to develop a framework allowing what we call “shape complexes,” which are “geodesic star convexity” constraints placed on a combination of multiple labels and a label ordering, rather than considering a single individual label (as demonstrated by Yuan et al.²⁰). This approach takes advantage of recent advances regarding multilabel topological constraints in continuous max-flow segmentation theory, and in directed acyclic graph continuous max-flow (DAGMF) segmentation²² in particular. By using both star convexity constraints and label orderings in tandem, much more expressive, yet still general-purpose, shape information can be encoded. As with their predecessors, shape complexes require minimal training data, relying solely on the definition of the label ordering and the placement of vantage points. This makes shape complexes ideal for medical image segmentation problems in which there is complex anatomy (either in the object-of-interest or in adequately modeling the background objects) or anatomy with “walls,” such as vessels and cardiac structures, where multiple star convexity constraints need to be applied in tandem.

2 Background and Methods

For shape complexes, we refine the approach taken by Yuan et al.²⁰ of incorporating geodesic star convexity along a single labeling and combine it with that of Baxter et al.²² of organizing labels in a rooted weighted directed acyclic graph in a strict generalization of hierarchical orderings.^{23,24} Both take a primal-dual optimization approach using augmented Lagrangian multipliers.

2.1 Prior Work on Geodesic Star Convexity Constraints in Max-Flow

In the model developed by Yuan et al.,²⁰ an additional set of multipliers is used to permit flow in a predefined direction, yielding:

$$\begin{aligned} \min_{p_s, p_t, q, \lambda} \int_{\Omega} p_s(x) dx \\ \text{s.t. } |q'(x)| \leq R(x), \\ \lambda(x) \geq 0, \\ p_s(x) \leq D_s(x), p_t(x) \leq D_t(x), \\ \text{div}[q'(x) + \lambda(x)e(x)] - p_s(x) + p_t(x) = 0, \end{aligned} \quad (1)$$

where $D_s(x)$ and $D_t(x)$ are foreground and background data terms, respectively, $R(x)$ is the regularization or smoothness term, e is the local direction of the geodesic path, and q , p_s , and p_t are all flow variables. In this formulation, the foreground background labelings, u and $1 - u$, respectively, are derived from the multipliers on the constraints $p_t(x) \leq D_t(x)$ and $p_s(x) \leq D_s(x)$, respectively. In the framework of Yuan et al., λ , the amount of flow along the predefined direction is explicitly stored and optimized over. This explicit representation can be problematic in that it requires additional memory to store, and that, if not implemented in a similar approach as the spatial flow variables, determining the divergence of the field $q'(x) + \lambda(x)e(x)$ may be difficult. Ultimately, these limitations necessitate a framework in which the variable λ is represented and optimized over implicitly, rather than explicitly.

2.2 Implementation of Geodesic Star Convexity on a Single Label

The additional optimization and storage of the λ variables is not required because $\lambda(x)$, assuming it is non-negative, can be determined exactly as follows:

$$\lambda(x)e(x) = \text{Proj}_{e(x)}[q'(x) + \lambda(x)e(x)],$$

where $q'(x) \cdot e(x) = 0$,

given that all other variables are fixed. By defining a grouped spatial flow term, $q(x) = q'(x) + \lambda(x)e(x)$, the optimization problem expressed in Eq. (1) is equivalent to the more computationally efficient formula:

$$\begin{aligned} \min_{p_s, p_t, q} \int_{\Omega} p_s(x) dx \\ \text{s.t. } |q(x) - \lambda(x)e(x)| \leq R(x), \\ \lambda(x) = \max[q(x) \cdot e(x)/|e(x)|^2, 0], \\ p_s(x) \leq D_s(x), p_t(x) \leq D_t(x), \\ \text{div } q(x) - p_s(x) + p_t(x) = 0, \end{aligned} \quad (2)$$

assuming that the vector field $e(x)$ is normalized to unit length if nonzero, which can be achieved through an initialization step. This memory saving is crucial as multiple star convexity constraints are likely to be present even in a simple shape complex.

2.3 Inclusion into Extendable Max-Flow Framework with Label Orderings

The single-label formulation can be generalized to a framework in which labels are organized in a rooted, weighted directed acyclic graph, as described by Baxter et al.,²² necessitating a novel optimization framework extending DAGMF. This framework optimizes the primal-dual equation:

Algorithm 1 Augmented Lagrangian solution algorithm for Eq. (3).

Topological sort the DAG into ordering $\mathbb{O}$ (begins with source label S)
with reverse ordering $\mathbb{O}^{-1}$ (ends with source label S);

InitializeSolution();

while not converged **do**

UpdateFlows();

for $\forall L$ **do**

$\forall x, u_L(x) \leftarrow u_L(x) - c(\text{div} q_L(x) - \zeta_L(x) + p_L(x));$

end

end

$$\min_u \max_{p,q} \int_{\Omega} p_S(x) dx + \sum_{\forall L \neq S} \int_{\Omega} u_L(x) G_L(x) dx$$

where $G_L(x) = \text{div} q_L(x) + p_L(x) - \sum_{\forall L' \in L.P} w_{(L',L)} p_{L'}(x)$

s.t. $|q_L(x) - \lambda_L(x) e_L(x)| \leq R_L(x),$
 $\lambda_L(x) = \max[q_L(x) \cdot e_L(x) / |e_L(x)|^2, 0]$ given $L \neq S,$
 $p_L(x) \leq D_L(x),$ given $L \in \text{leaves}.$ (3)

This optimization can be addressed using augmented Lagrangian multipliers, as shown in Algorithm 1, using the sub-routines presented in Algorithms 2 and 3. This algorithm is

Algorithm 2 InitializeSolution() subroutine in Algorithm 1.

InitializeSolution()

Clear $u_L(x), q_L(x)$ for all labels;

for each L **in order** $\mathbb{O}^{-1}$ **do**

$\forall x, e_L(x) \leftarrow e_L(x) / |e_L(x)|;$

$\forall x, p_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x);$

$\forall x, \zeta_L(x) \leftarrow \min_{L'.C=\emptyset} D_{L'}(x);$

if $L.C = \emptyset$ **then**

if $L \in \arg \min_{L'.C=\emptyset} D_{L'}(x)$ **then**

$\forall x, u_L(x) \leftarrow 1 / |\arg \min_{L'.C=\emptyset} D_{L'}(x)|;$

else

$\forall x, u_L(x) \leftarrow 0;$

end

end

for each $L' \in L.P \setminus \{S\}$ **do**

$\forall x, u_{L'}(x) \leftarrow u_{L'}(x) + w_{(L',L)} u_L(x);$

end

end

Algorithm 3 UpdateFlows() subroutine in Algorithm 1.

UpdateFlows()

for $\forall L \neq S$ **do**

$\forall x, q_L(x) \leftarrow q_L + \tau \nabla (\text{div} q_L(x) + p_L(x) - p_{L.P}(x) - u_L(x) / c);$

$\forall x, \lambda(x) \leftarrow \max\{0, q_L(x) \cdot e_L(x)\};$

$\forall x, q_L(x) \leftarrow q_L - \lambda(x) e_L(x);$

$\forall x, q_L(x) \leftarrow \text{Proj}_{|q_L(x)| \leq R_L(x)}(q_L);$

$\forall x, q_L(x) \leftarrow q_L + \lambda(x) e_L(x);$

end

Clear $\zeta_L(x)$ for all labels;

for each L **in order** $\mathbb{O}$ **do**

for each $L' \in L.C$ **do**

$\forall x, \zeta_{L'}(x) \leftarrow \zeta_{L'}(x) + w_{(L,L')} p_L(x);$

end

if $L.C \neq \emptyset$ and $L.P \neq \emptyset$ **then**

$\forall x, \sigma_L(x) \leftarrow \zeta_L(x) - \text{div} q_L(x) + u_L(x) / c;$

else if $L = S$ **then**

$\forall x, \sigma_S(x) \leftarrow 1 / c;$

end

end

for each L **in order** $\mathbb{O}^{-1}$ **do**

if $L.C = \emptyset$ **then**

$\forall x, p_L(x) \leftarrow \min[D_L(x), \zeta_L(x) - \text{div} q_L(x) + u_L(x) / c];$

for $L' \in L.P$ **do**

$\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} [\text{div} q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{(L',L)} p_L(x)];$

end

else if $L = S$ **then**

$\forall x, p_S(x) \leftarrow \frac{1}{\sum_{L' \in S.C} w_{(S,L')}^2} \sigma_S(x);$

else

$\forall x, p_L(x) \leftarrow \frac{1}{1 + \sum_{L' \in L.C} w_{(L,L')}^2} \sigma_L(x);$

for $L' \in L.P$ **do**

$\forall x, \sigma_{L'}(x) \leftarrow \sigma_{L'}(x) + w_{(L',L)} [\text{div} q_{L'}(x) + p_{L'}(x) - \zeta_{L'}(x) + w_{(L',L)} p_L(x)];$

end

end

end

trivially parallelizable, making it suitable for acceleration using general purpose programming on graphics processing units (GPGPU). More detailed technical information and a proof-of-correctness for Algorithm 1 are provided in the corresponding technical report.²⁵ GPGPU-accelerated MATLAB implementations of shape complexes are provided in the ASETS library²⁶ for both 2-D image and 3-D volume segmentation problems.

3 Validation

In order to validate the shape complex framework, while maintaining a general focus, several distinct segmentation experiments are employed, including:

1. Synthetic images created to validate the basic properties of the shape complexes framework in comparison to the corresponding DAGMF models. These experiments demonstrate several important features of the algorithm, such as improved accuracy and regularization parameter robustness in a controlled setting. The first of these synthetic experiments is designed to mimic the appearance of vessels in ultrasound;
2. Ultrasound images of the carotid artery were collected to verify that the behavior seen in the synthetic images is reproducible in a medical context. Thus, a similar accuracy and robustness evaluation is performed;

3. Synthetic images mimicking the mitral valve and corresponding transesophageal ultrasound images showing more complicated shape complexes; and
4. Cardiac CT images were collected to test an extreme case of the algorithm with the presence of a very highly heterogeneous background with star convexity constraints applied to a very thin object-of-interest, specifically the left atrial wall.

3.1 Synthetic Image Segmentation

To demonstrate this approach, synthetic volumes were constructed, consisting of a medium intensity background with an embedded structure and white Gaussian noise. This structure has a slightly hypointense center surrounded by a hyperintense boundary. The contrast-to-noise ratio (CNR) between the background and the center is $\text{CNR} \approx 10\%$ and between the hyperintense boundary and each other region is $\text{CNR} \approx 100\%$. Due to the low CNR, segmentation of these images with minimal prior information can be challenging. These images were segmented with a simple hierarchical model using DAGMF²² and using our novel shape complexes, as shown in Fig. 2, with a manually picked centroid or “vantage point” to define a simple star convexity constraint by. The same intensity-based data terms and uniform regularization were used in both images.

This experiment was repeated for a range of regularization values between 10^{-1} and 10^1 , and the dice similarity coefficient

Fig. 2 Synthetic image segmentation problem using DAGMF (e) and DAGMF augmented with shape complexes (f) according to the label ordering in (d) with α referring to the level of regularization. (* a simple star convexity constraint is applied to this label.) Any overlap between segmentations can cause false colors, e.g., green occurs when the result is 50% exterior (cyan) and 50% interior (yellow). The “X” marks the vantage point for the simple star convexity constraint. (a) Gold standard, (b) original image, (c) image + noise, (d) label ordering, (e) DAGMF segmentation, and (f) DAGMF + shape.

Fig. 3 Quantitative segmentation results for each region based on regularization strength. The DSCs are shown on a logarithmic scale approaching 100% DSC. (a) External (cyan) region, (b) ring (magenta) region, and (c) internal (yellow) region.

(DSC) measured for each of the tree labels, as shown in Fig. 3. As expected, the segmentation with shape complexes consistently outperforms the segmentation without shape complexes for all regularization values and displays an additional degree of robustness to the regularization value chosen. This is especially important as star complexes allow for much lower regularization values to be used without sacrificing segmentation quality, which is desirable in terms of preserving less smooth portions of an object.

A second synthetic experiment, shown in Fig. 4, was performed to illustrate the use of shape complexes in a segmentation problem with a distinctly nonhierarchical label ordering. The CNR between regions in Fig. 4(b) is $\text{CNR} = 25\%$. As shown in Fig. 4(e), the regularization parameter was too low to enforce region contiguity under such high noise, which is readily addressed through the addition of a series of star-convexity constraints shown in Fig. 4(f). Again, keeping the regularization value low is essential to capturing the more tortuous boundaries of a segmented object without over-smoothing. Efficiently determining the regularization parameter in these segmentation models, which neither oversmooths or under-smooths, is an area of open research.²⁷

3.2 Ultrasound Vessel Segmentation

To demonstrate the applicability of these shape complexes to medical image segmentation, they were applied to the segmentation of an individual vessel of interest, in particular, the carotid artery, from an ultrasound image. This image was manually segmented into three regions, the background, vessel lumen, and vessel blood pool. Similar to the first synthetic experiment in

Sec. 3.1, the shape complex applied creates a ring-shape prior on the vessel wall. However, this segmentation model takes into account a multicomponent background, with both vessel wall and blood pool components, handling background heterogeneity. The overall model consists of labels for the vessel blood pool B_V , vessel wall W_V , background blood B_K , background hyperintense tissue, such as other vascular walls W_K , and other background tissue K . Segmentation results are shown in Fig. 5.

The data terms used are derived from Bayes' theorem on the voxel intensity:

$$D_L(x) = -\ln P(I(x)|x \in L) \pm \text{bias}, \quad (4)$$

where $P(I(x)|x \in L)$ is the probability of voxel x having intensity $I(x)$ given that it is a part of label L . The constant bias term (positive for background components and negative for foreground components) controls for the shrinking bias, which is especially severe as the background super-label contains two components with the same intensity distribution as the vessel of interest, i.e., labels B_K and W_K with the same intensity as B_V and W_V , respectively. The user-provided seeds for the blood, vascular wall, and background tissue components are given in Fig. 5(b). The regularization is a constant applied to all labels and superlabels shown in Fig. 5(c) with the exception of the vessel wall label, V , which has zero regularization to avoid shrinking bias.

An experiment similar to that shown in Fig. 3 was performed on the vessel ultrasound dataset, varying both the regularization and bias parameters. Quantitative results are shown in Fig. 6. Not only did including the star convexity constraint improve the DSC compared to unconstrained DAGMF at their respective

Fig. 4 Venn diagram segmentation with and without shape complexes. The label ordering is given in (c) (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex was the centroid of the region. Similar to Fig. 2, any overlap between segmentations can cause false colors. (a) Original image, (b) image + noise, (c) label ordering, (d) gold standard, (e) DAGMF segmentation, and (f) DAGMF + shape.

optimal values, but the segmentation became more robust to parameterization, maintaining similarly high DSC over a much broader range of parameters, confirming our earlier observation on synthetic images. This is especially important for medical image segmentation problems in which an exhaustive search through or optimization of the parameter space is difficult to perform, or when parameters are selected interactively,²⁸ as the addition of shape complexes makes selection easier and less sensitive to operator variability.

3.3 Cardiac Valve Segmentation from Ultrasound

Even more complicated shapes can be created by combining simpler shapes. For example, Fig. 4 involved the overlapping of two star convex objects and Fig. 2 showed how star shape priors could be nested to form rings. By combining these two ideas as shown in Fig. 7(a), we can create a shape complex that describes the intersection of two ring-shaped objects. This shape occurs when segmenting structures such as cardiac valve annuli that are a shared boundary between more readily segmented objects, such as blood pools. This model contains six regions: K , the background region; TB and TW , representing the top blood pool and surrounding wall; BB and BW , representing the bottom; and V representing the intersection between them.

Figures 7(b) and 7(c) show a synthetic segmentation example for this shape complex. The image has $CNR = 1$ with no

contrast between the background and blood pools nor between the walls and the valve. With no contrast between similar objects, the shape information is necessary for segmentation. The segmentation result demonstrates that the shape complex can adequately localize the synthetic valve, that is, a minimal region separating the blood pools. Without contrast, there are few defining features of the valve region defined against the walls, representing a limitation of the use of shape complexes alone.

Figure 8 displays a similar experiment using a transesophageal ultrasound image of the anatomy surrounding the mitral valve. Similar to the synthetic image example, the segmentation algorithm cannot accurately segment the valve annulus where it is adjacent to the myocardium. Additionally, the ultrasound image included a partial view of the aorta and aortic valve, which was not accounted for in the model. Because this was not included in the model, the segmentation algorithm had difficulty segmenting said region as the signal intensities contradicted the expectation of the shape term. (The hyperintense valve did not allow the aortic blood pool to be easily grouped with the left ventricular blood pool yet its proximity and attachedness discouraged associating it fully with the background.) Ambiguity in the segmentation result is evidenced by partial coloring.

Both synthetic and ultrasound experiments used uniform regularization and a relatively simple data term:

Fig. 5 Vessel segmentation in ultrasound with and without shape complexes. The label ordering is given in (c) (* a simple star convexity constraint is applied to this label) with the vantage point for the shape complex is marked with an “X.” Similar to Fig. 2, any overlap between segmentations can cause false colors. (a) Original image, (b) user samples, (c) label ordering, (d) manual segmentation, (e) DAGMF segmentation, and (f) DAGMF + shape.

$$D_L(x) = \begin{cases} |I(x) - I_B|, & \text{if } L \in \{\text{BB}, \text{TB}, K\} \\ |I(x) - I_W|, & \text{if } L \in \{\text{BWO}, \text{TWO}\} \\ |I(x) - I_W| - \text{bias}_I, & \text{if } L \in \{I\} \end{cases}, \quad (5)$$

where I_B and I_K were the average intensity of the blood pools and walls, respectively, and bias_I was a bias term to encourage the expansion of the valve annulus label and avoid partitioning the valve region between the BWO and TWO labels. Spatial seeding was provided for the background and two blood pool labels and, for the latter, the centroid of the seed locations was used as the vantage point for the star convexity constraints.

3.4 Atrial Wall Segmentation from Cardiac Computed Tomography

This technique can be used to segment anatomy with boundary structures, such as the atrial wall in contrast-enhanced CT. A challenging aspect of this problem is that the atrial wall has no contrast with surrounding muscle, thus requiring shape complexes to constrain it around the more distinctive atrial blood pool. The segmentation was semiautomated, whereby seeds placed by the user in the atrial blood pool, muscle, fat, and lungs were employed to fit a normal intensity distribution

to each tissue type and define a geodesic star convexity prior for the atrium. A uniform smoothness term was created with a label-specific uniform regularization. The segmentation model is shown in Fig. 9. Best and worst case results are shown in Figs. 10(a)–10(c) and 10(d)–10(e), respectively. Quantitative results are provided in Table 1.

Despite having the same intensity distribution and, therefore, the same data terms, the segmentation successfully differentiated between the atrial and nonatrial blood pools, as well as atrial wall versus other muscular structures. Currently, this segmentation protocol is semiautomatic, requiring some user initialization. In an automated protocol, the user initialization could be replaced by prior knowledge about the Hounsfield distribution of different tissue intensities with some mechanism for automatically estimating the centroid of the atrial blood pool. The smoothness terms allow the blood pools and walls to closely follow perceptible edges in the image even at high regulation values.

There is currently a non-negligible degree of user variability in the method as shown by the difference in accuracy results with respect to the manual segmentation outlined in Table 1. This is to be expected as the seeding not only provides a spatial anchor for the segmentation and defines the vantage points for the shape complex, but also affects the probabilistic data terms used in the optimization process.

4 Discussion and Future Work

The addition of geodesic star convexity and related topological considerations in a general-purpose and application-agnostic manner improves the expressiveness of possible anatomical information that can be encoded in a segmentation problem.

This encoded knowledge can greatly improve segmentation without requiring higher regularization, which can obscure fine structures and detail on the segmented objects. The use of shape complexes to improve the robustness of the optimization algorithm to regularization parameter may be especially

Fig. 6 Quantitative results for the segmentation problem shown in Fig. 5 varying regularization and bias parameters. Blue indicates $DSC \approx 0$ and yellow indicates $DSC \approx 1$, as shown in (g). (a) DAGMF + shape: vessel wall, (b) DAGMF only: vessel wall, (c) DAGMF + shape: lumen, (d) DAGMF only: vessel lumen, and (g) DSC color scale.

Fig. 7 Synthetic valve annulus segmentation. The label K indicates the background (in cyan), TW and BW indicates the top and bottom walls, respectively (in magenta), TB and BB indicate the top and bottom blood pools, respectively (in yellow), and V indicates the valve annulus (in green). (a) DAG used for valve annulus segmentation, (b) original noisy image, and (c) segmentation results.

Fig. 8 Mitral valve labeling using transesophageal ultrasound images. The model [Fig. 7(a)] includes label K indicates the background (in cyan), TW and BW indicates the top and bottom walls, respectively (in magenta), TB and BB indicate the top and bottom blood pools, respectively (in yellow), and V indicates the valve annulus. (a) Original ultrasound image and (b) segmentation.

Fig. 9 Atrial wall segmentation DAGMF augmented with shape complexes with α representing the regularization strength. (* a simple star convexity constraint is applied to this label.)

useful in these scenarios in which numeric parameterization may be opaque and unintuitive for clinical users. The mitigation of variations in accuracy may also be helpful in further validation and comparison with other algorithms, as less effort is required to generate an acceptable performance.

Because of the emphasis on maintaining a single continuous-space, there is no additional computational expense or ambiguity due to coordinate system warping, permitting the segmentation of multiple objects simultaneously and well as branching objects, both of which complicate coordinate system warping approaches. This is in stark contrast to prior approaches²¹ in which such warping is necessary, preventing specific types of shape complexes to be solved due to coordinate system ambiguity. Generally speaking, these coordinate system difficulties are avoided by discrete approaches,^{12,13} but by maintaining a continuous-space image domain, issues of metrication are avoided entirely, especially those resulting from the placement of infinite-cost edges associated with discrete domain star-convexity constraints.

Fig. 10 Best and worst case atrial wall segmentation results. The atrial blood pool is shown in magenta and the atrial wall in cyan. The black regions are user-provided seed points for the atrial blood label. (a) Best case original contrast-enhanced CT image (cropped to cardiac region), (b) worst case original contrast-enhanced CT image (cropped to cardiac region), (c) best case manual segmentation, (d) worst case manual segmentation, (e) best case shape complex result, and (f) worst case shape complex result.

Table 1 Mean distance error results for the blood pool and whole atrium labels. These are reflective of the errors seen on the inner and outer boundary of the atrial wall label.

($n = 10$)	Mean distance error (Operator 1)	Mean distance error (Operator 2)
Blood pool (inner wall)	0.76 ± 0.51 mm	0.59 ± 0.36 mm
Whole atrium (outer wall)	1.51 ± 0.55 mm	1.27 ± 0.29 mm

These priors, however, require some form of intelligent initialization to infer the geodesic direction field, $e_L(x)$. In both the synthetic and medical image segmentation experiments, this information was provided by the user through picking the centroid (as in Secs. 3.2 and 3.1) or by seeding the region of interest (as in Sec. 3.4). Although these methods may be suitable for hierarchical continuous max-flow-based interactive segmentation,²⁸ different approaches will be required for fully automated segmentation pipelines. Currently, the use of manual region-of-interest seeding is likely a large cause of variability in complex problems such as atrial wall segmentation.

There are four immediate areas of future work for shape complexes:

- Incorporation of multiple star-convexity constraints into a single label, increasing the number of shape options for each label. For example, constraining a shape to be circular about a particular vantage point could be implemented with three geodesics, one pointing toward the center of the circle and two tangent to it, pointing in the clockwise and counter-clockwise directions, respectively. In theory, a similar combination of geodesics could be used for arbitrary shapes provided that the boundary is parametrically defined or in interactive scenarios, where edge information can be extracted and scale invariance is assumed.¹⁸
- Performance improvements including GPGPU acceleration and incorporation of C++ implementations into the ASETS library.²⁶ As stated in Sec. 2.3, shape complexes can be implemented in an inherently parallelizable manner suitable for GPGPU acceleration. Incorporation into C++ would more readily allow its integration into open source libraries for medical imaging processing and visualization, such as ITK and VTK.
- Incorporation into existing continuous max-flow based interactive segmentation interfaces²⁸ that use input sampling mechanisms amenable to defining star convexity constraints. Alternatively, vantage point placement could be implemented as another secondary mechanism for user interaction.
- Investigation of automated approaches for defining geodesics based on atlas registration suitable for fully automatic segmentation pipelines. By processing deformation field resulting from the registration, it may be possible to automatically derive geodesic, rather than simple, star convexity constraints.

While these initial results are promising, and the additional robustness to parameter selection widely desirable in an array of medical image segmentation tasks, more validation is required to establish clinical utility in a particular medical domain. For example, in order to be clinically applicable for atrial wall segmentation, shape complexes must be augmented with a cost estimation framework that is robust to CT artifacts, such as those generated by the presence of pacemakers, as well as variable CNRs due to the variable dose rates of cardiac CT across clinical centers and scanners.

5 Conclusions

This paper presents “shape complexes,” which allow label ordering and geometric considerations to be used in tandem for segmenting multiple objects with complex shape requirements. This novel general-purpose segmentation approach, which augments previous extendable max-flow formulations, allows lower regularization values to be employed, preserves the segmented object’s structure with greater fidelity, and is generally more robust to differing parameter values.

Disclosures

There are no competing interests.

Acknowledgments

The authors would like to acknowledge Dr. Eranga Ukwatta for providing vessel ultrasound data. Funding for John S.H. Baxter’s research was provided by the Natural Sciences and Engineering Research Council of Canada.

References

1. T. F. Cootes et al., “Use of active shape models for locating structures in medical images,” *Image Vision Comput.* **12**(6), 355–365 (1994).
2. T. F. Cootes et al., “Active shape models—their training and application,” *Comput. Vision Image Underst.* **61**(1), 38–59 (1995).
3. T. Heimann and H.-P. Meinzer, “Statistical shape models for 3D medical image segmentation: a review,” *Med. Image Anal.* **13**(4), 543–563 (2009).
4. T. Okada et al., “Automated segmentation of the liver from 3D CT images using probabilistic atlas and multilevel statistical shape model,” *Acad. Radiol.* **15**(11), 1390–1403 (2008).
5. A. Tsai et al., “A shape-based approach to the segmentation of medical imagery using level sets,” *IEEE Trans. Med. Imaging* **22**(2), 137–154 (2003).
6. X. Artaechevarria, A. Munoz-Barrutia, and C. Ortiz-de Solórzano, “Combination strategies in multi-atlas image segmentation: application to brain MR data,” *IEEE Trans. Med. Imaging* **28**(8), 1266–1277 (2009).
7. T. Rohlfing and C. R. Maurer, “Shape-based averaging,” *IEEE Trans. Image Process.* **16**(1), 153–161 (2007).
8. Y. Boykov, O. Veksler, and R. Zabih, “Fast approximate energy minimization via graph cuts,” *IEEE Trans. Pattern Anal. Mach. Intell.* **23**(11), 1222–1239 (2001).
9. Y. Boykov and V. Kolmogorov, “An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision,” *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(9), 1124–1137 (2004).
10. D. Freedman and T. Zhang, “Interactive graph cut based segmentation with shape priors,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Vol. 1, pp. 755–762, IEEE (2005).
11. J. Malcolm, Y. Rathi, and A. Tannenbaum, “Graph cut segmentation with nonlinear shape priors,” in *IEEE Int. Conf. on Image Processing*, Vol. 4, pp. IV–365, IEEE (2007).
12. O. Veksler, “Star shape prior for graph-cut image segmentation,” in *European Conf. on Computer Vision (ECCV)*, pp. 454–467, Springer (2008).
13. V. Gulshan et al., “Geodesic star convexity for interactive image segmentation,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR ‘10)*, pp. 3129–3136, IEEE (2010).
14. J. Egger et al., “Square-cut: a segmentation algorithm on the basis of a rectangle shape,” *PLoS One* **7**(2), e31064 (2012).
15. J. Egger et al., “Template-cut: a pattern-based segmentation paradigm,” *Nat. Sci. Rep.* **2**, 420 (2012).
16. J. Egger et al., “Nugget-cut: a segmentation scheme for spherically- and elliptically-shaped 3D objects,” in *Joint Pattern Recognition Symp.*, pp. 373–382, Springer (2010).
17. R. Schwarzenberg et al., “Cube-cut: vertebral body segmentation in MRI-data through cubic-shaped divergences,” *PLoS One* **9**(4), e93389 (2014).
18. J. Egger et al., “Interactive-cut: real-time feedback segmentation for translational research,” *Comput. Med. Imaging Graphics* **38**(4), 285–295 (2014).
19. J. Yuan, E. Bae, and X.-C. Tai, “A study on continuous max-flow and min-cut approaches,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR ‘10)*, pp. 2217–2224, IEEE (2010).
20. J. Yuan et al., “An efficient convex optimization approach to 3D prostate MRI segmentation with generic star shape prior,” in *Prostate MR Image Segmentation Challenge, Conf. on Medical Image Computing and Computer-Assisted Interventions (MICCAI)*, Vol. **7512** (2012).
21. E. Ukwatta et al., “3-D carotid multi-region MRI segmentation by globally optimal evolution of coupled surfaces,” *IEEE Trans. Med. Imaging* **32**(4), 770–785 (2013).
22. J. S. H. Baxter et al., “A continuous max-flow approach to multi-labeling problems under arbitrary region regularization,” arXiv preprint arXiv:1405.0892 (2014).

23. J. S. H. Baxter et al., "A continuous max-flow approach to general hierarchical multi-labeling problems," arXiv preprint arXiv:1404.0336 (2014).
24. M. Rajchl et al., "Hierarchical max-flow segmentation framework for multi-atlas segmentation with kohonen self-organizing map based Gaussian mixture modeling," *Med. Image Anal.* **27**, 45–56 (2016).
25. J. S. H. Baxter, J. Yuan, and T. M. Peters, "Shape complexes in continuous max-flow hierarchical multi-labeling problems," arXiv preprint arXiv:1510.04706 (2015).
26. <http://www.advancedsegmentationtools.org/>
27. J. S. Baxter et al., "Smoothness parameter tuning for generalized hierarchical continuous max-flow segmentation," *Proc. SPIE* **9034**, 903410 (2014).
28. J. S. H. Baxter et al., "Optimization-based interactive segmentation interface for multi-region problems," *Proc. SPIE* **9413**, 94133T (2015).

John S.H. Baxter is a PhD candidate in the Imaging Research Laboratories at the Robarts Research Institute, London, Canada. He completed his bachelor's degree in software engineering at the University of Waterloo, Waterloo, Canada, in 2012. His research focuses on the use of variational optimization in defining and solving

medical image processing tasks, such as segmentation and image enhancement.

Jiro Inoue is a postdoctoral fellow in the Imaging Research Laboratories at the Robarts Research Institute, London, Canada. His research interests are translational computing in medicine, and improved management of atrial fibrillation through the use of image-based metrics.

Maria Drangova is a scientist with the Imaging Research Laboratories, Robarts Research Institute and a professor in the Departments of Medical Biophysics, University of Western Ontario. Her research interests include developing novel cardiac imaging techniques for diagnosis and the guidance of interventional procedures.

Terry M. Peters is a scientist in the Imaging Research Laboratories at the Robarts Research Institute, London, Canada, and a professor in the Departments of Medical Imaging and Medical Biophysics at Western University London, Canada, as well as a member of the Graduate Programs in Neurosciences and Biomedical Engineering. He directs a research laboratory with a focus on research and development in the field of image-guided surgery and therapy.