Journal of Medical Imaging

MedicalImaging.SPIEDigitalLibrary.org

Optimization-based interactive segmentation interface for multiregion problems

John S. H. Baxter
Martin Rajchl
Terry M. Peters
Elvis C. S. Chen

SPIE.

John S. H. Baxter, Martin Rajchl, Terry M. Peters, Elvis C. S. Chen, "Optimization-based interactive segmentation interface for multiregion problems," *J. Med. Imag.* **3**(2), 024003 (2016), doi: 10.1117/1.JMI.3.2.024003.

Optimization-based interactive segmentation interface for multiregion problems

John S. H. Baxter,[a,b,*] Martin Rajchl,[a,c] Terry M. Peters,[a,b] and Elvis C. S. Chen[a]
[a]Western University, Robarts Research Institute, 1151 Richmond Street N., London, Ontario N6A 5B7, Canada
[b]Western University, Biomedical Engineering Graduate Program, 1151 Richmond Street N., London, Ontario N6A 5B7, Canada
[c]Imperial College London, Department of Computing, Exhibition Road, London SW7 2AZ, United Kingdom

Abstract. Interactive segmentation is becoming of increasing interest to the medical imaging community in that it combines the positive aspects of both manual and automated segmentation. However, general-purpose tools have been lacking in terms of segmenting multiple regions simultaneously with a high degree of coupling between groups of labels. Hierarchical max-flow segmentation has taken advantage of this coupling for individual applications, but until recently, these algorithms were constrained to a particular hierarchy and could not be considered general-purpose. In a generalized form, the hierarchy for any given segmentation problem is specified in run-time, allowing different hierarchies to be quickly explored. We present an interactive segmentation interface, which uses generalized hierarchical max-flow for optimization-based multiregion segmentation guided by user-defined seeds. Applications in cardiac and neonatal brain segmentation are given as example applications of its generality. [DOI: 10.1117/1.JMI.3.2.024003]

Keywords: hierarchical max-flow segmentation; interactive segmentation; optimization-based segmentation; convex optimization; ASETS library.

Paper 16025PR received Feb. 10, 2016; accepted for publication May 26, 2016; published online Jun. 14, 2016.

1 Introduction

Interactive segmentation is the middle-ground between fully manual segmentation, where a user manually contours slices of a three-dimensional medical image to define objects of interest, and automated segmentation, where, with minimal user input, an algorithm attempts the segmentation task with no guidance or interaction with the user. The former is widely known to be time-consuming and subject to inadequacies regarding the number of slices segmented and the consistency between segmentations. The latter is often very rigid, being specific to a particular anatomy of interest in a designated modality under specific conditions and difficult to incorporate anatomical knowledge into, especially in the presence of pathology. In interactive segmentation, the user and algorithm work together, with the user providing initial input and corrections while the algorithm ensures the consistency of the segmentation across slices.[1] Interactive segmentation has long been known to improve segmentation time and consistency for tasks in which manual segmentation would otherwise be necessary.[2]

From a purely input-output point of view, interactive segmentation programs differ only in terms of the mechanisms in which the user can provide information and the algorithms used to process said information. More specifically, interactive segmentation programs differ in terms of

- the number of labels allowed by the interface and their topology (label orderings),
- the sampling mechanisms available to the user, such as paint-brushes, contours, and so on,
- the algorithms that process these sampled data to derive a labeling, and
- the organization of the multiple processing components used, which together form the segmentation pipeline.

These variables are inter-related, with the number of labels and types of sampling mechanisms constraining the types of algorithms available, and the types of algorithms constrain what types of algorithm organizations are meaningful. These factors dictate the scope of segmentation problems that can be readily addressed.

The number of labels allowed by the interactive segmentation interface is arguably the simplest method for categorizing interactive segmentation interfaces. Early methods in interactive segmentation, such as Interactive graph-cuts,[3] Grab-cut,[4] and Intelligent Scissors,[5] were constrained to the use of only two labels: foreground and background. TurtleSeg[6,7] and ITKSnap[8] permit the use of an arbitrary number of labels, making them better suited for multiregion problems. As of yet, there are no prior interactive segmentations that consider label orderings as a form of input. However, several take advantage of a particular label ordering suited to a particular segmentation problem.[9,10]

In terms of sampling image data to build a descriptive data model, Interactive graph-cuts[3] used a paint-brush mechanism, Intelligent Scissors[5] and TurtleSeg[6,7] a contouring mechanism, and ITKSnap[8] provides mechanisms for both. Some interactive segmentation frameworks, such as the MIDAS framework,[11] provide limited direct user manipulation of labels, shifting its focus to the user definition of pipelines containing fundamental segmentation algorithms, such as thresholding and region growing, and morphological operators with an emphasis on segmentation reproducibility.

Interactive segmentation interfaces display a considerable amount of variability in terms of the algorithms available to extrapolate the user's sampling information and other indications. Several methods utilize optimization-based approaches

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

2329-4302/2016/$25.00 © 2016 SPIE

ranging from shortest path algorithms[5] to discrete graph cuts,[3,4] level-sets,[8] and random walk based segmentation.[6,7,12] These algorithms generally have a fairly rigid organization, with the exception of interfaces in which pipelines are the primary focus of user interaction.[11]

As stated earlier, one issue with general-purpose interactive segmentation programs is their overall lack of explicit incorporation of anatomical knowledge in an intuitive manner. Initially, the concept of incorporating anatomical knowledge into a general-purpose (and therefore application/modality agnostic) program may seem paradoxical. However, certain abstract forms of anatomical knowledge about the spatial arrangement between objects may be expressed quickly and easily while maintaining generality across applications. The application of hierarchies to segmentation and natural scene understanding has been well-studied,[13,14] but often treats the hierarchy as a structure that the algorithm must learn aside from user input.

In this work, we allow the user to explicitly define a segmentation hierarchy, which can be optimized globally using generalized hierarchical max-flow (HMF).[15,16] These hierarchies express object grouping behavior by way of partitioning. For example, one can think of a super-object, such as the heart, as being composed of several subobjects (the left and right ventricles and atria), which can be recursively subdivided (blood pool and wall). These partitioning relationships allow for nuanced regularization requirements to be described.

2 Methods

2.1 *Hierarchical Max-Flow Segmentation*

HMF models[15,16] extend the notion of orderings from the Ishikawa model[17,18] to hierarchies. In this case, collections of labels can be joined to create a super-label. This process repeats itself until the entire image is represented as a single label, denoted as S or the root label. (Such S labels are shown at the top of the hierarchies in Figs. 5 and 6.) Alternatively, one can take a top-down interpretation, recursively splitting objects in the image into their constituent parts. The formula for these models is

$$\min_{u} \sum_{\forall L \in \mathbb{L}} \int_{\Omega} D_L(x) u_L(x) \mathrm{d}x + \sum_{\forall L \in \mathbb{G}} \int_{\Omega} S_L(x) |\nabla u_L(x)| \mathrm{d}x$$
$$\text{s.t.} \forall L \in \mathbb{L}, [u_L(x) \geq 0], \sum_{\forall L \in \mathbb{L}} u_L(x) = 1$$
$$\forall L \in \mathbb{G}, \left[\sum_{L' \in L.C} u_{L'}(x) = u_L(x) \right]. \quad (1)$$

The operator $.C$ refers to parent/child relationships in the hierarchy, specifically, $L.C$ returns the set of child labels of label L. This formula is similar to that of the continuous Potts model[19] in that it contains a series of unary data terms, $D_L(x)$, and a set of regularization or weighting terms, $S_L(x)$, on the gradient magnitude of each labeling function, $|\nabla u_L(x)|$, and that these terms are summed over each label (and super-label). These hierarchical models are strictly more expressive than both Potts[20] and Ishikawa[17] models together,[15] allowing for a wider array of segmentation problems to be addressed. However, hierarchies are more difficult to specify.

This tree structure has previously been considered a hard-coded part of the image segmentation algorithm, encouraging the use of Potts[20] or Ishikawa[17] models and their continuous max-flow counterparts[18,19] to handle general-purpose segmentation. However, this poses fundamental limitations on what can be segmented. For example, in the Potts model, only a single smoothness parameter is assigned, which makes it difficult to simultaneously segment smooth structures alongside irregular ones. Ishikawa models allow for more parameterization but require the objects being segmented to satisfy a full ordering, which is not the case for complex anatomy.

The general HMF solver alleviates this problem by permitting any arbitrary hierarchy to be defined, allowing for more anatomical knowledge to be encoded. This intuitive form of anatomic knowledge can be readily incorporated into the optimization-based segmentation of multiple regions. Problems regarding constructing the largest meaningful hierarchy given label grouping information are NP-hard (see the Appendix), meaning that interactive methods, at least for hierarchy definition, may be required so as to make use of a user's anatomical knowledge.

Details of the precise implementation of the HMF solver can be found in the technical report.[15] The solver is provided open-source at Ref. 21 in both MATLAB® and C++ implementations.

2.2 *Definition of Cost Terms*

A crucial decision in optimization-based segmentation is the structure of the cost terms. Log-likelihood data terms, derived from Bayes' theorem, have been effective in interactive[3] and multiregion segmentation,[22] taking the form

$$D_L(x) = \begin{cases} \infty & \text{if } x \text{ is a seed for a label other than } L \text{ or element of } L.P^* \\ -\ln\{P[I(x)|x\epsilon L]\} & \text{else} \end{cases}, \quad (2)$$

where $P[I(x)|x\epsilon L]$ is the likelihood of a voxel in label L having the same intensity as x, $I(x)$ and $L.P^*$ is the set of ancestors (parents, grandparents, and so on) of label L. The probability, $P[I(x)|x\epsilon L]$, is estimated from the histogram of the seeded voxels, which approximates the true value when a large number of seeds is used. The infinite cost ensures that any voxels used to seed a particular object remain a part of said object in the segmentation and that any voxel that has been seeded as a particular label can only be assigned to said label if it is an end-label or to its child labels otherwise.

Smoothness or regularization terms are non-negative costs used to both smooth the labeling and to align edges in the segmentation with those visible in the underlying image. The smoothness terms used were

$$S_L(x) = \alpha_L * \exp[-\beta_L |\nabla I(x)|] + \gamma_L, \quad (3)$$

where the parameters α_L, β_L, and γ_L are specified by the user. The exponential term implies that lower costs are associated with label boundaries, which occur when there is a high gradient magnitude, encouraging nearby edges in the segmentation to migrate to said areas similar to the contrast-sensitive smoothness terms used by Boykov and Jolly.[3]

2.3 *Plane Selection*

To improve efficiency and accuracy while encouraging interactivity, plane selection can be used.[6] Such planes can be used by the algorithm to inform the user as to which areas of the segmentation would benefit the most from user interaction. Top et al.[6] introduced a notion of active learning in which the segmentation algorithm identifies areas of maximum uncertainty, the uncertainty of a segmentation expressed as

$$U(x,y) = \lambda_E U_E(x,y) + \lambda_B U_B(x,y) + \lambda_R U_R(x,y) + \lambda_S U_S(x,y), \quad (4)$$

where U_E is the entropy of the segmentation results, U_B is the uncertainty associated with boundaries in the segmentation, U_R is the uncertainty associated with the regional intensity, and U_S is the uncertainty associated with the tortuosity of the boundary around x. The λ's are constants with the majority (80%) of the weight given to λ_E.[6] Note that the U_R and U_B terms are explicitly handled by the segmentation algorithm itself by the definition of the cost functions. We assign all the weight to the U_E term and use only maximum axis-aligned planes. This ensures that the plane selection algorithm quickly produces planes in orientations to which the user is accustomed. The segmentation used in plane selection is the previous segmentation generated by the user. Thus, plane selection is only defined after the first segmentation is computed and remains available for all subsequent interactions.

3 Interface Description

The interface is implemented using Kitware's Visualization Tool-Kit (VTK) for image processing and visualization and the Qt framework for graphical user interface support. The generalized HMF solver was encapsulated into a VTK algorithm object and accelerated using NVIDIA's Compute Unified Device Architecture. The sampling mechanism is brush-based similar to that used by Boykov and Jolly[3] and ITK-Snap[8] allowing for large portions of an object's interior to be covered with relative ease while not requiring strong boundary contrast.

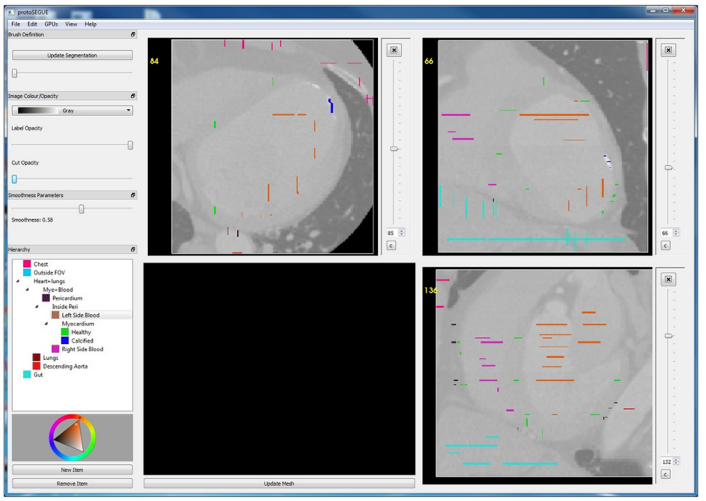

(a) Interface with seeding

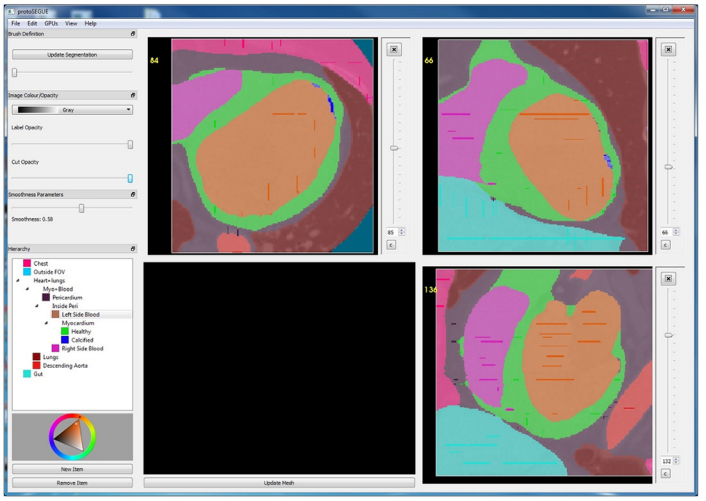

(b) Interface after labeling is performed with provided seeds

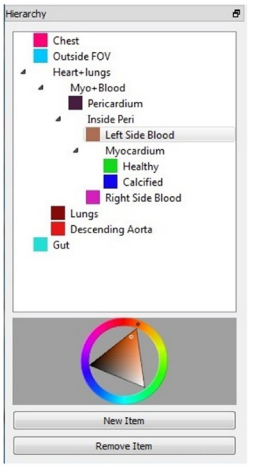

(c) Widget for hierarchy definition

Fig. 1 Segmentation interface with user seeds (a) before segmentation and (b) after segmentation. The hierarchy definition widget [bottom left corner of (a) and (b)] is shown enlarged in (c).

Table 1 Cardiac segmentation numerical results.

($n = 3$)	CT	MRA	TEE
Blood AVD (%)	6.6 ± 6.6	6.2 ± 3.6	14.2 ± 6.2
Myocardium AVD (%)	12.5 ± 11.3	16.7 ± 11.5	7.3 ± 4.5
Blood rMSE (mm)	1.14 ± 0.64	0.70 ± 0.21	1.080.27
Myocardium rMSE (mm)	1.31 ± 0.24	0.71 ± 0.24	1.48 ± 0.57
Blood DSC (%)	91.7 ± 2.6	94.3 ± 1.9	90.5 ± 4.3
Myocardium DSC (%)	83.8 ± 3.9	82.1 ± 3.7	91.8 ± 2.7
Weighted DSC (%)	87.5 ± 2.0	89.8 ± 2.7	91.2 ± 3.2
Interoperator variability Weighted DSC (%)	92.7 ± 4.9	93.6 ± 2.5	92.0 ± 2.1
Weighted DSC from Ref. 9 (%)	87.7 ± 3.7	89.3 ± 2.7	85.7 ± 2.0

The user can place seeds for any label or super-label using the brush, creating the data model described in Eq. (2).

The interface is shown in Fig. 1. Hierarchies are defined in a side bar as shown in Fig. 1(c), which also acts as a widget for selecting the active label or super-label of the brush. This widget also allows the hierarchy to be restructured quickly, operating in a drag-and-drop manner. Last, the user can save the hierarchy along with smoothness term parameters and the initial user-defined samples for later use.

4 Example Applications of Interactive Segmentation

4.1 *Cardiac Segmentation*

Because of the generality of the algorithm and the interface, several existing continuous max-flow based methods, such as those developed in Refs. 9 and 10, can be easily replicated. We reproduced the experiments performed in Ref. 9, which included three cardiac volumes from computed tomography (CT),

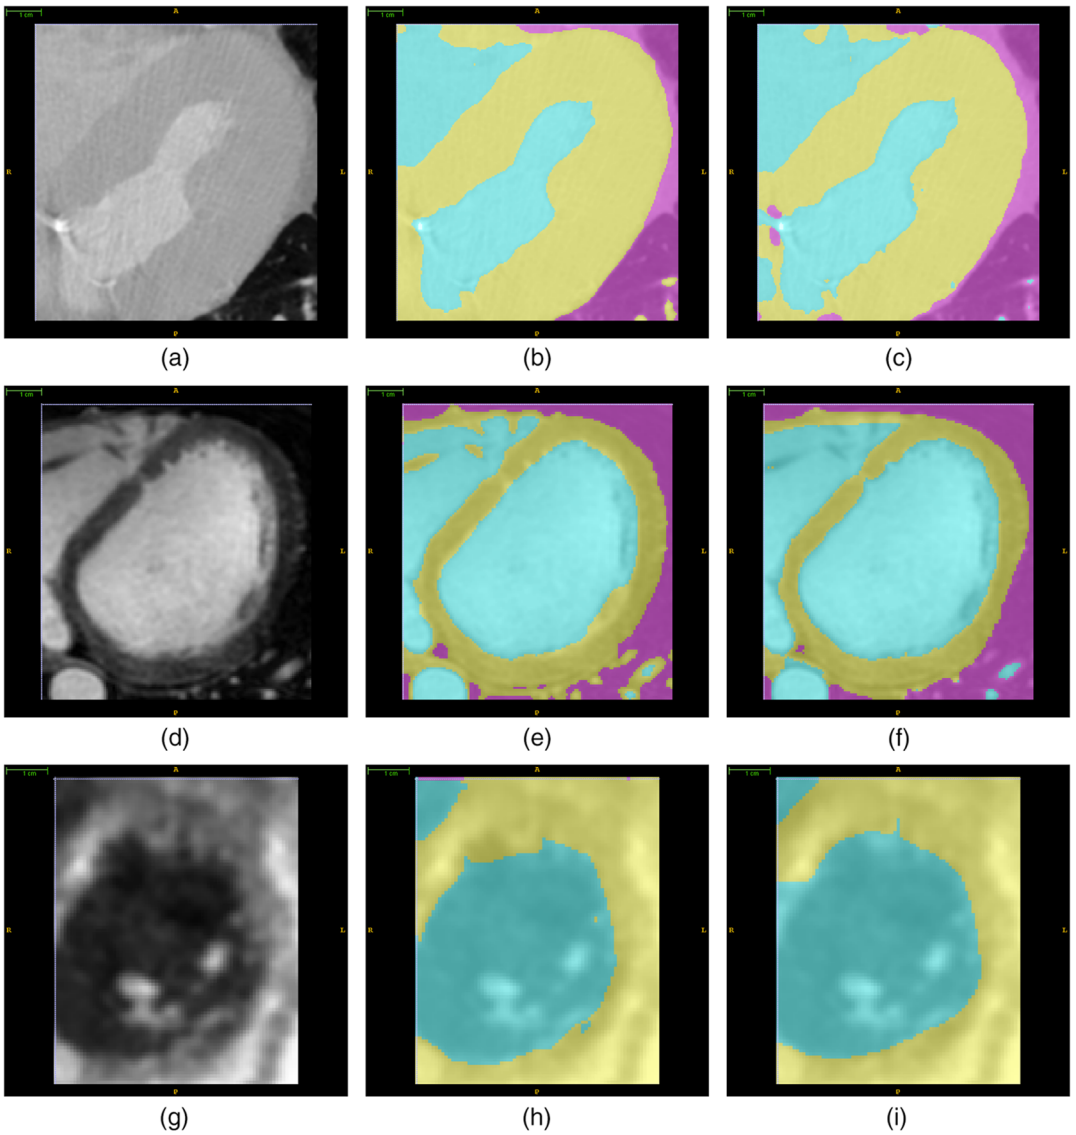

Fig. 2 Cardiac segmentation with underlying (a) CT, (d) MRA, and (g) TEE. Manual segmentations are in (b), (e), and (h), respectively, and interactive segmentation results in (c), (f), and (i).

magnetic resonance angiography (MRA), and trans-esophageal echocardiography (TEE).

Representative segmentation results are provided in Fig. 2. Numerical results in terms of average volume difference (AVD), root mean squared distance error (rMSE), and Dice similarity coefficient (DSC) are recorded in Table 1. These results are very consistent with interoperator variability above 90% and comparable with those presented by Rajchl et al.,[9] illustrating that our general-purpose segmentation interface can perform similarly to the one designed specifically for cardiac segmentation. Interestingly, the results for TEE indicate that the proposed interface outperforms the previous interactive segmentation interface.[9]

4.2 *Neonatal Cranial Magnetic Resonance Imaging Segmentation*

Neonatal brain images display some unique challenges for automated segmentation in that there are relatively few compared to adult brain images, making machine learning–based or atlas-based segmentation approaches infeasible. Additionally, bleeds in the ventricular system further complicated segmentation. In this context, interactive interfaces can be extremely useful since manual segmentation or correction is largely unavoidable. Figure 3 displays the visual results of neonatal ventricle segmentation using this interface.

To demonstrate the interactive segmentation interface's robustness to pathology, the previous experiment was extended to a neonatal MR image in which a severe ventricular bleed changes the intensity distribution of the ventricle to an extreme degree. The segmentation results are given in Fig. 4. Note that the hyperintense ventricular bleed is closer in intensity to white and gray matter than to the ventricles, and its appearance on the boundary of the ventricles would likely cause severe registration artifacts. The segmentation of the ventricle was achieved by partitioning it into two components: a healthy cerebral spinal fluid (CSF) component and the ventricular bleeding (Bl). The union of these components could then be regularized similar to the ventricle (Ve) in Fig. 3. The Ve label (the union of the CSF and Bl labels) for the pathological case is given in Fig. 4 In the hierarchies used in this segmentation problem, which are given in Fig. 5, the remaining labels are K, which refers to the background, He to the head, and Br to the brain.

5 Automatic Hierarchy Refinement

Although determining an appropriate hierarchy merely from grouping information is a computationally difficult problem, due to the mathematical formulation, it is possible to automatically refine a user-provided hierarchy for improved computational efficiency without compromising segmentation quality. This involves the contraction and removal of vertices in the hierarchy with zero regularization or where zero regularization can be induced without changing the optimization functional. One specific example is that when the source node has only two children, whereby one can be contracted by transferring its smoothness value to the other. To demonstrate this, we performed automatic hierarchy optimization on the method presented by Rajchl et al.[10] using late gadolinium enhanced magnetic resonance imaging (LGE-MRI). This segmentation problem involved partitioning the image into thoracic background (T) and cardiac (C) labels, the latter being subdivided into blood (Bl), healthy myocardium (M), and scar tissue (Sc) as shown in Fig. 6(a). The cardiac label, C, was automatically contracted, resulting in Fig. 6(b) and an estimated 20% improvement in speed.

As with the previous experiments, the results (recorded in Table 2) were comparable to those presented by Rajchl et al.[10] and was within the range of interoperator variability. However,

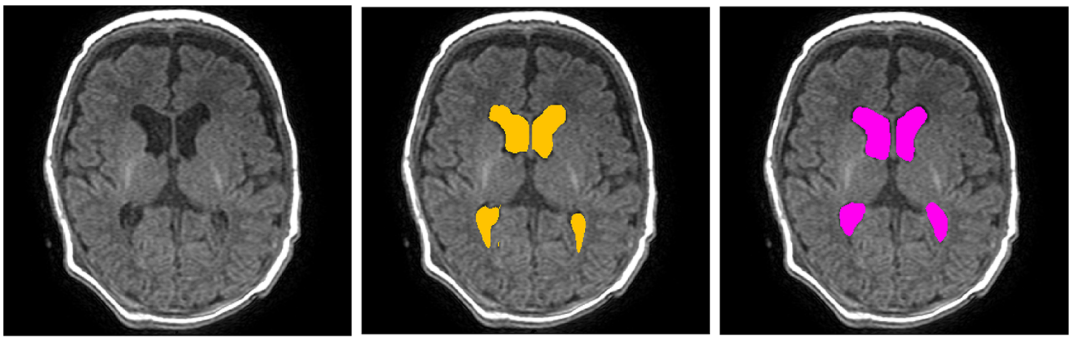

(a) Original MR Image (b) Manual Segmentation (c) Interactive Segmentation

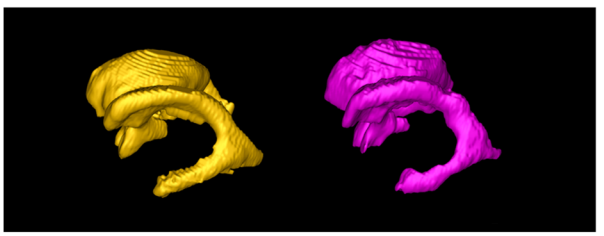

(d) Surface renderings

Fig. 3 Neonatal ventricle segmentation with (a) MR, (b) manual segmentation, and (c) interactive segmentation results. (d) Surface renderings of both the fully manual (left) and interactive (right) segmentation results.

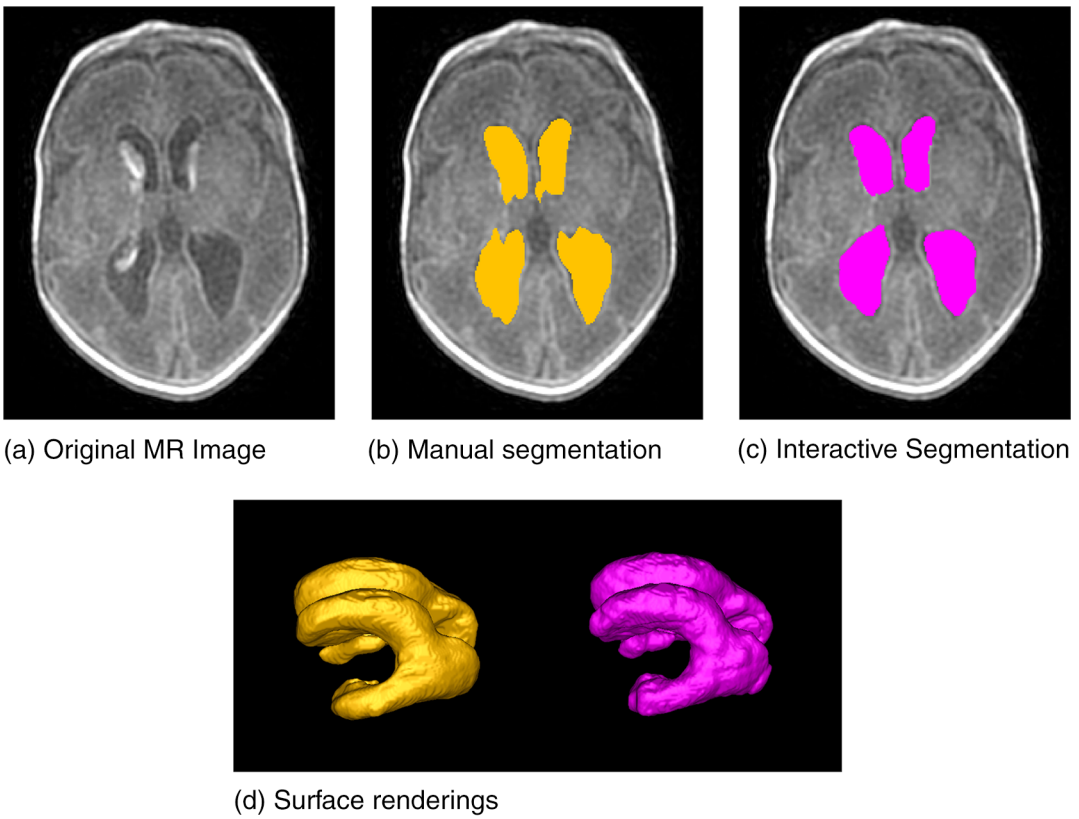

Fig. 4 Pathological neonatal ventricle segmentation with (a) MR, (b) manual segmentation, and (c) interactive segmentation results. (d) Surface renderings of both the fully manual (left) and interactive (right) segmentation results.

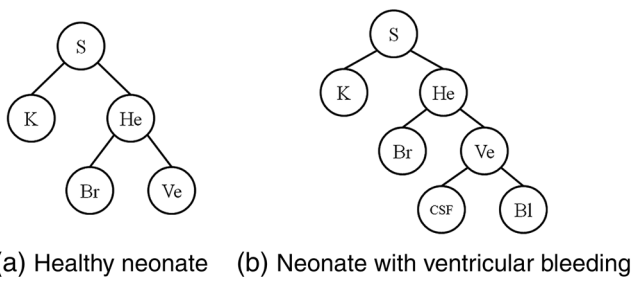

Fig. 5 Hierarchies used in (a) healthy and (b) pathological neonatal ventricle segmentation.

this level of accuracy was achieved without postprocessing steps, such as connected components analysis or other modifications that would make the interface specific to cardiac or LGE-MRI segmentation.

6 Discussion

Improvements in interactive segmentation interfaces can have a distinct impact in clinical contexts in which automated segmentation is not feasible. Several clinical applications require manual segmentation due to pathology such as tumors in radio-oncological applications or bleeds in neonatal cranial imaging. These applications require a user to manually delineate some anatomy in order to perform relevant measurements such as tumor volume. In these applications, accurate segmentation may be necessary for robust, correct measurements, and the use of interactive segmentation can have a distinct benefit, conserving user time while encouraging accurate results, which will in turn improve patient outcomes by improving the diagnostic capabilities of these measurements (compared to manual segmentation) in single acquisition and longitudinal studies.

The primary advantage of this interface over other interactive segmentation programs is that it allows the user to interactively specify both segmentation hierarchy and initial seeds. The former means that the interface is very general purpose, allowing for arbitrary regions to be defined, while incorporating anatomical knowledge in a direct manner. This gives it a distinct advantage over other interactive segmentation interfaces, which either limit the number or type of regions or do not allow the user to specify abstract anatomical knowledge. The latter takes advantage of a paint-brush mechanism, which allows for large regions of the interior of the object to be seeded with minimal user effort, thus improving the probabilistic data terms.

The second major advantage is that the algorithm is founded in optimization principles, ensuring robustness and repeatability across images. The formulation of the costs also allows for the regional and boundary uncertainty (U_R and U_B) identified in Ref. 6 to be actively addressed by the segmentation process, making plane selection simpler and more efficient. Plane selection is further improved by selecting only axis-aligned planes in which the user is accustomed.

7 Future Work

There are several future directions in which to take this work aside from general improvements to computational resource usage and performance. Specifically,

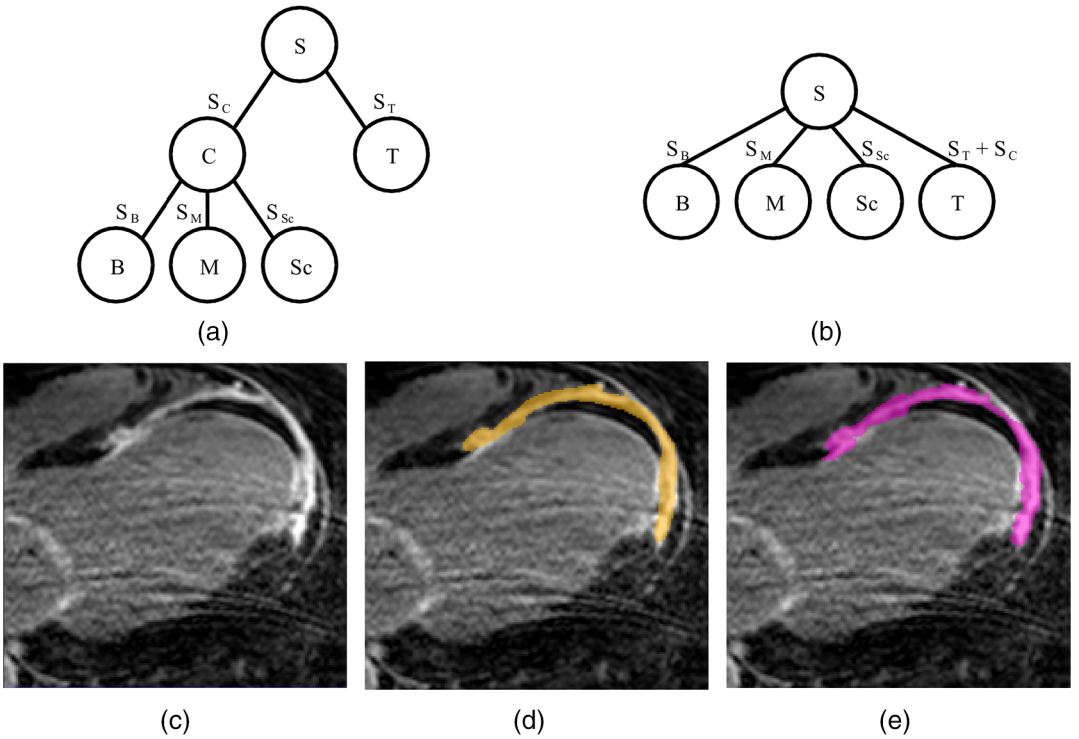

Fig. 6 Example of automatic hierarchy segmentation. (a) The original hierarchy reproduced from Ref. 10 and (b) the optimized version; (c) an LGE-MRI with (d) manual segmentation and (e) interactive segmentations results.

Table 2 Scar tissue segmentation results.

($n = 10$)	Accuracy
Scar AVD (%)	26.9 ± 15.6
Scar rMSE (mm)	1.30 ± 0.32
Scar DSC (%)	74.1 ± 3.5
Scar DSC from Ref. 10 (%)	76.0 ± 3.0
Interoperator variability from Ref. 10—scar DSC (%)	76.2 ± 2.6
Intraoperator variability from Ref. 10—scar DSC (%)	75.2 ± 2.8

- incorporation of a more extensive model of label organization,
- incorporation of geometric or shape constraints,
- improvements to the definition of the smoothness model, and
- improvements to the plane selection mechanism.

Recently, work has been performed that extends the possibility of label organization in continuous max-flow from hierarchical models[15] to models that allow for any possible label ordering.[23] However, there remain issues in terms of how these structures can be specified by a user in run-time in an intuitive manner as they are defined using a constrained set of rooted, weighted directed acyclic graphs, which do not have a user-friendly tool already in place.

There has also been increasing interest in the use of generic geometric or shape constraints such as star-shaped priors in both graph-cuts[24] and max-flow image segmentation.[25] Shape complexes have already been proposed, which combine the notions of label orderings and star-convex object constraints to develop complicated models of object geometry from the union and disjunction of star-convex objects.[26] Such frameworks can be readily incorporated into this interactive segmentation framework with minimal changes to the interface or usability, while contributing a significant improvement to the segmentation accuracy through the encoding of additional anatomical knowledge.

Currently, the interface allows the user to modify the parameters in the smoothness term, but does not permit any other manipulation. This could be incorporated through the addition of a contouring mechanism similar to that in Intelligent Scissors, TurtleSeg, and ITKSnap. These contours could supply specific information, which can improve the smoothness terms, as well as give the user complimentary ways to sample regions.

In terms of plane selection, future work could include defining a sequence of planes sensitive to the distance between them, rather than a single set. This would allow the algorithm to intelligently inform the user of multiple areas of uncertainty without reinvoking the continuous max-flow segmentation algorithm and allow the user to provide feedback on multiple high uncertainty planes in a single interaction cycle.

8 Conclusions

Interactive segmentation helps bridge the gap between manual and automatic segmentation, allowing each to address the weaknesses of the other. In this work, we present a general-purpose interactive segmentation interface and apply it to cardiac and

neonatal cranial segmentation with performance comparable to previously published methods specific to said applications.

This interface allows for the user to define a segmentation hierarchy in run-time, taking advantage of a fast, GPU-accelerated general HMF solver, which in turn allows for more knowledge of spatial relationships between anatomical regions to be encoded. This encourages the use of optimization techniques and interactive interfaces in which a user can quickly define and correct a segmentation, thereby increasing the speed, quality, and robustness of general segmentation tasks. The ability to modify the hierarchy in run-time allows for the interactive segmentation interface to account for extreme deviations, such as ventricular bleeds, by the addition of multiple labels to account for them. This interface is the first to allow the user to modify the abstract anatomic knowledge, i.e., label ordering, provided to the computer in run-time.

Appendix: NP-Hardness of Hierarchy Definition

Theorem: Consider $\mathbb{L}$ to be the set of objects in an image. Determining if there is a hierarchy with at least k elements from a specified set of group relationships ($\mathbb{G} \subseteq 2^{\mathbb{L}}$) is NP-complete, and specifying the largest hierarchy is NP-hard.

Proof. Any hierarchy is equivalent to an independent set in a particular, polynomial-time constructible graph. To prove this, we will show the construction of this graph and proceed through a proof by contradiction.

Let G be a graph in which each vertex represents a nonempty set of labels in the segmentation that are expected to have some regularization, that is, their union forms a meaningful structure or their grouping is meaningful. In this graph, edges represent conflicts, where the vertices refer to sets that are neither embedded (one is a subset of the other) nor disjoint. For the sake of notation, each vertex will be denoted via its corresponding element of $\mathbb{G}$, the grouping relationship it represents.

Assume there is a hierarchy where the nodes are selected from the vertices of G, but do not form an independent set. Consider the edge between two vertices that indicates a dependency $e = (g_1, g_2)$. The two adjacent vertices $g_1, g_2 \in \mathbb{G}$ refer to two sets of end-labels that are neither disjoint nor a subset of each other. (That is, both $g_1 \cap g_2$ and $g_1 \setminus g_2$ are nonempty.) Consider label A to be an end-label common to both sets. Note since each is a superset of $\{A\}$, they must correspond to ancestors in the hierarchy and both lie on the direct path from $\{A\}$ to the root of the hierarchy. This implies that one must be an ancestor of the other, which is a contradiction since neither is a superset of the other. Thus, any hierarchy must correspond to an independent set in G.

Without loss of generality, assume G is connected. Each independent set can be transformed into a hierarchy in polynomial time in a top-down manner. At each iteration, we want to grow the hierarchy by the vertices corresponding to the largest group of end-labels at the lowest tier possible. We do this by ordering the vertices in the independent set by the size of group they represent. Then, perform a breadth-first search through the current tree to find the lowest tier that is a superset of the node under consideration. We grow the hierarchy by adding the group under consideration to the identified part of the hierarchy. We repeat this for each node in the set, initializing the hierarchy as only the root node, equivalent to the full set of end-labels. Last, we augment the hierarchy with the end-labels to make it valid.

Since the maximum hierarchy and maximum independent set problems can be reduced to each other in polynomial time, determining the largest hierarchy must be NP-hard. □

Acknowledgments

The authors would like to acknowledge Dr. Sandrine de Ribaupierre and Jessica Kishimoto for her assistance with neonatal ventricular segmentation data. Funding for John S. H. Baxter's research was provided by the Natural Sciences and Engineering Research Council of Canada.

References

1. K. McGuinness and N. E. O'Connor, "A comparative evaluation of interactive segmentation algorithms," *Pattern Recognit.* **43**(2), 434–444 (2010).
2. A. X. Falcão et al., "User-steered image segmentation paradigms: live wire and live lane," *Graph. Models Image Process.* **60**(4), 233–260 (1998).
3. Y. Y. Boykov and M. P. Jolly, "Interactive graph cuts for optimal boundary & region segmentation of objects in ND images," in *Proc. Eighth IEEE Int. Conf. on Computer Vision*, Vol. **1**, pp. 105–112 (2001).
4. C. Rother, V. Kolmogorov, and A. Blake, "Grabcut: interactive foreground extraction using iterated graph cuts," *ACM Trans. Graph.* **23**, 309–314 (2004).
5. E. N. Mortensen and W. A. Barrett, "Interactive segmentation with intelligent scissors," *Graph. Models Image Process.* **60**(5), 349–384 (1998).
6. A. Top, G. Hamarneh, and R. Abugharbieh, "Active learning for interactive 3D image segmentation," *Lect. Notes Comput. Sci.* **6893**, 603–610 (2011).
7. A. Top, G. Hamarneh, and R. Abugharbieh, "Spotlight: automated confidence-based user guidance for increasing efficiency in interactive 3D image segmentation," in *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging*, M. Bjoern et al., Eds., pp. 204–213, Springer, Berlin, Germany (2011).
8. P. A. Yushkevich et al., "User-guided level set segmentation of anatomical structures with ITK-SNAP," *Insight J.* **1** (2005).
9. M. Rajchl et al., "Fast interactive multi-region cardiac segmentation with linearly ordered labels," in *9th IEEE Int. Symp. on Biomedical Imaging*, pp. 1409–1412 (2012).
10. M. Rajchl et al., "Interactive hierarchical max-flow segmentation of scar tissue from late-enhancement cardiac MR images," *IEEE Trans. Med. Imaging* **33**, 159–172 (2014).
11. P. A. Freeborough, N. C. Fox, and R. I. Kitney, "Interactive algorithms for the segmentation and quantitation of 3-D MRI brain scans," *Comput. Methods Programs Biomed.* **53**(1), 15–25 (1997).
12. L. Grady, "Random walks for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.* **28**(11), 1768–1783 (2006).
13. J.-M. Beaulieu and M. Goldberg, "Hierarchy in picture segmentation: a stepwise optimization approach," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(2), 150–163 (1989).
14. J. C. Tilton, "Analysis of hierarchically related image segmentations," in *IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data*, pp. 60–69 (2003).
15. J. S. H. Baxter et al., "A continuous max-flow approach to general hierarchical multi-labelling problems," 2014, pp. 1–19, http://www.insight-journal.org/browse/publication/15 (30 May 2016).
16. 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 2015).
17. H. Ishikawa, "Exact optimization for Markov random fields with convex priors," *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(10), 1333–1336 (2003).
18. E. Bae et al., "A fast continuous max-flow approach to non-convex multi-labeling problems," in *Efficient Algorithms for Global Optimization Methods in Computer Vision*, B. Andrés, P. Thomas, and T. Xue-Cheng, Eds., pp. 134–154, Springer, Berlin, Germany (2014).

19. J. Yuan et al., "A continuous max-flow approach to Potts model," *Lect. Notes Comput. Sci.* **6316**, 379–392 (2010).
20. R. B. Potts, "Some generalized order-disorder transformations," *Proc. Cambridge Philos. Soc.* **48**, 106–109 (1952).
21. http://www.advancedsegmentationtools.org/
22. F. van der Lijn et al., "Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts," *NeuroImage* **43**(4), 708–720 (2008).
23. 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).
24. O. Veksler, "Star shape prior for graph-cut image segmentation," in *10th European Conf. on Computer Vision*, pp. 454–467, Springer (2008).
25. J. Yuan et al., "An efficient convex optimization approach to 3D prostate MRI segmentation with generic star shape prior," 2012, https://www.researchgate.net/publication/232451017_An_Efficient_Convex_Optimization_Approach_to_3D_Prostate_MRI_Segmentation_with_Generic_Star_Shape_Prior (30 May 2016).
26. 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).

John S. H. Baxter is a PhD candidate in the Imaging Research Laboratories at the Robarts Research Institute (RRI), London, Canada. He completed his bachelor's 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 enhancement.

Martin Rajchl holds his PhD in biomedical engineering from the Robarts Research Institute, Western University, Canada. Currently, he is a research associate in computing at the Imperial College London, where he focus on the analysis of medical images in large databases. His research interests include machine learning and computer vision methods to address image analysis problems, to further extract useful information for clinical research.

Terry M. Peters is a scientist in the Imaging Research Laboratories at RRI, London, Canada, 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.

Elvis C. S. Chen obtained his PhD from the School of Computing, Queen's University, Kingston, Canada, in 2007. He works in the field of image-guided interventions, applying techniques from robotics, computer vision, and computer graphics to the field of surgery. His research interests include joint kinematics, ultrasound guided interventions, tool calibration and tracking, and vision-guided laparoscopy. Currently, he is a research associate at RRI with cross appointment (assistant professor) at Western University, London, Canada.