

Relaxed Image Foresting Transforms for Interactive Volume Image Segmentation

Filip Malmberg¹, Ingela Nyström¹, Andrew Mehnert², Craig Engstrom² and Ewert Bengtsson¹

¹ Centre for Image Analysis, Uppsala University, Sweden

² School of ITEE, The University of Queensland, Brisbane, Australia

ABSTRACT

The Image Foresting Transform (IFT) is a framework for image partitioning, commonly used for interactive segmentation. Given an image where a subset of the image elements (seed-points) have been assigned correct segmentation labels, the IFT completes the labeling by computing minimal cost paths from all image elements to the seed-points. Each image element is then given the same label as the closest seed-point. Here, we propose the *relaxed* IFT (RIFT). This modified version of the IFT features an additional parameter to control the smoothness of the segmentation boundary. The RIFT yields more intuitive segmentation results in the presence of noise and weak edges, while maintaining a low computational complexity. We show an application of the method to the refinement of manual segmentations of a thoracolumbar muscle in magnetic resonance images. The performed study shows that the refined segmentations are qualitatively similar to the manual segmentations, while intra-user variations are reduced by more than 50%.

Keywords: Seeded segmentation, Interactive segmentation, Minimum cost paths, Image Foresting Transform

1. INTRODUCTION

Image segmentation, the process of identifying and separating relevant objects and structures in an image, is a fundamental problem in image analysis. Accurate segmentation of objects of interest is often required before further processing and analysis can be performed. Despite years of active research, fully automatic segmentation of arbitrary images remains an unsolved problem.

Seeded segmentation methods attempt to solve the segmentation problem in the presence of prior knowledge in the form of a partial segmentation. Given an image where a small subset of the image elements (called *seed-points*) have been assigned correct segmentation labels (e.g., object or background), an automatic algorithm completes the labeling for all image elements. The seed-points may be provided either by some automatic pre-processing algorithm, or by a human user in an interactive setting. Many different algorithms for seeded segmentation have been proposed ranging from classical seeded region growing^{1,2} and the marker-based watershed transform,³ through to the more recent minimal graph cuts,⁴ random walks,⁵ and image foresting transform (IFT)⁶ approaches. Here, we focus on the IFT approach.

In the IFT, the image is interpreted as a weighted graph. Each image element corresponds to a node in the graph, and adjacent image elements are connected by graph edges. Segmentation is performed by assigning to each node the label of the closest seed-point, as determined by the minimum cost path from the node to the set of seed-points. The cost of a path is a function of local image image properties – such as intensity or gradient magnitude – along the path. By defining this cost in an appropriate way, popular image segmentation methods such as relative fuzzy-connectedness and watersheds can be implemented.⁷ The IFT has several favorable properties:

1. Multiple objects can be segmented simultaneously. The computation time is independent of the number of object categories.
2. Each segmented region is guaranteed to be connected to seed-points with the same label, i.e., there are no isolated regions of a particular label that contains no seed-points.

Send correspondence to: Filip Malmberg, E-mail: filip@cb.uu.se




Figure 1. (Left) A synthetic gradient-like image, with noise and weak edges. (Middle) Segmentation obtained from the IFT with two seed-points, shown here as circles. (Right) By applying 30 iterations of the proposed relaxation procedure a smoother, more intuitively correct segmentation boundary is obtained.

3. When applied to a completely homogeneous image, the IFT produces a neutral segmentation corresponding to the Voronoi regions (according to a path-based distance⁸) of the seed-points. Thus, the IFT is not subject to the “small-cut” problem associated with minimal graph-cuts.⁴

The IFT can be computed in linear (in the number of image elements) time using Dijkstra’s algorithm,⁹ slightly modified to allow multiple seed-points.¹⁰ In interactive segmentation applications, a user often adds or removes seed-points to refine an existing segmentation. Falcaõ et al.⁷ showed that seed-points can be added to, or removed from, an existing IFT solution, without re-computing the entire solution. This modified algorithm, called the *differential* IFT, significantly reduces the total time required for interactive segmentation.⁷

Numerous studies^{7,11–13} have shown that the IFT, and similar methods based on minimal cost paths, are capable of producing high quality segmentations in a wide range of contexts. However, in images with weak or missing boundaries the IFT tends to produce irregular segmentation boundaries. An explanation for this is that the IFT propagates information from the seed-points only along minimum cost paths. Since two adjacent image elements may receive their information from different seed-points, regularity of the segmentation boundary is not enforced.

In this paper, we address this weakness of the IFT by proposing the *relaxed* IFT (RIFT). This modified version of the IFT features an additional parameter that controls the smoothness of the segmentation boundary, thereby making the results more predictable in the presence of noise and weak edges. The method works by applying an iterated relaxation procedure to the segmentation labels. We show that these computations can be restricted to a narrow band around the segmentation boundary, yielding a fast segmentation algorithm suitable for interactive applications. The efficacy of the relaxation procedure is demonstrated in Figure 1. An early version of this work is presented in Ref. 14.

In addition, we present a study on the application of the relaxed IFT method to the problem of segmenting individual trunk muscles in magnetic resonance (MR) volume images of human athletes (javelin throwers). In these images, contrast between adjacent muscles is poor. The original IFT therefore produces segmentation results with noisy boundaries. Our tests indicate that the relaxed IFT produces more predictable segmentation results for these images.

2. BACKGROUND

2.1 Images and graphs

An *image* \mathbf{I} is a pair (\mathcal{I}, I) consisting of a set \mathcal{I} of image elements and a mapping I that assigns to each image element $p \in \mathcal{I}$ an element in some arbitrary set, typically a subset of \mathbb{Z}^n or \mathbb{R}^n (e.g., $\mathcal{I} \subset \mathbb{Z}^n$ and $I : \mathcal{I} \rightarrow [0, 255]$). We associate an image with an *adjacency* function \mathcal{N} that maps each image element $p \in \mathcal{I}$ to a set $\mathcal{N}(p) \subset \mathcal{I}$ of

adjacent image elements. We require the adjacency function to be symmetric, so that $p \in \mathcal{N}(q) \iff q \in \mathcal{N}(p)$ for all $p, q \in \mathcal{I}$. The image, together with the adjacency function, defines a directed graph, whose nodes are the image elements and whose edges are all ordered pairs of image elements $p, q \in \mathcal{I}$ such that $q \in \mathcal{N}(p)$.

For each ordered pair of adjacent nodes p and q , we assign a real valued, non-negative, *edge weight* $w(p, q)$. Although the proposed method does not require that $w(p, q) = w(q, p)$, this is commonly the case in applications. The edge weights represent local *dissimilarity*, i.e., p and q are strongly connected if $w(p, q)$ is close to 0. Moreover, we define the (*pseudo*) *inverse edge weight* $w^{-1}(p, q)$ as

$$w^{-1}(p, q) = \frac{1}{1 + \beta w(p, q)},$$

where $\beta \in \mathbb{R}$ is a constant. The inverse edge weights represent local *similarity*, i.e., p and q are strongly connected if $w^{-1}(p, q)$ is close to 1.

2.2 Paths and path costs

A *path* $\pi = \langle p_1, p_2, \dots, p_k \rangle$ of length $|\pi| = k - 1$ is a sequence p_1, p_2, \dots, p_k such that $p_{i+1} \in \mathcal{N}(p_i)$. We denote the *origin* p_1 and the *destination* p_k of π by $org(\pi)$ and $dst(\pi)$, respectively. The *cost* of a path is denoted $C(\pi)$. This cost is typically a function of the edge weights along the path, e.g., the sum of all the edge weights along the path or the maximum edge weight along the path. In all examples shown here, the sum of the edge weights was used.

A path π is a *minimum cost path* if $C(\pi) \leq C(\tau)$ for any other path τ with $org(\tau) = org(\pi)$ and $dst(\tau) = dst(\pi)$. Note that in general, the minimum cost path is not unique. The set of minimum cost paths between two nodes p and q is denoted $\pi_{min}(p, q)$.

The definition of a minimum cost path between two sets of nodes is analogous. For two sets $A \subseteq \mathcal{I}$ and $B \subseteq \mathcal{I}$, π is a path between A and B if $org(\pi) \in A$ and $dst(\pi) \in B$. If $C(\pi) \leq C(\tau)$ for any other path τ between A and B , then π is a minimum cost path between A and B . The set of minimum cost paths between A and B is denoted $\pi_{min}(A, B)$.

2.3 Image segmentation

A *segmentation* of a set $A \subseteq \mathcal{I}$, $A \neq \emptyset$ into k classes is a mapping \mathcal{L}_A that assigns to each image element $p \in A$ a *label vector* $\mathcal{L}(p) = \mathbf{l} = (l_1, l_2, \dots, l_k)$, where $l_i \in [0, 1]$ and

$$\sum_{i=1}^k l_i = 1.$$

If $A = \mathcal{I}$ the segmentation is *complete*, otherwise it is *partial*. A complete segmentation $\mathcal{L}_{\mathcal{I}}$ will hereinafter be denoted \mathcal{L} , i.e., we omit the subscript for ease of notation. Each element in the label vector indicates the degree to which the node belongs to the corresponding class. If all $l_i \in \{0, 1\}$, then the segmentation is *crisp*, otherwise it is *fuzzy*.¹⁵ The output of the RIFT is a crisp segmentation, although intermediate steps of the algorithm make use of a fuzzy segmentation. The *fuzziness* $f(\mathcal{L}_A) \in [0, 1]$ of a segmentation \mathcal{L}_A indicates the degree to which \mathcal{L}_A is not crisp. Here, we define $f(\mathcal{L}_A)$ as

$$f(\mathcal{L}_A) = \frac{\sum_{p \in \mathcal{I}} \phi(\mathcal{L}_A(p))}{|A|},$$

where $\phi(\mathbf{l})$, the fuzziness of an individual label vector, is defined as

$$\phi(\mathbf{l}) = 1 - \frac{k \sum_{i=1}^k |l_i - \frac{1}{k}|}{2(k-1)}.$$

A crisp segmentation has fuzziness 0, and a segmentation where all $l_i = \frac{1}{k}$ has fuzziness 1.

3. METHOD

Given an image \mathbf{I} , a path cost function f , an adjacency function \mathcal{N} and a set of seed-points $S \in \mathcal{I}$ with a corresponding partial segmentation \mathcal{L}_S , we seek a complete segmentation. We use the IFT to compute an initial complete segmentation \mathcal{L}^0 . This segmentation has the property that $\mathcal{L}^0(p) = \mathcal{L}_S(q)$ for some $q \in S$ such that $\pi_{\min}(p, q) \subseteq \pi_{\min}(\{p\}, S)$, i.e., each element is assigned the same label as the closest seed-point in the weighted graph. We refine this segmentation by applying an iterative relaxation procedure, described below.

3.1 Relaxation procedure

Starting from \mathcal{L}^0 , we compute a sequence of segmentations $\mathcal{L}^1, \mathcal{L}^2, \dots, \mathcal{L}^N$ by iterative *relaxation*:

$$\mathcal{L}^{i+1}(p) = \begin{cases} \frac{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q) \cdot \mathcal{L}^i(q)}{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q)} & \text{if } p \notin S \\ \mathcal{L}^i(p) & \text{otherwise} \end{cases} .$$

The second case is needed to ensure that the labels of the seed-points are not changed. The number of iterations N , corresponds to the amount of relaxation.

After relaxation has been performed, the segmentation is no longer guaranteed to be crisp. To achieve a final crisp segmentation, \mathcal{L}^N is therefore *defuzzified*. Given a label vector \mathbf{l} with a maximum element l_j , the defuzzified label vector $\hat{\mathbf{l}}$ is given by

$$\hat{l}_i = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} .$$

The maximum element of a label vector is not necessarily unique, and therefore the above definition may give rise to ambiguities. In practice, however, such ambiguities are rare, and may be resolved using an arbitrary additional criterion.

3.2 Implementation details

In a naive implementation of the relaxation procedure, the labels of all image elements are updated at each iteration. This yields an algorithm for the relaxation procedure with computational complexity $O(N|\mathcal{I}|)$. For interactive segmentation of large volume images, this is usually unreasonably slow unless N is very small. The key to a more efficient implementation of the relaxation procedure is to take advantage of the fact that for most iterations, a large subset of the image elements retain their labels. Formally, we say that a node p is *stationary* at iteration i if $\mathcal{L}^i(p) = \mathcal{L}^{i-1}(p)$. A stationary node may be excluded from the computations, without affecting the result. We will now describe a procedure for identifying a set of elements that are guaranteed to be stationary at an iteration, without having to compute the relaxed values. Restricting the domain of computation in this way greatly reduces the computation time.

Given a set $A \subseteq \mathcal{I}$, we define the *expansion* A^+ of A as

$$A^+ = A \cup \{q \mid q \in \mathcal{N}(p) \text{ for some } p \in A\} .$$

To provide some intuition to this concept, we note that if \mathcal{I} is an 8-connected 2D image, then the expansion of a set in \mathcal{I} is equivalent to the dilation of the set by a 3×3 square structuring element.

Lemma 1. If $A \subseteq \mathcal{I}$ and $p \in \mathcal{I} \setminus A^+$, then $\mathcal{N}(p) \subset \mathcal{I} \setminus A$.

Proof. The proof is by contradiction. If $\mathcal{N}(p) \not\subset \mathcal{I} \setminus A$, then there exists a node q such that $q \in A$ and $q \in \mathcal{N}(p)$. Since the adjacency function is symmetric, this means that $p \in \mathcal{N}(q)$, and thus $p \in A^+ \Rightarrow p \notin \mathcal{I} \setminus A^+$. \square

Theorem 1. If $A \subseteq \mathcal{I}$ is a set such that all nodes $p \in \mathcal{I} \setminus A$ are stationary at iteration i , then all nodes $p \in \mathcal{I} \setminus A^+$ are stationary at iteration $i + 1$.

Proof. Consider a node $p \in \mathcal{I} \setminus A^+$. If $p \in \mathcal{S}$, then p is stationary at all iterations by definition. If $p \notin \mathcal{S}$, then $\mathcal{N}(p) \subset (\mathcal{I} \setminus A)$ by Lemma 1. Therefore, any node $q \in \mathcal{N}(p)$ is stationary at iteration i . This gives us

$$\mathcal{L}^{i+1}(p) = \frac{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q) \cdot \mathcal{L}^i(q)}{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q)} = \frac{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q) \cdot \mathcal{L}^{i-1}(q)}{\sum_{q \in \mathcal{N}(p)} w^{-1}(p, q)} = \mathcal{L}^i(p),$$

and thus p is stationary at iteration $i + 1$. \square

The practical implication of Theorem 1 is that if all nodes that are non-stationary at iteration i are contained in a set $\Omega \subseteq \mathcal{I}$, representing the domain of computation, then all nodes that are non-stationary at iteration $i + 1$ are contained in Ω^+ . Therefore, by expanding the domain of computation at each iteration, we can guarantee that all non-stationary nodes are included in the computations. To determine an initial domain of computation, we note that if for some node p it holds that $\mathcal{L}^i(p) = \mathcal{L}^i(q)$ for all $q \in \mathcal{N}(p)$, then p is stationary at iteration $i + 1$. We can now formulate the following computational procedure for the relaxed IFT:

1. Compute the IFT to obtain \mathcal{L}^0 .
2. Define the initial domain of computation $\Omega = \mathcal{I} \setminus \{p \mid \mathcal{L}^0(p) = \mathcal{L}^0(q) \text{ for all } q \in \mathcal{N}(p)\}$.
3. For $i = 1, \dots, N$
 - (a) Compute $\mathcal{L}^i(p)$ for all $p \in \Omega$.
 - (b) $\Omega \leftarrow \Omega^+$
4. Defuzzify \mathcal{L}^N to obtain a final crisp segmentation.

For a typical segmentation scenario on an image with 512×512 image elements, the above procedure reduced the total number of relaxation operations by 99% for $N = 30$ and 90% for $N = 40$, compared to a naive implementation.

4. CASE STUDY

4.1 Segmentation of trunk muscles in MR images

In athletes such as cricket fast bowlers and javelin throwers, asymmetries of trunk muscles can be observed and measured using MR imaging. These asymmetries are of interest in sports medicine and science because they may be predictive of future sports injury. In the case of cricket fast bowlers, for example, substantial volume asymmetries of the quadratus lumborum (QL) muscle have been associated with a significantly increased risk of developing lumbar spine stress fractures.¹⁶ Measurement of these asymmetries requires that the muscles be segmented in the MR images. Jurcak et al.¹⁷ presented an automatic segmentation method for this purpose. The method uses a shape atlas, built from a large number (≈ 20) of manually produced reference segmentations, to segment a specific muscle. Manual segmentation of volume images is a tedious and time-consuming task, and the results are subject to intra- and inter-user variations. We are investigating whether this process can be accelerated, and made more robust, by refining the manual segmentations using the RIFT.

For reasonably well defined objects, a few seed-points are often sufficient for the RIFT to produce good segmentation results. The QL muscle, however, is particularly challenging to segment, because of the lack of contrast between it and adjacent muscles. Careful placement of a large number of seed-points is therefore needed to obtain a good segmentation. We therefore ask a user to provide an approximate delineation of the boundary in each slice of the volume. The boundary of this approximate segmentation is dilated to produce a narrow band, in which the true segmentation boundary is assumed to be contained. All image elements outside this narrow band are used as seed-points. The seed-points are labeled as *background* if they are connected to the image border, and *object* otherwise. A refined segmentation is then obtained by applying the RIFT. The procedure is illustrated in Figure 2. Compared to fully manual delineation, this procedure allows the user to be less precise, and therefore work faster.




Figure 2. Segmentation procedure. (a) A manual delineation of the QL muscle. (b) Object seeds (white) and background seeds (black) are defined by dilating the boundary of the manual delineation with a 8×8 square structuring element. The true boundary of the QL is assumed to be contained within the gray region. (c) Segmentation produced by the IFT. At the weak boundaries between adjacent muscles, the segmentation is dominated by noise. (d) Segmentation produced by the relaxed IFT. The segmentation is qualitatively similar to the manual delineation, but is less sensitive to intra-user variation.

4.2 Experiment

To validate the procedure described above, we undertook a study on a single slice from an MR volume image of the lumbar region of an elite javelin thrower. MR imaging was performed on a 1.5T Siemens Symphony scanner (Siemens AG, Munich) with the subject in the supine position. Images were acquired in the axial plane using a 2D spin echo sequence (TR/TE = 477/11 ms, flip angle = 150 degrees, slice thickness = 3.5mm, slice spacing = 3.5mm, and in-plane resolution = 1mm). The slice is shown in Figure 3. Two experiments were performed. In the first experiment a user performed 20 manual delineations of the left QL muscle, instructed to be as accurate as possible. The task was completed in 7 minutes and 27 seconds. We denote the set of these 20 segmentations MA (for manual accurate). In the second experiment, the same user again performed the same task, but this time with the instruction to be as fast as possible. This task was completed in 4 minutes and 4 seconds. We denote the set of these 20 segmentations MF (for manual fast). For each segmentation in these two sets, a segmentation was produced using the RIFT, with seeds generated according to the procedure described above. The same number of iterations, $N = 30$, was used for all segmentations. This number was determined empirically, by visual inspection of the segmentation results. For reference, we also produced the same segmentations without applying the relaxation procedure. These four sets of segmentations are denoted RIFT(MA), RIFT(MF), IFT(MA), and IFT(MF), respectively.




Figure 3. (Left) The image used in the experiment, a slice from an MRI volume. (Right) Gradient magnitude of the same image, showing the lack of contrast between the QL and adjacent muscles.

Table 1. Segmentation robustness.

Γ	$f(\bar{\Gamma})$ (%)
MA	0.9282
MF	1.1569
IFT(MA)	0.4669
IFT(MF)	0.6812
RIFT(MA)	0.3957
RIFT(MF)	0.6100

4.2.1 Evaluation of robustness

To measure the amount of variation within a set of segmentations Γ , we constructed *one* fuzzy segmentation $\bar{\Gamma}$ by averaging all segmentations within each set, i.e.,

$$\bar{\Gamma}(p) = \frac{\sum_{\mathcal{L} \in \Gamma} \mathcal{L}(p)}{|\Gamma|}$$

for all $p \in \mathcal{I}$. The fuzziness of $\bar{\Gamma}$, as defined in Section 2.3, represents the amount of variation within Γ . In this experiment, $k = 2$ classes were used. Since the elements of the label vector must sum to unity this means that $l_2 = 1 - l_1$, and thus the expression for $\Phi(\mathbf{l})$ may be reduced to

$$\Phi(\mathbf{l}) = 1 - |2l_1 - 1|.$$

If all segmentations within Γ are identical, then the fuzziness of $\bar{\Gamma}$ is 0. If, on the other hand, the segmentations within Γ are completely disjoint, then the fuzziness of $\bar{\Gamma}$ is 1.

The sets $\overline{\text{MA}}$, $\overline{\text{MF}}$, $\overline{\text{IFT(MA)}}$, $\overline{\text{IFT(MF)}}$, $\overline{\text{RIFT(MA)}}$ and $\overline{\text{RIFT(MF)}}$ are shown in Figure 4. The fuzziness of each of these sets is given in Table 1. As the table shows, applying the RIFT reduces the fuzziness by a factor 2. A slightly smaller reduction is obtained with the original IFT. It should also be noted that $\overline{\text{RIFT(MF)}}$ is less fuzzy than $\overline{\text{MA}}$, i.e., when we apply the RIFT to the fast manual segmentations, the result is more robust than the accurate manual segmentations.




Figure 4. Fuzzy segmentations obtained by taking the mean of the segmentation sets obtained in the experiment. The sets are defined in the text.

4.2.2 Qualitative evaluation

Given that no ground truth exists for the image in the experiment, it is difficult to make a quantitative evaluation of the segmentation results. Instead, we performed a qualitative study. To compare two sets of segmentations Γ_1 and Γ_2 , we randomly paired each segmentation in Γ_1 with a segmentation in Γ_2 . Each pair of segmentations was then shown to an observer, who was asked to select the best segmentation of the two. The result of this comparison was coded by a number $\tau \in [0, 1]$. A value of $\tau = 1$ indicates that Γ_2 was preferred in all cases, while $\tau = 0$ indicates that Γ_1 was preferred in all cases.

Pairwise comparisons between the sets $\overline{IFT(MF)}$ and $\overline{RIFT(MF)}$, as well as $\overline{IFT(MA)}$ and $\overline{RIFT(MA)}$ were performed by 8 observers. The results of these comparisons are shown in Figure 5(a–b). The segmentations obtained by the RIFT were preferred over the segmentations obtained by the IFT without relaxation. This validates our statement that in the presence of noise and weak edges, segmentations produced by the RIFT are more intuitively correct than those produced by the IFT.

Additionally, pairwise comparisons between the sets MA, MF, $\overline{RIFT(MA)}$ and $\overline{RIFT(MF)}$ were performed by 12 observers. The results of these comparisons are shown in Figure 5(c–f). The comparisons 5(c) and 5(d) show that the segmentations obtained by applying the RIFT were judged to be as good as, or slightly better than, the manual segmentations. Moreover, the comparisons 5(e) and 5(f) show that applying the RIFT to the segmentations in MF improves the perceived quality of the segmentation results to a level where they are comparable to the segmentations in MA.




Figure 5. Results of the qualitative evaluation. For each comparison, the gray box shows the mean τ -value of the 8 (12) observers, ± 1 standard deviation. The black dots indicate the minimum and maximum τ -values.




Figure 6. 3D surface renderings of a QL muscle, segmented without relaxation (Left) and after 20 iterations of the proposed relaxation procedure (Right). The relaxed result corresponds better to the true appearance of the muscle. The segmentations were computed in 12.0s and 15.7s, respectively.

4.3 3D segmentation

Although the experiment was performed on a 2D slice, our preliminary tests indicate that similar results can be expected for 3D segmentation. To extend the seed-generation procedure described above to 3D, approximate delineations must be obtained in all slices of the volume. A 3D segmentation result is shown in Figure 6. In this example, approximate delineations were produced manually in all 38 slices. In many cases, however, it may be sufficient to perform the manual delineations in a subset of the slices, and use interpolation to provide seed-points for the entire volume. Thereby, the user-time required for segmentation would be reduced even further.

5. CONCLUSIONS AND FUTURE WORK

We have introduced the RIFT, a modified version of the IFT method for seeded segmentation. We have showed that the RIFT produces more intuitively correct segmentation results in the presence of noise and weak boundaries than the original IFT algorithm.

In our case study, we apply the RIFT to refine manual segmentations of a thoracolumbar muscle in MR images. Based on the results of the study, we conclude that the refined segmentations are qualitatively comparable to the manual segmentations, while intra-user variation is reduced by more than 50%. We have validated the method on 120 segmentations of a muscle in a 2D slice from an MR volume. As the next step, we will extend this study to 3D.

In our current implementation of the method, the initial computational domain Ω is found by iterating over all image elements. In future work, we intend to investigate whether this set can be obtained as a by-product of the IFT computation. This would allow a differential implementation of the RIFT, similar to the differential IFT,⁷ which significantly reduces the total time required for interactive segmentation.

ACKNOWLEDGMENTS

The authors wish to thank Dr. Robin Strand at the Centre for Image Analysis, Uppsala University, Sweden for valuable comments on this manuscript.

REFERENCES

- [1] Adams, R. and Bischof, L., “Seeded region growing,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**(6), 641–647 (1994).
- [2] Mehnert, A. J. H. and Jackway, P. T., “An improved seeded region growing algorithm,” *Pattern Recognition Letters* **18**, 1065–1071 (1997).
- [3] Meyer, F., “Un algorithme optimal de ligne de partage des eaux,” in [*Proceedings of the 8^e Congres Reconnaissance des Formes et Intelligence Artificielle*], 847–857 (1991).
- [4] Boykov, Y. and Funka-Lea, G., “Graph cuts and efficient N-D image segmentation,” *International Journal of Computer Vision* **70**(2), 109–131 (2006).
- [5] Grady, L., “Random walks for image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(11), 1768–1783 (2006).
- [6] Falcão, A. X., Stolfi, J., and Lotufo, R. A., “The image foresting transform: Theory, algorithms, and applications,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(1), 19–29 (2004).
- [7] Falcão, A. X. and Bergo, F. P., “Interactive volume segmentation with differential image foresting transforms,” *IEEE Transactions on Medical Imaging* **23**(9), 1100–1108 (2004).
- [8] Strand, R., Malmberg, F., and Svensson, S., “Minimal cost-path for path-based distances,” in [*Proceedings of Image and Signal Processing and Analysis (ISPA)*], 379–384 (2007).
- [9] Dijkstra, E. W., “A note on two problems in connexion with graphs,” *Numerische Mathematik* **1**, 269–271 (1959).
- [10] Falcão, A. X., Udupa, J. K., and Miyazawa, F. K., “An ultra-fast user-steered image segmentation paradigm: Live wire on the fly,” *IEEE Transactions on Medical Imaging* **19**(1), 55–62 (2000).
- [11] Vidholm, E., Nilsson, S., and Nyström, I., “Fast and robust semi-automatic liver segmentation with haptic interaction,” in [*Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*], 774–781 (2006).
- [12] Bai, X. and Sapiro, G., “A geodesic framework for fast interactive image and video segmentation and matting,” in [*Proceedings of the International Conference on Computer Vision (ICCV)*], IEEE (2007).
- [13] Kauffman, C. and Piché, N., “Seeded ND medical image segmentation by cellular automaton on GPU,” *International Journal of Computer Assisted Radiology and Surgery* (2009).
- [14] Malmberg, F. and Nyström, I., “Interactive segmentation with relaxed image foresting transforms,” in [*Proceedings of Swedish Symposium on Image Analysis (SSBA)*], Bigun, J. and Verikas, A., eds., 17–20 (2009).

- [15] Zadeh, L. A., "Fuzzy sets," *Information and Control* **8**, 338–353 (1965).
- [16] Engstrom, C. M., Walker, D. G., Kippers, V., and Mehnert, A. J. H., "Quadratus lumborum asymmetry and L4 pars injury in fast bowlers: a prospective MR study.," *Medicine and Science in Sports and Exercise* **39**(6), 910–917 (2007).
- [17] Jurcak, V., Fripp, J., Engstrom, C., Walker, D., Salvado, O., Ourselin, S., and Crozier, S., "Automated segmentation of the quadratus lumborum muscle from magnetic resonance images using a hybrid atlas based - geodesic active contour scheme," in [*Proceedings of the International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*], 867–870, IEEE (2008).