AUTOMATIC SKELETAL MUSCLE SEGMENTATION THROUGH RANDOM WALKS AND GRAPH-BASED SEED PLACEMENT

P-Y. Baudin[1−7]*, N. Azzabou*[5−7]*, P. G. Carlier*[5−7]*, N. Paragios*[2−4]

[1]SIEMENS Healthcare, Saint Denis, FR ; [2]Center for Visual Computing, Ecole Centrale de Paris, FR ;
[3]Université Paris-Est, LIGM (UMR CNRS), Center for Visual Computing, Ecole des Ponts ParisTech, FR ;
[4]Equipe Galen, INRIA Saclay, Ile-de-France, FR ; [5]Institute of Myology, Paris, FR ;
[6]CEA, I[2]BM, MIRCen, IdM NMR Laboratory, Paris, FR ; [7]UPMC University Paris 06, Paris, FR.

ABSTRACT

In this paper we propose a novel skeletal muscle segmentation method driven from discrete optimization. We introduce a graphical model that is able to automatically determine appropriate seed positions with respect to the different muscle classes. This is achieved by taking into account the expected local visual and geometric properties of the seeds through a pair-wise Markov Random Field. The outcome of this optimization process is fed to a powerful graph-based diffusion segmentation method (random walker) that is able to produce very promising results through a fully automated approach. Validation on challenging data sets demonstrates the potentials of our method.

Index Terms— muscle, image segmentation, graphical models

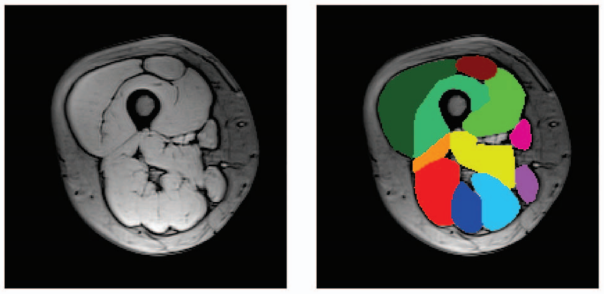


Fig. 1. (left) Cross-section of an MR volume of the thigh. (right) Manual segmentation of the muscles.

1. INTRODUCTION

In the field of medical image processing, skeletal muscle segmentation is a challenging problem. This is due to the fact that the basic segmentation assumption of visual discriminability between different classes is violated. Muscle regions consist of the same tissue and therefore their visual appearance is not distinctive. User-aided segmentation and delineation could address such a problem but is a time consuming and tedious process. Despite enormous progress on automatic segmentation methods, their direct application to such a clinical setting is far from being trivial. Skeletal muscles have large shape variation and no specific texture, while the contours separating them are often nonexistent or confused with noise (blood vessels, fat infiltrations).

The above issues are the predominant reason for limited prior work in the field. In [1, 2], the task of segmenting all the muscles in one limb was addressed. The proposed method belongs to the Deformable Model family of segmentation methods ([3, 4]), that rely on minimizing an energy balancing a data term - which pushes the model towards the target contours - and a regularization term - which imposes a smooth solution along the curve. The main drawback of such method is that only a local optimum is reached during optimization.

In [5], an alternative approach was considered that aimed at introducing prior knowledge in the process. The allowed deformations of the muscle were modeled through hierarchical wavelets, resulting on a sparse local-to-global representation. Gradient descent was also used to determine the lowest potential of the designed energy. In [6], the problem was approached using a graph-based method. Higher-order pose-invariant priors were considered to model shape variability and classifications methods were used to determine model-to-image correspondences. Linear programming and dual-decomposition were adopted towards optimization of the defined objective function. However the extension of these methods in the multi-object setting is not trivial. Moreover, the drawback of all knowledge-based models is their dependency on robust statistical learning methods, which, to function properly, generally require: a) a large annotated database and b) consistent object features to be learned. Both items are difficult to observe in our application in practice. Indeed, the complexity of the skeletal muscle structure implies that the database contains dozens of segmented subjects (a long and tedious work), and, besides, muscles have few detectable consistent landmark points, while their shape varies widely among individuals.

Our approach builds upon the observation that certain region-based methods can achieve excellent results, providing a substantial manual initialization. In [7, 8], the user draws "seeds", i.e. annotates a few pixels, from which the algorithm extrapolates the full segmentation, through high performance discrete optimization processes. The method proposed in [8] has the advantage of addressing the issue of incomplete contours (because of the random walks process) which is one of the main difficulties of our task. However, the performance of these methods depends heavily on the position of the initial seeds, that is also a time consuming process.

In this paper, we propose a method for generating the seeds automatically. Without loss of generality, we assume a rigid pre-registration step between a segmented atlas and the target volume. Through a sampling process, a number of seeds are generated in the image domain. The clustering of these seeds with respect to the different muscle classes is achieved through a graph-based approach.

This paper is organized as follows: in section 2 we detail the formulation of our method. Then, in section 3, we indicate how to solve numerically the problem, and present results obtained on 3D MR data. Section 4 concludes the paper.

2. AUTOMATIC SEED GENERATION

2.1. Seeds Sampling

The first step of our algorithm is to generate the seeds that will be used as initialization by the region-based algorithm - e.g. Random Walks (RW) - once labeled. For efficiency, we need to place the seeds well inside the muscles, and not too close to the contours. To achieve this, we first compute an edge map of our volume - e.g. by computing the local variance at each pixel for a given radius. Then, we generate regularly placed seeds, with a high enough density to ensure all the muscles contain several seeds. Finally, we move the seeds in the opposite direction of the gradient of the edge map for a few iterations. This way, seeds are moved away from the close-by edges by a few pixels (see figure 2).

2.2. MRF formulation

We now describe how to assign a label to the seeds we generated, according to the similarities between the test image and the reference segmentation.

2.2.1. Energy Form

We formulate our problem as a labeling problem, that we will solve by minimizing a first-order Markov Random Fields (MRF) energy function. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph, where $\mathcal{V}$ is the set of nodes and $\mathcal{E}$ is the set of edges. Given a set of labels $\mathcal{L}$, we want to assign a label $l \in \mathcal{L}$ to each node $p \in \mathcal{V}$. In this framework, the nodes in $\mathcal{V}$ are our unlabeled seeds, the labels in $\mathcal{L}$ are the indices of the muscles in the reference segmentation. We denote x_p the label assigned to p, and $\mathbf{x}$ the collection of all assignments. A first order MRF energy is of the form:

$$E(\mathbf{x}) = \sum_{p \in \mathcal{V}} \theta_p(x_p) + \sum_{(p,q) \in \mathcal{V}} \theta_{p,q}(x_p, x_q) \quad (1)$$

where $E(\cdot)$ is the energy (or cost function) we want to minimize, the unary potential $\theta_p(\cdot)$ depends only on the label assigned to the node p, and the binary potential $\theta_{p,q}(\cdot,\cdot)$ depends on the labels assigned to each nodes of the edge (p, q).

2.2.2. Unary Potential

We define a unary cost $\theta_p(\cdot)$ based on the knowledge that the test image is rigidly registered to the reference segmentation:

$$\theta_p(m) = d(p, M_m)^\alpha / \Theta^\alpha$$

where M_m is the set of pixels with label m in the reference segmentation, α and Θ are manually chosen parameters, and $d(\cdot,\cdot)$ is a distance function between a point and a set:

$$d(p, R) = \begin{cases} \min\{\|p - r\|, r \in R\} & \text{if } p \notin R \\ 0 & \text{otherwise} \end{cases}$$

The rigid registration ensures that, if we superimpose the test image and the reference segmentation, an unknown muscle in the test image is not very far from the corresponding muscle in the reference segmentation. This term will privilege the assignment to a node of a label corresponding to a geometrically close muscle, over the label of a more remote muscle. Let us now proceed with the definition of the pair-wise terms, driven from the topology and the geometry of the expected solution.

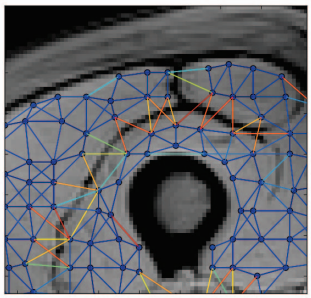


Fig. 2. Sampled seeds and graph edges. Edge color represent the geodesic distance - from blue: short distances, to red: long distances.

2.2.3. Geodesic Distance Potential

The unary term does not take into account the structure of the set of muscles. While the shape of the muscles can vary quite extensively between individuals. The topology of the set of muscles is consistent: muscle a is *always* "close to" muscle b, but "far from" muscle c, there is *most of the time* a visible contour between muscle a and b. We designed binary terms to account for this knowledge. The binary potential is the sum of two potentials:

$$\theta_{p,q}(\cdot,\cdot) = w_{geod}\theta^{geod}_{p,q}(\cdot,\cdot) + w_{orient}\theta^{orient}_{p,q}(\cdot,\cdot)$$

with weighting parameters w_{geod} and w_{orient}.

The term $\theta^{geod}_{p,q}(\cdot,\cdot)$ is derived from the geodesic distance $g(p, q)$ between the seeds p and q:

$$\theta^{geod}_{p,q}(x_p, x_q) = \begin{cases} 1/(1 + \exp(\gamma - g(p,q))) & \text{if } x_p = x_q \\ 1/(1 + \exp(g(p,q) - \gamma)) & \text{otherwise} \end{cases}$$

where γ is a parameter. For large geodesic distances $g(p, q)$, the case $x_p = x_q$ will be more penalized than $x_p \neq x_q$, as it is likely that there is a contour between nodes p and q . For small $g(p, q)$, this potential will favor $x_p = x_q$ over $x_p \neq x_q$ (cf. figure 2).

We compute the geodesic distance between the seeds by running the Fast Marching algorithm [9] on the test image. We supply the algorithm with the unlabeled sampled seeds. We obtain in return a partition of the image, where each seed in included in a different region, and where the boundary between two regions is equidistant to the seeds in the regions, in the geodesic sense (i.e. a geodesic Voronoi partition).

Note that we do not get the geodesic distance between every pair of seeds, but only between geometrically proximal seeds. This is not a problem, since $\theta^{geod}_{p,q}$ is designed to account for the presence or the absence of a contour between two seeds, and thus is mostly a local term. We also use this newly computed edge set to determine the connectivity of the graph, obtaining a graph with mainly local edges. This is motivated by the fact that our unary cost already handle the problem at a large scale. However, we add a few random edges to the graph, to strongly reduce the diameter of the graph while not raising the complexity too much. For these non local edges, the binary cost is made only of the orientation cost.

2.2.4. Relative Orientation Potential

The term $\theta^{orient}_{p,q}(\cdot,\cdot)$ is intended to insure that the relative position of a pair of muscles is maintained, e.g. muscle a is located "top right of" muscle b. In a pre-processing stage, for each pair (m, n) of muscles, we computed the estimate of the probability density (the

histogram) $h_{m,n}(\cdot)$ of random variable:

$$\mathbf{U} = (P_m - P_n) / \|P_m - P_n\|$$

where P_m is a randomly chosen point in muscle M_m, and $P_n = \arg\min_{p \in M_n} \|P_m - p\|$. Thus, $h_{m,n}(\mathbf{u})$ can be seen as the probability that, for any pixel in muscle m, the direction of the closest pixel in muscle n is $\mathbf{u}$. We now can define the orientation cost as:

$$\theta_{p,q}^{orient}(x_p, x_q) = \begin{cases} -\log h_{x_p,x_q}(p-q) & \text{if } x_p \neq x_q \\ 0 & \text{otherwise} \end{cases}$$

This cost will be large if the orientation of the edge between two nodes is unlikely, given the labels to which the nodes are assigned.

3. EXPERIMENTAL VALIDATION

3.1. MRF Optimization

Minimizing the energy (1) is a NP-hard problem. However, it can be efficiently minimized to a good approximation with the convergent Tree-Reweighted (TRW) Message Passing algorithm ([10]). The TRW algorithm uses a linear programming relaxation technique and operates on the dual of the relaxed problem in order to find a lower bound to the energy of the original problem. Due to its state of the art performance among discrete optimization methods, TRW has been applied to solve many computer vision and medical imaging tasks.

3.2. Experimental Results

To evaluate our method, we used data set composed of 3D volumes of the right thigh of 15 healthy subjects, covering a wide range of morphologies (8 females, 7 males, ages range: 14 to 60), acquired with a 3T Siemens scanner and using 3pt Dixon sequence (TR=10ms, TE1=2.75 ms TE2=3.95 ms TE3=5.15 ms, rf flip angle =3° ; we used the out-of-phase image), with the following resolution: 1mm×1mm×5mm. Every volume was manually segmented to obtain the ground truth against which the segmentation results are compared. We focused our evaluation on clinically relevant muscles of the thigh (13 muscles).

Sampling the seeds, as described in 2.1, with grid spacing 7px×7px×5px, and 2 iterations to bring the seeds further away from the edges (filter radius = 1px), gave us between 2000 and 2500 seeds per test image. At the MRF stage (2.2), we obtained our best results with the following cross-validation estimated parameters: $\alpha = 3$, $\Theta = 0.2 \times \max_{p \in \mathcal{V}, m \in \mathcal{L}} d(p, M_m)$, $w_{geod} = 0.1$, $w_{orient} = 0.01$, $\gamma = 2 \times \text{median}_{p,q \in \mathcal{V}} g(p,q)$, and U contains 26 regularly spread unitary orientation vectors. On a 2.8 GHz Intel® processor with 4 GB of RAM, total processing time is under 5 min, which adds up to the RW segmentation time (20 min).

All our results are based on a leave-one-out validation protocol: each volume is used as the reference segmentation for segmenting all the other volumes. First, we evaluated the automatic labeling process by computing the labeling error rate for each muscle (figure 3). For testing the efficiency of the binary potential, we also computed the error rate for an energy only composed of the unary potential. We observe that using the binary potential significantly reduce the error rate.

Then, we computed Dice coefficients to evaluate the segmentation results given by the RW with the labeled seeds (figure 4). These results are compared with the Dice coefficients of segmentation results with the same generated seeds with ground truth labeling. Not surprisingly, the segmentation results are consistently not as good as

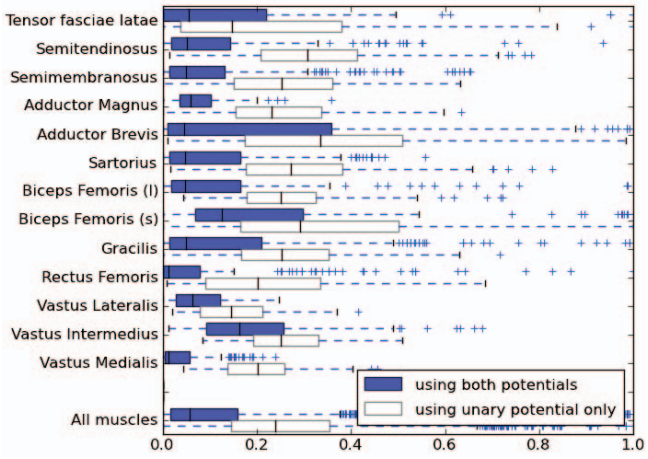


Fig. 3. Labeling error rate per muscle (number of mis-labeled seeds over total seed number). Box-plot presentation: the boxes contain the middle 50% of the data and the median value, and the extremities of the lines indicate the min and max values, excluding the outliers (for more details, see the documentation of Matplotlib). Using the binary potentials significantly reduce the error rate.

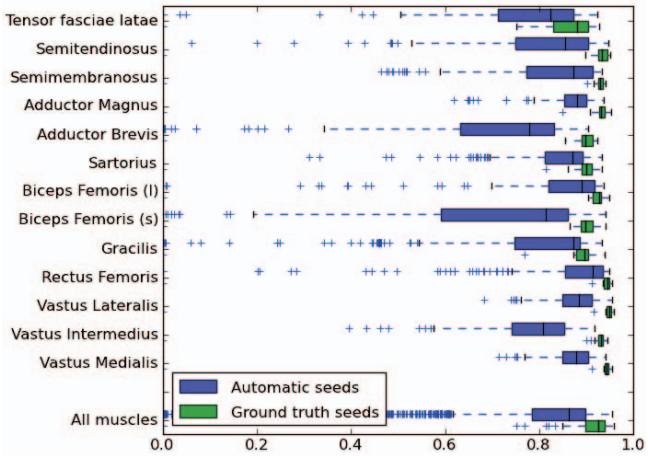


Fig. 4. Dice coefficients on segmentation results. The expression of the Dice coefficient is: $D = 2|T \cap R| / (|T| + |R|)$, where T and R are the pixel sets for the same muscle in the computed segmentation and the reference segmentation respectively.

the ones obtained with the ground truth seeds. However, the performance difference is not that significant in most of the cases, with exceptions of large labeling errors (and thus segmentation errors), which are most probably due to the large shape variability of the corresponding muscles. In figure 5, we show cross-sections of segmentation results. We observe that small muscles tend to be more affected by segmentation errors than large ones, which points out the limitations of our model. In particular, the unary term, based on the surimposition of the registered reference atlas with the target image, is bound to be less effective for small muscles, because the distance term $d(.,.)$ is likely to be larger than in the case of large muscles. Besides, many errors are due to the absence of reliable contours between parts of muscles, which, in the case of large errors, indicates that the topology constraints are failing to insure topology or shape correctness.

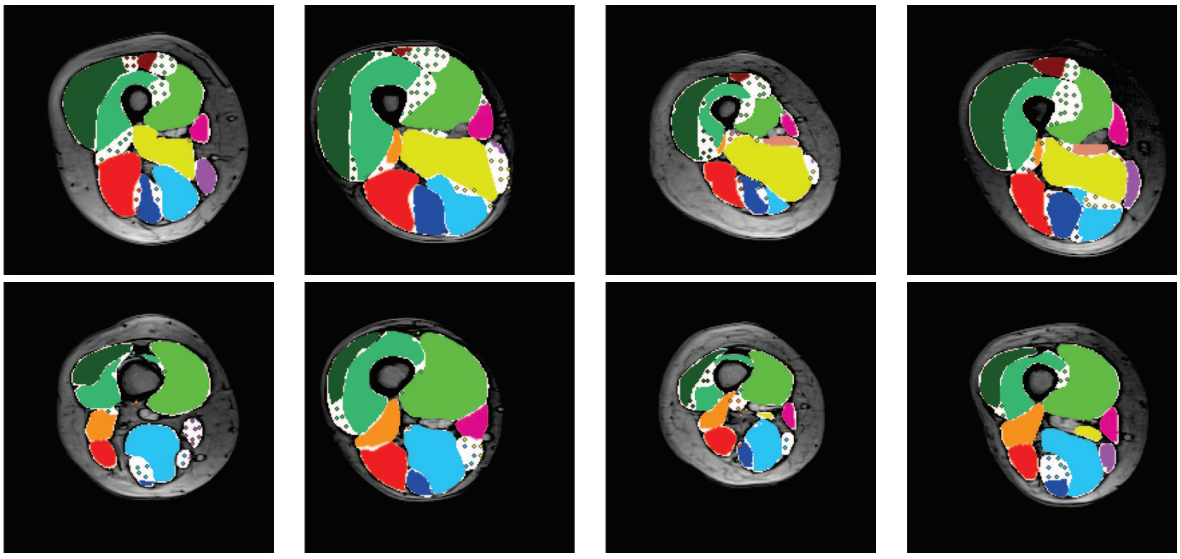


Fig. 5. Segmentation results using automatic seeds. Segmentation errors are shown in white. The mis-labeled seeds are also represented. Please refer to the text for interpretation.

4. CONCLUSION

Segmentation of muscles is a too often neglected problem. Standard methods are often ill-adapted to the complexity and specificity of the full segmentation of one limb, whether they be contour-based, or knowledge-based. Indeed, muscles have partial contours, no discernible texture differences, large variation inter-individuals and unremarkable shapes. However, Seed/Region-based methods show excellent contour accuracy providing the seeds are well placed. We propose a method to replace, the manual initialization of such algorithms. Since we rely on the RW to find the existent contours, our model does not need to focus on targeting the contours in the test image, but rather on identifying the muscles thanks to their relative positions and global organization. Promising results demonstrate the potentials of our method.

Future consists on alleviating the bias introduced from the rigid registration step. This can be achieved by the use of deformable registration methods. Furthermore, the proposed seed-labeling/placement process can have a statistical interpretation and therefore being associated with measures of uncertainties. The use of such measures (min-max marginals) could further enhance the performance of the random walk segmentation method when associated with seed placement confidence. Introducing soft prior knowledge to the Random Walker algorithm is also a promising direction; at least for the muscles with important volumetric mass.

5. REFERENCES

[1] Benjamin Gilles and Dinesh K Pai, "Fast musculoskeletal registration based on shape matching.," *MICCAI*, vol. 11, no. Pt 2, pp. 822–9, Jan. 2008.

[2] Benjamin Gilles and Nadia Magnenat-Thalmann, "Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations.," *Medical image analysis*, vol. 14, no. 3, pp. 291–302, June 2010.

[3] J Montagnat, H Delingette, and N Ayache, "A review of deformable surfaces : topology , geometry and deformation," *Image and Vision Computing*, vol. 19, no. 2001, 2004.

[4] T McInerney and D Terzopoulos, "Deformable models in medical image analysis: a survey.," *Medical image analysis*, vol. 1, no. 2, pp. 91–108, June 1996.

[5] Salma Essafi, Georg Langs, and Nikos Paragios, "Hierarchical 3D diffusion wavelet shape priors," in *CVPR*. Sept. 2009, pp. 1717–1724, IEEE.

[6] Chaohui Wang, Olivier Teboul, Fabrice Michel, Salma Essafi, and Nikos Paragios, "3D knowledge-based segmentation using pose-invariant higher-order graphs.," *MICCAI*, vol. 13, no. Pt 3, pp. 189–196, 2010.

[7] Y.Y. Boykov and M.-P. Jolly, "Interactive graph cuts for optimal boundary & region segmentation of objects in N-D images," *ICCV*, vol. 1, no. July, pp. 105–112, 2001.

[8] Leo Grady, "Random walks for image segmentation.," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 11, pp. 1768–1783, 2006.

[9] J a Sethian, "A fast marching level set method for monotonically advancing fronts.," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, no. 4, pp. 1591–5, Feb. 1996.

[10] Vladimir Kolmogorov, "Convergent tree-reweighted message passing for energy minimization.," *IEEE transactions on pattern analysis and machine intelligence*, vol. 28, no. 10, pp. 1568–83, Oct. 2006.