

Prior Knowledge, Random Walks and Human Skeletal Muscle Segmentation

P.-Y. Baudin¹⁻⁷, N. Azzabou⁵⁻⁷, P.G. Carlier⁵⁻⁷, and Nikos Paragios²⁻⁴

¹SIEMENS Healthcare, Saint Denis, FR

²Center for Visual Computing, Ecole Centrale de Paris, FR

³Université Paris-Est, LIGM (UMR CNRS), Center for Visual Computing,
Ecole des Ponts ParisTech, FR

⁴Equipe Galen, INRIA Saclay, Ile-de-France, FR

⁵Institute of Myology, Paris, FR

⁶CEA, I²BM, MIRCen, IdM NMR Laboratory, Paris, FR

⁷UPMC University Paris 06, Paris, FR

Abstract. In this paper, we propose a novel approach for segmenting the skeletal muscles in MRI automatically. In order to deal with the absence of contrast between the different muscle classes, we proposed a principled mathematical formulation that integrates prior knowledge with a random walks graph-based formulation. Prior knowledge is represented using a statistical shape atlas that once coupled with the random walks segmentation leads to an efficient iterative linear optimization system. We reveal the potential of our approach on a challenging set of real clinical data.

1 Introduction

Segmentation of the skeletal muscles is of crucial interest when studying myopathies. Diseases understanding, patient monitoring, etc. rely on discriminating the muscles in anatomical images. However, delineating the muscle contours manually is an extremely long and tedious task, and thus often a bottleneck in clinical research. Simple automatic segmentation methods rely on finding discriminative visual properties between objects of interest, accurate contour detection or clinically interesting anatomical points. However, skeletal muscles show none of these features and as a result, automatic segmentation is a challenging problem. In spite of recent advances on segmentation methods, their application in clinical settings is difficult, and most of the times, manual segmentation/correction is still the only option.

Among the limited amount of work on this specific subject, in [1,2] a method based on deformable models was proposed to perform the segmentation of all the muscles in one limb. Deformable models are surface models which are fitted to the target image by minimizing a functional balancing a data term - pushing the model towards the target contours - and a regularization term - which imposes a smooth solution along the curve. Such models only reach a local optimum of the functional, which can be far from the desired solution, and depend heavily on

their initial position. In [3], a more efficient shape representation was introduced where prior knowledge was encoded through diffusion wavelets to reduce the space of solutions and thus relax the smoothing constraints. The surface of one muscle was modeled through a hierarchical representation, and the set of allowed deformations at each scale was learned from an annotated training set. Modeling the surface of one muscle with landmark points is an efficient alternative, as proposed in [4] with a graph-based method. The shape variability was modeled through high-order pose-invariant priors and the data term relied on classification and detection of the landmark points. The graph-based framework allowed to perform an efficient non-local optimization, without initialization. However, such method requires to be able to learn consistent image features in order to detect the landmarks, which is difficult to insure in practice in the case of muscles. More generally, all surface models suffer from the absence of reliable contours in MRI images of skeletal muscles. Recently, in [5], a model-based method operating in the image domain with promising results was proposed. This approach consisted in modeling a segmentation through Principal Component Analysis in an Isometric Log-ratio space. Then, a gradient descent was performed with respect to the PCA coefficients to minimize an energy functional which allows label transition only along detected contours. In this method, contour detection is explicit, achieved in a pre-processing stage, and could be a weak link in the chain in cases of undetected or spurious contours.

Our approach builds upon the general Random Walker Segmentation algorithm proposed in [6]. The strength of this method relies on its robustness in the case of incomplete contours and its efficient optimization. While originally this method required manual interaction - an user had to annotate a few pixels of each desired object - the possibility of using prior knowledge based on intensity distribution was introduced in [7]. In this paper, we propose to build a prior model of the shape of the thigh muscles from a training data set, to be used in the RW framework. The prior term of our functional is derived from learning a Gaussian model of the RW unknown probability vector. We also propose to modulate the strength of the model constraints according to the strength of the contours found in the segmented image.

This paper is organized as follows: in section 2 we briefly recall the principle of the RW segmentation and detail the formulation of our model. Then, in section 3, we present segmentation results obtained on 3D MR volumes of the right thigh. Section 4 concludes the paper.

2 Random Walks Segmentation With Prior Knowledge

Notations Let us consider an image I with N pixels, and I_i the gray-level of pixel i . The segmentation is formulated as a labeling problem of an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of edges. Given \mathcal{S} , a set of labels, we want to assign a label $s \in \mathcal{S}$ to each node $p \in \mathcal{V}$. In this framework, the node v_i is the i -th pixel, and to each label corresponds a muscle.




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

2.1 Random Walks Formulation

Let us proceed with a review of the random walks algorithm for image segmentation. We refer the reader to [6] for an extensive description of this method.

The RW approach provides the probability x_i^s that the node $v_i \in \mathcal{V}$ is assigned to the label s . In its original formulation, one has to provide the algorithm with a few already labeled (marked) nodes, also called “seeds”. Typically, the user will manually mark some pixels of each object to be segmented with a different label. Lets denote \mathcal{V}_M the set of marked nodes and \mathcal{V}_U the set of unmarked nodes, such that $\mathcal{V}_U \cap \mathcal{V}_M = \emptyset$ and $\mathcal{V}_U \cup \mathcal{V}_M = \mathcal{V}$. It was shown [6] that all *unknown* entries of $x^s = [x_1^s, x_2^s, \dots, x_N^s]^T$ - i.e. the probabilities that each node $v_i \in \mathcal{V}_U$ is assigned to label s - can be obtained through the minimization of:

$$E_{RW}^s(x^s) = x^{sT} L x^s \tag{1}$$

where the *known* entries of x^s (the seeds) are set as follow:

$$\forall v_i \in \mathcal{V}_M, x_i^s = \begin{cases} 1 & \text{if pixel } i \text{ is marked with label } s \\ 0 & \text{if pixel } i \text{ is marked with another label} \end{cases} \tag{2}$$

and where L is the combinatorial Laplacian matrix of the graph, defined as:

$$L_{i,j} = \begin{cases} \sum_k w_{kj} & \text{if } i = j \\ -w_{ij} & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

with

$$w_{ij} = \omega + \exp -\beta (I_i - I_j)^2 \tag{4}$$

where β is a scaling parameter to be set according to the contrast of the image, and ω is a regularization parameter which amounts to penalizing the gradient norm of x^s (no regularization if $\omega = 0$). After minimizing E_{RW}^s for each label s , the segmentation is obtained by retaining the label of maximum probability: $l_i = \arg \max_s x_i^s$.

2.2 Prior Knowledge

In [7], prior appearance knowledge to the RW formulation was introduced through an estimate of the probability distribution of the gray-level intensity for each label. A prior appearance term is simply added to the RW cost function, balanced by a parameter γ :

$$E_{\text{RWP}}^s(x^s) = x^{s\ T} L x^s + \gamma (x^{s\ T} D x^s - 2x^{s\ T} d^s) \tag{5}$$

where $d^s(i)$ is the probability that the intensity at pixel i belongs to the intensity distribution for label s , and $D = \text{diag}(\sum_s d^s)$ (we refer the reader to [7] for details). In the context of muscle segmentation, the intensity distributions of the labels (the muscles) are extremely similar resulting in an inefficient prior. Moreover, we could think of no other discernible and discriminative features (textures, remarkable points, etc.) to use within this framework. Thus, we decided to learn a pixel-based model of the shape based on previous segmentations of images in a training set \mathcal{D} .

Assume we know \bar{x}_i^s and $\sigma_i^{s\ 2}$, respectively the mean and the variance of x_i^s . Our model simply penalizes the deviation of vector x_i^s from \bar{x}_i^s , weighted by the inverse of $\sigma_i^{s\ 2}$. In vector form, we obtain the following functional:

$$E_{\text{model}}^s(x^s) = (x^s - \bar{x}^s)^T A_\sigma^s (x^s - \bar{x}^s) \tag{6}$$

where A_σ^s is a diagonal matrix such that $A_\sigma^s(i, i) = 1/\sigma_i^{s\ 2}$.

This is equivalent to modeling x_i^s as a random variable with Normal distribution $\mathcal{N}(\bar{x}_i^s, \sigma_i^{s\ 2})$, and maximizing the log probability of x_i^s . Since x_i^s is a probability, such Gaussian modeling can only be a rough approximation. The mean and variance are estimated by computing, respectively, the empirical mean and the empirical variance over a training base of non-rigidly registered segmented images. When one owns only a small number of training examples, the empirical variance is known to be a particularly inefficient estimator. In [8], an improved locally-smooth estimator was proposed, for using as a similar shape prior in the level-set framework. The new estimate is computed, through a gradient descent, as the minimum of a functional which combines the log-likelihood of the training data and a spatial regularization term :

$$\tilde{\sigma}^s = \arg \min_{\sigma} \sum_{i=1}^N \left(\sum_{d \in \mathcal{D}} \log \sigma_i^2 + \frac{(x_{d,i}^s - \bar{x}_i^s)^2}{\sigma_i^2} \right) + \alpha \sum_{i,j=1}^N \delta_{i,j} (\sigma_i^2 - \sigma_j^2)^2 \tag{7}$$

where $x_{d,i}^s = 1$ if pixel i of training data d has label s ($x_{d,i}^s = 0$ otherwise), $\delta_{i,j} = 1$ if pixels i and j are neighbors ($\delta_{i,j} = 0$ otherwise), and α is a weighting parameter setting the degree of smoothing.

We combine energy functionals (1) and (6) by introducing a balancing parameter λ^s :

$$E_{\text{total}}^s(x^s) = E_{\text{RW}}^s(x) + \lambda^s E_{\text{model}}^s(x) \tag{8}$$

It is possible to set a different value of λ^s for each label s , as some muscles may require a stronger influence from the prior model than others. The solution which minimizes (8) verifies:

$$(L + \lambda^s A_\sigma^s) x = \lambda^s A_\sigma^s \bar{x}^s \tag{9}$$

As noted in [7], when one adds such a prior term to the RW functional, it is no longer necessary to own pre-labeled nodes (seeds) in order to compute the segmentation. Indeed, the system of equations (9) is directly invertible, even when all entries of x are unknown. However, it is still possible - and useful - to use seeds to obtain more robust segmentations.

2.3 Confidence Map

As we saw previously, the functional $E_{\text{model}}^s(x^s)$ penalizes the deviation of x^s from the mean \bar{x}^s . Such prior is all the more useful as the local uncertainty of contour presence is large. One can impose such a condition by adjusting the influence of the model according to the strength of the contours in the test image: the *stronger* the contours, the *least* we should rely on the model. Assume we possess such a “confidence map” c , with values close to 0 on strong contours, and values close to 1 in homogeneous regions, we replace the term (6) by the following:

$$E_{\text{model}}^s(x^s) = (x^s - \bar{x}^s)^T A_c A_\sigma^s (x^s - \bar{x}^s)^T \tag{10}$$

where A_c is a diagonal matrix with c on the diagonal.

The local confidence of the image can be easily determined using a decreasing function inversely proportional to the image variance (see figure 1):

$$c_i = \exp -k_v \sigma_r^2(i) \tag{11}$$

where $\sigma_r^2(i)$ is the variance at pixel i computed on a patch with radius r , and k_v is a free parameter. The system to solve is now:

$$(L + \lambda^s A_c A_\sigma^s) x^s = \lambda^s A_c A_\sigma^s \bar{x}^s \tag{12}$$

3 Experimental Validation

Our data set comprises 14 3D volumes of the right thigh of healthy subjects, covering a wide range of morphologies, 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°) of resolution: 1mm×1mm×5mm. We manually segmented each volume in order to evaluate the quality of the segmentation results. We focused our evaluation on clinically relevant muscles of the thigh (13 muscles). In order to compute the empirical mean \bar{x}^s and the empirical variance σ_i^s ², we non-rigidly registered all the volumes and their segmentation map in the




Fig. 2. Box-plot presentation of the Dice coefficients of our segmentation algorithm. (Right) Individual muscles; (Left) All muscles. Average dice values for: registration only: 0.81 ± 0.08 ; RW + shape prior: 0.84 ± 0.08 ; RW + shape prior + c. map: 0.86 ± 0.07 ; automatic seeds ([9]): 0.80 ± 0.19 .

training set to the same target volume. The registration process is achieved using the method presented in [10] and the related registration software (Drop[©], www.mrf-registration.net). We adopted a leave-one-out cross-validation protocol: each test volume is used as the target volume for the registration of the 13 other volumes. Then we computed the estimates of \bar{x}^s and σ_i^s ² on the 13 registered volumes, and perform the segmentation of the test volume. After the registration process, all volumes had the size: $191 \times 178 \times 63$. For solving the linear systems, we used iterative algorithms, such as Bi-conjugate Gradient. Computing the segmentation takes around 5 min on a 2.8 GHz Intel[®] processor with 4 GB of RAM. From a series of tests, we computed the best value for parameter $\lambda^s = 10^{-3}$, except for the Gracilis muscle which we had to constrain more: $\lambda^s = 10^{-1}$.

The quality of the segmentation is measured by computing Dice coefficients with the box-plot presentation¹ (See figure 2). The expression of the Dice coefficient is: $D = 2|T \cap R| / (|T| + |R|)$, where T and R are the pixel sets for the algorithm's output and the ground truth segmentation respectively. We also compared the different methods with p-values obtained using the non-parametric statistical test Wilcoxon rank-sum (cf. scipy.stats).

In figure 2, for comparison with a simple segmentation by atlas registration method, we computed the Dice coefficients of the segmentation which we obtain when retaining the label of maximum probability of the mean probability: $l_i =$

¹ 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).

$\arg \max_s \bar{x}_i^s$. This method yielded inferior results as compared to our method without confidence map with a p-value of 2×10^{-10} . Adding the confidence map slightly improves the segmentation results as compared to not using it (p-value: 6.6×10^{-2}), yielding an average Dice coefficient value of 0.86 ± 0.07 . We also compare our results with a previous method of ours [9]. This method consisted in automatically determining appropriate seed positions with respect to the different muscle classes. The output of this optimization process was then fed to the standard RW algorithm. We obtained inferior results to the method presented here (average Dice: 0.80 ± 0.19 ; p-value when compared to confidence map method: 7.0×10^{-4}).

In figure 3, we show cross-sections of segmentation results obtained with the RW method using the prior model and the confidence map. Segmentation errors tend to affect primarily small muscles (e.g. Gracilis) and muscles located on the extreme upper part of the volumes (e.g Tensor Fasciae Latae) which reveals the limitations of the mean model. These errors are due to the large registration errors on the same muscles. This shows that our model is too constraining, as it does not allow the segmentation to deviate enough from the mean. Due to the few number of training examples, we noted that the variance estimate had little influence on the results: replacing Λ_σ^s with the identity matrix gave us no significantly different results. This suggests that we should add more data to the training set in order to improve the statistical validity of our estimates.




Fig. 3. Segmentation results obtained with the RW algorithm with shape prior and confidence map. Segmentation errors are shown in white.

4 Conclusion

The inherent difficulties of segmenting the skeletal muscles in MR images - namely: partial contours, no discernible texture differences, large variation inter-individuals and unremarkable shapes - render standard segmentation methods

inoperative on this issue. One has to search for methods which perform robustly on inconsistent images cues (partial contours), and flexible models, allowing enough freedom to account for large inter-subject shape variations. We propose a prior model method, resting on the strengths of the Random Walks segmentation algorithm. Due to its robustness when faced with missing contours, the RW algorithm appears to be a good candidate for combination with a trained shape model. We believe to have achieved promising results which demonstrate the potential of our fully automatic approach.

Future work will consist in designing a model which allows more shape variability. We could obtain such model by building a low dimensional space through computing a PCA on the training base. Furthermore, the similarities between the RW algorithm and Markov Random Fields formulations let us envision applying the recent advances in MRF learning to the estimation of the Laplacian matrix.

References

1. Gilles, B., Pai, D.K.: Fast Musculoskeletal Registration Based on Shape Matching. In: Metaxas, D., Axel, L., Fichtinger, G., Székely, G. (eds.) MICCAI 2008, Part II. LNCS, vol. 5242, pp. 822–829. Springer, Heidelberg (2008)
2. Gilles, B., Magnenat-Thalmann, N.: Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations. *Medical Image Analysis* 14, 291–302 (2010)
3. Essafi, S., Langs, G., Paragios, N.: Hierarchical 3D diffusion wavelet shape priors. In: CVPR, pp. 1717–1724. IEEE (September 2009)
4. Wang, C., Teboul, O., Michel, F., Essafi, S., Paragios, N.: 3D Knowledge-Based Segmentation Using Pose-Invariant Higher-Order Graphs. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010, Part III. LNCS, vol. 6363, pp. 189–196. Springer, Heidelberg (2010)
5. Andrews, S., Hamarneh, G., Yazdanpanah, A., HajGhanbari, B., Reid, W.D.: Probabilistic Multi-shape Segmentation of Knee Extensor and Flexor Muscles. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part III. LNCS, vol. 6893, pp. 651–658. Springer, Heidelberg (2011)
6. Grady, L.: Random walks for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28(11), 1768–1783 (2006)
7. Grady, L.: Multilabel Random Walker Image Segmentation Using Prior Models. In: CVPR, vol. 1, pp. 763–770 (2005)
8. Rousson, M., Paragios, N.: Prior knowledge, level set representations & visual grouping. *International Journal of Computer Vision* 76(3), 231–243 (2008)
9. Baudin, P.-Y., Azzabou, N., Carlier, P.G., Paragios, N.: Automatic Skeletal Muscle Segmentation Through Random Walks and Graph-Based Seed Placement. In: ISBI 2012 (in press, 2012)
10. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. *Medical Image Analysis* 12, 731–741 (2008)