Using Texture Features for the Segmentation of Skeletal Muscles of the Thigh in Dixon MRI

Rafael Rodrigues
rafael.rodrigues@ubi.pt
Antonio M. G. Pinheiro
pinheiro@ubi.pt

Instituto de Telecomunicações
Universidade da Beira Interior
Covilhã, Portugal

Abstract

Segmentation of skeletal muscles in Magnetic Resonance Images (MRI) is essential for the study of muscle physiology and the diagnosis of muscular pathologies. However, manual segmentation of large MRI volumes is a time-consuming task. The state-of-the-art on algorithms for muscle segmentation in MRI is still not very extensive and, more recently, uses mostly learning-based solutions, which require large amounts of data. This work proposes an automated segmentation method for Dixon scans of the thigh based on pixel-wise classification of local texture features using AdaBoost. The descriptor includes features from the Histogram of Oriented Gradients (HOG), *Haar* Wavelet filtering, and statistical measures from both the original MRI and the Laplacian of Gaussian (LoG) filtering. An atlas-based approach is then applied to the resulting muscle tissue segmentation to provide individual muscle labeling.

1 Introduction

Muscle segmentation in medical imaging enables quantitative measurements of muscle tissue and fat infiltration, which are crucial to most physiological and/or pathological evaluations. An automated framework would increase the ability to deal with large whole-body Magnetic Resonance Image (MRI) datasets, as manual segmentation of muscles is still a common approach in radiology services. However, achieving a robust automated method for muscle segmentation in MRI is a challenging task.

Earlier automated methods for the segmentation of individual muscles of muscle regions in MRI of the thigh include a graph-based approach using the Random Walks algorithm [4], a probabilistic model using Gabor Features [3], and atlas-based methods [9]. More recently, deep learning-based approaches have been proposed [2, 6, 8, 12], which rely on U-Net models and typically require large amounts of training data.

In this paper, we present a fully automated method for the segmentation of skeletal muscle in Dixon MRI scans of the human thigh that relies on supervised pixel-wise classification using texture-based features. An atlas is then used to assign muscle labels to the classifier result. A more thorough description of the used methods may be found in [10].

2 Materials and Methods

The used dataset included 10 out-of-phase 3pt Dixon MRI volumes, acquired with a 3T Siemens scanner, with spatial resolution of $1 \times 1 \times 5mm$, $TR = 10ms$, $TE_1 = 2.75ms$, $TE_2 = 3.95ms$, $TE_3 = 5.15ms$, and RF flip angle $= 3$. Each volume contained 80 scans of the right thigh (Fig. 1(a)), with an image size of 256x256 pixels. To avoid including slices with too little muscle tissue, slices near the knee and the ankle were discarded *a priori*, resulting in a subset of 40 MRI slices. From this initial selection, five consecutive slices were randomly selected from each volume to simultaneously enable the proposed muscle atlas construction approach and address the memory limitations of the experimental setup.

2.1 Feature extraction and classification

The texture descriptor (Fig. 1) included the Histogram of Oriented Gradients (HOG) [5], computed using the Dalal-Triggs variant with 16x16 blocks and 9 orientations (36-bin histogram), as well as statistical measures from the original grayscale MRI and a filtered counterpart, using a 5x5 Laplacian of Gaussian (LoG) filter, with $\sigma = 1.5$ [1]. The mean, variance, skewness, and kurtosis were computed within non-overlapping 16x16 patches of those two images. The descriptor also included the coefficients from the *LH*, *HL*, and *HH* sub-bands from 3 levels of decomposition using *Haar* wavelets [13], as well as from the first *LL* sub-band.

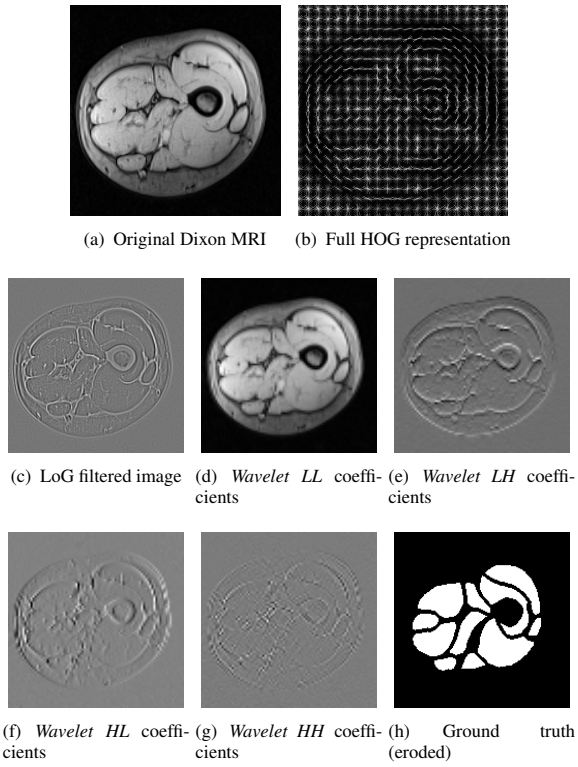

(a) Original Dixon MRI (b) Full HOG representation

(c) LoG filtered image (d) *Wavelet LL* coefficients (e) *Wavelet LH* coefficients

(f) *Wavelet HL* coefficients (g) *Wavelet HH* coefficients (h) Ground truth (eroded)

Figure 1: Out-of-phase Dixon MRI scan (a). Image preprocessing results for texture feature extraction (b) to (g). Muscle ground truth (h)

Image pixels were classified using an AdaBoost classifier [7, 11] for 500 iterations. The ground truth was obtained from the manual segmentation of thigh muscles by a radiologist. For each image, every individual muscle region was eroded with a 2-pixel radius disk (Fig. 1(h)) to avoid including undesired inter-muscular texture patterns into the positive training dataset. For the final label mask used to train the AdaBoost classifier, the majority label inside each 16x16 patch was considered. In the case of a tie, which is more likely to occur in regions near non-muscle tissue, a negative label was chosen to further reduce the inclusion of non-muscle texture into the positive training set. Algorithm validation was performed in a leave-one-out cross-validation, and the training data to classify slices from each volume was obtained from the remaining 9 MRI volumes.

2.2 Muscle labeling

The same images involved in each cross-validation iteration were used to construct the respective muscle atlas. These were registered to a common reference, selected randomly among the training set. First, the bone centroids of both the reference and target MRIs were aligned via translation. The bone was roughly segmented through histogram thresholding. Then, considering $\vec{D}_r$ and $\vec{D}_t$ as the vectors connecting the bone centroid to the most distal point on the convex hull of the muscle labels, in the reference and target images, respectively, all target images were rotated by the angle between $\vec{D}_r$ and $\vec{D}_t$. Finally, the target images were scaled to match the convex hull of the reference image.

An individual muscle probability map was obtained by overlapping the respective binary masks (Figs. 2(a) and (b)). The results were then truncated at 50% of the peak value (Fig. 2(c)). The AdaBoost segmentation output was aligned with the atlas for the respective cross-validation iteration, using the transformations described above. The final result was obtained by reverting all transformations to the target MRI.

(a) *Vastus medialis* (b) *Vastus intermedius* (c) Full muscle atlas

Figure 2: Muscle atlases. In (a) and (b), overlapping maps of two different muscles are shown. In (c), an example of a full atlas is presented.

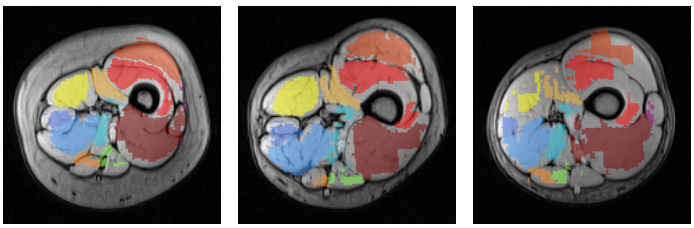

(a) (b) (c)

Figure 4: Examples of individual skeletal muscle segmentation results.

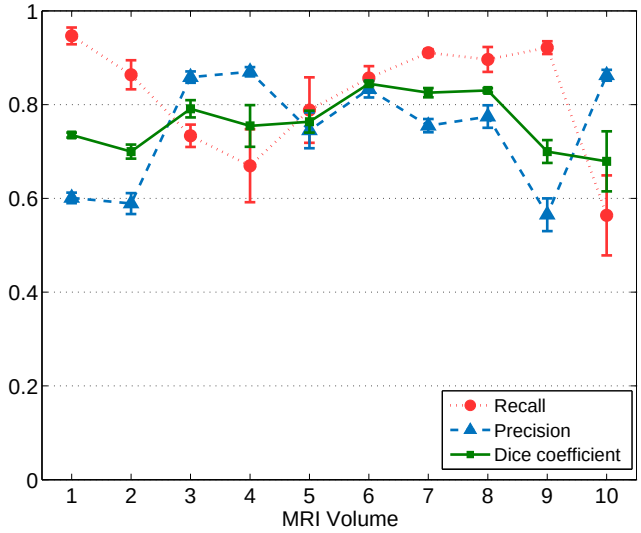

Figure 3: Performance measures per MRI volume.

3 Results and Discussion

The AdaBoost performance in full muscle segmentation was assessed using recall ($\frac{TP}{TP+FN}$) and precision rates ($\frac{TP}{TP+FP}$), and the Dice overlap coefficient ($\frac{2|A \cap B|}{|A|+|B|}$), where TP refers to pixels correctly classified as muscle, whereas FN and FP refer to pixels incorrectly classified as non-muscle and muscle. A and B refer to the segmentation output and ground truth. Fig. 3 shows the mean and standard deviation of these performance measures taken over a 10-fold cross-validation, with recall, precision, and dice overlap coefficients of 0.8150, 0.7454, and 0.7623, respectively. It should be noted that the AdaBoost algorithm was trained with a *downsampled* ground truth mask to address the descriptor resolution, but the final performance was measured considering the original ground truth resolution.

From the total 50 AdaBoost results obtained in cross-validation, the majority presented a correct identification of the muscle region. Results with high recall rates tend to also increase the number of false positives (lower precision). This may be seen in Fig. 4(a), where the inclusion of non-muscle tissue into the AdaBoost output led to a shrinkage of the atlas because of the used registration algorithm. In these cases, the atlas transformation is biased and the labels are slightly offset, even though the AdaBoost segmentation recall rate is high. Using more features could improve the discriminative properties of the proposed classification scheme. In other cases (e.g., Fig.4(b)), given the tradeoff between recall and precision, muscle labels are placed correctly more often, despite the total segmented muscle area being affected. Fig. 4(c) shows the worst overall result (from volume 10 in Fig. 3). Muscles with a smaller cross-sectional area, such as the *Sartorius* or the *Rectus femoris*, tend to not appear on the final segmentation if true positives decay.

The obtained results suggest that the proposed method may perform well in a variety of different scenarios, particularly in cases with more inter-muscular tissue or near articulations, where the muscle region of interest is more dispersed. Nonetheless, it should be noted that a suitable selection of reference MRI for the atlas construction is critical for robustness, given the variations in muscle geometry and relative positioning, which are quite stable in consecutive slices.

4 Conclusions

The results of this research show that texture features may contribute to a global and automated skeletal muscle segmentation tool, which would greatly improve the state-of-the-art in analysis and diagnosis based on muscle MRI. One key advantage of this model is that it requires far less training data than deep-learning approaches, making it easily adaptable to different MRI systems and acquisition protocols.

Acknowledgment

The authors would like to acknowledge Dr. Pierre Carlier, at the NMR Laboratory of the Institut de Myologie, Paris, for providing the MRI database. This work was funded by FCT/MCTES (UIDB/50008/2020 and SFRH/BD/130858/2017).

References

[1] S. Agaian and A. Almuntashri. Noise-resilient edge detection algorithm for brain MRI images. In *Proc EMBC IEEE*, pages 3689–3692, 2009.

[2] A. Agosti et al. Deep learning for automatic segmentation of thigh and leg muscles. *Magn Reson Mater Phys Biol Med*, pages 1–17, 2022.

[3] S. Andrews and G. Hamarneh. The generalized log-ratio transformation: learning shape and adjacency priors for simultaneous thigh muscle segmentation. *IEEE Trans Med Imaging*, 34(9):1773–1787, 2015.

[4] P.-Y. Baudin, N. Azzabou, P. G. Carlier, and N. Paragios. Prior knowledge, random walks and human skeletal muscle segmentation. In *MICCAI*, pages 569–576. Springer, 2012.

[5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proc CVPR IEEE*, pages 886–893, 2005.

[6] J. Ding, P. Cao, H.-C. Chang, Y. Gao, S. H. S. Chan, and V. Vardhanabhuti. Deep learning-based thigh muscle segmentation for reproducible fat fraction quantification using fat-water decomposition MRI. *Insights into Imaging*, 11(1):1–11, 2020.

[7] Y. Freund et al. Experiments with a new boosting algorithm. In *ICML*, volume 96, pages 148–156, 1996.

[8] S. Gaj et al. Deep learning-based automatic pipeline for quantitative assessment of thigh muscle morphology and fatty infiltration. *Magn Reson Med*, 2023.

[9] S. Mesbah et al. Novel stochastic framework for automatic segmentation of human thigh MRI volumes and its applications in spinal cord injured individuals. *PloS One*, 14(5):e0216487, 2019.

[10] R. Rodrigues and A. M. G. Pinheiro. Segmentation of skeletal muscle in thigh dixon MRI based on texture analysis. *arXiv preprint arXiv:1904.04747*, 2019.

[11] P. Viola, M. J. Jones, and D. Snow. Detecting pedestrians using patterns of motion and appearance. *Int J Comput Vis*, 63(2):153–161, 2005.

[12] C. K. Wong et al. Training CNN classifiers for semantic segmentation using partially annotated images: with application on human thigh and calf MRI. *arXiv preprint arXiv:2008.07030*, 2020.

[13] Y. Zhang, Z. Dong, L. Wu, and S. Wang. A hybrid method for MRI brain image classification. *Expert Syst Appl*, 38(8):10049–10053, 2011.