MR Image Segmentation Using Phase Information and a Novel Multiscale Scheme

Pierrick Bourgeat[1], Jurgen Fripp[1,3], Peter Stanwell[2], Saadallah Ramadan[2], and Sébastien Ourselin[1]

[1] BioMedIA Lab, Autonomous Systems Laboratory, CSIRO ICT Centre, Sydney, Australia
{pierrick.bourgeat, jurgen.fripp, sebastien.ourselin}@csiro.au
[2] Institute for Magnetic Resonance Research, Sydney, Australia
{stan, saad}@imrr.usyd.edu.au
[3] School of ITEE, University of Queensland, Australia

Abstract. This paper considers the problem of automatic classification of textured tissues in 3D MRI. More specifically, it aims at validating the use of features extracted from the phase of the MR signal to improve texture discrimination in bone segmentation. This extra information provides better segmentation, compared to using magnitude only features. We also present a novel multiscale scheme to improve the speed of pixel based classification algorithm, such as support vector machines. This algorithm dramatically increases the speed of the segmentation process by an order of magnitude through a reduction of the number of pixels that needs to be classified in the image.

1 Introduction

Segmentation of textured organs is a difficult problem, with the most successful approaches relying on the combination of shape and texture information [1,2,3]. However, texture features do not always provide a sufficiently good discrimination to allow an accurate final segmentation [4]. This work aims at validating the use of features extracted from the phase of the MR image to improve texture discrimination in bone segmentation, which has broad clinical applications in diagnosis, detection of changes in longitudinal studies and surgical planning.

For any type of MR image acquisition, magnitude is only one part of the MR signal, the latter being a complex signal acquired in k-space, with phase and magnitude components as shown Fig. 1. For conventional anatomical imaging methods, the phase information is non-coherent, with dephasing caused by chemical shift variation and local magnetic susceptibility. The latter effect is due to differing magnetic susceptibilities within the body and/or instrumental imperfections. Some imaging sequences are more sensitive to this effect than others, and with a properly chosen sequence, phase can be used to provide information about tissue interfaces, not available in the magnitude of the signal. Using phase to improve tissue contrast on magnitude images has prompted a growing interest in the MRI community, as highlighted by recent work from Haacke and Sehgal on susceptibility-weighted imaging [5,6,7].

R. Larsen, M. Nielsen, and J. Sporring (Eds.): MICCAI 2006, LNCS 4191, pp. 920–927, 2006.

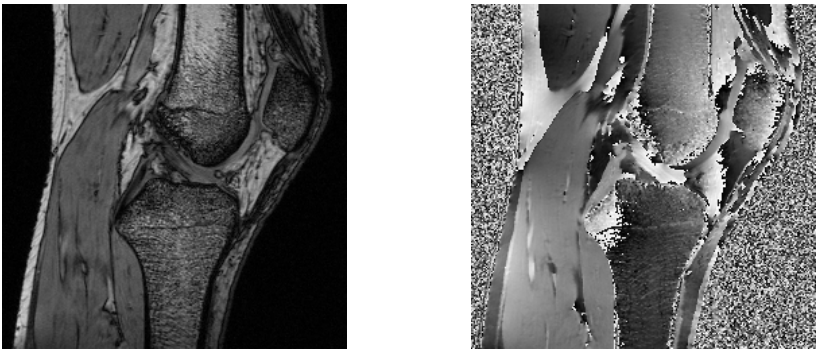

Fig. 1. Magnitude and phase image with $T_E = 8.6$ms

For the remainder of the paper, we first describe the MR acquisition protocol, followed by the feature extraction method that enables features to be extracted from the phase without phase unwrapping, and we present a novel multiscale segmentation scheme designed to speed-up the pixel-wise classification step. In the final section, we present results on segmentation accuracy using phase information, and the performances of the multiscale scheme in terms of speed and accuracy.

2 MR Acquisition

The images used in this paper were acquired on a whole-body 3T clinical scanner (Magnetom Trio, Siemens AG, Germany) using the manufacturer's transmit-receive quadrature extremity coil. Raw image data were acquired as a 3D-volume with the use of a simple gradient-echo sequence (FLASH). Raw data was then processed to produce separate phase and magnitude images. Acquisition parameters were chosen to maximise signal-to-noise ratio while enhancing the inherent phase contrast of the knee joint. The following acquisition parameters were used : echo times (T_E) - 4.9 & 8.6 ms, repetition time - 28 ms, flip angle - 15°, FOV 150 mm, matrix = 256 x 256, 1.5 mm slice thickness, and 64 partitions. The knees of 18 healthy volunteers were scanned, producing a set of 2 complex images per knee, corresponding to the 2 echo times. As presented in Fig. 1, the wrapped phase image shows strong textural information in the bone and the background, compared to the relatively smooth areas in the cartilage and the muscles. In areas of low intensity, such as the background, the signal does not contain enough information to produce an accurate measure of the phase.

3 Features Extraction

In Reyes-Aldasoro approach [8], a subband filtering is applied in the original k-space to extract texture features from the complex image, corresponding to a region of the spatial and frequency domain. The amplitude of the filter response

measures the signal energy within the selected region, and is strongly dependent on the local amplitude of the signal. Therefore, in the areas of low amplitude, the phase information will not have any effect on the filtering. However, if the complex image is normalized to produce a complex phase image of constant unity magnitude, as we proposed in [9], the filter response will be directly related to the phase content of the image. Therefore, the same filters can be applied on the magnitude and the normalized phase data. The normalization is performed by dividing the complex image $I(x, y, z)$ by the amplitude $A(x, y, z)$, to generate a complex phase image $I_\varphi(x, y, z)$ of constant unity amplitude, and therefore only composed of phase information :

$$I_\varphi(x, y, z) = \frac{I(x, y, z)}{A(x, y, z)} = e^{j\varphi(x,y,z)}. \tag{1}$$

$I_\varphi(x, y, z)$ is a complex image that can be Fourier transformed and then filtered in order to extract phase information. As phase wraps appears when the phase is mapped from the complex space to the real space, keeping the phase information in the complex domain removes the need of phase unwrapping. Variations in the phase induce variations in the frequency of the signal, and therefore, a frequency analysis performed using Gabor filters on $I_\varphi(x, y, z)$ can extract textural information about the phase.

Our subband filtering is similar to that of Zhan and Shen [2], with a bank of non symmetric Gabor filters applied in the coronal and sagittal plane. Because of the anisotropic nature of the images, the Gabor filter bank must be designed in the real space instead of the pixel space. Therefore, a bank of filters with 6 orientations and 5 scales is used in the sagittal plane, and only 3 scales in the coronal plane to accommodate the poor high frequency content in that plane. A 3D Gaussian filter is then applied on the magnitude of the output of each filter to smooth the response across slices. As it is important to preserve rotation invariance, the magnitude of the output of the filters is summed across all orientations in order to obtain a set of features that is rotation invariant. Fig. 2 presents the features obtained in the sagittal plane from the magnitude and the phase images with $T_E = 8.6$ ms. The feature images illustrate the discrimination power of the phase between bones and surrounding tissues.

4 Contour Refining Segmentation

The extracted features are classified using the support vector machines (SVM) classifier [10] in a pixel-wise fashion. The implementation relies on SVMLIB [11]. We use an RBF kernel, with the parameters (C, γ) optimised using a five fold cross validation. SVM are among the most effective methods for classification, but they usually require a large number of support vectors to address complicated non-linear problems, resulting in a high computational cost. Training methods have been designed to improve their efficiency through a smaller number of support vectors used [12]. However, they usually require to sacrifice some of the accuracy in order to significantly increase the decision speed. In our

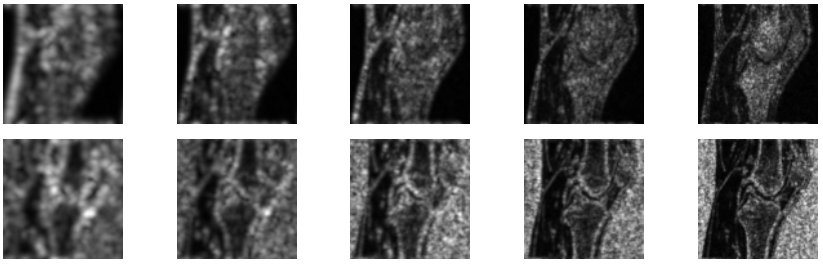

Fig. 2. Magnitude features (top row) and phase features (bottom row). The features represent a 5 level decomposition, from low to high frequencies (presented from left to right)

approach, we consider the classification problem in terms of image processing, where bones are large compact objects, and not all of the pixels need to be classified to get a good representation of their shape. Therefore, we propose a novel *coarse-to-fine* scheme that recursively refines the boundary of the objects, with a reduced number of pixels classified inside/outside of the object.

The standard multiscale approach is aimed at identifying homogeneous regions on the coarser scale, and recursively refining this estimation using finer scales [13,14]. However, these techniques are designed to improve segmentation results while reducing sensitivity to noise, but do not typically target the speed of the segmentation process. Our implementation is specifically designed to reduce the number of pixels being classified by the SVM, and improve the segmentation speed. We use a *coarse-to-fine* approach to get a representation of the objects at the lowest resolution, and use this representation to identify the pixels lying on the contour. This contour is recursively refined until the highest resolution is obtained, thus only classifying the pixels lying on this contour.

Consider a set of feature images $\mathbf{F}(x, y, z)$, a coarse segmentation of the bones can be obtained by subsampling $\mathbf{F}(x, y, z)$ by a scaling factor 2^k ($k \in \mathbb{N}$), such that $\mathbf{F}_k = \mathbf{F} \downarrow 2^k$, and classify $\mathbf{F}_k$ through SVM to produce a segmented mask $S_k = SVM(\mathbf{F}_k)$.

S_k is then upsampled by a factor 2 to produce an upsampled mask U_{k-1} such that $U_{k-1} = S_k \uparrow 2$. The upsampled mask U_{k-1} is then eroded and dilated using a cubic structuring element A of 1 pixel radius to produce two masks, respectively $E_{k-1} = U_{k-1} \ominus A$ and $D_{k-1} = U_{k-1} \oplus A$, and obtain a contour mask $C_{k-1} = D_{k-1} - E_{k-1}$. The contour mask C_{k-1} identifies the pixels where the real boundary of the object is most likely to be located, and therefore needs to be refined.

The feature images $\mathbf{F}(x, y, z)$ are then subsampled by a factor 2^{k-1} to produce $\mathbf{F}_{k-1}$ and match the scale of C_{k-1}. The segmented image S_{k-1} corresponding to the scale $k-1$ is obtained as a combination of the eroded image E_{k-1}, with the newly classified pixels of $\mathbf{F}_{k-1}$ lying within C_{k-1}, such that $S_{k-1} = SVM(\mathbf{F}_{k-1} \cap C_{k-1}) \cup E_{k-1}$.

This process is repeated k times until the full resolution is reached. Fig. 3 illustrates the process with $k = 4$, showing the intermediate segmentations at each scale Fig. 3(a-e), and the corresponding contour masks Fig. 3(k-n) generated

through erosion and dilation of these segmented images. This approach reduces dramatically the number of pixels that need to be classified to obtain the full resolution segmented image. Moreover, the number of pixels classified is further reduced by keeping a list of the pixels already classified at previous scales, and prevent reclassifying the same pixel several time. A diagram of the complete scheme is presented in Fig. 4. This type of approach could be easily adapted to a multiclass problem by computing the contour masks for each class, and merge these masks into a single mask containing all the pixels that need to be refined for all classes.

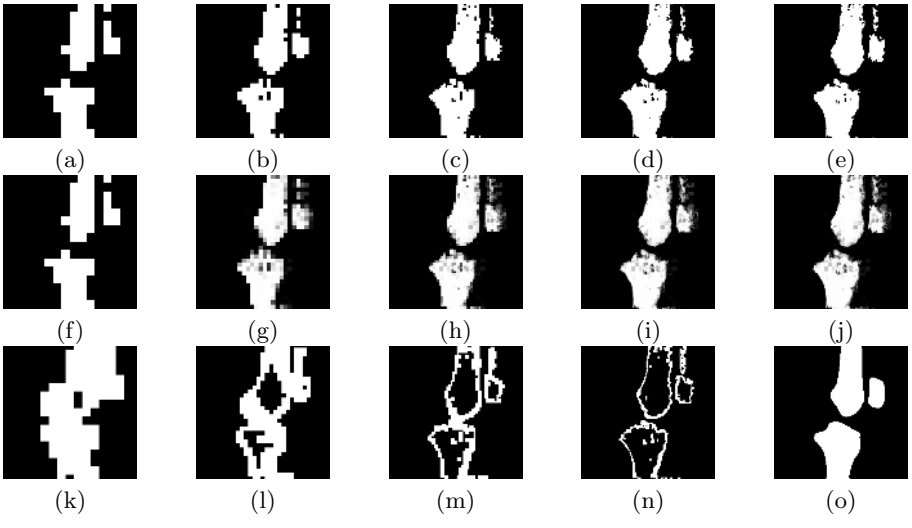

Fig. 3. Segmented images with $k = 4$ presented from scale 4 to 0 (a-e), along with the corresponding probability images (f-j), the contour masks showing where the contour is refined at each scale (k-n), and the manual segmentation (o). The images are all scaled at the same size for clarity.

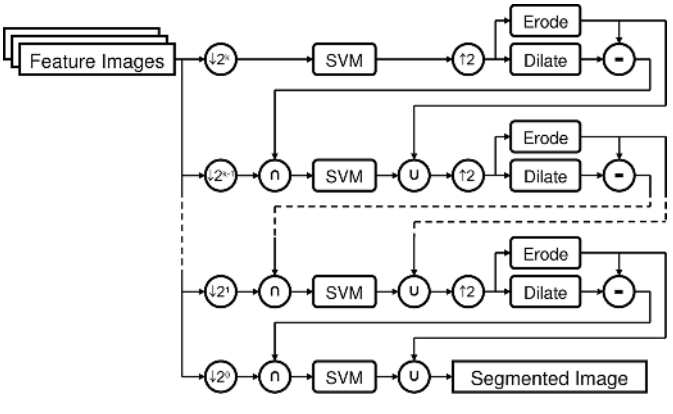

Fig. 4. Diagram of the multiscale scheme

The SVM can be trained to output a probability estimate, along with the hard segmentation. Our *coarse-to-fine* approach can also be used to generate probability maps where the initial probability is set by the segmentation at the coarse scale, and the probability estimate on the contour of the object is then recursively refined within the contour masks C_k. An example of the probability maps at each scale for $k = 4$ is presented Fig. 3(f-j). These probability maps can then be used as appearance model to drive an active shape model (ASM) with the shape constraints used to generate a more anatomically correct representation of the bones, as presented in [1] and [2].

5 Experimental Results and Discussion

From the database of 18 knees, where the bones have been manually segmented, 4 knees are used for training of the classifier with 10.000 points extracted from each image, and 14 knees are used for testing.

For the first experiment validating the usefulness of the phase in bone segmentation, the classifier was trained with 3 different sets of features (Gabor filtering applied on magnitude only, phase only, or magnitude and phase) using either the first echo only, the second echo only, or both echos.

Table 1. Mean (standard deviation) of the DSC on the bone segmentation over the 14 testing datasets

	T_{E1}	T_{E2}	$T_{E1}T_{E2}$
Magnitude	0.77 (0.05)	0.80 (0.04)	0.88 (0.03)
Phase	0.67 (0.08)	0.73 (0.05)	0.83 (0.05)
Magnitude and Phase	0.83 (0.04)	0.85 (0.03)	0.90 (0.02)

Since we are looking at segmenting 4 large bones, a size filter is applied to the segmented image to preserve the 4 largest connected components, and remove small misclassified volumes. The mean and standard deviation of the Dice similarity coefficient (DSC)[1] after filtering are presented in Table 1 for the 14 test images. The combination of phase and magnitude information improves the results by increasing the mean DSC, and reducing the standard deviation. This leads to more accurate results with less variation across the dataset. This is an important result as the phase is always acquired during any type of acquisition and can be readily included in segmentation algorithms. Also, in the case of an acquisition with 2 echos, including the 2 echos dramatically improves the results (+0.08 on the magnitude, +0.1 on the phase and +0.05 on the magnitude and phase), to produce an excellent DSC of 0.9 with the phase-magnitude combination.

For the second set of experiments, the multiscale approach is compared with regards to the number of scales k used ($k = 0$ being the full resolution) in terms

[1] DSC $= \frac{2(A \cap B)}{|A|+|B|}$ where A and B are two binary objects.

of percentage of pixels classified (Fig. 5(a)), and DSC of the segmented images (Fig. 5(b)) keeping only the 4 largest connected components. The experiments are conducted over the 14 testing dataset using the two echos, and the three sets of features. Fig. 5(a) shows the exponential decrease in the number of pixels classified when k is increased. The improvement is even more pronounced when the segmentation is more accurate, as it creates large areas with fewer contours. Interestingly, the DSC is also improved when k is increased as presented Fig. 5(b), to reach a peak for $k = 3$. This corresponds to an optimum where small misclassified volumes inside the bone will be correctly classified as they are not detected at the coarser scales, and therefore not refined. However, if k is increased above 3, small objects such as the thin fibula will not be detected at the coarser level, and will be missed out in the final segmentation. Therefore, k must be set according to the size of the smallest object to be segmented. With $k = 3$, only 8% of the pixels of the image are classified using magnitude, or combined phase and magnitude features. As a result, a classification that would take over 1 hour on 3.2GHz desktop computer if all the pixels were classified, can be performed in 5 minutes (the extra cost of the contour extraction accounting for only 10 seconds of this time).

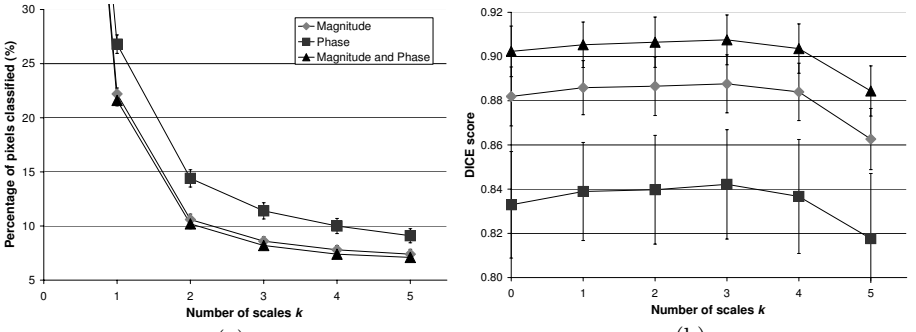

(a) (b)

Fig. 5. Performance results of the multiscale scheme for different k values in terms of speed (a), and accuracy (b)

6 Conclusion

Phase can be easily integrated in a feature extraction scheme to provide superior tissue discrimination compared to magnitude only based features. Phase information is always acquired and its incorporation into current image analysis techniques provides an exciting new field of research. In the case of bone segmentation, the use of phase improves the accuracy and robustness of the segmentation, resulting in a higher DSC and smaller standard deviation across the 14 testing dataset. The multiscale scheme dramatically improves the speed of the segmentation, speed being the most prominent drawback of SVM-based image segmentation algorithm. The segmented images can be used as an automatic initialisation of ASM models through registration with an atlas, or can be

directly used via the probability estimate to drive the ASM deformation. Future work will look into other approaches to extract phase information without phase unwrapping, and study other organs where phase can assist segmentation and analysis.

References

1. van Ginneken, B., Frangi, A., Staal, J., Romeny, B., Viergever, M.: Active shape model segmentation with optimal features. IEEE Trans Med Imaging **21**(8) (2002) 924–933
2. Zhan, Y., Shen, D.: Automated segmentation of 3D US prostate images using statistical texture-based matching method. In: MICCAI'03. Volume 2878 of LNCS., Montreal, Canada, Springer Verlag (2003) 688–696
3. Xie, J., Jiang, Y., t. Tsui, H.: Segmentation of kidney from ultrasound images based on texture and shape priors. IEEE Trans Med Imaging **24**(1) (2005) 45–57
4. Lorigo, L., Faugeras, O., Grimson, W., Keriven, R., Kikinis, R.: Segmentation of bone in clinical knee MRI using texture-based geodesic active contours. In: MICCAI'98. Volume 1496 of LNCS., Cambridge, MA (1998) 1195–1204
5. Haacke, E., Xu, Y., Cheng, Y.C., Reichenbach, J.: Susceptibility weighted imaging (SWI). Magn Reson Med **52**(3) (2004) 612–618
6. Haacke, E., Xu, Y.: The role of voxel aspect ratio in determining apparent vascular phase behavior in susceptibility weighted imaging. Magn Reson Med **24**(2) (2006) 155–160
7. Sehgal, V., Delproposto, Z., Haacke, E., Tong, K., Wycliffe, N., Kido, D., Xu, Y., Neelavalli, J., Haddar, D., Reichenbach, J.: Clinical applications of neuroimaging with susceptibility-weighted imaging. J Magn Reson Imaging **22**(4) (2005) 439–450
8. Reyes-Aldasoro, C.C., Bhalerao, A.: Volumetric texture description and discriminant feature selection for MRI. In: IPMI'03. Volume 18., Ambleside, UK, Springer (2003) 282–293
9. Bourgeat, P., Fripp, J., Janke, A., Galloway, G., Crozier, S., Ourselin, S.: The use of unwrapped phase in MR image segmentation : a preliminary study. In: MICCAI'05. Volume 3750 of LNCS., Palm Springs, CA, Springer Verlag (2005) 813–820
10. Vapnik, V.: The nature of statistical learning theory. Springer Verlag, New York (1995)
11. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines (2001) Software available at http://www.csie.ntu.edu.tw/∼cjlin/libsvm.
12. Zhan, Y., Shen, D.: An adaptive error penalization method for training an efficient and generalized SVM. Pattern Recognition **39**(3) (2006) 342–350
13. Liang, K.H., Tjahjadi, T.: Adaptive scale fixing for multiscale texture segmentation. IEEE Trans Image Process **15**(1) (2006) 249–256
14. Choi, H., Baraniuk, R.G.: Multiscale image segmentation using wavelet-domain hidden Markov models. IEEE Trans Image Process **10**(9) (2001) 1309–1321