Detection and Registration of Ribs in MRI Using Geometric and Appearance Models

Golnoosh Samei*, Gábor Székely, and Christine Tanner

Computer Vision Laboratory, ETH Zurich, 8092 Zurich, Switzerland

Abstract. Magnetic resonance guided high intensity focused ultrasound (MRgHIFU) is a new type of minimally invasive therapy for treating malignant liver tissues. Since the ribs on the beam path can compromise an effective therapy, detecting them and tracking their motion on MR images is of great importance. However, due to poor magnetic signal emission of bones, ribs cannot be entirely observed in MR. In the proposed method, we take advantage of the accuracy of CT in imaging the ribs to build a geometric ribcage model and combine it with an appearance model of the neighbouring structures of ribs in MR to reconstruct realistic centerlines in MRIs. We have improved our previous method by using a more sophisticated appearance model, a more flexible ribcage model, and a more effective optimization strategy. We decreased the mean error to 2.5 mm, making the method suitable for clinical application. Finally, we propose a rib registration method which conserves the shape and length of ribs, and imposes realistic constraints on their motions, achieving 2.7 mm mean accuracy.

1 Introduction

MRgHIFU is a relatively new and promising therapy, which locally ablates the tissues by focusing ultrasound beams on the target region. This form of therapy has already been successfully applied to static organs such as brain tumours and uterine fibroids. Yet, its applicability to organs in the abdomen, such as the liver is limited due to the motion caused by respiration.

There has been extensive work on modeling the respiratory motion of abdominal and thoracic organs e.g. [12,8]. Yet, for an effective therapy, knowledge of the accurate position of the organs on the beam path, as well as the position of the target, is required. In particular, the ribcage poses an important challenge, as it encloses parts of the liver and the absorption and reflections of ultrasound at bones can cause overheating at the ribs and their surrounding [7,14].

Currently, no imaging device is fast enough to allow real-time acquisition and quantification of the 3D motion of the ribcage and the target organ. Hence, partial observations need to be complemented by prior knowledge about the expected 3D motion. MR is the available image modality in MRgHIFU and also has the advantage that long acquisitions can be performed without any known

* We acknowledge the grant n^o 270186 from the EU's Seventh Framework Programme.

P. Golland et al. (Eds.): MICCAI 2014, Part I, LNCS 8673, pp. 706–713, 2014.

adverse effects on the volunteers. Therefore, to quantify the respiratory motion of the ribcage and its correlation with that of the liver in 4D MRIs, it is necessary to detect and register the ribs in this modality.

While hardly investigated in MRI, detecting ribs in CT has been extensively explored, e.g. [6,11]. Staal et al. [11], find local CT image structures, by means of ridge detection, and classify these into ribs and non-ribs. Eventually, full ribs are formed by grouping these rib structures. The general steps of this method have been applied to MR images [9]. However, the results from 5 volunteers on high resolution MRIs show quite noisy extractions and confusion between ribs.

Recently, we proposed a method [10] where the idea of combined use of a geometric ribcage model from CT and an MR appearance model was presented. Being built on a relatively small dataset, the models could not capture the variation in the population adequately, which led to high mean errors relative to the intercostal distances. In this work, we improve our previous method by using more sophisticated appearance features, reducing the redundancy between the different attribute models, increasing the flexibility of the ribcage model in adapting to a new subject, and enhancing the strategy for finding optimal ribcage parameters. These improvements reduced the errors from 8.1 mm to 2.5 mm, making the rib detection method suitable for clinical application.

In addition to rib detection, we require an MR rib registration method to study the ribs motion. Initial attempts with an elastic image registration method lead to unrealistic rib deformations. Avoiding these was another motivation for creating the geometric ribcage model. In this work we propose the incorporation of the ribcage model into a rib registration method, such that shape and length of the ribs are conserved and anatomically plausible results are achieved.

2 Materials

We used two abdominal image datasets. The **CT-Dataset** consists of high resolution CT images of end-inspiration (EI) breath-holds (BH). The images were acquired from 20 healthy volunteers, with a resolution of $1.37 \times 1.37 \times 1\,\mathrm{mm}^3$, in the anterior-posterior (AP), left-right (LR), and inferior-superior (IS) direction, respectively. The images were acquired and used in accordance with the local institutional review board regulations and were made available to us by the authors of [3]. The **MR-Dataset** consists of free-breathing 4D MRIs [13] obtained from 10 healthy volunteers ($1.33 \times 1.33 \times 5\,\mathrm{mm}^3$ in AP, IS, LR direction). Additionally, we have 5 BH 3D MR images in EI, and end-expiration (EE) with half (2.5 mm) slice-thickness. We applied our rib detection method to the EE images, as they show less artifacts in the reconstructed 4D MRIs.

3 Method

We built a ribcage model for ribs enclosing the liver, namely the right ribs number 7 to 10. Each rib was defined by its shape, length and orientation. The attributes of each rib are highly correlated with each other and with those of

other ribs. By building a statistical population model of these attributes, we extracted the main modes of variation of the ribcage geometry. Moreover, this model allowed us to generate anatomically plausible ribs using a few parameters. For detection all of these attributes had to be determined while for registration only the rotation angles were optimized to ensure shape and length preservation.

3.1 Ribcage Model

We built our ribcage model from CT data similar to [10]. First, the ribs were segmented in the CT images using a simple region growing algorithm in the 3D Slicer software [4]. To include full ribs, they were separated from the vertebral column at their joints and not as previously at the tubercle.

The centerline and natural coordinate system (NCS) of each rib were found, and a rotation from the NCS of each rib to the world coordinate system (WCS) was computed [10]. The corresponding 3 Euler angles constitute the rotation parameters. Finally, we established inter-subject correspondence between all ribs by locating the anatomical landmark, *angle point*, as the most posterior point in WCS, and using it to divide the ribs into two segments. Each segment was uniformly subdivided into a fixed number of points (100 points in total).

Shape Model. All ribs were normalized to have equal length and the 3D position vectors of all the 100 points of a rib were concatenated to form a 300D vector. The resulting vectors for ribs number 7 to 10 of all subjects formed a population of 80 ribs in total. By performing principle component analysis (PCA) on these vectors, we captured the high correlation between the elements of these vectors, with the first 2 PCs covering 96% of the variation. Hence, we chose our shape parameters to be the projection of the 300D vector of each rib onto these 2 PCs.

Ribcage Parameters. Each rib has 6 parameters, namely 2 shape parameters, 3 Euler angles and 1 length value. To model the correlation between these parameters for all the ribs, we performed another PCA on 20 instances of the resulting 24D vectors, with 6 PCs capturing 96% of the population variation. This allowed us to generate the centerlines of the ribs 7-10 approximately by 6 parameters $\mathbf{h} = [h_1..h_6]$. However, due to the relatively small training dataset, the population model might not capture the whole variability. Therefore, we reintroduce some local flexibility to the model to further adjust the generated ribcage to the given MR image, see Sec. 3.2 Optimization Strategy.

3.2 Rib Detection

The parameter vector $\mathbf{h}$ defined in Sec. 3.1, is used to represent a ribcage, and for detection we need to determine values for $\mathbf{h}$ such that the reconstructed ribs $r_{\mathbf{h}}$, are located on the rib centerlines in the given MR image. Therefore, we need a method which determines the likelihood of a given image location being on the

centerline of a rib (referred to as *ribness* hereafter), a cost function for a set of generated ribs based on these likelihoods, and a strategy to find the optimal set of ribs according to the cost function.

Appearance Model. For a set of generated rib centerlines, we want to determine their ribness probability given an MRI. Hence, we create an appearance model which computes the ribness probability of a point based on the intensities of its neighbourhood. To that end, we used RF classifiers [1] trained on MRI sagittal patches similar to [10]. However, in [10], feature vectors were based merely on image intensities, which make the model sensitive to intensity range changes. Moreover, slight changes in the structures (such as the size of the rib cross sections), would result in very dissimilar features, necessitating a large training set to learn all the variations. To overcome these shortcoming, we used more robust features namely $\mathcal{F}_{\mathbf{p}} = (F^1_{\mathbf{p}}...F^6_{\mathbf{p}})$, with 6 channels for a point $\mathbf{p} = [x_{\mathbf{p}}, y_{\mathbf{p}}, z_{\mathbf{p}}]$, as in [2]. These channels were comprised of raw intensity values, normalized intensity values, Canny and Soble edge detector responses, outputs of morphological operators (erosion and dilation) on the intensity, and the response of a Gabor filter bank with eight different rotations and four different phase shifts. Each $\mathcal{F}_{\mathbf{p}}$ was constructed from neighbourhood $N^w_{\mathbf{p}}(i,j) = I(x_{\mathbf{p}} + i, y_{\mathbf{p}} + j, z_{\mathbf{p}})$, around point $\mathbf{p}$ with $i,j \in \{-w,..,w\}$ and $I(x,y,z)$ the interpolated MR intensity at position $[x,y,z]$. Furthermore, we assumed that the ribness probability of $\mathbf{p}$ conditioned on patch $N^{15}_{\mathbf{p}}$, is independent of the rest of the image I, and therefore obtained $P(\mathbf{p}|I) \approx P(\mathbf{p}|N^w_{\mathbf{p}})$. Additionally we used RF thresholds on comparison features [2], rather than on the original features as done previously in [10]. We trained our RF on positive features, $\mathcal{F}_{\mathbf{p}_R}$, extracted from points on the manually determined rib centerlines ($\mathbf{p}_R$) and negative features, $\mathcal{F}_{\mathbf{p}_N}$, from randomly selected points ($\mathbf{p}_N$) with a distance of ≥ 10 mm from all $\mathbf{p}_R$. To obtain more discriminative classifiers, we trained 4 RFs for regions along the ribs with different appearances on sagittal slices. The first one (C1) was built for the points that had an AP distance of 5 mm from the angle point. The second one (C2), was trained on the points within an LR distance of 5 mm (slice thickness) from the most lateral rib point. The third and fourth classifiers (C3, C4) were trained on the points between the angle point and the most lateral rib points, and the remaining points, respectively.

Cost Function. Similar to [10], we adapt a Bayesian framework where the cost function, $C(\mathbf{h})$, is based on the conditional negative log-likelihood of the parameters $\mathbf{h}$, given the MR image I. However, the cost function utilized in [10], sharply increases with distance to rib centers, making it very hard to optimize. An ideal cost measure for efficient solving of our optimization problem should be proportional to the distance to rib centers with a smooth gradient. Therefore we employed a cost function that incorporates the ribness evidence from the neighbouring voxels by spatially smoothing the RF log-likelihood with a Gaussian kernel:

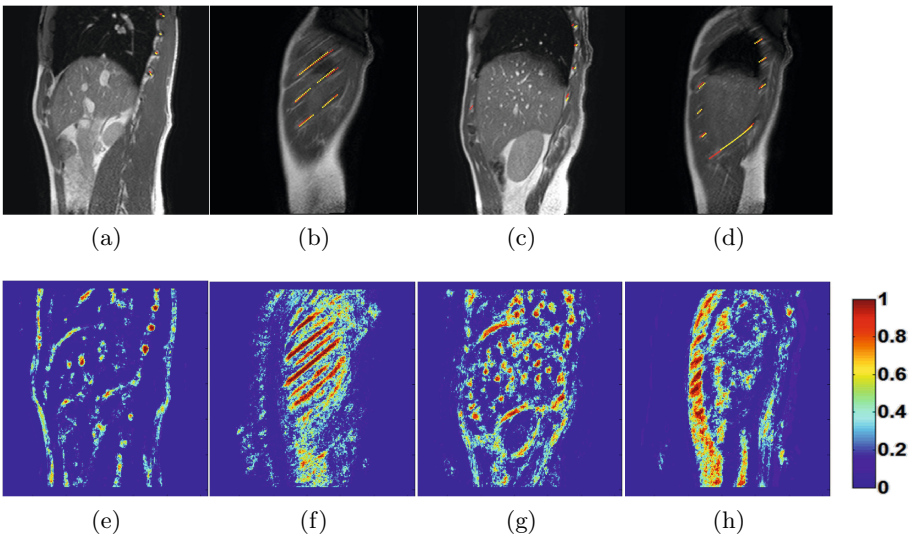

Fig. 1. (a-d) The detected (yellow) and the ground truth centerlines (red) on 4 sagittal slices. (e-h) The probability heat maps for the slices shown in (a-d) obtained from classifiers (e) C1, (f) C2, (g) C3 and (h) C4.

$$C(\mathbf{h}|I,\sigma) = -\sum_{\mathbf{p}\in r_{\mathbf{h}}}\sum_{\mathbf{q}\in S^{3\sigma}_{\mathbf{p}}} log(P(\mathbf{q}|N^{w}_{\mathbf{q}}))\cdot G^{\sigma}(\mathbf{q}-\mathbf{p}), \tag{1}$$

where G^{σ} is a Gaussian kernel with variance σ^2, $S^{3\sigma}_{\mathbf{p}}$ is a spherical neighbourhood around $\mathbf{p}$ with radius 3σ, and $r_{\mathbf{h}}$ is the rib associated with parameters $\mathbf{h}$.

Optimization Strategy. We positioned a generated centerline in the image using the location of a so-called starting point, which is manually selected in the vicinity of the rib's joint to the vertebral disk. It defines the position of one of the first rib points (exactly which one is determined during optimization). Furthermore, we reject any solution which exceeds the bounding box of the ribcage defined as in [10]. Optimization is based on grid search and a multi-resolution approach. First, we draw hypotheses h by uniformly sampling the interior of a 6-D hyper-ellipsoid which covers 95% of the multidimensional Gaussian distribution associated with the PCA model. From these, further hypotheses are created by uniformly changing the proposed orientation for all ribs within $\pm 20^o$ in steps of 5^o, by assigning the starting point to one of the first points up to the angle point. Then we calculate $C(r_{\mathbf{h}}|I,5)$ for all hypotheses in the bounding box and keep the $k = 5$ best results. In the next steps, we divide σ and the rotation range by 3 and modify the rib orientations and the starting point. We repeat

this process until either the change in cost function is below a threshold $\tau = 12$, or a maximum number of 3 iterations has been reached.

3.3 Rib Registration

A global rigid registration of the image volume cannot capture the individually different rotations of the ribs around their joints to the vertebral column. Also, an unconstrained non-rigid registration would introduce too many degrees of freedom. Therefore, we use our ribcage model to track the ribs by exploiting the fact that the shape and length of the ribs do not change. Hence, our detection framework and the optimization scheme of Sec. 3.2 can be simplified to find only the optimal rotation angles of the ribs and fixing all other parameters.

As a faster alternative to computing the RF cost measure, we used registration after the first detection step as follows. For a given rotation of the moving rib, we determined for each point on the fixed (i.e detected) rib the corresponding point on the moving rib. Then we calculate the negative normalized cross correlation (-NCC) between their associated two 2D patches. Finally, we chose the rotation with the lowest mean -NCC value over all rib points.

4 Experiments and Results

We generated our ground truth (GT) rib centerlines from manually selected points as in [10]. Since we used the same MR dataset for building the appearance models as for evaluating our method, we performed leave-one-out experiments. An example result is shown in Fig. 1, where the closeness of GT and detected centerlines can be observed. Moreover, the RF classifier falsely assigned high probabilities to the intercostal space (Fig. 1b,1f) and the liver vessels (Fig. 1c,1g), which indicates the importance of the constraints from the geometric model.

We report three error measures to describe the location accuracy (DistanceError), the accuracy in detecting the whole rib (LengthError), and the clinically relevant error (OutOfPlaneError). DistanceError was defined as the closest Euclidean distance between a point on the detected centerline and the GT. LengthError denotes the difference in length of the detected centerline and the GT. DistanceError and LengthError are 3D errors, while only the 2D error projected onto the ribcage surface matters clinically when deciding which elements of the HIFU transducer should be active for treatments through the ribcage [5]. Hence, we define OutOfPlaneError as the projection of the DistanceError on the z axis (LR direction) of the GT rib's NCS, see Table 1 for detection results summary.

The rib registration algorithm was evaluated on BH images at EE and EI. The LengthError is the same for registration and detection, as the ribs' length do not change in registration. The mean DistanceError is marginally lower for RF (see Table 2), however, NCC requires substantially lower computation time. Note that DistanceError of NCC also includes the detection error, as the registration is initialized with the detected and not the GT centerline.

Table 1. Statistics (mean ± std (95%)) of rib detection errors in mm

Subject	OutOfPlaneError	DistanceError	LengthError
1	0.64 ± 0.51 (1.64)	1.66 ± 0.91 (3.36)	14.70 ± 7.98 (24.66)
2	1.04 ± 0.93 (2.60)	2.27 ± 1.38 (5.05)	13.70 ± 12.03 (29.86)
3	1.03 ± 1.20 (3.08)	2.08 ± 1.23 (3.86)	9.52 ± 8.45 (19.52)
4	0.97 ± 1.19 (3.53)	2.80 ± 1.60 (5.85)	25.38 ± 7.29 (30.91)
5	1.37 ± 1.37 (4.14)	2.31 ± 1.53 (5.36)	23.22 ± 7.74 (31.20)
6	1.50 ± 1.34 (4.55)	3.15 ± 1.66 (6.33)	11.00 ± 9.65 (19.78)
7	1.67 ± 2.13 (6.62)	2.74 ± 2.08 (7.38)	16.49 ± 12.49 (27.93)
8	1.02 ± 1.24 (3.86)	2.22 ± 1.40 (5.18)	7.45 ± 5.29 (13.49)
9	1.34 ± 1.61 (5.14)	2.40 ± 1.65 (5.74)	7.54 ± 9.44 (18.40)
10	1.87 ± 3.16 (8.85)	3.58 ± 3.09 (9.87)	37.88 ± 4.56 (41.76)
11	0.93 ± 0.97 (3.06)	2.65 ± 1.98 (7.30)	21.12 ± 4.77 (26.49)
12	1.37 ± 1.37 (4.15)	2.63 ± 1.81 (5.71)	7.07 ± 5.14 (12.46)
13	0.99 ± 1.00 (3.20)	2.10 ± 1.41 (4.80)	9.95 ± 4.89 (16.44)
14	1.29 ± 1.08 (3.56)	2.64 ± 1.55 (5.66)	7.95 ± 8.58 (20.57)
15	1.03 ± 0.92 (3.04)	2.27 ± 1.01 (4.20)	9.78 ± 7.78 (20.35)
all	1.20 ± 1.48 (4.09)	2.49 ± 1.73 (5.67)	14.87 ± 11.09 (35.09)

Table 2. Statistics (mean ± std (95%)) of rib registration error and motion in mm

Sbj	NCC		RF	Respiratory
	OutOfPlaneError	DistanceError	DistanceError	motion
11	1.04 ± 0.96 (3.41)	4.60 ± 2.20 (7.89)	3.33 ± 2.58 (9.65)	4.06 ± 2.91 (9.85)
12	0.73 ± 0.65 (2.10)	2.14 ± 1.11 (4.23)	2.21 ± 1.08 (4.37)	5.82 ± 4.16 (13.02)
13	0.81 ± 0.77 (2.43)	2.47 ± 1.16 (4.30)	2.84 ± 3.43 (9.61)	9.46 ± 5.31 (17.75)
14	1.65 ± 1.53 (5.17)	2.98 ± 1.32 (5.06)	2.85 ± 1.45 (5.47)	6.06 ± 3.68 (12.54)
15	1.78 ± 1.71 (4.51)	3.54 ± 2.35 (8.62)	2.56 ± 2.01 (6.01)	9.73 ± 5.79 (18.78)
all	1.20 ± 1.28 (3.90)	3.07 ± 1.87 (6.81)	2.72 ± 2.26 (5.90)	6.27 ± 4.94 (16.39)

5 Discussion and Conclusion

In this study we showed that, despite their poor visibility, ribs can be detected in MRI by combining geometric and appearance ribcage models learned from CT and MR respectively. We substantially improved our previous results [10], reducing mean errors from 7.2 mm to 2.4 mm for the shared subset of the two studies (subject 1-10).

We also proposed a constrained non-rigid rib registration method based on our ribcage model, and implemented two registration schemes with clinically suitable errors using the RF and -NCC cost measures.

Future work will be on improving the implementation of our algorithm for faster detection and registration. One approach can be to build a regression RF to predict the distance of points to rib centerlines, which would give $C(\mathbf{h}|I, \sigma)$ directly without the need for convolution.

While the mean errors along the ribs were low (2.5 mm), the anterior part of the detected ribs were on average 14 mm shorter than the ground truth. This was

possibly caused by the fact that due to the difficulty in distinguishing cartilage from the bones in MR, some of the manual rib landmarks may have been placed on the cartilage, whereas the CT geometric model did not include them. In future, cartilage will be included in both models for creating an accurate method for the entire ribcage. Obtaining a dataset of MR and CT of the same subjects, will provide a better ground truth for evaluating our method.

References

1. Breiman, L.: Random forests. Mach. Learn. 45(1), 5 (2001)
2. Dantone, M., Gall, J., Fanelli, G., Van Gool, L.: Real-time facial feature detection using conditional regression forests. In: CVPR, pp. 2578–2585. IEEE (2012)
3. Donner, R., Menze, B., Bischof, H., Langs, G.: Global localization of 3D anatomical structures by pre-filtered Hough forests and discrete optimization. Med. Image. Anal. 17(8), 1304–1314 (2013)
4. Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., et al.: 3D slicer as an image computing platform for the quantitative imaging network. Magn. Reson. Imaging 30(9), 1323–1341 (2012)
5. Gao, J., Volovick, A., Pekelny, Y., Huang, Z., Cochran, S., Melzer, A.: Focusing through the rib cage for MR-guided transcostal FUS. In: AIP Conf. Proc., vol. 1481(1), pp. 94–99 (2012)
6. Lee, J., Reeves, A.P.: Segmentation of individual ribs from low-dose chest CT. In: SPIE Med. Imaging, p. 76243J (2010)
7. Li, F., Gong, X., Hu, K., Li, C., Wang, Z.: Effect of ribs in HIFU beam path on formation of coagulative necrosis in goat liver. In: AIP Conf. Proc., vol. 829(1), pp. 477–480 (2006)
8. McClelland, J.R., Hawkes, D., Schaeffter, T., King, A.: Respiratory motion models: A review. Med. Image. Anal. 17(1), 19–42 (2013)
9. Noorda, Y.H., Bartels, L.W., Viergever, M.A., Pluim, J.P.W.: Rib detection in 3D MRI using dynamic programming based on vesselness and ridgeness. In: Yoshida, H., Warfield, S., Vannier, M.W. (eds.) Abdominal Imaging 2013. LNCS, vol. 8198, pp. 212–220. Springer, Heidelberg (2013)
10. Samei, G., Tanner, C., Székely, G.: Rib detection in MR images using shape priors and appearance models. In: ISBI, pp. 798–801. IEEE (2014)
11. Staal, J., van Ginneken, B., Viergever, M.A.: Automatic rib segmentation and labeling in computed tomography scans using a general framework for detection, recognition and segmentation of objects in volumetric data. Med. Image. Anal. 11(1), 35 (2007)
12. Tanner, C., Boye, D., Samei, G., Szekely, G.: Review on 4D models for organ motion compensation. Crit. Rev. Biomed. Eng. 40(2), 135 (2012)
13. Von Siebenthal, M., Székely, G., Gamper, U., Boesiger, P., Lomax, A., Cattin, P.: 4D MR imaging of respiratory organ motion and its variability. Phys. Med. Biol. 52, 1547 (2007)
14. Wu, F., Wang, Z., Chen, W., Zhu, H., Bai, J., Zou, J., Li, K., Jin, C., Xie, F., Su, H.: Extracorporeal high intensity focused ultrasound ablation in the treatment of patients with large hepatocellular carcinoma. Ann. Surg. Oncol. 11(12), 1061–1069 (2004)