MRI Visualisation by Digitally Reconstructed Radiographs

Antoine Serrurier [*a1], Andrea Bönsch[b], Robert Lau[a], Thomas M. Deserno (né Lehmann)[a]
[a]Department of Medical Informatics, Uniklinik RWTH Aachen, Germany;
[b]Virtual Reality Group, RWTH Aachen University, Germany

ABSTRACT

Visualising volumetric medical images such as computed tomography and magnetic resonance imaging (MRI) on picture archiving and communication systems (PACS) clients is often achieved by image browsing in sagittal, coronal or axial views or three-dimensional (3D) rendering. This latter technique requires fine thresholding for MRI. On the other hand, computing virtual radiograph images, also referred to as digitally reconstructed radiographs (DRR), provides in a single two-dimensional (2D) image a complete overview of the 3D data. It appears therefore as a powerful alternative for MRI visualisation and preview in PACS. This study describes a method to compute DRR from T1-weighted MRI. After segmentation of the background, a histogram distribution analysis is performed and each foreground MRI voxel is labelled as one of three tissues: cortical bone, also known as principal absorber of the X-rays, muscle and fat. An intensity level is attributed to each voxel according to the Hounsfield scale, linearly related to the X-ray attenuation coefficient. Each DRR pixel is computed as the accumulation of the new intensities of the MRI dataset along the corresponding X-ray. The method has been tested on 16 T1-weighted MRI sets. Anterior-posterior and lateral DRR have been computed with reasonable qualities and avoiding any manual tissue segmentations. This proof-of-concept holds for research application for use in clinical PACS.

Keywords: Visualization, Preview, Picture Archiving and Communication Systems (PACS), Digitally Reconstructed Radiograph (DRR), Magnetic Resonance Imaging (MRI), X-ray Propagation, Virtually Physiological Human (VHP)

1. INTRODUCTION

Computed tomography (CT) and magnetic resonance imaging (MRI) constitute two major three-dimensional (3D) imaging modalities largely used in clinical routine. Once recorded, a dataset is composed of a set of parallel cross-sectional two-dimensional (2D) images, which frequently are difficult to visualize, manipulate and interpret. Having an overview of the content of a volumetric set on a picture archiving and communication system (PACS) client is often achieved by browsing the images in sagittal, coronal or axial views and by intensity-based 3D volume rendering. While volume rendering appears efficient for CT datasets, MRI requires finer tuning difficult to use in practise.

Computing virtual radiographs, also referred to as digitally reconstructed radiographs (DRR), provides in a single 2D image an overview of the entire 3D data and appears therefore as a powerful alternative to volume rendering.

Computing DRR from CT is straightforward as the CT intensities follow the Hounsfield unit, linearly related to the X-ray attenuation coefficients of the imaged tissue. On the other side, MRI intensities relate to hydrogen proton densities without any direct link to the attenuation coefficients. Previous attempts to compute DRR from MRI have been reported in the literature for the cranial [1-3] and the pelvis [4-5] regions. For the cranium region, Ramsey *et al.* [1-2] assume that the MRI intensities in a certain fixed range correspond to bone tissues for their protocols, set them to a value related to CT intensities and calculate DRR by ray-casting. For the same region, Yin *et al.* [3] propose a semi-automatic segmentation of the MRI to extract the skin contours and other tissues related to their study (brain tumour), fill these regions with different values related to the CT intensities and calculate DRR by ray-casting. More recently, Chen *et al.* [4] emphasize the possibility for the pelvis region to calculate DRR from MRI with a quality similar to DRR calculated from CT, based on manual segmentation of the bone tissues on the MRI and ray-casting. Finally, again after a manual segmentation of the bones in the pelvis region, Kapanen *et al.* [5] propose a model to estimate the Hounsfield values from the MRI intensities within the segmented bone region and calculate the DRR by ray-casting. All the methods require however prior segmentations of anatomical structures on the images, or assume thresholds values for well-defined regions of interest and recording protocols.

*aserrurier@mi.rwth-aachen.de
[1]The author is now with the Clinic for Phoniatrics, Padaudiology and Communication Disorders, Uniklinik RWTH University Aachen

Medical Imaging 2015: PACS and Imaging Informatics: Next Generation and Innovations,
edited by Tessa S. Cook, Jianguo Zhang, Proc. of SPIE Vol. 9418, 94180I

Our general objective is therefore to develop an automatic method to estimate CT intensities and calculate DRR from MRI. This paper was initially motivated by subject-specific volumetric data visualisation in correspondence to virtually physiological human (VHP) models. In our research-driven project, the objective is to develop a general method for MRI visualization, which is not region-dependant and not requiring anatomical segmentations. The application is in the field of simulating and assisting regional anaesthesia in virtual environments [6,7]. Within this framework, this paper intends to present the proof-of-concept for an automatic intensity-based approach to compute DRR from T1-weighted MRI.

Three types of tissue have distinct X-ray attenuation coefficients and are usually visible on conventional radiographs: cortical bone, muscle and subcutaneous fat and bone marrow. These tissues have also distinct proton densities and appear therefore with distinct intensity on T1-weighted MRI: the cortical bone in dark, the muscles in dark grey and the fat and bone marrow in bright grey. Our specific objective will therefore consist in developing a method to classify the voxels in these three categories and in simulating X-ray propagation and attenuation according to this volume classification (section 2), and then present the results of our proof-of-concept study (section 3).

2. MATERIAL AND METHODS

2.1 Data

In total, 16 T1-weighted MRI series recorded between 2003 and 2008 at Uniklinik RWTH Aachen (Germany) have been collected for the study. The axial or coronal images cover the hip or pelvis region with a resolution ranging from 0.44 to 1.86 mm/px and an inter-slice gap from 1.2 to 7 mm (Tab. 1).

2.2 Voxel classification

Like the cortical bone, air appears also dark on the MRI. Air however has an X-ray attenuation coefficient of almost zero and must therefore be differed from bone. Hence, the first step in our processing chain consists in segmenting the foreground (body) from the background (air) on the MRI according to the following process (Fig. 1): (1) denoising by applying simple thresholding and bilateral filtering [8], (2) detecting edges by applying Canny detection [9] (for axial slices) or thresholding using the 5-level Otsu method [10] and selecting the lower threshold (for coronal slices), (3) dilating, filling hole and eroding first on the whole image and second on each connected component in the image to avoid the fusion between several components like the two legs and (4) removing the smaller connected components.

The voxel classification is then based on the histogram of the foreground voxels (Fig. 2 for one example). Depending on the data, the histogram presents one or two major peaks, the first one corresponding to the dark grey voxels of the muscle tissues. The intensity of this first peak is automatically computed by detecting the highest peak in the histogram and looking for another potential peak until 66% of its Y value. The peak with the lowest intensity (X value) is selected, as pointed by the arrow on Figure 2.

Table 1. Characteristics of the MRI series used in the study.

Number of images	Inter-slice gap (mm)	Image size (px)	Image Resolution (mm/px)	Orientation
48	4.95	512×512	0.88×0.88	coronal
60	4.40	512×512	0.89×0.89	coronal
64	5.50	512×512	0.88×0.88	coronal
70	3.30	480×480	0.80×0.80	axial
48	5.00	480×480	0.79×0.79	axial
48	5.00	480×480	0.79×0.79	axial
100	3.00	480×480	0.94×0.94	axial
100	3.00	480×480	0.94×0.94	axial
100	3.00	480×480	0.94×0.94	axial
100	3.00	480×480	0.94×0.94	axial
100	3.00	480×480	0.94×0.94	axial
100	2.00	512×512	0.88×0.88	axial
140	3.00	1024×1024	0.44×0.44	axial
45	3.00	256×256	1.60×1.60	axial
192	1.20	256×256	1.52×1.52	axial
272	7.00	208×256	1.86×1.86	axial

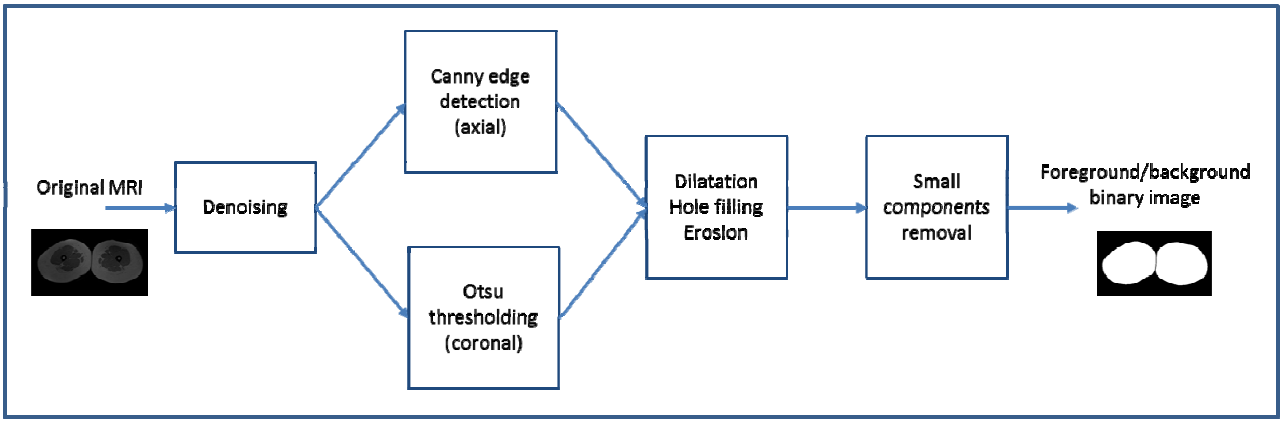

Figure 1. Foreground/background segmentation process on original MRI.

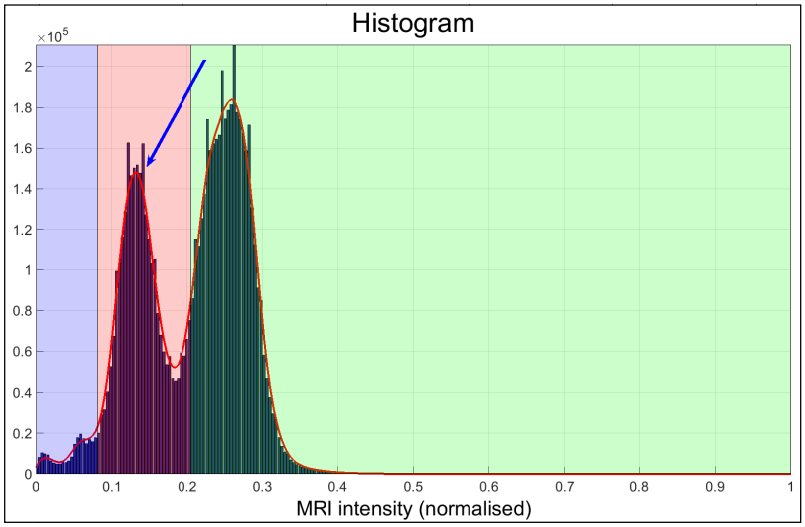

Figure 2. Histogram (bars) and its envelope (red solid line) of the foreground voxels of one MRI set; the blue, red and green zones denote the MRI intensities identified respectively as cortical bone, muscles, and fat. The arrow points on the first major peak, used for the calculations in our study.

The voxel classification is then based on the histogram of the foreground voxels (Fig. 2 for one example). Depending on the data, the histogram presents one or two major peaks, the first one corresponding to the dark grey voxels of the muscle tissues. The intensity of this first peak is automatically computed by detecting the highest peak in the histogram and looking for another potential peak until 66% of its Y value. The peak with the lowest intensity (X value) is selected, as pointed by the arrow on Figure 2.

Although very important for the computation of radiographs due to its high attenuation coefficient, the volume of cortical bone is nevertheless limited, not leading to any particular peak on the histogram. Similarly, the volume of fat did not appear significantly on the histograms of some series. The thresholds between the cortical bone and the muscle voxels and between the muscle and the fat voxels, located on each side of the detected peak, could however be estimated directly from the intensity of this peak. To draw such inferences, the thresholds have been manually identified by an expert for the 16 MRI sets and the intensities values of these thresholds *vs.* the intensity of the peak have been plotted (Fig. 3); a correlation of 0.97 has been observed in both cases.

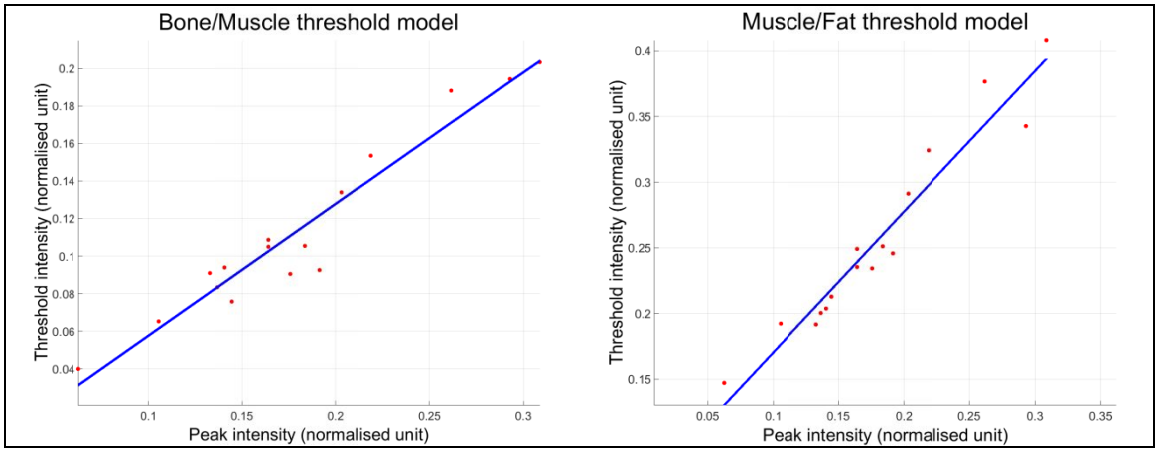

Figure 3. Values of the bone-muscle intensity threshold (left) and muscle-fat intensity threshold (right) vs. intensity of the first major histogram peak (points) and regression lines (solid line).

Linear regressions, also plotted on Figure 3, performed on these points provide prediction models to estimate the bone-muscle and muscle-fat thresholds from the first peak intensity according to equations (1) and (2):

$$Th_{bm} = 0.7015 \times I_p - 0.01246 \quad (1)$$

$$Th_{mf} = 1.072 \times I_p + 0.06318 \quad (2)$$

where Th_{bm} and Th_{mf} represent the bone-muscle and muscle-fat thresholds and I_p the intensity of the first major peak.

Using these linear models, the bone-muscle and muscle-fat thresholds could be respectively estimated with root mean square (RMS) errors of 0.01 and 0.02 on the normalised MRI intensity (Tab. 2).

The Hounsfield scale, used for the CT intensity, is linearly related to the attenuation coefficient. Each voxel intensity can therefore be changed to the normalised Hounsfield value corresponding to its label. In practise however, the relative intensity can be adjusted depending on the tissue to be highlighted on the DRR. Typically, the intensity for the bone, muscle and fat voxels, thereafter I_b, I_m and I_f, have been set respectively to 1, 0.1 and 0.05 in our simulations. The remaining voxels (*i.e.* background), have been set to 0.

From the voxel block filled with desired intensities, DRR can finally be simulated for any orientation using a ray-caszing technique. For each DRR pixel, the intensity I is computed as:

$$I = T_f \times I_f + T_m \times I_m + T_b \times I_b \quad (3)$$

where T_f, T_m and T_b represent the thickness of the fat, muscle and bone tissues crossed by the ray.

2.3 Proof-of-concept evaluation

The method has been evaluated on the 16 MRI sets described in Table 1. As these sets have been already used to design the models in equations (1) and (2), evaluation has been made through a leave-one-out process: after discarding one series, new equations (1) and (2) have been computed on the 15 remaining series as presented in the previous section. The bone-muscle and muscle-fat thresholds of the discarded series have then been determined from these equations. The process has been repeated 15 times, discarding each series one after the other. By this method, the models of equations (1) and (2) are replaced alternatively by 16 different models for our evaluation. The standard deviations of the two coefficients for the bone-muscle threshold models are respectively 0.019 and 0.04 and for the muscle-fat threshold models 0.025 and 0.04.

Table 2. Goodness of fit values of the bone-muscle and muscle-fat threshold estimations in normalised intensity units.

Threshold	Correlation coefficient	Sum of squared residuals	Root mean square error
Bone-Muscle	0.9656	0.002318	0.01287
Muscle-Fat	0.9666	0.005242	0.01935

For each series, the background/foreground segmentations have been qualitatively checked, the voxel classification has also been qualitatively evaluated and anterior-posterior DRR have been computed and qualitatively checked.

3. RESULTS

The background/foreground segmentations have been checked qualitatively, 9 series presenting excellent segmentations, 4 were reasonably good despite some segmentations failures on some slices and 3 presenting either significant tissue regions labelled as background or the contrary. Figure 4 presents the results for an axial MRI image and Figure 5 for a coronal one.

The voxel classification on the 16 series has also been evaluated qualitatively. In regions such as the legs where the three tissues are dominant, a good general classification was observed (Fig. 6). In other regions such as the pelvis area presenting thinner cortical bones and other types of tissues, more classification failures were observed.

Anterior-posterior DRR have been computed for the 16 series and qualitatively checked. Three series were hardly interpretable, while the other were interpretable, but with varying quality. Figure 7 middle and right depict two examples with higher and lower quality. Even though not completely similar to conventional radiograph (Fig. 7 left), characteristics such as skeleton projection are visible. A lateral DRR has also been computed in the region of the proximal femur to compare the result of our method with the method proposed by Kapanen *et al.* [5] (see Fig. 8). The proximal femur projection presents the same characteristics although the general quality remains slightly lower.

Figure 4. Axial raw MRI of the legs (left) and results of the Canny detection (middle) and final foreground-background segmentation (right).

Figure 5. Coronal raw MRI of the legs (left) and results of the Otsu thresholding (middle) and final foreground-background segmentation (right).

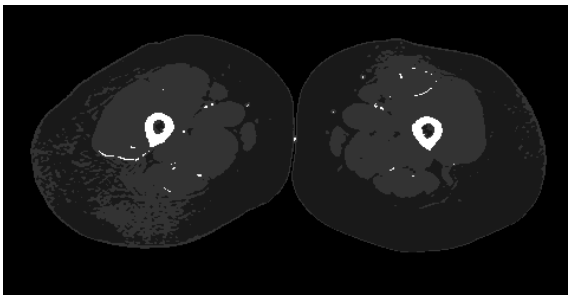

Figure 6. Axial MRI of the legs (original visible in Figure 4 left) where the intensities of the pixels identified as background, cortical bone, muscles and fat have been set respectively to 0, 1, 0.2 and 0.1.

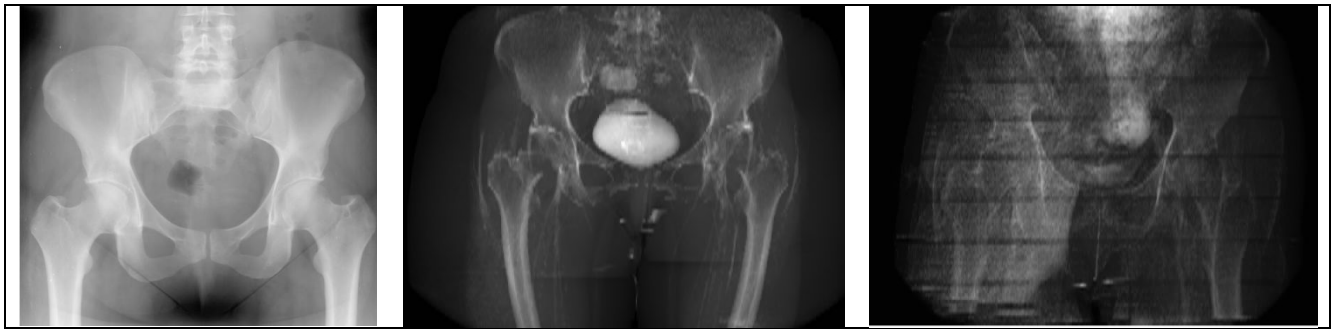

Figure 7. Conventional anterior-posterior radiograph (left) and two anterior-posterior DRR computed from two different MRI series for the pelvis region with intensities for the bone, muscle and fat voxels respectively set to 1, 0.1 and 0.05 (middle and right).

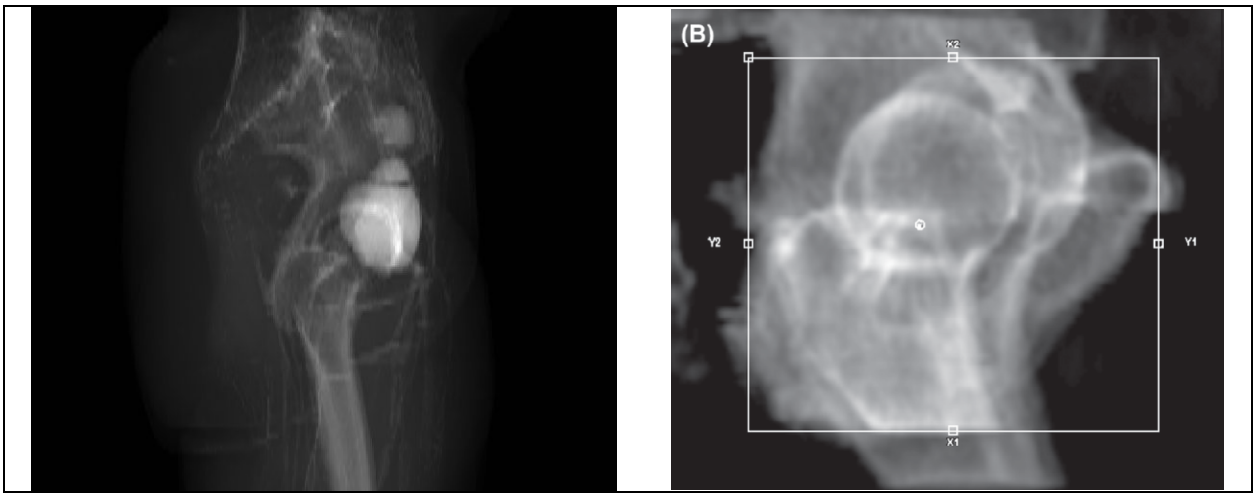

Figure 8. Lateral DRR computed from MRI with the method proposed in this study, with intensities for the bone, muscle and fat voxels respectively set to 1, 0.08 and 0.04 (left) and lateral DRR extracted from Kapanen *et al.* [5] (right).

4. DISCUSSION

The first important point to highlight is that the quality of the DRR relies directly on the quality and resolution of the underlying MRI. Sets with high noise and high thickness between slices are likely to produce low quality DRR, being an intrinsic limitation of this visualisation. This limitation could be observed on our case, where MRI sets with low resolution or high noise produced hardly interpretable DRR (Fig. 7 right, in the case of high noise). Related to this issue, our approach estimates also unique thresholds for a whole MRI set. This can be problematic when brightness is not homogeneous across all the images or within the images themselves, which happens regularly for MRI. This problem could be solved either by pre-processing the images to homogenize the brightness or by developing adaptive thresholding estimations by regions based on our method. In our case, inhomogeneous brightness has indeed been observed on several MRI sets (see Fig. 5 left). Although the segmentations of the background/foreground seemed of good quality in general, some failures were reported due to this issue. Additionally, the manual expert thresholding had to be a trade-off between the bright and dark regions, leading to voxels wrongly classified. These observations limit our modelling and affect the quality of the generated DRR at the current state.

Another intrinsic limitation is the region of interest of the MRI set. If the region is too limited and does not encompass a body region where conventional X-ray could be calculated, the resulting DRR could be difficult to interpret as not referring to any known imaging of the body.

Finally, the method described in this study applies only for T1-weighted MRI. Further work is required to extend the principle presented here to other types of MRI and to validate the method on sets not used to design the model.

We decided to classify the tissues into three types corresponding to identifiable MRI intensities and having distinct attenuation coefficients, usually identifiable on conventional radiographs. While the three tissues appeared in general well classified, all the other tissues were classified in one of the three categories, such as the small white artefacts on Figure 6. As another illustration, bladder fluid, dark on T1-weighted MRI, is mainly classified as bone, leading to bright intensities on the DRR, as visible on Figure 7 (middle). This is a limitation of intensity-based approach as proposed in this study. Morphological considerations may complement this approach to classify more robustly the various tissues relevant for DRR calculation.

Despite these limitations, our study shows that bone-muscle and muscle-fat thresholds could interestingly be related to the intensity of the muscle tissue peak on the histogram. The models proposed in the article show a good reconstruction level, with RMS errors below 0.02 in normalised intensity unit. As DRR and X-rays in general rely primarily on the cortical bone tissue, this ensures interpretable DRR. Indeed, the DRR computed as visible on Figure 7 middle approaches interestingly conventional radiographs. It is moreover possible to emphasize one type of tissue by adapting the relative intensities of the modified MRI. Promising results are also observed in comparison to results proposed by Kapanen *et al.* [5], based on manual segmentations of the bone (Fig. 8).

5. CONCLUSIONS AND FUTURE WORK

Despite limitations, this study validates the proof-of-concept to simulate radiographs from MRI without manual segmentations. As mentioned earlier, this method may not be applicable for every MRI set recorded in clinical routine, but opens an interesting complementary visualisation for PACS servers. Future work consists in validating more deeply the method, making it more robust and generalise it other types of MRI as discussed earlier in order to provide an alternative tool for PACS clients.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 610425 for the project RASimAs.

REFERENCES

[1] Ramsey CR, Oliver AL. Magnetic resonance imaging based digitally reconstructed radiographs, virtual simulation, and three-dimensional treatment planning for brain neoplasms. Med Phys. 1998; 25(10): 1928-34.
[2] Ramsey CR, Arwood D, Scaperoth D, Oliver AL. Clinical application of digitally-reconstructed radiographs generated from magnetic resonance imaging for intracranial lesions. Int J Radiat Oncol Biol Phys. 1999; 45(3): 797-802.
[3] Yin FF, Gao Q, Xie H, Nelson DF, Yu Y, Kwok W, Totterman S, Schell MC, Rubin P. MR image-guided portal verification for brain treatment field. Int J Radiat Oncol Biol Phys. 1998; 40(3): 703-11.
[4] Chen L, Nguyen TB, Jones E, Chen Z, Luo W, Wang L, Price RA Jr, Pollack A, Ma CM. Magnetic resonance-based treatment planning for prostate intensity-modulated radiotherapy: creation of digitally reconstructed radiographs. Int J Radiat Oncol Biol Phys. 2007; 68(3): 903-11.
[5] Kapanen M, Tenhunen M. T1/T2*-weighted MRI provides clinically relevant pseudo-CT density data for the pelvic bones in MRI-only based radiotherapy treatment planning. Acta Oncol. 2013; 52(3): 612-18.
[6] Ullrich S, Grottke O, Fried E, Frommen T, Liao W, Rossaint R, Kuhlen T, Deserno TM. An intersubject variable regional anaesthesia simulator with a virtual patient architecture. Int J Comput Assist Radiol Surg. 2009; 4(6): 561-70.
[7] Grottke O, Ntouba A, Ullrich S, Liao W, Fried E, Prescher A, Deserno TM, Kuhlen T, Rossaint R. Virtual reality-based simulator for training in regional anaesthesia. Br J Anaesth. 2009; 103(4): 594-600.
[8] Chaudhury K. Acceleration of the shiftable O(1) algorithm for bilateral filtering and non-local means. IEEE Trans Image Process. 2013; 22(4): 1291-300.
[9] Canny J. A computational approach to edge detection. IEEE Trans Pattern Anal Mach Intell. 1986; 8(6): 679-98.
[10] Otsu N. A threshold selection method from gray-level histogram. IEEE Trans Syst Man Cybern. 1979; 9: 62-6.