

Robust CT Synthesis for Radiotherapy Planning: Application to the Head & Neck Region

Ninon Burgos¹, M. Jorge Cardoso^{1,2}, Filipa Guerreiro³, Catarina Veiga⁴,
Marc Modat^{1,2}, Jamie McClelland⁵, Antje-Christin Knopf³,
Shonit Punwani^{6,7}, David Atkinson⁶, Simon R. Arridge⁵, Brian F. Hutton^{8,9},
and Sébastien Ourselin^{1,2}

¹ Translational Imaging Group, CMIC, University College London, London, UK

² Dementia Research Centre, Institute of Neurology, UCL, London, UK

³ The Institute of Cancer Research, Radiotherapy Imaging Department, London, UK

⁴ Radiation Physics Group, UCL Medical Physics and Bioengineering, London, UK

⁵ Centre for Medical Image Computing, UCL, London, UK

⁶ Centre for Medical Imaging, UCL, London, UK

⁷ Department of Radiology, University College London Hospitals, London, UK

⁸ Institute of Nuclear Medicine, UCL, London, UK

⁹ Centre for Medical Radiation Physics, University of Wollongong, NSW, Australia

Abstract. In this work, we propose to tackle the problem of magnetic resonance (MR)-based radiotherapy treatment planning in the head & neck area by synthesising computed tomography (CT) from MR images using an iterative multi-atlas approach. The proposed method relies on pre-acquired pairs of non-rigidly aligned T2-weighted MRI and CT images of the neck. To synthesise a pseudo CT, all the MRIs in the database are first registered to the target MRI using a robust affine followed by a deformable registration. An initial pseudo CT is obtained by fusing the mapped atlases according to their morphological similarity to the target. This initial pseudo CT is then combined with the target MR image in order to improve both the registration and fusion stages and refine the synthesis in the bone region.

Results showed that the proposed iterative CT synthesis algorithm is able to generate pseudo CT images in a challenging region for registration algorithms. We demonstrate that the robust affine decreases the overall absolute error compared to a single affine transformation, mainly in images with small axial field-of-view, whilst the bone refinement process further reduces the error in the bone region, increasing image sharpness.

1 Introduction

Magnetic resonance imaging (MRI) is often preferred over computed tomography (CT) as a structural imaging modality, mainly for its excellent soft-tissue contrast. However, MRI does not provide photon density information, which is essential for several applications such as performing dosimetry for MR-based radiotherapy treatment planning (RTP) or attenuation correction in the context of Positron Emission Tomography (PET)/MR scanners. To overcome this limitation, a solution is to recreate a CT image from the available MR images.

Several methods exist to obtain synthetic CT images and many have been applied to RTP. Results presented in [1] are obtained from a Gaussian mixture regression model linking the MRI intensity values to the CT Hounsfield units (HU). In [2], bones are manually segmented and the MRI intensity values are converted to HU using a dual model, within and outside of the bone class. Other approaches, called registration- or atlas-based methods, rely on a single [3] or a database [4–6] of MRI and CT image pairs and consist of registering each atlas MR image to the target MRI, applying the same transformation to the associated CT image and finally, for multi-atlas methods, fusing the registered CT images. The fusion can be obtained by computing the voxelwise median [4], using a probabilistic Bayesian framework [5], an arithmetic mean process or pattern recognition with Gaussian process [6]. Atlas-based methods have also been applied to MR-based attenuation correction. In [7], the fusion is obtained via a voxelwise weighting scheme where the weights are defined according to the morphological similarity between the deformed atlas and target images.

Most of these methods address the problem in the brain [1, 4–6], some in the pelvic area [2, 3], but none in the neck. CT synthesis in the neck is challenging for three main reasons. First, the mixture of bone and air present in this area makes the conversion of MRI intensity values to HU difficult as both bone and air have usually low intensities in MR images. A second challenge comes from the wide range of target image fields of view (FOV). Images with large axial FOV can cover an area starting from the top of the lungs to the middle of the head (Fig 1, bottom), while images with small axial FOV (Fig 1, top) only focus on a target area, such as a tumour. These disparities are a problem when aligning the atlas images to the target. The last challenge arises from the large-scale postural changes, such as flexion or extension of the neck, or the position of the jawbone, which reduces the performances of registration methods.

We propose to redesign the synthesis process and present an iterative multi-atlas algorithm to synthesise CT images in the neck region. The contribution of this paper, compared to the current atlas-based methods, consists of four points.

1. To form the database, the CT and MR images are non-rigidly aligned to account for the different positions between the two acquisitions. We also align all the MR-CT pairs to a common coordinate frame.
2. The first step to synthesise a pseudo CT is to register all the MRI atlases to the target MRI. Large differences in the FOV were observed between the subjects, which can hinder the inter-subject registration. To overcome this problem, we propose a robust alignment exploiting the fact that all the atlases are pre-aligned in the same coordinate frame.
3. Similarly to [7], the pseudo CT is obtained by fusing the mapped atlases according to their morphological similarity to the target. A similarity measure taking into account mismatches between FOVs is proposed.
4. The algorithm relies on the ability to accurately map T2 images from different subjects, a process that can be challenging, particularly in low-contrast areas. We propose to combine multiple modalities, an initial pseudo CT and the target MR image to improve the synthesis accuracy in the bone region.

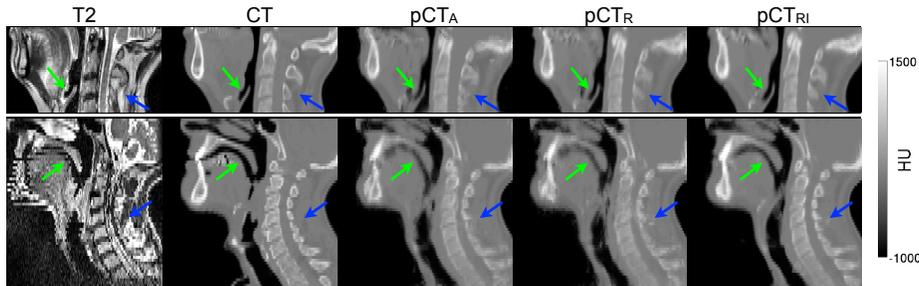


Fig. 1. Examples of T2-w MRI, CT, and pseudo CT images for a subject from dataset 1 (top) and dataset 2 (bottom). Note that the method is able to generate a pseudo CT even in the presence of motion and large-scale postural changes such as neck flexion.

2 Method

The proposed iterative multi-atlas algorithm aims at synthesising CT from MR images of the neck region using a database of non-rigidly aligned MRI and CT pairs. In this section, we apply a robust process to align the atlas images to the target image and propose an iterative process aiming at improving the synthesis in the bone region, critical for radiotherapy planning.

2.1 MR-CT database building

The CT synthesis method relies on pre-acquired pairs of T2-weighted MRI and CT images of the neck. For each subject, the MRI is mapped to the CT using a non-rigid deformation algorithm [8] to compensate for neck flexion. The MRI from the atlases are then mapped to a common coordinate frame via an affine groupwise registration [9]. The transformations are then applied to the MR-CT pairs by updating their image coordinate system, thus forming a database of MR and CT images aligned in a common space.

2.2 Inter-subject mapping

The first step to synthesise a CT for a given MRI is to register all the MRIs in the atlas database to the target MRI. This inter-subject coordinate mapping can be obtained using a symmetric global registration followed by a non-rigid registration [7]. However, when the FOV of the target MRI is limited, which can occur when a tumour is the sole target of the imaging protocol, the affine alignment can be inaccurate. To overcome this weakness, we propose a robust affine which relies on the fact that the atlases are aligned to the same space.

Robust affine The robust affine alignment consists of three steps. Each atlas MR is first affinely mapped to the target. The average transformation, computed in the log space, is obtained from all the pairwise affine transformations. This average affine is then used to initialise a second affine registration step to refine

the alignment. As all the atlases are aligned to the same space, applying the average affine transformation to all of them guarantees that each atlas is initially aligned with the target.

Non-rigid registration The robust affine alignment between atlas and target images is used to initialise a cubic B-spline parametrised non-rigid registration, using normalised mutual information as a measure of similarity [8]. These non-rigid transformations are also applied to all the CTs in the atlas database. Through this registration and resampling procedure, one obtains a series of MRI/CT pairs aligned to the target MRI.

2.3 CT synthesis

Similarly to [7], the pseudo CT is obtained by fusing the mapped atlases according to their morphological similarity to the target. A local image similarity measure between the target MRI and the set of registered MRIs from the database is used as a surrogate of the underlying morphological similarity, under the assumption that if two MRIs are similar at a certain spatial location, the two CTs will also be similar at this location.

Image/Morphological Similarity Due to different acquisition FOVs, the inter-subject mapping and resampling processes introduce areas where no information is available. Those areas have to be accounted for when the similarity measure is computed and during the intensity fusion process. We extend the convolution-based local normalised correlation coefficient (LNCC) method by Cachier *et al.* [10] to irregular regions-of-interest (ROI-LNCC).

Let the target subject’s MRI be denoted by I^{MRI} and, for each of the N atlases in the database, let the mapped MR image of atlas n be denoted by J_n^{MRI} . The ROI-LNCC between I^{MRI} and J_n^{MRI} at voxel x is then given by

$$\text{ROI-LNCC}_n(x) = \frac{\langle I^{MRI}(x), J_n^{MRI}(x) \rangle}{\sigma(I^{MRI}(x)) \sigma(J_n^{MRI}(x))} . \quad (1)$$

Let Ω be a density function equal to 1 where the fields of view overlap, and 0 otherwise. The standard deviation σ and the correlation \langle, \rangle are defined as in [7], but instead of a normal convolution process, a density normalised convolution is used, i.e. the mean will be estimated as $(G * I) / (G * \Omega)$, where G represents a Gaussian kernel, $*$ denotes the convolution operator, and $G * \Omega$ represents a density normalisation term that compensates for areas with missing information. The ROI-LNCC values outside the FOV are set to $-\infty$.

Intensity Fusion The ROI-LNCCs at each voxel are ranked across all atlas images in the database and the ranks, noted as $r_n(x)$, are converted to weights by applying an exponential decay function $w_n(x) = e^{-\beta r_n(x)}$ with $w_n(x)$ being the weight associated with the n^{th} mapped atlas image at voxel x and $\beta = 0.5$ [7]. As in [7], the final pseudo CT (I^{pCT}) is obtained by applying a spatially varying weighted averaging

$$I^{pCT}(x) = \frac{\sum_{n=1}^N w_n(x) \cdot J_n^{CT}(x)}{\sum_{n=1}^N w_n(x)} . \quad (2)$$

2.4 Iterative bone refinement process

In contrast to CT images, T2-weighted MR images do not provide a good estimate of the bone location. We propose to combine multiple modalities to regularise the registration in low-contrast areas, providing more realistic mappings.

An initial pseudo CT (pCT) is synthesised following the method described previously. This pCT is then combined with the MR image to form a MRI-pCT pair. The MRI-pCT pair is registered to all the MRI-CT pairs from the database. The inter-subject coordinate mapping is obtained using the robust affine followed by a multichannel non-rigid registration. Normalised mutual information is used as a similarity measure for the T2-weighted MR channel while the sum of squared differences is computed for the second channel, exploiting the quantitative property of the CT intensities. A refined pseudo CT is obtained by fusing the mapped MRI-CT pairs according to their morphological similarity to the target MRI-pCT pair, which is assessed using a multivariate LNCC (MV-ROI-LNCC) defined as $MV\text{-ROI-LNCC}_n = \text{ROI-LNCC}(I^{MRI}, J_n^{MRI}) + \text{ROI-LNCC}(I^{pCT}, J_n^{CT})$.

3 Validation & Results

Data Images from two retrospective studies were used to build the MR-CT database and validate the proposed method.

Dataset 1: Small axial FOV. Six subjects with a small axial FOV were used to assess the robustness of the method. They have both T2-weighted MRI (voxel size $1 \times 1 \times 2 \text{ mm}^3$, matrix size $192 \times 192 \times 40$) and CT (voxel size $1 \times 1 \times 2 \text{ mm}^3$, matrix size $512 \times 512 \times 200$) images.

Dataset 2: Large axial FOV. Seventeen subjects with a large axial FOV were used to build the database and validate the method. They have both T2-weighted MRI (voxel size $0.7 \times 0.7 \times 3.3 \text{ mm}^3$, matrix size $256 \times 256 \times 60$) and CT (voxel size $1 \times 1 \times 2.5 \text{ mm}^3$, matrix size $512 \times 512 \times 160$) images.

3.1 Algorithmic comparison

Subjects from dataset 2 were used to build the database. Using the proposed technique, we first synthesised pseudo CTs for all subjects in dataset 1 to study the benefits of the robust affine compared with a single affine. We then performed a leave-one-out cross validation using all the subjects from dataset 2 to study the impact of the bone refinement process on the pseudo CT images. For each subject in both datasets, three pseudo CTs were synthesised:

- pCT_A, obtained using a single affine between the atlases and target;
- pCT_R, obtained using the robust affine;
- pCT_{RI}, obtained after the robust affine and the iterative bone refinement.

The mean absolute error (MAE) and the standard deviation of the absolute error (SAE) were calculated for every subject between the reference CT non-rigidly aligned to the MR (R^{CT}) and each of the pseudo CTs (I^{CT}), in a region of interest comprising V voxels: $\text{MAE} = \frac{1}{V} \sum_x |I^{CT}(x) - R^{CT}(x)|$, $\text{SAE} =$

Table 1. Average \pm SD of the MAE and SAE computed between the reference CT and the pseudo CTs, and of the FDE computed for each pseudo CT.

		Dataset 1			Dataset 2		
MAE (HU)	Bone	204.5 ± 19.2	193.6 ± 23.2	196.8 ± 16.0	208.0 ± 19.6	205.2 ± 18.2	192.1 ± 13.4
	Soft	41.2 ± 1.7	40.2 ± 2.2	39.1 ± 1.9	41.1 ± 5.3	40.6 ± 5.1	40.4 ± 4.6
	Air	63.4 ± 8.4	55.6 ± 8.4	56.6 ± 10.3	133.8 ± 29.4	133.0 ± 23.0	126.7 ± 25.4
SAE (HU)	Bone	201.8 ± 40.0	189.5 ± 45.1	186.9 ± 34.9	184.4 ± 14.8	182.0 ± 14.1	169.7 ± 10.9
	Soft	45.8 ± 1.6	44.5 ± 2.1	43.4 ± 2.0	46.5 ± 3.7	46.1 ± 3.7	46.0 ± 3.4
	Air	114.7 ± 11.4	106.0 ± 11.1	108.6 ± 13.7	145.4 ± 18.5	143.8 ± 13.1	150.2 ± 20.6
FDE	All	18.40 ± 0.12	18.11 ± 0.13	17.77 ± 0.13	19.73 ± 0.32	19.69 ± 0.32	19.44 ± 0.34

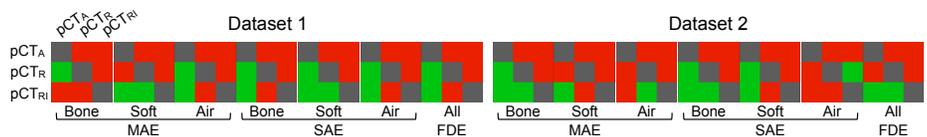


Fig. 2. Results of the one-tailed Wilcoxon signed-rank test at the 5% significance level for the MAE, SAE and FDE metrics. The colour green indicates a significant decrease for the row method when compared to the column method, while the red indicates that the difference is not significant.

$\frac{1}{V-1} \sqrt{\sum_x (|I^{CT}(x) - R^{CT}(x)| - \text{MAE})^2}$. The MAE gives information on the amount of error while the SAE gives information on the amount of dispersion of the error. We focus the validation on three regions: air (< -100 HU), soft-tissue (between -100 and 300 HU), and bone (>300 HU). As inaccuracies exist in the registration between the real CT and T2-weighted images, we also provide a reference-free metric using the entropy of the pseudo CT images in the frequency domain as a measure of sharpness: $\text{FDE} = -\sum (f \log_2(f))$ where f is the normalised absolute amplitudes of frequencies of the pseudo CT image [11].

3.2 Robust affine: validation with small axial FOV

The average and standard deviation (SD) of the MAE, SAE and FDE metrics are presented in Table 1. Using the robust instead of a single affine transform decreases the MAE, SAE and FDE in all the regions. Colour coded statistical test results are presented in Fig 2. These results show significant improvements when comparing synthesis using the robust affine to synthesis using a single affine. An example of MR, CT and pCT images are displayed in Fig 1 (top).

3.3 Iterative bone refinement: validation with large axial FOV

After one iteration of the bone refinement process, the MAE and SAE in the bone region and the FDE are significantly decreased (Table 1 & Fig 2). An example of MR, CT and pCT images are displayed in Fig 1 (bottom).

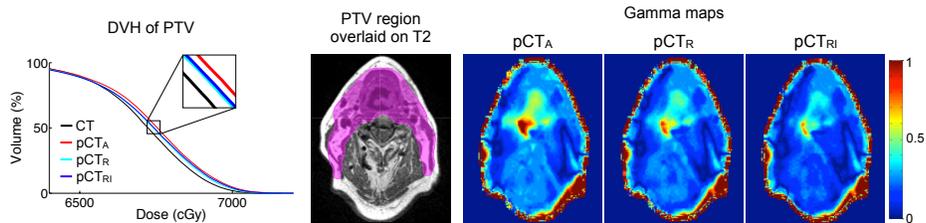


Fig. 3. Results of the dose calculations performed for a test subject from dataset 1.

3.4 Case study: dosimetry calculations

To test the impact of the proposed improvements in RTP, dose calculations were performed for a test subject from dataset 1. We compared the cumulative dose volume histogram (DVH) obtained for the pseudo CTs to the DVH obtained for the reference CT image in the planning target volume (PTV). The DVHs displayed in Fig 3 show a close agreement between the pseudo CTs and the CT. A 2D gamma analysis [12], evaluated at 2%/2 mm dose difference/distance to agreement, was used to compare axial dose distributions. A gamma index greater than 1 means that the calculation does not meet the acceptance criteria. The gamma maps displayed in Fig 3 show that the robust affine improves the results compared with a single affine and that the bone refinement process further reduces the error to an acceptable range ($\gamma < 1$).

4 Discussion & Conclusion

This paper presents a CT synthesis algorithm based on a multi-atlas information propagation scheme and iterative process able to generate pseudo CT images in the neck, a region challenging for registration algorithms. The results displayed in Table 1 demonstrate that the robust affine decreases the absolute error compared to the single affine, particularly when the axial FOV of the target image is limited, and that the bone refinement process further reduces the error in the bone region, producing sharper images. Most of the errors observed when comparing with results obtained in the brain [7] can be explained by inaccuracies in the inter-subject registrations due to a complex mixture of tissues present in the neck area, and by large epiglottis/tongue mismatch between MR and CT data.

Note that while the main focus of the proposed technique is to estimate dose deposition in the context of MR-based RTP, the proposed method could also be used for attenuation correction in PET/MR imaging studies of the neck. As a proof of concept, dose calculations were performed for a test subject (Fig 3), showing close agreement between the proposed pseudo CT and the reference CT, within acceptable error margins ($\gamma < 1$). Further experiments are required to verify the suitability of this proposed framework for dose calculation and to study the effect of anatomical abnormalities, such as tumours, on the results. We will also explore the effect of the number of iterations on the synthesis accuracy.

Acknowledgements This work was supported by an IMPACT studentship funded by Siemens and the UCL FES. Funding was received from the EPSRC (EP/H046410/1, EP/J020990/1, EP/K005278), the MRC (MR/J01107X/1), the EU-FP7 project VPH-DARE@IT (FP7-ICT-2011-9-601055), Alzheimer’s Research UK, the Brain Research Trust, the NIHR BRU (Dementia), the UCL Leonard Wolfson Experimental Neurology Centre (PR/ylr/18575), the Fundação para a Ciência e a Tecnologia (SFRH/BD/76169/2011), and by researchers at the NIHR UCLH BRC (including the High Impact Initiative).

References

1. Jonsson, J., Akhtari, M., Karlsson, M., Johansson, A., Asklund, T., Nyholm, T.: Accuracy of inverse treatment planning on substitute CT images derived from MR data for brain lesions. *Radiation Oncology* **10**(1) (2015)
2. Korhonen, J., Kapanen, M., Keyriläinen, J., Seppälä, T., Tenhunen, M.: A dual model HU conversion from MRI intensity values within and outside of bone segment for MRI-based radiotherapy treatment planning of prostate cancer. *Medical Physics* **41**(1) (2014) 011704
3. Dowling, J.A., Lambert, J., Parker, J., Salvado, O., Fripp, J., Capp, A., Wratten, C., Denham, J.W., Greer, P.B.: An atlas-based electron density mapping method for magnetic resonance imaging (MRI)-alone treatment planning and adaptive MRI-based prostate radiation therapy. *International Journal of Radiation Oncology · Biology · Physics* **83**(1) (2012) e5–e11
4. Sjölund, J., Forsberg, D., Andersson, M., Knutsson, H.: Generating patient specific pseudo-CT of the head from MR using atlas-based regression. *Physics in Medicine and Biology* **60**(2) (2015) 825
5. Gudur, M.S.R., Hara, W., Le, Q.T., Wang, L., Xing, L., Li, R.: A unifying probabilistic Bayesian approach to derive electron density from MRI for radiation therapy treatment planning. *Physics in Medicine and Biology* **59**(21) (2014) 6595
6. Uh, J., Merchant, T.E., Li, Y., Li, X., Hua, C.: MRI-based treatment planning with pseudo CT generated through atlas registration. *Med Phys* **41**(5) (2014)
7. Burgos, N., Cardoso, M.J., Thielemans, K., Modat, M., Pedemonte, S., Dickson, J., Barnes, A., Ahmed, R., Mahoney, C.J., Schott, J.M., Duncan, J.S., Atkinson, D., Arridge, S.R., Hutton, B.F., Ourselin, S.: Attenuation Correction Synthesis for Hybrid PET-MR Scanners: Application to Brain Studies. *Medical Imaging, IEEE Transactions on* **33**(12) (2014) 2332–2341
8. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* **98**(3) (2010) 278–84
9. Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging* **1**(2) (2014) 024003
10. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: the PASHA algorithm. *Computer Vision and Image Understanding* **89**(2-3) (2003) 272–298
11. Kristan, M., Pernus, F.: Entropy based measure of camera focus. In: *Proc. ERK2004*. (2004) 179–182
12. Low, D.A., Harms, W.B., Mutic, S., Purdy, J.A.: A technique for the quantitative evaluation of dose distributions. *Medical Physics* **25**(5) (1998) 656–661