

Joint Segmentation and CT Synthesis for MRI-only Radiotherapy Treatment Planning

Ninon Burgos¹(✉), Filipa Guerreiro², Jamie McClelland³, Simeon Nill²,
David Dearnaley⁴, Nandita deSouza⁵, Uwe Oelfke², Antje-Christin Knopf⁶,
Sébastien Ourselin^{1,7}, and M. Jorge Cardoso^{1,7}

¹ Translational Imaging Group, CMIC, University College London, London, UK
n.burgos@ucl.ac.uk

² Joint Department of Physics, Institute of Cancer Research and Royal Marsden
NHS Foundation Trust (ICR/RMH), London, UK

³ Centre for Medical Image Computing, University College London, London, UK

⁴ Academic Urology Unit, ICR/RMH, Sutton, UK

⁵ CRUK Centre for Cancer Imaging, ICR/RMH, Sutton, UK

⁶ Department of Radiation Oncology, University of Groningen, University Medical
Center Groningen, Groningen, The Netherlands

⁷ Dementia Research Centre, Institute of Neurology, UCL, London, UK

Abstract. Accurate knowledge of organ location and tissue attenuation properties are the two essential components to perform radiotherapy treatment planning (RTP). Computed tomography (CT) has been the modality of choice for RTP as it easily provides electron density information. However, its low soft tissue contrast limits the accuracy of organ delineation. On the contrary, magnetic resonance (MR) provides images with excellent soft tissue contrast but its use for RTP is limited by the fact that it does not readily provide tissue attenuation information.

In this work we propose a multi-atlas information propagation scheme that jointly segments the organs at risk and generates pseudo CT data from MR images. We demonstrate that the proposed framework is able to automatically generate accurate pseudo CT images and segmentations in the pelvic region, bypassing the need for CT scan for accurate RTP.

1 Introduction

The aim of radiotherapy treatment planning (RTP) is to deliver an optimal dose of radiation over the target area while sparing the normal tissues. RTP first requires contouring the target and organs at risk (OARs). Once these volumes have been defined, the optimal dose distribution for treating the tumour is determined according to the attenuation properties of the different tissues. Most radiotherapy treatments are planned using an X-ray computed tomography (CT) scan of the patient. The acquisition of a CT is fast and the tissue attenuation coefficients can easily be derived from the CT intensity values in Hounsfield unit (HU). However, CT images have low soft tissue contrast, which can lead to large variations when delineating the organs, particularly when located in

the brain, head & neck, or pelvic regions. Magnetic resonance (MR) imaging is often preferred over CT as a structural imaging modality, mainly for its excellent soft tissue contrast. Although increasingly used in clinical practice, the role of MR in RTP is currently limited by the fact that it does not readily provide electron density information, hampering the calculation of dose distributions. In this work, we propose to tackle the problem of RTP from MR images by developing a multi-atlas propagation framework to jointly delineate the OARs and estimate the tissue attenuation properties.

Multi-atlas propagation was first introduced for segmentation purposes [1, 2]. The technique relies on a database of atlases composed of an intensity image and a segmented image. To segment the target image, a first step consists of registering the atlas intensity images to the target intensity image and to apply the same transformations to the associated segmented images. A second step consists of fusing the propagated segmented images to generate the target segmentation. The technique was later extended to the synthesis of images by propagating intensity images instead of segmented images, for example CT images [3]. Multi-atlas CT synthesis methods have been developed for RTP [4–6] but only a few have been applied outside of the brain [7, 8].

Even though atlas-based segmentation and CT synthesis methods have been successfully applied to RTP independently, a key to expand MR-based planning is to guarantee that the segmentations and pseudo CT generated from the MR images match each other, *i.e.* a voxel labelled as bone should have a bone density value in the pseudo CT image. This is not guaranteed if the segmentation and CT synthesis tasks are handled separately. Dowling et al. [8] proposed to combine CT synthesis and segmentation using a database of atlases each composed of an MR, a CT and a segmented image. The atlas MR images were first registered to the target MR image and the same transformations were applied to both the atlas CT and segmented images. A local weighted voting method was then used to generate the target CT and segmented images. In this paper, we develop an iterative multi-atlas propagation framework that combines in a single pipeline segmentation and CT synthesis, with the aim to improve both the segmentation and synthesis accuracy, and guarantee consistent results.

2 Methods

2.1 Database Building

A multi-atlas database consisting of a T2-weighted and a T1-weighted MR image, a CT image, and a segmented image for each atlas was created by non-rigidly registering each subject's CT and T1 images to their T2 image, after resampling the MR images to isotropic resolution and intensity non-uniformity correction. The segmentations were obtained by manually contouring the T2 image. The number of atlases was artificially increased by left-right flipping the images.

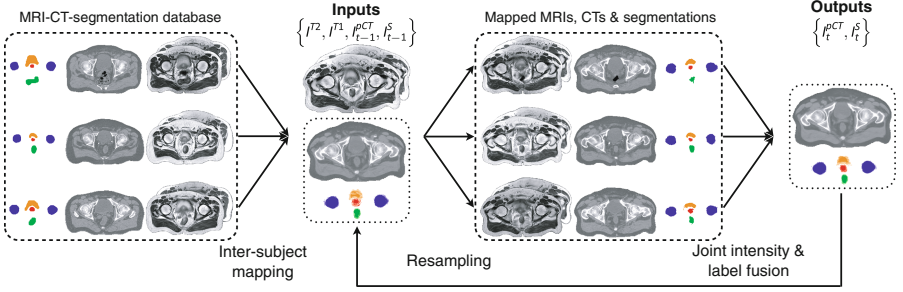


Fig. 1. Joint segmentation and CT synthesis at iteration t . All the atlases are non-rigidly registered to the target. A local similarity measure between the mapped atlases and the target is used to jointly generate a pseudo CT and a segmented image.

2.2 Joint Iterative Segmentation and Image Synthesis

In the proposed iterative framework, a set of probabilistic segmentations and pseudo CT (pCT) images is jointly generated from the target’s MR images by registering the atlases to the target, and fusing the atlas segmentations and CT images according to the similarity between the target and atlases. Similarly to [9], this set is then combined with the target MR images and used to improve the subsequent multi-atlas registration and fusion. A diagram illustrating the proposed method is shown in Fig. 1.

Joint Iterative Label and Intensity Fusion. Let the target subject’s dataset at iteration t be denoted by $\mathbb{I}_t = \{I^{T2}, I^{T1}, I_t^{pCT}, I_t^S\}$ with I^{T2} and I^{T1} the T2- and T1-weighted MR images, and I_t^{pCT} and I_t^S the pCT and segmented images obtained at iteration $t - 1$. Let the dataset of the n th atlas in the database be denoted by $\mathbb{I}_n = \{I_n^{T2}, I_n^{T1}, I_n^{CT}, I_n^S\}$ where I_n^{CT} corresponds to the real CT and I_n^S to the manual segmentation.

The first step of the method consisted of registering each atlas to the target subject. This inter-subject coordinate mapping was obtained using a symmetric global registration followed by a cubic B-spline parametrised non-rigid multi-channel registration as implemented in NiftyReg [10]. The local normalised cross-correlation was used as a similarity measure for the MR channels while the sum of squared differences was computed for the CT channel, exploiting the quantitative property of the CT intensities, and the Kullback-Leibler divergence for the segmentation channels. At each iteration t , a new transformation $\mathcal{T}_{n,t}$ that maps atlas n to the target was defined, generating a set of images aligned to the target: $\mathbb{J}_{n,t} = \{J_{n,t}^{T2}, J_{n,t}^{T1}, J_{n,t}^{CT}, J_{n,t}^S\}$ where $J_{n,t}^X = I_n^X(\mathcal{T}_{n,t}(x))$.

The pCT and probabilistic segmentations were then obtained by fusing the N atlases mapped to the target subject as follows:

$$I_t^{pCT}(x) = \frac{\sum_{n=1}^N w_{n,t}(x) \cdot J_{n,t}^{CT}(x)}{\sum_{n=1}^N w_{n,t}(x)}, \quad (1)$$

$$I_t^{S_i}(x) = \frac{\sum_{n=1}^N w_{n,t}(x) \cdot V_{n,t}^l(x)}{\sum_{n=1}^N w_{n,t}(x)} \text{ with } V_{n,t}^l(x) = \begin{cases} 1 & \text{if } J_{n,t}^{S_i}(x) = l, l \in \{1..L\} \\ 0 & \text{otherwise} \end{cases}.$$

l indexes through the labels and L is the number of all possible labels. V_n^l is the vote for label l produced by the n th atlas. If required, the categorical label result \mathcal{L} at location x can thus be obtained by estimating $\mathcal{L}(x) = \max_l(I^{S_i}(x))$.

The weighting factor $w_{n,t}(x)$ was obtained by applying an exponential decay function to the rank $r_{n,t}(x)$ of the measure used to assess the similarity between atlas n and the target: $w_{n,t}(x) = e^{-\beta_t r_{n,t}(x)}$. After each iteration, the registration for all the atlases improves and more atlases can contribute to the fusion. As a smaller β means that more atlases contribute to the average, we set β to decrease with the number of iterations (by 0.125 starting from $\beta_1 = 0.75$).

Note that at the first iteration both the inter-subject mapping and fusion steps were based on the MR images only.

Convolution-Based Local Similarity Measures. To locally select the atlases used in the fusion, a combination of two similarity measures computed between the target and atlases was used. The structural similarity [11] extended to irregular regions-of-interest (ROI) was computed on the MR and CT channels. The ROI-SSIM between images I and J at voxel x is given by

$$\text{ROI-SSIM}(I(x), J(x)) = \frac{2\mu_I(x)\mu_J(x) + C_1}{\mu_I^2(x) + \mu_J^2(x) + C_1} \frac{2\sigma_{I,J}(x) + C_2}{\sigma_I^2(x) + \sigma_J^2(x) + C_2}. \quad (2)$$

C_1 and C_2 are two constants used to improve the stability of the structural similarity [11]. Let Ω be a density function equal to 1 where the fields of view (FOV) overlap, and 0 otherwise. The means and standard deviations at voxel x were calculated using a Gaussian kernel G_{σ_G} with standard deviation σ_G through density normalised convolution

$$\mu_I(x) = \frac{[G_{\sigma_G} * I](x)}{[G_{\sigma_G} * \Omega](x)}, \sigma_I^2(x) = \mu_{I^2}(x) - \mu_I^2(x), \sigma_{I,J}(x) = \mu_{I \cdot J}(x) - \mu_I(x) \cdot \mu_J(x),$$

where $*$ denotes the convolution operator and $G_{\sigma_G} * \Omega$ represents a density normalisation term that compensates for areas with missing information. As the values of ROI-SSIM are only valid within the bounds of the FOV, values outside the FOV were set to $-\infty$.

A local fuzzy dice score coefficient (DSC) defined per label l and summed over all labels was used to assess the local overlap between the segmented images I^S and J^S

$$\text{LDSC}(I^S(x), J^S(x)) = \sum_{l \in \{1..L\}} \frac{2 \min(\mu_{I^{S_l}}(x), \mu_{J^{S_l}}(x))}{\mu_{I^{S_l}}(x) + \mu_{J^{S_l}}(x)}. \quad (3)$$




Fig. 2. The categorical segmentations of the prostate (red), bladder (orange), rectum (green) and femur heads (blue), and the pCTs were obtained from the T1-weighted and T2-weighted MR images of the patient. Note the increased smoothness of the segmentations and sharpness of the pCTs between *global fusion* and *local fusion iter 4*.

$\mu_I^{S_t}$ was obtained by convolving the segmentation density I^{S_t} with the Gaussian kernel G_{σ_G} : $\mu_I^{S_t}(x) = [G_{\sigma_G} * I^{S_t}](x)$.

The final local similarity measure (LSIM) computed at iteration t between the target and the n th atlas was obtained by summing the ROI-SSIM computed on the MR and CT channels and the LDSC computed on the segmentation channel: $\text{LSIM}(\mathbb{I}_{t-1}, \mathbb{J}_{n,t}) = \text{ROI-SSIM}(I^{T2}, J_{n,t}^{T2}) + \text{ROI-SSIM}(I^{T1}, J_{n,t}^{T1}) + \text{ROI-SSIM}(I_{t-1}^{pCT}, J_{n,t}^{CT}) + \text{LDSC}(I_{t-1}^S, J_{n,t}^S)$.

3 Validation and Results

Data. The proposed framework was evaluated on 15 prostate cancer patients. Each subject had a T2-(2D spin echo; TE/TR: 80/2500 ms; $1.46 \times 1.46 \times 5 \text{ mm}^3$) and a T1-weighted (2D spin echo; TE/TR: 10/400 ms; $1.64 \times 1.64 \times 5 \text{ mm}^3$) MR image, and a CT image (140 kVp, voxel size $0.98 \times 0.98 \times 1.5 \text{ mm}^3$) acquired the same day. The delineation of the OARs was performed by a qualified clinician.

The performance of the proposed methodology was compared with reference data for the 15 subjects following a leave-one-out cross-validation strategy. Four iterations of the pipeline were computed. They are referred to as *local fusion iter t* with $t = 1, 2, 3, 4$. An additional method (*global fusion*) equivalent to the first iteration but with a global atlas selection instead of per voxel, *i.e.* the weights $w_{n,t}$ in Eq. (1) are not a function of x , was also implemented for comparison with state-of-the-art [2]. An example of reference data and results obtained with the *global fusion* and *local fusion iter 1* and *4* are displayed in Fig. 2.

3.1 Segmentation Accuracy

The segmentation accuracy was assessed by computing the fuzzy DSC between the manual and atlas-based segmentations for the different OARs considered (bladder, prostate, rectum, and left (L) and right (R) femur heads). Results displayed in Fig. 3 show a statistically significant improvement for most OARs




Fig. 3. Boxplots displaying the median, lower and upper quartiles, and minimum and maximum of the DSC calculated between the manual and atlas-based segmentations (left); and the MAE computed between the reference and pseudo CTs (right). The stars indicate a significant improvement between the current and previous measurements.




Fig. 4. 3D representation of the prostate manually delineated and of the probabilistic segmentations (thresholded at 0.5) obtained with the proposed method. Note how the smoothness of the segmentations increases with the number of iterations.

between the global and local fusion methods, and between the first and second iterations. Statistical significance was assessed using the paired one-tailed Wilcoxon signed-rank test, with a 5% significance level. On average for the 4th iteration of the proposed method, the DSC was 0.90 for the bladder, 0.74 for the prostate, 0.76 for the rectum and 0.90 for both femur heads. Although the DSC does not significantly improve after the 2nd iteration, we observed an increase in the segmentation smoothness when we kept iterating, as illustrated in Fig. 4.

3.2 CT Synthesis Accuracy

The mean absolute error ($MAE = \frac{1}{V} \sum_x |I^{pCT}(x) - R^{CT}(x)|$) and the mean error ($ME = \frac{1}{V} \sum_x I^{pCT}(x) - R^{CT}(x)$) were calculated for every subject between the reference CT non-rigidly aligned to the MR (R^{CT}) and each of the pseudo CTs (I^{pCT}) in an ROI comprising V voxels. In both the full image and the bone region (delineated by thresholding the reference CT), the MAE showed a significant decrease in synthesis error between the global and local approaches, and between each iteration (Fig. 3). The low ME obtained for all the methods (-3 ± 10 HU on average in the full image and -2 ± 32 HU in the bone region) demonstrates an absence of bias in the methodology.




Fig. 5. Left: DVHs obtained for the reference CT (solid lines) and pseudo CT (dash-dot lines) in the PTV and OARs for a representative subject. Right: Boxplots of the percentage differences evaluated for several DVH points obtained for 15 subjects.

3.3 Dosimetry Calculations

Dose calculations were performed using the RayStation treatment planning system to assess the applicability of the proposed framework for RTP. We compared the cumulative dose volume histogram (DVH) obtained for the pseudo CT obtained after the 4th iteration (I_4^{pCT}) to the DVH obtained for the reference CT image in the planning target volume (PTV) located in the prostate, and in the OARs. The same contours were used for the pseudo and reference CT images. The DVHs displayed in Fig. 5 for a representative subject shows a close agreement between the doses calculated from the reference and pseudo CT images. For all the subjects, percentage differences were evaluated for several DVH points: $D_{98\%}$, D_{mean} and $D_{2\%}$ for the PTV, and D_{mean} and $D_{2\%}$ for the OARs (bladder, rectum and femur heads). D_x is the dose given to $x\%$ of the structure volume and D_{mean} is the mean dose given to the evaluated volume. Results are displayed in Fig. 5. On average for all the DVH points considered, the percentage difference was less than $\pm 0.15\%$ for the PTV and all the OARs.

4 Discussion and Conclusion

This paper presents a joint segmentation and CT synthesis framework for MRI-only RTP able to automatically generate accurate pseudo CT images and OAR contours in the pelvic region. Solving the segmentation and synthesis tasks simultaneously results in not only having the solutions in agreement, but the joint estimation aids in improving the accuracy of each aspect, as seen in Fig. 3, where we observe that an increase in segmentation overlap corresponds to a decrease in CT synthesis error.

When evaluating the segmentation accuracy obtained with their method, Dowling et al. [8] reported mean DSCs of 0.80, 0.86 and 0.84 for the prostate, bladder and rectum, respectively. Wong et al. [12] assessed the performance of several multi-atlas segmentation methods. With their recommended setting, the median DSC for the prostate, bladder, rectum and femurs was 0.84, 0.90, 0.77

and 0.95, respectively. Taking into account the lower resolution of the images used in this work (5 mm slice thickness versus 2 mm in [8] and 2.5 mm in [12]), we obtained comparable results with mean DSCs of 0.74, 0.89, 0.76 and 0.90.

Regarding the CT synthesis, the proposed method generates pseudo CTs as accurate as previously reported for the pelvic region, even though they were obtained from lower resolution MR images. Dowling et al. [8] reported an average MAE of 40.5 ± 8.2 HU while we obtained an average MAE of 42.9 ± 4.0 HU.

As future work, the probabilistic property of the segmentations obtained with the proposed framework could be used to automatically define margins when contouring organs and improve MR-based RTP.

Acknowledgements. Funding was received from the CMIC-EPSRC platform grant, the EPSRC (EP/H046410/1, EP/K005278) and the NIHR BRC UCLH/UCL High Impact Initiative-BW.mn.BRC10269. Research at ICR is supported by CRUK (C33589/A19727), NHS funding to the NIHR BRC at ICR/RMH, and CRUK and EPSRC support to the Cancer Imaging Centre at ICR/RMH in association with MRC and Department of Health (C1060/A10334, C1060/A16464).

References

1. Cabezas, M., Oliver, A., Lladó, X., Freixenet, J., Cuadra, M.B.: A review of atlas-based segmentation for magnetic resonance brain images. *Comput. Methods Programs Biomed.* **104**(3), e158–e177 (2011)
2. Klein, S., van der Heide, U.A., Lips, I.M., van Vulpen, M., Staring, M., Pluim, J.P.: Automatic segmentation of the prostate in 3D MR images by atlas matching using localized mutual information. *Med. Phys.* **35**(4), 1407–1417 (2008)
3. 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. *Med. Imaging IEEE Trans.* **33**(12), 2332–2341 (2014)
4. 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. *Phys. Med. Biol.* **59**(21), 6595 (2014)
5. 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), 51711 (2014)
6. Sjölund, J., Forsberg, D., Andersson, M., Knutsson, H.: Generating patient specific pseudo-CT of the head from MR using atlas-based regression. *Phys. Med. Biol.* **60**(2), 825 (2015)
7. Burgos, N., Cardoso, M.J., Guerreiro, F., Veiga, C., Modat, M., McClelland, J., Knopf, A., Punwani, S., Atkinson, D., Arridge, S.R., Hutton, B.F., Ourselin, S.: Robust CT synthesis for radiotherapy planning: application to the head and neck region. *Med. Image Comput. Comput. Assist. Interv. MICCAI* **2015**, 476–484 (2015)

8. Dowling, J.A., Sun, J., Pichler, P., Rivest-Hénault, D., Ghose, S., Richardson, H., Wratten, C., Martin, J., Arm, J., Best, L., Chandra, S.S., Fripp, J., Menk, F.W., Greer, P.B.: Automatic substitute computed tomography generation and contouring for magnetic resonance imaging (MRI)-alone external beam radiation therapy from standard mri sequences. *Int. J. Radiat. Oncol. Biol. Phys.* **93**(5), 1144–1153 (2015)
9. Bai, W., Shi, W., O'Regan, D.P., Tong, T., Wang, H., Jamil-Copley, S., Peters, N.S., Rueckert, D.: A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: Application to cardiac MR images. *Med. Imaging IEEE Trans.* **32**(7), 1302–1315 (2013)
10. 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. *Comput. Methods Programs Biomed.* **98**(3), 278–284 (2010)
11. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *Image Process. IEEE Trans.* **13**(4), 600–612 (2004)
12. Wong, W.K., Leung, L.H., Kwong, D.L.: Evaluation and optimization of the parameters used in multiple-atlas-based segmentation of prostate cancers in radiation therapy. *Brit. J. Radiol.* **89**(1057), 20140732 (2015)