

Multi-Atlas based pseudo-CT Synthesis using Multimodal Image Registration and Local Atlas Fusion Strategies

Johanna Degen^{1,2}, Mattias P. Heinrich¹

¹ Institute of Medical Informatics, University of Lübeck, Germany

[degen,heinrich]@imi.uni-luebeck.de

² Graduate School for Computing in Medicine and Life Sciences, University of Lübeck, Germany

Abstract

The synthesis of pseudo-CT images is of particular interest in the development of hybrid PET-MRI devices and MRI-guided radiotherapy. These images can be used for attenuation correction during PET image reconstruction. Furthermore, using MRI-based radiotherapy planning would enable a more accurate dosimetry planning due to the superior soft tissue contrast of the scans. The previously proposed methods for pseudo-CT synthesis are characterised by mainly two drawbacks. First, most proposed methods are limited to the head and neck region and therefore not feasible in case of whole body applications. Second, the presence of aligned training pairs of both MRI and CT scans for a number of subjects is assumed. In this work, we present preliminary results for atlas-based approaches using multiple CT atlas scans (from different patients) to synthesise a pseudo-CT image for a new patient using only their MRI data. This application requires accurate and robust deformable multimodal registration. We employed a recent discrete optimisation registration framework together with a self-similarity-based metric to accurately match the CT atlases to the anatomy of the patient. The registered atlases are then jointly combined by means of local fusion strategies. We apply our method to different 3D whole body MRI scans and a total of 18 3D whole body CT atlases. In addition to intensity fusion, the proposed methods can also be used for label fusion. Since evaluation based directly on synthesised intensity values is problematic, we use the Dice overlap after the fusion of segmentation labels as a proxy measure. Our proposed new method, which uses MIND descriptors for multimodal label fusion shows overall the best results.

1. Introduction

During the last two decades image synthesis has evolved to a versatile methodology in the field of medical image pro-

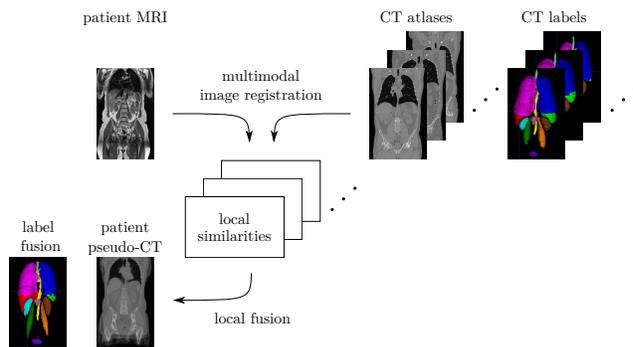


Figure 1: Basic overview of the proposed method for multimodal pseudo-CT synthesis. All atlas CTs are registered to the patient’s MRI using efficient multimodal image registration. The local fusion of the CT atlases is according to the local similarities based on self-similarities descriptors and normalized cross correlations.

cessing. It is for example used to generate missing scans in a series of multispectral magnetic resonance imaging (MRI) data sets or for reducing the number of scans to be acquired [12, 20]. Besides these synthesised MRI sequences, pseudo computer tomography (CT) images are also of great interest, since they can be used in MRI-guided radiotherapy. Owing to the superior soft tissue contrast of MRI, an improvement in dosimetry planning could be achieved. However, since information about the electron density is necessary for the calculation of radiation dose, a treatment planning based on MRI also necessitates the synthesis of a pseudo-CT.

Similarly, for the current developments of realising a hybrid combination of positron emission tomography (PET) and MRI within the same device, there is a compelling need of pseudo-CTs. This is due to the essential attenuation correction step during the PET image reconstruction. For this purpose so-called attenuation maps are created, which are based on the absorption behaviour of the individual struc-

tures. But in contrast to conventional PET-CT no CT scan is available and information for attenuation correction must thus be derived using the MRI scan. However, due to the different underlying physical phenomena the intensities of a MRI image can not be readily converted to typical Hounsfield units (HU) of a CT image. Especially, bony structures or air, which are both represented by low MRI intensities have entirely different HU responses in the CT image. Therefore, other ways must be found to overcome the lack of functional relationship.

To date, several research groups have proposed methods for synthesising pseudo-CTs and thus for creating attenuation maps, including segmentation based approaches (see Sec. 2.1), learning or transfer function based techniques (see Sec. 2.2) or algorithms using atlas-based registration (see Sec. 2.3). In particular the former two methods are usually limited to specific anatomical sights (*e.g.* the head where rigid intra-patient alignment is sufficient) or even more importantly to the availability of perfectly (nonrigidly) aligned training pairs of CT and MRI scans for training, and may be limited to a small number of differentiable classes (such as CSF, white and gray matter). We, in contrast, propose here a new direction, which results in a generic approach that is applicable to any anatomic sights including whole-body scans and does not require pre-aligned training data and is integrally based on multimodal deformable registration. We believe that while accurate multimodal deformable registration of whole body scans is still very challenging, combining it with powerful local atlas fusion strategies will lead to advances in pseudo-CT synthesis and open up new possible applications in radiotherapy or PET reconstruction. Our new image synthesis method (see Figure 1) employs robust image registration to generate a patient specific pseudo-CT based on the patient MRI scan and multiple CT atlases of other subjects. The atlas fusion is based on a local weighting, which corresponds to the multimodal similarity between the patient MRI and a certain atlas CT. We tested the local normalised cross correlation (LNCC) and the modality independent neighbourhood descriptor (MIND) [7] as similarity measure for label fusion. We show experimental results, compare them to the classic majority vote and reference CT data sets, and evaluate them in terms of Dice overlap of segmented structures and visually for pseudo-CT synthesis.

2. Related Work

The majority of the previously presented methods for pseudo-CT synthesis respectively attenuation map generation can be divided into three main classes.

2.1. Segmentation based synthesis

The principle idea of this method is to segment various structures in the patient MRI scan, which are then assigned

to the corresponding HUs. In [17] this principle has been used by Le Goff-Rougetet *et al.* for the first time. They used MRI and PET transmission scans of the head to generate simple attenuation maps. The three classes skin, brain and bone were segmented within the aligned MRT scan. Zaidi *et al.* [23] introduced a more discriminative method with the help of fuzzy-logic. With their approach, a classification in air, brain, skull, paranasal sinuses and scalp has been made possible. However, both methods are limited to the head and are not readily applicable in real PET-MRI devices, since there is usually not enough space for a transmission source.

Moreover, the classification of the bone is a challenging problem because the cortex is difficult to distinguish from air. To solve this problem, the usage of special MRI sequences has been proposed in [15],[1]. Due to the particularly short T2 relaxation time of the cortex, ultrashort echo time sequences can be used for a improved separation of bone structures from air and fat can be better differentiated from other soft tissue with the help of the Dixon technique.

Martinez-Müller *et al.* extended the segmentation based approach to whole body data sets. In [18] they also proposed to perform segmentations of lungs, fat, other soft tissue and background with suitable thresholds on images, which were acquired using the Dixon technique. The corresponding attenuation coefficients are assigned based on a bilinear transform of previously co-registered CT scans.

2.2. Synthesis with transfer functions and machine learning

A variety of machine learning methods enable the training of transfer functions for conversion of MRI intensities to corresponding HUs.

In the context of the registration of brain images Guimond *et al.* suggested in [5] the reduction of the multimodal problem to a monomodal one by means of transferring one modality into the other modality. In the course of this either mono- or bifunctional relationships between the intensities were supposed. The parameters of the polynomial model were estimated robustly with a combination of least trimmed squares and binary re-weighted least squares.

The work of Johansson *et al.* paved the way for using transfer functions also in the context of attenuation correction in hybrid PET-MRI devices. They performed a Gaussian mixture regression to train and extract model parameters for the synthesis of pseudo-CTs. The special feature of the proposed method in [14] is the usage of different MRI sequences and their smoothed versions. This work points up the promising potential of multispectral image data for the image synthesis with transfer functions.

Further machine learning approaches have been developed to tackle the problem of image synthesis. The general assumption of [12, 13, 20] is that small image patches

sufficiently characterise (and discriminate) local anatomical content. Therefore classifiers, such as random forests [13], patch-matching [12, 22] or restricted Boltzmann machines [20] can be employed to assign a likely intensity of a new modality to a voxel. These approaches, however, heavily rely on the availability of well-aligned training data. They may also have disadvantages for the use in whole-body scans, where the anatomical context is usually not well-represented by a small image patch and ambiguities arise from non-functional relations between CT and MRI for a wider range of tissue classes (than found in head scans).

2.3. Atlas based synthesis

Another form of image synthesis is constituted by atlas based registration methods. In contrast to segmentation based approaches, they may enable the estimation of spatially varying attenuation coefficients within the same organ. The aim of the atlas based methods is to find the best possible registration to an patient specific MRI scan. To adapt anatomical variations of individual subjects, several atlases are usually used. This automatically leads to the question in what way the results of image registration are fused into a single pseudo-CT.

In the method presented in [19], one high-quality CT atlas is adjusted by means of multimodal image registration to a patient’s MRI scan. For this purpose, an extended version of optical flow is used, so that it is suitable for multimodality registration. A disadvantage of this approach, however, is that larger transformations can not be estimated, and therefore the method was tested only on data sets of the head. Moreover, it is not possible to represent pathologies using only one atlas which is a general problem in the field of atlas based image synthesis.

Another approach, which combines machine learning and atlas based registration for synthesis has been proposed by Hofmann *et al.* in [11]. They pursue a combination of pattern recognition and atlas registration in order to incorporate both local and global information optimally in the synthesis. The underlying assumption is that a single voxel may not be significant enough to decide to which tissue it belongs. A mapping of MRI patches to HUs is learned with the help of a Gaussian process. Registered CT-MRI pairs of the head are taken as a basis for training.

A method which focuses on multi-atlas registration, was presented by Burgos *et al.* in [3]. Based on the concept of morphological similarity, a locally varying weighting of the individual atlases is derived. Thereby more similar atlases should be more involved into the image synthesis in comparison to more dissimilar atlases. The morphological similarity is expressed by means of a similarity measure. In [2], the use of the local normalised cross correlation (LNCC) is proposed. This method is based on paired and aligned CT-MRI scans. Therefore only a monomodal inter-subject

registration must be carried out, which can then be used for the deformation of the CT atlases. Hence the registered CT atlases can be fused according to the weighting depending on the local monomodal similarity.

In this work, we adapt this method and extend it to multimodal problems. Only the MRI scan of the patient and a series of CT atlases of other patients are assumed to be given. The used method for multimodal image registration is described in the next section. The local weighting is determined by comparison the multimodal data sets using LNCC or MIND as similarity measures. This approach and the extension of the methods to label fusion are explained following the registration method. A comparison of the results with the majority vote and a non-negative least squared regression is given in section 5.

3. Multimodal Image Registration

The image registration method is in addition to the local fusion strategy the most critical part of our approach since its quality has a direct impact on the synthesis result. For as robust as possible registration of the CT atlases to the patient MRI data set, the dense displacement sampling algorithm (deeds) [8],[9] was selected as multimodal image registration method. Within this algorithm, the self-similarity context is used as matching criterion. The main advantage of this registration method is due to the fact that the solution is found by a discrete optimization scheme according to the theory of Markov-Random-Field labeling. By applying dynamic programming, the deformation field can be found very efficiently while at the same time a dense displacement sampling can be employed. Similar to many other registration methods, a multi level approach is applicable. In this case, however, only the image at the highest resolution is used. For the different levels this high-resolution image is divided into non-overlapping groups of cubic pixels. Thus, the described registration method is determined by a number of parameters: The distance between the nodes of a regular B-spline grid is defined by the parameter g and l_{max} sets the search radius and is thus related to the maximum number of possible labels. A coarse to fine search for a wide range of possible displacements is performed over several scale.

4. Multimodal Local Atlas Fusion

The aim of our multimodal local atlas fusion is to merge within a specific localized area only those atlases to a pseudo-CT, which are morphologically similar to the patient’s MRI. Therefore we assess at each image position the similarity between the two modalities by using two different measures. In the following we describe the patient’s MRI image with I and one of the k CT atlases with J^k . The voxels of both data sets are referred by an index i . Fol-

lowing the method proposed in [3], the normalised cross correlation (NCC) can be defined as

$$\text{NCC}_i^k = \frac{\langle I, J^k \rangle_i}{\sigma_i(I)\sigma_i(J^k)}. \quad (1)$$

Here, a local definition of the normalized cross correlation (LNCC) is used where the averages are calculated by means of a convolution with Gaussian kernel. If we denote the Gaussian kernel with G , the mean values μ can be determined with

$$\mu_i(I) = G * I \quad \text{and} \quad \mu_i(I \cdot J^k) = G * (I \cdot J^k). \quad (2)$$

Thus, the local covariance follows with

$$\langle I, J^k \rangle_i = \mu_i(I \cdot J^k) - \mu_i(I) \cdot \mu_i(J^k) \quad (3)$$

and the local variance σ results in

$$\sigma_i(I) = \sqrt{\mu_i(I^2) - \mu_i(I)^2}. \quad (4)$$

The resulting similarity values are in the intervall between -1 and 1 . Because both positive and negative covariances may occur, we use the absolute value of LNCC, which is close to 1 for locations that are considered to be particularly similar even across modalities. After the similarity values are determined, the CT atlases can be ranked at each voxel. For this purpose, an array H_i^k corresponding to the size of the data set is extended by a dimension corresponding to the number of atlases. At each voxel all atlases are now assigned to a particular rank. That atlas, which has the greatest similarity to the patient's MRI data set, receives the lowest rank. The resulting ranking array can be used directly for the calculation of the weights W_i^k . In [3] the weights are calculated with

$$W_i^k = \begin{cases} e^{-\beta H_i^k} & \text{if } H_i^k \leq \kappa \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

wherein the parameter β controls the exponential decay and the parameter κ specifies the number of atlases to be retained. Finally the pseudo-CT is synthesized by a weighted averaging of the CT atlases.

In addition to the LNCC, another measure for assessing similarity between different modalities was tested in this work. In our opinion, a perhaps more appropriate way to determine the similarity is the use of the image descriptor MIND, since they achieve good results for challenging multimodal registration tasks [7]. In a preprocessing step, the descriptors of each atlas and the patient MRI has to be pre-calculated independently. In contrast to the use of LNCC the similarity can no longer be determined directly based on the image data, but rather by the distances of the descriptors. As described in [7], a descriptor of an image X can be calculated with the following Gaussian function:

$$\text{MIND}(X, G_P, \mathbf{r}) = \frac{1}{z} \exp\left(-\frac{D_P(X, G_P, \mathbf{r})}{V(X, G_P)}\right), \quad (6)$$

wherein z is a normalization constant to restrict the descriptor entries to the intervall of 0 and 1 . The numerator comprises patch-based distances D_P , which can be easily calculated with the aid of convolution with a certain Gaussian kernel G_P . The denominator consist of an estimation of the local variance V . The variable \mathbf{r} fixes the neighbourhood within the descriptor is determined. Now the patch-based sum of absolute differences of the MIND descriptors can be used to determine local similarity and thereby the local ranking of atlases. The patch size is equivalent to the calculation of LNCC.

The weights, which have been obtained from these two described methods were also used for label fusion. We were guided out of the method in [21]. In addition, the MIND descriptors of the CT atlases and the patient MRI were used in a different way for label fusion. Inspired by [10], the descriptors were used to determine weights for the label fusion using multimodal regression. The descriptor patches of the CT atlases are arranged in the matrix A , where the n -th column represents the patch of the atlas n . The corresponding descriptor patch of the patient's MRI data set is considered as a column vector b . By solving the non-negative least squares problem

$$\arg \min_c \|Ac - b\|_2^2 \quad \text{subject to } c_i \geq 0 \quad \forall i, \quad (7)$$

the weights c for the label fusion can be determined. In contrast to a weighting based on similarity ranks, this approach estimates the best linear combination of patches of MIND descriptors from all atlas CT scans that best resemble the local MIND patch in the patients MRI scan.

5. Experimental Evaluation

The presented methods were tested on 3D-CT and 3D-MRI whole body scans of the Visceral Concept Extraction Challenge in Radiotherapy [6]. A total of five different T1-weighted MRI scans were examined. For each of the MRI images 18 till 19 CT atlases were available. In addition to the gray scale images there were manual expert segmentations of 12 different organ structures for each of the MRI and CT images. For three of the five MRI data sets the corresponding CT image of the same subject was present.

The parameters for the image registration were determined empirically. For the number of grid points we set $g = [9, 8, 7, 6, 5]$. The search radius was fixed to $l_{max} = [7, 6, 5, 4, 3]$ with a label quantisation of $Q = [5, 4, 3, 2, 1]$ and five levels were adopted for the multi level method. As in [3] we choose $\beta = 0.5$. A variation of

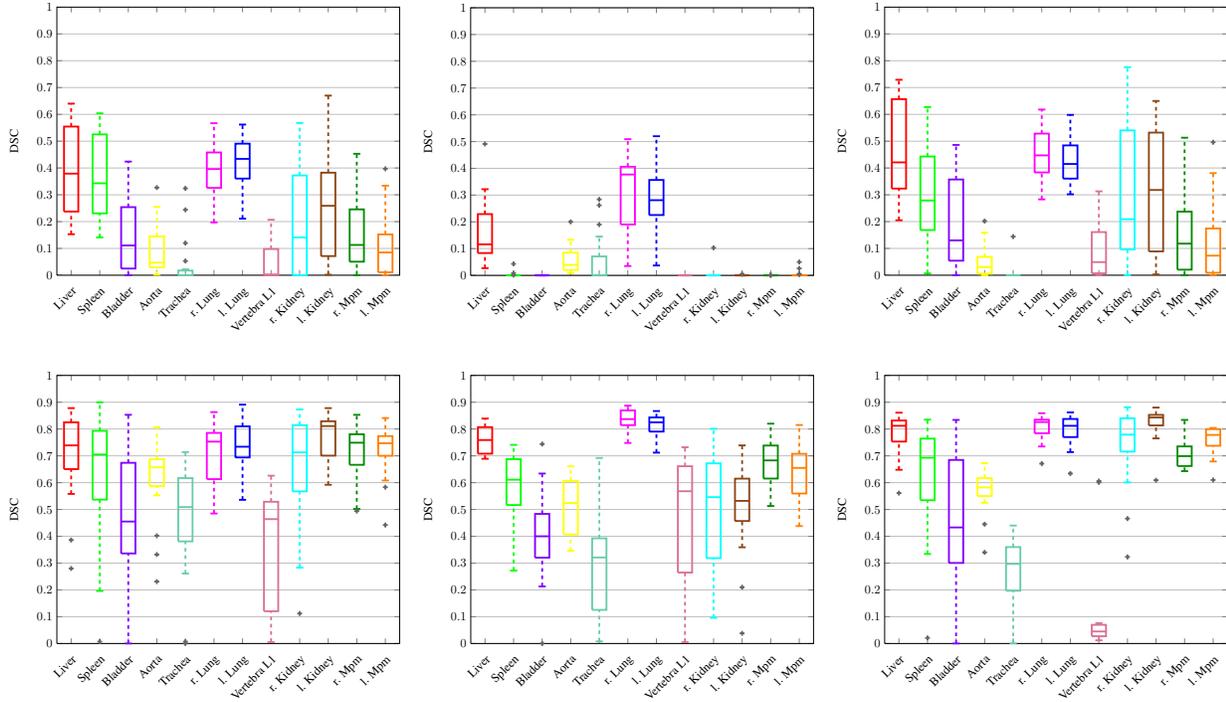


Figure 2: Evaluation of the Dice similarity coefficient before (top) and after (bottom) the multimodal image registration for three different patient-MRIs (MR 3, MR 12 and MR 15: left to right). Especially for small structures such as the trachea, the bladder and the vertebra L1 registration inaccuracies remain.

the parameter κ does not lead to any appreciable changes in the results, therefore we have set κ to the maximum value (number of CT atlases). The patch for the estimation of local similarities had a size of $9 \times 9 \times 9$ voxels.

Figure 2 shows the results of the multimodal image registration for three of the five studied data sets. For the evaluation of the registration quality, a comparison of the deformed atlas segmentations and the patient’s segmentation was done using the Dice similarity coefficient. The results for all 12 organ structures are summarized in box plots. The overall good quality of the multimodal registration method is clearly demonstrated. But for smaller organ structures such as the the trachea, the bladder and especially the vertebra L1, registration errors may still prevail.

The results of the pseudo-CT synthesis with LNCC and MIND as similarity measures are shown in Figures 3 and 4 by way of example for two different subjects. For the results of the top rows, all available CT atlases were considered within the synthesis. Somewhat refined results are shown in the bottom row, where only the top five atlases were taken into account. These were automatically selected by comparing the Dice overlap of any individual atlas with the majority vote (MV). The results of MR 3 in Figure 3 show undesirable intensity distributions in areas of large inter-

subject variability. These are mainly due to large residual registration errors of the multimodal registration of the CT atlases to the patient’s MRI. Slightly better results could be obtained for the data set MR 15 in Figure 4. For this subject, the registration method worked better, as already the results from the box plot on the bottom right in Figure 2 show. In both cases it is observed that the pseudo-CT synthesis with MIND achieved better results as compared to LNCC or (an unweighted) averaging.

In our further results, we compare four different approaches for local label fusion using the Dice similarity coefficient of 12 anatomical labels. First, a refined majority vote, which was obtained by selecting the five atlases that were most similar to an unweighted majority vote of all atlases (similar as done in SIMPLE [16]). Second, a weighted local fusion based on the rank of the patch-wise local normalised cross-correlation as done in [2]. Third, the weighting based on the rank of MIND similarity. Fourth, the use of multimodal regression using MIND descriptors (abbreviated with *Regr*). The results are further set in comparison to results obtained from registering a patient-specific CT scan to the MRI of the same subject. Note, that this scan would not be available in our envisaged application scenario. It does, however, highlight the challenges, which would oc-

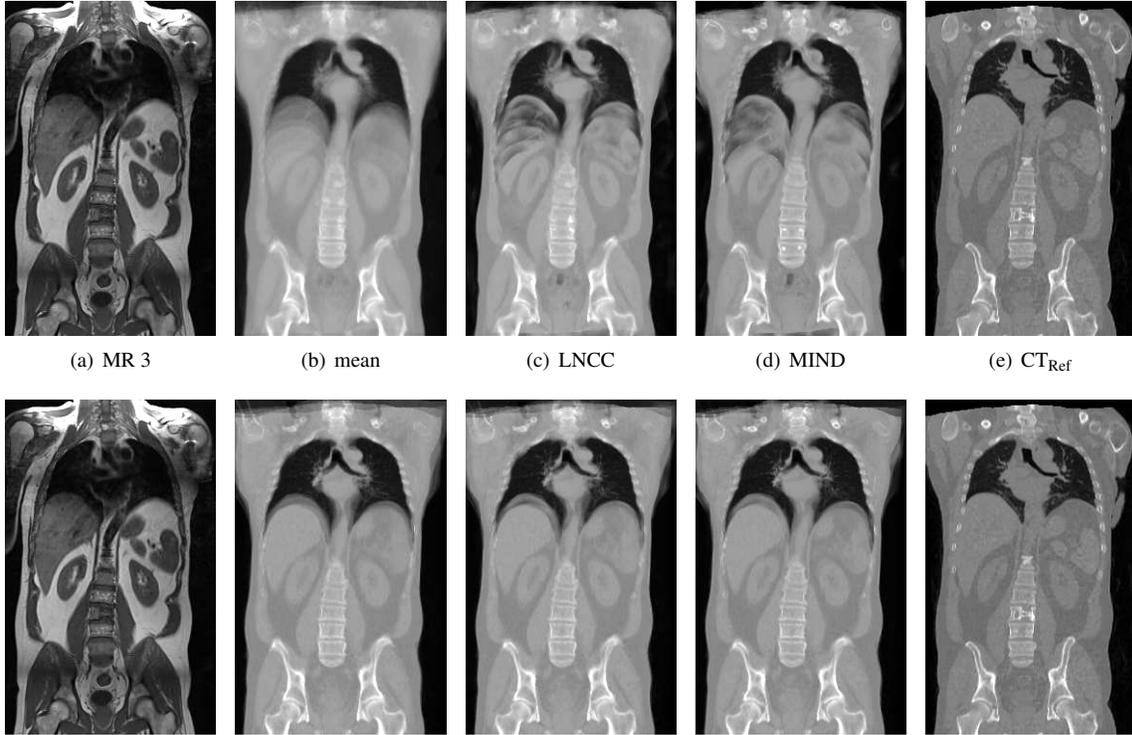


Figure 3: Results of the intensity fusion for MR 3. For the top row all available atlases and for the bottom row only the top five atlases in comparison to the MV are considered. In addition to the synthesis results, the registered reference CT (CT_{ref}) of the patient is shown. In regions of large inter-subject variability, undesired intensity distributions occur due to registration inaccuracies.

MR12	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	avg
MV	0.827	0.673	0.545	0.642	0.344	0.875	0.852	0.705	0.678	0.595	0.790	0.753	0.690
LNCC	0.805	0.701	0.552	0.654	0.380	0.869	0.854	0.701	0.702	0.598	0.798	0.769	0.698
MIND	0.830	0.719	0.551	0.666	0.372	0.882	0.856	0.718	0.688	0.594	0.793	0.766	0.703
Regr	0.777	0.663	0.568	0.654	0.362	0.876	0.853	0.726	0.689	0.607	0.790	0.755	0.693
CT_{ref}	0.698	0.424	0.744	0.659	0.483	0.770	0.759	0.645	0.561	0.532	0.820	0.815	0.659
MR3	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	avg
MV	0.855	0.802	0.607	0.731	0.654	0.782	0.813	0.590	0.850	0.898	0.820	0.788	0.766
LNCC	0.849	0.823	0.607	0.756	0.715	0.776	0.814	0.588	0.876	0.905	0.822	0.796	0.777
MIND	0.859	0.808	0.615	0.760	0.722	0.796	0.824	0.614	0.866	0.900	0.825	0.792	0.782
Regr	0.841	0.786	0.572	0.745	0.717	0.783	0.823	0.616	0.865	0.898	0.827	0.790	0.772
CT_{ref}	0.878	0.900	0.828	0.807	0.714	0.863	0.891	0.626	0.873	0.854	0.853	0.841	0.827

Table 1: Overview of the results of the label fusion for two different patient-MRIs (MR12 and MR3). The results of the four different approaches are listed individually for each of the 12 label (L1 to L12; the order corresponds to those of the box plots).

cur *e.g.* for MRI-guided radiotherapy even if CT imaging would be performed to obtain attenuation values for dose calculation. As evident from Table 1 the latter approach

does not necessarily yield the best outcome: in particular the spleen and the vertebra L1 of MR12 result in inferior overlap compared to all label fusion approaches. Finding

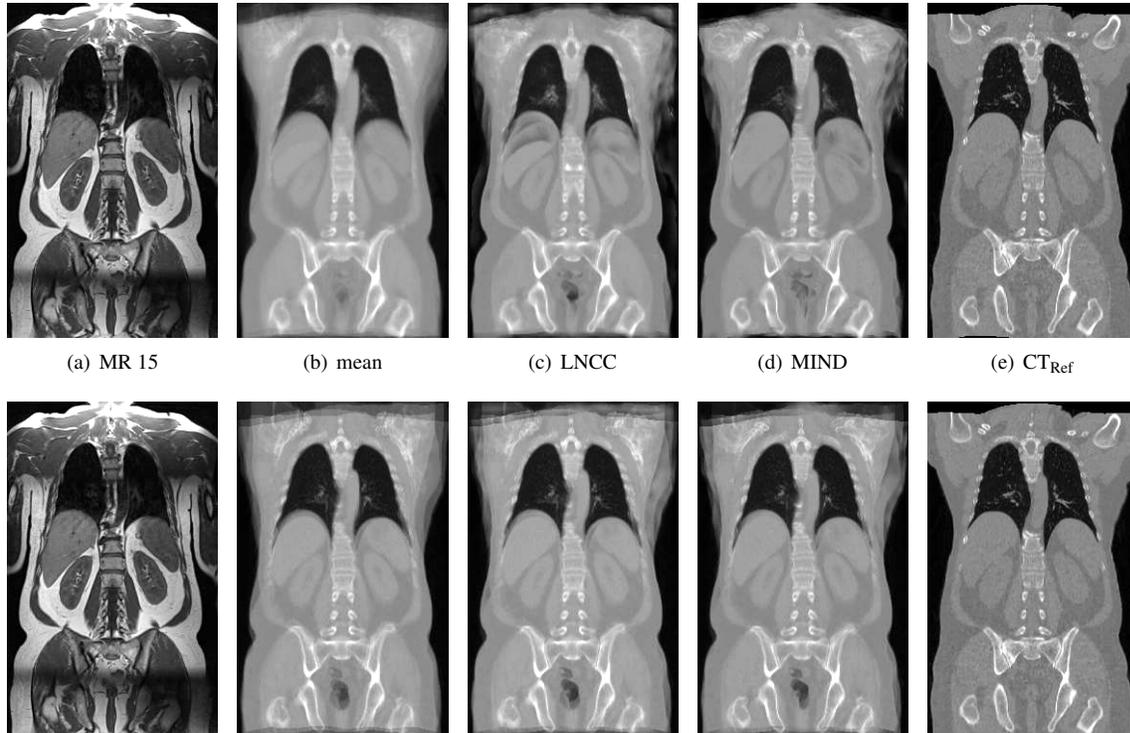


Figure 4: Results of the intensity fusion for MR 15. In the results from the top row all atlases have been considered. In the bottom row, only the five most similar atlases in comparison to the MV were considered. An improvement in the results through a more accurate registration can be seen.

the fusion weights by fitting a linear combination of MIND descriptors using regression, i.e. a generative approach that attempts to reconstruct the appearance of the target patch descriptors (see Eq. 7), yielded an unsatisfying outcome. The results for the MIND-based multi-atlas fusion are overall best showing a modest improvement over [2]. This may indicate that using the rank for determining fusion weights as originally proposed in [4] helps dealing with a potentially difficult choice of scaling parameters. While our approach achieves good alignment of 70% or more for most anatomical labels some errors remain. For a visual illustration of the results of our label fusion approaches, see Figure 5. There the results are shown as an overlay of the patient’s MRI data set. The exact influence of our accuracy on PET reconstruction and dose planning remains to be studied in future work.

6. Discussion

While performing the experiments, we observed that the multimodal registration is not well succeeded the same for all CT atlases. Especially larger deformations between the modalities state a problem, which could not be completely overcome by the used registration method. Even in the con-

text of intra-subject registration unsatisfactory results could be found (see average DSC of CT_{ref} for MR12 in Table 1 and the registered segmentation (CT_{ref}) in Figure 5). Another persistent problem is the evaluation of the results. An intensity-based comparison of pseudo-CTs and actually acquired CTs still lacking. In addition to an exact positioning of anatomical structures, it requires an assessment whether the synthesized intensities are alike realistic intensities. In the end no statement to what extent the atlas-based pseudo-CTs are superior to those results from machine learning techniques and paired training data sets can be made.

7. Conclusion

We have presented a new and very general approach for the synthesis of pseudo-CT images in a multimodal fashion. It combines efficient multimodal image registration and a new local atlas fusion based on self-similarity descriptors. We were able to tackle the very challenging tasks of extending modality synthesis to whole-body scans without requiring well-aligned pairs of CT and MRI scans from same subjects. The proposed method compares well with previous work and achieves modest improvements for both for inten-

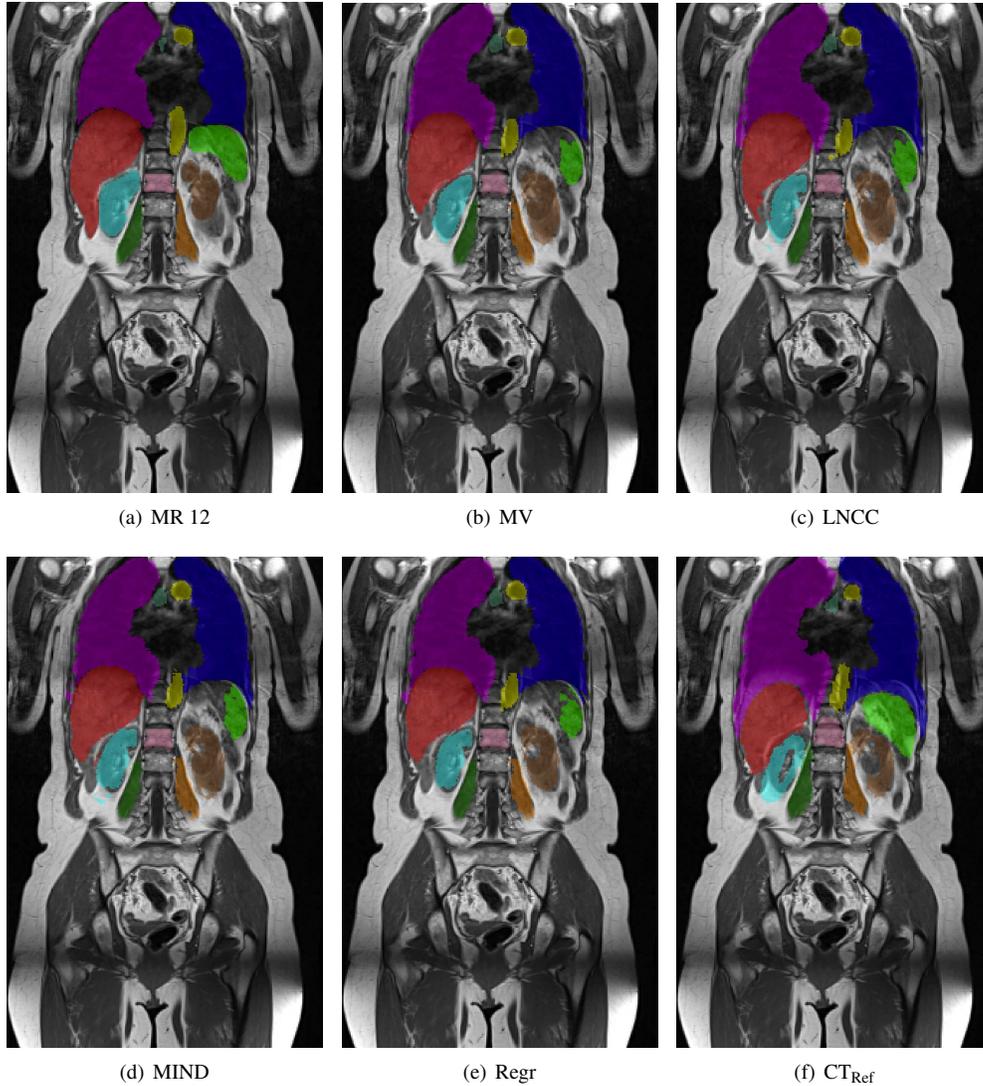


Figure 5: Comparative view of the results of the label fusion for MR 12. Again only the five most similar atlases compared to the majority vote were considered.

sity and label fusion. It also outperformed a simple majority vote over all registered atlases. In some cases, our approach was even superior to the direct registration of a patient-specific CT scan, which would normally not be available in MRI-only radiotherapy or PET-MRI hybrid scanners. However, some difficulties remain for regions, which are locally misaligned due to high variability between different subjects and where at the same time the structural image content does not provide enough information to select the best atlas(es) based on morphological similarity. Future work should focus on extending the contextual information captured by the image descriptors to enable a better treatment of ambiguous areas. The impact of modality synthesis on

image reconstruction and dose planning also requires further studies.

Acknowledgements

This work was supported by the Graduate School for Computing in Medicine and Life Sciences funded by Germany's Excellence Initiative [DFG GSC 235/2].

References

- [1] Y. Berker, J. Franke, A. Salomon, M. Palmowski, H. C. W. Donker, Y. Temur, F. M. Mottaghy, C. Kuhl, D. Izquierdo-Garcia, Z. A. Fayad, et al. Mri-based attenuation correc-

- tion for hybrid pet/mri systems: a 4-class tissue segmentation technique using a combined ultrashort-echo-time/dixon mri sequence. *Journal of Nuclear Medicine*, 53(5):796–804, 2012.
- [2] N. Burgos, M. J. Cardoso, F. Guerreiro, C. Veiga, M. Modat, J. McClelland, A.-C. Knopf, S. Punwani, D. Atkinson, S. R. Arridge, et al. Robust ct synthesis for radiotherapy planning: Application to the head and neck region. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, pages 476–484. Springer, 2015.
- [3] N. Burgos, M. J. Cardoso, M. Modat, S. Pedemonte, J. Dickson, A. Barnes, J. S. Duncan, D. Atkinson, S. R. Arridge, B. F. Hutton, et al. Attenuation correction synthesis for hybrid pet-mr scanners. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2013*, pages 147–154. Springer, 2013.
- [4] M. J. Cardoso, K. Leung, M. Modat, S. Keihaninejad, D. Cash, J. Barnes, N. C. Fox, S. Ourselin, A. D. N. Initiative, et al. Steps: Similarity and truth estimation for propagated segmentations and its application to hippocampal segmentation and brain parcellation. *Medical Image Analysis*, 17(6):671–684, 2013.
- [5] A. Guimond, A. Roche, N. Ayache, and J. Meunier. Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections. *IEEE Transactions on Medical Imaging*, 20(1):58–69, 2001.
- [6] A. Hanbury, H. Müller, G. Langs, M. A. Weber, B. H. Menze, and T. S. Fernandez. Bringing the algorithms to the data: cloud-based benchmarking for medical image analysis. In *Information Access Evaluation. Multilinguality, Multimodality, and Visual Analytics*, pages 24–29. Springer, 2012.
- [7] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Matin, F. V. Gleeson, M. Brady, and J. A. Schnabel. Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Medical Image Analysis*, 16(7):1423–1435, 2012.
- [8] M. P. Heinrich, M. Jenkinson, M. Brady, J. Schnabel, et al. Mrf-based deformable registration and ventilation estimation of lung ct. *IEEE Transactions on Medical Imaging*, 32(7):1239–1248, 2013.
- [9] M. P. Heinrich, M. Jenkinson, B. W. Papież, M. Brady, and J. A. Schnabel. Towards realtime multimodal fusion for image-guided interventions using self-similarities. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2013*, pages 187–194. Springer, 2013.
- [10] M. P. Heinrich, M. Wilms, and H. Handels. Multi-atlas segmentation using patch-based joint label fusion with non-negative least squares regression. In *Patch-Based Techniques in Medical Imaging*, pages 146–153. Springer, 2015.
- [11] M. Hofmann, F. Steinke, V. Scheel, G. Charpiat, J. Farquhar, P. Aschoff, M. Brady, B. Schölkopf, and B. J. Pichler. Mri-based attenuation correction for pet/mri: a novel approach combining pattern recognition and atlas registration. *Journal of Nuclear Medicine*, 49(11):1875–1883, 2008.
- [12] J. E. Iglesias, E. Konukoglu, D. Zikic, B. Glocker, K. V. Leemput, and B. Fischl. Is synthesizing mri contrast useful for inter-modality analysis? In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pages 631–638. Springer, 2013.
- [13] A. Jog, S. Roy, A. Carass, and J. L. Prince. Magnetic resonance image synthesis through patch regression. In *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*, pages 350–353. IEEE, 2013.
- [14] A. Johansson, M. Karlsson, and T. Nyholm. Ct substitute derived from mri sequences with ultrashort echo time. *Medical physics*, 38(5):2708–2714, 2011.
- [15] V. Keereman, Y. Fierens, T. Broux, Y. De Deene, M. Lonneux, and S. Vandenberghe. Mri-based attenuation correction for pet/mri using ultrashort echo time sequences. *Journal of nuclear medicine*, 51(5):812–818, 2010.
- [16] T. R. Langerak, U. A. Van Der Heide, A. N. T. J. Kotte, M. A. Viergever, M. Van Vulpen, and J. P. W. Pluim. Label fusion in atlas-based segmentation using selective and iterative method for performance level estimation (simple). *Medical Imaging*, 29(12):2000–2008, 2010.
- [17] R. Le Goff-Rougetet, V. Frouin, J. Mangin, and B. Bendriem. Segmented mr images for brain attenuation correction in pet. In *Medical Imaging 1994*, volume 2167, pages 725–736. International Society for Optics and Photonics, 1994.
- [18] A. Martinez-Mller, M. Souvatzoglou, G. Delso, R. A. Bundschuh, C. Ched’hotel, S. I. Ziegler, N. Navab, M. Schwaiger, and S. G. Nekolla. Tissue classification as a potential approach for attenuation correction in whole-body pet/mri: evaluation with pet/ct data. *Journal of nuclear medicine*, 50(4):520–526, 2009.
- [19] E. Schreiber, J. A. Nye, D. M. Schuster, D. R. Martin, J. Votaw, and T. Fox. Mr-based attenuation correction for hybrid pet-mr brain imaging systems using deformable image registration. *Medical physics*, 37(5):2101–2109, 2010.
- [20] G. van Tulder and M. de Bruijne. Why does synthesized data improve multi-sequence classification? In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, pages 531–538. Springer, 2015.
- [21] H. Wang, J. W. Suh, S. R. Das, J. B. Pluta, C. Craige, and P. A. Yushkevich. Multi atlas segmentation with joint label fusion. *Pattern Analysis and Machine Intelligence, IEEE Transaction on*, 35(3):611–623, 2013.
- [22] D. H. Ye, D. Zikic, B. Glocker, A. Criminisi, and E. Konukoglu. Modality propagation: coherent synthesis of subject-specific scans with data-driven regularization. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013*, pages 606–613. Springer, 2013.
- [23] H. Zaidi, M. Montandon, and D. O. Slosman. Magnetic resonance imaging-guided attenuation and scatter corrections in three-dimensional brain positron emission tomography. *Medical physics*, 30(5):937–948, 2003.