Non-rigid registration with position dependent rigidity for whole body PET follow-up studies

Kurt De Moor, Johan Nuyts, Lily Plessers, Sigrid Stroobants, Frederik Maes, and Patrick Dupont

Abstract— **A constrained non-rigid registration method is proposed for intra-patient registration of whole body images. Both the deformation field and the constraint operate on individual voxels. The deformation is implemented as a displacement vector for every voxel. The constraint is modeled as a position dependent non-linear resistance against deformation, using a spring model. This allows us to assign a different rigidity to different pairs of neighboring voxels, based on a segmented image, assuming that such an image is available.**

As a measure of image alignment, either mutual information or the sum of squared differences is used. The use of mutual information makes the software more widely applicable, while the sum of squared differences is expected to perform slightly better for intra-modality applications.

The method has been tuned for registration of ^{18}F FDG PET whole body images of the same patient in oncology. Thresholding is applied to segment regions with high tracer uptake. These regions are assigned a high degree of rigidity. As a result, the volume of lesions (and also of the bladder and the brain) will not be changed during registration. After registration, subtraction images are computed, which are expected to be helpful for the evaluation of treatment response. A first evaluation indicates that the registration reduces the time required for visual and quantitative analysis of the images.

I. Introduction

POSITRON emission tomography (PET) is being used increasingly in cancer treatment follow-up. PET whole body images acquired before and after therapy are compared to quantify changes in tracer uptake in existing lesions and to detect possible new lesions in the case of disease progression. To date, this involves a mental registration of the lesions in the two images, and manual or (semi) automatic creation of regions (ROI) or volumes (VOI) of interest. This process can be time consuming and difficult when multiple lesions are present, and the results are operator dependent.

Similar problems have been encountered in neurological imaging with PET, SPECT and/or MRI, where the creation of subtraction images after accurate rigid registration was found to be very helpful [1]. The application of a similar approach in whole body imaging is complicated by the need for non-rigid registration [2], [3]. An elastic registration is required, because the patient will lie in a different position in subsequent scans. However, the registration should not alter the shape or the intensity of the lesions, because physicians want to quantify the evolution of both the volume and the tracer uptake.

Work supported by F.W.O. grant G.0174.03.

KDM, JN, LP, SS and PD are with the department of Nuclear Medicine, K.U.Leuven, B3000 Leuven, Belgium; FM is with Medical Image Computing, K.U.Leuven, B-3000 Leuven, Belgium. E-mail: Johan.Nuyts@uz.kuleuven.ac.be

In this paper, we propose a non-rigid registration algorithm dedicated to whole body images of the same patient. The main difference with most existing methods is that different degrees of rigidity can be assigned to different voxels. By assigning high rigidity to lesions, one obtains a non-rigid registration of the patient body, combined with a (nearly) rigid registration of the lesions. Similarly, a high rigidity can be assigned to rigid structures such as bone, provided that an appropriate segmentation is available. In PET-PET registration, we use this feature to avoid deformation of the lesions and the brain. It is expected that this feature will be useful for PET-CT or CT-CT registration, where high rigidity can be assigned to bony structures.

The method is treated as an optimization problem, where the objective function is the combination of a similarity measure (favoring spatial alignment) and a penalty that discourages (excessive) deformations. With intra-patient registration, small deformations are expected, whereas very strong deformations are highly unlikely. For that reason, we have chosen for a penalty that increases faster than linearly with increasing deformation.

To obtain maximum flexibility in assigning local rigidity, both the deformation field and the penalty are implemented on a voxel basis. The deformation is coded as the 3D displacement of every voxel, and the penalty is a function of the distance between neighboring voxels.

Two similarity measures have been implemented: mutual information and the sum of squared differences. With mutual information [4], the method can be applied for inter-modality registration (e.g. PET-CT) and in intra-modality applications where the image intensities vary considerably between subsequent scans (e.g. different tracers or rapidly changing tracer uptake in PET studies). For PET-PET registration in oncology, the sum of squared differences might be more suitable, in particular because the accurate registration of regions with high intensity is requested.

The optimization is done with a gradient ascent algorithm. For that purpose, analytical expressions for the gradients of the similarity measure and the penalty are used.

II. Theory

Non-rigid registration deals with finding a transformation τ, that maps every voxel j from image F, the floating image, to a voxel $\tau(j) = \tau_j$ from image R, the reference image. To penalize inter-voxel distance changes, a function C has been chosen that models the energy of springs interconnecting the voxels. The objective function that will be maximized is

written as $O = I - \beta C$, where β is the weight of the spring cost and I is the similarity measure.

A. Mutual information

Mutual information has proved to be a reliable similarity measure in medical imaging [5]. The mutual information can be expressed in terms of the joint and marginal histograms h_{rf}, h_r, and h_f with the following relationship [4]:

$$I_1(R,F) = \sum_r \sum_f \left[\frac{h_{rf}}{N} \ln \frac{N.h_{rf}}{h_r.h_f} \right] \tag{1}$$

The histogram value h_{rf} equals the number of voxel pairs (j, τ_j) satisfying $F(j) = f$ and $R(\tau_j) = r$, N is the total number of voxel pairs, $h_r = \sum_f h_{rf}$ and $h_f = \sum_r h_{rf}$. The use of histograms assumes that the voxel intensities of both images have been binned into a finite number of bins r and f.

To ensure that the joint histogram varies smoothly with small changes of transformation τ, it is calculated with partial volume interpolation [4]. Because the histogram is directly computed from the (transformed) coordinates and the original image, no intermediate images have to be computed. Setting $\tau_j = (x_j, y_j, z_j)$, the gradient of I_1 in every voxel j is given by:

$$\frac{dI_1}{dx_j} = \sum_{r,f} \frac{dI_1}{dh_{rf}} \frac{dh_{rf}}{dx_j} = \sum_{r,f} \frac{1}{N} \left(\ln \frac{N h_{rf}}{h_r h_f} - I_1 \right) \frac{dh_{rf}}{dx_j} \tag{2}$$

Here $\frac{dh_{rf}}{dx_j}$ is a histogram with the same dimensions as the joint histogram h_{rf}. This histogram is very sparse, because only the entries near $(r, f) = (R(\tau_j), F(j))$ are affected. The derivatives in the y_j and z_j direction can be calculated in the same way.

B. Sum of squared differences

The sum of squared differences penalty is computed as

$$I_2(R,F) = -\frac{1}{2N} \sum_j (R(\tau_j) - F(j))^2 \tag{3}$$

and the gradient equals

$$\frac{dI_2}{dx_j} = -\frac{1}{N} \sum_j (R(\tau_j) - F(j)) \frac{dR(\tau_j)}{dx_j}. \tag{4}$$

The values of $R(\tau_j)$ and $dR(\tau_j)/dx_j$ are computed with linear interpolation.

C. Spring cost

The rigidity of the image is modeled by assuming that all voxels are connected to their neighbors by little rods or springs (fig. 1). Modifying the rigidity of these springs modifies the local rigidity of the image. The springs are given two parameters: the stiffness of the spring for small deformations, and the maximum relative deformation (either extension or compression) of the spring. To impose a maximum relative deformation, the stiffness is increased for increasing deformations, approaching

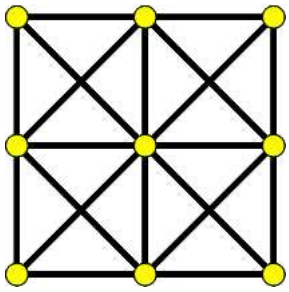

Fig. 1. Each voxel is connected to its neighboring voxels with little rods. If these rods are elastic springs, the resulting structure is deformable. If the rods are rigid, the structure is rigid as well.

$\pm\infty$ when the relative deformation approaches the predefined maximum.

To obtain this behavior, we propose the following non-linear function for the derivative of the spring cost function C:

$$\frac{dC}{dx_j} = \frac{2}{N\pi} \sum_{i=1}^{26} w_{ij}.B_{ij} \tan\left(\frac{\pi}{2.B_{ij}}.X_{ij} \right).\frac{dX_{ij}}{dx_j}, \tag{5}$$

and similar for dC/dy_j and dC/dz_j. $X_{ij} = (L_{ij} - L^0_{ij})/L^0_{ij}$, where L^0_{ij} is the distance between the neighboring voxels i and j in the undeformed image, and L_{ij} is the corresponding distance between τ_i and τ_j after deformation. The parameters w_{ij} and B_{ij} determine resp. the stiffness of the spring (for small deformations) and its maximum deformation. These parameters can be used to incorporate anatomical knowledge. Assuming all voxels have been attributed to an anatomical class (e.g. bone, soft tissue ...), the w_{ij} and B_{ij} can be adjusted to reflect the different physical properties of the classes. For a mixed pair the most rigid of the two classes is selected (maximum w_{ij} and minimum B_{ij}). The sum over i takes into account the neighboring voxels in a window centered around voxel j. We use the 26 nearest neighbors of voxel j, but the algorithm allows other configurations if necessary. A tangent function is chosen, because it is symmetrical around relative deformation zero, nearly linear for small deformations (with a slope of w_{ij}) and growing very quickly in the neighborhood of $\pi/2$, or X_{ij} close to $\pm B_{ij}$. Parameter B_{ij} can be considered as the maximum allowable relative deformation of the inter-voxel distance $i - j$. Figure 2 plots dC/dX_{ij} as a function of the relative deformation, for a few values of w_{ij} and B_{ij}. The penalty C is obtained by integrating (5), yielding:

$$C = -\frac{4}{N\pi^2} \sum_1^{26} w_{ij} B_{ij}^2 \ln(\cos(\frac{\pi}{2B_{ij}} X_{ij})). \tag{6}$$

The sum $I - \beta C$ can be computed to verify convergence.

D. Segmentation

In our application of whole body PET-PET registration, we currently consider three different regions with different rigidity: 1) lesions, which are characterized by high tracer uptake and which should be treated as rigid, 2) the rest of body, which may undergo non-rigid deformation and 3) the background, which is not clinically relevant and my be strongly deformed if necessary. Accordingly, very high rigidity

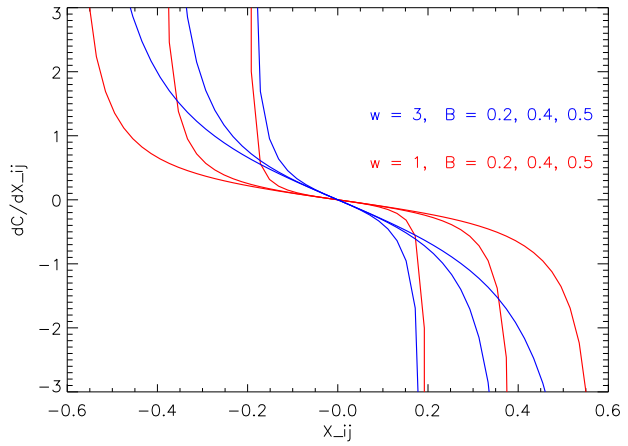

Fig. 2. The derivative of the spring cost function C, as a function of the relative deformation X_{ij}, for different values of w_{ij} and B_{ij}. B_{ij} defines the minimum and maximum deformation, w_{ij} determines the stiffness for small deformations.

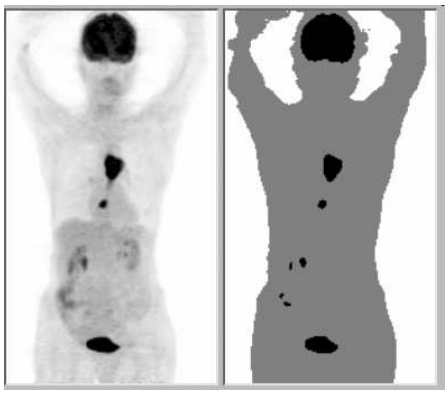

Fig. 3. Segmentation with simple thresholding to produce three classes: hot and rigid regions, non-rigid tissue and very non-rigid background.

was assigned to hot regions, an intermediate rigidity to the rest of the body, and very low rigidity to the background. Currently, simple thresholding with two thresholds was applied to identify these three classes. An example is shown in fig. 3. For ^{18}F-FDG images, a high threshold of standardized uptake value (SUV) = 3.5 was appropriate to identify all lesions and the brain, which were rightly treated as rigid objects. Also the bladder exceeded this threshold. Although the bladder shows considerable deformation, rigid registration of the bladder was not perceived as a problem by the clinicians. Obviously, a more sophisticated segmentation could be produced. In particular, when the images are acquired on a hybrid PET-CT system, the segmentation could use both the CT and the PET image.

E. Algorithm

The first step is a rigid registration (translation and rotation). Then non-rigid registration is applied in a multi-resolution scheme. Images of reduced resolution are computed by summing neighboring voxels, simple linear interpolation is used to transfer the deformation to a higher resolution level. The multi-resolution approach reduces the computation time and allows for relatively large displacements based on less detailed information.

In each iteration, the gradient ascent algorithm computes the displacement update Δu, which is added to the deformation field, using

$$\Delta u^{(k)} = S^{(k)}(\nabla I - \beta \nabla C)^{(k)} \tag{7}$$

where the superscript (k) denotes iteration k, ∇ denotes the gradient, and $S^{(k)}$ is the voxel dependent step size used to control convergence. The update Δu is always limited to $\pm$ 0.02 voxels per iteration step, whatever the result of the calculation. This is a rather small value, but we found that larger values may lead to oscillations, probably because of the non-linear spring stiffness. Because of the multi-resolution approach, a small per-iteration increment can still lead to large deformations in a reasonable amount of iterations. S is adjusted after every iteration: the value $S_j^{(k)}$ for voxel j is increased if the angle between the vectors $\Delta u_j^{(k)}$ and $\Delta u_j^{(k-1)}$ is small, and $S_j^{(k)}$ is decreased otherwise. For each resolution level, the calculations are repeated until the objective function changes less than a preset value, or until a maximum number of iterations has been reached.

When the rigidity is very high, the gradient of C dominates over that of I in (7) and as a result, rigid structures tend to move very slowly in response to ∇I. To accelerate the process, the registration at each resolution level is started with reduced rigidity (temporarily increasing the influence of ∇I), and completed with the actual high rigidity (undoing possible non-rigid deformations of the rigid structures).

The final deformed image is calculated by applying the deformation field to the floating image. In doing this, in general a voxel of the floating image ends up between eight voxels in the resulting deformed floating image. The intensity of this voxel is then distributed over the eight neighbors with trilinear interpolation, thus preserving the total image intensity. In PET, the images are usually scaled to SUV, and the interpolation should preserve the SUV rather than the total image intensity. In that case, the deformed image is further corrected based on the amount of compression or expansion of the deformed voxels.

III. Experiments

First a simple 2D simulation experiment is shown to illustrate the algorithm, in particular the use of local rigidity to avoid changes in lesion size and shape.

For a first clinical validation, the algorithm was applied retrospectively to the pre-therapy and first post-therapy whole body ^{18}F-FDG PET images of 18 patients. These patients received palliative treatment for esophageal cancer. The images had already been analyzed manually; for this evaluation they were reanalyzed to assess the increase in efficiency and the accuracy of the lesion matching obtained with the registration method.

IV. Results

Currently, the computation time for two images of 128 $\times$ 128 $\times$ 200 voxels on a PC with Intel Xeon 3.80 GHz cpu

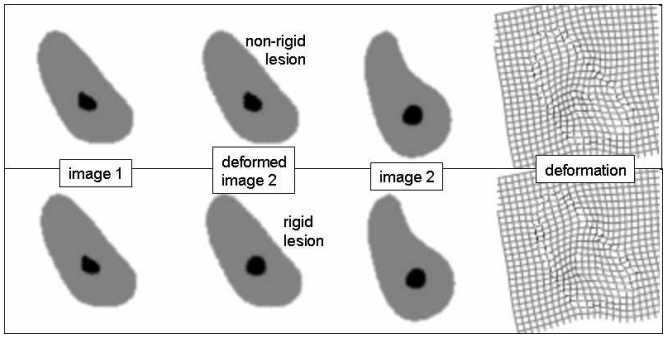

Fig. 4. Illustration of the local rigidity. The first row shows a registration result with a non-rigid "lesion". The second row shows a registration where a rigid "lesion" is surrounded by non-rigid tissue. After registration, the size and shape of the lesion are preserved.

is a few hours. However, the software was not optimized for speed, and we hope that the processing time can be reduced significantly. Because the registration requires no manual interaction, a computation time of an hour or less is considered acceptable for clinical use. The multi-resolution processing starts with coarse sampling obtained by summing $8 \times 8 \times 8$ voxels of the original image. However, most of the processing time is spent at the finest resolution, where the smallest displacements are contributed. The computation time is about half an hour if the registration is stopped at the coarser sampling of $64 \times 64 \times 100$ grid points. This produces image of good quality and was used in the results described below; a careful comparison is still to be done, to verify if the contribution at the finest resolution level has a noticeable effect on the registration quality.

The results shown below were obtained using the sum of squared differences. However, we found that very similar results could be obtained with mutual information.

Figure 4 shows the registration for a simulated object, once with a non-rigid lesion, and once with a rigid one. The deformation is visualized by applying it to an image of a square mesh.

Figure 5 shows an example from the patient study, illustrating the need for non-rigid registration. After registration, an excellent match of the indicated lesion is obtained. The isocontours also show a good alignment of the body contour, the brain, the lungs and vertebrae. The subtraction image shows changes in tracer uptake (white = reduction, black = increased uptake). The lesion is nearly invisible in the subtraction image, indicating that its SUV was hardly changed after therapy. The activity in the abdomen is unspecific tracer uptake in the bowel.

The clinician reported that the registration made the analysis significantly easier and faster, because in all cases the manual search for the corresponding lesion in the second image was eliminated by the registration. However, in two patients, the registration was considered less than ideal: the lesions in the two images only partially overlapped after registration. We believe this happened because the true local deformations can be slightly larger than the maximum deformation currently allowed by the spring function C. Figure 6 shows images from one of these two studies. Apparently, more local deformation

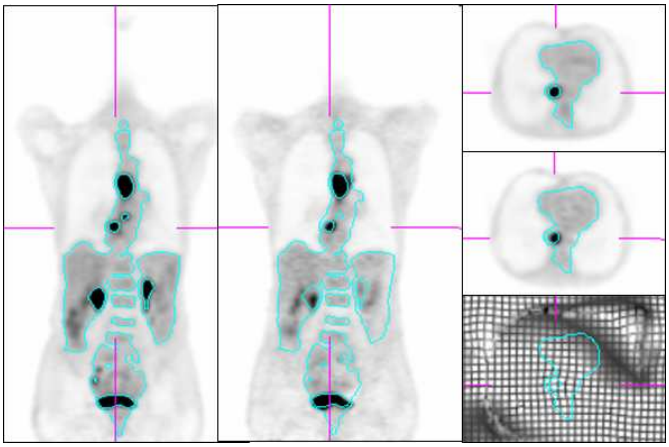

Fig. 6. Coronal and sagittal slices from the pre-therapy image and the registered post-therapy image. The isocontours were defined on the pre-therapy image (coronal image at the left, transaxial image at the top). Although the isocontours reveal good registration in the rest of the body, the lesion (at the center of the cross hairs) is not well aligned. The deformation image shows that considerable local deformation was already applied to obtain at least this partial overlap.

is required for perfect alignment of the lesions in both images. An alternative hypothesis, that the lesion has somehow evolved asymmetrically, was rejected after visual inspection of the corresponding CT-images.

V. Discussion

The use of position dependent rigidity has been studied by several groups. Little et al [6] combined rigid and non-rigid regions in a landmark based registration. More recently, Huesman et al [7] used rigid mutual information based registration to align rigid structures (bony structures in CT and MR scans), and landmark based registration to align the non-rigid tissue in between. These methods use different models for the rigid and non-rigid structures. The advantage is that the rigid objects can be represented efficiently (only specifying the position and orientation), which can lead to a computation efficient implementation. The disadvantage is that these methods cannot implement different degrees of rigidity. Claessens et al [8] proposed a pixel-based registration for 2D scintigraphic images, using a simple spring model to implement different degrees of position dependent rigidity. Loeckx et al [9] obtained different values for the local rigidity by constraining the local value of the Jacobian of the transformation.

The method of [8] is extended here to volume registration, using a somewhat more sophisticated non-linear spring model. Spring models have been used as a substitute for more accurate finite volume elements in tissue modeling algorithms [10]. In this kind of application, point masses and damping coefficients are assigned to the springs, in addition to the stiffness, in order to obtain realistic tissue motion in response to external forces. Because in our application, only the final position matters, there is no need to model a realistic motion. Consequently, the point masses and damping could be left out (although including them might lead to a viable optimization algorithm), and only the stiffness of the springs was modeled. By using a non-linear stiffness the maximum deformation of the springs

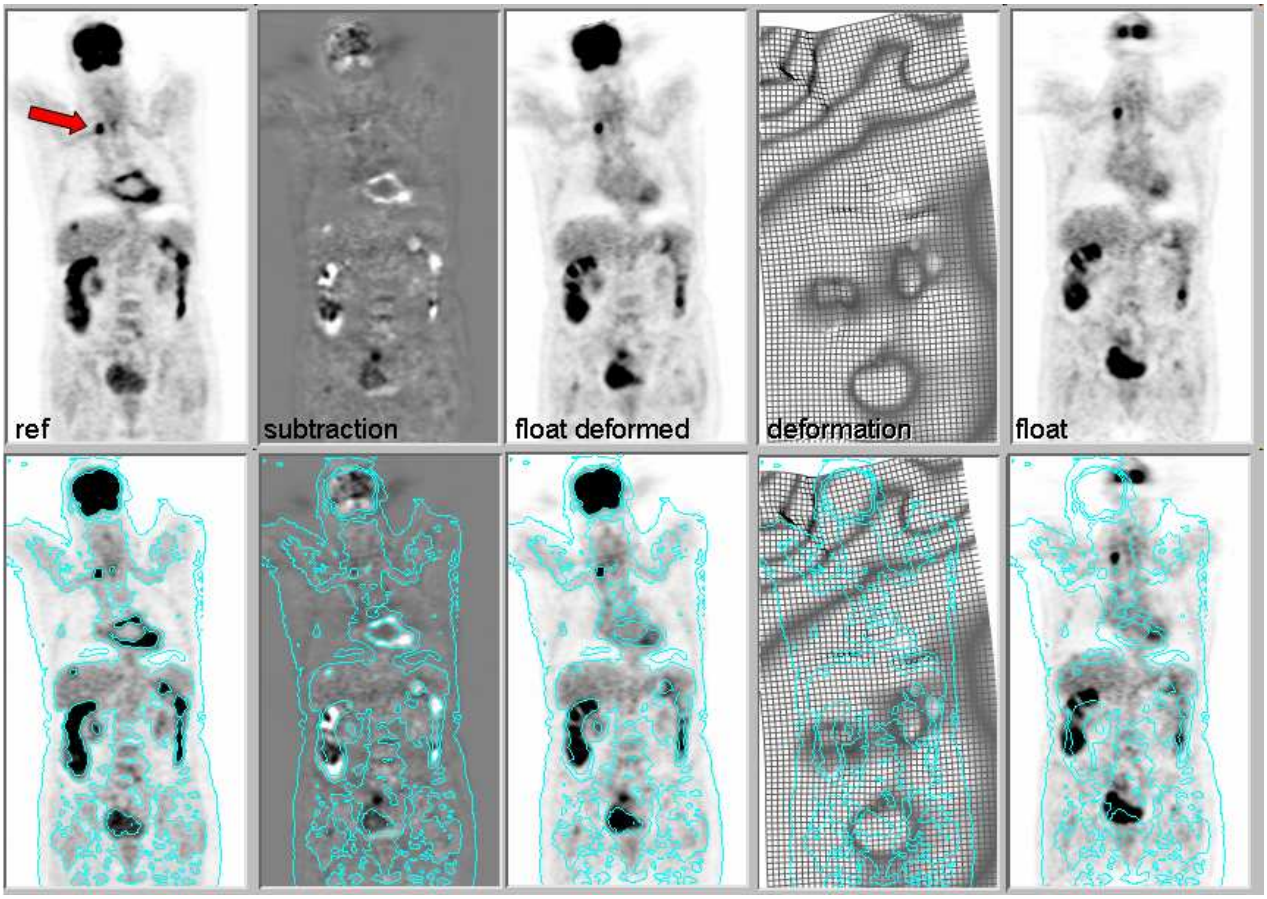

Fig. 5. PET whole body registration. First row from left to right: 1: pre-therapy image, 2: subtraction image, 3: deformed post-therapy image, 4: deformation applied to a mesh image, 5: original post-therapy image. The second row shows the same images, but now with isocontours defined on the pre-therapy image (the reference image) in overlay. The red arrow points to a lesion, which was well aligned after non-rigid registration.

was restricted. This not only reduces the possible set of solutions (eliminating unrealistically large deformations), it also ensures that the topology of the voxels is preserved. The advantage of a voxel-based method is that the implementation of local rigidity is simple and straightforward. A disadvantage is that the computational burden is proportional to the number of voxels in the images.

Based on the current evaluation, the parameters will be adjusted (allowing slightly larger deformations), and the method will be re-evaluated on an independent set of PET whole body images. Also the effect of the unavoidable interpolation on the quantitative analysis (e.g. mean and maximum SUV, total lesion glycolysis) will be studied.

VI. Conclusion

An algorithm for non-rigid registration with position dependent rigidity has been developed. The method was applied to a set of clinical PET whole body images: the registration was found to be accurate, and it made the therapy assessment easier and faster.

References

[1] TJ O'Brien, EL So, BP Mullan, MF Hauser, BH Brinkmann, NI Bohnen, D Hanson, GD Cascino, CR Jack, FW Sharbrough, "Subtraction ictal SPECT co-registered to MRI improves clinical usefulness of SPECT in localizing the surgical seizure focus." *Neurology*, vol 50, pp 445-454, 1998.

[2] R. K. Bajcsy and S. Kovacic, "Multi-resolution elastic matching," *Computer Vision, Graphics and Image Processing.*, vol 46, no 1, pp 1-21, 1989.

[3] G. Christensen, R. Rabbitt, and R. Miller, "Deformable templates using large deformation kinetics," *IEEE transactions on image processing.*, vol 5, no 10, pp. 1435-1447, 1996.

[4] F. Maes, A. Collignon, D. Vandermeulen et al. G. Marchal,and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE transactions on medical imaging.*, vol 16, no 2, pp. 187–198, April 1997.

[5] F Maes, D Vandermeulen, P Suetens. "Medical image registration using mutual information", *Proceedings of the IEEE*, vol 91, no 10, pp 1699-1722, 2003.

[6] JA Little, DLG Hill, DJ Hawkes. "Deformations incorporating rigid structures", *Comput Vision Image Understanding*, vol 66, no 2, pp 223-232, 1997.

[7] RH Huesman, GJ Klein, JA Kimdon, C Kuo, S Majumdar, "Deformable registration of multimodal data including rigid structures", *IEEE Trans Nucl Sci*, vol 50, no 3, pp 389-392, 2003.

[8] R Claessens, J Nuyts, S Stroobants, P Dupont, F Maes. "Non-rigid image registration for temporal subtraction of whole-body nuclear medicine images", *IEEE Nuclear Science Symposium Conference Record* Vol 5, pp 3173 - 3175 Vol.5, 2003

[9] D Loeckx, F Maes, D Vandermeulen, P Suetens. "Nonrigid image registration using free-form deformations with a local rigidity constraint", *Lecture notes in computer science (proceedings of MICCAI 2004*, vol. 3216, pp. 639-646, 2004

[10] H Delingette, "Toward realistic soft-tissue modeling in medical simulation", *Proceedings of the IEEE*, vol 86, no 3, pp 512-523, 1998.