SELECTIVE X-RAY RECONSTRUCTION AND REGISTRATION FOR POSE ESTIMATION IN 6 DEGREES OF FREEDOM

B. P. Selby [a, *], G. Sakas [b], S. Walter [a], W.-D. Groch [c], U. Stilla [d]

[a] Medcom GmbH, Rundeturmstr. 12, 64283 Darmstadt, Germany - pselby@medcom-online.de
[b] Fraunhofer IGD, Cognitive Computing and Medical Imaging, 64283 Darmstadt, Germany
[c] University of Applied Sciences, Fachbereich Informatik, 64295 Darmstadt, Germany
[d] Technische Universitaet Muenchen, Photogrammetry and Remote Sensing, Germany

KEY WORDS: Radiometry, Orientation, Detection, Rendering, Registration, Spatial, X-Ray, Accuracy

ABSTRACT:

Particle beams in radiological cancer treatment provide high accuracy in dose delivery. Thus approaches from image-guided radiotherapy (IGRT) are used to overcome accuracy limitations caused by the patient misalignment in the treatment device. By comparing stereoscopic X-ray images of the patient in treatment position to a reference Computed Tomography (CT) scan, a correction of the initial patient set-up can be computed. Automatic registration of the X-ray images with digital reconstructed radiographs (DRRs) from the CT and back-projection of the transformations gives a pose correction in 5 degrees of freedom (DOF). To obtain a 6 DOF correction, DRRs have to be generated for a large amount of hypothetical alignments to find the optimal match to the X-ray images. To accelerate this time consuming process and to reduce the disturbing influence of image contents that do not match correctly, we automatically exclude regions that may not improve the resulting pose correction from the rendering as well as from the matching process. Therefore these regions are identified in the X-ray images and transferred into the plane of the respective DRR. We then perform the radon transform for DRR generation only for a subset of possible pixel values and exclude the missing information from the registration process. As a result of this approach, the time needed for a full automatic pose correction computation in 6 DOF is decreased by means of 4 and more and additionally misregistrations caused by unsuitable image contents can be avoided.

KURZFASSUNG:

Die Verwendung von Partikelstrahlung in der radiologischen Krebsbehandlung ermöglicht eine sehr hohe Genauigkeit bezüglich der Dosisverteilung. Deshalb werden Verfahren aus der bildgestützten Radiotherapie (IGRT) eingesetzt, um Begrenzungen der Genauigkeit zu überwinden, welche durch die initiale Patientenlagerung zustande kommen. Eine Korrektur der initialen Patientenposition kann durch Vergleich von stereoskopisch aufgenommenen Röntgenbildern, die den Patienten in der aktuellen Lage zeigen, mit einer Vergleichs-Computertomographie (CT) Aufnahme ermittelt werden. Eine automatische Registrierung der echten Röntgenbilder mit aus der CT Aufnahme rekonstruierten Röntgenbildern (DRRs) und Rückprojektion der resultierenden Transformationen ergibt eine Lagekorrektur in 5 Freiheitsgraden (DOF). Um eine Korrektur in 6 Freiheitsgraden zu erhalten, müssen DRRs für eine große Anzahl an hypothetischen Lagerungen erzeugt werden um diejenige Lage zu finden, in welcher die Bilder am besten zu den Röntgenaufnahmen passen. Um diesen zeitintensiven Prozess zu beschleunigen und um den störenden Einfluss von Bildinhalten zu reduzieren, welche nicht in beiden Bildpaaren übereinstimmen, schließen wir bestimmte Bildregionen, welche die resultierende Patientenlage nicht verbessern, automatisch vom Rendering und dem Bildvergleich aus. Dazu werden die besagten Regionen in den Röntgenbildern identifiziert und in the Ebene der entsprechenden DRRs übertragen. Die Radon Transformation für die DRR Erzeugung wird nun nurnoch für einen Teilbereich des Gesamtbildes durchgeführt und die ausgeschlossene Information wird beim Rendering nicht mehr berücksichtigt. Im Ergebnis kann die Zeit die zum vollautomatischen Bestimmen der Patientenlage in 6 Freiheitsgraden benötigt wird um den Factor ¼ reduziert werden. Zusätzlich kann das Auftreten von Fehlregistrierungen, verursacht durch nicht zusammenpassende Bildregionen, verringert werden.

1. INTRODUCTION

Modern particle beam based cancer treatment methods allow highly accurate application of the treatment dose onto the carcinogen tissue. Today a rapid growing number of commercial health centers all over the world exploit the advances of particle beams. Due to the narrow dose maximum of the Bragg Peak particle beam based radiotherapy enables accurate delivery of the treatment dose onto the diseased tissue so that accuracies in the sub-millimeter domain are feasible even in depth of the patient body. However, to be able to achieve those accuracies a major issue is the exact alignment of the patient in the treatment machine (Verhey et al., 1982). Common strategies like tracking of external markers or fixation of the patients body do not suffice the requirement of high set-up precision and are not feasible for a variety of organs because of possible movements of the treatment target relative to the outer body shape.

It is common practice in image guided radiotherapy (IGRT) to align patients manually or semi-automatic in the treatment device. By visual evaluation of X-ray images and respective digital reconstructed radiographs (DRRs) the misalignment of the patient can be estimated (Thilmann et al., 2005; Yue et al., 2006). The Computed Tomography (CT) data recorded during the treatment-planning phase serves as reference position of the relevant body part. X-ray images acquired before treatment

from within the radiation device reflect the real patient alignment.

During the time consuming and error prone manual procedure the alignment of the respective body region may change, which leads to degradation of the treatment results. Furthermore, manual alignment cannot be done for six degrees of freedom (DOF), because rotational misalignments can hardly be detected and quantified accordingly by visual evaluation of the 2D images. For this reasons automatic patient alignment procedures are necessary.

For automatic alignment correction of patients in particle radiation treatment devices we register two stereoscopic digital radiographic images (DRs) taken from within the treatment device from different viewing angles with the respective projections of a high-resolution reference CT series, the planning CT. The projections are computed for an initial treatment set-up, starting with the expected patient alignment. The results of the rigid registration are then back-projected into the coordinate system of the patient table. The resulting correction vector in 6 degrees of freedom is used to move the table and to bring the patients tumor in the correct position for radiation treatment.

Two major problems occur when computing the patients pose in 6 DOF. One is that the rotation around the axis perpendicular to both central rays of the X-ray equipment axes (which is in most cases the table roll axis) cannot be computed directly from the 2D projections. This is because these rotations do not lead to detectable movements of the contents of the 2D images, but to implicit changes of the images, which cannot be interpreted by a registration process that relies purely on 2D images. The solution to this problem is to maximize the image similarity between the DRs and new projections of the CT, depending on the free parameters for the 6 degrees of freedom patient alignment (3 translations and 3 rotations). This approach implies a large number of CT projections, actually done by ray-tracing and therewith leads to high calculation times, even with optimized rendering techniques.

The other problem is that the comparison of the images suffers from image contents, that for example are present in the DRs, but not in the DRRs, e.g. parts of the patient fixation equipment.

To improve performance of the 6 DOF alignment detection and to reduce the influence of inherent deviations of the image contents on the registration process, we propose a modified approach, relying only on parts of the respective X-ray images and the CT scan. Therefore we initially perform a 5 DOF correction to gain a good estimation of the patient pose. Then we find regions in the X-ray images, which are expected to lead to stable and reliable registration results. All other regions are excluded from the DRR rendering process as well as from the image similarity maximization. This allows us to reduce computation time and to enhance the reliability of the pose estimation process.

2. RELATED WORK

In (Jeongtae et al., 2001) it is already shown, that mutual information is a suitable measure to find an estimation of the patient set-up error in radiotherapy. We make use of this measure to compare X-ray images with the respective reconstructed radiographs. However, to gain a full 6 DOF alignment, a large number of DRRs have to be rendered and must be compared to the X-ray images. This is normally done at high costs of computation time and reduces the acceptance of the full 6 DOF alignment in clinical applications.

One suggested solution to this problem is given in (Birkfellner et al., 2003). They propose to perform several 2D to 2D image registrations between DR images and DRRs. The resulting 3D transformation is then computed by inverse projection. The DRRs are updated as soon as the hypothetical 3D transformation of the patient, and therewith the CT scan, reaches a certain amount of translation or rotation. The complete pose estimation process can be sped-up by means of factor ⅓ at the cost of some tenths of a degree in rotation accuracy. However, only ±1.6 millimetres accuracy could be reached in translation.

In (Selby et al., 2008) full 6 DOF alignment correction is shown for a single X-ray image and a single DRR image. Using only a single X-ray image could reduce rendering time by ½ compared to the stereoscopic approach, but suffers from low translational accuracy in direction of the X-ray axis (axis from the X-ray source to the centre of the digital flat-panel).

(Rohlfinga et al., 2004) propose what they call progressive attenuation fields to speed-up the rendering process for pose estimation. As many DRRs with only slight deviations have to be created, each ray through the volume is computed only one time. Once computed, each ray represented by the result of a line integral is stored in a hash table and can be reused. To reduce the number of stored rays, interpolation is applied. Through this approach the pose detection could be sped-up by factor ½ compared to the same algorithm using standard ray-casting to render the DRRs. Unfortunately, accuracy and reliability suffer from large initial patient set-up errors and with 16 mm initial misalignment, only about 30% of the tested cases led to correct results.

For our approach we aim to achieve reliable results for at least 20 mm of initial misalignment. Thus we perform an initial 5+1D pose detection similar as in (Birkfellner et al., 2003), which is based on 2D to 2D image registrations with updating the DRRs at certain steps of the process, to reflect the real alignment of the patent relative to the X-ray equipment. This is done until the transformation cannot be further optimized.

To enhance accuracy in 6 DOF we then perform automatic image comparisons with new DRRs, rendered for each tested 3D patient pose. This is done as described in (Selby et al., 2008), but we use two X-ray images to ensure an acceptable 3D accuracy. To reduce rendering time and the influence of certain areas of the image, we select parts of the X-ray image, which are suitable for image matching. The DRR is ray-traced only in these areas and all other parts of the image are excluded from the processing.

In the selected areas, the rendering is done without reducing image resolution or radiometric quality to avoid degradation of the reachable accuracy.

3. METHODS

In figure 1 we first give a brief overview over the whole pose estimation process as performed by our approach. After that, the relevant working steps will be explained in more detail.

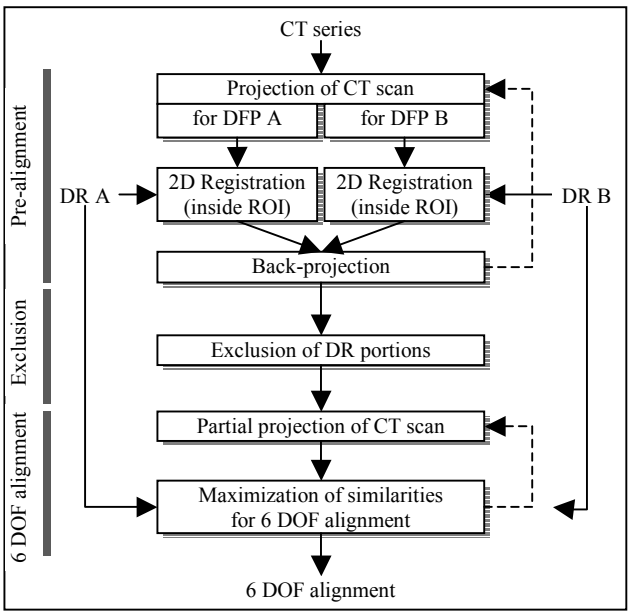

Figure 1. Procedure overview

The first step of finding the 6 DOF patient alignment is to perform an initial alignment (pre-alignment in figure 1). In this pre-alignment it is not possible to find the alignment in full 6 degrees of freedom, but the CT series has to be projected into the digital flat-panel (DFP) plane only a few times.

However, through the pre-alignment we gain a good approach to the real patient pose. Next, we identify portions of the DR images that can be excluded from further image comparison. We can identify these regions in the planes of the DRR images as well. They are then excluded from the rendering process, to speed-up the computation.

In the last step, the full 6 DOF alignment, we fine-tune the initially detected pose to achieve the desired accuracy (here 0.5 mm). Therefore, image comparisons between DRRs for different alignments are performed and the similarity is maximized. As this procedure requires a large number of consecutive CT projections, it benefits from the fact, that areas could be excluded from the rendering process.

3.1 Pre-alignment

To find the patient alignment we first perform a step that we call pre-alignment. This is done by 2D registrations of two X-ray images to the respective DRRs. The results of the registrations are inversely projected into 3D space and used to update the DRRs with the new alignment.

3.1.1 The Registration Process: There exists a wide range of gray-value based image comparators in the scope of registration. As methods like cross-correlation or usage of difference images are not applicable for images that differ in much more aspects than contrast and intensity, we decided to use Mutual Information (MI) as image similarity measure (PLUIM et al., 2003).

A joint histogram is built-up by reading the gray-values of both images at the position of two overlaid pixels. A cell of the two-dimensional histogram is then incremented by one at the respective coordinates, defined by the two gray-values. The Mutual Information value MI is calculated by equation 1:

$$MI(R,F) = H(R) + H(F) - H(R,F) \qquad (1)$$

where $MI(R,F)$ = Mutual Information value
R, F = reference (DR) and floating image (DRR)
$H(R)$, $H(F)$ = Shannon Entropies of the images
$H(R,F)$ = Joint Entropy of R and F

The negative MI value is minimized by a Downhill Simplex minimization algorithm as described in (Press et al., 1982). The three free transformation parameters are the floating image shifts in X- and Y- direction of the image plane and rotation of the image plane around its normal vector.

3.1.2 Inverse Projection: After each registration, the results are back-projected into a common coordinate system. The underlying geometry is shown in figure 2.

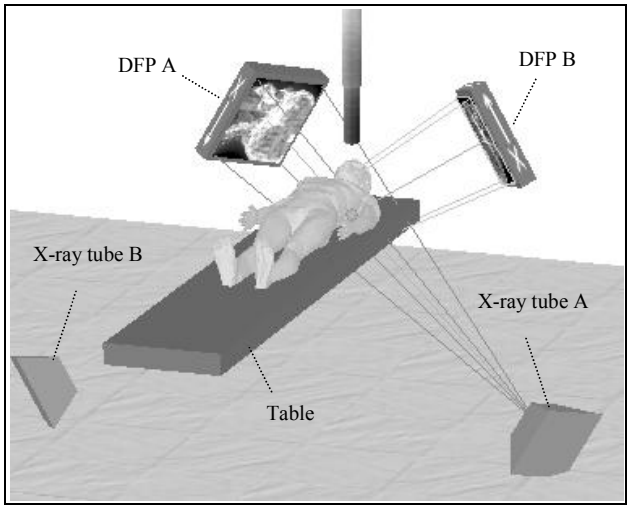

Figure 2. Geometry of the treatment equipment

Figure 2 depicts only the relevant parts of the equipment. The image detectors and the X-ray tubes determine the geometric properties that are of essential importance for the DRR rendering and for the inverse projection of the registration results. The patient table determines the coordinate system used for patient alignment.

3.1.3 DRR Update and User ROIs: The DRRs are created by ray-tracing. When scattering is neglected, the intensity of an X-ray passing through the respective object is given by the line integral along the virtual X-ray:

$$I = I_0 * e^{-\int_{-\infty}^{+\infty} f(x,mx+b)dx} \qquad (2)$$

where I_0 = intensity of the X-ray at the source
I = intensity of the DRR gay-value

we choose I_0 to normalize the expression in equation 2 to a resulting intensity range of

$$I \in [0.0, 1.0],$$

which can be scaled to discrete 12 bit pixel values. In our case, the object the ray runs through is the CT volume.

Usually all DRR pixels would be computed by the line integral. However, it is possible that the user defines a Region of Interest (ROI) in the X–ray image(s) to exclude certain regions from the matching process. This can be the case if an image contains features that only occur in one of the modalities (X-ray image or projected CT scan) as for example head fixation equipment or parts of the patient table, to name but two. In this case a validity mask is created to mark valid parts of the X-Ray image. The markings are transferred to the image plane of the respective DRR image and are used to mask out DRR pixels. The ray-tracing process for the pre-alignment then skips rays that are not included in the ROI (see figure 3).

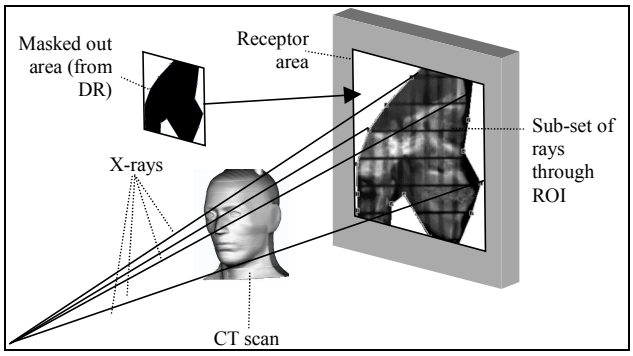

Figure 3. Ray-tracing and ROI

3.2 Automatic Exclusion of DR Portions

After the pre-alignment we mask out regions of the X-ray image and the DRR plane to speed-up the rendering process and to concentrate on relevant areas of pixels when performing the alignment fine-tuning. As shown in (Olson, 2001) aligning entropies of images is a stable way to perform a registration and can be considered more reliable than aligning gradients. This holds especially in cases (as in our case) where it is not trivial to find corresponding gradients in the image pairs, because of differences between the two image modalities. However, (Olson, 2001) aim a template based matching, whereas we want to find the maximal similarity between two different images. For this reason we still use MI as measure, but mask out parts of the images using the Block Entropy for a certain window *W* around each pixel of the X-ray image.

The Shannon Entropy is computed by equations 3:

$$H(X) = -\sum_{g=0}^{G} p_g * \ln(p_g) \qquad (3)$$

with $\lim_{p_g \to 0} p_g \ln(p_g) = 0$

where *H(X)* = entropy between 0 and 1
X = respective image (here the X-ray image)
p_g = probability for the occurrence of gray-value *g*

We compute the Block Entropies *H(W)* of the image for each pixel in a predefined window *W*. Therefore we use a window of size *7 x 7 = 49* pixels. To reduce computation time and to avoid overweighing of small pixel intensity fluctuations we apply histogram bins of size 32 gray-levels (the original images have 12 bit gray-value resolution and are downsampled). The entropies of the single pixels are stored in an entropy map with the same resolution as the original image. At the same time we compute the total entropy *H(X)* for the whole image by summation of the Block Entropies *H(W)*. Because pixels can occur twice in different windows, we additionally normalize by the number of blocks (which are the number of windows, for which the entropy could be computed).

To make the algorithm more tolerant against noise and to further decrease computation time, the entropies are computed in half image resolution and the image data is filtered by a *3 x 3* Gaussian kernel. The low-resolution entropy map is then resampled to the original image size. In this way, the *7 x 7* pixel area of the Block Entropy window, applied to the lower resolution image, covers *14 x 14* pixels in the original resolution. Enlargement of the entropy map additionally leads to a margin around the masked out areas. This can be advantageous because inaccuracies of the previous processing steps can still be corrected by aligning only the correct sup-parts of these areas in the final matching process.

The entropy gives the average information per pixel, normalized to the range 0 to 1. The Block Entropy *H(W)* gives the entropy for subparts of *H(X)*. Thus we compare the Block Entropies with the image entropy and mask out regions where

$$H(W) < H(X)$$

holds, to exclude image regions with relative low information per pixel (see figure 4).

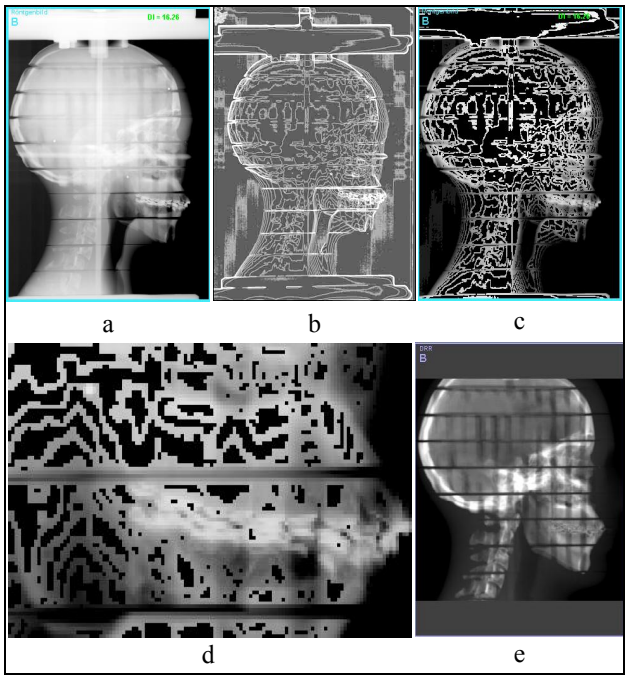

Figure 4. Exclusion of image regions: a) X-ray image of an anatomic head-phantom; b) map of Block Entropies for *7 x 7* pixel windows and 32 gray-value bins; c) X-Ray image where low Block Entropies have been

masked out; d) enlarged jaw area of the masked out X-ray image; e) corresponding DRR that will be masked by the validity mask from the DR image.

As visible in figure 4-d, areas containing high information density with respect to the Block Entropy are preserved. This enables us to perform a Mutual Information based comparison of the remaining gray-values.

3.3 6 DOF Alignment Computation

Relying on two 2D transformations, only 2 of the 3 possible rotations for a spatial full 6 DOF correction can be calculated. This is because object rotation around the vector vertical to the two X-ray beamlines results in no rotational component on the flat-panels after the object is projected. The rotation only results in a change of the image contents (see figure 5).

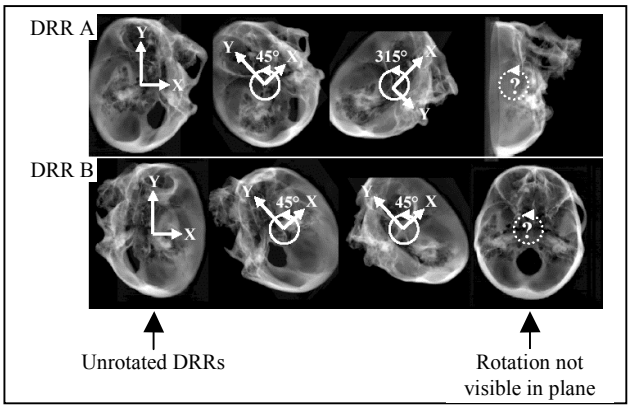

Figure 5. Original DRRs and DRRs rotated by 45° around different room axes

We use a Downhill Simplex approach with 6 free parameters (3 translations and 3 rotations) to maximize the combined Mutual Information between the DR and the DRR for both stereoscopic image pairs *A* and *B*. The single Mutual Information values are combined as given by equation 4:

$$MI(DR_A * VM_A, DRR_A(T) * VM_A(T))^2 + MI(DR_B * VM_B, DRR_B(T) * VM_B(T))^2 \rightarrow \max \quad (4)$$

where A = image pair for X-ray flat-panel A
B = image pair for X-ray flat-panel B
T = sextuplet of the free transformation parameters
VM = validity mask for the respective image pair

For the optimization we need to re-render each DRR approximately 100 times. To reduce rendering time we mask out pixels of the DRR. This is done according to the validity mask built-up from the entropy map and if applicable, a user defined ROI. Finally we obtain our optimal transformation T and can now realign the patient support equipment or the radiation source for treatment, e.g. as proposed by (Yue et al., 2006).

4. RESULTS

To test our approach we compared the proposed algorithms to a reference algorithm, which uses the full image domain to perform 6 DOF alignment correction. The reference algorithm is described in (Selby et al., 2007). We did not apply a user defined ROI for the performed tests, to stay independent from human input. However, using regions of interest usually does not degrade the pose estimation results, except the ROI has been defined in inappropriate areas of the images or contains not enough data for a stable registration.

A head-phantom CT series with 1 mm slice distance and a human pelvis-phantom CT with 0.8 mm slice distance, each with 0.49 mm in slice X- and Y- direction have been used to test the pose correction approach. The respective kilovoltage X-ray images where acquired with Varian PaxScan 4030R flat-panels (40 cm x 30 cm X-ray image size at 3200 x 2304 pixels resolution) from within a real treatment environment.

The performance of the alignment estimation could be increased by means of 4 up to 8, depending on the number of DRR creations necessary (the more iterations with DRR renderings had to be computed, the higher was the speed-up). The DRR creation itself could be sped-up by approximately factor 10, by masking out the pixels which possess low average information compared to the whole image. The creation of the entropy map, which in contrast to the rendering and matching algorithms, was not optimized for a multi-processor system required several seconds of additional computation time. Yet the additional time expense could be compensated easily by the accelerated rendering.

For the head dataset we achieved the same accuracy as with the reference algorithm. The differences between the resulting pose parameters were $\sigma = \pm 0.2$ mm for translations and $\sigma = \pm 0.1°$ for rotations.

For the pelvis phantom, containing more unwanted artifacts than the head data, we could increase the accuracy and reliability of the pose estimation and were even able to compute an acceptable patient alignment at a large initial translation error of 20 mm, where the reference algorithm failed.

In figure 6 (left) the remaining translational alignment errors are shown for the 6 DOF reference algorithm (approach A) and our new approach (approach B). These tests have been performed with the human head-phantom dataset on a 2,66 GHz Dual Core workstation for a range of initial misalignments between 1.0 mm and 30 mm. Concerning the remaining misalignments it is not obvious, which approach to prefer. However, the computation time can be reduced dramatically (figure 6, right) due to the reduced rendering efforts coming with the methods introduced with this paper.

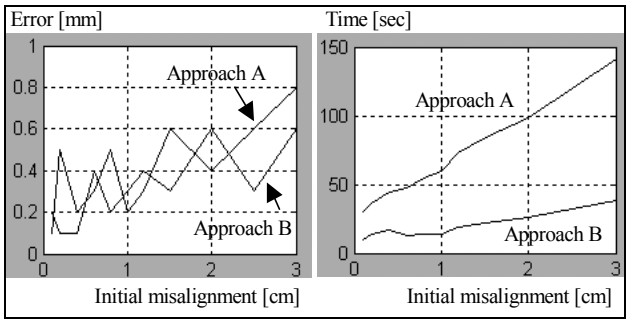

Figure 6. Error in dependence of the initial misalignment (left) and computation time (right); Approach A is the reference algorithm, approach B is the approach contributed in this paper

5. CONCLUSIONS

We introduced an optimized approach for the pose correction of patients in 6 degrees of freedom. First we relied on the whole image domains to compute a 5 DOF estimation of the alignment based on two 2D registrations and a low number of DRR renderings. Then we improved accuracy and found the 6th degree of freedom by optimizing the correlation of certain image regions. By reduction of the data used in the X-ray reconstruction process and by exclusion of image regions that do not improve the matching result, we were able to reduce pose estimation time and to preserve the accuracy of the matching process. This can be beneficial especially for real time matching applications, e.g. for alignment surveillance with fluoroscopic imaging devices during the treatment. To achieve even higher performance we suggest to combine our approach with other techniques, for example the progressive lightfield rendering introduced in (Rohlfinga et al., 2004).

For images that contain large areas of image contents that do not occur in either the CT series or the X-ray image(s) our approach successfully increased reliability of the pose detection outcome. However, the method contributed here still has to undergo additional testing, especially with respect to radiometric differences of the image modalities and for the variety of possible imaged body parts (e.g. for thorax or abdomen datasets).

Further investigations should be made, if our approach combined with a multi-resolutional matching could lead to further improvements.

After all, we showed that the solution to achieve lower pose detection times with similar accuracies can be to reduce the image registration to the relevant sub-sets instead of improving the rendering or matching technique itself. As we still use the mutual information measure for registration of these sub-sets, we still benefit from the stability of this measure.

REFERENCES

Birkfellner, W; Wirth1, J.; Burgstaller, W.; Baumann, B.; Staedele, H.; Hammer, B.; Gellrich, N. C.; Jacob, A. L.; Regazzoni1, P.; Messmer, P., 2003. A faster method for 3D/2D medical image registration - a simulation study. Phys. Med. Biol., 48, pp. 2665-2679.

Jeongtae, K.; Fessler, J. A.; Lamy, K. L.; Baltery, J. M.; Hakeny, R. K. T., 2001. A Feasibility Study on Mutual Information based Set-up Error Estimator for Radiotherapy, Medical Physics, 28(12), pp. 2507-2517.

Olson, C. F., 2001. Image registration by aligning entropies. Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2, pp. 331-336.

Pluim, J.; Maintz, J.; Viergever, M., 2003. Mutual information based registration of medical images: A survey. IEEE transactions on medical imaging, 22(8), pp. 986-1004.

Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P., 1992. Numerical Recipes in C 2, Cambridge University Press, pp. 408-411.

Rohlfinga, T; Russakoff, D. B.; Denzler, J.; Mori, K.; Maurer, C. R., 2004. Progressive attenuation fields: Fast 2D-3D image registration without precomputation. In Proc. of Medical Image Computing and Computer-Assisted Intervention, 3216, pp. 631-638.

Selby, B. P.; Sakas, G.; Walter, S.; Groch, W.-D.; Stilla, U., 2008. Patient Alignment Estimation in Six Degrees of Freedom Using a CT-scan and a Single X-ray Image. Proceedings of Bildverarbeitung für die Medizin 2008, pp. 127-132.

Selby, B. P.; Sakas, G.; Walter, S.; Groch, W.-D.; Stilla U., 2007. Detection of pose changes for spatial objects from projective images. In: Stilla, U.; Meyer, H.; Rottensteiner, F.; Heipke, C.; Hinz, S. (eds), PIA07 Photogrammetric Image Analysis 2007. International Archives of Photogrammetry and Remote Sensing, 36(3), pp. 105-110.

Thilmann, C.; Nill, S.; Tücking, T.; Höss, A.; Hesse, B.; Dietrich, L.; Bendl, R.; Rhein, B.; Häring, P.; Thieke, C.; Oelfke, U.; Debus, J.; Huber, P., 2005. Correction of patient positioning errors based on in-line cone beam CTs: clinical implementation and first experiences. International Journal of Radiation Oncology Biology Physics, 63(1), pp. 550-551.

Verhey, L. J.; Goitein, M.; McNulty, P.; Munzenrider, J. E.; Suit, H. D., 1982. Precise positioning of patients for radiation therapy. International Journal of Radiation Oncology Biology Physics, 8(2), 289-294.

Yue, N. J.; Knisely, J. P. S.; Song, H.; Nath, R., 2006. A method to implement full six-degree target shift corrections for rigid body in image-guided radiotherapy. Medical Physics, 33(1), pp. 21-31.