Fast Intensity-based 2D-3D Image Registration of Clinical Data Using Light Fields

Daniel B. Russakoff
Dept. of Computer Science
Stanford University
Stanford, CA 94305
daniel.russakoff@cs.stanford.edu

Torsten Rohlfing
Image Guidance Labs
Stanford University
Stanford, CA 94305
rohlfing@igl.stanford.edu

Calvin R. Maurer, Jr.
Image Guidance Labs
Stanford University
Stanford, CA 94305
maurer@igl.stanford.edu

Abstract

Registration of a preoperative CT (3D) image to one or more X-ray projection (2D) images, a special case of the pose estimation problem, has been attempted in a variety of ways with varying degrees of success. Recently, there has been a great deal of interest in intensity-based methods. One of the drawbacks to such methods is the need to create digitally reconstructed radiographs (DRRs) at each step of the optimization process. DRRs are typically generated by ray casting, an operation that requires $O(n^3)$ time, where we assume that n is approximately the size (in voxels) of one side of the DRR as well as one side of the CT volume. We address this issue by extending light field rendering techniques from the computer graphics community to generate DRRs instead of conventional rendered images. Using light fields allows most of the computation to be performed in a preprocessing step; after this precomputation, very accurate DRRs can be generated in $O(n^2)$ time. Another important issue for 2D-3D registration algorithms is validation. Previously reported 2D-3D registration algorithms were validated using synthetic data or phantoms but not clinical data. We present an intensity-based 2D-3D registration system that generates DRRs using light fields; we validate its performance using clinical data with a known gold standard transformation.

1. Introduction

In order to use preoperatively acquired three-dimensional (3D) images for intraoperative navigation, the images must be registered to a coordinate system defined in the operating room. The image-to-physical registration is commonly performed using stereotactic frames and fiducial markers. Alternatively, the preoperative 3D image can be registered to an intraoperative two-dimensional (2D) image. Registration of an X-ray computed tomography (CT) image to one or more X-ray projection images (e.g. simulator images, portal images, fluoroscopy images, amorphous silicon detector images) is a particularly interesting example of 2D-3D registration that has a number of possible applications, including patient placement for radiotherapy planning and treatment verification [1, 4], radiosurgery [13], cranial neurosurgery [9], neurointerventions [6, 7], spinal surgery [8, 21], orthopedic surgery [5], and aortic stenting procedures [2, 14, 21].

The 2D-3D registration problem involves taking one or more X-ray projection (2D) images of the patient's anatomy and using those projections to determine the rigid transformation $\mathbf{T}$ (rotation and translation) that aligns the coordinate system of the CT (3D) image with that of the operating room. Figure 1 shows a schematic representation of the 2D-3D registration process. In general, most of the proposed solutions to this problem fit into this general framework. We are interested in intensity-based 2D-3D image registration [3, 9, 14, 15, 21]. In this case, the reference image is one or more X-ray projection images, and the floating image is a CT image. The method involves computing synthetic X-ray images, which are called digitally reconstructed radiographs (DRRs), by casting rays using a known camera geometry through the CT image (Fig. 2). The DRR pixel values are simply the summations of the CT values encountered along each projection ray. The pose (position and orientation) of the CT image (given by the transformation $\mathbf{T}$) is adjusted iteratively until the DRR it produces is most similar to the reference X-ray projection image. A variety of similarity measures have been used, including cross correlation, entropy, mutual information, gradient correlation, pattern intensity, and gradient difference [15].

DRRs are computationally expensive to create, and their generation is typically a bottleneck in the execution of the registration process. In this paper, we address this issue by using an extension of light field rendering from the com-

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

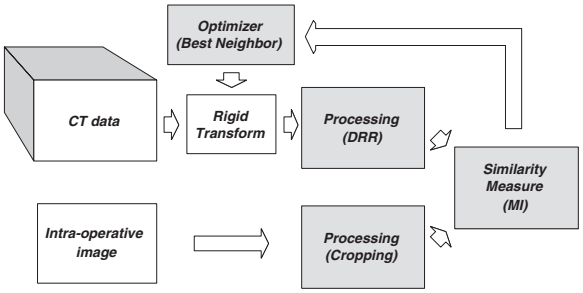

Figure 1. Schematic overview of the 2D-3D registration process. In our case, the reference image is an intraoperative X-ray projection (2D) image. It is used as is with no processing. The floating image is a CT (3D) image. It is processed by generating DRRs (synthetic X-ray projection images) for various orientations of the CT image relative to the X-ray imaging system. The optimizer searches for the rigid transformation $\mathbf{T}$ that produces the DRR that is most similar to the real X-ray projection image. The optimal transformation is used to align the CT's coordinate system with that of the operating room.

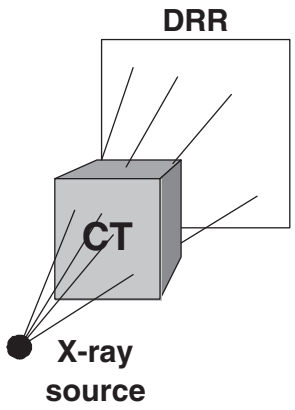

Figure 2. Schematic geometry of DRR creation.

puter graphics community [10] to generate DRRs instead of conventional rendered images [17]. Using light fields allows most of the computation to be performed in a preprocessing step. After this precomputation, very accurate DRRs can be generated substantially faster than with conventional ray casting.

Another important issue for 2D-3D registration algorithms is validation. Previously reported 2D-3D registration algorithms were validated using synthetic data or phantoms [3, 9, 14, 15, 19, 20, 21] but not clinical data. This problem is due mainly to the fact that it is particularly difficult to obtain a gold standard transformation from clinical data. We address this issue by using clinical spine image data from a radiation oncology system that performs 2D-3D registration using bone-implanted fiducial markers. This system (CyberKnife Stereotactic Radiosurgery System, Accuray, Sunnyvale, CA), uses two intraoperative orthogonal X-ray projection images of known geometry to triangulate the marker positions in 3D. It then registers the marker positions to their corresponding locations in a preoperative CT. For every patient treated, the CyberKnife system archives not only the preoperative CT and the intraoperative X-ray projection images, but also the transformation $\mathbf{T}$ it calculates. We ignore the fiducial markers and use the CT and X-ray projection images as input to our 2D-3D registration system. We validate our results by comparing them to the transformation $\mathbf{T}$ calculated by the CyberKnife system using the markers, which we use as ground truth.

The rest of this paper is organized as follows. Section 2 discusses the creation and use of light field DRRs, Section 3 details our registration algorithm and its validation, Section 4 presents our results, and Section 5 presents some discussion and conclusions.

2. Light Field DRRs

In traditional light field rendering, a pixel is a value indicating how much light is reflected off the first surface its ray intersects in the direction of that ray. By contrast, in a DRR, each pixel is a representation of the sum of the CT attenuation coefficients its ray encounters along the path from the source to the destination.

To take these differences into account we modify the light field generation process by introducing the *virtual image plane* (Fig. 3). The virtual image plane is placed exactly where the (s, t) plane would be if we considered the CT data to be a 3D scene and were performing standard light field rendering. The two-plane parameterization of the rays is thus unchanged from the normal case.

In the light field generation, however, instead of creating images with the standard definition of pixels, we associate each sample $\mathbf{p_i} = (u_i, v_i, s_i, t_i)$ with a scalar function $\mathbf{p}_i \mapsto q(\mathbf{p}_i)$, which is the sum of the CT attenuation coefficients encountered along the ray $R_{\mathbf{p}_i}$:

$$q(\mathbf{p}_i) = \sum_{\mathbf{x}_j \epsilon R_{\mathbf{p}_i}} \mathrm{CT}(\mathbf{x}_j). \qquad (1)$$

In standard ray casting, computation along a ray stops as soon as it intersects an opaque surface. In our DRR formulation, we need to continue to trace the ray through the CT to determine its sum. To do so while maintaining the same parameterization of rays in space, we must cast the rays beyond the virtual image plane onto the effective image plane. The values we use to generate the light field are those that lie on the image created on the effective image plane. In both cases we create a skewed perspective image. The main difference is that in regular light field generation,

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

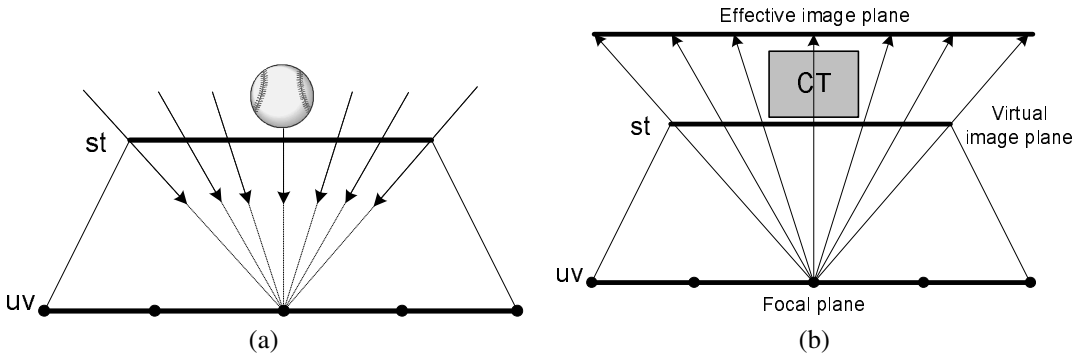

Figure 3. (a) Standard light field geometry. Each pixel on the image plane represents the amount of light reflected to the center of projection from the surface its ray intersects. (b) DRR light field geometry. The virtual image plane allows us to maintain the same two-plane parameterization as in the normal case. In this case, however, the rays are extended to an effective image plane situated beyond the CT image (the scene) so that each pixel may be associated with the sum of CT attenuation coefficients along its ray.

the effective image plane remains fixed and lies between the scene and the focal plane. In DRR light field generation, the virtual image plane remains fixed while the effective image plane can move and lies on the other side of the scene from the focal plane.

Using this extension to the light field technique, we can create large DRR light fields from particular viewpoints (anterior-posterior (AP) and lateral, for example), compress them using vector quantization, and use them to generate new DRRs using the same interpolation of the 4D ray space used by Levoy & Hanrahan [10]. This works because for each ray of every new DRR we create, we can find its corresponding sum of attenuations by sampling the light slabs we generate. In addition, two of the principle problems with normal light field rendering, occlusions and lighting variations, are not issues in the DRR domain. In particular, the scene (the CT volume) is always a rectangular parallelepiped and, thus, always convex and free from occluders. Also, there is inherently no lighting to consider.

3. 2D-3D Registration

With an apparatus for fast DRRs in place, we can now begin to work on the other important features of a 2D-3D registration system. In particular, we need a reference image, a similarity measure, an optimization scheme and, most importantly, a gold standard with which to compare our results.

3.1. Gold standard

Radiation oncology has helped many people beat cancer. Unfortunately, there are a number of people who have tumors that can neither be operated on nor irradiated in the traditional sense as they are too close to important structures, e.g. the speech center of the brain or the spinal cord. To help these people, physicians and engineers have developed a system that uses a robot to shoot focused beams of radiation at a tumor from many different directions. Each beam individually is not harmful enough to hurt living tissue, but the intersection of all beams receives a lethal dose of radiation. If the beam directions are calculated carefully, one can ensure that the tumor is the only region that receives such a lethal dose.

The only remaining issue is that the set of all beam directions must be precalculated using a CT of the patient with the tumor segmented out. The robot and the patient, however, are in different reference frames so there needs to be a 2D-3D registration to bring them into alignment. The current system (depicted in Fig. 4) uses a point-based registration, which involves the insertion of several radio-opaque bone-implanted fiducial markers. These markers show up well in the CT, making them easy to localize. During the procedure, the patient is imaged with two orthogonal X-ray sources with amorphous silicon display (ASD) detectors (Fig. 5a). If the same marker can be localized in both images and the imaging geometry is known, the marker's 3D position can be calculated in the reference frame of the robot by triangulation. If this can be done for at least 3 non-colinear markers, their 3D positions in the operating room can be registered to their corresponding 3D CT positions to arrive at $\mathbf{T}$. It should be noted that this registration procedure is rigid (with only 6 degrees of freedom) and it is designed to be used for tumors on or near rigid or bony structures.

We obtained four archived data sets of patients who have undergone this procedure, each with a tumor close to the spinal cord. Two of the patients have tumors near cervical vertebrae (C3 and C5), and two have tumors near thoracic vertebrae (T1 and T8). This data is actual, clinical data with a built-in gold standard, namely the transformation calculated using the fiducial markers. In particular, for each pa-

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

Figure 4. Photograph of the CyberKnife radiosurgery system showing: (1) ceiling mounted diagnostic X-ray sources, (2) compact linear accelerator mounted on computer-controlled robot arm, and (3) amorphous silicon detectors.

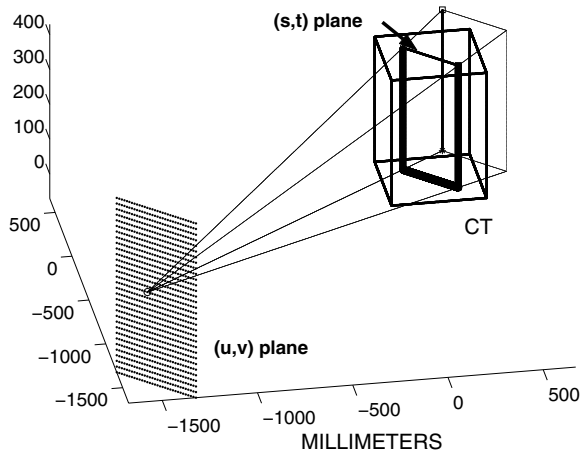

Figure 6. (u, v) and (s, t) planes for one view of the patient used in our experiments. These parameters did not change from patient to patient and were chosen based on our knowledge of the projection geometry of the X-ray system.

tient we have the preoperative CT, the ASD X-ray images from both sources, and the transformation **T** calculated by the system. What we propose is to perform the registration using only the image intensities. In the long run, this would save the physicians the time and trouble of inserting the fiducials and save the patient the morbidity associated with such a procedure.

Our registration proceeds exactly as depicted in Fig. 1. In particular, we focus on the shaded boxes of the figure, those dealing with the processing of the reference image and the floating image, the similarity measure, and the optimization.

3.2. Reference image

We use the ASD images as our references. Figure 5b shows a schematic of the system with two examples of ASD images. Since each procedure is limited to a specific area, e.g. a tumor near a vertebra, we crop the reference images to include only this region. This cropping allows us to speed up our computations and saves us from having to account for the fact that the CT and the intraoperative images were created at different times. By restricting the registration to a specific region of interest, we minimize the chance that the structures have moved relative to each other between the time the CT is taken and the time the procedure is performed. This cropping is performed manually and with minimal effort.

3.3. Floating image

Given a preoperative CT, we first generate two separate light fields, each corresponding to one of the fixed ASD image views. Following [17], we use a (u, v, s, t) resolution of $32 \times 32 \times 256 \times 256$. Given a CT image, we generate the light fields such that the distance between the (u, v) and (s, t) planes are roughly equal to the distance between the radiation sources and their respective detectors. This heuristic is likely to give us a higher sampling density in the light field where we need it the most. Figure 6 shows the light field geometry for one of the ASD imagers in our experiments.

Next, we pick a point in the CT as an origin around which to center our registration. This point does not need to be particularly precise. In general, a point near the structure of interest (the tumor, in this case) works best. We pick the point manually, by looking at CT slices and choosing a point in the center of one of the vertebral bodies closest to the tumor.

3.4. Similarity

As a similarity measure, we use mutual information [11]. Though it had not performed well in previous 2D-3D spine image registration work [15], we assume that this performance was due to the low resolution of the images used

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

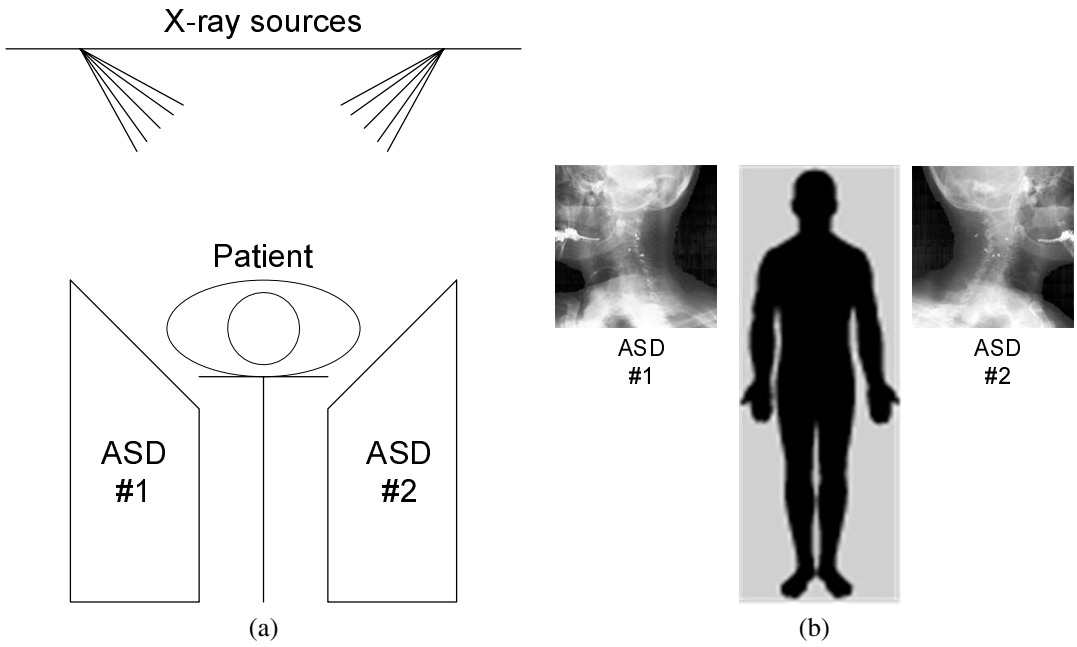

Figure 5. (a) Head-on view of the system. The two X-ray sources are fixed into the ceiling and calibrated so that both their intrinsic and extrinsic parameters are known. (b) A birds-eye view of system with two example ASD images.

and very restricted regions of interest, which made for extremely sparse joint histograms. Given our faster DRR creation, we are able to use higher resolution images that, in turn, allow us to better populate a joint histogram for the purposes of calculating mutual information. This fact, combined with mutual information's obvious success in 3D-3D registrations, both rigid and non-rigid, made it our measure of choice. Our cost function then becomes the sum of the mutual information between the two reference ASD images and their corresponding DRRs calculated with $\mathbf{T}$.

3.5. Optimization

As our optimization strategy, we use Studholme et al.'s best neighbor search [18]. This procedure takes $\mathbf{T}_s$ as an initial transformation and expands it by varying each parameter by a given step size. The expansion generates the 12 closest neighboring transformations $\mathbf{T}_i$ which are each evaluated by our cost function. The evaluation consists of generating DRRs using each $\mathbf{T}_i$ and the geometry of each ASD imager and comparing them to their corresponding reference ASD images. The neighbor with the best value as defined by our cost function is then expanded, and so on until no more improvements can be made. Then, the current best transformation is taken and expanded with a smaller step size. This continues until some predetermined resolution that we set at the beginning.

We choose the step size in the following way. For each parameter, a scaling factor is computed such that when added to its corresponding parameter, the average motion of all projected voxels in the projection plane is equal [16].

Additionally, we augmented this algorithm by performing two passes of it, one with smoothed versions of the reference images to get a good approximation of the optimal $\mathbf{T}$ and then one with the actual reference images. This procedure has the effect of smoothing the cost function in order to avoid local minima.

4. Results

4.1. Regular DRRs vs. light field DRRs

We performed 2D-3D image registration using our data set of four archived CyberKnife patient procedures complete with gold standard. Two of the procedures were on tumors close to cervical vertebrae, and two were on tumors near the thoracic vertebrae. Specifically, we ran our registration algorithm using initial configurations derived by randomly perturbing the gold standard by a maximum of 8 mm in translation and 5° in rotation. These perturbations are on the order of those suggested by Penney et al. as being a reasonable estimate for how close one can get using a very approximate registration procedure [15]. One way to compute an approximate initial transformation is to manually pick corresponding anatomical landmarks in the two X-ray views, triangulate their positions in 3D, and then perform a point-based registration with the corresponding positions of

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

	Patient #1		
	Gold Standard	*Regular DRRs*	*LF DRRs*
x	-284.1	-284.0 ± 0.1	-284.0 ± 0.1
y	-247.0	-247.6 ± 0.1	-247.5 ± 0.1
z	-226.7	-227.0 ± 0.1	-227.2 ± 0.2
θ	-90.4	-89.3 ± 0.2	-89.4 ± 0.4
ϕ	-0.1	0.0 ± 0.2	0.6 ± 0.3
ψ	-89.9	-89.4 ± 0.1	-89.3 ± 0.2
TRE	0.0	0.85	0.92
Time		11,612	151

Table 1. Comparison data between 10 registrations with regular DRRs and light field DRRs for the patient with a tumor near the thoracic vertebrae. The TREs are very close, while the timing differences are significant. Translation (x,y,z) and TRE values are in millimeters. Rotation (θ, ϕ, ψ) angles are in degrees. Execution times are in seconds.

	Patient #2		
	Gold Standard	*Regular DRRs*	*LF DRRs*
x	-254.8	-254.2 ± 0.3	-254.3 ± 0.2
y	-219.9	-220.0 ± 0.3	-220.0 ± 0.5
z	-158.0	-157.9 ± 0.2	-157.5 ± 0.2
θ	-91.7	-87.0 ± 0.2	-87.8 ± 1.4
ϕ	-0.5	3.9 ± 1.2	2.9 ± 1.1
ψ	-90.2	-89.5 ± 0.5	-89.9 ± 1.1
TRE	0.0	1.80	1.54
Time		11,631	147

Table 2. More comparison data between 10 registrations with regular DRRs and light field DRRs for a patient with a tumor near the cervical vertebrae. Again, the TREs are very close, while the timing differences are significant. Units are the same as in Table 1.

those landmarks in the CT.

Using these initial random perturbations to initialize the system, we then perform the registration as described, once using ray casting to generate the DRRs and once using light fields. We performed these experiments on a PC with a 2 GHz Intel Pentium 4 processor. Table 1 shows the results for 10 registrations on a patient with a thoracic tumor, and Table 2 shows the same for a patient with a cervical tumor. We display only the results for the runs that converged close to the correct solution. All 10 runs converged in Table 1, while only 9 of 10 converged in Table 2. Fortunately, given the small standard deviations of the parameters on the runs that did converge, it is relatively straightforward to decide whether or not a registration has converged. To validate our results, we used target registration error (TRE) [12] of a

	Regular DRRs	Light Field DRRs
Avg. TRE	1.35 mm	1.31 mm
Avg. Time	16,108 sec	158 sec

Table 3. Summary of our results for all four patients in our data set. The total TRE is calculated against our known gold standard. The light field DRRs perform just as well as standard DRRs but do so two orders of magnitude faster.

region centered near the tumor and extending to dimensions roughly equal to those of the vertebral bodies in question ($40 \times 40 \times 40$ mm for thoracic vertebrae and $30 \times 30 \times 30$ mm for cervical vertebrae). We used the gold standard to compare against our results.

A few interesting observations follow from this data. First of all, the TRE values for the regular DRR registrations are very close to those of the light field DRR registrations. In fact, in Table 2 we actually see that the TRE for the light field DRRs is lower than that of the regular DRRs. We attribute this mainly to the smoothing that is inherent in the light field DRR construction, which sometimes prevents the cost function from overfitting noise in the optimization. Taking all four patients in our data set together, Table 3 summarizes our results. In particular, the TRE values are quite similar, while the timings are more than two orders of magnitude apart.

4.2. Capture range

Having validated the use of light fields to create fast DRRs, we now look at the accuracy of our entire registration procedure. In particular, we are interested in the capture range of our algorithm. Towards this end, we took two of our patients and performed the following experiment. We randomly perturbed the ground truth transformation until we had 30 transformations with TRE between 0 and 2 mm, 30 between 2 and 4 mm, and so on until the final increment of 14 to 16 mm. We then performed our registrations and counted how many times they converged for each range in the scale. Figure 7 illustrates our results. It shows the percentage of registrations that converged to the true parameters (TRE $<$ 2mm) vs. the initial TRE of the start configuration. As we see, the capture range extends out at least to around 6 or 7 mm of initial TRE. After that, the percentage of success falls.

5. Conclusions

We have demonstrated the effectiveness of using light field DRRs in the application of 2D-3D image registration on real, clinical data with an inherent ground truth. To our knowledge this is the first time light fields have been used

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 IEEE

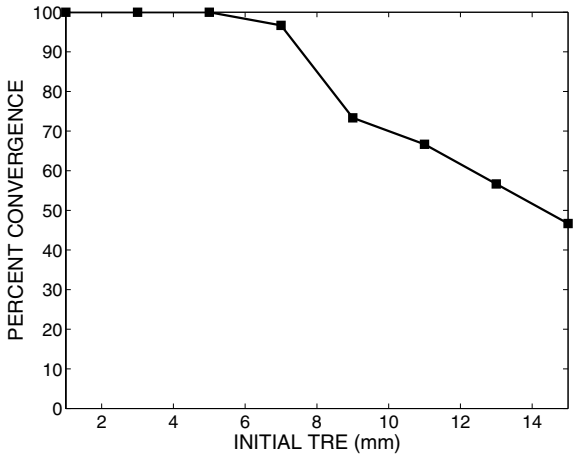

Figure 7. Plot of percentage of successful registrations vs. initial TRE. Each data point represents a range of TRE of 1 mm on either side of its x-coordinate.

to speed up the bottleneck of the DRR formation in traditional 2D-3D intensity-based registration schemes. Also, our work is the first to be validated on real, clinical data with a known gold standard. Our results indicate that light fields are an extremely effective way to dramatically improve the performance of intensity-based 2D-3D registration. Using light fields to generate our DRRs, we achieve a speedup of approximately two orders of magnitude with no accompanying loss of accuracy.

Currently, we are exploring new optimization schemes to help avoid the problem of getting stuck in local minima as well as looking into new similarity measures that may smooth out the cost functions and avoid local minima altogether. In addition, we are looking into the tradeoff between registration accuracy and the resolution of the light fields.

References

[1] J. Bijhold. Three-dimensional verification of patient placement during radiotherapy using portal images. *Med. Phys.*, 20:347–356, 1993.

[2] M. Breeuwer, J. P. Wadley, H. L. T. de Bliek, et al. The EASI Project—Improving the effectiveness and quality of image-guided surgery. *IEEE Trans. Inform. Technol. Biomedicine*, 2:156–168, 1998.

[3] L. M. G. Brown and T. E. Boult. Registration of planar film radiographs with computed tomography. *Proc. IEEE MMBIA 1996*, pages 42–51, 1996.

[4] K. G. A. Gilhuijs, P. J. H. van de Ven, and M. van Herk. Automatic three-dimensional inspection of patient setup in radiation therapy using portal images, simulator images, and computed tomography data. *Med. Phys.*, 23:389–399, 1996.

[5] A. Gueziec, P. Kazanzides, B. Williamson, et al. Anatomy-based registration of CT-scan and intraoperative X-ray images for guiding a surgical robot. *IEEE Trans. Med. Imaging*, 17:715–728, 1998.

[6] E. Kerrien, M.-O. Berger, E. Maurincomme, et al. Fully automatic 3D/2D subtracted angiography registration. *MICCAI 1999*, pages 664–671. Springer-Verlag, Berlin, 1999.

[7] Y. Kita, D. L. Wilson, and J. A. Noble. Real-time registration of 3D cerebral vessels to X-ray angiograms. *MICCAI 1998*, pages 1125–1133. Springer-Verlag, Berlin, 1998.

[8] S. Lavallee, J. Troccaz, P. Sautot, et al. Computer-assisted spinal surgery using anatomy-based registration. *Computer-Integrated Surgery: Technology and Clinical Applications*, pages 425–449. MIT Press, Cambridge, MA, 1996.

[9] L. Lemieux, R. Jagoe, D. R. Fish, et al. A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs. *Med. Phys.*, 21:1749–1760, 1994.

[10] M. Levoy and P. Hanrahan. Light field rendering. *Comput. Graph. (SIGGRAPH '96)*, 30:31–42, 1996.

[11] F. Maes, A. Collignon, D. Vandermeulen, et al. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imaging*, 16:187–198, 1997.

[12] C. R. Maurer, Jr., R. J. Maciunas, and J. M. Fitzpatrick. Registration of head CT images to physical space using a weighted combination of points and surfaces. *IEEE Trans. Med. Imaging*, 17:753–761, 1998.

[13] M. J. Murphy. An automatic six-degree-of-freedom image registration algorithm for image-guided frameless stereotaxic radiosurgery. *Med. Phys.*, 24:857–866, 1997.

[14] G. P. Penney, P. G. Batchelor, D. L. G. Hill, and D. J. Hawkes. Validation of two- to three-dimensional registration algorithm for aligning preoperative ct images and intraoperative fluoroscopy images. *Med. Phys.*, 28:1024–1032, 2001.

[15] G. P. Penney, J. Weese, J. A. Little, et al. A comparison of similarity measures for use in 2D-3D medical image registration. *IEEE Trans. Med. Imaging*, 17:586–595, 1998.

[16] T. Rohlfing, D. B. Russakoff, M. J. Murphy, and C. R. Maurer, Jr. An intensity-based registration algorithm for probabilistic images and its application for 2-D to 3-D image registration. *Medical Imaging 2002: Image Processing*, Proc. SPIE 4684:581–591, 2002.

[17] D. B. Russakoff, T. Rohlfing, D. Rueckert, R. Shahidi, D. Kim, and C. R. Maurer, Jr. Fast calculation of digitally reconstructed radiographs using light fields. *Medical Imaging 2003: Image Processing*, Proc. SPIE 5032, 2003 (in press).

[18] C. Studholme, D. L. G. Hill, and D. J. Hawkes. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Med. Phys.*, 24:25–35, 1997.

[19] R. Szeliski and S. Lavallee. Matching 3-D anatomical surfaces with non-rigid deformations using octree splines. *Int. J. Comput. Vision*, 18:171–186, 1996.

[20] P. Viola and W. M. Wells, III. Alignment by maximization of mutual information. *Int. J. Comput. Vision*, 24:137–154, 1997.

[21] J. Weese, G. P. Penney, T. M. Buzug, et al. Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery. *IEEE Trans. Inform. Technol. Biomedicine*, 1:284–293, 1997.

Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003) 2-Volume Set
0-7695-1950-4/03 $17.00 © 2003 **IEEE**