3D/2D Registration Via Skeletal Near Projective Invariance in Tubular Objects

Alan Liu[1], Elizabeth Bullitt[2], Stephen M. Pizer[2]

1. Presently at the Center for Information-Enhanced Medicine
National University of Singapore
liu@ciemed.nus.edu.sg

2. Medical Image Display & Analysis Group
The University of North Carolina at Chapel Hill

Abstract. We present a method of 3D/2D image registration. The algorithm is based on the property of near projective invariance in tubular objects. The skeletons of tubular anatomical structures (e.g., intracerebral blood vessels) are used as registration primitives. Experiments with Magnetic Resonance Angiogram (MRA) patient studies and both simulated and actual X-ray angiograms suggest that the algorithm is very accurate and robust. The algorithm requires only a small number of primitives. In addition, the algorithm is relatively insensitive to the choice of tubular structures used. Experimental results justifying these claims are included.

1 Introduction

The objective of 3D/2D image registration can be described thus: Given a 2D image and a 3D model, determine the position and orientation (i.e., *pose*) of the imaging device when the 2D image was taken. In this paper, we describe an algorithm for registering 3D Magnetic Resonance Angiogram (MRA) images with X-ray angiograms. The imaging device for our application is a digital fluoroscope. Many 3D/2D registration algorithms have the following paradigm. First, image structures common to both 3D and 2D imaging modalities are chosen as *registration primitives*. Examples include the images of external fiducial markers. The images are assumed registered when the projected 3D registration primitives are within a small distance of coinciding with their 2D counterparts. Second, for a given choice of primitives, an objective function is formulated. The objective function has a minimum value when the primitives coincide. Finally, an optimization algorithm to minimize the objective function is chosen.

Registration primitives include points, curves, and surfaces. The projection involved in registering 3D with 2D images complicates the choice of registration primitives. Not all types of primitives can be unambiguously projected. For example, 3D surfaces may overlap and project onto the same part of the 2D plane. In this paper, we show that the property of near projective invariance permits us to use the skeleton curves of tubular objects as registration primitives. With local exceptions, 3D curves

can be projected unambiguously as 2D curves. Section 2 describes this property in greater detail.

2 Method

A 3D tubular structure such as a blood vessel generally appears as a 2D tube under projection. This section describes the property of projective invariance and its application in 3D/2D registration. Section 2.1 presents the concept of projective invariance and applicable conditions. Section 2.2 describes conditions where projective invariance is inapplicable.

2.1 Near Projective Invariance

A 3D tube contains a central axis or *skeleton* such that cross sections of the object made perpendicular to the skeleton are circular. A 2D tube contains a skeleton which is equidistant from the tube's boundaries. Tubes are not required to have constant width. Generally, the projection of a 3D tube is a 2D tubular shadow. For a given projection, if the 3D skeleton projects onto the 2D skeleton of the projection's shadow, *projective invariance* is preserved. That is, a 3D tube is said to preserve projective invariance if the projection of the 3D skeleton and the 2D skeleton of the tube's projection are the same.

Strict projective invariance is preserved where the tubular object is not overlapped under projection. Liu [1] enumerates such conditions in detail. In practical situations, intensifier induced image distortions, the resolution of the imaging device, and the characteristics of X-ray image formation affect invariance. While strict invariance is not preserved, it is minimally affected. Tubular objects exhibit the property of *near projective invariance* in this situation. That is, the projected 3D skeleton only differs slightly from the projection's 2D skeleton. Tests using both simulated and actual X-ray angiograms suggest that this difference does not significantly affect registration results, as shown in section 4.

Near projective invariance simplifies the use of tubular objects for 3D/2D registration. The problem is reduced to that of registering sets of curves. Curves are computationally simple structures from which a fast, highly accurate registration algorithm can be developed. This algorithm is described in section 3.

2.2 Exceptions to Near Projective Invariance

A tube may not display near projective invariance throughout its length. The invariance property is not preserved when tubes overlap under projection. Such segments should not be used for registration. Two kinds of overlap are possible: local and non-local overlaps.

A local overlap or self-occlusion occurs when a contiguous portion of the same tube overlaps under projection. Fig. 1 illustrates this case. The object is a tubular helix. The helix's axis is perpendicular to the view direction. The left image shows the tube and its 3D skeleton. The right image is a projection of the tube with the projection's 2D skeleton. The tube's 3D skeleton forms a cusp after projection whereas the 2D skeleton is smooth and does not extend as far out the bend as its 3D complement.

A non-local overlap occurs when portions of two distinct tubes or when two non-contiguous portions of the same tube overlap under projection. Fig. 2 illustrates. The
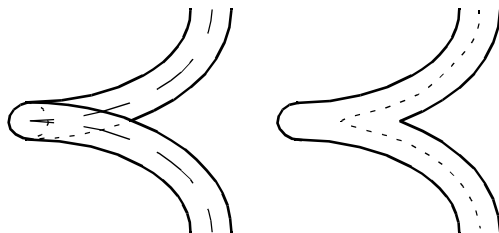
Fig. 1: The effect of self occlusion on projective invariance. Left: projection of the 3D skeleton. Right: The projected tube's 2D skeleton

middle image is an angiogram. The side images are magnified regions where ambiguity arises due to non-local overlaps. From these images, it is not clear whether the projected vessels cross or are just touching.

Fig. 2: A lateral X-ray angiogram of the head. Highlighted areas contain overlapping projections of distinct vessels

3 The Registration Algorithm

Our goal is to register 3D MRA studies with X-ray angiograms. By the property of projective invariance, registering blood vessels with their corresponding projections is equivalent to registering the projected 3D skeletons with their 2D analogues. A 3D/2D registration algorithm was developed that uses the skeletons directly. The algorithm accepts as input the set of 3D skeletons extracted from volume data, the set of 2D skeletons extracted from X-ray images, and the correspondence between the 3D and 2D skeletons. The algorithm returns the pose required to register the primitives. The set of 3D skeletons is assumed to be rigid. Liu [1] develops the registration algorithm at length. In this paper, we summarize its key aspects. Section 3.1 describes our method of extracting tubular skeletons. Section 3.2 describes the objective function. Section 3.3 briefly outlines the optimization method.

3.1 Extracting Tubular Skeletons

Tubular skeletons are extracted directly from digital images. Extraction is performed via cores [2], a method of multi-scale object description. For tubular structures, cores encode the object as a curve in $\Re^n \times \Re^+$ space, where n is the dimensionality of the tube. The *core-middle*, or spatial component of the core, corresponds to the tube's skeleton. Fritsch's algorithm [3] is used to extract cores from X-ray angiograms. The 3D extraction of core-middles from MRAs is performed using Aylward's algorithm [4].

3.2 Computing the Disparity Value

From the property of projective invariance, tubular structures are accurately registered with their projections when the projected 3D skeletons overlap their corresponding 2D equivalents. This section describes a method for measuring registration accuracy based on the degree of overlap or *disparity* between skeletons.

Let C be the skeleton of a 3D tubular blood vessel, and c be the skeleton extracted from the vessel's projection. Let P be the perspective projection function. Under perfect registration, projective invariance implies that $P(C)$ perfectly overlaps c. When they are misregistered, c and $P(C)$ will be misaligned. Given a projected 3D and 2D curve pair $(P(C), c)$, take a set of evenly spaced points $p_0, p_1, ..., p_n$ along $P(C)$. For each point p_i, the corresponding point q_i on c is located by computing the intersection of c with the line perpendicular to the tangent of $P(C)$ at p_i. If more than one intersection exists, the closest is selected. If there are no intersections, then p_i does not have an analogue on c. Fig. 3 illustrates a few cases. Note that either c or $P(C)$
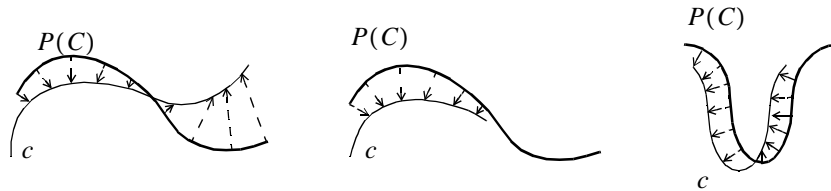

Fig. 3: Computing the disparity between 2D and projected 3D skeletons

may be incomplete. Fig. 3 (middle) illustrates point pairing when c is incomplete.

The disparity value is taken to be the mean-square value of all point pairs (p_i, q_i) for all curve pairs $(P(C), c)$. In our registration algorithm, the association of 3D skeletons with their 2D counterparts is performed by the user. Under perfect registration, the disparity is at most small positive value. When misregistered, a significantly larger non-zero value is obtained. The disparity value changes according to the projection P. We write $Cost(P)$ to denote the objective function.

3.3 Optimizing $Cost(P)$

Various methods for optimizing $Cost(P)$ exist [5]. Liu [1] describes an efficient method based on computing the partial derivatives of P with respect to the pose (registration) parameters. The method is similar to that described in [6]. This section summarizes the discussion in [1].

Since rigid registration is assumed, P can be expressed as a six-parameter function (3 rotations and 3 translations). Let X be the vector of parameters. Let P_X be the projection computed using the parameters in X. Given an approximation X_0 to the final solution, the algorithm uses Newton's method to compute the actual solution X_{sol}. From section 3.2, a set of point pairs (p_i, q_i) for each curve pair $(P(C), c)$ can be computed. The disparity function $Cost(P)$ is minimal when $p_i = q_i$ for all point pairs. A refinement X_1 to X_0 is obtained by computing a least squares solution for ΔX in $q_i = p_i + Jacobian(P)\big|_{\vec{0}} \bullet \Delta X$ for all points (p_i, q_i). $Jacobian(P)\big|_{\vec{0}}$ is the partial derivative of P w.r.t. X evaluated at $X = \vec{0}$ and ΔX is a 6-dimensional correction vector. Given ΔX, X_1 can be computed. The process is repeated for $X_2, X_3, \ldots, X_n$ or until $Cost(P_{X_n})$ falls below an arbitrary threshold.

4 Experiments

This section describes experiments to evaluate the algorithm's accuracy and performance under various conditions. Both simulated and actual X-ray angiograms are used. Using simulated X-ray angiograms with known poses permit registration accuracy to be quantified. The actual pose is perturbed by arbitrary amounts to derive the initial approximation. The effect of the perturbation is to displace all points in the 3D volume from their actual positions. The registration algorithm is used to correct the perturbation. The total number of 3D skeletons available is considerably larger than the set of 3D skeletons actually used for registration. By computing the difference between the initial and final displacements of *all* 3D skeletons, a measure of the algorithm's accuracy in recovering the initial pose can be determined. In the following discussion, the function *MAX(pose1,pose2)* computes the maximum displacement among all 3D skeletons between two states. For example, *pose1* may be the actual pose and *pose2* may be the computed pose returned by the algorithm. The function *MIN(pose1,pose2)* can be similarly defined.

Experiments in sections 4.1 through 4.3 used two 3D MRA studies of the head as input. Study A is a $256 \times 256 \times 61$ scan with voxel size $0.78mm \times 0.78mm \times 1.3mm$. Study B is a $256 \times 256 \times 48$ scan from a different patient. The voxel size for this study was $0.62mm \times 0.62mm \times 1mm$. In both studies, the imaging parameters were chosen to highlight intracerebral vessels. For each study, a set of 3D curves representing the central axis of intracerebral vessels was extracted using Aylward's algorithm (section 3.1). 204 curves were extracted from Study A while 223 vessels were extracted from

Study B. Simulated X-ray angiograms were generated by applying a perspective projection on these segmented vessels.

In contrast, the experiment in section 4.4 used actual patient angiograms. A digitally subtracted angiogram was acquired ($XRAY_1$). A simple algorithm was used to correct for image distortion. In addition, an MRA dataset (Study C) of size $256 \times 256 \times 72$ was acquired from the same patient. The voxel size was $0.86mm \times 0.86mm \times 1.1mm$.

In the following sections, details of each experiment are presented together with the results. Section 4.1 is a test of the algorithm's interactive performance. Section 4.2 evaluates the algorithm's ability to converge. Section 4.3 describes an experiment to test the algorithm's sensitivity to the choice of registration primitives. Sections 4.1 through 4.3 use synthetic 2D images. Section 4.4 describes an experiment to evaluate the algorithm using actual angiograms.

4.1 Interactive Performance Test

This experiment evaluates the algorithm's interactive performance. Three trials were conducted. Trials 1 and 2 were conducted using Study B whereas trial 3 was conducted using Study A. In each trial, a simulated X-ray angiogram was generated from an arbitrary viewpoint and the pose noted. A perturbation was introduced to the actual pose to produce the initial pose. The amount and nature of perturbation differed for each trial. To ensure fairness, the individual generating the angiograms and perturbations was different from the individual performing the registration. The actual pose was not known to the latter until after the experiment. The individual performing the registration was free to extract a number of 2D skeletons from each simulated angiogram. Between 27 and 37 curves per image were extracted using Fritsch's algorithm (section 3.1). Correspondence between the 2D and 3D skeletons was manually established. The registration algorithm was executed until $Cost(P)$ did not show any further improvement. The current solution at that point was noted. The program required approximately 3-5 minutes of run time on an HP 712/80 workstation with 64Mb of memory.

Table 1 shows the results. In each case, *MAX(actual,initial)* was in the range of a few centimeters. That is, the initial misregistration displaced all 3D skeletons on the order of centimeters. After registration, *MAX(actual,final)* was in the range of tenths of a millimeter. That is, the largest amount of misregistration among all 3D skeletons is in the tenths of a millimeter. We emphasize that we measure the misregistration of *all* 3D skeletons (more than 200), not just the skeletons used for the registration (27 to 37). The results indicate that the registration algorithm performs very well in an interactive environment. The set of 2D primitives was selected with no particular limitations other than ensuring that curves were chosen from all parts of the projected image. The algorithm is sufficiently robust to converge to the true solution in each case. This suggests that the algorithm is relatively insensitive to the choice of initial approximation as well as having a strong tolerance to the choice of registration primitives. The next two sections provide additional proof to these claims.

Exp.	Max. initial misregistration (cm)	Max. final misregistration (cm)
1	4.04×10^0	1.16×10^{-2}
2	2.85×10^0	9.60×10^{-2}
3	5.03×10^0	9.59×10^{-2}

Table 1: Results of registration accuracy test.

4.2 Test of Robustness

This is a test of the algorithm's ability to converge to the true solution given a range of initial poses. A good registration algorithm should be relatively insensitive to the choice of starting pose, quantified by the *capture radius*. The algorithm is said to have a capture radius of at least r if it converges to the correct solution from any choice of initial pose that displaces all 3D skeletons by *at least r*. In practice, there are an infinite number of possible starting positions where the initial misregistration is at least r. Moreover, the capture radius is partially dependent on the 2D and 3D images used. An approximation of the capture radius can be determined by using a large number of trials and datasets representative of typical cases.

Two experiments were conducted. The first used Study A, and the second used Study B. For each experiment a simulated angiogram was generated from an arbitrary pose. Thirteen 2D curves were extracted from each angiogram. These curves form the set of 2D registration primitives. Correspondence between 2D and 3D curves was established manually. For each experiment, 100 trials were performed. In each trial, the actual pose was given random perturbations to produce an initial pose. Up to $\pm 30°$ rotation and ± 10 cm translation in all three coordinate axis were used for experiments on Study A. Up to $\pm 25°$ rotation and ± 5 cm translation in all three coordinate axis were used for experiments on Study B. For each perturbation, $MIN(actual, initial)$ was determined. That is, the set of more than 200 3D skeletons was initially *at least* that distance away from their actual position. The registration algorithm was executed until $Cost(P)$ fell below 1.75×10^{-5} cm^2 or 100 until iterations have occurred. The computed solution was noted. $MAX(actual, comp)$ was determined. That is, the set of all 3D skeletons was displaced *at most* that distance from their final position after registration.

Fig. 4 are scatterplots of the results. The abscissa gives the initial minimum misregistration. That is, all 3D skeletons were displaced from their true position by *at least* that amount. The ordinate gives the maximum final misregistration. That is, all skeletons were displaced by *at most* that amount after registration. Apart from a single outlier in Study B, the capture radius is at least 15 cm in both studies. For each experiment, the set of trials formed two distinct groups. One group terminated with misregistrations in excess of 10 cm whereas the other has misregistrations less than 0.1 cm. This suggests that the algorithm's objective function is remarkably free of

local minima over a wide region surrounding the actual solution. Either the algorithm converged to the true solution with little residual error or it did not converge at all.
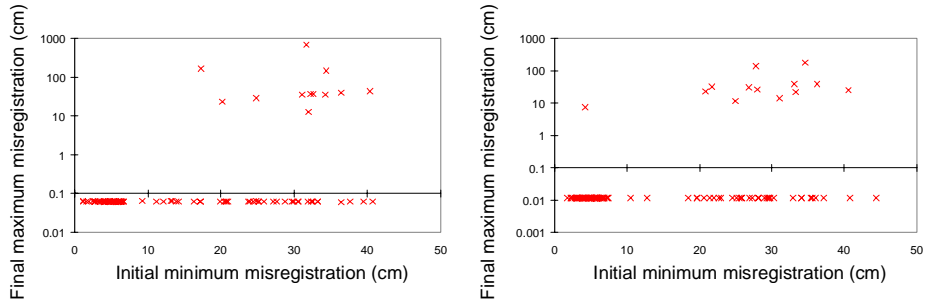

Fig. 4: The results of the capture radius experiment. Scatterplots for Study A (left) and Study B (right)

4.3 Sensitivity to Choice of Primitives

This experiment evaluates the algorithm's performance as a function of the number of curves used and their 3D spatial distribution. Two experiments were performed. The first used Study A whereas the other used Study B. The same angiograms and sets of 2D skeletons described in section 4.2 were used. In each experiment, an initial pose was generated by applying a perturbation to the actual pose. Each experiment has three series of trials. For each series, a random subset of k 2D curves were chosen of the 13 available. In every series, k was chosen to be 4, 8, and 10. Fifty trials were held for each series. Fig. 5 illustrates the organization of this experiment.
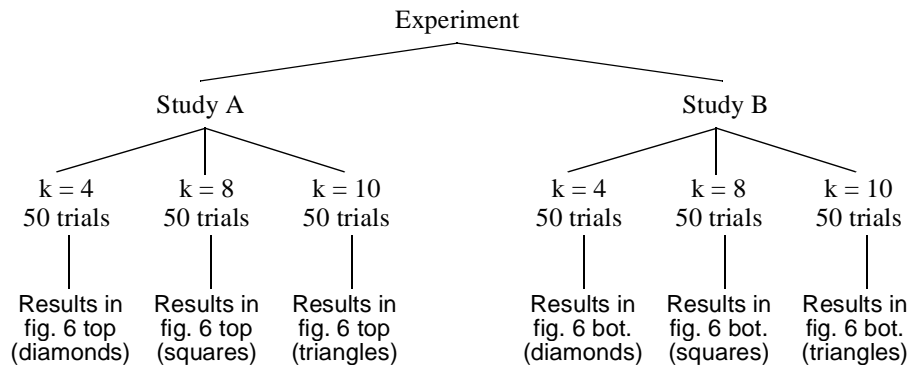

Fig. 5: Organization of experiment to evaluate sensitivity to choice of primitives

In each trial, the 3D spatial distribution of the 3D curve primitives were quantified using $Spread_{3D}$, a moment of inertia measure sensitive to the number of curves used as well as the spatial distribution or "spread" of 3D primitives [1]. The algorithm was

executed until $Cost(P)$ fell below 1.75×10^{-5} cm^2 or until 100 iterations have occurred. The computed pose was noted and $MAX(actual, comp)$ determined.

Fig. 6 plots the results. The absicca gives the value of $Spread_{3D}$. The ordinate gives the amount of residual misregistration when the algorithm halts. While four curves converged with submillimeter accuracy in some cases, the number of curves used is too small to be reliable. Using eight curves, the algorithm is very likely to succeed. With ten curves, there is virtual certainty. As the number of curves increased, $Spread_{3D}$ also increased. Having a larger value for $Spread_{3D}$ generally produced more accurate results. Spatial distribution is not just a function of the number of curves used. In both experiments, there is considerable overlap between the range of $Spread_{3D}$ achieved using 8 and 10 curves. This suggests that a well chosen but smaller set of curves can perform just as well as a larger set of poorly chosen curves.
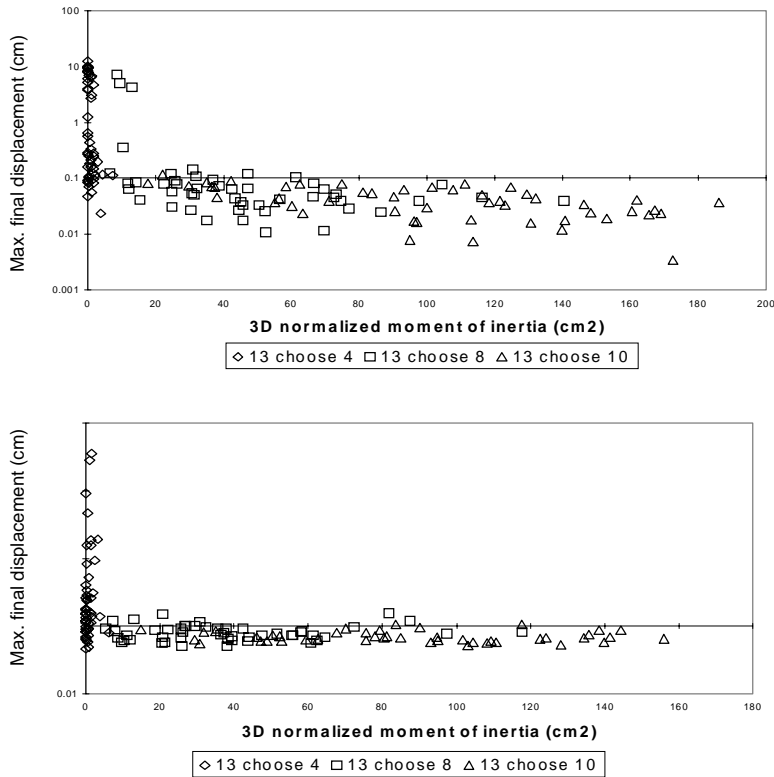

Fig. 6: Choice of registration primitives experiment. Top: Study A. Bottom: Study B

4.4 Clinical Test

Unlike the previous experiments which used synthetic angiograms, this experiment evaluates the algorithm's performance by registering an actual patient angiogram

($XRAY_1$) with its corresponding MRA (Study C). True angiograms have more sources of errors and present a greater challenge to the algorithm. With actual angiograms, the true pose is not known. We evaluated the algorithm's performance by comparing its performance with a manual registration using the same images. A neurosurgeon familiar with the patient registered the projected 3D skeletons against the 2D image. The manual attempt was then compared with the algorithm's solution. Fig. 7 (top) illustrates the results. The attempt used an anterior-posterior (AP) orientation as the starting pose. The effort took approximately one hour. The bottom figure was registered using our algorithm from the same starting pose. The attempt took approximately 10 minutes in total. Approximately 5 minutes of this time was spent by the user determining the correct correspondence between 2D and 3D skeletons.

The registration results appear very similar, indicating that the algorithm performs at least as well as manual registration but requiring only a fraction of the latter's time. In some places (notably vessel "B"), the algorithm performed noticeably better than the neurosurgeon. Since the entire intracerebral circulation was not highlighted, some 3D vessel skeletons did not have a corresponding 2D vessel. For example, vessels "A" and "C".
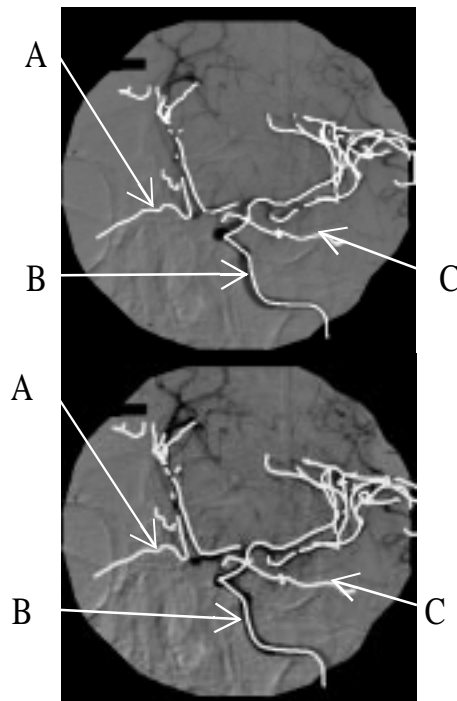

Fig. 7: Results of registration experiment on clinical data. Top: Manual registration by neurosurgeon. Bottom: Registration via our algorithm

5 Discussion

3D/2D registration algorithms have applications in surgical instrument guidance (e.g., [7], [8], and [9]) and in anatomical model synthesis (e.g., [10]). An accurate, robust registration algorithm significantly improves the accuracy of reconstruction and surgical guidance. Sections 4.1 through 4.4 demonstrated our algorithm's accuracy and robustness. The algorithm is capable of submillimeter accuracy over a wide range of starting poses and choice of primitives.

An advantage of our registration method is the use of curves as registration primitives. Point based algorithms (e.g., [11] and [12]) generally rely on a few landmarks for registration. These methods have the advantage of computational simplicity and thus speed. However, errors in locating fiducial points can result in decreased accuracy. Some surface based algorithms register the surface's silhouette with the 2D image. However, determining the silhouette may be computationally expensive or may require auxiliary data structures. In contrast, the projection of a curve is generally still a curve. Curves may be treated as a locus of points. By using curves as registration primitives in our algorithm, we retain the computationally simplicity of point based registration without sacrificing registration accuracy. By the principle of projective invariance, registering the skeletons of tubular structures such as blood vessels is equivalent to directly registering the vessels themselves. Thus, our algorithm is well suited to the task of registering X-ray angiograms to 3D MRA.

The use of cores as a method of extracting primitives enhances the accuracy and robustness of our algorithm. Cores are remarkably robust in the presence of image noise and differences in image resolution [13]. In addition, cores are little affected by variations in normal vessels which do not have perfectly circular cross-sections.

The algorithm as described requires both 2D and 3D tubular anatomical objects. For some applications, such structures may not be available in adequate numbers. We have proposed a novel method to overcome this problem [9].

One shortcoming of our present implementation is the need for manual correspondence between 2D and 3D vessel skeletons. Since angiograms are typically taken from standard poses, it may be possible to automatically associate the vessels based on their relative projected positions on the 2D image [14].

6 Conclusion

This paper described our method of 3D/2D registration. Our algorithm is based on the principle of projective invariance which permits the skeletons of tubular anatomical structures (e.g., blood vessels) to be used as registration primitives. Using curves as registration primitives is computationally straightforward and does not sacrifice accuracy. Experiments under test conditions where truth is known show that our method is capable of submillimeter accuracy. Elaborate preprocessing is unnecessary. Our algorithm is robust, and converges to the true solution even from large initial misregistrations. In addition, our algorithm requires only a small number (typically less than 10) of curve pairs to achieve submillimeter accuracy. An experiment conducted using actual X-ray and 3D MRA studies suggest that the algorithm is at least as accurate as a manual registration performed by an expert, but takes only 20% of the time required by the latter.

7 Acknowledgments

References

1. Liu A. "3D/2D registration and reconstruction in image-guided surgery." Ph.D. Dissertation, University of North Carolina at Chapel Hill, 1998.
2. Pizer S.M., Eberly D., Morse B.S., Fritsch D. "Zoom-invariant vision of figural shape: The mathematics of cores" Computer Vision and Image Understanding Vol. 69 No. 1, January, pp. 55-71, 1998, Article No. IV970563. Also available as UNC Technical Report TR96-004.
3. Fritsch D.S., Eberly D.H., Pizer S.M., McAuliffe M.J. "Stimulated cores and their applications in medical imaging" IPMI '95: Information Processing in Medical Imaging, 1995. pp. 365-368.
4. Aylward S., Pizer S., Bullitt E., Eberly D. "Intensity ridge and widths for tubular object segmentation and description." Proc IEEE WWMMBIA (IEEE96TB100056), 1996. Also available as UNC Technical Report TR96-018.
5. Press W.H., et. al. "Numerical recipes in C: The art of scientific computing." Cambridge [Cambridgeshire]: New York: Cambridge University Press, 1992.
6. Lowe D.G. "Three-dimensional object recognition from single two-dimensional images" Artificial Intelligence, Vol. 31, No. 3, 1987. pp. 355-395.
7. Grimson W.E.L., Ettinger G.J., White S.J., Gleason P.L., Lozano-Perez T., Wells W.M., III; Kikinis R. "Evaluating and validating an automated registration system for enhanced reality visualization in surgery." Computer Vision, Virtual Reality and Robotics in Medicine, 1995 pp. 3-12.
8. Lavallee S., Cinquin P., Szeliski R., Peria O., Hamadeh A., Champleboux G., Troccaz J. "Building a hybrid patient's model for augmented reality in surgery: A registration problem." Comput. Biol. Med. Vol 25, No. 2., 1995. pp. 149-164.
9. Liu A., Bullitt E., Pizer S.M. "Surgical instrument guidance using synthesized anatomical structures" CVRMed-MRCAS '97. First Joint Conference, Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery. pp. 99-108.
10. Bullitt E., Liu A., Pizer, S.M. "Three-dimensional reconstruction of curves from pairs of projection views in the presence of error. I. Algorithms." Medical Physics, vol. 24, no. 11, Nov. 1997. pp. 1671-1678.
11. Mellor J.P. "Realtime camera calibration for enhanced reality visualization." Computer Vision, Virtual Reality and Robotics in Medicine. First International Conference, CVRMed '95. Proceedings. pp. 471-475.
12. Uenohara M., Kanade T. "Vision-based object registration for real-time image overlay." Comput. Biol. Med. Vol. 25, No. 2. pp. 249-260.
13. Morse B., Pizer S., Fritsch D. "Robust object representation through object-relevant use of scale", Medical Imaging '94, Image Processing, SPIE 2167. pp. 104-115.
14. Fritsch D., Pizer S.M., Yu L., Johnson, V., Chaney E. "Segmentation of medical image objects using deformable shape loci." IPMI '97, Lecture Notes in Computer Science (Duncan J., Gindi G., eds), 1230, 1997. pp. 127-140.