

CT from an Unmodified Standard Fluoroscopy Machine Using a Non-reproducible Path

Chris Baker¹, Chris Debrunner¹, Mohamed Mahfouz^{2,3},
William Hoff¹, Jamon Bowen¹

¹Colorado School of Mines, Golden, Colorado

²University of Tennessee at Knoxville

³Oak Ridge National Laboratory

Abstract. 3D reconstruction from image data is required in many medical procedures. Recently, the use of fluoroscopy data to generate these 3D models has been explored. Most existing methods require knowledge of the scanning path either from precise hardware, or pre-calibration procedures. We propose an alternative of obtaining this needed pose information without the need of additional hardware or pre-calibration so that many existing fluoroscopes can be used.

Our method generates 3D data from fluoroscopy collected along a non-repeatable path using cone-beam tomographic reconstruction techniques. The novelty of our approach is its application to imagery from existing fluoroscopic systems that are not instrumented to generate pose information or collect data along specific paths. Our method does not require additional hardware to obtain the pose, but instead gathers the needed object to camera pose information for each frame using 2D to 3D model matching techniques [1-3]. Metallic markers are attached to the object being imaged to provide features for pose determination. Given the pose, we apply Grangeat's cone-beam reconstruction algorithm to recover the 3D data.

In developing this approach, several problems arose that have not been addressed previously in the literature. First, because the Radon space sampling is different for each scan, we cannot take advantage of a known Radon space discretization. Therefore we have developed a matching score that will give the best Radon plane match for the resampling step in Grangeat's approach [4]. Second, although we assume Tuy's condition [5] is satisfied, there are sometimes data gaps due to discretization. We have developed a method to correct for these gaps in the Radon data.

1 Introduction

Many medical procedures require three dimensional (3D) imaging of the human body. For example, when planning or executing surgical procedures, 3D imaging yields tremendous advantages over conventional 2D static radiographs. 3D image data can be obtained from CT and MRI, but these procedures may not be available in many clinics and operating rooms, and are therefore not as convenient as 3D from a fluoroscopy machine.

Fluoroscopy machines are much less expensive than CT or MRI instruments, and are widely available in operating rooms and clinics. However, a fluoroscopy image is

only a single plane projection, similar to a 2D X-ray. To obtain 3D data, the fluoroscope can be rotated in order to get multiple views of the object, and these projections can be processed using cone-beam tomographic reconstruction methods [4] to produce a 3D volumetric image similar to that obtained from CT.

Recently, Siemens has created a 3D fluoroscope using cone-beam tomography, called the SIREMOBIL Iso-C 3D. This instrument has a fluoroscope on a mobile C-arm that rotates automatically in a 190-degree arc around the patient. General Electric has produced a similar 3D fluoroscopic imaging and navigation system called FluoroCAT, an add-on for their OEC® 9800 Plus fluoroscope.

Although the method used in these systems gives good results, it requires a specialized fluoroscope that is designed to precisely and automatically rotate the source/detector (*SD*) around an axis, along with instrumentation to accurately measure the rotation. There are a great numbers of existing fluoroscopes in use that do not have this capability.

We have developed a method of generating 3D volumetric intensity data from fluoroscopy x-ray projections collected along a non-repeatable path using exact cone-beam reconstruction techniques. The novelty of our approach is that it can be applied to imagery from existing fluoroscopic systems that are not instrumented to reproduce specific paths or to generate accurate pose information for each image. Therefore our method does not require the addition of any external hardware to control or sense the motion of the fluoroscopy unit, but instead gathers the needed object to camera pose information for each frame using 2D to 3D model matching techniques [1-3]. In developing this method, several problems arose that have not been addressed previously in the literature. First, because the discrete sampling of the Radon space is different for each set of fluoroscopy images, we are not able to take advantage of a known discretization of our Radon space. Therefore we have developed a Radon plane matching score that will find the best matching plane during the resampling step in Grangeat's approach [4]. Second, although we assume that Tuy's condition [5] is satisfied, in practice there are sometimes gaps in the data due to discretization. We have developed a method to correct for any such missing Radon data. Although our implementation assumes that internal camera parameters are fixed, it could easily be modified to remove this assumption as was done in [6].

This approach allows much greater flexibility in imaging configurations, even allowing reconstruction from images collected as the patient moves in the field of view of a fixed fluoroscope. Variations of this approach could also be applied to 3D imaging using highly portable systems such as might be used in field hospitals in disaster relief or combat situations, as well as applications in non-destructive testing.

2 Previous Work

2.1 Pose Estimation

Tomographic reconstruction from projections requires accurate knowledge of the position and orientation (referred to collectively as the pose) of the *SD* relative to the object being scanned. A standard technique of determining the pose of a 3D object

from a two dimensional image is to incorporate a set of fiducial markers of known 3D structure into the object. Once the images of these fiducial markers are correctly associated with their 3D counterparts, the pose of the object with respect to the imaging sensor can be determined using an “absolute orientation” algorithm [7]. Examples of 2D-to-3D registration in X-ray images using the fiducial technique include [1-3].

2.2 Cone Beam Tomographic Reconstruction

The Siemens SIREMOBIL Iso-C^{3D} mentioned in Section 1 is an example of a commercial cone-beam tomography unit. In addition, many approaches to cone-beam tomography have been developed or enhanced in earlier work including many specific path algorithms [4, 5, 8-18]. Whereas most commercial systems use the Feldkamp method [8], we have selected the method of Grangeat [4] as the most suitable for our problem since it is an exact cone-beam reconstruction method. We have modified the common resampling step of Grangeat’s algorithm in a way similar to [19] to allow it to be applied to arbitrary discrete sets of SD locations. Because our path may be different for each scan, our sampling of the Radon space will be different for each scan. We therefore have developed a new method of choosing the best-matching Radon plane.

3 Approach

The processing required for reconstruction from an arbitrary path consists of two steps: recovery of the unknown pose of the SD , and cone-beam tomographic reconstruction. We assume that the images being used to perform this reconstruction have already been calibrated in such a way that the common pinhole camera model is an acceptable projection model [6] which removes the geometric distortions such as pincushion or barrel distortions. We also assume that the pixel values have been calibrated to be proportional to the integral of the density along the 3D ray projecting to each pixel. The details of these calibration methods are available in [20]. The first step, recovery of the unknown pose is described in Section 3.1, followed by the description of the reconstruction process in Section 3.2.

3.1 Pose Estimation

In order to estimate the pose of the sensor in each image from the image data, we use markers at known 3D locations relative to the object and perform 2D to 3D model matching techniques. We use 2mm steel ball bearings rigidly attached to the object and uniformly distributed to cover the entire field of view.

To compute the pose of the sensor, we first find the correct associations (or correspondences) between image points and the 3D fiducial locations using an interpretation tree search method [21]. The correspondences are used to calculate a pose solution which is used to project the object points back onto the image to determine how accurately the pose solution fits the data. Once a set of correct

correspondences is found, we compute the pose using a non-linear least squares method [22]. The solution minimizes the squared error between the observed image points and the expected position of those points, based on the projected model and the hypothesized pose. Once this process is complete, the pose of the SD for each image can be used in Grangeat’s tomographic reconstruction algorithm to reconstruct the 3D volume.

3.2 Tomographic Reconstruction

The second step, cone-beam tomographic reconstruction, uses the algorithm developed by Grangeat [4]. Tomographic reconstruction using Grangeat’s method has been well published and so we will not provide the details of his process here [4, 12, 13, 18, 19, 23]. Grangeat provides us with a method of computing 3D Radon data from the 2D detector data of a cone beam scan by way of the first derivative of the detector plane Radon data. This 3D Radon data must then be resampled to a spherical coordinate system. More specifically, it is resampled from the available planes (AP) in the sensor coordinate system, which is attached to the SD , to the desired planes (DP) in the 3D spherical object coordinate system where the backprojections can be computed. Using the pose found for the SD for each image in the previous pose estimation step, we can resample our Radon data collected on our sensor array into 3D Radon data using a method similar to [18, 19]. However, unlike methods with fixed paths, we cannot take advantage of any regularity in the path to reduce computation or to precompute the resampling coefficients. One issue that must be addressed in resampling is how to find the DP closest to a particular AP . Unlike [19], we define the closeness entirely in terms of the distance between the planes in the reconstruction volume. The available planes are parameterized by the 2D Radon parameters (A, α) in the image plane, and the image number, $inum$. The desired planes are given in the object coordinate system, whose axis are $(\hat{i}, \hat{j}, \hat{k})$ with coordinates (x, y, z) . In the object coordinate space, we parameterize the DP by its unit normal $\hat{n}$ and its distance, ρ , from the origin. In order to get accurate reconstruction, it is crucial to match these two planes in the reconstruction volume as well as possible.

In [19] the plane matching method is based on the assumption that the DP should pass as close as possible to the vertex of the SD . Once the SD vertex position has been identified for each frame of the fluoroscopic image sequence the distance to the DP can be calculated and minimized to obtain the 2D Radon plane giving the best estimate of the plane. Then, within that image a search will determine which 2D Radon plane will match best. However, due to the discretization of the available 2D Radon dataset (the 2D Radon is typically computed on a discrete set of (A, α) values), there may be cases where the best AP will not be associated with the closest to the SD vertex. Fig. 1 helps to illustrate one situation where this could arise. Two consecutive SD positions are shown looking along the edge of the planes of interest. The first SD position is labeled $SD - A$ and the second $SD - B$. Because of the discretization of the 2D Radon transforms, the closest two AP s for $SD - A$ are $AP - A1$ and $AP - A2$ shown as the dotted lines in Fig. 1, and similarly for $SD - B$

shown as the solid lines. Notice that in this case, the vertex for $SD-B$ is closer to the DP , however, $SD-A$ will give a better approximation of the DP by using $AP-A1$. Similar situations can arise due to the limited field of view of the sensor. Even if the desired plane passes close to the SD vertex, the sensor's field of view might not cover that desired plane. For these reasons we chose to develop a new plane matching method that finds the closest plane in the reconstruction volume rather than the SD vertex location [19].




Fig. 1. The best matching AP may not always be found from the SD vertex position which is closest to the DP . This figure illustrates a case in which the discretization of the actual planes may be such that a plane whose SD vertex is closest to the DP does not contain the best estimate of the DP in the reconstruction volume. Notice how the Radon planes for $SD-A$ are closer to the DP , but the vertex position of $SD-B$ is closer. Here, most existing algorithms would choose $SD-B$ as the source location from which to get the Radon data, however it is clear that $SD-A$ would provide more accurate data.

Depending on our path taken during the data collection, we will have a set of images providing available Radon planes, AP , indexed by $(A, \alpha, inum)$. From these we must determine the appropriate $(\rho\hat{n})$ of the closest DP . Using similar triangles we compute the distance ρ as

$$\rho = \frac{AD}{\sqrt{A^2 + D^2}}, \quad (1)$$

where D is the perpendicular distance between the SD vertex and the image plane. We can write the AP normal in the image coordinates, (u, v, w) as

$$[\rho\hat{n}_u, \rho\hat{n}_v, \rho\hat{n}_w] = \left[\frac{\rho^2}{A} \cos \alpha, \frac{\rho^2}{A} \sin \alpha, \frac{A\rho}{\sqrt{A^2 + D^2}} \right]. \quad (2)$$

With the AP defined in the image coordinate system, we can use the pose of the SD to transform this vector to AP_o , which is the available plane in object coordinates. Now we can directly compare AP_o and the DP , which is also in object coordinates.

Our criterion for measuring the closeness of two Radon planes is related to the squared distance between the two planes within the reconstruction volume. We define two planes we would like to match (both in object coordinates) as

$$DP_i : [x, y, z] \bullet \hat{n}_{DP_i} = \rho_{DP_i} \quad (3a)$$

$$AP_i : [x, y, z] \bullet \hat{n}_{AP_i} = \rho_{AP_i} \quad (3b)$$

We can write the signed perpendicular distance from an arbitrary point $[x, y, z]$ in DP_i to AP_i as

$$[x, y, z] \bullet \hat{n}_{AP_i} - \rho_{AP_i} \quad (4)$$

Solving Equation (3a) for z ¹, substituting into (4), squaring to get the squared distances, and integrating within the reconstruction volume over $-d$ to d in x and y , gives us our match score, (MS), correlating to how well the two planes are aligned.

$$MS = \int_{-d}^d \int_{-d}^d \left\{ x \hat{n}_{AP_x} + y \hat{n}_{AP_y} + \hat{n}_{AP_z} \left[\frac{\rho_{DP_i} - x \hat{n}_{DP_x} - y \hat{n}_{DP_y}}{\hat{n}_{DP_z}} \right] - \rho_{AP_i} \right\}^2 dx dy \quad (5)$$

Each time a new plane is found that will project into a particular DP , we calculate the MS given by (5), compare it to the currently stored MS in the destination array (the distances are initialized to infinity) and keep the best match.

Because our source data is generated along an arbitrary non-repeatable path and therefore, the discretization of the Radon space may be different for each scan, we may have cases where there is no AP to support a particular DP . Tuy's data sufficiency criterion [5] is not strictly met in this case, and we refer to this case as "missing data." In order to handle any missing data that may arise from an arbitrary non-repeatable path, we use a binary array to record whenever a DP bin is filled, defined as $N(\rho \hat{n})$. After the complete dataset has been resampled from the detector coordinate system to the world coordinate system, we can backproject the array N along with the resampled data as described in [4]. Once N has been processed it will be a reconstructed volume weighted according to which portions of the volume have the most supporting data. We can then use N as a normalization factor for our reconstructed volume to help account for missing data artifacts.

¹ In general we need to solve for whichever variable, (x, y, z) is most aligned with the normal $\hat{n}$. However, for simplicity of the discussion, we just use z . Note also that this will avoid division by zero in (5) since the coefficient in the divisor will always be the greatest in magnitude of the three (i.e. not equal to zero).

4 Results

Since the accuracy of our reconstruction relies heavily on the computation of the pose of the SD for each image, we first quantify the error in the pose calculation by comparing estimated poses to ground truth measurements collected with an optical tracker. Then we present the tomographic reconstructions of synthetic and real data sets to characterize the quality of the reconstructions.

4.1 Pose Estimation Compared to an Optical Tracker

To test the pose estimation accuracy, we used an Optotrak optical tracker² to record ground truth poses during an image data collection of a physical phantom constructed of PVC and acrylic. The various coordinate systems we used are illustrated in Fig.2. There are four coordinate systems that we must be concerned with in order to perform the verification in this way. First, we refer to the Optotrak probe coordinate system as P and the Optotrak base coordinate system as our world coordinate system, or W since it is the reference from for the ground truth measurements from the Optotrak. Finally, we refer to the SD coordinate system of the fluoroscope as F and the object coordinate system origin as O .

In order to measure the pose accuracy, we compare the motion of the object from frame 0 to frame i using two different transform paths. The paths are illustrated in Fig. 3. The first path (shown on the bottom of Fig. 3) takes us from the P coordinate system's initial pose through the world reference coordinate system (W) to the i^{th} frame of P . The second path (shown across the top of Fig. 3) takes us from the same starting reference coordinate system of P through the initial frame of O to the SD reference frame, (F) (which is fixed during the data collection) back through the i^{th} frame of O , and into the i^{th} frame of P .




Fig. 2. Top view of Pose estimation verification setup. The P is attached to the outside of the phantom whose origin, O , is defined in the center of all the beads.

² The Optotrak used is the Northern Digital model #3020.




Fig. 3. The two paths used to verify the pose finding algorithm.

The transforms between W and P are calculated from the Optotrak data. The transforms between F and O are given from the pose estimation algorithm, which is the transform we wish to verify. The transform from P to O can be found with physical measurements since P and O are rigidly attached as Fig. 2 suggests. We denote the transform between O and P as H_O^P . Note also that because the probe is rigidly attached to the object, we know that $H_O^P = H_{O_i}^{P_i} = H_{O_0}^{P_0}$ and these notations will be used interchangeably to make consecutive transform cancellations more clear. We define the transformation from the F to O at the time of the i^{th} image as $H_F^{O_i}$, and the transform between W and P at the time of the i^{th} image as $H_W^{P_i}$.




Fig. 4. The Euclidean distance between the probe origin found from the pose finding algorithm to the origin found using the Optotrak method (left). The error in the rotation angle between two consecutive image frames for 67 frames in the sequence. This error is calculated as the rotational error from the angle-axis representation of the residuals from the two pose calculation paths (right).

We measure our pose error by finding the change in the pose of P relative to its initial position, or $H_{P_0}^{P_i}$, in two ways, one using the Optotrak and the other using our

pose estimation algorithm, or the two paths shown in Fig. 3. Then we compare the results to measure our error. We call the result of the first (using the Optotrak) ${}^O H_{P_0}^{P_i}$, and readily see that it can be computed as

$${}^O H_{P_0}^{P_i} = H_W^{P_i} H_{P_0}^W \quad (6)$$

We call the result of the second method (using our pose estimation algorithm) ${}^A H_{P_0}^{P_i}$. Given the transform H_p^O from the physical measurements, we see that we can compute ${}^A H_{P_0}^{P_i}$ from the transforms $H_{O_0}^F$ and $H_F^{O_i}$ computed by our pose estimation algorithm as

$${}^A H_{P_0}^{P_i} = H_{O_i}^{P_i} H_F^{O_i} H_{O_0}^F H_{P_0}^{O_0}. \quad (7)$$

By comparing these two results we can analyze the error in our pose finding method. We compare these two paths by computing

$$\left[{}^O H_{P_0}^{P_i} \right]^{-1} {}^A H_{P_0}^{P_i}. \quad (8)$$

If the two transforms were equal this would produce an identity matrix. We convert the relative rotation from this transform to angle-axis form and use the magnitude of the angle as our orientation error measure, and the magnitude of the translation as our position error measure. The position errors are shown in Fig. 4 (left), while the orientation errors are shown in Fig. 4 (right) for 63 frames. The mean translation error is 3.37mm, with a standard deviation of 1.53. The mean error in the angular deviation is 0.023 degrees, with a standard deviation of 0.0104.

These results demonstrate that our pose recovery algorithm is accurate when compared to a commercial optical tracker. Because the accuracy of our reconstruction depends primarily on the accuracy of pose, this result is critical to the success of our approach. Next we consider the reconstruction of various objects using our proposed method.

4.2 Tomographic Reconstruction Results

The next step in the analysis of our method is to evaluate the results of the complete reconstruction process. First we evaluate reconstructions from a set of synthetic phantom images and then from real imagery of a cadaver knee.

Synthetic Phantom Reconstruction. In order to verify the reconstruction algorithm independently of the pose finding algorithm, we generated synthetic data by integrating densities along a cone-beam projection pattern through a synthetically created 3D volume. The synthetic volumetric data used was a set of ellipsoids. Fig. 5 shows the rendered reconstructed volumetric results on the right with the original data on the left. The boundaries of the various ellipsoids can be seen clearly.




Fig. 5. Images taken of the reconstructed synthetic phantom. There are two different views of the same reconstructed volume with the original volume displayed next to it. The original volumes are shown on the left while the reconstructed volumes are shown on the right.

We used these synthetic reconstruction results to analyze the accuracy of the recovered densities relative to the original volumetric data, and to measure the geometric accuracy of our reconstruction algorithm.

Density analysis. In order to test the accuracy of our reconstructed voxel density information we can compare the final reconstructed voxel values with the density values of the original data. Since the most common use of the 3D density data is segmentation, ensuring accurate reconstructed density values is not as important as verifying relative density information. Therefore we fit a linear function to transform our reconstructed density information to the original density information and then compare to the linearly transformed data. We found the mean signed deviation to be 3.64 (from a range of zero to 255), and the standard deviation was 2.62.

Geometric accuracy analysis. In order to determine whether the reconstruction algorithm introduced any geometric errors (e.g., position offsets and scaling), we compared the position and size of each reconstructed ellipsoid to those of the corresponding ellipsoid in the original volume. We first found the three axes of each reconstructed ellipsoid by measuring the extremities of the segmented ellipsoid in the x , y , and z -directions. We then found the center of the ellipsoid by averaging the axis endpoints. The average of the absolute deviations of the center positions for the ellipsoids was 0.163 voxels, while the average of the absolute deviations from the axis lengths was 1.303 voxels. The errors in the axis lengths are larger than the error in the positions which is probably due to the blurring effect of the algorithm (accurate segmentation is very difficult with blurred edges, and a different segmentation method might reduce these errors). However, note the positions of the ellipsoids were found very accurately.

Cadaver Reconstruction. Fig. 6 shows our algorithm’s reconstruction of a cadaver knee from 596 fluoroscopic images taken at 30 frames per second. Some noise can be seen at the top and bottom of the reconstruction due to parts of the cadaver being in the field of view during only part of the scan, thus these portions of the reconstructed volume do not satisfy Tuy’s condition and are not a part of our supported reconstruction area. This effect has been addressed in the literature [24] and we do not address it here as it is out of the scope of this particular paper. We also notice some

metal artifacts from the marker beads which are also addressed extensively in the literature (e.g. [25-27]). Further refinements to our algorithms could incorporate these enhancements and remove these artifacts. However, this demonstrates that our method can tomographically reconstruct a volume from real fluoroscopic imagery with quality sufficient to allow segmentation (e.g., to produce a 3D model of the bone for surgical planning or surgical navigation).




Fig. 6. The reconstructed volume from projections generated along an arbitrary path using our proposed algorithm.

5 Conclusions and Discussion

This paper presents results of a method of generating 3D volumetric data from fluoroscopy using image-based pose estimation and cone-beam reconstruction methods. The novelty of our approach is that it can be used with existing generic fluoroscopic systems that produce scans along a new path for each scan, and that have no hardware to control or sense the movement of the fluoroscopy system.

Because the discrete sampling of the Radon space is different for each scan of images, we are not able to take advantage of a known discretization of our Radon space, which has led us to develop a Radon plane matching score that will give the best plane match for the resampling step in Grangeat's approach. This allows any arbitrary and non-repeatable sensor path that meets Tuy's condition to be used to generate the source data for reconstructions. Because of the arbitrary discretization of the Radon space we have also developed a method to help correct for any missing Radon data. We have demonstrated the feasibility of our approach through accurate preliminary reconstructions of phantom data and a cadaver knee and have as well performed analysis of the accuracy of our pose finding method and the reconstruction.

The results from this work have the potential to provide an inexpensive and readily available 3D imaging capability that can benefit many medical applications. Advances in visualization and localization based on this method will allow minimally invasive procedures to be used more frequently, thus reducing the trauma to the patient. This 3D imaging capability can also make computer-aided procedures more accessible. It is expected that long term health care costs will be reduced because of

the high efficiency of computer-aided procedures, and from the shorter hospitalization and recovery times that will be required [28].

References

1. Kall, B.K., *Comprehensive multimodality surgical planning and interactive neurosurgery*, in *Computers in Stereotactic Neurosurgery*, P.J. Kelly and B.K. Kall, Editors. 1992, Blackwell Scientific: Boston. p. 209-229.
2. Penney, G., et al., *A comparison of similarity measures for use in 2D-3D medical image registration*. IEEE Trans. on Medical Imaging, 1998. **14**(6): p. 586-595.
3. Weese, J., et al., *Voxel-based 2D/3D registration of fluoroscopy images and CT scans for image-guided surgery*. IEEE Trans. Information Technology in Biomedicine, 1997. **1**(4): p. 284-293.
4. Grangeat, P., *Mathematical Framework of Cone Beam 3D Reconstruction via the First Derivative of the Radon Transform*, in *Mathematical Methods in Tomography*. 1992: Springer Verlag. p. 66.
5. Tuy, H.K., *INVERSION FORMULA FOR CONE-BEAM RECONSTRUCTION*. SIAM Journal on Applied Mathematics, 1983. **43**(3): p. 546-552.
6. Mitschke, M. and N. Navab, *Recovering the X-ray projection geometry for three-dimensional tomographic reconstruction with additional sensors: Attached camera versus external navigation system*. Medical Image Analysis, 2003. **7**: p. 65-78.
7. Horn, B.K.P., *Closed-form solution of absolute orientation using unit quaternions*. Journal Optical Soc. of America, 1987. **4**(4): p. 629-642.
8. Feldkamp, L.A., L.C. Davis, and J.W. Kress, *Practical cone-beam algorithm*. Journal of the Optical Society of America A: Optics and Image Science, 1984. **1**(6): p. 612-619.
9. Mueller, K., R. Yagel, and J.J. Wheller, *Anti-aliased three-dimensional cone-beam reconstruction of low-contrast objects with algebraic methods*. IEEE Trans. on Medical Imaging, 1999. **18**(6): p. 519-537.
10. Mueller, K., R. Yagel, and J.J. Wheller, *Fast implementations of algebraic methods for three-dimensional reconstruction from cone-beam data*. IEEE Trans. on Medical Imaging, 1999. **18**(6): p. 538-548.
11. Noo, F., M. Defrise, and R. Clack, *Direct reconstruction of cone-beam data acquired with a vertex path containing a circle*. IEEE Trans. on Image Processing, 1998. **7**(6): p. 854-867.
12. Schaller, S., T. Flohr, and P. Steffen, *Efficient Fourier method for 3-D Radon inversion in exact cone-beam CT reconstruction*. IEEE Transactions on Medical Imaging, 1998. **17**(2): p. 244-250.
13. Axelsson, C. and P.E. Danielsson, *Three-dimensional reconstruction from cone-beam data in $O(N^3 \log N)$ time*. Physics in Medicine and Biology, 1994. **39**(3): p. 477.
14. Jacobson, C., *Fourier Methods in 3D-Reconstruction from Cone-Beam Data*. Linkoping Studies in Science and Technology Dissertations, 1996. **427**.
15. Kudo, H., et al., *Performance of Quasi-Exact Cone-Beam Filtered Backprojection Algorithm for Axially Truncated Helical Data*. IEEE Transactions on Nuclear Science, 1999. **46**(3): p. 608-617.
16. Turbell, H. and D. Per-Erik, *Helical cone-beam tomography*. International Journal of Imaging Systems and Technology, 2000. **11**(1): p. 91-100.
17. Wang, G., et al., *A General Cone-Beam Reconstruction Algorithm*. IEEE Transactions on Medical Imaging, 1993. **12**(3): p. 486-496.
18. Wang, X. and R. Ning, *Cone-beam reconstruction algorithm for circle-plus-arc data-acquisition geometry*. IEEE Transactions on Medical Imaging, 1999. **18**(9): p. 815-824.

19. Noo, F., R. Clack, and M. Defrise, *Cone-beam Reconstruction from General Discrete Vertex Sets using Rdon Rebinning Algorithms*. IEEE Transactions on Nuclear Science, 1997. **44**(3): p. 1309-1316.
20. Baker, C., *Computed Tomography from Imagery Generated by Fluoroscopy along an Arbitrary Path*, in *Engineering*. 2004, Colorado School of Mines: Golden. p. 115.
21. Grimson, W.E.L., *Object recognition by Computer*. 1990, Cambridge, Massachusetts: MIT Press.
22. Haralick, R. and L. Shapiro, *Computer and Robot Vision*. 1993: Addison-Wesley Inc.
23. Kudo, H. and T. Saito, *Derivation and implementation of a cone-beam reconstruction algorithm for nonplanar orbits*. IEEE Transactions on Medical Imaging, 1994. **13**(1): p. 196-211.
24. Hsieh, J., *Reconstruction Algorithm for Single Circular Orbit Cone Beam Scans*. IEEE, 2002: p. 836-838.
25. Benac, J., *Alternating Minimization Algorithms for Metal Artifact Reduction in Transmission Tomography*, in *Electrical Engineering*. 2002, Washington University: St. Louis, Missouri.
26. Wang, G., et al., *Iterative Deblurring for CT Metal Artifact Reduction*. IEEE Transactions on Medical Imaging, 1996. **15**(5): p. 657-664.
27. Zhao, S., et al., *X-Ray CT Metal Artifact Reduction Using Wavelets: An Application for Imaging Total Hip Prostheses*. IEEE Transactions on Medical Imaging, 2000. **19**(12).
28. Jolesz, F., Kikinis, and F. Shtern, *The Vision of Image-Guided Computer Surgery*. The High Tech Operating Room in Computer Integrated Surgery - Technology and Clinical Applications, ed. e.a. R. Taylor. 1996, Cambridge, Massachusetts: MIT Press. 717-721.