Calibrating cone-beam tomography systems with generic trajectories

Thomas Blumensath, Neil O'Brien, Charles E. Wood

Abstract

Iterative reconstruction of tomographic data relies on a precise knowledge of the geometric properties of the system. We are here interested in tomographic reconstruction from generic trajectories, that is, scan trajectories other than circular or helical. In this general setting, the estimation of geometric system parameters becomes more complex. We propose an approach that uses a calibration scan in which a fiducial marker is attached to a sample. The fiducial marker is then moved to different locations in the cone-beam. Detecting the location of the projection of the marker at different locations on a flat panel detector then allows us to use an optimisation algorithm to find the relevant system parameters. We explore this approach using a scanning system with a generic sample manipulator with several degrees of freedom.

Index Terms

Laminography, geometric calibration, x-ray tomography

I. Introduction

X-ray tomographic imaging technology is a valuable tool in medical diagnosis, scientific investigations, non-destructive testing and dimensional metrology. Traditionally, x-ray tomography relies on circular x-ray scanning trajectories where the x-ray source and detector perform a circular motion around the object under investigation (or, equivalently, the object is rotated around a fixed axis whilst source and detector remain fixed). Projection images are then taken at regular angular intervals. With this scan geometry,

T. Blumensath and C.E. Wood are with the ISVR Signal Processing and Control Group, University of Southampton, SO17 1BJ, UK, Tel.: +44 (0) 23 8059 3224 ,e-mail: thomas.blumensath@soton.ac.uk, Neil O'Brien is with the μ-VIS X-ray Imaging Centre, University of Southampton, UK

This research was supported by Innovate UK, the UK's innovation agency and the Aerospace Technology Institute (ATI), through grant 101804 (ProjecCAN).

Thanks goes to the ProjectCAN consortium

filtered back-projection type algorithms, such as the FDK algorithm [1] are fast and work well if enough projections are available. Imaging systems are regularly calibrated to ensure the rotation axis is positioned appropriately relative to the source and detector locations, which is a prerequisite for the use of these algorithms. Additionally, software corrections can be performed after data has been acquired. For example, the nominal position of the rotation axis can be shifted artificially in relation to the detector plane to provide sharper reconstructed volumetric images. Some software packages use automated calibration routines that estimate the spatial orientation of the rotation axis.

The use of x-ray tomography for non-destructive testing is often limited by the shape of components. Whilst traditional, circular scan trajectories work well for approximately cylindrical objects, flat components such as circuit boards, fossils or flat panels are often difficult to scan, either because the aspect ratio prevents the full rotation of the component or, because x-rays might be unable to fully penetrate the material along the long component direction whilst providing too little contrast along the short direction. In these cases, laminographic scanning is a viable alternative [2].

Furthermore, for cone-beam computed tomography systems, magnification depends on the ratio between the source to detector distance and the source to object distance. To achieve high magnification, the source needs to be close to the object. It would thus be desirable for objects with complex shapes to be able to use scan trajectories where the source can stay close to the object's surface.

To investigate the use of micro-focus x-ray tomography systems that use generic scan trajectories, we have recently implemented two generic sample manipulator systems [3]. The first system uses a high precision hexapod. A hexapod consists of two plates connected by six linear actuators allowing the top plate to move, tilt and rotate relative to the bottom plate. To increase possible motions, the Hexapod in our system is mounted on a linear stage and an additional rotation stage was mounted on the top plate of the hexapod. The second system used a six axis robot arm manipulator as commonly used in the automotive industry.

To reconstruct volumetric images from data collected with these systems, filtered back-projection algorithms are no longer suitable if non-standard trajectories are used. Instead, iterative methods are required. Iterative (or algebraic) reconstruction methods require knowledge of the system and scan trajectory geometry, which is used to define the tomographic forward model. Whilst it is well known that these models might not be invertible or might be ill-conditioned unless the scan trajectories satisfy certain constraints [4], we could show that even highly under-sampled data can lead to useful reconstructions where full CT is impossible [3]. However, any uncertainty in the specification of the system geometry will lead to a reduction in image resolution. Methods to measure these parameters are thus crucial.

A. Contribution

This paper looks at the problem of geometric calibration of a generic tomographic sample manipulator. We here use the term calibration not in the metrological sense of a traceable chain of measurements with specific uncertainties [5], but as a technique to estimate the location and orientation of the sample, the x-ray source and the x-ray detector during scanning. We propose an approach where calibration scans are performed to provide geometric information and to link this information to the information available from the manipulator's position encoders. By detecting and tracking features in the projection images during the calibration scan, we describe optimisation algorithms that use the manipulator's reported position to estimate source and detector location, detector orientation and the position and orientation of linear and rotational axes of the manipulator system.

Calibration in traditional, rotation based, X-ray tomography systems has been studied before. Purely data driven approaches try to find the rotation axis alignment based on projections that are offset by 180 degree [6], or optimise quality measures such as entropy to improve the resolution of the reconstructed image [7], thus estimating system parameters in the process. An alternative is the work in [8] where a geometrical model of the object is used for calibration. Another set of approaches uses calibration phantoms or other measurement devices. For example, [9] and [10] use fixed test objects for calibration, whilst atomic force microscopy has been used in [11].

Our work builds on the work in [9] and [10]. As in [9], we use calibration scans. Our calibration scans will use a small ($\sim$3mm) sphere that will be attached to the object. We then use the sample manipulator to move the sphere to different positions within the x-ray cone-beam. This is equivalent to the use of a test phantom constructed of several spheres placed at fixed and known distances and angles. Uncertainty in the accuracy of sphere location now depends on the accuracy of the sample manipulator system. This uncertainty could be assessed using co-ordinate measurement machines, though in this work we rely on the encoder accuracy in the manipulator and thus assume that the reported locations are the actually achieved locations. An evaluation of biases, non-linearities and variances in the sample manipulator's acuracy remains to be undertaken in the future.

As in [10], we define the geometry of our system with all its degrees of freedom and use an optimisation method to optimise the geometric parameters based on an estimate of the centres of the projected spheres. The difference here is however that our geometry is less constrained than the traditional cone-beam tomography geometry with a single rotation axis as used in [10].

B. Notation

We will use bold face roman characters to specify vectors in three dimensional Euclidean space. In particular,

- $\mathbf{S}$ - is the source location
- $\mathbf{D}_0$ - is the detector location
- $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{n}$ are orthogonal vectors that define the detector orientation, $\mathbf{u}$ points along detector rows and $\mathbf{v}$ points along detector columns
- $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$ are orthonormal co-ordinate axis that define the main co-ordinate system we will be using.
- $\mathbf{P}$ will be a generic point in space, such as the location of the calibration sphere.
- $\mathbf{d}(\mathbf{P})$ will be a two dimensional vector that defines the location of the projection of point $\mathbf{P}$ in terms of the two row and column detector co-ordinate vectors.

All vectors will be in relation to a co-ordinate system with zero point at one of the sphere locations. The directions of the coordinate axis are then the directions of the sample manipulator movement axes, which are here assumed to be orthogonal.

II. System setup

The geometry of an x-ray cone-beam projection system can be defined by the source location $\mathbf{S}$, the location of the detector $\mathbf{D}_0$ and by a local co-ordinate system on the detector defined by orthogonal vectors $\mathbf{u}$ and $\mathbf{v}$, so that each point on the detector can be written as $\mathbf{D}_P = \mathbf{D}_0 + d_1\mathbf{u} + d_2\mathbf{v}$. Note that $\mathbf{u}$ and $\mathbf{v}$ also define the orthogonal vector $\mathbf{n} = \mathbf{u} \times \mathbf{v}$. In general, source and detector locations can change from projection to projection, though for simplicity, we here assume that the source and detector location and orientation remain fixed during the scan[1].

The manipulator is assumed to provide linear motions along three orthogonal axis $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$ as well as rotational motion around a rotation axis $\mathbf{r}$. We here restrict the derivation to a single rotation axis, though the extension to several axes can follow a similar approach to the one we use here. We assume the three linear axes $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$ form the basis of our global co-ordinate system, that is, we express all other vectors as combinations of $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$.

[1]To extend our approach to allow movement of source and detector, additional calibration scans that link the source and detector movement axes to the system co-ordinates can be introduced.

III. CALIBRATION SCANS

In order to perform geometric calibration, we either need a calibrated test object with several features distributed in space or we use an object with a single feature that is moved to different spatial locations using the manipulator. We here follow the second approach which requires some form of calibration of the manipulator itself. Whilst we have done some initial calibration experiments with one of our manipulators, we here assume that the accuracy of the manipulators is high enough and that the positional readouts from the internal positional encoders is accurate to the required precision.

As test object, we use spheres, either ruby spheres or high precision ball bearings that are either attached to the test object or some other low attenuating holder. Whilst the cone-beam projections of spheres are not exactly circular [12], we here make this assumption.

A. Calibration of linear axis

We use a single sphere to estimate $\mathbf{S}$, $\mathbf{D}_0$, $\mathbf{u}$ and $\mathbf{v}$ relative to the three linear axes of the manipulator $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$. This is done using a calibration scan where projection images are taken at different spatial positions with different values for $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$. We here use 18 projections, taken with the manipulator performing linear motions to points $x\mathbf{x} + y\mathbf{y} + z\mathbf{z}$, where $x \in \{-dx, 0, +dx\}$, $y \in \{-dy, 0, +dy\}$ and $z \in \{-dz, 0, \}$. The values for dx, dy and dz depend on the magnification used in the scan and are chosen so that all points are projected onto the detector. We here did not use $z = +dz$ as our $\mathbf{z}$ axis is towards the detector and we needed to avoid collision with the x-ray source as the sample was mounted relatively close to the detector already in the 0 position.

B. Calibration of rotation axis

To calibrate the rotation axes, we use the same sphere attached to the same location on the sample as in the calibration of the linear axis. We also use two additional spheres mounted to the manipulator in different positions. We then rotate the axis through a range of rotational offsets (including a rotation of 0). The position of the sphere used for both, linear and rotational calibration should be in the same position when the rotation is set to 0 or when the linear shifts are set to 0. This is important to allow us to link the position of the rotation axis to the co-ordinate system $\mathbf{x}$, $\mathbf{y}$ and $\mathbf{z}$. Both calibration scans are thus performed one after the other with at least one sphere staying in place.

C. Analysing the projection data

We estimate the centres of the balls by manually identifying the different features in the image. This is done using a graphic user interface in which the user clicks on the sphere in each image. A mask is

generated around this pixel that is large enough to contain the entire sphere. Once the different features are identified in each image, a neighbourhood of each feature is selected using the masks. This region is then filtered using a gaussian filter with a standard deviation of five pixels before canny edge detection detects the outline of the spheres. After outlier detection, the edge points are fitted with a circle whose centre is used as the centre of the projected sphere.

The results are shown in figure 1, where we show six different projection images after linear (top) and rotational (bottom) displacement of the sample. Overlaid are all estimated sphere centres for the linear (top) and rotational (bottom) motions[2]. For the linear motion, we have used a single sphere to trace a 3 by 3 grid of locations at two magnifications. For the rotation, we have use seven projections at different rotation angles and estimated the centres of three spheres whenever these were projected onto the detector (On occasion, a sphere might be rotated fully or partially outside the cone beam so that no data is available for this sphere for this projection). Centres for different spheres are here shown in different colour. Note that one sphere was attached directly to the sample, whilst the other two spheres were mounted on a low attenuating acrylic block that was attached to the sample.

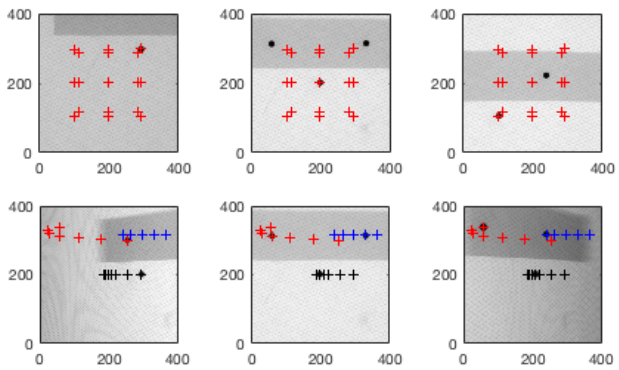

Fig. 1: Projection images at three different linear off-set locations (top) and three rotations (bottom). The (+) indicates the estimates for the centre of the spheres at all 18 linear locations (top) and the different rotations (bottom).

[2]Note that when plotting the crosses that identify sphere centres here, the crosses are not exactly centred on the spheres. This is an artefact of plotting the figures and does not reflect the actual estimates of the sphere centres.

IV. GEOMETRIC CALIBRATION

To estimate the system's geometric parameters, we model the centre of each sphere as a point in 3D space, which we parameterise as a function of the unknown system parameters and the linear shifts or rotation angles. We then model the projection of these points onto the detector plane. We optimise the unknown system parameters to minimise the distance between the predicted locations of the points on the plane and the actually observed centre positions of the projected spheres.

A. Cone-beam projection

Let us start by modelling the projection of a point in 3D space onto the detector. For any point $\mathbf{P}$ that lies between the source and the detector, the cone-beam projection of $\mathbf{P}$ onto the detector can then be written in two ways:

$$\mathbf{D}_0 + d_1(\mathbf{P})\mathbf{u} + d_2(\mathbf{P})\mathbf{v} \tag{1}$$

and

$$\mathbf{S} + \alpha(\mathbf{P} - \mathbf{S}), \tag{2}$$

where α is a positive scalar. This is shown in figure 2. Let us also define the point $\mathbf{D}_S$, which is the orthogonal projection of the source location onto the detector. To derive an expression for α, we can then consider the two triangles $\overline{\mathbf{D}_0\mathbf{S}\mathbf{D}_S}$ and $\overline{\mathbf{D}_P\mathbf{S}\mathbf{D}_S}$, which share the side $\overline{\mathbf{S}\mathbf{D}_S}$.

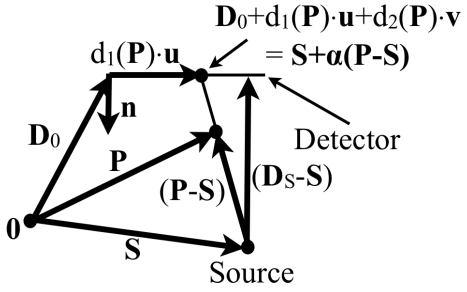

Fig. 2: 2D view along the direction $\mathbf{v}$ of the geometry of the cone-beam projection setting.

Note that the line $\overline{\mathbf{S}\mathbf{D}_S}$ is parallel to the normal of the detector plane, which is defined as $\mathbf{n} = \mathbf{u} \times \mathbf{v}$. Let us assume that $\mathbf{u}$ and $\mathbf{v}$ are orthonormal. This implies that $\mathbf{n}$ is of unit length as well. Let us also assume that $\mathbf{n}$ points from the detector towards the source, that is, $\langle \mathbf{D}_S - \mathbf{S}, \mathbf{n} \rangle < 0$. The length of the vector $\overline{\mathbf{S}\mathbf{D}_S}$ is thus $-\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle$. Similarly, the component of $\mathbf{P} - \mathbf{S}$ in direction $\mathbf{n}$ is $|\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle|$, where the negative sign comes from the assumption that $\mathbf{P}$ lies between the source and the detector. If

we instead assume that $\mathbf{n}$ points away from the source, then the signs are in the opposite direction. In either case, the ratio $\frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle}$ is always positive whenever the points $\mathbf{P}$ lie between the source and the detector.

As $\overline{\mathbf{D}_P \mathbf{S} \mathbf{D}_S}$ and $\overline{\mathbf{S} \mathbf{P} \mathbf{S} + \mathbf{n}\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle}$ are similar, we have the relationship:

$$\alpha \frac{\|\mathbf{P} - \mathbf{S}\|}{\|\mathbf{D}_S - \mathbf{S}\|} = \frac{\|\mathbf{P} - \mathbf{S}\|}{\|\mathbf{n}\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle\|} = \frac{\|\mathbf{P} - \mathbf{S}\|}{-\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle} \tag{3}$$

Which, using $\|\mathbf{D}_S - \mathbf{S}\| = -\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle$ means that

$$\alpha = \frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle}. \tag{4}$$

We thus have the following equation that links the detector co-ordinate location of a projected point to the location of this point:

$$d_1(\mathbf{P})\mathbf{u} + d_2(\mathbf{P})\mathbf{v} = (\mathbf{S} - \mathbf{D}_0) + \frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle}(\mathbf{P} - \mathbf{S}). \tag{5}$$

Taking inner products with $\mathbf{u}$ and $\mathbf{v}$ and using the fact that $\langle \mathbf{u}, \mathbf{v} \rangle = 0$ and $\langle \mathbf{u}, \mathbf{n} \rangle = 0$, we thus get the following equations:

$$\mathbf{d}(\mathbf{P}) = \begin{bmatrix} d_1(\mathbf{P}) \\ d_2(\mathbf{P}) \end{bmatrix} = \begin{bmatrix} \langle \mathbf{S} - \mathbf{D}_0, \mathbf{u} \rangle + \frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle} \langle \mathbf{P} - \mathbf{S}, \mathbf{u} \rangle \\ \langle \mathbf{S} - \mathbf{D}_0, \mathbf{v} \rangle + \frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n} \rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n} \rangle} \langle \mathbf{P} - \mathbf{S}, \mathbf{v} \rangle \end{bmatrix} \tag{6}$$

This links the location of any point $\mathbf{P}$ to the location of the projection of this point in terms of the detector co-ordinates $\mathbf{u}$ and $\mathbf{v}$. This mapping is a function of the unknown geometric parameters $\mathbf{S}$, $\mathbf{D}_0$, $\mathbf{u}$ and $\mathbf{v}$.

B. Mapping linear and rotational motion to points in 3D space

We assume that the linear motion defines a linear co-ordinate system with three orthonormal axes in 3D. We thus assume we know the locations of the points $\mathbf{P}$ for the different linear motions.

For the rotational motion, we know that one of the points $\mathbf{P}$ is at the centre of the co-odinate system. However we do not know the spatial position of the other points, all we know is that they are rotated versions of points in 3D space (some are rotated versions of the zero point. Importantly, the rotation is around an unknown axis of rotation which needs to be estimated.

For sphere i, let $\mathbf{P}_i^k$ be the actual location of the centre of the sphere in 3D space after a rotation by some angle θ_k. Each sphere is assume to be rotated around the same rotational axis. We here use the following way to describe the location of points $\mathbf{P}_i^k$ in terms of this axis

$$\mathbf{P}_i^k = \mathbf{p}_0 + \alpha_i \mathbf{r} + a_i \mathbf{q}_i^k. \tag{7}$$

Here $\mathbf{r}$ is the axis of rotation. $\mathbf{p}_0$ specifies the location of this axis in relation to our co-ordinate system. Vectors $\mathbf{q}_i^k(\theta_i^k)$ are unit vectors that are orthogonal to the rotation axis which are rotated around this axis by θ^k relative to some initial vector $\mathbf{q}_i^k(0)$. As this model is over-specified, we assume that $\mathbf{r}$ and $\mathbf{q}_i^k(\theta_i^k)$ are unit vectors. Furthermore, we fix one of the co-ordinates of $\mathbf{p}_0$. We here chose the direction which is roughly parallel to the expected axis of rotation. The parameter a_i thus specifies the distance of the sphere from the axis of rotation. Similarly, the parameters α_i specify the height of the plane of rotation along $\mathbf{r}$ above point $\mathbf{p}_0$.

To define the vectors $\mathbf{q}_i^k$, we define

$$\mathbf{q}_o(\mathbf{r}) = \frac{\begin{bmatrix}1\\0\\0\end{bmatrix} - \mathbf{r}/\|\mathbf{r}\|^2 \begin{bmatrix}1\\0\\0\end{bmatrix}^T \mathbf{r}}{\left\| \begin{bmatrix}1\\0\\0\end{bmatrix} - \mathbf{r}/\|\mathbf{r}\|^2 \begin{bmatrix}1\\0\\0\end{bmatrix}^T \mathbf{r} \right\|} \tag{8}$$

which is a unit norm vector orthogonal to $\mathbf{r}$ (We here assume that the rotation axis is not parallel to the x axis of our co-ordinate system). Let $R_\theta(\mathbf{P})$ be the affine transformation that describes the rotation of a point $\mathbf{P}$ around the axis defined by $\mathbf{p}_0$ and $\mathbf{r}$. The vectors $\mathbf{q}_i^k$ are then

$$\mathbf{q}_i^k = R_{\theta_i^0+\theta^k}(\mathbf{q}_o), \tag{9}$$

where θ^k is the rotation angle used when taking the k^{th} projection and where θ_0^i is an initial rotation that depends on the location of the sphere relative to the rotation axis.

With this model, we thus have to estimate the following parameters: $\mathbf{p}_0$, $\mathbf{r}$, α_i, a_i and θ_i^0. The first two vectors specify the location and orientation of the rotation axis, whilst the other parameters define the location of the spheres relative to this axis.

V. Estimation of parameters from linear motions

Assume we have conducted a calibration experiment in which a sphere is moved to different points along three orthogonal axes. Assume we know the distances between the sphere centres after each motion. We can then define a lab co-ordinate system, that is, we have a co-ordinate system in which the centre locations are known, that is, we know $\mathbf{P}_1, \mathbf{P}_2, \cdots, \mathbf{P}_N$. Assume all spheres have been projected onto the detector to points $\mathbf{d}(\mathbf{P}_1), \mathbf{d}(\mathbf{P}_2), \cdots, \mathbf{d}(\mathbf{P}_N)$ which have been estimated using, for example, our

circle fitting procedures. Assuming the detector rows and columns are orthogonal, we define $\mathbf{D}_0$ to be the bottom left (as seen from the source) pixel location in the lab co-ordinate system. We define $\mathbf{u}$ to point along the pixel rows and $\mathbf{v}$ along pixel columns. The estimated centres of the sphere projections are given in detector co-ordinates, that is, for sphere i, we have a vector $\mathbf{d}_i$.

To estimate the geometry of the cone-beam projection, we define the following cost function:

$$\frac{1}{N}\sum_i \langle \mathbf{d}_i - \mathbf{d}(\mathbf{P}_i), \mathbf{d}_i - \mathbf{d}(\mathbf{P}_i)\rangle,$$

where $\mathbf{d}_i$ are the locations of the sphere centres on the detector estimated with circle fitting. We would like to minimise this cost with respect to the vectors $\mathbf{S}$, $\mathbf{D}_0$, $\mathbf{u}$ and $\mathbf{v}$ under the constraints that

$$\langle \mathbf{u}, \mathbf{u}\rangle = 1, \tag{10}$$

$$\langle \mathbf{u}, \mathbf{v}\rangle = 0, \tag{11}$$

$$\mathbf{n} = \mathbf{u} \times \mathbf{v} \tag{12}$$

and

$$\frac{\langle \mathbf{D}_0 - \mathbf{S}, \mathbf{n}\rangle}{\langle \mathbf{P} - \mathbf{S}, \mathbf{n}\rangle} > 0. \tag{13}$$

Different approaches to the optimisation of this cost function are possible. We could use a gradient descend method, however, exploratory experiments have shown that a carefully designed co-ordinate descend method with line search procedure was significantly faster and this is the approach we present here.

We initialise all unknown geometric parameters with estimates taken from our scanner's geometric setup. We have set up our system so that two of the axis of the manipulator run roughly parallel to the two axes of the detector. We also have mounted the sphere so that it is projected close to the centre of the detector. We thus initialise $\mathbf{u}$ to be in the direction of one of our manipulator axis, whilst $\mathbf{v}$ is in the direction of the other manipulator axis. We initialise the position of the source location to be zero in these two directions and make the same assumption for the Detector position $\overline{\mathbf{D}_0}$. Whilst we also have a rough estimate of the distance between the source and detector we found that we could initialise the source to object and source to detector distances to arbitrary values.

We know that point $\mathbf{P}_0 = [0, 0, 0]^T$ projects onto the detector at point $\mathbf{d}_0$ say. We thus constrained our optimisation so that this is always satisfied. To do this, we start optimisation by moving the detector location $\mathbf{D}_0$ so that the point $[0, 0, 0]^T$ projects onto $\mathbf{d}_0$. We now re-parameterise the problem and define

point $\overline{\mathbf{D}_0}$ to be the point $\overline{\mathbf{D}_0} = \mathbf{D}_0 + \mathbf{d}_0[1]\mathbf{u} + \mathbf{d}_0[2]\mathbf{v}$. Let $\mathbf{d}_i$ be the adjusted coordinates of the measured projections such that

$$\mathbf{D}_0 + \mathbf{d}_i[1]\mathbf{u} + \mathbf{d}_i[2]\mathbf{v} = \overline{\mathbf{D}_0} + \overline{\mathbf{d}}_i[1]\mathbf{u} + \overline{\mathbf{d}}_i[2]\mathbf{v}. \tag{14}$$

With this parameterisation, we can then optimise system properties one at a time whilst keeping the projection of $\mathbf{P}_0$ fixed on point $\overline{\mathbf{D}_0}$. $\mathbf{S}$ and $\overline{\mathbf{D}_0}$ are defined through the vectors $\mathbf{S}$ and $\overline{\mathbf{D}_0} - \mathbf{S}$ which are parallel as the points $\mathbf{S}$, $\overline{\mathbf{D}_0}$ and $[0,0,0]^T$ lie on a line by definition. If we move $\mathbf{S}$ along this line closer to $\overline{\mathbf{D}_0}$, we in effect change the magnification of the projections. If we move both $\mathbf{S}$ and $\overline{\mathbf{D}_0}$ so that the ratio between the length of $\overline{(\mathbf{SP})}$ and the length of $\overline{\mathbf{PD}}$ stays the same, then the magnification stays the same but we change the effective cone angle used in the projection. We can also rotate $\mathbf{S}$ and $\overline{\mathbf{D}_0}$ around the co-ordinate system centre (this gives us two parameters to optimise) and we can rotate $\mathbf{u}$ and $\mathbf{v}$ around point $\overline{\mathbf{D}_0}$ (giving two more parameters.

We then iteratively use bracketing line search procedures to optimise

1) Magnification
2) Cone angle
3) Rotation of $\mathbf{S}$, $\overline{\mathbf{D}_0}$ and $\mathbf{v}$ around an axis parallel to $\mathbf{u}$ going through the coordinate centre.
4) Rotation of $\mathbf{S}$, $\overline{\mathbf{D}_0}$ and $\mathbf{u}$ around an axis parallel to $\mathbf{v}$ going through the coordinate centre.
5) Rotation of $\mathbf{u}$ and $\mathbf{v}$ around $\overline{\mathbf{D}_0}$ keeping $\mathbf{u}$ and $\mathbf{v}$ in the detector plane.
6) Rotation of $\mathbf{v}$ around $\overline{\mathbf{D}_0}$ and axis $\mathbf{u}$.

In the first iteration, we only optimise based on the projections for which the points $\mathbf{P}$ where in one plane, i.e. $\mathbf{P}_i = [\cdot, \cdot, 0]^T$ which provided a rough initial estimate of all parameters apart from the source to detector distance.

VI. Optimisation of rotation axis

Once a co-ordinate system has been established together with the source and detector locations as well as the detector co-ordinate axis, the use of rotations in many tomographic trajectories requires the estimation of a rotation axis.

We also use a bracketing line search to estimate the rotation axis location and orientation. We again parameterise the search for the parameters so that point $\mathbf{P}_0$ projects onto $\overline{\mathbf{D}_0}$. We get an initial estimate of the orientation and location of the axis by optimising only the projections from the ball that at θ_0^0 is at point $\mathbf{P}_0$. We do this assuming we know the direction of the rotation axis, which we set up to be parallel to one of the axis of the detector. Fine tuning of this axis will be performed later. Assuming we

know $\mathbf{r}$, we first estimate α_0 by projecting the initial estimate of the rotation axis onto the detector. We then project all points $\mathbf{d}(\mathbf{P}_i)$ on the projector onto this projection of the rotation axis. The mean hight of these projections acts as an initial estimate of α_0. (We do the same for the other spheres).

We then fix the position of the rotation axis by using the fact that one sphere is at point $[0, 0, 0]$ at a rotation of 0. This constraint defines the following equation

$$\mathbf{0} = \mathbf{p}_0 + \alpha_1 \mathbf{r} + a_1 \mathbf{q}_i^0, \tag{15}$$

which allows us to compute a_1, $\mathbf{p}_0[1]$ and $\mathbf{p}_0[3]$ for any α_1 a θ_0. Remember that θ_0 defines the vector $\mathbf{q}_i^0$ in relation to $\mathbf{r}$. In fact, we use the above equality every time we update θ_0 and $\mathbf{q}_i^0$ to update the location of the rotation axis.

We then estimate the vector $\mathbf{q}_0$ for this sphere which is parameterised through two parameters, an initial rotation around $\mathbf{r}$ and its length a_0. We perform a two stage line search, for each proposed value of a_0, we run a bracketed line search to find the optimal initial rotation. The cost is always evaluated such that for any combination of length and initial rotation, the vector $\mathbf{p}_0$ is shifted so that point $\mathbf{P}_0 = [0, 0, 0]^T$. Note that the initial rotation cost will be a circular cost function that is 2π periodic. The initial bracket for the line search is thus always made up of two points that are 2π apart plus a point between them that has a lower cost. This way we are able to explore all possible values of rotation. Also, sometimes, the line search for the length of the vector returns a negative value at which point we change the sign of this length together with the initial rotation angle to which we add π. Initially this optimisation is done looking only at the error between the locations of the projected sphere centres for this one sphere. Optimisation is stopped once this distance is at a minimum and does not change significantly any more.

After this initial optimisation, we have an estimate of $\mathbf{p}_o$ which we then use to do the same optimisation for the other spheres. The only difference here is that we also do an additional bracketed line search for the values of α_i, that is, for each value of a_i, we first optimise the initial rotation of the vector and then optimise α_i, both using bracketed line searches.

This gives us initial estimates of all parameters. Fine tuning for these parameters is performed using a cost function that takes account of the distance between each point $\mathbf{d}(\mathbf{P}_i^k)$ and its estimated position for a given set of parameters. We then use bracketed line search to optimise a_i, α_i ($i \neq 0$) and initial rotation θ_i^0. Note that when estimating a_0 and θ_0 we are automatically updating p_0 as well. The order of this optimisation is randomised from iteration to iteration. Once the cost function has converged with sufficient precision, we add bracketed line searches that rotate the rotation axis $\mathbf{r}$ around $\mathbf{q}_0(\mathbf{r})$ and $\mathbf{p}_o(\mathbf{r}) = \mathbf{r} \times \mathbf{q}_0(\mathbf{r})$.

VII. Results

To analyse the performance of the method, we performed several calibration scans using both manipulators. Several aspects of the method are of interest and we here explore the convergence, how well the method works over a range of scans, how repeatable the measurements are and, finally, how much the calibration improves image quality. The main performance measure we are using will be in terms of the average error between the measured locations of the projections of the spheres and the estimated projections for the assumed geometry. Detector pixels were 0.2mm wide and high and we express the error here also in mm. Thus, errors below 0.2mm are smaller than the pixel size.

A. Convergence

1) Calibrating linear axis: To see how well the method performs and converges, we show the difference between measured and estimated centre locations on the detector in figure 3. The first panel shows the difference after an initial estimation of the magnification and cone angle whilst the second panel shows the final result after 35 iterations. Figure 4 shows the speed with which the method converges. It only takes about 5 iterations to get a sub-pixel average error.

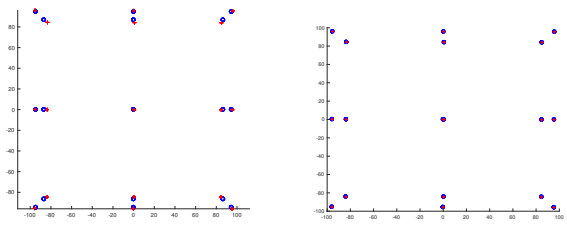

(a) After initial step (b) After 35 iterations

Fig. 3: Measured sphere centres on detector (red +) and estimated centres for optimised geometry (blue circle). Average distance between observed and estimated centres here is 0.092mm, which is half the width of a pixel.

2) Calibrating rotational axis: Results for the calibration of the rotation axis are similar. To show the difference between the initial estimates of the projected sphere centres and achieved accuracy after convergence, we show both the estimated (coloured circles) and measured (red crosses) centres in figure 5.

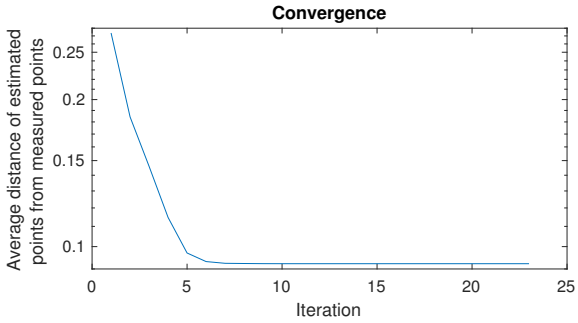

Fig. 4: Convergence speed of the method plotted on a logarithmic y-axis. An error of less than 0.1mm is achieved after only 4 iterations

.

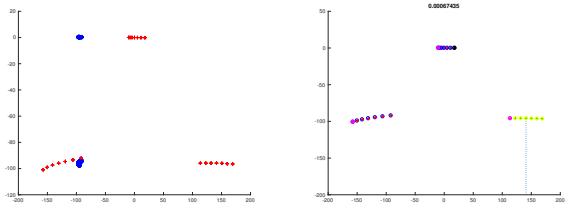

(a) After initial step (b) After convergence

Fig. 5: Measured sphere centres on detector (red +) and estimated centres for optimised geometry (blue circle) for the three spheres used to estimate the location of the rotation axis. Average distance between observed and estimated centres here is 0.033mm, which is a sixth of the width of a pixel.

B. Performance for different scans

The average accuracy in mm achieved in different scans is given in table I, which shows sub-pixel accuracy for all but one scan. Accuracy for the robot was in general worse than that for the Hexapod manipulator. Note that scan 17 was performed with a magnification of 5, whilst all other scans used a magnification of 10. At the lower magnification, positional errors are magnified less which explains the error for scan 17 which is less than the error observed in the other two robot arm scans.

Accuracy of the limited angle calibrations, which measure the average accuracy of the sphere centres from the rotational calibration scan is significantly smaller than the errors in the raster calibration scans. This is likely to be due to the more stable rotation axis in our system. The rotation stage is mounted on top of the hexapod. Thus, when using this axis, all other axis do not move. For the raster scan, we however used the linear axis together with the hexapod to perform the required motion leading potentially

to additional compound inaccuracies. We also assumed that our linear axes are linear and orthogonal and errors here will lead to further uncertainty in the results.

C. Repeatability

For some of the raster scans, we performed a calibration scan before and after the full scan. This allows us to compare repeatability of the calibration process and the repeatability of the manipulator. Table II shows the average repeatability in the estimation of the sphere centres. We here used two approaches to compare repeatability. For scans 1 to 5, we performed two calibration scans, where the sphere was attached to the sample and a 3 by 3 grid of points was imaged in a single plane. We then compared the repeatability of this grid of projected points. As the spheres were removed for the full scan, the spheres were not attached to the same points for both calibration. We thus shifted the second set of points to have

Scan	manipulator	motion	accuracy (mm)
1	hexapod	raster	0.0923
2	hexapod	raster	0.0981
3	hexapod	raster	0.0424
4	hexapod	raster	0.0627
5	hexapod	raster	0.0628
6	hexapod	raster	0.0639
7	hexapod	raster	0.0537
8	hexapod	raster	0.1187
9	robot	raster	0.2831
10	hexapod	raster	0.1433
11	hexapod	raster	0.1079
12	hexapod	raster	0.1324
13	hexapod	raster	0.1139
14	hexapod	raster	0.1430
15	hexapod	raster	0.1497
16	robot	raster	0.1736
17	robot	raster	0.0816
18	hexapod	lim. angl.	0.0316
19	hexapod	lim. angl.	0.0344

TABLE I: Average accuracy after optimisation for a range of scans with different manipulators and for different scan trajectories.

the same mean to the mean location of the first set of points. We then compute the difference between the two measurements and calculate a mean derivation as well as a standard deviation in this error. However, as the spheres were mounted in different locations, this introduced some additional errors due to the cone angle of the beam. For the last two scans (scans 6 and 7) we performed two full calibration scans, used these to estimate source and detector location/orientation. Once we have the two sets of estimated source and detector locations, we then used these to simulate projected points. This ensured that the same locations were used to compute the two sets of observations. Both results show that there was an error of between one and two pixels in the location of the estimated sphere centres.

This is thus the expected resolution limit for any reconstruction, as the uncertainties in the estimated source and detector locations will mean that we will have errors in the specification of our geometry of the order of a pixel.

It is interesting to observe that there do not seem to be significant differences here in the errors between the robot arm and the hexapod. Also of note is the fact that the error in the repeatability seems to be large than the error in the calibration.

To investigate the influence that these errors can have on the estimation of the location of the source and detector, we did a simulation study where we defined a geometry similar to the geometry in the experiments used above with a magnification of about 10. We then calculated the exact locations of a grid of points projected onto he detector. We added different amounts of gaussian noise to these points. The errors in the location of the source and the detector are shown in table III. These results are the standard derivation in the error over several repeated experiments. These errors are linear in the amount of noise added to the original locations. The error in the estimation of the source location has a standard

Scan	manipulator	accuracy (mm)
1	robot	0.2658
2	robot	0.4388
3	hexapod	0.2896
4	robot	0.3150
5	robot	0.1332
6	hexapod	0.2207
7	hexapod	0.4333

TABLE II: Error estimates in the location of the sphere centres when repeating the calibration scans.

deviation that is roughly ten times the error in the projected grid of points, whilst the error in the detector location is roughly of the same magnitude. The error in the orientation of the detector is here measured not in mm but in radians and measures the rotation of the estimated vector relative to the true vector.

D. Sphere centre detection performance

We also analysed the performance of the sphere fitting procedure. We here used canny edge detection and the parameters of the edge detection algorithm did not influence the performance significantly as long as we reliably estimated the edge of the sphere of interest without the inclusion of other edges, which we ensured with our masking and outlier detection approach.

The background on the projection image did however have an influence on the estimation of the sphere centre. As we attached the spheres to different samples, the variations in the sample's attenuation were observed to lead to edge estimates that were not always circular. Care should thus be taken when attaching the sphere to either use a part of the sample with relatively constant attenuation or to remove the sample and mount the sphere differently.

The applied Gaussian smoothing was observed to have some influence on the accuracy of the centre estimation process. To evaluate the influence of Gaussian smoothing, we looked at the error after optimisation of the geometry between the estimated sphere centres and the observed sphere centres. Results for two scans are shown in table IV for different amounts of smoothing. We see that a smoothing kernel with standard deviation of about 3mm (or 15 pixels) provides the most accurate results here.

E. Improvement in image quality

To compare the change in image quality with and without calibration, we did two reconstructions of a scan acquired from a carbon fibre composite panel using a raster scan trajectory. In the un-calibrated

Quantety	std=0.1	std=0.5	std=1	std=2
$\mathbf{S}$	1.1522	6.0109	12.5913	23.6311
$\mathbf{D}_0$	0.1336	0.6985	1.4583	2.7480
$\mathbf{u}$	0.0010	0.0055	0.0103	0.0182
$\mathbf{v}$	0.0009	0.0048	0.0080	0.0185

TABLE III: Standard deviation of the error in the different geometric vectors for different amounts of noise added to the estimated sphere centre locations.

reconstruction, we assumed that the motion in the raster scan was parallel to the detector and that the detector was orthogonal to the vector from the source to the detector centre. Note that the estimation of the source to detector distance is not important here as a change in this parameter simply re-scales the reconstructed volume in the source to detector distance. Results with and without calibration are shown in figure 6, where more detail is visible in the reconstruction with calibration. By measuring image sharpness by looking at the fraction of high frequency image components relative to low frequency components, we observed that calibration improved image quality for a range of images. Note that reconstruction of limited angle scans was not possible without an estimation of the location of the rotation axis.

F. Conclusions

The estimation of geometric properties for a generic tomographic sample manipulator system poses several challenges. Precise geometric information is required for good volumetric reconstruction with iterative algorithms. We have here explored an approach that uses one or several spheres that are imaged with the manipulator in different configurations. By choosing the configurations carefully, we could show that the projection images contain sufficient information for an accurate estimation of parameters such as source location, detector location and orientation and the alignment of linear and rotational manipulator axes. We explored this approach using two manipulator, a six axis robot arm as well as a hexapod system.

Scan						
	std	0.2	0.4	0.6	0.8	1
1		0.1207	0.1203	0.1207	0.1205	0.1187
2		0.0567	0.0557	0.0532	0.0540	0.0537
	std	1.2	1.4	1.6	1.8	2
1		0.1178	0.1172	0.1138	0.1094	0.1030
2		0.0542	0.0578	0.0529	0.0569	0.0577
	std	2.2	2.4	2.6	2.8	3
1		0.1022	0.1033	0.1008	0.1057	0.1038
2		0.0565	0.0559	0.0577	0.0548	0.0606
	std	3.2	3.4	3.6	3.8	4
1		0.1056	0.1083	0.1143	0.1207	0.1269
2		0.0568	0.0580	0.0592	0.0598	0.0603

TABLE IV: Average error (accuracy) after optimisation between observed and estimated sphere centres for different levels of Gaussian smoothing (std in mm).

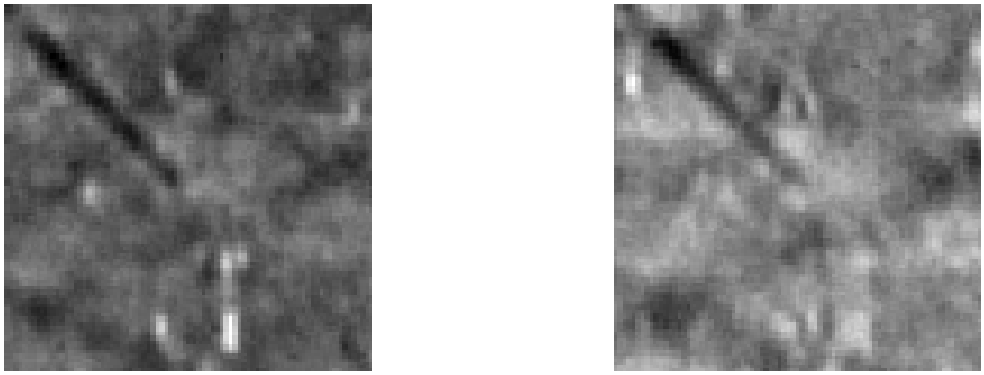

(a) Reconstruction with calibrated geometry. (b) Reconstruction with un-calibrated geometry.

Fig. 6: Detail in a slice of the reconstructed volume with (left) and without(right) use of exact calibration data.

Our approach allowed us to estimate geometric properties with sufficient accuracy to guarantee average geometric errors in the projected image to be below pixel size. This provided in plane resolutions in our laminography scans performed with ten times magnification, that were well below 100 micrometers.

REFERENCES

[1] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *Journal of the Optical Society of America A*, vol. 1, no. 6, pp. 612–619, 1984.

[2] N. O'Brian, R. Boardman, I. Sinclair, and T. Blumensath, "Recent advances in x-ray cone-beam computed laminography," *Journal of X-ray science and technology*, vol. 24, pp. 691–707, 2016.

[3] C. Wood, N. O'Brian, A. Denisov, and T. Blumensath, "Computed laminography of cfrp using an x-ray cone beam and robotic sample manipulator systems," *in preparation*, vol. 0, p. 0, 2018.

[4] H. G.T, L. A.K., and N. F, Eds., *Mathematical framework of cone beam 3D Reconstruction via the first derivative of the radon transform*, ser. Lecture Notes in Mathematics, vol. 1497. Berlin, Germany: Springer, 1991.

[5] M. Ferrucci, R. K. Leach, C. Giusca, S. Carmignato, and W. Dewulf, "Towards geometrical calibration of x-ray computed tomography systemsa review," *Measurement Science and Technology*, vol. 26, no. 9, 2015.

[6] D. Panetta, N. Belcari, A. D. Guerra, and S. Moehrs, "An optimization-based method for geometrical calibration in cone-beam ct without dedicated phantoms," *Physics in Medicine and Biology*, vol. 53, no. 14, pp. 3841–3861, 2008.

[7] Y. Kyriakou, R. Lapp, L. Hillebrand, D. Ertel, and W. Kalender, "Simultaneous misalignment correction for approximate circular cone-beam computed tomography," *Physics in Medicine and Biology*, vol. 53, no. 22, pp. 6267–6289, 2008.

[8] I. B. Tekaya, V. Kaftandjian, F. Buyens, S. Sevestre, and S. Legoupil, "Registration-based geometric calibration of industrial x-ray tomography system," *IEEE Transactions on Nuclear Science*, vol. 60, no. 5, 2013.

[9] F. Noo, R. Clackdoyle, C. Mennessier, T. White, and T. Roney, "Analytic method based on identification of ellipse parameters for scanner calibration in cone-beam tomography," *Physics in Medicine and Biology*, vol. 45, no. 11, pp. 3489–3508, 2000.

[10] C. Mennessier, R. Clackdoyle, and F. Noo, "Direct determination of geometric alignment parameters for cone-beam scanners," *Physics in Medicine and Biology*, vol. 54, no. 6, p. 16331660, 2010.

[11] P. Kalukin, B. Winn, Y. Wang, C. Jacobsen, Z. Levine, and J. Fu, "Calibration of high-resolution x-ray tomography with atomic force microscopy," *Journal of Research of the National Institute of Standards and Technology*, vol. 105, no. 6, pp. 867–874, 2000.

[12] R. Clackdoyle and C. Mennessier, "Centers and centroids of the cone-beam projection of a ball," *Physics in Medicine and Biology*, vol. 56, no. 23, pp. 7371–7391, 2011.