

Geometric calibration using bundle adjustment for cone-beam computed tomography devices

Alexander Ladikos and Wolfgang Wein
White Lion Technologies AG, Munich, Germany

ABSTRACT

In this paper we present a novel geometric calibration procedure for cone-beam computed tomography (CBCT) devices with arbitrary geometry using a calibration phantom containing steel beads. In contrast to typical calibration procedures the position of the beads does not have to be known precisely as it is also recovered during calibration. In addition, the arrangement of the beads inside the phantom is very flexible and does not have to follow hard constraints. The bead centers are extracted with subpixel precision from the projection images while taking the absorption properties of the calibration phantom into account. Based on the recovered center positions and phantom geometry, the projection geometry is computed for every projection image. This geometry can be arbitrary and does not have to lie on a specific path, e.g. a circle. This allows to calibrate devices with reproducible mechanical errors in the gantry movement. We present an evaluation of the point extraction and the calibration procedure on ground-truth data and show reconstruction results on a device calibrated using the proposed calibration method.

Keywords: Cone-Beam Computed Tomography, CBCT, Calibration, Bundle Adjustment

1. INTRODUCTION

Cone-beam computed tomography (CBCT) devices such as digital volume tomographs (DVT) are nowadays commonly used, for instance in maxillofacial and oral surgery. In order to obtain the best possible reconstruction results these devices need to be calibrated accurately, since an inaccurate calibration leads to artifacts and loss of detail in the reconstruction which can impair the diagnostic value. DVT manufacturers generally suggest to perform a calibration once a year, since the old calibration becomes less accurate over time due to vibrations and mechanical changes. Apart from commercial CBCT devices, there are also research-oriented setups which can have different geometries and are potentially often modified. For such devices a precision calibration phantom may not be available or too expensive. In this case it would be desirable to be able to use a low-cost hand-made phantom in order to calibrate the device. Since such a phantom will not reach micron accuracy on the placement of the beads, a method is required to optimize both for the geometry of the CBCT device and the geometry of the calibration phantom. To address this issue we propose a novel geometric calibration method for CBCT devices. Our calibration phantom design consists of an object (e.g. a cylinder) which contains several steel beads. These beads are extracted in the projection images and then used in a bundle adjustment procedure to optimize the geometry of the device. The position of the beads does not have to be known exactly since their exact position is also recovered during the calibration procedure. We use a general formulation for the projection geometry which does not make any assumptions on the movement of the gantry, e.g. that it moves on a circular path. Instead we compute the exact gantry configuration for every projection which allows us to also calibrate devices with repeatable mechanical aberrations.

For geometric calibration of CBCT devices we can roughly distinguish between two approaches: methods which estimate parameters specific to the cone-beam geometry¹⁻³ and methods which estimate a general projection matrix.⁴⁻⁶ One advantage of using a projection matrix is that the complex geometric relationships between the device parameters and the projection geometry need not be explicitly modeled. For instance, Navab *et al.*⁵ estimate the projection matrix for each projection image from correspondences between known 3D points in the calibration phantom and their image projections. We use a similar approach for initializing our calibration procedure. However, in contrast to previous work dealing with general gantry motions (i.e. not only circular ones), we do not assume that the calibration phantom is known to high accuracy as we are also optimizing the point positions.

Contact author: Alexander Ladikos (ladikos@wlt.ag)

2. METHOD

Our calibration procedure comprises the following steps: acquiring projection images of the calibration phantom, extracting the bead centers from the projection images, computing an initial calibration and performing the bundle adjustment procedure. Each of these steps is described in detail in the following sections. In addition, we discuss the parameterization of the projection geometry and the design guidelines for the calibration phantom.

2.1 Projection Geometry Parameterization

Instead of using a custom parameterization specific to a certain device geometry we chose to parameterize the geometry using the pinhole camera model.⁷ In this model the projection geometry is represented by a 3×4 projection matrix \mathbf{P} which can be decomposed as $\mathbf{P} = \mathbf{K} [\mathbf{R} \ \mathbf{t}]$ with \mathbf{K} a 3×3 matrix containing the camera intrinsic parameters and \mathbf{R} and \mathbf{t} the camera rotation and translation respectively. Since modern flat panel detectors have square pixels which are arranged in a grid, \mathbf{K} is of the form

$$\mathbf{K} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

where f is the focal length and $\mathbf{p} = [p_x \ p_y]^\top$ is the principal point. The geometric configuration of the device is then given by one set of global intrinsic parameters \mathbf{K} for the gantry and for each projection image I_j the spatial configuration (i.e. rotation \mathbf{R}_j and translation \mathbf{t}_j) of the gantry. Given the projection matrix \mathbf{P} a 3D-point $\mathbf{X} = [X \ Y \ Z \ 1]^\top$ given in homogeneous coordinates is projected to the homogeneous image point $\mathbf{x} = \mathbf{P}\mathbf{X} = [x \ y \ w]^\top$ which in inhomogeneous coordinates is $\hat{\mathbf{x}} = [\frac{x}{w} \ \frac{y}{w}]^\top$.

2.2 Phantom Design

One common choice for the design of the calibration phantom is to use a cylinder in which a number of steel beads are embedded. The cylinder itself should be made of a material with low X-ray absorption, for instance acrylic glass. The placement of the steel beads in the cylinder should be as balanced as possible, so that they cover the biggest possible detector area during a full revolution of the device. However, the projections of the steel beads should not intersect in the projection images since this will decrease the accuracy of the point extraction. Intersecting bead projections are in particular an issue for devices in which the tilt angle of the X-ray source is not zero. Our method requires at least 6 steel beads to work, but we recommend to use at least 8 and ideally as many as possible. However, care should be taken not to put them in any degenerate configuration, e.g. placing them all on a single line.⁷

2.3 Bead Center Extraction

After acquiring a series of images with the device, we first need to extract the bead centers from the projection images. This is done using the watershed segmentation algorithm⁸ which gives a first estimate of the blob positions and sizes. These initial blob regions are then dilated in order to obtain a larger blob which also contains some background. Next, the dilated regions are deleted from the input images and inpainted using an intensity profile approximated from the surrounding background. The result of the inpainting step is a projection image which only contains the cylinder, but not the beads anymore (see figure 1 (b)). Subtracting this image from the original input images removes the cylinder and only leaves the beads which are now much easier to extract precisely (see figure 1 (c)). The background subtraction step is also important to remove intensity gradients in the bead profile which are caused by the cylinder in which they are embedded. This correction for the background intensities significantly improves the extraction accuracy as we will show in the experimental section. Finally, the beads are extracted from the subtracted images using a sub-pixel precise contour tracing algorithm and a circle is fitted to each extracted contour. Figure 1 summarizes the steps during point extraction.

After extracting the bead centers in all images false positives are removed by grouping the centers into tracks based on their spatial movement and removing tracks which do not cover the whole sequence.




Figure 1. Steps during point extraction.

2.4 Calibration

The actual calibration comprises two steps: First an initial projection geometry is computed for every image, then the individual projection geometries are optimized with a set of common intrinsic parameters to obtain the final solution.

2.4.1 Initialization

For every projection image I_j an initial projection matrix $\tilde{\mathbf{P}}_j$ relating the 3D bead positions \mathbf{X}_i and the extracted bead centers \mathbf{x}_{ji} is estimated using the direct linear transform (DLT).⁷ It is assumed that the point correspondences are known. This information can for instance be obtained by ordering the extracted bead center positions in the images from top to bottom and assigning them the corresponding beads in the phantom.

The recovered $\tilde{\mathbf{P}}_j$ is then decomposed into $\tilde{\mathbf{K}}_j$, $\tilde{\mathbf{R}}_j$ and $\tilde{\mathbf{t}}_j$ as $\tilde{\mathbf{P}}_j = \tilde{\mathbf{K}}_j [\tilde{\mathbf{R}}_j \quad \tilde{\mathbf{t}}_j]$. Although ideally all $\tilde{\mathbf{K}}_j$ should be the same, they differ in practice due to noise in the measurements. In addition, in general $\tilde{\mathbf{K}}$ does not have the desired form given in equation 1. To bring $\tilde{\mathbf{K}}$ to the desired form we compute the mean of all $\tilde{\mathbf{K}}_j$ and use it as the initial estimate $\mathbf{K} = \frac{1}{M} \sum_{j=1}^M \tilde{\mathbf{K}}_j$ where M is the number of projection images. However, by changing \mathbf{K} the previously recovered \mathbf{R}_j and \mathbf{t}_j are not correct anymore since $\tilde{\mathbf{P}}_j \neq \mathbf{K} [\tilde{\mathbf{R}}_j \quad \tilde{\mathbf{t}}_j]$. We therefore have to perform a second DLT with \mathbf{K} fixed. This is achieved by first normalizing the bead centers according to $\mathbf{x}'_{ji} = \mathbf{K}^{-1} \mathbf{x}_{ji}$ thereby implicitly integrating the fixed intrinsic parameters \mathbf{K} into the measurements. Applying the DLT to the correspondences $\mathbf{x}' \leftrightarrow \mathbf{X}$ we would ideally expect the recovered projection matrix \mathbf{P}_j to be of the form $\mathbf{P}_j = [\mathbf{R}_j \quad \mathbf{t}_j]$ with \mathbf{R}_j a rotation matrix. This is not true in general since \mathbf{K} was only an approximation. We therefore replace \mathbf{R}_j by the closest orthogonal matrix as given by the Frobenius norm and modify \mathbf{t}_j accordingly. As the DLT is only a linear estimation method, we can further refine the recovered rotation and translation \mathbf{R}_j and \mathbf{t}_j using least-squares minimization over the reprojection error while holding \mathbf{K} fixed according to the cost function

$$\operatorname{argmin}_{\mathbf{R}_j, \mathbf{t}_j} \frac{1}{2} \sum_{i=1}^N \|h(\mathbf{K} [\mathbf{R}_j \quad \mathbf{t}_j] \mathbf{X}_i) - \mathbf{x}_{ji}\|^2 \quad (2)$$

where N is the number of steel beads in the calibration phantom, M is the number of projection images and $h(\mathbf{Y})$ transforms the homogeneous point \mathbf{Y} to inhomogeneous coordinates. The optimized rotation and translation matrices \mathbf{R}_j and \mathbf{t}_j for each frame together with the global intrinsic parameters \mathbf{K} are the input to the bundle adjustment step.

2.4.2 Bundle Adjustment

The initial estimates of \mathbf{K} , \mathbf{R}_j , \mathbf{t}_j and \mathbf{X}_i are refined using bundle adjustment.⁹ The cost function we are minimizing is given by

$$\operatorname{argmin}_{\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{X}_i} \frac{1}{2} \sum_{j=1}^M \sum_{i=1}^N \|h(\mathbf{K} [\mathbf{R}_j \quad \mathbf{t}_j] \mathbf{X}_i) - \mathbf{x}_{ji}\|^2 \quad (3)$$

We minimize this cost function using the Levenberg-Marquardt algorithm. In order to solve this optimization problem efficiently the problem structure has to be taken into account. It is clear that the Jacobian of our cost function will have a sparse structure, since for instance the derivative of \mathbf{R}_j with respect to any image other than I_j will be zero. This can be exploited by partitioning the optimization parameters into global intrinsic, extrinsic and point parameters which creates a block structure in the Jacobian.⁷ Using the structure of the Jacobian \mathbf{J} we can derive the structure of $\mathbf{A} = \mathbf{J}^T \mathbf{J}$ which is used in solving the normal equations. The normal equations can then be solved efficiently using a sparse LDL decomposition.¹⁰ In addition, the columns of \mathbf{A} should be reordered to have a minimum degree ordering.¹¹ This significantly increases the speed of the sparse LDL solver.


After the bundle adjustment step we obtain the optimized global intrinsic parameters \mathbf{K} , as well as the rotation \mathbf{R}_j and translation \mathbf{t}_j for every projection image I_j together with the optimized 3D positions \mathbf{X}_i of the steel beads in the phantom.

3. EVALUATION

We conducted experiments both for evaluating the accuracy of the point extraction and for evaluating the accuracy of the calibration procedure. These tests were conducted on synthetic data in order to have ground truth data. For a final assessment of the calibration quality we performed an experiment on real data.

In order to evaluate the accuracy of the point extraction we used a background image of our calibration phantom and simulated a steel bead by adding a 2D Gaussian blob to the image at subpixel precise locations along a predefined trajectory. We then ran the point extraction and compared the recovered center positions to the ground-truth. Figure 2(a) shows an evaluation of the background subtraction. It can be seen that the results with background subtraction are significantly better in areas of high background gradient while they are comparable in regions with no gradient. This is due to the fact that the background gradient also affects the intensity profile of the blob and therefore biases the extraction procedure. Figure 2(b) shows the accuracy when adding uniform noise to the synthetic images. It can be seen that even with significant amounts of noise the extracted points are very accurate. For real data we therefore expect an accuracy between 0.1 and 0.2 pixels given 5%-10% noise in the X-ray images.

To evaluate the accuracy of the calibration we first generated a phantom description and then added uniformly distributed noise to both the initial calibration phantom point coordinates and the measurements in the images. Our simulated setup had a detector of size 992x672 with 0.1 mm pixel size, a source-object distance of 340 mm, a source-detector distance of 530 mm and a 5 degree source tilt. We generated 360 projections over a range of 360 degrees. The calibration routine was then run on this input data, once with and once without the point optimization enabled. From the results we extracted the position of the X-ray source and the direction towards the detector center and compared them with the ground-truth values. The results we obtained are shown in table 3. It can be seen that the errors are quite small. With point optimization enabled the error is as expected almost independent of the noise in the calibration phantom points and only changes with the noise in the image



(a) Background subtraction

Noise	0%	2%	5%	10%
Mean	0.0563	0.0662	0.1284	0.2716
Std. Dev.	0.0326	0.0416	0.0896	0.1819

(b) Added Noise

Figure 2. Point extraction evaluation. (a) Top: phantom intensity profile. Bottom: Accuracy with background subtraction (blue) and without (red). (b) Error in pixels with noise added to the images. The percentage refers to the intensity range of the image.

<i>Without point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.48	0.95
0.5	1.90	1.58	2.25
1.0	3.62	3.62	3.61
2.0	5.92	6.16	10.88

(a) Error in source position (mm)

<i>Without point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.07	0.14
0.5	0.26	0.22	0.31
1.0	0.49	0.49	0.49
2.0	0.87	0.89	0.47

(b) Error in source angle (degrees)

<i>Without point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.59	1.16
0.5	4.19	4.62	4.70
1.0	8.29	7.96	8.68
2.0	17.34	15.99	16.79

(c) Reprojection error (pixels)

<i>With point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.48	0.95
0.5	4e-4	0.48	0.96
1.0	4e-4	0.48	0.96
2.0	4e-4	0.47	0.96

(d) Error in source position (mm)

<i>With point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.07	0.14
0.5	0.03	0.10	0.16
1.0	0.16	0.12	0.25
2.0	0.37	0.26	0.22

(e) Error in source angle (degrees)

<i>With point optimization</i>			
	0.0	1.0	2.0
0.0	0.00	0.58	1.16
0.5	3e-4	0.58	1.16
1.0	3e-4	0.58	1.16
2.0	1e-4	0.58	1.16

(f) Reprojection error (pixels)

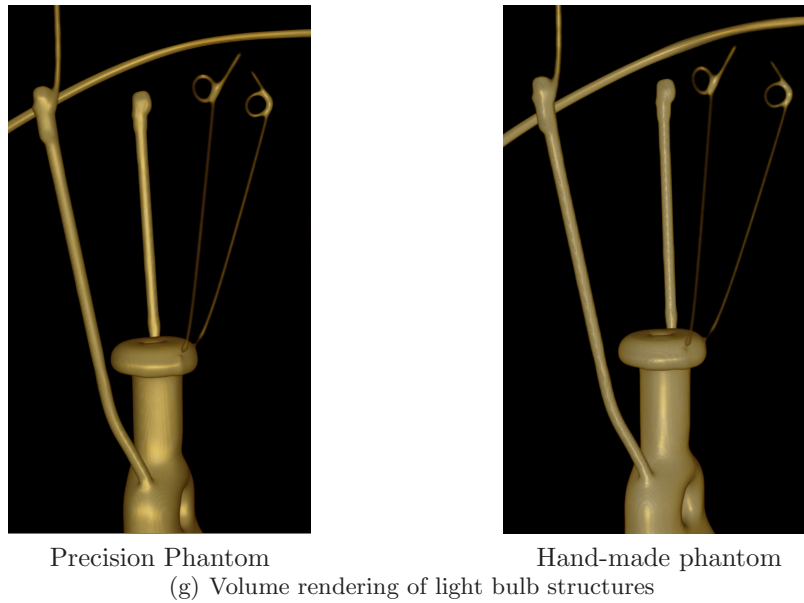


Figure 3. **Left:** Results of the accuracy evaluation. The row labels specify the noise added to the phantom point positions (in mm) and the column labels specify the noise added to the image measurements (in pixels). The first row of tables (a-c) shows the results without optimizing the phantom point positions, while the second row of tables (d-f) shows the results when also optimizing the phantom point positions. **Right:** Comparison of reconstruction results using precision-phantom (left) and hand-made phantom (right) on a volume rendering of a detail of a light bulb.

measurements. It is therefore advisable to optimize the phantom point positions. The only case in which this is undesirable is when the accuracy with which the position of the calibration phantom steel beads is known is at or above the spatial resolution of the imaging system. In this case the measurement noise would dominate the optimization of the point positions.

Finally, we ran a test on real data using both a precision-manufactured calibration phantom and a hand-made phantom (see figure 3(g)) which was built by manually placing little steel spheres on a cylinder made of cardboard. The manufactured phantom we used for our experiments contained 8 steel beads with a diameter of 1 mm placed on an acrylic cylinder with 30 mm diameter in a helical layout with an offset of 5 mm in height and 45 in angle between successive steel beads. The manufacturing accuracy of the phantom was 30 μm . The

accuracy of the hand-made phantom was not very high, with the bead positions known only in the millimeter range. We recorded CBCT sequences of both calibration phantoms and a light bulb. We then reconstructed the light bulb at a resolution of $256 \times 256 \times 400$ voxels with the parameters obtained from the precision and the hand made phantom and compared the reconstructions. Figure 3(e) and (f) each show a volume rendering of the fine wires inside the bulb. There is no visible difference. The calibration time (without point extraction) is in general around 1 second while the point extraction takes approximately 200 ms per image.

4. CONCLUSION

We presented a calibration method for CBCT devices with arbitrary geometry. Our method does not need a precision-manufactured phantom in order to calibrate the system since we also optimize the phantom point positions. Experimental results on ground-truth and real data show the accuracy and validity of our approach.

REFERENCES

- [1] Von Smekal, L., Kachelriess, M., Stepina, E., and Kalender, W. A., “Geometric misalignment and calibration in cone-beam tomography,” *Medical Physics* **31**(12), 3242–3266 (2004).
- [2] Cho, Y., Moseley, D. J., Siewerdsen, J. H., and Jaffray, D. A., “Accurate technique for complete geometric calibration of cone-beam computed tomography systems,” *Medical Physics* **32**(4), 968–983 (2005).
- [3] Mennessier, C. and Clackdoyle, R., “Automated geometric calibration and reconstruction in circular cone-beam tomography,” in [*2008 IEEE Nuclear Science Symposium*], (2008).
- [4] Rougee, A., Picard, C. L., Trouset, Y. L., and Ponchut, C., “Geometrical calibration for 3d x-ray imaging,” in [*Proc. SPIE 1897, 161*], (1993).
- [5] Navab, N., Bani-Hashemi, A., Nadar, M. S., Wiesent, K., Durlak, P., Brunner, T., Barth, K., and Graumann, R., “3d reconstruction from projection matrices in a c-arm based 3d-angiography system,” in [*MICCAI*], (1998).
- [6] Li, X., Da, Z., and Liu, B., “A generic geometric calibration method for tomographic imaging systems with flat-panel detectors—a detailed implementation guide,” *Medical Physics* **37**(7), 3844–3854 (2010).
- [7] Hartley, R. I. and Zisserman, A., [*Multiple View Geometry in Computer Vision*], Cambridge University Press, second ed. (2004).
- [8] Cousty, J., Bertrand, G., Najman, L., and Couprie, M., “Watershed cuts: Minimum spanning forests and the drop of water principle,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **31**, 1362–1374 (2009).
- [9] Triggs, B., McLauchlan, P., Hartley, R., and Fitzgibbon, A., “Bundle adjustment – a modern synthesis,” in [*Vision Algorithms: Theory and Practice*], *Lecture Notes in Computer Science* **1883**, 298–372 (2000).
- [10] Davis, T. A., “Algorithm 849: A concise sparse cholesky factorization package,” *ACM Trans. Math. Softw.* **31**, 587–591 (December 2005).
- [11] Davis, T. A., Gilbert, J. R., Larimore, S. I., and Ng, E. G., “A column approximate minimum degree ordering algorithm,” *ACM Trans. Math. Softw.* **30**, 353–376 (September 2004).