

Image-based kinematic calibration of a dual-robot X-ray computed tomography system via reprojection error optimization

Anton Weiss, Simon Zabler, Markus G. R. Sause

Angaben zur Veröffentlichung / Publication details:


Weiss, Anton, Simon Zabler, and Markus G. R. Sause. 2026. "Image-based kinematic calibration of a dual-robot X-ray computed tomography system via reprojection error optimization." *Results in Engineering* 30: 111176.
<https://doi.org/10.1016/j.rineng.2026.111176>.



ELSEVIER


Contents lists available at ScienceDirect

Results in Engineering

journal homepage: www.sciencedirect.com/journal/results-in-engineering

Research paper

Image-based kinematic calibration of a dual-robot X-ray computed tomography system via reprojection error optimization

Anton Weiss ^{a,*}, Simon Zabler ^a, Markus G.R. Sause ^b^a Deggendorf Institute of Technology, Deggendorf, 94469, Bavaria, Germany^b University of Augsburg, Augsburg, 86159, Bavaria, Germany

ARTICLE INFO

Keywords:

Robotic computed tomography
Kinematic calibration
Image-based optimization
Reprojection error
Direct linear transformation (DLT)
FlyBy acquisition

ABSTRACT

This paper presents an image-based calibration method for a dual-robot X-ray computed tomography system with an additional rotary table that requires no external measuring instruments. Using radiographic projections of a spherical-marker object with known geometry, correction parameters for both robot kinematic chains and the turntable are jointly estimated by minimizing a multi-view reprojection error under smooth bounded limitation. The resulting corrected forward-kinematic model enables generation of a single, globally consistent projection geometry, eliminating the need for subsequent per-projection refinement. The method is quantitatively validated using modulation transfer function (MTF) measurements on reconstructed data and is compared to projection-wise geometry correction based on the Direct Linear Transformation (DLT). Results demonstrate the viability of the proposed approach for both classical Stop-and-Go and continuous FlyBy CT acquisitions.

1. Introduction

Robotic X-ray computed tomography (CT) systems represent an emerging class of industrial CT machines that offer substantially greater flexibility than conventional stationary configurations [1]. A twin-robotic system employs two robots: one carrying the X-ray source and the other the detector. Current implementations typically use six-degree-of-freedom (6-DOF) industrial robots or collaborative robots (cobots). In some configurations, the robots are mounted on linear rails to extend their range of motion, and turntables are integrated to enable additional specimen manipulation.

The use of robots enables flexible view selection for two-dimensional (2D) X-ray inspection and arbitrary trajectory generation for three-dimensional (3D) CT acquisition. CT reconstruction requires a large number of views, typically several thousand projections, to sufficiently resolve internal structures. However, the inherent absolute positioning accuracy of industrial manipulators is limited and generally insufficient for high-precision CT acquisition. As a result, geometric deviations in robot motion directly translate into projection inconsistencies, causing artifacts and loss of resolution in the reconstructed volume, as shown in Fig. 1.

These challenges motivate the need for in-situ, image-based calibration methods that exploit the CT system's own imaging data to estimate and correct geometric deviations. Such self-consistent

calibration eliminates dependence on external measurement instruments and ensures alignment consistency under actual operating conditions.

Conventional CT calibration techniques assume static and well-defined source and detector geometries. In robotic CT, however, the continuously changing projection geometry requires pose-wise correction. Existing approaches such as the Direct Linear Transform (DLT) [2,3] rely on repeated acquisitions for reference measurements. Because DLT requires a fully visible, known arrangement of markers, it mandates replacing the specimen with a calibration object and executing a dedicated calibration run. Crucially, due to the path-dependent repeatability of industrial manipulators (e.g., hysteresis and joint backlash), achieving the exact same physical pose requires executing the exact same dynamic path leading up to it. This means the entire acquisition trajectory must be replicated perfectly, which effectively doubles the total scan time. Furthermore, executing this trajectory repetition in a continuous FlyBy manner requires exact temporal and spatial synchronization between the calibration and measurement runs, which is technically difficult and still imposes the severe time penalty. While interpolating DLT matrices between sparse calibration poses is theoretically possible to save time, it is practically flawed; DLT can introduce large, non-physical offsets, particularly in the detector normal direction, that drastically diminish the quality of the interpolated poses and result in worse reconstructed image quality [4].

* Corresponding author.

E-mail addresses: anton.weiss@th-deg.de (A. Weiss), simon.zabler@th-deg.de (S. Zabler), markus.sause@uni-a.de (M.G.R. Sause).<https://doi.org/10.1016/j.rineng.2026.111176>

Received 25 March 2026; Received in revised form 11 May 2026; Accepted 20 May 2026

Available online 22 May 2026

2590-1230/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. (Left) Specimen with highlighted region indicating the relevant cross section. (Middle and right) Reconstructed slices without and with DLT-based projection geometry correction, respectively [2].

Alternatively, fiducial marker-based methods [5] are effective for specific setups but lose robustness when the specimen exhibits high attenuation and when markers are occluded. External metrology-based calibration using laser trackers or photogrammetry achieves high accuracy but introduces complexity, cost, and line-of-sight limitations, making in-situ application less practical [6–8].

While some recent methods address robots by modeling internal deformations [9,10], our study focuses on the localized geometric calibration of under-constrained systems.

This work makes two primary contributions to robot kinematic as well as CT geometric calibration:

- An image-based optimization framework that derives the kinematic calibration of both robots directly from radiographic projections, from which the corresponding projection geometries can be generated for a limited imaging volume.
- An experimental validation demonstrating consistent geometric refinement, evaluated using the MTF on reconstructed data, comparing the proposed kinematic calibration with per-image DLT correction, and demonstrated across multiple CT acquisition trajectories.

The remainder of this paper is organized as follows: Section II reviews related work on robotic and CT calibration. Section III introduces the underlying kinematic and projection formulations. Section IV details the proposed calibration method. Section V presents the experimental setup and results, followed by conclusions in Section VI.

2. Related work

Research related to this study can be grouped into three main areas: robot kinematic calibration, calibration of industrial CT systems, and the observability of image-based geometric estimation methods. Each area addresses a different aspect of geometric accuracy but faces similar challenges when applied to dual-robot CT configurations.

2.1. Robot kinematic calibration

Robot kinematic calibration improves the absolute positioning accuracy of manipulators by identifying and compensating model errors in the kinematic chain. Most approaches extend the nominal model with static or functional correction terms that are estimated from measurements of end-effector poses obtained using external measurement instruments such as laser trackers or optical motion-capture devices [11]. Although these methods achieve high absolute accuracy, their reliance on external instrumentation and line-of-sight visibility limits their applicability to enclosed or multi-robot systems where direct observation of the end effector is infeasible. The topic of robot calibration has been extensively studied, as evidenced by several comprehensive review papers summarizing methodologies, error sources, and performance metrics reported in the literature [8,12].

2.2. Industrial CT system calibration

Calibration of industrial CT systems refines the projection geometry defined by the relative positions of the X-ray source, detector, and rotation axis. Deviations in these parameters introduce geometric distortions and artifacts in reconstructed volumes. Conventional techniques rely on reference objects with known geometries, often calibration phantoms, to estimate projection parameters through reprojection error minimization or analytical models [13]. While effective for static setups, these methods are difficult to extend to robotic or reconfigurable CT systems, where both the source and detector move along complex trajectories.

2.3. Observability and combined limitations

A fundamental limitation in image-based calibration of multi-body robotic CT systems is restricted parameter observability. Due to the coupled nature of the involved kinematic and geometric transformations, projection data alone constrain only the relative configuration of the system components. In the absence of external references or metrology, global translations and rotations that affect all components equally are not observable.

Consequently, image-based calibration methods recover an internally consistent system geometry rather than an absolute, world-referenced configuration [14]. This ambiguity is intrinsic to the problem formulation and does not affect reconstruction fidelity, provided that the relative geometry between the imaging components is accurately preserved.

3. Preliminaries

This section summarizes the theoretical foundations underlying the proposed image-based kinematic calibration framework. The main formulations of the forward kinematics, projection geometry, and calibration object are introduced to establish notation and provide the mathematical context required for the subsequent methodology. Implementation details and their application to experimental data are presented in Section 4.

In the proposed image-based approach, the conventional Euclidean position residuals used in robot calibration are replaced by the reprojection error, which quantifies the difference between measured and synthetically projected marker locations. The synthetic projections are generated from the known calibration body introduced in Section 3.4, whose spherical features are clearly visible in radiographic projections and whose internal geometry is precisely defined and easily modeled. The projection is computed using the geometry model defined in Section 3.3, which in turn depends on the forward kinematics described in Section 3.1. This formulation establishes the reprojection error as the principal quantity minimized during the optimization described in Section 4.2.

3.1. Forward kinematics

The forward kinematics describe the transformation from the robot base to the tool center point (TCP) as a function of the joint variables $q = [q_1, \dots, q_n]^T$ and a vector of correction parameters θ . The parameter vector $\theta = [\theta_1, \dots, \theta_n]^T$ collects the translational and rotational offsets of all link transitions. Each link transition is modeled as the product of a nominal transform N , a small correction $E(\theta)$, and the joint motion $M(q)$:

$$T(q, \theta) = \prod_{i=1}^n N_i^{i+1} E(\theta_i)^{i+1} M(q_i)^{i+1}, \quad (1)$$

The correction terms E represent small deviations between consecutive links and constitute the parameters to be estimated during calibration (see Section 4.2).



Fig. 2. The physical twin-robot CT system, featuring a bicycle to demonstrate system scale and a representative application.




Fig. 3. Digital twin of the twin-robotic CT system and an additional turntable.

3.2. Kinematic calibration formulation

Kinematic calibration seeks to identify small translational and rotational deviations in the robot model by minimizing the discrepancy between measured and predicted positions of a reference point rigidly attached to the robot flange or end effector.

The forward kinematic model $\mathbf{T}(q, \theta)$ defined in Section 3.1 (see Eq. (1)) provides the pose of this reference point as a function of the joint configuration q and the correction parameters θ . Its translational component $p(q, \theta)$ is compared with the corresponding measured TCP position p^{meas} .

The calibration problem is formulated as a nonlinear least-squares optimization:

$$\theta^* = \arg \min_{\theta} \sum_i \|p_i^{\text{meas}} - p_i(q_i, \theta)\|_2. \quad (2)$$

This expression corresponds to the conventional formulation used in external metrology-based robot calibration, where p_i^{meas} is typically obtained from laser-tracker or photogrammetric measurements. In the proposed approach, these measurements are replaced by image-derived marker coordinates, as detailed in Section 3.3.

3.3. Projection geometry

In X-ray projection geometry, the mapping from the 3D reference frame to the 2D detector plane is expressed by the projection matrix \mathbf{P} . It transforms a 3D point x_{3D} in homogeneous coordinates into its image-plane projection x_{2D} according to

$$\begin{bmatrix} x_{2D} \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} [d_h | d_v | d-f]^{-1} & [d_h | d_v | d-f]^{-1} [-f] \\ \hline & \end{bmatrix}}_{\mathbf{P}} \begin{bmatrix} x_{3D} \\ 1 \end{bmatrix}. \quad (3)$$

Here, d_h and d_v denote the horizontal and vertical unit vectors of the detector plane, d its center position, and f the X-ray focal spot. The projection matrix \mathbf{P} is therefore fully determined by the tool coordinate systems of the source and detector robots.




Fig. 4. Calibration body used for geometric evaluation. (Left) Radiograph showing the spherical arrangement of steel markers and the internal glass-bead packing. (Center) 3D-printing slicer rendering visualizing the internal structure. (Right) Magnified view of a steel ball-bearing projection, appearing as a circular feature.

3.4. Calibration object

The calibration object consists of κ spherical markers arranged in a known three-dimensional configuration. In our implementation, the markers are positioned along a helical structure comprising two turns with an overall diameter of approximately 150 mm and a height of approximately 200 mm, as shown in Fig. 4. This design provides uniform spatial coverage within a compact volume, ensuring observability across a wide range of projection angles. The rigid body thus offers a dense and spatially distributed set of reference points for estimating the projection geometry and optimizing the kinematic parameters.

The reference points are realized as radiopaque spherical markers made of high-attenuation materials such as steel or tungsten. In cone-beam X-ray imaging, these spheres project onto the detector as nearly circular regions; however, to maintain mathematical rigor, we account for the minor ellipticity introduced by oblique viewing angles. The centers of these elliptical projections can be reliably detected using image-based algorithms and correspond to the physical marker centers with subpixel accuracy, typically within one-tenth of a detector pixel.

The subsequent methodology (Section 4) applies the formulations above to reconstruct the marker coordinates from multiple projections and to minimize their reprojection error.

4. Methodology

In this section, the complete workflow for image-based kinematic calibration is described. The proposed procedure estimates the joint- and tool-level geometric errors that represent the deviation of both robots from their nominal kinematic models, as introduced in Section 3.1. To obtain the required observations, a set of radiographic projections is recorded in the joint-space region in which subsequent CT scans will be performed, using the calibration object described in Section 3.4. Based on the known 3D marker configuration and the projection geometry defined in Section 3.3, a reprojection error is formulated as the difference between the detected marker projections and their synthetic for-

ward projections computed through the robot kinematic chains. Using the forward kinematic model, we define a nonlinear optimization problem, see (2), conceptually similar to the formulation in Section 3.2, that minimizes this reprojection error jointly over all κ markers and all η projections acquired along a trajectory that covers the relevant portion of the joint space around the intended measurement pose. The individual components of this workflow are detailed in the following subsections.

4.1. Initial 3D localization of the calibration body

To initialize the calibration, the 3D positions of the calibration body's spherical markers are reconstructed directly from the acquired radiographic projections. This step provides an initial estimate of the calibration body pose that is consistent with the nominal system geometry and sufficiently accurate for the subsequent optimization.

4.1.1. Ray intersection modeling

Each detected ellipsoid center in projection j corresponds to the 2D projection of a spherical marker. The associated ray is constructed by connecting the focal spot with the back-projected detector coordinate of this ellipse center, representing the X-ray path that passed through the marker. Since rays from different views do not generally intersect in 3D space, the midpoint of the shortest connecting segment between two rays (k, l) is used as their intersection candidate \mathbf{x}_{kl} . Candidates obtained from shallow intersection angles are less reliable; therefore, all pairwise midpoints are fused into a weighted estimate of the marker position.

The weight

$$w_{kl} = \frac{\|\mathbf{u}_k \times \mathbf{u}_l\|_2}{\sigma_{\perp,k}^2 + \sigma_{\perp,l}^2}, \quad (4)$$

depends on the intersection angle through the cross-product magnitude and on the transverse positional uncertainty $\sigma_{\perp,*}$ of each ray. The final 3D marker estimate is computed as

$$\hat{\mathbf{x}} = \left(\sum_{i < j} w_{ij} \right)^{-1} \left(\sum_{i < j} w_{ij} \mathbf{x}_{ij} \right). \quad (5)$$

4.1.2. Rigid registration

The reconstructed 3D marker set is aligned to the known calibration-object geometry using a rigid transformation. The transformation is obtained from a standard orthogonal Procrustes solution, which provides a closed-form least-squares estimate of both rotation and translation. This step compensates for global drift in the initial ray-based reconstruction and yields the initial pose of the calibration body in the world coordinate frame.

4.1.3. 2D-3D correspondence

Correspondence between detected marker projections and the known 3D marker configuration is established by associating each predicted projection with the nearest detected ellipse center, subject to geometric consistency constraints. Projections exhibiting overlap or insufficient separation are excluded to avoid ambiguous assignments. The optimization therefore operates only on reliable correspondences, and missing markers are handled through per-sample masking. In practice, stable convergence is achieved when at least ten markers are visible per projection and a minimum of twenty distinct projection views are available.

4.2. Optimization problem

The calibration is formulated as a nonlinear optimization problem that adjusts the joint- and tool-level correction parameters such that the synthetic projections of the calibration-object markers match their detected image locations while ensuring that the estimated corrections remain within physically plausible limits.

4.2.1. Calibration data acquisition

The calibration parameters are estimated from a dedicated sinusoidal circular trajectory (Fig. 5) that encloses the intended measurement volume. The calibration object is positioned at 500 mm radial offset from the turntables rotation axis and 1200 mm above the turntables surface. The trajectory comprises 240 equally spaced projections over 360°, with a vertical sinusoidal modulation of 200 mm amplitude and 20 periods along the full rotation. These projections are used exclusively for the kinematic calibration; the reconstructed datasets and image-quality evaluation are based on the trajectories described in Section 5.3.

4.2.2. Reprojection error

Let \mathbf{x}_i^c denote the 3D position of marker i in the world frame. For projection j , the forward projection through the projection matrix \mathbf{P}_j yields the predicted 2D marker center coordinates in the image plane

$$\mathbf{u}_{i,j}^{\text{est}} = \mathbf{P}_j \mathbf{x}_i^c. \quad (6)$$

The corresponding detected ellipse center in the acquired radiograph is denoted by $\mathbf{u}_{i,j}^{\text{obs}}$. The mean reprojection error over all visible markers and all projections is

$$E_{\text{reprj}} = \frac{1}{\eta \kappa} \sum_{j=1}^{\eta} \sum_{i=1}^{\kappa} \|\mathbf{u}_{i,j}^{\text{est}} - \mathbf{u}_{i,j}^{\text{obs}}\|_2. \quad (7)$$

This term directly measures the geometric consistency of the projection model with respect to the observed marker locations.

4.2.3. Joint error limitation

To avoid unrealistic corrections, each joint's translational and rotational offsets are constrained by a smooth bounded limiter. A smooth bounded penalty function $f(\cdot)$ [15] penalizes deviations approaching the prescribed lower and upper bounds (l, u) while remaining near zero within the admissible interval. This differentiable formulation avoids the discontinuities associated with hard box constraints and provides stable gradient signals during optimization.

The limiting term is added element-wise to each joint's translational and rotational correction vectors, \mathbf{e}_j^t and \mathbf{e}_j^r , yielding

$$E_{\text{reg}} = \sum_{j=1}^n \left[f(\mathbf{e}_j^t; l_j^t, u_j^t) + f(\mathbf{e}_j^r; l_j^r, u_j^r) \right]. \quad (8)$$

4.2.4. Joint error modeling

Joint corrections are parameterized as small rotational perturbations in axis-angle form and translational offsets. This representation provides a minimal and well-conditioned update while avoiding Euler-angle ordering effects and quaternion normalization constraints. To ensure physical plausibility, the estimated joint corrections are limited to remain within predefined bounds expressed in a joint-local parameterization. Translational corrections are applied directly.

4.2.5. Calibration objective

The full calibration loss combines the reprojection term and the smooth joint-offset regularizer:

$$L_c = E_{\text{reprj}} + E_{\text{reg}}. \quad (9)$$

The optimization seeks the set of correction parameters that minimize L_c , thereby enforcing geometric consistency across all projections while keeping the estimated joint errors within physically meaningful limits.

4.2.6. Optimization details and reproducibility

The loss L_c is minimized using gradient-based optimization with automatic differentiation (JAX [16]) and the Adam optimizer (Optax [17], default parameters $\beta_1=0.9$, $\beta_2=0.999$, $\epsilon=10^{-8}$). All parameters are optimized jointly in full batch without mini-batching, as the dataset size

permits evaluation of all projections per iteration. The learning rate follows a linear schedule with a peak value of 10^{-4} with proportional decay to zero over T iterations, terminating early when the loss stabilizes below a threshold of 10^{-6} . All gradients are computed in double precision, and a fixed PRNG seed ensures reproducibility. Reported results correspond to the final optimized parameters and the resulting reprojection error statistics.

4.2.7. Observability metric

To characterize local parameter identifiability under the acquired trajectory, we analyze the Jacobian of the stacked reprojection residual vector $\mathbf{r}(\theta)$ with respect to the parameter vector θ . The Jacobian $\mathbf{J} = \partial \mathbf{r} / \partial \theta$ is evaluated at the final solution and the Gauss-Newton information matrix is formed as $\mathbf{H} = \mathbf{J}^T \mathbf{J}$, comparable to [18].

Eigenvalue decomposition $\mathbf{H} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T$ provides a basis of parameter directions; small eigenvalues λ_i indicate weakly observable combinations. To relate the near-null space to individual parameters, we compute the per-parameter participation

$$p_j = \sum_{i \in \mathcal{N}} v_{j,i}^2, \quad \mathcal{N} = \{i \mid \lambda_i < \tau\}, \quad (10)$$

where $v_{j,i}$ denotes entry (j, i) of \mathbf{V} and τ is a fixed threshold.

4.3. Metrics

Calibration performance is assessed through the achievable spatial resolution of the reconstructed CT volume, quantified by the modulation transfer function (MTF) of the calibration object. The MTF characterizes the system's ability to reproduce fine structural details and is widely used as an objective resolution metric in industrial CT.

Two calibration strategies are compared across all experiments: (i) the DLT-based per-projection refinement, and (ii) the proposed kinematic calibration.

The nominal, uncalibrated geometry is not included in the MTF analysis, as the resulting reconstructions exhibit severe geometric blur that prevents a stable estimation of the edge-based MTF. The calibration object contains a dense glass-bead packing that provides sufficient contrast and structure for reliable MTF computation.

Three FlyBy acquisition trajectories are evaluated, reflecting the intended operational use of the system, together with a reference Stop-and-Go trajectory that represents the most stable acquisition mode for benchmarking. Although DLT-based calibration can be applied to the Stop-and-Go acquisition, it is not compatible with true FlyBy operation due to its requirement for repeated views of identical poses, further emphasizing the practical advantage of the proposed approach.

5. Results

This section presents the calibration results and evaluates their impact on CT acquisition.

5.1. Calibration-trajectory results

To quantify calibration convergence on the acquisition trajectory described in Section 4.2.1, reprojection residuals between measured and forward-projected marker locations were evaluated before and after optimization. For each marker i , the two-dimensional residual vector is defined as $\chi_i = [\Delta u_i, \Delta v_i]^T$, with Euclidean norm $r_i = \|\chi_i\|_2$. Table 1 summarizes the resulting residual statistics.

As shown in Table 1, the optimization substantially reduced the reprojection error. The RMSE decreased from 3.64 mm to 0.38 mm, and both residual components were centered near zero with markedly reduced spread, indicating improved geometric consistency for the calibration trajectory.

Table 1

Reprojection residual statistics on the calibration trajectory (before/after optimization).

	χ_i [mm]		Norm r_i [mm]	Components [mm]	
	RMSE	MAE	$\mu_r \pm \sigma_r$	$\mu_u \pm \sigma_u$	$\mu_v \pm \sigma_v$
Bef.	3.64	3.27	3.27 ± 1.02	-2.28 ± 0.79	2.59 ± 2.05
Aft.	0.38	0.24	0.00 ± 0.39	0.00 ± 0.32	0.00 ± 0.22

5.2. Experimental setup

The experiments were conducted on a twin robotic X-ray CT system comprising two 7-DOF manipulators (KUKA KR120 R2900 extra) mounted on linear rails (KUKA KL4000) and equipped with a turntable (WEISS CR1000C) for specimen motion, see Figs. 2 and 3. The manipulators specify a repeatability of 50 μm . The rotary table exhibits an axial and radial runout of approximately 30 μm , measured at a radius of 400 mm. A laser-tracker survey at installation reported a normally distributed TCP position error with a standard deviation of 1.24 mm and 1.39 mm for the source and detector robot TCPs respectively. Subsequent measurements indicated a dependence of the TCP position on the linear rail coordinate, suggesting higher-order kinematic effects. During the experiments, the rails remained fixed, and continuous laser tracking was not feasible due to line-of-sight limitations.

The imaging chain employed a micro focus X-ray source (Comet VarioFocus) operated at 225 kV and 700 μA , with an effective focal spot size of approximately 300 μm . Projections were recorded using a flat-panel detector (iRay NDT1717M3) featuring a 3072×3072 pixel matrix, a pixel pitch of 139 μm , and an integration time of 250 ms.

Reconstruction was performed using an in-house implementation of the simultaneous algebraic reconstruction technique (SART) [19] at an isotropic voxel size of 100 μm .

5.3. CT trajectory design and acquisition

Each acquisition comprised 3000 projections recorded at a focus-detector distance of 2000 mm, and a magnification factor of 2. Four trajectories were evaluated: three FlyBy trajectories representing typical robotic CT motion, and one Stop-and-Go trajectory serving as a stable reference. The sinusoidal trajectory and its transformation into the world frame are shown as an example in Fig. 6.

- **Circular Stop-and-Go (SNG-C):** A conventional circular acquisition with static imaging poses, used as a high-stability reference dataset.
- **Circular FlyBy (FB-C):** A repetition of the conventional circular trajectory executed in continuous FlyBy motion.
- **Oscillatory FlyBy (FB-O):** A circular trajectory with an additional smoothed square-wave vertical modulation (see Fig. 5, left).
- **Sinusoidal FlyBy (FB-S):** A circular trajectory with a sinusoidal vertical offset (see Fig. 5, right).

5.4. Runtime and convergence

All computations were performed on an NVIDIA RTX A4500 GPU with 16 GB VRAM and an Intel Core i7-12800H CPU. The calibration was optimized in full batch using JAX/Optax in double precision (cf. Section 4.2). For the reported calibration run, optimization terminated after approx. 2 min and achieved a throughput of 56 s^{-1} iterations with a preceding JIT compilation time of 68 s. Convergence was assessed using the RMSE-based reprojection loss and its component statistics (μ_u , μ_v). As shown in Fig. 7, transient oscillations occurred during the early iterations (0–500) due to the 150-step cosine warmup. At the peak learning rate, large step sizes caused the optimizer to temporarily overshoot the descent path, amplifying gradient competition between the primary

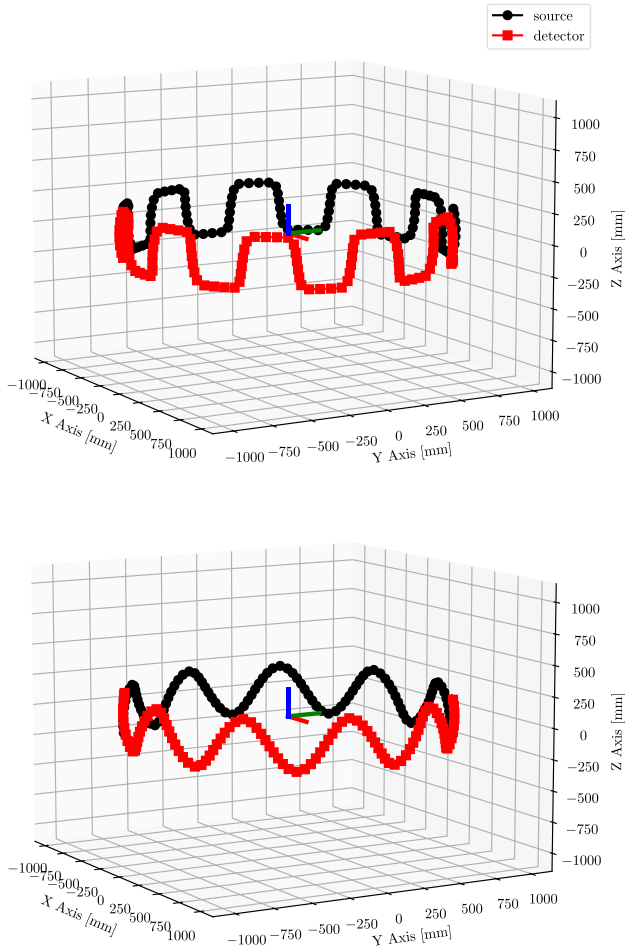


Fig. 5. Both plots illustrate a circular trajectory with an additional vertical offset. The upper plot shows an oscillatory offset generated by a smoothed square-wave signal, whereas the lower plot depicts a sinusoidal offset. For clarity, only 180° of the trajectory points are displayed to reduce visual overlap.

Table 2

Limit bounds for prismatic joints, revolute joints, and tool transformations.

	Prismatic	Revolute	Tool
Translational limit [μm]	1000.	100.	1000.
Rotational limit [mrad]	100.	10.	100.

loss and regularization terms. Once the learning rate decayed, these step sizes dampened, allowing the loss to decrease rapidly and stabilize below the stopping threshold of 10^{-6} .

5.5. Error limit compliance

The calibration converged successfully for all datasets, with estimated parameters remaining within predefined bounds. Revolute joint corrections were constrained to $\pm 100 \mu\text{m}$ translational and $\pm 2 \text{ mrad}$ rotational offsets, while prismatic joint parameters remained within $\pm 1 \text{ mm}$ translational and $\pm 20 \text{ mrad}$ rotational deviations.

Table 3

MTF values at 1 lp/mm and 2 lp/mm for on-axis and off-axis positions.

Dataset	On-Axis		Off-Axis	
	1 lp/mm	2 lp/mm	1 lp/mm	2 lp/mm
FB-O DLT	0.65	0.18	0.66	0.19
FB-O KC	0.46	0.05	0.49	0.06
FB-S DLT	0.64	0.17	0.69	0.23
FB-S KC	0.44	0.04	0.48	0.05
FB-C DLT	0.72	0.27	0.69	0.23
FB-C KC	0.69	0.23	0.66	0.19
SNG-C DLT	0.69	0.23	0.70	0.23
SNG-C KC	0.59	0.12	0.55	0.09

5.6. Modulation transfer function

A glass-bead phantom containing two bead populations with nominal diameters of 2 and 5 mm was used for quantitative evaluation. The two bead sizes provide spatial separation between bead classes and enable the same reference object to be used for both micro-focus and higher-resolution measurements. In the present analysis, only the larger beads were considered.

Radial intensity profiles across the bead edges were averaged to obtain the edge-spread function (ESF), which was differentiated to yield the line-spread function (LSF) [20]. A Gaussian model was fitted to the LSF to estimate its standard deviation σ , from which an analytical modulation transfer function (MTF) was derived as

$$\text{MTF}(f) = \exp[-2\pi^2(\sigma s)^2 f^2], \quad (11)$$

where s denotes the spatial scaling.

The MTF was evaluated at spatial frequencies of 1 lp/mm and 2 lp/mm for both on-axis and off-axis calibration object locations. Results are reported for continuous FlyBy (FB) and sequential Stop-and-Go (SNG) acquisition trajectories using direct linear transformation (DLT) and kinematic calibration (KC). The extracted modulation values, summarized in Table 3, provide an objective comparison of spatial resolution across the different trajectory and calibration configurations. Because DLT explicitly optimizes per-view projection consistency, edge-based MTF metrics inherently favor projection-wise correction approaches and therefore represent a conservative comparison for KC.

5.7. Observability results

Using the criterion in Section 4.2.7 with $\tau = 10^{-12}$, 72 eigenvalues fall into the near-null space, corresponding to an observable subspace dimension of $118 - 72 = 47$. Grouping parameters per joint (three axis perturbations, one rotation magnitude, three translations) shows that the rotation-magnitude component is consistently supported by the observable subspace, whereas axis-perturbation and translation components exhibit substantial participation in weak modes. This behavior is consistent with the expected redundancy of axis-angle updates for small rotations and limited translational excitation in the employed trajectory.

5.8. Laser tracker validation

To independently validate the proposed KC model, a reference point relative to the TCP was measured for both the source and detector robots using a laser tracker. By fitting the physical laser-tracked paths to the modeled KC paths, we evaluated the physical accuracy of the predicted trajectories.

Because the image-based calibration optimizes the three manipulators (source robot, detector robot, and turntable) simultaneously without relying on an external global reference frame, the calibrated system establishes an internally consistent geometric representation relative to the measurement volume. Consequently, aligning




Fig. 6. (Left) Sinusoidal trajectory prescribed in the object's reference frame (object omitted for clarity). (Right) Corresponding executed TCP trajectory measured during the robotic acquisition.




Fig. 7. Convergence of the full-batch Adam optimization for L_c Eq. (9). Shown are RMSE and MAE of $r_i = \|\chi_i\|_2$ and the residual statistics (μ_r, σ_r), indicating a near-zero mean residual ($\mu_r \approx 0$) at convergence.

the paths revealed a constant global translational offset of $(x, y, z) = (0.47 \text{ mm}, 0.83 \text{ mm}, 1.06 \text{ mm})$. This confirms that the calibration accurately corrects relative system motion but naturally deviates from absolute world coordinates due to the lack of rigid spatial anchoring, an expected behavior that does not impact CT reconstruction quality.

The calibration significantly improved the absolute positioning accuracy of the manipulators during the trajectory. By evaluating the standard deviation of the Euclidean distance between the physical laser-tracker measurements and the kinematic models, the positional error for the X-ray source robot was shown to reduce from 1.654 mm (nominal) to 0.462 mm (calibrated). Similarly, the detector robot exhibited an error reduction from 1.205 mm to 0.360 mm. These physical measurements validate the underlying assumptions of the KC model and demonstrate that the image-based geometric corrections successfully translate into highly stable physical trajectories.

5.9. Qualitative visual comparison

A qualitative visual comparison of reconstructed slices is shown in Fig. 8. While reconstructions obtained using kinematic calibration and DLT appear visually similar in representative slices, the quantitative MTF analysis in Table 3 provides a more sensitive assessment and reveals systematic resolution differences that are not readily discernible by visual inspection.

6. Discussion

The proposed image-based calibration substantially reduces reprojection error (Table 1) and yields a physically plausible set of bounded joint and tool corrections (Table 2). The resulting projection geometry is internally consistent and can be generated from a single kinematic parameter set. This is an advantage over per-projection DLT refinement that requires repeated poses and is therefore not compatible with FlyBy operation.

The reconstruction results show that kinematic calibration produces geometrically coherent volumes across all evaluated trajectories; however, the MTF comparison in Table 3 also indicates that KC does not uniformly match the DLT baseline. In particular, KC underperforms DLT for




Fig. 8. Central vertical slice through the calibration body (see Fig. 4) reconstructed using (l-r) DLT-based calibration, the proposed kinematic calibration, and the nominal robot geometry. The proposed calibration yields a reconstruction visually comparable to the DLT baseline, whereas the nominal geometry exhibits pronounced blur and loss of fine structure. Top: 20 mm scale bar; yellow box indicates the zoomed region. Bottom: zoomed view 10 mm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

FB-O, FB-S, and SNG-C, while FB-C shows closer agreement. This outcome is consistent with the different problem formulations: DLT optimizes an independent projection model per view and can absorb view-specific deviations, whereas KC estimates a single, globally consistent kinematic correction model that cannot represent per-view effects. Consequently, KC may yield slightly lower apparent resolution in trajectories where these effects are more pronounced, even though the overall geometry remains consistent. The estimated KC parameters should be interpreted as effective corrections that reproduce the observed projection geometry, not as a physically unique decomposition of individual joint errors.

The observability analysis (Section 5.7) provides an additional explanation. The calibration trajectory constrains only a low-dimensional subspace of the parameter vector, with weak modes primarily associated with axis-perturbation and translational components. Limited excitation of these components can lead to residual geometric errors that manifest as reduced MTF in some acquisitions, especially when the reconstruction trajectory deviates from the calibration trajectory or when the imaging volume probes regions of the workspace with different sensitivity.

From a practical standpoint, the method converges within minutes on a single GPU, making routine recalibration feasible. The results suggest a trade-off: KC enables projection-geometry generation for FlyBy acquisitions without repeated poses, at the cost of a modest reduction in MTF relative to per-projection DLT in some trajectories. Future work will address this gap by (i) designing calibration trajectories that improve excitation of weak parameter groups, (ii) extending the correc-

tion model to capture dominant non-kinematic effects, and (iii) quantifying run-to-run variability and trajectory-dependent generalization on a broader set of datasets.

CRediT authorship contribution statement

Anton Weiss: Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Simon Zabler:** Writing – review & editing; **Markus G.R. Sause:** Writing – review & editing.

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] X. Duan, J. Cai, Q. Ling, Y. Huang, H. Qi, Y. Chen, L. Zhou, Y. Xu, Knowledge-based self-calibration method of calibration phantom by and for accurate robot-based CT imaging systems, *Knowl. Based Syst.* 229 (2021) 107343. <https://doi.org/10.1016/j.knosys.2021.107343>
- [2] A. Weiss, S. Wittl, G. Herl, S. Zabler, Simulated and experimental evaluation of the accuracy of twin robotic CT systems, in: *12th Conference on Industrial Computed Tomography (iCT) 2023*, (2023).
- [3] R. Kang, G. Probst, M. Vlaeyen, G. Fei, P. Guerrero, H. Haitjema, P. Slaets, W. Dewulf, Geometric qualification for robot CT with flexible trajectories, in: *11th Conference on Industrial Computed Tomography (iCT) 2022*, (2022).
- [4] A. Weiss, S. Wittl, G. Herl, S. Zabler, A. Trauth, M.G.R. Sause, Safeguarding accuracy for CT imaging with industrial robots: efficient calibration methods for arbitrary trajectories: Anton Weiss, Simon Wittl, Gabriel Herl, Simon Zabler, Anna Trauth, Markus G. R. Sause, *e-J. Nondestruct. Test.* 30 (2) (2025). <https://doi.org/10.58286/30725>.
- [5] Y.Q. Ma, T. Reynolds, T. Ehtiati, C. Weiss, K. Hong, N. Theodore, G.J. Gang, J.W. Stayman, Fully automatic online geometric calibration for non-circular cone-beam CT orbits using fiducials with unknown placement, *Med. Phys.* (2024). <https://doi.org/10.1002/mp.17041>
- [6] M. Ulrich, C. Steger, F. Butsch, M. Liebe, Vision-guided robot calibration using photogrammetric methods, *ISPRS J. Photogramm. Remote Sens.* 218 (2024) 645–662. <https://www.sciencedirect.com/science/article/pii/S0924271624003757>. <https://doi.org/10.1016/j.isprsjsprs.2024.09.037>
- [7] H.M. Balanji, A.E. Turgut, L.T. Tunc, A novel vision-based calibration framework for industrial robotic manipulators, *Robot. Comput. Integr. Manuf.* 73 (2022) 102248. <https://doi.org/10.1016/j.rcim.2021.102248>
- [8] Z. Li, S. Li, X. Luo, An overview of calibration technology of industrial robots, *IEEE/CAA J. Autom. Sin.* 8 (1) (2021) 23–36. <https://doi.org/10.1109/JAS.2020.1003381>
- [9] H. Ye, J. Wu, T. Huang, Kinematic calibration of over-constrained robot with geometric error and internal deformation, *Mech. Mach. Theory* 185 (2023) 105345. <https://www.sciencedirect.com/science/article/pii/S0094114X23001167>. <https://doi.org/10.1016/j.mechmachtheory.2023.105345>
- [10] Y. Song, M. Liu, B. Lian, Y. Qi, Y. Wang, J. Wu, Q. Li, Industrial serial robot calibration considering geometric and deformation errors, *Robot. Comput. Integr. Manuf.* 76 (2022) 102328. <https://www.sciencedirect.com/science/article/pii/S0736584522000175>. <https://doi.org/10.1016/j.rcim.2022.102328>
- [11] A. Nubiola, I.A. Bonev, Absolute calibration of an abb irb 1600 robot using a laser tracker, *Robot. Comput. Integr. Manuf.* 29 (1) (2013) 236–245. <https://doi.org/10.1016/j.rcim.2012.06.004>
- [12] Z. Roth, B. Mooring, B. Ravani, An overview of robot calibration, *IEEE J. Robot. Autom.* 3 (5) (1987) 377–385. <https://doi.org/10.1109/JRA.1987.1087124>
- [13] A. Aichert, B. Bier, L. Rist, A.K. Maier, Projective invariants for geometric calibration in flat-panel computed tomography, in: *Proceedings of the CT-Meeting, 2018*, pp. 69–72. <https://www5.informatik.uni-erlangen.de/Forschung/Publikationen/2018/Aichert18-PIF.pdf>.
- [14] M. Ferrucci, R.K. Leach, C. Giusca, S. Carmignato, W. Dewulf, Towards geometrical calibration of X-ray computed tomography systems—a review, *Meas. Sci. Technol.* 26 (9) (2015) 092003. <https://doi.org/10.1088/0957-0233/26/9/092003>
- [15] Y. Wang, P. Praveena, D. Rakita, M. Gleicher, RangedIK: an optimization-based robot motion generation method for ranged-goal tasks (2023). [arXiv:2302.13935](https://arxiv.org/abs/2302.13935)
- [16] J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, JAX: composable transformations of Python + NumPy programs, (2018). <http://github.com/google/jax>.
- [17] D.P. Kingma, J. Ba, Adam: a method for stochastic optimization (2017). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- [18] Y. Sun, J.M. Hollerbach, Observability index selection for robot calibration, in: *2008 IEEE International Conference on Robotics and Automation, 2008*, pp. 831–836. <https://doi.org/10.1109/ROBOT.2008.4543308>
- [19] A. Andersen, Simultaneous algebraic reconstruction technique (sart): a superior implementation of the ART algorithm, *Ultrasound. Imaging* 6 (1) (1984) 81–94. [https://doi.org/10.1016/0161-7346\(84\)90008-7](https://doi.org/10.1016/0161-7346(84)90008-7)
- [20] A.D.A. Maidment, M. Albert, Conditioning data for calculation of the modulation transfer function, *Med. Phys.* 30 (2) (2003) 248–253. <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.1534111>. <https://doi.org/10.1118/1.1534111>