Projective Invariants for Geometric Calibration in Flat-Panel Computed Tomography

André Aichert, Bastian Bier, Leonhard Rist and Andreas K. Maier

Abstract—We present a new phantom design for geometric calibration of flat-panel detector CT systems. This work does not address a specific phantom but rather a toolbox for building very different phantoms for special trajectories and a general software to detect them and determine the geometric parameters of the scanner. Flexible robot trajectories and almost arbitrary distribution of metal beads in space are supported. A calibration algorithm is devised, which exploits a projectively invariant descriptor of four collinear points to solve the correspondence problem and determine the projection matrix for each projection. A proof-of-concept numerical study is presented with a randomly generated example phantom. We present a comparison to the frequently used PDS2 phantom.

Figure 1. Example 1 (left): The corss-ratio as a projectively invariant property $\mathrm{cr}\,(\mathbf{a}, \mathbf{b}; \mathbf{c}, \mathbf{d}) = \mathrm{cr}\,(\mathbf{a}', \mathbf{b}'; \mathbf{c}', \mathbf{d}')$. Example 2 (right): Projective scales. Three collinear points p_0, p_1 and the horizon p_∞ define a projective scale which allows us to take distance measurements directly in a photograph. Here, $\mathrm{cr}\,(\,p_\infty,\,p_0;\,p_1\,x\,) = 7$ is the distance of x to p_0 in units of the distance between p_0 and p_1.

I. INTRODUCTION

Robotic C-arms with a flat-panel detector are becoming increasingly flexible and support circle-line, saddle, rectangular and other trajectories. Using non-circular trajectories may be beneficial in some applications for more complete data, less redundancy or simply working with hardware and space constraints. However, the geometric calibration of a flat-panel CT system based on robot odometry alone is difficult, since small variations in joint angles have a large effect in terms of detector pixels. Combined approaches are feasible [8]. Image-based approaches are therefore preferable for reproducible trajectories and it remains common practice to calibrate scanner geometry using an X-ray compatible phantom prior to data acquisition. These calibration phantoms are typically comprised of radiopaque spherical markers, which are manufactured at high accuracy, although other forms exist [3]. In addition, not all phantoms work for general projections due to their shape [7]. See Mennessier et al. for some design considerations [6].

Most calibration algorithms determine only few parameters of the (perfectly elliptical, helical etc.) trajectory, instead of the projection for individual projections, e.g. [1], [9]. Other approaches, including this work, understand the trajectory more generally as a set of independent projections, in no particular order [7], [6], [3], which makes them generally applicable. In this context, calibration with a flat-panel detector is a standard computer vision problem, merely with X-ray images instead of visible light photography. It is well-known, that the determination of an 11 DOF projection matrix requires six images of known 3D points, no more than four of which may be coplanar. If we assume a decently manufactured X-ray detector has square pixels, only 9 DOF remain [2]. The reason why established methods for pose estimation, factorization and

auto-calibration are not directly applicable, is that they all rely on descriptors for matching corresponding points. Descriptors are usually based on texture, color or local gradients, none of which are salient for X-ray projections of a bead-phantom.

The contribution of this paper is a flexible and general phantom design comprised of metal beads for the determination of geometric parameters of an X-ray source and detector from its projections images. We suggest a descriptor based on four collinear metal beads, among them exactly one larger bead, in the following referred to as "pins". A phantom may be comprised of five to about thirty of such pins, arbitrarily arranged in space. Their detection is robust, since it is unambiguous and independent. We are able to solve the matching problem with as little as three correctly detected pins. This paper presents the underlying theory and a proof-of-concept based on numerical simulation, as well as comparison to the established PDS2 phantom [7]. We show that our phantom works equally well when projected from arbitrary angles and therefore supports not just circular, but also unusual (arbitrary) trajectories. Manufacturing accuracy does not limit calibration accuracy. We suggest cheap manufacturing using 3D-printing, since our solution of the matching problem allows standard Computer Vision algorithms to replace accurate manufacturing of phantom hardware with an accurate measurement process of both structure and motion.

II. ESTIMATION OF PROJECTIVE TRANSFORMATIONS

A. Objective

In the following, we outline the established gold-standard algorithm in computer vision for the estimation of the projection matrix from detected 2D points and a known 3D geometry. We understand geometric calibration in the sense that a single 3×4 projection matrix $\mathbf{P}$ must be estimated for each projection

image, using a pre-defined 3D phantom. The phantom consists of metal beads, which are easy to detect in 2D projection images (see Section IV-A). The problem is then to find the linear transformation between two unordered sets of points. The algorithms for both geometric calibration (i.e. estimation of a projection matrix) and registration (i.e. estimation of a linear 2D or 3D transformation) or fiducial-based rigid body tracking (i.e. rigid 3D pose) from points, are very similar and can be addressed by the algorithms in this work. For a concise mathematical notation, we will restrict ourselves to the calibration problem and present an overview of the process in this section.

Given detected 2D image points $\mathbf{x}_i \in \mathbb{P}^2$ in the projective plane (image) and known 3D points $\mathbf{y}_j \in \mathbb{P}^3$ in projective three-space (world), we seek to estimate a projection matrix $\mathbf{P}$, which minimizes the reprojection error

$$\underset{\mathbf{P}}{\text{argmin}} \ \frac{1}{|\mathcal{M}|} \sum_{(i,j) \in \mathcal{M}} \mathrm{d}\left(\mathbf{x}_i, \mathbf{P} \cdot \mathbf{y}_j\right), \qquad (1)$$

where $\mathcal{M} \subset \mathbb{N}_0^2$ is a set of index matches between the detected 2D points and the known 3D points and $\mathrm{d}(\cdot, \cdot)$ is the euclidian distance.

Given a set of at least 6 point matches one may obtain an algebraic estimate of $\mathbf{P}$ using the Direct Linear Transformation (DLT) [2, Ch. 7]. In order to be robust against outliers however, RANdom SAmpling Consensus (RANSAC) applies DLT to six randomly selected matches many times and determines the quality of the current estimate. A good measure of quality is the proportion of detected 2D points, for which a projected 3D point is close-by. The algorithm terminates when that proportion reaches a certain upper threshold, or, after a fixed number of iterations. Once RANSAC has produced a stable algebraic estimate of $\mathbf{P}$, bundle adjustment [2, Ch. 18] can be used to refine the solution by minimizing the non-linear geometric error according to Eq. 1.

B. Matching and Invariant Descriptors

The performance of RANSAC strongly depends on the frequency of outliers in the data, i.e. the chance that a random candidate match is incorrect. Unfortunately, there exist $\frac{N!}{(N-6)!}$ possibilities for a random point match with N points, so a better heuristic for establishing candidate matches than random guessing is mandatory. The common approach is to find a local descriptor of the points, typically based either on neighboring image data or local structure of the data. A good descriptor is (1) invariant to the observed transformations, so it will be possible to identify it across two images (2) local, so that occlusion or overlap in one part of the object does not affect the matching in other regions (3) salient, so that no two different points shall have a similar descriptor.

We can specifically design a calibration phantom, so prior knowledge reduces the number of possible point matches. We demand that any metal bead in the phantom is collinear with three other beads and that exactly one of these beads has a significantly larger radius. This enables us to determine a sequence of the points. We refer to a a set of four beads in this


Figure 2. Example pin configurations. The left pins have a positive descriptor (green background) while the right pins have negative descriptor (red backgr.) due to the position of the big bead. Note that the 11 rows at the bottom show pin configurations with identical descriptor to one of the top ten rows.

configuration as a "pin". Given four collinear points, the cross-ratio can be used as an invariant descriptor. By combining prior knowledge and a descriptor, we will show that the matching problem is effectively solved for a wide range of possible bead configurations in space. The main novelty of the paper lies in an elegant solution to the matching problem.

III. A Projective Invariant for Matching Points

A. The Cross-Ratio

Many familiar properties of an object change under projective transformation, notably, length, area, angles and ratio of lengths are not generally preserved under X-ray projection. However, some invariant properties prevail, notably incidence relations, such as two lines meet in a point and the cross-ratio. The cross-ratio bears its name since it is also the ratio of ratios of the distances between four points on a line. Find an illustrative example of its application to take distance measurements directly in photographs in Figure 1. We are interested in this quantity, because it remains constant under projective transformation, including translation, rotation, scaling and especially projective distortion. Using coordinates $a, \ldots, d$ on the line, the cross-ratio is defined

$$\lambda = \mathrm{cr}\left(a, b; c, d\right) \overset{\text{def}}{=} \frac{(a-c) \cdot (b-d)}{(a-d) \cdot (b-c)} \in \mathbb{R}. \qquad (2)$$

B. Practical Implementation in Arbitrary Dimension

In order to work with measured 2D (or, analogously 3D) points, we take the practical approach of projecting a set of approximately collinear points in arbitrary dimension $\mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{a} \in \mathbb{R}^n$ to the real line $\mathbb{R}$ with the scalar product and use Equation 2 to compute the cross-ratio. W.l.o.g. let the euclidean points $\mathbf{c}, \mathbf{d}$ lie between $\mathbf{a}$ and $\mathbf{b}$. The scalar product with the correctly scaled vector defines a coordinate frame

$$\mathbf{i} = \frac{\mathbf{a} - \mathbf{b}}{\|\mathbf{a} - \mathbf{b}\|^2}, \qquad (3)$$

between $b = 0$ and $a = 1$ on the line with

$$c = \mathbf{i}^\top \left(\mathbf{c} - \mathbf{b}\right) < d = \mathbf{i}^\top \left(\mathbf{d} - \mathbf{b}\right) \in \ ]\,0, 1\,[. \qquad (4)$$

Figure 3. Left: Example of a randomly generated phantom containing several randomly distributed pens (sets of four collinear beads each). Lines with different cross-ratios are color-encoded by their index. Right: Imperfect detection of the same cross-ratios in a simulated projection. In this case there were 30 correct and 8 incorrect initial point matches in $\mathcal{M}_\star$. After RANSAC, all 96 detected beads were correctly matched, while 6 beads were truncated at the top and 6 were not detected due to an overlap.

For points not exactly the measured cross-ratio is

$$\tilde{\mathrm{cr}}\,(\mathbf{a},\,\mathbf{b},\,\mathbf{c},\,\mathbf{d}) \overset{\text{def}}{=} \mathrm{cr}\,(a,b,c,d) = \frac{d - cd}{c - cd} > 0. \qquad (5)$$

Note that $\mathbf{i}^\top \cdot$ is itself a linear projection, to which the cross-ratio is invariant, so this is consistent with Equation 2 provided that the points are exactly collinear. We observe a hyperbolic growth from $\lambda = 0$ at $d = 0$ towards infinity at $d = 1$.

C. Phantom Design

We suggest a phantom design based on several sets of four collinear metal beads. We allow two different radii of beads $r_l > r_s$. The difference between the radii should be sufficient to make classification into large and small beads easy, even under projective distortion (size is *not* a projective invariant). Each collinear four-set shall contain exactly one large bead, because the cross-ratio is symmetric with respect to the order of points $\mathrm{cr}\,(a,b;c,d) = \mathrm{cr}\,(d,c;b,a)$. We define that the large bead shall be named either $\mathbf{b}$ or $\mathbf{c}$, thus resolving the ambiguity. In applications where differently sized beads are not an option, $p2$ invariants [5] provide an alternative.

IV. CALIBRATION ALGORITHM

A. Bead Detection and Approximate Collinearity

We employ the fast radial symmetry transform (FRST) [4] for bead detection, parametrized by two sets of radii in image pixels which correspond to the sizes of projected large and small beads. In experiments we found that running the algorithm twice sometimes produces spurious detections of small beads. We therefore remove detected small beads, when their center is less than one radius away from the center of a detected large bead. Selection of parameters for the FRST is done manually within this work. Next, we extract sets of collinear points. We loop over all four-sets of detected points which contain a large bead and test for approximate collinearity with a simple distance-threshold, called candidate pins. We suggest two other heuristic priors for better stability. First, we ignore all large beads which form part of more than three candidate pins. If two coplanar 3D pins project to the

same 2D line, this produces $2 \cdot \mathrm{binomial}(3, 6) = 40$ candidates of which only two are correct detections, so it is safer to just ignore all of them. Second, we assume that beads are spread out along the pin. If the length of a pin is larger than 20 times the shortest distance between any two of its beads, it is an unlikely candidate and we also ignore it.

B. Descriptor and Initial Matching

Since we require $b < c < d < a$, we always obtain cross-ratios $\mathrm{cr}\,(a,b,c,d) > 0$. This will be identical whether $\mathbf{b}$ or $\mathbf{c}$ is the large bead in the configuration. To differentiate between those two cases, we use λ as a descriptor if $\mathbf{b}$ is a large bead, and $-\lambda$ otherwise. The initial matching is defined by assignment of the best match of all detected sets of four collinear points, based on the descriptor. The initial matching is given by $\mathcal{M}_\star \subset \mathbb{N}_0^2$, which contains indices of two points which are candidate matches. However, $\mathcal{M}_\star$ may still contain a relatively large number of outliers. An example for this case is visualized in Figure 3, right.

C. Robust Estimation of Projection Geometry

A algebraic estimate of the projection matrix $\mathbf{P}$ is unreliable in the presence of outliers. We employ RANSAC to find, with high probability, a set $m \subset \mathcal{M}_\star$ with $|m| = 6$, whose estimated projection matrix $\mathbf{P}_m$ explains most detected points given the known 3D bead locations. By assignment of these beads to the closest projection of the known 3D beads, we can establish an improved and relatively complete set $\mathcal{M} \subset \mathbb{N}_0^2$, which likely contains many more point matches than $\mathcal{M}_\star$ but no outliers. The final step is a re-estimation of $\mathbf{P}$ using all points contained in $\mathcal{M}$. This work is restricted to a straight-forward algebraic estimation using DLT. A slightly better solution may be found using non-linear optimization of the geometric error and bundle-adjustment.

V. VALIDATION AND EXPERIMENTS

A. Validation with Source Positions on a Sphere

For validation we created a digital phantom based on an approximately equal distribution of points on a sphere surface using the Fibonacci series and the golden angle. It is randomized by misaligning collinear points from the ray though to the center of the sphere with $m = 27$ lines and $n = 108$ beads at a size of 3.2 mm for large beads and 1.6 mm for small beads (same bead sizes and number of points as the PDS2 phantom). An instance of such a phantom is shown in Figure 3. We present 3456 noise-free projections of 1240 px $\times$ 960 px of a phantom . We presents a validation with with a spacing of $0.308\,\frac{\text{mm}}{\text{px}}$ from all directions (sampled by equal longitudinal and latitudinal angles) with a source-to-detector-distance of 325 mm and a source-isocenter-distance of 200 mm, where the source positions are distributed on a sphere instead of a circle. We simulate the phantom instance from Figure 3, to compute projection matrix $\tilde{\mathbf{P}}$ with the proposed algorithm (without non-linear refinement) and compare to the ground truth projection $\mathbf{P}$. The target projection

Figure 4. Distribution of errors and number of points used in 3456 calibrations from Figure 5, top. There were no outliers.


Figure 5. Evaluation of matching algorithm. Color encodes the number of points correctly detected and used for estimation of the projection matrix. The displayed range is from 54 (deep blue) to 108 (bright green). Top row: Results for an evaluation with projection direction vectors on a sphere. See also Figure 4, right. Bottom row: Results for the same evaluation using a naive algorithm for the PDS2 phantom.

error (TPE) is evaluated on 500 random points within the field of view of ~ 20 cm radius according to

$$\text{TPE} = \sum_i \text{d} \left(\tilde{\mathbf{P}} \cdot \mathbf{y}_i, \, \mathbf{P} \cdot \mathbf{y}_i \right). \tag{6}$$

The results are themselves projected to a sphere for visualization, see Figure 5. A distribution of errors, residuals and number of points can be found in Figure 4. For all 3456 projections more than 54 points were correctly matched, with an average of 100.1 out of 108 points detected. A mean error of 0.73 px was achieved at a mean residual of just 0.177 px. The mean error of the estimated source position was 0.43 mm.

We repeated the experiment for projections of the PDS2 phantom. The PDS2 phantom is designed for circular trajectories so its detection for steep angles is more difficult and a comparison is fair only close to the equator. There were 1243 usable projections with more than 54 points correctly matched, compare Figures 5, bottom. The results for those views are comparable to the suggested phantom with a mean target projection error of 0.79 px. The residual of 0.21 px. The mean error of the estimated source position was 0.63 mm.

VI. DISCUSSION AND CONCLUSION

We present a new phantom design for geometric calibration of FD-CT systems. Users may quickly build phantoms of arbitrary size and shape, all of which can be analyzed with the same software. The phantoms are comprised of short pens that contain four metal beads each. The detection and matching is based on the cross-ratio, which allows robust detection, even when parts of the phantom are truncated. We present a proof of concept with a randomized numerical phantom. The suggested phantoms can work equally well if viewed from all spacial directions and support a more equal distribution of beads in space. Manufacturing accuracy does not limit the calibration accuracy, since the descriptor is the missing link to apply standard Computer Vision algorithms such as factorization and bundle-adjustment to recover both geometry of the phantom

and parameters of the projection. The long-term goal of the project is to 3D print low-cost elements (pens), which the user can freely distribute in space. Application with unusual trajectories or varying size of the scanned objects are ideal for an application in material testing, for instance. The algorithm may also be useful to research as a tool for marker based tracking and fiducial-based registration.

REFERENCES

[1] M. J. Daly, J. H. Siewerdsen, Y. B. Cho, D. A. Jaffray, and J. C. Irish. Geometric calibration of a mobile c-arm for intraoperative cone-beam ct. *Medical physics*, 35(5):2124–2136, 2008.

[2] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.

[3] M. W. Jacobson, M. D. Ketcha, Sa. Capostagno, A. Martin, A. Uneri, J. Goerres, T. De Silva, S. Reaungamornrat, R. Han, A. Manbachi, J. W. Stayman, S. Vogt, G. Kleinszig, and J. H. Siewerdsen. A line fiducial method for geometric calibration of cone-beam ct systems with diverse scan trajectories. *Physics in Medicine and Biology*, 2017.

[4] G. Loy and A. Zelinsky. Fast radial symmetry for detecting points of interest. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):959–973, Aug 2003.

[5] Peter Meer, Reiner Lenz, and Sudhir Ramakrishna. Efficient invariant representations. *International Journal of Computer Vision*, 26(2):137–152, 1998.

[6] C. Mennessier, R. Clackdoyle, and F. Noo. Direct determination of geometric alignment parameters for cone-beam scanners. *Physics in medicine and biology*, 54(6):1633, 2009.

[7] N. K. Strobel, B. Heigl, T. M. Brunner, O. Schuetz, M. M. Mitschke, K. Wiesent, and T. Mertelmeier. Improving 3 d image quality of x-ray c-arm imaging systems by using properly designed pose determination systems for calibrating the projection geometry. In *Proceedings of SPIE*, volume 5030, pages 943–954, 2003.

[8] S. Thürauf, O. Hornung, M. Körner, F. Vogt, M. A. Nasseri, and A. Knoll. Evaluation of a 9d-position measurement method of a c-arm based on x-ray projections. In *Proc. of MICCAI: 1st Interventional Microscopy Workshop. Munich, Germany*, pages 9–16, 2015.

[9] L. Von Smekal, M. Kachelrieß, E. Stepina, and W. A. Kalender. Geometric misalignment and calibration in cone-beam tomography. *Medical physics*, 31(12):3242–3266, 2004.