



Stochastic Optimization based Design of C-Arm Calibration Phantoms

Peter Zhang¹, Amandine Gout¹, Youyang Shen¹, Philipp Fürnstahl², and Hooman Esfandiari¹

¹ X23D AG, Zurich, Switzerland
peter.zhang@balgrist.ch

² Research in Orthopedic Computer Science, Balgrist University Hospital, University of Zurich

Abstract

C-arm imaging plays a crucial role in intraoperative surgical guidance, with many methods requiring precise information about the C-arm calibration parameters, achieved through calibration phantoms with known geometries. We propose leveraging a stochastic optimization approach to design calibration phantoms, ensuring accuracy in the presence of noise. Our method integrates performance and usability criteria within the optimization, enabling the design of specialized calibration phantoms.

1 Introduction

C-arm fluoroscopy is commonly used in Computer-Assisted Surgery for real-time imaging of the patient anatomy. Once these images are calibrated, techniques such as 2D/3D registration [6, 7] or recent reconstruction algorithms such as [4, 5] can be used for intraoperative guidance. Accurate calibration is critical for such methods, underscoring the importance of phantom design. Prior work [3, 2] demonstrated improved performance using phantoms with spheres, lines, and ellipses as compared to sphere-only designs. Nonetheless, more recent applications continue to exclusively utilize spheres as fiducial markers for practical reasons. To balance ease of use and performance, we focus on sphere-based phantoms and propose a stochastic optimization method to determine their placement. Our approach considers the following key criteria:

- Insensitivity to 2D localization (i.e. segmentation) errors
- Visibility of spheres in the clinically relevant viewing directions
- Adherence to phantom size constraints

We achieve these objectives using a tunable cost function, based on a mathematical representation of the intraoperative C-arm calibration process. By aligning phantom size and views with the intended procedure, our method ensures use case-specific optimization.




Figure 1: Rendering depicting how the spheres (orange) of a phantom would be projected onto the X-ray image plane (black) within the AP and Lateral views.

2 Calibration Phantom Optimization

2.1 Setup

We modeled the C-arm as a pinhole camera to mathematically represent the calibration process. The calibration phantom composed of n_f fiducial spheres can be placed near the anatomy of interest, as shown in Figure 1. For any relevant views (e.g. AP or Lateral), the 3D sphere coordinates were projected onto the simulated C-arm image plane. Using the Direct Linear Transform (DLT) algorithm, the C-arm can be calibrated given these 2D-3D point correspondences. To simulate realistic conditions, normally distributed noise was added to the 2D image coordinates, mimicking segmentation errors that lead to calibration inaccuracies.

2.2 Cost Function

Assuming a bounding constraint on the sphere placement during optimization, we tackled the remaining two criteria by optimizing the following cost function:

$$L = \frac{1}{n_v} \sum_{i=1}^{n_v} (w_{seg} L_{seg}^i + w_{occl} L_{occl}^i + w_{vis} L_{vis}^i).$$

As noise insensitivity and bead visibility can be conflicting, the cost comprises of a tunable weighted sum, averaged over all relevant views n_v . It includes three components:

- L_{seg} : reduces calibration error for noisy 2D coordinates.
- L_{occl} : penalizes occluding spheres in the image plane.
- L_{vis} : ensures all spheres remain in the field of view.

Segmentation Error Cost

We define the segmentation error cost L_{seg} as a weighted sum of errors, calculated by comparing noise-affected calibration parameters to ground truth values. Intrinsic (focal length, focal point, skew) and extrinsic (rotation, translation) errors were averaged across n_s samples, with Gaussian noise applied to the projected sphere coordinates for each sample.

$$L_{seg} = \frac{1}{n_s} \sum_{i=1}^{n_s} (w_{focl} L_{focl}^i + w_{focp} L_{focp}^i + w_{skew} L_{skew}^i + w_r L_r^i + w_t L_t^i)$$

Occlusion Cost

The occlusion cost L_{occl} was introduced a soft constraint, linear when the distance between two spheres d_{ij} was below a threshold t_{occl} :

$$L_{occl} = \sum_{i=1}^{n_f} \sum_{j=i}^{n_f} \begin{cases} |d_{ij} - t_{occl}| & \text{if } d_{ij} - t_{occl} < 0, \\ 0 & \text{otherwise.} \end{cases}$$

Visibility Cost

Similarly, the visibility cost L_{vis} quadratically penalized spheres projected outside of the image boundaries \mathbf{u}_{tol} :

$$L_{vis} = \sum_{i=1}^{n_f} \begin{cases} 0 & \text{if } \mathbf{u}_i \in \mathbf{u}_{tol}, \\ \|\mathbf{u}_i - \mathbf{u}_{tol}\|^2 & \text{otherwise.} \end{cases}$$

2.3 Optimizer

The 3D coordinates of each fiducial sphere served as optimization parameters to minimize the cost L . Given the high-dimensional ($3n_f$) and non-convex nature of the cost function, we employed a stochastic optimization approach, leveraging the SciPy [8] implementation of the Dual Annealing algorithm derived from [10].

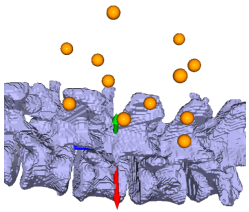
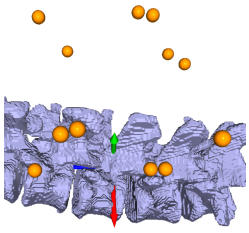
3 Results

Constraining 12 beads within a boundary of 7x7x3 cm, we compared the calibration accuracy of a phantom with randomly distributed spheres to that of an optimized one derived using the herein presented algorithm. The optimization strategically placed spheres near the boundaries, while avoiding occlusions. This improved the calibration performance under potential 2D segmentation errors (Table 1) and reduced occlusions across all relevant views.

4 Discussion

The proposed method for designing C-arm calibration phantoms improves calibration accuracy in the presence of segmentation noise, as compared to randomly placing the fiducial spheres within the available space. Previous works, report the accuracy directly on the intended application, such as e.g. the mean target registration error [1, 6, 9], or they do not go into detail on the sphere placement [3], and can thus not be quantitatively compared. Our approach offers

Table 1: Comparison of the average calibration accuracy between an optimized and a randomized phantom. The average is based on $n_v = 216$ views with $n_s = 100$ samples draw from the normally distributed noise with a mean of 0.2 mm and standard deviation of 0.15 mm.

	Focal Length Error [mm]	Focal Point Error [mm]	Skew Error [mm]	Rotation Error [degrees]	Translation Error [mm]
Random 	146.57 ± 9.16	117.05 ± 6.83	13.63 ± 1.40	8.48 ± 0.45	108.30 ± 8.59
Optimized 	53.90 ± 2.46	48.48 ± 1.98	4.19 ± 0.25	3.51 ± 0.14	36.58 ± 1.79

a tool to generate a view-specific phantoms satisfying clinically and technically relevant constraints, which could be integrated into more advanced calibration methods to further increase calibration robustness and accuracy.

References

- [1] Shahram Amiri, David R. Wilson, Bassam A. Masri, and Carolyn Anglin. A low-cost tracked C-arm (TC-arm) upgrade system for versatile quantitative intraoperative imaging. *International Journal of Computer Assisted Radiology and Surgery*, 9(4):695–711, July 2014.
- [2] Ameet Jain and Gabor Fichtinger. C-arm Tracking and Reconstruction Without an External Tracker. In Rasmus Larsen, Mads Nielsen, and Jon Sporring, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, pages 494–502, Berlin, Heidelberg, 2006. Springer.
- [3] Ameet Kumar Jain, Tabish Mustafa, Yu Zhou, Clif Burdette, Gregory S. Chirikjian, and Gabor Fichtinger. FTRAC—A robust fluoroscope tracking fiducial. *Medical Physics*, 32(10):3185–3198, 2005. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1118/1.2047782>.
- [4] Sascha Jecklin, Carla Jancik, Mazda Farshad, Philipp Fürnstahl, and Hooman Esfandiari. X23D-Intraoperative 3D Lumbar Spine Shape Reconstruction Based on Sparse Multi-View X-ray Data. *Journal of Imaging*, 8(10):271, October 2022.
- [5] Sascha Jecklin, Youyang Shen, Amandine Gout, Daniel Suter, Lilian Calvet, Lukas Zingg, Jennifer Straub, Nicola Alessandro Cavalcanti, Mazda Farshad, Philipp Fürnstahl, and Hooman Esfandiari. Domain adaptation strategies for 3D reconstruction of the lumbar spine using real fluoroscopy data. *Medical Image Analysis*, 98:103322, December 2024. arXiv:2401.16027 [cs].

- [6] Yoshito Otake, Mehran Armand, Robert S. Armiger, Michael D. Kutzer, Ehsan Basafa, Peter Kazanzides, and Russell H. Taylor. Intraoperative image-based multiview 2D/3D registration for image-guided orthopaedic surgery: incorporation of fiducial-based C-arm tracking and GPU-acceleration. *IEEE transactions on medical imaging*, 31(4):948–962, April 2012.
- [7] Barbara Postolka, Renate List, Benedikt Thelen, Pascal Schütz, William R. Taylor, and Guoyan Zheng. Evaluation of an intensity-based algorithm for 2D/3D registration of natural knee videofluoroscopy data. *Medical Engineering & Physics*, 77:107–113, March 2020.
- [8] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [9] J. B. West, J. M. Fitzpatrick, S. A. Toms, C. R. Maurer, and R. J. Maciunas. Fiducial point placement and the accuracy of point-based, rigid body registration. *Neurosurgery*, 48(4):810–816; discussion 816–817, April 2001.
- [10] Y Xiang, D. Y Sun, W Fan, and X. G Gong. Generalized simulated annealing algorithm and its application to the Thomson model. *Physics Letters A*, 233(3):216–220, August 1997.