

Camera Calibration for Synthetic Digital Tomography

Michael Albert, Ann-Katherine Carton, Andrew D.A. Maidment

Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104

ABSTRACT

A method for calibrating camera geometry is described. This method has been used to implement synthetic tomography on a commercially available full-field digital mammography system. The method utilizes a phantom containing six point-like calibration objects whose positions are approximately known. The image of five calibration objects in a given projection allows an associated projection matrix to be determined up to one free parameter. By using the positions of the shadows of the sixth calibration object in three or more views, one can fit the remaining free parameter associated with each view and the position of the sixth calibration object relative to the first five. Uncertainty in the position or geometry of the phantom does not affect the geometric consistency, thus tomograms produced by back-projection suffer no blurring from errors in the determination of camera geometry. Uncertainties in the position or geometry of the phantom result in proportionate translations or distortions of the tomograms. For a tomogram corresponding to a plane containing an object, the positions of the backprojections of the shadows of the object are consistent to the same precision as the measurements of the shadows in each projection, *i.e.*, the positions of the backprojections differ by about the size of the pixel spacing in the detector.

Keywords: Synthetic tomography, camera calibration, projective geometry, digital mammography

1. INTRODUCTION

A significant limitation on the efficacy of projection imaging is imposed by the inherent summation of multiple layers of anatomy. This “anatomical noise” can produce both false-negatives, as when a lesion is obscured by overlying anatomy, and false-positives or “summation artifacts,” as when a lesion is not properly distinguished from the clutter produced by the summation of various layers of normal anatomy. The importance of anatomical noise has been studied in the context of a variety of medical-imaging tasks including chest radiography¹ and mammography.²

Synthetic tomography combines multiple projections acquired at distinct angles to produces images focused on a given anatomical plane, while other anatomical planes are blurred. The utility of tomography in overcoming anatomical noise is discussed in the review by Samei, Eyler and Baron.³

In order to combine information from the various projections it is important to have accurate geometric information. To be explicit, let (x, y, z) represent the coordinates of a point in space, and let (u_α, v_α) represent the position on an area detector where the x-ray shadow of that point falls. The index α indicates which image, of a series of acquired images, is being considered. Accurate geometric knowledge allows one to determine functions f_α such that

$$(u_\alpha, v_\alpha) = f_\alpha(x, y, z). \quad (1)$$

The observation of the shadow of a point-like object in a projection gives two constraints on three unknowns, thus, as x-rays travel in straight lines, determining a line upon which the object must lie. Because of uncertainties in geometrical information, the lines determined from the various views will generally be skew and there will

Further author information: (Send correspondence to A.D.A.M.)

M.A. E-mail: michael.albert@uphs.upenn.edu

A.-K.C. E-mail: ann-katherine.carton@uphs.upenn.edu

A.D.A.M.: E-mail: andrew.maidment@uphs.upenn.edu, Telephone: 215-746-8763,

Address: Dept. of Radiology, 1 Silverstein, 3400 Spruce St., Philadelphia, PA 19104

be no spatial point (x, y, z) which, within uncertainties of measurement, lies on all of the lines so determined. For a simple back-projection algorithm, this would result in the object never being in focus, regardless of the plane which is nominally to be the focus of the tomosynthetic procedure. Clinically, small objects of high spatial-frequency such as small calcifications or the boundaries of some lesions might be obscured.

The problem of “camera calibration” has been extensively studied⁴ in the field of computer vision, although such research generally relates to the acquisition of stereo-pairs. Additionally, computer vision research is generally interested in cameras which work at or near visible wavelengths, so that the focal-plane and the object are on opposite sides of a lens; for x-ray imaging, of course, the object or patient being imaged is between the x-ray detector and the x-ray focus. In adapting the current generation of clinically available digital radiography systems to tomosynthesis, reports in the literature discuss a combination of the use of the nominal geometry specified by the manufacturer and various calibration techniques which incompletely determine the geometry. For example, one study using a digital-mammography system performed translations on projection images to align a single point-like marker between the several images,⁵ but this does not assure the geometric consistency of backprojections. Metz and Fencil^{6,7} have introduced a technique using eight calibration points, but this technique assumes that the intrinsic geometry of the camera (*i.e.*, the position of the x-ray focus relative to the camera) is known *a priori*.

Here we discuss a method of camera calibration applicable when at least three views are acquired. The method utilizes a phantom containing six point-like calibration objects. The positions of five of these objects determine a projective coordinate system for the volume being imaged. The sixth point-like object is used to determine the projection functions f_α in terms of this projective coordinate system. The technique results in a camera calibration which is consistent in the sense that, up to inaccuracies in the measurements of the positions of the shadows, for any shadow of a point-like object seen in multiple views there will exist a point in space whose projection accounts for all of those shadows. Experimental uncertainties in the positions (or rather, the relative positions) of the point-like objects used for calibration is absorbed into a projective transformation. If the positions of the point-like objects used for calibration are known accurately, this transformation will be approximately an identity transform over a volume of interest which includes the calibration objects. To the extent that the positions of the calibration objects are not precisely known, the camera calibration will still be consistent, but will result in a distortion corresponding to a projective transformation. Such transformations will be in general well behaved. This is supported first by the fact that projective transformations preserve collinearity between points. Second, the projective transformation will vary smoothly from the identity transformation as a function of the errors in the positions of the calibration objects, so that images will be clinically acceptable except for extremely large errors.

2. THEORY

Consider, as is conventional in computer vision,⁸ the parameterization of physical space in terms of projective coordinates $(x, y, z, w) \in \mathbb{R}^4$, where \mathbb{R}^4 is a four-dimensional real vector space. For $w \neq 0$, this corresponds to the physical point $(x/w, y/w, z/w) \in \mathbb{E}^3$, where \mathbb{E}^3 is a three-dimensional Euclidean space. Vectors for which $w = 0$ correspond to idealized “points at infinity.” Similarly, positions on an area detector are parameterized in terms of projective coordinates $(u, v, s) \in \mathbb{R}^3$, corresponding to a physical position on the detector $(u/s, v/s) \in \mathbb{E}^2$ for $s \neq 0$. Note that these vectors are “column vectors” but, for typographical reasons, will be written as row vectors in running text.

Given these conventions, the operation of projecting a point in space onto a point on the surface of an area detector can be written as

$$\begin{pmatrix} u \\ v \\ s \end{pmatrix} = \mathcal{A} \begin{pmatrix} x \\ y \\ z \\ w \end{pmatrix} \quad (2)$$

where \mathcal{A} is a 3×4 matrix (three rows and four columns). To connect with the notation of the introduction, given a spatial point with coordinates (x, y, z) , one would set $w = 1$ in equation 2 and then

$$(u/s, v/s) = f(x, y, z) \quad (3)$$

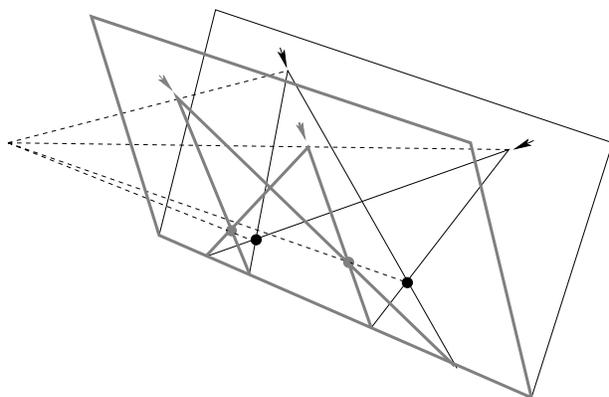


Figure 1. The nature of a projective transform can be illustrated in two dimensions. The two planes represent two possible configurations of objects (dots), “x-ray sources” (arrows) and a detector “line” (the line common to the two planes). The configurations in the two planes are not congruent (*e.g.*, the distances of the objects from the “detector” differ). However, the configurations can not be distinguished based upon the shadows of the objects observed on the detector. This follows from the observation that the dotted lines “project” one plane onto the other.

where the index α has been omitted for convenience. A general projective transformation of three-dimensional space can be written

$$\begin{pmatrix} x \\ y \\ z \\ w \end{pmatrix} = \mathcal{B} \begin{pmatrix} x' \\ y' \\ z' \\ w' \end{pmatrix} \quad (4)$$

where \mathcal{B} is a 4×4 matrix. Composition of a projective transformation of three dimensional space with a projection onto a two dimensional space yields an equation of the same form as the original projection, *i.e.*

$$\begin{pmatrix} u \\ v \\ s \end{pmatrix} = \mathcal{A}' \begin{pmatrix} x' \\ y' \\ z' \\ w' \end{pmatrix} \quad (5)$$

where

$$\mathcal{A}' = \mathcal{A}\mathcal{B} \quad (6)$$

is also a 3×4 matrix. The nature and use of the projective transform \mathcal{B} is illustrated in figure 1 for the corresponding two dimensional situation. The transformation, represented by the dotted lines, takes one configuration of points on one plane to a configuration of points in a second plane while preserving the relationship of collinearity. The transform \mathcal{B} works analogously, but in three dimensional space.

It is a fundamental theorem of three dimensional projective geometry that any five points in general position can be taken to any other five points in general position by a projective transformation, and that the resulting transformation is uniquely defined. In three dimensional space, a set of points are said to be in “general position” if no four points lie in a single plane. Projective transformations are precisely those transformations which are linear in terms of projective coordinates, as in equation 4, and which preserve the collinearity of sets of points. In terms of projective coordinates, this theorem follows readily from the fact that any set of four linearly independent vectors in \mathbb{R}^4 can be mapped into any other set of four vectors, and the fact that these vectors can be arbitrarily rescaled without changing the point in space which each represents. These results are discussed in several texts, *e.g.* Artin⁹ (section II.10 and theorem 2.28).

The phantom used for calibration in this technique contains six point-like objects in general position, with spatial coordinates given by \mathbf{r}_i , $i = 1, \dots, 6$, where each \mathbf{r}_i represents the spatial coordinates x_i , y_i , and z_i of the object in space. In general, these spatial positions are not precisely known. There is, of course, uncertainty due to the positioning of the phantom, which corresponds to a rigid translation and rotation of the coordinate system. The accuracy with which the relative positions of the objects in the phantom are known may also be limited. Independent of such knowledge, it is possible, as discussed in the previous paragraph, to choose a projective transformation which associates the positions of the first five objects \mathbf{r}_i with five arbitrarily assigned positions

\mathbf{r}'_i , where each \mathbf{r}'_i represents a position in the projectively transformed space. Such a completely arbitrary choice is probably of little use, and in general one would choose the values of the projectively-transformed positions \mathbf{r}'_i to be the nominal positions of the objects as determined by direct measurement. The relevant fact is that any uncertainty in the positions of the first five objects is absorbed into the projective transform. In the transformed space, the positions are known exactly.

It is now asserted that, for a given projection, knowledge of the positions of five point-like objects in space and their projections determines the projection matrix \mathcal{A}' up to a single free parameter. Note that the matrix \mathcal{A}' contains 12 entries. However, an overall rescaling of the matrix (*i.e.* multiplying each element by a non-zero constant) does not change the projection, so that the projection is specified by 11 free parameters. For each point, the coordinates of each shadow give two constraints, as the position of the shadow has two coordinates in the imaging plane. Thus, the observation of five shadows gives $5 \times 2 = 10$ constraints, leaving one free parameter. An explicit formula for this one-parameter family of 3×4 matrices can be easily derived using projective coordinates. Let

$$\rho'_i = \begin{pmatrix} \mathbf{r}'_i \\ 1 \end{pmatrix} = \begin{pmatrix} x'_i \\ y'_i \\ z'_i \\ 1 \end{pmatrix} \quad (7)$$

be the canonical representative of \mathbf{r}'_i in projective coordinates. Similarly, let

$$\pi_i = \begin{pmatrix} u_i \\ v_i \\ 1 \end{pmatrix} \quad (8)$$

be the projective coordinates of the positions of the shadows on the plane of the detector. As the points $\mathbf{r}'_1, \dots, \mathbf{r}'_4$ are in general position, the four-dimensional vectors ρ'_1, \dots, ρ'_4 are linearly independent. Thus it is possible to determine dual vectors $\hat{\rho}'_1, \dots, \hat{\rho}'_4$ such that

$$\hat{\rho}'_i{}^t \rho'_j = \delta_{ij} = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases} \quad (9)$$

for $i, j \in \{1, 2, 3, 4\}$. Similarly, one can choose vectors $\hat{\pi}_1, \hat{\pi}_2$, and $\hat{\pi}_3$ in \mathbb{R}^3 such that

$$\hat{\pi}_i{}^t \pi_j = \delta_{ij} \quad (10)$$

for $i, j \in \{1, 2, 3\}$. A direct calculation using the duality conditions expressed by equations 9 and 10 then shows that the 3×4 matrix

$$\mathcal{C}' = \frac{\hat{\pi}_1{}^t \pi_5}{\hat{\rho}'_1{}^t \rho'_5} \pi_1 \hat{\rho}'_1{}^t + \frac{\hat{\pi}_2{}^t \pi_5}{\hat{\rho}'_2{}^t \rho'_5} \pi_2 \hat{\rho}'_2{}^t + \frac{\hat{\pi}_3{}^t \pi_5}{\hat{\rho}'_3{}^t \rho'_5} \pi_3 \hat{\rho}'_3{}^t \quad (11)$$

projects ρ_1, ρ_2, ρ_3 and ρ_5 onto π_1, π_2, π_3 and π_5 , respectively. Explicitly, for $i \in \{1, 2, 3\}$ two of the three terms vanish and

$$\mathcal{C}' \rho'_i = \frac{\hat{\pi}_i{}^t \pi_5}{\hat{\rho}'_i{}^t \rho'_5} \pi_i \propto \pi_i \quad (12)$$

which is sufficient as projective coordinates can be arbitrarily rescaled. For π_5 , the denominators in each term cancel explicitly and

$$\mathcal{C}' \rho'_5 = (\hat{\pi}_1{}^t \pi_5) \pi_1 + (\hat{\pi}_2{}^t \pi_5) \pi_2 + (\hat{\pi}_3{}^t \pi_5) \pi_3 = \pi_5. \quad (13)$$

For ρ_4 , the conditions specified by equation 9 give

$$\mathcal{C}' \rho'_4 = 0. \quad (14)$$

Similarly it is seen that for the matrix

$$\mathcal{D}' = \frac{\hat{\pi}_1{}^t \pi_4}{\hat{\rho}'_1{}^t \rho'_5} \pi_1 \hat{\rho}'_1{}^t + \frac{\hat{\pi}_2{}^t \pi_4}{\hat{\rho}'_2{}^t \rho'_5} \pi_2 \hat{\rho}'_2{}^t + \frac{\hat{\pi}_3{}^t \pi_4}{\hat{\rho}'_3{}^t \rho'_5} \pi_3 \hat{\rho}'_3{}^t - \frac{1}{\hat{\rho}'_4{}^t \rho'_5} \pi_4 \hat{\rho}'_4{}^t \quad (15)$$

one readily obtains that for $i \in \{1, 2, 3\}$

$$\mathcal{D}' \rho'_i = \frac{\hat{\pi}_i^t \pi_4}{\hat{\rho}_i^t \rho'_5} \pi_i \propto \pi_i \quad (16)$$

while for $i = 4$

$$\mathcal{D}' \rho'_4 = -\frac{1}{\hat{\rho}_4^t \rho'_5} \pi_4 \propto \pi_4 \quad (17)$$

and

$$\mathcal{D}' \rho'_5 = (\hat{\pi}_1^t \pi_4) \pi_1 + (\hat{\pi}_2^t \pi_4) \pi_2 + (\hat{\pi}_3^t \pi_4) \pi_3 - \pi_4 = 0. \quad (18)$$

Using equations 12–14, and equations 16–18, it is seen that any linear combination of \mathcal{C}' and \mathcal{D}' will properly describe a projection from $\mathbf{r}'_1, \dots, \mathbf{r}'_5$ to $(u_1, v_1), \dots, (u_5, v_5)$ in terms of projective coordinates, so long as both \mathcal{C}' and \mathcal{D}' have non-vanishing coefficients. As an overall rescaling of \mathcal{A}' will not affect the projection, the coefficients of \mathcal{C}' and \mathcal{D}' really represent one degree of freedom, and the family of projection matrices \mathcal{A}' can be written

$$\mathcal{A}' = \cos \omega \mathcal{C}' + \sin \omega \mathcal{D}' \quad (19)$$

with the constraint that $\cos \omega \neq 0$ and $\sin \omega \neq 0$.

The sixth point-like calibration object can now be used to determine the remaining degree of freedom. Note that the coordinates of the sixth object can not be set arbitrarily, but the projectively transformed coordinates \mathbf{r}'_6 must themselves be determined. Thus there are three unknown parameters representing the coordinates of the point \mathbf{r}'_6 and one free parameter for each projection. However, in each projection one measures two coordinates for the position of the shadow, thus given three projections one has $3 + 3 = 6$ unknowns and $2 \times 3 = 6$ measurements, which is precisely sufficient to determine \mathbf{r}'_6 and the free parameter associated with each projection. If more than three projections have been acquired, then the system is over determined. In either case, the number of free parameters is sufficiently small that they can be obtained using numerical fitting algorithms such as those available from the Minuit¹⁰ package. To show this explicitly, it is convenient to re-introduce the index α to distinguish the various projections. Thus the positions of the i -th point in the projection indexed by α will be written $(u_{\alpha i}, v_{\alpha i})$, the projection matrices can be written

$$\mathcal{A}'_{\alpha} = \cos \omega_{\alpha} \mathcal{C}'_{\alpha} + \sin \omega_{\alpha} \mathcal{D}'_{\alpha} \quad (20)$$

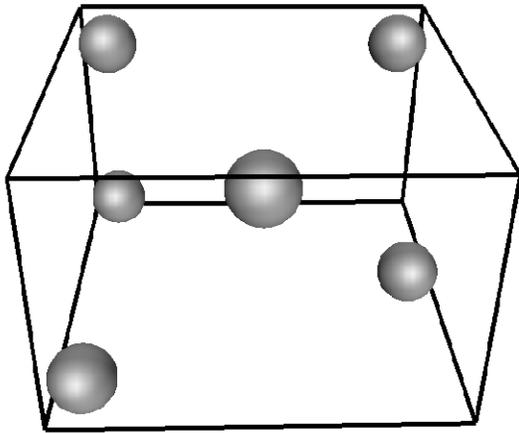
and each matrix \mathcal{A}'_{α} induces a projection f'_{α} as in equations 2 and 3. The values of $\mathbf{r}'_6 = (x'_6, y'_6, z'_6)$ and ω_{α} can then be determined by a least squares fit

$$\chi^2 = \sum_{\alpha=1}^N |(u_{\alpha 6}, v_{\alpha 6}) - f'_{\alpha}(x'_6, y'_6, z'_6)|^2 \quad (21)$$

assuming a total of N projections. In equation 21 “ $|\cdot|$ ” is the usual two-dimensional Euclidean length and $(u_{\alpha 6}, v_{\alpha 6})$ is the projection of \mathbf{r}'_6 in the image indexed by α ; the dependence of f'_{α} on ω_{α} is not explicitly shown, but it must be understood that a total of $N + 3$ values are being adjusted in the fitting procedure. This χ^2 function uses the Euclidean distances measured in the projections, thus avoiding inaccurate techniques sometimes associated with projective geometry such as solving for the “fundamental matrix” via the associated linear equation.¹¹

3. SIMULATION

Simulation studies were carried out using a phantom of 10 cm \times 10 cm \times 6 cm consisting of six point-like calibration objects, no four of which lie in a plane, as illustrated in figure 2. Geometries were simulated consisting of $N_{proj} = 3, 5, 7,$ or 9 projections, each projection corresponding to a unique position of the x-ray focus. The simulated x-ray focus was nominally 68 cm above the surface of the detector and traveled along a path corresponding to $\approx \pm 24^\circ$. During the simulation of each projection image, the position of the x-ray focus was randomly perturbed by adding to each coordinate a random displacement drawn from a Gaussian distribution with $\sigma_t = 5$ cm, in order to simulate uncertainties related to the nominal positions a vendor might specify. Before



id	x(cm)	y(cm)	z(cm)
1	0	10.00	0.00
2	10.16	4.92	0.00
3	0	-0.16	0.00
4	0	10.16	6.096
5	5.08	0	6.096
6	10.16	10.16	6.10

Figure 2. A three dimensional phantom containing six point-like objects, no four of which lie on a plane. For visualization the objects have been drawn as large spheres. The table gives the nominal positions of the objects relative to the coordinate system used in the simulation.

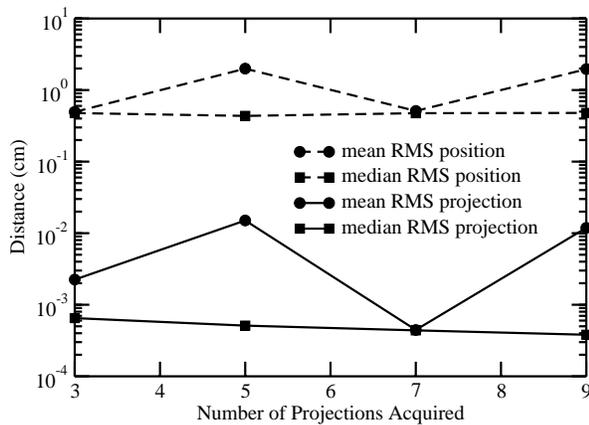


Figure 3. Simulated random variations (≈ 0.4 cm) in the positions of the calibration objects from their nominal positions result in errors of similar magnitude in the reconstructed positions, but the shadows of the estimated positions agree with the simulated shadows to 0.001–0.01 cm, indicating consistency of the back-projections.

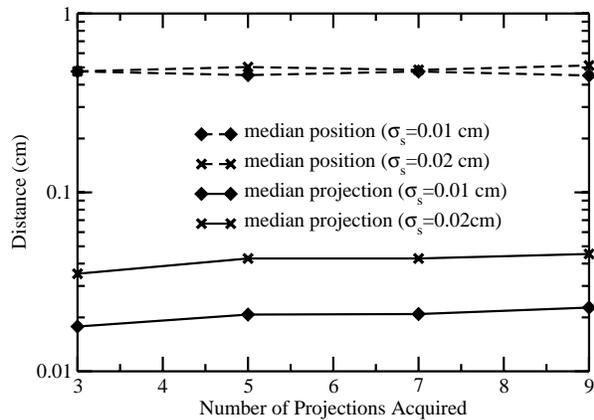


Figure 4. Simulated random errors in the measured positions of the projections, in addition to the uncertainties in the positions of the calibration objects and the position of the x-ray focus during the acquisition of each projection image. The errors in the measured positions of the shadows (σ_s) are comparable to the spacing of detector elements on commercially available devices. The consistency of the backprojections (“median projection”) are comparable to σ_s .

simulating each set of N_{proj} projections, the positions of the six point-like calibration objects were perturbed by adding to each nominal coordinate a random length drawn uniformly from the interval $[-0.4 \text{ cm}, 0.4 \text{ cm}]$, in order to simulate uncertainties in the construction of the phantom. The projections of the calibration objects were then simulated using these perturbed positions, and the projection matrices were estimated by minimizing the χ^2 of equation 21 as described in section 2.

To test these reconstructed projection matrices, fifty additional points were then generated uniformly for $x, y \in [0, 10 \text{ cm}]$, $z \in [0, 8 \text{ cm}]$. Using the estimated projection matrices and the calculated projections, the coordinates of each point in the transformed coordinate system, (x', y', z') , were estimated by minimizing

$$\chi_s^2 = \sum_{\alpha=1}^{N_{proj}} (u_\alpha - \tilde{u}_\alpha)^2 + (v_\alpha - \tilde{v}_\alpha)^2 \quad (22)$$

where (u_α, v_α) represents the shadow of one of these fifty points based upon the estimated coordinates (x', y', z') and the estimated projection matrices, while $(\tilde{u}_\alpha, \tilde{v}_\alpha)$ represents the “actual” position of the shadow based upon the simulated positions of the object and the x-ray focus.

The results of these simulations are shown in figure 3, where for each value of N_{proj} , 100 sets of projections were simulated. For each of the 100 sets of N_{proj} images, the root-mean-square (RMS) of the distances between the generated positions (x, y, z) and the estimated positions (x', y', z') was calculated for the set of 50 additional random positions. The figure shows the mean and the median of the 100 RMS “position” values for each value of N_{proj} . Additionally, for each of the 100 sets of images, the RMS of the distances between the “actual” projections (\tilde{u}, \tilde{v}) (calculated from the “true” point positions (x, y, z) and the “true” position of the x-ray focus) and the shadows of the estimated positions (u, v) (*i.e.*, (u, v) is the shadow of the estimated position (x', y', z') using the estimated projection matrices) was calculated for the set of 50 additional random positions. The figure shows the mean and median of the 100 RMS “projection” values for each value of N_{proj} . This measures the discrepancy between the “observed” projections and the projections of the estimated positions (x', y', z') . Significantly, these values are between one and three orders of magnitude lower than the errors in the positions. This indicates that, while absolute geometric positions can not be estimated with greater accuracy than allowed by knowledge of the positions of the calibration objects, it is still possible to estimate the projection matrices in such a manner as to obtain geometric consistency among the backprojections of the shadows of various objects.

In this graph both the means and the medians are quoted. There appears to be a long tail to these distributions which can provisionally be attributed to limitations of our technique for numerical optimization of χ^2 and χ_s^2 . The differences between the means and the medians generally represent one or two sets of projections for which the minimization procedure did not converge, and we feel the median values are likely to be the more representative statistic.

An additional source of uncertainty to be considered is in the estimation of the positions of the shadows on the detector. Modern digital radiography detectors are fabricated to tolerances comparable to those in the electronics industry. Thus variations in the spacing between detector elements, curvature of the detector surface, or misalignment of detector elements need not be considered. The spacing between detector elements can be $\approx 100 \mu\text{m}$, which does set a fundamental limit on how well the positions of the shadows of objects can be measured.

The effect of the uncertainty in the positions of the shadows is quantified in figure 4. In addition to the various uncertainties discussed above, each projected position was perturbed by adding to each of the two coordinates on the face of the detector a random length drawn from a Gaussian with $\sigma_s = 0.01 \text{ cm}$ or $\sigma_s = 0.02 \text{ cm}$. The errors in the absolute positions of the points are again comparable to the uncertainties in the positions of the calibration objects ($\approx 0.4 \text{ cm}$). The RMS distances between the perturbed projections and the projections of the estimated positions are on the order of 0.02 cm for $\sigma_s = 0.01 \text{ cm}$, and on the order of 0.04 cm for $\sigma_s = 0.02 \text{ cm}$. This is as expected, since the result of randomly perturbing the positions of the shadows is that, even with the freedom of the projective transform \mathcal{B} , the data will be fundamentally inconsistent. Nevertheless, for reasonable values of σ_s , the consistency of the backprojections remains on the order of twice σ_s . That is, the backprojections of points will be consistent to within margins comparable to the accuracy with which the shadows of the individual points were measured.

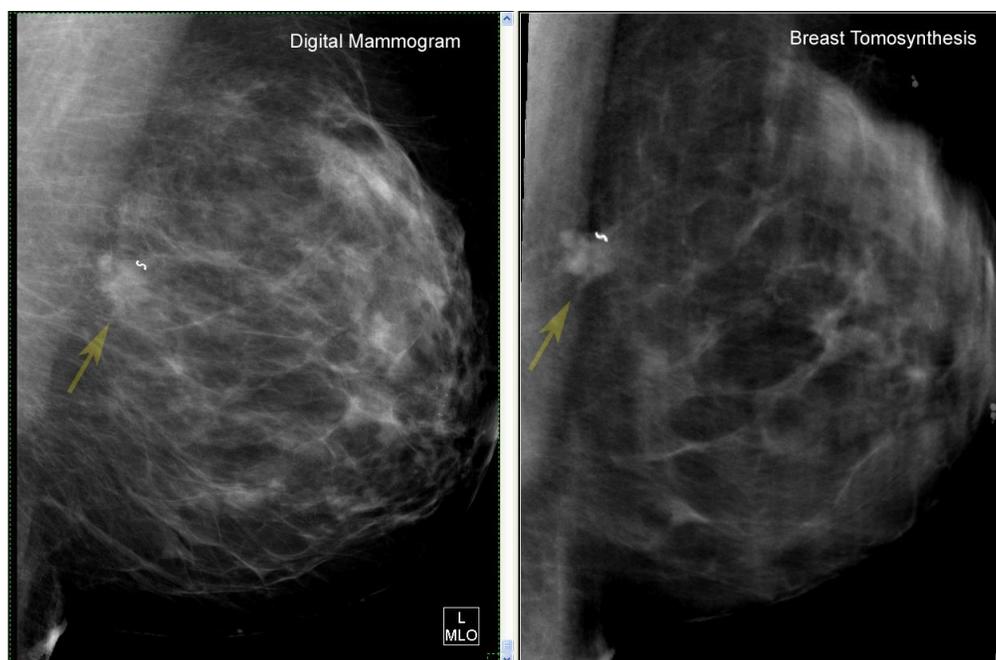


Figure 5. A single-projection radiograph (left) and a tomogram (right) acquired using a GE Senographe 2000D after geometric calibration as described.

4. DISCUSSION AND CONCLUSION

This paper describes a means of geometric calibration useful for performing synthetic tomography with digital-radiography devices of the type currently used for single-projection imaging. Indeed, this work has been part of our efforts to implement synthetic digital tomography for breast-imaging using a commercially available digital mammography device (GE Senographe 2000D, Milwaukee, WI). Figure 5 shows an example of a tomogram obtained using the calibration procedure described, along with a single-projection image for comparison. The reconstruction performed was a modification of simple backprojection, involving moderate filtering, and will be described in detail elsewhere.

The calibration technique avoids reliance upon the nominal geometry specified by the manufacturer. The phantom required is simple to construct and can be made easily from PMMA plates and ball-bearings. For our research we have used commercially available (McMaster-Carr, Atlanta, GA) aluminum and tungsten 1/32" ball-bearings. Shallow holes were drilled into the surface of the PMMA so that, once inserted, the ball-bearings were flush with the original surface. The high attenuation of tungsten proved useful in allowing the ball-bearings to be visualized despite the fact that, for extreme positions of the x-ray tube on the device currently available, some of the the ball-bearings were just beyond the edge of the collimated x-ray field.

The algorithm is tolerant of uncertainties in the construction of the phantom. After calibration using the phantom, the errors in the estimates of the positions of objects are comparable to the uncertainties in the positions of the calibration objects. More significantly, after calibration the backprojection lines will behave in a geometrically consistent manner. That is, there will be a point (x', y', z') in the transformed coordinate system which, in each projection, actually projects onto each of the shadows of the relevant object. Equivalently, when performing a reconstruction there will be a plane in which the object is actually "in focus". This last description applies most clearly to reconstruction algorithms such as simple backprojection, but similar remarks should apply to other reconstruction algorithms. Uncertainties in the positions of the calibration objects are "absorbed" into the projective transform \mathcal{B} , so that the volume to be reconstructed is deformed in a smooth manner, which preserves the relationship of collinearity, and the degree to which the volume is deformed is proportionate with the uncertainties in the construction of the calibration device. The resulting backprojections are consistent in a

manner proportionate with the accuracy with which the positions of the shadows are measured on the surface of the detector, which probably corresponds to somewhat below 100 μm on modern digital radiography devices.

We have encountered some rare problems in the convergence of the various minimizations which are required to implement the algorithm. These have not occurred clinically, and might be related to distortions of the phantom which can easily be avoided (*e.g.*, it is very easy to assure that the two groups of three calibration objects lie on planes which are nearly parallel, but this constraint was disregarded in the simulations). Clearly, for clinical use, it would be desirable if these exceptions could be eliminated, but for research purposes they are already sufficiently infrequent as to cause little concern.

The algorithm as stated treats the six calibration points asymmetrically. This seems esthetically undesirable, but our experience has been that this method is practical, as the simulations show. The method could be extended to use additional points, albeit still asymmetrically, by modifying equation 21 to be a sum over the additional points, but we have not had the need to do this, and instead use additional points on the calibration phantom as a check, independent of the calibration process.

REFERENCES

1. E. Samei, M. J. Flynn, W. R. Eyler, and E. Peterson, "The effect of local background anatomical patterns on the detection of subtle lung nodules in chest radiographs," in *Medical Imaging: Image Perception*, H. L. Kundel, ed., *Proc. SPIE* **3340**, pp. 44–54, (Bellingham, Washington), 1998.
2. F. O. Bochud, F. R. Verdun, J.-F. Valley, C. Hessler, and R. Moeckli, "Importance of anatomical noise in mammography," in *Medical Imaging: Image Perception*, H. Kundel, ed., *Proc. SPIE* **3036**, pp. 74–80, (Bellingham, Washington), 1997.
3. E. Samei, W. Eyler, and L. Baron, *Effects of Anatomical Structure on Signal Detection*, vol. 1 of *Handbook of Medical Imaging*, pp. 655–682. Bellingham, Washington, 2000.
4. O. Faugeras and Q.-T. Luong, *The Geometry of Multiple Images*, MIT Press, Cambridge, MA, 2001.
5. S. Suryanarayanan, A. Karellas, S. Vedantham, S. J. Glick, C. J. D'Orsi, S. P. Baker, and R. L. Webber, "Comparison of tomosynthesis methods used with digital mammography," *Academic Radiology* **7**, pp. 1085–1097, 2000.
6. C. E. Metz and L. E. Fencil, "Determination of three-dimensional structure in biplane radiography without prior knowledge of the relationship between the two views: Theory," *Medical Physics* **16**, pp. 45–51, 1989.
7. L. E. Fencil and C. E. Metz, "Propagation and reduction of error in three-dimensional structure determined from biplane views of unknown orientation," *Medical Physics* **17**, pp. 951–961, 1990.
8. O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, Cambridge, Massachusetts, 1999.
9. E. Artin, *Geometric Algebra*, Interscience Publishers, New York, 1957.
10. F. James, *MINUIT: Function Minimization and Error Analysis*. CERN, Geneva, version 94.1 ed., 1998.
11. B. Horn, "Projective geometry considered harmful." <http://people.csail.mit.edu/people/bkph/articles/Harmful.pdf>, 1999. Last accessed 2005-03-08.