

Enhancement of mobile C-arm cone-beam reconstruction using prior anatomical models

Ofri Sadowsky^a, Junghoon Lee^b, Edward Grant Sutter^c, Simon J. Wall^c, Jerry L. Prince^b and Russell H. Taylor^a

^aDepartment of Computer Science, The Johns Hopkins University

^bDepartment of Electrical and Computer Engineering, The Johns Hopkins University

^cInternational Center for Orthopaedic Advancement, Department of Orthopaedic Surgery, The Johns Hopkins Bayview Medical Center

ABSTRACT

We demonstrate an improvement to cone-beam tomographic imaging by using a prior anatomical model. A protocol for scanning and reconstruction has been designed and implemented for a conventional mobile C-arm: a 9 inch image-intensifier OEC-9600. Due to the narrow field of view (FOV), the reconstructed image contains strong truncation artifacts. We propose to improve the reconstructed images by fusing the observed x-ray data with computed projections of a prior 3D anatomical model, derived from a subject-specific CT or from a statistical database (atlas), and co-registered (3D/2D) to the x-rays.

The prior model contains a description of geometry and radiodensity as a tetrahedral mesh shape and density polynomials, respectively. A CT-based model can be created by segmentation, meshing and polynomial fitting of the object's CT study. The statistical atlas is created through principal component analysis (PCA) of a collection of mesh instances deformably-registered (3D/3D) to patient datasets.

The 3D/2D registration method optimizes a pixel-based similarity score (mutual information) between the observed x-rays and the prior. The transformation involves translation, rotation and shape deformation based on the atlas. After registration, the image intensities of observed and prior projections are matched and adjusted, and the two information sources are blended as inputs to a reconstruction algorithm.

We demonstrate reconstruction results of three cadaveric specimens, and the effect of fusing prior data to compensate for truncation. Further uses of hybrid reconstruction, such as compensation for the scan's limited arc length, are suggested for future research.

Keywords: CTCB, ALG, RECON, SIM

1. INTRODUCTION

Cone-beam computed tomography (CTCB or CBCT)¹ has the potential to provide intra-operative three-dimensional (3D) or volume images of patients using C-arms, which are x-ray imaging systems present in the operating room (OR). This potential, however, is not fully realized in practice, due to various limitations. Most existing CBCT systems are based on high-end C-arms, e.g. flat panels, and are permanently installed in an OR. This poses cost and mobility barriers to potential uses. Simpler C-arms, e.g. image intensifiers, are cheaper and mobile, though they typically require a special calibration. High- and low-end C-arms alike frequently limit the images that can be acquired, e.g. by a small field of view (FOV) or constraints on the scan trajectory. The reduced acquired image data leads to a degradation of reconstruction quality.

In a previous publication,² we demonstrated in *simulation* how information from a pre-operative anatomical atlas, deformably registered with a small set of observed patient projections, can augment a limited-trajectory scan to improve cone-beam reconstruction. We called the strategy of combining pre-operative and intra-operative images *hybrid reconstruction*.

This paper continues where the previous left off, demonstrating fusion of *real* x-ray images and projections of a pre-operative model that compensate for image truncation. We use a 9-inch x-ray image-intensifier (XRII)

Contact: Ofri Sadowsky, ofri@cs.jhu.edu

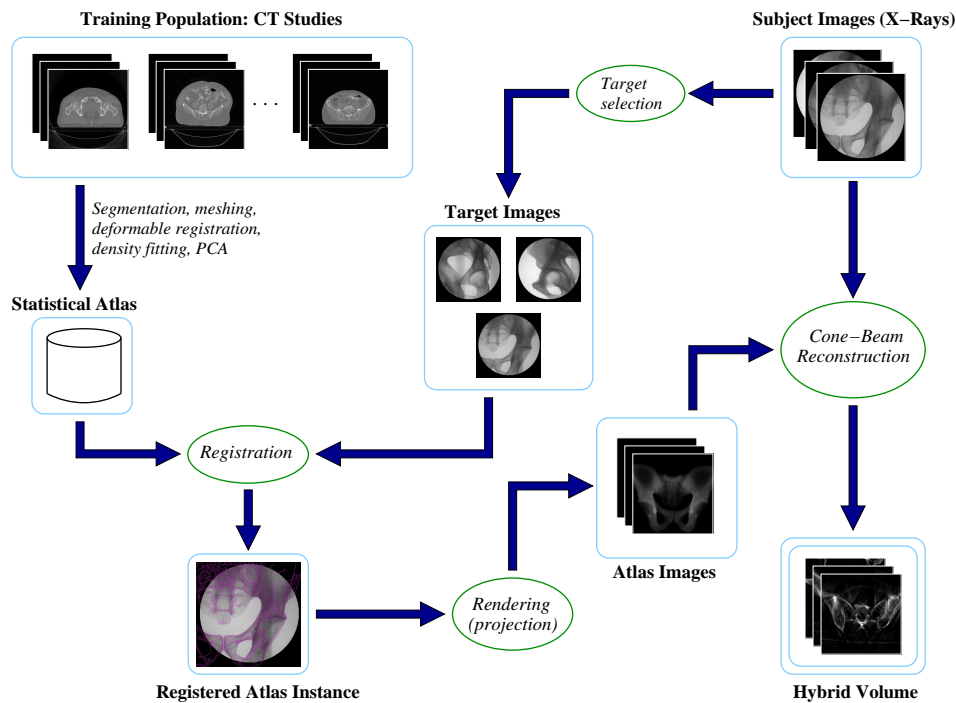


Figure 1. An overview flowchart of the hybrid reconstruction process.

OEC 9600 C-arm (c. 1993). The small FOV of this system allows for a small reconstruction volume, and causes artifacts when the object images are truncated. We show how these artifacts are reduced when prior information is injected into the reconstruction.

The new work we present, therefore, includes the development of a technique to match the pixel intensities in the two modalities (x-ray images and simulated x-rays from the prior), and the transition from simulated experiments to real-data. The design and implementation of the calibration and scan protocols with the C-arm posed a significant technical challenge, specifically as the imaging system was not designed for reconstruction. This can be counted as a secondary contribution, whose importance is in supporting intra-operative CBCT using ubiquitous and mobile imaging equipment, where so far most of the systems relied on high-end (e.g. “angiography units”) or permanently installed scanners.

2. METHODS

The elements and methods in the hybrid reconstruction process, illustrated schematically in Fig. 1, can be summarized as follows. A statistical atlas (database) of anatomical shapes is created pre-operatively by analysis of a training population of CT data-sets. Intra-operatively, a collection of x-ray images along an orbital scan is acquired with a C-arm. A subset of them are selected as targets, and co-registered with the atlas by matching pose and shape parameters. Then, new projections are computed from the atlas and combined with the observed x-rays as inputs to a reconstruction algorithm.³ The data from the atlas is inserted where the observed scan is insufficient. The result is a hybrid-reconstructed volume. Below is a closer look at the process.

2.1 Anatomical Models

Our anatomical models follow the work of Yao,⁴ which was continued by Sadowsky⁵ and others. The model description includes shape as a tetrahedral mesh, radiodensity as piecewise volumetric polynomials, and statistical modes of shape variation. We explain them briefly here.

The mesh is defined by a list of vertices given as 3-tuples, $V_0 = \{\mathbf{v}_i = (x_i, y_i, z_i) \in \mathbb{R}^3\}_{i=1}^{N_v}$, and a mesh topology given as a list of tetrahedrons, each defined by four vertex indexes: $T = \{\mathbf{t}_k = (t_{k,1}, \dots, t_{k,4}) \in \mathbb{N}^4\}_{k=1}^{N_t}$.

An initial mesh is created from a selected “template” CT dataset, on which the anatomy of interest — the pelvis or femur bones, in our case — is delineated and labeled manually. The mesh is automatically generated to fill the labeled region using the “tetsplit” utility.⁶

The density is defined locally in each tetrahedral cell using barycentric-form Bernstein polynomials as described briefly next.⁴ Barycentric coordinates parametrize the interior of a tetrahedron \mathbf{t}_k as follows. Let

$$\mathbf{T}_k = \begin{bmatrix} \mathbf{v}_{t_k,1} & \mathbf{v}_{t_k,2} & \mathbf{v}_{t_k,3} & \mathbf{v}_{t_k,4} \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

be the homogeneous vertex matrix of \mathbf{t}_k . If $\mathbf{p} = (x, y, z)^T$ is a point inside \mathbf{t}_k , its barycentric coordinates are: $\mathbf{u}(\mathbf{p}) = (u_1, u_2, u_3, u_4)^T = \mathbf{T}_k^{-1} \begin{bmatrix} \mathbf{P} \\ 1 \end{bmatrix}$. We note that the sum of the elements of \mathbf{u} , which we will write as $|\mathbf{u}|$, is equal to 1, and that the barycentric coordinates in the interior of \mathbf{t}_k are all non-negative values between 0 and 1. We define a density polynomial of order n over the elements of \mathbf{u} as

$$P_k(\mathbf{u}) = \sum_{|\mathbf{d}'|=n} \beta_{\mathbf{d}'}^k \binom{n}{\mathbf{d}'} u_1^{d_1} u_2^{d_2} u_3^{d_3} u_4^{d_4} . \quad (1)$$

Here, $\mathbf{d} = (d_1, d_2, d_3, d_4)^T$ is a four-element *power index*, $\beta_{\mathbf{d}'}^k$ is a free scalar coefficient (the index k identifies the tetrahedron), and $\binom{n}{\mathbf{d}'} = \frac{n!}{d_1!d_2!d_3!d_4!}$ is a multinomial factor. The polynomial is fitted to a sample of points taken from the template CT inside \mathbf{t}_k by minimization of square errors for the unknown vector of coefficients $(\beta_{(n,0,0,0)}^k, \beta_{(n-1,1,0,0)}^k, \dots, \beta_{(0,0,0,n)}^k)^T$. The combination of the template mesh shape and polynomial functions approximates the radiodensity of the anatomy of interest in the template CT.

After the template model is created, we compute a deformable mapping of the template CT to each member of a training population, consisting of CT scans from N_{trn} different subjects, using the Mjölknir registration software.^{7,8} The result is a population of training shapes, $\mathcal{S} = \{S_0, \dots, S_{N_{trn}}\}$, sharing the mesh topology T , and differing by their vertex lists $\{V_j\}_{j=0}^{N_{trn}}$.

The shapes are aligned by the best similarity transformations⁹ and processed using principal component analysis (PCA),¹⁰ where each V_j is expressed as a $(3N_v)$ -element long shape vector. This produces a *mean shape vector* V_{mean} and an array of *shape variation modes* $G = [\mathbf{g}_1 \dots \mathbf{g}_m]$ (where m is a user-specified number). New, unlearned patient shapes are approximated by this representation as the linear combination $V_{\text{new}} = \sigma (V_{\text{mean}} + G\boldsymbol{\lambda})$, where σ is a global scale factor, and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_m)^T$ is a vector of weights applied to the shape mode matrix.

2.2 Model Visualization and 3D-2D Deformable Registration

From the analytical approximation of radiodensity described in Sec. 2.1, we developed a closed-form line-integration formula for the polynomials in Eq. (1). This allowed us to create an efficient algorithm for computing simulated x-ray images — digitally reconstructed radiographs (DRRs) — from an instance of the atlas for any set of viewing parameters. Details of the method were published by Sadowsky et al.¹¹

During operation, a set of target x-ray images of a patient is acquired, along with the view parameters (see Section 2.3). These can be compared with the model’s DRR images, and used for co-registration with the atlas as follows.¹²

The registration parameters include translation \mathbf{t} , rotation R , isotropic scale σ , and deformation mode weights $\boldsymbol{\lambda}$, applied as

$$V_{\text{reg}}(R, \mathbf{t}, \sigma, \boldsymbol{\lambda}) = R\sigma (V_{\text{mean}} + G\boldsymbol{\lambda}) + \mathbf{t} .$$

The registration searches through the space of transformations in phases. In every phase a different subset of the parameters is chosen, and the Downhill Simplex (DS) algorithm¹³ is used as an optimizer to maximize a *mutual information*¹⁴ similarity score between the observed x-rays and the DRRs. Notice that under this notation, a special case of *rigid* registration is marked by $\sigma = 1$ and $m = 0$ (m is the number of shape modes).

In earlier simulation experiments,^{5,12} we succeeded in recovering rigid transformations to under 0.1 mm of translation error and under 0.1° of rotation error. Subsequent real data experiments produced corresponding error of about 0.7 mm and 0.7° . This establishes the accuracy of the rigid registration process.

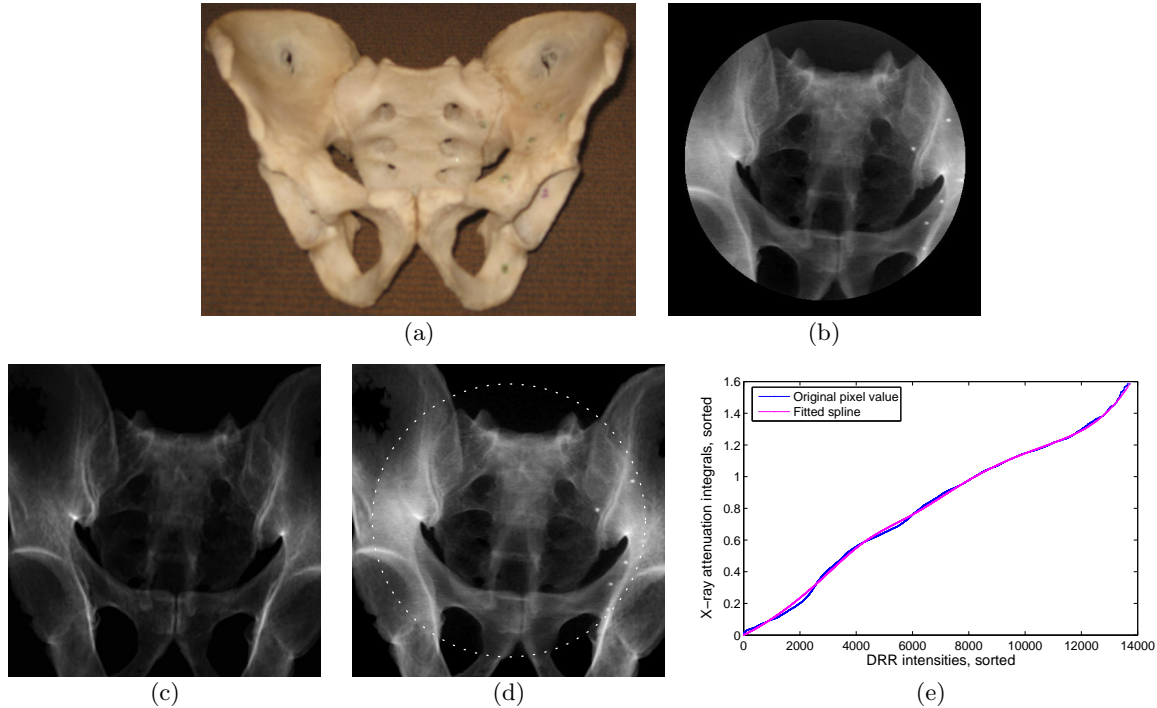


Figure 2. Registration of a tetrahedral model and an x-ray image. (a) A photograph of the target dry pelvis bone. (b) A rectified x-ray image of the target. (c) A DRR projection of a CT-based tetrahedral model of the target, registered with the x-ray. (d) Fusion of (b) and (c) used as input in hybrid reconstruction; the dotted circle indicates the boundary between observed x-ray data (inside) and prior model (outside). (e) A plot of the order-matched intensities in (b) and (c) (blue) and a spline fitted using eight control points (magenta); see Sec. 2.3.4 for more detail.

Simulation experiments were also used to measure the accuracy of deformable registration, this time as a surface distance in a leave- n -out test. We observed that the accuracy depended, among other factors, on the size of the target x-rays. With a nearly full view of the bone (256 mm detector), the mean errors were typically 2.0 mm to 2.1 mm; for truncated views of 160 mm and 128 mm, which are roughly the effective FOV of our C-arm, the mean errors increased to 2.25 mm and 2.5 mm, respectively. Comparatively, a “best case” 3D-3D deformable registration, based on accurately known shapes, produced mean errors on the order of 1.4 mm to 1.5 mm.

Similar experiments were subsequently conducted using real x-ray images of a dry pelvis specimen, shown in Fig. 2. Several x-rays of the bone in image (a) were rigidly registered with a tetrahedral model. Image (b) shows one of the x-rays, and (c) shows a DRR of the model after registration. Before the reconstruction is computed, we fuse images (b) and (c) to form image (d) (see Sec. 2.3.4). The accuracy of this rigid registration has a translation error of about 0.6 mm and a rotation error of about 0.6° relative to a fiducial-based registration.

To conclude this section, our results demonstrate a successful 3D-2D image-based registration of x-rays and a 3D mesh model using both simulation and real data. We specifically show that the registration is successful even when the FOV is relatively narrow, as in the actual C-arm images. Studies of the accuracy of deformable registration with real x-rays (as opposed to simulation) are planned as future research.

2.3 Image Acquisition, Calibration and Camera Tracking

Our x-ray images are acquired using a 9-inch XRII OEC 9600 C-arm, and recorded on a computer using a frame grabber, with an image of 640×480 8-bit pixels. This enables us to capture a continuous stream of images while the C-arm is rotated on its motorized (“L-arm” or “propeller”) axis. For every image, a calibration that includes camera pose, pin-hole projection parameters, and image rectification, is required. Newer C-arms sometimes provide this information through factory calibration, encoders, and better interfaces to external computer systems. Flat panel detectors do not even require image rectification. Our older, “legacy model”

C-arm, however, requires calibration, and so are many installed-base systems. To help readers who, like us, have legacy C-arms, in potential future research, we describe our imaging and calibration protocols next.

2.3.1 C-arm configuration

By default, many C-arms are set to display high contrast images with clearly distinguished details. In our case, however, we attempt to approximate as close as possible a quantitative measurement of the attenuation of x-ray energy in the imaged body, which may be in conflict with increasing image contrast. We therefore turn off all the functions of the C-arm that are intended to enhance contrast, and are switched on by default. In the OEC 9600 model, these include histogram equalization, automatic control of tube voltage and current, and automatic gain adjustment in the video camera. Instead, we set the voltage and current manually, and turn on the “tech lock” feature of the OEC 9600 for fixed gain.

In this configuration, we assume that the recorded 8-bit image intensity, $I(\mathbf{c})$ for a pixel \mathbf{c} , is proportional to the amount of x-ray energy that hits that pixel, and that a pixel value of 255 represents zero attenuation. Then, according to Beer’s attenuation law, the integral of attenuation for \mathbf{c} is computed as

$$\alpha(\mathbf{c}) = \int_0^l \mu(t) dt = -\ln(I(\mathbf{c})/255) , \quad (2)$$

where t is a scalar parametrization of the ray piercing \mathbf{c} , whose length is l , and μ is the linear attenuation coefficient we seek to reconstruct. This expression is simplifying reality, of course, ignoring effects such as detector response curve, beam hardening, scatter, and so on; but it is sufficient for our presentation here.

We set the C-arm to fire at pulse mode, in order to reduce motion blur in the images. The highest pulse rate is 8 per second (we have no control over pulse timing or duration). In our image capture software, we typically set the frame rate to double the C-arm’s pulse rate, so as to recover the complete input signal.

2.3.2 C-arm calibration

For reconstruction, every x-ray image must be calibrated for camera parameters, which include rectification, intrinsic projection parameters, and extrinsic C-arm pose parameters. The rectification and intrinsic parameters are measured using a calibration phantom; the extrinsic parameters are recorded using a Polaris[®] optical tracking system (NDI, Ontario, Canada).

Our custom-made calibration phantom, fabricated of aluminum, steel wire and acrylic glass plates, is shown in Fig. 3(a). The phantom contains features on three levels, or distances from the XRII surface. The most proximal is a grid pattern, used as a rectification fiducial. More distal are two levels with diamond patterns, whose intersection points are used for calibrating the pin-hole model parameters.

The phantom is mounted on the XRII for calibration, using thumb screws for fixation, and is released when actual object x-rays are acquired. It is equipped with 12 cone-shaped holes that are used as reference points in contact-based registration. A “rigid body” with reflective spheres is attached next to it, on the XRII casing, in order to track the motion of the arm using the Polaris.* We use a calibrated pointer tool (or probe) to mark the positions of these holes, then compute a rigid registration to model-based coordinates.⁹

Image rectification is performed by detecting the vertical and horizontal grid lines in the original image and using their intersections as control points. A global warp function is defined as a barycentric-form Bernstein polynomial, similar to Eq. (1) except for using three variables instead of four, including the X and Y coordinates of a pixel normalized to the $[0, 1]^2$ square, and a homogeneous variable that sets the total sum to 1. We found empirically that the fifth-order polynomials have the lowest residual error, compared with orders 1 through 6. After the coefficients of the function are determined, the pixels of a rectified image are transformed by it and sampled from the original x-ray. Typically, we produce an image of size 480×480 pixels, with a pixel size of $0.45 \times 0.45 \text{ mm}^2$. In the rectified image (Fig. 3(c)), the lines of the grid appear straight and parallel to the image axes, as opposed to curved and possibly slightly rotated lines in the original image (Fig. 3(b)).

We denote the local reference frame of the calibration phantom, which coincides with the rectified image, as F_{cp} . As shown in Fig. 3, its origin is at the center of the rectification grid, and its axes are parallel to the

*The rigid body is not shown in Fig. 3(a).

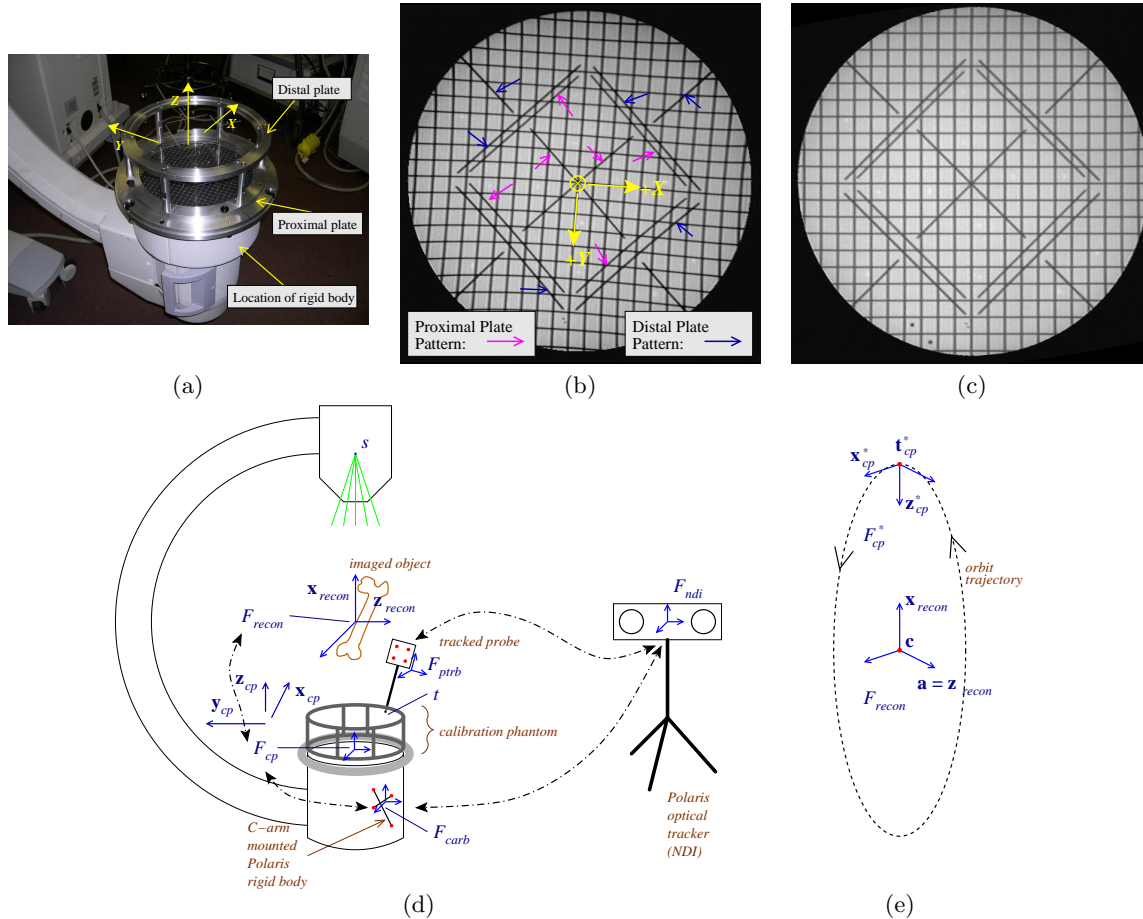


Figure 3. The C-arm calibration phantom and reference frames. (a) The phantom mounted on the XRII, with axis directions of F_{cp} shown in yellow. (b) An original, warped x-ray showing the rectification grid and the diamond patterns; the origin and axes of F_{cp} are highlighted in yellow. (c) A rectified x-ray. (d) Relations between the different reference frames in the reconstruction: curved dashed lines show the tracked frames $F_{ndi,carb}$ and $F_{ndi,ptrb}$, the rigid registration $F_{carb,cp}$, and the computed relation $F_{cp,recon}$. (e) The relation between F_{recon} and the base frame F_{cp}^* : \mathbf{z}_{recon} , the Z axis of F_{recon} , is perpendicular to the orbit plane; \mathbf{x}_{recon} is in opposite direction to \mathbf{z}_{cp}^* — the view direction of the base frame; and \mathbf{x}_{cp}^* is tangent to the trajectory.

image axes. The C-arm's *intrinsic parameters* follow the pin-hole model, and include the coordinates of the x-ray source \mathbf{s} relative to F_{cp} : $\mathbf{s}_{(cp)} = (x_s, y_s, z_s)^T$. To compute $\mathbf{s}_{(cp)}$, we detect the $\pm 45^\circ$ slanted lines of the diamond patterns in the rectified calibration images. Eighteen intersection points of these lines — nine on each level — are used to solve a linear system for the unknown $\mathbf{s}_{(cp)}$.

The *extrinsic parameters* include the position and orientation of the imaging system relative to a reference frame. In this study, we define the following frames, which are illustrated in Fig. 3(d): F_{ndi} is the reference frame of the Polaris camera; F_{carb} is the frame of the rigid body fixed on the XRII; and F_{cp} is the local frame of the calibration phantom. The transformation $F_{ndi,carb}$ includes the position and orientation of the rigid body relative to the Polaris camera. $F_{carb,cp}$ is the registration between the rigid body and the calibration phantom.[†] The pose of F_{cp} relative to F_{ndi} is obtained by the following composition:

$$F_{ndi,cp} = F_{ndi,carb} \cdot F_{carb,cp} .$$

[†]We will use the notation: $F_{cp,ndi} = F_{ndi,cp}^{-1}$, etc., for inverse transformations; a parenthesized subscript will denote a value relative to a named frame, as in $\mathbf{s}_{(cp)}$ above.

2.3.3 Interpolation of calibration

Our imaging protocol includes two parts: a calibration scan and an object scan. It is known that the image warp pattern inside the XRII depends in part on the surrounding magnetic field, and therefore changes when the C-arm is rotated. At the same time, the pin-hole intrinsic parameters change due to the mechanical flexing of the arm. Consequentially, a different calibration potentially exists for any imaging position along the C-arm's scan trajectory. The calibration scan is performed to acquire a dense sample of x-rays of the phantom in the trajectory, which subsequently provides the calibration parameters.

The calibration scan can be acquired at any time, before or after the object scan. During the object scan, which images the patient or another object to be reconstructed, the phantom is detached from the C-arm. We now face a repeatability problem: How to determine the calibration of each object image without the phantom?

Our solution, illustrated in Fig. 3(e), involves an analysis of the scan trajectory and two levels of interpolation. For the scan trajectory, we assume an orbital near-circular motion about the C-arm's mechanical axis. The axis, \mathbf{a} , is determined first by computing PCA over the list of $F_{ndi,cp}$ translation components in the scan, and choosing the direction of minimal variance. In an idealized system, any point on the axis is fixed throughout the entire scan relative to both F_{ndi} and F_{cp} . As an "iso-center" \mathbf{c} , we choose the point that minimizes the square error relative to the fixed-point constraint (i.e. it is near the axis) and to the distance from the motion plane. Expressed as $\mathbf{c}_{(ndi)}$ relative to the frame F_{ndi} , we use it as the origin of a computed *reconstruction frame*, denoted F_{recon} .

For the orientation of F_{recon} , we define the rotation axis \mathbf{a} as the frame's Z direction. The X direction is determined by the C-arm's operator, by recording a "base frame," $F_{ndi,cp}^*$, along the scan trajectory, typically where the arm is in upright position (0° rotation). We take the Z direction of $F_{ndi,cp}^*$ as the negated X direction of F_{recon} . Following this definition, we assume that the arm's motion relative to F_{recon} (which is approximately an arc in the XY plane) is repeatable, and therefore the calibration parameters can be expressed as a function of the arm's position in F_{recon} .

Both the pose parameters of $F_{ndi,carb}$ (from Polaris) and the x-ray images (from the frame grabber) are read continuously and asynchronously to a computer's memory by a multi-threaded program. Every pose instance and every image are time-stamped by a shared system clock. From the timestamps, we interpolate the C-arm's pose at the moment of image acquisition without requiring a synchronized sample that could introduce latency. Now, each image in the calibration scan and in the object scan has an associated set of extrinsic camera parameters.

We will denote the camera poses during the calibration scan as $\left\{ F_{ndi,cp}^i \right\}_{i=1}^{N_{cat}}$, and the camera poses during the object scan as $\left\{ G_{ndi,cp}^j \right\}_{j=1}^{N_{obj}}$.

To tie the reference frames of the two scans, we record two base frames: one, $F_{ndi,cp}^*$, is taken immediately next to the calibration scan; and the other, $G_{ndi,cp}^*$, next to the object scan. In both, the arm is set to the same base position. The motion trajectory of the C-arm, presumably, is repeatable not only relative to F_{recon} , but also relative to the base frame. In frame notation, we assume that $F_{recon,cp}^* = G_{recon,cp}^*$. This enables us to express the frames of the object scan relative to F_{recon} (which is determined in the calibration scan):

$$G_{recon,cp}^j = F_{recon,cp}^* \cdot G_{cp,ndi}^* \cdot G_{ndi,carb}^j \cdot F_{carb,cp}.$$

Connecting the scans this way has the advantage of tolerating a relocation of the Polaris tracker between the scans, and small relocations of the C-arm. Large relocations of the C-arm may affect the repeatability of the warp pattern, however, and therefore may add calibration errors.

Having the camera poses of the calibration scan and the object scan given relative to the same reference (F_{recon}) now allows us to define a neighborhood of calibration frames for each object frame. Denoting the translation component of $G_{recon,cp}^j$ as \mathbf{s}^j , we seek for the calibration camera poses in $\left\{ F_{recon,cp}^i \right\}$ whose translation component, \mathbf{t}^i , is in the ball $\|\mathbf{t}^i - \mathbf{s}^j\| < \delta$. A line fitted to this set of neighboring translations is used to compute interpolation weights, which we apply to the warp function coefficients and the pin-hole parameters to obtain the desired calibration of the object image at G^j .

2.3.4 Fusion of x-rays and model projections

Our ultimate goal is to inject model-based information where x-ray data is missing to improve a cone-beam reconstruction. After the anatomical model is co-registered with the x-rays (using the tracked camera pose and interpolated intrinsic parameters), we generate DRRs of the model under any imaging parameters desired. As a final step before reconstruction, the model and x-ray data must be blended.

The problem we address in this part is the following. Our anatomical model is based on a high accuracy CT, where radiodensity is given in Hounsfield units which are carefully calibrated. For the x-rays, however, the accuracy is lower: some image contrast is lost, the registration and calibration are imperfect, and our logarithmic expression of attenuation is a coarse approximation. Yet, a mapping of pixel values between the two modalities is needed in order to compute an effective blending.

A careful study to calibrate the intensity response of the C-arm, matching it against known attenuation to recover Hounsfield unit-equivalent voxels, is much desired. At this point, though, we keep it for future research. Instead, we construct an “empirical” calibration curve by matching the observed x-ray pixel intensities with corresponding intensities in DRRs of the model.

A most intuitive matching method is pixel-to-pixel, i.e. for every pixel \mathbf{c} , the x-ray image intensity $\alpha(\mathbf{c})$ (see Eq. (2)) is matched with the DRR intensity $D(\mathbf{c})$. This, however, produces a relatively high variance of the matched population of intensities; that is, every observed intensity in I is matched with a large range of intensities in D , and vice versa. Consequentially, resulting calibration curves are not very accurate.

Our approach, therefore, is to construct an *order matching* between the intensities. The underlying intuition is that even if pixel by pixel matching is not reliable, at least low-intensity values in D correspond to low intensities in I , and high intensities to high intensities. Following this, we select all the pixels in the observed viewport (i.e. the circular portion of the image where x-ray data appears) from each x-ray image and the corresponding DRR, and sort them to define the order match. An example of this is shown in Fig. 2(e) as the blue line. Then, we fit an interpolating spline function through this curve, and use it as our calibration curve. Such a spline is shown as the magenta line in the plot. An example of the images fused this way is shown in Fig. 2(d).

3. EXPERIMENTS AND RESULTS

We visually demonstrate reconstruction of three cadaveric specimens. The first is a fresh femur bone, in which bone cement has been injected. The second is a dry pelvis bone. And the third is a fresh male pelvis, denuded of most of the soft tissue, but retaining the spinal column. Each specimen represents a different setting of the reconstruction problem.

3.1 Fresh Cadaveric Femur

Our first example is reconstruction of a relatively simple object: a fresh cadaveric femur, seen in Fig. 4. The specimen’s diameter is small enough to fully fit within the C-arm’s FOV (its length is too large, but we do not attempt to reconstruct it fully), as demonstrated in image (a). We laid the bone on a wooden board around which the C-arm could be freely rotated, and acquired a dense scan covering 234° , sufficient for a complete reconstruction.³ Slices of the reconstructed volume, of size 152^3 mm³ and 320^3 voxels, are shown in images (e)–(g). For comparison, we reformatted a CT scan of the specimen, and we show corresponding slices in images (b)–(d). The basis for registration between the CT and the x-rays was steel beads used as fiducials, which appear as bright spots in image (a).

Both scans show the bed or board on which the specimen lay. The CT scan included two bones. Part of the second bone appears in image (b) at the top left. We leave it in to highlight the registration between the volumes. The original slice spacing in the CT was 3 mm, which produces a coarser resolution than our x-ray reconstruction, as can be seen in images (c) and (d).

Our reconstruction clearly shows the difference between bone cortex and the marrow in the medullary canal, and some of the internal structure of the bone. Bone cement (polymethyl methacrylate — PMMA) was injected into the specimen as part of a separate study. It appears as the bright blob in the femoral neck. By our judgment, the overall quality of reconstruction is good. It is somewhat blurry and has a somewhat lower contrast than the CT; likely causes for this are an imperfect calibration of the imaging process (covering geometric distortion,




Figure 4. Reconstruction of a fresh cadaveric femur. (a) A rectified x-ray of the specimen. (b)–(d) Transverse, “coronal” and “sagittal” slices in a CT scan of the specimen, aligned with the x-ray scan. (e)–(g) Transverse, “coronal” and “sagittal” slices in a volume reconstructed from the x-rays.




Figure 5. Tomographic reconstruction of the dry pelvis target. (a) A CT scan of the bone, registered with the x-ray scan and resampled. (b) A reconstruction from truncated x-ray images (such as Fig. 2 (b)). (c) A hybrid reconstruction from truncated x-rays fused with DRRs of registered CT-based model (such as Fig. 2 (d)).

intrinsic and extrinsic parameters, and intensity response, discussed in Sec. 2.3) and general noise in the x-ray images. Yet, the reconstruction shows fine and essentially accurate anatomical detail, comparable to the CT.

We note arc-shaped artifacts in the x-ray reconstructions, which outline the circular or spherical boundary of the scanner’s FOV. The artifacts result from the step-like behavior of the x-ray images at this boundary, as seen in image (a). They become more significant when the scanned object is large, as we show in the next sections.

3.2 Dry Pelvis Bone

Our next example, which first involves hybrid reconstruction, uses the dry pelvis specimen shown in Fig. 2. The prior model is an accurate representation of the shape and density of the examined specimen, which has been rigidly registered with x-ray images using the method of Sec. 2.2. Figure 5 shows a cross section from two cone-beam reconstructions of the specimen: image (b) was computed from the raw, truncated x-rays, and (c) was computed from the fused images. These can be visually compared with a CT sample shown in image (a).

We notice that the morphology of our reconstruction aligns well with the morphology in the CT, and that our reconstruction reflects the bone mineral density distribution as measured in the CT. Blurring, again, is

visible, as in the previous example. Importantly, the arc artifacts in image (b) are stronger than in Fig. 4. However, they mostly disappear in image (c). Image (b) also contains more visual noise and artifacts within the reconstructed region, which again are reduced in image (c). The reason for this improvement has to do with eliminating the step-like boundary of the FOV and filling the blank region with information that is *consistent* with the appearance of the specimen. This has a similar effect to having a scanner with a larger FOV and less image truncation.

3.3 Fresh Cadaveric Pelvis with Spine

The last example involves a fresh cadaveric pelvis specimen, shown in Fig. 6(a). This experiment differs from the previous in several important aspects. The specimen contains not only bone mineral, but bone marrow and some surrounding soft tissue. It also includes a portion of the spinal column. The soft tissue and the spine are not part of our anatomical model so far.

We demonstrate fusion with two types of images. One type is created as a DRR directly from a CT scan of the specimen,¹⁵ registered with the cone-beam scan through a set of implanted screw fiducials; a fusion of the modalities is shown in Fig. 6(b). The second is computed from our statistical atlas (Sec. 2.1) which includes shape and density information inside the pelvis bones. The atlas was deformably registered with x-rays of the bone, as shown in image (c). Notice that a successful registration was accomplished in spite of our incomplete model of the observed specimen. However, as image (d) shows, there are remaining discrepancies between the atlas and the specimen.

Two reconstructed slices of the specimen are shown in Fig. 7. Images (a) and (e) are sampled from the CT study; images (b) and (f) are reconstructed from the truncated images; images (c) and (g) are hybrids of x-rays and the CT; and images (d) and (h) are hybrids of the x-rays and the statistical atlas. Images (a)–(d) include blobs of PMMA that were injected into the specimen and appear as bright spots. All the images show a recess in the soft tissue on the right side, where it was resected to inject the cement.

As in the dry bone experiment, the cone-beam reconstruction recovers morphological details, such as the gap between the ilium and sacrum bones at the joint, and the spinal cavity in the sacrum. Again, the reconstruction from x-rays is blurred, compared with the CT, and arc artifacts and other types of noise are present in the reconstruction from truncated images and reduced in the hybrid reconstruction.

We notice differences between using the prior CT and the atlas to fill-in the missing pixels. The most visible difference is between images (g) and (h), where a bright artifact appears when the atlas is fused with the x-rays. Such differences are expected due to the discrepancies between the atlas and the actual specimen. Nevertheless, the two reconstructions appear very much alike inside the region covered by the XRII viewport, e.g. within the ball whose boundary is the bright arcs of image truncation. In this region, most of the information is based on the x-rays, and the pre-operative data contributes little to the reconstruction, so the differences are expected to be small. Still, either one of the prior volumes (CT or atlas) provides data that reduces the discrepancies between observations of the bone from different angles; this contributes to artifact reduction in both cases, compared with the raw reconstruction.

4. DISCUSSION AND CONCLUSIONS

We have demonstrated a calibration and imaging protocol for cone-beam reconstruction using a mobile C-arm with a relatively small field of view and a coarse estimation of the x-ray attenuation. Under these suboptimal conditions, we recover meaningful and correct details of the scanned anatomy, which are comparable, though not identical to a CT scan. Our results show that reconstruction from truncated x-rays is significantly improved by fusing the observed x-rays with synthetic images of a prior model to form a *hybrid reconstruction*. The prior model may be a CT scan or a statistical atlas. Importantly, the latter can substitute, to a degree, for a CT scan when one is not available. In other words, we accomplish improvement of the reconstruction without an exact knowledge of the anatomy, using an approximation given by the atlas.

While our demonstration used a “legacy” XRII C-arm, whose calibration is also described in this paper, the hybridization technique can assist in many cases of incomplete or distorted x-ray data, which also occur with more advanced scanners such as flat panel C-arms. Besides the image truncation problem presented here, we previously explored in simulation the examples of limited-arc scan trajectories² and reduction of camera tracking




Figure 6. Fresh cadaveric pelvis registration. (a) A photograph of the specimen. (b) Image fusion with a DRR projection of a CT volume, rigidly registered using implanted fiducials. (c) An overlay of the contours from a deformably registered statistical atlas (green pixels) on an x-ray of the specimen. (d) Image fusion with a DRR projection of the registered atlas.




Figure 7. Tomographic reconstruction of the fresh cadaveric pelvis specimen. (a), (e) Resampled cross sections of a CT volume. (b), (f) Reconstruction from truncated x-rays. (c), (g) Hybrid reconstruction using rigidly registered CT and x-rays. (d), (h) Hybrid reconstruction using deformably registered atlas.

errors.¹⁶ Other cases where a prior model can assist in reconstruction include occlusions and improvement of the estimation of attenuation. For example, the order-based spline fitting (Sec. 2.3.4) can be used in either direction, potentially to compute Hounsfield-like reconstructed voxels. We are currently evaluating this potential use.

After establishing a method to quantitatively compare the C-arm reconstruction with a “ground truth” CT, we plan to extend the initial results of this paper to the various scenarios above, and study the improvement in reconstruction for different ratios of x-ray data to prior data. Hybrid reconstruction may ultimately become a component in more widely available intra-operative volume imaging.

ACKNOWLEDGMENTS

We thank the International Center for Orthopaedic Advancement at the Johns Hopkins University Bayview Medical Center, headed by Dr. Stephen Belkoff, for hosting the cadaver experiments. The OEC 9600 C-arm was donated to us by GE Healthcare. The collection of CT studies used to build the statistical atlas was provided by Dr. Ted DeWeese and Dr. Lee Myers from the Johns Hopkins Department of Radiation Oncology. Other contributors to this research include Gouthami Chintalapani, Lotta M. Ellingsen, and Nathanael Kuo, all from the Johns Hopkins University Whiting School of Engineering. This research was supported in part by NSF ERC Grant EEC9731748, and by NIH/NIBIB research grant R21-EB003616.

REFERENCES

- [1] Feldkamp, L. A., Davis, L. C., and Kress, J. W., “Practical cone-beam algorithm,” *Journal of the Optical Society of America A (JOSA A)* **1**, 612–619 (June 1984).
- [2] Sadowsky, O., Ramamurthi, K., Ellingsen, L. M., Chintalapani, G., Prince, J. L., and Taylor, R. H., “Atlas-assisted tomography: registration of a deformable atlas to compensate for limited-angle cone-beam trajectory,” in [*IEEE International Symposium on Biomedical Imaging (ISBI)*], 1244–1247, IEEE (2006).
- [3] Noo, F., Defrise, M., Clackdoyle, R., and Kudo, H., “Image reconstruction from fan-beam projections on less than a short scan,” *Physics in Medicine and Biology* **47**, 2525–2546 (July 2002).
- [4] Yao, J., *A statistical bone density atlas and deformable medical image registration*, PhD thesis, Johns Hopkins University (2002).
- [5] Sadowsky, O., *Image Registration and Hybrid Volume Reconstruction of Bone Anatomy Using a Statistical Shape Atlas*, PhD thesis, The Johns Hopkins University (2008).
- [6] Mohamed, A. and Davatzikos, C., “An approach to 3d finite element mesh generation from segmented medical images,” in [*IEEE International Symposium on Biomedical Imaging (ISBI)*], 420–423 (2004).
- [7] Ellingsen, L. M. and Prince, J. L., “Mjolnir: Deformable image registration using feature diffusion,” in [*Proceedings of the SPIE Medical Imaging Conference*], **6144**, 329–337 (2006).
- [8] Chintalapani, G., Ellingsen, L. M., Sadowsky, O., Prince, J. L., and Taylor, R. H., “Statistical atlases of bone anatomy: Construction, iterative improvement and validation,” in [*MICCAI 2007*], Ayache, N., Ourselin, S., and Maeder, A., eds., 499–506, Springer-Verlag Berlin Heidelberg (2007).
- [9] Horn, B. K., “Closed-form solution of absolute orientation using unit quaternions,” *Journal of the Optical Society of America A* **4**(4), 629–642 (1987).
- [10] Cootes, T. and Taylor, C., “Statistical models of appearance for computer vision,” tech. rep., Imaging Science and Biomedical Engineering, University of Manchester (March 2004).
- [11] Sadowsky, O., Cohen, J. D., and Taylor, R. H., “Projected tetrahedra revisited: A barycentric formulation applied to digital radiograph reconstruction using higher-order attenuation functions,” *IEEE Transactions on Visualization and Computer Graphics (TVCG)* **12**, 461–473 (July/August 2006).
- [12] Sadowsky, O., Chintalapani, G., and Taylor, R. H., “Deformable 2d-3d registration of the pelvis with a limited field of view, using shape statistics,” in [*Medical Image Computing and Computer-Assisted Intervention – MICCAI 2007*], Ayache, N., Ourselin, S., and Maeder, A., eds., *LNCS 4792*, 519–526, Springer (2007).
- [13] Nelder, J. and Mead, R., “A simplex method for function minimization,” *Computer Journal* **7**, 308–313 (1965).
- [14] Wells, W., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R., “Multi-modal volume registration by maximization of mutual information,” *Medical Image Analysis* **1**(1), 35–51 (1996).
- [15] Muller-Merbach, J., “Simulation of x-ray projections for experimental 3d tomography,” tech. rep., Image Processing Laboratory Department of Electrical Engineering Linköping University, SE-581 83, Sweden (1996).
- [16] Sadowsky, O., Ramamurthi, K., Prince, J. L., and Taylor, R. H., “Using anatomical models and fast rendering algorithms for c-arm pose recovery and cone-beam tomographic reconstruction of bone anatomy,” in [*Computer Assisted Orthopedic Surgery (CAOS)*], Langlotz, F., Davies, B. L., and Schlenzka, D., eds., 394–397 (2005).