

# Hybrid Cone-Beam Tomographic Reconstruction: Incorporation of Prior Anatomical Models to Compensate for Missing Data

Ofri Sadowsky\*, Junghoon Lee, *Member, IEEE*, E. Grant Sutter, Simon J. Wall, Jerry L. Prince, *Fellow, IEEE*, and Russell H. Taylor, *Fellow, IEEE*

**Abstract**—We propose a method for improving the quality of cone-beam tomographic reconstruction done with a C-arm. C-arm scans frequently suffer from incomplete information due to image truncation, limited scan length, or other limitations. Our proposed “hybrid reconstruction” method injects information from a prior anatomical model, derived from a subject-specific computed tomography (CT) or from a statistical database (atlas), where the C-arm X-ray data is missing. This significantly reduces reconstruction artifacts with little loss of true information from the X-ray projections. The methods consist of constructing anatomical models, fast rendering of digitally reconstructed radiograph (DRR) projections of the models, rigid or deformable registration of the model and the X-ray images, and fusion of the DRR and X-ray projections, all prior to a conventional filtered back-projection algorithm. Our experiments, conducted with a mobile image intensifier C-arm, demonstrate visually and quantitatively the contribution of data fusion to image quality, which we assess through comparison to a “ground truth” CT. Importantly, we show that a significantly improved reconstruction can be obtained from a C-arm scan as short as 90° by complementing the observed projections with DRRs of two prior models, namely an atlas and a preoperative same-patient CT. The hybrid reconstruction principles are applicable to other types of C-arms as well.

**Index Terms**—Anatomical atlas, C-arm, computed tomography (CT), cone-beam reconstruction, hybrid reconstruction.

## I. INTRODUCTION

ONE-BEAM computed tomography (CBCT) [1] is an increasingly popular 3D medical imaging modality. One of its attractive properties is the ability to acquire the images intra-operatively using C-arms without the need for a

CT scanner. However, many such C-arm “scanners” may produce degraded results because of data loss during the imaging process. In this paper, we address two problems which may occur during a C-arm scan: image truncation and limited-arc, or reduced scans. We show that under such conditions, traditional reconstruction algorithms produce low quality images with strong artifacts.

To address such limitations, we propose to complement the observed X-ray images (or projections) with data from a prior anatomical model. An immediate candidate for the prior would be a same-patient CT volume registered with the acquired projections. But when a preoperative CT is not available, we propose to use a statistical “atlas” of the anatomy, which produces an approximation of the patient. We show that this statistics-based approximation may be sufficient for reconstructing a region of interest (ROI) that preserves the patient-specific information from the projections, yet reduces artifacts caused by the data loss.

The elements of this “hybrid” reconstruction, described in Section IV, are anatomical models, 2D-3D registration between them and X-ray images, fusion of the observed projections with computed projections of the registered model, and then reconstruction. In this paper, we discuss the techniques for registration and data fusion, and present reconstruction results from data collected with a C-arm (OEC 9600) in Section V. We study the effects of various sources of anatomical data, with varying similarity to the observed object, on the reconstruction. We also test varying proportions of blending between the X-ray images and the model to validate the extent to which the recovered volume represents the real scanned object rather than the prior model.

Importantly, our paper focuses on reconstruction from actual scans with the C-arm rather than on a theoretical analysis of the method or on simulation experiments (simulation results can be found in [2]). This goal faced specific calibration challenges, described in Section III-A, which also affect our quality assessment methods (Section III-C). We ultimately show that our reconstructions are visually similar and quantitatively correlated with actual CT reconstructions, and that the fusion of a prior model visually and quantitatively improves the reconstruction relative to the use of truncated data.

The paper is organized as follows. Section II relates this work to previous studies on data loss compensation in reconstruction. Section III presents the infrastructure, or “toolkit” used in the research: the equipment and protocols we use for C-arm

Manuscript received June 08, 2010; accepted July 11, 2010. Date of publication July 26, 2010; date of current version December 30, 2010. This work was supported in part by NSF ERC Grant EEC9731748, in part by NIH/NIBIB under Grant R21-EB003616 and Grant 5R21EB007747-02, and in part by Johns Hopkins University internal funds. *Asterisk indicates corresponding author.*

\*O. Sadowsky is with the Department of Computer Science, The Johns Hopkins University, Baltimore, MD 21218 USA.

R. H. Taylor is with the Department of Computer Science, The Johns Hopkins University, Baltimore, MD 21218 USA.

J. Lee and J. L. Prince are with the Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD 21218 USA.

E. G. Sutter and S. J. Wall are with the International Center for Orthopaedic Advancement, Department of Orthopaedic Surgery, The Johns Hopkins at Bayview Medical Center, Baltimore, MD 21224 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2010.2060491

CBCT and quality assessment. Section IV describes the key components of the hybrid reconstruction strategy: anatomical models, 2D-3D registration, and data fusion. In Section V we describe a series of experiments that evaluate the contribution of the fused data to reconstruction. In Section VI we provide a broader perspective on the experimental results. We conclude in Section VII.

## II. BACKGROUND

Cone-beam reconstruction algorithms have been studied for decades [1], and the impact of missing data, such as image truncation and reduced arcs, was recognized early on. Methods addressing data loss can be divided into the following strategies.

*Sinogram completion* methods, such as [3]–[6], typically address truncation of the projected images due to a narrow field-of-view (FOV). The truncation affects the sinusoidal path of a pixel in the reconstructed volume, and these methods attempt to recover the missing portions. Completion functions proposed were a linear continuation [3], cylinder of uniform density [4], minimal observed intensity [5], and others [6]. In comparison, our hybrid reconstruction can be thought of as completing the sinogram with a prior model which is registered with the observed data. However, we note that a true sinogram only exists on the midplane of the scan, while the hybrid reconstruction can be used for off-central slices as well.

*Total variation* (TV) methods, such as [7]–[10], focus on identifying object boundaries and fitting smooth densities inside them. Sidky *et al.* used TV with observed data only to improve limited-angle fan-beam [7] and later cone-beam [8] reconstructions. Chen *et al.* [9] used TV to address low angular sampling rate, with same-subject CT images as a reconstruction prior. Noël *et al.* [10] propose a simplification to TV based on a polygonal representation of the reconstructed objects. In comparison, this work uses conventional filtered back-projection (FBP) without assumption of interior smoothness. We do use prior models, which can be a same-patient CT or a generic atlas, and address the registration problem between the observed data and the prior.

Differentiated backprojection (DBP) applies the inverse Hilbert transform to a back-projection of derivative-filtered images to obtain the studied volume [11]–[15]. It was shown to address the reduced scan problem by exactly reconstructing the intersection of FOV and convex hull of the source trajectory [11], [13], and also a truncated ROI [12], [14] under some geometrical conditions. Our basic reconstruction method is DBP [11], [16]. However, the experiments in this paper show that without data completion, the boundary of the ROI still contains strong artifacts. These are significantly reduced by the truncation compensation in our method. Furthermore, we note that even the relatively mild condition in [11]—a scan arc of  $(180^\circ - \text{fan angle})$  is required for a non-empty exactly reconstructed ROI—is often unsatisfiable in practical C-arm imaging protocols. The hybrid reconstruction provides a way to estimate an extension even to such short arcs.

Within the DBP framework, Kudo *et al.* [15] showed that *a priori* knowledge of a small interior portion of the ROI can

provide a unique reconstruction. In comparison, in this work the ROI is regarded as an unknown region of change, whose reconstruction is assisted by prior knowledge of its exterior. This approach was demonstrated by Ramamurthi [17], who proposed to fuse the Hilbert transform of a prior model (e.g., a CT scan) of the object into the DBP framework. Left open in his work are the questions of how the prior information is obtained, how it is registered with the acquired images, and how the two sources are normalized to a meaningful dynamic range (see Section IV-C). This work continues Ramamurthi’s path, addressing these questions through statistical anatomical atlases, their registration with X-ray images, and an alternative method of fusing the data sources in the projected image space. Sadowsky *et al.* [2] demonstrated the technique in simulation experiments to compensate for reduced trajectory. Here, we extend this approach to real X-ray projections.

To summarize, this paper proposes and demonstrates a practical approach to compensate for missing X-ray data based on a prior anatomical model. Our approach differs from classical sinogram completion that relies on very little prior knowledge, and from DBP methods that focus only on an interior ROI regardless of its boundary or exterior. We regard the hybrid reconstruction as a potential complementary technique, that can be combined, for example, with sinogram completion or TV methods. We show that statistical anatomical models can be used as priors in the absence of an exact CT. And we demonstrate the method using real X-ray data, addressing registration and range normalization issues.

## III. RECONSTRUCTION TOOLKIT

In this section, we describe the reconstruction “toolkit”—basic methods of data acquisition and processing used in our experiments. They include the imaging devices, their calibration and acquisition protocols, the reconstruction algorithm and the tools for image quality assessment. These tools define the baseline for the reconstruction results, which hybrid reconstruction, described in the next section, should improve.

### A. System Calibration and Image Acquisition

The system used to acquire X-ray images is an OEC 9600 C-arm (GE OEC Medical Systems, Salt Lake City, UT, 1993) with a nominal 9-in X-ray image intensifier (XRII), which we fitted for reconstruction. We note that the methods proposed in Section IV are equally suitable for other types of C-arms, such as those incorporating a flat-panel detector.

To calibrate the acquired data, we developed a protocol involving a calibration phantom mounted on the XRII, whose X-ray images provide information on the warp pattern inside the XRII and the projection’s intrinsic (pin-hole) parameters (see Fig. 1). This type of calibration has been studied in many previous works, for example [18]–[21]. However, during the scan of a specimen, the phantom is taken off the XRII, and the parameters are estimated by lookup and interpolation of data in a separate calibration scan, using the arm’s pose as an “index.” We use an external system—the Polaris optical tracker (NDI, ON, Canada)—to acquire the extrinsic imaging parameters (camera




Fig. 1. The calibration phantom. (a) Mounted on the C-arm XR11. (b) An X-ray image: the grid is used to compute dewarp parameters, and the diamonds and diagonals for pin-hole parameters.

pose). The complete details of the calibration protocol are listed in [22].

Scans are executed by turning the C-arm on its motorized “L-arm” axis (alternatively called “U-arm”) in continuous motion. The X-ray tube is fired at eight pulses per second, and data is collected with a WinTV 2000 frame grabber (Hauppauge Digital, Inc., Hauppauge, NY). Typically, the warp-corrected (rectified) images are of size  $480 \times 480$  pixels, with a pixel size of  $0.45 \times 0.45 \text{ mm}^2$ , and 8 bits per pixel. These parameters define a nominal detector-plane FOV of 216 mm; but since the isocenter, where the scanned specimen is placed, is about 300 mm away from the detector, and the source-detector distance is about 980 mm, the effective FOV diameter around the isocenter is approximately  $0.7 \times 216 \approx 150$  mm. The actual X-ray data takes a roughly circular portion of the image, of about 440 pixels in diameter.

To provide consistent imaging conditions throughout the scan, we manually set the C-arm’s kVp and mA values, and disable the automatic contrast-enhancing features of histogram equalization and gain adjustment. Subsequently, we assume that the recorded intensity  $J(\mathbf{c})$  of a pixel  $\mathbf{c}$  is linearly related to the accumulated X-ray energy hitting the respective area of the detector, and therefore the attenuation integral for the pixel  $\mathbf{c}$  is  $I(\mathbf{c}) = -\ln(J(\mathbf{c})/255)$ . This is a simplifying assumption, which likely distorts the reconstruction to some degree (for example, we do not account for nonlinear detector responses, or for beam hardening effects). Better calibration of the detector response, as well as better detector sensitivity (as in [23]) may produce better reconstruction results overall.

## B. Reconstruction Algorithm

For CBCT volume reconstruction, we used a Feldkamp-type volume of interest (VOI) reconstruction algorithm [16] which is a variation of an FBP-type 2D fan-beam reconstruction algorithm originally proposed by Noo *et al.* [11] for ROI reconstruction. Noo’s method has merits when a full ( $360^\circ$ ) or short ( $\geq 180^\circ + \text{fan angle}$ ) scan is not achievable, which is the case for most mobile C-arms. It enables exact ROI reconstruction with very-short scan data ( $< 180^\circ + \text{fan angle}$ ) if all the lines passing through the ROI (rather than the whole object) are imaged without truncation. Generalized into a Feldkamp-type algorithm, it inherits its Feldkamp’s exact reconstruction proper-

ties. Its correctness and merits have been well verified in simulations [16]. Truncated X-ray projections, however, produce strong artifacts along the truncation edge, as shown in Section V.

The modified VOI reconstruction is formulated as follows. Let  $\Theta = [\theta_b, \theta_e]$  denote the angular support of the source trajectory and  $\theta$  be the angle of an individual source position. Let  $I(\theta, u, v)$  be the pixel value at the image point  $(u, v)$  as given in the X-ray image taken at  $\theta$ , where  $u$  is the image axis parallel to the scan plane; and let  $(u_\theta(\mathbf{x}), v_\theta(\mathbf{x}))$  be the projected coordinates of  $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$  for the view of  $\theta$ .

The voxel value  $V(\mathbf{x})$  is computed by

$$V(\mathbf{x}) = \frac{1}{2\pi} \cdot \int_{\Theta} \frac{1}{R_0 - x \cos \theta - y \sin \theta} \times w(\theta, u) I_F(\theta, u_\theta(\mathbf{x}), v_\theta(\mathbf{x})) d\theta \quad (1)$$

with

$$I_F(\theta, u, v) = \mathcal{H}_u \left[ \frac{D}{\sqrt{D^2 + u^2 + v^2}} \times \left( \frac{\partial}{\partial \theta} + \frac{D^2 + u^2 + v^2}{D} \frac{\partial}{\partial u} \right) I(\theta, u, v) \right] \quad (2)$$

where  $R_0$  is the distance between the source and the isocenter;  $D$  is the distance between the source and the detector plane; and  $\mathcal{H}_u$  is the Hilbert transform operator in the  $u$  direction.  $w(\theta, u)$  is a redundancy weighting function given by [11, eq. (40), (46)].

Note that our C-arm scanner does not perform a perfect orbital scan, and we typically approximate some of the parameters, such as  $R_0$  and  $D$ , by using the mean over all camera poses. In an exemplary scan, the standard deviation in  $D$  was on the order of about 2.0 mm with  $\text{mean}(D) \approx 976$  mm. Therefore, we expect the reconstruction errors introduced by this approximation to be small. We note also that cross sections other than the midplane are reconstructed only approximately, with errors increasing with the distance from the midplane [17].

## C. Quality Assessment

Ultimately, we wish to compare a CBCT volume with a CT scan of the same specimen. This requires a registration between the two volumes, which we describe next. We also address here the issue of incompatible units in the CT and CBCT volumes, and present a way for quantitatively comparing the volumes.

On each specimen, six to ten small, high contrast markers—metal beads or screws—were attached, to be used as fiducials. They were localized manually in the CT volume and 2D X-ray images, and their 3D positions relative to the CBCT volume were computed using the C-arm’s scan calibration. Then, a rigid registration [24] was computed between the two volumes of CT and CBCT. The CT volumes were resampled to the grid of C-arm reconstruction, and the result was regarded as “ground truth.” To exclude differences caused by the presence of different objects within the FOV of the compared volumes (such as the CT bed), we defined a region of interest (ROI) over which the numerical comparison took place.

In [2], Sadowsky *et al.* demonstrated the basic hybrid reconstruction technique in simulation, and were able to compare reconstructed data sets directly. In this paper, however, because of the simplified calibration model (see Section III-A), we could not precisely reconstruct linear attenuation coefficients (LAC) or Hounsfield units. Therefore, the quantitative comparison of the volumes comprised the following statistical measures of similarity, described next: 1) correlation coefficient, 2) mutual information (MI) [25], and 3) structural similarity index (SSIM) [26].

The correlation coefficient between two sets of samples,  $\mathbf{x} = (x_1, \dots, x_n)$  and  $\mathbf{y} = (y_1, \dots, y_n)$  is computed as

$$\text{corr}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n \frac{(x_i - \bar{x})(y_i - \bar{y})}{\sigma_x \sigma_y} \quad (3)$$

where  $\bar{x}$  and  $\bar{y}$  are the means, and  $\sigma_x, \sigma_y$  standard deviations of  $\mathbf{x}$  and  $\mathbf{y}$  respectively. Its value roughly indicates how “linear” the relation is between voxel values within the ROI; values closer to 1 indicate a stronger positive correlation.

MI indicates the ability to predict statistically the value of a voxel in one reconstruction knowing the value in the other. It is computed as

$$\text{MI}(\mathbf{x}, \mathbf{y}) = H(\mathbf{x}) + H(\mathbf{y}) - H(\mathbf{xy}) \quad (4)$$

where  $H$  is the entropy, computed for the marginal distribution of samples in  $\mathbf{x}$  and  $\mathbf{y}$  and the joint distribution of pairs  $(x_i, y_i)$ . In practice, we divide voxel intensities into  $N_{\text{bins}}$  discrete values (“bins”) and count the voxels falling in each bin to obtain an empirical distribution

$$P = \left\{ p_1, \dots, p_{N_{\text{bins}}} \mid \sum_{i=1}^{N_{\text{bins}}} p_i = 1 \right\}$$

then compute the entropy

$$H[P] = - \sum_{i=1}^{N_{\text{bins}}} p_i \ln(p_i). \quad (5)$$

Higher MI values for two images (or volumes) mean higher similarity between them.

SSIM [26] combines multiple similarity factors: luminance ( $l$ ), contrast ( $c$ ), and structure ( $s$ ), which are computed locally with a moving window and then averaged. For two image regions,  $\mathbf{x}$  and  $\mathbf{y}$ , the similarity factors are defined as

$$\begin{aligned} l(\mathbf{x}, \mathbf{y}) &= \frac{2\bar{x}\bar{y} + C_1}{\bar{x}^2 + \bar{y}^2 + C_1} \\ c(\mathbf{x}, \mathbf{y}) &= \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \\ s(\mathbf{x}, \mathbf{y}) &= \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3} \end{aligned}$$

where  $C_1, C_2$ , and  $C_3$  are constants included to avoid instability, and  $\sigma_{xy} = (1/(n-1)) \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})$ . The three components are combined as

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = l(\mathbf{x}, \mathbf{y}) \cdot c(\mathbf{x}, \mathbf{y}) \cdot s(\mathbf{x}, \mathbf{y}). \quad (6)$$

TABLE I  
NORMALIZATION PARAMETERS IN COMPARING  
CT AND CBCT RECONSTRUCTIONS

Parameter	CT	CBCT
Lower threshold	-1000	0
Upper threshold (quant.)	max	max
Upper threshold (vis.)	max	$0.4 \times \text{max}$
Additive bias	1000	None
Normalized dynamic range	[0,1]	[0,1]
Number of bins	256	256

Different upper thresholds are used in quantitative and visual comparisons.

Notice that each component, and ultimately their SSIM product, has a value between 0 and 1, with higher values indicating higher similarity, and 1 indicating equality.

Normalization of the reconstructed voxels is part of each of the similarity measures. Broadly, it may include thresholding, shifting (adding bias), scaling, and binning (coarse discretization). We summarize the normalization parameters we chose for our experiments in Table I.

#### IV. HYBRID RECONSTRUCTION METHODS

A hybrid reconstruction process includes the following steps, which are illustrated in Fig. 2. Preoperative anatomical data is collected from patients and analyzed. The outcome may be a statistical atlas or a patient-specific model. During surgery, the patient is scanned with a C-arm to acquire a set of orbital X-ray images. A subset of the images is registered (deformably or rigidly) with the pre-operative model. Then additional images can be rendered from the model to fill-in missing data in the X-ray projections. The two information sources are fused as inputs to a conventional cone-beam reconstruction algorithm. The outcome is a hybrid volume that combines the observed intra-operative data with the prior knowledge.

##### A. Anatomical Models and X-ray Simulation

Our representation of the anatomy of the scanned object, which is to be combined with the X-ray data, includes shape and radiodensity. We follow Yao [27] in using a tetrahedral mesh for the shape and barycentric-form Bernstein polynomials to approximate volumetric density. This representation has two important properties: topology-preserving deformations are relatively easy to apply, and an analytical line-integral formula for the density enables us to efficiently compute simulated projections of this deformable structure. Yao’s work includes further discussion of the tetrahedral mesh properties. The creation of the model is outlined in Fig. 3 and summarized below.

Statistical atlases are created by analysis of a training population. First, an initial shape is generated by manual labeling of a hand-picked template CT volume and tessellation of the anatomy of interest into a tetrahedral mesh [28]. Deformable registration is computed from the template CT to each training study [29]. The transformation is applied to the template mesh, and principal component analysis (PCA) is performed for the concatenated list of  $N_{\text{vert}}$  mesh vertices [30]. The outcome is a “mean shape”

$$S_{\text{mean}} = (x_1, y_1, z_1, \dots, x_{N_{\text{vert}}}, y_{N_{\text{vert}}}, z_{N_{\text{vert}}})^T \quad (7)$$

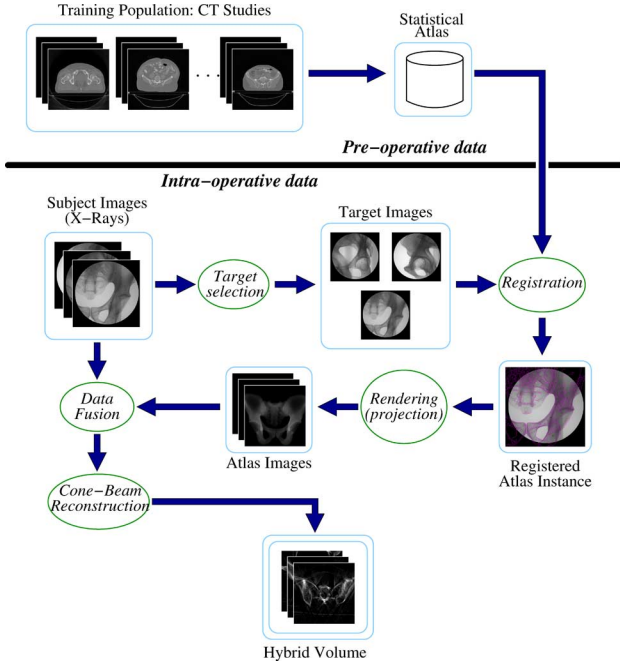


Fig. 2. Outline of the hybrid reconstruction process.

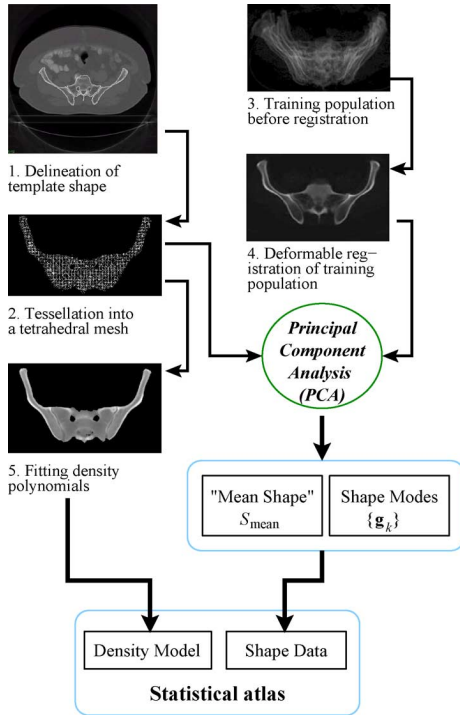


Fig. 3. An outline of the anatomical atlas creation process (images 3 and 4 courtesy of L. M. Ellingsen).

and shape modes,  $\{\mathbf{g}_k\}$ , which are displacements to be added to the mean vertex positions.

New shape instances can be synthesized as

$$S_{\text{new}} = S_{\text{mean}} + \sum_{k=1}^M \lambda_k \mathbf{g}_k \quad (8)$$

where  $M$  is a user-defined number of shape modes in use, and  $\lambda_1, \dots, \lambda_M$  are scalar weights applied to the modes. We treat the special case of a patient-specific mesh, derived directly from a CT without deformable registration, as having only a mean shape—the patient’s shape—and  $M = 0$  modes.

Within the space of each tetrahedron of the template mesh, a density polynomial is fitted to the template CT intensities. Given the vertices of a tetrahedron  $T$  as the ordered set  $\mathbf{v}_0, \dots, \mathbf{v}_3 \in \mathbb{R}^3$ , we form the matrix of homogeneous vertex coordinates

$$M_T = \begin{bmatrix} \mathbf{v}_0 & \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \\ 1 & 1 & 1 & 1 \end{bmatrix}. \quad (9)$$

Then, the barycentric coordinates of a point  $\mathbf{x} \in \mathbb{R}^3$  relative to  $T$  are

$$\mathbf{u}(\mathbf{x}) = (u_0, \dots, u_3)^\top = M_T^{-1}(\mathbf{x}, 1)^\top. \quad (10)$$

In this notation, the density polynomials take the form

$$P^n(\mathbf{u}) = \sum_{d_0 + \dots + d_3 = n} \beta_{\mathbf{d}} \binom{n}{\mathbf{d}} u_0^{d_0} u_1^{d_1} u_2^{d_2} u_3^{d_3} \quad (11)$$

where  $n$  is the degree of the polynomial;  $\mathbf{d} = (d_0, \dots, d_3)$  is a multi-index vector;  $\beta_{\mathbf{d}}$  is a “free” scalar coefficient; and  $\binom{n}{\mathbf{d}} = n! / d_0! d_1! d_2! d_3!$  is a multinomial factor. Fitting the polynomial to CT data means determining the set of free coefficients  $\{\beta_{\mathbf{d}}\}$  by which the polynomial best approximates (by least square error) a given CT volume. This is done independently for every tetrahedron in the template mesh to approximate the template CT.

We note that Yao [27] proposed to perform PCA on the density polynomials in addition to the shape statistical analysis. In contrast, we used a simple atlas whose density distribution was derived only from the template CT. Potentially, an atlas that incorporates density statistics may produce a closer estimation of an observed specimen. Yet we show here, and earlier in [2], that the coarser approximation of the atlas based on template intensities still assists in the hybrid reconstruction. The subject of density statistics is under study separately from this paper [31].

Simulated X-ray images, commonly known as digitally reconstructed radiographs (DRRs), can be computed efficiently from a shape instance and the density polynomials [32]. The rendering involves computation of integrals of the radiodensity function along the projection lines of image pixels. Encoding density as polynomials provides a closed-form equation for the integral, listed in [32]. The algorithm computes the projected silhouette of each tetrahedron, which is sent to a graphics processing unit (GPU) that computes the integral for each pixel and accumulates the result.

An example of a DRR of a dry pelvis specimen is shown in Fig. 4(a), and is comparable to an actual X-ray image seen in Fig. 4(b).

## B. 2D-3D Registration

To register a model with a C-arm scan, we first select a small subset of targets, e.g., four to six images, out of the full scan (see






Fig. 4. Simulated image and registration of model and C-arm X-ray. (a) A DRR image of a mesh model of a dry pelvis specimen. (b) An X-ray image of the specimen, to which the mesh model was rigidly registered by maximization of mutual information. (c) An overlay of the edges of image (a) on image (b). The full dynamic range of each image was normalized independently to  $[0, 1]$ .

more details below). Given the projection parameters of the images, our registration process repeatedly creates DRRs of the model and compares them with the targets. We use the Downhill Simplex (DS) optimizer [33], [34] to search through pose and shape parameters, including translation  $\mathbf{t}$ , rotation  $R$ , scale  $s$  and shape mode weights  $\{\lambda_k\}$ , to maximize mutual information (MI) [25], [35] [see (4)] between the targets and atlas projections.

Zöllei *et al.* [36] had a similar approach, but used stochastic samples to estimate the MI score between images; our fast method to compute DRRs enables us to use the full image instead. We also found empirically that the DS optimizer is fairly efficient, often better than gradient descent [36], for example.

Studholme *et al.* [37] suggested that normalized MI (NMI) as the image similarity score is more robust than MI to an initially small overlap between model and reference. Our independent tests support this hypothesis. Yet we also observed that the use of MI score can lead to a slightly higher final registration accuracy than NMI. For the purpose of this paper, the performance of both scores is roughly the same, and we use MI for 2D-3D registration.

As the registered prior model is fused with X-ray images for hybrid reconstruction, we expect that the accuracy of the registration will affect the quality of hybrid reconstruction. The accuracy of deformable 2D-3D registration between X-ray projections and a statistical anatomical atlas was measured in simulation and real images by Sadowsky *et al.* [38]. Additional studies were described in [22]. We summarize the results of these experiments below.

Simulation experiments of rigid registration show very small final translation and rotation errors, typically 0.1 mm and  $0.1^\circ$ , respectively, or less, compared with ground truth. Experiments with real X-ray images had translation differences of 0.6–1.3 mm and rotation differences of  $0.5^\circ$  to  $1.1^\circ$  between MI-based and fiducial-based registrations of a dry pelvis specimen. Both differences are roughly within the boundaries of the calibration errors of our imaging and tracking systems. For a visual example, the alignment of the model shown in Fig. 4(a) relative to the X-ray image in Fig. 4(b) was accomplished by such a rigid registration.

Deformable registration of a statistical atlas and simulated X-ray images was studied in [38]. Eleven test models of whole pelvis bones, segmented from CT volumes, were left out of the training population during atlas creation. The registration error




Fig. 5. Distribution of mean errors in deformable 2D-3D registration on the surface of a pelvis model.

was measured as distances between the outer surface vertices of the registered atlas shape and their nearest neighbors on the tested surfaces. Typical mean errors were 2.0–2.5 mm. The errors increased when narrower views of the test dataset were used as targets.

The study also observed that the error is distributed unevenly on the surface of the pelvis bone, with extremity points, such as the crest of the ilium, the superior iliac spine or the ischium having noticeably larger errors than broad surfaces (see Fig. 5). One may expect registration errors to be greater in regions of high curvature, given the averaging process in atlas creation which can erode sharp curves. It is also possible that properties of the deformable registration, used to create the training instances, affect errors in specific locations. This issue is still being studied.

Our experiments [22] show that rigid registration can succeed with a single anterior–posterior (AP) image of a pelvic specimen, under ideal (simulated) imaging conditions. However, for deformable registration additional views, e.g., lateral (LT) and  $45^\circ$  oblique (OB), contribute much to the registration quality. Typically, an LT view of the sacrum serves as an excellent guide for registration. When the imaging FOV is truncated, as with many C-arms, “out of plane” (OOP) X-ray images, involving camera motions outside a single scan’s orbit, effectively increase the visible portion of the specimen, improving the registration. The registration experiments in [38] typically used four to six target X-ray image out of the  $\{AP, LT, 45^\circ OB\}$  and OOP set.

As we note above (Section III-C), various normalization and binning operations are applied to the MI input images. It is known (cf. [39], [40]) that the choice of such parameters may affect the registration outcome. In the experiments referred above, the normalization parameters are the same as in Table I, i.e., we divide the full dynamic range of the compared images into 256 bins. Given the final registration accuracy, our parameter choice seems adequate.

### C. Image Fusion

For fusion of prior model information with observed X-ray projections, the units of image intensity in the two input sources must be compatible. For example, if we consider fusion in the Hilbert transform space [17], then the preoperative and CBCT volumes must both be given in equivalent units. That is, the original X-ray image data must be given as line integrals on a similar scale as the prior volume. This is easily achievable in simulation experiments by using Hounsfield unit-based attenuation coefficients.




Fig. 6. Fusion of a C-arm X-ray image (inside circle) with a DRR projection of a prior model (periphery). (a) Linear scaling of intensities in each modality to the range  $[0, 1]$ . (b) The DRR image intensities remapped by spline to the dynamic range of X-ray image. (c) Fusion of the X-ray image and the remapped DRR. (d) Fusion done by remapping the X-ray image intensities to the dynamic range of the DRR.

In practice, however, we are not always able to measure exact attenuation integrals. As noted in Section III, the C-arm only provides a coarse, hypothesized approximation of the attenuation, subject to many measurement and calibration errors. Data fusion under these conditions is particularly challenging.

An example of the differences between a computed DRR and a C-arm image is shown in Fig. 4, where the DRR in image (a) and the X-ray projection in (b) have been individually normalized by linear scale to the range of  $[0, 1]$ . In the DRR, only a few regions have the highest brightness, and most of the bone is relatively dark. In comparison, the X-ray image is bright in more regions. These differences are more noticeable in Fig. 6(a), which attempts, ineffectively, to fuse the normalized images.

The problem is further aggravated when the X-ray tube voltages in the scans differ. In the examples studied here, typical C-arm scans were done at about 45–60 kVp (which produced fair image quality with isolated specimens), while the CT scanners use, typically, the 100–120 kVp range. The attenuation profiles vary significantly between these different energy levels, and a calibration of the XRII response must consider this factor for volumetric fusion to work. To address this issue, and bypass (to a degree) the need for a complete calibration of the XRII, we fuse the prior and observed data as projected images, which is easier than fusing volumes.

When the model is registered (rigidly or deformably) with the C-arm images, there is, presumably, a strong dependence between pixel intensities in the DRRs and intensities in the X-ray images. A similar idea guides the multimodal mutual information-based registration algorithms [25]. In the case of X-ray images and DRRs, we can safely assume that the mapping between corresponding pixels is continuous and monotonic, up to (possibly spurious) registration errors and measurement noise.

Following this, we estimate the mapping by sorting the pixels of each modality independently and matching pixels by their sorted position. Notice that the direct spatial correspondence between intensity-sorted pixels is lost, but the effect of spurious errors is reduced. Then, we choose a small number (8–12) of corresponding representatives from the sorted lists, and fit a spline through them to approximate the hypothesized monotonic non-linear mapping. A graphic example of this process is shown in Fig. 7, and the outcome of the mapping in Fig. 6(b). The mapping can be applied either from DRR to X-ray image intensities, or in the opposite direction. Fig. 6(c) shows a fusion of an X-ray image and a remapped DRR, and Fig. 6(d)—a remapped X-ray




Fig. 7. Matching graph of sorted DRR intensities (on an arbitrary scale) and sorted X-ray image intensities (as normalized logarithm). The control points for spline mapping between them are highlighted in circles.

image and a DRR. We tested the effect of both mapping directions in reconstruction experiments.

This method somewhat resembles intensity standardization (for example, [41]) in setting the objective of bringing co-registered images to comparable intensities. But our solution is simpler because we are dealing with a narrower problem domain. We are not aware of other works that use a similar technique in the context of completing missing data in X-ray based imaging modalities.

If the missing data is a portion of each image, as with truncated images, then intensity mapping can be computed for each image independently. That is, a new spline is defined for each acquired X-ray image, and the mapping functions may differ for different views of the specimen. This makes the intensity mapping for the other type of information loss we study—the case of reduced scan—more complex, as it may not be uniform over the entire scan.

In this case, illustrated in Fig. 8, we choose an easier solution: we assume that the DRR intensities are “ground truth” and remap the observed X-ray image intensities to that range. Extrapolation of the intensity-mapping function is not required this way, but it may affect the reconstruction quality.

In addition, for reduced C-arm scans, we must extrapolate the camera trajectory to the portions missing from a short scan. This is done by fitting a circle to the C-arm positions recorded in the scan, and selecting new positions at regular intervals on it. Compensating DRR images are computed for these new positions.

## V. EXPERIMENTS AND RESULTS

In the experiments below, we scanned and reconstructed several *ex vivo* isolated bone specimens: a femur, a dry pelvis,

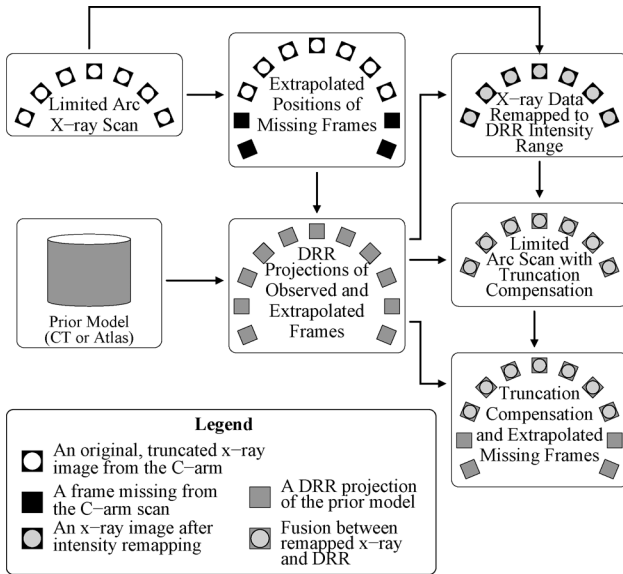


Fig. 8. Flowchart of the process for compensating limited-arc CBCT scans by “extrapolation” based on a prior anatomical model.

TABLE II  
STATISTICAL SIMILARITY MEASURES BETWEEN CBCT  
AND CO-REGISTERED CT VOLUMES

Each CBCT volume was reconstructed twice: first from the “raw” x-ray images, after taking logarithm, and second by using spline mapping of x-ray image intensities to DRR intensities. The similarity scores include: correlation coefficient (Corr), mutual information (MI) and structural similarity index (SSIM).

Specimen	Remap x-ray	Corr	MI	SSIM
Femur, truncated x-ray images	–	0.846	0.276	0.236
	+	0.792	0.206	0.223
Dry pelvis, truncated x-ray images	–	0.656	0.156	0.524
	+	0.660	0.100	0.575
Dry pelvis, truncation compensation by <i>exact model</i> projections	–	0.696	0.172	0.680
	+	0.659	0.136	0.722
Fresh pelvis, truncated x-ray images	–	0.761	0.434	0.482
	+	0.688	0.410	0.448
Fresh pelvis, truncation compensation by <i>CT</i> projections	–	0.827	0.537	0.501
	+	0.822	0.520	0.488
Fresh pelvis, truncation compensation by <i>deformable atlas</i> projections	–	0.786	0.496	0.522

and a fresh pelvis with some soft tissue and spinal vertebrae. A volume of  $152 \times 152 \times 152 \text{ mm}^3$  and  $320 \times 320 \times 320$  voxels was reconstructed from each. Quantitative similarity scores for the truncation-compensated reconstructions are summarized in Table II.

#### A. Femur Bone Specimen

This is a specimen of the left femur from a male human cadaver, into which bone cement (polymethyl methacrylate, or PMMA, mixed with  $\text{BaSO}_4$  as a contrast agent) was injected. While this specimen does not involve hybrid reconstruction, we used it as a baseline for reconstruction quality in the subsequent experiments.

CT and C-arm scans of the specimen and its right counterpart were acquired before and after the injection. Here,




Fig. 9. Transverse cross sectional reconstructions of femur specimen, with a rectangular ROI highlighted. (a) CT volume, resampled in registration with CBCT. (b) CBCT reconstruction using observed X-ray images. (c) CBCT reconstruction from a spline remapping of X-ray images to match CT projection intensities.

we focus on the postinjection reconstruction, of which cross sections are shown in Figs. 9 and 10. The figures include the rectangular ROI over which similarity measures were computed, selected to ignore other objects in the images, such as the right femur or the wooden board. The voxel size of the CT was  $0.503 \times 0.503 \times 3.0 \text{ mm}^3$ , which explains the coarse appearance of the images in Fig. 10.

Visually, the CBCT provides relatively fine detail of the scanned object. The boundaries of the bone cement, for example, are comparable in both CT and CBCT. Inner structures, such as the boundary between cortical and inner bone, are visible even where the cortex is relatively thin.

Visual comparison between the “raw” and “remapped” reconstructions [Fig. 9(b) and (c)] is subjective, i.e., differs between observers, and depends on the normalization applied to the image. The raw reconstruction seems to produce finer detail, while the remapped reconstruction may have less background noise. Quantitatively, a high correlation of about 0.8 exists between either CBCT ROI and the CT, although the raw reconstruction has a slightly higher score: 0.85 versus 0.79. A similar trend is found with the MI scores (0.28 versus 0.21) and the SSIM (0.24 versus 0.22). This argues in favor of using the raw X-ray images as reconstruction input, rather than attempting remap them to “DRR integrals.”

#### B. Dry Pelvis Specimen

The dry pelvis bone is shown in Fig. 11(a). Its X-ray images, which are truncated, are shown in Fig. 4(b). A cross section from its CT scan, whose voxel size was  $0.512 \times 0.512 \times 0.5 \text{ mm}$ , is shown in Fig. 11(b). The cortex is identifiable, and the pores of the spongy bone structure are filled with air. A C-arm scan with 280 projections covering  $232.5^\circ$  was acquired, and five of its images were selected as 3D-2D registration targets. They were spaced about  $50^\circ$ – $55^\circ$  apart from each other, and showed important anatomical features such as the lateral view of the sacrum, the frontal view of the obturator foramen, and the acetabulum.

We had limited success in registering the specimen with a deformable statistical atlas. General alignment was accomplished, but accurate scale and shape recovery failed. This may be due to various differences of structure between it and the live models, some of which are examined in Section VI.

Our image-based rigid registration with a tetrahedral model derived from the CT was successful, on the other hand. For






Fig. 10. Orthogonal cross sections of femur reconstruction. (a) “Coronal” slice from CT. (b) Coronal slice from CBCT. (c) “Sagittal” slice from CT. (d) Sagittal slice from CBCT.




Fig. 11. Transverse cross sectional midplane reconstructions of dry pelvis specimen. (a) A photograph of the specimen. (b) CT volume, resampled in registration with CBCT. (c) CBCT reconstruction using truncated X-ray images, such as the ones in Fig. 4(b); the ROI for quantitative comparison is highlighted. (d) CBCT reconstruction from a spline remapping of X-ray image intensities to match DRR intensities. (e) CBCT of observed X-ray images fused with projections of a CT-based model. (f) CBCT using X-ray images remapped to DRR intensities and fused with CT-based model projections.

nine metal-bead fiducials manually identified in the CT and X-ray images, the target registration error (TRE) between the MI-based and fiducial-based registrations had a mean of 1.38 mm and maximum of 2.14 mm (these numbers include tracking errors and fiducial localization errors). Following the image registration, we computed DRR images of the mesh, and fused them with the X-ray images, as shown in Fig. 6.

Fig. 11 shows also four midplane CBCT reconstructions. Images (c) and (d) show reconstructions from truncated X-ray images, where in image (c) we used the raw data, and in (d) we remapped it to the range of DRR intensities. The external shape and inner structures can be discerned in both. Notice the strong arc-shaped artifacts, caused by the loss of information in the truncated FOV. We defined the ROI for quantitative comparison, whose boundary is highlighted in image (c), as a truncated sphere within the “fully visible” portion of the volume. Some reconstruction artifacts are present even within this ROI. Image (d) contains fewer artifacts in the external air portion of the ROI, while the raw reconstruction shows better details of the inner structures. When fused X-ray and DRR images are used as reconstruction input, the artifacts are reduced, as seen in images (e) and (f).

Quantitative comparison demonstrates a similar trend. With the raw X-ray images, the correlation coefficient rises from 0.66 to 0.69, the MI from 0.16 to 0.17, and the SSIM index from 0.52 to 0.68. With the intensities remapped to DRR ranges, the scores are 0.66 for correlation, 0.10–0.14 for MI, and 0.58–0.72 for SSIM index. This demonstrates that artifact reduction takes place not only outside the FOV, but inside as well.

### C. Fresh Pelvis Specimen

The third specimen, shown in Fig. 12(a), was a fresh pelvis, which had been stripped of most of the soft tissue and separated from the femurs, though a part of the spinal column re-




Fig. 12. Images of the fresh pelvis specimen. (a) A photograph of the specimen. (b) An X-ray image, before taking logarithm, highlighting the fiducial screws (magenta circles) and the cement injected near the ilio-sacral joint (yellow outline). (c) An overlay of green edges from an atlas, deformably registered with the X-ray scan. (d) Fusion of X-ray image (at center, after taking logarithm) and a DRR computed from a CT scan of the specimen (periphery). (e) A similar fusion of the X-ray image and the registered atlas; notice the slight misalignment of the iliac crest and the absence of the spine from the atlas.

mained attached. Eight small stainless steel screws were driven into the bone as registration fiducials. Their heads were manually located in CT scans and X-ray images of the specimen to define a “ground truth” registration. In addition, PMMA was injected near the right ilio-sacral joint, simulating a sacroplasty. A rectified X-ray of the specimen, including these elements, is shown in Fig. 12(b).

This specimen poses more challenges than image truncation alone. Fig. 12(b) includes many white pixels, where the detector was saturated with X-ray photons, and the anatomy cannot be

identified. This demonstrates a practical difficulty to find a uniform tube setting that provides good contrast from all view directions. The anterior-posterior X-ray images included white (“saturated”) pixels, while lateral images of the same specimen, at the same tube setting, included black (“occluded”) pixels. In such cases, increasing or decreasing the kVp does not prevent signal loss.

In addition, the presence of some soft tissue and spinal vertebrae is a challenge to our atlas-based registration, since neither is modeled in the atlas. Furthermore, the fused reconstruction of the atlas with the X-ray images after registration is expected to be relatively degraded due to the same reason. Future atlases may include soft tissue models and multiple bones to address this. We discuss this further in Section VI.

To register the X-ray images and the atlas, we acquired additional static X-ray images from various directions, including some which were not along the reconstruction scan’s trajectory, i.e., “out of plane” (OOP) X-ray projections. Six of them were selected as registration targets: covering an arc of about  $179^\circ$  from one lateral view of the sacrum to the opposite lateral view, with  $\sim 45^\circ$  oblique views and two frontal views at different planes. Based on [38], we expect OOP views to provide additional shape information and thus improve the registration.

A visual example of the results of a deformable registration of our atlas to the X-ray images is shown in Fig. 12(c). We see that the overall shape was captured by the atlas, but misalignments, seen more clearly in image (e), still exist. As an error metric, we computed distances between the boundary vertices of the atlas mesh and a manually-segmented surface shape of the bone only in the specimen. The mean error was 3.95 mm, and the maximum was 17.78 mm. This is higher than earlier results, such as [38], where errors were on the order of 2.1 mm on average. But considering the confounding factors described above—signal loss and unmodeled tissue—this is arguably a good registration.

Fusion of the X-ray images with prior data highlights more challenges. Fig. 12(d) shows the X-ray image fused with a DRR made from a CT of the specimen. The fiducial-based alignment of the two is good. But a difference can be seen in the “background,” where the CT bed is projected as gray pixels, compared to mostly black pixels in the background of the X-ray image. Image (e) shows a fusion with the atlas. Here we see some misalignment of features, a difference in the intensities, and, conspicuously, the absence of the spine from the atlas.

Reconstruction was computed from the original, truncated X-ray images, from a fusion of X-ray images with CT projections, and from a fusion of X-ray images with atlas projection as before. Exemplary midplane slices from the original CT (re-sampled) and the different CBCT volumes are shown in Fig. 13. Image (a), the CT, clearly shows the difference between materials such as soft tissue, bone, and bone cement. Importantly, the different materials are visible in all the reconstructions, including inner structures such as the vertebral foramen or the gap between bones in the ilio-sacral joint. Two blobs of bone cement are distinctly visible, roughly with the same shape. A resection of the soft tissue was made in order to inject the bone cement, and is also visible on the posterior side of the specimen.

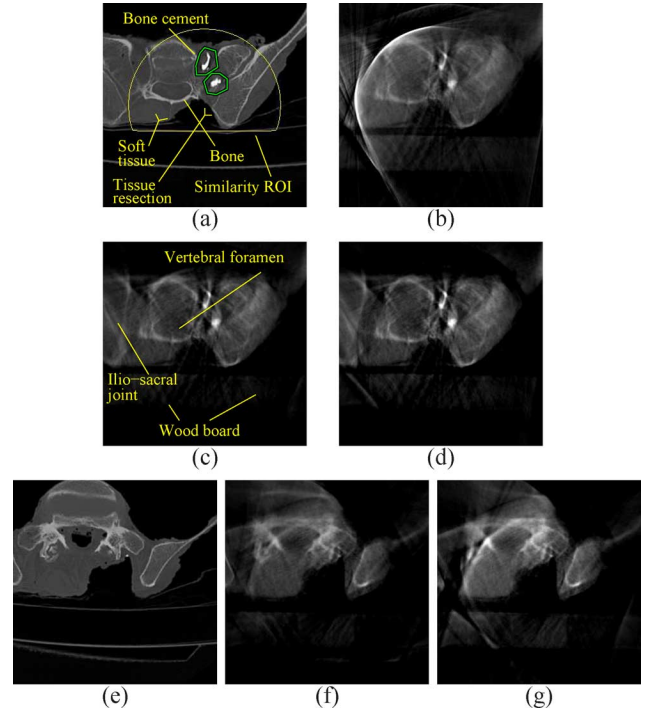


Fig. 13. Reconstruction of fresh pelvis specimen: the midplane slice. (a) Resampled CT volume, with the similarity ROI highlighted. (b) Reconstruction from truncated X-ray images. (c) Reconstruction from X-ray images fused with CT projections. (d) Reconstruction from X-ray images fused with atlas projections. (e)–(g) Off-midplane slices from the CT, CT-fused CBCT, and atlas-fused CBCT; notice the stronger artifacts in image (g), compared with (f).

Furthermore, the coarse approximation of the anatomy by the deformable atlas still improves the reconstructed volume. The artifacts were visibly reduced [compare Fig. 13(b) and (d)], although off the midplane, discrepancies between the CT and the atlas are noticeable, seen in Fig. 13(e)–(g). The image quality metrics increased from (0.76, 0.43, 0.48) in the truncated reconstruction to (0.83, 0.54, 0.50) when the CT was fused in, and to (0.78, 0.49, 0.52)—a more modest increase—when the atlas was used.

#### D. Reduced Arc Compensation

Our final example deals with a common limitation of C-arm cone-beam scans, which is independent of detector size or calibration issues. Many times, acquiring a “short scan” that covers an arc of  $(180^\circ + \text{fan angle})$  is impossible due to various physical constraints on the rotation of the arm. For instance, in the OEC 9600, the opening of the “C” is narrower than a typical surgical table, and this prevents its rotation about the patient using the motorized axis.

The hybrid reconstruction strategy may be especially useful in such cases. This time, in addition to truncation compensation in individual X-ray images, we also complete missing views by extrapolation of the scan trajectory with DRR projections of the prior model, as illustrated in Fig. 8. We demonstrate the process on the fresh pelvis specimen.

The experiment was conducted as follows. We trimmed the data of the original scan, covering over  $230^\circ$ , to narrower angles:  $90^\circ$ – $210^\circ$  in increments of  $20^\circ$  (compared with about  $193^\circ$

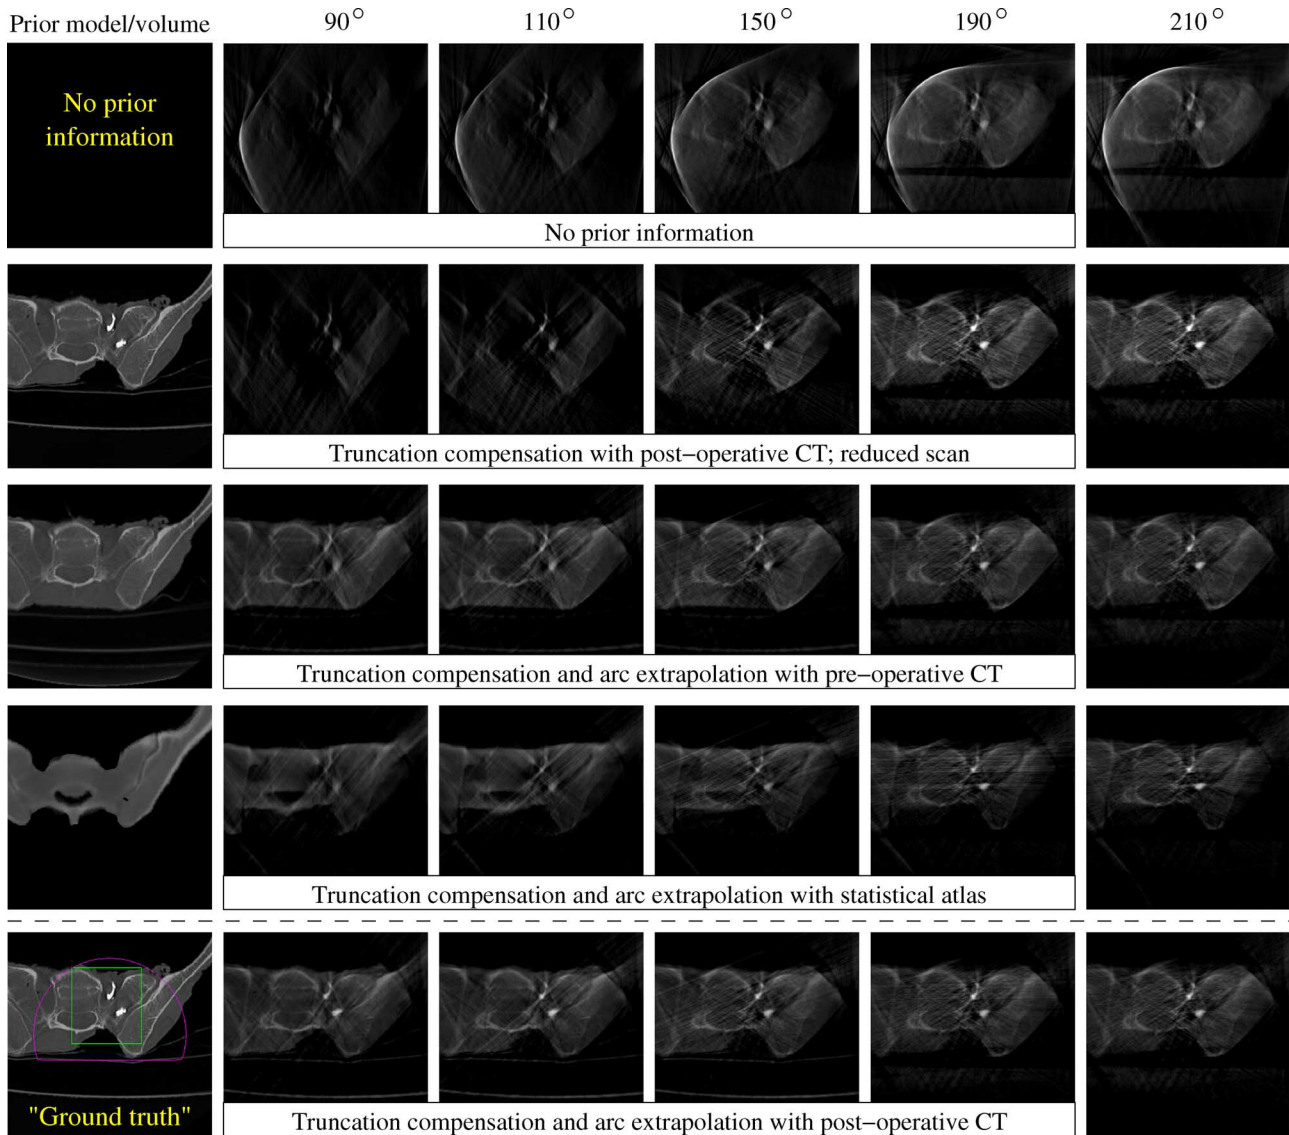


Fig. 14. Reconstruction of the fresh pelvis specimen with a reduced arc trajectory, and compensation by a preoperative scan, an atlas, and a postoperative CT scan, all registered to the original  $210^\circ$  C-arm scan. Two regions of interest, a magenta semicircle and a green rectangle, are highlighted on a slice from the “ground truth” postoperative CT. The semicircle is the “complete” FOV, and the rectangle is the volumetric bounding box of the injected cement. The dynamic range of the images was normalized based on the parameters in Table I.

required for a short scan). A circular trajectory was fitted to the original, slightly irregular C-arm trajectory. DRR images were computed at uniform intervals for this trajectory from three datasets: 1) a preoperative CT scan of the specimen, 2) a postoperative scan (which includes the bone cement), and 3) the statistical atlas. The CT volumes were rigidly registered to the C-arm scan by the fiducials, and the atlas by MI maximization as described in Section IV-A. The DRRs, with their respective projection matrices, substituted for the views that were trimmed off the original trajectory. Examples of slices reconstructed without the extrapolation and with extrapolation based on the three sources is shown in Fig. 14.

The first row shows the reconstruction based only on the truncated X-ray images, with no prior information. Arc-shaped truncation artifacts are visible. In reconstructions from scans shorter than  $190^\circ$ , the anatomical detail is degraded, although the high-contrast cement is visible even in the  $90^\circ$  reconstruction.

In the second row, the postoperative CT is used to compensate for truncation, but not for the reduced scan. The arc artifacts are reduced, and the overall image quality is better. Yet, in the shorter scans little but the bone cement can be discerned. The  $210^\circ$  scan, while still noisy, includes essentially the complete structure of the specimen.

Notice, in the first and second rows, that the shortest scans, of  $90^\circ$ – $110^\circ$ , were sufficient only to recover the high-contrast elements. The intensity window’s top was set at 0.4 of the maximal intensity, yet only the cement can be discerned. Adjusting the intensity window cannot add the missing information.

In the third row, the X-ray images are fused with a preoperative CT scan, which, presumably, is the most accurate prior available during surgery, yet does not include the bone cement. We notice that the cement is still reconstructed from the  $90^\circ$  arc, with a somewhat decreased contrast. Likewise, the resected soft tissue region is blurred where a reduced C-arm arc was used,



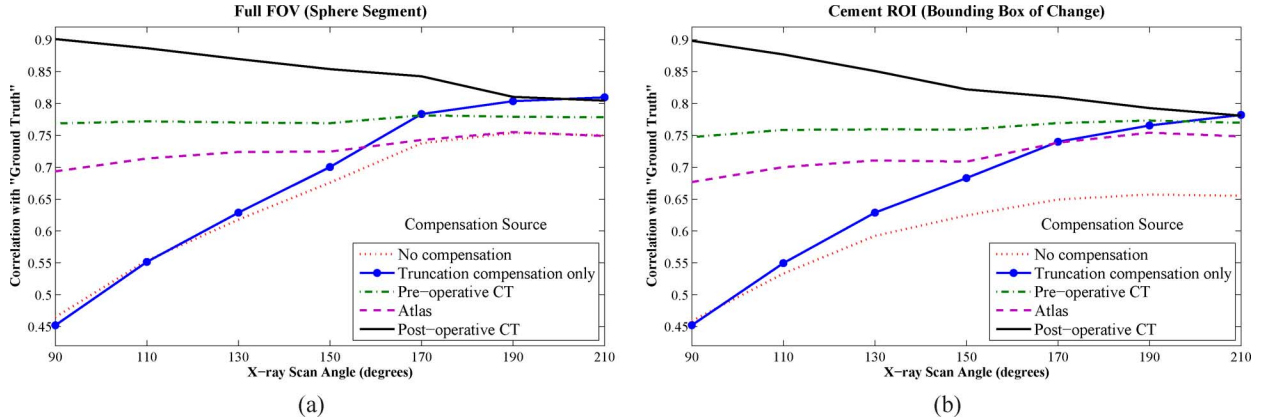


Fig. 15. Behavior of the correlation coefficient similarity score between the “ground truth” postoperative CT and a number of CBCT reconstructions, involving different sources of prior data: raw X-ray images without a prior, raw data compensated for truncation only, and compensation for reduced arc length using preoperative CT, statistical atlas and postoperative CT. Notice the contribution of the prior CT and the atlas to the reconstruction from the reduced scan arc of  $90^\circ$ . The regions of interest are (a) the full FOV, and (b) a box containing the injected cement. These are highlighted in Fig. 14.

and appears sharper as the angle increases. This is because the resection was made during the operation on the specimen, where preoperatively the tissue was whole. We can identify these elements of relatively high contrast even with the limited information of a  $90^\circ$  scan.

The fourth row represents the scenario where a preoperative CT is unavailable, and a statistical atlas is used to complete missing data. The alignment of the atlas and the CBCT volume is fair—the vertebral foramen, for example, aligns well in both volumes. We reconstructed the bone cement from the  $90^\circ$  scan, though not as sharply as when the preoperative CT was used. Since the atlas is a coarse approximation of the specimen, we see a transition between the reconstruction from a  $90^\circ$  X-ray scan (combined with  $120^\circ$  of DRR projections), and the  $210^\circ$  X-ray scan that uses only truncation compensation. The high dependency on the atlas to compensate the  $90^\circ$  scan creates a reconstruction that resembles the atlas more than the corresponding hybrid that used the preoperative CT.

In the last row, the postoperative CT was used to compensate for both truncation and limited scan. This represents the theoretical “best case,” where an exact model of the specimen is available preoperatively, though such a case is unrealistic. Visually the image quality is much improved, and the reconstructed shape is consistent with the ground truth CT (unlike the preoperative CT compensation). This reconstruction also demonstrates the effect of varying relative portions of prior and observed data. We notice, for example, that the reconstruction from the  $90^\circ$  C-arm scan includes the CT bed at the bottom of the image; the corresponding region in the  $210^\circ$  scan shows a wooden board on which the specimen lay.

These visual observations are reflected in the similarity scores as well. Fig. 15 plots the correlation coefficient score between the ground truth postoperative CT and the various CBCT reconstructions as a function of the portion of observed X-ray images in the reconstruction input. We cover two regions of interest, which were highlighted in Fig. 14: a spherical segment of the full FOV, and a box region that contains only the injected cement.

Without compensation, the correlation is consistently poor: about 0.75 for the large FOV and only 0.66 in the region of change at the maximum, which is attained for the  $210^\circ$  scan. Truncation compensation, without reduced-arc compensation, increases the correlations to 0.81 and 0.78 for the respective regions. However, the truncation compensation helps little for the reduced arc of  $90^\circ$ , where the correlation score is about 0.45 or 0.46 with or without the prior data.

When arc extrapolation is used, the correlation is consistently higher. With the atlas, it goes from 0.69 at the low end to 0.75 at the high end for the large FOV, and from 0.68 to 0.75 for the box. With the preoperative CT, the corresponding values are 0.77–0.78 and 0.75–0.77. Importantly, both the prior CT and the atlas achieve a better correlation than the truncation compensation alone for scan arcs of  $150^\circ$  or less, which are common in clinical use.

We also include the effect of using the postoperative CT for compensation. The similarity decreases with the increased length of the X-ray scan, because the hybrid volume is a blend of the “ground truth” and the X-ray data, and the portion of ground truth decreases when the arc length increases. The correlation scores for this set of hybrid reconstructions are the highest, simply because the prior model is identical to the ground truth.

To conclude, this experiment studies the use of hybrid reconstruction technique to compensate for reduced arc scans. We recover high contrast details from very short scan angles, even as low as  $90^\circ$ . The use of prior models to compensate for the reduced scan is beneficial even with a coarse approximation, such as the statistical atlas, and the reconstruction quality improves when the external data source is more similar to the observed specimen. This calls for creation of more advanced anatomical models as a potential tool for improving cone beam reconstruction with limited means.

## VI. DISCUSSION

Image truncation and reduced scan length are common in clinical uses of C-arm CBCT. Our experiments show that a



hybrid reconstruction with prior anatomical models can very significantly improve the reconstruction and reduce artifacts, while preserving meaningful information from the observed X-ray projections. This is of particular importance if a region of change, such as injected cement or resected tissue, needs to be reconstructed. Our similarity scores consistently show that the various forms of hybrid reconstruction are better than reconstruction from the limited data only.

There are a few points to pay attention to when studying the quality of the reconstructed images. As we noted earlier, we accomplished only a partial calibration of our XR22 detector, with a simplified logarithmic model for the intensity response. The recorded intensities suffer from various distortions, including nonuniformities, nonlinearity, limited discretization range, heel effect, saturation, full occlusion, photon scatter, and others, which we did not address in this work. A better-calibrated scanner or better detector technology (such as a flat-panel detector) may produce better reconstructions than ours. The incomplete calibration also means that we could only use indirect similarity measures, as opposed to difference-based metrics, which are most desirable.

With this said, we can draw several conclusions from the experiments. The quality of hybrid reconstruction improves if the prior model is more similar to the scanned specimen; but even a relatively coarse approximation of the anatomy—a statistical atlas that only includes parts of the specimen’s structure—still produces a better reconstruction than without any prior.

Our registration of the atlas to the dry bone X-ray images was not as good as with the fresh specimen. One possible reason may be the dissimilarity in radiodensity between whole bone, as modeled in the atlas, and the mostly hollow dry bone. Another is that the specific dry specimen could be an outlier relative to the space of training shapes. While future atlases may be more robust and accurate than the one used here, outliers will likely exist always. Therefore, an outlier detection method may be desired.

As to the question of whether the original X-ray intensities should be remapped based on CT projections, the general answer is negative, as our results, summarized in Table II, show. Almost always, the C-arm reconstructions directly from the X-ray images have a higher similarity score to the CT than reconstructions from remapped X-ray images. Although for the reduced-arc compensation (Section V-D) we remapped the X-ray images rather than the DRRs, it was done to simplify the fusion process rather than have the best possible correlation to CT. Better imaging systems, with better sensitivity and calibration, may produce more consistent attenuation measures, and not require intensity mapping. For the system studied here, more robust techniques for intensity matching may exist.

Better anatomical models and reliable registration methods should lead to better atlas-based hybrid reconstruction. An immediate extension would include multiple bones, e.g., the femurs and spinal vertebrae for the pelvis region. Importantly, inclusion of soft tissue in the atlas is essential if the method is applied to whole patients, as opposed to isolated specimens. Soft tissue models need not be necessarily as detailed or accurate as the bone models. Perhaps even a simple external measurement of the patient could provide enough information about

his shape, and an assumption of uniform attenuation in the soft tissue might be sufficient to model it. This idea relates back to Boyd *et al.* [42], and its potential in deformable registration was demonstrated by Sadowsky [22]. Overall, imaging and reconstruction of soft tissue pose new challenges for future research.

Algorithms tailored for reconstructing a region of change, especially high-contrast ones such as the bone cement or resected tissue, may also produce sharper reconstructions, compared to the filtered back-projection we used. For example, hybrid reconstruction may be used complementarily with sinogram completion methods such as [5] and [6], to obtain either an improved registration (relying on sinogram-completed images) or a better estimate of the volume under reconstruction.

We did not attempt to optimize the performance of our methods. Our current image registration toolkit, based on [32] and [38], has a performance bottleneck in exchanging data between CPU and GPU, and a registration can take several minutes. The reconstruction toolkit is implemented very simply with a serial CPU computation. Modern GPU technology and algorithms (e.g., [43]) can accelerate this component in the pipeline.

Yet, we demonstrate that even a modest toolbox can be used successfully for hybrid reconstruction. With the issues above highlighted, this work paves the way towards its clinical use.

## VII. CONCLUSION

This paper demonstrated the use of hybrid cone-beam reconstruction that fuses observed X-ray images from a C-arm with a prior anatomical model to compensate for different types of information loss. We addressed specific limitations of mobile C-arms: narrow FOV that causes image truncation, and reduced scan trajectories. We showed how different types of prior models—patient-specific CT or a statistical atlas—are registered with the X-ray images, and presented a practical method to combine the different information sources as projected images. Our results are applicable not only to our relatively old C-arm, but to more modern ones which still suffer from such limitations.

A series of experiments showed that a high correlation between the CBCT and CT volumes can be accomplished, that reconstruction artifacts are significantly reduced with the fusion of prior images, and that high-contrast details can be recovered from a very short scan. For example, truncation compensation increased the correlation coefficient between CT and CBCT from 0.69 to 0.79 for a fresh pelvis specimen fused with a statistical atlas. Reconstructions from reduced scans, as short as 90°, improved the correlation with CT from about 0.45 to about 0.7 with a statistical atlas and 0.77 with a prior CT. Fusion of data from the reduced X-ray scan with a prior model can provide important context for reconstructing a region of change under limited imaging conditions.

Other potential uses of hybrid reconstruction include occluded regions, e.g., where surgical tools or implants are, or saturated regions, where attenuation cannot be detected. The technique might also be useful in standard CT reconstruction, and in low-dose imaging protocols.

The quality of the reconstruction depends, among other factors, on the quality of the prior model. This calls for improved

forms of statistical anatomical atlases. The atlas used here included a single bone structure, and was just adequate for fusion with a more complex anatomical structure of a fresh specimen, and inadequate for fusion with a dry bone specimen. Future developments of the atlas may study the variations of density distribution inside the bone, and include models of other tissues, such as internal organs and soft body structures.

Independently from the imaging devices and the representation of the anatomical model, we showed hybrid reconstruction used with an actual X-ray system, and demonstrated how it improves reconstruction quality by compensating for lost information. This is the major result of our work.

#### ACKNOWLEDGMENT

The authors would like to thank the International Center for Orthopaedic Advancement at the Johns Hopkins University Bayview Medical Center, headed by Dr. S. 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. T. DeWeese and Dr. L. Myers from the Johns Hopkins Department of Radiation Oncology. Other contributors to this research include G. Chintalapani, L. M. Ellingsen, I. Iordachita, and N. Kuo, all from the Johns Hopkins University Whiting School of Engineering. The authors would also like to thank Dr. J. Siewerdsen for commenting on this work.

#### REFERENCES

- [1] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *J. Opt. Soc. Am. A (JOSA A)*, vol. 1, no. 6, pp. 612–619, Jun. 1984.
- [2] O. Sadowsky, K. Ramamurthi, L. M. Ellingsen, G. Chintalapani, J. L. Prince, and R. H. Taylor, "Atlas-assisted tomography: Registration of a deformable atlas to compensate for limited-angle cone-beam trajectory," in *IEEE Int. Symp. Biomed. Imag. (ISBI)*, 2006, pp. 1244–1247.
- [3] R. M. Lewitt, "Processing of incomplete measurement data in computed tomography," *Med. Phys.*, vol. 6, no. 5, pp. 412–417, Sep. 1979.
- [4] J. Hsieh, E. Chao, J. Thibault, B. Grekowitz, A. Horst, S. McOlash, and T. J. Myers, "A novel reconstruction algorithm to extend the CT scan field-of-view," *Med. Phys.*, vol. 31, no. 9, pp. 2385–2391, Sep. 2004.
- [5] R. Chityala, K. R. Hoffmann, S. Rudin, and D. R. Bednarek, "Artifact reduction in truncated CT using sinogram completion," in *Proc. SPIE Med. Imag.*, 2005, vol. 5747, pp. 2110–2117.
- [6] A. A. Zamyatin and S. Nakanishi, "Extension of the reconstruction field of view and truncation correction using sinogram decomposition," *Med. Phys.*, vol. 34, no. 5, pp. 1593–1604, May 2007.
- [7] E. Y. Sidky, C.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *J. X-ray Sci. Technol.*, vol. 14, pp. 119–139, 2006.
- [8] E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Phys. Med. Biol.*, vol. 53, pp. 4777–4807, Sep. 2008.
- [9] G.-H. Chen, J. Tang, and S. Leng, "Prior image constrained compressed sensing 'PICCS': A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets," *Med. Phys.*, vol. 35, no. 2, pp. 660–663, Feb. 2008.
- [10] P. B. Noël, J. Xu, K. R. Hoffmann, and J. J. Corso, "Geometric tomography: A limited-view approach for computed tomography," in *Annu. Symp. Computat. Geometry (SCG)*, 2009, pp. 98–99.
- [11] F. Noo, M. Defrise, R. Clackdoyle, and H. Kudo, "Image reconstruction from fan-beam projections on less than a short scan," *Phys. Med. Biol.*, vol. 47, no. 14, pp. 2525–2546, Jul. 2002.
- [12] F. Noo, R. Clackdoyle, and J. D. Pack, "A two-step hilbert transform method for 2D image reconstruction," *Phys. Med. Biol.*, vol. 49, no. 17, pp. 3903–3923, Sep. 2004.
- [13] Y. Zou, X. Pan, and E. Y. Sidky, "Image reconstruction in regions-of-interest from truncated projections in a reduced fan-beam scan," *Phys. Med. Biol.*, vol. 50, no. 1, pp. 13–27, 2005.
- [14] M. Defrise, F. Noo, R. Clackdoyle, and H. Kudo, "Truncated hilbert transform and image reconstruction from limited tomographic data," *Inverse Problems*, vol. 2, no. 3, pp. 1037–1053, Jun. 2006.
- [15] H. Kudo, M. Courdurier, F. Noo, and M. Defrise, "Tiny a priori knowledge solves the interior problem in computed tomography," *Phys. Med. Biol.*, vol. 53, no. 7, pp. 2207–2231, 2008.
- [16] H. Yu and G. Wang, "Feldkamp-type voi reconstruction from super-short-scan cone-beam data," *Med. Phys.*, 2004.
- [17] K. Ramamurthi, "Cone-beam tomography using C-arm X-ray projections: Complete trajectories and integration of prior CT information," Ph.D. dissertation, Johns Hopkins Univ., Baltimore, MD, 2006.
- [18] R. Fahrig, M. Moreau, and D. Holdsworth, "Three-dimensional computed tomographic reconstruction using a C-arm mounted XRRI: Correction of image intensifier distortion," *Med. Phys.*, vol. 27, no. 7, pp. 1097–1106, Jul. 1997.
- [19] N. Navab, M. Mitschke, and O. Schütz, C. Taylor and A. Colchester, Eds., "Camera-augmented mobile C-arm (CAMC) application: 3D reconstruction using a low-cost mobile C-arm," in *Proc. MICCAI '99*, 1999, LNCS, pp. 688–698, 1679.
- [20] K. Wiesent, K. Barth, N. Navab, P. Durlak, T. Brunner, O. Schuetz, and W. Seissler, "Enhanced 3D-reconstruction algorithm for C-arm systems suitable for interventional procedures," *IEEE Trans. Med. Imag.*, vol. 19, no. 5, pp. 391–403, May 2000.
- [21] 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," *Med. Phys.*, vol. 35, no. 5, pp. 2124–2136, May 2008.
- [22] O. Sadowsky, "Image registration and hybrid volume reconstruction of bone anatomy using a statistical shape atlas," Ph.D. dissertation, Johns Hopkins Univ., Baltimore, MD, 2008.
- [23] D. Ritter, M. Mitschke, and R. Graumann, "Intraoperative soft tissue 3D reconstruction with a mobile C-arm," in *Computer Assisted Radiology and Surgery (CARS)*, 2003. New York: Elsevier, 2003, pp. 200–206.
- [24] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Am. A*, vol. 4, no. 4, pp. 629–642, 1987.
- [25] W. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multimodal volume registration by maximization of mutual information," *Med. Image Anal.*, vol. 1, no. 1, pp. 35–51, 1996.
- [26] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [27] J. Yao, "A statistical bone density atlas and deformable medical image registration," Ph.D. dissertation, Johns Hopkins Univ., Baltimore, MD, 2002.
- [28] A. Mohamed and C. Davatzikos, "An approach to 3D finite element mesh generation from segmented medical images," in *IEEE Int. Symp. Biomed. Imag. (ISBI)*, 2004, pp. 420–423.
- [29] L. M. Ellingsen and J. L. Prince, "Mjolnir: Deformable image registration using feature diffusion," in *Proc. SPIE Med. Imag. Conf.*, 2006, vol. 6144, pp. 329–337.
- [30] T. F. Cootes and C. J. Taylor, Statistical models of appearance for computer vision Imag. Sci. Biomed. Eng. Univ. Manchester, Manchester, U.K., Tech. Rep., 2004.
- [31] G. Chintalapani, O. Sadowsky, L. M. Ellingsen, J. L. Prince, and R. H. Taylor, "Integrating statistical models of bone density into shape based 2D-3D registration framework," in *MICCAI 2009 Workshop: Probabilistic Models for Medical Image Analysis*, 2009, pp. 151–161.
- [32] O. Sadowsky, J. D. Cohen, and R. H. Taylor, "Projected tetrahedra revisited: A barycentric formulation applied to digital radiograph reconstruction using higher-order attenuation functions," *IEEE Trans. Vis. Comput. Graphics*, vol. 12, no. 4, pp. 461–473, Jul./Aug. 2006.
- [33] J. A. Nelder and R. Mead, "A simplex method for function minimization," *Computer J.*, vol. 7, pp. 308–313, 1965.
- [34] The VXL Software Package, Open Source Libraries. [Online]. Available: <http://vxl.sourceforge.net>

- [35] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, no. 2, pp. 187–198, Apr. 1997.
- [36] L. Zöllei, E. Grimson, A. Norbash, and W. Wells, "2D-3D rigid registration of X-ray fluoroscopy and CT images using mutual information and sparsely sampled histogram estimators," in *IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR) 2001*, 2001, vol. II, pp. II 696–II 703.
- [37] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recognit.*, pp. 71–86, 1999.
- [38] O. Sadowsky, G. Chintalapani, and R. H. Taylor, N. Ayache, S. Ourselin, and A. Maeder, Eds., "Deformable 2D-3D registration of the pelvis with a limited field of view, using shape statistics," in *Med. Image Computing Computer-Assisted Intervent. – MICCAI 2007*, 2007, vol. 4792, LNCS, pp. 519–526.
- [39] Z. F. Knops, J. B. A. Maintz, M. A. Viergever, and J. P. W. Pluim, "Normalized mutual information based registration using k-means clustering and shading correction," *Med. Image Anal.*, vol. 10, no. 3, pp. 432–439, 2006.
- [40] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Interpolation artefacts in mutual information-based image registration," *Comput. Vis. Image Understand.*, vol. 77, no. 9, pp. 211–232, 2000.
- [41] F. Jäger and J. Hornegger, "Nonrigid registration of joint histograms for intensity standardization in magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 28, no. 1, pp. 137–150, Jan. 2009.
- [42] J. E. Boyd and J. J. Little, "Complementary data fusion for limited-angle tomography," in *Proc. IEEE Comput. Vis. Pattern Recognit. (CVPR)*, 1994, pp. 288–294.
- [43] P. B. Noël, A. M. Walczak, J. Xu, J. J. Corso, K. R. Hoffmann, and S. Schafer, "GPU-based cone beam computed tomography," in *Computer Methods and Programs in Biomedicine (HP-MICCAI)*, Jun. 2010, vol. 98, no. 3, pp. 271–277.