Published in final edited form as:

IEEE Trans Med Imaging. 2012 April ; 31(4): 860–869. doi:10.1109/TMI.2011.2171498.

An Active Contour Method for Bone Cement Reconstruction From C-Arm X-Ray Images

Blake C. Lucas,
Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 USA

Yoshito Otake,
Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 USA

Mehran Armand, and
Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 USA

Russell H. Taylor [Fellow, IEEE]
Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 USA

Blake C. Lucas: blake@cs.jhu.edu

Abstract

A novel algorithm is presented to segment and reconstruct injected bone cement from a sparse set of X-ray images acquired at arbitrary poses. The sparse X-ray multi-view active contour (SxMAC —pronounced “smack”) can 1) reconstruct objects for which the background partially occludes the object in X-ray images, 2) use X-ray images acquired on a noncircular trajectory, and 3) incorporate prior computed tomography (CT) information. The algorithm’s inputs are preprocessed X-ray images, their associated pose information, and prior CT, if available. The algorithm initiates automated reconstruction using visual hull computation from a sparse number of X-ray images. It then improves the accuracy of the reconstruction by optimizing a geodesic active contour. Experiments with mathematical phantoms demonstrate improvements over a conventional silhouette based approach, and a cadaver experiment demonstrates SxMAC’s ability to reconstruct high contrast bone cement that has been injected into a femur and achieve sub-millimeter accuracy with four images.

Index Terms

Active contour; bone cement; deformable models; intra-operative imaging; reconstruction; segmentation

I. Introduction

The problem of recovering 3-D shape from a sparse set of 2-D projection images is common in interventional imaging. If prior information such as a statistical shape model is available, this information may be used to assist in reconstruction [1]. However, such information is

Correspondence to: Blake C. Lucas, blake@cs.jhu.edu.

not always available, especially if the object is highly deformable or its shape is created and/or substantially modified during the procedure. Examples include surgical procedures for injecting cement into bones, such as vertebroplasty [2], sacroplasty [3], and femoroplasty [4].

Our immediate, focusing clinical application is femoral bone augmentation, in which cement is injected into the neck and in-tertrochanteric area of the femur for patients with severe osteoporosis. The goal is to strengthen the femur to prevent fractures [4]. In a recently proposed procedure for bone augmentation (see Fig. 1) [5], pre-operative CT images are used for pre-operative planning, based on a 3-D finite element analysis of the patient's femur and planned cement injection [6]. The pre-operative model and plan are registered to the patient and intra-operative navigation system using 2-D-3-D registration from intra-operative X-rays, and a robotic device is used to inject cement containing an appropriate contrast agent. At various stages during the injection, a sparse set of intra-operative X-rays (at most 8, but preferably 4) are taken of the cement and the cement volume in the bone is estimated. This information is used to repeat the finite-element analysis of augmented bone strength and to support re-planning and optimization for further injections. Conventionally, the shape of the cement volume is estimated by intersecting cones formed from the silhouettes of the cement in the images. However, the resulting models do not accurately reflect the actual cement volumes [i.e., Fig. 7(b) as compared to ground truth shown in Fig. 7(a)]. Our goal in this work is to significantly improve the accuracy of this reconstruction while still using only a small number of intra-operative X-rays from a conventional C-arm.

II. Background

Techniques have been developed in computer vision to reconstruct objects observed from multiple viewpoints without prior information. One classical approach is to segment an object's silhouette in images, back-project the silhouettes into 3-D space, and compute the intersecting volume. This technique is known as silhouette reconstruction or visual hull [7] computation, and has been used in computer vision [8] and C-arm X-ray reconstruction [9]. It has been shown that the visual hull encloses all other reconstructions that can be explained by 2-D segmentations of an object [10]. The visual hull is unlikely to be consistent with observed image intensities of the object. However, the visual hull can be used to initialize more sophisticated approaches that generate reconstructions consistent with image intensities. In particular, geodesic active contours [11] reconstruct objects by optimizing an objective function on the image intensities that considers all observations of an object simultaneously [12], [13].

Active contour techniques have been extended to 2-D to-mographic reconstruction [14] and cone-beam CT (CBCT) [15]–[17]. CBCT is a volumetric reconstruction from a series of X-ray images (more than 100) acquired with a 2-D X-ray image scanner, such as a C-arm. The active contour segmentation process is formulated as the minimization of an objective function (1) that incorporates information from K X-ray images $\{I_k(\boldsymbol{x}_k)|\boldsymbol{x}_k \in D[k] \subset \mathbb{R}^2\}$ in log space and linear attenuation coefficients $\{\mu(\boldsymbol{X})|\boldsymbol{X} \in D \subset \mathbb{R}^3\}$ in patient coordinates (E_{data}). The objective function also incorporates information about geometric properties of

the object (E_{model}), such as a penalty against high curvature. See Fig. 2 for a depiction of the imaging scenario and Table I for definitions of all major terms

$$E=E_{\text{model}}(\boldsymbol{X})+E_{\text{data}}(\mu(\boldsymbol{X}),I_1(\boldsymbol{x}_1),\ldots,I_K(\boldsymbol{x}_K)) \quad (1)$$

The objective function (1) is optimized by alternating between minimization of E_{model} and E_{data} until convergence. One choice for the data term uses the log-likelihood estimator on the foreground $P_{\text{fg}}(\boldsymbol{x}_k)$ and background $P_{\text{bg}}(\boldsymbol{x}_k)$ probability densities [17]

$$E_{\text{data}}=\sum_k^K \int_{D_{fg}[k]}(-\log(P_{\text{fg}}(\boldsymbol{x}_k)))\,d\boldsymbol{x}_k+\int_{D_{bg}[k]}(-\log(P_{\text{bg}}(\boldsymbol{x}_k)))\,d\boldsymbol{x}_k \quad (2)$$

where the foreground $D_{fg}[k]$ and background $D_{bg}[k]$ are assumed to appear as disjoint regions in X-ray images. The objective function can be expressed using a Heaviside function $H(\cdot)$ on the silhouette's level set $\{\psi_k(\boldsymbol{x}_k)|\boldsymbol{x}_k\in D[k]\subset\mathbb{R}^2\}$ such that the level set is positive outside and negative inside the silhouette of the deformable model

$$E_{\text{data}}=-\sum_k^K \int_{D[k]}\log(P_{\text{fg}}(\boldsymbol{x}_k))\,(1-H(\psi_k(\boldsymbol{x}_k)))+\log(P_{\text{bg}}(\boldsymbol{x}_k))\,H(\psi_k(\boldsymbol{x}_k))\,d\boldsymbol{x}_k. \quad (3)$$

There may not exist prior information about the probability distributions of the X-ray intensities. Instead, the active contour may be driven by image gradient information [17] or dynamically estimate the object's appearance with a Mumford–Shah approach. Alvino and Yezzi [14] proposed the following Mum-ford–Shah objective function for CT reconstruction:

$$E_{\text{data}}=\sum_k^K \int(I(s,\theta_k)-f(s,\theta_k))^2 ds+\int_D\{\lambda_{fg}|\nabla\mu_{fg}(\mathbf{x})|^2 H_{fg}(\mathbf{x})+\lambda_{bg}|\nabla\mu_{bg}(\mathbf{x})|^2 H_{bg}(\mathbf{x})\}d\mathbf{x} \quad (4)$$

where

$$f(s,\theta)=\int(\mu_{fg}(\mathbf{x})H_{fg}(\mathbf{x})+\mu_{bg}(\mathbf{x})H_{bg}(\mathbf{x}))\delta(x\cos\theta+y\sin\theta-s)d\mathbf{x} \quad (5)$$

$f(s,\theta)$ is the radon transform of $\mu(\boldsymbol{X})$. The latter Mum-ford–Shah approach has the advantage that it does not require prior knowledge of the foreground and background appearance. However, it assumes the foreground and background densities are smoothly varying, which may not be the case. To achieve high accuracy, the objective function must be representative of the anatomical contents and imaging scenario for a specific application.

To serve the needs of the bone augmentation procedure previously described, we focus on reconstructing homogenous highly deformable objects (i.e., bone cement) from C-arm X-ray projection images. In this paper, the segmentation process is formulated as an optimization problem that permits the segmentation algorithm to 1) reconstruct deformable objects for which the back-ground partially occludes the object in X-ray images, 2)use X-ray images

acquired on a noncircular trajectory, and 3) incorporate prior CT information. Subsequently, we describe a method for optimizing the objective function and evaluate the feasibility and performance of the Sparse X-ray multi-view active contour algorithm (SxMAC—pronounced “smack”) to reconstruct injected bone cement. In particular, we are interested in knowing how many X-ray images, how much contrast, and how large a sweep angle is required to achieve acceptable accuracy.

III. Method

Following the approach by Alvino and Yezzi [14], the segmentation process is formulated as an optimization problem that minimizes the following objective function, which describes the disparity between X-ray images and digitally reconstructed radiographs (DRRs) of a deformable model:

$$E_{\text{data}}=\sum_{k}^{K}\int_{D[k]}(I_k(\boldsymbol{x}_k)-\text{DRR}_k\{\mu\}(\boldsymbol{x}_k))^2 d\boldsymbol{x}_k+\int_D\{\lambda_{fg}|\nabla\mu_{fg}(\boldsymbol{X})|^2H_{fg}(\boldsymbol{X})+\lambda_{bg}|\nabla\mu_{bg}(\boldsymbol{X})|^2H_{bg}(\boldsymbol{X})d\mathbf{X} \quad (6)$$

where

$$\text{DRR}_k\{\mu\}(\boldsymbol{x}_k)=\int_D\mu(\boldsymbol{X})\delta(\boldsymbol{x}_k-\wp_k(\boldsymbol{X}))\,d\boldsymbol{X} \quad (7)$$

and

$$\mu(\boldsymbol{X})=\mu_{fg}(\boldsymbol{X})H_{fg}(\boldsymbol{X})+\mu_{bg}(\boldsymbol{X})H_{bg}(\boldsymbol{X}) \quad (8)$$

and

$$H_{fg}(\boldsymbol{X})=(1-H(\omega_{bg}(\boldsymbol{X})))\,(1-H(\omega_{fg}(\boldsymbol{X}))) \quad (9)$$

and

$$H_{bg}(\boldsymbol{X})=(1-H(\omega_{bg}(\boldsymbol{X})))\,H(\omega_{fg}(\boldsymbol{X})). \quad (10)$$

The objective function E_{data} measures the L^2 norm of the difference between simulated X-rays of μ and the log of each X-ray image, subject to an L^2 penalty on the smoothness of μ. Alternatively, the objective function could be expressed with an L^1 instead of L^2 norm, for which there is evidence that the L^1 norm may have better performance if there exist a sparse representation of image intensity information [18], [19]. Likewise, the smoothness term [second term in (6)] could minimize the total variation of the foreground and background appearance $|\nabla\,\mu(\boldsymbol{X})|$. The following discourse will focus on L^2 optimization because the objective function can be efficiently solved with linear methods. The complete objective function for the L^2 norm, including data and geometric terms, is as follows:

$$E=\sum_{k}^{K}\int_{D[k]}(I_k(\boldsymbol{x}_k)-\mathrm{DRR}_k\{\mu\}(\boldsymbol{x}_k))^2 d\boldsymbol{x}_k + \int_D\{\lambda_{fg}|\nabla\mu_{fg}(\boldsymbol{X})|^2 H_{fg}(\boldsymbol{X}) + \lambda_{bg}|\nabla\mu_{bg}(\boldsymbol{X})|^2 H_{bg}(\boldsymbol{X}) + \lambda_\kappa\delta(\omega_{fg}(\boldsymbol{X}))|\nabla\omega_{fg}(\boldsymbol{X})|\}\, d\mathbf{X} \quad (11)$$

$\mathrm{DRR}_k(\cdot)$ can be discretized and expressed as a weighted linear combination of $\mu(\boldsymbol{X})$

$$\mathrm{DRR}_k\{\mu\}(\boldsymbol{x}_k)=\sum_{\{\boldsymbol{X}|\boldsymbol{x}_k=\wp_k(\boldsymbol{X})\}} w_k(\boldsymbol{X})\mu(\boldsymbol{X}) \quad (12)$$

or alternatively

$$\mathrm{vec}(\mathrm{DRR}_k\{\mu\})=\boldsymbol{M}_k\mathrm{vec}(\mu). \quad (13)$$

$\boldsymbol{M}_k$ is an $M \times N$ matrix where M is the number of pixels in the X-ray image and N is the number of voxels in μ. The matrix is completely defined by X-ray geometry (extrinsic and intrinsic parameters) and does not depend on the image or volume intensities. To efficiently solve (11), we let the foreground and background appearances be modeled as a constant (i.e., $\mu_{fg}(\boldsymbol{X}) = c_{fg}$ and $\mu_{bg}(\boldsymbol{X}) = c_{bg}$. This assumption is consistent with work by Chan and Vese [20]. The objective function simplifies to

$$E=\sum_{k}^{K}||\mathrm{vec}(I_k)-c_{fg}\boldsymbol{M}_k\mathrm{vec}(H_{fg})-c_{bg}\boldsymbol{M}_k\mathrm{vec}(H_{bg})||^2+\int_D\lambda_\kappa\delta(\omega_{fg}(\boldsymbol{X}))|\nabla\omega_{fg}(\boldsymbol{X})|d\mathbf{X}. \quad (14)$$

The SxMAC model can be augmented to incorporate prior CT information by replacing the background indicator $\mathrm{H}_{bg}(\boldsymbol{X})$, with the prior CT ($\mu_{\mathrm{CT}}$)

$$\tilde{H}_{bg}(\boldsymbol{X})=H_{bg}(\boldsymbol{X})\mu_{\mathrm{CT}}(\boldsymbol{X}). \quad (15)$$

This extension assumes μ_{CT} is properly registered and intensity calibrated so that the background DRR is highly correlated with the background observed in acquired X-ray images. $\mathrm{H}_{bg}(\boldsymbol{X})$ is a segmentation mask of the field of view that is common to both the prior CT and C-arm acquisition that encloses the foreground $\mathrm{H}_{fg}(\boldsymbol{X})$.

Alternatively, we can incorporate prior information by replacing I_k with the difference ΔI_k between X-ray images obtained before and after cement injection. Either way, the optimization procedure is the same.

Solving the Euler–Lagrange equation for c_{fg} and c_{bg}

$$0=\frac{\partial E}{\partial c_{fg}}\Big|_{c_{fg}} =-2c_{fg}\sum_{k}^{K}\mathrm{vec}^{T}(H_{fg})\boldsymbol{M}_{k}^{T}(\mathrm{vec}(I_{k})-\boldsymbol{M}_{k}\mathrm{vec}(\mu)). \quad (16)$$

From (16) we obtain

$$\hat{c}_{fg}(c_{bg})=\sum_{k}^{K}\frac{(\boldsymbol{M}_{k}\mathrm{vec}(H_{fg}))^{T}(\mathrm{vec}(I_{k})-c_{bg}\boldsymbol{M}_{k}H_{bg})}{||\boldsymbol{M}_{k}\mathrm{vec}(H_{fg})||^{2}} \quad (17)$$

and similarly

$$\hat{c}_{bg}(c_{fg})=\sum_{k}^{K}\frac{(\boldsymbol{M}_{k}\mathrm{vec}(H_{bg}))^{T}(\mathrm{vec}(I_{k})-c_{fg}\boldsymbol{M}_{k}(\mathrm{H}_{fg}))}{||\boldsymbol{M}_{k}\mathrm{vec}(H_{bg})||^{2}}. \quad (18)$$

The model's appearance can be optimized by alternating between $c_{fg} = \hat{c}_{fg}(c_{bg})$ and $c_{bg} = \hat{c}_{bg}(c_{fg})$. Evolution of the deformable model's level set $\omega_{fg}(\boldsymbol{X})$ is computed by gradient descent [11]

$$\omega_{fg}^{t+1}(\boldsymbol{X})=\omega_{fg}^{t}(\boldsymbol{X})-\gamma\frac{\partial E}{\partial\omega_{fg}^{t}} \quad (19)$$

where

$$\frac{\partial E}{\partial\omega_{fg}}\Big|_{\boldsymbol{X}=\boldsymbol{X}_{n}}=2(C_{fg}-C_{bg})\delta(\omega_{fg}(\boldsymbol{X}_{n}))\sum_{k}^{K}\mathbb{I}_{n}\boldsymbol{M}_{k}^{T}(\mathrm{vec}(I_{k})-\boldsymbol{M}_{k}\mathrm{vec}(\mu))+\lambda_{\kappa}\delta(\omega_{fg}(\boldsymbol{X}_{n}))\nabla\cdot\frac{\nabla\omega_{fg}(\boldsymbol{X}_{n})}{|\nabla\omega_{fg}(\boldsymbol{X}_{n})|}. \quad (20)$$

$\boldsymbol{M}_k$ and $\boldsymbol{M}_k^T$ are referred to as forward and backward projection operators, respectively. In the actual implementation, we do not store $\boldsymbol{M}_k$ and $\boldsymbol{M}_k^T$ because they are very large sparse matrices. Instead, the graphics card is used to compute elements of those matrices on-the-fly with OpenCL. We choose a voxel driven approach for its simplicity [21], [22], but more accurate methods have been developed that can be implemented on the graphics card [23]–[25].

The Euler–Lagrange optimization procedure is only guaranteed to find local optima. Therefore, it's important to initialize the model close to the globally optimal solution. When possible, we initialize SxMAC with the object's visual hull obtained from automated silhouette reconstruction.

IV. Results

A. Phantom Experiments

SxMAC was evaluated on mathematical phantoms of solid objects. Synthetic X-ray projections for the objects were generated with the same DRR operator used in the segmentation process. Table II lists parameters used in phantom experiments that reflect realistic C-arm X-ray imaging parameters. Fig. 3 depicts reconstruction of a metasphere from two X-ray images spaced 90° apart ($\lambda_\kappa = 0.05$). SxMAC recovers concavities that are not recoverable with silhouette reconstruction alone. Fig. 4 depicts a torus reconstructed from four images ($\lambda_\kappa = 0.1$) with different angular spacing. Notice that the deformable model better captured the hole in the middle of the torus when the acquisition trajectory was ±10° above and below the orbital plane. Fig. 5 shows a reconstruction of a dragon from six images ($\lambda_\kappa = 0.01$) acquired from a circular trajectory. The regularization parameter λ_κ should be chosen large enough to reduce directional bias, but small enough to preserve sharp object features. The dragon reconstruction demonstrates SxMAC has applications in the broader realm of 3-D model acquisition.

B. SxMAC Reconstruction With Synthesized Images

Cadaver experiments were conducted to evaluate a surgical procedure in which bone cement is injected into an osteoporotic femur. One objective of the procedure is to provide feedback to the surgeon via intra-operative reconstruction of the injected bone cement. Predictions about the post-operative structural properties of the femur can be made from its 3-D reconstruction and pre-operative CT [6]. The imaging scenario has been described in previous work [5] as follows.

1. Either a pre-operative diagnostic CT or CBCT is acquired.
2. C-arm X-ray images are acquired and registered to the pre-operative CT with 2-D/3-D registration.
3. Pre-injection X-ray images are taken to guide the procedure.
4. Bone cement is injected into the femur; after which, X-ray images are acquired with the C-arm.
5. SxMAC is used to reconstruct the injected bone cement.
6. Finite element analysis of the cement is conducted.
7. The injection procedure is replanned and repeated from step 4 until the surgeon is satisfied with the result.

This workflow is applicable to using either pre-operative CT or pre-injection X-ray images as priors. SxMAC does not address how to register pre-operative CT to intra-operative X-ray images nor how to calibrate intensities between DRRs and X-ray images, although solutions have been described in literature [26], [27]. We will later describe how registration and intensity calibration can be avoided by using pre-injection X-ray images as priors instead of pre-operative CT.

To evaluate the performance of just the reconstruction algorithm, we synthesized X-ray images from pre- and post-operative CBCTs (Fig. 6). An intra-operative CBCT would not be acquired in a real clinical scenario, which is why a sparse X-ray reconstruction method like SxMAC is necessary. The following preprocessing steps synthesize pre- and post-operative CBCT and X-ray images that are perfectly registered and intensity matched.

1. Pre- and post-operative CBCTs are acquired with a flat panel C-arm (Table III). Geometric calibration follows the method described by Daly *et al.* [28].
2. Pre-operative CBCT is registered to post-operative CBCT using intensity-based 3-D/3-D registration.
3. The femur is segmented in pre-operative CBCT.
4. The cement is segmented in post-operative CBCT and the volume of interest (VOI) is copied and pasted into the registered pre-operative CBCT.
5. DRRs for the CBCTs with and without the bone cement are generated.

SxMAC was evaluated on these synthetic X-ray images. The CBCT and deformable model's level set representation were sampled at 1 mm isotropic, and the regularization parameter was chosen to be ($\lambda_\kappa = 0.001$). The background $\omega_{bg}(\boldsymbol{X})$ corresponded to a segmentation of the femur. SxMAC was initialized with a sphere (10 mm radius) located at the center of the bone cement. Such initialization is clinically plausible because the cement injector's tip is tracked with a Polaris optical tracker as part of the bone augmentation procedure. We were unable to reliably initialize SxMAC with silhouette reconstruction because it failed in several experiments due to lack of contrast.

In the first set of experiments, the attenuation of the cement was varied from 1645 to 2070 HU at equal intervals, and the number of images was held constant at eight images (Fig. 7). In the second set of experiments, the number of images was varied from 2 to 10, and the cement's attenuation was held constant at 1900 HU (Fig. 8). Each collection of images was evenly sampled in plane from a 180° arc trajectory. In the third set of experiments, the sweep angle was varied between 20° and 180° (Fig. 9), and the number of images and cement attenuation were held constant at 8 and 1900 HU, respectively. In the final set of experiments, the pre-operative CBCT was misaligned with the post-operative CBCT, either by translating it in a random direction or rotated it around a random axis, to simulate the effect of registration error. Experiments measuring the accuracy of SxMAC reconstruction were repeated 10 times per translation (Fig. 10) and rotation (Fig. 11). We compensate for 2-D/3-D registration error with 2-D Affine registration of pre-operative and post-operative X-ray images, also plotted in Figs. 10 and 11.

Intra-operative reconstruction errors were measured in terms of distance relative to mesh vertices on the ground truth reconstructed surface (truth-to-reconstruction) or vice versa (reconstruction-to-truth). A 180° sweep angle was required to obtain sub-millimeter accuracy of 1.0±0.75 (four images) and 0.85±0.62 (eight images) mm for reconstruction-to-truth and 0.84±0.52 (four images) and 0.73±0.51 (eight images) for truth-to-reconstruction. Reconstructions are depicted in Fig. 12.

Results demonstrate SxMAC's ability to use intensity information to recover shape information that is not recoverable with silhouette reconstruction alone. Results from the cadaver experiment demonstrate SxMAC's reconstruction performance with prior CT. Fig. 7 suggests that SxMAC reconstruction can be greatly improved by increasing the bone cement attenuation from 1550 HU to at least 1900 HU. An increase in attenuation could be achieved by using a higher concentration of Barium in the PMMA bone cement. Figs. 8 and 9 demonstrate sub-millimeter accuracy is achievable with a minimum of four images and 180° sweep angle. Reconstruction accuracy improves with increasing sweep angle because there is less redundant information between views. This trend is evident in phantom experiments (Fig. 4). There was a 15% improvement between four and eight images, and only a 1% improvement between eight and ten images. A 180° sweep angle is necessary to obtain sub-millimeter reconstruction-to-truth accuracy with either four or eight images. Acquisition of four to eight images is clinically relevant for intra-operative procedures.

Algorithms for 2-D/3-D registration are reported to have less than 2 mm translation and 1° rotation error in the pelvic region [5], [29], [30]. With the addition of 2-D image registration, SxMAC achieves 1 mm accuracy in the range of these misalignment errors and is robust to even larger displacement errors (see Figs. 10 and 11).

C. SxMAC Reconstruction With Real Images

The first cadaver study produced many important findings; however, the image contrast between cement (1550 HU) and surrounding tissue was insufficient to use X-ray images acquired from that experiment for reconstruction. Observing that reconstruction accuracy is sensitive to the amount of image content that can be explained with prior information, we changed the imaging scenario and conducted a second cadaver study that uses pre-injection X-ray images as priors. This modification circumvents issues with synthesizing realistic X-ray images from pre-operative CT.

In the second study, X-ray images were acquired immediately before and after cement injection from the same poses. With modern C-arms, it is trivial to retake images from previous poses because the C-arm is motorized and position encoded. Care was taken not to move the cadaver between acquisition of pre- and post-injection images, although small patient motions could be corrected with 2-D/2-D registration of X-ray images. We then computed the log of the pre- and post-injection images and subtracted them to produce images that contained only the cement (Fig. 13). Additionally, we chose a cadaver with a hip implant in the other leg. This reduced the sweep angle in which the cement was visible to 130°. Occlusion is anticipated because the bone augmentation procedure is expected to occur in conjunction with hip replacement or revision surgery.

The cement was first automatically segmented with 2-D active contours [20]. A silhouette reconstruction was computed from 2-D segmentations and used to initialize SxMAC. To capture spiculated cement features visible in the second experiment, the level set resolution was increased to 0.6 mm isotropic. The regularization parameter ($\lambda_\kappa = 0.05$) was chosen roughly to match the feature size of the ground truth segmentation and held constant in all experiments. SxMAC was run on collections of images evenly sampled from a 130° arc. Figs. 14 and 16 show that sub-millimeter accuracy is achievable with as few as four images

(0.69±1.03 mm). We then held the number of images constant at 4 and varied the sweep angle for SxMAC reconstruction. Fig. 15 shows that 1.54±1.31 mm accuracy is achievable with at least a 90° sweep angle.

V. Discussion

It is interesting to observe that reconstruction accuracy is better with fewer images. This is because increasing the number of images increases the risk of a poor 2-D segmentation in at least one image, leading to less accuracy in silhouette and subsequent SxMAC reconstruction. Consider a pathological case where the cement is occluded in only one of many images. Since the object has no silhouette in that one image, the visual hull reconstruction will contain nothing. Therefore, it is important that the cement is not occluded in any images used for reconstruction. For a thorough discussion of visual hulls, see Laruentini [7].

It may seem surprising that compared to SxMAC, silhouette reconstruction is slightly more accurate when measuring error from truth-to-reconstruction. However, Fig. 16 shows that SxMAC reconstructions better resemble ground truth and are less biased towards particular imaging directions. The regularization term (λ_κ) is responsible for reducing directional bias and other aliasing artifacts.

SxMAC makes several assumptions that are particular to bone cement injection procedures. Namely, the object is assumed to be homogeneous, located in a particular region of space, and the image background resembles DRRs of a pre-operative CT or pre-injection image. These assumptions enable SxMAC to produce accurate reconstructions where there is soft tissue, truncation, and metal that would pose challenges for to-mographic reconstruction algorithms [31], [32], which operate under fewer assumptions and generally require more images than the four needed by SxMAC. Another important difference between SxMAC and cone-beam reconstruction approaches is that SxMAC produces a geometric representation of the object, whereas cone-beam reconstruction produces images that must still be segmented to extract the object's geometry. Segmenting objects in CBCT images can be challenging because of metal and streak artifacts [33], which are to be expected in intra-operative imaging scenarios.

The synthesized X-ray images used to evaluate SxMAC are representative of ideal results achievable with the following preprocessing pipeline: register the diagnostic CT to X-ray images, generate DRRs from diagnostic CT, and remap the X-ray intensities to DRR intensities. SxMAC's true performance during an intra-operative procedure will be affected by the performance of these preprocessing steps.

To circumvent challenges posed by 2-D/3-D registration, we have presented an alternative imaging scenario that uses image subtraction instead of registration. SxMAC's performance in this scenario is superior to the previous, and the requirements for bone cement contrast and sweep angle are less stringent. The procedure is independent of patient anatomy and thus adaptable to other injection regions, such as the spine. A standard CBCT with our C-arm requires 400 X-ray images. SxMAC achieves sub-millimeter accuracy with 4 images, or 1% of the dosage. This suggests SxMAC can be incorporated into a surgical procedure

without affecting patient safety. Moreover, the dosage is low enough that more than one reconstruction could be performed during the procedure to provide more feedback to the surgeon.

Although the intended application for SxMAC is to intra-operatively reconstruct bone cement, there are other surgical procedures that involve bone cement injection, such as vertebroplasty [2], sacroplasty [3], and femoroplasty [4], that could benefit from intra-operative reconstruction. In the same spirit as Varshney *et al.* [17], SxMAC could also be applied to metal implant reconstruction.

The dragon reconstruction suggests SxMAC has other imaging applications in 3-D model acquisition. The advantage of using a C-arm X-ray machine to acquire a 3-D model is that X-rays can image occluded regions that could not be imaged with a laser range scanner, structured light system, multi-view stereo system, or similar technology. C-arm X-ray machines are popular at medical institutions, making 3-D model acquisition with a C-arm X-ray machine potentially cheaper and more accessible than other technologies.

VI. Conclusion

This paper has presented an algorithm for 3-D segmentation of bone cement observed in a small number of X-ray images. The SxMAC algorithm provides a computationally efficient procedure for segmentation that can incorporate prior CT and an initial estimate for the object's shape and location. The algorithm is implemented within an automated pipeline whose inputs are preprocessed X-rays, their associated pose information, and prior CT, if available. Cadaver experiments demonstrate SxMAC can segment injected bone cement with sub-millimeter accuracy from as few as 4 X-ray images, and thus it is readily applicable to intra-operative surgical procedures.

The regularization parameter λ_κ was held constant for each cadaver experiment, but it is an important parameter that should be optimized in the future once the number of images, sweep angle, and cement viscosity are decided upon. Future work will also examine the accuracy and sensitivity of finite element analysis to SxMAC reconstruction of the cement.

Acknowledgments

This work was supported in part by the National Institutes of Health (NIH) under Grant 5R21EB007747-02, in part by a graduate student fellowship from the Johns Hopkins Applied Physics Laboratory, and in part by a research fellowship from Japan Society for the Promotion of Science (JSPS) Postdoctoral Fellowships for Research Abroad.

The authors would like to thank Dr. J. Siewerdsen and others at the laboratory of Imaging for Surgery, Therapy, and Radiology (ISTAR) at Johns Hopkins University for hosting the cadaver experiments.

References

1. Sadowsky O, Lee J, Sutter E, Wall S, Prince J, Taylor R. Hybrid cone-beam tomographic reconstruction: Incorporation of prior anatomical models to compensate for missing data. IEEE Trans Med Imag. Jan; 2011 30(1):69–83.
2. Garfin SR, Yuan HA, Reiley MA. New technologies in spine: Kyphoplasty and vertebroplasty for the treatment of painful osteoporotic compression fractures. Spine. 2001; 26:1511. [PubMed: 11462078]

3. Richards A, Mears S, Knight T, Dinah A, Belkoff S. Biomechanical analysis of sacroplasty: Does volume or location of cement matter? Am J Neuroradiol. 2009; 30:315. [PubMed: 18988683]
4. Beckmann J, Ferguson S, Gebauer M, Luering C, Gasser B, Heini P. Femoroplasty-augmentation of the proximal femur with a composite bone cement-feasibility, biomechanical properties and osteosynthesis potential. Med Eng Phys. 2007; 29:755–764. [PubMed: 17023189]
5. Otake, Y.; Armand, M.; Sadowsky, O.; Armiger, R.; Kutzer, M.; Mears, S.; Kazanzides, P.; Taylor, R. An image-guided femoroplasty system: Development and initial cadaver studies. Proc. SPIE; San Diego, CA. 2010. p. 76250P
6. Basafa, E.; Armiger, R.; Kutzer, M.; Sutter, E.; Mears, S.; Armand, M. Optimizing cement injection in osteoporotic femur augmentation. presented at the 9th Annu. Meeting CAOS-Int; Boston, MA. 2009.
7. Laurentini A. The visual hull concept for silhouette-based image understanding. IEEE Trans Pattern Anal Mach Intell. Feb; 1994 16(2):150–162.
8. Laurentini A. The visual hull concept for silhouette-based image understanding. IEEE Trans Pattern Anal Mach Intell. Feb; 1994 16(2):150–162.
9. Liebschner, M.; Templeton, A. Intra-operative 3-D reconstruction of bone cement boli using X-rays. Patent 7596254. Sep 29. 2009
10. Kutulakos K, Seitz S. A theory of shape by space carving. Int J Comput Vis. 2000; 38:199–218.
11. Caselles V, Kimmel R, Sapiro G. Geodesic active contours. Int J Comput Vis. 1997; 22:61–79.
12. Yezzi A, Soatto S. Stereoscopic segmentation. Int J Comput Vis. 2003; 53:31–43.
13. Kolev K, Brox T, Cremers D. Robust variational segmentation of 3-D objects from multiple views. Pattern Recognit. :688–697.
14. Alvino, C.; Yezzi, A. Tomographic reconstruction of piecewise smooth images. IEEE Computer Vision and Pattern Recognition (CVPR'04); Washingtion, DC. 2004. p. 576-581.
15. Keil A, Vogel J, Lauritsch G, Navab N. Dynamic cone beam reconstruction using a new level set formulation. Med Image Comput Computer-Assisted Intervent (MICCAI). 2009:389–397.
16. Yoon S, Pineda A, Fahrig R. Simultaneous segmentation and reconstruction: A level set method approach for limited view computed tomography. Med Phys. 37:2329. [PubMed: 20527567]
17. Varshney K, Paragios N, Deux J, Kulski A, Raymond R, Hernigou P, Rahmouni A. Post-arthroplasty examination using X-ray images. IEEE Trans Med Imag. Mar; 2009 28(3):469–474.
18. Chen G, Tang J, Leng S. Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. Med Phys. 2008; 35:660. [PubMed: 18383687]
19. Bresson X, Esedoglu S, Vandergheynst P, Thiran J, Osher S. Fast global minimization of the active contour/snake model. J Math Imag Vis. 2007; 28:151–167.
20. Chan T, Vese L. Active contours without edges. IEEE Trans Image Process. Feb; 2001 10(2):266–277. [PubMed: 18249617]
21. Zhuang W, Gopal S, Hebert T. Numerical evaluation of methods for computing tomographic projections. IEEE Trans Nucl Sci. Aug; 2002 41(4):1660–1665.
22. Kruger J, Westermann R. Acceleration techniques for GPU-based volume rendering. IEEE Visualization. 2003:287–292.
23. Spoerk J, Bergmann H, Wanschitz F, Dong S, Birkfellner W. Fast DRR splat rendering using common consumer graphics hardware. Med Phys. 2007; 34:4302. [PubMed: 18072495]
24. Li N, Zhao HX, Cho SH, Choi JG, Kim MH. A fast algorithm for voxel-based deterministic simulation of X-ray imaging. Comput Phys Commun. 2008; 178:518–523.
25. Long Y, Fessler JA, Balter JM. 3-D forward and back-projection for X-ray CT using separable footprints. IEEE Trans Med Imag. Nov; 2010 29(11):1839–1850.
26. Sadowsky, O. PhD dissertation. Dept. Comput. Sci., Johns Hopkins Univ; Baltimore, MD: 2009. Image registration and hybrid volume reconstruction of bone anatomy using a statistical Shape Atlas.
27. Markelj P, Tomazevic D, Likar B, Pernus F. A review of 3-D/2-D registration methods for image-guided interventions. Med Image Anal. 2011 to be published.

28. Daly M, Siewerdsen J, Cho Y, Jaffray D, Irish J. Geometric calibration of a mobile C-arm for intraoperative cone-beam CT. Med Phys. 2008; 35:2124. [PubMed: 18561688]
29. Livyatan H, Yaniv Z, Joskowicz L. Gradient-based 2-D/3-D rigid registration of fluoroscopic X-ray to CT. IEEE Trans Med Imag. Nov; 2003 22(11):1395–1406.
30. Zöllei, L.; Grimson, E.; Norbash, A.; Wells, W. 2-D-3-D rigid registration of X-ray fluoroscopy and CT images using mutual information and sparsely sampled histogram estimators. IEEE Comput. Vis. Pattern Recognit. (CVPR'01); Kauai, HI. 2001. p. 696
31. Sidky EY, Pan X. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. Phys Med Biol. 2008; 53:4777. [PubMed: 18701771]
32. Feldkamp L, Davis L, Kress J. Practical cone-beam algorithm. J Opt Soc Am A. 1984; 1:612–619.
33. Popilock R, Sandrasagaren K, Harris L, Kaser KA. CT artifact recognition for the nuclear technologist. J Nucl Med Technol. 2008; 36:79. [PubMed: 18483138]

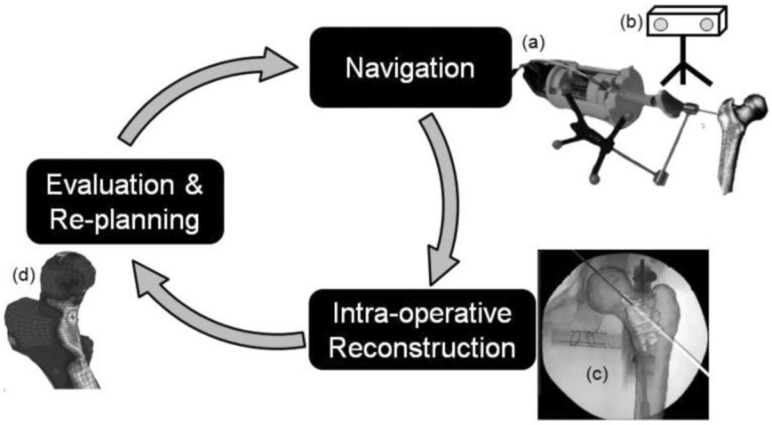

Fig. 1.

System overview showing (a) robotic cement injector, (b) optical tracker, (c) intra-operative reconstruction overlaid on X-ray images, and (d) finite-element analysis of femur.

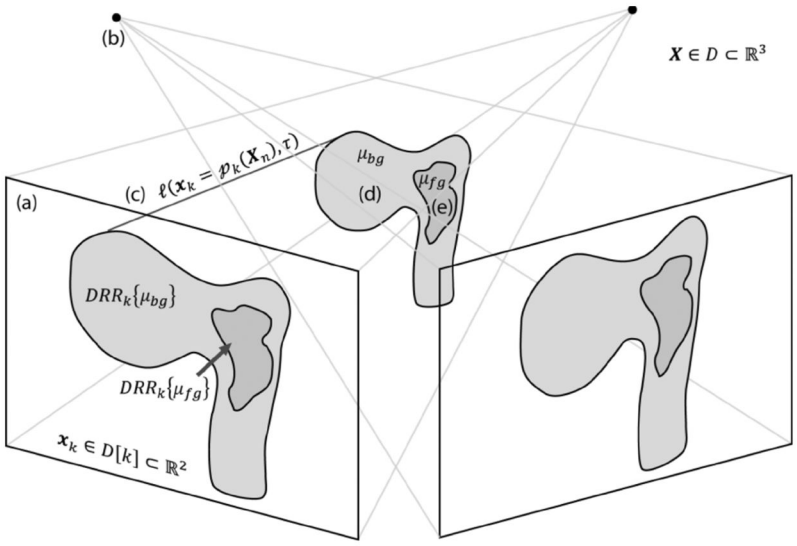

Fig. 2.

Imaging scenario depicting (a) X-ray images, (b) X-ray source, (c) projection lines, (d) background objects, and (e) foreground object.

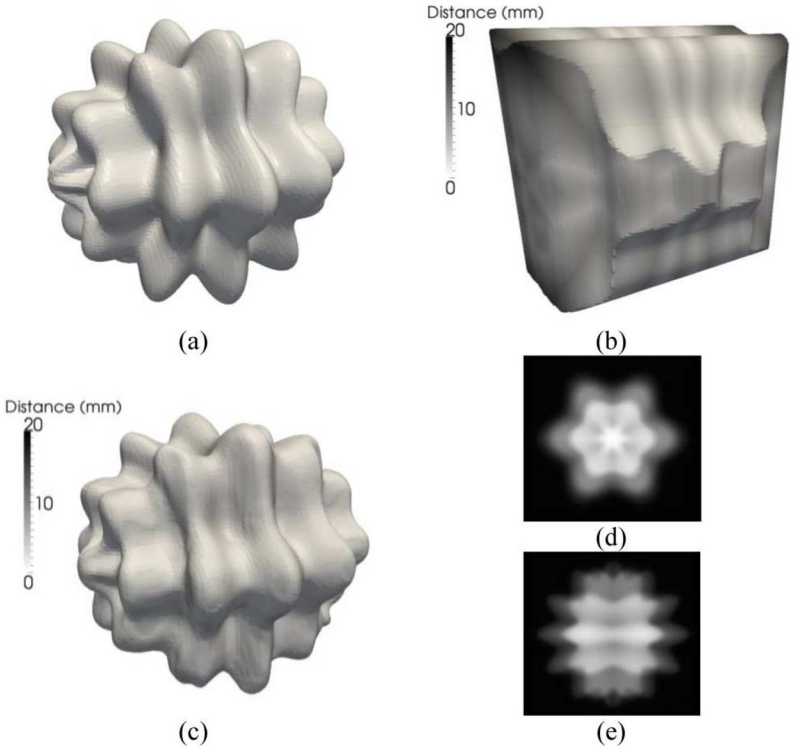

Fig. 3.

(a) Metasphere phantom with (b) corresponding silhouette, (c) SxMAC reconstruction, and (d), (e) DRRs (12).

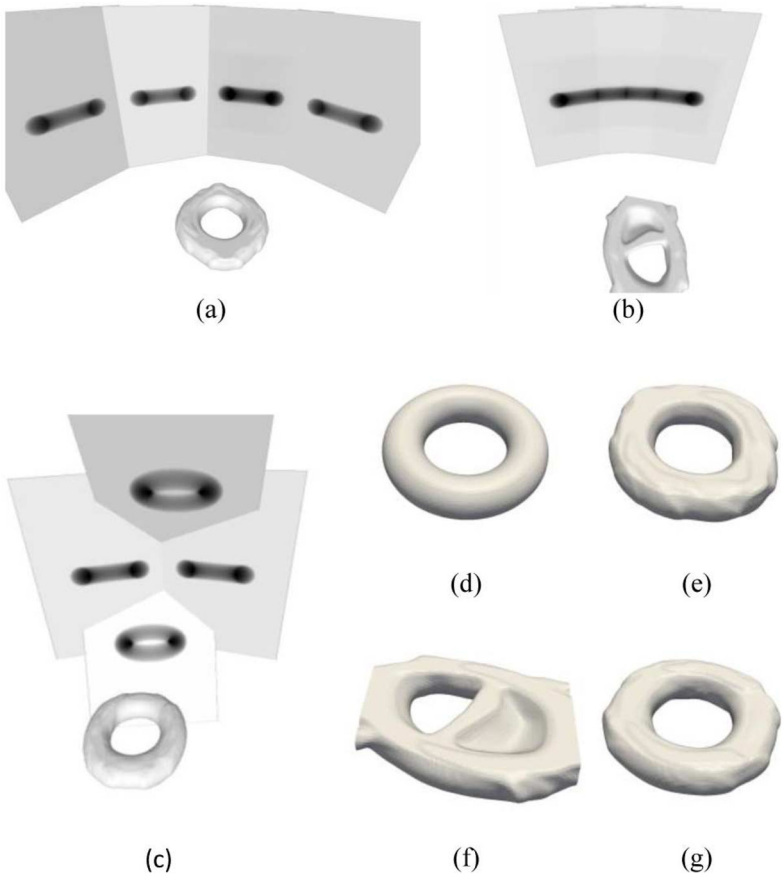

Fig. 4.

SxMAC reconstruction of a torus acquired from a [(a), (e)] 90° arc, [(b), (f)] 30° arc, and wobbled [(c), (g)] 30° arc. X-ray projection images are depicted in their proper pose relative to the reconstructed object. Reconstructions shown in (a), (b), and (c) are juxtaposed to (d) ground truth in (e), (f), and (g).

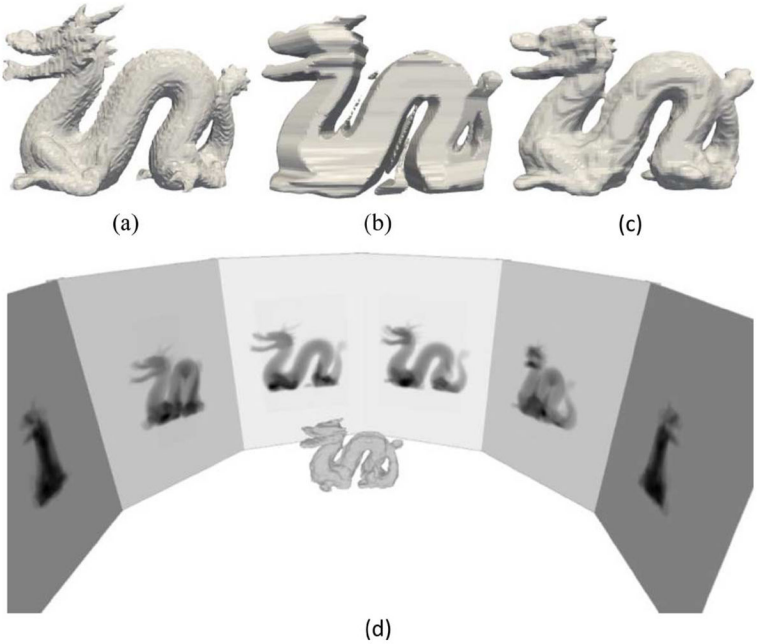

Fig. 5.

(a) Dragon phantom ground truth, (b) visual hull initialization, (c) SxMAC reconstruction, and (d) X-ray images shown in their relative pose.

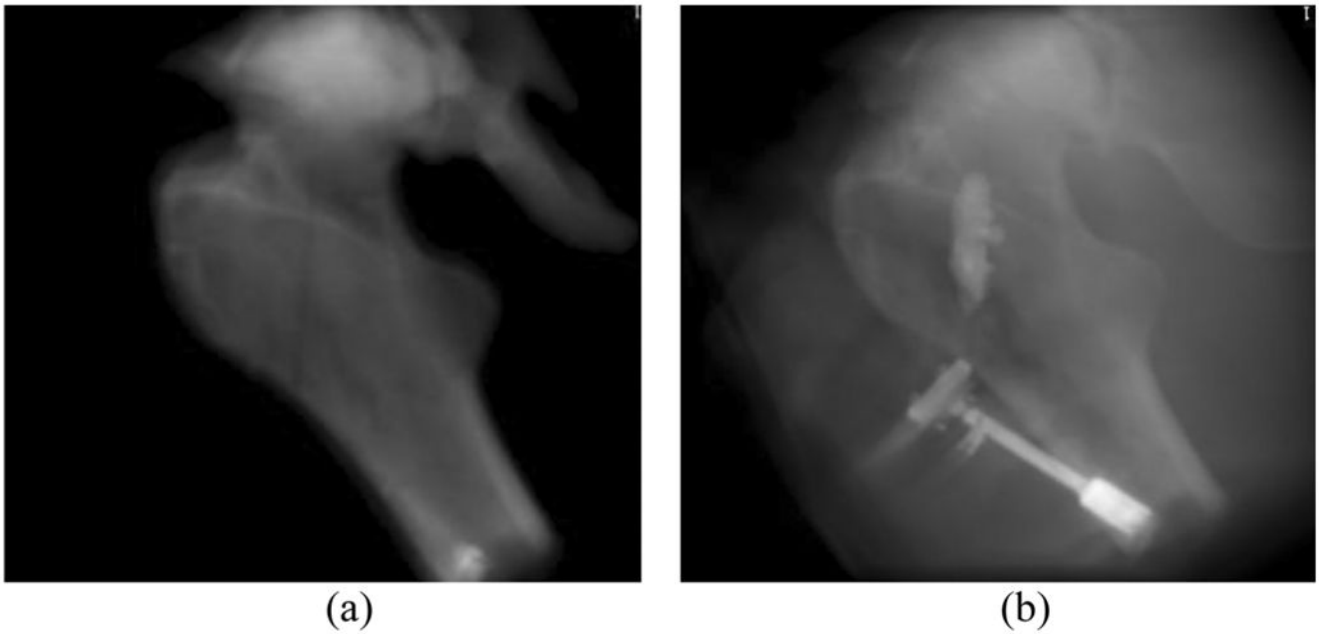

Fig. 6.
Synthesized (a) pre-operative X-ray image of dry femur and (b) postoperative X-ray image with cement attenuation of 1900 HU.

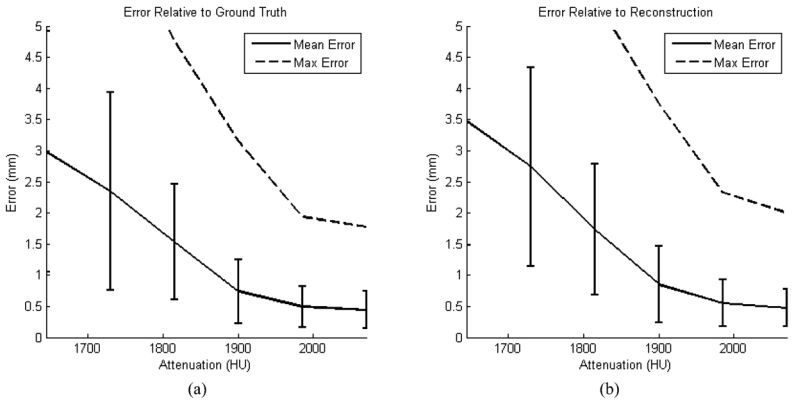

Fig. 7.

Error between ground truth and SxMAC reconstruction for varying bone cement attenuation. Error bars indicate one standard deviation from the mean.

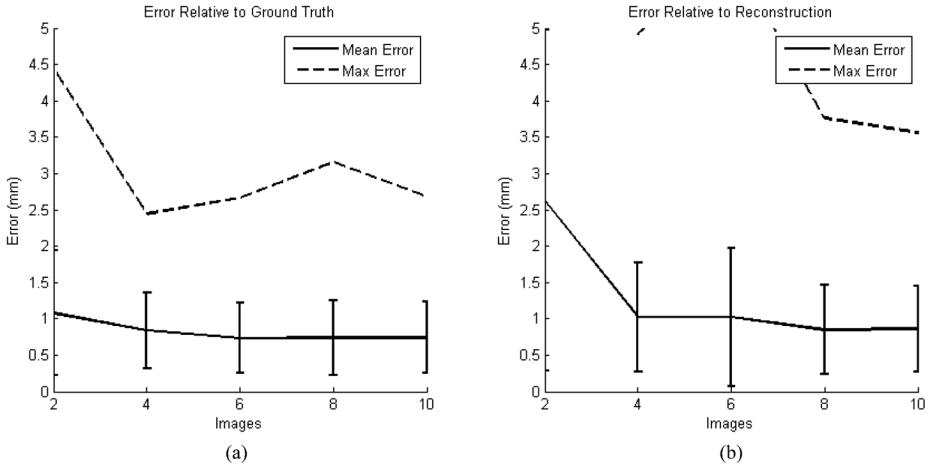

Fig. 8.

Error between ground truth and SxMAC reconstruction for different numbers of synthetic X-ray images.

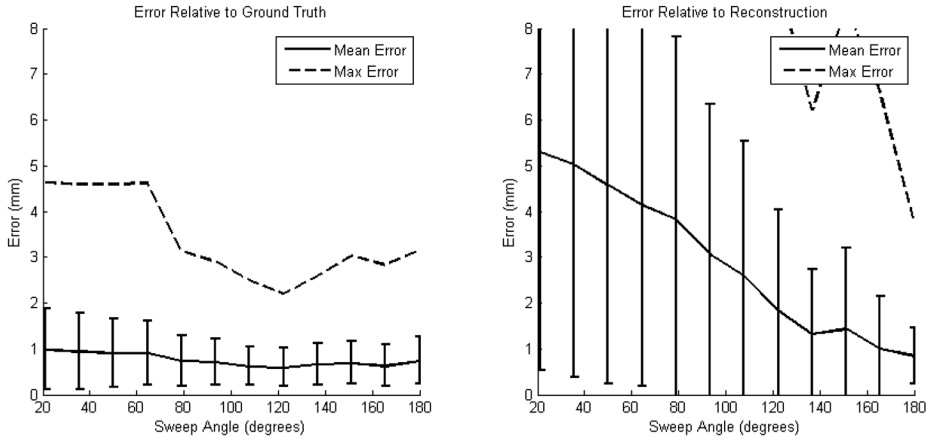

Fig. 9.

Error between ground truth and SxMAC reconstruction from eight images for different sweep angles.

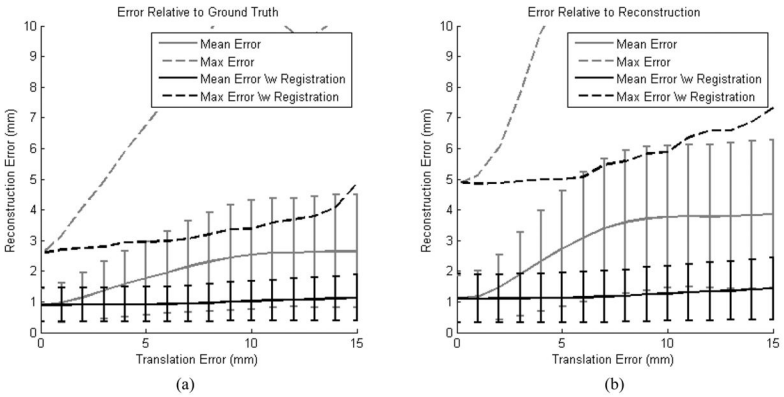

Fig. 10.

Error between ground truth and SxMAC reconstruction from four images and different magnitude translation errors.

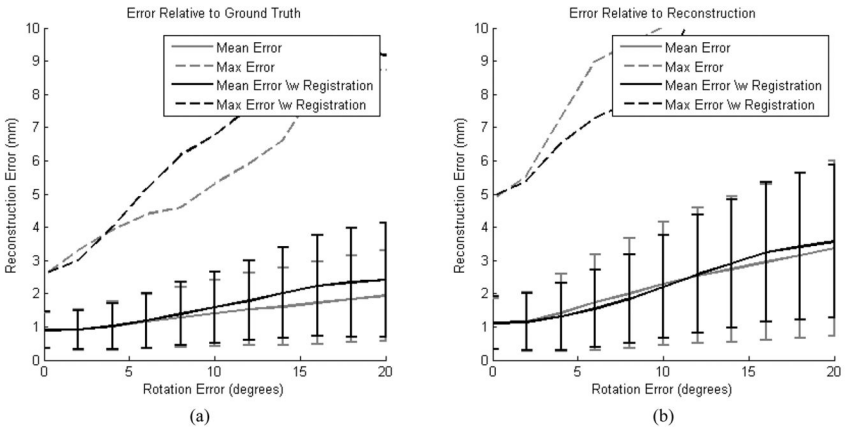

Fig. 11.

Error between ground truth and SxMAC reconstruction from four images and different magnitude rotation errors.

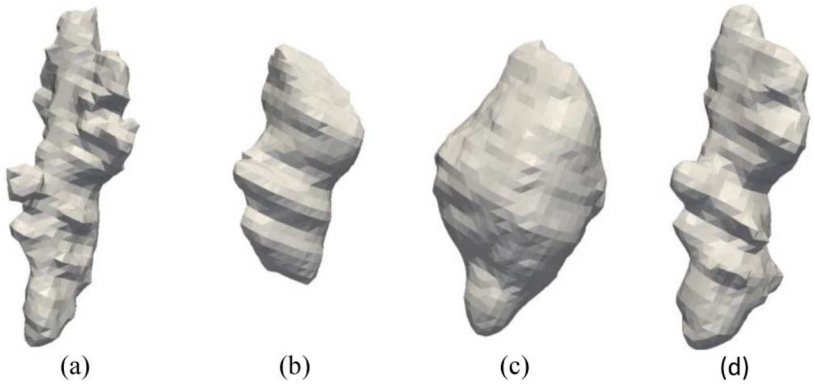

Fig. 12.

(a) Ground truth segmentation and (b) silhouette reconstruction from 8 images with cement attenuation of 1900 HU. SxMAC reconstruction from 8 images with cement attenuation of (c) 1730 HU with soft tissue; (d) 1900 HU with soft tissue.

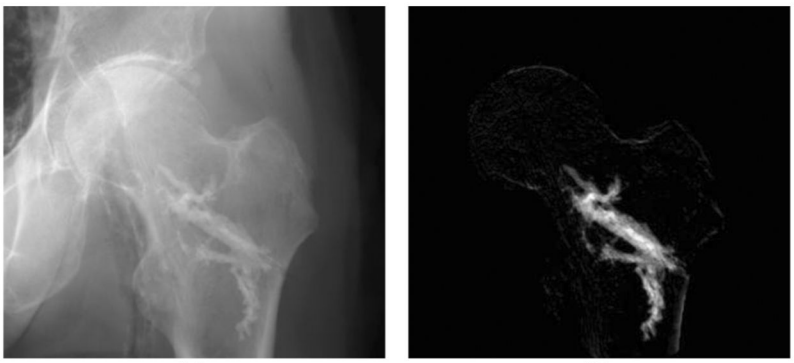

Fig. 13.
(a) Post-injection X-ray image and (b) Pre/Post-injection difference image.

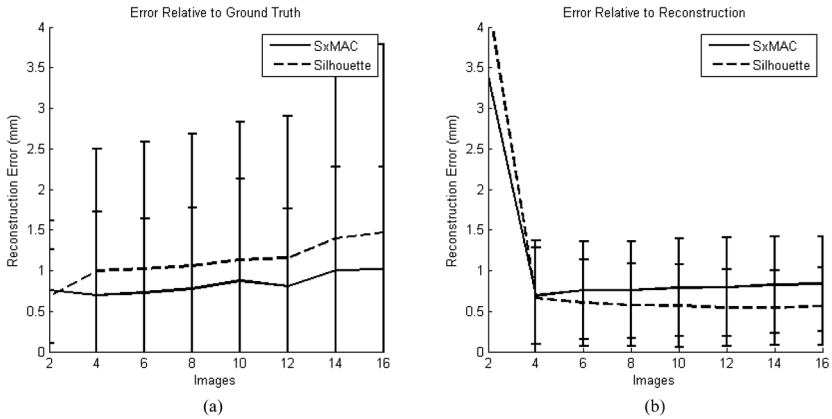

Fig. 14.

Error between ground truth and SxMAC reconstruction from for different numbers of real X-ray images.

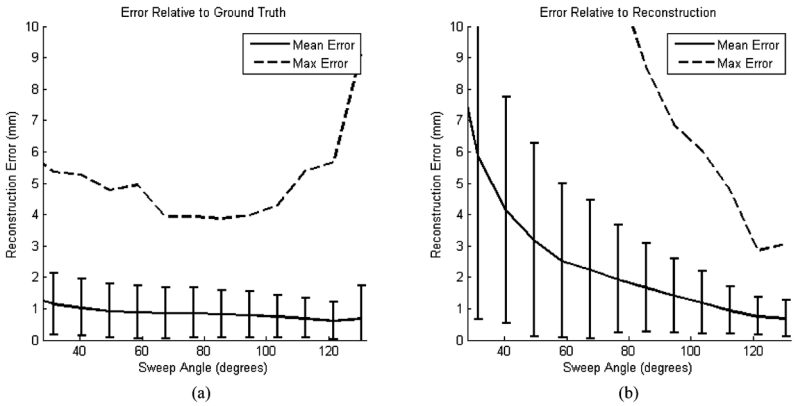

Fig. 15.

Error between ground truth and SxMAC reconstruction from four real X-ray images and different sweep angles.

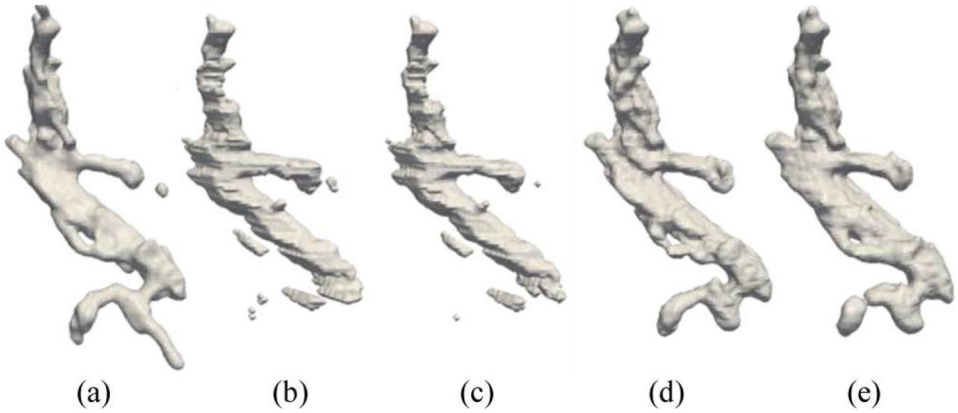

Fig. 16.

(a) Ground truth cement segmentation, (b) silhouette reconstruction from four images, (c) silhouette reconstruction from eight images, (d) SxMAC reconstruction from four images, (e) SxMAC reconstruction from 8 images.

TABLE I

Definition of Terms and Expressions

Symbol	Definition
$\boldsymbol{x}_k \in D[k] \subset \mathbb{R}^2$	Pixel coordinates in the domain of image k.
$\boldsymbol{X} \in D \subset \mathbb{R}^3$	Patient coordinates in 3D.
$\boldsymbol{X}_n \in \mathbb{R}^3$	Position of the n^{th} voxel in an image.
$\mathbb{I}_n \in \mathbb{R}^{1\times N}$	Indicator vector for which the n^{th} column is 1 and all other columns are 0.
$P_{fg}(\mathbf{x}_k)$ and $P_{bg}(\mathbf{x}_k)$	Probability distributions of linear attenuation coefficients in X-ray image k for pixel location $\boldsymbol{x}_k$.
$D_{fg}[k]$ and $D_{bg}[k]$	Foreground (fg) and background (bg) domains in X-ray image k.
$\mu_{fg}(\cdot)$ and $\mu_{bg}(\cdot)$	Linear attenuation coefficients for foreground and background regions in 3D.
$I_k(\cdot)$	The log X-ray intensity image.
$H(\cdot)$	Heaviside function.
$DRR_k\{\mu\}(\mathbf{x})$	Simulated X-ray image k generated from volume μ.
$vec(\boldsymbol{M})$	A column vector representing the elements of matrix $\boldsymbol{M}$.
$\delta(x)=\frac{d}{dx}H(x)$	Dirac delta corresponding to the Heaviside function $H(\cdot)$.
$\mathcal{P}_k(\boldsymbol{X})$	Mapping from patient coordinate space to pixel coordinates c
$\boldsymbol{M}_k \in \mathbb{R}^{M\times N}$	System matrix approximating the X-ray imaging equation.
$\omega_{fg}(\boldsymbol{X})$ and $\omega_{bg}(\boldsymbol{X})$	3D Level set corresponding to the foreground (fg) and background (bg) object. Level sets are negative inside the object and positive outside.
$\psi_k(\mathbf{x}_k)$	Silhouette's level set in image k.
λ_κ	Curvature weight.
$H_{fg}(\boldsymbol{X})$ and $H_{bg}(\boldsymbol{X})$	Volumes indicating the foreground and background.

TABLE II

Parameters for Phantom Experiments

Parameter	Value
Source Detector Distance	900 mm
Detector Dimensions	288 mm × 288 mm
Pixel Size	0.45 mm
λ_{fg}	1
λ_{bg}	1

TABLE III

Flat Panel C-Arm Specification

Parameter	Value
Manufacturer	Siemens Medical Solutions
Type	Mobile isocentric flat-panel cone-beam CT
Scan time	128 sec.
Voltage	120 kVp
Current	5.2 mAs
Pixel size	0.388 mm × 0.388 mm