

Robust automatic rigid registration of MRI and X-ray using external fiducial markers for XFM-guided interventional procedures

Ashvin K. George^{a)}

Translational Medicine Branch, Division of Intramural Research, National Heart, Lung, and Blood Institute, National Institutes of Health, Bethesda, Maryland 20892-1538

Merdim Sonmez

Translational Medicine Branch, Division of Intramural Research, National Heart, Lung, and Blood Institute, National Institutes of Health, Bethesda, Maryland 20892-1538 and Biomedical Engineering Institute, Bogazici University, Istanbul 34684, Turkey

Robert J. Lederman and Anthony Z. Faranesh^{b)}

Translational Medicine Branch, Division of Intramural Research, National Heart, Lung, and Blood Institute, National Institutes of Health, Bethesda, Maryland 20892-1538

(Received 22 June 2010; revised 24 September 2010; accepted for publication 11 November 2010; published 16 December 2010)

Purpose: In *X-ray fused with MRI*, previously gathered roadmap MRI volume images are overlaid on live X-ray fluoroscopy images to help guide the clinician during an interventional procedure. The incorporation of MRI data allows for the visualization of soft tissue that is poorly visualized under X-ray. The widespread clinical use of this technique will require fully automating as many components as possible. While previous use of this method has required time-consuming manual intervention to register the two modalities, in this article, the authors present a fully automatic rigid-body registration method.

Methods: External fiducial markers that are visible under these two complimentary imaging modalities were used to register the X-ray images with the roadmap MR images. The method has three components: (a) The identification of the 3D locations of the markers from a full 3D MR volume, (b) the identification of the 3D locations of the markers from a small number of 2D X-ray fluoroscopy images, and (c) finding the rigid-body transformation that registers the two point sets in the two modalities. For part (a), the localization of the markers from MR data, the MR volume image was thresholded, connected voxels were segmented and labeled, and the centroids of the connected components were computed. For part (b), the X-ray projection images, produced by an image intensifier, were first corrected for distortions. Binary mask images of the markers were created from the distortion-corrected X-ray projection images by applying edge detection, pattern recognition, and image morphological operations. The markers were localized in the X-ray frame using an iterative backprojection-based method which segments voxels in the volume of interest, discards false positives based on the previously computed edge-detected projections, and calculates the locations of the true markers as the centroids of the clusters of voxels that remain. For part (c), a variant of the iterative closest point method was used to find correspondences between and register the two sets of points computed from MR and X-ray data. This knowledge of the correspondence between the two point sets was used to refine, first, the X-ray marker localization and then the total rigid-body registration between modalities. The rigid-body registration was used to overlay the roadmap MR image onto the X-ray fluoroscopy projections.

Results: In 35 separate experiments, the markers were correctly registered to each other in 100% of the cases. When half the number of X-ray projections was used (10 X-ray projections instead of 20), the markers were correctly registered in all 35 experiments. The method was also successful in all 35 experiments when the number of markers was (retrospectively) halved (from 16 to 8). The target registration error was computed in a phantom experiment to be less than 2.4 mm. In two *in vivo* experiments, targets (interventional devices with pointlike metallic structures) inside the heart were successfully registered between the two modalities.

Conclusions: The method presented can be used to automatically register a roadmap MR image to X-ray fluoroscopy using fiducial markers and as few as ten X-ray projections.

[DOI: [10.1118/1.3523621](https://doi.org/10.1118/1.3523621)]

Key words: X-ray, MRI, fiducial markers, localization, registration, fusion

I. INTRODUCTION

In several areas of modern medicine, there has been interest in combining multiple, complementary imaging modalities in order to improve the diagnoses of diseases and the delivery of therapies. In oncology,¹ computed tomography (CT) has long been combined with positron emission tomography (PET) for improved diagnosis and monitoring of cancer. In image-guided radiotherapy,² multiple modalities are combined for better treatment planning. Morphological data from X-ray and perfusion data from PET have been combined³ to improve diagnoses in cardiology. 3D CT and MR images have been combined with electroanatomical maps⁴ of the heart to guide electrophysiological procedures. Ultrasound has been combined with MR (Ref. 5) to guide prostate biopsies. CT images have been fused with X-ray fluoroscopy⁶⁻⁸ to guide various interventional procedures.

In this paper, we are concerned with a technique known as X-ray fused with MRI (XFM),⁹⁻¹⁵ also known as XMR, in which live X-ray fluoroscopy (XF) projection images are displayed along with a previously gathered roadmap 3D MR image. In this paper, we present a novel, fully automatic method to register the two modalities. In an interventional procedure, the clinician attempts to avoid invasive surgery by using a small incision to insert and maneuver an instrument, such as a catheter, inside a patient. The ability to accurately visualize the interventional instrument in relation to adjacent tissue is essential to the success and safety of interventional procedures.

I.A. XFM guidance for interventional procedures

XF, the most commonly used interventional imaging modality, offers projection images of high temporal and spatial resolution. The contrast in an XF image depends on the opacity of the object to X-ray, in particular the attenuation coefficient of the substance. While XF provides excellent images of a catheter or blood that contains an iodinated contrast agent, most kinds of human tissue have very similar attenuation coefficients and are very poorly visualized under XF. Unlike XF, MRI can produce images that provide good soft-tissue contrast, with the nature of the contrast varied depending on the structures of interest to the clinician. The disadvantages of MRI are its slow acquisition time and the inability to use commonly used interventional tools with the associated magnetic fields and radiofrequency pulses.

In XFM, these two complementary imaging modalities are combined by displaying a previously gathered roadmap MR image along with the live XF projection images. This allows the interventionist to visualize soft-tissue structures that are invisible under XF. A previous work¹³ has demonstrated that this enhanced image guidance can enable procedure simplification and reduction of radiation exposure. Attempts have been made to more closely combine the XF and MR systems by integrating an X-ray system into an open “double-doughnut” MR system¹⁶ or placing a flat-panel XF system adjacent to a closed bore magnet.¹⁷ In this paper, we describe a method that can be used with the conventional X-ray-MR suites in which the patient has to be physically

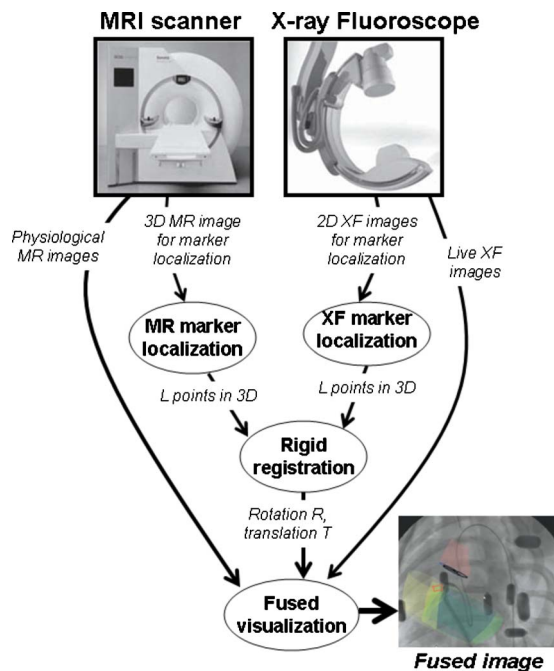


FIG. 1. Outline of XFM system.

moved between the two modalities and which requires no new hardware. Furthermore, a combined X-ray-MR suite is not necessary; our method can be used even when the patient is transferred between rooms.

This requires the registration of the roadmap MR image with the live XF images. A few different methods have been proposed. The method described by Rhode *et al.*,^{14,15} which optically tracks sensors placed on the X-ray table and C-arm, does not allow for the tracking of patient motion relative to the table and requires an extra optical tracking system. The disadvantage of using intrinsic anatomical landmarks to perform the registration is that the landmarks (such as ribs) are not necessarily visible and automatically identifiable under both MR and X-ray. Tomokowiak *et al.*¹⁸ propose using a left ventriculogram (filling the left ventricle with contrast agent under XF) in order to register the two modalities, but their method requires the injecting of contrast agent. Rhode *et al.*¹⁹ presented a promising method of combining the optical tracking system with external fiducial markers, similar to the markers we use, that are MR and X-ray visible. In their work, which used manual segmentation of the markers, they report target registration errors of a fraction of a mm, in a phantom, and of less than 6 mm *in vivo*.

For registration we use a set of stationary external fiducial markers (Beekley, Bristol, CT) that are visible under both imaging modalities (a belt of plastic beads filled with a mixture of gadolinium and iodinated contrast agent) attached to the patient while undergoing both MR and XF imaging. The 3D locations of these beads are identified separately under MR and XF and then the two sets of beads are registered to each other in order that the roadmap MR image may be correctly displayed on the live XF visualization (as seen in the lower right hand corner of Fig. 1). There are several

advantages of this method to those described above. It works even if the patient moves relative to the table, unlike in some previous methods;^{14,15} the belt of beads are noninvasive and do not require the injection of contrast agent, unlike in other methods;¹⁸ the system does not require new hardware and can be used with older image intensifier systems, unlike in some methods;^{15,16} and can be used in cases when the two modalities are in different locations and not in a single suite, unlike in many methods.^{14–17}

This procedure requires three distinct computational tasks: (a) Identifying the 3D locations of the beads from a full 3D MR volume, (b) identifying the 3D locations of the beads from a small number of 2D XF images, and (c) finding a rigid transformation (rotation and translation) that relates the MR and X-ray coordinate systems. Identifying the 3D location of the beads (henceforth referred to as “localizing the beads”) from MR data requires segmenting a full 3D volume. The MRI sequence (T_1 -weighted, gradient echo) that produces this image was chosen to provide a high contrast between the Gd-filled beads and the rest of the image. Consequently localizing the beads is a relatively straightforward image processing and segmentation task, described in Sec. II B. On the other hand, localizing the beads from a small number of XF images, described in Sec. II C, is more challenging because the data is limited and because of the presence of other objects in the XF images. Finally, the rigid transformation between the two sets of beads is achieved by a version of the well-known iterative closest point (ICP) algorithm²⁰ and is described in Sec. II D.

While previous publications^{9,12,13,21} have presented the XFM system as applied to endomyocardial injections,⁹ mitral cerclage annuloplasty,²¹ and the closure of ventricular septal defects¹³ (VSDs) in a swine model, none have presented the computational method that achieves the automatic registration. In early works,^{9,21} the procedure was minimally automated and required extensive, time-consuming, interaction by a skilled human operator to localize the MR and X-ray markers by manual examination of the data. A more recent work¹³ used a promising predecessor^{22,23} of the method that we present in this paper. That method achieved automatic MR and XF localization of the markers but suffered a few shortcomings including the restriction that the angular spacing of the XF projections be small ($\sim 3^\circ$) and was therefore not conducive to using fewer X-ray projections. The method we present in this paper uses an edge-detection-based strategy for XF projection segmentation that is similar to that used by its predecessor but is otherwise completely different. These differences include the use of a more robust XF projection segmentation method, a novel backprojection-based XF localization method that does not require the XF projections to be closely spaced in angle, and the use of the well-known iterative closest point algorithm for the registration (which replaces a more heuristic prior method).

I.B. Localizing fiducial markers from a few X-ray projection images

A problem similar to bead localization from XF images has been addressed in the area of prostate brachytherapy seed

localization. The problem in brachytherapy is to identify the 3D locations of radioactive seeds that have been implanted in the patient from a few XF images. A range of algorithms^{24–30} has been described to solve the problem. In all these methods, each XF image is first segmented to identify the pixels that represent the seeds, just as we do. The algorithms to then identify the 3D locations of the seeds from their segmented projection images may be divided into two classes. The first class of algorithm^{24–26} explicitly computes a correspondence (or matching) between seeds (i.e., they find which seed shadows on one projection image correspond to which other seed shadows on a different projection image). These kinds of methods are referred to in the brachytherapy literature as “seed-matching” approaches. The second class of algorithms^{27–30} avoids this explicit matching step and directly finds the 3D locations of the seeds, either using a backprojection-based method^{28,30} or setting up the problem as the minimization of a cost function.^{27,29} These *nonmatching* approaches have the important advantage of avoiding the difficult task of reliably finding correspondences from (sometimes erroneous, as explained below) segmented projection images.

There are differences between the localization problems in XFM and brachytherapy that make none of the above methods directly applicable. All the brachytherapy methods, while they make allowances for spurious and hidden/occluded seeds, assume that every true seed is identified in every projection. This assumption does not necessarily hold true in XFM for two reasons. First, unlike in brachytherapy where the seeds tend to be restricted to a small volume, inside the patient and close to the origin of the imaging system, the beads in XFM lie on a belt external to the patient and each projection image does not contain all the beads. Second, the segmentation of the projection images in XFM tend to be less perfect than in brachytherapy and therefore the segmented XF images can overidentify and underidentify beads. The imperfect segmentation is because of the presence of other structures (such as ribs, tissue, and interventional devices) and because, in order to reduce dosage, a lower power X-ray (70 kVp/102 mA) is used which results in reduced contrast. Consequently, in XFM, a bead may be identified in as few as two projection images.

In contrast with these aspects that are more difficult in XFM than in brachytherapy, the XFM problem is easier in some aspects. Specifically, the beads in XFM are larger than brachytherapy seeds, making the problem less sensitive to errors in the pose of the C-Arm, the number of markers is smaller in XFM than in brachytherapy, and because the belt of beads is fixed and taut, we have a minimum bound on the distances between beads that we can use to eliminate false positives.

Our algorithm falls into the second class of nonmatching methods. We first segment the XF projection images to produce bead masks to identify those pixels which represent beads and those which do not. We backproject these bead masks onto the volume of interest. We then process this backprojected volume to identify the groups of voxels that represent beads. The key problem is to identify, among the

large group of candidate beads, which ones are false positives. We explain this iterative method in Sec. II C 4. This general approach of discarding false positives from a large number of candidates is one that is also employed in the work of Lee *et al.*²⁸ While the general approach of their method is similar to ours, their method is not directly applicable to XFM for two reasons. First, their assumption that every seed is identified in every projection does not hold true for XFM. Second, the criteria that they use to eliminate false positives are based on an empirical observation of brachytherapy seeds, while the criteria and resultant method we use are based on the nature of beads and their XF projections.

II. METHODS

II.A. Outline of XFM system

The outline of the XFM imaging system is shown in Fig. 1. Before the interventional procedure, the belt of beads (numbering no more than $L=16$) is strapped around the torso of the patient and MR images are collected. In our experiments, the beads were distributed with a cluster of eight beads on the anterior portion of the belt and a cluster of eight beads on the posterior portion. For the purpose of bead localization, a T_1 -weighted gradient-echo MR sequence is used to gather a 3D image that displays a strong contrast between the gadolinium inside the beads and surrounding tissue. Other 2D and 3D physiological MR images are also gathered for later fusion during the procedure. The patient is then moved to the X-ray fluoroscope and a small number ($< \sim 20$) of XF images are collected for the purpose of bead localization. The automatic localization of the beads is performed on both the MR and XF data. Each bead localization produces a set of 3D coordinates of the centers of the beads. These two sets of 3D coordinates are then fed into a registration algorithm which computes the rigid-body transformation (the rotation and translation) that relates the two sets of markers. This rigid-body transformation is used to fuse the MR images onto the live XF image during the procedure. An example of a fused image is shown at the bottom right of the figure. The contours from the MR image are overlaid onto an XF projection image. Also visible on the XFM image are the beads, which appear dark as they contain iodinated contrast agent.

II.B. Automatic localization of markers in MR

The 3D MR images used for marker localization contain the whole volume of interest. A typical image is $128 \times 128 \times 144$ voxels with a resolution of $3.1 \times 3.1 \times 1.6$ mm³. As seen in Fig. 2, which shows two maximum-intensity projections (MIPs) of one such volume image, the contrast between the voxels representing the beads (light regions) and voxels representing other tissue is large as a result of the use of a T_1 -weighted sequence (a 3D gradient-echo sequence with echo time $TE=0.94$ ms, repetition time $TR=2.6$ ms, and flip angle $=17^\circ$).

The markers are identified by first thresholding the full 3D volume image to make a binary image which is zero




FIG. 2. Bead localization in MR. (a) Coronal MIP and (b) axial MIP. The white x's denote the locations of the markers found by the MR marker localization algorithm.

everywhere except for voxels whose value is greater than a fraction of the value of the maximally valued voxel (we use a threshold of 0.5). Nonzero voxels in the image that are connected to each other are grouped together. We define two voxels to be connected if they are both nonzero and one voxel is one of the 26 neighbors of the other voxel³¹ (p. 67). Groups that contain more than S voxels are discarded and then the remaining groups are sorted on the basis of the maximally valued voxel in each group. The best L remaining voxel groups are considered to be real beads. In our system, we use no more than 16 plastic beads of a finite volume so we choose $L=16$ and S corresponding to 0.9 cc. The centroid of each voxel group is computed to be the center of the bead.

To refine the values of this centroid, a second pass through the L identified beads is performed. The problem with the initial estimate is that the threshold, being conservative, sometimes segments beads partially. Furthermore, the value of the highest-valued voxel is not the same in every bead (because of the inconsistencies in the concentration of the Gd contrast agent or the sensitivity of the MR receiver coils). In this second pass, all the L beads are reconsidered and thresholded using a lower threshold (0.2) relative to the maximally valued voxel of the specific bead, not the maximally valued voxel of the whole volume. The black dots at the centers of the white circles in Fig. 2 represent the centers of the beads. A visual inspection suggests that all the markers have been correctly identified in Fig. 2.

II.C. Automatic localization of markers in X-ray

II.C.1. Segmenting beads in X-ray fluoroscopy images

Each XF projection image is first distortion-corrected, by the method described by Gutiérrez *et al.*,¹¹ to produce an image as shown in Fig. 3(a). This image is segmented to produce a bead mask, as shown in Fig. 3(c), an image that is zero everywhere except for the pixels corresponding to the beads. We refer to a contiguous nonzero region in a bead mask as a *bead region*. Note that while most bead regions correspond to the projection of a single bead, sometimes they correspond to more than one bead whose projections overlap (identified by the dashed circles). Single-bead and multibead regions can be distinguished on the basis of the convex hull of the bead region. It is easy to see that the area of a single-bead region would be equal to the area of its convex hull,




FIG. 3. Segmenting beads: (a) XF projection image. (b) Corresponding edge mask. (c) Bead mask extracted from edge mask. Dashed circles identify overlapping beads.

while the area of a multibead region would be less than the area of its convex hull. Multibead regions are identified as those for which the ratio of its area to the area of its convex hull is less than a threshold (that we set to be 0.9).

To create a bead mask from an XF projection image, a Canny edge detector³² [followed by skeletonization (Ref. 31, p. 543)] is first applied to each XF image to produce a binary-valued edge mask, as shown in Fig. 3(b). This edge mask is zero everywhere except for pixels that lie on edges where it is equal to 1. As shown in Fig. 3(b), the edge detector returns edges of several features in the XF image such as the ribs and the shadow of the heart. The distinctive feature of the beads is that their projections have a roughly ellipselike shape regardless of the orientation, or view angle, of the XF projection. Particularly, they consist of two parallel straight edges with rounded corners. In the best cases, the edge detector finds the complete, continuous, outline of the bead, but more often than not, the edge detector finds an incomplete edge that contains at least one corner of the outline. We iterate through every edge found by the edge detector and identify those that might represent the corner of a bead.

The corners of beads are identified as those edges that (i) contain a rounded corner (defined as changing direction by 90° within a sufficiently small distance; we use 40 pixel units) and (ii) have lower-valued pixels on the inside of the corner (because the iodine in the bead absorbs x rays). Figure 4(a) plots the direction vs distance of three different edges (cases I, II, and III). The XF images and the edge masks corresponding to the three cases are shown on the right. Both




FIG. 4. Identifying bead corner (a) Direction vs distance. (b) Case I: Good. (c) Case II: Nonellipsoidal. (d) Case III: Ellipsoidal but not a bead shadow. The thick line represents the segment of the trace that is identified as a bead corner.




FIG. 5. Fitting a bead outline to an edge. (a) Edge map. The thick line represents the bead corner that has been identified. (b) The two straight edges adjacent to the bead corner are identified. (c) The predicted bead (dashed line) outline is formed by extending the bead corners and edges in a polygonal manner. (d) The bead outline is formed by adjusting the bead-outline prediction by the closest edge pixels. The closed bead outline is filled to form the bead mask.

cases I and III qualify as bead corners by the first test. The points represented by the thicker line correspond to points on the rounded corner. Case II does not pass the bead-corner test and is correctly identified as not being a bead. Case III fails the pixel-value test and is therefore correctly rejected as a bead corner. Identifying beads by the change in direction of the edge allows for the segmentation to be robust to rotations of the bead outline within the XF image. Furthermore, identifying them purely by their distinctive corners makes the algorithm robust to cases when the bead projections overlap each other.

If an edge is identified to contain a corner of a bead outline, the bead mask corresponding to that bead is found in a two step template-matching process. First, the shape of the bead outline is predicted by using knowledge of the edge under consideration. Then, this prediction is corrected by the *other* edges in its neighborhood. In particular, if the predicted bead outline does not lie on an edge pixel, it is adjusted to lie on the closest edge pixel that lies on a line segment perpendicular to the predicted bead outline.

This process of fitting a bead outline to an edge is illustrated in Fig. 5. The edge identified as a bead corner is shown in Fig. 5(a). The bead corner is represented by the thicker line. A rough template of the bead is formed using the two straight edges [as shown in Fig. 5(b)] of the bead outline *adjacent* to the bead corner and closing the shape in a simple polygonal manner [as shown in Fig. 5(c)]. This predicted bead outline is then adjusted to lie on the closest edge pixel that lies on a line segment perpendicular to the predicted bead outline as shown in Fig. 5(d). The *corrected* bead outline is then filled³¹ (p. 535) to make the bead mask.

II.C.2. Backprojecting the XF projections

In anticipation of the second part of our algorithm, we describe the efficient tomographic backprojection of sparse

projection images (bead masks). The geometry of an XF projection (see Figs. 1–3 in Ref. 33) oriented at view angle $\vec{\theta}$ is encapsulated by a projection matrix $P_{\vec{\theta}}$. Note that because the C-arm of the XF system is able to rotate in 3D, the view angle $\vec{\theta}$ is a scalar pair consisting of a primary and secondary angle. As described by Shechter *et al.*,³³ the matrix $P_{\vec{\theta}}$ depends on the primary and secondary view angles, the distance of the source to the isocenter of the XF system, the size and resolution of the XF intensifier that detects the transmitted x rays, and the distance of the source to the intensifier. In addition, to account for an additional residual in-plane rotation and translation, the projection matrix is modified based on calibration data.¹⁰ A 3D point (q_x, q_y, q_z) in space is projected onto the 2D point (p_u, p_v) in the XF projection image. The point (p_u, p_v) is derived as follows. The vector $\vec{r} = (r_u, r_v, r_w)$ is computed from the vector $\vec{q} = (q_x, q_y, q_z, 1)$ by matrix multiplication: $\vec{r} = P_{\vec{\theta}}\vec{q}$. Then the point (p_u, p_v) in the XF projection image is computed as $p_u = r_u/r_w$ and $p_v = r_v/r_w$. The backprojection of a single XF projection image $g(p_u, p_v)$ onto a 3D volume $f(q_x, q_y, q_z)$ can be expressed as $f(q_x, q_y, q_z) = g(p_u, p_v)$, where \vec{p} is computed from \vec{q} as explained above.

Backprojecting L projections onto a 3D image volume with N voxels, using nearest neighbor interpolation, costs c_L arithmetic operations per voxel, i.e., $c_L LN$ total operations. But since the bead masks that we are backprojecting are very sparse, zero everywhere except for in the bead regions, this cost is reduced by using ray-driven backprojection³⁴ and only operating on the bead regions of the XF projections. If the fraction of bead-region pixels to the total number of pixels of each bead mask is α , then the cost of backprojecting L bead masks is $\alpha c_L LN$.

II.C.3. Identifying the volume of interest

Since the cost of backprojection is proportional to the number of voxels of the backprojected 3D volume image, it is computationally beneficial to use a small number of voxels. The number of voxels is, in turn, inversely proportional to the voxel size (which affects how finely the bead locations can be resolved) and proportional to the size of the volume of interest. A voxel size of 3 mm is found to be adequate in resolving bead locations using backprojection. In Sec. II E, we describe a method to refine the estimate of the bead locations, so the 3 mm voxel size used here does not limit the resolution of the final estimate of the bead locations.

In order to reduce the 3D volume of interest, a rough estimate of the volume occupied by the beads is made by the following simple multiresolution algorithm. First, a FOV-mask image is created from each bead-mask image to capture the vertical extent of the bead mask. In particular, a FOV-mask image is zero-valued on the rows above and below the most extremal bead regions of the corresponding bead mask and 1 everywhere else. These FOV-mask images are then backprojected onto a very sparse, large 3D volume centered at the origin of the system. A rough estimate of this volume is made by finding a cuboidal region in this backprojected volume outside of which the voxel values fall be-

low a threshold (we use half the number of projections). This gives a rough estimate of the volume of interest. To reduce the volume of interest a little more, we backproject the same projections onto a finer 3D grid within this cube-shaped region. We then find a smaller cube-shaped region within this backprojected volume outside of which the voxel values fall below the threshold.

II.C.4. Identifying beads in 3D volume

As described previously, we compute a bead mask for each XF projection. Since the field of view of the XF system is limited, each XF projection image does not contain all 16 beads. Furthermore, the segmentation of beads is often imperfect and the algorithm both overidentifies (i.e., contains false positives) and underidentifies (i.e., contains false negatives) beads in XF projection images.

A naïve approach to identifying the location of the beads would be to backproject all the bead masks onto the volume of interest, threshold the volume, and compute the centroids of connected components in this volume. While this works well when the number of XF projections is large and all the bead projections are correctly segmented, in the case of a small number of XF projections that have errors in their segmentation, it tends to produce false negatives (for high thresholds) and false positives (for low thresholds).

In order to improve the performance, we employ an iterative approach that exploits additional constraints. The constraints are based on the following four observations about the voxels of the backprojected volume that correspond to true beads.

First, we know that a voxel corresponding to a true bead would likely have a large magnitude. The magnitude of a voxel is equal to the number of bead regions that backproject onto it. This is equivalent to the number of projections for which the voxel lies on a ray between the X-ray source and a bead region (which we will henceforth refer to as *bead-region-rays*). It is unlikely that a nonbead voxel will lie on bead-region-rays from multiple projections. Conversely, if a voxel does not lie on a bead-region-ray from multiple projections, it is very unlikely to be a bead.

Furthermore, if we eliminate a bead region whose corresponding beads have been correctly identified, the likelihood of false positives decreases. This is because as the bead masks are made sparser, the likelihood that a nonbead voxel lies on bead-region-rays from several views decreases.

Second, we know that the beads have a known volume and are not arbitrarily small. In particular, if the backprojected volume contains a small number of connected voxels with a large magnitude, it is unlikely to be a bead. Third, we know that the beads are not arbitrarily close to each other (because they are affixed to a belt in a known configuration). When backprojecting bead masks from a limited range of view angles, the thresholded image sometimes contains a false-positive group of voxels that is close to a true bead. This phenomenon has been observed²⁸ in the brachytherapy

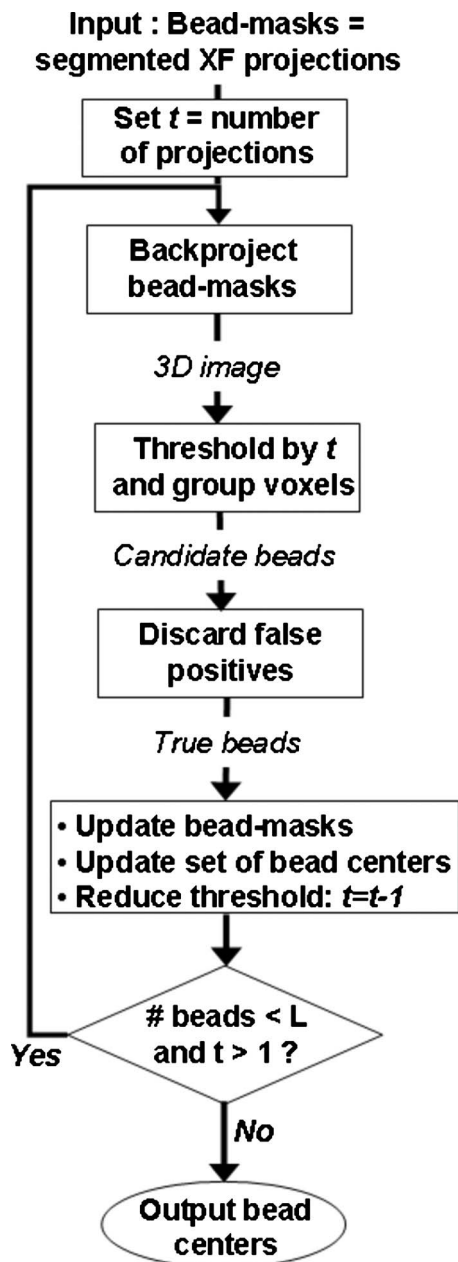


FIG. 6. XF bead finding iterative algorithm.

problem. These false positives can be eliminated by ensuring that the identified beads are no less than a certain distance from each other (we use 18 mm).

The fourth observation is about the nature of imperfections when segmenting the XF projection images to make the bead masks. Even when a bead has not been identified on a bead mask, the outline of its projection is always partially identified by the edge detector in every view. So if the 3D coordinate of a true bead is reprojected onto the edge mask of any XF view, it will lie within the partially detected outline of a bead. This observation is also used to eliminate false positives.

The outline of our iterative algorithm is shown in Fig. 6. Initially, a 3D volume image is created by backprojecting all the bead masks. The volume image is thresholded using a




FIG. 7. Scoring of candidate beads. (a) The template bead outline, denoted by open circles, is overlaid onto the edge mask. (b) The corrected bead outline, denoted by the black dots.

high threshold (observation I) and connected clusters of voxels that are of sufficient volume (we use 60 mm³) (observation II) are identified using image morphological operations³¹ (Chapter 9). The centroids of these connected clusters of voxels are computed to get a list of 3D coordinates of candidate beads.

False positives are eliminated from this list of candidate beads using two tests. In the first test, candidate beads that are too close to previously correctly identified beads (observation III) are eliminated. In the second test, candidate beads whose coordinates when reprojected onto any XF view do *not* lie within the outline of a projected bead (observation IV) are eliminated.

The details of this second test, illustrated in Fig. 7, are as follows. The test is based on the same principles used to create bead masks in Sec. II C 1. It assumes that the XF projection of a bead is likely to be associated with an edge in the edge map and the inside of the edge will be lower-valued than the outside of the edge. First a template bead outline is extracted from the existing bead masks (as will be explained shortly). This template bead outline is then reprojected onto all the other XF views and the outline is corrected to match the closest edge in the edge mask. Figure 7(a) shows the template bead outline, denoted by the open circles, overlaid onto the edge mask. The points of the template bead outline are corrected by considering a short line segment perpendicular to the template bead outline and resetting the point in the template bead to the closest edge pixel along that line.

The corrected bead outline, overlaid on the XF image, is shown in Fig. 7(b). Notice that the points of the corrected bead outline better match the edge pixels and the edge of the bead shadow. This corrected bead outline is scored on the basis of how likely it is to be a bead. This score, a number between 0.0 and 1.0, is the fraction of points in the estimated bead outline that are judged to lie on the edge of a bead shadow. A point is categorized as lying on the edge of a bead




FIG. 8. Segmented XF projection with the 3D coordinates of the localized beads (dots) reprojected and overlaid. The dashed lines denote the outlines of the segmented beads found by the method described in Sec. II C 1. Note that the segmentation misidentifies and underidentifies the beads in this particular projection. Despite the inaccurate segmentation, all 16 beads are correctly identified by the full marker localization method.

shadow if it (a) lies on an edge pixel in the edge mask and (b) if the pixels adjacent to the point on the “inside” of the bead outline are lower-valued than the pixels on the “outside” of the bead outline. The solid dots in Fig. 7(b) represent points on the corrected bead outline that have passed these tests and are classified as belonging to an edge of a bead shadow. This particular bead-outline scores 0.945, i.e., 105/111 of its points pass the test.

A candidate bead is labeled a false positive if the minimum score over all the XF views or the average score over all the XF views fall below certain thresholds (we use 0.05 and 0.5, respectively). The template bead outline is chosen as the bead region corresponding to the candidate bead that has the highest area-to-convex-hull-area ratio. As explained before, this makes it likely to be a single-bead region and therefore a good template when comparing to other projection views.

Once the false positives have been eliminated from the list of beads, the bead masks are made sparser by removing bead regions that correspond to beads that have been identified so far. Care is taken to ensure that only bead regions corresponding to a single bead (as tested by the area-to-convex-hull-area ratio) are removed. This is done so that bead regions that correspond to beads that have not yet been identified *but* share a bead region with a previously identified bead are not removed.

Then the iteration is repeated; the sparsified bead masks are backprojected onto the 3D volume, a list of candidate beads is made and reduced by removing false positives, and the bead masks are further sparsified by removing the bead regions corresponding to the beads that have been so far found, until all the 16 beads are found. Figure 8 shows the computed centroids (dots) of the identified beads overlaid on one of the XF projection images. The dotted lines denote the borders of the beads that were identified by the segmentation procedure (as described in Sec. II C 1). Notice that even though the segmentation of this particular projection is inaccurate, the centroids of the beads are correctly identified because they have been correctly identified in a sufficient number of other projections.

The iterative method we have described requires the back-projection of the bead masks onto a 3D volume once every iteration. This operation is made more efficient in two ways. First, as we have mentioned before, the cost of backprojecting sparse bead masks onto a 3D volume using ray-driven backprojection (ac_LLN) is a small fraction of the cost of backprojecting nonsparse projections (c_LLN). Second, we are able to exploit the structure of our iterative algorithm to guarantee that the total cost of all the backprojections over all the iterations is no more than $2ac_LLN$, regardless of how many iterations are performed. This is achieved by avoiding redundant operations. Between iterations, the change to the bead masks is that they are sparsified by zeroing out pixels corresponding to bead regions. Rather than sparsifying the bead masks and then rebackprojecting all of them, we instead subtract the contributions of the removed bead masks from the existing volume image. To determine the computational cost, note that in the worst case, every bead region will be first backprojected and then subtracted from the volume. Using ray-driven backprojection, this results in a total cost equivalent to two backprojections: $2ac_LLN$.

II.D. Registration of markers across modalities

The ICP method²⁰ is widely used to register shapes in 3D. It is a greedy algorithm that finds the rotation and translation that best registers two sets of 3D points $\{x_1, x_2, x_3, \dots, x_M\}$ and $\{y_1, y_2, y_3, \dots, y_N\}$, i.e., it attempts to find the rotation matrix R and translation vector T that minimizes $\sum_j \| (Ry_j - T) - x_{i(j)} \|^2$ [a greedy algorithm³⁵ (p. 329) is one that makes locally optimal choices that might not necessarily lead to the globally optimum solution]. In order to allow for missing points and false positives in both the MR and XF marker localization, the sum over j is performed not on all the markers but only on the \tilde{L} best pairs of markers (we use $\tilde{L}=10$). During each iteration of the ICP algorithm, it first associates each point y_j with the closest point in the other set $x_{i(j)}$ and then finds the optimal rotation and translation to minimize the registration error. The ICP algorithm is guaranteed to find a local minimum but not necessarily the global minimum.

In order to avoid local minimums, we modify the traditional algorithm as follows. We first run the ICP and check if the resulting registration is good enough. If not, we rerun the ICP with a perturbed starting guess. The criterion we use for judging the goodness of the estimate is that a minimum number of beads (we use two) in each of the anterior and posterior parts of the belt have been registered to within a certain threshold (we use $\delta=5$ mm). In other words, the registration is good enough if at least two beads in each of the anterior and posterior parts of the belt satisfy the condition that $\| (Ry_j - T) - x_{i(j)} \| < \delta$. The set of markers is automatically divided into anterior and posterior by using the K-means clustering algorithm.³⁶ If the registration is not good enough, we perturb the current solution of the ICP by associating a point y_j with, not the nearest point in the other set, but the second nearest point. We rerun the ICP until the registration meets the termination criteria stated above or the maximum number of iterations is surpassed. Figure 9 shows the registration

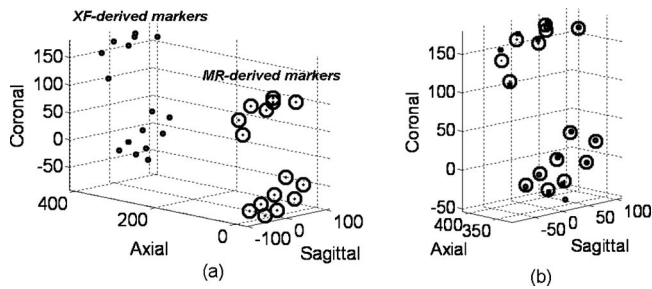


FIG. 9. Registration. (a) The two sets of markers before the ICP is applied. (b) The two sets of correctly registered markers after the application of the ICP.

method in action. Figure 9(a) shows both sets of points: The XF (small circles on the left) and MR (larger circles on the right). Figure 9(b) shows both sets after the MR points have been registered to the XF ones. Notice that the XF set has a false positive and therefore no MR counterpart.

II.E. Refining the registration

Now that the beads from the two modalities have been registered, we perform a final refinement of the solution. We (i) reduce the number of markers to a set in which we have high confidence, (ii) exploit the knowledge of the MR localization to refine the XF localization, and then (iii) perform a final registration of the set of high-confidence beads.

For part (i), we choose the two best registered beads in each of the anterior and posterior subsets and then all the remaining beads that fall below the $\delta=5$ mm threshold [i.e., $\|(Ry_j - T) - x_{i(j)}\| < \delta$]. For the rest of this process, we only use these high-confidence beads.

The reason for part (ii) is that initial localization of the beads from the XF data can be incorrect. This is due to the inexact segmentation of the beads and the limited view-angle range causing the backprojected volume to contain beads that are stretched out in the general direction of the x rays and therefore resulting in possibly incorrectly computed centroids. Our strategy for refining this estimate is twofold. First, we produce better bead masks of the XF projections whose centers (i.e., 2D coordinates) can be more reliably computed. Second, we analytically solve for the 3D coordinates of the centers of the beads from the 2D coordinates of their centers from their bead masks using the matrix $P_{\bar{\theta}}$ that describes the projective geometry of the X-ray system. The exact segmentation of the XF projection is more important here than in the initial segmentation used for the backprojection because the error in computing the refined 3D centroid of the bead is affected by the error in computing the 2D centroid of its associated region in the bead mask. Furthermore, incorrectly identifying multiple beads as a single bead can introduce errors in the associated 2D centroid.

To produce a better bead mask of the XF projection, we first create a MIP of the 3D MR volume at the equivalent orientation of the XF projection. Figure 10 shows an example of how the MR MIP is used to create a more reliable bead mask of the XF projection. Figure 10(a) shows a detail




FIG. 10. Refining the segmentation of XF projections. (a) The boundaries of the beads (dashed black lines) that have been computed from the MR volume, overlaid on a MIP of the MR volume. (b) Detail of the corresponding edge mask of the XF projection. (c) Correctly segmented bead outline (solid black line) found using a template from the MR data. Compare to the originally incorrectly segmented bead outline (lighter gray line).

of the MR MIP. Note that in creating the MIPs, we only consider voxels in the regions of the beads of interest, i.e., we interpolate along rays within the 3D MR volume only in the neighborhood of the high-confidence beads. This is not only computationally efficient but can create reliable 2D boundaries of the beads even when the beads appear overlapping in the 2D XF projection. By operating on voxels in the locality of the bead of interest, the problem of obstruction by another bead is avoided. Overlaid in Fig. 10(a) are the boundaries of the beads (dashed black lines) computed from the 3D MR volume. The arrows point to two beads which appear partially overlapped by each other in the 2D projection image but whose boundaries have been successfully detected. We use these reliable boundaries as templates, overlay them on the edge map of the XF projection, using the procedure described in Sec. II C 4 (and Fig. 7) to generate a better bead mask for the XF projection. In particular, we displace the template boundary to the nearest edge in the edge map. In Fig. 10(b), the edge map of the XF projection is shown in gray and the template boundary is shown by the black line. Figure 10(c) shows the underlying XF projection image with the successfully detected bead boundary overlaid. For comparison, the original bead boundary, computed by the method described in Sec. II C 1, is shown by the dotted gray line. Observe that the original boundary falsely includes a second bead within the bead boundary because of the obstruction problem but the newer method correctly segments the bead of interest. The original boundary would have resulted in an incorrect estimate of the 2D centroid of the bead shadow.

For part (iii) of registration refinement, we simply register the set of bead centroids of the high-confidence beads $\{x_1, x_2, x_3, \dots, x_M\}$ computed from the MR volume, as described in Sec. II B, with the refined set of corresponding bead centroids $\{y_1, y_2, y_3, \dots, y_M\}$ that is computed by the method described above. Since the two sets are exactly the same size and are matched to each other, we can directly apply the SVD-based method of Arun *et al.*³⁷ to compute the optimal rigid-body transformation.

II.F. Experiments

The method was implemented on an XFM system using a 1.5 T clinical MR scanner (Espree; Siemens Medical Sys-

tems, Erlangen, Germany) and a single-plane XF system (Axiom Artis; Siemens Medical Systems, Erlangen, Germany). Several experiments, in phantoms and *in vivo*, were conducted as described below. All animal experiments presented in this paper were approved by the NHLBI Animal Care and Use Committee.

II.F.1. Measuring distortions due to nonlinear MR gradients

Yu *et al.*³⁸ reported that the nonlinearity of the MR gradient can cause the MR image to be spatially distorted and consequently introduce significant errors in cross-modal registration. To measure the effect of this distortion, a phantom experiment was conducted. Six markers, each a toroid-shaped bead filled with iodinated contrast and Gd measuring about $15\text{ mm} \times 15\text{ mm} \times 5\text{ mm}$ in volume with a 5 mm diameter hole (Beekley, Bristol, CT), were mounted on the surface of a spherical Plexiglas phantom (Siemens Medical Systems, Erlangen, Germany) of diameter 250 mm and filled with water. This phantom was imaged in 3D under CT, with a voxel size of 0.5 mm using a 320-slice clinical CT scanner (Aquilion One; Toshiba Medical Systems, Tokyo, Japan) and under MRI, with a voxel size of $1.9 \times 1.9 \times 1.6\text{ mm}$ using the MR scanner that is used in the remainder of this paper. The MR image was reconstructed in the three different ways provided by the native Siemens reconstruction software: Without gradient distortion correction, with distortion correction in 2D, and with distortion correction in 3D.

The centers of the markers were identified in all four volume images (one CT volume image and three MR volume images) using the method described in Sec. II B. In all four cases, the centers of the six markers in the volume images were correctly identified and visually verified. The set of CT-derived marker centers was paired with each of the three MR-derived marker centers. The two sets of markers were registered to each other and the registration errors were computed by measuring the distance between the corresponding pairs of markers after registration. The mean and maximum registration errors over all six markers were computed.

II.F.2. *In vivo* experiments

The algorithm was tested on data from 35 separate *in vivo* experiments conducted over a span of 20 months. Distortion correction (in 2D) provided by the Siemens MR image reconstruction software was used on the MR image data. The angular range of the X-ray projections was approximately 60° (between 54° and 61°) in all the experiments, except for one experiment for which the range was 180° . In each experiment, there were between 13 and 16 beads used. The fully automatic method successfully localized the beads in MR and XF and registered the two sets in all the experiments. The automatic registration was repeated twice for each data set: In run A, the full set of projections (numbering between 16 and 20) was used and in run B, only 10 projections (automatically chosen to be uniformly distributed within the original set) were used.

The error of the XF localization alone was automatically computed. This was done by measuring the consistency of the 3D bead location across projections. The error of a particular 3D bead coordinate relative to a particular projection can be quantified by computing the perpendicular distance between the 3D coordinate of the center of the bead, computed by the localization method, and the ray that passes through the X-ray source and the center of the corresponding bead region in the projection image (as long as the bead region corresponds to a single bead). The RMS of this number over all projections that see this bead represents the error of the XF localization method in identifying this bead. This number was averaged for all the high-confidence beads in the 35 experiments. To judge fairly between run A and run B, for each experiment, only the beads common to both runs were used. Furthermore, in both runs, the error was computed over all the projections, even though the XF bead localization in run B was performed using only ten projections.

A leave-one-out approach was also used to measure the XF bead localization error. For each bead, the set of its corresponding bead regions, each from a separate projection, was considered. For each of these M bead regions, the 3D coordinate of the bead was computed from the remaining $M-1$ bead regions. Then the error of this 3D coordinate relative to the left-out bead region was computed as above. This error was averaged over all bead regions. Only those beads that had more than two corresponding bead regions were used as the computation of the 3D coordinate requires at least two projections.

Finally, the average error of the registration between the two point sets in each experiment was computed. In each of the 35 experiments, the registration was performed using the eight best beads in that experiment. We chose eight beads in order to allow for the motion of some of beads between modalities and yet maintain an overdetermined registration problem. The eight best beads were chosen by first performing the rigid-body registration of all the refined high-confidence beads and then picking the best beads on the basis of the error between corresponding pairs (i.e., MR and XF) of registered beads. The eight beads were chosen, for the reasons described in Sec. II D, as the two best beads from each of the anterior and posterior parts of the belt and the four best of the remaining beads. For each experiment, the error of the registration was computed over both the eight best beads and over *all* the high-confidence beads.

II.F.3. XF localization with fewer projections

In order to demonstrate the effect of reducing the number of XF projections on the localization of the markers, in each of the 35 experiments we reran the XF localization using a reduced numbers of X-ray projections (numbering four to ten projections). The reduced set of projections was automatically chosen to be as uniformly distributed within the original set as possible. The number of correctly localized beads in each experiment was determined by comparing the output of the XF localization method to the gold-standard set of beads for that particular experiment. The gold-standard set of

beads was formed by first considering the set formed from the union of the sets of markers identified from multiple subsets of projections. (The unions of multiple sets, rather than just the single set of markers identified by the full set of projections, were used in order to avoid the unlikely case in which a marker that is identified in one of the sets is not identified by the full set of projections). False positives were manually removed from the gold-standard set after visual verification upon reprojection onto the XF projections. Correctly identified beads were those that were less than 5 mm from a bead in the gold-standard set of beads. For a fixed number of projections the mean, standard deviation, minimum, and maximum (over the 35 experiments) of the number of correctly identified and false-positive beads were computed.

II.F.4. Retrospectively reduced set of markers

In order to demonstrate the effect of reducing the number of fiducial markers, we performed the automatic registration using a retrospectively reduced number of beads. We reran the registration to mimic the effect of using half the number of fiducial markers (i.e., a total of eight markers; four each from the anterior and posterior subsets). We used an automated method, described below, that chose this reduced set in a partially randomized manner. This method allows for the inclusion of false positives and false negatives (i.e., missing beads) and tries to ensure that the reduced set of beads form an asymmetric arrangement and are well-separated. The asymmetric, well-separated arrangement is chosen to reduce the probability that the ICP avoids local minima. If the reduced set of beads were chosen assuming a purely uniform probability distribution, the resulting arrangement might be highly symmetric and closely spaced. We believe that this would be an unrealistic simulation of an experiment conducted with a (prospectively) reduced set of markers, as in that case we would have complete freedom in choosing the arrangement of the markers.

We reran the registration using the reduced set and compared the error to the original registration computed from the high-confidence beads as described in Sec. III B. The registration was successful in all 35 experiments. Note that in applying the ICP, we did not use $\tilde{L}=10$ (Sec. II D) as when we used the full set of markers but instead used a range of \tilde{L} 's between $\tilde{L}=4$ and $\tilde{L}=8$ and automatically chose the correct registration as the one that produced the largest number of matched pairs of beads (a matched pair of beads was categorized as one in which the Euclidean distance between the pair was less than a threshold of 10 mm). The rigid-body transformations computed from registering the reduced set of beads were applied to the high-confidence beads from the full sets and the error calculated.

The simple automated method to choose the reduced subset of beads is as follows. The input of the method is the complete combined set of registered X-ray and MR beads found by the methods described in Secs. II B–II D. The beads are split into anterior and posterior sets using K-means clustering.³⁶ Using a threshold of 15 mm (to allow for sub-

stantial motion between modalities), we categorize the beads as matched between modalities or unmatched (i.e., in the case when one of the modalities has a missing bead or a false positive). Notice from Fig. 9 that the beads in each of the anterior and posterior subsets lies in a gridlike pattern on a rectangular plane that is oriented roughly in a coronal plane. We perform principal component analysis on the sagittal and axial coordinates of the markers. The first principal component vector can be assumed to correspond to the lengthwise direction of the rectangle. We then set up four anchor points in an asymmetric pattern (the four corners of a trapezoidal-like shape whose base is along the sagittal direction) in the sagittal-axial plane using the first principal component vector and its perpendicular. This pattern is then randomly flipped (i.e., multiplied by a discrete uniformly distributed random number that is either -1 or 1) in the axial direction about the center of mass of the (anterior or posterior) subset. The reduced set of beads is simply the set of beads that are closest to the four anchor points. If the closest bead to an anchor point is a matched pair of beads, then a bead from each modality is added to the reduced set. If the closest bead is an unmatched bead, then only that bead is added to the reduced set. Consequently, the reduced sets include both false positives and false negatives and beads that have undergone motion (i.e., less than 15 mm) between modalities.

II.F.5. Target registration error in a static phantom

In order to measure the error of registering a target between the two modalities, a simple phantom made of plastic was constructed to mimic a human torso. Two static targets, each a toroid-shaped bead filled with iodinated contrast and Gd measuring about 15 mm \times 15 mm \times 5 mm in volume with a 5 mm diameter hole (Beekley, Bristol, CT), were mounted inside the phantom to mimic a structure within the torso in the region of the heart. The belt of beads was fixed to the phantom. As described in Sec. II B, a T_1 -weighted 3D MR volume image (with voxels of dimensions 1.88 mm \times 1.88 mm \times 1.6 mm) was collected. The phantom was moved to the XF system and 20 low-power X-ray projections (with an angular range of 60°) were collected. The automatic registration experiment was performed twice. First, with all 20 projections (run A) and then with only 10 projections (i.e., every alternate projection, run B).

The target registration error was computed in each run. First the 3D coordinates of the target was computed under X-ray. The X-ray bead localization algorithm automatically identified the 3D coordinates of the target. In order to get a more accurate target registration error, we refined the estimate of the 3D coordinate of the targets alone by ensuring that the 2D coordinates of the target centers corresponded to the center of the toroid in the X-ray projections by simple image processing (thresholding and centroid computation) followed by a visual inspection of the X-ray projections of the target. The 3D coordinates of the targets computed from this refined estimate of the X-ray data was then transformed to the MR coordinate system using the automatically computed registration from each run. To allow for the motion of

TABLE I. Summary of results of 35 *in vivo* experiments.

Number of projections	Number of beads registered with high confidence			Marker localization error under X-ray (mm)					
				RMS error per bead			Leave-one-out error per bead region		
	Min	Max	Median	Mean \pm 1 sd	Min	Max	Mean \pm 1 sd	Min	Max
All (16–20)	10	16	14	0.32 \pm 0.19	0.00	1.70	0.43 \pm 0.32	0.01	5.20
10	10	16	13	0.43 \pm 0.49	0.00	6.95	0.47 \pm 0.34	0.00	3.17

some of the beads between modalities, the final rigid-body transformation was computed, as in Sec. II F 2, using the eight best beads. The true 3D coordinates of the target in the MR image was computed by picking a subimage of $15 \times 15 \times 15$ voxels surrounding the transformed target coordinates and finding the centroid of this subimage.

II.F.6. Target registration *in vivo*

Two *in vivo* (porcine) experiments were conducted to validate the automatic registration. In each experiment a T_1 -weighted 3D MR image (with voxel size $2.3 \times 2.3 \times 1.6$ mm) and 20 low-power X-ray projections (with an angular range of 60°) were collected. The rigid-body registration between the coordinate systems of the two imaging modalities was automatically computed, as described in Sec. II. Only ten X-ray projections were used for the X-ray localization.

The targets were chosen to be small metallic structures on an interventional device, while the device was present inside the heart of the animal and subject to cardiac motion. The devices used were an Amplatzer muscular VSD occluder in the first experiment and an ASD occluder in the second experiment (both from AGA Medical Corp., Minneapolis, MN). The 3D coordinates of the target in the X-ray coordinate system were determined from two sets of fluoroscopic cine images, each set at a fixed view angle. Each set was collected while the animal was subject to a breath-hold using a respirator. ECG information was used to sort the images according to cardiac phase. The coordinates of the targets during a heart cycle were computed semiautomatically, using simple image processing (peak finding), and verified visually. The 3D coordinates of the target over the whole heart cycle were computed from the two sets of 2D coordinates.

Under MR, 2D cine slices corresponding to a full heart cycle and containing the target were collected. To validate

the registration, the 3D coordinates of the targets were transformed to the MR coordinate system, using the previously computed rigid-body registration parameters, and displayed. ECG information was used to synchronize the heart-phase of the images from the two data sets.

III. RESULTS

III.A. Measuring distortions due to nonlinear gradients

For the experiment described in Sec. II F 1 the mean registration errors were 4.78, 0.97, and 0.48 mm, respectively, for the no-distortion correction, 2D-distortion-correction and 3D-distortion-correction cases. The maximum registration errors in the three cases were 5.03, 1.79, and 0.74 mm, respectively. It was concluded that errors in spatial distortion in the *distortion-corrected* images (i.e., the latter two cases) were sufficiently small for XFM.

III.B. *In vivo* experiments

The results of the *in vivo* experiments, performed as described in Sec. II F 2, are listed in Tables I and II for both run A (all projections) and run B (ten projections). The number of high-confidence beads was between 10 and 16 with a median of 14 in run A and was between 10 and 16 with a median of 13 in run B. The mean of the XF localization error was 0.32 mm in run A and 0.43 mm in run B. The mean leave-one-out error of the XF localization was 0.43 and 0.47 mm for runs A and B, respectively. Table II lists the error of registration between the two sets of beads. While the registration was performed using only the eight best beads in each experiment, the error was computed for both those eight beads and all the high-confidence beads. When averaged over the eight best beads alone, the error was 1.2 mm in both

TABLE II. Summary of XFM registration errors. In all cases, registration was performed using the best eight beads. The registration error was then computed in two ways: First, over the best eight beads and then over all the high-confidence beads.

Number of projections	Error computed over eight best beads (mm)			Error computed over all high-confidence beads (mm)		
	Mean \pm 1 sd	Min	Max	Mean \pm 1 sd	Min	Max
All (16–20)	1.16 \pm 0.52	0.07	3.18	1.50 \pm 0.87	0.07	6.26
10	1.18 \pm 0.56	0.27	3.53	1.52 \pm 0.94	0.27	5.78




FIG. 11. Fused XFM image. The rigid registration parameters found using our method are used to overlay information from the MR data onto the XF projection. The contours, representing tissue boundaries, are extracted from the MR volume image.

run A and run B and when averaged over all the high-confidence beads, the error was 1.5 mm.

Figure 11 shows the resulting fused image from one of the experiments. The colored contours, representing the aorta (red), left ventricular epicardium (green), left ventricular endocardium (blue), and right ventricular endocardium (yellow), extracted from the MR image are displayed on the grayscale XF projection image. The process of creating the contours from the MR image is outside the scope of this paper, but details can be found in the work of Ratanayak *et al.*¹² Notice that the contours line up with the faint shadow of the heart in the XF image.

III.C. XF localization with fewer projections

The results for the experiments, in which XF localization is performed with a reduced number of projections as described in Sec. II F 3, are listed in Table III. Table III lists the mean, standard deviation, minimum, and maximum number of correctly identified beads and false positives. The number of projections was varied from four to ten in each of the 35 experiments. As the number of projections is decreased, the number of correctly identified beads generally decreases and the number of false positives increases.

TABLE III. Success of bead identification vs number of projections.

Number of projections		10	9	8	7	6	5	4
Number of correct beads	Mean \pm sd	13.8 ± 1.8	13.6 ± 1.7	13.0 ± 1.6	12.0 ± 1.9	12.5 ± 2.1	10.3 ± 2.3	6.0 ± 1.6
	Min	10	10	9	7	8	6	2
	Max	16	16	16	16	16	15	10
Number of false positives	Mean \pm sd	0.9 ± 1.1	0.7 ± 0.8	1.5 ± 1.6	2.7 ± 2.0	1.6 ± 1.5	3.9 ± 2.5	9.9 ± 1.8
	Min	0	0	0	0	0	0	6
	Max	5	3	7	9	5	10	14

III.D. Retrospectively reduced set of markers

The error in registering a retrospectively reduced set of markers was computed as described in Sec. II F 4. The mean registration error (computed over all the beads as described in Sec. III and II) increased from 1.2 to 1.6 mm. Note that this reduced set of beads has no more than eight beads. The standard deviation, minimum, and maximum of the registration error was 1.1, 0.1, and 10.8 mm, respectively.

III.E. Target registration error in a static phantom

Figure 12(a) shows one of these X-ray projections with the targets identified by the dashed squares and their centers identified by black dots. Figures 12(b) and 12(c) show details of MIPs of the 3D MR volume image in the coronal and axial directions. The number of high-confidence beads was 13 in run A and 12 in run B. The $15 \times 15 \times 15$ voxels sub-images whose centroids are used to find the locations of the targets in MR are denoted by the dashed white squares in Figs. 12(b) and 12(c). The coordinates of these targets in the MR system are denoted by the white x's in Fig. 12(b) and 12(c). The coordinates of the targets transformed from the X-ray frame are shown by the white +'s. Figure 12 corresponds to run B in which only ten projections were used. The target registration errors for the two targets were 1.6 and 2.2 mm for run A and 1.9 and 2.4 mm for run B. As can be seen in Figs. 12(b) and 12(c), these differences are on the order of the size of the MR voxel (i.e., $1.9 \times 1.9 \times 1.6$ mm³, and whose longest diagonal is, therefore, 3.1 mm).

III.F. Target registration *in vivo*

Figure 13 shows images from the experiments. Figures 13(a) and 13(b) show, from the first and second experiment, respectively, an X-ray projection image with the location of the targets marked by the black circles. Figures 13(c) and 13(d) show, from the first and second experiment, respectively, a MR slice containing the target. While the targets are visible with high resolution under X-ray, under MRI the targets produce a magnetic susceptibility artifact (caused by distortions of the magnetic field) and are visible as extended dark regions. Since the targets do not appear sharp under MRI and are inherently spatially distorted, an exact computation of TRE is impossible, but in both experiments the transformed coordinates of the targets (marked by the white circles) are seen to lie in the region of the susceptibility artifacts. In Figs. 13(c) and 13(d), the transformed coordi-




FIG. 12. Phantom experiment to measure target registration error. (a) Representative X-ray projection. The centroids of the fiducial markers are marked by the white x's and the targets are marked by black dashed squares. Detail of MIPs of the MR volume (b) along the transverse direction and (c) the coronal direction, showing the two targets (within the dashed white boxes). The white x's mark the true centroid of the targets in the MR image and the white +'s mark the location of the centroids that have been transformed from the X-ray coordinate system using the registration parameters computed by our automatic method.

nates of the targets can be seen to lie within 0 and 3 pixel units of the susceptibility artifact, which, using the MR image pixel size of 1.56 mm, is between 0 and 4.7 mm within the 2D plane.

IV. DISCUSSION

We have demonstrated the fully automatic registration between MR and X-ray imaging systems using external fiducial markers. The robustness of the method has been demonstrated by its 100% success rate in 35 *in vivo* experiments. The most challenging part of the method is identifying the 3D locations of the fiducial markers from a small number of X-ray projections. The high accuracy of this method was




FIG. 13. *In vivo* target registration experiment. (a) Target (closure device) in X-ray projection. (b) Target in X-ray projection. (c) Target from (a) transformed into MR coordinates and overlaid on slice. (d) Target from (b) transformed into MR coordinates and overlaid on slice.

demonstrated by quantifying the consistency of the computed 3D locations relative to their 3D projections. The average RMS error of under 0.4 mm and the average leave-one-out error of under 0.5 mm validate both the accuracy of the reconstruction method and the distortion correction method¹¹ used to correct spatial distortions in the image intensifier projection images. The backprojection-based X-ray localization method makes no assumptions about the configuration of the beads or the view angles of the projections. Using backprojection also allows for beads that overlap with each other on an XF projection, and result in multibead regions as seen in Fig. 3, to still be successfully identified. Using a nonmatching method for localization avoids the problem of matching beads from erroneously segmented projections.

The clinically acceptable accuracy of the registration depends on the particular procedure performed. For VSD closure, the accuracy depends on the size of the VSD, which can range in diameter from a few mm to a few cm. A registration accuracy of 5 mm would be acceptable for the majority of such VSDs. Rhode *et al.*¹⁹ specify an acceptable accuracy of 5 mm for cardiac electrophysiological procedures. The TREs in the static phantom experiment presented in Sec. II F 5 are less than 2.4 mm. The initial results from *in vivo* experiments, presented in Sec. III F and Fig. 13, display registration errors in the plane of the MR slice of under 5 mm. As demonstrated in Sec. III A, using 3D instead of 2D-distortion correction on the MR localization volume image will likely reduce errors further.

We have demonstrated that the X-ray marker localization method is sufficiently robust when using ten low-power X-ray projections. This represents 1.2 μGy of total absorbed radiation dose which is an extremely small fraction ($<0.1\%$) of the total radiation³⁹ that the patient is exposed to during a typical interventional procedure. Table III shows that reducing the number of projections increases errors in

X-ray marker localization because the likelihood of a marker falling outside the field of view of a particular projection increases. If a small number of projections, say six, were used, the XF marker localization and registration would be successful in the majority of cases but the likelihood of failure would increase. If a small number of projections were used along with a larger number of markers, the registration would likely fail more often. This is because even though the number of correctly identified markers would increase, the number of false positives would likely increase more because the false-positive-elimination test described in Fig. 7 and Sec. II C 4 would fail more often because of the increased number of bead regions per projection.

The localization uncertainty from XF projections depends on the geometry of the system and the range of view angles used. The geometry of X-ray system (the size and resolution of the XF detector and distance from the X-ray source) results in an uncertainty of about $\Delta=0.5$ mm perpendicular to the projection direction at isocenter. The localization uncertainty along the direction of the central projection is approximately $\Delta/\sin(\theta/2)$, where θ is the full range of view angles. For $\theta=60^\circ$ that we use, the uncertainty in the projection direction is about 1 mm and is therefore sufficient. While the angular range could be decreased further without resulting in unacceptable localization uncertainty, the reason for using the full 60° range is the limited field of view of the XF detector (which has a side length of 330 mm). Reducing the view-angle range would make it more likely that some markers would fall outside the field of view of all the projections. From Fig. 9, it may be observed that lateral views are likely to contain many overlapping markers as the markers lie on roughly coronal planes. Consequently, we align the central XF view along the anterior-posterior direction.

The advantage of using the relatively large number of fiducial markers, as we do, is that the redundancy in the system makes the ICP registration more robust and allows for the motion of a small number of these beads during the transfer between imaging modalities. While the markers have been deliberately placed far apart from each other in order to minimize obstruction, when a marker does obstruct an important region of the anatomy, the view angle of the projection can be slightly modified to reveal the region. If a smaller number of markers were to be used, the strategy to efficiently place them would be to (a) place them such that they did not obstruct structures of interest, even adapting the arrangement to typical views used in a particular procedure, and to (b) use an asymmetric pattern in order that the iterative closest point method avoids local minimums. In Sec. III D, we have demonstrated that the modalities can be successfully registered with a smaller number of markers. While the likelihood of failure in registration increases as the number of markers is greatly reduced, the results presented in Secs. III B and III D suggest that using about eight to ten markers from the current marker configuration would be sufficient in the vast majority of cases. If smaller markers were used, the methods described here could be used with only minor modifications of parameters (such as the thresholds related to the size of

the markers used in localizing in MR and XF) as long as the MR image had a sufficiently small voxel size.

In order to minimize the movement of the markers, a specialized table system (Miyabi, Siemens Medical Systems, Erlangen, Germany) is used. It allows the subject to be transferred between the MRI scanner and the X-ray system without leaving the table. Furthermore, the belt is constructed from elastic and placed securely on the subject's torso. If the belt does shift relative to the patient, it can be detected by the misregistration of the MRI contours and corresponding anatomy in XF. If necessary, a left ventriculogram could be performed to adjust the registration. Other possible ways of detecting and compensating for this motion could be to fix markers directly to locations on the skin surface, as suggested by Rhode *et al.*,¹⁹ and to use the relative motion of the markers on the belt and the skin to correct for the shift.

One challenge in translating this technique to humans would be to deal with respiratory motion. In the *in vivo* experiments, we have presented, in which the animal is on a respirator, the movement of the beads during respiration has not been observed to be significant. In human patients, respiratory gated or breath-hold acquisitions can be used for the bead localization scans, so that the initial registration is done with images acquired at the same respiratory cycle for each modality. Work is currently underway to incorporate motion correction into XFM. On the other hand, performing XFM in humans will present the opportunity to place the external fiducial markers more securely by using the shoulders of human patient as an additional anchor for the belt of beads. In XF, the location of the human patient's arms is sometimes altered during the procedure which has the potential to cause the markers to move. But as long as the patient's arms are in the same position under both MR and XF for the marker localization scans, the registration will be unaffected.

The method was implemented, with minimal optimization of code, in MATLAB, on a PC (3.2 GHz Pentium D with 3.5 GB of RAM). The typical time taken for the whole method, using ten X-ray projections, was between 1 and 2 min. This time was divided among the various components of the method as follows. The MR marker localization takes about 1.0 s, the XF marker localization takes about 90 s, and the registration between the two data sets takes about 0.1 s and fine-tuning the registration takes 15 s (with about 10 s spent on simulating a MIP from a MR volume). The XF marker localization, the most time-consuming operation, spent its time on average as follows: 7 s on the dewarping (distortion correction) of the ten X-ray images, 18 s on the processing of the XF images to produce bead masks, and about 60 s on the iterative backprojection-based algorithm. The time taken by iterative algorithm varies depending on the particular experiment. It can take tens of seconds less or more in time depending on the number of iterations it needs and the number of false positives it needs to eliminate. The use of ray-driven backprojection produces a speedup of about $5\times$ in the actual backprojection but there is also an overhead in initially computing the reduced number of voxels of the volume to backproject onto, which takes about 20 s in the current implementation.

As the automatic registration is performed only once during the interventional procedure, these run times are not excessive but there is room for optimization of the code and speedups. The runtimes can be reduced if implemented on parallel processors. In the X-ray localization method, both the distortion correction and the creation of the bead masks can be parallelized by operating separately on individual projections, giving an order of magnitude potential reduction in runtime. Much work⁴⁰ has been done on using parallel processors to reduce the cost of backprojection by orders of magnitude. In our case, memory requirements, and consequently runtime, can be further reduced by using 8-bit integers instead of doubles when performing the backprojection, as our method is unaffected by limiting the voxel values in the backprojected to small integers (because the maximum voxel value is less than the number of projections used).

V. CONCLUSIONS

We have presented a fully automatic method, which uses a set of external fiducial markers strapped to the subject's torso, to register a roadmap MR image to an X-ray fluoroscopy system. The advantages of using external markers include its noninvasiveness, the ability to resolve patient motion relative to the table, no requirement for any extra hardware, and the ability to use it when the MR and X-ray systems do not share a suite. We have demonstrated the robustness of the registration method in more than 35 *in vivo* experiments using a conventional image intensifier X-ray system. We have further computed the target registration error of the method in a phantom to be less than 2.4 mm. The method allows us to use fewer X-ray projections and a smaller number of fiducial markers than previously used. Further work is underway to incorporate cardiac and respiratory motion into XFM.

ACKNOWLEDGMENTS

This work was supported by the intramural research program of the National Heart, Lung, and Blood Institute. The authors thank the following people for their help in collecting data and performing experiments: Israel Barbash, Jessica Colyer, Katherine Lucas, Kanishka Ratnayak, Bill Schenke, and Victor Wright.

^{a)}Telephone: 301496 1151; Fax: 301 402 2389; Electronic mail: georgeak@nhlbi.nih.gov

^{b)}Author to whom correspondence should be addressed. Electronic mail: faranesa@nhlbi.nih.gov

¹T. Beyer and D. W. Townsend, "Putting 'clear' into nuclear medicine: A decade of PET/CT development," *Eur. J. Nucl. Med. Mol. Imaging* **33**, 857–861 (2006).

²L. A. Dawson and D. A. Jaffray, "Advances in image-guided radiation therapy," *J. Clin. Oncol.* **25**, 938–946 (2007).

³T. H. Schindler, E. U. Nitzsche, N. Magosaki, M. Mix, A. D. Facta, J. O. Prior, U. Solzbach, H. R. Schelbert, and H. Just, "Myocardial viability in patients with ischemic cardiomyopathy—evaluation by 3-D integration of myocardial scintigraphic data—and coronary angiographic data," *Mol. Imaging Biol.* **6**, 160–171 (2004).

⁴B. J. Micaelian, Z. J. Malchano, P. Neuzil, J. Weichet, S. K. Doshi, J. N. Ruskin, and V. Y. Reddy, "Images in cardiovascular medicine. Integration of 3-dimensional cardiac computed tomography images with real-time electroanatomic mapping to guide catheter ablation of atrial fibrillation,"

Circulation **112**, e35–e36 (2005).

⁵A. K. Singh, J. Kruecker, S. Xu, N. Glossop, P. Guion, K. Ullman, P. L. Choyke, and B. J. Wood, "Initial clinical experience with real-time transrectal ultrasonography-magnetic resonance imaging fusion-guided prostate biopsy," *BJU Int.* **101**, 841–845 (2008).

⁶N. Amoretti, P. Y. Marcy, V. Lesbats-Jacquot, I. Hovorka, M. E. Fonquerne, C. Roux, O. Hericord, Y. Maratos, and L. Euler-Ziegler, "Combined CT and fluoroscopic guidance of balloon kyphoplasty versus fluoroscopy-only procedures," *Skeletal Radiol.* **38**, 703–707 (2009).

⁷A. Kerner, E. Ghersin, and A. Roguin, "Multimodality imaging of borderline left main coronary disease using fluoroscopy, IVUS and CT coronary angiography," *Int. J. Cardiovasc. Intervent.* **7**, 108–109 (2005).

⁸C. J. Lin, R. Blanc, F. Clarencon, M. Piotin, L. Spelle, J. Guillemic, and J. Moret, "Overlying fluoroscopy and preacquired CT angiography for road-mapping in cerebral angiography," *AJNR Am. J. Neuroradiol.* **31**, 494–495 (2009).

⁹R. de Silva, L. F. Gutierrez, A. N. Raval, E. R. McVeigh, C. Ozturk, and R. J. Lederman, "X-ray fused with magnetic resonance imaging (XFM) to target endomyocardial injections: Validation in a swine model of myocardial infarction," *Circulation* **114**, 2342–2350 (2006).

¹⁰L. F. Gutiérrez, "X-ray fused with MRI (XFM) for guidance of catheter-based interventions," Ph.D. thesis, Johns Hopkins University, 2005.

¹¹L. F. Gutiérrez, C. Ozturk, E. R. McVeigh, and R. J. Lederman, "A practical global distortion correction method for an image intensifier based X-ray fluoroscopy system," *Med. Phys.* **35**, 997–1007 (2008).

¹²L. F. Gutiérrez, R. Silva, C. Ozturk, M. Sonmez, A. M. Stine, A. N. Raval, V. K. Raman, V. Sachdev, R. J. Aviles, M. A. Waclawiw, E. R. McVeigh, and R. J. Lederman, "Technology preview: X-ray fused with magnetic resonance during invasive cardiovascular procedures," *Catheter. Cardiovasc. Interv.* **70**, 773–782 (2007).

¹³K. Ratnayaka, V. K. Raman, A. Z. Faranesh, M. Sonmez, J. H. Kim, L. F. Gutierrez, C. Ozturk, E. R. McVeigh, M. C. Slack, and R. J. Lederman, "Antegrade percutaneous closure of membranous ventricular septal defect using X-ray fused with magnetic resonance imaging," *JACC Cardiovasc. Interv.* **2**, 224–230 (2009).

¹⁴K. S. Rhode, D. L. Hill, P. J. Edwards, J. Hipwell, D. Rueckert, G. Sanchez-Ortiz, S. Hegde, V. Rahunathan, and R. Razavi, "Registration and tracking to integrate X-ray and MR images in an XMR facility," *IEEE Trans. Med. Imaging* **22**, 1369–1378 (2003).

¹⁵K. S. Rhode, M. Sermesant, D. Brogan, S. Hegde, J. Hipwell, P. Lambiase, E. Rosenthal, C. Bucknall, S. A. Qureshi, J. S. Gill, R. Razavi, and D. L. Hill, "A system for real-time XMR guided cardiovascular intervention," *IEEE Trans. Med. Imaging* **24**, 1428–1440 (2005).

¹⁶A. Ganguly, Z. Wen, B. L. Daniel, K. Butts, S. T. Kee, V. Rieke, H. M. Do, N. J. Pelc, and R. Fahrig, "Truly hybrid X-ray/MR imaging: Toward a streamlined clinical system," *Acad. Radiol.* **12**, 1167–1177 (2005).

¹⁷L. Brzozowski, A. Ganguly, M. Pop, Z. Wen, R. Bennett, R. Fahrig, and J. A. Rowlands, "Compatibility of interventional X-ray and magnetic resonance imaging: Feasibility of a closed bore XMR (CBXMR) system," *Med. Phys.* **33**, 3033–3045 (2006).

¹⁸M. Tomkowiak, K. Vigen, M. Speidel, N. Shah, M. Van Lysel, and A. N. Raval, "Multimodality image merge to guide catheter based injection of biologic agents," in Proceedings of the RSNA, Chicago, IL, 2008 (unpublished).

¹⁹K. Rhode, Y. Ma, A. Chandrasena, A. King, G. Gao, P. Chinchapatnam, M. Sermesant, D. Hawkes, T. Schaeffter, and J. Gill, "Evaluation of the use of multimodality skin markers for the registration of pre-procedure cardiac MR images and intra-procedure X-ray fluoroscopy images for image guided cardiac electrophysiology procedures," *Proc. SPIE* **6918**, 69181R.1–69181R.8 (2008).

²⁰P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Mach. Intell.* **14**, 239–256 (1992).

²¹J. H. Kim, O. Kocaturk, C. Ozturk, A. Z. Faranesh, M. Sonmez, S. Sam-path, C. E. Saikus, A. H. Kim, V. K. Raman, J. A. Derbyshire, W. H. Schenke, V. J. Wright, C. Berry, E. R. McVeigh, and R. J. Lederman, "Mitral cerclage annuloplasty, a novel transcatheter treatment for secondary mitral valve regurgitation: Initial results in swine," *J. Am. Coll. Cardiol.* **54**, 638–651 (2009).

²²M. Sonmez, "A fiducial-based automatic registration method for X-ray imaging fused with MRI," M.S. thesis, Bogazici University, 2004.

²³M. Sonmez, A. Z. Faranesh, R. J. Lederman, C. Lorenz, and C. Ozturk, "Automatic registration of X-ray fused with MRI system," in Proceedings

- of the Seventh Interventional MRI Symposium, Baltimore, MD, 2008 (unpublished).
- ²⁴M. D. Altschuler and A. Kassae, "Automated matching of corresponding seed images of three simulator radiographs to allow 3D triangulation of implanted seeds," *Phys. Med. Biol.* **42**, 293–302 (1997).
- ²⁵A. K. Jain, Y. Zhou, T. Mustafa, E. C. Burdette, G. S. Chirikjian, and G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)," *Med. Phys.* **32**, 3475–3492 (2005).
- ²⁶C. Labat, A. Jain, G. Fichtinger, and J. Prince, "Toward optimal matching for 3D reconstruction of brachytherapy seeds," *Med Image Comput Comput Assist Interv* **10**, 701–709 (2007).
- ²⁷J. Lee, X. Liu, A. K. Jain, J. L. Prince, and G. Fichtinger, "Tomosynthesis-based radioactive seed localization in prostate brachytherapy using modified distance map images," *Proc. IEEE Int Symp Biomed Imaging*, 680–683 (2008).
- ²⁸J. Lee, X. Liu, A. K. Jain, D. Y. Song, E. C. Burdette, J. L. Prince, and G. Fichtinger, "Prostate brachytherapy seed reconstruction with Gaussian blurring and optimal coverage cost," *IEEE Trans. Med. Imaging* **28**, 1955–1968 (2009).
- ²⁹X. Liu, A. K. Jain, and G. Fichtinger, "Prostate implant reconstruction with discrete tomography," *Med Image Comput Comput Assist Interv* **10**, 734–742 (2007).
- ³⁰I. B. Tutar, R. Managuli, V. Shandasani, P. S. Cho, S. D. Pathak, and Y. Kim, "Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy," *Med. Phys.* **30**, 3135–3142 (2003).
- ³¹R. C. Gonzalez and R. E. Woods, *Digital Image Processing* (Prentice-Hall, Upper Saddle River, NJ, 2002).
- ³²J. Canny, "A computational approach to edge-detection," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-8**, 679–698 (1986).
- ³³G. Shechter, B. Shechter, J. R. Resar, and R. Beyar, "Prospective motion correction of X-ray images for coronary interventions," *IEEE Trans. Med. Imaging* **24**, 441–450 (2005).
- ³⁴J. Hsieh, *Computed Tomography: Principle, Design, Artifacts, and Recent Advances* (SPIE, Bellingham, WA, 2003).
- ³⁵T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 1997).
- ³⁶S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory* **28**, 129–137 (1982).
- ³⁷K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-9**, 698–700 (1987).
- ³⁸H. Yu, R. Fahrig, and N. J. Pelc, "Co-registration of X-ray and MR fields of view in a hybrid XMR system," *J. Magn. Reson Imaging* **22**, 291–301 (2005).
- ³⁹D. Katritsis, E. Efstathopoulos, S. Betsou, S. Korovesis, K. Faulkner, G. Panayiotakis, and M. M. Webb-Peploe, "Radiation exposure of patients and coronary arteries in the stent era: A prospective study," *Catheter. Cardiovasc. Interv.* **51**, 259–264 (2000).
- ⁴⁰F. Xu and K. Mueller, "Real-time 3D computed tomographic reconstruction using commodity graphics hardware," *Phys. Med. Biol.* **52**, 3405–3419 (2007).