

Gradient-Based 2-D/3-D Rigid Registration of Fluoroscopic X-Ray to CT

Harel Livyatan, Ziv Yaniv, *Student Member, IEEE*, and Leo Joskowicz*, *Senior Member, IEEE*

Abstract—We present a gradient-based method for rigid registration of a patient preoperative computed tomography (CT) to its intraoperative situation with a few fluoroscopic X-ray images obtained with a tracked C-arm. The method is noninvasive, anatomy-based, requires simple user interaction, and includes validation. It is generic and easily customizable for a variety of routine clinical uses in orthopaedic surgery. Gradient-based registration consists of three steps: 1) initial pose estimation; 2) coarse geometry-based registration on bone contours, and; 3) fine gradient projection registration (GPR) on edge pixels. It optimizes speed, accuracy, and robustness. Its novelty resides in using volume gradients to eliminate outliers and foreign objects in the fluoroscopic X-ray images, in speeding up computation, and in achieving higher accuracy. It overcomes the drawbacks of intensity-based methods, which are slow and have a limited convergence range, and of geometry-based methods, which depend on the image segmentation quality. Our simulated, *in vitro*, and cadaver experiments on a human pelvis CT, dry vertebra, dry femur, fresh lamb hip, and human pelvis under realistic conditions show a mean 0.5–1.7 mm (0.5–2.6 mm maximum) target registration accuracy.

Index Terms—Fluoroscopic X-ray to CT registration, gradient based, image registration, 2D/3D rigid registration.

I. INTRODUCTION

REGISTRATION is the task of finding a transformation from the coordinate system of one data set to another so that all features that appear in both data sets are aligned. Registration is an essential step in most computer-aided surgery (CAS) systems, since it is necessary to match information from different data modalities obtained at different times. In image-guided surgery, it is required to match the preoperative images and plans to the intraoperative situation, and to determine the relative positions of surgical tools and anatomical structures. Examples of deployed CAS systems include preoperative planning, intraoperative navigation and robotics systems for orthopaedic surgery [1]–[4], for neurosurgery [5], [6], and for radiosurgery [7], among many others. Practical, accurate, and robust registration has emerged as one of the key technical challenges of CAS.

Manuscript received June 4, 2003; revised August 4, 2003. This research was supported in part by a grant from the Israel Ministry of Industry and Trade for the IZMEL Consortium on Image-Guided Therapy. *Asterisk indicates corresponding author.*

H. Livyatan and Z. Yaniv are with the School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem 91904, Israel.

*L. Joskowicz is with the School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem 91904, Israel (e-mail: josko@cs.huji.ac.il).

Digital Object Identifier 10.1109/TMI.2003.819288

One of the most sought after methods is anatomy image-based rigid registration between preoperative and intraoperative data sets. The goal is to enable surgeons to use preoperative plans and computed tomography (CT) and magnetic resonance imaging (MRI) data in the operating room for image-guided navigation and robot positioning. The registration can be performed with a few intraoperative fluoroscopic X-ray or ultrasound images, which are ubiquitous, noninvasive, and easy to acquire. Current CAS systems rely on implanted fiducials, which require an additional surgical procedure, or on points obtained by direct contact from the anatomy surface, which require additional exposure of the anatomy and can be time-consuming and error-prone. The alternative is to use the imaged bone shapes, which are rigid, to perform the registration. This allows for less invasive procedures, is faster and less human-error prone, and does not require surgeon training.

Anatomy image-based rigid registration is technically much harder than fiducial or contact-based registration because it requires analyzing the intraoperative images. The images may include foreign objects such as surgical tools and implants not present in the preoperative data. Fluoroscopic X-ray images have a small field of view, limited resolution, and orientation-dependent geometric and intensity distortions. Research on anatomy image-based registration started in 1994 [8], [9] and is very active [10]–[20]. However, with the exception of the Cyberknife radiosurgery system [7], none of these methods is in routine clinical use. The main obstacles are robustness, accuracy, computation time, and lack of integration.

In this paper, we present a new gradient-based method for rigid registration of a patient preoperative CT to its intraoperative situation with a few fluoroscopic X-ray images acquired by a tracked C-arm [21]. The method is noninvasive, requires simple user interaction, and includes validation. It is generic and easily customizable for a variety of routine orthopaedic procedures. It consists of three steps: 1) initial pose estimation; 2) coarse geometry-based registration using bone contours; and 3) fine gradient projection registration (GPR) using edge pixels.

This hybrid method optimizes speed, accuracy, and robustness. Its novelty resides in using the relationship between CT and fluoroscopic X-ray image gradients instead of geometric or intensity information. Volume gradients and their projections help eliminate foreign objects present in fluoroscopic X-ray images and achieve higher accuracy. Our method overcomes the drawbacks of intensity-based methods, which are slow and have a narrow convergence range, and those of geometry-based methods, which depend on the contour segmentation quality of the fluoroscopic X-ray and CT images.

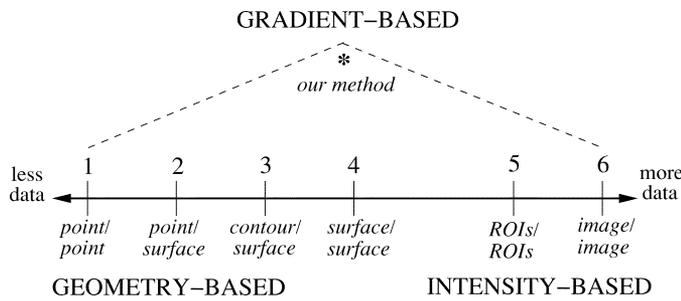


Fig. 1. Classification of rigid registration methods.

II. PREVIOUS WORK

For a comprehensive survey of medical image registration methods, see [22]. We classify rigid registration algorithms along a line, based on how much of the original data is used to compute the transformation (Fig. 1). On one side are geometry-based algorithms, which use a few selected points or features. On the other are intensity-based algorithms, which use most of the intensity information in both data sets.

Geometry-based registration algorithms match selected geometric features from each data set by finding the transformation that minimizes the sum of distances between paired features. Features can be implanted fiducials, anatomical landmarks, or surface contour features. The algorithms can be classified into four categories: 1) *point/point*: contact or image landmark points on anatomy surface to CT landmark points (5–10 points in each) [4]; 2) *point/surface*: contact “cloud of points” on anatomy surface to CT surface (10–30 points versus 10^5 – 10^6 surface points) [23]–[25]; 3) *contour/surface*: contours in fluoroscopic X-ray or ultrasound images to CT surface (10^2 – 10^3 points or 1–10 splines versus 10^5 – 10^6 surface points) [9], [11], [13], [26]; 4) *surface/surface*: CT or skin surface data from scanning laser to CT (10^5 – 10^6 points, 10^2 – 10^3 ridges/surfaces) [27].

Geometry-based registration consists of four steps: 1) feature extraction: choosing the features of interest in each data set; 2) feature pairing: establishing correspondences between features of each data set; 3) dissimilarity formulation and outliers removal: quantifying the dissimilarity between paired features, e.g., the sum of pairwise distances; and 4) dissimilarity reduction: finding the transformation that optimally minimizes the dissimilarity. Steps 2)–4) are repeated until convergence. Feature extraction requires segmenting the CT and fluoroscopic X-ray images. Features are paired by finding for each feature in one data set, the closest (distance-wise) feature on the other data set [23], [24]. Removal of outliers can be either explicit [24] or implicit by weighing the paired features [25]. Algorithms for geometry-based point to CT registration include the iterative closest point (ICP) algorithm [23], [24] and its variations [25], [28]–[30]. X-ray contours to CT surface mesh algorithms have been developed by Hamadeh *et al.* [9], [11] and Guéziec *et al.* [13], [26]. Fitzpatrick *et al.* [31] show how to estimate the registration error of point-based rigid registration.

The key characteristic of geometry-based methods is that they use a small fraction of the image data, usually fiducial centers and anatomy surface points, whose location is assumed to be

known very accurately. Geometry-based registration works best with a high quality segmentation, an efficient feature pairing scheme, and a good outlier removal. For point/point registration, an initial position independent, closed-form solution that minimizes the sum of distances is known [32]. In practice, robustness is achieved by first performing coarse registration with landmarks followed by fine registration with surfaces. Commercial systems rely on implanted fiducials and on “cloud of points” to perform the registration. Geometry-based registration between fluoroscopic X-ray images and CT has not yet reached the market, most likely because robustly segmenting fluoroscopic X-ray images is technically challenging.

Intensity-based algorithms match the intensities of one image data set with the intensity of the other by minimizing a measure of difference between them, such as histogram difference, normalized cross-correlation, or mutual information [33]–[35]. The matching can be restricted to regions of interest (ROIs) in the image, such as regions around bone surfaces in CT and fluoroscopic X-ray images. In this case, the matching is closer to geometry-based registration. Algorithms can be classified into two categories: 1) *ROIs/ROIs*: ROIs in both data sets, usually in the vicinity of the anatomy surface (10^5 – 10^6 pixels/voxels) for each data set [10], [12], [14], [36], [16]–[18], and 2) *image/image*: the entire CT or X-ray image is used (10^6 – 10^7 for CT data sets, 2–10 images, each about 10^6 pixels for X-ray images) [8].

Intensity-based registration consists of three steps: 1) generation of digitally reconstructed radiographs (DRRs) for each camera pose; 2) measurement of the pose difference by comparing the DRRs with the real fluoroscopic X-ray images; and 3) computation of a pose that reduces the difference. The first step requires precomputation and fast DRR generation [15], [36]. The second step requires computing a similarity measure which is not guaranteed to lead to an optimal solution [37]. Algorithms for intensity-based registration between X-rays and CT started with Lemieux *et al.* [8], which was followed by many others [10], [12], [14], [16], [17], [36], [37]. The Cyberknife radio-surgery system [7] is the only commercial system in routine clinical use that uses this registration method.

The key characteristic of intensity-based registration is that it does not require segmentation. The rationale is that using as much information as available and “averaging it out” reduces the influence of outliers and is, thus, more robust. However, this approach is computationally expensive since it requires generating high-quality DRRs and searching a six-dimensional (6-D) space with local minima which depend on the similarity measure employed. It requires an initial pose guess close to the final pose and the definition of ROIs.

Very recent work, conducted independently to ours, describes gradient-based registration between X-rays and CT or MR images [20], [38]. The idea is to compute projections of the volumetric data gradients, compare them with X-ray image gradients, and adjust the volumetric data set pose accordingly. The gradients need not be computed on all rays, but rather on selected rays in the vicinity of the anatomy contours, as proposed in [39]. While the idea of using gradients to establish the correspondence is similar to ours, the algorithm described in [20] relies on the pairing between rays emanating from the camera focal point and passing through image pixels and points on the

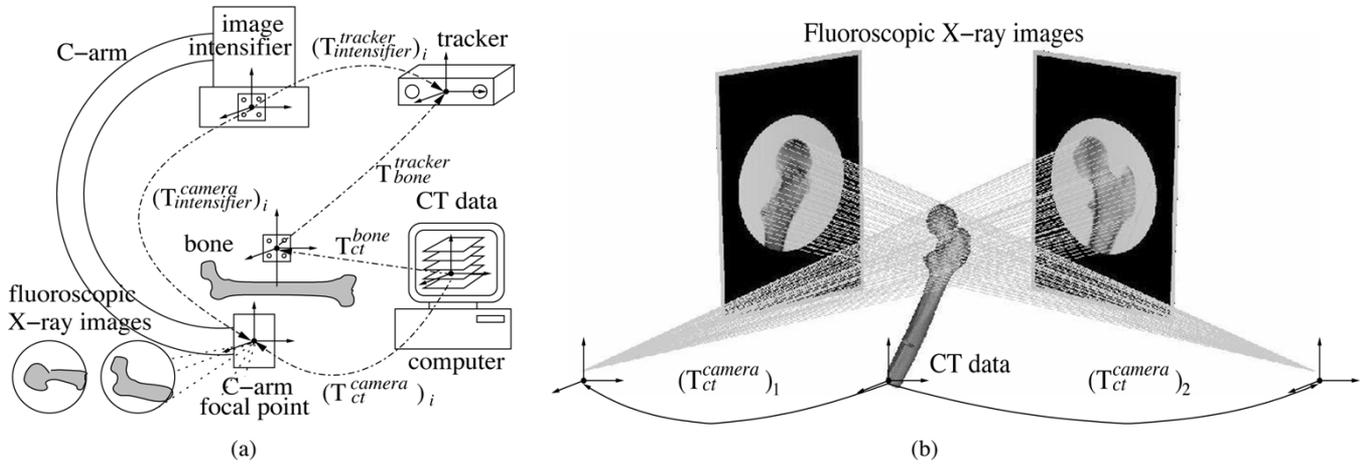


Fig. 2. (a) Registration chain between the actual bone and its preoperative CT and (b) 2-D/3-D image registration geometry between X-ray and CT.

bone surface. This pairing is very sensitive to discontinuities on both data sets and can produce outliers, which degrade the accuracy of the computed transformation. The experimental results reported in [20] assume distortion-less X-ray images, initial pose guesses near the final pose, and do not account for tracking errors, as our experiments do.

III. GOALS AND SPECIFICATIONS

Our goal is to develop a practical anatomical image-based rigid registration protocol and algorithm for preoperative CT to intraoperative fluoroscopic X-ray registration. The method should be generic and easily customizable to a variety of rigid anatomical structures (pelvis, vertebra, femur, tibia) and conditions (healthy, fractured, with tumors). Following a careful analysis of the most common orthopaedic procedures, we compiled the following specifications.

The system requirements are: 1) accuracy: a target registration error of 1–1.5 mm on average (2- to 3-mm worst case) measured on the bone surface; 2) robustness: the registration succeeds on the first try at least 95% of the time with an error of at most 2 mm; 3) speed: the registration process takes at most 1 min; 4) user interaction: simple and minimal preoperative and intraoperative user interaction; and 5) validation: both qualitative and quantitative after the registration.

The data characteristics are: 1) a CT data set with 0.5- to 1.5-mm-thick slices 1–3 mm apart, 12-bit gray-scale, and pixel size of $0.5 \times 0.5 \text{ mm}^2$ or less; 2) two to five fluoroscopic X-ray images 800×600 pixels, 8-bit gray-scale and pixel size of $0.5 \times 0.5 \text{ mm}^2$ or less, possibly including anatomy and surgical objects not present in the CT; and 3) C-arm position and orientation are computed using an optically tracked target which is rigidly attached to the C-arm image intensifier. Target position is known to an accuracy of 0.3–0.5 mm at a distance of 1–2 m (the current performance of commercial optical trackers).

The system consists of: 1) a PC with a monitor; 2) a video frame grabber; 3) a position sensor for the C-arm (e.g., an optical tracker); and 4) a calibration and distortion correction grid mounted on the C-arm.

IV. PROBLEM DEFINITION

The problem consists of finding the rigid transformation $T_{\text{ct}}^{\text{bone}}$ that relates the preoperative CT bone model to the intraoperative bone coordinate frame. This transformation can be obtained with a location tracker and a fluoroscopic X-ray imaging system by constructing the transformations chain shown in Fig. 2(a). The C-arm is modeled as a pinhole camera, with the camera focal point at the X-ray source and the image plane at the image intensifier. Since the imaging characteristics of the C-arm are orientation dependent, this calibration is computed anew for every orientation [40]. Fig. 2(b) illustrates the registration geometry.

The transformation chain consists of five transformations

$$(T_{\text{ct}}^{\text{camera}})_i = (T_{\text{intensifier}}^{\text{camera}})_i \cdot (T_{\text{intensifier}}^{\text{tracker}})_i^{-1} \cdot T_{\text{bone}}^{\text{tracker}} \cdot T_{\text{ct}}^{\text{bone}}$$

where i indicates the C-arm viewpoint. $(T_{\text{intensifier}}^{\text{tracker}})_i$ and $T_{\text{bone}}^{\text{tracker}}$ are given directly by the tracker. $(T_{\text{intensifier}}^{\text{camera}})_i$ is the orientation-dependent calibration transformation, which is computed as described in [40]. $(T_{\text{ct}}^{\text{camera}})_i$ are the desired transformations relating the camera poses with the bone model as shown in Fig. 2(b). $T_{\text{ct}}^{\text{bone}}$ is the initial bone pose estimate, which is successively refined with the two-dimensional (2-D)/three-dimensional (3-D) rigid registration algorithm described in Section V until convergence.

V. REGISTRATION PROTOCOL AND ALGORITHM

The registration protocol is as follows. Preoperatively, we obtain the CT data and automatically compute off-line from it three data structures: a bone surface mesh, a bounding sphere octree, and a volume gradient vector field. Intraoperatively, before the surgery starts, the tracking system is set up and the calibration grid is mounted on the C-arm. Then, the patient is prepared, and 2–5 fluoroscopic X-ray images from various orientations are taken. The fluoroscopic X-ray images are corrected for distortion and the camera parameters are computed for each pose [40]. The transformation $T_{\text{ct}}^{\text{bone}}$ is then computed with the algorithm described below. For validation, the algorithm shows

a projection of the volume gradients projected onto the fluoroscopic X-ray images. When there is close correspondence, the CT data set pose is very close to the actual patient pose.

The algorithm consists of three steps: 1) initial pose estimation; 2) coarse geometry-based registration on the bone contours; and 3). fine GPR on edge pixels. Each step has a funnelling effect: it brings the data sets closer, has a narrower convergence range, uses more information, and is more accurate than its predecessor. The first two steps, which we describe first, are based on previous work.

VI. VOLUME GRADIENT PROJECTIONS

The initial pose can be obtained in several ways, depending on the type of surgery and data available: 1) from the clinical setup, which usually indicates the position of the patient on the operating table (e.g., supine, on the side) and the C-arm imaging views (e.g., anterior-posterior, lateral); 2) intraoperatively, by having the surgeon touch implanted fiducials or by acquiring landmark points on the anatomy surface with a tracked pointer; 3) by placing skin markers prior to the CT scan and having the surgeon touch them with a tracked pointer intraoperatively; and 4) intraoperatively by having the user identify a few matching landmarks on X-ray images, estimating their actual location as the intersection of rays, and performing weighted point-based registration as described in [25]. Options 1, 3, and 4 are appropriate for percutaneous procedures. Regardless of the method employed, the initial pose estimate is usually within 10–20 mm and 5° – 15° of the final pose.

Coarse registration further reduces the distance between the bone surface mesh and sampled points on the fluoroscopic X-ray bone contours with the ICP method [23], [24]. It yields the best transformation that can be obtained from the segmented images, which provide an estimate of the real contour location and have occlusions and may contain foreign objects. GPR further reduces the difference by incorporating contour pixel and volume gradient data. It eliminates foreign objects which appear in the X-ray images and not in the CT data. It is more efficient than intensity-based registration with predefined ROIs, although it has a narrow convergence range.

Coarse geometry-based registration computes a transformation that positions the bone surface mesh such that the rays emanating from the camera focal point passing through the bone contours on the fluoroscopic X-ray images are tangent to it. It optimizes the distances between the rays and apparent bone surface mesh contour. The bone contour is extracted from the fluoroscopic X-ray images using a livewire segmentation algorithm [41]. Then, 100–500 contour points are sampled and matched to the corresponding bone surface mesh points with the 2-D/3-D ICP registration method [13]. The basic operation is to find for each ray the closest point on the bone apparent contour.

To speed up the search for the nearest point on the apparent contour, we construct a hierarchical structure, called the bounding sphere octree, in which we place the bone surface mesh edges (Fig. 3). Each edge holds the normal information of its coincident faces. The tree is recursively constructed as follows: initially, the entire bone surface mesh is enclosed in the smallest bounding sphere [42], [43]. The sphere's bounding

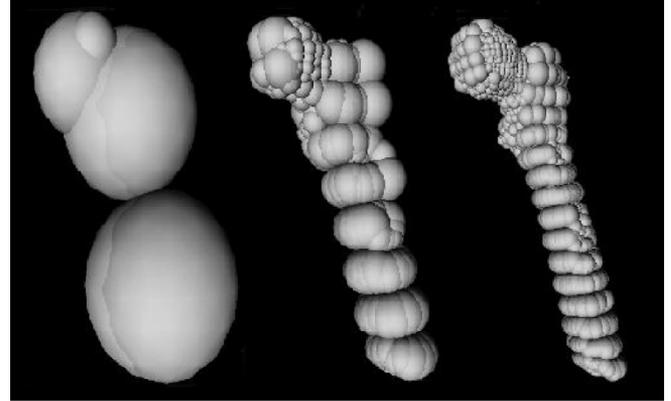


Fig. 3. Successive smallest bounding sphere approximations of a proximal femur surface mesh.

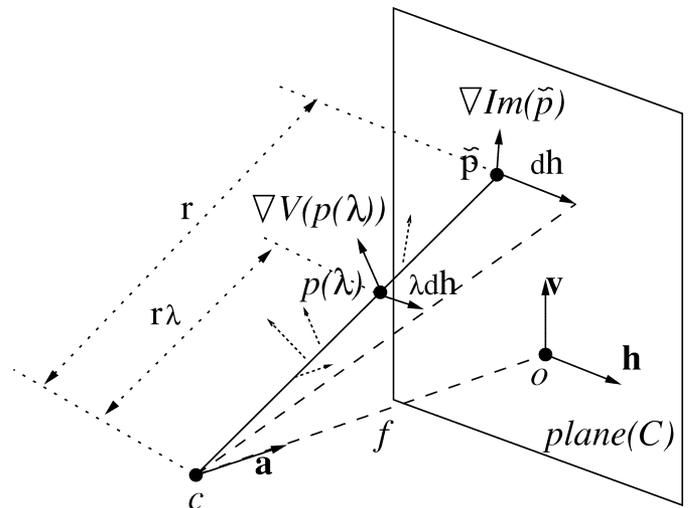


Fig. 4. The integral gradient projection property.

box is decomposed into eight cells and the edges are placed in the containing cells. Mesh edges which belong to several cells are split into segments, with each segment placed in its corresponding cell. The minimal bounding sphere for the edges in each cell is computed and is recursively subdivided until the sphere size is below a predefined threshold. The closest point to a ray is found by traversing the tree and using a priority queue according to the ray sphere distance, considering only the edges which are part of the apparent contour.

We now present the gradient projection property, which is the basis of the GPR step. We define the following entities:

$R \subset \mathbb{R}^3$	region of space containing the imaged anatomy;
V	CT image of the anatomy in region R ;
Im	X-ray image of the anatomy in region R ;
$p = (x, y, z)$	point in region R ;
$V(p)$	CT density value of point p ;
$Im(p)$	X-ray intensity value of the image point corresponding to p ;
$\nabla V(p)$	3-D gradient field of $V(p)$;
$\nabla Im(p)$	2-D gradient field of $Im(p)$.

We model the imaging process as follows (Fig. 4). The pin-hole camera C is defined by its focal length f positioned at focal

point c and oriented in viewing direction \mathbf{a} . The image plane of camera C is $plane(C) = (o, \mathbf{h}, \mathbf{v})$, where o is the image principle point and \mathbf{h} and \mathbf{v} are the horizontal and vertical directions of the image plane ($\mathbf{a} = \mathbf{v} \times \mathbf{h}$). We denote by \tilde{p} the projection of a point p on the image plane $plane(C)$. The ray $l(p)$ emanates from the camera focal point c and passes through point p . A point $p(\lambda)$ on this ray is given by the line equation $p(\lambda) = \lambda\tilde{p} + (1 - \lambda)c$ where $p(0) = c$ and $p(1) = \tilde{p}$. The distance between the point \tilde{p} and the camera focal point c is $r(\tilde{p}) = \|\tilde{p} - c\|$.

Gradient Projection Theorem: The image gradient of a point \tilde{p} in the image plane is linearly proportional to the integral of the weighted volume gradients of points $p(\lambda)$ along the ray emanating from the camera focal point and passing through the point \tilde{p}

$$\nabla Im(\tilde{p}) \propto \begin{pmatrix} \mathbf{h}^T \\ \mathbf{v}^T \end{pmatrix} \cdot r(\tilde{p}) \int_{\lambda} \lambda \cdot \nabla V(p(\lambda)) d\lambda. \quad (1)$$

Proof: Following the physical model of X-ray propagation presented in [44], the ratio between the number of photons that enter and exit the imaged object for a given ray is

$$N_{\text{out}} = N_{\text{in}} e^{-\int_{\text{ray}} \mu(x,y) ds}$$

where N_{out} and N_{in} are number of exiting and entering photons, $\mu(x, y)$ is the material attenuation coefficient, and ds is an element of length along the ray.

In our context, the number of exiting photons logarithmically corresponds to the pixel intensity value of an image point. The number of photons entering the imaged object corresponds to the initial intensity of the ray, which is constant and equal for all rays. The attenuation coefficient per element of length along the ray corresponds to the intensity value of the CT voxels along the ray, $V(p)$, where p is a point on the ray. Substituting into the above equation, we obtain

$$N_{\text{out}} = N_{\text{in}} e^{-\int_{p \in \text{ray}} V(p) ds}. \quad (2)$$

Let $s = \|p(\lambda) - c\|$ be the distance between the camera focal point and a point along the ray. s can be written as

$$s = r(\tilde{p})\lambda$$

differentiating, we get

$$ds = r(\tilde{p})d\lambda.$$

Substituting into (2) and omitting the constant N_{in} we get

$$N_{\text{out}} \propto e^{-r(\tilde{p}) \int_{\lambda} V(p(\lambda)) d\lambda}.$$

Assuming a standard logarithmic sensor response, the above equation becomes

$$Im(\tilde{p}) \propto r(\tilde{p}) \int_{\lambda} V(p(\lambda)) d\lambda$$

where $Im(\tilde{p})$ is the pixel intensity value of image point \tilde{p} . For every point along the projection ray, $p(\lambda) = \lambda\tilde{p} + (1 - \lambda)c$. Thus

$$Im(\tilde{p}) \propto r(\tilde{p}) \int_{\lambda} V(\lambda\tilde{p} + (1 - \lambda)c) d\lambda.$$

The image gradient is defined as the vector of partial image derivatives in the horizontal and vertical directions

$$\nabla Im(\tilde{p}) = \begin{pmatrix} \frac{\partial}{\partial \mathbf{h}} Im(\tilde{p}) \\ \frac{\partial}{\partial \mathbf{v}} Im(\tilde{p}) \end{pmatrix}.$$

The partial derivative of the X-ray image in direction \mathbf{h} is defined as

$$\frac{\partial}{\partial \mathbf{h}} Im(\tilde{p}) = \lim_{dh \rightarrow 0} \frac{Im(\tilde{p} + dh) - Im(\tilde{p})}{\|dh\|}.$$

Substituting, we obtain

$$\begin{aligned} \frac{\partial}{\partial \mathbf{h}} Im(\tilde{p}) &\propto \\ \lim_{dh \rightarrow 0} \int_{\lambda} \frac{r(\tilde{p} + dh)V(\lambda(\tilde{p} + dh) + (1 - \lambda)c) - r(\tilde{p})V(\lambda\tilde{p} + (1 - \lambda)c)}{\|dh\|} d\lambda &\propto \\ r(\tilde{p}) \lim_{dh \rightarrow 0} \int_{\lambda} \lambda \frac{V(p(\lambda) + \lambda dh) - V(p(\lambda))}{\|\lambda dh\|} d\lambda &\propto \\ r(\tilde{p}) \int_{\lambda} \lambda \mathbf{h}^T \cdot \nabla \mathbf{V}(p(\lambda)) d\lambda & \\ \mathbf{h}^T \cdot r(\tilde{p}) \int_{\lambda} \lambda \cdot \nabla V(p(\lambda)) d\lambda. & \end{aligned}$$

Similarly, the partial derivative of the X-ray image in direction \mathbf{v} is

$$\frac{\partial}{\partial \mathbf{v}} Im(\tilde{p}) \propto \mathbf{v}^T \cdot r(\tilde{p}) \int_{\lambda} \lambda \cdot \nabla V(p(\lambda)) d\lambda.$$

Combining the two expressions we obtain

$$\nabla Im(\tilde{p}) \propto \begin{pmatrix} \mathbf{h}^T \\ \mathbf{v}^T \end{pmatrix} \cdot r(\tilde{p}) \int_{\lambda} \lambda \cdot \nabla V(p(\lambda)) d\lambda.$$

The X-ray image gradient is equal to the integral over the weighted projections of the volume gradient onto the image plane, where the weight is the relative distance of the 3-D point from the focal point. Note that the weight increases as the 3-D point is further away from the focal point because variations in the 2-D image are a result of larger variations in the 3-D volume.

VII. FINE GRADIENT PROJECTION REGISTRATION (GPR)

We perform fine registration based on the gradient projection property. This step is based on the following observation: when the CT is aligned with the anatomy in the world, the rays emanating from the camera focal point that pass through contour pixels in the fluoroscopic X-ray images are tangent to the bone surface, as illustrated in Fig. 2(b). In this case, these rays pass through local magnitude maxima of the 3-D gradient vector field, since they are tangent to the surface. The desired transformation is, thus, the one that maximizes the sum of 3-D gradient magnitudes which are incident on these rays.

Formally, let T be a 6-D pose transformation matrix, \mathbf{h} and \mathbf{v} the horizontal and vertical directions of the image plane (Fig. 4). Let $p(\lambda)$ be a point on a ray emanating from the camera focal point and passing through an edge pixel \tilde{p} , in the

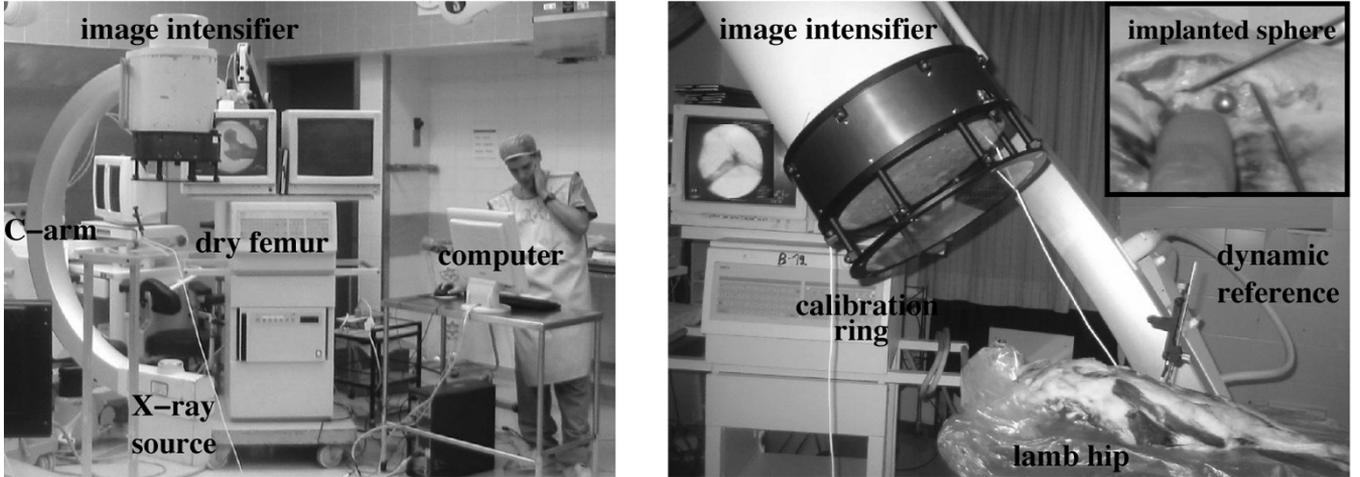


Fig. 5. Experimental setup: *in vitro* (left) and cadaver study (right). The insert on the upper right corner shows the implanted aluminum sphere.

fluoroscopic X-ray image. The expected fluoroscopic X-ray gradient, $GP(T, \tilde{p})$, at edge pixel \tilde{p} for a given pose T is, according to (1)

$$GP(T, \tilde{p}) \propto \begin{pmatrix} \mathbf{h}^T \\ \mathbf{v}^T \end{pmatrix} \cdot T \int_{\lambda} \nabla V(T^{-1}(p(\lambda))) d\lambda \quad (3)$$

(we omit $r(\tilde{p})$ and λ to speed up the computation since their influence on the optimization is minor).

The goal is to find the transformation \hat{T} that maximizes the sum of gradient projections over all image edges \tilde{p} , that is

$$\hat{T} = \arg \max_T \sum_{\tilde{p} \in \text{edges}} \|GP(T, \tilde{p})\|. \quad (4)$$

The fine GPR step is as follows. Preoperatively, we compute the CT volume gradient by convolving it with a Gaussian derivative kernel and up-sample it to a 0.5 mm^3 resolution to obtain a fast, high-quality nearest neighbor ray sampling. Intraoperatively, we extract edge pixels from each fluoroscopic X-ray image with the Canny edge detector [45] and construct the set of rays emanating from the camera focal point and passing through the pixels. We then apply the Downhill Simplex method [46] on the function defined in (4). The value of $GP(T, \tilde{p})$ is computed by sampling each ray in 1 millimeter intervals.

To achieve high accuracy, it is essential to filter out outlier edge pixels from the fluoroscopic X-ray images. Outlier edges are edges from foreign objects or from other anatomical structures. Outlier edges that are far from the anatomy of interest will be automatically filtered out: the gradient projection value of their rays will be small because there is no corresponding object in the CT. Outliers that are close to the anatomy of interest are eliminated by comparing the direction of the gradient projection and the actual gradient. When the directions diverge, the magnitude of the gradient projection is set to zero.

GPR combines the advantages of both geometry and intensity-based registration while overcoming their deficiencies. Like the geometry-based approach, it uses only edge pixels (from both the inner and outer bone contours). These pixels are only a small fraction of all pixels, so the computation time is significantly reduced. Unlike it, it does not rely on

segmentation. Like the intensity-based approach, it selectively uses all the CT information, without relying on segmentation or pairing between fluoroscopic X-ray pixels and CT voxels. Unlike it, it automatically defines focused ROIs, which speeds up the computation.

VIII. EXPERIMENTAL RESULTS

We have implemented the gradient based algorithm and have validated it with the proposed protocol on three types of situations: 1) simulation experiments with clinical CT data and simulated fluoroscopic X-rays; 2) *in vitro* experiments with dry bones; and 3) two cadaver experiments. The simulation experiments establish a lower bound on the error for ideal fluoroscopic X-ray imaging and tracking conditions and show how the algorithm copes with soft tissue and partial occlusions. The *in vitro* experiments establish a lower bound on the error for real CT and fluoroscopic X-ray images for ideal conditions. The cadaver experiments emulate the surgical situation and establish the expected error for intraoperative navigation with CT images. To demonstrate the generality of our method, we applied it to four different structures: human femur, spine, pelvis, and lamb hip.

We used a CT scanner, a 9 in BV29 C-arm (Phillips, The Netherlands), a Polaris optical tracking camera (NDI, Calgary, AB, Canada), a FluoroTrax C-arm calibration ring and active optical trackers (Traxtal, Toronto, ON, Canada), and a Matrox Meteor II digital frame grabber. Processing was on a 2.4-GHz, 1-GB RAM PC running Windows XP. Fig. 5 shows the experimental setup.

A. Registration Error Measurement and Validation

To quantify the registration error, we use the target registration error (TRE) as defined in [31]. The target registration error is defined as the distance between the actual and the computed position of selected target features, which can be landmark points or the bone surface itself. The difficulty in estimating the TRE lies in determining the actual position of the targets, which itself is prone to measurement errors. The most accurate but expensive and cumbersome method is to use a custom mechanical device which allows controlled precise positioning

of the anatomy. The next best option is to use implanted fiducials (spheres) and perform point-to-point registration between their centers on the CT images and their actual centers as measured by direct contact with a tracked pointer. This establishes a ground-truth transformation to which the computed transformation can then be compared. We opted, as many others, for this option in our experiments.

The fiducials TRE (fTRE) is defined as the distance between the actual and computed position of the fiducials [31]

$$\text{fTRE}(p, T_{\text{computed}}) = \|T_{\text{computed}}p_{\text{ct}} - p_{\text{measured}}\|$$

where p is a fiducial point in space, p_{ct} are its coordinates in the CT image coordinate frame, p_{measured} are its coordinates as measured by the tracking system, and T_{computed} is the computed transformation. The accuracy of the fTRE measurement depends on the accuracies of the tracking device, of the fiducial center localization on the CT, and on the computed transformation. Its advantages are that it can be computed with one or more fiducials with no restrictions on their relative positions and that it does not require the computation of a ground-truth transformation. Its disadvantages are that it requires fiducials, that it depends on very few points and their spatial distribution, and that it is only an indirect estimate of the error that the surgeon will observe when using the CT images for intraoperative navigation.

We propose an alternative measure, the surface TRE (sTRE), which we define as the distance between the actual and the computed position of points on the bone surface identified on the CT image

$$\text{sTRE}(p, T_{\text{computed}}, T_{\text{gold}}) = \|(T_{\text{computed}}p_{\text{ct}} - T_{\text{gold}}p_{\text{ct}})\|$$

where T_{gold} is the ground-truth (gold) rigid transformation computed by fiducial contact-based registration.

Although the sTRE is relative to the ground-truth registration, it does not require additional implanted fiducials, including instead many points uniformly distributed over the entire anatomy surface which do not require individual actual measurement. We quantify the expected error of the ground-truth registration with the method described in Fitzpatrick *et al.* [31] for all points on the bone surface. Although in the worst case this error should be added to the sTRE results, the error will be smaller in most cases since it also depends on the same optical tracker inaccuracies.

The sTRE includes all the errors in the registration chain, thus providing a faithful estimate of tool positioning errors during intraoperative navigation based on CT images. In our experiments, the ground-truth bone position for the *in vitro* and cadaver experiments are obtained from implanted fiducials by contact-based registration. The ground-truth transformation for the simulation experiments is known in advance. Note that the sTRE can also be used to quantify how far the initial guess is from the final one by substituting the initial guess transformation T_{initial} instead of the computed transformation T_{computed} .

For validation, we show the bone model position with respect to the actual bone position by overlaying the gradient projections magnitude directly onto the fluoroscopic X-ray images [Fig. 6(b) and (c)]. This shows how far the bone model is from

where it should be. The overlaid edges are computed as follows. For each pixel in the fluoroscopic X-ray image, we create a ray starting at the camera focal point and passing through the pixel. For each such ray, we compute its gradient projection using the CT volume gradient, which is the expected 2-D gradient direction and magnitude for the pixel. Nonmaxima suppression and thresholding on these pixels leave the expected edge pixels which are overlaid on the original fluoroscopic X-ray image. Note that this approach yields more accurate results than generating DRRs from the CT at the computed position and then extracting bone contours from them, since edges in the DRR are usually more blurred than actual fluoroscopic images [47].

B. Experiments

We performed a simulation experiment on a real clinical pelvis CT. We first generated DRRs at known poses and input these as fluoroscopic X-ray images, together with an initial guess transformation with a realistic sTRE of 9.5 mm, to the algorithm. We then computed the final sTRE error as described above. The first entry in Tables I and II shows the results averaged over ten runs. The mean measured error is 0.5 mm, about the size of X-ray image pixel.

Since the reported sTRE for all other cases depends on the accuracy of the ground-truth registration, we performed fiducial point-based registration and applied the method described in Fitzpatrick *et al.* [31] to compute the expected error over the entire imaged anatomy. We obtained an expected average accuracy of 0.3-0.5 mm (0.4-mm to 0.7-mm maximum) for all the points on the bone surface.

We performed *in vitro* experiments on a single vertebra of a dry spine and on a dry proximal femur. First, we implanted seven 6-mm aluminum spheres (Fig. 5 top right insert) and CT scanned them at 0.6-mm slice interval. We extracted from each data set the sphere centers at a resolution of 0.1 mm. In the operating room, we acquired two sets of three fluoroscopic X-ray images at various C-arm orientations, one with, and one without anatomy for optimal camera calibration. The fluoroscopic X-ray images were 800×600 pixels, 8-bit gray-scale with pixel size of 0.45 mm^2 . We performed C-arm calibration to mean accuracy of 0.3 mm (0.6-mm maximum), as described in our previous work [40]. We performed fiducial contact-based registration on the spheres and established the ground-truth registration. We then performed image-based registration with the gradient based algorithm and compared the resulting transformations.

The second and third entries in Tables I and II show the results. The 20-s to 50-s computation time for the ideal case increased to 60–270 s for the worst case when foreign objects and surrounding anatomy appeared in the fluoroscopic X-ray images. However, the accuracy error is still acceptable. We observed a small decrease in error when using three fluoroscopic X-ray images instead of two, and no further significant decrease beyond three. We observed little or no influence when foreign objects were present in the fluoroscopic X-ray images.

We performed cadaver experiments on a fresh lamb hip and a human pelvis following the same protocol as in the *in vitro* experiment, except that we implanted four spheres instead of seven. The last entries in Tables I and II show the results. For the lamb hip, the decrease in accuracy as compared to the *in vitro*

TABLE I

SUMMARY OF EXPERIMENTAL RESULTS. EACH SCENARIO (IDEAL, REALISTIC, AND BAD) IS DEFINED BY THE CT SLICE SPACING Δh (MM) AND THE PRESENCE OF FOREIGN OBJECTS IN THE FLUOROSCOPIC X-RAY IMAGES (NONE, SOME). EACH ENTRY SHOWS THE MEAN (MAXIMUM) SURFACE TRE IN MILLIMETERS. COMPUTATION TIMES ARE 20–370 s

Data set	# of X-ray images	Ideal $\Delta h = 0.6$ no noise	Realistic $\Delta h = 2.4$ some noise	Bad $\Delta h = 4.2$ some noise
Simulation				
1. real pelvis	5	—	0.5 (0.8)	0.6 (0.8)
In-vitro				
2. dry vertebra	3	0.4 (0.5)	0.5 (0.6)	0.6 (1.0)
3. dry femur	3	0.6 (0.8)	1.3 (1.5)	1.3 (1.6)
Cadaver				
4. lamb hip	4	1.2 (1.6)	1.4 (2.5)	1.7 (3.0)
5. human pelvis	4	0.9 (1.2)	1.7 (2.6)	2.6 (4.7)

TABLE II

DETAILED EXPERIMENTAL RESULTS FOR THE REALISTIC CASES (CT SLICE SPACING OF 2.4 MM, SOME FOREIGN OBJECTS IN THE FLUOROSCOPIC X-RAY IMAGES), WITH THE NUMBER OF IMAGES AS SHOWN IN TABLE I

Data set	Step	Final sTRE (mm)	Δx (Δy Δz) (mm)	$\Delta \theta_x$ ($\Delta \theta_y$ $\Delta \theta_z$) (degrees)	time (secs)
simulation					
1. real pelvis	initial	9.5	5.3 3.5 4.8	0.8 1.2 1.3	5.3 139.9
	coarse	1.7	2.7 1.3 2.4	0.6 0.7 0.6	
	fine	0.5	0.2 0.3 0.2	0.1 0.1 0.2	
in vitro					
2. dry vertebra	initial	9.6	3.7 6.4 4.5	3.1 3.9 3.1	4.6 268.9
	coarse	3.6	3.5 3.7 3.5	3.5 3.5 1.9	
	fine	0.5	0.5 0.1 0.4	0.2 0.1 0.2	
3. dry femur	initial	9.8	3.9 4.6 4.8	6.6 4.2 4.6	3.8 58.3
	coarse	2.3	1.6 1.3 1.4	3.1 1.0 3.2	
	fine	1.3	0.3 0.7 0.9	1.2 0.6 1.4	
cadaver					
4. lamb hip	initial	9.8	5.0 5.6 5.0	1.9 0.9 1.7	24.2 36.4
	coarse	2.1	4.1 1.5 4.1	4.1 4.1 8.0	
	fine	1.4	4.8 1.3 1.9	4.4 0.5 3.7	
5. human pelvis	initial	9.8	4.4 5.0 6.5	1.8 1.2 0.5	4.7 153.4
	coarse	6.5	5.5 2.9 5.5	5.5 5.5 2.5	
	fine	1.7	2.0 0.4 0.3	0.7 0.1 0.4	

For each data set, three rows of results are shown, one after the initial guess registration step (first row), one after the coarse geometry-based registration step (second row) following it, and one after the fine gradient-projection registration step following it (third row). Each row of results shows the average over ten runs sTRE error at the end of the step, the individual position (Δx , Δy , Δz) and orientation parameters ($\Delta \theta_x$, $\Delta \theta_y$, $\Delta \theta_z$) and the running time in seconds. The real pelvis, dry femur, and lamb hip have a success rate of 100% and the dry vertebra and human pelvis have a success rate of 70%. Note that a small angular deviation from the correct transformation can yield a large sTRE, depending on the location of the origin. The origins of the dynamic reference frames are at a distance of 80–150 mm from the surface of the anatomy.

case is most likely due to a less accurate ground-truth registration and the fact that the lamb femur has fewer salient features than the other anatomical structures. For the human pelvis, the decrease in accuracy is due to the larger size of the pelvic bone as compared to the other structures.

To better understand and quantify the various aspects of the proposed method, we conducted an extensive series of experiments on the dry femur data set. To determine the influence of

the initial pose guess on the coarse and fine registration steps, we applied the gradient based algorithm on many initial pose guesses between 1 and 85 mm sTRE from the ground-truth registration. For each 1-mm interval, we randomly generated 50 initial positions and performed the registration. To isolate registration error from the ground-truth error, we also computed the fTRE for 5–7 representative implanted fiducials. Fig. 7 summarizes the results of the 4250 runs. We conclude that the mean registration accuracy is nearly independent of the initial guess, but that the percentage of failures (registrations with sTRE > 2 mm) increases as the initial guess is further away. A mean failure rate of 5% occurs at 72-mm sTRE (16-mm maximum).

To better understand the characteristics of the fine gradient-based projection registration search space, we recorded the value of the optimization function $GP(T, \tilde{p})$ as the search converges toward the final transformation. Fig. 8 shows the results for the human pelvis cadaver study (the other cases are very similar). The plots show a unique minimum near the ground-truth value, an appropriate convergence range, and relatively smooth, monotonically decreasing values. This validates our choice of optimization function and search method.

IX. CONCLUSION

We have presented the GPR algorithm, a new method for rigid registration of a patient preoperative CT to its intraoperative situation with a few fluoroscopic X-ray images obtained with a tracked C-arm. The three-step hybrid method progressively brings the data sets closer with landmark point-based registration, coarse geometry-based registration on the bone contours, and fine GPR on edge pixels. Each step uses more information than its predecessor, has a narrower convergence range, is slower than its predecessor, but adds accuracy. The last step, GPR, which exploits the volume gradient projection property achieves good accuracy even in the presence of other anatomical structures and foreign objects, such as implants and surgical tools in the fluoroscopic X-ray images. It does not rely on the accuracy of segmentation, as do geometry-based approaches, and is more efficient than intensity-based registration, although it has a narrow convergence range.

We conclude from our experimental results that the desired goal, e.g., 1- to 1.5-mm mean target registration error (2- to 3-mm maximum), obtained within 60 s 95% of the time with simple and minimal user interaction, including validation, with standard imaging and tracking equipment in clinical conditions is within reach. We achieved on the cadaver studies, for the lamb hip 1.4 mm (2.5-mm maximum) in 60–120 s 100% of the time and for the human pelvis 1.7 mm (2.6-mm maximum) in 200–270 s 70% of the time with some user interaction (CT processing, initial pose estimation and livewire bone contour segmentation on fluoroscopic X-ray images). For validation, we show how far the bone model is from where it should be by overlaying the bone edge contours directly onto the fluoroscopic X-ray images.

We plan to further reduce the GPR algorithm computation time with space leaping techniques adapted from volume rendering. We also plan to improve the robustness of the GPR algorithm by implementing genetic or simulated annealing tech-

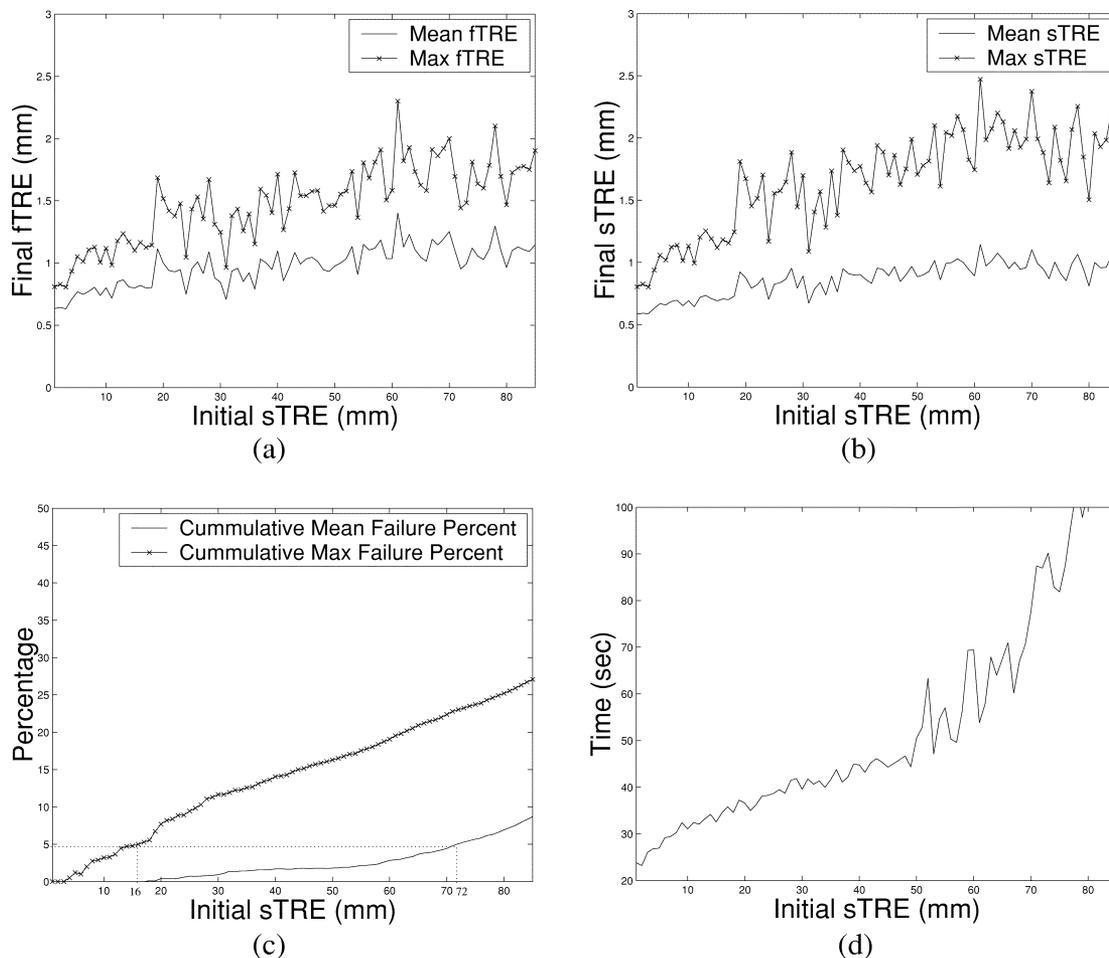


Fig. 7. Detailed results of *in vitro* dry femur experiments. The horizontal axis indicates the initial surface TRE from the ground-truth transformation. The vertical axes indicate (a) the final fiducial TRE; (b) the final surface TRE; (c) the cumulative percentage TRE failures (sTRE > 2 mm), and; (d) the average running time.

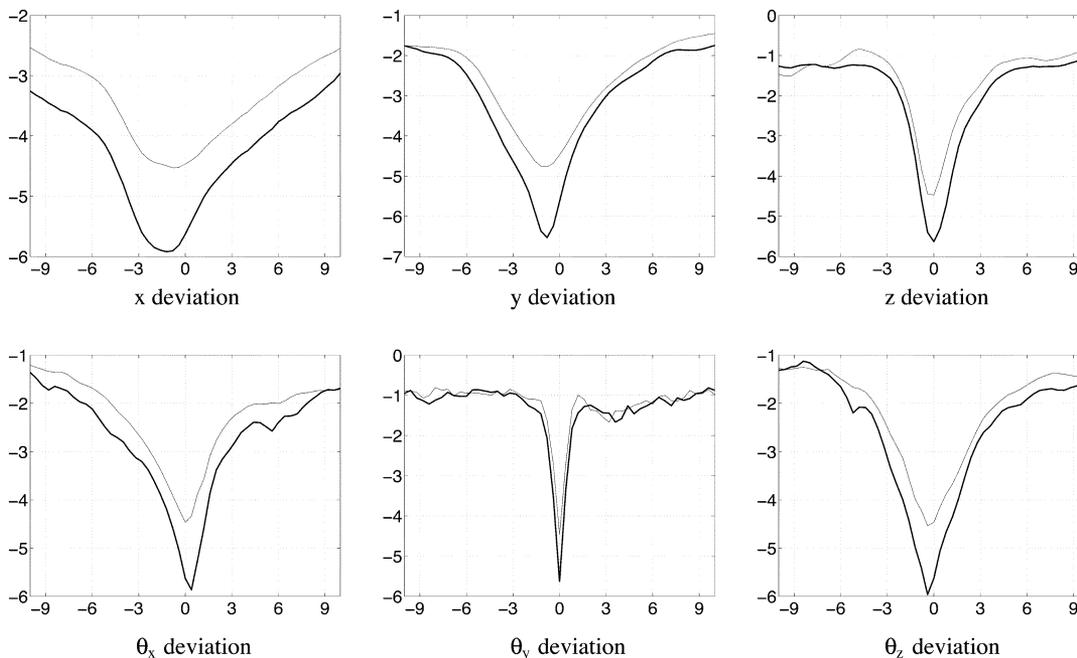


Fig. 8. Optimization function values (vertical axis) for individual pose parameters as a function of their deviation from the ground truth value, (zero for all parameters). Translational deviations are in millimeters, rotational deviations are in degrees. The lower dark curves correspond to the values for ideal conditions, the upper light curves correspond to the values for realistic conditions.

TABLE III
MAIN PARAMETERS AND THEIR SETTINGS USED IN THE EXPERIMENTS

Category	Parameter description	Default value
1. 3D model	Marching Cubes iso-value	1200
	Sphere Tree maximal depth	7
2. Livewire	# of points input by user	6–10
	# of contour points per image	200
3. ICP	Max # of iterations	1000
4. Fluoro X-ray Images	Gaussian mask (μ, σ)	(0,1)
	Lower threshold	lower 5% of grad. mag.
	Upper threshold	upper 10% of grad. mag.
	Min contour length	30 pixels
5. CT Volume	# of images	3–5
	Gaussian mask (μ, σ)	(0, $\frac{1}{1024}$)
6. Downhill Simplex Optimization	Ray sampling interval	1mm
	Contraction	0.5
	Expansion	2
	Reflection	1
	Max. # of iterations	5000
	Internal Simplex delta threshold	10^{-3}
	External Simplex delta threshold	10^{-3}

niques to help avoid local minima. We plan to embed the algorithm in navigation and positioning systems for minimally invasive and percutaneous orthopaedic procedures, including the FRACAS system for long bone fracture reduction [3], the MARS robot for percutaneous spinal pedicle screw insertion [48], and a new system for percutaneous pelvic fracture reduction.

APPENDIX

MAIN PARAMETERS AND THEIR SETTINGS

The proposed algorithm, as any of its kind, relies on tens of parameters with preset values. Table III lists the most important ones and their values. The parameters are classified in six categories, with the first four used in coarse geometry-based registration [Step 2]) and the last three used in fine gradient-based projection registration [Step 3]). No parameters are used for the initial registration.

The 3-D model parameters include the iso-value threshold used by the Marching Cubes algorithm to segment out the bone surface and the maximum depth of the hierarchical sphere tree described at the end of Section V. The Livewire parameters include the number of points input by the user on each fluoroscopic X-ray image and the number of points sampled on the segmented bone contours. The ICP parameter is the maximum number of iterations. For processing the fluoroscopic X-ray image parameters, we use the Canny edge detector. As input to the edge detector we specify a Gaussian mask which is used for computing the image gradient and two thresholds which are defined relative to the gradient magnitude data. As a post processing step to the edge detection we discard edges which are shorter than 30 pixels. For computing the gradient of the CT image, we use a Gaussian filter and sample the image with rays

at 1-mm intervals. Finally, in the gradient based registration step we use the downhill simplex optimization algorithm.

ACKNOWLEDGMENT

The authors would like to thank Dr. A. Khoury, Prof. M. Leibergall, and Dr. R. Mosheiff of the Hadassah Medical Center, Ein Karem, Jerusalem, for providing the CT data sets and for their invaluable assistance in preparing and performing the *in vitro* and the cadaver experiments.

REFERENCES

- [1] A. M. DiGioia, D. A. Simon, B. Jaramaz, M. Blackwell, F. Morgan, R. V. O'Toole, and B. Colgan, "Hipnav: Preoperative planning and intraoperative navigational guidance for acetabular implant placement in total hip replacement surgery," in *Computer Assisted Orthopedic Surgery*, L. P. Nolte and R. Ganz, Eds. Seattle, WA: Hogrefe & Huber, 1999.
- [2] R. Hofstetter, M. Slomczykowski, M. Sati, and L. P. Nolte, "Fluoroscopy as an imaging means for computer assisted surgical navigation," *Comput. Aided Surg.*, vol. 4, no. 2, pp. 65–76, 1999.
- [3] L. Joskowicz, C. Milgrom, A. Simkin, L. Tockus, and Z. Yaniv, "FRACAS: A system for computer-aided image-guided long bone fracture surgery," *Comput. Aided Surg.*, vol. 3, no. 6, pp. 271–288, 1998.
- [4] R. H. Taylor, B. D. Mittelstadt, H. A. Paul, W. Hanson, P. Kazanzides, J. F. Zuhars, B. Williamson, B. L. Musits, E. Glassman, and W. L. Bargar, "An image-directed robotic system for precise orthopaedic surgery," *IEEE Trans. Robot. Automat.*, vol. 10, pp. 261–275, 1994.
- [5] K. Smith, K. Frank, and R. Bucholz, "The neurostation: A highly accurate, minimally invasive solution to frameless stereotactic neurosurgery," *Comp. Med. Imag. Graphics*, vol. 18, no. 1, pp. 247–56, 1994.
- [6] Q. H. Li, L. Zamorano, A. Pandya, R. Perez, J. Gong, and F. Diaz, "The application accuracy of the neuromate robot – a quantitative comparison with frameless and frame-based surgical localization systems," *Comput. Aided Surg.*, vol. 7, no. 2, pp. 90–98, 2002.
- [7] J. R. Adler Jr, M. J. Murphy, S. D. Chang, and S. L. Hancock, "Image-guided robotic radiosurgery," *Neurosurgery*, vol. 44, no. 6, pp. 1299–1307, 1999.
- [8] L. Lemieux *et al.*, "A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs," *Med. Phys.*, vol. 21, no. 11, 1994.
- [9] A. Hamadeh, P. Sautot, S. Lavallée, and P. Cinquin, "Towards automatic registration between CT and X-ray images: Cooperation between 3D/2D registration and 2D edge detection," in *Proc. Medical Robotics and Computer Assisted Surgery Conf.*, 1995, pp. 39–46.
- [10] M. J. Murphy, "An automatic six-degree-of-freedom registration algorithm for image-guided frameless stereotactic radiosurgery," *Med. Phys.*, vol. 24, no. 6, 1997.
- [11] A. Hamadeh, S. Lavallée, and P. Cinquin, "Automated 3-Dimensional computed tomographic and fluoroscopic image registration," *Comput. Aided Surg.*, vol. 3, no. 1, 1998.
- [12] M. Roth, C. Brack, R. Burgkart, and A. Czopf, "Multi-view contourless registration of bone structures using a single calibrated X-ray fluoroscope," in *Proc. Computer-Assisted Radiology and Surgery Conf.*, 1999, pp. 756–761.
- [13] A. Guéziec *et al.*, "Providing visual information to validate 2D to 3D registration," *Med. Image Anal.*, vol. 4, no. 4, 2000.
- [14] G. P. Penney, P. G. Batchelor, D. L. G. Hill, D. J. Hawkes, and J. Weese, "Validation of a 2D to 3D registration algorithm for aligning preoperative CT images and intraoperative fluoroscopy images," *Med. Phys.*, vol. 28, no. 6, 2001.
- [15] D. A. LaRose, "Iterative X-ray/CT Registration Using Accelerated Volume Rendering," Ph.D. dissertation, Robotics Inst., Carnegie Mellon Univ., Pittsburgh, PA, 2001.
- [16] L. Zöllei, W. E. L. Grimson, A. Norbash, and W. M. Wells III, "2-D-3-D rigid registration of X-ray fluoroscopy and CT images using mutual information and sparsely sampled histogram estimators," in *Proc. IEEE Computer Vision and Pattern Recognition Conf.*, 2001, pp. 696–703.
- [17] D. Sarut and S. Clippe, "Geometrical transformation approximation for 2D/3D intensity-based registration of portal images and CT scan," in *Proc. Medical Image Computing and Computer Assisted Intervention Conf.*, 2001, pp. 532–540.

- [18] T. Rohlfing and C. R. Maurer Jr, "A novel image similarity measure for registration of 3-D MR images and X-ray projection images," in *Proc. Medical Image Computing and Computer Assisted Intervention Conf.*, 2002, pp. 467–476.
- [19] B. Brendel, S. Winter, A. Rick, M. Stockheim, and H. Ermert, "Registration of 3D CT and ultrasound datasets of the spine using bone structures," *Comput. Aided Surg.*, vol. 7, no. 3, pp. 146–155, 2002.
- [20] D. Tomaževič, B. Likar, and F. Pernuš, "Rigid 2D/3D registration of intraoperative digital X-ray images and preoperative CT and MR images," in *Proc. SPIE Medical Imaging Conf.*, 2002, pp. 507–517.
- [21] H. Livyatan, "Calibration and Gradient-Based Rigid Registration of Fluoroscopic X-ray to CT, for Intra-Operative Navigation," masters thesis, Sch. Eng. Comput. Sci., The Hebrew Univ. Jerusalem, Jerusalem, Israel, 2003.
- [22] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, no. 1, pp. 1–37, 1998.
- [23] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, pp. 239–255, Feb. 1992.
- [24] Z. Zhang, "Iterative point matching for registration of free-form curves and surfaces," *Int. J. Comput. Vis.*, vol. 13, no. 2, pp. 119–148, 1994.
- [25] C. R. Maurer Jr, G. B. Aboutanos, B. M. Dawant, R. J. Maciunas, and J. M. Fitzpatrick, "Registration of 3-D images using weighted geometrical features," *IEEE Trans. Med. Imag.*, vol. 15, pp. 836–849, Dec. 1996.
- [26] A. Guéziec, P. Kazanzides, B. Williamson, and R. H. Taylor, "Anatomy-based registration of CT-scan and intraoperative X-ray images for guiding a surgical robot," *IEEE Trans. Med. Imag.*, vol. 17, pp. 715–728, Oct. 1998.
- [27] S. Lavallée, R. Szeliski, and L. Brunie, "Anatomy-based registration of three-dimensional medical images, X-ray projections, and three-dimensional models using octree-splines," in *Computer-integrated surgery, Technology and clinical applications*, R. H. Taylor, S. Lavallée, G. C. Burdea, and R. Mösges, Eds. Cambridge, MA: MIT Press, 1995, ch. 7.
- [28] B. Ma, R. E. Ellis, and D. J. Fleet, "Spotlights: A robust method for surface-based registration in orthopedic surgery," in *Proc. Medical Image Computing and Computer Assisted Intervention Conf.*, 1999, pp. 936–944.
- [29] O. Sadowski, Z. Yaniv, and L. Joskowicz, "Comparative *in vitro* study of contact and image-based rigid registration for *Comput. Aided Surg.*," *Comput. Aided Surg.*, vol. 7, no. 4, 2002.
- [30] R. Bächler, H. Bunke, and L. P. Nolte, "Restricted surface matching-numerical optimization and technical evaluation," *Comput. Aided Surg.*, vol. 6, no. 3, 2001.
- [31] J. M. Fitzpatrick, J. B. West, and C. R. Maurer, Jr., "Predicting error in rigid-body, point-based registration," *IEEE Trans. Med. Imag.*, vol. 17, pp. 694–702, Oct. 1998.
- [32] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Amer. A*, vol. 4, no. 4, pp. 629–642, Apr. 1987.
- [33] W. M. Wells, III, P. Viola, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," in *Proc. Medical Robotics and Computer Assisted Surgery Conf.*, 1995, pp. 55–62.
- [34] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multi-modality image registration by maximization of mutual information," in *Proc. Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA '96)*, 1996, pp. 14–22.
- [35] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Image registration by maximization of combined mutual information and gradient information," *IEEE Trans. Med. Imag.*, vol. 19, pp. 809–814, Aug. 2000.
- [36] D. A. LaRose, J. Bayouth, and T. Kanade, "Transgraph: Interactive intensity-based 2D/3D registration of X-ray and CT data," *SPIE Image Processing*, 2000.
- [37] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. G. Hill, and D. J. Hawkes, "A comparison of similarity measures for use in 2-D-3-D medical image registration," *IEEE Trans. Med. Imag.*, vol. 17, pp. 586–595, Aug. 1998.
- [38] D. Tomaževič, B. Likar, and F. Pernuš, "Gradient-based registration of 3D MR and 2D X-ray images," *Comput. Assist. Radiol. Surg.*, pp. 327–332, 2001.
- [39] K. G. A. Gilhuijs, P. J. H. van de Ven, and M. van Herk, "Automatic three-dimensional inspection of patient set-up in radiation therapy using portal images, simulator images, and computed tomography data," *Med. Phys.*, vol. 23, pp. 389–399, 1996.
- [40] H. Livyatan, Z. Yaniv, and L. Joskowicz, "Robust automatic C-arm calibration for fluoroscopy-based navigation: A practical approach," in *Proc. Medical Image Computing and Computer Assisted Intervention Conf.*, T. Dohi and R. Kikinis, Eds., 2002, pp. 60–68.
- [41] E. N. Mortensen and W. A. Barrett, "Interactive segmentation with intelligent scissors," *Graph. Models Image Processing*, vol. 60, no. 5, pp. 349–384, 1998.
- [42] B. Gärtner, "Fast and robust smallest enclosing balls," in *Proc. Eur. Symp. Algorithms (ESA)*, 1999, pp. 325–338.
- [43] ———, (2002) Smallest Enclosing Balls of Points – Fast and Robust in C++. [Online]. Available: <http://www.inf.ethz.ch/personal/gaertner/miniball.html>
- [44] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Piscataway, NJ: IEEE Press, 1988.
- [45] R. Jain, R. Kasturi, and B. G. Schunk, *Machine Vision*. New York: McGraw-Hill, 1995.
- [46] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. Cambridge Univ. Press, Cambridge, U.K., 1999.
- [47] W. Cai, "Transfer functions in DRR volume rendering," in *Proc. Computer-Assisted Radiology and Surgery Conf.*, H. Lemke, et al., Ed., 1999, pp. 284–289.
- [48] L. Joskowicz, C. Milgrom, M. Shoham, Z. Yaniv, and A. Simkin, "Robot-guided long bone intramedullary distal locking: Concept and preliminary results," in *Proc. 3rd Int. Symp. Robotics and Automation*, Toluca, Mexico, 2002, pp. 463–468.