

2D-3D registration of coronary angiograms for cardiac procedure planning and guidance

Guy-Anne Turgeon

*Imaging Research Laboratories, Robarts Research Institute, London, Ontario, Canada
and Biomedical Engineering Program, The University of Western Ontario, London,
Ontario, Canada*

Glen Lehmann

Department of Medical Biophysics, The University of Western Ontario, London, Ontario, Canada

Gerard Guiraudon

*Canadian Surgical Technologies and Advanced Robotics (C-STAR), London Health Sciences Centre,
London, Ontario, Canada*

Maria Drangova

*Imaging Research Laboratories, Robarts Research Institute, London, Ontario, Canada
Department of Medical Biophysics, The University of Western Ontario, London, Ontario, Canada
and Department of Diagnostic Radiology and Nuclear Medicine, The University of Western Ontario,
London, Ontario, Canada*

David Holdsworth

*Imaging Research Laboratories, Robarts Research Institute, London, Ontario, Canada
Department of Diagnostic Radiology and Nuclear Medicine, The University of Western Ontario,
London, Ontario, Canada
and Department of Medical Biophysics, The University of Western Ontario, London, Ontario, Canada*

Terry Peters^{a)}

*Imaging Research Laboratories, Robarts Research Institute, London, Ontario, Canada
Department of Diagnostic Radiology and Nuclear Medicine, The University of Western Ontario,
London, Ontario, Canada
Department of Medical Biophysics, The University of Western Ontario,
London, Ontario, Canada
and Biomedical Engineering Program, The University of Western Ontario,
London, Ontario, Canada*

(Received 21 March 2005; revised 15 September 2005; accepted for publication 18 September 2005; published 21 November 2005)

We present a completely automated 2D-3D registration technique that accurately maps a patient-specific heart model, created from preoperative images, to the patient's orientation in the operating room. This mapping is based on the registration of preoperatively acquired 3D vascular data with intraoperatively acquired angiograms. Registration using both single and dual-plane angiograms is explored using simulated but realistic datasets that were created from clinical images. Heart deformations and cardiac phase mismatches are taken into account in our validation using a digital 4D human heart model. In an ideal situation where the pre- and intraoperative images were acquired at identical time points within the cardiac cycle, the single-plane and the dual-plane registrations resulted in 3D root-mean-square (rms) errors of 1.60 ± 0.21 and 0.53 ± 0.08 mm, respectively. When a 10% timing offset was added between the pre- and the intraoperative acquisitions, the single-plane registration approach resulted in inaccurate registrations in the out-of-plane axis, whereas the dual-plane registration exhibited a 98% success rate with a 3D rms error of 1.33 ± 0.28 mm. When all potential sources of error were included, namely, the anatomical background, timing offset, and typical errors in the vascular tree reconstruction, the dual-plane registration performed at 94% with an accuracy of 2.19 ± 0.77 mm. © 2005 American Association of Physicists in Medicine.
[DOI: 10.1118/1.2123350]

Key words: image-registration, cardiac surgery, mutual information, vascular

I. INTRODUCTION

Minimally invasive surgery (MIS) often involves the insertion of an endoscopic camera and laparoscopic surgical instruments through small incisions in the patient (port access). To improve dexterity and precision in MIS, surgeons may

employ robotic assistance.^{1,2} Unfortunately, the constrained manipulation of the robotic instruments, the limited field of view given by the endoscope, and the lack of anatomical information to allow the adequate planning of the procedure prevent the benefits of MIS from being fully applied to general patient populations.




FIG. 1. The registration of the preoperative 3D coronary model (3DCM) to the intraoperative patient environment will allow the accurate positioning of the virtual cardiac model to the beating heart in the patient's thorax.

The motivation of this work is to improve the minimally invasive robotic coronary artery bypass (MIRCAB)^{1,2} procedure and to provide support for the intracardiac treatment of atrial fibrillation,³ which involves the introduction of instruments into the beating heart. In spite of the progress in design and the complexity of the robotic systems used to assist the surgeon, current systems lack the capability to visualize the operative area in three dimensions (3D). This results in inaccurate port placement, excessive trauma to the patient, and increased patient risk. To address this issue, we have developed the virtual cardiac surgical planning (VCSP) platform.^{4–8} Within this environment, a patient-specific 3D representation of the heart and chest wall, derived from preoperative images, is registered to the patient in the operative position. The VCSP platform, similar to the simulation and transfer architecture for robotic surgery (STARS) system,^{9–11} will ultimately serve as a simulation, planning, and guidance tool to assist a variety of minimally invasive interventions.

The work described in this paper provides the link that will permit the accurate registration of a virtual cardiac model⁵ to the beating heart within the patient's thorax in the operating room (OR) (Fig. 1). We make the assumption that in the OR, a cardiac model derived from preoperative images can be approximately registered to the patient using markers placed on the skin,⁹ or by an ultrasound-based registration of the ribs of the patient to those in the model.¹² Such a registration step provides a suitable starting point for the automated 2D-3D coronary artery registration technique described in this paper.

X-ray angiography is commonly employed as a real-time imaging modality during surgery, and many researchers have fused intraoperative projection images with preoperatively acquired 3D computed tomography (CT) or magnetic resonance (MR) images. Most of the approaches used to register projection images to 3D volumes have been either feature-based or intensity-based. Feature-based methods^{13–16} are rapid and robust, but rely on accurate segmentation. In orthopedics, implant and bone contours are common features used for the registration,^{13–15} while in neurovascular interventions, both the vessel contours¹³ or the centerlines¹⁶ may be employed as registration features. Intensity-based techniques^{17–23} have the advantage of not depending on segmentation accuracy, but are computationally more demand-

ing. Penney *et al.*,^{19,20} Hipwell *et al.*²¹ and Byrne *et al.*²² worked extensively on the development of 2D-3D automated intensity-based registration algorithms to update the 3D position of interventional instruments using a single-plane angiogram registered to a 3D volume. These groups first implemented their algorithm to register preoperative CT vertebral images intraoperatively^{19,20} and then expanded their method to register 3D cerebral magnetic resonance angiography (MRA) (Ref. 21) and 3D x-ray digital subtraction angiography (3D-DSA) (Ref. 22) images. From their MRA study, 95% of the phantom and 82% of the clinical registrations were successfully registered with rms reprojection errors of 1.3 ± 0.6 and 1.5 ± 0.9 mm, respectively. For the 3D-DSA, the use of two registered images from the same modality improved the registration accuracy to 1.3 ± 0.6 mm in a clinical study. A third category of 2D-3D registration comprises hybrid techniques, which combine feature- and intensity-based properties.^{24–28} Recently, Mahfouz *et al.*²⁷ and Vermandel *et al.*²⁸ introduced hybrid registration approaches that bear similarities to the work described here. Prior to implementing their registration algorithms, the two groups derived 3D mesh-based surface models, which were then registered to one angiogram using binary projections rather than the digitally reconstructed radiographs (DRRs) used in intensity-based methods. Working on knee implant registration, Mahfouz *et al.* used an implant computer-aided design model to create their surface model and implemented a cost function combining intensity and contour matching scores.²⁷ From *in-vivo* experiments, a rms deviation from the optimal pose of 0.4° , 0.1 mm in-plane, and 1.4 mm out-of-plane was measured. For their part, Vermandel *et al.* worked with MRA/DSA images of a cerebral vessel phantom and employed the quadratic distance between pixel gray levels as similarity measure²⁸ to obtain an approximate 3D rms reprojected error of 1.8 mm.

In our 2D-3D registration method, a preoperative 3D mesh of the coronary arteries is registered to intraoperative angiograms. Heart motion creates difficulties for registration based on the coronary arteries prohibiting the use of DSA and complicating the validation procedure. The 2D-3D method presented here differs from previous methods, in that we are working with the moving coronary arteries and we ensure a robust and completely automated registration by

first processing the intraoperative angiograms into binary images using a multiscale segmentation. This segmentation step overcomes the low contrast of the coronary angiograms, by converting the images into a form similar to the projections of the 3D preoperative mesh model generated at each registration iteration. The motivation of our work is also different from much of the similar research reported in the literature,^{13–28} since the 2D-3D vascular registration is employed to register the 3D preoperative heart volume to the patient, rather than to map the 2D projection image to the 3D volume for the purpose of “road-mapping” during an interventional procedure.

In this paper, we describe our automated binary image-based 2D-3D registration approach, and demonstrate that a coronary artery registration at single time points within the cardiac cycle is robust and accurate. Using a realistic environment to simulate a clinical setting, we also compare the results of registering preoperative 3D images of the coronary arteries to both single- and dual-plane intraoperative angiograms.

II. METHODS

An accurate 2D-3D registration of the preoperative 3D coronary model with the intraoperative angiogram(s) (single- or dual-plane) intrinsically registers a preoperatively acquired geometric heart (or derived model) to the intraoperative situation. Since the vessels are constantly in motion, the registration is performed at end diastole, the most quiescent phase of the heart cycle. Both the preoperative and the intraoperative images must be acquired at this same time point within the cardiac cycle. The preoperative 3D vessel image has the form of a mesh dataset either extracted from gated MR, CT, or 3D x-ray rotational angiography images, using the marching cubes algorithm,²⁹ or reconstructed from two orthogonal angiograms as outlined by Chen and Carroll.³⁰ It is assumed in this study that the preoperative and intraoperative images are acquired at the same phase of the respiratory cycle.

A. Gold standard and clinical data simulation

In this section, we describe an environment that simulates the characteristics of a clinical setting, taking into account the different variables encountered in a typical cardiac angiographic suite. To simulate OR conditions for the intraoperative images, we included the effects of temporal accuracy, temporal sampling rate, and contrast agent injection.

The most important variable involved in the registration of cardiac images is motion. Electrocardiogram (ECG) gating can be used to ensure that the angiograms are acquired at the desired cardiac phase (end diastole). However, small temporal discrepancies can occur between the acquisitions. In addition, within the OR, single-plane angiography systems are more common than dual-plane, necessitating the acquisition of a second synchronized view.

For the purpose of this study we assumed that the x-ray detector on the clinical system was either fully corrected for distortion³¹ or that it had a flat panel detector. In addition, we

assumed that the angulation of the gantry can be determined accurately and can be eliminated as a potential source of significant error. In our study, the geometry of the angiographic system was used to calculate the four intrinsic parameters (u_o, v_o, k_1, k_2) describing skew and scaling. The skew parameters (u_o, v_o) correspond to the x and y coordinates of the normal from the imaging plane to the x-ray source. The perspective parameters (k_1, k_2) are the magnification in x and y generated by the projection. The field of view (FOV) of the image intensifier, the magnification factor (M), and the isocenter distance relative to the focal spot (source to object distance) were used for this calculation. Clinically, these intrinsic parameters can be calculated using a calibration cube³² or they can be estimated from the system geometry.²²

To validate our registration method, a 4D dataset derived from an excised human heart was used to create a simulation environment composed of realistic 2D and 3D images. This digital dataset consists of a 3D computed rotational angiography (CRA) image of an *ex-vivo*, contrast-perfused human heart,³³ animated using cine-angiograms acquired from a patient with similar coronary artery anatomy. The motion information derived from the cine-angiograms was applied to the 3D image using backprojection of homologous arterial bifurcation points. A nonlinear warping algorithm was then used to deform the 3D image between successive time points to create the complete heart cycle. The resulting digital heart model was composed of 30 frames per cardiac cycle, each of which was a $400 \times 400 \times 400$ volume with $540 \mu\text{m}$ isotropic voxels. This 4D dataset allowed us to extract a gold standard coronary model, simulate clinical timing offset, and digitally create angiograms that were combined with precontrast patient images to assure realistic clinical angiographic datasets.

1. 2D, 3D, and 4D models

The various models that are referred to in this paper are defined below:

- **4DBM** 4D base model—four-dimensional (space and time) model comprising separate left and right coronary arteries (LCA, RCA) acquired using CRA and animated with cine-angiograms.
- **3DCM** 3D coronary model—end diastole LCA or RCA mesh dataset described by the vessel centerlines and radius (the experimental gold standard).
- **3DCM+N** 3D coronary model with noise—3DCM with Gaussian distributed noise added to vessel centerlines to simulate clinical reconstruction inaccuracies.
- **SAM** Simulated angiogram no background mask—digitally reconstructed radiograph (DRR) generated from the 4DBM and segmented into a binary image.
- **SABM** Simulated angiogram with background mask—DRR from the 4DBM combined with a clinical patient angiogram background (precontrast image) and segmented into a binary image [Fig. 3(c)].

The digital heart was segmented into two distinct vascular models (4DBMs) containing the RCA and LCA branches,




FIG. 2. Generation of simulated angiogram with background by the addition of the DRR image to a real patient angiogram background retrieved from the patient 224 study. (a) Simulated angiogram with no background generated by the DRR technique. 60° LAO view of the LCA. (b) Real patient angiogram background or precontrast angiogram retrieved from patient 224 study. (c) Simulated angiogram with background resulting from the fusion of image (a) and (b).

respectively. In the simulation environment, the LCA and RCA 4DBMs represented the dynamic coronary branches of a patient. We selected the end-diastole frame (referred to as time T1 throughout this paper) to generate the preoperative 3DCM as described below.

2. 3DCM extraction

An underlying premise of this work is that a 3D coronary model of the vasculature has already been created preoperatively by vessel extraction of a 3D image,³⁴ or by angiographic reconstruction.³⁰ In our simulation and throughout our experiments, an equivalent model of such a vascular reconstruction, characterized by 3D centerline coordinates was used. To avoid confusion between the errors resulting from the 3D reconstruction and those due to the registration process, the 3DCM centerlines were semimanually segmented from the diastole time frame (T1) of the 4DBM. The marching cubes algorithm²⁹ was first applied to the eight-bit 3D image, with an intensity threshold chosen by visual inspec-

tion to minimize the number of unconnected vessel segments, while maximizing the vessel radii. Then, using the generated vessel surfaces, each centerline and its associated radius were manually extracted, resulting in the gold standard 3DCM used for validation.

3. Generation of simulated angiograms

Using the ray-casting method described by Siddon,³⁵ DRRs were generated from the LCA and RCA 4DBMs to create the simulated intraoperative angiograms [Fig. 2(a)]. Prior to ray casting, a user-selected threshold was applied to the 3D image to avoid the reprojection of background noise in the DRR. The images generated from the digital model approximated the right and left anterior oblique (RAO, LAO) views typically encountered in clinical practice.

To simulate realistic angiograms using these DRRs, we retrieved precontrast images from four typical patient studies [Fig. 2(b)] at acquisition angles similar to those from which the DRR images were obtained. These clinical angiograms

TABLE I. Nine datasets used to validate the registration. Dataset 0 is composed of DRR images without the addition of anatomical background. Datasets 1–8 are produced from the combination of DRR images and real patient angiogram backgrounds taken at matching acquisition angles (Fig. 2). Column 2 identifies the patient number from which the real patient angiogram backgrounds were taken, column 3 lists the arterial branch, and columns 4 and 5 present the RAO and the LAO angulation of the views used for the registration. The LAO view is used for the single-plane registration.

Datasets	Patient No.	Coronary artery branch	RAO angle (deg)	LAO angle (deg)
Dataset 0 (no background)	—	LCA	30	60
Dataset 1	240	LCA	30	60
Dataset 2	240	LCA	45	45
Dataset 3	240	LCA	0	90
Dataset 4	240	RCA	45	45
Dataset 5	226	LCA	30	60
Dataset 6	224	LCA	30	60
Dataset 7	222	LCA	30	60
Dataset 8	222	LCA	60	30

incorporate structured background such as soft tissues, ribs, catheterization instruments and stents, as well as background that embodies the effect of quantum noise, polychromatic beam energy, the heel effect, radiation scatter, and electronic noise in the simulated radiographs. Eight unique datasets were created (Table I) by combining these background images with the DRR images made from the model at similar acquisition geometries.

The fusion of the two images [Figs. 2(a) and 2(b)] was achieved by appropriately weighting their intensities prior to summing them and then adding an offset. Intensity histograms and pixel value comparisons between similar components were used to ensure that the simulated images [Fig. 2(c)] were indistinguishable from real clinical datasets. Figure 3(a) shows the simulated 60° LAO view of the LCA 4DBM added to the background from patient 224, while Fig. 3(b) shows the angle-matched real clinical angiogram of patient 224 after the injection of contrast agent.




FIG. 3. Pixel value similarity comparison between a simulated and a real angiogram. (a) Simulated angiogram with background of the LCA at 60° LAO view from dataset 6 created by the addition of a 4DBM DRR and patient 224 background (Fig. 2). (b) Clinical angiogram of patient 224 with contrast agent. (c) Resulting SABM of (a) after multiscale analysis. (d) Resulting binary mask image of (b) after multiscale analysis.

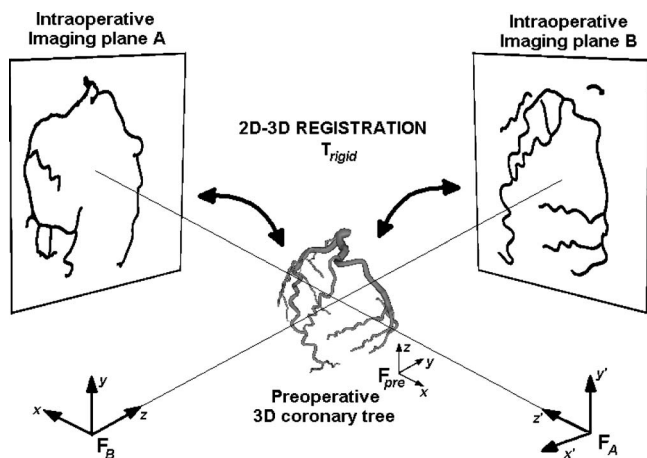


FIG. 4. 2D-3D registration of a preoperative 3DCM to two intraoperative angiograms that have been automatically segmented into binary masks.

B. Image registration

The 2D-3D registration algorithm repeatedly repositions the 3DCM in space and compares its projections to the intraoperative angiogram(s). At each iteration, a similarity measure is used to compare the 2D images, and the three translation and three rotation parameters are changed according to the results. This process continues until the similarity measure is maximized.

We considered registration based both on single- (1P) and dual-plane (2P) angiograms (Fig. 4). In the 2P case, at each iteration, the similarity results of both planes are added together to provide a global similarity measure. The unconventional aspect of our registration technique is its use of binary data instead of grey-scale images. To overcome the low vessel-to-background contrast ratio of only 3–15 % achieved in typical coronary angiograms, the intraoperative angiograms are processed with an automatic multiscale segmentation algorithm prior to registration. This segmentation results in a binary image where the detected vessels are 0's (black) and the background is 1's (white) [Figs. 3(c) and 3(d)]. This segmented image is then used as an intraoperative image for the registration.

1. Angiogram segmentation

The intraoperative simulated angiograms are automatically converted into binary masks³⁶ using the multiscale segmentation algorithm described by Koller *et al.*³⁷ and Sato *et al.*³⁸ Koller's filter searches for elongated structures of arbitrary shape, creating a sharp maximum at the centerline of the vessels. The multiscale nature of this filter makes it robust to different vessel sizes and also allows the detection of the vessel widths.

Koller's angiogram segmentation method proceeds in the following manner. First, the 2D image is processed with the Koller filter. The filter is applied at six different scales equally distributed between 2.5 and 4.0 pixels, which represents the range of radii of the searched vessels. For each scale, the filter's response is



FIG. 5. Comparison of a MBP as generated at each registration iteration and of a SABM. (a) MBP of the preoperative mesh 3DCM as generated at each registration iteration. LCA 60° LAO. (b) Simulated angiogram with background of the LCA at 60° LAO view from dataset 1 processed with the multiscale segmentation yielding the binary mask image or SABM. Note that the extra "branch" at the top of (b) (arrow) is a catheter that is part of the background, rather than part of the vascular system.

$$R_s(\vec{x}) = \min\{\text{Pos}[(\nabla f_s(\vec{x} + s\vec{d}) \cdot \vec{d})], \text{Pos}[-\nabla f_s(\vec{x} + s\vec{d}) \cdot \vec{d}]\}, \quad (1)$$

where the operator $\text{Pos}(\cdot)$ returns the positive values of (\cdot) , f is the original image, $f_s = f \otimes G_s$ is the scale-space image, \vec{x} is the pixel coordinate, \vec{d} the local direction vector, and s the scale parameter. Once the filtering is complete, a new image is created with the maximum response of each pixel throughout all the scales. The pixel scale value and local line orientation \vec{d} of each maximum are also tabulated. For each scale, \vec{d} is found such that the second derivative of the function $f_s(\vec{x}) = (f \otimes G_s)(\vec{x})$ is maximized.

The second step of the segmentation procedure suppresses nonmaximum values by detecting local maxima in the direction \vec{d} resulting in the creation of the thin centerlines. In step three, the images are thresholded and binarized. In step four, an "island removal" algorithm is used to accept only those interconnected segments comprising more than 65 pixels. Finally, in step five, the binary mask is created. The vessel diameters are determined by using the scale information of each centerline pixel and multiplying it by a factor of 1.5 (chosen by visual inspection), since as reported by Koller *et al.*³⁷ and as evident in our results, the filtering process is robust and efficient but tends to underestimate the vessel radii.

It is not necessary to compute DRRs for our 2D-3D registration, since the preoperative 3DCM is a mesh dataset. Instead, at each iteration, a binary projection of the 3DCM is generated using the intraoperative angiogram acquisition geometry. The resulting model binary projection (MBP) is a perspective view of the 3DCM mesh in which the vessels are 0's and the background is 1's [Fig. 5(a)]. While the 3D vessel thicknesses could be used to generate an intensity projection, the MBP is employed to maintain maximum compatibility between the two images being registered.

2. Similarity measure

Registration of the MBP and the SABM is performed with the entropy correlation coefficient (ECC) as the similarity

measure. ECC is a normalized form of the standard mutual information metric (MI) and uses information theory to quantify how well one image is explained by another. ECC can be defined as³⁹

$$\text{ECC} = \sqrt{2 \left(1 - \frac{2H(A,B)}{H(A) + H(B)} \right)}, \quad (2)$$

where $H(A)$ and $H(B)$ are the Shannon entropies of images A and B , respectively, and $H(A,B)$ is their joint entropy

$$H(A) = - \sum_a p_A(a) \log p_A(a), \quad (3)$$

$$H(A,B) = - \sum_a \sum_b p_{AB}(a,b) \log p_{AB}(a,b), \quad (4)$$

where $p_A(a)$ and $p_{AB}(a,b)$ are the single and joint probability densities of pixel values in image A , and in images A and B , respectively. Theoretically, the metric is maximized when the images are aligned or when they are maximally dependent. The square root in the definition is not necessary but it has the effect of magnifying the variations when the two images are independent.

We selected the ECC measure after evaluating the performance of six different similarity measures with respect to the 2P registration procedure. Mean absolute difference (MAD), mean square difference (MSD), and normalized cross correlation (NCC) were the three linear metrics⁴⁰ investigated, while we also employed the information theoretic metrics mutual information (MI), normalized mutual information (NMI), and ECC. A total of 80 registrations [eight patient datasets (Table I) \times 10 starting positions] were performed on full size images (512×512 pixels) using each similarity measure. Analysis of variance demonstrated that the statistical results from each metric were equivalent at the $p=0.05$ level. Nevertheless, ECC was chosen since it yielded the highest success rate (ECC=86%, NMI=79%, MI=76%, (MAD)=71%, (MSD)=69%, and NCC=68%) obtained in preliminary results where a single simplex optimization was performed (our notions of success and of single simplex optimization are explained in Secs. II C and III). While the ECC metric takes twice as long to compute as a simple linear metric, for our application the increase in performance outweighs the additional time required to complete the registration.

3. Optimization strategy

We employed the downhill simplex method^{41,42} to determine the optimal set of extrinsic parameters based on the chosen similarity metric. For a six degree of freedom (DOF) search strategy, where the DOFs represent the translations and rotations of the preoperative 3DCM (Fig. 4), the simplex is a geometric figure composed of seven points ($n\text{DOFs} + 1$) interconnected with line segments. At initialization, these seven points, or vertices, are defined based on P_o , an initial starting transform for the search. The six other vertices can then be derived relative to P_o using the unit vectors (\vec{e}_i) of the six degrees-of-freedom such as

$$P_i = P_o + \lambda \vec{e}_i, \quad (5)$$

where P_i is the new vertex and λ is the characteristic length scale of the problem. In our 2P registration algorithm, we set $\lambda=10$ mm for the three translations and $\lambda=10^\circ$ for the three rotations of the preoperative 3DCM. For the 1P registration, the same values are used except for the x and y translations where λ is increased to 40 mm. A larger search space must be set for these two axes since the single-plane centroid matching cannot evaluate the position of the 3DCM in the direction normal to the angiogram plane (x and y coordinates of the 3DCM). Starting from this initial position, the simplex makes its way through the complexity of the six dimensions until it finds at least a local minimum (defined by the negative of the similarity function). In this process, the geometric figure is iteratively altered according to the value of the similarity cost function at each of its seven points. By reflection, expansion, and contraction, the downhill simplex algorithm moves away from the highest points. All the parameters are considered simultaneously and convergence is reached when the seven points of the simplex are within a specified tolerance of each other.

The simplex algorithm, while found to be the best optimizer for our task, is nevertheless affected by the presence of local optima. For this reason, three successive reinitializations of the degenerated simplex^{42,43} were implemented, considerably increasing the percentage of registration successes. Each reinitialization starts at the last convergence point with the x , y , z translation and rotation λ s reduced to 5 mm and 5° since after the first registration, the model is usually closer to the global maximum. While the computational time is increased as a result of this implementation, as was noted when using the ECC similarity measure, the increased performance is worth the penalty.

4. Computing environment

Both our virtual visualization environment VCSP and our registration algorithm were implemented using the open-source Visualization Toolkit (VTK) (Ref. 44) class libraries, using C++ and the Python interpretive scripting language. Running on two 3 GHz Xeon processors with 1(GB) of RAM, the execution time of our registration algorithm averaged 17 sec for the 2P and 32 sec for the 1P registration. Single-plane registration took longer, since the optimizer was not given the constraint of the second image to reduce searching time around local optima. Also, in such circumstances it was necessary to reinitialize the simplex, which extended the processing time. To reduce the computational load and thus accelerate the algorithm by a factor of two, the iterative MBP images were generated from on-screen renderings by the NVIDIA GeForce FX 5900 graphic card through the use of VTK and OpenGL libraries.

The multiscale segmentation of the angiograms was processed independently from the registration, using the Insight Segmentation and Registration Toolkit (ITK) (Ref. 45) libraries in conjunction with VTK. Although not optimized, the generation of each binary mask took approximately one

minute when running on a desktop computer with a 2.6 GHz Intel Pentium 4 processor and 1 GB of RAM.

C. Algorithm implementation

One of the advantages of working with binary images is that it allows centroid matching prior to registration in order to bring the 3DCM closer to its final position. To optimize this centroid match, the FOV of the binary mask and of the MBP are doubled. This ensures that the model's projection falls within the FOV of the mask image even with gross initial misalignment of the source and target images. Once the 3D displacement is calculated, the FOVs are set back to their original values and the intensity-based registration begins. This centroid match increases the robustness of the registration and was applied for all of our experiments. However, this step is not required for studies where the position of the 3DCM is known to be very close (<10 mm) to the registered location.

The validation was performed by misaligning the 3DCM using ten randomly generated rigid transformation matrices, and then, recovering the original pose using our registration. The 3D translations and rotations for the ten matrices were randomly generated from Gaussian distributions whose standard deviation values were set by making the following assumptions: The translation error from the skin or rib registration is expected to be between 2–10 mm,¹¹ the center of the patient's heart is approximately 50–150 mm under the skin, and the heart translations in the thorax due to positioning and other effects can be as high as 35 mm. Furthermore, throughout the procedure, the images were expressed relative to an OR coordinate system defined using an optically tracked probe (NDI Polaris) with an accuracy of 0.5 mm. Based on these factors, we chose Gaussian distributions with $\sigma_t=25$ mm and $\sigma_r=6^\circ$ to create the ten misalignment matrices.

To evaluate the quality of a given registration with respect to the gold standard transformation, the 3D root mean square (3DRMS) measure between homologous points on the coronary tree centerlines was computed as

$$3DRMS = \sqrt{\frac{1}{N_{pts}} \sum_{i=1}^{N_{pts}} (\vec{p}_{i,gs} - \vec{p}_{i,reg})^2}, \quad (6)$$

where N_{pts} is the total number of pairs of homologous points, $\vec{p}_{i,gs}$ is the vector specifying 3D coordinates of the gold standard point i , and $\vec{p}_{i,reg}$ is the 3D position of the registered point i . The RCA gold standard 3DCM was described by six branches comprising 602 points, and the LCA gold standard 3DCM by 17 branches [eight branches for the left circumflex and nine branches for the left anterior descending] with a total of 1591 points. In each branch, the points were distributed uniformly by a spline generator used during the manual segmentation. For each registration, the 3DRMS was computed and its value determined whether or not the resulting transformation was successful. If the 3DRMS was greater than 5 mm, we considered the registration to have failed.

Note that the 3DRMS of our ten starting positions ranged between 26 and 68 mm.

D. Experimental trials

We investigated the effect of several independent error sources on registration accuracy. To examine the effect of the anatomical background, we created “dataset 0,” comprising two orthogonal DRR images without any “clinical background” to confuse the registration. Throughout our validation, dataset 0 was always compared with datasets 1–8 (Table I) generated from angiograms with patient background included, and in which a segmented catheter or rib was often identified as a vessel by our segmentation algorithm [Fig. 5(b)].

The next set of experiments evaluated the effect of temporal phase misregistration between the 3D preoperative image and the intraoperative angiograms; an important potential source of error in the registration. In our simulation, temporal mismatch between the pre- and the intraoperative image acquisitions was reproduced by selecting two images at times T1 and T2 spaced by ≈ 100 ms during the end-diastolic phase. We differentiated between these images by using the notation (') to designate an image acquired at time (T2), and no prime to designate an image acquired at T1. The 3DRMS deviation between the 18 homologous bifurcation points in the 4DBM and 4DBM' models was 3.0 mm.

Finally, acknowledging the fact that any 3D vascular model created from preoperative images exhibits its own errors, we added Gaussian noise with $\sigma=1.5$ mm to the centerline generated by our preoperative models in order to simulate the average 3D reconstruction errors reported in the literature.^{30,46–50}

III. RESULTS

The results of our experiments are presented in Table II and in Fig. 6. The 1P and the 2P registrations were investigated separately. For 1P registration tests, the LAO view of each dataset was used for the registration (see Table I).

On average, 1P registration needed one more reinitialization compared to the 2P approach to converge successfully to the correct solution (Table II). This indicated that the algorithm was more likely to converge to a local rather than global optimum when the second-plane constraint was missing.

A. Ideal situation

The first experiment investigated 1P and 2P ideal registrations with no added “patient background.” For the ten starting misalignments, the 1P registration performed successfully 90% of the time with an accuracy of 1.44 ± 0.15 mm, while the 2P method achieved a 100% success rate with an accuracy of 0.55 ± 0.04 mm.

For datasets 1–8, in which anatomically realistic backgrounds were added to the DRRs (leading to segmentation discrepancies in the SABM), the performance of the algorithm was slightly reduced. Out of the 80 1P registrations,

TABLE II. Results of the 2D-3D registration experiments. Column 3 gives the total number of registrations performed using Table I datasets and the ten starting positions. Column 4 lists the average number of restarts necessary for registration success. Column 5 gives the 3DRMS error of the successful registrations. Columns 6 and 7 present the time taken to perform the registration and the success rate, respectively.

Registration	Number of plane(s)	Number of registrations	Number of restart(s)	3DRMS error (mm)	Time (sec)	Success rate (%)
Ideal case						
3DCM—SAM	1	10	0.9	1.44±0.15	29	90%
3DCM—SAM	2	10	0.2	0.55±0.04	17	100%
3DCM—SABM	1	80	1.0	1.60±0.21	33	79%
3DCM—SABM	2	80	0.2	0.53±0.08	17	99%
10% Time offset						
3DCM—SAM'	2	10	0.3	1.37±0.17	23	100%
3DCM—SABM'	2	80	0.3	1.33±0.28	17	98%
10% Time offset+3D centerline 1.5 MM rms error						
3DCM+N—SAM'	2	10	0.3	2.42±0.82	22	90%
3DCM+N—SABM'	2	80	0.3	2.19±0.77	19	94%

79% were successful, with an average 3DRMS accuracy of 1.60 ± 0.21 mm. The 18 failed registrations converged into local optima mostly located along the out-of-plane axis. For the 2P registration, 99% of the 80 registrations were successful, with a mean 3DRMS errors of 0.53 ± 0.08 mm. The only failure occurred in dataset 3 in which the angiograms were acquired at 0° RAO and 90° LAO.

B. Anatomical “noise” and phase misregistration

Here we investigated the effect of cardiac phase misregistration between the 3D preoperative model and the intraoperative angiograms. The T1 preoperative 3DCM was registered to the SABM'(s), generated from the 4DBM' (at T2). For the 2P registration, the ten registrations of the mask images generated from angiograms without anatomical background, 3DCM to SAM', all converged with a final 3DRMS accuracy of 1.37 ± 0.17 mm. Of the 80 registrations with ana-

tomical background and segmentation discrepancies, 98% were successful with an average 3DRMS error of 1.33 ± 0.28 mm. The two failures that occurred in datasets 3 and 4 are discussed further in Sec. IV.

The 1P registration results are not presented in Table II since they failed to achieve a 3DRMS error < 5 mm for most trials. This result is explained by the deformation of the coronary tree between the two time frames due to motion (similar to a magnification), which results in the inability of the single plane to register the 3DCM accurately in the out-of-plane axis. For these experiments, when the optimizer converged successfully to the global optimum, the out-of-plane error was generally between 6 and 18 mm.

C. Full clinical simulation

In the above experiments, we employed the preoperative 3DCM as a gold standard model. Unfortunately, in a clinical




FIG. 6. 3DRMS error calculated between homologous points of the registered coronary tree and the gold standard. Accuracy results of the dataset without anatomical background and the eight “patient” datasets for: 1P and 2P ideal cases, 2P with time offset, and 2P with time offset and reconstruction noise.

case, the 3DCM has been shown to be less accurate.^{30,46–50}

To simulate the effect of reconstruction and other sources of error, Gaussian distributed noise was added to the vessel centerlines to generate the 3DCM+N dataset. This more realistic dataset was then registered to the phase misregistered SAM' and SABM' angiograms.

For the reasons described above, the results of the 1P registration are not tabulated. For the 2P registration, 9 of the 10 SAM' registrations were successful with an average 3DRMS of 2.42 ± 0.82 mm. For the 80 SABM' patient registrations, 94% converged toward the desired optima, with three failures occurring in dataset 3 and one in each of datasets 4 and 8. The 3DCM+N to SABM' target registration error was 2.19 ± 0.77 mm, which is approximately 1 mm greater than the registration results obtained without the added coronary model noise.

D. Comparison with 3D-3D registration

Another possible way to register the 3DCM to the patient's environment is to generate a second 3D vessel reconstruction in the OR, using two angiograms,³⁰ and perform a 3D-3D point-based registration. To test this method in our complete clinical simulation (phase misregistration and 3D reconstruction noise), we generated ten different 3DCM+N' models, selected the same misaligned matrices as were used in the previous section, and then performed a feature-based registration (homologous points) with the 3DCM+N model. From this experiment, a 3DRMS error of 2.16 ± 0.87 mm was obtained. While this accuracy is practically the same as the 2D-3D registration (3DCM+N to SABM') results, the drawbacks of the 3D-3D technique are the time involved prior to surgery to complete the coronary reconstruction (approximately 10 min), and the manual intervention needed to assure a good segmentation of the vessels from the 3D dataset.^{30,49–53}

IV. DISCUSSION

In this research we used binary images to overcome the problems caused by the low contrast of standard angiograms, in order to ensure a robust registration of the patient's pre-operative 3DCM to the its actual heart in the OR. The use of unsegmented images was investigated using our image-based method, but the success rate was very low due to the convergence of the optimizer into local optima. This outcome was expected due to the high anatomical background noise in the angiograms and the small number of voxels available for the similarity measure calculation (resulting in an irregular cost functions and small variations between the global and local optima).

From the registration success rate and accuracy results, the eight simulated patient datasets showed very similar results to those obtained with dataset 0, in which no anatomical background was added (Fig. 6). This demonstrates that the automatic segmentation is reliable and that the ECC similarity measure is robust to the segmentation discrepancies caused by the catheter, rib structures in the FOV, and missing vessels. Evaluating the ideal case results, the average accu-




FIG. 7. Plot of the cost function for each of the x , y , and z projection axes for LCA dataset 1, single plane 60° LAO view. For the in-plane axis (x and y), the cost function peak is at +2 mm and for the out-of-plane axis (z), it lies at -5 mm.

racy for the 1P registration between the nine datasets varied from 1.11 for dataset 3 to 2.98 mm for dataset 1. For the 2P registration, the performance was consistent across datasets with minimum and maximum accuracies of 0.44 and 0.60 mm, respectively. The absence of the second plane constraint caused the high variability of the 1P registration. Figure 7 shows a plot of projection axis x , y , and z cost functions for the 1P registered position (dataset 1 LAO view). For this dataset, the peak cost function of axis x and y is at +2 mm and the maximum of the z axis is at -5 mm. These results, representing the cost function variations versus the displacement of the volume in each of the three axes, illustrate the variation in the accuracy values. For the 2P registration, the addition of both plane cost functions almost eliminates this variability.

The accuracies reported are results from successful registrations, however some failures to register within 3DRMS of 5 mm did occur. For the 1P registration, especially when the timing offset was investigated, the algorithm had difficulty registering in the out-of-plane axis because of the motion of the vasculature (magnification effect). The other failures found in the ideal cases of the 1P registrations and the ideal and “noise-added” cases of the 2P registrations are explained by the presence of local optima. To illustrate this, we plot the cost function of dataset 8 LAO view as a function of x and y displacements (Fig. 8). The plot clearly shows the presence of sharp local optima of the ECC function surrounding the narrow peak of the misregistered position at a 20 mm out-of-plane displacement from the correct solution. When the second plane is used to constrain the search, the addition of the two-plane similarity measures generally reduces the amplitude and the number of local optima, resulting in accurate registration close to 100% of the time. However, even after three reinitializations, 3% of the 2P registrations still failed in local optima. One of these failures occurred in the 3DCM+N to SAM' registration where there was no anatomical background to affect the segmentation, but most of these failures (5) were found in dataset 3 where the angio-

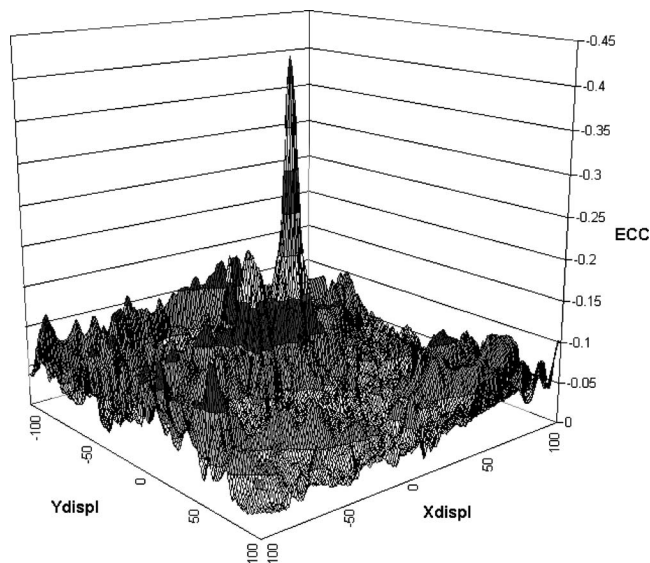


FIG. 8. 3D plot of dataset 8, single plane cost function for the 30° LAO view of the LCA. Cost function values for in-plane, x and y , at the out-of-plane offset of $z=20$ mm.

grams were acquired at 0° RAO and 90° LAO. The lower success rate with this dataset, compared to the other datasets, is explained by the following factors: the superimposition of the vessels in the 0° RAO view and the less efficient centroid matching of dataset 3. The compacted vessel data resulting from the segmentation of the superimposed vessels made the position search more difficult. In addition, the starting position of 17 mm from the correct solution after centroid match for dataset 3 (compared to the 9 mm average of the other eight datasets), increased the probability of the optimizer of converging to a local minimum.

All our validation results are expressed as 3DRMS accuracy. This error calculation is more significant than the agreement between the translation and rotation parameters recovered by the registration compared to the initial misalignment, since the 3DRMS measure accounts equally for the errors throughout the cardiac vascular system and is the expected source of registration error at the surface of the heart. Depending on the magnitude and the position relative to the origin of the evaluated points, the overall 3DRMS error is the result of different translations and rotations. Two typical examples comparing 3DRMS parameters to their respective 3D global translation and rotational components are taken from the average registration results of dataset 1. For the 2P ideal case, a $3DRMS=0.56\pm 0.10$ mm is equivalent to a 0.57 mm translation and a 0.09° rotation error. For the noise case, including both 10% time offset and 3D centerline reconstruction error, a $3DRMS=2.34\pm 0.77$ mm is equivalent to a 1.57 mm translation and a 1.55° rotation error. When our results are expressed either in terms of 3DRMS differences or differences in translations and rotations, the errors are always measured at the centerline points of the coronary vessels. These errors correspond to the target registration errors (TRE) in cases when the VCSP platform is used for coronary artery bypass procedures and the coronary vessels are the

surgical targets. In other cases, where the VCSP model is mapped to the patient by the use of the coronary registration but the surgical procedure is performed on another part of the heart (valves or atrium walls), the TRE will vary with the position within the thorax. However, inside the heart, the TRE will always be smaller than at the surface, since the origin of the rotation axes is located near the centroid of the heart.

V. CONCLUSION

In this paper we presented a technique for registering a patient's heart to preoperatively acquired images, based on a 2D-3D image registration technique employing intraoperatively acquired angiograms. The registration method was validated using a simulated clinical environment in which expected sources of noise, such as 3D reconstruction errors and 10% temporal errors between the pre- and the intraoperative acquisitions were modeled. The binary registration approach we employed was insensitive to the presence of imaged surgical and anatomical background such as interventional catheters, stents, soft tissue, and ribs. However, some intraoperative constraints must be imposed for this registration to be successful. An approximate registration of the preoperative images of the patient must initially be performed in order to position the preoperative cardiac model within the field of view of the intraoperative angiograms, the patient's heart rate must be sufficiently stable to minimize the temporal error from the ECG gating, and the system geometry must be well calibrated.

Both single- and dual-plane registration procedures were compared. For the ideal case, the 1P registration demonstrated an accuracy of 1.60 ± 0.21 mm while the 2P approach was accurate to within 0.53 ± 0.08 mm. However, when temporal error was added, the 1P method was unable to register the 3DCM along the out-of-plane axis. The 2P registration, was robust to the same simulated errors, giving consistent and accurate registration results of 1.33 ± 0.28 mm when the 3DCM was ideal, and 2.19 ± 0.77 mm when a Gaussian distribution 3DRMS reconstruction noise (standard deviation = 1.5 mm) was added. We believe that these results indicate that the 2P angiographic technique described here is an appropriate method for registering a preoperative model of the heart, including the vascular information, to a patient heart in the operating room.

ACKNOWLEDGMENTS

The authors would like to thank Marcin Wierzbicki, Dr. Mark Wachowiak, Dr. David Gobbi, and Atamai Inc. for helpful discussions, as well as Chris Wedlake for software assistance. G.A.T. acknowledges support from a studentship from the National Science and Engineering Research Council of Canada and a training fellowship in Vascular Imaging from The Canadian Institutes for Health Research (CIHR). This work was also supported by CIHR Grant No. MOP 14735, the Ontario Research and Development Challenge Fund, the Canadian Foundation for Innovation, and the Ontario Innovation Trust.

- ^a)Electronic address: tpeters@imaging.robarts.ca
- ¹D. Boyd, B. Kiaii, R. Novick, R. Rayman, S. Ganapathy, W. Dobkowski, G. Jablonsky, F. McKenzie, and A. Menkis, "RAVE-CAB: Improving outcome in off-pump minimal access surgery with robotic assistance and video enhancement," *Can. J. Surg.* **44**, 45–50 (2001).
 - ²D. Boyd, B. Kiaii, K. Kodera, R. Rayman, W. Abu-Khudair, S. Fazel, W. Dobkowski, S. Ganapathy, G. Jablonsky, and R. Novick, "Early experience with robotically assisted internal thoracic artery harvest," *Surg. Laparosc. Endosc. and Percutaneous Techniques* **1**, 52–57 (2002).
 - ³G. M. Guiraudon, D. L. Jones, A. C. Skanes, D. Bainbridge, C. M. Guiraudon, S. M. Jensen, X. Yuan, M. Drangova, and T. M. Peters, "En bloc exclusion of the pulmonary vein region in the pig using off pump, beating, intra-cardiac surgery: A pilot study of minimally-invasive surgery for atrial fibrillation," *Ann. Thorac. Surg.* **80**, 1417–1423 (2005).
 - ⁴A. Chiu, D. Dey, M. Drangova, W. Boyd, and T. Peters, "3-D image guidance for minimally invasive robotic coronary artery bypass (MIR-CAB)," *Heart Surgery Forum* **3**, 224–231 (2000).
 - ⁵M. Wierzbicki, M. Drangova, G. Guiraudon, and T. Peters, "Validation of dynamic heart models obtained using non-linear registration for virtual reality training, planning, and guidance of minimally invasive cardiac surgeries," *Med. Image Anal.* **8**, 387–401 (2004).
 - ⁶M. Wachowiak, X. Wang, A. Fenster, and T. Peters, "Compact support radial basis functions for soft tissue deformation," *Proc ISBI*, April 21–25, 2004 (IEEE, New York, 2004) pp. 1259–1262.
 - ⁷J. Moore, M. Drangova, M. Wierzbicki, and T. Peters, "A high-resolution averaged heart model based on averaged MRI data," *Medical Image Computing—Computer-assisted Intervention (MICCAI) Montreal*, LNCS 2878 (Springer-Verlag, Heidelberg, 2003), pp. 549–556.
 - ⁸M. Wachowiak and T. Peters, "Parallel optimization approaches for medical image registration," *Medical Image Computing—Computer-Assisted Intervention (MICCAI) St-Malo*, France, LNCS 3216 (Springer-Verlag, Heidelberg, 2004), pp. 781–788.
 - ⁹E. Coste-Manière, L. Adhami, F. Mourgues, and A. Carpentier, "Planning, simulation and augmented reality for robotic cardiac procedures: The STARS system of the chlR team," *Semin Thorac. Cardiovasc. Surg.* **15**, 141–156 (2003).
 - ¹⁰F. Mourgues, T. Vieville, V. Falk, and E. Coste-Manière, "Interactive guidance by image overlay in robot assisted coronary artery bypass," *Medical Image Computing—Computer-assisted Intervention (MICCAI) Montreal*, LNCS 2878 (Springer-Verlag, Heidelberg, 2003), pp. 173–181.
 - ¹¹E. Coste-Manière, L. Adhami, F. Mourgues, and O. Bantiche, "Optimal planning of robotically assisted heart surgery: Transfer precision in the operating room," *Int. J. Robot. Res.* **23**, 539–548 (2004).
 - ¹²X. Huang, N. A. Hill, and T. M. Peters, "Ultrasound-based technique for intra-thoracic surgical guidance," *Proceedings SPIE Medical Imaging Vol. 5744: Visualization, Image-Guided Procedures, and Display*, San Diego, CA, Feb. 13–15, 2005, edited by Robert L. Galloway, Jr. and Kevin R. Cleary (SPIE, Bellingham, WA, 2005), 822–830.
 - ¹³J. Feldmar, N. Ayache, and F. Betting, "3D-2D projective registration of free-form curves and surfaces," *Comput. Vis. Image Underst.* **65**, 403–424 (1997).
 - ¹⁴A. Guezec, P. Kazanzides, B. Williamson, and R. Taylor, "Anatomy-based registration of CT-scan and intraoperative x-ray images for guiding a surgical robot," *IEEE Trans. Med. Imaging* **17**, 715–728 (1998).
 - ¹⁵S. Lavallée and R. Szeliski, "Recovering the position and orientation of free-form objects from image contours using 3D distance maps," *IEEE Trans. Pattern Anal. Mach. Intell.* **17**, 378–390 (1995).
 - ¹⁶Y. Kita, D. Wilson, and J. Noble, "Real-time registration of 3D cerebral vessels to X-ray angiograms," *Medical Image Computing—Computer-assisted Intervention (MICCAI)*, Cambridge USA, LNCS 1496 (Springer-Verlag, Heidelberg, 1998), pp. 1125–1133.
 - ¹⁷L. Lemieux, R. Jagoe, D. Fish, N. Kitchen, and D. Thomas, "A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs," *Med. Phys.* **21**, 1749–1760 (1994).
 - ¹⁸L. Brown and T. Boulton, "Registration of planar film radiographs with computed tomography," *Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA)*, San Francisco, California, 42–51 July 1996 (IEEE, New York, 1996).
 - ¹⁹G. Penney, J. Weese, J. Little, P. Desmedt, D. Hill, and D. Hawkes, "A comparison of similarity measures for use in 2D-3D medical image registration," *IEEE Trans. Med. Imaging* **17**, 586–595 (1998).
 - ²⁰G. Penney, P. Batchelor, D. Hill, and D. Hawkes, "Validation of a two- to three-dimensional registration algorithm for aligning preoperative CT images and intraoperative fluoroscopy images," *Med. Phys.* **28**, 1024–1032 (2001).
 - ²¹J. Hipwell, G. Penney, R. McLaughlin, K. Rhode, P. Summers, T. Cox, J. Byrne, A. Noble, and D. Hawkes, "Intensity-based 2D-3D registration of cerebral angiograms," *IEEE Trans. Med. Imaging* **22**, 1417–1426 (2003).
 - ²²J. Byrne, C. Colominas, J. Hipwell, T. Cox, J. Noble, G. Penney, and D. Hawkes, "Assessment of a technique for 2D-3D registration of cerebral intra-arterial angiography," *Br. J. Radiol.* **77**, 123–128 (2004).
 - ²³R. McLaughlin, J. Hipwell, D. Hawkes, A. Noble, J. Byrne, and T. Cox, "A comparison of 2D-3D intensity-based registration and feature-based registration for neurointerventions," *Medical Image Computing—Computer-assisted Intervention (MICCAI) LNCS 2489* (Springer-Verlag, Heidelberg, 2002), pp. 517–524.
 - ²⁴D. Tomazevic, B. Likar, T. Slivnik, and F. Pernu, "3D-2D registration of CT and MR to X-ray images," *IEEE Trans. Med. Imaging* **22**, 1407–1416 (2003).
 - ²⁵H. Livyatan, Z. Yaniv, and L. Joskowicz, "Gradient-based 2-D/3-D rigid registration of fluoroscopic x-ray to CT," *IEEE Trans. Med. Imaging* **22**, 1395–1406 (2003).
 - ²⁶D. Russakoff, T. Rohlfing, R. Shahidi, D. Kim, J. Adler, and C. Maurer, "Intensity-based 2D-3D spine image registration incorporating one fiducial marker," *Medical Image Computing—Computer-Assisted Intervention (MICCAI) LNCS 2878* (Springer-Verlag, Heidelberg, 2003), pp. 287–294.
 - ²⁷M. Mahfouz, W. Hoff, R. Komistek, and D. Dennis, "A robust method for registration of three-dimensional knee implant models to two-dimensional fluoroscopy images," *IEEE Trans. Med. Imaging* **22**, 1561-287-2941574 (2003).
 - ²⁸M. Vermandel, N. Betrouni, G. Palos, J.-Y. Gauvrit, C. Vasseur, and J. Rousseau, "Registration, matching, and data fusion in 2D/3D medical imaging: Application to DSA and MRA," *Medical Image Computing—Computer-Assisted Intervention (MICCAI) LNCS 2878* (Springer-Verlag, Heidelberg, 2003), pp. 778–785.
 - ²⁹W. Lorensen and H. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," *Comput. Graph. (ACM)* **21**, 163–169 (1987).
 - ³⁰S. Chen and J. Carroll, "Kinematic and deformation analysis of 4-D coronary arterial trees reconstructed from cine angiograms," *IEEE Trans. Med. Imaging* **22**, 710–721 (2003).
 - ³¹R. Fahrig, M. Moreau, and D. Holdsworth, "Three-dimensional computed tomographic reconstruction using a C-arm mounted XR2I: Correction of image intensifier distortion," *Med. Phys.* **24**, 1097–1106 (1997).
 - ³²C. Maurer, J. Fitzpatrick, M. Wang, R. Galloway, R. Maciunas, and G. Allen, "Registration of head volume images using implantable fiducial markers," *IEEE Trans. Med. Imaging* **16**, 447–462 (1997).
 - ³³G. Lehmann, D. Habets, D. Holdsworth, T. Peters, and M. Drangova, "Simulation of intra-operative 3D coronary angiography for enhanced minimally invasive robot cardiac intervention," *Medical Image Computing—Computer-Assisted Intervention (MICCAI) Tokyo*, Japan, LNCS 2489 (Springer-Verlag, Heidelberg, 2002), pp. 268–275.
 - ³⁴S. Aylward and E. Bullitt, "Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction," *IEEE Trans. Med. Imaging* **21**, 61–75 (2002).
 - ³⁵R. Siddon, "Fast calculation of the exact radiological path for a three-dimensional ct array," *Med. Phys.* **12**, 252–255 (1985).
 - ³⁶C. Blondel, G. Malandain, R. Vaillant, F. Vernavay, E. Coste-Manière, and N. Ayache, "4D tomographic representation of coronary arteries from one rotational x-ray sequence," *Medical Image Computing—Computer-Assisted Intervention (MICCAI) Montreal LNCS 2878* (Springer-Verlag, Heidelberg, 2003), pp. 416–423.
 - ³⁷T. Koller, G. Gerig, G. Szekely, and D. Dettwiler, "Multiscale detection of curvilinear structures in 2-D and 3-D image data," *International Conference on Computer Vision (ICCV95)* (Academic Press, San Diego, CA, 1995), pp. 864–869.
 - ³⁸Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis, "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images," *Med. Image Anal.* **2**, 143–168 (1998).
 - ³⁹J. Astola and I. Virtanen, "Entropy correlation coefficient, a measure of statistical dependence on categorized data," *Proceedings of the University of Vassa, Finland, Discussion Papers 44* (University of Vassa, Vassa, Finland, 1982), pp. 1–12.
 - ⁴⁰J. Hajnal, D. Hill, and D. Hawkes, *Medical Image Registration*, The Bio-

- medical Engineering Series (CRC Press, Boca Raton, 2001).
- ⁴¹J. Nelder and R. Mead, "A simplex for function minimization," *Comput. J. (U.K.)* **7**, 308–313 (1965).
- ⁴²W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C*, 2nd ed. (Cambridge University Press, New York, 1993).
- ⁴³M. Luersen, R. Le Riche, and F. Guyon, "A constrained, globalized, and bounded Nelder-Mead method for engineering optimization," *Struct. Multidiscip. Optim.* **27**, 43–54 (2004).
- ⁴⁴W. Schroeder, K. Martin, and W. Lorensen, *The Visualization Toolkit*, 2nd ed. (Prentice-Hall, Upper Saddle River, New Jersey, 1998).
- ⁴⁵L. Ibanez, W. Schroeder, L. Ng, and J. Cates, *ITK Software Guide* (Kitware Inc., 2003).
- ⁴⁶G. Shechter, F. Devernay, E. Coste-Manière, A. Quyyumi, and E. McVeigh, "Three-dimensional motion tracking of coronary arteries in biplane cineangiograms," *IEEE Trans. Med. Imaging* **22**, 493–503 (April 2003).
- ⁴⁷F. Cheriet and J. Meunier, "Self-calibration of a biplane x-ray imaging system for an optimal three-dimensional reconstruction," *Comput. Med. Imaging Graph.* **23**, 133–141 (1999).
- ⁴⁸I. Liu and Y. Sun, "Fully automated reconstruction of three-dimensional vascular tree structures from two orthogonal views using computational algorithms and production rules," *Opt. Eng.* **31**, 2197–2207 (1992).
- ⁴⁹J. Klein, J. Hoff, J. Peifer, R. Folks, C. Cooke, I. S. King, and E. Garcia, "A quantitative evaluation of the three-dimensional reconstruction of patient's coronary arteries," *Int. J. Card. Imaging* **14**, 75–87 (1998).
- ⁵⁰C. Henri and T. Peters, "Three-dimensional reconstruction of vascular trees: experimental evaluation," *Med. Phys.* **23**, 617–627 (May 1996).
- ⁵¹J. Maddux, S. Chen, B. Groves, J. Messenger, O. Wink, and J. Carroll, "Rotational angiography and 3D coronary modeling: Revolutions in the cardiac cath lab," *Med. Mundi* **47**, 8–14 (2003).
- ⁵²C. Henri, *Three-Dimensional Reconstruction of Cerebral Vasculature*, Ph.D. thesis, McGill University (1993).
- ⁵³T. Kayikcioglu, A. Gangal, and M. Turhal, "Reconstructing coronary arterial segments from three projection boundaries," *Pattern Recogn. Lett.* **22**, 611–624 (2001).