

Nonrigid 2D/3D Registration of Coronary Artery Models With Live Fluoroscopy for Guidance of Cardiac Interventions

David Rivest-Hénault*, Hari Sundar, and Mohamed Cheriet

Abstract—A 2D/3D nonrigid registration method is proposed that brings a 3D centerline model of the coronary arteries into correspondence with bi-plane fluoroscopic angiograms. The registered model is overlaid on top of interventional angiograms to provide surgical assistance during image-guided chronic total occlusion procedures, thereby reducing the uncertainty inherent in 2D interventional images. The proposed methodology is divided into two parts: global structural alignment and local nonrigid registration. In both cases, vessel centerlines are automatically extracted from the 2D fluoroscopic images, and serve as the basis for the alignment and registration algorithms. In the first part, an energy minimization method is used to estimate a global affine transformation that aligns the centerline with the angiograms. The performance of nine general purpose optimizers has been assessed for this problem, and detailed results are presented. In the second part, a fully nonrigid registration method is proposed and used to compensate for any local shape discrepancy. This method is based on a variational framework, and uses a simultaneous matching and reconstruction process to compute a nonrigid registration. With a typical run time of less than 3 s, the algorithms are fast enough for interactive applications. Experiments on five different subjects are presented and show promising results.

Index Terms—Chronic total occlusions, computed tomography, image-guided interventions, interventional cardiology, X-ray fluoroscopy, 2D/3D Registration.

I. INTRODUCTION

X-RAY fluoroscopy is the modality of choice for the guidance of percutaneous coronary interventions (PCI) of chronic total occlusions (CTO). During these procedures, crossing CTOs using a guidewire is particularly hazardous, since the occlusion blocks the propagation of the contrast agent, and makes the occluded portion of the vessel invisible under fluoroscopy. In addition to the contrast issue, the projective nature of fluoroscopy results in ambiguities in the interpretation of 3D structures that further complicates the interventional procedure. Similar challenges arise in a vast array of minimally invasive procedures [28], [27].

Manuscript received January 30, 2012; revised March 28, 2012; accepted April 02, 2012. Date of publication April 17, 2012; date of current version July 27, 2012. Asterisk indicates corresponding author.

*D. Rivest-Hénault is with École de Technologie Supérieure, Montreal, QC, H3C 1K3 Canada, and also with Siemens Corporate Research, Princeton, NJ 08540 USA (e-mail: david.rivest-henault.1@ens.etsmtl.ca).

M. Cheriet is with École de Technologie Supérieure, Montreal, QC, H3C 1K3 Canada.

H. Sundar is with Siemens Corporate Research, Princeton, NJ 08540 USA.

Digital Object Identifier 10.1109/TMI.2012.2195009

Current preinterventional planning routinely includes the acquisition of a computed tomography angiography (CTA), or other 3D imaging modality, which is used to reduce visual uncertainty. The performance characteristics of CTA with blood pool contrast injection enables the calcifications causing the CTOs to be clearly distinguished, as opposed to what can be perceived with interventional X-ray fluoroscopy with direct contrast agent injection. This makes CTA a highly valuable tool at the planning stage, but establishing correspondence between these data and the interventional images can prove difficult.

To address this issue, a 3D model can be extracted preprocedure from the acquired volume, aligned with the 2D fluoroscopic views, and overlaid on top of the live images, thereby augmenting the interventional images. However, achieving this alignment is a challenge in itself, mainly because finding intermodal correspondence is a nontrivial problem, and also because of the nonlinear aspect of the underlying optimization problem. Furthermore, a simple rigid transformation might not be sufficient to provide a satisfying 2D/3D registration. Since the 3D planning image is acquired under a breathhold, there are significant shape changes as compared to the intraoperative images acquired under free breathing. This makes it extremely important to use a nonrigid registration method while aligning the preoperative model with the live intraoperative images. This is, however, a difficult ill-posed inverse problem. Nevertheless, 2D/3D registration methods have the potential to greatly reduce the uncertainty relative to interventional X-ray angiography, and to do so with only minimal modification to the existing clinical flows. The development of such methods is the main objective of this paper, and its most significant contribution.

From a broader perspective, 2D/3D registration methods have numerous applications in fields such as neurology [16], orthopaedics [2], and cardiology [35]. The associated body of literature is expanding rapidly, as is apparent in the thorough review of techniques recently published by Markelj *et al.* [22]. Below, the techniques most closely related to the one presented in this paper are discussed.

The maximal precision that can potentially be achieved by a registration process is directly linked to the complexity of the transformation model involved. Many 2D/3D registration approaches consider only a rigid transformations model [11], [45], [43], [35], or a slightly more flexible affine model [44]. In general, this is appropriate for rigid structures, such as bones, or to provide an initial alignment of the modalities, but might prove insufficient to account for the shape changes of flexible structures. Consequently, nonrigid deformation models have been

proposed for 2D/3D registration [13], [23], [21]. In [13], the transformation model is strongly constrained by a length conservation term. The method has mostly been demonstrated on synthetic examples, and its computational complexity, resulting in a computational time of around 5 min, makes it impractical in an interventional setting. The approach by Metz *et al.* [23] requires the acquisition of 4D CTA, that allows for the use of a $3D + t$ model for cardiac motion. However, it does not account for shape changes occurring through respiratory movement, and the acquisition of a 4D CTA is not possible in most interventions, which is a major shortcoming of this approach. Interesting results are presented in the paper by Liao *et al.* [21], but their technique makes use of features specific to the abdominal aorta.

In a clinical setting, one or two fluoroscopic planes are used for intervention guidance. Some authors [7], [13], [9] have investigated the monoplane scenario, but the reported errors are large in the out-of-plane direction. This lack of accuracy seriously limits confidence in the registration process. Biplane acquisition greatly reduces the ambiguity associated with these interventions, which is why this paper focuses on the alignment and registration of a 3D model with two fluoroscopic images. In our experience of five vastly different clinical sites located in Canada, Germany, and the Netherlands, we have observed that biplane acquisitions are now performed on a daily basis for complex cardiac interventions. Thus, the proposed methodology is practical in many clinical situations.

The method presented in this paper is designed to nonrigidly align a preoperative 3D centerline model of the coronary arteries with two intraoperative fluoroscopic images to visually augment the interventional images. It is composed of two steps: 1) a global transformation model is calculated to provide an initial rigid or affine alignment, and 2) a fully nonrigid model is used to compute the final registration. To estimate the global alignment parameters, a formulation derived from [42] is used. This formulation benefits from distance maps [10], [25] to measure the discrepancies between the projections of the 3D model and features automatically extracted from the 2D images. Such formulations have less computational complexity than those based on point matching approaches [40], [43] or on digitally reconstructed radiographs [33], [39]. A contribution of this paper is to generalize the cost function defined in [42] to cover the affine and biplane case.

The optimization surface in 2D/3D registration problems can be highly nonlinear for several reasons: the discretization of the image, the complexity of the structures to be registered, and the use of a rigid or affine transformation model. As a result, the minimization of the cost function is challenging. As we know that no single optimization method outperforms all the others, it is hard to understand why only one or two optimization algorithms have been evaluated in most related works [45], [42], [35]. In addition, work published by Lau and Chung [20] suggests that global optimizers might perform better than the more popular local optimizers for a related 2D/3D registration problem. An important contribution of this work is thus to present a rigorous comparison between seven local and two global optimizers. The data were gathered using one realistic simulated case and five clinical cases; 2D and 3D errors are re-

ported as well as runtime measurements. The results presented here can help implementers choose the best algorithm for their application. A similar study, but for intensity-based 2D/3D registration, is presented in [5].

Work recently published by Ruijters *et al.* [35] shares some similarities with the proposed 2D/3D alignment method. Their registration method makes use of the distance transform of the projection of a 3D centerline and of the output of the Frangi vesselness filter [12] to compute the cost associated with a certain pose. The optimization is carried out with either a Powell optimizer or a stochastic optimizer. There are, however, a few important differences between the approaches. While they specifically avoid 2D segmentation, we propose using a very recent 2D segmentation algorithm [37]. By doing so, it is possible to precompute the distance transform on the 2D images, instead of computing the distance transform of the projection of the 3D centerline for a certain pose at each iteration. This can lead to an improvement of up to two orders of magnitude in registration time. Nonetheless we acknowledge that computing the automatic 2D segmentation incurs some small overhead. We also make use of the affine transformation model in addition to the rigid one, and present a more thorough evaluation using additional optimizers. Finally, [35] does not consider nonrigid registration.

The 2D/3D registration method can also be used to capture the motion of the structure of interest across a sequence of frames [3], [44]. The demonstration of the suitability of the proposed global registration method in this multiframe setting is another contribution of this work. With respect to the nonrigid registration method, our main contribution is the formulation of this problem as a simultaneous matching, reconstruction and registration problem, that can, on modern hardware, be solved fast enough to be used intraoperatively during CTO procedures.

The rest of this paper is organized as follows. Section II presents the practical system considered here and other background information. The global alignment method is described in Section III. The proposed nonrigid registration method is introduced in Section IV. Experiments with both the global alignment method and the nonrigid registration method are presented in Section V. Finally, a discussion and the conclusion are presented in Section VII.

II. BACKGROUND INFORMATION

A 3D centerline representation of a coronary artery tree, segmented [14] from a preoperative CTA volume, is to be nonrigidly registered to two simultaneous fluoroscopic images. Starting from the default 3D location, computed from the calibration of the apparatus, the 3D alignment of the centerline is progressively refined using translation-only motions, a rigid transformation, and an affine transformation. Lastly, a nonrigid transformation is computed and provides the final registration.

The geometry of the system under consideration is represented schematically in Fig. 1 and is described using five coordinate systems (CS). The 3D centerline representing the coronary tree is described with respect to U , and a CS C is centered at the reference point of the 3D imaging device. The 3D centerline itself is described by using a set of S segments composed of Q_s control points $\mathbf{x}_{s,q} \in \mathbb{R}^3$ forming an undirected




Fig. 1. Geometry of the imaging system. Labeled frames represent coordinate systems. Dashed lines and bold symbols represent transformation matrixes.

acyclic graph. In term of 2D angiography, X is the reference CS of the biplane C-arm, and CSs P_1 and P_2 are centered on the 2D imaging planes. The transformations between C and U , and between X and $\{P_1, P_2\}$, are known from the calibration of the apparatus, and are encoded in the rigid transformation matrixes \mathbf{C} , \mathbf{P}_1 and \mathbf{P}_2 , respectively. The projective geometry of the two X-ray planes is also assumed to be available, and is encoded in the projection operators Ψ_1 and Ψ_2 that map a 3D point in homogeneous coordinates $\mathbf{x} = [x, y, z, 1]^T$ to a 2D point $\mathbf{x}' = [x, y]$ on the corresponding fluoroscopic plane.

The vessels are automatically segmented from the input fluoroscopic images, using the method proposed in [37], and represented as a binarization $B_n(i, j)$ of the input image. $B_n(i, j) = 1$, if the input image pixel at location $(i, j) \in \{0, \dots, W - 1\} \times \{0, \dots, H - 1\}$ corresponds to some structure of interest, and $B_n(i, j) = 0$ otherwise. It is assumed that the two fluoroscopic images are the same size $W \times H \in \mathbb{N}^2$. For the sake of completeness, the automatic segmentation method in [37] is briefly outlined here: 1) an input fluoroscopic image is processed using the Hessian-based Sato *vesselness* filter [36]; 2) seed points are generated by sparsely sampling a thresholding of the *vesselness* measure; 3) starting from those points, fibers are generated by integrating along the first eigenvector of the Hessian matrix computed at each pixel location; and 4) the fiber bundles are iteratively thinned to extract the centerlines of the vessels. The default parameters of this method have been used without modification throughout this paper. Examples of outputs are presented in Fig. 2.¹

III. TRANSLATIONAL, RIGID, AND AFFINE ALIGNMENT

In a typical clinical setting, the transformation that aligns the 3D coronary centerlines from the CTA acquisition to the images of the same structures on the biplane fluoroscopy can only be measured approximately, since the position of the patient cannot be controlled with a high degree of accuracy. A transformation \mathbf{T} , that is a mapping between the two imaging modality coordinate systems, can be estimated by a calibration of the apparatus, and refined by taking into account the geometry of the structures of interest. Starting from the initial conditions, the

improved alignment transformation is computed using a minimization process to estimate the parameters of a global translational or rigid transformation model. From that point, an affine transformation model can also be used to deform the 3D centerline in order to compensate for the shape discrepancy between the CTA and the biplane X-rays. A technique, inspired by [42], but extended to cover the biplane case and the use of affine transformations, is proposed for this purpose.

Formally, the total distance between the projection of the 3D centerline and the vessel centerlines segmented on each fluoroscopic plane is to be minimized. Let $\chi(\mathbf{x}) = \mathbf{T} \cdot \mathbf{C} \cdot \mathbf{x}$ be an affine transformation operator that maps a point $\mathbf{x} = [x, y, z, 1]^T$, relative to U , to the X coordinate system. Also, let $\Pi_n(\mathbf{x}) = \Psi_n \cdot \mathbf{P}_n \cdot \mathbf{x}$ be a projection that maps a point from X to the fluoroscopic plane P_n . We define the energy

$$E_{\text{Global}}(\chi) = \sum_{n=1}^2 \sum_{s,q} D_{B_n}(\Pi_n(\chi(\mathbf{x}_{s,q}))) \quad (1)$$

as the quantity that needs to be minimized. Here, $D_{B_n}(x)$ is the distance between a 2D point and the closest point where $B_n(i, j) = 1$. The distance $D_{B_n}(x)$ for each image position is computed beforehand, by creating the distance transform of each $B_n(i, j)$ segmentation, as presented in Fig. 2.

Optimizing the 12 values of \mathbf{T} directly is both ineffective and inefficient, and so a parametric approach has been used. The translation-only, rigid, or affine transformations \mathbf{T}_T , \mathbf{T}_R , and \mathbf{T}_A are represented using the parameter sets $\mathcal{T} \in \mathbb{R}^3$, $\mathcal{R} \in \mathbb{R}^6$, $\mathcal{A} \in \mathbb{R}^{12}$, respectively. The mappings from the parametric to the matrix representations of the \mathbf{T} s are defined as follows [38]:

$$\begin{aligned} \mathbf{T}_T(\mathcal{T}) &= M_T(\mathcal{T}_1^3) \\ \mathbf{T}_R(\mathcal{R}) &= M_T(\mathcal{R}_1^3) \cdot M_R(\mathcal{R}_4^6) \end{aligned}$$

and

$$\begin{aligned} \mathbf{T}_A(\mathcal{A}) &= M_T(\mathcal{A}_1^3) \cdot M_R(\mathcal{A}_4^6) \\ &\quad \cdot M_R(\mathcal{A}_{10}^{12}) \cdot M_S(\mathcal{A}_7^9) \cdot (M_R(\mathcal{A}_{10}^{12}))^T \end{aligned} \quad (2)$$

where \mathcal{S}_i^j is the subset containing the elements with indexes $\{i, \dots, j\}$ of \mathcal{S} and

$$M_T(\mathcal{T}) = \begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad M_S(\mathcal{S}) = \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

and

$$M_R(\theta) = \begin{bmatrix} c_y c_z, & c_z s_x s_y - c_x s_z, & s_x s_z + c_x c_z s_y, & 0 \\ c_y s_z, & s_x s_y s_z + c_x c_z, & c_x s_y s_z - c_z s_x, & 0 \\ -s_y, & c_y s_x, & c_x c_y, & 0 \\ 0, & 0, & 0, & 1 \end{bmatrix} \quad (3)$$

where $c_x = \cos(\theta_x)$ and $s_x = \sin(\theta_x)$. The transformation matrixes \mathbf{T}_T , \mathbf{T}_R , and \mathbf{T}_A are recentered in such a way that any change in the rotation parameters induces a motion that appears to occur around the origin of the U coordinate system, as opposed to the origin of C . This operation does not influence the

¹All 3D centerline projections presented in this work are color-coded as follows: if the 3D centerline is locally parallel to the view plane, it is green; if it is perpendicular, it is red. A linear interpolation is used for intermediary situations.

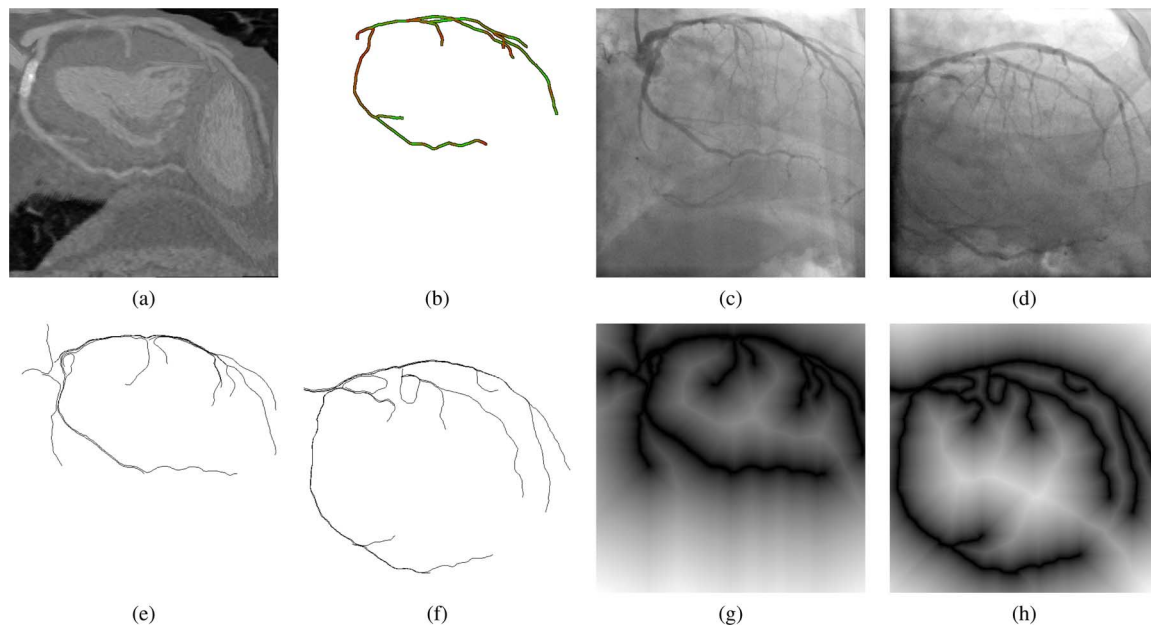


Fig. 2. Input data: (a) curved slice from a CT volume; (b) 3D centerlines; (c) and (d) fluoroscopic angiograms from two different points of view; (e) and (f) segmentation of the fluoroscopic images; and (g) and (h) distance transforms of the segmentations.

energy function (1), but is usually beneficial for the optimizers. Intuitively, this is because such recentering reduces the distance, in parameter space, associated with a rotation or a scaling operation that appears to occur at the centroid of the CT volume. The final definition of \mathbf{T} is then

$$\mathbf{T} = \mathbf{C} \cdot \mathbf{T}_J \cdot \mathbf{C}^{-1} \quad (4)$$

where \mathbf{C} is the transformation from U to C , and $J \in \{T, R, A\}$, depending on the transformation model chosen.

Any nonderivative general purpose optimization algorithm can be used to minimize (1). In this work, seven local (L) and two global (G) algorithms have been tested for this purpose. These algorithms are briefly described here.

Best Neighbor (L): At each iteration, a step in all principal directions of the parameter space is considered. The best move is applied, and a new iteration begins. When no move improves the solution, the step is halved.

Nelder-Mead² (L): A classic numerical optimization method that minimizes an m -dimensional function by evaluating the function value at the $m + 1$ vertices of a general simplex (or polytope). At each iteration, the vertex with the worst value is replaced by another one using a reflection operation followed by either an expansion or a contraction operation [24].

Sbplx³ (L): This method decomposes the problem in a low-dimensional subspace, and uses the Nelder-Mead algorithm to perform the search [34].

Cobyla (L): The *Constrained Optimization BY Linear Approximation* method works by constructing linear-approximations of the cost function and constraints using $m + 1$ vertex simplexes, and minimizes this approximation. The radius of the

simplex is progressively reduced, while maintaining a regular shape [30], [31].

Bobyqa (L): Each iteration of the *Bound Optimization BY Quadratic Approximation* method use a quadratic approximation of the cost function F constructed, typically by considering $2m + 1$ interpolation point. The trust region is progressively reduced until there is no further improvement [32].

Powell-Brent (L): At each iteration, a succession of exact 1D line optimizations is performed using Brent's method. The solution is updated using Powell's method of conjugate search directions [29], [8].

Praxis (L): Brent's *PRincipal-AXIS* method is a refinement of Powell's method of conjugate search directions [8].

Differential Evolution (G): This is a population-based stochastic global optimization method, the main feature of which is the use of the vector of the difference between pairs of individuals as the basis for the population evolution [41].

Direct (G): This global optimization algorithm is designed for problems with finite bound constraints, as is the case here. The parameter space is systematically and deterministically searched by dividing it into smaller and smaller hyperrectangles [18].

The free parameters of each algorithm have been adjusted in accordance with the recommendations of their original author or implementer. The implementations of the following algorithms are taken from the NLOpt library [17]: *Praxis*, *SBPLX*, *Cobyla*, *Bobyqa*, and *Direct*. The performance of these optimizers is discussed in detail in Section V. In order not to change the nature of the problem, generous bounds have been fed to the optimizers: ± 200 mm for the translation parameters, ± 45 deg for the rotation parameters, $[0.5, 1.5]$ for the scale parameters, and ± 360 deg for the scale-rotation parameters $[\mathcal{A}_{10}^{12}]$ in (2).

The same set of stopping criteria was used for all the optimizers, except *Differential Evolution*. Specifically, the optimizers were allowed to run until, after one optimization step:

²Also known as *Downhill Simplex*.

³The name of the original author's implementation is *Subplex*, the name of the NLOpt [17] implementation is *Sbplx*.

1) the value of the cost function was reduced by less than $1e^{-14}$, or 2) the value of the optimized parameter changed less than $1e^{-4}$ mm for the translation parameters, $1e^{-5}$ rad for the rotation parameters, and $1e^{-7}$ for the scaling parameters. The parameter tolerances have an intuitive interpretation. For example, a difference in translation of $1e^{-4}$ mm corresponds to a maximal displacement of 1/2000 of a pixel on an image plane. As to the *Differential Evolution* optimizer, since the convergence of stochastic algorithms is not regular, we decided to let the optimizer run until it had evaluated the cost function 100 000 times⁴, that corresponds to a median run of approximately 10 s.

Starting from the default patient position, the alignment is progressively refined by considering 1) translations only, 2) rigid transformations, and 3) affine transformations.

A. Multi Frame Alignment

The global alignment method can be expanded to cover the multi-frame scenario. In this case, the crucial hypothesis is that because of the temporal continuity, the estimated alignment transformation would only change little from one frame to the other. The global energy function can then be redefined for a sequence of F biplane frames

$$E_{\text{Multi}}(\chi_f) = \sum_{f=1}^F \sum_{n=1}^2 \sum_{s,q} D_n(\Pi_n(\chi_f(\mathbf{x}_{s,q}))) + \gamma \sum_{f=1}^{F-1} D(\chi_f, \chi_{f+1}) \quad (5)$$

where γ is a free parameter, and χ_f is defined in a way similar to χ , but with a different transformation matrix \mathbf{T}_f for each time point. The most delicate part of this energy function is the definition of the inter-frame distance $D(\chi_f, \chi_{f+1})$ since it is well known that the space of rigid, and affine, transformation matrixes is not linear. The rigid transformation space can be locally linearized for small changes in orientation [4], [26], as is assumed to be the case here. Using this framework, a rigid transformation is represented by a translation vector \vec{t} and a rotation vector \vec{r} , defined as the product of a unit vector \vec{n} and a rotation angle θ , such that $\vec{r} = \theta\vec{n}$. The translation vector simply corresponds to the translational part of the transformation matrix, and the mapping between the rotation matrix R and the rotation vector \vec{r} is given by [4]

$$\theta = \arccos\left(\frac{\text{Trace}(R) - 1}{2}\right)$$

and

$$\begin{pmatrix} 0 & -n_z & n_y \\ n_z & 0 & -n_x \\ -n_y & n_x & 0 \end{pmatrix} = \frac{R - R^T}{2 \sin(\theta)}.$$

Using this representation, the inter-frame distance is defined as follows [4]:

$$D(\chi_f, \chi_{f+1}) = N_\alpha(\mathbf{T}_{f+1}^{-1} \cdot \mathbf{T}_f) \quad (6)$$

⁴1 000 000 times, in the case of simulation.

with

$$N_\alpha(R) = \|\vec{r}\|^2 + \|\alpha\vec{t}\|^2. \quad (7)$$

Here, the transformations \mathbf{T}_f have the same meaning as the transformations \mathbf{T} defined in the previous section. Also, α is a real number that balances the respective contributions of the rotation and translation parts. It was set to 0.05, in accordance with the recommendation in [4].

Defining a linear distance between affine transformation matrixes is more involved and requires complex computation [1]. Since it is assumed in this work that rigid transformations capture the major part of the transformation, we decided to regularize only the rigid part of the transformation and leave the remaining affine parameters unconstrained. Formally, in the case of affine transformations, the transformations \mathbf{T}_f in the regularizer definition are composed using the parameters \mathcal{A}_1^6 only, whereas the full set of 12 parameters is used to compose the transformations χ_f in the energy definition.

This energy function can be useful for tracking the alignment transformation over a sequence of frames, or, if $\gamma \rightarrow \infty$, for computing an average transformation using, for example, three adjacent frames. As with the single frame energy, the transformations χ_f are progressively refined by sequentially increasing the complexity of the transformation models.

IV. NONRIGID REGISTRATION

Global affine transformation models cannot entirely compensate for the inter-modal shape discrepancy caused by breathing and by the beating of the heart. A nonrigid registration method is now introduced, that can greatly improve the visual correspondence between the 3D coronary tree centerlines and the two calibrated fluoroscopic images. The aim of this method is to be automatic, in the sense that it does not require the user to identify correspondences.

The nonrigid transformation is represented as a set \mathbf{r} of 3D translation vectors $\mathbf{r}_{s,q}$ that are to be applied to the corresponding centerline points $\mathbf{x}_{s,q}$ in CS U . Thus, a registered point $\bar{\mathbf{x}}_{s,p}$ is computed using $\bar{\mathbf{x}}_{s,q} = \mathbf{x}_{s,q} + \mathbf{r}_{s,q}$. An energy function is defined to measure the quality of a solution

$$E_{\text{NR}}(\mathbf{r}) = E_{\text{Image}}(\mathbf{r}) + E_{\text{Internal}}(\mathbf{r}). \quad (8)$$

An iterative minimization process is used to minimize (8), and, starting from the position $\mathbf{r}_{s,q} = [0, 0, 0]^T \forall \{s, q\}$, the transformations $\mathbf{r}_{s,q}$ are progressively refined by following a gradient descent approach. This results in a smooth motion of the registered centerline, as can be seen in Fig. 3.

A. Image Energy

A matching and reconstruction process is used to calculate a reconstructed point $\hat{\mathbf{x}}_{s,q}$ for every point $\mathbf{x}_{s,q}$ in the 3D centerline model. The matching is performed on each fluoroscopic plane separately. Starting from the projection \mathbf{x}' of $\bar{\mathbf{x}}_{s,q}$ on an image plane n , a search on the corresponding segmentation image B_n is performed in the directions that are perpendicular to the projection of the centerline—see Fig. 4(a) for an illustration. This direction is computed, at every location, based on the partial




Fig. 3. Progression of a nonrigid registration. From left to right: sample input image, and position of the centerline after 0, 50, 350, and 1200 iterations. This corresponds to subject 4 in Table III. See also the related video file, as described in Section VI.




Fig. 4. (a) Matching of a point from the centerline x' with a point from the 2D structure. Points x''_a and x''_b are identified as potential candidates. (b) Creation of three myocardium constraints (k_1 , k_2 , and k_3) at a junction.

derivatives $((\partial x')/(\partial x), (\partial x')/(\partial y))$. As a result, at most two matching points, $\mathbf{x}''_{n,a}$ and $\mathbf{x}''_{n,b}$, are identified on each fluoroscopic plane $n \in \{1, 2\}$. The pair $(\mathbf{x}''_{1,i}, \mathbf{x}''_{2,i})$, $i \in a, b$ that best satisfies the epipolar constraint is kept. For this purpose, the following quantity is considered:

$$d(\mathbf{x}''_{1,i}, \mathbf{x}''_{2,i}) = d(\mathbf{x}''_{1,i}, e_{2,i}) + d(\mathbf{x}''_{2,i}, e_{1,i}) \quad (9)$$

where $d(\mathbf{x}''_{n,i}, e_{m,i})$ is the Euclidean distance between point $\mathbf{x}''_{n,i}$ and the epipolar line induced by $\mathbf{x}''_{m,i}$, $\{n, m\} \in \{1, 2\}$, with $n \neq m$. The pair minimizing (9) is then used to calculate a reconstructed 3D point $\hat{\mathbf{x}}_{s,p}$ using the projective geometry of the system [15]. If only one point $\mathbf{x}''_{n,i}$ is found on an image plane, it is retained by default.

Searching for matching points in directions that are locally perpendicular to the projection of the centerline is an important aspect of this process. Since it is known *a priori* that the global alignment of the structures is correct, the perpendicular search direction prevents many unrealistic matches, compared to a nearest-point approach. This is especially important when the 2D automatic segmentation is not perfect. For example, when the 3D projection results in a curve that is longer than the 2D curve that has been automatically segmented, taking the nearest point would result in a significant shortening of the curve from its tip. In contrast, considering the perpendicular direction considerably reduces this effect, because no match would be found at the tip of the 3D centerline projection.

The quality of the reconstruction cannot be guaranteed and an indicator function $H(\mathbf{x}_{s,p}) \mapsto \{0, 1\}$ is used to reject outliers. $H(\mathbf{x}_{s,p}) = 0$ if: 1) no matching point is found in any image planes; 2) the 2D distance between \mathbf{x}' and \mathbf{x}''_n is greater than Max_{2D} ; and 3) the 3D distance between $\mathbf{x}_{s,p}$ and $\hat{\mathbf{x}}_{s,p}$ is greater than Max_{3D} . Otherwise, $H(\mathbf{x}_{s,p}) = 1$. Here, Max_{2D} and Max_{3D} are thresholding functions. More complex kernels, such as Huber functions, could be integrated in the proposed

scheme, but this possibility has not been explored. Their values were fixed to $\text{Max}_{2D} = 50$ pixels and $\text{Max}_{3D} = 6$ mm.

The image energy in (8) is defined as the sum of the squared error between the reconstructed points $\hat{\mathbf{x}}_{s,p}$ and the centerline point, for all valid points, as follows:

$$\begin{aligned} E_{\text{Image}}(\mathbf{r}) &= \sum_{s,q} H(\mathbf{x}_{s,q}) (\hat{\mathbf{x}}_{s,q} - (\mathbf{x}_{s,q} + \mathbf{r}_{s,q}))^2 \\ &= \sum_{s,q} H(\mathbf{x}_{s,q}) (\hat{\mathbf{x}}_{s,q} - \bar{\mathbf{x}}_{s,q})^2. \end{aligned} \quad (10)$$

This energy is at its minimum when all the registered centerline points are at the same position as the reconstructed points.

B. Internal Energy

Regularizers are used to keep the registered centerline visually coherent and geometrically plausible. The following three internal energy terms are considered for this purpose:

$$\begin{aligned} E_{\text{Internal}} &= E_{\text{Disp}}(\mathbf{r}) + E_{\text{Smooth}}(\mathbf{r}) \\ &\quad + E_{\text{Myocard}}(\mathbf{r}) \\ &= \mu \sum_{s,q} |\mathbf{r}_{s,q}|^2 + \nu \sum_{s,q} |\dot{\mathbf{r}}_{s,q}|^2 \\ &\quad + \lambda \sum_{\{i,j\} \in \mathcal{K}} \left(\frac{|\mathbf{r}_i - \mathbf{r}_j|}{|\mathbf{x}_i - \mathbf{x}_j|} \right)^2 \end{aligned} \quad (11)$$

where μ , ν , and λ are energy-balancing free parameters. The first term, $E_{\text{Disp}}(\mathbf{r})$, is minimal when the displacement owing to the nonrigid transformations is small. This regularizer is necessary to minimize 3D out-of-plane motions that are only poorly constrained by a pair of 2D angiograms. The second, $E_{\text{Smooth}}(\mathbf{r})$, is used to ensure smoothness over each vessel segment. In this section, $\dot{\mathbf{r}}_{s,q}$ and $\ddot{\mathbf{r}}_{s,q}$ represent, respectively, the first and second derivatives of the translation vector \mathbf{r} for the position (s, q) with respect to the segment parameter. At segment junctions, Neumann boundary conditions, $\dot{\mathbf{r}}_{s,q} = [0, 0, 0]^T$, are assumed.

With only the first two internal energy terms, the centerline model is flexible, and the motions of the vessels are independent of each other, except at junction nodes. In reality, the vessel motions are mechanically constrained not only by their own specific rigidity, but also because they are attached to the myocardium. In this respect, $E_{\text{Myocard}}(\mathbf{r})$ is intended to act as a minimalist model of the myocardium constraint, and is used to ensure a certain degree of rigidity at the vessel branches as it prevents small segments from collapsing onto bigger ones. This




Fig. 5. Total displacement of the centerline as a function of the number of iterations—see Section IV-C for details.

constraint is modeled by using artificial links around the junction of three segments, as depicted in Fig. 4(b). Let D_s be a distance parameter, with its value set to three times the average diameter of the vessels present at the junction. The node on each segments that is at a distance D_s from the junction is joined to the corresponding node of the other segments. In a way that is analogous to a mechanical spring, with each link generating a force that is proportional to its displacement ratio if it is either expanded or compressed during the registration process. In (11), \mathcal{K} is defined as the set of all pairs of node $\{i, j\}$ that have been linked together.

C. Energy Minimization

The complete energy function is defined using (8), (10), and (11), and it is minimized by computing its Euler–Lagrange equations and following a gradient descent approach. Starting from the point where $\mathbf{r}_{s,p} = [0, 0, 0]^T | \nabla \{s, p\}$ at iteration $z = 0$, the resulting discretized update equation is

$$\mathbf{r}_{s,q}^{z+1} = \mathbf{r}_{s,q}^z + \epsilon \left[H(x_{s,q}) (\hat{\mathbf{x}}_{s,q}^z - \bar{\mathbf{x}}_{s,q}^z) - \mu \mathbf{r}_{s,q}^z + \nu \ddot{\mathbf{r}}_{s,q}^z + \lambda \sum_{\{i\} \in \mathcal{K}_{s,q}} (\mathbf{r}_i^z - \mathbf{r}_{s,q}^z) \right] \quad (12)$$

where the numerical step ϵ is a small constant and $\mathcal{K}_{s,q} \subset \mathcal{K}$ is the possibly empty set of constraints acting on node (s, q) . Equation (12) is evaluated iteratively until the total displacement stagnates. More formally, the total displacement is denoted $\Delta^z \equiv \sum_{s,q} |\mathbf{r}_{s,q}^z|$, and the stopping criterion is defined, for $z > \beta$, as $\Delta^z - \Delta^{z-\beta} < \zeta \Delta^\beta$. An example of the progression of Δ^z is shown in Fig. 5. As we can deduce from the shape of the curve, the final result is not very sensitive to the specific values of β and ζ . In practice, we used $\beta = 100$ and $\zeta = 0.001$.

D. Parameter Selection

The three parameters μ , ν , and λ govern the behavior of (12) and must be selected for a specific task. Starting from a certain parameter set ($\mu = 0.10, \nu = 10.0, \lambda = 1.0$), some insight about the behavior of method can be acquired by varying each value individually, and by computing the mean 2D projection




Fig. 6. Illustration of the behavior of parameter λ in the nonrigid registration method: (left) $\lambda = 0$, (right) $\lambda = 0.3$.

error and the 3D error (described in Section V). This methodology has been applied to the patient 4 of the clinical dataset, following an affine alignment, and also to a dataset with synthetic deformations (see Section V-A). The former depicts a RCA, and the latter, a LCA. The results are presented in Fig. 7.

Parameter μ constrains the overall displacement, and the 2D error curves suggest that the computed error increases with its value. However, the 3D error curve clearly indicates a minimum when μ has a value in the 0.05–0.10 range. In addition, we found that the convergence rate increases with μ . For example, the registration used to produce the graph in Fig. 7 (top left) converged in 4880, 3320, and 1100 iterations with $\mu = \{0.001, 0.01, 0.1\}$, respectively. Parameter ν controls the rigidity of the model. In this case, the 2D projection error suggests using a value in the 0.50–2.00 range. However, the 3D error curve shows that higher values might result in more accurate registration. We thus choose to use a values of $\nu = 10.0$, which appears to be a good trade-off between the 2D and 3D error curves. Parameter λ balances the myocardium constraint, and helps to preserve the general integrity of the global shape. With the clinical RCA dataset, it was found that high value can overly constrain the motion, which results in increased errors. However, with the simulated LCA dataset, the situation appears almost inversed, as λ values of up to 10.0 are clearly beneficial. We believe that this situation can be explained by the differences in the nature of the deformations that affects the LCA and RCA datasets. A value $\lambda = 0.5$ was selected because it represents a good trade-off in the error measures, and because it improves the visual appearance in many cases, as depicted in Fig. 6. It might be better to use a different value for LCA and RCA datasets, but this idea was not pursued here. Finally, the parameters were thus fixed to the following values in all subsequent experiments: $\mu = 0.05$, $\nu = 10.0$, and $\lambda = 0.5$, unless stated otherwise.

V. EXPERIMENTAL RESULTS

The algorithms were implemented in C++ within a proprietary prototyping environment that allows us to show an overlay of the 3D centerline model on the live fluoroscopy images during interventions. All global optimizers are single threaded. The nonrigid registration method is partly multi-threaded using OpenMP. Run time given are for an i7 Q820 Quad Core Intel CPU. Presented run times do not account for 2D image segmentation. On average, this operation has a run time 0.200 s per image, with a standard deviation of 0.088 s. Experiments on both simulated and clinical data are presented.

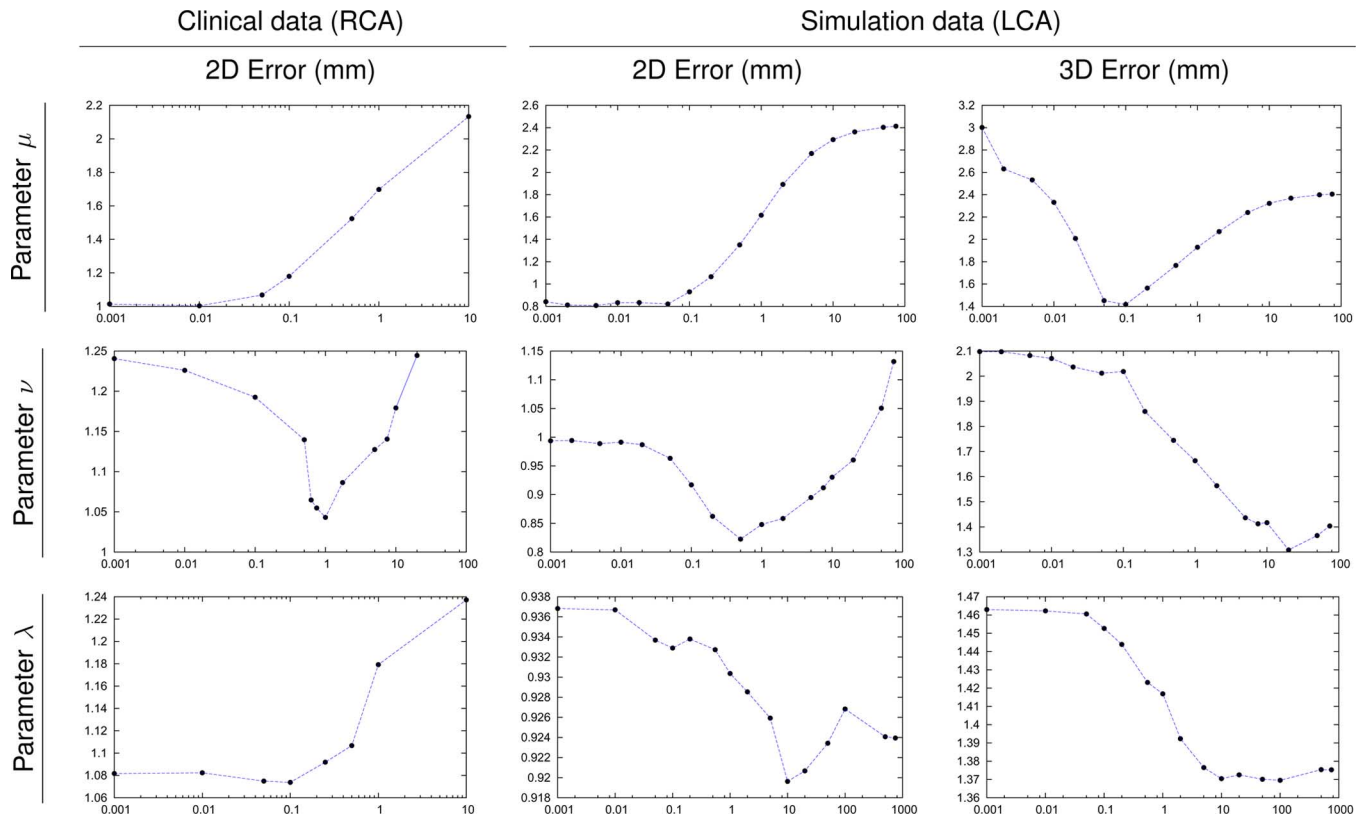


Fig. 7. Mean 2D and 3D error after nonrigid registration. For each graph, one of the parameters was varied, while the others two were kept fixed. The default parameter values are: $\mu = 0.1$, $\nu = 10.0$, and $\lambda = 1.0$. The x-axis scales are logarithmic.

A. Simulations

The proposed alignment and registration methods have been tested on a set of simulations. This has allowed us to characterize the performance of the various algorithms in a somewhat ideal scenario, and also to evaluate the 3D registration error, which is not possible using clinical data. The centerline was registered with a pair of digitally reconstructed radiographs (DRR), of size 512×512 pixels with a resolution of 0.345 pixel/mm, that were generated as follows: 1) the coronary arteries segmented in the CT volume were projected onto the simulated fluoroscopic planes using the technique presented in [19], and 2) the inverse of the resulting projections was subtracted from the reference background images. This process is illustrated in Fig. 8, and is similar to that proposed in [45]. The simulations were used in the three scenarios described below. Since all the transformations are known during the simulation, all the 3D errors presented in this section are exact point-to-point errors. It should be noted that this style of error will penalize the compression or expansion of a segment, even when this produces no visual effect. It is thus more strict than most TRE formulations. The 2D error is computed as described in Section V-B, using the projection of the true centerline as the reference segmentation.

1) *Dependence on the Initial Solution*: The centerline is registered with a pair of DRRs, starting from different initial positions. This serves to quantify the sensitivity of the global alignment methods to a perturbation of the initial position, and equally, to a miscalibration of the \mathbf{C} , \mathbf{P}_1 , \mathbf{P}_2 , or \mathbf{T} rigid transformation matrixes. The selected initial positions are as follows. Let \mathbf{T}_0 be the transformation given by the calibration of the




Fig. 8. Creation of the DRRs: The coronaries segmented in the CT volume are projected onto a simulated radiographic plane (a); this image is then subtracted from a real contrast-free angiography (b) to produce the final DRR (c). In one experiment, Gaussian noise is added (d), here with $\sigma = 50$.

apparatus. Then, the 12 initial points are \mathbf{T}_0 displaced by $\pm\alpha$ mm along the three principal axes, and \mathbf{T}_0 rotated by $\pm\beta$ deg around the three principal axes. We used: $\alpha \in 10 * \{0, 1, \dots, 6\}$

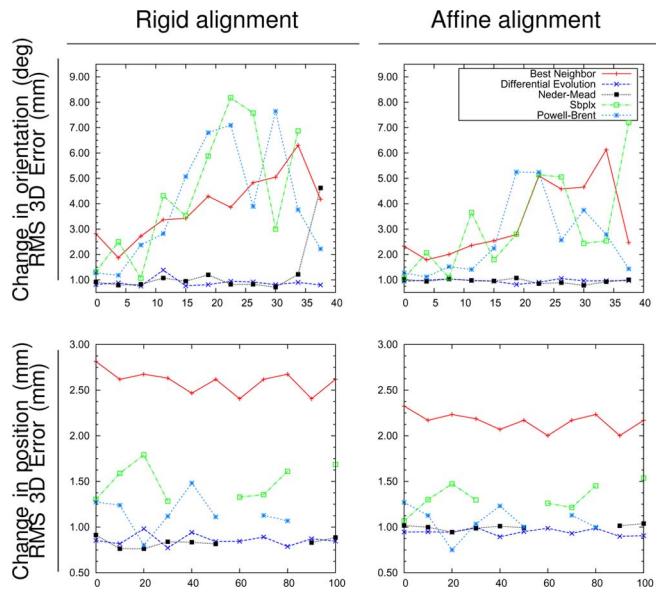


Fig. 9. Performance of the optimizers with respect to a perturbation of the initial position. *Bottom row*: missing points represent rms error > 3 mm.

mm, and $\beta \in 3.75 * \{0, 1, \dots, 6\}$ deg. The rms error was computed for all results with the same level of perturbation (translation or rotation). The results depicted in Fig. 9 show that all algorithms presented are satisfactory when the rotational perturbation is less than 10 deg. Except for the *Neder-Mead* and the *Differential Evolution* algorithms, the error level increases rapidly past this threshold. All the algorithms in Fig. 9 seem to be able to recover from perturbations of up to 30 mm in translation. At higher level, only the *Differential Evolution* and the *Best Neighbor* algorithms give consistent results, although the latter is less accurate. The *Neder-Mead* algorithm is more accurate and has a capture range of 50 mm, which makes it a practical choice. The error curves for the *Cobyla*, *Bobyqa*, *Praxis*, and *Direct* algorithms were removed, because they were significantly higher, which impaired the interpretation of the graphs.

2) *Robustness to Image Noise*: Various amounts of Gaussian noise with variance $\sigma \in \{5, 10, \dots, 60\}$ were added to the fluoroscopic images, in order to assess the impact of image quality on the performance of the method. Each optimizer was run from 13 initial positions for each noise level (as above, with $\alpha = 5$ mm, and $\beta = 7.5$ deg, plus \mathbf{T}_0). For each noise level, the rms errors are presented in Fig. 10. The results of this test tend to demonstrate that the proposed methodology is not really sensitive to the presence of Gaussian noise in the fluoroscopies. The 2D and 3D error level shows very little correlation with the amount of added noise, even at levels as high as $\sigma = 60$ that are seldom observed in a clinical setting. This is an indication of the good performance of the selected 2D segmentation algorithm. In this test, the *Neder-Mead* and the *Differential Evolution* algorithms gave the most accurate and most consistent results. Again, the error curves for the *Cobyla*, *Bobyqa*, *Praxis*, and *Direct* algorithms were removed, because they were significantly higher.

3) *Nonrigid Deformation*: In these simulations, the centerline was deformed using a $3 \times 3 \times 3$ node thin plate spline (TPS)

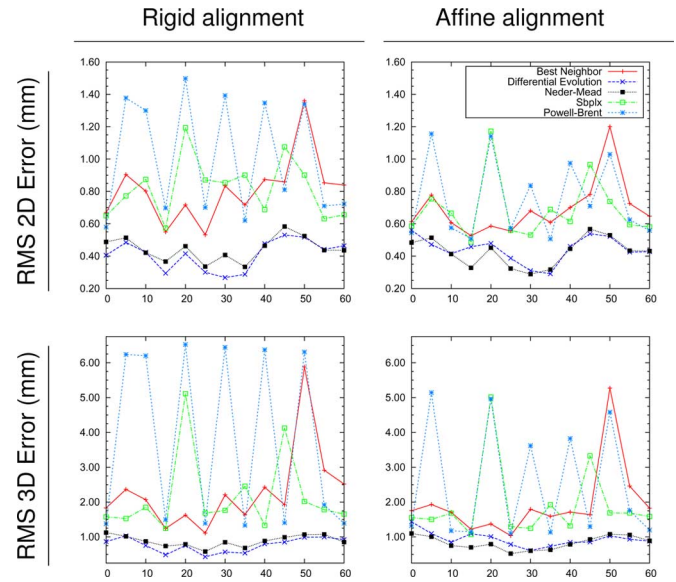


Fig. 10. Performance of the optimizers with respect to the standard deviation of the input DRR noise. Gaussian white noise model with $\sigma \in [0, 60]$.




Fig. 11. Sample nonrigid registration with simulated data. Top row: initial position; bottom row: final position.

deformation model [6] covering the CT-scanned region. Each TPS node is moved toward the center node by a factor depending on the deformation parameter ξ . The nine left nodes were shifted by a factor of ξ , the nine center nodes, by 0.5ξ , and the nine right nodes by 2ξ . Sample images are shown in Fig. 11. The nonrigid registration algorithm was tested for deformation levels $\xi \in \{0.020, 0.025, \dots, 0.060\}$, which resulted in mean and maximum 3D displacements of the centerline of $[1.209 - 3.638]$ mm and $[2.472 - 7.416]$ mm respectively. The nondeformed centerline curve was used as the initial solution, and no global alignment was performed before applying the nonrigid registration algorithm. A quantitative evaluation is presented in Fig. 12. It was found that this algorithm gives good results up to the level $\xi = 0.045$. At this point, it reduces the




Fig. 12. Residual 3D error with respect to the simulated nonrigid deformation level. *Baseline* correspond to the initial solution error.

average 3D error from 2.721 mm to 1.198 mm, which is qualitatively significant (Fig. 11). As will be demonstrated in the next sections, this level of performance is appropriate in a clinical scenario.

B. Clinical Data

Five datasets were used for the experiments presented in the following sections. Each dataset includes one CTA scan acquired at end-diastole (datasets 1 and 3) or end-systole (datasets 2, 4, and 5), and one biplane X-ray fluoroscopy recording. The two modalities are temporally aligned using ECG gating, and the coronary arteries were semi-automatically segmented in the CTA by a specialist using the technique described in [14]. The mean 3D inter-point distances of the segmented centerlines were in the [1.36 – 1.60] mm range. The angiograms image the left coronary arteries (LCA) in three cases and the right coronary artery (RCA) in the other two. The image size is 512×512 in all cases, and the image resolution is $\{0.345, 0.279, 0.279, 0.216, 0.279\}$ mm/pixel for dataset 1 to 5, respectively. The acquisition rate is 15 fps. Standard Siemens C-Arm calibration matrixes were used. Although this is not a requirement of the method, matrixes \mathbf{P}_1 and \mathbf{P}_2 were kept constant during acquisition (they are different for each dataset). The CTA and fluoroscopy exams had been prescribed to the patient for the treatment of a coronary disease.

The assessment of the performance of the alignment and registration algorithm was quantified using the mean 2D projection error

$$\text{Err}(\chi, \bar{\mathbf{x}}_{s,q}) = \frac{1}{2M} \sum_{n=1}^2 \sum_{s,q} D_{S_n}(\Pi_n(\chi(\mathbf{x}_{s,q}))) \quad (13)$$

where M is the total number of points in the 3D centerline model. Also, D_{S_n} is the distance transform of the reference 2D segmentation S_n , which was obtained by manual tracing on top of the corresponding fluoroscopy image.

C. Global Alignment: Evaluation of the Performance of the Optimizers

The performance of the nine optimizers in minimizing the global alignment energy (1) and in producing good quality 2D/3D alignment, were assessed using the following experimental setup. For each of the five patient datasets, three pairs of fluoroscopic images were considered: the images temporally aligned with the CTA acquisition, and the previous and next adjacent frames. Temporal alignment helps to minimize the observable differences between the two modalities, thereby

reducing the risk of failing a registration. This is reasonable clinically, as it would allow the registration matrix to be updated approximately once every second, in order to improve surgical guidance. Also, to test the robustness of the optimizers, 13 different initial points were used for the initialization. Let \mathbf{T}_0 be the transformation given by the calibration of the apparatus. The 13 initial points are \mathbf{T}_0 , \mathbf{T}_0 displaced by ± 5 mm along the three principal axes, and \mathbf{T}_0 rotated by ± 7.5 deg around the three principal axes. A total of $5 * 3 * 13 = 195$ experiments were thus conducted for each optimizer tested. As discussed in Section III, (1) is minimized by successively using: 1) translation-only transformation, 2) rigid transformation, and 3) affine transformation.

The mean residual energies left after minimizing with each of the three transformation models are presented in Table I, along with the measured mean 2D projection error. In addition, per patient box-and-whisker plots of the mean 2D error are displayed in Fig. 14, and sample results are shown in Fig. 13.

The total computational time needed to successively estimate the translational, rigid, and affine alignments is reasonable for all local optimizers with median values of 105 ms or less, as can be seen in Fig. 15. In all cases, the total time was under 1 s. The method is thus suitable for an interactive application with any of the local optimizers since the alignment appears to be computed almost instantaneously at the push of a button. In addition, the very short computational time of the *Neder-Mead* optimizer allows us to envision real time application. The computational time is much higher with the two global optimizers, as presented in Table II. This means that they are only usable in an offline low-interaction setting.

A look at the average values of Table I reveals that the best algorithms for minimizing the energy function with the affine transformation model are the following: *Differential Evolution*, *Neder-Mead*, *Powell-Brent*, *Best Neighbor*, *Sbplx*, and *Direct*. When the mean 2D projection error is considered, the order changes slightly: *Differential Evolution*, *Powell-Brent*, *Best Neighbor*, *Sbplx*, *Neder-Mead*, and *Direct*. The difference can probably be attributed to the discrepancy between the automatic segmentation used in the alignment process, and the manual segmentation used for computing the error. Nevertheless, the performance of the top ranking algorithms appears to be satisfying. The relatively poor performance of the *Bobyqa* and *Cobyqa* algorithms might be an indication that the shape of the energy function is not well represented by their quadratic and linear models, respectively. The *Praxis* method appears to be the least appropriate local optimizer for this problem. The box-and-whisker plots in Fig. 14 indicate that the performance of the optimizers varies from one dataset to another, even though the previously discussed observation hold. Thus, the alignment results can probably be boosted by using two or more optimizers in parallel.

The *Differential Evolution* algorithm was found to be the best performer in term of both the residual energy and the mean 2D error. However, the gains in mean 2D error were marginal for a computational time that is about 100 times that of the local algorithms. Still, the median value of the computational time of that optimizer, just under 11 s, might be reasonable for applications that are not time critical. In general,

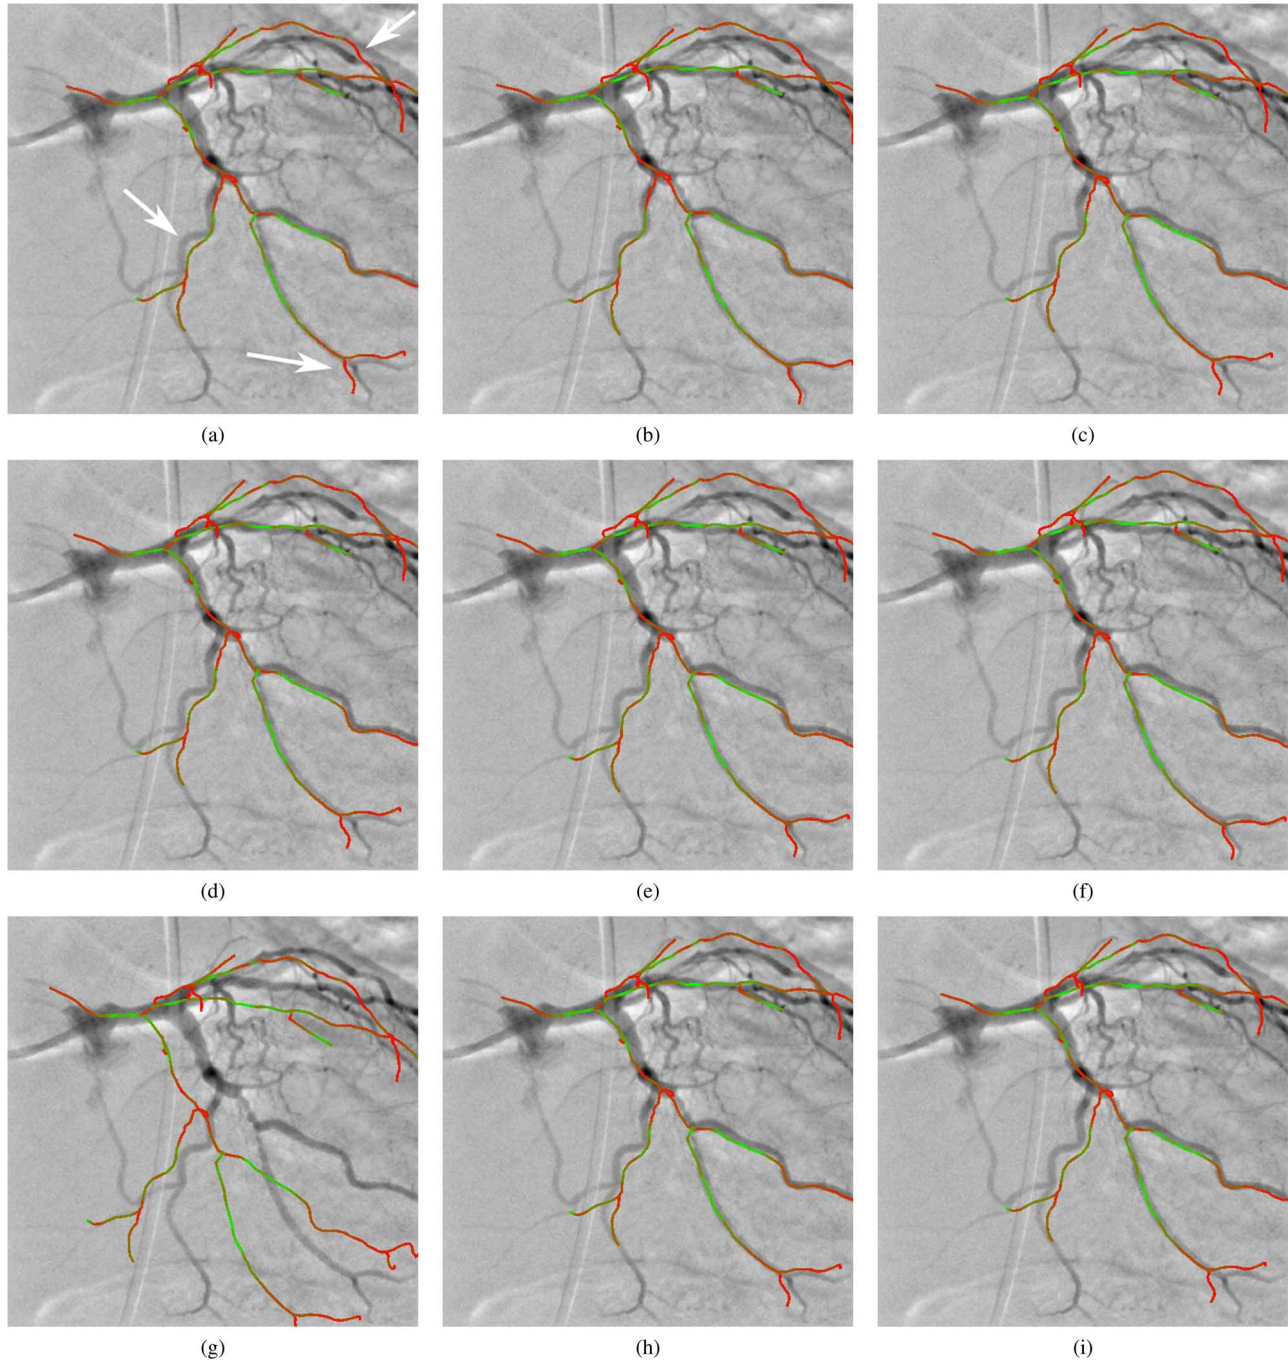


Fig. 13. Sample output obtained with the affine transformation model using nine different optimizers on the Patient 3 dataset. The values in brackets indicate the residual energy of (1). (a) *Best Neighbor* [31.86], (b) *Differential Evolution* [29.87], (c) *Neder-Mead* [30.70], (d) *Direct* [31.69], (e) *Bobyqa* [34.90], (f) *Cobyla* [36.21], (g) *Praxis* [59.68], (h) *Sbplx* [31.39], and (i) *Powell-Brent* [30.65]. Note that, although only one fluoroscopy plane is shown, the biplane pair was used. The arrows indicate some regions of interest.

it can be noted that the global optimizers only rarely lead to lower energy or error figures than the best-performing local optimizers. This is an indication that the latter actually finds solutions that are close to the global optimum. It is also worth noting that Fig. 14 shows an interesting characteristic of the global optimizers, in that they generally succeed in avoiding the worst solutions. However, because the computational time of the global optimizers is several orders of magnitude larger than that of the local optimizers, we found that they are less appropriate in our setting.

The sample image in Fig. 13 shows that there is a good correlation between the residual energy and the perceived visual correspondence. Nevertheless, even with the best alignment with the affine transformation model, a relatively large discrepancy exists between the projected centerline and the 2D fluoroscopies. This can probably be attributed to the presence of nonaffine deformation in the dataset due to the inaccuracy of the temporal alignment between the modalities, and to the shape change induced by the patient's position change during the various acquisitions steps.

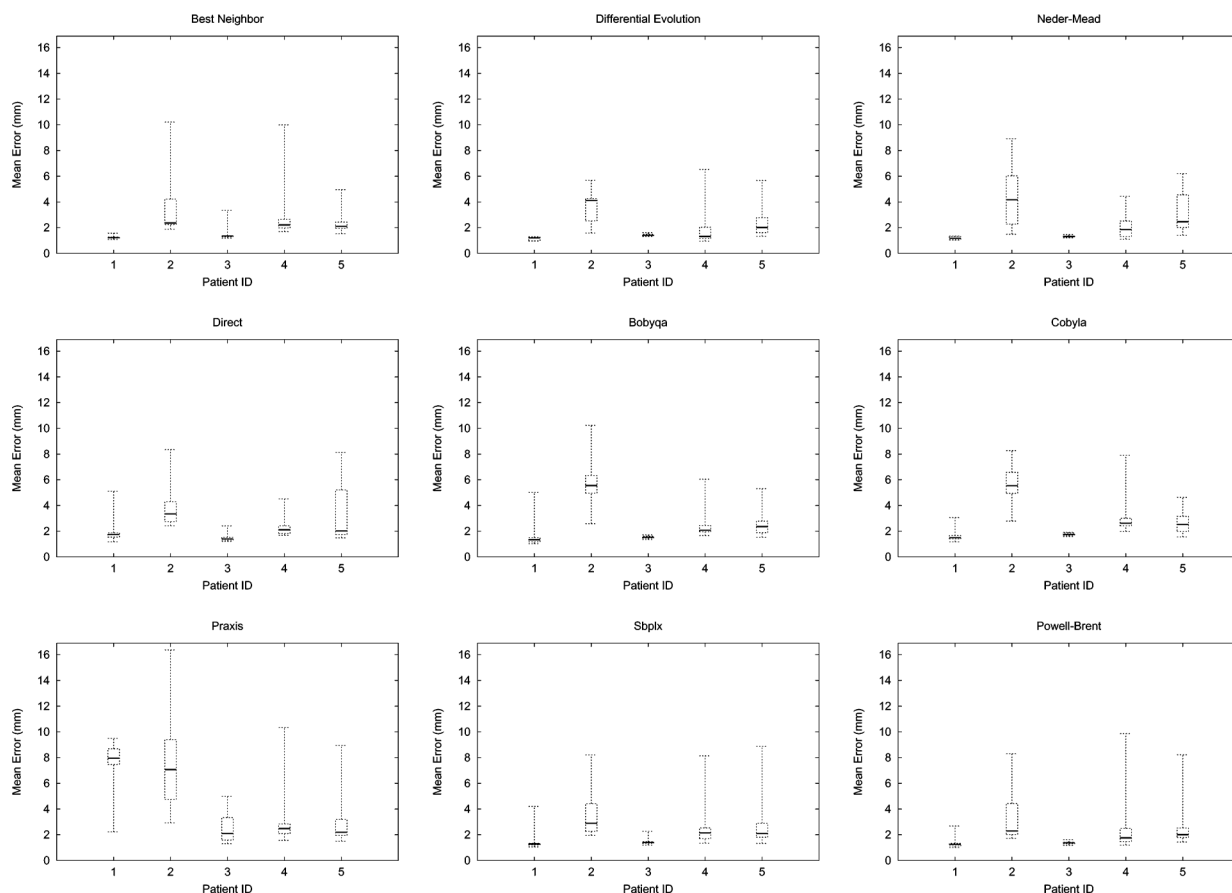


Fig. 14. Mean 2D error, per optimizer, and per patient, after alignment with the affine transformation model. Each box-and-whisker point correspond to the results of 39 experiments. The boxes represent the first, second, and third quartiles. The whiskers indicate the 5th and the 95th percentiles.




Fig. 15. Total computational time, in millisecond, for all local optimizers, when aligning the 3D centerline using the translation only, rigid, and affine transformation models, successively. The boxes represent the first, second, and third quartiles. The whiskers indicate the 5th and the 95th percentiles.

D. Comparison of the Global Alignment Method With Nonrigid Registration

In this section, detailed global alignment results obtained using the *Best Neighbor* optimizer are presented and compared with those obtained after nonrigid registration. In all cases, the centerline model was registered with the fluoroscopic frames that are gated in the same cardiac phase used for the CT reconstruction. The results after translational, rigid, affine,

TABLE I
rmsRESIDUAL ENERGY AND RESULTING 2D ERROR, IN FUNCTION OF THE TRANSFORMATION MODEL

Optimizer	Translation		Rigid		Affine	
	energy	mm	energy	mm	energy	mm
Best Neighbor	29.98	4.10	22.01	3.56	16.20	2.63
Differential Evolution	29.57	4.05	20.23	3.90	13.69	2.43
Nelder-Mead	29.74	4.03	21.24	3.76	15.42	2.78
Direct	29.73	4.07	23.85	3.76	17.38	2.85
Broyden-CG	31.20	4.67	22.58	3.55	20.51	3.29
Cobyla	30.12	4.03	22.90	3.43	22.39	3.37
Praxis	46.24	6.75	38.46	5.78	32.90	5.38
Soplex	30.00	4.15	22.22	3.65	16.69	2.74
Powell-Brent	30.23	4.16	18.05	3.58	15.79	2.60

Note: the values presented above correspond to the RMS results over 195 experiments for each optimizer.

TABLE II
TOTAL COMPUTATIONAL TIME FOR THE TWO GLOBAL OPTIMIZERS

Optimizer	Min	Q1	Median	Q3	Max
Differential Evolution	6 065	9 177	10 898	11 239	17 079
Direct	81.51	248.4	35 418	84 451	2.604e ⁵

Total computational time, in millisecond, when aligning the 3D centerline using the translation only, rigid, and affine transformation models, successively. Q1 and Q3 correspond to the first and third quartiles respectively.

and nonrigid registration are shown for the dataset from three patients in Figs. 3 and 16. The mean projection error of the




Fig. 16. From left to right: input image, rigid alignment, affine alignment, and nonrigid registration. Both biplane images are shown in all cases. Top and bottom rows correspond to Patients 1 and 2 in Table III, respectively. The arrows indicate some regions of interest.

TABLE III
MEAN RESIDUAL 2D PROJECTION ERROR, IN MM, CALCULATED AFTER RIGID, AFFINE, AND NONRIGID REGISTRATION FOR FIVE PATIENTS

	Subject	Translation			Rigid			Affine			Non-Rigid		
		mean	std	max	mean	std	max	mean	std	max	mean	std	max
LCA	1	3.248	3.547	15.661	1.168	1.149	7.031	1.174	1.155	7.031	0.663	0.742	7.031
	2	5.734	4.945	18.621	5.352	4.830	17.959	2.241	2.879	14.440	1.537	1.969	14.440
	3	2.430	2.325	14.398	1.731	1.588	8.700	1.204	1.088	7.886	0.838	0.856	7.886
RCA	4	2.653	2.565	10.205	2.558	2.458	10.037	2.210	2.376	9.190	1.049	1.281	8.243
	5	3.161	2.701	11.701	2.627	2.286	10.259	2.627	2.286	10.259	1.380	1.333	7.251

centerline on the two fluoroscopic planes with respect to 2D manual segmentations has also been calculated (in millimeters). Results are presented in Table III. The computational times needed for the nonrigid registration were in the range [730–3 300] ms, with an average of 1 591 ms.

The mean 2D projection errors and their standard deviations decrease, or stay approximately constant, as the complexity of the model increases. However, the relative contribution of the

rigid, affine, and nonrigid transformation model varies from one dataset to another. This can be explained by the nature of the deformation presented by each individual case, which is linked to the patient's position, the interval between the CTA and 2D fluoroscopy acquisitions, and the acquisition protocol. Also, patient respiration can cause significant nonrigid heart deformation. The accuracy of the temporal alignment by ECG gating is also important, since the beating of the heart is a significant




Fig. 17. Registration over a sequence of frames. Top row) rigid transformation model [$\gamma = 0.1$, $E_{\text{Multi}}^* = 120.284$]; middle row) affine transformation model [$\gamma = 0$, $E_{\text{Multi}}^* = 119.356$]; and bottom row) affine transformation model + regularization [$\gamma = 0.1$, $E_{\text{Multi}}^* = 116.905$]. The arrows indicate some regions of interest. E_{Multi}^* corresponds to the first righthand term of (5).

source of nonrigid deformation. The maximal error also tends to decrease with model complexity, but it sometimes stays stable. In those cases, the performance of the registration algorithm might be limited by an incomplete 2D segmentation, as computed by the automatic segmentation method [37]. In fact, in regions where little or no information is available, the nonrigid registration method should alter the centerline as little as possible. This characteristic is desirable in situations where parts of the structure are poorly visible on the X-ray fluoroscopic images, as is sometimes the case with CTO. For an example, see Fig. 16(bottom) rows: the centerline is well aligned over and under the CTO, but there is little deformation where the contrast is poor.

E. Global Alignment in the Multi Frame Scenario

Another potentially useful scenario, beyond the alignment of a centerline with a biplane pair, is the alignment of a centerline with a temporal sequence of frames from biplane angiography. The performance of the global alignment method in the multi-frames scenario is demonstrated using a sequence of 11 biplane frames from the Patient 1 LCA dataset. Global alignments were computed using the rigid and affine transformation models. The inter-frame regularization parameters γ in (5) varied between 0.0 and 1.0. The *Best Neighbor* optimizer was used in this case since it was found to be the most effective at minimizing the cost function (5). Sample results obtained are shown in Fig. 17. Computational times were 39 s, 89 s, and 91 s, for the experiments presented in the top, middle, and bottom row

respectively. As anticipated, it was found that using the affine transformation model led to better qualitative results than using the rigid one. Using regularization ($\gamma = 0.1$) also improved the results when using the affine transformation model. With large regularization values, $\gamma > 1.0$, the stiffness of the model increases, and the optimizers cannot find a local minima that is far from the initial position. Preliminary investigation suggests that domain specific optimizers might achieve better results, e.g., by allowing the parameters to change in a coordinated manner. Nevertheless, we found that the current setup was appropriate for tracking the alignment transformations of the LCA during most of the cardiac cycle. Experiments on the RCA datasets did not lead to convincing results. In fact, even though the global alignment procedure works, the amount of nonrigid deformation sustained by the RCA during the cardiac cycle renders the affine model ineffective except when the CTA and X-ray acquisition are in close temporal alignment.

F. Semiautomatic Tracking of the Right Coronary Artery

In this experiment, a semiautomatic procedure based on the proposed nonrigid registration method has been used to track the RCA. Starting from the gated frame, global alignment and nonrigid registration are performed. This deformed centerline model is then used by the operator as the initial model and position for the next pair of fluoroscopic frames. This process is repeated for all frames over one cardiac cycle, as presented in Fig. 18 and in the attached video. As can be seen, this procedure permits successful tracking of the RCA during one cardiac




Fig. 18. Tracking the RCA. From left to right: at the gated frame, and at +4, +6, +8, and +10 frames. The initial curve is green/red, and the registered curve is blue.

cycle. Because this experiment has been conducted before the simulation study, a slightly different set of parameters was used: $\mu = 0.05$, $\nu = 5.0$, and $\lambda = 0.1$. The required amount of time is about 15 min, which is reasonable in an offline setting.

VI. MULTIMEDIA MATERIAL

Three videos illustrating the method are available online.⁵ They present: 1) an iterative update of the nonrigid registration; 2) the full alignment and registration progress; and 3) the tracking of an RCA artery using nonrigid registration.

VII. DISCUSSION AND CONCLUSION

A new 2D/3D registration method has been proposed, and applied to the problem of registering a 3D centerline model of the coronary arteries with a pair of fluoroscopic images. The methodology is divided into two main parts: 1) global alignment, and 2) nonrigid registration.

In the first part, an energy depending on a global transformation model (translation-only, rigid, or affine) is defined. Nine general purpose optimization algorithms have been used to minimize this energy, which results in an estimation of the 2D/3D alignment transformation parameters. Based on the experiments on clinical data, it appears that the following local optimizers give good results in a short time (median < 105 ms, max < 950 ms): *Neder-Mead*, *Powell-Brent*, *Best Neighbor*, and *Sbplx*. Considering also the simulation results, the *Neder-Mead* algorithm was the best overall performer. The two global optimizers only rarely led to major improvements in the result and required a computational time orders of magnitude higher. Disregarding computational time, the *Differential Evolution* algorithm generally returned the best solution. When using a local optimizer, the alignment time was consistently under 1 s, which makes the method suitable for use during an intervention. The advantage of using an affine transformation instead of a rigid transformation is dependent on the nature of the dataset and can be significant in some cases. Overall, it was found that the global alignment procedure is appropriate for use on both LCA and RCA datasets, when the 3D and 2D modalities are temporally aligned using ECG gating. The experiment on the dataset from a CTO patient demonstrates the benefit of the proposed method when applied to similar clinical cases. The proposed method can also help with intervention guidance by augmenting the 2D imagery with the 3D geometry segmented from a CTA acquisition, thereby greatly reducing the ambiguities inherent in the interpretation of the 2D images.

⁵<http://www.synchromedia.ca/reg2D3D>

As for nonrigid registration, the proposed method uses a robust reconstruction strategy to compute the forces used to deform the 3D model. Regularization terms limit the total displacement of the segments, ensure that the displacements are smooth, and help preserve the relative orientation of the branches, making the nonrigid registration results plausible. The regularizers included in the deformation model enable the elegant management of regions with missing data (e.g., Patient 2 in Fig. 16) by deforming the centerline a minimum amount to ensure coherence with the rest of the structure. No excessive deformation or straightening will occur. Experiments on five different patients were presented with promising results. The total computational time, which was generally below 3 s, is acceptable for interactive applications. Nonetheless, better numerical algorithms with faster computational time will be researched. The proposed nonrigid registration method will make the centerline snap to adjacent 2D structures when started from an appropriate initial point, making interpretation of the 2D images easier, especially for difficult low contrast CTO cases. This also makes it possible to present an updated 3D model alongside the operational images to provide an improved perception of the 3D space.

Finally, experiments using multiple X-ray biplane angiography frames have also been presented. It was found that the multiframe global alignment method works well on LCA datasets, and the proposed interframe regularizer leads to improved results. A nonrigid semiautomatic tracking procedure has been devised to handle cases with more nonrigid deformation, such as with nontemporally aligned RCA datasets, and applied to one such dataset. The semi-automatic method required approximately 15 min of interaction for a sequence covering one heart beat, which seems reasonable for practical offline applications. Future work will focus on automating this nonrigid method for tracking coronary arteries.

ACKNOWLEDGMENT

The C++ code for DRR generation has been generously provided by Julien Couet and Luc Duong of École de technologie supérieure.

REFERENCES

- [1] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, "Log-Euclidean metrics for fast and simple calculus on diffusion tensors," *Magn. Reson. Med.*, vol. 56, pp. 411–421, 2006.
- [2] S. Benamer, M. Mignotte, S. Parent, H. Labelle, W. Skalli, and J. de Guise, "3D/2D registration and segmentation of scoliotic vertebrae using statistical models," *Comput. Med. Imag. Graph.*, vol. 27, pp. 321–337, 2003.

- [3] C. Blondel, G. Malandain, R. Vaillant, and N. Ayache, "Reconstruction of coronary arteries from a single rotational X-ray projection sequence," *IEEE Trans. Med. Imag.*, vol. 25, no. 5, pp. 653–663, May 2006.
- [4] J. Boisvert, F. Cheriet, X. Pennec, H. Labelle, and N. Ayache, "Geometric variability of the scoliotic spine using statistics on articulated shape models," *IEEE Trans. Med. Imag.*, vol. 27, no. 4, pp. 557–568, Apr. 2008.
- [5] I. M. J. van der Bom, S. Klein, M. Staring, R. Homan, L. W. Bartels, and J. P. W. Pluim, "Evaluation of optimization methods for intensity-based 2D-3D registration in X-ray guided interventions," in *SPIE Med. Imag.*, 2011.
- [6] F. Bookstein, "Principal warps: Thin-plate splines and the decomposition of deformations," *IEEE Pattern Anal. Mach. Intell.*, vol. 11, no. 6, pp. 567–585, Jun. 1989.
- [7] S. Bouattour, R. Arndt, and D. Paulus, "4D reconstruction of coronary arteries from monoplane angiograms," in *Proc. Comput. Analysis Images Patterns*, 2005, pp. 724–731.
- [8] R. Brent, *Algorithms for Minimization Without Derivatives*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [9] L. Duong, R. Liao, H. Sundar, B. Tailhades, and C. Xu, "Curve-based 2D-3D registration of coronary vessels for image guided procedure," in *Proc. SPIE Med. Imag.*, 2009.
- [10] R. Fabbri, L. da F. Costa, J. C. Torelli, and O. M. Bruno, "2D euclidean distance transform algorithms: A comparative survey," *ACM Comput. Surv.*, vol. 40, no. 1, pp. 1–44, 2008.
- [11] J. Feldmar, N. Ayache, and F. Betting, "3D-2D projective registration of free-form curves and surfaces," *Comput. Vis. Image Understand.*, vol. 65, pp. 403–424, 1997.
- [12] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI'98*, 1998, pp. 130–137.
- [13] M. Groher, D. Zikic, and N. Navab, "Deformable 2D-3D registration of vascular structures in a one view scenario," *IEEE Trans. Med. Imag.*, vol. 28, no. 6, pp. 847–860, Jun. 2009.
- [14] M. Gülsün and H. Tek, "Robust vessel tree modeling," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008*, 2008, pp. 602–611.
- [15] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [16] J. H. Hipwell, G. P. Penney, R. A. McLaughlin, K. Rhode, P. Summers, T. C. Cox, J. V. Byrne, J. A. Noble, and D. J. Hawkes, "Intensity-based 2-D-3-D registration of cerebral angiograms," *IEEE Trans. Med. Imag.*, vol. 22, no. 11, pp. 1417–1426, Nov. 2003.
- [17] S. G. Johnson, The Nlopt Nonlinear-Optimization Package [Online]. Available: <http://ab-initio.mit.edu/nlopt>
- [18] D. R. Jones, C. D. Perttunen, and B. E. Stuckmann, "Lipschitzian optimization without the lipschitz constant," *J. Optimizat. Theory Appl.*, vol. 79, pp. 157–181, 1993.
- [19] J. Couet, D. R. H., J. Miró, C. Lapierre, L. Duong, and M. Cheriet, "Intensity-based 3d/2d registration for percutaneous intervention of major aorto-pulmonary collateral arteries," in *Proc. SPIE Med. Imag.*, 2012.
- [20] K. K. Lau and A. C. S. Chung, "A global optimization strategy for 3D-2D registration of vascular images," in *Proc. 17th Br. Mach. Vis. Conf. (BMVC 06)*, 2006, pp. 489–498.
- [21] R. Liao, Y. Tan, H. Sundar, M. Pfister, and A. Kamen, "An efficient graph-based deformable 2D/3D registration algorithm with applications for abdominal aortic aneurysm interventions," in *Med. Imag. Augmented Reality*, 2010, pp. 561–570.
- [22] P. Markelj, D. Tomazevic, B. Likar, and F. Pernus, "A review of 3D/2D registration methods for image-guided interventions," *Med. Image Anal.*, to be published.
- [23] C. Metz, M. Schaap, S. Klein, L. Neeffjes, E. Capuano, C. Schultz, R. J. van Geuns, P. W. Serruys, T. van Walsum, and W. J. Niessen, "Patient specific 4D coronary models from ECG-gated CTA data for intraoperative dynamic alignment of CTA with X-ray images," in *MICCAI 2009*, 2009, pp. 369–376.
- [24] J. A. Nelder and R. Mead, "A simplex method for function minimization," *Comput. J.*, vol. 7, no. 4, pp. 308–313, 1965.
- [25] N. Paragios, M. Rousson, and V. Ramesh, "Non-rigid registration using distance functions," *Comput. Vis. Image Underst.*, vol. 89, pp. 142–165, Feb. 2003.
- [26] X. Pennec and J. P. Thirion, "A framework for uncertainty and validation of 3-D registration methods based on points and frames," *Int. J. Comput. Vis.*, vol. 25, no. 3, pp. 203–229, 1997.
- [27] T. Peters and K. Cleary, Eds., *Image-Guided Interventions: Technology and Applications*. New York: Springer Science, 2008.
- [28] T. M. Peters, "Image-guidance for surgical procedures," *Phys. Med. Biol.*, pp. R505–R540, 2006.
- [29] M. J. D. Powell, "An efficient method for finding the minimum of a function of several variables without calculating derivatives," *Comput. J.*, vol. 7, no. 2, pp. 155–162, 1964.
- [30] M. J. D. Powell, "Advances in Optimization and Numerical Analysis, chap," in *A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation*. Dordrecht: Kluwer Academic, 1994, pp. 51–67.
- [31] M. J. D. Powell, Direct search algorithms for optimization calculations Dept. Appl. Math. Theoretical Phys., Cambridge Univ., U.K., Tech. Rep., 1998.
- [32] M. J. D. Powell, The BOBYQA algorithm for bound constrained optimization without derivatives Centre Math. Sci., Cambridge Univ., Cambridge, U.K., Tech. Rep., 2009.
- [33] M. Prümmer, J. Hornegger, M. Pfister, and A. Dörfler, "Multi-modal 2D-3D non-rigid registration," in *Proc. SPIE Med. Imag.*, 2006.
- [34] T. Rowan, "Functional stability analysis of numerical algorithms," Ph.D. dissertation, Dept. Comput. Sci., Univ. Texas, Austin, 1990.
- [35] D. Ruijters, B. M. ter Haar Romeny, and P. Suetens, "Vesselness-based 2D-3D registration of the coronary arteries," *Int. J. CARS*, pp. 391–397, 2009.
- [36] Y. 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.*, vol. 2, no. 2, pp. 143–168, 1998.
- [37] M. Schneider and H. Sundar, "Automatic global vessel segmentation and catheter removal using local geometry information and vector field integration," in *IEEE Int. Symp. Biomed. Imag.*, Apr. 2010, pp. 45–48.
- [38] K. Shoemake and T. Duff, "Matrix animation and polar decomposition," in *Int. Proc. Conf. Graphics Interface*, 1992, pp. 258–264, Morgan Kaufmann.
- [39] L. S. Souha Aouadi, "Accurate and precise 2D-3D registration based on X-ray intensity," *Comput. Vis. Image Understand.*, vol. 110, pp. 134–151, 2007.
- [40] C. Stewart, "Handbook of Mathematical Models in Computer Vision," in *Point-Based Image Registration*. New York: Springer, 2006, ch. 14, pp. 221–235.
- [41] R. Storn and K. Price, "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optimizat.*, vol. 11, pp. 341–359, 1997.
- [42] H. Sundar, A. Khamene, C. Xu, F. Sauer, and C. Davatzikos, "A novel 2D-3D registration algorithm for aligning fluoro images with 3D pre-op CT/MR images," in *Proc. SPIE Med. Imag.*, 2006.
- [43] M. V. Truong, A. Aslam, C. A. Rinaldi, R. Razavi, G. P. Penney, and K. S. Rhode, "Preliminary investigation: 2D-3D registration of MR and X-ray cardiac images using catheter constraints," in *Workshop Cardiovascular Intervent. Imag. Biophys. Modell.*, 2009.
- [44] Y. Tsin, K. Kirchberg, G. Lauritsch, and C. Xu, "A deformation tracking approach to 4D coronary artery tree reconstruction," in *MICCAI*, 2009, pp. 68–75.
- [45] G. A. Turgeon, G. Lehmann, G. Guiraudon, M. Drangova, D. Holdsworth, and T. Peters, "2D-3D registration of coronary angiograms for cardiac procedure planning and guidance," *Med. Phys.*, vol. 32, pp. 3737–3749, 2005.