Auhtors' accepted manuscript
Paper presented at *Medical Image Computing and Computer-Assisted Intervention* (MICCAI 2012) (Nice, France, October 1-5, 2012)
Published in *Lecture Notes in Computer Science*, vol. 7510, p. 651-658.

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-642-33415-3_80

Stochastic 3D Motion Compensation of Coronary Arteries from Monoplane Angiograms

Jonathan Hadida*, Christian Desrosiers, and Luc Duong

École de technologie supérieure, 1100 rue Notre-Dame Ouest, Montreal, Canada
{jonathan.hadida.1, christian.desrosiers, luc.duong}@etsmtl.ca

Abstract. Image-based navigation during percutaneous coronary interventions is highly challenging since it involves estimating the 3D motion of a complex topology using 2D angiographic views. A static coronary tree segmented in a pre-operative CT-scan can be overlaid on top of the angiographic frames to outline the coronary vessels, but this overlay does not account for coronary motion, which has to be mentally compensated by the cardiologist. In this paper, we propose a new approach to the motion estimation problem, where the temporal evolution of the coronary deformation over the cardiac cycle is modeled as a stochastic process. The sequence of angiographic frames is interpreted as a probabilistic evidence of the succession of unknown deformation states, which can be optimized using particle filtering. Iterative and non-rigid registration is performed in a projective manner, and relies on a feature-based similarity measure. Experiments show promising results in terms of registration accuracy, learning capability and computation time.

1 Introduction

Percutaneous coronary intervention (PCI) is a routine procedure in cardiology aimed at providing revascularization of coronary vessels. Such interventions remain highly challenging since cardiologists must rely on occasional contrast injection to assess coronary motion and topology on 2D X-ray angiographic frames. CT-based navigation was proposed [1] to outline coronary vessels even when contrast agent is not present, but these techniques provide only static guidance. The task of intra-vascular navigation through the coronary arteries is impeded by the complexity of their topology, their constant motion composed of sudden movements, and the fact that they can often be observed only one angle at a time. An accurate estimation of this motion would be a step further towards efficient motion compensation techniques, or augmented reality solutions, where angiographic frames would be overlaid with interactive catheter roadmaps. This problem is difficult for several reasons: first, the coronary motion is a complex combination of rigid and non-rigid deformations caused by both cardiac and respiratory activities. Second, the topology of the coronary arteries is remarkably complex and varies significantly from one patient to another. Third, because

* This research was supported by NSERC Discovery Grant.

we aim for an intra-operative application, computational costs have to remain low. Last, the estimation of 3D motion from 2D temporal projections is notably known as an *ill-posed* problem, therefore difficult by essence.

1.1 Related works

The goal of 3D/2D registration methods is to find a geometric transformation that brings a 3D pre-operative patient dataset into the best spatial correspondence with 2D intra-operative datasets [2]. Such methods can be roughly grouped in three separate categories. The first category proposes simple rigid deformation models and focuses on the *registration criterion*, implicitly seeking a smooth criterion to be optimized [3]. The second category derives from the work on active shape models, and focuses on the *deformation model* using statistical atlases [4, 5]. These methods rely on the assumption that an average topology can be drawn from multiple segmentations, and that patient-specific structures can later be consistently described as deviations from this mean topology. In our case, this assumption does not apply since the topology of the coronary arteries varies too much from one patient to another. Recently, [6] proposed a graph-based non-rigid deformation model showing good results despite the important number of parameters and the heavy computational costs. Other attempts at motion compensation in the field of 3D reconstruction require multiple views and generally propose heavily parameterized deformation models [7–9]. Finally, the methods in the third category put their attention on the *optimization strategy* [10, 11], arguing that the ill-posedness of 3D/2D registration calls for the use of robust optimization techniques allowing multiple hypothesis to coexist during the search, such as techniques based on Monte-Carlo sampling. Nevertheless, very few of the works in the litterature explicitly consider temporal consistency in the problem of coronary motion estimation. Furthermore, when included, such temporal constraints are usually applied between consecutive frames only, instead of considering a stable and coherent deformation over the entire cardiac cycle – a point that is stressed out in the discussion of [8].

1.2 Contributions

We propose a novel probabilistic approach to describe and predict the 3D deformations of the coronary arteries from 2D monoplane angiogram. The originality and advantages of this approach, with respect to existing methods for the problem of 3D/2D registration of coronary arteries, are the following:

- As opposed to current methods, our approach enforces temporal consistency over the entire sequence by using a Hidden Markov Model, which enables a more robust and accurate registration;
- Unlike deterministic registration methods, our probabilistic approach considers multiple hypotheses simultaneously during the optimization process, thereby avoiding the problem of local optima;

– Since its deformation model has few parameters and the optimization process, based on particle filtering, is highly parallelizable, it is computationally efficient and could eventually be used in a real-time registration setting.

The rest of this paper is divided as follows. The next section presents our method and the probabilistic model on which it is based. In the experimental section, we evaluate the accuracy of this method on simulated and real patients angiograms. Finally, we conclude this paper by summarizing our contributions and experimental results, and provide possible extensions of this work.

2 Proposed method

The 3D/2D coronary registration problem can be formulated as follows. Let $\boldsymbol{I}_{1..K}$ be a sequence of K segmented binary angiographic images, where image $\boldsymbol{I}_k$ is taken at time t_k, and denote by $\boldsymbol{X}_0$ the set of points forming the centerline of the reference 3D coronary tree. As mentioned above, this reference data is acquired prior to the registration process, for instance, by segmenting vessel centerlines from a CT-scan. The goal is to find the sequence of deformed 3D centerline points $\boldsymbol{X}_{1..K}$ that best corresponds to the sequence of observed images, once projected in the image plane. In practice, optimizing every deformed centerline coordinate directly, at each instant, is not feasible due to the largeness of the resulting search space. Instead, we suppose that these coordinates are fully determined by the parameters $\boldsymbol{\theta}$ of a deformation model $f_d(\boldsymbol{X}_0; \boldsymbol{\theta})$, and search for the most likely sequence of parameters $\boldsymbol{\theta}_{1..K}$. In the next section, we describe a probabilistic model that offers a simple yet realistic description of the coronary motion.

2.1 Probabilistic generative model

The general tracking approach proposed by [12] was adapted to the problem of coronary motion compensation by including a model of cardiac motion based on Hidden Markov Models (HMM). This model, illustrated in Figure 1, receives at each instant t_k an estimate $\hat{\phi}_k$ of the 2π-period phase of the cardiac cycle, and describes the evolution of the unknown cardiac deformations as follows[1]:

$$p(\boldsymbol{\theta}_k \mid \boldsymbol{\theta}_{k-1}, \boldsymbol{v}_{k-1}) \sim \delta(\boldsymbol{\theta}_k, \boldsymbol{\theta}_{k-1} + (t_k - t_{k-1})\, \boldsymbol{v}_{k-1}) \quad (1)$$

$$p(\boldsymbol{v}_k \mid \boldsymbol{v}_{k-1}, \phi_k) \sim \mathcal{N}((1-\alpha)\, \boldsymbol{v}_{k-1} + \alpha\, \boldsymbol{v}_{\mathrm{e}}(\phi_k), \boldsymbol{\Sigma}_v) \quad (2)$$

$$p(\phi_k) \sim \mathcal{N}(\hat{\phi}_k, \sigma_\phi^2) \quad (3)$$

Instead of estimating directly the deformation parameters $\boldsymbol{\theta}_k$, we evaluate their instantaneous variations $\boldsymbol{v}_k$, using a time-linear approximation (Equation 1). By constraining the variations instead of the actual parameters, we obtain smoother and more realistic movements. Moreover, the transition model of these parameter variations, as defined in Equation 2, is constrained in two important ways. First, we consider a temporal constraint that acts as an *inertial factor* on the movement. Again, this allows to focus on regular, jitter-free movements.

[1] δ is the Kronecker symbol, and $\mathcal{N}$ is the multivariate normal law.

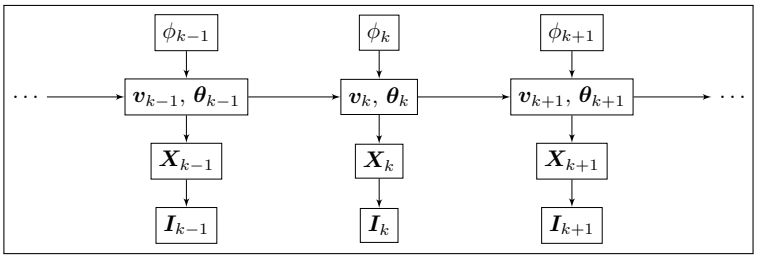

Fig. 1. Graphical model (HMM) of the image generation process.

However, because the cardiac motion contains sudden movements caused by ventricular contractions, we also inject *prior knowledge*, which enables the model to anticipate and follow such movements. Thus, $\boldsymbol{v}_k$ also depends on an empirical model (or *template*) of the cardiac motion, given by the continuous function $\boldsymbol{v}_\mathrm{e}$. Since the cardiac motion is cyclic, $\boldsymbol{v}_\mathrm{e}$ is defined as a 2π-periodic function of the cardiac phase ϕ_k. To control the trade-off between having a purely inertial model or a model based only on empirical knowledge, we use the parameter $\alpha \in [0, 1]$. Moreover, covariance matrix $\boldsymbol{\Sigma}_v$ is used to model the inherent uncertainty of the estimation and the possible temporal correlation between the deformation parameters. In Section 2.5, we describe how $\boldsymbol{v}_\mathrm{e}$ and $\boldsymbol{\Sigma}_v$ can be learned.

Finally, as expressed in Equation 3, we suppose the measured phase values to be noisy estimates of the real values, and use parameter σ_ϕ^2 to model this noise. In practice, the cardiac cycle phase can be estimated from the ECG signal by synchronizing $(\phi_k \bmod 2\pi) = 0$ with end-diastolic phases. Interestingly, this latter assumption allows for the exploration of a wider range of deformations at contraction time. Indeed, for a given phase uncertainty, a larger range of values can be found in the template function $\boldsymbol{v}_\mathrm{e}$ in regions where this function has greater variations. Additionally, slower component of the cardiac motion (e.g., the respiratory component) can be implicitely captured by this model, due to the first-order inertial constraints allowing small, potentially directed velocity increments throughout the sequence.

2.2 Cardiac deformation model

The unknown 3D deformed centerline points at time t_k are assumed to be fully determined once the parameters $\boldsymbol{\theta}_k$ of the deformation model f_d are estimated:

$$p(\boldsymbol{X}_k|\boldsymbol{\theta}_k) \;=\; \delta(\boldsymbol{X}_k,\, f_d(\boldsymbol{X}_0,\, \boldsymbol{\theta}_k)) \tag{4}$$

The choice of f_d is critical to the efficiency and accuracy of registrations. Such model should have the following two properties: *i*) be expressive enough to allow realistic deformations, and *ii*) have a limited number of parameters to reduce the search space at registration time. For this work, we chose the *planispheric* deformation model defined in [13], which is presented as a compromise between spherical and cylindrical models, and provides a convenient affine formulation to

non-rigid deformations in a transformed coordinates system. This system relies on a geometrical model of the left ventricle, and can be conceived as a spherical frame shifting along a portion of the long axis. Our deformation model comprised 6 rigid parameters (rotations and translations), and 5 non-rigid parameters.

2.3 Image generation and likelihood

The last component of our probabilistic model describes how 2D images $\boldsymbol{I}_k$ are generated from the deformed 3D centerline points $\boldsymbol{X}_k$, and how their likelihood $p(\boldsymbol{I}_k|\boldsymbol{X}_k)$ is evaluated. We considered a feature-based similarity measure involving a Frangi *et al.* vesselness filter [14] followed by a graylevel thresholding of angiographic frames $\boldsymbol{I}_{1..K}$ to obtain binary images of the coronary vessels. In order to diffuse the information of proximity to the target vessels, a distance transform was computed on the resulting sequences, using the distance $d(\gamma; \boldsymbol{di}) = ((1 + \boldsymbol{di})^{\gamma})^{-1}$, where $\gamma \geq 0$ is a diffusion factor and $\boldsymbol{di}$ is the distance to nearest centerline pixel. From preliminary tests, we found $\gamma = 0.5$ to give optimal results and used this value for our experiments. Given a predicted deformation $\boldsymbol{\theta}_k$, a 2D predicted image $\hat{\boldsymbol{I}}_k$ is then generated by projecting the 3D deformed centerlines points $\boldsymbol{X}_k$, and compared to the target image using a cosine similarity:

$$p(\boldsymbol{I}_k|\boldsymbol{X}_k) \;=\; \mathrm{sim}(\boldsymbol{I}_k, \hat{\boldsymbol{I}}_k) \;=\; \frac{\sum_{i,j} \boldsymbol{I}_k(i,j) \cdot \hat{\boldsymbol{I}}_k(i,j)}{\left(\sum_{i,j} \boldsymbol{I}_k^2(i,j)\right)^{1/2} \cdot \left(\sum_{i,j} \hat{\boldsymbol{I}}_k^2(i,j)\right)^{1/2}} \qquad (5)$$

2.4 Deformation parameters optimization

To find the most likely sequence of parameters $\boldsymbol{\theta}_{1..K}$, given the observed images and measured phase values, we use particle filtering [15]. Essentially, a set of candidate deformations (the *particles*) are extended using the transition model and resampled according to their observation likelihood (Equation 5). At each frame, the deformation used for registration is the one having the highest likelihood. While well-known approaches, such as Kalman filtering, exist to compute the *a posterori* distribution directly, we have selected particle filtering for three important reasons. First, unlike Kalman filters, particle filters can learn multimodal distributions, which allows to consider several hypotheses simultaneously. Secondly, it is highly parallelizable and, therefore, a good candidate for real-time registration. Finally, it allows one to use probability models that can not be expressed in closed form, such as our image generation process.

2.5 Learning $\boldsymbol{v}_{\mathrm{e}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{v}}$

The empirical knowledge $\boldsymbol{v}_{\mathrm{e}}$ and the covariance matrix $\boldsymbol{\Sigma}_v$ play important roles in our model, respectively guiding the search, and controling its range during optimization. However, when these values are unknown, our model can still be used with purely inertial transitions ($\alpha = 0$), compensating for the loss of *prior*

by increasing manually the search space controled by $\boldsymbol{\Sigma}_v$, and increasing the number of particles accordingly, at the expense of longer computation times. In turn, this allows one to estimate both parameters from multiple registrations performed without empirical knowledge.

3 Results and Discussion

Our dataset contained the 3D coronary centerlines – extracted from pre-operative CT-scans – and associated intra-operative biplane angiographic sequences of 7 patients (5 left coronary arteries, 2 right, sequences from 10 to 18 frames). The projection matrices were computed from the geometric parameters of the imaging systems. Although not necessary, covariance matrix $\boldsymbol{\Sigma}_v$ was assumed to be diagonal to simplify the configuration and learning steps of our experiments. The registration was performed in a coarse-to-fine manner, dividing the search of the deformation parameters in two steps. The first step looked for rigid deformations using N_1 particles, while the second step looked for both rigid and non-rigid deformations with N_2 particles. Target-to-registration errors were used to evaluate the accuracy of the compensation, in 3D (TRE 3D) during the simulations, and in 2D (TRE 2D) for both simulations and real patient datasets. For the latter, the TRE 2D was computed, with knowledge of the pixel-spacing, as the average pixel-wise product between an Euclidian distance transform computed on the target binary image, and the registered binary projection of the centerline.

3.1 Evaluation on synthetic sequences

For these experiments, matrix $\boldsymbol{\Sigma}_v$ and function $\boldsymbol{v}_\mathrm{e}$ were configured manually to randomly sample visually realistic deformation sequences with motion noise. The deformed centerlines were projected on binary images, and thickened prior to the distance transform, to better simulate thresholded, vesselness-filtered angiographic images. Four experiments aimed at quantifying the influence, respectively of: template confidence, noise, coronary topology, and number of particles on the registration accuracy (see Figure 2). The TRE 2D was less than 1.2 mm for all experiments, and as expected, degraded with increasing noise, and decreased with increasing prior knowledge or number of particles. Nevertheless, if TRE 3D globally followed these variations, the 3D errors remained high during simulations (average of 4 mm), and the results indicated that the 3D accuracy was sensitive to the template $\boldsymbol{v}_\mathrm{e}$. These observations can be explained by the fact that only one projection plane was used to assess the 3D motion.

3.2 Evaluation on real sequences

These experiments were designed to estimate both the applicability of our method to clinical datasets, and the ability of our method to learn parameters $\boldsymbol{\Sigma}_v$ and $\boldsymbol{v}_\mathrm{e}$ as presented in Section 2.5. For this purposes, we used a leave-one-out approach, registering six out of seven sequences using both planes, with $\alpha = 0$, $N_1 = 600$

and $N_2 = 1200$ particles. These registrations were used to estimate $\boldsymbol{\Sigma}_v$ and $\boldsymbol{v}_e$, which were then injected in our model (with $N_1 = 60$ and $N_2 = 150$) for the single-plane registration of the remaining sequence, with learned template ($\alpha = 0.8$) and without ($\alpha = 0$). The mean TRE 2D were respectively 2.75 ± 1.15 mm and 3.61 ± 2.10 mm (24% improvement in accuracy, 46% in deviation), for a mean computation time of 5.2 seconds per image using Matlab.

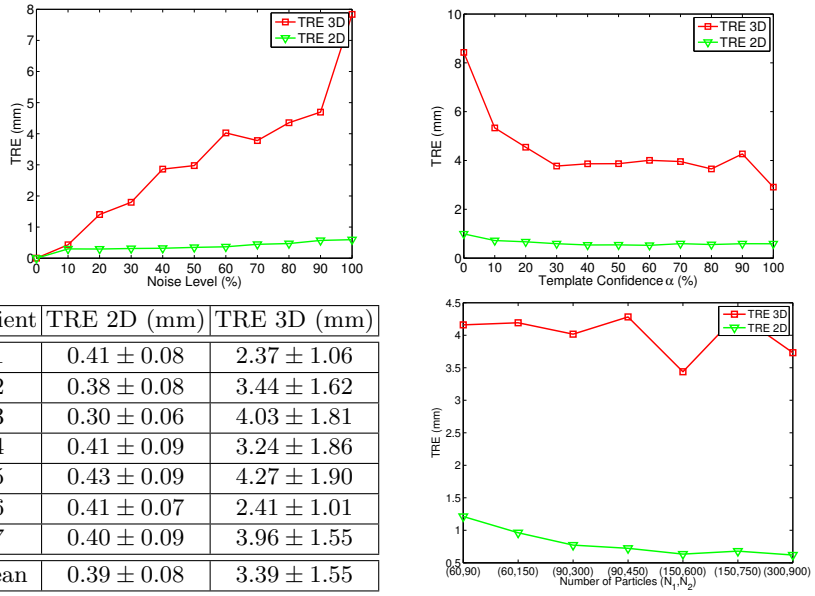

Patient	TRE 2D (mm)	TRE 3D (mm)
1	0.41 ± 0.08	2.37 ± 1.06
2	0.38 ± 0.08	3.44 ± 1.62
3	0.30 ± 0.06	4.03 ± 1.81
4	0.41 ± 0.09	3.24 ± 1.86
5	0.43 ± 0.09	4.27 ± 1.90
6	0.41 ± 0.07	2.41 ± 1.01
7	0.40 ± 0.09	3.96 ± 1.55
Mean	0.39 ± 0.08	3.39 ± 1.55

Fig. 2. TRE 3D (red squares) and 2D (green triangles) for simulations under the influence of: increasing noise level (top left, with $N_1 = 150$, $N_2 = 600$, $\alpha = 80\%$); increasing template confidence (top right, with $N_1 = 90$, $N_2 = 300$, 60% noise); coronary topology (bottom left, mean $\pm$ std, with $N_1 = 150$, $N_2 = 600$, $\alpha = 80\%$, 30% noise); and increasing number of particles (bottom-right, with $\alpha = 50\%$, 70% noise).

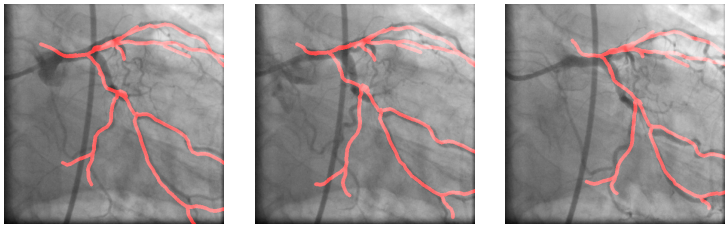

Fig. 3. Motion-compensated CT-centerline overlaid on top of angiographic images at different cardiac phases; end-diastole, early systole and systolic peak (from left to right).

4 Conclusion

This study presented a novel 3D stochastic motion compensation technique from single-plane angiograms based on a temporal model of 3D coronary centerline deformations over the cardiac cycle. Despite its generic formulation, this approach allows one to learn patient-specific knowledge on coronary motion, and use this knowledge to improve the accuracy and robustness of the registration with notably low computation times. Motion can be learned and generalized from a large database of patients undergoing coronary intervention, or measured using X-ray rotational angiography before the beginning of the intervention. The latter case could be further investigated in future works, to provide a more personalized model of coronary motion.

References

1. Duong, L., Liao, R., Sundar, H., Tailhades, B., Meyer, A., Xu, C.: Curve-based 2D-3D registration of coronary vessels for image guided procedure. In: Society of Photo-Optical Instrumentation Engineers (SPIE). Volume 7261. (2009)
2. Markelj, P., Tomazevic, D., Likar, B., Pernus, F.: A review of 3D/2D registration methods for image-guided interventions. Medical Image Analysis (2010)
3. McLaughlin, R., Hipwell, J., Hawkes, D., Noble, J., Byrne, J., Cox, T.: A comparison of 2D-3D intensity-based registration and feature-based registration for neurointerventions. In: MICCAI 2002. Volume 2489 of LNCS. Springer
4. Tang, T., Ellis, R.: 2D/3D deformable registration using a hybrid atlas. In: MICCAI 2005. Volume 3750 of LNCS. Springer 223–230
5. Chen, X., Graham, J., Hutchinson, C., Muir, L.: Inferring 3D kinematics of carpal bones from single view fluoroscopic sequences. In: MICCAI 2011. Volume 6892 of LNCS., Springer 680–687
6. Groher, M., Zikic, D., Navab, N.: Deformable 2D-3D registration of vascular structures in a one view scenario. Medical Imaging, IEEE Trans. on **28**(6) (2009)
7. Blondel, C., Malandain, G., Vaillant, R., Devernay, F., Coste-Manire, E., Ayache, N.: 4D tomographic representation of coronary arteries from one rotational X-ray sequence. In: MICCAI 2003. (2003) 223–230
8. Shechter, G., Devernay, F., Coste-Maniere, E., Quyyumi, A., McVeigh, E.: Three-dimensional motion tracking of coronary arteries in biplane cineangiograms. Medical Imaging, IEEE Trans. on **22**(4) (2003) 493 –503
9. Rohkohl, C., Lauritsch, G., Prümmer, M., Hornegger, J.: Interventional 4-D Motion Estimation and Reconstruction of Cardiac Vasculature without Motion Periodicity Assumption. In: MICCAI 2009. Volume 5761 of LNCS. 132–139
10. Florin, C., Williams, J., Khamene, A., Paragios, N.: Registration of 3D angiographic and X-ray images using sequential monte carlo sampling. In: Computer Vision for Biomedical Image Applications. Volume 3765 of LNCS. Springer (2005) 427–436
11. Ruijters, D., ter Haar Romeny, B., Suetens, P.: Vesselness-based 2D-3D registration of the coronary arteries. Int. Journal of Computer Assisted Radiology and Surgery **4**(4) (2009) 391–397
12. Sidenbladh, H., Black, M., Fleet, D.: Stochastic tracking of 3D human figures using 2D image motion. In: ECCV 2000. Volume 1843 of LNCS. Springer 702–718

13. Declerck, J., Feldmar, J., Ayache, N.: Definition of a four-dimensional continuous planispheric transformation for the tracking and the analysis of left-ventricle motion. Medical Image Analysis **2**(2) (1998) 197–213
14. Frangi, A., Niessen, W., Vincken, K., Viergever, M.: Multiscale vessel enhancement filtering. In: MICCAI 1998. Volume 1496 of LNCS. Springer 130–137
15. Djuric, P., Kotecha, J., Zhang, J., Huang, Y., Ghirmai, T., Bugallo, M., Miguez, J.: Particle filtering. Signal Processing Magazine, IEEE **20**(5) (Sep 2003) 19–38