

Nonrigid Registration of Myocardial Perfusion MRI Using Pseudo Ground Truth

Chao Li and Ying Sun

Department of Electrical and Computer Engineering,
National University of Singapore, Singapore 117576, Republic of Singapore
{chaoli, elesuny}@nus.edu.sg

Abstract. In this paper we present a method for nonrigid registration of myocardial perfusion MR images. Instead of registering pairs of images within the observed sequence, we register the observed sequence to a pseudo ground truth, which is a motion/noise-free sequence estimated from the observed one. As the corresponding images of the two sequences have almost identical intensity distributions, our method overcomes the challenges arising from rapidly varying image intensity and contrast. The pseudo ground truth and the deformation fields for the observed sequence are obtained simultaneously by minimizing an energy functional integrating both the registration error and the spatiotemporal constraints on the pseudo ground truth in an expectation-maximization fashion. We have tested the proposed nonrigid registration method on real cardiac MR perfusion scans, both qualitatively and quantitatively. Experimental results show that the proposed method is able to successfully compensate for the heart motion during contrast enhancement.

1 Introduction

Myocardial perfusion magnetic resonance imaging (MRI) has demonstrated great potential for diagnosing cardiovascular diseases. In a myocardial perfusion MRI study, the heart is scanned along short-axis slices repeatedly at the same phase of the cardiac cycle through electrocardiographic gating, following a bolus injection of a contrast agent. Patient breathing during image acquisition often causes large variations in the position of the left ventricle (LV) in different frames. At the same time, the shape of the heart may change during contrast enhancement due to variations in heart rate. Therefore, it is desirable to perform nonrigid registration on time-series perfusion images to account for local elastic deformation.

The challenges in perfusion image registration mainly arise from rapid intensity changes of the heart ventricles during the wash-in/wash-out of the contrast agent. Mutual information (MI) [1] and normalized gradient fields [2,3] have been used to account for time-varying intensity. However, these methods are either computationally expensive or pure gradient based.

To minimize the intensity difference between the reference and floating images, some approaches achieve registration of an image sequence by registering every two consecutive frames [4,2]. However, this serial registration scheme tends to




Fig. 1. Initial alignment of a myocardial perfusion MR image sequence

accumulate registration errors. In [3], Wollny et al. proposed to select the reference frame according to the breathing periodicity, but this method still requires registration of images from different perfusion phases. To avoid this problem, changes in image intensity were modeled by a Cluster-aware Active Appearance Model built from an annotated training set [5]. In another approach, Independent Component Analysis was used to generate a time-varying reference image for motion compensation [6]. However, in the above two approaches, it is not easy to choose the appropriate number of clusters or independent components.

In this paper, we propose a novel registration method that overcomes the challenges arising from time-varying intensity by using pseudo ground truth. The pseudo ground truth is an estimate of the image sequence that would have been acquired without being affected by motion or noise. We design an energy function that integrates both nonrigid registration and pseudo ground truth estimation, which can be minimized iteratively by solving a system of linear equations and applying nonrigid registration to corresponding pairs of images between the observed sequence and the pseudo ground truth sequence. In contrast to pairs of images within the observed sequence, these pairs of images have similar intensities, and the registration problem is therefore greatly simplified.

The rest of the paper is organized as follows. Section 2 describes our nonrigid registration algorithm. Section 3 presents the experimental results on real patient cardiac perfusion MRI data sets, followed by conclusion in Section 4.

2 Method

We propose to solve the nonrigid registration problem in three steps: 1) automatic selection of one reference frame and detection of a region of interest (ROI) that contains the LV; 2) initial alignment by identifying the global translation of the ROI in each frame; and 3) estimation of local elastic deformation. For the first two steps, we adopt the methods described in [7] and [8], respectively. In the following sections, we focus on the last step, i.e., nonrigid registration.

Fig. 1 displays the results after initial alignment for 5 selected frames from a myocardial perfusion MRI study, in which the bounding box of the ROI has been shifted to the best match location in each frame. By aligning the ROI from different frames, we obtain a roughly registered ROI sequence that still contains residual motion incurred by the local elastic deformation of the LV.

2.1 Energy Function for Nonrigid Registration

Instead of registering pairs of images within the image sequence, we propose to register each image to its counterpart in a pseudo ground truth image sequence that is assumed to be free of motion. This avoids the problem caused by unknown intensity variations between different frames, because registration is between images with almost identical intensity distributions.

Given an observed image sequence g , we solve for a pseudo ground truth sequence f and a nonrigid deformation function H that minimize the following energy functional:

$$E(H(g), f) = E_d(H(g), f) + \alpha E_s(f) + \beta E_t(f), \quad (1)$$

subject to the constraint on H that its underlying displacement field is zero for the reference frame, and is smooth for the rest frames. In (1), E_d is the data fidelity term that measures the difference between the pseudo ground truth sequence f and the nonrigidly deformed sequence $H(g)$; E_s is the spatial smoothness constraint penalizing the intensity difference between neighboring pixels of the same tissue type; and E_t is the temporal smoothness constraint penalizing the second order derivative of the intensity-time curve of each pixel; α and β are positive scalars that control the weights of different terms.

To rewrite the energy functional in matrix-vector form, each image sequence is represented as a column vector, e.g., if $g(i, j, t)$ is the intensity at pixel (i, j) in MRI frame t , then the column vector \mathbf{g} is given by:

$$\mathbf{g} = \text{vec}(g) = [g(1, 1, 1) \dots g(N_i, 1, 1) \dots g(1, N_j, 1) \dots \\ g(N_i, N_j, 1) \dots g(1, 1, N_t) \dots g(N_i, N_j, N_t)]^T,$$

where N_i and N_j are respectively the number of rows and columns of each image, and N_t is the number of frames. Similarly, we define $\mathbf{f} = \text{vec}(f)$ as the column vector of the pseudo ground truth f .

Data Fidelity Term. Let $\tilde{g} = H(g)$ denote the image sequence obtained by deforming g with H . The data fidelity term E_d is the sum-squared intensity difference between \tilde{g} and f :

$$E_d(H(g), f) = E_d(\tilde{\mathbf{g}}, \mathbf{f}) = (\tilde{\mathbf{g}} - \mathbf{f})^T (\tilde{\mathbf{g}} - \mathbf{f}), \quad \tilde{\mathbf{g}} = \text{vec}(\tilde{g}) = \text{vec}(H(g)). \quad (2)$$

Spatial Smoothness Constraint. Motivated by the observation that pixels of the same tissue type have similar intensities in each frame of the sequence, we impose a spatial smoothness constraint on f by penalizing the sum of weighted intensity differences between each pixel and its nearest neighbors:

$$E_s(f) = E_s(\mathbf{f}) = \sum_{k=1}^K (\mathbf{D}_k^s \mathbf{f})^T \mathbf{W}_k (\mathbf{D}_k^s \mathbf{f}), \quad (3)$$

where K is the number of neighboring pixels being considered, and in this work we use $K = 4$; \mathbf{D}_k^s is the first order spatial derivative operator along the direction

between each pixel and its k th neighbor, and \mathbf{W}_k is the corresponding weight matrix derived from the intensity-time curves, such that intensity differences between neighboring pixels of different tissue types are not penalized.

Let $\mathbf{D}_f(N)$ and $\mathbf{D}_b(N)$ respectively denote the first order forward and backward derivative operators of size $N \times N$. The four derivative operators are then respectively given by $\mathbf{D}_1^s = \mathbf{I}_{N_t} \otimes \mathbf{I}_{N_j} \otimes \mathbf{D}_f(N_i)$, $\mathbf{D}_2^s = \mathbf{I}_{N_t} \otimes \mathbf{I}_{N_j} \otimes \mathbf{D}_b(N_i)$, $\mathbf{D}_3^s = \mathbf{I}_{N_t} \otimes \mathbf{D}_f(N_j) \otimes \mathbf{I}_{N_i}$, and $\mathbf{D}_4^s = \mathbf{I}_{N_t} \otimes \mathbf{D}_b(N_j) \otimes \mathbf{I}_{N_i}$, where \otimes represents the Kronecker operator and \mathbf{I}_N denotes the identity matrix of size $N \times N$.

Another observation is that pixels of the same tissue type have intensity-time curves that exhibit similar temporal dynamics. Therefore, we use the correlation coefficients between intensity-time curves of neighboring pixels to set the weight matrix \mathbf{W}_k . Here the intensity-time curves are obtained from the deformed image sequence \tilde{g} . Let ρ_{ijk} denote the correlation coefficient between the intensity vectors of pixel (i, j) and its k th neighbor, thus $\mathbf{W}_k = \text{diag}(\text{vec}(w_k))$ and

$$w_k(i, j, t) = \begin{cases} \rho_{ijk}, & \text{if } \rho_{ijk} > \rho_{th} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $\rho_{th} \in [0, 1)$ is a user-defined threshold to ensure that only intensity differences between neighboring pixels of the same tissue type are penalized.

Temporal Smoothness Constraint. As the contrast agent gradually perfuses through the heart, generally the intensity increase/decrease step-size does not vary much during the same perfusion phase. Therefore we impose a temporal smoothness constraint on f by penalizing the second order time derivative of the pseudo ground truth:

$$E_t(f) = E_t(\mathbf{f}) = (\mathbf{D}_2^t \mathbf{f})^T (\mathbf{D}_2^t \mathbf{f}). \quad (5)$$

Let $\mathbf{D}_{tt}(N)$ denote the second order derivative operator of size $N \times N$. The second order time derivative operator is given by $\mathbf{D}_2^t = \mathbf{D}_{tt}(N_t) \otimes \mathbf{I}_{N_j} \otimes \mathbf{I}_{N_i}$.

2.2 Energy Minimization

In our implementation, the energy functional defined in (1) is minimized by iteratively solving for the optimal H and \mathbf{f} , in an expectation-maximization fashion. In each iteration, first we keep H fixed, and estimate the pseudo ground truth \mathbf{f} by minimizing

$$E = (\tilde{\mathbf{g}} - \mathbf{f})^T (\tilde{\mathbf{g}} - \mathbf{f}) + \alpha \sum_{k=1}^K (\mathbf{D}_k^s \mathbf{f})^T \mathbf{W}_k (\mathbf{D}_k^s \mathbf{f}) + \beta (\mathbf{D}_2^t \mathbf{f})^T (\mathbf{D}_2^t \mathbf{f}). \quad (6)$$

This minimization requires solving a system of linear equations:

$$\left[\mathbf{I} + \alpha \sum_{k=1}^K \left(\mathbf{D}_k^s{}^T \mathbf{W}_k \mathbf{D}_k^s \right) + \beta \mathbf{D}_2^t{}^T \mathbf{D}_2^t \right] \mathbf{f} = \tilde{\mathbf{g}},$$

which can be solved by using Gaussian elimination and results in an estimate $\hat{\mathbf{f}}$.

Then we perform registration between the reference image (static image) in g and its counterpart image (moving image) in \hat{f} and the resulting deformation field is then applied to all the images in \hat{f} , which leads to a calibrated pseudo ground truth image sequence f so as to ensure that the underlying deformation field for the reference frame is zero. Next, we keep \mathbf{f} fixed, and register each image in g to its counterpart. The resultant deformation fields completely and uniquely define the deformation function H , and hence $\tilde{\mathbf{g}}$ can be updated accordingly. Here we use the demons algorithm described in [9] to register corresponding images. Nevertheless, this pseudo ground truth based approach can be used in conjunction with many other registration methods.

The initial condition is $\tilde{\mathbf{g}} = \mathbf{g}$, i.e., the original observed image sequence, which implies the initial deformation function H is the identity. The weight matrices $\mathbf{W}_k, k = 1, \dots, K$ are re-estimated using the updated $\tilde{\mathbf{g}}$ at each iteration. The iteration continues until E cannot be further reduced.

3 Experimental Results

We have tested the proposed method on 9 slices of perfusion images from 7 patients. The images were acquired on Siemens Sonata MR scanners following bolus injection of Gd-DTPA contrast agent. For all the data sets in our study, we consistently observed monotonic decrease of the energy function, which had little change after 3 or 4 iterations. Based on visual validation from video, the heart in the registered sequence was well stabilized. Next, we present the experimental results for a representative data set in two aspects: contour propagation and comparison of intensity-time curves.

3.1 Contour Propagation

Given the contour that delineates the boundary of the LV blood pool in the reference frame, one can propagate this contour to other frames according to the deformation field obtained using the method described in Section 2. Therefore, one way to evaluate the performance of the nonrigid registration algorithm is to verify whether the propagated contour well delineates the boundaries of the LV blood pool in other frames.

The first two rows in Fig. 2 compares the propagated contours before and after applying nonrigid registration for 4 consecutive frames from a real cardiac MR perfusion scan. These frames are chosen because the LV in these images undergoes noticeable shape changes. As shown in the top row, the contours before nonrigid registration do not lie exactly at the boundaries between the LV blood pool and the LV myocardium, especially at the regions that are indicated by the bright arrows. In contrast, the contours in the second row, after nonrigid registration using the proposed method, delineate well the boundaries between the LV blood pool and the LV myocardium.

To further demonstrate the performance of the proposed method, we compare our results with those obtained by serial registration using the demons algorithm [9], in which every two consecutive frames are registered to propagate




Fig. 2. Contour propagation for 4 consecutive frames from a real patient cardiac MR perfusion scan: contours before applying nonrigid registration (top row); contours propagated by our method (second row), serial demons registration (third row), and MI-based registration (bottom row)

from the reference frame to the rest of the sequence. As shown in the third row in Fig. 2, the propagated contours in all four frames are away from the true boundaries in the bottom right region due to the accumulation of registration errors. We have also applied the free-form deformation (FFD) registration method [10] using MI as the similarity measure, in which the reference frame is used to register all other frames in the sequence. As shown in the bottom row in Fig. 2, the propagated contours do not delineate the boundaries as accurately as those obtained by our method (see the second row in Fig. 2).

For quantitative evaluation, we manually drew the contours that delineate the LV blood pool in each frame for one slice, and measured the maximum and mean distances between the propagated contours and the corresponding manually-drawn contours. For the four frames shown in Fig. 2, the maximum distances were 2.9, 2.5, 1.7, and 1.7 pixels with only translation, as compared to 0.8, 0.9, 0.7 and 1.1 pixel(s) after applying our registration method. Over the entire sequence, the average mean distance was decreased from 0.70 pixel to 0.58




Fig. 3. Comparison of the average intensity-time curves of four neighboring pixels from the myocardium for the original (raw) data set, after compensating for global translation, and after performing nonrigid registration

pixel. In contrast, using the serial demons registration and the MI-based method, the average mean distance was increased to 1.40 and 1.95 pixels, respectively, due to accumulation of registration errors and large registration errors in frames having distinct intensity distributions from that of the reference frame.

3.2 Comparison of Intensity-Time Curves

In cardiac perfusion MRI, changes in pixel intensity corresponding to the same tissue across the image sequence are closely related to the concentration of the contrast agent. As the concentration of the contrast agent changes smoothly during the perfusion process, the intensity-time curves should be temporally smooth at pixels that are located on the LV myocardium.

We use a 2×2 pixel window to select a region in the LV myocardium, and the window is chosen to be close to the LV blood pool such that the motion of the LV can be reflected in the intensity-time curve of the window. Fig. 3 plots the average intensity-time curves within the window, for the original/raw data set (green triangle), after compensating for global translation (blue cross), and after performing nonrigid registration (magenta circle). For the original image sequence, the intensity-time curve exhibits very big oscillations in the second half of the sequence, in which the global translation is large due to patient breathing. After compensating for the global translation, the intensity-time curve becomes temporally smoother but still has some oscillations, e.g., in frames 37 – 40 and 52 – 58. This is because global translation is not capable of describing shape changes in these frames. However, after compensating for the local deformation by performing nonrigid registration, the intensity-time curve becomes smoother at frames when the LV undergoes noticeable local deformation, with only small local oscillations due to noise.

4 Conclusion

This paper presents a novel nonrigid registration algorithm for cardiac perfusion MR images. Unlike most registration methods that estimate the deformation

between pairs of images within the observed perfusion sequence, we introduce a pseudo ground truth to facilitate image registration. The pseudo ground truth is a motion-free sequence estimated from the observed perfusion data. Since the intensity distributions of the corresponding images between the pseudo ground truth sequence and the observed sequence are almost identical, this method successfully overcomes the challenges arising from intensity variations during perfusion. Our experimental results on real patient data have shown that our method is able to effectively compensate for the elastic deformation of the heart, and that it significantly outperforms the serial demons registration method and an MI-based method when registering myocardial perfusion images.

Acknowledgments. The authors thank Siemens Corporate Research, NJ, USA, for providing the datasets, and acknowledge the support by NUS grant R-263-000-470-112.

References

1. Ólafsdóttir, H.: Nonrigid registration of myocardial perfusion MRI. In: Proc. Svenska Symposium i Bildanalys (March 2005)
2. Wollny, G., Ledesma-Carbayo, M.J., Kellman, P., Santos, A.: A new similarity measure for non-rigid breathing motion compensation of myocardial perfusion MRI. In: Proc. IEEE EMBS, August 2008, pp. 3389–3392 (2008)
3. Wollny, G., Ledesma-Carbayo, M.J., Kellman, P., Santos, A.: Non-rigid motion compensation in free-breathing myocardial perfusion magnetic resonance imaging. In: Proc. Computers in Cardiology, September 2008, pp. 465–468 (2008)
4. Breeuwer, M., Spreuwers, L., Quist, M.: Automatic quantitative analysis of cardiac MR perfusion images. In: Proc. SPIE, February 2001, pp. 733–742 (2001)
5. Stegmann, M.B., Larsson, H.B.W.: Motion-compensation of cardiac perfusion MRI using a statistical texture ensemble. In: Magnin, I.E., Montagnat, J., Clarysse, P., Nenonen, J., Katila, T. (eds.) FIMH 2003. LNCS, vol. 2674, pp. 151–161. Springer, Heidelberg (2003)
6. Milles, J., van der Geest, R.J., Jerosch-Herold, M., Reiber, J., Lelieveldt, B.: Fully automated motion correction in first-pass myocardial perfusion MR image sequences. *IEEE Trans. Medical Imaging* 27(11), 1611–1621 (2008)
7. Lorenz, C.H., Sun, Y., Meng, T., Luk, C., Kraitchman, D.L., Kirchberg, K.J.: Generation of MR myocardial perfusion maps without user interaction. In: Proc. Joint Annual Meeting ISMRM-ESMRMB 2007, Berlin, Germany (May 2007)
8. Sun, Y., Jolly, M.P., Moura, J.M.F.: Contrast-invariant registration of cardiac and renal MR perfusion images. In: Proc. 7th International Conference on Medical Image Computing and Computer Assisted Intervention, France (September 2004)
9. Wang, H., Dong, L., O’Daniel, J., Mohan, R., Garden, A.S., Ang, K.K., Kuban, D.A., Bonnen, M., Chang, J.Y., Cheung, R.: Validation of an accelerated ‘demons’ algorithm for deformable image registration in radiation therapy. *Phys. Med. Biol.* 50(12), 2887–2905 (2005)
10. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Medical Imaging* 18(8), 712–721 (1999)