

PHASE-BASED NON-RIGID REGISTRATION OF MYOCARDIAL PERFUSION MRI IMAGE SEQUENCES

Lennart Tautz¹, Anja Hennemuth¹, Mats Andersson², Achim Seeger³, Hans Knutsson², Ola Friman¹

¹Fraunhofer MEVIS Institute for Medical Image Computing, Bremen, Germany

²Linköping University, Linköping, Sweden

³Eberhard Karls University, Tübingen, Germany

ABSTRACT

The condition of the heart muscle tissue can be inferred by analyzing the time-intensity curves obtained with myocardial perfusion MRI. Specifically, identifying tissue that is under-supplied with blood is important when choosing a suitable therapy for patients with coronary heart disease. Before an analysis can be carried out, the images must be registered to compensate for cardiac and respiratory motion. This is a difficult problem, as the motion is non-rigid and because the image contrast varies strongly over time due to the injection of a contrast agent into the blood stream. To address these problems, an automatic non-rigid registration approach is presented that utilizes local phase instead of intensity or gradient information.

Index Terms— perfusion, registration, local phase

1. INTRODUCTION

In a medical context, the term *perfusion* means the through-flow of blood in biological tissue. When an arterial blood vessel is partially or completely blocked, the supported tissue becomes underperfused and may die from a lack of oxygen. The situation is especially critical when a coronary artery is blocked, as this may cause a heart infarction and lead to sudden death. Myocardial perfusion MRI is an imaging technique with which the blood supply to the heart muscle, the myocardium, can be assessed. Such an examination is performed to identify underperfused muscle tissue that is still alive but at risk of dying. This tissue, and the patient, may be saved through revascularization, i.e., an intervention which removes or bypasses the blockage in the vessel. The myocardial perfusion is measured by injecting a bolus of contrast agent into the blood stream followed by a continuous MRI acquisition during the contrast agent's first pass through the heart. Typically, around 40 time frames are acquired, see Fig. 1 for example images. The local delivery of blood to the myocardium can be quantified by examining the time-intensity curve in each voxel, see Fig. 3. Healthy muscle tissue exhibits a characteristic perfusion curve, whereas underperfused and dead tissue have abnormal curves, i.e., the tissue state can be inferred based on the perfusion curves. However, before such an analysis can be carried out, the images must




Fig. 1. Example cardiac perfusion MRI images at three different time points in the cardiac cycle. Note the varying contrast as the contrast agent flows through the heart.

be registered to compensate for cardiac and respiratory motion to insure that a voxel represents the same tissue location throughout the image sequence. The image deformations are non-rigid, and the registration problem is further complicated by the varying image intensity caused by the contrast agent, see Fig. 1. Moreover, to attain the necessary temporal resolution, spatial coverage of the heart must be sacrificed so that only a few sparsely distributed 2D slices are acquired.

Image registration has long been an active research area, and a comprehensive review of proposed methods is outside the scope of this work. The special problem of aligning myocardial perfusion MRI sequences has also been addressed before, overviews can be found in [1, 2, 3]. As noted by Milles et al. [2], the proposed methods are limited when dealing with the contrast variations in the image sequences. Clearly, an image similarity measure based on voxel differences or correlation is suboptimal in the presence of contrast variations. Alternative features for registration include, for example, landmarks [1] and mutual information [3]. As a basis for registration in this work, we instead propose to use local phase, which represents image features such as edges and lines but is invariant to their magnitude. Seminal work on phase-based displacement estimation was presented by Fleet and Jepson [4], and more recent work on this subject was presented by Mellor and Brady [5]. The algorithm presented in this work also draws on the Demons algorithm by Thirion [6]. A similar algorithm, known as the Morphon [7], has previously been applied in other contexts.

2. PHASE-BASED NON-RIGID REGISTRATION

Phase-based image registration is directly related to the Fourier Shift Theorem, which states that the Fourier transforms of a signal $f(x)$ and a globally shifted signal $f(x-d)$ are related via a phase factor $\mathcal{F}\{f(x-d)\} = e^{-jd\omega}\mathcal{F}\{f(x)\}$. That is, for two signals $f_1(x) = f(x)$ and $f_2(x) = f(x-d)$, we have $d \propto \arg\left(\mathcal{F}\{f_1(x)\}\overline{\mathcal{F}\{f_2(x)\}}\right)$, where the $\overline{}$ denotes the complex conjugate. By using the *local* (or instantaneous) phase $\phi(x)$, derived from the complex analytical signal $f_a(x) = A(x)e^{j\phi(x)}$ of $f(x)$, the above approach can be extended to estimate non-stationary or non-rigid shifts in 1D. In practice, the analytic signal is estimated by applying a quadrature filter $q(x)$, $\hat{f}_a(x) = (f * q)(x)$, which has a band-pass characteristic that determines the scale of the structures of interest [8].

One way of generalizing the analytic signal, which is inherently a 1D construct, to images of higher dimensions, is to apply a set of quadrature filters $q^{(i)}(\mathbf{x})$ with different orientations $\hat{\mathbf{n}}_i$. The generalized analytic signal in direction $\hat{\mathbf{n}}_i$ for an image $I(\mathbf{x})$ is then obtained as $I_a^{(i)}(\mathbf{x}) = (I * q^{(i)})(\mathbf{x})$. Assume now that we have a deformed image $J(\mathbf{x}) = I(\mathbf{x} + \mathbf{d}(\mathbf{x}))$, where $\mathbf{d}(\mathbf{x})$ is an unknown dense deformation field that we wish to estimate. Following the 1D procedure above, the displacement $\hat{d}_i(\mathbf{x})$ along the orientation $\hat{\mathbf{n}}_i$ can be estimated by the local phase difference of the complex product

$$p_{IJ}^{(i)}(\mathbf{x}) = I_a^{(i)}(\mathbf{x})\overline{J_a^{(i)}(\mathbf{x})}. \quad (1)$$

Specifically, $\hat{d}_i(\mathbf{x}) \propto \arg\left(p_{IJ}^{(i)}(\mathbf{x})\right)$. For each oriented quadrature filter $q^{(i)}(\mathbf{x})$, a dense displacement estimate is obtained. A confidence measure $c_i(\mathbf{x})$ for the displacement estimate in each filter direction, as well as a combined confidence measure $c(\mathbf{x})$, can be formulated as

$$c_i(\mathbf{x}) = \sqrt{|p_{IJ}^{(i)}(\mathbf{x})|} \left[1 + \cos\left(\arg\left(p_{IJ}^{(i)}(\mathbf{x})\right)\right)\right] \quad (2)$$

and

$$c(\mathbf{x}) = \sum_i c_i(\mathbf{x}). \quad (3)$$

The rationale behind this confidence measure is that the magnitude of $p_{IJ}^{(i)}(\mathbf{x})$ is large when there is a strong structure seen by the filter $q^{(i)}(\mathbf{x})$ in both images, indicating a higher confidence in the estimate. If the phase of $p_{IJ}^{(i)}(\mathbf{x})$, i.e., the phase difference between $I_a^{(i)}(\mathbf{x})$ and $J_a^{(i)}(\mathbf{x})$, is large, the quadrature filter has likely locked on different structures and we should have less confidence in such an estimate. Taking the confidence measures into account, a first estimate of the complete deformation field is the weighted sum

$$\mathbf{d}(\mathbf{x}) = \frac{\sum_i c_i(\mathbf{x})\hat{d}_i(\mathbf{x})\hat{\mathbf{n}}_i}{\sum_i c_i(\mathbf{x})}. \quad (4)$$

In the medical imaging context, the deformation field $\mathbf{d}(\mathbf{x})$ is generally smooth, and a spatial regularization should

be applied to reflect this prior knowledge. In this work, we adopt a non-parametric approach and smooth with a Gaussian kernel $g(\mathbf{x}; \sigma^2)$ to the desired smoothness [6]. The confidence measure in Eq. 2 should also be taken into account in the regularization process, and a so-called normalized averaging is used to produce a regularized deformation field

$$\mathbf{d}_{reg}(\mathbf{x}) = \frac{[\mathbf{d}(\mathbf{x})c(\mathbf{x})] * g(\mathbf{x}; \sigma^2)}{c(\mathbf{x}) * g(\mathbf{x}; \sigma^2)}, \quad (5)$$

where the division is taken voxel-wise. The result of Eq. 5 is a weighted average displacement vector in each voxel, the weight being the product of the confidence and the Gaussian kernel.

Finally, to estimate large non-rigid deformations, the displacement estimation outlined above must be implemented in a scale space, and it may also be necessary to iterate the estimation several times on each scale to refine the estimation. To avoid repeated resampling and interpolation of the images $I(\mathbf{x})$ and $J(\mathbf{x})$, the deformation estimates are instead accumulated in $\mathbf{d}_{tot}(\mathbf{x})$ as

$$\mathbf{d}_{tot}(\mathbf{x}) \leftarrow \mathbf{d}_{tot}(\mathbf{x}) + \frac{c(\mathbf{x})}{c_{tot}(\mathbf{x}) + c(\mathbf{x})} \mathbf{d}_{reg}(\mathbf{x}), \quad (6)$$

where $c_{tot}(\mathbf{x})$ is an accumulated confidence measure that is updated for each iteration as

$$c_{tot}(\mathbf{x}) \leftarrow \frac{c_{tot}^2(\mathbf{x}) + c^2(\mathbf{x})}{c_{tot}(\mathbf{x}) + c(\mathbf{x})}. \quad (7)$$

After convergence, $\mathbf{d}_{tot}(\mathbf{x})$ is the final estimate of the true deformation field $\mathbf{d}(\mathbf{x})$.

3. METHOD

This section provides implementation-specific details of the phase-based non-rigid registration of perfusion MRI sequences. Due to the sparse 3D information, the registration is performed slice-wise (2D). High quality quadrature filters are important for estimating the analytic signal and the ensuing 1D displacements. In this work, log-normal quadrature filters are used, which in the Fourier domain are expressed as polar separable functions:

$$Q_i(\mathbf{u}) = R(\|\mathbf{u}\|)D_i(\hat{\mathbf{u}}) \quad (8)$$

with

$$R(\|\mathbf{u}\|) = e^{C \ln^2(\|\mathbf{u}\|/u_0)}, \quad C = -\frac{4}{B^2 \ln(2)}. \quad (9)$$

$R(\|\mathbf{u}\|)$ is a Gaussian function on a logarithmic scale (hence the name log-normal) and gives the filters a bandpass characteristic. u_0 is the center frequency and B is the width of the passband in octaves. We use two different u_0 to capture different scales, and the bandwidth is fixed to $B = 2$ octaves. $D_i(\hat{\mathbf{u}})$ in Eq. 8 gives the filters a direction and the quadrature

property by setting one half of the Fourier domain to zero as follows:

$$D_i(\hat{\mathbf{u}}) = \begin{cases} (\hat{\mathbf{u}}^T \hat{\mathbf{n}}_i)^2 & \text{if } \hat{\mathbf{u}}^T \hat{\mathbf{n}}_i > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

where $\hat{\mathbf{n}}_i$ is the filter direction. We use four 2D directions $\hat{\mathbf{n}}_i = [\cos(\phi_i), \sin(\phi_i)]^T$ with $\phi_i \in \{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$. A filter optimization procedure was applied to obtain finite filter kernels with good spatial localization and a frequency response that closely match the ideal shape given by Eq. 8.

A coarse-to-fine dyadic scale-space approach over 3 octaves is adopted to capture large deformations and to increase efficiency. In addition, on each scale, two different quadrature sets with center frequencies u_0 of $\frac{\pi}{4}$ and $\frac{\pi}{2\sqrt{2}}$, respectively, are used, so that, effectively, structures over six different scales are considered in total. 3 registration iterations (cf. Eq. 6) are carried out on each scale. The deformation field regularization in Eq. 5 is performed with a Gaussian kernel with $\sigma = 3$ pixels.

To align a perfusion MRI sequence, the registration must be repeated for successive images in the temporal sequence. To this end, the middle slice $I_k(\mathbf{x})$ of the sequence is taken as reference image. $I_{k-1}(\mathbf{x})$ (and $I_{k+1}(\mathbf{x})$) is then registered to $I_k(\mathbf{x})$ yielding the deformation field $\mathbf{d}_{k-1}(\mathbf{x})$. To avoid unnecessary resampling and loss of detail in the images, $I_{k-2}(\mathbf{x})$ is then registered to $I_{k-1}(\mathbf{x})$, yielding an intermediate deformation field $\tilde{\mathbf{d}}_{k-2}(\mathbf{x})$. The final deformation field that aligns $I_{k-2}(\mathbf{x})$ with the reference $I_k(\mathbf{x})$ is then obtained as $\mathbf{d}_{k-2}(\mathbf{x}) = \tilde{\mathbf{d}}_{k-2}(\mathbf{x}) + \mathbf{d}_{k-1}(\mathbf{x})$. This procedure is repeated until the ends of the temporal sequence are reached.

The phase-based registration algorithm was implemented in C++ and made available as a module in the free software package MeVisLab (<http://www.mevislab.de>) for easy incorporation in software applications and future evaluation.

4. COMPARISON METHOD

The aforementioned local phase registration method was compared with a recently published registration method for cardiac perfusion MRI [3]. In [3], an affine registration step is followed by a non-rigid registration based on B-splines and a Normalized Mutual Information (NMI) similarity measure. Additional manual steps are required in this method for robustness: a ROI around the heart must be manually defined in one slice for each data set, and outliers (large movements) should be removed before registration.

5. IMAGE DATA

Cardiac perfusion MRI data sets from 8 patients were used for validation and comparison. Each data set consisted of three 2D short-axis slices of the heart, i.e., 24 sequences in total. The perfusion images were acquired after an injection of Gadolinium contrast agent bolus and show the first pass of the contrast agent through the heart and myocardium. The

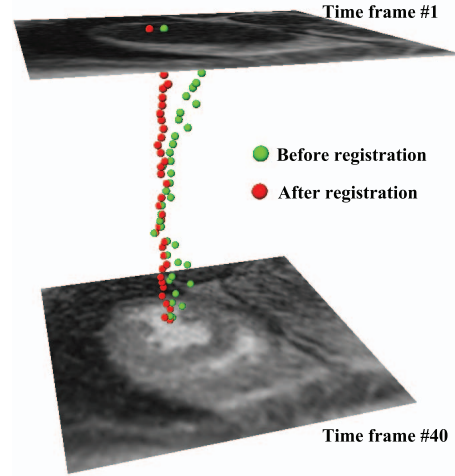


Fig. 2. To evaluate the registration, the left ventricle was manually segmented, and its center of gravity for each time frame was calculated before (green spheres) and after registration (red spheres). As a measure of motion, the standard deviation of the center of gravity is used.

image slices have 8 mm thickness, an in-plane resolution of $1.7 \text{ mm} \times 1.7 \text{ mm}$ and a sequence of 40 time frames. The data sets were acquired with parallel imaging (GRAPPA 2) with a balanced SSFP technique and ECG-triggering. The imaging parameters were: TR 2.8 s, TE 1.08 ms, FOV $340 \text{ mm} \times 300 \text{ mm}$ and matrix size $160 \text{ pixels} \times 192 \text{ pixels}$.

6. RESULTS

The local-phase based and NMI-based registration algorithms were applied to all 24 perfusion MRI sequences. For the local-phase based algorithm, no pre-processing or pre-segmentation were made - the registration is fully automatic. For the NMI-based algorithm, ROIs covering the heart were defined manually, and for three data sets (number 1, 4 and 8), a few slices had to be excluded due to large movements. To quantify the performance of the methods, the left ventricle (cf. Fig. 1) was manually delineated in all $24 \times 40 = 960$ images. The center of gravity of the left ventricle was then calculated before and after registration, see Fig. 2 for an example. As a measure of motion, the standard deviation (in pixels) of the center of gravity over the time frames in the sequences was calculated. Table 1 shows the motion quantification before and after registration for the 8 perfusion data sets (the results for the 3 image slices within each data set have been averaged for clarity). The average motion across all sequences before registration is 3.04 pixels. After the NMI registration, the average is 2.20 pixels, which is significantly better than the motion in the uncorrected images (paired t -test significance $< 10^{-4}$). Note, however, that for data set 6, the displacement is larger after the registration. With the proposed local-phase registration, the residual standard deviation averages 1.66 pixels, which in turn is significantly better than

Data set					
Method	#1	#2	#3	#4	#5
None	3.22	3.51	2.62	3.86	2.18
NMI	1.92*	2.87	2.43	1.53*	1.87
Local phase	2.22	2.04	1.48	1.46	1.46
Method	#5	#6	#7	#8	Average
None	2.18	2.99	3.34	2.61	3.04
NMI	1.87	3.18	2.21	1.59*	2.20
Local phase	1.46	1.46	1.65	1.49	1.66

Table 1. Standard deviation (in pixels) of the center of gravity of the left ventricle. An average of the 3 slices within each data set is presented. For the data sets marked with a *, outlier images were manually removed prior to the registration.

the NMI method (paired t -test significance $< 10^{-3}$). The residual displacement error is also significantly lower than in the original images for all data sets. The computational time to register two image slices was approximately 1.5 seconds with the NMI-based algorithm and 5 seconds with the phase-based algorithm.

Example perfusion curves after registration are shown in Fig. 3c. A map of the blood inflow can be calculated by estimating the slope of the curve in each voxel. Fig. 3b shows such a map before registration and Fig. 3d after registration using the proposed phase based registration. An infarcted area, denoted by the yellow arrow, can clearly be seen in this image, whereas the up-slope parameter gives no reliable information for the unregistered images in Fig. 3b.

7. DISCUSSION

Local phase is a magnitude-insensitive measure of image structure and thus a good fit for registering cardiac perfusion MRI images, in which in the image intensity is not a reliable feature. Future avenues of research include the incorporation of anatomical information in the displacement field estimation, to combine the confidence-weighted displacement field smoothing with more advanced regularization methods and criteria such as diffeomorphic transformations, and to make use of the weak, but still existent, 3D information, e.g., through a 3D displacement field regularization and removal of images with strong out-of-plane motion.

8. REFERENCES

[1] M. B. Stegmann, H. Olafsdóttir, and H. B. W Larsson, "Unsupervised motion-compensation of multi-slice cardiac perfusion MRI," *Medical Image Analysis*, vol. 9, no. 4, pp. 394–410, 2005.

[2] J. Milles, R. J. van der Geest, M. Jerosch-Herold, J. H. C. Reiber, and B. P. F. Lelieveldt, "Fully automated motion correction in first-pass myocardial perfusion MR image sequences," *IEEE Transactions on Medical Imaging*, vol. 27, no. 11, pp. 1611–1621, 2008.




Fig. 3. **a)** A cardiac perfusion MRI sequence. **b)** Up-slope parameter map of the perfusion curves *before* registration. **c)** Perfusion curves from a healthy (green) and an infarcted (red) area of the myocardium. The up-slope parameter shows the inflow rate of contrast agent. **d)** Up-slope parameter map of the perfusion curves *after* registration with the proposed method.

[3] A. Hennemuth, A. Seeger, O. Friman, S. Miller, B. Klumpp, S. Oeltze, and H.-O. Peitgen, "A comprehensive approach to the analysis of contrast enhanced cardiac MR images," *IEEE Transactions on Medical Imaging*, vol. 27, no. 11, pp. 1592–1610, 2008.

[4] D. J. Fleet and A. D. Jepson, "Computation of component image velocity from local phase information," *International Journal of Computer Vision*, vol. 5, no. 1, pp. 77–104, 1990.

[5] M. Mellor and M. Brady, "Phase mutual information as a similarity measure for registration," *Medical Image Analysis*, vol. 9, no. 4, pp. 330–343, 2005.

[6] J.-P. Thirion, "Image matching as a diffusion process: An analogy to Maxwells demons," *Medical Image Analysis*, vol. 2, no. 3, pp. 243–260, 1998.

[7] H. Knutsson and M. Andersson, "Morphons: Segmentation using elastic canvas and paint on priors," in *IEEE International Conference on Image Processing (ICIP'05)*, 2005.

[8] D. Boukerroui, J. Noble, and M. Brady, "On the choice of band-pass quadrature filters," *Journal of Mathematical Imaging and Vision*, vol. 21, no. 1, pp. 53–80, 2004.