

Unsupervised Inline Analysis of Cardiac Perfusion MRI

Hui Xue^{1,*}, Sven Zuehlsdorff², Peter Kellman³, Andrew Arai³,
Sonia Nielles-Vallespin⁴, Christophe Chefdhotel¹, Christine H. Lorenz¹,
and Jens Guehring¹

¹Imaging and Visualization, Siemens Corporate Research, Princeton, NJ, USA

²CMR R&D, Siemens Medical Solutions USA, Inc., Chicago, IL, USA

³Laboratory of Cardiac Energetics, National Institutes of Health, Bethesda, MD, USA

⁴MED MR PLM AW Cardiology, Siemens AG Healthcare Sector, Erlangen, Germany

hui-xue@siemens.com

Abstract. In this paper we first discuss the technical challenges preventing an automated analysis of cardiac perfusion MR images and subsequently present a fully unsupervised workflow to address the problems. The proposed solution consists of key-frame detection, consecutive motion compensation, surface coil inhomogeneity correction using proton density images and robust generation of pixel-wise perfusion parameter maps. The entire processing chain has been implemented on clinical MR systems to achieve unsupervised inline analysis of perfusion MRI. Validation results are reported for 260 perfusion time series, demonstrating feasibility of the approach.

1 Introduction

Myocardial first pass perfusion magnetic resonance imaging (MRI) has proven its clinical significance in the diagnosis of known and suspected ischemic heart disease, particularly in combination with cardiac delayed enhancement imaging [1]. However, the clinical routine to evaluate myocardial perfusion still relies on radionuclide imaging, such as Single Photon-Emission Computed Tomography (SPECT).

Despite many advantages over radionuclide techniques, MR perfusion imaging is still not widely used in clinical routine. Certain technical challenges prevent this technique to be added to the clinical workflow. Among them is complex cardiac motion caused by respiration, irregular heart rates, and imperfect cardiac gating. Due to these factors, a motion compensation procedure has to be applied prior to computing the myocardial signal intensity (SI) curves (Fig. 1) on a pixel-by-pixel basis. Another imperfection is B1-field inhomogeneity caused by non-uniform characteristics of the receiver coils. While qualitative visual reading is often not compromised by this effect [2], inhomogeneity can result in errors of quantitative or semi-quantitative analysis [1], which aims to estimate perfusion parameters, such as up-slope (SLOPE), area-under-curve (AUC) or myocardial blood flow.

The long and labor-intensive analysis process is also an important barrier to clinical utilization of MR stress perfusion imaging. Typically, clinical assessment requires sufficient coverage of the left ventricle (LV), while, due to the limited imaging efficiency of

* Corresponding author.

current MR scanners, a typical MR perfusion sequence acquires several two-dimensional (2D) images to approximate the coverage of LV. A minimum of three short-axis slices covering the basal, mid-ventricular, and apical portions of the LV is recommended [3]. To maintain sufficient temporal resolution during the first-pass of the contrast bolus, each slice should be imaged every one to two heart beats for a total duration of approximately 60s. As a result, a typical perfusion study can produce ~250 images. To generate signal time-intensity curves in order to estimate perfusion parameters, manual delineation of the endo- and epicardium is required due to cardiac motion. This process is time-consuming, and can be even worse in patients who are unable to hold their breath and require the data acquisition to be performed in a free breathing fashion.

To overcome these challenges, various image processing approaches are being investigated, mainly focusing on motion compensation of myocardium, including rigid-body image registration [4], ICA (independent component analysis) based correction [5], active contour [6] and active shape models [7]. However, few of these studies have yet obtained widespread usage and even fewer efforts were made so far to automate more comprehensive perfusion analysis workflows. Therefore, we propose an unsupervised perfusion analysis system, consisting of key-frame detection, consecutive motion compensation (MOCO), surface coil inhomogeneity correction (SCC) using proton density images, an Expectation-maximization (EM) algorithm, and robust parameter map generation. To maximize the clinical applicability, the proposed solution does not rely on assumptions about myocardial anatomy; therefore it can cope with different slice positions (basal, mid-ventricular and apical). Also, the processing pipeline can handle three widely used MR perfusion sequences including TurboFLASH (Turbo Fast Low Angle Shot), TrueFISP (True Fast Imaging with Steady state Precession), and GRE-EPI (Gradient Echo type Echo Planar Imaging). Finally, to provide clinicians with access to the proposed techniques, the algorithms are implemented within the Image Calculation Environment (ICE) as an inline processing chain on Siemens MR system.

2 Methods

2.1 Sequence Design

A flexible MR perfusion pulse sequence was implemented and tested on two clinical 1.5T scanners (Siemens MAGNETOM Espree and Avanto). The sequence supports different acquisition techniques, such as TurboFLASH, TrueFISP, and GRE-EPI hybrid. In order to enable a fully integrated inhomogeneity correction, the pulse




Fig. 1. Signal intensity curve and perfusion parameters. For each pixel the signal-time curve is analyzed and parameters like upslope (SLOPE), time-to-peak (TTP), peak time (PT) and area-under-curve (AUC, area under the curve between foot and peak) can be calculated. t_f and t_p are foot time and peak time.

sequence was modified to first acquire a small number (e.g. 2) of proton density (PD) weighted images before the start of the conventional first pass perfusion acquisition (Fig. 2). The slice prescription for the PD and perfusion images is identical. With the approximation that the proton density across the myocardial anatomy is constant [1], the intensity changes of PD images can be positively related to local surface coil sensitivity.




Fig. 2. (a) Example proton density image acquired before the normal perfusion acquisition. (b) Intensity profile across the heart region (yellow line). The intensity bias can be clearly observed.

2.2 Key-Frame Detection

The first step of the proposed workflow aims at detecting a *key-frame* for a perfusion series. This key-frame will be defined as the reference image, and relative motion between other phases and this reference image will be corrected. To improve the motion compensation, this key-frame should be a frame in which the myocardium has good contrast compared to the blood pool and surrounding tissues. We propose a key-frame selection approach which is based on the observation that during the contrast uptake the image intensity in regions where the contrast bolus enters will have higher standard deviation (SD) along the time dimension. As the first step, the standard deviation image for the perfusion series is computed. Although inconsistent myocardial motion can degrade the sharpness of myocardium, the contrast between myocardium and surrounding tissues in the SD image is found to be consistently noticeable. This observation holds true for the described perfusion MR pulse sequences. The next step is to select a frame having similar contrast as the SD image. For this purpose, the cross correlation ratios (CC) between every phase in the perfusion series and the SD image are computed. During the passing of contrast bolus, the CC ratio continues to increase and reaches its peak around the time point where the myocardium blood perfusion is maximized. We therefore pick the phase corresponding to the maximal CC ratio as the key-frame.

2.3 Consecutive Motion Compensation

We have found that the registration is more robust if two slices to be aligned have similar contrast. Therefore, a consecutive motion compensation strategy is developed to improve the performance of registration. As shown in Fig. 3, motion compensation starts from the key-frame and its direct neighbors (previous and next). After the first registration is finished, the next image is registered to its warped neighbor that has been transformed into the key-frame coordinate system. The complete series is corrected by consecutively performing multiple 2D-2D registrations between temporally adjacent slices. Considering the temporal resolution of perfusion studies is usually one to two heart beats, adjacent frames consistently show similar contrast, even during the first pass of contrast agents.

A fast variational non-rigid registration algorithm [8] is applied as the working-engine of perfusion motion compensation. This approach can be considered as an extension of the classic optical flow method. In this framework, a dense deformation field is estimated as the solution to a calculus of variation problem, which is solved by performing a compositional update step corresponding to a transport equation. The regularization is added by low-pass filtering the gradient images which are in turn used as velocity field to drive the transport equation. To speedup the convergence and avoid local minima, a multi-scale image pyramid is created. We selected the local cross correlation as the image similarity measure, as its explicit derivative can be more efficiently calculated than mutual information and still general enough to cope with intensity fluctuation and imaging noise between two adjacent perfusion frames.

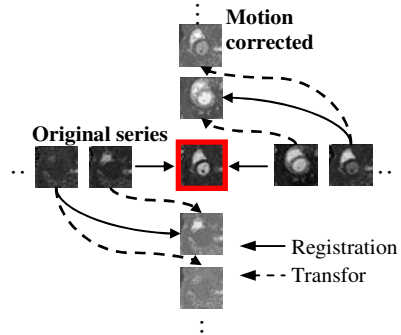


Fig. 3. An illustration of consecutive motion compensation. Motion compensation starts from the key-frame and its direct neighbors. Every image is aligned to its transformed previous neighbor. In this scheme, registration is performed between two perfusion phases with similar contrast.

2.4 Surface Coil Inhomogeneity Correction

To correct the variation in intensity due to the surface coil sensitivities, the proton density (PD) images are first registered to the key-frame. All motion compensated PD images are then averaged to improve the signal-to-noise ratio (SNR). To estimate the inhomogeneity field from the averaged PD images, we propose an algorithm to interleave tissue classification and bias estimation using Expectation-Maximization (EM) algorithm and B-Spline Free Form Deformation (BFFD).

The EM algorithm [9] consists of an expectation step (E-step) which performs the classification and a maximization step (M-step) which updates the parameter estimation. Assuming a Gaussian distribution and given initial parameters, the algorithm iteratively maximizes the data likelihood and updates the tissue classification. In the context of PD image based bias correction, we decided based on experimental evidence to classify PD images into three classes: background (BG), tissue with low intensity (TL) and tissue with high intensity (TH) because the contrast level in PD images is not sufficient to delineate specific tissue classes and the purpose here is not to get a detailed segmentation. We found this three-class assumption is robust for separating the regions of background and lung from organ tissues. To improve the accuracy of inhomogeneity estimation, all background pixels are then excluded from further computations.

We assume a multiplicative bias field. For a pixel i , its measured intensity is x_i . Defining the bias field at location i as b_i , the unbiased signal r_i can be estimated by $x_i = r_i \cdot b_i$. Using the notation $\tilde{x}_i = \log(x_i)$, the image formation model can become additive $\tilde{x}_i = \tilde{r}_i + \tilde{b}_i$. Then corresponding mean and sigma become $\tilde{\mu}_k$ and $\tilde{\sigma}_k$.

A B-Spline Free-Form Deformation is applied to approximate the bias field. In this representation, a dense 2D bias field is parameterized at a sparse control point lattice. Define the field-of-view (FOV) of the PD image as follows: $\Omega_s = \{(x, y) | 0 \leq x \leq X, 0 \leq y \leq Y\}$ and ϕ_s denotes a grid of control points $\phi_{p,q}$ with the grid spacing being $\Delta_x \times \Delta_y$. This spacing between adjacent control points is uniform for each coordinate direction. The 2D tensor of uniform 1D cubic B-splines is used to represent the spatial-variant bias ratio \tilde{b}_i :

$$\mathbf{T}_{local}(\tilde{b}_i) = \sum_{m=0}^3 \sum_{n=0}^3 B_m(u) B_n(v) \phi_{p+m, q+n} \tag{1}$$

Where (x, y) is the coordinate of pixel i , and $p = \lfloor x / \Delta_x \rfloor - 1$, $q = \lfloor y / \Delta_y \rfloor - 1$, $u = x / \Delta_x - \lfloor x / \Delta_x \rfloor$, and $v = y / \Delta_y - \lfloor y / \Delta_y \rfloor$. B_m represents the m-th basis function of the B-spline. The basis functions of cubic B-splines have limited support. Therefore changing a control point in the grid affects only a 4x4 region around that control point.

Unlike the formula in [11] where a polynomial with infinite support is used to approximate the bias field, we choose not to explicitly optimize the control point value during M-step, because it leads to solving a linear system for every pixel in the image due to the local support of B-Spline. To find the optimal control point value, we estimate a ‘bias-free’ image, similar to [11]:

$$\tilde{r}_i^{(m)} = \frac{\sum_{k=1}^3 p^{(m)}(k|x_i) \cdot \tilde{\mu}_k^{(m)}}{\sum_{k=1}^3 p^{(m)}(k|x_i)} \tag{2}$$

where $\tilde{r}_i^{(m)}$ denotes the estimated real signal at pixel location i for iteration m . Then the bias for this iteration can be estimated as $\tilde{b}_i^{(m)} = approx(\tilde{x}_i^{(m)} - \tilde{r}_i^{(m)})$. $approx(\cdot)$ is the FFD approximation step, which calculates the optimal control point value $\tilde{\phi}_{i,j}^{(m)}$ (a detailed formula can be found in [10]). Given the estimated bias field, the corrected signal can be updated as $\tilde{x}_i^{(m+1)} = \tilde{x}_i^{(m)} - \tilde{b}_i^{(m)}$. Once the iteration converges or a maximum number of iterations is reached, the final bias field and corrected PD image are calculated by an exponential operator. As an illustration, Fig. 4 shows the estimated multiplicative bias field and corrected intensity profile for the PD image in Fig. 2.




Fig. 4. Estimated (a) bias field and (b) corrected intensity profile (b) for Fig. 2. This bias field is used to correct the entire perfusion time series.




Fig. 5. An illustration of multi-scale perfusion parameter estimation. Although the pixel-wise time-intensity curve shows drastic fluctuation due to imaging noise, the multi-scale strategy is able to detect the first uptake of contrast bolus. Left: a perfusion time series with clear perfusion deficit; Right: the SI curve for a selected point (marked by the green cross) and detected contrast uptake.

2.5 Perfusion Parametric Map Generation

The final step of the proposed workflow is to calculate perfusion parameter maps. Instead of computing perfusion parameters for every myocardial segment, we developed a robust map generation algorithm based on the scale-space theory, which estimates perfusion parameters for each pixel in the image. Given a SI curve $s(t)$, a series of smoothed curves $s_i(t)$ are generated by convolution with Gaussian kernels $s_i(t) = s(t) * g_i(t)$, where $g_i(t)$ is a Gaussian function with the variance being $\sigma_i, i = 1, \dots, N$ and $\sigma_i < \sigma_j, 0 < i < j \leq N$. Similarly the first-order derivative $s'_i(t)$ can be computed by convolution with the derivative of the Gaussian kernel. As no segment averaging is performed, the pixel-wise intensity curve $s(t)$ can be quite noisy (Fig. 5). To obtain a robust detection of first pass contrast bolus uptake, only the stable features that consistently appear across the whole scale space are kept. As the first step, all local maxima and zero-crossings of gradient $s'_i(t)$ are found as curve feature points. For each feature point, its appearance across all scales is checked using the concept of non-maximum suppression. The stable maximum with largest gradient is picked as the time-point corresponding to maximal up-slope. The foot time t_f is defined by the 20% of the maximal gradient, while the peak time t_p is determined by the first stable zero-crossing point after the maximal up-slope, corresponding to the maximal intensity of first bolus uptake. Once the bolus uptake region is found, a weighted least-square fit is applied to this part of time-intensity curve and the optimal up-slope is estimated. The weight for intensity point t is defined as its gradient magnitude computed from $s'_i(t)$. A scatter interpolation strategy is finally applied to fill the ‘holes’ which often appear on the noisy background where the algorithm can not find enough stable features across the scale space.

3 Results and Discussion

Validation was performed on anonymized data from 40 subjects (3 institutions), with a total of 260 perfusion series. Three different MR perfusion imaging sequences (74 TurboFLASH, 12 TrueFISP, and 174 GRE-EPI) were used in these scans. All scans

were performed with a minimum of three slice positions (basal, mid-ventricular and apical) and 2 PD images were acquired before the perfusion acquisition. We visually reviewed all datasets and classified them into two categories (no significant motion and with significant motion) according to the maximal motion magnitude presented in the series. Significant motion was not found in 92 series (35%) while the other 168 series (65%) clearly require motion compensation.

The proposed analysis workflow was applied to all series and outputs of every processing step were inspected. For the key-frame detection, we found that the detected frames consistently showed good contrast between blood pool and myocardium. Fig. 6 shows the distribution of the key-frames, which is defined as the ratio between the index of key-frame and the length of the corresponding perfusion series. The mean of the key-frame ratio is 0.44 ± 0.19 .

As demonstrated in Fig. 7 where an example of MR perfusion motion compensation is given, the jitter motion of myocardium was largely eliminated for those cases with significant motion. For cases where hearts remain stationary, we found no discernible errors were introduced by the algorithm. To quantitatively verify the motion compensation, we selected 30 series with significant motion (In-plane resolution: $1.4 \sim 2.4 \text{ mm}^2$), covering all pulse sequences and slice positions. The quantitative evaluation was performed by manually delineating the left ventricle and myocardium. For every selected series, two single 2D frames

were picked. One is the detected key-frame and the other frame was chosen when the myocardium motion was clearly discernible. Four statistical measures are computed to give a comprehensive quantification: a) T_{A-I} (the relative motion of the left ventricle center point along the Anterior-Inferior direction); b) T_{S-L} (the same motion along the Septal-Lateral direction); c) *Dice ratio* (the myocardium overlap ratio); d) *MBE* (the myocardium boundary errors, defined as the minimal distances between myocardium contours (endo and epi) extracted from the template and the registered slice). Table 1 summarizes the results, showing noticeable improvement after motion compensation.

Table 1. The quantitative measures of motion compensation

	T_{A-I} [mm]		T_{S-L} [mm]		<i>Dice ratio</i>		<i>MBE</i> [mm]	
	original	moco	original	moco	original	moco	original	moco
<i>Mean</i>	3.97	1.39	6.68	1.12	0.56	0.87	3.84	1.29
<i>STD</i>	3.26	1.03	4.69	0.83	0.20	0.04	3.19	1.24

moco: motion compensation.




Fig. 6. Distribution of detected key-frames for the whole data cohort




Fig 7. Motion compensation of a free-breathing series. First row: TrueFISP 2D slice overlaid with myocardium contour of key-frame and the intensity-time profile. Second row: the same slice after motion compensation and the corrected intensity-time profile.

The inhomogeneity correction fields estimated from the PD images were applied to the entire perfusion series to correct for the bias introduced by the surface coils. Visual inspection showed the reduction of intensity inhomogeneity that was consistently discernible throughout the datasets. To quantitatively verify the effects of bias correction, we selected the first frame of the perfusion acquisition and measured the intensity profile across the heart. We then fit a straight line to the data and estimated the absolute slope (AS) with and without bias correction. Because saturation recovery or inversion recovery pulses are normally applied to null the pre-contrast blood and tissue in the perfusion imaging, the intensity profile of the first phase often shows bias through the heart. For a group of 20 randomly selected series from the whole data cohort, the mean AS was originally 0.17 ± 0.13 and reduced to 0.06 ± 0.07 after the bias correction (Paired t-test, $P < 0.005$, 95% CI: 0.05-0.18), which is consistent with our visual impression.

As the final step, four parameter maps (SLOPE, TTP, PT and AUC) were calculated for every test case. We found that the detection of bolus uptake was highly robust, which was verified by manually sampling the SI curves from blood pool and myocardium and comparing to the automated results. As an example, for patients without clear perfusion deficits, the maps show uniform myocardium, while real perfusion deficits can be clearly visualized (Fig. 8).

To make the proposed techniques accessible to clinicians, all processing steps were implemented within the Image Calculation Environment (ICE) of Siemens MRI systems. The inline perfusion analysis starts immediately after the reconstruction of MR signals. Original time series, corrected time series, and derived parameter maps are stored in the image database after processing. The entire processing chain was implemented to support concurrent calculation via simultaneous multi-threading. The processing time for a perfusion study with three series is typically less than 1min. The proposed system is under the clinical evaluation currently taken at four institutes (Langone Medical Centre, University of New York; Laboratory of Cardiac Energetics, National Institute of Health; University Medical Centre of Ohio State, Department of Cardiology, HELIOS-Klinikum Berlin-Buch). Among the overall positive feedback, clinicians strongly favor the robustness of proposed motion compensation, as many patients were scanned in the free/shallow breathing fashion.




Fig. 8. Perfusion parameter maps (a-d: SLOPE, TTP, PT and AUC) for the series in Fig. 5. The deficit is clearly visualized.

References

1. Kellman, P., Arai, A.E.: Imaging Sequences for First Pass Perfusion-A Review. *J. Cardio. MR* 10, 525–537 (2007)
2. Guillemaud, R., Brady, M.: Estimating bias field of MR images. *IEEE Trans. Med. Imaging* 16, 238–251 (1997)

3. Puvaneswary, M.: Cardiac MR Imaging. Anshan Ltd. (2005)
4. Bracoud, L., Vincen, F., Pachai, C., Canet, E., Croisille, P., Revel, D.: Automatic Registration of MR First-Pass Myocardial Perfusion Images. In: Magnin, I.E., Montagnat, J., Clarysse, P., Nenonen, J., Katila, T. (eds.) FIMH 2003. LNCS, vol. 2674, pp. 215–223. Springer, Heidelberg (2003)
5. Milles, J., van der Geest, R.J., Jerosch-Herold, M., Reiber, J.H., Lelieveldt, B.P.: Fully automated registration of first-pass myocardial perfusion MRI using independent component analysis. *Inf. Process Med. Imaging* 20, 544–555 (2007)
6. Yang, G.Z., Burger, P., Panting, J., Gatehouse, P.D., Rueckert, D., Pennell, D.J., Firmin, D.N.: Motion and deformation tracking for short-axis echo-planar myocardial perfusion imaging. *Med. Image Anal.* 2(3), 285–302 (1998)
7. Stegmann, M.B., Olafsdottir, H., Larsson, H.B.: Unsupervised motion-compensation of multi-slice cardiac perfusion MRI. *Med. Image Anal.* 9(4), 394–410 (2005)
8. Chéfd'hotel, C., Hermosillo, G., Faugeras, O.: Flows of diffeomorphisms for multimodal image registration. In: *Proceedings of the IEEE ISBI* (2002)
9. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* 39, 1–38 (1977)
10. XLee, S.Y., Wolberg, G., Shin, S.Y.: Scattered data interpolation with Multilevel B-Splines. *IEEE Trans. Visualization and Computer Graphics* 3(3), 228–244 (1997)
11. Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P.: Automated model-based bias field correction of MR images of the brain. *IEEE Trans. Med. Imaging* 18, 885–896 (1999)