Published in final edited form as:

J Magn Reson Imaging. 2017 October ; 46(4): 1060–1072. doi:10.1002/jmri.25659.

A Robust Universal Nonrigid Motion Correction Framework for First-Pass Cardiac Magnetic Resonance Perfusion Imaging

Mitchel Benovoy, PhD^{1,2}, Matthew Jacobs, MSc^{1,3}, Farida Cheriet, PhD², Nagib Dahdah, MD⁴, Andrew E. Arai, MD¹, and Li-Yueh Hsu, DSc¹

¹National Heart, Lung and Blood Institute, National Institutes of Health, Bethesda, USA

²Department of Biomedical Engineering, Polytechnique Montreal, Montreal, Canada

³Department of Electrical Engineering and Computer Science, Catholic University of America, Washington DC, USA

⁴Sainte-Justine University Hospital Research Center, Montreal, Canada

Abstract

Purpose—To present and assess an automatic nonrigid image registration framework that compensates motion in cardiac MR perfusion series and auxiliary images acquired under a wide range of conditions to facilitate myocardial perfusion quantification.

Materials and Methods—Our framework combines discrete feature matching for large displacement estimation with a dense variational optical flow formulation in a multithreaded architecture. This framework was evaluated on 291 clinical subjects to register 1.5T and 3.0T steady-state free precession (FISP) and fast low angle shot (FLASH) dynamic contrast myocardial perfusion images, arterial input function (AIF) images and proton density (PD) weighted images acquired under breath hold (BH) and free breath (FB) settings.

Results—Our method significantly improved frame-to-frame appearance consistency compared to raw series, expressed in correlation coefficient ($R^2=0.996\pm3.735E-3$ vs. $0.978\pm2.024E-2$, $p<0.0001$) and mutual information ($3.823\pm4.098E-1$ vs. $2.967\pm4.697E-1$, $p<0.0001$). It is applicable to both BH ($R^2=0.998\pm3.217E-3$ vs. $0.990\pm7.527E-3$) and FB ($R^2=0.995\pm3.410E-3$ vs.

Corresponding Author: Li-Yueh Hsu, DSc, Bldg 10, Rm B1D416, MSC 1061, 10 Center Drive, National Heart, Lung and Blood Institute, National Institutes of Health, Bethesda, MD 20892–1061, U.S.A. Tel: (301) 435–5839, Fax: (301) 402–2389, lyhsu@mail.nih.gov.

SUPPLEMENTARY MATERIAL

Comparisons of pre- (RAW) vs. post-motion corrected (MOCO) images with the proposed framework for first-pass cardiac MR perfusion imaging are shown in supplemental movie files. In the first video (Breath Hold) when the patient was instructed to hold his breath during the perfusion imaging, two respiratory gasps occurred during the contrast passage and resulted in large displacements of the heart due to chest wall and diaphragmatic motion. The results show LDOF-MOCO did an excellent job on correcting this cardiac motion.

In the second video (Free Breathing) when the perfusion imaging was performed while the patient breathes normally, there was quasi-periodic cardiac motion in the series due to normal tidal breathing. LDOF-MOCO also performed excellently in correcting the cardiac motion. Both examples show that the system can largely eliminate cardiac motion on the arterial input function (AIF) images, proton-density (PD) images, and perfusion (MYO) images. The system also worked well for all three slice locations (basal, mid-ventricular, and apical) of the stress perfusion series shown in both video files.

Disclosures

Dr. Arai is the principal investigator on a Cooperative Research and Development Agreement with Siemens. The remaining authors declare that they have no competing interests.

0.968±2.257E-3) paradigms as well as FISP and FLASH sequences. The method registers PD images to perfusion T1 series (9.70% max increase in R^2 vs. no registration, $p<0.001$) and also corrects motion in low resolution AIF series ($R^2=0.987\pm1.180E-2$ vs. $0.964\pm3.860E-2$, $p<0.001$). Finally, we showed the myocardial perfusion contrast dynamic was preserved in the motion corrected images compared to the raw series ($R^2=0.995\pm6.420E-3$).

Conclusion—The critical step of motion correction prior to pixel-wise cardiac MR perfusion quantification can be performed with the proposed universal system. It is applicable to a wide range of perfusion series and auxiliary images with different acquisition settings.

Keywords

Cardiac Magnetic Resonance; Myocardial Perfusion; Contrast Enhancement; Quantitative Perfusion; Motion Correction; Nonrigid Image Registration

INTRODUCTION

First-pass contrast-enhanced cardiac magnetic resonance (MR) imaging is an increasingly reliable and informative tool for diagnosing coronary artery disease (1). Newer techniques can generate pixel-wise myocardial blood flow maps for high-resolution quantitative analysis capable of differentiating transmural perfusion gradients in patients with varying severities of ischemia (2, 3). However, despite the development and optimization of acquisition paradigms and parameters, many studies are hindered by motion of the heart and surrounding tissue during the acquisition (4). As myocardial perfusion analyses increasingly target individual pixel resolutions, the necessity of frame-to-frame pixel-wise signal intensity correspondence becomes critical and is hindered by labor-intensive and subjective manual operator interaction (5).

The physical sources of cardiac motion are respiration, voluntary patient displacement, involuntary thoracic organ movement, and the pumping action of the heart chambers (4). Of interest to myocardial perfusion imaging, the left ventricle (LV) and right ventricle (RV) present motion patterns characterized by longitudinal shortening, radial contractions, and rotations. Electrocardiogram gating is widely used in cardiac MR imaging to synchronize the cardiac phase, but it is prone to errors, especially under stress conditions or due to arrhythmias. Motion in myocardial perfusion imaging can also be mitigated, by instructing the patient to hold their breath (typically lasting 60 heartbeats). However, this can be taxing for patients and results in respiratory gasps that can induce strikingly large motion events in the perfusion images. Conversely, while free breathing acquisition paradigms lower patient and clinician constraints, they introduce quasi-periodic thoracic movement throughout the acquired series which must be taken into consideration at the analysis stage. Retrospective image-based motion correction (MOCO) is therefore a critical step in the quantitative myocardial perfusion analysis pipeline to align the imaged anatomical structures throughout the series that invariably present different appearances due to the issues outlined above.

For practical clinical considerations, an ideal MOCO system must be capable of handling cardiac MR images produced from different magnetic field strengths and sequence types that may have different signal-to-noise ratios (SNR), spatial resolutions, and contrast

characteristics. As cardiac MR perfusion series are typically acquired at three different short-axis slice locations (basal, mid-ventricular, and apical) under both rest and stress conditions, the MOCO algorithm must also be robust to diverse anatomical appearances and pacing conditions, respectively. Moreover, the compatibility with myocardial perfusion images acquired under a free breathing paradigm would be an interesting asset in terms of augmenting patient comfort and clinical applicability.

The use of auxiliary images is an increasingly important aspect for cardiac MR perfusion image diagnosis and quantification. These include proton density (PD) weighted images, which are used to normalize myocardial intensity inhomogeneities to enhance perfusion quantification (6, 7). Their appearance is inherently distinct from conventional perfusion T1-weighted images by fast imaging with steady-state free precession (FISP) and fast low angle shot (FLASH) imaging sequences. For example, PD-T1-weighted configurations may include: FISP-FISP, FLASH-FISP, or FLASH-FLASH (8). For optimal bias field normalization, intra-PD MOCO followed by PD-to-T1 perfusion series registration is required to align the anatomical structures, particularly the myocardial region, between the two types of images. Arterial input function images are another class of auxiliary series specifically designed to maintain the linearity of the LV signal intensity during contrast bolus passage and are used to calculate myocardial blood flow (9). Despite their interleaved acquisition with the main perfusion images, they are differentiated by an optimized dynamic range of contrast information in the LV cavity at the cost of a significantly lower spatial resolution and lower SNR in the myocardium.

The purpose of this work is to present and assess an automatic motion correction procedure that addresses the requirements stated above for clinical applicability. This work improves the automated nonrigid motion correction framework for first-pass cardiac MR perfusion imaging presented in our previous work (10) by evaluating a more refined and robust system on a large clinical cohort over a wide range of image acquisitions types and settings.

MATERIALS AND METHODS

Study Population

First-pass cardiac MR perfusion images of 291 clinical subjects were evaluated in this study. All studies were performed under procedures and protocols approved by our institutional review board, and all subjects gave written informed consent. Indications for stress myocardial perfusion included diagnosis of coronary artery disease and chest pain, assessment of coronary stenosis severity, and technical development based on asymptomatic patients and healthy volunteers.

Image Acquisition

Cardiac MR imaging was conducted on 1.5 and 3.0 Tesla scanners (Siemens Healthcare, Erlangen, Germany) with a saturation recovery dual-sequence technique (9) based on a FISP or FLASH pulse sequences. Perfusion studies were obtained by administering a 0.05 mmol/kg bolus of gadolinium-based contrast agent (Magnevist; Berlex Laboratories, Wayne, NJ, USA) followed by a saline flush. Stress and rest studies were performed at the basal,

mid-ventricular, and apical slice positions during either breath hold or free breathing for 60 heartbeats. Typical imaging parameters for perfusion imaging included: 90° composite saturation preparation pulse, 50° (FISP) or 12° (FLASH) flip angle, 90 ms (FISP) or 100 ms (FLASH) inversion time, 2.3 ms repetition time, 1.2 ms echo time, 360 × 270 mm field of view, 128 × 80 to 192 × 144 acquisition matrix, 192 × 128 to 256 × 192 image matrix after interpolation, 2.8 × 3.4 to 1.9 × 1.9 mm raw pixel resolution, 8 mm slice thickness, and parallel imaging factor of 2 to 3. The low resolution AIF series were acquired at the beginning of each RR cycle at the basal location with the same field of view but a reduced acquisition matrix of 64 × 48 using a FLASH readout (8). Additionally, two or three PD-weighted images were acquired in each perfusion series using a small 5° magnetization flip angle and no saturation preparation pulse prior to both FISP and FLASH perfusion imaging.

Motion Correction Framework

The schematic diagram of the proposed motion correction framework is shown in Figure-1. Briefly summarized, the system uses post-hoc interpolation warping following nonrigid displacement estimation based on an optical flow formulation specifically amenable to perfusion series dynamics. The purpose of this framework is to register contrast-enhanced cardiac MR perfusion image series with varying characteristics, including different field strengths, sequence types and breathing paradigms, and be compatible with auxiliary series such as PD images and independent AIF acquisitions used to facilitate fully quantitative myocardial perfusion analysis. The implementation also features automatic reference frame and PD image detection and a multithreaded architecture to improve processing speed.

Image Pre-Processing

A proxy copy of the raw image series was used to estimate the motion, while the original raw series is used exclusively at the later registration stages during geometrical transformations. To standardize processing parameters and optimize computation time, the proxy images are first rescaled to a default size (long dimension = 128 pixels for perfusion series, 96 pixels for AIF series) and quantized to a lower dynamic range (8 bits per pixel). As the image warping following motion estimation uses the original raw frame, there is no reduction in image quality besides the expected smoothing caused by the interpolation.

Reference Perfusion Frame and PD Frame Detection

A reference frame is used as an initial target to which all other images will be registered. The optimal reference was identified as the baseline frame immediately preceding contrast arrival in the RV. This image depicts contrast-free anatomical structures compared to the later stages, where frame-to-frame pixel fluctuations are higher and large respiration events are prone to occur. This frame ρ is automatically detected by first denoising the proxy series using principle component analysis (PCA) decomposition with 85% variance retention to reduce transient image artifacts and noise, in addition to smoothing out structural movement across the proxy series. Sequential neighbors of the PCA-reduced series are then cross-correlated together until a large correlation coefficient drop is detected with differential analysis, which ostensibly corresponds to the onset of contrast arrival in the RV. The PD frames preceding the perfusion T1 series are automatically detected in a similar way by identifying the large correlation variation corresponding to the PD-to-T1 image transition.

Large Displacement Optical Flow Motion Estimation

Nonrigid deformation maps between pairwise sequential frames are computed in the perfusion series using an optical flow (OF) formulation robust to large displacement (LDOF) (11). This characteristic is essential to accommodate the wide range of motion found in the perfusion series and the fast ventricular bolus arrivals without causing motion estimation artifacts. As illustrated in Figure-2, the LDOF formulation couples discrete feature points tracking computed from histograms of oriented gradients (HOG) (12) with a continuous variational optical flow step (13) in a coarse-to-fine optimization scheme (14). Given two images I_1, I_2 to be aligned and a two-dimensional point $\mathbf{x}$ in the image domain $\omega: \mathbf{x} = (x, y)^T$, the optical flow field $\mathbf{w} := (u, v)^T$ is given by the pixel intensity energy term:

$$E_{int}(\mathbf{w}) = \int_{\omega} \psi(|I_2(\mathbf{x} + \mathbf{w}(\mathbf{x})) - I_1(\mathbf{x})|^2) dx \quad (1)$$

where $\psi(s^2) = \sqrt{s^2 + \varepsilon^2}$ with $\varepsilon = 1\text{E-}3$ is a robust function to mitigate occlusions. Equation 1 penalizes deviations from the assumption that corresponding points should have the same intensity, but this is seldom encountered in practice so a gradient (∇) constraint invariant to additive illumination is defined:

$$E_{grad}(\mathbf{w}) = \int_{\omega} \psi(|\nabla I_2(\mathbf{x} + \mathbf{w}(\mathbf{x})) - \nabla I_1(\mathbf{x})|^2) dx \quad (2)$$

Equations 1 and 2 match relatively weak features (intensity and gradients), and may result in non-unique solutions. A regularization term applied to the computed flow fields is thus introduced:

$$E_{smooth}(\mathbf{w}) = \int_{\omega} \psi(|\nabla u(\mathbf{x})|^2 + |\nabla v(\mathbf{x})|^2) dx \quad (3)$$

Hence, the variational optical flow model is given by:

$$E(\mathbf{w}) = E_{int}(\mathbf{w}) + \gamma E_{grad}(\mathbf{w}) + \alpha E_{smooth}(\mathbf{w}) \quad (4)$$

where α , and γ are tuning parameters. Large displacement estimation is achieved by adding the HOG-based landmark correspondence energy term:

$$E_{match}(\mathbf{w}) = \int \delta(\mathbf{x}) \tau(\mathbf{x}) \psi(|\mathbf{w}(\mathbf{x}) - \mathbf{w}_1(\mathbf{x})|) dx \quad (5)$$

where $\mathbf{w}_1(\mathbf{x})$ is a correspondence vector constructed by matching a given point $\mathbf{x}$, $\delta(\mathbf{x})$ is a binary variable indicating the presence of a match, and τ is a matching score related to the distance between matches descriptor vectors. The task of discrete descriptor matching can be formulated in a continuous approach to make it compatible with the continuous variation model:

$$E_{desh}(\mathbf{w}) = \int \delta(\mathbf{x}) |f_2(\mathbf{x} + \mathbf{w}_1(\mathbf{x})) - f_1(\mathbf{x})|^2 dx \quad (6)$$

where $f_1(\mathbf{x})$, $f_2(\mathbf{x})$ are the fields of descriptor vectors in I_1 , I_2 . Thus, the final LDOF model can be represented as:

$$E(\mathbf{w}) = E_{int}(\mathbf{w}) + \gamma E_{grad}(\mathbf{w}) + \alpha E_{smooth}(\mathbf{w}) + \beta E_{match}(\mathbf{w}, \mathbf{w}_1) + E_{desc}(\mathbf{w}_1) \quad (7)$$

The optical flow field ($\mathbf{w}$) that minimizes (7) is therefore the optimal solution. The HOG tracking component handles large displacements of arbitrarily moving anatomical structures, while the variational module estimates dense flow fields to measure relatively subtle deformations such as the elastic motion of the myocardial wall at sub-pixel accuracy. The LDOF engine is controlled with three primary parameters: α , the smoothness of the flow fields; β , the weight of the HOG descriptor matching term; and γ , the weight of the gradient consistency term. These parameters are set to optimized values (small α , large β and γ) automatically learned from a small subset of patients using a grid-search approach (15) minimizing the propensity to generate artefacts in the flow fields, while maximizing the mean sequential-frame correlation score.

Instead of immediately registering the current frame before computing the subsequent motion, the stepwise deformation fields (f^x , f^y) of each image pair are saved for warping at a later stage. Instead of immediately registering all images to a reference frame or to consecutively registered images (16, 17), this post-hoc approach avoids successive smoothing. This technique may help image series that have large motion events due to incomplete patient breath holding or high tissue contrast changes during late perfusion. Furthermore, as the estimation of a given flow field does not depend on pre-registered frames, the MOCO pipeline is designed as a parallel processing architecture where groups of neighboring image pairs are handled by multiple LDOF processing engines and executed by independent CPU threads.

Following motion estimation on every sequential image pair, all raw perfusion images are registered to the reference frame ρ using 2D bicubic interpolation warping with the cumulative displacements computed from the stepwise estimates. This procedure fully resolves the pixel-wise $[x, y]$ excursions from ρ to a given image frame t :

$$[F_t^x F_t^y] = \begin{cases} \sum_{i=\rho}^{t-1} [f_i^x, f_i^y], & t > \rho \\ \sum_{i=t}^{\rho-1} [f_i^x, f_i^y], & t < \rho \end{cases} \quad (8)$$

PD-to-Perfusion Image Registration

Once motion correction for the T1-weighted perfusion images is completed, the PD-weighted images are aligned to the last PD frame. These images are then registered to the perfusion series via two bridge images of the PD and T1 frames to match their photometric properties and increase their structural information, as shown in Figure-3. The T1 bridge

image (I_{brT1}) is constructed by first taking the median (I_{T1}) of the motion-corrected baseline T1 frames leading up to the reference frame p as these are not influenced by ventricular contrast dynamics. An edge-enhanced image (I_{edgeT1}) of the I_{T1} is computed using a weighted combination of its gradient image obtained with the Sobel operator. Similarly, the PD bridge image (I_{brPD}) is constructed with the median (I_{PD}) and edge-enhanced (I_{edgePD}) images. Next, histogram equalization and unsharp masking are performed on the I_{edgeT1} and I_{edgePD} images to further enhance the anatomical boundary gradients. This is followed by exact histogram specification (18) to form the final bridge image pair. The flow field mapping I_{brPD} to I_{brT1} is computed with the LDOF process outlined above using a high gradient consistency parameter γ and a large descriptor matching weight β to emphasize the discrete HOG feature correspondence component. The resulting deformation estimation is finally applied to the motion-corrected PD images to re-register them with the motion-corrected perfusion T1 series.

Image Post-Processing

Post-processing of the PD and perfusion T1 motion-corrected images is limited to matching their photometric profiles with their corresponding raw frames using exact histogram specification (18). This recovers the original signal intensity distributions that have invariably been altered during the 2D interpolation warping stage.

Performance Evaluation

The LDOF motion estimation framework was implemented in the C/C++ programming language with Intel's (Santa Clara, U.S.A.) Math Kernel Library under Microsoft (Redmond, U.S.A.) Windows 7 64-bit operation system. The results of the proposed approach are labeled "LDOF-MOCO", whereas the series without motion compensation are designated "raw". The comparative analyses were carried out on the full frame images as well as a square region of interest encompassing the RV, LV and myocardium (henceforth labelled "heart region") that was automatically delimited using an autonomous heart ventricle finder (19).

As a measure of frame-to-frame appearance consistency and structural alignment, Pearson's squared correlation coefficients (R^2) and mutual information (MuInfo) (20) were computed on sequential image pairs of the perfusion T1 series across the full cohort of patients.

The R^2 between the last PD frame and the first perfusion T1 frame was used as a measure of registration quality for different PD-T1 perfusion series combinations.

To study whether the MOCO process distorts the pixel intensity profile of the highly dynamic signal changes in the myocardial regions undergoing perfusion, stress and rest time-intensity curves of the LV and myocardium were measured on all three slices of the LDOF-MOCO and raw perfusion series in a subset of 17 healthy subjects. On the MOCO series, the myocardial region of interest (ROI) was manually drawn on the epicardium and endocardium walls by a trained cardiac MR imager with two years of experience. The anterior septal RV insertion point was also positioned to divide the myocardium into six equiangular sectors. The mean intensities of the pixels encompassed by the endocardial and sector boundaries at each frame were then computed to produce the LV and myocardium

perfusion time signal intensity curves, respectively. To extract the signal intensity dynamics of the corresponding myocardial regions from the raw series, the myocardial borders and the RV insertion point delimited above were automatically *backwarped* from the MOCO series throughout the raw series using the flow fields computed from the LDOF engine. The corresponding myocardial pixels in both MOCO and raw series are thus included in the corresponding sectoral ROIs for time-intensity curve comparisons. This method allows the non-linear tracking of the ROI boundaries from the MOCO series back to the raw series and the measurement of the signal intensity dynamics from the native image pixels. Both R^2 and normalized root-mean-squared-error (NRMSE) were used to measure the LV and sectoral perfusion dynamic similarities between the LDOF-MOCO and raw series.

Finally, translational motion correction was evaluated locally in the heart region by computing the LV centroid displacement on the three slices of the subset cohort series that underwent ROI backwarping. This analysis assesses the amount of global translational movement in the heart region present in the raw series that was subsequently corrected by our method.

Statistical Analysis

The mean and standard deviation (mean $\pm$ SD) are reported unless stated otherwise for the evaluation metrics used throughout the analysis. Non-parametric Mann-Whitney U-tests were carried-out to compare different test groups. A p-value < 0.05 was considered statistically significant.

RESULTS

As summarized in Table-1, all stress and rest perfusion studies from 291 clinical subjects were processed with the proposed motion correction framework without exclusion. A total of 1746 myocardial perfusion series (three stress and three rest perfusion slices per patient) and 582 AIF series (one stress and one rest slice per patient) were processed using the same LDOF parameter settings. The computation time to process a 60 frame standard perfusion series including PD-to-perfusion registration averaged 17.96 ± 1.63 seconds and 10.13 ± 2.58 seconds for the AIF series on an Intel Core i7 6950X 3.0 GHz CPU.

The qualitative results shown in the supplemental videos illustrate our approach's effectiveness to compensate cardiac and respiratory motions in both breath-hold and free-breathing series for the AIF, PD and perfusion images.

PD-to-Perfusion Image Registration

For the automatic PD-to-T1 registration, Figure-4 shows results from different sequence combinations. An RGB composite color image is used to display the superposition of the images before and after registration. The PD image is presented in the green channel, the T1 image is encoded in red, and the overlay of the two images is fused in yellow color. Qualitatively, there is a consistent improvement of myocardial wall alignment after the registration in all three PD-T1 combinations. Moreover, R^2 in both full frame and heart region images were increased in all registered images compared to raw in these examples. For the overall cohort, the largest improvement was obtained in the PD_{FLASH}-to-T1_{FISP}

configuration with a 9.70% increase of R^2 score compared to no registration ($p = 0.0005$). Such improvement is followed by PD_{FLASH}-to-T1_{FLASH} (3.76%, $p=0.0472$) and PD_{FISP}-to-T1_{FISP} (1.44%, $p=0.0497$) configurations. However, despite statistically significant improvements in R^2 in all the PD-T1 configurations for both full-frame and heart region comparisons, the metric used to assess the quality of the registration is hindered by the large photometric and textural differences of these two types of images. As evidence, there is a modest increase of the R^2 scores in all comparisons and drastic qualitative improvement is apparent in the accompanied Figure-4 comparison.

Sequential Perfusion Frame Consistency

Figure-5 shows the mean consecutive-frame R^2 and MuInfo curves of the raw and LDOF-MOCO perfusion series across the entire cohort. All R^2 curves show a drop around frames five-to-eight which corresponds to high signal intensity changes in the ventricles due to contrast bolus arrival. For the full frame comparison, the R^2 curve of the LDOF-MOCO presents higher scores than the raw throughout the series and also exhibits noticeably less jitter. These fluctuations seen in the raw curves after approximately frame 30 is noticeably improved in the LDOF-MOCO curves. For a stricter comparison in the heart region encompassing the RV, LV, and myocardium, the improvement is further pronounced and shows a similar trend with an increased R^2 and reduced jitter over the entire range.

The MuInfo curves also show large improvements during baseline and early perfusion until reaching a plateau near peak myocardial enhancement. This reflects the structural enhancement in the ventricles and myocardium during the arrival of the bolus, which improves the similarity between adjacent frames. This observation is more pronounced in the heart region than in the full frame comparisons.

For group statistics the mean correlation coefficient shows improved frame-to-frame consistency in LDOF-MOCO vs. raw on the full frame series. Our approach scored $R^2 = 0.996 \pm 3.735E-3$ and raw followed with $R^2 = 0.978 \pm 2.024E-2$ in the full frame comparison ($p < 0.0001$). In the heart region LDOF-MOCO obtained $R^2 = 0.985 \pm 1.459E-2$, whereas raw trailed with $R^2 = 0.933 \pm 5.205E-2$ ($p < 0.0001$). Of clinical interest, the heart region is well managed by LDOF-MOCO, but shows a generally lower R^2 score compared to full frame as this area undergoes large intensity fluctuations during perfusion. The MuInfo shows a similar improvement with an increase of 28.89% in full frame and 17.25% in the heart region (both $p < 0.0001$).

The breathing paradigm used during acquisition has a modest, but statistically significant effect on the performance of the MOCO system (Figure-6). In breath hold series, the LDOF-MOCO approach has a $2.764E-3$ higher R^2 than in free breathing for full frame comparison ($p < 0.05$), whereas in the heart region the R^2 increases even further by $8.480E-3$ ($p < 0.05$).

Finally, there was no significant difference between 1.5T and 3.0T results in the full frame ($R^2 = 0.996 \pm 3.735E-3$ vs. $0.996 \pm 3.55E-3$, $p=0.214$) or in the heart region ($R^2 = 0.984 \pm 1.549E-2$ vs. $0.987 \pm 1.353E-3$, $p=0.146$).

AIF Series Image Registration

The motion in the AIF series was also well compensated by our approach. The average full frame R^2 scores were improved to $0.987 \pm 1.180E-2$ after LDOF-MOCO vs. $0.964 \pm 3.860E-2$ for the raw, whereas the MuInfo scores increased by 18.54% from $2.325 \pm 9.663E-1$ to $2.756 \pm 8.798E-1$ (both $p < 0.0001$). Likewise, R^2 scores in the heart region reached $0.974 \pm 2.310E-2$ after LDOF-MOCO compared to $0.944 \pm 5.240E-2$ for the raw, and the MuInfo increased to $3.349 \pm 8.907E-1$ from $3.051 \pm 9.730E-1$ (both $p < 0.0001$).

LV and Myocardial Perfusion Dynamics

The MOCO series produced by our approach conserves the signal intensity dynamics incurred by the contrast agent transit and subsequent myocardial tissue perfusion. This is illustrated in Figure-7 for a stress perfusion series where the effects of motion on the time signal intensity curves are evident for the raw series, but well-controlled after LDOF-MOCO. Across our test cohort, as shown in Table-2a, the correlation between the time signal intensity curves of the LDOF-MOCO and the raw-backwarped series shows excellent agreement with R^2 values of $0.999 \pm 5.070E-4$ in the LV cavity and an average of $0.995 \pm 6.420E-3$ in the myocardial sectors. The average correlation between the raw vs. raw-backwarped series was $0.997 \pm 2.811E-3$ in the LV and only $0.894 \pm 1.501E-1$ in the myocardium.

The NRMSE of the LDOF-MOCO and the raw-backwarped signal intensity curves is also well managed, as shown in Table-2b, with a mean value of $7.580E-3 \pm 3.906E-3$ for the LV signal and $5.245E-2 \pm 2.139E-2$ in the myocardial sectors. These scores represent a 34.40% increase for the LV and up to 51.04% in the myocardial sectors (both $p < 0.01$) compared to the static ROI measures.

LV Centroid Displacement

Translational movement of the LV region is largely corrected by the proposed method. The raw mid-ventricular and apical slices show progressively more motion compared to the basal slice, particularly in the y direction due to chest wall movement. In the test cohort, the mean cumulative Euclidean displacement of the LV region corrected by the LDOF-MOCO is 59.11 ± 34.21 mm in the basal slices; 62.26 ± 33.89 mm in the mid-ventricular slices; and 66.03 ± 33.00 mm in the apical slices.

DISCUSSION

The work presented in the study addresses several important requirements for a robust automated nonrigid motion correction scheme in cardiac MR perfusion series. Although our approach uses an optical-flow based algorithm as the core non-linear motion estimator which is similar to the previous work by Xue et al. (16), this work advances the technique in several aspects. The novel contributions of the proposed system include:

- Robustness to multiple perfusion imaging sequence configurations (FISP, FLASH, AIF) and magnetic field strengths (1.5T and 3.0T).

- Multi-contrast PD-weighted to T1-weighted perfusion image registration, including various PD-T1 configurations.
- Robustness to both free breathing and breath hold paradigms.
- Explicit post-MOCO photometric specification to preserve perfusion dynamics.
- Multithreaded architecture designed to increase computational efficiency.

Various image registration schemes have been developed to compensate for myocardial motion in cardiac MR perfusion imaging. As a substantial portion of motion in the series is composed of in-plane translational and rotational movement, a large body of the earlier work was focused on rigid registration (21 – 23). Although these techniques improve qualitative and some aspects of quantitative analysis, they fail in precisely aligning the epicardium and endocardium tissue undergoing elastic deformations, which is needed for high-resolution pixel-wise myocardial perfusion analysis. Nonrigid registration techniques have thus garnered recent research interest.

Notable contributions include techniques that build a motion-free synthetic reference series to which the perfusion frames are registered. Wollny et al. (24) devised a multi-pass scheme that uses independent component analysis and quasi-periodic breathing patterns to identify motion components throughout the series. Motion-free synthetic images were then created by removing the primary independent components related to movement. Likewise, Li et al. (25) constructed a pseudo ground-truth synthetic series in the heart region that was free of translational motion, image noise, and conserved pairwise image intensity profiles by using a spatiotemporal smoothness constraint to ostensibly help the registration stage. Following semi-automatic segmentation of the LV and RV, elastic myocardium deformations were corrected by aligning the real and synthetic image sequences together. Conversely, some techniques use a unique reference frame to which all other images are registered. Khalifa et al. (26) employed this paradigm by first performing rigid registration followed by nonrigid motion correction. They started by manually delineating the LV walls to linearly align the segments to a single reference image. This was followed by two stages of nonrigid registration where a B-spline transformation model was used to register the target and reference myocardial walls. A refinement step was then performed by matching the iso-contours of the walls formed using the Laplace-potential approach (27). Tarroni et al. (28) also non-linearly registered the myocardial walls using a semi-automatic approach requiring manual LV seed point placement. Following automatic reference frame detection, the cavity (including the papillary muscles) was segmented using a region-based level set approach (29). The combination of a correlation-based rigid motion estimation followed by boundary refinement was then used to register the frames to the reference target. Finally, as mentioned above, Xue et al. (16) used a variational optical-flow based approach to consecutively register neighboring images to an automatically determined key frame. However, unlike our approach, this technique and others that perform serial or sequential registration (17, 30) tend to accumulate alignment errors and may result in large motion estimation artifacts for frames temporally distant from the reference.

The issues complicating motion correction compiled from this body of work originate from the patient specific displacement and deformation of different anatomical structures during

dynamic contrast enhancement. It is thus difficult to devise population-based statistical models or rely on anatomical assumptions. The dramatic contrast variations observed in the cardiac MR perfusion series during LV and RV contrast bolus passages, and subsequent tissue perfusion must also be taken into account to achieve robust performance by any motion correction systems. Furthermore, we note the scarcity of methods tested with varying sequence types, under both traditional breath hold and free breathing paradigms, and correcting motion within/between auxiliary/perfusion images used to improve perfusion quantification. Consideration for computational speed, which is required for practical clinical workflow, is also seldom discussed.

In our quantitative evaluation over a large and diverse cohort, we showed motion reduction reflected by both R^2 and MuInfo metrics in full frame and heart region series. Temporal analysis of the MOCO performance of our approach showed smoother frame-to-frame appearance consistency and noticeably less jitter caused by diaphragm movement from tidal breathing or respiration gasps. Robustness to respiration was also validated by our results. This highlights the strength of the discrete feature matching component of the LDOF engine as large chest wall and diaphragm displacements caused by breath hold failure or free breathing were effectively corrected.

The motion in the low resolution AIF series was also corrected by the LDOF-MOCO. Despite a drastically reduced image quality and resolution in the AIF series, similar improvement in R^2 scores to the higher resolution perfusion T1 series was shown in full frame and particularly in the heart region, which may facilitate the automation of the LV signal intensity measurements used for perfusion quantification and signal calibration (19). In addition, our LDOF approach was successfully generalized for PD-to-T1 image registration. The HOG feature matching component was shown to be robust to the varying appearances of the PD-T1 combinations as correlation scores improved modestly. However, in our opinion this numerical score may not adequately reflect the drastic improvement in myocardial alignment between PD and T1 series as shown in the qualitative comparison.

In terms of clinical applicability, our approach preserves tissue perfusion dynamics as measured in the LV cavity and the myocardium. Using the ROI backwarping approach, excellent agreement of perfusion dynamics was found in the LV and myocardial region between the two series. With regards to clinical workflow, an adequate processing was obtained on a standard quad-core processor desktop computer and will likely improve as our multithreaded pipeline can scale favorably with increasingly sophisticated processor architectures.

Lastly, the amount of translational motion, which is the major detrimental factor for cardiac MR perfusion analysis, corrected by our approach was measured by tracking the LV cavity centroid across the perfusion T1 series. The large displacement of the LV cavity subsequently corrected highlights the benefits of automatic motion correction and its necessity for pixel-wise perfusion quantification applications.

The work presented here added several layers of robustness and improvement upon our previously presented method (10). These include the locally-weighted HOG feature

matching, dynamic range and pyramid decomposition optimization, and an decreased bias to the regularization function in the intensity matching energy minimization term. However, it does present some limitations. Foremost, some residual movement of diffuse, structure-poor regions may still be present in areas peripheral to the myocardium. Although these regions are of no consequence to myocardial perfusion analysis, their persistent motion is noteworthy. Another issue is the optimization of various parameters controlling the quality of the motion estimates. We have used an automatic grid-search approach to find optimal combinations of the parameters globally, but this step is dependent on the variability of the data used for training. Re-estimation of these parameters may be needed for optimal performance with different MR sequence types or image dataset. An important study design limitation is the relatively small cohort of subjects that were included in the LV and myocardial perfusion dynamics comparisons. While our results showing an excellent agreement of perfusion dynamics between the raw and MOC images, further studies are required to evaluate such comparison in a larger patient cohort and under various vendor platforms.

In conclusion, this paper presents an automatic and universal nonrigid motion correction system for first-pass cardiac MR perfusion series. It was shown to be robust to varying acquisition parameters including different sequence type and field strength (FISP, FLASH, 1.5T, 3.0T), breathing paradigm (BH, FB), SNR and image resolution (AIF, T1). It was also demonstrated to successfully register proton density weighted images to T1-weighted perfusion images in varying sequence types. Importantly, our approach was shown to preserve the signal intensity dynamics crucial for qualitative and quantitative diagnosis. The proposed system showed adequate motion correction on a large cohort of clinical patients. Finally, its multi-threaded post-hoc warping is another notable contribution. Designed around strict operational requirements, the system's quantitative and qualitative performance, combined with its multithreaded architecture to deliver reasonable processing time promotes its integration in routine cardiac MR perfusion analysis protocols.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This project was funded and supported by the intramural research program of the National Heart Lung and Blood Institute (ZIA HL006137-05) and the Bourses de Recherche MÉDITIS.

Abbreviations

AIF	Arterial Input Function
CPU	Central Processing Unit
BH	Breath Hold
FB	Free Breath
FISP	Fast Imaging with Steady-state Free Precession

FLASH	Fast Low Angle Shot
HOG	Histograms of Oriented Gradients
LDOF	Large Displacement Optical Flow
LV	Left Ventricle
MOCO	Motion Correction
MR	Magnetic Resonance
MuInfo	Mutual Information
NRMSE	Normalized Root-Mean-Square Error
OF	Optical Flow
PCA	Principal Component Analysis
PD	Proton Density
R²	Squared Pearson's Correlation Coefficient
ROI	Region of Interest
RV	Right Ventricle
SD	Standard Deviation
SNR	Signal-to-Noise Ratio

References

1. Coelho-Filho OR, Rickers C, Kwong RY, Jerosch-Herold M. MR myocardial perfusion imaging. *Radiology*. 2013; 266:701–15. [PubMed: 23431226]
2. Hsu LY, Groves DW, Aletras AH, Kellman P, Arai AE. A quantitative pixel-wise measurement of myocardial blood flow by contrast-enhanced first-pass CMR perfusion imaging: microsphere validation in dogs and feasibility study in humans. *JACC: Cardiovascular Imaging*. 2012; 5:154–66. [PubMed: 22340821]
3. Zarinabad N, Chiribiri A, Hautvast GL, et al. Voxel-wise quantification of myocardial perfusion by cardiac magnetic resonance. Feasibility and methods comparison *Magnetic Resonance in Medicine*. 2012; 68:1994–2004. [PubMed: 22354744]
4. Scott AD, Keegan J, Firmin DN. Motion in Cardiovascular MR Imaging 1. *Radiology*. 2009; 250:331–51. [PubMed: 19188310]
5. Zarinabad N, Chiribiri A, Breeuwer M. Myocardial blood flow quantification from MRI – an image analysis perspective. *Current Cardiovascular Imaging Reports*. 2014; 7:1–9.
6. Hsu, LY., Rhoads, KL., Aletras, AH., Arai, AE. *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer; Berlin Heidelberg: 2003. Surface coil intensity correction and non-linear intensity normalization improve pixel-resolution parametric maps of myocardial MRI perfusion; p. 975-976.
7. Miller CA, Hsu LY, Ta A, Conn H, Winkler S, Arai AE. Quantitative pixel-wise measurement of myocardial blood flow: The impact of surface coil-related field inhomogeneity and a comparison of methods for its correction. *Journal of Cardiovascular Magnetic Resonance*. 2015; 17:11. [PubMed: 25827156]

8. NIELLES-VALLESPIN S, KELLMAN P, HSU LY, ARAI AE. FLASH proton density imaging for improved surface coil intensity correction in quantitative and semi-quantitative SSFP perfusion cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*. 2015; 17:16. [PubMed: 25827180]
9. Gatehouse PD, Elington AG, Ablitt NA, Yang GZ, Pennell DJ, Firmin DN. Accurate assessment of the arterial input function during high-dose myocardial perfusion cardiovascular magnetic resonance. *Journal of Magnetic Resonance Imaging*. 2004; 20:39–45. [PubMed: 15221807]
10. Benovoy, M., Jacobs, M., Cheriet, F., Dahdah, N., Arai, AE., Hsu, LY. Automatic nonrigid motion correction for quantitative first-pass cardiac MR perfusion imaging. *Proceedings of International Symposium on Biomedical Imaging (ISBI); IEEE; 2015. p. 1588-1591.*
11. Brox T, Malik J. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2011; 33:500–513. [PubMed: 20714020]
12. Dalal, N., Triggs, B. Histograms of oriented gradients for human detection. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR); IEEE; 2005. p. 886-893.*
13. Horn BK, Schunck BG. "Determining optical flow": a retrospective. *Artificial Intelligence*. 1993; 59:81–7.
14. Brox, T., Bruhn, A., Papenberg, N., Weickert, J. High accuracy optical flow estimation based on a theory for warping. *Proceedings of European Conference on Computer Vision (ECCV); Berlin Heidelberg: Springer; 2004. p. 25-36.*
15. Staelin, C. Hewlett-Packard Company, Tech. Rep. HPL-2002-354R1. 2003. Parameter selection for support vector machines.
16. Xue, H., Zuehlsdorff, S., Kellman, P., et al. Unsupervised inline analysis of cardiac perfusion MRI. *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI); Berlin Heidelberg: Springer; 2009. p. 741-749.*
17. Hennemuth A, Seeger A, Friman O, et al. A comprehensive approach to the analysis of contrast enhanced cardiac MR images. *IEEE Transactions on Medical Imaging*. 2008; 27:1592–1610. [PubMed: 18955175]
18. Coltuc D, Bolon P, Chassery JM. Exact histogram specification. *IEEE Transactions on Image Processing*. 2006; 15:1143–1152. [PubMed: 16671295]
19. Jacobs, M., Gorbachev, M., Benovoy, M., Chang, LC., Arai, AE., Hsu, LY. Automated measurement of arterial input function in first-pass myocardial perfusion magnetic resonance images using independent component analysis. *Proceedings of International Symposium on Biomedical Imaging (ISBI); IEEE; 2015. p. 1332-1335.*
20. Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*. 1997; 16:187–98. [PubMed: 9101328]
21. Hamrouni, S., Rougon, N., Prêteux, F. Multi-feature information-theoretic image registration: application to groupwise registration of perfusion MRI exams. *Proceedings of International Symposium on Biomedical Imaging (ISBI); IEEE; 2011. p. 574-577.*
22. Adluru G, DiBella EV, Schabel MC. Model-based registration for dynamic cardiac perfusion MRI. *Journal of Magnetic Resonance Imaging*. 2006; 24:1062–1070. [PubMed: 17031818]
23. Milles J, van der Geest RJ, Jerosch-Herold M, Reiber JH, Lelieveldt BP. Fully automated motion correction in first-pass myocardial perfusion MR image sequences. *IEEE Transactions on Medical Imaging*. 2008; 27:1611–1621. [PubMed: 18955176]
24. Wollny G, Kellman P, Santos A, Ledesma-Carbayo MJ. Automatic motion compensation of free breathing acquired myocardial perfusion data by using independent component analysis. *Medical image analysis*. 2012; 16:1015–1028. [PubMed: 22465078]
25. Li C, Sun Y, Chai P. Pseudo ground truth based nonrigid registration of myocardial perfusion MRI. *Medical image analysis*. 2011; 15:449–459. [PubMed: 21376656]
26. Khalifa, F., Beache, GM., Firjani, A., Welch, KC., Gimel'farb, G., El-Baz, A. A new nonrigid registration approach for motion correction of cardiac first-pass perfusion MRI. *Proceedings of International Conference on Image Processing (ICIP); IEEE; 2012. p. 1665-1668.*

27. Khalifa, F., Beache, GM., Gimel'farb, G., El-Baz, A. A new nonrigid registration framework for improved visualization of transmural perfusion gradients on cardiac first-pass perfusion MRI. Proceedings of International Symposium on Biomedical Imaging (ISBI); IEEE; 2012. p. 828-831.
28. Tarroni, G., Patel, AR., Veronesi, F., et al. Proceedings of Computing in Cardiology. IEEE; 2010. MRI-based quantification of myocardial perfusion at rest and stress using automated frame-by-frame segmentation and non-rigid registration; p. 1-4.
29. Sethian JA. Level set methods and fast marching methods. Journal of Computing and Information Technology. 2003; 11:1–2.
30. Gupta SN, Solaiyappan M, Beache GM, Arai AE, Foo TK. Fast method for correcting image misregistration due to organ motion in time-series MRI data. Magnetic Resonance in Medicine. 2003; 49:506–514. [PubMed: 12594754]

Figure 1. Automatic Nonrigid Motion Correction Framework

Image processing pipeline of our proposed system using a robust large displacement optical flow (LDOF) approach for first-pass cardiac MR perfusion image series and auxiliary series.

Figure 2. Displacement Flow Fields Computation

Longitudinal (x) and transversal (y) displacement flow fields between pairwise sequential frames are computed using an optical flow (OF) formulation robust to large displacement (LDOF). This approach couples discrete feature point tracking computed from histograms of oriented gradients (HOG) with a continuous variational optical flow step. The discrete HOG matches (red vectors) are displayed and superimposed on the first cardiac MR image of the pair. The resulting dense pixel-wise x and y displacement flow fields are shown at bottom.

Figure 3. PD-to-Perfusion Image Registration

The median PD-weighted image and median T1-weighted perfusion image are computed from the series and processed to enhance gradient information and match photometric profiles. These are used as bridge images for image registration via the LDOF processing engine.

Figure 4. Comparison of PD-to-Perfusion Image Registration

Top two rows show examples of PD (red) to T1 perfusion (green) image registration with corresponding correlation coefficient (R^2) for three different sequence combinations. Percent change in correlation scores between the PD and T1 images after registration are shown in brackets next to the scores.

Figure 5. Comparison of Sequential Perfusion Frame Mean Consistency Curves
Temporal variability of the mean correlation coefficient (R^2 ; top row) and mutual information (MuInfo; bottom row) metrics computed across the T1-weighted perfusion series of the entire cohort for both full frame (left column) and heart region (right column).

Figure 6. Comparison of sequential perfusion frame consistency for breathing paradigms
 In both full-frame and heart region comparison, mean correlation coefficient (R^2 ; top row) and mutual information (MuInfo; bottom row) metrics are improved in all LDOF-MOCO series compared to the raw series (all $P < 0.001$) regardless of breathing paradigm. However, the breath-hold (BH) series shows higher R^2 and MuInfo measurements than the free-breathing (FB) series on both LDOF-MOCO and raw series. Each interval plot shows the mean $\pm$ SD.

Figure 7. Comparison of LV and Myocardial Perfusion Dynamics

Top three image rows show the change of contrast dynamic during a stress perfusion series as well as LV and myocardium sectors defined by different regions of interest (ROI). The ROI were traced on motion corrected images (LDOF-MOCO) and then backwarped to the raw perfusion images (Raw-BW) to generate the corresponding sectoral signal intensity (SI) curves as shown. As a comparison, SI curves of the raw perfusion images without ROI backwarping (Raw) were also generated to show the impact of image motion in different myocardial sectors. In contrast to large motion-related signal intensity fluctuations in the raw SI curves (sector 4 and sector 5), LDOF-MOCO and Raw-BW curves are much smoother and are almost indistinguishable from each other.

Table 1

Patient numbers according to acquisition characteristics.

Number of Patients	291		
Perfusion Image Sequence Type	FISP (255)	FLASH (36)	
AIF Image Sequence Type	FLASH (291)		
PD-Perfusion Image Configuration	FISP-FISP (100)	FLASH-FISP (155)	FLASH-FLASH (36)
Magnetic Field Strength	1.5T (201)	3.0T (90)	
Breathing Paradigm	Breath Hold (100)	Free Breathing (191)	
Pacing Condition	Rest (291)	Stress (291)	
Slice Location	Basal (291)	Mid-ventricular (291)	Apical (291)

Table 2

Perfusion Dynamics Comparison

a) Correlation coefficient (R^2) and b) normalized root-mean-square error (NRMSE) measurements of sector-wise time signal intensity curves in LDOF-MOCO vs. Raw-backwarped (Raw-BW), and Raw vs. Raw-BW image series. An increase of R^2 and a decrease of NRMSE are shown in the LDOF-MOCO series for the LV and all 6 sectors (S1 to S6). Results are shown as mean $\pm$ SD.

a)									
R^2	LV	S1	S2	S3	S4	S5	S6		
LDOF-MOCO vs. Raw-BW	$0.999 \pm 5.069E-4$	$0.997 \pm 4.700E-3$	$0.997 \pm 2.284E-3$	$0.996 \pm 4.727E-3$	$0.996 \pm 5.479E-3$	$0.995 \pm 7.78E-3$	$0.995 \pm 1.357E-2$		
Raw vs. Raw-BW	$0.997 \pm 2.811E-3$	$0.917 \pm 9.645E-2$	$0.923 \pm 1.005E-1$	$0.929 \pm 1.072E-1$	$0.855 \pm 1.817E-1$	$0.828 \pm 2.308E-1$	$0.908 \pm 1.895E-1$		
b)									
NRMSE	LV	S1	S2	S3	S4	S5	S6		
LDOF-MOCO vs. Raw-BW	$0.008 \pm 3.906E-3$	$0.059 \pm 2.943E-2$	$0.067 \pm 2.441E-2$	$0.044 \pm 1.939E-2$	$0.044 \pm 2.125E-2$	$0.064 \pm 2.895E-2$	$0.037 \pm 2.239E-2$		
Raw vs. Raw-BW	$0.022 \pm 1.248E-2$	$0.146 \pm 1.049E-1$	$0.131 \pm 9.307E-2$	$0.119 \pm 1.066E-1$	$0.289 \pm 2.518E-1$	$0.290 \pm 2.139E-1$	$0.148 \pm 1.337E-1$		