

©2010 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Digital Object Identifier: 10.1109/TMI.2010.2049270

Exploiting Quasiperiodicity in Motion Correction of Free-Breathing Myocardial Perfusion MRI

Gert Wollny, *Member, IEEE*, Maria J. Ledesma-Carbayo, *Member, IEEE*, Peter Kellman, and Andres Santos, *Senior Member, IEEE*,

Abstract—Free-breathing image acquisition is desirable in first-pass gadolinium-enhanced magnetic resonance imaging (MRI), but the breathing movements hinder the direct automatic analysis of the myocardial perfusion and qualitative readout by visual tracking. Nonrigid registration can be used to compensate for these movements but needs to deal with local contrast and intensity changes with time. We propose an automatic registration scheme that exploits the quasiperiodicity of free breathing to decouple movement from intensity change. First, we identify and register a subset of the images corresponding to the same phase of the breathing cycle. This registration step deals with small differences caused by movement but maintains the full range of intensity change. The remaining images are then registered to synthetic references that are created as a linear combination of images belonging to the already registered subset. Because of the quasiperiodic respiratory movement, the subset images are distributed evenly over time and, therefore, the synthetic references exhibit intensities similar to their corresponding unregistered images. Thus, this second registration step needs to account only for the movement. Validation experiments were performed on data obtained from six patients, three slices per patient, and the automatically obtained perfusion profiles were compared with profiles obtained by manually segmenting the myocardium. The results show that our automatic approach is well suited to compensate for the free-breathing movement and that it achieves a significant improvement in the average Pearson correlation coefficient between manually and automatically obtained perfusion profiles before (0.87 ± 0.18) and after (0.96 ± 0.09) registration.

Index Terms—image registration, heart, myocardial perfusion.

I. INTRODUCTION

FIRST-PASS gadolinium-enhanced, myocardial perfusion magnetic resonance imaging (MRI) is used to observe and quantify blood supply to the different regions of the myocardium. Ultimately, such observations can lead to the diagnosis of coronary artery disease, which causes narrowing of the coronary arteries leading to reduced blood supply to the myocardium.

A typical imaging protocol includes some precontrast baseline images, and the full cycle of contrast agent first enters the right ventricle (RV), then the left ventricle (LV), and finally, the agent perfuses the LV myocardium (Fig. 1). ECG-triggered

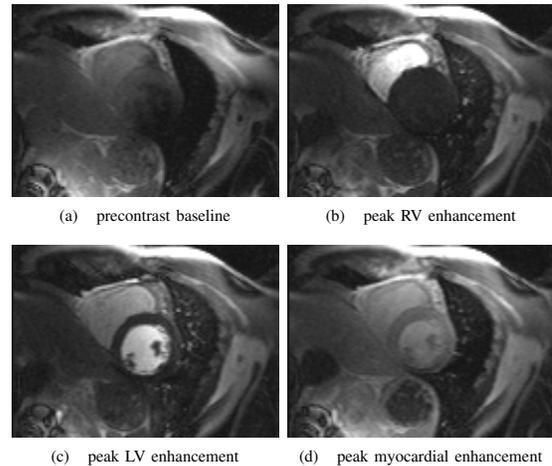


Figure 1. Images from a first-pass gadolinium-enhanced, myocardial perfusion MRI study.

images are acquired to ensure that the acquisition always takes place at the same contraction phase of the heart. Mistriggering may occur, resulting in nonrigid deformations of the heart.

The image sequence is acquired over 60 seconds and covers the perfusion from precontrast through the complete first pass, which is too long for most patients to hold their breath. Therefore, respiratory motion results in a nonrigid misalignment of the sequence of images through the whole acquisition. If the patient tries to hold the breath and fails, normally a sudden deep gasp occurs and the imaging sequence will not contain movement in the first part of the sequence but a strong movement will be present in the second part. If the patient is allowed to breathe freely, the respiratory motion is shallower and more repetitive, almost periodic. In both cases, proper alignment of the heart structures over the whole sequence is desired to enable an automatic analysis. To achieve such alignment, a matching procedure such as image registration may be performed.

A. State of the art

The major challenge in correcting the motion problem in contrast-enhanced perfusion imaging is that the local tissue contrast in the image sequence changes locally with time, especially in the region of interest, the left ventricular myocardium. In addition, although the imaging protocol triggers the acquisition at the same cardiac cycle phase, resulting in a nearly rigid representation of the heart, the breathing

G. Wollny, M.J. Ledesma-Carbayo, and A. Santos are with Ciber BBN, Spain, and the Group of Biomedical Imaging Technologies, Department of Electronic Engineering, ETSIT, Universidad Politécnic de Madrid, Spain, e-mail: {gert,mledesma,andres}@die.upm.es

P. Kellman is with the Laboratory of Cardiac Energetics, National Heart, Lung and Blood Institute, National Institutes of Health, DHHS, Bethesda, MD, USA.

Manuscript received ??????????, 2009; revised ??????????, 20XX.

movement occurs locally and the final sequence deforms nonrigidly.

Various registration methods have been proposed to achieve the alignment of myocardial perfusion heart images. For brevity, we will mention only a few publications; for further information, we point the reader to the references given in [1], [2].

Some approaches rely on rigid registration only and employ masks to restrict the registration to the area of the nearly rigid motion of the heart [3]–[8]. To overcome the problem of intensity change, one may optimize the similarity measures drawn from information theory, e.g., (normalized) *mutual information* (MI) [6], or (normalized) *cross correlation* (CC) [3], [5]. Other options include the use of contour masks obtained from gradient images and potential maps [7] or the removal of the area of high intensity change by masking [8].

Another proposed approach [4] is to identify three feature images as a vector base (baseline, peak RV enhancement, peak LV enhancement) using independent component analysis (ICA) of the intensity curve within the LV and RV. This vector base is then used to create a reference image for each time step by a weighted linear combination, and hopefully this image exhibits a similar intensity distribution to the corresponding original image to be registered. Image registration of the original image to the composed reference image is then achieved by a rigid transformation that minimizes CC. Because the motion may also affect the ICA base images, this approach was later extended to run the registration in two stages [2].

Because rigid registration alone does not account for the nonrigid deformations of the heart and requires some form of masking or feature extraction, other authors employ nonrigid registration optimizing MI to overcome the intensity change, e.g., [9]. MI is a global measure in the sense that it relies on a consistent material-intensity mapping over the whole image domain and does not account for the local intensity change. Methods have been proposed to minimize the effects of these local intensity variations on MI during registration [10], [11]. However, these methods are tailored only to accommodate slowly varying intensities that may result from field inhomogeneities or tissue degeneration and not the strong local changes resulting from a contrast agent passing through the heart ventricles and the myocardium. In addition, the evaluation of MI is quite expensive in computational terms.

To overcome these limitations, in our previous work [12], we used a combination of *normalized gradient fields* [13] and the *sum of squared differences* for registration. Because here only image pairs in temporal succession were registered, the registration method needed only to accommodate small differences. The final alignment of the whole image series to a common reference was achieved by accumulating the transformations resulting from the paired registrations. In [14] another serial registration scheme was proposed, that would first identify a key frame as the common reference, and then apply registration by optimizing CC between images in succession.

However, it is important to note that in serial schemes the accumulation of small registration errors may produce considerable errors in the overall alignment for distant time

frames with respect to the common reference.

B. Our contribution

Most of the methods described in the previous section are applied to data that are normally acquired during breath holding, and only some of the proposed methods dealt with data acquired during free breathing [8], [14], [15] or by simulation to exhibit the motion found in free-breathing data [2]. However, none of these methods was tailored specifically to free-breathing data. When a patient breathes freely, the breathing movement is quasiperiodic. As an example, Fig. 2 shows the movement of the manually tracked center of the LV in one example sequence.

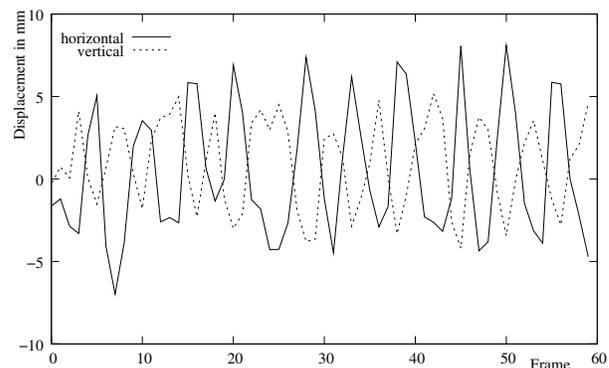


Figure 2. Displacements of the manually obtained center of the LV in a free-breathing acquisition. Note the quasiperiodicity of the breathing movement. In this graph, horizontal corresponds to a movement in the direction *septal* \leftrightarrow *lateral*, and vertical to *anterior* \leftrightarrow *inferior*.

Our contribution exploits this quasiperiodicity to correct the breathing movement. The registration scheme is divided in three stages starting with the automatic selection of a global reference, followed by the identification and registration of a subset of images that correspond to the same respiratory phase and, finally, the alignment of the rest of the images in the sequence with respect to the synthetic reference images of similar intensity distribution derived from the previously aligned subset. Following this scheme, the two elements that introduce differences in the sequence of images—the intensity change and the movement—are decoupled as follows. In the first registration step, only small movements have to be corrected, but the registration method has to deal with strong intensity changes. In the second registration step, the process deals with images that have similar intensity distributions, but the full amplitude of the breathing movement needs to be accounted for. The general image alignment framework uses nonrigid registration to account for deformations caused by the respiratory movement, possible out-of-plane movement, and to provide a fully automatic procedure that does not require user interaction. The nonrigid registration scheme uses a semilocal B-spline parametric transformation to optimize the similarity metric based on either normalized gradient fields or the sum of squared differences, or a combination of both as proposed in [12]. Robustness and efficiency are achieved using a multiresolution approach. This approach was presented partially with preliminary results in [16]. In

this article, we describe our completely automated method, which we validated thoroughly using data from six patients (18 sequences). We performed a number of experiments to determine the registration parameters that produced the best registration results. Intensity evolution curves from manually tracked myocardial segments were used as the gold standard.

In the remainder of the paper, we first give a short overview of the proposed nonrigid registration scheme, and describe our choice of a *normalized gradient fields* (NGF)-based similarity measure and our modifications. We then describe in detail our approach to identifying and using the respiratory phase-aligned image sequence to achieve motion correction. We then present our experiments and results to validate our motion-correction scheme by applying it to patient data. Finally, we discuss our findings and point to future work.

II. METHODS

A. Proposal overview

Given an image domain $\Omega \subset \mathbb{R}^d$ in the d-dimensional Euclidean space, an intensity range $\mathbb{V} \subset \mathbb{R}$, an image at time step $i \in \Theta$ is defined as mapping $I_i := I(\mathbf{x}, i) : \Omega \times \Theta \rightarrow \mathbb{V}$. Let us now consider an image sequence $\mathcal{I} := \{I_i | i \in \Theta\}$ obtained following a first-pass gadolinium MR acquisition protocol over N time steps $\Theta := \{1, 2, \dots, N\}$. Given this sequence, our goal is to obtain a motion-corrected sequence $\hat{\mathcal{I}}$. To correct for the existing motion, these images are registered directly or indirectly to a certain global reference I_{ref} . The proposed registration scheme is divided in three stages as outlined in Fig. 3. First, the global reference I_{ref} for the whole registration process and a subset of images \mathcal{I}' are selected automatically to correspond to the same breathing phase of the quasiperiodic motion. Second, this subsequence \mathcal{I}' is registered nonrigidly to the reference I_{ref} obtaining a registered subset $\hat{\mathcal{I}}'$. Because the images from the \mathcal{I}' sequence correspond to the same breathing phase, they are almost aligned, and nonrigid registration needs only to account for small deformations. However, the images in \mathcal{I}' will exhibit the full range of possible intensity change resulting from the contrast agent passing through the heart. In the third step, the remaining images $I_i \in \mathcal{I} \setminus \mathcal{I}'$ are registered. To do so, for each registration process of images $I_i \in \mathcal{I} \setminus \mathcal{I}'$, a reference image R_i is generated synthetically from the aligned subsequence $\hat{\mathcal{I}}'$ using a weighted linear combination. Hence, the reference images R_i will have a similar intensity distribution to that of the referring test image. Therefore, in this registration step, the intensity change has less influence than it would have without using the prealigned subset \mathcal{I}' , but the full range of the breathing movement still needs to be accounted for.

In the following subsections, the image registration framework is described in detail, and a thorough description of the different steps of the presented approach is given, including discussion of the method to select the prealigned subset, the global reference image, and the particularities of the different registration processes.

B. Image registration

Image registration can be defined as follows. Consider an image domain $\Omega \subset \mathbb{R}^d$ in the d-dimensional Euclidean space, a test image S , a reference image R , and a transformation of an image as a mapping $T : \Omega \rightarrow \Omega$ from a set of allowed transformations Ψ , and $S_T(\mathbf{x}) := S(T(\mathbf{x}))$ the test image deformed by applying transformation T . Then, the registration of S to R aims at finding a transformation T_{reg} according to:

$$T_{\text{reg}} := \min_{T \in \Psi} (F(S_T, R) + \kappa E(T)). \quad (1)$$

F measures the similarity between the deformed test image S_T and the reference, E ensures a steady and smooth transformation T , and κ is a weighting factor between smoothness and similarity. In nonrigid registration, the transformation T needs only to be neighborhood preserving, a restriction that is enforced by the selection of a proper term E . In our application, the cost function F is derived from a so-called voxel-similarity measure that takes into account the intensities of the whole image domain. As a consequence, the driving force of the registration is calculated directly from the given image data.

1) *Image similarity measures*: MR first-pass gadolinium perfusion studies show a strong local change in intensity because of the dynamics of the contrast agent, which results in local changes in the material-intensity mapping over the image domain and with time. Global similarity measures that rely on a consistent material-intensity mapping over the whole image domain are, therefore, not well suited for registering these images. Instead, a proper image similarity measure should rely on local features to drive the registration. One example of such a measure is the normalized gradient fields (NGF) proposed by Haber and Modersitzki [13].

Given an image $I(\mathbf{x})$ $\mathbf{x} \in \Omega$ and its noise level η , a measure ϵ for boundary ‘‘jumps’’ (locations with a high gradient) can be defined as:

$$\epsilon := \eta \frac{\int_{\Omega} |\nabla I(\mathbf{x})| d\mathbf{x}}{\int_{\Omega} d\mathbf{x}}. \quad (2)$$

Note, that the denominator in (2) represents the volume of Ω . With:

$$\|\nabla I(\mathbf{x})\|_{\epsilon} := \sqrt{\sum_{i=1}^d (\nabla I(\mathbf{x}))_i^2 + \epsilon^2}, \quad (3)$$

the NGF of an image I is defined as follows:

$$\mathbf{n}_{\epsilon}(I, \mathbf{x}) := \frac{\nabla I(\mathbf{x})}{\|\nabla I(\mathbf{x})\|_{\epsilon}}. \quad (4)$$

NGF-based similarity measures for the image registration of a test image S to a reference image R have been formulated based on either the scalar product $\langle \cdot, \cdot \rangle$ or the cross product \times of the vectors of the NGF [13]:

$$F_{\text{NGF}}^{(\times)}(S, R) := \frac{1}{2} \int_{\Omega} \|\mathbf{n}_{\epsilon}(R) \times \mathbf{n}_{\epsilon}(S)\|^2 d\mathbf{x}, \quad (5)$$

$$F_{\text{NGF}}^{(\cdot)}(S, R) := -\frac{1}{2} \int_{\Omega} \langle \mathbf{n}_{\epsilon}(R), \mathbf{n}_{\epsilon}(S) \rangle^2 d\mathbf{x}. \quad (6)$$

However, both similarity measures exhibit problems when applied to nonrigid registration. The formulation $F_{\text{NGF}}^{(\times)}$ (5) that

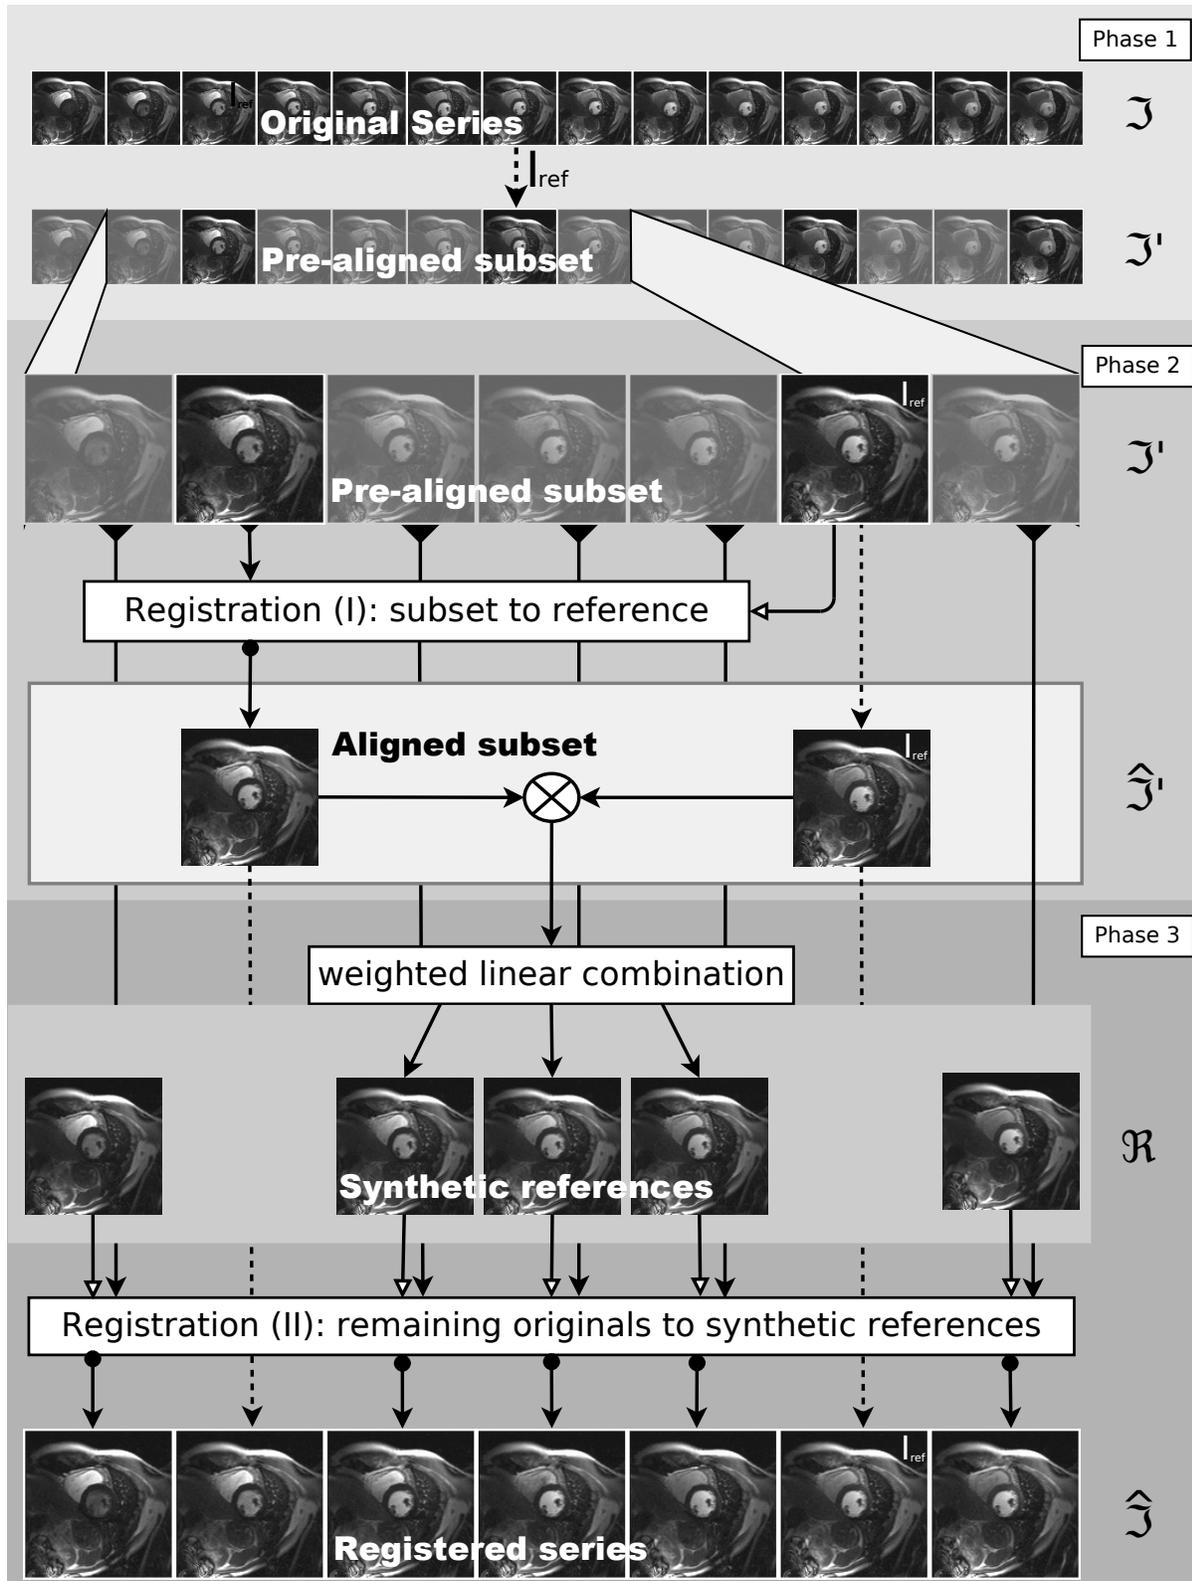


Figure 3. Scheme of the registration process. In the first phase, the overall reference I_{ref} and its prealigned subset \mathcal{J}' are selected, in the second phase \mathcal{J}' images are registered to obtain an aligned subset $\hat{\mathcal{J}}$, and in the third phase the remaining images $I_i \in \mathcal{J} \setminus \mathcal{J}'$ are registered using the corresponding synthetic references R_i generated from the aligned subset $\hat{\mathcal{J}}$.

is based on the cross product is zero at its minimum. However, $\mathbf{n}_\epsilon(R, \mathbf{x}) \times \mathbf{n}_\epsilon(S, \mathbf{x})$ is zero when $\mathbf{n}_\epsilon(R, \mathbf{x}) \parallel \mathbf{n}_\epsilon(S, \mathbf{x})$ (as desired) and when either $\mathbf{n}_\epsilon(R, \mathbf{x})$ or $\mathbf{n}_\epsilon(S, \mathbf{x})$ has a zero norm. Hence, $F_{\text{NGF}}^{(\times)}$ may have various global minima that do not coincide with the best solution for the image registration task. $F_{\text{NGF}}^{(\cdot)}$, on the other hand, has only one global minimum when $\mathbf{n}_\epsilon(R, \mathbf{x}) \parallel \mathbf{n}_\epsilon(S, \mathbf{x}) \forall \mathbf{x} \in \Omega$. Yet, even though the gradient $\nabla F_{\text{NGF}}^{(\cdot)}$ is analytically zero at this optimum for practical implementations of the gradient evaluation, e.g., by using finite differences, the gradient may evaluate to nonzero even if $S = R$, thereby making the optimization using gradient-based methods difficult.

To prevent these problems, the measure:

$$d(\mathbf{a}, \mathbf{b}) := \begin{cases} \mathbf{a} - \mathbf{b}, & \text{if } \langle \mathbf{a}, \mathbf{b} \rangle > 0, \\ \mathbf{a} + \mathbf{b}, & \text{otherwise} \end{cases} \quad (7)$$

$$F_{\text{NGF}}(S, R) := \frac{1}{2} \int_{\Omega} \langle \mathbf{n}_\epsilon(R), d(\mathbf{n}_\epsilon(R), \mathbf{n}_\epsilon(S)) \rangle^2 dx, \quad (8)$$

was proposed in [12], but this introduces an unsteady first derivative when $\mathbf{n}_\epsilon(R, \mathbf{x}) \perp \mathbf{n}_\epsilon(S, \mathbf{x})$, thereby making the registration unstable. Therefore, we refined this approach and propose to replace (7) by:

$$d^*(\mathbf{a}, \mathbf{b}) := \mathbf{a} - \frac{\langle \mathbf{a}, \mathbf{b} \rangle}{\|\mathbf{a}\| \|\mathbf{b}\|} \mathbf{b}, \quad (9)$$

resulting in:

$$F_{\text{NGF}}(S, R) := \frac{1}{2} \int_{\Omega} \left(\|\mathbf{n}_\epsilon(R)\|^2 - \frac{\langle \mathbf{n}_\epsilon(R), \mathbf{n}_\epsilon(S) \rangle^2}{\|\mathbf{n}_\epsilon(R)\| \|\mathbf{n}_\epsilon(S)\|} \right)^2 dx, \quad (10)$$

as another NGF-based cost function.

This cost function is always differentiable and both its evaluation and the evaluation of its derivatives are straightforward, making it easy to use for image registration. $F_{\text{NGF}}(S, R)|_{\mathbf{x}}$ is minimized when $\mathbf{n}_\epsilon(R, \mathbf{x}) \parallel \mathbf{n}_\epsilon(S, \mathbf{x})$. In the optimal case, $S = R$ the cost function and its first-order derivatives are zero, and their evaluation is numerically stable. In addition, as outlined by Haber and Modersitzki [13], NGF-based similarity measures have less local minima than e.g., MI-based measures, are easier to implement, and have a low computational complexity. However, in homogeneous areas of the reference image, where $\mathbf{n}_\epsilon(R, \mathbf{x})$ has a zero norm, $F_{\text{NGF}}(S, R)$ also has a zero value and a zero gradient. Therefore, the measure is best applied using a reference image with many gradients. In addition, if large deformations are to be accommodated and the images contain large homogeneous areas, nonrigid registration can generally not be achieved using an NGF-based measure only, and improvements are required to obtain a robust measure.

Pluim et al. [17] proposed combining MI and gradient information to achieve better registration results. In our application, most of the imaged area, i.e., everything outside the heart, will exhibit similar intensities in the test and the reference images. In addition, as outlined below, in the final step of our series registration, we will create synthetic references that will exhibit a similar material-intensity mapping as the corresponding test image over the whole image domain. Therefore, as discussed in [12], we will also consider combining the

proposed NGF-based measure (10) with the sum of squared differences (SSD) as a registration criterion in parts of the algorithm:

$$F_{\text{SSD}}(S, R) := \frac{1}{2} \int_{\Omega} (S(\mathbf{x}) - R(\mathbf{x}))^2 dx. \quad (11)$$

This combined similarity measure is defined as:

$$F_{\text{Sum}} := \lambda_{\text{ngf}} F_{\text{NGF}} + \lambda_{\text{ssd}} F_{\text{SSD}}, \quad (12)$$

with λ_{ngf} and λ_{ssd} to weight the two measures appropriately. As the validation will show, it is possible, to actually use F_{NGF} as the sole measure in one part of the algorithm and F_{SSD} in the other.

2) *Transformation space, regularization & optimization:* In our registration approach, the transformation is formulated in terms of B-spline basis functions located on a regular grid of control points $\mathbf{i} \in \mathbf{J} \subset \mathbb{Z}^2$ [18]:

$$T(\mathbf{x}) := \mathbf{x} + \sum_{\mathbf{i} \in \mathbf{J}} \mathbf{c}_i \beta^D \left(\frac{\mathbf{x}}{h} - \mathbf{i} \right), \quad (13)$$

where $\beta^D(\mathbf{x})$ is the tensor product of centered uniform B-splines of degree D , and $\mathbf{c}_i \in \mathbb{R}^2$ are the vector value-weighting coefficients. The parameter h governs the grid knot spacing and, therefore, the number of registration parameters and partially the smoothness of the solution.

To enforce smoothness of the transformation and to improve the stability of the solution in homogeneous areas, our nonrigid registration approach employs a regularization that is based on the separate norms of the second derivative of each of the deformation components [19] as energy term $E(T)$ in (1):

$$E(T) := \int_{\Omega} \sum_i^d \sum_j^d \left\| \frac{\partial^2}{\partial x_i \partial x_j} T(\mathbf{x}) \right\|^2 dx. \quad (14)$$

As given in eq. (1) the regularization term is weighted against the similarity measure by a factor κ .

To solve the registration problem by optimizing (1), generally any gradient-based optimizer could be used. Specifically, we considered a *gradient descent optimizer with quadratic step size estimation* (GD) [18] and a variant of the *Levenberg–Marquardt* (LM) optimizer [20] using a Hessian approximation as outlined in [21]. To reduce the computational load, only a maximum number n_{max} of the n_{total} parameters of the deformation function is updated during every single iteration. Here, the n_{max} B-spline coefficients corresponding to the steepest gradient values of $\nabla (F(S_T, R) + \kappa E(T))$ are chosen.

Speed and robustness are improved by the use of a multiresolution approach in both the image and the transformation space. Its implementation is based on [18]. The multiresolution strategy makes use of a pyramid of subsampled images that are optimal in the L2-sense, taking advantage of the spline representation [22].

C. The automatic motion-correction algorithm

A scheme of the automatic motion-correction algorithm as summarized in subsection II-A is shown in Fig. 3. In this section, we describe the three stages of the algorithm in more

detail: first, the selection of the global reference I_{ref} and the subset of prealigned images \mathcal{I} ; second, the registration process of the prealigned subset; and finally the alignment of the rest of the images in the sequence.

1) *Selecting a global reference and the prealigned subset:* The motion-correction procedure requires that all the images from the sequence be transformed to the same spatial reference framework. This process is accomplished by registering all images directly or indirectly to a certain global reference image I_{ref} .

The selection of the global reference I_{ref} is based on finding an image with high contrast to ensure that NGF-based similarity measures generally give a good response and on finding an image that provides strong periodicity in its sequence-similarity profile:

$$\mathfrak{F} := \{F_{\text{NGF}}(I_i, I_{\text{ref}}) | I_i \in \mathcal{I}\}. \quad (15)$$

The application of sequence-similarity profiles for selecting a good reference has been used previously in [23]. Because the LV does not exhibit any usable gradients before the contrast agent passes through it, the search range for I_{ref} will be restricted to the sequence interval starting at the LV enhancement peak.

Once the global reference is selected, its corresponding sequence-similarity profile \mathfrak{F} comprises a representation of the quasiperiodic movement and allows the selection of images corresponding approximately to the same breathing phase. To perform this step automatically, an extreme phase of the breathing movement has been selected.

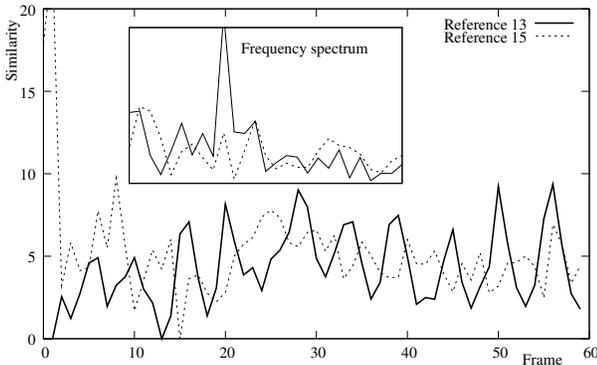


Figure 4. Sequence-similarity profiles with respect to two different global reference image candidates and their respective overall frequency spectra. Note, that reference image 13 is a better candidate because its sequence-similarity profile exhibits stronger periodicity, which is reflected by the high energy value in the frequency spectrum.

In more detail, to obtain I_{ref} , for each image $I_i \in \mathcal{I}$, we first estimate the standard deviation $\sigma(I_i)$ of its pixel intensities as an indicator of the image's contrast. We next select a subset of the images with the highest intensity standard deviations; these comprise the candidate set \mathcal{I}^* for the overall reference image. Then, for all $I_k \in \mathcal{I}^*$, we evaluate the sequence-similarity profiles $f_k = \{F_{\text{NGF}}(I_i, I_k) | I_i \in \mathcal{I}\}$. The profile f_k that exhibits the strongest periodicity is identified based on its Fourier transform. Because f_k comprises real numbers only, its Fourier transform is symmetric and we only need to take into account the positive frequency coefficient. Here,

we select the profile $\mathfrak{F} = f_{k^*}$ with index k^* that exhibits the largest absolute value in the positive frequency coefficients of its Fourier transform. It corresponds to the profile where most of the signal energy is concentrated in one frequency, which in turn points to a strong periodic component (Fig 4). Therefore, our global reference image will be $I_{\text{ref}} = I_{k^*}$.

Next, we choose all $I_i \in \mathcal{I}$ for which $F_{\text{NGF}}(I_i, I_{\text{ref}}) \in \mathfrak{F}$ exhibits a local minimum:

$$F_{\text{NGF}}(I_i, I_{k^*}) < F_{\text{NGF}}(I_{i \pm m}, I_{k^*}) \forall m = 1, 2. \quad (16)$$

These images form the phase-aligned subsequence of the quasiperiodic motion.

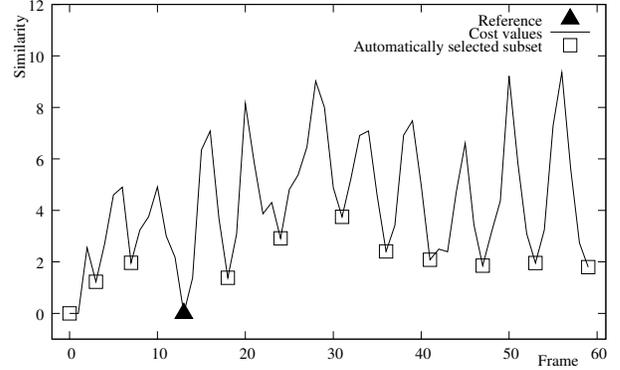


Figure 5. Series reference and prealigned subset obtained by the first step of the algorithm. Note the even distribution of the time points selected for the subset.

Because of the quasiperiodic nature of the breathing movement, the images of this subset will be evenly distributed over the time series (Fig. 5). To ease the implementation of the last registration step, we also add the first and the last image of the series to this set to obtain the subset $\mathcal{I}' \subset \mathcal{I}$ for the initial registration step.

2) *Aligning the prealigned subset:* In this step, we register all images $I \in \mathcal{I}'$ to I_{ref} following the registration method described in subsection II-B. Because these images will exhibit strong differences in their intensity, using a combined similarity measure that also employs F_{SSD} (11) would be counterproductive. In addition, because these images all stem from the same breathing phase and are, therefore, already well aligned, using F_{NGF} (10) as the registration criterion should suffice. As a result of this step, we obtain a phase-aligned-registered subset $\hat{\mathcal{I}}'$ of the original image series \mathcal{I} .

3) *Aligning the remaining images:* In the last processing step, we pick up on the idea of Milles et al. [2] to create synthetic reference images to overcome the intensity change. For each unregistered image $I_i \in \mathcal{I} \setminus \mathcal{I}'$, we first select the two registered images from the phase-aligned subset $I_{i^-}, I_{i^+} \in \hat{\mathcal{I}}'$ that enclose I_i in the temporal continuity:

$$i^- < i^+ \wedge i \in (i^-, i^+) \wedge (i^+ - i^-) \rightarrow \min. \quad (17)$$

Because we added the first and the last image of the series to the subset \mathcal{I}' , these two images I_{i^-}, I_{i^+} always exist. Then, a reference image R_i is created as a weighted linear combination:

$$R_i(\mathbf{x}) := \frac{i - i^-}{i^+ - i^-} I_{i^+}(\mathbf{x}) + \frac{i^+ - i}{i^+ - i^-} I_{i^-}(\mathbf{x}). \quad (18)$$

Because of the quasiperiodic motion, one of the preregistered images can be found about every five frames. Therefore, the images used to create the weighted combination are very close to each other, and for most of the series, linear interpolation should suffice to create reference images R_i that will exhibit a similar intensity distribution like I_i , the image to be registered to.

Compared with Milles et al. [2], who reported that the synthetic reference images created from the ICAs were blurred and thus needed a multipass scheme, in our approach, the interpolated images used to create the reference images are already registered and, therefore, no multipass scheme is required.

In the optimal case, creating synthetic reference images make it possible to use F_{SSD} as the only registration criterion. However, at the beginning of the image series, when the contrast agent passes through the right and left heart ventricles, the intensity change is not modeled well by the linear interpolation, and using F_{Sum} (12) as the registration criterion may be a better choice. Therefore, we investigated the best weighting between F_{SSD} and F_{NGF} in F_{Sum} during validation.

D. Approaching validation

Various methods have been proposed for validation including tracking the LV center [2], [24]; correlation and/or the mean square error between manually obtained intensity profiles and automatically obtained profiles [2], [15]; measuring the per-frame intensity variation within the myocardium [2]; counting the number of false-positive/false-negative pixels of masks of the myocardium obtained from a reference in the registered sequence, with respect to a manual segmentation [24]; and the myocardial boundary error defined as the minimum distance between the myocardial contours obtained from the hand segmentation and the registered slices [24].

Considering these validation techniques, tracking the LV center is not an appropriate approach to validate our methodology because this method does not account for differences in rotation and nonrigid deformation and, therefore, we did not use it for validation purposes.

Segmentation based methods could be also used. However, the accurate tracking of corresponding anatomical features through time in perfusion studies is not an easy and repeatable task. Firstly, at the beginning of the series the myocardium and the left ventricular cavity exhibit the same intensities. Secondly, the papillary muscles and the myocardium often also exhibit the same intensities. In order to assess the reliability of segmentations we repeated the series segmentation in two slices by two observers obtaining a mean Hausdorff distance of 3.6mm and a maximal Hausdorff distance of 5.7mm between the segmented boundaries of the myocardium. These errors in the segmentation would show up as misregistrations when segmentation based statistics are used that analyze the myocardial boundary error or false-positive/false-negative pixels. Therefore, segmentation based methods were not used to assess the quality of the correspondence of the boundaries or the section overlap in the registered series.

However, when assessing the performance of the intended task – perfusion analysis through intensity profiles – we

confirmed that the influence of the segmentation differences was quite small, and hence, using intensity profiles obtained from manually segmented series as a gold standard for this comparison is still an effective approach for validation. Therefore, we generally provide methods to assess the quality of the image registration by analyzing time–intensity curves.

For a first automatic assessment of the registration quality, we measure how well the structures are aligned by focusing on the pixel-wise intensity curves over time. Here, the second-order intensity derivative is of interest. In a registered series, only the perfusion by the contrast agent will induce an intensity change, and its temporal gradient changes rapidly thereby producing outliers when the contrast agent enters, reaches saturation, or leaves the area corresponding to the pixel. By contrast, in an unregistered series, the temporal gradient may also change rapidly at tissue boundaries. Therefore, we propose to use the median of the absolute values of the second-order derivative of the pixel intensity over time as quality measure. It is expected to be much lower in a registered series than in its corresponding unregistered series. To obtain this gradient-based measure, we first reduce the interacquisition noise by smoothing the pixel-wise intensities over time by applying a Gaussian filter G of width 5:

$$I_{G_5}(\mathbf{x}, i) := \sum_{k=-2}^2 G_5(k)I(\mathbf{x}, i - k). \quad (19)$$

The smoothed image series $\mathfrak{J}_{G_5} = \{I_{G_5}(i) | i \in \Theta\}$ is then used to evaluate the median second-order derivative image pixel-wise:

$$I_{\mathfrak{D}^2}(\mathbf{x}) := \text{median}_{i \in \Theta} \left(\left. \frac{d^2 I_{G_5}(\mathbf{x}, t)}{dt^2} \right|_{t=i} \right). \quad (20)$$

For every sequence, we restrict the evaluation of these images to the region around the LV, which was obtained by enlarging the LV bounding box obtained from the manual segmentation of the LV myocardium that is required for the analysis of the intensity profiles described below. Finally, to obtain a single measure of image alignment, we evaluate the mean intensity

$$\mathfrak{D}^2 := \frac{\int_{\Omega} I_{\mathfrak{D}^2}(\mathbf{x}) d\mathbf{x}}{\int_{\Omega} d\mathbf{x}}, \quad (21)$$

of this median time-derivative image $I_{\mathfrak{D}^2}$. For a better interpretation of this value \mathfrak{D}^2 , we compared the results evaluated from the free-breathing acquired series before and after registration with values obtained from the nonmoving initial images of a series of three slices acquired during breath holding.

However, the above measure is useful only for preliminary assessment of how well gray-scale values are aligned by the registration method. For a more thorough validation, we turn to the comparison of perfusion profiles obtained from the registered image series with profiles obtained by manually segmenting the myocardium in the original sequences. Therefore, as a prerequisite, in all images of every sequence, the myocardium of the LV was segmented manually into six transmural segments $\mathbb{S} = \{s_1, s_2, \dots, s_6\}$ (Fig. 6) counting clock-wise and starting at the LV and RV anterior intersection.

The reference intensity profiles $K_{\text{ref}}^{(s)}$ of the segments $s \in \mathbb{S}$ over the image series were obtained by evaluating the average

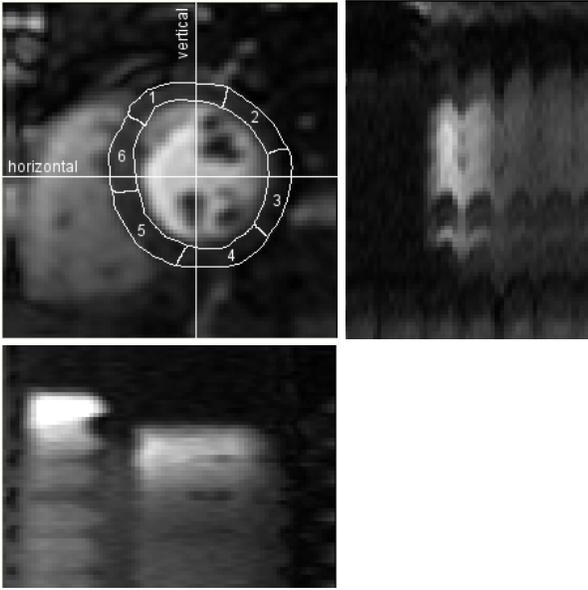


Figure 6. Segmentation of the LV myocardium into six segments and horizontal and vertical profiles of the original image series before motion correction. Note that, in the profiles, the breathing movement is evident in the zigzag lines that visualize the movement at tissue boundaries.

intensities in these regions and plotting these over the time of the sequence. By using only the segmentation of the reference image I_{ref} as a mask to evaluate the intensities in all registered images, the test intensity profiles $K_{\text{reg}}^{(s)}$ were obtained. The intensity profiles $K_{\text{orig}}^{(s)}$ for the unregistered series were evaluated similarly.

To make it possible to compare the sequences of different image series, the intensity curves K were normalized based on the reference intensity range $[v_{\min}, v_{\max}]$, with $v_{\min} := \min_{s \in \mathbb{S}, t \in \Theta} K_{\text{ref}}^{(s)}(t)$ and $v_{\max} := \max_{s \in \mathbb{S}, t \in \Theta} K_{\text{ref}}^{(s)}(t)$:

$$\hat{K} := \left\{ \frac{v - v_{\min}}{v_{\max} - v_{\min}} \mid v \in K \right\}. \quad (22)$$

To quantify the effect of the motion correction, the quotient of the average absolute distance between the registered and reference curves, $\hat{K}_{\text{reg}}^{(s)}$ and $\hat{K}_{\text{ref}}^{(s)}$ and the average absolute distance between the unregistered and reference curves, $\hat{K}_{\text{orig}}^{(s)}$ and $\hat{K}_{\text{ref}}^{(s)}$, are evaluated, producing the value Q_s as a quality measure of the registration with respect to each section $s \in \mathbb{S}$:

$$Q_s := \frac{\sum_{t \in \Theta} |\hat{K}_{\text{reg}}^{(s)}(t) - \hat{K}_{\text{ref}}^{(s)}(t)|}{\sum_{t \in \Theta} |\hat{K}_{\text{orig}}^{(s)}(t) - \hat{K}_{\text{ref}}^{(s)}(t)|}. \quad (23)$$

As a result, we obtain $Q_s > 0$ with smaller values indicating better registration.

We also evaluated the squared Pearson correlation coefficient R^2 of the manually obtained intensity profiles and the corresponding postregistration profiles. The range of this coefficient is $R^2 \in [0, 1]$ with higher values indicating a better correlation between the data sets. However, because the correlation accounts only for linear dependencies, neither an error in scaling nor an intensity shift is caught by this measure.

Finally, we consider the average standard deviation of the intensity in the six segments s_i of the myocardium $\sigma_{s_i, t} :=$

$\sum_{t \in \Theta} \sigma(s_i)$ as obtained by using the myocardial mask of the reference image I_{ref} to the motion-corrected sequence. Because the intensity in these regions is relatively homogeneous with a good alignment of the images, only noise and the intensity differences caused by disease should influence this value. In particular, in the first part of the perfusion image series, when the contrast agent passes through the RV and LV, this approach makes it possible to assess the registration quality without comparing it with the manual segmentation. Any misalignment between the section mask of the reference image and the corresponding section of the analyzed series frame will add pixels of the ventricular cavity to one or more segments, increasing the intensity range, and hence its standard deviation. Therefore, with proper alignment, this value will decrease.

For all measures, the statistical measures *average*, *standard deviation*, *median*, *minimum*, and *maximum* were evaluated and used to compare the results.

E. Comparison with rigid registration

To assess the differences between applying rigid registration and nonrigid registration for motion compensation, we used two approaches. In one approach, we implemented the ICA-based registration scheme [2]. In the other approach, we implemented our motion-compensation scheme by replacing the nonrigid registration by an implementation for rigid registration [25] that makes use of The Insight Toolkit (ITK) [26]. As discussed later in section III-D3, rigid registration including rotation turned out to be unstable. Therefore, the rigid movement was restricted to translation only. Optimization was achieved by using the regular gradient descent optimizer provided in the ITK library. This approach was run on cropped images series containing only the LV.

III. EXPERIMENTS AND RESULTS

A. Image data

First-pass contrast-enhanced myocardial perfusion imaging data sets were acquired and processed for six subjects under clinical research protocols approved by the Institutional Review Boards of the National Heart, Lung, and Blood Institute and Suburban Hospital. The patients provided written informed consent, and the analysis was approved by the NIH Office of Human Subject Research. Two distinct pulse sequences were used for image acquisition: a hybrid GRE-EPI sequence and a true-FISP sequence. Both sequences were ECG-triggered and used 90-degree-saturation recovery imaging of several slices per R–R interval acquired for 60 heartbeats. The pulse sequence parameters for the true-FISP sequence were 50-degree readout flip angle, 975 Hz/pixel bandwidth, TE/TR/TI = 1.3/2.8/90 ms, 128×88 matrix, 6 mm slice thickness. The GRE-EPI sequence parameters were 25-degree readout flip angle, echo train length = 4, 1500 Hz/pixel bandwidth, TE/TR/TI = 1.1/6.5/70 ms, 128×96 matrix, 8 mm slice thickness. The spatial resolution was about $2.8 \text{ mm} \times 3.5 \text{ mm}$. Parallel imaging using the TSENSE [27] method with acceleration factor = 2 was used to improve temporal resolution and spatial coverage.

Motion correction was performed for three short-axis slices of the six patients $\{A, B, \dots, F\}$ (18 distinct slices total) covering different levels of the LV myocardium (basal, mid, and apical levels). For all but one patient, a single dose of contrast agent (Gd-DTPA, 0.1 mmol/kg) was administered at 2.5 ml/s, followed by saline flush. For Patient C the dose was 5 ml/s instead. Images were generally reconstructed to a final matrix size of 256×192 (3/4 phase FOV) using zero filling for interpolation. In the case of Patient A, the final matrix size was 256×196 , and in the case of Patient B, 128×128 . The data for Patient D were acquired with a shallow breathing protocol, and all other series were acquired using free breathing. Mistriggering occurred in only one time step in the data set of Patient A. Of the slices obtained, the first two were baseline images taken with a proton density-weighted protocol and, therefore, exhibited a different intensity distribution compared with the rest of the series. For that reason, these images were omitted from the analysis.

B. Parameters

In the first step of the algorithm, we set the search range for the global reference to $[20, N]$. The nonrigid registration approach used uniform cubic B-Splines for the transformation with Dirichlet boundary conditions. To investigate the best parameter combination, we tested the performance of the algorithm in terms of the quality measure Q using different parameter values for the knot spacing $h \in \{5, 7, 10, 16, 22, 26\}$ given in pixels for the B-spline grid, and the weight $\kappa \in \{1.0, 5.0, 10.0, 20.0\}$ of the regularization term. For the last part of the registration scheme, i.e., the registration of the images not corresponding to the respiratory phase-aligned subset, we used $(\lambda_{ngf}, \lambda_{ssd}) \in \{(0, 1), (1, 1), (1, \frac{1}{2}), (1, \frac{1}{10}), (1, 0)\}$ as weights between the F_{NGF} and F_{SSD} . Estimating the noise level of images is difficult, and we only approximated it by using the standard deviation of the intensity gradient norm: $\eta \approx \sigma(\|\nabla I(x)\|)$.

After some initial experiments, the number of multiresolution levels l was fixed to 3.

C. Implementation and run time

The analysis software was implemented as a mix of Python for selecting the reference image and the prealigned subset, and a set of C++ programs to run the registration and to evaluate similarity measures. The software was run on a Linux workstation equipped with an Intel Pentium Core2 6600 and 4 GB of working memory. The complete run time of the nonrigid registration scheme of a series with 60 images of size 256×192 pixels was about 3 min. This time could be reduced easily by running the registration in parallel.

D. Results

Following the scheme described above, a good reduction of the breathing motion was achieved for all image data sets. A first visual assessment of the registration results was obtained by observing videos of the registered images and by analyzing the time profiles of the image series (Fig. 7). Examples of these videos are available online [28].

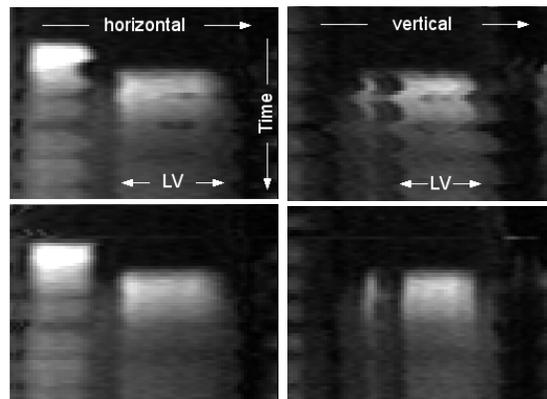


Figure 7. Profiles (Patient E) cut through the time stack at the locations indicated in Fig. 6. In the original series (upper row), the breathing movement is clearly visible. In the registered series (lower row), this movement is nearly completely removed.

Analyzing the $I_{\mathcal{D}^2}$ images (Fig. 8) before and after registration confirms these results. A visual inspection shows a clear reduction in the median second-order time-intensity gradient after registration in the areas of the moving heart (Fig. 8). Table I shows the summary over all time series of the analysis of the \mathcal{D}^2 registration quality measure corresponding to the mean time-intensity deviation over the region of interest, which results in a clear improvement. The results after registration are similar to the ones obtained from nonmoving image series, indicating that the remaining gradients over time result mainly from noise.

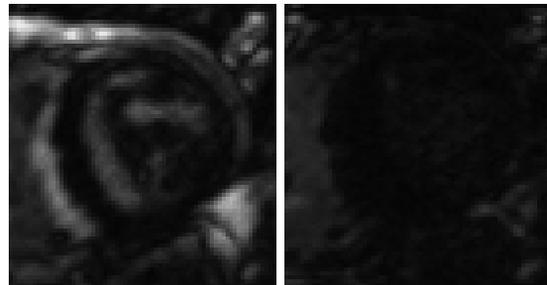


Figure 8. $I_{\mathcal{D}^2}$ image (Patient A): Note the high values before registration (left) that correspond to the intensity changes at the tissue boundaries resulting from movement and that are represented by the zigzag line in the profile images (Fig. 7). The absence of these high values after registration marks the successful alignment (right).

Table I
GENERAL REGISTRATION QUALITY BASED ON \mathcal{D}^2

	Mean	SD	Median	Min	Max
unregistered	2.55	1.25	2.35	1.05	5.12
registered	1.32	0.35	1.29	0.77	1.99
nonmoving reference	1.66	0.08	1.65	1.58	1.75

Statistical analysis of the registration quality based on the median of the second-order derivative of the pixel-wise intensity over time (\mathcal{D}^2) (a smaller number is better).

The registration results were also confirmed by observing the signal intensity time courses in different regions of the myocardium and by comparing these with the manually segmented intensity curves. Fig. 9 shows the intensity curves of the corresponding sections of the LV myocardium represented

in Fig. 6. In the example given (Patient A, basal), a clear improvement is observed for all regions, enabling further automatic analysis of the myocardial perfusion and easier visual inspection.

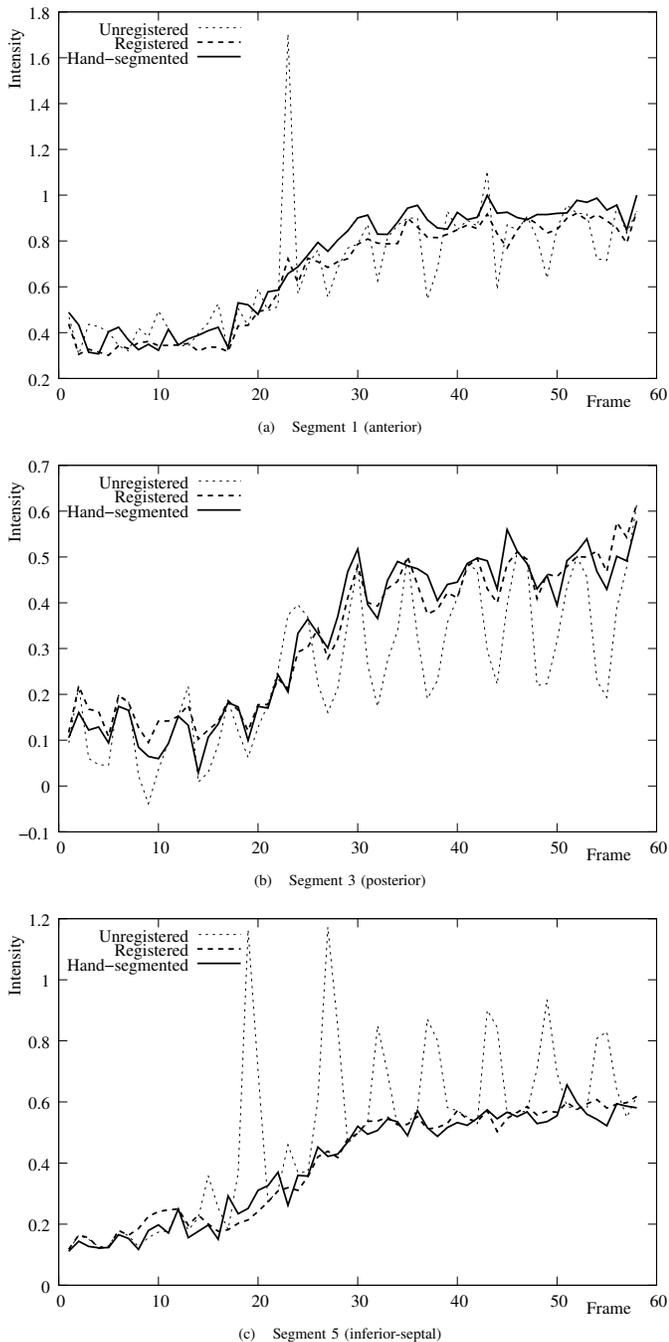


Figure 9. Intensity curves before and after registration compared with the manually obtained curves in three segments of the basal slice of Patient A. Note the periodic intensity variations in the unregistered series caused by the breathing movement and how the registered series more closely resembles the manually obtained intensity curves. The spike in the intensity curve in (a) is caused by an error in triggering, which is amplified by the breathing movement but damped in the other sections. As one can see, the nonrigid registration can also compensate for this error.

1) *Optimal parameters set:* Using F_{SSD} as the only criterion in the second phase of the registration, employing $l = 3$ multiresolution levels, a value of $\kappa = 15.0$ as the regularization

weight, a knot spacing of $h = 5$ pixels, and the modified Levenberg–Marquardt optimizer results in a good registration in all cases with a slight deviation for patient C that will be discussed later. This parameter set corresponds to an averaged quality measure of $Q = 0.62$, an averaged correlation coefficient of $R^2 = 0.96$, and an average intensity variation of $\sigma = 0.53$ as given in Table II. Nevertheless, based on our statistical analysis of the parameter set given above, using regularizer weights of $\kappa \in [5.0, 20.0]$, $(\lambda_{ngf}, \lambda_{ssd}) \in \{(0, 1), (1, 1)\}$, and $h \in [5, 16]$, all give comparable registration results, i.e., according to the Welch two-sample t test ($p = 0.05$) the average values for Q , σ_s , and R^2 do not differ significantly. As for the comparison between the application of LM and GD as optimizers and as suggested in [18], the modified implementation of the LM generally needed fewer iterations to converge, and it resulted in a slightly higher registration accuracy. The gradient descent, on the other hand, needed less time and less working memory to achieve similar registration results.

For patient C the dose of contrast agent used was twice as high as for the other patients, this resulted in a stronger intensity change before LV peak enhancement, and therefore, the linear interpolation used to create synthetic reference images in the last registration step didn't model the intensity change well. As a result, the registration resulted in artifacts at the beginning of the series that could be compensated by using higher values for the regularization weight and a larger B-spline knot spacing at the cost of more residual movement after registration.

Table II
MEASURES OF REGISTRATION QUALITY

	$\lambda_{ngf}, \lambda_{ssd}$	Mean	SD	Median	Min	Max
Q smaller is better	0, 1	0.62	0.26	0.62	0.14	1.38
R^2	unreg	0.87	0.18	0.94	-0.17	1.00
larger is better	0, 1	0.96	0.09	0.98	0.09	1.00
σ_s	unreg	0.66	0.34	0.58	0.15	2.33
smaller is better	manual	0.45	0.14	0.45	0.17	0.84
	0, 1	0.53	0.23	0.49	0.14	1.26

The registration quality Q , the correlation R^2 , and the average intensity variation σ , analyzed for all 18 data sets (108 myocardial segments) of the full series F_{SSD} is the only registration criterion in the second registration step. The remaining registration parameters are fixed to $l = 3$ multiresolution levels, a B-spline knot spacing of $h = 5$, and a regularizer weight of $\kappa = 15.0$.

2) *Slice location-based analysis:* One may also compare the results for the different slices with respect to the position relative to the heart. The results shown in Table III suggest that the registration at the apex performs worse than that in the other imaging planes. This is not surprising because at the apex, the out-of-plane motion is usually greater than in the other planes, and this motion cannot be corrected by a 2D registration algorithm.

The intensity–time curves for the hand-segmented, unregistered, and registered series, as well as the quality measures for each slice are given in a separate technical report [28]. At request we also provide the raw time–intensity data and the scripts used to obtain the reported results.

3) *Comparison with rigid registration-based methods:* Comparing the registration results with other published results is difficult because different data sets have been used, and the

Table III
SEPARATED BY SLICE LOCATION

	Q	R^2		σ		
		unreg	reg	unreg	manual	reg
apex	0.65	0.87	0.95	0.50	0.37	0.40
mid	0.56	0.87	0.98	0.51	0.35	0.41
basal	0.53	0.85	0.97	0.72	0.54	0.60

The registration quality Q analyzed by slice location for registration parameters fixed to $l = 3$ multiresolution levels, a B-spline knot spacing of $h = 5$, and a regularizer weight of $\kappa = 15.0$ and $\lambda_{ngf} = 0, \lambda_{ssd} = 1$ is the last registration phase.

exact methodologies used to obtain the statistical results are not described thoroughly. Therefore, selecting a measure to compare results fairly is almost impossible. Milles et al. [2] used various measures to describe registration accuracy, among them Pearson’s correlation coefficient R^2 between manually obtained intensity curves and automatically derived ones. They reported R^2 values of 0.88 ± 0.16 before and 0.92 ± 0.10 after registration [2].

Before registration our image data is comparable to theirs as it exhibit a similar mean correlation R^2 of 0.87 ± 0.18 . After registration, the obtained intensity curves correlate better on average with a lower standard deviation, suggesting better motion correction with our method, i.e. we obtained an average R^2 of 0.96 ± 0.09 (Table II). This is not surprising because in terms of registration accuracy, nonrigid motion correction should outperform rigid registration.

To validate this finding, we ran the ICA-based registration approach as reported in [2]. For data with initial breath holding, this approach worked as expected, which confirmed that we had implemented the method properly. However, in free-breathing data, the periodic movement made the estimation of the RV and LV peak enhancement time point unstable, and thus automatic mask creation was difficult. Applying the ICA-based method to the cropped images, thereby circumventing the mask creation, also did not produce reliable registration results.

We then modified our scheme presented here by replacing the nonrigid registration with rigid registration and then applied the method to cropped images containing only the LV region. However, the use of the ITK framework for rigid registration required a fine-tuning of the scaling between the rotational and translational parameters that highly influenced the registration result and made the procedure unstable. By further restricting the transformation space Ψ to translations only, this parameter tuning was no longer needed and a more stable rigid registration could be achieved. Note, that the rigid registration method used in [2] also optimizes only the translation.

As a result, a reasonably good and stable registration could be obtained for 16 of the 18 slices, but for these 16 slices, only the subset of the slices beginning at and after LV peak enhancement registered well. In some series, before LV peak enhancement the cropped images did not contain sufficient gradient information to obtain a reliable rigid registration of the prealigned subset. In these cases it could happen that the gradient-based rigid registration actually aligned a part of the outer wall of the myocardium in one image to a part of

the inner wall in the other image, which in turn resulted in misaligned synthetic references for the final registration step. This effect does not occur with the non-rigid registration of the full images, because here the smoothness constraint combined with the non-moving parts surrounding the heart limit the freedom of the transformation.

In summary, employing rigid registration generally performed less optimally than using nonrigid registration. This is even the case when the instabilities described above are avoided by running the analysis only for images after LV peak enhancement and by restricting the transformation space Ψ to translations only (Table IV).

Table IV
RESULTS COMPARING RIGID AND NONRIGID REGISTRATION

	Measure	Mean	SD	Median	Min	Max
rigid	R^2	0.88	0.22	0.97	-0.43	1.00
	Q	0.85	0.43	0.79	0.13	2.88
	σ_s	0.59	0.30	0.59	0.20	2.27
nonrigid	R^2	0.96	0.07	0.98	0.53	1.00
	Q	0.63	0.26	0.64	0.11	1.37
	σ_s	0.56	0.25	0.50	0.15	1.35

Results obtained using rigid registration allowing translations only instead of nonrigid registration in our scheme compared with results obtained by the nonrigid registration scheme. Compared with Table II the analysis was restricted to the time after LV peak enhancement. All three measures indicate that nonrigid registration performs better.

IV. DISCUSSION AND FUTURE WORK

Allowing a patient to breathe freely during myocardial perfusion image acquisition has some advantages over asking the patient to hold his or her breath. On one hand, the procedure is easier for the patient. On the other hand, normal free breathing typically results in smaller, more periodic, and hence, predictable movements, which are easier to correct for respiratory motion than breath holding in situations when the patient can no longer hold the breath and takes a large gasp.

Following this protocol, the obtained image series pose two problems: the strong local difference in image intensities and the breathing movement. In this case, registering all images to one reference directly is usually unreliable, and serial registration can fail because of error propagation.

Based on these assumptions, we propose a fully automated approach based on nonrigid image registration to correct for the breathing movement by taking into account the advantages of free-breathing acquisition protocols as follows. First, by aligning a subset of the image series acquired at a similar breathing phase, we can establish a baseline of registered images that first ensures that no error propagation occurs. Second, because these images are already well aligned, the challenge in their registration is reduced to dealing with the strong local intensity change. Third, by using intensity-interpolated images as references for the registration of the remaining images, which are very close in the temporal series, the problem of intensity change is reduced significantly in the last phase of the algorithm.

In the first phase of the registration, the image pairs will naturally exhibit strong differences of the intensity-material mapping. Therefore, we rely on the normalized gradient-based

similarity measure F_{NGF} as the registration criterion. The normalized gradient-based similarity measure F_{NGF} has some advantages over well-known measures based on information theory, e.g., MI. The evaluation of F_{NGF} can be implemented easily, and it has a low computational footprint. In addition, it is a local measure in the sense that contributions to the overall cost function are based only on the very close vicinity of each pixel. Therefore, it is well suited to accommodate the local intensity change induced by the contrast agent passing through the heart ventricles and the myocardium. However, because of this local nature of the NGF-based measure, it is not advisable to use it as the sole criterion when large differences between images must be accommodated.

Therefore, in the second phase, learning from previous experience [12], we performed the analysis using a weighted combination of F_{NGF} and F_{SSD} as the registration criterion. The validation showed that one can actually omit F_{NGF} from this measure because the synthetically generated references generally model the intensity distribution of the corresponding unregistered images well enough. Only at the beginning of the series, when the contrast agent passes through the right heart ventricle, the linear interpolation sometimes did not produce good reference images resulting in worse motion compensation in this phase of the sequence. This is the case for high contrast data (Patient C). In addition, when applying the method to shallow breathing data (Patient D), the selection of the prealigned subset in the first step of the algorithm may result in a poor distribution of references. While it is possible to compensate for these problems by applying a higher weight to the regularization term and a larger knot spacing (as we did in the case of patient C), this usually results in more residual movement after registration and, hence, in less over-all alignment. Therefore, to make the algorithm more robust, other interpolation strategies could be used to model the intensity change better when creating the synthetic references. By comparing the intensity–time curves obtained from our registered image sets with manually acquired images, we were able to prove that our proposed method yields good results for motion correction. Compared with the nonregistered data, the corrected intensity curves align well with the manually acquired sets (Figure 9), improving both the visual and quantitative analysis significantly. With the validation method in place, we were also able to optimize the parameters to obtain the best registration results.

Finally, a comparison of our validation results with those reported by Milles et al. [4], who relied on rigid registration only, suggests that the nonrigid registration described above produces greater registration accuracy. We ran their ICA-based registration scheme for comparison and, although the method worked well for breath-holding data, thereby confirming their results and validating our implementation, its application to free-breathing data was not robust. Replacing the nonrigid registration method by rigid registration in our method did not provide a stable motion-compensation scheme. In the cases where it worked reasonably well, the rigid registration method performed worse than the proposed nonrigid registration scheme.

Future work will include comparing this method with other

approaches. We also plan to exploit the quasiperiodicity of free breathing even more by targeting an optimization process that would register all images in one step to describe the quasiperiodic transformation using a spatiotemporal transformation model [29].

ACKNOWLEDGMENT

This study was supported partially by research projects TIN2007-68048-C02-01, CDTI-CDTEAM, and SINBAD (PS-010000-2008-1) from Spain’s Ministry of Science and Innovation. Experimental data were provided with the support of the Intramural Research Program of the NIH, National Heart, Lung and Blood Institute.

REFERENCES

- [1] T. Makela, P. Clarysse, O. Sipila, N. Pauna, Q. Pham, T. Katila, and I. Magnin, “A review of cardiac image registration methods,” *IEEE Trans. Med. Imag.*, vol. 21, no. 9, pp. 1011–1021, 2002.
- [2] J. Milles, R. van der Geest, M. Jerosch-Herold, J. Reiber, and B. Lelieveldt, “Fully automated motion correction in first-pass myocardial perfusion mr image sequences,” *IEEE Trans. Med. Imag.*, vol. 27, no. 11, pp. 1611–1621, 2008.
- [3] M. Breeuwer, L. Spreeuwers, and M. Quist, “Automatic quantitative analysis of cardiac MR perfusion images,” in *Proceedings of the SPIE - The International Society for Optical Engineering*, vol. 4322(1-3), 2001, pp. 733–742.
- [4] J. Milles, R. J. van der Geest, M. Jerosch-Herold, J. H. Reiber, and B. P. Lelieveldt, “Fully automated registration of first-pass myocardial perfusion MRI using independent component analysis,” *Inf Process Med Imaging*, vol. 20, pp. 544–55, 2007.
- [5] S. N. Gupta, M. Solaiyappan, G. M. Beache, A. E. Arai, and T. K. Foo, “Fast method for correcting image misregistration due to organ motion in time-series mri data,” *Magnetic Resonance in Medicine*, vol. 49, pp. 506–514, 2003.
- [6] K. K. Wong, E. S. Yang, E. X. Wu, H.-F. Tse, and S. T. Wong, “First-pass myocardial perfusion image registration by maximization of normalized mutual information,” *J. Magn. Reson. Imaging*, vol. 27, pp. 529–537, 2008.
- [7] T. Delzescaux, F. Frouin, A. D. Cesare, S. Philipp-Foliguet, A. Todd-Pokropek, A. Herment, and M. Janier, “Using an adaptive semiautomated self-evaluated registration technique to analyze mri data for myocardial perfusion assessment,” *Journal of Magn. Reson. Imag.*, vol. 18, pp. 681–690, 2003.
- [8] C. Dornier, M. K. Ivancevic, P. Thevenaz, and J.-P. Vallee, “Improvement in the quantification of myocardial perfusion using an automatic spline-based registration algorithm,” *Journal of Magn. Reson. Imag.*, vol. 18, no. 2, pp. 160–168, 2003.
- [9] H. Ólafsdóttir, “Nonrigid registration of myocardial perfusion MRI,” in *Proc. Svenska Symposium i Bildanalys, SSBA 2005, Malmö, Sweden*. SSBA, 2005. [Online]. Available: <http://www2.imm.dtu.dk/pubdb/p.php?3599>
- [10] B. Likar and F. Pernus, “A hierarchical approach to elastic registration based on mutual information,” *Image Vision Comput.*, vol. 19, no. 1-2, pp. 33–44, 2001.
- [11] C. Studholme, C. Drapaca, B. Iordanova, and V. Cardenas, “Deformation-based mapping of volume change from serial brain MRI in the presence of local tissue contrast change,” *Medical Imaging, IEEE Transactions on*, vol. 25, no. 5, pp. 626–639, May 2006.
- [12] G. Wollny, M. J. Ledesma-Carbayo, P. Kellman, and A. Santos, “A New Similarity Measure for Non-Rigid Breathing Motion Compensation of Myocardial Perfusion MRI,” in *Proc. of the 30th Int. Conf. of the IEEE-EMBS*, Vancouver, BC, Canada, 2008, pp. 3389–3392.
- [13] E. Haber and J. Modersitzki, “Beyond mutual information: A simple and robust alternative,” in *Bildverarbeitung für die Medizin 2005*, ser. Informatik Aktuell, A. H. Hans-Peter Meinzer, Heinz Handels and T. Tolxdorff, Eds. Springer Berlin Heidelberg, 2005, pp. 350–354.
- [14] H. Xue, S. Zuehlsdorff, P. Kellman, A. Arai, S. Nielles-Vallespin, C. Chedfdhotel, C. H. Lorenz, and J. Guehring, “Unsupervised Inline Analysis of Cardiac Perfusion MRI,” in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2009*, ser. LNCS, vol. 5762/2009. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 741–749.

- [15] L. M. Bidaut and J. P. Vallée, "Automated Registration of Dynamic MR Images for the Quantification of Myocardial Perfusion," *Journal of Magn. Reson. Imag.*, vol. 13, no. 4, pp. 648–655, 2001.
- [16] G. Wollny, M. J. Ledesma-Carbayo, P. Kellman, and A. Santos, "Non-Rigid Motion Compensation in Free-Breathing Myocardial Perfusion Magnetic Resonance Imaging," in *Proc. of the 34th Int. Conf. of Computers in Cardiology*, vol. 35, Bologna, Italy, 2008, pp. 465–468.
- [17] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Image registration by maximization of combined mutual information and gradient information," in *International Conference on Medical Image Computing and Computer-Assisted Intervention MICCAI'2000*, ser. Lecture Notes in Computer Science. Cambridge, England: Springer Verlag, 2000, pp. 452–461.
- [18] J. Kybic and M. Unser, "Fast parametric elastic image registration," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1427–1442, 2003.
- [19] T. Rohlfing, C. R. M. Jr., D. A. Bluemke, and M. A. Jacobs, "Volume-preserving nonrigid registration of mr breast images using free-form deformation with an incompressibility constraint," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 730–741, 2003.
- [20] D. Marquardt, "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *SIAM J. Appl. Math.*, vol. 11, pp. 431–441, 1963.
- [21] C. Sánchez Sorzano, P. Thévenaz, and M. Unser, "Elastic registration of biological images using vector-spline regularization," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 4, pp. 652–663, 2005.
- [22] M. Unser, A. Aldroubi, and M. Eden, "The L_2 -polynomial spline pyramid," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 4, pp. 364–379, 1993.
- [23] M. Ledesma-Carbayo, P. Kellman, A. Arai, and E. McVeigh, "Motion corrected free-breathing delayed enhancement imaging of myocardial infarction using non-rigid registration," *Journal of Magn. Reson. Imag.*, vol. 26, pp. 184–190, 2007.
- [24] H. Xue, J. Guehring, L. Srinivasan, S. Zuehlsdorff, K. Saddi, C. Chedhotel, J. V. Hajnal, and D. Rueckert, "Evaluation of rigid and non-rigid motion compensation of cardiac perfusion mri," in *MICCAI '08: Proceedings of the 11th International Conference on Medical Image Computing and Computer-Assisted Intervention, Part II*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 35–43.
- [25] G. Wollny, M. J. Ledesma-Carbayo, and A. Santos, "An ITK implementation of the Normalized Gradient Field Image to Image Metric," *The Insight Journal*, 2010, <http://hdl.handle.net/10380/3150>.
- [26] T. Yoo, M. J. Ackerman, W. E. Lorensen, W. Schroeder, V. Chalana, S. Aylward, D. Metaxes, and R. Whitaker, "Engineering and Algorithm Design for an Image Processing API: A Technical Report on ITK - The Insight Toolkit," in *Medicine Meets Virtual Reality*, J. Westwood, Ed. IOS Press Amsterdam, 2002, pp. 586–592.
- [27] P. Kellman, F. H. Epstein, and E. R. McVeigh, "Adaptive sensitivity encoding incorporating temporal filtering (TSENSE)," *Magn Reson Med.*, vol. 45, no. 5, pp. 846–852, 2001.
- [28] G. Wollny, M. J. Ledesma-Carbayo, P. Kellman, and A. Santos, "Exploiting Quasi-Periodicity in Motion Correction of Free-Breathing Myocardial Perfusion MRI *supplemental*," <http://www.die.upm.es/im/videos/TMI2009/>, 2010.
- [29] M. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Suhling, P. Hunziker, and M. Unser, "Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation," *IEEE Trans. Med Imag.*, vol. 24, no. 9, pp. 1113–1126, 2005.