Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3D echocardiography

Mathieu de Craene, Gemma Piella, Oscar Camara, Nicolas Duchateau, Etelvino Silva, Adelina Doltra, Jan D'hooge, Josep Brugada, Marta Sitges, Alejandro Frangi

Temporal Diffeomorphic Free-Form Deformation: Application to Motion and Strain Estimation from 3D Echocardiography

Mathieu De Craene[*,a,b], Gemma Piella[a,b], Oscar Camara[a,b], Nicolas Duchateau[a,b], Etelvino Silva[e],
Adelina Doltra[e], Jan D'hooge[d], Josep Brugada[e], Marta Sitges[e], Alejandro F. Frangi[a,b,c]

[a]Center for Computational Imaging & Simulation Technologies in Biomedicine; Information & Communication Technologies
Department, Universitat Pompeu Fabra, Barcelona, Spain.
[b]Centro de Investigación Biomédica en Red en Bioingeniería, Biomateriales y Nanomedicina (CIBER-BBN), Spain.
[c]Department of Mechanical Engineering, The University of Sheffield, Sheffield, UK.
[d]Department of Cardiovascular Diseases, Cardiovascular Imaging and Dynamics, Katholieke Universiteit Leuven, Belgium.
[e]Hospital Clínic i Provincial de Barcelona; IDIBAPS; Universitat de Barcelona, Spain.

Abstract

This paper presents a new registration algorithm, called Temporal Diffeomorphic Free Form Deformation (TDFFD), and its application to motion and strain quantification from a sequence of 3D ultrasound (US) images. The originality of our approach resides in enforcing time consistency by representing the 4D velocity field as the sum of continuous spatiotemporal B-Spline kernels. The spatiotemporal displacement field is then recovered through forward Eulerian integration of the non-stationary velocity field. The strain tensor is computed locally using the spatial derivatives of the reconstructed displacement field. The energy functional considered in this paper weighs two terms: the image similarity and a regularization term. The image similarity metric is the sum of squared differences between the intensities of each frame and a reference one. Any frame in the sequence can be chosen as reference. The regularization term is based on the incompressibility of myocardial tissue. TDFFD was compared to pairwise 3D FFD and 3D+t FFD, both on displacement and velocity fields, on a set of synthetic 3D US images with different noise levels. TDFFD showed increased robustness to noise compared to these two state-of-the-art algorithms. TDFFD also proved to be more resistant to a reduced temporal resolution when decimating this synthetic sequence. Finally, this synthetic dataset was used to determine optimal settings of the TDFFD algorithm. Subsequently, TDFFD was applied to a database of cardiac 3D US images of the left ventricle acquired from 9 healthy volunteers and 13 patients treated by Cardiac Resynchronization Therapy (CRT). On healthy cases, uniform strain patterns were observed over all myocardial segments, as physiologically expected. On all CRT patients, the improvement in synchrony of regional longitudinal strain correlated with CRT clinical outcome as quantified by the reduction of end-systolic left ventricular volume at follow-up (6 and 12 months), showing the potential of the proposed algorithm for the assessment of CRT.

Key words: Spatiotemporal registration, diffeomorphism, FFD, strain, 3D ultrasound.

1. Introduction

1.1. Motion and strain quantification from 3D ultrasound

The quantification of cardiac motion and strain provides insight into cardiac function through the assessment of how a given pathology affects global and local deformation of the myocardium. In clinical routine, ultrasonography (US) is the standard imaging modality for the quantification and the analysis of regional

motion and strain. 3D US and Magnetic Resonance Imaging (MRI) offer in clinical practice similar volume rates (about 20-40 frames per heart cycle). Standard 3D US acquisition protocols divide the field of view in sectors, each sector being acquired separately. Within a sector, 3D US provide real-time images, while MRI requires gating. Besides, MRI is most often acquired per-slice and not as a complete volume. Even though MRI has better signal-to-noise ratio (SNR), US remains the most commonly used imaging modality. Indeed, US allows image analysis of follow-ups with implanted devices, has lower cost and higher availability than MRI at clinical sites. With respect to 2D US, 3D US enables motion quantification in the entire myocardium and does not suffer from out-of-plane motion artifacts. Nonetheless, several factors can induce artifacts in US images, such as signal attenuations, shadows, inhomogeneous resolution etc. Thus, 3D US images have relatively low SNR, making their processing more challenging.

Several approaches (Duan et al. (2007); Elen et al. (2008); Kawagishi (2008); Abe et al. (2008); Wang et al. (2010); Andrade et al. (2011)) have been proposed to extend 2D speckle tracking techniques to 3D with the objective of recovering myocardial motion from a 3D US image sequence. Duan et al. (2007) used an optical flow technique to quantify myocardial motion but limited their validation to simulated US data. Elen et al. (2008) applied B-Spline registration technique guided by mutual information to estimate 3D myocardial strain from 3D US sequences. They obtained absolute errors up to 21% for the radial strain and up to 4% for the longitudinal and circumferential strains on simulated US images, and recovered strain curves on *in vivo* data with physiologically plausible magnitude and dynamics. Kawagishi (2008) described the methods for applying speckle tracking on CRT candidates using 3D block matching and further computed regional strain curves from the resulting sparse displacement field. Radial strain curves obtained from a 3D speckle tracking technique were plotted for a single patient, but the authors did not look at the improvement of the strain pattern over therapy. No validation nor comparison to other techniques was made in their paper. Earlier work by the same authors (Abe et al. (2008)) on 3D speckle tracking only reported strain quantification results on a synthetic left ventricle (LV) model. Andrade et al. (2011) used 3D speckle tracking to quantify LV twist in 50 patients and highlighted statistical differences between 2D and 3D rotation measurements without looking at strain. Finally, Wang et al. (2010) recently proposed an unified Bayesian framework fusing multiple features (speckle patterns, border detection, intensity gradient and motion prediction) to recover 3D motion and strain. As this framework uses a training set to encode motion priors, an important issue is to know whether a single training set can handle all possible pathological cases. Such an approach could also introduce a bias towards deformation patterns that are more predominant among the training population. This approach is still pairwise since it estimates motion at every frame in a sequential manner without optimizing a joint objective function on the entire sequence.

1.2. Temporal consistency in image registration

One of the main drawbacks of speckle-tracking techniques based on *pairwise* registrations is that they do not make use of the temporal information embedded in the 3D US sequences.

In the literature, several image registration approaches, were proposed to exploit the temporal coherence of the input image sequence and they are summarized in Table 1. Such approaches can be organized into two main application fields: inter-subject sequence matching and deformation/motion quantification. Table 1 gives for each approach the application domain, the image modality, and its most distinctive algorithmic feature. It also specifies the input and output dimensionality of the transformation model ($d_{in} \rightarrow d_{out}$). The two last columns of Table 1 specify if the spatiotemporal transformation is represented as a displacement or a velocity field, and if it ensures temporal continuity.

In the context of inter-subject sequence matching, spatiotemporal matching attempts to map the cardiac geometry and dynamics of two patients for statistical atlas construction (Perperidis et al. (2005)) or to compare cardiac function of a patient before and after treatment (Peyrat et al. (2010)). Perperidis et al. (2005) proposed a Free Form Deformation (FFD) (Rueckert et al. (1999)) separable parameterization of the spatiotemporal transformation mapping two subjects. Peyrat et al. (2010) proposed to enforce the spatiotemporal mapping to match a given material point at different time points using trajectory constraints. Sundar et al. (2009) applied an inter-sequence alignment strategy for motion quantification. They created a static 3D+t sequence by replicating the first image of the sequence and then aligned the original one with the static 3D+t sequence.

2

In the context of motion quantification, spatiotemporal registration algorithms are based on the simultaneous alignment of all input images in the sequence. Ledesma-Carbayo et al. (2005) were pioneers in introducing a parametric spatiotemporal model, smooth in space and time, to take advantage of the temporal information contained in 2D US image sequences. Their algorithm was recently extended in two different directions. First, Metz et al. (2011) extended the output dimension of the transformation model to 3D+t. They used the sum of intensity variances over time as a groupwise similarity metric. Since uniqueness of the solution was not guaranteed, the authors constrained the average transformation in time to be the identity transformation. Second, Yigitsoy et al. (2011) kept the transformation model unchanged but took the mean squared intensity differences between all possible combinations of warped images as similarity metric. They referred to this metric as accumulated pairwise estimates. These three algorithms (Ledesma-Carbayo et al. (2005); Metz et al. (2011); Yigitsoy et al. (2011)) guarantee to recover smooth transformations in time by representing the displacement as a sum of smooth and continuous kernels. This displacement is expressed in the space of coordinates of a fixed reference frame. Such a representation is Lagrangian if the reference space is interpreted as the space of coordinates of material points. The advantage of this representation is that strain curves can be computed in a Lagrangian way, giving the deformation of one material point over time. The disadvantage, however, is that the motion field is not constructed as the composition of small incremental transformations. The displacement is smooth in time but does not functionally depend on the motion at earlier time instants. As it is expected that the transformation is not only smooth but also correlated between consecutive time points, this means that plausible transformations are restricted to a subset of the optimization space. The convergence of the optimization process can therefore be slowed down because of a suboptimal representation of the motion field.

One way of enforcing the motion to depend on earlier time points is to switch from a representation based on displacement to a representation based on velocity. Although the transformation model does not encode directly the trajectories, the Lagrangian displacement can still be reconstructed from the integration in time of the velocity field. Lagrangian strain or the image-based similarity to any reference frame in the sequence can then be computed from the Lagrangian displacement field. The benefit of basing the transformation representation on velocity rather than displacement is that by construction, motion at every time is obtained as the composition of velocity at all previous times. We refer to this functional dependence as *temporal consistency* and give experimental evidence in this paper of the resulting gain in accuracy and robustness against image noise.

1.3. Temporal diffeomorphic registration

Diffeomorphic registration algorithms (e.g. Beg et al. (2005); Vercauteren et al. (2007); Rueckert et al. (2006); Hernandez et al. (2009); Ashburner (2007); Avants et al. (2008)) ensure a continuous and differentiable correspondence with continuous inverse between the features to register. They are therefore particularly well suited to handle medical image sequences as they preserve topology and orientation of the observed anatomical structures over time. By integrating a velocity field over time, they provide an elegant way of enforcing temporal consistency.

This concept was applied by Khan and Beg (2008) to monitor growth processes by extending the Large Deformation Diffeomorphic Metric Mapping (LDDMM) image registration algorithm (Beg et al. (2005)). Velocities were computed using a dense grid, which did not guarantee their spatiotemporal continuity, unlike registration methods that rely on parametric transformation models. A regularization term was added to the image similarity metric and involved a smoothing kernel thus enforcing the spatial continuity of the computed velocities. The *temporal* continuity of the velocities was then fully conditioned by the implicit assumption that the observed image features preserved their topology along the image sequence. Such an assumption does not hold for noisy image sequences, such as those arising from cardiac echocardiography, as speckle and image artifacts can generate discontinuities from one frame to the next.

The concept of exploiting the diffeomorphic registration framework to enforce temporal consistency was applied to 2D contours and 3D shapes by Durrleman et al. (2009) and Qiu et al. (2009). However, while the computational cost of reconstructing dense velocity fields is acceptable for contours, paths or surfaces, its extension to dense volumetric spatiotemporal data remains critical. Moreover, all previous approaches

3

extending diffeomorphic image registration to temporal sequences only output a piecewise geodesic mapping (Durrleman et al. (2009)), for which the continuity of velocities at each data time point is not guaranteed. Trouvé and Vialard (2010) recently presented a framework for smoothly interpolating the trajectories of a set of landmarks over time by minimizing the norm of a forcing term in the evolution of the momentum equation. While this approach is suitable for handling landmark data, its computational cost would become prohibitive when applied to 3D images.

1.4. Incompressibility of myocardial tissue and application to cardiac motion quantification

Non-rigid image registration algorithms often require the inclusion of some prior information to compensate for the noise in the input data. In the case of cardiac motion quantification, several authors proposed the use of the incompressibility constraint to regularize the registration result.

The rationale behind incompressibility is that the various constituents of myocardial tissue as well as the blood perfusing it are incompressible. Experimental evidence of incompressibility was provided by Tsuiki and Ritman (1980) on a canine LV metabolically supported and monitored using angiography images. Although the volume quantified over the cardiac cycle was found to be globally constant, they reported variations about $\pm 5\%$. One explanation provided by the authors is the possible change in coronary blood volume between each phase of the cardiac cycle. Since the myocardium is perfused with distinsible vessels, and because of the direct contact between muscles and vessels, one can expect that local changes in stiffness due to tissue contraction would modify myocardial volume (Yin et al. (1996)). Westerhof et al. (2006) reported about mechanical interactions between cardiac muscle and coronary vasculature. In particular, during systole, thickening of the cardiac muscle can take place at the expense of the vascular volume. Muscle contraction is augmented when vascular emptying is facilitated. This corroborates with data showing that vascular volumes decreases with contraction (Spaan et al. (1981)). Yin et al. (1996) experimentally showed the change of vascular volume to vary from 2 to 4 ml/100 g tissue when modifying the perfusion pressure from 0 to 120 mmHg.

In regard to these experimental results, myocardial tissue can be considered as *quasi-incompressible*, with a slight deviation from the fully incompressible case varying over the cardiac cycle.

In the context of motion/deformation quantification algorithms, incompressibility was first introduced as a constraint combined with the constraint of mass conservation in an optical flow framework (Song and Leahy (1991)). It led to the joint minimization of optical-flow, divergence-free and smoothness constraints. This approach was applied by Song and Leahy (1991) and Gorce et al. (1997) to the estimation of myocardial velocities from Computed Tomography (CT) images. Bistoquet et al. (2007) proposed the use of divergence-free radial basis kernels to represent myocardial displacements, implementing a first-order approximation of incompressibility. The model was validated by comparing the recovered motion obtained from cine MRI to motion information provided by tagged MRI images for healthy subjects and ventricular dyssynchrony patients. Haber and Modersitzki (2004) introduced a volume-preserving non-rigid registration algorithm by constraining locally the determinant of the transformation to equal unity. A specific discretization strategy was combined with a variant of the Sequential Quadratic Programming algorithm for solving the optimization problem. The method was applied to regularize motion quantification from contrast-enhanced data, being robust to contrast uptake registration artifacts. Saddi et al. (2007) also focused on motion quantification from perfusion sequences, but from CT images of the liver. In their approach, the non-rigid displacement field was projected onto the space of divergence-free vector fields using a multigrid solver.

Mansi et al. (2010, 2011) extended the diffeomorphic Demons registration method introduced by Vercauteren et al. (2007) by integrating elasticity and incompressibility for soft-tissue tracking. The basic idea in this method, as well as in Hinkle et al. (2009), is that diffeomorphisms parametrized by divergence-free velocity fields are incompressible. The proposed algorithm was applied to cine MRI for myocardial strain quantification. The introduction of the incompressibility constraint improved the strain estimation in cine MRI, taking the strain values computed from tagged MRI as ground truth.

1.5. Proposed approach

In this paper, we present a novel temporal diffeomorphic registration algorithm that models the velocities continuously in time and space and can be applied to 3D sequences with a reasonable computational cost.

Reference	Application domain	Modality	Principle	Transformation representation		Temporal continuity
				Vel. / Disp.	$d_{in} \rightarrow d_{out}$	
Inter-subject sequence matching						
Peyrat et al. (2010)	Cardiac inter-sequence registration	CT	Trajectory constraints, multi-channel diffeomorphic Demons	Disp.	$4 \rightarrow 4$	Yes
Perperidis et al. (2005)	Cardiac inter-sequence registration	MRI	Separable smooth spatiotemporal matching using FFD	Disp.	$4 \rightarrow 4$	Yes
Sundar et al. (2009)	Cardiac inter-sequence registration	tagged MRI	Attribute vector matching	Disp.	$4 \rightarrow 4$	Yes
Quantification of motion or longitudinal anatomical changes						
Khan and Beg (2008)	Neurological follow-up	synth. images & MRI	Sequence of dense non-stationary velocity fields	Vel.	$3 \rightarrow 3$	No
Durrleman et al. (2009)	Skull & Amygdala evolution	CT & MR	Sequence of non-stationary velocity fields represented by moments applied to a current-based shape representation	Vel.	$3 \rightarrow 3$	No
Trouvé and Vialard (2010)	Shape interpolation	synth. images & CT	Sequence of non-stationary coupled velocity and momentum applied to a current-based shape representation	Vel.	$3 \rightarrow 3$	Yes
Ledesma-Carbayo et al. (2005)	Cardiac motion and deformation	2D US	Non-stationary disp. field represented by a FFD model	Disp.	$4 \rightarrow 3$	Yes
Metz et al. (2011)	Lung, coronary and cardiac motion	CT, MRI, and US	Non-stationary disp. field represented by a FFD model with zero average displacement constraint	Disp.	$4 \rightarrow 4$	Yes
Yigitsoy et al. (2011)	Respiratory motion compensation	MRI	Non-stationary disp. field represented by a FFD model with accumulated pairwise registration	Disp.	$4 \rightarrow 4$	Yes
De Craene et al. (2010)	Cardiac motion and deformation	3D US	Non-stationary velocity field represented by a FFD model	Vel.	$4 \rightarrow 3$	Yes

Table 1: Proposed approaches introducing temporal consistency. Algorithms are classified according to application domain, imaging modality and their main algorithmic features. The last three columns indicate if the representation of the transformation uses a velocity (Vel.) or displacement (Disp.) field; the input, d_{in}, and output, d_{out} dimensions of the transformation model; and if temporal continuity is ensured.

Preliminary results on this methodology were presented in De Craene et al. (2010). In this work, we additionally introduce the possibility of using any time point as a reference, as well as the inclusion of quasi-incompressibility as a soft constraint in the optimization problem, and extend the clinical application of our algorithm regarding the analysis and the number of patients. We refer to our approach as Temporal Diffeomorphic Free Form Deformation (TDFFD) algorithm, since we extend the popular parametric FFD registration technique (Rueckert et al. (2006)) by summing spatiotemporal B-Spline kernels to model the 3D+t velocity field. An important contribution of this paper is the enforcement of the continuity of the velocity field by using a continuous parametric representation. As a result, the velocity and displacement fields can be computed at any time within the temporal interval captured by the image sequence. Another advantage of solving for velocity rather than for displacement is to endow the transformation model with temporal consistency. Indeed, if the velocity is changed locally at a given spatiotemporal location, the entire trajectory will be modified at subsequent times. This introduces a coupling between time steps that is expected to improve the robustness of the image registration algorithm against image noise.

Although the TDFFD algorithm can be applied to any imaging modality, its clinical applicability is demonstrated here on synthetic and real 3D US sequences. For assessing the robustness of the algorithm with respect to the presence of noise, we generated ground truth data with a known motion field at several Contrast to Noise Ratios (CNRs) levels. The underlying clinical objective is to accurately estimate the impact of Cardiac Resynchronization Therapy (CRT) on 3D myocardial strain. CRT aims to restore synchrony of motion and strain in the four chambers. When looking specifically at the LV, CRT responders are expected to show at follow-up a higher peak strain value at the end of systole (Delgado et al. (2009)) and an improved synchronization of the strain curves in the different LV segments (Delgado et al. (2008); Bertola et al. (2009)). Our algorithm was applied to a database of 13 CRT patients including both responders and non-responders before and after CRT, and at 6- and 12-months follow-ups, representing a total of 51 image sequences. For each of the 51 sequences, longitudinal strain curves were computed in different segments of the LV. The improvements in amplitude and synchrony were quantified for all patients and related to CRT outcome.

The remainder of this paper proceeds as follows. In Section 2, we describe the algorithmic framework implemented in this paper. We first detail the computation of a continuous spatiotemporal velocity field and how it generates trajectories in the 3D+t image domain. We then focus on the computation of the metric and its derivatives. Subsequently, we describe the computation of strain and briefly detail some implementation aspects. Section 3 is devoted to experiments on synthetic and clinical 3D US datasets, the latter including both volunteers and CRT patients. On synthetic images, we compare the TDFFD algorithm to the state-of-the-art pairwise FFD algorithm and to a 3D+t groupwise algorithm on the 3D+t displacement field, demonstrating an improvement in robustness in case of low CNRs. We then quantify the incidence of the main parameters of our algorithm on registration accuracy. On clinical datasets, we examine for healthy volunteers the uniformity of the strain pattern among different segments of the LV. On CRT patients, we quantify the correlation between strain curves and therapy outcome, as quantified by reverse remodeling. Finally Section 4 and 5 discuss the findings of this paper in regard to the literature and potential improvements for future work.

2. Methods

2.1. Registration framework and notations

In this paper, we denote a sequence of $N_1 + N_2 + 1$ images as $\{I_n\}$, with $n \in \mathcal{N} = \{n \in \mathbb{Z}, -N_1 \leq n \leq N_2\}$. The image I_0 designates the image taken as reference. All images before the reference image are indexed by $n = -N_1 \ldots -1$ while all images after the reference are indexed by $n = 1 \ldots N_2$. Each image, I_n, is defined on a spatial domain $\Omega \subset \mathbb{R}^d$ where d stands for the spatial dimension. Each image I_n is associated to a time stamp $t_n \in \mathcal{T} = [t_{-N1}, t_{N2}]$, with $t_{-N1}, t_{N2} \subset \mathbb{R}$.

A flowchart of the overall algorithmic registration framework is given in Fig. 1. The purpose of the registration algorithm is to solve for the diffeomorphic mapping $\boldsymbol{\varphi_0} : \Omega \times \mathcal{T} \to \mathbb{R}^d$ that transports a material point $\mathbf{x} \in \Omega$ from $t = 0$ to a target time $t \in \mathcal{T}$. In this paper, a continuous velocity field is parameterized using a regular 4D grid and smooth spatiotemporal B-Spline kernels, thus ensuring spatial and temporal continuity.
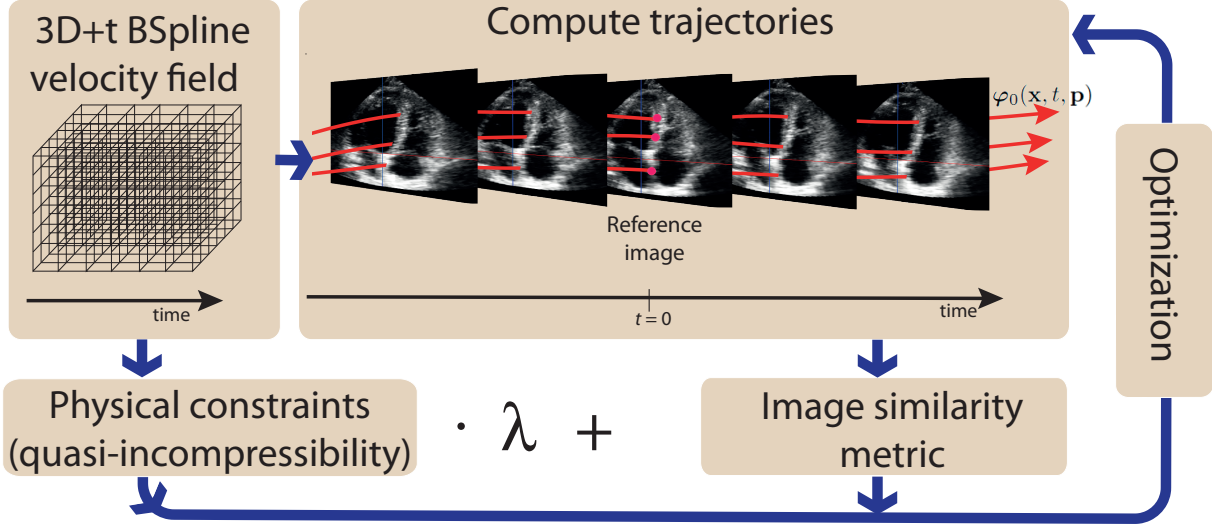
6

Figure 1: Overview of the algorithmic framework. Continuous spatiotemporal trajectories are computed from the 3D+t velocity field, parameterized by a 3D+t grid of control points with B-Spline kernels. Image similarity is computed from image intensities for the current set of trajectories. Quasi-incompressibility is added as a biophysical constraint for regularizing the velocity field.

Temporal continuity of the velocity field is not guaranteed when working with temporal diffeomorphic registration algorithms as described in Khan and Beg (2008) and in Durrleman et al. (2009). Recent extensions (Trouvé and Vialard (2010)) do ensure continuity when interpolating sparse trajectories defined on a set of landmarks by performing a variational perturbation on the temporal derivative of the momentum rather than on the velocity. This second-order approach supposes a significant increase in the computational cost. Since this paper addresses the problem of estimating dense trajectories from sequences of 3D images, the introduction of a continuous kernel in time was preferred to the introduction of higher orders.

The inclusion of the velocity in the registration framework adds one intermediate step in comparison with other spatiotemporal registration approaches that directly model the displacement field (Ledesma-Carbayo et al. (2005) and Chandrashekara et al. (2004)). The velocity field $\mathbf{v} : \Omega \times \mathcal{T} \rightarrow \mathbb{R}^d$ transports a given material point $\mathbf{x}$ from its original location at time $t = 0$ to its mapped location by a continuous trajectory, $\boldsymbol{\varphi}_0(\mathbf{x}, t)$, whose tangent vector at any time is given by $\mathbf{v}(\boldsymbol{\varphi}_0(\mathbf{x}, t), t)$. The numerical solution of this transport equation, described in Section 2.3, is obtained by the discretization of the temporal domain $\mathcal{T}$ using a non-uniform sampling $\{t_k\}$.

The velocity field, $\mathbf{v}$, is obtained through an optimization procedure guided by a cost function composed of an image-driven term depending on image intensities and a regularization term representing prior information about the motion field. For the data-driven term, we used the Mean Squared Error (MSE) between the intensities of the reference frame and all other frames. For the regularization term, a constraint imposing incompressibility of the myocardial tissue was chosen. Both of them are described in Section 2.4.

2.2. A continuous velocity field in the spatiotemporal domain.

The temporal dimension is introduced into the diffeomorphic registration problem by relating the mapping $\boldsymbol{\varphi}_0$ at any time t to a time-varying velocity field. In this paper, the velocity field is represented as a sum of spatiotemporal B-Spline kernels. The B-Spline velocity coefficients assigned to all the control points are concatenated in a vector of parameters $\mathbf{p}$, the velocity being then denoted as $\mathbf{v}(\mathbf{x}, t; \mathbf{p})$ and computed as

$$\mathbf{v}(\mathbf{x}, t; \mathbf{p}) = \sum_{i,j,k,l} \beta\left(\frac{x - q^i}{\Delta_i}\right)\beta\left(\frac{y - q^j}{\Delta_j}\right)\beta\left(\frac{z - q^k}{\Delta_k}\right)\beta\left(\frac{t - q^l}{\Delta_l}\right)\mathbf{p}^{i,j,k,l} \tag{1}$$

7

where $\mathbf{x} = (x, y, z)$, $\beta(\cdot)$ is a 1D cubic B-Spline kernel, $\mathbf{q}^{i,j,k,l} = (q^i, q^j, q^k, q^l)$ is a sparse grid of 3D+t control points, and $\boldsymbol{\Delta} = (\Delta_i, \Delta_j, \Delta_k, \Delta_l)$ is the spacing between control points in each dimension. Hence, the trajectory $\boldsymbol{\varphi}_0$ that maps a material point of coordinates $\mathbf{x}$ at time $t = 0$, being the time associated to the reference frame, to time t_n is obtained by adding to the initial position the integration of the velocity field from time $t = 0$ to time $t = t_n$, i.e.,

$$\boldsymbol{\varphi}_0(\mathbf{x}, t_n; \mathbf{p}) = \mathbf{x} + \int_0^{t_n} \mathbf{v}(\boldsymbol{\varphi}_0(\mathbf{x}, t, \mathbf{p}), t; \mathbf{p}) \mathrm{d}t \,. \tag{2}$$

Registration on 2D echocardiographic sequences may require the use of drift correction techniques due to out-of-plane motion. The clinical modality targeted in this work is 3D US, and taking into account that images are acquired under breath-hold in real time, no significant drift effect is expected. For this reason, drift correction was not included in the transformation model.

2.3. Trajectory computation.

The transport equation for computing $\boldsymbol{\varphi}_0$ in Equation (2) is solved numerically using a forward Euler integration scheme in which the continuous integral is replaced by a discrete summation. The continuous time interval is now split into a collection of $t_k \in \mathcal{T}$ values where the time increment between consecutive time-steps Δt_k is adapted to ensure invertibility as further described. By convention, we will use the subscript index k to refer to the discretization of the temporal dimension for solving the transport equation, whereas the subscript index n will refer to the temporal instants at which imaging data is available. The k-sample scheme is assumed to be denser than the temporal sampling of the image grid to ensure accurate estimation of the mapping $\boldsymbol{\varphi}_0$. For the simplicity of notation, we will assume that the n-sample scheme is a subset of the k-sample scheme.

Using this discretization, Equation (2) can be approximated by

$$\boldsymbol{\varphi}_0(\mathbf{x}, t_n; \mathbf{p}) = \mathbf{x} + \sum_{k=0}^{n-1} \mathbf{v}(\boldsymbol{\varphi}_0(\mathbf{x}, t_k; \mathbf{p}), t_k; \mathbf{p}) \Delta t_k \,, \tag{3}$$

when n takes a positive value. When n takes a negative value, the same expression can be used but the upper limit of the sum has to be changed to $n + 1$.

If we define $\mathbf{x}_0^k(\mathbf{p}) = \boldsymbol{\varphi}_0(\mathbf{x}, t_k; \mathbf{p})$ as the transport of the coordinate $\mathbf{x}$ from time $t = 0$ to t_k and $\mathbf{v}^k(\mathbf{p}) = \mathbf{v}(\mathbf{x}_0^k(\mathbf{p}), t_k; \mathbf{p})$, we can recursively write $\mathbf{x}_0^k(\mathbf{p})$ as follows:

$$\mathbf{x}_0^k(\mathbf{p}) = \mathbf{x}_0^{k-1}(\mathbf{p}) + \mathbf{v}^{k-1}(\mathbf{p}) \Delta t_{k-1} \,. \tag{4}$$

Equation (4) is still valid for a time $t_k < 0$ giving a negative sign to Δt_{k-1}. It can also be generalized to the transport from any frame in the sequence by taking an initial time different from 0. The integration of Equation (2) using the discrete approximation of Equation (4) requires to select a time-step sufficiently small for ensuring accurate computation and invertibility of the mapping $\boldsymbol{\varphi}_0$. In our method, we start off with a uniform sampling of the temporal domain $\mathcal{T}$, arbitrarily chosen as half of the temporal spacing of the image sequence. To ensure invertibility at any time t_n, one needs to consider the Jacobian of the mapping $\mathbf{x}_0^n$ with respect to the set of parameters $\mathbf{p}$, here denoted as $\mathbf{D}\mathbf{x}_0^n(\mathbf{p})$ and computed from Equation (4) using

$$\mathbf{D}\mathbf{x}_0^n(\mathbf{p}) = \Delta\boldsymbol{\varphi}_{n-1}^n \dots \Delta\boldsymbol{\varphi}_{k-1}^k \dots \Delta\boldsymbol{\varphi}_0^1 \text{ with } \Delta\boldsymbol{\varphi}_{k-1}^k = \mathbf{I} + \mathbf{D}\mathbf{v}^{k-1}(\mathbf{p}) \Delta t_{k-1} \,, \tag{5}$$

where $\mathbf{I}$ stands for the identity matrix. This Jacobian must be positive definite everywhere to ensure invertibility of the transformation. A necessary condition to accomplish this goal is to have $\det(\mathbf{D}\mathbf{x}_0^n(\mathbf{p})) > 0$ for all $\mathbf{x} \in \Omega$. Computing the product over k of all $\det(\Delta\boldsymbol{\varphi}_{k-1}^k)$ gives the determinant of the Jacobian matrix in Equation (5). When a negative value of $\det(\Delta\boldsymbol{\varphi}_{k-1}^k)$ is detected, the value of Δt_k is reduced by a factor of 2 until enforcing a positive determinant. Note that the zero threshold could easily be substituted by a small positive value for not hindering the stability of the trajectories computation.

8

2.4. Similarity metric and non-linear optimization.

2.4.1. Image similarity

Tissue and blood pool intensities are globally conserved over the cardiac cycle in 3D US images. Therefore, we use the Mean Squared Error (MSE) metric for capturing the optimal set of B-Spline velocity coefficients $\mathbf{p}$ of Equation (1). More complex similarity measures as described by Cohen and Dinstein (2002) have been proposed to take into account the statistical properties of the speckle noise. However, in the case of an independent additive gaussian noise, motion estimation based on a maximum likelihood approach leads to the minimization of the MSE metric.

The similarity metric is computed between an arbitrarily chosen reference frame I_0 in the sequence and all the consecutive time frames according to

$$M(\mathbf{p}) = \sum_{n=-N_1}^{N_2} \int_\Omega \left(I_n(\mathbf{x}_0^n(\mathbf{p})) - I_0(\mathbf{x}) \right)^2 \mathrm{d}\mathbf{x}. \tag{6}$$

The framework presented in this paper can be generalized to other similarity metrics such as information theoretic measures as long as they are continuous and differentiable. The derivative will be a function of $\mathrm{d}\mathbf{x}_0^k/\mathrm{d}\mathbf{p}$ whose computation is given in Section 2.4.3.

2.4.2. Quasi-incompressibility

Quasi-incompressibility is equivalent, in this framework, to imposing a zero divergence of the velocity field at every spatiotemporal location. Since the divergence is the trace of the Jacobian, the quasi-incompressibility constraint can be formulated as

$$C(\mathbf{p}) = \sum_{n=-N_1}^{N_2} C_n(\mathbf{p}), \tag{7}$$

where $C_n(\mathbf{p})$ is given by

$$C_n(\mathbf{p}) = \int_{\Omega_M} \left(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big) \right)^2 \mathrm{d}\mathbf{x}. \tag{8}$$

Here, Ω_M stands for the subset of the Ω domain that delineates the myocardial wall in the reference image I_0, *i.e.*, the subdomain defined by all the material points in the myocardium undergoing the deformations we are interested in analyzing. In this paper, Ω_M is obtained by segmenting the LV in the first frame of the sequence as described in Section 2.5. Since the quasi-incompressibility constraint is computed over a fixed domain in the reference frame, Ω_M can be considered as constant when computing derivatives. The domain is propagated using the current estimate of the transformation, thus avoiding the definition of a region of interest at every frame.

2.4.3. Complete metric and derivation

As there is some controversy in the literature (see Section 1.4) on the fully incompressible nature of the myocardial tissue, we opted in this paper for enforcing quasi-incompressibility as soft constraint. This differs from previous approaches (see Mansi et al. (2010, 2011) and Hinkle et al. (2009)), which project the velocity field at every iteration on a divergence-free velocity space.

In this paper, the incompressibility constraint of Equation (7) is added to the image similarity metric using a Lagrange multiplier $\lambda \in \mathbb{R}$. The complete cost function is therefore

$$F(\mathbf{p}) = M(\mathbf{p}) + \lambda \cdot C(\mathbf{p}). \tag{9}$$

The choice of an optimal λ is discussed in Section 3.1.4. Since the number of parameters characterizing the transformation is large, and the metric is explicitly differentiable, gradient-based optimization methods are well indicated for minimizing Equation (9). In this paper, the L-BFGS-B method of Byrd et al. (1995),

9

which searches the optimum based on the gradient and a low-rank approximation of the Hessian of the metric, was chosen.

The total derivative of the MSE metric in Equation (6) with respect to the parameters $\mathbf{p}$ of the velocity kernels is equal to

$$\frac{\mathrm{d}M(\mathbf{p})}{\mathrm{d}\mathbf{p}} = \sum_{n=-N_1}^{N_2} \int_{\Omega} 2 \left(I_n(\mathbf{x}_0^n(\mathbf{p})) - I_0(\mathbf{x}) \right) \mathbf{D}I_n(\mathbf{x}_0^n(\mathbf{p})) \frac{\mathrm{d}\mathbf{x}_0^n}{\mathrm{d}\mathbf{p}} \mathrm{d}\mathbf{x} \,, \tag{10}$$

where $\mathbf{D}I_n$ is the spatial gradient of the n^{th} image in the sequence and $\frac{\mathrm{d}\mathbf{x}_0^k}{\mathrm{d}\mathbf{p}}$ can be obtained by differentiation of Equation (4):

$$\frac{\mathrm{d}\mathbf{x}_0^k}{\mathrm{d}\mathbf{p}} = \frac{\mathrm{d}\mathbf{x}_0^{k-1}}{\mathrm{d}\mathbf{p}} + \frac{\mathrm{d}\mathbf{v}^{k-1}}{\mathrm{d}\mathbf{p}}\Delta t_{k-1} \quad \text{with} \quad \frac{\mathrm{d}\mathbf{v}^{k-1}}{\mathrm{d}\mathbf{p}} = \frac{\partial\mathbf{v}^{k-1}}{\partial\mathbf{x}_0^{k-1}}\frac{\mathrm{d}\mathbf{x}_0^{k-1}}{\mathrm{d}\mathbf{p}} + \frac{\partial\mathbf{v}^{k-1}}{\partial\mathbf{p}} \,. \tag{11}$$

Hence, $\frac{\mathrm{d}\mathbf{x}_0^k}{\mathrm{d}\mathbf{p}}$ can be obtained from the following recursive equation:

$$\frac{\mathrm{d}\mathbf{x}_0^k}{\mathrm{d}\mathbf{p}} = \left(\mathbf{I} + \mathbf{D}\mathbf{v}^{k-1}(\mathbf{p})\Delta t_{k-1} \right)\frac{\mathrm{d}\mathbf{x}_0^{k-1}}{\mathrm{d}\mathbf{p}} + \frac{\partial\mathbf{v}^{k-1}}{\partial\mathbf{p}}\Delta t_{k-1} \,, \tag{12}$$

where the first term propagates the derivative $\frac{\mathrm{d}\mathbf{x}_0^{k-1}}{\mathrm{d}\mathbf{p}}$ computed at the previous time step, weighted by the volume change the transformation introduces at the previous time step (factor between brackets).

The derivative of $C_n(\mathbf{p})$ in Equation (8) with respect to one element of $\mathbf{p}$ is given by

$$\frac{\mathrm{d}C_n(\mathbf{p})}{\mathrm{d}\mathbf{p}^{i,j,k,l}} = \int_{\Omega_M} 2\,\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\frac{\mathrm{d}}{\mathrm{d}\mathbf{p}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big)\mathrm{d}\mathbf{x} \,. \tag{13}$$

The detailed computation of $\frac{\mathrm{d}}{\mathrm{d}\mathbf{p}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big)$ is given in A.

2.5. Myocardial strain computation

The myocardial strain tensor is directly estimated using the spatial derivatives of the displacement field from Equation (5). In this paper, strain is computed in the space of coordinates of the first frame, which corresponds to the definition of Lagrangian strain. The strain tensor is then computed as

$$\boldsymbol{\epsilon}(\mathbf{x}, t_n) = \frac{1}{2}\Big((\mathbf{D}\mathbf{x}_0^n)^T\mathbf{D}\mathbf{x}_0^n - \mathbf{I}\Big) \,. \tag{14}$$

Note that in the cardiac literature (D'hooge et al. (2000)), it is common to define strain as a linear change in length, whereas Equation (14) is quadratic with respect to the spatial gradient of the displacement field. Both expressions are equivalent under the assumption of small displacements, i.e. that both the displacements and the displacements gradient are small compared to unity. In the case of Lagrangian strain, as strain is computed with respect to a fixed reference (end of diastole), the assumption of small displacement is invalid. This justifies why we opted for Equation (14) rather than the linear formulation. However, it is important to keep in mind that the coexistence of the two strain definitions might be responsible for some discrepancies in the strain values reported in the literature.

This tensor is projected along a specific direction $\mathbf{h}$ using $\epsilon_{\mathbf{h}}(\mathbf{x}, t_n) = \mathbf{h}^T \cdot \boldsymbol{\epsilon}(\mathbf{x}, t_n) \cdot \mathbf{h}$. The set of $\mathbf{h}$ directions considered here are the three vectors (radial, longitudinal, circumferential) of a local coordinate system related to the anatomy of the LV (D'hooge et al. (2000)). For computing this set of directions, the LV is segmented in the first frame of the sequence. According to Section 3.1.3, this corresponded to the optimal reference choice. This segmentation also defines Ω_M in Equation (8) and is obtained from an active shape and appearance model whose construction and application to image segmentation is described in Butakoff et al. (2007). This algorithm was initialized by defining the aortic valve, the mitral valve and the apex on

10

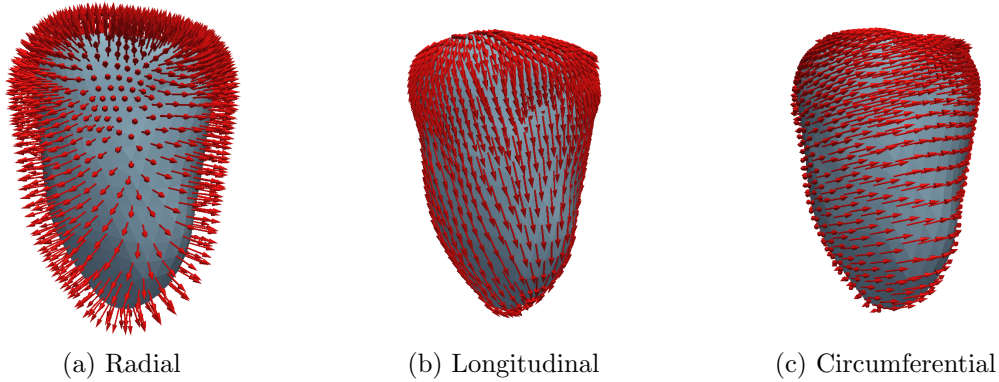(a) Radial   (b) Longitudinal   (c) Circumferential

Figure 2: The first frame in the image is segmented using a triangular surface mesh for defining AHA regions and local directions: (a) Radial vectors, (b) Longitudinal vectors, (c) Circumferential vectors.

the endocardial side. This provides semi-automated definition of the right ventricle insertion points. The longitudinal direction, $\mathbf{l}$, is defined uniformly by drawing a line from the apex to the mitral valve. These two landmarks are directly stored on the mesh, which makes the long axis computation automatic and robust to any image orientation. The radial direction is obtained from the normal $\mathbf{e}$ to the mesh at each node and the $\mathbf{l}$ vector using

$$\mathbf{r} = \mathbf{e} - (\mathbf{e} \cdot \mathbf{l})\,\mathbf{l}. \tag{15}$$

The circumferential direction is then obtained by the cross product of longitudinal and radial directions

$$\mathbf{c} = \mathbf{l} \times \mathbf{r}. \tag{16}$$

The LV was divided using the 17-segment model as proposed by the American Heart Association (AHA, Cerqueira et al. (2002)). These 17 regions are directly defined on the surface mesh for computing regional average strain curves. In this paper, strain is computed in a Lagrangian space of coordinates, hence both local coordinate system and AHA segments need to be defined on the first image only.

2.6. Implementation

The TDFFD registration algorithm was implemented in C++ within the ITK[1] toolkit. The B-Spline velocity field computation was inspired from the B-Spline deformable transformation object and mainly differ on allowing different dimensions for the input point and the output velocity. The computation of the matrix product and summation in Equation (12) was implemented through sparse matrices. All experiments were run on a 24-node double quad-core Intel Xeon (2.66 GHz CPU, 16 GB RAM) Linux cluster for distributing the load of processing different input image sequences. Regarding the computation times, the FFD algorithm required 20 minutes to process the whole sequence. The TDFFD algorithm, due to its higher complexity in the computation of metric derivatives, took 3h40 min. Interestingly, because of a very slow convergence; the Ledesma-Carbayo *et al.*'s algorithm took 8h40 min and 10 hours (for the two initialization strategies described in Section 3.1.2). The convergence of the TDFFD algorithm was regular and the objective function value uniformly decreased over the iterations, which explains a shorter execution time.

The code is currently not parallelized and uses a single CPU. Conceptually, every sample from the reference image can be processed independently for computing its contribution to the metric and its derivative. Dividing the set of image samples among different threads is therefore expected to significantly speed-up the computation time.

---

[1]http://www.itk.org

11

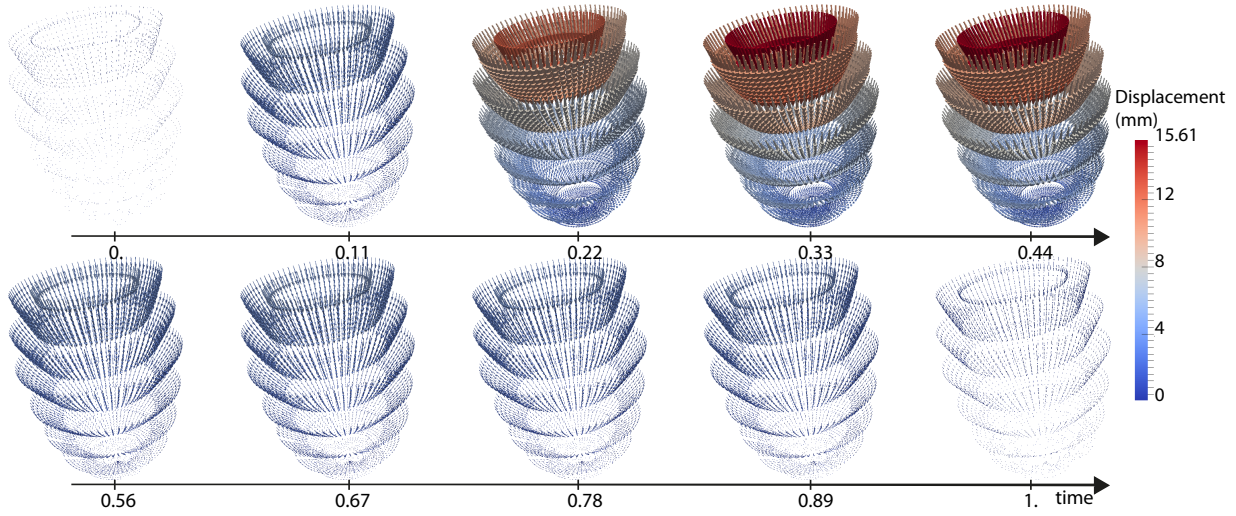Figure 3: 3D rendering of the ground truth displacement field used to generate the synthetic ultrasound images.

3. Experiments

Experiments were performed first on a set of synthetic 3D US sequences to evaluate the accuracy and the robustness of the TDFFD algorithm with respect to known ground truth displacement. Then, the clinical applicability of our method was tested for two populations on clinical 3D US sequences: a population of 9 healthy volunteers and a population of 13 CRT patients including responders and non-responders acquired before and after CRT, as well as 6 and 12 months follow-up.

3.1. Registration accuracy on simulated ultrasound data

The evaluation on synthetic US images consisted first in comparing the accuracy of the TDFFD algorithm to two state-of-the-art algorithms. The first one is the pairwise FFD (Rueckert et al. (1999)). The second one is the 3D+t algorithm proposed by Ledesma-Carbayo et al. (2005), guaranteeing a temporally smooth displacement field. In this figure, as in the sequel of this manuscript, the time was normalized by the duration of one cardiac period and is referred to as normalized cardiac time. This comparison was performed both in terms of displacement and strain, considering several imaging noise levels. Then, an optimal set of parameters for the B-Spline grid resolution and the weighting of the incompressibility term were determined by optimizing the error between true and recovered displacements.

3.1.1. Experimental setup

Elen et al. (2008) simulated 3D US images of the deforming LV for a normal contraction. This dataset had an isotropic voxel size of $0.36\,mm$. Concerning the temporal resolution, 20 frames per cycle were generated, which is slightly higher than the average number of frames per cycle for the set of healthy volunteers used in Section 3.2.1. In the anatomical model employed to generate the simulated 3D US images, the LV was represented as a thick-walled ellipsoid with end-diastolic dimensions in the physiological range.

Synthetic US data sets were generated from this mechanical model using a previously described methodology mimicking the whole image formation process (Gao et al. (2009)). In addition, a simplified kinematic model with an ejection fraction of 60 % over a cardiac cycle was developed to generate the ground truth displacement field. LV mechanics was based on a kinematic model as described in Arts et al. (1992). In this model, only transformations based on the changes of the cavity volume and the torsion were considered while the myocardium was assumed to be incompressible. The resulting synthetic displacement field is plotted in Fig. 3 and has an amplitude reaching $15.5\,mm$ in the base at the end of systole.

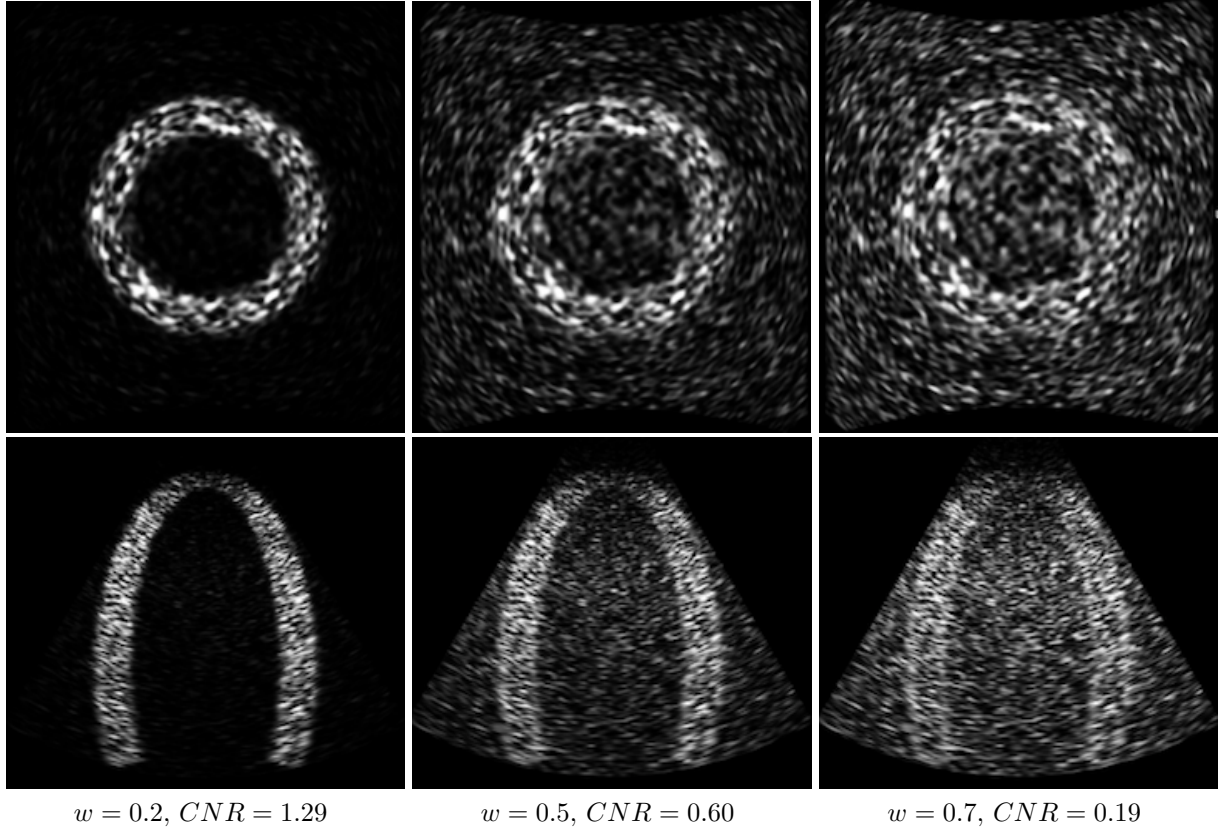$w = 0.2, CNR = 1.29$ $w = 0.5, CNR = 0.60$ $w = 0.7, CNR = 0.19$

Figure 4: Axial (top) and longitudinal (bottom) views of the synthetic datasets with the three different CNR generated for the evaluation of registration accuracy.

This ground truth data is then used in this paper to evaluate the accuracy of the proposed algorithm and compare it to a pairwise registration strategy at different noise levels. Various CNRs were generated by multiplying the blood pool speckle acoustic impedance (i.e. the amplitude of the intensities) by a factor w. In this paper, three values were considered for w: 0.2, 0.5 and 0.7. An increasing w means an increasing noise power, thus a reduced CNR. Snapshots of axial and longitudinal slices for the three levels of noise considered here are plotted in Fig. 4.

This set of values for w can be related to CNR values using the equation given in van Wijk and Thijssen (2002):

$$CNR = \frac{\mu_F - \mu_B}{\sqrt{(\sigma_B^2 + \sigma_F^2)/2}} \quad , \tag{17}$$

where μ_F and μ_B are the average intensity values of the foreground and background, σ_B^2 and σ_F^2 designating the foreground and background intensity variances. Computing μ_F, μ_B, σ_F and σ_B for the three values of w listed above returns CNR values of 1.29, 0.60, 0.19, respectively. These values roughly cover the range used in van Wijk and Thijssen (2002) for nominal echo levels going from -6 to 3 dB.

3.1.2. Comparison between TDFFD and state-of-the-art algorithms

In this Section, we compare our TDFFD algorithm to the FFD (Rueckert et al. (1999)) and Ledesma-Carbayo et al. (2005) (LC) algorithms in terms of accuracy for both displacement and strain quantification. All algorithms considered in this section were compared without regularization metric, i.e. without the incompressibility constraint for the TDFFD. This choice was made to avoid biasing the analysis of the relative performance of the different methods under evaluation.
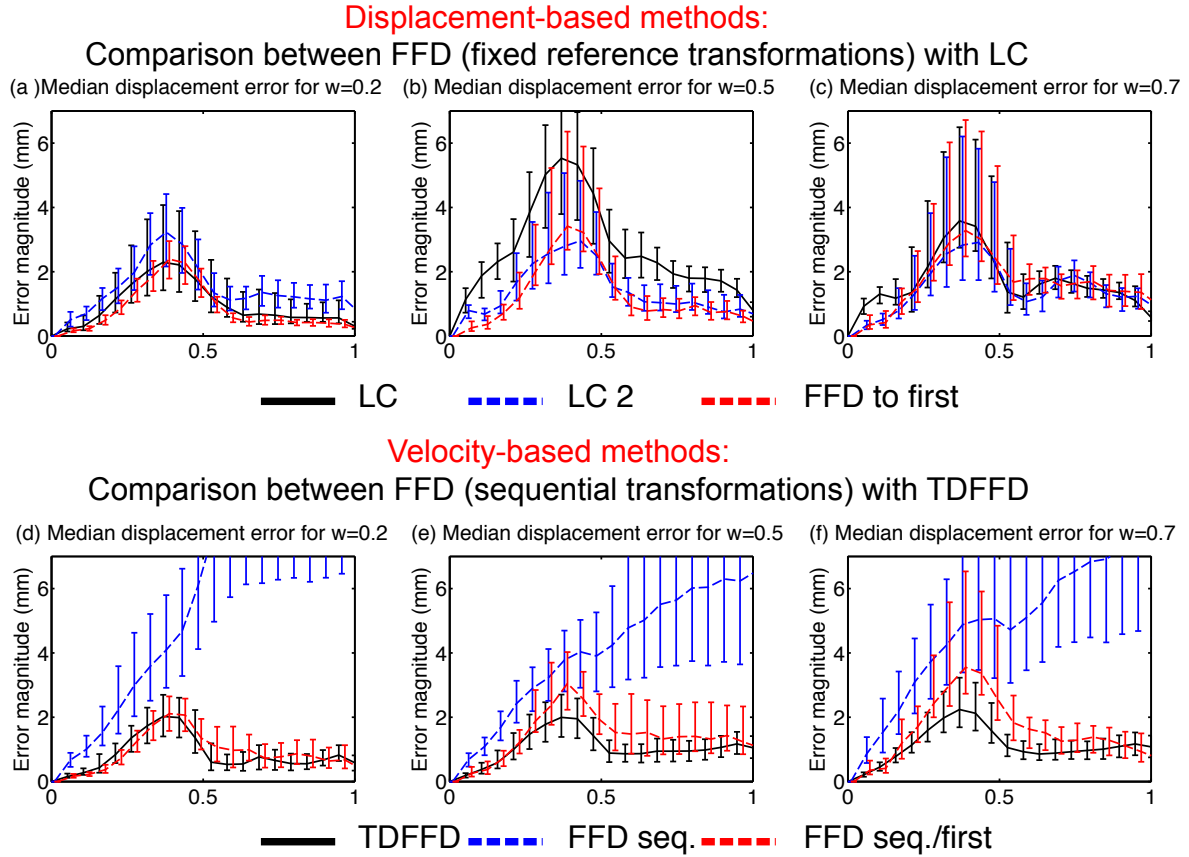
13

Displacement-based methods:
Comparison between FFD (fixed reference transformations) with LC

(a) Median displacement error for w=0.2 (b) Median displacement error for w=0.5 (c) Median displacement error for w=0.7

LC —— LC 2 ---- FFD to first ----

Velocity-based methods:
Comparison between FFD (sequential transformations) with TDFFD

(d) Median displacement error for w=0.2 (e) Median displacement error for w=0.5 (f) Median displacement error for w=0.7

TDFFD —— FFD seq. ---- FFD seq./first ----

Figure 5: Robustness of the 3D displacement quantification to several noise levels, as measured by the proposed TDFFD, LC and FFD algorithms. Vertical bars show the dispersion of the error values over the entire myocardium by plotting second and third quartiles limits. The horizontal axis is the normalized cardiac time (from 0 to 1). The error is measured in millimeters over the entire myocardium. Three levels of noise are considered (see Fig. 4): $w = 0.2$ (left), $w = 0.5$ (center) and $w = 0.7$ (right). First row: comparison between 3D FFD and 3D+t FFD on the displacement fields. *FFD to first* refers to applying FFD between the first frame and every frame of the sequence, *LC* refers to Ledesma-Carbayo et al. (2005) with zero initialization, *LC 2* refers the same algorithm initialized by a set of pairwise transformations. Second row: comparison between 3D FFD and 3D+t FFD on velocities. *FFD seq.* refers to applying FFD sequentially, both for the metric and the transformation, *FFD seq./first* refers to a sequential representation of the transformation but defining the metric to the first frame.

Displacement accuracy versus noise level. Fig. 5 plots the median magnitude of the difference between the ground truth displacement field for all considered algorithms. Vertical bars show the dispersion on the entire myocardium of this displacement error values by plotting second and third quartiles limits. For all algorithms, the B-Spline grid had a resolution of three control points in the longitudinal direction and five in the two transverse directions. This resolution was the optimal initial resolution of the TDFFD algorithm for this set of images, as described in Section 3.1.3.

Two configurations of the FFD algorithm were considered: computing the pairwise transformations between consecutive frames (sequential transformations) and computing them relative to the first frame (fixed reference transformations).

For sequential transformations, two variants for quantifying the image metric were considered. In the first, *FFD seq.* in Fig. 5, the metric was quantified between consecutive frames. In the second, *FFD seq./first* in Fig. 5, the metric was quantified with respect to the first frame. We used the mutual information metric as in the original version of the FFD algorithm (Rueckert et al. (1999)). For each frame, the concatenation of all previously computed transformations was taken as fixed transformation, and was kept constant during

14

	Algorithm	Initialization	Sequential transformation	Sequential metric	Vel. / Disp.
3D	FFD first	Previous registration	No	No	Disp.
	FFD seq.	Zero	Yes	Yes	Vel.
	FFD seq./first	Zero	Yes	No	Vel.
3D+t	LC	Zero	No	No	Disp.
	LC2	FFD first	No	No	Disp.
	TDFFD	Zero	Yes	No	Vel.

Table 2: Summary of the different algorithms under comparison in Section 3.1.2 and Fig. 5. The table specifies for each algorithm: the initialization strategy, if the transformation representation is sequential, if similarity is measured sequentially and if the algorithm quantifies a velocity or a displacement field.

the optimization process. Only the last transformation in the chain, between the previous and current frames, was optimized. This combination of sequential transformations but fixed reference is conceptually the closest choice to the TDFFD approach proposed in this paper.

In Ledesma-Carbayo et al. (2005), the transformation represents the displacement field to the first image as a smooth 3D+t B-Spline field. The sum of squared intensity differences between each image and the first image in the sequence was chosen as similarity measure. Two initialization strategies were compared for the LC method: initializing the transformations with null parameters (LC in Fig. 5) and initializing the 3D+t displacement field by a set of pairwise transformations relative to the first frame ($LC\ 2$ in Fig. 5).

Results are plotted in two rows for incrementally evaluating the impact of temporal continuity and temporal consistency as defined in Section 1.2. First, the advantage of using a continuous representation of the displacement field is evaluated by comparing FFD with fixed reference transformations with Ledesma-Carbayo et al. (2005). The purpose of this first comparison is to evaluate the advantage of going from a 3D to a 3D+t representation of the *displacement field*. Second, sequential FFD is compared to TDFFD, the purpose being to compare a 3D with the respective 3D+t representation of the *velocity field*. The main features of the different algorithms tested in this comparison are summarized in Table 2.

The top row in Fig. 5 shows that for the lowest amount of noise ($w = 0.2$), the LC algorithm and FFD perform similarly. In this case, there is no interest in initializing the 3D+t transformation by a set of pairwise transformations (LC 2) as the initialization by the identity transform (LC) provides a better accuracy. On the contrary, when working at higher noise levels, Ledesma-Carbayo *et al.*'s approach with pairwise initialization gave higher accuracy than the two other algorithms. At the intermediate noise level, the LC algorithm without initialization did not converge and produced higher median error values than FFD and the LC 2 algorithm .

The bottom row in Fig. 5 compares FFD with sequential transformations with the TDFFD algorithm. Applying FFD sequentially for the metric and the transformation resulted in the generation of a significant drift effect and overall divergence. TDFFD and *FFD seq./first* performed similarly at the lowest noise level. However, at the intermediate and highest noise levels, the TDFFD algorithm showed increased robustness, and produced smaller errors. The dispersion was also clearly reduced when using the TDFFD algorithm for all noise levels, as observable from Fig. 5. For $w = 0.7$, the upper limit of the third quartile goes from $6.5\,mm$ using pairwise FFD and the LC algorithm to $3.23\,mm$ using TDFFD.

At all levels of noise, the convergence of the LC algorithm was slower and less regular than TDFFD and all variants of FFD. One possible explanation is that the transformation model used in the LC approach does not exploit the high correlation between the displacements at consecutive times. Searching for optimal directions in the optimization space is therefore more challenging than for the sequential FFD and TDFFD algorithms that naturally represent the displacement at a time $t + \Delta t$ as a composition of the transformation at time t and an incremental contribution over the Δt time interval. As previously mentioned, we used mutual information as similarity metric for the FFD algorithm (Rueckert et al. (1999)). While this has the potential of introducing some bias in the comparison between the different approaches, one can observe in the top row of Fig. 5 that FFD did not show worse performance than LC, despite of using a more generic

15

metric. This suggests that the choice of mutual information does not alter significantly FFD accuracy in this application.

It can be observed from Fig. 5 that all methods compared in this section produce error curves following cardiac cycle phases since they increase over systole and reach a peak of maximum error at the end-systolic phase, being reduced over diastole. This indicates that for the three methods, the error increases when estimating larger displacements.

Myocardial strain accuracy versus noise level. The impact of imaging noise level on radial, longitudinal and circumferential strain curves was evaluated using the same synthetic datasets for the three algorithms (FFD, LC and TDFFD). For each noise level, the configuration of FFD and the LC algorithm giving the best displacement accuracy was taken for comparing the resulting strain outputs.

Fig. 6 shows on the top row, the ground truth strain curves. Vertical bars show the dispersion on the entire myocardium of the strain values by plotting second and third quartiles limits. This dispersion models transmural changes from endocardium to epicardium. It can be neglected for longitudinal strain (inferior to 3 %) but becomes substantial (up to 10 %) for radial and circumferential components. The remainder rows show the strain curves as recovered by the three algorithms for the settings described in Section 3.1.1.

At a low level of noise ($w = 0.2$), TDFFD showed the largest accuracy and smallest dispersion for quantifying longitudinal strain. All three methods underestimated the peak value of radial strain. FFD slightly shifted in time the peak of radial strain. For circumferential strain, FFD and TDFFD performed better (lower dispersion) than the LC algorithm. At intermediate and high noise levels, TDFFD qualitatively preserved better the pattern of the strain curves than the two other algorithms, but at the expense of an underestimation of the peak values and an increased dispersion of strain values. In all three algorithms, the longitudinal strain was the most sensitive to the change in CNR. FFD registration failed to recover meaningful radial and longitudinal strain curves at $w = 0.5$, and at $w = 0.7$.

The FFD strain curves at $w = 0.5$ and $w = 0.7$ showed a more pronounced drift effect (the curves do not come back to zero) in comparison with TDFFD and LC algorithm at similar noise levels. This seems to indicate that spatiotemporal registration techniques are more robust to drift than pairwise algorithms (Elen et al. (2008)). As incompressibility was not used in this experiment, this drift robustness may come from temporal consistency.

In order to quantify the statistical agreement between the different registration algorithms and the ground truth, we performed a set of Wilcoxon rank sum tests for equality of medians at the lowest imaging noise level (w=0.2). These tests were performed on radial, circumferential and longitudinal strain values averaged per AHA segment at the end of systole. All obtained p-values are given in Fig. 7 together with a box plot giving the quartiles of the distributions under comparison. For all algorithms, longitudinal strain returned the lowest p-values. This mainly stems from the high dispersion of the estimated strain values compared to the ground truth. Regarding radial strain, LC and TDFFD gave similar performance for quantifying radial strain at end of systole. FFD in turn underestimated radial strain. Finally for circumferential strain, LC returned incorrect strain patterns in some regions with positive values. This resulted in a high dispersion compared to the ground truth and a lower p-value than FDD and TDFFD.

Displacement accuracy versus time decimation. In addition to the deterioration of spatial CNR, examined in the former experiment, another potential source of image deterioration is the loss of temporal resolution. To evaluate the robustness of all compared algorithms to time resolution of the input data, the synthetic sequence at $w = 0.2$ was undersampled by using one out of every two frames. We then compared the performance of each algorithm on the undersampled data (red curves in Fig. 8) to two alternative situations: running the same algorithm on the full dataset using 1) one control point per time frame (blue curves) and 2) using one temporal control point every two time frames (green curves).

Fig. 8 shows error curves in the same format as in Fig. 5 for the three algorithms. In the case of the TDFFD algorithm, the peak systolic error showed a slight increase (from 1.58 to 1.85 mm) while the average diastolic error went from 0.4 to 0.6 mm. In the case of the LC algorithm, peak systolic error increased from 2.29 to 4 mm. The diastolic error increased in a much higher proportion than for TDFFD (more than a factor 2). For FFD, systolic error was practically unchanged after decimation but the pairwise algorithm failed
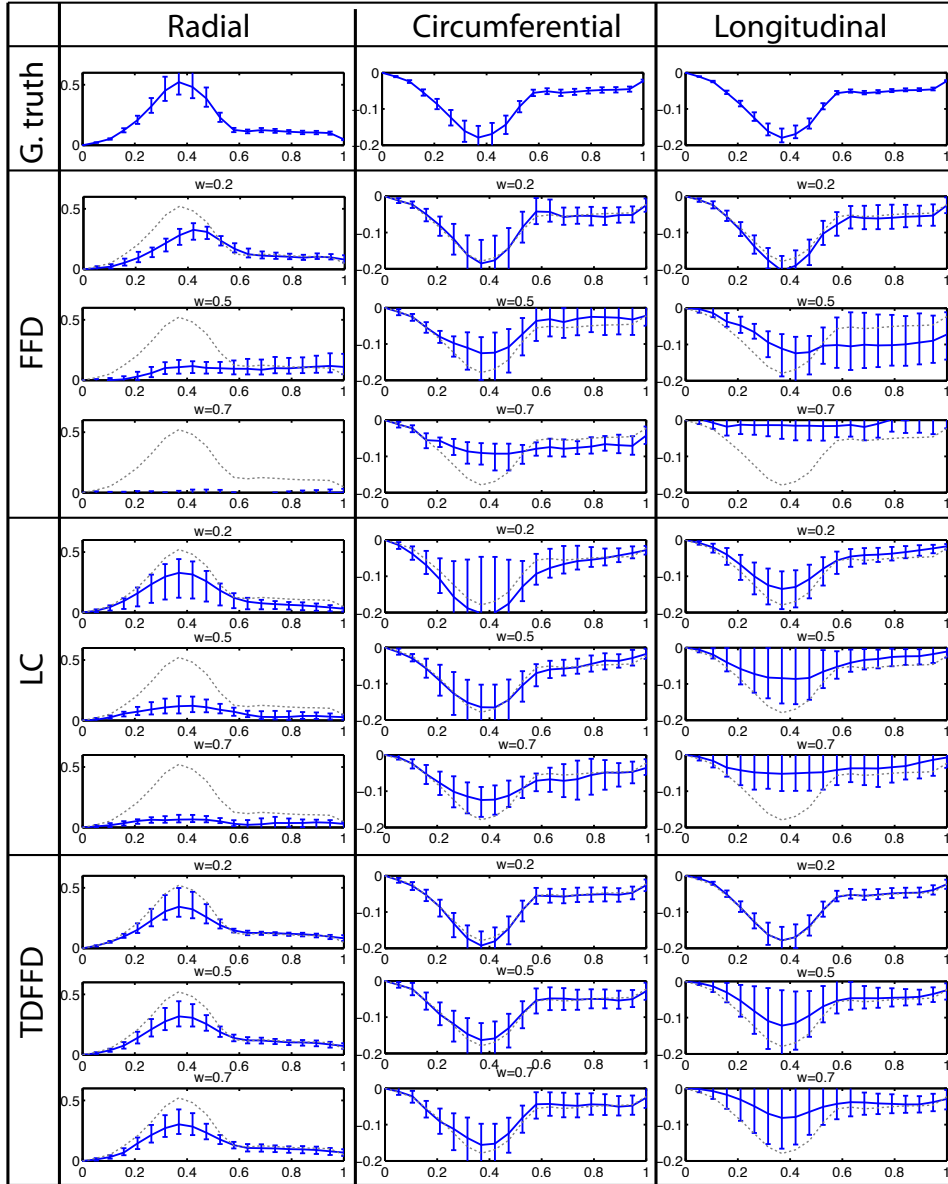
16

Figure 6: Robustness of the 3D strain quantification to several noise levels, as measured by the proposed TDFFD, LC and FFD algorithms. Ground truth curves are overlaid in dashed format. Vertical bars show the dispersion of the strain values over the entire myocardium by plotting second and third quartiles limits. The horizontal axis is the normalized cardiac time (from 0 to 1). Strain is measured in mm over the entire myocardium. Three levels of noise are considered (see Fig. 4): $w = 0.2$, $w = 0.5$ and $w = 0.7$.

to properly recover diastolic motion at the lower temporal resolution. This can be observed in Fig. 8(c), showing a significant increase of error dispersion in the second half of the heart cycle.

3.1.3. Optimization of TDFFD parameters

In this section, we quantify the influence of several parameters of our algorithm on the displacement accuracy. Median displacement values and their dispersion are plotted as in Section 3.1.2.
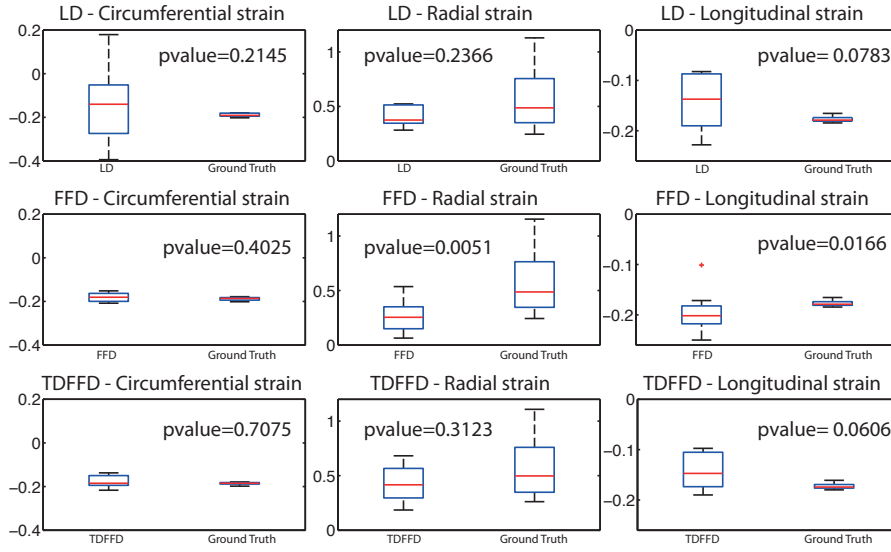
17

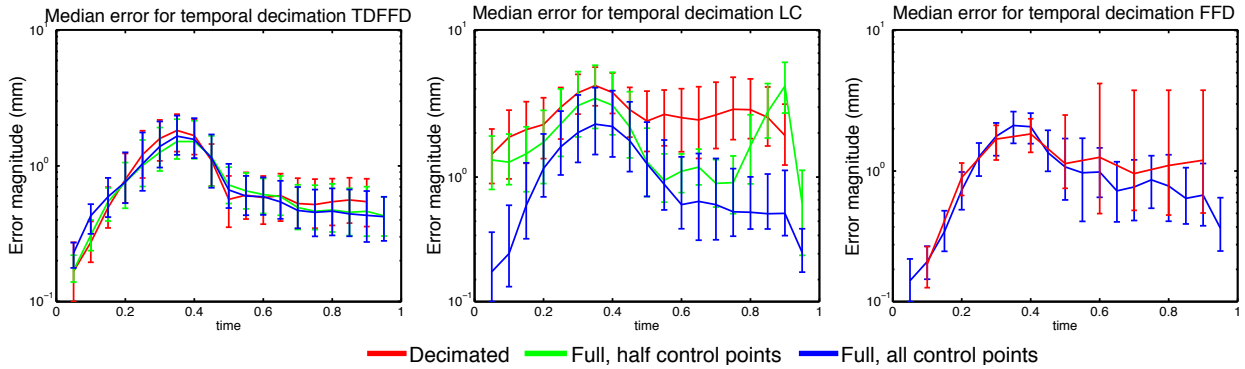Figure 7: Statistical test of equality of median at $w = 0.2$ using the different algorithms.

Figure 8: Robustness of the 3D displacement quantification to temporal decimation of the input sequence, as measured by the proposed TDFFD, LC and FFD algorithms. Vertical bars show the dispersion of the error values over the entire myocardium by plotting second and third quartiles limits. The horizontal axis is the normalized cardiac time (from 0 to 1). The error is measured in millimeters over the entire myocardium. This experiment was performed for $w = 0.2$. For FFD and LC algorithms, we retained the configuration that gave the lowest error from Figure 5. The red curves shows the error magnitude after decimation. The error at full temporal resolution is plotted in blue. For LC and TDFFD algorithms, the error obtained by running the registration using a grid at half temporal resolution on the full dataset is also shown (green curve).

Influence of initial resolution of control points. The purpose of this experiment was to determine the optimal setting for the initial resolution of the 3D+t B-Spline grid. For ensuring convergence of the global optimization procedure, it is important to have an optimal setting of the number of control points in the 3D+t B-Spline grid to avoid solving the optimization problem in an unnecessary low or high dimensional space. In the synthetic dataset described in Section 3.1.1, X and Y axis define the short axis planes. Because of the symmetry of the geometry, the motion is expected to have the same magnitude in X and Y. Therefore, the same number of control points was used in these two directions. Fig. 9(a) shows the displacement accuracy when spanning from 3 to 5 control points in the short and long axis directions. The number of control points in the temporal dimension was fixed to the number of frames in the input sequence. From these curves, it

18

(a) Influence of density
of spatial control points

(b) Influence of density
of temporal control points

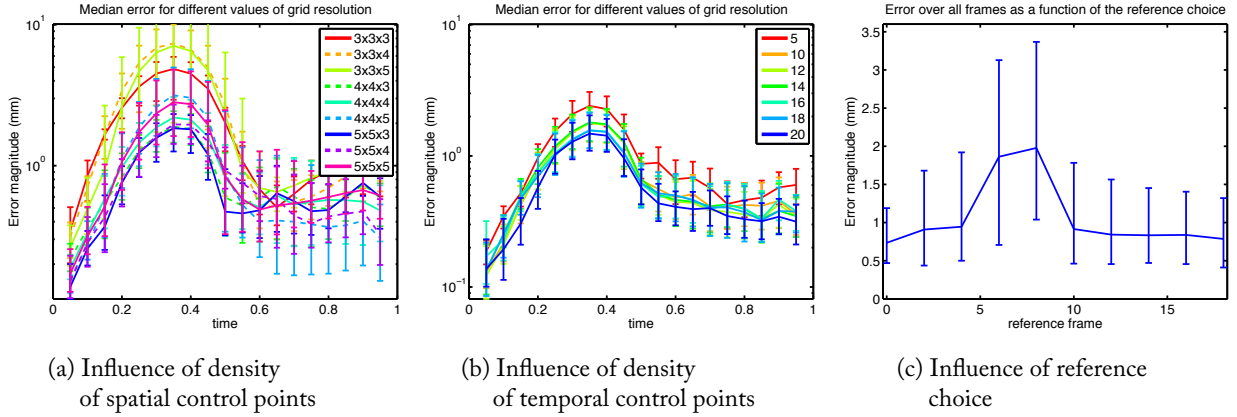(c) Influence of reference
choice

Figure 9: Sensitivity analysis to some algorithmic parameters: Median error over the entire LV volume for (a) different values of spatial grid resolution, (b) different temporal grid resolutions and (c) different choices for the reference frame. Vertical bars show the dispersion of the error values over the entire myocardium by plotting second and third quartiles limits. The horizontal axis is the normalized cardiac time (from 0 to 1). The error is measured in millimeters over the entire myocardium. All these plots were computed from the synthetic images with $w = 0.2$.

can be observed that 5 control points in the short axis direction and 3 control points along the long axis directions returned the lowest error over the whole myocardium.

To analyze the temporal resolution of the TDFFD grid, we varied the temporal resolution by decreasing it from the total number of frames (20 in this case) to 18, 16 14, 12, 10 and 5 control points. The resulting accuracy curves on the displacement errors are plotted in Fig. 9(b). This confirms that an accurate quantification of the motion pattern requires an elevated number of control points. Since taking more control points than the number of frames would result in oversampling the velocity field with respect to the temporal resolution of the input data, we kept the number of control points to its maximum, i.e. the number of frames, in all experiments.

Influence of reference choice. In this experiment, the reference image in the sequence was changed from the first to the last image to assess its influence on the registration results. This corresponds to adjusting the position of the image indexed with 0 in the similarity metric computation with Equation (6). Fig. 9(c) shows the averaged displacement accuracy for every reference choice. Taking the first image in the sequence as reference image (dashed red curve in Fig. 9) provides the lowest error considering all time frames. By symmetry, an equivalent accuracy is obtained when taking the last image of the sequence as reference. Obtaining a similar accuracy when taking the first or the last image as reference suggests that errors introduced by the backward integration of the velocity in time do not significantly alter the accuracy of the estimated trajectories.

3.1.4. Evaluation of the quasi-incompressibility constraint

The influence of the incompressibility weight λ (Equation (9)) on the displacement accuracy was quantified at two resolutions of the B-Spline grid. It is expected that the impact of the incompressibility constraint should be higher when the grid resolution increases. Indeed, the higher the dimensionality of the optimization space, the higher the impact of the regularization term.

The value of λ was first modified in a range from 0 to 2000 for a grid of $5 \times 5 \times 3$ control points (Resolution 1). This grid size was the optimal grid resolution obtained from the previous experiment. Fig. 10(a) plots the displacement accuracy for every λ value. Choosing $\lambda = 1250$ offers a good compromise between reducing systolic and diastolic errors. Indeed, the corresponding curve is the lowest during systole and within the low range of other curves during diastole. The absence of incompressibility constraint gives the worst result in terms of accuracy. During diastole, $\lambda \in [500, 1500]$ gives 30 % less error than $\lambda = 0$.

19

The impact of extending the grid resolution by a factor 2 in all dimensions was then studied. At the refined resolution, it is expected that the impact of the incompressibility would be stronger. We therefore extended the range of incompressibility weights to allow for higher values than the initial resolution. Fig. 10(b) shows the displacement accuracy obtained at this second resolution, namely using a grid of $10 \times 10 \times 6$ control points (Resolution 2). If we assume that 10 % of the total volume is occupied by myocardial tissue, as it is the case in the synthetic dataset used here, then 60 of the 600 cells covered by the control points grid correspond to the myocardium. Since we aim at quantifying strain in each of the 17 segments, it means that each AHA segment is divided in more than 3 cells, corresponding to 16 control points. This can be considered a good trade-off between regularization of the transformation model and required accuracy since for clinical applications, strain will be locally averaged within each segment. At this second resolution, incompressibility was found to have more impact in diastole than in systole. Working without the incompressibility constraint led to worse accuracy at the fine than at the coarse resolution. Giving a weight of 5000 provided a good promise between the reduction of the diastolic error without significantly compromising the end-systolic accuracy.

To verify that the incompressibility constraint is satisfied for the synthetic dataset, we plotted the volume curve (relatively to end of diastole), averaging local volume change on the entire myocardium geometry. The resulting curves at the two consecutive resolutions are given in Fig. 10(c). Averages over time are given for each curve in the caption. These numbers fit in the 5 to 10 % range, in accordance to the volume changes reported by Westerhof et al. (2006); Yin et al. (1996). For further evaluating the effect of the incompressibility constraint, we measured the local volume change of the LV. Fig. 10(d) shows the local volume ratio between end-systole and end-diastole for different incompressibility weights λ at Resolution 2. Light colors indicate small or no volume changes while dark colors correspond to a volume increase (red) or decrease (blue). For incompressible tissue, this value would be equal to 1 on the entire LV. It can be observed how dark color regions fade away when giving more weight to the incompressibility constraint. For instance, dark blue areas (indicating a local compression rate of 50%) are attenuated when introducing the incompressibility constraint.

(a) Resolution 1

(b) Resolution 2

(c) Volume conservation
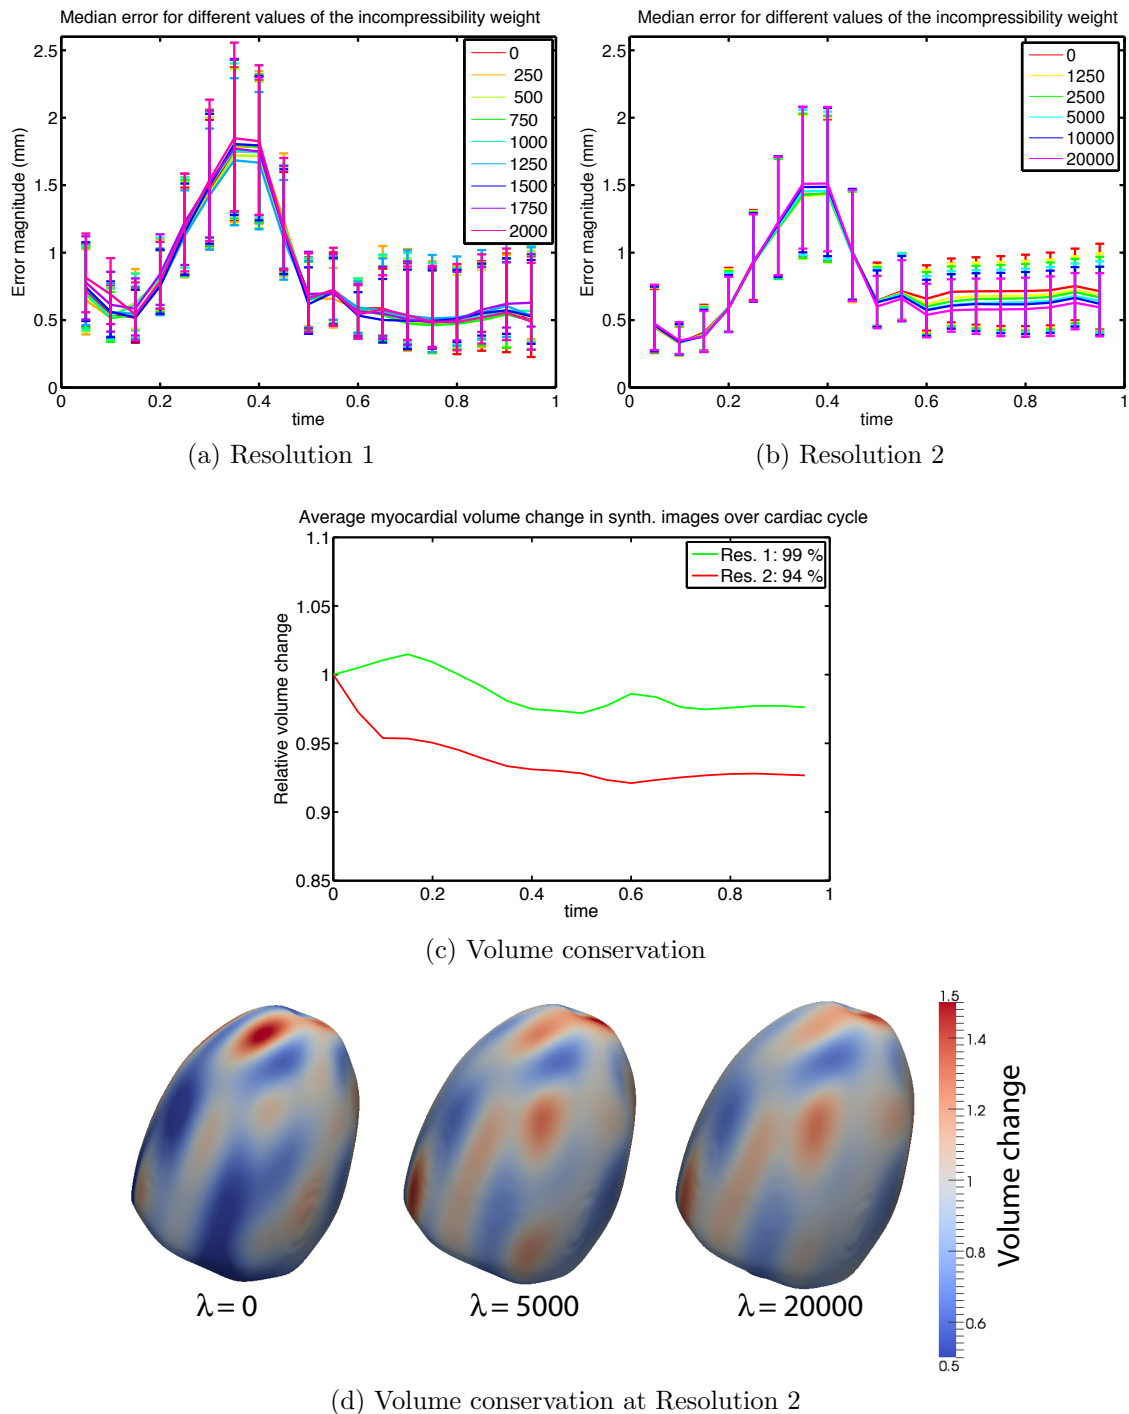
(d) Volume conservation at Resolution 2

Figure 10: Evaluation of the incompressibility constraint at two spatial resolutions of the B-Spline grid: (a) Resolution 1 is the first optimal resolution found from Fig. 9(a), (b) Resolution 2 is obtained from Resolution 1 by dividing the spacing of control points by 2. (c) Average volume change over the entire myocardium domain at Resolutions 1 and 2. (d) Local volume ratio between end of systole and end of diastole for different λ weights at Resolution 2. Exact volume conservation correspond to a volume ratio equal to 1.

21

3.2. Experiments on clinical data

Experiments were performed on two different sets of clinical 3D US sequences. Myocardial strain was first quantified for healthy subjects whose strain pattern and amplitude are known from clinical literature. Then, the TDFFD algorithm was applied to 3D US sequences of CRT candidates and the obtained strain curves were analyzed in conjunction with the clinical outcome of the therapy.

3.2.1. Data acquisition and database

We acquired 3D US sequences in an apical view for two populations, using a GE Healthcare Vivid 7 (Milwaukee, WI, USA) ultrasound system. The first population was composed of 9 healthy volunteers (31 $\pm$ 6 years old), and the second population was composed of 13 patients selected as CRT candidates (61 $\pm$ 8 years old). All 13 patients underwent CRT at a single tertiary care university Hospital (Hospital Clínic i Provincial de Barcelona (HCPB)), based on current European Society of Cardiology clinical guidelines (a recent update of these guidelines was recently published in Dickstein et al. (2010)): left ventricular ejection fraction (LVEF) $< 35\%$, QRS duration $> 120 ms$, and NYHA classification III-IV or NYHA II who covered less than 500 meters in the 6 minutes walking test.

The average number of images per cardiac cycle was 18 ± 3 for the healthy subjects and 17 ± 3 for CRT candidates. The pixel size was on average of $0.9 \times 0.6 \times 0.9\,mm^3$ for the healthy volunteers and $1.0 \times 0.7 \times 1.0\,mm^3$ for the CRT candidates.

Among these patients, there were 7 responders and 6 non-responders according to echocardiographic response. Echocardiographic response quantifies the amount of reverse remodeling by computing the changes in LV volume at the end of systole (LVESV) between baseline and follow-up. Using this criterion, a patient is classified as responder in case of reducing more than 15 % the LVESV, conditioned to patient survival (Yu et al. (2005)). All 3D echocardiographic data were acquired at three different stages: before CRT (OFF), 24-72 hours after implantation (ON), at 6 months (M6), and at 12 months (M12) follow-up. No data was available in the hospital record at M12 for Patient # 11. This patient was admitted to intensive care one week prior to the M6 follow-up and was a non-responder to CRT at M6.

All CRT patients had dilated geometry before implantation, thus the LV did not entirely fit in the field of view. The image acquisition protocol then gave priority to the full inclusion of the septum at the center of the 3D US sector. Therefore, for these patients, myocardial strain in the inferior segments was excluded from the analysis because of their partial inclusion or absence of the field of view.

The exact definition of CRT response is known to be difficult since the application of a single threshold cannot properly render the different grades of response. The reader is invited to refer to Bleeker et al. (2006) for the different definitions of CRT response and their discrepancies in a large population of patients. Because of the complexity of defining CRT response, we distinguished between the following categories: large responder for a LVESV reduction higher than 30 %; moderate responder for a reduction between 15 and 30 %; moderate non-responder for a reduction between 10 and 15 %; and large non-responder when either no reduction was observed or the reduction was below 10 %. The LVESV reduction together with the classification between large/moderate responder/non-responder is given in Table 3. Improvements in both the magnitude and the synchronicity of myocardial strain among the different LV segments are expected for CRT responders (Delgado et al. (2008); Bertola et al. (2009)). This is illustrated for one responder (Patient # 1) in Fig. 11, showing a color map of longitudinal strain at the end of systole before and after CRT. In this patient, the average longitudinal peak strain goes from -6 to -11 %. Before CRT, areas of longitudinal stretching in the lateral wall (in red) coincide with shortening areas (in blue) at end of systole. At follow-up, the spatial distribution of strain becomes uniform over all the LV.

3.2.2. Strain quantification in a population of healthy volunteers

Fig. 12 shows the recovered longitudinal strain curves for the database of 9 healthy subjects at mid and basal segments of the AHA 12 segments (not counting apical segments). The segments either not totally included in the field of view of the 3D US images or suffering from image artifacts (low visibility of the lateral wall, reflections of surrounding anatomical structures, lower spatial resolution on the lateral sides of the US sector), were excluded from the analysis after visual inspection of the images by clinical experts. On

22

Patient	LVESV OFF (mm^3)	LVESV M6 (mm^3)	LVESV M12 (mm^3)	LVESV reduction (%)
Large responders (LVESV reduction > 30 %)				
1	147	83	85	- 42.18
5	164	67	45	- 72.56
7	179	108	113	- 36.87
8	226	60	69	- 69.47
Moderate responders (15 % < LVESV reduction ≤ 30 %)				
4	90	75	67	- 25.56
10	118	77	85	- 27.97
12	144	95	113	- 21.53
13	227	190	164	- 16.30
Moderate non-responders (10 % < LVESV reduction ≤ 15 %)				
3	266	212	233	- 12.40
9	225	196	200	- 11.11
Large non-responders (LVESV reduction < 15 %)				
2	340	307	332	- 02.35
6	165	158	169	+02.42
11	62	140	—	—

Table 3: Values of end-systolic left ventricular volume (LVESV) at baseline (OFF) and follow-up in the population of 13 CRT candidates. The reduction of LV volume defines the echocardiographic response.

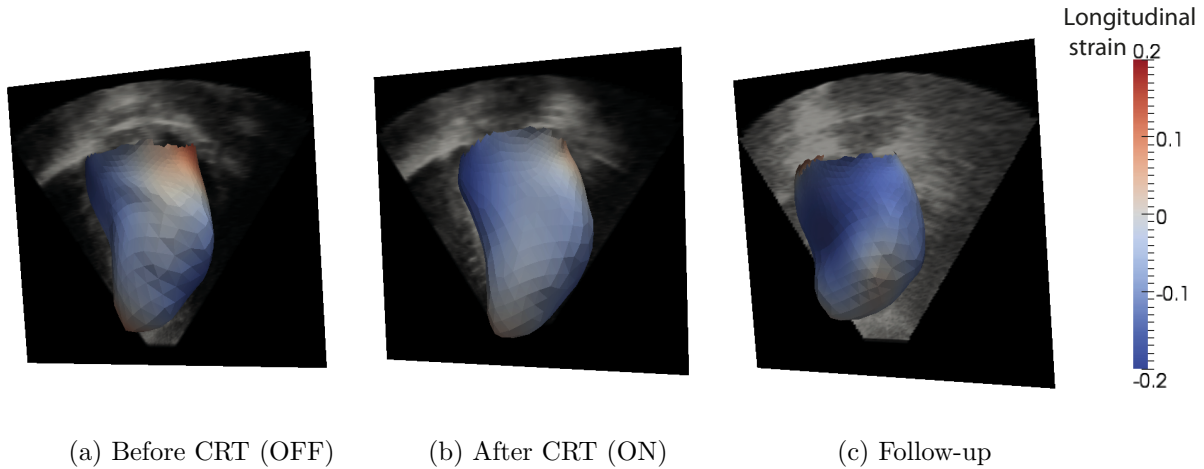(a) Before CRT (OFF)   (b) After CRT (ON)   (c) Follow-up

Figure 11: Longitudinal strain as a color map at end of systole for one CRT responder (Patient # 1) before the therapy (OFF), immediately after CRT (ON), and at M12 follow-up. See an animated version of this figure on `http://mathieu.decraene.info/4strain/Anim.gif`.

average, 78 % of the segments over the 9 volunteers were conserved. TDFFD parameters were set according to the optimal values found for the synthetic dataset from Section 3.1.3. For incompressibility, the optimal λ value found for the synthetic dataset was rescaled according to the initial value of the MSE metric. This ensures that the relative weight of the incompressibility constraint is conserved, despite linear changes of intensity between real and synthetic images.

The recovered strain curves showed a similar pattern in all volunteers, with values in good agreement with clinical literature (Edvardsen et al. (2002)). Typical phases of diastole such as the isovolumetric relaxation and the atrial contraction periods (resulting in an acceleration of myocardial strain at the end
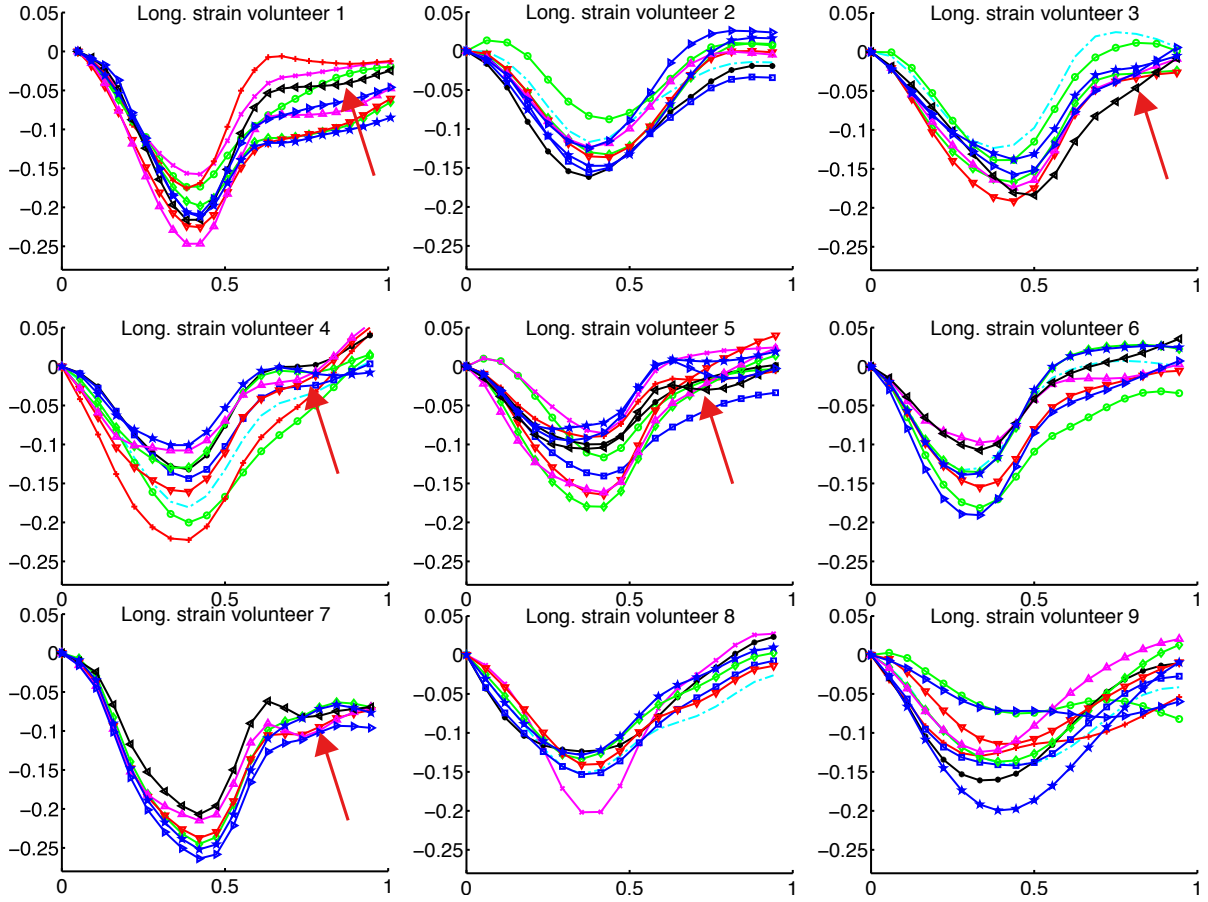
Figure 12: Longitudinal myocardial strain quantified for 9 healthy volunteers in basal and mid segments. The AHA segments are labelled according to the legend given in C. The horizontal axis is the normalized cardiac time (from 0 to 1). The red arrows indicate the acceleration of the strain curve due to atrial contraction, visible in some of the volunteers.

of the diastole, see red arrows in Fig. 12) periods can be distinguished in cases with higher image quality or improved temporal accuracy such as Volunteer #4, #7 and #8. The observed peak systolic longitudinal strain was, on average, of 16.3 %, with a standard deviation of 4.7 %. Edvardsen et al. (2002) (see Table 1 in that reference) reported an average of -17.5 % and a standard deviation of 4 % from tagged MRI images, when considering all but apical segments. If we assume that the values reported by Edvardsen et al. (2002) are the true average and standard deviation, one can compute a normalized $Z \sim N(0,1)$ statistic by subtracting the true mean to the observed average and dividing by the true standard deviation. In this case, we obtained a $p_{value} = p(|Z| > |z_0|) = 0.76$. The average value obtained in this paper is also in agreement with the study by Saito et al. (2009) reporting a mean longitudinal strain of -17.4 ± 5.0 % (yielding a $p_{value} = 0.83$) by speckle tracking on 3D US sequences.

The averaged volume change was computed in a similar way as in Section 3.1.4. The average volume change was computed relatively to end of systole over the whole myocardium and is plotted in Fig. 13. It clearly appears that the total volume change exceeds the expected 5 to 10 % of volume change as reported in the literature (Westerhof et al. (2006); Yin et al. (1996)). While volunteers #1, #2, #3 and #5 show an average volume change close to 10 % range, the other four volunteers exhibit an average volume variation in the range of 14-18%. This point is further commented in Section 4.
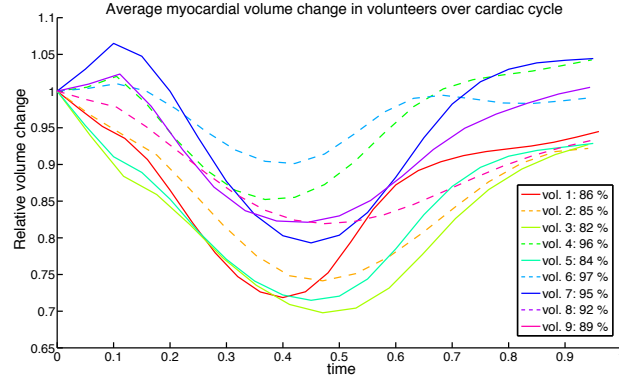
24

Figure 13: Evaluation of myocardial volume conservation over the cardiac cycle for all volunteers.

3.2.3. Strain quantification in CRT candidates before and after implantation

Longitudinal strain curves per segment are reported in Fig. 14 for large echocardiographic responders and in Fig. 15 for large non-responders at OFF, ON, M6 and M12, respectively.

On the patient datasets, TDFFD parameters were set according to the optimal values found for the synthetic dataset from Section 3.1.3, except for the incompressibility weight that was left to 0. This choice was made to avoid local convergence problems close to the border of the imaging domain when computing the incompressibility constraint. This tends to happen more frequently in patients with dilated LV than in volunteers. Moderate responders and non-responders are plotted in Fig. 16 and Fig. 17, respectively. For each patient and each time point, the minimum value of the mean strain and the normalized time corresponding to this minimum are indicated (dashed gray lines). This average value roughly describes the global contraction in the longitudinal direction over the entire LV. Fig. 14 and 15 shows that this value increases at follow-up in all large responders while it stays constant in large non-responders.

Fig. 18 plots the same information using a 3D glyph visualization for Patients #8 and #11, using the same AHA segments as in Figs. 14 and 15. Each ellipsoidal glyph represents the eigen values and eigen vectors of the Cauchy-Green (CG) tensor defined as $\mathbf{D}(\mathbf{x}_0^n)^T \mathbf{D}\mathbf{x}_0^n$. Physically, the CG tensor gives the square of the stretching part of the transformation as the rotational component is compensated when being multiplied by its transposed. A zero deformation corresponds to a CG tensor equal to the identity and a sphere in the glyph representation. The colormap used in Fig. 18 represents longitudinal strain. Negative longitudinal strain is plotted in green, indicating compression, while positive longitudinal strain, indicating expansion, is plotted in red. Patient #8 is a large responder, for which we can observe at follow-up both the reduction of LVESV and a more uniform distribution of longitudinal strain with negative values, compared to baseline. Darker green color also indicates higher contraction at follow-up, reflecting improvements in the cardiac function. Negative and positive values for longitudinal strain coexist at both baseline and follow-up for Patient #11, who is a large non-responder.

Qualitative agreement is observed between the evolution of myocardial strain curves and clinical outcome as quantified by echocardiographic response. As a global trend, it can be observed that responders show an increase between baseline and follow-up both in the amplitude and synchronicity of strain curves among different LV regions. Quantitative measurement of the synchronicity of strain was computed using a covariance index (CI) based on the approach proposed in Silva et al. (2009). The CI between two regional strain curves ϵ_1 and ϵ_2 is computed as

$$CI = \int_{t_{-N_1}}^{t_{N_2}} \epsilon_1(t)\epsilon_2(t)\mathrm{d}t\,, \tag{18}$$

where t_{-N_1} and t_{N_2} represent the limits of the time domain of the input sequence as defined in Section 2.1. A high value of the CI means highly correlated strain curves, while negative or zero value indicates either
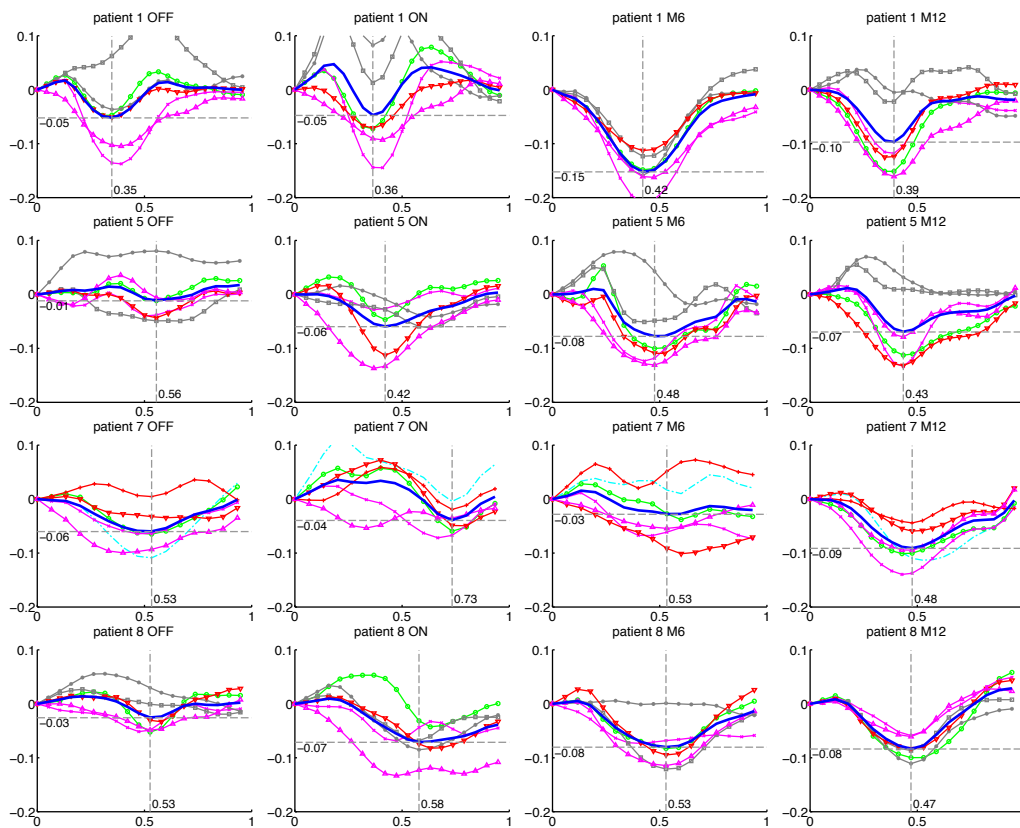
25

Figure 14: Longitudinal myocardial strain quantified for 4 large CRT responders. The curves are plotted before CRT (OFF), right after CRT (ON), and at M6 or M12 follow-up. Each curve corresponds to one AHA segment (see definition in C). Dashed gray lines indicate the minimum value of the average strain and the normalized time corresponding to this minimum. The horizontal axis is the normalized cardiac time (from 0 to 1).

low strain amplitude or that the two strain curves have opposite signs. The CI is sensitive to an overall improvement of strain amplitudes in all segments and to the synchronization of the strain curves among the different segments, which are the two parameters one is interested to quantify after CRT. All possible combinations of segments were considered to generate a set of CI for each input image sequence. The median value and dispersion of the CI are descriptors of the average strain amplitude and homogeneity over all considered segments. The CI was measured at OFF and at M6 and M12. The resulting values are shown using a box plot next to the corresponding strain curves in Fig. 19. The null hypothesis of equal median was computed using the Wilcoxon rank sum test. This non-parametric test was chosen because the assumption of normal distribution was not valid for our population. The test was performed between OFF and M6 and between OFF and M12. Rejections of the null hypothesis, indicating a shift in CI at the 95 % significance level, are indicated in Fig. 19 by an horizontal line between either OFF and M6, either OFF and M12. Interestingly, all responders showed a significant change in median CI at M6 or M12. This corroborates with the fact that all these patients showed reverse remodeling at follow-up, which is a strong indicator of response. Reverse remodeling was quantified using the volumes given in Table 3. In the group of non-responders, no significant shift in CI is observed for all large non-responders (namely Patients #2, #6 and #11). This is in agreement with the fact that these patients did not show any reverse remodeling (LVESV reduction of approximately 2 % for Patients #2 and #6 and a doubling in LVESV at M6 for Patient # 11). The only significant change was detected for Patient #9, who is a moderate non-responder. Hence, CI discriminated in our population large non-responders from responders.
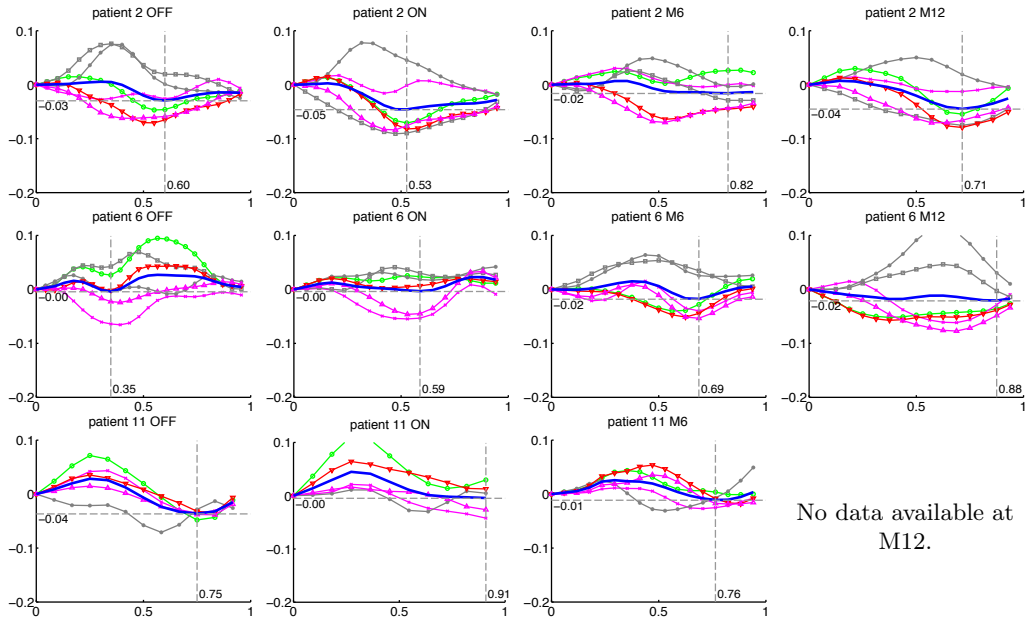
26

Figure 15: Longitudinal myocardial strain quantified for 3 large non-responders to CRT. The curves are plotted before CRT (OFF), right after CRT (ON), and at M6 or M12 follow-up. Each curve corresponds to one AHA segment (see definition in C). Dashed gray lines indicate the minimum value of the average strain and the normalized time corresponding to this minimum. The horizontal axis is the normalized cardiac time (from 0 to 1).

4. Discussion

In this paper, we presented a diffeomorphic registration algorithm that enforces temporal consistency in the recovered displacement field by imposing the temporal continuity of the underlying velocity field. All speckle tracking or registration-based approaches currently in the literature either do not exploit the temporal consistency of myocardial motion over time and treat every frame in a sequential manner (Wang et al. (2010); Kawagishi (2008)), or do not guarantee the transformation to be diffeomorphic (Ledesma-Carbayo et al. (2005)). By enforcing temporal consistency, our algorithm does not require any drift correction as usually applied in sequential approaches (Elen et al. (2008)). With respect to other temporal diffeomorphic registration methods, our approach ensures recovering a velocity field continuous in time, unlike piecewise geodesic solutions given in Khan and Beg (2008); Durrleman et al. (2009). Temporal continuity of the velocity is obtained through regularization by a spatiotemporal kernel while current approaches regularized on the spatial domain only. This avoids solving differential equations at higher orders as proposed in Trouvé and Vialard (2010), which would result in a prohibitive computational cost when applied to dense 3D images.

The image similarity metric referenced in this work uses a fixed reference. Possible extensions to the metric in Equation (6) could consider other frame combinations. For example, taking into account the similarity metric between consecutive frames, or all possible pairwise combinations of frames. The first option can be highly sensitive to drift but prone to capture the motion of speckle patterns that only remain constant for a limited number of frames. The second has the advantage of taking the best of sequential and non-sequential approaches, but is associated to a higher computational cost.

Although it is not the purpose of this paper, we believe that the algorithm we propose has a clear potential for CRT application, namely for improving patient selection. Its improved robustness and accuracy compared to other methods makes it an excellent approach to improve the quantification of mechanical dyssynchrony, currently more based on myocardial velocity timings or in 2D analysis (Fornwalt et al. (2009)). The potential of analyzing the real 3D nature of myocardial mechanics provides an excellent tool to understand and recognize patterns of dyssynchrony that could be eventually corrected with CRT. The link
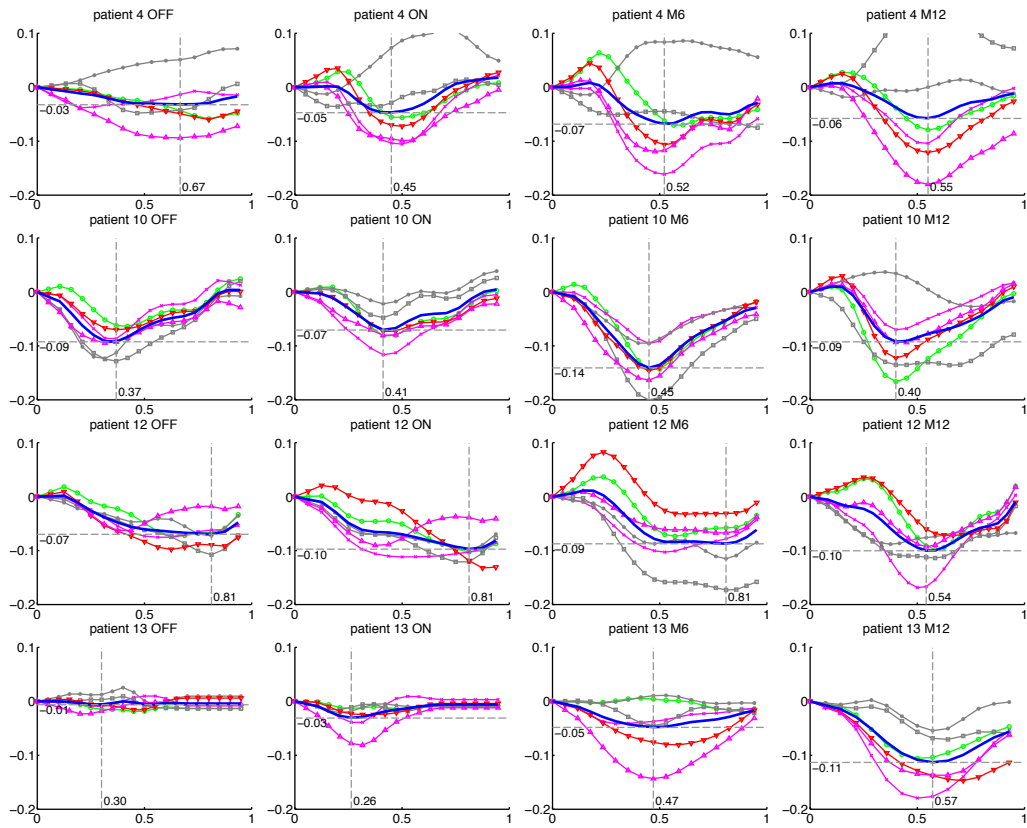
27

Figure 16: Longitudinal myocardial strain quantified for 4 moderate CRT responders. The curves are plotted before CRT (OFF), right after CRT (ON), and at M6 or M12 follow-up. Each curve corresponds to one AHA segment (see definition in C). Dashed gray lines indicate the minimum value of the average strain and the normalized time corresponding to this minimum. The horizontal axis is the normalized cardiac time (from 0 to 1).
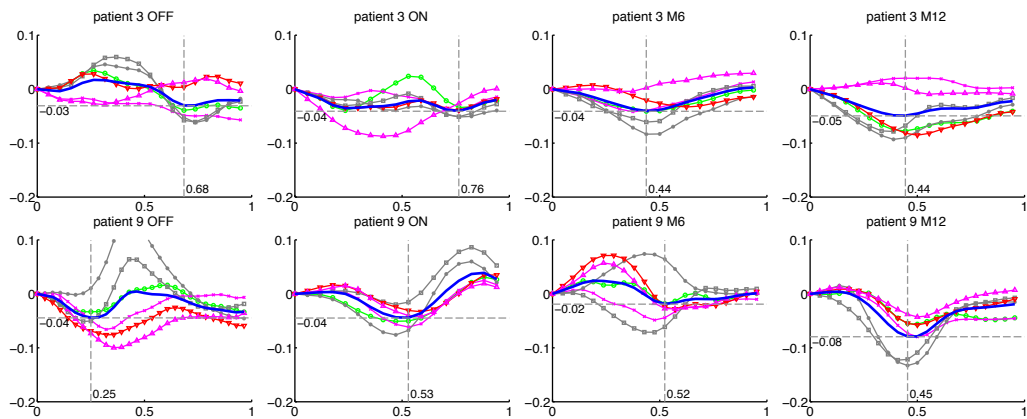

Figure 17: Longitudinal myocardial strain quantified for 2 moderate CRT non-responders. The curves are plotted before CRT (OFF), right after CRT (ON), and at M6 or M12 follow-up. Each curve corresponds to one AHA segment (see definition in C). Dashed gray lines indicate the minimum value of the average strain and the normalized time corresponding to this minimum. The horizontal axis is the normalized cardiac time (from 0 to 1).
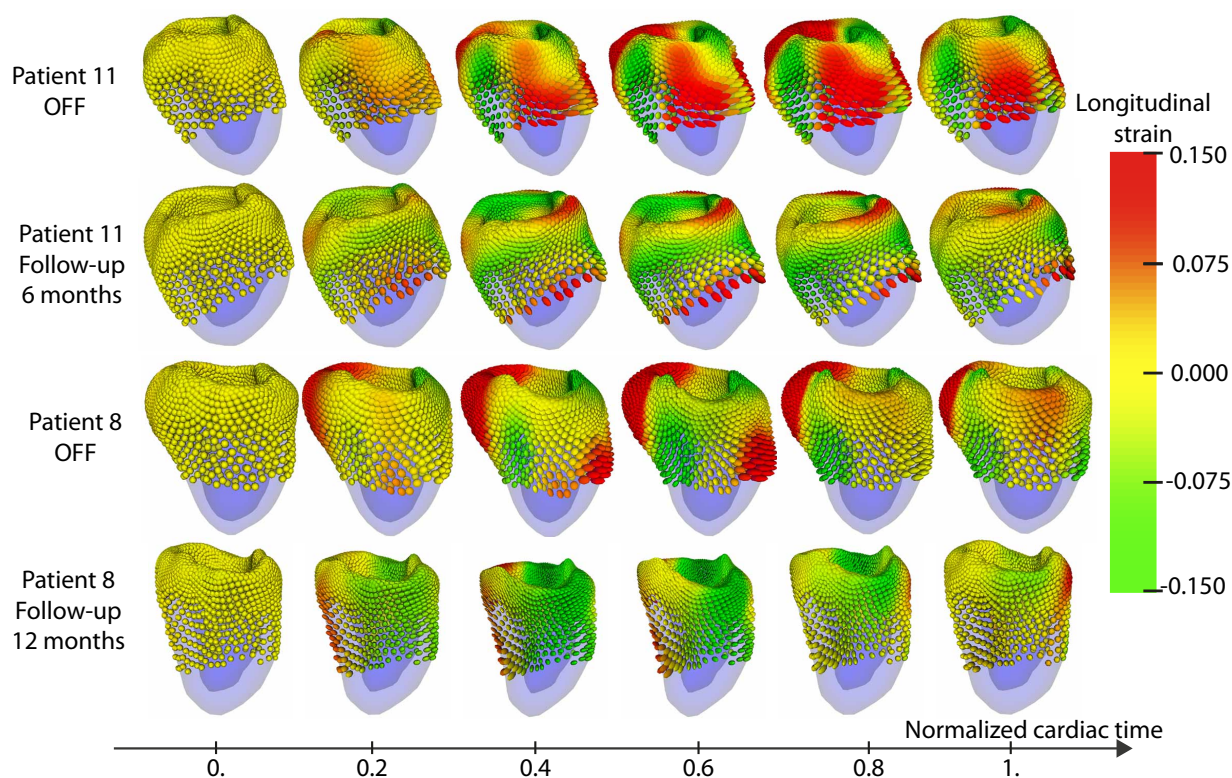
Figure 18: Visualization of the strain tensor using 3D glyphs for Patients #8 and #11 (see text for details). Spherical glyphs as observed at the beginning of the cycle indicate no deformation. Glyphs are colored according the longitudinal strain values, for the AHA regions used in Figures 14 and 15.

between mechanical dyssynchrony and response to CRT is still a controversial issue (Delgado and Bax (2011); Sung and Foster (2011)), particularly limited if a single parameter (measured strain) is considered in the analysis. Comparing the measured strain to the deformation patterns of specific mechanisms involved in CRT response may be a more relevant strategy (Parsai et al. (2009)) and this could be achieved with the use of 3D strain analysis by using such a tool as the one we present.

Experiments were performed on a set of 13 patients with 4 acquisition times each (OFF, ON, M6 and M12), representing a total of 51 sequences. Results showed that strain curves obtained from the TDFFD algorithm were in accordance with CRT outcome at M6 and M12. This was quantified using the CI computed from the longitudinal strain curves recovered by the TDFFD algorithm. In all responders, significant differences between CI at OFF and follow-up were found. In the group of non-responders, the only patient showing a significant CI difference at follow-up was Patient # 9. This patient was classified as moderate non-responder. Although this patient does improve strain-wise, its LVESV reduction is insufficient (11 %) for being classified as responder.

All clinical images used in our experiments were acquired over the last 3 years at HCPB. Data acquisition was not optimized for 3D motion and strain assessment but was representative of a standard, routinely available, image quality in terms of spatiotemporal resolution and SNR. As previously commented, all CRT responders showed an improvement in strain between OFF and follow-up. Improved strain synchronicity between OFF and ON was qualitatively observable in Patients #4, #5, #10 and #13. However, this observation could not be generalized to all responders. This probably indicates that such a differentiation is too challenging for the quality and resolution of datasets currently available in clinical routine. Additionally, it is well established that time-course response to CRT may vary among patients: some of them are acute
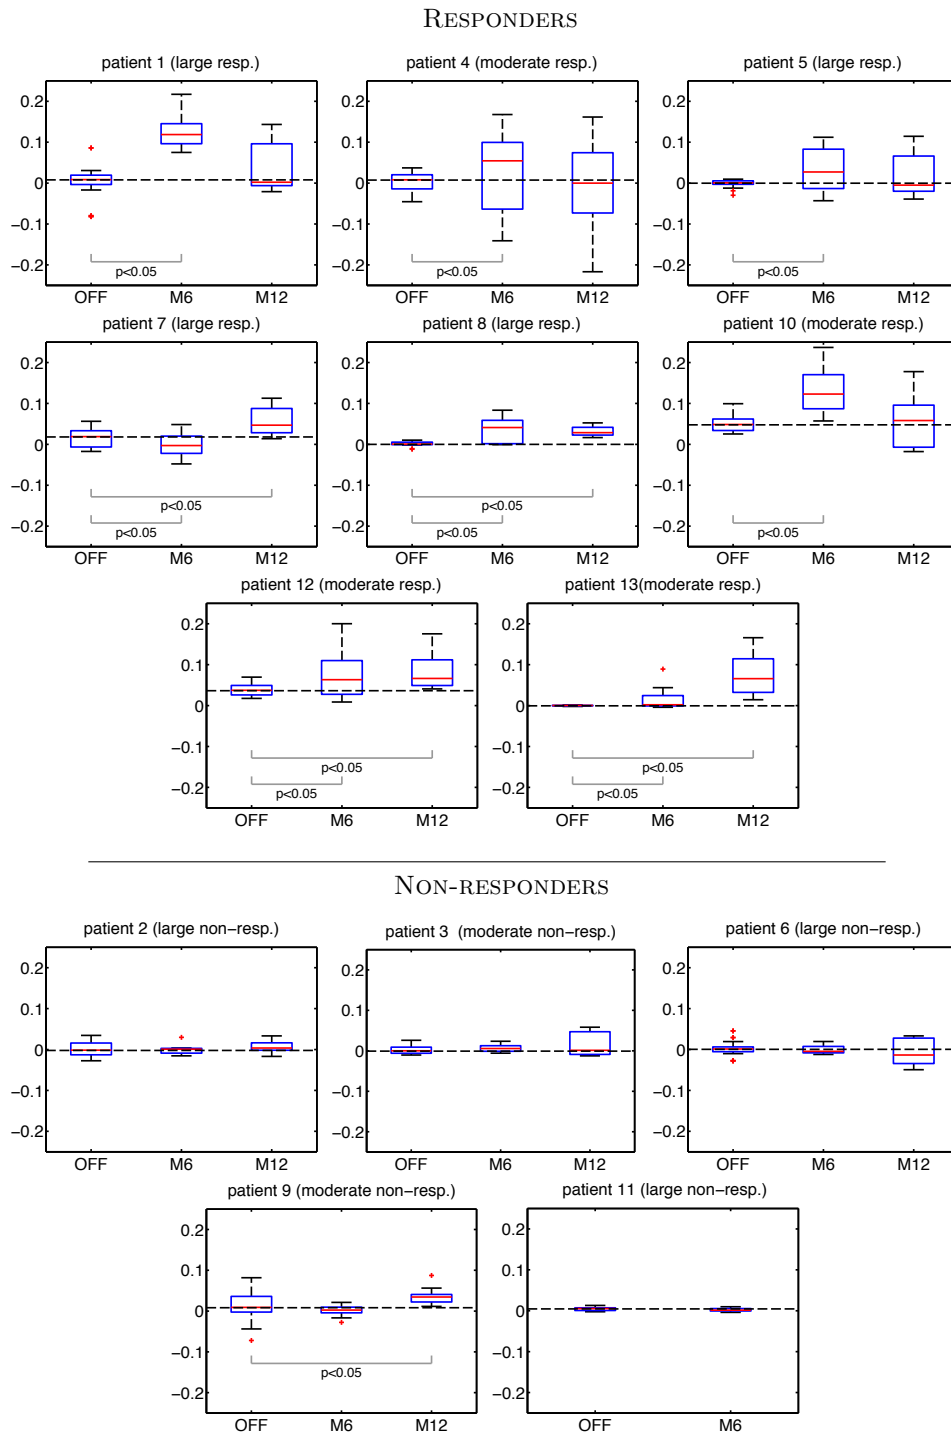
29

Figure 19: Ranges of the covariance index for all the patient population. Average differences between OFF and M6 or OFF and M12 are highlighted when significant at the 5% level.

responders while others do show a later response (Ypenburg et al. (2009)). Thus, we would expect the highest amplitude and best synchronicity at M12, as it is the case for Patients #3, #5, #7, #8, #9, #12 and #13. However, image quality at M6 can be higher than the M12 one, and may temper these conclusions.

The incidence of the different registration parameters on the myocardial motion and strain accuracy were quantified in Section 3.1.3. On the synthetic dataset, incompressibility had a reduced influence on the registration accuracy in comparison with the choice of the reference frame or the initial resolution of the control points grid. One possible reason is that, on this dataset, signal to noise ratio was similar in radial, circumferential and longitudinal directions. As radial thickening is clearly observable from the input images, radial strain is correctly estimated and myocardial volume is relatively constant over time (see Fig. 13). On healthy volunteers, volume change was found to exceed 11% in 4 volunteers over 9. The reason for this is the underestimation of radial strain by our method. This indicates that the soft constraint scheme used in this paper was insufficient to efficiently impose incompressibility over the myocardium volume. Nonetheless, imposing incompressibility as a hard constraint raises several questions. First, Yin et al. (1996) provided experimental evidence that compliance to the incompressibility constraint is changing over the cardiac cycle. Second, some parts of the myocardial tissue, such as papillary muscles are known to be highly compressible. Third, incompressibility has many local optima since any constant velocity field is incompressible. For all these reasons, we made the choice of imposing incompressibility as a weak constraint. Although a hard constraint would be more accurate, it requires to know at which spatial and temporal locations incompressibility is fully reliable.

As visible from Fig. 9, the end-systolic frame was the poorest choice of reference. This can be interpreted in two ways. First, this frame corresponds to the shortest path on average when computing trajectories over time. It appears from Equation (11) that any control point propagates derivatives with respect to **p** to posterior time points using the first term of this equation. The longer the temporal path, the stronger this propagation mechanism. Accuracy was increased when taking the first frame as a reference, which suggests that the higher the coupling between time points, the better with respect to the overall accuracy. Second, the end-systolic frame is the most dissimilar to the other frames and therefore represents a poorest reference choice for the metric computation. Hence, a reference that is closer to the average of the sequence, as the end-diastolic frame, gives higher accuracy.

One limitation of using a smooth temporal kernel is to introduce temporal smoothing in the recovered myocardial motion and strain curves. Our current implementation uses a fixed size kernel that introduces the same amount of smoothing in the entire cardiac cycle. One improvement of the proposed method would consist in tuning the width of the temporal kernel when faster changes in myocardial velocity are expected to occur (e.g. end of systole). Another limitation of this study is to focus the clinical observations only on the longitudinal component of strain. This choice was made because all clinical 3D image datasets used in this study were acquired from an apical 4 chambers view. In such a setting, CNR is better in the longitudinal direction, and longitudinal deformation can be more accurately estimated than radial deformation.

5. Conclusion

We have described a new temporal diffeomorphic registration method called TDFFD and applied it to recover motion and strain from 3D echocardiography images. By introducing in the velocity computation a continuous kernel over time, we enforce temporal consistency in the recovered displacements. This can be particularly useful in image sequences with substantial amount of noise and artifacts. The TDFFD algorithm was applied to the quantification of strain in 3D US using synthetic datasets, and clinical images from healthy subjects and patients undergoing CRT.

Experiments on synthetic 3D US datasets showed an improved robustness and accuracy in the quantification of motion and strain for low values of CNR compared to a classical pairwise approach and one of its extension to recover 3D+t displacement fields. On healthy volunteers, the method provided physiologically meaningful longitudinal strain curves with small dispersion among LV segments. On CRT patients, the quantification performed before and after CRT, including follow-ups at 6 and 12 months showed agreement between the quantified myocardial strain curves and clinical outcome. This agreement was quantified using

31

covariance indexes of average and segmental strain curves. Future extension of the TDFFD algorithm will focus on incorporating similarity metrics adapted to statistical characteristics of US speckle noise (Myronenko et al. (2009)) and its extension to incorporate compounding strategies to improve the limited field-of-view in 3D US sequences of heart failure patients with dilated LV similar to Piella et al. (2011).

Acknowledgments

A. Derivation of the incompressibility constraint

In this appendix, we detail the computation of $\frac{\mathrm{d}}{\mathrm{d}\mathbf{p}^{i,j,k,l}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big)$ required to compute the derivative of the incompressibility constraint from Equation (13).

If $\mathbf{x}_0^n(\mathbf{p}) = (x, y, z)$ and $\widetilde{x}_i = \frac{x-q^i}{\Delta_i}$, $\widetilde{y}_j = \frac{y-q^j}{\Delta_j}$, $\widetilde{z}_k = \frac{z-q^j}{\Delta_k}$, and $\widetilde{t}_l = \frac{t-q^l}{\Delta_l}$, the trace of the Jacobian in Equation (13) can be written as

$$
\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big) = \sum_{i,j,k,l} \frac{\beta'(\widetilde{x}_i)\beta(\widetilde{y}_j)\beta(\widetilde{z}_k)\beta(\widetilde{t}_l)\ p_x^{i,j,k,l}}{\Delta_i}
$$
$$
+ \frac{\beta(\widetilde{x}_i)\beta'(\widetilde{y}_j)\beta(\widetilde{z}_k)\beta(\widetilde{t}_l)\ p_y^{i,j,k,l}}{\Delta_j} \tag{19}
$$
$$
+ \frac{\beta(\widetilde{x}_i)\beta(\widetilde{y}_j)\beta'(\widetilde{z}_k)\beta(\widetilde{t}_l)\ p_z^{i,j,k,l}}{\Delta_k} \quad,
$$

where β' designates the derivative of the 1D kernel defined in Equation (1). The derivative of Equation (19) to compute the derivative of Equation (13) is

$$
\frac{\mathrm{d}}{\mathrm{d}\mathbf{p}^{i,j,k,l}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big) = \frac{\partial}{\partial\mathbf{p}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big) + \frac{\partial}{\partial\mathbf{x}_0^n}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big)\frac{\mathrm{d}\mathbf{x}_0^n}{\mathrm{d}\mathbf{p}}. \tag{20}
$$

The first term in Equation (20) can be easily obtained from Equation (19):

$$
\frac{\partial}{\partial\mathbf{p}}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big) = \Big(\frac{\beta'(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)\cdot\beta(\widetilde{t}_l)}{\Delta_i}, \frac{\beta(\widetilde{x}_i)\cdot\beta'(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)\cdot\beta(\widetilde{t}_l)}{\Delta_j}, \frac{\beta(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta'(\widetilde{z}_k)\cdot\beta(\widetilde{t}_l)}{\Delta_k}\Big).
$$
$$\tag{21}$$

The second term in Equation (20) is more complex and involves second-order terms. It can be expanded as

$$
\frac{\partial}{\partial\mathbf{x}_0^n}\Big(\mathrm{tr}\big(\mathbf{D}\mathbf{v}^n(\mathbf{p})\big)\Big)\frac{\mathrm{d}\mathbf{x}_0^n}{\mathrm{d}\mathbf{p}} =
$$
$$
\Big(\frac{\beta''(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)p_x^{i,j,k,l}}{\Delta_i^2} + \frac{\beta'(\widetilde{x}_i)\cdot\beta'(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)p_y^{i,j,k,l}}{\Delta_i\Delta_j} + \frac{\beta'(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta'(\widetilde{z}_k)p_z^{i,j,k,l}}{\Delta_i\Delta_k},
$$
$$
\frac{\beta'(\widetilde{x}_i)\cdot\beta'(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)p_x^{i,j,k,l}}{\Delta_i\Delta_j} + \frac{\beta(\widetilde{x}_i)\cdot\beta''(\widetilde{y}_j)\cdot\beta(\widetilde{z}_k)p_y^{i,j,k,l}}{\Delta_j^2} + \frac{\beta(\widetilde{x}_i)\cdot\beta'(\widetilde{y}_j)\cdot\beta'(\widetilde{z}_k)p_z^{i,j,k,l}}{\Delta_j\Delta_k},
$$
$$
\frac{\beta'(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta'(\widetilde{z}_k)p_x^{i,j,k,l}}{\Delta_i\Delta_k} + \frac{\beta(\widetilde{x}_i)\cdot\beta'(\widetilde{y}_j)\cdot\beta'(\widetilde{z}_k)p_y^{i,j,k,l}}{\Delta_j\Delta_k} + \frac{\beta(\widetilde{x}_i)\cdot\beta(\widetilde{y}_j)\cdot\beta''(\widetilde{z}_k)p_z^{i,j,k,l}}{\Delta_k^2}\Big)\cdot\beta(\widetilde{t}_l)\cdot\frac{\mathrm{d}\mathbf{x}_0^n}{\mathrm{d}\mathbf{p}},
$$
$$\tag{22}$$

where the factor $\frac{\mathrm{d}\mathbf{x}_0^n}{\mathrm{d}\mathbf{p}}$ of Equation (22) can be obtained from Equation (12) and β'' stands for the second order derivative of the 1D kernel β defined in Equation (1).

33

B. Notations and acronyms

Symbol	Meaning
n	Index of an image in the sequence
$\mathcal{N} = \{n \in \mathbb{Z}, -N_1 \leq n \leq N_2\}$	Set of possible values of n. Index 0 designates the reference
I_n	Intensity function of the n^{th} image in the sequence
$\Omega \subset \mathbb{R}^d$	Spatial domain for each image in the sequence
$\Omega_M \subset \mathbb{R}^d$	Spatial domain with incompressibility constraint
$\mathcal{T} = [t_{-N1}, t_{N2}] \subset \mathbb{R}$	Temporal domain
t_n	Time associated to image I_n
$\{t_k\}$	Sampling of the temporal domain used for computing trajectories
$\mathbf{v}^k(\mathbf{p}) = \mathbf{v}(\mathbf{x}, t_k, \mathbf{p})$	Velocity field at spatiotemporal location $\mathbf{x}, t$ parameterized by $\mathbf{p}$
$\mathbf{x}_0^k(\mathbf{p}) := \boldsymbol{\varphi}_0(\mathbf{x}, t_k, \mathbf{p})$	Trajectory of a point $\mathbf{x}$ from time $t = 0$ to time t_k
$\mathbf{q}^{i,j,k,l} := (q^i, q^j, q^k, q^l)$	Grid of control points
$\mathbf{D}$	Jacobian operator on spatial dimensions
$\boldsymbol{\epsilon} = \begin{pmatrix} \epsilon_{rr} & \epsilon_{rc} & \epsilon_{rl} \\ \epsilon_{cr} & \epsilon_{cc} & \epsilon_{cl} \\ \epsilon_{lr} & \epsilon_{lc} & \epsilon_{ll} \end{pmatrix}$	Strain tensor (r=radial, c=circumferential, l=longitudinal)

Acronym	Meaning
AHA	American Heart Association
CI	Covariance Index
CNR	Contrast to Noise Ratio
CRT	Cardiac Resynchronization Therapy
ES	End of Systole
ED	End of Diastole
FFD	Free Form Deformation
LV	Left Ventricle
LVESV	Left Ventricular End of Systole Volume
MRI	Magnetic Resonance Imaging
MSE	Mean Squared Error Metric
TDFFD	Temporal Diffeomorphic Free Form Deformation
US	Ultrasound

C. AHA segments symbols in strain plots

Symbol	AHA segment	Name
.	1	Basal anterior
∘	2	Basal anteroseptal
×	3	Basal inferoseptal
+	4	Basal inferior
∗	5	Basal inferolateral
□	6	Basal anterolateral
◇	7	Mid anterior
▽	8	Mid anteroseptal
△	9	Mid inferoseptal
◁	10	Mid inferior
▷	11	Mid inferolateral
⋆	12	Mid anterolateral

References

Abe, Y., Kawagishi, T., Ohuchi, H., 2008. Accurate detection of regional contraction using novel 3-dimensional speckle tracking technique. Journal of the American College of Cardiology 51, Supplement A, A116.

Andrade, J., Cortez, L., Campos, O., Arruda, A., Pinheiro, J., Vulcanis, L., Shiratsuchi, T., Kalil-Filho, R., Cerri, G., 2011. Left ventricular twist: comparison between two-and three-dimensional speckle-tracking echocardiography in healthy volunteers. European Journal of Echocardiography 12, 76.

34

Arts, T., Hunter, W., Douglas, A., Muijtjens, A., Reneman, R., 1992. Description of the deformation of the left ventricle by a kinematic model. Journal of biomechanics 25, 1119–1127.

Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. Neuroimage 38, 95–113.

Avants, B., Epstein, C., Grossman, M., Gee, J., 2008. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis 12, 26–41.

Beg, M.F., Miller, M.I., Trouvé, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61, 139–157.

Bertola, B., Rondano, E., Sulis, M., Sarasso, G., Piccinino, C., Marti, G., Devecchi, P., Magnani, A., Francalacci, G., Marino, P., 2009. Cardiac dyssynchrony quantitated by time-to-peak or temporal uniformity of strain at longitudinal, circumferential, and radial level: implications for resynchronization therapy. Journal of the American Society of Echocardiography 22, 665–671.

Bistoquet, A., Oshinski, J., Skrinjar, O., 2007. Left ventricular deformation recovery from cine MRI using an incompressible models. IEEE Transactions on Medical Imaging 26, 1136–1153.

Bleeker, G.B., Bax, J.J., Fung, J.W.H., van der Wall, E.E., Zhang, Q., Schalij, M.J., Chan, J.Y.S., Yu, C.M., 2006. Clinical versus echocardiographic parameters to assess response to cardiac resynchronization therapy. The American Journal of Cardiology 97, 260 – 263.

Butakoff, C., Balocco, S., Ordas, S., Frangi, A.F., 2007. Simulated 3D ultrasound LV cardiac images for active shape model training, in: Pluim, J.P.W., Reinhardt, J.M. (Eds.), Medical Imaging 2007: Image Processing, SPIE. p. 65123U.

Byrd, R., Lu, P., Nocedal, J., Zhu, C., 1995. A limited memory algorithm for bound constrained optimization. SIAM Journal on Scientific Computing 16, 1190–1208.

Cerqueira, M.D., Weissman, N.J., Dilsizian, V., Jacobs, A.K., Kaul, S., Laskey, W.K., Pennell, D.J., Rumberger, J.A., Ryan, T., Verani, M.S., 2002. Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart: A Statement for Healthcare Professionals From the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association. Circulation 105, 539–542.

Chandrashekara, R., Mohiaddin, R., Rueckert, D., 2004. Cardiac motion tracking in tagged MRI images using a 4D B-spline motion model and nonrigid image registration, in: International Symposium on Biomedical Imaging, IEEE. pp. 468 – 471.

Cohen, B., Dinstein, I., 2002. New maximum likelihood motion estimation schemes for noisy ultrasound images. Pattern Recognition 35, 455 – 463.

De Craene, M., Piella, G., Duchateau, N., Silva, E., Doltra, A., Gao, H., D'hooge, J., Camara, O., Brugada, J., Sitges, M., Frangi, A.F., 2010. Temporal diffeomorphic free-form deformation for strain quantification in 3D-US images, in: Jiang, T., Navab, N., Pluim, J., Viergever, M. (Eds.), Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 1–8.

Delgado, V., Ypenburg, C., van Bommel, R., Tops, L., Mollema, S., Marsan, N., Bleeker, G., Schalij, M., Bax, J., 2008. Assessment of left ventricular dyssynchrony by speckle tracking strain imaging:: Comparison between longitudinal, circumferential, and radial strain in cardiac resynchronization therapy. Journal of the American College of Cardiology 51, 1944–1952.

Delgado, V., Ypenburg, C., Zhang, Q., Mollema, S., Fung, J., Schalij, M., Yu, C., Bax, J., 2009. Changes in global left ventricular function by multidirectional strain assessment in heart failure patients undergoing cardiac resynchronization therapy. Journal of the American Society of Echocardiography 22, 688 – 694.

Delgado, V., Bax, J., 2011. Assessment of systolic dyssynchrony for cardiac resynchronization therapy is clinically useful. Circulation 123, 640.

D'hooge, J., Heimdal, A., Jamal, F., Kukulski, T., Bijnens, B., Rademakers, F., Hatle, L., Suetens, P., Sutherland, G., 2000. Regional strain and strain rate measurements by cardiac ultrasound: Principles, implementation and limitations. European Journal of Echocardiography 1, 154–170.

Authors/Task Force Members, Dickstein, K., Vardas, P.E., Auricchio, A., Daubert, J.C., Linde, C., McMurray, J., Ponikowski, P., Priori, S.G., Sutton, R., van Veldhuisen, D.J., for Practice Guidelines (CPG), E.C., Vahanian, A., Auricchio, A., Bax, J., Ceconi, C., Dean, V., Filippatos, G., Funck-Brentano, C., Hobbs, R., Kearney, P., McDonagh, T., Popescu, B.A., Reiner, Z., Sechtem, U., Sirnes, P.A., Tendera, M., Vardas, P., Widimsky, P., Reviewers, D., Tendera, M., Anker, S.D., Blanc, J.J., Gasparini, M., Hoes, A.W., Israel, C.W., Kalarus, Z., Merkely, B., Swedberg, K., Camm, A.J., 2010. 2010 focused update of ESC guidelines on device therapy in heart failure. Europace 12, 1526–1536.

Duan, Q., Angelini, E., Homma, S., Laine, A., 2007. Validation of optical-flow for quantification of myocardial deformations on simulated RT3D ultra-sound, in: International Symposium on Biomedical Imaging, IEEE. pp. 944–947.

Duckett, S.G., Ginks, M., Knowles, B.R., Chiribiri, A., Ma, Y.L., Razavi, R., Schaeffter, T., Carr-White, G., Rinaldi, C.A., Rhode, K., 2010. A Novel Cardiac MRI Protocol To Guide Successful Cardiac Resynchronization Therapy Implantation. Circulation: Heart Failure 3, e18–e21.

Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., Ayache, N., 2009. Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets, in: Yang, G.Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (Eds.), Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 297–304.

Edvardsen, T., Gerber, B.L., Garot, J., Bluemke, D.A., Lima, J.A., Smiseth, O.A., 2002. Quantitative assessment of intrinsic regional myocardial deformation by doppler strain rate echocardiography in humans: Validation against three-dimensional tagged magnetic resonance imaging. Circulation 106, 50–56.

Elen, A., Choi, H., Loeckx, D., Gao, H., Claus, P., Suetens, P., Maes, F., D'hooge, J., 2008. Three-dimensional cardiac strain estimation using spatio-temporal elastic registration of ultrasound images: a feasibility study. IEEE Transactions on Medical Imaging 27, 1580–1591.

Fornwalt, B.K., Delfino, J.G., Sprague, W.W., Oshinski, J.N., 2009. It's time for a paradigm shift in the quantitative evaluation of left ventricular dyssynchrony. Journal of the American Society of Echocardiography 22, 672 – 676.

35

Gao, H., Choi, H., Claus, P., Boonen, S., Jaecques, S., van Lenthe, G., Van der Perre, G., Lauriks, W., D'hooge, J., 2009. A fast convolution-based methodology to simulate 2D/3D cardiac ultrasound images. IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control 56, 404–409.

Gorce, J., Friboulet, D., Magnin, I., 1997. Estimation of three-dimensional cardiac velocity fields: assessment of a differential method and application to three-dimensional CT data. Medical Image Analysis 1, 245–261.

Grau, V., Becher, H., Noble, A., 2007. Registration of multiview real-time 3-d echocardiographic sequences. IEEE Transactions on Medical Imaging 26, 1154–1165.

Haber, E., Modersitzki, J., 2004. Numerical methods for volume preserving image registration. Inverse problems 20, 1621.

Hernandez, M., Bossa, M., Olmos, S., 2009. Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. International Journal of Computer Vision 85, 291–306.

Hinkle, J., Fletcher, P., Wang, B., Salter, B., Joshi, S., 2009. 4D map image reconstruction incorporating organ motion, in: Prince, J., Pham, D., Myers, K. (Eds.), Information Processing in Medical Imaging, Springer. pp. 676–687.

Kapetanakis, S., Kearney, M., Siva, A., Gall, N., Cooklin, M., Monaghan, M., 2005. Real-time three-dimensional echocardiography: a novel technique to quantify global left ventricular mechanical dyssynchrony. Circulation 112, 992–1000.

Kawagishi, T., 2008. Speckle tracking for assessment of cardiac motion and dyssynchrony. Echocardiography 25, 1167–1171.

Khan, A., Beg, M., 2008. Representation of time-varying shapes in the large deformation diffeomorphic framework, in: International Symposium on Biomedical Imaging, IEEE. pp. 1521–1524.

Ledesma-Carbayo, M.J., Kybic, J., Desco, M., Santos, A., Suhling, M., Hunziker, P., Unser, M., 2005. Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation. IEEE Transactions on Medical Imaging 24, 1113–1126.

Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N., 2010. Logdemons revisited: Consistent regularisation and incompressibility constraint for soft tissue tracking in medical images, in: Jiang, T., Navab, N., Pluim, J., Viergever, M. (Eds.), Medical Image Computing and Computer-Assisted Intervention. Springer. volume 6362 of *Lecture Notes in Computer Science*, pp. 652–659.

Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N., 2011. ilogdemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. International Journal of Computer Vision 92, 92–111.

Metz, C., Klein, S., Schaap, M., van Walsum, T., Niessen, W., 2011. Nonrigid registration of dynamic medical imaging data using nD+t B-splines and a groupwise optimization approach. Medical Image Analysis 15 (2), 238–249.

Muraru, D., Badano, L., Gianfagna, P., Ermacora, D., Proclemer, A., Iliescu, C., 2010. Sources of variation and bias in assessing left ventricular dyssynchrony using three-dimensional echocardiography, in: European Heart Journal, pp. 876–876.

Myronenko, A., Song, X., Sahn, D., 2009. Maximum likelihood motion estimation in 3D echocardiography through non-rigid registration in spherical coordinates, in: Functional Imaging and Modeling of the Heart, Springer. pp. 427–436.

Parsai, C., Bijnens, B., Sutherland, G.R., Baltabaeva, A., Claus, P., Marciniak, M., Paul, V., Scheffer, M., Donal, E., Derumeaux, G., Anderson, L., 2009. Toward understanding response to cardiac resynchronization therapy: left ventricular dyssynchrony is only one of multiple mechanisms. European Heart Journal 30, 940–949.

Perperidis, D., Mohiaddin, R.H., Rueckert, D., 2005. Spatio-temporal free-form registration of cardiac MR image sequences. Medical Image Analysis 9, 441 – 456.

Peyrat, J.M., Delingette, H., Sermesant, M., Xu, C., Ayache, N., 2010. Registration of 4D cardiac ct sequences under trajectory constraints with multichannel diffeomorphic demons. IEEE Transactions on Medical Imaging 29, 1351–1368.

Piella, G., De Craene, M., Yao, C., Penney, G., Frangi, A.F., 2011. Multiview diffeomorphic registration for motion and strain estimation from 3D ultrasound sequences, in: Metaxas, D., Axel, L. (Eds.), Functional Imaging and Modeling of the Heart. Springer. volume 6666 of *Lecture Notes in Computer Science*, pp. 6666, 375–383.

Qiu, A., Albert, M., Younes, L., Miller, M.I., 2009. Time sequence diffeomorphic metric mapping and parallel transport track time-dependent shape changes. NeuroImage 45, S51–S60.

Rueckert, D., Aljabar, P., Heckemann, R.A., Hajnal, J.V., Hammers, A., 2006. Diffeomorphic registration using b-splines, in: Larsen, R., Nielsen, M., Sporring, J. (Eds.), Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 702–709.

Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J., 1999. Nonrigid registration using free-form deformations: Application to breast MRI images. IEEE Transactions on Medical Imaging 18, 712–721.

Rutz, A.K., Manka, R., Kozerke, S., Roas, S., Boesiger, P., Schwitter, J., 2009. Left ventricular dyssynchrony in patients with left bundle branch block and patients after myocardial infarction: integration of mechanics and viability by cardiac magnetic resonance. European Heart Journal 30, 2117–2127.

Saddi, K.A., Chefd'hotel, C., Cheriet, F., 2007. Large deformation registration of contrast-enhanced images with volume-preserving constraint, SPIE. p. 651203.

Saito, K., Okura, H., Watanabe, N., Hayashida, A., Obase, K., Imai, K., Maehama, T., Kawamoto, T., Neishi, Y., Yoshida, K., 2009. Comprehensive evaluation of left ventricular strain using speckle tracking echocardiography in normal adults: Comparison of three-dimensional and two-dimensional approaches. Journal of the American Society of Echocardiography 22, 1025–1030.

Silva, E., Sitges, M., Mont, L., Delgado, V., Tamborero, D., Vidal, B., Godoy, M., Poyatos, S., Pare, C., Azqueta, M., Brugada, A, J., 2009. Quantification of left ventricular asynchrony throughout the whole cardiac cycle with a computed algorithm: Application for optimizing resynchronization therapy. Journal of Cardiovascular Electrophysiology 20, 1130–1136.

Song, S., Leahy, R., 1991. Computation of 3D velocity fields from 3D cine ct images of a human heart. IEEE Transactions on Medical Imaging 10, 295–306.

Spaan, J., Breuls, N., Laird, J., 1981. Diastolic-systolic coronary flow differences are caused by intramyocardial pump action in the anesthetized dog. Circulation Research 49, 584.

Sundar, H., Litt, H., Shen, D., 2009. Estimating myocardial motion by 4D image warping. Pattern recognition 42, 2514–2526.

Sung, R., Foster, E., 2011. Assessment of systolic dyssynchrony for cardiac resynchronization therapy is not clinically useful. Circulation 123, 656.

Trouvé, A., Vialard, F., 2010. Shape splines and stochastic shape evolutions: A second order point of view. ArXiv e-prints 1003.3895.

Tsuiki, K., Ritman, E., 1980. Direct evidence that left ventricular myocardium is incompressible throughout systole and diastole. The Tohoku journal of experimental medicine 132, 119.

Vercauteren, T., Pennec, X., Perchant, A., Ayache, N., 2007. Non-parametric diffeomorphic image registration with the demons algorithm, in: Medical Image Computing and Computer-Assisted Intervention, pp. 319–326.

Wang, Y., Georgescu, B., Houle, H., Comaniciu, D., 2010. Volumetric myocardial mechanics from 3d+t ultrasound data with multi-model tracking, in: Camara, O., Pop, M., Rhode, K., Sermesant, M., Smith, N., Young, A. (Eds.), Statistical Atlases and Computational Models of the Heart. Springer Berlin / Heidelberg. volume 6364 of *Lecture Notes in Computer Science*, pp. 184–193.

van Wijk, M.C., Thijssen, J.M., 2002. Performance testing of medical ultrasound equipment: fundamental vs. harmonic mode. Ultrasonics 40, 585 – 591.

Westerhof, N., Boer, C., Lamberts, R.R., Sipkema, P., 2006. Cross-talk between cardiac muscle and coronary vasculature. Physiological Reviews 86, 1263–1308.

Yigitsoy, M., Wachinger, C., Navab, N., 2011. Temporal Groupwise Registration for Motion Modeling, in: Information Processing in Medical Imaging. 648–659.

Yin, F.C., Chan, C.C., Judd, R.M., 1996. Compressibility of perfused passive myocardium. American Journal of Physiology - Heart and Circulatory Physiology 271, H1864–H1870.

Ypenburg, C., van Bommel, R.J., Borleffs, C.J.W., Bleeker, G.B., Boersma, E., Schalij, M.J., Bax, J.J., 2009. Long-Term Prognosis After Cardiac Resynchronization Therapy Is Related to the Extent of Left Ventricular Reverse Remodeling at Midterm Follow-Up. J Am Coll Cardiol 53, 483–490.

Yu, C., Bleeker, G., Fung, J., Schalij, M., Zhang, Q., van der Wall, E., Chan, Y., Kong, S., Bax, J., 2005. Left ventricular reverse remodeling but not clinical improvement predicts long-term survival after cardiac resynchronization therapy. Circulation 112, 1580–1586.

37