

3D Harmonic Phase Tracking with Anatomical Regularization

Y. Zhou^{a,b}, O. Bernard^b, E. Saloux^{c,e}, A. Manrique^{d,e}, P. Allain^a, S. Makram-Ebeid^a, M. De Craene^a

^aPhilips Research Medisys, Suresnes, France,

^bCREATIS, CNRS UMR 5220, INSERM U1044, INSA-Lyon, Université Lyon 1, France,

^cDepartment of Cardiology, CHU Caen, France,

^dDepartment of Human Investigations, GIP Cyceron, Caen, France.

^eEA4650, Normandie Université, Caen, France

Abstract

This paper presents a novel algorithm that extends HARP to handle 3D tagged MRI images. HARP results were regularized by an original regularization framework defined in an anatomical space of coordinates. In the meantime, myocardium incompressibility was integrated in order to correct the radial strain which is reported to be more challenging to recover. Both the tracking and regularization of LV displacements were done on a volumetric mesh to be computationally efficient. Also, a window-weighted regression method was extended to cardiac motion tracking which helps maintain a low complexity even at finer scales. On healthy volunteers, the tracking accuracy was found to be as accurate as the best candidates of a recent benchmark. Strain accuracy was evaluated on synthetic data, showing low bias and strain errors under 5 percents (excluding outliers) for longitudinal and circumferential strains, while the second and third quartiles of the radial strain errors are in the $(-5\%, 5\%)$ range. In clinical data, strain dispersion was shown to correlate with the extent of transmural fibrosis. Also, reduced deformation values were found inside infarcted segments.

Keywords: 3D tagged MR, HARP, Anatomical regularization, Myocardium incompressibility, Monogenic phase

1. Introduction

The non invasive quantification of local myocardial mechanics remains a central challenge for diagnosing heart diseases. Among motion and deformation indexes, strain is a prominent measure. Indeed, strain is less sensitive to tethering from surrounding segments, and is therefore preferred for detecting abnormal segments (Voigt and Flachs-kampf, 2004).

Tagged MR is a classical technique used to quantify regional myocardial deformation. It creates non-invasive markers that move with the myocardium along the cardiac cycle. Tracking these markers (tags) allows precise recovery of myocardial deformation. The introduction of CSPAMM (Fischer et al., 1993) efficiently improved tag fading issues that were traditionally hampering the analysis of tagged MR data in diastole. Breathing artifacts producing unaligned slices have been corrected by the introduction of navigator-driven protocols (Rutz et al., 2008) and the move to fully 3D acquisitions. In (Rutz et al., 2008), a rapid CSPAMM-based 3D protocol was proposed by performing three acquisitions sequentially with line tag preparation in each orthogonal direction. Each of these three acquisitions is performed in a single breath-hold of 18 heartbeats duration and retrospectively corrected for misalignment using a respiratory navigator.

1.1. State of the art in tagged MR quantification

Since its introduction in the late eighties, cardiac tagged MR has triggered the development of a wide spectrum of strain quantification methodologies. Some of them were specifically tailored for tagged MR. Some others are generic methodologies that were readily applied to tagged MR images.

1.1.1. Generic methods

Among generic methods, intensity-based registration by optical flow or Free Form Deformation (FFD) were successfully applied to tagged MR images. Chandrashekara *et al.* (Chandrashekara et al., 2004) applied multilevel FFD to the post-processing of tagged MR. De Craene *et al.* (De Craene et al., 2012) applied temporal diffeomorphic FFD to the analysis of tagged MR, resulting in a temporal regularization of the motion field. Shi *et al.* (Shi et al., 2012) proposed a FFD-based tracking method that combines CINE and tagged MR images to analyze cardiac deformation, with the possibility of adding incompressibility into the framework. Prince *et al.* (Prince and McVeigh, 1992) developed an optical flow method with considerations about compensating the “tag fading” effect. Florack *et al.* (Florack et al., 2007) developed an optical flow method on phase images that automatically selects an optimal scale.

Generic registration methods are robust and have been applied to a large number of cardiac modalities. Most of registration-based methods were applied to a grid tagged

image combining all tagging acquisitions in a single sequence. **However, they do not take advantage of separate directional encodings often performed in tagged acquisitions.** Tagged MR images are acquired in several directions that provide different “channels” for recovering the motion field. This feature can potentially overcome the aperture problem inherent to monochannel optical flow or any parametric registration method. By analogy, this is equivalent to reducing a color image to a black-and-white picture where any segmentation task will not exploit color differences between the objects in the scene.

1.1.2. Tagged MR-specific methods

All methods that exploit tagged MR specificities can be divided into four categories. These categories are rather hermetic, opening up the possibility of combining their respective advantages. Also, most of these methods were only applied to 2D tagged MR and adaptations to 3D are currently in an early stage (Wang et al., 2013)(Rutz et al., 2008), or focus on multiplanar (Pan et al., 2005)(Liu et al., 2012) rather than the volumetric acquisition protocol of Rutz *et al.* (Rutz et al., 2008).

Detection-based methods. A first option retained by Young *et al.* (Young et al., 1995) and Amini *et al.* (Amini et al., 2001) was to first detect the tags planes or intersections and then to track these extracted features. Amini *et al.* (Amini et al., 2001) detected tag planes by B-spline surfaces and tracked sparse myocardial beads. Young *et al.* (Young et al., 1995) used an active contour model to track the tag lines in images and then interpolated a 3D dense motion by a finite-element model. In (Young, 1999), the tagging stripes were detected by a Gaussian-shaped filter and matched by a line searching algorithm. The obtained sparse displacements were further regularized by a finite-element model defined in the local *radial, longitudinal and circumferential* directions. Chen *et al.* (Chen et al., 2010) applied Gabor filters to detect tag line intersections which are further used to initialize a meshless deformable model.

HARP. HARP (Osman et al., 1999) is another tagged MR-specific method that became a reference in this modality. A first step in HARP is to compute phase images. They are computed by applying a band-pass filter in the Fourier domain. Back to the spatial domain, the tracking is then performed on phase (rather than intensity) images. The main rationale for substituting intensity by phase is to improve robustness against tag fading. The overall tracking procedure can be interpreted as an optical flow performed jointly on two “channels”, each channel corresponding to one tagging direction. Because the number of tagging directions equals the number of components of the displacement field, the aperture problem simply disappears. The estimation of HARP displacements was further improved when introducing the CSPAMM protocol, enabling the use of a wider k-space HARP filter (Kuijjer

et al., 2001). However, unlike traditional optical flow algorithms, HARP does not integrate any spatial regularization. As strain is computed from the spatial derivatives of trajectories, one expects that outliers in the reconstructed motion field will yield high errors in the strain field.

The HARP community also proposed several strategies for addressing the limited capture range of tagged images. Indeed, the periodicity of tagged images prevents to recover displacements exceeding one-half of the tag spacing. Osman (Osman et al., 1999) proposed an initialization scheme by sorting the list of points to be tracked. Points with expected smaller motion are tracked first and the tracking result is used as initialization to neighboring points. This concept was further developed in (Liu and Prince, 2010) where the order is optimized through solving for the shortest path on a graph where each edge has a cost related to phase continuity between neighboring voxels in the image. However, the integration of an order to the point-wise tracking is somehow dangerous. Indeed, “tag jump” errors will be propagated to all subsequent points in the list.

SinMod. The SinMod (Arts et al., 2010) algorithm has recently emerged as an alternative to HARP. In the local environment of each pixel, the intensity distribution is modeled as a sinusoid wave orthogonal to the tagging direction. Then both the local frequency and the inter-frame phase shift are computed. The ratio of these two quantities gives an estimate of the displacement when combining the different tagging directions. SinMod was shown to be as fast as HARP and have advantage in accuracy and robustness to image noise. Yet, the ability of tracking large myocardial motion (superior than half tag spacing) remains to be further verified. Also, the absence of any spatial regularization makes strain estimates very sensitive to tracking errors when extracting spatial derivatives.

Gabor-based methods. Qian *et al.* (Qian et al., 2011) proposed to compute strain directly from the image without tracking the myocardial motion. The idea is to filter the intensity images through a Gabor filter bank for retrieving the local orientation and spacing of the tags. This gives a direct access to the spatial derivatives of the displacement field, without the need of computing displacements. This concept was further extended by Kaue *et al.* in (Kaue et al., 2014) where the deformation gradient was directly evaluated from the frequency covector fields. The authors claim that it makes their strain estimate independent to any tracking result. However, they overlook that a tracking is always required for reporting strain evolution at all time points per material point.

1.2. Open question and contributions of this paper

The extension and combination of 2D tagged MR quantification methods to the 3D protocol of Rutz *et al.* (Rutz et al., 2008) remains an open challenge in the tagged MR

community. Several attempts have been made for extending HARP to 3D. They are all based on interpolating 2D HARP results from sparse short- and long-axis images (SLIs). Pan *et al.* (Pan et al., 2005) obtain 3D motion by Gaussian diffusing 2D motions from SLIs on a mesh model. Liang *et al.* (Liang et al., 2007) use a nonuniform rational B-spline volumetric model to densely interpolate the 2D motions from SLIs. Liu *et al.* (Liu et al., 2012) apply the divergence-free vector spline as the interpolation kernel. Recently, SinMod was also extended to volumetric tagged MR in a straightforward manner (Wang et al., 2013), but without embedding any spatial regularization.

This calls for a novel method that integrates the benefits of available strategies in the literature. From registration-based approaches, regularization strategies have the potential to extract robustly motion fields from noisy image data. From tagged MR-specific techniques, extracting phase images and taking full benefit from the multi-channel nature of the input data are crucial aspects to preserve in the design of novel 3D tagged MR quantification tools. In this paper, we present a HARP-based method suitable for fast quantification of 3D *radial, longitudinal and circumferential* strains on volumetric tagged MR images (Rutz et al., 2008). Our main contributions are twofold.

First, we propose to add a spatial regularization step implemented in an anatomical space of coordinates to a straight extension of HARP to 3D. This regularization is performed on a volumetric mesh of the left ventricle (LV) to maintain a low computational cost. The volumetric mesh and anatomical coordinates follow the local directions (*radial, longitudinal and circumferential*). Local windows are defined on the LV myocardial domain, and HARP results are regularized by affine transformations in anatomical coordinates at a regional level. Heyde *et al.* (Heyde et al., 2013) introduced the concept of anatomical regularization in ultrasound but implemented it by resampling the images and doing the tracking in anatomical coordinates. For this application, this would introduce non-linear distortions in the shape of the tags.

Second, we propose to enforce myocardial incompressibility by correcting strain values in the radial direction only. Indeed, this direction is by experience the most challenging to recover meaningful strains. This is likely due to the fact that the number of tags from endocardium to epicardium is small in the radial direction compared to circumferential and longitudinal directions. The divergence-free constraint is imposed at a regional level through the direct application of the divergence theorem.

Besides, we quantitatively compared two phase computation schemes: the original HARP phase and the Monogenic phase (MP) as introduced by (Alessandrini et al., 2012) to the field of tagged MR quantification. As earlier, the impact of using HARP vs. MP was evaluated on synthetic tagged MR images.

Additionally to synthetic images, our method has been evaluated *in vivo* on healthy volunteers and patients. For

healthy volunteers, we compare our accuracy against manually tracked trajectories on a public database (Tobon-Gomez et al., 2013). This evaluation includes the comparison to state-of-the-art tracking algorithms as available from (Tobon-Gomez et al., 2013). On clinical cases, we quantified the ability of our software to discriminate healthy from diseased myocardium by comparing our results between a control group and an ischemic group showing typical subendocardial late-enhancement MR.

2. Data Acquisition & Preprocessing

2.1. 3D Tagged MR acquisition

The 3D tagged MR images were acquired on a Philips 3T Achieva System (Philips Healthcare, Best, The Netherlands) and a 32-channel cardiac coil. A survey scan was used to define the position and axis of the LV. The images were obtained using a 3D CSPAMM sequence (Rutz et al., 2008) based on three volumetric datasets with **sinusoidal** line tagging preparation (tag distance=7 mm) in three orthogonal directions (3DTAG, GyroTools, Zurich, Swiss). **The reader is referred to (Rutz et al., 2008) for the details of the tag preparation pulse.** For data acquisition, a modified hybrid multishot, segmented EPI sequence was employed (TR/TE= 7.0/3.2 ms), with ramped flip angles (19-25°) to prevent tag fading during the cardiac cycle. Images were acquired with a reduced field-of-view (108x108x108 mm) enclosing the LV. **The acquisition matrix size was 28x14x14. The voxel size for each of the three datasets was 0.96x0.96x7.71mm. The number of reconstructed frames was 24, and the temporal resolution ranged from 30ms to 45ms.** An echo-navigator technique was used to compensate for respiratory motion. As the three orthogonal sequences are aligned, the whole acquisition can be interpreted as a multi-channel sequence of 3D volumes where at each voxel, three intensities are available (one for each direction).

2.2. Segmentation of LV

LV segmentation was done manually. The LV was segmented at end-diastole only (last frame). Segmenting the first frame is not feasible because tissue and blood are both tagged and cannot be distinguished. For segmentation purposes, an untagged image was produced by taking modulo for the complex image obtained from HARP bandpass filtering (described later in Sect. 3.3.3) on each channel. Untagged images from the 3 channels were then averaged to produce a single image with isotropic voxel resolution. A template mesh was first positioned on the untagged image by clicking landmarks: 1 landmark was placed in the apex, 4 in the basal planes and 1 at mid-level in the septum. The template mesh was rotated and scaled to match the set of input landmarks. The mesh was then converted to an implicit function and adjusted by adding pointwise penalties to the algorithm described in (Mory et al., 2012).

2.3. Mesh model

The output of the segmentation is a closed surface mesh that encompasses LV endocardium and epicardium. The purpose of the resampling procedure (described in Section 2.3.1) is twofold. First, the surface mesh is converted to a volumetric *dense* mesh that also contains an inner layer between endocardium and epicardium. Second, the LV domain is sampled regularly in *radial*, *longitudinal* and *circumferential* directions. This facilitates regularizing the displacement field on the mesh topology.

2.3.1. Resampling rules and directional indexes

The resampling is done by casting rays from the long axis to the endocardial and epicardial surfaces as described in Fig. 1. The mid part of the LV is discretized using a cylindrical sampling. The long axis is sampled uniformly. From each long axis point, N_c rays are then **cast** with equal angular binning in the plane orthogonal to the long axis. For each ray, the intersection with the surface mesh gives two points: one at endocardium and one at epicardium. The average of these two points produces the inner layer. Similarly, apical and basal parts are sampled in a spherical manner. Azimuth and elevation angles are binned uniformly (see Fig. 1(a) and 1(b)).

Following this discretization scheme, each point in the generated mesh can be associated to three indexes. First, each point is either endocardial, mid or epicardial. Second, each point is intercepted by a single ray in the *circumferential* direction. Finally, each ray is either casted from one source point on the long axis, being indexed from the sequential cylindrical discretization or angled with respect to the long axis, being indexed from the sequential spherical discretization (Fig. 1(a)). In this way, every node in the mesh is related to three indexes (m_r, m_l, m_c) where $m_d \in [1, N_d]$ and N_d ($d=r,l,c$) is the number of divisions in each of the three directions (*radial*, *longitudinal* or *circumferential*). Then the spatial coordinates of every node can be written as $\mathcal{P}(m_r, m_l, m_c)$.

In our experiments, $N_r = 3$, $N_l = 35$ and $N_c = 50$. As the number of tags is low in the radial direction, we chose to use only three radial layers. N_l and N_c were chosen to produce isotropic cells in the volumetric mesh, which implies similar cell resolutions in longitudinal and circumferential directions.

By convention, radial indexes are ordered from endocardium to epicardium, longitudinal indexes are ordered from apex to base, and circumferential indexes go counter-clockwise when looking from the apex.

2.3.2. Local directions

Based on the mapping \mathcal{P} (Sect. 2.3.1), local *radial*, *longitudinal* and *circumferential* directions $\hat{\mathbf{e}}_r(\mathbf{x})$, $\hat{\mathbf{e}}_l(\mathbf{x})$ and $\hat{\mathbf{e}}_c(\mathbf{x})$ can be computed by normalizing the edges of the volumetric mesh. However, this set of directions would not be orthogonal. As it will later be seen that the orthogonality assumption must be verified at all locations except

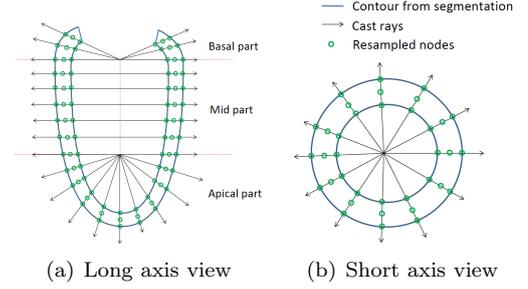


Figure 1: Resampling of LV surface mesh to volumetric dense mesh

the purely apical ones during the regularization procedure (Sect. 3.3.5), local directions were defined as follows.

For a non-apical mesh point \mathbf{x} whose directional indexes are (m_r, m_l, m_c) , its longitudinal direction can be computed as:

$$\begin{aligned} \mathbf{e}_l(\mathbf{x}) &= \mathcal{P}(m_r, m_l, m_c) - \mathcal{P}(m_r, m_l - 1, m_c) \quad \text{if } m_l > 1 \\ \hat{\mathbf{e}}_l(x) &= \frac{\mathbf{e}_l(x)}{\|\mathbf{e}_l(x)\|} \end{aligned} \quad (1)$$

Its circumferential direction $\hat{\mathbf{e}}_c(\mathbf{x})$ was then computed similarly but further corrected to be orthogonal to $\hat{\mathbf{e}}_l(\mathbf{x})$:

$$\begin{aligned} \mathbf{e}_c(\mathbf{x}) &= \mathcal{P}(m_r, m_l, \text{mod}(m_c + 1, N_c)) - \mathcal{P}(m_r, m_l, m_c) \\ &\quad \text{if } m_l > 1 \\ \hat{\mathbf{e}}_c(\mathbf{x}) &= \frac{\mathbf{e}_c(\mathbf{x}) - (\mathbf{e}_c(\mathbf{x}) \cdot \hat{\mathbf{e}}_l(\mathbf{x}))\hat{\mathbf{e}}_l(\mathbf{x})}{\|\mathbf{e}_c(\mathbf{x}) - (\mathbf{e}_c(\mathbf{x}) \cdot \hat{\mathbf{e}}_l(\mathbf{x}))\hat{\mathbf{e}}_l(\mathbf{x})\|} \end{aligned} \quad (2)$$

Finally, the radial direction $\hat{\mathbf{e}}_r(\mathbf{x})$ was obtained from the cross product between $\hat{\mathbf{e}}_c(\mathbf{x})$ and $\hat{\mathbf{e}}_l(\mathbf{x})$.

$$\hat{\mathbf{e}}_r(\mathbf{x}) = \hat{\mathbf{e}}_c(\mathbf{x}) \times \hat{\mathbf{e}}_l(\mathbf{x}) \quad \text{if } m_l > 1 \quad (3)$$

For purely apical points ($m_l = 1$), $\hat{\mathbf{e}}_c(\mathbf{x})$ has no definition due to the singularity on the apex. $\hat{\mathbf{e}}_r(\mathbf{x})$ was defined in a similar manner by normalizing the edges in the radial direction. $\hat{\mathbf{e}}_l(\mathbf{x})$ was computed by taking the mean of its N_c nearest neighbors located in the same radial layer and normalizing the result.

2.3.3. Anatomical coordinates

Our goal is to map Cartesian (\mathbf{x}) to anatomical (r, l, c) coordinates while satisfying the following properties:

$$\begin{aligned} \forall \mathbf{x} \in \Omega, \quad \nabla r(\mathbf{x}) &= \hat{\mathbf{e}}_r(\mathbf{x}), \\ \nabla l(\mathbf{x}) &= \hat{\mathbf{e}}_l(\mathbf{x}), \\ \nabla c(\mathbf{x}) &= \frac{1}{\rho(\mathbf{x})}\hat{\mathbf{e}}_c(\mathbf{x}) \quad . \end{aligned} \quad (4)$$

where Ω is the LV myocardium domain, and $\rho(\mathbf{x})$ is the distance of point \mathbf{x} to the LV long axis.

For each direction $d \in \{r, l, c\}$, we start from a group of reference points $\{X_0^d\}$ where the anatomical coordinates

are set to be zero. For longitudinal coordinates, the apex is taken as a reference point. In the circumferential direction, points located at the center of the local windows (described later in Sect. 3.3.2) are chosen as reference points. Finally, for the radial direction, $\{X_0^d\}$ is defined as the set of mid-layer points. From $\{X_0^d\}$, the coordinates are obtained by curvilinear integration. Taking one direction $d \in \{r, l, c\}$, $d(\mathbf{x})$ is defined as

$$d(\mathbf{x}) = \begin{cases} \min_{\mathcal{C}_d(\mathbf{x})} \int_{\mathcal{C}_d(\mathbf{x})} \hat{\mathbf{e}}_d(\mathbf{s}) \cdot d\mathbf{s}, & \text{if } d = r, l \\ \min_{\mathcal{C}_d(\mathbf{x})} \int_{\mathcal{C}_d(\mathbf{x})} \frac{1}{\rho(\mathbf{s})} \hat{\mathbf{e}}_d(\mathbf{s}) \cdot d\mathbf{s}, & \text{if } d = c \end{cases} \quad (5)$$

where $\mathcal{C}_d(\mathbf{x})$ is a geodesic path on the mesh starting at a node $\in X_0^d$.

The $r(\mathbf{x})$ and $l(\mathbf{x})$ coordinates are computed globally for the whole LV. However, practically, the $c^{(i)}(\mathbf{x})$ coordinate is defined locally for each window i (described later in Sect. 3.3.2) in order to ensure its continuity over the window domain. Examples of r, l over the whole myocardial domain and $c^{(i)}$ over the local window i are plotted in Fig. 2.

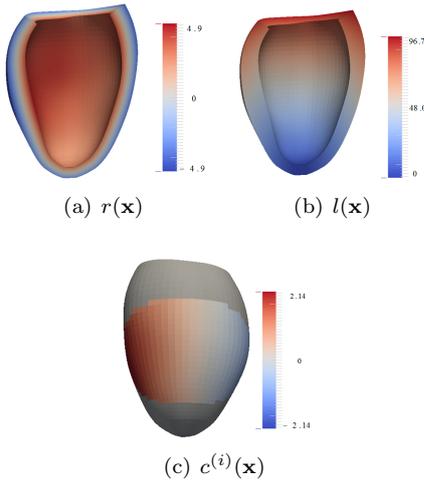


Figure 2: Anatomical coordinates $\{r, l, c^{(i)}\}$ defined on the LV volumetric mesh (in (c), $c^{(i)}(x)$ has no definition in the surrounding grey region since it is beyond the range of the current window i)

3. Tracking methodology

After the pre-processing step, the LV is tracked frame-by-frame backwardly. As a first step, the HARP tracking method is used to generate initial estimates of the motion field. HARP-based displacements are further regularized following a parametric motion model defined in an anatomical space of coordinates. Myocardial incompressibility is integrated as a *soft* constraint that helps retrieve robust radial strain estimations. For achieving fast convergence of the optimization, a window-weighted regression

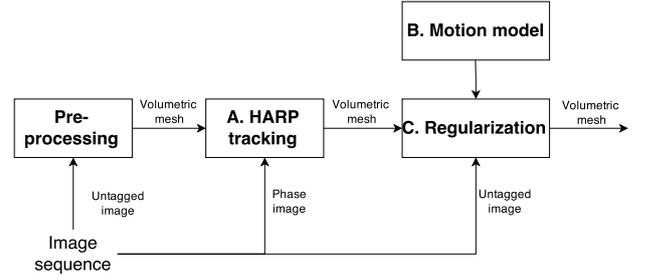


Figure 3: Pipeline of the proposed tracking method

technique is used to decompose the complex global problem into simpler local ones. Also, a multiscale framework was developed where strong regularization is imposed at first and gradually decreased.

3.1. HARP tracking

HARP (Osman et al., 1999) estimates the displacements by first computing two phase image sequences and then performing a non-regularized optical flow on these two *channels*.

3.1.1. Phase computation

HARP (Osman et al., 1999) pre-processes the original intensities and computes phase images. This pre-processing is meant to compensate for tag fading artifacts inherent to tagged MR images.

As a preliminary, it should be noted that the acquired 3D sequences do not have isotropic spacing (inter-slice thickness is about seven times the in-plane pixel spacing). Therefore, although the combination of the three channels can be seen as 3D (since the three stacks are acquired in orthogonal directions), we opted for computing phase images in 2D slice by slice.

HARP phase is computed from the result of a band-pass filter in the frequency domain. As in our case, the tagging directions coincide with the m and n axes of each 2D slice, the 2D ellipsoid-shape bandpass filter $f_k(\boldsymbol{\omega})$ used in (Osman et al., 2000) is therefore defined as

$$f_k(\boldsymbol{\omega}) = \begin{cases} 1, & \text{if } \delta_k(\boldsymbol{\omega}) \leq 1 \\ e^{-(\delta_k(\boldsymbol{\omega})-1)^2/(2\sigma_h^2)}, & \text{otherwise} \end{cases} \quad (6)$$

with

$$\delta_k(\boldsymbol{\omega}) = \begin{cases} \sqrt{\frac{(\omega_m - \omega_0)^2}{(0.5\omega_0)^2} + \frac{\omega_n^2}{\omega_0^2}}, & \text{if } m \text{ tagging direction} \\ \sqrt{\frac{(\omega_m)^2}{\omega_0^2} + \frac{(\omega_n - \omega_0)^2}{(0.5\omega_0)^2}}, & \text{if } n \text{ tagging direction} \end{cases} \quad (7)$$

where $\boldsymbol{\omega} = [\omega_m \ \omega_n]^T$, and ω_0 is the spatial tagging frequency. In essence, $f_k(\boldsymbol{\omega})$ is an elliptic-shape bandpass filter centered at the spatial tagging frequency, with a smooth transition around the ellipse border as shown in Fig. 4. In our experiments, we set σ_h to 0.05 as suggested in (Osman et al., 2000). It can be seen from Eq. 7 that

the filter keeps positive frequencies only, around the first harmonic peak in the frequency domain (Osman et al., 1999).

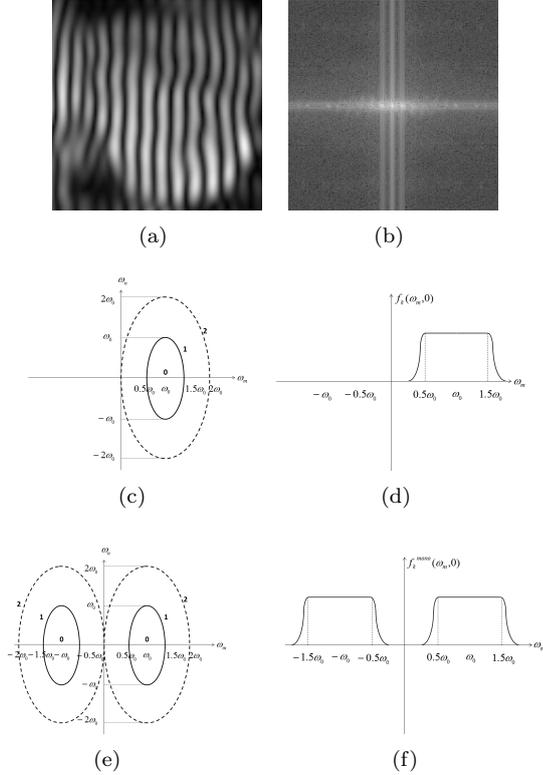


Figure 4: Bandpass filter used in HARP and MP with vertical line tagging (tagging direction in m -axis). ((a): $I_k(m, n)$; (b): $|\mathcal{F}\{I_k(m, n)\}|$; (c): isovalues of $\delta_k(\omega)$ in HARP; (d): $f_k(\omega)$ along ω_m in HARP; (e): isovalues of $\delta_k^{mono}(\omega)$ in MP; (f): $f_k^{mono}(\omega)$ along ω_m in MP)

Back to the spatial domain, the phase $a_k(m, n)$ is then computed as

$$a_k(m, n) = \arg(\mathcal{F}^{-1}\{\mathcal{F}\{I_k(m, n)\}f_k(\omega)\}) \quad (8)$$

Since the 2D slices are sparsely distributed, the in-slice phase needs to be interpolated to 3D. In our experiments, the complex image $\mathcal{F}^{-1}\{\mathcal{F}\{I_k(m, n)\}f_k(\omega)\}$ was linearly interpolated before extracting the phase. This avoids artifacts related to phase wrapping discontinuities.

In this paper, we compare the phase definition of HARP with the alternative brought by the monogenic signal theory (Alessandrini et al., 2013) with a proposed bandpass filter. We opted for a symmetric ellipsoid-shape bandpass filter similar to the one used in HARP except that it consists of two ellipses centered at $\pm\omega_0$ and therefore is symmetric. This profile is obtained by modifying the $\delta_k(\omega)$ distance function introduced in Eq. 6 by:

$$\delta_k^{mono}(\omega) = \begin{cases} \sqrt{\frac{(|\omega_m| - \omega_0)^2}{(0.5\omega_0)^2} + \frac{\omega_0^2}{\omega_0^2}} & , \text{ if } m \text{ tagging direction} \\ \sqrt{\frac{(\omega_m)^2}{\omega_0^2} + \frac{(|\omega_n| - \omega_0)^2}{(0.5\omega_0)^2}} & , \text{ if } n \text{ tagging direction} \end{cases} \quad (9)$$

Replacing $\delta_k(\omega)$ by $\delta_k^{mono}(\omega)$ in Eq. 6 produces the bandpass filter as shown in Fig. 4(e) and 4(f). Both phase definitions are compared experimentally in Sect. 5.1.

3.1.2. Dense motion estimation

HARP tracking can be seen as a multi-channel optical flow on the true phases (phases without wrapping artifacts) (Osman et al., 1999). Here we denote the true phase by $\phi_k(\mathbf{x})$ with $k = 1, 2, 3$ in the 3D case where k indexes the channel number (tagging direction). The tracking of a material point from $t \rightarrow t-1$ can be solved iteratively by updating the displacement $\tilde{\mathbf{u}}_t$

$$\tilde{\mathbf{u}}_t \leftarrow \tilde{\mathbf{u}}_t + \mathbf{v}_t \quad (10)$$

where \mathbf{v}_t is computed by solving the following 3x3 linear system:

$$A \cdot \mathbf{v}_t = b \quad (11)$$

$$A = \begin{bmatrix} \nabla\phi_1^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t)^T \\ \nabla\phi_2^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t)^T \\ \nabla\phi_3^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t)^T \end{bmatrix} \quad (12)$$

$$b = \begin{bmatrix} \phi_1^t(\mathbf{x}) - \phi_1^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t) \\ \phi_2^t(\mathbf{x}) - \phi_2^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t) \\ \phi_3^t(\mathbf{x}) - \phi_3^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t) \end{bmatrix} \quad (13)$$

Spatial derivatives were corrected when the phase gets wrapped from $-\pi$ to π (Osman et al., 1999). Similarly, b was computed from the wrapped phases a_k as in (Osman et al., 1999).

3.2. Motion model

Let $\mathbf{u}_t^{(i)}(\mathbf{x})$ be a local parametric motion within a window i (described later in Sect. 3.3.2) from $t \rightarrow t-1$. As tracking aims at recovering both motion and deformation values, we opted for an affine transformation in the anatomical coordinates $\{r, l, c^{(i)}\}$ that combines translations, stretching (or shortening) and shearing. This 12-parameter model was defined in the anatomical system of coordinates introduced in Sect. 2.3.2 and 2.3.3:

$$\begin{aligned} \mathbf{u}_t^{(i)}(\mathbf{x}) &= L_r^{(i)}(\mathbf{x})\hat{\mathbf{e}}_r(\mathbf{x}) + L_l^{(i)}(\mathbf{x})\hat{\mathbf{e}}_l(\mathbf{x}) + L_c^{(i)}(\mathbf{x})\hat{\mathbf{e}}_c(\mathbf{x}) \\ \text{with } L_r^{(i)}(\mathbf{x}) &= a_{rr}r(\mathbf{x}) + a_{rl}l(\mathbf{x}) + a_{rc}c^{(i)}(\mathbf{x}) + b_r \\ L_l^{(i)}(\mathbf{x}) &= a_{lr}r(\mathbf{x}) + a_{ll}l(\mathbf{x}) + a_{lc}c^{(i)}(\mathbf{x}) + b_l \\ L_c^{(i)}(\mathbf{x}) &= (a_{cr}r(\mathbf{x}) + a_{cl}l(\mathbf{x}) + a_{cc}c^{(i)}(\mathbf{x}) + b_c)\rho(\mathbf{x}) \end{aligned} \quad (14)$$

The set of zero-order $\{b_r, b_l, b_c\}$ (related to motion), first-order $\{a_{rr}, a_{ll}, a_{cc}\}$ (related to *radial, longitudinal and circumferential* strains) and first-order $\{a_{rl}, a_{rc}, a_{lr}, a_{lc}, a_{cr}, a_{cl}\}$ (related to shearing) parameters fully define the motion in the local window i .

The tracking accuracies between using the above motion model and a simplified one without the shearing elements $\{a_{rl}, a_{rc}, a_{lr}, a_{lc}, a_{cr}, a_{cl}\}$ will be further discussed in Sect. 5.1.2.

3.3. Regularization by window-weighted regression

McLeod *et al.* (McLeod et al., 2013) integrated local window functions defined in prolate spheroidal coordinates within the polyaffine registration framework. In this section, we will explain how we define the window functions similarly and further couple it with the motion model (Sect. 3.2) for regularizing HARP results.

3.3.1. The Partition of unity method

As introduced in (Makram-Ebeid and Somphone, 2007), a global regression problem can be simplified by solving a set of local ones with the help of the *Partition of unity* technique.

A partition of the LV domain is defined by a set of control points $\{\mathbf{x}_i\}$ and by window functions $\varphi^{(i)}(\mathbf{x})$ associated to each control point. The latter are assumed to satisfy the following conditions:

$$\begin{aligned} \arg \max_{\mathbf{x}} \varphi^{(i)}(\mathbf{x}) &= \mathbf{x}_i, \quad \forall i \\ \varphi^{(i)}(\mathbf{x}) &\geq 0, \quad \forall \mathbf{x} \in \Omega, \\ \sum_i \varphi^{(i)}(\mathbf{x}) &= 1 \quad \forall \mathbf{x} \in \Omega. \end{aligned} \quad (15)$$

where Ω is the LV myocardium. The computation of $\varphi^{(i)}(\mathbf{x})$ is described later in Sect. 3.3.2. The regularized displacement $\mathbf{u}_t(\mathbf{x})$ is computed as

$$\mathbf{u}_t(\mathbf{x}) = \sum_i \varphi^{(i)}(\mathbf{x}) \mathbf{u}_t^{(i)}(\mathbf{x}) \quad (16)$$

where $\mathbf{u}_t^{(i)}(\mathbf{x})$ is the motion model introduced in Sect. 3.2.

Under Eq. 15 and 16, if given certain measurement motion estimates $\tilde{\mathbf{u}}_t(\mathbf{x})$ (HARP tracking results in Sect. 3.1.2) and confidence weight $\tilde{w}_t(\mathbf{x})$ (described later in Sect. 3.3.3), the sum of the local errors $E^{(i)}$ then gives an upper bound to the global error E .

$$\begin{aligned} E &= \int_{\Omega} \tilde{w}_t(\mathbf{x}) \left\| \tilde{\mathbf{u}}_t(\mathbf{x}) - \sum_i \varphi^{(i)}(\mathbf{x}) \mathbf{u}_t^{(i)}(\mathbf{x}) \right\|^2 d\mathbf{x} \\ &\leq \sum_i \int_{\Omega} \tilde{w}_t(\mathbf{x}) \varphi^{(i)}(\mathbf{x}) \left\| \tilde{\mathbf{u}}_t(\mathbf{x}) - \mathbf{u}_t^{(i)}(\mathbf{x}) \right\|^2 d\mathbf{x} = \sum_i E^{(i)} \end{aligned} \quad (17)$$

The reader is referred to (Makram-Ebeid and Somphone, 2007) for further details. Bounding all $E^{(i)}$ to a small enough quantity also minimizes the global error, thus making it possible to decompose a global optimization problem into regional ones that are easier to solve.

3.3.2. Window functions

The LV domain was first divided into small windows as illustrated in Fig. 5. The purpose of this section is to define how $\varphi^{(i)}$ functions are defined for a given partition of the LV domain. This partition is obtained using a multi-resolution strategy as explained in Section 3.5.

Given a partition of the LV, control points were defined as the center of these windows. Window functions are centered at each control point and must decrease when the distance to the control point increases. Since we opted for tracking a volumetric mesh and defined a system of coordinates on the mesh, the geodesic distance was chosen rather than the Cartesian one for computing the window function. Considering a standard Gaussian kernel involving the geodesic distance ξ between \mathbf{x} and the center of the window \mathbf{x}_i

$$g^{(i)}(\mathbf{x}) = \frac{1}{2\pi\sigma_i} e^{-\frac{\xi(\mathbf{x}, \mathbf{x}_i)^2}{2\sigma_i^2}} \quad (18)$$

To use a single σ parameter for various window sizes, we expressed σ_i as

$$\sigma_i = \sigma \chi_i \quad (19)$$

where χ_i is computed for each window by averaging its circumferential and longitudinal extents. In Eq. 19, σ is an auxiliary parameter named as *kernel bandpass*. It controls how neighboring windows will overlap for producing the final displacement estimate in Eq. 16. The influence of σ on the final tracking accuracy is further discussed in Sect. 5.1.1. Note that the σ parameter is normalized and can thus be used across several window resolutions (see Sect. 3.5).

Window functions are defined by normalizing $g^{(i)}$ values for summing to 1:

$$\varphi^{(i)}(\mathbf{x}) = \frac{g^{(i)}(\mathbf{x})}{\sum_j g^{(j)}(\mathbf{x})} \quad (20)$$

As seen in Sect. 2.3.2, $\hat{\mathbf{e}}_c(\mathbf{x})$ is not defined for purely apical points. Therefore these points should be excluded from the regularization process. This is done by setting all $\varphi^{(i)}(\mathbf{x})$ to 0 for those points.

3.3.3. Confidence weight

It was shown in Eq. 8 that the HARP phase is computed by taking the argument of the complex signal. In fact, computing its **amplitude** $|\mathcal{F}^{-1}\{\mathcal{F}\{I_k(m, n)\}f_k(\omega)\}|$ (note as \mathcal{M}_k) corresponds to untagging the original image (Osman et al., 1999). This information can be used as a confidence weight map $\tilde{w}_t(\mathbf{x})$: a high value of \mathcal{M}_k indicates a higher probability of belonging to the myocardium, hence an increased confidence in the HARP result $\tilde{\mathbf{u}}_t(\mathbf{x})$. In our implementation, the confidence weights $\tilde{w}_t(\mathbf{x})$ were obtained by averaging the three channels and taking the geometric mean over the two frames t and $t-1$:

$$\tilde{w}_t(\mathbf{x}) = \sqrt{\frac{1}{3} \left(\sum_{k=1}^3 \mathcal{M}_k^t(\mathbf{x}) \right) \cdot \frac{1}{3} \left(\sum_{k=1}^3 \mathcal{M}_k^{t-1}(\mathbf{x} + \tilde{\mathbf{u}}_t(\mathbf{x})) \right)}, \quad (21)$$

3.3.4. Incompressibility constraint

Imposing a zero divergence is a classical way of enforcing incompressibility. This constraint can be imposed

either as a soft (global) or hard (local) level. In our case, computing the divergence locally would require complex spatial derivations of the $r(\mathbf{x}), l(\mathbf{x})$ and $c^{(i)}(\mathbf{x})$ functions. Therefore, we preferred to impose incompressibility within each window, enforcing the overall volume change to be zero. For doing so, we applied the divergence theorem. If $\partial S^{(i)}$ stands for the surface enclosing the window i , the outward flow of $\mathbf{u}_t^{(i)}(\mathbf{x})$ through $\partial S^{(i)}$ equals the volume integral of the divergence over that window. Therefore imposing myocardial incompressibility is equivalent to forcing the flow through $\partial S^{(i)}$ to be zero. For each local window i defined on the volumetric mesh, $\partial S^{(i)}$ consists of quadrilaterals. Each of the quadrilateral j can be represented by its center \mathbf{x}_j , its normal (pointing outward) \vec{n}_j and its area A_j . The flow of $\mathbf{u}_t^{(i)}(\mathbf{x})$ through $\partial S^{(i)}$ can then be written as

$$\begin{aligned} Q_t^{(i)} &= \sum_{j \in \partial S^{(i)}} \mathbf{u}_t^{(i)}(\mathbf{x}_j) \cdot A_j \vec{n}_j \\ &= \sum_{j \in \partial S^{(i)}} \sum_{d=r,l,c} L_d^{(i)}(\mathbf{x}_j) (\hat{\mathbf{e}}_d(\mathbf{x}_j) \cdot A_j \vec{n}_j) \end{aligned} \quad (22)$$

Because $L_d^{(i)}(\mathbf{x}_j)$ is linear in $\{b_d, a_{rd}, a_{ld}, a_{cd}\}$ ($d = r, l, c$), so is $Q_t^{(i)}$. Imposing $Q_t^{(i)} = 0$ amounts to add a linear equality constraint. In other words, a_{rr} can be written as a linear combination of all other parameters $\{b_r, b_l, b_c, a_{rl}, a_{rc}, a_{lr}, a_{ll}, a_{lc}, a_{cr}, a_{cl}, a_{cc}\}$.

3.3.5. Solving local system and reforming global motion

Let us show how, with the motion model defined in Sect. 3.2, we can obtain a quadratic form for the local $E^{(i)}$ errors. As $\{\hat{\mathbf{e}}_r(\mathbf{x}), \hat{\mathbf{e}}_l(\mathbf{x}), \hat{\mathbf{e}}_c(\mathbf{x})\}$ define an orthogonal basis, we can project HARP displacements $\tilde{\mathbf{u}}_t(\mathbf{x})$ on this basis:

$$\begin{aligned} \tilde{\mathbf{u}}_t(\mathbf{x}) &= \sum_{d=r,l,c} p_d(\mathbf{x}) \hat{\mathbf{e}}_d(\mathbf{x}) \\ \text{with } p_d(\mathbf{x}) &= \tilde{\mathbf{u}}_t(\mathbf{x}) \cdot \hat{\mathbf{e}}_d(\mathbf{x}) \end{aligned} \quad (23)$$

From Eq. 14 and 23, $E^{(i)}$ in Eq. 17 can be expanded as:

$$E^{(i)} = \int_{\Omega} \varphi^{(i)}(\mathbf{x}) \tilde{w}_t(\mathbf{x}) \left\| \sum_{d=r,l,c} \left(L_d^{(i)}(\mathbf{x}) - p_d(\mathbf{x}) \right) \hat{\mathbf{e}}_d(\mathbf{x}) \right\|^2 d\mathbf{x} \quad (24)$$

Using orthogonality between $\{\hat{\mathbf{e}}_r(\mathbf{x}), \hat{\mathbf{e}}_l(\mathbf{x}), \hat{\mathbf{e}}_c(\mathbf{x})\}$, it is easy to further obtain:

$$E^{(i)} = \int_{\Omega} \varphi^{(i)}(\mathbf{x}) \tilde{w}_t(\mathbf{x}) \sum_{d=r,l,c} \left(L_d^{(i)}(\mathbf{x}) - p_d(\mathbf{x}) \right)^2 d\mathbf{x} \quad (25)$$

Because $L_d^{(i)}(\mathbf{x})$ is linear in b_d, a_{dr}, a_{dl} and a_{dc} ($d = r, l, c$), $E^{(i)}$ is a quadratic form and can be minimized by solving a linear system. Since the radial strain is reported to be the most challenging to recover, we integrated the incompressibility constraint and in Eq. 25 we substituted a_{rr}

by the linear combination of $\{b_r, b_l, b_c, a_{rl}, a_{rc}, a_{lr}, a_{ll}, a_{lc}, a_{cr}, a_{cl}, a_{cc}\}$. The minimization of $E^{(i)}$ then gives a 11x11 linear system which is solved by pseudo-inverse. As a final step, a_{rr} was computed from the linear combination mentioned above. Once the local motions $\mathbf{u}_t^{(i)}(\mathbf{x})$ are computed, the global motion is reformed by Eq. 16.

3.4. Computing apex motion a posteriori

Since purely apical points do not take part in the regularization process (see Sect. 3.3.2), we need to compute their motions by other ways. For each of the N_r purely apical points, its displacement was computed by averaging the final motions $\mathbf{u}_t(\mathbf{x})$ of its N_c nearest neighbors located in the same radial layer.

3.5. Multi-resolution framework

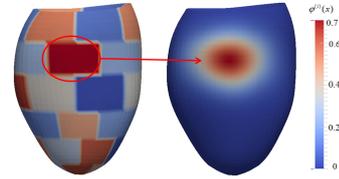


Figure 5: Definition of LV local windows and an example of the window function for the red-circled window

We need to define a partition of the LV domain at increasing levels of resolutions. For simplification purposes, we partitioned the LV using a single parameter: the number of divisions in the circumferential and longitudinal directions.

Additionally, to obtain a more homogeneous mix between windows, they were shifted by half their circumferential extent over contiguous longitudinal layers (see Fig. 5 for an illustration).

For improving the convergence of the motion and deformation coefficients, we started with 3 divisions in longitudinal and circumferential directions. This number was then doubled over the two next resolutions.

At each resolution, the HARP tracking (Sect. 3.1.2) was initialized with the result of the previous resolution before proceeding to the regularization described earlier in this Section.

The influence of the number of resolutions n on the final tracking accuracy is further discussed in Sect. 5.1.1.

4. Generation of synthetic images

For evaluating the impact of several parameters on tracking and strain accuracy, a simple strategy was implemented for generating synthetic images. The pipeline is shown in Fig. 6 and detailed in this Section. The required inputs are 1) one sequence of 3D+t tetrahedral LV meshes (noted as \mathcal{S}^t where t is the frame index) as made publicly available from (De Craene et al., 2013); and 2)

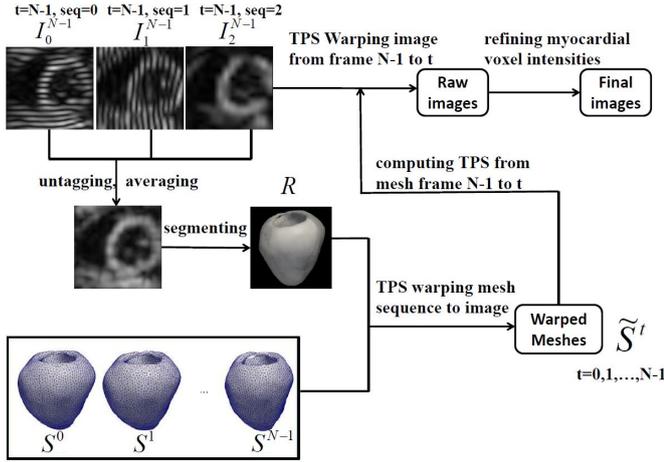


Figure 6: Pipeline for generating synthetic tagged MR sequence

three CSPAMM images with orthogonal tagging orientations (noted as \mathcal{I}_k^{N-1} where N is the number of frames and $k \in \{1, 2, 3\}$), as made publicly available from (Tobon-Gomez et al., 2013). The LV meshes used here come from an electro-mechanical model (Marchesseau et al., 2013) which takes into account fiber distribution. Fibers have a varying orientation from endocardium to epicardium that follows values published in the literature. The quasi incompressibility of the myocardial motion is ensured. However, on the simulation generated in (De Craene et al., 2013), the apex was fixed using a hard constraint on displacements. This limits torsion on this simulation.

4.1. Registering LV meshes to real tagged MR images

We assume that the last frame in the image sequence and the last electro-mechanical mesh both correspond to the end of diastole. As described in Sect. 2.2, the untagged images were computed for the last frame and averaged to an isotropic volume. This volume was segmented by first extracting the surface of the last mesh of the electro-mechanical model as the template and manually position/deform it through the same segmentation technique as in Sect. 2.2. The segmentation result will be noted as \mathcal{R} . A Thin plate spline (TPS) (Bookstein, 1989) transformation \mathcal{T} was then computed between the template and \mathcal{R} . This transformation was then applied to all meshes \mathcal{S}^t , creating a new sequence of tetrahedral LV meshes $\tilde{\mathcal{S}}^t = \mathcal{T}(\mathcal{S}^t)$.

4.2. Warping real tagged MR images by forces computed from meshes

Synthetic images were obtained by deforming the last frame images \mathcal{I}_k^{N-1} to match the sequence of $\tilde{\mathcal{S}}^t$ meshes. This was done in two steps. First, a TPS transformation $\mathcal{T} : \tilde{\mathcal{S}}^{N-1} \mapsto \tilde{\mathcal{S}}^t$ was computed for extending displacements defined at the mesh nodes only to the whole volume and warping the \mathcal{I}_k^{N-1} image. Second, this image was corrected inside the myocardial for removing TPS-induced

artifacts as follows. For every voxel in \mathcal{I}_k^t , we checked if it belonged to the myocardial domain $\tilde{\mathcal{S}}^t$. If it did, we found the corresponding tetrahedral cell and the barycentric coordinates of the voxel within that cell. We then interpolated the intensity of \mathcal{I}_k^{N-1} for the same cell and barycentric coordinates. This intensity was used instead of the one obtained through TPS warping.

5. Results

5.1. Synthetic images

A synthetic tagged MR image sequence was generated for a normal geometry using the approach described in Sect. 4. Fig. 7 shows the ground truth meshes and the synthetic images. The grid pattern was obtained by multiplying the three input channels (with orthogonal line tagging pattern) in a single image after resampling. This resampled image is used for visualization only and not for processing.

The tetrahedral simulation mesh at the last frame was resampled to a hexahedral mesh as described in Sect. 2.3. Over this procedure, special attention was paid to have all nodes of the resampled mesh within the elements of the original mesh. Using barycentric coordinates of the simulation meshes, node displacements can be interpolated at all times on the resampled mesh, hence providing ground truth displacement values. After propagating the last mesh through tracking, measured displacements can be compared to ground truth for quantifying tracking accuracy.

Ground truth strain values were computed on the set of resampled meshes. Our mesh model makes the computation of 1D engineering strain trivial since all edges are following *radial*, *longitudinal* and *circumferential* directions. Relative length changes were computed for each hexahedral cell according to

$$\epsilon_d = \frac{1}{4} \sum_{j=1}^4 \frac{l_{dj}^t}{l_{dj}^0} - 1 \quad (26)$$

where l_{dj}^t is the length of edge j in the d direction (being either *radial*, *longitudinal* or *circumferential*) at time t . Strain values were then averaged per AHA segment for all cells belonging to that segment. Similarly to tracking accuracy, strain accuracy was measured by doing the same computation on tracked meshes and comparing ground truth with measured values.

Tracking and strain accuracies on the synthetic sequence were used as a reference for tuning key parameters, comparing the two phase computation techniques (HARP and MP) and discussing the influence of shearing elements. Fig. 8 and 9 show the evolution of motion and circumferential strain errors when spanning different n (Sect. 3.5) and σ (Sect. 3.3.2) values. Results for longitudinal and radial strains are provided as supplementary material at <http://bit.ly/1xsreCa>.

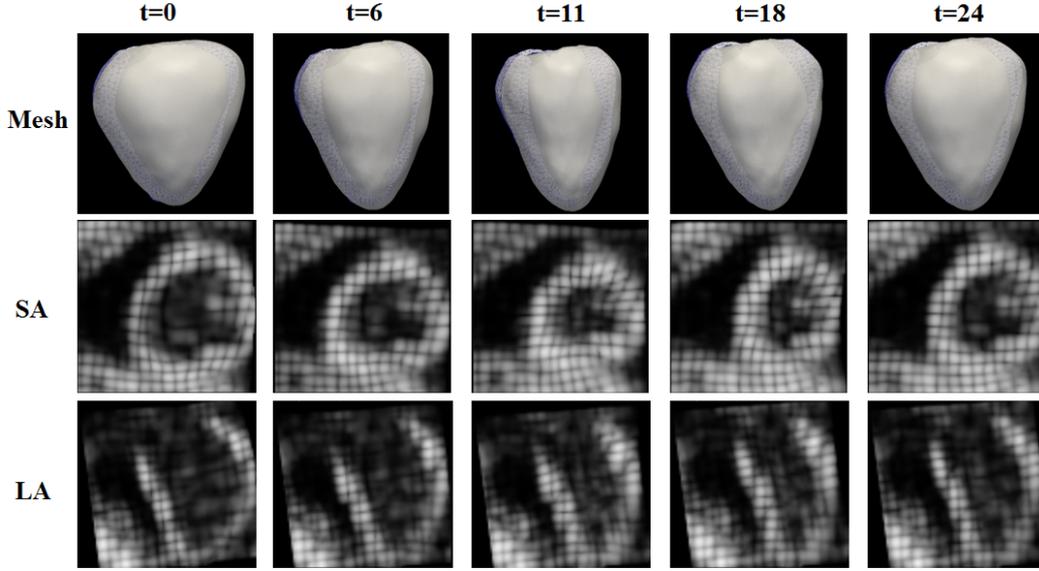
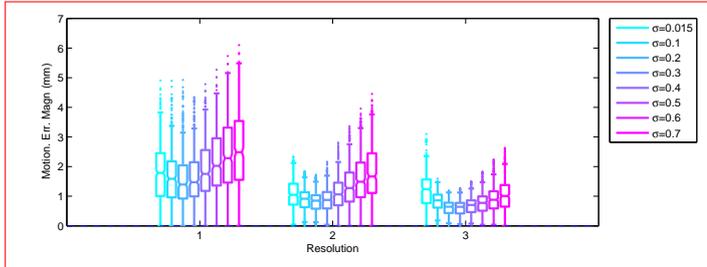


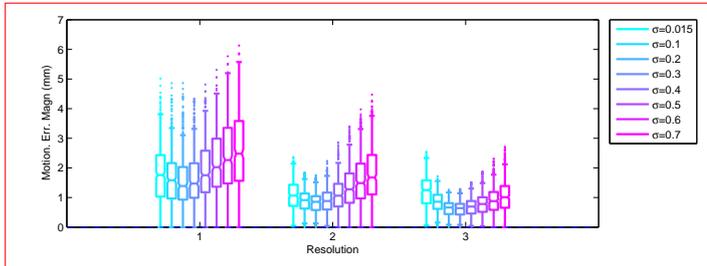
Figure 7: Ground truth meshes and short- and long-axis slices of synthetic tagged MR images (3 sequences of tagged MR images with line tagging patterns are multiplied for better visualization)

5.1.1. Tuning the regularization parameters

From Fig. 8 and 9, it can be seen that for both HARP and MP, setting either $n = 3$ and $\sigma = 0.2$, or $n = 3$ and $\sigma = 0.3$ outperform other configurations with tracking errors below 1.5 mm . When also considering strain accuracy, $n = 3$ and $\sigma = 0.3$ is the best choice on the synthetic data. Bland Altman plots of *radial*, *longitudinal* and *circumferential* strains of this case are shown in Fig. 10.

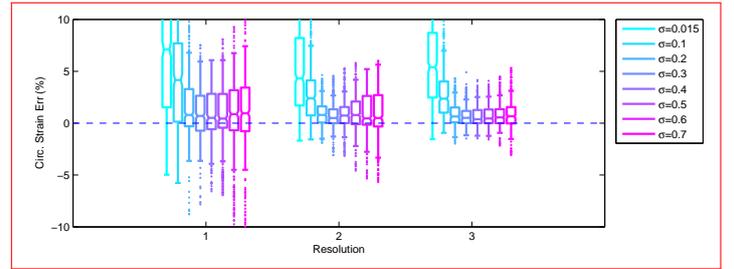


(a) HARP

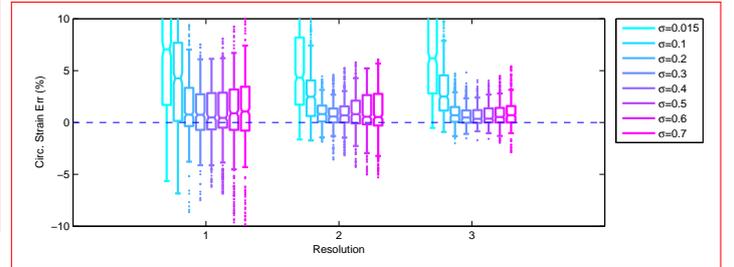


(b) MP

Figure 8: Quartile plots showing the evolution of tracking accuracy with the number of resolutions n and the kernel bandwidth σ when using HARP (a) and MP (b) on synthetic images (each data point represents the motion error of certain AHA segment at certain time frame).



(a) HARP

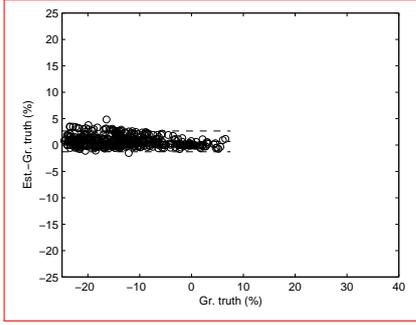


(b) MP

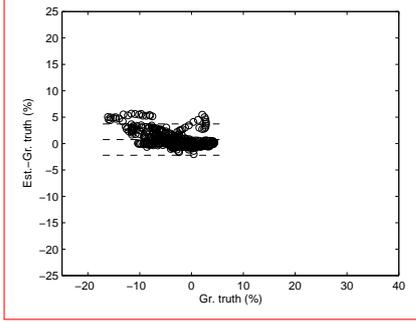
Figure 9: Quartile plots showing the evolution of circumferential strain accuracy with the number of resolutions n and the kernel bandwidth σ when using HARP (a) and MP (b) on synthetic images (each data point represents the circ. strain error of certain AHA segment at certain time frame).

5.1.2. Comparison of HARP and MP

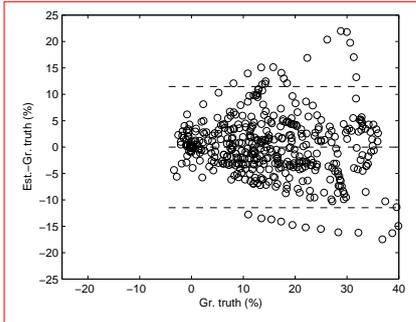
Fig. 11 summarizes the comparison between MP and HARP in terms of tracking and strain accuracies. One can see that both phase computation methods give very similar results. We computed the point-wise end-systolic motion and strain accuracies when using HARP and MP, and applied the two sample t-test (Cressie and Whitford, 1986) on them. The returned p-values are all above 0.05 when



(a) Circ. strain



(b) Long. strain



(c) Rad. strain

Figure 10: Bland Altman plots on Circ. Long. and Rad. strains when $n = 3$ and $\sigma = 0.3$ using MP on synthetic images (each data point represents the strain of certain AHA segment at certain time frame)

comparing motion, radial, longitudinal and circumferential strain accuracies. Therefore no significant differences were found between HARP and MP.

5.1.3. Influence of the shearing elements in the motion model

In our previously defined motion model (Sect. 3.2), $\{a_{rl}, a_{rc}, a_{lr}, a_{lc}, a_{cr}, a_{cl}\}$ are parameters related to shearing. To understand the role they play in the tracking, we also considered the use of a simplified model where all shearing coefficients are set to zero. With this 6-parameter motion model, we regenerated the tracking on synthetic images. The magnitude of displacement errors are compared to those obtained by the original model (Eq. 14) in Fig. 12. As seen from Fig. 12, both the median and

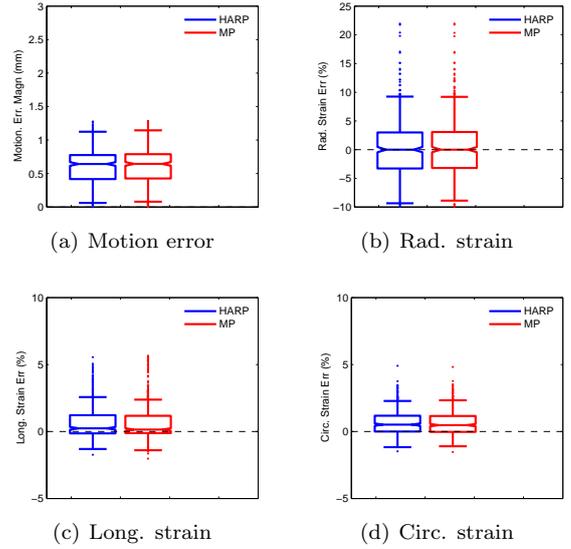


Figure 11: Comparisons of HARP and MP with respect to end-systolic tracking, rad. strain, long. strain and circ. strain accuracies for $n = 3$ and $\sigma = 0.3$ on synthetic images (each data point represents the motion, rad. strain, long. strain or circ. strain errors of certain AHA segment at the end-systole).

the dispersion of the displacement errors are reduced when taking into account the shearing elements. This result is further confirmed by two statistical tests. Applying the Levene's test (Miller Jr, 1997) to these two datasets of displacement errors returns a p-value below 0.05, rejecting the null hypothesis that their variances are equal. Also, the Wilcoxon signed-rank (Woolson, 2007) test on these two datasets reveals that their median values are statistically different.

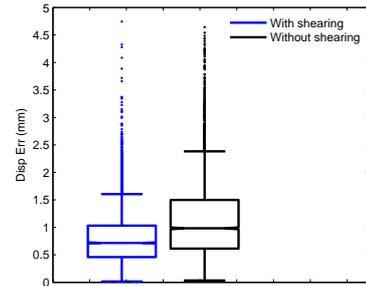


Figure 12: End-systolic motion error magnitudes with/without shearing elements using MP (parameters $n = 3$, $\sigma = 0.3$) on synthetic images (The data points represent the motion error magnitudes of all mesh points at end-systole)

5.2. Volunteer data sets

This section compares quantitatively our method to other state-of-the-art tagged MR algorithms. For doing so, we used the data provided by (Tobon-Gomez et al., 2013) to compare our tracking accuracy using manually

tracked landmarks located in the basal, mid and apical myocardium. In total, 15 volunteer datasets were used for the evaluation. The 12 landmarks were warped forwardly in time by computing barycentric coordinates in the first frame and propagating them through the sequence of volumetric meshes. This could slightly penalize our method in this evaluation as image-based tracking was performed backward in time. Fig. 13 shows box plots of motion tracking accuracy for all the 15 volunteers at end-systole. It can be seen that our method matches accuracies obtained using MEVIS and UPF algorithms while having less outliers. However, it should be noted, that the computation time of our method is in the range of minutes while UPF and MEVIS are in the range of hours as reported in (Tobon-Gomez et al., 2013) for this same dataset.

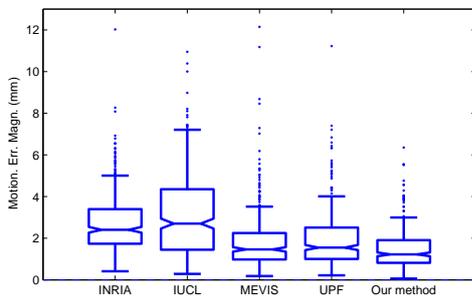


Figure 13: Motion tracking accuracy using MP (parameters $n = 3, \sigma = 0.3$) on volunteer data sets compared to four other state-of-art methods: INRIA, IUCL, MEVIS and UPF (for details of these methods please refer to (Tobon-Gomez et al., 2013)).

5.3. Patient data sets

In this paper, we show quantification results for four patients showing grading levels of fibrosis confirmed by late gadolinium-enhanced MR images. The patient datasets were acquired at GIP Cyceron. All patients gave an informed consent and the study was approved by the regional ethics committee (CPP Nord Ouest III). For comparison purposes, we also quantified three healthy volunteers acquired at the same institution and enrolled in the same study. Fig. 14 plots the dispersion of end-systolic circumferential strain values for the volunteers (v1 to v5) and the four patients (p1 to p4). The two first volunteers were taken from (Tobon-Gomez et al., 2013) while the others were scanned at Caen-CHU/Cyceron. The four patients were ordered by the number of segments where transmural fibrosis was visible on late-enhancement MR. It is expected that for diseased cases with grading levels of infarction, more segments with abnormally low deformation values will coexist with segments overcompensating for infarcted segments. Therefore, the overall dispersion of strain values should increase for more extreme cases of infarction. Fig. 14 shows that this tendency can be observed in our database: the circ. strain dispersion increases from healthy volunteers (v1-v5) to patients (p1-p4) classified

by degree of fibrosis. Also, the median circ. strain decreases from volunteers (around -20%) to patients (around -15%). We took the example of circumferential strain as we observed this component to be more robust when image quality is lower. Indeed, the four patients we report here showed very different image qualities, mainly depending on their ability to hold their breath over all acquisitions in a similar manner.

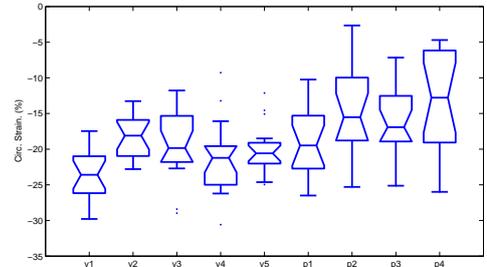


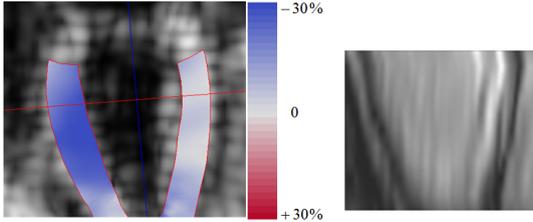
Figure 14: End-systolic Circ. strain at AHA 17 segment level, v1 and v2 are healthy volunteers from Sect. 5.2, v3,v4,v5 are healthy volunteers from GIP Cyceron, and p1, p2, p3, p4 are patients from GIP Cyceron increasingly classified by the myocardial fibrosis degree (results obtained using MP)

To look at strain at a more detailed level, in Fig. 15 we plot for the second patient p2 end-systolic circumferential strain as a colormap and display circ. strain curves at a segmental level. This patient had myocardial fibrosis at the entire inferior wall, basal and mid levels of the inferolateral wall, and basal inferoseptal wall (different degrees of fibrosis at AHA segments 3,4,5,10,11,15). Fig. 15(a)(left) and Fig. 15(b)(left) show a clear difference of circ. strain values between the inferior wall (infarcted) and the anterior wall (normal). This is coherent with the results by late-enhancement MR, as shown in Fig. 15(a)(right) and Fig. 15(b)(right) where high gray levels indicate the existence of fibrosis. A clear correlation could be seen between the low deformation regions and the infarcted areas (those with brighter image intensities). However, here the observed correlation is quite qualitative and limited because the two modalities were merely manually aligned. Indeed, these two kinds of images had different spatial resolutions, making their registration a difficult task. Besides, from the strain curves in Fig. 15(c) infarcted segments can be easily distinguished from the healthy ones.

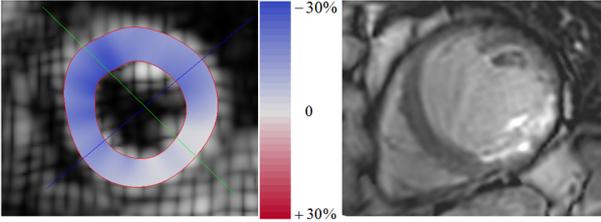
6. Discussion

The concept of anatomical regularization applied to cardiac motion tracking was directly or implicitly introduced in several previous publications (Young and Axel, 1992; Young, 1999; McLeod et al., 2013; Heyde et al., 2013; Pan et al., 2005). We here discuss the differences between those methods and the proposed one.

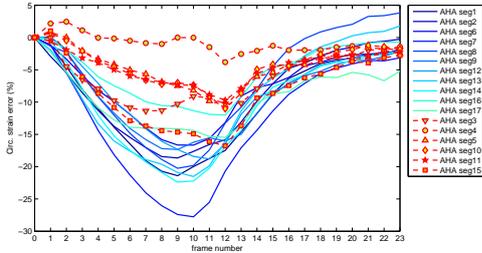
Young *et al.* (Young and Axel, 1992; Young, 1999) employed a finite-element model for regularizing the sparse



(a) End-systolic Circ. strain colormap displayed on a LA slice and late-enhancement MR image of the same location



(b) End-systolic Circ. strain colormap displayed on a mid-level SA slice and late-enhancement MR image of the same location



(c) Circ. strain evolution of 17 AHA segs

Figure 15: Circ. strain plotted as a colormap {(a),(b)} and over time {(c)} for patient p2 in Fig. 14. This patient had fibrosis at the entire inferior wall, part of the inferolateral wall and part of the inferoseptal wall (corresponding to AHA segments 3,4,5,10,11,15). The colormap is shown at end of systole. Circ. strain curves are shown for all 17 segments: solid lines display normal segments while curves with markers show segments with fibrosis (results obtained using MP)

image-based displacements. The variation of the deformation gradient tensor in *radial*, *longitudinal* and *circumferential* directions was used as the penalty term. The model was recently used to improve registration of a generic template to patient specific biventricular geometry (Gilbert et al., 2014). Although the mesh is defined in anatomical coordinates, the projection of sparse input displacements on the mesh is done in Cartesian coordinates. Our approach differs by projecting displacements in anatomical coordinates which we supposed to be more appropriate for the proposed regularization scheme.

McLeod *et al.* (McLeod et al., 2013) used an affine regional approximation of displacements. Affine transformations were mixed through weight functions that are following anatomy. However, the affine transformations were defined in the Cartesian space, rather than exploiting the anatomical system of coordinates. We preferred to use

anatomical coordinates for both defining window functions and the transformation model. This allows to regularize and constrain the motion field in the system used for reporting strain values. In our model, first order parameters are directly linked to *radial*, *longitudinal* and *circumferential* strains. Pan *et al.* (Pan et al., 2005) interpolated 2D HARP displacements by a Gaussian kernel defined using the geodesic distance on the mesh. This is quite similar to our window function. However, as mentioned above, we differ in not only defining window functions that are following the anatomy, but also regularizing and constraining HARP displacements in anatomical coordinates. Besides, we processed the recent 3D CSPAMM data introduced in (Rutz et al., 2008). Compared with the tagging data used in both (Pan et al., 2005) and (Young, 1999), each volume has more slices and images have better in-plane resolution.

Heyde *et al.* (Heyde et al., 2013) performed the tracking of ultrasound images in anatomical coordinates. However, this was achieved by resampling all images in this space prior to the tracking. Such a resampling would introduce non-linear distortion in image intensities. Also, it would modify the spectrum of tagged MR images, modifying the band-pass nature of the signal. This would violate HARP assumptions. Thus, we preferred to use anatomical coordinates in the representation of the transformation.

This study has several technical limitations that are listed and discussed hereafter. First, the method is dependent of the segmentation of the LV done on the last frame as it serves both for defining the set of points to track and a local coordinate system used for regularization. This limitation also exists in the original HARP algorithm and in any method using an LV mask as region of interest. Second, we assumed the phase images to be available as 3D information while each image stack has a low interslice spatial resolution. However, each stack has a higher voxel spacing in a different dimension, hence reducing the overall impact of this low resolution. Also, since the local windows have a full transmural extent, the motion model will not capture the increase in circumferential shortening and radial thickening occurring from epicardium to endocardium (Clark et al., 1991; Rademakers et al., 1994). Such a choice is justified by the limits of the acquisition scheme where the number of tags is low from endocardium to epicardium, making the information contained in the radial direction too poor for capturing radial deformation. **Admittedly, defining multiple windows from endocardium to epicardium would allow to quantify transmural changes in deformation thanks to the good tagging resolution in the circumferential direction. Transmural changes in circumferential strain would then propagate to radial strain with the incompressibility constraint. This would require a more detailed validation and is currently left to future work.**

Besides, our method depends on HARP accuracy and could be prone to tag jumps. We observed that tag jumps mostly occur in the longitudinal direction and that the multiscale strategy helps limiting tag jumps. Furthermore,

at the end of systole, the myocardial incompressibility does not hold anymore due mainly to blood being squeezed out of the myocardium. The myocardial volume conservation is therefore an approximation (of about 8%) at the end of systole. Nonetheless, as explained earlier in this section, image intensities only will not suffice to capture the radial deformation. The incompressibility constraint was therefore used to “help” estimating the radial deformation, rather than as a hard constraint influencing all the deformation components. Moreover, using engineering strains as obtained from length changes in the volumetric mesh edges neglects shear-induced rotations. However, note that this approximation is standard in clinical echocardiography as reported by (Voigt et al., 2014). It results from applying 1D concepts, that are more intuitive to the clinician. **In addition**, for simplifying the parameter tuning, σ was set to be homogeneous for all resolutions. Nonetheless, σ is multiplied by the window extent to obtain σ_i (Eq. 19). As the window extent decreases with the resolution, so does σ_i . **Finally, the tagged images were not apodized before computing the Fourier transform. As a result, some artifacts (the vertical lines) appear in Fig. 4(b). However, those artifacts’ influence on the tracking is rather limited because most of them will be filtered out by the bandpass filter (Fig. 4(c),4(e)). Nonetheless, a study of the apodization’s impact on the tracking would help to determine its benefits. It is currently left for future work.**

There are also several limitations with respect to the validation. We performed a comparison to the state-of-the-art on a reduced number of cases and landmarks, benchmarking tracking accuracy only. Besides the relatively limited number of methods present in this evaluation, we also see the need for a tagged MR validation framework yielding dense ground truth displacement *and strain* values. Clearly, comparing the accuracy on displacements only overlooks many challenges related to strain, including the impact of the regularization method on strain accuracy and sensitivity. Besides, our strategy for generating synthetic images was too simplistic for being representative of challenges inherent to clinical images. Images were produced through warping of a single frame, ignoring artifacts induced by valve motion or tag contrast variations over the cardiac cycle. Moreover, regarding simulation aspects, an electromechanical simulation should be done specifically on the patient geometry. Also, MR simulators would likely be more realistic than simple image warping. Finally, the experiments performed in this paper are preliminary and only aim at demonstrating the applicability of our method to process clinical images. Image quality for patient datasets was found to be more heterogeneous than volunteer data. Evaluating the clinical impact of 3D tagged MR in general - and this method in particular - requires a thorough study on the correlation between the low deformation regions quantified by our method with the fibrosis areas revealed by late-enhancement MR. The interaction mechanisms between the fibrosed and the surrounding contracting regions need to be further studied on

a larger patient population, taking into account the different tissue properties.

7. Conclusion

This paper introduced a novel algorithm that extends HARP to handle 3D tagged MR data. While doing so, we proposed an original regularization method done in an anatomical space of coordinates. HARP results were regularized according to a windows-weighted regression method that maintains a low computational complexity. Our implementation performs LV tracking and strain computation in less than a minute, a time range compatible with clinical practice requirements. The algorithm was evaluated at three levels: on synthetic, healthy volunteers and patient data. On healthy volunteers, tracking accuracy was found to be similar to the best candidates of a recent benchmark. Strain accuracy was evaluated on synthetic data (as no public data for strain was available), showing low bias and strain errors under 5 percents (excluding outliers) for longitudinal and circumferential strains. As reported in the literature, radial strain was found to be more challenging but had the second and third quartiles of strain errors in the $(-5\%, 5\%)$ range, thanks to an implementation of the incompressibility constraint that solely corrects this strain component. Our method was shown to be applicable to clinical data by first correlating strain dispersion with the extent of transmural fibrosis. Lower deformation values were also observed inside and around the infarcted region. In synthetic data, we also compared the effectiveness of tracking the HARP phase and MP, although no significant differences were found between them. As future work, we aim at extending clinical findings by applying this algorithm to more patients. We will also develop a more elaborated validation framework for the validation of tagged MR tracking algorithms by extending our pipeline for the generation of synthetic images.

Acknowledgements

This work was in part supported by an ANRT grant (N° 2013/1534) and the 7th framework program (FP7/2007-2013, N° 611823).

References

- Alessandrini, M., Basarab, A., Liebgott, H., Bernard, O., 2013. Myocardial motion estimation from medical images using the monogenic signal. *Image Processing, IEEE Transactions on* 22, 1084–1095.
- Alessandrini, M., Liebgott, H., Basarab, A., Clarysse, P., Bernard, O., 2012. Monogenic signal for cardiac motion analysis from tagged magnetic resonance image sequences, in: *Computing in Cardiology (CinC)*, 2012, IEEE. pp. 685–688.
- Amini, A., Chen, Y., Elayyadi, M., Radeva, P., 2001. Tag surface reconstruction and tracking of myocardial beads from spamm-mri with parametric b-spline surfaces 20, 94–103.

- Arts, T., Prinzen, F., Delhaas, T., Milles, J., Rossi, A., Clarysse, P., 2010. Mapping displacement and deformation of the heart with local sine-wave modeling 29, 1114–1123.
- Bookstein, F.L., 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence* 11, 567–585.
- Chandrashekar, R., Mohiaddin, R., Rueckert, D., 2004. Analysis of 3-d myocardial motion in tagged mr images using nonrigid image registration. *Medical Imaging, IEEE Transactions on* 23, 1245–1250.
- Chen, T., Wang, X., Chung, S., Metaxas, D., Axel, L., 2010. Automated 3d motion tracking using gabor filter bank, robust point matching, and deformable models. *Medical Imaging, IEEE Transactions on* 29, 1–11.
- Clark, N.R., Reichel, N., Bergey, P., Hoffman, E.A., Brownson, D., Palmon, L., Axel, L., 1991. Circumferential myocardial shortening in the normal human left ventricle. assessment by magnetic resonance imaging using spatial modulation of magnetization. *Circulation* 84, 67–74.
- Cressie, N., Whitford, H., 1986. How to use the two sample t-test. *Biometrical Journal* 28, 131–148.
- De Craene, M., Marchesseau, S., Heyde, B., Gao, H., Alessandrini, M., Bernard, O., Piella, G., Porras, A., Saloux, E., Tautz, L., et al., 2013. 3d strain assessment in ultrasound (straus): A synthetic comparison of five tracking methodologies. *IEEE Transactions on Medical Imaging* .
- De Craene, M., Tobon-Gomez, C., Butakoff, C., Duchateau, N., Piella, G., Rhode, K.S., Frangi, A.F., 2012. Temporal diffeomorphic free form deformation (TDDFFD) applied to motion and deformation quantification of tagged mri sequences, in: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*, Springer Berlin / Heidelberg. pp. 68–77. URL: http://dx.doi.org/10.1007/978-3-642-28326-0_7.
- Fischer, S.E., McKinnon, G., Maier, S., Boesiger, P., 1993. Improved myocardial tagging contrast. *Magnetic resonance in medicine* 30, 191–200.
- Florack, L., van Assen, H., Suinesiaputra, A., 2007. Dense multi-scale motion extraction from cardiac cine mr tagging using harp technology, in: *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, IEEE. pp. 1–8.
- Gilbert, K., Cowan, B.R., Suinesiaputra, A., Occlshaw, C., Young, A.A., 2014. Rapid d-affine biventricular cardiac function with polar prediction, in: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2014*. Springer, pp. 546–553.
- Heyde, B., Barbosa, D., Claus, P., Maes, F., D’hooge, J., 2013. Three-dimensional cardiac motion estimation based on non-rigid image registration using a novel transformation model adapted to the heart, in: *STACOM*, Springer. pp. 142–150.
- Kause, H.B., Filatova, O.G., Duits, R., Bruurmijn, L.M., Fuster, A., Westenberg, J.J., Florack, L.M., van Assen, H.C., 2014. Direct myocardial strain assessment from frequency estimation in tagging mri, in: *Statistical Atlases and Computational Models of the Heart. Imaging and Modelling Challenges*. Springer, pp. 212–219.
- Kuijper, J., Jansen, E., Marcus, J.T., van Rossum, A.C., Heethaar, R.M., 2001. Improved harmonic phase myocardial strain maps. *Magnetic Resonance in Medicine* 46, 993–999.
- Liang, J., Wang, Y., Jia, Y., 2007. Cardiac motion estimation from tagged mri using 3d-harp and nurbs volumetric model, in: *Computer Vision–ACCV 2007*. Springer, pp. 512–521.
- Liu, X., Abd-Elmoniem, K.Z., Stone, M., Murano, E.Z., Zhuo, J., Gullapalli, R.P., Prince, J.L., 2012. Incompressible deformation estimation algorithm (idea) from tagged mr images. *Medical Imaging, IEEE Transactions on* 31, 326–340.
- Liu, X., Prince, J.L., 2010. Shortest path refinement for motion estimation from tagged mr images. *Medical Imaging, IEEE Transactions on* 29, 1560–1572.
- Makram-Ebeid, S., Somphone, O., 2007. Non-rigid image registration using a hierarchical partition of unity finite element method, in: *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, IEEE. pp. 1–8.
- Marchesseau, S., Delingette, H., Sermesant, M., Ayache, N., 2013. Fast parameter calibration of a cardiac electromechanical model from medical images based on the unscented transform. *Biomechanics and modeling in mechanobiology* 12, 815–831.
- McLeod, K., Seiler, C., Toussaint, N., Sermesant, M., Pennec, X., 2013. Regional analysis of left ventricle function using a cardiac-specific polyaffine motion model, in: *Functional Imaging and Modeling of the Heart*. Springer, pp. 483–490.
- Miller Jr, R.G., 1997. *Beyond ANOVA: Basics of applied statistics*. CRC Press.
- Mory, B., Somphone, O., Prevost, R., Ardon, R., 2012. Real-time 3d image segmentation by user-constrained template deformation, in: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012*. Springer, pp. 561–568.
- Osman, N.F., Kerwin, W.S., McVeigh, E.R., Prince, J.L., 1999. Cardiac motion tracking using cine harmonic phase (harp) magnetic resonance imaging. *Magnetic resonance in medicine: official journal of the Society of Magnetic Resonance in Medicine/Society of Magnetic Resonance in Medicine* 42, 1048.
- Osman, N.F., McVeigh, E.R., Prince, J.L., 2000. Imaging heart motion using harmonic phase mri. *Medical Imaging, IEEE Transactions on* 19, 186–202.
- Pan, L., Prince, J.L., Lima, J.A., Osman, N.F., 2005. Fast tracking of cardiac motion using 3d-harp. *Biomedical Engineering, IEEE Transactions on* 52, 1425–1435.
- Prince, J., McVeigh, E., 1992. Motion estimation from tagged mr image sequences. *Medical Imaging, IEEE Transactions on* 11, 238–249.
- Qian, Z., Liu, Q., Metaxas, D.N., Axel, L., 2011. Identifying regional cardiac abnormalities from myocardial strains using nontracking-based strain estimation and spatio-temporal tensor analysis. *Medical Imaging, IEEE Transactions on* 30, 2017–2029.
- Rademakers, F.E., Rogers, W.J., Guier, W.H., Hutchins, G.M., Siu, C.O., Weisfeldt, M.L., Weiss, J.L., Shapiro, E.P., 1994. Relation of regional cross-fiber shortening to wall thickening in the intact heart. three-dimensional strain analysis by nmr tagging. *Circulation* 89, 1174–1182.
- Rutz, A.K., Ryf, S., Plein, S., Boesiger, P., Kozerke, S., 2008. Accelerated whole-heart 3d cspamm for myocardial motion quantification. *Magnetic Resonance in Medicine* 59, 755–763.
- Shi, W., Zhuang, X., Wang, H., Duckett, S., Luong, D.V., Tobon-Gomez, C., Tung, K., Edwards, P.J., Rhode, K.S., Razavi, R.S., et al., 2012. A comprehensive cardiac motion estimation framework using both untagged and 3-d tagged mr images based on nonrigid registration. *Medical Imaging, IEEE Transactions on* 31, 1263–1275.
- Tobon-Gomez, C., Craene, M.D., McLeod, K., Tautz, L., Shi, W., Hennemuth, A., Prakosa, A., Wang, H., Carr-White, G., Kapetanakis, S., Lutz, A., Rasche, V., Schaeffter, T., Butakoff, C., Friman, O., Mansi, T., Sermesant, M., Zhuang, X., Ourselin, S., Peitgen, H.O., Pennec, X., Razavi, R., Rueckert, D., Frangi, A., Rhode, K., 2013. Benchmarking framework for myocardial tracking and deformation algorithms: An open access database. *Medical Image Analysis* 17, 632 – 648. URL: <http://www.sciencedirect.com/science/article/pii/S1361841513000388>, doi:<http://dx.doi.org/10.1016/j.media.2013.03.008>.
- Voigt, J.U., Flachskampf, F., 2004. Strain and strain rate. *Zeitschrift für Kardiologie* 93, 249–258.
- Voigt, J.U., Pedrizzetti, G., Lysyansky, P., Marwick, T.H., Houle, H., Baumann, R., Pedri, S., Ito, Y., Abe, Y., Metz, S., et al., 2014. Definitions for a common standard for 2d speckle tracking echocardiography: consensus document of the eacvi/ase/industry task force to standardize deformation imaging. *European Heart Journal–Cardiovascular Imaging* , jcu184.
- Wang, H., Stoeck, C.T., Kozerke, S., Amini, A.A., 2013. 3d left-ventricular deformation analysis from 3d cspamm with 3d sinmod. *Proc. Intl. Soc. Mag. Reson. Med.* .
- Woolson, R., 2007. Wilcoxon signed-rank test. *Wiley Encyclopedia of Clinical Trials* .
- Young, A.A., 1999. Model tags: direct three-dimensional tracking of heart wall motion from tagged magnetic resonance images. *Med-*

ical image analysis 3, 361–372.

Young, A.A., Axel, L., 1992. Three-dimensional motion and deformation of the heart wall: estimation with spatial modulation of magnetization—a model-based approach. *Radiology* 185, 241–247.

Young, A.A., Kraitchman, D.L., Dougherty, L., Axel, L., 1995. Tracking and finite element analysis of stripe deformation in magnetic resonance tagging. *Medical Imaging, IEEE Transactions on* 14, 413–421.