



Libraries and Learning Services

University of Auckland Research Repository, ResearchSpace

Version

This is the publisher's version. This version is defined in the NISO recommended practice RP-8-2008 <http://www.niso.org/publications/rp/>

Suggested Reference


Nitzken, M., Beache, G. M., Ismail, M., Gimel'farb, G., & El-Baz, A. (2016). Improving full-cardiac cycle strain estimation from tagged CMR by accurate modeling of 3D image appearance characteristics. *Egyptian Journal of Radiology and Nuclear Medicine*, 47(1), 83-94. doi: 10.1016/j.ejnm.2015.10.014

Copyright

Items in ResearchSpace are protected by copyright, with all rights reserved, unless otherwise indicated. Previously published items are made available in accordance with the copyright policy of the publisher.

This is an open-access article distributed under the terms of the [Creative Commons Attribution NonCommercial NoDerivatives License](#).

For more information, see [General copyright](#), [Publisher copyright](#), [SHERPA/RoMEO](#).



ORIGINAL ARTICLE

Improving full-cardiac cycle strain estimation from tagged CMR by accurate modeling of 3D image appearance characteristics



Matt Nitzken^a, Garth M. Beache^b, Marwa Ismail^a, Georgy Gimel'farb^c,
Ayman El-Baz^{a,*}

^a BioImaging Laboratory, Bioengineering Department, University of Louisville, Louisville, KY 40292, USA

^b Radiology Department, University of Louisville, Louisville, KY 40202, USA

^c Department of Computer Science, University of Auckland, Auckland 1142, New Zealand

Received 26 August 2015; accepted 27 October 2015

Available online 14 November 2015

KEYWORDS

Tagged CMR image;
MGRF;
LCDG;
Strain parameters

Abstract To improve the tagged cardiac magnetic resonance (CMR) image analysis, we propose a 3D (2D space + 1D time) energy minimization framework, based on learning first- and second-order visual appearance models from voxel intensities. The former model approximates the marginal empirical distribution of intensities with two linear combinations of discrete Gaussians (LCDG). The second-order model considers an image of a sample from a translation–rotation invariant 3D Markov–Gibbs random field (MGRF) with multiple pairwise spatiotemporal interactions within and between adjacent temporal frames. Abilities of the framework to accurately recover noise-corrupted strain slopes were experimentally evaluated and validated on 3D geometric phantoms and independently on *in vivo* data. In multiple noise and motion conditions, the proposed method outperformed comparative image filtering in restoring strain curves and reliably improved HARP strain tracking during the entirety of the cardiac cycle. According to these results, our framework can augment popular spectral domain techniques, such as HARP, by optimizing the spectral domain characteristics and thereby providing more reliable estimates of strain parameters.

© 2015 The Authors. The Egyptian Society of Radiology and Nuclear Medicine. Production and hosting by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Tagged cardiac magnetic resonance (CMR) imaging is widely used for detailed and non-invasive visualization of myocardium motion/deformation, with a full spatial geometric concordance (1). MR tagging places on body tissues a pre-defined pattern of virtual fiducial markers (tags) created by radiofrequency modulated magnetic spins (2) to

* Corresponding author at: Lutz Hall, University of Louisville, Room 423, Louisville, KY 40292, USA. Tel.: +1 (502) 852 5092; fax: +1 (502) 852 6806.

E-mail address: aselba01@louisville.edu (A. El-Baz).

Peer review under responsibility of Egyptian Society of Radiology and Nuclear Medicine.

<http://dx.doi.org/10.1016/j.ejrnmm.2015.10.014>

0378-603X © 2015 The Authors. The Egyptian Society of Radiology and Nuclear Medicine. Production and hosting by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

complement traditional anatomical images with functional data from the tagged images, e.g., strain and strain rate in the heart (3). Local diseases, such as coronary atherosclerosis, as well as global ones, e.g., heart failure and diabetes, result in wall contractile dysfunction that manifests itself on the tagged CMR images.

1.1. Spatial vs. spectral methods

Images have equivalent dual spatial and spectral (frequency) representations, related by Fourier transforms, the tagged MR image analysis exploiting both the domains (4). Model-based or model-free spatial techniques estimate tissue motion and strain by identifying grid intersection points in a given image. Finite element analysis, employed typically to track actual pixels or voxels over a time series, is often computationally expensive due to required substantial image preprocessing and segmentation (5). Spectral techniques exploring the magnitudes and phases of spectrum harmonics build a tissue motion field using motion-invariant phases of material points and information from magnitude. Spectral analysis has become popular due to fast and arguably more robust tracking of harmonics over the cardiac cycle than spatial pixel or voxel tracking (6–8). Spatial processing proposed in this paper enhances the reliability of spectral analysis techniques, such as the popular Harmonic Phase (HARP).

1.2. HARP and spectral tracking

HARP computes phase images from the sinusoidal tagged MR images by band-pass filtering in the spectral domain (4,6–8), where smooth or sharp spatial transitions concentrate related information in central lobes or increase the number of the higher-order lobes, respectively. Due to duality, the spatial noise and corruption manifest themselves in the spectral domain, too. Noise in a spatial image can cause false phase values. For a time series, the corrupted reference time frame affects the subsequent frames (4,9). Spectral tracking assumes that points found in a tissue do not move substantially between two successive time frames (10). High movement rates of the tissue, and/or a low temporal image resolution, and/or non-optimal MR tag parameters increase data corruption and noise (2,4). The resulting spectral tracking failures often require manual identification and correction of mistracked points (11,12), although this can become time-inefficient over multiple frames that constitute the cardiac cycle.

Our goal is to modify spatial images for more reliable spectral tracking by reducing noise across a tag line, making tag edges less sharp, and amplifying tag-to-background contrast. Some prior works that improve the accuracy of the spectral tracking are discussed below.

1.3. Related work on improving tagged MR images

To directly improve the tagged CMR images in the spatial domain, Yang and Murase (13) explored impacts of histogram modification and local contrast enhancement on the tag-to-background contrast and automatic segmentation of left ventricle in 2D tagged MR images. The amplified contrast made the tag grid lines appear sharper and mostly unbroken. Filtering was added then before the histogram equalization

(14). Although the spatial tracking has been improved to some extent, the increased spectral noise due to sharper lines was able to hamper spectral tracking techniques, such as the HARP. Complex band-pass filtering by Yuan et al. (15) using wavelet decomposition was computationally inefficient and also sharpened edges of tags, limiting the spectral analysis. Li (16) used intensity thresholds and template matching to classify tag and background points in the image. However, the expert knowledge was required to construct the tag templates (masks). Li and Yu (9) attempted to improve tag characteristics only on the initial frame of a time series.

1.4. Related work on direct HARP modification

Khalifa et al. (17) modified components of an active contour with the HARP, rather than increased the image quality. Several other HARP modifications have been proposed in (2,10,18,19).

1.5. Modeling 3D image appearance

Our framework refines the spatial MR images to optimize their traditional spectral analysis, such as HARP, by reducing the noise across a tag line and sharpness of the tag edges. The refinement slightly improves the tag-to-background contrast, too.

For this purpose, a special 3D energy function, which accounts for both the noise and contrast, is derived using the probabilistic first- and second-order models of visual appearance of the tagged CMR images. The first-order model approximates marginal probability distributions of the voxel-wise tag and background signals with the LCDGs to initially classify the signals. Then the second-order MGRF model with analytically learned multiple pairwise spatiotemporal (in-plane space/time) signal interactions facilitates the accurate final tag/background separation.

Fig. 1 and Algorithm 1 detail processing steps of the proposed novel stochastic framework applying an energy minimization principle to optimize spectral representations of the tagged cardiac MR images and improve current spectral tracking methods. Unlike other generalized MR image restoration frameworks, such as those developed by Raj et al. (20,21), our approach requires no prior information. Also, due to its modularity, the images are refined directly, thereby permitting to use existing commercial packages. Simultaneous noise removal and MR tagging optimizations tend also to the spatial domain smoothing that further favorably differentiates our framework from those that restore all image characteristics (21).

2. Methods

2.1. Visual appearance modeling

The appearance of a tagged CMR image is described with two probabilistic models: a first-order LCDG and a second-order generic translation–rotation invariant central-symmetric MGRF with voxel-wise and multiple characteristic pairwise voxel dependencies. The MGRF potentials for the interdependent voxel pairs are functions of signal differences.

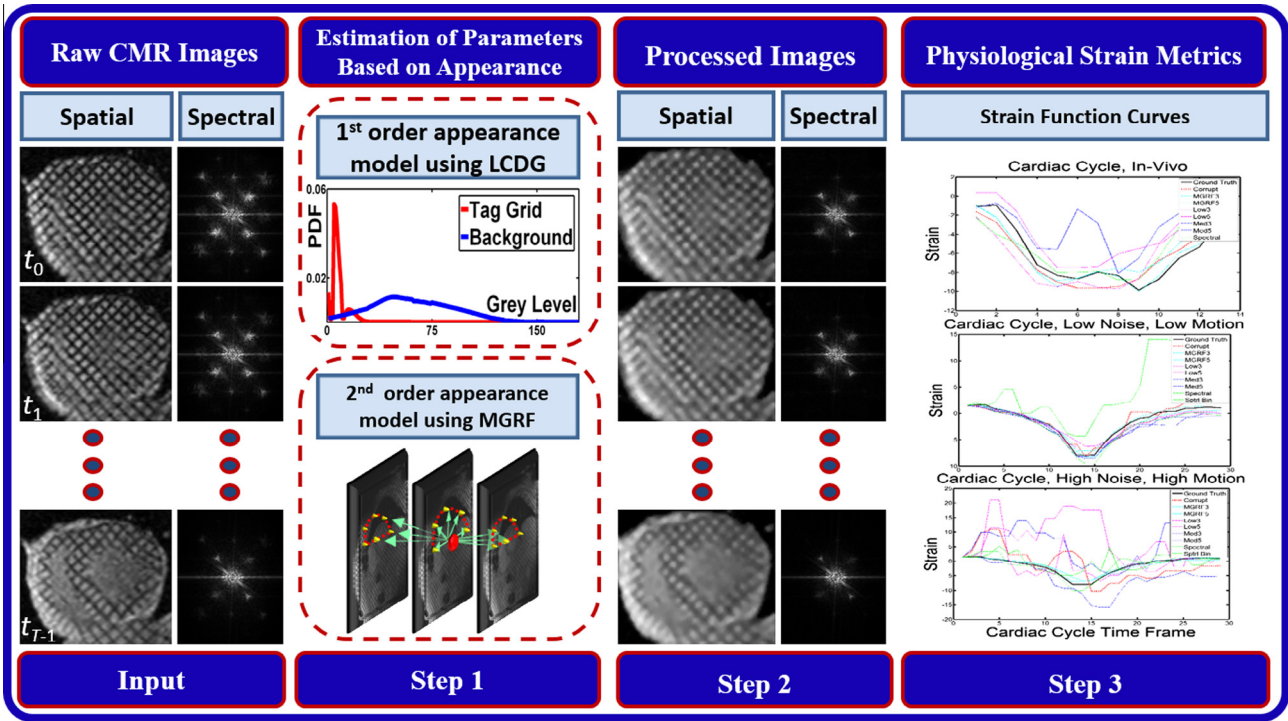


Fig. 1 Refining tagged CMR images before their quantitative HARP-based analysis. An input 3D (2D + time) image series is explored at Step 1 (detailed in Algorithm 1) using both the first- and second-order appearance models. Then it is refined at Step 2 and analyzed at Step 3 with the traditional HARP for computing the goal descriptive metrics.

2.1.1. Basic notation

Let $\mathbf{R} = [(x, y, t) : x = 0, \dots, X-1; y = 0, \dots, Y-1; t = 0, \dots, T-1]$ be a finite 3D spatiotemporal lattice of size $R = XYT$ with two planar, (x, y) , and one temporal, t , coordinates of voxels $\mathbf{r} = (x, y, t)$. Let $\mathbf{Q} = \{0, \dots, Q-1\}$ be a finite set of voxel-wise signals – gray levels, or intensities. The lattice \mathbf{R} supports 3D CMR images $\mathbf{g} = [g(\mathbf{r}) : \mathbf{r} \in \mathbf{R}; g(\mathbf{r}) \in \mathbf{Q}]$, being each a series of 2D planar cross-sections, taken at successive time instants.

2.2. The first-order appearance model: LCDG

A discrete Gaussian (DG) $\Psi_\theta = (\psi(q|\theta) : q \in \mathbf{Q})$ is a Q -component 1D probability distribution. Each component $\psi(q|\theta)$ is obtained by integrating a continuous 1D Gaussian probability density with given parameters θ (the mean and variance) over the interval associated with q (22,23). To improve the tag-to-background contrast in a CMR image to be refined, the empirical marginal distribution of the image signals is closely approximated with an LCDG $P_{w,\theta} = [P_{w,\theta}(q) : q \in \mathbf{Q}]$; $\sum_{q \in \mathbf{Q}} P_{w,\theta}(q) = 1$, with two positive dominant DGs and multiple sign-alternate subordinate DGs:

$$P_{w,\theta}(q) = \sum_{k=1}^{K_p} w_{p,k} \psi(q|\theta_{p,k}) - \sum_{l=1}^{K_n} w_{n,l} \psi(q|\theta_{n,l}) \quad (1)$$

where K_p ; $K_p \geq 2$, and K_n ; $K_n \geq 2$ are total numbers of the positive and negative DGs. The non-negative weights $\mathbf{w} = [w_{p,k} : k = 1, \dots, K_p; w_{n,l} : l = 1, \dots, K_n]$ meet the obvious

common constraint $\sum_{k=1}^{K_p} w_{p,k} - \sum_{l=1}^{K_n} w_{n,l} = 1$. The subordinate DGs closely approximate deviations of the empirical distribution from the conventional mixture of the dominant positive DGs.

This first-order LCDG model is built and separated into the two LCDG sub-models for the tag lines and their background, respectively, with Expectation–Maximization (EM)-based techniques introduced previously in (23). The marginal signal distributions for the tagged MR images have two dominant modes – one mode for the tag lines and other for their background. Given the number K of the dominant DGs (here, $K = 2$), the numbers $K_p - K$ and K_n of the positive and negative subordinate DGs, respectively, as well as the weights \mathbf{w} and parameters θ (the means and variances) of all DGs are estimated first to produce an initial LCDG, closely approximating the empirical distribution. Then the weights and parameters of the two found dominant DGs and the found K_p and K_n subordinate DGs are refined with due account of their alternate signs. The refined LCDG $P_{w,\theta}$ is finally partitioned into the sub-models $P_{\text{vox},\alpha} = [P_{\text{vox},\alpha}(q) : q \in \mathbf{Q}]$, one per class $\alpha \in \{\text{tag, background}\}$, by associating the subordinate DGs with the dominant ones as to minimize the tag-to-background misclassification rate.

Fig. 2 illustrates basic steps of building both the LCDG sub-models after collecting the marginal empirical probability distribution of gray values for the input images:

- (1) The obtained empirical marginal is approximated with a mixture of two positive DGs, relating each to the dominant mode.

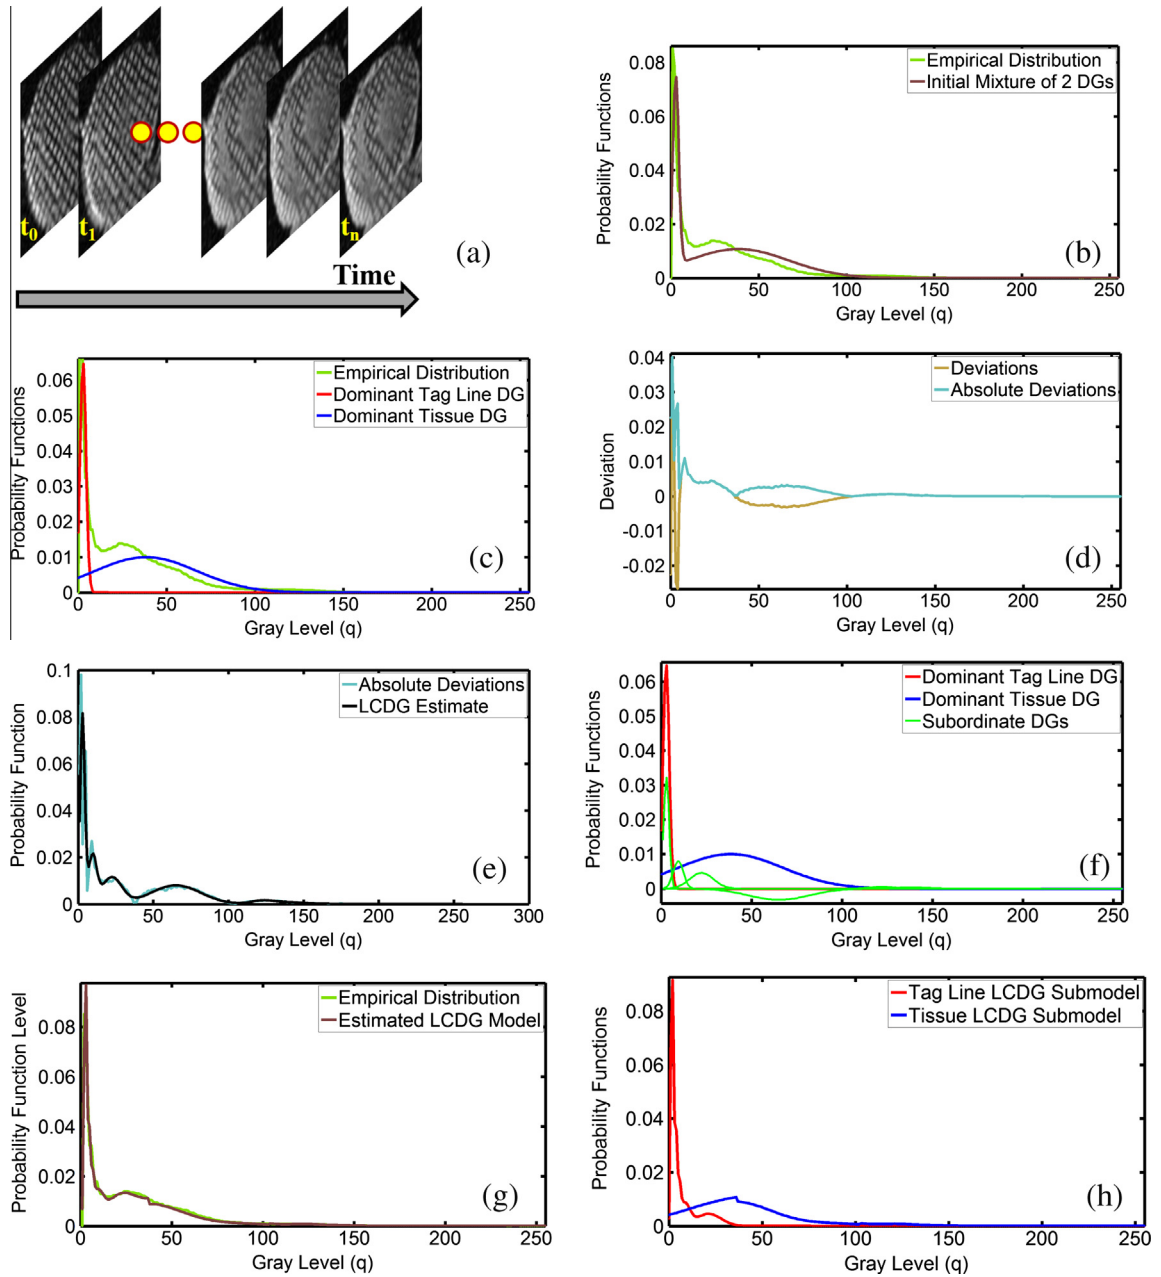


Fig. 2 Typical tagged MR images (a); the empirical marginal probability distribution (b) of their signals; the estimated mixture (c) of the two dominant DGs; deviations (d) between the empirical distribution and the estimated dominant mixture; the estimated subordinate LCDG model (e) of the absolute deviations; the individual estimated DGs (f); the final refined LCDG (g), and the LCDG sub-models (h) for the tagged lines and background tissues.

- (2) The absolute deviations between the empirical marginal and the estimated dominant mixture are approximated with the mixture of the scaled DGs and the latter is separated into the positive and negative subordinate DGs as described in (22,23).
- (3) The obtained rescaled positive and negative subordinate DGs are added to the dominant mixture.
- (4) The resulting mixed LCDG is refined to increase the approximation accuracy.
- (5) Finally, the refined LCDG is partitioned into the two LCDG sub-models, one per class, by relating the subordinate DGs to the proper dominant ones.

2.3. The second-order appearance model: MGRF

MGRFs (Markov random fields with Gibbs probability distributions) on rectangular lattices that have been widely used in image modeling since the 1980s (24–28) mostly assume translation-invariant conditional signal dependencies. A few models also have limited rotational invariance (commonly, the same pairwise dependencies between rotated at angular steps of 90° or 45°).

To account for sizable local geometric deviations between relevant areas of successive CMR images, our framework describes a time series of images \mathbf{g} with a more general




Fig. 3 Central-symmetric 3D neighborhoods for a particular spatiotemporal voxel indicated by the central red sphere: (a) ideal nested rotation-invariant continuous neighborhoods and (b) their actual lattice-wide realization in an acquired temporal stack of planar CMR cross-sections (the neighbors are indicated by yellow dots). The “red-sphere” voxel of the 2D snapshot at the instant t interacts with its neighbors in the same, preceding, $t - 1$, and subsequent, $t + 1$, snapshots. By symmetry considerations, only the upper half of the voxel neighborhood is depicted.

translation–rotation invariant second-order MGRF. This descriptive model assumes a certain number, N , of nested characteristic central-symmetric neighborhoods $\mathbf{n}_v \subset \mathbf{R}$; $v = 1, \dots, N$, of each voxel $\mathbf{r} = (x, y, t)$ within the volume \mathbf{R} (Fig. 3). The neighborhood \mathbf{n}_v of the voxel \mathbf{r} consists of the voxels \mathbf{r}' , located at distances $d(\mathbf{r}, \mathbf{r}')$ from an indexed semi-open interval $[d_{v,\min}, d_{v,\max})$, such that $d_{v-1,\max} = d_{v,\min}$ for $v = 2, \dots, N$. The distances are measured by the L2 norms of the coordinate offsets $\mathbf{r}' - \mathbf{r}$.

The tag intensity fades with T1, as the heart progresses through the cardiac cycle. Oriented filtering of a series of tagged images along the tag evolution in time has been previously used to account for the random noise and relative informational tag reliability in the images tracked (29). Such a stack of 2D image planes, being built after scaling the DICOM images to unit values at successive temporal locations, forms an oriented 3D (2D space + time) volume of intensities to be modeled with the aforementioned second-order MGRF. Each index v in this model specifies a family of the neighboring voxel pairs, $\mathbf{C}_v = \{(\mathbf{r}, \mathbf{r}') : \mathbf{r}' - \mathbf{r} \in \mathbf{n}_v; \mathbf{r}, \mathbf{r}' \in \mathbf{R}\}$, which are considered second-order cliques of the neighborhood graph with nodes in the voxels. Fig. 3 exemplifies the nested spherical neighborhood system for the latter MGRF.

Cliques of each family \mathbf{C}_v support the same real-valued Gibbs potentials, which quantify this particular voxel interactions: $V_v(|q - q'|)$ where $q = g(\mathbf{r})$ and $q' = g(\mathbf{r}')$. The potential depends on absolute in-clique signal differences $\Delta \equiv |q - q'| \in \mathbf{D} = \{0, 1, \dots, Q - 1\}$ to account for possible local brightness, or signal offset deviations. For brevity, the potentials are represented below by the column vectors, $\mathbf{V}_v = [V_v(\Delta) : \Delta \in \mathbf{D}]^T$ where T indicates the transposition. Such an MGRF has the Gibbs probability distribution (30):

$$P(\mathbf{g}) = \frac{1}{Z_V} \exp(|\mathbf{R}| \mathbf{V}^T \mathbf{F}(\mathbf{g})) \\ \equiv \frac{1}{Z_V} \exp\left(|\mathbf{R}| \left(\mathbf{V}_{\text{vox}}^T \mathbf{F}_{\text{vox}}(\mathbf{g}) + \sum_{v=1}^N \rho_v \mathbf{V}_v^T \mathbf{F}_v(\mathbf{g}) \right)\right) \quad (2)$$

where Z_V is the normalizing factor (the partition function) depending on the first- and second-order potentials; $\mathbf{V} = [\mathbf{V}_{\text{vox}}; \mathbf{V}_v : v = 1, \dots, N]$ is the vector-column of these

potentials (with the second-order clique families $\{\mathbf{C}_v : v = 1, \dots, N\}$); $\rho_v = \frac{|\mathbf{C}_v|}{|\mathbf{R}|}$ is the relative clique family size with respect to the lattice cardinality $|\mathbf{R}| = XYT$, i.e., the relative number of cliques in the family \mathbf{C}_v , and the vector-column $\mathbf{F}(\mathbf{g}) = [\mathbf{F}_{\text{vox}}(\mathbf{g}); \rho_v \mathbf{F}_v(\mathbf{g}) : v = 1, \dots, N]$ contains empirical probabilities $f_{\text{vox}}(q|\mathbf{g})$ of signals $q \in \mathbf{Q}$ in the voxels and scaled empirical probabilities $f_v(\Delta|\mathbf{g})$ of absolute signal differences $\Delta \in \mathbf{D}$ in the cliques from the family \mathbf{C}_v over the image \mathbf{g} :

$$\mathbf{F}_{\text{vox}}(\mathbf{g}) = \left[f_{\text{vox}}(q|\mathbf{g}) = \frac{|R_q(\mathbf{g})|}{|\mathbf{R}|}; q \in \mathbf{Q} \right]; \quad (3)$$

$$\mathbf{F}_v(\mathbf{g}) = \left[f_v(\Delta|\mathbf{g}) = \frac{|\mathbf{C}_{v,\Delta}(\mathbf{g})|}{|\mathbf{C}_v|}; \Delta \in \mathbf{D} \right];$$

$$\sum_{q \in \mathbf{Q}} f_{\text{vox}}(q|\mathbf{g}) = 1; \quad \sum_{\Delta \in \mathbf{D}} f_v(\Delta|\mathbf{g}) = 1$$

Here, the sublattice $R_q(\mathbf{g})$ contains all the voxels \mathbf{r} , such that $g(\mathbf{r}) = q$, and the subfamily $\mathbf{C}_{v,\Delta}(\mathbf{g})$ contains all the pairwise cliques $(\mathbf{r}, \mathbf{r}')$ of the family \mathbf{C}_v , such that $|g(\mathbf{r}) - g(\mathbf{r}')| = \Delta$. The maximum likelihood estimates of the potentials have analytical first approximations (30):

$$V_{\text{vox}}(q) = \lambda(f(q|\mathbf{g}) - f_{\text{irf:vox}}(q)); \quad q \in \mathbf{Q}; \quad (4)$$

$$V_v(\Delta) = \lambda(f_v(\Delta|\mathbf{g}) - f_{\text{irf:dif}}(\Delta)); \quad \Delta \in \mathbf{D}; \quad v = 1, \dots, N$$

where the scaling factor λ is also estimated analytically and $f_{\text{irf:vox}}(q) = \frac{1}{Q}$ and $f_{\text{irf:dif}}(\Delta)$ are the probabilities of the voxel signal q and inter-voxel signal difference Δ , respectively, for an independent random field of equiprobable signals: $f_{\text{irf:dif}}(\Delta) = \frac{1}{Q}$ if $\Delta = 0$ and $\frac{2(Q-\Delta)}{Q^2}$ otherwise. The factor λ can be omitted (i.e., set to $\lambda = 1$) if only relative interaction energies are computed for the clique families in order to select the most characteristic ones.

In order to account for the tag fading, an extra step beyond the formal MGRF modeling is performed, namely, the Gibbs potentials computed for the found characteristic neighborhood system are adjusted further by weighting them in the temporal direction. The selected empirically linearized weights are used to improve the tag reliability by preferring the information from the earlier tagged images, rather than from the faded ones.

2.4. Energy minimization as a stopping constraint

The tagged CMR image series, \mathbf{g} , under consideration are modified by applying the voxel-wise Iterative Conditional Mode (ICM) relaxation to search for a closest local minimum of the Gibbs energy function for the second-order MGRF appearance model:

$$\hat{\mathbf{g}} = \arg \min_{\mathbf{g}' \in \mathcal{Q}^{xyt}} \left\{ V_{\text{vox}}^T \mathbf{F}(\mathbf{g}') + \sum_{v=1}^N \rho_v V_v^T \mathbf{F}_v(\mathbf{g}') \right\} \quad (5)$$

where the empirical probabilities $\mathbf{F}_{\text{vox}}(\mathbf{g}')$ and $\mathbf{F}_v(\mathbf{g}')$ are collected over each current tagged CMR image series \mathbf{g}' .

To better classify each modified signal as either the tag or the background, the voxel signals are nudged additionally toward their most appropriate grouping by incrementing or decrementing by a small bias δ . The latter was selected experimentally at 1.5% of the maximum gray value, in accord with the discriminant threshold determined from the LCDG sub-models in Section 2.2. The larger the bias, the wider the tag-to-background signal gap in the modified image series. While smaller biases may be useful to optimize gradually the tags-background separation, larger ones sharpen signal gradients and increase the overall spectral noise. Basic steps of the proposed image preprocessing are outlined in Algorithm 1. Our experiments used the following parameters: $K = 2$; $N = 33$; $d_{v:\text{min}} = v - 1$, and $d_{v:\text{max}} = v$.

Algorithm 1. Preprocessing of a time series of tagged CMR images.

Step 1:

i. Find the empirical marginal probability distribution of image signals (gray values) by normalizing the signal histogram collected over a given tagged CMR time series, \mathbf{g} , for the cardiac cycle.

ii. Closely approximate the found distribution using the LCDG model of Eq. (1) in Section 2.2 and partition the latter model into two LCDG sub-models of signals for the tag lines and the background.

iii. Determine the discriminant threshold, τ , ensuring the best separation between the tag and background voxel signals.

Step 2:

iv. Estimate the first- and second-order potentials, \mathcal{V} , and characteristic second-order neighborhoods

$\mathbf{n} = \{\mathbf{n}_v : v = 1, \dots, N\}$ in the MGRF appearance model for the time series \mathbf{g} , using Eq. (4) in Section 2.3.

v. Modify the time series \mathbf{g} with the Iterative Conditional Mode (ICM) algorithm determining the series $\hat{\mathbf{g}}$ corresponding to the closest local energy minimum for the MGRF model.

Step 3 (may be omitted if contrast enhancement is not desired or required in a set of images):

vi. Compare the modified voxel signals, $\hat{\mathbf{g}}(\mathbf{r})$; $\mathbf{r} \in \mathbf{R}$, to the threshold, τ , found at Step 2. iii, and nudge voxels that may be close to the tag line-background boundary by either adding a small fixed bias δ to or subtracting it from the modified signals.

2.5. Validation

The ability of the proposed framework to restore the noise-corrupted strain curves had been experimentally tested and

validated using both synthetic phantoms and *in vivo* data sets in the DICOM format. A commercial version of the HARP (Diagnosoft, Inc., Morrisville, NC, USA) was used to perform the strain analysis.

To evaluate the accuracy of recovering strain measurements, exemplified by strain curves, in all the tests the curves were measured with the HARP software for both the proposed and all the compared techniques. All analyses were done using the automated HARP processing pipeline.

2.5.1. Strain curve recovery

A full cardiac cycle strain curve is an important functional indicator in the cardiac strain analysis. The strain curve is characterized by well-defined cardiac performance indices, such as peak contraction and relaxation rate, and its accurate recovery is useful to measure physiologically meaningful effects of restoring noise-corrupted tagged MR images and the HARP reliability in tracking tag intersection points through the cardiac cycle. After our preprocessing, the CMR images are input to third-party software for band-pass filtering and automatic analysis using the commercial HARP (6–8). The circumferential full cardiac cycle strain curves are then derived from the tagged CMR images, along with the time required to process the data.

For evaluating the strain curve recovery, we defined the absolute total strain error (ATSE) and the absolute strain slope error (ASSE) metrics, measuring the accuracy of a recovered strain curve and of its systolic and diastolic peak slopes, respectively. Let $\mathbf{S}_{\text{obs}} = \{\mathbf{s}_{\text{obs}:i} : i = 1, \dots, m\}$ and $\mathbf{S}_{\text{ini}} = \{\mathbf{s}_{\text{ini}:i} : i = 1, \dots, m\}$ denote a set of m data points, comprising the full cardiac cycle curve of the observed CMR time series to be analyzed, i.e., the noisy or processed observations, and their ground truth known for the initial data set (phantom), respectively. Then the ATSE is defined as follows:

$$\text{ATSE} = 100 \sum_{i=1}^m |\mathbf{s}_{\text{obs}:i} - \mathbf{s}_{\text{ini}:i}| \quad (6)$$

The ATSE specifies the closeness of the observed strain curve to the ground truth.

The ASSE is defined as an absolute relative slope change in the noisy data with respect to the ground truth:

$$\text{ASSE} = 100 \left| \frac{\text{SL}_{\text{obs}}}{\text{SL}_{\text{ini}}} \right| \quad (7)$$

where SL_{obs} and SL_{ini} denote the slope (systolic or diastolic) of the observed data and the ground truth, respectively.

2.6. Synthetic phantoms

The phantoms were constructed in accord with a descriptive mathematical model accounting for physical features of the left ventricle (LV) and physiological LV responses as the heart progresses through the cardiac cycle (31,32). The model describes the LV deformation by mapping locations of each particular material LV point to a corresponding spatial point at a certain time instant during the cardiac cycle. The mapping performs a geometric transformation incorporating shearing, rotation, translation, torsion, and compression of the LV.

Using this transformation, an inverse deformation map is calculated analytically to establish correspondences between two points at any two time instants for simulating tagged




Fig. 4 Simulated slice (a) of the original MR tagged phantom (squares of tag intersection grids localize material points); its corruption (b) with a medium-level Rician noise and a high-level motion, and the preprocessed phantom (c). The Rician noise makes the tag profile across a given grid less uniform.

MR images. In brief, each phantom is built by transforming a generated geometric 3D LV model. Then, an image is formed by selecting an image plane that intersects the LV and assigning every point on the plane a gray value that depends on whether the point lies inside or outside the LV wall.

Fig. 4 depicts an ideal phantom constructed using this model. In addition to computing the deformation, the phantom model also incorporates tag fading such that the tags fade based on the T1 and T2 decay as defined in (32). To simulate different *in vivo* noise scenarios, our phantoms were corrupted using both a Rician distribution (33) and motion filters to simulate scanner noise and motion distortions, respectively.

The Rician distribution describes the actual tagged MR image noise more accurately than other statistical models (34). Three levels of the Rician noise with $\sigma = 1$ and $\nu = 0.15, 0.25, 0.35$ and two motion levels were used: the low level (L) with a 4–6-pixel variable motion filter and the high level (H) with an 8–12-pixel variable motion filter. Motion directions were selected at random. The noise levels were selected to accurately represent typical noise, being observed in acquisitions in different MR imaging environments. The motion levels were selected to represent small (still-laying patient) or larger natural movements during acquisition. To investigate the recovery behavior under the different scenarios, our tests have been performed at the different noise combinations, each test being run 20 times for each noise distribution. In all the tests the preprocessed images were analyzed using the automated third-party HARP software.

2.7. Comparison to other techniques

The proposed framework primarily performs image filtering and does not employ intricate image processing method, such as deformable models, to restore and replace missing and broken tag lines. Therefore, it is most appropriate to compare it against other filters, rather than complex HARP restoration methods that do not function as filters. We examined the performance of the nine filters: our framework with the radius of 0.5 (MGRF3 in the subsequent tables) and radius of 1 (MGRF5); the low-pass filter (35) with the 3×3 (Low3) and 5×5 (Low5) window; the median filter (36) with the 3×3 (Med3) and 5×5 (Med5) window; the Gabor filter (Gabor) (37,38), and the spectral (Spectral) and binary spectral filters (SpecBin) (39,40).

To form the noisy phantoms, the images were corrupted with four levels of independent noise having the Rician

distribution (33) with parameters $\sigma = 1$ and $\nu = 0.15, 0.25, 0.35$. The chosen distribution describes the actual tagged MR image noise more accurately than other statistical models (34). The original phantom corrupted with $\nu = 0.35$ and preprocessing results are shown in Fig. 4.

We also compare the results against the state-of-the-art and publicly available tagged MR tracking, namely the InTag Cardiac MRI tagging analysis toolbox (<https://www.creatis.insa-lyon.fr/inTag/>) (41–43), a commercial plugin for the OsiriX Imaging Software platform. The InTag processing was performed using the software-optimized default parameters and settings.

2.8. Clinical imaging

We examined 20 *in vivo* data sets of the tagged cardiac MR images, acquired by breath hold 1–1 SPAMM imaging in the cardiac short axis plane at basal, mid, and apical levels using a 1.5-T Siemens Espree scanner and phased-array cardiac coil reception with the maximum gradient amplitude of $33 \frac{mT}{m}$ and maximum slew rate of $100 \frac{mT}{s \cdot m}$. An ECG-triggered segmented k -space fast gradient echo sequence was performed with typical grid tag spacing of 10 mm; echo time of 4.0 ms; repetition time of 44.0 ms; flip angle of 14 degrees; voxel size of $1.48 \times 1.48 \times 10 \text{ mm}^3$; bandwidth of $184 \frac{Hz}{\text{voxel}}$; 12–20 cardiac cycle frames, and typical total acquisition time of 15–20 s (the breath held imaging). The ground truth strain curves for examining *in vivo* images were drawn by an expert.

3. Results

3.1. Validation on synthetic phantoms

The phantom’s ground truth strain curves over the cardiac cycle were used to determine the absolute total strain error (ATSE) and absolute strain slope error (ASSE) detailed in Section 2.5.1. As shown in Fig. 5, the well-defined physiological morphology (the solid black line) of the strain curves, including manifesting peak systolic and diastolic slopes, is lost in the presence of noise. The morphology is recovered reliably after our pre-processing, whereas the comparative filters are much less reliable.

The ATSE and ASSE indices quantifying the strain curves morphologies and the recovery of strain slopes, respectively, are summarized in Tables 1 and 2. Our preprocessing produces the best results.

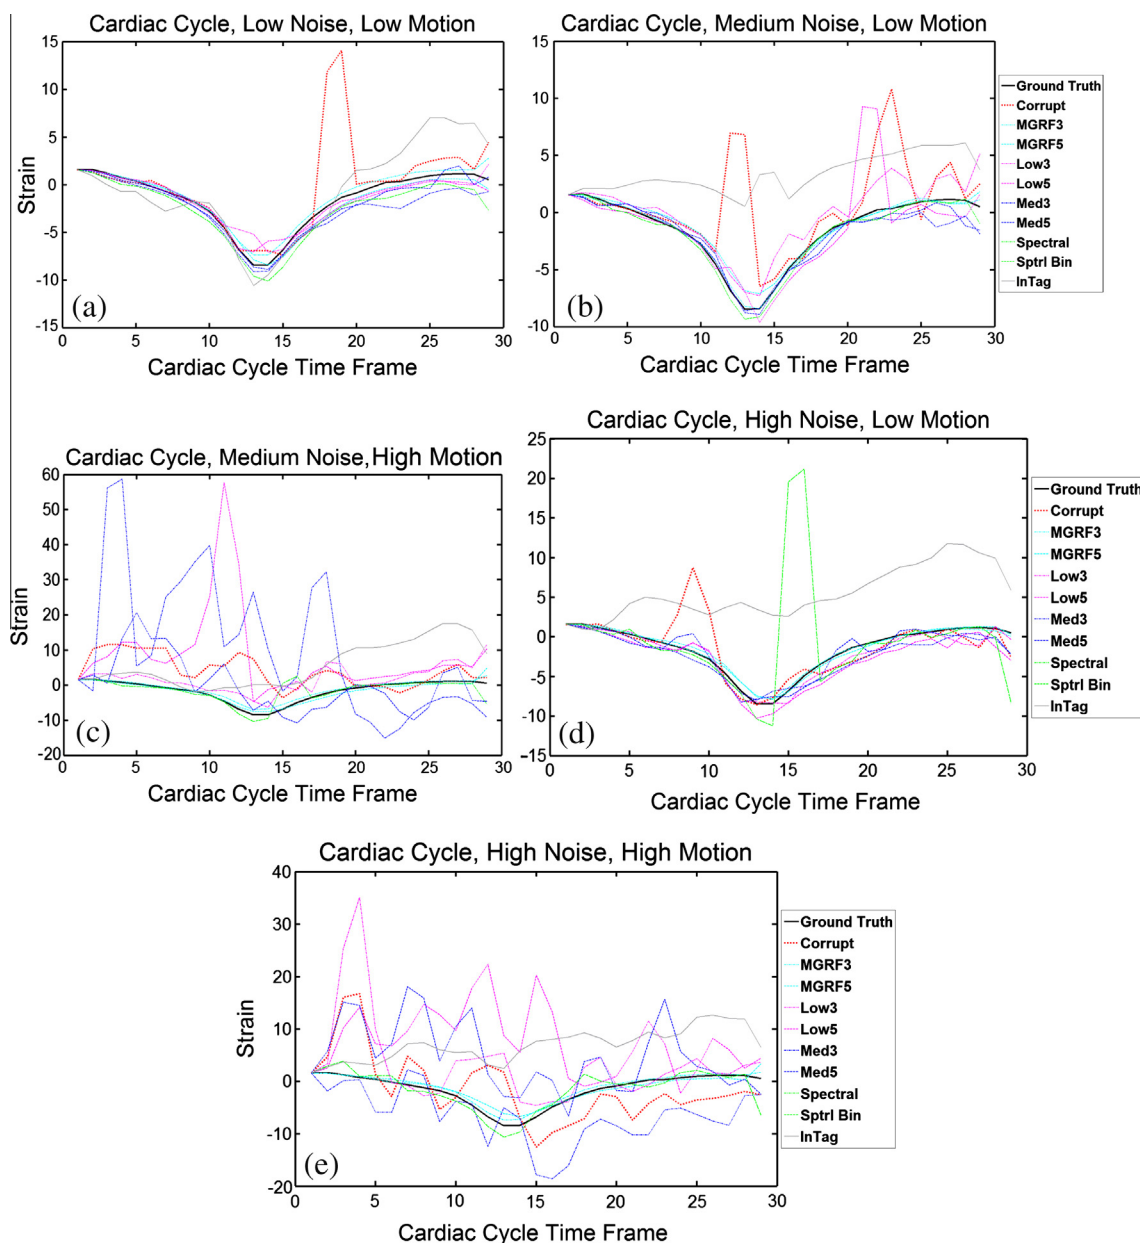


Fig. 5 Full cardiac cycle strain curves for the original and noise-corrupted phantoms, as well as for each comparative filtering and processing method. Note the unreliable slope of the corrupted phantom’s strain curve differing from the original phantom, while after our pre-processing the strain curve is notably closer to the ground truth.

3.2. Validation on *in vivo* data

A ground truth strain curve for examining *in vivo* images was manually constructed by an expert who selected points on tagged MR images over the cardiac cycle. The strain is then computed by tracking these manually placed points. Fig. 6 and Table 3 present both qualitatively and quantitatively the experimental strain curve recovery results for a representative *in vivo* data set (see Fig. 7).

3.3. Computational efficiency

We estimated additional time for our preprocessing in a durational context of manual corrections of mistracked tag

intersection points based on the HARP image analysis. To process an *in vivo* set of 12 frames of 256×256 DICOM images over the cardiac cycle, an expert spent on the average 922 ± 70 sec to correct for residual mistracked points across the 20 sets of the original HARP-processed data. After combining our data pre-processing with the HARP, it required only 334 ± 41 sec to correct for the residual mistracked points, giving approximately 50% (or 5 min) reduction. The average time of our pre-processing for one 256×256 DICOM image was 1.2 s, which amounts to about 14.4 s on the average for a typical data set of 12 CMR temporal frames over the cardiac cycle. Therefore, combining our preprocessing method with HARP increases the net cost efficiency of tracking the myocardial data.

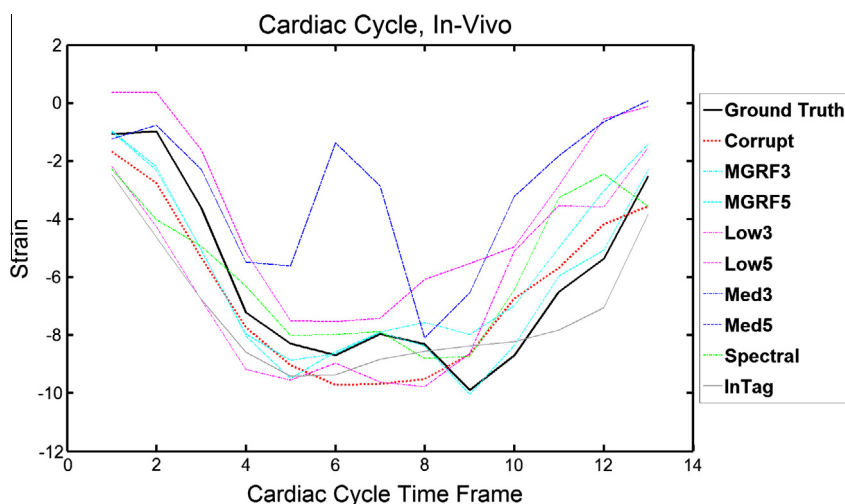


Fig. 6 Full cardiac cycle strain curves representing the expert-determined ground truth, the original data, our preprocessing, and best performing low-pass filtering for the *representative in vivo* data set. The Gabor filter is not shown due to its poor accuracy (its data graph hinders viewing other results). While the difference is not as dramatic as in the phantoms, the proposed preprocessing demonstrates the most accurate recovery of the ground truth strain curve. The ATSE and ASSE in Table 3 support these graphical results quantitatively.

Table 1 The ATSE (the lower, the better) for the corrupted and restored phantoms w.r.t. the ground truth: 0.15...0.35 and L/H indicate the noise level and low/high motion, respectively, and means and standard deviations (st.d.) are computed across all the noise levels for each method.

Absolute total strain error							
Noise level	0.15	0.25L	0.25H	0.35L	0.35H	Mean	St.d.
Corrupted	22	45	120	34	130	71	52
MGRF3	9.2	7.2	11	8.4	11	9.5	1.8
MGRF5	13	12	17	6.9	19	14	4.6
InTag (42)	60	140	210	200	240	170	72
Low3	20	29	47	20	100	43	34
Low5	20	25	230	41	270	120	120
Med3	28	19	380	25	120	120	160
Med5	26	14	170	23	180	83	85
Gabor	310	390	1100	550	1100	690	380
Spectral	24	11	16	23	31	21	7.8
SpecBin	158	90	75	24	59	81	49

4. Discussion

4.1. Strain curve recovery index

In terms of the ATSE and ASSE performance metrics, the proposed preprocessing is consistently more reliable than any other filtering technique that took part in our experimental comparisons. In the cases of lower motion the median and low pass filters provided modest improvements; however, their performance during diastole considerably degraded. The spectral filter performed the best among all other filters selected for these comparisons. In several, primarily low-motion scenarios its performance during the systolic phase was similar to our preprocessing. However, just as all other compared techniques, the spectral filter could not adequately restore the images during the diastolic phase. The Gabor filtering is often incorporated alongside more complex methods (deformable models), but by itself it produces inadequate results.

As compared to the InTag Cardiac MRI tagging analysis toolbox (41–43), our approach performed more consistently.

Similar to the other filters, it met with difficulties in the cases of high noise and motion. But in total the proposed framework has been more reliable in a wide range of noise and motion scenarios. In part it benefits largely from incorporating in the MGRF model the temporal information, which is not used by the conventional filters. Additionally, the conventional filtering treats all the image voxels uniformly, whereas our LCDG-based first-order modeling allows for more accurate voxel classification into a tag line or background to modify their signals accordingly.

In the *in vivo* data, our approach shows the overall best strain curve recovery, including the slopes, albeit less dramatic than in the phantom data.

4.2. Limitations

Although our preprocessing consistently outperforms the conventional filtering in terms of the metrics examined, its performance gets worse with the growing noise level and thus, in principle, should considerably deteriorate under a very high

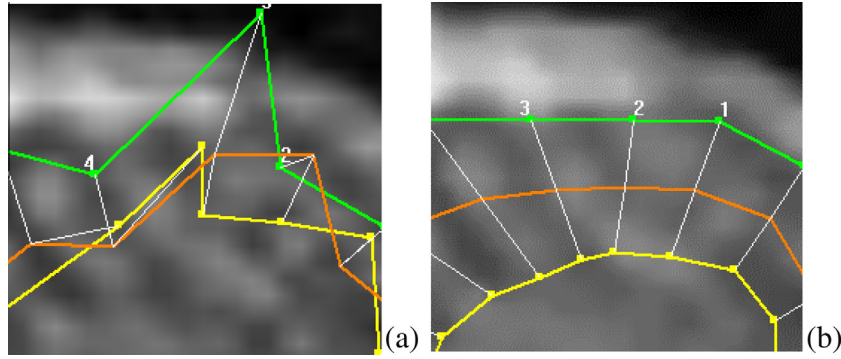


Fig. 7 The HARP tracking screenshots for the *in vivo* data showing an incorrectly tracked point (a) in the original data during late diastole and the same correctly tracked point (b) after the proposed preprocessing. The tracked outer and inner wall and the interpolated mid-wall of the heart are green, yellow, and orange, respectively, and the numbers indicate the points.

Table 2 The ASSE (the lower, the better) for the corrupted and restored phantoms w.r.t. the ground truth: 0.15...0.35; L/H, and s/d indicate the noise level; low/high motion, and systole/diastole, respectively, and means and standard deviations (st.d.) are computed across all the noise levels for each method in both the systole and diastole cases.

Noise level	0.15	0.25L	0.25H	0.35L	0.35H	Mean	St.d.
Corrupted	2.6s	7.8s	30s	10s	26s	16	12
	6.7d	5.9d	15d	6.0d	18d	10	5.8
MGRF3	1.7s	0.9s	3.8s	3.4s	2.5s	2.5	1.2
	1.9d	1.9d	3.8d	0.9d	3.7d	2.4	1.3
MGRF5	1.7s	4.7s	1.1s	2.4s	5.6s	3.1	2
	4.4d	1.4d	4.2d	1.8d	2.5d	2.9	1.4
InTag (42)	4.6s	29s	20s	37s	39s	26	14
	4.8d	33d	36d	40d	60d	35	20
Low3	4.7s	3.7s	8.2s	2.9s	22s	8.3	7.8
	4.5d	7.8d	9.7d	4.7d	7.9d	6.9	2.3
Low5	1.8s	4.1s	110s	3.9s	87s	42	54
	4.1d	7.1d	27d	10d	62d	22	24
Med3	2s	2.1s	120s	2.4s	10s	27	50
	5.1d	2.5d	66d	7.1d	43d	25	28
Med5	2.4s	1.1s	38s	4.4s	44s	18	21
	5.8d	3.9d	13d	6.3d	13d	8.4	4.3
Gabor	77s	74s	230s	110s	100s	120	65
	90d	91d	210d	170d	180d	150	54
Spectral	2.9s	2.6s	3.1s	1.2s	4.8s	2.9	1.3
	6.2d	2.1d	4.0d	6.2d	7.6d	5.2	2.2
SpecBin	11s	18s	6.6s	1.8s	5.8s	8.7	6.3
	28d	6.9d	20d	7.6d	20d	9.2	9.2

noise. Nonetheless, no typical real patient data set in our study demonstrated too high levels of the observed intrinsic noise, suggesting that the proposed framework could be reasonably adequate for clinical applications.

It is also important to note that an improved power spectrum and reduced spectral noise will blur the tags and decrease their contrast and sharpness in the spatial domain. While this trade-off is arguably beneficial for spectral analysis, it reduces visual quality of the images. As mentioned earlier, the proposed technique will also perform more advantageously on the *in vivo* data sets, which are acquired with advanced techniques, such as a high-order SPAMM (spatial modulation of magnetization) or DANTE (delay alternating with nutations for tailored excitation) that may lead to an increased number of multi-dimensional signal peaks.

All the analysis methods, including ours, should consider the acquired tag data quality, e.g., possible broken or missing tag lines in the initial temporal frames. However, our preprocessing cannot recover such missing information.

4.3. MGRF-based processing vs. conventional filtering

The traditional linear low-pass and nonlinear median filtering, along with many other simple moving-window transforms, treats each individual frame of the MRI time series independently and uniformly. Such filtering may lower noise in each frame, but also tends to lose information contained in the tagged image lines due to their excessive frame-wise smoothing. This loss of information affects the accuracy of tracking the tagged line via the MRI time

Table 3 The ATSE and ASSE for the representative *in vivo* data: our preprocessing vs. the conventional filtering (s/d indicate systole/diastole).

Processing	ATSE	ASSE	
Original	16	4.1	4.4
MGRF3	14	2.8	4.4
MGRF5	5.9	2.8	0.6
InTag (42)	19	8.2	2.2
Low3	25	7.6	6.3
Low5	31	4.8	10
Med3	40	1.7	9.1
Med5	40	1.7	9.1
Gabor	260	59	61
Spectral	19	5.6	3.9

series and hence leads to inaccurate estimates of the heart wall strain.

The main advantage of using our translation- and rotation-invariant MGRF model of tagged 3D MR images (tagged lines) is that its characteristic clique families and their potentials are learned analytically from the whole time series of images. As a result, the corrupted MR image is restored without excessive blurring of edges of the tagged line by searching for the closest local minimum of the learned Gibbs energy. Because transforming the corrupted image into the energy minimizer accounts for both the spatial and temporal signal dependencies, it is not necessarily uniform over each individual 2D MR frame.

5. Conclusion

The proposed preprocessing framework was examined in the context of augmenting existing techniques for strain computation from the tagged CMR images. Their spectral analysis, e.g., the HARP, exploiting symmetries in the spectral domain, is arguably of the highest computational efficiency. To the best of our knowledge, our approach is the first to explicitly examine both the corresponding spatial and spectral factors that affect strain computations.

Unlike all the known techniques, our preprocessing is a separate front module, requiring no modification of the subsequent algorithms, such as HARP. In addition to being useful for restoring the HARP grid, it is also applicable to other types of spatiotemporal images that exhibit “grid-like” patterns, e.g., with only horizontal or vertical lines, as well as their combinations.

Conflict of interest

Matthew Nitzken, Garth M. Beache, Marwa Ismail, Georgy Gimel'farb, and Ayman El-Baz declare that they have no conflict of interest.

Ethical approval

This type of study requires no formal consent. This article does not contain any studies with animals performed by any of the authors.

References

- (1) Axel L, Dougherty L. Heart wall motion: improved method of spatial modulation of magnetization for MR imaging. *Radiology* 1989;172(2):349–50.
- (2) Liu X, Prince JL. Shortest path refinement for motion estimation from tagged MR images. *Med Imag, IEEE Trans* 2010;29(8):1560–72.
- (3) Axel L, Dougherty L. MR imaging of motion with spatial modulation of magnetization. *Radiology* 1989;171(3):841–5.
- (4) Ibrahim ESH. Myocardial tagging by Cardiovascular Magnetic Resonance: evolution of techniques—pulse sequences, analysis algorithms, and applications. *J Cardiovasc Magn Reson* 2011;13(1):36.
- (5) Denney TS, Prince JL. Reconstruction of 3-D left ventricular motion from planar tagged cardiac MR images: an estimation theoretic approach. *IEEE Trans Med Imag* 1995;14(4):625–35.
- (6) Osman NF, Kerwin WS, McVeigh ER, Prince JL. Cardiac motion tracking using CINE harmonic phase (HARP) magnetic resonance imaging. *Magn Reson Med: Off J Soc Magn Reson Med/Soc Magn Reson Med* 1999;42(6):1048.
- (7) Osman NF, McVeigh ER, Prince JL. Imaging heart motion using harmonic phase MRI. *IEEE Trans Med Imag* 2000;19(3):186–202.
- (8) Osman NF, Prince JL. Visualizing myocardial function using HARP MRI. *Phys Med Biol* 2000;45(6):1665–82.
- (9) Li W, Yu X. Quantification of myocardial strain at early systole in mouse heart: restoration of undeformed tagging grid with single-point HARP. *J Magn Reson Imag* 2010;32(3):608–14.
- (10) Liu X, Murano E, Stone M, Prince JL. HARP tracking refinement using seeded region growing. In: 4th IEEE international symposium on biomedical imaging: from nano to macro; 2007. p. 372–5.
- (11) Levine WS, Essex-Torcaso C, Stone ML, Murano EZ, Prince JL, Parthasarathy V, et al. A nonlinear incompressible model of the human tongue. *SIAM News* 2006;39(2).
- (12) Liu X, Stone M, Prince JL. Tracking tongue motion in three dimensions using tagged MR image. In: 3rd IEEE international symposium on biomedical imaging: nano to macro; 2006. p. 1372–5.
- (13) Yang X, Murase K. Tagged cardiac MR image segmentation by contrast enhancement and texture analysis. In: 9th International conference on electronic measurement & instruments; 2009. p. 4–210.
- (14) Yang X, Murase K. A multi-scale phase-based optical flow method for motion tracking of left ventricle. In: 4th International conference on bioinformatics and biomedical engineering (iCBBE); 2010. p. 1–4.
- (15) Yuan X, Zhang J, Buckles BP. A multiresolution method for tagline detection and indexing. *IEEE Trans Inf Technol Biomed* 2010;14(2):507–13.
- (16) Li J. Tag line tracking and cardiac motion modeling from tagged MRI. PhD Dissertation Auburn University, Alabama, USA; 2006.
- (17) Khalifa AM, Youssef A, Osman NF. Improved harmonic phase (HARP) method for motion tracking a tagged cardiac MR images. In: 27th Annual international conference of the engineering in medicine and biology society, 2005; 2006. p. 4298–301.
- (18) Cordero-Grande L, Vegas-Sanchez-Ferrero G, Casaseca-de-la-Higuera P, Alberola-Lopez C. Improving harmonic phase imaging by the windowed Fourier transform. In: IEEE international symposium on biomedical imaging: from nano to macro; 2011. p. 520–3.
- (19) Tecelão SR, Zwanenburg JJ, Kuijjer J, Marcus JT. Extended harmonic phase tracking of myocardial motion: improved coverage of myocardium and its effect on strain results. *J Magn Reson Imag* 2006;23(5):682–90.

- (20) Raj A, Singh G, Zabih R. MRF's for MRI's: Bayesian reconstruction of MR images via graph cuts. In: IEEE computer society conference on computer vision and pattern recognition, vol. 1; 2006. p. 1061–8.
- (21) Raj A, Singh G, Zabih R, Kressler B, Wang Y, Schuff N, et al. Bayesian parallel imaging with edge-preserving priors. *Magn Reson Med* 2007;57(1):8–21.
- (22) El-Baz A, Elnakib A, Khalifa F, El-Ghar MA, McClure P, Soliman A, et al. Precise segmentation of 3-D magnetic resonance angiography. *IEEE Trans Biomed Eng* 2012;59(7):2019–29.
- (23) El-Baz A, Gimel'farb G. EM based approximation of empirical distributions with linear combinations of discrete Gaussians. In: IEEE international conference on image processing, vol. 4; 2007. p. IV–373.
- (24) Besag J. On the statistical analysis of dirty pictures. *J Roy Stat Soc: Ser B (Methodol)* 1986;259–302.
- (25) Chellappa R, Jain A. Markov random fields. Theory and application. In: Chellappa Rama, Jain Anil, editors. Boston: Academic Press; 1993. p. 1.
- (26) Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images*. *J Appl Stat* 1993;20(5–6): 25–62.
- (27) Stan ZL, editor. Markov random field modeling in computer vision. Secaucus, NJ, USA: Springer-Verlag New York, Inc.; 1995.
- (28) Poggio T, Torre V, Koch C. Computational vision and regularization theory. *Image Und* 1989;3(1–18):111.
- (29) Chen Y, Amini AA. A MAP framework for tag line detection in SPAMM data using Markov random fields on the B-spline solid. *IEEE Trans Med Imag* 2002;21(9):1110–22.
- (30) Gimel'farb GL. Image textures and Gibbs random fields. Dordrecht: Kluwer Academic; 1999.
- (31) Arts T, Hunter WC, Douglas A, Muijtjens AMM, Reneman RS. Description of the deformation of the left ventricle by a kinematic model. *J Biomech* 1992;25(10):1119–27.
- (32) Waks E, Prince JL, Douglas AS. Cardiac motion simulator for tagged MRI. In: Proceedings of the workshop on mathematical methods in biomedical image analysis; 1996. p. 182–91.
- (33) Rice SO. Mathematical analysis of random noise. *Bell Syst Tech J* 1944;23(3):282–332.
- (34) Nowak RD. Wavelet-based Rician noise removal for magnetic resonance imaging. *IEEE Trans Image Process* 1999;8(10): 1408–19.
- (35) Tomasi C, Manduchi R. Bilateral filtering for gray and color images. In: Sixth international conference on computer vision; 1998. p. 839–46.
- (36) Wang Z, Zhang D. Progressive switching median filter for the removal of impulse noise from highly corrupted images. *IEEE Trans Circ Syst II: Analog Digit Signal Process* 1999;46(1):78–80.
- (37) Qian Z, Montillo A, Metaxas DN, Axel L. Segmenting cardiac MRI tagging lines using Gabor filter banks. In: Proceedings of the 25th annual international conference of the engineering in medicine and biology society, vol. 1; 2003. p. 630–3.
- (38) Weldon TP, Higgins WE, Dunn DF. Efficient Gabor filter design for texture segmentation. *Pattern Recogn* 1996;29(12):2005–15.
- (39) Davis CA, Li J, Denney TS. Analysis of spectral changes and filter design in tagged cardiac MRI. In: 3rd IEEE international symposium on biomedical imaging: nano to macro; 2006. p. 137–40.
- (40) Shynk JJ. Frequency-domain and multirate adaptive filtering. *IEEE Signal Process Mag* 1992;9(1):14–37.
- (41) Arts T, Prinzen FW, Delhaas T, Milles JR, Rossi AC, Clarysse P. Mapping displacement and deformation of the heart with local sine-wave modeling. *IEEE Trans Med Imag* 2010;29(5):1114–23.
- (42) Frouin F, Croisille P, Donal E, Manrique A. Cardiac imaging research group. Results and future works. *IRBM* 2013;34(1):21–3.
- (43) Schaerer J, Casta C, Pousin J, Clarysse P. A dynamic elastic model for segmentation and tracking of the heart in MR image sequences. *Med Image Anal* 2010;14(6):738–49.