



ELSEVIER

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Flow network tracking for spatiotemporal and periodic point matching: Applied to cardiac motion analysis



Nripesh Parajuli^{a,*}, Allen Lu^b, Kevinminh Ta^b, John Stendahl^c, Nabil Boutagy^c,
Imran Alkhalil^c, Melissa Eberle^c, Geng-Shi Jeng^e, Maria Zontak^f, Matthew O'Donnell^e,
Albert J. Sinusas^{c,d}, James S. Duncan^{a,b,d}

^a Department of Electrical Engineering, Yale University, New Haven, CT 06520, USA

^b Department of Biomedical Engineering, Yale University, New Haven, CT 06520, USA

^c Department of Internal Medicine, Yale University, New Haven, CT 06520, USA

^d Department of Radiology & Biomedical Imaging, Yale University, New Haven, CT 06520, USA

^e Department of Bioengineering, Washington University, Seattle 98195, WA, USA

^f College of Computer and Information Science, Northeastern University, Seattle 98195, WA, USA

ARTICLE INFO

Article history:

Received 9 July 2018

Revised 16 February 2019

Accepted 17 April 2019

Available online 18 April 2019

Keywords:

Echocardiography

Motion

Flow network

Neural networks

ABSTRACT

The accurate quantification of left ventricular (LV) deformation/strain shows significant promise for quantitatively assessing cardiac function for use in diagnosis and therapy planning. However, accurate estimation of the displacement of myocardial tissue and hence LV strain has been challenging due to a variety of issues, including those related to deriving tracking tokens from images and following tissue locations over the entire cardiac cycle. In this work, we propose a point matching scheme where correspondences are modeled as flow through a graphical network. Myocardial surface points are set up as nodes in the network and edges define neighborhood relationships temporally. The novelty lies in the constraints that are imposed on the matching scheme, which render the correspondences one-to-one through the entire cardiac cycle, and not just two consecutive frames. The constraints also encourage motion to be cyclic, which an important characteristic of LV motion. We validate our method by applying it to the estimation of quantitative LV displacement and strain estimation using 8 synthetic and 8 open-chested canine 4D echocardiographic image sequences, the latter with sonomicrometric crystals implanted on the LV wall. We were able to achieve excellent tracking accuracy on the synthetic dataset and observed a good correlation with crystal-based strains on the *in-vivo* data.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Cardiovascular diseases (CVDs) were the leading cause of death in the world in 2013 according to the AHA (as stated in the heart disease and stroke statistics - 2017 update Benjamin et al., 2017). Among CVDs, ischemic heart diseases are the most common, which occur as a result of the formation of atherosclerotic plaques in the coronary arteries. This results in the narrowing of the arteries that can lead to under-supply of blood to the left ventricle (myocardial ischemia). A sudden blockage of the arteries, for instance, due to plaque ruptures, can lead to irreversible tissue damage (myocardial infarction) and ultimately heart failure, which can be fatal.

Analysis of the left ventricle (LV), which is the main pumping chamber of the heart, can provide invaluable insights into cardiac

health. An ischemic event in the ventricle is manifest as a reduction of the contractility of the LV wall muscle (myocardium). However, these wall motion abnormalities can be localized to a specific area of the LV and therefore, hard to detect. Global measures of wall motion abnormalities, such as ejection fraction (EF), cannot account for such local changes. The 2013 ACC/AHA guideline on heart failure states that approximately 50% of heart failure cases present themselves with preserved ejection fraction (Yancy et al., 2013). Therefore, it is crucial that a more local and informative measure be developed and adopted.

Visual wall motion scoring is a popular clinical technique for assessing such local deformation behavior and has been shown to be more predictive of clinical outcomes than EF (Galasko et al., 2001; Eek et al., 2010). Stress imaging based wall motion scoring is also of interest and is shown to be of great utility in identifying risk factors associated with mortality and stratifying them (Yao et al., 2003). However, visual wall motion scoring is prone

* Corresponding author.

E-mail address: nripesh.parajuli21@yale.edu (N. Parajuli).

to a high level of uncertainty because it is a qualitative metric, and hence the interobserver variability has been shown to be high in a study involving multiple imaging modalities (Hoffmann et al., 2006). In this context, Lagrangian strain analysis has emerged as a viable method for wall motion quantification and can assist detection and diagnosis of disease, as well as track the therapy, recovery, and management process (Pellerin et al., 2003; Götte et al., 2006).

Among the popular imaging modalities that could be useful for assessing myocardial strain such as magnetic resonance imaging (MRI, e.g., Le Folgoc et al., 2017; Papademetris et al., 2002), computed tomography (CT, e.g., Cury et al., 2008; Sugeng et al., 2006) and echocardiography (echo, e.g., Compas et al., 2014; Heyde et al., 2013), echo is of special interest because of its affordability, higher frame rates compared to MR and CT, and portability. However, the trade-off is that ultrasound imaging is prone to artifacts such as bone shadows, attenuation, signal dropouts, incomplete geometry due to wrong imaging angle and location. These issues call for robust image analysis algorithms and processing pipeline.

While previous efforts have focused on establishing the efficacy of strains and other locally sensitive measures quantifying LV motion (Ledesma-Carbayo et al., 2005; De Craene et al., 2012; Compas et al., 2014), in this work we focus on a novel method aimed at improving this by capturing the myocardial dynamics with a fully spatiotemporal model. A spatiotemporal viewpoint is consistent with the manner in which clinical readings of echo is done – as a movie as opposed to viewing still frames. Most cardiac motion analysis methods perform frame-to-frame displacement estimation and combine the series of deformations to obtain Lagrangian displacements. Uncertainties and ambiguities, that arise at each step in time, get compounded and propagated as frame-to-frame estimations are aggregated. Therefore, significant drift can occur while tracking voxels through the cardiac cycle, particularly past the systolic phase. Another aspect of cardiac motion that is ignored by most methods is the periodicity of the deformation estimates over the cardiac cycle.

Therefore, in this work, we propose a method where the motion model accounts for global spatiotemporal consistency and correspondence as well as periodicity. We build a graphical network where myocardial surface points are set up as nodes and each node is connected to a few other nodes that are its candidate matches in the next time frame, via edges. The edges are associated with weights that capture the likelihood of a particular match. The flow f through the network – a binary-valued variable that captures whether or not a particular match amongst the candidates was chosen – is then solved via optimization and subject to a variety of constraints.

We reported a preliminary version of the flow network tracking (FNT) method in (Parajuli et al., 2017). In that work, we introduced the graphical network model for motion analysis. Here, we expand that model further to get binary-valued flow solution in order to obtain non-overlapping and complete motion trajectories through the entire cardiac cycle. Instead of defining edge relationships by a nearest neighbor search using spatial distance, we do this by feature distance. Furthermore, by introducing additional constraints in the optimization, we are now able to encourage trajectories to be closed-looped and thereby model the periodic aspect of cardiac motion. Also in our previous effort, we used a supervised learning based Siamese network for feature learning. While that performed well with synthetic data, it was not easy to use a similar strategy with *in vivo* data due to a lack of training samples. Therefore, in this work, we use an unsupervised method involving convolutional autoencoders to derive features.

We validate the application of our FNT shape tracking method on 8 synthetic 3D+t ultrasound image sequences developed by Alessandrini et al. (2015) and on 8 open-chested canines imaged at baseline, ischemic and under dobutamine stress (24 studies in

total). Validation is done by comparing with strains obtained from implanted sonomicrometric crystals in the LV. Sonomicrometry derived strains were available for 7 baseline canines and 5 ischemic and stressed canines. We perform correlation analysis to compare echocardiography based strains and crystal based strains.

2. Related work

2.1. Speckle/image-based tracking methods

2.1.1. Non-rigid registration

Non-rigid registration methods typically consist of a model where the motion (displacement) field is parametrized by smooth functions such as B-splines. Ledesma-Carbayo et al. applied it to 3D ultrasound sequences in a frame-to-frame manner (Ledesma-Carbayo et al., 2005). Heyde et al. proposed a 3D deformation model where the LV image is transformed from Cartesian coordinate to an anatomical LV shaped coordinate system (Heyde et al., 2012). However, most methods of this class do not use a fully spatiotemporal motion model. The optimizations are highly non-convex and are prone to get stuck in local minimums that can yield non-optimal solutions.

Some work has been done towards addressing the spatiotemporal alignment issue. Ledesma-Carbayo et al. proposed a 3d+t B-spline spatiotemporal model which parameterized the Lagrangian motion of a point at end diastole (ED) through the cardiac cycle (Ledesma-Carbayo et al., 2005). However, their model does not explicitly capture any notion of global spatiotemporal correspondence and consistency. This is because their cost function accounts for the dissimilarity with the ED frame but not with any other frames, including the neighboring frames. Also, this method does not capture large deformations effectively. De Craene et al. proposed a 3D+t diffeomorphic map based registration method, where a B-spline parameterization over the velocity field is used (De Craene et al., 2012). While this explicitly models the notion of spatiotemporal correspondence, the concern with velocity-based parametrization, in general, is that it is prone to error accumulation as Lagrangian displacements are calculated by integrating the velocities.

2.1.2. Block matching

Block matching involves taking a patch of image in a frame and searching for the best match in a spatial window in the next time frame. Langeland et al. implemented this on 2D echo RF (radio frequency) images (Langeland et al., 2005). Lubinski et al. also implemented this on RF images and further refined displacement estimation in the beam direction using zero-crossings of the phase of the complex correlation function (Lubinski et al., 1999). Boukerroui et al. implemented block-matching on 2D B-Mode images for velocity estimation (Boukerroui et al., 2003). A maximum likelihood based similarity function was used as distance metric and smoothing was implicitly applied by encouraging velocity vectors of neighboring pixels to be similar. These methods can be time-consuming due to a large search space and also lack a regularization term in general.

2.1.3. Optical flow

Optical flow methods assume that the intensity of a point in a moving image is consistent across time and that motion is responsible for temporal intensity variation. Song and Leahy applied this to model cardiac motion in 3D CT images (Song and Leahy, 1991). Optical flow methods suffer from the fact that cardiac ultrasound speckle patterns do not display such consistency in intensity across frames and also that they are significantly more noisy than CT or MR images.

2.2. Shape matching/tracking methods

Shape-based methods try to match shape/image descriptors derived from a point set. Pre-processing is necessary to generate the points either by simple edge/feature detection algorithms or by a more sophisticated segmentation algorithm. Post-processing is also generally required for smoothing and dense field generation as the solutions are sparse (Papademetris et al., 2002).

2.2.1. Frame-to-frame matching

Chui et al. proposed a point matching algorithm that modeled deformation using non-rigid thin plate spline parameterization and used it to align point sets derived from brain imaging in (Chui and Rangarajan, 2003). The correspondences that map point sets are fuzzy (non-binary) initially and are refined iteratively to obtain one-to-one binary correspondences. Belongie et al. introduced the shape context feature, which is more global than the local curvature, and used it to match point sets (Belongie et al., 2000). They solve a weighted bipartite graph matching problem using the Hungarian algorithm to obtain one-to-one correspondences (Belongie et al., 2000).

2.2.2. Temporal tracking

We previously proposed a method that tracks individual points located at myocardial surfaces through time in (Parajuli et al., 2016). Points on the myocardial surfaces form nodes in a graph and edges exist between points and their spatial neighbors in the next time frame. The motion of an individual point is then posed as the shortest path through this graph. Our current work improves this by modeling the motion of all points on the surface together as opposed to individually. Berclaz et al. (2011) used a flow network structure to build a fully spatiotemporal model for an object tracking problem. We expand upon their work, as will be seen below, by providing a probabilistic mechanism of outlier handling and accounting for periodic motion. Furthermore, while their work handles the uncertainty in the nodes of the graph, we handle uncertainty in the edges to solve for correspondences.

2.3. Post hoc regularization models

Methods lacking in inherent regularization, or producing a sparse set of displacements as our method does, rely on post hoc regularization of the initial tracking results to produce smooth displacement fields. Papademetris et al. first estimated initial correspondences between myocardial surfaces using a shape matching approach (Papademetris et al., 2002). The initial estimation was then regularized by using a biomechanically inspired finite element method approach.

Compas et al. proposed the use of radial basis functions to generate smooth and dense displacements from the integration of sparse sets of shape and speckle tracking displacements in (Compas et al., 2014). We expanded this strategy to impose further smoothness and biomechanical constraints on the displacement fields in (Parajuli et al., 2015). Lu et al. learned how to regularize noisy motion by training a neural network to filter noisy 4D Lagrangian displacement vector fields (Lu et al., 2017).

3. Methods

Many point matching methods that model rigid or non-rigid deformations try to achieve one-to-one or symmetric matches (Belongie et al., 2000; Chui and Rangarajan, 2003). Even within the free-form deformation framework, diffeomorphic transformations are a popular choice, which captures a similar notion. The resulting displacement fields from such mappings are more realistic and robust to noise and artifacts. However, matching algorithms leading

Table 1
Notations.

T	Number of frames
$N(t)$	Number of points per frame, $t \in [1: T]$
x_i^t	i th point of frame t , $i \in [1: N(t)]$
e_{ij}^t	Edge from point i in frame t to point j in frame $t + 1$ $i \in [1: N(t)]$ $j \in [1: N(t + 1)]$
w_{ij}^t	Weight associated with e_{ij}^t
f_{ij}^t	Flow through e_{ij}^t
$\eta(t, i)$	Indices of points in the neighborhood of x_i^t in frame $t + 1$

to one-to-one correspondence at the frame level cannot guarantee that the composition of those frame level one-to-one correspondences leads to globally optimal one-to-one correspondence across all frames. Thus, in the work presented in this paper below, we develop an approach that (i.) incorporates global one-to-one correspondence for all points being tracked and (ii.) accounts for the cyclical nature of cardiac motion.

3.1. Constrained flow optimization

First, point clouds are obtained by uniformly sampling the endocardial and epicardial surfaces (see Fig. 2). The sequence of point clouds, through the cardiac cycle, is then set up as nodes in a graph with directed edges between points and their match candidates in the next time frame. Each node is endowed with a feature vector representing local appearance characteristic and the match candidates are chosen based on feature distances between nodes and their spatial neighbors. This is illustrated in Fig. 3.

The edges capture particle (tissue) motion possibilities, and their weights capture the likelihood of the motion. We have T time frames in total with $N(t)$ ($t \in [1: T]$) points per frame. Each node is defined as x_i^t ($i \in [1: N(t)]$), where an edge e_{ij}^t exists between x_i^t at time t and its neighbor x_j^{t+1} (based on feature distances) at time $t + 1$ ($i \in [1: N(t)]$ and $j \in [1: N(t + 1)]$). The flow through an edge e_{ij}^t in this network is captured by the binary-valued variable f_{ij}^t , and the corresponding edge weight is w_{ij}^t . $f_{ij}^t = 1$ implies that the point x_i^t and x_j^{t+1} are a match (Table 1).

We would like to solve for flow f that is proportional to the edge weights w . This amounts to maximizing the inner product $w^T f$, subject to the following constraints at each node x_i^t ($\eta(t, i)$ indexes neighbors of x_i^t in frame $t + 1$):

1. Flow is either 0 or 1:

$$\forall t, i \quad f \in \{0, 1\}. \quad (1)$$

2. Sum of outgoing flows is less than or equal to one (C_{out} , see Fig. 4a).

$$\forall t, i \quad \sum_{j \in \eta(t, i)} f_{ij}^t \leq 1 \quad (2)$$

3. Sum of outgoing and incoming flows is equal (C_{bal} , see Fig. 4b).

$$\forall t, i \quad \sum_{h: i \in \eta(t-1, h)} f_{hi}^{t-1} = \sum_{j \in \eta(t, i)} f_{ij}^t \quad (3)$$

C_{out} and C_{bal} are constraint matrices that are later explicitly used in solving an optimization (see Eq. (4)). C_{out} ensures that total flow is preserved from one frame to another. By itself, it would result in complete trajectories that traverse through the best possible path but without any consideration for spatial consistency with other trajectories. C_{bal} is enforced so that incoming and outgoing




Fig. 1. Overall method outline.




Fig. 2. Preprocessing steps for *in vivo* data to get endocardial and epicardial surface points.




Fig. 3. Nodes, edges and other relationships in the network. The point sets are sampled from the myocardial surface sequence at each frame.




Fig. 4. A node and different edges/flows visualized.

flows through nodes are equal. This is helpful in preventing many-to-one correspondences, which can create uncharacteristic stretching and shrinking in the displacement fields. One-to-many correspondences are also avoided for the same reason.

Even though we would like to solve an integer programming (IP) optimization to get a binary-valued solution for f , we start with a linear programming (LP) relaxation.

The constraint that $f \in \{0, 1\}$ is dropped for a looser constraint that $f \geq 0$. We solve the following optimization instead (all variables are considered in their vectorized forms here for simplicity):

$$\begin{aligned} & \text{Maximize} && w'f \\ & \text{subject to} && f \geq 0, \quad C_{out}f \leq 1, \quad C_{bal}f \leq 0 \end{aligned} \quad (4)$$

Despite the relaxation of an IP into an LP, we still obtain a binary-valued solution. This is of great consequence because solving IPs directly is an NP-hard problem. Due to the special nature of the constraint matrices that define our LP, our solutions remain integer-valued. This builds upon work that proves that inequalities in LP with constraint matrices that satisfy the total unimodularity property, and that have an integral right-hand side, lead to an integral solution (Hoffman and Kruskal, 2010). A unimodular matrix is a square integer matrix with determinant $+1$ or -1 . This results in the inverse of these matrices also being an integer matrix. Berclaz et al. show that such flow balance constraint matrices satisfy total unimodularity in (Berclaz et al., 2011). This allows us to obtain fully connected non-intersecting trajectories throughout the cardiac cycle.

It is instructive to think of the effects of the constraints in terms of an equivalent graph-based matching/tracking methods. For instance, only imposing C_{out} results in shortest path tracking for each point. In addition to C_{out} , if we also impose C_{in} , such that the sum of incoming flows at a node is also constrained ($C_{in}f \leq 1$), we achieve a maximum bipartite match, with one-to-one correspondence, between point sets in consecutive time frames. Finally, if C_{out} and C_{bal} are both enforced, we obtain a complete maximum bipartite match, which results in one-to-one correspondence that extends through the entire cardiac cycle (C_{in} gets implicitly imposed in this case). The effects of these constraints on the resulting trajectories are illustrated in Fig. 5(a)–(c).

Since our constraint matrices are very sparse, the LP can be solved very efficiently. We used the CVX package (in MATLAB), for specifying the LP and other optimizations in this work (CVX Research, 2012; Grant and Boyd, 2008). CVX was used in conjunction with the MOSEK solver for solving the LP, which uses the interior-point method (ApS, 2014).

3.1.1. Constraint matrices

At a node associated with x_i^t , let e_1, e_2, e_3 be the incoming edges and let e_{11}, e_{12}, e_{13} be the outgoing edges (in vectorized form). Fig. 6 illustrates this relationship and the form of the equation corresponding to a row of the C_{bal} matrix appears as follows:




Fig. 5. Outcomes of applying different constraints on 1D+t point sets, with points stacked vertically in space and horizontally in time.

$$\begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ \vdots \\ f_{11} \\ f_{12} \\ f_{13} \end{bmatrix} \leq 0. \quad (5)$$

Similarly, C_{out} has 1's where C_{bal} also has 1 and zeros elsewhere. Conversely, C_{in} has 1's where C_{bal} has -1 .

Fig. 5 displays how imposing C_{out} , C_{out} and C_{in} , C_{out} and C_{bal} affects tracking outcomes in a toy 1D+t problem. The points are stacked vertically for each time step. We can clearly see how just imposing C_{out} leads to the qualitatively worst result in Fig. 5(a). This is because there is nothing preventing two trajectories from merging and occupying the same nodes.

We can also notice, in Fig. 5(c), as C_{bal} is imposed strictly, one-to-one point correspondence leads to non-overlapping complete trajectories. However, nodes on the top areas in Fig. 5(c) have no trajectories passing through. In an earlier implementation of this algorithm, we attempted to remedy situations like this by relaxing this strict balance constraint (Parajuli et al., 2017). In addition to that, we also added spatiotemporal smoothness constraints. However, in this work, we impose C_{bal} strictly and develop a different strategy of regularization suitable to cardiac motion. We discuss this strategy next.

3.1.2. Imposing loop constraints

A simple extension to the algorithm described so far encourages trajectories to end up in close proximity to their start location. This is done by adding one more set of edges (and constraints) between the points in the last frame and the first frame (see Fig. 7.). First, we impose the following constraint between the flow from source node x_{src} (let the flow be $f_i^{x_{src}}$) and the flow from the first frame to second (f_{ij}^1):

$$\forall i, i \in [1 : N(1)], \quad f_i^{x_{src}} = \sum_{j \in \eta(1,i)} f_{ij}^1. \quad (6)$$

Here, since each node in frame 1 is connected directly to x^{src} , we are saying that if there is a flow into that node from the source, there has to be a flow out of the node going into frame 2 as well. Next, we impose a balance between flow from the last frame to the first frame (f_{hi}^T via the loop edges) and flow from first to second (f_{ij}^1):

$$\forall i, i \in [1 : N(1)], \quad \sum_{h: i \in \eta(T,h)} f_{hi}^T = \sum_{j \in \eta(1,i)} f_{ij}^1. \quad (7)$$

This is the key flow balance constraint that encourages trajectories to be closed-looped. Since edges always exist between nodes that are close spatially, and we already have a mechanism for obtaining non-overlapping complete trajectories, this helps us ensure that our trajectories are roughly closed-looped. Finally, to make sure that the total flow leaving the source node is also equal to the total flow leaving the last frame (and going to the first

implicitly), we also impose the following constraint:

$$\sum_{i=1}^{N(1)} f_i^{x_{src}} = \sum_{i=1}^{N(1)} \sum_{h:i \in \eta(T,h)} f_{hi}^T \quad (8)$$

We shall call these constraints C_{loop} , which is applied in the same way as C_{bal} and is incorporated within it during the optimization (see Eq. (4)). In Fig. 5(d), we can see how the loop constraint helps us recover trajectories that lead to the end points of the trajectories being in close proximity of their start points, which is what we expect due to the periodicity of the LV motion/displacement.

3.1.3. Outlier handling

Because we are maximizing the flow through the network, our algorithm always solves for the maximum number of possible trajectories, even if they are of very low quality. Fig. 8 displays such an example. The ideal match would have been if point 1 was matched to point 4 and point 3 to 6, but instead, the opposite has happened. We tackle this by probabilistic thresholding. All edges with weights below a certain threshold P_{th} are omitted. An example is shown in Fig. 9 with randomly generated points for illustration.

3.1.4. Edge weight calculation.

Each edge e_{ij}^t has the following weight:

$$w_{ij}^t = \exp\left(\frac{-\|x_i^t - x_j^{t+1}\|^2}{2\sigma_x^2(t)}\right) \exp\left(\frac{-\|F(x_i^t) - F(x_j^{t+1})\|^2}{2\sigma_f^2(t)}\right) \quad (9)$$

The first component of the weight (Eq. (9)) captures the likelihood that an edge gets selected based on spatial distance. The second component captures the same likelihood, but based on feature similarity. The two effects are assumed to impact the likelihood independently, hence the two components are simply multiplied. F can be any shape or appearance-based feature associated with x_i^t and x_j^{t+1} . $\sigma_x(t)$ and $\sigma_f(t)$ act as normalization constants and are the standard deviations of Euclidean and feature distances, respectively, for each image frame. For every point in a frame, once the NK best neighbors have been established, Euclidean distances between the neighbors are calculated along with the feature distances. All these $N(t) \times NK$ values are aggregated and their standard deviations calculated to obtain $\sigma_x(t)$ and $\sigma_f(t)$.

3.2. Neural network based appearance features

Since we do not have ground truth trajectories for *in vivo* data, we simply use a convolutional autoencoder to learn appearance features here, unlike our previous work in (Parajuli et al., 2017). Convolutional autoencoders are known to produce a state-of-the-art solution in unsupervised learning problems. The structure of the network is shown in Fig. 10 and has the standard encoder-decoder format. We use Euclidean distance to quantify the similarity between two learned embeddings. The learned embedding F is used to calculate edge weights in Eq. (9).

3.3. Dense field generation and strains

We use radial basis functions (RBFs) as interpolants and calculate dense displacement fields, which is necessary for generating dense motion trajectories and Lagrangian strain which helps us assess myocardial function. This is based on our previous work and that of (Compas et al., 2014; Parajuli et al., 2015).

The most important parameter of the interpolation process was the placement resolution of the RBFs. This was set at 10 mm for level 1 and 5 mm at level 2 for our 2-stage coarse-to-fine interpolation process. The decrease in spacing by a factor of 2 was held




Fig. 6. Incoming and outgoing edges at a node.




Fig. 7. A simple flow network displaying the additional loop edges between the last frame and the first (not all shown) which helps us obtain closed-looped trajectories. The source node and edges are also shown.





Fig. 8. An unlikely, but a valid scenario where balance and closed loop constraints are satisfied but the result is poor qualitatively.



(a) Trajectories computed from random $1D + t$ data. (b) Trajectories computed from random $1D + t$ data after thresholding.

Fig. 9. Outcome of applying thresholding (of $P_{th} = 0.3$) on the edge weights.

constant but different starting values - 10 mm, 8 mm, 12 mm - were experimented with. Also, all the data points per frame ($N(t)$) were not used for computation, since that would be computationally very intensive. A finer grid of resolution 1.75 mm was used to sample the original displacement values into representative data-points. Resolution of 1.00 mm, 1.50 mm, 1.75 mm and 2.00 mm were experimented with.




Fig. 10. Autoencoder network: A low dimensional embedding of image patches that captures images statistics is learned.

These parameters, and the regularization constant described below, were chosen by conducting a parameter sweep on the Normal data and then the ideal set of parameters were used for all data.

We regularize our motion field (U) by imposing sparsity on the weights associated with the basis functions representing our motion field that account for tissue incompressibility. We do so under the assumption that the cardiac tissue is roughly incompressible and therefore, the motion vector field should be roughly divergence-free (i.e., $\nabla \cdot U = 0$) (Song and Leahy, 1991). Finally, we also mildly penalize the norm of the spatial derivatives (∇U) to discourage jumps and discontinuities in the motion fields. We use a compactly supported basis function, which results in a sparse basis matrix unlike other popular choices such as Gaussian and thin plate spline based RBFs. Sparse matrices are more conducive to numerical optimization (Compas et al., 2014; Wendland, 1995).

Once dense displacements are obtained, Lagrangian strains were calculated. For the synthetic data, the ground truth mesh data also contained information on each mesh point, regarding which neighbor of it was oriented in which cardiac coordinate (e.g., endocardial-epicardial pairs defined the radial direction). The ratio of distances between these directionally oriented neighboring pairs, as they were propagated across time using the dense displacement fields, were used to directly calculate Lagrangian radial, circumferential and longitudinal strains (Alessandrini et al., 2015).

For, *in vivo* data, first, dense displacements were composed through the cardiac cycle to obtain dense Lagrangian displacements. Then derivatives of the (U , V , W) components of the displacements were calculated across the regular (x , y , z) Cartesian axes. Then a 3×3 matrix consisting of the strain components in e_{xx} , e_{yy} , e_{zz} , e_{xy} , e_{yz} , e_{zx} directions are calculated using formulas given in (Yan et al., 2007). These are projected to the local cardiac geometry, obtained from the segmentation mask, to obtain radial, circumferential and longitudinal strains.

Excluding image segmentation, which takes about 2 h including manual first frame initialization, but is not the focus of this work, the overall method takes about 20 min. The time consuming portions are the dense displacement field generation where large matrix multiplications are being done by breaking them into chunks, as opposed to at once, and the neural network based feature generation, which were run on the CPU and not the GPU due to memory restrictions. All of this was run on a 8-core, 32gb RAM machine. Most of the code was in MATLAB and the neural network portion in Python.

4. Experiments and results

We explored how changes in key parameters affect performance via a synthetic dataset where ground truth was available. Although the synthetic data do not fully capture all the variation present




Fig. 11. Synthetic data image example with endocardial and epicardial contours (normal data).

in *in vivo* echocardiographic images, it helped us establish best practices. The autoencoder-derived features are generated using models trained in a leave-one-out fashion where, for every synthetic data sequence, a different model was trained without using that sequence. For the *in vivo* data, we obtained endocardial and epicardial surfaces using an automated (except the first frame) level-set segmentation method (Huang et al., 2014). In the synthetic case, segmentation masks were generated from the ground truth mesh masks, hence a segmentation step was not required during pre-processing. Other than that, the steps are the same for the *in vivo* data as well. As a reminder, the overall method is depicted in Fig. 1.


4.1. Synthetic data

4.1.1. Data description

We used 8 synthetic 3D+t ultrasound image sequences developed by (Alessandrini et al., 2015). The dataset consisted of 3 categories of image sequences. First consisted of just one normal sequence (normal). Second consisted of ischemic sequences with ischemia in the distal and proximal left anterior descending artery (LADDIST and LADPROX), right circumflex artery (RCA) and left circumflex artery (LCX). The third consisted of dilated myocardium sequences - one synchronous (SYNC) and two dyssynchronous, induced by left bundle branch block (LBBB and LBBBSMALL). Examples of the synthetic data - 2D slices from one image sequence at end diastole and corresponding contours - are given in Fig. 11. Image volumes were of size $224 \times 176 \times 208$ voxels with a resolution of $0.67 \text{ mm} \times 0.85 \text{ mm} \times 0.58 \text{ mm}$. Sequences were of length 34 except for LBBB and LBBBSMALL, which were of length 40, and captured 1 cardiac cycle. During computation, volumes were sampled such that voxels were of $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$ for computational efficiency.

Table 2
Parameters that were tuned for the FNT algorithm.

Name	Description
NK	Number of nearest neighbors (by feature distance) in consideration for next frame.
Z_{fr}	Number of slices sampled in the long (z) axis per frame.
θ_{fr}	Angular sampling along the circumference. (roughly along the short axis.)
P_{th}	Probabilistic threshold for outlier edges removal.



(a) Along the circumference of the surfaces, θ_{fr} points are uniformly sampled in terms of angle.

(b) Along the long (z) axis of the surface, Z_{fr} points are uniformly sampled.

Fig. 12. Sampling scheme.

We report tracking errors on 2250 myocardial mesh points, for which the positions through the sequence were provided as ground truth. These points were evenly distributed in the endocardium, epicardium and in the mid-wall. We calculated distances between ground truth mesh points and mesh points from our tracking algorithm by propagating the known first frame mesh points through time. The errors are summarized using median and interquartile range (IQR). Overall errors, aggregated over all time frames, are displayed separately from end diastolic (last frame) and ES errors because the number of data points is different. Analyzing errors at ES and ED is important as many clinical readings are made at these time points.

4.1.2. Parameter selection for FNT

We vary Z_{fr} , θ_{fr} and NK , which control the level of sampling of our surface masks and graph neighbor assignment (see Table 2 for description). We adopted a cylindrical sampling strategy where we sample uniformly along the z axis and along the circumference (see Fig. 12). We present the median tracking errors (MTE) under different combination of these parameters on the normal data in Table 3. The combination of $Z_{fr} = 40$, $\theta_{fr} = 30$ and $NK = 3$ provided the lowest overall MTE on the normal data.

A special case of our algorithm is when $NK = 1$. In this scenario, the algorithm essentially becomes a simple nearest neighbor assignment problem without any optimization. We also check the performance for this condition in Table 4. For normal data, even though the overall error, in the best case, when $Z_{fr} = 50$, $\theta_{fr} = 15$, is not significantly higher than for when $NK = 3$ or $NK = 5$ (0.94 ± 0.77 compared to 0.86 ± 0.65), the same does not apply to the entire dataset. The overall error for all dataset aggregated is 1.25 ± 1.45 , magnitude of which is about 50% higher than obtained with $NK = 3$. Performance is significantly worse for end di-

Table 3
Result of changing Z_{fr} , θ_{fr} and NK on the normal data and median tracking errors (MTE).

Z_{fr}	θ_{fr}	NK	Overall/mm	ES/mm	ED/mm
30	15	3	0.96 ± 0.74	1.13 ± 0.75	1.01 ± 0.82
30	30	3	0.94 ± 0.71	1.25 ± 0.87	0.77 ± 0.56
30	40	3	0.95 ± 0.73	1.30 ± 0.87	0.82 ± 0.53
30	15	5	0.90 ± 0.71	1.09 ± 0.68	0.90 ± 0.79
30	30	5	0.97 ± 0.73	1.24 ± 0.83	0.98 ± 0.77
30	40	5	0.93 ± 0.70	1.20 ± 0.80	0.90 ± 0.65
40	15	3	0.90 ± 0.68	1.17 ± 0.76	0.78 ± 0.66
40	30	3	0.86 ± 0.65	1.09 ± 0.73	0.82 ± 0.58
40	40	3	0.90 ± 0.66	1.21 ± 0.73	0.77 ± 0.53
40	15	5	0.91 ± 0.66	1.11 ± 0.71	0.87 ± 0.64
40	30	5	0.95 ± 0.74	1.24 ± 0.75	0.87 ± 0.75
40	40	5	0.95 ± 0.71	1.16 ± 0.73	0.92 ± 0.79
50	15	3	0.96 ± 0.75	1.15 ± 0.72	1.02 ± 0.84
50	30	3	0.93 ± 0.74	1.21 ± 0.79	0.86 ± 0.78
50	40	3	0.91 ± 0.72	1.09 ± 0.72	0.80 ± 0.72
50	15	5	0.90 ± 0.67	1.13 ± 0.82	0.86 ± 0.59
50	30	5	0.88 ± 0.73	1.31 ± 0.75	0.75 ± 0.67
50	40	5	0.91 ± 0.75	1.34 ± 0.80	0.83 ± 0.66

Table 4
Result of changing Z_{fr} , θ_{fr} for the special case $NK = 1$ on the normal data and median tracking errors (MTE).

Z_{fr}	θ_{fr}	NK	Overall/mm	ES/mm	ED/mm
40	10	1	1.21 ± 1.03	1.21 ± 0.99	1.62 ± 1.43
40	15	1	1.11 ± 0.91	1.01 ± 0.94	1.60 ± 1.38
40	30	1	1.08 ± 0.88	0.94 ± 0.76	1.61 ± 1.36
40	40	1	1.06 ± 0.87	0.94 ± 0.83	1.58 ± 1.39
50	10	1	1.05 ± 0.85	1.04 ± 0.97	1.49 ± 1.25
50	15	1	0.94 ± 0.77	0.86 ± 0.70	1.39 ± 1.18
50	30	1	1.00 ± 0.86	0.93 ± 0.78	1.51 ± 1.40
50	40	1	0.97 ± 0.81	0.91 ± 0.73	1.47 ± 1.29
60	10	1	1.09 ± 0.90	1.08 ± 0.97	1.47 ± 1.19
60	15	1	0.96 ± 0.79	0.91 ± 0.84	1.32 ± 1.16
60	30	1	1.03 ± 0.82	0.97 ± 0.88	1.49 ± 1.22
60	40	1	1.03 ± 0.85	0.94 ± 0.85	1.48 ± 1.23

Table 5
Outcome of changing P_{th} on the normal data and MTE.


P_{th}	Overall/mm	ES/mm	ED/mm
0.1	0.92 ± 0.75	1.27 ± 0.79	0.83 ± 0.55
0.3	0.90 ± 0.70	1.24 ± 0.74	0.80 ± 0.49
0.5	0.87 ± 0.63	1.23 ± 0.74	0.81 ± 0.51

astole with the overall error magnitude about 100% higher than with $NK = 3$. All of this makes sense since motion is particularly hard to track post end systole and in pathology cases, this is even more pronounced.

We examined how performance changed as we changed P_{th} (see Table 2 for description). Each edge in our flow network is associated with a probabilistic weight that represents how likely that edge transition is. Dropping edges whose weights are below P_{th} from consideration is a method of handling outliers. The results are shown in Table 5. MTE values were typically the lowest for $P_{th} = 0.5$. With high P_{th} , low-quality edges get removed and the resulting trajectories are better probabilistically.


4.1.3. Effect of different constraints

Next, we see how the algorithm performs under different constraints. We start by applying C_{out} only, then C_{in} , C_{bal} and C_{loop} in an incremental fashion. C_{out} does not enforce any one-to-one correspondence constraint. It is equivalent to tracking all points independently using a shortest path formulation. C_{out} and C_{in} together enforce one-to-one correspondence at a frame-to-frame level. C_{out}



(a) Overall MTE for all data.

(b) MTE for ES and ED for all data.

Fig. 13. MTE for all data, for different constraint setting. C_{in} , C_{bal} and C_{loop} were added incrementally.

(a) Overall MTE for all data.

(b) MTE for ES and ED for all data.

Fig. 14. MTE for all data, comparing different features. The same tracking method (FNT) was used for all of these.

Table 6

Result of changing constraints applied to the optimization and MTE.

Constraint configuration	Overall/mm	ES/mm	ED/mm
C_{out} only	1.22 ± 1.15	1.50 ± 1.19	1.24 ± 1.25
C_{out} and C_{in}	1.17 ± 1.04	1.39 ± 1.16	1.19 ± 1.02
C_{out} and C_{bal}	1.09 ± 0.92	1.24 ± 0.93	1.23 ± 1.03
C_{out} , C_{bal} and C_{loop}	0.84 ± 0.68	1.13 ± 0.73	0.72 ± 0.56

and C_{bal} together enforce one-to-one correspondence throughout the cardiac cycle. Finally, applying C_{loop} in addition to C_{out} and C_{bal} also enforces a balance between first and last frames, thereby encouraging trajectories to start and end at nearby positions. The findings are summarized in Table 6 and Fig. 13.

Not surprisingly, there is an improvement with each additional constraint. The jump from having no loop constraint to the loop constraint is especially significant. This validates our intuition that accounting for the cyclical nature of cardiac motion is necessary.

4.1.4. Comparing features

We also compared the performance of the learned (unsupervised) feature against other features (and/or metric). A comparison to the shape context (Belongie et al., 2000) feature, gradient histogram feature and intensity cross-correlation metric-based approaches is shown in Table 7 and Fig. 14.

The learned feature using an Autoencoder provides better tracking results in comparison to other features generated using shape and image information. It was fascinating to find that the cross-correlation of intensity patches also performed relatively well. This

is likely because speckle de-correlation is probably not substantial across time in this synthetic dataset.

4.1.5. Comparing methods

We compared the FNT method against other point tracking methods. Table 8 and Fig. 15 summarize the findings. For shape context matching (Belongie et al., 2000), we had to use a lower sampling rate ($Z_{fr} = 20$, $\theta_{fr} = 10$) because the algorithm took over an hour to run per frame and therefore was not tractable at a higher sampling rate. The Dynamic Shape Tracking (DST) algorithm is run with the same setting as the FNT algorithm ($Z_{fr} = 40$, $\theta_{fr} = 30$ and $NK = 3$). A free-form deformation (FFD, Rueckert et al., 1999) implementation available from the Bioimagesuite package (Papademetris et al., 2005) was also applied to our data for reference. FFD was applied by registering all frames to the next frame in the sequence (FFD fr-to-fr) and also by registering all frames to the first frame (FFD fr-to-ED). Control point spacing of $13 \text{ mm} \times 13 \text{ mm} \times 11 \text{ mm}$ was used for FFD. Affine registration was applied before the non-rigid part.

We first note that the shape context matching and FFD applied frame-to-frame seem to provide the worst results overall. Since these methods did not have any temporal aspect to them, their tracking suffers a lot post systole, which is evident from the high ED errors. The DST method performs only slightly better as it does not track all points together but provides improved tracking post systole and therefore results in lower ED errors. The frame-to-ED FFD, where images are registered directly to the first frame and FNT seem to provide similar results overall. But FNT is better at ES. This is because the deformation from ED to ES is large and the FFD algorithm was most likely not able to account for that. Good

Table 7
Average MTE for different feature generation methods.

Feature	Overall/mm	ES/mm	ED/mm
Shape context (Belongie et al., 2000)	1.19 ± 0.99	1.41 ± 1.15	1.19 ± 0.91
Gradient histograms	1.20 ± 1.03	1.47 ± 1.16	1.18 ± 1.00
Intensity Cross-correlation	0.98 ± 0.81	1.24 ± 0.91	0.86 ± 0.72
Learned feature (Autoencoder)	0.84 ± 0.68	1.13 ± 0.73	0.72 ± 0.56

Table 8
MTE for different tracking methods.

Method	Overall/mm	ES/mm	ED/mm
Shape Context Matching	1.60 ± 1.22	2.03 ± 1.35	1.52 ± 1.01
DST	1.22 ± 1.11	1.54 ± 1.27	1.21 ± 1.12
FFD fr-to-fr	1.30 ± 1.05	1.41 ± 0.94	1.64 ± 1.28
FNT	0.84 ± 0.68	1.13 ± 0.73	0.72 ± 0.56
FFD fr-to-ED	0.83 ± 0.80	1.37 ± 1.08	0.56 ± 0.50

Table 9
Median absolute difference of Lagrangian strains: Comparing FNT and ground truth.

Data group	Radial (%)	Circumferential (%)	Longitudinal (%)
Normal	2.19 ± 2.94	1.37 ± 2.02	0.16 ± 0.21
Ischemic	2.39 ± 3.58	1.55 ± 2.24	0.18 ± 0.27
Dilated	6.93 ± 9.69	3.79 ± 6.83	0.65 ± 1.12
Overall	3.42 ± 5.85	2.12 ± 3.43	0.27 ± 0.53

performance at ES is crucial since peak strains typically occur around ES and are widely evaluated and reported clinically as a measure of function.

4.1.6. Regional strain analysis

Since we are ultimately interested detecting regional (within the LV) changes in strain values in order to isolate areas with injury, we also test whether or not the strain values we calculate can help us do this. Radial strain curves, obtained using the FNT method and ground truth positions, for the normal dataset are shown for basal, mid and apical areas of the LV in Fig. 16. Basal and mid areas are divided into 6 sectors - anterior (Ant), antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) sectors. Apical area is divided into 4 sectors - anterior (Ant), septal (Sept), inferior (Inf) and lateral (Lat) sectors. For the normal data, we can see that strain values do not differ significantly across different sectors. Compared to the ground truth, our method seems to underestimate peak radial strain values as evident from Fig. 16. However, the broader trends seem to be consistent with the ground truth strain values.


We show circumferential and longitudinal strains for the Normal data in Figs. 17 and 18 respectively. Similar trends are observed here as well. There is no underestimation or overestimation apparent in these curves, consistent with the values reported in Table 10.

In Fig. 19, radial strain curves for the LADPROX data are shown (FNT and ground truth position based). There appears to be a significant reduction of strain values around the infero-septal and

infero-lateral sectors in the basal and mid regions. In the apical region, the anterior and lateral region RCA strains are also lower. Fig. 20 shows radial strain curves for the RCA data (FNT and ground truth position based). Infero-lateral and inferior strain values are reduced in basal and mid regions. There appears to be no significant changes in the apical region. Similar to the normal data, there also appears to be an underestimation of strains for both the ischemic data. However, there is also a broad agreement on the strain trends.

It should be clear from Figs. 16, 19 and 20 that, we are able to discern changes in strains across different regions of the LV. For instance, from normal to RCA, there is barely any change in the strain values in the apical area, whereas there is a significant reduction in inferior and infero-lateral strain values in both the basal and mid regions. From normal to LADPROX, there is a significant change in the apical strains. The observation that different injury profiles lead to different regional strain patterns, as demonstrated here, is of great value.

To quantify this more stringently, we directly compared strains obtained from FNT and the ground truth positions by analyzing the differences in strain values and cross-correlations. Table 9 summarizes absolute difference in radial, circumferential and longitudinal strain values for normal, Ischemic (LADPROX, LADDIST and RCA) and Dilated (LBBB, LBBB SMALL and SYNC) data groups. While the median absolute difference in strain values are within reasonable range for the normal and ischemic data, it is rather large for the dilated data group. This is perhaps indicative of the fact that we



(a) Overall MTE for all data.

(b) MTE for ES and ED for all data.

Fig. 15. MTE for all data, comparing different methods.





Fig. 16. Radial strain curves in the basal, mid and apical area of the LV for the normal Leuven data (our method and ground truth). Curves indicating mean strains for anterior (Ant) antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) regions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)






Fig. 17. Circumferential strain curves in the basal, mid and apical area of the LV for the normal Leuven data (our method and ground truth). Curves indicating mean strains for anterior (Ant) antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) regions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

are not able to account for their complex motion patterns, which are different than that of the normal and ischemic group.


To understand whether these differences in strains were random or systematic, we also summarized median differences (not absolute) in Table 10. This helps us identify the direction in which the algorithm is biased. First thing of note is that, overall, there is no substantial bias in the circumferential and longitudinal strains. There is bias in all strain types for the dilated data group. Radial strains seem to be underestimated for all data groups, which is

consistent with our findings from earlier as we observed the strain curves.

A reason behind this systematic bias in the radial strains is our use of segmented surfaces. Since points were constrained to move between surfaces, the maximum radial displacement was constrained. At a very small spatial scale, if we assume surfaces are flat (planes in 3D) and consecutive surfaces are parallel, the maximum possible radial motion is fixed - the projection of the normal vector between the two surfaces along the radial




(a) Normal data longitudinal strains (FNT)




(b) Normal data longitudinal strains (ground truth)

Fig. 18. Longitudinal strain curves in the basal, mid and apical area of the LV for the normal Leuven data (our method and ground truth). Curves indicating mean strains for anterior (Ant) antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) regions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



(a) LADPROX data radial strains (FNT)



(b) LADPROX data radial strains (Ground Truth)

Fig. 19. Radial strain curves in the basal, mid and apical area of the LV for the LADPROX Leuven data (our method and ground truth). Curves indicating mean strains for anterior (Ant) antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) regions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

direction. However, because point correspondences are not perfect during optimization, there is noise in the displacement vectors. As these noisy vectors are regularized and smoothed, the final radial displacements are lower on aggregate than the ground truth. Such constraints do not exist circumferentially or longitudinally as long as points are sampled densely and regularly.

Finally, we compare how well the trends agree with ground truth for the three strain types in Table 11 by looking at the summary of correlation values of individual sector strain curves. Again, the dilated data seem to have the worst correlations, which is con-

sistent with the findings so far. Longitudinal strains are slightly worse and noisier overall. This is partly due to the fact that we use a sparse set of displacements located on the myocardial surfaces. Longitudinal motion is hard to quantify towards the basal and apical regions in this setting. Overall, the results with synthetic data were good both in terms of point tracking and strain analysis.

In Fig. 21 we also display radial strain maps at ES calculated using our method and the ground truth for 3 ischemic datasets. Areas with injuries have low strains and this can be seen in the strain maps. Again, this is to illustrate that we could reliably




Fig. 20. Radial strain curves in the basal, mid and apical area of the LV for the RCA Leuven data (our method and ground truth). Curves indicating mean strains for anterior (Ant) antero-septal (Ant-Sept), infero-septal (Inf-Sept), inferior (Inf), infero-lateral (Inf-Lat) and antero-lateral (Ant-Lat) regions are shown. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 10
Median difference in Lagrangian strains: Comparing FNT and ground truth.

Data group	Radial (%)	Circumferential (%)	Longitudinal (%)
Normal	2.13 ± 3.56	0.00 ± 2.75	0.06 ± 0.33
Ischemic	1.99 ± 4.19	0.00 ± 3.10	0.04 ± 0.36
Dilated	5.18 ± 10.79	-1.33 ± 7.69	-0.03 ± 1.29
Overall	2.79 ± 6.13	-0.20 ± 4.54	0.02 ± 0.54

Table 11
Median correlations of Lagrangian strains: Comparing FNT and ground truth.

Data group	Radial	Circumferential	Longitudinal
Normal	0.99 ± 0.02	0.96 ± 0.05	0.98 ± 0.04
Ischemic	0.98 ± 0.04	0.96 ± 0.07	0.96 ± 0.15
Dilated	0.87 ± 0.16	0.74 ± 0.35	0.42 ± 0.43
Overall	0.96 ± 0.10	0.93 ± 0.17	0.90 ± 0.56

localize injuries since the maps from our method and ground truth appear fairly similar.

4.2. In vivo data

4.2.1. Echocardiography data

We applied our method and explored physiological variations in the heart by analyzing *in vivo* canine 4D (3D+t) echocardiography (4DE) data from 8 canine studies. The imaging was done on anesthetized open chested animals with a transducer suspended in a water bath. The animals were imaged in baseline condition (BL), an ischemic condition (HO), induced by occluding the left anterior descending artery (LAD), and stress condition (HODOB), induced by injecting dobutamine in low dosage: 5 µg/kg/min. These conditions were tested due to our interest in ultimately developing a 4DE strategy for rest-stress imaging/analysis. Echocardiographic images were available for all 8 studies in all conditions but sonomicrometry data was available for 7 studies in BL and 5 studies in HO and HODOB (Table 12).




Fig. 21. Epicardial surfaces displaying radial strains for three different types of ischemia, induced by occlusion at left anterior descending artery (LADPROX), right coronary artery (RCA) and left circumflex artery (LCX).

Table 12
Different physiological conditions of imaging (*in vivo* data).

Condition	Description
BL	Baseline.
HO	High LAD occlusion.
HODOB	High LAD occlusion with low dobutamine stress.

Philips iE33 ultrasound system (Philips Medical Systems, Andover, MA), with the X7-2 phase array transducer and a hardware attachment that provided RF data, were used for






Fig. 22. Example of *in vivo* images and segmentation contours in one image sequence. I_1 , I_2 and I_3 are three images in the systolic cycle.



(a) Crystals arranged in 3 cuboidal lattices.

(b) Crystal alignment in the LV.

Fig. 23. Crystals and their relative position in the LV.



(a) A single tetrahedron

(b) An edge before (left) and after deformation (right)

Fig. 24. Components of the elementary tetrahedron.

acquisition. Imaging frequency ranged from 50–60 fps, which typically gave us 20–30 volumes per 4D image sequence. Typical volume dimensions were $400 \times 120 \times 110$ voxels of resolution $0.26 \text{ mm} \times 1.10 \text{ mm} \times 1.10 \text{ mm}$. Images were captured along with ECG gating and capture a single cardiac cycle. All experiments were conducted in compliance with the Institutional Animal Care and Use Committee policies.

Once these images were acquired, they were segmented using a semi-automated scheme. Endocardial and Epicardial surfaces were manually traced for the first frame of all images (see Fig. 22). Then we used a dictionary learning based level set algorithm to propagate these surfaces through the cardiac cycle (Huang et al., 2014). The FNT algorithm was then applied to these data with some adjustments. Since the extent of LV captured by imaging is different for different sequence in the long axis, Z_{fr} is set as $Z_{fr} = \max(25, \text{total } z \text{ slices available})$. θ_{fr} is then set as $\theta_{fr} = Z_{fr}/1.3$ since 1.3 was the ratio of the best (Z_{fr}, θ_{fr}) combination for the synthetic data - (40, 30). We used the unsupervised learning derived feature here as well, by training an autoencoder with *in vivo* data. Other parameters were the same - $NK = 3$, $P_{th} = 0.5$. Radial basis function based interpolation method was used post FNT tracking and Lagrangian strains were calculated based on the techniques outlined in (Yan et al., 2007).

We compared these strains with the ones obtained from sonomicrometric crystals implanted close to the mid-anterior LV wall during the same imaging studies as described above. Even though the strain values calculated from the sonomicrometric crystals cannot be considered to be ground truth since the crystal-based strain calculation method itself is prone to noise and uncertainties, they do provide a helpful validation benchmark. We focused on analyzing if the trends were consistent using correlation analysis. We describe the sonomicrometric crystal processing and calculations next.

4.2.2. Sonomicrometer data

We used sonomicrometric transducer crystals (crystals), recording instrument and processing software *SonoSoft* and *SonoView* (Sonometrics Corporation, London, Ontario, Canada) to acquire signal from crystals and process them. We implanted 19 crystals in the heart, 16 in the targeted areas and 3 in the basal and apical areas for reference. 3 more were placed in the transducer to align the crystals in Echocardiography's LV coordinate.

The 16 crystals were arranged in such a way that they formed three adjacent cuboidal structures. The three cubes were roughly located in: (i) the ischemic region of LV (ISC), which was caused by the aforementioned LAD occlusion (ii) the remote region (away




Fig. 25. FNT and crystal strains in BL condition for a data across the 3 cubic regions (ISC, BOR and REM top to bottom) for 1 dataset. Radial (red), circumferential (cyan) and longitudinal (green) strains from left to right. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 13
Cubical areas of crystal placement.

Area	Description
ISC	Ischemic area.
BOR	Borderline area between ischemic and remote.
REM	Remote area.

Table 14
Comparing FNT and FFD strain correlations with sonomicrometric crystal strains (median and IQR) by physiological conditions.

Method	BL	HO	HODOB
FNT	0.91 ± 0.19	0.80 ± 0.51	0.85 ± 0.28
FFD	0.90 ± 0.30	0.76 ± 0.37	0.82 ± 0.60

from the ischemic area) and (iii) the border region between the two, as shown in Fig. 23 (also Table 13).

We adapted the 2D sonomicrometry-based strain calculation method outlined in (Waldman et al., 1985) for 3D. We considered a tetrahedral element to be the elementary 3D unit to obtain strain information from. Each cube was composed of over 50 different possible tetrahedrons. We calculated strain tensors for each cube separately and averaged them to obtain a strain tensor for the overall cube. Fig. 24(a) shows a single tetrahedral unit and Fig. 24(b) shows representations of initial and post deformation widths of an edge of the tetrahedron.

We calculated radial, circumferential and longitudinal strains by projecting onto the cardiac geometry defined using the apical and basal crystals. We could then compare these strains with echocardiography (echo) based strains after aligning them in the LV coordinate system using rigid registration.

4.3. Comparing FNT and FFD

We compared the strains obtained from the FNT and FFD (the FFD fr-to-ED version specifically), since they were the best performing methods amongst the ones we tested, with the strains obtained from sonomicrometry analysis, across a few categories -

imaging conditions (BL, HO, HODOB), ischemic regions (ISC, BOR, REM) and strain types. BL data was available for 8 volumes and HO and HODOB data were available for 5 volumes.

The median correlations (with IQR), of FNT and FFD with the crystal strains across BL, HO and HODOB, are given in Table 14. BL data was available for 8 volumes, hence there were 72 total correlation data points for this - radial, circumferential and longitudinal strains for each cube giving 9 correlations per data. There were 45 data points for the other two conditions as there were 5 sequences each (with 9 correlations per data similarly).

The correlations were the best for the BL condition and somewhat lower for HO and HODOB, with HO being the worst. Since the magnitude of deformation is typically the smallest in HO, perhaps the signal-to-noise ratio of the displacement data pertaining to HO was lower, thereby resulting in lower median correlation.

FFD and FNT strain correlations with the crystal strains, in the 3 cubic regions of interest, are given in Table 15. 5 data with BL, HO and HODOB sequences, and 3 more with only BL resulted in 18 sequences in total for consideration, giving us 54 correlation data points. Table 15 suggests that there was very little difference in performance, between FFD and FNT, across the 3 cubic regions. The ISC area correlations were slightly lower.




Fig. 26. FNT and crystal strains in HO condition for a data across the 3 cubic regions (ISC, BOR and REM top to bottom) for 1 dataset. Radial (red), circumferential (cyan) and longitudinal (green) strains from left to right. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 15

Comparing FNT and FFD strain correlations with sonomicrometric crystal strains (median and IQR) by cubic regions.

Method	ISC	BOR	REM
FNT	0.83 ± 0.38	0.87 ± 0.26	0.86 ± 0.22
FFD	0.84 ± 0.35	0.83 ± 0.33	0.84 ± 0.69

Table 16

Comparing FNT and FFD strain correlations with sonomicrometric crystal strains (median and IQR) by strain type.

Method	Radial	Circumferential	Longitudinal
FNT	0.93 ± 0.08	0.90 ± 0.14	0.54 ± 0.90
FFD	0.91 ± 0.11	0.88 ± 0.30	0.51 ± 0.77

Finally, comparison of FFD and FNT strain correlations, in 3 cardiac axes - radial, circumferential and longitudinal, is given in Table 16. 54 data points were available for each strain type (3 - ISC, BOR and REM for each image sequence). The FFD and FNT correlations appeared to be roughly the same for all three categories. Longitudinal strain correlations were the lowest for both methods. As images were acquired such that the beams are parallel to the short axis, full volume images were created by stitching smaller slices in the long axis. This can lead to stronger speckle decorrelation in the long axis and therefore low quality longitudinal strains.

Overall, the two methods provided similar trends and similar strain correlation values were observed as a result.

4.3.1. Changes across physiological conditions

We continued rest of our analysis by just using the FNT method. First, we explored how strains change from BL to HO and then to HODOB. In BL condition, we expected normal heart function and strain values. During HO, which induces ischemia, we expected an overall decrease in the heart function and strain magnitudes. From HO to HODOB, we expected to be a recovery of function in non-ischemic areas, whereas the outcome in ischemic area was uncertain. The ischemic (ISC) and non-ischemic areas were defined by the location of the crystals that defined the 3 cubic regions ISC, BOR and REM (see 23). Figs. 25–27 show strains for BL, HO and HODOB conditions respectively for 1 dataset.

There was a decrease in overall strains, across all regions, going from BL (Fig. 25) to HO (Fig. 26). There was also recovery in all regions going from HO (Fig. 26) to HODOB (Fig. 26). Although we expected to find a discernible difference between the level of strain changes across ISC, BOR and REM, particularly between ISC and REM, these strain curves do not suggest any obvious differences.

Fig. 28 shows the changes in peak radial, circumferential and longitudinal strains from BL to HO to HODOB (for 8 BL data and 5 HO and HODOB data). Median values across different groups are shown, along with the IQR. The trend of decrease in peak strains from BL to HO and then the subsequent increase from HO to HODOB was strongly shown by radial strains. The corresponding radial strains can be found in Table 17 as well. The change in HO to HODOB is more substantial for both crystal and FNT-based strains. There is a subtle difference in the magnitude of increase, from BL to HO, and decrease, from HO to HODOB, of radial strains across ISC, BOR and REM regions for the crystal-based strains. Such a difference was not found for the echo-based strains.




Fig. 27. FNT and crystal strains in HODOB condition for a data across the 3 cubic regions (ISC, BOR and REM top to bottom) for 1 dataset. Radial (red), circumferential (cyan) and longitudinal (green) strains from left to right. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 17

Median of peak **radial** strains for data across BL, HO and HODOB, also broken down by regions - ISC, BOR and REM. See Fig. 28 for pictorial representation.

Method	BL (%)			HO (%)			HODOB (%)		
	ISC	BOR	REM	ISC	BOR	REM	ISC	BOR	REM
Crys	12.3 ± 9.7			11.6 ± 7.1			30.2 ± 15.0		
Echo	13.7 ± 7.3			11.5 ± 4.4			22.0 ± 16.1		
Crys	12.3	11.8	20.6	8.9	11.6	16.2	29.5	23.8	34.6
Echo	16.0	15.7	13.0	12.0	11.3	12.0	22.0	25.0	19.5

Table 18

Median of peak **circumferential** strains for data across BL, HO and HODOB, also broken down by regions - ISC, BOR and REM. See Fig. 28 for pictorial representation.

Method	BL (%)			HO (%)			HODOB (%)		
	ISC	BOR	REM	ISC	BOR	REM	ISC	BOR	REM
Crys	-11.2 ± 2.3			-10.2 ± 7.1			-16.4 ± 3.7		
Echo	-7.0 ± 2.8			-8.3 ± 3.4			-10.8 ± 3.1		
Crys	-12.8	-10.0	-10.7	-12.1	-12.1	-10.4	-19.3	-14.5	-16.4
Echo	-7.4	-6.8	-7.0	-8.3	-8.7	-7.4	-10.8	-9.6	-12.1

Circumferential strain values across the regions are reported in Table 18. Both crystal-based and echo-based (FNT) circumferential strains did not display changes in increase/decrease magnitudes across ISC, BOR and REM (see Fig. 28 and Table 18). Changes in longitudinal strains are also shown in Fig. 28 and values reported in Table 19. The crystal-based strains display strong increase, from BL to HO, and decrease, from HO to HODOB, in strain magnitudes respectively. The echo strain magnitudes seemed very small and therefore not very informative since the IQRs were fairly high. Just like radial strains, the crystal-based longitudinal strains also point

to the existence of a subtle difference in the magnitude of increase from BL to HO, and decrease from HO to HODOB, across ISC, BOR and REM regions.

A source of uncertainty in the crystal-based analysis was the challenge involved in registering crystals with LV. While the positions of the transducer were available (from the reference crystals in the transducer), there still remained the task of rotational alignment that required some manual intervention. Furthermore, the transducer position crystals, which served as references, were themselves subject to noise and uncertainties. The ischemia that




Fig. 28. Peak strain bar graphs (with median and IQR) for radial (top), circumferential (middle) and longitudinal (bottom) strains at BL, HO and HODOB - shown across ISC, BOR and REM regions for echo and crystal-based strains.

Table 19

Median of peak **longitudinal** strains for data across BL, HO and HODOB, also broken down by regions - ISC, BOR and REM. See Fig. 28 for pictorial representation.

Method	BL (%)			HO (%)			HODOB (%)		
	ISC	REM	BOR	ISC	REM	BOR	ISC	REM	BOR
Crys	-9.1 ± 2.93			-3.2 ± 2.9			-11.6 ± 6.9		
Echo	-5.1 ± 3.2			-3.0 ± 5.2			-4.2 ± 3.0		
Crys	-8.1	-9.8	-9.4	-2.6	-5.4	-4.9	-11.3	-14.1	-13.3
Echo	-5.1	-4.9	-5.7	-3.0	-2.6	-5.8	-3.9	-4.7	-3.3

we induced was also possibly not severe enough to cause highly localized functional difference. Therefore, we also conducted analysis by regions defined by the local cardiac geometry (as defined by the AHA standard (Cerqueira et al., 2002)). For Lagrangian strain calculation, the axes were defined based on the segmentation mask of the first frame of the echo sequences. Fig. 29 displays radial, circumferential and longitudinal strains across BL, HO and HODOB conditions across anterior (Ant), antero-lateral (Ant-Lat), infero-lateral (Inf-Lat), inferior (Inf), infero-septal (Inf-Sept) and antero-septal (Ant-Sept) sectors. These plots are for the same data as in the earlier plots comparing echo and crystal strains across cubic regions through BL, HO and HODOB.

Similar to the crystal-based analysis, we were able to detect global changes from BL to HO and HO to HODOB. However, local variations were difficult to detect. If we look at radial strains in Fig. 29, the anterior and antero-septal regions seem to have the highest strains in HO and HODOB, while we expected the opposite. Since the ischemic region was close to the mid-anterior sector, we expected lower strains along this area. Summary of peak radial

strains across different sectors at BL, HO and HODOB can be found in Table 20.

Overall, the results of *in vivo* testing were also highly encouraging with room to expand and improve for enhanced local sensitivity.

5. Conclusion

While spatiotemporal methods do exist for full volume tracking in 4DE (Ledesma-Carbayo et al., 2005; De Craene et al., 2012), we have developed, to the best of our knowledge, one of the first fully spatiotemporal 4D cardiac surface tracking method. This was then used to generate dense displacements using a regularized radial basis function framework. Our method was also able to account for the cyclical nature of the cardiac motion. Strains calculated from these displacements were then used to analyze the changes in global and local deformation pattern of the LV. We also proposed an unsupervised neural network-based feature generation method for motion tracking.




Fig. 29. Radial, circumferential and longitudinal strains across 6 sectors along the mid-LV slice. BL, HO and HODOB conditions show expected behavior. The sectors are - anterior (Ant), antero-lateral (Ant-Lat), infero-lateral (Inf-Lat), inferior (Inf), infero-septal (Inf-Sept) and antero-septal (Ant-Sept).

Table 20

Median of peak **radial** strain percentages using echo across the 6 standard cardiac sectors (at the mid-LV area) at BL, HO and HODOB.

Condition	Ant	Ant-lat	Inf-lat	Inf	Inf-sept	Ant-Sept	Overall
BL	26.7	15.9	14.3	13.4	18.0	18.7	17.9 ± 12.4
HO	22.7	16.4	11.3	9.8	14.9	21.1	15.4 ± 11.5
HODOB	37.4	17.4	18.6	31.3	34.9	39.4	28.0 ± 25.9

We obtained very good tracking accuracy with our method and features on the Leuven synthetic dataset. Even though we only tracked endocardial and epicardial surfaces, we obtained comparable results to the FFD method, which accounts for motion throughout the myocardium. We were also able to detect local regional changes in strain patterns and obtain different strain profiles for different infarct locations and extent.

In the future, the tracking method can be expanded to model dense myocardial displacements as well. This would have the advantage of not requiring perfect segmentations. A rough region-of-interest would suffice as long as it contains the LV. Since doing this would make the problem more ill-posed, the motion model would also have to be expanded to incorporate more constraints. The neural network-based feature generation strategy can be expanded further to develop supervised or transfer learning based methods applicable to *in vivo* data as well. For instance, the Siamese neural

network based approach we explored in our previous work can be extensively leveraged with more training data (Parajuli et al., 2017).

Conflict of interest

None.

Acknowledgements

We are immensely thankful for the efforts of many past and present members of Dr. Albert Sinusas group that were involved in the image acquisitions. We would also like to thank Dr. Hemant Tagare and Dr. Lawrence Staib of the Image Processing and Analysis Group at Yale for many fruitful discussions. This work was supported in part by the National Institute of Health (NIH) grant number R01HL121226.

References

- Alessandrini, M., De Craene, M., Bernard, O., Giffard-Roisin, S., Allain, P., Waechter-Stehle, I., Weese, J., Saloux, E., Delingette, H., Sermesant, M., et al., 2015. A pipeline for the generation of realistic 3d synthetic echocardiographic sequences: methodology and open-access database. *IEEE Trans. Med. Imaging* 34 (7), 1436–1451.
- ApS, M., 2014. The MOSEK optimization toolbox for MATLAB manual. Version 7.0. <http://docs.mosek.com/7.0/toolbox/index.html>.
- Belongie, S., Malik, J., Puzicha, J., 2000. Shape context: A new descriptor for shape matching and object recognition. In: *NIPS*, vol. 2, p. 3.
- Benjamin, E.J., Blaha, M.J., Chiuve, S.E., Cushman, M., Das, S.R., Deo, R., de Ferranti, S.D., Floyd, J., Fornage, M., Gillespie, C., et al., 2017. Heart disease and stroke statistics - 2017 update: a report from the american heart association. *Circulation* 135 (10), e146–e603.
- Berclaz, J., Fleuret, F., Turetken, E., Fua, P., 2011. Multiple object tracking using k-shortest paths optimization. *IEEE Trans. Pattern Anal. Mach. Intell.* 33 (9), 1806–1819.
- Boukerroui, D., Noble, J.A., Brady, M., 2003. Velocity estimation in ultrasound images: a block matching approach. In: *Biennial International Conference on Information Processing in Medical Imaging*. Springer, pp. 586–598.
- Cerqueira, M.D., Weissman, N.J., Dilsizian, V., Jacobs, A.K., Kaul, S., Laskey, W.K., Pennell, D.J., Rummerger, J.A., Ryan, T., Verani, M.S., et al., 2002. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Circulation* 105 (4), 539–542.
- Chui, H., Rangarajan, A., 2003. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Understanding* 89 (2), 114–141.
- Compas, C.B., Wong, E.Y., Huang, X., Sampath, S., Lin, B.A., Pal, P., Papademetris, X., Thiele, K., Dione, D.P., Stacy, M., et al., 2014. Radial basis functions for combining shape and speckle tracking in 4d echocardiography. *IEEE Trans. Med. Imaging* 33 (6), 1275–1289.
- Cury, R.C., Nieman, K., Shapiro, M.D., Butler, J., Nomura, C.H., Ferencik, M., Hoffmann, U., Abbara, S., Jassal, D.S., Yasuda, T., et al., 2008. Comprehensive assessment of myocardial perfusion defects, regional wall motion, and left ventricular function by using 64-section multidetector ct. *Radiology* 248 (2), 466–475.
- CVX Research, I., 2012. CVX: Matlab software for disciplined convex programming, version 2.0. <http://cvxr.com/cvx>.
- De Craene, M., Piella, G., Camara, O., Duchateau, N., Silva, E., Doltra, A., Dhooge, J., Brugada, J., Sitges, M., Frangi, A.F., 2012. Temporal diffeomorphic free-form deformation: application to motion and strain estimation from 3d echocardiography. *Med. Image Anal.* 16 (2), 427–450.
- Eek, C., Grenne, B., Brunvand, H., Aakhus, S., Endresen, K., Hol, P.K., Smith, H.-J., Smiseth, O.A., Edvardsen, T., Skulstad, H., 2010. Strain echocardiography and wall motion score index predicts final infarct size in patients with non-st-segment-elevation myocardial infarction: clinical perspective. *Circulation* 3 (2), 187–194.
- Galasko, G., Basu, S., Lahiri, A., Senior, R., 2001. A prospective comparison of echocardiographic wall motion score index and radionuclide ejection fraction in predicting outcome following acute myocardial infarction. *Heart* 86 (3), 271–276.
- Götte, M.J., Germans, T., Rüssel, I.K., Zwanenburg, J.J., Marcus, J.T., van Rossum, A.C., van Veldhuisen, D.J., 2006. Myocardial strain and torsion quantified by cardiovascular magnetic resonance tissue tagging: studies in normal and impaired left ventricular function. *J. Am. Coll. Cardiol.* 48 (10), 2002–2011.
- Grant, M., Boyd, S., 2008. Graph implementations for nonsmooth convex programs. In: *Blondel, V., Boyd, S., Kimura, H. (Eds.), Recent Advances in Learning and Control*. Springer-Verlag Limited, pp. 95–110. http://stanford.edu/~boyd/graph_dcp.html.
- Heyde, B., Barbosa, D., Claus, P., Maes, F., Dhooge, J., 2012. Three-dimensional cardiac motion estimation based on non-rigid image registration using a novel transformation model adapted to the heart. In: *International Workshop on Statistical Atlases and Computational Models of the Heart*. Springer, pp. 142–150.
- Heyde, B., Bouchez, S., Thieren, S., Vandenhevel, M., Jasaityte, R., Barbosa, D., Claus, P., Maes, F., Wouters, P., Dhooge, J., 2013. Elastic image registration to quantify 3-d regional myocardial deformation from volumetric ultrasound: experimental validation in an animal model. *Ultrasound Med. Biol.* 39 (9), 1688–1697.
- Hoffman, A.J., Kruskal, J.B., 2010. Integral boundary points of convex polyhedra. In: *50 Years of Integer Programming 1958–2008*. Springer, pp. 49–76.
- Hoffmann, R., von Bardeleben, S., Kasprzak, J.D., Borges, A.C., ten Cate, F., Firschke, C., Lafitte, S., Al-Saadi, N., Kuntz-Hehner, S., Horstick, G., et al., 2006. Analysis of regional left ventricular function by cineventriculography, cardiac magnetic resonance imaging, and unenhanced and contrast-enhanced echocardiography: a multicenter comparison of methods. *J. Am. Coll. Cardiol.* 47 (1), 121–128.
- Huang, X., Dione, D.P., Compas, C.B., Papademetris, X., Lin, B.A., Bregasi, A., Sinusas, A.J., Staib, L.H., Duncan, J.S., 2014. Contour tracking in echocardiographic sequences via sparse representation and dictionary learning. *Med. Image Anal.* 18 (2), 253–271.
- Langeland, S., Dhooge, J., Wouters, P.F., Leather, H.A., Claus, P., Bijmens, B., Sutherland, G.R., 2005. Experimental validation of a new ultrasound method for the simultaneous assessment of radial and longitudinal myocardial deformation independent of insonation angle. *Circulation* 112 (14), 2157–2162.
- Le Folgoc, L., Delingette, H., Criminisi, A., Ayache, N., 2017. Sparse Bayesian registration of medical images for self-tuning of parameters and spatially adaptive parametrization of displacements. *Med. Image Anal.* 36, 79–97.
- Ledesma-Carbayo, M.J., Kybic, J., Desco, M., Santos, A., Suhling, M., Hunziker, P., Unser, M., 2005. Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation. *IEEE Trans. Med. Imaging* 24 (9), 1113–1126.
- Lu, A., Zontak, M., Parajuli, N., Stendahl, J.C., Boutagy, N., Eberle, M., Alkhalil, I., O'Donnell, M., Sinusas, A.J., Duncan, J.S., 2017. Learning-based spatiotemporal regularization and integration of tracking methods for regional 4d cardiac deformation analysis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 323–331.
- Lubinski, M.A., Emelianov, S.Y., O'Donnell, M., 1999. Speckle tracking methods for ultrasonic elasticity imaging using short-time correlation. *Ultrason. Ferroelectr. Freq. Control/IEEE Trans.* 46 (1), 82–96.
- Papademetris, X., Jackowski, M.P., Rajeevan, N., Constable, R.T., Staib, L.H., 2005. Bioimage suite: an integrated medical image analysis suite. Section of Biomedical Sciences, Dept. of Diagnostic Radiology, Yale School of Medicine. <http://www.bioimagesuite.org>.
- Papademetris, X., Sinusas, A.J., Dione, D.P., Constable, R.T., Duncan, J.S., 2002. Estimation of 3-d left ventricular deformation from medical images using biomechanical models. *IEEE Trans. Med. Imaging* 21 (7), 786–800.
- Parajuli, N., Compas, C.B., Lin, B.A., Sampath, S., O'Donnell, M., Sinusas, A.J., Duncan, J.S., 2015. Sparsity and biomechanics inspired integration of shape and speckle tracking for cardiac deformation analysis. In: *International Conference on Functional Imaging and Modeling of the Heart*. Springer, pp. 57–64.
- Parajuli, N., Lu, A., Stendahl, J.C., Zontak, M., Boutagy, N., Alkhalil, I., Eberle, M., Lin, B.A., O'Donnell, M., Sinusas, A.J., et al., 2017. Flow network based cardiac motion tracking leveraging learned feature matching. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 279–286.
- Parajuli, N., Lu, A., Stendahl, J.C., Zontak, M., Boutagy, N., Eberle, M., Alkhalil, I., O'Donnell, M., Sinusas, A.J., Duncan, J.S., 2016. Integrated dynamic shape tracking and rf speckle tracking for cardiac motion analysis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 431–438.
- Pellerin, D., Sharma, R., Elliott, P., Veyrat, C., 2003. Tissue doppler, strain, and strain rate echocardiography for the assessment of left and right systolic ventricular function. *Heart* 89 (suppl 3), iii9–iii17.
- Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J., 1999. Non-rigid registration using free-form deformations: application to breast mr images. *IEEE Trans. Med. Imaging* 18 (8), 712–721.
- Song, S.M., Leahy, R.M., 1991. Computation of 3-d velocity fields from 3-d cine ct images of a human heart. *IEEE Trans. Med. Imaging* 10 (3), 295–306.
- Sugeng, L., Mor-Avi, V., Weinert, L., Niel, J., Ebner, C., Stenger-Mascherbauer, R., Schmidt, F., Galuschky, C., Schummers, G., Lang, R.M., et al., 2006. Quantitative assessment of left ventricular size and function: side-by-side comparison of real-time three-dimensional echocardiography and computed tomography with magnetic resonance reference. *Circulation* 114 (7), 654–661.
- Waldman, L.K., Fung, Y., Covell, J.W., 1985. Transmural myocardial deformation in the canine left ventricle. normal in vivo three-dimensional finite strains. *Circ. Res.* 57 (1), 152–163.
- Wendland, H., 1995. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv. Comput. Math.* 4 (1), 389–396.
- Yan, P., Sinusas, A., Duncan, J.S., 2007. Boundary element method-based regularization for recovering of lv deformation. *Med. Image Anal.* 11 (6), 540–554.
- Yancy, C.W., Jessup, M., Bozkurt, B., Butler, J., Casey, D.E., Drazner, M.H., Fonarow, G.C., Geraci, S.A., Horwich, T., Januzzi, J.L., et al., 2013. 2013 AHA/ACC guideline for the management of heart failure: a report of the american college of cardiology foundation/american heart association task force on practice guidelines. *J. Am. Coll. Cardiol.* 62 (16), e147–e239.
- Yao, S.-S., Qureshi, E., Sherrid, M.V., Chaudhry, F.A., 2003. Practical applications in stress echocardiography: risk stratification and prognosis in patients with known or suspected ischemic heart disease. *J. Am. Coll. Cardiol.* 42 (6), 1084–1090.