Tracking Using Motion Estimation With Physically Motivated Inter-Region Constraints

Omar Arif, Ganesh Sundaramoorthi*, Byung-Woo Hong, and Anthony Yezzi

***Abstract*—We propose a method for tracking structures (e.g., ventricles and myocardium) in cardiac images (e.g., magnetic resonance) by propagating forward in time a previous estimate of the structures using a new physically motivated motion estimation scheme. Our method estimates motion by regularizing only within structures so that differing motions among different structures are not mixed. It simultaneously satisfies the physical constraints at the interface between a fluid and a medium that the normal component of the fluid's motion must match the normal component of the medium's motion and the No-Slip condition, which states that the tangential velocity approaches zero near the interface. We show that these conditions lead to partial differential equations with Robin boundary conditions at the interface, which couple the motion between structures. We show that propagating a segmentation across frames using our motion estimation scheme leads to more accurate segmentation than traditional motion estimation that does not use physical constraints. Our method is suited to interactive segmentation, prominently used in commercial applications for cardiac analysis, where segmentation propagation is used to predict a segmentation in the next frame. We show that our method leads to more accurate predictions than a popular and recent interactive method used in cardiac segmentation.**

***Index Terms*—Cardiac magnetic resonance image (MRI) segmentation, image registration, motion estimation, tracking.**

I. Introduction

ACCURATE segmentation of deforming structures from medical image sequences [e.g., cardiac magnetic resonance imaging (MRI)] is an important step in many clinical applications that study structure and function of organs. While many methods have been proposed to segment a deforming structure by dividing each frame based on image intensity statistics and incorporating training data (see Section I-A), it is sometimes easier to exploit the temporal coherence of the structure, and apply a tracking framework. One matches a current estimate of the structure of interest at time t to the image at time $t+1$ to detect the structure at time $t+1$. This requires an accurate registration between images.

One difficulty in registration stems from the nonuniqueness of registrations that are able to explain two images, and therefore, regularization is needed to constrain the set of possible solutions. Typically, global regularization, uniform regularization across the whole image, is used. However, the image consists of many structures, each moving with different motions/deformations, and global regularization smooths across multiple structures. Therefore, motion/deformation information from surrounding structures is used in the registration estimate within the structure of interest. This leads to errors in the registration and therefore also, in the structure segmentation. It may seem that this problem can be avoided by registering only the organ at time t (not the entire image) to a subset of the image at time $t+1$, thus regularizing only within the organ. However, that approach is problematic as the background registration helps limit the possible registrations of the organ of interest, aiding the registration of the organ. See Figs. 6 and 7 for an experiment.

In this work, we derive a new motion estimation method that estimates motion separately in each structure by performing only within structure regularization (so that motions from heterogeneous structures are not mixed) while satisfying physical motion constraints across the boundary between a fluid and a medium. Specifically, we model the physical constraints that the motion of a fluid and a medium at the interface is such that the normal components of the motions are the same, and the No-Slip condition, which states that the tangential motions approach zero at the interface.

Our new motion estimation scheme can be used in structure segmentation from image sequences. Given the initial segmentation in the first frame, our algorithm propagates the segmentation to the next frame. Although our method is not restricted to a particular imaging modality or structure, we focus on an application where the physically motivated considerations are natural—segmentation of the left (LV) and right ventricle (RV), and the surrounding heart muscle from cardiac MRI. RV and LV segmentation are important and challenging problems. In fact, there were recent challenges for segmentation of the LV [54] and RV [28], held in the course of the Medical

Manuscript received April 02, 2014; revised May 05, 2014; accepted May 07, 2014. Date of publication May 16, 2014; date of current version August 28, 2014. This work was supported in part by KAUST Baseline and Visual Computing Center funding, Korea NRF-2010-220-D00078 and NRF-2011-0007898, and the National Science Foundation (NSF) under Grant CCF-1347191 and Grant CMMI-1068624. *Asterisk indicates corresponding author.*

O. Arif is with the Department of Electrical Engineering, King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia (e-mail: omar.arif@kaust.edu.sa).

*G. Sundaramoorthi is with the Department of Electrical Engineering and Department of Applied Mathematics and Computational Science, King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia (e-mail: ganesh.sundaramoorthi@kaust.edu.sa).

B.-W. Hong is with the Department of Computer Science, Chung-Ang University, Seoul 156-756, Korea (e-mail: hong@cau.ac.kr).

A. Yezzi is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ayezzi@ece.gatech.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TMI.2014.2325040

0278-0062 © 2014 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

Image Computing and Computer Assisted Intervention Conference (MICCAI). Our new motion estimation scheme leads to more accurate segmentation of the LV and RV than motion estimated with global regularization.

The main motivation for our new method is segmentation propagation in interactive segmentation of image sequences. Segmentation propagation predicts the segmentation in the next frame from the current estimate. Interactive methods are the norm in commercial medical applications for cardiac MRI as full automation (see Section I-A) is still at the research stage. Our new physically viable motion estimation scheme leads to better segmentation propagation than recent commercially available software for cardiac segmentation.

A. Related Work: Cardiac Segmentation

We briefly review the literature on cardiac segmentation. See [39], [59] for a more thorough review.

There are two types of methods for cardiac segmentation, **fully automatic** and **interactive**. Early methods for fully automatic cardiac segmentation use image partitioning algorithms (e.g., active contours [22] implemented via level sets [34], graph cuts [4], or other optimization methods [5]), which optimize energies that integrate basic image features such as edges [40], intensity statistics [10], motion cues [14], and smoothness of the partition. These methods partition the image into regions of homogeneous statistics, however, they do not always select regions that correspond to physical objects/structures.

One way to improve basic partitioning algorithms is to use training data to construct a model of the heart and then use the model to aid segmentation. Some approaches (e.g., [19], [25], [50], [52], [56], [57], [60]) make use of active shape and appearance models [8], [9] where manual landmarks around the boundary of the object model the shape, and texture descriptors describing a neighborhood around the landmark are used to model object appearance. Landmarks and descriptors allow for the use of PCA to generate a statistical model from training data. More precise models of shape are based on performing PCA of segmented objects using mesh-based approaches [16] or level set representations [38], [51]. These methods construct static models of the heart. However, the heart is a dynamic object, and thus, in [44], [58], dynamic models of shape are constructed by considering the shape from multiple frames in the sequence. Once the heart model is constructed, object segmentation is performed by fitting the model to the image. This can be done in a number of ways, the most common method (e.g., [6], [13], [51]) is to restrict the optimization of the energies based on basic image partitioning to the shapes determined by the parametric shape model. Alternative approaches register an atlas, which has both shape and appearance information, to the target image, thereby determining the object segmentation [21], [24], [29], [60]. Multi-atlas approaches (e.g., [1], [21], [61]) register several atlases to a target image and fuse the results.

The above approaches aim at fully automated segmentation. Although fully automatic segmentation is the ideal goal, these methods are not accurate enough, especially when there is deviation from the training set (e.g., in cases of disease), to be used in many cardiac applications (e.g., [23], [26], [27], [30], [31], [37], [48], [51], [52], [59]). Therefore, in commercial applications, interactive approaches to segmentation, in which user input is integrated, are the norm. Various computer vision techniques (e.g., [3], [12], [55]) have been designed to incorporate user specified seed points. These methods modify energies based on simple image features to incorporate constraints from the seed points entered by the user.

Other interactive methods allow for a manual or semi-automated segmentation of the first frame in the cardiac sequence and then attempt to propagate the segmentation to subsequent frames, thereby predicting a segmentation that needs little interaction to correct [33], [47]. Several methods propagate the initial segmentation (e.g., by registration [33], [47]) and/or use manual segmentation in the first frame as initialization to an automated segmentation algorithm. Several commercial software for interactive heart segmentation have been designed. For example, the recent software Medviso [18], [49] allows the user to input an initial segmentation, which is then propagated to subsequent frames in order to segment various structures including the ventricles and myocardium. The software also allows for other manual interactions to correct any errors in the propagation. The algorithm is a culmination of many techniques including registration to propagate the segmentation as well as the use of shape priors constructed from training data.

Our contribution lies in improving the prediction step that is present in interactive methods for cardiac image sequence segmentation. We derive a segmentation propagation method that is more accurate than existing schemes.

B. Related Work: Registration

We construct a segmentation propagation scheme by deriving a physically motivated registration method. Therefore, we give a brief review of recent related work in registration.

The goal of registration is to find pixel-wise correspondence between images. Regularization is needed as multiple registrations can explain the images. The pioneering work [20] from computer vision uses uniform smoothness applied to the entire image domain (global regularization) to estimate registration under small pixel displacements. Larger deformations computed using global regularization that lead to diffeomorphic registrations, a property of a valid registration in typical medical images, have been considered by [2], [7], [53]. For some medical images (e.g., cardiac), there are multiple structures each of which have different motion characteristics. Global regularization across adjacent objects (mixing heterogeneous motion characteristics) is not desired in this case. Moreover, there are physical constraints between adjacent regions that can be exploited. In cardiac applications, the ventricles and the heart muscle (myocardium), which are a fluid (blood) interacting with a medium, are such that the normal component on the boundary of the motions (velocity) of both structures are equal. The previous constraint arises from the fact that the ventricle and the surrounding muscle do not separate during motion. Further, the No-Slip condition [32] from fluid mechanics for viscous fluids states that the motion of the fluid, at the interface of a medium, relative

to the boundary is zero, i.e., the tangent component of the motion along the boundary is zero. The scale at which this happens may be small compared with the resolution of the imaging device, and thus we allow an option for the tangent components to be arbitrary across the boundary.

Recent work has considered nonglobal regularization. In lung registration, the lung slides along the rib-cage, and the motion of both these structures is different and thus global regularization is not desired. In [35], [36], [45], [46] anisotropic regularization is used to favor smoothing in the tangential direction near organ boundaries. In [42], [43], Log-Demons [53] is generalized so that smoothing is performed on the tangent component of the registration within organs, and the normal component is globally smoothed. Our approach differs as it is motivated by the physical constraints of fluid motion present in the heart. Further, our method applies regularization only within homogeneous structures, different than [42], [43], which smooths the normal component across structures, mixing inhomogeneous motions. In [35], [36], [45], [46] the regularization across the interface in the normal direction can be chosen nonzero, thus mixing motions of heterogeneous structures, which our algorithm is designed to avoid as the cardiac application requires *within* structure regularization. A choice of zero regularization in the normal direction in [35], [36], [45], [46] implies only smoothing along the tangential direction, not ensuring that the normal motions are equal on both sides of the boundary. Such equality is designed in our algorithm. Both choices for the normal direction regularization in [35], [36], [45], [46] do not simultaneously satisfy within structure regularization and matching normal velocities on each side of the boundary, which our application requires. While [42], [43] achieves equal normal motions, it does so by smoothing across structures. Our method also implements the No-Slip condition, not present in other works.

C. Organization of Paper

In Section II, we specify the physical motion constraints between structures, and formulate a registration model. In Section III, we use the model to setup a variational problem for estimating the motion given the two images assuming infinitesimal motion, and then show how to estimate the motion. In Section IV, we show how to estimate noninfinitesimal motion and simultaneously propagate the segmentation across frames. In Section V, we show how to track multiple structures together. Finally, in Section VI we show a series of experiments to verify our method as well as compare it to recent methods.

II. Motion and Registration Model

In this section, we state the assumptions for motion estimation. We assume the brightness constancy plus noise model

$$J_1(x) = J_0(w^{-1}(x)) + \eta(x),\ x \in \Omega \tag{1}$$

where $\Omega \subset \mathbb{R}^n$ $(n = 2, 3)$ is the image domain, $J_0, J_1 : \Omega \to \mathbb{R}$ are two consecutive images from the sequence, $w : \Omega \to \Omega$ is the registration between frames J_0 and J_1, and η is a noise process, which is assumed to be Gaussian. The structure of interest in J_0 is denoted $R \subset \Omega$ and the background is denoted $R^c := \Omega \backslash R$. Our model can be extended to any number of structures, but we forgo the details for simplicity of presentation. The structure in frame J_1 is $w(R)$.

The registration is an invertible map, and thus we represent it as an integration of a time varying velocity field (following standard representation from fluid mechanics):

$$w(x) = \phi_T(x),\ \phi_s(x) = x + \int_0^s v_\tau(\phi_\tau(x))\,\mathrm{d}\tau,\ s \in [0, T] \tag{2}$$

where $T > 0$, $v_\tau : \Omega \to \mathbb{R}^n$ is a velocity field, and $\phi_\tau : \Omega \to \Omega$ for every $\tau \in [0, T]$. The map ϕ_τ is such that $\phi_\tau(x)$ indicates the mapping of x after it flows along the velocity field for time τ, which is an artificial time parameter.

We assume that the motion/deformation of the structure of interest R and the surrounding region have different characteristics and therefore the registration w consists of two components, w^i and w^o defined inside the structure of interest R and outside the structure R^c, resp. This can be achieved with a velocity field that has two components v_τ^i, v_τ^o (both smooth within their domains)

$$w(x) = \begin{cases} w^i(x) & x \in R \\ w^o(x) & x \in R^c \end{cases}, \quad v_\tau(x) = \begin{cases} v_\tau^i(x) & x \in R_\tau \\ v_\tau^o(x) & x \in R_\tau^c \end{cases} \tag{3}$$

where $R_\tau = \phi_\tau(R)$. This implies that w^i and w^o are smooth and invertible. When the structure R contains a fluid (as the ventricles) and R^c is the surrounding medium (e.g., myocardium), as in our case of interest, there are physical constraints from fluid mechanics imposed on v_τ^i and v_τ^o at the boundary ∂R_τ. One constraint is that the normal component of v_τ^i and v_τ^o are equal on ∂R_τ

$$v_\tau^i(x) \cdot N(x) = v_\tau^o(x) \cdot N(x),\ x \in \partial R_\tau \tag{4}$$

where N indicates the surface normal of ∂R_τ. The condition implies that R and R^c do not separate when deformed by the infinitesimal motion. Further, the No-Slip Condition from fluid mechanics [32] implies that the motion of the fluid relative to the surrounding medium is zero at the interface, thus, the tangent component of v_τ^i and v_τ^o are zero, i.e.,

$$v_\tau^i - (v_\tau^i \cdot N)N = v_\tau^o - (v_\tau^o \cdot N)N = 0,\ \text{on}\ \partial R_\tau. \tag{5}$$

Modeling the velocity with two separate components allows for different deformations for the structure of interest and surrounding medium, and the constraint (4) couples the two components in a physically plausible manner as does the No-Slip condition. This implies *continuity* of the normal component of v_τ across ∂R_τ, but not necessarily differentiability.

III. Energy for Infinitesimal Deformations

In this section, we consider the case when the registration is $w(x) = x + v(x)$ where v is an infinitesimal deformation, and we show how one computes v such that (4) and (5) are satisfied while applying regularization only within regions.

In tracking, R is given (e.g., at the initial frame or the estimate from a previous frame), and the goal is to determine v and $R +$

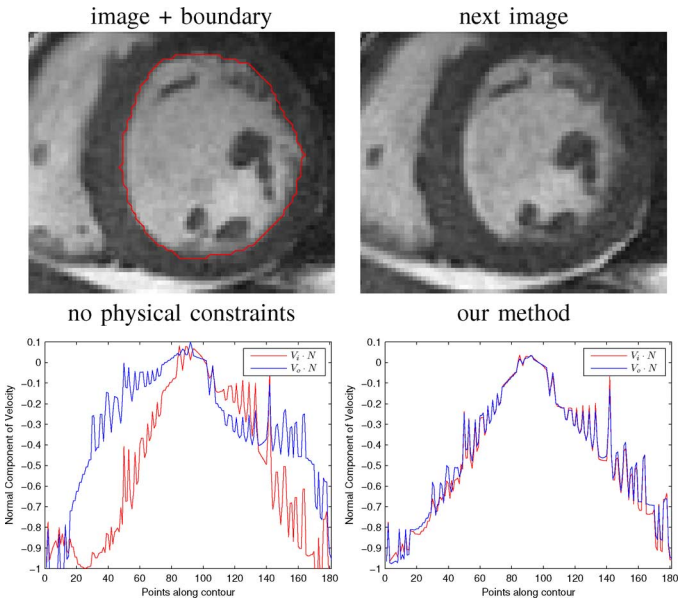

Fig. 1. Regularization within regions alone does not induce matching normal motions. [Top, left]: J_0 and ∂R in red. [Top, right]: J_1. [Bottom, left]: Normal components of the velocities just inside and outside ∂R using optical flow computed separately within R and R^c. They are different, and thus not physically viable. [Bottom, right]: Our approach (regularization only within regions, while satisfying the normal matching constraint): normal motions along ∂R are equal.

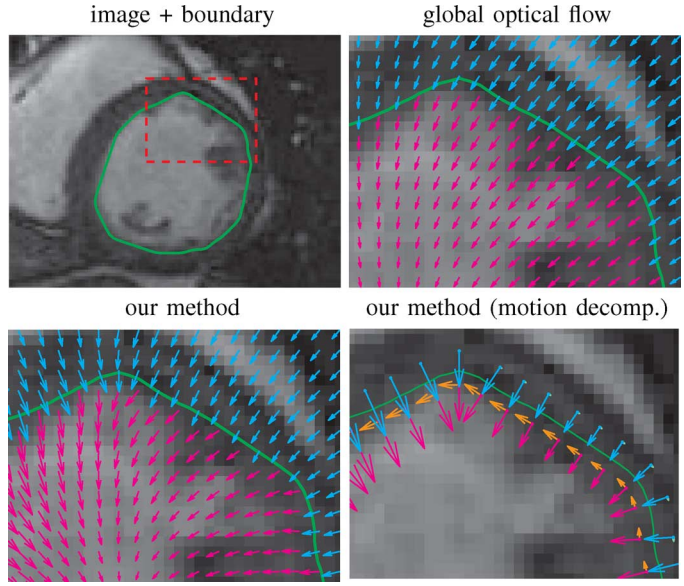

Fig. 2. Why regularization within regions? Heart is contracting. [Top, left]: J_0, ∂R in green, and the red box is zoomed in next images. [Top, right]: velocity field with global regularization shows smoothing across the boundary. [Bottom, left]: velocities v^i, v^o with within region regularization and normal constraint shows difference to global regularization. [Bottom, right]: normal and tangent components of inside/outside velocities v^i, v^o; notice the matching normal components, but discontinuity of tangent components (amplified by a factor of 2 for visualization) across boundary (No-Slip not enforced). Boundary velocity indicates contraction of myocardium, and tangent component inside shows some circulation near boundary: cannot be captured with global regularization.

$v = \{x + v(x) : x \in R\}$, the object in the next frame. The energy, which applies to both 2-D and 3-D data, is

$$E(v_i, v_o; R, I) = \frac{1}{2}\int_R \left(|J_1 - I + \nabla I \cdot v_i|^2 + \alpha_i |\nabla v_i|^2\right) \mathrm{d}x + \frac{1}{2}\int_{R^c} \left(|J_1 - I + \nabla I \cdot v_o|^2 + \alpha_o |\nabla v_o|^2\right) \mathrm{d}x \tag{6}$$

subject to

$$v_i - (v_i \cdot N)N = v_o - (v_o \cdot N)N = 0, \text{ and}$$
$$v_i \cdot N = v_o \cdot N \text{ on } \partial R. \tag{7}$$

The velocities are defined $v_i : R \to \mathbb{R}^n$ and $v_o : R^c \to \mathbb{R}^n$ (as in (3)). The operator ∇ is the spatial gradient defined within R or R^c (without crossing ∂R). The first term in each of the above integrals arises from a linearization of (1) (when $I = J_0$). The second terms in each of the integrals are regularization. Regularization is done *separately* within each of the regions. The weights $\alpha_i, \alpha_o > 0$ indicate the amount of regularity desired within each region. I and R are given variables, and in the final deformation scheme to be described in Section IV, they will initially be J_0 and the initial region, resp.; they will then be updated as part of an iterative scheme.

The problem above is a generalization of Horn–Schunck optical flow [20]. Solving for the Horn–Schunck optical flow within each region separately does not lead to velocities that have equal normal components (see Fig. 1), whereas the solution of (6) and (7) to be presented next does. Computing Horn–Schunck optical flow in each region requires boundary conditions, and typically they are chosen to be Neumann boundary conditions: $\nabla v_i \cdot N = 0$ and $\nabla v_o \cdot N = 0$ on ∂R. Replacing those conditions with the constraints (7) does not specify a unique solution. While Horn–Schunck optical flow computed on the whole domain Ω gives a globally smooth motion, which by default satisfies matching normals at the interface, this is not natural for the ventricles/myocardium, where different motions exist in the regions (see Fig. 2), and the motions should not be smoothed across regions.

In the next subsections, we show how the constraints (7) can be enforced. The methods apply to both 2-D and 3-D data.

A. Energy Optimization With a Soft Constraint

The first approach to optimize (6) subject to (7) is to add a term into (6) that penalizes deviations away from (7). This favors deformations that satisfy the constraint, but does not satisfy the constraint exactly. The energy is

$$E_s(v_i, v_o; R, I) = E(v_i, v_o; R, I) + \frac{\beta}{2}\int_{\partial R} (v_i \cdot N - v_o \cdot N)^2 \,\mathrm{d}s + \frac{\gamma}{2}\int_{\partial R} \left(|v_i - (v_i \cdot N)N|^2 + |v_o - (v_o \cdot N)N|^2\right) \mathrm{d}s \tag{8}$$

where $\beta, \gamma > 0$. The space of v_i, v_o is linear, the energy E_s is convex, and thus any local optimum is a global optimum. To find the necessary conditions, we compute the first variation of E_s. Let $h_i : R \to \mathbb{R}^2$ and $h_o : R^c \to \mathbb{R}^2$ be perturbations of v_i

and v_o, then the first variation (after simplification) is

$$\begin{aligned} &dE(v_i, v_o) \cdot (h_i, h_o) \\ &= \int_R [(J_1 - I + \nabla I \cdot v_i)\nabla I - \alpha_i \Delta v_i] \cdot h_i \, \mathrm{d}x \\ &\quad + \int_{R^c} [(J_1 - I + \nabla I \cdot v_o)\nabla I - \alpha_o \Delta v_o] \cdot h_o \, \mathrm{d}x \\ &\quad + \int_{\partial R} [\alpha_i (\nabla v_i \cdot N) \cdot h_i + \alpha_o (\nabla v_o \cdot N) \cdot h_o] \, \mathrm{d}s \\ &\quad + \int_{\partial R} [\beta (v_i \cdot N - v_o \cdot N) N + \gamma (v_i - (v_i \cdot N) N)] \cdot h_i \, \mathrm{d}s \\ &\quad + \int_{\partial R} [\beta (v_i \cdot N - v_o \cdot N) N + \gamma (v_o - (v_o \cdot N) N)] \cdot h_o \, \mathrm{d}s \end{aligned} \tag{9}$$

where the boundary integrals arise from integration by parts, and Δ denotes the Laplacian operator. The necessary conditions for a minimum are obtained by choosing v_i, v_o such that $dE(v_i, v_o) \cdot (h_i, h_o) = 0$ for all h_i, h_o, and thus, we have

$$-\alpha_i \Delta v_i + \nabla I \nabla I^T v_i = -(J_1 - I)\nabla I, \quad \text{in } R \tag{10}$$

$$-\alpha_o \Delta v_o + \nabla I \nabla I^T v_o = -(J_1 - I)\nabla I, \quad \text{in } R^c \tag{11}$$

$$\begin{aligned} &\alpha_i \nabla v_i \cdot N + \beta (v_i \cdot N - v_o \cdot N) N \\ &\quad + \gamma (v_i - (v_i \cdot N) N) = 0, \text{ on } \partial R \end{aligned} \tag{12}$$

$$\begin{aligned} &\alpha_o \nabla v_o \cdot N + \beta (v_i \cdot N - v_o \cdot N) N \\ &\quad + \gamma (v_o - (v_o \cdot N) N) = 0, \text{ on } \partial R \end{aligned} \tag{13}$$

where ^{T} indicates transpose, and $\nabla v_i \cdot N := (\nabla v_i)^T N$ (∇v_i is a matrix). The last two conditions are Robin boundary conditions, which are conditions on a linear combination of the normal derivatives of the functions and the function values on the boundary. These boundary conditions specify a unique solution of the PDE. Note that v_i and v_o are linked to each other by the boundary conditions (unlike separate solution of v_i and v_o using Neumann boundary conditions in traditional optical flow). Such a link between v_i and v_o is expected given that the normal components on the boundary are to be close.

B. Energy Optimization With a Hard Constraint

We now optimize E in (6) subject to (7) by enforcing (7) exactly. We first do the derivation without the No-Slip condition and then enforce the No-Slip condition at the end. The space of v_i, v_o satisfying the normal constraint is a linear space and the energy is convex, so any local optimum must be a global optimum. We compute the first variation of E evaluated at (v_i, v_o) applied to a perturbation h_i, h_o in the permissible space (those that perturb v_i, v_o so that the constraint is satisfied). The space of permissible perturbations satisfy the normal matching constraint: $h_i \cdot N = h_o \cdot N$ on ∂R (this is obtained by differentiating the constraint in the direction of h_i, h_o). The variation is $dE(v_i, v_o) \cdot (h_i, h_o) =$

$$\begin{aligned} &\int_R [(J_1 - I + \nabla I \cdot v_i)\nabla I - \alpha_i \Delta v_i] \cdot h_i \, \mathrm{d}x \\ &\quad + \int_{R^c} [(J_1 - I + \nabla I \cdot v_o)\nabla I - \alpha_o \Delta v_o] \cdot h_o \, \mathrm{d}x \\ &\quad + \int_{\partial R} \alpha_i (\nabla v_i \cdot N) \cdot h_i \, \mathrm{d}s - \int_{\partial R} \alpha_o (\nabla v_o \cdot N) \cdot h_o \, \mathrm{d}s. \end{aligned} \tag{14}$$

One can decompose h_i and h_o on ∂R into its normal and tangential components

$$\begin{aligned} h_i &= \pi_N(h_i) + \pi_T(h_i) \\ h_o &= \pi_N(h_o) + \pi_T(h_o) \end{aligned}$$

where $\pi_N(h) = (h \cdot N)N$ and $\pi_T(h) = h - \pi_N(h)$ for h defined on ∂R. Note that $\pi_N(h_i) = \pi_N(h_o)$ by permissibility, and thus, we will set $h^N = \pi_N(h_i) \cdot N = \pi_N(h_o) \cdot N$. One can similarly decompose $\nabla v_i \cdot N$ and $\nabla v_o \cdot N$ on ∂R; then

$$\begin{aligned} (\nabla v_i \cdot N) \cdot h_i &= h^N (\nabla v_i \cdot N) \cdot N + \pi_T(h_i) \cdot \pi_T(\nabla v_i \cdot N) \\ (\nabla v_o \cdot N) \cdot h_o &= h^N (\nabla v_o \cdot N) \cdot N + \pi_T(h_o) \cdot \pi_T(\nabla v_o \cdot N). \end{aligned}$$

Substituting these formulas into the variation (14) yields

$$\begin{aligned} &dE(v_i, v_o) \cdot (h_i, h_o) \\ &\quad = \int_R [(J_1 - I + \nabla I \cdot v_i)\nabla I - \alpha_i \Delta v_i] \cdot h_i \, \mathrm{d}x \\ &\qquad + \int_{R^c} [(J_1 - I + \nabla I \cdot v_o)\nabla I - \alpha_o \Delta v_o] \cdot h_o \, \mathrm{d}x \\ &\qquad + \int_{\partial R} h^N [\alpha_i (\nabla v_i \cdot N) \cdot N - \alpha_o (\nabla v_o \cdot N) \cdot N] \, \mathrm{d}s \\ &\qquad + \int_{\partial R} \alpha_i \pi_T(h_i) \cdot \pi_T(\nabla v_i \cdot N) \, \mathrm{d}s \\ &\qquad - \int_{\partial R} \alpha_o \pi_T(h_o) \cdot \pi_T(\nabla v_o \cdot N) \, \mathrm{d}s. \end{aligned}$$

Since the h^N and $\pi_T(h_i)$, $\pi_T(h_o)$ may be chosen independently, the necessary conditions for an optimum are

$$-\alpha_i \Delta v_i + \nabla I \nabla I^T v_i = -(J_1 - I)\nabla I, \quad \text{in } R \tag{15}$$

$$-\alpha_o \Delta v_o + \nabla I \nabla I^T v_o = -(J_1 - I)\nabla I, \quad \text{in } R^c \tag{16}$$

$$\alpha_i (\nabla v_i \cdot N) \cdot N = \alpha_o (\nabla v_o \cdot N) \cdot N, \quad \text{on } \partial R \tag{17}$$

$$\nabla v_i \cdot N = \pi_N (\nabla v_i \cdot N), \quad \text{on } \partial R \tag{18}$$

$$\nabla v_o \cdot N = \pi_N (\nabla v_o \cdot N), \quad \text{on } \partial R \tag{19}$$

$$v_i \cdot N = v_o \cdot N, \quad \text{on } \partial R. \tag{20}$$

The above PDE is uniquely specified, and thus the solution specifies a global optimum. The boundary conditions indicate that the normal derivatives of v_i, v_o only have normal components, and the normal components of the normal derivative of v_i and v_o differ by a scalar factor α_o/α_i. Like the case of enforcing the normal continuity constraint using the soft penalty, v_i and v_o are related by the boundary conditions, which enforce the normal continuity constraint exactly while regularizing only within regions R and R^c separately.

Remark 1 (No-Slip Condition): The No-Slip condition can be easily enforced as a hard constraint. The optimizing conditions (15) and (16) remain the same, however, some of the boundary conditions are modified as (since $\pi_T(h_i) = \pi_T(h_o)$)

$$\alpha_i (\nabla v_i \cdot N) \cdot N = \alpha_o (\nabla v_o \cdot N) \cdot N, \quad \text{on } \partial R \tag{21}$$

$$v_i = \pi_N (v_i), \quad \text{on } \partial R \tag{22}$$

$$v_o = \pi_N (v_o), \quad \text{on } \partial R \tag{23}$$

$$v_i \cdot N = v_o \cdot N, \quad \text{on } \partial R. \tag{24}$$

C. Numerical Solution for Infinitesimal Deformation

The operators on the left-hand side of (15) and (16) (and similarly (10) and (11) in the previous subsection) that act on v_i and v_o are (with the given boundary conditions) positive semi-definite, and thus, one may use the conjugate gradient algorithm for a fast numerical solution.

Since the boundary conditions are not standard in medical imaging, we now show one possible scheme for the numerical discretization of the PDE in the previous sections. We apply a finite difference discretization. Consider a regular grid, we apply finite differences for the Laplacian

$$\Delta v_i(x) = \sum_{\substack{y\sim x\\ y\in R}}(v_i(y)-v_i(x)) + \sum_{\substack{y\sim x\\ y\in R^c}}(v_i(y)-v_i(x)) \quad (25)$$

$$\Delta v_o(x) = \sum_{\substack{y\sim x\\ y\in R^c}}(v_o(y)-v_o(x)) + \sum_{\substack{y\sim x\\ y\in R}}(v_o(y)-v_o(x)) \quad (26)$$

where $y \sim x$ indicates that y is a four-neighbor (in 2-D) or six-neighbor (in 3-D) of x. $v_i(y)$ is not defined for $y \in R^c$ (also, $v_o(y)$ is not defined for $y \in R$). Therefore, we extrapolate these quantities by discretizing the boundary conditions.

1) Soft Constraint: We consider discretization of (12) and (13). Let $x \in R$ and $y \in R^c$. Applying a one-sided first order difference to approximate $\nabla v_i(x)\cdot N$ and $\nabla v_o(y)\cdot N$, we find

$$\begin{aligned}\alpha_i(v_i(y)-v_i(x)) + \beta\pi_N(v_i(x)-v_o(x))\\ +\gamma[v_i(x)-\pi_N(v_i(x))] = 0\\ \alpha_o(v_o(y)-v_o(x)) + \beta\pi_N(v_i(y)-v_o(y))\\ +\gamma[v_o(x)-\pi_N(v_o(x))] = 0\end{aligned}$$

where N, the outward normal can be approximated simply by the unit vector pointing from x to y, or more accurately as the gradient of the level set function of R. We employ the latter approximation in determining π_N.

Solving for $v_i(y)$ and $v_o(x)$ in terms of $v_o(y)$ and $v_i(x)$, one obtains

$$\begin{aligned}v_i(y) &= v_i(x) - \frac{\gamma}{\alpha_i}\pi_T v_i(x) + \frac{\beta(\beta-\alpha_o)}{\beta^2-\alpha_i\alpha_o}\pi_N(v_o(y)-v_i(x))\\ v_o(x) &= v_o(y) - \frac{\gamma}{\alpha_o}\pi_T v_o(x) - \frac{\beta(\beta-\alpha_i)}{\beta^2-\alpha_i\alpha_o}\pi_N(v_o(y)-v_i(x)).\end{aligned}$$

Let $v = (v_i, v_o)$ be the velocity on Ω and A_s be the operator on the left-hand side of (10) and (11), then the discretization is $A_s v(x) =$

$$\begin{cases}\sum\limits_{\substack{y\sim x\\ y\in R^c}}\left[\frac{\alpha_i\beta(\beta-\alpha_o)}{\beta^2-\alpha_i\alpha_o}\pi_N(v(y)-v(x)) - \frac{\gamma}{\alpha_i}\pi_T v(x)\right]\\ -\alpha_i\sum\limits_{\substack{y\sim x\\ y\in R}}(v(y)-v(x)) + \nabla_i I(x)\nabla_i I(x)^T v(x), \qquad x\in R\\ \sum\limits_{\substack{y\sim x\\ y\in R}}\left[\frac{\alpha_o\beta(\beta-\alpha_i)}{\beta^2-\alpha_i\alpha_o}\pi_N(v(y)-v(x)) - \frac{\gamma}{\alpha_i}\pi_T v(x)\right]\\ -\alpha_o\sum\limits_{\substack{y\sim x\\ y\in R^c}}(v(y)-v(x)) + \nabla_o I(x)\nabla_o I(x)^T v(x), \quad x\in R^c\end{cases} \quad (27)$$

where ∇_i and ∇_o are gradient operators approximated with central differences for interior points of R and R^c, respectively, and one-sided differences are applied at the boundary points so that differences do not cross the boundary.

2) Hard Constraint: By similar methodology, one can discretize the PDE (15)–(20). The discretization of the boundary conditions is

$$\begin{aligned}\alpha_i(v_i(y)-v_i(x))\cdot N &= \alpha_o(v_o(y)-v_o(x))\cdot N\\ v_i(y)-v_i(x) &= \pi_N(v_i(y)-v_i(x))\\ v_o(y)-v_o(x) &= \pi_N(v_o(y)-v_o(x))\\ v_i(y)\cdot N &= v_o(x)\cdot N\end{aligned}$$

for $y \in R^c$ and $x \in R$. Since $v_i(y)$ and $v_o(x)$ are not defined, we derive extrapolation formulas for these quantities by solving the above system for $v_i(y)$ and $v_o(x)$ in terms of $v_i(x)$ and $v_o(y)$. This yields

$$v_i(y) = v_i(x) + \frac{\alpha_o}{\alpha_i+\alpha_o}\pi_N(v_o(y)-v_i(x)) \quad (28)$$

$$v_o(x) = v_o(y) - \frac{\alpha_i}{\alpha_i+\alpha_o}\pi_N(v_o(y)-v_i(x)). \quad (29)$$

The discretization of the operators on the left-hand side of (15) and (16), which we denote A_h, is then $A_h v(x) =$

$$\begin{cases}\frac{\alpha_i\alpha_o}{\alpha_i+\alpha_o}\sum\limits_{\substack{y\sim x\\ y\in R^c}}\pi_N(v(y)-v(x))\\ -\alpha_i\sum\limits_{\substack{y\sim x\\ y\in R}}(v(y)-v(x))\nabla_i I(x)\nabla_i I(x)^T v(x) \qquad x\in R\\ \frac{\alpha_i\alpha_o}{\alpha_i+\alpha_o}\sum\limits_{\substack{y\sim x\\ y\in R}}\pi_N(v(y)-v(x))\\ -\alpha_o\sum\limits_{\substack{y\sim x\\ y\in R^c}}(v(y)-v(x))\nabla_o I(x)\nabla_o I(x)^T v(x) \quad x\in R^c.\end{cases} \quad (30)$$

Remark 2 (No-Slip Condition): It turns out that, up to first order discretization of the boundary conditions, discretization of the No-Slip condition (21)–(24) leads to the same discretization as (17)–(20) (boundary conditions from the hard normal matching constraint without the No-Slip condition). Therefore, the No-Slip condition is automatically enforced (up to the first-order discretization error) by the discretization scheme (30).

3) Solution Using Conjugate Gradient: The solution for $v = (v_i, v_o)$ is obtained by solving

$$Av(x) = -(J_1(x)-I(x))\begin{cases}\nabla_i I(x) & x\in R\\ \nabla_o I(x) & x\in R^c\end{cases} \quad (31)$$

where A is A_s(27) or A_h(30) using conjugate gradient.

We note that the numerical solution for our method has similar computational cost as global regularization using the traditional Horn–Schunck method. The operators A_s and A_h slightly differ from the operator in Horn–Schunck. The simple modification and the computational speed make our method an easy and costless alternative to Horn–Schunck.

Both the hard and soft constraint solution lead to boundary conditions that are not traditional, and both are solved using a similar numerical scheme. An advantage of the soft constraint is that it is more general, i.e., choosing $\beta = \alpha_i = \alpha_o$ yields

the hard formulation. By tuning of the parameter β, it is possible to obtain better segmentation than the hard formulation. The advantage of the hard formulation is that it requires fewer parameters (no choice of β or γ is needed).

IV. Larger Deformation and Shape Tracking

In this section, we consider noninfinitesimal deformations between frames, typical in realistic sequences. We derive a technique for the registration that satisfies (2) and (3). This is accomplished by using the results of the previous section to estimate an initial infinitesimal deformation v_0 between the two given images J_0 and J_1, then J_0 and the region R are deformed infinitesimally by v_0, then the process is repeated on the deformed region and image until convergence.

The scheme mentioned above can be formulated with PDEs. This formulation estimates R_τ, ϕ_τ^{-1}, and w^{-1} defined in (2) and (3). One solves for the incremental deformation v_τ by one of the methods (hard or soft formulation) presented in the previous section, the image J_0 is warped by the accumulated warp ϕ_τ^{-1}, and the procedure is repeated, but this time solving for the velocity to deform $J_0 \circ \phi_\tau^{-1}$. This procedure is

$$v_\tau = \arg\min_v E(v; R_\tau, I_\tau) \tag{32}$$

$$\partial_\tau \phi_\tau^{-1} = -\nabla \phi_\tau^{-1} \cdot v_\tau,\ \phi_0^{-1}(x) = x \tag{33}$$

$$\partial_\tau \Psi_\tau = -\nabla \Psi_\tau \cdot v_\tau,\ \Psi_0(x) = \mathrm{dist}_R(x) \tag{34}$$

$$R_\tau = \{\Psi_\tau \le 0\}, \quad I_\tau = J_0 \circ \phi_\tau^{-1} \tag{35}$$

where dist_R is the signed distance function of R, $R_\tau = \phi_\tau(R)$ is the region formed by flowing R along the velocity field for time τ, and E is defined in (6). The region R_τ is represented by a level set function $\Psi_\tau : \Omega \to \mathbb{R}$, which makes the computation of the region R_τ convenient and allows sub-pixel accuracy, although the level set is not required (R_τ may be directly computed from R and ϕ_τ^{-1}). The level set function satisfies a transport equation shown in (34). The backward map ϕ_τ^{-1} satisfies a transport PDE: the identity map is transported along integral curves of v_τ to determine ϕ_τ^{-1}. While in tracking, only R_τ is desired, the backward map is computed to aid in accurate numerical computation of I_τ(35), which is required to estimate v_τ(32). At the time of convergence of the region R_τ, T, I_T approximates J_1, and the registration between J_1 to J_0 is $w^{-1} = \phi_T^{-1}$.

The approach described to obtain a noninfinitesimal deformation by accumulating infinitesimal deformations relates to [7], which also accumulates infinitesimal deformations. Christensen *et al.* [7] use the conservation of momentum to obtain the infinitesimal velocity, while our method uses a generalization of Horn–Schunck optical flow to obtain the velocity. The main difference of our work from [7] is that our method uses within region regularization and the physical constraints at the interface between a fluid and a medium.

Integrating a sufficiently smooth vector field v_τ to form ϕ_τ as in (33) [and (2)] guarantees that ϕ_τ is a diffeomorphism within R and R^c (see classical results [15]). Since v_τ in each region

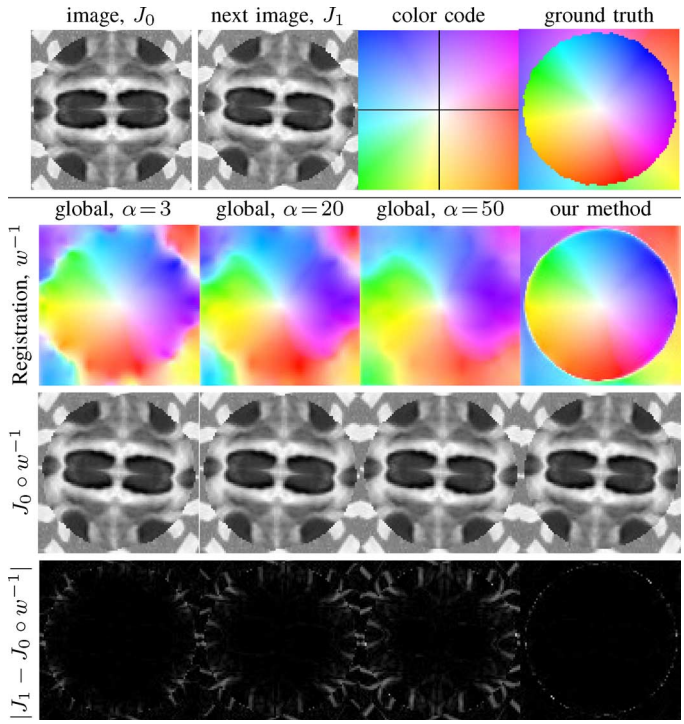

Fig. 3. Global regularization versus physically motivated registration on a synthetic sequence. [Top]: First two frames of the sequence, registration color code (color indicates displacement direction, intensity indicates magnitude), and ground truth registration between images (two regions of differing motion and matching normals on the boundary). [Second row]: Registration computed with global smoothness $\alpha = \{3, 20, 50\}$, and by the proposed method with smoothness $\alpha_i = \alpha_o = 3$. [Third row]: warped images under the registration computed. [Bottom]: Absolute difference between the image warped by the registration and the next image (black—near zero values, white—high values). The proposed method recovers the true registration while no amount of global regularization recovers the true registration.

R_τ and R_τ^c are solutions of Poisson equations, v_τ is differentiable with Sobolev H^2 regularity in each region, and so sufficiently smooth. A surface (namely the zero level set) deformed according to this velocity will remain smooth since the normal component of the velocity is smooth along the boundary and the normal component is the only component that affects the geometry of the surface. Although, the tangent component is discontinuous across the surface, the tangent component does not affect the geometry of the surface. Further, since level sets are used, any tangential components are automatically annihilated in the surface evolution (34).

To track a structure in multiple frames, I_t, $t = 0, \ldots, K-1$, the initial given region $R_{t=0}$ is deformed using (32)–(35) to obtain $R_{t=1}$, the segmentation at frame $t = 1$. Proceeding iteratively, the estimate of the structure in frame t, R_t, is propagated using (32)–(35) to determine R_{t+1} in frame $t+1$.

V. Tracking Multiple Regions

In cardiac image analysis, multiple structures (e.g., the right and left ventricles, and myocardium) should be segmented. Our method is easily adaptable to this case. Computation of v_τ in Section III-C can be readily generalized. Multiple level sets should be used to represent multiple regions. However, in our

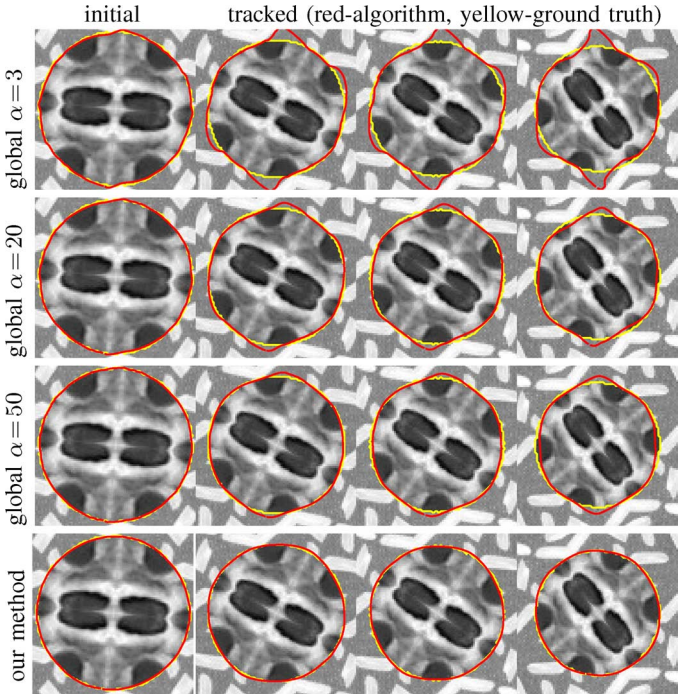

Fig. 4. Global regularization versus physically motivated tracking on synthetic sequence. [First three rows]: Region tracking using global regularization $\alpha = \{3, 20, 50\}$. [Last row]: proposed method. Since the proposed method is able to better capture the underlying registration than global regularization, it also has higher accuracy in segmentation of the circular object.

case of interest (ventricles and surrounding epicardium), the regions form a topology that can be represented with a single level set (ventricles—negative sign, myocardium—positive, everything else—negative; see Fig. 10).

While theoretically ϕ_τ for each τ will be an invertible/onto map in each individual region, and thus regions cannot change topology. However, numerically, between close by structures (e.g., epicardium and RV), self-intersections may occur due to finite step sizes and/or noise. Since we know that a topology change is not physically possible, we enforce that the level set evolution does not induce topology change. This is now standard using discrete topology preserving techniques [17]. The original level set evolution is augmented with a step that looks for *nonsimple points* that change sign in a level set update, i.e., locations of topology change. Such points are not allowed to change sign, and this preserves topology. Nonsimple points are easily detected with local pixel-wise operations, and this adds almost no computational cost. The reader is referred to [17] for details.

The heart has approximately periodic behavior, and enforcing such a constraint may be helpful. While our method does not incorporate this constraint, using a sequential tracking framework does not preclude incorporating a periodic constraint (see for example, [11]).

VI. Experiments

This section consists of five sets of experiments. The first three experiments illustrate and verify that our technique works as expected. The fourth and fifth experiments illustrate the main motivation of our algorithm, which is to improve the prediction step in interactive segmentation algorithms used in commercial applications for cardiac analysis. These latter experiments thus compares our technique to the recent interactive cardiac image segmentation software Medviso [18], [49]. We also compare to automatic segmentation methods.

TABLE I
Quantitative Evaluation of Hard Versus Soft Constraint Formulations on the MICCAI 2009 LV Validation Dataset [41]. The Table Shows a Summary of All Cases. APD is Average Perpendicular Distance and DM is the Dice Metric. Lower APD and Higher DM Indicates Better Fit to Ground Truth

mean $\pm$ std	APD	DM
Hard Constraint	2.39$\pm$0.30	0.88$\pm$0.02
Soft Constraint	2.41$\pm$0.30	0.88$\pm$0.03
Soft Constraint with No Slip	2.31$\pm$0.29	0.89$\pm$0.02

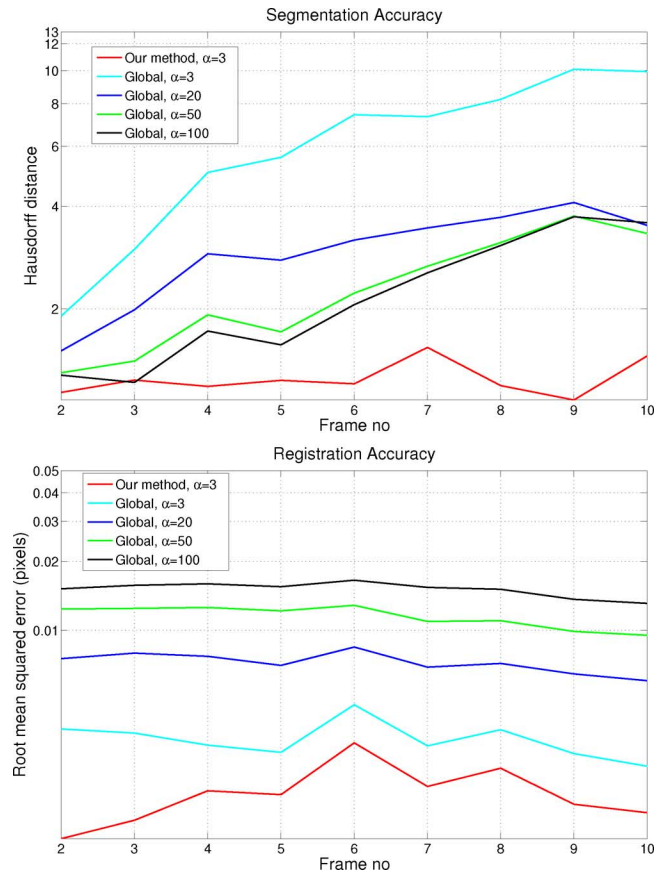

Fig. 5. Global regularization versus physically motivated tracking on synthetic sequence: Quantitative evaluation. [Top]: Graph of Hausdorff distance (lower indicates better fit to ground truth) between the ground truth segmentation and the estimated contour using regularization $\alpha = \{3, 20, 50, 100\}$ and our method, which yields the same curve for a wide range of parameters. [Bottom]: Graph of registration accuracy as measured by root mean squared difference between the deformed image and the next image. Lower rms indicates better fit to ground truth.

A. Synthetic Experiment

We start by verifying our technique on a synthetic sequence designed to mimic the piecewise deformation with matching normals for which our technique is designed. We consider a sequence composed of images with two textures (one for the

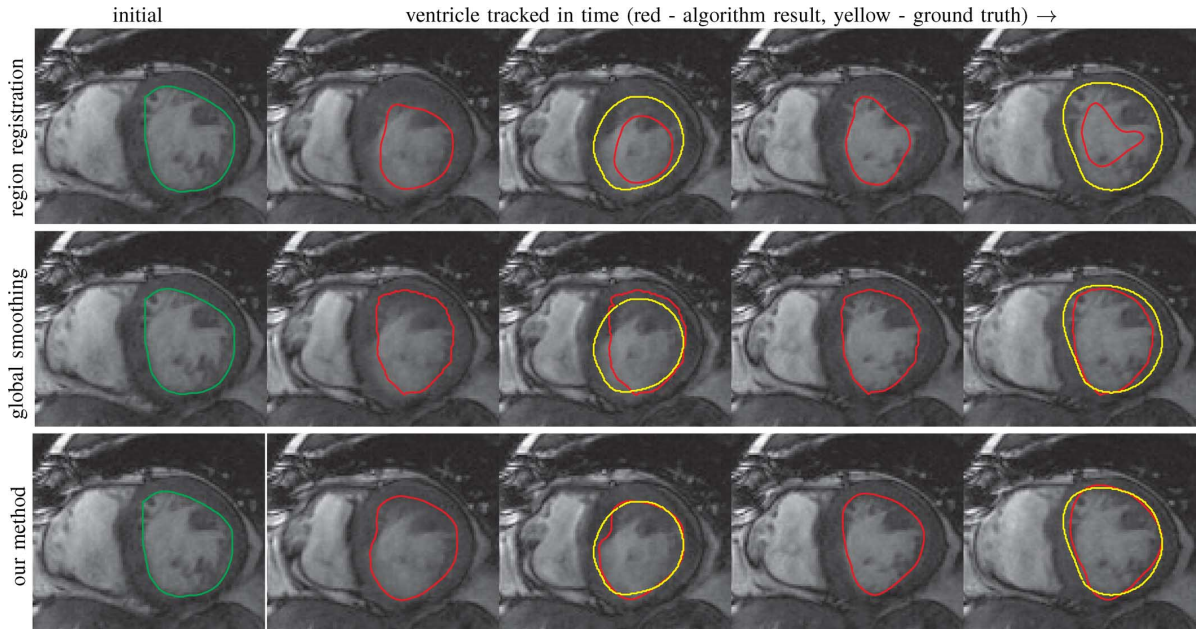

Fig. 6. Illustrative differences of registration schemes on the LV. [Top]: registering only the inside of the ventricle leads to inaccurate segmentation in subsequent frames. [Middle]: registering the whole image with global regularization (LDDMM) smooths motion across different structures and leads to inaccurate segmentation. [Bottom]: registering the entire image with proposed technique leads to the most accurate segmentation. Green: initialization, red: algorithm result, yellow: ground truth.

object, and the other for the background). Textures are needed so that registration can be determined. The region of interest is the disc, and it along with the background contracts. In addition to the contraction, there is small rotational motions of the disc and the background in opposite directions. The sequence is constructed so that the normal component across the boundary matches in both regions. This causes the true flow to be non-smooth across the boundary. Ground truth registration between consecutive images is known.

The first row of Fig. 3 shows the first two frames of the synthetic sequence, the registration color code (color indicates direction and intensity indicates magnitude) and the ground truth registration for the first two frames. The second row shows registrations computed by using (32)–(35) with Horn–Schunck optical flow (global regularization) with smoothness $\alpha = \{3, 20, 50\}$ in (32), and the proposed method (hard constraint) with $\alpha_i = \alpha_o = 3$. Global regularization smooths across the boundary, mixing inhomogeneous motions, leading to an inaccurate registration. Our proposed registration, which does not smooth across the boundary while satisfying physical constraints, is able to accurately recover the registration.

Fig. 4 displays the results of tracking the whole synthetic sequence of 10 frames (only four are shown) with global regularization and our method. The first three rows display the result of tracking using global regularization ($\alpha = \{3, 20, 50\}$). Notice that global regularization of the deformation with small regularization ($\alpha = 3$) leads to less smoothing across the boundary, but an inaccurate segmentation due to small regularization which traps the contour in small scale structures. Larger regularization smooths more across the boundary leading to an inaccurate registration. The segmentation improves, but the boundary is still not captured accurately. No amount of global regularization is able to detect an accurate boundary. Finally, our method (last row), which smooths within regions while simultaneously satisfying the normal matching constraint is able to capture both an accurate registration and segmentation, and obtains similar results for a wide range of regularity parameters. Fig. 5 quantifies the tracking results. We include both the segmentation accuracy (left) as measured by Hausdorff measure (HD) to ground truth, and the accuracy of the warped image under the registration (right) to the next image using root mean square error. Both verify the visual results.

B. Datasets for Ventricle Segmentation

The evaluations in the next sub-sections were carried out on public data sets, the MICCAI Left Ventricle Dataset [41] and the MICCAI Right Ventricle Dataset [28]. The validation and training dataset from [41] consists of 15 sets of cardiac cine-MRI images each. Each set contains images obtained during 10–15 s breath-holds with a temporal resolution of 20 cardiac phases over the heart cycle. The ground truth is provided for two out of 20 images for each slice and each set has eight to 20 slices. Similarly, the training and test data set from [28] contains 16 sets of cardiac cine-MRI images each, and 16 test sets (Test1Set). Each set contains about 10 slices. There are 20 images per cardiac cycle.

C. Comparison of Hard and Soft Constraints

This experiment on cardiac data shows the differences between the hard and soft formulations for enforcing matching normal velocities at the interface and the No-Slip condition. As stated in Section III-C, the soft formulation is more general than

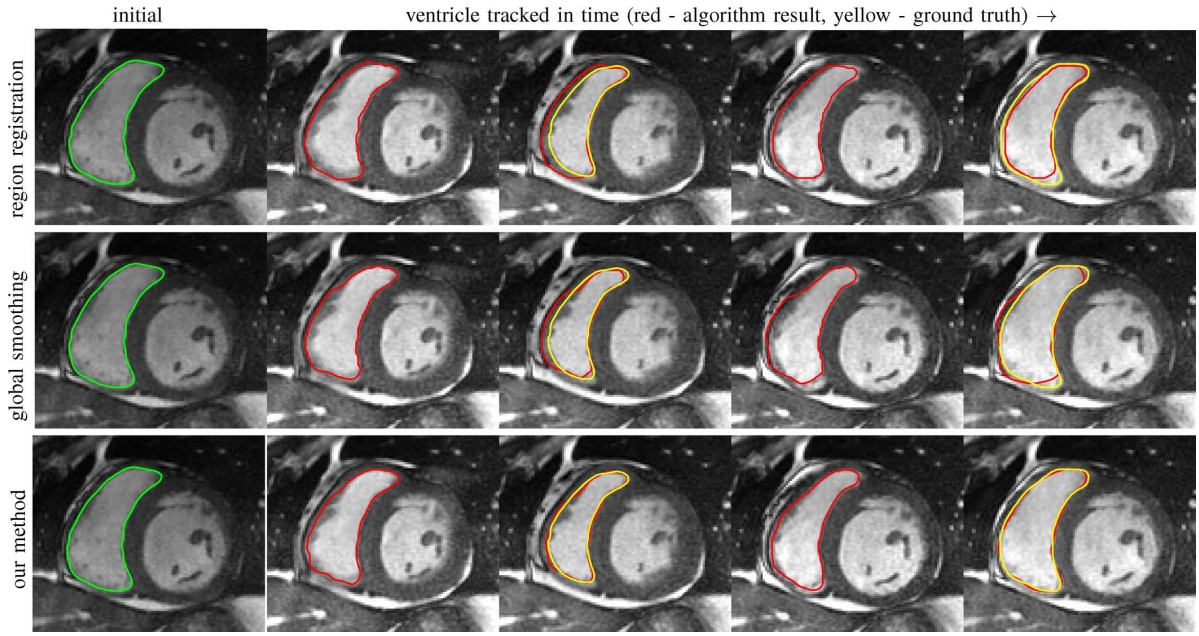

Fig. 7. Illustrative differences of registration schemes on the RV. [Top]: registering only the inside of the ventricle leads to inaccurate segmentation in subsequent frames. [Middle]: registering the whole image with global regularization (LDDMM) smooths motion across different structures and leads to inaccurate segmentation. [Bottom]: registering the entire image with proposed technique leads to the most accurate segmentation. Green: initialization, red: algorithm result, yellow: ground truth.

TABLE II
QUANTITATIVE COMPARISON TO MEDVISO AND LDDMM (GLOBAL REGULARIZATION). RESULTS ARE ON THE MICCAI LV VALIDATION [41] AND MICCAI RV TRAINING [28] DATASETS. LOW APD/HD AND HIGH DM INDICATE GOOD MATCHES

	MICCAI LV		MICCAI RV	
mean ± std	APD	DM	HD	DM
Ours	2.39± .31	.88±.02	6.72 ± 3.22	.83±.15
LDDMM	2.48± .35	.87±.02	7.15 ± 3.64	.81±.14
MedViso	4.68± 1.2	.78±.09	15.19 ± 6.08	.73±.19

the hard formulation (specific values of β and γ in the soft formulation lead to the hard formulation). Quantitative assessment to ground truth is performed by using the APD (average perpendicular distance) and DM (dice metric). Table I shows the results on the MICCAI LV validation set [41] using the hard formulation (No-Slip condition automatically enforced), the soft formulation without the No-Slip condition ($\beta = 2, \gamma = 0$), and the soft formulation with the No-Slip condition implemented ($\beta = 2, \gamma = 2$). We have tuned β based on the training set and used this to evaluate on the test set. The soft formulation without the No-Slip condition enforced performs the least accurate, followed by the hard formulation, and finally the most accurate is the soft formulation with the No-Slip condition implemented. Although the soft formulation with the No-slip condition can be tuned to obtain more superior results than the hard formulation for a single image, as this experiment shows, that value of β may not lead to superior performance on the entire dataset. The experiment also shows that the No-Slip condition, implicit in the hard formulation, and explicit in the soft formulation leads to better performance. In the next experiments, to avoid choosing the parameters β and γ, we use the hard formulation.

TABLE III
QUANTITATIVE COMPARISON TO METHODS FROM MICCAI CHALLENGES. EVALUATION ON THE MICCAI LV VALIDATION+TRAINING [41] AND MICCAI RV TEST1SET [28] DATASETS. LOW APD/HD AND HIGH DM INDICATE GOOD MATCHES. OUR METHOD DOES NOT MAKE USE OF TRAINING DATA IN CONTRAST TO OTHER METHODS, BUT IT DOES REQUIRE INITIALIZATION IN THE FIRST FRAME, WHICH OTHER METHODS DO NOT REQUIRE

	MICCAI 2009 LV	
mean ± std	APD	DM
Ours	2.24± .41	.89 ± .02
Marak et al.	3.00 ± .59	.86 ± .04
Constantinides et al.	2.04 ±.47	.89 ± .04
Jolly et al.	2.26 ± .59	.88 ± .04
Huang et al.	2.06 ± .39	.89 ± .04
Brien et al.	3.73	-

	MICCAI 2012 RV	
mean ± std	HD	DM
Ours	8.61 ± 3.72	.73 ± .18
Team 1	11.16 ± 5.53	.69 ± .25
Team 2	20.44 ± 17.8	.53 ± .32
Team 3	10.63 ± 5.59	.69 ± .22
Team 4	23.19 ± 9.71	.48 ± .25
Team 5	29.30 ± 23.33	.48 ± .33
Team 6	11.41 ± 10.49	.72 ± .27
Team 7	14.49 ± 7.56	.71 ± .23

D. LV and RV: Comparison of Three Registration Schemes

In this experiment, we compare registration methods used for segmentation of the LV and RV. We visually compare the tracking results given by our method to (M1) registration of only the interior of current estimate of the ventricle to a subset of next image (to show whole image registration is needed), and

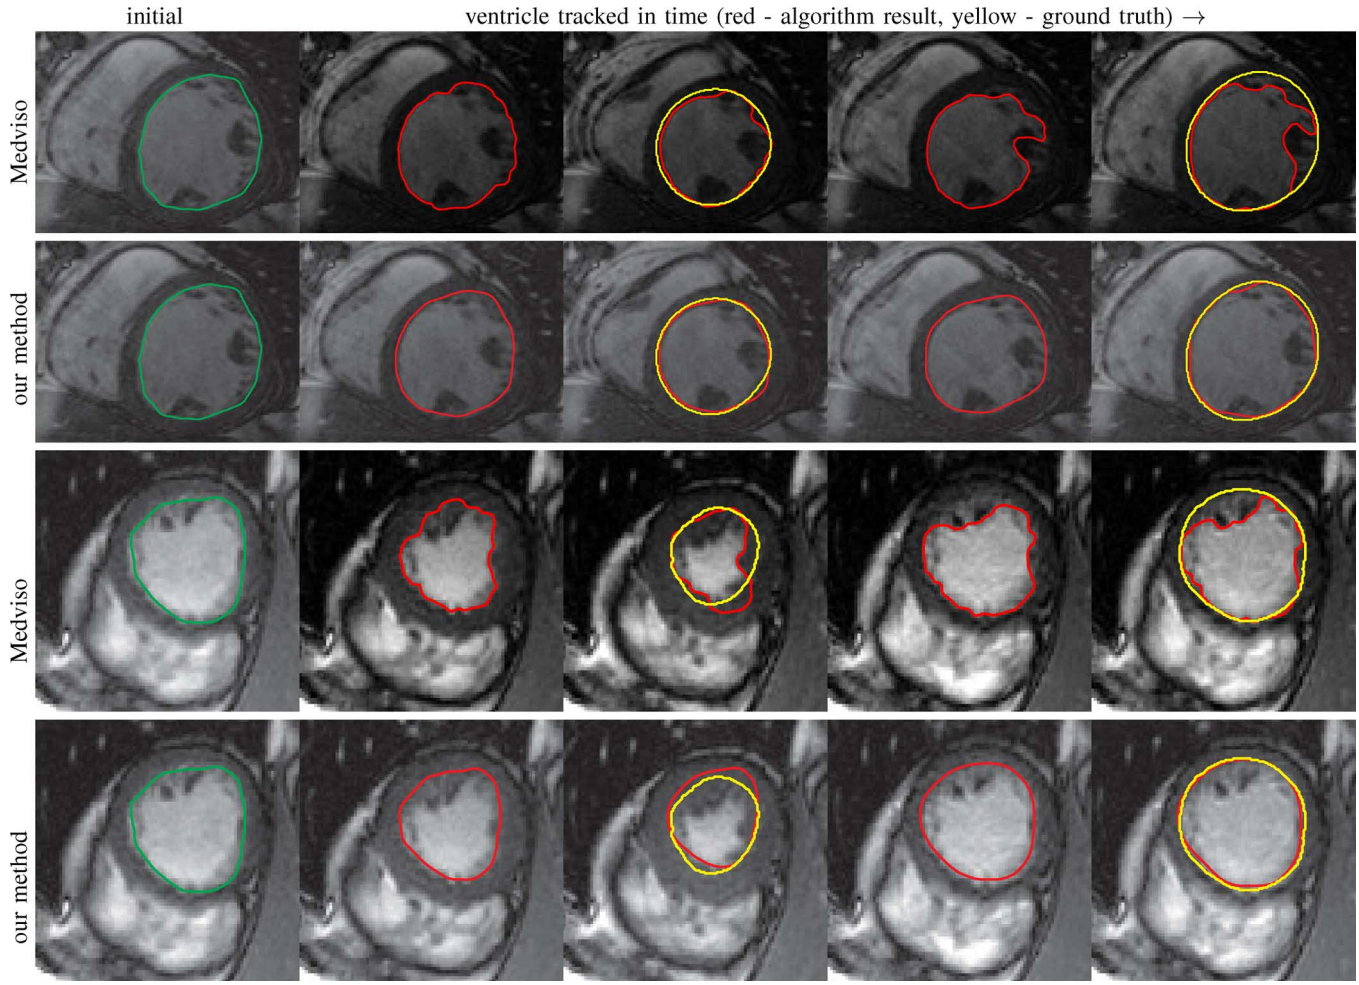

Fig. 8. Comparison on tracking the left ventricle. Two sample results on the MICCAI LV Dataset [41] of full cardiac cycles (only five out of 20 images are shown) of the proposed algorithm and Medviso. Ground truth when available is superimposed in yellow, and the red contour is the result obtained by the indicated algorithm. Visual results indicate that our algorithm is better able to handle inhomogeneous appearance, and is thus more accurate.

to (M2) full image registration with global regularization (we use LDDMM [2] to ensure an advanced global regularization method). M1 is achieved by computing just the inside velocity with Neumann boundary conditions on ∂R_τ. The best results with respect to ground truth are chosen by choosing the optimal smoothness parameter in all methods. Results on tracking a full cardiac cycle are given in Fig. 6 for the LV and Fig. 7 for the RV. Registering only the organ (M1) results in errors (as the background registration is helpful in restricting undesirable registrations of the foreground). LDDMM registration (M2) smooths motion from irrelevant background structures into the ventricles, which results in drifting from the desired boundary. Our method, which smooths within regions while satisfying physical constraints, achieves the most accurate results.

E. LV and RV Segmentation: Quantitative Comparison

We show experiments verifying that our algorithm improves the prediction step of interactive segmentation methods. We show that less interaction is needed with our approach than with a recent commercial cardiac segmentation software, Segment from Medviso [18], [49]. We also compare it to LDDMM to show that the physically motivated constraints lead to superior results quantitatively. Both methods start with the same initial hand segmentation, and subsequent frames are segmented via propagation. No manual interaction is used as we wish to show that our method would require less interaction. The regularity parameters $\alpha_i = \alpha_o = 3$ in our method are found by choosing them so that the results are closest to ground truth in a few training cases. The same parameter is then used for all other cases. LDDMM uses regularity parameter ($\alpha = 7$), which is similarly tuned.

Figs. 8 and 9 shows some sample tracking results of the proposed method and Medviso on full cardiac cycles of two different cases on both the LV dataset and the RV dataset. The ground truth (yellow) is superimposed when available. A summary of the results on the entire datasets is shown in Table II, and results of LDDMM (global regularization) are also provided. The accuracy with respect to ground truth is measured using average perpendicular distance (APD) and dice metric (DM) for left ventricle, and Hausdorff distance (HD) and DM for the right ventricle. These metrics are chosen since they are the standard ones used on these datasets. Both qualitative and quantitative

TABLE IV

QUANTITATIVE ANALYSIS OF PERIODIC CONSISTENCY. QUANTIFICATION OF DIFFERENCE BETWEEN THE INITIAL SEGMENTATION AND THE SEGMENTATION FROM THE ALGORITHM INDICATED AFTER ONE CARDIAC CYCLE. EVALUATION ON THE MICCAI LV TRAINING DATASET. LOWER APD AND HIGHER DM INDICATES CLOSER MATCH OF INITIAL AND FINAL SEGMENTATIONS

	Epicardium		RV		LV	
mean $\pm$ std.	APD	DM	APD	DM	APD	DM
Ours	1.40 ± 0.22	0.97 ± 0.01	1.69 ± 0.18	0.95 ± 0.02	1.05 ± 0.27	0.97 ± 0.02
Medviso	9.19 ± 0.89	0.82 ± 0.05	3.65 ± 1.95	0.90 ± 0.05	4.86 ± 2.42	0.82 ± 0.05

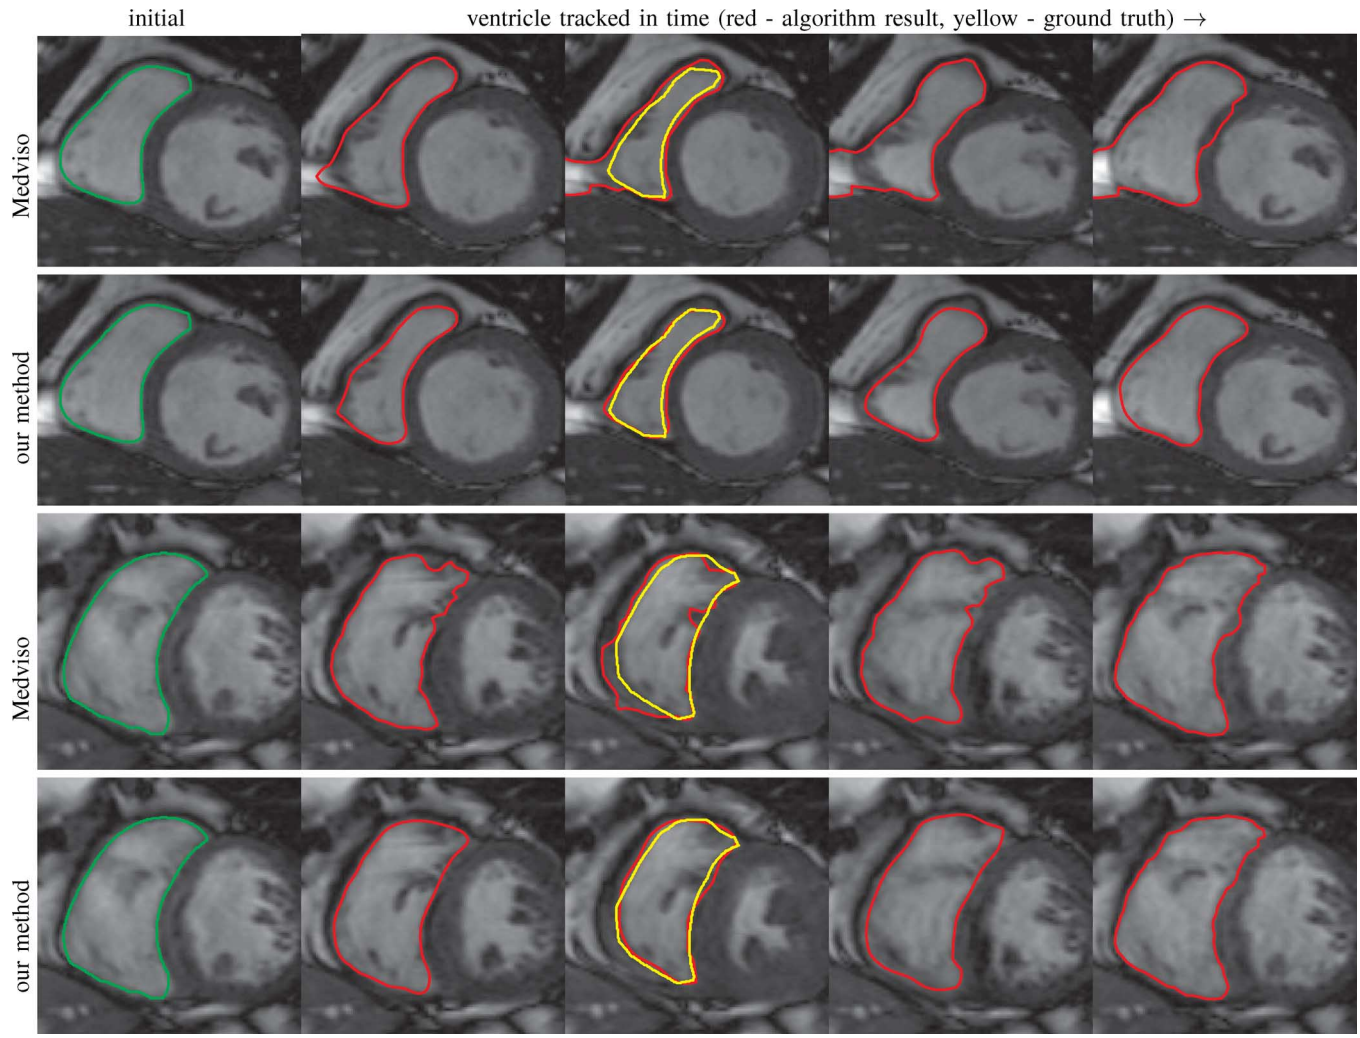

Fig. 9. Comparison on tracking the right ventricle. Two sample results on the MICCAI RV Dataset [28] of full cardiac cycles (only five out of 20 images are shown) of the proposed algorithm and Medviso. Ground truth when available is superimposed in yellow (not available on last frame), and the red contour is the result obtained by the indicated algorithm. Our method is less susceptible to clutter nearby the RV, and thus is able to capture the RV more accurately.

results show that our proposed method leads to more accurate segmentation of the ventricles than Medviso and LDDMM, and thus less interaction than segmentation propagation in Medviso.

We also compare the results with other methods in the MICCAI 2009 LV and MICCAI 2012 RV challenges. Table III shows the results. Other methods use training data or an atlas, while our method requires no training data. Our method does require an initial segmentation from the first frame, which other methods reported on the challenges do not require. For MICCAI 2009 LV dataset, we used training and validation dataset (30 sets). For MICCAI 2012 RV we used Test1Set (16 sets) for evaluation. The algorithm results were sent to the RVSC organizers for blind statistical evaluation since ground truth was not provided. Results are compared with MICCAI 2012 RVSC participating teams as reported on RVSC website. We compare only the end systole (ES) results as we used the end diastole (ED) for initialization. Results on the MICCAI LV indicate that our method performs well (third best), and our method performs the best on the MICCAI RV dataset.

F. Multiple Region Segmentation: Full Heart Segmentation

We now demonstrate our approach in performing challenging full heart segmentation: segmentation of the ventricles and epicardium all at once. Both the RV and epicardium are especially

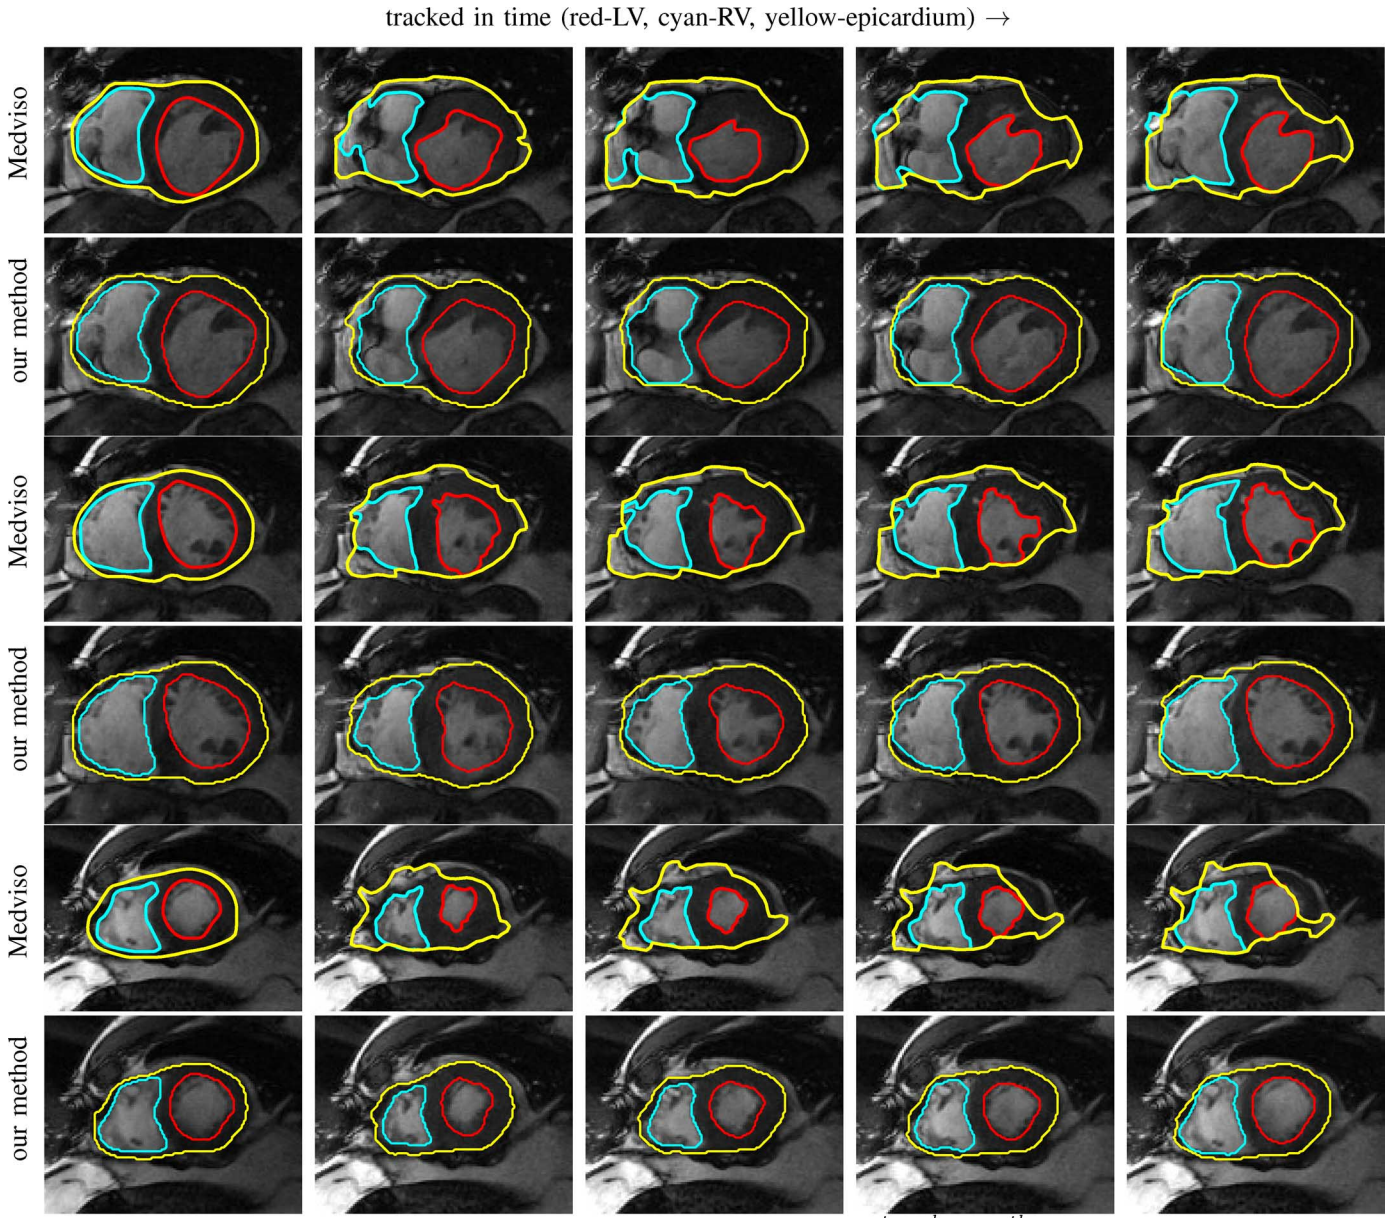

Fig. 10. Comparison on multiple region segmentation. [Top to Bottom]: first, third, and sixth slices shown. The LV (red), RV (cyan), and myocardium outer boundary (yellow) are simultaneously segmented using our proposed technique. Comparison is shown to Medviso. Visual results indicate that our technique is more accurate in segmenting all structures.

challenging as the contrast of the RV and background is subtle in comparison to the LV, and the myocardium wall near parts of the RV is very thin. We are not aware of another interactive method that is able to segment all structures, and so we compare to Medviso even though the method is not specifically tailored to the myocardium, but the method is generic and is able to propagate a segmentation. Medviso does not segment multiple regions all at once and thus we perform separate segmentation of the LV, RV, and epicardium.

Fig. 10 shows the slice-wise results of our method and Medviso on a full 3-D cardiac MRI sequence for a full cardiac cycle. Results indicate that our method is more accurate in capturing the shape of the ventricles and epicardium. Medviso gives segmentations that stray far from the epicardium, and sometimes physically impossible configurations (intersections of the RV and epicardium). Fig. 11 shows visualization of the results in 3-D, and that our method more accurately resembles the structure of the heart.

Ground truth is not available for the outer wall of the myocardium in any standard dataset to the best of our knowledge. Thus, to perform quantitative analysis, we evaluate periodic consistency. After a complete cardiac cycle, the position of the epicardium is approximately the same. Thus, we quantify the differences of the initial segmentation of the epicardium from the segmentation after one cardiac cycle. We also include LV and RV to quantify the periodic consistency. Table IV gives a summary of the results on 3-D data from the MICCAI RV Dataset. Our method is more accurate than Medviso.

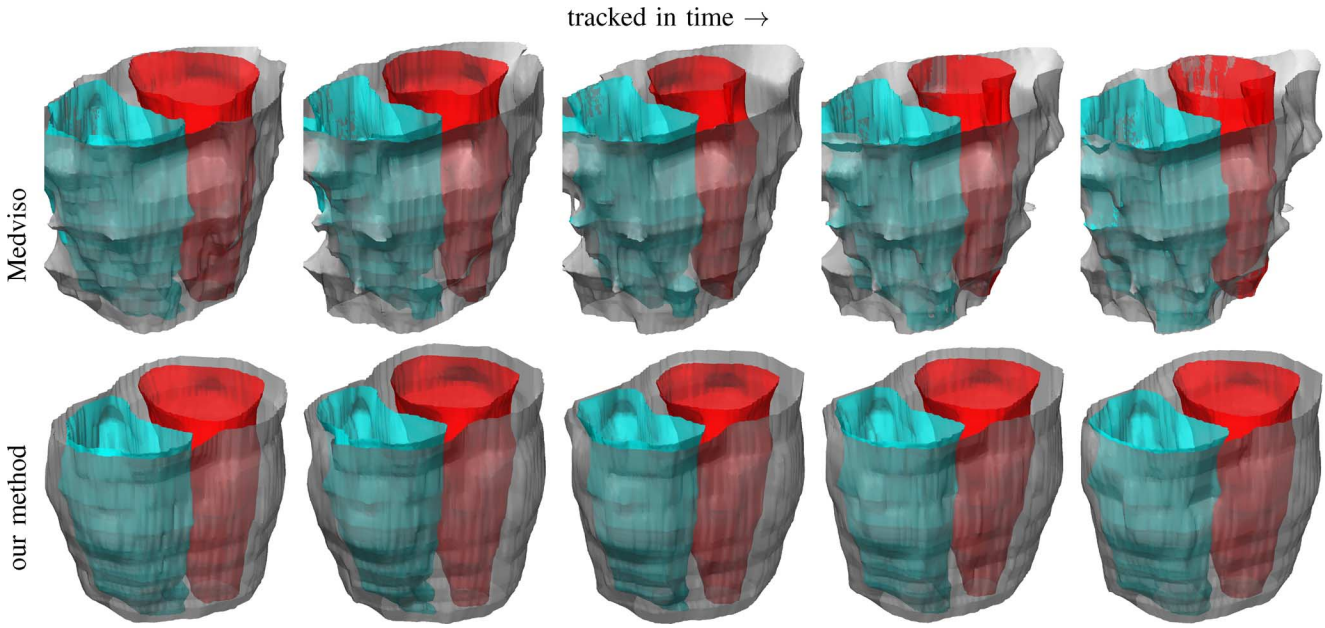

Fig. 11. Comparison on Multiple Region Segmentation (3-D Visualization). [Top row]: Medviso result, [Bottom row]: our result. Grey: Myocardium, cyan: right ventricle, red: left ventricle. Only 5 of 20 frames shown. Results indicate that the output of our method better resembles heart structure.

VII. Conclusion

We have presented an algorithm for propagating the segmentation from one frame in an image sequence to another via a novel motion estimation algorithm. The registration is physically motivated by the heterogeneous motions among substructures and the physical constraints between motions in adjacent regions, specifically the matching condition of normal velocities and the No-Slip condition at the interface. Traditional registration algorithms apply regularization that smooths across boundaries, mixing motions of differing substructures, do not incorporate the aforementioned physical constraints, and therefore yield inaccurate registrations and hence segmentation propagations. Our motion estimation technique is computationally efficient and has nearly the same cost as traditional Horn–Schunck optical flow.

Experiments have shown that our method is more effective than global regularization in propagating segmentations in cardiac MRI data of the heart. Moreover, we improved the propagation step in interactive segmentation techniques, which are used commercially for cardiac MRI segmentation. We have compared our technique both qualitatively and quantitatively against a recent commercial software, Medviso, and results indicate that our method would require less manual interaction, specifically in LV, RV, and epicardium segmentation.

References

[1] W. Bai, W. Shi, H. Wang, N. S. Peters, and D. Rueckert, "Multi-atlas based segmentation with local label fusion for right ventricle MR images," *Image*, vol. 6, p. 9, 2012.

[2] M. Beg, M. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. Comput. Vis.*, vol. 61, no. 2, pp. 139–157, 2005.

[3] Y. Boykov and M. Jolly, "Interactive organ segmentation using graph cuts," in *Proc. MICCAI*, 2000, pp. 147–175.

[4] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1222–1239, Nov. 2001.

[5] X. Bresson, S. Esedoáÿalu, P. Vandergheynst, J.-P. Thiran, and S. Osher, "Fast global minimization of the active contour/snake model," *J. Math. Imag. Vis.*, vol. 28, no. 2, pp. 151–167, 2007.

[6] Y. Chen, H. D. Tagare, S. Thiruvenkadam, F. Huang, D. Wilson, K. S. Gopinath, R. W. Briggs, and E. A. Geiser, "Using prior shapes in geometric active contours in a variational framework," *Int. J. Comput. Vis.*, vol. 50, no. 3, pp. 315–328, 2002.

[7] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Process.*, vol. 5, no. 10, pp. 1435–1447, Oct. 1996.

[8] T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 681–685, Jun. 2001.

[9] T. Cootes *et al.*, "Active shape models-their training and application," *Comput. Vis. Image Understand.*, vol. 61, no. 1, pp. 38–59, 1995.

[10] L. Cordero-Grande, G. Vegas-Sánchez-Ferrero, P. Casaseca-de-la Higuera, J. A. San-Román-Calvar, A. Revilla-Orodea, M. Martín-Fernández, and C. Alberola-López, "Unsupervised 4D myocardium segmentation with a Markov random field based deformable model," *Med. Image Anal.*, vol. 15, no. 3, pp. 283–301, 2011.

[11] D. Cremers, "Dynamical statistical shape priors for level set-based tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 8, pp. 1262–1273, Aug. 2006.

[12] D. Cremers, O. Fluck, M. Rousson, and S. Aharon, "A probabilistic level set formulation for interactive organ segmentation," in *Med. Imag.*, 2007, pp. 65120V–65120V.

[13] D. Cremers, T. Kohlberger, and C. Schnörr, "Shape statistics in kernel space for variational image segmentation," *Pattern Recognit.*, vol. 36, no. 9, pp. 1929–1943, 2003.

[14] D. Cremers and S. Soatto, "Motion competition: A variational approach to piecewise parametric motion segmentation," *Int. J. Comput. Vis.*, vol. 62, no. 3, pp. 249–265, 2005.

[15] D. G. Ebin and J. Marsden, "Groups of diffeomorphisms and the motion of an incompressible fluid," *Ann. Math.*, vol. 92, no. 1, pp. 102–163, 1970.

[16] O. Ecabert *et al.*, "Automatic model-based segmentation of the heart in CT images," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1189–1201, Sep. 2008.

[17] X. Han, C. Xu, and J. L. Prince, "A topology preserving level set method for geometric deformable models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 6, pp. 755–768, 2003.

[18] E. Heiberg, J. Sjögren, M. Ugander, M. Carlsson, H. Engblom, and H. Arheden, "Design and validation of segment-freely available software for cardiovascular image analysis," *BMC Med. Imag.*, vol. 10, no. 1, p. 1, 2010.

[19] T. Heimann and H.-P. Meinzer, "Statistical shape models for 3-D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, no. 4, pp. 543–563, 2009.

[20] B. K. Horn and B. G. Schunck, "Determining optical flow," *Artif. Intell.*, vol. 17, no. 1, pp. 185–203, 1981.

[21] I. Isgum, M. Staring, A. Rutten, M. Prokop, M. Viergever, and B. van Ginneken, "Multi-atlas-based segmentation with local decision fusion–application to cardiac and aortic segmentation in CT scans," *IEEE Trans. Med. Imag.*, vol. 28, no. 7, pp. 1000–1010, Jul. 2009.

[22] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1988.

[23] M. Kaus, J. Berg, J. Weese, W. Niessen, and V. Pekar, "Automated segmentation of the left ventricle in cardiac MRI," *Med. Image Anal.*, vol. 8, no. 3, pp. 245–254, 2004.

[24] H. A. Kirisli, M. Schaap, S. Klein, L. A. Neefjes, A. C. Weustink, T. Walsum, and W. J. Niessen, "Fully automatic cardiac segmentation from 3-D CTA data: A multi-atlas based approach," in *Proc. SPIE*, 2010, vol. 7623, pp. 762305–1.

[25] J. Koikkalainen, T. Tolli, K. Lauerma, K. Antila, E. Mattila, M. Lilja, and J. Lotjonen, "Methods of artificial enlargement of the training set for statistical shape models," *IEEE Trans. Med. Imag.*, vol. 27, no. 11, pp. 1643–1654, Nov. 2008.

[26] A. Lalande, L. Legrand, P. M. Walker, F. Guy, Y. Cottin, S. Roy, and F. Brunotte, "Automatic detection of left ventricular contours from cardiac cine magnetic resonance imaging using fuzzy logic," *Invest. Radiol.*, vol. 34, no. 3, pp. 211–217, 1999.

[27] X. Lin, B. R. Cowan, and A. A. Young, "Automated detection of left ventricle in 4D MR images: Experience from a large study," in *Proc. MICCAI*, 2006, pp. 728–735.

[28] RV segmentation challenge in cardiac MRI Litis, 2012.

[29] M. Lorenzo-Valdes, G. Sanchez-Ortiz, R. Mohiaddin, and D. Rueckert, "Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm," in *Proc. MICCAI*, 2004, vol. 1, p. 440.

[30] J. Lötjönen, S. Kivistö, J. Koikkalainen, D. Smutek, and K. Lauerma, "Statistical shape model of atria, ventricles and epicardium from short-and long-axis MR images," *Med. Image Anal.*, vol. 8, no. 3, pp. 371–386, 2004.

[31] M. Lynch, O. Ghita, and P. F. Whelan, "Automatic segmentation of the left ventricle cavity and myocardium in MRI data," *Comput. Biol. Med.*, vol. 36, no. 4, pp. 389–407, 2006.

[32] B. R. Munson, D. F. Young, and T. H. Okiishi*, Fundamentals of Fluid Mechanics*. New York: Wiley, 1990.

[33] N. M. Noble, D. L. Hill, M. Breeuwer, J. A. Schnabel, D. J. Hawkes, F. A. Gerritsen, and R. Razavi, "Myocardial delineation via registration in a polar coordinate system," in *Proc. MICCAI*, 2002, pp. 651–658.

[34] S. Osher and J. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. Comput. Phys.*, vol. 79, no. 1, pp. 12–49, 1988.

[35] D. F. Pace, A. Enquobahrie, H. Yang, S. R. Aylward, and M. Niethammer, "Deformable image registration of sliding organs using anisotropic diffusive regularization," in *Proc. Int. Symp. Biomed. Imag.: From Nano to Macro*, 2011, pp. 407–413.

[36] D. F. Pace, M. Niethammer, and S. R. Aylward, "Sliding geometries in deformable image registration," in *Abdominal Imaging: Computational and Clinical Applications*. New York: Springer, 2012, pp. 141–148.

[37] N. Paragios, "A variational approach for the segmentation of the left ventricle in cardiac image analysis," *Int. J. Comput. Vis.*, vol. 50, no. 3, pp. 345–362, 2002.

[38] N. Paragios, "A level set approach for shape-driven segmentation and tracking of the left ventricle," *IEEE Trans. Med. Imag.*, vol. 22, no. 6, pp. 773–776, Jun. 2003.

[39] C. Petitjean and J. Dacher, "A review of segmentation methods in short axis cardiac MR images," *Med. Image Anal.*, vol. 15, no. 2, 2011.

[40] Q. Pham, F. Vincent, P. Clarysse, P. Croisille, and I. Magnin, "A FEM-based deformable model for the 3-D segmentation and tracking of the heart in cardiac MRI," in *Proc. 2nd Int. Symp. Image Signal Process. Anal.*, 2001, pp. 250–254.

[41] P. Radau, Y. Lu, K. Connelly, G. Paul, A. Dick, and G. Wright, "Evaluation framework for algorithms segmenting short axis cardiac MRI," *Midas J.*, 2009.

[42] L. Risser, H. Baluwala, and J. A. Schnabel, "Diffeomorphic registration with sliding conditions: Application to the registration of lungs CT images," in *Proc. 4th Int. Workshop Pulmonary Image Anal. MICCAI*, 2011, pp. 79–90.

[43] L. Risser, F.-X. Vialard, H. Y. Baluwala, and J. A. Schnabel, "Piecewise-diffeomorphic image registration: Application to the motion estimation between 3-D CT lung images with sliding conditions," *Med. Image Anal.*, vol. 17, no. 2, pp. 182–193, 2013.

[44] J. Schaerer, C. Casta, J. Pousin, and P. Clarysse, "A dynamic elastic model for segmentation and tracking of the heart in MR image sequences," *Med. Image Anal.*, vol. 14, no. 6, pp. 738–749, 2010.

[45] A. Schmidt-Richberg, J. Ehrhardt, R. Werner, and H. Handels, "Fast explicit diffusion for registration with direction-dependent regularization," *Biomed. Image Registrat.*, pp. 220–228, 2012.

[46] A. Schmidt-Richberg, R. Werner, H. Handels, and J. Ehrhardt, "Estimation of slipping organ motion by registration with direction-dependent regularization," *Med. Image Anal.*, vol. 16, no. 1, pp. 150–159, 2012.

[47] J. A. Schnabel *et al.*, "A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations," in *Proc. MICCAI*, 2001, pp. 573–581.

[48] J. Senegas, C. A. Cocosco, and T. Netsch, "Model-based segmentation of cardiac MRI cine sequences: A Bayesian formulation," in *Proc. SPIE Med. Imag.*, 2004, pp. 432–443.

[49] J. Sjögren *et al.*, "Semi-automatic segmentation of myocardium at risk in T2-weighted cardiovascular magnetic resonance," *J. Cardiovasc. Magn. Reson.*, vol. 14, no. 10, 2012.

[50] M. Stegmann, H. Ólafsdóttir, and H. Larsson, "Unsupervised motion-compensation of multi-slice cardiac perfusion MRI," *Med. Image Anal.*, vol. 9, no. 4, pp. 394–410, 2005.

[51] A. Tsai *et al.*, "A shape-based approach to the segmentation of medical imagery using level sets," *IEEE Trans. Med. Imag.*, vol. 22, no. 2, pp. 137–154, Feb. 2003.

[52] H. Van Assen, M. Danilouchkine, A. Frangi, S. Ord s, J. Westenberg, J. Reiber, and B. Lelieveldt, "SPASM: A 3-D-ASM for segmentation of sparse and arbitrarily oriented cardiac MRI data," *Med. Image Anal.*, vol. 10, no. 2, pp. 286–303, 2006.

[53] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Symmetric log-domain diffeomorphic registration: A demons-based approach," in *Proc. MICCAI*, 2008, pp. 754–761.

[54] J. Wijnhout, D. Hendriksen, H. Assen, and R. Geest, "LV challenge LKEB contribution: Fully automated myocardial contour detection," in *MICCAI Workshop Cardiac MR LV Segment. Challenge*, London, U.K., 2009, p. 683.

[55] X. Xue, "Interactive 3-D heart chamber partitioning with a new marker-controlled watershed algorithm," in *Adv. Vis. Comput.*, G. Bebis, R. Boyle, D. Koracin, and B. Parvin, Eds. Berlin, Germany: Springer, 2005, vol. 3804, Lecture Notes in Computer Science, pp. 92–99.

[56] H. Zhang, A. Wahle, R. Johnson, T. Scholz, and M. Sonka, "4-D cardiac MR image analysis: Left and right ventricular morphology and function," *IEEE Trans. Med. mag.*, vol. 29, no. 2, pp. 350–364, Feb. 2010.

[57] S. Zhang, Y. Zhan, M. Dewan, J. Huang, D. Metaxas, and X. Zhou, "Deformable segmentation via sparse shape representation," in *Proc. MICCAI*, 2011, pp. 451–458.

[58] Y. Zhu, X. Papademetris, A. J. Sinusas, and J. S. Duncan, "Segmentation of the left ventricle from cardiac MR images using a subject-specific dynamical model," *IEEE Trans. Med. Imag.*, vol. 29, no. 3, pp. 669–687, Mar. 2010.

[59] X. Zhuang, Challenges and methodologies of fully automatic whole heart segmentation: A review.

[60] X. Zhuang, K. Rhode, R. Razavi, D. Hawkes, and S. Ourselin, "A registration-based propagation framework for automatic whole heart segmentation of cardiac MRI," *IEEE Trans. Med. Imag.*, vol. 29, no. 9, pp. 1612–1625, Sep. 2010.

[61] M. Zuluaga, M. Cardoso, and S. Ourselin, "Automatic right ventricle segmentation using multi-label fusion in cardiac MRI," presented at the MICCAI Workshop RV Segmentat. Challenge Cardiac MRI, Nice, France, 2012.