



HAL
open science

Model-based generation of large databases of cardiac images: synthesis of pathological cine MR sequences from real healthy cases

Nicolas Duchateau, Maxime Sermesant, Hervé Delingette, Nicholas Ayache

► To cite this version:

Nicolas Duchateau, Maxime Sermesant, Hervé Delingette, Nicholas Ayache. Model-based generation of large databases of cardiac images: synthesis of pathological cine MR sequences from real healthy cases. *IEEE Transactions on Medical Imaging*, 2018, 37, pp.755-766. 10.1109/TMI.2017.2714343 . hal-01533788

HAL Id: hal-01533788

<https://inria.hal.science/hal-01533788v1>

Submitted on 6 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Model-based generation of large databases of cardiac images: synthesis of pathological cine MR sequences from real healthy cases

Nicolas Duchateau, Maxime Sermesant, Hervé Delingette, Nicholas Ayache

Abstract—Collecting large databases of annotated medical images is crucial for the validation and testing of feature extraction, statistical analysis and machine learning algorithms. Recent advances in cardiac electromechanical modeling and image synthesis provided a framework to generate synthetic images based on realistic mesh simulations. Nonetheless, their potential to augment an existing database with large amounts of synthetic cases requires further investigation. We build upon these works and propose a revised scheme for synthesizing pathological cardiac sequences from real healthy sequences. Our new pipeline notably involves a much easier registration problem to reduce potential artifacts, and takes advantage of mesh correspondences to generate new data from a given case without additional registration. The output sequences are thoroughly examined in terms of quality and usability on a given application: the assessment of myocardial viability, via the generation of 465 synthetic cine MR sequences (15 healthy and 450 with pathological tissue viability [random location, extent and grade, up to myocardial infarct]). We demonstrate that our methodology (i) improves state-of-the-art algorithms in terms of realism and accuracy of the simulated images, and (ii) is well-suited for the generation of large databases at small computational cost.

Index Terms—Image synthesis and simulation, large database for machine learning, validation, myocardial infarct, biomechanical modeling, magnetic resonance.

I. INTRODUCTION

Despite decades of designing algorithms for the analysis of medical images, the collection of proper databases of annotated medical images remains crucial to validate or just correctly use such algorithms.

On the one hand, standard feature extraction methods from single images/sequences such as segmentation and registration require ground truth imaging data for validation purposes. This validation relies on manual annotations by experts, invasive tracers attached to the anatomy, or phantom data. However, these processes are cumbersome and have limited accuracy or realism. New types of ground truth data are therefore required.

On the other hand, data-based statistical analysis and machine learning algorithms are optimally exploited with large sample size. Computer vision research now provides access to datasets of 100,000+ 2D images¹. In contrast, collecting

relevant sets of medical images undergoes limitations from healthcare data management (permission and storage size) [1] and acquisition (image quality and availability) [2]. In our field of interest—cardiac imaging, and especially motion and deformation analysis—current datasets with 3D sequences and coherent pathophysiological condition comprise around 1000 subjects [3]–[5]. The recent UK Biobank project intends to gather an unprecedented database of 100,000 cardiac sequences [6]. Nonetheless, it targets an outpatient or normal population, which means that the study of specific pathologies in coherent subgroups might still have limited sample size or be subject to a large variety of confounding factors [7].

Thus, it is urgent to develop alternative approaches to collect (i) large datasets with (ii) coherent clinical conditions and (iii) ground truth data. For this, we propose to build upon recent initiatives for simulating and synthesizing realistic medical images [8], [9], and to extend the existing literature by specifically investigating database generation issues for a given clinical problem—in our case, localizing myocardial infarct.

A. Methods for simulating realistic cardiac images

Unlike modality propagation approaches, which transfer anatomical information from one modality to another [10], [11], most cardiac synthesis methods include the creation of new data—new myocardial shape or motion, obtained from models of cardiac function. In the quest for higher realism, state-of-the-art methods base this anatomical and functional knowledge on individual patients’ segmentations [12]–[15], and simulated motion from realistic electromechanical models of the heart [12]–[16].

In terms of image synthesis, a first branch of methods investigated the *physics of imaging* in echocardiography [17], [18], computed tomography [19] and magnetic resonance [20], [21]. However, the simulated images were still far from real data as texture, artifacts, and surrounding structures were not taken into account. A significant breakthrough came from the incorporation of real image properties into the synthesized images. Highly realistic results were obtained for 3D echocardiography, with scatterers distributed as in the real sequence and displaced according to the electromechanical simulation [14]. This scheme ideally combines the physics of image formation and real data, but is modality-specific. Besides, generating a new sequence with different myocardial motion requires simulating again the image formation, which may be costly for the creation of large synthetic databases.

ND, MS, HD and NA, are with the INRIA Asclepios team, Sophia Antipolis, France. ND is also with the CREATIS team, CNRS UMR 5220, Lyon, France. Address for correspondence: Nicolas Duchateau, INRIA Asclepios, 2004 route des Lucioles BP 93 06902 Sophia Antipolis Cedex, France. Tel: +33.49238.5024; Fax +33.49238.7669.

Email: nicolas.duchateau@creatis.insa-lyon.fr ; nicholas.ayache@inria.fr

¹https://en.wikipedia.org/wiki/List_of_datasets_for_machine_learning_research

In parallel, a second branch of methods investigated the *warping of existing images* according to a given motion model. First attempts on real data warped a single static CT image with motion from a registered 3D echocardiographic sequence [22], and a single static MR image with a simple kinematic model [23]. Substantial improvements were obtained from the use of a dynamic sequence, previously stabilized to a given instant by registration along the sequence, and realistic simulated motion (Fig.1) [13]. Although this approach can be applied to any modality, it is better suited for cine MRI and CT. Indeed, undesirable warping artifacts may be observed for highly textured data as in echocardiography [13] and tagged-MRI [15].

In the perspective of generating large databases of synthetic cases, we would like to build upon this approach for its portability, provided that specific aspects are improved. In particular, artifacts were reported in [13] when the simulated motion is too different from the original motion. Besides, inaccuracies in the registration along the sequence may propagate and alter the realism of the outcome, and bias its correspondence with the simulated motion used as ground truth. Thus, part of this paper investigates two substantial improvements to this scheme: updating the processing pipeline and therefore solving a much easier registration problem, and directly re-using mesh correspondences to generate new data from a given case without additional registration.

B. Assessment of tissue viability and infarct diagnosis

In terms of application, our objective is to evaluate the potential of this approach for computer-aided diagnosis algorithms, through the generation of large synthetic databases containing ground truth shape and motion. We would like to focus on a concrete clinical problem under active investigation: the assessment of myocardial viability from cardiac image sequences.

Coronary heart disease is actually the leading cause of death in occidental countries (1.8 million deaths per year in Europe [24]). Imaging data is crucial for clinical routine: it facilitates the reading of cardiac function up to shape, motion and deformation descriptors, which are directly related to changes in the myocardial viability [25]–[27]. These descriptors are complex and in high dimension, and therefore require powerful computational tools that properly integrate them to the diagnosis process. Recent works using machine learning showed promising results for infarct diagnosis from image [28] or shape data [29]. In the meantime, we presented a learning method that can estimate the infarct localization [30]. Validation was performed on a large set of synthetic data, which consisted of meshes from realistic electromechanical simulations—but not images. Focus was on the accuracy of the localization, and the database therefore included infarcts of varying extent, location and shape.

In the present work, we would like to move forward and design a database of synthetic images. Besides, we propose to challenge the limits of diagnosis by refining the database design and incorporating different grades of tissue viability. Such cases can be obtained by altering parameters of the

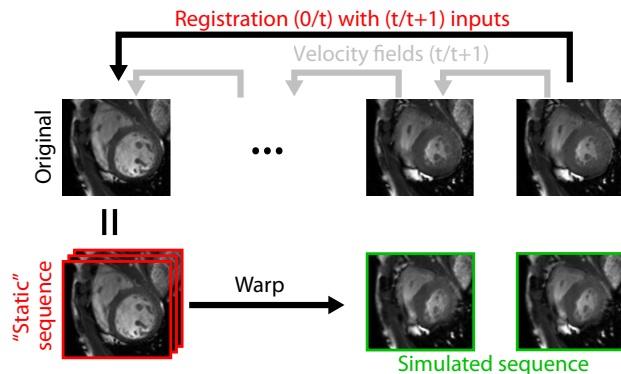


Fig. 1. Processing pipeline used in [13]. The original cardiac sequence is “frozen” to the first instant via image registration (top row). Better performance is reached when initializing the registration between frames 0 and t by the composition of transformations first obtained from the registration of consecutive frames t and $t + 1$ (in practice, handled in the tangent space as the sum of velocity fields linking frames t and $t + 1$). Then, this “static” image is warped according to the simulation displacement field (bottom row). Consistency between the myocardium and the inside/outside anatomy is ensured by diffusion on this displacement field.

electromechanical simulation from healthy to locally diseased. This strategy inherits from [31], which varied the electrical activation parameters to generate up to 180 scenarios from three given pathological image sequences. However, this design modified the parameters around existing personalized values, but did not investigate the creation of pathological cases from healthy data, and therefore the limits of detectable disease and image realism. We tackle this issue in the present paper.

C. Proposed method and contributions

We propose a revised scheme for the generation of realistic 3D cine MR cardiac sequences, which makes substantial improvements to the current state-of-the-art outlined in [13], [14]. Our contributions are both on technical and applicative aspects.

First, our revised processing only involves the registration of rather similar image sequences. This easier registration problem reduces the amount of potential artifacts. In contrast, registration along the sequence involves larger transformations and potential artifacts that could propagate along the sequence [13], [15]. Besides, this strategy biases the validation of the synthesis methodology, if evaluated via tracking along the sequence by a similar algorithm [13], [14].

Then, our new pipeline takes advantage of mesh correspondences to generate variants of a given simulation. No additional registration is involved and correspondences are exact. This opens the door to the generation of large databases of synthetic cases at small computational cost, demonstrated on a given application: the assessment of myocardial viability.

Finally, we propose a thorough qualitative and quantitative evaluation of the quality and usability of the generated images, in the context of the generation of large databases. We notably investigate whether the designed database is balanced below and above the diagnosis limit of a given algorithm, within the

limits of detectable disease and image realism, as illustrated in Fig.2.

II. MATERIALS AND METHODS

A. Template sequences: real database from STACOM 2011

We used the cine MR sequences from 15 volunteers collected for the 1st Cardiac Motion Analysis (cMAC) challenge held at STACOM 2011. These volunteers had no clinical history of cardiac diseases (age 28 ± 5 years, 12(80%) male). Details about data acquisition and the subjects' characteristics can be found in the challenge paper [32]. Images had standard quality, with a spatial resolution of 1.15 or 1.25 mm in the x and y directions, and 8 mm in the z direction (11 to 16 slices, average 14 slices), and 30 frames along the cycle.

Data was already segmented at end-diastole by the semi-automatic protocol described in [33], which involved a biventricular deformable model and manual corrections. As illustrated in Fig.3, the available meshes were volumetric with tetrahedral elements, and included the two ventricles (23770 ± 5219 tetrahedra and 6633 ± 1212 vertices for the LV, of average edge length 3.6 ± 0.2 mm, for a myocardial volume and mass of 95 ± 27 mL and 101 ± 28 g, respectively). Rule-based fibre directions were added over this geometry with an elevation angle between -70 and 70° .

B. Mesh simulations

We applied a realistic electromechanical model to each of the 15 cases to simulate the cardiac function. The model computations were updated each 2 ms of a full cardiac cycle of 1 s. Volumetric meshes were exported at uniformly sampled instants of the cycle (each 17 steps, resulting in 30 meshes along the cycle).

The model was implemented within the SOFA framework², and derives from the Bestel-Clement-Sorine model for the active contraction, and the Mooney-Rivlin model for passive elasticity. It consists of energy-preserving equations during contraction and relaxation, and includes the Frank-Starling effect. We refer the reader to previous publications from our group for further details about this model, including validation issues and a full sensitivity analysis of its parameters [33], [34]. It was retained for its ability to reproduce the behavior of real cases, as evaluated in [33] and more recently in [35], [36]. The realism of simulated pathological behaviors was also tested through the generation of a large database of meshes with spatially varying infarcts [30], and single infarcted or dyssynchronous cases used as input to image synthesis algorithms [13], [14].

1) *Healthy cases*: The simulation parameters were adapted from the ranges in [14] and [30], and are detailed in Tab.I. Balance between the right and left ventricles was slightly updated to reach a more synchronous relaxation on these new geometries. The conductivity and the Rayleigh stiffness were respectively set to $500 \text{ mm} \cdot \text{s}^{-1}$ and 0.06 for the same reason ($400 \text{ mm} \cdot \text{s}^{-1}$ and 0.1 in [30]). Values were also adapted

²<https://www.sofa-framework.org/>




Fig. 2. Schematic view of the tackled challenges for database generation, illustrated on the 30 variants obtained from case v2, arranged according to the disease grade (myocardial viability) and extent (diseased zone volume). Generating a synthetic database should provide a balanced amount of cases below and above the diagnosis limit of a given algorithm, within the limits of detectable disease and image realism.




Fig. 3. Sample of the volumetric meshes and images used in this study (first cases from the cMAC-STACOM 2011 database), at end-diastole.

over the mitral annulus to reduce its deformation and better fit the imaging data. Case-based personalization [35], [37] was not targeted here, and the same parameters were used for all cases. Nonetheless, the retained values led to realistic behaviors close to the original sequences. The personalization error was evaluated in terms of volumes and distance between surfaces. At end-systole, the distance between the

TABLE I
PARAMETERS DEFINING THE HEALTHY-INFARCT RANGE USED FOR THE MESH SIMULATIONS (DETAILED SENSITIVITY ANALYSIS IN [34]).

Category	Symbol	Description	Units	Mitral annulus	[this paper]			[Duchateau et al. 2016]	
					RV	LV Healthy	LV Infarct	Healthy	Infarct
Stiffness	c_1	Mooney-Rivlin modulus	Pa	40e4	2e4	4e4	4e4	2e4	2e4
	c_2	Mooney-Rivlin modulus	Pa	40e4	2e4	4e4	4e4	2e4	2e4
	K	Bulk modulus	Pa	3e6	3e6	3e6	3e6	3e6	3e6
Contraction	σ_0	maximum contraction	Pa	7.4e6	3.7e6	7.4e6	1e5	3.7e6	5e4
	k_0	maximum stiffness	Pa	6e6	6e6	6e6	1e5	6e6	1e5
	k_{atp}	contraction rate	s^{-1}	30	30	30	20	30	20
	k_{rs}	relaxation rate	s^{-1}	70	70	70	80	70	80

manually-segmented and simulated LV endocardial surfaces³ was 3.6 ± 1.7 mm. The agreement between the manually-segmented and simulated LV volumes is plotted in Fig.4a, and a representative example is represented in Fig.5.

2) *Diseased cases*: Variants of each case in the database were generated by modifying the simulation parameters in a given zone, labeled as diseased.

This zone was randomly located within the LV territory irrigated by the left-anterior-descending coronary —territory of higher infarct prevalence [38]— and had random extent and shape. It was defined iteratively by an algorithm similar to the one used in [30], which was improved on two aspects: (i) the algorithm now handles mesh labels as cell attributes instead of vertex attributes for better compatibility with the SOFA mesh format; (ii) for better realism with the mechanisms of ischemia, the zone generation is now constrained to start from the endocardial surface. At each iteration, a cell is randomly selected within the region already labeled as diseased, and a spherical neighborhood of random radius between 5 and 15 mm is picked out around this point and added to the diseased region. The algorithm repeats this process a random number of times (values from 1 to 16). On our data, this led to diseased zones of 5.8 ± 3.0 mL, corresponding to $6.3 \pm 3.4\%$ of the LV myocardium.

Then, the simulation parameters with major influence on the deformation [34] were altered in this diseased region, to mimic changes in active force and tissue elasticity. Contrary to our previous work [30], values were randomly modified within a given interval to simulate varied grades of tissue viability. This was achieved by setting all the selected parameters at a random portion of the healthy-infarct range in Tab.I, 100% standing for healthy and 0% for the infarcts modeled here. This range was adapted from previous publications [14], [30] through new tests on the cMAC-STACOM 2011 data, aiming at realistic motion close from the imaging data. This approach might not correspond to a fine biophysical modeling of the

³Manual segmentation was performed on 2D slices through the Segment software (Medviso, Lund, SE). Intra-operator reproducibility was 0.58 and 0.57 mm at end-diastole and end-systole, respectively. We used a non-symmetric distance between contours, as we compare a 3D surface Ω_{simu} (the endocardium from the mesh simulation) to a stack of 2D contours (the segmented endocardium in each slice). The distance was defined as $\frac{1}{K} \sum_{k=1}^K d(\mathbf{x}_k, \Omega_{simu})$, where \mathbf{x}_k denotes the k -th point of the segmentation stack. Thus, these surfaces did not perfectly match at end-diastole: their distance was 2.3 ± 0.4 mm, and the volume difference was 11.3 ± 6.6 mL.




Fig. 4. Comparison of LV volumes at end-systole: the 3D mesh simulation compared to the 3D segmentation from 2D slices (a) from the original sequence to evaluate the personalization error, and (b) from the synthesized healthy sequence to evaluate the accuracy of the generated images with respect to the mesh simulation. The black arrows point out a subject with representative personalization quality (v7, average personalization error in terms of distance between surfaces at end-systole: 3.0 mm).




Fig. 5. Simulated mesh superimposed to the original sequence of the representative subject v7 pointed out in Fig.4. An animated version of the figure is available as [Supplementary Material](#)⁶.

infarct, but this would require a more in-depth study, outside the scope of this paper.

We created 30 realistic pathological variants of each of the 15 cases in the database, resulting in 450 diseased cases that complement the 15 healthy simulations described in Sec.II-B1. The characteristics of all the generated cases (extent, location, and tissue viability) are summarized in Fig.6.

C. Image synthesis

1) *Generation of a reference synthetic sequence*: This part of the methodology is synthesized in Algorithm 1. It corresponds to the generation of a synthetic sequence close from the original one (the first two rows in Fig.7) —in our case, a healthy sequence.

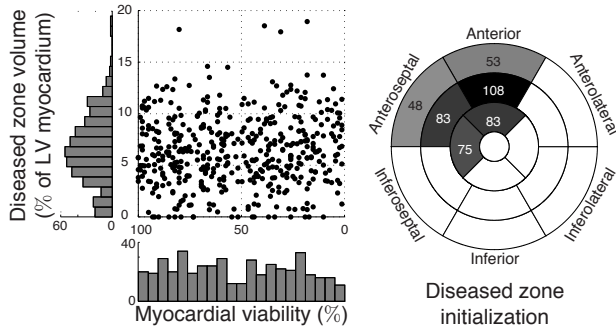


Fig. 6. Characteristics of the generated infarcts. *Left*: grade vs. volume. *Right*: starting location on the endocardial surface.




Fig. 7. Processing pipeline proposed in this paper. The image at $t = 0$ is first warped according to the simulation (second row). The original sequence is then registered to this warped sequence (red color). This registration is easier if the simulated motion is close to the one in the original sequence. The obtained displacement field is masked and temporally smoothed before applying it to the original sequence (Sec.II-C1). Variants of this sequence with pathological motion are easily obtained by using correspondences between the simulated meshes (blue color).

The original sequence is cropped around the heart and upsampled in the z direction to reach an isotropic resolution (step 1, linear interpolation via the ITK image resampling filter). The first frame in the sequence is then warped according to the healthy mesh simulation (steps 2 to 4). In practice, the displacement field $\varphi_0(t)$ from the simulation between times 0 and t is only available inside the myocardium. Diffusion is therefore applied to expand this displacement field before applying it to the frame at time 0 (Fig.8a-b). This guarantees continuity in the motion of the anatomical structures surrounding the heart, with decreasing effect according to the distance from the myocardium. As in [13], this was achieved via the iterative convolution with a Gaussian kernel (the ITK




Fig. 8. Spread of the vector field $\varphi_0(t)$ from the electromechanical simulation (a), after diffusion to ensure continuity in the motion around the heart (b), and with cropping below the mitral annulus before the diffusion to reduce artifacts due to potential discrepancies between the original and simulated motion at this level (c).

recursive Gaussian filter with a standard deviation of 1 voxel) followed by the reinitialization of the vector field values over the myocardium and forcing them to 0 at the image borders. The process was iterated at most 500 times, and stopped if the difference in vectors between two iterations is lower than 0.001 voxel. Despite modifying the parameters value at the mitral annulus (Sec.II-B1), the mesh simulations may still differ from the original sequence at this level. Thus, the displacement field from the simulation is also cropped 1 cm below the mitral annulus before the diffusion step (Fig.8c). This arbitrary cropping limit roughly corresponds to the myocardial thickness, and our tests showed that it improves the realism of the simulated images.

This process provides a sequence that moves as the simulated motion, but keeps the texture of the first frame, which is still unrealistic [22], [23]. Nonetheless, if the simulated motion is similar to the one in the original sequence, registering the original sequence to this warped sequence involves small displacements and should therefore be relatively easy to solve (step 5, also marked with red color in Fig.7). We achieved this by applying the Demons algorithm with sum-of-squared-

Algorithm 1: Generation of a healthy sequence.

1. *Crop and upsample sequence to isotropic resolution*
for $t \in [0, T]$ do
 2. *Extract mesh local displacement: $\varphi_0(t)$*
 3. *Crop at the mitral annulus and diffusion: $D(\varphi_0(t))$*
 4. *Warp image at time 0 by $D(\varphi_0(t))$*
 5. *Register warped and original images at time t*
Output: vector field $\phi_{original \rightarrow warped}(t)$
 6. *Mask this field by a diffused mask of the myocardium, and temporally smooth over 2 instants*
Output: vector field $\phi_{original \rightarrow warped}^*(t)$
 7. *Warp original image at t by this field*
 8. *Downsample sequence to the original resolution*
-

differences similarity metric⁴.

To prevent artifacts outside the heart, the vector field $\phi_{original \rightarrow warped}(t)$ obtained from this registration is multiplied by a diffused mask of the myocardium. This mask corresponds to a binary image of the myocardium that is expanded outside the myocardium by a diffusion process similar to the one used in Sec.II-C1. The masked vector field is then smoothed temporally to force consistency in the temporal motion (step 6, via the ITK recursive Gaussian filter applied along the temporal dimension of the 4D vector field, with a standard deviation of $2\Delta_t$, where Δ_t is arbitrarily set to 1), and applied to the original sequence (step 7). Finally, the simulated sequence is downsampled to the resolution of the original sequence (step 8, ITK filter similar to the one in step 1).

2) *Generation of additional image sequences from the reference synthetic sequence:* Given a simulated sequence close from the original data (in our case, the healthy sequence generated in Sec.II-C1 from the healthy mesh simulation), we propose to exploit mesh correspondences to generate new sequences corresponding to the infarct configurations described in Sec.II-B2. This part of the methodology corresponds to the blue-squared part of Fig.7, and is synthesized in Algorithm 2.

For a given infarct configuration, we first extract the vector field $\phi_{healthy \rightarrow infarct}(t)$ describing the mesh correspondences at a given instant, and expand it outside the myocardium via the diffusion process of Sec.II-C1. This vector field is then composed with $\phi_{original \rightarrow warped}^*(t)$, the one used to warp the original sequence in the healthy case, and finally applied to the original sequence.

3) *Parallelization of the processes and computational time:* All the mesh simulations were performed in parallel on our local cluster platform⁵. Each mesh simulation from SOFA

⁴Chosen parameters: smoothing of the displacement field by 3 voxels, smoothing of the update field by 0.5 voxel, weight between the similarity and regularization terms of 0.1, 3 levels of resolution (30, 20 and 10 iterations, respectively), Gaussian smoothing of the similarity criteria by 3 voxels, and Gaussian regularization. Code v1.2, April 2016 from [39], available at: <https://team.inria.fr/asclepios/software/lcclodemons/>

⁵<https://wiki.inria.fr/ClustersSophia/Hardware>

Algorithm 2: Generation of a diseased sequence.

```

for  $i \in N_{infarct}$  do
  for  $t \in [0, T]$  do
    1. Extract mesh correspondences
      Output: vector field  $\phi_{healthy \rightarrow infarct}(t)$ 
    2. Diffusion:  $D(\phi_{healthy \rightarrow infarct}(t))$ 
    3. Compose
       $D(\phi_{healthy \rightarrow infarct}(t)) \circ \phi_{original \rightarrow warped}^*(t)$ 
    4. Warp original image at time  $t$  by this field
    5. Downsample sequence to the original resolution

```

takes approximatively one hour to complete. Each registration takes at most 30 min. The diffusion is the longest process in the pipeline, and takes around 30 min. The remaining processes take less than 1 min.

For the generation of the reference sequence, all the instants can be processed independently and therefore parallelized except the temporal smoothing, which involves the whole 4D sequence. A similar parallelization was adopted for the generation of the variants of the reference sequence. Thus, simulating one or several variants of the same reference sequence takes the same amount of time—within the limits of parallel processes allowed for a given user on our cluster platform—and is therefore a substantial asset for the generation of large databases of synthetic sequences.

III. EXPERIMENTS AND RESULTS

*Supplementary Material*⁶ is provided for a better assessment of the synthetic sequences quality. It includes animated views of Fig.5 and 9, and animated views of the synthetic sequences for the best, medium and worst personalized cases for the healthy and strong infarct (0% viability) simulations, which complement the illustration in Fig.11.

A. Qualitative evaluation

1) *Mimicking an ideal personalization:* We first designed a toy experiment to highlight the limits of the method from [13]. The first frame in the sequence was warped along the cycle, as in step 4 of Algorithm 1 and was used as input to the image synthesis algorithm illustrated in Fig.1. By definition, this toy sequence matches the mesh simulation. It therefore mimics an ideal personalization, namely a perfect match between the original sequence and the mesh simulation (Fig.9a-b). The method from [13] fails at reproducing the toy sequence (Fig.9c-f), due to the mismatch between the displacement field from the sequence registration and the one from the mesh simulation. In contrast, our method does not suffer these limitations on this toy experiment: the registration part of our method involves already matched images, as the warped sequence is used in lieu of the original sequence, and therefore $\phi_{original \rightarrow warped}(t) = 0$.

2) *Matching with the mesh contours:* Figure 10 illustrates the representative differences between the output of our method and the one from [13], on the best and worst personalized cases (similar observations were made on all cases). Our method better matches the mesh contours. In contrast, for the method from [13] the myocardium spreads inside the LV and RV mesh blood pool, stretching is higher near the epicardium, and other anatomical zones such as the mitral annulus are insufficiently deformed. Figure 11 complements these observations by showing that the matching with the mesh contours is conserved when simulating strongly diseased tissue, despite substantial changes in the simulation (contraction is lower and mesh contours are thinner at the infarct level). The animated views

⁶<http://www-sop.inria.fr/asclepios/docs/ImageSynthesisExamples.zip>




Fig. 9. Comparison between the method from [13] and our approach on a toy example (subject v7), demonstrating the limits of [13] by the generated MR image at end-systole. (a) and (b): first frame warped along the sequence (step 4 of Algorithm 1), referred to as the toy sequence. By definition, it matches the mesh simulation and mimics an ideal personalization. Thus, it also corresponds to the result of our method on this ideal case, as the warped sequence is used in lieu of the original sequence and therefore $\phi_{original \rightarrow warped}(t) = 0$. (c) and (d): sequence simulated from the toy sequence using the method from [13], which fails at reproducing the toy sequence. (e) and (f): signed difference between the toy sequence and the result of [13]. An animated version of the figure is available as [Supplementary Material](#)⁶.

in the [Supplementary Material](#)⁶ lead to similar conclusions on the best, medium and worst personalized cases. Slight mismatch is observed along the lateral wall of the medium personalized case. This results from mis-aligned slices in the original sequence, at the infarct level where images are shown, which are not taken into account in the 3D mesh used for the simulations.

3) *Examining the limits to the image realism:* The warping applied to the original sequence might alter the realism of the generated synthetic images, although this depends on many factors, among which the quality of the personalization for the healthy simulation, the characteristics of the local disease, and the texture of the original sequence in the warped regions (myocardium and surroundings). Figure 13 provides a qualitative assessment of the simulated images with the highest tissue damage (8 cases with tissue viability < 10% and diseased zone volume > 10%), against the original and healthy simulated images. Warping artifacts are hardly noticeable for the healthy simulations, although changes in the myocardial thickness and shape can be noted at the apical level of some cases (v2 infarct #7 and v5 infarct #23). Due to the diffusion of the vector fields around the myocardium, the image realism is preserved in the myocardial surroundings, even for the cases with low




Fig. 10. Qualitative comparison of our method to the one from [13], on the best and worst personalized cases (v6 and v12, average personalization errors at end-systole: 2.1 and 8.3 mm, respectively). Our method better fits the mesh simulation, as highlighted on the blood pool and epicardium along the LV lateral wall, on the RV contours, and at the mitral annulus (white arrows). Similar observations were made on all cases.

tissue viability. The main limits to the image realism actually come from the motion prescribed in the mesh simulations, with sometimes noticeable myocardial shape differences due to highly contrasted myocardial border in the original sequence (v4 infarct #2 and #18, v5 infarct #25 and v15 infarct #22).

B. Quantitative evaluation

1) *Manual assessment: segmentation:* We manually segmented a representative subset of the simulated images at end-systole, made of the healthy one and the ones with tissue viability closest to 50 and 0%, for each of the 15 cases in the database. Then, we compared the segmentations (stacked 2D contours) with the mesh simulations (3D surfaces, taken as ground truth). These two steps followed the protocol adopted for the original sequence, as described in Sec.II-B1. Agreement is observed on volumes (Fig.4b), and reflects the accuracy of the synthesized images with respect to the prescribed motion. This agreement is naturally better than the match between the original sequence and the simulated

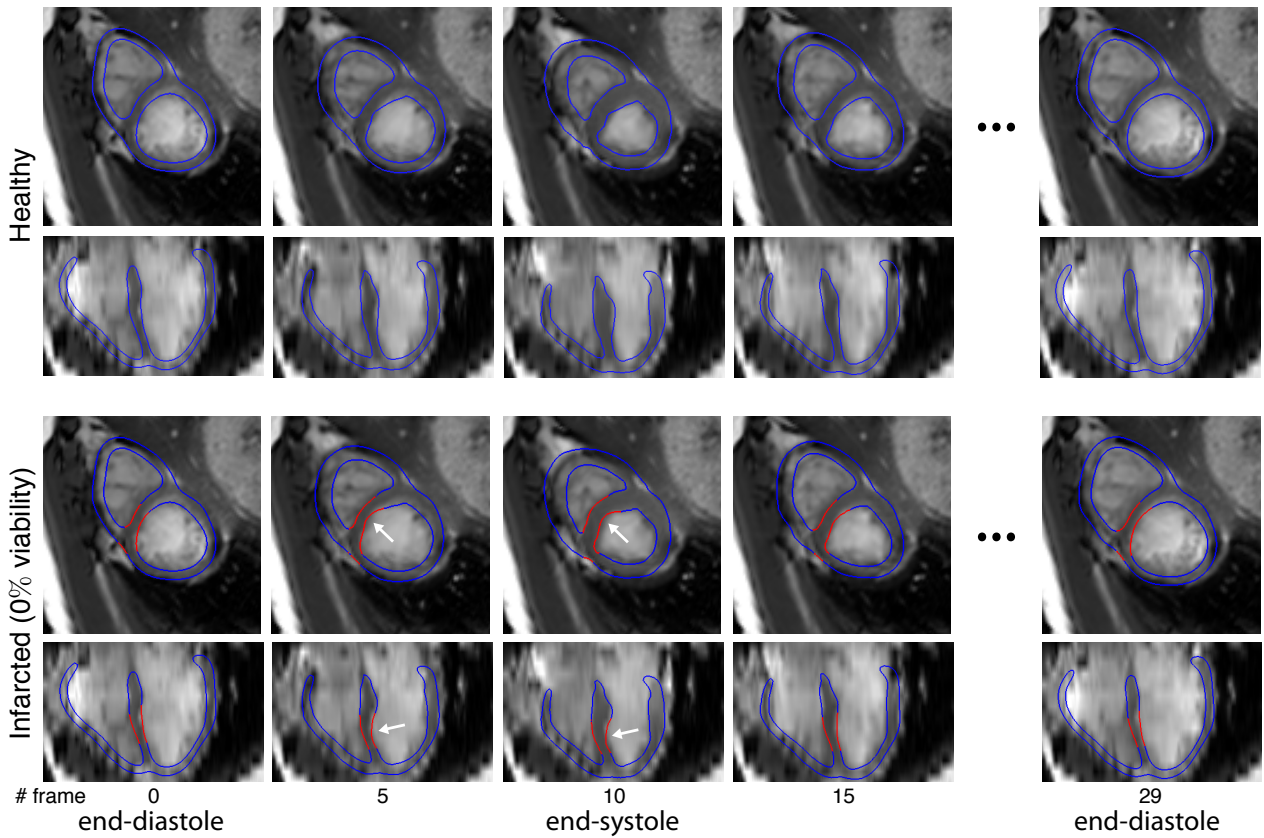


Fig. 11. Comparison of healthy and infarcted synthetic sequences along the cycle, for the best personalized case v6. The matching with the mesh contours is conserved, despite substantial changes in the simulation at the infarct level (red color), where contraction is lower and mesh contours are thinner. An animated version of the figure is available as *Supplementary Material*⁶, together with similar views of the medium and worst personalized cases.

contours (Fig.4a), which reflected the personalization error. Similar observations can be made in terms of distance between the surfaces³, without substantial differences between the different grades of tissue viability: values were 2.1 ± 0.4 , 2.2 ± 0.5 and 2.3 ± 0.5 mm for the healthy simulation and the tissue grades of 50 and 0%, respectively. These measurements lie in the range of differences between the 3D surface and the stack of 2D contours at end-diastole (Sec.II-B1). One common outlier was observed, which actually corresponded to a case in the cMAC-STACOM database (V11) with two mis-aligned slices at the papillary muscles level, while the 3D mesh used for the simulations does not take this into account.

2) *Automated assessment: registration*: We finally evaluated the sequence quality against state-of-the-art motion estimation along the sequence. We applied the Demons algorithm with sum-of-squared-differences similarity metric. We followed a two-step strategy similar to the one in Fig.1, where pairs of consecutive frames are first registered, and the resulting transformations are composed (in practice, handled in the tangent space as the sum of velocity fields linking frames t and $t+1$) and used as initialization for a second registration between frames at $t=0$ and t . Parameters⁴ were similar to the ones detailed in Sec.II-C1, except the smoothing of the

displacement field —by 2 voxels here. Then, the first frame was propagated along the sequence according to the estimated displacement field.

The simulated contraction deficiency due to the altered tissue viability was assessed through the LV regional volume change along the sequence, on the simulated meshes and the ones propagated along the simulated images via the above-described registration. The regional volume of a given region is defined as the volume formed by the LV barycenter and the endocardial surface of this region [33]. This measure partially relates to the myocardial deformation, and is more robust to small registration errors since it averages values over a given zone, and does not involve a spatial derivative.

These measurements are plotted in Fig.12 against the tissue viability. Differences between the healthy and diseased behaviors get more contrasted as the tissue viability decreases. They become significant for medium-grade tissue viability, although results are less contrasted for the simulated images. The latter mainly comes from the information loss when computing these values on meshes (high resolution) or images (lower resolution, in particular on the z direction), and from the accuracy of the motion estimation along the sequence: lower regional volume change is also observed for the healthy myocardium; this should not be the case, as segmentation




Fig. 12. Regional volume change from end-diastole to end-systole computed on the simulated meshes and images. Comparison between the infarcted (red) and healthy (blue) simulations over the region labeled as infarct. Stars on the x -axis indicate statistically significant differences between the infarcted and healthy simulations (p -value < 0.01 from the Mann-Whitney U test).

showed that the simulated sequences agree with the prescribed motion (Sec.III-B1).

IV. DISCUSSION

We presented a new approach to synthesize cardiac image sequences from existing real sequences and a realistic electromechanical model, demonstrated on cine MR data. Our pipeline improves state-of-the-art methods by revisiting specific aspects of the processing, resulting in higher realism of the simulated images and better suitability for the generation of large databases of pathological cases. In this context, we designed a large database of 15 healthy and 450 diseased cases, and evaluated the quality and usability of the generated images.

Our method provides higher image realism and better agreement with the ground truth motion (Fig.10). This mainly results from avoiding registration along the original sequence in our pipeline. Such a registration is a potential source of artifacts, as demonstrated on a toy example (Fig.9). There is potentially a mixture of registration inaccuracies, and of mismatch between the displacement fields from the sequence registration and from the mesh simulation. The former might be addressed by better algorithms, although registering cine MR sequences is still highly challenging [32] due to the limited z -axis resolution and the lack of clear texture patterns attached to the myocardium, contrary to echocardiography or tagged-MRI. The latter might be improved by the personalization of the electromechanical model, using more advanced schemes [35], [37] and incorporating the match of displacement fields in the personalization metrics. In our study, the use of well-adjusted consensual values for all cases still provided simulations close to the original sequence (Fig.4 and 5). Taking advantage of this, our pipeline involved a much easier registration problem at each instant of the sequence. Better personalization might help but is not a prerequisite, as demonstrated on our data. Nonetheless, this strategy may be limited by artifacts from warping the first frame, and texture or gray-level differences with the original sequence.

This registration quality was carefully checked and was not a major issue in our study.

Infarct detection and localization implies the use of dedicated algorithms [28]–[30], whose performance depends on the quality of the features extraction (e.g. shape, motion and deformation), and therefore segmentation or registration techniques. Thus, we preferred to focus on the validation of the accuracy, quality, and usability of the simulated images assessed from the segmentation and registration point-of-view, as testing the pipeline on more advanced and application-specific algorithms may deviate from the original focus of the paper.

Segmentation was not used as a diagnosis tool as it does not allow to propagate the labeling of the myocardial zones (and in particular the diseased region) along the cycle, unlike registration. It only served for evaluation of the simulated images accuracy, and was therefore kept manual to prevent additional uncertainties coming from the use of an algorithm.

The registration used for the quantitative evaluation actually suffers the above-discussed limitations, and results in less contrasted differences between healthy and diseased behaviors, in contrast with the mesh simulations (these differences are also attenuated by the resolution differences between meshes and images). The optimization of a registration strategy was not the core of this paper and we preferred to test our images with a state-of-the-art algorithm. Nonetheless, the generated population still allows to see significant differences when the tissue viability decreases. As synthetic cases lie below and above the limit of the retained diagnosis method, this highlights a critical aspect for the generation of large databases (Fig.2): better understanding the limits of the evaluated algorithms should take precedence over the race for the largest possible population.

The variety of the textures in the generated sequences was not addressed here. All the generated variants of a given sequence re-use the same texture, which is warped according to the differences between the original and the simulated motion. This may be a potential limitation for testing methods that actually make use of the texture information, such as the recent neural networks-based learning algorithms. The risk of unrealistic texture warping still exists with our method when deviating too far from the original motion, but this highly depends on the texture information of the original sequence and the prescribed motion of the simulated meshes (Fig.13). On our dataset, this limit was substantially less problematic than in the results presented for highly textured modalities such as echocardiography [13] and tagged-MRI [15]. Simulating the physics of imaging would be a robust way to tackle these issues, but at high computational cost. On the contrary, our method has the advantage of easily providing variants of a given case. We notably take advantage of existing mesh correspondences to generate new data without additional registration, at reduced computational cost and through highly parallelizable algorithms, which are key assets when targeting the design of large databases. Future improvements might incorporate texture synthesis on top of it, inspired from physics-based [14] or patch synthesis [11] strategies.

Due to such texture modeling limits, we preferred to

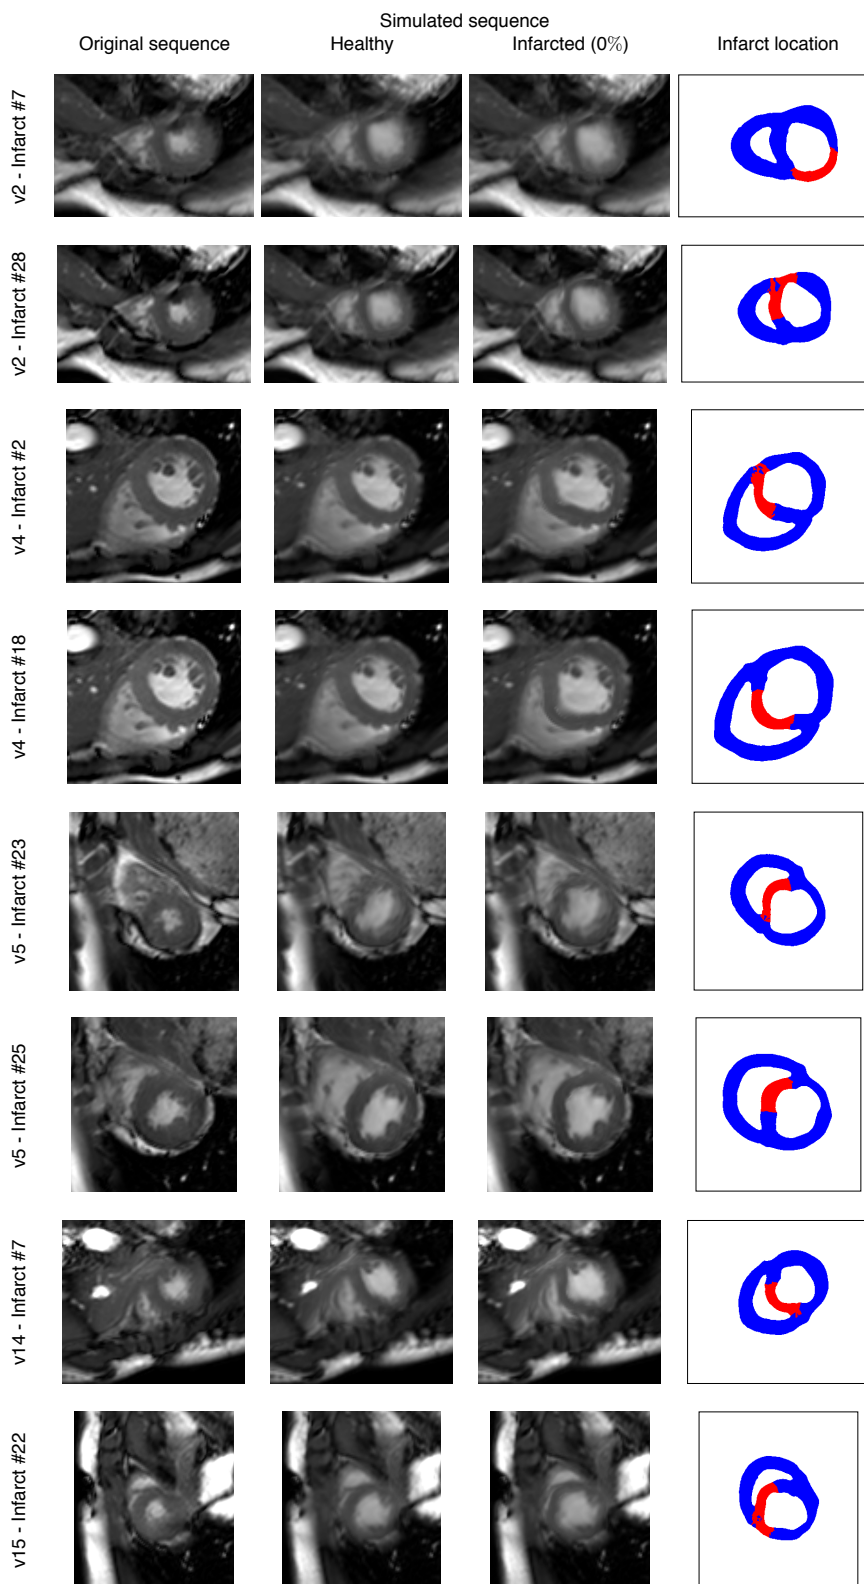


Fig. 13. Qualitative assessment of the simulated images with the highest tissue damage (8 cases with tissue viability < 10% and diseased zone volume > 10%), against the original and healthy simulated images. Warping artifacts are hardly noticeable for the healthy simulations. Slight myocardial shape differences can be noticed on the diseased cases with highly contrasted myocardial border in the original sequence.

demonstrate the methodology on cine MR sequences. From the technical side, the interest is to support the design of computer-assisted methods that will tackle the processing of large databases with such sequences, such as the UK Biobank [6], eventually for the diagnosis of myocardial infarct. From the clinical side, applying such algorithms to cine MR would have potential for a first opinion diagnosis, which if positive and possible on the studied patient would lead to the recommendation of an extra LGE scan. Moreover, the proposed methodology is rather generic and could be used for echocardiography once better textures can be synthesized, for higher clinical impact.

The STACOM 2011 challenge data were already segmented in a format compatible with SOFA simulations. Focus was therefore set on the image synthesis steps coming after the segmentation of the initial frame. Atlas-based segmentation techniques could provide correspondences between the generated cases, of relevance for larger databases. Nonetheless, mesh fitting needs to be carefully checked to prevent the propagation of artifacts. Atlas-based template sequences [40] might also be used, although this raises texture-related issues (variety of textures and potential warping artifacts).

Our strategy is in line with the recent initiatives for simulating and synthesizing realistic medical images [8], [9]. On the clinical problem chosen here for illustration purposes (assessment of the tissue viability and infarct diagnosis), expectations are high to bridge the gap between synthetic mesh simulations and real imaging data [30]. In this sense, we incorporated here different grades of tissue viability, and higher realism in the diseased zone generation algorithm — now initiated at the endocardial surface. Yet, aspects of higher realism of the pathological simulations (e.g. the definition of a neighboring partially affected region, and incorporating ventricular remodeling consecutive to an infarct) were outside the scope of our objectives here. Nonetheless, our method is readily applicable to these more realistic simulations, once available. Mixing synthetic and real data should be considered as a future target, as demonstrated on easier problems [41], provided issues of transfer learning are also taken into account.

V. CONCLUSION

We presented a new approach for the generation of synthetic cardiac image sequences from existing real sequences, which improves current methods in terms of image realism. Our pipeline, demonstrated on cine MR sequences, is highly suitable for the generation of large databases at reduced computational cost, and paves the ground for bridging the gap between synthetic mesh simulations and real imaging data [30].

ACKNOWLEDGMENTS

The authors acknowledge the partial support from the European Research Council (MedYMA ERC-AdG-2011-291080) and the European Union 7th Framework Programme (VP2HF FP7-2013-611823). They also thank R. Mollero and S. Giffard-Roisin (INRIA Asclepios, Sophia-Antipolis, France) for fruitful discussions about SOFA simulations, and L. Le Folgoc (Microsoft Research, Cambridge, UK) for his initial help on the implementation of [13].

SUPPLEMENTARY MATERIAL

Filename	Description
Fig5.avi	Animated view of Fig.5
Fig9.avi	Animated view of Fig.9
Fig11_Best.avi	Animated synthetic sequences for the best/medium/worst personalized cases that complement Fig.11
Fig11_Medium.avi	
Fig11_Worst.avi	

<http://www-sop.inria.fr/asclepios/docs/ImageSynthesisExamples.zip>

REFERENCES

- [1] A. Suinesiaputra, P. Medrano-Gracia, B. Cowan, and A. Young, "Big heart data: advancing health informatics through data sharing in cardiovascular imaging," *IEEE J Biomed Health Inform*, vol. 19, pp. 1283–90, 2015.
- [2] V. Klinke, S. Muzzarelli, N. Lauriers, D. Locca, G. Vincenti, P. Monney, C. Lu, D. Nothnagel, G. Pilz, M. Lombardi, A. van Rossum, A. Wagner, O. Bruder, H. Mahrholdt, and J. Schwittler, "Quality assessment of cardiovascular magnetic resonance in the setting of the European CMR registry: description and validation of standardized criteria," *J Cardiovasc Magn Reson*, vol. 15, p. 55, 2013.
- [3] C. Fonseca, M. Backhaus, D. Bluemke, R. Britten, J. Chung, B. Cowan, I. Dinov, J. Finn, P. Hunter, A. Kadish, D. Lee, J. Lima, P. Medrano-Gracia, K. Shivkumar, A. Suinesiaputra, W. Tao, and A. Young, "The Cardiac Atlas Project—an imaging database for computational modeling and statistical atlases of the heart," *Bioinformatics*, vol. 27, pp. 2288–95, 2011.
- [4] P. Lancellotti, "Normal reference ranges for echocardiography: do we really need more?" *Eur Heart J Cardiovasc Imaging*, vol. 15, pp. 253–4, 2014.
- [5] W. Bai, W. Shi, A. de Marvao, T. Dawes, D. O'Regan, S. Cook, and D. Rueckert, "A bi-ventricular cardiac atlas built from 1000+ high resolution MR images of healthy subjects and an analysis of shape and motion," *Med Image Anal*, vol. 26, pp. 133–45, 2015.
- [6] S. Petersen, P. Matthews, F. Bamberg, D. Bluemke, J. Francis, M. Friedrich, P. Leeson, E. Nagel, S. Plein, F. Rademakers, A. Young, S. Garratt, T. Peakman, J. Sellors, R. Collins, and S. Neubauer, "Imaging in population science: cardiovascular magnetic resonance in 100,000 participants of UK Biobank - rationale, challenges and approaches," *J Cardiovasc Magn Reson*, vol. 15, p. 46, 2013.
- [7] A. Ganna and E. Ingelsson, "5 year mortality predictors in 498,103 UK Biobank participants: a prospective population-based study," *Lancet*, vol. 386, pp. 533–40, 2015.
- [8] T. Glatard, C. Lartizien, B. Gibaud, R. da Silva, G. Forestier, F. Cervenansky, M. Alessandrini, H. Benoit-Cattin, O. Bernard, S. Camarasa-Pop, N. Cerezo, P. Clarysse, A. Gaignard, P. Hugonnard, H. Liebgott, S. Marache, A. Marion, J. Montagnat, J. Tabary, and D. Friboulet, "A virtual imaging platform for multi-modality medical image simulation," *IEEE Trans Med Imaging*, vol. 32, pp. 110–8, 2013.
- [9] S. Tsaftaris, A. Gooya, A. Frangi, and J. Prince, *Simulation and Synthesis in Medical Imaging, MICCAI-SASHIMI, LNCS 9968*. Springer International Publishing, 2016.
- [10] D. Ye, D. Zikic, B. Glocker, A. Criminisi, and E. Konukoglu, "Modality propagation: coherent synthesis of subject-specific scans with data-driven regularization," in *MICCAI, LNCS 8149*, 2013, pp. 606–13.
- [11] N. Cordier, H. Delingette, M. Le, and N. Ayache, "Extended modality propagation: image synthesis of pathological cases," *IEEE Trans Med Imaging*, vol. 35, pp. 2598–608, 2016.
- [12] M. De Craene, S. Marchesseau, B. Heyde, H. Gao, M. Alessandrini, O. Bernard, G. Piella, A. Porras, L. Tautz, A. Hennemuth, A. Prakosa, H. Liebgott, O. Somphone, P. Allain, S. Makram Ebeid, H. Delingette, M. Sermesant, J. D'hooge, and E. Saloux, "3D strain assessment in ultrasound (Straus): a synthetic comparison of five tracking methodologies," *IEEE Trans Med Imaging*, vol. 32, pp. 1632–46, 2013.
- [13] A. Prakosa, M. Sermesant, H. Delingette, S. Marchesseau, E. Saloux, P. Allain, N. Villain, and N. Ayache, "Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images," *IEEE Trans Med Imaging*, vol. 32, pp. 99–109, 2013.

- [14] M. Alessandrini, M. De Craene, O. Bernard, S. Giffard-Roisin, P. Allain, J. Weese, E. Saloux, H. Delingette, M. Sermesant, and J. D'hooge, "A pipeline for the generation of realistic 3D synthetic echocardiographic sequences: methodology and open-access database," *IEEE Trans Med Imaging*, vol. 34, pp. 1436–51, 2015.
- [15] Y. Zhou, M. De Craene, O. Somphone, M. Sermesant, and O. Bernard, "Generation of realistic 4D synthetic CSPAMM tagged MR sequence for benchmarking cardiac motion tracking algorithms," in *MICCAI-SASHIMI, LNCS 9968*, 2016, pp. 108–117.
- [16] Q. Duan, P. Moireau, E. Angelini, D. Chapelle, and A. Laine, "Simulation of 3D ultrasound with a realistic electro-mechanical model of the heart," in *FIMH, LNCS 4466*, 2007, pp. 463–73.
- [17] H. Gao, H. Choi, P. Claus, S. Boonen, S. Jaecques, G. Van Lenthe, G. Van der Perre, W. Lauriks, and J. D'hooge, "A fast convolution-based methodology to simulate 2-D/3-D cardiac ultrasound images," *IEEE Trans Ultrason Ferroelectr Freq Control*, vol. 56, pp. 404–9, 2009.
- [18] O. Shakil, F. Mahmood, and R. Matyal, "Simulation in echocardiography: an ever-expanding frontier," *J Cardiothorac Vasc Anesth*, vol. 26, pp. 476–85, 2012.
- [19] W. Segars, M. Mahesh, T. Beck, E. Frey, and B. Tsui, "Realistic CT simulation using the 4D XCAT phantom," *Med Phys*, vol. 35, pp. 3800–8, 2008.
- [20] P. G. Kwan RK, Evans AC, "MRI simulation-based evaluation of image-processing and classification methods," *IEEE Trans Med Imaging*, vol. 18, pp. 1085–97, 1999.
- [21] C. Tobon-Gomez, F. Sukno, B. Bijnens, M. Huguette, and A. Frangi, "Realistic simulation of cardiac magnetic resonance studies modeling anatomical variability, trabeculae, and papillary muscles," *Magn Reson Med*, vol. 65, pp. 280–8, 2011.
- [22] F. Li, J. White, M. Rajchl, A. Goela, and T. Peters, "Generation of synthetic 4D cardiac CT images by deformation from cardiac ultrasound," in *MICCAI-AECAI, LNCS 7815*, 2013, pp. 132–41.
- [23] P. Clarysse, J. Tafazzoli, P. Delachartre, and P. Croisille, "Simulation based evaluation of cardiac motion estimation methods in tagged-MR image sequences," *J Cardiovasc Magn Reson*, vol. 13, p. 360, 2011.
- [24] M. Nichols, N. Townsend, P. Scarborough, and M. Rayner, *European cardiovascular disease statistics*. European Society of Cardiology, European Heart Network and British Heart Foundation, 2012.
- [25] J. Bax and V. Delgado, "Detection of viable myocardium and scar tissue," *Eur Heart J Cardiovasc Imaging*, vol. 16, pp. 1062–4, 2015.
- [26] B. Bijnens, P. Claus, F. Weidemann, J. Strotmann, and G. Sutherland, "Investigating cardiac function using motion and deformation analysis in the setting of coronary artery disease," *Circulation*, vol. 116, pp. 2453–64, 2007.
- [27] M. Konstam, D. Kramer, A. Patel, M. Maron, and J. Udelson, "Left ventricular remodeling in heart failure: current concepts in clinical significance and assessment," *JACC Cardiovasc Imaging*, vol. 4, pp. 98–108, 2011.
- [28] V. Sudarshan, U. Acharya, E. Ng, C. Meng, R. Tan, and D. Ghista, "Automated identification of infarcted myocardium tissue characterisation using ultrasound images: a review," *IEEE Rev Biomed Eng*, vol. 8, pp. 86–97, 2015.
- [29] A. Suinesiaputra, X. Zhang, P. Ablin, X. Alba, M. Alessandrini, J. Allen, W. Bai, S. Cimen, B. Cowan, P. Claes, J. D'hooge, N. Duchateau, J. Ehrhardt, A. Frangi, V. Grau, A. Gooya, K. Lekadir, A. Lu, A. Mukhopadhyay, I. Oksuz, N. Parajuli, P. Piras, X. Pennec, M. Pereanez, D. Rueckert, M. Rohe, D. Saring, M. Sermesant, K. Siddiqi, M. Tabassian, L. Teresi, S. Tsaftaris, M. Wilms, A. Young, and P. Medrano-Gracia, "Statistical shape modelling of the left ventricle: myocardial infarct classification challenge," *IEEE J Biomed Health Inform*, 2017, in press.
- [30] N. Duchateau, M. De Craene, P. Allain, E. Saloux, and M. Sermesant, "Infarct localization from myocardial deformation: prediction and uncertainty quantification by regression from a low-dimensional space," *IEEE Trans Med Imaging*, vol. 35, pp. 2340–52, 2016.
- [31] A. Prakosa, M. Sermesant, P. Allain, N. Villain, C. Rinaldi, K. Rhode, R. Razavi, H. Delingette, and N. Ayache, "Cardiac electrophysiological activation pattern estimation from images using a patient-specific database of synthetic image sequences," *IEEE Trans Biomed Eng*, vol. 61, pp. 235–45, 2014.
- [32] C. Tobon-Gomez, M. De Craene, K. McLeod, L. Tautz, W. Shi, A. Hennemuth, A. Prakosa, H. Wang, G. Carr-White, S. Kapetanakis, A. Lutz, V. Rasche, T. Schaeffter, C. Butakoff, O. Friman, T. Mansi, M. Sermesant, X. Zhuang, S. Ourselin, H. Peitgen, X. Pennec, R. Razavi, D. Rueckert, A. Frangi, and K. Rhode, "Benchmarking framework for myocardial tracking and deformation algorithms: an open access database," *Med Image Anal*, vol. 17, pp. 632–48, 2013.
- [33] S. Marchesseau, H. Delingette, M. Sermesant, R. Cabrera-Lozoya, C. Tobon-Gomez, P. Moireau, R. Figueras i Ventura, K. Lekadir, A. Hernandez, M. Garreau, E. Donal, C. Leclercq, S. Duckett, K. Rhode, C. Rinaldi, A. Frangi, R. Razavi, D. Chapelle, and N. Ayache, "Personalization of a cardiac electromechanical model using reduced order unscented Kalman filtering from regional volumes," *Med Image Anal*, vol. 17, pp. 816–29, 2013.
- [34] S. Marchesseau, H. Delingette, M. Sermesant, M. Sorine, K. Rhode, S. Duckett, C. Rinaldi, R. Razavi, and N. Ayache, "Preliminary specificity study of the Bestel-Clement-Sorine electromechanical model of the heart using parameter calibration from medical images," *J Mech Behav Biomed Mater*, vol. 20, pp. 259–71, 2013.
- [35] R. Mollero, X. Pennec, H. Delingette, N. Ayache, and M. Sermesant, "A multiscale cardiac model for fast personalisation and exploitation," in *MICCAI, LNCS*, 2016.
- [36] S. Giffard-Roisin, T. Jackson, L. Fovargue, J. Lee, H. Delingette, R. Razavi, N. Ayache, and M. Sermesant, "Non-invasive personalisation of a cardiac electrophysiology model from body surface potential mapping," *IEEE Trans Biomed Eng*, pp. 292–9, 2016, in press.
- [37] D. Neumann, T. Mansi, L. Itu, B. Georgescu, E. Kayvanpour, F. Sedaghat-Hamedani, A. Amr, J. Haas, H. Katus, B. Meder, S. Steidl, J. Hornegger, and D. Comaniciu, "A self-taught artificial agent for multi-physics computational model personalization," *Med Image Anal*, vol. 34, pp. 52–64, 2016.
- [38] J. Ortiz-Perez, J. Rodriguez, S. Meyers, D. Lee, C. Davidson, and E. Wu, "Correspondence between the 17-segment model and coronary arterial anatomy using contrast-enhanced cardiac magnetic resonance imaging," *JACC Cardiovasc Imaging*, vol. 1, pp. 282–92, 2008.
- [39] M. Lorenzi, N. Ayache, G. Frisoni, X. Pennec, and A. D. N. I. (ADNI), "LCC-Demons: a robust and accurate symmetric diffeomorphic registration algorithm," *Neuroimage*, vol. 81, pp. 470–83, 2013.
- [40] C. Hoogendoorn, N. Duchateau, D. Sanchez-Quintana, T. Whitmarsh, F. Sukno, M. De Craene, K. Lekadir, and A. Frangi, "A high-resolution atlas and statistical model of the human heart from multislice CT," *IEEE Trans Med Imaging*, vol. 32, pp. 28–44, 2013.
- [41] T. Heimann, P. Mountney, M. John, and R. Ionasec, "Real-time ultrasound transducer localization in fluoroscopy images by transfer learning from synthetic training data," *Med Image Anal*, vol. 18, pp. 1320–8, 2014.