



HAL
open science

A High-Resolution Atlas and Statistical Model of the Human Heart From Multislice CT

C. Hoogendoorn, Nicolas Duchateau, D. Sanchez-Quintana, T. Whitmarsh, F. Sukno, Mathieu de Craene, K. Lekadir, A. Frangi

► **To cite this version:**

C. Hoogendoorn, Nicolas Duchateau, D. Sanchez-Quintana, T. Whitmarsh, F. Sukno, et al.. A High-Resolution Atlas and Statistical Model of the Human Heart From Multislice CT. *IEEE Transactions on Medical Imaging*, 2013, 32 (1), pp.28-44. 10.1109/TMI.2012.2230015 . hal-02282416

HAL Id: hal-02282416

<https://hal.science/hal-02282416>

Submitted on 11 Sep 2019

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.

A High-Resolution Atlas and Statistical Model of the Human Heart From Multislice CT

C. Hoogendoorn, Nicolas Duchateau, D. Sanchez-Quintana, T. Whitmarsh, F. Sukno, Mathieu de Craene, K. Lekadir, A. Frangi

► **To cite this version:**

C. Hoogendoorn, Nicolas Duchateau, D. Sanchez-Quintana, T. Whitmarsh, F. Sukno, et al.. A High-Resolution Atlas and Statistical Model of the Human Heart From Multislice CT. IEEE Transactions on Medical Imaging, Institute of Electrical and Electronics Engineers, 2013, 32 (1), pp.28-44. 10.1109/TMI.2012.2230015 . hal-02282416

HAL Id: hal-02282416

<https://hal.archives-ouvertes.fr/hal-02282416>

Submitted on 11 Sep 2019

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.

A High-Resolution Atlas and Statistical Model of the Human Heart from Multislice CT

Corné Hoogendoorn, Nicolas Duchateau, Damián Sánchez-Quintana, Tristan Whitmarsh, Federico M. Sukno, Mathieu De Craene, Karim Lekadir, and Alejandro F. Frangi* *Senior Member, IEEE*

Abstract—Atlases and statistical models play important roles in the personalization and simulation of cardiac physiology. For the study of the heart, however, the construction of comprehensive atlases and spatio-temporal models is faced with a number of challenges, in particular the need to handle large and highly variable image datasets, the multi-region nature of the heart, and the presence of complex as well as small cardiovascular structures.

In this paper, we present a detailed atlas and spatio-temporal statistical model of the human heart based on a large population of 3D+time multi-slice computed tomography sequences, and the framework for its construction. It uses spatial normalization based on non-rigid image registration to synthesize a population mean image and establish the spatial relationships between the mean and the subjects in the population. Temporal image registration is then applied to resolve each subject-specific cardiac motion and the resulting transformations are used to warp a surface mesh representation of the atlas to fit the images of the remaining cardiac phases in each subject.

Subsequently, we demonstrate the construction of a spatio-temporal statistical model of shape such that the inter-subject

and dynamic sources of variation are suitably separated. The framework is applied to a 3D+time data set of 138 subjects. The data is drawn from a variety of pathologies, which benefits its generalization to new subjects and physiological studies. The obtained level of detail and the extendability of the atlas present an advantage over most cardiac models published previously.

Index Terms—Computed tomography, heart, atlases, population analysis, probabilistic and statistical methods, registration, segmentation, computational physiology

I. INTRODUCTION

INTEREST in cardiac electrophysiological and mechanical simulation has risen significantly over the past decade, enabled by increased power and availability of computing resources and developments in distributed computing. This interest has been partly structured around initiatives to develop computational physiology models of the (human) body and its systems, collectively known as Virtual Physiological Human (VPH) initiatives [1]. The cardiovascular system is one of twelve systems identified within the International Union of Physiological Sciences (IUPS) Physiome Project [2], and is subdivided further into cardiac and vascular modeling.

Atlases play an important role in computational physiology of any organ, including the heart [3], [4], [5]. They provide insight regarding the division into and localization of substructures within the body, within an organ, or within a structure. For computational physiology studies to return accurate results, the use of such a ‘map’ is of great importance, as different structures within an organ have different electrical and mechanical properties.

Statistical atlases provide not only an average layout of structures within an encapsulating structure, but also encode deviations from this average. This provides a means to deform the encapsulating structure, within statistically justified bounds, and have the substructures deform and move accordingly based on their statistical correlation. This enables two very important applications of computational physiology: first, simulation studies can be personalized geometrically. This is in addition to incorporating subject-specific measurements to parameterize models of electrical conduction and of mechanics. For geometrical personalization, the atlas can be matched to medical imaging data. Typically, only a subset of atlas structures can be matched explicitly to the image, and the remaining structures are moved by virtue of their statistical correlation. Secondly, the statistics learned from the population can be used to generate populations for virtual population

This work was funded by the Spanish Ministerio de Economía y Competitividad (MINECO) through project STIMATH (TIN2009-14536-C02-01), and through project cvREMOD (CEN-20091044) under the CENIT programme of the MINECO and the Centro para el Desarrollo Tecnológico Industrial (CDTI), and by the European Commission through its 7th Framework Programme project euHeart (FP7/ICT-2007-224495). T. Whitmarsh was funded by the Instituto de Salud Carlos III through PFIS predoctoral grant FI09/00757. K. Lekadir is funded by a Juan de la Cierva grant of the Spanish MINECO. Asterisk indicates corresponding author.

C. Hoogendoorn and K. Lekadir are with the Center for Computational Imaging & Simulation Technologies in Biomedicine (CISTIB), Information and Communication Technologies Department, Universitat Pompeu Fabra, Barcelona, Spain, and with the Networking Biomedical Research Center on Bioengineering, Biomaterials and Nanomedicine (CIBER-BBN), Barcelona, Spain.

N. Duchateau was with CISTIB and CIBER-BBN, Barcelona, Spain. He is now with Hospital Clínic i Provincial de Barcelona, Institut d’Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Barcelona, Spain.

D. Sánchez-Quintana is with the Department of Anatomy, Faculty of Medicine, Universidad de Extremadura, Badajoz, Spain.

T. Whitmarsh was with CISTIB and CIBER-BBN, Barcelona, Spain. He is now with the Department of Engineering, University of Cambridge, Cambridge, United Kingdom.

F.M. Sukno was with CIBER-BBN and CISTIB, Barcelona, Spain. He is now with the Vision Systems Group, Dublin City University, Dublin, Ireland.

M. De Craene was with CISTIB and CIBER-BBN, Barcelona, Spain. He is now with Medisys, Philips Research, Paris, France.

* A.F. Frangi is with CISTIB, Department of Mechanical Engineering, The University of Sheffield, Sir Frederick Mappin Building, Sheffield S1 3JD, United Kingdom. He is also with CIBER-BBN, Barcelona, Spain (e-mail: a.frangi@sheffield.ac.uk).

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

Copyright (c) 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

studies. Such a population can be controlled by the user, either to statistically match an existing population, to increase sample size, or to generate more extreme cases.

For both applications, a statistical atlas provides a final advantage in the postprocessing and analysis of any simulation results that are generated. The encoding of population variation implies that spatial relationships between instances are known. This ensures a straightforward and reliable way of warping study outputs into a common reference frame in which they can be further presented, compared and analyzed.

The construction of a statistical atlas from a population of images requires that each of the structures in the atlas be segmented (labeled) in each of the images in the database. To do this manually is generally considered an impossibility for 3D and 3D+time atlases. To bypass this problem, atlas-based segmentation methods provide a solution. First, one applies *spatial normalization* to the population. This is the synthesis of an average image from the population, usually based on image registration techniques. Additionally, it provides the spatial relationships between the population and this average. As one then *labels this atlas*, one may consider all the population images segmented through these spatial relationships. By representing the atlas as a surface mesh, one uses point distribution analysis as the approach to *statistical analysis*.

The spatial normalization of a population of 3D+time cardiac images is faced with a number of challenges. Some of these have a counterpart in the construction of brain atlases; others are specific to the cardiac case:

- Cardiac and vascular structures vary significantly in size and geometry, with large global variations observed throughout the populace and both global and local variations due to pathology. In order to extract and represent these accurately, high-resolution image data is required, and the algorithms used must be able to handle this.
- Correct topological relationships between structures are of the utmost importance for the atlas to be used in cardiac simulations. Seamless transitions between structures while maintaining mesh quality are required. Mesh extraction, simplification and smoothing algorithms must be able to provide this.
- High anatomical variability in the heart requires a large population for its statistical modeling. This makes it very desirable to minimize the greatest difference to be resolved, which can be achieved by selecting a suitable initial estimate of the spatially normalized image.
- A large population also means that case-by-case parameter tuning becomes increasingly unfeasible. Consequently, this must be addressed in an automated manner or the algorithms must be robust to potentially suboptimal parameterization.
- Cardiac motion presents a challenge in statistical analysis, as it means there are at least two sources of shape variation: inter-subject and temporal. Additionally, it introduces an increase in data set size with respect to 3D imaging. Together with the large population requirement and the high resolution requirement, this means that the algorithms must be able to handle very large data sets.

In summary, the algorithms must be robust to suboptimal

parameterization and large variations, and able to handle very large data sets. It means that the atlas representation must have the flexibility to permit the statistical modeling as well as the simulation studies. Finally, the statistical modeling must be able to handle spatio-temporal data in a correct manner.

This work presents a statistical atlas of the human heart with the detail and flexibility necessary for personalized cardiac simulation, and the framework through which we constructed it. Elements within the framework are freely exchangeable for more advanced techniques, yet in the current form we have opted for some well-established approaches. In addition to the framework for the atlas construction, we present a framework for its validation. Also, we have made the atlas itself publicly available. An additional minor contribution is the automated selection of a reference image as an initial estimate of the mean.

The remainder of this paper is organized as follows. The next section will provide an overview of related work, on cardiac statistical atlases and on spatial normalization of a population of images. Section III will delineate the steps taken in the construction of the atlas, and the considerations implied in the automation of this process at a large scale. Section IV outlines the properties of the atlas and their relationship to the atlas' applications, whereas V covers experiments regarding the validation of the atlas, and in Sec. VI we discuss possible improvements and extensions.

II. RELATED WORK

The traditional atlas in the medical field consists of a collection of commented illustrations—drawings or photos—of the structure of choice, either intact or dissected, as a whole or in close-up. Invariably these are two-dimensional representations of the structures *ex vivo*, without the possibility to see variations, active or passive movement, or points of view other than those presented. Computerized medical imaging techniques provide solutions to each of these issues. In this context, we see related work in the development of digital three- and four-dimensional atlases, as well as efforts in computational anatomy to characterize population variability.

Digital cardiac atlases have been developed for a variety of purposes. Recently, atlas-based segmentation using voxel-based atlases was introduced in the cardiac domain [6], [7], [8], [9]. The construction of a voxel-based atlas of cardiac structures was only outlined in [6], using only three cardiac labels (two ventricular blood pools and left ventricular myocardium).

More recently, voxel-based atlases of cardiac fiber orientation have been constructed, first from canine [10] and later from human [11] *ex vivo* diffusion tensor MRI data. These atlases provide an important component for simulation of both cardiac electrophysiology and mechanics, despite the limitation of being acquired *ex vivo*.

A special case of cardiac atlases is the Cardiac Atlas Project (CAP) [12]. It adheres more to the traditional idea of an atlas in that it is a collection of maps—a database—rather than a single mean map, or a statistical description of the shapes and their relationships. The construction of a statistical shape

model of the left ventricle from part of this data (200 training shapes) was demonstrated in [13].

A. Cardiac Statistical Atlases

In [14], first a voxel-based and then a surface-based mean atlas was constructed of the same structures used later in [6]. The surface-based atlas, together with deformations obtained during construction of its voxel-based counterpart, was used to generate point correspondence across the population such that a point distribution model could be built. Virtually the same technique was used to generate a four-chamber model in [15].

Similarly, Perperidis used voxel-based atlases of images in the first cardiac phase to generate surface-based training data for a 3D+time statistical model covering both inter-subject and intra-subject (functional) deformations [16], by applying PCA first to the mean shapes of each subject, and then to all phases of all subjects, minus their mean shapes. Another 3D+time statistical model of cardiac shape was based on direct bilinear decomposition of the surfaces in the training set [17], using a higher order Singular Value Decomposition. A similar model was used for segmentation of cardiac MR images by Zhu *et al.* [18].

A slightly different approach was followed in [19], where a single surface-represented segmentation was fitted to other segmentations to generate the point correspondence. With the surface sampling depending on the curvature observed in the initial shape, a bias in local mesh density may be introduced. Later, [20] used a similar technique but required full manual segmentation only of the initially chosen image.

In [21], sampling of the curves obtained by intersection of manual delineations with cut planes is employed to generate point correspondence.

A recurring requirement in nearly all of these approaches is the need for (manually) segmented images, as is also apparent from the review in [22]. The same holds true for the various methods suggested to achieve point correspondence on shapes represented by meshes. Overall, this has had the effect of limiting training set sizes in statistical shape modeling. Some effort was made to enlarge such training sets by adding synthetic variations [23], however, one could debate the plausibility of these variations. For further reading on surface-based atlases, including methods that generate surface point correspondence without prior volumetric point correspondence, please see [22].

In this work, we extend the works of [14] and [24], in that we construct an atlas to generate correspondences throughout our training population, and we use a surface-based representation so that this correspondence is carried by the surface mesh vertices. Only one volume needs to be segmented explicitly, creating the possibility to use arbitrarily large training sets.

B. Spatial Normalization

When a synthesized anatomy is used for an atlas, it is usually generated to represent a central tendency of a population. Estimating this central tendency is generally approached as a registration problem. This field has been a very active one since the landmark paper for average image construction by Guimond *et al.* [25]. Over the period since this

work, a shift can be observed from image averaging towards transformation averaging, initiated by Christensen *et al.* [26]. Physical constraints have been introduced in the registration approaches, including diffeomorphic and other, more tissue-specific, constraints. Imaging modalities for the construction of atlases have become more varied as development on the mathematics of their deformation and similarity progressed.

Within the field of spatial normalization we can distinguish two main approaches: true groupwise registration on the one hand, and a processing of pairwise registrations on the other hand. The contributions for the latter are mostly strategies to process the outcomes of the pairwise registrations.

Using pairwise registrations starts out from a single reference image. Typically this one is chosen from the population at hand, either by an expert [27], or by an automated method; this can be based on subspace exploration [28] or on an information-theoretical basis [29].

Once the reference has been established, registration of each member of the population to this reference is carried out, producing a deformation for each population member. The combination of these deformations typically leads to a mean deformation to be applied to the reference image, ideally transforming it into the mean image directly. As demonstrated in [25], this is not necessarily the case. Recently, various interpretations of the use of intermediate means have been proposed. Jongen *et al.* [30] suggested the construction of a mean from a small subset of the population, before registering the entire population to this mean estimate.

This is different from intermediate means used in manifold-based approaches. Jia *et al.* [31] proposed the derivation of a tree to guide registration step by step. The registration of a population member to the population mean is subdivided into steps of registration to neighbors nearer to the mean, thus constructing a path through the population. Child-parent relationships are determined through the stabilization of a clustering method. This contrasts with the approach of Wu *et al.* [32], who computes a Minimum Spanning Tree based exclusively on image intensity differences between images.

Other approaches explore the number of clusters that a population could be divided into. Atlas stratification [33] was an initial approach to only do this clustering, based on k -means. Sabuncu *et al.* [34] introduced the automated atlas into a similar framework, but used a generalization of Estimation Maximization to discover the clusters.

True groupwise approaches tend to start from an average over the non-deformed images in the population. From that point onwards, Joshi *et al.* [35] update the deformations by explicit minimization of deformation magnitudes with respect to the Fréchet mean of the population, thus obtaining that mean. Lorenzen *et al.* [36] extended this work to handle multimodality image data, while Fletcher *et al.* [37] replaced the Fréchet mean with the Fréchet median to ensure population membership of the atlas.

The above approaches use an explicit minimization of the average deformations required to deform a population member to the population-representative atlas. Other true groupwise approaches construct the atlas by maximization of a groupwise similarity criterion. This is either similarity of the population

to the atlas via information theoretic criteria [38], [39], or a single measure quantifying similarity between all pairs [40]. In either case, an additional constraint is required to enforce zero average deformation.

Our framework can be changed to incorporate any of the above methods. However, in its current form we present it using the standard approach of Guimond *et al.* [25], with multi-scale diffeomorphic B-spline registration [41], [42].

III. ATLAS CONSTRUCTION

This section covers the steps required to construct a population average image using image registration and the subsequent generation of an atlas represented using a surface mesh. Although various strategies exist for registering a set of images to an initially unknown average [43], [44], [38], [35], [40], [36], [33], [45], [46], [31], we chose the classical approach of registration to a chosen reference, as did Guimond *et al.* [25]. The simplicity of this approach holds some advantages:

- 1) The process is transparent. Unlike stratification [33] and manifold-based methods [46], it does not introduce more parameters than those required for the pairwise registration and the steps to update the mean image. Additionally, there is no interaction between these sets of parameters.
- 2) The complexity is kept to only one registration per subject per iteration, rather than a full cross-registration of all or a subset of subjects. Also, the optimization in a groupwise approach is a much higher-dimensional optimization problem, with ramifications for the robustness.
- 3) The strategies of stratification [33] and of using intermediate local averaging [31], [46], [47], [32] employ distance measures to determine which pairwise registrations must be carried out. It is unclear whether cardiac image volumes would be at an advantage or disadvantage for stratification compared to brain image volumes, given the increased number of disjoint regions of clearly different appearance.

The process is outlined in Fig. 1, with parameters listed in the last subsection of this section (III-F). It was implemented on a grid computing facility using Fura software (Grid Systems S.A., Palma de Mallorca, Spain) to handle the distribution. The facility was equipped with 12 nodes with 2 64-bit quad core processors each, with 16 Gb of shared memory per node.

Note that we synthesize a mean first cardiac phase image only. This is because of the greater stability of this phase in retrospective gating. Thus, our means are computed over the 138 subjects rather than the 2070 image volumes.

For further reading on cardiac image registration, we refer to Mäkelä *et al.* [48]; for aspects of medical image registration in general, we refer to [49], [50], [51], [52].

A. Imaging Data

The database of imaging data used in this study was retrospectively collected from a clinical cohort of 138 consecutive patients that underwent a CT examination as part of their routine diagnostic protocol for suspected coronary artery disease, and follow-up. The resulting population distribution,




Fig. 1. Flow diagram of the mean image synthesis. The affine transformation obtained initially is re-used as an initial transform in the non-rigid registration.

as well as imaging parameters, are outlined in Table I. All information was anonymized before its transfer from the clinic to our group.

While our database is not as large as that of the CAP [12], our image resolution is 2.5 to 5 times higher in-plane [53], [12], and 3 to 4 times higher axially [53], [54], [12]. This enables us to capture more anatomical detail than would be possible with clinical resolution MR data. Similarly, ultrasound has the advantages of noninvasiveness and higher temporal resolution, but these are undone by the speckle patterns which make automated further processing—specifically inter-subject registration—extremely difficult.

B. Reference Selection

The selection of a suitable reference is one of multiple approaches to minimize bias in the synthesized mean volume. The first box in Fig. 1 is therefore an attempt to select the reference already close to the unknown mean. For this we use a heuristic based on a Groupwise Mutual Information (GWMI) score, an extension of Mutual Information [55] that describes the amount of information between a single image and a set of images [29]. We compute the GWMI score over affinely registered volumes. The affine registration is required to remove the detrimental influence of pose and size on the score.

TABLE I
 DETAILS OF DATA. CAD: CORONARY ARTERY DISEASE; MI:
 MYOCARDIAL INFARCTION

<i>Imaging</i>	Machine	Toshiba Aquilion 64 ¹
	Tube voltage	120 kV
	Tube current	400-430 mA
	Contrast agent	Xenetix 350
	- quantity	80-100 ml
	- rate	5 ml/s
<i>Reconstruction</i>	Workstation	Vitrea ²
	Resolution	0.4 × 0.4 × 2.0 mm
	In-plane grid	512 × 512 pixels
	Axial grid	65.3 ± 11.3 slices
	Volumes per cycle	15
<i>Population</i>	Size	138
	Men/women	94/44 (64% men)
	Age	59.4 ± 13.0
	Pathology CAD only	20%
	Pathology CAD+MI	20%

¹ Toshiba Medical Systems, Tochigi, Japan

² Vital Images Inc., Minnetonka, MN, USA

In short, the GWMI between image \mathbf{I} and image set \mathcal{J} is defined using the Shannon entropy H over voxel intensities as

$$\text{GWMI}(\mathbf{I}, \mathcal{J}) = H(\mathbf{I}) + H(\mathcal{J}) - H(\mathbf{I}, \mathcal{J}). \quad (1)$$

With the probability of an intensity j in the set defined as

$$p(j) = \sum_{\mathbf{J} \in \mathcal{J}} p(j|\mathbf{J})p(\mathbf{J}), \quad (2)$$

and with an *a priori* uniform probability distribution over the candidate images ($p(\mathbf{J}) = \frac{1}{|\mathcal{J}|}$ with $|\cdot|$ denoting cardinality), we are effectively summing the histograms of the images in the set \mathcal{J} defined on the same grid of bins:

$$\text{Histogram}(\mathcal{J}) = \sum_{\mathbf{J} \in \mathcal{J}} \text{Histogram}(\mathbf{J}). \quad (3)$$

The joint probabilities of an intensity i in the candidate reference with an intensity j in the set \mathcal{J} are defined analogously as

$$p(i, j) = \sum_{\mathbf{J} \in \mathcal{J}} p(i, j|\mathbf{J})p(\mathbf{J}), \quad (4)$$

also coming down to a summation of the joint histograms:

$$\text{JointHistogram}(\mathbf{I}, \mathcal{J}) = \sum_{\mathbf{J} \in \mathcal{J}} \text{JointHistogram}(\mathbf{I}, \mathbf{J}). \quad (5)$$

In the case where \mathcal{J} consists of only one image, this reduces to the standard MI. The method takes one parameter, the number of histogram bins.

The score represents the amount of information each volume carries with respect to the remainder of the volumes, taking into account the variation observed in this set. To make the process even more robust to the reference volume used in the affine registration, we registered all images affinely to 11 randomly selected volumes, and applied GWMI to each of the 11 resulting sets of 138 volumes, obtaining 11 rankings. The reference ultimately chosen for the atlas construction was the volume with the best mean rank. Figure 2 illustrates the varying rankings of subjects presenting a rank standard

deviation of less than 10 positions with the dash-dot lines. The solid lines (connected to the associated dash-dot lines using dotted lines) illustrate the mean ranks for these subjects. It shows that high and low ranks are relatively stable [29].

C. Global Registration

With the reference subject chosen from the population as described in the previous section, we register the volume corresponding to each subject's first cardiac phase to the first phase of the reference subject. Corresponding to the second box in Fig. 1, the first step of this registration is a global registration, requiring few parameters to optimize. It has been shown that such an approach increases the robustness of the overall scheme [41].

We aim to model shape variation through a statistical shape model rather than a statistical deformation model (as in [56], [57]). As the inverses of both the global and local transformations are applied to the mesh corresponding to the average image, it is of no importance which global transformation model is used. Hence, we use affine registration to remove global variations, minimizing as much as possible the deformation remaining to be resolved in non-rigid registration.

The small number of degrees of freedom of an affine transformation make the registration suitable for a gradient descent optimization. The similarity metric is mutual information (MI) [55], of which the gradient is computed using Mattes' method [58]. As it was developed for inter-modality image registration, MI is capable of handling varying intensities for corresponding structures, or in other words, to register images with different image intensity profiles. In our dataset, we observed image intensity variations throughout the population in the blood pools, due to varying contrast agent concentrations (note the 25% variation reported in Table I). Additionally, some subjects presented with pacing devices, leading to some very strong local intensity variations.

D. Diffeomorphic Nonrigid Registration

Following the affine registration, in the third box in Fig. 1, nonrigid registration is used to resolve the remaining variations between the reference and population images. With a large number of registrations to be carried out, a fast registration approach is desirable. The multiscale approach using B-splines [41] fulfils this criterion.

A later work by Rueckert *et al.* [42] focused on imposing diffeomorphic constraints on B-splines, meaning that the resulting deformations are smooth and invertible. From an anatomical as well as a technical viewpoint, this is a very important property. Borrowing from the work of Choi and Lee [59], the injectivity of a B-spline-based deformation is guaranteed if the local deformation is limited to 0.4 times the spacing of the control points. As the group of diffeomorphic transformations is closed for the composition operator (\circ) [60], this allows for a combination of the multiscale and the diffeomorphic B-spline approaches. That is, the composition of diffeomorphic B-splines, regardless of equal or varying control point spacings, produces a diffeomorphic transformation. However, it is generally not possible to represent this composition as a single diffeomorphic B-spline.

The multiscale approach is robust and allows for fast registration. The constraints on control point movement are implemented straightforwardly using the Insight Toolkit (Kitware, Inc., Clifton Park, NY, USA) [61] and a bounded quasi-Newton optimizer [62]. Klein *et al.* [50] showed that this optimizer provides a good balance between speed and accuracy when used for registration using B-splines and MI.

In the original formulation of the diffeomorphic B-splines, new splines were added until convergence of the registration. However, we choose to maintain a fixed set of transforms. This results in control over the maximum possible local deformation, and therefore prevents the construction of outliers that could corrupt the computation of the mean deformation.

Using T for transformations obtained by composition of transforms, A^s for the affine transformation from the reference subject to subject s , B^s for a B-spline based deformation from the reference subject to subject s , $B^{s,\tau}$ for the B-spline based deformation for the nonrigid deformation from frame $\tau - 1$ to frame τ of subject s , the total transformation to frame τ of subject s is given by

$$T_{\text{total}}^{s,\tau} = T_{\text{intra}}^{s,\tau} \circ T_{\text{inter}}^s \circ A^s \quad (6)$$

with

$$T_{\text{intra}}^{s,\tau} = \bigcirc_{t=2}^{\tau} B^{s,t} \quad (7)$$

and

$$T_{\text{inter}}^s = \bigcirc_{i=1}^{n_b} B_i^s. \quad (8)$$

Here, n_b is the number of B-spline transforms we use to compose the deformation from the reference to subject s , and \bigcirc is the big version of the composition operator \circ . After the affine registration has brought the images into a global alignment, by controlling the spline control point spacing and n_b we control the maximum total local deformation. We will show later that we observed a convergence rate similar to that reported by Guimond *et al.* [25].

After the non-rigid registration, we represent the inter-subject transformations T_{inter}^s as displacement vector fields φ^s , thus simplifying further processing in the absence of the composition operators. The use of B-splines to model our deformation reduces the number of degrees of freedom in the optimization problems by multiple orders of magnitude as compared to a parameter-free approach that directly produces a vector field, like the diffeomorphic demons [63].

E. Mean Deformation

The vector fields $\varphi(s)$ representing $T_{\text{inter}}(s)$, and defined on the same grid as the reference image, belong to the group of diffeomorphic vector fields. This group is closed only for the composition operator (\circ) [60], and therefore the euclidean metrics of addition and division cannot be applied directly. If the deformation is a geodesic in the metric space defined by the deformation model used for the registration, one can use the log-Euclidean framework of Arsigny *et al.* [64] to work in the tangent space of these diffeomorphic vector fields, in which Euclidean metrics can be applied. The mapping to and from the tangent space is provided by the logarithm and exponential operators, respectively. Due to the multi-scale approach to

the registration, we should assume that we approximate the geodesic closely enough for the log-Euclidean framework to be applicable. We will later comment on the accuracy of the logarithm computation (Sec. III-F). This leads to the definition of the mean vector field as

$$\bar{\varphi} = \exp\left(\frac{1}{n_s} \sum_{i=1}^{n_s} \log \varphi^s\right), \quad (9)$$

corresponding to the 5th through 7th boxes of Fig. 1. Note again that the mean is computed over the first-phase images only.

Given the mean vector fields obtained in the iterations up to i , and the current approximation of the mean image $\tilde{\mathbf{I}}_{i-1}$, the approximation is updated in boxes 9 and 10 of Fig. 1 using

$$\tilde{\mathbf{I}}_i = \bar{\varphi}_{0\dots i}^{-1} \circ \tilde{\mathbf{I}}_0, \quad (10)$$

where

$$\bar{\varphi}_{0\dots i}^{-1} = \bar{\varphi}_i^{-1} \circ \bar{\varphi}_{0\dots i-1}^{-1}. \quad (11)$$

For completeness, $\bar{\varphi}_{0\dots 0} \equiv \bar{\varphi}_0$. The development of $\tilde{\mathbf{I}}_i$ is illustrated in Fig. 3.

Arsigny's algorithm [64] for fast computation of the vector field logarithm is an inverse scaling and squaring method, exploiting the equivalence

$$\log \varphi = 2^n \log \varphi^{2^{-n}}. \quad (12)$$

For sufficiently large n , the logarithm can be approximated by

$$\log \varphi^{2^{-n}} \approx \varphi^{2^{-n}} - Id, \quad (13)$$

where Id stands for the identity. The square roots are computed using a gradient descent approach, which uses the inverse of the transform. The vector field must extend beyond the image domain and smoothly reduce to zero magnitude to avoid discontinuities in the image domain in the inverse vector field. Effectively, the actual deformations must be contained in a box, rigid under logarithm, averaging and exponential. The finite support of B-splines provides this without any additional cost in registration complexity.

F. Parameters and Settings

Each step in the process of synthesizing the mean image volume depends heavily on good parameters, with the image registrations a prime example. However, optimizing each of these parameters for each instance is unfeasible. Therefore it is imperative that only a few, if multiple, distinct configurations are used. As a consequence, for the majority of registrations these configurations are suboptimal.

In order to obtain the best possible registrations under these conditions, a small set of six first-phase volumes was used to obtain a configuration suitable for all pairwise registrations in the subset, which was the first instance of human intervention. Using a larger subset of 20 first-phase volumes, the useability of this configuration was verified. Table II lists the values obtained for the registration parameters.

For the inter-subject registrations in the mean image synthesis, image masks are defined based on the image intensities of the fixed and moving images, ranging from air (-1000




Fig. 3. The development of \tilde{I}_i , for $0 \leq i \leq 7$. After five iterations, the composed deformations, although diffeomorphic and less biased, become implausible. This is visible in the lateral wall of the left ventricle (framed in the white square for $i = 7$, but also visible for $i = 6$).




Fig. 2. Rankings and mean ranking for subjects with a ranking standard deviation below 10. One could consider the six subjects with both $\sigma_{\text{rank}} < 10$ and $\mu_{\text{rank}} < 10$ to be suitable subjects for the initial reference.

Hounsfield Units (HU)) to cortical bone (1500 HU). The mutual information similarity metric is computed from 50-bin histograms generated from voxels randomly sampled within the mask regions, at a rate of one per 30 voxels.

The non-rigid registration uses a composition of B-splines of increasing resolution. Due to the hard limit on the local deformation of each B-spline, we control the maximum local deformation. Using only one each of 32, 16 and 8 mm control point spacing, this would amount to a maximum of 22.4 mm $((32 + 16 + 8) \times 0.4)$. We assumed this to be insufficient, and added an additional allowance of 12.8 mm through an

additional B-spline with 32 mm spacing.

The segmentation of the atlas is used to provide the masks for the intra-subject registrations, while the greater histogram resolution is used to improve sensitivity to smaller differences.

An alternative would be to apply a small-range parameter sweep around the values found, for every registration. However, this would increase the computational load prohibitively, even if only one parameter is varied at any given moment.

For the elements of the framework that are not registration steps, we list the parameters in Table III. Error reductions are computed with respect to the error at initialization (identity vector field), and the iteration limits supersede this criterion. For the root order, Arsigny *et al.* [64] empirically reported convergence of the log computation with $n = 7$ in Eq. 12. We found this value suitable for our data as well, yielding vector norm errors under 5% for 93.7% of vectors when comparing φ^s to $\exp(\log(\varphi^s))$ in 20 subjects.

Note that the computation of the mean log vector field, the composition of the mean estimate with earlier estimates, and the update of the mean image (boxes 6, 9 and 10 in Fig. 1 require no parameters).

G. Atlas Segmentation

Through registration, we have obtained a mapping from the atlas coordinate system to each of the subjects in our population. Thus, by segmenting the atlas volume, we obtain a segmentation of each member of the population as well. We represent the segmentation not by a mask image but by a triangulated mesh. This allows us to define other structures in local coordinate systems, structures which are typically obtained through literature description or separate modeling approaches.

TABLE II
 REGISTRATION PARAMETERS AND VALUES

	<i>Parameter</i>	<i>Value</i>
All intersubject	Mask—lower threshold	-1000 HU
	Mask—upper threshold	1500 HU
	Histogram bin count	50
	Voxels per sample	30
Affine	Gaussian smoothing kernel σ	1.0 mm
	Min. optimization step length	0.001
	Max. optimization step length	1.0
32 mm inter (twice)	Gaussian smoothing kernel σ	1.5 mm
16 mm inter	Gaussian smoothing kernel σ	1.0 mm
8 mm inter	Gaussian smoothing kernel σ	0.5 mm
40 mm intra	Histogram bin count	60
	Voxels per sample	30
	Gaussian smoothing kernel σ	0.5 mm
	Mask	warped atlas, 10 mm dilated

TABLE III
 PARAMETERS AND VALUES FOR NON-REGISTRATION ELEMENTS

	<i>Parameter</i>	<i>Value</i>
Reference selection	Histogram bin count	50
Vector field logarithm	Inversion—step size	1.0
	Inversion—stopping criteria	99% error reduction
		10 iterations
	Squaring—step size	1.0
Squaring—stopping criteria	95% error reduction	
	10 iterations	
	Root order (n in Eq. 12)	7
Mean vector field inversion	Step size	0.3
	Stopping criterion	99% error reduction 100 iterations

By far most human interaction in the procedure is concentrated at this step. An anatomist specialized in cardiac anatomy (D. S.-Q.) segmented the synthesized mean image using GIMIAS v1.2.0b software [65]. In each slice of the image, the structures listed in Table IV were outlined using a free-hand polygon tool, which labeled the voxels on the same grid as the image.

For the coronary arteries, their intersections with the image planes were marked using a pointer tool. Subsequently these sets of points were connected and the arteries themselves were modeled by fitting knotted cones to the segments.

H. Mesh Construction and Processing

The manual segmentation of the mean image is converted to a mesh representation. There is a sizeable body of literature on the creation of ‘reasonable’ non-manifold meshes from multi-label medical imaging data. These meshes have a smooth surface, high accuracy with respect to the non-smooth (i.e., with staircase artifacts) surface, and with good element quality (usually polygon regularity). In addition to the challenge to achieve this in binary (foreground+background) data, the multi-label property of the input data presents the added challenge of smooth surfaces between materials while maintaining smooth outer surfaces. Examples of approaches to

this problem are multi-material extensions to Marching Cubes (MC) [66], [67], marching tetrahedra [68], dynamic particle systems [69], non-manifold dual contouring [70], [71], and multi-level partitions of unity [72].

We use a more straightforward approach, schematically presented in Fig. 4. We first generate a new segmentation by merging labels from the original segmentation. In this case, into three groups: left blood, right blood, and myocardium. The resulting masks are resampled using shape-based interpolation [73], [74], and then triangulated using MC [75]. Distance transforms of the original segmentations are then used to assign labels to the faces of the resulting mesh.

Subsequently, the mesh is coarsened using an Approximate Centroidal Voronoi Tessellation [76]. That is, each Voronoi cell is represented by a set of mesh faces in their entirety, and the seed of each cell coincides with the cell centroid. This approach has the advantage of allowing full control over the number of vertices in the coarse mesh, and as outlined in [76], is guaranteed to converge to a global optimum. Additionally, the resulting mesh is independent from the original MC mesh.

With the coarsening, the mesh is smoothed somewhat due to a lower sampling rate. However, staircase artifacts may remain especially when a large number of vertices is used. Therefore the meshes were submitted to a smoothing step [77]. Overall, the coarsening and smoothing introduced some local distortions, though in general these remained either within acceptable bounds, or they smoothed out undesired features which appeared as a result of local inconsistencies in the manual segmentation.

The transformations T_{inter}^s and A^s were used to warp the atlas mesh to the first cardiac phase of each subject s . Subsequently the transformations $T_{\text{intra}}^{s,\tau}$ were applied to these meshes to obtain the meshes pertaining to the remaining 14 phases of each subject.

IV. ATLAS PROPERTIES AND STATISTICAL MODEL

We developed the statistical atlas with the aim of using it in a computational physiology setting. In this section we outline the properties of the atlas that enable this. The atlas is available for downloading at our website: <http://www.cistib.upf.edu/cistib/index.php/downloads>.

A. Structures

The atlas has been designed to be instantiated in varying compositions, ranging from the entire heart to a single sub-structure. This serves the purpose of extracting or generating knowledge from images and simulation studies that only employ the structure of interest, without changing the atlas itself. This in turn enables the seamless incorporation of this new information into the atlas.

The structures outlined are listed in Table IV and colored separately in Fig. 5. For the left ventricle, the myocardium and blood pool were segmented, while for both atria, for the right ventricle and for the trunk of the aorta, only the blood pool was segmented. On CT images, the walls of these structures are not discernible and as such are estimated from the boundaries of the blood pool. We did this by extrusion of the surface along

TABLE IV
 ATLAS SUBPARTS (3D) AND ELEMENT STATISTICS

	Name	Vertex count	Vertex density (cm^{-2})	Face surface area (mm^2)
1.	Left Ventricular Myocardium	2032	10.66	4.85 ± 1.88
2.	Intraventricular Septum	1492	12.40	4.22 ± 2.19
3.	Right Ventricular Myocardium	3454	10.72	4.77 ± 2.34
4.	Left Atrial Myocardium	2010	11.68	4.43 ± 1.61
5.	Right Atrial Myocardium	1962	10.90	4.80 ± 1.98
6.	Aorta	2654	9.40	5.35 ± 1.46
7.	Vena Cava (inferior and superior)	742	11.25	4.69 ± 1.81
8.	Pulmonary Trunk and Artery	1032	11.46	4.44 ± 1.47
9.	Pulmonary Veins (left and right, inferior and superior)	1232	10.27	5.02 ± 1.79
10.	Anterior descending artery			
11.	Circumflex artery			
12.	Diagonal artery			
13.	Right coronary artery			
	Total	16113	10.44	4.78 ± 1.93

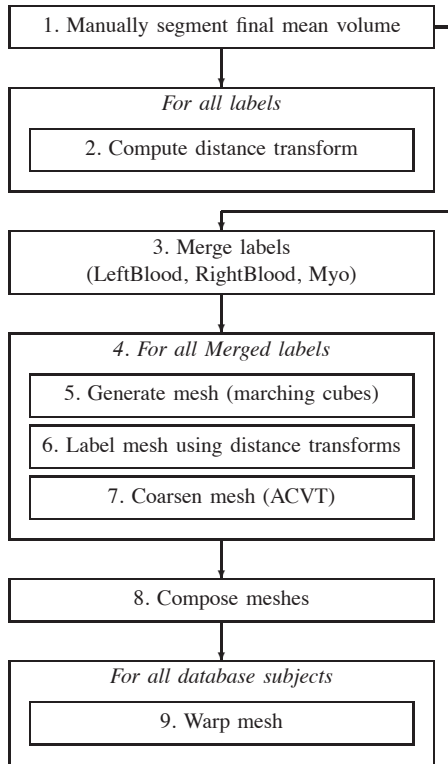


Fig. 4. Flow diagram of the meshing steps. ACVT: Approximate Centroidal Voronoi Tessellation. Initially the nine labels are merged into three; using the nine distance transforms, the vertices on the marching cubes meshes are labeled before coarsening. The atlas is finally composed by putting the three labeled, coarsened meshes together.

its normal direction, to generate the thicknesses reported in the literature. For the RV wall, the surface was extruded by 4 mm [78], for the aorta this was 3 mm [79] while the atria were extruded by 2 mm [80]. The vertex counts in Table IV already include the extruded parts.




Fig. 5. The full atlas mesh, corresponding to the synthesized mean image. Different colors indicate different structures.

B. Statistical Modeling

Paramount to the clinical value of computational physiology is the possibility to perform patient-specific simulations, where the computational domain is a patient-specific geometry, or to generate large virtual populations following the model population statistics. Approaches based on statistical shape models to segment patient geometries from image data have shown their worth many times over. For inherently dynamic shapes, naturally one should consider using a dynamic statistical shape model, such as those described in [16], [17], [18].

Here we reiterate some background on the bilinear spatio-temporal statistical shape model, constructed as per [17] from the framework's output.

In Sec. III-H we described how we obtained the meshes corresponding to the 2070 image volumes we began with. From these, we construct a spatio-temporal statistical model of shapes, decomposing our data along the axes of anatomical variation and motion-induced shape change using a bilinear model as described by Tenenbaum and Freeman [81]. They used the terms *style* and *content* to define the two sources of variation. We could consider here the style to be the temporal element, whereas the content is the anatomical variation that exists between subjects. By using a bilinear model, we maintain temporal information which would be lost if the more traditional Principal Component Analysis (PCA) [82] had been used.

However, for a rudimentary outlier detection, we used PCA [82] to generate a simple statistical model of the first-phase shapes, explaining 50% of total variation, amounting to five modes. The model parameters of these shapes were used to identify potential misregistration. A shape (and the associated subject) was marked as suspicious if any of its parameters exceeded four standard deviations from the mean. This led to the flagging of four subjects. In each of these cases, the affine registration had failed, and thus the subjects were discarded. We generated the spatio-temporal model from the remaining 134 subjects (2010 shapes).

Training of a bilinear shape model is done as described in [81], [17]. With 134 subjects and 15 phases each, the model reduces the data dimensionality to at most 147. For comparison, using Perperidis *et al.*'s approach [16] one would need to reduce the dimensionality from 2025, also from two sets of parameters. For the temporal synchronization, necessary due to the electrocardiogram-based retrospective gating of the CT data, we did use Perperidis' method [83] to identify the end-systolic and end-diastolic phases. More detail on the application of this method to surface meshes is provided in [17].

Figure 6 shows three phases of four subjects, together with bar plots for their subject and phase parameters. The applications of cardiac bilinear shape models have already been demonstrated in image segmentation [18] and motion analysis [84], [85]. In earlier works we have also demonstrated pipelines towards electrophysiological [86] and mechanical simulations [87] from surface-based cardiac segmentations.

Figures 7 and 8 show the first two modes of variation as mean plus and minus one and a half and three standard deviations of both the full heart and the left ventricular endocardium, of the traditional PCA-based shape model. Coincidentally in both models, the first two modes contain a strong component of elongation versus roundness.

V. VALIDATION

The quality of the atlas depends on the quality of the registration. In turn, the registration quality depends on each of the basic components of a registration algorithm: similarity metric, transform, and optimizer, and on the initial similarity between volumes [88].




Fig. 6. Three phases of two subjects, with their bilinear parameter sets derived from the full 134-subject data set. The dimensionality of the phase and subject parameters (11 and 115, respectively) are based on explaining 95% of total variance for each axis.




Fig. 7. The first two modes of variation in the linear statistical model of the full heart at end diastole, showing the mean plus and minus one and a half and three standard deviations.




Fig. 8. The first two modes of variation in the linear statistical model of the left ventricle at end diastole, showing the mean plus and minus one and a half and three standard deviations.

In this section we validate the atlas through evaluation of the registration. This consists of three elements:

- Random selection of volumes and slices. The dataset size does not allow the evaluation of all registrations. It is also undesirable to fully segment volumes for all the selected subjects. Thus, a random selection is made, based on some heuristics which make the subjects and slices therein useful.
- Evaluation of vector field inversion. In order to evaluate everything in one common space to avoid bias to the specific subjects involved, all segmentations and landmarks are warped to the atlas space. This is done using the inverse of the vector fields we obtain during registration. As the inversion is an optimization problem in itself, the inversion error presents a localization uncertainty. We report the inversion error at the locations of the landmarks and contours.
- Distance-based and overlap-based evaluation. After random slice selection and quantification of uncertainty in landmark localization, we report contour-to-surface and point-to-point errors and overlap accuracy to evaluate the registration quality.

Furthermore, we illustrate the convergence of the synthesis of the mean image.

A. Registration—Selection of Volumes and Slices

Both the optimization and the evaluation of all registrations carried out is unfeasible with large datasets. Therefore it is of importance that a meaningful subset of the data is selected for the evaluation. We used random sampling after identification of volumes and slices satisfying four simple qualification criteria: two for volumes, and two for slices.

A volume qualifies if the image intensity histogram peak pertaining to contrast enhanced blood could be identified, and the modal intensity and width of the peak fall within two interquartile distances from the population median for these two quantities. Additionally, the volume must contain a sequence of at least 20 qualified slices.

A slice qualifies if at least 3.5% of its pixels fall in the intensity range of contrast enhanced blood, as determined via the volume's intensity histogram. Additionally, the slice has to be part of a sequence of at least 20 in the volume.

The resulting sets of qualified slices are nearly guaranteed to intersect the cardiac and supracardiac blood pool. The threshold of 3.5% contrast enhanced blood pixels excludes slices where the enhanced blood is provided exclusively by the vena cava and aorta. The 20 slices cover 38 mm, which is much less than the cardiac blood pool should extend for, but excludes slices that were incorrectly qualified by the first slice criterion.

After determining qualifying volumes and slices, we draw randomly from the volumes. The sequences of qualifying slices in these volumes are divided into three equal parts and we randomly draw a slice from each part.

B. Registration—Measurements

1) *Contours*: From the qualifying subjects, 36 were selected randomly. Six researchers each segmented the slices

of six subjects in GIMIAS v1.2.0 [65] using the same tool as used to outline the atlas structures. The set of structures to be segmented correspond either to one or to multiple structures of the final atlas:

- 1) Aortic root
- 2) Left ventricular cavity
- 3) Left ventricular myocardium (constituting the left ventricular free wall and the interventricular septum)
- 4) Right ventricular cavity
- 5) Left atrial cavity (including pulmonary vein trunks)
- 6) Right atrial cavity (including trunk of vena cava)

For each of these structures, the boundary of the segmentation in the slice was extracted and warped to the reference space using the inverse of the vector fields obtained during registration. Subsequently the mean unsigned distance from the warped contour to the corresponding structure's surface in the atlas was recorded. These distances are reported for each of the structures in the top half of Table V.

2) *Overlap*: For each of the structures listed above, we also measured the overlap with the atlas segmentations. We warped the voxel-based segmentation into the atlas space using nearest neighbor interpolation. By using three labels (foreground, background in segmented slice, background elsewhere), we could identify and count false negative voxels ('fn') in addition to true and false positive voxels ('tp' and 'fp', respectively). These counts were combined into F-scores for the registrations in Table VI. The F-scores are defined as $F = \frac{2pr}{p+r}$, with precision $p = \frac{tp}{tp+fp}$ and recall $r = \frac{tp}{tp+fn}$. Thus, perfect matchings would yield an F-score of 1.

3) *Landmarks*: The contour-based evaluations have drawbacks in that they are not measuring the registration accuracy in a truly three-dimensional manner. By measuring the point-to-point accuracy of a set of anatomical landmarks, such evaluation is possible.

From the pool of subjects, 36 were selected for landmark-based evaluation. The same six researchers each processed 12 subjects; each subject was processed by two researchers and each pair of researchers shared either two or three subjects.

For each subject, five landmarks were identified and marked using GIMIAS v1.2.0 [65]: the endocardial apex, the aortic and mitral valve centers, and the origins of the left and right coronary arteries. Like the contours, the landmarks were warped to the reference space. We report the interobserver variations both prior to warping and after warping, and the point-to-point registration error in the reference space. These results are presented in the bottom half of Table V.

4) *Vector field inversion*: To warp the contours and landmarks from the subject space to the reference space, we require the inverse of the vector fields obtained during registration. The computation of this inverse is an optimization problem, and as such will introduce some uncertainty in the valuation. We quantify the error in the estimation of the inverse as the vector magnitude of $\varphi^s \circ (\varphi^s)^{-1}$ at the landmark and contour locations in the subject image space. The last column of Table V lists these errors.

TABLE V
 REGISTRATION AND VECTOR FIELD INVERSION ACCURACY OF SELECTED ANATOMICAL LANDMARKS AND CONTOURED STRUCTURES

	Interobserver variability (mm)		Registration error (mm)	Inversion error (mm)	
	(image space)	(ref. space)			
Contours	1. Aorta		2.19 ± 1.81	0.33 ± 0.19	
	2. LV cavity		1.83 ± 0.64	0.45 ± 0.35	
	3. LV myocardium		1.59 ± 0.45	0.48 ± 0.38	
	4. RV cavity		2.45 ± 0.87	0.72 ± 0.68	
	5. LA cavity		2.31 ± 1.73	0.43 ± 0.45	
	6. RA cavity		2.76 ± 1.21	0.51 ± 0.39	
Landmarks	1. LV endocardial apex	5.03 ± 2.97	4.34 ± 2.32	9.28 ± 10.12	0.47 ± 1.04
	2. Aortic valve center	4.40 ± 2.88	4.69 ± 3.14	5.40 ± 8.47	0.38 ± 0.32
	3. Mitral valve center	7.31 ± 3.87	7.68 ± 3.84	9.82 ± 12.04	0.27 ± 0.19
	4. Left coronary origin	3.76 ± 2.18	3.40 ± 1.84	5.50 ± 6.86	0.33 ± 0.20
	5. Right coronary origin	2.52 ± 1.30	2.54 ± 1.31	9.22 ± 7.88	0.89 ± 1.76

C. Registration—Results

The contour-based evaluation reported in Table V provides insight into the agreement on boundaries. The average distance of these contours ranges from 1.59 mm for the left ventricular area to 2.76 mm for the right atrial area. This compares favorably with the slice spacing of 2.0 mm. The difference in performance between the left ventricular area and the atrial and right ventricular areas is easily explained by the low complexity of the LV compared to the atria, and the much smaller influence of trabeculation on the LV blood-myocardium boundary strength.

The landmark errors compare well with the interobserver variation for the valves and the left coronary origin. The greater error in the apex is due to a sliding along the LV cavity wall; it is on the boundary but not at the apical point. During registration, boundaries are matched yet boundary curvature, by which observers would identify the apex, is not taken into account. The difference between interobserver variation and registration error of the right coronary artery (RCA) origin is due to the size of the RCA. While its position is stable at the base of the aorta, its appearance is often weak, thus offering no boundary for the registration to match.

The left ventricular apex and mitral valve center show a greater interobserver variability than the other landmarks. For the apex, this can be attributed to a dependence on image plane orientation; the observers are free to rotate the image planes to best visualize the landmark to be located, which for the apex tends to be a pure short-axis view. The apparent apex location varies along the endocardial wall with the deviation from this view. The interobserver difference in the mitral valve is due to its structure; whereas the closure of the aortic valve is defined by a curve where the three leaflets meet, the mitral valve's closure is a curved surface on which the valve centroid is not easily identified. In addition, its oblique location with respect to the image axes results in a high probability that the observers mark the centroid in different slices.

The overlap-based accuracy is reported in Table VI as F-scores in the atlas space. We have listed both the mean and standard deviations per subject per structure as well as an aggregate F-score per structure (column 'overall') and over all structures (last row of column 'overall'). Relatively little

change in accuracy between the first (reference) and fifth (final) iterations can be observed. A slight deterioration of the registration quality may be due to the inevitable reduction in image quality when the reference image is warped, which has affected all structures but the aorta, the structure with the most homogeneous blood pool. Seemingly 'difficult' structures to match properly have complex geometry (atria) or are thin (myocardium). A difficulty with the aorta was found to be the proximity of the descending part to the spinal column, and of the ascending part to the atria. 'Bleeding' into these structures was observed in some of the test subjects.

D. Bias Removal

The choice of a reference image introduces a bias. The iterative approach to synthesizing a mean image is designed to remove this bias. At full and perfect conversion, the mean vector field with which the mean image is updated should represent an identity transformation. We can therefore quantify the bias by measuring φ_i for each i .

The ideal mean vector field—the identity transformation—would contain exclusively zero vectors. Thus, its vectors' x , y and z components would follow a zero mean Dirac delta function with the peak at 0. In Fig. 9's left panel we show the deviation from 0 of the mean x , y and z component over seven iterations. It is clear that this deviation decreases strongly in the first iterations and then starts oscillating. In the right panel we plot the sample kurtosis of the distributions over the seven iterations, quantifying the 'peakedness' of the distributions. While the Dirac delta distribution's kurtosis is undefined, one can easily see that kurtosis increases with better approximation of the delta function.

In Fig. 10 we show the distributions per component in full over the iterations; after only three iterations the changes are no longer perceivable, which is similar to the findings of Guimond *et al.* [25].

In Fig. 11 we show how the distribution of the vector magnitude of φ_i changes over the iterations. The pattern of convergence is clear, yet after five iterations the percentage of longer vectors starts to grow again. This has a detrimental effect on the synthesized mean image and confirms the qualitative observation of increased artifact visibility from Fig. 3.

TABLE VI
 REGISTRATION ACCURACY QUANTIFIED BY F-SCORE OF OVERLAP OF SELECTED STRUCTURES

	<i>F-score</i>			
	<i>Initial</i> <i>mean ±</i> <i>standard deviation</i>	<i>overall</i>	<i>Final</i> <i>mean ±</i> <i>standard deviation</i>	<i>overall</i>
1. Aorta	0.850 ± 0.152	0.833	0.871 ± 0.105	0.867
2. LV cavity	0.925 ± 0.021	0.925	0.916 ± 0.023	0.917
3. LV myocardium	0.773 ± 0.055	0.775	0.773 ± 0.049	0.774
4. RV cavity	0.823 ± 0.070	0.825	0.817 ± 0.043	0.817
5. LA cavity	0.723 ± 0.131	0.723	0.699 ± 0.104	0.700
6. RA cavity	0.643 ± 0.167	0.636	0.628 ± 0.106	0.622
7. Left blood (1+2+5)	0.904 ± 0.056	0.900	0.895 ± 0.042	0.893
8. Right blood (4+6)	0.821 ± 0.079	0.821	0.806 ± 0.058	0.804
Overall		0.838		0.827




Fig. 10. Histograms of vector components per iteration. Qualitatively, it shows that the first three iterations resolve the greatest amount of deformation. After this, the majority of change occurs in the *y* component (sagittal plane).




Fig. 9. The absolute values of the bias (left) observed in the mean vector field decrease with iterations, down to less than half of the shortest side of a voxel in all directions after three iterations. The ‘peakedness’ of the distributions displayed in Fig. 10 are quantified through their kurtosis (right). It is clear that this measure continues to increase.

Therefore, we consider the synthesized mean image at $i = 5$ to be the final image.

VI. DISCUSSION

We have presented a mesh-represented statistical atlas of the human heart, and a framework for its construction from a large set of 3D+time multi-slice CT image volumes. We validated

the atlas using measures of registration accuracy and evaluated the iterative bias removal.

We have made the atlas publicly available through our group’s website.

A. Framework Performance

We have demonstrated the construction of a spatio-temporal statistical model of the atlas shape over a population of 134 subjects at 15 cardiac phases. To our knowledge this represents the largest 3D+time population to date to be used for cardiac statistical model generation.

The number of pairwise registrations in the framework scales linearly with the number of subjects. This makes the framework applicable to larger data sets than the set used in this work, such as the Cardiac Atlas Project database. The framework derives robustness from automation of the reference selection and the imposition of diffeomorphic constraints on the registrations.

Human interaction is limited to two distinct instances during the application of the framework. The first, tuning of registration parameters, is based on a subset of the population of a user-defined size. In the second instance, the manual labeling of the mean image, the level of detail and the number of structures can be controlled by the user. Existing atlases can be employed to further reduce the human input.




Fig. 11. The distributions of vector magnitudes in the mean vector fields (displayed as cumulative histograms) show convergence of the registration. In (b) we zoom in on the first three millimeters in (a) to better visualize that after only three iterations the majority of vectors is shorter than the voxel diagonal, and after five iterations around 70 percent of the vectors has a length shorter than the voxel diagonal. In (c) it shows that after five iterations, the tail of the vector magnitude distribution starts to widen. The synthesized mean image at this iteration is considered the end result of the spatial normalization.

Steps in spatial normalization that were previously typically done by human interaction, such as reference selection and assessment of convergence, were automated. An image-to-group similarity criterion was used to select a suitable reference, while convergence was based on the distribution of vector magnitudes in each iteration’s mean vector field.

We validated the atlas through measures of registration accuracy. Random sampling was employed to balance the need to cover a sufficiently large portion of the population, while maintaining the required human effort at a feasible level. In addition to random selection of subjects, within the subjects the selection of image slices was randomized as well. Domain knowledge was used to encourage the selection of slices presenting cardiac structures, while also enforcing an even distribution along the transversal axis of the body.

To complement the sample-based explicit registration evaluation, a PCA-based outlier detection system was applied to the first-phase shapes. While conceptually simple, it allowed the removal of a very small (3%) subset of the population for which registration had failed, contributing to the robustness of the overall framework.

By applying the spatial normalization only on the first frame of each subject’s volume sequence, we explicitly resolve geometric variations due to inter-subject variation separately from deformation due to cardiac function. This reduced computational requirements with respect to registration using an integrated spatio-temporal transformation model, while improving temporal consistency compared to normalizing the full set of 2070 image volumes.

B. Limitations

The atlas comprises 13 structures, of which 9 are outlined and 4 are traced; this represents only a small fraction of structures present in the heart but reflects a limitation resulting from the imaging data used. Populations imaged using complementary imaging techniques can be used to provide more structures, further subdivision or level of detail of the current structures, and increase the population size for the statistical modeling. Similarly, the aorta and pulmonary trunk are better modeled based on dedicated imaging studies.

Further limitations due to the imaging data are related to the price of the high spatial resolution of MSCT. Firstly, the temporal resolution is rather low at 15 frames per cycle. This is a well-known tradeoff that exists in all imaging modalities. Secondly, the use of CT data means the subjects were submitted to ionizing radiation, in addition to having contrast agent injected. As a result, all subjects are pathological, which influences both the population mean and the observed variation.

The coronary trees can only be traced over a short section; through the oblique orientation of the heart with respect to the image coordinate system, the coronary arteries quickly disappear ‘between slices’. Additionally, the structures are too small and are subject to topological variations, such that accurate registration of these becomes impossible.

Another problem of visibility, though not by obliqueness but by appearance, is that of the thin walls of the atria, right ventricle and trunks of great vessels. The appearance of these structures coincides with that of the pericardium, and thus our atlas only estimates these structures by extrusion of their blood pool boundaries.

A limitation of the framework is that it is currently unable to model the topological variations that are known to occur in the complex of left atrium and pulmonary veins [89], and in the coronary trees [90]. It deserves further study to unify a discrete statistical model of topological variations with the continuous statistical model of shape. The separation of the pulmonary vein trunks from the left atrium is a step towards accommodating such an extension.

The approach of iteratively removing bias from a chosen reference image may require more constraints than only the condition that the registrations produce diffeomorphic deformations. This may be related to the phenomenon we observe in the later iterations in Fig. 3. There, even the diffeomorphic transformations start producing seemingly implausible deformations, potentially as a result of shifting from larger to smaller differences to be resolved. As Hamm *et al.* [46] suggested, it may be desirable to explicitly constrain the deformations to generate only images that lie on the manifold generated by the input images.

Furthermore, this manifold may allow discontinuities in the deformations due to different positions of the heart relative to neighboring structures. Consider here the lungs and pulmonary vasculature laterally, the diaphragm inferiorly, and the sternum and rib cage frontally. Such discontinuities cannot be captured using the B-spline deformation model. Providing a region of interest that comprises only the heart may alleviate such a problem, as the B-spline needs only model deformation within this region. This could work under the assumption that the reference image can serve as an initial atlas, or by using an automated detection and/or rough segmentation method.

By using pairwise registrations to resolve the cardiac motion (the intra-subject registration), the time points at which labeled shapes are obtained are dictated by the imaging protocol. Currently, work on spatio-temporal registration [91], [92], [6] with a spatially and temporally continuous deformation model has not yet been demonstrated to handle high-resolution CT image data. However, the framework allows the pairwise approach to be replaced with a continuous approach when required and feasible.

The high degree of automation of the framework calls for methods, preferably also automated, to quantify accuracy. We used a rudimentary means to identify suspicious shapes by outlier detection. However, as it was rudimentary, it was limited to identifying only very obvious cases, and thus a more sophisticated approach suitable also for early detection is desired.

Finally, the framework presents the user with many decisions to be made; there are many steps involved, one can interchange methods, and each method has its parameters. This combines to a vast number of possible combinations that will not all produce the exact same atlas, and some combinations may even fail. However it is beyond the scope of this paper to present a quantification of the effect of each decision exhaustively.

C. Future Perspectives

The various steps in the framework are independent of each other; the methods used in each step can be replaced by other methods according to requirements without affecting the workflow ‘downstream’. One should note stratification [33] as an exception to the latter; while it can be plugged into the framework, it introduces branches in the workflow as it will lead to multiple atlases.

The atlas was developed with the aim of using it in computational physiology studies. We demonstrated the incorporation of fiber orientations in the left ventricular myocardium, which is of paramount importance in the simulation of both electrical and mechanical cardiac function.

The combination of an atlas with a statistical model facilitates personalized computational physiology studies by connecting the image analysis and simulation elements described by Young and Frangi [3].

The flexibility of the statistical model allows fitting the shape to image data to personalize geometry [18]; although the model is built from CT data, other modalities can be used as well with the help of simulated image appearance models

[93], [94]. Alternatively, the generative nature of the model can be exploited to generate virtual populations [86], [95], [96] to study the influence of certain parameters—including geometry itself—on simulation results.

Further personalization of the mesh in preparation for simulation can be based on invasive measurements, by taking into account statistical relationships between geometry and fiber orientation [97], or draw upon the literature to introduce normal, abnormal, or expected parameter values.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their constructive comments, and Mr. Milton Hoz and Ms. Silvina Ré for their help with the atlas website and the computing cluster.

REFERENCES

- [1] Virtual Physiological Human Initiative. VPH Network of Excellence. [Online]. Available: www.vph-noe.eu
- [2] International Union of Physiological Sciences. IUPS Physiome Project. [Online]. Available: www.physiome.org.nz
- [3] A. A. Young and A. F. Frangi, “Computational cardiac atlases: From patient to population and back,” *Exp. Physiol.*, vol. 94, no. 5, pp. 578–596, May 2009.
- [4] O. Camara, M. Pop, K. Rhode, M. Sermesant, N. Smith, and A. Young, Eds., *Proceedings of the First International Workshop on Statistical Atlases and Computational Models of the Heart (STACOM)*, ser. Lecture Notes in Computer Science. Springer, 2010, vol. 6364.
- [5] O. Camara, E. Konukoglu, M. Pop, K. Rhode, M. Sermesant, and A. Young, Eds., *Proceedings of the Second International Workshop on Statistical Atlases and Computational Models of the Heart (STACOM): Imaging and Modelling Challenges*, ser. Lecture Notes in Computer Science. Springer, 2011, vol. 7085.
- [6] M. Lorenzo-Valdés, G. I. Sanchez-Ortiz, A. G. Elkington, R. Mohiaddin, and D. Rueckert, “Segmentation of 4D cardiac MR images using a probabilistic atlas and the EM algorithm,” *Med. Image Anal.*, vol. 8, no. 3, pp. 255–265, Sep. 2004.
- [7] I. Išgum, M. Staring, A. Rutten, M. Prokop, M. A. 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.
- [8] E. M. van Rikxoort, I. Išgum, Y. Arzhaeva, M. Staring, S. Klein, M. A. Viergever, J. P. W. Pluim, and B. van Ginneken, “Adaptive local multi-atlas segmentation: Application to the heart and the caudate nucleus,” *Med. Image Anal.*, vol. 14, no. 1, pp. 39–49, Feb. 2010.
- [9] H. Kirşli, M. Schaap, S. Klein, T. van Walsum, S. L. Papadopoulos, A. C. Weustink, N. R. Mollet, M. Bonardi, C. H. Chen, E. J. Vonken, R. J. van der Geest, and W. J. Niessen, “Evaluation of a multi-atlas based method for segmentation of cardiac CTA data: A large-scale, multicenter and multivendor study,” *Med. Phys.*, vol. 37, no. 12, pp. 6279–6291, Dec. 2010.
- [10] J.-M. Peyrat, M. Sermesant, X. Pennec, H. Delingette, C. Xu, E. R. McVeigh, and N. Ayache, “A computational framework for the statistical analysis of cardiac diffusion tensors: Application to a small database of canine hearts,” *IEEE Trans. Med. Imag.*, vol. 26, no. 11, pp. 1500–1514, Nov. 2007.
- [11] H. Lombaert, J.-M. Peyrat, P. Croisille, S. Rapacchi, L. Fanton, F. Cheriet, P. Clarysse, I. Magnin, H. Delingette, and N. Ayache, “Human atlas of the cardiac fiber architecture: Study on a healthy population,” *IEEE Trans. Med. Imag.*, vol. 31, no. 7, pp. 1436–1447, Jul. 2012.
- [12] C. G. Fonseca, M. Backhaus, D. A. Bluemke, R. D. Britten, J. D. Chung, B. R. Cowan, I. D. Dinov, J. P. Finn, P. J. Hunter, A. H. Kadish, D. C. Lee, J. A. C. Lima, P. Medrano-Gracia, K. Shivkumar, A. Suinesiaputra, W. Tao, and A. A. Young, “The cardiac atlas project—an imaging database for computational modeling and statistical atlases of the heart,” *Bioinformatics*, vol. 27, no. 16, pp. 2288–2295, Aug. 2011.
- [13] P. Medrano-Gracia, B. R. Cowan, J. P. Finn, C. G. Fonseca, A. H. Kadish, D. C. Lee, W. Tao, and A. A. Young, “The Cardiac Atlas Project: Preliminary description of heart shape in patients with myocardial infarction,” in *Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM)*, 2010, pp. 46–53.

- [14] A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen, "Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling," *IEEE Trans. Med. Imag.*, vol. 21, no. 9, pp. 1151–1166, Sep. 2002.
- [15] 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, Sep. 2004.
- [16] D. Perperidis, R. Mohiaddin, and D. Rueckert, "Construction of a 4D statistical atlas of the cardiac anatomy and its use in classification," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 3750, 2005, pp. 402–410.
- [17] C. Hoogendoorn, F. M. Sukno, S. Ordás, and A. F. Frangi, "Bilinear models for spatio-temporal point distribution analysis: Application to extrapolation of left ventricular, biventricular and whole heart cardiac dynamics," *Int. J. Comput. Vis.*, vol. 85, no. 3, pp. 237–252, Dec. 2009.
- [18] 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.
- [19] M. R. Kaus, V. Pekar, C. Lorenz, R. Truyen, S. Lobregt, and J. Weese, "Automated 3-D PDM construction from segmented images using deformable models," *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 1005–1013, Aug. 2003.
- [20] C. Lorenz and J. von Berg, "A comprehensive shape model of the heart," *Med. Image Anal.*, vol. 10, no. 4, pp. 657–670, Aug. 2006.
- [21] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuering, and D. Comaniciu, "Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features," *IEEE Trans. Med. Imag.*, vol. 27, no. 11, pp. 1668–1681, Nov. 2008.
- [22] T. Heimann and H.-P. Meinzer, "Statistical shape models for 3D medical image segmentation: A review," *Med. Image Anal.*, vol. 13, no. 4, pp. 543–563, Aug. 2009.
- [23] J. Koikkalainen, T. Tölli, K. Lauerma, K. Antila, E. Mattila, M. Lilja, and J. Lötjönen, "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.
- [24] S. Ordás, E. Oubel, R. Leta, F. Carrera, and A. F. Frangi, "A statistical shape model of the heart and its application to model-based segmentation," in *Proc. SPIE Medical Imaging*, vol. 6511, 2007, article no. 65111K.
- [25] A. Guimond, J. Meunier, and J.-P. Thirion, "Average brain models: A convergence study," *Comput. Vis. Image Understand.*, vol. 77, no. 2, pp. 192–210, Feb. 2000.
- [26] G. E. Christensen, H. J. Johnson, and M. W. Vannier, "Synthesizing average 3D anatomical shapes," *NeuroImage*, vol. 32, no. 1, pp. 146–158, Aug. 2006.
- [27] H. Park, P. H. Bland, and C. R. Meyer, "Construction of an abdominal probabilistic atlas and its application in segmentation," *IEEE Trans. Med. Imag.*, vol. 22, no. 4, pp. 483–492, Apr. 2003.
- [28] H. Park, P. H. Bland, A. O. Hero III, and C. R. Meyer, "Least biased target selection in probabilistic atlas construction," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 3750, 2005, pp. 419–426.
- [29] C. Hoogendoorn, T. Whitmarsh, N. Duchateau, F. M. Sukno, M. De Craene, and A. F. Frangi, "A groupwise mutual information metric for cost efficient selection of a suitable reference in cardiac computational atlas construction," in *Proc. SPIE Medical Imaging*, vol. 7623, 2010.
- [30] C. Jongen, J. P. W. Pluim, P. J. Nederkoorn, M. A. Viergever, and W. J. Niessen, "Construction and evaluation of an average CT brain image for inter-subject registration," *Comput. Biol. Med.*, vol. 34, no. 8, pp. 647–662, Dec. 2004.
- [31] H. Jia, G. Wu, Q. Wang, and D. Shen, "ABSORB: Atlas building by self-organized registration and bundling," *NeuroImage*, vol. 51, no. 3, pp. 1057–1070, Jul. 2010.
- [32] G. Wu, H. Jia, Q. Wang, and D. Shen, "SharpMean: Groupwise registration guided by sharp mean image and tree-based registration," *NeuroImage*, vol. 56, no. 4, pp. 1968–1981, Jun. 2011.
- [33] D. J. Blezek and J. V. Miller, "Atlas stratification," *Med. Image Anal.*, vol. 11, no. 5, pp. 443–457, Oct. 2007.
- [34] M. R. Sabuncu, S. K. Balci, M. E. Shenton, and P. Golland, "Image-driven population analysis through mixture modeling," *IEEE Trans. Med. Imag.*, vol. 28, no. 9, pp. 1473–1487, Sep. 2009.
- [35] S. Joshi, B. Davis, M. Jomier, and G. Gerig, "Unbiased diffeomorphic atlas construction for computational anatomy," *NeuroImage*, vol. 23, no. S1, pp. S151–S160, Sep. 2004.
- [36] P. Lorenzen, M. Prastawa, B. Davis, G. Gerig, E. Bullitt, and S. Joshi, "Multi-modal image set registration and atlas formation," *Med. Image Anal.*, vol. 10, no. 3, pp. 440–451, Jun. 2006.
- [37] P. T. Fletcher, S. Venkatasubramanian, and S. Joshi, "The geometric median on Riemannian manifolds with application to robust atlas estimation," *NeuroImage*, vol. 45, no. 1, pp. S143–S152, Mar. 2009.
- [38] K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert, "Consistent groupwise non-rigid registration for atlas construction," in *Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI)*, 2004, pp. 908–911.
- [39] S. Marsland, C. J. Twining, and C. J. Taylor, "A minimum description length objective function for groupwise non-rigid image registration," *Image Vis. Comput.*, vol. 26, no. 3, pp. 333–346, Mar. 2008.
- [40] C. Studholme and V. Cardenas, "A template free approach to volumetric spatial normalization of brain anatomy," *Pattern Recogn. Lett.*, vol. 25, no. 10, pp. 1191–1202, Jul. 2004.
- [41] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, no. 8, pp. 712–721, Aug. 1999.
- [42] D. Rueckert, P. Aljabar, R. A. Heckemann, J. V. Hajnal, and A. Hammers, "Diffeomorphic registration using B-splines," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 4191, 2006, pp. 702–709.
- [43] S. Marsland, C. J. Twining, and C. J. Taylor, "Groupwise non-rigid registration using polyharmonic clamped-plate splines," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 2879, 2003, pp. 771–779.
- [44] M. De Craene, A. du Bois d'Aische, B. Macq, and S. K. Warfield, "Multi-subject registration for unbiased statistical atlas construction," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 3216, 2004, pp. 655–662.
- [45] F. Wang, B. C. Vemuri, A. Rangarajan, and S. J. Eisenschenk, "Simultaneous nonrigid registration of multiple point sets and atlas construction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 11, pp. 2011–2022, Nov. 2008.
- [46] J. Hamm, D. H. Ye, R. Verma, and C. Davatzikos, "GRAM: A framework for geodesic registration on anatomical manifolds," *Med. Image Anal.*, vol. 14, no. 5, pp. 633–642, Oct. 2010.
- [47] H. Jia, P.-T. Yap, G. Wu, Q. Wang, and D. Shen, "Intermediate templates guided groupwise registration of diffusion tensor images," *NeuroImage*, vol. 54, no. 2, pp. 928–939, Jan. 2011.
- [48] T. Mäkelä, P. Clarysse, O. Sipilä, N. Pauna, Q. C. Pham, T. Katila, and I. E. Magnin, "A review of cardiac image registration methods," *IEEE Trans. Med. Imag.*, vol. 21, no. 9, pp. 1011–1021, Sep. 2002.
- [49] M. Holden, "A review of geometric transformations for nonrigid body registration," *IEEE Trans. Med. Imag.*, vol. 27, no. 1, pp. 111–128, Jan. 2008.
- [50] S. Klein, M. Staring, and J. P. W. Pluim, "Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 2879–2890, Dec. 2007.
- [51] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Mutual-information based registration of medical images: A survey," *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 986–1004, Aug. 2003.
- [52] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, no. 1, pp. 1–36, Mar. 1998.
- [53] H. Tandri, S. K. Daya, K. Nasir, C. Bomma, J. A. C. Lima, H. Calkins, and D. A. Bluemke, "Normal reference values for the adult right ventricle by magnetic resonance imaging," *Am. J. Cardiol.*, vol. 98, no. 12, pp. 1660–1664, Dec. 2006.
- [54] A. H. Kadish, D. Bello, J. P. Finn, R. O. Bonow, A. Schaechter, H. Subacius, C. Albert, J. P. Daubert, C. G. Fonseca, and J. J. Goldberger, "Rationale and design for the defibrillators to reduce risk by magnetic resonance imaging evaluation (DETERMINE) trial," *J. Cardiovasc. Electrophysiol.*, vol. 20, no. 9, pp. 982–987, Sep. 2009.
- [55] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, no. 2, pp. 187–198, Apr. 1997.
- [56] D. Rueckert, A. F. Frangi, and J. A. Schnabel, "Automatic construction of 3-D statistical deformation models of the brain using nonrigid registration," *IEEE Trans. Med. Imag.*, vol. 22, no. 8, pp. 1014–1025, Aug. 2003.
- [57] T. F. Cootes, C. J. Twining, K. O. Babalola, and C. J. Taylor, "Diffeomorphic statistical shape models," *Image Vis. Comput.*, vol. 26, no. 3, pp. 326–332, Mar. 2008.

- [58] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT image registration in the chest using free-form deformations," *IEEE Trans. Med. Imag.*, vol. 22, no. 1, pp. 120–128, Jan. 2003.
- [59] Y. Choi and S. Lee, "Injectivity conditions of 2D and 3D uniform cubic B-spline functions," *Graph. Model.*, vol. 62, no. 6, pp. 411–427, Nov. 2000.
- [60] A. Trounev, "Diffeomorphisms groups and pattern matching in image analysis," *Int. J. Comput. Vis.*, vol. 28, no. 3, pp. 213–221, Jul. 1998.
- [61] Kitware inc. The Insight Segmentation and Registration Toolkit. [Online]. Available: www.itk.org
- [62] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM J. Sci. Comput.*, vol. 16, no. 5, pp. 1190–1208, 1995.
- [63] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Diffeomorphic demons: Efficient non-parametric image registration," *NeuroImage*, vol. 45, no. 1, pp. S61–S72, Mar. 2009.
- [64] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache, "A log-euclidean framework for statistics on diffeomorphisms," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 4190, 2006, pp. 924–931.
- [65] I. Larrabide, P. Omedas, Y. Martelli, X. Planes, M. Nieber, J. A. Moya, C. Butakoff, R. Sebastián, O. Camara, M. De Craene, B. H. Bijmens, and A. F. Frangi, "GIMIAs: An open source framework for efficient development of research tools and clinical prototypes," in *Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH)*, ser. Lecture Notes in Computer Science, vol. 5528, 2009, pp. 417–426.
- [66] H. Hege, M. Seebaß, D. Stalling, and M. Zöckler, "A generalized marching cubes algorithm based on non-binary classifications," Konrad-Zuse-Zentrum für Informationstechnik, Berlin, Germany, Tech. Rep. ZIB-Report SC 97-05, 1997.
- [67] Z. Wu and J. M. Sullivan jr., "Multiple material marching cubes algorithm," *Int. J. Numer. Meth. Eng.*, vol. 58, no. 2, pp. 189–207, Sep. 2003.
- [68] S. E. Dillard, J. F. Bingert, D. Thoma, and B. Hamann, "Construction of simplified boundary surfaces from serial-sectioned metal micrographs," *IEEE Trans. Vis. Comput. Graphics*, vol. 13, no. 6, pp. 1528–1535, Nov./Dec. 2007.
- [69] M. Meyer, R. Whitaker, R. M. Kirby, C. Ledergerber, and H. Pfister, "Particle-based sampling and meshing of surfaces in multimaterial volumes," *IEEE Trans. Vis. Comput. Graphics*, vol. 14, no. 6, pp. 1539–1546, Nov./Dec. 2008.
- [70] M. Bertram, G. Reis, R. van Lengen, S. Köhn, and H. Hagen, "Non-manifold mesh extraction from time-varying segmented volumes used for modeling a human heart," in *Proc. Joint EuroGraphics/IEEE VGTC Symp. on Visualization (EuroVis)*, 2005, pp. 199–206.
- [71] Y. Zhang, T. J. R. Hughes, and C. L. Bajaj, "An automatic 3D mesh generation method for domains with multiple materials," *Comput. Meth. Appl. Mech. Eng.*, vol. 199, no. 5–8, pp. 405–415, Jan. 2010.
- [72] V. d'Otreppe, R. Boman, and J.-P. Ponthot, "Generating smooth surface meshes from multi-region medical images," *Int. J. Numer. Meth. Biomed. Eng.*, vol. 28, no. 6/7, pp. 642–660, Jun/Jul. 2012.
- [73] S. P. Raya and J. K. Udupa, "Shape-based interpolation of multidimensional objects," *IEEE Trans. Med. Imag.*, vol. 9, no. 1, pp. 32–42, Jan. 1990.
- [74] G. T. Herman, J. Zheng, and C. A. Bucholtz, "Shape-based interpolation," *IEEE Comput. Graph. Appl.*, vol. 12, no. 3, pp. 69–79, May 1992.
- [75] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," in *ACM SIGGRAPH Int. Conf. on Computer Graphics and Interactive Techniques*, 1987, pp. 163–169.
- [76] S. Valette, J.-M. Chassery, and R. Prost, "Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams," *IEEE Trans. Vis. Comput. Graphics*, vol. 14, no. 2, pp. 369–381, Mar./Apr. 2008.
- [77] G. Taubin, T. Zhang, and G. Golub, "Optimal surface smoothing as filter design," in *European Conf. on Computer Vision (ECCV)*, ser. Lecture Notes in Computer Science, vol. 1064, 1996, pp. 283–292.
- [78] R. Prakash, "Determination of right ventricular wall thickness in systole and diastole," *Br. Heart J.*, vol. 40, no. 11, pp. 1257–1261, Nov. 1978.
- [79] R. Erbel and H. Eggebrecht, "Aortic dimensions and the risk of dissection," *Heart*, vol. 92, no. 1, pp. 137–142, Jan. 2006.
- [80] B. Hall, V. Jeevanantham, R. Simon, J. Filippone, G. Vorobiof, and J. Daubert, "Variation in left atrial transmural wall thickness at sites commonly targeted for ablation of atrial fibrillation," *J. Interv. Card. Electrophysiol.*, vol. 17, no. 2, pp. 127–132, Nov. 2006.
- [81] J. B. Tenenbaum and W. T. Freeman, "Separating style and content with bilinear models," *Neural Comput.*, vol. 12, no. 6, pp. 1247–1283, Jun. 2000.
- [82] I. T. Jolliffe, *Principal Component Analysis*, 2nd ed. Springer-Verlag, New York, 2002.
- [83] D. Perperidis, "Spatio-temporal registration and modelling of the heart using cardiovascular MR imaging," Ph.D. dissertation, Imperial Coll. London, London, UK, 2005.
- [84] R. M. Figueras i Ventura, C. Hoogendoorn, F. M. Sukno, and A. F. Frangi, "Bilinear point distribution models for heart motion analysis," in *Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI)*, 2010, pp. 476–479.
- [85] A. R. Porras, G. Piella, A. Berrueto, C. Hoogendoorn, D. Andreu, M. Sitges, J. Fernandez-Armenta, and A. F. Frangi, "Interventional endocardial motion estimation from electro-anatomical mapping data: Application to scar characterization," *IEEE Trans. Biomed. Eng.*, accepted for publication.
- [86] C. Hoogendoorn, A. Pashaei, R. Sebastián, F. M. Sukno, O. Cámara, and A. F. Frangi, "Influence of geometric variations on LV activation times: A study on an atlas-based virtual population," in *Proc. Workshop Statistical Atlases and Computational Models of the Heart (STACOM)*, ser. Lecture Notes in Computer Science, vol. 6364, 2010, pp. 242–251.
- [87] V. Y. Wang, C. Hoogendoorn, G. Engelbrecht, A. F. Frangi, A. A. Young, P. J. Hunter, and M. P. Nash, "Unsupervised segmentation and personalised FE modelling of in vivo human myocardial mechanics based on an MRI atlas," in *Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI)*, 2012, pp. 1360–1363.
- [88] M. Wierzbicki, M. Drangova, G. Guiraudon, and T. Peters, "Validation of dynamic heart models obtained using non-linear registration for virtual reality training, planning and guidance of minimally invasive cardiac surgeries," *Med. Image Anal.*, vol. 8, no. 3, pp. 387–401, Sep. 2004.
- [89] E. M. Marom, J. E. Herndon, Y. H. Kim, and H. P. McAdams, "Variations in pulmonary venous drainage to the left atrium: Implications for radiofrequency ablation," *Radiology*, vol. 230, pp. 824–829, Mar. 2004.
- [90] E. L. Alderman and M. Stadius, "The angiographic definitions of the Bypass Angioplasty Revascularization Investigation," *Coron. Artery Dis.*, vol. 3, no. 12, pp. 1189–1207, Dec. 1992.
- [91] C. T. Metz, S. Klein, M. Schaap, T. van Walsum, and W. J. Niessen, "Nonrigid registration of dynamic medical imaging data using $nd + t$ B-splines and a groupwise optimization approach," *Med. Image Anal.*, vol. 15, no. 2, pp. 238–249, Apr. 2011.
- [92] M. De Craene, G. Piella, O. Camara, N. Duchateau, E. Silva, A. Doltra, J. D'Hooge, J. Brugada, M. Sitges, and A. F. Frangi, "Temporal diffeomorphic free-form deformation: Application to motion and strain estimation from 3D echocardiography," *Med. Image Anal.*, vol. 16, no. 2, pp. 427–450, Feb. 2012.
- [93] C. Tobon-Gomez, C. Butakoff, S. Aguade, F. M. Sukno, G. Moragas, and A. F. Frangi, "Automatic construction of 3D-ASM intensity models by simulating image acquisition: Application to myocardial gated SPECT studies," *IEEE Trans. Med. Imag.*, vol. 27, no. 11, pp. 1655–1667, Nov. 2008.
- [94] C. Tobon-Gomez, F. M. Sukno, B. H. Bijmens, M. Huguet, and A. F. Frangi, "Realistic simulation of cardiac magnetic resonance studies modeling anatomical variability, trabeculae, and papillary muscles," *Magn. Reson. Med.*, vol. 65, no. 1, pp. 280–288, Jan. 2011.
- [95] C. Hoogendoorn, A. Pashaei, R. Sebastián, F. M. Sukno, O. Cámara, and A. F. Frangi, "Sensitivity analysis of mesh warping and subsampling strategies for generating large scale electrophysiological simulation data," in *Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH)*, ser. Lecture Notes in Computer Science, vol. 6666, 2011, pp. 418–426.
- [96] A. Pashaei, C. Hoogendoorn, R. Sebastián, D. Romero, O. Cámara, and A. F. Frangi, "Infarction and fast electrophysiology of human heart: In-silico study on atlas-based virtual population," in *Int. Conf. on Functional Imaging and Modeling of the Heart (FIMH)*, ser. Lecture Notes in Computer Science, vol. 6666, 2011, pp. 427–436.
- [97] K. Lekadir, B. Ghafaryasl, E. Muñoz-Moreno, C. Butakoff, C. Hoogendoorn, and A. F. Frangi, "Predictive modeling of cardiac fiber orientation using the Knutsson mapping," in *Int. Conf. Medical Image Computing and Computer Assisted Intervention (MICCAI)*, ser. Lecture Notes in Computer Science, vol. 6892, 2011, pp. 50–57.