

Current-based 4D shape analysis for the mechanical personalization of heart models

Loïc Le Folgoc¹, Hervé Delingette¹, Antonio Criminisi², and Nicholas Ayache¹

¹ Asclepios Research Project, INRIA Sophia Antipolis, France.

² Machine Learning and Perception Group, Microsoft Research Cambridge, UK.

Abstract. Patient-specific models of the heart may lead to better understanding of cardiovascular diseases and better planning of therapy. A machine-learning approach to the personalization of an electro-mechanical model of the heart, from the kinematics of the endo- and epicardium, is presented in this paper. We use 4D mathematical currents to encapsulate information about the shape and deformation of the heart. The method is largely insensitive to initialization and does not require on-line simulation of the cardiac function. In this work, we demonstrate the performance of our approach for the joint estimation of three parameters on one heart geometry. We manage to retrieve parameters such that the model matches the 4D observations with an accuracy below the voxel size, in less than three minutes of computation.

Keywords: patient-specific heart model, mechanical personalization, currents, machine-learning

1 Introduction

Patient-specific models may help better understand the role of biomechanical and electrophysiological factors in cardiovascular pathologies. They may also prove to be useful in predicting the outcome of potential therapeutic interventions for individual patients. In this paper we focus on the mechanical personalization of the Bestel-Clement-Sorine (BCS) model, as described in [2][4].

Model personalization aims at optimizing model parameters so that the behaviour of the personalized model matches the acquired patient-specific data (e.g. cine-MR images). Several approaches to the problem of cardiac model personalization have been suggested in the recent years, often formulating the inverse problem via the framework of variational data assimilation[6] or that of optimal filtering theory[14][13][3]. The output of these methods is dependent on the set of parameters used to initialize the algorithm; for this reason calibration procedures are introduced as a preprocessing stage, such as the one developed in [16]. Furthermore these approaches rely on on-line simulations, as an accurate estimation of the effect of parameter changes along several directions in the parameter space is required to drive the parameter estimation. Due to the complexity of the direct simulation these approaches are costly in time and computations.

In this paper, we explore a novel machine-learning approach, in which the need for initialization and on-line simulation is removed, by moving the analysis of the parameter effects on the kinematics of the model (and thus the bulk of the computations) to an off-line learning phase. In this work we assume the tracking of the heart motion from images to be given (e.g. via [15]) and focus on the mechanical personalization of the cardiac function from meshes. Our work makes use of currents, a mathematical tool which was originally introduced to the medical imaging community in the context of shape registration[18][8] and offers a unified, correspondence-free statistical representation of geometrical objects. Our main contributions include the construction of 4D currents to represent, and perform statistics on $3D + t$ beating hearts and the proposal of a machine-learning framework to personalize electromechanical cardiac models.

The remaining of this article is organized as follows. In the first part we introduce the background on currents necessary to present the rest of our work. We develop our method in the following section, then present and discuss experimental results in the final sections.

2 Currents for shape representation

2.1 A statistical shape representation framework

Currents provide a unified representation of geometrical objects of any dimension, embedded in the Euclidean space \mathbb{R}^n , that is fit for statistical analysis. The framework of currents makes use of geometrically rich and well-behaved data spaces allowing for the proper definition of classical statistical concepts. Typically the existence of an inner product structure provides a straightforward way to define the mean and principal modes of a data set for instance, as in the Principal Component Analysis (PCA). These comments motivate an approach of currents from the perspective of kernel theory in this section, although currents are formally introduced in a more general way *via* the field of differential topology. The connection to differential topology is particularly relevant to outline the desirable properties of currents when dealing with discrete approximations of continuous shapes, in terms of convergence and consistence of the representation [7].

A well-known theorem due to Moore and Aronszajn[1] states that for any symmetric, positive definite (p.d.) kernel on a set \mathcal{X} , there exists a unique Hilbert space $\mathcal{H}_K \subset \mathbb{R}^{\mathcal{X}}$ for which K is a reproducing kernel. This result suggests a straightforward way of doing statistics on \mathcal{X} as long as a p.d. kernel K can be engineered on this set, by mapping any point $x \in \mathcal{X}$ to a function $K(x, \cdot) \in \mathcal{H}_K$ and exploiting the Hilbert space structure in \mathcal{H}_K . Furthermore, practical computations can be efficiently tractated thanks to the *reproducing kernel* property - namely, for any $x, y \in \mathcal{X}$, we have

$$(K(x, \cdot) | K(y, \cdot))_{\mathcal{H}_K} = K(x, y), \quad (1)$$

and more generally yet, for any $f \in \mathcal{H}_K$, $(K(x, \cdot) | f)_{\mathcal{H}_K} = f(x)$. Expanding on this, one can compute statistics on pairs of points and m -vectors $(x, \eta) \in \mathbb{R}^n \times \Lambda_m \mathbb{R}^n$ by mapping them to functions $K(x, \cdot)\eta$ and making use of the reproducing property

$$(K(x, \cdot)\eta | K(y, \cdot)\nu) = \eta^\top \nu K(x, y). \quad (2)$$

Eq. 2 simply extends Eq. 1 to vector-valued functions, making use of the fact that the tensor product of two kernels is again a kernel over the product space. Expanding the framework even further, we can regard a discrete shape as a finite set $\{(x_i, \eta_i)\}_{1 \leq i \leq p}$, where η_i describes the tangent space at x_i , and associate to it a signature function $\sum_{1 \leq i \leq p} K(x_i, \cdot) \eta_i$. The correlation between two discrete shapes $\{(x_i, \eta_i)\}_{1 \leq i \leq p}$ and $\{(y_j, \nu_j)\}_{1 \leq j \leq q}$ can then be measured by the inner product

$$\left(\sum_i K(x_i, \cdot) \eta_i \mid \sum_j K(y_j, \cdot) \nu_j \right) = \sum_{i,j} \eta_i^\top \nu_j K(x_i, y_j). \quad (3)$$

This construction may in fact be acknowledged as a special case of the convolution kernel on discrete structures described in [11] and [10]. The above defines a correspondence-free way to measure proximity between shapes, trading hard correspondences for an aggregation of the measures of proximity between each simplex of one shape with every simplex of the other shape in the sense of a kernel $K(\cdot, \cdot)$. We have yet to specify a choice of kernel K . In the following, we will consider the multivariate Gaussian kernel with variance Σ :

$$K_\Sigma(x, y) = \frac{1}{\{(2\pi)^n |\Sigma|\}^{1/2}} \exp -\frac{1}{2} (x - y)^\top \Sigma^{-1} (x - y).$$

The choice of kernel width Σ can be interpreted as a choice of scale at which the shape of interest is observed: shape variations occurring at a lower scale are likely to be smoothed by the convolution and go unnoticed. This mechanism naturally introduces some level of noise insensitivity in the analysis. This parameter should thus be decided with regard to the mesh resolution and the level of noise in the data.

Finally, the linear pointwise-evaluation functional $\delta_x^\eta: \omega \mapsto \omega(x)(\eta)$ is continuous and dual to $K(x, \cdot)\eta$ by the reproducing kernel property. In the following we will refer to δ_x^η as a *delta-current* or a *moment*. To summarize, the discretized m -manifold $\{(x_i, \eta_i)\}_{1 \leq i \leq p}$ admits equivalent representations as the current $\sum_i \delta_{x_i}^{\eta_i}$, its dual differential m -form $\sum_{1 \leq i \leq p} K(x_i, \cdot) \eta_i^\top$ or its dual vector field $\sum_{1 \leq i \leq p} K(x_i, \cdot) \eta_i$.

2.2 Computational efficiency and compact approximate representations

This framework lends itself to an efficient implementation. Firstly, the inner product between two discrete shapes can be computed in linear time with respect to the number of momenta through the use of a translation invariant kernel. Indeed $\gamma(\cdot) = \sum_i K(x_i, \cdot) \eta_i$ may then be precomputed at any desired accuracy on a discrete grid by convolution, and rewriting $\sum_{i,j} \eta_i^\top \nu_j K(x_i, y_j)$ as $\sum_j \gamma(y_j)^\top \nu_j$ demonstrates the claim.

Secondly, if the mesh diameter is small with respect to the scale Σ , the initial delta-current representation will be highly redundant. Durrleman et al.[9] introduced an iterative method to obtain compact approximations of currents at a chosen scale and with any desired accuracy. We rely on this procedure at training time to fasten computations and reduce the memory load. This algorithm is inspired from the Matching Pursuit method[5]. A compact current is built from the current S to approximate (of

dual field γ) by iteratively adding a single delta current $\delta_{x_n}^{\eta_n}$ to the previous approximation S_{n-1} , in such a way that the difference $\|S - S_n\|_{\mathcal{H}'_\Sigma}$ steadily decreases. This is achieved by greedily placing the moment at the maximum (in $\|\cdot\|_2$ norm) x_n of the residual field $\gamma(\cdot) - \gamma_{n-1}(\cdot)$, then choosing the optimal η , i.e. the one that minimizes $\|\gamma - \{\gamma_{n-1} + K(x_n, \cdot)\eta\}\|_{\mathcal{H}_\Sigma}^2$. It is shown in [9] that this algorithm is greedy in $\|\cdot\|_{\mathcal{H}_\Sigma}$ norm, and converges both in $\|\cdot\|_{\mathcal{H}_\Sigma}$ norm and $\|\cdot\|_\infty$ norm. The stopping criterion is on the residual norm $\|\gamma(\cdot) - \gamma_n(\cdot)\|_{\mathcal{H}_\Sigma}^2$. Our implementation uses a discrete kernel approximation of the Gaussian kernel, rather than an FFT based scheme, for fast local updates of the residual field.

3 Method

The workflow for the proposed machine-learning based parameter estimation method couples three successive processing steps: the first one aims at generating a current from an input sequence of meshes, so as to obtain a statistically relevant representation; the second one consists in a dimensionality reduction step, so as to derive a reduced shape representation in \mathbb{R}^k , which leads to computationally efficient statistical learning; the third step tackles the matter of finding a relationship between the reduced shape space and the (biophysical) model parameters. The three modules are mostly independent and can easily be adjusted in their own respect. As a machine learning based method, our work involves an off-line learning stage and an on-line testing stage: all three modules of the pipeline are involved during each stage. Fig. 1 gives a visual overview of our approach. The rest of this section describes the three afore-mentioned processing steps and their use during learning and testing stages.

3.1 Current generation from mesh sequences

Let us briefly describe the way we build a current from a time sequence of 3D meshes. We first extract surface meshes from the volumetric meshes. This choice derives from the assumption that the displacement of surface points can be recovered more easily than the displacement of all points within the myocardium, given a sequence of images; thus learning from surface meshes may be more relevant for real applications. In this work we assume the trajectory of surface points to be entirely known, as opposed to the displacement in the direction normal to the contour only (aperture problem). Several variants to derive currents for 4D object representation can be discussed (e.g. [7]), but their relevance largely depends on the application and complete processing work flow from the original data.

In this work, we rely on the remark that the concatenation of smoothly deformed surface meshes can be visualized as a (3D) hyper-surface in 4D (Fig. 2). The i th simplex of this hyper-surface generates a current $\delta_{x_i}^{\eta_i}$, where x_i is its barycenter and η_i is the vector of \mathbb{R}^4 normal to its support and of length the volume of the simplex. The current associated to the series of meshes is the aggregation of such delta currents, $\sum_i \delta_{x_i}^{\eta_i}$. This construction captures both the geometry of the heart and its motion.

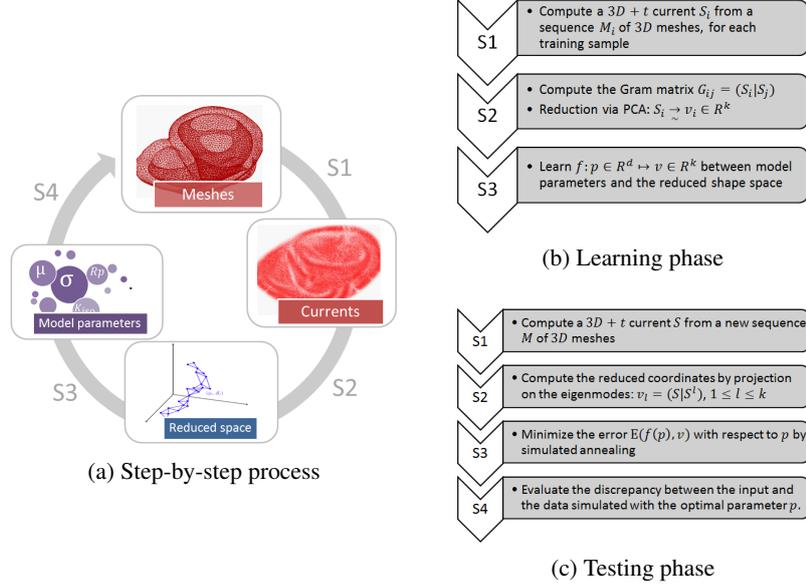


Fig. 1: Overview of the learning and testing phases.

3.2 Shape space reduction

Since learning a direct mapping between the space of model parameters and the space of $3D+t$ currents is a cumbersome task, we introduce an intermediate step of dimensionality reduction via PCA. During the learning stage, we compute the mean current and principal modes of variation from the learning database of N currents $\{S_i\}_{1 \leq i \leq N}$ generated from the N training mesh sequences $\{M_i\}_{1 \leq i \leq N}$ as described in §3.1. This is achieved efficiently by computing the Gram matrix of the data $\mathbf{G}_{ij} = (S_i | S_j)$ column by column and using the "kernel trick"[17]. Each column of \mathbf{G} is computed in $\mathcal{O}(N \cdot P)$, where P is the maximum number of momenta among all currents S_j (cf. §2.1). Finally, we compute an approximate compact representation at the scale Σ of the mean current \bar{T} and of the K first modes of variation $\{T_k\}_{1 \leq k \leq K}$ to accelerate computations of inner products involving these currents[9].

At testing time and given a new current S , we derive its coordinates $v = (v_1, \dots, v_K)$ in the reduced shape space by projection on the principal modes of variation, $v_k = (S - \bar{T} | T_k)$.

3.3 Regression problem for model parameter learning

It remains to link the physiological (model) parameters to the reduced shape space. Although we are ultimately interested in finding an optimal set of parameters $p \in \mathbb{R}^d$ from an observation $v \in \mathbb{R}^K$ we will actually learn a mapping in the other direction, $f: p \in \mathbb{R}^d \mapsto v \in \mathbb{R}^K$. We motivate this choice by three arguments. Firstly, the observation v is a deterministic output of the cardiac model given a parameter set p and thus

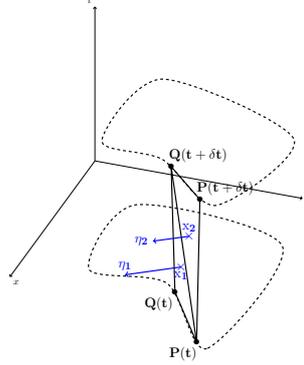


Fig. 2: Current generation from a mesh element, illustrated on an element of contour in 2D deformed in time. The simplex PQ is followed over two consecutive timesteps, which gives a quad embedded in 3D. The quad is divided into two triangles, from which we get two current deltas, applied at each triangle barycenter, orthogonal to the support of their corresponding triangles and of norm the area of the triangle. For a surface in 3D deformed over time, each element of the triangulation followed over two consecutive timesteps generates a hyper-prism embedded in 4D, which is in turn decomposed in three tetrahedra from which we obtain three momenta.

the mapping f is well-defined; however there may be several parameter sets resulting in the same observable shape and deformation, as parameter identifiability is not *a priori* ensured. Secondly, the parameter space is expected to be of smaller dimensionality than the reduced shape space and therefore easier to sample for combinatorial reasons. Finally, we can also expect that the set of biologically admissible model parameters be relatively well-behaved; on the other hand few points in the shape space may actually relate to anatomically reasonable hearts: thus mapping every $v \in \mathbb{R}^k$ to a parameter set could be impractical.

The regression function f is learned by kernel ridge regression using a Gaussian kernel[12], and admits a straightforward close-form expression. During the testing phase, given a new observation v , we solve the optimization problem $\arg \min_p \|f(p) - v\|^2$ by Simulated Annealing[19]. This optimization problem involves an analytical mapping between low-dimensional spaces, as opposed to optimizing directly over the 4D meshes or currents. Thus it will not constitute a computational bottleneck regardless of the chosen optimization scheme. Naturally, if a prior on the likelihood of a given parameter set $p \in \mathbb{R}^d$ were known (e.g. via a biophysical argument), it could be integrated in the cost function in the form of a prior energy term $\lambda \cdot R(p)$.

4 Experimental results

In our first experiment we focus on the prediction of the maximum contractility parameter σ_0 of the BCS model, defined globally for the whole cardiac muscle. Building on

the sensitivity analysis from [16], we consider that σ_0 covers the range of values from 10^6 to $2 \cdot 10^7$ in an anatomically plausible way. We form a training base of ten cases $\{p_i, \mathcal{M}_i\}$ by sampling this range deterministically and launching simulations with the corresponding parameter sets, for a single heart geometry from the STACOM'2011 dataset. Following the PCA, the first principal mode of variation is found to explain 81% of the variance, thus we set the reduced shape space to be of dimension 1 ($K = 1$); the regression function ($\sigma = 0.3$, $\lambda = 10^{-5}$) bijectively maps the model parameter space and the reduced shape space. In all experiments, the model parameters are affinely mapped to $[-1, 1]$ for convenience, for the regression and optimization stages. We use an isotropic Gaussian kernel of width 1cm in space and 50ms in time.

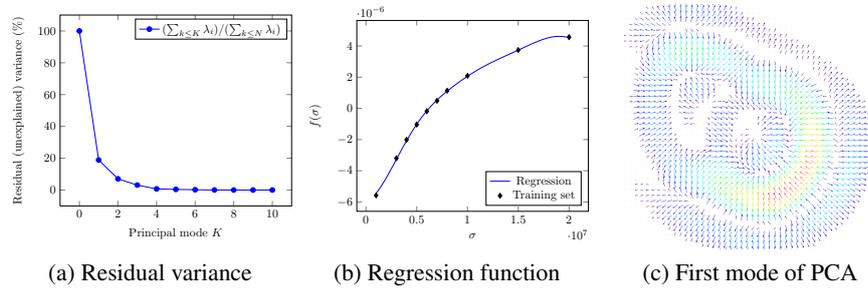


Fig. 3: PCA results for the first experiment. The projection of the first mode of variation on a plane orthogonal to the z -axis at a fixed time step is shown in (c), and can be interpreted as capturing the variability in the extent of the contraction of the muscle.

In the spirit of cross-validation procedures, we evaluate the performance of our approach on an independent test set $\{p_j, \mathcal{M}_j\}_{0 \leq j < N'}$ by randomly choosing parameter sets in the admissible range of parameters and launching the corresponding simulations. We thereafter refer to p_j as the *real* parameter (value) and to the output of our approach p_j^* as the *optimal* parameter (value). Our test set is of size $N' = 100$ samples. The whole personalization pipeline, from the current generation to the parameter optimization phase, takes roughly 2 minutes per sample on a regular laptop. We define the relative error on the parameter value for a given test sample j as $\epsilon_r p_j = |p_j^* - p_j|/p_j$. In addition to the relative error, we consider the absolute error over the range of admissible parameters, $\epsilon_a p_j = |p_j^* - p_j|/|p_{\max} - p_{\min}|$. We refer to $\epsilon_a p$ as an absolute error but express it for convenience as a percentage of the admissible parameter variation. Over the test set, we found a mean relative (resp. absolute) error of 9.2% (resp. 4.5%) and a median relative (resp. absolute) error of 6.8% (resp. 2.3%).

We are also interested in a preliminary evaluation of the robustness of our approach with respect to geometry changes. Ten samples are generated following the same procedure as before, but using another heart geometry of the STACOM dataset. The 10 mesh sequences are manually registered (*via* a similarity transform) to the training geometry

based on the end-diastole mesh before applying the normal pipeline, as described in Section 3. The mean relative (resp. absolute) error on the contractility parameter over our sample is 25% (9.3%), with 15% (resp. 7.5%) median relative (absolute) error.

The second experiment proceeds similarly to the first one, but we simultaneously estimate the contractility σ_0 , the relaxation rate k_{rs} and the viscosity μ . For the training phase, the parameter space is sampled on a $7 \times 7 \times 7$ grid with σ_0 in the range $[10^6, 2 \cdot 10^7]$, k_{rs} in $[5, 50]$ and μ in $[10^5, 8 \cdot 10^5]$. The explained variance with 1 eigenmode of the PCA (resp. 2 to 5) out of the $N = 343$ modes equals 63.2% of the total variance (resp. 80.3%, 89.5%, 94.1%, 96.7%). We set the dimension of the reduced shape space to $K = 3$. The performance is tested on $N' = 100$ random samples. Because we can no longer assume the parameter set to be identifiable *a priori*, we introduce another measure of the goodness of fit of our personalization by directly evaluating the error on the observations. Given two surface mesh sequences $\mathcal{M} = \{\mathcal{M}_i\}_{1 \leq i \leq T}$ and $\mathcal{M}' = \{\mathcal{M}'_i\}_{1 \leq i \leq T}$, we define the pseudo-distance $d_{\text{sur}}(\mathcal{M}, \mathcal{M}') = \max_i d_s(\mathcal{M}_i, \mathcal{M}'_i)$ where $d_s(\mathcal{M}_i, \mathcal{M}'_i)^2$ is the mean square distance of the points of the surface \mathcal{M}_i to the surface \mathcal{M}'_i . Additionally given one-to-one correspondences between \mathcal{M} and \mathcal{M}' , we can define the distance $d_{\text{nod}}(\mathcal{M}, \mathcal{M}') = \max_i d_p(\mathcal{M}_i, \mathcal{M}'_i)$, where $d_p(\mathcal{M}_i, \mathcal{M}'_i)$ is the mean distance between corresponding nodes of \mathcal{M}_i and \mathcal{M}'_i . While d_{sur} intuitively relates to an upper bound for the matching between surface meshes at any time step, d_{nod} conveys more information about the quality of the matching of point trajectories. The results for this experiment are reported in Table 1. As a comparison, two mesh sequences corresponding to extreme values in the parameter set will yield a value for $d_{\text{sur}}(\mathcal{M}, \mathcal{M}')$ (resp. $d_{\text{nod}}(\mathcal{M}, \mathcal{M}')$) of the order of 6mm (resp. 8mm).

Table 1: Experiment 2 - results

	$\epsilon_r \sigma_0$ ($\epsilon_a \sigma_0$)	$\epsilon_r k_{rs}$ ($\epsilon_a k_{rs}$)	$\epsilon_r \mu$ ($\epsilon_a \mu$)	d_{sur} (mm)	d_{nod} (mm)
Mean	15.2% (8.0%)	48.8% (26.4%)	40.5% (20.0%)	0.92mm	1.42mm
Median	13.2% (6.3%)	44.7% (19.6%)	32.1% (17.5%)	0.80mm	1.32mm

In addition we compute the optimal parameters and performance indicators for a different choice of the reduced space dimension K , obtaining quasi-identical statistics for $K = 4$. Finally, we test here again the robustness with respect to changes of the heart geometry. Using the same procedure as before on 10 test samples on a different geometry, we find a mean error of 1.4mm and a median value at 1.3mm for d_{sur} (respectively, 1.8mm and 1.6mm for d_{nod}).

5 Discussion

Despite working around the bias and error introduced by the model and image processing in real applications, our synthetic experiments show promising performance for our framework in terms of accuracy, tolerance to non-linear effects of parameters, robustness and computational efficiency. The accuracy of our approach was found to be below

the typical voxel dimension (1mm), while a priori optimizing among a very wide range of parameter values at test time, and using a reasonable number of training samples at learning time. Although a single geometry is used for the training phase, the accuracy was of the same order on similar (non-pathological) heart geometries. Naturally, further work should handle geometry variability in a proper way, taking it into account at the training stage, and adding "shape factors" to the model parameter space capturing 3D shape variability. Moreover the addition in the pipeline of a pre-clustering stage with respect to the heart geometry, so as to distinguish very different geometries and treat them separately, should reduce the number of samples required to cover the whole parameter space while achieving better model personalization.

The proposed framework also brings an interesting perspective on the issue of parameter identifiability. It should be noticed that we achieve good results in terms of spatial distance between the matched model and observations while significant differences in the parameter space may still be observed. Parameter identifiability encompasses two distinct aspects. Firstly, small variations of the parameter values may result in changes that are not noticeable at the scale of reference. This sensitivity to parameters partially explains the error on the retrieved set of parameters. In our approach, the kernel width for currents impacts the ability of the algorithm to discern shape differences. In the future we will experiment with smaller kernel widths and improve algorithms to handle increased computational cost. Secondly in joint parameter estimation, a whole subset in the parameter space may result in identical observations, which also affects parameter identifiability. Such considerations can be analyzed in depth at the regression or optimization steps: several parameter sets with similar costs along with a measure of local sensitivity around these values may be additionally output by the Simulated Annealing algorithm. Biophysical priors may also be introduced at the optimization step by penalizing unlikely parameter sets without adding significant computational cost.

Finally more efficient machine learning algorithms should be tested in lieu of PCA, so as to capture non-linear 4D shapes variation, and to obtain and exploit precise information about the manifold structure of 4D heart shapes. Not only will this be of help with parameter identifiability and to derive efficient representations in the reduced shape space, but it could also provide valuable feedback for "smart" sampling of the parameter space.

6 Conclusion

A machine-learning current-based method has been proposed in this paper for the personalization of electromechanical models of the heart from patient-specific kinematics. A framework to encapsulate information regarding shape and motion in a way that allows the efficient computation of statistics via 4D currents has been described. This approach has been evaluated on synthetic data using the BCS model, with the joint estimation of the maximum contraction, relaxation rate and viscosity. It is found that the proposed method is accurate, computationally efficient and robust.

Acknowledgments. This work was partly funded by Microsoft Research through its PhD Scholarship Programme and by the ERC Advanced Grant MedYMA.

References

1. Aronszajn, N.: Theory of reproducing kernels. Harvard University (1951)
2. Bestel, J., Clément, F., Sorine, M.: A biomechanical model of muscle contraction. In: Niessen, W.J., Viergever, M.A. (eds.) MICCAI. LNCS, vol. 2208. pp. 1159–1161. Springer (2001)
3. Chabiniok, R., Moireau, P., Lesault, P.F., Rahmouni, A., Deux, J.F., Chapelle, D.: Estimation of tissue contractility from cardiac cine-MRI using a biomechanical heart model. *Biomechanics and Modeling in Mechanobiology* 11(5), 609–630 (2012)
4. Chapelle, D., Le Tallec, P., Moireau, P., Sorine, M.: An energy-preserving muscle tissue model: formulation and compatible discretizations. *IJMCE*, 10(2):189-211 (2012)
5. Davis, G., Mallat, S., Avellaneda, M.: Adaptive greedy approximations. *Constructive approximation* 13(1), 57–98 (1997)
6. Delingette, H., Billet, F., Wong, K., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C., Razavi, R., Ayache, N., et al.: Personalization of Cardiac Motion and Contractility from Images using Variational Data Assimilation. *IEEE Trans. Biomed. Eng.* 59(1), 20 (2012)
7. Durrleman, S.: Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution. Ph.D. Thesis, INRIA (March 2010)
8. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Measuring brain variability via sulcal lines registration: a diffeomorphic approach. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791. pp. 675–682. Springer, Heidelberg (2007)
9. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Sparse approximation of currents for statistics on curves and surfaces. In: Metaxas, D., Axel, L., Szekely, G., and Fichtinger, G. (eds.) Proceedings MICCAI, Part II, LNCS, vol. 5242. pp. 390–398. Springer (2008)
10. Gärtner, T., Flach, P., Kowalczyk, A., Smola, A.: Multi-instance kernels. In: Proceedings of the 19th International Conference on Machine Learning. pp. 179–186 (2002)
11. Haussler, D.: Convolution kernels on discrete structures. Tech. rep., Technical report, UC Santa Cruz (1999)
12. Hoerl, A., Kennard, R.: Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* pp. 55–67 (1970)
13. Imperiale, A., Chabiniok, R., Moireau, P., Chapelle, D.: Constitutive parameter estimation using tagged-mri data. In: Metaxas, D., Axel, L. (eds.) Proceedings of FIMH’11, LNCS 6666. pp. 409–417. Springer (2011)
14. Liu, H., Shi, P.: Maximum a posteriori strategy for the simultaneous motion and material property estimation of the heart. *IEEE Trans. Biomed. Eng.* 56(2), 378–389 (2009)
15. Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N.: ilogdemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. *International journal of computer vision* 92(1), 92–111 (2011)
16. Marchesseau, S., Delingette, H., Sermesant, M., Rhode, K., Duckett, S., Rinaldi, C., Razavi, R., Ayache, N.: Cardiac Mechanical Parameter Calibration based on the Unscented Transform. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI, LNCS. vol. 7511. Springer (2012)
17. Schölkopf, B., Smola, A.: Learning with kernels: Support vector machines, regularization, optimization, and beyond. the MIT Press (2002)
18. Vaillant, M., Glaunes, J.: Surface matching via currents. In: Christensen, G.E., Sonka, M. (eds.) IPMI 2005. LNCS, vol. 3565. pp. 381–392. Springer, Heidelberg (2005)
19. Xiang, Y., Gubian, S., Suomela, B., Hoeng, J.: Generalized simulated annealing for efficient global optimization: the GenSA package for R. *The R Journal* (2012), forthcoming