

Longitudinal Analysis using Personalised 3D Cardiac Models with Population-Based Priors: Application to Paediatric Cardiomyopathies

Roch Mollero^{*1}, Hervé Delingette¹, Manasi Datar³, Tobias Heimann³, Jakob A. Hauser², Dilveer Panesar², Alexander Jones⁶, Andrew Taylor², Marcus Kelm⁴, Titus Kuehne⁴, Marcello Chinali⁵, Gabriele Rinelli⁵, Nicholas Ayache¹, Xavier Pennec¹, and Maxime Sermesant¹

¹ Université Cote d'Azur, Inria, Asclepios Team, Sophia Antipolis, France

² UCL Centre for Cardiovascular Imaging, London, United Kingdom

³ Imaging and Computer Vision, Siemens Corporate Technology, Erlangen, Germany

⁴ Deutsches Herzzentrum Berlin, Berlin, Germany

⁵ Ospedale Pediatrico Bambino Gesù. DMCCP Cardiology, Rome Italy

⁶ Department of Paediatrics, University of Oxford, Oxford, United Kingdom

Abstract. Personalised 3D modelling of the heart is of increasing interest in order to better characterise pathologies and predict evolution. The personalisation consists in estimating the parameter values of an electromechanical model in order to reproduce the observed cardiac motion. However, the number of parameters in these models can be high and their estimation may not be unique. This variability can be an obstacle to further analyse the estimated parameters and for their clinical interpretation. In this paper we present a method to perform consistent estimations of electromechanical parameters with prior probabilities on the estimated values, which we apply on a large database of 84 different heartbeats. We show that the use of priors reduces considerably the variance in the estimated parameters, enabling better conditioning of the parameters for further analysis of the cardiac function. This is demonstrated by the application to longitudinal data of paediatric cardiomyopathies, where the estimated parameters provide additional information on the pathology and its evolution.

1 Introduction

Tridimensional Personalised Cardiac Models are of increasing interest for clinical applications. They compute the myocardial motion under the influence of simulated electromechanics and haemodynamics, in order to simulate a heartbeat. Their equations usually depend on a large number of parameters, so after extracting a patient's heart geometry from clinical imaging, the first step to build a *personalised simulation* is to estimate parameter values for which the simulation matches the measured heartbeat.

* rochmollero@hotmail.com

Recent works have shown that the *personalised parameter values* can capture intrinsic properties of the heart [1]. In particular, the simulations can help predict the possible behavior of the heart to some changes associated to some specific conditions (such as exercise or drug treatment), leading to applications in therapy planning [2]. Although these studies are very promising, progress is facing three main obstacles:

- the difficulty to find large homogeneous cardiac databases where the same information is available for all the cases
- the non-uniqueness of the parameter values due to the sparsity of clinical measurements compared to the high number of parameters of the models
- the computational time required to run models (from few minutes to several days for the most complex ones). This can be a burden in the parameter estimation requiring many simulations, and even worse for large databases

Here we present a cardiac modelling study overcoming these difficulties by first building a homogeneous cohort of more than 61 patients including 22 controls and 39 children with various cardiomyopathies. For each case, MRIs were acquired, together with pressure and heart rate measurements, resulting in 84 heart mesh geometry and haemodynamic conditions. On the modelling side, we performed the estimation of 6 parameters to reproduce the stroke volume and pressure measurements. This was performed with prior probabilities on the parameter values, in order to overcome the problem of the parameters uniqueness. Finally, the personalisation for the full cohort was performed in a relatively short time (around 2 days), thanks to a "multi-fidelity" optimization scheme which predicts changes in simulations of the 3D model with a much faster and simpler 0D model.

This led to more consistent parameter values across the 84 cases, on which we studied the relationship to clinical condition and its evolution. In particular, using the follow-up data patients with cardiomyopathy we show that the evolution of parameters naturally suggest and improvement of the heart condition under therapy. Finally we demonstrate that these estimated parameters could also be complementary to the clinical measurements in order to characterise better the difference between healthy and cardiomyopathy cases.

2 Clinical Data

We used two different cohorts (C1 and C2) in this study. The two protocols were approved by the local Research Ethics Committees. First 22 volunteers (C1) who participated to a clinical study to assess the cardiovascular response after the ingestion of a high-energy (1635 kcal), high-fat (142g) meal after fasting for 12h, closely following the protocol in [3]. In this study, short axis cardiac cine MRI sequences were acquired before the ingestion and at one or more time points within 1h of the ingestion of the meal, in order to study the evolution of the blood flow in the arteries. After the meal ingestion, both the heart rate and the cardiac output increased by around 15%. However no substantial changes were

observed in the mean, diastolic and systolic pressures during digestion (compared to the intra-patient variability of the measurement).

The second cohort (C2) consists in 39 children with various cardiomyopathies, ranging from class I to IV on the Ross and NYHA classifications [4] for heart failure symptoms, from two different clinical centers. The cine MRI was acquired at their enrollment and for 4 of them, at follow-up (after few months). The most common symptom among this cohort is a dilation of the left ventricle (Dilated CardioMyopathy) with low ejection fraction.



Fig 1a: Typical mesh geometry of a heart in the cohort (C1)




Fig 1b: Mesh geometry of a paediatric heart with Dilated Cardiomyopathy (DCM) in the cohort(C2)

We performed the cardiac modelling of a total of 41 different instants across the 22 volunteers of the first cohort, and at the 39 enrollment times and 4 available follow-up times of the second cohort. This lead to a total of 84 complete set of cine MRI, cuff pressure measurements and heart rates. See Figs.1a,1b for typical heart geometries from each cohort.

3 Personalised Cardiac Modeling

3.1 3D Electromechanical Cardiac Model

From each MRI, a high-resolution biventricular tetrahedral mesh of the patient's heart morphology (around 15 000 nodes) is generated with a method similar to the one in [5]. On this mesh, a *myocardial fibre direction* is defined at each node of the mesh by varying the elevation angles of the fibre across the myocardial wall from $\alpha_1 = -80$ on the epicardium to $\alpha_2 = 80$ on the endocardium.

The depolarization times in the myocardium were computed with the Eikonal model using default values of conductivities and the APD was computed from the Heart Rate with classical values of the restitution curve. Myocardial forces are computed based on the Bestel-Clement-Sorine model as detailed in [6]. It models the forces as the combination of an active contraction force in the direction of the

fibre, in parallel with a passive anisotropic hyperelasticity driven by the Mooney-Rivlin strain energy. In this paper, we only consider two main parameters of the model: the *Maximal Contractility* σ and the *Passive Stiffness* c_1 .

Finally the mechanical equations are coupled with a haemodynamic model which implements the 4 phases of the cardiac cycle, and describes the pressure in the cardiac chambers with global values (see [6] for implementation details). In particular, the pressure of the aortic artery in particular is modeled with a 4-parameter Windkessel model [7], which main parameters are *blood inertia* L , the *arterial compliance* C and the *proximal and distal (peripheral) resistances* Z_C and R . A *mean venous pressure* P_{ve} has to be set as well. In the following, Z_C and L are set at a default value, while C , R and P_{ve} are estimated parameters.

3.2 Parameter Estimation with Priors

A typical *parameter estimation* problem is composed of simulated quantities called the "outputs" O (such as the simulated Stroke Volume and Mean Pressure), and a set of model parameters P . The estimation consists in finding adequate values x of the parameters such that the output values $O(x)$ in the 3D model simulation fit the "observed values" from the clinical measurements \mathcal{O} of interest.

This is done by minimizing a cost function (or **score**) $S(x, \mathcal{O})$ between the simulated values $O(x)$ and the target values \mathcal{O}

$$S(x, \mathcal{O}) = \|O(x) - \mathcal{O}\|^2 + \lambda R(x)$$

where $R(x)$ is a *penalty* (or *regularisation*) term, weighted with λ , that can be formulated as a quadratic form with mean value μ_R and covariance matrix Δ :

$$R(x) = (x - \mu_R)\Delta^{-1}(x - \mu_R)^T$$

In Bayesian Inference, this is equivalent to finding the *maximum a posteriori* with a Gaussian prior and a Gaussian likelihood.

Finally, since the parameters are positive, it is more meaningful to consider that the logarithm of the parameter values follow a Gaussian distribution rather than the parameter themselves, so the optimisation is performed over the logarithm of the parameters values.

In this scope, we focus on the set of **5 parameters** P : c_1 , σ , R , C and P_{ve} , in order to fit **3 target outputs** to their clinical measurements, which are the *Stroke Volume* **SV** the *Aortic Diastolic Pressure* **DP** and *Mean Pressure* **MP**.

We then performed two different personalisations: first, one (P1) without priors on the parameter values during the optimization ($\lambda=0$). This allowed us to

have a first assessment of the variability of the parameters and of the values which lead to the best simulations. It was then followed by a second personalisation (P2), with priors on the values of both c_1 and C equipped with diagonal covariance matrix:

$$\mu_R = \begin{pmatrix} 1 \\ \mu_{c_1} \\ \mu_C \end{pmatrix}, \Delta = \begin{pmatrix} 1 & & \\ & \delta_{c_1} & 0 \\ & 0 & \delta_C \end{pmatrix} \quad (1)$$

We observed in $P1$ that simulations with c_1 around 500 000 Pa lead to good behavior compared to the dynamics observed in echocardiographic images. For the arterial compliance C , we used a prior based on the mean value of $1.8e^{-8}$ $S.I.$ as reported in [8].

Therefore $\mu_c = \ln(500\ 000) = 13.12$, $\delta_a = 0.5$ and $\mu_C = \ln(1.8e^{-8}) = -17.83$, $\delta_c = 1$.

3.3 Efficient Multi-Fidelity Optimization

The optimization in each personalisation problem was performed with a "multi-fidelity" approach [9] based on our recent work in cardiac model personalisation [10], where the outputs of the 3D model are approximated during the optimization by simulations from a very fast low-fidelity model called "0D model", made of around twenty equations only.

As explained in [10], a small number of 3D "sigma-simulations" can indeed be used to approximate many 3D simulations in a large parameter space. This is done by first finding similar 0D simulations with the same outputs of interest (such as the pressure and stroke volume), then building a mapping between the parameters of the corresponding 3D and 0D simulations.

This multifidelity method was here adapted for the two personalisation problems, where the 0D and the 3D models share the same haemodynamic variables R , C and P_{ve} from the Windkessel model. In this specific case, the variations of pressure and stroke volume with respect to these variables are very similar, and we found to have better approximation results by only computing a mapping between c_1 and σ (as opposed to all the personalised parameter in our original approach) to the 0D model parameters.

As a result, our Multi-Fidelity Optimization Method performs the optimization of 3D parameters based on 0D model approximations which are very fast to compute. It only requires the computation of successive sets of 5 simulations of the 3D model (the sigma-simulations). In particular, the simultaneous personalisation (P1) of the 84 hearts was completed in around 36 hours, and in 48 hours for (P2).

4 Results

In Table 1 we report the mean of both the estimated values (*Mean*) and logarithmic estimated value (*Log-Mean*), as well as the standard deviation of the logarithm of the estimated values (*Log-Std*).

As expected, we can notice that the standard deviation of all the parameters in the population is reduced between (P1) and (P2). Interestingly, the goodness of fit was not impacted in the personalisation by the use of prior probabilities on c_1 and σ in (P2). Most cases are fitted under 1.2ml for the stroke volume and 0.5 mmHg for the pressure measurements with few outliers. This means that the prior could be stronger in order to further reduce the variability while maintaining simulations which match the clinical measurements.

Table 1: Statistics of the estimated parameters in the estimations (P1) and (P2)

	ρ (MPa)	c_1 (kPa)	R (MPa.m ³ .s)	C (MPa ⁻¹ .m ⁻³)	P_{ve} (mmHg)
<i>Mean</i> (P1)	156	$37.8e^1$	48.0	$6.13e^{-3}$	44.6
<i>Log-Mean</i> (P1)	18.51	12.60	17.62	-18.96	8.62
<i>Log-Std</i> (P1)	0.98	0.66	0.37	0.32	0.44
<i>Mean</i> (P2)	$71.1e^7$	$4.44e^1$	84.1	$7.29e^{-3}$	32.01
<i>Log-Mean</i> (P2)	18.03	12.98	18.2	-18.77	8.32
<i>Log-Std</i> (P2)	0.41	0.22	0.26	0.25	0.27

4.1 Application to Longitudinal Analysis of the Cardiac Function

From a clinical point view, an interesting application of the modeling is to characterise the state of the heart function, beyond the information given by the clinical measurements and the imaging. The underlying idea is that some of the estimated parameters values can capture properties of the heart which cannot be directly measured from the imaging (such as the myocardial contractility). This additional information on the heart could contribute to the diagnosis, by comparing the estimated parameters with the parameters of other known cases.

To analyse the relationship between the parameter values and the clinical condition, we performed here a linear discriminant analysis (LDA) over the parameter values and the heart rate, in order to classify between the two cohorts. This leads to the computation of two vectors w and b such that given a vector X of parameter x , the predicted cohort is C2 if $A^T X + b > 0$, and C1 otherwise.

The vector w corresponds to the most discriminative direction in the population between healthy and cardiomyopathy cases. In this context, this axis could be a candidate to characterise whether the cardiac function at a given time is closest to a healthy heart, or a heart with cardiomyopathy, based on the parameters values observed in the two cohorts.

For example, we display in Fig ?? the projection of the parameters on this vector (x-axis) and a orthogonal direction to w . Most healthy cases (dark blue dots) are on left side of the black line ($w < 0$) and most cardiomyopathy cases (red dots) are on the right ($w > 0$).

Interestingly, this could also help to quantify the evolution of the patient's heart condition under the influence of the pathology and the therapy. Indeed, for all the cardiomyopathy cases for which we have the follow up data, we can notice a decrease in the coordinates along the horizontal axis (see the pairs of brown, light blue, green and orange dots. The larger dot is the follow-up). This could be interpreted as *an improvement of the cardiac function with the therapy*, which is at least becoming closer to the condition of an healthy heart. One of the cases (in light blue) is on the "healthy" side of the classification at follow-up. Finally, the predictive power of such a classifier can be assessed, through *leave-one out cross-validation*. This is done by training the classifier on all the cases but one, and predicting the diagnosis for the remaining case. If we perform the LDA over the 5 estimated parameters and the heart rate, the number of




Fig. 2: Projection of the parameter on the main direction w of a LDA classifier between the healthy cases (dark blue dots) and cardiomyopathy (other dots) cases (x-axis) and an principal orthogonal direction of this vector (y-axis). The dots in light blue, brown, orange and green correspond to 4 patients for which the data was available both at baseline (small dot) and follow-up (larger dot).

prediction errors is 11. The same classifier trained on the 3 outputs (stroke volume and pressures) and the heart rate makes 9 prediction errors. However, if we train the classifier with *both the 5 estimated parameters and the clinical measurements* and the heart rate, it only makes **6 classification errors**.

In this context, this can mean that the estimated parameters were able to capture a more complex information on the cardiac function than the clinical measurements of volume and pressure only, both through the *3D personalisation* and the *comparison with the other values estimated in the population*. This information was then used by the simple linear classifier to improve its accuracy in the diagnosis of a patient.

5 Conclusion and Discussion

In this manuscript, we presented a cardiac modelling study based on the estimation of 5 model parameters from 3 clinical measurements of stroke volume and pressures, on a large cohort of 61 patients. We used recent ideas developments in "multi-fidelity" personalisation, to drive a very fast and computationally efficient estimation of these parameters with priors. Both the personalisations with and without priors were performed simultaneously for all the patients on our cluster,

and converged respectively in less than 36 and 48 hours. We showed that the use of priors during optimization reduces the variability of the estimated values in the population, leading to more consistency for further applications.

We then analyzed the estimated parameter values with respect to the clinical conditions of the patients. A linear discriminant analysis (LDA) was used to characterise the cardiac function the cases along the most discriminative axis between the two cohorts. For cardiomyopathy patients, we showed that the evolution in time along this axis suggests that their cardiac function is improving under therapy. Finally, we also demonstrated how the estimated parameter values could be complementary to clinical measurements in the context of diagnosis.

A direct extension of this study is to estimate values for more model parameters, from a larger set of measurements such as the flow or the myocardial strain. On the cardiomyopathy point of view, this could help to further discriminate between the various types of cardiomyopathy, with applications in risk stratification of heart failure. Another interesting direction of study is the further analysis of longitudinal data in order to better understand both the short-term and long-term variabilities in cardiac function, with applications on the prediction of disease evolution and therapy planning.

Acknowledgements This work has been partially funded by the EU FP7-funded project MD-Paedegree (Grant Agreement 600932) and contributes to the objectives of the ERC advanced grant MedYMA (2011-291080).

References

1. Chabiniok, R., et al.: Estimation of tissue contractility from cardiac cine-MRI using a biomechanical heart model. *Biomech Model Mechanobiol* **11**(5) (2012) 609–630
2. Crozier, W.A.: Personalised electromechanical modelling of cardiac resynchronisation therapy. PhD thesis, King's College London (2015)
3. Hauser, J.A., et al.: Comprehensive assessment of the global and regional vascular responses to food ingestion in humans using novel rapid mri. *Am. J. Physiol. Regul. Integr. Comp. Physiol.* **310**(6) (2016) R541–R545
4. Hsu, D.T., et al.: Heart failure in children. *Circ. Heart. Fail.* **2**(1) (2009) 63–70
5. Molléro, R., et al.: Propagation of Myocardial Fibre Architecture Uncertainty on Electromechanical Model Parameter Estimation: A Case Study. In: FIMH, LNCS. (2015) 448–456
6. Marchesseau, S., et al.: Fast parameter calibration of a cardiac electromechanical model from medical images based on the unscented transform. *Biomech Model Mechanobiol*, **12**(4), 815–831 (2013)
7. Westerhof, N., et al.: The arterial windkessel. *Med Biol Eng Comput* **47**(2) (2009) 131–141
8. Laskey, W.K., et al.: Estimation of total systemic arterial compliance in humans. *J Appl Physiol* **69**(1) (1990) 112–119
9. Peherstorfer, B., et al.: Survey of multifidelity methods in uncertainty propagation, inference, and optimization. ACDL Technical Report TR-16-1 (2016)
10. Molléro, R., et al.: A multiscale cardiac model for fast personalisation and exploitation. In: MICCAI, Springer (2016) 174–182