

Population-Based Priors in Cardiac Model Personalisation for Consistent Parameter Estimation in Heterogeneous Databases

Roch Molléro | Xavier Pennec | Hervé Delingette | Nicholas Ayache | Maxime Sermesant*

¹Inria, Epione Research Project, Sophia Antipolis, France

Correspondence

Email: maxime.sermesant@inria.fr

Personalised cardiac models are a virtual representation of the patient heart, with parameter values for which the simulation fits the available clinical measurements. Models usually have a large number of parameters while the available data for a given patient is typically limited to a small set of measurements, thus the parameters cannot be estimated uniquely. This is a practical obstacle for clinical applications, where accurate parameter values can be important. Here we explore an original approach based on an algorithm called Iteratively Updated Priors (IUP), in which we perform successive personalisations of a full database through Maximum A Posteriori (MAP) estimation, where the prior probability at an iteration is set from the distribution of personalised parameters in the database at the previous iteration. At the convergence of the algorithm, estimated parameters of the population lie on a linear subspace of reduced (and possibly sufficient) dimension in which for each case of the database, there is a (possibly unique) parameter value for which the simulation fits the measurements. We first show how this property can help the modeler select a relevant parameter subspace for personalisation. In addition, since the resulting priors in this subspace represent the population statistics in this subspace, they can be used to perform consistent parameter estimation for cases where measurements are possibly different or missing in the database, which we illustrate with the personalisation of a heterogeneous database of 811 cases.

KEYWORDS:

Cardiac Electromechanical Modeling, Parameter Estimation, Parameter Selection, Personalised modeling

1 | INTRODUCTION

Personalised cardiac models are of increasing interest for clinical applications^{13,1,22}. To that end, parameter values of a cardiac model are estimated to get a *personalised* simulation which reproduces the available measurements for a clinical case. Then the personalised simulations can be used for advanced analysis of pathologies. In particular, recent works have been successful in predicting haemodynamic changes in cardiac resynchronization therapy²¹, ventricular tachycardia inducibility and dynamics⁶, as well as in detecting and localising infarcts⁸ using 3D personalised models.

Most cardiac models depend on many parameters, especially in 3D models where each parameter can take a different value at each node or region of the mesh,^{16,13,2,9}. The number of parameters can be up to tens of thousands while on the other hand for an individual patient, the available clinical data is usually sparse and typically consists of a set of global measurements (such as ejection fraction, systolic and diastolic pressure) and possibly some imaging data with a considerable degree of noise or blurriness. Consequently, many sets of parameter values can lead to a simulation which can reproduce the available clinical data. The parameter estimation problem is an *ill-posed inverse problem* and all the parameters cannot be uniquely estimated from the measurements only.

A classic technique to estimate relevant parameters in this context is the use of *prior probabilities distributions* over the parameter values^{19,14,15}. In this framework, parameters values come from a (usually Gaussian) *prior* probability distribution, which represents some knowledge or beliefs about the distribution of parameters. Then, given a set of measurements, a vector of *Maximum A Posteriori* (MAP) parameter values is estimated, which realises an optimal trade-off between its likelihood in the prior probability distribution and the error of fit of the simulation. The set of *Maximum A Posteriori* values is usually smaller than the set of values for which the simulation fits the measurements and possibly unique. A challenge to estimate relevant parameters is then to define accurately the *prior probability distribution*, which should be ideally as close as possible to the "true" distribution of parameters.

Here we explore an original approach which we call **Iteratively Updated Priors** (IUP), to address the chicken-and-egg problem of defining an accurate "prior" when not other information is available beside the currently considered data. It consists in performing successive personalisations with priors, where the priors distribution is set from the distribution of personalised parameters in the previous personalisation. The rationale is twofold: first, if the algorithm converges, all the cases are personalised through a MAP where the prior distribution is the distribution of the population parameters itself. Intuitively this opens the possibility, for cases where the value of a parameter is unobservable, to *regularise the fitting by using population information* from other cases in the database where the value was observable. Secondly, when there are directions in the parameter space in which the simulated measurements do not vary (i.e. an *unobservable* parameter direction), the use of priors promotes the estimation of parameter values which are closer to the prior mean, which in turn makes the penalty (from the prior) stronger in this direction at the next personalisation. Intuitively, this process will tend to group parameters onto directions which are observable, and reduce the spread into directions which are not. In practice, we will see this algorithm leads the parameters to lie on a *linear subspace* of smaller dimension (See Figure 1).

In Section 2 we present the IUP algorithm in relation to two different mathematical tools. First, we express parameter estimation as a Maximum A Posteriori (MAP) estimation with prior probability distributions. Secondly, we use the *Iteratively Reweighted Least Square* (IRLS) algorithm from sparse regression to show that the IUP algorithm, when it converges, performs the minimisation of *population-wide cost function* with a sparse regulariser on the *number of dimensions of the parameter space*. We introduce two types of updates of the parameters of the (Gaussian) prior probability: "**Full Matrix**", where no assumption is made on the distribution and the covariance matrix, and "**Diagonal Matrix**" where the covariance matrix

is supposed diagonal, which assumes independence of parameters in the distribution of the population. We demonstrate that in the first case the algorithm leads to the automatic selection of *parameter directions* in which parameter values are set to a constant value in the population, and in the second case to the automatic selection of *parameters* which are set constant.

In Section 3 we present results of the IUP algorithm on a database of 137 *complete cases* where the same 4 measurements are available. We perform the estimation of 6 parameters of a cardiac model likely to vary in this population, which also makes the problem ill-conditioned. We observe that the IUP algorithm leads to a trade-off between the number of dimensions of the final parameter set, and the mean error of fit of the 4 measurements across the whole database of personalised simulations. Then we show that if we impose a high goodness of fit of the personalised simulations, the algorithm leads to the selection of a *linear subspace of minimal dimension* in which *for each case of the database* there are parameter values where the simulation fits exactly the measurements. We discuss how this algorithm can be used to find a relevant subspace for personalisation from the modeling point of view. Finally we extract a relevant subspace of dimension 4, based on 5 parameters out of the original 6, in which there is a unique set of parameter values for each case of the database.

In the last section we apply the proposed framework to the consistent personalisation of a larger database of 811 cases, in which the set of measurements reported by the clinicians are heterogeneous and sometimes incomplete. Applied to this database, the algorithm leads to personalised simulations with high goodness of fit for all the cases of the population while grouping the parameters onto a parameter space of minimal dimension. Simultaneously, the resulting *population-based priors* computed in this parameter space lead to consistent parameter estimation for the cases where the parameters are not all observable, and are in particular influenced by their values in cases of the population where they are observable. Finally we show how *external parameters* such as the weight and the height of the patient can also be added to the prior distribution in order to constrain the parameter estimation with more consistency in case of unobservability.

2 | ITERATIVELY UPDATED PRIORS AND POPULATION-BASED PRIORS

In this section we introduce the personalisation framework, based on the estimation of a *Maximum A Posteriori* (MAP) with prior probabilities. In order to get relevant values for each case through MAP estimation, the prior probability has to be relevant. Ideally, it should be the "true" underlying probability distribution of parameters in the population. We introduce




FIGURE 1 Schematic representation (with synthetic data) of parameter estimation when the problem is ill-posed: in this example both the contractility and the stiffness are estimated from values of the *stroke volume* (SV). Both have an influence on the *stroke volume* (SV) so there are *isolines* of stroke volume (in grey) on which the value of the stroke volume is the same. **Image a:** estimation without priors, the estimated values (green) for each case can be anywhere on an *isoline* (grey). **Image b:** the estimation is performed with a (gaussian) prior (the gaussian covariance is in blue), the estimated values are grouped closer to the prior mean. **Image c:** the **Iteratively Updated Priors** (IUP) algorithm performs successive estimations where the prior is set from the distribution of estimated parameters at the previous iteration. This leads the parameters to lie on a *reduced linear subspace* (orange).

here the *Iteratively Updated Priors* (IUP) algorithm, which performs successive personalisations of the whole population, where the prior at each step is set from the distribution of the personalised parameters at the previous step.

First we present the MAP estimation of personalised parameters or a single case (Section 2.1). Then we formulate the *Iteratively Updated Priors* algorithm for a population of cases (Section 2.2). We then give the equations of the prior parameter update in two different cases with different assumptions on the shape of the distribution (Section 2.3) and the equations which are minimised within our optimisation framework (Section 2.4). Finally we present a link between the IUP algorithm and the *Iteratively Reweighted Least Square* algorithm for sparse regularisation, and show that the IUP algorithm performs the minimisation of a *population-wide* cost function with a penalty on the *rank* of the set of personalised parameters (Section 2.5). Finally we present a modification of the IUP algorithm called **IUP*** to impose a high goodness of

fit to the model outputs to the target values for each case and at each iteration (Section 2.6).

2.1 | Maximum A Posteriori Estimation of Personalised Parameters

We consider a cardiac model M and a set of simulated quantities called the *outputs* \mathcal{O} , such as the ejected volume and the mean ventricular pressure. We then consider a subset \mathcal{P}_M of varying parameters (such as the contractility, the aortic resistance) of the model, while the other model parameters \mathcal{P}'_M are supposed fixed. Given a vector x of values of the parameters \mathcal{P}_M , we note $\mathcal{O}_M(x)$ the values of the outputs \mathcal{O} of the model M .

With these notations, personalisation consists in estimating parameter values X for which the outputs values $\mathcal{O}_M(X)$ are consistent with some observed values $\hat{\mathcal{O}}$ of these outputs. This is an inverse problem, which can be tackled by different methods (see the review of Chabiniok et al.⁴).

A classic approach is to formulate the problem as the estimation of a *Maximum A Posteriori* with prior probabilities on the parameter values. With Bayes' theorem, the *posterior probability* $P(x|\hat{\theta})$ of parameter values considering the observed output values $\hat{\theta}$ is proportional to the product of the *conditional likelihood* $P(\hat{\theta}|x)$ and the *prior probability* $P(x)$:

$$P(x|\hat{\theta}) \propto P(\hat{\theta}|x)P(x) \quad (1)$$

In this work, we use Gaussian distributions for both the conditional likelihood and the prior probability:

$$\begin{cases} P(\hat{\theta}|x) \propto e^{-\frac{1}{2}(O(x)-\hat{\theta})^T \Sigma^{-1}(O(x)-\hat{\theta})} \\ P(x) \propto e^{-\frac{1}{2}(x-\mu)^T \Delta^{-1}(x-\mu)} \end{cases} \quad (2)$$

which correspond to the assumptions that first, the observed values $\hat{\theta}$ are the sum of the outputs values $O(x)$ and a Gaussian noise with covariance Σ . Second, the underlying probability distribution of the model parameters is a Gaussian with mean μ and covariance Δ . Finding a *Maximum A Posteriori* (MAP) consists in finding (one of) the parameter values which has the highest posterior probability:

$$X = \operatorname{argmax}_x P(x|\hat{\theta}) \quad (3)$$

This can also be interpreted as finding a maximum of the joint probability $P(x, \hat{\theta})$ of x and $\hat{\theta}$ with a uniform prior $P(\hat{\theta})$ on $\hat{\theta}$, thanks to the formula:

$$P(x, \hat{\theta}) = P(x|\hat{\theta})P(\hat{\theta}) \quad (4)$$

2.2 | Iteratively Updated Priors (IUP)

In the following we consider a population of $i = 1..n$ cases to be personalised, we note $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$ and $\mathcal{X} = (x_1, \dots, x_n)$. We introduce the algorithm called *Iteratively Updated Priors* (IUP) which consists in successively estimating personalised parameters X_i with the MAP estimation, and re-estimating the prior parameters (μ, Δ) as the maximum likelihood of these parameters considering the estimated personalised parameters X_i :

- **Step 1: Personalisation**

For each case i , find $X_i = \operatorname{argmax}_{x_i} P(x_i, \hat{\theta}_i | \mu, \Delta)$. This consists in performing the MAP estimation in Equation 3 for each case with current μ and Δ , which is the estimation of personalised parameters with a given prior.

- **Step 2: Re-estimation of the Prior**

Find $(\mu, \Delta) = \operatorname{argmax}_{(\mu, \Delta)} P(\mathcal{X} | \mu, \Delta)$.

This consists in finding the most likely values (maximum likelihood estimation) of μ and Δ given a set of personalised parameters \mathcal{X} .

When the algorithm converges, the prior probability distribution with the resulting parameters (μ^*, Δ^*) is called a *population-based prior* based on the $\hat{\theta}_i$.

2.3 | Explicit update formulas for μ and Δ in two cases.

We explicit here the update in the **Step 2** of the IUP algorithm in which μ and Δ are set to the solutions of $(\mu, \Delta) = \operatorname{argmax}_{(\mu, \Delta)} P(\mathcal{X} | \mu, \Delta)$. We consider two formulations of the Gaussian prior probability:

- First, a case (**Full Matrix**) where no assumptions are made on the covariance matrix Δ of the Gaussian distribution.
- Second, a case (**Diagonal Matrix**) where we suppose that Δ is diagonal. This is equivalent to the (simplifying) assumption that parameter values are independent from each other in the population.

The solutions of the maximum likelihood estimation of μ and Δ in the two cases are¹:

- **Update of μ :** In both cases, μ is set to the mean of personalised parameters X_i :

$$\mu^k = \frac{1}{n} \sum_{i=1}^n X_i^{k-1} \quad (5)$$

- **"Full Matrix" update of Δ :** In this case the minimisation is reached when Δ is the covariance matrix of the personalised parameters:

$$\Delta^k = \frac{1}{n} \sum_{i=1}^n (X_i^{k-1} - \mu^k)(X_i^{k-1} - \mu^k)^T \quad (6)$$

- **"Diagonal Matrix" update of Δ :** In this case the minimisation is reached when Δ is the *diagonal part* of the covariance matrix of the personalised parameters:

$$\Delta^k = \operatorname{Diag}\left(\frac{1}{n} \sum_{i=1}^n (X_i^{k-1} - \mu^k)(X_i^{k-1} - \mu^k)^T\right) \quad (7)$$

2.4 | Practical Implementation of the MAP Estimation.

In the following, we use a constant $\Sigma = \gamma \operatorname{diag}(\mathcal{N})$ for the noise model in 2, where \mathcal{N} is a normalisation vector whose coefficients are explicitated later.

¹A practical derivation of the maximum likelihood estimation of the covariance matrix can be found in <https://people.eecs.berkeley.edu/~jordan/courses/260-spring10/other-readings/chapter13.pdf>

We find the MAP of Equation 3 through the minimization of the *negative log-likelihood* $-\frac{2}{\lambda} \log(P(x|\hat{O}), \mu, \Delta)$, a *regularised cost-function* which we call \hat{S} :

$$\hat{S}(x, \hat{O}, \mu, \Delta) = S(x, \hat{O}) + \gamma R(x) \quad (8)$$

with

$$\begin{cases} R(x) = (x - \mu)^T \Delta^{-1} (x - \mu) \\ S(x, \hat{O}) = \|(O(x) - \hat{O}) \oslash \mathcal{N}\|^2 \end{cases} \quad (9)$$

where \oslash is the Hadamard (coordinate-by-coordinate) division. We call $S(x, \hat{O})$ the *data-fit term* and $R(x)$ the *regularisation term*.

We perform optimisation of this function with a derivative-free algorithm called *CMA-ES*¹¹, which stands for Covariance Matrix Adaptation Evolution Strategy. It asks at each iteration n for the scores of m points x_i (a *generation*), drawn from a multivariate distribution with covariance I_n^c and mean I_n^m . Then, it combines Bayesian principles of *maximum likelihood* with *natural gradient descent* on the ranks of the points scores in the generation to update both I_n^c and I_n^m .

For each MAP estimation, we perform from 50 to (maximum) 250 iterations of CMA-ES with a population size of 20, and the algorithm also stops if all the values of $\sqrt{\hat{S}}$, (which we call the *data-fit term*) within the search space are in an interval smaller than 0.01.

The algorithm is stochastic (because the parameter values of the generation samples are drawn randomly). It iteratively enlarges its search space until it contains a minimal solution, then reduces the search space according to the best members of each generation (drawn randomly). As a consequence it does not converge to the same solution when the possible solutions are not unique. Because of this it can be used to give an (empirical) evaluation of the "non-uniqueness variability" by repeating the personalisation process multiple times.

Finally for practical reasons, since all the parameters of our model are positive, we always consider in the following the logarithm of parameter values instead of their values. This enables in particular to not have the optimisation algorithm test negative (non-physical) values in this step.

2.5 | Link with Sparse Regularization

The IUP algorithm as presented in Section 2.2 has strong links with the *Iteratively Reweighted Least Squares (IRLS)* algorithm in sparse regression. When the IUP converges, we can show (see a more complete derivation in APPENDIX A) that it performs the minimization of a *population-wide cost function*, with a penalty on the *rank of the set of personalised parameters*:

$$\mathcal{S}(\mathcal{X}, \hat{\Theta}) = \frac{1}{n} \sum_{i=1}^n S(x_i, \hat{O}_i) + \gamma D \quad (10)$$

Where D is the rank (number of non-zero eigenvalues) of Δ^* , which is the covariance of the set of personalised parameters at convergence. The minimisation of D thus correspond to the minimisation of the *number of dimensions of the set of personalised parameters*. With the **Diagonal Matrix** assumption, the covariance matrix eigenvectors are aligned with the coordinate axes so in this case, the IUP algorithm minimises the *number of parameters which do not have a constant value in the population of personalised parameters*. In order to illustrate the difference between the two types of updates, we display in Figure 2 a schematic representation of the two different behaviours.

2.6 | IUP* algorithm: Equations in the Limit Case where the Noise is Considered Negligible

The MAP estimation of Section 2.1 and the cost function minimisation in Section 2.4 makes the assumption that the measured values of an output is the sum of its true value and a Gaussian noise (Equation 2). Though this is formulation is practical and can be especially appropriate in cardiac personalisation since clinical data is often noisy, it has drawbacks. In particular, it makes the questionable assumption of having the same Gaussian noise model for all the cases in the database, whose parameter values has to be estimated or manually imposed. In addition, we might want in some context to consider the noise to be negligible or null, and have the model outputs fit exactly the target values in the clinical data (This will be useful in Section 4).

To that end, we introduce here a variation of the IUP algorithm called **IUP***, where the data-fit term is progressively reweighted during parameter estimation, until it reaches a value lower than a specific threshold. Namely, we consider a *modified cost function* \hat{S}^* from the Equation 8 of Section 2.4, with an additional term α :

$$\hat{S}^*(x, \hat{O}, \mu, \Delta, \alpha) = \alpha S(x, \hat{O}) + \gamma R(x) \quad (11)$$

In the IUP algorithm, we first initialise α to 1 at the first iteration. Then for each parameter estimation of each case in the database we perform the following algorithm, which ensures that the data-fit term of every case is below 0.01:

1. **optimise** $\hat{S}^*(x, \hat{O}, \mu, \Delta, \alpha)$
2. **if** $S(x, \hat{O}) > 0.01$ set $\alpha = 2\alpha$ and go to step 1.
3. **else** stop.

Then, in each subsequent iterations of the IUP algorithm, we use the final value of α from the previous iteration as an initialisation. It is important to note that this formulation makes the assumption that there exists at least one set of parameter values x for which $S(x, \hat{O}) < 0.01$ for which measurements are fitted very closely.




FIGURE 2 Behaviour of the algorithm for the two different updates (See legend in Figure 1). **Image a:** with the **Full Matrix** updates, there is no assumption on the covariance matrix so the prior distribution at the next iteration is the closest gaussian distribution (in terms of *maximum likelihood*) to the estimated parameters distribution. **Image b:** with the **Diagonal Matrix** updates, the prior distribution at each iteration is the closest gaussian distribution *with the axes parallel to the coordinate axes*. In the final population of personalised parameters, this leads some *parameters* to have a constant value (here the stiffness).

From a mathematical point of view, this algorithm is equivalent to solving the following problem:

$$\min_x \gamma R(x) \quad s.t. \quad S(x, \hat{O}) \leq 0.01 \quad (12)$$

which is a relaxed formulation of the *constrained optimisation problem* with the constraint that measurements are perfectly fitted:

$$\min_x \gamma R(x) \quad s.t. \quad S(x, \hat{O}) = 0 \quad (13)$$

We can also use the sparse formulation in Section 2.5 to exhibit the optimisation of a cost function over the whole database: by considering the mean of all *modified cost functions* over all cases, we observe that this is equivalent to solving the following problem over the complete database:

$$\min_{\mathcal{X}} \gamma \frac{1}{n} \sum_{i=1}^n R(x_i) \quad s.t. \quad \forall i = 1..n, \quad S(x_i, \hat{O}) \leq 0.01 \quad (14)$$

which is a relaxed version of the constrained optimisation problem over the whole database:

$$\min_{\mathcal{X}} \gamma \frac{1}{n} \sum_{i=1}^n R(x_i) \quad s.t. \quad \forall i = 1..n, \quad S(x_i, \hat{O}) = 0 \quad (15)$$

At convergence of the IUP algorithm, we have $R(x) = (x - \mu^*)^T (\Delta^*)^{-1} (x - \mu^*)$ where μ^* and Δ^* are respectively the mean and covariance of the set of personalised parameters \mathcal{X} (see APPENDIX A). We then have:

$$\frac{1}{n} \sum_{i=1}^n R(x_i) = D \quad (16)$$

where D is the rank of Δ^* . Overall we are then solving:

$$\min_{\mathcal{X}} \gamma D \quad s.t. \quad \forall i = 1..n, \quad S(x_i, \hat{O}) \leq 0.01 \quad (17)$$

which is a relaxed version of the constrained optimisation problem:

$$\min_{\mathcal{X}} \gamma D \quad s.t. \quad \forall i = 1..n, \quad S(x_i, \hat{O}) = 0 \quad (18)$$

Since at convergence Δ^* is the covariance of the set of personalised parameters, in other terms, we are minimising the *number of dimensions of the set of personalised parameters* under the constraint that all cases have a high goodness of fit to their corresponding measurements.

3 | RESULTS ON 137 COMPLETE CASES AND APPLICATION TO PARAMETER SELECTION.

In this section we present results on the personalisation of 137 *complete cases* with a cardiac model. The cardiac model we use is a fast "0D" model which we introduced in¹⁸. It is a reduced version of our in-house 3D electromechanical model, based on the implementation of the Bestel-Clement-Sorine (BCS) model⁵ by Marchesseau et al.^{17,16} in SOFA². As described in¹⁸, both the 3D and 0D model share the same mechanical and haemodynamic equations, but simplifying assumptions (see APPENDIX B) are made on the geometry of the 0D model. This leads to a very fast model made of 18 equations, which can simulate 15 beats per second at a heart rate of 75 bpm. In our context it is particularly practical for experiments which involve repeated personalisations of many cases, since the personalisation of a single case with CMA-ES as described in Section 2.4 takes around 3 minutes.

We consider a set \mathcal{P}_{OD} of 6 parameters of the model, and a set of 4 measurements \mathcal{O} listed in Table 1 for which *the 4 values are available for all the 137 cases*.

The normalisation coefficients for this problem (in the vector \mathcal{N} defined in Section 2.4) are **10 ml** for the Maximal Volume (MV) and the Stroke Volume (SV), **200 Pa** for the Diastolic Aortic Pressure (DP) and the Mean Aortic Pressure (MP). With this normalisation, an *error of fit* $\sqrt{S(x, \hat{\mathcal{O}})}$ lower than 1 (resp 0.1), means a personalised simulation matches volume measurements within 10 ml (resp 1 ml) and pressure measurements within 200 Pa (resp 20 Pa), which can qualitatively be considered as *acceptable (very good)*.

3.1 | Trade-off between Data-fit, Regularisation and Dimension of the Set of Personalised Parameters

We provide results of the IUP algorithm for at least 20 iterations, with both the *Diagonal Matrix* and the *Full Matrix* assumptions for values of γ of 0.5, 0.1, 0.05 (see names of the run in Table 2).

In Figure 3 , we report for each run the mean value $\bar{S}(\mathcal{X}, \hat{\mathcal{O}})$ of the *data-fit term* $S(x, \hat{\mathcal{O}})$ and the mean value $\bar{R}(\mathcal{X})$ of the *regularization term* $\gamma R(x)$ at convergence of

TABLE 1 Sets of *0D model parameters* and *outputs* for which the measurements are available in the example.

Personalised Parameters \mathcal{P}
Contractility σ_0
Resting Radius R_0
Stiffness c_1
Peripheral Resistance R_p
Windkessel Characteristic Time τ
Venous Pressure P_{ve}
Outputs \mathcal{O}
Maximal Volume <i>MV</i>
Stroke Volume <i>SV</i>
Mean Aortic Pressure <i>MP</i>
Diastolic Aortic Pressure <i>DP</i>

Name	γ	Update
FULL-IUP-005	0.05	Full Matrix
FULL-IUP-01	0.1	
FULL-IUP-05	0.5	
DIAG-IUP-005	0.05	Diagonal Matrix
DIAG-IUP-01	0.1	
DIAG-IUP-05	0.5	

TABLE 2 Names and parameters of the 6 runs of the IUP algorithm

the MAP estimations (bottom). We also perform the Singular Value Decomposition (SVD) of the set of personalised parameters and report the eigenvalues (Figure 3 , top). Then we report (rows of Table 3) the eigenvectors at the last iteration ($e_i, i = 1..6$), and their corresponding eigenvalue ($\lambda_i, i = 1..6$). The eigenvectors coordinates are reported as 0 in bold when the value was below 0.01.

First we observe the classic phenomenon that the higher γ is, the higher the *data-fit term* (which characterizes the distance between simulated and measured values) is (Figure 3 , bottom, blue). A higher γ means that the noise (i.e. uncertainty on the measurements) is considered bigger in Equation 2 so estimated parameters values then tend to be closer to the prior mean than from a value which perfectly fits the measurements.

Second, that in every run of the algorithm, there is one or more eigenvalue which converges to 0 and is under 10^{-4} at convergence. When the value is zero, this means that the parameters lie on a subspace of lower dimension. Here for each run, we report in Table 4 the number of eigenvectors with a eigenvalue above 10^{-4} , the mean value $\bar{S}(\mathcal{X}, \hat{\mathcal{O}})$ of the

²www.sofa-framework.org




FIGURE 3 Top: value of the eigenvalues of the SVD of personalised parameters at each IUP iteration. **Bottom:** mean values $\bar{S}(\mathcal{X}, \hat{\theta})$ and $\bar{R}(\mathcal{X})$ of the *data-fit term* $S(x, \hat{\theta})$ (blue) and the *regularization term* $\gamma R(x)$ (green) across the population at each IUP iteration.

FULL-IUP-05								DIAG-IUP-05							
	σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i		σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i
e_1	0.54	-0.31	-0.26	0.13	0.53	-0.50	3.62	e_1	0.16	-0.18	0	-0.35	0.90	0	3.91
e_2	-0.12	0.16	0.06	-0.74	-0.17	-0.62	2.73	e_2	-0.22	0.23	0	-0.91	-0.27	0	2.96
e_3	-0.45	-0.69	-0.47	-0.25	0.03	0.15	1.34	e_3	-0.92	0.20	0	0.20	0.28	0	1.88
e_4	-0.21	0.63	-0.66	-0.03	0.34	0.08	7.42e-5	e_4	-0.28	-0.93	0	-0.11	-0.18	0	1.11
e_5	-0.49	0	-0.05	0.61	-0.22	-0.58	5.42e-6	e_5	0	0	1.00	0	0	0.04	2.08e-13
e_6	-0.45	0	0.51	-0.05	0.73	0	2.87e-9	e_6	0	0	-0.04	0	0	1.00	3.71e-30

FULL-IUP-01								DIAG-IUP-01							
	σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i		σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i
e_1	0.38	0.09	-0.04	0.02	-0.55	0.74	4.80	e_1	0.12	-0.08	0	0	0.54	0.83	7.15
e_2	0.68	-0.30	-0.41	0.33	0.42	-0.04	3.37	e_2	-0.46	-0.14	0	0	0.76	-0.44	4.44
e_3	0.55	0.17	0.07	-0.64	-0.23	-0.45	2.31	e_3	0.86	-0.24	0	0	0.29	-0.33	3.04
e_4	0.08	0.69	-0.14	0.57	-0.25	-0.34	1.41	e_4	-0.16	-0.96	0	0	-0.23	0.08	1.06
e_5	0.03	0.63	-0.06	-0.29	0.61	0.37	6.65e-5	e_5	0	0	-0.31	-0.95	0	0	3.06e-13
e_6	0.30	0	0.90	0.27	0.19	0.02	3.13e-7	e_6	0	0	-0.95	0.31	0	0	3.64e-17

FULL-IUP-005								DIAG-IUP-005							
	σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i		σ_0	R_0	c_1	P_{ve}	R_p	τ	λ_i
e_1	0.09	0.15	0.37	-0.68	-0.37	-0.49	5.03	e_1	0.13	-0.08	0	0.02	0.50	0.85	7.06
e_2	0.46	0.11	-0.27	-0.06	-0.67	0.50	4.49	e_2	0.51	0.13	0	0	-0.76	0.39	4.69
e_3	-0.72	0.29	0.38	0.11	-0.34	0.36	2.99	e_3	0.84	-0.24	0	-0.01	0.34	-0.35	3.18
e_4	0.31	0.74	0.20	-0.14	0.47	0.27	1.12	e_4	0.16	0.96	0	0	0.23	-0.07	1.06
e_5	0.14	-0.57	0.51	-0.31	0.21	0.50	5.73e-5	e_5	0	0	0	-1	0.01	0.01	0.17
e_6	-0.39	0	-0.59	-0.64	0.20	0.23	1.10e-6	e_6	0	0	-1.00	0	0	0	5.21e-14

TABLE 3 For each run of the IUP algorithm, final eigenvectors of the SVD of personalised parameters (e_i) and their corresponding eigenvalue λ_i . In bold we emphasize the coordinates of eigenvectors which are lower than 10^{-3} and the eigenvalues which are lower than 10^{-4} .

Run	γ	$\lambda \geq 10^{-4}$	$\bar{S}(\mathcal{X}, \hat{\theta})$	$\bar{R}(\mathcal{X})$
FULL-IUP-05	0.5	3	1.85	1.5
FULL-IUP-01	0.1	4	2.15e-2	0.40
FULL-IUP-005	0.05	4	6.18e-3	0.20
DIAG-IUP-05	0.5	4	0.65	2.04
DIAG-IUP-01	0.1	4	0.23	0.41
DIAG-IUP-005	0.05	5	1.96e-2	0.25

TABLE 4 For each run at the final iteration : number of eigenvectors for each run with a eigenvalue above 10^{-4} (column ' $\lambda \geq 10^{-4}$ '), mean *data-fit* value $\bar{S}(\mathcal{X}, \hat{\theta})$ across all cases and mean value $\bar{R}(\mathcal{X})$ of the regularization term.

data-fit term across all cases and the mean value $\bar{R}(\mathcal{X})$ of the regularization term.

From this table, we first observe that the mean of the *regularization term* is very close to γD where D is the number of eigenvectors which do not have a value close to 0 (we use $\lambda \geq 10^{-4}$ as the threshold for *close to 0*), which is consistent with Equation 10 of a cost function with a sparse regulariser on the dimensionality of the set of personalised parameters. We also observe that for the same γ , the data-fit term is lower for the runs with the **Full Matrix** updates than the **Diagonal Matrix** while having a lower number of non-zero eigenvalues.

More interestingly, we can observe the shape of the eigenvectors at convergence of the algorithm. In every run with **Diagonal Matrix** updates, we observe that the smaller the eigenvalue associated to eigenvector is, the more it is aligned to a coordinate (a vector with shape $(\mathbf{0}, \dots, \mathbf{1}, \dots, \mathbf{0})$). This means that for these runs, the parameters associated with these coordinates have a constant values in the final set of personalised parameters. In particular, the contractility c_1 has always a constant value, and either τ (in DIAG-IUP-05) or P_{ve} (in DIAG-IUP-01) is constant as well. On the other hand, there is no such phenomenon in the runs with **Full Matrix**.

The behaviour of the algorithm can be understood from the sparse formulation explicated in 2.5 and the cost functions which is minimised. The algorithm indeed tries to find a prior for which there is an optimal trade-off between the number of dimensions of the set of personalised parameters, and the mean value $\bar{S}(\mathcal{X}, \hat{\theta})$ of the *data-fit term*. Given a specific γ , the "cost" of lowering the dimension is an increase of γ in $\bar{S}(\mathcal{X}, \hat{\theta})$. With the **Diagonal Matrix** updates, this is done with a constraint on the prior covariance matrix to be diagonal so having a lower dimension means fixing the value of a parameter. On the other hand, with the **Full Matrix** updates, the algorithm can find any direction.

In particular we can observe than for $\gamma = 0.1$, both algorithms find a parameter subspace of dimension 4, but the data-fit $\bar{S}(\mathcal{X}, \hat{\theta})$ is higher with **Diagonal Matrix** updates

($0.23 \geq 2.15e-2$), where both c_1 and P_{ve} are fixed. By comparing to the DIAG-IUP-005 where only the c_1 is fixed, we can interpret that the "cost" of fixing P_{ve} is a loss of around 0.23 on $\bar{S}(\mathcal{X}, \hat{\theta})$, which also means that the quality of some personalisations in the database is impacted.

Indeed, for the final personalisation of DIAG-IUP-01 there are at least 3 cases for which the personalisation is highly impacted ($\sqrt{S} \geq 3.5$) because of the fixed value of P_{ve} . These cases have an *aortic diastolic pressure* which is particularly low compared to the rest of the database $DP \leq 5400Pa$. To fit this measurement, P_{ve} (which is the asymptotic and minimal value of the aortic pressure in the blood flow model) needs to be at least below this value (in particular in DIAG-IUP-005 and FULL-IUP-01, all measurements of these cases are almost perfectly fitted and the estimated P_{ve} for these cases is $\leq 5380Pa$), but at this step the prior value of P_{ve} in DIAG-IUP-01 (which is thus the fixed constant value) is $5873Pa$, which makes the fitting of the Diastolic Pressure impossible in these cases.

3.2 | Parameter Existence and Uniqueness in the Resulting Hyperplanes

A classical question in modeling and inverse problems, is to determine which parameters are *observable* with respect to a specific set of measurements. For example, here we estimate 6 parameters from 4 observed outputs. If there is a linear relationship between the parameters and the outputs (i.e there exist a matrix M such as $\mathcal{O}(x) = Mx$), the size of the kernel of the matrix is at least 2. Considering some measurements \mathcal{O} , some parameters x such as such as $\mathcal{O} = Mx$, then for any vector y in the kernel $\mathcal{O} = M(x + y)$ as well. This means that there are at least two orthogonal parameter directions in which the parameter values are *unobservable*. With a non linear-model (such as cardiac models), a similar phenomenon of non-uniqueness exists locally around some personalised parameters, to the extent that the model can be approximated by its gradient. In general, there is an entire manifold of parameters for which the outputs are the same.

Using Gaussian priors on such underconstrained linear models leads a unique solution to the inverse problem, because in this case, the cost function 8 is strictly convex. To a certain extent, this can be locally true for a non-linear model and a unique specific value is promoted within the manifold of parameters for which the outputs are the same.

Finding a (possibly unique) value which is *the most consistent* considering a *prior knowledge* on the distribution of parameters is then the most interesting consideration of the MAP estimation, but this supposes to have an accurate prior. Since we are deriving the prior probability distribution from successive personalisations over the dataset itself, it is thus not

possible to rely on its interpretation as a "prior knowledge" to set parameter values in the unobservable directions of *this dataset*.

On the other hand when no information on the statistics of the parameters is available, the only possibility is to perform a reduction of the parameter space, by forcing some parameter or parameter directions to have a fixed value. This can be done through PCA, PLS²⁰ or sensitivity analysis¹⁶ on the parameters with respect to the clinical data, and estimate coordinates of the first modes only. This also leads to a few questions regarding the reduced space of parameters, particularly on the existence and uniqueness of parameter values for which the simulation fits the target data.

Here our parameters converge into a reduced subspace in all runs of the algorithm. In particular in the FULL-IUP-005, FULL-IUP-01, DIAG-IUP-005 run, there is also a very low data-fit term across the cases. This suggests that among the undetermined directions of a parameter space, **with a small enough prior, the algorithm selects an hyperplane of minimal dimension in which for each case, there is at least a set of parameter values which fits exactly the measurements.**

We demonstrate this claim by analysing the resulting parameter subspaces, in these terms of parameter existence and uniqueness. To that end we perform multiple personalisations without priors (i.e. an uniform prior, which is equivalent to solving Equation 8 without regulariser, or setting $\gamma = 0$) where parameter are taken from within these subspaces. Because the CMA-ES algorithm is stochastic, it usually converges toward different values of the parameters at each run if there are multiple set of parameter values which minimise the cost function. We compare the following parameter subspaces:

1. The complete space \mathcal{H}_0 of 6 parameters.
2. The subspace \mathcal{H}_{c_1} of 5 parameters of all parameters except the stiffness c_1 , which is set to its final (constant) in the DIAG-IUP-005 run.
3. The 5 subspaces $\mathcal{H}_{(c_1, \sigma_0)}$, $\mathcal{H}_{(c_1, R_0)}$, $\mathcal{H}_{(c_1, R_p)}$, $\mathcal{H}_{(c_1, \tau)}$ and $\mathcal{H}_{(c_1, P_{ve})}$ of 4 parameters where both the stiffness and another parameter are set to the final prior mean value in the DIAG-IUP-005 run.
4. The 4 subspaces \mathcal{H}_5 , \mathcal{H}_4 , \mathcal{H}_3 , \mathcal{H}_2 of dimensions respectively 5,4,3 and 2, which are the hyperplanes defined by their center at the prior mean of the FULL-IUP-01, and the l largest eigenvectors of the prior covariance for respectively $l=5,4,3,2$.

We then report both the mean *error of fit* $\sqrt{S(\mathcal{X}, \hat{\theta})}$ across all cases in the database, and the *variability* of estimated parameters across different personalisations, estimated by averaging across all cases the standard deviation of the 5 estimated values in the 5 personalisations.

TABLE 5 Mean error of fit and variability of personalised parameters in 5 personalisations for the parameter subspaces. The value is reported as 0 if it is lower than 10^{-3} .

	$\overline{\sqrt{S}}$	σ_0	R_0	c_1	P_{ve}	R_p	τ
\mathcal{H}_0	0	0.03	0	0.13	0.02	0.03	0.04
\mathcal{H}_{c_1}	0	0	0	-	0.02	0.02	0.03
$\mathcal{H}_{(c_1, \sigma_0)}$	0.99	-	0	-	0.02	0.02	0.14
$\mathcal{H}_{(c_1, R_0)}$	3.90	0	-	-	0.05	0.02	0.04
$\mathcal{H}_{(c_1, P_{ve})}$	0.06	0	0	-	-	0	0.06
$\mathcal{H}_{(c_1, R_p)}$	0.07	0	0	-	0.07	-	0
$\mathcal{H}_{(c_1, \tau)}$	0.01	0	0	-	0.02	0	-
\mathcal{H}_5	0	0	0	0	0.015	0.017	0.03
\mathcal{H}_4	0	0	0	0	0	0	0
\mathcal{H}_3	1.99	0.03	0	0.002	0	0.004	0.01
\mathcal{H}_2	3.03	0	0	0	0	0	0

We first observe that **the mean error of fit is 0 (lower than 10^{-3}) for only four parameter spaces**: the original space \mathcal{H}_0 , the parameter space \mathcal{H}_{c_1} with all the parameters except c_1 , and the two parameter spaces \mathcal{H}_4 and \mathcal{H}_5 with respectively the 4 and 5 largest eigenvectors of the FULL-IUP-01 run. This shows that they are the only subspaces which contain parameter values which fit the measurements for all cases.

Then among these subspaces, we observe that **the only subspace for which both the mean error of fit and the variability of all parameters is 0 (lower than 10^{-3}) is \mathcal{H}_4** , the hyperplane with the 4 largest eigenvectors of FULL-IUP-01. In the other subspaces there is a variability from one personalisation to the other for at least the haemodynamic parameters P_{ve} , R_p and τ showing that the parameters are not unique.

This shows that around the prior which was found in the DIAG-IUP-005 run, c_1 was the only parameter which is possible to set to a constant value without the personalisation of some cases being impacted (such as the 3 cases described in the previous section). However, once c_1 is set, the variability in \mathcal{H}_1 show that there is still an unobservable direction (especially in the haemodynamic parameters), but it is then necessarily a *combination of parameters* which is unobservable.

Finally, because of the possibility to find subspaces which are not necessarily aligned with the coordinates, the FULL-IUP-01 run was able to find a subspace of lower dimension, 4, in which there are parameters fitting all the cases and no variability, thus reducing the space in both directions of uncertainty.

3.3 | Selection of a Parameter Subspace for Personalisation

The resulting parameter subspaces in the previous section exhibit interesting properties in terms of existence and uniqueness of parameters for personalisation, so a question is *how relevant are they in the context of personalisation?* First, it is important to observe that there is no guarantee of uniqueness of the parameter directions which are selected by the algorithm (and the personalised parameter either). Indeed, in addition of the model being non-linear (which makes the data-fit term not convex, thus the whole Equation 8), sparse regularisations are usually not convex and have many possible solutions. In particular here, there are possibly multiple parameter subspaces on which the algorithm could converge. Visually, this can be observed in the schema of Figure 2, where two different subspaces can be selected, each containing a unique set of parameter values for each of the 5 cases.

Secondly, there are no guarantees that the parameter spaces selected by the algorithm are relevant from a modeling or physiological point of view. In particular, when two parameters are not completely observable from some measurements in a database, their 'actual' value (if they correspond to physical parameters) in a population is likely variable. In addition, a drawback of fixing a parameter to a specific value, is that it might force other estimated parameters to vary more within the population to account for variations which would have come from the fixed parameter.

Parameter selection in the sense of setting a specific parameter direction to a constant value is then an imperfect approach in modeling, but it is unavoidable to not have variability in the estimation when no other information or statistics is available on this specific parameter direction.

In this context, though the selection of relevant parameters for personalisation cannot be performed entirely automatically (because the physical meaning of parameters is ignored by the algorithms), we believe that our algorithm can help the modeler by revealing unobservable parameter and parameter directions. To that end, we recommend the following general approach, considering a set of clinical data for which observable parameters are unclear:

1. First, perform one (or multiple) runs of IUP with **Diagonal Matrix** updates. This is because the resulting subspace is easier to analyze, since the parameters usually have a physical meaning in the context of modeling. Then analyze the parameters which end up with a eigenvalue close to 0, and set them to a constant value if it is not incompatible with physiological considerations.
2. Second, perform one (or multiple) IUP with **Full Matrix** updates to further select a lower dimensional subspace.

TABLE 6 Final eigenvectors and eigenvalues of the IUP run with **Full Matrix** and $\gamma = 0.05$ on \mathcal{H}_5 .

Final-IUP-005						
	σ_0	R_0	P_{ve}	R_p	τ	λ_i
e_1	0.58	0.07	0.05	-0.51	0.63	4.77467
e_2	0.46	-0.32	0.50	0.65	0.10	3.41794
e_3	0.65	-0.06	-0.44	-0.06	-0.61	2.95973
e_4	0.16	0.94	0.15	0.25	-0.05	1.03746
e_5	0.01	-0.01	0.73	-0.50	-0.47	4.30e-6

TABLE 7 Mean error of fit and variability of personalised parameters in 5 personalisations in \mathcal{H}^* . The value is reported as 0 if it is lower than 10^{-3} .

	\sqrt{S}	σ_0	R_0	P_{ve}	R_p	τ
\mathcal{H}^*	0	0	0	0	0	0

We apply this approach to the current problem. First it seems physically likely that we do not have so much information on the stiffness c_1 from only the 4 measurements in Table 1, so we can reasonably decide to set its value to a constant, which we set at its value in `DIAG-IUP-005`. Then we perform a new run of the IUP algorithm, which we call `Final-IUP-005`, with **Full Matrix** and $\gamma = 0.05$ on the resulting parameters (thus in the hyperplane \mathcal{H}_5). This leads to the final eigenvectors and eigenvalues:

We then select the hyperplane \mathcal{H}^* made of the 4 largest eigenvectors (e_1, e_2, e_3, e_4) at the end of the run, and test the existence and uniqueness of personalised parameters in this hyperplane, with the same method than in the previous section. Results are reported in Table 7.

As expected, both the variability of parameters and the mean error of fit are 0 (lower than 10^{-3}). This means that for each case, the hyperplane contains a unique (to the extent that we can evaluate it through this algorithm) set of parameter values for which the simulation fits the measurements.

To conclude, we found a **minimal parameter subspace of dimension 4, based on 5 parameters**, consistent from a physiological point of view, in which to perform consistent parameter estimation. More importantly, the final parameters μ^* and Δ^* of the prior at the end of the `Final-IUP-005` run respectively correspond to the mean and covariance of the estimated parameters, which we then call the *population-based priors*.

Finally, we point out that the resulting population-based priors will not necessarily correspond to the "real" underlying

parameter distribution. Indeed, considering a population with a known parameter distribution, the priors obtained through the IUP method will correspond to this distribution only if the data enables to completely determine the distribution parameters (in which case the influence of the prior will be minimal). However, if some parameters or parameter space directions are unobservable, the algorithm will choose arbitrary directions within the space, and build priors on these directions (which will result in dimensionality reduction to a minimal subspace). It is then up to the modeller to check that the selected subspace is relevant, possibly add additional output values to personalised or, as we will see in the next section, enhance the dataset with other elements which can help determine the remaining unobservable directions.

4 | CONSISTENT PARAMETER ESTIMATION IN A DATABASE WITH MISSING OR HETEROGENEOUS MEASUREMENTS

In this section we present the application of the proposed framework to the **consistent personalisation of a large database of cases with heterogeneous (or missing) measurements**. We present the personalisation of all the cases in such database with the **IUP*** algorithm which leads to the following properties of the personalised simulations:

1. All the measurements of all cases are well fitted in the corresponding personalised simulations.
2. Parameters lie on a reduced subspace of minimal and sufficient dimension.
3. Parameters for cases where measurements are missing are constrained by the *population-based* priors in this subspace which means in particular that unobservable parameters for these cases are guided by their values in the other cases of the database where they are available.

In addition we also show in this section how the proposed framework can be extended to integrate *external parameters* such as the *height* and the *weight* of the patient in order to guide the estimation of parameters for cases where measurements are missing, leading to an improved consistency of parameter estimation for these cases.

4.1 | Heterogeneous Database of Clinical Cases

We consider a larger database of 811 cases from different studies, hospitals and protocols. Depending on the protocol, the same measurements are not available for all the patients and even within a single study with the same protocol, some measurements can be missing. To build this database, we focused

Measurements \mathcal{O}

Maximal Volume $MaxV$

Minimal Volume $MinV$

Stroke Volume SV

Ejection Fraction EF

Mean Aortic Pressure MP

Diastolic Aortic Pressure DP

TABLE 8 Measurements considered in the heterogeneous database

on gathering patients and acquisitions for which the heart rate and *at least one of the 6 following measurements in Table 8* was available:

Within this database of 811 cases, we have the following statistics:

- Pressure Measurements are available for 651 cases only.
- The Maximal Volume and Minimal Volume are both available for only 340 cases.
- Among the 471 other cases, either the Ejection Fraction (63 cases), Stroke Volume (386 cases) or no volume measurement at all (21 cases) are available.
- Ejection Fraction is the only measurement available in 38 cases.
- Stroke Volume is the only measurement available in 45 cases.
- 258 cases have a 'complete' set of 4 measurements: Maximal Volume, Minimal Volume, Mean Aortic Pressure and Mean Diastolic Pressure (we do not report Ejection Fraction and Stroke Volume if the Maximal Volume and Minimal Volume are already reported).

In order to accommodate the heterogeneous nature of the database, we use a *heterogeneous data-fit term*. Instead of using the fixed formulation $S(x, \hat{O}) = \|(O(x) - \hat{O}) \oslash \mathcal{N}\|^2$ with O being the vector of outputs ($MaxV$, $MinV$, MP , DP) and \hat{O} the corresponding measurements, we build a different vector of observations for each patient. Depending on the available measurements, we use in the order of priority:

1. ($MaxV$, $MinV$, MP , DP), 7. ($MaxV$, $MinV$),
2. ($MaxV$, SV , MP , DP), 8. ($MaxV$, SV),
3. ($MaxV$, EF , MP , DP), 9. ($MaxV$, EF),
4. ($MinV$, MP , DP), 10. ($MinV$),
5. (SV , MP , DP), 11. (SV),
6. (EF , MP , DP), 12. (EF),

and a corresponding normalisation vector \mathcal{N} . The normalisation coefficient for the Ejection Fraction (EF) (which is a percentage) is **5%**, it is **10 ml** for the *volume* values (MaxV,MinV,SV) and **200Pa** for the *pressure* values (MP, SP and DP).

As we explained in Section 3.1, the IUP algorithm with a constant noise as used in Section 3 realises a trade-off, driven by the constant γ , between the number of dimension of the set of personalised parameters and the mean *goodness of fit* over the whole database. In particular by using a small γ , a high goodness of fit was reached for all cases, leading to the selection of a *minimal and sufficient* subspace in where for each case of the database there is a set of parameter values for which the simulation fits exactly the measurements.

However on this heterogeneous database, selecting a small γ does not lead to a high goodness of fit for all cases. As we show in the next section (Section 4.2), an IUP run on the database of 811 patients with the heterogeneous data-fit term and $\gamma = 0.05$ (and even as small as 0.02) leads to a parameter space of 3 directions instead of 4, and many *complete cases* (with 4 measurements available) are not well fitted. This is because there are cases in the database with less than 4 measurements which can be well fitted with parameters in a parameter space of dimension 3 only. Consequently the contribution to the *mean error of fit* of badly fitted cases is small, which leads the algorithm to remove a parameter direction which is necessary to fit other cases. A possibility would be to use an even lower γ , but this leads to practical numerical problems during the MAP estimation, because the value of the regularization term becomes too small in Equation 8 compared to the data-fit term. Instead, we use in the next section the **IUP*** version with reweighting of the data-fit term in the of the algorithm presented in Section 2.6.

4.2 | Comparison between the IUP and IUP* Algorithms on the Database

Here we demonstrate differences in the behaviour of the **IUP algorithm** and the extended **IUP* algorithm** on this database. We compare two runs of these algorithms, respectively called FULL-IUP-05 and FULL-IUP*-05, based on the 5 parameters of \mathcal{H}_{c_1} (the stiffness is set to a constant value), a value of $\gamma = 0.05$ and the **Full Matrix** updates. We first report as in Section 3.1 the number of eigenvectors with a eigenvalue above 10^{-4} at convergence, the mean *data-fit* value $\bar{S}(\mathcal{X}, \hat{\theta})$ across all cases and mean value $\bar{R}(\mathcal{X})$ of the regularization term, for each run in Table 9 .


We observe that the IUP run leads to 3 eigenvectors with an eigenvalue above 10^{-4} while the IUP* run leads to 4. The mean data-fit term value $\bar{S}(\mathcal{X}, \hat{\theta})$ for the IUP* run is very low (1.35e-3) and the IUP run is also low (1.01e-2). However we

Run	$\lambda \geq 10^{-4}$	$\bar{S}(\mathcal{X}, \hat{\theta})$	$\bar{R}(\mathcal{X})$
FULL-IUP-05	3	1.01e-2	0.15
FULL-IUP*-05	4	1.35e-3	0.19

TABLE 9 For each run at convergence: Number of eigenvectors for each run with a eigenvalue above 10^{-4} (column ' $\lambda \geq 10^{-4}$ '), mean *data-fit* value $\bar{S}(\mathcal{X}, \hat{\theta})$ across all cases and mean value $\bar{R}(\mathcal{X})$ of the regularization term.

now show a boxplot which compares the repartition of the values of the data-fit term *error of fit* $\sqrt{S(x_i, \hat{\theta}_i)}$ in the population in Figure 4 :

FIGURE 4 Comparison of the values of the error of fit $\sqrt{S(x_i, \hat{\theta}_i)}$ in the FULL-IUP-05 and FULL-IUP*-05 runs.



As we can see with the IUP run, there are many cases (almost one fifth of the population actually) which are not well personalised ($\sqrt{S(x, \hat{\theta})} \geq 0.1$). On the other hand in the IUP* run, by design of the algorithms all the cases have a data-fit term below 0.01.

Finally in Table 10 we also report, for the various possible values of the weight of the data-fit term α (in Equation 11) at convergence, the percentage of cases in the population with these values. As we can see, this term doubled up to 4 times in some cases (thus leading to $\alpha = 16$).

Overall, this section shows that the extended version **IUP*** enables a personalisation of the complete database where **all**

α	1	2	4	8	16
%	81.79	0.49	1.47	9.85	6.4

TABLE 10 Percentage of cases in the population with the corresponding final value of α in Equation 11 at convergence of the FULL-IUP*-05 run.

the individual cases have their measurements fitted within a specific value of the *data-fit term*, while parameters are gathered on a **parameter subspace of minimal and sufficient dimension**.

4.3 | Parameter Estimation in cases of Unobservability

Here we demonstrate the impact of the population-based priors computed through the **IUP* algorithm** in this database on cases with unobservable parameters. To that end, we compare the estimated parameter values for cases where measurements are missing in two different parameter estimations for this database: first, the estimation of the 5 parameters *without priors* and second, the estimation at convergence of the FULL-IUP*-05 run of the **IUP* algorithm** over the complete database and the 5 parameters.

The most interesting set of values is the following: in Figure 5, we display the (log-)values of estimated parameters as a function of Ejection Fraction (resp Stroke Volume) for the cases where only the Ejection Fraction (resp Stroke Volume) is available. We observe a classic phenomenon with priors: personalised values not only have less variance with the use of priors (blue points), but because one measurement has to be fitted, they also lie onto a space of (local) dimension 1. Indeed, this space is defined (in the case of the ejection fraction) by the equation

$$x(EF) = \operatorname{argmin}_x \gamma R(x) \text{ s.t. } S(x, \hat{\mathcal{O}}) = 0 \quad (19)$$

It is a space of local dimension 2 (resp 3) when 2 (resp 3) measurements have to be fitted (to the extent that the model is locally approximable to its gradient).

On an interpretative level, when some measurements are missing in our database and many parameter values are possible, the personalisation with priors in \mathcal{H}^* leads to the selection of a set of parameter values which maximises its likelihood in the probability distribution of the priors, or equivalently minimizes the distance $(x_i - \mu^*)^T (\Delta^*)^{-1} (x_i - \mu^*)$. To that extent, it performs a form of imputation in the parameter space by choosing the *most likely* set of parameters according to the distribution defined by μ^* and Δ^* . Since μ^* and Δ^* are the statistics of the whole population of personalised parameter values, we argue that values of unobservable parameters and

parameter directions in some cases are then largely determined by their values in the other cases of the database where they are observable.

4.4 | Integration of External Parameters in the Prior Distribution for Improved Estimation of Unobserved Parameters.

The key idea behind the use of priors is to model the distribution of parameters in the population. Then with the MAP estimation, the goal is to find the *most likely parameters* according to the correlations in this distribution. Here we explore the possibility of integrating parameters which are not estimated into the prior distribution (we call them *external parameters*), to influence the value of estimated parameters.

Namely, for all our 811 cases, the *height* and *weight* of the patients were available. We can also consider the *heart rate* which is not an estimated parameter. In order to add these three parameters to the prior distribution, we perform the following modifications to the method: in Equation 8, instead of considering a vector x which only contains the parameters to be estimated, we use a *concatenation vector* of dimension 9 which contains the 6 estimated parameters and the 3 *external parameters*. Formulations of the prior covariance and mean in Equations 2 and in the equations of Section 2.3 are adapted as well to accommodate this *concatenation vector*.

We then perform the following simple experiment: from the estimated parameters in Section 4.3, we estimate the covariance matrix and the mean of the *concatenation vector*, which we then use as a prior for a new estimation. This is equivalent to perform only one iteration of the IUP algorithm with **Full Matrix** with this *concatenation vector*.

We report the most interesting result here in Figure 6. For both the estimation *without external parameters* **E1** (blue points) and *with external parameters* **E2** (red points), we display the estimated values of the *resting radius* R_0 for the cases where the Ejection Fraction is the only measurement, as a function of **a/** the *height* of the patient (left), **b/** the *weight* of the patient (middle) and **c/** the *Ejection Fraction* of the patient.

The results are the following: first in both estimations the goodness of fit for the cases in the database is similar and high (a mean *error of fit* of around 0.06). For the estimation without external parameters **E1**, the resting radius is well correlated to the Ejection Fraction (this was already observed in Figure 5), but **not at all to the weight and the height**. However with the external parameters **E2**, the values of the resting radius is **very correlated to the height and weight of the patients**, and tends to increase with both these measurements. Since the resting radius, from a physical point of view, is related to the *size* of the heart, this correlation makes sense from a physical point

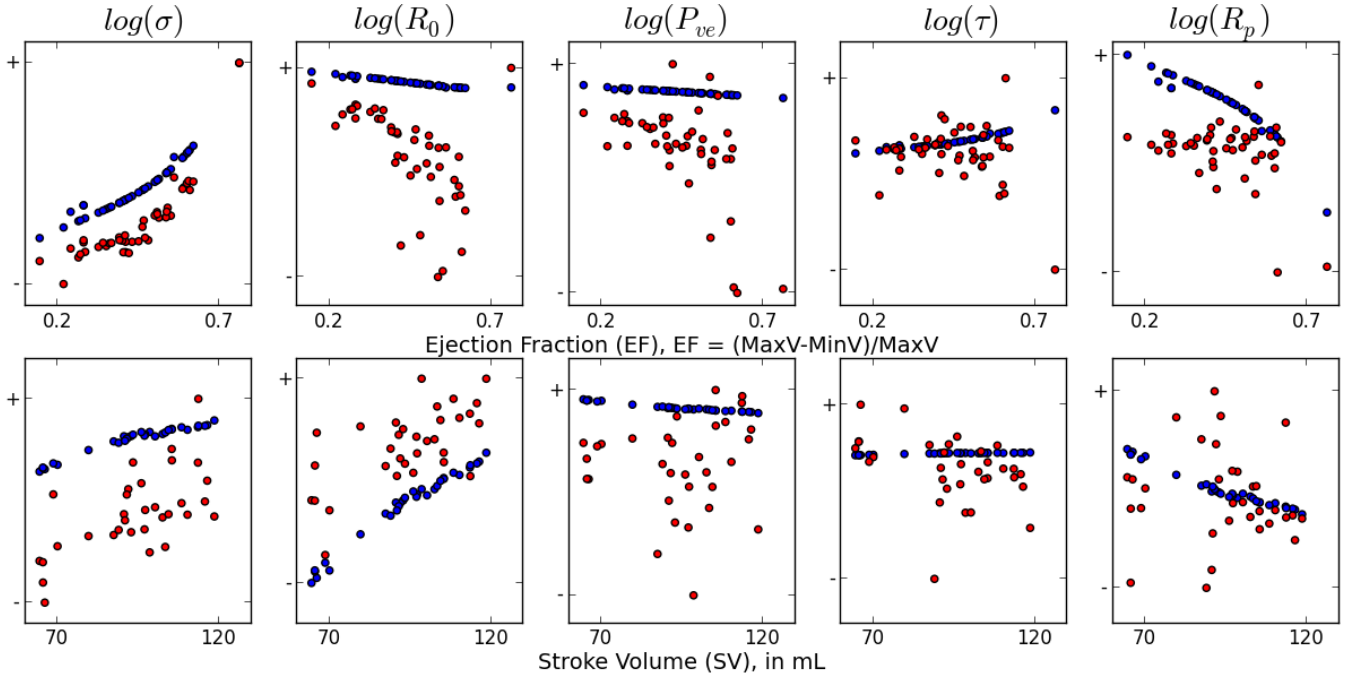


FIGURE 5 Top: (resp **Bottom:**) Log-value of estimated parameters value as a function of the Ejection Fraction (resp Stroke volume) for the 38 (resp 45) cases where only the Ejection Fraction (resp Stroke Volume) is available, in blue with the **IUP*** algorithm, in red without priors.




FIGURE 6 Estimated values of the *resting radius* R_0 in the estimation *without external parameters* (blue points) and *with external parameters* (red points) in the prior distribution, for cases where only the Ejection Fraction was available (the resting radius was thus not completely observable).

of view and leads to an improved consistency of the estimated values in **E2** than in **E1**.

5 | CONCLUSION

In this manuscript we introduced a method called **Iteratively Updated Priors** which performs successive personalisations of the cases in a population, where the prior probabilities on parameter values at each iteration are set from the statistics

of personalised parameters in the previous iteration. We presented two different updates of prior parameters depending on the assumptions made on the distribution of parameters in the population, then derived a mathematical link between this algorithm and the minimisation of a *population-wide cost function* with a *sparse* penalty on the *dimension of the set of personalised parameters*. Depending on the type of covariance matrix update, this leads either to the selection of *parameters* which are set to a constant value in the population, or *parameter directions* which are set constant.

In Section 3 we demonstrated that the algorithm leads to a trade-off between the mean goodness of fit of personalised simulations over the database, and the dimension of the set of personalised parameters, and showed that this can be used to select a consistent subspace of *minimal and sufficient dimension* in which for each case of the population, there is a unique set of parameter values for which the simulation fits the corresponding measurements.

In Section 4, we presented an extended version of this algorithm applied to the personalisation of a larger database of 811 cases with heterogeneous measurements. This algorithm minimises the number of dimensions of the set of personalised parameters while ensuring a high (almost exact) goodness of fit over the whole database. We showed that the resulting *population-based priors* computed through this algorithm simultaneously leads to consistent parameter estimation for the cases where the parameters are not all observable, and are in particular influenced by their values in cases where they are observable. Finally, we show it is possible to add *external parameters* such as the *height* and the *weight* of the patients in the prior probability distribution in order to further guide the estimation of a consistent value of parameters which are unobservables in some cases.

A first extension is that beside the two types of updates presented here (**Full Matrix** and **Diagonal Matrix**), many other assumptions on the probability distribution of parameters in the population could be used, resulting in different formulations of the prior parameters. For example, the assumption could be made that the contractility σ_0 is independent of the other parameters, resulting in a block matrix covariance. Similarly, if the mean of a parameter value is known, then it is also easy to change the updates to accommodate this case (the derivations of the maximum likelihood updates for parameters of Gaussian priors are easily tractable even with such constraints). Finally, other type of prior shapes (other convex functions, mixture of gaussians...) could be used and the IUP algorithm would lead a dimensionality reduction of the parameter space and the creation of population-based statistics along the degrees of the freedom of the prior function (in our example it was the along the principal directions of the covariance matrix of the Gaussian prior).

Secondly, a last extension of this work would be to apply the IUP to a 3D cardiac model. However in the case of our 3D model¹⁶ whose simulations can be up to 1 hour of computation, the computational burden associated with performing a dozen or more repeated personalisations of the whole database can be too high. To tackle this and the 3D model, we believe a possibility would be to use the personalisation in a reduced subspace of 0D model parameters (such as \mathcal{H}^*) to influence the estimation of a 3D model. Another possibility could be to investigate the use of specific 0D / 3D *multifidelity parameter couplings* as presented in¹⁸ for multiple iterations of the IUP algorithm at once to lower the burden of repeated personalisations. We believe this could enable the use of the IUP algorithm on databases of 100 to 200 cases with the 3D model in around 4-5 days.

We conclude with a general comment on priors and the relevance of parameter selection in the context of personalisation. In the ideal case where we have prior statistics on all the parameters, every parameter which has an influence on a specific measurement will likely vary from the mean during parameter estimation, in accordance to the probability defined by the prior. This is why all the parameters vary at least slightly from the mean when only the ejection fraction or the stroke volume are available in our example. However if no statistics or external information is available on a set of parameters, a specific set of values of these parameters does not make *more sense* than another one that also fits the measurements. There is thus no other choice than arbitrarily choosing a direction, until more information can be given on this parameter direction from more experimental or physiological knowledge. Within these considerations, the best approach consists in our opinion in starting with a high number of parameters, and iteratively reducing the parameters space until a sufficient subspace is found in which all directions can be observed. We argue that our *Iteratively Update Prior* method, though it absolutely requires supervision for the selection of parameters with physiological relevance, gives a complete framework for consistent parameter estimation through the joint selection of a parameter subspace of reduced dimension and the creation of relevant priors on its directions.

Compliance with Ethical Standards : Informed consent was obtained from the subjects and the protocol was approved by the local Research Ethics Committee.

Funding: The research leading to these results has received European funding from the FP7 project MD-Paedigree (600932), the ERC starting grant ECSTATIC (715093) and the ERC advanced grant MedYMA (2011-291080) and was supported by the grant AAP Santé 06 2017-260 DGA-DSH and by the Inria Sophia Antipolis - Méditerranée NEF computation cluster.

Conflict of Interest: The authors declare that they have no conflict of interest.

6 | APPENDIX A: LINK BETWEEN THE IUP ALGORITHM AND THE ITERATIVELY REWEIGHTED LEAST SQUARE (IRLS) ALGORITHM

Here we demonstrate a link between the IUP algorithm and the *Iteratively Reweighted Least Square* in sparse regression. In particular, we exhibit a population-wide cost function with *sparse regularisation*. The sparse penalty correspond to the rank of the covariance matrix of the estimated set of personalised parameters, which has two different interpretation depending on the covariance matrix update.

For a given iteration k of the IUP algorithm, we note $\mathcal{X}_k = (x_1, \dots, x_n)$ the concatenation of some vectors of parameter values $x_i, i = 1..n$ for all the n cases, and $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$. We define the **population-wide cost function** $\mathcal{S}(\mathcal{X}_k, \hat{\theta})$ as the mean of all the individual cost functions for all the $i = 1..n$ cases:

$$\begin{aligned}\mathcal{S}(\mathcal{X}_k, \hat{\theta}) &= \frac{1}{n} \sum_{i=1}^n \hat{S}(x_i, \hat{\theta}_i, \mu^k, \Delta^k), \\ \mathcal{S}(\mathcal{X}_k, \hat{\theta}) &= \frac{1}{n} \left(\sum_{i=1}^n S(x_i, \hat{\theta}_i) + \gamma \sum_{i=1}^n R_k(x_i) \right).\end{aligned}$$

The first sum of this expression is the sum of the data-fit terms over all the cases and is the same expression at each iteration. The second sum has a different expression at each step depending on the updates which are performed, which impact the behavior of the algorithm. We explicit here these expressions and explain the resulting behaviour of the IUP algorithm with these updates.

"Diagonal Matrix" updates.

At each iteration k we have:

$$\gamma \sum_{i=1}^n R_k(x_i) = \gamma \sum_{i=1}^n (x_i - \mu^k)^T (\Delta^k)^{-1} (x_i - \mu^k),$$

where Δ^k is a diagonal matrix computed from the personalised parameters X_i^{k-1} at iteration $k-1$ with Equation 7. We can break down the sum along the $j = 1..N$ coordinates (corresponding to the N personalised parameters): we note Δ_{jj}^k the j -th component of the diagonal of Δ^k which formula is:

$$\Delta_{jj}^k = \frac{1}{n} \sum_{i=1}^n ((X_i^{k-1})_j - (\mu^k)_j)^2.$$

We then have:

$$\begin{aligned}\gamma \sum_{i=1}^n R_k(x_i) &= n\gamma \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^N \frac{((x_i)_j - (\mu^k)_j)^2}{\Delta_{jj}^k}, \\ \gamma \sum_{i=1}^n R_k(x_i) &= n\gamma \sum_{j=1}^N \frac{\frac{1}{n} \sum_{i=1}^n ((x_i)_j - (\mu^k)_j)^2}{\Delta_{jj}^k}.\end{aligned}$$

To understand the behaviour of our algorithm with these updates, we can look at the similarity between our updates and a classic method called *Iteratively Reweighted Least Square*, used to solve the minimisation of cost functions involving L_p norms such as

$$X^* = \operatorname{argmin}_X \|Y - f(X)\|^p = \operatorname{argmin}_X \sum_i (Y - f(X))_i^p. \quad (20)$$

The IRLS algorithm works by successively optimising the following (and usually easier) "weighted L_2 " problems where in the k -th problem, each of the l coordinates are reweighted by a vector W_l^k derived from the norm of the coordinate in the solution of the previous problem:

$$\begin{cases} (X^*)^k = \operatorname{argmin}_X \sum_l W_l^k (Y_l - f_l(X))^p, \\ \text{where } W_l^k = |Y_l - f_l((X^*)^{k-1})|^{p-2}. \end{cases} \quad (21)$$

In this context, the **Diagonal Matrix** updates in our algorithm are reweighting the penalty of each coordinate (or personalised parameters) j with the weight $(\Delta_{jj}^k)^{-1}$ which is inverse of the variance of the parameter in the current personalisation. To understand how the IRLS method applies in our case, we define the random vector D as the difference $(X_i^k - \mu^k)$ of the personalised parameters X_i^k to the mean of the population μ^k , for which all the cases are samples. The variance of each personalised parameter is then the (L_2) norm of the coordinate of this random vector. In this formulation, each update correspond to the IRLS formulation with $p=0$ for the minimisation of the L_0 norm of D .

As a consequence, if the successive μ^k and x_i^k converge, according to the IRLS method the value of regularization term converges to:

$$\gamma \sum_{i=1}^n R_k(x_i^*) = n\gamma \|D\|_0,$$

where $\|D\|_0$ is the number of coordinates of D which are non-zero, which is also the number of personalised parameters which do not have a unique value in the population.

"Full Matrix" updates.

At each iteration k we have:

$$\gamma \sum_{i=1}^n R_k(x_i) = \gamma \sum_{i=1}^n (x_i - \mu^k)^T (\Delta^k)^{-1} (x_i - \mu^k).$$

As Δ^k is the covariance matrix of the X^{k-1} it can be orthogonally diagonalised and expressed as $\Delta^k = (O^k)(H^k)(O^k)^T$ with O^k an orthogonal matrix which columns are the principal directions of the set of X^{k-1} , H^k diagonal which diagonal coefficients are the variances of the X^{k-1} in these principal directions. We can write:

$$\gamma \sum_{i=1}^n R_k(x_i) = \gamma \sum_{i=1}^n (O^k(x_i - \mu^k))^T (H^k)^{-1} O^k(x_i - \mu^k),$$

and in this case, each update consists in an IRLS step in the basis b_k , defined by the principal directions of the set of X_i^{k-1} (the columns of O^k) in which the covariance matrix Δ^k is diagonal.

Each of these updates thus correspond to an optimisation step of the L_0 norm of the random vector D^k of the difference $O^k(X_i^k - \mu^k)$ of the personalised parameters X_i^k to the mean of the population μ^k , expressed in the basis b_k .

As a consequence, if the successive μ^k and x_i^k respectively converge to μ^* and x_i^* , and we can build a random vector D^* on a basis b^* to which the D^k and b_k respectively converge, the regularization term converges to:

$$\gamma \sum_{i=1}^n R_k(x_i) = n\gamma \|D^*\|_0,$$

where $\|D^*\|_0$ is the number of coordinates of D^* which are non-zero expressed in the basis b^* (made of principal directions of the set of x^*). It is the rank of \mathcal{X} or more simply the number of principal directions in which parameters do not have a unique value in the population.

7 | APPENDIX B: THE 0D CARDIAC MODEL

The 0D cardiac model is a reduced version of our in-house 3D electromechanical model^{17,16}, derived with the method presented by Caruel et al.³. In this section we first present the mechanical model itself, then we give the equations of the 0D model.

7.1 | The Mechanical Model

The electromechanical equations are based on the Bestel-Clement-Sorine model (BCS) of sarcomere contraction for 3D models as extended by Chapelle et al.⁵, in conjunction with a Mooney-Rivlin energy for the passive hyperelasticity. Hemodynamics are represented through global values of pressures and flows in the cardiac chambers, and coupled to the mechanical equations with the Windkessel model of blood pressure for the after-load (aortic pressure).

7.1.1 | The BCS model: Active Contraction and Passive Material

The BCS model describes the sarcomere forces as the sum of an active contraction force in the direction of the fibre, in parallel with a passive isotropic visco-hyperelastic component. It is compatible with the laws of thermodynamics, and allows to model physiological phenomena at the sarcomere scale which translate at the macroscopic scale (such as the Starling Effect).

The active force in the sarcomere is modeled by the filament model of Huxley¹², which describes the binding/unbinding process of the actin and myosin in the sarcomere at the nanoscopic scale. At the mesoscopic scale, it results³ in a differential equation which relates the active stress τ_c , the stiffness k_c and the strain e_c of the filament within the sarcomere:

$$\begin{cases} \dot{k}_c = -(|u|_+ + |u| + \alpha|e_c|)k_c + k_0|u|_+, \\ \dot{\tau}_c = -(|u|_+ + |u| + \alpha|e_c|)\tau_c + e_c k_c + \sigma_0|u|_+, \end{cases} \quad (22)$$

where α is a constant related to the cross-bridge destruction during contraction, k_0 and σ_0 are respectively the maximum stiffness and contraction. The values of $|u|_+$ and $|u|_-$ are respectively the rate of build-up k_{ATP} and decrease k_{RS} of the force during contraction and relaxation, which depends on the depolarisation and repolarisation times T_d and T_r of the sarcomere:

$$u = \begin{cases} k_{ATP} & \text{when } T_d \leq t \leq T_r \\ -k_{RS} & \text{otherwise} \\ |u|_+ = \max(u, 0), \\ |u|_- = -\min(u, 0). \end{cases} \quad (23)$$

This active force is applied in the direction of the fibre through the visco-elastic component, made of a spring E_s and a dissipative term μ . As derived in Caruel et al.³, the resulting stress σ_{1D} in the fibre direction is given by:

$$\begin{cases} \sigma_{1D} = E_s \frac{e_{1D} - e_c}{(1 + 2e_c)^2}, \\ (\tau_c + \mu \dot{e}_c) = E_s \frac{(e_{1D} - e_c)(1 + 2e_{1D})}{(1 + 2e_c)^3}, \end{cases} \quad (24)$$

where $e_{1D} = \underline{\tau_1} \cdot \underline{e} \cdot \underline{\tau_1}$ is the strain in the fibre direction τ_1 (\underline{e} is the Green-Lagrange strain tensor).

Finally for the passive component the isotropic Mooney Rivlin model of hyper-elastic material is used, driven by the following strain energy:

$$W_e = c_1(I_1 - 3) + c_2(I_2 - 3) + \frac{K}{2}(J - 1)^2, \quad (25)$$

where I_1 , I_2 and J are the invariants of the Cauchy-Green deformation tensor, c_1 , c_2 and K are the parameters of the material.

7.1.2 | Haemodynamic Coupling

To model the influence of blood dynamics during the cardiac cycle, the mechanical equations are coupled with a basic circulation model implementing the 4 phases of the cardiac cycle. For a given ventricle, if we note P_{at} the pressure in the atrium, P_{ar} the pressure in the artery and P_V the pressure in the ventricle, the phases are the following:

- *Diastolic Filling*: when $P_V \leq P_{at}$, the atrial valve is open and the ventricle fills up with blood.
- *Isovolumetric contraction*: when contraction starts, P_V rises. $P_{at} \leq P_V \leq P_{ar}$ and all the valves are closed.
- *Systolic Ejection*: when $P_V \geq P_{ar}$, the arterial valve opens and the blood is ejected into the artery.
- *Isovolumetric relaxation*: when the contractile forces disappear, P_V finally decreases. $P_{at} \leq P_V \leq P_{ar}$ again and all the valves are closed.

We use the haemodynamic model introduced by Chapelle et al.⁵ which links the blood flow q to the ventricular, atrial and arterial pressures with the following equations:

$$q = \begin{cases} K_{at}(P_V - P_{at}) & \text{for } P_V \leq P_{at} \\ K_{iso}(P_V - P_{at}) & \text{for } P_{at} \leq P_V \leq P_{ar} \\ K_{ar}(P_V - P_{at}) + K_{iso}(P_{ar} - P_{at}) & \text{for } P_V \geq P_{ar} \end{cases} \quad (26)$$

Here the atrial pressure $P_{at}(t)$ (cardiac preload) is imposed at a constant value P_{at_lower} except for a *pressure bump* up to P_{at_upper} at the beginning of cardiac cycle, to account for the contraction of the atrium before the ventricular contraction. Finally the pressure of the artery P_{ar} (cardiac afterload) is modeled with the 3-parameters Windkessel model²³ and coupled to the ventricular outflow q through the equation:

$$R_p C \dot{P}_{ar} + P_{ar} - P_{ve} = (R_p + Z_c)q + R_p Z_c C \dot{q}, \quad (27)$$

where R_p is the *Peripheral resistance*, Z_c is the *Characteristic impedance*, C is the *Arterial compliance* and P_{ve} is the *Central Venous Pressure*.

7.2 | The equations of the 0D model

We derived the equations of the *0D model* of the heart from our 3D model with the method of Caruel et al.³. With this approach, both the 0D and the 3D model rely on the same BCS equations of myocardial forces, but the following simplifying assumptions are made on the geometry, the electrical activation and the properties of the material from the 0D model:

1. The ventricle has a spherical shape.

2. The material is incompressible.

3. The electrical activity is synchronous and homogeneous over the sphere.

With these assumptions, myocardial forces and motion are spherically symmetric and can be entirely described by the inner radius r of the ventricle. Deformation and stress tensors can also be reduced to a simple form³, which leads to a system of a dozen equations which we report in Table 11 and Table 12 and detail in Section 7.2.1 and 7.2.2.

The equations are implemented into C code and the system of equations is solved using an explicit *Forward Euler* method with a temporal discretisation of 0.01 milliseconds. This leads to the simulation of 15 beats per second. Finally, the 0D model was also encoded in the CellML format⁷, an open standard based on the XML markup language to store and exchange computer-based mathematical models. It can be downloaded from the Physiome Model Repository³ and easily exploited through the software OpenCOR¹⁰.

7.2.1 | Mechanical Equations

The list of simplified equations of our 0D model is reported in Table 11. Equations (a), (b), (c) and (f) are the same sarcomere and visco-elastic equations than Equations 22 & 24, which are calculated once for the whole sphere. C in equations (d), (e), (g) and (h) denotes a component of the simplified Cauchy-Green deformation tensor which depends only on $y = R - R_0$. $\sigma_{passive}$ in equation (g) is the stress due to the passive law and $\sigma_{viscosity}$ in equation (h) is the stress due to an additional viscous damping η , both expressed as a simple function of C (see³ for the full derivations). In equation (i), Σ_{sph} is the sum of all the stresses applied to the sphere. Equation (j) is the resulting equation of motion which, coupled with the haemodynamic model (k) and the windkessel equation (l), gives the full system of 3 equations to be solved at each iteration.

7.2.2 | Electrophysiology Equations

Assuming synchronous and homogeneous electrical activation (and thus sarcomere force) means that all of the ventricle is depolarised simultaneously. This leads to a rate of ventricular pressure rise during the isovolumetric contraction (resp. isovolumetric relaxation) which is very close to the rate of build-up k_{ATP} (resp. decrease k_{RS}) of the active stress τ_c . However in 3D, this rate is also very dependent on the time for the ventricle to be fully depolarised, which is roughly the QRS duration.

In order to correct this discrepancy between the models, we adapted the electrical parameter u to take into account the QRS duration. We model the fraction f_{depo} of the ventricle which is currently depolarised as a piecewise linear function of time which depends on $T_{d,global}$, $T_{r,global}$ and $QRS_{duration}$. Then the values of $|u|_+$ and $|u|_-$ in Equation (a) are adapted to depend on the value of f_{depo} as described in Table 12.

References

1. Brian Baillargeon, Nuno Rebelo, David D Fox, Robert L Taylor, and Ellen Kuhl. The living heart project: a robust and integrative simulator for human heart function. *European Journal of Mechanics-A/Solids*, 48:38–47, 2014.
2. Gabriel Balaban, Henrik Finsberg, Hans Henrik Odland, Marie E Rognes, Stian Ross, Joakim Sundnes, and Samuel Wall. High-resolution data assimilation of cardiac mechanics applied to a dyssynchronous ventricle. *International journal for numerical methods in biomedical engineering*, 33(11), 2017.
3. Matthieu Caruel, Radomir Chabiniok, Philippe Moireau, Yves Lecarpentier, and Dominique Chapelle. Dimensional reductions of a cardiac model for effective validation and calibration. *Biomechanics and modeling in mechanobiology*, 13(4):897–914, 2014.
4. Radomir Chabiniok, Vicky Y. Wang, Myrianthi Hadjicharalambous, Liya Asner, Jack Lee, Maxime Sermesant, Ellen Kuhl, Alistair A. Young, Philippe Moireau, Martyn P. Nash, Dominique Chapelle, and David A. Nordsletten. Multiphysics and multiscale modelling, data-model fusion and integration of organ physiology in the clinic: ventricular cardiac mechanics. *Interface Focus*, 6(2), 2016. ISSN 2042-8898. URL <http://rfs.royalsocietypublishing.org/content/6/2/20150083>.

³<https://models.physiomeproject.org/e/470>

$$\begin{cases}
\dot{k}_c = -(|u|_+ + |u|_- + \alpha|e_c|)k_c + k_0|u|_+ & (a) \\
\dot{\tau}_c = -(|u|_+ + |u|_- + \alpha|e_c|)\tau_c + \dot{e}_c k_c + \sigma_0|u|_+ & (b) \\
(\tau_c + \mu\dot{e}_c) = E_s \frac{(e_{1D} - e_c)(1 + 2e_{1D})}{(1 + 2e_c)^3} & (c) \\
C = \left(1 + \frac{y}{R_0}\right)^2 & (d) \\
e_{1D} = \frac{C - 1}{2} & (e) \\
\sigma_{1D} = E_s \frac{e_{1D} - e_c}{(1 + 2e_c)^2} & (f) \\
\sigma_{\text{passive}} = 4(1 - C^{-3})(c_1 + c_2 C) & (g) \\
\sigma_{\text{viscosity}} = 4\eta(1 + C^{-6})\dot{C} & (h) \\
\Sigma_{\text{sph}} = \sigma_{1D} + \sigma_{\text{passive}} + \sigma_{\text{viscosity}} & (i) \\
\rho d_0 \ddot{y} = P_v \left(1 + \frac{y}{R_0}\right)^2 - \frac{d_0}{R_0} \left(1 + \frac{y}{R_0}\right) \Sigma_{\text{sph}} & (j) \\
q = 4\pi R_0^2 \left(1 + \frac{y}{R_0}\right)^2 \dot{y} = \begin{cases} K_{\text{at}}(P_V - P_{\text{at}}) & \text{for } P_V \leq P_{\text{at}} \\ K_{\text{iso}}(P_V - P_{\text{at}}) & \text{for } P_{\text{at}} \leq P_V \leq P_{\text{ar}} \\ K_{\text{ar}}(P_V - P_{\text{at}}) + K_{\text{iso}}(P_{\text{ar}} - P_{\text{at}}) & \text{for } P_V \geq P_{\text{ar}} \end{cases} & (k) \\
R_p C \dot{P}_{\text{ar}} + P_{\text{ar}} - P_{\text{ve}} = (R_p + Z_c)q + R_p Z_c C \dot{q} & (l)
\end{cases}$$

TABLE 11 Mechanical equations of the 0D model

$$f = \begin{cases} \frac{t - T_{d,global}}{QRS_{\text{duration}}} & \text{when } T_{d,global} \leq t \leq T_{d,global} + QRS_{\text{duration}} \\ 1 & \text{when } T_{d,global} + QRS_{\text{duration}} \leq t \leq T_{r,global} \\ \left(1 - \frac{t - T_{r,global}}{QRS_{\text{duration}}}\right) & \text{when } T_{r,global} \leq t \leq T_{r,global} + QRS_{\text{duration}} \\ 0 & \text{otherwise} \end{cases} \quad (28)$$

$$\begin{aligned}
|u|_+ &= k_{\text{ATP}} \cdot f_{\text{depo}} \\
|u|_- &= k_{\text{RS}} \cdot (1 - f_{\text{depo}})
\end{aligned}$$

TABLE 12 Electrical activation in the 0D model

5. D. Chapelle, P. Le Tallec, P. Moireau, and M. Sorine. Energy-preserving muscle tissue model: formulation and compatible discretizations. *International Journal for Multiscale Computational Engineering*, 10(2), 2012.
6. Zhong Chen, Rocio Cabrera-Lozoya, Jatin Relan, Manav Sohal, Anoop Shetty, Rashed Karim, Herve Delingette, Jaswinder Gill, Kawal Rhode, Nicholas Ayache, Peter Taggart, Christopher Aldo Rinaldi, Maxime Sermesant, and Reza Razavi. Biophysical modelling predicts ventricular tachycardia inducibility and circuit morphology: A combined clinical validation and computer modelling approach. *Journal of Cardiovascular Electrophysiology*, 27(7): 851–860, 2016. . URL <https://hal.archives-ouvertes.fr/hal-01301426>.
7. Autumn A. Cuellar, Catherine M. Lloyd, Poul F. Nielsen, David P. Bullivant, David P. Nickerson, and Peter J. Hunter. An overview of cellml 1.1, a biological model description language. *SIMULATION*, 79(12):740–747, 2003. . URL <http://dx.doi.org/10.1177/0037549703040939>.
8. Nicolas Duchateau, Mathieu De Craene, Pascal Allain, Eric Saloux, and Maxime Sermesant. Infarct localization from myocardial deformation: Prediction and uncertainty quantification by regression from a low-dimensional space. *Transactions on Medical Imaging*, 2016.
9. Henrik Finsberg, Gabriel Balaban, Stian Ross, Trine F Håland, Hans Henrik Odland, Joakim Sundnes, and Samuel Wall. Estimating cardiac contraction through high resolution data assimilation of a personalized mechanical model. *Journal of Computational Science*, 2017.
10. Alan Garny and Peter J Hunter. Opencor: a modular and interoperable approach to computational biology. *Frontiers in Physiology*, 6(26), 2015. ISSN 1664-042X. . URL http://www.frontiersin.org/computational_physiology_and_medicine/10.3389/fphys.2015.00026/abstract.
11. Nikolaus Hansen. The cma evolution strategy: a comparing review. In *Towards a new evolutionary computation*, pages 75–102. Springer Berlin Heidelberg, 2003.

2006.



12. A Huxley. Muscle structure and theories of contraction. *Progress in biophysics and biophysical chemistry*, 7:255–318, 1957.
13. Elham Kayvanpour, Tommaso Mansi, Farbod Sedaghat-Hamedani, Ali Amr, Dominik Neumann, Bogdan Georgescu, Philipp Seegerer, Ali Kamen, Jan Haas, Karen S Frese, Maria Irawati, Emil Wirsz, Vanessa King, Sebastian Buss, Derliz Mereles, Edgar Zitron, Andreas Keller, Hugo A Katus, Dorin Comaniciu, and Benjamin Meder. Towards Personalized Cardiology: Multi-Scale Modeling of the Failing Heart. *PLoS ONE*, 10(7), 2015. .
14. E. Konukoglu, J. Relan, U. Cilingir, B. Menze, P. Chinchapatnam, A. Jadidi, H. Cochet, M. Hocini, H. Delingette, P. Jaïs, M. Haïssaguerre, N. Ayache, and M. Sermesant. Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology. *Progress in Biophysics and Molecular Biology*, 107(1):134–146, Oct 2011. . URL <http://www.inria.fr/sophia/asclepios/Publications/Maxime.Sermesant/KonukogluPBMB2011.pdf>.
15. Brodie AJ Lawson, Christopher C Drovandi, Nicole Cusimano, Pamela Burrage, Blanca Rodriguez, and Kevin Burrage. Unlocking data sets by calibrating populations of models to data density: A study in atrial electrophysiology. *Science advances*, 4(1):e1701676, 2018.
16. S. Marchesseau, H. Delingette, M. Sermesant, and N. Ayache. Fast parameter calibration of a cardiac electromechanical model from medical images based on the unscented transform. *Biomechanics and modeling in mechanobiology*, 12(4), 815–831, 2013.
17. Stéphanie Marchesseau, Tobias Heimann, Simon Chatelin, Rémy Willinger, and Hervé Delingette. Multiplicative jacobian energy decomposition method for fast porous visco-hyperelastic soft tissue model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 235–242. Springer, 2010.
18. Roch Molléro, Xavier Pennec, Hervé Delingette, Alan Garny, Nicholas Ayache, and Maxime Sermesant. Multifidelity-cma: a multifidelity approach for efficient personalisation of 3d cardiac electromechanical models. *Biomechanics and Modeling in Mechanobiology*, Sep 2017. ISSN 1617-7940. . URL <https://doi.org/10.1007/s10237-017-0960-0>.
19. Dominik Neumann, Tommaso Mansi, Bogdan Georgescu, Ali Kamen, Elham Kayvanpour, Ali Amr, Farbod Sedaghat-Hamedani, Jan Haas, Hugo Katus, Benjamin Meder, et al. Robust image-based estimation of cardiac tissue parameters and their uncertainty from noisy data. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2014*, pages 9–16. Springer, 2014.
20. Marc-Michel Rohé, Roch Molléro, Maxime Sermesant, and Xavier Pennec. Highly Reduced Model of the Cardiac Function for Fast Simulation. In *IEEE - IVMSWP Workshop 2016, Image, Video, and Multidimensional Signal Processing Workshop (IVMSWP)*, 2016 IEEE 12th, page 5, Bordeaux, France, July 2016. IEEE. . URL <https://hal.inria.fr/hal-01373715>.
21. Maxime Sermesant, Radomir Chabiniok, Phani Chinchapatnam, Tommaso Mansi, Florence Billet, Philippe Moireau, Jean-Marc Peyrat, Ken C.L. Wong, Jatin Relan, Kawal S. Rhode, Matt Ginks, Pier Lambiase, Hervé Delingette, Michel Sorine, C. Aldo Rinaldi, Dominique Chapelle, Reza Razavi, and Nicholas Ayache. Patient-Specific Electromechanical Models of the Heart for Prediction of the Acute Effects of Pacing in CRT: a First Validation. *Medical Image Analysis*, 16(1):201–215, January 2012. . URL <https://hal.inria.fr/inria-00616191>.
22. Nic Smith, Adelaide de Vecchi, Matthew McCormick, David Nordsletten, Oscar Camara, Alejandro F. Frangi, Hervé Delingette, Maxime Sermesant, Jatin Relan, Nicholas Ayache, Martin W. Krueger, Walther H. W. Schulze, Rod Hose, Israel Valverde, Philipp Beerbaum, Cristina Staicu, Maria Siebes, Jos Spaan, Peter Hunter, Juergen Weese, Helko Lehmann, Dominique Chapelle, and Reza Rezavi. euheart: personalized and integrated cardiac care using patient-specific cardiovascular modelling. *Interface Focus*, 1(3):349–364, 2011. ISSN 2042-8898. . URL <http://rfs.royalsocietypublishing.org/content/1/3/349>.
23. Nicolaas Westerhof, Frederik Bosman, Cornelis J De Vries, and Abraham Noordergraaf. Analog studies of the human systemic arterial tree. *Journal of biomechanics*, 2(2), 1969.