Detailed evaluation of a population-wise personalization approach to generate synthetic myocardial infarct images

Anastasia Konik, Patrick Clarysse, Nicolas Duchateau

▶ **To cite this version:**

Anastasia Konik, Patrick Clarysse, Nicolas Duchateau. Detailed evaluation of a population-wise personalization approach to generate synthetic myocardial infarct images. Pattern Recognition Letters, 2024, 188, pp.8-14. ⟨10.1016/j.patrec.2024.11.017⟩. ⟨hal-04805237v2⟩

HAL Id: hal-04805237

https://hal.science/hal-04805237v2

Submitted on 11 Dec 2025

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

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

Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Detailed evaluation of a population-wise personalization approach to generate synthetic myocardial infarct images

Anastasia Konik[a], Patrick Clarysse[a], Nicolas Duchateau[a,b]

[a]*INSA-Lyon, Université Claude Bernard Lyon 1, CNRS, Inserm, CREATIS, UMR 5220, U1294, 21 avenue Jean Capelle, Villeurbanne, 69100, France*
[b]*Institut Universitaire de France (IUF), 103 boulevard Saint-Michel, Paris, 75005, France*

ABSTRACT

Personalization of biophysical models to real data is essential to achieve realistic simulations or generate relevant synthetic populations. However, some of these models involve randomness, which poses two challenges: they do not allow the standard personalization to each individual's data and they lack an analytical formulation required for optimization.
In previous work, we introduced a population-based personalization strategy which overcomes these challenges and demonstrated its feasibility on simple 2D geometrical models of myocardial infarct. The method consists in matching the distributions of the synthetic and real populations, quantified through the Kullback-Leibler (KL) divergence. Personalization is achieved with a gradient-free algorithm (CMA-ES), which generates sets of candidate solutions represented by their covariance matrix, whose coefficients evolve until the synthetic and real data are matched. However, the robustness of this strategy regarding settings and more complex data was not challenged.
In this work, we specifically address these points, with (i) an improved design, (ii) a thorough evaluation on crucial aspects of the personalization process, including hyperparameters and initialization, and (iii) the application to 3D data. Despite some limits of the simple geometrical models used, our method is able to capture the main characteristics of the real data, as demonstrated both on 2D and 3D segmented late Gadolinium images of 123 subjects with acute myocardial infarction.

Keywords: Model personalization; image synthesis; myocardial infarction; cardiac magnetic resonance; gradient-free optimization.

1. Introduction

Biophysical simulations are an invaluable support to medical image analysis. Apart from better understanding of physiological processes, they are a source of knowledge for data-driven approaches such as machine learning [1]. Such prior physiological knowledge can be either integrated into the learning process or used through physiologically-based synthetic data for training or validation of learning algorithms. The latter approach is particularly interesting as many variants of a simulation can be generated by modifying a small set of parameters in a relevant range. However, simulations remain simplified models of the reality, and a substantial gap may therefore exist between synthetic and real data. The process of *personalizing* a biophysical model explicitly targets this: it consists in finding the optimal parameters of this model such that it gets as close as possible to the real data [2]. While personalization to a given subject's data is the standard approach but for deterministic models, we specifically target in this paper personalization to a population of subjects, which is suited for models including randomness.

1.1. Personalization of biophysical models

As biophysical models can be complex and computationally demanding, a personalization strategy based on grid search or heuristic approaches should be avoided for efficiency and robustness reasons. However, personalization can be automated by means

of more adapted optimization methods. Among these, gradient-free approaches are particularly interesting when an analytical formulation of the models is complex or almost impossible, as they do not require specific knowledge on the internal behavior of the model. In particular, the Covariance Matrix Adaptation – Evolution Strategy (CMA-ES) [3] stands as state-of-the-art on many non-linear non-convex optimization problems, as demonstrated for the personalization of complex biophysical models of the heart to individuals' data [4].

Nonetheless, *individual* personalization to each subject's data is not always feasible, in particular for models that include randomness (e.g., for the generation of vascular [5] or electrical [6] networks, or as in our application, the generation of infarct patterns within the myocardium). In those cases, *population-wise* personalization should be preferred. It consists of an optimization process that adapts the distribution of a synthetic population generated by the model to the real population. The KL divergence can be a relevant choice for quantifying differences between these distributions, provided the data complexity is correctly taken into account. In previous work, we have shown a proof-of-concept of this population-wise approach for the personalization of simple geometrical models of myocardial infarction in 2D [7]. However, several crucial challenges remain regarding the robustness of the personalization procedure, its ability to consider subtle physiological characteristics of the infarcts and its generalization to 3D data. The first objective of the proposed paper is to explicitly address these issues.

1.2. Geometrical models of myocardial infarction

Myocardial infarct is mostly assessed by late Gadolinium enhancement in magnetic resonance imaging [8], which provides a stack of 2D images covering the left ventricle from base to apex (with between-slice resolution around 5 to 8 mm and intra-slice resolution around 1 to 2 mm). Infarcts have complex 3D shapes, which vary depending on the occluded coronary artery [9]. Across the course of occlusion, infarcts propagate from the inner surface of the ventricle (endocardium) and progressively extend towards the outer surface (epicardium), following a complex wavefront propagation [10].

In contrast, the few geometrical models proposed to represent infarct zones are rather simple compared to those for infarct biomechanics [11]. Several works generating synthetic databases for regional quantification of the myocardial function considered a set of segments from the American Heart Association definition [12], whose properties are altered to mimic the diseased tissue [13]. Partial transmural extent of the infarct was better considered with ellipses centered on the endocardium [14]. A step further in generating complex infarct shapes was achieved through the union of random spheres, the first one starting at the endocardium [15]. Much more complex models considering infarct progression from the coronary artery occlusion were also proposed [16].

However, all these models have been set heuristically and are therefore hardly representative of a given population. A proper personalization of such models would not only reduce the domain gap between the synthetic and real data, but would also allow stating on the relevance of such geometrical models and if further refinements are needed. As some of these models include randomness, they cannot be personalized to the data of a given individual, and a population-wise approach is therefore needed. The second objective of the proposed paper is to thoroughly evaluate the population-wise personalization strategy for such models of myocardial infarct.

1.3. Deep learning-based medical image synthesis

State-of-the-art methods for medical image synthesis from a given data distribution mostly consist of GANs (generative adversarial networks) and diffusion models. For example, different GAN architectures were compared on 2D cardiac MRI images (on the public ACDC dataset) in [17], the highest performance was achieved by the StyleGAN method [18]. The work in [19] proposed Medfusion, a conditional latent denoising diffusion probabilistic model (DDPM), which outperformed StyleGAN for medical non-cardiac datasets. Still on the ACDC dataset, the work in [20] showed that a transformer-based DDPM performed better than 2D GANs. On 3D data, first attempts on brain MRIs [21] showed comparable performance between DDPMs and 3D GANs. A review of other previously popular deep learning architectures for medical image synthesis can be found in [22].

Despite these achievements, these algorithms are still extremely long to train. Besides, our underlying interest is in biophysical models, which provide physiologically-relevant simulations of a given subpart of the body, for example, for better understanding pathophysiological processes and eventually performing in-vitro digital experiments. These models can also be used into an image synthesis pipeline, with the asset of adding physiological knowledge to the simulations. Recent state-of-the-art examples in cardiac imaging are on the echocardiographic sequences [23] and MR images [24] generation. Our work demonstrates that such models can be personalized to a whole population (here, we chose simple geometrical models to better focus on the personalization process), therefore opening the way to the use of more complex simulations for the generation of synthetic but physiologically-realistic simulations.

1.4. Contributions

In this paper, we detail a population-wise personalization method which is particularly relevant for the generation of realistic synthetic medical images with models involving randomness. We focus on geometrical models of myocardial infarct which are simple enough to lay stress on the evaluation of the personalization procedure in 2D and 3D.

We presented a proof-of-concept of this approach at the FIMH conference in 2021 [7]. It was limited to 2D infarcts with no evaluation of the method robustness against the personalization design, hyperparameters and initialization and coded in Matlab.

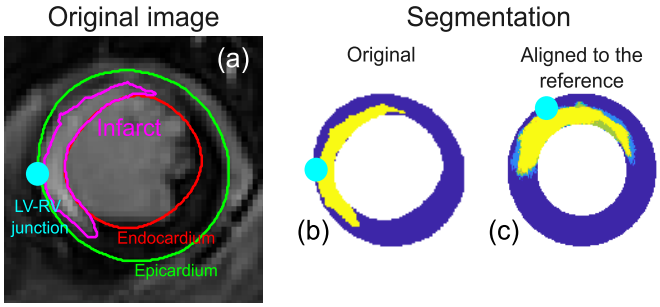

Fig. 1. Representative real magnetic resonance image (2D slice from a 3D stack) and corresponding infarct segmentation (before and after spatial alignment to a common reference geometry) used in this work.

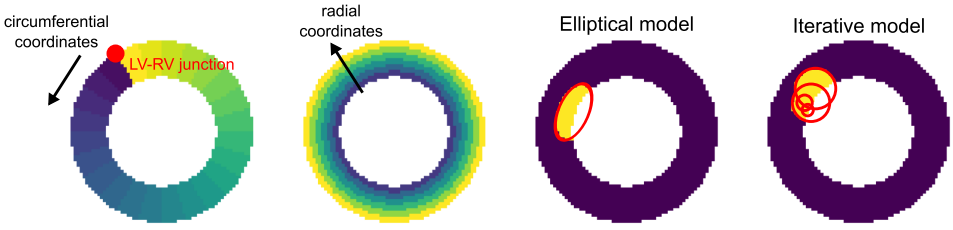

Fig. 2. Left (two columns): anatomical coordinates used to parameterize the myocardium in 2D, with the left/right ventricular junction indicated by the red dot. Right (two columns): 2D examples of the geometrical models tested in this paper.

The present paper significantly extends this work by explicitly considering these aspects and targeting both 2D and 3D data, with a publicly available and updated implementation in Python.

2. Material and methods

The following sections summarize the data and preprocessing steps (Sec. 2.1), the way our geometrical models of myocardial infarct are constructed (Sec. 2.2), the core of the population-wise personalization strategy we propose (Sec. 2.3), and our evaluation procedure (Sec. 2.4).

2.1. Data and preprocessing

The real data consisted of segmentations of acute infarcts from the Minimalist Immediate Mechanical Intervention study (MIMI / Clinical Trials ID: NCT01360242) [25]. This study comprised segmented late Gadolinium enhancement images from 123 patients. The myocardial borders (endocardium and epicardium) and the infarct were segmented in each slice by experts using commercial software (CVI42 v.5.1.0 Circle Cardiovascular Imaging, Calgary, CA), as illustrated in Figure 1a. The data from each individual were spatially aligned to a common reference geometry (a semi-ellipsoid with maximal endocardial and epicardial radii of 30 and 50 pixels, respectively), as described in [9] and illustrated in Figure 1b and c. This alignment included a rotation along the circumference such that all infarcts match with the territory of the left anterior descending artery (LAD), therefore providing a larger dataset at a similar location. Of note, this process may be arguable as infarct patterns slightly differ depending on the occluded coronary artery. We refer the reader to [25] for more details on these data and to [9] for its detailed statistical analysis.

After spatial alignment, the images consisted of 3D stacks of 11 slices of 80×80 pixels each. We focused on two subsets of images: (i) the 2D slices at mid-cavity exhibiting a myocardial infarct (118 images), (ii) the 3D stacks for which the left ventricle was fully covered and segmented (77 images).

2.2. Models of infarct

We investigated two simple geometrical models of acute myocardial infarction (Figure 2) which allow generating synthetic lesions of varying size, location along the circumference, transmurality and shape:

1. *Elliptical:* a random ellipsoid centered on the endocardium and with one axis tangent to the endocardium. The infarct consists of the intersection between this ellipsoid and the myocardium. In [7], this model was controlled by the length of the ellipsoid axes (4 parameters for 2D data: for each axis, the distribution center and its extension). Here, we improve the model by using means and standard deviations to have better correspondence with the individual fitting of ellipses (4 and 6 parameters for 2D and 3D data, respectively).

2. *Iterative spherical:* the union of a random number of spheres of random size, the first one being centered on the endocardium. Here too, the infarct consists of the intersection between this region and the myocardium. This model is controlled by the number of spheres and the maximum possible sphere radius (2 parameters for 2D and 3D data).

In both cases, the circumferential coordinates of the starting point (center of the ellipsoid or of the first sphere) were set in the neighborhood of the weighted average circumferential angle of all infarcts from the real population.

2.3. Population-wise personalization method

2.3.1. Individual vs. population-based personalization

Let's denote $\{\mathbf{x}_p\}_{p\in[1,P]}$ and $\{\mathbf{y}_q\}_{q\in[1,Q]}$ the sets of P synthetic and Q real infarct images, represented by the distributions $\mathcal{P}$ and $\mathcal{Q}$ respectively.

Individual personalization (under the condition that $P = Q$) means finding the optimal parameters $\theta_q \in \mathbb{R}^{m_d}$ such that $\mathbf{x}_q = f_d(\theta_q) \approx \mathbf{y}_q$, $\forall q \in [1, Q]$, where m_d is the number of parameters of a given (deterministic) model f_d. For example, $m_d = 4$ for the 2D elliptical model, and parameters stand for the means and standard deviations of the distribution of the elliptical axes.

In contrast, *population-based personalization* consists in finding the optimal parameters $\theta \in \mathbb{R}^{m_r}$ which lead to comparable synthetic and real distributions $\mathcal{P}$ and $\mathcal{Q}$, quantified by the loss L, where f_r is the (random) model generating the synthetic samples $\{\mathbf{x}_p = f_r(\theta)\}_{p\in[1,P]}$ and m_r is the number of parameters of the random model. In other words, this amounts at solving:

$$\hat{\theta} = \underset{\theta\in\mathbb{R}^{m_r}}{\operatorname{argmin}} \quad L\Big(\{\mathbf{x}_p = f_r(\theta)\}_{p\in[1,P]}, \{\mathbf{y}_q\}_{q\in[1,Q]}\Big). \tag{1}$$

2.3.2. Personalization with CMA-ES

We implemented the personalization with the Covariance Matrix Adaptation – Evolution Strategy (CMA-ES) [3]. This algorithm consists in iteratively adapting a covariance matrix representing candidate samples of the parameters to optimize (in our case, each sample of parameters serves to generate a given synthetic population). At each iteration (referred to as generation), CMA-ES randomly samples N_{tot} parameter sets from a multi-normal distribution in the space of possible parameters and retains the N_{best} parameter sets with the lowest loss values (see definition of the loss in Sec. 2.3.3). Then, the covariance matrix is updated with these samples until convergence. The personalization stops when one of three following criteria is met: the loss is below a given threshold ϵ, or M iterations have been performed, or the condition number of the covariance matrix exceeds 10^{14}.

2.3.3. Loss function

The KL divergence is relevant to compare the distributions $\mathcal{P}$ and $\mathcal{Q}$, approximated from the known samples $\{\mathbf{x}_p\}_{p\in[1,P]}$ and $\{\mathbf{y}_q\}_{q\in[1,Q]}$ using a Gaussian kernel density estimation of bandwidth σ. We used two variants of the loss, corresponding to the way we consider samples in the kernel density estimation:

1. Directly comparing samples using the L^2-norm on images of the infarct segmentations, leading to the loss KL_{full},
2. Comparing samples based on two values encoding infarct-related attributes estimated from the infarct segmentations: infarct size and transmurality, leading to the loss KL_{attr}.

2.4. Evaluation of the personalization approach

The optimization was monitored across iterations using:

1. The decrease of the loss function,
2. The evolution of the covariance eigenvalues, which may increase until the optimum is within the search range encoded by the covariance and then decrease until a stable search range,
3. The evolution of the optimized parameters.

Gradient-free methods such as CMA-ES may be sensitive to initialization; we evaluated the stability of the final parameters estimated by the personalization with respect to the CMA-ES initialization (Sec. 3.2.1, 3.2.2 and 3.3) and its parameters (Sec. 3.2.3). Also, as the elliptical model is deterministic, we compared the population-based personalization of this model against the values found by personalizing it to each real sample (Sec. 3.2.1).

Finally, for interpretation purposes, we encoded the synthetic and real populations into a low-dimensional latent space (one for each population) and reconstructed high-dimensional infarct images corresponding to $\{-2, -1, 0, +1, +2\}$ standard deviations along each of the first three dimensions. For the low-dimensional encoding, we used a standard manifold learning algorithm for which the Euclidean distance in the latent space has a statistical meaning (Isomap [26]). It performs the eigendecomposition of an affinity matrix encoding the shortest path between pairs of samples, therefore leading to a low-dimensional latent space where the Euclidean distance approximates the geodesic distance between high-dimensional samples. This algorithm remains simple and offers relevant statistical meaning to navigate across the latent space, which would not be fully guaranteed with other non-linear methods (e.g., t-SNE, UMAP, PaCMAP, variational autoencoders), and does not suffer the limitations of linear methods (e.g., PCA) to analyze such data. As decoding (namely, reconstructing high-dimensional samples from the low-dimensional latent space) is not

intrinsically part of the Isomap method, we reconstructed high-dimensional samples with a multiscale version of the kernel ridge regression [27], which is robust to the non-uniform density of samples in the latent space, with equally weighted regularization and data terms. Note that this regression may provide intermediate values between the healthy and infarcted regions of the myocardium (dark blue and yellow in 2D images), which we kept on purpose as thresholding would degrade the lesion patterns and lead to misinterpretations. We only thresholded the values above/below the healthy/infarct limits.

3. Results

The following experiments evaluate the performance and robustness of our personalization on 2D (Sec. 3.2) and 3D (Sec. 3.3) data, for the iterative (2D and 3D) and elliptical (2D, serving for benchmarking, being deterministic) models.

3.1. Personalization hyperparameters

In the following experiments, we set the sampling size to $N_{\text{tot}} = 15$ samples generated at each iteration of the personalization algorithm (except Sec. 3.2.3 where we explicitly evaluate its influence) and $N_{\text{best}} = 3$ best samples were retained to update the covariance matrix (initialized by the identity matrix).

The other CMA-ES parameters consisted of the maximum number of iterations and the functional threshold, which were set to $M = 34$ (leading to $N_{\text{tot}} \times M = 510$ samplings and loss evaluations) and $\epsilon = 0.004$, respectively. Finally, the bandwidth σ used for the kernel density estimation in the KL divergence was set to the average distance between each sample and its 9 nearest neighbors.

3.2. 2D personalization

3.2.1. Elliptical model

Individual personalization. For benchmarking, we first personalized this deterministic model to the data of each individual. This means estimating the ellipse whose intersection with the myocardium best matches each subject's infarct pattern. The candidate ellipses were fit by selecting their center (in the endocardial neighborhood, 20 candidate circumferential coordinates around the weighted average circumferential angle of all infarcts) and their axes (grid search, 160 candidate values in the interval [0.5, 80] pixels). Similarity between the synthetic and real infarcts was quantified by the Dice coefficient.

The individual personalization to one image took around 200 seconds and around 6 hours for all 118 images[1]. After this process, the estimated parameters of ellipse axes were $(22, 16, 16, 11)$ pixels, in the order: mean 1^{st} axis, standard deviation 1^{st} axis, mean 2^{nd} axis, standard deviation 2^{nd} axis across the 118 individual personalizations.

Population-based personalization. To evaluate the stability of the personalization with respect to the initial parameter values, we launched the population-based personalization 10 times, for different initial values of each model parameter, and examined both the final loss and final parameter values. As this can be computationally expensive, we separated the evaluation of the 4 parameters of the elliptical model in two steps: the means and then the standard deviations encoding the distribution of the ellipse axes. Results are summarized in Tables 1 and 2.

The best initialization consisted of the parameters $(22, 10, 16, 20)$ pixels, the final parameters after the personalization being $(18, 10, 13, 18)$ pixels, with standard deviations across the 10 launches of $(1, 1, 2, 1)$ pixels.

Figure 3 displays the first three modes of variation of synthetic infarct patterns generated with this optimal solution, to be compared to those obtained for the real population. The modes of the synthetic population appear to encode the same infarct characteristics as in the real population (1^{st}: size, 2^{nd}: transmurality, 3^{rd}: angular position), indicating that personalization to the real population was achieved.

Table 1. Sensitivity to initial guess of the 2D elliptical model: average solutions (mean values for 4 model parameters: mean 1^{st} axis, std 1^{st} axis, mean 2^{nd} axis, std 2^{nd} axis) for variations in initial axes means for KL_{full}. The smallest loss corresponds to yellow, the highest to red, the best solution is in bold.

Mean 1st axis / Mean 2nd axis	10	15	20	25	30
5	(18,18,11,16)	(18,16,10,16)	(22,16,8,15)	(22,13,5,14)	(24,13,5,12)
10	(16,18,14,13)	(18,14,14,11)	(20,13,12,12)	(21,13,9,11)	(23,12,7,13)
15	(15,17,18,12)	(17,16,16,11)	(18,11,12,8)	(18,11,12,12)	**(19,11,11,11)**
20	(15,16,21,15)	(16,15,19,11)	(18,12,15,12)	(18,11,14,11)	(18,11,12,11)
25	(14,16,24,13)	(15,14,24,14)	(16,13,21,14)	(18,11,17,14)	(18,11,14,18)
30	(13,17,28,14)	(15,13,29,12)	(16,13,24,14)	(16,13,21,17)	(17,13,18,15)

[1] On a computer with an Intel(R) Core(TM) i7-10610U CPU@1.80GHz, RAM 16 GB.

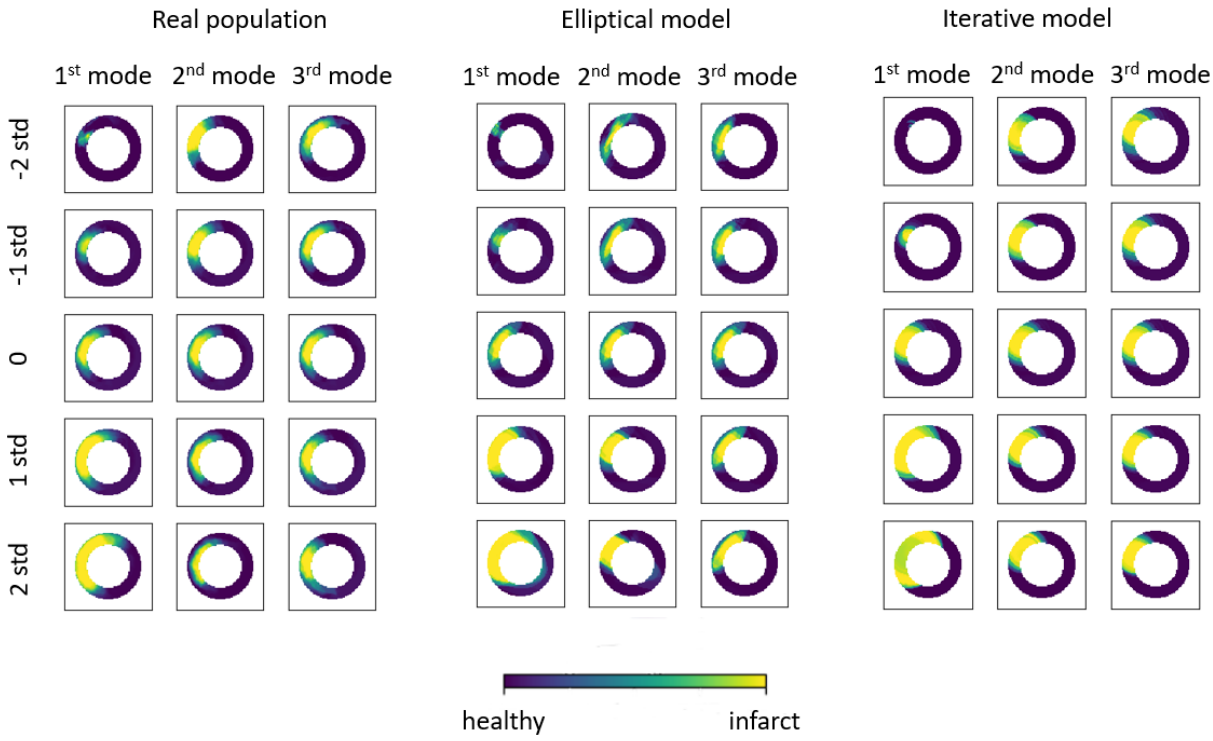

Fig. 3. First modes of variation for the 2D real population and the 2D synthetic populations generated using the elliptical and the iterative models, with the best parameters learned by the population-based personalization with KL_{full}.

Table 2. Sensitivity to initial guess of the 2D elliptical model: average solutions (mean values for 4 model parameters: mean 1^{st} axis, std 1^{st} axis, mean 2^{nd} axis, std 2^{nd} axis) for variations in initial axes standard deviations for KL_{full}. Displayed similar to Tab. 1.

Std 1st axis / Std 2nd axis	5	10	15	20	25
5	(19,11,12,15)	(19,11,12,11)	(19,13,12,10)	(18,17,13,9)	(17,19,12,9)
10	(19,10,14,16)	(19,11,13,12)	(19,12,13,11)	(20,15,12,9)	(18,20,11,7)
15	(19,10,14,17)	(19,10,14,16)	(19,11,12,12)	(18,17,12,14)	(18,20,12,13)
20	(18,9,12,18)	**(18,10,13,18)**	(19,11,13,18)	(19,16,12,16)	(18,19,12,16)

3.2.2. Iterative model

The population-based personalization procedure of this model was similar to the elliptical model's one. The best initialization consisted in starting with 5 iterations and radii of 35 pixels, the final values after the personalization being 1 and 32 (10-launches standard deviations: 0 and 3), as visible in Table 3.

Figure 3 displays the first three modes of variation of the synthetic infarct patterns generated with this optimal solution. Although the personalization process was completed, the modes of variation match less with the real population than those of the elliptical model, as confirmed by the loss after personalization: $KL_{\text{full}} = 0.0056$ and 0.0155 for the elliptical and iterative models, respectively. The main reason for this is the limited design of the iterative model, and in particular the way spheres are progressively added to the existing infarct pattern, with potentially too much randomness in 2D.

3.2.3. Influence of hyperparameters

Figure 4 displays the evolution of the loss across iterations for two values of the sampling size: $N_{\text{tot}} = 15$ and 50 (average curves across 10 repeated launches of the personalization), with proportionally corresponding $N_{\text{best}} = 3$ and 12. The personalization converges in both cases for both 2D models. It is faster and more stable for $N_{\text{tot}} = 50$, but also more computationally expensive, which is why we retained the faster procedure with $N_{\text{tot}} = 15$ (around 4 sec for one iteration compared to 15 sec for $N_{\text{tot}} = 50$). Note that in this experiment, almost the same solutions for both sampling sizes were obtained:

1. for $N_{\text{tot}} = 15$: $(19, 10, 9, 22)$ pixels (10-launches standard deviations: $(1, 1, 2, 3)$ pixels) for the elliptical model, and $(1, 31)$ (10-launches standard deviations: $(0, 2)$) for the iterative model,
2. for $N_{\text{tot}} = 50$: $(19, 9, 11, 19)$ pixels (10-launches standard deviations: $(0, 0, 3, 2)$ pixels) for the elliptical model, and $(1, 30)$ (10-launches standard deviations: $(0, 1)$) for the iterative model.

Table 3. Sensitivity to initial guess of the 2D iterative model: average solutions (number of iterations and maximal radius) for variations in initial parameters for KL_{full}. Displayed similar to Tab. 1.

Iterations / Max radius	1	5	10	15	20
5	(13,14)	(11,14)	(11,13)	(17,11)	(21,11)
10	(7,16)	(6,15)	(14,12)	(15,12)	(19,11)
15	(3,19)	(4,17)	(9,13)	(14,12)	(17,11)
20	(2,21)	(2,20)	(4,17)	(11,14)	(13,13)
25	(1,27)	(2,25)	(4,20)	(5,17)	(10,14)
30	(1,30)	(1,28)	(2,21)	(12,15)	(11,15)
35	(1,34)	**(1,32)**	(2,25)	(5,18)	(14,14)

We did the initialization in the neighborhood of the best initial values reported in earlier sections. Note that $N_{\text{tot}} = 15$ with $N_{\text{best}} = 7$, $N_{\text{tot}} = 50$ with $N_{\text{best}} = 25$ provided slightly faster convergence. From Figure 4 we see that $M = 34$ was enough to guarantee the convergence of the optimization. In all experiments the optimization procedure ended when the maximum number of iterations was reached, before reaching the functional threshold, thus we decided to keep without testing the reasonably high $\epsilon = 0.004$. Different bandwidth choices ($\sigma = 5, 9, 15$) gave similar solutions with a slight tendency towards decreasing final stabilized loss values.

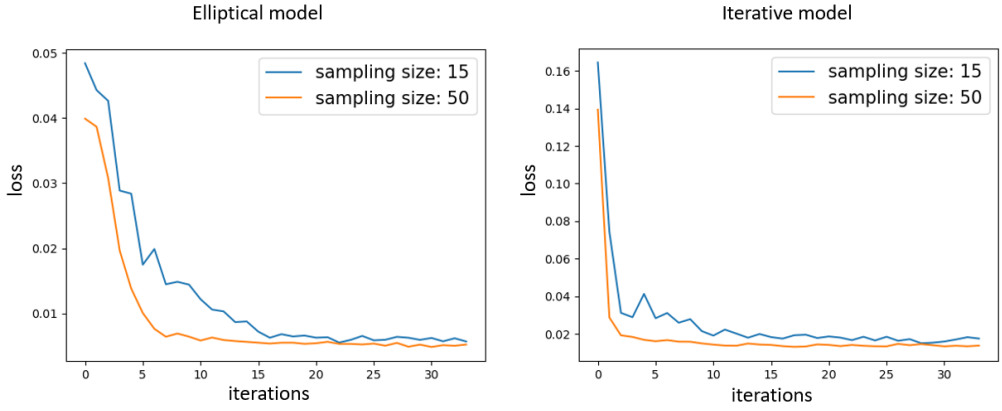

Fig. 4. Influence of the sampling size ($N_{\text{tot}} = 15$ and 50) on the convergence of the 2D elliptical and iterative models with KL_{full}.

3.2.4. Influence of the loss

The first modes of variation obtained for the KL_{attr} loss (meaning samples are compared according to infarct-related attributes) are provided as Supplementary Material for the elliptical and iterative models. The modes are similar to the ones obtained with the KL_{full}, although the final parameter values are slightly different: $(18, 9, 13, 17)$ pixels vs. $(18, 10, 13, 18)$ pixels for the elliptical model, 2 iterations and radii of 21 pixels vs. 1 iterations and radii of 32 pixels for the iterative model.

3.3. 3D personalization (iterative model)

Table 4 reports the final parameters obtained for the 3D personalization of the iterative model with different initialization parameters. KL_{full} final parameters (34 iterations and radii of 9 pixels, 10-launches standard deviations: 2 and 0) were comparable to those for KL_{attr} (33 iterations and radii of 8 pixels, 10-launches standard deviations: 5 and 0). We only tested a limited range of initialization parameters and repeated 5 times the optimization, due to the substantially higher computational cost for 3D data (around 9 hours for the experiments with KL_{full} and around 20 hours for the experiments with KL_{attr}).

Figure 5 illustrates the population mean (non-linear average reconstructed after low-dimensional embedding by Isomap) for the real population and the synthetic population after personalization. The pattern of the synthetic mean is more compact, also the examination of the main modes of variation (not displayed due to space constraints) indicates limited matching of some infarct characteristics, potentially due to the limits of the iterative model used.

4. Discussion

We presented a detailed evaluation of an original strategy for the personalization of biophysical models with randomness, in the perspective of using these models to generate populations of realistic synthetic images. We focused the study on simple geometrical

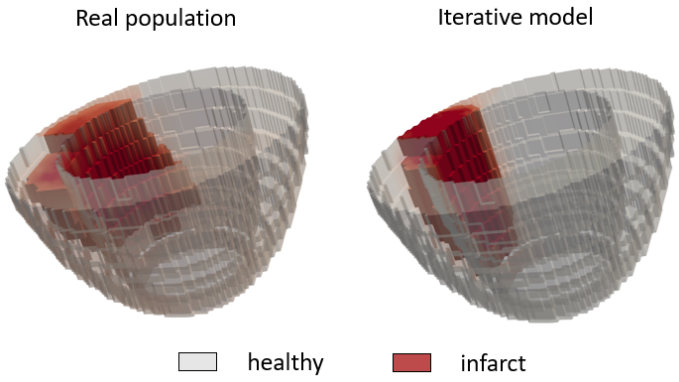

Fig. 5. Non-linear average infarct patterns for the real and synthetic (3D iterative model with KL_{full}) populations.

Table 4. Sensitivity to initial guess of the 3D iterative model: average solutions (number of iterations and maximal radius) for variations in initial parameters for KL_{full} and KL_{attr}. The smallest loss corresponds to yellow, the highest to red, the best solution is in bold.

Initial values / Loss types	(1,30)	(10,20)	(20,10)	(30,1)
KL_{full}	(2,30)	(25,10)	(21,10)	**(34,9)**
KL_{attr}	(1,27)	(41,6)	(20,8)	**(33,8)**

models of myocardial infarct, enhancing interpretation in 2D and 3D, including comparisons with a deterministic model for which individual personalization is also possible.

Our results confirm that population-based personalization is feasible, within the limits of the realism of the retained models. In 2D, the main variations in the synthetic data correspond to those in the real data, although the personalization favors very simple infarcts (1 to 2 spheres for the iterative model, with slightly lower fit performance than the elliptical model). However, in 3D, the iterative process leads to more complex unions of spheres to represent infarcts and therefore to a better fit to the real population. No substantial differences were observed with a loss focused on specific attributes of the infarct patterns (i.e., transmurality and size).

Finer geometrical models of infarcts may lead to better realism and therefore better personalization. In particular, mimicking the wavefront propagation of infarcts [10] with a diffusion process or more constraints on the aggregation of spheres in the iterative model. However, the design of such models goes beyond the scope of our paper, whose primary purpose was to demonstrate that population-based personalization is effective even with simple geometrical models.

This personalization strategy paves the way for two types of extensions. First, in the context of myocardial infarct, one could consider adding finer aspects of the lesions [8] such as reperfusion lesions (micro-vascular obstruction) that can be represented as smaller areas within the infarct, border zone (e.g., with intermediate values between healthy and diseased) and potential changes from acute to chronic infarcts. These could directly be used for generating realistic synthetic images [13, 28], for the training or the evaluation of learning algorithms (e.g., to estimate myocardial deformation). Then, this work also has high potential for extensions to other medical imaging applications involving random models (e.g., vascular [5] or electrical [6] network generation), although for these applications quantitative comparisons between the synthetic and real populations may be challenging to formulate.

5. Conclusion

We have demonstrated that a population-wise strategy is relevant for the personalization of models that include randomness, which would not be possible with standard personalization to each individual sample. On simple geometrical infarct models, it allowed generating synthetic images whose distribution matches a population of real clinical samples. It paves the ground for the generation of realistic synthetic populations from more complex models, of relevance for training and validation in the context of machine learning for medical image analysis.

CRediT authorship contribution statement.
AK: Methodology, Software, Validation, Investigation, Visualization, Writing - Original Draft.
PC: Supervision, Writing - Review & Editing.
ND: Conceptualization, Data Curation, Supervision, Writing - Review & Editing.

Declaration of competing interest. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability. Code and demo data will be made accessible upon acceptance, in the form of a ready-to-launch Jupyter notebook.

Acknowledgments. The authors acknowledge the support from the French National Research Agency (ANR) through the LABEX PRIMES of Université de Lyon (ANR-11-LABX-0063) and the Institut Universitaire de France.

The authors are grateful to the H2P platform[2] (CREATIS, Villeurbanne, France) for the data storage and management support, and to Pierre Croisille, Magalie Viallon and Lorena Petrusca (CREATIS, CHU Saint Etienne, France) for providing the imaging data and segmentations for the MIMI population.

References

[1] N. Duchateau, O. Camara, Machine learning and biophysical models: how to benefit each other?, Reduced order models for the biomechanics of living organs, Elsevier (2023) 147–64.
[2] R. Gray, P. Pathmanathan, Patient-specific cardiovascular computational modeling: diversity of personalization and challenges, J Cardiovasc Trans Res 11 (2018) 80–8.
[3] N. Hansen, A. Ostermeier, Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation, Proc. ICEC (1996) 312–7.
[4] R. Mollero, X. Pennec, H. Delingette, et al., Multifidelity-CMA: a multifidelity approach for efficient personalisation of 3D cardiac electromechanical models, Biomech Model Mechanobiol 17 (2018) 285–300.
[5] L. Cury, G. Maso Talou, M. Younes-Ibrahim, et al., Parallel generation of extensive vascular networks with application to an archetypal human kidney model, R Soc Open Sci 8 (2021) 210973.
[6] L. Berg, B. Rocha, R. Oliveira, et al., Enhanced optimization-based method for the generation of patient-specific models of Purkinje networks, Sci Rep 13 (2023) 11788.
[7] K. Mom, P. Clarysse, N. Duchateau, Population-based personalization of geometric models of myocardial infarction, Proc. FIMH, LNCS 12738 (2021) 3–11.
[8] H. Bulluck, R. Dharmakumar, A. Arai, et al., Cardiovascular magnetic resonance in acute ST-segment-elevation myocardial infarction: recent advances, controversies, and future directions, Circulation 137 (2018) 1949–64.
[9] N. Duchateau, M. Viallon, L. Petrusca, et al., Pixel-wise statistical analysis of myocardial injury in STEMI patients with delayed enhancement MRI, Frontiers: Cardiovasc Med 10 (2023) 1136760.
[10] K. Reimer, J. Lowe, M. Rasmussen, et al., The wavefront phenomenon of ischemic cell death. 1. Myocardial infarct size vs duration of coronary occlusion in dogs, Circulation 56 (1977) 786–94.
[11] W. Li, Biomechanics of infarcted left ventricle: a review of modelling, Biomed Eng Lett 10 (2020) 387–417.
[12] M. Cerqueira, N. Weissman, V. Dilsizian, et al., Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart, Circulation 105 (2002) 539–42.
[13] M. Alessandrini, M. De Craene, O. Bernard, et al., A pipeline for the generation of realistic 3D synthetic echocardiographic sequences: methodology and open-access database, IEEE Trans Med Imaging 34 (2015) 1436–51.
[14] G. Rumindo, N. Duchateau, P. Croisille, et al., Strain-based parameters for infarct localization: evaluation via a learning algorithm on a synthetic database of pathological hearts, Proc. FIMH, LNCS 10263 (2017) 106–14.
[15] N. Duchateau, M. De Craene, P. Allain, et al., Infarct localization from myocardial deformation: prediction and uncertainty quantification by regression from a low-dimensional space, IEEE Trans Med Imaging 35 (2016) 2340–50.
[16] P. Saez, E. Kuhl, Computational modeling of acute myocardial infarction, Comput Methods Biomech Biomed Engin 19 (2016) 1107–15.
[17] Y. Skandarani, P.-M. Jodoin, A. Lalande, GANs for medical image synthesis: an empirical study, J Imaging 9 (2023) 69.
[18] T. Karras, T. Aila, S. Laine, et al., Progressive growing of GANs for improved quality, stability, and variation, Proc. ICLR (2017).
[19] G. Müller-Franzes, J. Niehues, F. Khader, et al., A multimodal comparison of latent denoising diffusion probabilistic models and generative adversarial networks for medical image synthesis, Sci Rep 13 (2023) 12098.
[20] S. Pan, T. Wang, R. Qiu, et al., 2D medical image synthesis using transformer-based denoising diffusion probabilistic model, Phys Med Biol 68 (2023) 105004.
[21] Z. Dorjsembe, H.-K. Pao, S. Odonchimed, et al., Conditional diffusion models for semantic 3D brain MRI synthesis, IEEE J Biomed Health Inform 28 (2024) 4084–4093.
[22] T. Wang, Y. Lei, Y. Fu, et al., A review on medical imaging synthesis using deep learning and its clinical applications, J Appl Clin Med Phys 22 (2020) 11–36.
[23] N. Burman, C. A. Manetti, S. V. Heymans, et al., Large-scale simulation of realistic cardiac ultrasound data with clinical appearance: methodology and open-access database, IEEE Access 12 (2024) 117040–55.
[24] S. Amirrajab, Y. Al Khalil, C. Lorenz, et al., A framework for simulating cardiac MR images with varying anatomy and contrast, IEEE Trans Med Imaging 42 (2023) 726–38.
[25] L. Belle, P. Motreff, L. Mangin, et al., Comparison of immediate with delayed stenting using the Minimalist Immediate Mechanical Intervention approach in acute ST-segment-elevation myocardial infarction: the MIMI Study, Circ Cardiovasc Interv 9 (2016) e003388.
[26] J. Tenenbaum, V. De Silva, J. Langford, A global geometric framework for nonlinear dimensionality reduction, Science 290 (2000) 2319–23.
[27] N. Duchateau, M. De Craene, M. Sitges, et al., Adaptation of multiscale function extension to inexact matching. Application to the mapping of individuals to a learnt manifold, Proc. SEE-GSI 8085 (2013) 578–86.
[28] N. Duchateau, M. Sermesant, H. Delingette, et al., Model-based generation of large databases of cardiac images: synthesis of pathological cine MR sequences from real healthy cases, IEEE Trans Med Imaging 37 (2018) 755–66.

[2] `http://humanheart-project.creatis.insa-lyon.fr/`