



ELSEVIER

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

A self-taught artificial agent for multi-physics computational model personalization



Dominik Neumann^{a,c,*}, Tommaso Mansi^b, Lucian Itu^{d,e}, Bogdan Georgescu^b,
Elham Kayvanpour^f, Farbod Sedaghat-Hamedani^f, Ali Amr^f, Jan Haas^f, Hugo Katus^f,
Benjamin Meder^f, Stefan Steidl^c, Joachim Hornegger^c, Dorin Comaniciu^b

^a Medical Imaging Technologies, Siemens Healthcare GmbH, Erlangen, Germany

^b Medical Imaging Technologies, Siemens Healthcare, Princeton, USA

^c Pattern Recognition Lab, FAU Erlangen-Nürnberg, Erlangen, Germany

^d Siemens Corporate Technology, Siemens SRL, Brasov, Romania

^e Transilvania University of Brasov, Brasov, Romania

^f Department of Internal Medicine III, University Hospital Heidelberg, Germany

ARTICLE INFO

Article history:

Received 19 January 2016

Revised 8 April 2016

Accepted 19 April 2016

Available online 21 April 2016

Keywords:

Computational modeling

Model personalization

Reinforcement learning

Artificial intelligence

ABSTRACT

Personalization is the process of fitting a model to patient data, a critical step towards application of multi-physics computational models in clinical practice. Designing robust personalization algorithms is often a tedious, time-consuming, model- and data-specific process. We propose to use artificial intelligence concepts to learn this task, inspired by how human experts manually perform it. The problem is reformulated in terms of reinforcement learning. In an off-line phase, Vito, our self-taught artificial agent, learns a representative decision process model through exploration of the computational model: it learns how the model behaves under change of parameters. The agent then automatically learns an optimal strategy for on-line personalization. The algorithm is model-independent; applying it to a new model requires only adjusting few hyper-parameters of the agent and defining the observations to match. The full knowledge of the model itself is not required. Vito was tested in a synthetic scenario, showing that it could learn how to optimize cost functions generically. Then Vito was applied to the inverse problem of cardiac electrophysiology and the personalization of a whole-body circulation model. The obtained results suggested that Vito could achieve equivalent, if not better goodness of fit than standard methods, while being more robust (up to 11% higher success rates) and with faster (up to seven times) convergence rate. Our artificial intelligence approach could thus make personalization algorithms generalizable and self-adaptable to any patient and any model.

© 2016 Published by Elsevier B.V.

1. Introduction

Computational modeling attracted significant attention in cardiac research over the last decades (Frangi et al., 2001; Noble, 2002; Hunter and Borg, 2003; Kerckhoffs et al., 2008; Clayton et al., 2011; Kuijpers et al., 2012; Krishnamurthy et al., 2013). It is believed that computational models can improve patient stratification and therapy planning. They could become the enabling tool for predicting disease course and therapy outcome, ultimately leading to improved clinical management of patients suffering from cardiomyopathies (Kayvanpour et al., 2015). A crucial prerequisite for achieving these goals is precise model personalization:

the computational model under consideration needs to be fitted to each patient. However, the high complexity of cardiac models and the often noisy and sparse clinical data still hinder this task.

A wide variety of manual and (semi-)automatic model parameter estimation approaches have been explored, including Augenstein et al. (2005); Schmid et al. (2006); Wang et al. (2009); Sermesant et al. (2009); Aguado-Sierra et al. (2010); Konukoglu et al. (2011); Aguado-Sierra et al. (2011); Delingette et al. (2012); Chabiniok et al. (2012); Xi et al. (2013); Marchesseau et al. (2013); Le Folgoc et al. (2013); Prakosa et al. (2013); Wallman et al. (2014); Zettinig et al. (2014); Neumann et al. (2014a); 2014b); Itu et al. (2014); Seegerer et al. (2015); Wong et al. (2015). Most methods aim to iteratively reduce the misfit between model output and measurements using optimization algorithms, for instance variational (Delingette et al., 2012) or filtering (Marchesseau et al., 2013)

* Corresponding author. Tel.: +4915206275313.

E-mail address: dominik.neumann@siemens.com (D. Neumann).

approaches. Applied blindly, those techniques could easily fail on unseen data, if not supervised, due to parameter ambiguity, data noise and local minima (Neumann et al., 2014a; Wallman et al., 2014; Konukoglu et al., 2011). Therefore, complex algorithms have been designed combining cascades of optimizers in a very specific way to achieve high levels of robustness, even on larger populations, i.e. 10 or more patients (Kayvanpour et al., 2015; Seegerer et al., 2015; Neumann et al., 2014b). However, those methods are often designed from tedious, trial-and-error-driven manual tuning, they are model-specific rather than generic, and their generalization to varying data quality cannot be guaranteed. On the contrary, if the personalization task is assigned to an experienced human, given enough time, he almost always succeeds in manually personalizing a model for any subject (although solution uniqueness is not guaranteed, but this is inherent to the problem).

There are several potential reasons why a human expert is often superior to standard automatic methods in terms of personalization accuracy and success rates. First, an expert is likely to have an intuition of the model's behavior from his prior knowledge of the physiology of the modeled organ. Second, knowledge about model design and assumptions, and model limitations and implementation details certainly provide useful hints on the "mechanics" of the model. Third, past personalization of other datasets allows the expert to build up experience. The combination of prior knowledge, intuition and experience enables to solve the personalization task more effectively, even on unseen data.

Inspired by humans and contrary to previous works, we propose to address the personalization problem from an artificial intelligence (AI) perspective. In particular, we apply reinforcement learning (RL) methods (Sutton and Barto, 1998) developed in the AI community to solve the parameter estimation task for computational physiological models. With its roots in control theory on the one hand, and neuroscience theories of learning on the other hand, RL encompasses a set of approaches to make an artificial agent learn from experience generated by interacting with its environment. Contrary to standard (supervised) machine learning (Bishop, 2006), where the objective is to compute a direct mapping from input features to a classification label or regression output, RL aims to learn *how* to perform tasks. The goal of RL is to compute an optimal problem-solving strategy (agent behavior), e.g. a strategy to play the game "tic-tac-toe" successfully. In the AI field, such a behavior is often represented as a policy, a mapping from states, describing the current "situation" the agent finds itself in (e.g. the current locations of all "X" and "O" on the tic-tac-toe grid), to actions, which allow the agent to interact with the environment (e.g. place "X" on an empty cell) and thus influence that situation. The key underlying principle of RL is that of reward (Kaelbling et al., 1996), which provides an objective means for the agent to judge the outcome of its actions. In tic-tac-toe, the agent receives a high, positive reward if the latest action led to a horizontal, vertical or diagonal row full of "X" marks (winning), and a negative reward (punishment) if the latest action would allow the opponent to win in his next move. Based on such rewards, the artificial agent learns an optimal winning policy through trial-and-error interactions with the environment.

RL was first applied to game (e.g. [Tesauro, 1994](#)) or simple control tasks. However, the past few years saw tremendous breakthroughs in RL for more complex, real-world problems (e.g. [Nguyen-Tuong and Peters, 2011](#); [Kveton and Theodorou, 2012](#); [Barreto et al., 2014](#)). Some noteworthy examples include [Mülling et al. \(2013\)](#), where the control entity of a robot arm learned to select appropriate motor primitives to play table tennis, and [Mnih et al. \(2015\)](#), where the authors combine RL with deep learning to train an agent to play 49 Atari games, yielding better performance than an expert in the majority of them.




Fig. 1. Overview of Vito: a self-taught artificial agent for computational model personalization, inspired by how human operators approach the personalization problem.

Motivated by these recent successes and building on our previous work (Neumann et al., 2015), we propose an RL-based personalization approach, henceforth called *Vito*, with the goal of designing a framework that can, for the first time to our knowledge, learn by itself how to estimate model parameters from clinical data while being model-independent. As illustrated in Fig. 1, first, like a human expert, Vito assimilates the behavior of the physiological model under consideration in an off-line, one-time only, data-driven exploration phase. From this knowledge, Vito learns the optimal strategy using RL (Sutton and Barto, 1998). The goal of Vito during the on-line personalization phase is then to sequentially choose actions that maximize future rewards, and therefore bring Vito to the state representing the solution of the personalization problem. To setup the algorithm, the user needs to define what observations need to be matched, the allowed actions, and a single hyper-parameter related to the desired granularity of the state-space. Then everything is learned automatically. The algorithm does not depend on the underlying model.

Vito was evaluated on three different tasks. First, in a synthetic experiment, convergence properties of the algorithm were analyzed. Then, two tasks involving real clinical data were evaluated: the inverse problem of cardiac electrophysiology and the personalization of a lumped-parameter model of whole-body circulation. The obtained results suggested that Vito can achieve equivalent (or better) goodness of fit as standard optimization methods, increased robustness and faster convergence rates.

A number of novelties and improvements over (Neumann et al., 2015) are featured in this manuscript. First, an automatic, data-driven state-space quantization method is introduced that replaces the previous manual technique. Second, the need to provide user-defined initial parameter values is eliminated by employing a new data-driven technique to initialize personalization of unseen data. Third, a stochastic personalization policy is introduced, for which the previously used standard deterministic policy is a special case. Fourth, the convergence properties are evaluated in parameter space using a synthetic personalization scenario. In addition, thorough evaluation of Vito's performance with increasing amount of training samples was conducted and personalization of the whole-body circulation model was extended to several variants involving two to six parameters. Finally, the patient database used for experimentation was extended from 28 to 83 patients for the cardiac electrophysiology experiments, and from 27 to 56 for the whole-body circulation experiments.

The remainder of this manuscript is organized as follows. [Section 2](#) presents the method. In [Section 3](#), the experiments are described and the results are presented. [Section 4](#) concludes the manuscript with a summary and discussions about potential limitations and extensions of the method.

2. Method

This section presents the reinforcement-learning (RL) framework for computational model personalization. [Section 2.1](#) introduces Markov decision process (MDP). [Section 2.2](#) defines the personalization problem and how it can be reformulated in terms of an MDP. [Section 2.3](#) describes how the artificial agent, Vito, learns

how the model behaves. Next, Section 2.4 provides details about state-space quantization, and Section 2.5 describes how the model knowledge is encoded in the form of transition probabilities. All steps mentioned so far are performed in an off-line training phase. Finally, Section 2.6 explains how the learned knowledge is applied on-line to personalize unseen data.

2.1. Model-based reinforcement learning

2.1.1. MDP definition

A crucial prerequisite for applying RL is that the problem of interest, here personalization, can be modeled as a Markov decision process (MDP). An MDP is a mathematical framework for modeling decision making when the decision outcome is partly random and partly controlled by a decision maker (Sutton and Barto, 1998). Formally, an MDP is a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where:

- \mathcal{S} is the finite set of states that describe the agent's environment, $n_{\mathcal{S}}$ is the number of states, and $s_t \in \mathcal{S}$ is the state at time t .
- \mathcal{A} is the finite set of actions, which allow the agent to interact with the environment, $n_{\mathcal{A}}$ is the number of actions, and $a_t \in \mathcal{A}$ denotes the action performed at time t .
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0; 1]$ is the stochastic transition function, where $\mathcal{T}(s_t, a_t, s_{t+1})$ describes the probability of arriving in state s_{t+1} after the agent performed action a_t in state s_t .
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the scalar reward function, where $r_{t+1} = \mathcal{R}(s_t, a_t, s_{t+1})$ is the immediate reward the agent receives at time $t + 1$ after performing action a_t in state s_t resulting in state s_{t+1} .
- $\gamma \in [0; 1]$ is the discount factor that controls the importance of future versus immediate rewards.

2.1.2. Value iteration

The value of a state, $V^*(s)$, is the expected discounted reward the agent accumulates when it starts in state s and acts optimally in each step:

$$V^*(s) = E \left\{ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s \right\}, \quad (1)$$

where $E\{\}$ denotes the expected value given the agent always selects the optimal action, and t is any time step. Note that the discount factor γ is a constant and the superscript k its exponent. V^* can be computed using value iteration (Sutton and Barto, 1998), an iterative algorithm based on dynamic programming. In the first iteration $i = 0$, let $V_i : \mathcal{S} \rightarrow \mathbb{R}$ denote an initial guess for the value function that maps states to arbitrary values. Further, let $Q_i : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denote the i th “state-action value function”-guess, which is computed as:

$$Q_i(s, a) = \sum_{s' \in \mathcal{S}} \mathcal{T}(s, a, s') [\mathcal{R}(s, a, s') + \gamma V_i(s')]. \quad (2)$$

Value iteration iteratively updates V_{i+1} from the previous Q_i :

$$\forall s \in \mathcal{S} : V_{i+1}(s) = \max_{a \in \mathcal{A}} Q_i(s, a), \quad (3)$$

until the left- and right-hand side of Eq. 3 are equal for all $s \in \mathcal{S}$; then $V^* \leftarrow V_{i+1}$ and $Q^* \leftarrow Q_{i+1}$. From this equality relation, also known as the Bellman equation (Bellman, 1957), one can obtain an optimal problem-solving strategy for the problem described by the MDP (assuming that all components of the MDP are known precisely). It is encoded in terms of a deterministic optimal policy $\pi^* : \mathcal{S} \rightarrow \mathcal{A}$:

$$\pi^*(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a), \quad (4)$$

i.e. a mapping that tells the agent in each state the optimal action to take.




Fig. 2. A computational model f is a dynamic system that maps model input parameters \mathbf{x} to model state (output) variables \mathbf{y} . The goal of personalization is to tune \mathbf{x} such that the objectives \mathbf{c} , defined as the misfit between \mathbf{y} and the corresponding measured data \mathbf{z} of a given patient, are optimized (the misfit is minimized).

2.1.3. Stochastic policy

In this work not all components of the MDP are known precisely, instead some are approximated from training data. Value iteration, however, assumes an exact MDP to guarantee optimality of the computed policy. Therefore, instead of relying on the deterministic policy π^* (Eq. 4), a generalization to stochastic policies $\tilde{\pi}^*$ is proposed here to mitigate potential issues due to approximations. Contrary to Eq. 4, where for each state only the one action with maximum Q^* -value is considered, a stochastic policy stores several candidate actions with similar high Q^* -value and returns one of them through a random process each time it is queried. To this end, the $Q^*(s, \cdot)$ -values for a given state s are first normalized:

$$\tilde{Q}_s^*(a) = \frac{Q^*(s, a) - \min_{a' \in \mathcal{A}} [Q^*(s, a')]}{\max_{a' \in \mathcal{A}} [Q^*(s, a')] - \min_{a' \in \mathcal{A}} [Q^*(s, a')]} \quad (5)$$

All actions whose normalized \tilde{Q}_s^* -value is below a threshold of $\epsilon = \frac{4}{5}$ (set empirically and used throughout the entire manuscript) are discarded, while actions with large values are stored as potential candidates. Each time the stochastic policy is queried, $a = \tilde{\pi}_s^*(s)$, it returns one of the candidate actions a selected randomly with probability proportional to its \tilde{Q}_s^* -value: $\tilde{Q}_s^*(a) / \sum_{a'} \tilde{Q}_s^*(a')$; the sum is over all candidate actions a' .

2.2. Reformulation of the model personalization problem into an MDP

2.2.1. Problem definition

As illustrated in Fig. 2, any computational model f is governed by a set of parameters $\mathbf{x} = (x_1, \dots, x_{n_x})^\top$, where n_x denotes the number of parameters. \mathbf{x} is bounded within a physiologically plausible domain Ω , and characterized by a number of n_y (observable) state variables $\mathbf{y} = (y_1, \dots, y_{n_y})^\top$. The state variables can be used to estimate \mathbf{x} . Note that some parameters may be pre-estimated or assigned fixed values. The goal of personalization is to optimize a set of n_c objectives $\mathbf{c} = (c_1, \dots, c_{n_c})^\top$. The objectives are scalars defined as $c_i = d(y_i, z_i)$, where d is a measure of misfit, and z_i denotes the patient's measured data (\mathbf{z}) corresponding to y_i . In this work $d(y_i, z_i) = y_i - z_i$. Personalization is considered successful if all user-defined convergence criteria $\boldsymbol{\psi} = (\psi_1, \dots, \psi_{n_c})^\top$ are met. The criteria are defined in terms of maximum acceptable misfit per objective: $\forall i \in \{1, \dots, n_c\} : |c_i| < \psi_i$.

2.2.2. Problem reformulation

Personalization is mapped to a Markov decision process as follows:

States: An MDP state encodes the misfit between the computed model state (outcome of forward model run) and the patient's measurements. Thus, MDP states carry the same type of information as objective vectors \mathbf{c} , yet the number of MDP states has to be finite (Section 2.1), while there are an infinite number of different objective vectors due to their continuous nature. Therefore the space of objective vectors in \mathbb{R}^{n_c} is reduced to a finite set of *representative states*: the MDP states \mathcal{S} , each $s \in \mathcal{S}$ covering a small region in that space. One of those states, $\hat{s} \in \mathcal{S}$, encodes personalization success as it is designed such that it covers exactly the region where all convergence criteria are satisfied. The goal of Vito is to learn how to reach that state.

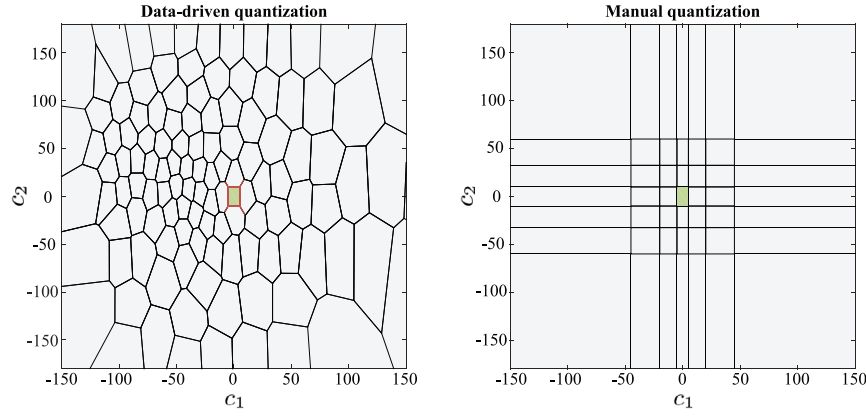


Fig. 3. State-space quantization. **Left:** Example data-driven quantization of a two-dimensional state-space into $n_s = 120$ representative states. The states are distributed according to the observed objective vectors \mathbf{c} in one of the experiments in Section 3.2. The objectives were QRS duration [ms] (c_1) and electrical axis [deg] (c_2). The center rectangle (green region) denotes the success state \hat{s} where all objectives are met ($\forall i: |c_i| < \psi_i$); see text for details. **Right:** Manual quantization as used in Neumann et al. (2015). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Actions: Vito’s actions modify the parameters \mathbf{x} to fulfill the objectives \mathbf{c} . An action $a \in \mathcal{A}$ consists in either in- or decrementing one parameter x_i by $1\times$, $10\times$ or $100\times$ a user-specified reference value δ_i with $\delta = (\delta_1, \dots, \delta_{n_x})^\top$. This empirically defined quantization of the intrinsically continuous action space yielded good results for the problems considered in this work.

Transition function: \mathcal{T} encodes the agent’s knowledge about the computational model f and is learned automatically as described in Section 2.5.

Rewards: Inspired by the “mountain car” benchmark (Sutton and Barto, 1998), the rewards are defined as always being equal to $\mathcal{R}(s, a, s') = -1$ (punishment), except when the agent performs an action resulting in personalization success, i.e. when $s' = \hat{s}$. In that case, $\mathcal{R}(\cdot, \cdot, \hat{s}) = 0$ (no punishment).

Discount factor: The large discount factor $\gamma = 0.99$ encourages policies that favor future over immediate rewards, as Vito should always prefer the long-term goal of successful personalization to short-term appealing actions in order to reduce the risk of local minima.

2.3. Learning model behavior through model exploration

Like a human operator, Vito first learns how the model “behaves” by experimenting with it. This is done through a “self-guided sensitivity analysis”. A batch of sample transitions is collected through model exploration episodes $\mathcal{E}^p = \{e_1^p, e_2^p, \dots\}$. An episode e_i^p is a sequence of $n_{e\text{-steps}}$ consecutive transitions generated from the model f and the patient p for whom the target measurements \mathbf{z}^p are known. An episode is initiated at time $t = 0$ by generating random initial model parameters \mathbf{x}_t within the physiologically plausible domain Ω . From the outputs of a forward model run $\mathbf{y}_t = f(\mathbf{x}_t)$, the misfits to the patient’s corresponding measurements are computed, yielding the objectives vector $\mathbf{c}_t = d(\mathbf{y}_t, \mathbf{z}^p)$. Next, a random exploration policy π_{rand} that selects an action according to a discrete uniform probability distribution over the set of actions is employed. The obtained $a_t \in \mathcal{A}$ is then applied to the current parameter vector, yielding modified parameter values $\mathbf{x}_{t+1} = a_t(\mathbf{x}_t)$. From the output of the forward model run $\mathbf{y}_{t+1} = f(\mathbf{x}_{t+1})$ the next objectives \mathbf{c}_{t+1} are computed. The next action a_{t+1} is then selected according to π_{rand} , and this process is repeated $n_{e\text{-steps}} - 1$ times. Hence, each episode can be seen as a set of consecutive tuples:

$$e = \{(\mathbf{x}_t, \mathbf{y}_t, \mathbf{c}_t, a_t, \mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{c}_{t+1}), \quad t = 0, \dots, n_{e\text{-steps}} - 1\}. \quad (6)$$

In this work, $n_{e\text{-steps}} = 100$ transitions are created in each episode as a trade-off between sufficient length of an episode to cover a

real personalization scenario and sufficient exploration of the parameter space.

The model is explored with many different training patients and the resulting episodes are combined into one large training episode set $\mathcal{E} = \bigcup_p \mathcal{E}^p$. The underlying hypothesis (verified in experiments) is that the combined \mathcal{E} allows to cancel out peculiarities of individual patients, i.e. to abstract from patient-specific to model-specific knowledge.

2.4. From computed objectives to representative MDP state

As mentioned above, the continuous space of objective vectors is quantized into a finite set of representative MDP states \mathcal{S} . A data-driven approach is proposed. First, all objective vectors observed during training are clustered according to their distance to each other. Because the ranges of possible values for the individual objectives can vary significantly depending on the selected measurements (due to different types of measurements, different units, etc.), the objectives should be normalized during clustering to avoid bias towards objectives with relatively large typical values. In this work the distance measure performs implicit normalization to account for these differences: the distance between two objective vectors $(\mathbf{c}_1, \mathbf{c}_2)$ is defined relative to the inverse of the convergence criteria $\boldsymbol{\psi}$:

$$\|\mathbf{c}_1 - \mathbf{c}_2\|_{\boldsymbol{\psi}} = \sqrt{(\mathbf{c}_1 - \mathbf{c}_2)^\top \text{diag}(\boldsymbol{\psi})^{-1} (\mathbf{c}_1 - \mathbf{c}_2)}, \quad (7)$$

where $\text{diag}(\boldsymbol{\psi})^{-1}$ denotes a diagonal matrix with $(\frac{1}{\psi_1}, \frac{1}{\psi_2}, \dots)$ along its diagonal. The centroid of a cluster is the centroid of a representative state. In addition, a special “success state” \hat{s} representing personalization success is created, which covers the region in state-space where all objectives are met: $\forall i: |c_i| < \psi_i$. The full algorithm is described in Appendix A. Finally, an operator $\phi: \mathbb{R}^{2c} \rightarrow \mathcal{S}$ that maps continuous objective vectors \mathbf{c} to representative MDP states is introduced:

$$\phi(\mathbf{c}) = \arg \min_{s \in \mathcal{S}} \|\mathbf{c} - \boldsymbol{\xi}_s\|_{\boldsymbol{\psi}} \quad (8)$$

where $\boldsymbol{\xi}_s$ denotes the centroid corresponding to state s . For an example state-space quantization see Fig. 3.

2.5. Transition function as probabilistic model representation

In this work, the stochastic MDP transition function \mathcal{T} encodes the agent’s knowledge about the computational model f . It is learned from the training data \mathcal{E} . First, the individual samples $(\mathbf{x}_t, \mathbf{y}_t, \mathbf{c}_t, a_t, \mathbf{x}_{t+1}, \mathbf{y}_{t+1}, \mathbf{c}_{t+1})$ are converted to state-action-state transition tuples $\hat{\mathcal{E}} = \{(s, a, s')\}$, where $s = \phi(\mathbf{c}_t)$, $a = a_t$ and




Fig. 4. Vito's probabilistic on-line personalization phase. See text for details.

$s' = \phi(\mathbf{c}_{t+1})$. Then, \mathcal{T} is approximated from statistics over the observed transition samples:

$$\mathcal{T}(s, a, s') = \frac{|\{(s, a, s') \in \hat{\mathcal{E}}\}|}{\sum_{s'' \in \mathcal{S}} |\{(s, a, s'') \in \hat{\mathcal{E}}\}|}, \quad (9)$$

where $|\{\cdot\}|$ denotes the cardinality of the set $\{\cdot\}$. If n_S and n_A are large compared to the total number of samples it may occur that some state-action combinations are not observed: $|\{(s, a, \cdot) \in \hat{\mathcal{E}}\}| = 0$. In that case uniformity is assumed: $\forall s'' \in \mathcal{S} : \mathcal{T}(s, a, s'') = 1/n_S$.

\mathcal{M} is now fully defined. Value iteration (Section 2.1) is applied and the stochastic policy $\tilde{\pi}_\epsilon^*$ is computed, which completes the off-line phase.

2.6. On-line model personalization

On-line personalization, as illustrated in Fig. 4, can be seen as a two-step procedure. First, Vito initializes the personalization of unseen patients from training data. Second, Vito relies on the computed policy $\tilde{\pi}_\epsilon^*$ to guide the personalization process.

2.6.1. Data-driven initialization

Good initialization can be decisive for a successful personalization. Vito's strategy is to search for forward model runs in the training database \mathcal{E} for which the model state $f(\mathbf{x}) = \mathbf{y} \approx \mathbf{z}^p$ is similar to the patient's measurements. To this end, Vito examines all parameters $\Xi = \{\mathbf{x} \in \mathcal{E} \mid f(\mathbf{x}) \approx \mathbf{z}^p\}$ that yielded model states similar to the patient's measurements. Due to ambiguities induced by the different training patients, data noise and model assumptions, Ξ could contain significantly dissimilar parameters. Hence, picking a single $\mathbf{x} \in \Xi$ might not yield the best initialization. Analyzing Ξ probabilistically instead helps Vito to find likely initialization candidates. The details of the initialization procedure are described in Appendix B. Given the patient's measurements \mathbf{z}^p , the procedure outputs a list of initialization candidates $\mathcal{X}_0 = (\mathbf{x}'_0, \mathbf{x}''_0, \dots)$. The list is sorted by likelihood with the first element, \mathbf{x}'_0 , being the most likely one.

2.6.2. Probabilistic personalization

The first personalization step initializes the model parameter vector \mathbf{x}_0 with the most likely among all initialization candidates, $\mathbf{x}_0 \in \mathcal{X}_0$ (see previous section for details). Then, as illustrated in Fig. 4, Vito computes the forward model $\mathbf{y}_0 = f(\mathbf{x}_0)$ and the misfit between the model output and the patient's measurements $\mathbf{c}_0 = d(\mathbf{y}_0, \mathbf{z}^p)$ to derive the first state $s_0 = \phi(\mathbf{c}_0)$. Given s_0 , Vito decides from its policy the first action to take $a_0 = \tilde{\pi}_\epsilon^*(s_0)$, and walks through state-action-state sequences to personalize the computational model f by iteratively updating the model parameters through MDP actions. Bad initialization could lead to oscillations between states as observed in previous RL works (Kveton and Theodorou, 2012; Neumann et al., 2015). Therefore, upon detection of an oscillation, which is done by monitoring the parameter traces to detect recurring sets of parameter values, the personalization is re-initialized at the second-most-likely $\mathbf{x}_0 \in \mathcal{X}_0$, etc. If all $|\mathcal{X}_0|$ initialization candidates have been tested, a potential re-initialization defaults to fully random within the physiologically

plausible parameter domain Ω . The process terminates once Vito reaches state \hat{s} (success), or when a pre-defined maximum number of iterations is reached (failure).

3. Experiments

Vito was applied to a synthetic parameter estimation problem and to two challenging problems involving real clinical data: personalization of a cardiac electrophysiology (EP), and a whole-body-circulation (WBC) model. All experiments were conducted using leave-one-out cross-validation. The numbers of datasets and transition samples used for the different experiments are denoted n_{datasets} and n_{samples} , respectively.

3.1. Synthetic experiment: the Rosenbrock function

First, Vito was employed in a synthetic scenario, where the ground-truth model parameters were known. The goals were to test the ability of Vito to optimize cost functions generically, and to directly evaluate the performance in the parameter space.

3.1.1. Forward model description

The Rosenbrock function (Rosenbrock, 1960), see Fig. 5, left panel, is a non-convex function that is often used to benchmark optimization algorithms. It was treated as the forward model in this experiment:

$$f^\alpha(x_1, x_2) = (\alpha - x_1)^2 + 100 \cdot (x_2 - x_1^2)^2, \quad (10)$$

where $\mathbf{x} = (x_1, x_2)^\top$ were the model parameters to estimate for any α , and $f^\alpha : \Omega \rightarrow \mathbb{R}$. As described in Section 2.2.2, each of Vito's actions $a \in \mathcal{A}$ in- or decrements a parameter value by multiples (1×, 10×, 100×) of parameter-specific reference values. The reference values were set to $\delta = (0.01, 0.01)^\top$, determined as 0.1% of the defined admissible parameter space per dimension, $\Omega = [-5; 5]^2$. The parameter $\alpha \in \mathbb{R}$ defines a family of functions $\{f^\alpha\}$. The goal was to find generically $\arg \min_{x_1, x_2} f^\alpha(x_1, x_2)$.

The Rosenbrock function has a unique global minimum at $\mathbf{x} = (\alpha, \alpha^2)^\top$, where both terms $T_1 = (\alpha - x_1)$ and $T_2 = (x_2 - x_1^2)$ evaluate to 0. The personalization objectives were therefore defined as $\mathbf{c} = (|T_1 - 0|, |T_2 - 0|)^\top$, with the measured data $\mathbf{z} = (0, 0)^\top$ were zero for both objectives and the computed data $\mathbf{y} = (T_1, T_2)^\top$. The convergence criteria were set empirically to $\psi = (0.05, 0.05)^\top$.

3.1.2. Evaluation

Vito was evaluated on $n_{\text{datasets}} = 100$ functions f^α with randomly generated $\alpha \in [-2, 2]$. In the off-line phase, for each function, $n_{\text{samples}} = 10 \cdot n_{\text{e-steps}} = 1000$ samples, i.e. ten training episodes, each consisting in $n_{\text{e-steps}} = 100$ transitions (Section 2.3), were generated to learn the policy. The number of representative states was set to $n_S = 100$. To focus on Vito's on-line personalization capabilities, both the data-driven initialization and the re-initialization on oscillation (Section 2.6) were disabled. In total, 441 experiments with different initializations (sampled on a 21×21




Fig. 5. Synthetic experiment. **Left:** Contour plot of the Rosenbrock function $f^{\alpha=1}$ with global minimum at $\mathbf{x} = (1, 1)^\top$ (red dot). The color scale is logarithmic for visualization purposes: the darker, the lower the function value. **Mid:** Maximum L_2 -error in parameter space after personalization over all functions for varying initial parameter values. See text for details. Yellow represents errors ≥ 5 (maximum observed error ≈ 110). **Right:** Same as mid panel, except the extended action set was used. The red dots are the 100 ground-truth parameters $\mathbf{x} = (\alpha, \alpha^2)^\top$ generated for random α . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

uniform grid spanned in Ω) were conducted. For each experiment all 100 functions were personalized using leave-one-family-function-out cross validation, and the error value from the function exhibiting the maximum L_2 -error (worst-case scenario) between ground-truth (α, α^2) and estimated parameters was plotted. As one can see from the large blue region in Fig. 5, mid panel, for the majority of initial parameter values Vito always converged to the solution (maximum L_2 -error < 0.25 ; the maximum achievable accuracy depended on the specified convergence criteria ψ and on the reference values δ , which “discretized” the parameter space). However, especially for initializations far from the ground-truth (near border regions of Ω), Vito was unable to personalize some functions properly, which was likely due to the high similarity of the Rosenbrock function shape in these regions.

To investigate this issue, the experiment was repeated after additional larger parameter steps were added to the set of available actions: $\mathcal{A}' = \mathcal{A} \cup \{\pm 500\delta_1; \pm 500\delta_2\}$. As shown in Fig. 5, right panel, Vito could now personalize successfully starting from any point in Ω . The single spot with larger maximum error (bright spot at approximately $\mathbf{x} = (-1, 2)^\top$) can be explained by Vito’s stochastic behavior: Vito may have become *unlucky* if it selected many unfavorable actions in sequence due to the randomness introduced by the stochastic policy. Enabling re-initialization on oscillation solved this issue entirely. In conclusion, this experiment showed that Vito can learn how to minimize a cost function generically.

3.2. Personalization of cardiac electrophysiology model

Vito was then tested in a scenario involving a complex model of cardiac electrophysiology coupled with 12-lead ECG. Personalization was performed for real patients from actual clinical data. A total of $n_{\text{datasets}} = 83$ patients were available for experimentation. For each patient, the end-diastolic bi-ventricular anatomy was segmented from short-axis cine magnetic resonance imaging (MRI) stacks as described in Zheng et al. (2008). A tetrahedral anatomical model including myofibers was estimated and a torso atlas affinely registered to the patient based on MRI scout images. See Zettinig et al. (2014) for more details.

3.2.1. Forward model description

The depolarization time at each node of the tetrahedral anatomical model was computed using a shortest-path graph-based algorithm, similar to the one proposed in Wallman et al. (2012). Tissue anisotropy was modeled by modifying the edge costs to take into account fiber orientation. A time-varying voltage map was then derived according to the depolarization time: at a given time t , mesh nodes whose depolarization time was higher than t were assigned a trans-membrane potential of -70 mV, 30 mV

otherwise. The time-varying potentials were then propagated to a torso model where 12-lead ECG acquisition was simulated, and QRS duration (QRSd) and electrical axis (EA) were derived (Zettinig et al., 2014). The model was controlled by the conduction velocities (in m/s) of myocardial tissue and left and right Purkinje network: $\mathbf{x} = (v_{\text{Myo}}, v_{\text{LV}}, v_{\text{RV}})^\top$. The latter two domains were modeled as fast endocardial conducting tissue. The admissible parameter space Ω was set to $[200; 1000]$ for v_{Myo} and $[500; 5000]$ for both v_{LV} and v_{RV} . Reference increment values to build the action set \mathcal{A} were set to $\delta = (5, 5, 5)^\top$ m/s for the three model parameters. The goal of EP personalization was to estimate \mathbf{x} from the measured QRSd and EA. Accounting for uncertainty in the measurements and errors in the model, a patient was considered personalized if QRSd and EA misfits were below $\psi = (5 \text{ ms}, 10^\circ)^\top$, respectively.

3.2.2. Number of representative states

In contrast to Neumann et al. (2015), where state-space quantization required manual tuning of various threshold values, the proposed approach relies on a single hyper-parameter only: n_S , the number of representative states (Section 2.4). To specify n_S , eight patients were selected for scouting. Exhaustive search was performed for $n_S \in \{10, 20, \dots, 490, 500\}$ representative states. The goodness of a given configuration was evaluated based on the success rate (relative number of successfully personalized cases according to convergence criteria ψ) over five independent, consecutive, leave-one-patient-out cross-validated personalization runs of the eight patients. Furthermore, the average number of required forward model runs was considered. To this end, 100 training episodes ($100 \cdot n_{e\text{-steps}} = 10^4$ transition samples) per patient were generated for each personalization run as described in Section 2.3. As one can see from Fig. 6, good performance was achieved from 50 to 300 representative states.

The large range of well performing n_S indicates a certain level of robustness with respect to that hyper-parameter. A slight performance peak at 120 representative states was observed. Therefore, $n_S = 120$ was selected for further experimentation as compromise between maintaining a low number of states and sufficient state granularity. An example quantization with $n_S = 120$ is visualized in Fig. 3. The eight scouting datasets were discarded for the following experiments to avoid bias in the analysis.

3.2.3. Reference methods

Vito’s results were compared to two standard personalization methods based on BOBYQA (Powell, 2009), a widely-used gradient-free optimizer known for its robust performance and fast convergence. The first approach, “BOBYQA simple”, mimicked the most basic estimation setup, where only the minimum level of model and problem knowledge were assumed. The objective function was




Fig. 6. Hyper-parameter scouting. Vito's performance for varying number of representative states n_s on eight scouting datasets. The solid and dashed curves represent success rate and average number of forward runs until convergence, respectively, aggregated over five personalization runs with varying training data.

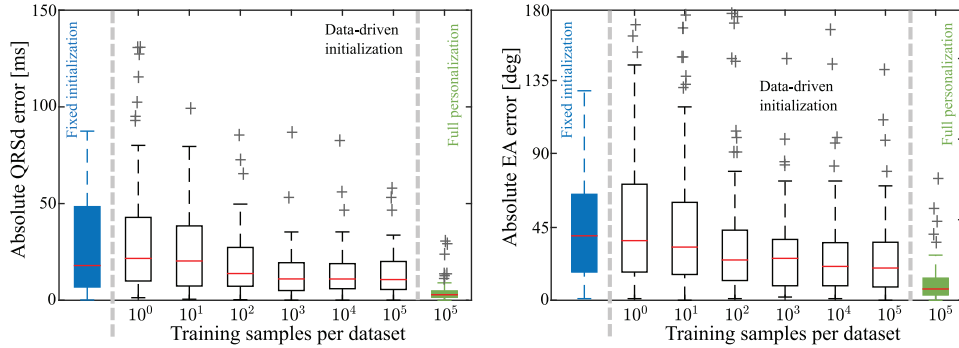


Fig. 7. Absolute errors over all patients after initialization with fixed parameter values (blue), after data-driven initialization for increasing amount of training data (white), and after full personalization with Vito (green). Data-driven initialization yielded significantly reduced errors if sufficient training data were available ($>10^2$) compared to initialization with fixed values. Full personalization further reduced the errors by a significant margin. The red bar and the box edges indicate the median absolute error, and the 25 and 75 percentiles, respectively. **Left:** QRS duration errors. **Right:** Electrical axis errors. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the sum of absolute QRSd and EA errors: $\sum_{i=1}^n |c_i|$. It was minimized in a single optimizer run where all three parameters in \mathbf{x} were tuned simultaneously. The algorithm terminated once all convergence criteria ψ were satisfied (success) or if the number of forward model evaluations exceeded 100 (failure). The second approach, “BOBYQA cascade”, implemented an advanced estimator with strong focus on robustness, which computed the optimum parameters in a multi-step iterative fashion. It is based on Seegerer et al. (2015), where tedious manual algorithm and cost function tuning was performed on a subset of the data used in this manuscript. In a first step, the myocardial conduction velocity was tuned to yield good match between computed and measured QRS duration. Second, left and right endocardial conduction velocities were optimized to minimize electrical axis error. Both steps were repeated until all parameter estimates were stable.

To remove bias towards the choice of initial parameter values, for each of the two methods all datasets were personalized 100 times with different random initializations within the range of physiologically plausible values Ω . The differences in performance were striking: only by changing the initialization, the number of successfully personalized cases varied from 13 to 37 for BOBYQA simple, and from 31 to 51 for BOBYQA cascade (variability of more than 25% of the total number of patients). These results highlight the non-convexity of the cost function to minimize.

3.2.4. Full personalization performance

First, Vito's overall performance was evaluated. The full personalization pipeline consisting in off-line learning, initialization, and on-line personalization was run on all patients with leave-one-patient-out cross-validation using 1000 training episodes ($n_{\text{samples}} = 1000 \cdot n_{\text{e-steps}} = 10^5$ transition samples) per patient. The maximum number of iterations was set to 100. The green box plots in the two panels of Fig. 7 summarize the results.

The mean absolute errors were 4.1 ± 5.6 ms and $12.4 \pm 13.3^\circ$ in terms of QRSd and EA, respectively, a significant improvement over the residual error after initialization. In comparison to the reference methods, the best BOBYQA simple run yielded absolute errors of 4.4 ± 10.8 ms QRSd and $15.5 \pm 18.6^\circ$ EA on average, and the best BOBYQA cascade run 0.1 ± 0.2 ms QRSd and $11.2 \pm 15.8^\circ$ EA, respectively. Thus, in terms of EA error all three methods yielded comparable performance, and while BOBYQA simple and Vito performed similarly in terms of QRSd, BOBYQA cascade outperformed both in this regard. However, considering success rates, i.e. successfully personalized patients according to the defined convergence criteria (ψ) divided by total number of patients, both the performance of Vito (67%) and BOBYQA cascade (68%) were equivalent, while BOBYQA simple reached only 49% or less. In terms of runtime, i.e. average number of forward model runs until convergence, Vito (31.8) almost reached the high efficiency of BOBYQA simple (best: 20.1 iterations) and clearly outperformed BOBYQA cascade (best: 86.6 iterations), which means Vito was $\approx 2.5\times$ faster.

3.2.5. Residual error after initialization

A major advantage over standard methods such as the two BOBYQA approaches is Vito's automated, data-driven initialization method (Section 2.6.1), which eliminates the need for user-provided initial parameter values. To evaluate the utility of this step, personalization using Vito was stopped directly after initialization (the most likely \mathbf{x}_0 was used) and the errors in terms of QRSd and EA resulting from a forward model run f with the computed initial parameter values were quantified. This experiment was repeated for increasing number of transition samples per dataset: $n_{\text{samples}} = 10^0 \dots 10^5$, and the results were compared to the error after initialization when fixed initial values were used (the initialization of the best performing BOBYQA experiment was used). As one can see from Fig. 7, with increasing amount of

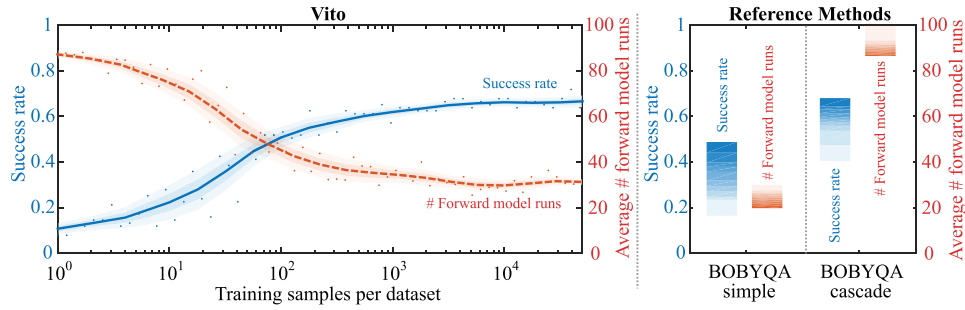


Fig. 8. EP personalization results. Personalization success rate in blue and average number of iterations in red. **Left:** Vito's performance for increasing number of training transition samples per dataset. Each dot represents results from one experiment (cross-validated personalization of all 75 datasets), solid/dashed line is low-pass filtered mean, shaded areas represent $0.5\times$ and $1\times$ standard deviation. **Right:** Performance of both reference methods. Each shade represents 10% of the results, sorted by performance.

training data both errors decreased notably. As few as 10^2 transitions per dataset already provided more accurate initialization than the best tested fixed initial values. Thus, not only does this procedure simplify the setup of Vito for new problems (no user-defined initialization needed), this experiment showed that it can reduce initial errors by a large margin, even when only few training transitions were available. It should be noted that Vito further improves the model fit in its normal operating mode (continue personalization after initialization), as shown in the previous experiment.

3.2.6. Convergence analysis

An important question in any RL application relates to the amount of training needed until convergence of the artificial agent's behavior. For Vito in particular, this translates to the amount of transition samples required to accurately estimate the MDP transition function \mathcal{T} to compute a solid policy on the one hand, and to have enough training data for reliable parameter initialization on the other hand. To this end, Vito's overall performance (off-line learning, initialization, personalization) was evaluated for varying number of training transition samples per dataset. As one can see from the results in Fig. 8, with increasing amount of training data the performance increased, suggesting that the learning process was working properly. Even with relatively limited training data of only $n_{\text{samples}} = 10^2$ samples per patient, Vito outperformed the best version of BOBYQA simple (49% success rate). Starting from $n_{\text{samples}} \approx 3000$, a plateau at $\approx 66\%$ success rate was reached, which remained approximately constant until the maximum tested number of samples. This was almost on par with the top BOBYQA cascade performance (68% success rate). Also the runtime performance increased with more training data. For instance, Vito's average number of iterations was 36.2 at 10^3 samples, 31.5 at 10^4 samples, or 31.8 at 10^5 samples.

These results suggested that not only Vito can achieve similar performance as an advanced, manually engineered method, but also the number of required training samples was not excessive. In fact, a rather limited and thus well manageable amount of data, which can be computed in a reasonable time-frame, sufficed.

3.3. Personalization of whole-body circulation model

Next, Vito was asked to personalize a lumped-parameter whole-body circulation (WBC) model from pressure catheterization and volume data. A subset of $n_{\text{datasets}} = 56$ patients from the EP experiments were used for experimentation. The discrepancy was due to missing catheterization data for some patients, which was required for WBC personalization only. For each patient, the bi-ventricular anatomy was segmented and tracked from short-axis cine MRI stacks throughout one full heart cycle using shape-constraints,

Table 1

WBC parameters \mathbf{x} , their default values and domain Ω . The last column denotes the experiment setups in which a parameter was personalized (e.g. "5": parameter was among the estimated parameters in 5p experiment). Default values were used in experiments where the respective parameters were not personalized.

\mathbf{x}	Default value	Ω	Setups
Initial volume	400 mL	[200; 1000] mL	6, 5, 3, 2
LV max. elastance	2.4 mmHg/mL	[0.2; 5] mmHg/mL	6, 5, 3, 2
Aortic resistance	1100 g/(cm ⁴ s)	[500; 2500] g/(cm ⁴ s)	6, 5, 3
Aortic compliance	$1.4 \cdot 10^9$ cm ⁴ s ² /g	[0.5; 6] $\cdot 10^9$ cm ⁴ s ² /g	6, 5
Dead volume	10 mL	[-50; 500] mL	6, 5
Time to E_{max}	300 ms	[100; 600] ms	6

learned motion models and diffeomorphic registration (Wang et al., 2013). From the time-varying endocardial meshes, ventricular volume curves were derived. Manual editing was performed whenever necessary.

3.3.1. Forward model description

The WBC model to personalize was based on Itu et al. (2014). It contained a heart model (left ventricle (LV) and atrium, right ventricle and atrium, valves), the systemic circulation (arteries, capillaries, veins) and the pulmonary circulation (arteries, capillaries, veins). Time-varying elastance models were used for all four chambers of the heart. The valves were modeled through a resistance and an inertance. A three-element Windkessel model was used for the systemic and pulmonary arterial circulation, while a two-element Windkessel model was used for the systemic and pulmonary venous circulation. We refer the reader to Itu et al. (2014); Neumann et al. (2015); Westerhof et al. (1971) for more details. Personalization was performed with respect to the patient's heart rate as measured during catheterization.

The goal of this experiment was to compare Vito's personalization performance for the systemic part of the model in setups with increasing number of parameters to tune and objectives to match. To this end, Vito was employed on setups with two to six parameters (2p, 3p, 5p, 6p): initial blood volume, LV maximum elastance, time until maximum elastance is reached, total aortic resistance and compliance, and LV dead volume. The reference values δ to define Vito's allowed actions \mathcal{A} were set to .5% of the admissible parameter range Ω for each individual parameter, see Table 1 for details. The personalization objectives were MRI-derived end-diastolic and end-systolic LV volume, ejection time (time duration during which the aortic valve is open and blood is ejected), and peak-systolic, end-diastolic, and mean aortic blood pressures as measured during cardiac catheterization, see Fig. 9. To account for measurement noise, personalization was considered successful if the misfits per objective were below acceptable threshold values ψ as listed in Table 2.

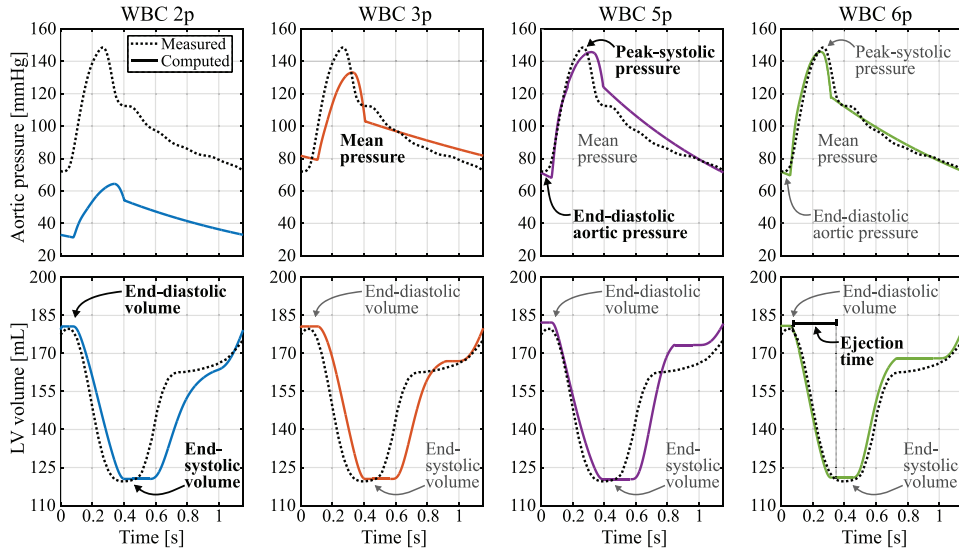


Fig. 9. Goodness of fit in terms of time-varying LV volume and aortic pressure for Vito personalizing an example patient based on the different WBC setups. The added objectives per setup are highlighted in the respective column. With increasing number of parameters and objectives Vito manages to improve the fit between model and measurements.

Table 2

WBC objectives c , their convergence criteria ψ and range of measured values in the patient population used for experimentation.

c	ψ	Measured range	Setups
End-diastolic LV volume	20 mL	[129; 647] mL	6, 5, 3, 2
End-systolic LV volume	20 mL	[63; 529] mL	6, 5, 3, 2
Mean aortic pressure	10 mmHg	[68; 121] mmHg	6, 5, 3
Peak-systolic aortic pressure	10 mmHg	[83; 182] mmHg	6, 5
End-diastolic aortic pressure	10 mmHg	[48; 99] mmHg	6, 5
Ejection time	50 ms	[115; 514] ms	6

3.3.2. Number of representative states

Along the same lines as Section 3.2.2, the hyper-parameter for state-space quantization was tuned based on the eight scouting patients. The larger the dimensionality of the state-space, the more representative states were needed to yield good performance. In particular, for the different WBC setups, the numbers of representative states (n_S) yielding the best scouting performance were 70, 150, 400 and 600 for the 2p, 3p, 5p and 6p setup, respectively. The scouting datasets were discarded for the following experiments.

3.3.3. Reference method

A gradient-free optimizer (Lagarias et al., 1998) based on the simplex method was used to benchmark Vito. The objective function was the sum of squared differences between computed and measured values, weighted by the inverse of the convergence criteria to counter the different ranges of objective values (e.g. due to different types of measurements and different units): $\|c\|_\psi$ (Eq. 7). Compared to non-normalized optimization, the algorithm converged up to 20% faster and success rates increased by up to 8% under otherwise identical conditions. Personalization was terminated once all convergence criteria were satisfied (success), or when the maximum number of iterations was reached (failure). To account for the increasing complexity of optimization with increasing number of parameters n_x , the maximum number of iterations was set to $50 \cdot n_x$ for the different setups.

As one can see from Fig. 10, right panels, with increasing number of parameters to be estimated, the performance in terms of success rate and number of forward model runs decreased slightly. This is expected as the problem becomes harder. To suppress bias originating from (potentially poor) initialization, the reference

method was run 100 times per setup (as in EP experiments), each time with a different, randomly generated set of initial parameter values. The individual performances varied significantly for all setups.

3.3.4. Convergence analysis

For each WBC setup the full Vito personalization pipeline was evaluated for increasing training data ($n_{\text{samples}} = 10^0 \dots 10^5$) using leave-one-patient-out cross-validation. The same iteration limits as for the reference method were used. The results are presented in Fig. 10, left panels. With increasing data, Vito's performance, both in terms of success rate and run-time (iterations until convergence), increased steadily until reaching a plateau. As one would expect, the more complex the problem, i.e. the more parameters and objectives involved in the personalization, the more training data was needed to reach the same level of performance. For instance, Vito reached 80% success rate with less than $n_{\text{samples}} = 50$ training samples per dataset in the 2p setup, whereas almost $90 \times$ as many samples were required to achieve the same performance in the 6p setup.

Compared to the reference method, given enough training data, Vito reached equivalent or better success rates (e.g. up to 11% higher success rate for 6p) while significantly outperforming the reference method in terms of run-time performance. In the most basic setup (2p), if $n_{\text{samples}} \geq 10^3$, Vito converged after 3.0 iterations on average, while the best reference method run required 22.6 iterations on average, i.e. Vito was seven times faster. For the more complex setups (3p, 5p, 6p), the speed-up was not as drastic. Yet, in all cases Vito outperformed even the best run of the reference method by a factor of 1.8 or larger.

4. Conclusion

4.1. Summary and discussion

In this manuscript, a novel personalization approach called Vito has been presented. To our knowledge, it is the first time that biophysical model personalization is addressed using artificial intelligence concepts. Inspired by how humans approach the personalization problem, Vito first learns the characteristics of the computational model under consideration using a data-driven approach. This knowledge is then utilized to learn how to

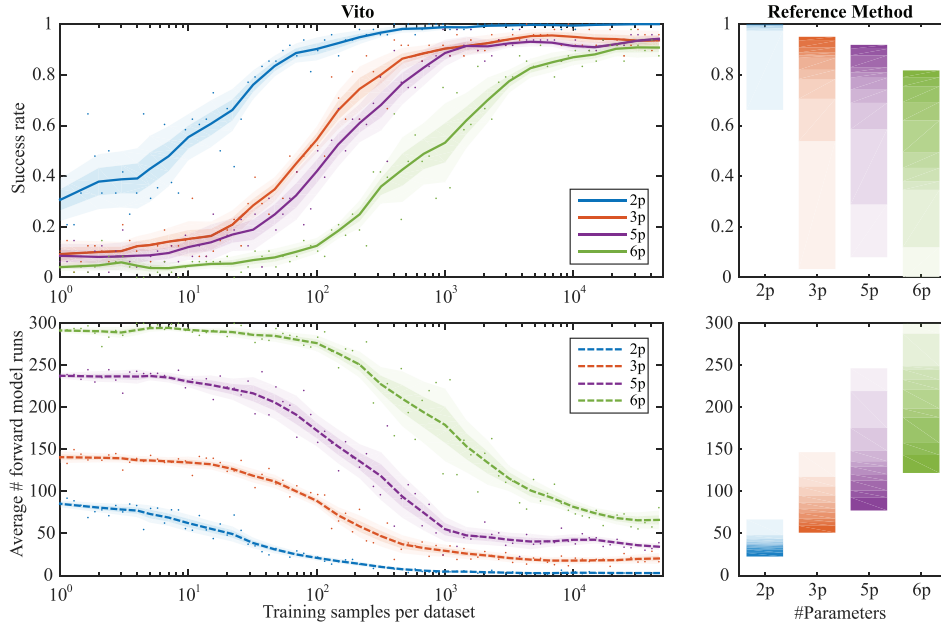


Fig. 10. WBC model personalization results (top: success rate, bottom: average number of forward model runs until convergence) for various estimation setups (different colors), see text for details. **Left:** Vito’s performance for increasing number of training transition samples per dataset. Each dot represents results from one experiment (cross-validated personalization of all 48 datasets), solid/dashed lines are low-pass filtered mean, shaded areas represent $0.5 \times$ and $1 \times$ standard deviation. **Right:** Performance of reference method. Each shade represents 10% of the results, sorted by performance; darkest shade: best 10%. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3

Applications considered in this manuscript described in terms of the number of parameters (n_x), objectives (n_c) and datasets (n_{datasets}) used for experimentation (in brackets: excluding scouting patients, if applicable); and Vito’s hyper-parameters: the number of representative MDP states (n_s) and the number of actions per parameter (n_A/n_x). The last column (n_{plateau}) denotes the approximate number of samples needed to reach the performance “plateau” (see convergence analyses in Section 3.2.6 and Section 3.3.4).

Application	n_x	n_c	n_{datasets}	n_s	n_A/n_x	n_{plateau}
Rosenbrock	2	2	100	100	6	<i>n/a</i>
Rosenbrock ext.	2	2	100	100	8	<i>n/a</i>
EP	3	2	83 (75)	120	6	3 000
WBC 2p	2	2	56 (48)	70	6	450
WBC 3p	3	3	56 (48)	150	6	2 000
WBC 5p	5	5	56 (48)	400	6	3 500
WBC 6p	6	6	56 (48)	600	6	20 000

personalize the model using reinforcement learning. Vito is generic in the sense that it requires only minimal and intuitive user input (parameter ranges, authorized actions, number of representative states) to learn by itself how to personalize a model.

Vito was applied to a synthetic scenario and to two challenging personalization tasks in cardiac computational modeling. The problem setups and hyper-parameter configurations are listed in Table 3. In most setups the majority of hyper-parameters were identical and only few (n_s) required manual tuning, suggesting good generalization properties of Vito. Another key result was that Vito was up to 11% more robust (higher success rates) compared to standard personalization methods. Vito’s ability to generalize the knowledge obtained from a set of training patients to personalize unseen patients was shown as all experiments reported in this manuscript were based on cross-validation. Furthermore, Vito’s robustness against training patients for whom we could not find a solution was tested. In particular, for about 20% of the patients, in none of the electrophysiology experiments in Section 3.2 any personalization (neither Vito nor the reference methods) could pro-

duce a result that satisfied all convergence criteria. Hence, for some patients no solution may exist under the given electrophysiology model configuration.¹ Still, all patients were used to train Vito, and surprisingly Vito was able to achieve equivalent success rate as the manually engineered personalization approach for cardiac EP.

Generating training data could be considered Vito’s computational bottleneck. However, training is *i*) performed off-line and one-time only, and *ii*) it is independent for each training episode and each patient. Therefore, large computing clusters could be employed to perform rapid training by parallelizing this phase. On-line personalization, on the contrary, is not parallelizable in its current form: the parameters for each forward model run depend on the outcome of the previous iteration. Since the forward computations are the same for every “standard” personalization method (not including surrogate-based approaches), the number of forward model runs until convergence was used for benchmarking: Vito was up to seven times faster compared to the reference methods. The on-line overhead introduced by Vito (convert data into an MDP state, then query policy) is negligible.

As such, Vito could become a unified framework for personalization of any computational physiological model, potentially eliminating the need for an expert operator with in-depth knowledge to design and engineer complex optimization procedures.

4.2. Challenges and outlook

Important challenges still remain, such as the incorporation of continuous actions, the definition of states and their quantization. In this work we propose a data-driven state-space quantization strategy. Contrary to Neumann et al. (2015), where a threshold-based state-quantization involving several manually tuned threshold values (Fig. 3) was employed, the new method is based on a single hyper-parameter only: the number of representative states.

¹ Potential solution non-existence may be due to possibly invalid assumptions of the employed EP model for patients with complex pathologies.

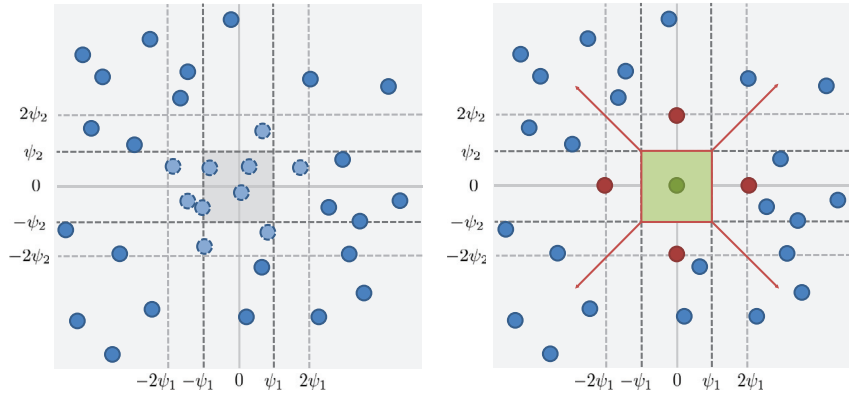


Fig. A.11. Preprocessing of k -means input data to enforce the success state $\hat{\mathcal{S}}$. **Left:** Continuous state-space with observed objective vectors \mathbf{c} (blue points). The points with dashed outline will be canceled out. **Right:** Delineation of $\hat{\mathcal{S}}$ in green, enforced by inserted vectors (green / red points) with large weights. See text for details. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Although it simplifies the setup of Vito, this quantization strategy may still not be optimal, especially if only little training data is available. Therefore, advanced approaches for continuous reinforcement learning with value function approximation (Sutton and Barto, 1998; Mnih et al., 2015) could be integrated to fully circumvent quantization issues.

At the same time, such methods could improve Vito’s scalability towards high-dimensional estimation tasks. In this work we showed that Vito can be applied to typical problems emerging in cardiac modeling, which could be described as medium-scale problems with moderate number of parameters to personalize and objectives to match. In unreported experiments involving >10 parameters, however, Vito could no longer reach satisfactory performance, which is likely due to the steeply increasing number of transition samples needed to sample the continuous state-space of increasing dimensionality sufficiently during training. The trends in Section 3.3 confirm the need for more data. In the future, experience replay (Lin, 1993; Adam et al., 2012) or similar techniques could be employed to increase training data efficiency. Furthermore, massively parallel approaches (Nair et al., 2015) are starting to emerge, opening up new avenues for large-scale reinforcement learning.

Although the employed reinforcement learning techniques guarantee convergence to an optimal policy, the computed personalization strategy may not be optimal for the model under consideration as the environment is only partially observable and the personalization problem ill-posed: there is no guarantee for solution existence or uniqueness. Yet, we showed that Vito can solve personalization more robustly and more effectively than standard methods under the same conditions. However, a theoretical analysis in terms of convergence guarantees and general stability of the method would be desirable, in particular with regards to the proposed re-initialization strategy. As a first step towards this goal, in preliminary (unreported) experiments on the EP and the WBC model we observed that the number of patients which do not require re-initialization (due to oscillation) to converge to a successful personalization consistently increased with increasing training data.

The data-driven initialization proposed in this work simplifies Vito’s setup by eliminating the need for user-provided initialization. However, currently there is no guarantee that the first initialization candidate is the one that will yield the “best” personalization outcome. Therefore, one could investigate the benefits of a fuzzy personalization scheme: many personalization processes could be run in parallel starting from the different initialization candidates. Parameter uncertainty quantification techniques (Neumann et al., 2014a) could then be applied to compute a prob-

ability density function over the space of model parameters. Such approaches aim to gather complete information about the solution-space, which can be used to study solution uniqueness and other interesting properties.

An important characteristic of any personalization algorithm is its stability against small variations of the measured data. A preliminary experiment indicated good stability of Vito: the computed parameters from several personalization runs, each involving small random perturbations of the measurements, were consistent. Yet in a small group of patients some parameter variability was observed, however, it was below the variability of the reference method under the same conditions. To what extent certain degrees of variability will impact other properties of the personalized model such as its predictive power will be subject of future research. We will also investigate strategies to improve Vito’s stability further. For instance, the granularity of the state-space could provide some flexibility to tune the stability: less representative states means a larger region in state space per state, thus small variations in the measured data might have less impact on personalization outcome. However, this could in turn have undesirable effects on other properties of Vito such as success rate or convergence speed (see Section 3.2.2).

Beyond these challenges, Vito showed promising performance and versatility, making it a first step towards an automated, self-taught model personalization agent. The next step will be to investigate the predictive power of the personalized models, for instance for predicting acute or long-term response in cardiac resynchronization therapy (Sermesant et al., 2009; Kayvanpour et al., 2015).

Appendix A. Data-driven state-space quantization

This section describes the details of the proposed data-driven quantization approach to define the set of representative MDP states \mathcal{S} (see Section 2.4). It is based on clustering, in particular on the weighted k -means algorithm described in Arthur and Vassilvitskii (2007). To this end, all objective vectors $\mathcal{C} = \{\mathbf{c} \in \mathcal{E}\}$ are extracted from the training data (Section 2.3). $\mathcal{C} \subset \mathbb{R}^{n_c}$ represents all observed “continuous states”. The goal is to convert \mathcal{C} into the finite set of representative MDP states \mathcal{S} while taking into account that Vito relies on a special “success state” $\hat{\mathcal{S}}$ encoding personalization success.

The success state $\hat{\mathcal{S}}$ does not depend on the data, but on the maximum acceptable misfit ψ . In particular, since personalization success implies that all objectives are met, $\hat{\mathcal{S}}$ should approximate a hyperrectangle centered at $\mathbf{0}$ and bounded at $\pm\psi$, i.e. a small

region in \mathbb{R}^{n_c} where $\forall i: |c_i| < \psi_i$. To enforce \hat{s} , the input to weighted k -means is preprocessed as follows.

First, the $\mathbf{0}$ -vector is inserted into \mathcal{C} , along with two vectors per dimension i , where all components are zero, except the i th component, which is set to $\pm 2\psi_i$. These $2n_c + 1$ inserted vectors are later converted into centroids of representative states to delineate the desired hyperrectangle for \hat{s} as illustrated in Fig. A.11.

Furthermore, to avoid malformation of \hat{s} , no other representative state should emerge within that region. Therefore, all vectors $\mathbf{c} \in \mathcal{C}$, where $\forall i: |c_i| < 2\psi_i$ (except for the inserted vectors) are canceled out by assigning zero weight, while the inserted vectors are assigned large weights $\rightarrow \infty$ and all remaining vectors weights of 1.

Next, k -means is initialized by placing a subset of the initial centroids at the locations of the inserted states, and the remaining $n_s - 2n_c - 1$ centroids at random vectors in \mathcal{C} . Both the large weight and the custom initialization enforce the algorithm to converge to a solution where one cluster centroid is located at each inserted vector, while the other centroids are distributed according to the training data. To ensure equal contribution of all objectives (cancel out different units, etc.), similarity is defined relative to the inverse of the user-defined convergence criteria (Eq. 7).

Finally, after k -means converged, the resulting centroids, denoted ξ_s , are used to delineate the region in \mathbb{R}^{n_c} assigned to a representative state s .

Appendix B. Data-driven initialization

This section describes the details of the proposed data-driven initialization approach to compute a list of candidate initialization parameter vectors $\mathcal{X}_0 = (\mathbf{x}'_0, \mathbf{x}''_0, \dots)$ for a new patient p based on the patient's measurements \mathcal{Z}^p and the training database \mathcal{E} (see Section 2.6.1).

First, all model states are extracted from the training database: $\Upsilon = \{\mathbf{y} \in \mathcal{E}\}$. Next, Υ is fed to a clustering algorithm (e.g. k -means). As in Appendix A, the distance measure is defined relative to the inverse of the convergence criteria (Eq. 7). The output is a set of centroids (for simplicity, in this work the number of centroids was set to n_s), and each vector is assigned to one cluster based on its closest centroid. Let $Y_p \subseteq \Upsilon$ denote the members of the cluster whose centroid is closest to \mathbf{z}_p and $\Xi_p = \{\mathbf{x} \in \mathcal{E} \mid f(\mathbf{x}) \in Y_p\}$ the set of corresponding model parameters. For each cluster, an approximation of the likelihood over the generating parameters is computed in terms of a probability density function. In this work a Gaussian mixture model is assumed:

$$\text{GMM}_p(\mathbf{x}) = \sum_{m=1}^M \nu_m \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_m, \Sigma_m) . \quad (\text{B.1})$$

The parameter vectors in Ξ_p are treated as random samples drawn from GMM_p . Its properties, namely the number of mixture components M , their weights ν_m , and their means $\boldsymbol{\mu}_m$ and covariance matrices Σ_m , are estimated from these samples using a multivariate kernel density estimator with automated kernel bandwidth estimation, see Kristan et al. (2011) for more details. Finally, the M estimated means are selected as initialization candidates and stored in a list $\mathcal{X}_0 = (\boldsymbol{\mu}_{m'}, \boldsymbol{\mu}_{m''}, \dots)$. The elements of \mathcal{X}_0 are sorted in descending order according to their corresponding ν_m -values to prioritize more likely initializations: $\boldsymbol{\mu}_{m'}$ is the mean with $m' = \arg \max_m \nu_m$.

References

Adam, S., Busoniu, L., Babuska, R., 2012. Experience replay for real-time reinforcement learning control. *IEEE Sys. Man. Cybern.* 42 (2), 201–212.

Aguado-Sierra, J., Kerckhoffs, R.C.P., Lionetti, F., Hunt, D., Villongco, C., Gonzales, M., Campbell, S.G., McCulloch, A.D., 2010. A computational framework for patient-specific multi-scale cardiac modeling. In: Kerckhoffs, R.C. (Ed.), *Patient-Specific Modeling of the Cardiovascular System*. Springer, pp. 203–223.

Aguado-Sierra, J., Krishnamurthy, A., Villongco, C., Chuang, J., Howard, E., Gonzales, M.J., Omens, J., Krummen, D.E., Narayan, S., Kerckhoffs, R.C.P., McCulloch, A.D., 2011. Patient-specific modeling of dyssynchronous heart failure: a case study. *Prog. Biophys. Mol. Bio.* 107 (1), 147–155.

Arthur, D., Vassilvitskii, S., 2007. k -means++: The advantages of careful seeding. In: *ACM-SIAM Symp. Discrete Algorithm.*, pp. 1027–1035.

Augenstein, K.F., Cowan, B.R., LeGrice, I.J., Nielsen, P.M., Young, A.A., 2005. Method and apparatus for soft tissue material parameter estimation using tissue tagged magnetic resonance imaging. *J. Biomech. Eng.* 127 (1), 148–157.

Barreto, A., Precup, D., Pineau, J., 2014. Practical kernel-based reinforcement learning. *arXiv preprint arXiv:1407.5358*.

Bellman, R., 1957. *Dynamic Programming*. Princeton University Press, Princeton, New Jersey.

Bishop, C.M., 2006. *Pattern Recognition and Machine Learning*, 4. Springer, New York.

Chabiniok, R., Moireau, P., Lesault, P.-F., Rahmouni, A., Deux, J.-F., Chapelle, D., 2012. Estimation of tissue contractility from cardiac cine-mri using a biomechanical heart model. *Biomech. Model. Mech.* 11 (5), 609–630.

Clayton, R., Bernus, O., Cherry, E., Dierckx, H., Fenton, F., Mirabella, L., Panfilov, A., Sachse, F., Seemann, G., Zhang, H., 2011. Models of cardiac tissue electrophysiology: progress, challenges and open questions. *Prog. Biophys. Mol. Bio.* 104 (1), 22–48.

Delingette, H., Billet, F., Wong, K.C.L., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C.A., Razavi, R., Ayache, N., 2012. Personalization of cardiac motion and contractility from images using variational data assimilation. *IEEE T. Biomed. Eng.* 59 (1), 20–24.

Frangi, A.F., Niessen, W.J., Viergever, M., 2001. Three-dimensional modeling for functional analysis of cardiac images, a review. *IEEE T. Med. Imaging* 20 (1), 2–5.

Hunter, P.J., Borg, T.K., 2003. Integration from proteins to organs: the physiome project. *Nat. Rev. Mol. Cell Bio.* 4 (3), 237–243.

Itu, L., Sharma, P., Georgescu, B., Kamen, A., Suciu, C., Comaniciu, D., 2014. Model based non-invasive estimation of PV loop from echocardiography. In: *Proceeding of the 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. Chicago, IL, pp. 6774–6777. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6945183>

Kaelbling, L.P., Littman, M.L., Moore, A.W., 1996. Reinforcement learning: A survey. *J. Artif. Intell. Res.* 4, 237–285.

Kayvanpour, E., Mansi, T., Sedaghat-Hamedani, F., Amr, A., Neumann, D., Georgescu, B., Seegerer, P., Kamen, A., Haas, J., Frese, K.S., Irawati, M., Wirsz, E., King, V., Buss, S., Mereles, D., Zitron, E., Keller, A., Katus, H.A., Comaniciu, D., Meder, B., 2015. Towards personalized cardiology: Multi-scale modeling of the failing heart. *PLoS ONE* 10 (7), e0134869.

Kerckhoffs, R.C.P., Lumens, J., Vernooy, K., Omens, J., Mulligan, L., Delhaas, T., Arts, T., McCulloch, A., Prinzen, F., 2008. Cardiac resynchronization: insight from experimental and computational models. *Prog. Biophys. Mol. Bio.* 97 (2), 543–561.

Konukoglu, E., Relan, J., Cilingir, U., Menze, B.H., Chinchapatnam, P., Jadidi, A., Cochet, H., Hocini, M., Delingette, H., Jais, P., Haïssaguerre, M., Ayache, N., Sermesant, M., 2011. Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology. *Prog. Biophys. Mol. Bio.* 107 (1), 134–146.

Krishnamurthy, A., Villongco, C.T., Chuang, J., Frank, L.R., Nigam, V., Belezouli, E., Stark, P., Krummen, D.E., Narayan, S., Omens, J.H., McCulloch, A.D., Kerckhoffs, R.C.P., 2013. Patient-specific models of cardiac biomechanics. *J. Comput. Phys.* 244, 4–21.

Kristan, M., Leonardis, A., Skočaj, D., 2011. Multivariate online kernel density estimation with gaussian kernels. *Pattern Recognit.* 44 (10), 2630–2642.

Kuijpers, N.H., Hermeling, E., Bovender, P.H., Delhaas, T., Prinzen, F.W., 2012. Modeling cardiac electromechanics and mechano-electrical coupling in dyssynchronous and failing hearts. *J. Cardiovasc. Transl. Res.* 5 (2), 159–169.

Kveton, B., Theodorou, G., 2012. Kernel-based reinforcement learning on representative states. In: *Association for the Advancement of Artificial Intelligence*, pp. 977–983.

Lagarias, J.C., Reeds, J.A., Wright, M.H., Wright, P.E., 1998. Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM J. Optimiz.* 9 (1), 112–147.

Le Folgoc, L., Delingette, H., Criminisi, A., Ayache, N., 2013. Current-based 4D shape analysis for the mechanical personalization of heart models. In: *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging*. In: LNCS, 7766. Springer, pp. 283–292.

Lin, L.-J., 1993. Reinforcement learning for robots using neural networks. Technical Report. DTIC Document.

Marchesseau, S., Delingette, H., Sermesant, M., Cabrera-Lozoya, R., Tobon-Gomez, C., Moireau, P., Figueras i Ventura, R.M., Lekadir, K., Hernandez, A., Garreau, M., Donal, E., Leclercq, C., Duckett, S.G., Rhode, K., Rinaldi, C.A., Frangi, A.F., Razavi, R., Chapelle, D., Ayache, N., 2013. Personalization of a cardiac electromechanical model using reduced order unscented kalman filtering from regional volumes. *Med. Image Anal.* 17 (7), 816–829.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Hiedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumara, D., Wierstra, D., Legg, S., Hassabis, D., 2015. Human-level control through deep reinforcement learning. *Nature* 518 (7540), 529–533.

Mülling, K., Kober, J., Kroemer, O., Peters, J., 2013. Learning to select and generalize striking movements in robot table tennis. *Int. J. Robot. Res.* 32 (3), 263–279.

- Nair, A., Srinivasan, P., Blackwell, S., Alciçek, C., Fearon, R., De Maria, A., Panneershelvam, V., Suleyman, M., Beattie, C., Petersen, S., Legg, S., Mnih, V., Kavukcuoglu, K., Silver, D., 2015. Massively parallel methods for deep reinforcement learning. arXiv:1507.04296.
- Neumann, D., Mansi, T., Georgescu, B., Kamen, A., Kayvanpour, E., Amr, A., Sedaghat-Hamedani, F., Haas, J., Katus, H., Meder, B., Hornegger, J., Comaniciu, D., 2014a. Robust image-based estimation of cardiac tissue parameters and their uncertainty from noisy data. In: MICCAI. In: LNCS, 8674. Springer, pp. 9–16.
- Neumann, D., Mansi, T., Grbic, S., Voigt, I., Georgescu, B., Kayvanpour, E., Amr, A., Sedaghat-Hamedani, F., Haas, J., Katus, H., Meder, B., Hornegger, J., Kamen, A., Comaniciu, D., 2014b. Automatic image-to-model framework for patient-specific electromechanical modeling of the heart. In: IEEE Int. Symp. Biomed. Imaging, pp. 935–938.
- Neumann, D., Mansi, T., Itu, L., Georgescu, B., Kayvanpour, E., Sedaghat-Hamedani, F., Haas, J., Katus, H., Meder, B., Steidl, S., Hornegger, J., Comaniciu, D., 2015. Vito – a generic agent for multi-physics model personalization: Application to heart modeling. In: MICCAI. In: LNCS, 9350. Springer, pp. 442–449.
- Nguyen-Tuong, D., Peters, J., 2011. Model learning for robot control: a survey. *Cogn. Process.* 12 (4), 319–340.
- Noble, D., 2002. Modeling the heart – from genes to cells to the whole organ. *Science* 295 (5560), 1678–1682.
- Powell, M.J., 2009. The BOBYQA algorithm for bound constrained optimization without derivatives. NA Report NA2009/06. Cambridge.
- Prakosa, A., Sermesant, M., Allain, P., Villain, N., Rinaldi, C., Rhode, K., Razavi, R., Delingette, H., Ayache, N., 2013. Cardiac electrophysiological activation pattern estimation from images using a patient-specific database of synthetic image sequences. *IEEE T. Biomed. Eng.*
- Rosenbrock, H., 1960. An automatic method for finding the greatest or least value of a function. *Comput. J.* 3 (3), 175–184.
- Schmid, H., Nash, M., Young, A., Hunter, P., 2006. Myocardial material parameter estimation – a comparative study for simple shear. *J. Biomech. Eng.* 128 (5), 742–750.
- Seegerer, P., Mansi, T., Jolly, M.-P., Neumann, D., Georgescu, B., Kamen, A., Kayvanpour, E., Amr, A., Sedaghat-Hamedani, F., Haas, J., Katus, H., Meder, B., Comaniciu, D., 2015. Estimation of regional electrical properties of the heart from 12-lead ECG and images. In: *Statistical Atlases and Computational Models of the Heart – Imaging and Modelling Challenges*. In: LNCS, 8896. Springer, pp. 204–212.
- Sermesant, M., Billet, F., Chabiniok, R., Mansi, T., Chinchapatnam, P., Moireau, P., Peyrat, J.-M., Rhode, K., Ginks, M., Lambiase, P., Arridge, S., Delingette, H., Sorine, M., Rinaldi, C.A., Chapelle, D., Razavi, R., Ayache, N., 2009. Personalised electromechanical model of the heart for the prediction of the acute effects of cardiac resynchronisation therapy. In: *Functional Imaging and Modeling of the Heart*. In: LNCS, 5528. Springer, pp. 239–248.
- Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning: An Introduction, 1. MIT press Cambridge, London, England.
- Tesauro, G., 1994. Td-gammon, a self-teaching Backgammon program, achieves master-level play. *Neural Comput.* 6 (2), 215–219.
- Wallman, M., Smith, N.P., Rodriguez, B., 2012. A comparative study of graph-based, eikonal, and monodomain simulations for the estimation of cardiac activation times. *IEEE T. Biomed. Eng.* 59 (6), 1739–1748.
- Wallman, M., Smith, N.P., Rodriguez, B., 2014. Computational methods to reduce uncertainty in the estimation of cardiac conduction properties from electroanatomical recordings. *Med. Image Anal.* 18 (1), 228–240.
- Wang, V.Y., Lam, H., Ennis, D.B., Cowan, B.R., Young, A.A., Nash, M.P., 2009. Modelling passive diastolic mechanics with quantitative mri of cardiac structure and function. *Med. Image Anal.* 13 (5), 773–784.
- Wang, Y., Georgescu, B., Chen, T., Wu, W., Wang, P., Lu, X., Ionasec, R., Zheng, Y., Comaniciu, D., 2013. Learning-based detection and tracking in medical imaging: a probabilistic approach. In: *Deformation Models*. In: LNCVB, 7. Springer, pp. 209–235.
- Westerhof, N., Elzinga, G., Sipkema, P., 1971. An artificial arterial system for pumping hearts. *J. Appl. Physiol.* 31 (5), 776–781.
- Wong, K.C., Sermesant, M., Rhode, K., Ginks, M., Rinaldi, C.A., Razavi, R., Delingette, H., Ayache, N., 2015. Velocity-based cardiac contractility personalization from images using derivative-free optimization. *J. Mech. Behav. Biomed.* 43, 35–52.
- Xi, J., Lamata, P., Niederer, S., Land, S., Shi, W., Zhuang, X., Ourselin, S., Duckett, S.G., Shetty, A.K., Rinaldi, C.A., Rueckert, D., Razavi, R., Smith, N.P., 2013. The estimation of patient-specific cardiac diastolic functions from clinical measurements. *Med. Image Anal.* 17 (2), 133–146.
- Zettinig, O., Mansi, T., Neumann, D., Georgescu, B., Rapaka, S., Seegerer, P., Kayvanpour, E., Sedaghat-Hamedani, F., Amr, A., Haas, J., Steen, H., Katus, H., Meder, B., Navab, N., Kamen, A., Comaniciu, D., 2014. Data-driven estimation of cardiac electrical diffusivity from 12-lead ECG signals. *Med. Image Anal.* 1361–1376.
- Zheng, Y., Barbu, A., Georgescu, B., Scheuering, M., Comaniciu, D., 2008. Four-chamber heart modeling and automatic segmentation for 3-D cardiac CT volumes using marginal space learning and steerable features. *IEEE T. Med. Imaging* 27 (11), 1668–1681.