

Bayesian optimisation for efficient parameter inference in a cardiac mechanics model of the left ventricle

Agnieszka Borowska  | Hao Gao  | Alan Lazarus | Dirk Husmeier 

School of Mathematics and Statistics,
University of Glasgow, Glasgow, UK

Correspondence

Dirk Husmeier, School of Mathematics and Statistics, University of Glasgow, Glasgow G12 8QQ, UK.
Email: dirk.husmeier@glasgow.ac.uk

Funding information

Engineering and Physical Sciences Research Council, Grant/Award Numbers: EP/T017899/1, EP/R018634/1

Abstract

We consider parameter inference in cardio-mechanic models of the left ventricle, in particular the one based on the Holzapfel-Ogden (HO) constitutive law, using clinical in vivo data. The equations underlying these models do not admit closed form solutions and hence need to be solved numerically. These numerical procedures are computationally expensive making computational run times associated with numerical optimisation or sampling excessive for the uptake of the models in the clinical practice. To address this issue, we adopt the framework of Bayesian optimisation (BO), which is an efficient statistical technique of global optimisation. BO seeks the optimum of an unknown black-box function by sequentially training a statistical surrogate-model and using it to select the next query point by leveraging the associated exploration-exploitation trade-off. To guarantee that the estimates based on the in vivo data are realistic also for high-pressures, unobservable in vivo, we include a penalty term based on a previously published empirical law developed using ex vivo data. Two case studies based on real data demonstrate that the proposed BO procedure outperforms the state-of-the-art inference algorithm for the HO constitutive law.

KEYWORDS

biomechanical model calibration, global optimisation, Holzapfel-Ogden constitutive law, statistical inference

1 | INTRODUCTION

The last decade has seen impressive progress in the mathematical modelling of soft-tissue mechanics.^{1,2} In particular, the specific application to cardiac mechanics³ promises to improve our understanding of cardiac physiology and pathology. A recent case-control study⁴ involving patients diagnosed with myocardial infarction (MI) and healthy volunteers has demonstrated that the biophysical parameters of a cardio-mechanic model of the left ventricle (LV) of the heart can be non-invasively estimated from cardiac magnetic resonance (CMR) scans, and that the parameters thus inferred allow for an accurate classification of subject disease status. However, the cardio-mechanic equations do not admit a closed-form solution and require computationally expensive numerical procedures based on finite element method (FEM).⁵ These procedures have to be repeated iteratively as part of a numerical optimisation or sampling procedure, leading to potentially excessive computational run times that discourage uptake of the methodology in the clinical practice.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *International Journal for Numerical Methods in Biomedical Engineering* published by John Wiley & Sons Ltd.

We aim to address this difficulty by adopting the framework of Bayesian optimisation (BO), which has been particularly designed for optimisation problems where each evaluation of the objective function is expensive and the number of such evaluations has to be kept as low as possible. The main idea is to sequentially approximate the unknown objective function by a statistical model and use this model to identify the next query point, typically according to some exploration-versus-exploitation trade-off criterion. In the literature the said statistical model is typically referred to as a *surrogate function*^{6,7} and Gaussian processes⁸ are commonly used in this context. (Surrogates are also called approximations, metamodels, response surface models or emulators.⁹ We reserve the latter term for pre-trained approximations based on considerably larger numbers of points compared to those used in BO.) The trade-off criterion is quantified by the so-called *acquisition function*. Several strategies on how to specify the details of BO, in particular the acquisition function, have been proposed in the literature.⁶ In this study, we use a popular criterion called Expected Improvement¹⁰ as it admits a closed-form expression for Gaussian process surrogates, is conceptually simple and performs well for a range of problems.¹¹

To date, there have only been a few applications of BO to physiology and clinical decision support, noticeably published only recently. BO has been applied to improve personalised time series models for patient monitoring¹² or for efficient human-in-the-loop optimisation of wearable devices, for example powered prosthetics for walking assistance¹³ or soft exosuits.¹⁴ BO has also allowed for rapidly mapping residual network function in stroke¹⁵ or an efficient design of transcatheter aortic valve leaflets and optimisation of their geometry.¹⁶ BO has been also used in the context of computer-aided (or robot-assisted) surgery.^{17,18} Another strand of literature has applied BO for efficient parameter inference for finite-element problems, for example in haemodynamics,¹⁹ for learning material property of atherosclerotic carotid arteries,²⁰ or in cardiology.²¹ The latter study of Fan et al.²¹ is most closely related to our work as it has combined finite element modelling and BO to characterise exercise-induced myocardium growth to infer two unknown growth parameters. The main difference to our work is that Fan et al.²¹ focused on the estimation of two growth parameters affecting myocardial growth, whereas our work focuses on constitutive parameters determining the passive filling process in diastole.

Myocardial stiffness has been considered an important factor in heart failure patients with preserved ejection fraction,^{22,23} which is thought to be caused by myocyte stiffening and interstitial fibrosis.²⁴ To this end, a robust and fast estimation of myocardial passive stiffness is a key step for improved patient characterisation and treatment planning, which remains a great challenge in the cardiac modelling community despite continuous efforts in decades.^{25,26,27,28,29,30,31} The purpose of the present paper is to adapt the BO framework to the problem of inferring passive myocardial properties of healthy human LVs from in vivo measurements. In this study, the passive myocardial properties are characterised by an invariant-based hyperelastic material model proposed by Holzapfel and Ogden³² (HO model), which has been widely used in cardiac modelling communities.^{33,4,34} The first attempt to address parameter estimation, also known as the inverse problem, in this context was made by Gao et al.²⁸ These authors proposed an algorithmic, multi-step procedure to infer the parameters of the HO model, in which in each step a subset of parameters was optimised with possibly different objective functions. The choice of particular subsets was based on expert knowledge of the myocardial properties of the LV. Our proposed BO approach is conceptually different, as BO iterations are run with respect to an a priori chosen parametrisation and objective function, thus requiring little expert knowledge.

To the best of our knowledge we are the first to introduce BO into cardiac mechanics for inferring unknown constitutive parameters. Our work contributes both to the global optimisation literature and to the application specific literature. Our first main contribution is developing a BO framework for the challenging problem of parameter inference in the HO model. Specifically, we propose modifications of the two components of BO, the target function and the acquisition function. The former modification consists in augmenting the target function by incorporating prior knowledge of the end-diastolic volume-pressure relationship from an empirical law established by Klotz et al.³⁵ based on ex vivo data. This allows us to obtain more realistic myocardial properties not only at physiologically normal pressures but also at high pressures, usually not observed in vivo. Generally, the myocardium response to increasing pressure is nonlinear: it appears linear for typical low pressures observed in vivo for healthy subjects, and exponential for increasing high pressures, such as those observed for example for patients with hypertension. It can be challenging to infer those parameters that characterise the nonlinear stiffening effects based solely on in vivo data.

Our second main contribution is adapting the standard Expected Improvement criterion¹⁰ to the challenges of inference in the HO model. First, following previous studies,^{36,37} we allow for unknown constraints in the parameter space by training a classifier based on a probity model with Gaussian process priors. Unknown constraints are related to those parameter regions that violate the physiological assumptions of cardio-mechanic models, which is then demonstrated as crashes of the associated forward simulators. Second, next to the standard approach of developing a surrogate for the

objective function as a whole, we develop an alternative paradigm, based on a *partial error surrogate*, in which we develop surrogates for individual squared error terms of the target function. This is motivated by a previous study³⁸ in the context of parameter inference in the HO model, which reported a substantial improvement in accuracy in the emulation framework resulting from emulating a vector of outputs of the cardio-mechanic model. Our proposed approach based on using a surrogate for the squared error terms instead of outputs has an advantage of being directly applicable to the Expected Improvement acquisition function, without requiring any approximation.

Our third main contribution is related to two empirical studies. First, we demonstrate that our proposed BO framework outperforms the state-of-the-art algorithm of Gao et al.²⁸ for parameter estimation in cardio-mechanic models. Second, we show how incorporating the prior knowledge about volume-pressure relationship (based on an empirical law established by Klotz et al.³⁵ using ex vivo data) allows for physiologically more realistic parameter estimation consistent with the LV's characteristics in high pressure regimes, relevant for example for patients with hypertension or patients with diastolic heart failure.²²

We note that in this paper we are concerned with Bayesian optimisation, which should not be confused with Bayesian inference. Bayesian inference is a statistical approach that quantifies uncertainty about model parameters via posterior distributions. Bayesian optimisation, on the other hand, is a computationally efficient global optimisation method. The word “Bayesian” in “Bayesian optimization” is related to the fact that methods from non-parametric Bayesian statistics are used to approximate the objective function and quantify the uncertainty of the approximation. As opposed to Bayesian inference, uncertainty quantification in Bayesian optimisation is not an end in itself, but a means to another end: optimally trading off exploration and exploitation in the quest for a computationally efficient optimisation path. In fact, conceptually Bayesian optimisation is closer to classical, or frequentist, statistics, than to Bayesian statistics, as the frequentist approach is based on obtaining point estimates via optimising an underlying loss function.

The structure of this paper is as follows. Section 2 provides necessary details of the cardio-mechanic model of interest and introduces the state-of-the-art inference method.²⁸ We discuss BO methodology, including the proposed extensions, in Section 3. Our three empirical studies are presented in Section 4. Section 5 concludes the paper and presents the outline for further research.

2 | CARDIO-MECHANIC MODEL OF THE LEFT VENTRICLE IN DIASTOLE

Biomechanical modelling of the LV in diastole is concerned with the process of diastolic filling, starting at early diastole and finishing at the end of diastole, which can be considered as a purely passive response. The nonlinear myocardial passive response means the myocardium is soft under lower pressure but becomes stiffer and stiffer under higher pressures because of gradually engaged collagen fibres.³² A few models of the LV biomechanics have been proposed in the literature, see a recent review by Chabiniok et al.² In this work, we follow previous studies^{33,28} and focus on the constitutive law proposed by Holzapfel and Ogden³² to describe passive myocardial response, in other words the myocardial stiffness. The adopted HO model is an invariant-based incompressible hyperelastic constitutive law which takes into account the layered myofibre structure of the myocardium. Studies have demonstrated that it can realistically capture the key properties of the myocardium and at the same time is relatively easy to implement numerically.³⁹ For these reasons it has been widely used in cardiac modelling communities^{40,43} The incompressible HO model is based on the following strain energy function

$$\begin{aligned} \Psi = & \frac{a}{2b} \{ \exp[b(I_1 - 3)] - 1 \} + \sum_{i \in \{f,s\}} \frac{a_i}{2b_i} \{ \exp[b_i(I_{4i} - 1)^2] - 1 \} \\ & + \frac{a_{fs}}{2b_{fs}} [\exp(b_{fs} I_{8fs}^2) - 1], \end{aligned} \quad (1)$$

where $I_i, i \in \{1, 4f, 4s, 8fs\}$ are invariants describing the deformation (corresponding to the matrix and fibre structure of the myocardium), and $a, b, a_f, b_f, a_s, b_s, a_{fs}, b_{fs}$ are the eight constitutive parameters, which we collect in vector ϑ . The invariants in (1) are given as

$$\begin{aligned} I_1 &= \text{tr}(\mathbb{C}), & I_{4f} &= \mathbf{f}_0 \cdot (\mathbb{C}\mathbf{f}_0), \\ I_{4s} &= \mathbf{s}_0 \cdot (\mathbb{C}\mathbf{s}_0), & I_{8fs} &= \mathbf{f}_0 \cdot (\mathbb{C}\mathbf{s}_0), \end{aligned} \quad (2)$$

where $\mathbb{C} = \mathbb{F}^T \mathbb{F}$ is the right Cauchy–Green deformation tensor, with \mathbb{F} being the deformation gradient, while \mathbf{f}_0 and \mathbf{s}_0 are the myocyte and sheet directions, respectively, determined via a rule based approach, details can be found in Wang et al.³³ Note that \mathbf{f}_0 and \mathbf{s}_0 are initial values hence are known before running the forward simulator. Given that the myocardium is incompressible, the Cauchy stress tensor can be determined from (1) as

$$\boldsymbol{\sigma} = \sum_{i=1,4f,4s,8fs} \frac{\partial W}{\partial I_i} \frac{\partial I_i}{\partial \mathbb{F}} \mathbb{F}^T - p \mathbf{I}, \quad (3)$$

in which p is a Lagrange multiplier to enforce the incompressibility constraint, and \mathbf{I} is the identity matrix. The eight parameters in ϑ have the following interpretations: a and b describe the matrix response, a_f and b_f describe the contribution from the fibres along myocytes; a_s and b_s represent the contribution of the fibres along the sheet direction; a_{fs} and b_{fs} account for the shear effects in the fibre–sheet plane. Interested readers are referred to Holzapfel and Ogden³² for a detailed explanation of the HO model.

The quasi-static cardio-mechanics model in diastole is given by⁴¹

$$\left. \begin{aligned} \nabla \cdot \boldsymbol{\sigma} + \mathbf{b} &= 0, & \text{in } \Omega, \\ \boldsymbol{\sigma} \cdot \mathbf{n} &= -p_{\text{endo}} \mathbf{n}, & \text{in } \Gamma^N, \\ \mathbf{u} &= \bar{\mathbf{u}}, & \text{in } \Gamma^D, \end{aligned} \right\} \quad (4)$$

where \mathbf{b} is the body force density per unit volume, p_{endo} represents the LV cavity pressure, $\bar{\mathbf{u}}$ is the vector of displacements at boundaries, Ω is the computational domain defined by the LV mesh with the boundary denoted as $\partial\Omega$, \mathbf{n} is the normal direction of $\partial\Omega$, Γ^N and Γ^D are the Neumann and Dirichlet boundaries, respectively, with $\Gamma^N \cup \Gamma^D = \partial\Omega$ and $\Gamma^N \cap \Gamma^D = \emptyset$. Solving (4) gives us the LV dynamics in response to the applied LV blood pressure, from which relevant quantities of interest can be readily extracted, such as the end-diastolic LV volume and circumferential strains. Figure 1 illustrates the essential components of the cardio-mechanics model in diastole.

The boundary value problem (4) does not admit a closed-form solution, hence numerical solutions are necessary in this context. The standard approach to solving biomechanical models is via the finite element method (FEM).⁵ In this work we use the nonlinear finite element software Abaqus FEA (Simulia, Providence, RI, USA) to solve the discretized cardio-mechanics model. For a given parameter vector ϑ , the LV mesh, that is the discretized LV geometry at early diastole, and the desired end diastolic LV pressure, the Abaqus forward simulator will determine the deformed mesh at the end of diastole by balancing the internal and external mechanical energies under prescribed boundary conditions, see Figure 1. The deformed mesh is then used to extract segmental circumferential strains at corresponding segments defined in four short-axial CMR cine images following clinical convention as specified by the American Heart Association (AHA),⁴² as well as to compute the cavity volume at the end of diastole. Figure 2 illustrates an example of the reconstructed LV geometry derived from in vivo cine images with manually defined wall boundaries.

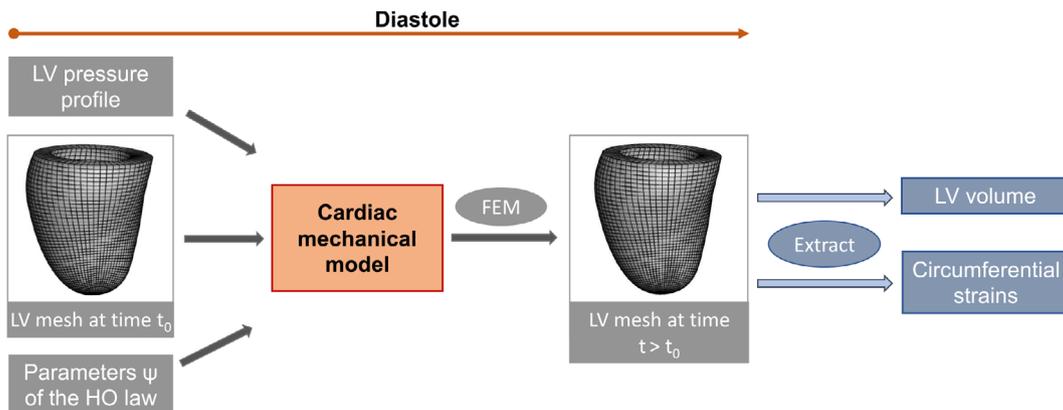


FIGURE 1 Inference pipeline in cardiac mechanics models

Details of the LV geometry reconstruction can be found in Gao et al.,⁴ including the definition of circumferential segmental strains.

The parameter vector ϑ describes the myocardial material properties and hence determines its mechanical responses through the HO model. Table 1 presents two sets of reference values of ϑ used in the literature. However, a subject-specific value of ϑ can neither be directly observed nor measured in vivo. Thus, we are interested in inferring it using routinely-available non-invasive in vivo data obtained from CMR images. Figure 2, left, presents an instance of such a scan together with manually annotated LV boundaries for the epicardium (red) and endocardium (blue). To the best of our knowledge, Gao et al.²⁸ were the first to address the issue of parameter inference in the HO model from in vivo clinical measurements (CMR scans). We discuss their algorithm, together with its updated version,⁴³ in Section 2.1. Still, parameter estimation for the HO model using in vivo data remains a challenging problem due to the potential non-identifiability of the constitutive parameters and their strong correlation,^{28,38} as well as the high cost of solving the cardio-mechanics model numerically.

2.1 | Existing estimation methods

Gao et al.²⁸ proposed a three step algorithm based on matching the responses from the forward simulator associated with the HO model with the in vivo measurements obtained from CMR scans (circumferential segmental strains and the end-diastolic cavity volume). We refer to their approach as the HGO algorithm. The key idea behind the HGO algorithm is not estimating ϑ directly but via rescaling a reference value ϑ_0 , taken from an earlier study.³³ Moreover, each step rescales a different subset of HO parameters.

Listing 1 presents steps of the HGO algorithm, where the scaling scalars C_a , C_b and C_3 are defined in the following way. Scalars C_a and C_b are used to match the parameters in two parameter groups, $a^{\text{group}} = \{a, a_f, a_s, a_{fs}\}$ and $b^{\text{group}} = \{b, b_f, b_s, b_{fs}\}$, respectively, as

$$a^{\text{group}} = C_a a_0^{\text{group}}, \quad b^{\text{group}} = C_b b_0^{\text{group}},$$

while scalar C_3 matches a and a_{fs}

$$a = C_3 a^{(2)}, \quad a_{fs} = C_3 a_{fs}^{(2)},$$

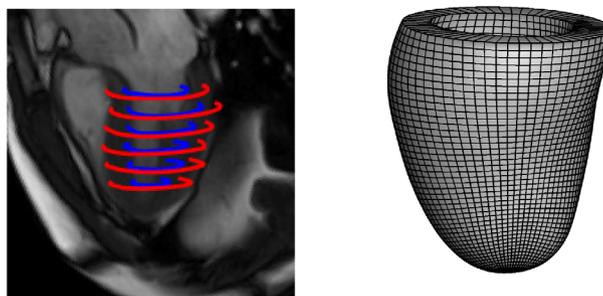


FIGURE 2 LV mesh reconstruction from a CMR image for a healthy volunteer. Left: segmented ventricular boundaries (blue: endocardium, red: epicardium) superimposed on a long axis CMR; right: the reconstructed LV geometry discretized with hexahedron elements

TABLE 1 Reference parameters ϑ_0 proposed in the literature

Work	a_0 (kPa)	b_0	a_{f0} (kPa)	b_{f0}	a_{s0} (kPa)	b_{s0}	a_{fs0} (kPa)	b_{fs0}
Wang et al. ³³	0.236	10.810	20.037	14.154	3.724	5.164	0.411	11.300
Gao et al. ⁴	0.18	2.6	3.34	2.73	0.69	1.11	0.31	2.58

Note: Wang et al.³³ estimated their parameters based on ex vivo data, while Gao et al.⁴ used in vivo data.

where $a^{(2)}$ and $a_{fs}^{(2)}$ are values of a and a_{fs} , respectively, optimised in the second step of the algorithm. The two objective functions matching the simulated values, depending on the constitutive parameter $\boldsymbol{\vartheta}$ and geometry \mathcal{H} , to the measurement are defined as follows

$$f_{O1}(\boldsymbol{\vartheta}, \mathcal{H}) = (V(\boldsymbol{\vartheta}, \mathcal{H}) - V^*)^2 + \sum_{i=1}^K (\varepsilon_i(\boldsymbol{\vartheta}, \mathcal{H}) - \varepsilon_i^*)^2, \quad (5)$$

$$f_{O2}(\boldsymbol{\vartheta}, \mathcal{H}) = \frac{(V(\boldsymbol{\vartheta}, \mathcal{H}) - V^*)^2}{V^*} + \sum_{i=1}^K (\varepsilon_i(\boldsymbol{\vartheta}, \mathcal{H}) - \varepsilon_i^*)^2, \quad (6)$$

where V^* and ε_i^* , $i = 1, \dots, K = 24$ are the measurements of the LV cavity volume and 24 circumferential strains, respectively. Note the 24 circumferential strains are measured from 4 short-axis slices from the base to the middle ventricle with 6 regions per slice. While $V(\boldsymbol{\vartheta}, \mathcal{H})$ and $\varepsilon_i(\boldsymbol{\vartheta}, \mathcal{H})$ are the corresponding values obtained from the forward simulator using the parameter vector $\boldsymbol{\vartheta}$ and the LV geometry \mathcal{H} .

2.1.1 | Updated algorithm

Gao et al.⁴³ have developed an updated version of the original HGO algorithm,²⁸ which we present in [Algorithm 2](#). We can see that the main difference between the two versions is that the new one includes an additional step based on refining C_a and C_b with $f_{O2, Klotz}$, which we discuss in [Section 2.2](#). In short, $f_{O2, Klotz}$ augments f_{O2} with an additional term which, for the given parameters $\boldsymbol{\vartheta}$ and LV geometry \mathcal{H} , quantifies the discrepancy between the LV volume predicted for pressure of 30 mmHg and—since no in vivo data is observed at such a high pressure—the prediction from the so-called Klotz-curve,³⁵ which we discuss in [Section 2.2](#). This aims at inferring parameters which not only agree with the in vivo data observed at low pressures, but also give realistic prediction for high pressure scenarios. Another difference between [Algorithms 1](#) and [2](#) can be observed in the last step: the updated version matches a and b while the original algorithm optimises a and a_{fs} . There are two reasons for the latter modification. First, the sensitivity of a_{fs} is much lower compared to a and b ²⁸ and it is more related to the shearing response, which does not directly link to the measured circumferential strains.³² Second, a and b are two highly sensitive parameters and directly related to the extracellular matrix response. Thus, in the final step a and b are updated to match the measured data.

2.1.2 | Parameter reduction

The eight parameters in $\boldsymbol{\vartheta}$ are highly correlated and not identifiable from limited data based on CMR scans.²⁸ Therefore, we follow Davies et al.³⁸ and use a parametrisation reduced to a 4-dimensional manifold

$$\left. \begin{aligned} a &= \theta_1 a_0, & b &= \theta_1 b_0, \\ a_f &= \theta_2 a_{f0}, & a_s &= \theta_2 a_{s0}, \\ b_f &= \theta_3 b_{f0}, & b_s &= \theta_3 b_{s0}, \\ a_{fs} &= \theta_4 a_{fs0}, & b_{fs} &= \theta_4 b_{fs0}, \end{aligned} \right\} \quad (7)$$

ALGORITHM 1 Original HGO algorithm of Gao et al.²⁸

- 1 Optimise C_a, C_b with f_{O1} on a grid $[0.1, 0.2, \dots, 1.0]^2$.
- 2 Optimise a_f, b_f with f_{O2} , s.t. $a_f \geq 2a_s$ and $b_f \geq 2b_s$.
- 3 Optimise C_3 with f_{O1} .

ALGORITHM 2 Updated HGO algorithm of Gao et al.⁴³

- 1 Optimise C_a, C_b with f_{O1} on a grid $[0.1, 0.2 \dots, 1.0]^2$.
- 2 Refine C_a, C_b with $f_{O2, Klotz}$, s.t. $a \geq 0.1, a_f \geq 0.1,$
 $a_s \geq 0.1$ and $a_{fs} \geq 0.1$.
- 3 Optimise a_f, b_f with f_{O2} , s.t. $a_f \geq 2a_s$ and $b_f \geq 2b_s$.
- 4 Optimise a, b with f_{O2} , s.t. $a \geq 0.1$.

where $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4)^T$ is the vector of scalings of the reference parameters $\boldsymbol{\vartheta}_0 = (a_0, b_0, a_{f0}, b_{f0}, a_{s0}, b_{s0}, a_{fs0}, b_{fs0})^T$. Wang et al.³³ proposed a popular reference parameter vector, used in several later studies.^{28,38} However, those values were based on ex vivo data, making them potentially of limited relevance to our problem (inference based on in vivo data). Hence, we use different values for $\boldsymbol{\vartheta}_0$ based on in vivo data population-wide averages.⁴ The only exception is the implementation of the HGO algorithm,²⁸ which originally uses the earlier reference values.³³ We have implemented this algorithm using both reference parameter vectors.^{33,4}

2.2 | Klotz-curve

In vivo data are collected at physiologically normal (low) ventricular pressures. These pressures, however, are not recorded in healthy volunteers due to the invasive nature of the procedure. We assume a population-based end-diastolic (ED) pressure of 8 mmHg.²⁸ Working with “low pressure data” implies we can only estimate those parameters of the HO model that describe myocardium properties at such low pressures. However, parameters leading to a similar low pressure behaviour (in terms of the implied stress–strain relationships) may result in a very different behaviour at high pressures, such as 30 mmHg. These high pressures are of interest as they may reveal LV stiffness with impaired relaxation, which characterises diastolic heart failure.²² Hence, we want to avoid that parameters which have been inferred from in vivo pressures exhibit physiologically unrealistic behaviour at high pressures.

However, high pressure in vivo volume measurements are not available, therefore we propose to predict them using the empirical law found by Klotz et al.,³⁵ which we will refer to as the Klotz-curve. This empirical relationship is valid for a wide range of hearts, including healthy and diseased (with congestive heart failure or left ventricular assist device support) human hearts as well animal (canine and rat) hearts. The normalised ED volume \tilde{V}^* is defined as³⁵

$$\tilde{V}^* = \frac{V^* - V_0}{V_{30} - V_0}, \quad (8)$$

where V^* is the measured unnormalised volume at P^* , V_0 is the zero-pressure volume called the load-free volume. An empirical law based on population-wide ex vivo data is then established³⁵ relating the normalised LV volume to the corresponding measurement pressure P^* with the specific form of

$$P^* = A(\tilde{V}^*)^B, \quad (9)$$

where A and B are parameters. In our case V^* is the in vivo measured volume at pressure $P^* = 8$ mmHg. Combining (9) and (8) results in the following formula for the high-pressure ED volume

$$V_{30}^{KL} = V_0 + \frac{V^* - V_0}{\widehat{V}^*} = V_0 + \frac{V^* - V_0}{\left(\frac{P^*}{A}\right)^{1/B}}, \quad (10)$$

where V_{30}^{KL} denotes V_{30} predicted using the Klotz-curve. Note that the ED pressure-volume relationship in (9) as well as the derived formula for the high-pressure ED volume (10) are based on *normalised* volumes: volume normalisation was crucial to obtain a widely-holding pressure-volume relationship.³⁵

Note that formula (10) depends on the load-free volume V_0 , which can only be measured *ex vivo*. With only *in vivo* measurements available, we need to make some assumption regarding how to estimate it. We use the early diastolic volume as a proxy for the load-free volume, since it is known that the LV pressure is the lowest at early diastole.^{28,44} Thus, we predict V_{30}^{KL} using the approximated \widehat{V}_0 and the *in vivo* measured volume V^* at $P^* = 8$ mmHg as

$$\widehat{V}_{30}^{KL} = \widehat{V}_0 + \frac{V^* - \widehat{V}_0}{\left(P^*/\widehat{A}\right)^{1/B}},$$

where $\widehat{A} = 27.78$ and $\widehat{B} = 2.76$ are the estimates.³⁵

We use the Klotz-curve predicted high-pressure volume \widehat{V}_{30}^{KL} to add a high-pressure penalty term to the objective function (6) as follows

$$f_{O2,Klotz}(\boldsymbol{\vartheta}, \mathcal{H}) = \left(\frac{V_8(\boldsymbol{\vartheta}, \mathcal{H}) - V_8^*}{V_8^*} \right)^2 + \sum_{i=1}^K (\varepsilon_i(\boldsymbol{\vartheta}, \mathcal{H}) - \varepsilon_i^*)^2 + \left(\frac{V_{30}(\boldsymbol{\vartheta}, \mathcal{H}) - \widehat{V}_{30}^{KL}}{\widehat{V}_{30}^{KL}} \right)^2, \quad (11)$$

where V_8 and V_{30} are LV cavity volumes from the forward simulator, at 8 mmHg and 30 mmHg, respectively; ε_i , $i = 1, \dots, K = 24$, are regional circumferential strains from the forward simulator at 8 mmHg; V_8^* is the *in vivo* measured LV cavity volume (at the normal pressure of 8 mmHg).

3 | BAYESIAN OPTIMISATION

Suppose we are interested in finding the global optimum of a smooth function $f(\mathbf{x}) : \mathbb{R}^d \supset \mathcal{X} \rightarrow \mathbb{R}$. We will refer to f as the objective or target function and without loss of generality assume it is a loss function to be minimised. Further, suppose computing f at any given \mathbf{x} is expensive and/or time-consuming so that we want to find the optimiser $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} f(\mathbf{x})$ in as few steps as possible. BO proceeds in an iterative fashion, at each iteration t selecting an input \mathbf{x}_t and querying the black-box function $y_t = f(\mathbf{x}_t)$. Input selection is done by optimising the so-called acquisition function (AF), which is based on an approximation of f with a surrogate and models the desired trade-off between *exploration* (searching for new promising regions) and *exploitation* (increasing the resolution around beneficial regions already discovered).

3.1 | Gaussian process surrogates

Typically, the AF is based on a *surrogate model* of the objective function f , with a popular choice for such a surrogate being Gaussian processes (GPs). As opposed to alternative nonlinear methods, such as for example splines and neural networks, GPs allow exact interpolation rather than regression, which is what is conceptually required for deterministic objective functions, such as the one based on a deterministic cardiomechanic model. However, due to technical reasons discussed below, restating the interpolation problem as a regression problems is more practical, so that below we focus on the GP regression setting.

In BO GP regression is trained on a sequentially expanding set of inputs $\mathbf{X}_t = [\mathbf{x}_j]_{j=1}^t$ and outputs $\mathbf{Y}_t = [y_j]_{j=1}^t$. We refer to Rasmussen and Williams⁸ for an extensive treatment of GPs. GP regression is a nonparametric model in which

a GP prior is put over the unknown function of interest to represent our initial beliefs about its smoothness. The target values at points $\mathbf{x}_j, j = 1, \dots, t$ observed up to iteration t are given by

$$y_j = f(\mathbf{x}_j) + \varepsilon_j, \quad \varepsilon_j \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2), \quad (12)$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the Gaussian distribution with mean μ and variance σ^2 (and *iid* stands for independently and identically distributed). Specification (12) is a general one, allowing for noisy observations. Our forward simulator of interest is deterministic, so mathematically there is no noise ($\sigma^2 = 0$). However, including a non-trivial noise term ε_j is still convenient for at least two reasons. First, in practice there may be numerical noise, due to the finite precision of the numerical solution of the forward problem. Second, when working with GPs it is a common practice to allow for a small *jitter* (or nugget) term (e.g. 10^{-9}), which is added to the diagonal of the training covariance matrix for numerical stability of the matrix inversion. Effectively, jitter is imposed noise, which slightly dilutes the informativeness of our deterministic data.

In (12) the latent values $f(\mathbf{x})$ follow a GP

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')),$$

where

$$\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))]$$

are the prior mean and the covariance function (kernel) of the process f , respectively. Hence, it is assumed that the responses y_j are conditionally independent given the latent values $f(\mathbf{x}_j)$. Below we adopt a standard assumption that $\mu(\mathbf{x}) = 0$ so the latent process $f(\mathbf{x})$ is fully specified by its kernel function.

The kernel can be specified in many different ways, including numerous standard functional forms.⁸ We consider the Matérn 3/2 kernel with automatic relevance determination (ARD) given as

$$k_0(\mathbf{x}, \mathbf{x}') = \sigma_k^2 \left(1 + \sqrt{3}r\right) \exp\left(-\sqrt{3}r\right), \quad r = \left(\sum_{i=1}^d \frac{(x_i - x'_i)^2}{l_i^2}\right)^{1/2},$$

where $l_i, i = 1, \dots, d$ is the length scale parameter for the i th input variable and ARD refers to this allowing for a different length scale for each dimension. Thus our GP regression model in (12) is parametrised by a vector of hyperparameters $\phi = (\sigma^2, \sigma_k^2, l_1, \dots, l_d)^T$, so that we consider the observation noise variance σ^2 a hyperparameter.⁸ The Matérn 3/2 kernel produces continuous trajectories with continuous derivatives at the same time providing a considerable degree of roughness (it is once differentiable) allowing us to model functions with cliffs and plateaus. Although this choice of kernel rules out standard second-order optimisation methods, like quasi-Newton, we note that second-order methods are of limited use in the presence of extensive multimodality (which is a typical problem when optimising AFs, see Section 3.2.1). Moreover, based on our experimentation this low-level Matérn class leads to numerically more stable covariance matrix inversion than a higher-level Matérn class.

For the collected inputs \mathbf{X}_t we obtain the GP prior over function values $p(\mathbf{f}_t | \mathbf{X}_t, \phi) = \mathcal{N}(\mathbf{f}_t | \mathbf{0}, \mathbf{K}_t)$, with $\mathbf{f}_t = [f(x_j)]_{j=1}^t$ and $\mathbf{K}_t = k(\mathbf{X}_t, \mathbf{X}_t)$. The likelihood is given by $p(\mathbf{Y}_t | \mathbf{f}) = \mathcal{N}(\mathbf{Y}_t | \mathbf{f}_t, \sigma^2 \mathbb{I})$, where \mathbb{I} is the identity matrix, and marginalising over the latent variables \mathbf{f}_t gives the formula for the marginal likelihood

$$p(\mathbf{Y}_t | \mathbf{X}_t, \phi) = \mathcal{N}(\mathbf{Y}_t | \mathbf{0}, \mathbf{K}_t + \sigma^2 \mathbb{I}).$$

Under the adopted Gaussian observation model (12) the conditional posterior distribution of the latent variables also becomes Gaussian, with the following form

$$p(\mathbf{f}_t | \mathbf{Y}_t, \mathbf{X}_t, \phi) = \mathcal{N}\left(\mathbf{K}_t (\mathbf{K}_t + \sigma^2 \mathbb{I})^{-1} \mathbf{Y}_t, \mathbf{K}_t - \mathbf{K}_t (\mathbf{K}_t + \sigma^2 \mathbb{I})^{-1} \mathbf{K}_t\right).$$

3.2 | Acquisition function

The performance of BO algorithms depends on the particular choice of AF. Various AFs have been proposed in the literature,⁶ but the Expected Improvement¹⁰ criterion remains one of the most popular choices. EI has gained popularity as it is conceptually simple (as opposed to e.g. information theoretic AFs), it admits a closed form expression for GP surrogates and performs well for a range of problems.¹¹ As its name suggests, EI quantifies the expected improvement, that is the amount of improvement over the current *incumbent* f^* , that is the best (lowest) value of f found so far, weighted by the probability of it occurring. EI is defined as

$$\text{EI}(\mathbf{x}) = \mathbb{E}_{p(y|\mathbf{x},\mathcal{D})}[\min(f^* - f(\mathbf{x}), 0)],$$

where \mathcal{D} denotes the set of inputs and outputs recorded so far and $p(y|\mathbf{x},\mathcal{D})$ is the predictive posterior output distribution given the chosen evaluation point \mathbf{x} and the previously observed \mathcal{D} . With a GP surrogate $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}))$ EI can be expressed^{10,6} as

$$\text{EI}(\mathbf{x}) = (f^* - \mu(\mathbf{x}))\Phi(z) + \sqrt{k(\mathbf{x})}\phi(z),$$

where $z = (f^* - \mu(\mathbf{x})) / \sqrt{k(\mathbf{x})}$ and Φ and ϕ are the CDF and PDF of the standard normal distribution, respectively. The EI is made up of two terms. The first term is increased by decreasing the predictive mean $\mu(\mathbf{x})$, the second term is increased by increasing the predictive uncertainty $k(\mathbf{x})$. This shows how EI automatically balances exploitation and exploration.

3.2.1 | Inner optimisation

BO selects the next query point based on maximising the AF, which is often referred to as *inner optimisation* as it is within the *outer optimisation* (optimising the unknown objective function f) of the BO algorithm. Even though inner optimisation does not require any computationally expensive forward simulations, it can be a challenging task.⁷ This is often due to the observed multimodality of EI, which typically is picked around several points and flat between them. A popular approach to overcome this problem is the multi-starts method,⁴⁵ which performs standard numerical optimisation using for example quasi-Newton methods such as BFGS multiple times, from different initial points. However, a pilot study we conducted revealed that the multi-starts method performs poorly in our application, irrespective of a particular kernel choice (we tested several standard kernels, such as squared exponential, Matérn 5/2, Matérn 3/2 and neural network), often indicating suboptimal points and thus leading to unstable BO runs. That was the case even with an atypically large number of restarts (set to 1000 and drawn from a Latin hypercube design).

For these reasons we use a different approach based on a global optimisation algorithm called OQNLP⁴⁶ (or Global Search in its implementation in MATLAB's Global Optimisation toolbox that we use), which led to very stable results. In short, the OQNLP is a heuristic algorithm also based on multiple starting points. Those starting points, however, are not randomly (or quasi-randomly) selected, but rather calculated from the scatter search algorithm.⁴⁷ They are then scored based on the corresponding value of the objective function and, if present, constraint satisfaction. Next, gradient-based nonlinear programming (NLP) solvers are run from the starting point with the highest score value. Thereafter, the remaining points are analysed and those that are unlikely to improve the best local minimum found so far are rejected (which is different to the standard multistart approach in which solvers are run for all the initial points). The filtering analysis takes into account the corresponding value of the objective function, constraint violations and basins of attraction. The latter means that the points in the same basins of attraction as previously analysed points are excluded from further NLP optimisation. This feature, that is analysis of basin of attraction, contributes to the efficiency of the OQNLP algorithm as it allows for discarding points which would converge to the already found local optima.

3.2.2 | Unknown constraints

It may happen during BO iterations that the proposed point, in spite of maximising the AF based on the current surrogate, will violate the biomechanical assumptions behind the HO law (we refer the reader to Appendix F) for an analysis

of simulator crashes. Such a situation will demonstrate itself as a crash of the associated forward simulator, which is inefficient as some computations will be carried out in vain. The problem of a simulator failing to terminate or crashing was addressed in the literature^{36,37} by weighting EI by the probability of the constraint being satisfied

$$\text{EI}_{\text{con}}(\mathbf{x}) = \text{EI}(\mathbf{x})\mathbb{P}(\mathbf{x} \in \mathcal{C}),$$

where $\mathbb{P}(\mathbf{x} \in \mathcal{C})$ is the probability of \mathbf{x} being a valid point not leading to a crash of the forward simulator.

To model the probability of the unknown constraint being satisfied we use a GP classifier with a probit likelihood³⁶ with approximate inference carried out using expectation propagation.⁴⁸ Similarly as for the regression model, we use the Matérn 3/2 kernel. A probit likelihood is a conceptually sound approach, however it means that we lose the analytic tractability of the standard EI function as for non-Gaussian likelihoods approximate inference methods (such as Laplace approximation,⁴⁹ variational inference,⁵⁰ expectation propagation⁴⁸) are necessary.⁸ A simpler though heuristic classifier could be provided by the so-called label regression⁵¹ (LR). LR is an ordinary linear regression model fitted to explain a binary outcome variable and its advantage is that $\mathbb{P}(\mathbf{x} \in \mathcal{C})$ can be computed explicitly, yet based on a sub-optimal probability model (a Gaussian, rather than a logit or probit model). LR has been successfully employed to account for hidden constraints in the context of BO applied to pulmonary circulation.⁵² However, in our preliminary study we found that in our application LR led to unstable results, generally performing poorly.

We note that the probability of a crash happening during a BO run is LV geometry and material parameter dependent. For the investigated LV geometries crashes are not very frequent, see Sections 4.3 and 4.4, however, we note that the reported numbers of crashes, if nonzero, are already based on training a classifier. In fact, our pilot study revealed that not allowing for any learning from past unsuccessful simulations can lead to repeatedly querying very similar “crash points”, resulting in much higher numbers of crashes.

3.2.3 | Partial error surrogates

Notice that both target functions, f_{O_2} and $f_{O_2, \text{Klotz}}$, are sums of individual squared error terms. Thus, instead of developing a surrogate for the target as a whole we propose to construct a surrogate for each error term separately. In the context of emulation³⁸ a similar approach led to a substantial improvement in accuracy by emulating the vector of outputs from the cardiac mechanic model. That approach is efficient in the context of emulation when the developed emulator is aimed to be used for a wide range of different subjects. However, in our context using surrogate functions for the outputs would require us to *approximate* EI, as implied by the predictive means from individual output surrogates. Considering surrogates for the individual error terms instead allows us to obtain a closed form expression for EI. Note that considering partial error surrogates instead of output surrogates does not cause an efficiency loss as in our case optimisation is carried out for subject-specific measurements, known before running computations (which is different than for emulators).

We refer to the proposed approach as *partial error surrogates*.

Let

$$f(\boldsymbol{\vartheta}, \mathcal{H}) = \sum_{i=1}^{K^*} f^{(i)}(\boldsymbol{\vartheta}, \mathcal{H}), \quad (13)$$

where f is either f_{O_2} in (6) so that $K^* = K + 1 = 25$, or $f_{O_2, \text{Klotz}}$ in (11) so that $K^* = K + 2 = 26$, and each $f^{(i)}$ refers to the squared error terms on the right-hand side in those formulae. We assume that $f^{(i)}$ s are mutually independent and model each $f^{(i)}$ using a GP regression parametrised by hyperparameters $\boldsymbol{\phi}^{(i)}$. Allowing for correlated $f^{(i)}$ s would require us to use computationally much more expensive multi-output GPs,^{53,54,55} which would increase the computational overhead of each BO iteration, diminishing the posited “cheapness” of the surrogate compared to the forward simulator. Each partial GP regression has the conditional posterior mean $\mu_t^{(i)}(\mathbf{x})$ and covariance $k_t^{(i)}(\mathbf{x})$

$$\mathbf{f}_t^{(i)} | \mathbf{Y}_t^{(i)}, \mathbf{X}_t, \boldsymbol{\phi}^{(i)} \sim \mathcal{GP}\left(\mu_t^{(i)}(\mathbf{x}), k_t^{(i)}(\mathbf{x})\right),$$

where $\mathbf{f}_t^{(i)} = \left[f^{(i)}(\mathbf{x}_j) \right]_{j=1}^t$ are the latent variables for the i th output from the cardiomechanic model and $\mathbf{Y}_t^{(i)} = \left[y_j^{(i)} \right]_{j=1}^t$ are the observed values for the i th output. Then the conditional posterior mean μ_t of the full target f is just the sum of conditional posterior means $\mu_t^{(i)}$ from the partial surrogates and the conditional posterior variance k_t is the sum of the partial variances $k_t^{(i)}$:

$$\mathbf{f}_t \mid \mathbf{Y}_t^{(1)}, \dots, \mathbf{Y}_t^{(K^*)}, \mathbf{X}_t, \boldsymbol{\phi}^{(1)}, \dots, \boldsymbol{\phi}^{(K^*)} \sim \mathcal{GP}(\mu_t(\mathbf{x}), k_t(\mathbf{x})),$$

with

$$\mu_t(\mathbf{x}) = \sum_{i=1}^{K^*} \mu_t^{(i)} \quad \text{and} \quad k_t(\mathbf{x}) = \sum_{i=1}^{K^*} k_t^{(i)}.$$

3.2.4 | Standardisation

Standardising inputs and outputs is a common practice when working with GPs, aimed to provide numerical stability of the computations. In our case input standardisation is not an issue as all the partial surrogates use the same inputs. Predictions for the outputs, however, enter the formulae for the conditional posterior distribution (14) so we need to “de-standardise” the predictions to obtain correct partial moments. Let $\tilde{\mathbf{Y}}_t^{(i)}$ denote the i th, $i = 1, \dots, K^*$, standardised partial squared error

$$\tilde{\mathbf{Y}}_t^{(i)} = \frac{\mathbf{Y}_t^{(i)} - \mathbf{m}_0^{(i)}}{\sigma_0^{(i)}},$$

where $\mathbf{m}_0^{(i)}$ and $\sigma_0^{(i)}$ are the mean and standard deviation, respectively, used to standardise the i th partial squared errors (e.g. calculated on the initial design). The reverse transformation is given as

$$\mathbf{Y}_t^{(i)} = g\left(\tilde{\mathbf{Y}}_t^{(i)}\right) = \mathbf{m}_0^{(i)} + \sigma_0^{(i)} \tilde{\mathbf{Y}}_t^{(i)},$$

so if the i th GP surrogate predicts $\mathcal{N}\left(\tilde{\mu}_t^{(i)}, \tilde{k}_t^{(i)}\right)$ for the standardised i th partial squared error, this corresponds to predicting $\mathcal{N}\left(g\left(\tilde{\mu}_t^{(i)}\right), \left(\sigma_0^{(i)}\right)^2 \tilde{k}_t^{(i)}\right)$ for the original i th partial squared error. Then the moments (14) used in the acquisition function are given by

$$\mu_t(\mathbf{x}) = \sum_{i=1}^{K^*} \left(\mathbf{m}_0^{(i)} + \sigma_0^{(i)} \tilde{\mu}_t^{(i)} \right) \quad \text{and} \quad k_t(\mathbf{x}) = \sum_{i=1}^{K^*} \left(\sigma_0^{(i)} \right)^2 \tilde{k}_t^{(i)}.$$

3.3 | Emulator for high-pressure volume

Computing $f_{O_2, Klotz}$ in (11) requires running the forward simulator twice, at 8 mmHg and 30 mmHg, which requires double computing time/power. Each of these runs may be unsuccessful, that is may lead to a crash. However, parameters leading to crashes may differ between the 8 and 30 mmHg cases. Hence, we would need two classifiers, one for the 8 mmHg case and one for the 30 mmHg case, and we could only update the GP surrogate for $f_{O_2, Klotz}$ when both cases terminate successfully.

To overcome these problems and to reduce the computational costs we propose to obtain the value for V_{30} using an emulator rather than by running the forward simulator. Emulators of cardiac mechanics models often include constitutive parameters as inputs.³⁸ However, the outputs from the forward simulator depend not only on the constitutive parameters but also on the LV geometry. Since we want the V_{30} emulator to approximate V_{30} for a broad range of LV geometries (as opposed to developing a geometry-specific emulator³⁸) we want to include some information about the

LV geometry as input. However, we cannot use the LV mesh directly as it is a high-dimensional (with thousands of coordinates) object and such a dimension is prohibitively high for any emulator. Our preliminary experimentation revealed that accurately predicting only the LV volume at 30 mmHg, and not the strains, does not require much details about a subject's specific shape of the LV and can be done based on a single “summary statistic” of the LV mesh—an approximation \hat{V}_0 to the load-free LV cavity volume V_0 (see Section 2.2)—without a substantial loss in the accuracy of the emulator. We note that if we were also after emulating circumferential strains at 30 mmHg, accounting for more subtle features of the LV geometry would be necessary. To sum up, we train the V_{30} emulator on the scaling parameters θ and the load-free volume \hat{V}_0 , in which we approximate the true function $V_{30}(\theta, \mathcal{H})$ by $V_{30}(\theta, V_0)$.

3.4 | Final algorithm

Listings 3 and 4 summarise the initialisation and iterations, respectively, of the BO algorithm with the constraint-weighted EI, Klotz-curve objective (11), partial error surrogates and V_{30} emulator.

4 | APPLICATIONS

This section presents three applications of BO to the problem of parameter inference in the HO model of the LV. We start with comparing BO *without* the Klotz-curve prior with the original HGO algorithm,²⁸ which we refer to as the *basic study*. We demonstrate that BO attains a similar or a lower value of the objective function f_{O_2} in considerably fewer iterations (Section 4.3). Next, in the *Klotz-curve study* we compare the versions of both algorithms that include the Klotz-curve (Section 4.4). We show that also in this case BO typically outperforms the multi-step algorithm, even though the updated HGO algorithm⁴³ works better than its original version. In both these studies we consider four healthy volunteers (HVs), same in both applications, which we label HV A, HV B, HV C and HV D. Finally, in Appendix C we investigate theoretical accuracy of the proposed method by means of a synthetic data study based on HV A from the basic study.

4.1 | Set-up

In all our applications we compare two variants of BO, with target and partial error surrogates (13) (see Section 3.1 and 3.2.3). We use 500 iterations for BO, which are run after 40 iterations corresponding to an initial design obtained with Latin hypercube sampling,⁵⁶ in which we follow a commonly followed rule-of-thumb¹⁰ to use $10d$ points to initialise optimisation (for a d -dimensional problem). We use a recent value for θ_0 (defined below (7)) based on in vivo data⁴ for both BO and the HGO algorithm, except the original HGO algorithm in the *basic study* for which we use the reference value from Wang et al.,³³ originally used in the HGO algorithm.

4.2 | Evaluation

In our result analysis we focus on three aspects. First, we consider the value of the objective function, as this is what both algorithms directly target. In the *basic study* this is f_{O_2} given in (6), in the two later studies this is $f_{O_2, \text{Klotz}}$ given in (11). Visually, we analyse the convergence of the *incumbent*, that is the best value of the objective function recorded so far, which is standard in the BO literature.^{6,37} The convergence is with respect to the invocations of the forward simulator, which are the main computational burden and which we want to limit. The physical computing time is a less reliable measure of computing costs as it necessarily depends on the machine used, number of task performed in parallel, etc., and does not generalise to other mechanistic models with different computational costs of the forward problem. Note that the HGO algorithm is based on optimisation solvers implemented as “black-boxes” (such as the trust-region-reflective algorithm or the sequential quadratic programming algorithm in MATLAB), so that tracking the incumbent over the individual iterations is not directly available. Moreover, each step in HGO has its own stopping criterion determining the number of iterations carried out in each step. Therefore, for HGO we just report the final value of the objective y_{\min}^{HGO} and the corresponding number of iterations i_{\min}^{HGO} (defined as the sum of iterations performed in all steps).

ALGORITHM 3 Initialisation of Bayesian optimisation with partial error surrogates

Input : d – input dimension (here: $d = 4$);
 K^* – number of outputs to be emulated ($K^* = 26$);
 P^* – measurement (low) pressure (8 mmHg);
 V_8^* and ε_i^* , $i = 1, \dots, K$ – volume and strain measurements at P^* ;
 \hat{V}_0 – load free volume estimate.

Output: \hat{V}_{30}^{KI} – Klotz-curve predicted volume at 30 mmHg;
 $\mathbf{X}(\tilde{\mathbf{X}})$ – explanatory variables (standardised) for regressions;
 $\mathbf{Y}^{(i)}(\tilde{\mathbf{Y}}^{(i)})$, $i = 1, \dots, K^*$ – response var. (standardised) for regressions;
 $\phi^{(i)}$, $i = 1, \dots, K^*$ – regression hyperparameters on the initial design;
 m_x, σ_x – inputs mean and st. dev. on the initial design;
 $m_0^{(i)}, \sigma_0^{(i)}$, $i = 1, \dots, K^*$ – outputs mean and st. dev. on the initial design;
 \mathbf{Z} (optional) – explanatory variables for classification on the initial design;
 \mathbf{U} (optional) – response variables for classification on the initial design;
 ϕ^c (optional) – classification hyperparameters on the initial design.

- 1 Calculate $\tilde{V}_8^* = (P^*/\hat{A})^{1/\hat{B}}$ (normalised volume at P^* from (9)).
- 2 Calculate $\hat{V}_{30}^{KI} = \hat{V}_0 + (V_8^* - \hat{V}_0)/\tilde{V}_8^*$.
- 3 Generate $t_0 = 10d$ points $\mathbf{X} = [\mathbf{x}_j]_{j=1}^{t_0}$ from an initial design scheme (e.g. Latin hypercube).
- 4 Run the forward simulator on \mathbf{X} to obtain partial squared errors $\mathbf{Y}^{(i)}$, $i = 1, \dots, K^*$.
- 5 **if** crashes for $\mathbf{X}_B = [\mathbf{x}_j]_{j \in B}$, $B \subset \{1, \dots, t_0\}$ **then**
- 6 Set $\mathcal{G} = \{1, \dots, t_0\} \setminus B$ (indices of good points).
- 7 Set $\mathbf{Z}_G = \mathbf{X}_G = [\mathbf{x}_j]_{j \in \mathcal{G}}$, $\mathbf{Z}_B = \mathbf{X}_B$ and $\mathbf{Z} = [\mathbf{Z}_G^T, \mathbf{Z}_B^T]^T$.
- 8 Set $\mathbf{U}_G = \mathbf{1}_{|\mathcal{G}| \times 1}$, $\mathbf{U}_B = \mathbf{0}_{|B| \times 1}$ and $\mathbf{U} = [\mathbf{U}_G^T, \mathbf{U}_B^T]^T$.
- 9 Set $\mathbf{X} := \mathbf{X}_G$ (only good points for regression).
- 10 Fit an initial GP probit classifier on \mathbf{Z} and \mathbf{U} (obtain ϕ^c).
- 11 **end**
- 12 Calculate m_x and σ_x of $\mathbf{X}^{(i)}$.
- 13 Standardise \mathbf{X} with m_x, σ_x to obtain $\tilde{\mathbf{X}}$.
- 14 **for** $i = 1$ **to** K^* **do**
- 15 Calculate $m_0^{(i)}$ and $\sigma_0^{(i)}$ of $\mathbf{Y}^{(i)}$.
- 16 Standardise $\mathbf{Y}^{(i)}$ with $m_0^{(i)}, \sigma_0^{(i)}$ to obtain $\tilde{\mathbf{Y}}^{(i)}$.
- 17 Fit an initial GP regression model on $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}^{(i)}$ (obtain $\phi^{(i)}$).
- 18 **end**

Quantitatively, we compare BO and HGO based on (y_{\min}, i_{\min}) and $(y(i_{\min}^{\text{HGO}}), i(i_{\min}^{\text{HGO}}))$. The former pair is concerned with the lowest value attained by BO over all 540 iterations y_{\min} and the number of iterations completed to obtain y_{\min} denoted i_{\min} . These values are compared with their HGO counterparts, y_{\min}^{HGO} and i_{\min}^{HGO} . The latter pair relates BO iterations to the final result from the HGO algorithm: $y(i_{\min}^{\text{HGO}})$ is the value of a BO run after the number of iterations used by the HGO algorithm i_{\min}^{HGO} , while $i(i_{\min}^{\text{HGO}})$ is the number of iterations required to achieve at least as good a performance as the HGO algorithm. If after 540 iterations the value of the objective obtained with BO is still higher than the final value from the HGO algorithm we set $i(i_{\min}^{\text{HGO}}) = \emptyset$. Table 2 summarises the notation used in both empirical studies.

Second, we analyse the *stress-stretch curves* corresponding to the final parameter estimates. We consider two virtual uni-axial stretch experiments, one is along the myofibre direction (myocytes), and the other one is along the sheet direction (transmurally), often adopted in the literature for characterising nonlinear myocardial property.^{28,39,57} Stress-stress curves present the relationship between stress and stretch along two axes in the myocardium, along the myocyte and along the sheet direction. They provide insights into the predicted stiffness of the myocardium, with higher values of stress for a given stretch value indicating a stiffer material. They are of interest to cardiac mechanics specialist and clinicians as the parameters of the HO model cannot always be uniquely identified using experimental data²⁸ and further due to the nonlinear myocardium response. The latter makes the differences in the predicted parameters not fully

ALGORITHM 4 Iterations of Bayesian optimisation with partial error surrogates

Input : t_0 – number of valid points on the initial design;
 K^* – number of outputs ($K^* = 26$);
 V_8^* and ϵ_i^* , $i = 1, \dots, K$ – volume and strain measurements;
 \hat{V}_{30}^{KI} – Klotz-curve predicted volume at 30 mmHg;
 \mathbf{X} ($\tilde{\mathbf{X}}$) – explanatory variables (standardised) for regressions;
 $\mathbf{Y}^{(i)}$ ($\tilde{\mathbf{Y}}^{(i)}$), $i = 1, \dots, K^*$ – response var. (standardised) for regressions;
 $\phi^{(i)}$, $i = 1, \dots, K^*$ – regression hyperparameters on the initial design;
 m_x, σ_x – inputs mean and st. dev. on the initial design;
 $m_0^{(i)}, \sigma_0^{(i)}$, $i = 1, \dots, K^*$ – outputs mean and st. dev. on the initial design;
 \mathbf{Z} (optional) – explanatory variables for classification on the initial design;
 \mathbf{U} (optional) – response variables for classification on the initial design;
 ϕ^c (optional) – classification hyperparameters on the initial design.

Output: \mathbf{X} – valid inputs;
 $\mathbf{Y}^{(i)}$, $i = 1, \dots, K^*$ – partial squared errors;
 \mathbf{Z} (optional) – crash points.

```

1 for  $t = 1$  to  $T$  do
2   if exists  $\mathbf{Z}$  then
3     Optimise  $\text{EI}_{con}$  with GlobalSearch and set  $\tilde{\mathbf{x}}^* = \arg \min(-\text{EI}_{con})$ .
4   else
5     Optimise EI with GlobalSearch and set  $\tilde{\mathbf{x}}^* = \arg \min(-\text{EI})$ .
6   end
7   Re-standardise  $\tilde{\mathbf{x}}^*$  to obtain  $\mathbf{x}^*$ .
8   Run the forward simulator with  $\mathcal{G}$  implied by  $\mathbf{x}^*$ .
9   if crash for  $\mathbf{x}^*$  then
10    Add  $\mathbf{x}^*$  as the last row of  $\mathbf{Z}_B$ .
11    Set  $\mathbf{Z}_G := \mathbf{X}$  and  $\mathbf{Z} = [\mathbf{Z}_G^T, \mathbf{Z}_B^T]^T$ .
12    Set  $\mathbf{U}_B := [\mathbf{U}_B^T, 1]^T$ .
13    Denote the number of rows of  $\mathbf{X}$  by  $g$  and set  $\mathbf{U}_G = \mathbf{1}_{g \times 1}$  and  $\mathbf{U} = [\mathbf{U}_G^T, \mathbf{U}_B^T]^T$ .
14    Update the classifier using  $\mathbf{Z}$  and  $\mathbf{U}$  (obtain  $\phi^c$ ).
15  else
16    Add  $\mathbf{x}_*$  as the last row of  $\mathbf{X}$ .
17    Standardise  $\mathbf{x}_*$  with  $m_x, \sigma_x$  to obtain  $\tilde{\mathbf{x}}_*$ .
18    Add  $\tilde{\mathbf{x}}_*$  as the last row of  $\tilde{\mathbf{X}}$ .
19    Use the emulator to predict  $V_{30}$ .
20    for  $i = 1$  to  $K^*$  do
21      Calculate  $y^{(i)}$  using the  $i$ th output and the corresponding measurement/the Klotz-curve prediction  $\hat{V}_{30}^{KI}$ .
22      Add  $y^{(i)}$  as the last row of  $\mathbf{Y}^{(i)}$ .
23      Standardise  $y^{(i)}$  with  $m_0^{(i)}, \sigma_0^{(i)}$  to obtain  $\tilde{y}^{(i)}$ .
24      Add  $\tilde{y}^{(i)}$  as the last row of  $\tilde{\mathbf{Y}}^{(i)}$ .
25      Update the GP regression model on  $\tilde{\mathbf{X}}$  and  $\tilde{\mathbf{Y}}^{(i)}$  (obtain  $\phi^{(i)}$ ).
26    end
27  end
28 end

```

informative about the implied output differences. For the reproducibility of stretch-stress curves, we report the underlying parameters in Appendix A.

Third, we consider deformations of the 3D LV geometry implied by the final parameter estimates. The underlying rationale is that uni-axial stretch-stress curves provide only a simplified test and do not capture the full picture of the deformation of the whole LV geometry. Following a cardiac mechanics benchmark paper,⁵⁸ we illustrate the deformation of a line in the epicardial surface (see Figure 3), with a specific focus on two regions: the inflexion point near the middle ventricle and the apical region. In this section we discuss deformations at the standard pressure of 8 mmHg, while Appendix D provides additional figures related to deformation at 30 mmHg. The latter, much higher pressure may be useful for comparing estimates as it leads to exaggerated differences in the implied deformations because of the intrinsic nonlinearity in the HO model.

Both BO and HGO provide point estimates. Given an increasing awareness in the cardiac modelling community of the importance of uncertainty quantification and sensitivity analysis for model parameter estimates,^{59,60,61} we discuss in Appendix E how to quantify uncertainty of BO estimates in a straightforward way. We then illustrate one of the available methods, the residual bootstrap approach, in a computational experiment.

4.3 | Basic study

Gao et al.²⁸ have established a benchmark for parameter inference in the HO model. Their algorithm (see Listing 1) does not involve the Klotz-curve prior discussed in Section 2.2, hence for comparability we first consider a version of BO without the Klotz-curve prior, with the objective function given by f_{O_2} from (6).

We appreciate that the HGO algorithm runs with respect to different parameters than the proposed BO framework and that the number of parameters used by both algorithms is different (in particular, in HGO the number of parameters with respect to which the optimisation is carried out may vary between the algorithm steps). This may raise concerns about the fairness of the BO–HGO comparison. Therefore, in Appendix G we present a comparison of BO against an established iterative optimisation scheme in the same 4-dimensional parameter space and demonstrate the superiority of BO also in this context.

Figure 4 presents the convergence of the *incumbent* for the four considered HVs. We observe that BO converges very quickly, typically in around 200 iterations after the initial design, so that the graph stays flat over the remaining

TABLE 2 Notation used in the empirical studies

Symbol	Meaning
y_{\min}	The lowest value of the $f_{O_2, \text{Klotz}}$ objective function
y_{\min}^S	y_{\min} but with the Klotz error computed using the simulator (not emulator)
i_{\min}	Number of iterations leading to y_{\min} ($y(i_{\min}) = y_{\min}$)
$y(i_{\min}^{\text{HGO}})$	Value of the objective function after the number of iterations used by HGO
$y^S(i_{\min}^{\text{HGO}})$	$y(i_{\min}^{\text{HGO}})$ but with the Klotz error computed using the simulator (not emulator)
$i(i_{\min}^{\text{HGO}})$	Number of iterations required to achieve at least as good a performance as the HGO algorithm (\emptyset if that did not happen within 540 iterations)
V_{30}	LV volumes at 30 mmHg, computed with emulator or simulator at y_{\min} or $y(i_{\min}^{\text{HGO}})$: $E(y_{\min})$ when computed for y_{\min} with the emulator $S(y_{\min})$ when computed for y_{\min} with the simulator $E(i_{\min}^{\text{HGO}})$ when computed for $y(i_{\min}^{\text{HGO}})$ with the emulator $S(i_{\min}^{\text{HGO}})$ when computed for $y(i_{\min}^{\text{HGO}})$ with the simulator
#Cr.	Total number of crashes of the forward simulator

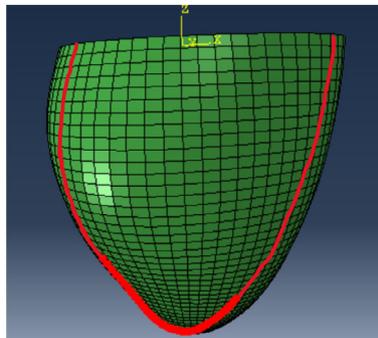


FIGURE 3 A defined reference line in the epicardial surface for illustrating the deformed LV shape⁵⁸

iterations. The HGO algorithm typically requires more iterations to terminate. Moreover, in all the cases BO converges to a lower value of the objective function than HGO. In the most extreme case of HV C, the lowest value found on the BO initial design is lower than the final value for HGO. The two variants of BO, with target and partial error surrogates

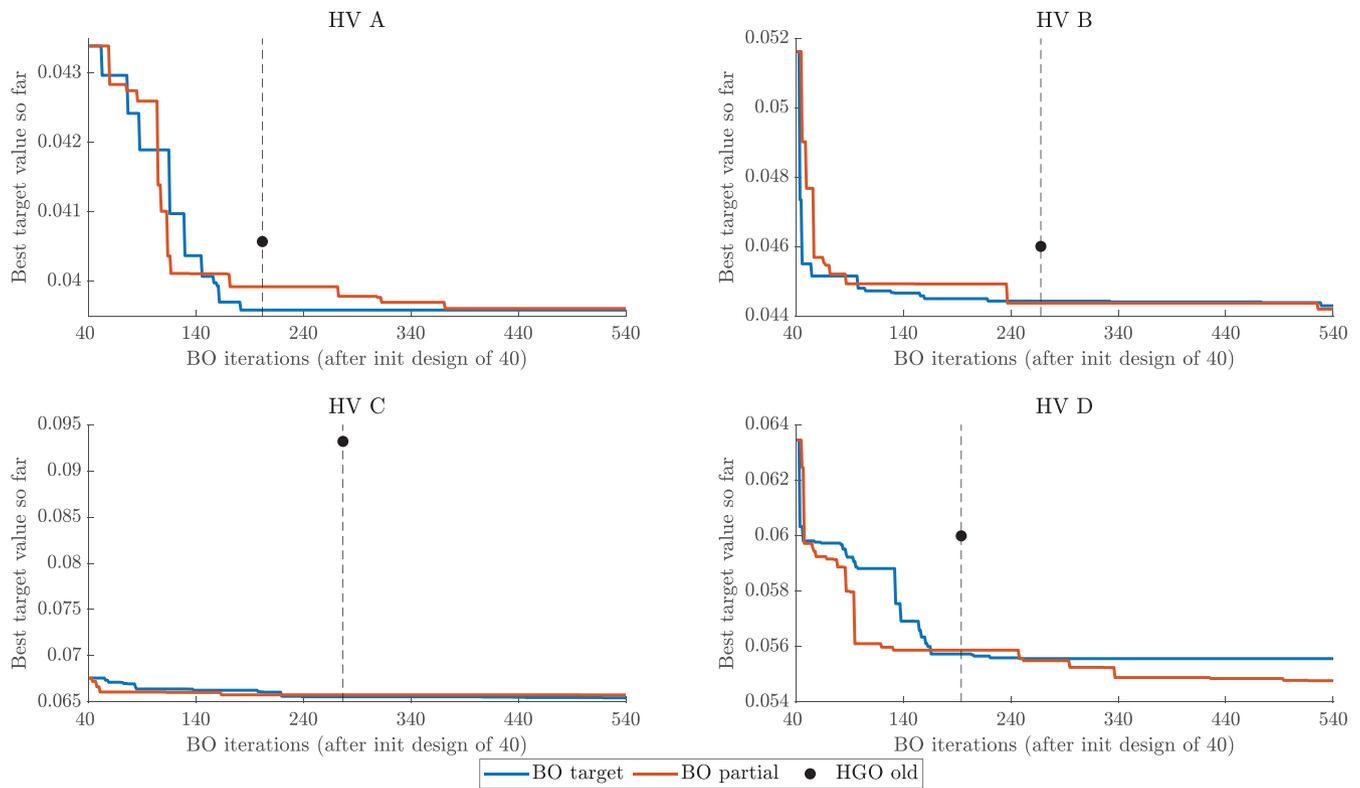


FIGURE 4 Basic study: convergence of the objective function f_{O2} for Bayesian optimisation and the original HGO algorithm for four LV geometries (HV A, HV B, HV C, HV D). Horizontal axis: Bayesian optimisation iterations starting from 40 iterations for the initial design. Vertical axis: best value of the objective function f_{O2} recorded so far. Black dot and horizontal dashed line: the final value of the objective function f_{O2} for the HGO algorithm and the associated number of iterations. Bayesian optimisation with a target surrogate (target) and a partial error surrogate (partial) together with the old version of the HGO algorithm (HGO old)

TABLE 3 Basic study: convergence of the objective function for Bayesian optimisation and the original HGO algorithm for four LV geometries (HV A, HV B, HV C, HV D)

HV	Algorithm	y_{\min}	i_{\min}	$y(i_{\min}^{\text{HGO}})$	$i(y_{\min}^{\text{HGO}})$	#Cr.
HV02	BO targ.	0.0396	181	0.0396	129	1
	BO part.	0.0396	371	0.0399	113	1
	HGO old	0.0406	201	0.0406	201	0
HV B	BO targ.	0.0443	528	0.0444	45	0
	BO part.	0.0442	525	0.0444	56	6
	HGO old	0.0460	267	0.0460	267	0
HV C	BO targ.	0.0655	518	0.0656	40	3
	BO part.	0.0658	499	0.0658	40	0
	HGO old	0.0932	276	0.0932	276	0
HV D	BO targ.	0.0556	247	0.0557	46	7
	BO part.	0.0548	516	0.0559	47	0
	HGO old	0.0600	193	0.0600	193	0

Note: Bayesian optimisation with a target surrogate (targ.) and a partial error surrogate (part.) together with the old version of the HGO algorithm (HGO old).

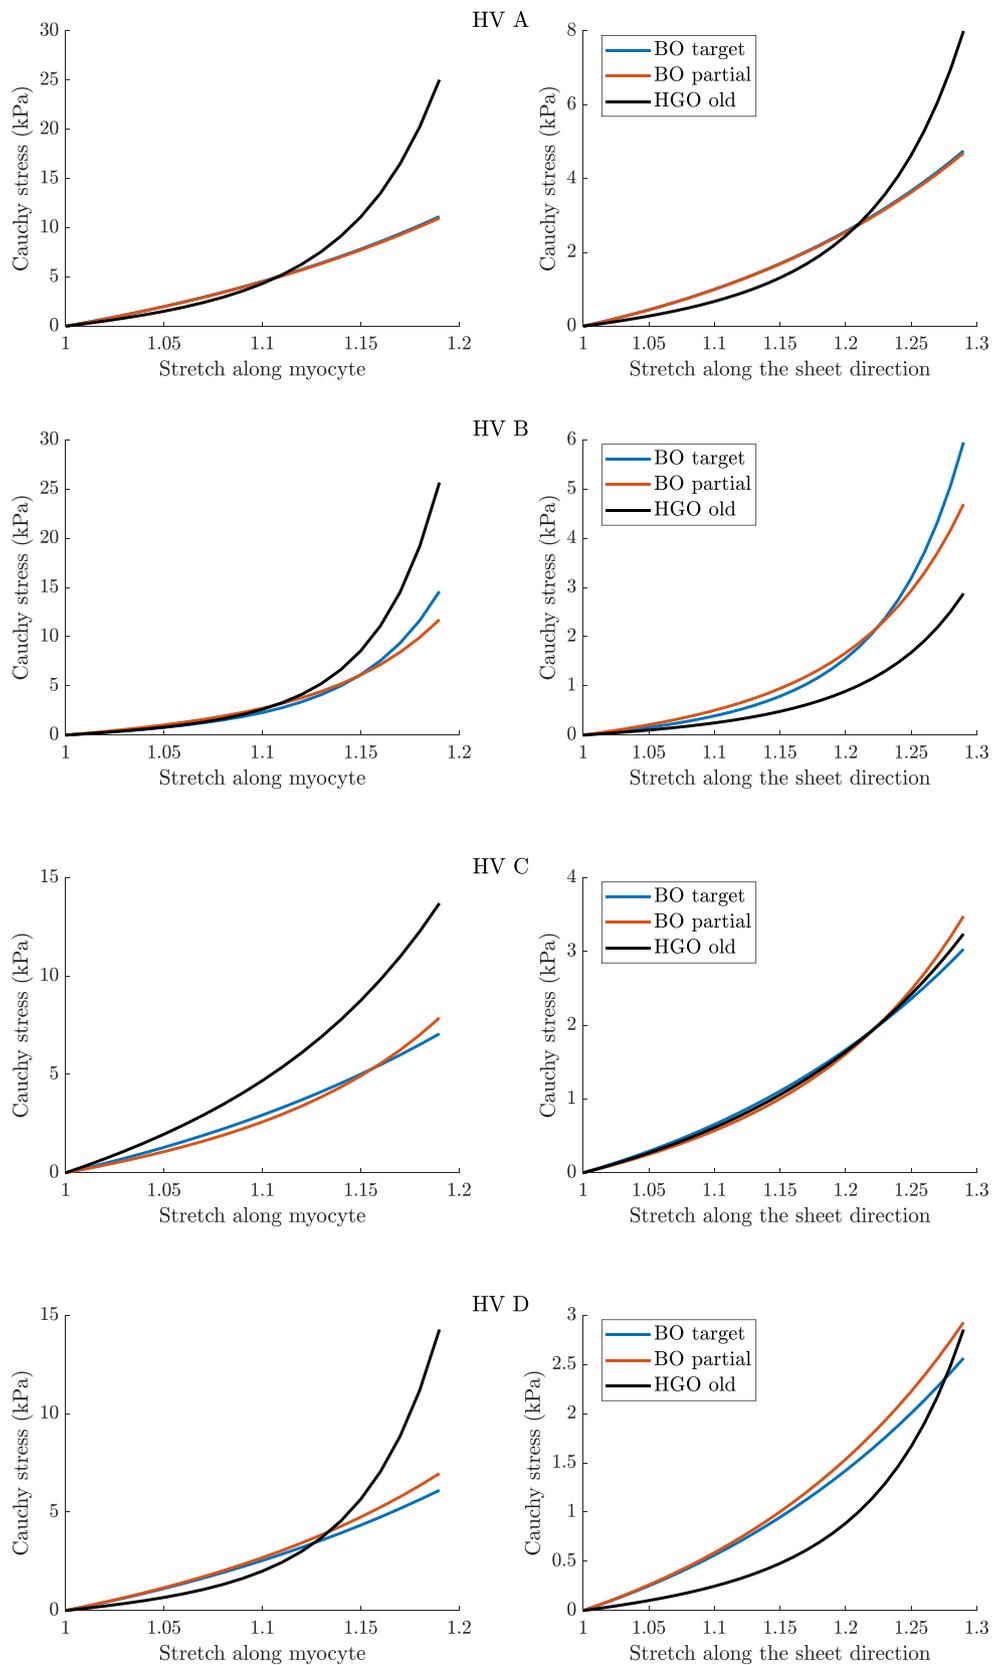
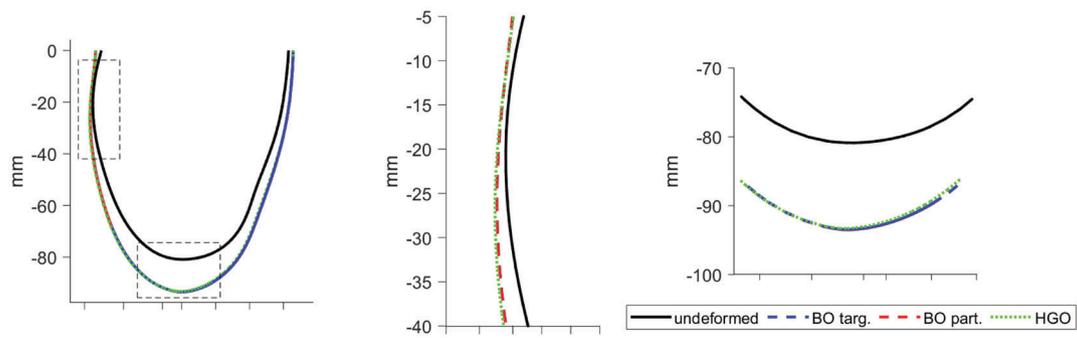
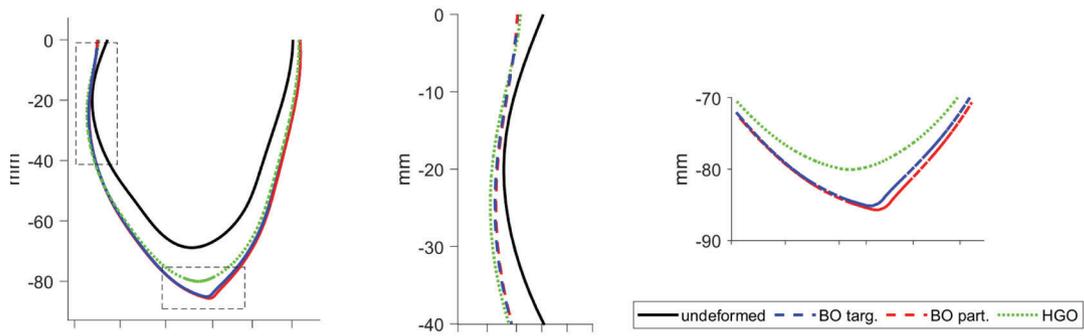


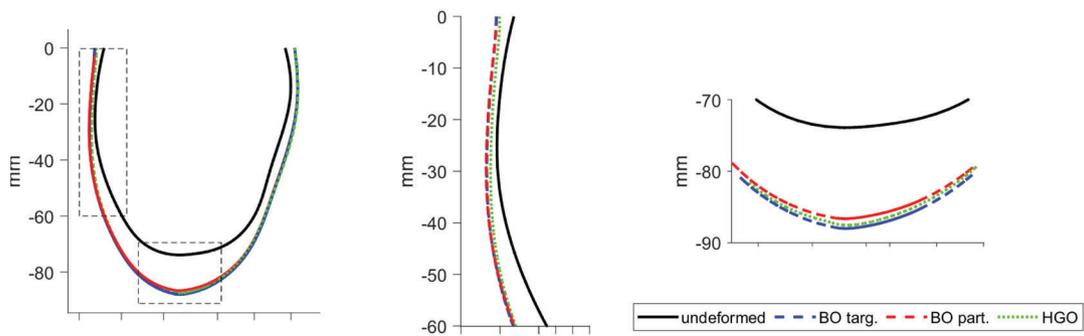
FIGURE 5 Basic study: stretch-stress curves for four LV geometries (HV A, HV B, HV C, HV D). Left: responses to stretches along the myocyte direction λ_0 , right: responses to stretches along the sheet direction λ_s (see (2)). Bayesian optimisation with a target surrogate (target) and a partial error surrogate (partial) together with the old version of the HGO algorithm (HGO old)



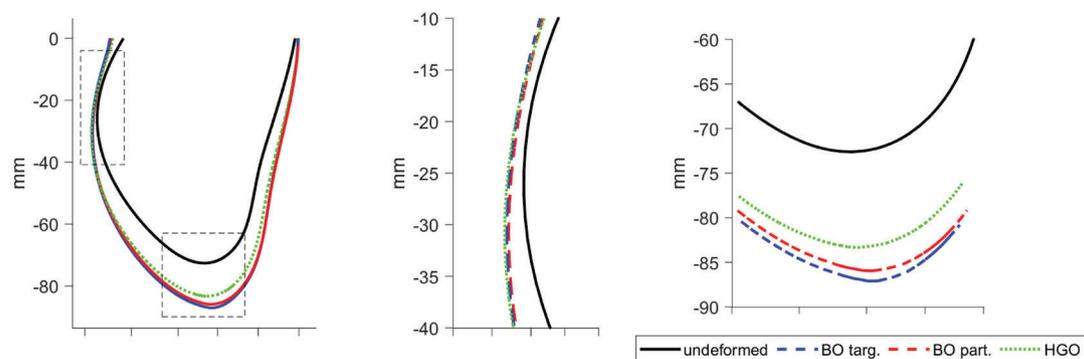
(A) HV A.



(B) HV B.



(C) HV C.



(D) HV D.

FIGURE 6 Basic study: deformations of the LV geometry at 8 mmHg implied by the final estimates from different algorithms. Left panels show a deformed reference line in the epicardial surface, see the red path in the LV geometry in Figure 3; middle and right panels show details of the inflexion point and the apical region, respectively

(13) (see Section 3.1 and 3.2.3) generally perform comparably, except for HV D, for which BO with a partial error surrogate converges faster and to a better point than BO with a target surrogate.

The main quantitative insights from Figure 4 are summarised in Table 3. We can see that y_{\min} obtained with BO are always lower than their HGO counterparts. However, in most cases (except HV A and BO with a target surrogate) they are obtained using more iterations than HGO. A fairer comparison is hence provided by y_{\min}^{HGO} . Even after less iterations BO still provides lower values of the objective function. This better performance of BO can be alternatively assessed using $i(y_{\min}^{\text{HGO}})$ which shows that for HV B, HV C and HV D we record at least as good performance of BO as for HGO using 4–7 times less iterations.

As mentioned at the beginning of this section, we focus on analysing stretch-stress curves and deformations of the LV geometry, presented in Figures 5 and 6, and not parameters estimates explicitly (the latter are nevertheless reported in Appendix A). We observe that the HGO algorithm predicts stiffer myocardium for high stretches along myocyte, while for low stretches the responses from BO and HGO are typically very comparable. The only exception is HV C for which HGO predicts much higher stiffness than BO along the whole range of stretches. On the other hand, for HV C both algorithms predict almost identical responses for stretches along the sheet direction. For the remaining subjects sheet direction stretches are associated with higher stiffness predicted by BO compared with HGO. The deformations of the LV geometry implied by the parameter estimates from different algorithms are generally very similar. The biggest differences between BO and HGO can be observed in the apical region (especially for HV B), while both BO versions lead to comparable deformations.

We conclude that the original HGO algorithm is easily outperformed by both variants of BO. Moreover, the mechanical response predicted by both BO and HGO for high stretches (i.e. > 1.15) can differ substantially. This could be partially explained by the limitation of using in vivo measurements, which are less informative about stiffening effects under high stretch as the myocardium is mostly operating in the linear regime of the stress-stretch curve in vivo,

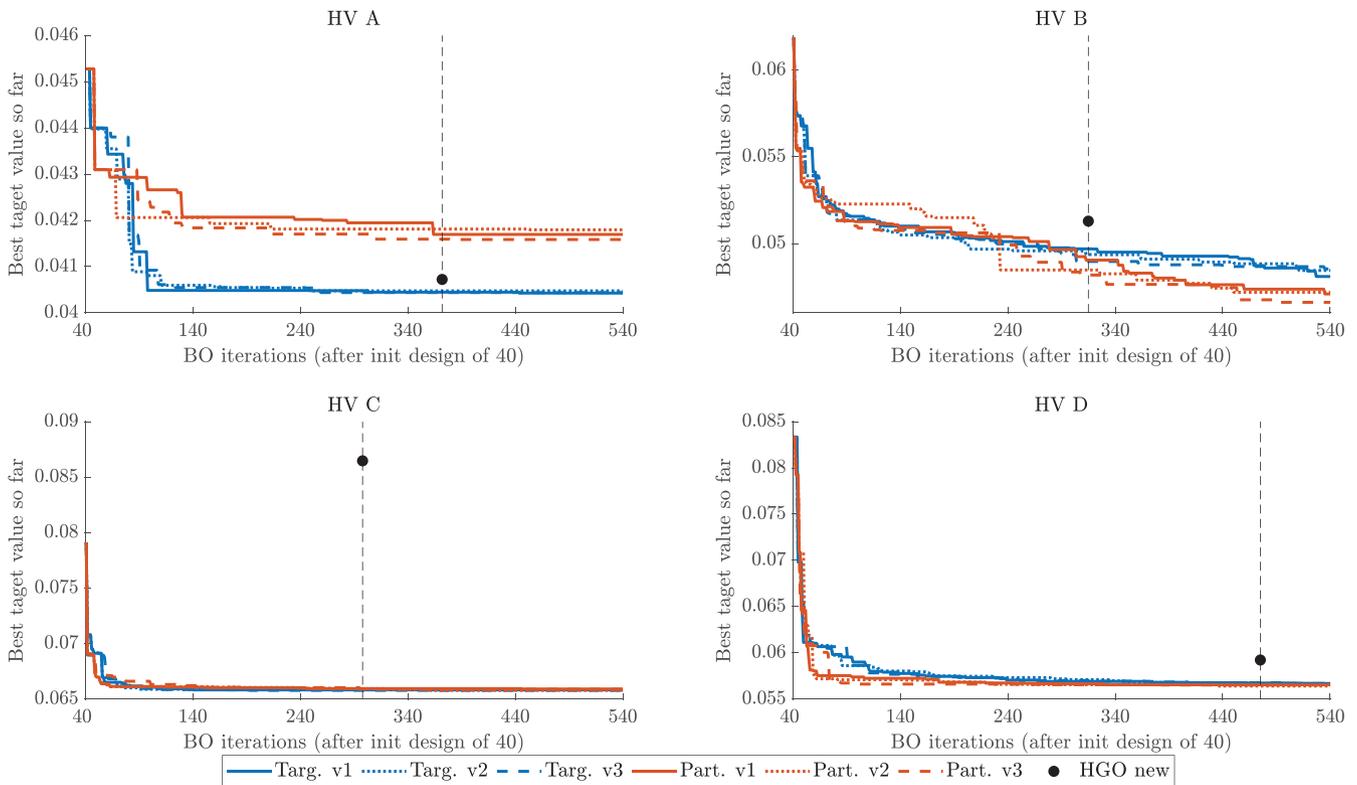


FIGURE 7 Klotz-curve study: convergence of the objective function $f_{O_2, \text{Klotz}}$ for Bayesian optimisation and the original HGO algorithm for four LV geometries (HV A, HV B, HV C, HV D). Horizontal axis: Bayesian optimisation iterations after 40 iterations for the initial design. Vertical axis: best value of the objective function f_{O_2} recorded so far. Black dot and horizontal dashed line: the final value of the objective function f_{O_2} for the HGO algorithm and the associated number of iterations. Bayesian optimisation with a target surrogate (targ.) and a partial error surrogate (part.), three independent runs (v1, v2, v3) in each version, together with the new version of the HGO algorithm (HGO new). For HGO, the Klotz curve error was computed using the forward simulator (not the emulator)

in particular for healthy LVs with a maximum stretch below 1.2.⁶² The latter confirms that inferring parameters characterising the nonlinear stiffening effects is challenging when based on in vivo data only. Therefore, introducing extra information about pressures higher than the normal end-diastolic pressure becomes necessary in order to characterise nonlinear responses at higher pressures. In the clinic, this could be done through stress CMR.⁶³ In this study, in our next application we will rely on the Klotz-law discussed in Section 2.2, which provides an alternative way to capture

TABLE 4 Klotz-curve study: convergence of the objective function for Bayesian optimisation and the updated HGO algorithm for four LV geometries (HV A, HV B, HV C, HV D)

HV	Algorithm/ run	\mathbf{y}_{\min}	\mathbf{y}_{\min}^S	i_{\min}	$\mathbf{y}(i_{\min}^{\text{HGO}})$	$\mathbf{y}^S(i_{\min}^{\text{HGO}})$	$i(\mathbf{y}_{\min}^{\text{HGO}})$	$E(\mathbf{y}_{\min})$	$S(\mathbf{y}_{\min})$	$E(i_{\min}^{\text{HGO}})$	$S(i_{\min}^{\text{HGO}})$	#Cr.
Target surrogate												
HV A	BO v1	0.040	0.040	443	0.040	0.040	97	223	222	224	223	3
	BO v2	0.040	0.040	248	0.040	0.040	110	225	224	225	224	2
	BO v3	0.040	0.040	258	0.040	0.040	107	224	223	224	223	3
	HGO new	0.041	0.041	371	0.041	0.041	371	–	224	–	224	0
HV B	BO v1	0.048	0.048	526	0.050	0.050	116	189	187	188	188	0
	BO v2	0.048	0.049	510	0.049	0.050	114	189	186	188	186	0
	BO v3	0.049	0.049	534	0.049	0.049	111	189	187	187	186	0
	HGO new	0.051	0.051	314	0.051	0.051	314	–	179	–	179	0
HV C	BO v1	0.066	0.066	140	0.066	0.066	40	229	227	229	227	1
	BO v2	0.066	0.066	327	0.066	0.066	40	230	227	230	227	1
	BO v3	0.066	0.066	192	0.066	0.066	40	230	228	230	228	0
	HGO new	0.086	0.086	297	0.086	0.086	297	–	200	–	200	0
HV D	BO v1	0.057	0.057	534	0.057	0.057	95	152	149	152	149	0
	BO v2	0.057	0.057	464	0.057	0.057	85	153	149	153	149	2
	BO v3	0.057	0.057	500	0.057	0.057	92	153	150	153	150	1
	HGO new	0.059	0.059	474	0.059	0.059	474	–	144	–	144	0
Partial error surrogate												
HV A	BO v1	0.042	0.042	363	0.042	0.042	∅	223	223	223	223	1
	BO v2	0.042	0.042	453	0.042	0.042	∅	224	224	222	222	2
	BO v3	0.042	0.042	379	0.042	0.042	∅	224	225	223	223	1
	HGO new	0.041	0.041	371	0.041	0.041	371	–	224	–	224	0
HV B	BO v1	0.047	0.050	534	0.049	0.052	97	185	179	182	178	1
	BO v2	0.047	0.048	449	0.048	0.052	208	188	184	183	177	1
	BO v3	0.047	0.049	486	0.048	0.050	88	185	179	184	180	0
	HGO new	0.051	0.051	314	0.051	0.051	314	–	179	–	179	0
HV C	BO v1	0.066	0.066	446	0.066	0.066	40	229	227	229	227	5
	BO v2	0.066	0.066	357	0.066	0.066	40	229	227	228	226	1
	BO v3	0.066	0.066	295	0.066	0.066	40	231	228	231	228	0
	HGO new	0.086	0.086	297	0.086	0.086	297	–	200	–	200	0
HV D	BO v1	0.057	0.057	330	0.057	0.057	55	152	150	152	150	0
	BO v2	0.056	0.057	434	0.056	0.057	58	151	148	151	148	2
	BO v3	0.057	0.056	382	0.057	0.056	73	153	150	153	150	3
	HGO new	0.059	0.059	474	0.059	0.059	474	–	144	–	144	0

Note: Three independent runs for BO (v1, v2, v3) together with the new version of HGO (HGO new). Missing values for simulator-based values (\mathbf{y}_{\min}^S , V_{30} with $S(\mathbf{y}_{\min})$ and $S(i_{\min}^{\text{HGO}})$) when the simulator crashed.

the nonlinear relationship between the LV cavity volume and the loading pressure within a large range, that is up to as high as 30 mmHg.

4.4 | Klotz-curve prior study

In our analysis of the basic study we have demonstrated that our proposed BO framework outperforms the original HGO algorithm²⁸ in terms of delivering lower values of the objective function using less invocations of the forward simulator. In this section we want to replace that benchmark with its updated version.⁴³ We can now use the full BO algorithm described in Listing 4 as the new HGO algorithm includes the Klotz-curve prior. This time we run each version of BO (with a target surrogate and a partial error surrogate (13), see Sections 3.1 and 3.2.3), three times (with independent runs based on the same initial design) to verify the robustness and stability of the procedure.

Figure 7 shows that BO again outperforms the HGO algorithm in terms of achieving lower values of the objective function $f_{O_2, \text{Klotz}}$ in fewer iterations. We can see that the updated HGO algorithm performs better than the original one *relatively* to BO, for example for HV C and HV D the new HGO algorithm converges closer to BO than in the basic study. Appendix B provides additional insights into the behaviour of the objective function $f_{O_2, \text{Klotz}}$ by decomposing it into the f_{O_2} component (i.e. errors for the 8 mmHg volume and strains) and the Klotz component (i.e. error for the 30 mmHg volume).

Table 4 summarises the results related to the convergence of the objective function. Since in the present study the performance of the HGO algorithm is more comparable to the one of BO, we decided to investigate whether the low values of the objective function for BO were not a consequence of a potential bias in the V_{30} emulator. The latter is used as a part of our framework aimed at achieving computational savings as discussed in Section 3.3, however the new HGO algorithm still uses the forward simulator to obtain the volume at 30 mmHg. We thus run the forward simulator for each run of BO at points θ corresponding to y_{\min} and $y(i_{\min}^{\text{HGO}})$, to obtain the true values of V_{30} at these points and hence true values of $f_{O_2, \text{Klotz}}$. We report the volumes in the columns indicated with V_{30} , with the $E(\cdot)$ values being the volumes from the emulator and the $S(\cdot)$ values being the volumes from the simulator. Similarly, we have two sets of the objective function values, obtained with the simulator (with superscript S) and the emulator. These results reveal that our V_{30} emulator provides very accurate predictions, with hardly any differences in the high pressure volumes and consequently in the values of the objective function compared to the simulator. Figure 8 illustrates this point using a Bland–Altman plot. The first four sets of points (from the left) show the differences between the simulator-based and emulator-based values of $f_{O_2, \text{Klotz}}$ against the average over both these values. Note that the scale of the y-axis is 10^{-3} . We can see that all the differences are indeed negligible so comparing the HGO algorithm against emulator-based runs of BO is valid.

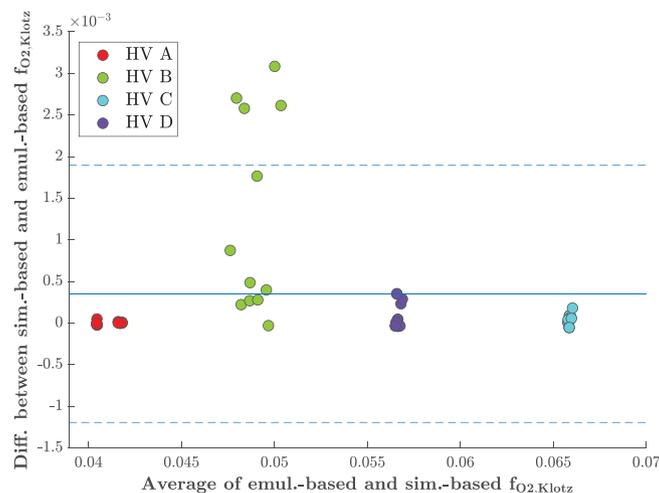


FIGURE 8 Klotz-curve: Bland–Altman plot comparing the Klotz objective function with V_{30} computed using the emulator and the forward simulator

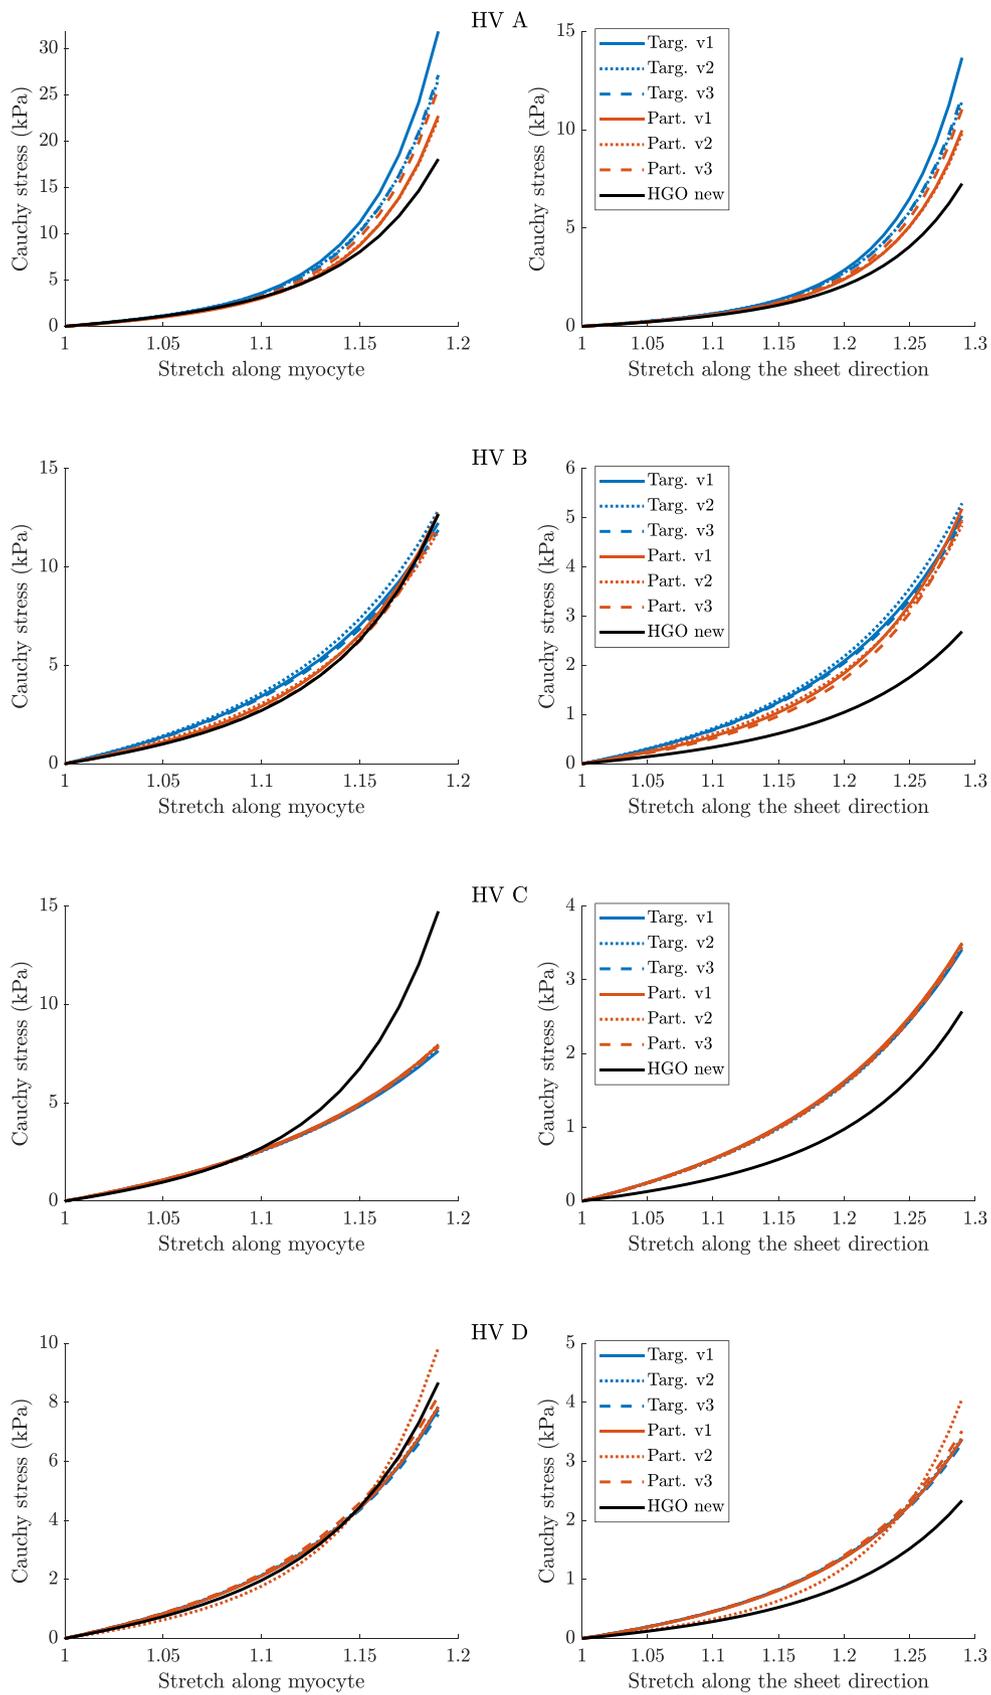
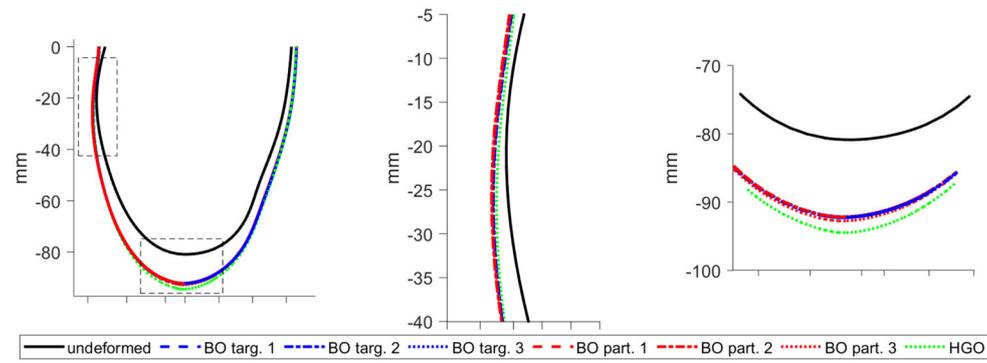
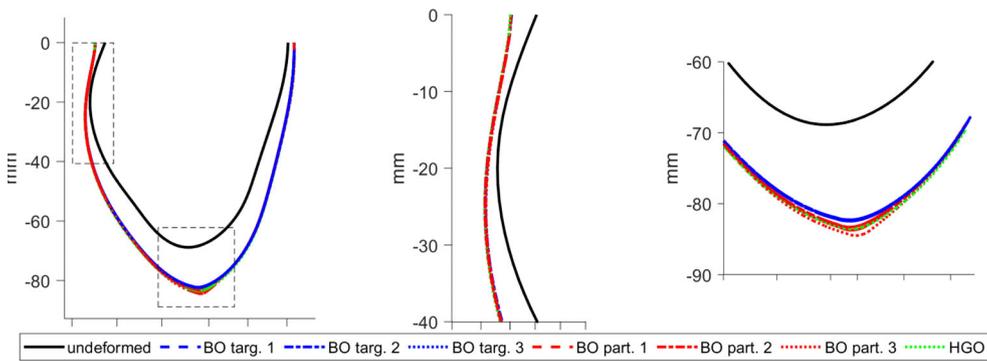


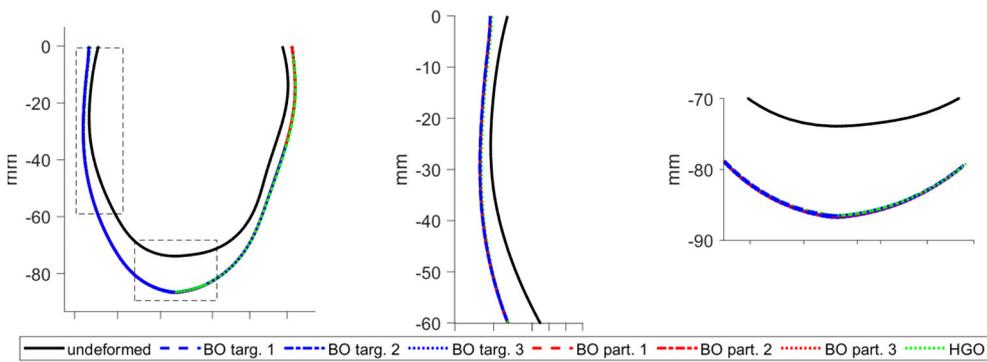
FIGURE 9 Klotz-curve study: stretch-stress curves for four LV geometries (HV A, HV B, HV C, HV D). Left: responses to stretches along the myocyte direction \mathbf{f}_0 , right: responses to stretches along the sheet direction \mathbf{s}_0 (see (2)). Bayesian optimisation with a target surrogate (targ.) and a partial error surrogate (part.), three independent runs (v1, v2, v3) in each version



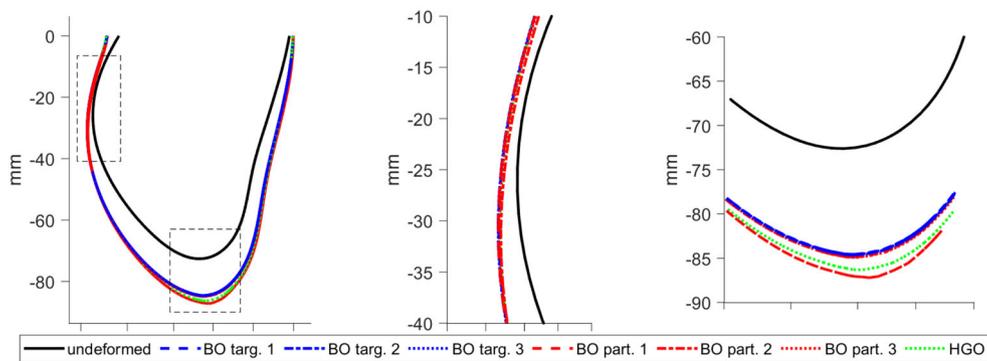
(A) HV A.



(B) HV B.



(C) HV C.



(D) HV D.

FIGURE 10 Klotz-curve study: deformations of the LV geometry at 8 mmHg implied by the final estimates from different algorithms. Left panels show a deformed reference line in the epicardial surface, see the red path in the LV geometry in Figure 3; middle and right panels show details of the inflexion point and the apical region, respectively

The stretch-stress curves in Figure 9 show that including the Klotz-curve prior leads to the results from BO and HGO being more comparable. For all the curves, except HV C along the sheet direction, the differences in predictions for small stretches (observable in vivo) observed in the basic study now with the Klotz-curve prior are considerably reduced. Moreover, comparing for example the curves along the myocyte direction for HV B with those from Figure 5 we observe that now the predicted stiffness for high stretches is noticeably reduced. A similar observation can be made for the deformations of the LV geometry, see Figure 10. For instance the difference in the apical region for HV B visible in Figure 6 have now disappeared.

5 | DISCUSSION

5.1 | Surrogates performance and heterogeneity

Comparing the performance of target and partial error surrogates does not reveal a clear pattern. In fact, the preferred surrogate type turns out to be case-specific (depending on the LV geometry and the data), so that one cannot know a priori which surrogate to use. Therefore, for robustness and extra uncertainty quantification, we would suggest trying both approaches, preferably in parallel.

Finally, we note that we would expect partial error surrogates to be superior to target surrogates when there is considerable heterogeneity in the strain data, which would be the case for MI patients. Logically, in a perfectly homogeneous case K partial surrogates would just learn the same mapping, which would lead to K same surrogates. The more heterogeneity in the strain data, the more need for flexibility in modelling individual errors.

5.2 | Acquisition function

The particular choice of the acquisition function is pivotal for the performance of BO. The Expected Improvement (EI) criterion applied in this paper has a number of advantages, most importantly it admits a closed form formula under Gaussian process surrogates and is widely used in the machine learning community. However, we note that the analytic tractability of EI becomes less advantageous as we do not use just EI as our AF, but EI_{con} , in which EI is augmented with the probit-based probability of unknown constraint satisfaction, making the whole AF not tractable. Moreover, EI might perform in a greedy way.⁶⁴

The development of better specifications for acquisition functions is an active research area in the BO literature.⁶ Two classes of alternative acquisition functions seem to be particularly relevant in our context. The first group consists of information-based policies, such as the Entropy Search^{65,66} (ES) criterion, that aim to infer the location of the unknown minimum by focusing on the posterior distribution of the minimiser. When deciding where to query ES takes into account what we expect to *learn* about the unknown minimum from evaluating at a particular point. This is different to EI, which evaluates where we believe the minimum is, without taking into account the effect of such a query on learning. ES is thus conceptually very appealing, however it is more challenging to implement in practice than EI. The second group of potentially useful AF is related the so-called portfolios of AFs.^{67,6} As demonstrated by our empirical studies, it is unlikely to develop a globally preferable acquisition strategy, that is performing best regardless of the given LV geometry and measurements. The portfolio approach considers multiple AFs and uses a meta-criterion (a high level AF) to select the final query point from those proposed by individual AFs.

5.3 | Computing time

Depending on the hardware used and the number of parallel tasks performed on a given clusters, computing time per BO iteration may vary considerably. For our studies, the mean computing time was approximately 9.5 minutes (on a typical Linux workstation being Dual 10 core Intel(R) Xeon(R), with a 2.50 GHz CPU and 64 GB RAM), while the median computing time was slightly over 7 minutes (the 10%, 25%, 75% and 90% quantiles were 5.3, 5.8, 10.8, 16.9 minutes, respectively). In our studies we considered 500 BO iterations, which for the median computing time would correspond to 2.5 days to complete optimisation. However, we note that premature convergence might be satisfactory in the clinic, as we observe rapid convergence of the objective function in the initial iterations. If we

conservatively assume 200 BO iterations, that would correspond to 1 day for optimisation with the median computing time.

Therefore, despite considerable efficiency gains compared to the HGO algorithm, BO is still likely to be too time-consuming to provide a viable tool for the clinical practice if optimisation is supposed to be performed independently for each subject, starting from scratch. Multi-task BO⁶⁸ could address this issue by leveraging prior knowledge acquired from searches over comparable domains, in our case from optimisation performed for previous subjects. Multi-task BO uses the data coming from different searches by processing them simultaneously and constructing a multi-task GP⁵⁵ to capture correlation between related tasks. This approach would allow the user to focus on parameter regions that have been found to be promising.

An alternative approach for fast parameter estimation would be to estimate parameters directly from clinical images without first obtaining strain or displacements using cardiac models, however, to our best knowledge, we are not aware of such studies. Cutting-edge physics-informed machine learning approaches may be the answer to that question.^{69,70,71} Nevertheless, we note that the goal of our group's research agenda is to make the whole inference pipeline automatic, that is to estimate the myocardial properties directly from CMR images. One of the stepping stones towards that goal was our work on an automatic framework for LV geometry prediction directly from CMR images based on convolutional neural networks.⁷²

5.4 | Biomechanics cardiac modelling

In this study, we have focused on the passive parameter estimation of in vivo hearts by only using conventional cine CMR images. The main reason is that knowing the mechanical properties of myocardial tissue is essential to understand the underlying mechanisms of various heart diseases, such as diastolic heart failure,⁷³ and it can further lead to personalised patient treatments, for example de-stiffening the heart.⁷⁴ In this study, we have adapted the framework of Bayesian optimisation to infer unknown parameters in the HO model, which outperformed the multi-step algorithm of Gao et al.²⁸ Future studies shall take into account the stiffness heterogeneity and material properties of different constituents that will provide a closer link to cardiac pathology compared to the global stiffness estimation. We also expect that the developed procedure in this study can be readily applied to other inverse problems by simply updating the loss function according to the available measurements, that is nonlinear-mechanics in general, or any parametric mathematical model.

A standalone single ventricle model during diastolic filling is chosen here by only including the essential features, such as subject-specific geometry, layered myofibre architecture, etc. Such standalone LV model is still being widely used in cardiac modelling communities.^{75,30,34} It shall be mentioned that the LV model can be further improved by including the right ventricle, the valves, the pericardium with more realistic boundary conditions, and further coupled with the blood flow and ectrophysiology. Such a complex model will inevitably introduce extra unknown parameters, which make the inference problem more challenging. For example, Strocchi et al.⁷⁶ used spatially varying robin boundary conditions to represent the pericardium in a four-chamber heart model; their results have shown that even though the pericardium can significantly affect ventricular motion, it has little effect on the local strains and wall thickness. The basal plane of the LV model is fixed because of lack of measured cardiac motion on the basal plane in our 2D cine CMR imaging dataset. This could be overcome by 3D tagged CMR imaging,²⁹ while it requires complex processing and extra scan time that is not routinely available.

Studies have demonstrated the high descriptive and predictive capabilities of the HO model in modelling cardiac dynamics.^{33,39,34} Due to the poor identifiability of some parameters in the original HO model,³² a few reduced HO models have been proposed in the context of parameter inference, such as the reduced formula with transverse isotropy used in Asner et al.²⁹ For example, by using a reduced HO model with only the first two terms, Hadjicharalambous et al.³⁰ inversely estimated a and a_f by fixing b and b_f , and further proposed to learn the ratio a/a_f . Instead, other studies introduced scaling parameters to simplify the parameter inference problem for the HO model. For example, Peirlinck et al.³⁴ introduced two scaling factors for a 's and b 's, respectively. This is the first step in the HGO algorithms. Considering that the HO-type material models are still being used widely in the cardiac modelling community, the Bayesian optimisation framework developed in this study can be readily extended to other HO-type strain energy functions, and further to other types of constitutive laws in soft tissue mechanics, such as the Fung-type strain energy function.⁵⁸

A further future research direction is to estimate spatially varied material stiffness for a given patient, especially for a patient with localised disease, such as myocardial infarction. It is believed that even for a healthy heart, myocardial stiffness will vary spatially. However, there is a lack of controlled experimental data to inform how heterogeneous stiffness is distributed in the myocardium. Should it be based on the widely used AHA-17 divisions or other segmental divisions (free wall, septum, etc.)? Potentially, CMR T1 mapping may provide key information about interstitial fibrosis,⁷⁷ which could further inform myocardial stiffness. For patients with myocardial infarction, experiments have demonstrated that the infarcted region, in general, has much higher stiffness than healthy myocardium. To model that, regional varied stiffness is usually used when modelling myocardial infarction.³ In a recent study, Balaban et al.⁷⁸ have demonstrated that heterogeneous elastic material properties in infarcted hearts could be estimated by using an adjoint-based data assimilation approach with a reduced HO model. Further validation studies with controlled experiments are needed. For patients with diastolic heart failure,^{23,73} the heart remodels its function and structure more globally compared to myocardial infarction, including the remodelling of myocardial stiffness, thus the global myocardial stiffness, as estimated in this study, is still relevant to those patients.

To sum up, this study should be considered as a step forward for passive parameter identification in cardiac mechanics using the HO model. Still many challenges need to be overcome before addressing the passive parameter identification in full for modelling cardiac mechanics, such as detailed cardiac motion measurements, unloaded residual-stressed geometry, micro-structure informed spatial heterogeneity, interactions with active contraction and electrophysiology, in vivo experiments, etc. Interested readers are referred to recent reviews in computational cardiac mechanics.^{1,79}

6 | CONCLUSIONS

We have proposed an efficient, Bayesian optimisation-based framework for parameter inference in a cardiac mechanic model of the left ventricle in diastole. We have demonstrated that BO converges to lower values of the two objective functions considered and requires less invocations of the associated forward simulator than established state-of-the-art iterative optimisation algorithms. We have developed a new approach to minimising a target function given as a sum of error terms based on approximating each of these terms individually via partial error surrogates. The findings of our empirical studies suggest that this approach may outperform the standard one based on target surrogates. However, which of the two variants is preferred in a specific case could be subject-dependent (depending on e.g. the LV geometry).

ACKNOWLEDGEMENTS

This work was funded by the UK Engineering and Physical Sciences Research Council (EPSRC), grant numbers EP/T017899/1 and EP/R018634/1.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

ORCID

Agnieszka Borowska  <https://orcid.org/0000-0001-9123-6227>

Hao Gao  <https://orcid.org/0000-0001-6852-9435>

Dirk Husmeier  <https://orcid.org/0000-0003-1673-7413>

REFERENCES

1. Mangion K, Gao H, Husmeier D, Luo XY, Berry C. Advances in computational modelling for personalised medicine after myocardial infarction. *Heart*. 2018;104(7):550-557.
2. Chabiniok R, Wang VY, Hadjicharalambous M, et al. Multiphysics and multiscale modelling, data-model fusion and integration of organ physiology in the clinic: ventricular cardiac mechanics. *Interface Focus*. 2016;6(2):20150083.
3. Gao H, Mangion K, Carrick D, Husmeier D, Luo XY, Berry C. Estimating prognosis in patients with acute myocardial infarction using personalized computational heart models. *Sci Rep*. 2017;7(13527):1.
4. Gao H, Aderhold A, Mangion K, Luo X, Husmeier D, Berry C. Changes and classification in myocardial contractile function in the left ventricle following acute myocardial infarction. *J R Soc Interface*. 2017;14(132):20170203.

5. Bathe KJ. *Finite Element Procedures*. 2nd ed. Prentice-Hall; 2014.
6. Shahriari B, Swersky K, Wang Z, Adams RP, de Freitas N. Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE*. 2016;104(1):148-175.
7. Wilson J, Hutter F, Deisenroth M. Maximizing acquisition functions for Bayesian optimization. *Adv Neural Inf Proces Syst*. 2018;31:9884-9895.
8. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning*. MIT Press; 2006.
9. Wang C, Duan Q, Gong W, Ye A, Di Z, Miao C. An evaluation of adaptive surrogate modeling based optimization with two benchmark problems. *Environ Model Softw*. 2014;60:167-179.
10. Jones DR, Schonlau M, Welch WJ. Efficient global optimization of expensive black-box functions. *J Glob Optim*. 1998;13(4):455-492.
11. Snoek J, Larochelle H, Adams RP. Practical Bayesian optimization of machine learning algorithms. *Adv Neural Inf Proces Syst*. 2012;25:2951-2959.
12. Colopy GW, Roberts SJ, Clifton DA. Bayesian optimization of personalized models for patient vital-sign monitoring. *IEEE J Biomed Health Inform*. 2018;22(2):301-310.
13. Kim M, Ding Y, Malcolm P, et al. Human-in-the-loop Bayesian optimization of wearable device parameters. *PLoS One*. 2017;12(9):e0184054.
14. Kim M, Liu C, Kim J, et al. Bayesian optimization of soft exosuits using a metabolic estimator stopping process. *2019 International Conference on Robotics and Automation (ICRA)*. IEEE; 2019:9173-9179.
15. Lorenz R, Johal M, Dick F, Hampshire A, Leech R, Geranmayeh F. A Bayesian optimization approach for rapidly mapping residual network function in stroke. *Brain*. 2021;144(7):2120-2134.
16. Travaglino S, Murdock K, Tran A, et al. Computational optimization study of Transcatheter aortic valve leaflet design using porcine and bovine leaflets. *J Biomech Eng*. 2019;142(1):1-8. doi:[10.1115/1.4044244](https://doi.org/10.1115/1.4044244)
17. Ayvali E, Srivatsan RA, Wang L, Roy R, Simaan N, Choset H. Using Bayesian optimization to guide probing of a flexible environment for simultaneous registration and stiffness mapping. *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE; 2016:931-936.
18. Yan Y, Pan J. Fast localization and segmentation of tissue abnormalities by autonomous robotic palpation. *IEEE Robot Automat Lett*. 2021;6(2):1707-1714. doi:[10.1109/LRA.2021.3058870](https://doi.org/10.1109/LRA.2021.3058870)
19. Perdikaris P, Karniadakis GE. Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *J R Soc Interface*. 2016;13(118):20151107.
20. Guvenir Torun S, Torun HM, Hansen HH, et al. Multicomponent material property characterization of atherosclerotic human carotid arteries through a Bayesian optimization based inverse finite element approach. *J Mech Behav Biomed Mater*. 2022;126:104996. doi:[10.1016/j.jmbbm.2021.104996](https://doi.org/10.1016/j.jmbbm.2021.104996)
21. Fan Y, Coll-Font J, van den Boomen M, et al. Characterization of exercise-induced myocardium growth using finite element modeling and Bayesian optimization. *Front Physiol*. 2021;12:694940.
22. Zile MR, Brutsaert DL. New concepts in diastolic dysfunction and diastolic heart failure: part I. *Circulation*. 2002;105(11):1387-1393.
23. Zile MR, Baicu CF, Gaasch WH. Diastolic heart failure—abnormalities in active relaxation and passive stiffness of the left ventricle. *N Engl J Med*. 2004;350(19):1953-1959.
24. Paulus WJ, Tschöpe C. A novel paradigm for heart failure with preserved ejection fraction: comorbidities drive myocardial dysfunction and remodeling through coronary microvascular endothelial inflammation. *J Am Coll Cardiol*. 2013;62(4):263-271.
25. Guccione J, McCulloch A, Waldman L. Passive material properties of intact ventricular myocardium determined from a cylindrical model. *J Biomech Eng*. 1991;113(1):42-55.
26. Remme EW, Hunter PJ, Smiseth O, et al. Development of an in vivo method for determining material properties of passive myocardium. *J Biomech*. 2004;37(5):669-678.
27. Xi J, Lamata P, Niederer S, et al. The estimation of patient-specific cardiac diastolic functions from clinical measurements. *Med Image Anal*. 2013;17(2):133-146. doi:[10.1016/j.media.2012.08.001](https://doi.org/10.1016/j.media.2012.08.001)
28. Gao H, Li WG, Cai L, Berry C, Luo XY. Parameter estimation in a Holzapfel–Ogden law for healthy myocardium. *J Eng Math*. 2015;95(1):231-248.
29. Asner L, Hadjicharalambous M, Chabiniok R, et al. Estimation of passive and active properties in the human heart using 3D tagged MRI. *Biomech Model Mechanobiol*. 2016;15(5):1121-1139.
30. Hadjicharalambous M, Asner L, Chabiniok R, et al. Non-invasive model-based assessment of passive left-ventricular myocardial stiffness in healthy subjects and in patients with non-ischemic dilated cardiomyopathy. *Ann Biomed Eng*. 2017;45(3):605-618.
31. Zhang Y, Wang VY, Morgan AE, et al. Finite-element based optimization of left ventricular passive stiffness in normal volunteers and patients after myocardial infarction: utility of an inverse deformation gradient calculation of regional diastolic strain. *J Mech Behav Biomed Mater*. 2021;119:104431. doi:[10.1016/j.jmbbm.2021.104431](https://doi.org/10.1016/j.jmbbm.2021.104431)
32. Holzapfel GA, Ogden RW. Constitutive modelling of passive myocardium: a structurally based framework for material characterization. *Philos Trans Royal Soc Lond A Math Phys Eng Sci*. 1902;2009(367):3445-3475.
33. Wang HM, Gao H, Luo XY, et al. Structure based finite strain modelling of the human left ventricle in diastole. *Int J Numer Methods Biomed Eng*. 2013;29(1):83-103.
34. Peirlinck M, Sahli Costabal F, Sack KL, et al. Using Machine Learning to Characterize Heart Failure across the Scales. *Biomech Model Mechanobiol*. 2019;18(6):1987-2001. doi:[10.1007/s10237-019-01190-w](https://doi.org/10.1007/s10237-019-01190-w)

35. Klotz S, Hay I, Dickstein ML, et al. Single-beat estimation of end-diastolic pressure-volume relationship: a novel method with potential for noninvasive application. *Am J Physiol Heart Circ Phys Ther.* 2006;291(1):H403-12.
36. Snoek JR. *Bayesian Optimization and Semiparametric Models with Applications to Assistive Technology.* PhD thesis. University of Toronto; 2013.
37. Gelbart MA, Snoek J, Adams RP. Bayesian optimization with unknown constraints. *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence.* AUAI Press; 2014:250-259.
38. Davies V, Noè U, Lazarus A, et al. Fast parameter inference in a biomechanical model of the left ventricle by using statistical emulation. *J R Stat Soc Ser C Appl Stat.* 2019;68(5):1555-1576.
39. Guan D, Ahmad F, Theobald P, Soe S, Luo XY, Gao H. On the AIC-based model reduction for the general Holzapfel–Ogden myocardial constitutive law. *Biomech Model Mechanobiol.* 2019;18(4):1213-1232. doi:10.1007/s10237-019-01140-6
40. Sack KL, Aliotta E, Ennis DB, et al. Construction and validation of subject-specific biventricular finite-element models of healthy and failing swine hearts from high-resolution DT-MRI. *Front Physiol.* 2018;9:539.
41. Noè U, Lazarus A, Gao H, et al. Gaussian process emulation to accelerate parameter estimation in a mechanical model of the left ventricle: a critical step towards clinical end-user relevance. *J R Soc Interface.* 2019;16(156):20190114.
42. Cerqueira MD, Weissman NJ, Dilsizian V, et al. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart: a statement for healthcare professionals from the cardiac imaging Committee of the Council on clinical cardiology of the American Heart Association. *Circulation.* 2002;105(4):539-542.
43. Gao H, Mangion K, Berry C, Luo XY. *Mathematical Modelling Acute Myocardial Infarction Using In Vivo Magnetic Resonance Imaging.* Proceeding of virtual physiological human conference; 2018.
44. Mojsejenko D, McGarvey JR, Dorsey SM, et al. Estimating passive mechanical properties in a myocardial infarction using MRI and finite element simulations. *Biomech Model Mechanobiol.* 2015;14(3):633-647.
45. Kirschner J, Mutný M, Hiller N, Ischebeck R, Krause A. Adaptive and safe Bayesian optimization in high dimensions via one-dimensional subspaces. *Proceedings of the 36th International Conference on Machine Learning.* PMLR; 2019:3429-3438.
46. Ugray Z, Lasdon L, Plummer J, Glover F, Kelly J, Martí R. Scatter search and local NLP solvers: a multistart framework for global optimization. *INFORMS J Comput.* 2007;19(3):328-340.
47. Glover F. A template for scatter search and path relinking. *European Conference on Artificial Evolution.* Springer; 1998:1-51.
48. Minka TP. Expectation propagation for approximate Bayesian inference. *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence.* Morgan Kaufmann Publishers Inc; 2001:362-369.
49. Rue H, Martino S, Chopin N. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J R Stat Soc Series B Stat Methodology.* 2009;71(2):319-392.
50. Jordan MI, Ghahramani Z, Jaakkola TS, Saul LK. An introduction to variational methods for graphical models. *Mach Learn.* 1999;37(2):183-233.
51. Nickisch H, Rasmussen CE. Approximations for binary Gaussian process classification. *J Mach Learn Res.* 2008;9(Oct):2035-2078.
52. Noè U, Chen W, Filippone M, Hill N, Husmeier D. In: Bracciali A, Caravagna G, Gilbert D, Tagliaferri R, eds. *Inference in a Partial Differential Equations Model of Pulmonary Arterial and Venous Blood Circulation Using Statistical Emulation. Computational Intelligence Methods for Bioinformatics and Biostatistics* Springer International Publishing; 2017.
53. Conti S, O'Hagan A. Bayesian emulation of complex multi-output and dynamic computer models. *J Stat Plan Infer.* 2010;140(3):640-651.
54. Conti S, Gosling JP, Oakley JE, O'Hagan A. Gaussian process emulation of dynamic computer codes. *Biometrika.* 2009;96(3):663-676.
55. Bonilla EV, Chai KM, Williams C. Multi-task Gaussian process prediction. *Advances in Neural Information Processing Systems.* Vol 21. Curran Associates, Inc; 2008:153-160.
56. McKay MD, Beckman RJ, Conover WJ. Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Dent Tech.* 1979;21(2):239-245.
57. Sommer G, Haspinger DC, Andrä M, et al. Quantification of shear deformations and corresponding stresses in the biaxially tested human myocardium. *Ann Biomed Eng.* 2015;43(10):2334-2348.
58. Land S, Gurev V, Arens S, et al. Verification of cardiac mechanics software: benchmark problems and solutions for testing active and passive material behaviour. *Proc Royal Soc A Math Phys Eng Sci.* 2015;471(2184):20150641.
59. Campos J, Sundnes J, Dos Santos R, Rocha B. Uncertainty quantification and sensitivity analysis of left ventricular function during the full cardiac cycle. *Phil Trans R Soc A.* 2020;378(2173):20190381.
60. Campos JO, Sundnes J, Dos Santos RW, Rocha BM. Effects of left ventricle wall thickness uncertainties on cardiac mechanics. *Biomech Model Mechanobiol.* 2019;18(5):1415-1427.
61. Levrero-Florencio F, Margara F, Zacur E, et al. Sensitivity analysis of a strongly-coupled human-based electromechanical cardiac model: effect of mechanical parameters on physiologically relevant biomarkers. *Comput Methods Appl Mech Eng.* 2020;361:112762.
62. Mangion K, Gao H, McComb C, et al. A novel method for estimating myocardial strain: assessment of deformation tracking against reference magnetic resonance methods in healthy volunteers. *Sci Rep.* 2016;6(1):38774.
63. Hundley WG, Morgan TM, Neagle CM, Hamilton CA, Rerkpattanapipat P, Link KM. Magnetic resonance imaging determination of cardiac prognosis. *Circulation.* 2002;106(18):2328-2333.
64. Qin C, Klabjan D, Russo D. Improving the expected improvement algorithm. *Advances in Neural Information Processing Systems.* Curran Associates Inc; 2017:5381-5391.
65. Hennig P, Schuler CJ. Entropy search for information-efficient global optimization. *J Mach Learn Res.* 2012;13(1):1809-1837.

66. Hernández-Lobato JM, Hoffman MW, Ghahramani Z. Predictive entropy search for efficient global optimization of black-box functions. *In: Advances in Neural Information Processing Systems*. 2014;27:918-926.
67. Hoffman M, Brochu E, de Freitas N. Portfolio allocation for Bayesian optimization. *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*. AUAI Press; 2011:327-336.
68. Swersky K, Snoek J, Adams RP. Multi-task Bayesian optimization. *Advances in Neural Information Processing Systems*. Curran Associates, Inc; 2013:2004-2012.
69. Alber M, Buganza Tepole A, Cannon WR, et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *Npj Digit Med*. 2019;2(1):115. doi:10.1038/s41746-019-0193-y
70. Corral-Acero J, Margara F, Marciniak M, et al. The ‘digital twin’ to enable the vision of precision cardiology. *Eur Heart J*. 2020;41:4556–4564. doi:10.1093/eurheartj/ehaa159
71. Buoso S, Joyce T, Kozerke S. Personalising left-ventricular biophysical models of the heart using parametric physics-informed neural networks. *Med Image Anal*. 2021;71:102066.
72. Romaszko L, Borowska A, Lazarus A, et al. Neural network-based left ventricle geometry prediction from CMR images with application in biomechanics. *Artif Intell Med*. 2021;119:102140. doi:10.1016/j.artmed.2021.102140
73. Zile MR, Baicu CF, Ikonomidis JS, et al. Myocardial stiffness in patients with heart failure and a preserved ejection fraction: contributions of collagen and titin. *Circulation*. 2015;131(14):1247-1259.
74. Dabiri Y, Sack KL, Shaul S, et al. Intramyocardial injections to De-stiffen the heart: a subject-specific in Silico approach. *Mol Cell Biomech*. 2019;16(3):185-197.
75. Maso Talou GD, Babarenda Gamage TP, Sagar M, Nash MP. Deep learning over reduced intrinsic domains for efficient mechanics of the left ventricle. *Front Phys*. 2020;8:30. doi:10.3389/fphy.2020.00030
76. Strocchi M, Gsell MA, Augustin CM, et al. Simulating ventricular systolic motion in a four-chamber heart model with spatially varying robin boundary conditions to model the effect of the pericardium. *J Biomech*. 2020;101:109645. doi:10.1016/j.jbiomech.2020.109645
77. Everett R, Stirrat C, Semple S, Newby D, Dweck M, Mirsadraee S. Assessment of myocardial fibrosis with T1 mapping MRI. *Clin Radiol*. 2016;71(8):768-778.
78. Balaban G, Finsberg H, Funke S, et al. In vivo estimation of elastic heterogeneity in an infarcted human heart. *Biomech Model Mechanobiol*. 2018;17(5):1317-1329.
79. Peirlinck M, Costabal FS, Yao J, et al. Precision medicine in human heart modeling. *Biomech Model Mechanobiol*. 2021;20:1-29.
80. Tibshirani RJ, Efron B. An introduction to the bootstrap. *Monogr Stat Appl Probab*. 1993;57:1-436.

How to cite this article: Borowska A, Gao H, Lazarus A, Husmeier D. Bayesian optimisation for efficient parameter inference in a cardiac mechanics model of the left ventricle. *Int J Numer Meth Biomed Engng*. 2022; 38(5):e3593. doi:10.1002/cnm.3593

APPENDIX A: RESULTS ON PARAMETERS

A.1 | Basic study

Table A1.

TABLE A1 Basic setting: final optimised values of the eight parameters of the HO law for Bayesian optimisation and the original HGO algorithm (HGO old) for four different LV geometries (HV A, HV B, HV C, HV D), Bayesian optimisation with a target surrogate (targ.) and a partial error surrogate (part.)

HV	Algorithm	a	b	a_f	b_f	a_s	b_s	a_{fs}	b_{fs}
HV A	BO targ.	0.257	3.715	8.636	0.409	1.784	0.166	0.045	0.371
	BO part.	0.258	3.728	8.581	0.360	1.773	0.146	0.074	0.619
	HGO old	0.100	5.405	6.172	7.114	1.117	2.582	0.123	5.650
HV B	BO targ.	0.018	0.260	3.174	7.859	0.656	3.196	1.444	12.016
	BO part.	0.018	0.260	4.344	4.780	0.898	1.944	1.458	12.131
	HGO old	0.100	5.766	3.090	11.252	0.372	2.582	0.041	5.650
HV C	BO targ.	0.207	2.991	5.567	0.294	1.150	0.119	1.546	12.864
	BO part.	0.245	3.532	4.481	2.158	0.926	0.877	1.550	12.900
	HGO old	0.100	1.960	8.447	1.811	1.117	0.516	0.123	1.130
HV D	BO targ.	0.162	2.342	4.834	0.293	0.999	0.119	0.676	5.623
	BO part.	0.174	2.519	4.982	0.878	1.029	0.357	0.169	1.406
	HGO old	0.100	5.405	2.626	8.791	0.372	2.582	0.041	5.650

A.2 | Klotz-curve study

Table A2.

TABLE A2 Klotz-curve study: final optimised values of the parameters of the HO law for Bayesian optimisation and the updated HGO algorithm for four LV different geometries (HV A, HV B, HV C, HV D)

HV	Algorithm/run	a	b	a_f	b_f	a_s	b_s	a_{fs}	b_{fs}
Target surrogate									
HV A	BO v1	0.279	4.033	4.326	10.536	0.894	4.284	0.034	0.287
	BO v2	0.288	4.154	4.530	9.225	0.936	3.751	0.119	0.994
	BO v3	0.299	4.319	4.307	9.624	0.890	3.913	0.031	0.258
	HGO	0.205	6.404	4.411	7.123	0.819	2.912	0.368	6.767
HV B	BO v1	0.166	2.395	5.727	3.382	1.183	1.375	0.139	1.156
	BO v2	0.166	2.398	5.977	3.432	1.235	1.395	0.039	0.324
	BO v3	0.177	2.554	5.605	3.329	1.158	1.353	0.120	0.999
	HGO	0.100	2.113	4.126	5.506	0.555	1.655	0.250	3.847
HV C	BO v1	0.254	3.673	4.463	1.995	0.922	0.811	1.550	12.899
	BO v2	0.241	3.476	4.357	2.287	0.900	0.930	1.519	12.640
	BO v3	0.242	3.496	4.515	2.156	0.933	0.877	1.548	12.886
	HGO	0.126	4.098	3.897	6.689	0.481	1.750	0.216	4.067
HV D	BO v1	0.213	3.075	3.458	3.591	0.714	1.460	0.087	0.727
	BO v2	0.207	2.989	3.479	3.641	0.719	1.480	0.062	0.515
	BO v3	0.213	3.070	3.473	3.446	0.718	1.401	0.061	0.511
	HGO	0.108	3.618	3.043	5.041	0.458	1.684	0.206	3.915

(Continues)

TABLE A2 (Continued)

HV	Algorithm/run	a	b	a_f	b_f	a_s	b_s	a_{fs}	b_{fs}
Partial error surrogate									
HV A	BO v1	0.319	4.605	3.819	9.279	0.789	3.773	0.619	5.150
	BO v2	0.313	4.516	3.977	8.964	0.822	3.645	0.660	5.489
	BO v3	0.270	3.897	4.046	9.684	0.836	3.937	0.586	4.875
	HGO	0.205	6.404	4.411	7.123	0.819	2.912	0.368	6.767
HV B	BO v1	0.105	1.521	4.564	4.942	0.943	2.009	0.515	4.287
	BO v2	0.139	2.007	4.958	4.018	1.024	1.634	0.452	3.764
	BO v3	0.076	1.101	4.274	5.112	0.883	2.078	0.628	5.227
	HGO	0.100	2.113	4.126	5.506	0.555	1.655	0.250	3.847
HV C	BO v1	0.246	3.555	4.536	2.121	0.937	0.862	1.433	11.930
	BO v2	0.238	3.432	4.398	2.335	0.909	0.950	1.550	12.900
	BO v3	0.239	3.446	4.407	2.225	0.910	0.905	1.509	12.562
	HGO	0.126	4.098	3.897	6.689	0.481	1.750	0.216	4.067
HV D	BO v1	0.201	2.897	3.421	3.734	0.707	1.518	0.175	1.457
	BO v2	0.087	1.258	2.534	6.858	0.523	2.789	0.304	2.531
	BO v3	0.194	2.797	3.534	3.821	0.730	1.554	0.130	1.080
	HGO	0.108	3.618	3.043	5.041	0.458	1.684	0.206	3.915

Note: For each BO variant (target surrogate, partial error surrogate three independent runs are reported (v1, v2, v3).

APPENDIX B: DECOMPOSITION OF $f_{O_2, \text{Klotz}}$ INTO f_{O_2} AND THE KLOTZ COMPONENT

Figure B1 decomposes the incumbent convergence plot from Figure 7 into the f_{O_2} component, that is errors for the 8 mmHg volume and strains, (top) and the Klotz component, that is error for the 30 mmHg volume (bottom). When compared with Figure 4 from the Basic study in Section 4.3 we can see that the f_{O_2} component is larger in the Klotz study (Section 4.4). Besides, the magnitude of the Klotz component quickly becomes negligible compared to the f_{O_2} component

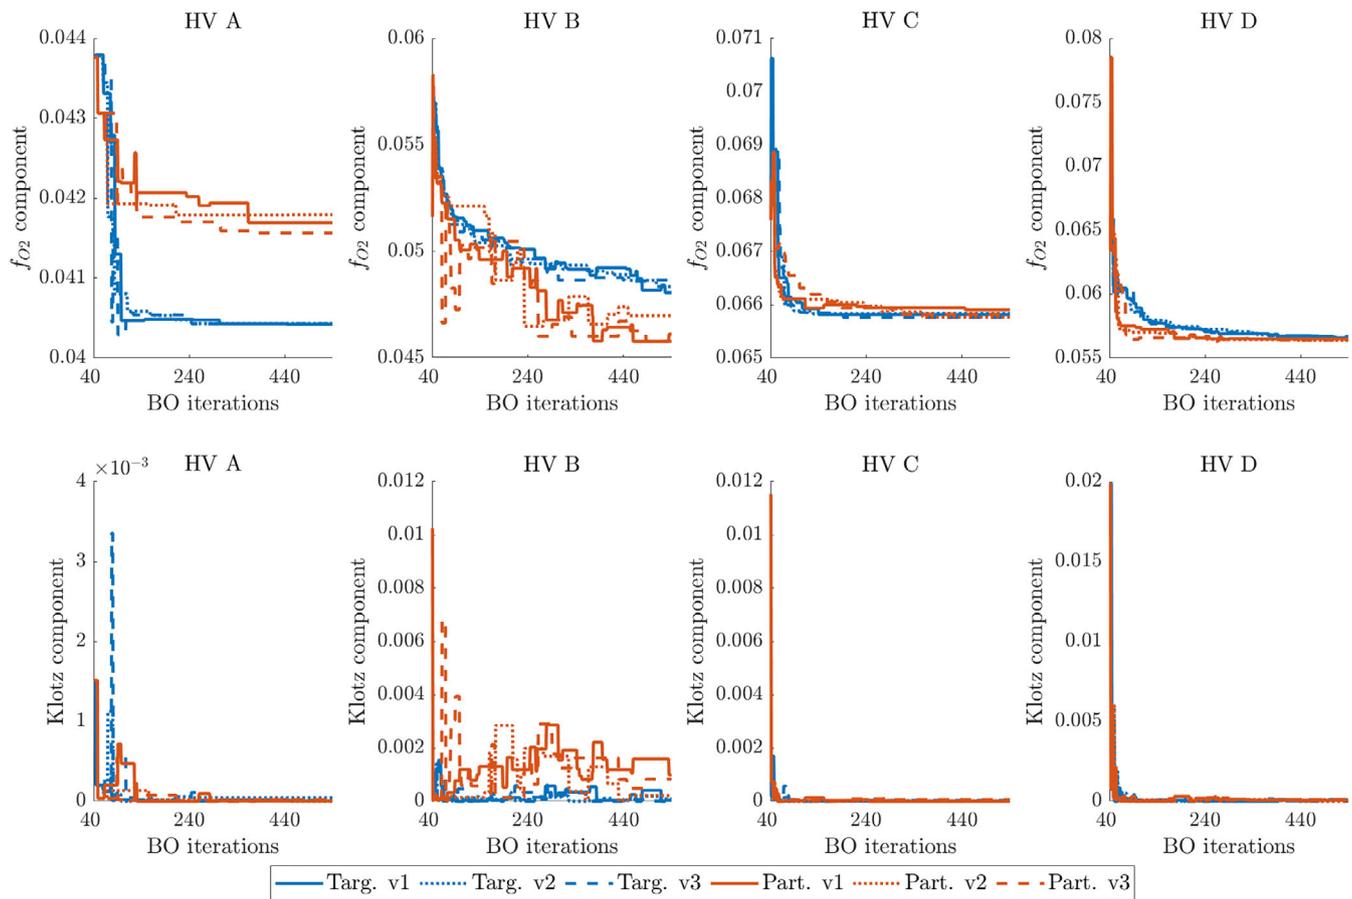


FIGURE B1 Klotz study: decomposition of the incumbent trajectories from Figure 7 based on $f_{O_2, \text{Klotz}}$ from (11) into f_{O_2} from (6) (top) and the Klotz component (bottom). Bayesian optimisation with a target surrogate (targ.) and a partial error surrogate (part.), three independent runs (v1, v2, v3) in each version

APPENDIX C: SYNTHETIC DATA STUDY

In this appendix, we report the results of a synthetic data study aimed at verifying the theoretical accuracy of the proposed method. We consider the geometry of HV A and generate ten synthetic data sets by perturbing the estimate for HV A corresponding to the lowest value of the objective function with Gaussian noise. This way we obtain synthetic data sets with known underlying ground-truth (GT) parameter values, which we then estimate with BO.

Figure C2 presents the results from this experiment in terms of errors in the output space (by means of the objective function) and in the parameter space (by means of the difference between the GT value and its estimate). Figure C3 illustrates the convergence of the incumbent values. We can see that the convergence of the objective function is rapid and to very low values (compared with the real data studies in Sections 4.3 and 4.4), of order 10^{-5} . Moreover, we notice that on average the difference between the GT parameter values and their estimates is zero. As expected, θ_1 and θ_2 are estimated with high precision, while θ_3 and θ_4 are harder to identify, hence there is more variance in the corresponding estimation errors.

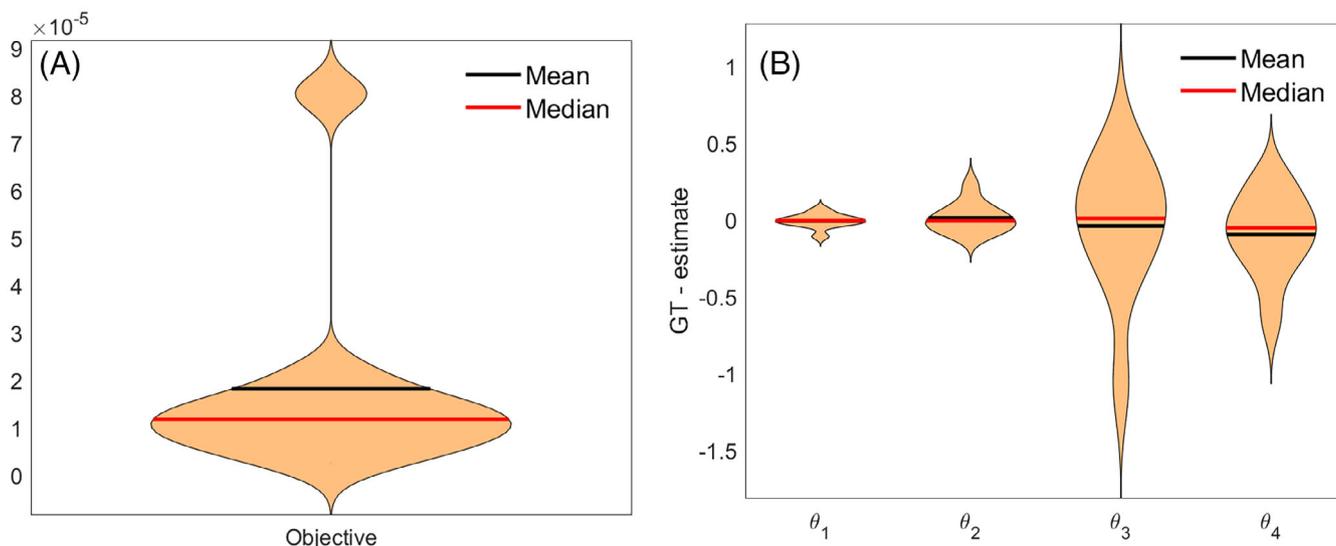


FIGURE C2 Synthetic data study: estimation errors. Left: objective function; right: parameters

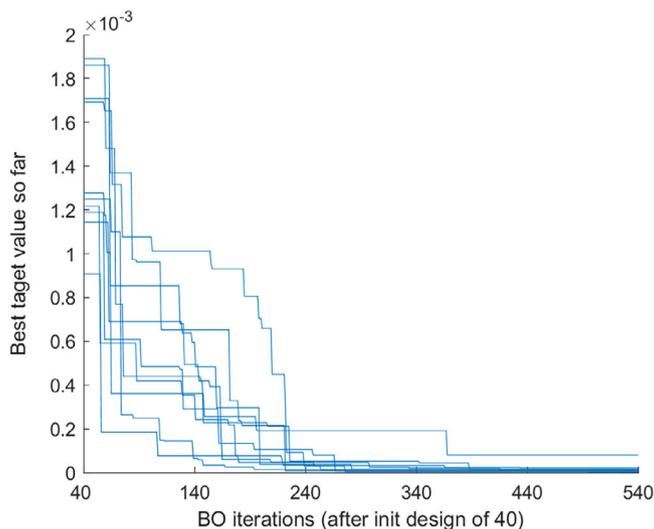
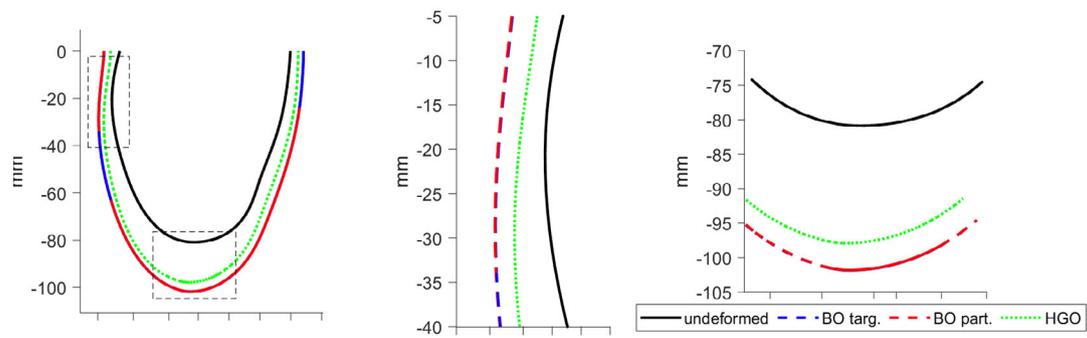


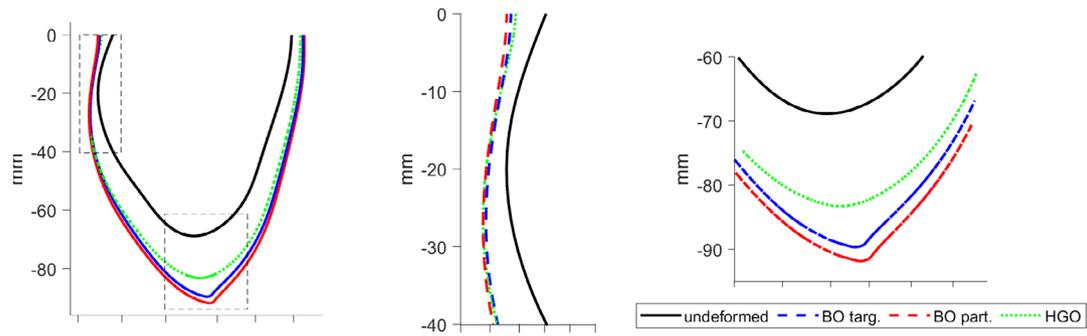
FIGURE C3 Synthetic data study: incumbent convergence for ten synthetic data sets

APPENDIX D: LV GEOMETRY DEFORMATIONS

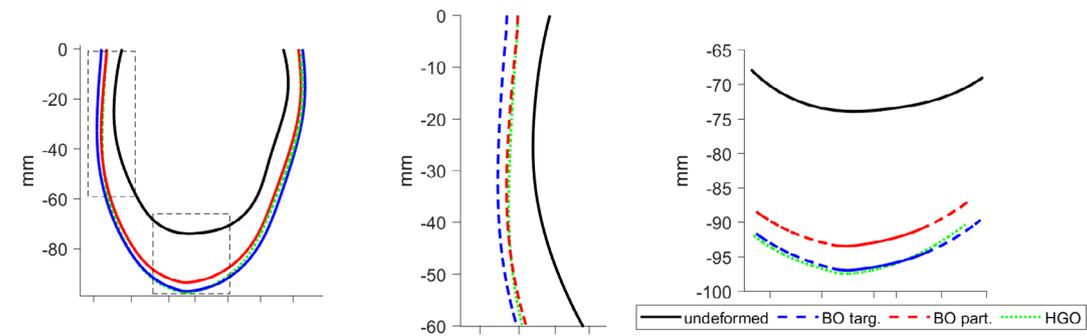
Figures D4 and D5 provide high-pressure (at 30 mmHg) counterparts of low pressure (8 mmHg) Figures 6 and 10, respectively. Applying much higher pressure results in differences between deformations implied by estimates from different algorithms being exaggerated. Still, the deformations implied by BO and HGO estimates are generally consistent, with most noticeable differences appearing in the apical region.



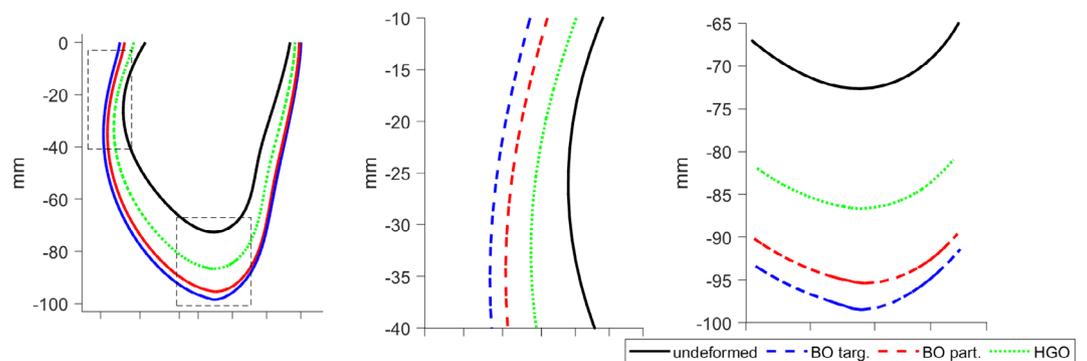
(A) HV A.



(B) HV B.

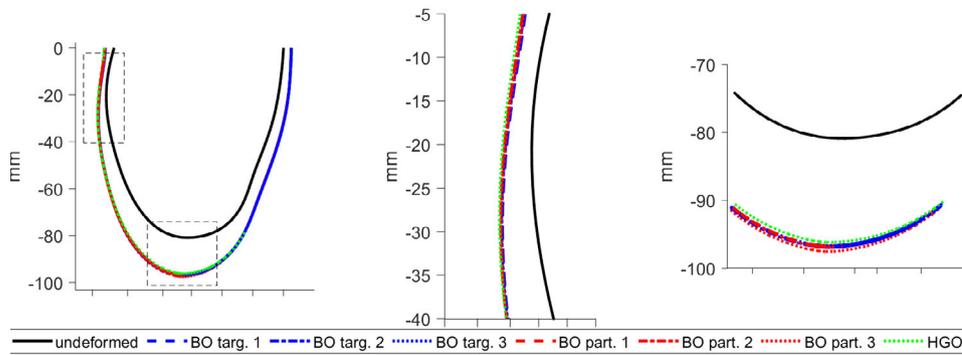


(C) HV C.

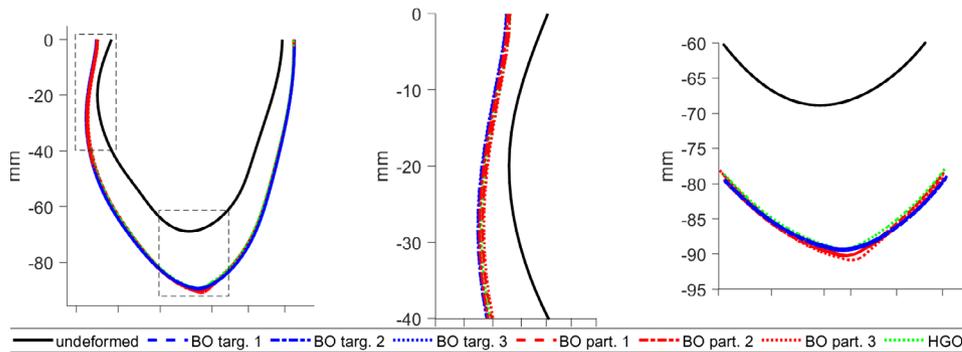


(D) HV D.

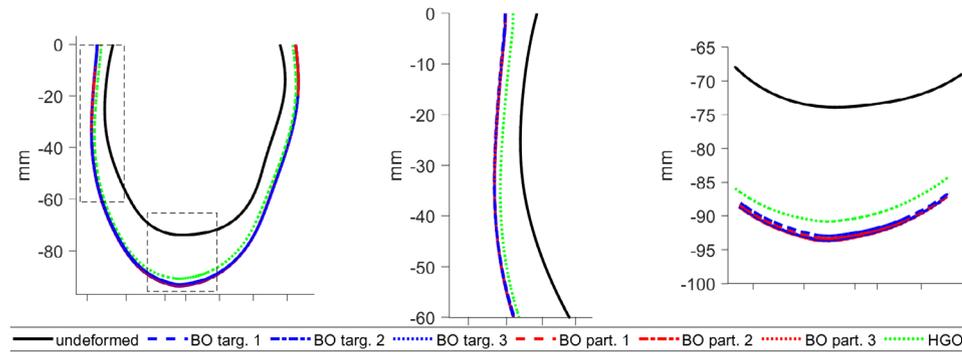
FIGURE D4 Basic study: deformations of the LV geometry at 30 mmHg implied by the final estimates from different algorithms. Left panels show a deformed reference line in the epicardial surface, see the red path in the LV geometry in Figure 3; middle and right panels show details of the inflexion point and the apical region, respectively



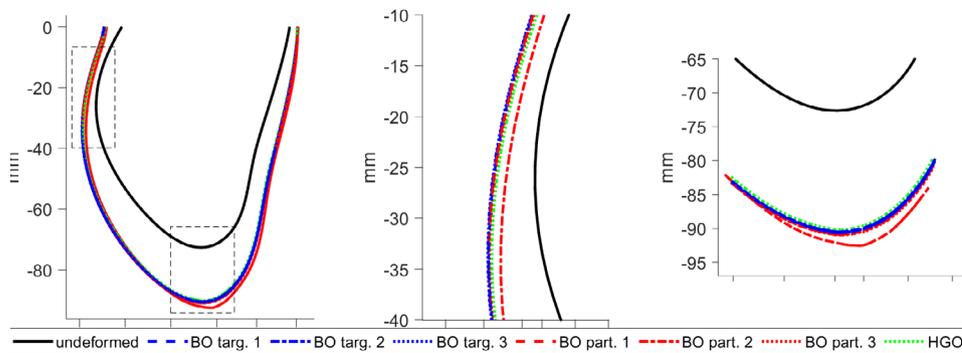
(A) HV A.



(B) HV B.



(C) HV C.



(D) HV D.

FIGURE D5 Klotz-curve study: deformations of the LV geometry at 30 mmHg implied by the final estimates from different algorithms. Left panels show a deformed reference line in the epicardial surface, see the red path in the LV geometry in Figure 3; middle and right panels show details of the inflexion point and the apical region, respectively

APPENDIX E: UNCERTAINTY QUANTIFICATION

Recently, there has been an increasing awareness in the cardiac modelling community of the importance of uncertainty quantification and sensitivity analysis for model parameter estimates.^{59,60,61} Below we describe three principled directions one can take to assess the uncertainty of parameter estimates: an asymptotic approach, a computational frequentist approach, and a computational Bayesian approach. We also discuss the advantage of Gaussian processes (see Section 3.1) used in this work over polynomial chaos methods in the context of BO.

E.1 | Asymptotic approach

The asymptotic approach is based on the Fisher information inequality and the Cramer-Rao bound, whereby the variance of an estimated parameter is lower-bounded by the inverse of the negative second derivative of the log likelihood with respect to this parameter. For additive independent Gaussian noise, the log likelihood is just a rescaled version of our sum-of-squares objective function; so it can easily be obtained. To generalise this to more than one parameter, we need to get the diagonal elements of the inverse of the negative Hessian, that is the matrix of negative second derivatives of the log likelihood. To compute the second derivatives, we can exploit the fact that our objective function is approximated by a Gaussian process, and that a Gaussian process is closed under differentiation, that is if the kernel function of the Gaussian process is twice differentiable, we will get these derivatives in closed form. The advantage of this approach is that it is analytically tractable and fast. The disadvantage is that it is based on asymptotic assumptions. These may not hold for sparse data, with the consequence that the lower bound may not be tight. For that reason, we prefer computationally more expensive non-asymptotic methods.

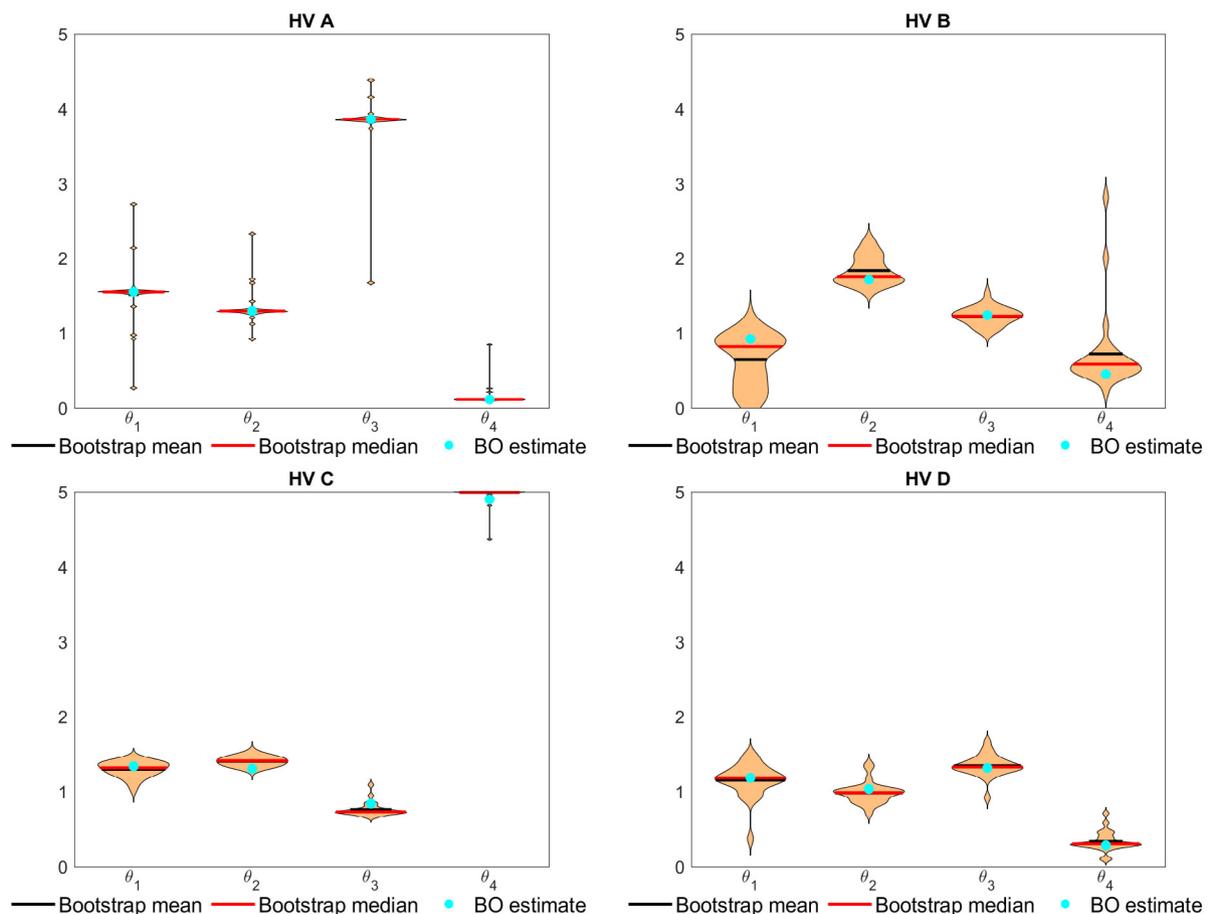


FIGURE E6 Klotz-curve study: violin plots for the bootstrapped parameter estimates from run v1 with target surrogates

TABLE E3 Klotz-curve study: median absolute deviations (MADs) for the bootstrapped parameter estimates from run v1 with target surrogates

Case	θ_1	θ_2	θ_3	θ_4
HV A	0.204	0.140	0.176	0.070
HV B	0.302	0.159	0.087	0.313
HV C	0.081	0.052	0.062	0.071
HV D	0.145	0.100	0.097	0.094

E.2 | Computational frequentist approach

This approach is based on a residual bootstrapping procedure⁸⁰ to estimate the parameter estimation uncertainty. Take the predictions (standardised circumferential strains and end-diastole LV volume) obtained from the final parameter vector estimate θ^* , $\xi(\theta^*)$, and compute the residuals by comparison with the data. Let us call the set of residuals E . We then generate an ensemble of hypothetical data, D_1, \dots, D_K , by drawing residuals with replacement from E and adding these values to $\xi(\theta^*)$. For each of these bootstrap replica, D_1, \dots, D_K , we carry out a local iterative optimisation (e.g. using conjugate gradients), starting from θ^* . This procedure returns a set of bootstrap estimates, $\{\theta_1, \theta_2, \dots, \theta_K\}$, which can be used for uncertainty quantification. We use the median absolute deviation (MAD) as a simple summary statistic; this is a robust measure of variability that is less susceptible to outliers than the standard deviation. We multiply the MAD with a constant scale factor of $k = 1.4826$, which makes it a consistent estimator for the estimation of the standard deviation of a normally distributed random variable.

Figure E6 presents examples of bootstrap results by means of violin plots for BO estimates (with target surrogates, run v1) from the Klotz-curve study. Each figure is based on $K = 25$ bootstrap samples. Recall that the parameters come from a compact domain $[0.1, 5]^4$ hence any apparent truncation of a violin plot (e.g. θ_4 for HV C) is caused by the underlying parameter being estimated close to its boundary. We can see that the uncertainty in parameter estimates for HV A is very small as visualised by hardly visible “violins”; we can observe a few outliers, however. In contrast, the uncertainty for HV B is much larger, as depicted by larger “violin” areas. Table E3 reports the associated MADs

E.3 | Computational Bayesian approach

We can use the emulated objective function to approximate the log likelihood, define appropriate prior distributions for the parameters and apply standard Markov chain Monte Carlo techniques to approximately sample from the posterior distribution. The practical disadvantage is that this is an intrinsically sequential approach, whereas the optimisation runs for the bootstrap approach can be parallelised.

E.4 | Gaussian processes versus polynomial chaos

Polynomial chaos methods have been widely used as an alternative to Gaussian processes for emulation and surrogate modelling.^{59,60,61} However, for Bayesian optimisation, the emulation of the objective function is not sufficient. We also need a reliable estimate of the variance of the approximated objective function, which is needed for the acquisition function (see Section 3.2). Figure 4 in Shahriari et al.⁶ shows several examples of methods that emulate the objective function well, but provide poor estimates of its variance. This will lead to a poor trade-off of exploration versus exploitation and is unsuitable for Bayesian optimisation. For that reason, we have used Gaussian processes rather than polynomial chaos expansion in our work.

APPENDIX F: SIMULATOR CRASHES

In our applications, crashes are caused by implausible parameter sets, by which we mean that the “crash” parameter sets result in an excessive distortion of the underlying LV geometry, for example, when all parameters are close to the lower bounds, suggesting a very soft myocardium. An increased number of crashes for BO compared to HGO is a

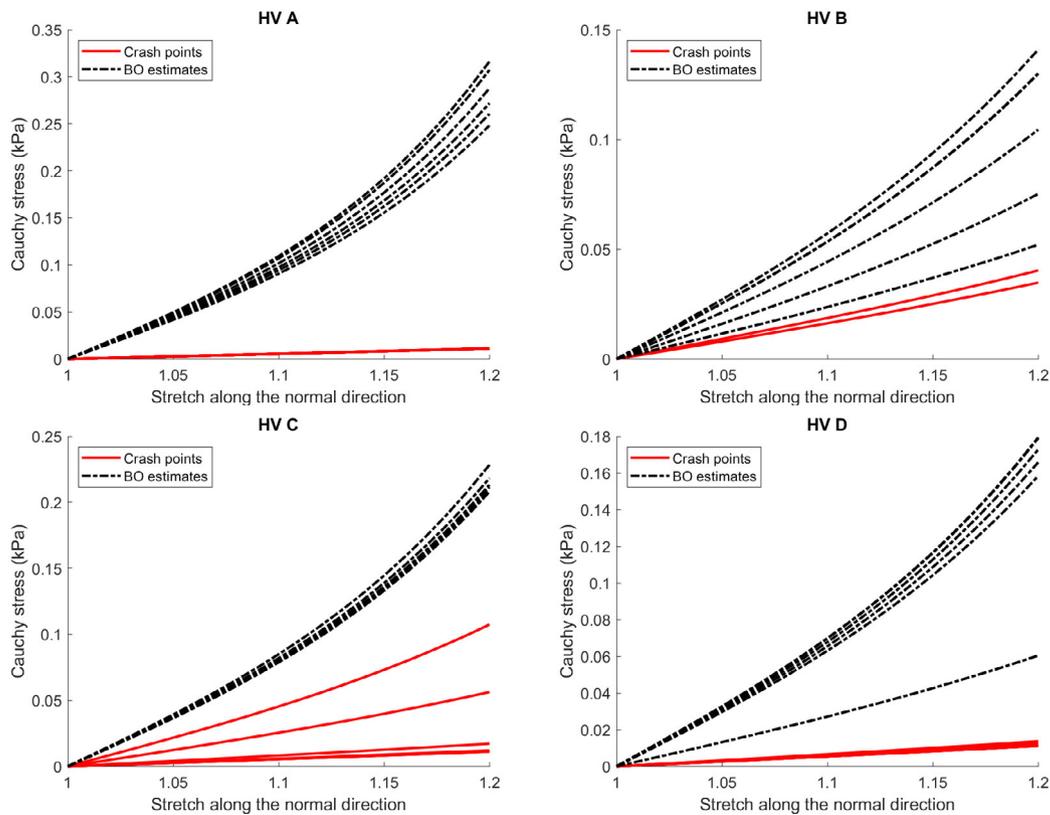


FIGURE F7 Uni-axial stretch-stress responses along the sheet-normal direction for crash points (solid red) and the BO estimates (dashed black)

consequence that BO is a global optimisation technique, in which exploitation is balanced with *exploration*, while HGO performs mostly local, gradient-based steps and hence focuses on exploitation. From this point of view, despite their name, crashes are not necessarily something negative, provided an algorithm has a feature that allows for accounting for them (like the classifier in our proposed BO framework, see Section 3.2.2), as they allow us to learn about the feasible parameter set. We emphasise again that the advantage of our global optimisation scheme is a wider exploration of the configuration space and that our model thus becomes aware of these critical regimes. A local greedy optimisation scheme, on the other hand, is by construction “blinkered”, that is so focused on exploitation that it never explores these critical regimes.

Figure F7 presents uni-axial stretch-stress responses along the sheet-normal direction between the crash points (red) and the BO estimates (black). We can see that the material stiffness along the sheet-normal direction is dramatically lower in crashed simulations than in successful simulations corresponding to BO estimates, in other words, very soft along the sheet-normal directions for those crashed simulations. This would explain why we observed crashes for certain parameter sets.

APPENDIX G: COMPARISON WITH A STATE-OF-THE-ART NUMERICAL OPTIMISER

We appreciate that the HGO algorithm runs with respect to different parameters than the proposed BO framework and that the number of parameters used by both algorithms is different (in particular, in HGO the number of parameters with respect to which the optimisation is carried out may vary between the algorithm steps). This may raise concerns about the fairness of the BO–HGO comparison. Therefore, in this appendix, we present a comparison of BO against an established iterative optimisation scheme in the same 4-dimensional parameter space. Specifically, we consider the SQP algorithm (sequential quadratic programming), which is a state-of-the-art gradient-based numerical optimiser, implemented in the `fmincon` function from MATLAB's Optimisation Toolbox. Since the performance of numerical optimisation depends on the initial point, we have run SQP from 4 different initial points: $\theta_0^{(1)} = [1, 1, 1, 1]^T$ and three

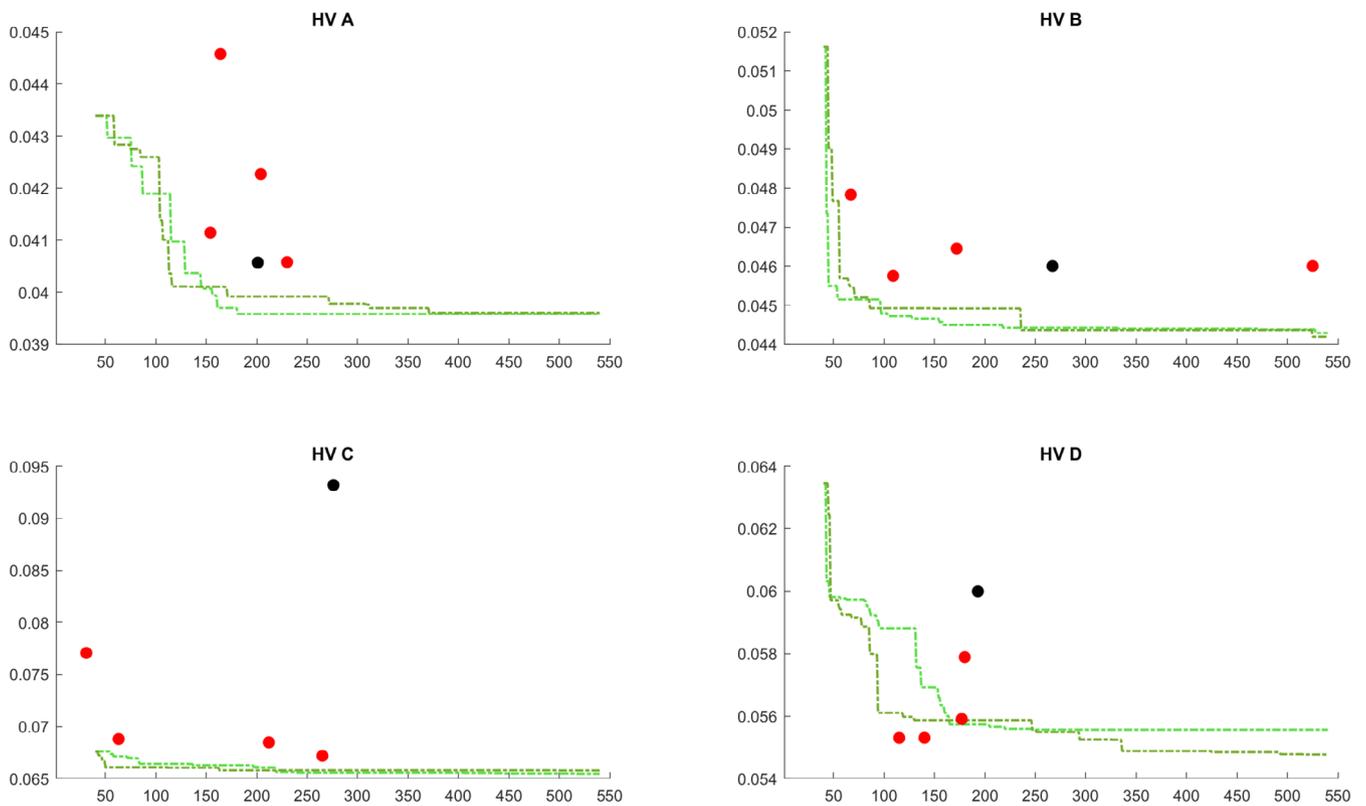


FIGURE G8 Basic study: comparison of the BO convergence (dashed lines, light green: target surrogates; dark green: partial error surrogates) with the final optima from the SQP (four red dots, one for each initial point) and HGO (black dot) algorithms

points closest to $\theta_0^{(1)}$ in the Latin hypercube initial design used to initialise BO. Figure G8 reveals that BO converges to lower values of the objective function for all the HVs. Moreover, BO typically requires fewer forward simulations achieve the same level of accuracy as the numerical optimizer.