Annals of Biomedical Engineering
https://doi.org/10.1007/s10439-026-04269-5

BMES BIOMEDICAL ENGINEERING SOCIETY

ORIGINAL ARTICLE

Automated Tuning of Cardiovascular Boundary Conditions via Differentiable Surrogate Modeling

Shridhar Thakar[1] · **Mehran Mirramezani**[1]

Received: 16 February 2026 / Accepted: 23 June 2026

Abstract

We present an end-to-end *differentiable* framework that facilitates three core capabilities in cardiovascular simulations: hemodynamic surrogate modeling, automated model calibration, and stochastic parameter tuning. Central to this approach is the introduction of a hybrid mechanistic and data-driven reduced order model (ROM) that represents each vascular domain through a nonlinear parametrization of lumped parameter networks. By exploiting the native differentiability of the pipeline, we calibrate the ROM parameters against a single high-fidelity 3D CFD simulation. The resulting optimized ROM serves as an efficient surrogate for both gradient-based deterministic and gradient-informed stochastic boundary condition calibration. With its computational efficiency and high fidelity, the framework directly addresses critical bottlenecks that currently limit the clinical adoption of cardiovascular simulations.

Keywords Cardiovascular simulation · Surrogate modeling · Model calibration · Differentiable computing

Introduction

Image-based 3D cardiovascular simulations [1] have become a crucial component of fundamental clinical and basic science research in surgical planning [2, 3], medical device design [4, 5], disease mechanisms [6, 7], and in the FDA-approved noninvasive clinical diagnosis [8]. These simulations enable detailed patient-specific analysis of hemodynamics [9] through both anatomical and physiological personalized model-twinning. The anatomical twinning stage uses medical imaging data to construct a patient-specific geometric model that defines the 3D computational fluid dynamics (CFD) domain, with the workflow accelerated and automated utilizing state-of-the-art machine learning techniques [10, 11]. The physiological twinning stage mainly relies on calibrating 0D lumped parameter networks (LPNs), which provide boundary conditions to the 3D domain [12], to match simulated hemodynamics with patient-specific clinical measurements. Tuning these parameters constitutes an inverse problem that is labor-intensive, operator-dependent, and complicated by a high-dimensional parameter space, making manual calibration extremely challenging.

Recently, researchers have proposed frameworks to automate boundary condition calibrations, broadly categorized into deterministic and stochastic approaches. Deterministic approaches primarily rely on gradient-based optimization, where gradients are typically obtained through adjoint methods [13]. Although adjoint methods offer a computationally efficient means of evaluating gradients, they require manually deriving sensitivities with respect to each input parameter and carefully addressing nonlinearities and constraints, rendering their implementation non-trivial for many applications [14–16]. Derivative-free optimization methods [17] provide an alternative to gradient-based techniques; however, they generally exhibit slow convergence, poor scalability in high-dimensional parameter spaces, and substantial sensitivity to noise [18]. Whereas deterministic approaches generally yield only single-valued estimates of the calibrated parameters, stochastic approaches enable probabilistic model calibration through quantification of uncertainties in clinical observations and modeling assumptions. Such

Associate Editor Ender A. Finol oversaw the review of this article.

✉ Mehran Mirramezani
mmirram@ncsu.edu

Shridhar Thakar
ssthakar@ncsu.edu

[1] Mechanical and Aerospace Engineering Department, NCSU, Raleigh, NC 27606, USA

Published online: 21 July 2026

stochastic formulations are commonly implemented using Bayesian techniques [19], which cast calibration as a statistical inference problem and employ variants of Monte Carlo (MC) sampling [20] to estimate posterior distributions over parameters and predictions, often requiring a large number of simulations for convergence [21].

Nevertheless, the effectiveness of current deterministic and stochastic approaches depends critically on accurate surrogate modeling, whose development remains challenging and not readily accessible in many complex cardiovascular applications [22, 23]. Reduced order models (ROMs) of blood flow can be broadly categorized into physics-based and data-driven approaches. Physics-based models [22–24] are typically derived by reducing the dimensionality of the 3D governing equations under simplifying assumptions, which can limit their ability to represent the full complexity and nonlinearity of cardiovascular flow dynamics. In contrast, data-driven and physics-informed machine learning approaches [25, 26] offer enhanced expressivity and accuracy of the underlying physics; however, they often face challenges related to scalability, numerical stability in handling complex boundary conditions, and reliance on large amounts of computationally expensive offline simulations for training projection-based or learned surrogates [25, 27].

To address these challenges, we introduce a framework that integrates three fundamental capabilities in cardiovascular simulations within a single end-to-end *differentiable* pipeline: (i) accurate surrogate modeling, (ii) automated model calibration, and (iii) efficient stochastic parameter tuning. This is achieved by proposing a mechanistic and data-driven ROM of blood flow, representing each vessel with LP elements parametrized by nonlinear resistances and inductances. We exploit the native differentiability of the pipeline to calibrate resistances and inductances of the ROM through optimization against a single high-fidelity 3D CFD simulation. The optimized ROM then serves as the surrogate model for both deterministic and stochastic model calibration. We deterministically tune boundary conditions to match a set of synthetic clinical observations using direct gradient-based optimization techniques, thereby eliminating the customized algorithmic and implementation burdens associated with adjoint methods. The differentiability of our pipeline also makes it well suited for Hamiltonian MC [28], a gradient-informed sampling technique capable of exploring posterior distributions with substantially higher efficiency than MCMC [29] or Sequential MC [30] methods, accelerating stochastic parameter calibration. The framework is implemented entirely in JAX [31], a scientific computing environment with native state-of-the-art automatic differentiation [32].

Differentiable programming represents a transformative shift in computational sciences, forging a powerful bridge between traditional scientific computing and modern machine learning through leveraging automatic differentiation. In contrast to traditional approaches, where simulations proceed only forward and gradient calculation requires cumbersome manual extraction of adjoint [13], differentiable programming allows derivatives to propagate naturally through the entire computational workflow. This capability enables seamless integration with data sciences, while facilitating efficient gradient-based optimization, inverse modeling, and data-driven discovery in scientific and engineering applications [33–36].

The paper is structured as follows. We first introduce the differentiable framework for surrogate modeling and automated boundary condition tuning in Sect. "Methods." Next, we validate the performance of the proposed framework across different cardiovascular geometries in Section Results. Finally, the paper concludes with a discussion of limitations and future outlook in Sect. "Discussion."

Methods

Estimating boundary conditions from data constitutes an inverse problem, wherein the underlying parameters are inferred indirectly from available observational data rather than measured directly. We propose two approaches, both leveraging differentiability, to solve this inverse problem: a deterministic method that provides single-valued estimates and a stochastic method that yields probabilistic distributions. Both approaches rely on an accurate ROM of blood flow that serves as a fast and accurate surrogate for full 3D CFD simulations. We begin by introducing a parametric ROM and subsequently describe the deterministic and stochastic techniques used to calibrate the boundary conditions. An overview of our end-to-end differentiable framework is demonstrated in Fig. 1.

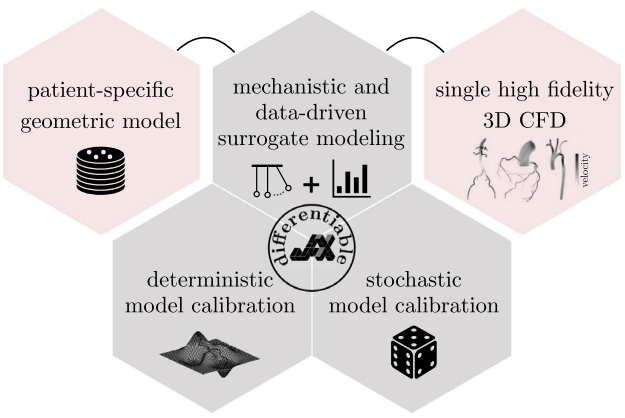

Fig. 1 An overview of our end-to-end differentiable framework for parametric reduced order modeling of blood flow, including deterministic and stochastic boundary condition tuning

Parametric ROM of Blood Flow

Here, we propose a hybrid physics-based and data-driven surrogate simulator of cardiovascular hemodynamics. First, we construct an LPN of the vascular domain by representing each vessel as a resistance-inductance ($\mathcal{R}$-$\mathcal{L}$) circuit element, capturing hydraulic resistance and blood flow inertia under the assumption of rigid vessel walls. Building on classical fluid mechanics, we derive a reduced order relationship between cross-sectionally averaged pressure P and volumetric flow rate Q that satisfies conservation of mass at vessel junctions and balance of momentum along any vascular segment, given by

$$\sum Q_{\text{in}} = \sum Q_{\text{out}} \quad , \tag{1}$$

$$\mathcal{L}\frac{\partial Q}{\partial t} + \mathcal{R}Q + \Delta P = 0. \tag{2}$$

To improve physical fidelity, we adopt a hybrid strategy that combines mechanistic insights with data-driven inference for estimating hydraulic resistance $\mathcal{R}$ and inductance $\mathcal{L}$. In particular, we previously developed a distributed lumped parameter (DLP) modeling approach [23, 37] in which $\mathcal{R}$ is decomposed into distinct physically and physiologically motivated components commonly observed in cardiovascular systems: viscous losses due to blood viscosity ($\mathcal{R}_v$), geometric effects from curvature and tortuosity ($\mathcal{R}_c$) that contribute to increased viscous dissipation and are captured through the augmented term $\mathcal{R}_{vc}$, sudden expansions or pathological changes such as stenoses and aneurysms ($\mathcal{R}_s$), and energy dissipation at vascular junctions ($\mathcal{R}_j$). Although DLP substantially outperforms conventional 1D Navier–Stokes [22] and classical LP ROMs [6, 24], it still relies on semi-empirical pressure loss formulas whose applicability is limited by poor generalizability across diverse anatomical geometries and hemodynamic regimes.

To overcome these limitations, we introduce a parametric ROM ($\boldsymbol{\alpha}$-ROM), in which each vessel is assigned a generalized resistance expressed as a nonlinear function of flow. This $\boldsymbol{\alpha}$-ROM formulation retains mechanistic structure by incorporating nonlinearities inspired by known sources of energy dissipation while enabling greater flexibility and data-alignment than existing semi-empirical models. Namely, the generalized resistance as a nonlinear function of flow rate for each vessel j is defined as

$$\mathcal{R}\left(\boldsymbol{\alpha}_j^{\text{res}}, Q_j\right) = \mathcal{R}_{\text{P}}\left(1 + a_0^j\right) + \sum_{i=1}^{d} a_i^j Q(t)_j^i, \tag{3}$$

where $\boldsymbol{\alpha}_j^{\text{res}} \in \mathbb{R}^{d+1}$ is the vector containing the polynomial coefficients $(a_0, a_1, \ldots, a_d)$ for blood vessel j and d is the degree of the polynomial nonlinearity, and $\mathcal{R}_{\text{P}}$ is the Poiseuille resistance

$$\mathcal{R}_{\text{P}} = \frac{8\mu}{\pi}\int_0^L \frac{1}{R(x)^4}\,\mathrm{d}x, \tag{4}$$

which accounts for the baseline viscous losses along the vessel. In this study, we set $d = 2$, which models nonlinear hydraulic resistance through a quadratic dependence on the flow rate, as commonly observed in physics-based flow models [23]. Similarly, the inductance of each vessel is parameterized as

$$\mathcal{L} = \mathcal{L}_{\text{g}}\left(1 + \alpha_j^{\text{ind}}\right), \tag{5}$$

where α_j^{ind} represents the modifier that parametrizes the baseline inductance $\mathcal{L}_{\text{g}}$, defined from the geometric and material properties of the vessel as

$$\mathcal{L}_{\text{g}} = \frac{\rho}{\pi}\int_0^L \frac{1}{R(x)^2}\,\mathrm{d}x. \tag{6}$$

As a result, the parameter set characterizing the nonlinear $\mathcal{R}$-$\mathcal{L}$ network for each vessel j is given by

$$\boldsymbol{\alpha}_j = \left[\boldsymbol{\alpha}_j^{\text{res}},\ \alpha_j^{\text{ind}}\right]. \tag{7}$$

The constructed LPN of $\mathcal{R}$-$\mathcal{L}$ is coupled with physiologically consistent boundary conditions to represent both the upstream and downstream vasculature. Specifically, at each boundary, the pressure-flow relationship is enforced through a boundary condition operator

$$\mathcal{F}\left(Q_j, P_j, \boldsymbol{\beta}_j\right) = 0, \tag{8}$$

whose form depends on the selected boundary model and its associated parameter set $\boldsymbol{\beta}_j$. In this work, we incorporate $\mathcal{RCR}$ Windkessel model and coronary-specific LPNs at the outlets, along with a time-varying inflow waveform prescribed at the inlet of the model. However, the proposed pipeline is readily adaptable and can incorporate alternative or additional boundary condition formulations commonly used in cardiovascular simulations. For example, when the inflow waveform is not directly available, as is often the case in clinical data, a template inflow waveform can be introduced and parameterized through a multiplicative correction represented by a truncated Fourier series. Specifically, the prescribed inflow can be written as

$$Q_{\text{in}}(t;\boldsymbol{\beta}^{\text{in}}) = \tilde{Q}(t)\exp\left[a_0 + \sum_{k=1}^{K}\left(a_k\cos\left(\frac{2\pi kt}{T}\right) + b_k\sin\left(\frac{2\pi kt}{T}\right)\right)\right], \tag{9}$$

where $\tilde{Q}(t)$ is a template waveform, T is the cardiac cycle length, K is a hyperparameter that defines the number of

Fourier modes, and $\boldsymbol{\beta}^{\text{in}} = [a_1, b_1, \ldots, a_K, b_K]$ denotes the Fourier coefficients defining the waveform correction. With this construction, the inflow is treated as an additional parameterized boundary condition, and the full boundary parameter set is written as

$$\boldsymbol{\beta}_j = \left[\boldsymbol{\beta}_j^{\text{out}}, \boldsymbol{\beta}_j^{\text{in}}\right], \tag{10}$$

so that the unknown inflow shape can be inferred jointly with the outlet LP boundary condition parameters $\boldsymbol{\beta}_j^{\text{out}}$.

In summary, the $\boldsymbol{\alpha}$-ROM model is defined as

$$\mathcal{M}_{\boldsymbol{\alpha}\text{-ROM}} : (\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{u}, t) \mapsto \mathbf{y}, \tag{11}$$

where $\boldsymbol{\alpha}$ denotes the resistance and inductance parameters of the vascular domain, $\boldsymbol{\beta}$ represents the boundary condition parameters, $\mathbf{u}$ specifies the model inputs (e.g., inlet flow or pressure waveforms), $\mathbf{y}$ defines the model outputs, i.e., pressure at all junction nodes and boundaries and flow rate through all vessels, and t is time. We will evaluate the performance of the proposed $\boldsymbol{\alpha}$-ROM against both 3D CFD simulations and our previous DLP approach, which is defined as

$$\mathcal{M}_{\text{DLP}} : (\boldsymbol{\beta}, \mathbf{u}, t) \mapsto \mathbf{y}. \tag{12}$$

While the governing equations and the boundary condition implementation are identical for both ROMs, the definition of the hydraulic resistance $\mathcal{R}$ and inductance $\mathcal{L}$ differs between them. The $\boldsymbol{\alpha}$-ROM estimates $\mathcal{R}$-$\mathcal{L}$ through optimizing $\boldsymbol{\alpha}$ for each model by leveraging data extracted from a single 3D CFD simulation, whereas the DLP estimates the inductance using $\mathcal{L}_{\text{g}}$ and computes resistance $\mathcal{R}$ using analytical expressions for various sources of energy dissipation as

$$\mathcal{R} = \underbrace{\frac{8\mu}{\pi}\int_0^L \gamma\,\frac{1}{R^4}\,\mathrm{d}x}_{\mathcal{R}_{vc}} + \underbrace{\sum_{i=1}^{n}\frac{\rho K_t}{2A_{0,i}^2}\left(\frac{A_{0,i}}{A_{s,i}}-1\right)^2|Q|}_{\mathcal{R}_s} + \underbrace{\frac{1}{2Q}\,\rho\,\frac{Q_{\text{dat}}^2}{A_{\text{dat}}^2}\left(1+\lambda_j^2\psi_j^2-2\lambda_j\psi_j\cos\phi_j\right)}_{\mathcal{R}_j}, \tag{13}$$

with fixed semi-empirical parameters that are prescribed a priori and not calibrated against corresponding 3D CFD simulations. For an in-depth derivation of the individual resistances and their corresponding parameters, the reader is referred to [23]. The governing equations in both ROMs are discretized in time using the implicit second-order BDF2 scheme [38], and all nonlinear terms are linearized using solutions from previous time steps, resulting in a linear system that is solved at each step to obtain $\mathbf{y}$.

Calibrating $\boldsymbol{\alpha}$-ROM with Differentiable Simulator

We develop an end-to-end differentiable framework to simulate blood flow and to efficiently solve the inverse problem for model calibration and parameter tuning. Our implementation is built entirely in JAX [31], a scientific computing framework in Python with automatic differentiation (AD). AD enables the computation of exact derivatives of functions expressed as compositions of primitive operations with known derivatives, applying the chain rule to machine precision. JAX implements AD by tracing functions and storing the flow of data in a static and typed computational graph using an intermediate representation called `jaxpr` [31]. Specifically, JAX computes the Jacobian $\mathcal{J}_{\mathbf{x}}[\mathbf{f}] = \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \in \mathbb{R}^{m\times n}$ for a vector-valued function $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^m$ with input $\mathbf{x} \in \mathbb{R}^n$ using two fundamental operations: Jacobian-vector product (JVP) and vector-Jacobian product (VJP). The JVP is used for forward-mode differentiation and computes $\text{JVP}(\mathbf{f}, \mathbf{x}, \mathbf{v}) = \mathcal{J}_{\mathbf{x}}[\mathbf{f}]\mathbf{v} \in \mathbb{R}^m$, where $\mathbf{v} \in \mathbb{R}^n$ is a *tangent* vector. The VJP on the other hand is used for reverse-mode differentiation, whereas for an *adjoint* vector $\mathbf{w} \in \mathbb{R}^m$ the VJP calculates $\text{VJP}(\mathbf{f}, \mathbf{x}, \mathbf{w}) = \mathbf{w}^{\text{T}}\mathcal{J}_{\mathbf{x}}[\mathbf{f}] \in \mathbb{R}^n$. Both JVP and VJP exploit the linearity of differentiation within computational graph structures to calculate the matrix–vector or vector–matrix product implicitly rather than calculating and realizing the full Jacobian matrix explicitly, which becomes prohibitively expensive for high-dimensional functions.

We will leverage the differentiability of our simulator to calibrate the $\boldsymbol{\alpha}$-ROM for each patient-specific model. We obtain the optimal parameters $\boldsymbol{\alpha}^*$ by solving a nonlinear least-squares inverse problem defined by the following objective function:

$$\begin{aligned}\boldsymbol{\alpha}^* &= \text{argmin}_{\boldsymbol{\alpha}} \left\|\mathbf{r}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{u}, \mathbf{y}_{3\text{D}})\right\|_2^2 \\ &\triangleq \text{argmin}_{\boldsymbol{\alpha}} \left\|\mathcal{M}_{\boldsymbol{\alpha}\text{-ROM}}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{u}, t) - \mathbf{y}_{3\text{D}}\right\|_2^2.\end{aligned} \tag{14}$$

This objective minimizes the residual $\mathbf{r}$ between the pressure and flow predicted by the $\boldsymbol{\alpha}$-ROM and those computed from 3D CFD simulations ($\mathbf{y}_{3\text{D}}$) described in Appendix 1, within the same geometric domain and under identical boundary conditions. Access to exact derivatives from the differentiable solver, along with the relatively small parameter dimension $\mathcal{O}(10^2)$, allows us to efficiently employ the curvature-aware Levenberg–Marquardt (LM) algorithm to solve the inverse problem and iteratively obtain the optimal model ($\boldsymbol{\alpha}^*$-ROM) via

$$\boldsymbol{\alpha}^{k+1} = \boldsymbol{\alpha}^{k} + \Delta\boldsymbol{\alpha}. \tag{15}$$

This update is computed using the LM step obtained by solving the linear system

$$\left(\mathbf{J}^{\mathrm{T}}\mathbf{J} + \lambda\mathbf{I}\right)\Delta\boldsymbol{\alpha} = -\mathbf{J}^{\mathrm{T}}\mathbf{r}, \tag{16}$$

where $\mathbf{J}$ is the Jacobian of the residual vector $\mathbf{r}$ in Eq. 14, given by

$$\mathbf{J} = \mathcal{J}_{\boldsymbol{\alpha}}[\mathbf{r}] = \frac{\partial \mathbf{r}\left(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{u}, \mathbf{y}_{\mathrm{3D}}\right)}{\partial \boldsymbol{\alpha}}. \tag{17}$$

The LM iterations proceed until the residual stagnates, defined as a change of less than 1.0% between successive iterations. Here, λ is the damping parameter that allows the LM algorithm to blend between a gradient descent algorithm when λ is large, and a Gauss-Newton algorithm when λ is small. The choice and scale of λ are often specific to the optimization problem and are determined through trial and error. We initialize the LM iterations with $\boldsymbol{\alpha} = \mathbf{0}$, which by construction (Eqs. 3 and 5) recovers the Poiseuille resistance $\mathcal{R}_{\mathrm{P}}$ and geometric inductance $\mathcal{L}_{\mathrm{g}}$. This reduces the $\boldsymbol{\alpha}$-ROM to a linear LPN whose elements are determined directly from the parameters of the 3D geometric domain, providing a physics-informed initialization for the subsequent parameter estimation process. Once the $\boldsymbol{\alpha}^*$ is identified, the resulting $\boldsymbol{\alpha}^*$-ROM is subsequently used to tune the boundary condition parameters $\boldsymbol{\beta}$ through both deterministic and stochastic optimization approaches to match a set of target observations.

Deterministic Parameter Tuning

In this section, we describe our deterministic approach to calibrate boundary condition parameters $\boldsymbol{\beta}$ to match a set of target hemodynamic measurements within the vascular domain. Concretely, we solve an inverse problem similar to $\boldsymbol{\alpha}$-ROM calibration in Eq. 14 with the addition of terms that penalizes quantities of interest like pressure and flow rate from deviating from prescribed physiological ranges

$$\boldsymbol{\beta}^* = \mathrm{argmin}_{\beta}\, \mathcal{L}(\boldsymbol{\beta}) \triangleq \mathrm{argmin}_{\beta} \left\| \mathcal{M}_{\boldsymbol{\alpha}^*\text{-ROM}}(\boldsymbol{\alpha}^*, \boldsymbol{\beta}, \mathbf{u}, t) - \mathbf{y}_{\mathrm{obs}} \right\|_2^2 + \sum_{k=1}^{N_p} \lambda_k \mathcal{P}_k(\boldsymbol{\beta}), \tag{18}$$

where $\mathbf{y}_{\mathrm{obs}}$ represents the target hemodynamic values, generally from clinical data observation. $\mathcal{P}_k(\boldsymbol{\beta})$ represent penalty terms that constrain the model to remain within physiological limits when needed, N_p are the number of penalty terms, and λ_k is a hyperparameter that represents the weight of each penalty terms in the objective function. For the purpose of this study, we first validate our framework using synthetic observations generated from 3D CFD simulations across a diverse range of geometries. We then apply the framework to patient-specific clinical data to demonstrate its usage to infer patient-specific parameters from noisy measurements representative of clinical practice. Specifically, we construct the observed data vector $\mathbf{y}_{\mathrm{obs}}$ from either a subset of 3D CFD results or clinical measurements, such as mean pressure, systolic pressure, diastolic pressure, mean flow, and other quantities explicitly detailed in the Results section for each vascular model. With the inverse problem defined, we adopt the same procedure in Eqs. 15–17 used for estimating $\boldsymbol{\alpha}$, directly leveraging the computed gradients from our differentiable pipeline within the LM algorithm to obtain the optimized boundary condition parameters $\boldsymbol{\beta}^*$.

Stochastic Parameter Tuning

While deterministic parameter tuning yields single-valued estimate for boundary conditions, the stochastic process outlined below provides a posterior distribution $P\left(\boldsymbol{\beta} \mid \mathbf{y}_{\mathrm{obs}}\right)$ for $\boldsymbol{\beta}$ while quantifying uncertainties in the observed data $\mathbf{y}_{\mathrm{obs}}$. Using Bayes' rule, we define the posterior for $\boldsymbol{\beta}$ as

$$P\left(\boldsymbol{\beta} \mid \mathbf{y}_{\mathrm{obs}}\right) = \frac{P\left(\mathbf{y}_{\mathrm{obs}} \mid \boldsymbol{\beta}\right) P(\boldsymbol{\beta})}{P\left(\mathbf{y}_{\mathrm{obs}}\right)}. \tag{19}$$

The term $P\left(\mathbf{y}_{\mathrm{obs}} \mid \boldsymbol{\beta}\right)$ representing the probability of observing the data $\mathbf{y}_{\mathrm{obs}}$ given a particular choice of $\boldsymbol{\beta}$. The term $P(\boldsymbol{\beta})$ is the prior over $\boldsymbol{\beta}$ and encodes the probabilistic distribution of boundary condition parameters before incorporating any observational data. $P\left(\mathbf{y}_{\mathrm{obs}}\right)$ is the marginal likelihood, which acts as a normalizing constant and is defined as

$$P(\mathbf{y}_{\mathrm{obs}}) = \int P(\mathbf{y}_{\mathrm{obs}} \mid \boldsymbol{\beta}) P(\boldsymbol{\beta})\, d\boldsymbol{\beta}, \tag{20}$$

the evaluation of which becomes computationally intractable in high-dimensional parameter space and typically lacks a closed form expression for general distributions. Therefore, $P\left(\boldsymbol{\beta} \mid \mathbf{y}_{\mathrm{obs}}\right)$ is often approximated by $\pi(\boldsymbol{\beta})$, defined as

$$P\left(\boldsymbol{\beta} \mid \mathbf{y}_{\mathrm{obs}}\right) \propto \pi(\boldsymbol{\beta}) \triangleq P\left(\mathbf{y}_{\mathrm{obs}} \mid \boldsymbol{\beta}\right) P(\boldsymbol{\beta}). \tag{21}$$

In this work, we assume the likelihood follows a Gaussian distribution

$$P(\mathbf{y}_{\text{obs}} \mid \boldsymbol{\beta}) = \frac{1}{\sqrt{(2\pi)^n \det \boldsymbol{\Sigma}}} \exp\left[-\frac{1}{2}\left(\mathbf{y}_{\text{obs}} - \mathcal{M}_{\boldsymbol{\alpha}^*\text{-ROM}}(\boldsymbol{\beta})\right)^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} \left(\mathbf{y}_{\text{obs}} - \mathcal{M}_{\boldsymbol{\alpha}^*\text{-ROM}}(\boldsymbol{\beta})\right)\right], \tag{22}$$

where $\boldsymbol{\Sigma}$ is the covariance matrix for the observed data $\mathbf{y}_{\text{obs}}$. For this study, we assume that the observations are uncorrelated; hence, $\boldsymbol{\Sigma}$ is diagonal. The parameter $\boldsymbol{\beta}$ is sampled from a uniform distribution, representing non-informative prior. Since boundary condition parameters are strictly positive, the resulting constrained search space renders MC-based sampling inefficient. In addition, the $\boldsymbol{\beta}$ encompasses parameters that span various orders of magnitude, ranging from compliances of $\mathcal{O}(10^{-5})$ to microcirculatory resistances of $\mathcal{O}(10^{5})$. This extreme variation in parameter scales leads to a poorly conditioned optimization problem, leading to convergence inefficiency and instability. To mitigate these limitations and facilitate sampling in an unconstrained space, we will employ a log-space transformation T

$$T : \mathbb{R}_+^n \to \mathbb{R}^n; \quad \boldsymbol{\theta} = T(\boldsymbol{\beta}) := \log \boldsymbol{\beta}, \tag{23}$$

which leads to the posterior $\pi(\boldsymbol{\theta}) \propto P(\mathbf{y}_{\text{obs}} \mid \boldsymbol{\theta}) P(\boldsymbol{\theta})$ defined over $\boldsymbol{\theta}$ and expressed explicitly using Eq. 21

$$\pi(\boldsymbol{\theta}) \triangleq P\left(\mathbf{y}_{\text{obs}} \mid T^{-1}(\boldsymbol{\theta})\right) P\left(T^{-1}(\boldsymbol{\theta})\right) \left|\det\left(J_{T^{-1}}(\boldsymbol{\theta})\right)\right|, \tag{24}$$

where $J_{T^{-1}}$ is the Jacobian of the inverse transformation T^{-1}. Finally, we will approximate the posterior distribution by implementing three different MC-based methods leveraging the $\mathcal{M}_{\boldsymbol{\alpha}^*\text{-ROM}}$ surrogate simulator within the open-source sampling library BlackJAX [39], which is implemented in JAX and fully compatible with our computational pipeline. Single-valued estimates of the boundary condition parameters are then obtained as the expected value

$$\boldsymbol{\beta}^* = \frac{1}{N} \sum_{i=1}^{N} T^{-1}\left(\boldsymbol{\theta}^i\right), \; \boldsymbol{\theta}^i \sim \pi(\boldsymbol{\theta}), \tag{25}$$

where N is the number of samples.

We will summarize three MC-based methods used in this study in the following sections.

Random-Walk Metropolis-Hastings (RMH)

We will approximate the posterior distribution by utilizing the RMH sampling techniques. RMH constructs a Markov chain by iteratively proposing candidate samples $\boldsymbol{\theta}'$ from a symmetric Gaussian distribution centered at the current sampling state $\boldsymbol{\theta}^n$

$$\boldsymbol{\theta}' \sim \mathcal{N}(\boldsymbol{\theta}^n, \mathbf{C}),$$

where $\mathbf{C}$ is a user-defined proposal covariance matrix. In this study, we adopt a Gaussian random-walk proposal in which the standard deviation of each parameter is set to 10% of the corresponding component of the current state $\boldsymbol{\theta}^n$. The proposed sample $\boldsymbol{\theta}'$ is accepted with probability

$$A(\boldsymbol{\theta}', \boldsymbol{\theta}^n) = \min\left(1, \frac{\pi(\boldsymbol{\theta}')}{\pi(\boldsymbol{\theta}^n)}\right),$$

where $\pi(\boldsymbol{\theta})$ is the posterior. If the proposal is accepted, it becomes the next state in the chain; otherwise, the current state is retained. The RMH implementation is detailed in Algorithm 1.

Sequential Monte Carlo (SMC)

The SMC method [40], also known as particle filtering, estimates the posterior by evolving a set of particles from the prior toward the posterior through a sequence of intermediate distributions. Specifically, the n-th intermediate distribution $\pi_n(\boldsymbol{\theta})$ is approximated using

$$\pi_n(\boldsymbol{\theta}) \approx \sum_{i=1}^{k} w_n^i \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_i), \tag{26}$$

with normalized w_n^i weights associated with each particle location $\boldsymbol{\theta}_i$, representing the relative importance of the i-th particle. Here, δ denotes the Dirac delta function, and k is the total number of particles. This particle-based construction enables efficient exploration of the parameter space and accurate estimation of complex and multi-modal posterior distributions. Through tempering of π_n, reweighting of w_n^i, and resampling of $\boldsymbol{\theta}_i$, the particle population is progressively concentrated in regions of high posterior probability, thereby refining the posterior distribution. Algorithm 2 summarizes the steps employed in SMC for solving the Bayesian calibration problem.

Hamiltonian Monte Carlo (HMC)

The HMC [41] algorithm is a Markov chain MC method that approximates the posterior distribution by formulating the sampling problem as a Hamiltonian system through the introduction of an auxiliary momentum variable $\boldsymbol{\psi}$. Unlike the previous two methods, HMC requires gradients of the posterior, which are readily available via our differentiable framework. The gradient-informed dynamics significantly improve exploration of the parameter space, yielding higher acceptance rates and more efficient sampling compared to RMH and SMC. As a result, HMC requires

substantially fewer samples to achieve performance comparable to RMH and SMC.

We briefly outline the HMC formulation, beginning with the log-posterior

$$l(\boldsymbol{\theta}) = \log \pi(\boldsymbol{\theta}), \tag{27}$$

and its gradient $\mathbf{g}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} l(\boldsymbol{\theta})$. The potential energy of the Hamiltonian system is then defined as

$$U(\boldsymbol{\theta}) = -l(\boldsymbol{\theta}). \tag{28}$$

To define the kinetic energy, we introduce the auxiliary momentum variable $\boldsymbol{\psi}$, sampled from a multivariate normal distribution with zero mean $\boldsymbol{\psi} \sim \mathcal{N}(\mathbf{0}, \mathbf{M})$. The covariance matrix $\mathbf{M}$ serves as the mass matrix of the Hamiltonian system. We assume the calibration parameters are uncorrelated; accordingly, the matrix $\mathbf{M}$ is taken to be diagonal. The kinetic energy is therefore defined as

$$K(\boldsymbol{\psi}) = \frac{1}{2}\boldsymbol{\psi}^{\mathrm{T}}\mathbf{M}^{-1}\boldsymbol{\psi}, \tag{29}$$

resulting in the Hamiltonian $H(\boldsymbol{\theta}, \boldsymbol{\psi}) = U(\boldsymbol{\theta}) + K(\boldsymbol{\psi})$. The corresponding Hamiltonian dynamics of the system can be expressed as a coupled system of differential equations:

$$\begin{aligned} \dot{\boldsymbol{\theta}} &= \frac{\partial H(\boldsymbol{\theta}, \boldsymbol{\psi})}{\partial \boldsymbol{\psi}} = \mathbf{M}^{-1}\boldsymbol{\psi} \\ \dot{\boldsymbol{\psi}} &= -\frac{\partial H(\boldsymbol{\theta}, \boldsymbol{\psi})}{\partial \boldsymbol{\theta}} = \mathbf{g}(\boldsymbol{\theta}). \end{aligned} \tag{30}$$

These equations are integrated numerically using a symplectic leapfrog integration scheme detailed in Algorithm 4. During an initial warmup run, the mass matrix $\mathbf{M}$ and the step size Δt that are required by the leapfrog integrator are tuned using the windowed adaptation scheme introduced with the No-U-Turn Sampler [42]. The tuned parameters are then used in the sampling phase, where the number of leapfrog steps is adaptively chosen. The complete HMC algorithm is presented in Algorithm 3.

Results

To evaluate the utility of the proposed differentiable pipeline, we have applied this framework to a diverse range of patient-specific cardiovascular anatomies exhibiting complex hemodynamics, including aortic, aortofemoral, and coronary models. In each vascular model, we first employ a single 3D CFD simulation to calibrate a parametric $\boldsymbol{\alpha}$-ROM. We then use the calibrated ROM to deterministically and stochastically tune the boundary condition parameters so that the model reproduces a prescribed set of hemodynamic targets representing clinical data observations. Across all models, we investigate multiple types of target responses to evaluate the efficacy and robustness of the proposed approach.

$\boldsymbol{\alpha}$-ROM Calibration

Here, our goal is to obtain the calibrated $\boldsymbol{\alpha}^*$-ROM for each vascular model by optimizing $\boldsymbol{\alpha}$ so that the resulting pressure and flow waveforms from ROM match those computed from a single 3D CFD simulation performed on the same geometric domain and under identical boundary conditions. To ensure that the $\boldsymbol{\alpha}^*$-ROM generalizes robustly across a broad range of LP boundary conditions and inflow profiles, we evaluated its accuracy using 10 distinct 3D CFD simulations for each vascular model. In each simulation, the initial boundary condition parameters were perturbed multiplicatively by a scalar randomly and independently drawn from a uniform distribution $\mathcal{U}[0.8,\ 1.2]$.

We use the 3D CFD results as reference solutions and evaluate the accuracy of the proposed ROM against them. We compare the performance of the $\boldsymbol{\alpha}^*$-ROM with that of the DLP approach by quantifying the reduction in the maximum percentage pressure error at each outlet. The percentage error decrease at each outlet i is defined as

$$\Delta\varepsilon_P^{(i)} = \left[\frac{\max_t \left|P^{(i)}_{\boldsymbol{\alpha}^*\text{-ROM}}(t) - P^{(i)}_{\text{3D}}(t)\right|}{P^{(i)}_{\text{3D}}(t)}\right]\% - \left[\frac{\max_t \left|P^{(i)}_{\text{DLP}}(t) - P^{(i)}_{\text{3D}}(t)\right|}{P^{(i)}_{\text{3D}}(t)}\right]\%. \tag{31}$$

To provide a statistical summary across the perturbation study, we also report the mean and standard deviation of the maximum pressure error over the $N_{\text{pert}} = 10$ perturbation cases. For each perturbation case k at outlet i, the scalar pressure error is defined as

$$\varepsilon^{(i)}_{P,k} = \max_t \frac{\left|P^{(i)}_{\text{model},k}(t) - P^{(i)}_{\text{3D},k}(t)\right|}{\left|P^{(i)}_{\text{3D},k}(t)\right|} \times 100, \qquad k = 1, \ldots, N_{\text{pert}}, \tag{32}$$

where P_{model} denotes either the $\boldsymbol{\alpha}^*$-ROM or DLP prediction. The corresponding mean and sample standard deviation across perturbation cases are then computed using $\varepsilon^{(i)}_{P,k}$. It is important to note that pressure error serves as a sufficient metric for assessing model accuracy because pressure and flow rate are tightly coupled through boundary condition operators defined in Eq. 8 within the underlying hemodynamic system. Consequently, achieving close agreement between the ROM and the 3D CFD simulation in inlet and

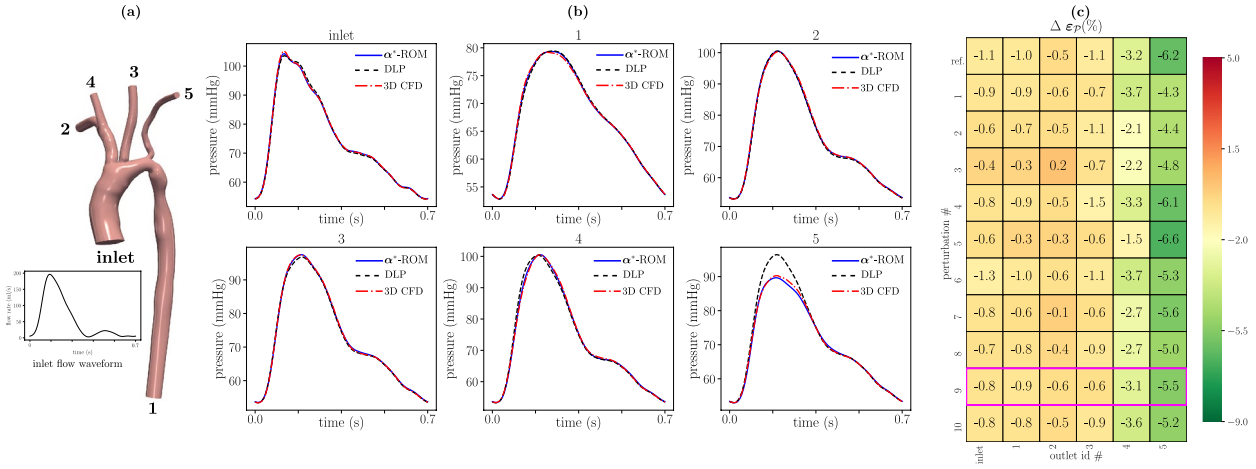

Fig. 2 **a** Aortic model with mild coarctation. **b** Pressure waveforms from 3D CFD, $\boldsymbol{\alpha}^*$-ROM, and DLP methods at the inlet and outlets using boundary conditions from perturbation 9. **c** Percentage error changes between $\boldsymbol{\alpha}^*$-ROM and DLP prediction against 3D CFD across outlets and perturbation cases, with negative values indicating $\boldsymbol{\alpha}^*$-ROM outperforming DLP

outlet pressures inherently yields a correspondingly accurate match in the associated flow rates at those locations.

Aortic model

Here, we evaluate the ability of the proposed $\boldsymbol{\alpha}$-ROM to accurately predict temporal pressure and flow waveforms in an aortic model with mild coarctation (Fig. 2a), a configuration that yields complex flow physics and poses challenges for CFD simulation. The inlet flow waveform, shown in Fig. 2a, was obtained from the Vascular Model Repository [43] and prescribed at the inlet; the corresponding aortic model identifier is provided in Table 6 of Appendix 3. The outlets were coupled to three-element $\mathcal{RCR}$ LP boundary conditions. We first calibrated the $\boldsymbol{\alpha}$-ROM to a single reference 3D CFD simulation using the LM algorithm with the damping parameter λ set to 10^{-6}, and then compared pressure waveforms at all outlets and the inlet predicted by the calibrated $\boldsymbol{\alpha}^*$-ROM and the DLP model against those from 10 additional 3D CFD simulations with perturbed boundary conditions described in Section $\boldsymbol{\alpha}$-ROM Calibration. The subclavian artery outlet (outlet id #5) is of particular interest because the DLP model was unable to reproduce the corresponding 3D pressure waveform, as shown in Fig. 2b. The $\boldsymbol{\alpha}^*$-ROM produced pressure waveforms in strong agreement with 3D CFD across all outlets, with a maximum error of 1.0% against reference 3D CFD results and 2.4% at the subclavian outlet for perturbation 7. By comparison, the DLP model exhibited a maximum error of 7.5% at the same location. The overall improvement of the proposed $\boldsymbol{\alpha}$-ROM relative to the DLP approach is summarized in the heat map shown in Fig. 2c and in Table 1, which reports the percentage error differences between the two models as defined in Eq. 31, and mean and standard deviation of maximum-in-time pressure errors across the perturbation cases, respectively. Notably, the $\boldsymbol{\alpha}$-ROM achieves substantial accuracy gains in the subclavian artery, corresponding to the last column of Fig. 2c, where average of maximum-in-time pressure error reduces from 6.69% to 1.42%, as shown in Table 1.

Aortofemoral Model

We considered an extensive aortofemoral model, shown in Fig. 3a, which includes the descending, abdominal, iliac, and femoral arteries, regions prone to severe vascular diseases such as aneurysms and stenoses. The individualized inflow waveform, shown in Fig. 3, was determined

Table 1 Mean and standard deviation of maximum-in-time pressure errors across the perturbation cases for the aortic model

	Outlet id					
Error statistics	Inlet	1	2	3	4	5
Mean_{DLP} (%)	1.56	1.46	1.18	1.42	3.98	6.69
STD_{DLP} (%)	0.19	0.18	0.22	0.28	0.45	0.63
$\text{Mean}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.80	0.75	0.79	0.55	1.13	1.42
$\text{STD}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.19	0.11	0.17	0.12	0.34	0.52

based on the Baker equation in previous study and is available from the Vascular Model Repository [43] and prescribed at the inlet; the corresponding model identifier is provided in Table 6 of Appendix 3. All outlets were coupled to three-element $\mathcal{RCR}$ LP boundary conditions. The $\boldsymbol{\alpha}^*$-ROM was achieved using the LM-based procedure with the damping parameter λ set to 10^{-3}. The calibrated model accurately captured the flow dynamics across the aortofemoral model over the full range of perturbed boundary conditions, with a maximum error of 0.6% relative to the reference 3D CFD simulation for the reference calibration, and a maximum of 2.1% at outlet 4 for perturbation 6. There were some instances in which the parametric ROM showed slightly lower accuracy than the DLP at specific outlets, for example, outlet 4 for perturbations 1, 6, and 8 (Fig. 3c). Nevertheless, when averaging over all outlets, the $\boldsymbol{\alpha}$-ROM generally outperformed the DLP

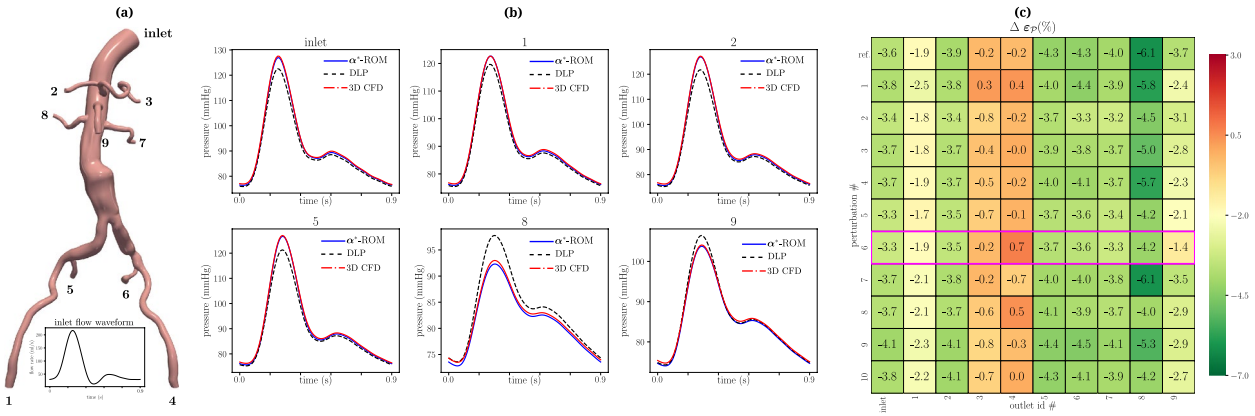

Fig. 3 **a** Aortofemoral model. **b** Pressure waveforms from 3D CFD, $\boldsymbol{\alpha}^*$-ROM, and DLP methods at the inlet and outlets using boundary conditions from perturbation 6. **c** Percentage error changes between $\boldsymbol{\alpha}^*$-ROM and DLP prediction against 3D CFD across outlets and perturbation cases, with negative values indicating $\boldsymbol{\alpha}^*$-ROM outperforming DLP

Table 2 Mean and standard deviation of maximum-in-time pressure errors across the ten perturbation cases for the aortofemoral model

	Outlet id									
Error statistics	Inlet	1	2	3	4	5	6	7	8	9
Mean_{DLP} (%)	4.26	2.70	4.53	1.41	1.16	4.80	4.69	4.46	5.60	3.36
STD_{DLP} (%)	0.19	0.18	0.21	0.32	0.23	0.22	0.20	0.22	0.56	0.58
$\text{Mean}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.62	0.66	0.81	0.95	1.17	0.83	0.76	0.79	0.71	0.76
$\text{STD}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.12	0.16	0.10	0.29	0.50	0.08	0.26	0.18	0.24	0.22

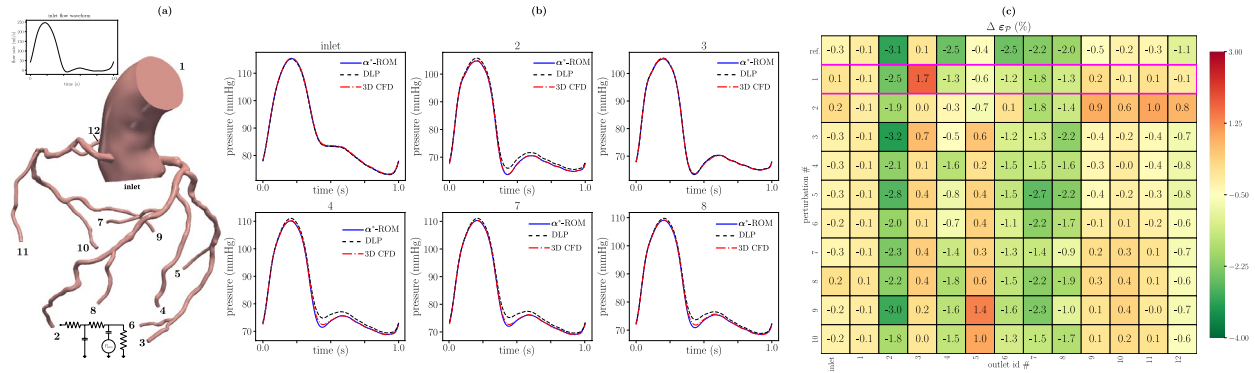

Fig. 4 **a** Coronary model. **b** Pressure waveforms from 3D CFD, $\boldsymbol{\alpha}^*$-ROM, and DLP methods at the inlet and outlets using boundary conditions from perturbation 1. **c** Percentage error changes between $\boldsymbol{\alpha}^*$-ROM and DLP prediction against 3D CFD across outlets and perturbation cases, with negative values indicate $\boldsymbol{\alpha}^*$-ROM outperforming DLP

approach significantly (Table 2). For example, at outlet 8, the average maximum-in-time error was reduced markedly from 5.60% using the DLP approach to 0.71% with the proposed $\boldsymbol{\alpha}$-ROM.

Coronary Model

Finally, we calibrated a parametric ROM to study hemodynamics in a complex coronary model featuring a severe stenotic lesion in the main left circumflex artery, characterized by a 75% reduction in cross-sectional area. Unlike systemic arteries, coronary circulation exhibits a phase shift between pressure and flow waveforms, primarily due to extravascular myocardial compression. We selected this model to evaluate the effectiveness of our framework in capturing such physiologically complex behavior, given its challenging geometry, large number of outlets, and non-trivial temporal dynamics. The ascending aorta was coupled to an $\mathcal{RCR}$ model, while all coronary outlets were coupled to a five-element coronary LPN (shown in Fig. 4a). Intramyocardial pressure was modeled by scaling representative left and right ventricular pressure waveforms to account for extravascular myocardial compression [23]. The same calibration procedure described for the aortic and aortofemoral models was employed, with the damping parameter λ set to 10^{-6}. Across the majority of outlets and perturbed cases, particularly during early diastole, the $\boldsymbol{\alpha}^*$-ROM outperformed the DLP model, as illustrated in Fig. 4b. The maximum pressure error in diastole for the DLP is 2.95%, which reduces to 1.25% using the proposed $\boldsymbol{\alpha}$-ROM method. Overall, a consistent improvement in maximum-in-time pressure errors is observed for all 10 perturbations in Table 3, along with reduced percentage errors between $\boldsymbol{\alpha}^*$-ROM and DLP predictions relative to the 3D CFD results across nearly all outlets indicated in Fig. 4c. Averaged over all systemic and coronary outlets, the mean maximum-in-time pressure error decreased from 1.77% for the DLP model to 1.17% for the $\boldsymbol{\alpha}^*$-ROM. The largest improvement occurs at outlet 2, where the maximum-in-time pressure error decreases from 3.34% for the DLP model to 0.95% for the calibrated $\boldsymbol{\alpha}^*$-ROM.

Deterministic Boundary Condition Tuning

Here, we demonstrate our framework's ability to deterministically tune boundary condition parameters to match hemodynamic observations. We use *noiseless* synthetic observations generated from 3D CFD simulations to verify our framework's robustness across aortic, aortofemoral, and coronary models, followed by a demonstration of our frameworks effectiveness to tune boundary conditions to match actual *noisy* clinical data for an aortic case. To evaluate the robustness of the pipeline, we considered multiple types of observational data $\mathbf{y}_{\text{obs}}$, including scalar hemodynamic metrics such as average pressure and average flow rates (P_{avg} and Q_{avg}), as well as full temporal waveforms ($P(t)$ and $Q(t)$) extracted from a 3D CFD simulation results with known boundary conditions. Because the nonlinear hemodynamic system does not necessarily admit a unique set of boundary conditions that reproduce a given target response, the calibrated boundary conditions are subsequently employed in a 3D CFD simulation to assess whether the resulting hemodynamics from 3D CFD match the target quantities predicted by the optimized ROM.

Aortic model. We consider a set of synthetic scalar target observations, including average, maximum, and minimum pressure and flow rates at the boundaries of the geometric domain, i.e., $\mathbf{y}_{\text{obs}} = [P^{\text{ao}}_{\text{avg}}, P^{\text{ao}}_{\text{max}}, P^{\text{ao}}_{\text{min}}, \{Q^{\text{i}}_{\text{avg}}, Q^{\text{i}}_{\text{max}}, Q^{\text{i}}_{\text{min}}$ for $\text{i} \in [1, \cdots, n_o]\}]$, where P^{ao} denotes the pressure at the aortic root, Q^{i} is the flow rate at outlet i, and n_o is the number of outlets. While average pressure and flow are of primary clinical importance, we intentionally included maximum and minimum values to assess the robustness of the proposed calibration framework. In addition, several clinical indicators depend explicitly on extrema rather than mean quantities; for example, the ankle-brachial index (ABI) in peripheral arterial disease is defined using maximum systolic pressure measurements. Figure 5 illustrates the performance of our gradient-based optimization procedure, which leverages the calibrated $\boldsymbol{\alpha}^*$-ROM to estimate the optimal boundary condition parameters $\boldsymbol{\beta}^*$ that best match the target observations. The $\boldsymbol{\beta}$ comprises the $\mathcal{RCR}$ parameters of the Windkessel model at all outlets. While the initial parameter vector $\boldsymbol{\beta}_0$ produces hemodynamic predictions that deviate substantially from the prescribed targets, the proposed direct

Table 3 Mean and standard deviation of maximum-in-time pressure errors across the ten perturbations for the coronary model

	Outlet id												
Error statistics	Inlet	1	2	3	4	5	6	7	8	9	10	11	12
Mean_{DLP} (%)	0.67	0.51	3.34	1.01	2.63	1.19	2.90	2.94	2.64	0.62	0.50	0.71	1.35
STD_{DLP} (%)	0.24	0.25	0.78	0.78	0.44	0.61	1.43	0.48	0.41	0.15	0.22	0.25	0.51
$\text{Mean}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.54	0.44	0.95	1.42	1.47	1.54	1.70	1.08	1.06	0.68	0.65	0.72	0.86
$\text{STD}_{\boldsymbol{\alpha}^*\text{-ROM}}$ (%)	0.44	0.27	0.75	1.22	0.46	0.46	1.45	0.38	0.42	0.40	0.40	0.37	0.47

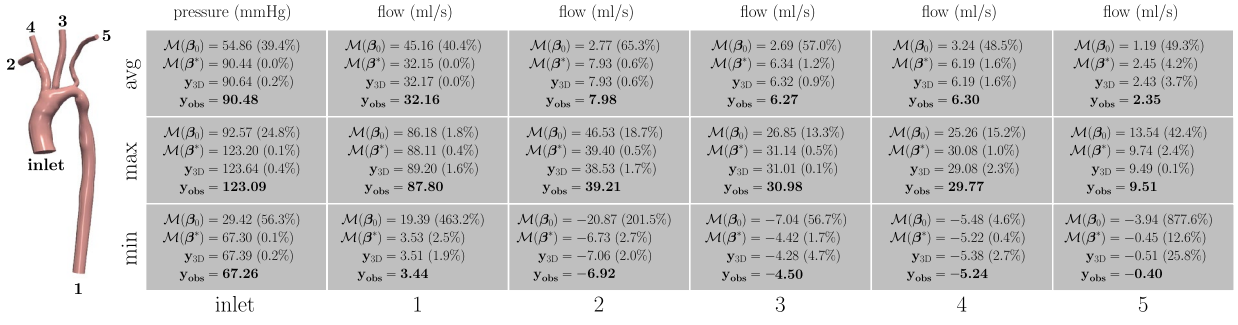

Fig. 5 Predicted pressure and flow rates obtained from $\boldsymbol{\alpha}^*$-ROM (defined as $\mathcal{M}$) using initial parameters $\boldsymbol{\beta}_0$ and calibrated boundary condition parameters $\boldsymbol{\beta}^*$, and also from 3D CFD with $\boldsymbol{\beta}^*$ boundary conditions, together with their corresponding percentage errors relative to the synthetic target observed data $\mathbf{y}_{obs}$ for the aortic model

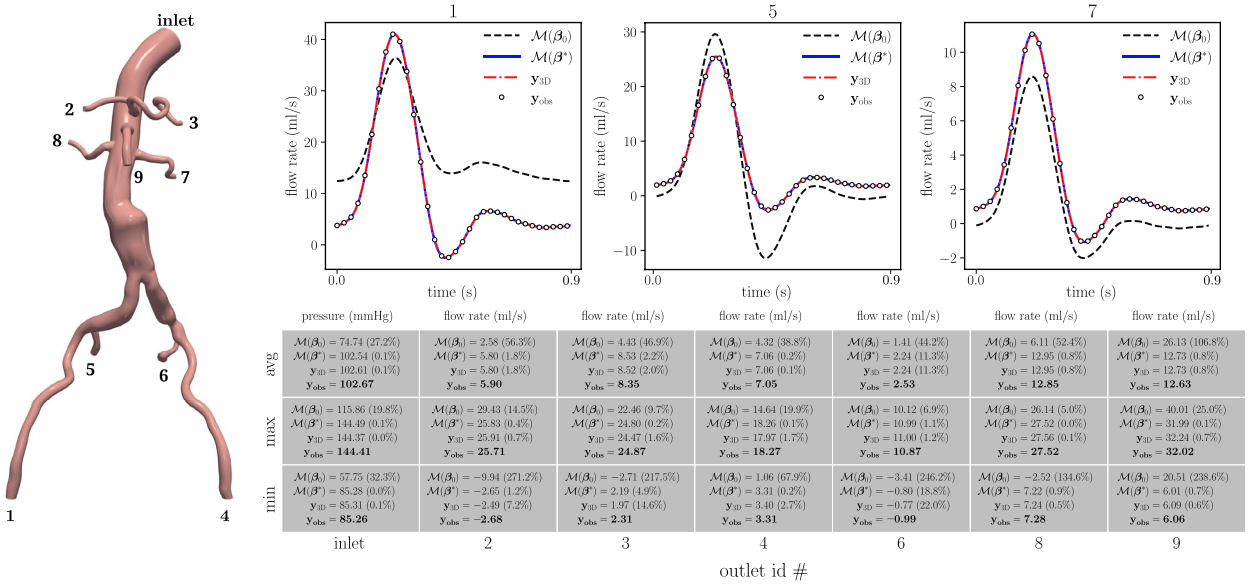

Fig. 6 Predicted flow rate waveforms obtained from $\boldsymbol{\alpha}^*$-ROM (defined as $\mathcal{M}$) using initial parameters $\boldsymbol{\beta}_0$ and calibrated boundary condition parameters $\boldsymbol{\beta}^*$, and also from 3D CFD with $\boldsymbol{\beta}^*$ boundary conditions (top row). Predicted pressure and flow rate scalars, together with their corresponding percentage errors relative to the synthetic target data $\mathbf{y}_{obs}$ for the aortofemoral model (bottom row)

optimization framework identifies an optimal parameter set $\boldsymbol{\beta}^*$ which, when deployed within the $\boldsymbol{\alpha}^*$-ROM and full 3D CFD simulations, both yield close agreement with the target hemodynamic responses across all quantities of interest. After boundary condition calibration, the $\boldsymbol{\alpha}^*$-ROM predictions were within 13% of the target scalar observations across all quantities with the highest relative error seen for outlet 5. When aiming for average flow and pressure values, the maximum errors of both $\boldsymbol{\alpha}^*$-ROM and 3D CFD simulations are less than 4.2% for all the outlets.

Aortofemoral model. We aim to match a combination of synthetically generated scalar and temporal hemodynamic metrics within this geometric model by calibrating all $\mathcal{RCR}$ parameters at the outlets. To this end, we define the observation vector as $\mathbf{y}_{obs} = [P^{in}_{avg}, P^{in}_{max}, P^{in}_{min}, \{Q^{i}_{avg}, Q^{i}_{max}, Q^{i}_{min} \text{ for } i \in [1, \cdots, n_s]\}, \{Q^{j}(t) \text{ for } j \in [1, \cdots, n_t]\}]$, where n_s is the number of outlets for which scalar flow metrics are available and n_t denotes the number of outlets with prescribed temporal flow waveforms. Figure 6 compares the prediction from $\boldsymbol{\alpha}^*$-ROM obtained using the initial boundary condition parameters $\boldsymbol{\beta}_0$ and the calibrated parameters $\boldsymbol{\beta}^*$, together with the corresponding 3D CFD results using $\boldsymbol{\beta}^*$, against the target hemodynamic responses. The proposed framework enables systematic optimization of the initial parameter vector $\boldsymbol{\beta}_0$, yielding a calibrated surrogate model

$\mathcal{M}(\boldsymbol{\beta}^*)$ that accurately reproduces both the target flow waveform dynamics and the scalar hemodynamic quantities across multiple outlets of the aortofemoral model. The calibrated predictions exhibit strong agreement with both 3D CFD simulations and the observed data. The largest relative discrepancies were observed for the minimum flow rates, reaching approximately 22% in the 3D CFD simulations and decreasing to 13% for the average flow predictions. These relatively high error values are primarily attributable to normalization by small reference values. This high values in error is primarily attributable to normalization by small reference values. Previous studies [44] reported flow rate errors normalized with respect to the maximum flow over the cardiac cycle to mitigate this issue. In contrast, the present study computes errors using the instantaneous flow rate at each time step similar to Eq. 31. If the errors were instead normalized by the maximum flow, the reported errors associated with matching the minimum flow would be comparable to those for matching the target maximum flow, with a maximum error of approximately 2%.

Coronary model. Similar to the aortofemoral model, the synthetic target hemodynamic metrics in the coronary model include a combination of scalar values and temporal waveforms of pressure and flow rate. Calibrating the coronary-specific LPN models is particularly challenging because the coronary circulation exhibits distinct temporal dynamics between the right and left coronary arteries. Specifically, flow in the right coronary arteries behaves like systemic circulation, with maximum flow occurring during systole, whereas flow in the left coronary arteries is out-of-phase, reaching its peak during diastole due to elevated left ventricular pressure during systole. To evaluate the efficacy of our method, we define $\mathbf{y}_{\text{obs}} = [P^{\text{in}}_{\text{avg}}, P^{\text{in}}_{\text{max}}, P^{\text{in}}_{\text{min}}, \{Q^{\text{i}}_{\text{rcor}}(t) \text{ for } \text{i} \in [1, \cdots, n_r]\}, \{Q^{\text{j}}_{\text{lcor}}(t) \text{ for } \text{j} \in [1, \cdots, n_l]\}]$, where n_r and n_l denote number of right and left coronary outlets. The target flow waveforms are specified as $Q^{\text{i}}_{\text{rcor}}(t) = Q^{\text{i}}_{\text{avg}} \tilde{Q}_{\text{rcor}}(t)$ and $Q^{\text{j}}_{\text{lcor}}(t) = Q^{\text{j}}_{\text{avg}} \tilde{Q}_{\text{lcor}}(t)$, where $\tilde{Q}_{\text{rcor}}(t)$ and $\tilde{Q}_{\text{lcor}}(t)$ are the normalized template waveforms for the right and left coronaries, respectively, shown in Fig. 7. To assess the robustness of our pipeline, we initially set the coronary-specific LPN parameters such that the ROM predicts in-phase flow behavior across both left and right coronary outlets (Fig. 7). After calibrating the boundary condition parameters, the ROM produces physiologically consistent coronary flow waveform dynamics with the target average flow rates, where left coronary outlets exhibiting out-of-phase flow and right coronary outlets showing in-phase flow, as illustrated in Fig. 7. The $\boldsymbol{\alpha}^*$-ROM reproduced the prescribed inlet pressure and coronary flow targets with errors below 1.0% for the scalar quantities shown in Fig. 7, while the corresponding 3D CFD simulation results using $\boldsymbol{\beta}^*$ also remained within 1.0% of the targets.

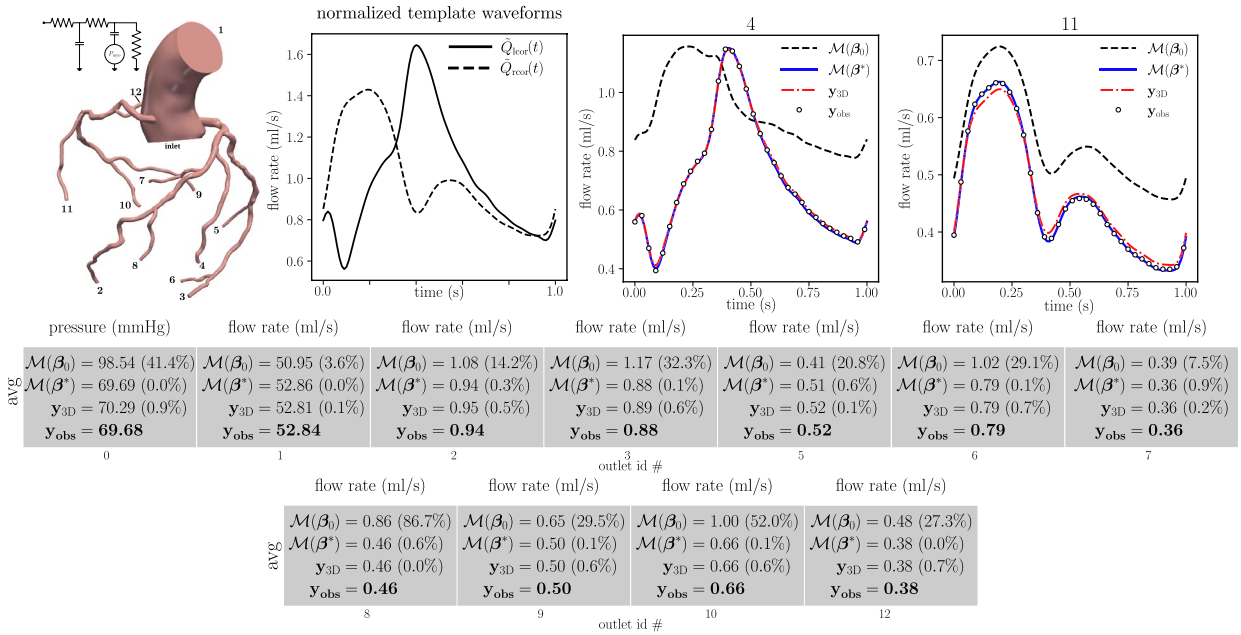

Fig. 7 Normalized template flow rate waveforms for right and left coronaries (top left). Predicted flow rate waveforms from $\boldsymbol{\alpha}^*$-ROM (defined as $\mathcal{M}$) using initial parameters $\boldsymbol{\beta}_0$ and calibrated boundary condition parameters $\boldsymbol{\beta}^*$, and also from 3D CFD with $\boldsymbol{\beta}^*$ boundary conditions (top right). Predicted pressure and flow rates, together with their corresponding percentage errors relative to synthetic the target observed data $\mathbf{y}_{\text{obs}}$ (bottom) for the coronary model

Aortic model (clinical data). Finally, we consider a patient-specific aortic model constructed from CT imaging. First, we apply the previously described calibration procedure to obtain the corresponding $\boldsymbol{\alpha}$-ROM calibrated model, which is subsequently used to estimate optimal $\mathcal{RCR}$ Windkessel and inflow boundary conditions that best match the available clinical targets. The calibration targets were obtained from clinical measurements in a 77-year-old female patient following transcatheter aortic valve implantation (TAVI). The target data include a temporal waveform of aortic root pressure measured via invasive catheterization and post-procedural cardiac output estimated from echocardiographic assessment. The pressure waveform was digitized over five cardiac cycles and reduced to a representative single cycle target by temporally averaging over the cycles. A pointwise $\pm 1\sigma$ uncertainty band was then constructed from the cycle to cycle standard deviation across the five cardiac cycles. There was no patient-specific inflow waveform available for this patient; therefore, the inlet flow was prescribed using a representative flow waveform at the root of aorta defined by a multiplicative correction of a template waveform, as mathematically described in Eq. 9. In summary, we define $\mathbf{y}_{\text{obs}} = [P^{\text{ar}}(t), CO]$, where $P^{\text{ar}}(t)$ is the representative single cycle aortic root pressure waveform, while CO represents the cardiac output of 4.86 l/min for this patient.

The calibration parameters were defined as $\boldsymbol{\beta} = [\boldsymbol{\beta}^{\text{out}}, \boldsymbol{\beta}^{\text{in}}]$, where $\boldsymbol{\beta}^{\text{in}}$ represents the Fourier coefficients parameterizing the inflow waveform, while $\boldsymbol{\beta}^{\text{out}}$ represents the $\mathcal{RCR}$ boundary conditions enforced at the outlets of the geometry. The multiplicative correction waveform was constructed using 12 Fourier modes. Compared with the previous calibration studies using synthetic target data, where pressure and flow constraints were enforced at all boundaries, the clinical dataset only provides measurements at the aortic root. Consequently, the estimation of optimal boundary conditions becomes an underdetermined nonconvex optimization problem, with non-unique solution. To ensure physiologically admissible solutions during optimization, we introduce additional penalty terms that restrict non-physiological backflow behavior at the outlets, specifically limiting the maximum reverse flow to -2 mL/s. Following calibration, the predicted aortic pressure and cardiac output from $\boldsymbol{\alpha}^*$-ROM with the calibrated parameters $\boldsymbol{\beta}^*$ are in excellent agreement with the clinical observations and are well within the estimated uncertainty intervals, as shown in Fig. 8a. The corresponding 3D CFD results, where the 3D CFD simulation was performed using the optimized inflow waveform and $\mathcal{RCR}$ outlet boundary condition parameters, indicate a maximum error of 2.59% in the aortic pressure measurement, and 1% in the cardiac output prediction, while ensuring the flow reversal is not more than -2 ml/s at all outlets shown in Fig. 8d.

Stochastic Boundary Condition Tuning

In this section, we perform Bayesian calibration of the boundary conditions in the aortic model while accounting

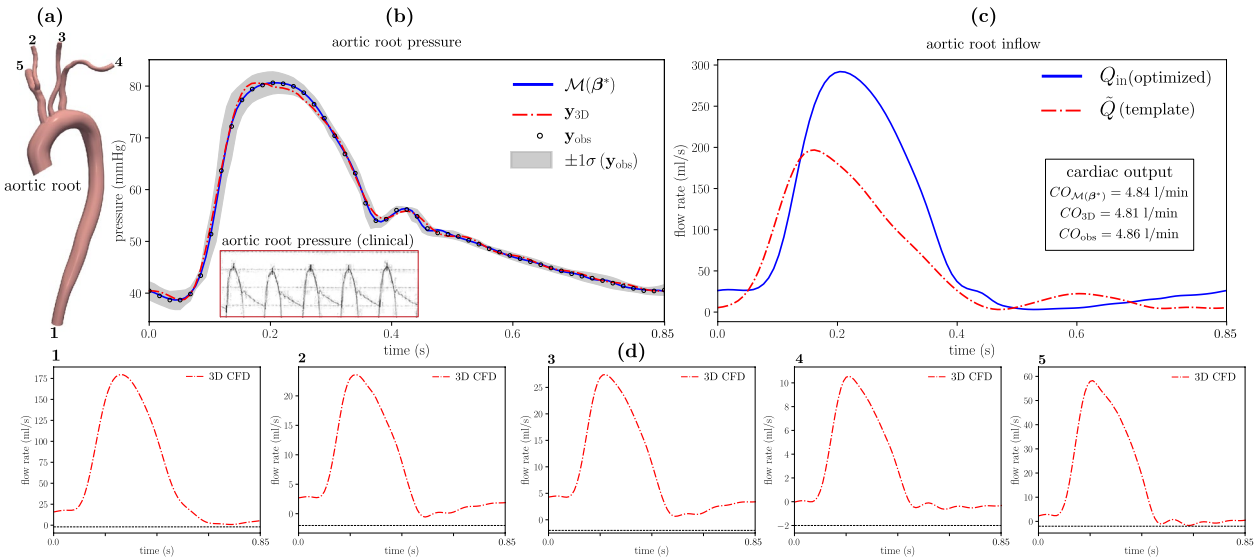

Fig. 8 **a** Patient-specific aortic geometry from CT imaging. **b** Predicted aortic pressure from $\boldsymbol{\alpha}^*$-ROM (defined as $\mathcal{M}$) using calibrated boundary conditions $\boldsymbol{\beta}^*$ to match clinical observations $\mathbf{y}_{\text{obs}}$ obtained by digitizing pressure catheterization data (shown in red box), as well as 3D CFD simulation results ($\mathbf{y}_{\text{3D}}$) using $\boldsymbol{\beta}^*$. **c** Optimized inflow Q_{in} compared to template initial inflow waveform $\bar{Q}$ along with the comparison of cardiac output from $\boldsymbol{\alpha}^*$-ROM and 3D CFD predictions using $\boldsymbol{\beta}^*$ and observed clinical data (black box). **d** Flow waveforms at all outlets showing the threshold of -2 ml/s (black dashed line) imposed through penalty terms

for uncertainties in the observed data. We consider the average inlet pressure and the average flow rates at the outlets as the observational data, defined as $\mathbf{y}_{\text{obs}} = [P^{\text{in}}_{\text{avg}}, \{Q^{\text{i}}_{\text{avg}}, \text{ for } \text{i} \in [1, \cdots, n_o]\}]$, where these observations are generated by adding noise to the 3D CFD simulation results. Specifically, $\mathbf{y}^{\text{noisy}}_{\text{obs}} = \mathbf{y}_{\text{3D}} + \boldsymbol{\epsilon}$, with $\boldsymbol{\epsilon}$ sampled from a zero-mean Gaussian distribution $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, where Σ is the same diagonal covariance matrix used in the likelihood function in Eq. 22. Since the observations consist of average pressure and flow rates, which are primarily determined by the total resistance of the $\mathcal{RCR}$ Windkessel boundary conditions, we focus on calibrating the total resistance $\mathcal{R}_{\text{t}}$. A fixed ratio is then used to distribute $\mathcal{R}_{\text{t}}$ between the proximal and distal resistances at each outlet, while the boundary condition compliances are held constant. This reduces calibration strategy also yields substantial computational efficiency. We utilized the RMH, SMC, and HMC sampling strategies described in Sect. "Stochastic Parameter Tuning" to approximate the posterior distribution over $\boldsymbol{\beta}$ and subsequently estimate single-valued parameters using Eq. 25. To determine an appropriate sample size for each method, we conducted a sample size study and identified a threshold beyond which additional samples provided less than 1% improvement in accuracy.

The No-U-Turn sampler (NUTS) [42] was used to tune the HMC sampling parameters, including an adapted step size of 0.66 and a diagonal mass matrix $\mathbf{M} = \text{diag}(0.0063, 0.013, 0.014, 0.012, 0.014)$. For SMC, the effective sample size threshold was set to 0.85, with rejuvenation and resampling performed whenever the effective sample size fell below this threshold. This improved posterior exploration and reduced sample degeneracy, although at the expense of increased computational cost. Table 4 summarizes the comparison of estimated average pressures and flow rates from the posterior distributions obtained using different sampling methods, with observational data perturbed by up to 20% relative uncertainty. All three sampling methods show consistent agreement with the observations for both pressure and average flow rates, with HMC approach achieving accuracy comparable to RMH and SMC while requiring substantially fewer samples. Moreover, when $\boldsymbol{\beta}^*$ estimated from HMC samples is used to run 3D CFD simulations, the resulting hemodynamic quantities show strong agreement with both the α-ROM predictions and the observed target data. We note that because the observational data are randomly and independently perturbed, mass conservation is not necessarily satisfied for the target observed values, which explains the deviation observed in Q^1_{avg} between the target observations and the physics-based simulation, in which mass conservation is always enforced.

Among the methods, SMC required the longest computational time, primarily due to the numerous tempering steps needed to capture the sharply peaked posterior. In contrast, HMC achieved the lowest computational cost, requiring 90 s for window adaptation warmup and 40 s for the actual inference phase on a workstation equipped with an Intel Core i7-14700 CPU (up to 5.4 GHz) running Ubuntu 24.04 LTS. This efficiency stems from HMC's gradient-aware updates, which enable rapid posterior estimation with relatively few samples ~ 500. The differentiable nature of our framework allows seamless integration with the HMC sampling pipeline, as the gradients required for Hamiltonian dynamics are computed on-the-fly and do not need to be precomputed prior to posterior estimation.

The joint distribution of all $\mathcal{R}_{\text{t}}$ parameters is shown in Fig. 9. As illustrated, the HMC-derived distributions closely match those obtained from RMH, which we consider the ground truth among the MC-based methods, despite HMC requiring only approximately 1/100 of the RMH sample size.

To assess the effect of relative uncertainty on HMC performance, we inferred posterior distributions under three uncertainty levels in the observed data: 5%, 10%, and 20%. The corresponding results are presented in Fig. 10. As the relative uncertainty decreases, the posterior distributions

Table 4 Predicted hemodynamic values obtained from the $\boldsymbol{\alpha}^*$-ROM using the expected value $\boldsymbol{\beta}^*$, estimated via HMC, SMC, and RMH sampling methods and Eq. 25, along with the required number of samples for convergence and the computational time for posterior estimation

	$P^{\text{in}}_{\text{avg}}$ (mmHg)	Q^1_{avg} (ml/s)	Q^2_{avg} (ml/s)	Q^3_{avg} (ml/s)	Q^4_{avg} (ml/s)	Q^5_{avg} (ml/s)	Avg error (%) –	# of samples –	Time (s)
$\mathbf{y}_{\text{obs}}$	87.18	36.08	8.76	6.73	6.76	2.46	–	–	–
$\mathbf{y}_{\text{3D}}(\boldsymbol{\beta}^*_{\text{HMC}})$	85.58	32.03	8.01	6.09	6.22	2.19	8.35	–	–
HMC	85.84	32.31	8.05	6.19	6.29	2.20	7.61	500	90 + 40
RMH	85.48	32.55	8.08	6.08	6.10	2.25	7.91	100000	900
SMC	85.67	32.60	8.01	6.12	6.13	2.18	8.28	20000	6000

$\mathbf{y}_{\text{3D}}$ represents hemodynamic values obtained from 3D CFD using the expected value $\boldsymbol{\beta}^*$ from HMC samples. The average error (avg error) is defined as the mean relative percentage error computed across the six target observations, $\mathbf{y}_{\text{obs}}$

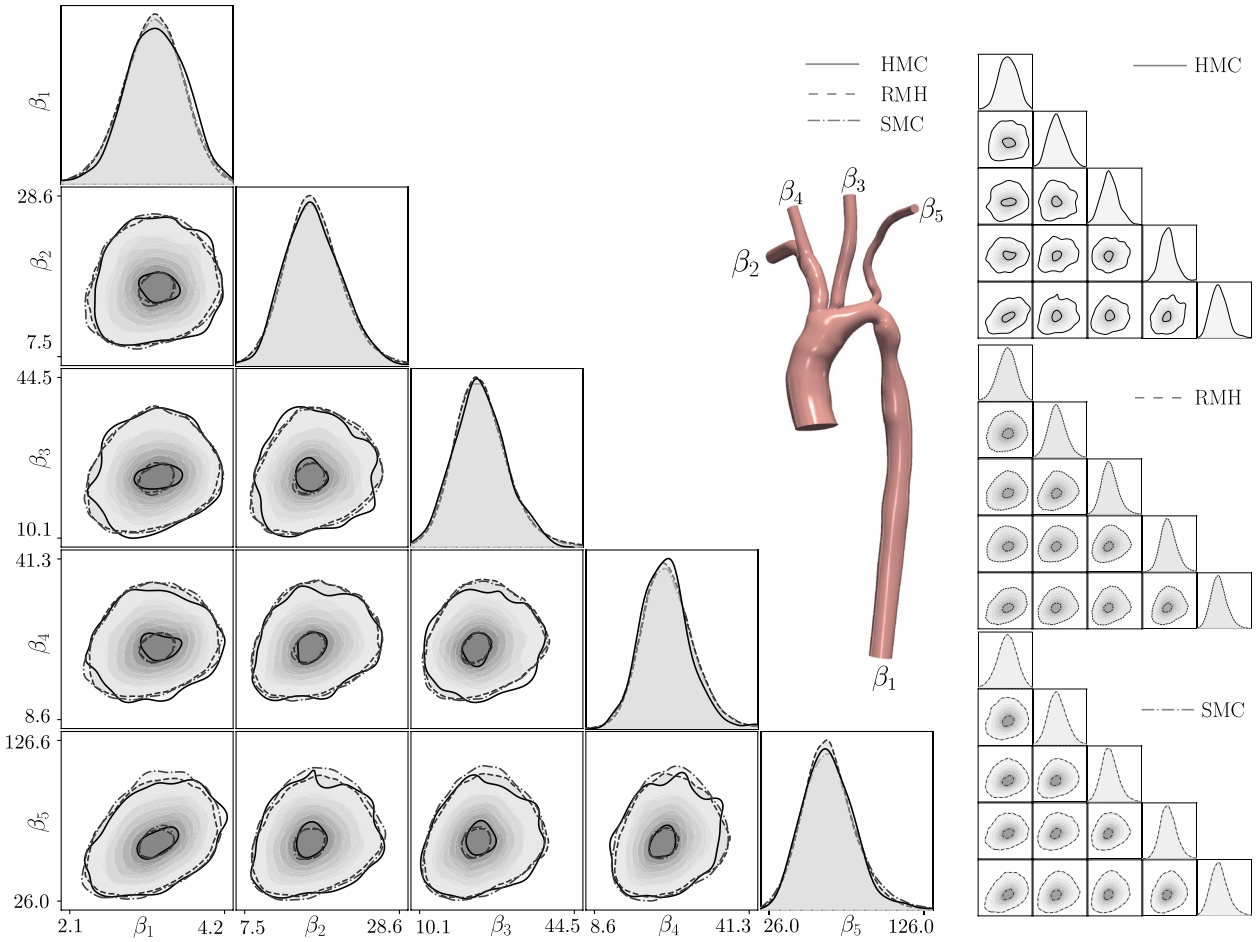

Fig. 9 Posterior distributions between total resistances β_i (defined as $\mathcal{R}_t \times 10^{-3}$) with corresponding one-dimensional kernel density estimates obtained using HMC, RMH, and SMC sampling methods

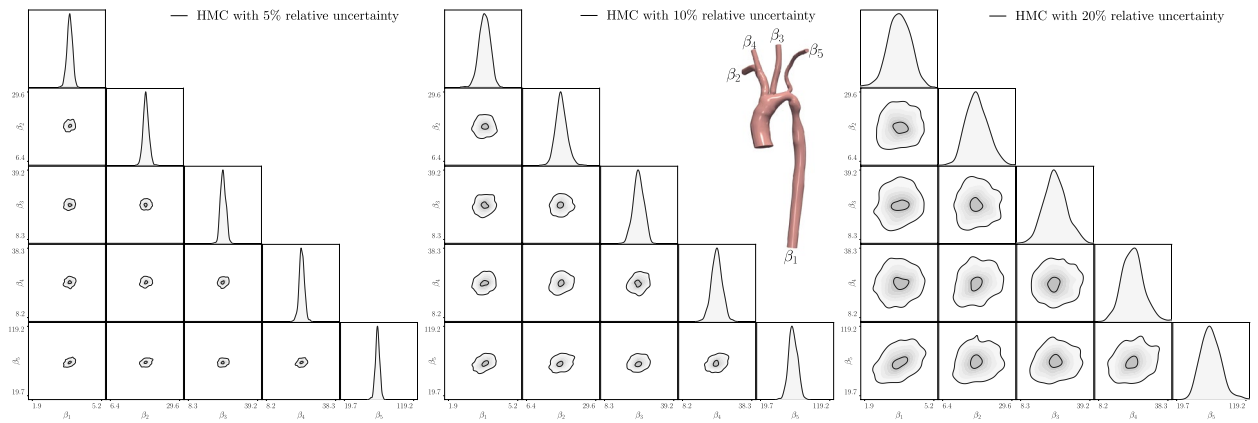

Fig. 10 Posterior distributions between total resistances β_i (defined as $\mathcal{R}_t \times 10^{-3}$) with corresponding one-dimensional kernel density estimates obtained using HMC sampling method for observational uncertainties of 5%, 10%, and 20%

become increasingly sharp. HMC is able to accurately capture these sharper posteriors with relatively few samples, whereas RMH would require a substantially larger number of samples to resolve the narrow peaks, and SMC would require many tempering steps to transition from an uninformative prior to the sharply concentrated posterior.

Discussion

We have developed an end-to-end differentiable framework that enables (i) accurate surrogate modeling, (ii) rapid model calibration, and (iii) automated deterministic and stochastic boundary condition tuning—thereby unifying three fundamental capabilities that substantially advance personalized cardiovascular modeling and simulation. The inherent differentiability of the pipeline allows us to exploit state-of-the-art gradient-based optimization methods, eliminating the need for custom adjoint implementations and avoiding the scalability limitations commonly encountered with derivative-free approaches in cardiovascular applications. Furthermore, this differentiability facilitates seamless integration with gradient-informed Monte Carlo sampling strategies, enabling efficient Bayesian boundary condition calibration in the presence of observational uncertainties.

A recent study [21] employed a Bayesian framework to optimize cardiovascular boundary conditions using surrogate models. While our approach similarly leverages surrogate modeling calibrated to a single 3D CFD simulation, it addresses key limitations of this work. First, their framework primarily focused on stochastic boundary condition tuning, where deterministic calibration emerged only implicitly through posterior expectation estimates. Although uncertainty-aware calibration is valuable in clinical decision-making, the associated SMC methodology imposes substantial computational overhead, limiting its applicability in clinical and translational settings that require rapid patient-specific model calibration. In contrast, our framework explicitly supports both deterministic and stochastic calibration as complementary functionalities within a unified differentiable framework. Specifically, we achieve substantial computational efficiency gains enabling timely model calibration through both gradient-based deterministic optimization and gradient-informed HMC sampling for stochastic inference. Second, their study considered calibration of only outlet resistances, with compliance constrained through a fixed resistance-compliance relationship, for a single aortic model involving five parameters and a limited set of *scalar* hemodynamic targets under varying signal-to-noise ratios. This restricted formulation largely reflects the limited scalability of SMC-based stochastic calibration, where increasing parameter dimensionality substantially exacerbates convergence and computational costs. In contrast, our framework provides a significantly more comprehensive and scalable formulation, supporting diverse boundary conditions including $\mathcal{RCR}$ Windkessel model, coronary-specific LPNs, and inflow waveform parameterization. Moreover, the framework can readily incorporate additional boundary condition models with minimal implementation effort due to its native differentiability and modular design. We further evaluated the framework using both scalar and *temporally* resolved hemodynamic targets derived from synthetic CFD simulations and clinical datasets, demonstrating robustness across heterogeneous calibration settings.

Table 5 Number of calibrated parameters in the parametric $\boldsymbol{\alpha}$-ROM and the corresponding average per-iteration optimization time for aortic, aortofemoral, and coronary models

Vascular model	# of parameters	Optimization time/iter (s)
Aortic model	32	0.012
Aortofemoral model	64	0.034
Coronary model	88	0.056

The parametric ROM developed in this study demonstrates reliable prediction of temporal hemodynamic waveforms across cardiovascular domains with distinct geometric complexity and flow characteristics, and it consistently outperforms the previously proposed DLP approach. Whereas the DLP model is a purely physics-based ROM that depends on analytical and semi-empirical pressure loss formulations with predefined nonlinearities, whose accuracy degrades across heterogeneous vascular anatomies and flow characteristics, the proposed parametric surrogate integrates mechanistic insights with data-driven techniques to approximate a generalized resistance as a nonlinear function of flow rate alongside a tunable inductance term over a range of hemodynamic conditions, both calibrated efficiently to a single high-fidelity 3D CFD simulation. For instance, in the coronary model, where calibration involves 88 parameters, the average computational cost per optimization iteration is approximately 0.056 s as reported in Table 5, with convergence typically achieved in fewer than 15 iterations. All surrogate model calibrations and boundary condition optimizations were performed on a workstation equipped with an Intel Core i7-14700 CPU (up to 5.4 GHz) running Ubuntu 24.04 LTS. This substantial computational efficiency is enabled primarily by just-in-time (JIT) compilation in JAX, the framework used for model development, with performance benefits becoming increasingly pronounced as the number of model parameters grows.

The surrogate model was subsequently leveraged to estimate optimal boundary conditions to match a set of target hemodynamic observations through a deterministic approach

formulated as a least-squares inverse problem that is solved using gradient-based LM optimization. The method demonstrated consistent and robust performance in calibrating boundary conditions across diverse cardiovascular domains by accommodating a wide range of synthetically generated and clinical observations, including time-averaged scalar metrics, time-dependent waveform targets, and combinations of the two. Importantly, even with poor initialization of the boundary conditions, particularly in the coronary circulation, where proximity to the heart and myocardial compression create unique hemodynamics that make calibration extremely challenging [45], the optimizer successfully recovered boundary condition parameters that yielded pressure and flow rates in strong agreement with the target observations. This robustness stems from the curvature-aware LM algorithm used to solve the inverse problem: LM adaptively blends gradient descent, which provides stability when far from the solution, with Gauss–Newton updates, which accelerate convergence near the optimum. The damping parameter enables dynamic transitioning between these regimes, enabling reliable convergence even from weak or highly inaccurate initializations.

To complement the single value estimates of optimal boundary conditions obtained via the deterministic method, we employed Monte Carlo-based approaches to estimate posterior distributions over boundary condition parameters by propagating uncertainties from the observed data, which typically requires large ensembles of simulations to achieve convergence. By leveraging gradient information, HMC drastically reduced the number of samples required for convergence, resulting in a total computational speedup of at least 4× for cases tested in this study. This efficiency makes Bayesian calibration significantly more computationally tractable compared to standard Monte Carlo approaches such as RMH and SMC. The key advantage of our framework, however, lies not merely in the use of HMC, which has recently been applied for parameter estimation and uncertainty quantification [46], but in the seamless integration of gradient computations, eliminating the need for manual derivation or cumbersome approximations in stochastic model calibration.

In this study, we assumed that the hydraulic resistance is a parabolic function of flow, consistent with physics-based energy dissipation models, particularly at sudden expansions and junctions [23]. Table 7 in Appendix 4 compares the average maximum outlet pressure errors obtained using different polynomial degree during $\boldsymbol{\alpha}$-ROM calibration for the aortic, aortofemoral, and coronary models. Our observation demonstrated that the cubic representation ($d = 3$) exhibits calibration difficulties, particularly for the aortofemoral model. This behavior is likely attributable to the introduction of non-physiological nonlinearities, where resistance scales proportionally to Q^3, leading to increased optimization complexity and potential overfitting. Among all tested polynomial degrees, $d = 2$ consistently produced the most accurate and robust results across the considered models. The differentiability and modularity of our pipeline enable straightforward adaptation to any form of nonlinearity beyond polynomial representation in defining the system dynamics including hydraulic resistances and inductances, which is not readily achievable in prior works that require substantial reformulation and redevelopment of adjoint-based derivations [21].

Despite these promising results, a few limitations remain. The use of LM algorithm, while robust, does not scale efficiently to problems with a large number of parameters for parametric ROM calibration and boundary condition tuning. For instance, we evaluated the parametric ROM approach on an extensive pulmonary model with 508 parameters, which required approximately 1,200 s per LM optimization iteration, making its use prohibitively expensive. Extending this framework to high-dimensional settings will necessitate the use of first-order gradient-based optimizers, such as ADAM [47], which can be seamlessly integrated into our differentiable pipeline. For example, we utilized the ADAM optimizer to perform a preliminary $\boldsymbol{\alpha}$-ROM calibration for a pulmonary model consisting of 508 parameters, demonstrating the scalability of the proposed framework. Each optimization iteration required approximately 3.0 s computational time, and approximately 1,000 iterations were needed to achieve agreement with the 3D CFD results. Computational efficiency can be further improved through advanced optimization strategies, such as adaptive learning rate scheduling. The corresponding results are indicated in Appendix 5.

Another limitation of this study is that the stochastic calibration results obtained using HMC and SMC are validated against RMH using ROM predictions, rather than directly against full 3D CFD simulations, which would be computationally prohibitive. Given the demonstrated accuracy of the proposed parametric ROM in reproducing 3D CFD hemodynamics, this surrogate-based validation is a reasonable and practically justified approximation, as further supported in our study, where $\boldsymbol{\beta}^*$ estimated from HMC samples is used to perform 3D CFD simulations for reproducing target hemodynamics shown in Table 4. Finally, although gradient-informed HMC sampling effectively reduces the number of samples compared to conventional Monte Carlo methods, it is sensitive to disparities in parameter scales, which can lead to numerical instabilities. This limitation is not specific to our framework but inherent to HMC itself and can be mitigated through reparameterization strategies or by constructing non-dimensionalized physics-based models.

Finally, the proposed pipeline does not account for vessel wall deformability. Previous studies have demonstrated that pathologies such as aortic aneurysms involve substantial wall motion, and that predicted flow and pressure waveforms can differ between rigid wall and deformable wall assumptions in cardiovascular flow simulations [48]. In addition,

physiological phenomena such as pulse wave propagation, which play an important role in cardiovascular physiology and diagnostics, are not captured within the current framework. However, owing to the modular and differentiable structure of the proposed pipeline, the framework can be readily extended to incorporate vessel compliance, by representing vascular segments with $\mathcal{RLC}$ lumped parameter models.

Appendix 1 3D CFD Simulation

The open-source SimVascular [1] was used for 3D CFD simulation of blood flow using numerical solution of incompressible Navier–Stokes equations

$$\begin{aligned} \rho \mathbf{v}_{,t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} &= -\nabla p + \nabla . \tilde{\tau} \\ \nabla . \mathbf{v} &= 0 \\ \tilde{\tau} &= \mu \big(\nabla \mathbf{v} + (\nabla \mathbf{v})^{\mathrm{T}} \big), \end{aligned} \tag{33}$$

where $\mathbf{v}$ is the fluid velocity, p is the pressure, and τ is the viscous stress tensor. The vessel walls were assumed to be rigid.

Appendix 2 MCMC Algorithms

Random-Walk Metropolis-Hastings (RWMH)

Algorithm 1 Random-Walk Metropolis-Hastings (RWMH)

Require: Initial state $\boldsymbol{\theta}^{(0)}$, number of samples N, proposal covariance $\mathbf{C}$, unnormalized target density $\pi(\boldsymbol{\theta})$
for $n = 0, 1, \ldots, N-1$ **do**
  Sample $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{C})$
  Propose $\boldsymbol{\theta}' = \boldsymbol{\theta}^{(n)} + \boldsymbol{\varepsilon}$
  $A = \min\left(1, \frac{\pi(\boldsymbol{\theta}')}{\pi(\boldsymbol{\theta}^{(n)})}\right)$
  Sample $u \sim \mathcal{U}(0, 1)$
  if $u \leq A$ **then**
    $\boldsymbol{\theta}^{(n+1)} \leftarrow \boldsymbol{\theta}'$
  else
    $\boldsymbol{\theta}^{(n+1)} \leftarrow \boldsymbol{\theta}^{(n)}$
  end if
end for
return $\{\boldsymbol{\theta}^{(n)}\}_{n=1}^{N}$

Adaptive Tempered Sequential Monte Carlo (SMC)

Algorithm 2 Adaptive Tempered Sequential Monte Carlo (SMC)

Require: Prior $P(\boldsymbol{\theta})$, likelihood $P(\mathbf{y} \mid \boldsymbol{\theta})$ in the log-transformed space $\boldsymbol{\theta} = T(\boldsymbol{\beta}) = \log \boldsymbol{\beta}$ (Eq. 22), with the Jacobian $|\det \mathbf{J}_{T^{-1}}(\boldsymbol{\theta})|$ carried in $P(\boldsymbol{\theta})$, particles N, target ESS ratio $\gamma \in (0, 1)$

$n \leftarrow 0$, $\lambda^{(0)} \leftarrow 0$

Sample $\boldsymbol{\theta}^{(i,0)} \sim P(\boldsymbol{\theta})$ for $i = 1, \ldots, N$

Set $\tilde{w}^{(i,0)} \leftarrow 1/N$ for $i = 1, \ldots, N$

while $\lambda^{(n)} < 1$ **do**

For a trial temperature $\lambda \in (\lambda^{(n)}, 1]$, define the incremental weights

$$w^{(i)}(\lambda) \leftarrow \tilde{w}^{(i,n)} \left[P(\mathbf{y} \mid \boldsymbol{\theta}^{(i,n)}) \right]^{\lambda - \lambda^{(n)}}, \qquad \tilde{w}^{(i)}(\lambda) \leftarrow \frac{w^{(i)}(\lambda)}{\sum_{j=1}^{N} w^{(j)}(\lambda)}$$

Choose $\lambda^{(n+1)} \in (\lambda^{(n)}, 1]$ by bisection such that

$$\mathrm{ESS}\left(\lambda^{(n+1)}\right) = \left(\sum_{i=1}^{N} \left(\tilde{w}^{(i)}(\lambda^{(n+1)}) \right)^2 \right)^{-1} = \gamma N,$$

or set $\lambda^{(n+1)} \leftarrow 1$ if $\mathrm{ESS}(1) \geq \gamma N$

Set $\tilde{w}^{(i,n+1)} \leftarrow \tilde{w}^{(i)}(\lambda^{(n+1)})$ for $i = 1, \ldots, N$

Resample $\{\boldsymbol{\theta}^{(i,n)}\}_{i=1}^{N}$ according to $\{\tilde{w}^{(i,n+1)}\}_{i=1}^{N}$ to obtain $\{\bar{\boldsymbol{\theta}}^{(i,n)}\}_{i=1}^{N}$

Reset $\tilde{w}^{(i,n+1)} \leftarrow 1/N$ for $i = 1, \ldots, N$

for $i = 1, \ldots, N$ **do**

Run MCMC targeting $\pi_{n+1}(\boldsymbol{\theta}) \propto P(\boldsymbol{\theta})\, P(\mathbf{y} \mid \boldsymbol{\theta})^{\lambda^{(n+1)}}$ initialized at $\bar{\boldsymbol{\theta}}^{(i,n)}$

Set $\boldsymbol{\theta}^{(i,n+1)}$ to the MCMC output

end for

$n \leftarrow n + 1$

end while

return $\{\boldsymbol{\theta}^{(i,n)}\}_{i=1}^{N}$

Hamiltonian Monte Carlo (HMC)

Algorithm 3 Hamiltonian Monte Carlo (HMC) with Leapfrog Integrator

Require: Initial state $\boldsymbol{\theta}^{(0)}$, mass matrix $\mathbf{M}$, step size Δt, steps L, number of samples N
Require: Log target $\ell(\boldsymbol{\theta}) = \log \pi(\boldsymbol{\theta})$, gradient $\nabla \ell(\boldsymbol{\theta})$
for $n = 0, 1, \ldots, N-1$ **do**
  Sample $\boldsymbol{\psi} \sim \mathcal{N}(\mathbf{0}, \mathbf{M})$
  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}^{(n)}$, $\boldsymbol{\psi} \leftarrow \boldsymbol{\psi}$
  $(\boldsymbol{\theta}', \boldsymbol{\psi}') \leftarrow \textsc{Leapfrog}(\boldsymbol{\theta}, \boldsymbol{\psi}, \Delta t, L, \mathbf{M}, \nabla \ell)$
  $U(\boldsymbol{\theta}) \leftarrow -\ell(\boldsymbol{\theta})$
  $K(\boldsymbol{\psi}) \leftarrow \frac{1}{2} \boldsymbol{\psi}^\top \mathbf{M}^{-1} \boldsymbol{\psi}$
  $\Delta H \leftarrow \big(U(\boldsymbol{\theta}') + K(\boldsymbol{\psi}')\big) - \big(U(\boldsymbol{\theta}^{(n)}) + K(\boldsymbol{\psi})\big)$
  $A \leftarrow \min\big(1, \exp(-\Delta H)\big)$
  Sample $u \sim \mathcal{U}(0, 1)$
  if $u \leq A$ **then**
    $\boldsymbol{\theta}^{(n+1)} \leftarrow \boldsymbol{\theta}'$
  else
    $\boldsymbol{\theta}^{(n+1)} \leftarrow \boldsymbol{\theta}^{(n)}$
  end if
end for
return $\{\boldsymbol{\theta}^{(n)}\}_{n=1}^{N}$

Algorithm 4 Leapfrog integrator

Require: Initial $(\boldsymbol{\theta}, \boldsymbol{\psi})$, step size Δt, steps L, mass matrix $\mathbf{M}$, gradient $\nabla \ell(\boldsymbol{\theta})$
for $\ell = 1, 2, \ldots, L$ **do**
  $\boldsymbol{\psi} \leftarrow \boldsymbol{\psi} + \dfrac{\Delta t}{2} \nabla \ell(\boldsymbol{\theta})$
  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \Delta t \, \mathbf{M}^{-1} \boldsymbol{\psi}$
  $\boldsymbol{\psi} \leftarrow \boldsymbol{\psi} + \dfrac{\Delta t}{2} \nabla \ell(\boldsymbol{\theta})$
end for
return $(\boldsymbol{\theta}, \boldsymbol{\psi})$

Appendix 3 Patient and Model Information

See Table 6

Table 6 Summary of vascular models used in this study

Model	Age	Sex	Disease/Clinical condition	Source	Identifier
Aortic model	8	Female	Aortic coarctation	Vascular Model Repository	0020_H_AO_COA
Aortofemoral model	70	Male	Abdominal aortic aneurysm	Vascular Model Repository	0038_H_ABAO_AAA
Coronary model	N/A	NA	Coronary artery disease	Perelman School of Medicine at UPenn	Non-public institutional dataset
Pulmonary model	22	Female	Healthy	Vascular Model Repository	0092_H_PULM_H
Aortic model	77	Female	Post-TAVI	Long Beach Memorial Care	Non-public institutional dataset

Appendix 4 Ablation Study for Polynomial Degree of α-ROM

To assess the sensitivity of the proposed $\boldsymbol{\alpha}$-ROM formulation to the choice of polynomial degree in Eq. 3, we performed an ablation study in which the degree d of the nonlinear resistance model was varied while keeping the remaining calibration procedure fixed. For each vascular model and each value of d, the $\boldsymbol{\alpha}$-ROM was calibrated against the corresponding 3D CFD simulation data. We then evaluated the maximum relative pressure error between the calibrated ROM prediction and the 3D CFD pressure waveform at each pressure measurement location. The mean of these maximum errors, calculated across all outlets, $\bar{\varepsilon}_P$, was used to compare the accuracy obtained with different polynomial degrees. As shown in Table 7, the quadratic model generally provides a good balance between accuracy and robustness across the vascular models considered.

$$\varepsilon_P^{(i)} = \max_t \frac{\left|P^{(i)}_{\boldsymbol{\alpha}^*\text{-ROM}}(t) - P^{(i)}_{\text{3D}}(t)\right|}{\left|P^{(i)}_{\text{3D}}(t)\right|}\%, \tag{34}$$

Table 7 Mean maximum relative pressure error for different polynomial degrees across vascular models

Model	$\bar{\varepsilon}_{P,d=1}$	$\bar{\varepsilon}_{P,d=2}$	$\bar{\varepsilon}_{P,d=3}$
Aortic model	0.8397%	0.7087%	0.6416%
Aortofemoral model	0.24%	0.24%	9.76%
Coronary model	1.43%	0.51%	0.46%

$$\bar{\varepsilon}_P = \frac{1}{N_o}\sum_{i=1}^{N_o} \varepsilon_P^{(i)}, \tag{35}$$

where N_o is the number of outlets.

Appendix 5 Pulmonary Model Calibration Using ADAM Optimization

For higher-dimensional vascular networks, such as pulmonary models with extensive branching and many terminal vessels, the computational cost of the LM algorithm becomes prohibitive. For the patient-specific pulmonary geometry obtained from Vascular Model Repository [43], each $\boldsymbol{\alpha}$-ROM calibration step using LM required approximately 1, 200 seconds, making the approach computationally impractical. This cost is primarily associated with forming and solving curvature-based systems, which scale poorly as the number of calibration parameters increases. To address this issue, we calibrated the pulmonary $\boldsymbol{\alpha}$-ROM using the first-order ADAM optimizer, which avoids explicit Hessian or Gauss–Newton matrix construction and is therefore more scalable for large parameter spaces. The calibration was performed against a single reference 3D CFD simulation using a learning rate of 10^{-3} requiring around 2,000 iterations to achieve agreement with 3D CFD simulation results. As shown in Fig. 11, the calibrated $\boldsymbol{\alpha}^*$-ROM closely reproduces the 3D CFD pressure waveforms across the reported locations.

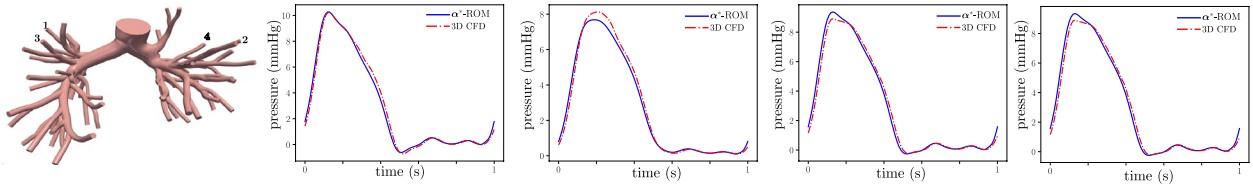

Fig. 11 Patient-specific pulmonary model and representative pressure waveforms from 3D CFD and the $\boldsymbol{\alpha}^*$-ROM calibrated using the ADAM optimizer

Acknowledgements The authors thank Dr. David Shavelle (Chief of Cardiology, MemorialCare, CA, US) for providing CT imaging and pressure catheterization data for TAVI patient used in this study. We acknowledge the computing resources provided by North Carolina State University High-Performance Computing Services Core Facility.

References

1. Updegrove, Adam, Nathan M. Wilson, Jameson Merkow, Hongzhi Lan, Alison L. Marsden, and Shawn C. Shadden. SimVascular: an open source pipeline for cardiovascular simulation. *Ann. Biomed. Eng.*. 45(3):525–541, 2017.
2. Sankaran, Sethuraman, Mahdi Esmaily Moghadam, Andrew M. Kahn, Elaine E. Tseng, Julius M. Guccione, and Alison L. Marsden. Patient-specific multiscale modeling of blood flow for coronary artery bypass graft surgery. *Ann. Biomed. Eng.*. 40(10):2228–2242, 2012.
3. LaDisa Jr, John F., Ismail Guler, Lars E. Olson, Douglas A. Hettrick, Judy R. Kersten, David C. Warltier, and Paul S. Pagel. Three-dimensional computational fluid dynamics modeling of alterations in coronary wall shear stress produced by stent implantation. *Ann. Biomed. Eng.*. 31(8):972–980, 2003.
4. Zhong, Liang, Jun-Mei. Zhang, Su. Boyang, Ru San Tan, John C. Allen, and Ghassan S. Kassab. Application of patient-specific computational fluid dynamics in coronary and intra-cardiac flow simulations: challenges and opportunities. *Front. Physiol.*. 9:742, 2018.
5. Gundert, Timothy J., Alison. L. Marsden, Weiguang. Yang, and John.F. LaDisa Jr. Optimization of cardiovascular stent design using computational fluid dynamics. *J. Biomech. Eng.*. 134(1):011002, 2012.
6. Itu, Lucian, Puneet Sharma, Kristóf Ralovich, Viorel Mihalef, Razvan Ionasec, Allen Everett, Richard Ringel, Ali Kamen, and Dorin Comaniciu. Non-invasive hemodynamic assessment of aortic coarctation: validation with in vivo measurements. *Ann. Biomed. Eng.*. 41(4):669–681, 2013.
7. Arzani, Amirhossein, and Shawn C. Shadden. Characterization of the transport topology in patient-specific abdominal aortic aneurysm models. *Phys. Fluids*.24(8):081901, 2012.
8. Taylor, Charles A., Timothy A. Fonte, and James K. Min. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis. *J. Am. Coll. Cardiol.*. 61(22):2233–2241, 2013.
9. Taylor, Charles A., and David A. Steinman. Image-based modeling of blood flow and vessel wall dynamics: applications, methods and future directions. *Ann. Biomed. Eng.*. 38(3):1188–1203, 2010.
10. Nazir, Anam, Muhammad Nadeem Cheema, Bin Sheng, Huating Li, Ping Li, Po. Yang, Younhyun Jung, Jing Qin, Jinman Kim, and David Dagan Feng. OFF-eNET: an optimally fused fully end-to-end network for automatic deense volumetric 3D intracranial blood vessels segmentation. *IEEE Trans. Image Proc.*. 29:7192–7202, 2020.
11. Cepero, Numi Sveinsson, and Shawn C. Shadden. Seqseg: learning local segments for automatic vascular model construction. *Ann. Biomed. Eng.*. 53(1):158–179, 2025.
12. Vignon-Clementel, Irene. E., C. Alberto. Figueroa, Kenneth. E. Jansen, and A. Charles. Taylor Outflow boundary conditions for three-dimensional finite element modeling of blood flow and pressure in arteries. *Comput. Methods Appl. Mech. Eng.* 195(29):3776–3796, 2006.
13. Givoli, Dan. A tutorial on the adjoint method for inverse problems. *Comput. Methods Appl. Mech. Eng.*.380:113810, 2021.
14. Ismail, Mahmoud, WolfgangA. Wall, and MichaelW. Gee. Adjoint-based inverse analysis of windkessel parameters for patient-specific vascular models. *J. Comput. Phys.*. 244:113–130, 2013.
15. Itu, Lucian, Puneet Sharma, Constantin Suciu, Florin Moldoveanu, and Dorin Comaniciu. Personalized blood flow computations: a hierarchical parameter estimation framework for tuning boundary conditions. *Int. J. Numer. Methods Biomed. Eng.*.33(3):e02803, 2017.
16. Itu, Lucian, Puneet Sharma, Tiziano Passerini, Ali Kamen, Constantin Suciu, and Dorin Comaniciu. A parameter estimation framework for patient-specific hemodynamic computations. *J. Comput. Phys.*. 281:316–333, 2015.
17. Andrew, R. Conn, Katya Scheinberg, and Luis N. Philadelphia: Vicente. Introduction to derivative-free optimization. Society for Industrial and Applied Mathematics: Philadelphia, 2009.
18. Paris Perdikaris and George Em Karniadakis. Model inversion via multi-fidelity Bayesian optimization: a new paradigm for parameter estimation in haemodynamics, and beyond. *J. Royal Soc. Interface*. 13(118):20151107, 2016.
19. Congdon, Peter. Bayesian statistical modelling. Chichester: Wiley, 2001.
20. Christian, P. Robert and George Casella, 2nd ed. New York: Monte Carlo statistical methods. Springer, 2004.
21. Richter, Jakob, Jonas Nitzler, Luca Pegolotti, Karthik Menon, Jonas Biehler, Wolfgang A. Wall, Daniele E. Schiavazzi, Alison L. Marsden, and Martin R. Pfaller. Bayesian Windkessel calibration using optimized zero-dimensional surrogate models. *Philos. Trans. Royal Soci. A: Math. Phys. Eng. Sci.*. 383(2292):20240223, 2025.
22. Hughes, T. J., and J. Lubliner. On the one-dimensional theory of blood flow in the larger vessels. *Math. Biosci.* 18(1):161–170, 1973.
23. Mirramezani, Mehran, and Shawn C. Shadden. A distributed lumped parameter model of blood flow. *Ann. Biomed. Eng.*. 48(12):2870–2886, 2020.
24. Chnafa, C., K. Valen-Sendstad, O. Brina, V. M. Pereira, and D. A. Steinman. Improved reduced-order modelling of cerebrovascular flow distribution by accounting for arterial bifurcation pressure drops. *J. Biomech.* 51:83–88, 2017.
25. Siena, Pierfrancesco, Caterina Balzotti, Michele Girfoglio, Annalisa Quaini, and Gianluigi Rozza. A machine learning-based reduced order model for the investigation of the haemodynamics in coronary artery bypass grafts.
26. Natalia L Rubio, Eric F Darve, and Alison L Marsden. Data-driven bifurcation handling in physics-based reduced-order vascular hemodynamic models. *Comput. Methods Prog. Biomed.* 276:109230, 2025.
27. Tenderini, Riccardo, and Simone Deparis. Model order reduction of hæmodynamics by space-time reduced basis and reduced fluid-structure interaction. *Comput. Methods Appl. Mech. Eng.*.447:118347, 2025.
28. Neal, Radford M. Mcmc using hamiltonian dynamics. *Handbook Markov Chain Monte Carlo*. 2(11):113–162, 2011.

29. Andrieu, Christophe, Nando De Freitas, Arnaud Doucet, and Michael I. Jordan. An introduction to mcmc for machine learning. *Mach. Learn.*. 50(1–2):5–43, 2003.
30. Del Moral, Pierre, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *J. Royal Stat. Soc.: Series B (Stat. Methodol.)*. 68(3):411–436, 2006.
31. Bradbury, J., R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne. JAX: composable transformations of python+numpy programs. 2018.
32. Baydin, Atilim Gunes, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *J. Marchine Learn. Res.*. 18:1–43, 2018.
33. Bezgin, Deniz A., Aaron B. Buhendwa, and Nikolaus A. Adams. Jax-fluids: a fully-differentiable high-order computational fluid dynamics solver for compressible two-phase flows. *Comput. Phys. Commun.*.282:108527, 2023.
34. Samuel Schoenholz and Ekin Dogus Cubuk. Jax md: a framework for differentiable physics. *Adv. Neural Inf. Proc. Syst.*. 33:11428–11441, 2020.
35. Pin Wen Guan. Differentiable thermodynamic modeling. *Scripta Materialia*.207:114217, 2022.
36. Du, H., and QiZhi He. JAX-MPM: A learning-augmuented differentiable meshfree framework for gpu-accelerated lagrangian simulation and geophysical inverse modeling. 2025.
37. Mirramezani, Mehran, and Shawn C. Shadden. Distributed lumped parameter modeling of blood flow in compliant vessels. *J. Biomech.*.140:111161, 2022.
38. Hairer, E., and G. Wanner. Solving ordinary differential qquations II: Stiff and differential-algebraic Problems, volume 14 of Springer Series in Computational Mathematics. Springer, 2nd edition, 1996.
39. Cabezas, Alberto, Adrien Corenflos, and Junpeng Lao. and Rémi Louf. Blackjax: Composable Bayesian inference in JAX, 2024.
40. Doucet, Arnaud, Nando de Freitas, and Neil J. Gordon (eds.). Sequential Monte Carlo methods in practice. New York: Springer, 2001.
41. Betancourt, Michael, and Mark Girolami. Hamiltonian monte Carlo for hierarchical models. *Curr. Trends Bayesian Methodol. Appl.*. 79(30):2–4, 2015.
42. Hoffman, Matthew D., Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* 15(1):1593–1623, 2014.
43. Wilson, Nathan M., Ana K. Ortiz, and Allison B. Johnson. The vascular model repository: a public resource of medical imaging data and blood flow simulation results. *J. Med. Devices*.7(4):040923, 2013.
44. Alastruey, Jordi, Ashraf W. Khir, Koen S. Matthys, Patrick Segers, Spencer J. Sherwin, Pascal R. Verdonck, Kim H. Parker, and Joaquim Peiró. Pulse wave propagation in a model human arterial network: assessment of 1-D visco-elastic simulations against in vitro measurements. *J. Biomec.*. 44(12):2250–2258, 2011.
45. Tran, Justin S., Daniele E. Schiavazzi, Abhay B. Ramachandra, Andrew M. Kahn, and Alison L. Marsden. Automated tuning for parameter identification and uncertainty quantification in multi-scale coronary simulations. *Computers & Fluids*. 142:128–138, 2017. Selected papers of 18th International Conference on Finite Elements in Flow Problems.
46. Salvador, Matteo, Francesco Regazzoni, Luca Dede', and Alfio Quarteroni. Fast and robust parameter estimation with uncertainty quantification for the cardiac function. *Computer Methods and Programs in Biomedicine*.231:107402, 2023.
47. Diederik, P., Jimmy Kingma, and Ba. ADAM: A method for stochastic pptimization. In: International Conference on Learning Representations, 2015. arXiv:1412.6980
48. Alberto Figueroa, C., Irene E. Vignon-Clementel, Kenneth E. Jansen, Thomas J.R.. Hughes, and Charles A. Taylor, A coupled momentum method for modeling blood flow in three-dimensional deformable arteries. Comput. Methods Appli. Mech. Eng. 195(41):5685–5706, John H. Part II: Argyris Memorial Issue, 2006.