

Personalized and uncertainty-aware coronary hemodynamics simulations: From Bayesian estimation to improved multi-fidelity uncertainty quantification

Karthik Menon ¹, Andrea Zanoni ^{2, 3}, Owais Khan⁴, Gianluca Geraci ⁵,
Koen Nieman ^{6, 7}, Daniele E. Schiavazzi ⁸, Alison L. Marsden ^{2, 3, 9}

¹ *Woodruff School of Mechanical Engineering and Coulter Department of Biomedical Engineering,
Georgia Institute of Technology, Atlanta, GA, USA*

² *Department of Pediatrics (Cardiology), Stanford School of Medicine, Stanford, CA, USA*

³ *Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA, USA*

⁴ *Department of Electrical, Computer, and Biomedical Engineering, Toronto Metropolitan University, Toronto, ON, Canada*

⁵ *Center for Computing Research, Sandia National Laboratories, Albuquerque, NM, USA*

⁶ *Division of Cardiovascular Medicine, Stanford School of Medicine, Stanford, CA, USA*

⁷ *Department of Radiology, Stanford School of Medicine, Stanford, CA, USA*

⁸ *Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame, IN, USA*

⁹ *Department of Bioengineering, Stanford University, Stanford, CA, USA*

Email for correspondence: karthik.menon@me.gatech.edu

Background: Non-invasive simulations of coronary hemodynamics have improved clinical risk stratification and treatment outcomes for coronary artery disease, compared to relying on anatomical imaging alone. However, simulations typically use empirical approaches to distribute total coronary flow amongst the arteries in the coronary tree, which ignores patient variability, the presence of disease, and other clinical factors. Further, uncertainty in the clinical data often remains unaccounted for in the modeling pipeline.

Objective: We present an end-to-end uncertainty-aware pipeline to (1) personalize coronary flow simulations by incorporating vessel-specific coronary flows as well as cardiac function; and (2) predict clinical and biomechanical quantities of interest with improved precision, while accounting for uncertainty in the clinical data.

Methods: We assimilate patient-specific measurements of myocardial blood flow from clinical CT myocardial perfusion imaging to estimate branch-specific coronary artery flows. Simulated noise in the clinical data is used to estimate the joint posterior distributions of the model parameters using adaptive Markov Chain Monte Carlo sampling. Additionally, the posterior predictive distribution for the relevant quantities of interest is determined using a new approach combining multi-fidelity Monte Carlo estimation with non-linear, data-driven dimensionality reduction. This leads to improved correlations between high- and low-fidelity model outputs.

Results: Our framework accurately recapitulates clinically measured cardiac function as well as branch-specific coronary flows under measurement noise uncertainty. We observe substantial reductions in confidence intervals for estimated quantities of interest compared to single-fidelity Monte Carlo estimation and state-of-the-art multi-fidelity Monte Carlo methods. This holds especially true for quantities of interest that showed limited correlation between the low- and high-fidelity model predictions. In addition, the proposed multi-fidelity Monte Carlo estimators are significantly cheaper to compute than traditional estimators, under a specified confidence level or variance.

Conclusions: The proposed pipeline for personalized and uncertainty-aware predictions of coronary hemodynamics is based on routine clinical measurements and recently developed techniques for CT myocardial perfusion imaging. The proposed pipeline offers significant improvements in precision and reduction in computational cost.

1. INTRODUCTION

Patient-specific computational modeling of coronary blood flow is an emerging technique for non-invasive diagnosis and treatment planning for coronary artery disease (CAD) [1, 2]. In large-scale randomized clinical trials, the use of patient-specific computational fluid dynamics (CFD) simulations to guide treatment planning for CAD resulted in improved diagnostic accuracy and fewer unnecessary invasive procedures compared to clinical decisions guided by anatomical imaging alone [3, 4, 5]. Computational modeling has also been used to correlate disease severity, progression, and treatment outcomes with biomechanical stimuli which are often inaccessible from imaging alone. Some examples include the use of computational modeling to correlate vessel wall shear stress with coronary artery bypass graft surgery outcomes [6, 7], the prediction of coronary artery bypass graft failure [8, 9, 10, 11], and simulation-derived metrics to quantify the risk of myocardial infarction [12, 13].

A crucial element in using computational models to guide clinical decisions is their personalization based on available patient-specific clinical data. However, there remain two major challenges in this context. First, current models are personalized primarily in terms of the anatomy, i.e., simulations use anatomical models constructed from patient-specific clinical imaging data. While additional personalization has been explored, generally in the form of tuning boundary conditions to match gross hemodynamic measurements of total flow and/or blood pressure, this does not identify personalized flow splits to each vessel in the coronary tree. Instead, flow splits are generally prescribed using empirical Murray's law-based relationships between vessel diameters and the associated flow [14, 15]. This does not account for vascular disease and variability amongst individuals.

This deficiency was highlighted in previous work by Menon *et al.* [16], Xue *et al.* [17] and Gao *et al.* [18]. These studies achieved personalized coronary flow distributions by leveraging dynamic stress CT myocardial perfusion imaging (MPI_{CT}), a non-invasive imaging technique to quantify the distribution of blood flow in the myocardium [19]. Menon *et al.* showed significant differences in coronary flows between models that used empirical diameter-based flow in each coronary artery versus those based on MPI_{CT}. The accurate assessment of coronary flow is an important factor in the diagnosis of microvascular dysfunction [20] as well as in coronary artery bypass graft surgery outcomes [21]. Moreover, all three studies showed differences in the estimation of *fractional flow reserve* (FFR) [22] between models based on vessel-specific personalization using MPI_{CT} versus the empirical diameter-based distribution. While invasive measurement of FFR is currently the gold-standard for the clinical assessment of CAD severity, non-invasive estimation of FFR using computational modeling has improved outcomes and reduced unnecessary invasive procedures in large randomized clinical trials [3, 4, 5]. In this context, Xue *et al.* [17] and Gao *et al.* [18] showed improved accuracy, specificity and sensitivity in detecting hemodynamically significant coronary artery lesions in simulations incorporating MPI_{CT} versus empirical vessel diameter-based techniques, when using invasively measured FFR as the reference. This emphasizes the importance of more fine-grained personalization to accelerate the reliable and accurate adoption of hemodynamic models in the clinic.

Second, while the assimilation of clinical data into personalized computational models is promising, it introduces another challenge. Namely, models that are personalized based on clinical measurements usually do not account for the uncertainty in this clinical data. The assimilation of noisy data into computational models naturally induces uncertainty in the model parameters and output quantities of interest (QoIs) [23]. Previous work to address this challenge has largely focused either on the estimation of probabilistic model parameters based on noisy clinical data, or on the propagation to probabilistic hemodynamic predictions, assuming known parameter distributions. The former has been performed by combining model reduction with Bayesian estimation through adaptive Markov chain Monte Carlo (MCMC) [24, 25, 26]. The primary hurdle in the latter has been the prohibitive computational cost of the large number of expensive model evaluations required to estimate uncertain QoIs with reasonable precision. This has previously been tackled with polynomial chaos-based propagation [27, 28, 29], where the degree of success strongly depends on the problem's dimensionality and the smoothness of the underlying stochastic response. More recently, multi-level and multi-fidelity approaches for uncertainty propagation [30] have shown encouraging variance reduction in predicted QoIs for high-dimensional problems. These methods combine outputs from a few expensive high-fidelity simulations with a large number of cheap low-fidelity simulations, and the variance reduction directly depends on the correlation between low- and high-fidelity model outputs [31, 32]. However, depending on assumptions relating to modeling the physics of the problem in the low- and high-fidelity contexts, possibly dissimilar parametrization between these models, etc., this correlation may be limited. Moreover, only a few studies have developed end-to-end, uncertainty-aware pipelines spanning clinical data to model predictions (see, e.g., a study in congenital heart disease [33]), particularly for CAD [32].

This work combines novel clinical imaging and probabilistic modeling to address two challenges: (1) the

estimation of patient-specific and *vessel-specific* coronary flows based on clinical MPI_{CT} , under uncertainty; and (2) computation of posterior predictive clinical and biomechanical indicators using a new multi-fidelity approach for uncertainty propagation that leverages non-linear dimensionality reduction to better correlate models and, in turn, increase estimator confidence. One significant contribution of this work is the development of an end-to-end pipeline for personalized quantification of coronary hemodynamics based on both routine and CT perfusion-based clinical measurements that is able to account for cardiac function as well as vessel-specific coronary flow. A second key contribution is the demonstration of multi-fidelity uncertainty propagation combined with non-linear dimensionality reduction to significantly improve prediction confidence for coronary hemodynamics simulations. The method presented here significantly extends previous efforts relying on linear dimension reduction strategies for multi-fidelity sampling [34, 35, 31, 36] by generalizing the recent work of Zanoni *et al.* [37, 38, 39], and investigates its application to personalized coronary hemodynamics. This framework is especially useful for applications where well-correlated low-fidelity models are unavailable, and consequently, the current state-of-the-art in multi-fidelity uncertainty quantification does not yield significant reductions in variance or computational cost. We will show that this is true for various important clinical and biomechanical hemodynamic metrics, and that the proposed framework leads to significantly lower computational cost in uncertainty-aware model predictions.

2. METHODS

In this section, we describe the framework developed in this work. We discuss the clinical data employed, the computational modeling techniques for coronary flows, the estimation of personalized model parameters, and finally uncertainty propagation for clinical and biomechanical QoIs. An overview of the entire pipeline is provided in Figures 1 and 2.

2.1. Clinical imaging and patient data

We constructed a patient-specific three-dimensional coronary artery anatomical model from coronary computed tomography angiography (CCTA) for a single patient with coronary artery disease who underwent CCTA and CT myocardial perfusion imaging (MPI_{CT}) at Stanford University School of Medicine, Stanford CA, USA. The imaging was part of an ongoing study (NCT03894423) that was approved by the Institutional Review Board at the Stanford University School of Medicine. Written informed consent was received prior to patient participation. Dynamic stress CT myocardial perfusion imaging (MPI_{CT}) and CCTA were performed on a third-generation dual-source CT scanner (SOMATOM Force, Siemens Healthineers) following a standard protocol [40, 19] with images acquired at systole and mid-diastole, respectively. For MPI_{CT} , reconstructed images with a 0.5×0.5 mm in-plane resolution, 3.0 mm slice thickness, and 2.0 mm reconstruction interval were processed on a dedicated workstation (Syngo.CT Myocardial Perfusion, Siemens Healthineers). For CCTA, the spatial resolution was approximately 0.5 mm in all three dimensions. More details on the imaging protocols and processing is available in [16]. Each patient also underwent an echocardiography study to evaluate their left ventricle (LV) ejection fraction and stroke volume. Finally, we collected cuff blood pressure measurements for each patient.

The CCTA and MPI_{CT} images were co-registered using an affine registration method in the open-source 3D *Slicer* (www.slicer.org) software. Details on the imaging and registration protocol are available in Menon *et al.* [16] and Khan *et al.* [41]. We then used the open-source *SimVascular* software package to segment and reconstruct a three-dimensional patient-specific model of the coronary artery anatomy from the CCTA imaging. The LV myocardium was segmented from the MPI_{CT} images using thresholding in *ParaView* (www.paraview.org). This produced co-registered three-dimensional models of the coronary tree and LV myocardium, as shown in Figure 1(b).

We then determined the flow corresponding to each coronary artery branch using the coronary tree anatomical model, segmented LV volume and measured myocardial blood flow (MBF) from MPI_{CT} . To compute the branch-specific coronary flows, we first performed a Voronoi tessellation on the LV volume, where each point in the LV was assigned to its closest coronary artery outlet. The assignment of each point in the LV to a coronary artery perfusing the LV resulted in the LV being divided into distinct non-overlapping territories. Each territory was associated with the closest coronary artery branch, which physiologically represents the LV perfusion territory of each coronary artery. This is shown in Figure 1(c), where the different colors on the LV denote distinct perfusion territories. Finally, since the MBF data from MPI_{CT} contained volumetric flow rates at each voxel in the LV, we determined the total volumetric flow rate within each perfusion territory by integrating the voxel MBF values within each territory. The

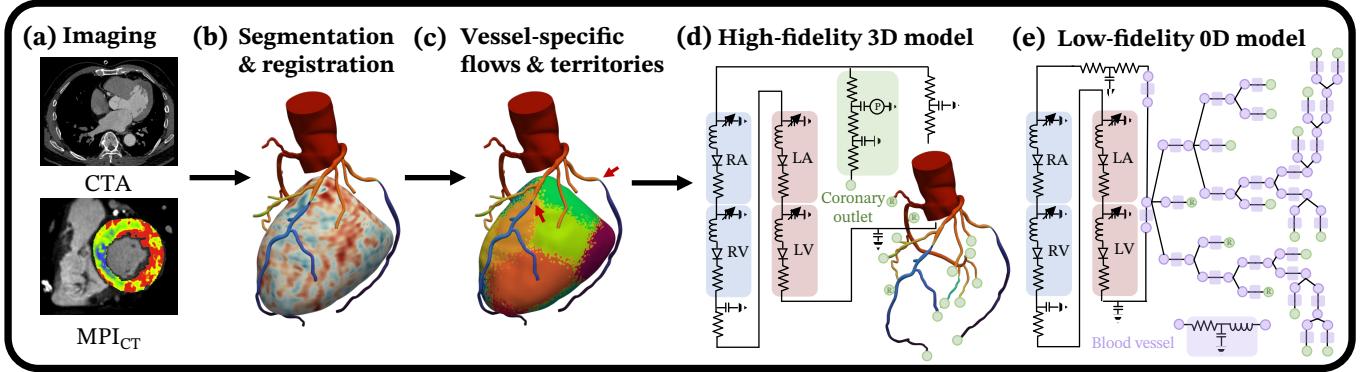


Figure 1: An overview of the image analysis and computational modeling pipeline. (a) Snapshots of the CTA and MPI_{CT} scans. (b) Segmentation and co-registration of the CTA and MPI_{CT} , resulting in three-dimensional models of the coronary tree and LV myocardium. (c) Visualization of the distinct LV perfusion territories corresponding to each coronary artery, shown in different colors on the LV surface. (d)-(e) The high- and low-fidelity computational models. The non-LV coronary outlets, where branch specific flows from MPI_{CT} were not available, are indicated with the letter 'R' within the green circles in (d)-(e).

flow rate associated with each LV territory thus yielded the flow rate through the corresponding coronary artery branch perfusing that territory.

Note that not all the coronary arteries perfuse the LV because some branches perfuse the right ventricle, septum, etc. Since MBF from MPI_{CT} was only collected in the LV, we only computed branch-specific flows for those branches perfusing the LV. The method used here for partitioning the LV into perfusion volumes is in line with several prior studies which have based LV perfusion territory selections on the unweighted or weighted distance of each point on the LV to either the closest coronary outlet or the coronary centerline [42, 43]. We showed in previous work that our personalization procedure is robust to the choice of method to compute branch-specific LV perfusion territories [16]. We note that this study patient was right-dominant and the arteries perfusing the LV were determined in consultation with an expert clinician. The model consisted of 17 coronary artery outlets, of which 14 outlets perfused the LV. The outlets that did not perfuse the LV are indicated in Figure 1(d).

2.2. Sources of uncertainty

As discussed above, we used several modalities of clinical data collected for the patient. While we consider only primary sources of uncertainty stemming from boundary condition assignments, the proposed framework allows the general inclusion of other sources of uncertainty.

We specifically focused on sources of noise directly related to the incorporation of MPI_{CT} into computational models of coronary flow. To account for coronary flow uncertainty from the imaging and analysis described above, we introduced three sources of noise into the procedure. One was in the computation of the LV perfusion territories corresponding to each coronary artery, where we introduced noise into the distance between each point on the LV and the coronary artery outlets. This accounts for anatomical uncertainty in reconstructing the locations of the arteries and the shape of the LV myocardium from the CT imaging using the segmentation and registration procedures outlined above. In our data, smaller coronary arteries were resolved by approximately 1-2 voxels in the imaging. Therefore, we chose a noise level such that the uncertainty in the distance between each point on the LV and the corresponding coronary artery outlets was approximately this size. We specified Gaussian noise with zero mean and a standard deviation of 10% of the baseline (deterministic) distance. Note that we did not consider anatomical uncertainty stemming from the segmentation of the coronary anatomy from the CTA. However, prior work has introduced methods to handle this additional source of uncertainty, which could be incorporated into future studies [44].

The second source of noise relates to the MBF data, and accounts for uncertainty in the underlying image acquisition and analysis that is used to compute MBF from dynamic contrast imaging of the LV [40]. This was specified at each voxel of the MBF data as Gaussian noise with zero mean and a standard deviation of 20% of the baseline (clinical) MBF at each voxel. Due to the absence of repeated clinical measurements of MPI_{CT} and its variability, this noise level was chosen based on consulting with the clinical expert on our research team, Dr. Nieman, about prior experience and unpublished clinical work on analyzing MPI_{CT} data.

Both the above sources of noise influence the relative distribution of flow amongst coronary arteries in the model – via the assignment of LV perfusion volumes to each coronary artery in the first case, and the MBF within each of those volumes in the second case. Therefore, we used these two sources of noise to generate uncertainty-aware distributions for the measured flow corresponding to each coronary branch. We computed statistics by simulating 2500 realizations of the branch-specific MBF and LV perfusion volume tessellation.

In addition to the above sources, we introduced a third source of uncertainty related to the *total* coronary flow. This was motivated by the challenge of accurately quantifying total coronary flow using MPI_{CT} [19]. Assuming that the coronary flow is in the range of 3% to 5% of the total cardiac output, we introduced an additional uncertain uniformly distributed parameter with limits computed to yield 3%-5% coronary-to-systemic flow splits. The systemic flow used here was based on the cardiac output from echocardiography. We also assumed the uncertainty in the total coronary flow was independent of the uncertainty in the relative flow between each coronary artery branch.

We describe how these three sources of uncertainty are incorporated into our modeling framework in sections 2.4 and 2.5. We also note that the framework demonstrated here does not depend on the specific noise levels chosen. Clinically-informed noise distributions could be easily incorporated into the current framework and would provide more realism.

2.3. Computational modeling

The multi-fidelity framework developed in this work relied on a combination of high-fidelity (HF) three-dimensional (3D) simulations of coronary artery flow combined with low-fidelity (LF) zero-dimensional (0D) flow models. For the three-dimensional patient-specific simulations, we used the *svSolver* flow solver, which is a part of the *SimVascular* software suite. *svSolver* combines a stabilized finite element spatial discretization using linear tetrahedral elements with a second-order accurate generalized- α time stepping [45, 46] scheme. The governing equations are the three-dimensional incompressible Navier-Stokes equations,

$$\nabla \cdot \mathbf{u} = 0 ; \rho \frac{\partial \mathbf{u}}{\partial t} + \rho \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \mu \nabla^2 \mathbf{u}, \quad (1)$$

where \mathbf{u} and p are the blood velocity and pressure, respectively. Blood was assumed to be Newtonian with density $\rho = 1.06 \text{ g/cm}^3$ and viscosity of $\mu = 0.04 \text{ dynes/cm}^2$. Note that non-Newtonian effects begin to appear for blood vessels with diameters below $300 \mu\text{m}$ [47], which was well below the size of the vessels modeled in this study. For simplicity, we treated coronary artery walls as rigid, motivated by previous work showing minimal effect of wall elasticity on coronary artery hemodynamics [48]. Mesh convergence was assessed using a similar coronary flow modeling setup in prior work [13]. Each 3D simulation consisted of 5 cardiac cycles, with 1000 time-steps for each cycle, requiring approximately 7 hours on 96 processors of a high performance computing cluster and approximately 1 hour for postprocessing on a single processor.

Zero-dimensional flow simulations were performed using *svZeroDSolver*, an open-source solver for lumped parameter hemodynamic simulation, provided with the *SimVascular* software suite. Zero-dimensional models approximate vascular hemodynamics using an electric circuit analogy, specifically, using resistors to model major viscous losses, capacitance to model vessel compliance, and inductance to represent blood flow inertia. The anatomical layout of the 0D models, consisting of vessel lengths, diameters and connectivity, mirrored the 3D anatomy. The resistance, capacitance, and inductance of each vessel was simple determined from the geometry of each vessel, by considering the flow through a cylinder of equivalent diameter and length. This was performed using an automated procedure to extract these anatomical details from the 3D model described in Pfaller *et al.* [49]. Each zero-dimensional simulation was run for 5 cardiac cycles, with 1000 time-step per cycle, and took approximately 10 seconds, including post-processing, on a single processor.

Boundary conditions at the inlet and outlet of the aorta and coronary outlets in both 3D and 0D simulations were prescribed through a closed-loop 0D representation of the systemic circulation and distal vascular resistance. Figure 1(d) shows a schematic of the 3D computational model, consisting of the 3D anatomy coupled with 0D closed-loop boundary conditions. Figure 1(e) shows a schematic of the corresponding 0D computational model. The simplified closed-loop circulation model consisted of the four heart chambers in the systemic circulation, labeled RA, RV, LA and LV in Figures 1(d)-(e), as well as pulmonary resistance. We used Windkessel boundary conditions at the aortic outlet [50] and modeled the four heart chambers in the systemic circulation using time-varying elastance. Coronary outflow boundary conditions included the effect of distal resistance as well as the intra-myocardial

pressure, which is essential to capture the diastolic nature of coronary flow [51]. The intramyocardial pressure was obtained by coupling the outlet of each coronary artery with the pressure within the ventricular chambers of the 0D boundary condition model. The venous sides of the coronary outflow boundary conditions as well as the Windkessel boundary condition at the outlet of the aorta were all connected back to the inlet of the right atrium, thus creating a closed-loop system. In Figures 1(d)-(e), the coronary outlets are highlighted using green circles.

Flow simulations were performed under hyperemic conditions by scaling the resistance distal to coronary arteries by 0.24 compared to their baseline at rest [52]. This was done to match the acquisition of MPI_{CT} under hyperemia. The estimation of the boundary condition parameters, based on matching clinical measurements, is described in Section 2.4.

2.4. Parameter estimation using Markov Chain Monte Carlo

The parameters of the closed-loop 0D boundary conditions were estimated so that the simulations recapitulated clinically measured patient-specific data. The parameters of the model included the total systemic and coronary resistance as well as capacitance. These parameters affect how flow leaving the heart is split between coronary flow and systemic flow, as well as the pulsatility of this flow. The model also has parameters related to the elastance and volume of each heart chamber, which determined the pulsatility, cardiac output, ventricular/atrial pressures, etc. Details of these parameters can be found in [13] and [16]. Finally, the distribution of the total coronary flow amongst branches in the coronary artery tree was determined by the distribution of total coronary resistance. The resistances at each coronary artery outlet therefore represented additional parameters that needed to be estimated.

Parameter estimation was performed in two stages. In the first stage, we estimated the total systemic and coronary resistance and capacitance, the elastance and volumes of the heart chambers, cardiac activation times and intramyocardial pressure so that the simulations produced satisfactory agreement with the clinically measured systolic and diastolic blood pressure cuff measurements, as well as LV stroke volume and ejection fraction from echocardiography. In addition, these patient-specific clinical measurements were augmented with literature-based population-averaged target data, including the waveforms for cardiac activation and coronary flow, as well as pulmonary pressure. This first stage of the estimation consisted of 36 parameters. We determined their (deterministic) point estimates using the 0D model and derivative-free Nelder-Mead optimization [53]. This step of the parameter estimation followed directly from previous work. Details on the parameters, the initial guess, and optimization setup are provided in [13] and [16], and an analysis of the sensitivity of the model to the parameters is in [24].

The second stage, which was the primary focus in this work, aimed to personalize the simulated branch-specific flow in each coronary artery based on the branch-specific flows determined from the MBF measured in the LV by MPI_{CT} . As mentioned earlier, the distribution of flow amongst the branches of the coronary artery tree is determined primarily by the distribution of outlet resistances amongst the branches. The estimation of these coronary artery outlet resistances was therefore the aim of the second stage of the parameter estimation.

As noted earlier, since not all coronary arteries perfuse the LV, we focus here only on those that do. For the other coronary artery outlets, such as those supplying the RV or septum, the distal resistance was specified based on Murray's law with an exponent of 2.6 [14, 15]. These were therefore not included in the parameter estimation. To account for the noise in the target branch-specific flows from MPI_{CT} (described in Section 2.1), the posterior distribution for the distal resistance at each coronary outlet perfusing the LV was determined using Bayesian estimation.

While keeping the 36 parameters optimized in the first stage fixed, we considered a set of outlet resistances $\mathbf{r} = \{r_1, r_2, \dots, r_{N_c}\}$, where N_c is the number of coronary artery outlets perfusing the LV. We generated samples from the joint probability density of \mathbf{r} , given a set of target flow rates, $\mathbf{f}_{\text{CT}} \in \mathbb{R}^{N_c}$ specified at each of the N_c outlets. In other words, the estimation yields a sample-based characterization of the posterior distribution, $p(\mathbf{r}|\mathbf{f}_{\text{CT}})$. As in the first stage of the optimization, we used the 0D model as a surrogate to map the input parameters, \mathbf{r} , to the simulated flows at the outlets of each coronary artery, $\mathbf{f} \in \mathbb{R}^{N_c}$.

According to Bayes theorem, the posterior distribution can be expressed as,

$$p(\mathbf{r}|\mathbf{f}_{\text{CT}}) = p(\mathbf{f}_{\text{CT}}|\mathbf{r})p(\mathbf{r})/p(\mathbf{f}_{\text{CT}}). \quad (2)$$

In Equation (2), the likelihood $p(\mathbf{f}_{\text{CT}}|\mathbf{r})$ quantifies the ability of a given set of boundary condition parameters, \mathbf{r} , to produce model outputs matching the measured targets, \mathbf{f}_{CT} , with uncertainty induced by a known probabilistic characterization of the measurement noise. Following the discussion on how noise was added to MBF data in

Section 2.1, each component of f_{CT} had a mean equal to the clinically measured MBF within the LV perfusion volume associated with a given coronary branch, and a standard deviation given by 20% of the mean. Interestingly, the uncertain target coronary flows computed through the process described above were characterized by a non-diagonal covariance matrix, Σ , which we estimated from 2500 realizations of the noisy MBF data. Note that the non-diagonal nature of the covariance comes primarily from the inter-dependence between adjacent LV perfusion territories. The likelihood is expressed as,

$$p(f_{CT}|\mathbf{r}) = \frac{1}{\sqrt{(2\pi)^{N_c} \det \Sigma}} \exp \left\{ -\frac{1}{2} [\mathbf{f}_{CT} - \mathbf{f}(\mathbf{r})]^T \Sigma^{-1} [\mathbf{f}_{CT} - \mathbf{f}(\mathbf{r})] \right\}, \quad (3)$$

where $\mathbf{f}(\mathbf{r})$ is the vector of simulated coronary outflows. Furthermore, we assumed the prior, $p(\mathbf{r})$ in Equation (2), to be uniformly distributed with a range $[0.5\hat{r}, 2.0\hat{r}]$, where \hat{r} is determined for each artery, or each component of \mathbf{r} , by simply distributing the total resistance distal to the LV-perfusing arteries (which was estimated in the first stage of parameter estimation) based on the relative target flows for each artery.

We then generated samples from the posterior distribution of personalized coronary outlet resistances, $p(\mathbf{r}|\mathbf{f}_{CT})$, using Markov Chain Monte Carlo. In particular, due to the slow convergence of the classical Metropolis-Hastings algorithm, we utilized a Differential Evolution Adaptive Metropolis (DREAM) sampler with adaptive proposal distribution [54]. This method leverages self-adaptive differential evolution and runs multiple parallel Markov chains, which enhances its computational performance and convergence assessment. In this work, the number of estimated parameters, corresponding to the number of coronary arteries perfusing the LV, was $N_c = 14$. These are indicated by the green circles at the coronary artery outlets in Figure 1(d)-(e). The green circles inscribed with 'R' are those that did not perfuse the LV and were therefore not included in the parameter estimation. We also performed a Sobol sensitivity analysis to verify that changes in all outlet resistances produced sensible variations in the corresponding outlet flows. Details on this are in Appendix B. As expected, the sensitivity of the outlet flows to outlet resistances was very high. We ran 24 parallel chains with a total of 10,000 generations each. We used the Gelman-Rubin statistic [55] with a threshold of 1.1 to assess convergence. We discarded 50% of the samples as burn-in samples, based on a convergence study. Finally, at each generation we rescaled the components in \mathbf{r} so that the total coronary resistance was preserved, consistent with the value determined at the first stage of estimation.

2.5. Multi-fidelity uncertainty propagation

The next step of the framework was estimating the posterior predictive distribution propagating the uncertain model parameters to relevant clinical and biomechanical QoIs. In other words, given a vector of uncertain model parameters, $\theta \in \mathbb{R}^{N_p}$, whose joint distribution was determined from the solution of the above inverse problem, we aimed to efficiently estimate the corresponding distribution for the QoI, $Q(\theta) : \mathbb{R}^{N_p} \rightarrow \mathbb{R}$.

In this work, $Q(\theta)$ represents a relevant output from a 3D patient-specific simulation of coronary blood flow, with random inputs $\theta = \{r_1, r_2, \dots, r_{N_c}, s\} = \{\mathbf{r}, s\}$. Here, the first N_c components are outlet resistances (see discussion in Section 2.4) having joint distribution equal to the posterior $p(\mathbf{r}|\mathbf{f}_{CT})$, and s is a scaling factor for the total coronary resistance which accounts for the uncertainty in quantifying total coronary flow, which is a documented challenge in MPICT [19, 17]. Assuming the coronary flow is in the range of 3% to 5% of the total cardiac output, we specify s to be uniformly distributed with $s \sim \mathcal{U}(0.7, 1.25)$. The limits were computed to yield 3%-5% coronary-to-systemic flow splits and s was assumed to be independent of \mathbf{r} . We denote the joint distribution of $\theta = \{\mathbf{r}, s\}$ by $p(\theta)$.

The simplest approach to quantify posterior predictive moments is through standard Monte Carlo sampling, where, for example, the mean of $Q(\theta)$ can be estimated as,

$$\hat{Q}_N^{MC} = \frac{1}{N} \sum_{i=1}^N Q(\theta_i). \quad (4)$$

Here, θ_i is a realization of $\theta \sim p(\theta)$, and N is the selected number of samples. This estimator is unbiased, but its variance scales as $\text{Var}[\hat{Q}^{MC}] \sim 1/N$, leading to a computationally intractable estimation for applications, such as the current one, where each simulation is computationally expensive. This motivates our use of recently proposed multi-fidelity Monte Carlo estimators.

2.5.1 Multi-fidelity Monte Carlo

In this section, we briefly introduce the multi-fidelity Monte Carlo (MFMC) formulation employed in this work. Although we limit our presentation to MFMC, for simplicity of exposure, we note here that the proposed strategy is

general and could be adopted with any of the existing multi-fidelity estimators recently introduced in literature; see, e.g., multi-level Monte Carlo [56, 57, 58], multi-index Monte Carlo [59, 60], multi-level multi-fidelity Monte Carlo [61, 62, 63], Approximate Control Variate [64, 65], and multi-level best linear unbiased estimators [66, 67, 68].

Multi-fidelity Monte Carlo estimators are designed to leverage the low computational cost of running a large number of cheap low-fidelity (LF) model evaluations to improve the variance reduction in estimating QoIs from expensive high-fidelity (HF) simulators [69, 30, 70]. Here, the high-fidelity and low-fidelity models are 3D and 0D coronary flow simulations, respectively, both of which are discussed in Section 2.3. We refer to QoIs computed by each of these models as $Q^{3D}(\theta)$ and $Q^{0D}(\theta)$, respectively. Note that in the standard multi-fidelity Monte Carlo formulation, the low- and high-fidelity models share the same parametrization, which we will discuss further in Section 2.5.2.

In the two-model case, i.e., where a single low-fidelity model is used as control variate, the MFMC estimator for the mean of $Q(\theta)$ takes the form [69, 30]

$$\hat{Q}_{N_{3D}, N_{0D}}^{MFMC} = \frac{1}{N_{3D}} \sum_{i=1}^{N_{3D}} Q^{3D}(\theta_i) + \alpha \left(\frac{1}{N_{0D}} \sum_{i=1}^{N_{0D}} Q^{0D}(\theta_i) - \frac{1}{N_{3D}} \sum_{i=1}^{N_{3D}} Q^{0D}(\theta_i) \right). \quad (5)$$

Here, N_{3D} and N_{0D} are the number of 3D and 0D simulations, respectively. In particular, $N_{0D} = N_{3D} + \Delta_{0D}$, where N_{3D} is referred to here as the pilot sample and Δ_{0D} are additional low-fidelity simulations evaluated at independently drawn samples of θ . It can be shown that \hat{Q}^{MFMC} is an unbiased estimator for $Q(\theta)$ owing to the telescoping sum on the right hand side of Equation (5) [30].

The coefficient α in Equation (5) is chosen to minimize the variance of the estimator and is given by [30]

$$\alpha = \frac{\text{Cov}(Q^{3D}, Q^{0D})}{\text{Var}[Q^{0D}]} \cdot \frac{1}{N_{3D}}. \quad (6)$$

With this value of α , the variance of the estimator is given by [30, Lemma 3.3]

$$\text{Var}[\hat{Q}_{N_{3D}, N_{0D}}^{MFMC}] = \frac{\text{Var}[Q^{3D}]}{N_{3D}} \left[1 - \left(1 - \frac{N_{3D}}{N_{0D}} \right) \rho^2 \right] = \text{Var}[\hat{Q}_{N_{3D}}^{MC}] \left[1 - \left(\frac{\Delta_{0D}}{N_{0D}} \right) \rho^2 \right], \quad (7)$$

where ρ is the Pearson's correlation coefficient between the low-fidelity and high-fidelity (0D and 3D) models. Note that this is the expression of the variance for general values of N_{3D} and N_{0D} , and one can compute optimal values that minimize the variance for a given computational budget [30]. However, in this work we used a smaller number of samples to maintain a tractable computational cost. This is discussed further in Section 3.3.

Equation (7) suggests a significant variance reduction compared to the standard Monte Carlo estimator ($\hat{Q}_{N_{3D}}^{MC}$) when $N_{0D} \gg N_{3D}$ and $\rho \approx 1$. In other words, having an inexpensive and well-correlated low-fidelity model is crucial to improve confidence in the estimates. However, this is challenging in applications where certain physical phenomena are under-resolved or even neglected by low-fidelity representations. In such a case, alternative strategies are required to construct well-correlated low-fidelity models, as discussed in the next section.

2.5.2 Non-linear dimensionality reduction with autoencoders and normalizing flows

In this section, we describe a new approach for effectively reducing variance in multi-fidelity Monte Carlo estimates. In particular, we modify the sampling of low-fidelity models, such that their resulting outputs are well-correlated with the corresponding high-fidelity model outputs. This is accomplished by determining an optimal shared space between high- and low-fidelity representations, and performing a data-driven re-parameterization of the low-fidelity model. This method achieves higher correlation with the high-fidelity response, and is also applicable to low- and high-fidelity models having dissimilar parametrizations, generalizing classical multi-fidelity estimators, e.g., [34, 35, 36, 37], or surrogate-based approaches, e.g., [71]. A schematic of the framework is provided in Figure 2.

We used autoencoders to identify *active* low-dimensional non-linear manifolds for each QoI. Autoencoders are a data-driven approach to finding low-dimensional representations of data using neural networks. Each autoencoder consisted of a dense (fully connected) encoder $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}^{d_r}$ and a dense decoder $\mathcal{D} : \mathbb{R}^{d_r} \rightarrow \mathbb{R}^d$. The encoder compresses the original d -dimensional data to a d_r -dimensional *latent-space* representation, and the decoder reconstructs the d -dimensional data from the latent representation.

Let $\theta_i^{3D} \sim \pi^{3D}$ and $\theta_i^{0D} \sim \pi^{0D}$ be samples from the high-fidelity and low-fidelity random inputs, respectively, with corresponding QoIs given by $Q^{3D}(\theta^{3D}) : \mathbb{R}^{d^{3D}} \rightarrow \mathbb{R}$ and $Q^{0D}(\theta^{0D}) : \mathbb{R}^{d^{0D}} \rightarrow \mathbb{R}$. We trained separate

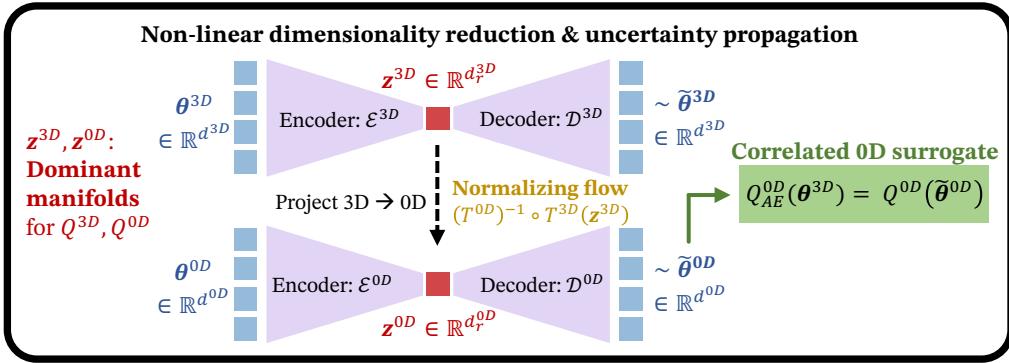


Figure 2: A schematic of the workflow for the non-linear dimensionality reduction with autoencoders and normalizing flows.

autoencoders, i.e. $\mathcal{E}^{3D} : \mathbb{R}^{d^{3D}} \rightarrow \mathbb{R}^{d_r^{3D}}$, $\mathcal{D}^{3D} : \mathbb{R}^{d_r^{3D}} \rightarrow \mathbb{R}^{d^{3D}}$ for the 3D model, and $\mathcal{E}^{0D} : \mathbb{R}^{d^{0D}} \rightarrow \mathbb{R}^{d_r^{0D}}$, $\mathcal{D}^{0D} : \mathbb{R}^{d_r^{0D}} \rightarrow \mathbb{R}^{d^{0D}}$ for the 0D model, with corresponding latent variables expressed as

$$z^{3D} = \mathcal{E}^{3D}(\theta^{3D}) \text{ and } z^{0D} = \mathcal{E}^{0D}(\theta^{0D}). \quad (8)$$

Note that, unlike classical unsupervised autoencoders, which discover manifolds of reduced dimensionality directly in the input space, we instead focus on determining a latent space capturing the full variation of a given QoI for both the low- and high-fidelity model response. Therefore, we minimized a loss function of the form (shown here only for the high-fidelity model for brevity),

$$\begin{aligned} \mathcal{L}^{3D}(\phi^{3D}) = & \text{MSE}[Q^{3D}(\theta^{3D}), Q_{NN}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D})] \\ & + \text{MSE}[Q^{3D}(\theta^{3D}), Q_{NN}^{3D} \circ \mathcal{E}^{3D} \circ \mathcal{D}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D})] \\ & + \text{MSE}[Q_{NN}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D}), Q_{NN}^{3D} \circ \mathcal{E}^{3D} \circ \mathcal{D}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D})] \\ & + \text{MSE}[\mathcal{D}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D}), \mathcal{D}^{3D} \circ \mathcal{E}^{3D} \circ \mathcal{D}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D})]. \end{aligned} \quad (9)$$

Here, ϕ^{3D} represents the trainable autoencoder parameters, and \circ indicates function composition. In addition, Q_{NN}^{3D} represents a fully-connected neural network surrogate of Q^{3D} with inputs in the latent space, trained simultaneously with the autoencoder. We used Q_{NN}^{3D} to circumvent the prohibitive cost of evaluating Q^{3D} every time the loss function is computed, while also providing a supervised loss to identify the latent variables describing the function's variability. This choice is motivated by prior work Ref. [39], though here we added the third term to reinforce the first two loss terms.

In order to map between latent spaces corresponding to the high- and low-fidelity models, we also needed to ensure that their respective latent space representations had the same probability distribution (so-called *shared space*). This was achieved by estimating their probability density through normalizing flows, which determine invertible transformations, \mathcal{T}^{3D} and \mathcal{T}^{0D} , from the *a priori* unknown latent space probability densities to a common standard Gaussian, $\mathcal{N}(0, 1)$.

We then evaluate the low-fidelity model at new sampling locations, which we expected to exhibit improved correlation with respect to Q^{3D} . To do so, every input realization for the high-fidelity model is mapped to its latent space and then the shared space. Each input is then mapped to the latent space of the low-fidelity model and finally to the input space of the low-fidelity model. Formally, this is expressed as,

$$Q_{AE}^{0D}(\theta^{3D}) = Q^{0D} \circ \mathcal{D}^{0D} \circ (\mathcal{T}^{0D})^{-1} \circ \mathcal{T}^{3D} \circ \mathcal{E}^{3D}(\theta^{3D}). \quad (10)$$

Intuitively, we sample the low- and high-fidelity models at the same locations within the *shared space*. Thus, by evaluating each model at equivalent locations within their respective influential manifolds (for a specific QoI), we expect well-correlated model outputs. Moreover, quantitative results have been shown under idealized assumptions in [36, Proposition 4.4] for the linear case, and in [39, Theorem 3.3] for the nonlinear approach presented here. Specifically, these references show that the new sampling locations for the low-fidelity model provided by the transformations in (10) yield a correlation with respect to the high-fidelity model that is always larger than the original correlation.

In this work, we used a one-dimensional latent space, i.e., $d_r^{3D} = d_r^{0D} = 1$. For this special case, the normalizing flow can be simplified. We can convert any distribution, i.e. the arbitrarily distributed z^{3D} and z^{0D} , to the standard Gaussian distribution using inverse transform sampling. Here, we used the empirical distribution function to first transform the latent space distribution to $\mathcal{U}(0, 1)$ and then to $\mathcal{N}(0, 1)$ by means of an invertible analytical map.

In addition, since we aim to increase the correlation between the low- and high-fidelity models by linking them through a common shared space, we performed a second stage of training for the autoencoders prior to the final construction of Q_{AE}^{0D} . Here, we combined the losses from low- and high-fidelity autoencoders, and added a loss term to boost the correlation, ρ , between the high-fidelity model evaluated at the original samples and the *resampled* low-fidelity model outputs. In practice, this second stage of training used a loss function of the form

$$\mathcal{L}(\boldsymbol{\phi}^{3D}, \boldsymbol{\phi}^{0D}) = \mathcal{L}^{3D} + \mathcal{L}^{0D} - \left| \rho [Q^{3D}, Q_{NN}^{0D} \circ (\mathcal{T}^{0D})^{-1} \circ \mathcal{T}^{3D} \circ \mathcal{E}^{3D}(\boldsymbol{\theta}^{3D})] \right|. \quad (11)$$

The neural network weights during this second stage were initialized with those determined from the first stage discussed above with loss function in Equation (9).

Using the modified low-fidelity model given by Equation (10), the proposed multi-fidelity estimator is expressed as

$$\hat{Q}_{N_{3D}, N_{0D}}^{MFMC, AE} = \frac{1}{N_{3D}} \sum_{i=1}^{N_{3D}} Q^{3D}(\boldsymbol{\theta}_i^{3D}) + \alpha \left(\frac{1}{N_{0D}} \sum_{i=1}^{N_{0D}} Q_{AE}^{0D}(\boldsymbol{\theta}_i^{3D}) - \frac{1}{N_{3D}} \sum_{i=1}^{N_{3D}} Q_{AE}^{0D}(\boldsymbol{\theta}_i^{3D}) \right). \quad (12)$$

The estimator $\hat{Q}^{MFMC, AE}$ is unbiased following the same reasoning as with the standard multi-fidelity Monte Carlo estimator in Equation (5) [30]. Moreover, this was demonstrated for various analytical examples in our previous work [37].

Note that while the framework developed here was inspired by Zanoni *et al.* [37, 39], we introduced several features here that significantly improved its performance. For one, while the previous work parameterized the normalizing flows using neural networks, our use of a one-dimensional latent space with the analytical transformations between probability distributions reduced the computational complexity and required dataset size. In addition, the inclusion of the correlation between low- and high-fidelity outputs in the loss, combined with the two-stage training of the autencoders, using Equations (9) and (11), significantly improved the correlation between the models as well as the reproducibility of the framework.

Details on the training of the neural networks, the hyper-parameters, and other aspects of the method are available in Appendix A.

3. RESULTS

We now discuss the performance of the proposed data-to-prediction pipeline, from the estimation of personalized parameters to the prediction of clinical and biomechanical QoIs under clinical data uncertainty.

3.1. Parameter estimation

As described in Section 2.4, our estimation consisted of a two-step procedure – a deterministic optimization-based step to assimilate patient-specific and literature-based measurements of cardiac function, and Bayesian estimation step to personalize vessel-specific coronary flows. Since the main focus of this work is on demonstrating the uncertainty-aware pipeline for personalized assessment of coronary hemodynamics, this section focuses on the Bayesian parameter estimation for vessel-specific coronary flows. We summarize the cardiac function estimation results in Figure B.1 of Appendix B.

Figure 3(a) shows the estimated marginals for the $N_c = 14$ -dimensional posterior distribution, $p(\mathbf{r})$. The unimodal character of all estimated marginals and their limited variance confirms that the estimation problem is well-posed, with identifiable resistance parameters.

In addition, Figure 3(b) shows a comparison for the clinically-measured target coronary flow in each artery, the assumed measurement noise, and the posterior predictive marginals. Our results (in blue) closely match both the mean target flows and their distribution induced by the noise model in Equation (3). We were also able to match the covariance of the target flows, as shown in Figure B.3 in Appendix B.

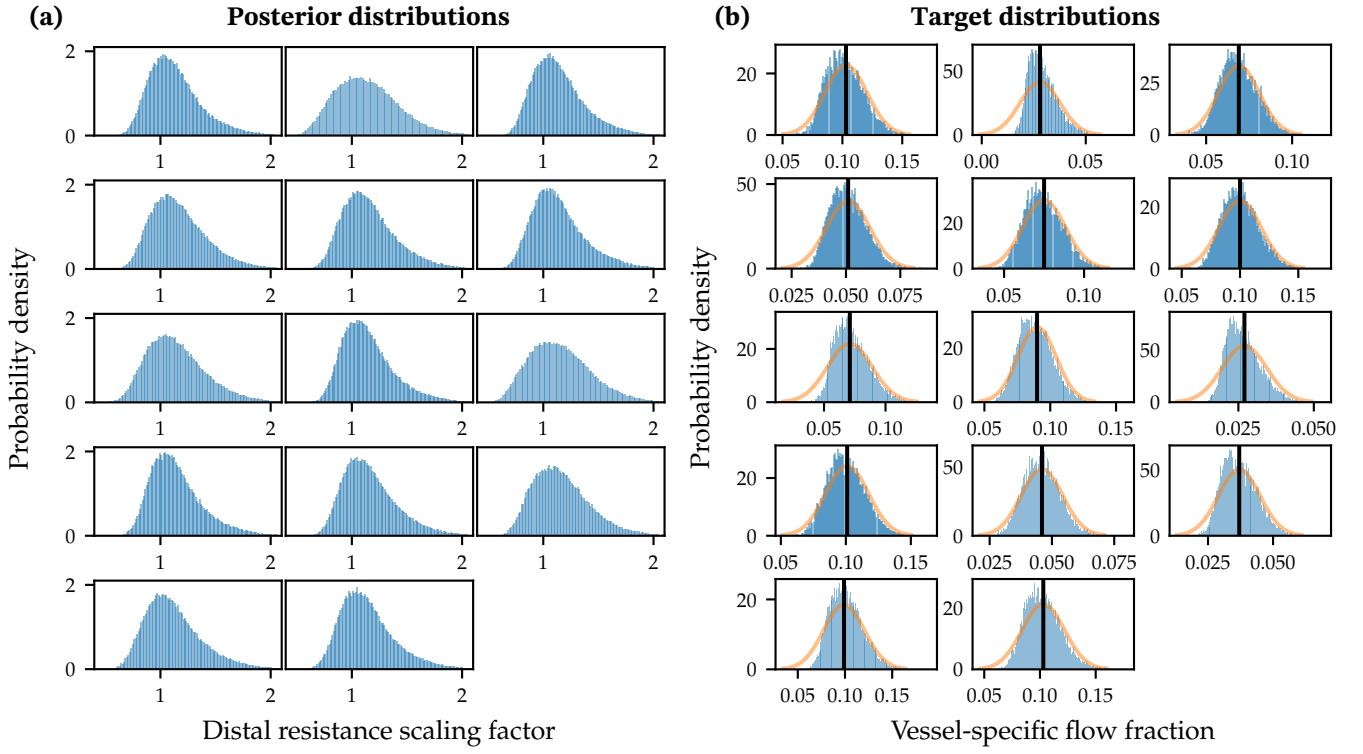


Figure 3: (a) Slices of the estimated posterior distributions of the distal resistance for all 14 coronary outlets. (b) Distributions of predicted coronary flows compared with clinically measured targets (vertical line) and noise distribution for each coronary artery.

3.2. Uncertainty propagation

We now discuss the estimation of various clinical and biomechanical QoIs with a focus on important metrics in the diagnosis and progression of CAD. We estimated these metrics in the three primary coronary artery branches - the left anterior descending (LAD), the left circumflex (LCx), and the right coronary artery (RCA). Furthermore, these quantities of interest were computed in the vicinity of stenoses present in each of these branches, indicated by red arrows in Figure 1(c). For each QoI, we compared the performance of the standard Monte Carlo (MC) estimator, the multi-fidelity Monte Carlo (MFMC) estimator based on the original low-fidelity model, and the proposed improved multi-fidelity Monte Carlo estimator with non-linear dimensionality reduction (MFMC-AE).

In this work, we used $N_{3D} = 500$ and $N_{0D} = 10,000$. Note that the estimated mean of each quantity of interest reported in this section was from one realization of each estimator. The confidence intervals were estimated using Chebyshev's inequality with the analytical variance, discussed in section 2.5, for each estimator. Therefore, while the estimators are unbiased, the offsets seen in the mean values are expected from a single realization. In addition, for this application, we do not necessarily need a computational budget as large as our chosen N_{3D} . We show the convergence of the method for smaller values of N_{3D} in Section 3.3.

3.2.1 Lumen wall shear stress

The biomechanical quantities of interest we report in this paper are the time-averaged wall shear stress (TAWSS) and oscillatory shear index (OSI). TAWSS is the time-average of the shear stress on the vessel wall at every point on the vessel wall and OSI is a metric between 0 and 0.5, which measures the degree of oscillation/direction-change in the shear-stress experienced by the vessel wall at a given point. Pathological values for both of these metrics have been correlated with the progression of vascular disease (see, e.g., [72]). We note that the 3D simulations used here assumed the coronary artery walls to be rigid, which may affect estimates of TAWSS and OSI.

From the 3D simulations, we computed the shear stress on the vessel wall as,

$$\tau = \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^T) \cdot \mathbf{n}|_{x=\text{wall}}, \quad (13)$$

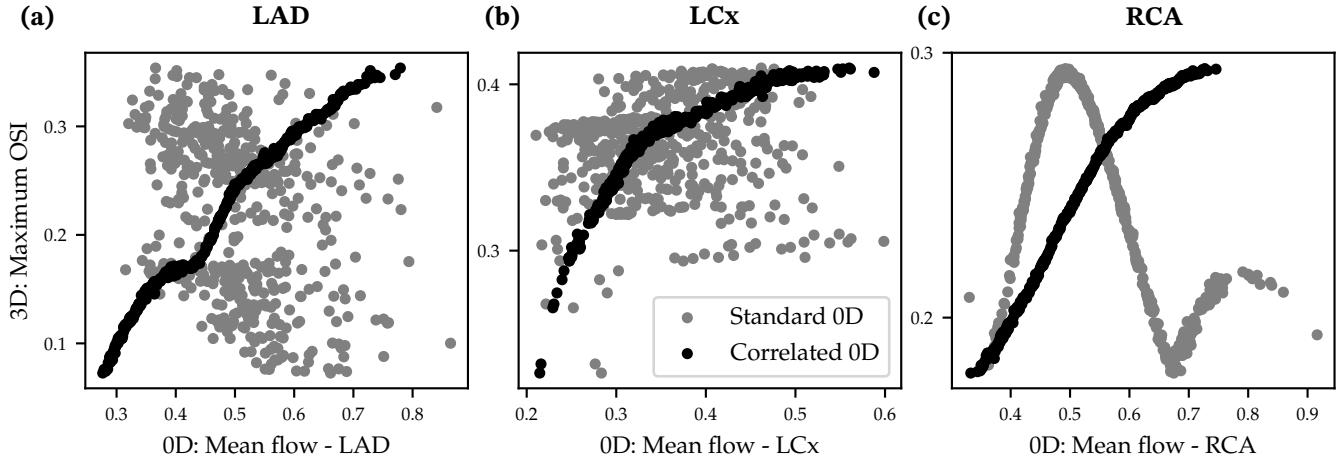


Figure 4: Correlation between maximum OSI computed from 3D simulations and mean outlet flow computed from 0D simulations. Estimates labeled as standard 0D models are 0D models evaluated using the same input parameters as the 3D models. Estimates labeled as correlated 0D models are from the MFMC-AE-reparameterized 0D models in equation (10). The three panels show the correlations for the LAD, LCx and RCA branches.

where \mathbf{n} is the unit vector normal to the vessel wall at any point, and the subscript “ $x = \text{wall}$ ” denotes evaluation at the lumen. TAWSS and OSI were then computed as,

$$\text{TAWSS} = \int_0^T \tau dt, \quad (14)$$

$$\text{OSI} = \frac{1}{2} \left(1 - \frac{|\int_0^T \tau dt|}{\int_0^T |\tau| dt} \right). \quad (15)$$

Here, T denotes the time-period of a cardiac cycle. We report the minimum and maximum TAWSS and OSI, respectively, within the stenotic region of each of the three primary coronary arteries.

Note that 0D low-fidelity models cannot capture spatially-resolved flow-fields, and therefore can only provide coarse approximations of TAWSS and OSI based on analytic solutions. Therefore, we used the mean flow at each outlet as the low-fidelity predictor (Q^{0D}) for both TAWSS and OSI. Figure 4 shows the correlation between the maximum OSI in the three primary coronary arteries evaluated from the 3D model versus the corresponding mean flow computed by the 0D model. Figure 4 shows 0D estimates using the same input parameters as the 3D model, labeled standard 0D, as well as 0D estimates using the reparameterized inputs obtained from the MFMC-AE framework, as in equation (10). The 0D model using the MFMC-AE framework produces remarkably higher correlations.

The Pearson correlation coefficient between the maximum OSI and the mean flow was -0.3635 for the LAD, 0.0998 for the LCx, and -0.5493 for the RCA. The current framework improved these correlations to 0.9906 for the LAD, 0.9018 for the LCx, and 0.9840 for the RCA. Error-bars with 95% and 99% confidence intervals for the maximum OSI in each of the three branches are shown in Figures 5(a)-(c) for the MC, MFMC and MFMC-AE estimators, respectively. The proposed MFMC-AE estimator achieves a roughly three-fold reduction in the variance.

For the case of minimum TAWSS, the Pearson correlation coefficient between the minimum TAWSS from 3D simulations (Q^{3D}) and the 0D-estimated mean flow (Q^{0D}) was 0.6294 for the LAD, 0.7622 for the LCx, and 0.7349 for the RCA. The current framework improved this to 0.9977 for the LAD, 0.9873 for the LCx, and 0.9979 for the RCA. Figures 5(d)-(f) show estimates of minimum TAWSS from the MC, MFMC, and MFMC-AE estimators for the three coronary branches. In this case too, we observed a three-fold improvement in the variance of the estimator.

3.2.2 Fractional flow reserve

Catheter-based measurement of fractional flow reserve (FFR) is the current clinical gold-standard for the assessment of functional CAD severity and decision-making regarding coronary interventions [22]. FFR is measured as the ratio of the pressure immediately distal to a lesion divided by the aortic pressure under hyperemia. Consistent with this

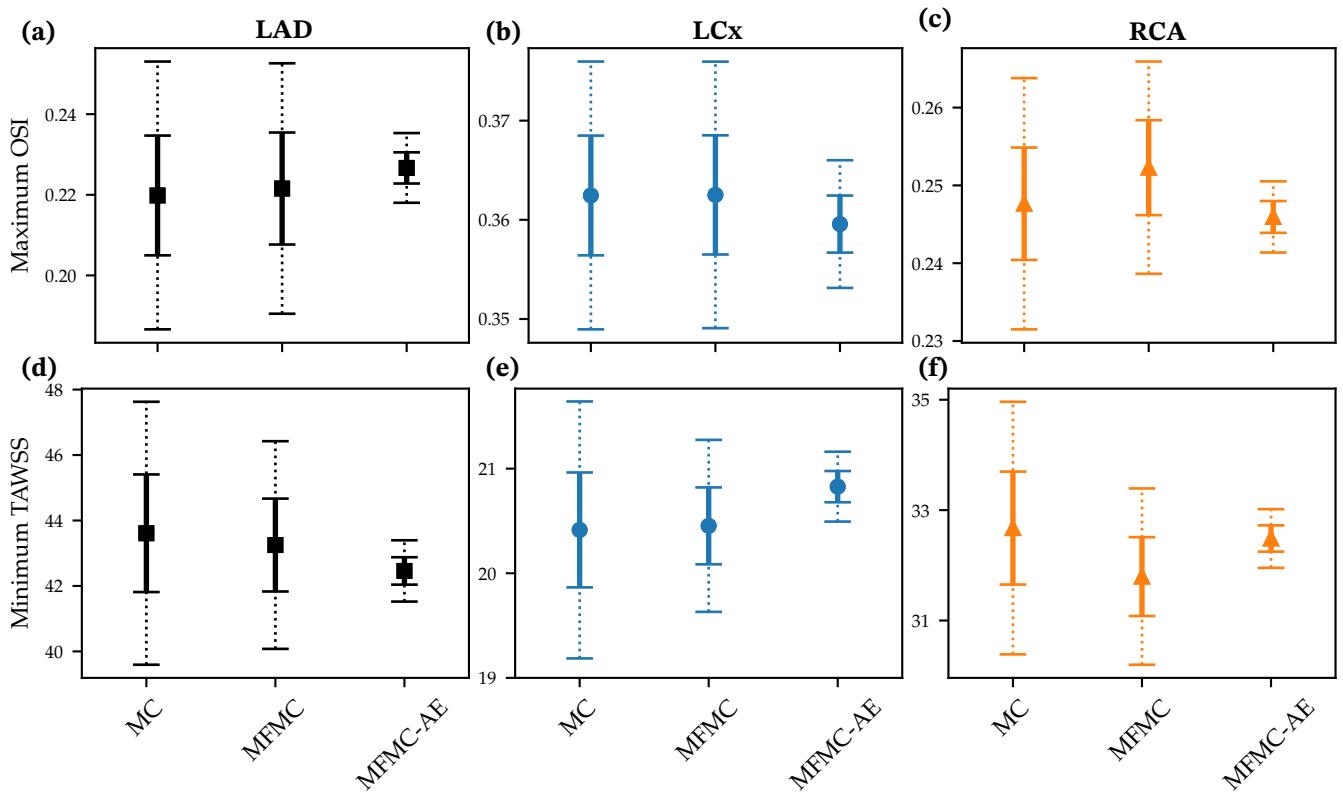


Figure 5: Maximum OSI in (a)-(c) and minimum TAWSS in (d)-(f) for LAD, LCx and RCA branches estimated using one realization of MC, MFMC and MFMC-AE estimators. Solid and dashed lines show 95% and 99% confidence intervals, respectively.

definition, we computed FFR from 3D simulations as the ratio of the cross-section-average pressure downstream of the stenotic region for each of the three primary vessels, divided by the aortic pressure.

The computation of FFR from 0D simulations is generally not possible in an analogous manner because these models do not include a spatial description of the flow and anatomy. Therefore, there is generally no output that is equivalent to the pressure at an anatomical location immediately distal to a stenosis. Moreover, these models usually assume that pressure loss along vessels is linear, inspired by Poiseuille flow. We included two specific features to address this. Each length of artery, separated by bifurcations, was modeled as a distinct vessel. This allowed us to isolate the location of each stenosis more precisely. We also modeled the non-linear pressure loss due to vascular stenoses as $\Delta P = S|Q|Q$, where ΔP is the pressure loss, Q is the flow rate, and S is a coefficient determined primarily by the degree of stenosis [49]. These feature allowed us to compute FFR from the 0D simulations in a similar manner to its computation from the 3D simulations. However, to demonstrate the efficacy of the current uncertainty quantification technique when using more standard (and widely used) 0D models that do not include these specialized features, we also report results where the low-fidelity estimator is simply the mean flow in each branch.

The Pearson correlation coefficient between FFR computed from the 3D and 0D simulations was 0.9951 for the LAD, 0.9954 for the LCx, and 0.9940 for the RCA. This high correlation is the result of the 0D model features described above. The current MFMC-AE framework was able to maintain similar correlations, specifically 0.9983 for the LAD, 0.9967 for the LCx, and 0.9868 for the RCA. Figures 6(a)-(c) show estimated FFR with 95% and 99% confidence intervals using MC, MFMC, and MFMC-AE estimators. We observed significant improvement in the confidence intervals with both MFMC and MFMC-AE, compared to MC.

When using the 0D mean flow as a low-fidelity predictor for FFR, the Pearson correlation between the 3D and 0D outputs dropped to -0.7612 for the LAD, -0.7820 for the LCx, and -0.9884 for the RCA. As expected, the mean flow was a less adequate low-fidelity surrogate for FFR, especially for the LAD and LCx. In this case, MFMC-AE was able to restore correlations to 0.9901 for the LAD, 0.9853 for the LCx, and 0.9495 for the RCA. The improved estimate for FFR in the LAD and LCx is evident in the confidence intervals shown in Figures 6(d)-(f).

A summary of all the computed quantities, the correlation between low and high-fidelity models, and the

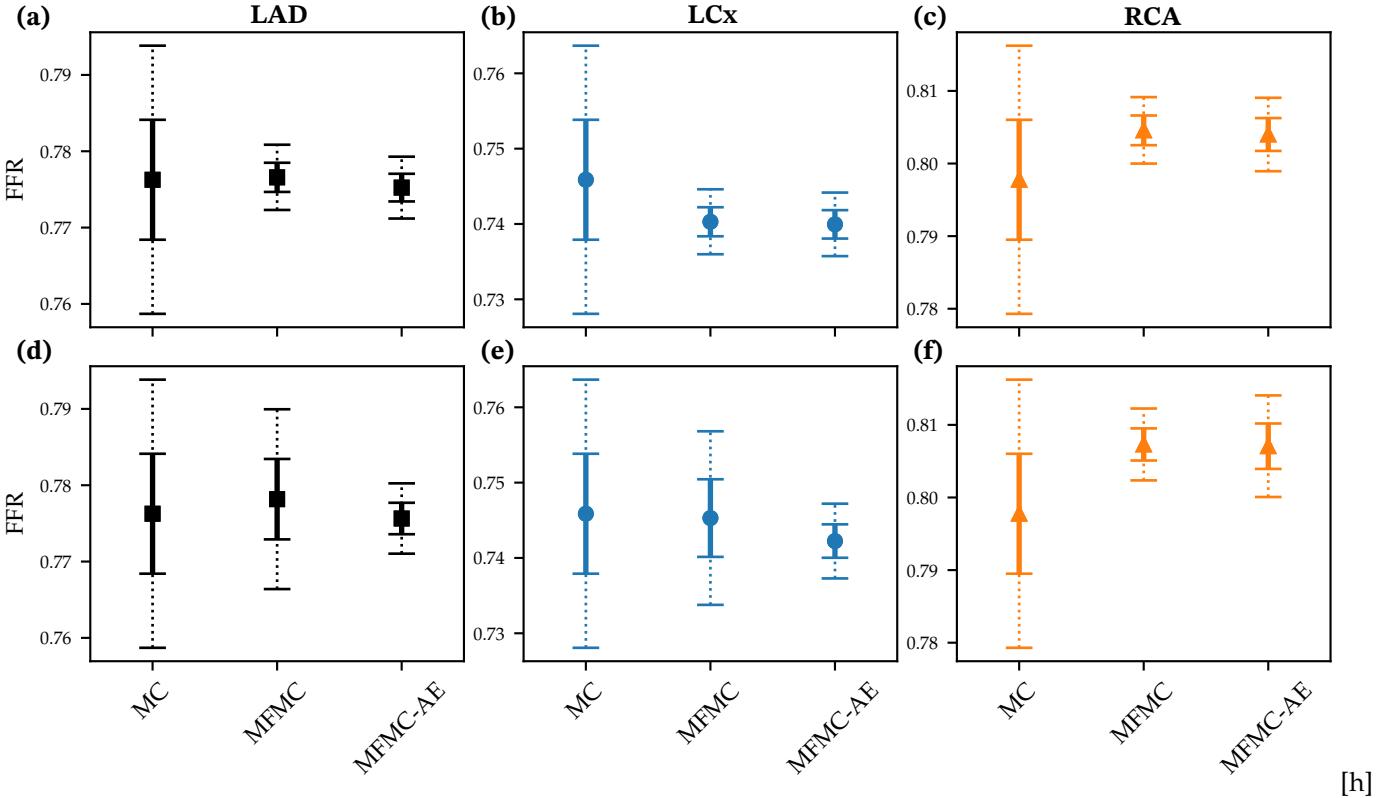


Figure 6: FFR estimated with 0D FFR in (a)-(c), and 0D mean flow in (d)-(f), using one realization of MC, MFMC and MFMC-AE. The three panels show data for the LAD, LCx and RCA branches. Solid and dashed lines show 95% and 99% confidence intervals.

standard deviation of the three estimators is provided in Table 1.

3.3. Computational budget

We now discuss the computational cost of the proposed framework, using the maximum OSI as an example QoI. We discuss the convergence of the framework with respect to the pilot sample N_{3D} , as well as the cost of estimating QoIs using the current MFMC-AE approach compared with standard MC and MFMC techniques.

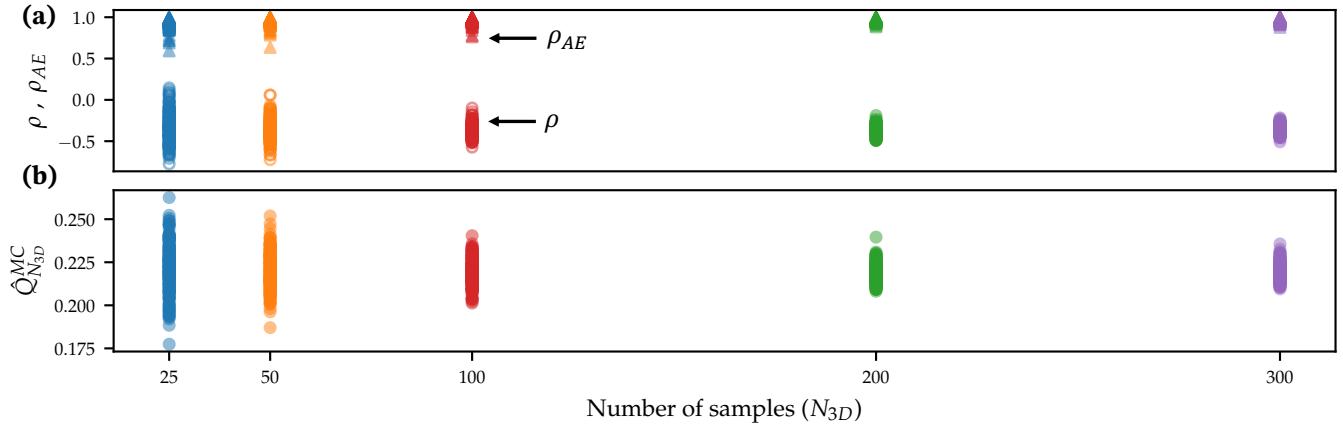
We begin with a discussion of the number of 3D samples required for convergence. As noted earlier, we used $N_{3D} = 500$ high-fidelity simulations in this work. To assess convergence for lower values of N_{3D} , we sub-sampled from the original $N_{3D} = 500$ pilot sample for the inputs and performed 200 independent trials of the MFMC-AE framework for each N_{3D} . Figure 7 shows the convergence of the framework for increasing N_{3D} in terms of the original correlation (ρ) between 3D and 0D outputs, the modified correlation (ρ_{AE}), and the Monte Carlo estimate of the mean ($\hat{Q}_{N_{3D}}^{MC}$). Remarkably, the framework is able to increase the correlation between 0D and 3D evaluations even for as low as $N_{3D} = 25$. However, the repeatability between independent trials significantly improves for increasing N_{3D} and converges for approximately $N_{3D} = 100$.

We now analyze the computational budget for the estimation of QoIs using the standard MC and MFMC techniques, compared to the current MFMC-AE framework. Note that the computational cost is reported in terms of equivalent HF simulations by dividing the total cost by the cost for a single 3D simulation (see Section 2.3).

Figure 8(a) compares the optimal/minimum variance of the MC, MFMC and MFMC-AE estimators as a function of the computational budget. For a given computational budget, the optimal variance was obtained following the work of Peherstorfer *et al.* [30]. The optimal number of high-fidelity and low-fidelity simulations for a given budget, \mathcal{B} , is,

$$N_{3D} = \frac{\mathcal{B}}{1 + w\gamma} \quad ; \quad N_{0D} = \gamma N_{3D} \quad ; \quad \gamma = \sqrt{\frac{\rho^2}{w(1 - \rho^2)}}, \quad (16)$$

Quantity of interest		Branch	Correlation		Standard deviation		
			Original	Modified	MC	MFMC	MFMC-AE
Max. OSI	Mean flow	LAD	-0.3635	0.9906	0.0033	0.0031	0.0009
		LCx	0.0998	0.9018	0.0013	0.0013	0.0006
		RCA	-0.5493	0.9840	0.0016	0.0014	0.0005
Min. TAWSS	Mean flow	LAD	0.6294	0.9977	0.4017	0.3173	0.0936
		LCx	0.7622	0.9873	0.1228	0.0822	0.0334
		RCA	0.7349	0.9979	0.2289	0.1597	0.0532
FFR	FFR	LAD	0.9951	0.9983	0.0018	0.0004	0.0004
		LCx	0.9954	0.9967	0.0018	0.0004	0.0004
		RCA	0.9940	0.9868	0.0018	0.0005	0.0005
FFR	Mean flow	LAD	-0.7612	0.9901	0.0018	0.0012	0.0005
		LCx	-0.7820	0.9853	0.0018	0.0012	0.0005
		RCA	-0.9884	0.9495	0.0018	0.0005	0.0007

Table 1: Summary of the computed quantities of interest.**Figure 7:** Convergence of MFMC-AE with respect to N_{3D} using maximum OSI as the QoI. Top panel shows original and modified correlations between 3D and 0D outputs. Bottom panel shows the Monte Carlo mean. Each marker is an independent trial.

where $w \ll 1$ represents the ratio of low-fidelity to high-fidelity model computational cost. The optimal variance for the MFMC and MFMC-AE estimators is given by using Equation (16) in Equation (7). Figure 8(a) shows a reduction in estimator variance of roughly $\mathcal{O}(10^{-2})$ resulting from the MFMC-AE estimator.

We note that in this work, $N_{0D} = 10,000$. This is approximately 50 times fewer low-fidelity evaluations than the optimal number suggested by Equation 16. We therefore also report the incremental value of increasing the current computational budget to reduce the variance of the current estimate of the maximum OSI (Figure 8(b)). We obtain an order of magnitude reduction in the estimator variance for $\sim 10\%$ of the cost of a single three-dimensional simulation. On the other hand, the MC and MFMC methods would require over 10^3 additional 3D simulations.

Since computational costs are described in terms of equivalent HF simulations, we reiterate that the cost of a single HF simulation was approximately 7 hours on 96 processors, plus 1 hour on a single processor for post-processing. Therefore, each HF simulation accounted for ~ 673 CPU hours. The $N_{3D} = 500$ simulations, which we showed in this section was vastly more than the required budget, accounted for $\sim 336,500$ CPU hours. In contrast, the $N_{0D} = 10,000$ LF simulations accounted for ~ 28 CPU hours for each QoI.

4. DISCUSSION

This paper provides a proof-of-concept demonstration of a novel uncertainty-aware framework for simulating patient-specific and vessel-specific coronary artery flows, informed by non-invasive MPI_{CT} and routine measurements of cardiac function. The main outcomes of this work were two-fold. First, we developed an end-to-end pipeline to

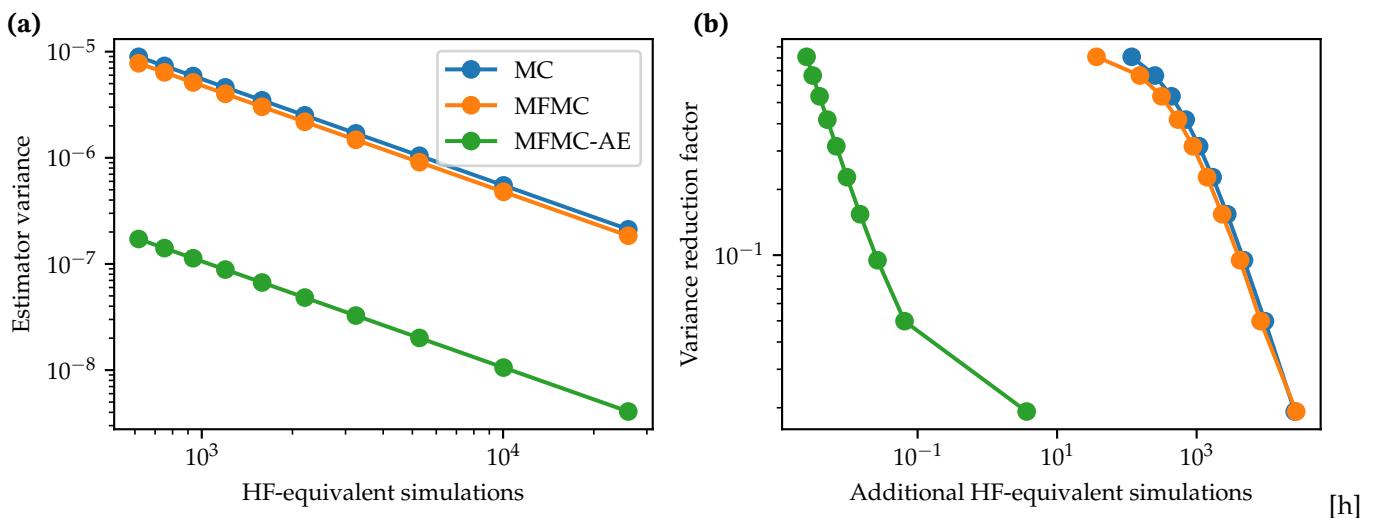


Figure 8: Computational budget analysis using maximum OSI as an example QoI. (a) Optimal variance reduction versus computational cost. (b) Additional computational budget required to reduce current (non-optimal) estimator variance.

estimate model parameters with the ability to match measured cardiac function and coronary flows, and to predict clinical and biomechanical quantities under this data-informed uncertainty. Second, we demonstrated significantly reduced variance and computational cost in the prediction of QoIs using a novel strategy for multi-fidelity uncertainty propagation relying on non-linear dimensionality reduction.

4.1. Personalized coronary flows

This study is motivated by prior work by Menon *et al.* [16] which highlighted the need for personalized vessel-specific flows in modeling coronary artery hemodynamics. Their work showed significant differences in coronary flows as well as FFR predicted by models that distribute the flow amongst coronary arteries based on patient-specific myocardial blood flow distributions from MPI_{CT} versus empirical diameter-based rules. The incorporation of MPI_{CT} into simulations of coronary flow is especially promising for clinical translation because current image-based coronary flow models are also based on CT imaging. Therefore, the high-resolution quantification of myocardial blood flow and full coverage of the myocardium from MPI_{CT} can potentially be incorporated into existing CT-based imaging protocols – streamlining the combination of these imaging techniques to develop more personalized hemodynamics simulations. In fact, the quantification of patient-specific hemodynamics through either simulations or MPI_{CT} have already been shown to provide improved clinical value in the diagnosis and treatment of CAD compared to traditional anatomical imaging via CT or invasive angiography alone [73, 74, 2]. Therefore, we expect increased personalization via the combination of novel imaging and modeling techniques to improve the accuracy and clinical value of these tools.

The inclusion of hemodynamic measurements in clinical CAD risk stratification is especially valuable in lesions that have borderline severity [75, 76]. Therefore, for the robust clinical translation of these simulation-based techniques, it is imperative to assess confidence in the predicted clinical risk. To that end, we extended our previous work on deterministic coronary flow personalization [16] to account for uncertainty in the clinical data.

4.2. Parameter estimation and multi-fidelity uncertainty quantification

While previous work involving uncertainty quantification for cardiovascular hemodynamics has either focused on the estimation of uncertain model parameters or the prediction of uncertain quantities of interest, here we demonstrated an end-to-end pipeline that goes from the clinical imaging to predicted clinical and biomechanical QoIs. For the estimation of personalized model parameters, we combined deterministic optimization for the parameters governing cardiac function with Bayesian estimation for the parameters that dictated vessel-specific coronary flows. Using a DREAM MCMC sampler [54], combining parallel Markov chains with differential evolution and adaptive proposal distributions, we demonstrated the ability to sample from a relatively high-dimensional

posterior distribution with 14 input parameters. This method is also known to show improvement over other MCMC techniques when sampling from heavy-tailed or multi-modal posterior distributions.

We also demonstrated a novel framework to significantly improve the confidence in predictions from simulations that incorporate noisy clinical data. While multi-fidelity Monte Carlo estimators are a mainstay for applications which can leverage computationally inexpensive low-fidelity surrogate models to reduce the variance of quantities estimated by expensive high-fidelity simulations [30], they rely on the availability of correlated low-fidelity surrogates. However, this is not always the case – especially for applications that aim to capture non-linear physics that is difficult to represent in low-fidelity models. We highlighted this deficiency in the estimation of TAWSS and OSI, which are governed by non-linear flow physics in the vicinity of the stenosis where their effect is most useful to estimate.

Using non-linear dimensionality reduction, we modified the sampling for existing low-fidelity 0D models to significantly improve their correlation with respect to high-fidelity 3D models. This was achieved by creating a *shared space* between the models – a one-dimensional non-linear manifold in the input space of each model that could capture most of the variability in the output QoI, and where the one-dimensional representation of the inputs for each model had the same distribution (the standard Gaussian). We then transformed each parameter vector for the 3D model to the shared space, and reconstructed the corresponding 0D input vector for each 0D evaluation. Intuitively, we ensured that the low-fidelity model was evaluated at the same locations as the high-fidelity model along the axes which are most influential for each of them, thus increasing their correlation.

Through the novel combination of multi-fidelity Monte Carlo with non-linear dimensionality reduction demonstrated here, we showed significant improvements in the confidence intervals for our predictions for both biomechanical as well as clinical quantities of interest. We also showed that for cases when we already have correlated low-fidelity surrogate models, such as for the FFR computation which included a model for the non-linear pressure drop within stenoses, our method preserves the effectiveness of the standard multi-fidelity Monte Carlo technique. Finally, we achieved significant variance reduction in the multi-fidelity Monte Carlo estimator with orders of magnitude reduction in the computational cost. This was demonstrated for the optimal number of simulations as well as with a more practical smaller computational budget.

4.3. Limitations

While this study demonstrates a promising direction in uncertainty-aware modeling of personalized coronary artery flows, there are several limitations and future directions that deserve attention. Importantly, while the current work represents a proof-of-concept, the methods presented here should be evaluated in large patient populations to demonstrate their clinical value.

We only considered the effect of uncertainty in MPI_{CT} , and to some extent, the image-based anatomical reconstruction. While the main focus of this work was on vessel-specific coronary flows, model personalization will be influenced by uncertainty in all the other clinical measurements utilized. These include features of the image-based anatomical models such as vessel lengths, diameters, and the number of branches in the coronary tree. Indeed, this significant source of uncertainty has been examined in recent work [44] and could be incorporated into future studies using the current framework. The boundary conditions also depend on the clinically measured cardiac output, stroke volume, and other measures of cardiac function determined by echocardiography. Finally, the modeling itself involves assumptions relating to the behavior of vessel walls, properties of blood, and the 0D modeling of boundary conditions. A particular aspect of this is the rigid wall assumption adopted in this study, which is known to influence the computation of TAWSS and OSI. In addition, the computation of the myocardial perfusion territories corresponding to each coronary artery was purely based on the distance of each point in the myocardium to the closest coronary artery, following previous work [42, 43]. However, the actual perfusion territories likely depend on factors influenced by coronary microvasculature that were not included in our modeling. Future work should evaluate the sensitivity of the computational modeling to these myriad sources of uncertainty in the clinical data and imaging to identify the most influential effects. While this work assumed the availability of certain clinical data, future work should account for potentially missing or gappy clinical data, as well as the impact of population-based parameters when modeling cases with disease.

Moreover, due to the absence of repeated measurements, this work used Gaussian noise to simulate clinical uncertainty. In the context of MPI_{CT} , it is unrealistic to obtain repeated measurements on the same patient due to clinical constraints and concerns related to the radiation dose for each scan. Therefore, repeated measurements to obtain realistic noise distributions would have to rely on animal or phantom studies. On the other hand,

realistic noise distributions in the anatomical reconstructions could be obtained by performing inter- and intra-operator variability studies [44]. However, since the aim of this work was to demonstrate the utility of the proposed framework for uncertainty-aware coronary flow models, this was considered beyond the scope of the study. Nevertheless, robust clinical applications would benefit from realistic clinically evaluated variability in the measurements.

The inclusion of additional sources of uncertainty would likely lead to very high-dimensional stochastic inputs. Extensions of this work would therefore benefit from recent developments in the area of Bayesian estimation, such as data-driven simulation-based inference techniques which combine variational inference with identifiability and sensitivity analysis for high-dimensional problems [77, 78]. In addition, while we used the zero-dimensional model as a surrogate for the parameter estimation, as has been done in previous work [24, 13, 16], multi-fidelity approaches would improve the accuracy of the parameter estimation [26]. Moreover, we assumed uniform priors for the estimation of parameter distributions. However, more informative priors would likely result in improved performance and reduced variance in the estimation of parameters.

Finally, although not demonstrated here, our novel multi-fidelity Monte Carlo technique allows the use of high-fidelity and low-fidelity models which have dissimilar input parameters [37]. This is enabled by the construction of the shared reduced-dimensional space between the models. Therefore, future work should explore promising avenues to further reduce computational cost by utilizing pre-computed simulation libraries with dissimilar inputs as surrogate models for relevant quantities of interest. Another future direction in this context is the effect of input parameter sensitivity and identifiability on the performance of the proposed multi-fidelity uncertainty propagation pipeline.

5. CONCLUSIONS

We demonstrated an end-to-end pipeline for personalized prediction of coronary hemodynamics under clinical data uncertainty. We utilized novel and routine clinical measurements – of myocardial blood flow from CT myocardial perfusion imaging and cardiac function from echocardiography – to personalize computational models of coronary hemodynamics in terms of gross hemodynamics as well as vessel-specific coronary flows. To account for uncertainty in the clinical data, Bayesian estimation was performed on personalized model parameters using adaptive Markov chain Monte Carlo. By combining multi-fidelity Monte Carlo uncertainty propagation with data-driven non-linear dimensionality reduction, we determined clinical and biomechanical posterior predictive QoIs, under uncertainty stemming from the clinical measurements. We demonstrated significantly improved correlations with respect to reference multi-fidelity estimators. For relevant quantities of interest, we showed that improved correlations greatly reduces estimator variance compared to standard multi-fidelity estimators, offering orders-of-magnitude computational cost savings.

6. DATA AVAILABILITY

All computational models built for this study have been anonymized and made publicly available through the Vascular Model Repository (<https://vascularmodel.com>).

7. CODE AVAILABILITY

The patient-specific modeling pipeline and computational fluid dynamics solvers are part of the SimVascular open-source project (<https://simvascular.github.io/>). Image analysis and postprocessing was performed using the open-source tools specified in section 2.1. The software to perform the non-linear dimensionality reduction and uncertainty quantification is available on Github at <https://github.com/StanfordCBCL/NeurAM>.

8. ACKNOWLEDGMENTS

This work was supported by NIH grant R01HL141712, NSF CAREER award 1942662, and NSF CDS&E awards 2104831 and 2105345. High performance computing resources were provided by the Stanford Research Computing Center. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC (NTESS), a wholly owned subsidiary of Honeywell International Inc.,

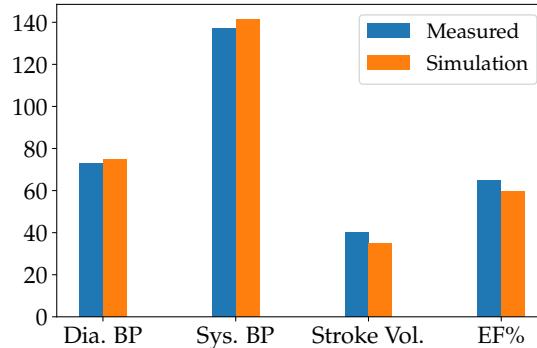


Figure B.1: Comparison of measured versus simulated metrics of cardiac function after model personalization. Dia. BP: Diastolic blood pressure. Sys. BP: Systolic blood pressure. Stroke Vol.: Stroke volume. EF: Ejection fraction.

for the U.S. Department of Energy’s National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525. This written work is authored by an employee of NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan.

APPENDIX

A. NEURAL NETWORK ARCHITECTURE AND HYPERPARAMETERS

The neural networks for both the autoencoders ($\mathcal{E}^{3D}, \mathcal{D}^{3D}, \mathcal{E}^{0D}, \mathcal{D}^{0D}$) were parameterized as fully-connected neural networks. The encoders and decoders both had 5 layers with 13 neurons each. The first and fifth layers were linear layers, and the other layers had hyperbolic tangent activation functions. The decoder had an additional sigmoidal layer as its last layer to ensure that the reconstructed inputs lay within the input parameter space.

The neural network surrogates for both the high- and low-fidelity quantities of interest, Q_{NN}^{3D} and Q_{NN}^{0D} , were fully connected networks with 3 layers and 10 neurons each. The first and last layer were linear, while the middle layer used a hyperbolic tangent activation.

All the neural networks were trained using the Adam optimizer [79] and an exponential scheduler for the learning rate. Hyper-parameters were tuned using the *Optuna* [80].

B. DETAILS ON PARAMETER ESTIMATION

As a first step of the parameter estimation process we used deterministic optimization to tune the parameters of the closed-loop lumped-parameter network. We used gradient-free Nelder-Mead optimization, with the primary targets being measurements of cardiac function from echocardiography and blood pressure cuff measurements. Further details on this parameter tuning are available in Menon *et al.* [16]. Figure B.1 shows the results of this tuning for the patient-specific data used in this study. We see that we successfully recapitulated the measurements of blood pressure as well as stroke volume and ejection fraction.

The second stage focused on estimating the distal resistance at each coronary outlet to match the target branch-specific coronary flows derived from MPI_{CT} . Prior to the parameter estimation, we first assessed the sensitivity of the outputs, i.e. the flow at each coronary outlet, to the input parameters, i.e. the distal resistance at each coronary outlet. The results from the global sensitivity analysis, in terms of the total-effect Sobol indices, are shown in Figure B.2(a). As expected, the flow at each outlet is very sensitive to the resistance at the corresponding outlet. Figure B.2(b) shows the logarithm of the sensitivity indices.

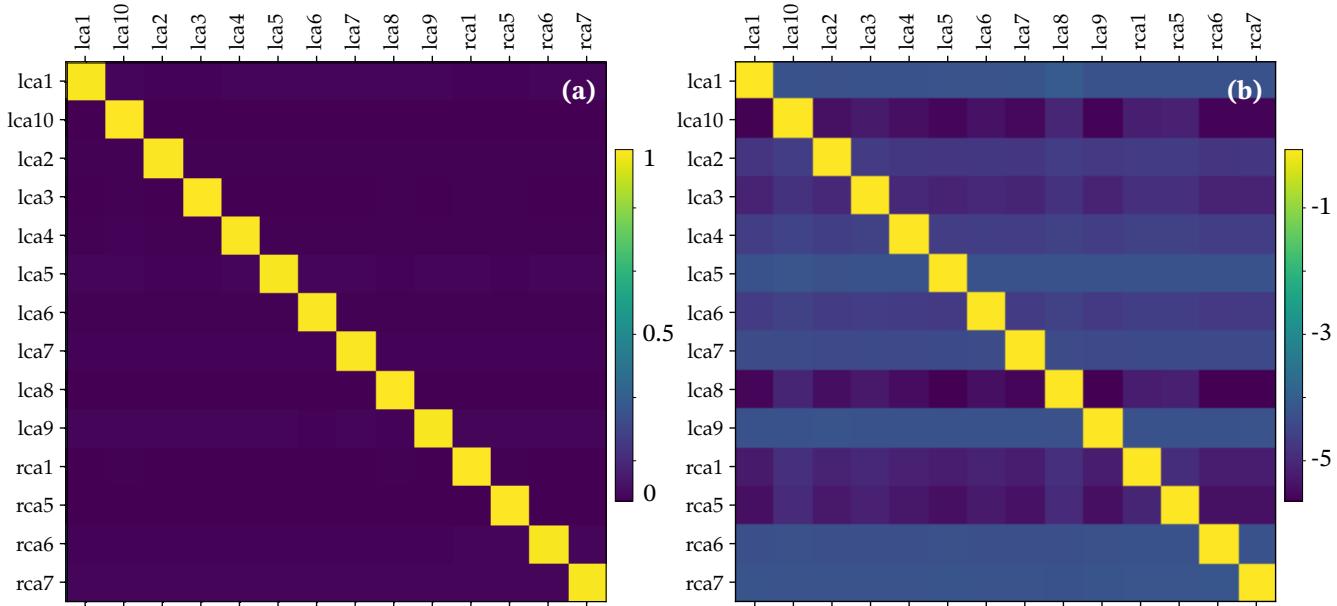


Figure B.2: (a) Total-effect Sobol sensitivity indices, with inputs (resistance at each coronary outlet) on the horizontal axis and the outputs (flow at each coronary outlet) on the vertical axis. (b) Logarithm of sensitivity indices.

When extracting the clinical targets for the second stage of the parameter estimation, i.e. the branch-specific coronary flows with simulated Gaussian noise, we also obtained the covariance of the targets (Σ). This covariance was a result of the interdependence amongst nearby branches when partitioning of the LV myocardium into non-overlapping perfusion territories corresponding to each coronary artery branch. The covariance was included in the Bayesian parameter estimation, as shown in the likelihood function in Equation (3). For each pair of coronary branches, Figure B.3 compares the predicted covariance from the parameter estimation, shown as scatter plots, with the target covariance, shown as ellipses. We see that the parameter estimation was able to capture the target covariance reasonably well.

REFERENCES

- [1] Karthik Menon, Zinan Hu, and Alison L Marsden. Cardiovascular fluid dynamics: a journey through our circulation. *Flow*, 4:E7, 2024.
- [2] Charles A Taylor, Kersten Petersen, Nan Xiao, Matthew Sinclair, Ying Bai, Sabrina R Lynch, Adam UpdePac, and Michiel Schaap. Patient-specific modeling of blood flow in the coronary arteries. *Computer Methods in Applied Mechanics and Engineering*, 417:116414, 2023.
- [3] Bon Kwon Koo, Andrejs Erglis, Joon Hyung Doh, David V. Daniels, Sanda Jegere, Hyo Soo Kim, Allison Dunning, Tony Defrance, Alexandra Lansky, Jonathan Leipsic, and James K. Min. Diagnosis of ischemia-causing coronary stenoses by noninvasive fractional flow reserve computed from coronary computed tomographic angiograms: Results from the prospective multicenter DISCOVER-FLOW (Diagnosis of Ischemia-Causing Stenoses Obtained Via Noni). *Journal of the American College of Cardiology*, 58(19):1989–1997, 2011.
- [4] James K. Min, Jonathon Leipsic, Michael J. Pencina, Daniel S. Berman, Bon Kwon Koo, Carlos Van Mieghem, Andrejs Erglis, Fay Y. Lin, Allison M. Dunning, Patricia Apruzzese, Matthew J. Budoff, Jason H. Cole, Farouc A. Jaffer, Martin B. Leon, Jennifer Malpeso, G. B. John Mancini, Seung Jung Park, Robert S. Schwartz, Leslee J. Shaw, and Laura Mauri. Diagnostic accuracy of fractional flow reserve from anatomic CT angiography. *JAMA - Journal of the American Medical Association*, 308(12):1237–1245, 2012.
- [5] Júlia Karády, Thomas Mayrhofer, Alexander Ivanov, Borek Foldyna, Michael T. Lu, Maros Ferencik, Amit Pursnani, Michael Salerno, James E. Udelson, Daniel B. Mark, Pamela S. Douglas, and Udo Hoffmann. Cost-

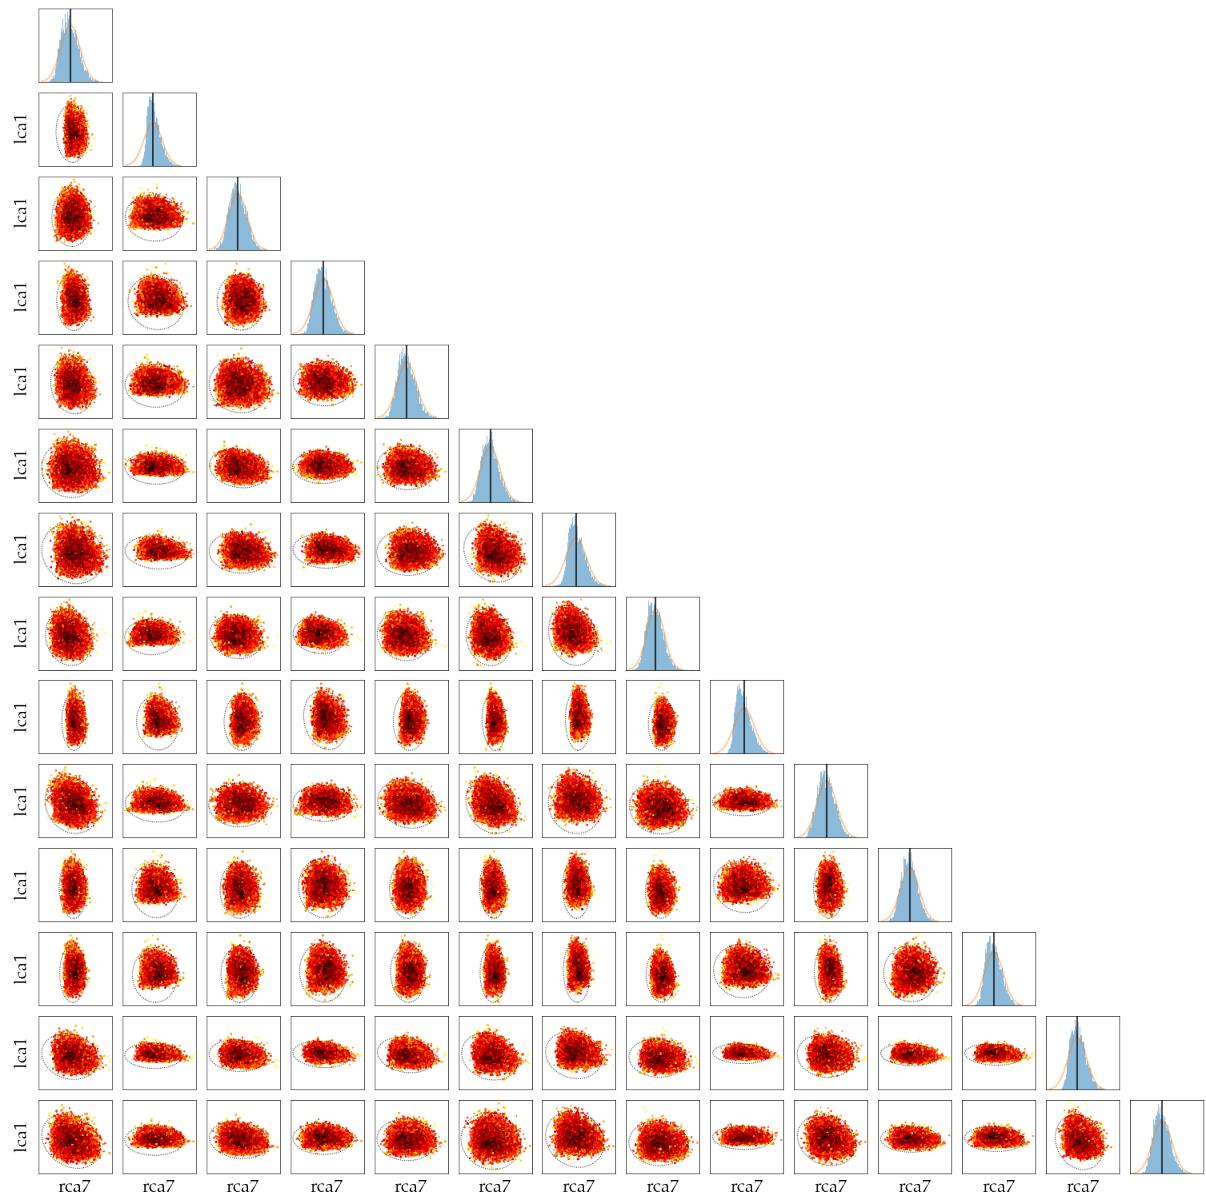


Figure B.3: Distributions of simulated (scatter) versus measured (dashed ellipse) covariance between branch-specific coronary flows resulting from Bayesian parameters estimation.

effectiveness Analysis of Anatomic vs Functional Index Testing in Patients With Low-Risk Stable Chest Pain. *JAMA Network Open*, 3(12):e2028312–e2028312, 12 2020.

- [6] Umberto Morbiducci, Raffaele Ponzini, Mauro Grigioni, and Alberto Redaelli. Helical flow as fluid dynamic signature for atherogenesis risk in aortocoronary bypass. A numeric study. *Journal of Biomechanics*, 40(3):519–534, 1 2007.
- [7] Muhammad Owais Khan, Justin S. Tran, Han Zhu, Jack Boyd, René R. Sevag Packard, Ronald P. Karlsberg, Andrew M. Kahn, and Alison L. Marsden. Low Wall Shear Stress Is Associated with Saphenous Vein Graft Stenosis in Patients with Coronary Artery Bypass Grafting. *Journal of Cardiovascular Translational Research*, 2020.
- [8] Abhay B. Ramachandra, Sethuraman Sankaran, Jay D. Humphrey, and Alison L. Marsden. Computational Simulation of the Adaptive Capacity of Vein Grafts in Response to Increased Pressure. *Journal of Biomechanical Engineering*, 137(3):1–10, 2015.
- [9] Abhay B. Ramachandra, Andrew M. Kahn, and Alison L. Marsden. Patient-Specific Simulations Reveal Significant Differences in Mechanical Stimuli in Venous and Arterial Coronary Grafts. *Journal of Cardiovascular Translational Research*, 9(4):279–290, 2016.
- [10] Abhay B. Ramachandra, Jay D. Humphrey, and Alison L. Marsden. Gradual loading ameliorates maladaptation in computational simulations of vein graft growth and remodelling. *Journal of the Royal Society Interface*, 14(130), 2017.
- [11] Abhay B. Ramachandra, Hanjay Wang, Alexa Wnorowski, Erica L. Schwarz, Joshua Pickering, Joseph C. Heiler, Haley J. Lucian, Camille E. Hironaka, Nicholas A. Tran, Yu Liu, Muhammad Owais Khan, Oluwatomisin Obafemi, Yuko Tada, Andrew M. Kahn, Nazish Sayed, Joseph C. Wu, Jay D. Humphrey, Jack H. Boyd, and Alison L. Marsden. Biodegradable external wrapping promotes favorable adaptation in an ovine vein graft model. *Acta Biomaterialia*, 151:414–425, 2022.
- [12] Alessandro Candreva, Mattia Pagnoni, Maurizio Lodi Rizzini, Takuya Mizukami, Emanuele Gallinoro, Valentina Mazzi, Diego Gallo, David Meier, Toshiro Shinke, Jean Paul Aben, Sakura Nagumo, Jeroen Sonck, Daniel Munhoz, Stephane Fournier, Emanuele Barbato, Ward Heggermont, Stephane Cook, Claudio Chiastra, Umberto Morbiducci, Bernard De Bruyne, Oliver Muller, and Carlos Collet. Risk of myocardial infarction based on endothelial shear stress analysis using coronary angiography. *Atherosclerosis*, 342:28–35, 2 2022.
- [13] Karthik Menon, Jongmin Seo, Ryuji Fukazawa, Shunichi Ogawa, Andrew M. Kahn, Jane C. Burns, and Alison L. Marsden. Predictors of Myocardial Ischemia in Patients with Kawasaki Disease: Insights from Patient-Specific Simulations of Coronary Hemodynamics. *Journal of Cardiovascular Translational Research*, 16(5):1099–1109, 2023.
- [14] Cecil D. Murray. The Physiological Principle of Minimum Work 1: The Vascular System and the Cost of Blood Volume. *Proceedings of the National Academy of Sciences*, 12(3):207–214, 3 1926.
- [15] Yifang Zhou, Ghassan S. Kassab, and Sabee Molloj. On the design of the coronary arterial tree: A generalization of Murray's law. *Physics in Medicine and Biology*, 44(12):2929–2945, 1999.
- [16] Karthik Menon, Muhammed Owais Khan, Zachary A. Sexton, Jakob Richter, Patricia K. Nguyen, Sachin B. Malik, Jack Boyd, Koen Nieman, and Alison L. Marsden. Personalized coronary and myocardial blood flow models incorporating CT perfusion imaging and synthetic vascular trees. *npj Imaging*, 2(9), 2024.
- [17] Xiaofei Xue, Xiujuan Liu, Zhifan Gao, Rui Wang, Lei Xu, Dhanjoo Ghista, and Heye Zhang. Personalized coronary blood flow model based on CT perfusion to non-invasively calculate fractional flow reserve. *Computer Methods in Applied Mechanics and Engineering*, 404:115789, 2023.
- [18] Xuelian Gao, Rui Wang, Zhonghua Sun, Hongkai Zhang, Kairui Bo, Xiaofei Xue, Junjie Yang, and Lei Xu. A Novel CT Perfusion-Based Fractional Flow Reserve Algorithm for Detecting Coronary Artery Disease. *Journal of Clinical Medicine*, 12(6):2154, 2023.
- [19] Koen Nieman and Sujana Balla. Dynamic CT myocardial perfusion imaging. *Journal of Cardiovascular Computed Tomography*, 14(4):303–306, 2020.

[20] Bernard De Bruyne, Nico H.J. Pijls, Emanuele Gallinoro, Alessandro Candreva, Stephane Fournier, Danielle C.J. Keulards, Jeroen Sonck, Marcel van't Veer, Emanuele Barbato, Jozef Bartunek, Marc Vanderheyden, Eric Wyffels, Annemiek De Vos, Mohamed El Farissi, Pim A.L. Tonino, Olivier Muller, Carlos Collet, and William F. Fearon. Microvascular Resistance Reserve for Assessment of Coronary Microvascular Function: JACC Technology Corner. *Journal of the American College of Cardiology*, 78(15):1541–1549, 2021.

[21] Mario Gaudino, Sigrid Sandner, Gabriele Di Giammarco, Antonino Di Franco, Hirokuni Arai, Tohru Asai, Faisal Bakaeen, Torsten Doenst, Stephen E. Fremes, David Glineur, Teresa M. Kieser, Jennifer S. Lawton, Roberto Lorusso, Nirav Patel, John D. Puskas, James Tatoulis, David P. Taggart, Michael Vallely, and Marc Ruel. The Use of Intraoperative Transit Time Flow Measurement for Coronary Artery Bypass Surgery: Systematic Review of the Evidence and Expert Opinion Statements. *Circulation*, 144(14):1160–1171, October 2021.

[22] Nico H.J. Pijls, Bernard de Bruyne, Kathinka Peels, Pepijn H. van der Voort, Hans J.R.M. Bonnier, Jozef Bartunek, and Jacques J. Koolen. Measurement of Fractional Flow Reserve to Assess the Functional Severity of Coronary-Artery Stenoses. *New England Journal of Medicine*, 334(26):1703–1708, 6 1996.

[23] G Ninos, V Bartzis, N Merlemis, and I E Sarris. Uncertainty quantification implementations in human hemodynamic flows. *Computer Methods and Programs in Biomedicine*, 203:106021, 2021.

[24] Justin S. Tran, 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 and Fluids*, 142:128–138, 2017.

[25] Matteo Salvador, 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.

[26] Jakob Richter, 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 0D surrogate models. *arXiv*, 2404.14187, 2024.

[27] Vinzenz Gregor Eck, Wouter Paulus Donders, Jacob Sturdy, Jonathan Feinberg, Tammo Delhaas, Leif Rune Hellevik, and Wouter Huberts. A guide to uncertainty quantification and sensitivity analysis for cardiovascular applications. *International Journal for Numerical Methods in Biomedical Engineering*, 32(8):e02755, 2016.

[28] A. Boccadifuoco, A. Mariotti, S. Celi, N. Martini, and M. V. Salvetti. Impact of uncertainties in outflow boundary conditions on the predictions of hemodynamic simulations of ascending thoracic aortic aneurysms. *Computers and Fluids*, 165:96–115, 2018.

[29] Justin S. Tran, Daniele E. Schiavazzi, Andrew M. Kahn, and Alison L. Marsden. Uncertainty quantification of simulated biomechanical stimuli in coronary artery bypass grafts. *Computer Methods in Applied Mechanics and Engineering*, 345:402–428, 2019.

[30] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Optimal model management for multifidelity Monte Carlo estimation. *SIAM Journal on Scientific Computing*, 38(5):A3163–A3194, 2016.

[31] Casey M. Fleeter, Gianluca Geraci, Daniele E. Schiavazzi, Andrew M. Kahn, and Alison L. Marsden. Multilevel and multifidelity uncertainty quantification for cardiovascular hemodynamics. *Computer Methods in Applied Mechanics and Engineering*, 365:113030, 2020.

[32] Jongmin Seo, Casey Fleeter, Andrew M. Kahn, Alison L. Marsden, and Daniele E. Schiavazzi. Multi-fidelity estimators for coronary circulation models under clinically-informed data uncertainty. *International Journal for Uncertainty Quantification*, 10(5):449–466, 2019.

[33] D E Schiavazzi, G Arbia, C Baker, A M Hlavacek, T Y Hsia, A L Marsden, I E Vignon-Clementel, and The Modeling of Congenital Hearts Alliance (MOCHA) Investigators. Uncertainty quantification in virtual surgery hemodynamics predictions for single ventricle palliation. *International Journal for Numerical Methods in Biomedical Engineering*, 32(3):e02737, 2016.

[34] G. Geraci and M.S. Eldred. Leveraging intrinsic principal directions for multifidelity uncertainty quantification. *SAND2018-10817*, 2018.

[35] G. Geraci, M.S. Eldred, A.A. Gorodetsky, and J.D. Jakeman. Leveraging active directions for efficient multifidelity uncertainty quantification. In *6th European Conference on Computational Mechanics (ECCM 6)*, pages 2735–2746, 2018.

[36] Xiaoshu Zeng, Gianluca Geraci, Michael S. Eldred, John D. Jakeman, Alex A. Gorodetsky, and Roger Ghanem. Multifidelity uncertainty quantification with models based on dissimilar parameters. *Comput. Methods Appl. Mech. Engrg.*, 415:Paper No. 116205, 36, 2023.

[37] Andrea Zanoni, Gianluca Geraci, Matteo Salvador, Karthik Menon, Alison L Marsden, and Daniele E Schiavazzi. Improved multifidelity Monte Carlo estimators based on normalizing flows and dimensionality reduction techniques. *Computer Methods in Applied Mechanics and Engineering*, 429:117119, 2024.

[38] Andrea Zanoni, Gianluca Geraci, Matteo Salvador, Karthik Menon, Alison L. Marsden, and Daniele E. Schiavazzi. Linear and nonlinear dimension reduction strategies for multifidelity uncertainty propagation of nonparametric distributions. *AIAA SciTech Forum*, 2024.

[39] Andrea Zanoni, Gianluca Geraci, Matteo Salvador, Alison L. Marsden, and Daniele E. Schiavazzi. Neuram: nonlinear dimensionality reduction for uncertainty quantification through neural active manifolds. Preprint arXiv:2408.03534, 2024.

[40] Andreas H. Mahnken, Ernst Klotz, Hubertus Pietsch, Bernhard Schmidt, Thomas Allmendinger, Ulrike Haberland, Willi A. Kalender, and Thomas Flohr. Quantitative whole heart stress perfusion ct imaging as noninvasive assessment of hemodynamics in coronary artery stenosis: Preliminary animal experience. *Investigative Radiology*, 45(6):298–305, 2010.

[41] M Owais Khan, Anahita A Seresti, Karthik Menon, Alison L Marsden, and Koen Nieman. Quantification and Visualization of CT Myocardial Perfusion Imaging to Detect Ischemia-Causing Coronary Arteries. *IEEE Transactions on Medical Imaging*, page 1, 2024.

[42] Simone Di Gregorio, Marco Fedele, Gianluca Pontone, Antonio F. Corno, Paolo Zunino, Christian Vergara, and Alfio Quarteroni. A computational model applied to myocardial perfusion in the human heart: From large coronaries to microvasculature. *Journal of Computational Physics*, 424, 1 2021.

[43] Lazaros Papamanolis, Hyun Jin Kim, Clara Jaquet, Matthew Sinclair, Michiel Schaap, Ibrahim Danad, Pepijn van Diemen, Paul Knaapen, Laurent Najman, Hugues Talbot, Charles A. Taylor, and Irene Vignon-Clementel. Myocardial Perfusion Simulation for Coronary Artery Disease: A Coupled Patient-Specific Multiscale Model. *Annals of Biomedical Engineering*, 49(5):1432–1447, 2021.

[44] Gabriel Maher, Casey Fleeter, Daniele Schiavazzi, and Alison Marsden. Geometric Uncertainty in Patient-Specific Cardiovascular Modeling with Convolutional Dropout Networks. *Computer Methods in Applied Mechanics and Engineering*, 386:114038, 2020.

[45] KE Jansen, CH Whiting, and G.M. Hulbert. A generalized-a method for integrating the filtered Navier-Stokes equations with a stabilized finite element method. *Computer Methods in Applied Mechanics and Engineering*, 190(3-4):305–320, 1999.

[46] Christian H. Whiting and Kenneth E. Jansen. A stabilized finite element method for the incompressible Navier-Stokes equations using a hierarchical basis. *International Journal for Numerical Methods in Fluids*, 35(1):93–116, 2001.

[47] Timothy W. Secomb. Blood Flow in the Microcirculation. *Annual Review of Fluid Mechanics*, 49:443–461, 2017.

[48] Parastou Eslami, Justin Tran, Zexi Jin, Julia Karady, Romina Sotoodeh, Michael T. Lu, Udo Hoffmann, and Alison Marsden. Effect of Wall Elasticity on Hemodynamics and Wall Shear Stress in Patient-Specific Simulations in the Coronary Arteries. *Journal of Biomechanical Engineering*, 142(2):1–10, 2020.

[49] Martin R. Pfaller, Jonathan Pham, Aekaansh Verma, Luca Pegolotti, Nathan M. Wilson, David W. Parker, Weiguang Yang, and Alison L. Marsden. Automated generation of 0D and 1D reduced-order models of patient-specific blood flow. *International Journal for Numerical Methods in Biomedical Engineering*, 38(10):e3639, 7 2022.

[50] Irene E. Vignon-Clementel, C. Alberto Figueroa, Kenneth E. Jansen, and Charles A. Taylor. Outflow boundary conditions for three-dimensional finite element modeling of blood flow and pressure in arteries. *Computer Methods in Applied Mechanics and Engineering*, 195(29-32):3776–3796, 2006.

[51] H. J. Kim, I. E. Vignon-Clementel, J. S. Coogan, C. A. Figueroa, K. E. Jansen, and C. A. Taylor. Patient-specific modeling of blood flow and pressure in human coronary arteries. *Annals of Biomedical Engineering*, 38(10):3195–3209, 2010.

[52] Charles A. Taylor, Timothy A. Fonte, and James K. Min. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: Scientific basis. *Journal of the American College of Cardiology*, 61(22):2233–2241, 2013.

[53] J. A. Nelder and R. Mead. A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313, 1 1965.

[54] Jasper A. Vrugt, C. J.F. Ter Braak, C. G.H. Diks, Bruce A. Robinson, James M. Hyman, and Dave Higdon. Accelerating Markov chain Monte Carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *International Journal of Nonlinear Sciences and Numerical Simulation*, 10(3):273–290, 2009.

[55] Andrew Gelman and Donald B Rubin. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457–472, 1992.

[56] Michael B Giles. Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617, 2008.

[57] K. A. Cliffe, M. B. Giles, R. Scheichl, and A. L. Teckentrup. Multilevel monte carlo methods and applications to elliptic pdes with random coefficients. *Computing and Visualization in Science*, 14(1):3, 2011.

[58] Michael B. Giles. Multilevel monte carlo methods. *Acta Numerica*, 24:259–328, 2015.

[59] Abdul-Lateef Haji-Ali, Fabio Nobile, and Raúl Tempone. Multi-index Monte Carlo: when sparsity meets sampling. *Numerische Mathematik*, 132(4):767–806, Apr 2016.

[60] A. Haji-Ali, F. Nobile, L. Tamellini, and R. Tempone. Multi-index stochastic collocation for random pdes. *Computer Methods in Applied Mechanics and Engineering*, 306:95 – 122, 2016.

[61] Fabio Nobile and Francesco Tesei. A multi level monte carlo method with control variate for elliptic pdes with log-normal coefficients. *Stochastic Partial Differential Equations: Analysis and Computations*, 3(3):398–444, Sep 2015.

[62] H.R. Fairbanks, A. Doostan, C. Ketelsen, and G. Iaccarino. A low-rank control variate for multilevel monte carlo simulation of high-dimensional uncertain systems. *Journal of Computational Physics*, 341:121–139, 2017.

[63] Gianluca Geraci, Michael S Eldred, and Gianluca Iaccarino. A multifidelity multilevel monte carlo method for uncertainty propagation in aerospace applications. In *19th AIAA non-deterministic approaches conference*, page 1951, 2017.

[64] Alex A. Gorodetsky, Gianluca Geraci, Michael S. Eldred, and J.D. Jakeman. A generalized approximate control variate framework for multifidelity uncertainty quantification. *Journal of Computational Physics*, 408:109257, 2020.

[65] G.F. Bomarito, P.E. Leser, J.E. Warner, and W.P. Leser. On the optimization of approximate control variates with parametrically defined estimators. *Journal of Computational Physics*, 451:110882, 2022.

[66] Daniel Schaden and Elisabeth Ullmann. On multilevel best linear unbiased estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):601–635, 2020.

[67] Daniel Schaden and Elisabeth Ullmann. Asymptotic analysis of multilevel best linear unbiased estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 9(3):953–978, 2021.

[68] M. Croci, K.E. Willcox, and S.J. Wright. Multi-output multilevel best linear unbiased estimators via semidefinite programming. *Computer Methods in Applied Mechanics and Engineering*, 413:116130, 2023.

[69] Leo W. T. Ng and Karen E. Willcox. Multifidelity approaches for optimization under uncertainty. *International Journal for Numerical Methods in Engineering*, 100(10):746–772, 2014.

[70] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Siam Review*, 60(3):550–591, 2018.

[71] Xiaoshu Zeng, Gianluca Geraci, Alex Gorodetsky, John Jakeman, Michael S Eldred, and Roger Ghanem. Improving bayesian networks multifidelity surrogate construction with basis adaptation. In *AIAA Scitech 2023 Forum*, pages AIAA 2023–0917, 2023.

[72] David N. Ku. Blood flow in arteries. *Annual Review of Fluid Mechanics*, 29:399–434, 11 1997.

[73] Nico H.J. Pijls, William F. Fearon, Pim A.L. Tonino, Uwe Siebert, Fumiaki Ikeno, Bernhard Bornschein, Marcel Van'T Veer, Volker Klauss, Ganesh Manoharan, Thomas Engstrøm, Keith G. Oldroyd, Peter N. Ver Lee, Philip A. MacCarthy, and Bernard De Bruyne. Fractional flow reserve versus angiography for guiding percutaneous coronary intervention in patients with multivessel coronary artery disease: 2-Year follow-up of the FAME (fractional flow reserve versus angiography for multivessel evaluation) study. *Journal of the American College of Cardiology*, 56(3):177–184, 7 2010.

[74] Fay M.A. Nous, Tobias Geisler, Mariusz B.P. Kruk, Hatem Alkadhi, Kakuya Kitagawa, Rozemarijn Vliegenthart, Michaela M. Hell, Jörg Hausleiter, Patricia K. Nguyen, Ricardo P.J. Budde, Konstantin Nikolaou, Cezary Kepka, Robert Manka, Hajime Sakuma, Sachin B. Malik, Adriaan Coenen, Felix Zijlstra, Ernst Klotz, Pim van der Harst, Christoph Artzner, Admir Dedic, Francesca Pugliese, Fabian Bamberg, and Koen Nieman. Dynamic Myocardial Perfusion CT for the Detection of Hemodynamically Significant Coronary Artery Disease. *JACC: Cardiovascular Imaging*, 15(1):75–87, 1 2022.

[75] Pim A.L. Tonino, Bernard De Bruyne, Nico H.J. Pijls, Uwe Siebert, Fumiaki Ikeno, Marcel T. Van Veer, Volker Klauss, Ganesh Manoharan, Thomas Engstrøm, Keith G. Oldroyd, Peter N. Ver Lee, Philip A. MacCarthy, and William F. Fearon. Fractional flow reserve versus angiography for guiding percutaneous coronary intervention. *New England Journal of Medicine*, 360(3):213–223, 1 2009.

[76] Ryo Nakazato, Hyung Bok Park, Daniel S. Berman, Heidi Gransar, Bon Kwon Koo, Andrejs Erglis, Fay Y. Lin, Allison M. Dunning, Matthew J. Budoff, Jennifer Malpeso, Jonathon Leipsic, and James K. Min. Noninvasive fractional flow reserve derived from computed tomography angiography for coronary lesions of intermediate stenosis severity results from the DeFACTO study. *Circulation: Cardiovascular Imaging*, 6(6):881–889, 2013.

[77] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences of the United States of America*, 117(48):30055–30062, 2020.

[78] Guoxiang Grayson Tong, Carlos A Sing Long, and Daniele E Schiavazzi. InVAErt networks: A data-driven framework for model synthesis and identifiability analysis. *Computer Methods in Applied Mechanics and Engineering*, 423:116846, 2024.

[79] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv*, arXiv:1412, 2017.

[80] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25th {ACM} {SIGKDD} International Conference on Knowledge Discovery and Data Mining*, 2019.