




Real-Time Peripheral Revascularization Planning in Chronic Limb Threatening Ischemia Using HarVI: A Digital Twin Approach

Cyrus Tanade¹ · Christopher W. Jensen² · Guinevere Ferreira¹ · Amanda Randles¹ 

Received: 27 June 2025 / Accepted: 14 February 2026
© The Author(s) under exclusive licence to Biomedical Engineering Society 2026

Abstract

Background Peripheral artery disease (PAD) is a leading cause of limb loss and morbidity worldwide, with chronic limb-threatening ischemia (CLTI) representing its most severe presentation. Although image-guided endovascular interventions are routinely performed, clinicians currently lack tools that provide real-time, patient-specific predictions of hemodynamic outcomes to guide revascularization decisions. Existing computational fluid dynamics (CFD) approaches can recover pre-operative hemodynamics but are typically too slow or insufficiently integrated into clinical workflows to support interactive, intraoperative planning.

Methods We extend HarVI (HARVEY Virtual Intervention), a previously established digital twin framework, to the peripheral circulation and evaluate its use for real-time prediction of postoperative blood flow in patients with superficial femoral artery (SFA) lesions. HarVI integrates one-dimensional CFD with machine learning to enable rapid assessment of patient-specific revascularization strategies. Key components include: (1) automated boundary condition tuning using patient-averaged and optimization-based approaches; (2) simulation of a wide range of endovascular interventions via a machine-learned surrogate model; and (3) validation of predicted postoperative hemodynamics against clinical duplex ultrasound measurements. Performance was evaluated retrospectively in a cohort of seven patients with SFA disease.

Results HarVI accurately predicted postoperative peak systolic velocities and reproduced full 1D CFD results across a synthetic revascularization landscape. Surrogate model predictions closely matched high-fidelity simulations while enabling rapid exploration of intervention scenarios, supporting near-real-time evaluation of treatment options.

Conclusions These results establish HarVI as a promising digital twin platform for real-time, patient-specific intervention planning in PAD. By enabling rapid, data-driven prediction of postoperative hemodynamics, HarVI opens the door to interactive intraoperative decision support with the potential to improve revascularization outcomes in patients with CLTI.

Keywords Treatment planning · Peripheral artery disease · Computational fluid dynamics · Machine learning

Associate Editor Keefe B. Manning, PhD oversaw the review of this article.

✉ Amanda Randles
amanda.randles@duke.edu

Cyrus Tanade
cyrus.tanade@duke.edu

Christopher W. Jensen
christopher.w.jensen@duke.edu

Guinevere Ferreira
guinevere.ferreira@duke.edu

¹ Department of Biomedical Engineering, Duke University, 534 Research Dr., Durham, NC 27705, USA

² Department of Surgery, Duke University, 2301 Erwin Rd., Durham, NC 27707, USA

Introduction

Real-time hemodynamic modeling has the potential to improve peripheral artery disease (PAD) treatment by enabling patient-specific planning of endovascular interventions. Recent developments in digital twins, personalized simulations that use observational data to predict disease states, have shown substantial promise in cardiovascular applications [1–3]. However, most of the existing advances are focused on coronary artery disease [4–7], leaving PAD relatively understudied. Peripheral disease often presents with multi-lesion or multi-level vessel involvement, complicating the identification of flow-limiting lesions. As a result, digital twin frameworks that can deliver patient-specific real-time predictions of postoperative hemodynamics are

particularly needed to guide intervention planning in these complex conditions.

PAD is primarily caused by the accumulation of atherosclerotic plaque, leading to narrowing or occlusion of the arterial walls and impaired blood flow to the limbs, most commonly the legs. This impaired circulation commonly leads to leg discomfort during physical activity (claudication), as well as numbness, muscle weakness, and, in advanced stages, ulceration or gangrene. In the United States alone, PAD affects roughly 7% of adults, or about 8.5 million people [8]. Chronic limb-threatening ischemia (CLTI) is a severe form of PAD characterized by persistent rest pain due to inadequate blood circulation. Between 2003 and 2008, its prevalence reached approximately 1.3% among adults over 40 years old [9], and it confers a 25% risk of amputation within the first year of diagnosis [10]. The healthcare burden of PAD currently exceeds \$21 billion annually, with projections indicating that this figure may double by 2050 [11]. Surgical revascularization, through traditional open bypass or minimally invasive methods such as angioplasty or arterial stenting, remains the mainstay therapy for CLTI. Yet, limb loss still occurs in more than 15% of cases [12], and in-stent restenosis (ISR) is highly prevalent in the superficial femoral artery (SFA), with incidences of 15–32% [13]. Diabetes, smoking, excessive stent length, overlapping stents, and stent overexpansion (leading to neointimal hyperplasia) have been implicated in driving ISR [14–16]. Consequently, for patients with CLTI, it is crucial to determine not only the optimal stenting configuration to achieve full functional revascularization but also the stent size that minimizes overexpansion and the associated risk of ISR.

Digital twin frameworks based on computational fluid dynamics (CFD) provide a mechanism to assess the severity of the lesion and predict the risk of restenosis in patients with CLTI [10, 17–22]. While 3D CFD approaches can resolve complex flow structures, they remain computationally prohibitive for intra-operative use [20, 23]—particularly in PAD where rapid evaluation of multiple treatment scenarios is essential. In contrast, one-dimensional (1D) CFD models offer a clinically viable alternative: they significantly reduce computation time while still accurately capturing key hemodynamic phenomarkers. Despite their anatomical simplifications, 1D models have been validated with clinical ground truths [24, 25]. However, most existing femoral digital twin methods do not predict postoperative flow dynamics or fail to align with standard clinical workflows. In the pre-operative stage, it is routine clinical practice to first evaluate disease burden via ankle–brachial index (ABI) [11, 26–28] and peak systolic velocity (PSV) from duplex ultrasound (DUS) [11, 29–36], each endorsed by multiple global guidelines [11]. PSV and PSV_r are clinically adopted, interpretable markers of flow-limiting stenosis and post-intervention patency, with guideline-based thresholds used to identify

clinically significant restenosis in CLTI surveillance. Guidelines for clinical practice on management of CLTI indicate that PSV values $> 30 \text{ cms}^{-1}$ and PSV_r > 3.5 are explicitly indicative of negative outcomes [11]. While late complications such as restenosis are multifactorial and influenced by factors beyond conduit-level flow, PSV provides a practical and actionable target for real-time hemodynamic feedback during intervention planning. While the pathological progression of restenosis occurs at both macroscopic and microscopic scales, it is important to acknowledge that the hemodynamic impact on mechanosensitive signaling at the vascular bed level may not be fully captured. Perturbations in wall shear stress (WSS) can influence both microscopic signaling pathways and macroscopic remodeling patterns, which are not directly represented by flow-based metrics such as PSV. Despite the clinical utility of these metrics, a degree of separation exists between such simplified hemodynamic descriptors and the local flow phenomena resolved by higher-fidelity models. Consequently, the predictive power of one-dimensional (1D) models for complications such as restenosis is inherently limited by their inability to resolve local WSS gradients, secondary flow structures, or recirculation zones that have been implicated in maladaptive vascular responses. Nevertheless, by providing rapid, anatomically informed estimates of flow and pressure, 1D models remain valuable for identifying global hemodynamic trends and for stratifying patients according to systemic or conduit-level risk factors. In this context, the present work focuses on PSV because it remains a clinically adopted and widely interpretable surrogate for luminal narrowing.

When an intervention is deemed necessary, intra-operative angiography typically occurs days to weeks after these initial evaluations. Because PSV is central to SFA post-stenting surveillance [29, 30], our work focuses on validating digital twin predictions against postoperative PSV.

Real-time planning is particularly critical in femoral disease, where multi-lesion, multi-level disease leads to a large combinatorial space of treatment options. For example, a patient with three lesions at three vessel levels and three stenting options per lesion, along with three potential restenosis scenarios, generates more than 500,000 possible configurations. Testing even a small subset of these with conventional physics based simulation methods, each requiring approximately 10 minutes [4, 5, 24, 25]—is computationally infeasible within an operative timeframe, even with parallelization.

Although machine learning–accelerated intervention planners have been developed for coronary disease [4, 5] and some have been clinically validated [6, 7], a comparable system for femoral circulation is lacking. Existing models interpolate across the intervention state space, the parameter space of possible treatment scenarios, to compute coronary-specific metrics, such as fractional flow reserve. A

PAD-specific system must accommodate preoperative imaging inputs, predict postoperative PSV, and explore treatment permutations in real time.

To meet these needs, we extend the HARVEY Virtual Intervention (HarVI) framework, originally developed for and established in coronary artery disease [4, 5], to peripheral artery disease for the first time. HarVI is a real-time digital twin platform built on HARVEY [37] that combines 1D CFD modeling with machine learning (ML) to automate boundary condition tuning, rapidly simulate revascularization strategies, and predict clinically relevant postoperative hemodynamics. In this work, we test the hypothesis that preoperative arterial flow measurements and anatomical models derived from angiography can forecast postoperative PSV in patients with superficial femoral artery (SFA) lesions. In a retrospective cohort of seven CLTI patients, we demonstrate that HarVI achieves clinically concordant predictions and

reproduces full 1D CFD results in a large synthetic intervention space. This study represents the first application of HarVI to PAD and the first clinically validated digital twin framework to predict postoperative flow in this patient population, establishing a foundation for planning patient-specific intervention in real-time in peripheral vascular care.

Methods

We established a framework in this study, HarVI for CLTI, to enable real-time intervention planning and validate the accurate capture of postoperative hemodynamics. A high-level overview of HarVI and how it integrates into the standard clinical workflow for the femoral arteries is shown in Fig. 1. The process begins with a preoperative evaluation, where patients typically present with symptoms such as

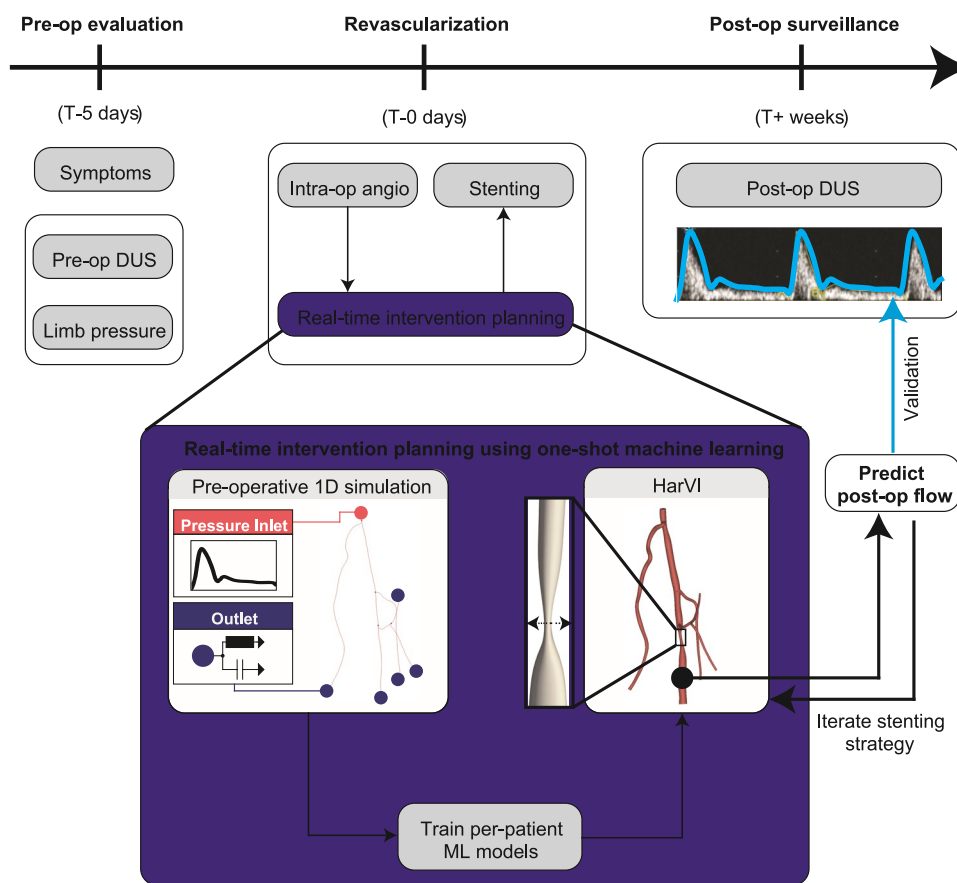


Fig. 1 HarVI framework for real-time revascularization planning within the context of standard clinical workflow for the treatment and management of patients with femoral disease. Initial patient screens are performed using duplex ultrasound as well as computing metrics such as ABI [38] and PSV. Moreover, segmental limb pressures are also measured here that are used to inform boundary conditions. Normally, a few days after the initial point of contact with the patient, intraoperative angiography is performed in conjunction with treat-

ment (such as endovascular stenting). 1D models are parameterized with a pressure inlet condition and 2-element Windkessel models at the outlet boundary condition. These 1D models are used to train HarVI for real-time virtual intervention. HarVI can predict postoperative distal velocities that are normally measured on the timescales of weeks to months after the intervention. Physicians can use HarVI to predict post-op flow metrics and use real-time hemodynamic feedback to iteratively guide clinical decisions intra-operatively

claudication. The clinicians then collect quantitative metrics, including preoperative DUS images and limb pressures, to evaluate measures such as PSV and ABI. When an intervention is necessary, they perform intraoperative angiography in the operating room. We used a semi-automated algorithm, detailed in “[Personalized Femoral Digital Twins](#)” section of “[Methods](#)” section to create patient-specific geometries. Measurements from the preoperative stage served as inputs for the boundary conditions of our blood flow models. Next, we generated preoperative 1D CFD models and trained patient-specific ML models to interpolate the potential intervention space (see “[Interpolating the Intervention State Space Using the HarVI Framework](#)” section of “[Methods](#)” section). After training a patient’s ML model (a one-time process per patient), simulations can be repeated to obtain realtime predictions of postoperative hemodynamics under various endovascular stenting strategies. This intervention planning step can be repeated as many times as needed to refine and iterate the stenting approach. Finally, we performed a validation step by comparing held-out interventions—those actually carried out in the clinic—to HarVI predictions. Clinically relevant metrics, such as postoperative PSV, were measured and compared against model outputs, ensuring that the framework accurately reflects real-world hemodynamic changes.

Patient Data

This study did not involve direct patient experimentation, interaction, or human tissue samples. This study was conducted under Duke University IRB approval (Pro00091585). Data used for this analysis were acquired at Duke University Medical Center. Seven patients with severe CLTI and SFA lesions, documented by digital subtraction angiography, were retrospectively identified and recruited. All patients underwent angioplasty and/or endovascular stenting, with a repeat angiography performed to confirm vessel patency. Angiograms comprised at least four standard orthogonal views of the femoral circulation, and physiological measurements, such as upper to lower thigh pressure drops and heart rate, were collected during initial DUS screening and intraoperative angiography to inform personalized digital twins. Preoperative and postoperative DUS images were available for all seven patients, and postoperatively, each patient was placed on dual anti-platelet therapy.

Personalized Femoral Digital Twins

SFA geometric models were reconstructed from pairs of angiograms using algorithms described in [39, 40] and [10], which have also been extensively applied to the coronary arteries [1, 2, 4, 5, 23–25]. Briefly, the two-dimensional vessel centerlines were first calculated along with the

corresponding cross-sectional diameters. A fully automated process then generated 3D femoral models from these 2D centerlines. This method was used to reconstruct patient-specific 3D femoral arteries before and after endovascular intervention. Reconstruction accuracy was validated topologically and anatomically using ImageJ v1.52 k (NIH, Bethesda, MD, USA). Specifically, the minimal luminal diameter (MLD) of each stenotic lesion in the 3D model was compared with that of the same diseased segment in the original angiograms. To create 1D geometries, the vessel centerlines and hydraulic diameters were extracted from the reconstructed 3D models using Mimics (Materialise, Leuven, BE). The lengths of the vessels were calculated at a resolution of 100 micrometers. These 1D femoral models were used as inputs for personalized CFD simulations. Finally, post-intervention 3D geometries were measured for the radius and length of the stenosis and served as a test set to validate postoperative hemodynamics. The stenotic radius is defined as the average radius across the stenosed segment. The location of the segment was identified through the relative percentage drop in radius above and below the area of interest. The reconstructive process took approximately 20–30 s per geometry. These 1D femoral models were used as inputs for personalized CFD simulations. Finally, post-intervention 3D geometries were measured for the radius and length of the stenosis and served as a test set to validate postoperative hemodynamics. Although this technique has been previously validated, it does not provide fully volumetric reconstructions comparable to CT angiography-based segmentation which could present as a source of error in stenotic and vessel radius calculations. 1D hemodynamic simulations were performed using the 1D blood flow simulator discussed in Feiger et al. [41]. This model has been clinically validated and its use is extended to PAD in this work. The primary inputs to the 1D simulator includes length and radius of each segment of a patient’s geometry, informed by parent-daughter relationships that indicate segment junction connections. These metrics are extracted from centerlines informed by stereolithography (STL) files of patient-specific segmented geometries. The radii of non-stenosed vessels are averaged along each centerline segment as input into the model. The boundary conditions were tuned to recover preoperative hemodynamics. Pulsatile pressure conditions were used at the inlet boundary condition. The estimated pressure drop from measurements of the upper and lower thigh cuffs prior to surgery was used to scale the pressure waveforms to a patient-specific level, discretely.

Two-element Windkessel models, consisting of peripheral resistance (R_p) and compliance (C), were used at the outlets to account for the effect of downstream hemodynamics. There are multiple ways to distribute R_p and C for terminal branches, and these separate strategies were evaluated for success in the HarVI framework. We first used

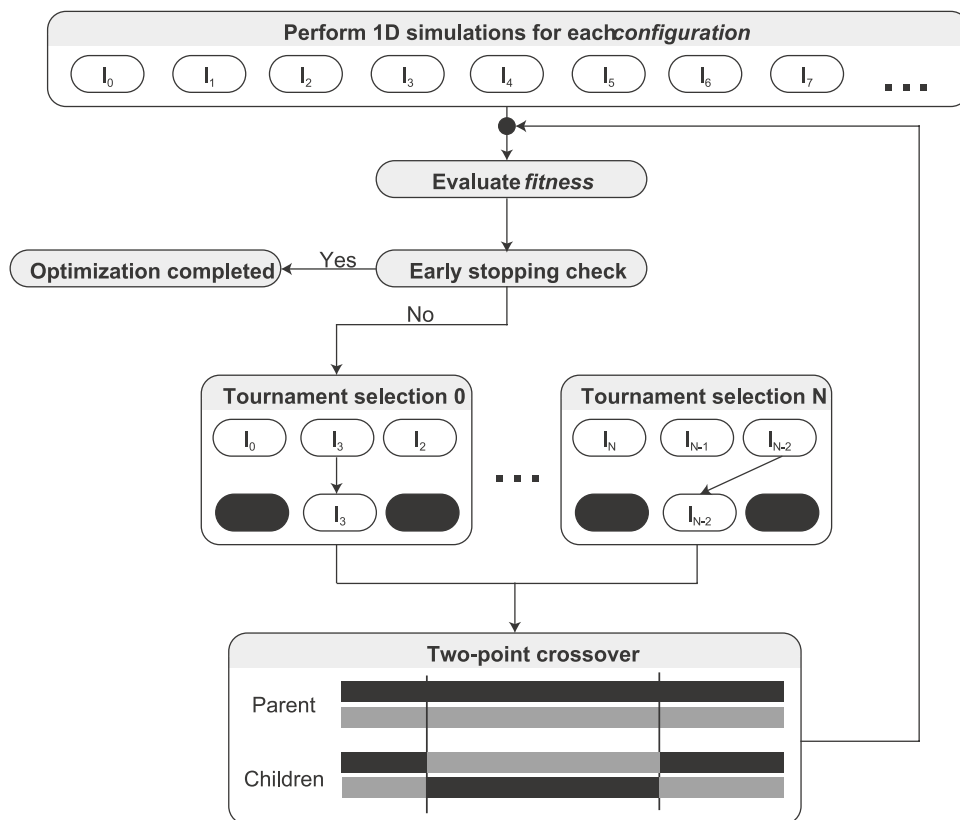
resistance-radius relationships, primarily Murray’s law, which is commonly used for coronary simulations [1, 2, 23–25]. The mean arterial pressure and inlet flow rate were obtained from clinical measurements at the time of initial patient screening using DUS and the terminal branch radii were calculated via the centerlines of the vessel.

Secondly, we developed a genetic algorithm to automatically tune boundary conditions to physiological values (see Fig. 2). These physiological values were derived from patient-specific DUS imaging at the pre-operative stage. Specifically, the algorithm optimized three parameters: (1) inlet pressure waveform amplitude, (2) arterial stiffness, and (3) the components of the two-element Windkessel models. The objective function minimized the mean-squared errors for two metrics: (1) the difference between distally measured preoperative peak velocities and the corresponding computational outputs at each time point, and (2) the difference between measured pressure drops across the SFA (upper vs. lower thigh) and those obtained computationally. These distal velocities were recorded at a distal SFA location preceding the vascular outlet but outside of the stenotic range and not immediately following the stenosis. This corresponds to approximately a superior thigh location, distal to the groin and proximal to the popliteal artery junction. The simulation of each patient had a population size of 1024 configurations. A sample size of 1024 was selected to enhance

computational efficiency as this number is a power of two (2^{10}) ensuring ease of binary manipulation. This value was an adequate sample space to balance efficiency with accuracy, based on iterative outputs of the GA. After computing the overall fitness for all configurations, a tournament selection process was applied. The population was divided into sets of three configurations each, with the “winning” configuration in each tournament generating new “child” configurations through two-point crossover, where the parameters were combined using weighted averages of the parent simulations, where weighted averages were based on the efficacy of the tournament prediction. Early stopping was used to terminate the genetic algorithm if no improvement in fitness was observed for three successive generations.

To ensure clinical translatability, we implemented a third set of boundary conditions which was applied to all patients based on the outputs of the genetic algorithm, which were averaged to obtain these resistances and compliances. The patient specific components of each simulations included the pressure waveforms and geometry. This approach was justified by the similar anatomy and morphology of the SFA lesion observed in all study patients. Specifically, we used an average peripheral resistance of $2406 \text{ gs}^{-1} \text{ cm}^{-4}$ for the side branches and $200 \text{ gs}^{-1} \text{ cm}^{-4}$ for the distal SFA. The pressure waveforms were still scaled to match the preoperatively measured thigh pressure drops.

Fig. 2 The genetic algorithm used to tune simulation boundary conditions. Pre-operative peak velocities, measured distally, and pressure drops across the superficial femoral artery were used as the feedback metrics that were minimized per reduced RMSE. 1024 1D simulations per patient were performed to sample adequate simulation inputs and results. Overall fitness was evaluated and if optimization was not initially achieved, the configurations entered a tournament-style selection process by which child configurations were developed from the winning condition configurations



This study treated blood flow in the femoral circulation as an incompressible Newtonian fluid, with a density of 1060 kg/m^3 and a dynamic viscosity of 4.2 cP , values that are typically used to approximate the properties of adult human blood. A pulsatile pressure waveform was prescribed at the inlet boundary condition, while two-element Windkessel models were applied at the outlets to capture the resistive and capacitive behavior of the downstream vasculature. These boundary conditions were kept constant for both the preoperative and postoperative states. This modeling approach assumes that the primary site of pressure recovery is the SFA lesion itself, allowing the isolation of the local hemodynamic impact of endovascular stenting. This simplification neglects potential changes in downstream resistance or other distal effects.

Interpolating the Intervention State Space Using the HarVI Framework

To fit within standard clinical workflows, it was crucial to quickly predict how different intervention strategies could affect postoperative hemodynamics. In this study, we addressed that need by developing ML models that can rapidly and accurately estimate post-intervention flow metrics based on a training set of 1D simulations (Fig. 3). These models were trained on geometric modifications randomly

sampled across stenosis radius and length, with each 1D simulation producing a corresponding set of post-intervention hemodynamic results.

We used Latin Hypercube Sampling (LHS) from the SALib library, implementing methods from [42, 43], to efficiently sample the intervention space and minimize the required training set. Quasi-random sampling was used over pseudo-random techniques to promote a more uniform distribution throughout the parameter space [4, 5]. Each LHS for each patient was initialized between 90% of the Minimum Luminal Diameter (MLD) and 110% of the unstenosed radius ($R_{unstenosed}$). This interval buffers the sampling range to reduce the uncertainty at its boundaries. Stenosis lengths were similarly sampled between 50 and 150% of the original lesion length. Radii and lengths were uniformly drawn within these bounds to construct the LHS sample matrix, which served as input for a series of 1D pulsatile simulations. We selected XGBoost for interpolation based on their training efficiency and prior promising results [5]. The training matrices were applied locally to modify the femoral geometries according to the scenarios sampled. To run the 1D simulations, we transferred all input files and SFA geometries to the Duke Compute Cluster, where each training instance (T_N) was assigned one CPU (Intel(R) Xeon(R) Gold 6336Y). This parallel workflow enabled the entire training set to be completed in the same time it would take to run

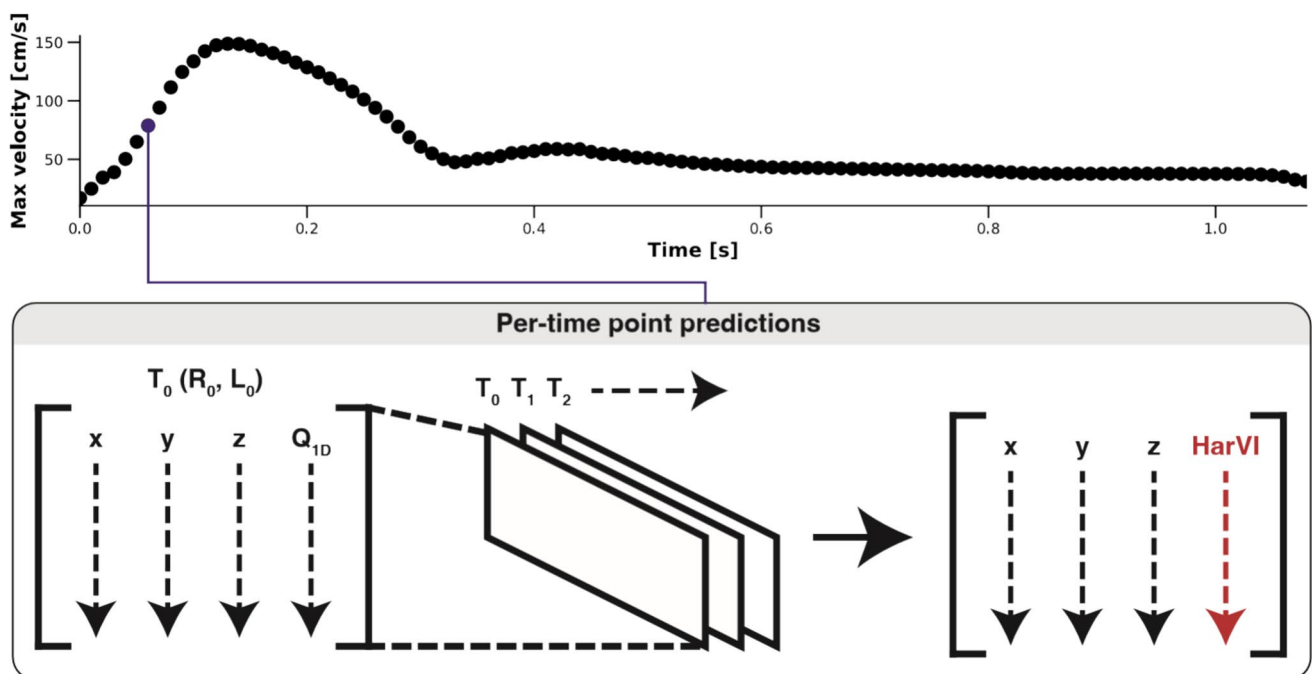


Fig. 3 Machine learning architecture for HarVI applied to femoral disease. Centerline data was first extracted from 3D patient-specific anatomy for 1D simulations for training the HarVI framework. This included (R_0) and (L_0), the baseline segment radius and length. Each predefined intervention scenario (T_N) for training consisted of creat-

ing a 2D matrix with 3D centerline data and corresponding flow results from 1D simulations. These matrices are then concatenated in-time over a cardiac cycle to predict pulsatile responses due to intervention

a single simulation sequentially. Post-intervention hemodynamics—including peak velocity at all centerline points for every timestep of the pulsatile waveform—were then returned to the local machine for ML model training. The trained XGBoost models predict hemodynamics at every centerline position and at every timepoint, as a function of stenosis radius and length, facilitating rapid assessments of potential intervention outcomes.

To establish HarVI, the initial step was determining the sample size required to accurately capture post-intervention hemodynamics. The goal was to identify the point of convergence, where adding more samples would not significantly change the flow predictions. We performed convergence studies across all patients, systematically varying the number of samples while using an 80–20% train–validation split. The validation set was used to estimate the average error across patients and time points at each sampling level. Specifically, we computed the root mean squared error (RMSE) over every time point in the peak velocity waveform at the distal SFA. RMSE was chosen as it heavily penalizes large deviations, making it well-suited for capturing discrepancies in both amplitude and timing of waveform features that are critical for accurate hemodynamic assessment. Based on these analyses (Fig. 4), we found that 40 samples were sufficient to capture post-intervention hemodynamics without further reduction in error.

Statistical Analysis

The Kolmogorov–Smirnov test was applied to verify that the data followed a normal distribution. Comparisons of computational model output with clinical measurements were then assessed using ordinary least-squares correlation and Bland–Altman analysis.

Results

Validating Blood Flow Digital Twins for Pre-operative Hemodynamics

To enable real-time virtual intervention planning, we needed to first validate our patient-specific modeling pipeline for accurate preoperative hemodynamics. Because these 1D simulations would be set up during intra-operative angiography, we aimed to minimize turnaround time through an automated approach. However, since standard methods for prescribing boundary conditions in SFA models are limited, we tested three tuning strategies: resistance-radius relationships via Murray’s Law, automated tuning using a genetic algorithm, and patient-averaged boundary conditions (detailed in “Personalized Hemodynamic Simulations” section of “Methods” section). We evaluated each method by calculating preoperative PSV and comparing the results to clinical

DUS measurements (Fig. 5). Pearson’s correlation coefficients were 0.63 ($p=0.12$), 0.94 ($p<0.05$), and 0.87 ($p<0.05$) for Murray’s Law, genetic algorithm, and patient-averaged conditions, respectively. The corresponding Bland–Altman biases and imprecisions were -32 ± 14.6 , -8.2 ± 8.2 , and -7.6 ± 9.6 cm/s. Given that both the genetic algorithm and patient-averaged approaches produced strong correlations and agreement, we selected patient-averaged boundary conditions as our first-line method for clinical translation because the genetic algorithm requires running thousands of simulations for each patient intra-operatively. The genetic algorithm was retained as a fallback if patient-averaged models fail to adequately capture preoperative hemodynamics, such as when the preoperative PSV is outside the Bland–Altman limits of agreement (differences larger than ± 9.6 cm/s). All subsequent results




Fig. 4 RMSE assessed at each time point to determine the number of samples needed to sufficiently represent postoperational hemodynamics. 40 samples was determined to be sufficient with low average RMSE as indicated by the red interrupted vertical line

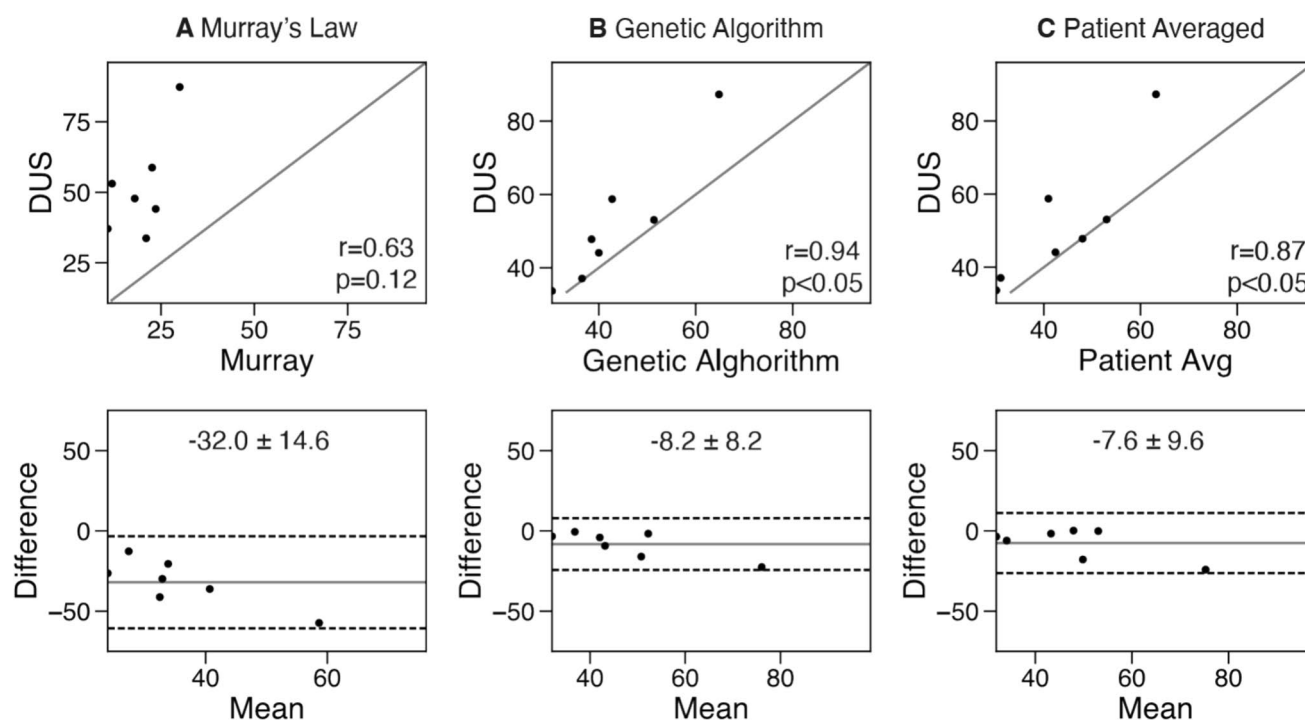


Fig. 5 Validating 1D pre-op models against DUS ground truths for PSV. Three methods: **A** Murray's Law, **B** Genetic algorithm—detailed in Fig. 2, and **C** patient-averaged from the genetic algorithm were compared to DUS PSV ground truths. The top row are scatter-plots comparing the different boundary conditions with DUS ground truths, with the solid diagonal line representing ideal correlation.

Pearson's correlation scores are displayed in the bottom right corner of each plot. The bottom row represents Bland-Altman plots showing displaying bias (solid horizontal line) and imprecision. Bias and imprecision scores are noted on the top half of each plot. Interrupted horizontal lines are 95% limits of agreement

presented here use patient-averaged boundary conditions. Overlays of preoperative peak velocity waveforms measured at the distal SFA from the 1D CFD models and DUS images are provided in Fig. 6.

HarVI Predicted Peak Systolic Velocities Agree with Clinical Ground Truths

To validate the accuracy of HarVI and 1D CFD in capturing postoperative hemodynamics in the patient set, we measured the stenosis radius and length from the reconstructed postoperative geometry and applied these values to the corresponding preoperative geometry to emulate real-life interventions. These details were applied to the preoperative imaging to create a more realistic surgical scenario where postoperative changes to other aspects of the geometry would be unknown. Overlays comparing peak velocity waveforms between HarVI and clinical ground truths are provided in Fig. 7. We focused on validating PSV due to its clinical relevance for assessing stenosis severity, post-stenting patency, and restenosis risk [11]. Specifically, we compared 1D CFD versus DUS and HarVI versus DUS measurements for both the postoperative PSV and the ratio of postoperative to preoperative PSV (Fig. 8).

For absolute postoperative PSV, the correlation coefficient between 1D CFD and DUS was of 0.82 ($p < 0.05$), with a Bland-Altman bias and imprecision of -15.7 ± 17.2 cm/s. The corresponding validation scores for HarVI versus DUS were 0.86 ($p < 0.05$) with -14.4 ± 14.5 cm/s. For the postoperative/preoperative PSV ratio, 1D CFD versus DUS had a correlation of 0.91 ($p < 0.05$) and a bias and imprecision of -0.03 ± 0.48 , whereas HarVI versus DUS had similar validation scores of 0.93 ($p < 0.05$) and -0.01 ± 0.40 . These results demonstrated that both 1D CFD and HarVI had a tendency to slightly underestimate postoperative hemodynamics. It is important to note a sample size of 7 patients was used ($n = 7$) in these calculations.

Furthermore, we substratified the validation scores by degree of stenosis and presence of downstream disease to understand whether HarVI performed better in certain subsets of the patient cohort (Table 1). The correlation and agreement scores for PSV were 0.88 and -19.9 ± 14.6 cm/s for lesions with $\leq 50\%$ stenosis, and 0.95 and -7.0 ± 12.9 cm/s for lesions with $> 50\%$ stenosis. Similarly, the validation scores for the PSV ratio were 0.93 and -0.26 ± 0.27 for lesions with $\leq 50\%$ stenosis, and 0.99 and 0.33 ± 0.25 for lesions with $> 50\%$ stenosis. In short, HarVI more closely matched clinical

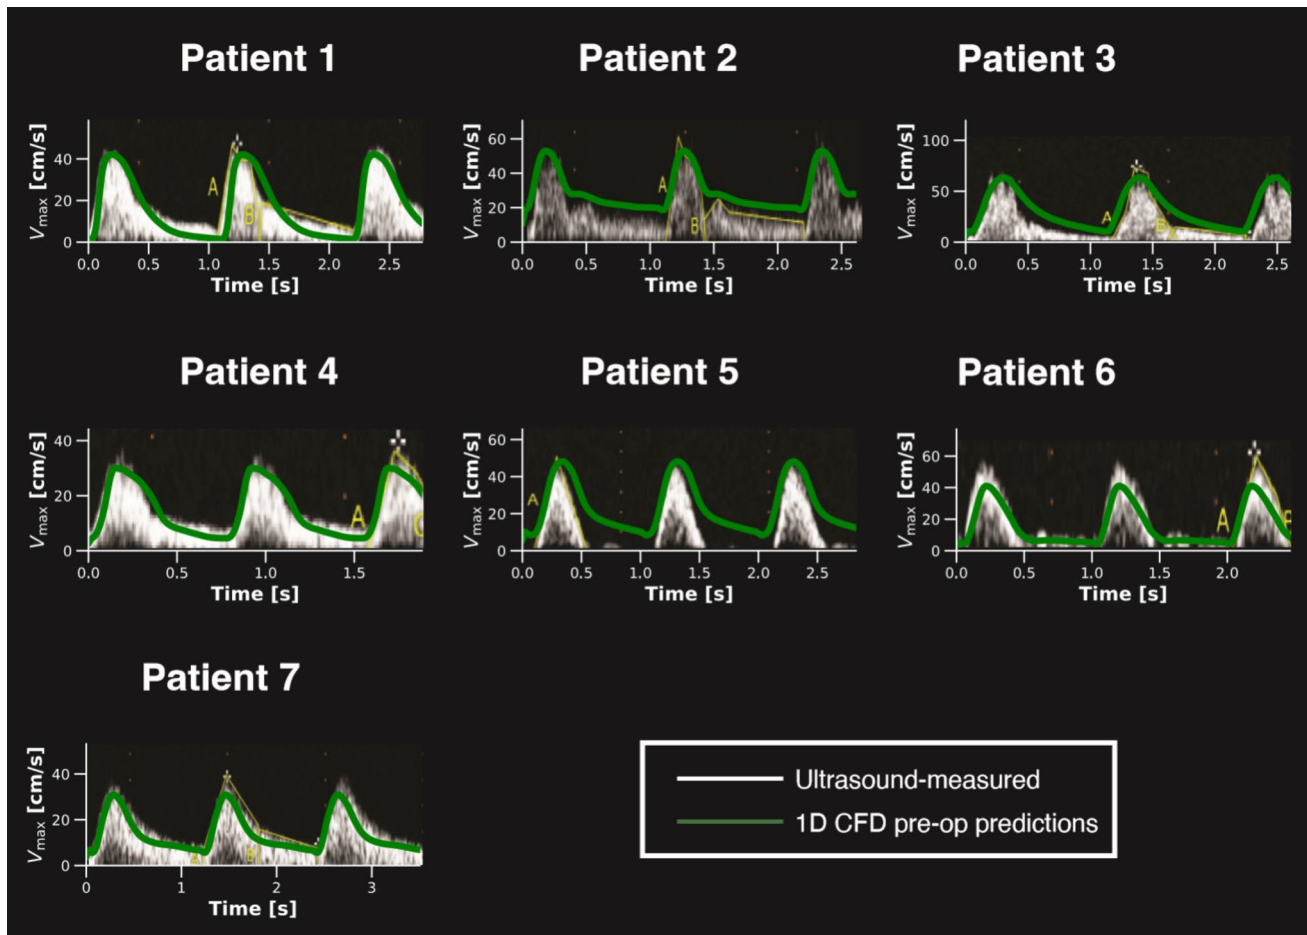


Fig. 6 Results from these simulations are overlaid (green) on patient-measured ultrasound flows. Of the three boundary condition strategies evaluated, patient-averaged boundary conditions were selected

and are shown. This figure highlights the agreement between the 1D pre-operative predictions and ultrasound measurements, demonstrating the validity of the 1D training data used in the HarVI pipeline

ground truths for lesions with greater than 50% stenosis degree. In terms of the PSV ratio, HarVI had a slight tendency to underestimate the ratio for lesions $\leq 50\%$ and overestimate it for lesions $> 50\%$.

When stratifying the cohort according to the presence of downstream disease, HarVI had a lower correlation and agreement for both PSV and the PSV ratio in cases with downstream disease, measuring 0.33 and -15.7 ± 18.6 cm/s for PSV and 0.85 and 0.06 ± 0.53 for the PSV ratio. In cases without downstream disease, the corresponding validation scores were 0.97 and -12.6 ± 10.0 cm/s for PSV, and 0.99 and -0.10 ± 0.19 for the PSV ratio. In general, HarVI matched clinical ground truths more closely in cases where the degree of stenosis was $> 50\%$ and there was no downstream disease.

Verifying HarVI with 1D Blood Flow Predictions in Synthetic Intervention Cases

After verifying that HarVI accurately reproduced clinical ground truths for postoperative hemodynamics, we conducted a synthetic verification study. From the initial cohort of 7 cases, we generated 25 randomly selected intervention scenarios (in terms of radius and length of stenosis) that were not part of the original training or clinical testing sets. To create this set of synthetic interventions, we first over-sampled ($n=5000$), corresponding to $100\times$ the number of samples required for convergence, and then selected 25 configurations that maximized the Euclidean distance from all previously used training and testing cases, thereby approximating a worst-case generalization challenge. We compared




Fig. 7 Overlays of HarVI postoperational predictions with DUS ground truths. White points are DUS-measured and blue lines are HarVI predicted maximum velocity waveforms

the postoperative PSV predicted by HarVI with the predictions of 1D CFD. Figure 9A illustrates both the preoperative and postoperative PSV values, while Fig. 9B shows the Bland–Altman plot. Bland–Altman bias and imprecision were -1.0 ± 2.1 cm/s (or 98.9% concordance on average), indicating that HarVI closely reproduced the 1D CFD results in these synthetic interventions.

HarVI Fits Within Current Clinical Workflow with Turnaround Time Within 10 Minutes

The clinical utility of HarVI relies on its ability to provide real-time predictions within a short lead time. PAD revascularization procedures vary in duration depending on the type of intervention. Angioplasty and stenting are typically 0.5–2 h procedures, while more complex bypass surgeries can take approximately 2–6 h. Patients who present with CLTI require intervention and revascularization within

hours to prevent amputation, emphasizing the need for rapid intervention frameworks [44]. As shown in Fig. 10, the entire setup requires under 10 min before intervention planning can begin. Most of this time is spent on the one-shot training process, which involves sampling the intervention space via pulsatile 1D models. When run in parallel, these simulations took an average of 220 s across all patients. Training the interpolation model with the resulting 1D CFD outputs required only 77 s on average, and generating postoperative hemodynamic predictions for any given geometric modification by querying HarVI only took 34 milliseconds on average. Training and subsequent querying was performed on an Apple Macbook Pro (2023) with Apple M4 chip (10-core CPU, 10-core GPU, 16-core Neural Engine), 128 GB of unified memory, and running macOS 15.1 (Sequoia). In a clinical workflow, a single laptop or workstation would be required to utilize the HarVI framework, underscoring its clinical translatability.

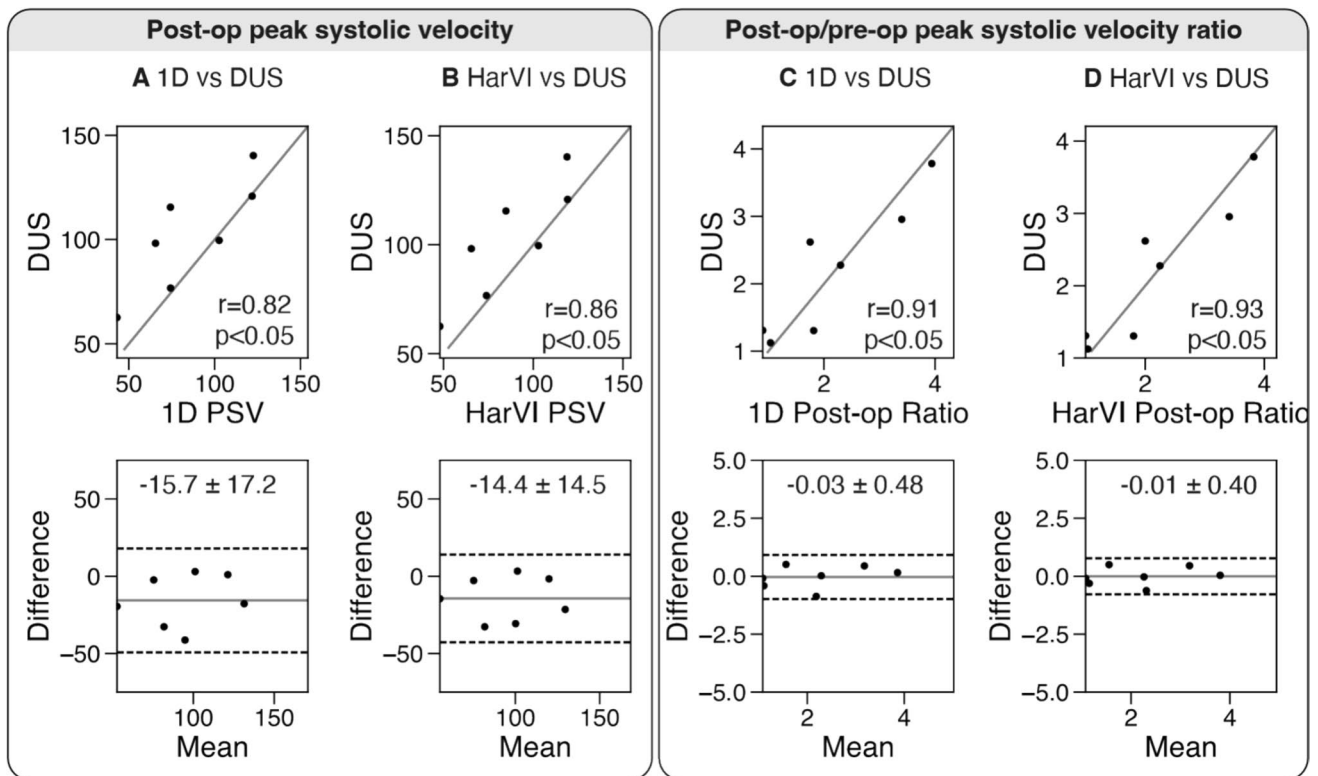


Fig. 8 Correlation and agreement between 1D CFD and HarVI with DUS ground truths for evaluating postoperative PSV and post-op/pre-op PSV ratio. Post-op PSVs are compared between **A** 1D CFD vs. DUS and **B** HarVI vs. DUS. Post-op/pre-op PSV ratios are compared between **C** 1D CFD vs. DUS and **D** HarVI vs. DUS. The top row are scatterplots. The solid diagonal lines represent ideal correlation.

Pearson’s correlation scores are displayed in the bottom right corner of each plot. The bottom row represents Bland-Altman plots displaying bias (solid horizontal line) and imprecision. Bias and imprecision scores are noted on the top half of each plot. Interrupted horizontal lines are 95% limits of agreement

Table 1 Correlation and agreement for HarVI postoperative PSV and PSV ratio predictions subdivided by percent stenosis and presence of downstream disease (PAD distal to investigated stenosis in the SFA)

Cohort	Peak systolic velocity		Peak systolic velocity ratio		Number of lesions
	Corr.	Bias ± Imprec	Corr.	Bias ± Imprec.	
All	0.86	-14.4 ± 14.5	0.93	-0.01 ± 0.40	7
Stenosis degree					
≤50%	0.88	-19.9 ± 14.6	0.93	-0.26 ± 0.27	4
>50%	0.95	-7.0 ± 12.9	0.99	0.33 ± 0.25	3
Presence of downstream disease					
Yes	0.33	-15.7 ± 18.6	0.85	0.06 ± 0.53	4
No	0.97	-12.6 ± 10.0	0.99	-0.10 ± 0.19	3

Discussion

The primary objective of this work was to establish a clinically viable framework for real-time preoperative planning in patients with CLTI, with a focus on SFA lesions. Using a retrospective cohort of seven patients as a pilot study, we demonstrated that the HarVI framework can accurately predict postoperative PSV compared to ground-truth DUS measurements. This pipeline was made feasible through a two-phase approach: (1) preoperative boundary condition tuning at the patient-specific level, and (2) training HarVI, an ML-accelerated framework, to efficiently sample the intervention state space for real-time intervention planning. Once trained, HarVI returned postoperative flow predictions with millisecond-level query times, achieving an average query time of just 34 milliseconds per scenario. We validated both preoperative and postoperative PSV predictions against clinical DUS data, confirming the accuracy of the framework in capturing hemodynamic changes after revascularization. To further contextualize HarVI predictions within established clinical practice, we examined model outputs relative to commonly used duplex ultrasound surveillance thresholds.

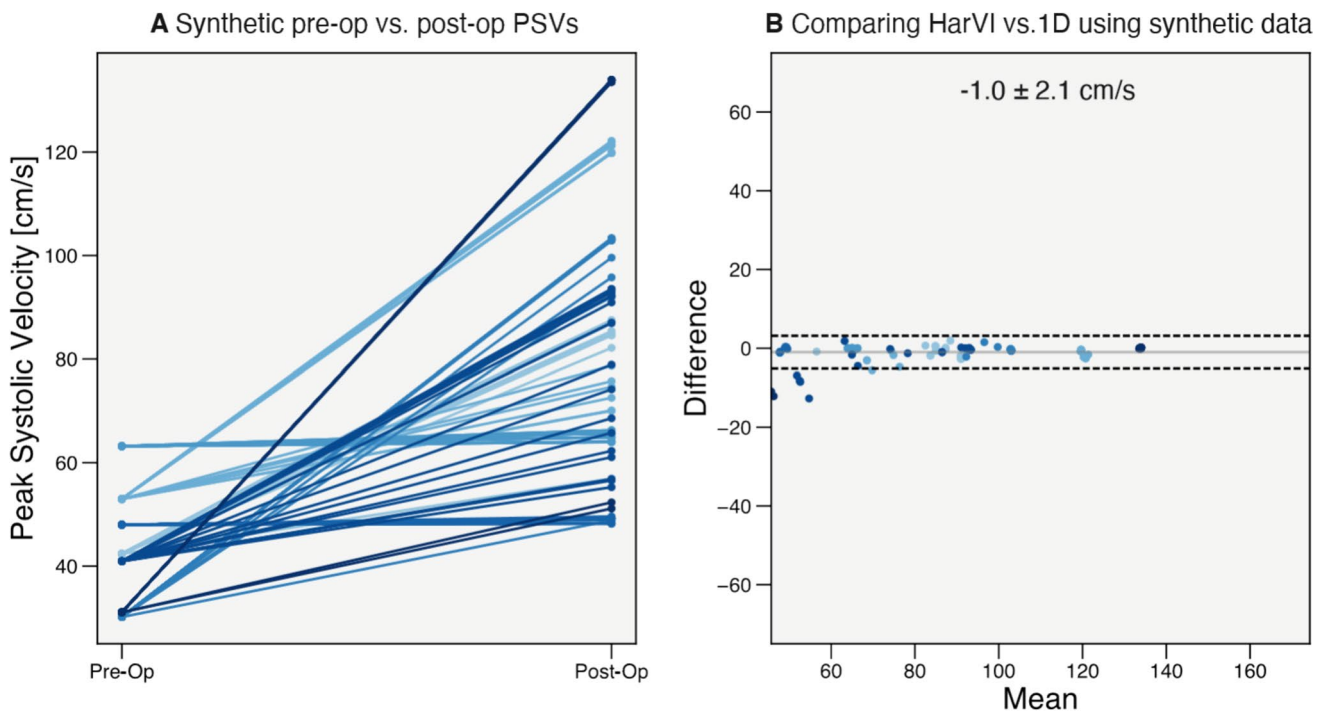


Fig. 9 Comparison of HarVI and 1D CFD in synthetic endovascular interventions. For each base patient, we perform 25 random interventions that are not part of the training set and the testing set used for validation with measured DUS. **A** PSV line plots showing the differ-

ence between pre-op and post-op. **B** Bland-Altman plot comparing HarVI versus 1D CFD. Each shade of blue corresponds to a set of synthetic interventions from one patient




Fig. 10 Time-to-solution for different aspects of the HarVI framework. (Left) Sampling refers to the time taken to run 1D CFD simulations to sample the intervention state space. (Middle) After these runs are completed, the outputs are used as inputs to train our machine

learning models to interpolate the intervention parameter space. (Right) Timing to perform one inference in response to a geometric modification

In this cohort of seven patients, one case (~14%) exhibited a postoperative PSVr greater than 3.5 (3.782), an accepted criterion associated with increased risk of hemodynamically significant restenosis. HarVI correctly

classified this case above the clinical cutoff, while all remaining patients were predicted below the threshold.

Predicted postoperative PSV values similarly fell on the appropriate side of the commonly referenced 30 cm/s benchmark used to assess adequacy of revascularization.

Although this study is not powered to prescribe specific intra-operative actions, these results demonstrate that HarVI outputs can be directly interpreted using existing duplex-based criteria, supporting their use for real-time assessment of whether a revascularization result is hemodynamically adequate. Larger cohorts will be required to quantify misclassification rates and to link intra-operative predictions to downstream management decisions. Additionally, HarVI's predictions closely matched full 1D CFD simulations across a synthetic space of 175 revascularization scenarios, reinforcing its robustness and potential for broader clinical deployment.

HarVI is designed for intra-operative decision support, where the goal is to rapidly compare candidate revascularization strategies based on their immediate hemodynamic impact and patency. We therefore focus on peak systolic velocity (PSV) and PSV ratio (PSVr) because they are clinically adopted, interpretable, and guideline linked markers used to assess flow-limiting stenosis and post-intervention surveillance in CLTI (e.g., $PSV > 30$ cm/s and $PSVr > 3.5$) [11]. Predicting longer-term outcomes such as restenosis or limb salvage will require integrating additional biological, mechanical, and clinical factors beyond conduit-level hemodynamics and is an important direction for future work.

Capturing accurate preoperative hemodynamics is foundational to HarVI's predictive power. A number of strategies exist for tuning boundary conditions, including resistance–radius relationships based on Murray's Law [45], which have shown success for coronary arteries [23–25], as well as patient-averaged parameters [25] and automated tuning approaches [46–52]. However, no clear consensus has emerged on which method is best suited for femoral arteries. Our results showed that Murray's Law-based approaches failed to sufficiently capture preoperative SFA hemodynamics, likely due to flow regime differences between coronary and femoral systems [53]. To address this gap, we employed a genetic algorithm to infer Windkessel parameters from available clinical measurements at the preoperative stage. This method not only offered improved fidelity over rule-based assumptions but also leveraged distributed computing for rapid optimization [54, 55]. While the genetic algorithm successfully recovered preoperative hemodynamics, running thousands of simulations for each patient intra-operatively remains impractical for widespread clinical translation. Fortunately, we observed that our seven patients shared similar anatomy and physiology, enabling the use of patient-averaged boundary conditions as a more clinically feasible first-line approach. This method provided accurate preoperative

physiology, effectively reducing computational burden without sacrificing fidelity. While not a universal solution, prior modeling studies have demonstrated success using patient averaged boundary conditions [56], suggesting that this approach holds promise for broader applicability, including cases involving novel phenotypes. As the available patient population expands, phenotype clustering could further enhance performance by enabling more refined subgrouping into stenosis-specific phenotypic bins, thereby increasing the clinical utility, generalizability, and adaptability of the framework.

The second prerequisite for HarVI was the accurate prediction of postoperative hemodynamics following geometric modifications introduced during revascularization. In this study, we used a pressure inlet boundary condition and two-element Windkessel models, holding them constant from the preoperative to the postoperative stage. This setup assumes that central pressure remains stable and that the primary pressure drop occurs locally at the SFA lesion, allowing increased flow into the system once revascularization is performed. Similar assumptions have been adopted for other anatomies transitioning from preoperative to postoperative states [57]. For instance, Antonuccio et al. [56] showed that fine-tuning outlet Windkessel models was unnecessary for post-stenting hemodynamics, provided the main flow-limiting lesion resided within the simulated domain rather than further downstream. A factorial sensitivity analysis was conducted to explore the impact of maintaining unchanged boundary conditions from the preoperative to postoperative state, as shown in Fig. 11 in supplementary. Heart rate was varied between 65 and 120 beats per minute to reflect a physiologically plausible range, and distal resistance was scaled by $\pm 50\%$ relative to the patient-averaged values to encompass the tuned resistances observed across the cohort. The largest deviations occurred at resistance scaling factors of 0.50 and 1.50 and at the lowest and highest heart rate values. Across all cases, these variations resulted in a maximum deviation of approximately 4.5 cm/s in distal preoperative PSV. While deviations near critical diagnostic thresholds have the potential to influence surgical planning, the magnitude of these changes is comparable to the reported measurement uncertainty in duplex ultrasound assessments [58]. Because proximal and distal PSV values scale proportionally under these variations, the PSVr remains unchanged as a ratio metric. Other studies have similarly highlighted that pre- to post-stenting physiological changes largely occur at the lesion itself [59, 60]. Our results support these findings by showing that postoperative 1D CFD models, using unchanged inlet and outlet boundary conditions, reliably reproduced postoperative PSV measurements. In particular, we observed that our boundary-condition strategy is most effective when the SFA lesion is the dominant source of pressure drop. Indeed, the highest accuracy in predicted PSV

was noted for cases with stenosis $> 50\%$ and no downstream disease, underscoring the validity of our approach in scenarios where the lesion is the primary hemodynamic obstacle.

After establishing that 1D CFD models can accurately capture preoperative and postoperative femoral hemodynamics, we next sampled the intervention state space to predict the hemodynamic response of revascularization in real time. HarVI successfully replicated the hemodynamic results of explicit 1D CFD simulations in rigorous, worst-case synthetic verification studies, as well as in comparisons against clinically measured postoperative PSVs. Although this work focuses on femoral artery disease, HarVI is broadly generalizable. In prior research, it has also demonstrated strong agreement with explicit 1D CFD models in CAD, accurately predicting fractional flow reserve, wall shear stress, and flow rate [5]. The core strength of HarVI lies in its ability to efficiently sample probable intervention scenarios using LHS and then to interpolate those outcomes with computationally efficient ML models. Our previous findings demonstrated the ability of HarVI to capture cardiac cycle-averaged metrics, and in this study, we show that the framework not only extends to a different vascular territory (femoral circulation) but also reproduces pulsatile flow behavior that is closely aligned with clinical measurements. A notable advance in this work is HarVI's demonstrated ability to reproduce pulsatile behavior and clinical PSV values in a new vascular territory, showcasing its adaptability across anatomies and disease types.

HarVI enables users to query changes in stenosis diameter and length in order to examine the surrogate effects of different stent dimensions on postoperative vessel geometry and resulting hemodynamics. In this way, the framework can be used intraoperatively to compare candidate repair scenarios, such as shorter versus longer stents or whether extending an intervention is likely to provide additional hemodynamic benefit. Because revascularization goals vary across patients, owing to factors such as stent availability and vessel integrity, HarVI is intended to allow clinicians to efficiently test multiple feasible intervention strategies and identify the option that best aligns with patient-specific therapeutic objectives.

Although additional flow metrics such as WSS can be approximated in one-dimensional models and may provide complementary insight, WSS is not currently incorporated into routine PAD intervention guidelines as a gold-standard hemodynamic metric. Accordingly, PSV and PSVr are used as the primary outputs for intra-operative risk stratification in this work. Future extensions of the framework may incorporate WSS as a supplementary output alongside established metrics such as PSV and PSVr.

More broadly, because maximizing vessel patency remains a central aim of revascularization procedures, HarVI provides a mechanism for rapid, iterative assessment

of alternative diameter and length dilation strategies to support stent selection and procedural planning.

While this study takes foundational steps toward real-time digital twin intervention planning in PAD, several challenges must be addressed to support full clinical translation. First, although our assumption that SFA lesions are the main contributor to the generally recovered postoperative flow pressure drop, it underestimated PSV when the lesions were $\leq 50\%$ or when downstream disease was present. These findings suggest that boundary conditions may need refinement—especially for smaller or multi-level lesions—and also indicate the need for expanding our computational domain to incorporate regions beyond the SFA (e.g., popliteal or tibial arteries). In more complex multi-lesion scenarios, future work should also extend HarVI to jointly simulate multiple vascular territories and validate against postoperative ground truths. A second limitation is that, by construction, the 1D model only recovers flow rate at the level of each vessel segment rather than locally within a stenosis. Because we collapse a lesion into a single pressure drop term and assume a constant radius across regular segments, rather than explicitly tapering the geometry, the model enforces a constant flow rate (and by extension velocity if walls are rigid) through each segment—even though the pressure distribution shows the characteristic drop. This segment-level approximation obscures velocity spikes and local hemodynamic variations across a stenosis, which may be critical for accurately predicting wall shear stress and focal flow dynamics. Incorporating explicit geometric tapering, subdividing vessels into smaller segments around stenoses, or coupling with 3D models where we model only the stenosis in 3D could help capture these localized effects. A third limitation is that we have only examined feasibility of HarVI for intra-operative use but have not yet fully automated our geometry reconstruction workflow. Streamlining geometry processing will be critical to seamlessly integrating HarVI into the clinical setting. Furthermore, this study evaluated a relatively small cohort of seven patients, each with a broadly similar SFA anatomy. Although we observed differences in outcomes based on the severity of stenosis and subsequent disease, a larger and more diverse cohort is needed to confirm generalizability. Such an expansion would also enable the development of a patient averaged boundary condition repository, allowing new patients to be matched with suitable baseline conditions based on anatomical and physiological similarity. A limitation of this retrospective study is the absence of patient demographic information due to the use of a fully de-identified dataset.

Lastly, while our genetic algorithm served as a fallback for cases where patient averaged approaches were inadequate, fully validating this workflow in a larger population is an essential next step. By addressing these points, boundary condition refinement, multi-lesion predictions,

automated geometry processing, and broader patient recruitment, HarVI can move closer to routine clinical implementation for real-time peripheral revascularization planning.

In conclusion, this study establishes the first application of HarVI to PAD and validates its potential for patient-specific real-time planning in CLTI. By combining 1D CFD with ML, HarVI provides clinically accurate post-operative hemodynamic predictions within milliseconds and requires only a finite set of precomputed simulations. These results lay the groundwork for more efficient, scalable, and clinically translatable digital twins for peripheral vascular disease. With future work focused on boundary condition adaptation, automation, and multi-lesion modeling, HarVI could support a paradigm shift in personalized vascular care—transforming how revascularization strategies are selected during intervention for patients with advanced PAD.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s13239-026-00825-1>.

Acknowledgements The authors thank Wendy Wu, Ayman Yousef, and Nusrat Sadia Khan for fruitful discussions. Computing support for this work came from the Duke Compute Cluster.

Author Contributions All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by Cyrus Tanade and Christopher Jensen. The first draft of the manuscript was written by Cyrus Tanade and all authors commented on previous versions of the manuscript. Rebuttal changes and analyses were performed by Guinevere Ferreira. All authors read and approved the final manuscript.

Funding The work of Cyrus Tanade was supported by the National Science Foundation Graduate Research Fellowship under Grant No. NSF GRFP DGE 1644868. The work of Amanda Randles was supported by the National Institute On Aging of the NIH under Award Number DP1AG082343 and the Duke/Duke-NUS Research Collaboration Pilot Project program. The content does not necessarily represent the official views of the NSF or NIH.

Data Availability Anonymized, de-identified data is available upon reasonable request. Please contact the corresponding author.

Declarations

Competing interests The authors have no relevant financial or non-financial interest/Competing interest to disclose.

Ethical Approval This research study was conducted retrospectively from data obtained for clinical purposes. This study was determined to be IRB exempt by the IRB.

Informed Consent This study involved a secondary analysis of fully deidentified data and did not entail any direct contact with human subjects.

Consent for Publication Note that because the data are fully de-identified and the study is exempt, no additional consent to publish individual data was required.

References

1. Tanade, C., et al. Establishing the longitudinal hemodynamic mapping framework for wearable-driven coronary digital twins. *NPJ Digit. Med.* 7(1):236, 2024.
2. Tanade, C., et al. Cloud computing to enable wearable-driven longitudinal hemodynamic maps. *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.* pp. 1–14, 2023.
3. Hernandez-Boussard, T., et al. Digital twins for predictive oncology will be a paradigm shift for precision cancer care. *Nat. Med.* 27(12):2065–2066, 2021.
4. Tanade, C., and A. Randles. HarVI: real-time intervention planning for coronary artery disease using machine learning. In: *International Conference on Computational Science*, Cham: Springer, 2024, pp. 48–62.
5. Tanade, C., and A. Randles. Real-time virtual intervention for simple and serial coronary artery disease using the HarVI framework. *J. Comput. Sci.* 2025:102570, 2025.
6. Sankaran, S., et al. Physics driven real-time blood flow simulations. *Comput. Methods Appl. Mech. Eng.* 364:112963, 2020.
7. Sonck, J., et al. Clinical validation of a virtual planner for coronary interventions based on coronary CT angiography. *Cardiovasc. Imaging.* 15(7):1242–1255, 2022.
8. Aday, A. W., and K. Matsushita. Epidemiology of peripheral artery disease and polyvascular disease. *Circ. Res.* 128(12):1818–1832, 2021.
9. Nehler, M. R., et al. Epidemiology of peripheral arterial disease and critical limb ischemia in an insured national population. *J. Vasc. Surg.* 60(3):686–695, 2014.
10. Vemulapalli, S., et al. Leveraging computational fluid dynamics for next-generation preoperative planning in vascular surgery. In: *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, IEEE, 2024, pp. 1–4.
11. Conte, M. S., et al. Global vascular guidelines on the management of chronic limb-threatening ischemia. *Eur. J. Vasc. Endovasc. Surg.* 58(1):S1–S109, 2019.
12. Khan, M. U. N., et al. Predictors of limb loss despite a patent endovascular-treated arterial segment. *J. Vasc. Surg.* 49(6):1440–1445, 2009.
13. Kim, W., and D. Choi. Treatment of femoropopliteal artery in-stent restenosis. *Korean Circ. J.* 48(3):191–197, 2018.
14. Goldberg, S. L., et al. Predictors of diffuse and aggressive in-stent restenosis. *J. Am. Coll. Cardiol.* 37(4):1019–1025, 2001.
15. Russo, R. J., P. D. Silva, and M. Yeager. Coronary artery overexpansion increases neointimal hyperplasia after stent placement in a porcine model. *Heart.* 93(12):1609–1615, 2007.
16. Wang, K. A. I., et al. Overlapping coronary stents result in an increased neointimal hyperplasia: insight from a porcine coronary stent model. *J. Interv. Cardiol.* 13(3):173–177, 2000.
17. Ninno, F., et al. Modelling lower-limb peripheral arterial disease using clinically available datasets: impact of inflow boundary conditions on hemodynamic indices for restenosis prediction. *Comput. Methods Program Biomed.* 251:108214, 2024.
18. Shahrolakmar, U. Z., M. N. Omar, and N. H. Johari. Brief review on recent advancement of computational analysis on hemodynamics in peripheral artery disease. In: *International Conference on Mechanical Engineering Research*, Cham: Springer, 2021, pp. 555–572.
19. Hossain, S. S., et al. Magnetic resonance imaging-based computational modelling of blood flow and nanomedicine deposition in patients with peripheral arterial disease. *J. R. Soc. Interface.* 12(106):20150001, 2015.
20. Gounley, J., et al. Computing the ankle-brachial index with parallel computational fluid dynamics. *J. Biomech.* 82:28–37, 2019.

21. van de Velde, L., et al. Computational fluid dynamics for the prediction of endograft thrombosis in the superficial femoral artery. *J. Endovasc. Ther.* 30(4):615–627, 2023.
22. McGah, P. M., et al. Hemodynamic conditions in a failing peripheral artery bypass graft. *J. Vasc. Surg.* 56(2):403–409, 2012.
23. Vardhan, M., et al. Diagnostic performance of coronary angiography derived computational fractional flow reserve. *J. Am. Heart Assoc.* 13(13):e029941, 2024.
24. Tanade, C., et al. Global sensitivity analysis for clinically validated 1d models of fractional flow reserve. In: 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), IEEE, 2021, pp. 4395–4398.
25. Tanade, C., et al. Analysis identifying minimal governing parameters for clinically accurate in silico fractional flow reserve. *Front. Med. Technol.* 4:1034801, 2022.
26. Khan, T. H., F. A. Farooqui, and K. Niazi. Critical review of the ankle brachial index. *Curr. Cardiol. Rev.* 4(2):101–106, 2008.
27. Potier, L., et al. Use and utility of ankle brachial index in patients with diabetes. *Eur. J. Vasc. Endovasc. Surg.* 41(1):110–116, 2011.
28. Aboynans, V., et al. Measurement and interpretation of the ankle-brachial index: a scientific statement from the American Heart Association. *Circulation.* 126(24):2890–2909, 2012.
29. Gao, M., et al. Optimal ultrasound criteria for grading stenosis of the superficial femoral artery. *Ultrasound Med. Biol.* 44(2):350–358, 2018.
30. Elhossieny Selim, N. M. E., et al. Role of ankle peak systolic velocity as a hemodynamic predictor following infrainguinal arterial angioplasty. *Egypt. J. Surg.* 40(3):1–10, 2022.
31. Idu, M. M., et al. Impact of a color-flow duplex surveillance program on infrainguinal vein graft patency: a five-year experience. *J. Vasc. Surg.* 17(1):42–53, 1993.
32. Calligaro, K. D., et al. Should duplex ultrasonography be performed for surveillance of femoropopliteal and femorotibial arterial prosthetic bypasses? *Ann. Vasc. Surg.* 15(5):520–524, 2001.
33. Stoner, M. C., et al. Reporting standards of the Society for vascular surgery for endovascular treatment of chronic lower extremity peripheral artery disease. *J. Vasc. Surg.* 64(1):e1–e21, 2016.
34. Shrikhande, G. V., et al. Determining criteria for predicting stenosis with ultrasound duplex after endovascular intervention in infrainguinal lesions. *Ann. Vasc. Surg.* 25(4):454–460, 2011.
35. Baril, D. T., et al. Duplex criteria for determination of in-stent stenosis after angioplasty and stenting of the superficial femoral artery. *J. Vasc. Surg.* 49(1):133–139, 2009.
36. Kawarada, O., et al. Peak systolic velocity ratio derived from quantitative vessel analysis for restenosis after femoropopliteal intervention: a multidisciplinary review from endovascular Asia. *Cardiovasc. Interv. Therap.* 35:52–61, 2020.
37. Randles A. et al. Massively parallel models of the human circulatory system. *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.* pp. 1–11, 2015.
38. Gornik, H. L., et al. 2024 ACC/AHA/AACVPR/APMA/ABC/SCAI/SVM/SVN/SVS/SIR/VESS guideline for the management of lower extremity peripheral artery disease: a report of the American College of Cardiology/American Heart Association Joint Committee on Clinical Practice Guidelines. *J. Am. Coll. Cardiol.* 83(24):2497–2604, 2024.
39. James Chen, S., and J. D. Carroll. 3-D reconstruction of coronary arterial tree to optimize angiographic visualization. *IEEE Trans. Med. Imaging.* 19(4):318–336, 2000.
40. Green, N. E., et al. Angiographic views used for percutaneous coronary interventions: a three-dimensional analysis of physician-determined vs. computer-generated views. *Catheter. Cardiovasc. Interv.* 64(4):451–459, 2005.
41. Feiger, B., et al. Determining the impacts of venoarterial extracorporeal membrane oxygenation on cerebral oxygenation using a one-dimensional blood flow simulator. *J. Biomech.* 104:109707, 2020.
42. McKay, M. D., R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics.* 42(1):55–61, 2000.
43. Iman, R. L., J. C. Helton, and J. E. Campbell. An approach to sensitivity analysis of computer models: part I—Introduction, input variable selection and preliminary variable assessment. *J. Qual. Technol.* 13(3):174–183, 1981.
44. Beckman, J. A., P. A. Schneider, and M. S. Conte. Advances in revascularization for peripheral artery disease: revascularization in PAD. *Circ. Res.* 128(12):1885–1912, 2021. <https://doi.org/10.1161/CIRCRESAHA.121.318261>.
45. Sherman, T. F. On connecting large vessels to small. The meaning of Murray's law. *J. Gen. Physiol.* 78(4):431–453, 1981.
46. Spilker, R. L., and C. A. Taylor. Tuning multidomain hemodynamic simulations to match physiological measurements. *Ann. Biomed. Eng.* 38:2635–2648, 2010.
47. Richter, J., et al. Bayesian Windkessel calibration using optimized zero-dimensional surrogate models. *Philos. Trans. A.* 383(2292):20240223, 2025.
48. Nolte, D., and C. Bertoglio. Inverse problems in blood flow modeling: a review. *Int. J. Numer. Methods Biomed. Eng.* 38(8):e3613, 2022.
49. Itu, L., et al. A parameter estimation framework for patient-specific hemodynamic computations. *J. Comput. Phys.* 281:316–333, 2015.
50. Itu, L., et al. Personalized blood flow computations: a hierarchical parameter estimation framework for tuning boundary conditions. *Int. J. Numer. Methods Biomed. Eng.* 33(3):e02803, 2017.
51. Fevola, E., et al. An optimal control approach to determine resistance-type boundary conditions from in-vivo data for cardiovascular simulations. *Int. J. Numer. Methods Biomed. Eng.* 37(10):e3516, 2021.
52. MacDonald Black, S., et al. Calibration of patient-specific boundary conditions for coupled CFD models of the aorta derived from 4D Flow-MRI. *Front. Bioeng. Biotechnol.* 11:1178483, 2023.
53. Fossan, F. E., et al. Uncertainty quantification and sensitivity analysis for computational FFR estimation in stable coronary artery disease. *Cardiovasc. Eng. Technol.* 9:597–622, 2018.
54. Liu, H., et al. The application of non-linear flow resistance in cerebral artery: compared with windkessel model based on genetic algorithm. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE, 2019, pp. 2285–2288.
55. El Zahab, Z., E. Divo, and A. Kassab. Minimisation of the wall shear stress gradients in bypass grafts anastomoses using meshless CFD and genetic algorithms optimisation. *Comput. Methods Biomed. Eng.* 13(1):35–47, 2010.
56. Antonuccio, M. N., et al. Effects of uncertainty of outlet boundary conditions in a patient-specific case of aortic coarctation. *Ann. Biomed. Eng.* 49(12):3494–3507, 2021.
57. Gamage, P. T., et al. Hemodynamic alternations following stent deployment and post-dilation in a heavily calcified coronary artery: in silico and ex-vivo approaches. *Comput. Biol. Med.* 139:104962, 2021.
58. Franz, R., et al. Accuracy of duplex ultrasonography in estimation of severity of peripheral vascular disease. *Int. J. Angiol. Off. Publ. Int. Coll. Angiol. Inc.* 22:155, 2013. <https://doi.org/10.1055/S-0033-1336830>.
59. Williams, A. R., et al. Local hemodynamic changes caused by main branch stent implantation and subsequent virtual side branch balloon angioplasty in a representative coronary bifurcation. *J. Appl. Physiol.* 109(2):532–540, 2010.

60. Gundelwein, L., et al. Personalized stent design for congenital heart defects using pulsatile blood flow simulations. *J. Biomech.* 81:68–75, 2018.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.