



Establishing the longitudinal hemodynamic mapping framework for wearable-driven coronary digital twins



Cyrus Tanade¹, Nusrat Sadia Khan¹, Emily Rakestraw¹, William D. Ladd¹, Erik W. Draeger² & Amanda Randles¹✉

Understanding the evolving nature of coronary hemodynamics is crucial for early disease detection and monitoring progression. We require digital twins that mimic a patient's circulatory system by integrating continuous physiological data and computing hemodynamic patterns over months. Current models match clinical flow measurements but are limited to single heartbeats. To this end, we introduced the longitudinal hemodynamic mapping framework (LHMF), designed to tackle critical challenges: (1) computational intractability of explicit methods; (2) boundary conditions reflecting varying activity states; and (3) accessible computing resources for clinical translation. We show negligible error (0.0002–0.004%) between LHMF and explicit data of 750 heartbeats. We deployed LHMF across traditional and cloud-based platforms, demonstrating high-throughput simulations on heterogeneous systems. Additionally, we established LHMF_C, where hemodynamically similar heartbeats are clustered to avoid redundant simulations, accurately reconstructing longitudinal hemodynamic maps (LHMs). This study captured 3D hemodynamics over 4.5 million heartbeats, paving the way for cardiovascular digital twins.

Understanding the evolving physiological state of a patient, particularly during their daily activities, is key to early detection and monitoring of disease. In this context, the creation of a human digital twin has become a primary objective in computational biomedicine. Recent advances have brought this goal closer to reality¹, yet these models still face limitations in capturing extended timescales. Digital twins are intricate, personalized, computer simulations of biophysical processes that rely on observational data to accurately forecast disease indicators². These simulations refer to modeling 3D blood flow in realistic patient anatomy. The growing prevalence of wearable technology offers a unique opportunity to continuously refine these digital twins to mirror the changing physical state of an individual.

A prime area of focus is the development of 3D blood flow models for cardiovascular diseases. These digital twins provide a detailed examination of hemodynamics tailored to an individual's specific anatomy. Significant strides have been made in using these models to assess complex 3D hemodynamic biomarkers^{3–5}. Nevertheless, fully capturing blood flow in 3D is computationally demanding, and as a result, current state-of-the-art methods are only able to reconstruct a few seconds of blood flow^{3,6}. Considering that cardiovascular diseases evolve over months to years, improving our ability to track blood flow over longer periods with 3D digital twins

could reveal new biomarkers, potentially revolutionizing standard of care practices.

In coronary artery disease, revascularization strategies have become routine clinical procedures to treat the full range of lesion presentations⁷. However, complications such as in-stent restenosis (ISR)—the gradual reduction in lumen diameter due to arterial damage following revascularization procedures⁸, occur in around 10% of all percutaneous interventions 3–6 months after the initial procedure⁹. However, over 50% of ISR cases are only treated >2 years after the baseline intervention because patients are often asymptomatic⁹. ISR screening currently depends on performing angiographic checks at discrete time points over months or when patients present with clinical symptoms. Unsurprisingly, estimating the true incidence of ISR is challenging, as many patients are likely to fall between the cracks. The burden of ISR is substantial, as 1 in 4 patients presents with acute myocardial infarction due to the ISR lesion⁹. In the US, the prevalence of interventions for ISR has remained relatively unchanged in the past two decades despite significant advances in stent technology⁹. Employing digital twins would allow us to track a patient's hemodynamic profile not just at discrete time points as in routine clinical practice but continuously as complications gradually progress. For a disease that progresses over timescales in months to years, analyses at single points in time are likely to be less

¹Department of Biomedical Engineering, Duke University, Durham, NC, 27708, USA. ²Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA. ✉e-mail: amanda.randles@duke.edu

representative and predictive of a patient's condition than a longitudinal assessment of hemodynamic biomarkers.

Computational fluid dynamics (CFD), the branch of fluid mechanics that numerically solves flow equations, has been shown to accurately predict a wide selection of personalized 3D hemodynamic biomarkers³. Nonetheless, capturing longitudinal 3D hemodynamics poses computational challenges for even state-of-the-art CFD methods^{3,6,10–15}. While CFD methods have been optimized to utilize some of the world's most powerful supercomputers^{11,16}, simulating even one month of activity would require decades to centuries of computing time. Lighter models make simplifying assumptions about patient anatomy^{17,18}; however, they fail to capture the full complexity of 3D hemodynamics. Although reduced-order models and machine learning-based approaches do not fully capture 3D blood flow¹⁹, they still hold the potential to identify novel biomarkers. In this study, our objective was to fully capture the intricacies of blood flow dynamics to broaden our search for new longitudinal biomarkers. Although these reduced-order models (0D and 1D) have been shown to accurately recapitulate some scalar metrics in some coronary indications, 3D models are required to cover the full gamut of disease manifestations and complexities. 0D models assume that spatially varying properties are lumped into discrete components such as resistors²⁰. This approach assumes that blood flow is steady, unidirectional, and axisymmetric. However, these assumptions are broken down in stenoses where blood flow is complex and can form complex vortices^{21–23}. 1D models overcome some of these assumptions, where blood flow is assumed to travel only along the longitudinal direction of flow and neglect non-axial components²⁴. The assumptions of the 1D model are valid in healthy and straight vessels. At a stenosis, pressure loss occurs via viscous, turbulent, and inertial components of blood flow and typically necessitates empirical models for accuracy. These empirical pressure loss models of stenoses have been shown to be clinically accurate, especially in focal lesions¹⁸, but explicit models for complex disease, such as serial or bifurcation lesions, still remain an open question. It is important to be able to account for all coronary indications. For example, bifurcation lesions represent 20% of all cases²⁵ and cause significant mortality and morbidity. 3D models do not make simplifying assumptions on anatomy and are able to capture the influence of complex geometry on local flow field patterns. While reduced-order models have shown their accuracy in clinically-relevant metrics such as fractional flow reserve^{17,18,26}, they fail to capture the spatial heterogeneity of complex metrics and can limit patient eligibility. Therefore, the aim here is to develop a framework for identifying novel biomarkers, focusing on 3D metrics to avoid constraining potential discoveries.

To our knowledge of research currently published in the literature, the longest patient-specific 3D simulations have only recovered less than 1 minute of biological time (Supplementary Table 1)^{6,27,28}. Efforts have been made to extend the temporal domain that could be feasibly analyzed, primarily using numerical^{12,13,29,30} or deep learning methods^{31–33}. However, these approaches are still intractable and, at best, could capture only minutes of hemodynamics. The computational intractability issue is twofold: even if it is possible to evaluate months of 3D hemodynamics promptly, relying on large supercomputers is a barrier to clinical translation. Computing resources must be accessible on demand and generalizable to diverse hardware. To enable longitudinal hemodynamic studies, we require advances to extend the duration of biological time we can efficiently capture, while also considering accessibility for eventual clinical translation.

In this study, we establish the longitudinal hemodynamic mapping framework (LHMF) to capture personalized 3D blood flow on the time scale of months (Fig. 1). We hypothesize that for any given longitudinal patient-specific hemodynamic data, deconstructing time order into discrete units of work that could all be deployed in parallel would enable accurate reconstruction of an explicit 3D hemodynamics simulation using established methods. Multimodal wearable data are used to continuously inform the 3D coronary digital twin of a patient's evolving physical state. High-throughput simulations are performed using heterogeneous computing resources, particularly cloud computing, to demonstrate the accessibility of the LHMF.

To our knowledge, this is the first study to combine digital twins of 3D blood flow, wearable devices, and cloud computing to enable the study of months of complex hemodynamic biomarkers. Here, we perform verification studies comparing LHMF with explicit ground-truths in a 750 heartbeat problem—that in of itself exceeds literature records by over an order of magnitude (Supplementary Table 1), and evaluate common 3D hemodynamic biomarkers over 4.5 million heartbeats or 1.5 months of activity to identify potential sites of disease progression. In this paper, we take the first critical steps to establish the computational foundations needed to investigate and eventually uncover novel temporal biomarkers that could be more predictive of coronary disease progression than currently established single-heartbeat analyses.

Results

We established a framework in this work, the LHMF, to capture 3D longitudinal hemodynamic behavior. A high-level overview of the LHMF is provided in Fig. 1. The framework begins with acquiring continuous heart rate signals from wearable devices (Fig. 1a). Heart rate signals were used to personalize inlet coronary velocity waveforms on the fly (see sub-sections “Driving coronary digital twins from continuous wearable data” of “Results” and “Wearable-driven boundary conditions for hemodynamic simulations” of “Methods”). Image-based anatomical models were reconstructed from pairs of coronary angiograms, and the quality of reconstruction was evaluated by expert interventional cardiologists (Fig. 1b). At the time of imaging, we also acquired clinical measurements, such as arterial pressure, cardiac output, flow dominance, and hematocrit to tune boundary conditions to a patient-specific level (Fig. 1c). These measurements were collected in one instance under clinical settings. We employed pulsatile velocity waveforms at the inlet and two-element Windkessel models at the outlet to simulate microvascular resistance (see sub-section “Personalized coronary digital twins” of “Methods”). To capture longitudinal hemodynamics, we deconstructed a long sequence of inlet coronary velocity waveforms into a representative spanning set of simulations (Fig. 1d) (see sub-sections “Deconstructing temporal dependence to enable parallelism in time” of “Results” and “Clustering algorithm to minimize running redundant simulations” of “Methods”). 3D blood flow simulations were performed using HARVEY, a massively parallel fluid dynamics solver^{11,16}. Simulation outputs for the representative waveforms were used to reconstruct the complete temporal domain and thereby capture longitudinal hemodynamic behavior (Fig. 1e).

Driving coronary digital twins from continuous wearable data

To enable longitudinal hemodynamic assessment, it was necessary to augment current approaches to parameterize coronary digital twin boundary conditions to capture varying activity states longitudinally (Fig. 1a). Current approaches capture a few cardiac cycles of coronary velocity waveforms at the inlet boundary condition at most. Our goal here was to propose a new method to leverage data from wearable devices to inform a patient's continuously changing inlet velocity profile (Fig. 1c). Coronary velocity waveforms are conventionally derived from representative velocity waveform tracings and scaled to a patient-specific level via cardiac output^{3,10,18,34,35}. Although such waveforms may be clinically accurate on a single-heartbeat scale (as in Fig. 2a) and are generally used in clinical settings, we incorporated instantaneous physiological information to modify the velocity waveforms on demand to reflect the evolving physical states seen through daily activity. Specifically, we integrated continuous heart rate data and electrocardiography (ECG) data into our inlet velocity boundary condition. Using a representative left coronary velocity waveform, continuous heart rate data was used to tune the amplitude and period. ECG data was used to expand or contract the diastolic fraction of the waveform, specifically by modifying the waveform to reach a desired diastolic and systolic fraction in the cardiac cycle as a function of the heart rate recorded from wearables. As the heart rate increases and the patient reaches a more exercise state, the time spent in diastole decreases. Furthermore, we modeled the natural variability in coronary velocity waveforms as seen in vivo³⁶ by introducing randomness

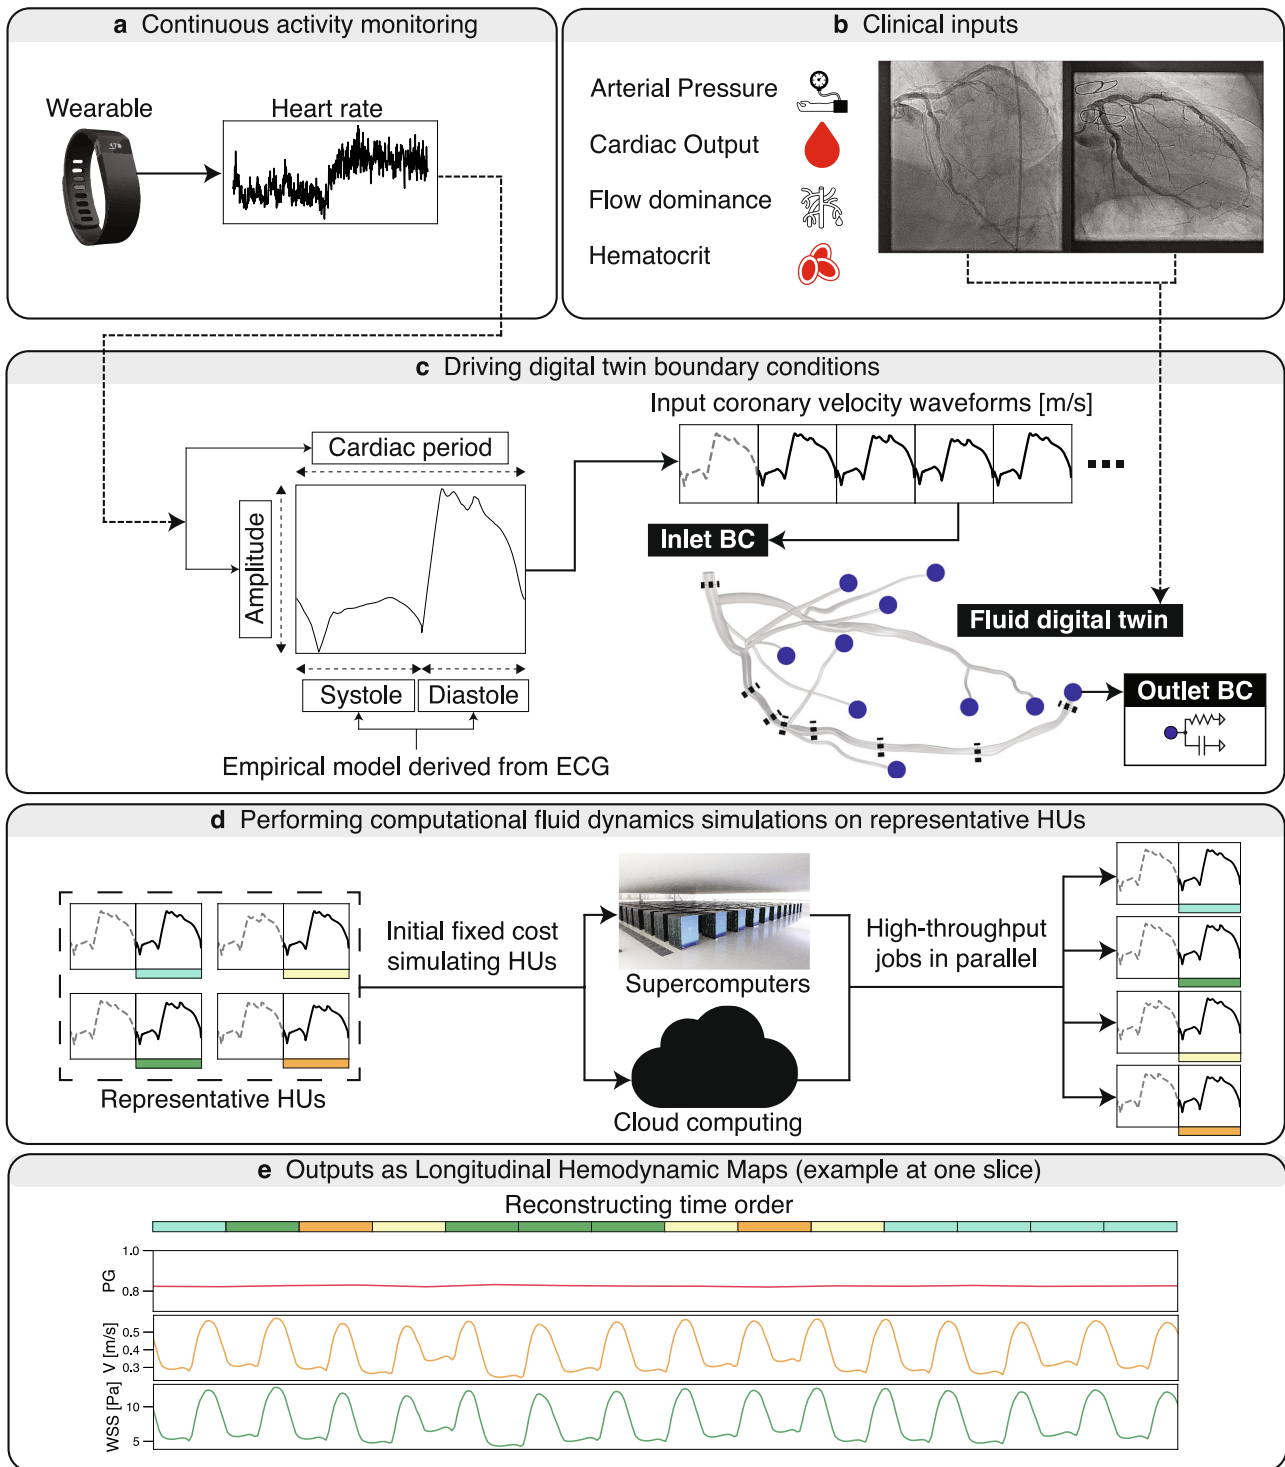


Fig. 1 | Overview of the longitudinal hemodynamic mapping framework (LHMF). This framework uses **a** data from wearable devices and **b** clinical measurements to create digital twins that reflect patient-specific physiology and anatomy. **c** Data from wearables (heart rate) inform input coronary velocity waveforms, and two-element Windkessel models are used at the outlet boundary condition. Empirical models derived from patient-averaged ECG data are used to inform systolic and diastolic fractions of the cardiac cycle. 3D patient-specific anatomy is reconstructed from pairs of 2D coronary angiograms. Interrupted lines on the coronary tree delineate slice locations used for analysis. **d** LHMF enables fluid simulations to be broken into a finite set of hemodynamic units (HUs) that can efficiently be run on heterogeneous systems, such as traditional supercomputers or

cloud resources. Heartbeats with colored boxes represent the heartbeat of interest, and those with interrupted lines represent pre-cardiac cycles needed for temporal convergence. The different colors represent hemodynamically unique heartbeats. Each HU is its own simulation and can be run in parallel independent of other HUs. **e** Results of HUs are used to reconstruct the long sequence of heartbeats. This approach allows us to visualize spatiotemporal hemodynamic maps, which we are introducing as a patient-specific longitudinal hemodynamic map (LHM). We show an illustrative sequence of pressure gradient (PG), velocity (V), and wall shear stress (WSS) averaged over one slice. Our LHMs can display hemodynamic metrics over the entire 3D anatomy. Figure is adapted from ref. 85; supercomputer illustration is from public domain.

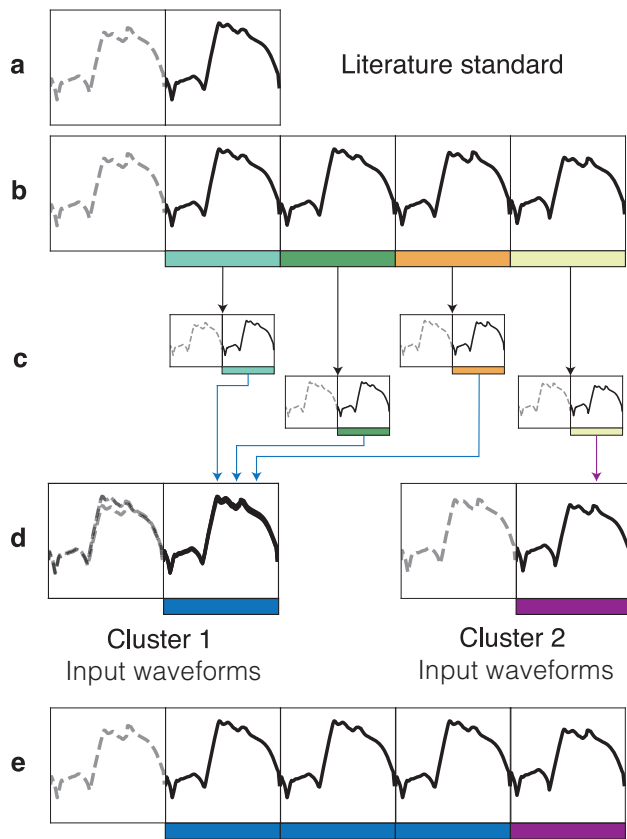


Fig. 2 | Deconstructing temporal dependence using the longitudinal hemodynamic mapping framework (LHMF). Solid lines indicate the heartbeat of interest used for analysis, and interrupted gray lines delineate pre-flow cardiac cycles for temporal convergence. **a** Literature standard for simulating transient flow. **b** Explicit approach for simulating multiple heartbeats. **c** The sequences of heartbeats are deconstructed into hemodynamic units (HUs). HUs consist of the input waveforms of interest and the preceding flow cycles used for temporal convergence. These HUs are simulated in parallel via cloud instances. **d** Simulating all HUs through brute-force leads to many simulations. Instead, capturing representative sets of hemodynamically equivalent HUs can minimize redundancy. Here, we present two resulting clusters of input waveforms deconstructed from a sequence of four heartbeats. **e** Clustered HUs are deployed on cloud instances and used to re-assemble a longitudinal sequence of cardiac cycles. Figure is reproduced from ref. 85.

into pertinent parts of the coronary velocity profile. We share the mathematical details of how we incorporate multimodal wearable data in “Wearable-driven boundary conditions for hemodynamic simulations” of “Methods”.

Deconstructing temporal dependence to enable parallelism in time

Simulating every cardiac cycle sequentially is an intractable approach to capturing 3D longitudinal hemodynamics. One of the salient contributions of this work is how LHMF enables deconstructing the sequential order of heartbeats to perform simulations over time in parallel (Fig. 1d). Our method draws inspiration from DNA genome shotgun sequencing, which reconstructs whole genome sequences by first sequencing shorter constituent DNA fragments in parallel before reassembling them to recover the original DNA sequence in its entirety^{37–39}. In LHMF, an analogous approach is to deconstruct long sequences of heartbeats into a finite set of representative heartbeat sequences, which we define as hemodynamic units (HU) (Fig. 1d). HUs consist of the heartbeat of interest for analysis and the preceding cardiac cycles that ensure temporally converged results. We perform high-resolution simulations on each of these HUs, before reconstructing and reassembling longitudinal patient-specific hemodynamic

behavior (Fig. 1e). Furthermore, while the primary challenge in shotgun sequencing is the efficient and accurate reconstruction of a larger DNA sequence from its constituent fragments, the main challenge in LHMF is to robustly deconstruct a long sequence of heartbeats into a ‘spanning set’ of HUs.

In the field of computational hemodynamics, the standard in the literature is to capture the preceding cardiac cycles (N_{pre}) for convergence and the heartbeat of interest. 3D flow is simulated over $N_{pre} + 1$ beats to ensure that analysis is performed on the converged heartbeat of interest (Fig. 2a). Solutions are considered to have temporally convergent solutions when the percentage difference is within 1% for the hemodynamic variables of interest at the outlet boundaries^{14,40,41}. With this conventional approach, longitudinal simulations could be performed in an explicit manner in which the flow across a long series of heartbeats is simulated continuously, as demonstrated in Fig. 2b. However, implementing this approach in the clinic would be impossible because it would take years to decades of wallclock time just to capture one day of continuous activity, even on a leadership class system.

To address this challenge, we established the LHMF where each heartbeat (with sufficient pre-flow cardiac cycles) is considered an HU and simulated in parallel (Fig. 2c). To verify the LHMF, we performed two experiments. We examined varying heart rates to elucidate whether the number of pre-flow cycles required for convergence differed between heart rates and heart rate transitions and to ensure that the solutions converged temporally for all HUs analyzed. First, we quantified the number of pre-flow cycles needed for temporal convergence for multiple hemodynamic endpoints over four heart rates (46, 91, 128, 176) spanning 6 weeks of continuous wearable data⁴². Namely, we focus on the resting pressure gradient (PG), a metric that is correlated with the lack of oxygen delivery due to atherosclerosis and has been correlated with the location and progression of vascular disease⁴³, blood velocity (V) and wall shear stress (WSS), the friction force of blood acting on the inner walls of blood vessels. Ensuring that the hemodynamics converged temporally was the first key pertinent step in the feasibility of LHMF.

Second, we performed a simulation of 750 heartbeats using explicit methods such as in Fig. 2b and compared the three hemodynamic metrics with the results obtained using LHMF. However, LHMF is not scalable in that capturing longer temporal domains would linearly scale in the number of HUs to simulate. Given a sufficiently large sequence of heartbeats, even the LHMF approach becomes intractable to simulate. As a refinement of LHMF, we established LHMF_C where we clustered hemodynamically similar HUs to avoid performing redundant simulations (Fig. 2d). Representative HUs could be simulated as a one-time, fixed cost to reconstruct a long sequence of heartbeats to enable the generation of longitudinal hemodynamics on demand (Fig. 2e). By capturing the full range of behavior over a substantial period of time, we enable the first 3D patient-specific longitudinal hemodynamic maps (LHMs) using LHMF. Unlike conventional spatial distributions of hemodynamics across patient anatomy, LHMs enable visualization of how hemodynamic metrics develop over time and indicate whether certain anatomic locations are predisposed to disease if adverse hemodynamic biomarkers are present frequently. A larger data set of 4.5 million heartbeats was used to generate the first LHMs (Supplementary Table 2).

Few cardiac cycles are required for temporal convergence in varying activity states and at the transition between activity states

To deconstruct temporal dependence, we needed to perform temporal convergence studies over a representative set of activity states and in between activity states. Performing such studies was imperative in ensuring accurate solutions for estimating the number of pre-flow cycles needed for all ranges of heartbeats and for all heartbeat transitions. Simulations were assumed to temporally converge after 5 cardiac cycles and used as ground truth⁵. We computed the percentage differences matched per time-step between the ground truth 5th cardiac cycle and the preceding cardiac cycles. Figure 3a demonstrates that the exponential decay curve of percent error vs.

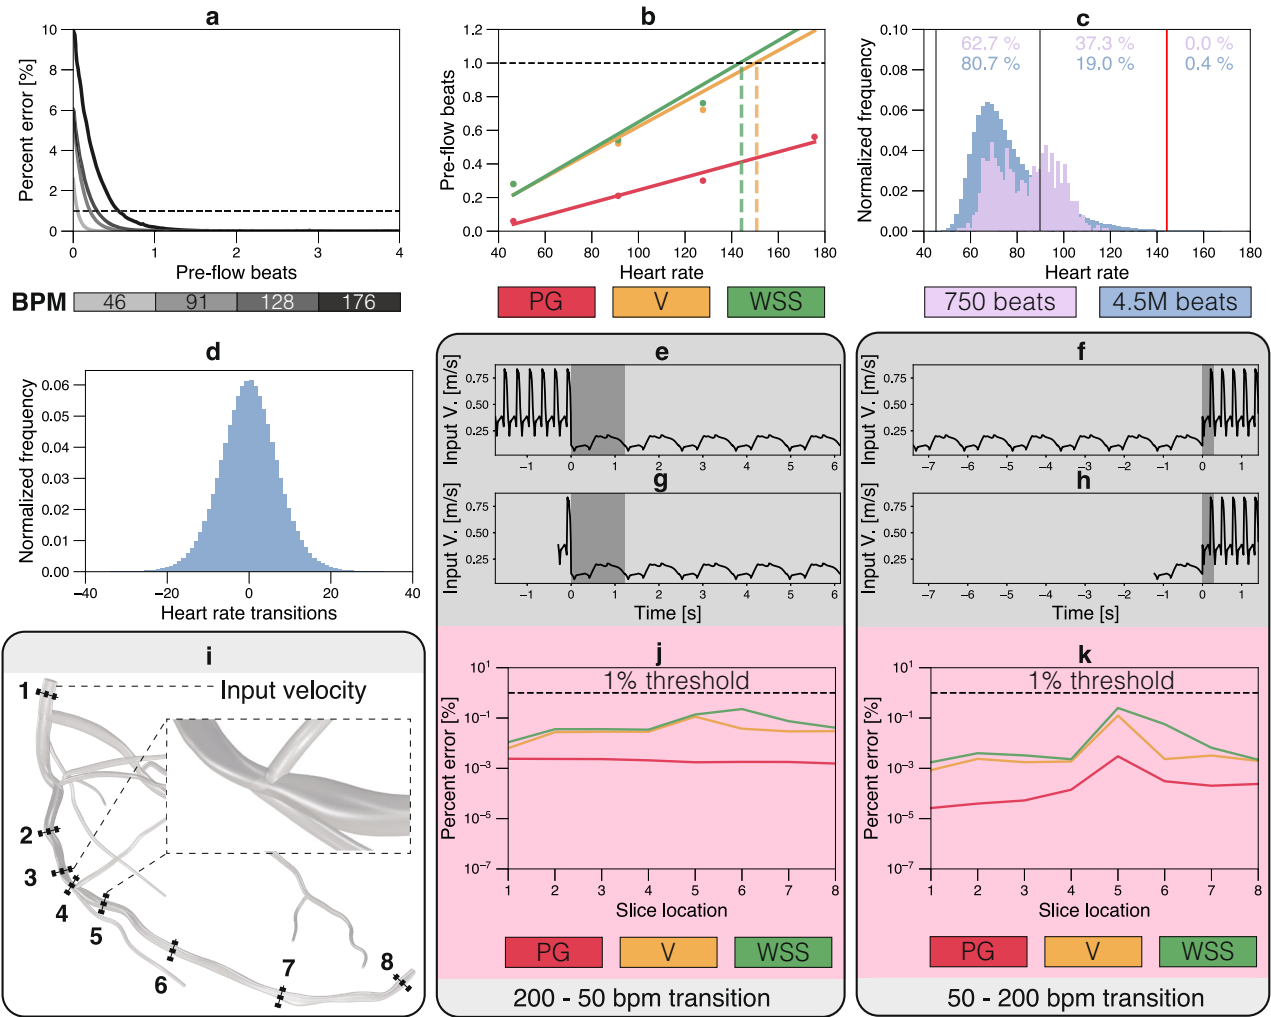


Fig. 3 | Evaluating temporal convergence for varying physical states. Temporal convergence was analyzed for resting pressure gradient (PG), velocity (V), and wall shear stress (WSS). **a, b** Temporal convergence studies with repeated input waveforms per heartbeat. **a** Per-time-step percentage discrepancy curves for PG. The percentage discrepancy was computed between timestep-matched data for an increasing number of pre-flow cardiac cycles and an assumed ground truth of 5 heartbeats. **b** Linear regressions yield R^2 values that exceed 0.97 for the three hemodynamic variables. Vertical interrupted lines delineate the heart rate at which more than one pre-flow cycle is needed. Heart rates to the left of the threshold converge in less than one pre-flow cycle, but heart rates to the right require two pre-flow cycles. **c** Histograms displaying the distribution of heart rates for the 750-beat

and 4.5 million-beat cases, divided into three bins showing prevalence. **d** Histogram showing heartbeat transitions in the 4.5 million heartbeat data. Although most of the transitions are ± 20 bpm, the largest changes were $+62$ and -78 bpm. **e-k** Temporal convergence studies for HUs consisting of varying physical states, specifically at the interface of high-to-low (**e, g, j**) and low-to-high heart rates (**f, h, k**). Input velocity waveforms are demarcated with a gray background (**e-h**), and errors in output hemodynamic metrics are shown in the pink background (**j, k**). Examples showing ground-truths in (**e, f**) and one pre-flow cases in (**g, h**). Gray regions in (**e-h**) are the heartbeats used for the analysis. **i** Locations of slices where temporal convergence was evaluated. **j, k** Two pre-flow cardiac cycles for both transitions were required for convergence at all locations. (**a-c**) is reproduced from ref. 85.

number of pre-flow beats was shifted to the right with increasing heart rate. Although only error curves are presented for the pressure gradient, the velocity and WSS also demonstrated the same characteristics. The number of pre-flow beats needed to cross the 1% threshold for temporal convergence was correlated for all hemodynamic variables (Fig. 3b). Linear regressions indicated near-perfect correlations with $R^2 > 0.97$. Across all heart rates, resting pressure gradient only required one pre-flow cycle whereas velocity and WSS required two pre-flow cycles after ~ 140 bpm. Histograms of the 750-beat and 4.5 million-beat long heart rate data indicated that all three hemodynamic metrics would converge under one pre-flow cycle for the 750-beat data and 0.4% of all beats in the 4.5 million-beat data (Fig. 3c).

The current standard in the literature is to resolve flow over single activity states, where input waveforms only consider rest or medically induced exercise states. Conventionally, temporal convergence is evaluated by simulating a sufficient number of preceding replicated heartbeats before the heartbeat of interest. This approach is insufficient for generating LHMs as HUs could capture heartbeats at the transition point between two activity

states. There may be instances of rapidly changing heart rates, such as when a patient abruptly starts to run from resting conditions, resulting in an extreme heart rate transition occurring within a HU. In the 4.5 million heartbeat data, the most extreme heart rate shifts were $+62$ and -78 bpm within an HU. The majority of the transitions were ± 20 bpm, as indicated by the histogram in Fig. 3d. To ensure temporal convergence for all HUs, we considered the extreme scenario with transitions of $50 \rightarrow 200$ and $200 \rightarrow 50$ bpm (Supplementary Table 3). For both high-to-low (Fig. 3e, g, j) and low-to-high (Fig. 3f, h, k) transitions, we evaluated temporal convergence by comparing heartbeats of interest after a transition between ground-truth waveforms and waveforms with an increasing number of preceding cycles. To ensure convergence across the left anterior descending artery, we collected results in eight different anatomic positions (Fig. 3i). We found that two pre-flow cycles were needed to reach temporally converged solutions for pressure gradient, velocity, and WSS with an error threshold of 1% (Fig. 3j, k) for both types of transitions. This result is similar to our previous temporal convergence study with increasing heart rates, where

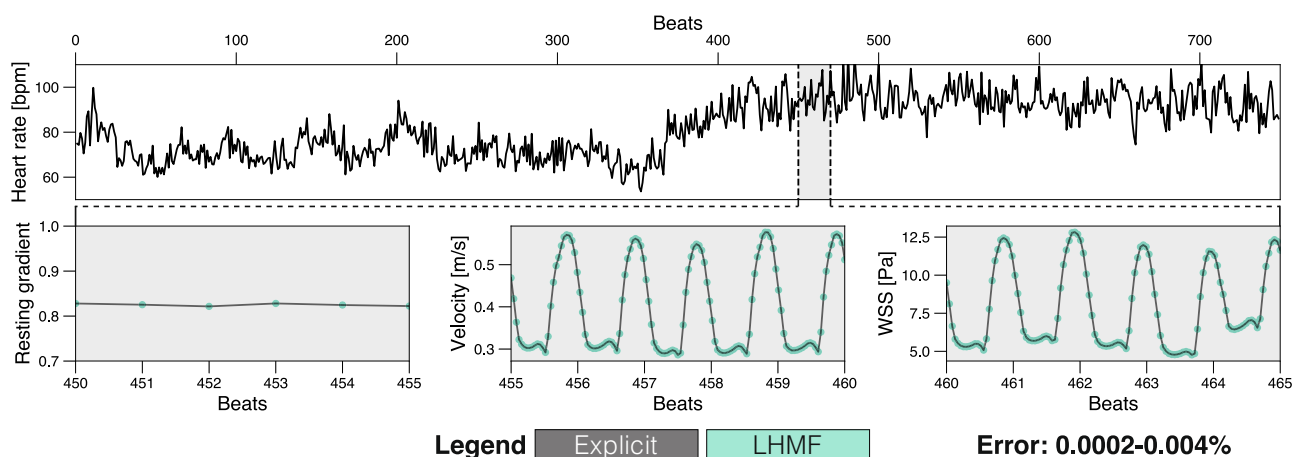


Fig. 4 | Longitudinal hemodynamic mapping framework (LHMF) recapitulates ground-truth simulations with over 99.995% accuracy. (Top) Continuous wearable heart rate data over 750 beats. (Bottom) Insets between beats 450–465 compare resting pressure gradient, velocity, and wall shear stress (WSS) using the explicit approach vs. LHMF. At outlets, the percent errors were 0.0002%, 0.003%, and

0.004% for resting pressure gradient, velocity, and WSS, respectively. Results from LHMF are depicted as points, while results from explicit methods are as lines to show overlap. Depicting both sets of results as points or lines would perfectly coincide such that the two sets of data would be indiscernible. Figure is reproduced from ref. 85.

HUs exceeding 140 bpm required two pre-flow cycles. Including a transition in heart rate within HUs does not appear to change the temporal convergence characteristics.

LHMF agrees with ground-truth 3D CFD solutions evaluated with 750 heartbeats

To verify that temporal dependence was deconstructed accurately, we ran a ground-truth explicit 3D CFD simulation spanning 750 heartbeats for comparison with results from the LHMF. Most studies tend to evaluate hemodynamic models based on a single output hemodynamic metric at one anatomical position^{3,4,17,18,44}. To ensure robust verification throughout the anatomical domain, we averaged data over the left anterior descending coronary artery with a lesion and incorporated three key hemodynamic metrics to demonstrate that LHMF accurately reconstructs longitudinal hemodynamics. LHMF was verified with ground-truth 750-beat data. Applying LHMF, we decomposed the 750 beats into 750 HUs and ran these simulations not only on the cloud, but also on traditional systems to show that the framework generalizes to heterogeneous systems. The objective was to show that LHMF accurately reconstructs a brute-force explicit longitudinal simulation and to further support the results of temporal convergence.

The resting pressure gradient, the velocity, and the WSS had percent errors of 0.0002%, 0.003%, and 0.004%, respectively. These errors were computed at the outlets, but verification across the entire left anterior descending artery also showed negligible error. Verification results are presented in Fig. 4.

LHMF_C drastically reduces the number of HUs needed for accurate re-assembly of longitudinal hemodynamics

LHMF was shown to recapitulate the results of an explicit 3D CFD simulation. As a refinement to LHMF, we further improved the tractability of generating LHMs by avoiding redundant simulations. To this end, we established LHMF_C, where we only performed simulations on representative HUs. Our goal here was to identify clusters of input conditions that are similar enough so that the resulting 3D flow simulations are not statistically different, thereby reducing the need for redundant simulations (see “Clustering algorithm to minimize running redundant simulations” of “Methods”). However, finding hemodynamically similar HUs is a nontrivial task, as hemodynamic solutions are not known a priori, yet matching the hemodynamic solution is the objective given varying velocity waveforms. Clustering nearly identical velocity waveforms may exaggerate the number

of HUs as relatively variable inlets may result in the same hemodynamic output^{17,18,45,46}. Bland-Altman bias (i.e., mean differences) and imprecision (i.e., standard deviation) are common metrics for quantifying agreement, which have been used extensively to verify computational hemodynamic models with in vivo measurements. For example, fractional flow reserve is a hemodynamic variable whose calculation by digital twins has been validated through comparison with invasive measurements. Standard software in the literature has reported a bias of 0.01 and an imprecision of 0.05 (provided normalization between 0 and 1)^{4,26,47–58}.

After establishing precision metrics, we used the 750-beat ground-truth data as a means to systematically explore clustering with the objective of keeping error to the Bland-Altman limits while minimizing the number of HUs. We first clustered all three hemodynamic outputs, averaged per heartbeat, to a stricter tolerance of 1%. Figure 5a illustrates an example of clustering WSS from 7 heartbeats into 3 clusters. We visually compared the inlet velocity waveforms of different HUs that belonged to the same group when clustering the hemodynamic outputs to correlate similarities in hemodynamics with similarities in inlet velocity waveforms (Fig. 5b). We identified three features that consistently identified HUs belonging to the same cluster that also revealed differences in HUs between clusters. We found that the phase-matched L^2 error, spanning the pre-flow cycles and heartbeat of interest, achieved effective clustering. L^2 is a metric used to compute the Euclidean distance. The waveforms were phase-corrected via dynamic time warping (DTW)—a common algorithm to align temporally mismatched data⁵⁹. Additionally, the velocity time integral (VTI), or area under the velocity waveform, for the diastolic (VTI_{dia}) and systolic phases (VTI_{sys}) each also identified hemodynamically similar HUs.

To determine the thresholds needed for clustering, we performed sensitivity studies between the three features used for clustering: VTI_{dia} , VTI_{sys} , and DTW. 15,625 combinations of features were considered by varying VTI_{dia} and VTI_{sys} between 1–50% discrepancy and DTW between 0.01–2.00, all with 25 regularly spaced increments. We used the resulting clusters to verify against the 750-beat ground truth by computing Bland-Altman bias and imprecision. We determined that DTW was the most sensitive feature and accounted for ~80% of the variance explained (Supplementary Fig. 1). Consequently, we simplified the threshold selection process by keeping VTI_{dia} and VTI_{sys} thresholds constant while varying DTW (Fig. 5c). The objective was to select the largest thresholds that

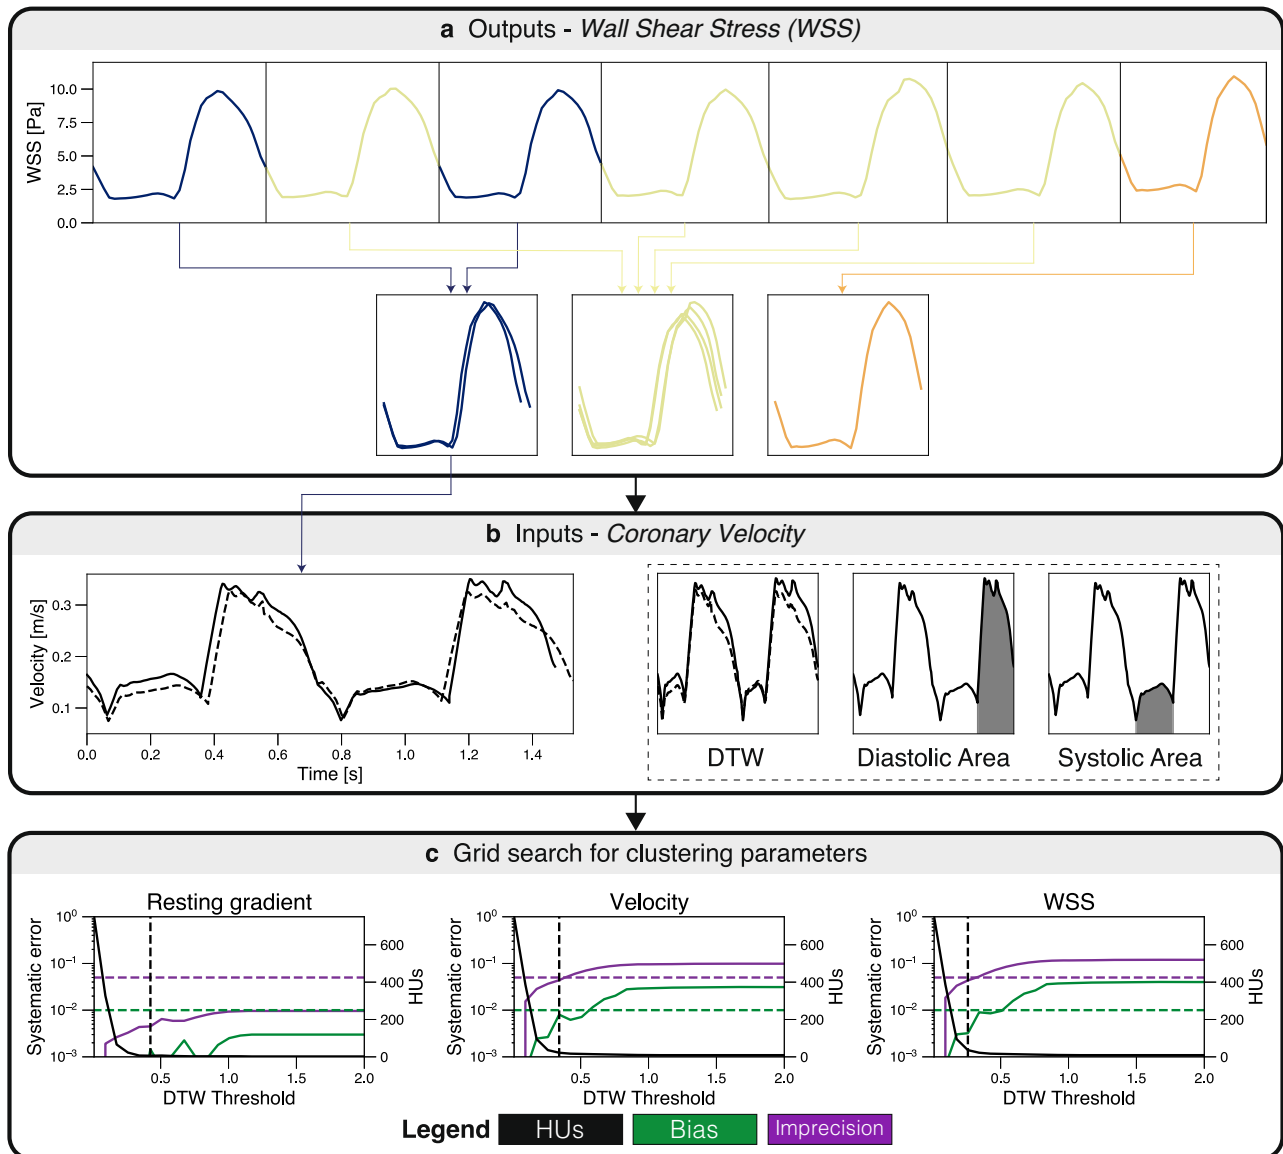


Fig. 5 | Clustering hemodynamically similar hemodynamic units (HUs) minimizes redundant simulations. **a** Clustering the outputs to inform hemodynamically important features. In this example, we find that the wall shear stress (WSS) outputs could be clustered into 3 groups with less than a 1% discrepancy. **b** Clustering outputs to identify features in the inlet velocity waveform that could be used for LHMFC. Solid and interrupted lines represent two example input waveforms in which the outputs were clustered to belong to the same group. This approach identified three salient features: dynamic time warped Euclidean distance (DTW), velocity time integral (VTI) or area of

diastole, and VTI of systole. **c** Clustering resting pressure gradient, and WSS over a fixed diastolic and systolic VTI. Interrupted horizontal lines delineate tolerances for Bland-Altman bias (0.01) and imprecision (0.05). Black vertical interrupted lines indicate DTW thresholds selected for each hemodynamic metric. For the same Bland-Altman tolerances, DTW thresholds decreased from resting pressure gradient to velocity and WSS. VTI of diastole and systole thresholds were kept constant for each subplot. **c** is reproduced from ref. 85.

adhere to the 0.01 bias and 0.05 imprecision limits while also minimizing the number of HUs. The resulting cluster thresholds and the number of HUs out of 750 beats are summarized in Supplementary Table 4. For the resting pressure gradient, even less stringent DTW thresholds could have been selected, but we decided to select 0.42 as increasing DTW within our parameter bounds no longer reduced HUs. The number of HUs was reduced by over an order of magnitude for all three hemodynamic metrics. While the DTW thresholds were relatively stringent, with L^2 errors ranging from 0.26-0.42, we found that the diastolic and systolic VTI could vary as much as 50% for the resting pressure gradient and still accurately obtain hemodynamic outputs. As one patient data set was used in this study, the cluster thresholds determined here were used for all clustering

results in subsequent sections (Supplementary Table 4). We compared output hemodynamics from LHMFC to the explicit ground truth in Supplementary Fig. 2.

Evaluating the performance of LHMFC for increasing problem sizes and across heterogeneous systems

We evaluated the performance of LHMFC and LHMFC in terms of the resulting number of HUs and time to solution for increasing problem sizes across heterogeneous computing resources. We used these two metrics as a way to compare both frameworks against the explicit ground truth. To test the resulting number of HUs for longer problem sizes, we performed clustering on increasing proportions of the 4.5 million-beat data, ranging from 1 day to 6 weeks' worth of heartbeats (Fig. 6a). The number of HUs

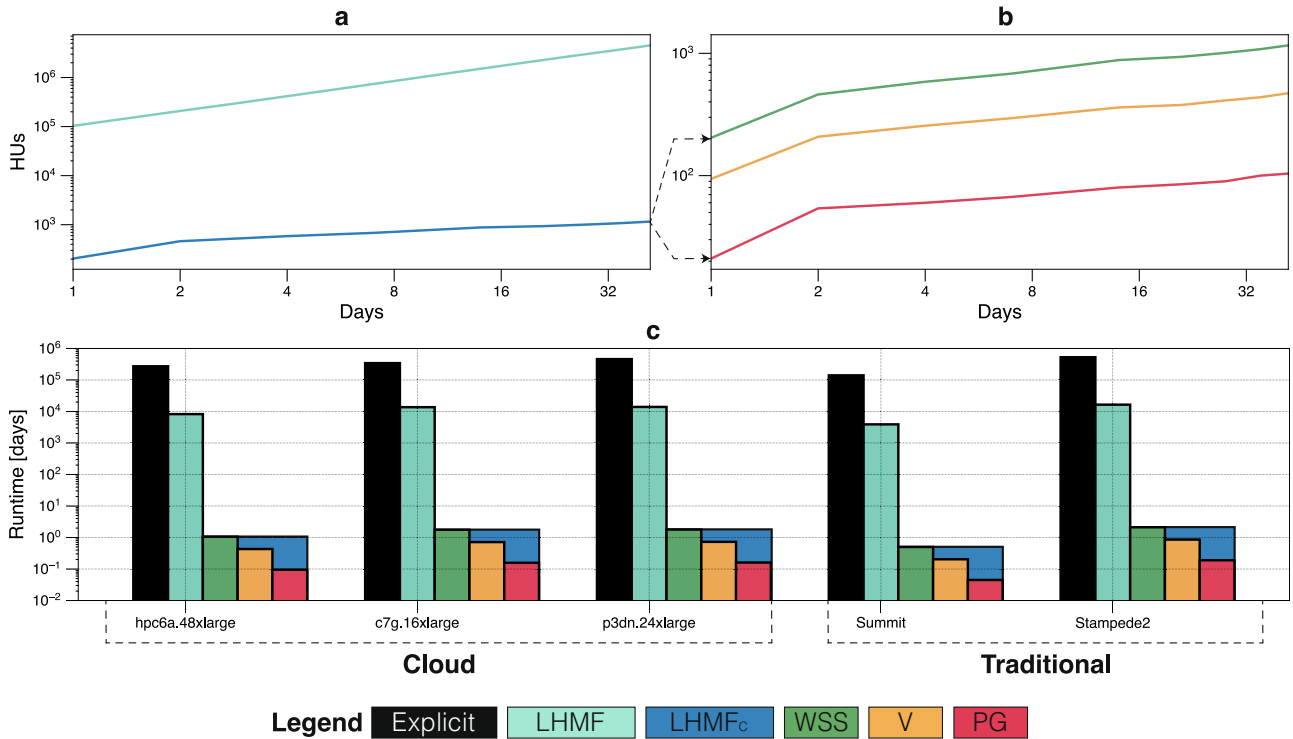


Fig. 6 | Longitudinal hemodynamic mapping framework (LHMF) and longitudinal hemodynamic mapping framework with clustering (LHMF_C) applied to different problem sizes and deployed on multiple computing resources.

a Number of HUs for LHMF vs. LHMF_C. For LHMF, the number of hemodynamic units (HUs) is equal to the number of HUs in the data set. Clustering achieves around four orders of magnitude decrease in the number of HUs required for the initial fixed set of simulations. **b** is an inset of the LHMF_C curve in **a**. Number of HUs for wall shear stress (WSS), velocity (V), and resting pressure gradient (PG). Across all problem sizes, WSS required the most HUs, followed by velocity and then resting pressure gradient. As the length of the

problem size increases, the rate of increase in the number of HUs diminishes. **c** Performance in terms of wallclock runtime across multiple types of cloud instances and traditional clusters. Each grouped bars represent a different system, including hpc6a.48xlarge (x86 CPUs), c7g.16xlarge (ARM CPUs), and p3dn.24xlarge (GPUs) instances. On the traditional cluster side, we present results on OLCF Summit and TACC Stampede2. The results are projections for a 4.5 million heartbeat sequence, comparing explicit vs LHMF vs LHMF_C for WSS, V, and PG. The results of LHMF_C are subdivided by hemodynamic metric. Figure is adapted from ref. 85.

decreased by around 4 orders of magnitude through clustering (Fig. 6b). From initially requiring 4.5 million HUs (and simulating over 9 million cardiac cycles) using LHMF to only needing 1160 HUs using LHMF_C for a problem size of 6 weeks. Clustering improved the tractability of performing longitudinal simulations. LHMF_C results were consistent with the temporal convergence results that demonstrated that WSS required the most pre-flow cardiac cycles at all heart rates, followed by velocity and resting pressure gradient.

Compared to traditional platforms, cloud computing resources are ideal for performing high-throughput digital twin simulations. Cloud resources do not have upfront capital costs or maintenance infrastructure, which allows flexible configuration with simulations provisioned on a per-usage basis^{60,61}. Cloud facilities are easily accessible to clinicians and can rapidly perform high-resolution 3D simulations in HIPAA-compliant applications with low latency⁶². Since access to compute resources likely varies between clinics, we require digital twins to be performant on both traditional and cloud-based systems.

The benefits of LHMF and LHMF_C were compared to the explicit approach across cloud facilities and traditional clusters. Using a fixed problem size of 6 weeks, we projected the wallclock time (Fig. 6c). The wallclock times for the explicit approach exceeded 10⁵ days. Across all systems, LHMF resulted in a two-orders-of-magnitude reduction in wallclock time. LHMF_C decreased wallclock times to ~43 hours across systems. Consistent with previous results, WSS required the most wallclock time, followed by velocity, and resting pressure gradient.

While OLCF Summit would have the fastest time-to-solution, cloud facilities were comparable in runtime. x86 CPUs in hpc6a.48xlarge

instances provided nearly identical wallclock time to ARM CPUs in c7g.16xlarge instances for the same CPU count. GPUs in p3dn.24xlarge instances had comparable wallclock time vs. cloud-based CPU instances.

Application of LHMF_C deployed across heterogeneous systems to generate LHMs spanning over 4.5 million heartbeats

We applied LHMF_C to a 4.5 million heartbeat problem to illustrate the need for LHMs over established single-heartbeat metrics. Hemodynamic maps are conventionally visualized on data collected within the clinic. With LHMs, we can monitor hemodynamically significant metrics following daily activity states longitudinally. LHMF_C enables an accurate re-assembly of longitudinal hemodynamics by just simulating representative HUs. To this end, we clustered 4.5 million heartbeats, representing 1.5 months of continuous activity. 1160 HUs were required to capture this sequence.

LHMF_C enabled tractable simulation of HUs to capture longitudinal behavior. Our focus was to generate the first LHM for WSS since low WSS (<1 Pa) has been linked to the development of atherosclerotic plaque⁶³. Persisting low WSS regions over time may indicate sites predisposed to further disease. We generated spatial maps of WSS for single heartbeats at rest and exercise state. For comparison, we generated a spatial map averaged over 1.5 months of activity—establishing the first spatial LHM of WSS (Fig. 7a). To quantify the differences between the spatial LHM and the established single-heartbeat maps, we calculated the percentage difference per-fluid point (Fig. 7b). The per-point error distributions demonstrated that both single-heartbeat maps differed from the LHM. The rest and exercise state maps had percent discrepancies between 10–1000% compared to spatial

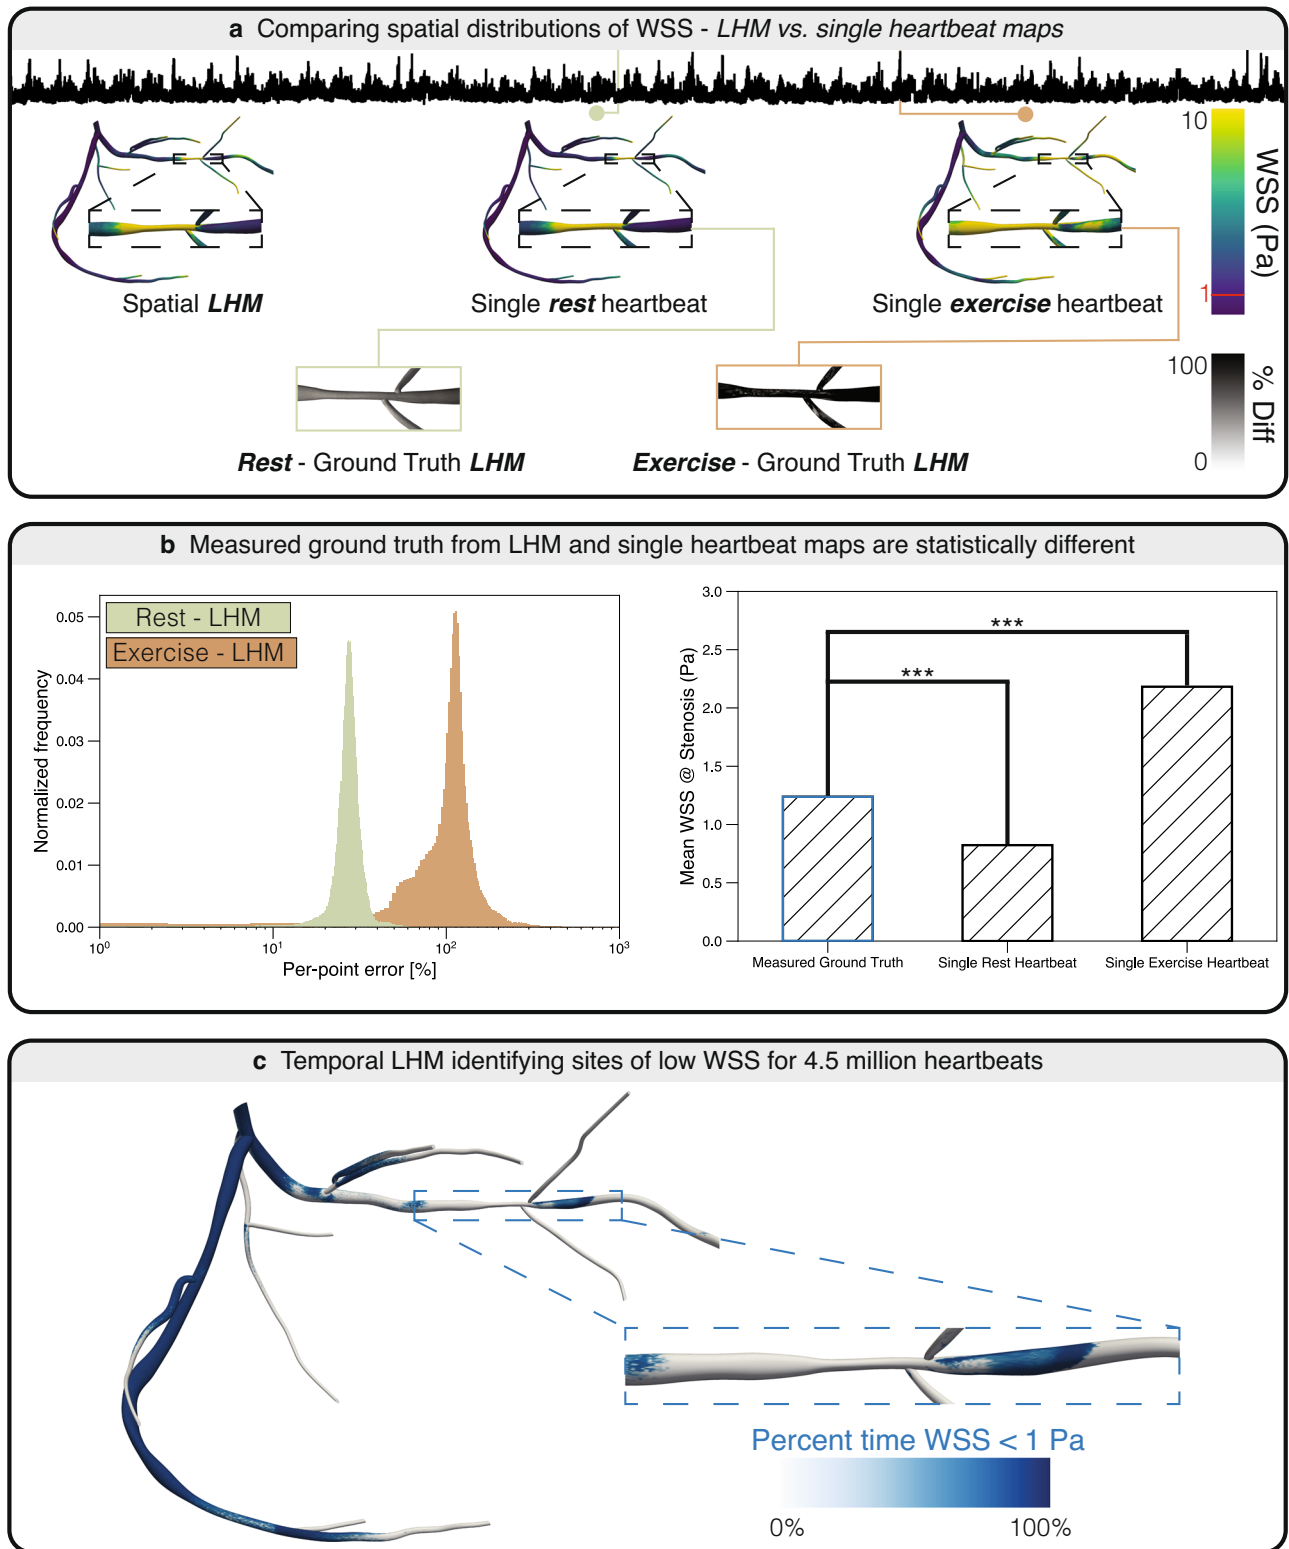


Fig. 7 | Application of longitudinal hemodynamic mapping framework with clustering (LHMF_C) for capturing 1.5 months of daily activity (or 4.5 million heartbeats). Using clustering, only 1160 hemodynamic units (HUs) were required for accurate re-assembly of hemodynamics. **a** Comparison of wall shear stress (WSS) spatial distributions between an average WSS spatial map for 4.5 million heartbeats, WSS map of the rest state single-heartbeat, and WSS map of the exercise state single-heartbeat. Heart rate signals are shown on top for the 4.5 million heartbeats, and moments that were used for the single rest and exercise heartbeats are demarcated with solid vertical lines. Differences between established single-heartbeat maps and the spatial longitudinal hemodynamic map (LHM) are shown in the bottom row. We

display percent error maps comparing rest vs. spatial LHM and exercise vs. spatial LHM. **b** Quantifying the differences between single-heartbeat maps and LHM. (Left) Taking the difference between the established single-heartbeat maps and the spatial LHM, representing the measured ground truth, for every fluid point, we show that the exercise state differed more than the rest state, but both differed from the LHM. (Right) At the location of the stenosis, we show that the mean WSS is statistically different compared to single-heartbeat maps ($p < 0.001$). **c** Temporal LHM of WSS. The fraction of time spent with WSS < 1 Pa over 1.5 months of activity was visualized in the temporal LHM, highlighting data differences that would not be apparent with single-heartbeat maps.

LHM. The exercise state had, on average, almost an order of magnitude higher error than the rest state. When isolating fluid points in the stenotic region, where flow is more likely to be disturbed and result in disease progression, the mean WSS calculated from the LHM was statistically different from the results of a single rest and exercise heartbeat ($p < 0.001$). These quantitative comparisons show that the established single-heartbeat maps cannot capture longitudinal hemodynamic behavior.

Beyond spatial hemodynamic maps, we created a temporal map for the first time that displays the percentage of time spent with low WSS, providing a novel way to identify local regions with potential risks for disease development. The temporal LHM of WSS (Fig. 7c) revealed that immediately after the stenosis, low WSS occurred frequently during the six weeks, which could indicate potential progression of the disease.

Discussion

The central goal of this work was to enable 3D patient-specific longitudinal hemodynamic simulations driven by wearables. This study presented the first LHM of WSS that spans six weeks of activity. Compared to single-heartbeat WSS maps in varying activity states, we demonstrated in a patient-specific case that LHMs provide additional hemodynamic information that could not be captured by established methods. The LHM was made feasible by integrating continuous wearable data as input to our CFD framework to account for realistic changes in physical states. We proposed LHMF and LHMF_C that take advantage of the rapid temporal convergence characteristics to decompose long sequences of heartbeats into representative HUs. The validity of this approach was compared with ground-truth data generated for 750 beats simulated contiguously for three hemodynamic metrics—the largest and most comprehensive verification of longitudinal 3D dynamics in a patient-specific geometry to our knowledge. We further demonstrated that the set of HUs needed to reconstruct longitudinal phenomena accurately could be drastically reduced via clustering hemodynamically equivalent heartbeats, thereby minimizing redundant simulations to improve the tractability of our framework further.

Establishing LHMF as it pertains to longitudinal blood flow simulations in the coronary is but one application. LHMF is generalizable to other anatomies and even time series applications beyond hemodynamics. The integral part of the feasibility of LHMF is to accurately deconstruct the temporal dependence. Time-series problems adhere to the causality principle, where solutions later in time are influenced, or even determined, by solutions occurring earlier³⁰. LHMF would be translatable to other applications if the window of influence of preceding solutions is determined. In our case, we performed temporal convergence studies to determine how many heartbeats would be influenced by earlier heartbeats. We sampled the entire 4.5 million heartbeat data set by performing convergence studies at four representative increments. Also, we simulated potential edge cases where a HU had a significant transition in heart rate. While other works¹⁴ have proposed heuristics to estimate the number of cardiac cycles required for temporal convergence, these other works have focused mainly on problems where heartbeats are replicated in time. Our problem is unique, where no two inlet coronary velocity waveforms would be perfectly identical, thereby necessitating robust empirical temporal convergence studies that consider varying physical states. To this end, we discovered a pattern where heart rates below 140 bpm only required simulating one pre-flow cardiac cycle, and those above 140 bpm required two pre-flow cardiac cycles. The same pattern was also found in the study that simulates the edge cases of large heart rate fluctuations at the interface of different physical states. Although we focused on one patient in this study, LHMF is generalizable, and expanding to other patients would be straightforward. When targeting a new patient, we would only need to run a series of temporal convergence studies to sample different activity levels to understand how to deconstruct temporal dependence. Future work could look at developing a standardized calibration protocol to enable seamless use of LHMF for any given patient. Given that the patient we selected is the average representative case from a data set of 160 others, we anticipate similar flow regimes in the other coronary geometries and, therefore, comparable convergence characteristics.

Temporal convergence studies have been performed for pressure gradient for a larger subset of patients of the 160 and they have all converged similarly⁷. For eventual clinical translation, this proof-of-concept study indicated that we require a systematic and standardized protocol to evaluate patient temporal convergence characteristics before effectively applying LHMF or LHMF_C.

The objective of LHMF_C was to minimize the number of HUs by ensuring hemodynamically equivalent outputs while allowing as much variability as possible in input coronary waveforms. Although no two coronary velocity waveforms would be perfectly identical, which we model by introducing some stochasticity into our wearable-driven inlet boundary condition, the resulting hemodynamics could still be equivalent. Taking advantage of this property, we established LHMF_C—a refinement of LHMF in which we clustered hemodynamically similar HUs. LHMF_C reduced the number of HUs required to capture longitudinal hemodynamics by around four orders of magnitude compared to LHMF. Cluster analysis demonstrated that WSS required the most HUs, followed by velocity and resting pressure gradient (Fig. 6b). Similarly, WSS required the most pre-flow cardiac cycles for temporal convergence for the same heart rate, followed by velocity and resting pressure gradient. This result could be attributed to error propagation. The pressure gradient is derived from the density of LBM, calculated as the macroscopic moment of 0th order of the LBM distribution function. Velocity is derived from the 1st order moment of the distribution function, and WSS is a function of both fluid density and velocity. While our work was verified by capturing the entire flow field of hemodynamic metrics, not all applications require capturing the entire set of data. For example, if one is interested in only capturing the pressure gradient, the number of HUs needed would decrease by another order of magnitude (five orders of magnitude when compared to LHMF). From requiring over 380 years in wallclock time using the entirety of OLCF Summit (the world's fifth fastest supercomputer⁶³) to simulate six weeks of activity using the explicit approach, we enable tractable simulation of longitudinal hemodynamics. Using 400 instances of `hpc6a.48xlarge` cloud resources, which is orders of magnitude less performant than the OLCF Summit, LHMF requires just over 22 years of compute time while LHMF_C only requires 2.3 hours to capture resting pressure gradient. Therefore, not only did we enable feasible computation of longitudinal blood flow, but depending on the hemodynamic variable of interest, selective clustering for accuracy in just one hemodynamic metric could further reduce compute time.

In recent years, cloud computing facilities have become increasingly performant and accessible. Furthermore, the computing hardware is rapidly improving. Beyond using conventional x86 CPUs, hardware acceleration via GPUs and ARM CPUs is becoming popular. Accelerating code by off-loading compute-intensive parts to the GPU is now a standard model. ARM CPUs are shifting the mobile computing paradigm and are even used in Fugaku, the world's second-fastest supercomputer, in 2023. For the eventual clinical translation of LHMF and LHMF_C, it was essential to future-proof our framework so that we could fully utilize the continually improving hardware. To test the feasibility of LHMF on these heterogeneous hardware systems, we deployed the framework on cloud facilities and traditional clusters for comparison. For a heartbeat sequence length of 6 weeks, executing high-throughput simulations of HUs on state-of-the-art instances for x86 CPUs, ARM CPUs, and GPUs as well as on Summit and Stampede2 confirmed the improvements in wallclock time, ranging from centuries when using established CFD methods to just under 43 hours via LHMF_C when capturing the complete set of hemodynamic outputs. Furthermore, cloud resources produced noninferior wallclock time to traditional clusters, highlighting that on top of providing key functionality established by traditional clusters, cloud platforms would become integral to our pipeline for future clinical translation as it avoids the need for on-premise systems and high upfront costs while also offering flexible, on-demand deployment features. Ultimately, this study enables longitudinal 3D hemodynamic monitoring that matches established brute-force methods for key hemodynamic metrics.

To generate the first spatial and temporal LHMs, we deployed LHMFC on heterogeneous systems to reconstruct 4.5 million heartbeats (or 1.5 months of activity). The need for LHMs was demonstrated by comparing them with established single-heartbeat maps. When comparing the full coronary tree and only the stenotic region, the spatial WSS LHM was statistically different from the single-heartbeat maps at rest and exercise state (Fig. 7b). Therefore, only capturing single-heartbeat maps would not be representative of longitudinal hemodynamics. Moreover, we generated the first temporal LHM of WSS, where we superimposed the incidence of time spent with low WSS on the 3D coronary tree model. Prolonged exposure to low WSS can cause endothelial cell dysfunction and thickening of the blood vessel wall, which has been linked to the development of plaque and hypertension⁶⁴. Additionally, this extended exposure to low WSS increases the risk of developing cardiovascular disease. Therefore, it is essential to identify when the WSS falls below a certain threshold and to measure how much time it spends at that level to better understand the potential risk to cardiovascular health. This case example demonstrates that LHMs capture additional hemodynamic information and nuances that can only be captured over time and that single-heartbeat analyses cannot provide.

Although this work takes foundational steps to establish the feasibility of LHMs, further advancements are still needed for clinical translation. Determining boundary conditions is a critical part of recapitulating physiologically accurate flow conditions in a digital twin. While we inform the inlet boundary condition from multimodal wearable data, further investigation is required to ensure that the translation of wearable data to personalized coronary velocity waveforms is clinically accurate. As we shift to analyzing 3D hemodynamics over weeks to months, there are further hurdles that need to be overcome. In this temporal regime, we would need to account for phenomena such as growth and remodeling. Patient-specific boundary conditions are likely to change over time. For example, changes due to vascular remodeling could alter patient anatomy, especially at the site of stenosis, which could modify the flow distribution down the coronary tree. Additionally, the effect of treatment should also be taken into account. Common drugs used to treat coronary disease can change blood viscosity, stroke volume, as well as microvascular resistance⁶⁵ over time. Incorporating boundary conditions that adapt to microvascular characteristics and the physiological state of the patient can provide even greater personalized precision⁶⁶. Sources of error from the wearable device itself need to be explicitly taken into account. LHMFC would require data quality measures that could account for edge cases that are not representative of patient physiology. In future work, it would be necessary to explore boundary conditions that are more physiologically accurate and adaptable to individual patients as well as implement on-the-fly data quality mechanisms. A tight feedback loop would be required to update the accurate anatomy and changes to hemodynamic conditions over time. In this work, we captured every single heartbeat to monitor the amount of time a patient spends at a disease-prone regime. Understanding the temporal resolution needed to capture longitudinal dynamics is an important question to be investigated in future work. We focused on capturing full-resolution temporal fluid dynamics in pursuit of uncovering novel longitudinal hemodynamic biomarkers. For eventual clinical translation, studying the minimum temporal resolution needed would be an essential next step. Our current implementation serves as a means to capture varying physical states over time to establish and verify LHMFC. The long-term goal of LHMFC is to apply it to phenomena over months. While more development is required for clinical translation and to study complications such as ISR, this work moves the needle for being able to tractably compute 3D hemodynamics beyond one heartbeat and sets the stage for the study of longitudinal 3D flow dynamics.

Uncertainty quantification, especially variance-based approaches^{17,18,67}, could not only be used to improve our clustering methodology but also

provide error bounds into our time-series predictions. While LHMFC will not have error bounds because we use the maximum number of HUs, uncertainty quantification should be considered when applying LHMFC. In future work, we can use global uncertainty quantification to determine the relative importance of different features of the input coronary velocity waveform on the resulting 3D hemodynamics to further reduce the number of HUs. Simultaneously, we could also incorporate uncertainties into our velocity waveforms to account for potential sources of error, such as noisy data from wearables. Accounting for potential sources of error from wearables, such as noisy data or faulty data, are important aspects that require careful consideration before making LHMs clinically translatable. Issues such as motion artefacts, signal crossover, and gaps in recording (such as due to charging) have been explored in^{68–70} and will need to be explicitly accounted for before clinical translation. Incorporating uncertainty quantification into our predictions and implementing data-cleaning measures will be explored in future work.

In conclusion, this study has established both LHMFC and LHMFC as a means to leverage cloud-based resources to produce LHMs on the order of months for the first time. We have shown that using a finite set of representative HUs is sufficient to accurately reconstruct the solution of a 3D CFD simulation that captures an entire long sequence of heartbeats. This study represents the first critical step to allow longitudinal hemodynamics calculation, enabling additional studies to accurately reflect longitudinal biological changes. With further development, LHMs are expected to advance our understanding and treatment of cardiovascular disease, creating a step-change improvement in the number of heartbeats that can be accurately simulated.

Methods

Patient data

This study did not involve direct patient experimentation, interaction, or human tissue samples. The study was approved by the Massachusetts General Brigham Institutional Review Board (IRB Protocol #2015P001084). The IRB did not require signed informed consent from individual patients because no patient interaction or intervention was performed, and the study was not prospective. Although we acquired coronary angiograms and clinical measurements for 160 patients, we selected a representative patient for analysis in terms of geometry size. The patient had angiographically documented coronary artery disease as evaluated by an expert interventional cardiologist at Brigham and Women's Hospital, Boston, MA, USA. Coronary angiograms comprised at least four standard orthogonal views of the left coronary circulation. Physiological measurements were taken during coronary angiography to inform personalized digital twins, including aortic blood pressure, cardiac output, heart rate, and hematocrit.

Personalized coronary digital twins

Coronary geometry reconstruction. The complete coronary tree models were reconstructed from pairs of coronary angiograms using an algorithm described in^{71,72}. 2D vessel centerlines were first computed with the corresponding cross-sectional diameters in the images; a fully automated computation generated 3D coronary tree models from 2D centerlines. As a feedback loop, an expert interventional cardiologist verified the reconstructions topologically and anatomically using ImageJ v1.52k (NIH, Bethesda, MD, USA). All identifiable (>1 mm) main and side branch vessels were reconstructed in the 3D model. This reconstruction algorithm has been extensively applied to the coronary arteries^{3,10,17,18}.

CFD solver. Hemodynamic simulations were performed using HARVEY, our in-house massively parallel fluid dynamics application^{11,16}. HARVEY is based on the lattice Boltzmann method (LBM), a mesoscopic approach to numerically solving the Navier-Stokes equations. The fluid is represented as a set of discrete particles with time evolution described with a particle distribution function. The lattice Boltzmann equation (Equation (1)) describes the time

evolution of the particle distribution function $f_i(x, t)$:

$$f_i(x + c_i \Delta t, t + \Delta t) = f_i(x, t) + \Omega_i(x, t) \tag{1}$$

The distribution function gives the probability of particles located at lattice point x and time t with discrete velocity c_i . HARVEY employs the standard D3Q19 velocity discretization and single relaxation time Bhatnagar-Gross-Krook collision kernel. Finite difference boundary conditions are implemented at the inlet and outlets based on ref. 73 and no-slip conditions are enforced at rigid walls using the halfway bounce-back boundary condition. A more detailed discussion of LBM can be found in refs. 74,75. HARVEY is parallelized using MPI and has been shown to scale linearly on up to 1.5 million CPUs¹¹ and 3 thousand GPUs⁷⁶. There are multiple versions of HARVEY, and we focused on using CPU-only and CPU-GPU implementations. For the CPU-GPU approach, compute-intensive kernels are performed on the GPU.

Assumptions of the digital twin. The coronary circulation was modeled as an incompressible Newtonian fluid with a density of 1060 kg/m³ and a patient-specific dynamic viscosity of 2.32 cP. Dynamic viscosity was computed from clinically measured hematocrit based on an empirical relationship linking the two metrics⁷⁷:

$$\mu = \frac{\mu_0}{1 - \phi} \tag{2}$$

where μ is dynamic viscosity, μ_0 is dynamic viscosity of plasma, and ϕ is hematocrit. We assumed a plasma dynamic viscosity of 1.2 cP^{3,10,77}. The vessels in the coronary tree were modeled as rigid walls with no-slip boundary conditions at the walls. This is a common assumption for coronary vessels because of low shear rates of blood that are unlikely to cause significant lumen deformation^{3,10,18,34,78}.

Analysis of hemodynamics. This work focused on extracting WSS, PG, and velocity to verify LHMf and generate the first LHMs. WSS was computed at fluid points directly adjacent to walls. Based on⁷⁹, the WSS vector (τ_i) was computed as follows:

$$\tau_i = -\frac{\mu\omega}{c_s^2\rho} f_\alpha^{neq} c_{\alpha j} n_j (c_{\alpha j} - c_{\alpha k} n_i n_k) \tag{3}$$

where f_α^{neq} is the non-equilibrium particle distribution function, μ is dynamic viscosity of blood, ω is relaxation rate, ρ is fluid density, $c_{\alpha i}$ are components of the discrete velocity vector c_i , and n is the outward normal vector. WSS was computed as the magnitude of the WSS vector. Further details on computing WSS inside HARVEY are described in ref. 3.

PG is a pressure-based metric. To convert lattice units of pressure to physical pressure, we employed the following:

$$\left. \frac{\rho_{diastole} - \rho_0}{\rho_{systole} - \rho_0} \right|_{lbm} = \left. \frac{P_{diastole,manometric}}{P_{systole,manometric}} \right|_{physical} \tag{4}$$

$$p[Pa] = c_s^2(\rho - \rho_0) \frac{\Delta x^2}{\Delta t^2} \rho_{blood} \tag{5}$$

where $\rho_{diastole}$, $\rho_{systole}$ are the lattice units of pressure at diastole and systole, respectively, $P_{diastole}$, $P_{systole}$ are clinically measured manometric pressures, c_s is the lattice speed of sound, Δx is the grid spacing, and Δt is the size of the timestep. ρ_{nought} is a constant used to calibrate LBM to physical units of pressure and has been applied and validated in ref. 3.

PG was computed as beat-averaged, cross-sectionally averaged pressure downstream of a stenosis (P_d)—a local narrowing of vessel lumen

because of disease, to that at the coronary inlet (P_a):

$$PG = P_d/P_a \tag{6}$$

Outputs per timestep consist of point cloud data where each fluid point ($n = 15.7$ million) in the domain contains pressure, velocity, and WSS data. Further, we stored cross-sectional slices of fluid points normal to the direction of flow for analysis. Slice data was primarily used for the temporal convergence studies and verifying LHMf, and the point cloud data for generating LHMs and comparing them with single-heartbeat maps.

Wearable-driven boundary conditions for hemodynamic simulations

Transient velocity waveforms were used at the inlet boundary condition and coupled to multimodal continuous wearable data (Fig. 1).

Baseline clinical characteristics were collected for one patient under resting conditions. Stroke volume (Q_{SV}) was assumed to be constant longitudinally and approximated via clinically measured cardiac output and a baseline resting heart rate ($HR_{baseline}$). Coronary fraction (α) is the proportion of cardiac output entering the coronary circulation. We used empirical coronary fraction patient averages from refs. 34,80. Baseline coronary flow rate ($Q_{baseline}$) was calculated as follows:

$$Q_{baseline} = \alpha Q_{SV} HR_{baseline} \tag{7}$$

We used a representative left coronary waveform from ref. 81 as a template and scaled the velocity amplitude (v_{max}) and period ($\tau_{cardiac}$) on demand. $Q_{baseline}$ was first converted into velocity using the coronary ostial area (A_{ostial}) and then modified with instantaneous heart rate data ($HR_{wearable}$):

$$v_{max} = \frac{2Q_{baseline}}{A_{ostial}} \frac{HR_{wearable}}{HR_{baseline}} + \sigma_{amplitude} \tag{8}$$

Further, we incorporated an amplitude noise term ($\sigma_{amplitude}$) to account for natural variability in coronary velocity waveforms. $\sigma_{amplitude}$ was modeled as a normal distribution with a coefficient of variance obtained by superimposing transthoracic Doppler coronary tracings from ref. 36. The cardiac period was scaled according to the following:

$$\tau_{cardiac} = \frac{60}{HR_{wearable}} + \sigma_{period} \tag{9}$$

where σ_{period} is a period noise term also informed by transthoracic Doppler coronary tracings.

As a proof of concept of leveraging ECG data to inform timings of systole and diastole, we used physician-labeled ECG tracings from ref. 82 and linearly correlated the RT segment duration to heart rate ($f_{ECG}(HR_{wearable})$). The RT segment varies linearly with heart rate ($R^2 = 0.99$), and this relationship was used to introduce phase shifts into the coronary velocity waveform by modifying diastolic ($\tau_{diastolic}$) and systolic timings ($\tau_{systolic}$). Specifically, we modified the diastolic and systolic timings as follows:

$$\tau_{diastole} = f_{ECG}(HR_{wearable}) + \sigma_{phase} \tag{10}$$

$$\tau_{systolic} = \tau_{cardiac} - \tau_{diastolic} \tag{11}$$

where σ_{phase} was computed as phase differences observed in transthoracic Doppler coronary tracings. In summary, multimodal continuous data, heart rate, and ECG data were used to inform the generation of coronary velocity waveforms.

We employed constant 2-element Windkessel models for the outlet boundary condition, consisting of peripheral resistance (R_p) and compliance (C). The Windkessel models were applied to the ends of each

terminal branch to account for microvascular hemodynamics⁸³:

$$Q = \frac{P}{R_p} + C \frac{dP}{dt} \quad (12)$$

C was assumed to be constant at $1 \times 10^{-12} \text{ m}^3 \text{ mm Hg}^{-1}$ based on ref. 84. R_p was distributed among the terminal vessels using relationships commonly used for patient-specific coronary simulations^{3,10,34,35}:

$$R_i = P_{\text{mean}} \frac{\sum_{j=1}^{N_{\text{terminal}}} r_j^3}{Q_{\text{ostial}} r_i^3} \quad (13)$$

where R_i is the peripheral resistance in each terminal vessel, P_{mean} is the mean arterial pressure, Q_{ostial} is ostial flow rate, r is the radius of the average terminal branch, and N_{terminal} is the number of terminal branches. We used a constant ostial flow rate in the resting state to parameterize the outflow conditions.

Clustering algorithm to minimize running redundant simulations

Our algorithm to determine hemodynamically similar clusters based on inlet velocity waveforms is summarized in Supplementary Algorithm 1. When clustering N_{beats} , we first determined a cluster centroid based on the first index of a set for all unclustered heartbeats. For this centroid, we computed the number of pre-flow cycles needed for temporal convergence (N_{pre}), and the VTI_{dia} and VTI_{sys} intended for the heartbeat intended for hemodynamic analysis. Next, we iterated every successive non-clustered heartbeat. We compared the $DTWL^2$ errors, and VTI metrics to empirically determined pre-set tolerances, where t_{dia} , t_{sys} , and t_{dtw} correspond to thresholds in VTI_{dia} , VTI_{sys} , and DTW , respectively. The algorithm was repeated until all beats had been clustered. Once all clusters had been determined, the representative heartbeats (or cluster centroids) were all deployed to heterogeneous systems for analysis. If new time-series data were provided beyond the initial set used for clustering, each new unclustered heartbeat would be sorted into pre-existing clusters on demand. Further, if hemodynamic results for the initial fixed set of heartbeats have been analyzed, the new set of time-series data would be re-assembled without running any more simulations. LHMF_C enables deconstructing long time-series data into representative HUs and then simulating just an initial fixed cost of the representative HUs to re-assemble longitudinal hemodynamics. The edge case of when a new unclustered heartbeat does not fit into any existing clusters was also considered, which would result in the deployment of new simulations on demand.

Deploying simulations on cloud and traditional systems

We wanted to ensure the generalizability of LHMF across multiple types of computing resources. In this study, we focused on cloud (`hpc6a.48xlarge`, `c7g.16xlarge`, and `p3dn.24xlarge` instances) and traditional supercomputers (OLCF Summit and TACC Stampede2). To verify LHMF, we ran the ground-truth simulation of 750 beats on the Duke Compute Cluster using 1440 tasks (or 36 nodes) of Intel Xeon E5-2695V4 ‘Broadwell’ processors. Applying LHMF, we decomposed the 750 beats into 750 HUs and launched the HUs in parallel in 6 instances of `hpc6a.48xlarge` instances, 64 nodes of the OLCF Summit, and 32 nodes of TACC Stampede2 per HU. To test LHMF_C, clustered HUs were run on multiple cloud instance types (Supplementary Table 5). `hpc6a.48xlarge` instances with AMD EPYC 7003 CPUs were used to run production simulations to generate LHMs; specifically, 400 instances with 38,400 CPUs were used to perform high-throughput simulations of clustered HUs. `c7g.16xlarge` instances containing ARM Graviton3 CPUs and `p3dn.24xlarge` instances containing Tesla V100 GPUs were used for comparison. In addition, traditional clusters such as OLCF Summit and TACC Stampede2 were used for comparisons with cloud facilities. When applying LHMF_C to generate LHMs spanning 4.5 million heartbeats,

the representative 1160 HUs were deployed in the cloud with 400 `hpc6a.48xlarge` instances, each job using six nodes of AMD EPYC 7003 CPUs. We chose this particular configuration for high-throughput simulations because scaling was nearly linear up to 6 nodes or 576 CPU cores (Supplementary Fig. 3). Furthermore, we also deployed simulations on the OLCF Summit and the Duke Compute Cluster to demonstrate that LHMF_C is generalizable across heterogeneous systems on both CPU-only and hybrid CPU-GPU architectures.

Statistical analysis

The Kolmogorov–Smirnov test was used to first confirm that the data followed the central limit theorem. The correlation between the number of pre-flow beats and heart rate was determined using ordinary least squares linear regression. Temporal convergence was evaluated using beat-averaged and slice-averaged percent differences for all three hemodynamic metrics. Similarly, LHMF was verified against the explicit ground-truth simulation of 750 cardiac cycles using slice-averaged percent differences per timestep and beat-averaged. Absolute agreement between LHMF_C and explicit results was evaluated using Bland–Altman analysis. Bias was computed as the mean difference, and imprecision as the standard deviation between the results of the two approaches. WSS in the stenotic region followed a log-normal distribution. To compare LHM with single-heartbeat maps in the stenotic region, we performed a logarithmic transform and confirmed normality using Kolmogorov–Smirnov tests. We performed a one-way ANOVA to compare the three maps and afterward performed a post hoc analysis via Tukey’s Honest Significant Difference test. A value of $p < 0.05$ was considered significant. All statistical tests were performed on JMP Pro 16 (JMP Statistical Discovery LLC, Cary, NC, USA).

Data availability

Clinical and imaging data cannot be shared due to the privacy of the individuals whose data was used in this work. Continuous data were used from <https://physionet.org/content/bigideaslab-step-hr-smartwatch/1.0/>, <https://zenodo.org/record/53894#.ZD4kXi-B2gQ>, and <https://archive.physionet.org/physiobank/database/qtddbto> to inform the generation of coronary velocity waveforms. The output 3D hemodynamic data are available at <https://doi.org/10.7924/r4f76jd8n>.

Code availability

The code is available for academic license through the Duke University Office of Licensing and Ventures.

Received: 13 November 2023; Accepted: 5 August 2024;

Published online: 06 September 2024

References

- Coveney, P. V., Hoekstra, A., Rodriguez, B. & Viceconti, M. Computational biomedicine. part ii: organs and systems, *Interface Focus* (2021).
- Hernandez-Boussard, T. et al. Digital twins for predictive oncology will be a paradigm shift for precision cancer care. *Nat. Med.* **27**, 2065–2066 (2021).
- Vardhan, M. et al. Non-invasive characterization of complex coronary lesions. *Sci. Rep.* **11**, 1–15 (2021).
- Nørgaard, B. L. et al. Diagnostic performance of noninvasive fractional flow reserve derived from coronary computed tomography angiography in suspected coronary artery disease: the NXT trial (analysis of coronary blood flow using CT angiography: next steps). *J. Am. Coll. Cardiol.* **63**, 1145–1155 (2014).
- Vardhan, M. et al. Diagnostic performance of coronary angiography derived computational fractional flow reserve. *J. Am. Heart Assoc.* **13**, e029941 (2024).
- Grande Gutiérrez, N. et al. Computational modeling of blood component transport related to coronary artery thrombosis in kawasaki disease. *PLoS Comput. Biol.* **17**, e1009331 (2021).

7. Alkhouli, M. et al. Trends in characteristics and outcomes of hospital inpatients undergoing coronary revascularization in the united states, 2003-2016. *JAMA Netw. Open* **3**, e1921326 (2020).
8. Dangas, G. D. et al. In-stent restenosis in the drug-eluting stent era. *J. Am. Coll. Cardiol.* **56**, 1897–1907 (2010).
9. Moussa, I. D. et al. Trends and outcomes of restenosis after coronary stent implantation in the United States. *J. Am. Coll. Cardiol.* **76**, 1521–1531 (2020).
10. Vardhan, M. et al. The importance of side branches in modeling 3d hemodynamics from angiograms for patients with coronary artery disease. *Sci. Rep.* **9**, 1–10 (2019).
11. Randles, A., Draeger, E. W., Ooppelstrup, T., Krauss, L. & Gunnels, J. A. Massively parallel models of the human circulatory system. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–11 (2015).
12. Randles, A. & Kaxiras, E. Parallel in time approximation of the lattice Boltzmann method for laminar flows. *J. Comput. Phys.* **270**, 577–586 (2014).
13. Randles, A. & Kaxiras, E. A spatio-temporal coupling method to reduce the time-to-solution of cardiovascular simulations. In: *2014 IEEE 28th International Parallel and Distributed Processing Symposium*, 593–602 (IEEE, 2014).
14. Pfaller, M. R., Pham, J., Wilson, N. M., Parker, D. W. & Marsden, A. L. On the periodicity of cardiovascular fluid dynamics simulations. *Ann. Biomed. Eng.* **49**, 3574–3592 (2021).
15. Bäumler, K. et al. Fluid–structure interaction simulations of patient-specific aortic dissection. *Biomech. Model. Mechanobiol.* **19**, 1607–1628 (2020).
16. Randles, A. P., Kale, V., Hammond, J., Gropp, W. & Kaxiras, E. Performance analysis of the lattice Boltzmann model beyond navier-stokes. In: *2013 IEEE 27th International Symposium on Parallel and Distributed Processing*, 1063–1074 (IEEE, 2013).
17. Tanade, C. et al. Global sensitivity analysis for clinically validated 1d models of fractional flow reserve. In: *43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 4395–4398 (IEEE, 2021).
18. Tanade, C., Chen, S. J., Leopold, J. A. & Randles, A. Analysis identifying minimal governing parameters for clinically accurate in silico fractional flow reserve. *Front. Med. Technol.* **4**, 1034801 (2022).
19. Taylor, C. A. et al. Patient-specific modeling of blood flow in the coronary arteries. *Comput. Methods Appl. Mech. Eng.* **417**, 116414 (2023).
20. Sagawa, K., Lie, R. K. & Schaefer, J. Translation of otto frank’s paper “die grundform des arteriellen pulses” zeitschrift für biologie 37: 483-526 (1899). *J. Mol. Cell. Cardiol.* **22**, 253–254 (1990).
21. Vardhan, M. et al. Evaluation of intracoronary hemodynamics identifies perturbations in vorticity. *Front. Syst. Biol.* **2**, 930396 (2022).
22. Tomizawa, N. et al. Coronary flow disturbance assessed by vorticity as a cause of functionally significant stenosis. *Eur. Radiol.* **32**, 6859–6867 (2022).
23. Chu, M. et al. Quantification of disturbed coronary flow by disturbed vorticity index and relation with fractional flow reserve. *Atherosclerosis* **273**, 136–144 (2018).
24. Hughes, T. J. & Lubliner, J. On the one-dimensional theory of blood flow in the larger vessels. *Math. Biosci.* **18**, 161–170 (1973).
25. Katritsis, D. G. et al. Flow patterns at stented coronary bifurcations: computational fluid dynamics analysis. *Circ. Cardiovasc. Interv.* **5**, 530–539 (2012).
26. Pellicano, M. et al. Validation study of image-based fractional flow reserve during coronary angiography. *Circ. Cardiovasc. Interv.* **10**, e005259 (2017).
27. Lantz, J., Gårdhagen, R. & Karlsson, M. Quantifying turbulent wall shear stress in a subject specific human aorta using large eddy simulation. *Med. Eng. Phys.* **34**, 1139–1148 (2012).
28. Christopher, J., Gao, X., Guzik, S. M., Falgout, R. & Schroder, J. Parallel in time for a fully space-time adaptive mesh refinement algorithm. In: *AIAA Scitech 2020 Forum*, 0340 (2020).
29. Gander, M. J. 50 years of time parallel time integration. In: *Multiple Shooting and Time Domain Decomposition Methods: MuS-TDD, Heidelberg*, 69–113 (Springer, 2015).
30. Tompson, J., Schlachter, K., Sprechmann, P. & Perlin, K. Accelerating Eulerian fluid simulation with convolutional networks. In: Precup, D. & Teh, Y. W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 3424–3433 (PMLR, 2017).
31. Um, K., Hu, X. & Thuerey, N. Liquid splash modeling with neural networks. In: *Computer Graphics Forum*, vol. 37, 171–182 (Wiley Online Library, 2018).
32. Sanchez-Gonzalez, A. et al. Learning to simulate complex physics with graph networks. In: Ill, H. D. & Singh, A. (eds.) *Proceedings of the 37th International Conference on Machine Learning*, vol. 119, 8459–8468 (PMLR, 2020).
33. Taylor, C. A., Fonte, T. A. & Min, J. K. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: scientific basis. *J. Am. Coll. Cardiol.* **61**, 2233–2241 (2013).
34. Sharma, P. et al. A framework for personalization of coronary flow computations during rest and hyperemia. In: *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 6665–6668 (IEEE, 2012).
35. Sunyecz, I. L., McCallinhart, P. E., Patel, K. U., McDermott, M. R. & Trask, A. J. Defining coronary flow patterns: comprehensive automation of transthoracic doppler coronary blood flow. *Sci. Rep.* **8**, 1–12 (2018).
36. Weber, J. L. & Myers, E. W. Human whole-genome shotgun sequencing. *Genome Res.* **7**, 401–409 (1997).
37. Chen, K. & Pachter, L. Bioinformatics for whole-genome shotgun sequencing of microbial communities. *PLoS Comput. Biol.* **1**, e24 (2005).
38. Brenchley, R. et al. Analysis of the bread wheat genome using whole-genome shotgun sequencing. *Nature* **491**, 705–710 (2012).
39. Kim, H. J., Vignon-Clementel, I., Figueroa, C., Jansen, K. & Taylor, C. Developing computational methods for three-dimensional finite element simulations of coronary blood flow. *Finite Elem. Anal. Des.* **46**, 514–525 (2010).
40. Fleeter, C. M., Geraci, G., Schiavazzi, D. E., Kahn, A. M. & Marsden, A. L. Multilevel and multifidelity uncertainty quantification for cardiovascular hemodynamics. *Comput. Methods Appl. Mech. Eng.* **365**, 113030 (2020).
41. Furberg, R., Brinton, J., Keating, M. & Ortiz, A. Crowd-sourced Fitbit datasets. <https://doi.org/10.5281/zenodo.53894> (2016).
42. Hennigan, B. et al. Discordance between resting and hyperemic indices of coronary stenosis severity: the verify 2 study (a comparative study of resting coronary pressure gradient, instantaneous wave-free ratio and fractional flow reserve in an unselected population referred for invasive angiography). *Circ. Cardiovasc. Interv.* **9**, e004016 (2016).
43. Blanco, P. J. et al. Comparison of 1d and 3d models for the estimation of fractional flow reserve. *Sci. Rep.* **8**, 17275 (2018).
44. Sankaran, S., Kim, H. J., Choi, G. & Taylor, C. A. Uncertainty quantification in coronary blood flow simulations: impact of geometry, boundary conditions and blood viscosity. *J. Biomech.* **49**, 2540–2547 (2016).
45. 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).
46. Morris, P. D. et al. Virtual fractional flow reserve from coronary angiography: modeling the significance of coronary lesions: results from the virtu-1 (virtual fractional flow reserve from coronary angiography) study. *Cardiovasc. Interv.* **6**, 149–157 (2013).

47. Tu, S. et al. Fractional flow reserve calculation from 3-dimensional quantitative coronary angiography and timi frame count: a fast computer model to quantify the functional significance of moderately obstructed coronary arteries. *JACC: Cardiovasc. Interv.* **7**, 768–777 (2014).
48. Tröbs, M. et al. Comparison of fractional flow reserve based on computational fluid dynamics modeling using coronary angiographic vessel morphology versus invasively measured fractional flow reserve. *Am. J. Cardiol.* **117**, 29–35 (2016).
49. Tu, S. et al. Diagnostic accuracy of fast computational approaches to derive fractional flow reserve from diagnostic coronary angiography: the international multicenter favor pilot study. *Cardiovasc. Interv.* **9**, 2024–2035 (2016).
50. Xu, B. et al. Diagnostic accuracy of angiography-based quantitative flow ratio measurements for online assessment of coronary stenosis. *J. Am. Coll. Cardiol.* **70**, 3077–3087 (2017).
51. Westra, J. et al. Diagnostic performance of in-procedure angiography-derived quantitative flow reserve compared to pressure-derived fractional flow reserve: the favor ii europe-japan study. *J. Am. Heart Assoc.* **7**, e009603 (2018).
52. Yazaki, K. et al. Applicability of 3-dimensional quantitative coronary angiography-derived computed fractional flow reserve for intermediate coronary stenosis. *Circ. J.* **81**, 988–992 (2017).
53. Westra, J. et al. Evaluation of coronary artery stenosis by quantitative flow ratio during invasive coronary angiography: the wifii study (wire-free functional imaging ii). *Circ. Cardiovasc. Imaging* **11**, e007107 (2018).
54. Legutko, J. et al. P2378 correlation between quantitative flow ratio (qfr) and fractional flow reserve (ffr). *Eur. Heart J.* **38**, ehx502.P2378 (2017).
55. Van Rosendaal, A. et al. Accuracy and reproducibility of fast fractional flow reserve computation from invasive coronary angiography. *Int. J. Cardiovasc. Imaging* **33**, 1305–1312 (2017).
56. Masdjedi, K. et al. Vessel fractional flow reserve (vffr) for the assessment of stenosis severity: the fast ii study. *Eurointervention* **17**, 1498–1505 (2021).
57. Li, J. et al. Accuracy of computational pressure–fluid dynamics applied to coronary angiography to derive fractional flow reserve: FLASH FFR. *Cardiovasc. Res.* **116**, 1349–1356 (2020).
58. Giorgino, T. Computing and visualizing dynamic time warping alignments in r: the dtw package. *J. Stat. Softw.* **31**, 1–24 (2009).
59. Gupta, A. et al. The who, what, why, and how of high performance computing in the cloud. In: *2013 IEEE 5th International Conference On Cloud Computing Technology And Science*. **1**, 306–314 (IEEE, 2013).
60. Evangelinos, C. & Hill, C. Cloud computing for parallel scientific hpc applications: feasibility of running coupled atmosphere-ocean climate models on amazons EC2. *Paper presented at the CCA-08 in Chicago* **2**, 2–34 (2008).
61. Yimam, D. & Fernandez, E. B. A survey of compliance issues in cloud computing. *J. Internet Serv. Appl.* **7**, 1–12 (2016).
62. Samady, H. et al. Coronary artery wall shear stress is associated with the progression and transformation of atherosclerotic plaque and arterial remodeling in patients with coronary artery disease. *Circulation* **124**, 779–788 (2011).
63. Top500 list june 2023 <https://www.top500.org/lists/top500/list/2023/06/> (2023).
64. Cecchi, E. et al. Role of hemodynamic shear stress in cardiovascular disease. *Atherosclerosis* **214**, 249–256 (2011).
65. Bairey Merz, C. N., Pepine, C. J., Shimokawa, H. & Berry, C. Treatment of coronary microvascular dysfunction. *Cardiovasc. Res.* **116**, 856–870 (2020).
66. Vignon-Clementel, I. E., Figueroa, C., Jansen, K. & Taylor, C. Outflow boundary conditions for 3d simulations of non-periodic blood flow and pressure fields in deformable arteries. *Comput. Methods Biomech. Biomed. Eng.* **13**, 625–640 (2010).
67. Eck, V. G. et al. A guide to uncertainty quantification and sensitivity analysis for cardiovascular applications. *Int. J. Numer. Methods Biomed. Eng.* **32**, e02755 (2016).
68. Bent, B., Goldstein, B. A., Kibbe, W. A. & Dunn, J. P. Investigating sources of inaccuracy in wearable optical heart rate sensors. *NPJ Digit. Med.* **3**, 18 (2020).
69. Hicks, J. L. et al. Best practices for analyzing large-scale health data from wearables and smartphone apps. *NPJ Digit. Med.* **2**, 45 (2019).
70. Nelson, B. W. et al. Guidelines for wrist-worn consumer wearable assessment of heart rate in biobehavioral research. *NPJ Digit. Med.* **3**, 90 (2020).
71. Chen, S. J. & Carroll, J. D. 3-d reconstruction of coronary arterial tree to optimize angiographic visualization. *IEEE Trans. Med. Imaging* **19**, 318–336 (2000).
72. 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**, 451–459 (2005).
73. Latt, J., Chopard, B., Malaspinas, O., Deville, M. & Michler, A. Straight velocity boundaries in the lattice boltzmann method. *Phys. Rev. E* **77**, 056703 (2008).
74. Chen, S. & Doolen, G. D. Lattice boltzmann method for fluid flows. *Annu. Rev. Fluid Mech.* **30**, 329–364 (1998).
75. Krüger, T. et al. The lattice Boltzmann method: principles and practice. *Springer* **690**, 1868–4513 (2017).
76. Ames, J. et al. Multi-gpu immersed boundary method hemodynamics simulations. *J. Comput. Sci.* **44**, 101153 (2020).
77. Pirofsky, B. et al. The determination of blood viscosity in man by a method based on Poiseuille’s law. *J. Clin. Investig.* **32**, 292–298 (1953).
78. Eslami, P. et al. Effect of wall elasticity on hemodynamics and wall shear stress in patient-specific simulations in the coronary arteries. *J. Biomech. Eng.* **142**, 024503 (2020).
79. Matyka, M., Koza, Z. & Miroslaw, Ł. Wall orientation and shear stress in the lattice Boltzmann model. *Comput. Fluids* **73**, 115–123 (2013).
80. Wilson, R. F., Wyche, K., Christensen, B. V., Zimmer, S. & Laxson, D. D. Effects of adenosine on human coronary arterial circulation. *Circulation* **82**, 1595–1606 (1990).
81. Hadjiloizou, N. et al. Differences in cardiac microcirculatory wave patterns between the proximal left mainstem and proximal right coronary artery. *Am. J. Physiol. Heart Circ. Physiol.* **295**, H1198–H1205 (2008).
82. Laguna, P., Mark, R. G., Goldberg, A. & Moody, G. B. A database for evaluation of algorithms for measurement of qt and other waveform intervals in the ECG. In: *Comput. Cardiol.* 673–676 (IEEE, 1997).
83. Westerhof, N., Lankhaar, J.-W. & Westerhof, B. E. The arterial Windkessel. *Med. Biol. Eng. Comput.* **47**, 131–141 (2009).
84. 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).
85. Tanade, C., Rakestraw, E., Ladd, W., Draeger, E. & Randles, A. Cloud computing to enable wearable-driven longitudinal hemodynamic maps. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 1–14 (2023).

Acknowledgements

The authors thank Drs. Christopher W. Jensen, Justen Geddes, and Arash Ghorbannia for fruitful discussions. Computing support for this work came from Amazon Web Services (AWS) Cloud Credits for Research, Duke Compute Cluster, and Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS), which is supported by NSF grant number ACI-1548562. An award of compute time was provided by the

INCITE program. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This research used resources from the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. The AWS team supported this work through architecture reviews to optimize for the cloud and provide access to high-performance computing cores. Research reported in this work was supported by the National Institute On Aging of the NIH under Award Number DP1AG082343, NSF GRFP under Grant no. DGE 164486, and the Coulter Foundation. The content does not necessarily represent the official views of the NIH or NSF.

Author contributions

C.T.: conceptualization, methodology, software development, analysis, draft writing. N.K.: analysis and draft writing. E.R.: conceptualization. W.L.: cloud computing setup. E.D.: conceptualization, supervision, and draft writing. A.R.: funding acquisition, resources, supervision, conceptualization, and draft writing. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41746-024-01216-3>.

Correspondence and requests for materials should be addressed to Amanda Randles.

Reprints and permissions information is available at <http://www.nature.com/reprints>

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

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024