HarVI: Real-time intervention planning for coronary artery disease using machine learning

Cyrus Tanade[1 [0000−0002−2395−6908]] and Amanda Randles[1 [0000−0001−6318−3885]]

Dept of Biomedical Engineering, Duke University, Durham NC 27705, USA
{cyrus.tanade, amanda.randles}@duke.edu

Abstract. Virtual planning tools that provide intuitive user interaction and immediate hemodynamic feedback are crucial for cardiologists to effectively treat coronary artery disease. Current FDA-approved tools for coronary intervention planning require days of preliminary processing and rely on conventional 2D displays for hemodynamic evaluation. Immersion offered by extended reality (XR) has been found to benefit intervention planning over traditional 2D displays. To bridge these gaps, we introduce HarVI, a coronary intervention planner that leverages machine learning for real-time hemodynamic analysis and extended reality for intuitive 3D user interaction. The framework uses a predefined set of 1D computational fluid dynamics (CFD) simulations to perform one-shot training for our machine learning-based blood flow model. In a cohort of 50 patients, we calculated fractional flow reserve (FFR), the gold standard biomarker of ischemia in coronary disease, using HarVI (FFR_{HarVI}) and 1D CFD models (FFR_{1D}). HarVI was shown to almost perfectly recapitulate the results of 1D CFD simulations through continuous and categorical validation scores. In this study, we establish a machine learning-based process for virtual coronary treatment planning with an average turnaround time of just 74 minutes, thus reducing the required time for one-shot training to less than a working day.

Keywords: coronary artery disease · machine learning · virtual reality.

1 Introduction

Enabling intervention planning can give physicians an intuitive approach to performing virtual interventions and receiving real-time hemodynamic feedback to guide clinical decision making. By integrating patient-specific computational fluid dynamics (CFD) models, our planning tool offers a seamless and intuitive platform to perform virtual interventions in extended reality (XR). The role of immersion beyond traditional 2D displays has been explored and shown to improve user interaction for intervention planning [1, 2, 3, 4, 5, 6, 7]. More importantly, our tool also provides real-time hemodynamic feedback, which is crucial for informed clinical decision making. Advances in personalized CFD have already demonstrated the ability to non-invasively and accurately determine key hemodynamic metrics vital for decision-making processes. For example,

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

fractional flow reserve (FFR) [8] is a pressure-based metric that indicates coronary ischemia [9]. There is even FDA-approved coronary intervention planning software that predicts FFR in response to treatment [10, 11]. However, these state-of-the-art models face limitations in turnaround time and user interaction.

In coronary artery disease (CAD), stent implantation is the leading percutaneous coronary intervention (PCI) to treat functionally significant lesions (i.e., FFR $\leq$ 0.80), with more than 600,000 stents implanted annually in the United States alone [12]. However, 25 % of the patients evaluated for a successful PCI procedure still have residual ischemia associated with long-term adverse outcomes [13]. This issue of incomplete functional revascularization could be mitigated by allowing interventional cardiologists to interactively experiment with a variety of possible PCI strategies and their impact on the hemodynamics of that patient to determine how best to relieve ischemia. Another use case is in complex coronary lesions. Bifurcation lesions represent 20 % of cases and occur in arterial bifurcations or branch points [14]. Determining how best to treat these lesions remains a particular challenge, as they can affect the main branch (MB), the side branch (SB), or both vessels, and bifurcation stenting is associated with a higher risk of adverse cardiac events [15]. Using intervention planning, we could rapidly test strategies in MB, SB, or both lesions to help identify options to achieve complete revascularization and improve outcomes.

Machine learning has emerged as a promising option to predict FFR without explicitly running CFD simulations [10]. The only FDA-approved coronary intervention planning tool uses an interpolation model, which requires running a series of 3D CFD simulations as part of model training [10]. Although accurate, these models require 24-48 hours of processing time before a clinician can apply the planning tool. 3D [8] and 1D CFD models of FFR [16] have both been shown to recover invasive FFR accurately, and we hypothesized that real-time predictors of post-intervention FFR derived from reduced-order models would accurately recapitulate virtual PCI hemodynamics with short turnaround times. Enabling a faster turnaround time is helpful for some patient subgroups with unstable coronary disease, such as patients with ST-elevated myocardial infarction, where door-to-balloon time – the time from hospital admission to intervention – is recommended to be less than 90 minutes [17].

In addition to efficiently training machine learning models, another essential part of intervention planning is enabling intuitive interaction with 3D geometries to simulate intervention. The utility of intervention planning tools can benefit from the immersion offered by XR devices. The role of immersion in intervention planning and evaluating hemodynamic feedback has been extensively explored through user studies. Visualizing complex anatomy in traditional 2D displays could be challenging with vessel overlap and foreshortening effects [18]. These user studies [1, 2, 3, 4] have demonstrated that immersive displays are beneficial over 2D displays in the analysis of hemodynamic maps and performing some intervention planning tasks. There is well-established geometry editing software in the literature [5, 19], but all of these frameworks do not provide instantaneous hemodynamic updates for geometric modification. The only FDA-approved coro-

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

nary intervention planning tool that provides instantaneous hemodynamic feedback uses 2D displays. To our knowledge, no intervention planning tool exists with a turnaround time of one working day and is compatible with commonly available XR headsets.

To enable intervention planning with short turnaround times that leverage XR, we present HarVI (**HAR**VEY [20] **V**irtual **I**ntervention). HarVI is the backend that predicts post-PCI FFR. On the user interaction side, we used Harvis [2] – an established computational platform to modify geometries, deploy massively parallel simulations, and visualize hemodynamic results [2]. We introduce an innovative tool for intuitive clinical decision support by leveraging the real-time flow prediction we establish here with HarVI alongside the immersive anatomical

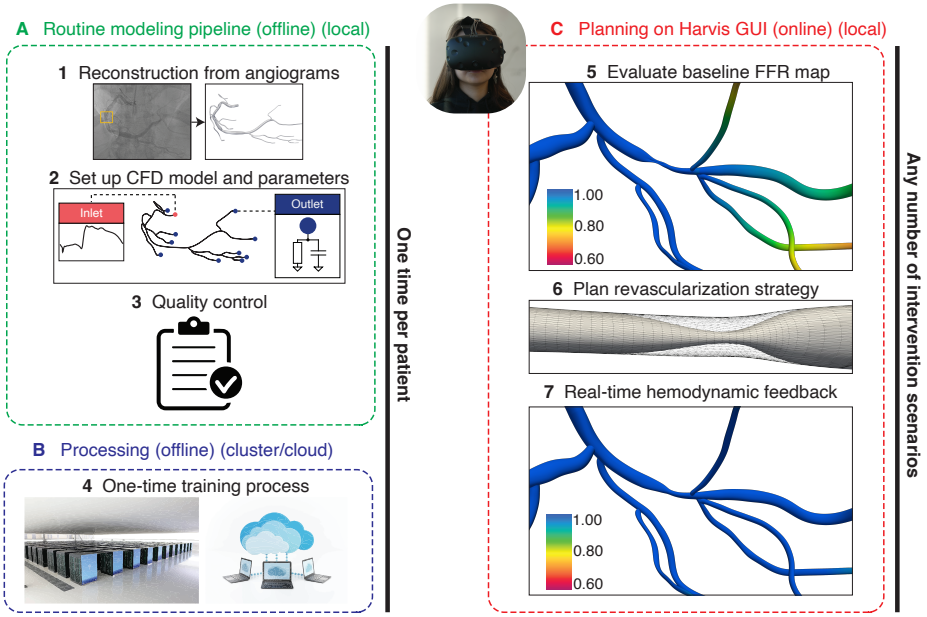

Fig. 1. Overview of the HarVI pipeline. (**A**) Clinically validated patient-specific modeling pipeline for coronary arteries. The computational domain was derived from coronary angiogram reconstructions. Boundary conditions were informed by clinical measurements. Quality control was performed to ensure reconstruction accuracy compared to coronary angiogram analysis. (**B**) A set of modifications was made to the reconstructed geometries to sample probable intervention scenarios. This step is a one-shot learning process and was invoked once per patient. The results of this training process were used to train a machine learning model to enable real-time prediction of post-intervention FFR. (**C**) After training the HarVI model, users could plan a revascularization strategy using extended reality headsets and receive real-time hemodynamic feedback in the Harvis GUI. HarVI refers to the machine learning model used for intervention planning. We coupled HarVI with Harvis – a virtual reality platform for patient-specific modeling.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

editing capabilities from Harvis (**Figure 1**). HarVI uses a clinically validated 1D FFR modeling pipeline [16] to train the machine learning model for real-time flow prediction (**Figure 1A-B**). This process is invoked only once per patient as part of a training routine. The results of the training process are used to create a machine learning model to predict post-PCI FFR in real-time. Finally, Harvis allows users to perform a virtual intervention in an immersive environment and visualize updated hemodynamic results, and HarVI provides instantaneous hemodynamic feedback (**Figure 1C**). This paper takes a foundational step in the establishment of immersive, intuitive, and real-time techniques for clinical decision support for CAD interventions.

2 Methods

2.1 Personalized hemodynamic analysis using reduced-order models

Image-derived data provided the complex 3D geometries needed for this study. Using anonymized and deidentified data from previous research [16], our analysis incorporated clinical measurements from 50 patients who had confirmed coronary artery disease by angiography. These data were originally collected during invasive coronary angiography procedures, to create patient-specific blood flow models. The key clinical measurements used for this purpose included cardiac output, cuff pressure, heart rate, and hematocrit.

We applied our in-house 1D blood flow simulator to create patient-specific 1D blood flow models of the coronary arteries [16, 21]. Our approach differs from FDA-approved software [16]. For flow simulations, image-derived 3D anatomies were required. In this study, we use anonymized 3D meshes from our previous work [16]. These 3D anatomies were reconstructed from coronary angiograms using the SnakeTree3D software described in [18]. The 1D computational domain was defined by computing centerlines and corresponding hydraulic diameters from the 3D reconstructed geometry using Mimics (Materialise, Leuven, BE).

Reduced-order blood flow simulations to compute pre-intervention FFR were performed using a well-established 1D blood flow simulator [16, 22]. Pre-intervention FFR calculated from the 1D simulator has been clinically validated against invasively measured FFR in all 50 patients in this study [16]. We maintain the same model assumptions as in [16] where blood was modeled as an incompressible Newtonian fluid with a density of 1060 kg/m^3 and dynamic viscosity evaluated per patient. Pulsatile flow rate waveforms were incorporated at the inlet, and 2-element Windkessel models (resistance and compliance) were employed at the outlets. All simulations were tuned to simulate hyperemic conditions.

2.2 Integrating with Harvis and editing 3D stenosis geometry

The first part of establishing a new virtual intervention tool was to enable users to easily edit 3D coronary geometries to emulate intervention. Within Harvis, the user would place two dots to delineate the endpoints of the stenosis and

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

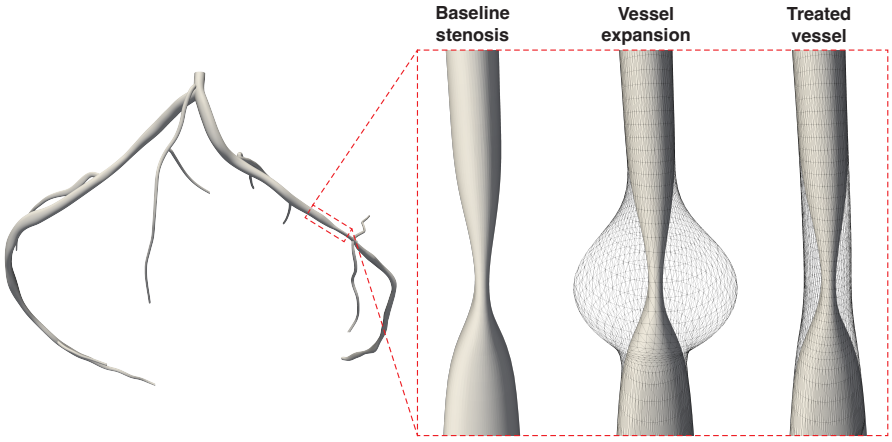

Fig. 2. Geometry modification process. Solid 3D geometry represents baseline anatomy and wireframes show mesh transformations used to treat the vessel. Vertices belonging to the stenosis were selected by the user and then translated following a sinusoidal function to simulate intervention. For the vessel expansion case, we show an extreme case displaying the sinusoidal nature of the modification process.

intervention to be simulated. The user-defined dots act as landmarks to identify all vertices in the mesh to be edited. The distance between the two dots was taken as the length of the stenosis for the intervention. Once the two dots are placed, a slider pops up that controls the radius of the stenosis. The mesh vertices between the dots would be displaced radially outward, normal to the vector between the two dots. The displacement of the vertices follows a sinusoidal function such that the vertices at the minimal luminal diameter, typically at the center, expand radially outward the most and the vertices closer to the endpoints do not expand much [2]. To illustrate this point, **Figure 2** demonstrates the vessel expansion process for one patient in our study cohort. As the geometry modification process involves only displacing vertices, no further post-processing was required for blood flow simulations.

From the user's perspective, stenosis modification occurs within the Harvis GUI (**Figure 3**). Users would first load the baseline 3D geometry of the coronary arteries before intervention in the format of `.off` into Harvis. At this point, the one-shot training process shown in **Figure 1A-B** would be completed. The pre-intervention FFR map could then be displayed. If the FFR map shows locations with FFR < 0.80, an intervention would be required, as shown in **Figure 3**. To perform a virtual intervention, two dots would be placed at the endpoints of a stenosis. A slider would appear to allow the user to control the diameter of the stenosis, and the distance between the two dots would control the length of the stenosis. The corresponding stenosis length and radius specified by the user would be relayed from the Harvis GUI to the HarVI machine learning backend to query the resulting post-intervention FFR map. This planning stage corresponds

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

to **Figure 1C** and can be performed any number of times without running any new 1D computational fluid dynamics simulations.

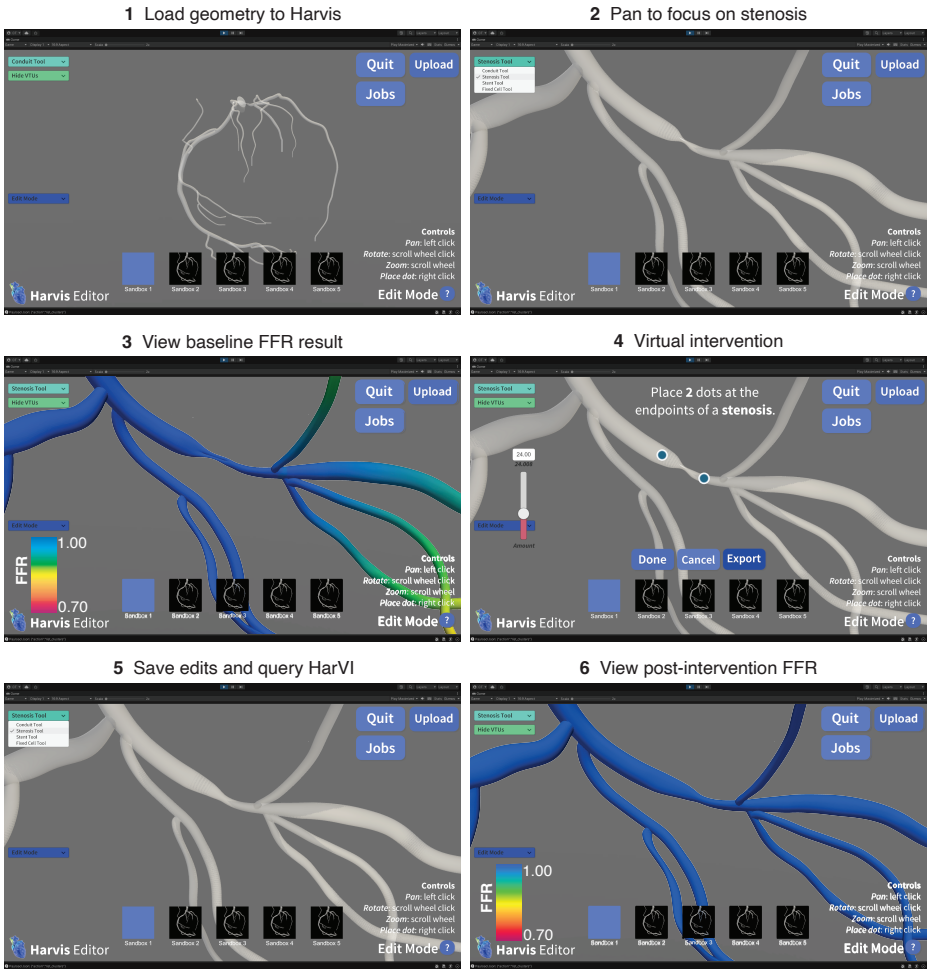

Fig. 3. Intervention planning workflow for the end user. Geometry modification was enabled using Harvis. The steps for virtual coronary intervention include: **1** loading baseline geometry into Harvis, **2** panning around the geometry to focus on the stenosis of interest, **3** evaluating baseline FFR results, **4** modifying stenosis to plan for intervention, **5** saving geometry modification and querying HarVI for hemodynamic feedback, and **6** evaluating and visualizing the resulting post-intervention FFR.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

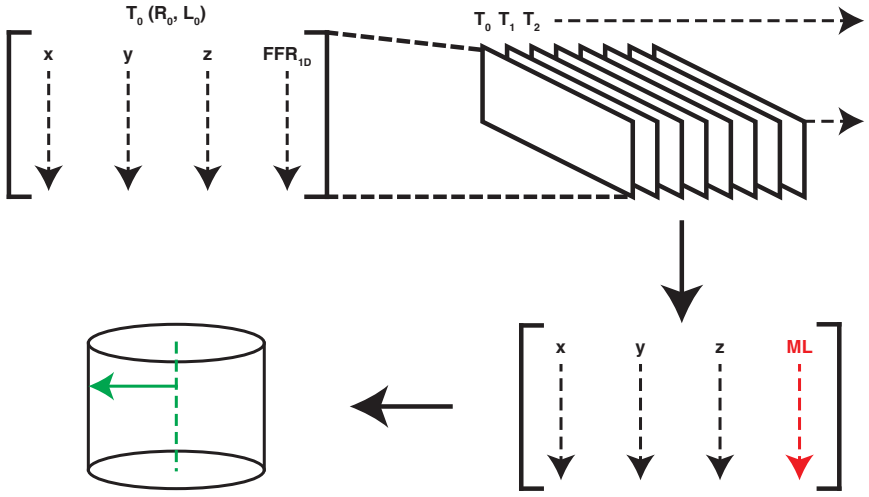

Fig. 4. Machine learning architecture for HarVI. Each predefined intervention scenario (T_N) for training consisted of creating a 2D matrix with 3D centerline data and corresponding FFR results from 1D simulations. Each of these intervention scenarios ran on 1 CPU, and all the 1D simulations were deployed simultaneously on a compute cluster in parallel. Each centerline point was fitted with a bivariate cubic-spline interpolation model. When queried a test radius and length, the resulting predictors would return a post-intervention FFR prediction per centerline point. The centerline points (interrupted green line) with post-intervention FFR predictions were then projected to the triangles of the modified geometry mesh (solid green arrow) to create an FFR map. Note that the cylinder shown is just an idealized case for illustration.

2.3 Establishing a real-time prediction model of post-intervention hemodynamics using machine learning

To improve clinical decision support, it is critical to quickly evaluate how various intervention strategies can influence the hemodynamics of that particular patient. In this work, we developed a machine learning model to quickly and accurately predict post-intervention FFR maps from a set of 1D CFD simulations used for training, directly addressing this need. The training set was constructed from a series of randomly sampled geometric modifications, in terms of stenosis length and radius, and then 1D CFD simulations were performed to obtain post-intervention FFR maps. Intuitively, the larger the set of 1D CFD simulations reserved for training, or in other words, the more samples used for training, the better the machine learning model would perform. Our objective in designing the interpolation model was to sample the parameter space of the radius and length of the stenosis more efficiently to minimize the training set while ensuring that the interpolation model would adequately capture the variance in the FFR distribution in response to intervention.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

To this end, we applied Latin Hypercube Sampling (LHS) from the `SALib` library that implements the algorithms presented in [23, 24]. To minimize sampling and achieve faster convergence, we opted for a quasirandom over pseudorandom sampling method so that samples would be more evenly distributed in the parameter space. We initialized the LHS for all patients from 90 % the minimal luminal diameter to 110 % the unstenosed radius (or radius if there was no stenosis). The radius not stenosed ($R_{unstenosed}$) was calculated according to the following equation:

$$\%_{stenosis} = \left(1 - \frac{R_{stenosed}}{R_{unstenosed}}\right) \times 100\% \tag{1}$$

where the degree of stenosis ($\%_{stenosis}$) and $R_{stenosis}$ were both labeled by intervention cardiologists. We added these bounds to buffer our predictions and minimize uncertainty at the extreme bounds. Similarly, lengths were sampled from 50 % to 150 % of the pre-intervention stenosis length labeled by interventional cardiologists. Radii and lengths were uniformly sampled between the standardized bounds we set for all patients and used as input to generate an LHS sample matrix.

The training matrix was applied to automatically modify the coronary geometries to reflect the scenarios in the training set. So far, all operations have been performed locally. To run the 1D simulations and generate the training set, all input files and coronary geometries were transferred to the Duke Compute Cluster (**Figure 4**). For each training instance (T_N), each 1D simulation was allocated one CPU and all simulations were deployed simultaneously. This parallel approach allowed us to complete the training process in the time it takes to run one 1D simulation. The post-intervention FFR was computed at all centerline points for all training scenarios. Centerline-parsed FFR data was transferred back to the local machine to construct the machine learning model. Bivariate cubic splines were fitted for each centerline point as a function of the radius of the stenosis and the lengths of the LHS. The specific bivariate cubic spline interpolation we applied was from `SciPy` and based on implementations from `FITPACK` [25, 26]. The centerline points were spaced at 100 μm resolution and compact enough to run efficiently on local machines. To generate the post-intervention FFR map, K-d trees were constructed to efficiently project the FFR data at the centerline to the nearest surface triangles on the 3D mesh.

2.4 Experimental protocol to establish and validate HarVI

To establish HarVI, the first pertinent step was to determine the level of sampling needed to accurately capture the hemodynamic changes in response to intervention. The number of samples impacts the accuracy of the machine learning model – the more samples, the better the interpolation. The objective was to find the minimum level of sampling that would result in the same FFR 2D heatmaps predicted using HarVI as the ground-truth with high sampling. We determined the level of sampling needed for all patients through convergence studies, where

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

we assumed a ground-truth sampling with $n = 250$ and compared it with interpolation results with progressively fewer samples. We created a $5,000 \times 5,000$ grid to uniformly evaluate the 2D parameter space between the minimum and maximum bounds of stenosis radius and length. In short, we made 25 million predictions for each sampling level. To compare with the ground-truth, we computed the maximum and average percentage error over all grid points. We set the error tolerances at 5 % for maximum error and 1 % for average error.

Once we determined the level of sampling needed across patients, we evaluated how closely FFR_{HarVI} compares to FFR_{1D}. We performed two test intervention cases that were held out during the training process. Specifically, we took treatment 1 (Rx 1) as $0.8R_{unstenosed}$ and $1.2L_{stenosis}$ and treatment 2 (Rx 2) as $R_{unstenosed}$ and $0.8L_{stenosis}$. All baseline geometries were modified to match these test intervention configurations. Post-intervention FFR was validated using continuous metrics and categorical metrics. The post-intervention FFR_{1D} ground-truth was generated by running 1D simulations with the corresponding test geometry modifications.

Lastly, we measured the total turnaround time for HarVI to compare with state-of-the-art FDA approved software. The turnaround time consisted of averages for all 50 patients for the patient-specific modeling pipeline (reconstruction and quality control) and the processing time needed for one-shot model training approach. Comparing the turnaround time of HarVI with established methods was important to gauge clinical translatability.

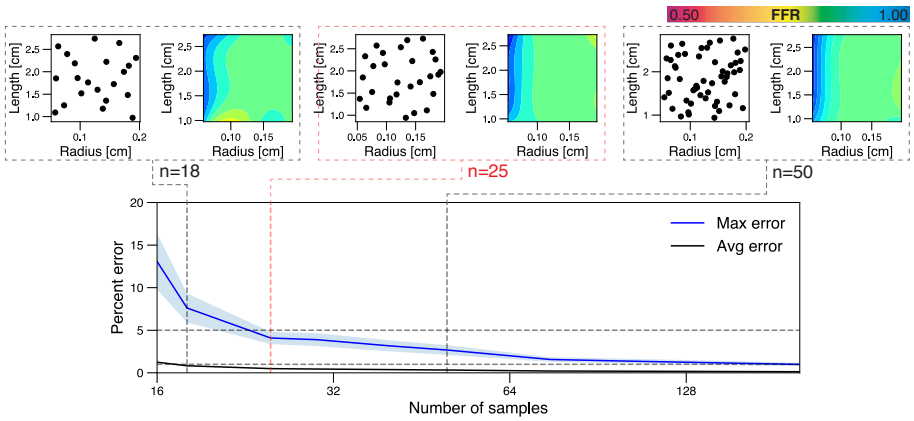

Fig. 5. Evaluating convergence in sampling stenosis radius and length. Ground-truth was assumed to have 250 samples. For all patients (n=50), we evaluated the maximum and average percentage errors for increasing number of samples as compared to the ground-truth. The area around the maximum error curve represents standard errors across all patients. For a representative patient, we showcase three insets at 18, 25, and 50 samples that show sampling on the left and resulting FFR 2D prediction heatmaps on the right.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

3 Results and Discussion

3.1 Few samples are needed to capture post-intervention fractional flow reserve accurately

To establish HarVI, it was important to determine the minimum level of sampling needed to capture changes in hemodynamics as a result of intervention. **Figure 5** demonstrates that the maximum error curve and average error curve decreases to below 5 % and 1 %, respectively, when using 25 samples. The 2D heatmaps presented in **Figure 5** show that sampling with 18 intervention configurations was insufficient. Although we show heatmaps for one representative case, there were apparent aberrations in the 2D heatmap that correspond to sites with a paucity of samples. 2D heatmaps for the 25 and 50 sample cases show nearly identical patterns without aberrations. Inter-patient variability was small for the maximum error curve and negligible for the average error curve. This result demonstrates that the convergence in sampling was predictable between patients. All lesions investigated in this study were intermediate focal stenoses, which may explain why a similar level of sampling was found to be sufficient for all patients.

3.2 Post-intervention fractional flow reserve predicted using HarVI agrees with 1D ground-truths

After determining the level of sampling needed for convergence, we validated HarVI against 1D ground-truths for two test intervention configurations. The

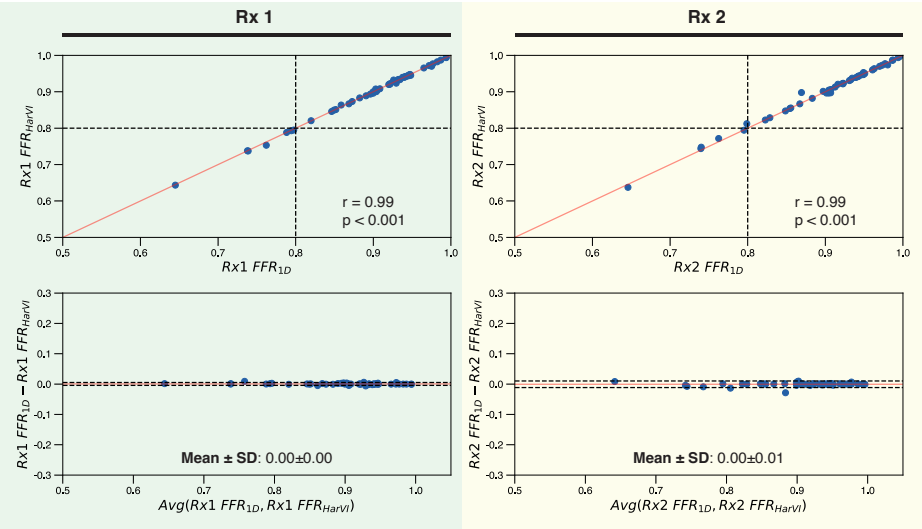

Fig. 6. Continuous validation of post-intervention FFR. For validation, two test cases (Rx 1 - left and Rx 2 - right) were used to compare $\mathrm{FFR_{HarVI}}$ against $\mathrm{FFR_{1D}}$. The top row show scatter plots with correlation scores and the bottom row show Bland-Altman plots with bias and imprecision scores.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

correlation coefficient was 0.99 ($p < 0.001$) and the bias was 0.0004 ± 0.0020 for Rx 1. The correlation coefficient was 0.99 ($p < 0.001$) and the bias was -0.0006 ± 0.0060 for Rx 2. In short, FFR_{HarVI} almost perfectly recovers FFR_{1D} for continuous metrics with negligible bias (**Figure 6**) in both test treatments. For context, FDA-approved software in the literature has reported a bias of around 0.01 ± 0.05 [27, 8] when comparing FFR computed using CFD to invasively measured FFR. The percentage discrepancy between FFR_{HarVI} and FFR_{1D} was on average 0.2 % and 0.3 % for Rx 1 and Rx 2, respectively.

The diagnostic performance of FFR_{HarVI} to discern stenoses with ischemia is summarized in **Table 1**. Diagnostic performance perfectly matched FFR_{1D} for Rx 1. Although there was a drop in sensitivity for Rx 2, there was, in fact, only one case misclassified in the Rx 2 test set. One case had post-intervention FFR_{1D} of 0.799 – right at the 0.80 threshold. The FFR_{HarVI} prediction was slightly above the threshold at 0.812, which resulted in a false negative. Except for the single case located right at the ischemic threshold, HarVI classified patients in full agreement with 1D CFD ground truths for both test cases.

3.3 End-to-end turnaround time to enable intervention planning within one working day

The clinical translatability of HarVI depends on the turnaround time required for the one-shot training process before real-time predictions could be made. FDA-approved software was estimated to require 24-48 hours of processing time, including the entire segmentation workflow, CFD simulations, and quality control, before allowing intervention planning [28, 29, 30, 31, 11, 10]. For the HarVI end-to-end pipeline, reconstructions took 15-30 minutes [18], quality control to ensure adequate reconstruction required another 30 minutes, and the processing time for all patients was on average 14 minutes (**Figure 7**). Therefore, the total turnaround time for HarVI was only 74 minutes on average, which fits in the door-to-balloon clinical standards of 90 minutes for patients with myocardial infarction [17]. After the one-shot training process, updating FFR_{HarVI} in response to intervention took an average of 0.4 s for all centerline points using local machines. If interventional cardiologists only want to view the resulting post-intervention FFR at the distal location of pressure measurement,

Table 1. Categorical validation of post-intervention FFR using 0.80 as the cut-off. Ground-truth was taken as post-intervention FFR_{1D}.

Metric	Rx 1 FFR_{HarVI}	Rx 2 FFR_{HarVI}
Sensitivity	100.0 (59.0-100.0)	83.3 (35.9-99.6)
Specificity	100.0 (91.8-100.0)	100.0 (92.0-100.0)
Positive predictive value	100.0 (59.0-100.0)	100.0 (47.8-100.0)
Negative predictive value	100.0 (91.8-100.0)	97.8 (88.0-99.6)
Accuracy	100.0 (92.9-100.0)	98.0 (89.4-99.9)
Area under the curve	1.00	1.00

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

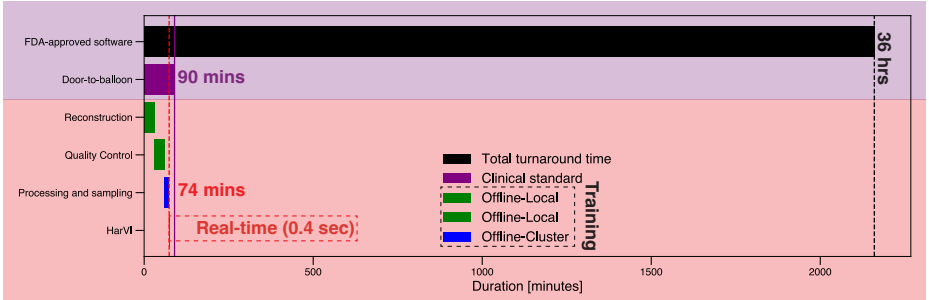

Fig. 7. Gantt chart showing timing breakdown for HarVI compared to clinical standards and FDA approved software. Top section (purple) shows estimated total turnaround time for FDA-approved intervention planning software and clinical standards for door-to-balloon time. Bottom section (pink) shows timing breakdowns (reconstruction, quality control, processing and sampling) for the one-shot training process needed for HarVI. After the training process, with an average turnaround time of 74 minutes, intervention planning in HarVI is near-instantaneous, requiring only 0.4 seconds per query (red box with interrupted lines).

FFR_{HarVI} updates in only 0.3 ms. HarVI would enable clinicians the ability to test a series of intervention scenarios pre-operatively within the same working day a patient is admitted.

4 Conclusion

Real-time virtual treatment technology is a recent innovation [11, 10, 31] and can give interventional cardiologists the option to interactively experiment with a variety of possible strategies to determine how best to relieve ischemia. However, the current turnaround time required to train FDA-approved software per patient ranges from 24-48 hours [29, 8]. This long turnaround time hinders the ability of cardiologists to test intervention strategies within the same day a patient is admitted and, perhaps more pertinently, prevents planning for patients who require immediate treatment, such as those with ST-elevated myocardial infarction. Moreover, current software in the literature relies on conventional 2D displays and is not deployed on immersive devices that help with intuitive evaluation of resulting post-intervention hemodynamics. To address both unmet needs, we established HarVI, a coronary intervention planning software that minimizes turnaround time to enable virtual intervention within one working day and leverages extended reality hardware.

In this study, we developed a machine learning model to predict FFR in response to intervention. We employed LHS to efficiently capture the variance in FFR in response to any possible intervention scenario in terms of radius and length of stenosis. For intuitive user interaction, interventions were captured as geometric modifications to 3D coronary meshes in virtual reality. We coupled the HarVI backend to the Harvis GUI, which is deployable on multiple XR headsets

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

with varying levels of immersion [2]. Enabling intuitive and immersive anatomical editing is important for more accurate treatment planning. Coronary arteries have high spatial complexity with many points of arterial branching and vessels with high tortuosity. The use of XR headsets for geometric modification and flow visualization improves intervention planning compared to traditional 2D displays [4, 32]. On the backend, we applied bivariate cubic spline interpolation models to predict post-intervention FFR from the training set of pre-configured intervention scenarios. Through convergence studies, we found that 25 samples were sufficient for the machine learning model to accurately predict FFR. To evaluate HarVI, we tested two held-out stenting procedures and compared them with 1D CFD ground truths. HarVI was able to recapitulate the 1D CFD results accurately (0.2-0.3 % error) as evaluated using continuous and categorical validation metrics for the 50 patients in the cohort. Finally, we recorded the time needed in critical parts of the HarVI pipeline and found that the average end-to-end turnaround time required to enable virtual intervention was 74 minutes.

In this study, we have taken initial steps toward integrating machine learning and extended reality for real-time planning of coronary interventions, focusing on focal lesions. The findings suggest a promising direction for expanding this approach to more complex diseases, such as serial and bifurcation lesions. HarVI offers a more streamlined method for interventional cardiologists to evaluate different treatment options within hours instead of days. The bivariate cubic spline model used here demonstrates strong accuracy. Future exploration of other machine learning-based approaches could further improve accuracy and reduce turnaround time. While the developments presented here are in their early stages, they lay the critical foundation for improving the efficiency of clinical decision making in cardiology.

5 Acknowledgements

The authors thank Samreen Mahmud, Aristotle Martin, and Runxin Wu for fruitful discussions and Juliet Jiang for her picture with the immersive headset. Computing support for this work came from the Duke Compute Cluster. The work of Cyrus Tanade was supported by the National Science Foundation Graduate Research Fellowship under Grant No. NSF GRFP DGE 1644868. The work of Amanda Randles was supported by the National Institute On Aging of the NIH under Award Number DP1AG082343 and the Duke/Duke-NUS Research Collaboration Pilot Project program. The content does not necessarily represent the official views of the NSF or NIH.

References

1. Vardhan, M., Shi, H., Gounley, J., James Chen, S., Kahn, A., Leopold, J., *et al.*: Investigating the role of VR in a simulation-based medical planning system for coronary interventions. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22, pp. 366–374 (2019)

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

2. Shi, H., Ames, J., Randles, A.: Harvis: an interactive virtual reality tool for hemodynamic modification and simulation. Journal of Computational Science **43**, 101091 (2020)
3. Vardhan, M., Shi, H., Urick, D., Patel, M., Leopold, J.A., Randles, A.: The role of extended reality for planning coronary artery bypass graft surgery. In: 2022 IEEE Visualization and Visual Analytics (VIS), pp. 115–119 (2022)
4. Shi, H., Vardhan, M., Randles, A.: The Role of Immersion for Improving Extended Reality Analysis of Personalized Flow Simulations. Cardiovascular Engineering and Technology **14**(2), 194–203 (2023)
5. Luffel, M., Sati, M., Rossignac, J., Yoganathan, A.P., Haggerty, C.M., Restrepo, M., *et al.*: SURGEM: A solid modeling tool for planning and optimizing pediatric heart surgeries. Computer-Aided Design **70**, 3–12 (2016)
6. Leo, J., Zhou, Z., Yang, H., Dass, M., Upadhayay, A., C Slesnick, T., *et al.*: Interactive cardiovascular surgical planning via augmented reality. In: Asian CHI Symposium 2021, pp. 132–135 (2021)
7. Yang, H., Mehta, P.D., Leo, J., Zhou, Z., Dass, M., Upadhayay, A., *et al.*: Evaluating Cardiovascular Surgical Planning in Mobile Augmented Reality. arXiv preprint arXiv:2208.10639 (2022)
8. Nørgaard, B.L., Leipsic, J., Gaur, S., Seneviratne, S., Ko, B.S., Ito, H., *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). Journal of the American College of Cardiology **63**(12), 1145–1155 (2014)
9. Tonino, P.A., De Bruyne, B., Pijls, N.H., Siebert, U., Ikeno, F., vant Veer, M., *et al.*: Fractional flow reserve versus angiography for guiding percutaneous coronary intervention. New England Journal of Medicine **360**(3), 213–224 (2009)
10. Sankaran, S., Lesage, D., Tombropoulos, R., Xiao, N., Kim, H.J., Spain, D., *et al.*: Physics driven real-time blood flow simulations. Computer Methods in Applied Mechanics and Engineering **364**, 112963 (2020)
11. Modi, B.N., Sankaran, S., Kim, H.J., Ellis, H., Rogers, C., Taylor, C.A., *et al.*: Predicting the physiological effect of revascularization in serially diseased coronary arteries: clinical validation of a novel CT coronary angiography–based technique. Circulation: cardiovascular interventions **12**(2), e007577 (2019)
12. Update, A.S.: Heart disease and stroke statistics–2017 update. Circulation **135**, e146–e603 (2017)
13. Fournier, S., Ciccarelli, G., Toth, G.G., Milkas, A., Xaplanteris, P., Tonino, P.A., *et al.*: Association of improvement in fractional flow reserve with outcomes, including symptomatic relief, after percutaneous coronary intervention. JAMA cardiology **4**(4), 370–374 (2019)
14. Katritsis, D.G., Theodorakakos, A., Pantos, I., Gavaises, M., Karcanias, N., Efstathopoulos, E.P.: Flow patterns at stented coronary bifurcations: computational fluid dynamics analysis. Circulation: Cardiovascular Interventions **5**(4), 530–539 (2012)
15. Antoniadis, A.P., Mortier, P., Kassab, G., Dubini, G., Foin, N., Murasato, Y., *et al.*: Biomechanical modeling to improve coronary artery bifurcation stenting: expert review document on techniques and clinical implementation. JACC: Cardiovascular Interventions **8**(10), 1281–1296 (2015)
16. Tanade, C., Chen, S.J., Leopold, J.A., Randles, A.: Analysis identifying minimal governing parameters for clinically accurate in silico fractional flow reserve. Frontiers in Medical Technology **4**, 1034801 (2022)

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4

17. Menees, D.S., Peterson, E.D., Wang, Y., Curtis, J.P., Messenger, J.C., Rumsfeld, J.S., *et al.*: Door-to-balloon time and mortality among patients undergoing primary PCI. New England Journal of Medicine **369**(10), 901–909 (2013)
18. Chen, S.J., Carroll, J.D.: 3-D reconstruction of coronary arterial tree to optimize angiographic visualization. IEEE transactions on medical imaging **19**(4), 318–336 (2000)
19. Pham, J., Wyetzner, S., Pfaller, M.R., Parker, D.W., James, D.L., Marsden, A.L.: svMorph: Interactive geometry-editing tools for virtual patient-specific vascular anatomies. Journal of Biomechanical Engineering **145**(3), 031001 (2023)
20. 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, pp. 1063–1074 (2013)
21. Tanade, C., Feiger, B., Vardhan, M., Chen, S.J., Leopold, J.A., Randles, A.: Global Sensitivity Analysis For Clinically Validated 1D Models of Fractional Flow Reserve. In: 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 4395–4398 (2021)
22. Feiger, B., Kochar, A., Gounley, J., Bonadonna, D., Daneshmand, M., Randles, A.: Determining the impacts of venoarterial extracorporeal membrane oxygenation on cerebral oxygenation using a one-dimensional blood flow simulator. Journal of biomechanics **104**, 109707 (2020)
23. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics **42**(1), 55–61 (2000)
24. Iman, R.L., Helton, J.C., Campbell, J.E.: An approach to sensitivity analysis of computer models: Part I—Introduction, input variable selection and preliminary variable assessment. Journal of quality technology **13**(3), 174–183 (1981)
25. Dierckx, P.: An algorithm for surface-fitting with spline functions. IMA Journal of Numerical Analysis **1**(3), 267–283 (1981)
26. Dierckx, P.: Curve and surface fitting with splines. Oxford University Press (1995)
27. Pellicano, M., Lavi, I., De Bruyne, B., Vaknin-Assa, H., Assali, A., Valtzer, O., *et al.*: Validation study of image-based fractional flow reserve during coronary angiography. Circulation: Cardiovascular Interventions **10**(9), e005259 (2017)
28. Torii, R., Yacoub, M.H.: CT-based fractional flow reserve: development and expanded application. Global Cardiology Science & Practice **2021**(3) (2021)
29. HeartFlow FFRCT for estimating fractional flow reserve from coronary CT angiography — Guidance — NICE, eng (2017). Publisher: NICE.
30. Gaur, S., Achenbach, S., Leipsic, J., Mauri, L., Bezerra, H.G., Jensen, J.M., *et al.*: Rationale and design of the HeartFlowNXT (HeartFlow analysis of coronary blood flow using CT angiography: NeXt sTeps) study. Journal of cardiovascular computed tomography **7**(5), 279–288 (2013)
31. Taylor, C.A., Petersen, K., Xiao, N., Sinclair, M., Bai, Y., Lynch, S.R., *et al.*: Patient-specific modeling of blood flow in the coronary arteries. Computer Methods in Applied Mechanics and Engineering **417**, 116414 (2023)
32. Pekkan, K., Whited, B., Kanter, K., Sharma, S., De Zelicourt, D., Sundareswaran, K., *et al.*: Patient-specific surgical planning and hemodynamic computational fluid dynamics optimization through free-form haptic anatomy editing tool (SURGEM). Medical & biological engineering & computing **46**, 1139–1152 (2008)

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63749-0_4