Journal of Computational Applied Mechanics 2027, 58(1): 90-110

DOI: 10.22059/jcamech.2026.426509.2016

RESEARCH PAPER

A Hybrid Variational Quantum Circuit–Physics-Informed Transformer for Prediction of Cerebral Aneurysm Hemodynamics

Amir Kheiri [a], M. Barzegar Gerdroodbary [b, c, *], Seyedeh Zahra Sajadi [d], Yousef Azizi [e, f], Peiman Valipour [g, †]

[a] *Department of Mechanical Engineering, K.N. Toosi University of Technology (KNTU), Tehran, Iran*

[b] *Department of Mechanical Engineering, Iranian Research Organization for Science and Technology (IROST), Tehran, Iran*

[c] *Department of Electromechanical Engineering, C-MAST-Center for Mechanical and Aerospace Science and Technology, Universidade da Beira Interior, Covilha, Portugal*

[d] *Department of Computer Science, University of Mazandaran, Iran*

[e] *Department of Mechanical Engineering, Bandar Anzali Branch, Islamic Azad University, Bandar Anzali, Iran*

[f] *Foulad Sepehr Sabalan, Ardabil, Iran*

[g] *Department of Textile Engineering, Apparel and Fashion, QaS.C., Islamic Azad University, Qaemshahr, Iran*

Abstract

Accurate prediction of cerebral aneurysm hemodynamics is important for understanding complex blood-flow patterns and evaluating wall-related hemodynamic characteristics. Although computational fluid dynamics (CFD) provides detailed spatial and temporal information, its computational cost can limit its application in repeated patient-specific simulations. This study proposes a hybrid quantum-classical framework, termed the VQC-PI-Transformer, for prediction of cerebral aneurysm hemodynamics. The proposed framework integrates singular value decomposition (SVD)-based dimensionality reduction with a trainable variational quantum circuit (VQC) and a physics-informed Transformer. Initially, SVD is applied to high-dimensional CFD snapshots to obtain a compact representation based on dominant spatial modes and their corresponding temporal coefficients. The reduced-order features are subsequently processed by a trainable VQC, which employs parameterized quantum rotation gates and entangling operations to generate quantum-derived features. These features are integrated into a classical Transformer to learn temporal dependencies and predict the evolution of reduced-order hemodynamic representations. The training process incorporates data-fitting objectives and relevant physical constraints and boundary conditions to guide the predictions toward physically consistent solutions. The proposed framework is evaluated using four unruptured, patient-specific middle cerebral artery aneurysm geometries, considering wall shear stress (WSS), time-averaged wall shear stress (TAWSS), oscillatory shear index (OSI), and pressure. The influence of qubit count and variational circuit depth is systematically investigated, with

* Corresponding author. *E-mail address:* m.barzegargerdroodbary@ubi.pt

† Corresponding author. *E-mail address:* pe.valipour@iau.ac.ir

favorable prediction performance obtained using an eight-qubit circuit with two variational layers. Comparisons with CFD reference results and a classical physics-informed Transformer demonstrate that the proposed hybrid architecture provides closer agreement with CFD in most of the evaluated case–parameter combinations. Improvements are particularly encouraging for pressure and TAWSS, while OSI exhibits relatively stable prediction errors across the investigated geometries. Overall, the findings demonstrate the potential of integrating trainable variational quantum circuits into physics-informed reduced-order models for patient-specific hemodynamic prediction and motivate further investigation of hybrid quantum-classical approaches in computational biomechanics.

Keywords: Quantum machine learning; Cerebral aneurysm; computational fluid dynamics; reduced-order modeling; singular value decomposition; variational quantum circuit;

1. Introduction

Intracranial aneurysms are localized dilations of cerebral arteries whose geometry and surrounding flow conditions influence the distribution of hemodynamic forces along the vessel wall. Computational fluid dynamics (CFD) is widely used to investigate these conditions because it provides spatially and temporally resolved estimates of velocity, pressure, wall shear stress, and related quantities [1, 2]. Such information can help characterize flow features that are difficult to measure directly in patient-specific vascular geometries.

Among the quantities commonly investigated in aneurysm hemodynamics, wall shear stress (WSS), time-averaged wall shear stress (TAWSS), and oscillatory shear index (OSI) describe complementary aspects of the blood-flow environment at the vessel wall. WSS represents the tangential traction exerted by flowing blood, whereas OSI characterizes the directional oscillation of the WSS vector during a cardiac cycle. These quantities have been used to study disturbed flow and variations in aneurysm-wall hemodynamics [2, 3]. Accurate estimation requires the underlying velocity field and, for cycle-averaged quantities, adequate temporal information. Despite their value, patient-specific CFD simulations require geometry preparation, mesh generation, boundary-condition specification, and numerical solution of the governing flow equations. These steps can become computationally demanding when several geometries, boundary conditions, or time-dependent scenarios must be evaluated. This motivates surrogate models capable of approximating high-fidelity CFD solutions at lower computational cost [4, 5].

Reduced-order modeling offers one such approach. Proper orthogonal decomposition (POD), often implemented through singular value decomposition (SVD), identifies a low-dimensional basis from a collection of high-dimensional snapshots [6, 7]. Instead of predicting every spatial degree of freedom directly, a model can learn the temporal evolution of a smaller set of modal coefficients. The spatial field is then reconstructed using the retained basis. This can reduce the dimensionality of the learning task while preserving dominant structures present in the CFD data [8, 9].

Machine learning models can be used to predict these reduced-order dynamics [10]. Transformer architectures employ self-attention to represent relationships across input sequences [11]. Physics-informed neural networks introduce equations, boundary conditions, or other physical information into the training objective [11,15]. Combining temporal prediction with physics-informed constraints offers a way to use both CFD-derived observations and prior physical knowledge [12, 13].

Variational quantum circuits provide a trainable quantum feature transformation. Classical input features are encoded into quantum states, transformed by parameterized quantum operations, and mapped back to classical features through measurements. The circuit parameters can be optimized using a classical loss function and a compatible gradient-computation method [14]. Previous work has examined hybrid quantum-classical physics-informed networks for partial differential equations and flow problems, reporting promising but problem-dependent results [14-16]. These studies suggest that additional quantum expressivity does not automatically guarantee improved prediction accuracy.

In this study, a trainable VQC is integrated with a physics-informed Transformer following SVD/POD-based dimensionality reduction. The proposed VQC-PI-Transformer uses the compact representation of the CFD data for temporal learning, while quantum-derived features are fused with classical inputs. Boundary conditions and physical constraints are included in training in accordance with the implemented objective.

The main contributions are: (i) a reduced-order hemodynamic prediction workflow connecting SVD, a trainable VQC, and a physics-informed Transformer; (ii) a controlled investigation of qubit count and circuit depth for

predicting WSS, OSI, and pressure; (iii) evaluation on four patient-specific aneurysm geometries; and (iv) comparison against a classical physics-informed Transformer and the CFD reference. The objective is to test whether trainable quantum-derived features provide a measurable benefit under comparable evaluation conditions, rather than assuming a quantum component must improve every target.

2. Related work

1.1. ***CFD-based hemodynamic analysis***

CFD has been widely used to evaluate three-dimensional flow patterns and wall-related quantities in intracranial aneurysms. Patient-specific computational workflows combine vascular geometry, numerical discretization, boundary conditions, and fluid-flow equations to calculate hemodynamic fields [17, 18]. The AneuRisk project [19] has supported multidisciplinary investigations of aneurysm morphology, hemodynamics, and vascular-wall mechanics and provides geometries used in computational studies. WSS-related quantities are particularly relevant to aneurysm-wall flow. Arzani and Shadden investigated the spatial and temporal characteristics of WSS in aneurysmal flow and correlations between WSS-derived measures [3]. In surrogate modeling, CFD serves both as the source of training data and as the numerical reference for evaluation. Conclusions drawn from a surrogate therefore depend on the quality and consistency of the CFD data, including its physical assumptions, numerical settings, and boundary conditions [20]. CFD modeling was conducted via ANSYS-Fluent 23.0 software [21].

1.2. Reduced-order modeling and physics-informed machine learning

SVD/POD is an established approach for reducing the dimensionality of flow data. Berkooz, Holmes, and Lumley reviewed POD for fluid-flow analysis, while Sirovich developed the method of snapshots for extracting dominant coherent structures [6]. Retained modes form a low-dimensional basis that approximates the high-dimensional field.

Physics-informed neural networks combine observed data with physical equations or constraints through a loss function [22]. Rather than minimizing only discrepancies between predicted and reference values, a physics-informed objective can penalize violations of governing equations, boundary conditions, or initial conditions. The exact form depends on which constraints are implemented. Transformers were introduced as attention-based sequence models and have since been adapted to time-dependent prediction tasks [23]. In the present work, the Transformer processes sequences of reduced-order and quantum-derived features. It predicts the next reduced-order state, while physics-informed terms guide training was supported by the implementation.

1.3. Quantum and hybrid quantum-classical physics-informed models

Variational quantum algorithms use parameterized quantum circuits whose parameters are updated through classical optimization [14]. In hybrid machine learning, such circuits can provide trainable feature transformations while classical neural networks perform subsequent processing and prediction. Trahan, Loveland, and Dent investigated quantum and hybrid quantum-classical PINNs using a PennyLane simulator and reported that selected hybrid configurations achieved accuracy comparable to or better than classical models in some settings [15]. Berger, Hosters, and Möller introduced trainable quantum embeddings for nonlinear PDEs [16]. Leong and colleagues evaluated hybrid quantum physics-informed networks for flow problems of varying complexity and found that performance depended on the flow characteristics and architecture [24]. Together, these studies motivate controlled evaluation rather than an assumption of universal improvement.

Hybrid quantum-classical Transformers have also been explored for time-series forecasting [25]. Their circuit placement, targets, and training objectives differ from the present work. Here, the circuit transforms reduced-order and physical inputs into quantum-derived features, fuses them with classical features, and uses a physics-informed Transformer for cerebral aneurysm hemodynamic prediction.

3. Methodology

3.1. Overview of the proposed framework

The proposed VQC-PI-Transformer is a hybrid quantum-classical framework designed to predict hemodynamic characteristics in patient-specific cerebral aneurysm geometries. The framework combines a trainable variational

quantum circuit (VQC) with a physics-informed Transformer to integrate quantum-derived features into a temporal prediction architecture. As illustrated in Figures 1 and 2, the overall methodology consists of several interconnected stages, beginning with the dimensionality reduction of CFD data and followed by quantum feature extraction, feature fusion, temporal prediction, and reconstruction of the predicted hemodynamic fields.

Figure 1 illustrates the architecture of the trainable VQC incorporated into the proposed framework. The circuit consists of a sequence of parameterized single-qubit rotation gates and entangling operations. First, the input features are encoded into the quantum states through rotation gates. In the implemented circuit, trainable R_X, R_Y, and R_Z gates are applied to the qubits, allowing the quantum state to be adjusted through learnable parameters during model training. These parameterized gates provide the circuit with the flexibility to transform the encoded input states into different quantum representations.

To enable interactions between qubits, controlled-NOT (CNOT) gates are arranged in a ring-connected entanglement structure. This arrangement introduces correlations between neighboring qubits and allows information encoded in individual qubits to be combined through multi-qubit operations. Following the parameterized and entangling operations, measurements are performed using the Pauli-X, Pauli-Y, and Pauli-Z operators. The resulting expectation values are extracted as classical numerical features and form the output of the quantum circuit. Therefore, the VQC acts as a trainable feature-transformation component, mapping the input data into a quantum-derived feature representation that can subsequently be processed by the classical neural network.

Figure 2 presents the complete workflow of the proposed VQC-PI-Transformer. Initially, the CFD snapshots are processed using singular value decomposition (SVD) to obtain a reduced representation of the original hemodynamic fields. This decomposition provides a set of spatial modes and their corresponding time-dependent coefficients, allowing the dominant flow information to be represented using a smaller number of variables. The resulting reduced coefficients, together with the selected physical input features, are then provided to the trainable VQC.

The quantum circuit processes the encoded inputs and generates quantum-derived features through its parameterized operations and measurements. These features are subsequently combined with the corresponding classical input features through a feature-fusion stage. The resulting hybrid representation is passed to the physics-informed Transformer, which learns the temporal relationships within the input sequences and predicts the future modal coefficients. During training, the model parameters are optimized according to the implemented data-driven and physics-informed objectives.

Finally, the predicted modal coefficients are combined with the corresponding spatial modes to reconstruct the hemodynamic fields in their original spatial representation. The reconstructed predictions, including OSI, WSS, and pressure, are then evaluated against the reference CFD results using quantitative error measures and spatial contour comparisons. Through this integration, the proposed framework combines the feature-transformation capability of a trainable quantum circuit with the temporal-learning capability of a physics-informed Transformer, while retaining the ability to reconstruct and evaluate the predicted physical fields.

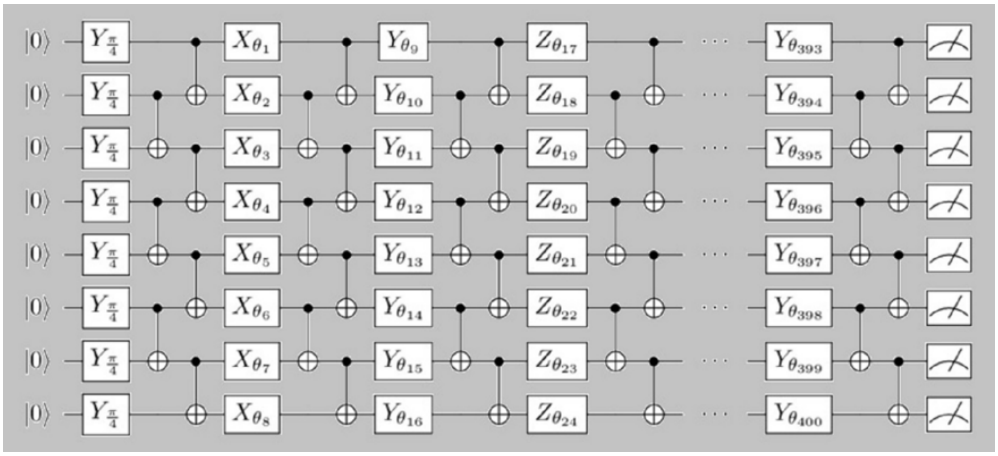


Fig. 1: Architecture of the trainable variational quantum circuit used in the VQC-PI-Transformer, showing trainable rotation gates, ring-connected CNOT entanglement, and quantum measurements.

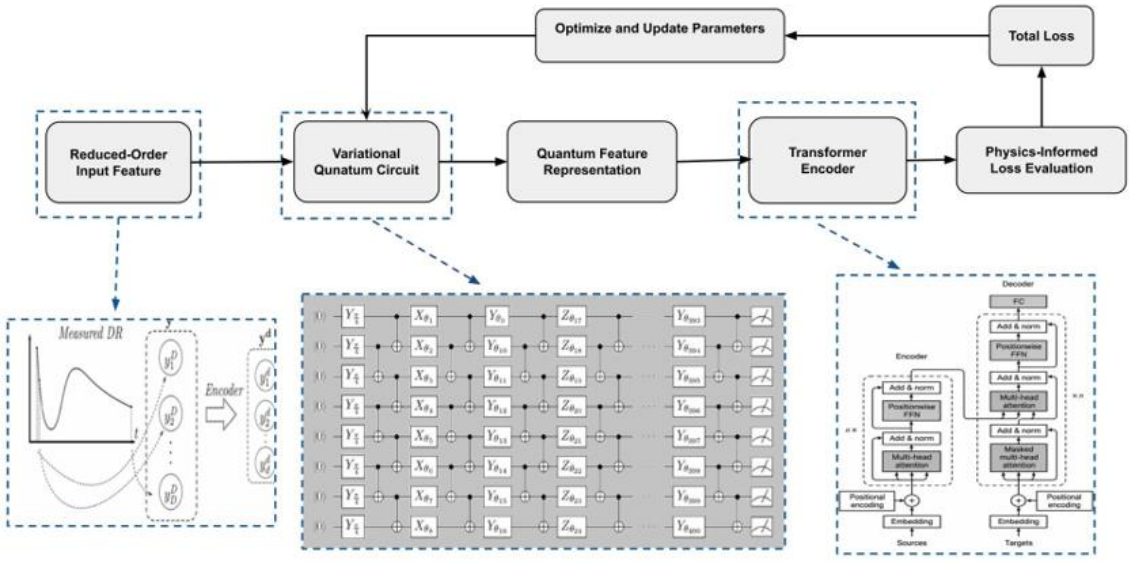


Fig. *2*: Workflow of the VQC-PI-Transformer. CFD snapshots are reduced using SVD; reduced coefficients and physical features are encoded by a trainable VQC; quantum-derived and classical features are fused and processed by a physics-informed Transformer.

3.2. SVD/POD-based dimensionality reduction

Let the spatial CFD field at time t_k be represented by $\mathbf{x}(t_k) \in \mathbb{R}^{N_x}$, where N_x is the number of spatial degrees of freedom. Collecting N_t snapshots gives the snapshot matrix

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}(t_1) & \boldsymbol{x}(t_2) & \cdots & \boldsymbol{x}(t_{N_t}) \end{bmatrix} \qquad (1)$$

The singular value decomposition is

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T \qquad (2)$$

where $\mathbf{U}$ contains spatial modes, $\boldsymbol{\Sigma}$ contains singular values, and $\mathbf{V}$ describes their temporal contributions. Retaining the first r modes gives the approximation

$$\boldsymbol{X} \approx \boldsymbol{U}_r\boldsymbol{\Sigma}_r\boldsymbol{V}_r^T \qquad (3)$$

The associated reduced coefficients can be written as

$$\boldsymbol{A}_r = \boldsymbol{U}_r^T\boldsymbol{X}, \qquad \widehat{\boldsymbol{X}} = \boldsymbol{U}_r\boldsymbol{A}_r \qquad (4)$$

For centered snapshots, the temporal mean must be added back during reconstruction. The fraction of snapshot energy retained by the first r modes is

$$E_r = \frac{\sum_{i=1}^{r} \sigma_i^2}{\sum_{i=1}^{R} \sigma_i^2} \qquad (5)$$

Where R is the number of singular values in the decomposition. Cumulative energy is a useful rank-selection criterion, but it should be interpreted together with reconstruction error because localized features may remain sensitive to truncation. A normalized reconstruction error at time t_k can be written as

$$e_{ROM}(t_k) = \frac{\| \boldsymbol{x}(t_k) - \widehat{\boldsymbol{x}}(t_k) \|_2}{\| \boldsymbol{x}(t_k) \|_2 + \varepsilon} \qquad (6)$$

where ε is a small stabilizing constant. This measures error from dimensionality reduction before temporal prediction is applied.

3.3. Trainable variational quantum circuit

3.3.1. Feature encoding

At time t, the circuit input includes the reduced-order coefficients and relevant physical features. Denote this vector by

$$\boldsymbol{s}_t = [\boldsymbol{a}(t),\ t,\ \boldsymbol{b}(t)] \qquad (7)$$

where $\mathbf{a}(t)$ is the reduced-order state and $\mathbf{b}(t)$ denotes the additional inputs used in the implemented model, such as prescribed inlet or pressure information. A classical projection maps this vector to a bounded set of n_q encoding angles:

$$\boldsymbol{\alpha}_t = f_{enc}(\boldsymbol{s}_t) \in \mathbb{R}^{n_q} \qquad (8)$$

Using Y-axis angle embedding, the initial state is

$$\left|\psi_t^{(0)}\right\rangle = U_{enc}(\boldsymbol{\alpha}_t)|0\rangle^{\otimes n_q} \quad (9)$$

3.3.2. Variational gates and entanglement

The VQC uses trainable rotation gates R_X, R_Y, and R_Z and ring-connected controlled-NOT (CNOT) gates. A parameterized rotation about axis $j \in \{x, y, z\}$ is defined by

$$R_j(\theta) = exp\left(-\frac{i\theta\sigma_j}{2}\right) \qquad (10)$$

where σ_j is the corresponding Pauli operator. For n_q qubits and L_q variational layers, the complete circuit can be written as

$$U(\boldsymbol{\theta}) = U_{L_q}\left(\boldsymbol{\theta}_{L_q}\right)\cdots U_2(\boldsymbol{\theta}_2)U_1(\boldsymbol{\theta}_1)U_{enc}(\boldsymbol{\alpha}_t) \quad (11)$$

The resulting state is

$$|\psi_t(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes n_q} \qquad (12)$$

Unlike a fixed quantum reservoir, the variational rotations are trainable and are optimized jointly with the classical network.

3.3.3. Quantum measurements and feature output

The circuit outputs expectation values of Pauli X, Y, and Z observables for each qubit:

$$z_i^X(t) = \langle\psi_t|\sigma_i^X|\psi_t\rangle, \qquad z_i^Y(t) = \langle\psi_t|\sigma_i^Y|\psi_t\rangle, \quad z_i^Z(t) = \langle\psi_t|\sigma_i^Z|\psi_t\rangle \qquad (13)$$

The complete quantum feature vector is

$$\boldsymbol{z}(t) = \left[z_1^X, \dots, z_{n_q}^X,\ z_1^Y, \dots, z_{n_q}^Y,\ z_1^Z, \dots, z_{n_q}^Z\right]^T \in \mathbb{R}^{3n_q} \qquad (14)$$

For eight qubits, the circuit produces 24 measured features. With three trainable rotation angles per qubit per layer, the number of trainable circuit parameters is

$$N_{VQC} = 3n_q L_q \qquad (15)$$

The eight-qubit, two-layer configuration therefore contains 48 trainable rotation parameters.

3.4. Physics-informed Transformer

3.4.1. Fusion of quantum and classical features

The VQC measurements do not directly represent the complete spatial hemodynamic field. They form an additional feature representation that is combined with the reduced-order and physical inputs. Let

$$\boldsymbol{h}_t^c = f_c(\boldsymbol{s}_t), \qquad \boldsymbol{h}_t^q = f_q(\boldsymbol{z}_t) \qquad (16)$$

The fused input token is

$$\boldsymbol{h}_t^{(0)} = f_{fuse}\left(\left[\boldsymbol{h}_t^c; \boldsymbol{h}_t^q\right]\right) + \boldsymbol{p}_t \qquad (17)$$

where [;] denotes concatenation and $\mathbf{p}_t$ is positional encoding. This provides the Transformer with both the original input information and quantum-derived features.

3.4.2. Temporal attention

For a sequence representation **H**, the query, key, and value matrices are

$$\boldsymbol{Q} = \boldsymbol{H}\boldsymbol{W}_Q \qquad \boldsymbol{K} = \boldsymbol{H}\boldsymbol{W}_K \qquad \boldsymbol{V} = \boldsymbol{H}\boldsymbol{W}_V \qquad (18)$$

Scaled dot-product attention is

$$Attention(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) = softmax\left(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d_k}}\right)\boldsymbol{V} \qquad (19)$$

where d_k is the key dimension. Multi-head attention uses multiple learned projections to represent different relationships in the temporal sequence [23]. The encoder combines attention and feed-forward operations with residual connections, normalization, and dropout. The prediction head maps the temporal representation to the next reduced-order state:

$$\hat{\boldsymbol{a}}(\boldsymbol{t}+\boldsymbol{1}) = \boldsymbol{f}_{\boldsymbol{\phi}}\left(\boldsymbol{h}_{\boldsymbol{t}-\boldsymbol{L}+\boldsymbol{1}}^{(\boldsymbol{0})}, \ldots, \boldsymbol{h}_{\boldsymbol{t}}^{(\boldsymbol{0})}\right) \qquad (20)$$

where ϕ represents the classical Transformer parameters and L is the input window length. The predicted modal coefficients are then reconstructed using the retained SVD basis.

3.5. Physics-informed training objective

Physics-informed learning combines a data-fitting objective with penalties associated with the physical conditions represented in the model [22]. The exact loss terms must match the implemented boundary conditions and residuals. A reduced-order data-fitting term can be written as

$$\mathcal{L}_{data} = \frac{1}{N_s r}\sum_{k=1}^{N_s} || \hat{\boldsymbol{a}}(t_k) - \boldsymbol{a}(t_k)||_2^2 \qquad (21)$$

where N_s is the number of samples and r is the number of retained modes. If the implementation computes the data loss in physical space, the loss should instead be defined over the reconstructed fields. A general boundary-condition penalty is

$$\mathcal{L}_{BC} = \frac{1}{N_b}\sum_{j=1}^{N_b} || \hat{\boldsymbol{y}}(\boldsymbol{x}_j, t_j) - \boldsymbol{y}_{BC}(\boldsymbol{x}_j, t_j)||_2^2 \qquad (22)$$

where the prescribed values correspond to the actual inlet, outlet, and wall conditions used in the CFD model and the learning implementation. For incompressible Newtonian flow with constant viscosity, the continuum equations are commonly written as

$$\nabla \cdot \boldsymbol{u} = 0 \qquad (23)$$

$$\rho\left(\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u} \cdot \nabla \boldsymbol{u}\right) = -\nabla p + \mu \nabla^2 \boldsymbol{u} \qquad (24)$$

where u is velocity, p is pressure, ρ is density, and μ is dynamic viscosity. These equations are included here as the governing-flow formulation. Continuity and momentum residuals are implemented; a corresponding term can be expressed as

$$\mathcal{L}_{phys} = \frac{1}{N_f}\sum_{j=1}^{N_f}\left(\| r_c(\boldsymbol{x}_j, t_j) \|_2^2 + \| \boldsymbol{r}_m(\boldsymbol{x}_j, t_j) \|_2^2\right) \qquad (25)$$

where r_c and $\mathbf{r}_m$ denote the continuity and momentum residuals. Initial conditions are explicitly enforced, and their penalty is

$$\mathcal{L}_{IC} = \frac{1}{N_i}\sum_{j=1}^{N_i} \| \hat{\boldsymbol{y}}(\boldsymbol{x}_j, 0) - \boldsymbol{y}_0(\boldsymbol{x}_j) \|_2^2 \qquad (26)$$

A general combined objective is

$$\mathcal{L}_{total} = \lambda_d \mathcal{L}_{data} + \lambda_b \mathcal{L}_{BC} + \lambda_i \mathcal{L}_{IC} + \lambda_p \mathcal{L}_{phys} \qquad (27)$$

where the weights balance the terms. The quantum and classical parameters are optimized jointly:

$$(\boldsymbol{\theta}^*, \phi^*) = \underset{\boldsymbol{\theta},\phi}{argmin}\mathcal{L}_{total} \qquad (28)$$

The quantum circuit is evaluated using PennyLane's simulator and the PyTorch interface [26][24].

3.6. Reconstruction and hemodynamic quantities

The predicted modal coefficients are mapped back to the physical field using the retained SVD basis. For a field represented by snapshots in $\mathbf{X}$, the reconstructed prediction is obtained from the corresponding truncated basis expansion. The resulting spatial-temporal fields can then be compared with the CFD reference. Under the usual no-slip, constant-viscosity formulation, WSS magnitude can be related to the near-wall tangential velocity gradient by

$$\tau_w = \mu \left\| \frac{\partial \boldsymbol{u}_t}{\partial n} \right\| \qquad (29)$$

where $\mathbf{u}_t$ is tangential velocity and n denotes the wall-normal direction. TAWSS over a cardiac cycle of duration T is

$$TAWSS(\boldsymbol{x}) = \frac{1}{T}\int_0^T \| \boldsymbol{\tau}_w(\boldsymbol{x}, t) \| \, dt \qquad (30)$$

and OSI is

$$OSI(\boldsymbol{x}) = \frac{1}{2}\left[1 - \frac{\left\|\int_0^T \boldsymbol{\tau}_w\,(\boldsymbol{x},t)\,dt\right\|}{\int_0^T \|\,\boldsymbol{\tau}_w(\boldsymbol{x},t)\,\|\;dt}\right] \qquad (31)$$

These measures follow commonly used definitions of wall-shear indices [17, 18, 27, 28].

4. Implementation and training

4.1. Patient-specific cases and CFD data

The study considers four unruptured, patient-specific terminal-type middle cerebral artery aneurysm geometries, designated Cases A–D, as selected from the Aneurisk project [19]. CFD provides time-dependent snapshots for the hemodynamic variables evaluated in the study, including WSS, OSI, and pressure [29].

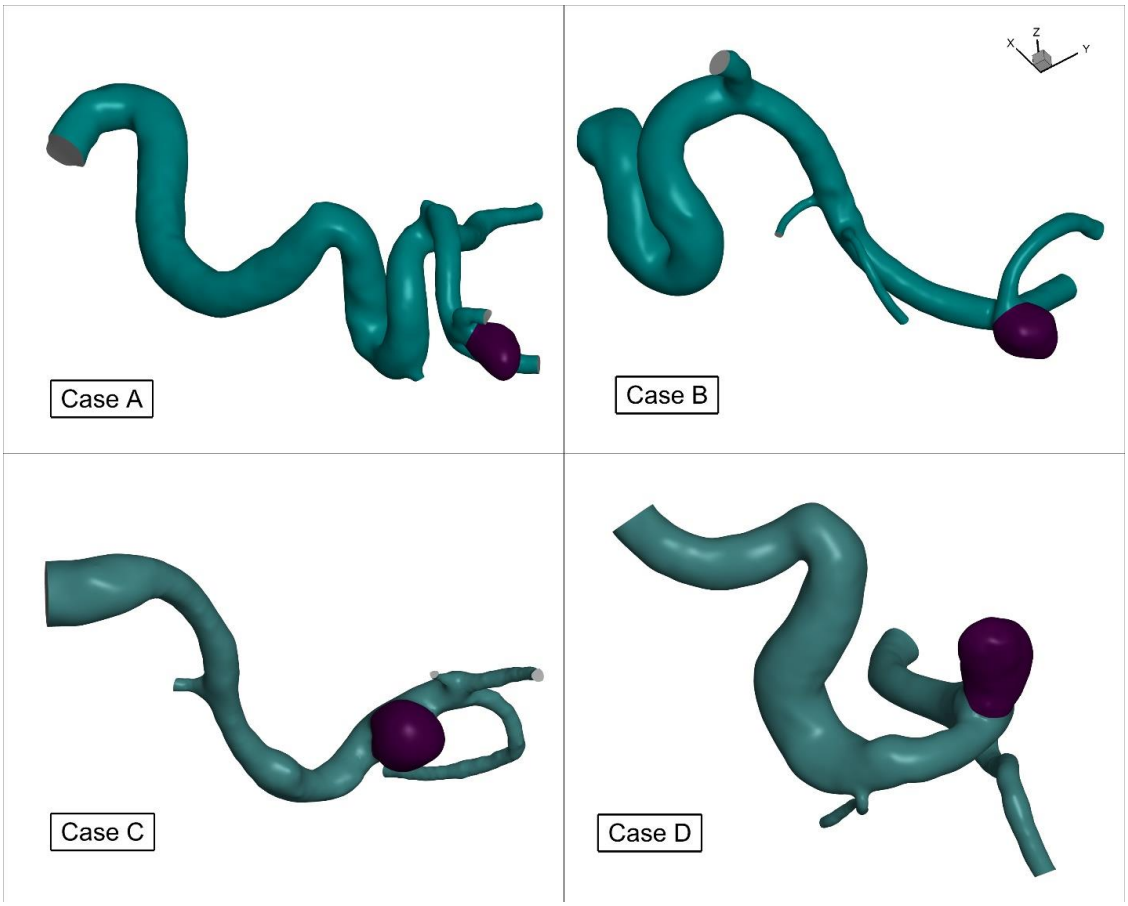


Figure 3. Patient-specific three-dimensional geometries of the four unruptured terminal-type middle cerebral artery aneurysm cases(A–D)

4.2. Reduced-order preprocessing

SVD is applied to the CFD data for each case and hemodynamic variable. Cumulative retained energy and reconstruction error are used to evaluate the trade-off between compression and representation accuracy. The resulting modal coefficients form the reduced-order representation used by the temporal predictor.

4.3. Quantum circuit and Transformer configuration

The VQC uses Y-axis angle embedding, trainable R_X, R_Y, and R_Z rotations, ring-connected CNOT gates, and Pauli X, Y, and Z expectation values. The nominal configuration described in Table 1 uses eight qubits and two variational layers, producing 24 quantum features and 48 trainable rotation parameters. The Transformer uses a model dimension of 256, eight attention heads, and four encoder layers. The feed-forward dimension is 1,024, and dropout is 0.10. Quantum-derived features are fused with classical reduced-order and physical inputs before temporal processing.

4.4. Training protocol and evaluation metrics

AdamW is used as the optimizer with an initial learning rate of 2×10^{-4} and weight decay of 10^{-5}. Training is limited to 120 epochs with a batch size of 32 and early-stopping patience of 20 epochs. Gradient clipping uses a maximum norm of 1.0. For reference values y_i and predictions $\hat{y}_i$, RMSE and MAE are defined as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \qquad (32)$$

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i| \qquad (33)$$

where $\bar{y}$ is the mean reference value over the evaluation samples. Relative L_2 error is calculated as

$$E_{L_2} = \frac{\| \boldsymbol{y} - \hat{\boldsymbol{y}} \|_2}{\| \boldsymbol{y} \|_2 + \varepsilon} \qquad (34)$$

5. Results

5.1. SVD reduction and reconstruction

Figure 4 presents cumulative retained energy as a function of the number of modes for a representative case. The curve illustrates how the selected rank determines the fraction of snapshot energy retained in the reduced representation. Figure 5 shows the time-dependent reconstruction error for Case A for the evaluated hemodynamic variables. These analyses characterize the error introduced by dimensionality reduction before the temporal prediction stage.

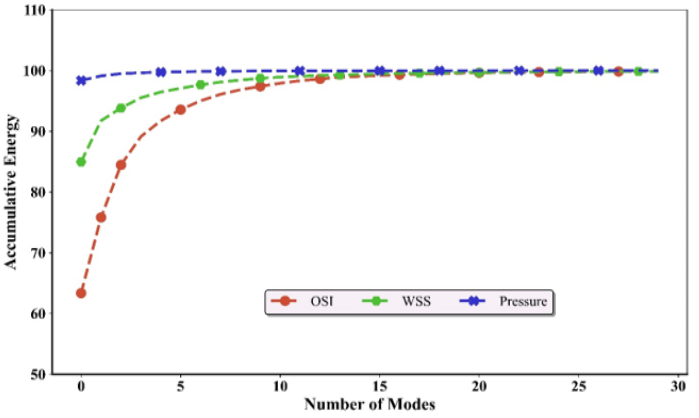


Figure 4. Cumulative energy as a function of retained modes for Case A

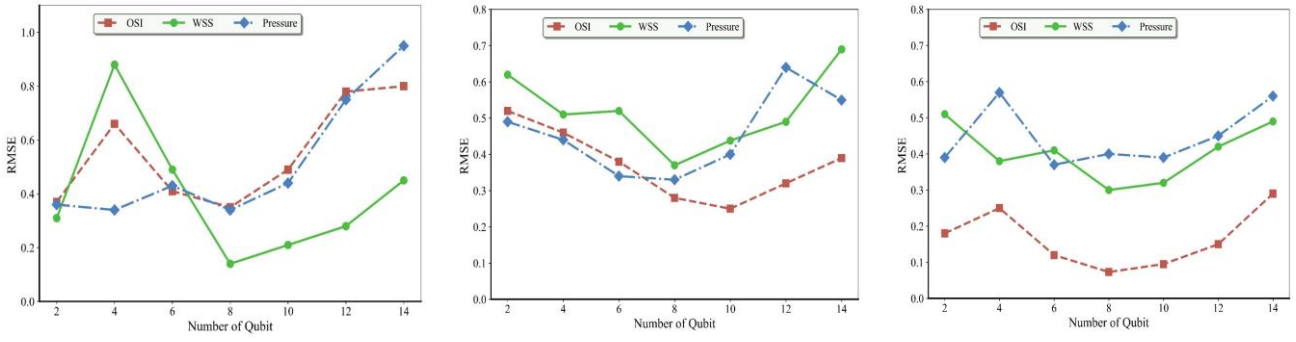


Figure 5. Reconstruction error for the reduced-order representation in Case A, shown for OSI, WSS, and pressure.

5.2. Effect of qubit count and circuit depth

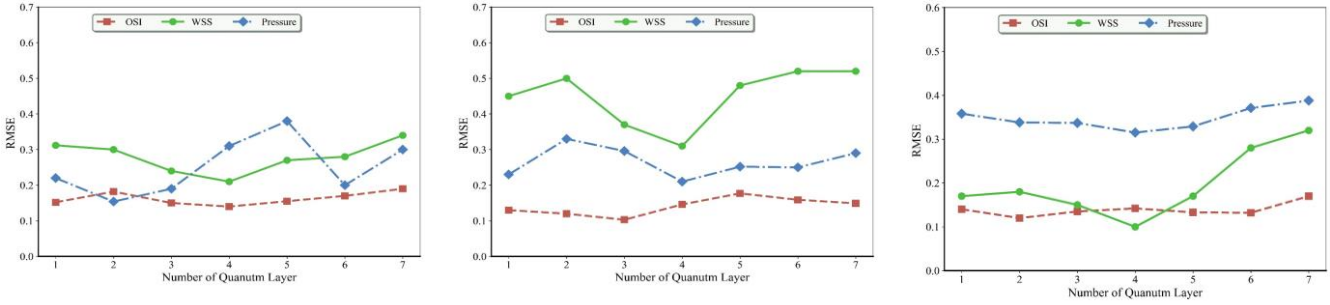
Figure 6 illustrates the effect of the number of qubits on the RMSE for WSS, OSI, and pressure under different variational circuit depths: one quantum layer in Figure 6(a), two layers in Figure 6(b), and four layers in Figure 6(c). The results indicate a non-monotonic relationship between the number of qubits and prediction error, suggesting the existence of an optimal range of circuit sizes for the investigated cases. The lowest RMSE values are generally obtained with approximately 8–10 qubits, whereas increasing the number of qubits beyond this range tends to increase the prediction error. This behavior suggests that increasing the dimensionality of the quantum feature representation does not necessarily improve predictive performance. Although additional qubits can provide a richer feature space, they may also increase the complexity of the trainable quantum-classical model without providing sufficiently useful information for the prediction task. Consequently, the observed results highlight the importance of selecting an appropriate circuit size rather than maximizing the number of qubits.

Figure 7 investigates the influence of the number of variational quantum layers on RMSE for different qubit configurations: two qubits in Figure 7(a), four qubits in Figure 7(b), and eight qubits in Figure 7(c). The results suggest that circuit depth also exhibits an optimal range, with the most favorable performance generally observed for approximately two to four quantum layers in the investigated configurations. Beyond the effective depth, adding more variational layers tends to increase the prediction error rather than provide further improvements. One possible explanation is that deeper circuits introduce additional trainable parameters and increase the complexity of the optimization problem. Given the limited size of the available training dataset, this additional complexity may not translate into improved generalization and may make the optimization process more difficult. Overall, Figures 6 and 7 demonstrate that both qubit count and circuit depth should be treated as important hyperparameters in the proposed VQC-PI-Transformer. Their joint selection is necessary to balance the expressivity of the quantum feature transformation against model complexity and predictive accuracy.



a) **b)** **c)**

Figure 6. Effect of qubit count on RMSE. a) one quantum layer, b)two quantum layers, c) four quantum Layers



a) b) c)

Figure 7. Effect of VQC depth on RMSE. a) two qubits, b) four qubits, c) eight qubits

5.3. Prediction error across the four aneurysm cases

Figures 8A–8D illustrate the temporal evolution of the prediction error for three hemodynamic parameters (OSI, WSS, and pressure) across the four patient-specific aneurysm cases. The horizontal axis represents normalized time over the prediction interval, while the vertical axis indicates the percentage prediction error. This representation allows the temporal behavior of the prediction error to be examined independently for each hemodynamic parameter and aneurysm geometry. The principal hyperparameters and model architecture are summarized in Table 1.

The results for Case A demonstrate different prediction-error behaviors among the three hemodynamic parameters. OSI exhibits a relatively stable error throughout the prediction interval, remaining approximately within the range of 3–4%. A slight reduction is observed during the initial portion of the interval, followed by relatively small fluctuations toward the end. This behavior indicates that the model maintains comparatively consistent predictive performance for OSI throughout the evaluated period. In contrast, WSS shows a progressive increase in prediction error. Its error begins at approximately 1% and increases considerably during the prediction interval, reaching values above 20% toward the final time steps. Pressure also exhibits an increasing error, although its growth is less pronounced than that of WSS, reaching approximately 12% near the end of the interval. The increasing errors in WSS and pressure suggest that predicting their temporal evolution becomes more challenging as the prediction interval progresses. In particular, WSS represents the most challenging variable for this case, whereas OSI maintains the lowest and most stable error.

Considering the four aneurysm cases together, several consistent patterns can be identified. First, OSI generally exhibits the lowest prediction error and the least variation over normalized time. This observation suggests that the proposed model maintains comparatively stable predictive performance for OSI across the investigated geometries. Second, WSS and pressure demonstrate a more pronounced dependence on normalized prediction time, with their errors generally increasing toward the end of the prediction interval. This behavior suggests that the prediction of these quantities is more sensitive to the temporal evolution of the hemodynamic fields. However, the relative behavior of WSS and pressure is case-dependent. WSS represents the largest source of prediction error in Cases A, C, and particularly D, whereas pressure becomes comparable to or exceeds WSS in the later portion of Case B. This variation demonstrates that the prediction difficulty is influenced not only by the type of hemodynamic parameter but also by the specific aneurysm geometry and its associated flow characteristics.

Another important observation is the difference in prediction-error magnitude among the four cases. Case D exhibits particularly high errors in WSS, while Case C shows relatively similar error trends for WSS and pressure. These differences indicate that the predictive performance of the VQC-PI-Transformer is geometry-dependent and that certain aneurysm configurations may require additional attention during model development and validation. The progressive increase in error observed for WSS and pressure may be associated with the temporal complexity of the underlying hemodynamic behavior and the accumulation of prediction deviations over the evaluated interval.

Overall, Figures 8A–8D demonstrate that the VQC-PI-Transformer provides parameter-dependent prediction performance across the four patient-specific aneurysm cases. While OSI exhibits relatively stable prediction behavior, WSS and pressure remain more sensitive to temporal evolution and aneurysm-specific flow characteristics.

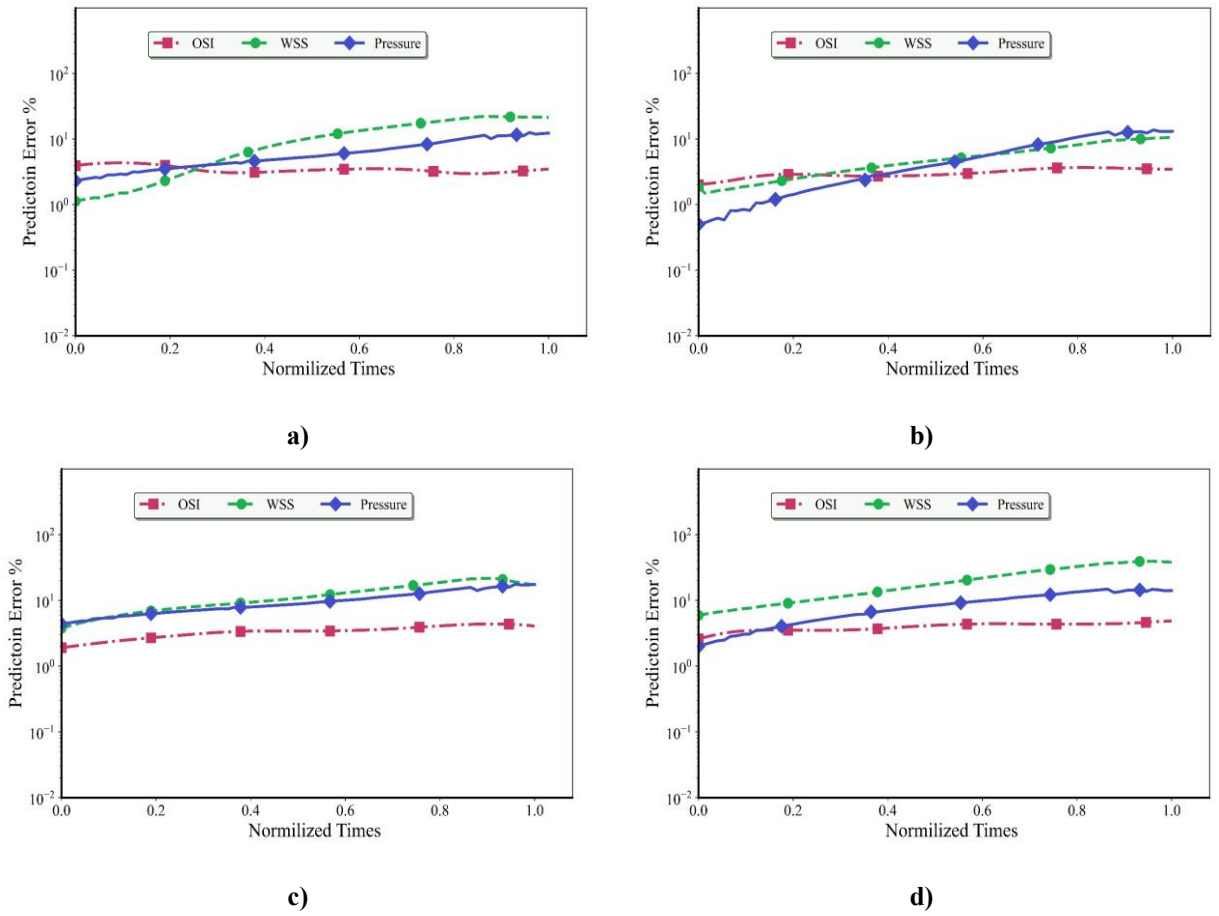


Figure 8. Prediction error over the held-out temporal interval in all cases. a)Case A, b)Case B, c)Case C, d)Case D

5.4. Comparison of VQC-PI-Transformer with the physics-informed Transformer

Figures 9A–9C compare the hemodynamic quantities obtained from the CFD simulations with those predicted by the classical physics-informed Transformer and the proposed VQC-PI-Transformer across the four patient-specific aneurysm cases. The comparisons include OSI, pressure, and time-averaged wall shear stress (TAWSS). The CFD results are considered the reference values, while the differences between the two prediction models indicate the potential contribution of the trainable variational quantum circuit to the classical Transformer architecture.

Figure 9A presents the comparison of OSI values for the four aneurysm cases. For Case 1, the CFD reference is approximately 0.32, whereas the classical Transformer and VQC-PI-Transformer predict values of approximately 0.35 and 0.34, respectively. Both models slightly overestimate the reference value; however, the VQC-enhanced model provides a closer prediction. Similar behavior is observed for Case 3, where the reference OSI is approximately 0.42. The classical Transformer predicts approximately 0.45, while the VQC-PI-Transformer produces a value close to 0.44. In this case, the quantum-enhanced model improves the agreement with CFD. For Case 3, the VQC-PI-Transformer provides a prediction close to the CFD reference (approximately 0.39), whereas the classical Transformer slightly underestimates the value. This represents a particularly favorable result for the quantum-enhanced architecture. Finally, for Case 4, both models underestimate the CFD reference of approximately 0.45. The VQC-PI-Transformer predicts approximately 0.40, compared with approximately 0.42 from the classical Transformer, indicating a larger deviation from the reference. Overall, the results demonstrate that the VQC integration improves the OSI prediction in Cases 1, 2, and 3.

Figure 9B illustrates the predicted pressure values in comparison with the CFD results. The pressure levels differ considerably among the four cases, with the highest values observed in Case 2 and the lowest in Case 3. Despite these differences in magnitude, both prediction models reproduce the general variation in pressure among the aneurysm geometries. In Case 1, the CFD pressure is approximately 1120 Pa. The classical Transformer slightly overestimates

this value, whereas the VQC-PI-Transformer produces a prediction closer to the reference. In Case 2, the CFD pressure is approximately 1325 Pa, and both models slightly overestimate the reference, with the VQC-enhanced prediction showing a marginally larger deviation. For Case 3, the CFD reference is approximately 550 Pa. Both models underestimate the pressure, although the VQC-PI-Transformer provides a slightly closer prediction. A similar improvement is observed in Case 4, where the CFD pressure is approximately 1025 Pa. Both models underestimate the reference, but the quantum-enhanced model produces a value closer to the CFD result. Overall, the pressure comparison indicates that the VQC-PI-Transformer improves the agreement with CFD in three of the four investigated cases. Although the differences between the two prediction models are relatively small, their consistent direction in several cases suggests that the quantum feature transformation can provide a beneficial contribution to the temporal representation learned by the Transformer for pressure prediction.

Figure 9C compares the predicted TAWSS values with the corresponding CFD reference results. In Case 1, the CFD value is approximately 3.3 Pa, while the classical Transformer predicts approximately 3.7 Pa. The VQC-PI-Transformer reduces this overestimation slightly, producing a value of approximately 3.6 Pa. In Case 2, the reference value is approximately 5.4 Pa. Both prediction models overestimate TAWSS, and the VQC-enhanced model produces a slightly higher value than the classical Transformer. Consequently, no improvement is observed for this case. For Case 3, the CFD reference is approximately 4.4 Pa, whereas the classical and quantum-enhanced models predict approximately 4.8 Pa and 4.7 Pa, respectively. The VQC-PI-Transformer therefore provides a closer estimate of the reference value. A similar behavior is observed in Case 4. The CFD TAWSS is approximately 7.2 Pa, while the classical Transformer predicts approximately 6.7 Pa. The VQC-enhanced model produces a slightly higher value of approximately 6.85 Pa, reducing the underestimation relative to CFD. These results indicate that the integration of the VQC improves TAWSS prediction in Cases 1, 3, and 4, although the improvement is not observed in Case 2. The results also demonstrate that the proposed architecture is capable of reproducing the differences in TAWSS magnitude among the four aneurysm geometries, while the remaining discrepancies are dependent on the specific case.

Considering Figures 9A–9C together, the results indicate that the proposed VQC-PI-Transformer generally provides closer agreement with CFD than the classical physics-informed Transformer in several of the investigated comparisons. An improvement is observed in approximately nine of the twelve case-variable combinations. However, the magnitude of improvement varies among the parameters and aneurysm geometries. The improved predictions observed in several cases suggest that the trainable quantum feature transformation may provide additional representational capability to the classical Transformer, allowing it to capture certain relationships in the reduced-order input features more effectively.

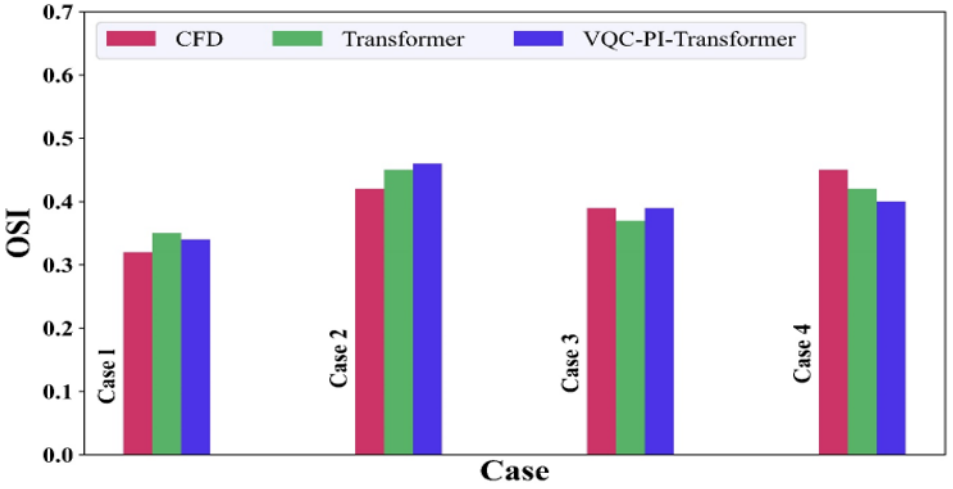


Figure 9A. Maximum OSI for the CFD, physics-informed Transformer, and VQC-PI-Transformer across Cases A–D

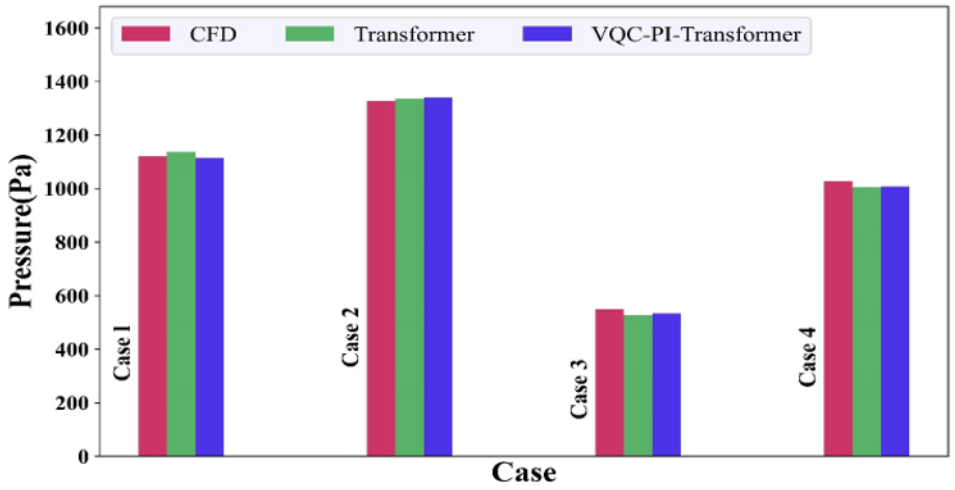


Figure 9B. Maximum pressure for the CFD, physics-informed Transformer, and VQC-PI-Transformer across Cases A–D

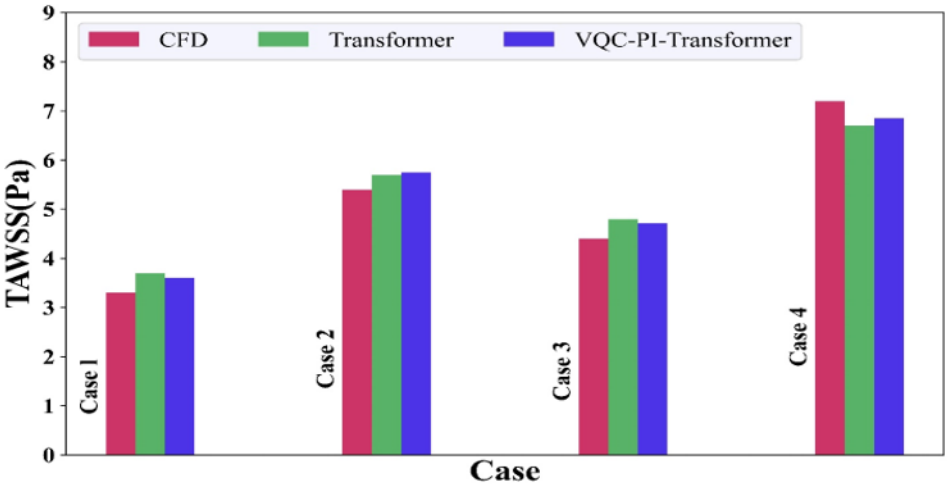


Figure 9C. TAWSS for the CFD, physics-informed Transformer, and VQC-PI-Transformer across Cases A–D

Figure 10 presents representative spatial distributions of the OSI and WSS for Case A. To further evaluate the predictive performance of the proposed VQC-PI-Transformer, the spatial contours generated by the proposed model are compared with those obtained from the classical Physics-Informed Transformer and the reference CFD simulations. The comparison is performed at a selected time instant within the prediction interval, allowing the spatial patterns predicted by the different approaches to be assessed under the same temporal conditions.

For both OSI and WSS, the VQC-PI-Transformer predictions exhibit a high degree of spatial agreement with the corresponding CFD reference contours. In particular, the predicted distributions reproduce the main spatial patterns and the overall variations across the aneurysm geometry. The contours also show considerable similarity to those generated by the classical Physics-Informed Transformer, indicating that the integration of the variational quantum circuit maintains the model's ability to capture the underlying spatial characteristics of the hemodynamic fields. These visual comparisons provide additional qualitative evidence of the proposed model's predictive capability.

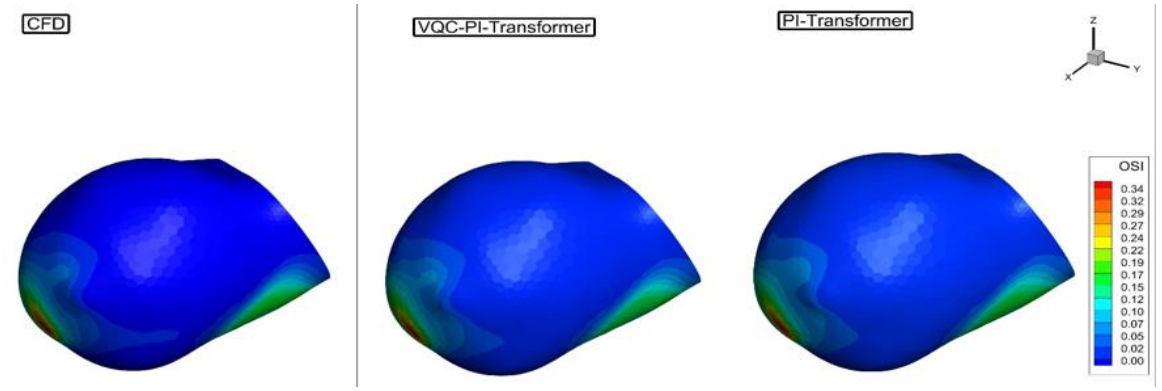


Figure 10A. OSI contour comparison for Case A

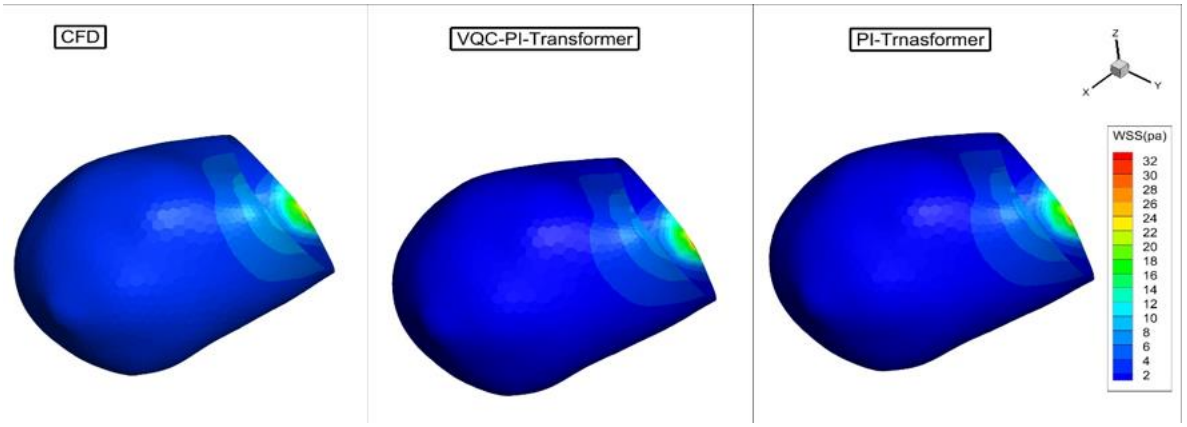


Figure 10B. WSS contour comparison for Case A

Table 1 summarizes the hyperparameters and principal configuration settings employed in the proposed VQC-PI-Transformer framework. The table includes the settings of both the classical Transformer architecture and the variational quantum circuit, providing a complete overview of the model configuration used in the experiments. The Transformer-related parameters include the embedding dimension, number of attention heads, number of encoder layers, feed-forward network dimension, and dropout rate. The quantum circuit configuration specifies the number of qubits and quantum layers, the parameterized rotation gates, the data-encoding strategy, the entanglement structure, and the measurement operators. In addition, the table reports the relevant training settings used to optimize the hybrid architecture.

Table 1. Hyperparameters of the VQC-PI-Transformer architecture.

Parameter	**Value**
Optimizer	AdamW
Initial learning rate	0.0002
Weight decay	0.00001
Maximum epochs	120
Batch size	32
Early-stopping patience	20

Gradient clipping	1.0
Scheduler	ReduceLROnPlateau
Dropout	0.10
Model dimension	256
Attention heads	8
Encoder layers	4
Feed-forward dimension	1,024
Quantum simulator	PennyLane default.qubit
QNode differentiation	PyTorch interface, backprop
Qubits	8
Variational quantum layers	2
Rotation gates	RX, RY, RZ
Encoding	Y-axis angle embedding
Entanglement	CNOT gates
Quantum feature dimension	24
Trainable rotation parameters	48

6. Discussion

The present study introduces a hybrid quantum-classical framework, termed the VQC-PI-Transformer, for the prediction of cerebral aneurysm hemodynamics. The proposed approach combines SVD-based dimensionality reduction, a trainable variational quantum circuit, and a physics-informed Transformer to predict time-dependent hemodynamic quantities from high-dimensional CFD data. The results obtained across four patient-specific middle cerebral artery aneurysm geometries demonstrate the potential of integrating quantum-derived features into classical physics-informed temporal prediction models. In particular, the comparison with CFD reference results indicates that the proposed hybrid architecture can improve prediction accuracy for several hemodynamic quantities compared with the classical physics-informed Transformer.

6.1. Effectiveness of the reduced-order prediction framework

One of the main advantages of the proposed framework is the combination of dimensionality reduction with temporal machine learning. Patient-specific CFD simulations produce large spatial datasets containing detailed information about blood-flow behavior. Directly processing these high-dimensional fields using a neural network can increase the computational and optimization requirements. The application of SVD provides a compact representation based on the dominant spatial modes, allowing the subsequent learning model to focus on the temporal evolution of a substantially reduced set of variables.

The cumulative-energy and reconstruction-error results presented in Figures 4 and 5 demonstrate the importance of selecting an appropriate number of retained modes. This stage provides a balance between reducing the original data dimensionality and preserving the spatial characteristics of the hemodynamic fields. The retained modal coefficients subsequently provide the input representation for the quantum-classical prediction architecture. Consequently, the proposed approach separates the spatial compression problem from the temporal prediction problem, making the overall learning task more manageable.

6.2. Influence of quantum circuit configuration

An important finding of this study concerns the relationship between quantum circuit configuration and prediction accuracy. The results presented in Figures 6 and 7 indicate that both the number of qubits and the number of variational quantum layers have a considerable influence on the prediction performance of the VQC-PI-Transformer.

The experiments reveal that increasing the number of qubits does not continuously reduce the prediction error. Instead, the investigated configurations exhibit an approximate optimal range, with the most favorable results generally obtained using approximately 8–10 qubits. A similar behavior is observed when examining the circuit depth, where configurations containing approximately two to four variational layers provide favorable prediction

performance. These observations indicate that the predictive capability of the hybrid architecture depends on an appropriate balance between quantum feature expressivity and model complexity.

The improvement observed within the selected qubit and layer configurations may be associated with the capacity of the trainable circuit to generate useful nonlinear feature representations from the reduced-order inputs. Increasing the number of qubits or variational layers introduces additional trainable parameters and expands the possible transformations of the encoded quantum states. However, beyond an appropriate level of complexity, these additional degrees of freedom may not provide useful information for the prediction task.

One possible explanation for the increase in prediction error observed in deeper or larger circuits is the limited size of the available training dataset. Increasing the circuit complexity may make the optimization problem more difficult and reduce the ability of the model to generalize to the unseen temporal interval. Nevertheless, this explanation should be interpreted as a possible mechanism rather than a directly established cause of the observed error trends.

These findings highlight the importance of systematic quantum-circuit design in hybrid scientific machine learning. Rather than simply increasing the number of qubits or variational layers, selecting an appropriate circuit configuration is necessary to obtain a useful quantum feature representation. The observed optimal ranges also suggest that relatively compact variational circuits can provide an effective integration strategy for reduced-order hemodynamic prediction.

6.3. Temporal prediction performance

The prediction-error distributions presented in Figures 8A–8D demonstrate that the proposed model maintains different levels of accuracy depending on the hemodynamic variable and aneurysm geometry. Despite these differences, several consistent patterns can be identified across the four investigated cases. Among the evaluated parameters, OSI generally exhibits comparatively stable prediction errors throughout the normalized prediction interval. This behavior is observed across all four aneurysm geometries, suggesting that the model can reproduce the temporal behavior associated with OSI with relatively consistent accuracy. This result is encouraging because OSI is a time-dependent wall-related parameter whose calculation depends on the variation of the wall shear stress vector throughout the cardiac cycle.

In contrast, WSS and pressure exhibit more pronounced variations in prediction error, particularly toward the later portion of the prediction interval. The magnitude of this behavior depends on the aneurysm geometry. For example, Case D shows a particularly pronounced increase in WSS prediction error, whereas Case B demonstrates a more comparable error distribution between WSS and pressure. These differences emphasize the importance of evaluating the proposed model across multiple patient-specific geometries rather than relying on a single aneurysm configuration.

The observed variation among cases may be associated with differences in the underlying spatial and temporal characteristics of the hemodynamic fields. Localized shear patterns, pressure gradients, and complex flow structures can influence the difficulty of reconstructing and predicting different physical quantities. Although the current results do not establish a direct causal relationship between geometric complexity and prediction error, they demonstrate that model performance is influenced by the specific aneurysm configuration.

An additional positive observation is that the model retains relatively stable OSI prediction behavior even when WSS and pressure become more challenging to predict. This suggests that the proposed reduced-order framework can represent different hemodynamic quantities with different levels of predictive accuracy, highlighting the value of parameter-specific evaluation.

6.4. Contribution of the variational quantum circuit

The comparison between the classical physics-informed Transformer and the proposed VQC-PI-Transformer, presented in Figures 9A–9C, provides an important assessment of the contribution of the trainable quantum circuit. The bar-chart comparisons demonstrate that the VQC-PI-Transformer produces values closer to the CFD reference in most of the evaluated case-parameter combinations. Improvements are particularly encouraging for pressure and TAWSS, for which the quantum-enhanced model provides closer estimates in three of the four aneurysm cases. For OSI, improvements are observed in three cases.

The results are particularly interesting because the quantum circuit is integrated as a feature-transformation component within an otherwise classical physics-informed Transformer. The VQC does not replace the temporal prediction architecture. Instead, it provides an additional trainable feature representation that is combined with the classical input features before temporal processing. The observed improvements therefore suggest that this quantum-derived representation can provide useful information for the subsequent Transformer prediction in several of the investigated cases.

The performance differences also indicate that the contribution of the VQC is dependent on the target hemodynamic quantity and aneurysm geometry. This is expected to be important in patient-specific modeling, where different geometries can produce different flow structures and spatial distributions.

From a methodological perspective, these results support the investigation of trainable variational quantum circuits as feature-learning components in physics-informed scientific machine learning. Unlike fixed quantum reservoir computing, the circuit parameters in the proposed architecture are optimized jointly with the classical Transformer parameters. This provides an opportunity for the quantum feature transformation to adapt to the hemodynamic prediction objective during training.

It is important, however, to distinguish improved predictive performance from quantum computational advantage. The present implementation uses PennyLane's ideal state-vector simulator with analytic expectation values. Consequently, the observed improvements demonstrate the potential predictive contribution of the VQC-based representation under the investigated simulation and training conditions, rather than a computational speedup over classical methods.

6.5. Physical relevance and potential applications

The ability to predict spatially distributed hemodynamic quantities is an important consideration when developing surrogate models for cerebral aneurysm analysis. Maximum or average values alone cannot describe the complete spatial distribution of wall-related parameters. For this reason, the comparison of reconstructed OSI and WSS contours presented in Figures 10A and 10B provides a complementary assessment of the model beyond the scalar comparisons shown in Figure 10.

The combination of reduced-order reconstruction and physics-informed temporal learning provides a framework for estimating these quantities without repeatedly performing a complete high-fidelity CFD simulation for every prediction. Once trained and validated, such a surrogate could potentially be useful for rapidly exploring alternative hemodynamic conditions, examining different aneurysm geometries, and conducting repeated computational analyses.

Although the present work focuses on four patient-specific aneurysm geometries, the proposed architecture provides a foundation for further investigations involving additional vascular configurations, different inlet flow conditions, and geometrical variations such as aneurysm blebs or vessel dilation. These extensions could help establish whether the hybrid model can generalize to a broader range of patient-specific hemodynamic scenarios.

6.6. Limitations and future directions

Despite the encouraging results, several aspects require further investigation. First, the present study considers four patient-specific aneurysm geometries. Although the results demonstrate the behavior of the proposed model across different cases, additional geometries will be required to assess its generalization to a wider population of aneurysms. Second, the quantum circuit is currently evaluated using an ideal simulator. Future work could investigate the influence of finite-shot measurements and realistic quantum noise, as well as the feasibility of implementing the circuit on quantum hardware. Such investigations would provide a more complete understanding of the practical behavior of the proposed architecture.

Finally, further research should examine alternative quantum encodings, entanglement structures, circuit depths, and feature-fusion strategies. Increasing the number of training cases would also help establish the consistency of the observed improvements. Overall, the findings of this study demonstrate the potential of combining trainable variational quantum circuits and physics-informed Transformers for patient-specific hemodynamic prediction. The improvements observed in several of the CFD comparisons provide motivation for further development of hybrid quantum-classical architectures in computational biomechanics.

7. Conclusion

This study presents a hybrid quantum-classical framework, termed the VQC-PI-Transformer, for prediction of cerebral aneurysm hemodynamics. The proposed methodology combines SVD-based dimensionality reduction with a trainable variational quantum circuit and a physics-informed Transformer. The reduced-order representation enables the model to process compact temporal features instead of the complete high-dimensional CFD fields, while the VQC provides an additional trainable feature transformation for the classical Transformer.

The proposed framework was evaluated using four unruptured, patient-specific terminal-type middle cerebral artery aneurysm geometries, with OSI, WSS, TAWSS, and pressure considered as the principal hemodynamic quantities. The investigation of quantum circuit configurations demonstrated that prediction performance depends on

both qubit count and circuit depth. The most favorable performance was generally obtained within the investigated ranges of approximately 8–10 qubits and two to four variational layers, highlighting the importance of selecting an appropriate balance between quantum circuit expressivity and model complexity.

The temporal prediction results demonstrated comparatively stable error behavior for OSI across the four aneurysm cases, while WSS and pressure exhibited stronger variations, particularly toward the later portion of the prediction interval. These findings emphasize the importance of considering both the physical quantity and aneurysm geometry when evaluating the accuracy of reduced-order hemodynamic prediction.

The comparison with CFD reference results provided encouraging evidence regarding the contribution of the variational quantum circuit. The VQC-PI-Transformer model produced closer estimates than the classical physics-informed Transformer in most of the evaluated comparisons. Although the improvement was not uniform across every case and parameter, the results indicate that trainable quantum-derived features can provide a beneficial contribution to classical temporal modeling for several hemodynamic prediction tasks.

In conclusion, the proposed VQC-PI-Transformer offers a promising direction for integrating variational quantum circuits into physics-informed modeling of cerebral aneurysm hemodynamics. The findings support further investigation of hybrid quantum-classical architectures for computational biomechanics, particularly in applications where high-dimensional CFD data, temporal dependencies, and physical constraints must be considered simultaneously. Future work will focus on expanding the number of patient-specific cases, improving the evaluation of physical consistency, and investigating more advanced quantum circuit configurations and implementations.

References

[1] M. B. Gerdroodbary, L.-C. Zhao, Investigation of pulsatile blood flow dynamics in anterior cerebral artery aneurysms under different physiological conditions: hemodynamic analyses, *Journal of Applied and Computational Mechanics*, 2025.

[2] I. Shiryanpoor, A. Kheiri, M. B. Gerdroodbary, P. Valipour, R. Moradi, Using computational fluid dynamic for evaluation of rupture risk of micro cerebral aneurysms in the growth process: hemodynamic analysis, *International Journal of Modern Physics C,* Vol. 36, pp. 2450184, 2025.

[3] A. Arzani, S. C. Shadden, Characterizations and correlations of wall shear stress in aneurysmal flow, *Journal of Biomechanical Engineering,* Vol. 138, No. 1, pp. 014503, 2016.

[4] S. V. Mousavi, M. B. Gerdroodbary, A. Sabernaeemi, S. Salavatidezfouli, P. Valipour, Impacts of the aneurysm deformation induced by stent on hemodynamic of blood flow in saccular internal carotid artery aneurysms, *AIP Advances,* Vol. 14, pp. 095035, 2024.

[5] S. V. Mousavi, M. Barzegar Gerdroodbary, I. Shiryanpoor, P. Valipour, Hemodynamic Assessment of Bleb Formation in Patient-Specific Intracranial Aneurysms using Computational Fluid Dynamics, *International Journal of Modern Physics C*, 2027.

[6] G. Berkooz, P. Holmes, J. L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annual Review of Fluid Mechanics,* Vol. 25, pp. 539–575, 1993.

[7] M. B. Gerdroodbary, S. Salavatidezfouli, A predictive surrogate model based on linear and nonlinear solution manifold reduction in cardiovascular FSI: a comparative study, *Computers in Biology and Medicine,* Vol. 189, pp. 109959, 2025.

[8] M. B. Gerdroodbary, S. Salavatidezfouli, A predictive surrogate model of blood haemodynamics for patient-specific carotid artery stenosis, *Journal of the Royal Society Interface,* Vol. 22, No. 224, pp. 20240774, 2025.

[9] W. Aich, J. Djuansjah, A. B. M. Ali, A. a. Alizadeh, M. A. Hussein, N. S. S. Singh, K. Hajlaoui, W. Rajhi, Pulsatile hemodynamic prediction in cerebral fusiform aneurysms using proper orthogonal decomposition model and long short-term memory networks, *International Communications in Heat and Mass Transfer,* Vol. 169, pp. 109833, 2025.

[10] A. Kheiri, S. Hariri, M. Barzegar Gerdroodbary, M. Sorkhi, P. Valipour, A CFD–machine learning framework for efficient prediction of cerebral aneurysm hemodynamic under vessel dilation, *Chinese Journal of Physics,* Vol. 102, pp. 748–767, 2026/08/01/, 2026.

[11] N. G. Hammouda, K. Kriaa, A. a. Alizadeh, H. Rajab, A. M. Sadeq, N. S. S. Singh, K. Hajlaoui, W. Aich, Physically informed machine learning for predicting non-Newtonian blood flow in giant MCA aneurysms after endovascular treatment, *International Communications in Heat and Mass Transfer,* Vol. 172, pp. 110296, 2026.

[12] N. G. Hammouda, J. Djuansjah, A. B. M. Ali, A. a. Alizadeh, M. A. Hussein, N. S. S. Singh, K. Hajlaoui, W. Aich, Patient-specific deep learning reduced-order model of blood hemodynamic indices in thoracic aortic aneurysms via dynamic mode decomposition, *Chinese Journal of Physics,* Vol. 99, 2026.

[13] M. Barzegar Gerdroodbary, S. Z. Sajadi, Reduced-Order Learning of Cerebral Aneurysm Hemodynamics: Comparing POD, PCA, and Autoencoder Models Coupled with TCN, *Journal of Computational Applied Mechanics,* Vol. 57, No. 4, pp. 731–750, 2026.

[14] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, P. J. Coles, Variational quantum algorithms, *Nature Reviews Physics,* Vol. 3, pp. 625–644, 2021.

[15] C. Trahan, M. Loveland, S. Dent, Quantum physics-informed neural networks, *Entropy,* Vol. 26, No. 8, pp. 649, 2024.

[16] S. Berger, N. Hosters, M. Möller, Trainable embedding quantum physics informed neural networks for solving nonlinear PDEs, *Scientific Reports,* Vol. 15, pp. 18823, 2025.

[17] R. Bozorgpour, P. Kim, M. Soltany Sadrabadi, Hemodynamic variables in cerebral aneurysms: Integrating CFD-based computational modeling and clinical insight, *Brain Hemorrhages*, 2026.

[18] M. Barzegar Gerdroodbary, X. Wang, A computational analysis of sac volume effects on hemodynamics and rupture risk in cerebral saccular aneurysms, *Journal of Applied and Computational Mechanics,* Vol. 12, No. 3, pp. 1084–1092, 2026.

[19] AneuriskWeb Project. *AneuriskWeb project website*, Accessed; http://ecm2.mathcs.emory.edu/aneuriskweb.

[20] M. Piccinelli, A. Veneziani, D. A. Steinman, A. Remuzzi, L. Antiga, A framework for geometric analysis of vascular structures: Application to cerebral aneurysms, *IEEE Transactions on Medical Imaging,* Vol. 28, No. 8, pp. 1141–1155, 2009.

[21] I. ANSYS, 2020, *ANSYS Fluent User's Guide, Release 2020 R2*, ANSYS, Inc., Canonsburg, PA

[22] M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics,* Vol. 378, pp. 686–707, 2019.

[23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in *Advances in Neural Information Processing Systems 30*, 2017, pp. 5998–6008.

[24] F. Y. Leong, W.-B. Ewe, S. B. Q. Tran, Z. Zhang, J. Y. Khoo, Hybrid quantum physics-informed neural network: Towards efficient learning of high-speed flows, *Computers & Fluids,* Vol. 301, pp. 106782, 2025.

[25] S. Ranilla-Cortina, D. A. Aranda, J. Ballesteros, J. Bonilla, N. Monrio, et al., E. F. Combarro, Quantum-enhanced architectures for multivariate time-series forecasting, *The Journal of Supercomputing,* Vol. 82, pp. 155, 2026.

[26] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, et al., PennyLane: Automatic differentiation of hybrid quantum-classical computations, arXiv, 2018.

[27] J. Ding, P. Valipour, Impacts of coiling technique on hemodynamic of the MCA aneurysms: Computational study, *International Journal of Modern Physics C,* Vol. 35, No. 5, pp. 1–14, 2024.

[28] R. Bozorgpour, Hemodynamic markers: CFD-based prediction of cerebral aneurysm rupture risk, *Vascular Pharmacology,* Vol. 162, pp. 107578, 2026.

[29] T. Hussain, O. A. Bég, S. Kuharat, T. A. Bég, M. M. Bhatti, Numerical simulation and visualization of Newtonian and non-Newtonian hemodynamics in a stenotic blood vessel, *Journal of Computational Applied Mechanics,* Vol. 57, No. 3, pp. 443–472, 2026.