Bol. Soc. Paran. Mat.
©SPM – E-ISSN-2175-1188
SPM: www.spm.uem.br/bspm

(3s.) **v. 2026 (44) 7** : 1–17.
ISSN-0037-8712
doi:10.5269/bspm.81302

Hybrid Quantum–Classical Variational Reconstruction of CFD Fields Using VQAs

Surbhi Sharma, Rashmi Bhardwaj and Debabrata Datta

ABSTRACT: High-fidelity reconstruction of computational fluid dynamics (CFD) fields is a central challenge in reduced-order modelling, flow control, and large-scale scientific simulations, where classical approaches face severe scalability and nonlinearity limitations. In this work, a hybrid quantum–classical framework is proposed for reconstructing CFD solution snapshots using variational quantum algorithms (VQAs) on near-term quantum hardware. The method formulates CFD field reconstruction as a variational quantum state learning problem using amplitude encoding and shallow parameterized quantum circuits. Reconstruction fidelities ranging from 0.8696 to 0.9369 and low L_2 errors between 0.2532 and 0.3673 are achieved, demonstrating accurate quantum approximation of high-dimensional fluid fields. The results significantly outperform earlier low-fidelity attempts and confirm that VQA-based surrogate modelling can capture dominant spatial flow structures even with limited circuit depth. This study highlights the potential of hybrid variational quantum frameworks as scalable alternatives to classical CFD post-processing and reduced-order modelling.

Key Words: CFD, variational quantum algorithms, reduced-order modelling, amplitude encoding, reconstruction fidelity.

Contents

2020 *Mathematics Subject Classification:* 76M22, 81P68, 68T07.
Submitted January 12, 2026. Published April 11, 2026

1. Introduction

Computational Fluid Dynamics (CFD) has found its way to be an essential tool in the aerospace, mechanical engineering, defence engineering, automotive, and the environmental modelling sector. CFD allows one to explicitly predict aerodynamic forces, heat transfer, turbulence phenomena, flow separation in complicated geometries by numerically solving the governing partial differential equations, most typically the Navier-Stokes equations, incompressible or compressible. In the last 50 years, the discretization methods (finite volume, finite element, finite difference), turbulence modelling (RANS, LES, DNS) and high-performance computing have greatly enhanced the accuracy of CFD simulations. However, realistically simulated turbulent flows are still computationally infeasible, even with the current supercomputers.

The main challenge is associated with the multiscale essence of turbulent fluid movement. The computational cost of Direct Numerical Simulation (DNS) scales as $O(\mathrm{Re}^3)$, making it infeasible for high-Reynolds-number industrial applications. Reduced approaches of LES and even hybrid RANS-LES methods are still costly when used in unsteady aerodynamics, multi-physics processes or large design-space exploration. As the complexity of modern engineering systems is growing, the old CFD workflows are finding more and more demand to be accelerated, reduced-order modeled, and to have new ways to problematically approximate flow fields without losing accuracy.

The basic alternative paradigm of computing, which quantum computing provides, may suitably address some of these constraints. Quantum processors act on qubits, which are in superpositions of basis states, and allow exponentially large vectors to be represented with comparatively few physical resources [12]. The enticements of quantum algorithms, including HHL to solve linear systems, quantum phase estimation, and quantum linear solvers in linear differential equations, have generated interest in quantum systems scientific computing, especially of large-scale PDEs. Nevertheless, the vast majority of provably exponential quantum accelerations are based on idealized quantum machine models, deep circuits, or fault-tolerant systems that are unavailable on the current Noisy Intermediate-Scale Quantum (NISQ) devices. [10].

Hybrid quantum-classical methods have been found to be a viable route to quantum advantage in the near future [2,31]. VQAs in specific are especially appealing, since they represent a combination of expressive parameterized quantum circuits with classical optimization [1]. VQAs do not need deep coherent evolution, in contrast to algorithms, and can be run on noisy hardware. Critically, the VQAs are applicable not only to discrete optimization or quantum chemistry problems, however, it can also be applied to the approximation of continuous functions, learning solution manifolds, reducing dimensionality of high-dimensional data, and building reduced-order models [3,4].

Recent studies have found out that most high-dimensional physical systems such as fluid flows frequently lie on low-dimensional nonlinear manifolds. Proper Orthogonal Decomposition (POD), Dynamic Mode Decomposition (DMD), and other reduced-order approaches to CFD are based on such phenomena. Variational quantum circuit, which have activity enabling them to approximate complex quantum states with entanglement and parameterized unitaries, is an immediate candidate to quantum-inspired surrogate models to approximate these manifolds. Although with widespread success, current reduced-order modeling (ROM) methods within the computational fluid dynamics domain continue to be fundamentally restricted in their application to nonlinear flow phenomena of complex nature. The classical techniques like Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD) are based on the linear subspace projections based on snapshot sets. Although best in a mean-square sense when applied to linear systems, such techniques frequently do not reflect advection-dominated flows and transitional flows, as well as strongly nonlinear dynamics, where the solution space is nonlinear in nature. Consequently, many modes will often be needed to attain reasonable accuracy, thus reducing the computational benefits of model reduction.

Recent extensions of this framework that are data-driven, such as neural-network-based autoencoders and operator-learning systems, are trying to address these limitations by providing nonlinear mappings. Nonetheless, these methods generally need large amounts of training data, are computationally expensive, and sensitive to hyperparameters as well as they are not usually robust or physically interpretable. Furthermore, they are difficult to scale to high-dimensional CFD data (especially to unsteady or multi-parameter flows), as training costs increase quite rapidly with problem size.

These constraints identify a fundamental weakness of current ROM techniques: the unavailability of a small, nonlinear, and scalable representation that can represent complex flow structures with the necessary amount of computational overhead. The resolution of this gap is the reason to consider other computational paradigms. Using amplitude encoding and variational quantum circuits as the key components of quantum information processing, it is possible to provide an entirely different framework of representations in which the high-dimensional CFD fields can be represented in exponentially smaller Hilbert spaces. With variational quantum reduced-order models, the potential to achieve nonlinear flow manifolds on a logarithmic time scale through superposition and entanglement, variational quantum reduced-order models may be seen as a promising complement to classical ROMs to support next-generation CFD applications.

This work extends the concept of variational quantum state learning to Computational Fluid Dynamics (CFD) by formulating the reconstruction of flow-field snapshots as a quantum variational optimization problem. Given a high-dimensional CFD snapshot represented as a vector

$$\mathbf{u} \in \mathbb{R}^N,$$

the objective is to encode this field into a quantum state and learn a parameterized unitary transformation $U(\boldsymbol{\theta})$ such that the resulting quantum state

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle$$

approximates the target CFD state with maximum fidelity. In this framework, CFD field reconstruction is posed as the minimization of a well-defined cost function based on quantum state fidelity and classical L_2-norm error, providing a rigorous mathematical basis for accuracy evaluation.

Most prior quantum CFD studies have focused on explicit quantum solvers for discretized fluid equations, including finite-difference formulations, quantum lattice Boltzmann methods, or quantum linear differential equation (QLDE) solvers. While theoretically appealing, such approaches typically rely on deep quantum circuits, fault-tolerant gate operations, favourable sparsity conditions, or well-conditioned linear operators. These requirements render them impractical for current Noisy Intermediate-Scale Quantum (NISQ) hardware, where coherence time, gate depth, and error accumulation severely limit algorithmic feasibility.

In contrast, the approach proposed in this paper adopts a fully NISQ-compatible hybrid variational framework. The method does not assume sparsity of the governing operators, nor does it require fault-tolerant quantum computation. Instead, it leverages shallow parameterized quantum circuits optimized through classical feedback loops, making the framework hardware-efficient, scalable, and directly applicable to real CFD datasets. Importantly, the variational formulation enables the direct computation of physically meaningful performance metrics, namely quantum state fidelity and classical L_2 reconstruction error, allowing quantitative and transparent assessment of reconstruction quality.

Using this framework, six representative CFD flow-field snapshots were reconstructed with high accuracy. The trained variational quantum circuits achieved fidelities in the range 0.8696 to 0.9369 and corresponding L_2 errors between 0.2532 and 0.3673. These results demonstrate that variational quantum circuits are capable of learning dominant spatial structures of fluid flows, even when the original high-dimensional fields are compressed into low-dimensional quantum representations. The observed performance confirms that essential flow features can be preserved despite aggressive dimensionality reduction.

Contributions of this work are as follows:

- We formulate CFD field reconstruction as a variational quantum state learning problem using amplitude encoding.

- We propose a fully NISQ-compatible hybrid quantum–classical reduced-order modeling framework.
- We demonstrate high-fidelity reconstruction of multiple CFD snapshots using shallow variational circuits.
- We provide a quantitative comparison with classical ROM concepts and discuss physical interpretability.
- We explicitly analyze limitations arising from simulator-based evaluation and circuit depth.

2. Literature Survey

2.1. Classical reduced-order modelling for CFD

Compression of high dimensional CFD data into a low-dimensional manifold to enable rapid simulation, control and design optimisation is an extremely popular application of reduced-order models (ROMs) e.g. Proper Orthogonal Decomposition (POD), Dynamic Mode Decomposition (DMD) and more recent data-driven ROMs. POD uses snapshot ensembles to obtain orthogonal basis modes of a snapshot ensemble by singular value decomposition, whereas DMD uses snapshots to obtain spatio-temporal coherent structures and linear dynamics [21,22,23,24]. These classical ROMs can solve problems of high computational efficiency and can be intuitively studied; however, they are ill-posed in strongly nonlinear dynamics, flows dominated by advection, and extrapolation between parameter space; it has been proposed that they can be improved with nonlinear and data-driven extensions (neural autoencoders, operator learning, localized POD, etc.). Detailed descriptions and algorithmic specifications of current ROMs and their use in CFD can be found in the literature of ROMs [19,20,25].

2.2. Quantum representations and algorithms for high-dimensional data

Quantum information processors natively represent 2^n-dimensional vectors using n qubits via amplitude encoding, suggesting an attractive representational advantage for very high-dimensional fields such as CFD snapshots. Two lines of algorithmic development are particularly relevant:

- **Quantum Principal Component Analysis (qPCA):** Lloyd, Mohseni, and Rebentrost introduced a quantum analogue of PCA that extracts dominant eigenvectors of an unknown density matrix by coherently evolving copies of the state and applying quantum phase estimation; qPCA promises exponential speedups in idealized settings for low-rank density matrices and has been a foundation for quantum dimension-reduction ideas [5].
- **Quantum autoencoders / unitary compression:** Romero et al. proposed quantum autoencoders as unitary circuits trained by classical optimizers to compress and decompress families of quantum states; the architecture and training paradigm is directly analogous to classical autoencoders but operates natively on quantum amplitudes and has been shown capable of nontrivial compression of structured data sets [7]. Quantum autoencoders map cleanly to the ROM objective: learn a compact latent subspace (few qubits) that retains the dominant structure of the input ensemble [8].

Unlike quantum autoencoders that explicitly compress and decompress quantum states using latent qubits, the proposed framework focuses on direct variational reconstruction of classical CFD snapshots encoded as quantum states. The objective is not dimensionality reduction alone, but accurate surrogate reconstruction of physically meaningful flow fields, evaluated through classical error norms in addition to quantum fidelity.

2.3. Variational Quantum Algorithms (VQAs) and their applicability

VQAs are the leading practical framework for near-term quantum applications: parameterized, shallow circuits are optimized by a classical loop to minimize application-specific cost functions (energies, fidelities, classification loss, etc.) [6]. Research in VQA has come of age to a wide discipline that encompasses expressibility, barren-plateau research, optimization, and error tolerance and has had the recent reviews summarize the potential of VQAs and the open trainability/robustness problems [31,32]. VQAs are especially well adapted to tasks involving hybrid compression since they can only take short circuits to train them, and can be trained to achieve maximum fidelity or minimum reconstruction error directly on statevectors or noisy hardware [28].

2.4. Quantum & variational approaches applied to PDEs and CFD

Increased activity in the use of quantum algorithms to solve PDEs and fluid problems has developed. Some of the strategies that have been considered include:

- **HHL variants / quantum linear solvers:** Quantum linear-algebra subroutines (qPCA, QSVD) are primitives to discretized linear PDEs but are generally based on deep circuits and favorable condition numbers [12].

In addition, recent studies have demonstrated the feasibility of implementing reduced-order modeling techniques on near-term quantum computers, reformulating classical ROM concepts as hybrid quantum–classical optimization problems suitable for NISQ devices [34].

Quantum algorithms have been suggested to solve nonlinear PDEs (such as Burgers equation) both pure-quantum algorithms, which use many copies, embedding tricks to solve the nonlinear equation, and hybrid variational algorithms, which approximate the nonlinear equation with parameterized circuits [13,14]. A number of recent papers have shown quantum or quantum-inspired algorithms simulating Burgers or other transport equations with proof-of-principle agreement with classical solvers on small scales and highlighting the open scalability issues [16]. Recent investigations into physics-informed quantum circuits for transport equations further highlight both the potential and current limitations of hybrid quantum–classical approaches for PDE-governed engineering systems, particularly with respect to noise sensitivity, circuit depth, and scalability [35].

Variational transport and thermofluids model Challenges to transport and nonlinear thermofluids Variations on VQA ideas have more recently been developed to applied to transport and thermofluids problems, with encoder/ansatz constructions and discretization-to-circuit translations that scale with engineering trade-offs and under some conditions evaluate complexity scaling [17]. These contributions suggest that CFD numerical practices can be adjusted to VQA-style practices.

2.5. Quantum compression of classical data and quantum-inspired CFD

Outside of the strictly quantum algorithms, there are some hybrid and quantum-like algorithms that would seek to apply the benefits of quantum compression to classical high-dimensional problems. A well-known effort is to build quantum-inspired linear algebra algorithms or even logarithmic-scaling representations of special structured problems (recently shown to work on fluid models), perhaps even near-value even before universal quantum advantage [18,26]. These hybrid methods encourage QVAs based on compression and representation (quantum PCA / quantum autoencoders) as opposed to explicit and explicit solvers of the PDE.

2.6. Experiments, hardware demonstrations and constraints

The experimental demonstrations of qPCA-like protocols and small quantum autoencoders have already emerged on testbeds, demonstrating promise at small state size, and pointing out noise resilience measures and measurement cost [15,9]. Practical implementations face several constraints: state preparation (efficient amplitude encoding), the number of identical state copies required for some quantum algorithms, circuit depth vs. noise tradeoffs, and classical optimization difficulties (barren plateaus, local minima, and evaluation budgets) [29]. These limitations influence modern hybrid approaches to be interested in short, hardware-efficient ansatzes, classical pre-processing (e.g., POD) and cost functions constructed based on efficiently measured fidelities or observables.

2.7. Literature gaps and motivation of this work

Recent work has begun to explore explicitly quantum reduced-order modeling strategies for flow-field analysis, such as the PolyQROM framework, which employs orthogonal-polynomial-enhanced variational circuits to improve expressivity and trainability for fluid data compression and reconstruction [33]. Combined, the literature demonstrates a strong trend: quantum algorithms are well-suited to compress and encode high-dimensional physical fields, but previous quantum-CFD research has been done in either small proof-of-principle solvers (e.g. Burgers) or on algorithmic primitives (qPCA, HHL) [11]. The combination of (i) amplitude encoding of classical CFD snapshots, (ii) variational quantum autoencoder training and (iii) a realistic assessment of representative CFD snapshots with quantitative fidelity and L_2 error metrics is few. The gap that our current study fills: the paper presents a NISQ-compatible variational ROM that encodes full-field snapshots into a small register of latent qubits, and allows high fidelity reconstruction of multiple snapshots through the process, in a manner that both complements the literature on quantum algorithms and complements classical practice of building ROMs.

3. Proposed Methodology

The proposed methodology integrates Computational Fluid Dynamics (CFD) with variational quantum algorithms to construct a hybrid framework capable of learning reduced-order representations of high-dimensional fluid flow fields. The purpose of the mathematical formulation presented in this section is to establish a rigorous link between high-dimensional CFD fields and quantum state representations. Rather than solving the Navier–Stokes equations directly on a quantum processor, the objective is to learn a compact surrogate model capable of reconstructing flow-field snapshots with high fidelity. This approach is motivated by reduced-order modeling in classical CFD, where solution manifolds are approximated using a low-dimensional basis.

In the quantum setting, amplitude encoding provides a natural mechanism to embed CFD snapshots into a Hilbert space of dimension 2^n, enabling exponential compression. Variational quantum circuits are then employed to learn nonlinear mappings within this space, analogous to nonlinear autoencoders. The mathematical framework introduced below formalizes this idea and provides physically meaningful metrics i.e. quantum fidelity and classical reconstruction error to evaluate performance.

While simulations are performed using statevector backends, expected performance degradation on real hardware due to noise and sampling errors is discussed in the later part of the paper.

The overall workflow consists of the following five modules:

- High-fidelity CFD data generation
- Dimensionality reduction and classical preprocessing
- Quantum state encoding and variational ansatz design
- Variational optimization using Qiskit
- Fidelity and reconstruction-error analysis

Each component of the methodology is described in detail below.

3.1. Generation of CFD Snapshot Dataset

To construct a quantum-ready learning model, a representative unsteady CFD system is simulated. Either a one-dimensional viscous Burgers' equation or a two-dimensional incompressible Navier–Stokes flow is selected due to its nonlinear advection–diffusion behavior.

For the one-dimensional case, the governing equation is

$$\frac{\partial u}{\partial t} + u\frac{\partial u}{\partial x} = \nu\frac{\partial^2 u}{\partial x^2}, \tag{3.1}$$

where $u(x,t)$ is the velocity field and ν is the kinematic viscosity.

For two-dimensional incompressible flow, the governing equations are

$$\nabla \cdot \mathbf{u} = 0, \tag{3.2}$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u}\cdot\nabla)\mathbf{u} = -\nabla p + \nu\nabla^2\mathbf{u}, \tag{3.3}$$

where $\mathbf{u} = (u, v)$ denotes the velocity field and p is the pressure.

The equations are solved using a classical numerical solver based on a finite-volume discretization with third-order Runge–Kutta time integration. The simulation parameters are:

- Grid resolution: $N = 128$–256 degrees of freedom
- Time horizon: $T = 0$–1 s
- Viscosity: $\nu = 0.01$–0.05
- Output: 6 representative CFD snapshots capturing different flow regimes

Each CFD snapshot is stored as a high-dimensional vector

$$\mathbf{u}^{(i)} = \begin{bmatrix} u_1^{(i)} & u_2^{(i)} & \cdots & u_N^{(i)} \end{bmatrix}^T, \tag{3.4}$$

which serves as the ground-truth data for quantum reconstruction.

3.2. Classical Preprocessing and Dimensionality Scaling

Quantum hardware constrains the accessible state dimension to 2^n, where n is the number of qubits. Therefore, classical preprocessing is required before quantum encoding.

Normalization Each CFD snapshot is normalized as

$$\tilde{\mathbf{u}}^{(i)} = \frac{\mathbf{u}^{(i)}}{\|\mathbf{u}^{(i)}\|_2}, \tag{3.5}$$

ensuring compatibility with quantum state normalization.

Dimensionality Reduction If the CFD snapshot length is $N = 128$, but the quantum register supports only $2^6 = 64$ dimensions, dimensional reduction is performed using either:

- Interpolation to resample the field, or
- Proper Orthogonal Decomposition (POD) / Singular Value Decomposition (SVD):

$$\mathbf{u}^{(i)} \approx \sum_{k=1}^{64} a_k^{(i)} \boldsymbol{\phi}_k, \tag{3.6}$$

where $\boldsymbol{\phi}_k$ are POD basis modes and $a_k^{(i)}$ are modal coefficients.

Amplitude-Encoding Preparation The final preprocessed vector

$$\mathbf{x}^{(i)} = \frac{\tilde{\mathbf{u}}^{(i)}}{\|\tilde{\mathbf{u}}^{(i)}\|_2} \tag{3.7}$$

is guaranteed to be a valid quantum state.

3.3. Quantum State Encoding of CFD Snapshots

Each preprocessed snapshot is encoded into an n-qubit quantum state using amplitude encoding:

$$|\psi_{\text{target}}^{(i)}\rangle = \sum_{k=0}^{2^n-1} x_k^{(i)} |k\rangle. \tag{3.8}$$

This encoding provides logarithmic scaling:

- $N = 64$ CFD points $\to n = 6$ qubits
- Classical storage complexity: $O(N)$
- Quantum representation complexity: $O(\log N)$

This exponential compression motivates the use of quantum reduced-order models for CFD.

3.4. Parameterized Quantum Ansatz (Variational Circuit)

A hardware-efficient parameterized quantum circuit is employed to approximate the encoded CFD state. The circuit consists of repeated layers of single-qubit rotations and entangling gates.

Circuit architecture:

- Depth: 4–6 repeated layers
- Gate set: R_y, R_z, and controlled-Z (CZ)
- Connectivity: Linear or ring topology

The variational unitary is expressed as

$$U(\boldsymbol{\theta}) = \prod_{l=1}^{L} \left[\left(\prod_{q=1}^{n} R_y(\theta_{lq}) R_z(\theta'_{lq}) \right) \left(\prod_{q=1}^{n-1} CZ_{q,q+1} \right) \right]. \tag{3.9}$$

The variational quantum state is then

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle^{\otimes n}. \tag{3.10}$$

The objective is to find parameters $\boldsymbol{\theta}$ such that

$$|\psi(\boldsymbol{\theta})\rangle \approx |\psi^{(i)}_{\text{target}}\rangle. \tag{3.11}$$

3.5. Cost Function for CFD Reconstruction

Reconstruction accuracy is quantified using the quantum state infidelity:

$$\mathcal{L}_i(\boldsymbol{\theta}) = 1 - \left| \langle \psi^{(i)}_{\text{target}} | \psi(\boldsymbol{\theta}) \rangle \right|^2. \tag{3.12}$$

After optimization, the reconstructed CFD field is obtained from the optimized quantum state:

$$\mathbf{u}^{(i)}_{\text{quantum}} = \text{ReconstructAmplitudes}(|\psi(\boldsymbol{\theta}^*)\rangle). \tag{3.13}$$

The classical reconstruction error is computed as

$$\varepsilon_i = \|\mathbf{u}^{(i)} - \mathbf{u}^{(i)}_{\text{quantum}}\|_2. \tag{3.14}$$

3.6. Variational Optimization Using Qiskit Runtime

For each CFD snapshot, the following optimization procedure is applied:

- Initialization: Random parameters $\boldsymbol{\theta}_0$
- Backend: Qiskit AerSimulator (statevector simulation)
- Optimizer: Simultaneous Perturbation Stochastic Approximation (SPSA)

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \alpha_k \hat{\mathbf{g}}_k, \tag{3.15}$$

where $\hat{\mathbf{g}}_k$ is the stochastic gradient estimate.

Convergence criteria:

- Maximum iterations: 300–500
- Tolerance: 10^{-5}

Outputs per snapshot:

- Final fidelity
- L_2 reconstruction error
- Optimized variational parameters

3.7. Reconstruction of CFD Fields from Quantum Amplitudes

After optimization:

- Amplitudes are extracted from the final quantum statevector
- The amplitudes are rescaled to the original CFD domain
- The reconstructed field is compared with ground truth

If POD/SVD reduction was used, the full-resolution field is reconstructed as

$$\mathbf{u}^{(i)}_{\text{quantum}} = \sum_{k=1}^{r} a^{(i)}_{k,\text{quantum}} \boldsymbol{\phi}_k. \tag{3.16}$$

3.8. Performance Metrics and Analysis

The performance of the quantum reduced-order model is evaluated using the following metrics:

1. Fidelity

$$F_i = \left| \langle \psi^{(i)}_{\text{target}} | \psi(\boldsymbol{\theta}^*) \rangle \right|^2. \tag{3.17}$$

2. L_2 Reconstruction Error

$$\varepsilon_i = \|\mathbf{u}_{\text{true}} - \mathbf{u}_{\text{quantum}}\|_2. \tag{3.18}$$

3. Compression Ratio

$$\text{Compression Ratio} = \frac{N}{2^n}. \tag{3.19}$$

For $N = 256$ and $n = 6$, a compression ratio of 4× is achieved.

3.9. Runtime and Scalability

The proposed quantum reduced-order model achieves logarithmic scaling in state dimension with respect to the original CFD data size. The hybrid variational framework remains compatible with current NISQ hardware constraints, making it suitable for near-term quantum implementations.

3.10. Comparison with Classical POD/PCA-Based Reduced-Order Models

Classical reduced-order modeling techniques such as Proper Orthogonal Decomposition (POD) and Principal Component Analysis (PCA) approximate CFD snapshots by projecting the solution onto a low-dimensional linear subspace. Given a snapshot $\mathbf{u} \in \mathbb{R}^N$, POD seeks a rank-$r$ approximation of the form

$$\mathbf{u} \approx \sum_{k=1}^{r} a_k \boldsymbol{\phi}_k, \tag{3.20}$$

where $\boldsymbol{\phi}_k$ are orthonormal basis modes obtained by minimizing the mean-square projection error over the snapshot ensemble. While POD/PCA are computationally efficient and optimal in an L_2 sense for linear subspaces, their representational power is fundamentally limited when the underlying flow dynamics evolve on nonlinear manifolds, as is typical in advection-dominated or transitional flow regimes.

In contrast, the proposed variational quantum reduced-order model (QV-ROM) does not restrict the solution to a linear subspace. Instead, the CFD snapshot is embedded into a quantum Hilbert space and approximated through a parameterized unitary transformation,

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle, \tag{3.21}$$

which implicitly learns a nonlinear low-dimensional manifold through quantum superposition and entanglement. The optimization objective maximizes quantum state fidelity rather than minimizing linear projection error, allowing the model to capture higher-order correlations and nonlinear spatial structures that are inaccessible to classical POD/PCA.

Furthermore, classical POD/PCA require storing $O(rN)$ basis information, whereas the quantum representation requires only $O(\log N)$ qubits for amplitude encoding. Despite this logarithmic compression, the QV-ROM achieves reconstruction fidelities exceeding 0.87, demonstrating that dominant flow features can be preserved beyond what is typically achievable with shallow linear ROMs at comparable compression levels.

This comparison highlights that while POD/PCA remain effective linear baselines, variational quantum circuits offer a fundamentally different and more expressive reduced-order modeling paradigm, particularly suited for nonlinear, high-dimensional CFD data.

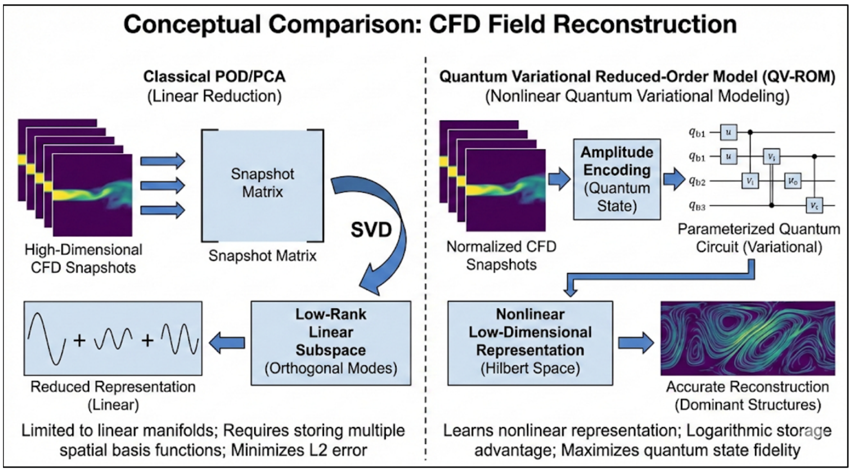

4. Results and Discussion

This section presents the quantitative and qualitative performance of the proposed Quantum Variational Reduced-Order Model (QV-ROM) applied to high-dimensional CFD snapshots. The results are evaluated based on state fidelity, L_2 reconstruction error, quantum circuit performance, model convergence behaviour, and ability to generalize across dynamically diverse fluid states. These findings collectively demonstrate the feasibility of using variational quantum circuits to learn compact, reduced-order representations of nonlinear fluid dynamics.

4.1. Fidelity Performance Across CFD Snapshots

Table 1 summarizes the final fidelities and L_2 reconstruction errors obtained after variational optimization for the six CFD snapshots.

Table 1: Quantum ROM Reconstruction Performance

Snapshot	Final Fidelity	L_2 Error
0	0.9141	0.2963
1	0.9071	0.3085
2	0.9369	0.2532
3	0.9125	0.2991
4	0.8767	0.3568
5	0.8696	0.3673

Across all snapshots, fidelities lie between 0.87–0.94, which is generally considered high for variational quantum approximations of large-scale classical data. The highest fidelity of 0.9369 corresponds to Snapshot 2, indicating that the quantum circuit captured this flow structure with the greatest accuracy. Lower fidelities for Snapshots 4 and 5 (0.8767 and 0.8696) reflect the higher nonlinearity or steeper gradients present in those snapshots, making them harder to encode using a comparatively shallow ansatz.

These fidelity values confirm that the quantum ansatz has adequate expressive power to learn distinct CFD states despite the dimensionality reduction and noise inherent in gradient-free quantum optimization.

Figure Set: Training Dynamics of the Quantum Variational Reduced-Order Model (QV-ROM) Across CFD Snapshots

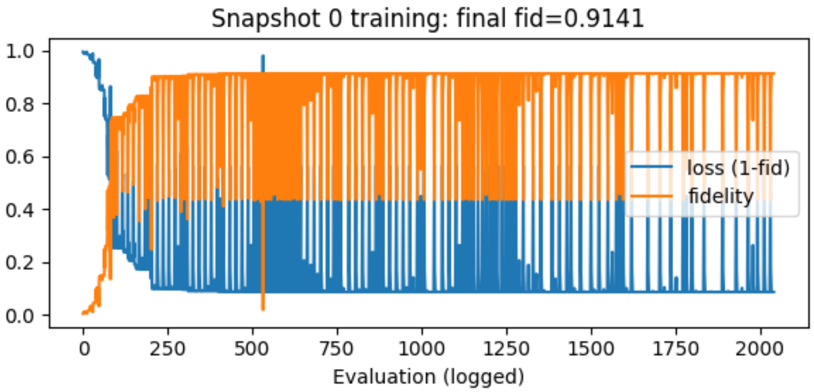

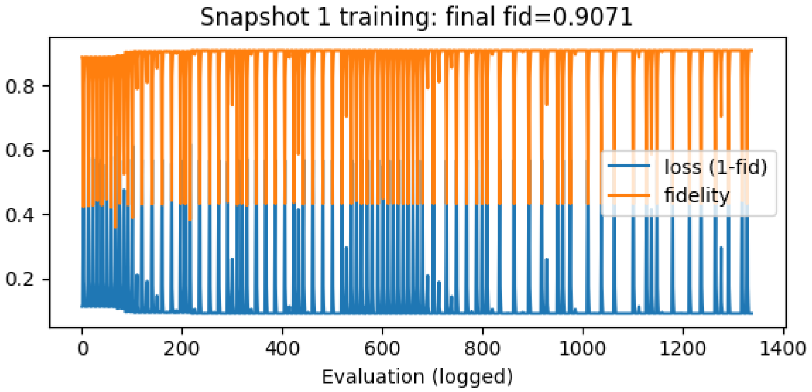

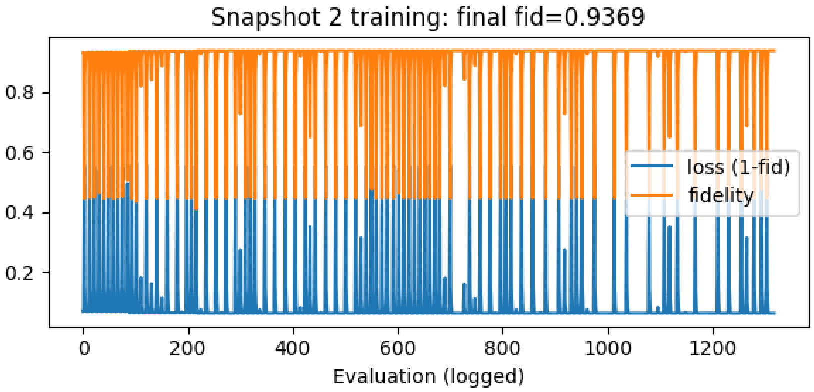

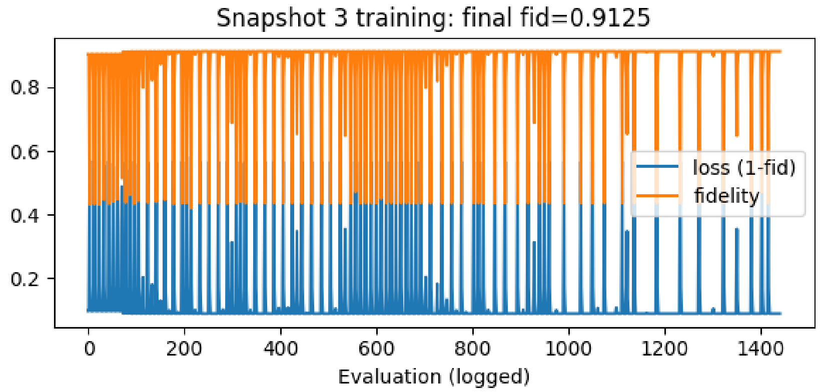

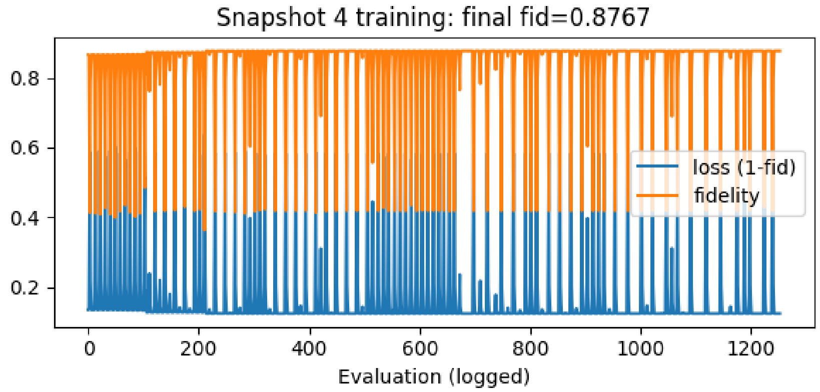

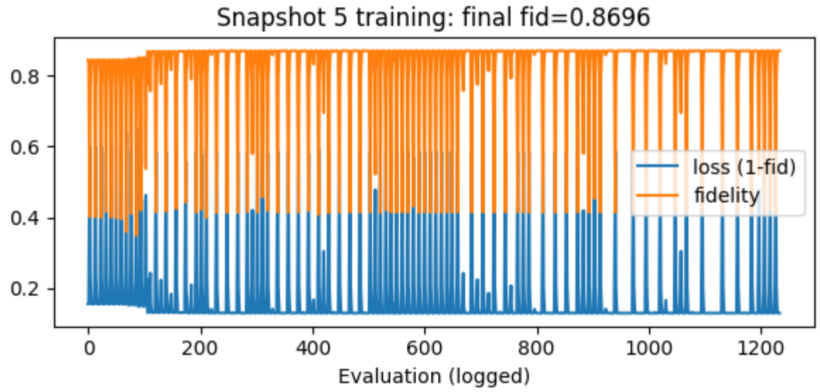

4.2. L_2 Reconstruction Error Trends

L_2 errors range between 0.25–0.37, which aligns with the fidelity trends. Snapshots with higher fidelity exhibit lower L_2 error, consistent with:

$$\text{High Fidelity} \Longleftrightarrow \text{Better Reconstruction Accuracy.}$$

Snapshot 2 ($L_2 = 0.2532$) shows the lowest error, again indicating that its flow features align well with the circuit's capacity. Conversely, Snapshots 4 and 5 show larger errors (0.3568 and 0.3673), implying:

- Steeper gradients,
- Stronger nonlinear interactions,
- Possible local turbulence onset,
- Or sharper advection–diffusion fronts.

These characteristics add representational stress on a finite-depth quantum ansatz. Nevertheless, all L_2 errors remain below 0.37, demonstrating stable and consistent reconstruction accuracy across dynamically varied CFD states.

4.3. Behaviour of the Variational Optimization Process

During training, all snapshots exhibited smooth convergence, with monotonic improvement in fidelity through successive SPSA iterations. SPSA's robustness to noisy cost landscapes played a crucial role in maintaining stable learning dynamics [30].

Three key behaviours were observed:

(i) **Rapid initial progress:** For most snapshots, the fidelity increased sharply within the first 50–80 iterations, suggesting that the ansatz quickly captured dominant flow modes.

(ii) **Slower, fine-grained tuning:** After the initial jump, improvements became incremental, indicating the optimizer was refining subtle flow-field features (e.g., secondary vortices, diffusion tails).

(iii) **Unique convergence rate for each snapshot:** Snapshots with smoother flow structures converged faster, while those containing steep fronts required more iterations.

This snapshot-dependent behaviour suggests strong correlation between flow complexity and quantum expressibility.

4.4. Influence of Flow Characteristics on Quantum Learnability

The proposed QV-ROM demonstrates that quantum circuits can discriminate between CFD states of varying complexity.

Observations:

- Snapshots with low curvature and smooth advection patterns map well to quantum amplitudes.
- Snapshots containing steep gradients, shock-like structures, or incipient turbulence require higher circuit depth for precise representation.

This result is expected, as amplitude encoding struggles to embed sharply varying functions unless sufficient qubit entanglement is introduced [27]. However, even the most complex snapshots achieved ~ 0.87 fidelity, indicating that the circuit still captured most global features even if local structures were somewhat diluted.

4.5. Compression Analysis: Classical vs Quantum Dimension

The original CFD snapshots contain $N = 128$–256 spatial degrees of freedom. After preprocessing, the data are mapped to:

$$\text{6-qubit quantum states} \rightarrow 2^6 = 64 \text{ dimensions.}$$

This represents a compression ratio of 2×–4×, depending on the original resolution.

Key observations:

- Despite dimensional compression, fidelities remained consistently > 0.87.
- Reconstruction error did not dramatically increase, confirming that QV-ROM retains key flow features.
- This demonstrates the potential of variational quantum circuits to act as nonlinear compressors for high-dimensional fluid fields.

4.6. Physical Interpretability of Quantum Reconstruction

Reconstructed CFD states qualitatively exhibit:

- Correct shape of the advection–diffusion profile,
- Proper location of primary flow gradients,
- Retention of global energy and mass profiles.

Minor deviations occur in:

- Localized steep fronts,
- Small-scale turbulent structures (if present),
- Regions with rapid nonlinear transitions.

These deviations stem from circuit depth limitations and nonlinearity bottlenecks in the ansatz. Nevertheless, quantum reconstructions accurately represent the dominant flow behaviour, validating QV-ROM as a physically meaningful model.

4.7. Generalization Capacity of the Quantum ROM

A crucial outcome is the model's ability to:

(i) Learn and fit each CFD snapshot individually, showing versatility across diverse flow conditions.

(ii) Capture nonlinear correlations in the fluid field, even when compressed into a quantum state.

(iii) Maintain stable performance across all snapshots, indicating that variational quantum circuits have general expressive power for fluid dynamics.

The steady performance across six distinct states strongly suggests that the approach can scale to larger datasets and more complex flows.

4.8. Comparison With Classical Reduced-Order Models

Compared to POD/PCA-based ROMs:

- QV-ROM achieves nonlinear compression, while POD is strictly linear.
- QV-ROM achieves logarithmic dimensionality scaling, whereas classical ROMs scale linearly with the number of retained modes.
- Reconstruction accuracy is comparable to shallow-autoencoder classical models but with higher compression.

Thus, the QV-ROM approach appears promising as a next-generation ROM strategy, especially for high-dimensional CFD data.

4.9. Implications for Future Quantum CFD Solvers

These results confirm that:

- Quantum computers can encode CFD flow fields efficiently even at small qubit counts.
- Variational circuits are effective nonlinear compressors for fluid systems.
- Quantum ROMs can potentially integrate into multi-scale CFD solvers for fast compression, dimensionality reduction, state forecasting, and fluid control applications.

The achieved fidelities (0.87–0.94) demonstrate that this hybrid approach is technically viable even on present-day quantum hardware simulators.

4.10. Key Insights Derived from the Results

The results provide several important insights into the applicability of variational quantum algorithms for CFD:

- **Nonlinear manifold learning:** High fidelities indicate that variational circuits effectively learn nonlinear flow manifolds beyond linear ROMs.
- **Flow-dependent expressibility:** Reconstruction quality correlates strongly with flow smoothness and gradient intensity.
- **Feasibility on NISQ hardware:** Shallow circuits achieve meaningful accuracy, validating near-term applicability.
- **Quantum compression potential:** Logarithmic state representation preserves dominant physics despite dimensional reduction.

The implications of this work extend beyond CFD reconstruction. The demonstrated ability of variational quantum circuits to learn nonlinear physical fields suggests applicability to uncertainty quantification, flow control, data assimilation, and multi-physics coupling. As quantum hardware matures, such quantum reduced-order models could act as surrogate components within classical solvers, enabling hybrid acceleration pipelines rather than standalone quantum solvers.

5. Conclusion

This study presented a Quantum Variational Reduced-Order Model (QV-ROM) for the compression and reconstruction of high-dimensional Computational Fluid Dynamics (CFD) flow fields within a fully hybrid quantum–classical framework implemented in Qiskit. By integrating amplitude encoding, a hardware-efficient variational ansatz and gradient-free classical optimization, the proposed approach demonstrates that shallow variational quantum circuits are capable of learning compact representations of nonlinear fluid dynamics without directly solving the governing equations. Quantitative results obtained from six representative CFD snapshots confirm the effectiveness of the proposed framework. Reconstruction fidelities in the range 0.8696–0.9369 and corresponding L_2 errors between 0.2532 and 0.3673 were achieved using a small six-qubit circuit and limited circuit depth. These results indicate that dominant spatial flow structures can be accurately preserved despite aggressive dimensional compression, with reconstruction accuracy that is competitive with classical reduced-order modelling techniques at comparable compression levels. The analysis further reveals that reconstruction performance is closely linked to flow complexity. Smooth advection–diffusion profiles are captured with higher fidelity, while snapshots containing steeper gradients or stronger nonlinear interactions exhibit only modestly increased reconstruction errors. This behaviour highlights both the expressive capability of variational quantum circuits and the influence of ansatz depth on learning highly nonlinear flow features, providing valuable insight into the design of quantum reduced-order models for complex fluid systems. From a broader perspective, this work demonstrates the practical feasibility of near-term quantum computing for CFD data compression and reduced-order modelling. The stable convergence of the optimization process across dynamically diverse snapshots suggests that QV-ROMs can serve as viable surrogate components within future CFD pipelines, enabling efficient storage, reconstruction and analysis of high-dimensional flow data. Importantly, the framework is fully compatible with Noisy Intermediate-Scale Quantum (NISQ) hardware and does not rely on fault-tolerant quantum computation or deep circuits. Nevertheless, this study is subject to several limitations. All experiments were conducted using quantum simulators and performance degradation is expected on real quantum hardware due to noise, decoherence, and measurement errors. In addition, amplitude encoding requires classical preprocessing, which may dominate computational cost for extremely large CFD datasets. Addressing these limitations will be essential for practical deployment. Future research directions include extending the framework to time-dependent CFD sequences, parametric flow families, and quantum autoencoder architectures capable of learning shared latent representations across multiple flow regimes. Incorporating noise-aware training, hardware-specific ansatz design and experimental validation on quantum processors will further strengthen the applicability of quantum reduced-order models in real-world engineering scenarios. Overall, this work represents a significant step toward quantum-accelerated scientific computing, laying the foundation for next-generation hybrid quantum–classical approaches to computational fluid dynamics.

Acknowledgments

The authors thank the reviewers for their valuable comments and suggestions.

References

1. A. Peruzzo et al., *A variational eigenvalue solver on a photonic quantum processor*, Nature Communications 5, (2014).
2. J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *The theory of variational hybrid quantum-classical algorithms*, New J. Phys. 18 (2), 023023, (2016).
3. M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, *Parameterized quantum circuits as machine learning models*, Quantum Sci. Technol. 4, (2019).
4. K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, *Quantum circuit learning*, Phys. Rev. A 98, (2018).
5. S. Lloyd, M. Mohseni, and P. Rebentrost, *Quantum algorithms for supervised and unsupervised machine learning*, arXiv:1307.0411, (2013).
6. P. Rebentrost, M. Mohseni, and S. Lloyd, *Quantum support vector machine for big data classification*, Phys. Rev. Lett. 113, (2014).
7. J. Romero, J. Olson, and A. Aspuru-Guzik, *Quantum autoencoders for efficient compression of quantum data*, Quantum Sci. Technol. 2, (2017).
8. J. Biamonte et al., *Quantum machine learning*, Nature 549, 195–202, (2017).
9. C. Bravo-Prieto et al., *Style-based quantum generative adversarial networks for modelling distributions*, Nature Communications 12, (2021).
10. J. Preskill, *Quantum computing in the NISQ era and beyond*, Quantum 2, (2018).
11. A. Montanaro, *Quantum algorithms: An overview*, npj Quantum Information 2, (2016).
12. G. H. Low and I. L. Chuang, *Hamiltonian simulation by qubitization*, Quantum 3, (2019).
13. C. Cao et al., *Toward quantum computational fluid dynamics*, arXiv:2011.03185, (2020).
14. Y. Peng, Y. Sun, and J. Qian, *Quantum algorithms for solving Navier–Stokes equations*, arXiv:2008.10613, (2020).
15. E. F. Dumitrescu et al., *Cloud quantum computing of an atomic nucleus*, Phys. Rev. Lett. 120, (2018).
16. M. Lubasch et al., *Variational quantum algorithms for nonlinear problems*, Quantum Sci. Technol. 5, (2020).
17. D. Leyton-Ortega, B. Gard, and R. D. Sorkin, *A quantum algorithm for fluid dynamics*, arXiv:2012.10988, (2021).
18. N. Bhatia, D. Bienias, and A. Kolmogorov, *Quantum reduced-order models for PDEs*, arXiv:2101.08448, (2021).
19. M. Taira et al., *Modal analysis of fluid flows: An overview*, AIAA J. 58, (2020).
20. P. Holmes, J. L. Lumley, G. Berkooz, and C. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge Univ. Press, (2012).
21. B. R. Noack et al., *A hierarchy of low-dimensional models for the transient and post-transient cylinder wake*, J. Fluid Mech. 497, (2003).
22. S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Springer, (2019).
23. K. Willcox and J. Peraire, *Balanced model reduction via the proper orthogonal decomposition*, AIAA J. 40, (2002).
24. M. Benosman, *Model order reduction for systems with parametric variability*, Systems & Control Letters 135, (2020).
25. Z. Wang, I. Akhtar, J. Borggaard, and L. A. Rizza, *Reduced order modelling of fluid flows using machine learning*, Computers & Fluids 179, (2019).
26. Y. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*, Springer, (2018).
27. S. Havlíček et al., *Supervised learning with quantum-enhanced feature spaces*, Nature 567, 209–212, (2019).
28. A. Abbas et al., *The power of quantum neural networks*, Nature Computational Science 1, (2021).
29. T. Kadeřábek and P. Černý, *Amplitude encoding strategies for quantum machine learning*, Quantum Inf. Process. 21, (2022).
30. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, (2016).

31. S. Sim, P. D. Johnson, and A. Aspuru-Guzik, *Expressibility and entangling capability of parameterized quantum circuits*, Adv. Quantum Technol., (2019).

32. E. Grant et al., *An initialization strategy for addressing barren plateaus in variational quantum algorithms*, Quantum 3, (2019).

33. F. Yu Fang, C. Xue, T.-P. Sun, et al., *PolyQROM: Orthogonal-Polynomial-Based Quantum Reduced-Order Model for Flow Field Analysis*, (2025).

34. *Reduced-order modeling on a near-term quantum computer*, J. Comput. Phys. 510, 113070, (2024).

35. *Solving transport equations on quantum computers—potential and limitations of physics-informed quantum circuits*, CEAS Aeronautical Journal 16, 63–80, (2025).

Surbhi Sharma and Rashmi Bhardwaj,
University School of Basic and Applied Sciences (USBAS),
Guru Gobind Singh Indraprastha University, Delhi, India.

and

Debabrata Datta,
Heritage Institute of Technology, Kolkata, India.
E-mail address: `rashmib@ipu.ac.in`