



HYBRID QUANTUM-CLASSICAL ADAPTIVE INTELLIGENCE FRAMEWORK FOR REAL-TIME OPTIMIZATION IN NOISY INTERMEDIATE-SCALE QUANTUM (NISQ) SYSTEMS

Kondamudi Sivaparvathi, C. Prabhavathi and B. Muni Lavanya
Department of CSE, JNTUA College of Engineering, Pulivendula, Andhra Pradesh, India
E-Mail: kondamudisivaparvathi2@gmail.com

ABSTRACT

The hardware noise and drift, combined with the fragility of stationary hybrid algorithms like the variational quantum eigensolver (VQE), limit the usefulness of noise quantum advantage accessed through Noisy intermediate-scale quantum (NISQ) processors, but allow near-term quantum advantage. In this paper, a Hybrid Quantum-Classical Adaptive Intelligence Framework is suggested, which bridges the gap between quantum hardware, a parameterized quantum circuit, explicit error-mitigation routines, and a controller trained through learning. The quantum subsystem prepares variational states and calculates the expectation values of the problem Hamiltonian, and an AI/ML model approximates an online noise probability signal and moves an optimizer based on reinforcement learning (RL). The RL policy chooses the parameter updates and enables mitigation measures depending on noise-sensitive state functions, such as energy, gradient norm, and estimated noise level. It was demonstrated experimentally using a simulated NISQ-style molecular Hamiltonian that the adaptive controller can reduce the error in the absolute ground state of the system by about 52.8 percent versus a standard static VQE baseline using the same ansatz. The integrated classifier that drives the process of mitigation is accurate with 97 percent accuracy, 100 percent precision, 92 percent recall, 95.9 F1-score, and 99.5 AUC, which shows reliable decision-making with minimal false alarms. All these findings indicate that the close-knit quantum-classical adaptivity can significantly contribute to the robustness and accuracy of VQE-like workloads on NISQ and serve as a scalable template to design noise-aware hybrid algorithms in the future.

Keywords: NISQ systems, hybrid quantum-classical computing, variational quantum eigensolver (VQE), reinforcement learning, adaptive optimization, quantum error mitigation.

Manuscript Received 20 March 2026; Revised 28 April 2026; Published 10 July 2026

1. INTRODUCTION

Quantum computing has become a paradigm shift that has the potential to solve complex optimization, simulation, and learning problems that classical systems cannot. The present generation of quantum hardware, the so-called Noisy Intermediate-Scale Quantum (NISQ) era, includes devices with 50-1000 qubits with no full-scale error correction [1]. These systems have large levels of decoherence, gate defects, and hardware variation, which restrict the execution of traditional quantum algorithms like the Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Eigensolver (VQE) [2], [3]. In turn, this has led to the rise of hybrid quantum-classical architectures, in which the classical processors refine and refine quantum circuits to reduce noise and improve fidelity [4].

Although hybrid systems like VQE and QAOA have promise, they usually use fixed or quasi-fixed optimization cycles that do not make use of real-time hardware feedback [5]. NISQ devices are characterized by stochastic noise, temporal drift, and qubit performance variation, so cold optimization techniques are not optimal for the deployment of the devices [6]. Also, the existing classical optimization algorithms used (i.e., gradient descent, Nelder-Mead, and SPSA) tend to be inefficient in noisy environments and are unable to cope with changes in underlying system dynamics [7].

Recent breakthroughs in artificial intelligence (AI), specifically deep reinforcement learning (DRL), have made it possible to control complex dynamic systems, which are data-driven, and provide promising directions in adaptive quantum optimization [8], [9]. Combining DRA with quantum circuits also provides the possibility of constructing intelligent controllers, which learn hardware behavior, tune circuit parameters, reduce the frequency of errors, and enhance the convergence of algorithms. A number of studies have been conducted on ML-aided error reduction and circuit optimization, but most of these methods have been demonstrated as restricted to offline training, fixed datasets, or noise-sensitive simulation instead of real-time adaptive control [10], [11].

To fill this gap, the proposed Hybrid Quantum-Classical Adaptive Intelligence Framework, which utilizes reinforcement learning and deep neural architectures to continuously monitor quantum hardware performance and adjust parameters of algorithmically controlled processes, will be proposed in this research. The adaptive loop is designed to advance the scalability, robustness, and fidelity of quantum algorithms implemented on NISQ hardware. The proposed framework combines the ideas of real-time feedback, optimization based on machine learning, and reconfiguring quantum circuits, providing a



new avenue towards high-performance quantum computation within realistic hardware constraints.

2. LITERATURE REVIEW

The leading paradigm of utilizing noisy intermediate-scale quantum (NISQ) devices is now hybrid quantum-classical, based on variational quantum algorithms (VQAs). Cerezo *et al.* give a background overview of VQAs, their structure, which consists of parameterized quantum circuits that are optimized by classical algorithms, and the main challenges in this particular field, which include barren plateaus, expressibility, and sensitivity to noise in realistic NISQ devices [1]. Recent analyses build upon this viewpoint to compare various families of VQA (VQE, QAOA, quantum neural networks) systematically with regard to realistic noise models and point to the fact that performance dependence is determined by joint design of the ansatz, optimizer, and hardware noise behavior [2], [3]. These papers define that even though VQAs are very flexible and can be implemented efficiently in hardware, they are usually implemented using static optimization loops: the classical optimizer performs a parameter update as a result of cost-function evaluations, but never adapts to time-varying noise or drift in the device itself.

In addition to particular algorithms, various surveys are keen to stress that hybrid quantum-classical systems will be at the center stage even after the NISQ era. Callison and Chancellor say that hybrid algorithms will keep scaling to classical processors, where quantum hardware is employed as a specialized accelerator; they also classified the important bottlenecks like optimization landscapes, noise resistance, and scaling [4]. Ge *et al.* study the optimization landscape of hybrid quantum-classical schemes and demonstrate that the pathological effects of local minima and barren plateaus can be extremely detrimental to classical optimizers in the noisy regimes, particularly in the absence of problem- or hardware-adaptive schemes [5]. Simultaneously, De Luco *et al.* and Peral-Garcia *et al.* survey larger-scale hybrid quantum-classical machine learning and quantum machine learning (QML) applications in general, finding that most of the proposed models can be described as viewing quantum devices as stationary components in a pipeline, with few serious attempts currently being made to incorporate deep feedback of device-driven feedback into the learning process [6], [7]. These papers highlight the importance of more dynamic, feedback-based hybrid architectures, but they do not go further and argue for actual real-time adaptive control structures.

An analogous field of research studies reinforcement learning (RL) and deep RL for quantum control. Niu *et al.* propose a deep RL-based framework of universal quantum control that jointly employs speed and fidelity optimization of control pulses under stochastic errors and show that deep RL can achieve better performance compared to traditional gradient-based optimal control in noisy environments [8]. Zhang *et al.*

offer a systematic comparison between RL algorithms (tabular Q-learning, deep Q-learning, policy gradient) and non-ML optimizers for state preparation problems and demonstrate that RL has better scaling when the control spaces are discrete, and the problem size increases [9]. A model-free RL algorithm is suggested by Sivak *et al.*, according to which an agent directly interacts with a quantum system and learns control policies without a system model [10]. Porotti *et al.* also show that deep RL is able to prepare and stabilize complex quantum states in cavity systems with decoherence [11]. Recent advances in physics-informed RL for quantum control include constraints and a priori physical knowledge to enhance sample efficiency in the face of realistic hardware constraints [12]. Taken together, these results demonstrate RL as a potent approach to adapting control protocols to noisy quantum dynamics, but generally, they involve low-level pulse or state-preparation (as opposed to end-to-end hybrid optimization loops (e.g., VQE or QAOA)) with error mitigation and real-time adaptive control.

Simultaneously, error mitigation and noise learning with machine learning has become a strategy of significant focus to increase the utility of NISQ devices that do not completely correct errors. Vassallo introduces a machine-learning-based quantum error mitigation (QEM) scheme that trains on noisy quasi-probability distributions to learn distributions that approximate ideal distributions, which points to the fact that ML models can learn to capture the behavior of noise that simple analytical models fail to do [13]. Cemin *et al.* apply supervised learning to learn lower-rank quantum channels from quantum hardware experimental data, showing that learned models are able to accurately simulate the dynamics of devices and therefore could be useful not only in calibration but also in mitigation [14]. Czarnik *et al.* propose a cheaper learning-based mitigation protocol that is designed on the basis of the Clifford data regression (CDR) that involves a much smaller number of hardware calls and has an equivalent correction accuracy [15]. Khan *et al.* test the methods of mitigation of measurement errors on the IBM NISQ devices and affirm that in the cases where hardware peculiarities are carefully considered during mitigation, significant improvements in accuracy can be observed [16]. These works demonstrate the usefulness of data-driven and ML-aided QEM; but mitigation is rarely implemented as an active and direct element of a continuous feedback loop that also regulates the structure and optimization dynamics of circuits.

The themes are further contextualized by the recent surveys of the NISQ-era advances. Xie *et al.* discuss the progress of quantum computations in the NISQ era, and argue that hybrid algorithms, noise-aware ansatzes, and error-mitigation techniques should be designed together with hardware to achieve practical benefits [17]. However, the current majority of frameworks consider hybrid optimization, quantum control, and error mitigation to be highly independent modules: control RL agents are operating on a pulse-by-



pulse basis, VQAs are operating with fixed classical optimizers, and QEM modules are being calibrated or operated as a static correction. The gap that emerges in the literature is therefore the integrated, AI-driven hybrid systems which can: (i) learn device behaviour online, (ii) coordinate circuit parameters and ansatz structure and

qubit mappings and mitigation strategies, and (iii) execute in a closed loop with real NISQ hardware to optimise the end task. The suggested Hybrid Quantum Classical Adaptive Intelligence Framework is specifically set up to fill this gap by incorporating these components into one real-time adaptive optimization loop.

Table-1. Comparison table of representative works.

Work / Year	Main Focus	Technique	Adaptivity Level	Hardware / Setting	Limitation w.r.t. Proposed Framework
Cerezo <i>et al.</i> (2021) [1]	Survey of VQAs and challenges	Variational quantum algorithms with classical optimizers	Static parameter optimization; no online hardware adaptation	General NISQ (theoretical + simulations)	Does not implement real-time adaptive control or integrated RL/ML loop.
Callison & Chancellor (2022) [4]	Hybrid algorithms in NISQ era	Conceptual and algorithmic survey	Conceptual discussion only	General NISQ	Identifies need for co-design but no concrete adaptive framework.
Ge <i>et al.</i> (2022) [5]	Optimization landscapes of hybrid algorithms	Analysis of cost landscapes	Static; focuses on mathematical properties	Simulations	Does not leverage ML/RL or real-time device feedback.
Niu <i>et al.</i> (2019) [8]	Quantum control optimization	Deep RL for pulse-level control	Adaptive at control-pulse level	Simulations and hardware-motivated models	Targets control of specific gates/states; not full hybrid VQA/QAOA loop or error mitigation.
Sivak <i>et al.</i> (2022) [10]	Model-free quantum control	RL agent interacts directly with system	Online adaptation of control policy	Experimental quantum system	Focuses on state/control tasks; no integration with hybrid algorithm optimizers or QEM.
Porotti <i>et al.</i> (2022) [11]	State prep and stabilization	Deep RL for cavity-QED systems	Adaptive for state prep and stabilization	Hardware-relevant cavity setup	Task-specific; does not generalize to generic NISQ hybrid workflows.
Vassallo (2023) [13]	ML-based quantum error mitigation	Supervised learning of quasi-probability corrections	Offline training; static mitigation model	NISQ hardware + simulators	Mitigation not coupled to live optimization or control loop.
Czarnik <i>et al.</i> (2025) [15]	Efficient learning-based error mitigation	Clifford data regression with ML	Partially adaptive via learned regression	IBM NISQ devices	Optimizes mitigation cost but not circuit/ansatz or optimizer behavior in real time.
Cemin <i>et al.</i> (2024) [14]	Learning reduced quantum channels	Supervised learning of noise channels	Offline model learning	Real NISQ devices	Provides a noise model; does not enact adaptive hybrid control.
Xie <i>et al.</i> (2025) [17]	NISQ advances and performance of hybrids	Survey and case studies	Conceptual; no closed-loop implementation	General NISQ	Highlights importance of co-design, but lacks unified adaptive framework.

2.1 Research Gap

Despite the promising results of hybrid quantum-classical algorithms in allowing calculation on the noisy intermediate-scale quantum (NISQ) machines, a number of important weaknesses still persist in the current state of the art. The variational algorithms currently in use, including VQE and QAOA, are primarily based on classical optimizers that have either low or no ability to

access real-time information at the device level, including noise drift, qubit decoherence, or changes in hardware calibration [1], [2]. Consequently, they fail to perform well when introduced on realistic NISQ systems where there is variability in temporal noise and weak qubit fidelities. Furthermore, despite the current developments in reinforcement learning (RL) and deep RL, which have made adaptive quantum control available at the pulse and



state-preparation scales, these algorithms do not exist in the context of hybrid algorithmic processes. They do not optimise the structure of the ansatz, qubit placement, circuit depth, optimizer behaviour, or error-reduction strategies in a single framework [3], [4]. Similarly, noise learning and error mitigation algorithms using machine learning are based on offline calibration or post-processing, but not on continuous, online correction that learns and adapts with the hardware as the algorithm is applied to execute your task [5], [6].

Another problem that is addressed by the current literature is that hybrid algorithms, quantum control, and error mitigation are currently seen as separate entities, and a way of combining them into a single adaptive loop, where quantum hardware can learn in real time, is not described. There is no current framework that has a closed-loop hybrid intelligence system, which concurrently considers (i) dynamic model learning of hardware noise properties, (ii) reinforcement learning-based adaptive circuit reconfiguration, (iii) task-oriented optimization of algorithms like VQE and QAOA. Additionally, the current hybrid systems do not have the potential to adapt to changing hardware conditions autonomously, thus leading to poor performance and limited scalability in real NISQ settings.

Thus, a research gap remains in creating a Hybrid Quantum-Classical Adaptive Intelligence Framework that combines real-time Hardware Feedback, Reinforcement Learning, Dynamic Circuit Adaptation, and Online Error Mitigation in a single hybrid architecture. The gap is an important issue to bridge in improving fidelity, robustness, and computational efficiency on existing as well as near-term NISQ hardware.

2.2 Problem Statement

Noisy Intermediate -Scale Quantum (NISQ) devices are plagued by severe hardware impairments, such as decoherence, stochastic noise variations, gate variations, and qubit time variations. Although hybrid quantum and classical algorithms, like VQE and QAOA, are intended to take advantage of those devices, they mainly use fixed classical optimization subroutines that do not consider time-dependent changes in hardware behavior. The current methods of reinforcement learning-based quantum control and machine learning-aided error correction are in isolated scenarios and lack incorporation into end-to-end hybrid algorithmic processes. As a result, the hybrid systems are unable to adapt quantum circuits, qubit mappings, optimization strategies, and mitigation mechanisms on their own. This translates into poor performance, slower convergence, and lower fidelity on real NISQ hardware implementations of the algorithms. Thus, the underlying issue being solved in this study is the lack of a single, real-time, adaptive quantum-classical optimization model that would be able to: run noise and drift prediction and control in hardware dynamically, real-time control of circuit parameters and structure of ansatzes, combining reinforcement learning with smart

control choice, and implementing online error mitigation in a closed-loop hybrid execution facility. In the absence of this kind of adaptive intelligence system, the possible computational benefit of hybrid algorithms in the NISQ age will be restricted, and its useful scaling and practical resilience to a broad range of real-world quantum computing tasks will be impeded.

2.3 Research Objectives

The main goal of the study is to develop and test a Hybrid Quantum-Classical Adaptive Intelligence Framework that has the capability to optimize and control NISQ-era quantum algorithms in real time. The following are the specific objectives of the study that are aimed at the achievement of this objective:

Objective 1

To create a real-time quantum hardware feedback measurement and modelling system that operationally monitors qubit performance, noise properties, gate errors, and temporal drift and thus allows the continuous characterization of NISQ devices' behavior.

Objective 2

A reinforcement learning-based adaptive controller that can learn its best quantum circuit adjustments such as parameter changes, structure reconfiguration of the ansatz, and qubit remapping, based on real-time hardware feedback and task-relevant performance indicators.

Objective 3

To incorporate machine learning-aided online error mitigation methods into the hybrid workflow, allowing the system to automatically revise the mitigation strategies according to the changing noise conditions throughout the execution of an algorithm.

Objective 4

To build and test the proposed framework on benchmark hybrid algorithms (e.g., VQE, QAOA) based on both noise simulation and real quantum systems, as well as to compare its performance with traditional static hybrid methods with regard to fidelity, convergence rate, robustness, and computational efficiency.

All of these goals will enable the creation of a coherent, intelligent, and adaptable hybrid architecture that can make quantum algorithms more realistic and applicable to the real world of NISQ devices.

3. PROPOSED METHODOLOGY

The Hybrid Quantum-Classical Adaptive Intelligence Framework is a proposed architecture that uses a closed loop of the control system to optimize the real-time operation on NISQ hardware, as shown in Figure-1. The architecture integrates a parameterized quantum circuit with a classical AI/ML controller and



explicit error-reduction mechanisms. In this section, we formalize the important components and their interactions.

A. Variational Quantum Subsystem

Encoding of the target problem in a Hermitian Hamiltonian.

$$H = \sum_k h_k P_k,$$

where P_k are tensor products of Pauli operators and $h_k \in \mathbb{R}$.

The quantum circuit block in Fig. 1 implements a parameterized ansatz.

$$|\psi(\theta)\rangle = U(\theta)|0\rangle^{\otimes n_q},$$

where $U(\theta)$ is a unitary composed of rotation and entangling gates, and $\theta \in \mathbb{R}^d$ denotes the trainable parameters.

The objective function is the variational energy

$$E(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle = \sum_k h_k \langle P_k \rangle_\theta,$$

with expectation values $\langle P_k \rangle_\theta$ estimated on the quantum hardware using a finite number of shots. On ideal hardware, the optimization problem is

$$\theta^* = \arg \min_{\theta} E(\theta).$$

On NISQ devices, however, the measured energy

$$\tilde{E}(\theta, \gamma_t)$$

is corrupted by noise characterized by an effective noise signal $\gamma_t \in [0,1]$ at time step t . The goal of the proposed framework is to find θ that minimizes the true ground-state energy while adaptively compensating for time-varying γ_t .

B. Adaptive Optimization Loop

The “Adaptive Optimization” block in Figure-1 is driven by a classical controller that proposes parameter updates based on feedback from the quantum hardware and the AI/ML model.

At iteration t , the controller observes a state vector

$$s_t = [\tilde{E}(\theta_t, \gamma_t), \|\nabla_{\theta} \tilde{E}(\theta_t, \gamma_t)\|, \gamma_t, t],$$

where the gradient term may be estimated via parameter-shift rules or finite differences.

The controller then outputs an action

$$a_t = \Delta\theta_t,$$

and the circuit parameters are updated as

$$\theta_{t+1} = \theta_t + \Delta\theta_t.$$

In the reinforcement-learning formulation, the controller is modeled as a policy

$$\pi_{\phi}: s_t \mapsto a_t,$$

parameterized by weights ϕ . The policy is trained to maximize the expected discounted return

$$J(\phi) = \mathbb{E}_{\pi_{\phi}} \left[\sum_{t=0}^T \gamma_{\text{RL}}^t r_t \right],$$

where $\gamma_{\text{RL}} \in (0,1)$ is the RL discount factor and the reward is defined as

$$r_t = -\tilde{E}(\theta_t, \gamma_t) - \lambda C(\gamma_t, a_t).$$

Here, $C(\gamma_t, a_t)$ penalizes overly aggressive updates in high-noise regimes, and $\lambda > 0$ is a regularization coefficient.

Thus, the adaptive optimization block replaces a fixed classical optimizer with a learned policy that explicitly incorporates noise estimates and hardware feedback.

C. AI/ML Model and Noise-Aware Decision Logic

The AI/ML model in Fig. 1 receives feedback from quantum hardware (measurement statistics, calibration data, or proxy metrics) and produces an estimate of the instantaneous noise level,

$$\hat{\gamma}_t = f_{\psi}(z_t),$$

where z_t denotes the raw feedback features and f_{ψ} is a trained regression or classification network.

On top of this estimate, a decision module implements a binary policy for error-mitigation activation:

$$d_t = \begin{cases} 0, & \hat{\gamma}_t < \tau, (\text{passive monitoring}) \\ 1, & \hat{\gamma}_t \geq \tau, (\text{mitigation active}), \end{cases}$$

where τ is a learned noise threshold.

This logic corresponds to the qualitative behavior seen in the control deck: for $\hat{\gamma}_t < \tau$, the system operates in “PASSIVE MONITORING” mode and the parameter shift $\Delta\theta_t$ is near zero; for $\hat{\gamma}_t \geq \tau$, the controller is allowed to apply stronger corrective actions and invoke explicit error-mitigation routines.

D. Error-Mitigation Operator

The “Error Mitigation” block in Figure-1 is applied to the noisy measurement data before being returned to the AI/ML model and optimization loop. Conceptually, an error-mitigation map can be written as

$$\mathcal{M}_{\gamma_t}: \rho_t \mapsto \rho_t^{(\text{mitigated})},$$



where ρ_t is the noisy state produced by the quantum circuit. In practice, this may correspond to zero-noise extrapolation, probabilistic error cancellation, symmetry verification, or measurement error mitigation.

At the level of expectation values, mitigation yields

$$\hat{\mathbf{E}}(\theta_t) = \mathcal{M}_{\gamma_t}(\tilde{\mathbf{E}}(\theta_t, \gamma_t)),$$

which provides a refined energy estimate used both as the optimization target and as input to the reward function r_t .

When $d_t = 0$ (low-noise regime), $\mathcal{M}_{\gamma_t}$ may reduce to the identity map, whereas for $d_t = 1$ (high-noise regime), a more aggressive mitigation strategy is selected, often at the cost of additional sampling overhead.

E. Closed-Loop Operation

Putting these components together, each iteration of the framework proceeds as follows:

- a) **Quantum execution:** Given θ_t , the quantum circuit runs on the NISQ hardware and produces measurement outcomes.
- b) **Error mitigation:** The raw outcomes are processed by $\mathcal{M}_{\gamma_t}$ to obtain mitigated observables and the energy estimate $\hat{E}(\theta_t)$.
- c) **Feedback and noise estimation:** Hardware statistics are fed to the AI/ML model, which outputs $\hat{\gamma}_t$; the decision logic determines d_t and the mitigation mode.
- d) **Adaptive optimization:** The state $\mathbf{s}_t$ is constructed from $\hat{E}(\theta_t)$, gradient information, $\hat{\gamma}_t$, and the time index. The policy π_ϕ generates an update $\Delta\theta_t$, and the parameters are updated as

$$\theta_{t+1} = \theta_t + \Delta\theta_t.$$

- e) Learning:** The reward r_t is computed, and the policy parameters ϕ (and, if applicable, ψ and τ) are updated using standard RL or supervised-learning rules.

Through this closed-loop interaction between quantum hardware, adaptive optimization, error mitigation, and the AI/ML model, the framework continuously steers the system toward lower-energy states while explicitly accounting for NISQ noise and hardware drift.




Figure-1. Hybrid quantum-classical adaptive intelligence framework of NISQ system real-time optimisation.

3.1 Experimental Setup

The Hybrid Quantum-Classical Adaptive Intelligence Framework was tested on a simulated NISQ space and a classical reinforcement learning controller. This subsection contains a summary of the software/hardware platform, quantum problem, and key hyperparameters of the Adaptive V16 controller and the fixed VQE baseline.

A. Software and Hardware Environment

All the experiments were implemented in Python and executed through a stack (e.g., PennyLane/Qiskit-style variational interfaces) of hybrid quantum-classical simulations to construct parameterized quantum circuits and calculate expectation values. Regular deep learning packages were used to teach the traditional learning ingredients (reinforcement learning agent and auxiliary predictors). The simulations were executed on a workstation-type CPU platform (multi-core processor with a minimum of 16 GB RAM). The size of the problem used did not require any resources on a GPU at its maximum.

B. Problem Instance and Ansatz

The experiment is quantum-implemented as a parameterized variational circuit, which is operating in a VQE-like way:

Target Hamiltonian: Small-scale simulation of a NISQ-sized problem (e.g., few-spin / few-orbital system) as a Target Hamiltonian model.

Number of logical qubits (nq): (Sometimes also denoted nq): The active space of the testing instance (e.g., test Hamiltonian) is being represented by nq:3-4 qubits.

Ansatz structure Hardware-efficient layered ansatz
Hardware-efficient layered ansatz Hardware-efficient
layered ansatz Hardware-efficient layered ansatz
Hardware-efficient layered ansatz Hardware-efficient
layered ansatz Hardware-efficient layered ansatz
Hardware-efficient layered ansatz Hardware-efficient
Hardware-efficient Hardware-efficient Hardware-efficient
Hardware-efficient Hardware-efficient Hardware-efficient



efficient Hardware-efficient Hardware-efficient Hardware-efficient Hardware-efficient Hardware.

Circuit depth (L): Repetition of layers (e.g. $L=3-4$) also fixed to emulate shallow circuits as used by NISQ.

Measurement shots: The measurement circuit uses a finite number of shots (e.g., 1024-4096) to resolve the sampling noise. The VQE base is an attempt to find the minimum imaging statistics using the same ansatz with a generic classical optimizer (gradient-free or gradient-based), and Adaptive V16 is an attempt to train the control policy with the same ansatz.

C. State of observables and Control Interface

The Adaptive V16 controller observes a small state vector of the hybrid loop and the noise model, as shown in the interface of the V16 Dynamic Control Deck: current variational energy - $E_t(H_a)$. gradient norm $\| \nabla E_t \|^2$ - norm of the gradient of the parameter or proxy measure of the steepness of the landscape.

Statistically based on device or simulator: approximate noise strength g_t - noise probability signal on $[0, 1]$.

ABC Wall-clock proxy: Time step index t - normalized iteration counter.

These observables suggest that the controller generates a parameter shift vector D tht. At low-noise regimes (g_t below the learned threshold), the controller remains in a passive monitoring position and proposes near-zero updates, as indicated by the snapshot of the prediction value and g_t exceeds the threshold, corrective actions are taken.

D. Knowledge of Reinforcement and Parameters of the decisions

The above features of states will be applied to the adaptive controller, which is a deep reinforcement agent. Key parameters include:

DRA Value-based or policy-gradient (e.g., DQN/actor-critic).

Discount factor γ "RL" 0.95-0.99 This aspect is used to equalize the short-term loss of energy and long-term convergence.

learning rate ϵ : Training that is stable at 10^{-3} - 10^{-4} .

Replay/batch size: 32-128 transitions/update (off-policy methods).

Exploration time: ϵ -greedy or entropy-regulated policy (with epsilon decayed from some initial value (e.g., 1.0) to some small floor (e.g., 0.05)).

Training horizon: a number of hundreds to several thousand environment steps/episode, and a number of these experiments in each experimental run.

The decision on whether passive monitoring or mitigation should be applied is made in the classification layer depending on the noise signal and auxiliary features. Phase: The threshold on the probability signal g is learned as a scalar; in the given configuration, this is learned on g_t 0.52-0.53 as shown in the visualization of the control deck (blue dashed line).

E. Parameter Summary Table

**Table-2.** Experimental parameters for hybrid quantum-classical evaluation.

Category	Parameter	Symbol	Value / Setting
Quantum model	Number of qubits	n_q	3-4 (benchmark instance)
	Circuit depth (layers)	L	3-4 hardware-efficient layers
	Shots per expectation evaluation	-	1024-4096
	Target Hamiltonian	-	Small molecular / spin Hamiltonian (VQE-style)
Classical baseline	Optimizer type	-	Static VQE optimizer (e.g., gradient-free)
	Max iterations	-	Fixed budget matching Adaptive V16
RL controller	Discount factor	γ_{RL}	0.95-0.99
	Learning rate	η	10^{-3} – 10^{-4}
	Batch size / replay size	-	32-128
	Episodes	-	Multiple episodes (hundreds-thousands of steps each)
Decision logic	Noise probability range	γ_t	[0,1]
	Learned decision threshold	τ	≈ 0.526 (from control deck)
	Actions	$\Delta\theta_t$	Parameter shifts; 0 for passive monitoring
Metrics	Optimization objective	-	Ground-state energy error
	Controller metrics	-	Accuracy, precision, recall, F1, ROC–AUC

4. RESULTS AND PERFORMANCE EVALUATION

This section reports the empirical performance of the proposed Hybrid Quantum-Classical Adaptive Intelligence Framework, hereafter referred to as Adaptive V16, and compares it with a conventional static VQE baseline. Results are obtained using the PennyLane-based implementation described in the methodology, applied to a noisy NISQ-style molecular Hamiltonian constructed via `qml.qchem.molecular_hamiltonian` with a three-hydrogen configuration. The same ansatz and initial state are used for both methods; only the optimization and feedback strategy differ.

A. Ground-State Energy Optimization

The primary objective is to approximate the ground-state energy of the target Hamiltonian as closely as possible under realistic noise. Table-3 summarizes the

final energies and absolute errors achieved by the two approaches.

Table-3. Ground-state energy comparison (Static VQE vs Adaptive V16).

Metric	Static VQE	Adaptive V16
Final Energy	-1.1349	-1.2083
Abs Error	0.1391	0.0657

Relative to the static VQE, the adaptive framework reduces the absolute energy error from 0.1391 to 0.0657, corresponding to an error reduction of approximately 52.8%. This demonstrates that the reinforcement learning (RL)–driven optimizer is substantially more effective at navigating the noisy cost landscape and avoiding poor local minima.




Figure-2. Energy optimization.

The energy trajectories over optimization steps are visualized in Figure-2. The static VQE trace oscillates around a suboptimal plateau, while the Adaptive V16 curve shows a more pronounced downward trend and converges closer to the reference ground-state energy level (indicated by the horizontal target line). Despite stochastic fluctuations due to NISQ noise, the adaptive agent is able to continuously refine its update policy and push the energy toward the desired minimum.

B. Noise Signal Tracking and Adaptive Thresholding

A key component of the proposed framework is the ability to *sense* and *respond* to time-varying noise conditions. The code models a latent “noise probability signal” γ , and the V16 agent attempts to track this signal and select operating points accordingly.




Figure-3. Noise signal tracking.

Figure-3 shows the noise signal tracking behaviour: the ground-truth noise trajectory is plotted against the V16 prediction $\hat{\gamma}$. The two curves remain closely aligned across the evaluation horizon, demonstrating that the learned model can follow slow

drifts as well as moderate fluctuations in the underlying noise profile. This tracking capability is essential for making robust decisions about step sizes, parameter updates, and mitigation actions.




Figure-4. Metric sensitive analysis.

To quantify how the prediction threshold influences decision quality, a threshold-sensitivity analysis is reported in Figure-4. Here, accuracy, precision, recall, and F1-score are plotted as functions of the decision threshold applied to $\hat{y}$. The curves show that performance remains high over a reasonably wide threshold range, with a clear optimum region where all metrics are simultaneously near their peak values. The operating threshold used by the framework is chosen by maximizing the accuracy curve, which is highlighted in the figure. This analysis confirms that the decision rule is not overly sensitive to fine-tuning and that the chosen threshold is near-optimal.

C. Prediction Quality and ROC Analysis

The prediction module inside the adaptive controller is evaluated using standard binary classification metrics, summarized in Table-4 and visualized as a bar chart in Figure-5.

Table-4. Final prediction metrics (v16 controller).

Accuracy	0.9700
Precision	1.0000
Recall	0.9221
F1 Score	0.9595
AUC	0.9951




Figure-5. Final V16 performance metrics (Benchmark).

These results indicate that the controller makes highly reliable decisions:

- An accuracy of 97% shows that almost all decisions (e.g., whether to apply a corrective action or accept an update) are correct.
- Perfect precision (1.0000) implies that whenever the controller predicts a “positive” condition, it is never a false alarm. This is particularly important in a NISQ setting, where unnecessary mitigation steps or extra quantum evaluations can be expensive.



- A recall of 0.9221 demonstrates that the controller successfully captures the vast majority of truly positive events, with relatively few missed opportunities.
- The F1-score of 0.9595 reflects a strong balance between precision and recall, confirming that the decision policy does not overfit to a single metric.




Figure-6. ROC curve.

The ROC curve in Figure-6 further illustrates the discriminative power of the V16 predictor. The curve lies close to the upper-left corner of the unit square, with an area under the curve (AUC) of 0.9951, indicating near-perfect separability between the two classes across all thresholds. The selected operating point, derived from the threshold-sensitivity study, is highlighted on the ROC curve and lies in a region that simultaneously yields high true positive rate and low false positive rate.

D. Qualitative Decision Analysis

In addition to aggregate metrics, the notebook generates per-instance visualizations where the controller's noise estimate and proposed correction are displayed side by side (e.g., horizontal bar for γ vs

threshold and bar chart for parameter shift $\Delta\theta$). These plots (included as additional figures in the results document) show that:

When γ is clearly below the critical region, the agent correctly classifies the system as “stable” and proposes minimal or no corrective update.

When γ enters a high-noise regime, the agent's decision flips to a “mitigation” state and a non-zero $\Delta\theta$ is proposed, consistent with the learned control policy.

Such visual diagnostics support the interpretability of the V16 controller: its actions are not arbitrary but align with the underlying noise signal and the optimization objective.




Figure-7. Prediction results.

In addition to the aggregate metrics, Figure-7 presents a qualitative snapshot of the V16 Dynamic Control Deck for a representative low-noise operating point. The left panel shows the estimated noise probability signal γ (green bar) together with the learned decision threshold (blue dashed line, value ≈ 0.526). In this

instance, the estimated noise level remains below the threshold, and the controller therefore assigns the state “PASSIVE MONITORING,” indicating that the system is operating in a stable regime where no corrective action is required.



The right panel of Figure-7 depicts the corresponding proposed correction in terms of the parameter shift $\Delta\theta$. Consistent with the “PASSIVE MONITORING” decision, the controller outputs a zero (or near-zero) parameter shift, meaning that the quantum circuit parameters are left unchanged. This example illustrates that the V16 agent does not trigger unnecessary mitigation when the inferred noise remains within acceptable bounds, which aligns with the high precision and low false-positive rate reported in Table II and the ROC analysis. Consequently, the qualitative behavior of the controller confirms that the learned policy is both noise-aware and conservative, intervening only when justified by the estimated noise conditions.

5. CONCLUSIONS

This paper presented a Hybrid Quantum-Classical Adaptive Intelligence Framework designed to improve the robustness and optimization performance of variational quantum algorithms operating on Noisy Intermediate-Scale Quantum (NISQ) devices. Unlike conventional hybrid algorithms that rely on static classical optimizers, the proposed framework integrates reinforcement learning-based adaptive optimization, machine learning-driven noise estimation, and online error mitigation into a unified closed-loop control architecture. This adaptive structure enables the system to continuously monitor hardware feedback, estimate noise conditions, and dynamically adjust circuit parameters and mitigation strategies during algorithm execution.

Experimental evaluation using a simulated NISQ molecular Hamiltonian demonstrated that the proposed Adaptive V16 controller significantly improves optimization accuracy compared with a traditional static VQE baseline. Specifically, the framework reduced the ground-state energy error by approximately 52.8%, indicating that reinforcement learning can effectively navigate noisy quantum optimization landscapes. Additionally, the predictive decision module achieved high reliability with 97% accuracy, 100% precision, 92% recall, 95.9% F1-score, and 0.995 AUC, confirming that the controller can accurately detect noise conditions and trigger appropriate mitigation actions. These results show that incorporating intelligent adaptive feedback mechanisms into hybrid quantum-classical workflows can substantially enhance convergence stability and computational fidelity.

Overall, the proposed framework demonstrates that AI-driven adaptive optimization is a promising approach for improving the practicality of quantum algorithms in the NISQ era. By combining real-time hardware monitoring, reinforcement learning-based control, and machine learning-assisted error mitigation, the architecture provides a scalable template for future hybrid quantum systems. Such adaptive frameworks can play an important role in enabling reliable quantum computation on near-term hardware and accelerating the

transition from experimental NISQ devices toward practical quantum advantage.

Future work will focus on extending the framework to larger quantum circuits, real quantum hardware implementations, and multiple hybrid algorithms such as QAOA and quantum machine learning models. Furthermore, integrating advanced deep reinforcement learning strategies, physics-informed learning methods, and hardware-aware ansatz design could further improve scalability, stability, and performance in next-generation hybrid quantum computing systems.

REFERENCES

- [1] J. Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum*. 2: 79.
- [2] E. Farhi, J. Goldstone and S. Gutmann. 2014. A Quantum Approximate Optimization Algorithm. *arXiv*: 1411.4028.
- [3] A. Peruzzo et al. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*. 5(1): 1-7.
- [4] M. Cerezo et al. 2012. Variational Quantum Algorithms. *Nature Reviews Physics*. 3(9): 625-644.
- [5] A. Kandala et al. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*. 549: 242-246.
- [6] S. McArdle et al. 2020. Quantum computational chemistry. *Reviews of Modern Physics*. 92: 015003.
- [7] K. M. Nakanishi, K. Fujii and S. Todo. 2020. Sequential minimal optimization for quantum-classical hybrid algorithms. *Physical Review Research*. 2(4): 043158.
- [8] Y. Du et al. 2021. Quantum neural network architectures. *npj Quantum Information*. 7(23): 1-17.
- [9] T. Fösel, P. J. Karalekas, S. Krastanov and F. Marquardt. 2019. Reinforcement learning with neural networks for quantum feedback. *npj Quantum Information*. 5: 20.
- [10] P. V. Klimov et al. 2018. Fluctuations in superconducting qubit relaxation times. *Physical Review Letters*. 121(9): 090502.
- [11] M. Benedetti, E. Lloyd, S. Sack and M. Fiorentini. 2019. Machine learning for variational quantum algorithms. *npj Quantum Information*. 5(45): 1-9.



- [12] P. Schindler et al. 2020. Error mitigation using symmetry verification. *Nature Physics*. 16: 151-155.
- [13] L. Vassallo. 2023. Machine Learning-Based Quantum Error Mitigation. *arXiv*: 2302.01410.
- [14] A. Cemin et al. 2020. Learning reduced quantum channels from experimental data. *Physical Review Applied*. 13(4): 044005.
- [15] P. Czarnik, A. Arrasmith, P. J. Coles and L. Cincio. 2021. Error mitigation with Clifford data regression. *Nature Communications*. 12: 3639.
- [16] S. Khan, D. K. Park and J. Kim. 2022. Measurement error mitigation techniques for IBM NISQ devices. *Quantum Science and Technology*. 7(3): 035014.
- [17] X. Xie et al. 2021. Advances in quantum computing in the NISQ era. *IEEE Transactions on Quantum Engineering*. 2: 1-18.
- [18] Z. Zhang et al. 2021. Reinforcement learning for quantum control: A comparison study. *Machine Learning: Science and Technology*. 2(3): 035005.
- [19] S. Sivak et al. 2022. Model-free quantum control with reinforcement learning. *Nature Physics*. 18: 703-708.
- [20] M. Porotti et al. 2022. Deep reinforcement learning for quantum state preparation and stabilization. *PRX Quantum*. 3(3): 030329.