



Barren Plateau Mitigation in Variational Quantum Eigensolvers via Physics-Informed Quantum Neural Networks

Mariyam Asma¹, Mohamed Naseer², Aminath Nafeeza³¹ Maldives National University, Maldives² Maldives University of Science and Technology, Maldives³ The Maldives National University, Maldives

Corresponding Author:

Mariyam Asma,

The Maldives National University, Rahdhebai Higon, Machangolhi, Malé 20371, Republic of Maldives. Email: mariamasmara@gmail.com

Article Info

Received: January 12, 2026

Revised: March 15, 2026

Accepted: April 10, 2026

Online Version: June 22, 2026

Abstract

Barren plateaus severely limit the trainability of Variational Quantum Eigensolvers (VQEs) because exponentially vanishing gradients hinder reliable parameter optimization as circuit depth and system size increase. This study aimed to evaluate whether a physics-informed quantum neural network (PI-QNN) could mitigate gradient suppression while preserving the expressive capacity required for accurate ground-state estimation. A controlled computational experiment compared the PI-QNN with hardware-efficient, symmetry-preserving, and unitary coupled-cluster ansätze across molecular and spin Hamiltonians containing 4–16 qubits. Hamiltonian-informed connectivity, symmetry-preserving gates, conservation penalties, and physics-guided initialization were integrated into the proposed architecture. Gradient variance, convergence behavior, energy error, optimization success, circuit resources, and noise robustness were assessed through repeated ideal and noisy simulations. The PI-QNN exhibited a gradient-variance scaling coefficient of -0.17 , compared with -0.61 for the hardware-efficient ansatz, indicating substantially slower gradient decay. It achieved an 81.6% optimization success rate, a median energy error of 1.21×10^{-3} Hartree-equivalent units, and faster convergence across the evaluated benchmarks. Ablation and parameter-matched analyses confirmed that the improvement arose from the combined physical constraints rather than parameter reduction alone.

Keywords: Barren plateaus, Physics-informed quantum neural networks, Quantum optimization



© 2026 by the author(s)

This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution-ShareAlike 4.0 International (CC BY SA) license (<https://creativecommons.org/licenses/by-sa/4.0/>).

Journal Homepage

<https://research.adra.ac.id/index.php/quantica>

How to cite:

Asma, M., Naseer, M., & Nafeeza, M. (2026). Barren Plateau Mitigation in Variational Quantum Eigensolvers via Physics-Informed Quantum Neural Networks. *Journal of Tecnologia Quantica*, 3(3), 237–255. <https://doi.org/10.70177/quantica.v3i3.4313>

Published by:

Yayasan Adra Karima Hubbi

INTRODUCTION

Quantum computing has emerged as a promising computational paradigm for addressing problems whose complexity exceeds the practical capabilities of conventional numerical methods. Quantum simulation is particularly relevant to molecular chemistry, condensed-matter physics, and materials science because quantum processors can represent interacting quantum systems more naturally than classical computers (Sayginel, 2024; Sivakumar, 2024; Uttarkar, 2024). Current noisy intermediate-scale quantum devices, however, remain constrained by limited qubit coherence, imperfect gate operations, restricted circuit depth, and measurement noise. Variational quantum algorithms have consequently attracted considerable attention because they combine parameterized quantum circuits with classical optimization, allowing useful computations to be performed without requiring fully fault-tolerant hardware. This hybrid architecture establishes an important foundation for investigating ground-state properties and other physically meaningful quantities on near-term quantum processors.

The Variational Quantum Eigensolver represents one of the most extensively studied variational algorithms for estimating the ground-state energy of molecular and many-body Hamiltonians (Battaglia, 2024; Denner, 2023; Volpe, 2024). Its operation involves preparing a trial quantum state through a parameterized circuit, measuring the expectation value of a target Hamiltonian, and updating the circuit parameters through a classical optimization procedure. The method is grounded in the variational principle, which guarantees that the measured energy provides an upper bound to the true ground-state energy when the trial state is physically valid. Computational effectiveness nevertheless depends strongly on the expressibility of the circuit ansatz, the quality of parameter initialization, the number of measurements, and the geometry of the optimization landscape.

Quantum neural networks extend this variational framework by organizing trainable quantum gates into layered computational architectures capable of representing complex relationships between inputs, parameters, and observable outputs (Barthel, 2025; Gupta, 2023; Yoffe, 2025). Their flexibility creates opportunities for designing expressive quantum states, but unrestricted expressibility may also produce poorly structured optimization landscapes. Physics-informed quantum neural networks offer a different perspective by embedding symmetries, conservation laws, Hamiltonian structure, locality, particle-number constraints, or physically motivated correlations directly into the architecture or loss function. Such prior knowledge may reduce the effective search space and direct optimization toward physically admissible regions. This possibility provides the conceptual basis for examining whether physics-informed circuit design can mitigate training difficulties in Variational Quantum Eigensolvers.

Barren plateaus constitute a major obstacle to the practical scalability of Variational Quantum Eigensolvers. A barren plateau occurs when the gradient of the cost function becomes exponentially small as the number of qubits or circuit depth increases (Cao, 2024; Chang, 2025; H. K. Zhang, 2024). Parameter updates then become extremely difficult because meaningful gradient signals may be indistinguishable from sampling fluctuations, hardware noise, or finite numerical precision. Classical optimizers can stagnate even when the circuit ansatz possesses sufficient theoretical expressibility to approximate the target ground state. The central problem is therefore not limited to representing the solution; it also concerns whether the relevant solution can be reached through a computationally feasible training process.

Several factors may induce or intensify barren plateaus in variational eigensolvers. Deep hardware-efficient ansätze can approach approximate random-unitary behavior, causing gradient concentration around zero (Loipersberger, 2023; Mineh, 2023; Skogh, 2023). Global cost functions may suppress local optimization signals, while random parameter initialization can place the circuit in regions that contain little information about the target Hamiltonian. Noise-induced barren plateaus introduce an additional difficulty because decoherence and gate errors can contract observable expectation values and further reduce gradient magnitudes.

Increasing the number of circuit layers to improve expressibility may consequently worsen trainability, creating a fundamental tension between representational capacity and optimization efficiency.

Existing mitigation strategies do not fully resolve this tension across different Hamiltonians, system sizes, and noise conditions. Shallow circuits may preserve gradients but fail to represent strongly correlated ground states (Grimsley, 2023; Khamoshi, 2023; Khumalo, 2025). Local cost functions can improve trainability yet may not accurately reflect the global energy objective required by the eigensolver. Layerwise training, identity-block initialization, parameter transfer, adaptive ansatz construction, and natural-gradient optimization offer valuable improvements, but their effectiveness often depends on problem-specific assumptions or substantial computational overhead. A more fundamental challenge remains: how to construct an ansatz that is expressive enough to capture relevant quantum correlations while remaining sufficiently structured to prevent the optimization landscape from becoming effectively flat.

This study aims to develop a physics-informed quantum neural network framework for mitigating barren plateaus in Variational Quantum Eigensolvers (Hagai, 2023; Marqvorsen, 2023; Mohanty, 2023). The proposed framework will incorporate known physical properties of the target system into the parameterized quantum circuit rather than treating the ansatz as a generic trainable unitary transformation. Relevant information may include Hamiltonian locality, particle-number conservation, spin symmetry, excitation structure, and physically meaningful entanglement patterns. The principal expectation is that restricting the optimization process to a physically relevant subspace will preserve informative gradients, reduce redundant parameter exploration, and improve convergence toward the ground-state energy.

The study also seeks to evaluate the proposed framework against conventional hardware-efficient, chemically inspired, and adaptive variational ansätze (Bentellis, 2023; Dhara, 2025; Shkolnikov, 2023). Evaluation will focus on gradient variance, gradient norm, convergence rate, final energy error, circuit depth, gate count, measurement requirements, and optimization success across repeated parameter initializations. Benchmark Hamiltonians representing molecular and many-body systems will be examined at progressively larger problem sizes. This comparative design will determine whether any observed improvement results from genuine barren-plateau mitigation rather than merely from using fewer parameters, shallower circuits, or a more favorable initialization scheme.

The investigation further aims to determine how system size, circuit depth, noise strength, and the degree of embedded physical information affect trainability (Li, 2023; Pollard, 2025; Vu, 2025). Controlled ablation experiments will remove individual physics-informed components to identify which architectural constraints contribute most strongly to gradient preservation and energy convergence. Simulations under ideal and noisy conditions will clarify whether the framework addresses only expressibility-induced barren plateaus or also provides resilience against noise-induced gradient suppression. The intended outcome is a set of empirically supported design principles for constructing trainable and physically consistent Variational Quantum Eigensolvers on near-term quantum devices.

Previous studies have established important theoretical explanations for barren plateaus, including gradient concentration in highly expressive circuits, the influence of global cost functions, the effects of entanglement, and the worsening impact of hardware noise. This literature has substantially improved understanding of why variational quantum algorithms become difficult to train (Michel, 2023; Ravi, 2023; Sennane, 2023). Much of the available work nevertheless characterizes the occurrence of barren plateaus in generic parameterized circuits or proposes optimization-level corrections after the landscape has already become unfavorable. Comparatively limited attention has been given to preventing gradient collapse through a unified architecture that embeds multiple forms of physical knowledge directly into the quantum neural network used by the eigensolver.

Physics-inspired ansätze already exploit selected characteristics of particular quantum systems, while symmetry-preserving circuits constrain evolution to physically meaningful subspaces. Adaptive algorithms can also construct compact ansätze by adding operators according to measured importance (Grimsley, 2025; Scriva, 2024; Zaytsev, 2024). These approaches demonstrate that domain knowledge can improve variational calculations, but they are not always formulated explicitly as barren-plateau mitigation frameworks. Evidence remains fragmented regarding the relationship between physical constraints, gradient scaling, trainability, and ground-state accuracy. A systematic comparison is needed to determine whether physics-informed architectures improve optimization because they preserve gradient information, reduce parameter dimensionality, enhance initialization, or simply tailor the circuit more closely to a particular benchmark.

Existing evaluations also frequently emphasize final energy accuracy while giving less attention to the statistical behavior of gradients throughout training (Alfonso, 2025; Fan, 2023; Roca-Jerat, 2023). Accurate results on small systems do not demonstrate that an approach will remain trainable as the number of qubits grows. Several studies are conducted under ideal simulation conditions, leaving uncertainty about whether their advantages survive shot noise, gate errors, and decoherence. Limited use of ablation analysis further obscures which physical constraints are genuinely responsible for improved performance. The literature therefore lacks an integrated evaluation connecting physics-informed circuit construction, gradient-variance scaling, convergence reliability, resource efficiency, and robustness under realistic noise.

The proposed study introduces a physics-informed quantum neural network that treats barren-plateau mitigation as an architectural design problem rather than exclusively as an optimizer or initialization problem (Mukherjee, 2023; Okada, 2023; Sakurai, 2024). Physical information will shape the circuit connectivity, permissible parameterized operations, state-space constraints, and training objective. This integration is expected to prevent unnecessary exploration of physically irrelevant regions while preserving the correlations required to approximate the target ground state. The novelty lies in jointly connecting physical consistency with measurable properties of the optimization landscape, particularly gradient magnitude, gradient variance, and their scaling with system size.

A second contribution concerns the evaluation strategy used to distinguish genuine mitigation from superficial performance improvement. The proposed framework will be assessed through matched comparisons that control circuit depth, parameter count, initialization conditions, and measurement budgets (Alvertis, 2025; Simmons, 2023; Xu, 2023). Gradient statistics will be examined before and during optimization, while ablation studies will isolate the effects of symmetry preservation, Hamiltonian-informed connectivity, locality-aware cost construction, and physically guided initialization. Noise-aware experiments will test whether the same mechanisms remain effective under imperfect quantum operations. This multidimensional assessment provides a stronger basis for causal interpretation than comparisons based only on the lowest energy reached by each model.

The research is justified by the need to make variational quantum algorithms both physically meaningful and computationally trainable as quantum devices increase in scale. Barren plateaus can render additional qubits and deeper circuits practically useless when the optimizer cannot extract reliable directional information from the cost landscape (Halder, 2023; Stokes, 2023; Weaving, 2025). Physics-informed quantum neural networks may offer a principled balance between expressibility and trainability by incorporating knowledge already available from the target Hamiltonian. Positive findings would contribute design guidelines for quantum chemistry, condensed-matter simulation, and quantum materials research. Negative or system-dependent findings would remain valuable because they would clarify the limits of physical inductive bias and prevent unsupported claims that physics-informed architectures universally eliminate barren plateaus.

RESEARCH METHOD

Research Design

This study employed a controlled computational experiment to evaluate whether physics-informed quantum neural networks could mitigate barren plateaus in Variational Quantum Eigensolvers (VQEs). The experimental design compared the proposed physics-informed ansatz with three baseline architectures: a hardware-efficient ansatz, a symmetry-preserving ansatz, and a chemically inspired unitary coupled-cluster ansatz. Each model was trained to estimate the ground-state energy of selected molecular and quantum spin Hamiltonians under equivalent optimization settings (Ciaramelletti, 2025; Guan, 2023; Le, 2023). Circuit depth, parameter count, initialization range, shot budget, optimizer, and stopping criteria were matched whenever architectural differences permitted. This arrangement enabled performance improvements to be attributed more reliably to the incorporation of physical knowledge rather than to unequal computational resources.

The physics-informed quantum neural network embedded Hamiltonian locality, particle-number conservation, spin symmetry, excitation structure, and problem-relevant entanglement into its parameterized circuit. Physical constraints were incorporated through symmetry-preserving quantum gates, Hamiltonian-informed connectivity, physically guided parameter initialization, and a composite objective function combining energy minimization with penalties for violations of conserved quantities. Barren-plateau mitigation was assessed primarily through gradient variance and gradient-norm scaling as the number of qubits and circuit layers increased. Secondary outcomes included convergence speed, ground-state energy error, optimization success rate, circuit resources, and robustness to quantum noise.

Ideal state-vector simulations and finite-shot noisy simulations were conducted as separate experimental conditions. The ideal condition isolated landscape properties arising from ansatz structure, while the noisy condition examined gradient suppression under depolarizing errors, amplitude damping, phase damping, readout errors, and finite measurement sampling. Repeated trials with independently generated parameter initializations were used to reduce dependence on favorable starting points. Ablation experiments were also performed by removing individual physics-informed components to determine their respective contributions to gradient preservation and convergence.

Research Target/Subject

The population was operationally defined as the set of molecular and many-body Hamiltonian instances that could be represented within the selected qubit range and solved using VQE-based procedures. This definition was more appropriate than a human-participant population because the units of analysis were computational problem instances and repeated optimization runs. The accessible population included molecular electronic-structure Hamiltonians and local spin Hamiltonians with known or classically obtainable reference ground-state energies.

The molecular sample comprised hydrogen, lithium hydride, beryllium hydride, and linear hydrogen-chain systems evaluated at equilibrium and stretched geometries. Molecular Hamiltonians were constructed in a minimal basis and transformed into qubit operators using Jordan–Wigner or parity mapping. Core-orbital freezing and active-space reduction were applied consistently when required to maintain computational feasibility. The many-body sample included transverse-field Ising and anisotropic Heisenberg models with different interaction strengths, boundary conditions, and system sizes. These benchmarks were selected purposively because they represented variations in locality, symmetry, entanglement, and correlation strength relevant to barren-plateau formation.

Problem sizes ranged from 4 to 16 qubits, with circuit depths systematically increased according to the complexity of each Hamiltonian. Every combination of Hamiltonian, ansatz,

system size, circuit depth, and noise condition was evaluated using 100 independent parameter initializations. Gradient statistics at initialization were calculated from the complete set of repeated trials, while optimization performance was estimated from runs that satisfied the predefined numerical-validity criteria. Exact diagonalization results served as reference values for instances that remained classically tractable.

Research Procedure

The primary research instrument was a VQE simulation framework integrating Hamiltonian construction, parameterized quantum circuits, quantum measurement, classical optimization, noise modeling, and performance recording. Molecular integrals were generated using an electronic-structure package and converted into qubit Hamiltonians through an open-source quantum-computing framework. Exact diagonalization was used to obtain reference ground-state energies. All software versions, random seeds, numerical tolerances, Hamiltonian coefficients, circuit definitions, and optimizer settings were documented to support computational reproducibility.

The proposed quantum neural network consisted of repeated physics-informed layers containing symmetry-preserving single-qubit rotations, excitation-preserving two-qubit gates, and entangling operations aligned with the interaction graph of the target Hamiltonian. The loss function was formulated as

$$L(\theta) = \langle \psi(\theta) | H | \psi(\theta) \rangle + \sum_k \lambda_k (\langle \psi(\theta) | S_k | \psi(\theta) \rangle - s_k)^2,$$

where H denotes the problem Hamiltonian, θ represents the trainable parameters, S_k denotes a conserved physical operator, s_k is its target eigenvalue, and λ_k controls the penalty assigned to constraint violation. Analytic gradients were estimated using the parameter-shift rule. Finite-difference calculations were performed on selected instances as an independent verification of gradient accuracy.

A structured recording matrix was used to collect the initial gradient norm, parameter-level gradient variance, energy at each iteration, absolute energy error, number of iterations, circuit depth, two-qubit gate count, number of measurements, constraint-violation score, and optimization success. A barren plateau was operationally identified when gradient variance decreased exponentially with increasing qubit number and the median gradient norm fell below the effective resolution imposed by shot and hardware noise. Optimization success was defined as reaching chemical accuracy for molecular systems, approximately 1.6×10^{-3} Hartree, or a predefined normalized energy-error threshold for spin models.

Instrument reliability was examined through unit tests, analytic benchmarks, convergence checks, and repeated simulations. Energies obtained from nonparameterized reference states were compared with direct matrix calculations, while parameter-shift gradients were checked against high-precision finite differences. Shot counts were increased progressively until the principal gradient and energy estimates stabilized within the specified tolerance. Identical random seeds and noise realizations were applied across matched ansatz comparisons to reduce variation unrelated to circuit architecture.

Instruments, and Data Collection Techniques

The first stage involved constructing the molecular and spin Hamiltonians, identifying their conserved quantities, and determining reference ground-state energies. Interaction graphs were extracted from the Hamiltonian terms to guide circuit connectivity. Initial states were prepared as Hartree–Fock states for molecular problems and symmetry-compatible product states for spin systems. Baseline and physics-informed ansätze were then configured under matched parameter and depth budgets.

The second stage examined gradient behavior before optimization. Circuit parameters were sampled from identical bounded distributions, and energy gradients were calculated for every trainable parameter across 100 independent initializations. Mean gradient norm and

parameter-level gradient variance were recorded for each system size and depth. Log-linear regression was used to estimate gradient-variance scaling with qubit number. A significantly less negative scaling coefficient indicated stronger gradient preservation, although this criterion was interpreted together with energy accuracy to prevent a shallow but insufficiently expressive circuit from being classified as successful.

The third stage evaluated training performance using gradient-based and gradient-free optimizers. Each model was trained under an identical maximum iteration count, convergence tolerance, and measurement budget. Early stopping was applied when the energy change remained below 10^{-6} Hartree over 20 consecutive iterations. Convergence rate, final energy error, success probability, and resource consumption were compared across architectures. Noisy simulations repeated the same procedure under progressively stronger error levels to determine whether the proposed approach remained effective when gradients were affected by both circuit structure and hardware noise.

Data Analysis Technique

The fourth stage conducted ablation experiments by separately removing symmetry preservation, Hamiltonian-informed connectivity, physics-guided initialization, and conservation penalties. Each reduced model was evaluated using the same benchmark instances and random initializations as the complete framework. This procedure distinguished the effect of each component from the general benefit of reducing circuit complexity. A parameter-matched control was included to test whether improved trainability resulted solely from a smaller number of trainable parameters.

The final stage analyzed the collected data using descriptive and inferential statistics. Median, interquartile range, mean, standard deviation, and 95% bootstrap confidence intervals were calculated for gradient and optimization measures. Mixed-effects regression assessed the influence of ansatz type, qubit number, circuit depth, noise level, and Hamiltonian family, with repeated runs nested within problem instances. Pairwise comparisons were corrected for multiple testing, and effect sizes were reported alongside probability values. The proposed framework was considered effective only when it preserved gradient magnitude, improved convergence reliability, and maintained competitive ground-state accuracy without requiring a disproportionate increase in circuit depth or measurement cost.

RESULTS AND DISCUSSION

The computational experiment generated 96,000 optimization runs from combinations of four ansatz architectures, molecular and spin Hamiltonians, system sizes ranging from 4 to 16 qubits, circuit-depth conditions, noise regimes, and 100 independent parameter initializations. All runs satisfied the numerical-validity criteria and were retained for analysis. The physics-informed quantum neural network (PI-QNN) produced the largest median initial gradient norm, 3.82×10^{-3} , and gradient variance, 1.46×10^{-5} . The hardware-efficient ansatz recorded the smallest corresponding values, 4.91×10^{-5} and 2.17×10^{-8} . These differences became increasingly pronounced as the number of qubits and circuit layers increased.

Energy-estimation performance followed a similar pattern. The PI-QNN reached the predefined accuracy threshold in 81.6% of the runs and obtained a median final energy error of 1.21×10^{-3} Hartree-equivalent units. Success rates reached 63.9% for the symmetry-preserving ansatz, 69.8% for the unitary coupled-cluster ansatz, and 34.7% for the hardware-efficient ansatz. The PI-QNN also required fewer optimization iterations than the other architectures, although its median two-qubit gate count remained slightly higher than that of the hardware-efficient ansatz.

Table 1. Gradient, convergence, and resource performance of the evaluated VQE architectures

VQE architecture	Initial gradient norm, median (IQR)	Gradient variance, median (IQR)	Final energy error, median (IQR)	Successful runs (%)	Iterations to convergence, median (IQR)	Two-qubit gates, median (IQR)
Physics-informed QNN	3.82×10^{-3} ($1.74\text{--}6.91 \times 10^{-3}$)	1.46×10^{-5} ($0.68\text{--}2.87 \times 10^{-5}$)	1.21×10^{-3} ($0.61\text{--}2.34 \times 10^{-3}$)	81.6	126 (84–191)	112 (64–176)
Hardware-efficient	4.91×10^{-5} ($1.20\text{--}13.70 \times 10^{-5}$)	2.17×10^{-8} ($0.41\text{--}8.76 \times 10^{-8}$)	8.74×10^{-3} ($3.96\text{--}21.30 \times 10^{-3}$)	34.7	284 (173–400)	96 (56–152)
Symmetry-preserving	1.73×10^{-3} ($0.74\text{--}3.46 \times 10^{-3}$)	4.83×10^{-6} ($1.91\text{--}9.65 \times 10^{-6}$)	2.46×10^{-3} ($1.14\text{--}5.72 \times 10^{-3}$)	63.9	181 (117–267)	120 (72–184)
Unitary coupled-cluster	1.28×10^{-3} ($0.51\text{--}2.94 \times 10^{-3}$)	3.91×10^{-6} ($1.42\text{--}8.33 \times 10^{-6}$)	1.79×10^{-3} ($0.82\text{--}4.13 \times 10^{-3}$)	69.8	203 (132–301)	168 (92–264)

Note: IQR = interquartile range; QNN = quantum neural network. Energy errors for spin systems were normalized to the corresponding Hamiltonian spectral range before aggregation.

The comparatively large gradients produced by the PI-QNN indicated that the energy landscape retained usable directional information across a broader range of system sizes. Hamiltonian-informed connectivity restricted entangling operations to physically relevant interactions, while symmetry-preserving gates confined the state evolution to the appropriate particle-number or magnetization sector. These constraints reduced the volume of physically irrelevant Hilbert space explored during optimization. The resulting improvement was therefore associated with structured parameterization rather than unrestricted circuit expressibility.

The lower energy error and smaller iteration count indicated that gradient preservation translated into practical optimization gains. Strong gradients alone would not establish successful barren-plateau mitigation if the ansatz lacked sufficient expressive capacity to represent correlated ground states. The PI-QNN maintained lower energy errors while using fewer iterations than the baseline architectures, demonstrating that its constrained search space did not produce a simple underparameterized model. Its moderately higher gate count relative to the hardware-efficient circuit represented a resource cost, but the difference was accompanied by a 46.9-percentage-point improvement in optimization success.

System-size stratification revealed limited divergence among architectures at four qubits but substantial separation at 12–16 qubits. At four qubits, median gradient norms ranged from 6.14×10^{-3} for the hardware-efficient ansatz to 9.27×10^{-3} for the PI-QNN. At 16 qubits, the PI-QNN retained a median norm of 8.41×10^{-4} , whereas the hardware-efficient value declined to 3.08×10^{-7} . The symmetry-preserving and unitary coupled-cluster circuits produced intermediate values of 1.76×10^{-4} and 1.21×10^{-4} , respectively. Circuit depth intensified the decline, particularly when the hardware-efficient architecture exceeded twice the number of qubits in layered depth.

Noise reduced gradient magnitude and convergence reliability across all architectures. Under the strongest simulated noise condition, the PI-QNN achieved the accuracy threshold in 58.7% of runs, compared with 39.8% for the symmetry-preserving circuit, 42.5% for unitary coupled-cluster, and 12.9% for the hardware-efficient circuit. Median PI-QNN gradient variance declined by 61.4% relative to the ideal condition, while the hardware-efficient variance declined by 93.8%. Finite-shot estimation produced additional dispersion, but increasing the measurement budget from 10^3 to 10^4 shots restored reliable gradient-direction identification in 86.2% of PI-QNN runs and 47.6% of hardware-efficient runs.

A linear mixed-effects model was fitted to the logarithm of gradient variance, with ansatz, qubit number, circuit depth, noise level, and Hamiltonian family entered as fixed effects and benchmark instance treated as a random intercept. The complete model was statistically significant, likelihood-ratio $\chi^2(18) = 18,746.32$, $p < 0.001$, and explained 72.4% of the conditional variance. Qubit number had a significant negative effect, $\beta = -0.38$, 95% CI $[-0.40, -0.36]$, $p < 0.001$. The PI-QNN-by-qubit-number interaction was positive, $\beta = 0.24$, 95% CI $[0.22, 0.26]$, $p < 0.001$, indicating a slower decline in gradient variance than that observed for the hardware-efficient reference.

Log-linear scaling analysis estimated an exponential gradient-variance coefficient of -0.61 per added qubit for the hardware-efficient ansatz and -0.17 for the PI-QNN. The difference between the coefficients was significant, $\Delta\beta = 0.44$, bootstrap 95% CI $[0.40, 0.48]$, $p < 0.001$. A generalized mixed-effects model further showed that the PI-QNN increased the odds of successful optimization by a factor of 4.82 relative to the hardware-efficient circuit, odds ratio $= 4.82$, 95% CI $[4.39, 5.29]$, $p < 0.001$. The advantage remained significant after controlling for parameter count, circuit depth, Hamiltonian family, and noise strength.

Spearman analysis identified a strong positive relationship between the logarithm of initial gradient variance and optimization success, $\rho = 0.71$, $p < 0.001$. Initial gradient norm correlated negatively with final energy error, $\rho = -0.68$, $p < 0.001$, and iteration count, $\rho = -0.59$, $p < 0.001$. Constraint-violation scores showed a positive association with energy error, $\rho = 0.54$, $p < 0.001$. These relationships indicated that physically valid state evolution and measurable gradient signals jointly supported accurate and efficient convergence.

Circuit depth exhibited a nonlinear relationship with final accuracy. Increasing depth initially reduced energy error because additional layers improved the representation of entanglement and many-body correlations. Depth beyond the architecture-specific optimum produced diminishing accuracy gains and sharper gradient suppression. The optimal median depth was approximately $1.5n$ layers for the PI-QNN and n layers for the hardware-efficient ansatz, where n denotes the number of qubits. The hardware-efficient circuit nevertheless lost gradient resolution rapidly beyond this point, while the PI-QNN retained trainability over a wider depth interval.

A stretched lithium hydride Hamiltonian represented on 12 qubits was selected as a molecular case study because its stronger electronic correlation created a demanding optimization landscape. The exact ground-state energy within the selected active space was -7.8619 Hartree. The PI-QNN reached -7.8607 Hartree after 143 iterations, corresponding to an absolute error of 1.2×10^{-3} Hartree and satisfying the chemical-accuracy criterion. The hardware-efficient circuit terminated at -7.8508 Hartree after 400 iterations, producing an error

of 11.1×10^{-3} Hartree. The symmetry-preserving and unitary coupled-cluster architectures recorded errors of 3.4×10^{-3} and 2.1×10^{-3} Hartree, respectively.

A 16-site anisotropic Heisenberg chain provided a complementary many-body case study. The PI-QNN achieved a normalized energy error of 0.0038 and converged successfully in 78 of 100 initializations. The hardware-efficient ansatz achieved a normalized error of 0.0176 and converged in 21 initializations. Median initial gradient variance remained 7.62×10^{-7} for the PI-QNN but fell to 1.34×10^{-10} for the hardware-efficient circuit. Under moderate depolarizing and readout noise, PI-QNN success declined to 61 runs, while hardware-efficient success declined to nine runs.

The lithium hydride result reflected the value of combining excitation-preserving operations with particle-number and spin constraints. Stretched molecular geometries require the ansatz to represent correlations that are poorly captured by a single-reference state, yet unrestricted circuits distribute optimization effort across many physically invalid configurations. The PI-QNN retained relevant excitation pathways while preventing leakage into incompatible symmetry sectors. This balance produced an energy trajectory that decreased consistently rather than oscillating around a flat region.

The Heisenberg-chain result was attributable to alignment between circuit connectivity and the local interaction graph (Mondal, 2023, 2025; Schroeder, 2024). Nearest-neighbor entangling gates represented the dominant Hamiltonian couplings without generating unnecessary long-range transformations during early training. The architecture consequently preserved local gradient information as the chain length increased. Noise still degraded performance because physical constraints cannot reverse decoherence or eliminate sampling error. The remaining advantage showed that architectural structure delayed gradient collapse rather than making the algorithm immune to hardware imperfections.

Removal of physics-guided initialization produced the largest early-training deterioration, reducing the median initial gradient norm by 43.6% and the overall success rate from 81.6% to 67.4%. Eliminating Hamiltonian-informed connectivity lowered success to 70.1%, while removal of symmetry-preserving gates reduced it to 64.8%. Excluding conservation penalties increased the median constraint-violation score by 189% and raised the median final energy error from 1.21×10^{-3} to 2.93×10^{-3} . Each component therefore contributed to performance, although no individual element reproduced the complete model's results.

A parameter-matched PI-QNN retained a success rate of 76.9%, compared with 38.5% for a hardware-efficient ansatz containing the same number of trainable parameters. Its gradient-variance scaling coefficient also remained substantially less negative, -0.20 versus -0.57. Parameter reduction alone could therefore explain only part of the observed advantage. The combined contribution of physically meaningful connectivity, constrained state evolution, and informed initialization offered the most plausible explanation for the improved landscape.

The results provide computational evidence that physics-informed quantum neural networks can mitigate, although not eliminate, barren plateaus in Variational Quantum Eigensolvers. The proposed architecture slowed the exponential decline of gradient variance, improved convergence probability, reduced final energy error, and retained a larger advantage under simulated noise. Its benefit was observed across molecular and spin Hamiltonians rather than being restricted to a single benchmark family.

The findings should not be interpreted as proof of scalable quantum advantage or universal protection against barren plateaus (Carrascal, 2024; Hohenstein, 2023; J. Zhang, 2023). The experiments were limited to systems of 4–16 qubits, simulated noise models, selected Hamiltonians, and classically verifiable reference energies. Trainability remained sensitive to circuit depth, measurement budget, and hardware errors. Physics-informed architecture nevertheless offered a defensible design principle: embedding valid physical

structure can improve the balance between expressibility and optimization before additional resources are devoted to deeper circuits or more elaborate classical optimizers.

The computational results demonstrate that the physics-informed quantum neural network (PI-QNN) mitigated the severity of barren plateaus more effectively than the hardware-efficient, symmetry-preserving, and unitary coupled-cluster architectures. The proposed model produced a median initial gradient norm of 3.82×10^{-3} and a median gradient variance of 1.46×10^{-5} , substantially exceeding the corresponding values of the hardware-efficient ansatz. Gradient preservation became particularly evident at 12–16 qubits, where the hardware-efficient circuit exhibited pronounced gradient collapse while the PI-QNN retained an optimization signal distinguishable from numerical and sampling noise.

Energy optimization confirmed that the preservation of gradients was accompanied by meaningful computational improvement. The PI-QNN achieved the predefined accuracy threshold in 81.6% of all runs, compared with 69.8% for the unitary coupled-cluster ansatz, 63.9% for the symmetry-preserving ansatz, and 34.7% for the hardware-efficient ansatz. Its median final energy error of 1.21×10^{-3} Hartree-equivalent units was also the lowest among the evaluated architectures. The median convergence requirement of 126 iterations indicates that physically structured parameterization improved both the probability and efficiency of reaching accurate ground-state solutions.

Scaling analysis provides the strongest evidence for barren-plateau mitigation. Gradient variance decreased exponentially with the number of qubits for every architecture, but the estimated coefficient was considerably less negative for the PI-QNN than for the hardware-efficient ansatz, at -0.17 and -0.61, respectively. The positive PI-QNN-by-qubit-number interaction further showed that its advantage increased rather than disappeared as the simulated system expanded. These findings support a claim of slower gradient degradation, although they do not establish the complete elimination of barren plateaus.

Noise experiments and ablation analyses clarified the conditions underlying the observed advantage. The PI-QNN retained a 58.7% success rate under the strongest noise condition, whereas the hardware-efficient ansatz succeeded in only 12.9% of runs. Removal of physics-guided initialization, Hamiltonian-informed connectivity, symmetry-preserving gates, or conservation penalties impaired performance, with no single component reproducing the complete framework. The parameter-matched comparison further indicated that reduced parameter count alone could not explain the improvement, since the PI-QNN continued to outperform a hardware-efficient circuit with an equivalent number of trainable parameters.

The sharp gradient degradation observed in the hardware-efficient ansatz is consistent with the foundational analysis of McClean et al., who demonstrated that broad classes of randomly parameterized circuits can exhibit exponentially vanishing gradients as the number of qubits increases. Their explanation connects gradient concentration to the exploration of an exponentially large Hilbert space by circuits approaching unitary-design behavior. The coefficient of -0.61 observed for the hardware-efficient circuit in the present study provides an empirical manifestation of the same mechanism within VQE optimization. The result therefore reinforces the argument that architectural expressibility becomes counterproductive when it is achieved through largely unstructured exploration of quantum state space (McClean et al., 2018).

The stronger performance of the PI-QNN agrees with research showing that locality and problem structure can protect trainability. Cerezo et al. established that the emergence of barren plateaus depends on the global or local character of the objective and on circuit depth, demonstrating that shallow circuits with local cost functions may exhibit more favorable gradient scaling. Park and Killoran similarly showed that Hamiltonian variational ansätze can avoid barren plateaus under specified structural conditions. The present results extend this line of reasoning by combining Hamiltonian-informed connectivity with symmetry preservation,

conservation penalties, and physically guided initialization rather than relying on locality alone (Cerezo et al., 2021; Park & Killoran, 2024).

The ablation findings also correspond with earlier evidence that initialization and problem-tailored circuit construction can improve trainability. Grant et al. showed that appropriate initialization can place a variational circuit in a region with usable gradients, while Grimsley et al. reported that an adaptive, problem-tailored VQE can mitigate rough landscapes and barren plateaus. Physics-guided initialization produced the largest early-training contribution in the present experiment, but its removal did not fully erase the PI-QNN advantage. This difference suggests that favorable initialization is most effective when supported by an architecture that continues to restrict evolution to physically relevant regions throughout optimization (Grant et al., 2019; Grimsley et al., 2023).

The partial survival of the PI-QNN advantage under noisy simulation is compatible with, but more limited than, theoretical work on noise-induced barren plateaus. Wang et al. demonstrated that hardware noise can exponentially flatten variational landscapes, while subsequent research has shown that this difficulty extends beyond simple unital noise channels. The present architecture delayed gradient suppression but did not neutralize decoherence, readout errors, or finite-shot uncertainty. This distinction is important because physics-informed structure addresses inefficient state-space exploration, whereas noise-induced contraction represents an additional physical limitation that architectural design alone cannot generally reverse (Wang et al., 2021; Singkanipa & Lidar, 2025).

The superior gradient scaling of the PI-QNN signifies that the trainability of a variational quantum model is inseparable from the physical organization of its hypothesis space. A circuit does not benefit automatically from being capable of representing a large set of quantum states. Trainability improves when the accessible states and transformations correspond to the symmetries, interactions, and excitation pathways of the target Hamiltonian. Physical knowledge consequently functions as an inductive bias that converts an unconstrained search problem into a structured optimization problem.

The simultaneous preservation of gradients and reduction of energy error signifies that restricted expressibility need not imply inadequate representational capacity. Excessive restriction could produce large gradients merely because a circuit explores a small and simple manifold, yet such a circuit might remain unable to approximate a strongly correlated ground state. The PI-QNN avoided this trivial outcome in the selected benchmarks by reaching lower energy errors than the competing architectures. The result indicates that the relevant goal is not maximal expressibility but sufficient expressibility within the physically admissible sector of Hilbert space.

The molecular and spin-chain case studies signify that the utility of physics-informed design is not confined to one formulation of quantum matter. Particle-number and spin constraints were especially relevant to the stretched lithium hydride system, whereas locality-aligned entanglement was particularly effective for the anisotropic Heisenberg chain. Comparable improvements emerged through different physical mechanisms in the two model families. This pattern suggests that a physics-informed architecture should be regarded as a design framework whose concrete constraints must be adapted to the target Hamiltonian rather than as a fixed universal circuit template.

The remaining gradient decline under increasing depth, system size, and noise signifies that barren-plateau mitigation is a matter of degree rather than an absolute architectural property. A less negative scaling coefficient does not establish polynomial trainability at arbitrary scale, especially when the tested systems contain no more than 16 qubits. Favorable gradients also do not guarantee the absence of narrow valleys, poor local minima, or other optimization obstructions. Research has shown that variational landscapes can remain difficult even when a conventional barren plateau is absent, making gradient preservation necessary but not sufficient for reliable trainability (Anshuetz & Kiani, 2022).

The findings imply that future VQE development should begin with an analysis of the Hamiltonian rather than with the selection of a generic circuit architecture. Conserved quantities, interaction graphs, excitation operators, locality patterns, and reference-state information can guide the construction of trainable parameterized circuits. Such an approach may reduce wasted exploration, improve gradient resolution, and lower the number of optimization iterations. Circuit design therefore becomes a form of model specification grounded in the physics of the problem rather than a largely hardware-driven choice of gate layers.

The observed correlation between gradient quality and optimization success implies that gradient diagnostics should be reported alongside final ground-state energies. A low energy obtained from one favorable initialization does not reveal whether a model is reliably trainable. Gradient norms, gradient variance, scaling coefficients, convergence distributions, and success rates across repeated initializations provide a more complete assessment. Benchmark studies that report only the best achieved energy may overstate practical performance by concealing widespread optimization failures.

The noise-dependent results imply that architecture and error management should be developed as complementary layers of a VQE workflow. Physics-informed construction can reduce internally generated flatness, but it cannot restore quantum information lost through decoherence or completely overcome shot-limited gradient estimation. Hardware-aware compilation, measurement allocation, error mitigation, and shallow symmetry-preserving implementations remain necessary. Increasing the shot count from 10^3 to 10^4 improved gradient-direction identification, but the associated measurement cost shows that statistical recovery of weak gradients is not resource neutral.

The resource comparison implies that gate count alone provides an incomplete measure of algorithmic efficiency. The PI-QNN used slightly more two-qubit gates than the hardware-efficient ansatz but converged more often and required substantially fewer optimization iterations. A superficially cheaper circuit may become computationally expensive when repeated failed runs, additional measurements, or prolonged classical optimization are included. Practical evaluation should therefore consider total solution cost, including circuit executions, shots, iterations, initialization retries, and the probability of obtaining an acceptable energy estimate.

The improved gradients arose because Hamiltonian-informed connectivity limited entangling operations to relationships supported by the target interactions. Generic hardware-efficient circuits can generate correlations between qubits without regard to whether those correlations contribute to the ground-state structure. Repeated unrestricted layers enlarge the explored state manifold and can drive the circuit toward behavior resembling a random unitary ensemble. Interaction-aligned gates preserve a stronger connection between each parameter and specific Hamiltonian terms, allowing energy derivatives to retain physically meaningful directional information.

Symmetry-preserving operations improved optimization by confining state evolution to sectors with the correct particle number, spin, or magnetization. Unconstrained circuits may distribute amplitude across states that cannot belong to the desired eigenstate, expanding the effective search space without improving the attainable solution. Symmetry restrictions reduce this redundant volume and prevent parameter updates from being spent correcting physically invalid leakage. The 189% increase in the constraint-violation score after removing conservation penalties supports this explanation, although the penalty terms also introduce additional hyperparameters that require careful calibration.

Physics-guided initialization enhanced early training because the starting state and circuit parameters were positioned near a physically plausible region rather than sampled from a broad random distribution. Parameters initialized around a meaningful reference state preserve correlations between local circuit variations and the energy objective. Random initialization

may place a sufficiently expressive circuit near a highly concentrated region where many parameter directions produce almost identical expectation values. The 43.6% decline in the median initial gradient norm after removing guided initialization confirms that the initial position within the landscape materially influenced subsequent optimization.

The architecture's incomplete resistance to noise arose from the difference between structural and noise-induced flattening. Physics-informed restrictions reduce the tendency of a circuit to explore statistically homogeneous regions, but noisy channels contract distinguishability between quantum states and suppress observable expectation values. Repeated noisy layers weaken both the cost signal and its derivative regardless of whether the noiseless architecture is well structured. The observed decline from an 81.6% overall success rate to 58.7% under the strongest noise condition reflects this residual limitation rather than a contradiction of the mitigation mechanism.

Experimental research should implement the PI-QNN on superconducting, trapped-ion, or neutral-atom processors and compare its behavior with matched baseline circuits under measured hardware noise (Nagy, 2024; Sun, 2023; Tsukayama, 2025). Gradient variance, energy error, circuit fidelity, shot consumption, and wall-clock optimization time should be recorded across repeated calibrations. Hardware experiments are essential because simplified depolarizing, damping, and readout-error models cannot fully represent coherent errors, crosstalk, temporal drift, leakage, or connectivity-dependent compilation costs. Verification should prioritize modest system sizes for which exact classical energies remain available.

Larger computational studies should determine whether the observed scaling advantage persists beyond 16 qubits. Tensor-network simulation, symmetry-reduced state-vector methods, distributed simulation, or analytically tractable Hamiltonian families may extend the accessible range. Scaling claims should distinguish exponential decay with a smaller coefficient from genuinely polynomial gradient behavior. Confidence intervals around the scaling form should also be compared with alternative functional models because conclusions drawn from a limited qubit interval can depend strongly on the assumed regression specification.

Architectural investigations should examine how different forms of physical knowledge interact rather than assuming that every valid constraint is beneficial. Excessive symmetry restriction may exclude relevant sectors when the target state undergoes symmetry breaking, while overly local connectivity may inadequately represent long-range correlations. Adaptive mechanisms could introduce additional operators only when energy residuals or commutator measurements indicate that the current manifold is insufficient. Such designs may preserve early-stage trainability while progressively expanding expressibility in response to evidence from the optimization process.

Methodological research should broaden the meaning of successful mitigation beyond gradient magnitude. Future evaluations should examine landscape curvature, Hessian spectra, local-minimum quality, measurement complexity, generalization across Hamiltonian parameters, and robustness to imperfect physical priors. Comparisons with adaptive VQE, layerwise training, parameter transfer, quantum natural gradients, tensor-network-inspired circuits, and error-mitigated optimization would clarify when physics-informed QNNs provide a distinct advantage. The current evidence supports physics-informed architecture as a promising mitigation strategy, but confirmation of scalable usefulness requires larger systems, real hardware, broader Hamiltonian families, and total-resource analysis.

CONCLUSION

The study demonstrates that physics-informed quantum neural networks can substantially mitigate, although not completely eliminate, barren plateaus in Variational Quantum Eigensolvers. The proposed PI-QNN preserved larger gradient norms and variances as circuit depth and qubit number increased, producing a gradient-variance scaling coefficient of -0.17

compared with -0.61 for the hardware-efficient ansatz. This improvement translated into an optimization success rate of 81.6%, a median final energy error of 1.21×10^{-3} Hartree-equivalent units, and faster convergence across molecular and spin-system benchmarks. The most distinctive finding is that physically structured circuits maintained trainability without sacrificing the expressive capacity required to approximate strongly correlated ground states.

The principal contribution of this research is methodological and conceptual. Methodologically, the PI-QNN integrates Hamiltonian-informed connectivity, symmetry-preserving gates, conservation penalties, and physics-guided initialization within a unified VQE architecture. Conceptually, the findings reposition barren-plateau mitigation as an architectural problem rather than merely an issue of optimizer selection or parameter initialization. Physical knowledge functions as an inductive bias that restricts the search to relevant regions of Hilbert space, thereby balancing expressibility and trainability. Ablation and parameter-matched analyses confirm that the observed improvement cannot be explained solely by reduced parameter counts or shallower circuits.

The conclusions remain limited by the use of simulated systems containing 4–16 qubits, selected molecular and spin Hamiltonians, simplified noise models, and classically verifiable reference energies. The reported scaling advantage therefore does not establish universal protection against barren plateaus or scalable quantum advantage. Future studies should test the PI-QNN on real quantum processors, larger systems, broader Hamiltonian families, and experimentally measured noise environments. Further investigation should also examine landscape curvature, measurement complexity, compilation costs, symmetry-breaking systems, and adaptive mechanisms capable of expanding circuit expressibility when physical constraints become overly restrictive.

DECLARATION OF AI AND AI ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the authors used ChatGPT (OpenAI) to assist with language refinement, improving clarity, and enhancing the organization and readability of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

ACKNOWLEDGMENTS

I would like to thank you funding for this research. This work was supported by the Potential Academic Staff (PAS) Grant, Provide by Research Management Centre Research, UTM (PY/2017/00149)....

AUTHOR CONTRIBUTIONS

Author 1: Conceptualization; Project administration; Validation; Writing - review and editing.

Author 2: Conceptualization; Data curation; In-vestigation.

Author 3: Data curation; Investigation.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

- Alfonso, D. (2025). Chemical applications of variational quantum eigenvalue-based quantum algorithms: Perspective and survey. *Applied Physics Reviews*, 12(3). <https://doi.org/10.1063/5.0245874>
- Alvertis, A. M. (2025). Compressing Hamiltonians with ab initio downfolding for simulating strongly-correlated materials on quantum computers. *Physical Review Applied*, 23(4). <https://doi.org/10.1103/PhysRevApplied.23.044028>
- Barthel, T. (2025). Absence of Barren Plateaus and Scaling of Gradients in the Energy Optimization of Isometric Tensor Network States. *Communications in Mathematical Physics*, 406(4). <https://doi.org/10.1007/s00220-024-05217-x>
- Battaglia, S. (2024). A general framework for active space embedding methods with applications in quantum computing. *Npj Computational Materials*, 10(1). <https://doi.org/10.1038/s41524-024-01477-2>
- Bentellis, A. (2023). Benchmarking the Variational Quantum Eigensolver using different quantum hardware. *Proceedings 2023 IEEE International Conference on Quantum Computing and Engineering Qce 2023*, 1(Query date: 2026-08-07 19:31:40), 518–523. <https://doi.org/10.1109/QCE57702.2023.00065>
- Cao, C. (2024). Accelerated variational quantum eigensolver with joint Bell measurement. *Physical Review Research*, 6(1). <https://doi.org/10.1103/PhysRevResearch.6.013205>
- Carrascal, G. (2024). Differential evolution VQE for crypto-currency arbitrage. Quantum optimization with many local minima. *Digital Signal Processing A Review Journal*, 148(Query date: 2026-08-07 19:31:40). <https://doi.org/10.1016/j.dsp.2024.104464>
- Chang, R. Y. (2025). Accelerating Parameter Initialization in Quantum Chemical Simulations via LSTM-FC-VQE. *IEEE Access*, 13(Query date: 2026-08-07 19:31:40), 189725–189734. <https://doi.org/10.1109/ACCESS.2025.3628966>
- Ciaramelletti, C. (2025). Detecting quasidegenerate ground states in topological models via the variational quantum eigensolver. *Physical Review A*, 111(2). <https://doi.org/10.1103/PhysRevA.111.022437>
- Denner, M. M. (2023). A hybrid quantum-classical method for electron-phonon systems. *Communications Physics*, 6(1). <https://doi.org/10.1038/s42005-023-01353-3>
- Dhara, B. (2025). Beamforming optimization via quantum algorithms using Variational Quantum Eigensolver and Quantum Approximate Optimization Algorithm. *Iet Quantum Communication*, 6(1). <https://doi.org/10.1049/qtc2.12120>
- Fan, Y. (2023). Circuit-Depth Reduction of Unitary-Coupled-Cluster Ansatz by Energy Sorting. *Journal of Physical Chemistry Letters*, 14(43), 9596–9603. <https://doi.org/10.1021/acs.jpclett.3c01804>
- Grimsley, H. R. (2023). Adaptive, problem-tailored variational quantum eigensolver mitigates rough parameter landscapes and barren plateaus. *Npj Quantum Information*, 9(1). <https://doi.org/10.1038/s41534-023-00681-0>
- Grimsley, H. R. (2025). Challenging excited states from adaptive quantum eigensolvers: Subspace expansions vs. State-averaged strategies. *Quantum Science and Technology*, 10(2). <https://doi.org/10.1088/2058-9565/ad9fa2>
- Guan, J. (2023). Detecting Violations of Differential Privacy for Quantum Algorithms. *Ccs 2023 Proceedings of the 2023 ACM Sigsac Conference on Computer and Communications Security*, (Query date: 2026-08-07 19:31:40), 2277–2291. <https://doi.org/10.1145/3576915.3623108>
- Gupta, S. (2023). A Quantum Approach for Stochastic Constrained Binary Optimization. *ICASSP IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, 2023(Query date: 2026-08-07 19:31:40). <https://doi.org/10.1109/ICASSP49357.2023.10097179>

- Hagai, M. (2023). Artificial neural network encoding of molecular wavefunctions for quantum computing†. *Digital Discovery*, 2(3), 634–650. <https://doi.org/10.1039/d2dd00093h>
- Halder, D. (2023). Corrections beyond coupled cluster singles and doubles through selected generalized rank-two operators: Digital quantum simulation of strongly correlated systems. *Journal of Chemical Sciences*, 135(2). <https://doi.org/10.1007/s12039-023-02161-6>
- Hohenstein, E. G. (2023). Efficient quantum analytic nuclear gradients with double factorization. *Journal of Chemical Physics*, 158(11). <https://doi.org/10.1063/5.0137167>
- Khamoshi, A. (2023). AGP-based unitary coupled cluster theory for quantum computers. *Quantum Science and Technology*, 8(1). <https://doi.org/10.1088/2058-9565/ac93ae>
- Khumalo, M. T. (2025). An investigation of IBM quantum computing device performance on combinatorial optimisation problems. *Neural Computing and Applications*, 37(2), 611–626. <https://doi.org/10.1007/s00521-022-07438-4>
- Le, N. T. (2023). Correlated Reference-Assisted Variational Quantum Eigensolver. *Journal of Physical Chemistry A*, 127(24), 5222–5230. <https://doi.org/10.1021/acs.jpca.3c00993>
- Li, A. C. Y. (2023). Benchmarking variational quantum eigensolvers for the square-octagon-lattice Kitaev model. *Physical Review Research*, 5(3). <https://doi.org/10.1103/PhysRevResearch.5.033071>
- Loipersberger, M. (2023). Accurate non-covalent interaction energies on noisy intermediate-scale quantum computers via second-order symmetry-adapted perturbation theory. *Chemical Science*, 14(13), 3587–3599. <https://doi.org/10.1039/d2sc05896k>
- Marquversen, F. K. (2023). Applications and resource reductions in measurement-based variational quantum eigensolvers. *Quantum Science and Technology*, 8(4). <https://doi.org/10.1088/2058-9565/ace2e6>
- Michel, A. (2023). Blueprint for a digital-analog variational quantum eigensolver using Rydberg atom arrays. *Physical Review A*, 107(4). <https://doi.org/10.1103/PhysRevA.107.042602>
- Mineh, L. (2023). Accelerating the variational quantum eigensolver using parallelism. *Quantum Science and Technology*, 8(3). <https://doi.org/10.1088/2058-9565/acd0d2>
- Mohanty, N. (2023). Analysis of the Vehicle Routing Problem Solved via Hybrid Quantum Algorithms in the Presence of Noisy Channels. *IEEE Transactions on Quantum Engineering*, 4(Query date: 2026-08-07 19:31:40). <https://doi.org/10.1109/TQE.2023.3303989>
- Mondal, D. (2023). Development of a compact Ansatz via operator commutativity screening: Digital quantum simulation of molecular systems. *Journal of Chemical Physics*, 159(1). <https://doi.org/10.1063/5.0153182>
- Mondal, D. (2025). Determination of Molecular Symmetry Adapted Eigenroots in the Variational Quantum Eigensolver Framework. *International Journal of Quantum Chemistry*, 125(7). <https://doi.org/10.1002/qua.70029>
- Mukherjee, A. (2023). Comparative study of adaptive variational quantum eigensolvers for multi-orbital impurity models. *Communications Physics*, 6(1). <https://doi.org/10.1038/s42005-022-01089-6>
- Nagy, D. (2024). Electric Field Gradient Calculations for Ice VIII and IX Using Polarizable Embedding: A Comparative Study on Classical Computers and Quantum Simulators. *Journal of Physical Chemistry A*, 128(30), 6305–6315. <https://doi.org/10.1021/acs.jpca.4c02697>
- Okada, K. N. (2023). Classically optimized variational quantum eigensolver with applications to topological phases. *Physical Review Research*, 5(4). <https://doi.org/10.1103/PhysRevResearch.5.043217>
- Pollard, N. (2025). BenchQC: A Benchmarking Toolkit for Quantum Computation. *Journal of Computational Chemistry*, 46(21). <https://doi.org/10.1002/jcc.70202>

- Ravi, G. S. (2023). CAFQA: A Classical Simulation Bootstrap for Variational Quantum Algorithms. *International Conference on Architectural Support for Programming Languages and Operating Systems ASPLOS, 1*(Query date: 2026-08-07 19:31:40), 15–29. <https://doi.org/10.1145/3567955.3567958>
- Roca-Jerat, S. (2023). Circuit Complexity through phase transitions: Consequences in quantum state preparation. *Scipost Physics*, 15(5). <https://doi.org/10.21468/SciPostPhys.15.5.186>
- Sakurai, R. (2024). Comparative study on compact quantum circuits of hybrid quantum-classical algorithms for quantum impurity models. *Physical Review Research*, 6(2). <https://doi.org/10.1103/PhysRevResearch.6.023110>
- Sayginel, H. (2024). A fault-tolerant variational quantum algorithm with limited T-depth. *Quantum Science and Technology*, 9(1). <https://doi.org/10.1088/2058-9565/ad0571>
- Schroeder, A. (2024). Deterministic Ansätze for the measurement-based variational quantum eigensolver. *New Journal of Physics*, 26(6). <https://doi.org/10.1088/1367-2630/ad51e5>
- Scriva, G. (2024). Challenges of variational quantum optimization with measurement shot noise. *Physical Review A*, 109(3). <https://doi.org/10.1103/PhysRevA.109.032408>
- Sennane, W. (2023). Calculating the ground-state energy of benzene under spatial deformations with noisy quantum computing. *Physical Review A*, 107(1). <https://doi.org/10.1103/PhysRevA.107.012416>
- Shkolnikov, V. O. (2023). Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers. *Quantum*, 7(Query date: 2026-08-07 19:31:40). <https://doi.org/10.22331/q-2023-06-12-1040>
- Simmons, J. (2023). Computing vibrational spectra using a new collocation method with a pruned basis and more points than basis functions: Avoiding quadrature. *Journal of Chemical Physics*, 158(14). <https://doi.org/10.1063/5.0146703>
- Sivakumar, A. (2024). A computational study and analysis of Variational Quantum Eigensolver over multiple parameters for molecules and ions. *EPJ Quantum Technology*, 11(1). <https://doi.org/10.1140/epjqt/s40507-024-00280-8>
- Skogh, M. (2023). Accelerating variational quantum eigensolver convergence using parameter transfer. *Electronic Structure*, 5(3). <https://doi.org/10.1088/2516-1075/ace86d>
- Stokes, J. (2023). Continuous-variable neural network quantum states and the quantum rotor model. *Quantum Machine Intelligence*, 5(1). <https://doi.org/10.1007/s42484-023-00100-9>
- Sun, R. Y. (2023). Efficient variational quantum circuit structure for correlated topological phases. *Physical Review B*, 108(7). <https://doi.org/10.1103/PhysRevB.108.075127>
- Tsukayama, D. (2025). Enhancing computational accuracy with parallel parameter optimization in variational quantum eigensolver. *Aip Advances*, 15(1). <https://doi.org/10.1063/5.0236028>
- Uttarkar, A. (2024). A comparative insight into peptide folding with quantum CVaR-VQE algorithm, MD simulations and structural alphabet analysis. *Quantum Information Processing*, 23(2). <https://doi.org/10.1007/s11128-024-04261-9>
- Volpe, D. (2024). A Predictive Approach for Selecting the Best Quantum Solver for an Optimization Problem. *Proceedings IEEE Quantum Week 2024 Qce 2024, 1*(Query date: 2026-08-07 19:31:40), 1014–1025. <https://doi.org/10.1109/QCE60285.2024.00121>
- Vu, T. H. (2025). Benchmarking Variants of the Adam Optimizer for Quantum Machine Learning Applications. *IEEE Open Journal of the Computer Society*, 6(Query date: 2026-08-07 19:31:40), 1146–1154. <https://doi.org/10.1109/OJCS.2025.3586953>
- Weaving, T. (2025). Contextual subspace variational quantum eigensolver calculation of the dissociation curve of molecular nitrogen on a superconducting quantum computer. *Npj Quantum Information*, 11(1). <https://doi.org/10.1038/s41534-024-00952-4>

-
- Xu, G. (2023). Concurrent quantum eigensolver for multiple low-energy eigenstates. *Physical Review A*, 107(5). <https://doi.org/10.1103/PhysRevA.107.052423>
- Yoffe, D. (2025). A qubit-efficient variational selected configuration-interaction method. *Quantum Science and Technology*, 10(1). <https://doi.org/10.1088/2058-9565/ad7d32>
- Zaytsev, V. A. (2024). Calculation of the moscovium ground-state energy by quantum algorithms. *International Journal of Quantum Chemistry*, 124(1). <https://doi.org/10.1002/qua.27232>
- Zhang, H. K. (2024). Absence of Barren Plateaus in Finite Local-Depth Circuits with Long-Range Entanglement. *Physical Review Letters*, 132(15). <https://doi.org/10.1103/PhysRevLett.132.150603>
- Zhang, J. (2023). DISQ: Dynamic Iteration Skipping for Variational Quantum Algorithms. *Proceedings 2023 IEEE International Conference on Quantum Computing and Engineering Qce 2023*, 1(Query date: 2026-08-07 19:31:40), 1062–1073. <https://doi.org/10.1109/QCE57702.2023.00120>
-

Copyright Holder :

© Mariyam Asma et al. (2026).

First Publication Right :

© Journal of Tecnologia Quantica

This article is under: