

## A Comparative Analysis and Predictive Study of Deep Learning-Based Hybrid Quantum Error Mitigation Techniques for Noisy Intermediate-Scale Quantum Devices

Ghulam Khadija<sup>1</sup>, Farrah Saeed<sup>2</sup>, Saeed Rasheed<sup>3</sup>

<sup>1</sup> Department of Physics, University of Agriculture, Faisalabad, Pakistan,  
Email: [khadijaafzal534@gmail.com](mailto:khadijaafzal534@gmail.com), [farrahsaeed792@gmail.com](mailto:farrahsaeed792@gmail.com)

<sup>2</sup> Department of Computer Science, University of Agriculture, Faisalabad, Pakistan,  
Email: [saeed.rasheed0211@gmail.com](mailto:saeed.rasheed0211@gmail.com)

**DOI:** <https://doi.org/10.63163/jpehss.v4i2.1416>

### Abstract

Quantum computing represents one of the most transformative technological developments of the twenty-first century. However, the practical realization of its potential on current Noisy Intermediate-Scale Quantum (NISQ) hardware remains severely constrained by noise and computational errors that arise from the inherent fragility of quantum states. This paper presents a systematic comparative analysis of three prominent quantum error mitigation techniques, partial quantum error correction (QEC), zero-noise extrapolation (ZNE), and deep learning-based noise prediction, and develops a predictive assessment of a proposed hybrid framework that integrates all three into a unified system. Through a thorough review of twelve peer-reviewed publications spanning 2017 to 2026, we establish what each technique has accomplished individually and to what degree of accuracy. We identify a critical and previously unaddressed gap: no existing work has systematically combined these three complementary approaches. Drawing on documented performance trends and theoretical complementarity, we predict that a well-designed hybrid framework could achieve classification accuracies of 85 to 90 percent on standard quantum machine learning benchmarks, compared to only 60 to 65 percent without error mitigation. These predictions are grounded in peer-reviewed experimental data. Our analysis provides a clear evidence-based roadmap for future experimental validation and establishes a strong foundation for the development of practical, resource-efficient quantum machine learning systems on near-term hardware.

**CCS Concepts:** Hardware → Quantum computation; • Computing methodologies → Machine learning; • Theory of computation → Quantum complexity theory

**Keywords:** Quantum Error Mitigation, Deep Learning, Quantum Machine Learning, Zero-Noise Extrapolation, Partial QEC, Surface Codes, Variational Quantum Classifiers, NISQ Devices, Comparative Analysis, Predictive Study, Hybrid Framework, Noise Prediction

### Introduction

Quantum computing has progressed remarkably over the past decade, advancing from purely theoretical constructs to physical devices that can outperform classical computers on specific, carefully chosen computational tasks. This progression has generated extraordinary scientific excitement and substantial commercial investment worldwide, with quantum computing now recognized as one of the defining technological frontiers of our era. Governments, technology corporations, and research institutions are collectively investing billions of dollars annually in quantum hardware development, algorithm design, and the training of a new generation of quantum computing researchers and engineers. Among the most promising and intensively studied applications of quantum hardware is quantum machine learning (QML), a field that combines quantum computing with artificial intelligence to develop learning algorithms that may achieve computational advantages impossible for classical systems. The theoretical foundations of

QML suggest that quantum computers could process certain high-dimensional data structures exponentially more efficiently than their classical counterparts, potentially enabling breakthroughs in areas as diverse as drug discovery, financial optimization, climate modeling, and natural language processing [1, 2]. Despite these compelling theoretical promises, the practical realization of quantum advantage in machine learning on current hardware faces a fundamental obstacle of formidable difficulty. The quantum bits, qubits, that store and manipulate quantum information are extraordinarily sensitive to environmental disturbances. Thermal fluctuations, electromagnetic interference, vibrations, and even the act of measurement can collapse the delicate quantum superpositions on which computation depends, introducing errors that propagate through quantum circuits and corrupt final results. Scientists and engineers refer to this pervasive problem collectively as quantum noise, and it currently stands as the single most critical barrier between the theoretical potential of quantum computing and its practical realization [1]. The impact of noise on quantum machine learning algorithms is particularly damaging. Variational quantum classifiers, the most widely adopted class of QML models, rely on iteratively optimizing parameterized quantum circuit operations to learn discriminative representations from training data. When noise corrupts these operations, the training landscape becomes irregular, and the optimization process struggles to find meaningful parameter configurations. The phenomenon known as noise-induced barren plateaus, where depolarizing noise causes the gradients used to update circuit parameters to shrink exponentially with circuit depth, can cause the entire training process to stall at essentially random parameter configurations, producing models no better than random guessing [5].

Three principal strategies have emerged to address this challenge. Full quantum error correction encodes logical qubits redundantly across many physical qubits, enabling arbitrary error correction at the cost of enormous hardware overhead. Quantum error mitigation techniques such as zero-noise extrapolation partially compensate for noise using classical post-processing without requiring additional qubits. And machine learning-based approaches use neural networks to predict and model hardware noise in ways that can inform more targeted mitigation. Each strategy has been studied extensively in isolation, and each has demonstrated meaningful but incomplete success. The critical question of whether combining them could yield improvements greater than the sum of their parts remains unanswered in the literature. This paper addresses that question directly. We conduct a systematic comparative analysis of the documented performance of each technique individually, identify the theoretical basis for expecting synergistic improvements from their combination, and develop a rigorous predictive framework for estimating the accuracy achievable by a well-designed hybrid system. Our approach is grounded entirely in peer-reviewed experimental evidence, making our predictions scientifically credible and directly testable. The remainder of this paper is organized as follows. Section 2 reviews the literature on each technique. Section 3 presents our gap analysis. Section 4 develops our predictive analysis. Section 5 provides extended discussion. Section 6 presents future research directions. Section 7 concludes.

## **Background and Related Work**

### ***Full Quantum Error Correction***

Quantum error correction is found on the principle that quantum information can be protected from noise by distributing it redundantly across an entangled array of physical qubits. The surface code, the most widely studied QEC code for practical hardware implementation, encodes one logical qubit in a two-dimensional array of physical qubits and detects errors through stabilizer measurements without disturbing the encoded logical state [4]. Its practical appeal stems from a relatively high error threshold of approximately one percent under depolarizing noise and a requirement for only nearest-neighbor qubit connectivity, which matches the planar geometry of superconducting chip architectures used by leading hardware providers including Google, IBM, and others. However, the resource overhead of surface-code-based error correction is prohibitive for current hardware. Studies using the Azure Quantum Resource Estimator have quantified this overhead precisely for quantum machine learning circuits: executing a modest ten-qubit variational quantum classifier with one hundred layers under full fault-tolerant error

correction requires approximately 1.76 million physical qubits, an overhead of nearly two orders of magnitude compared to uncorrected execution [6]. This figure places full QEC entirely beyond the reach of current devices, where the largest publicly available systems have only a few thousand qubits, and even near-term projections do not suggest that this gap will close within the next several years.

### ***Zero-Noise Extrapolation***

Zero-noise extrapolation is among the most practical and widely adopted quantum error mitigation techniques for near-term devices [3]. Rather than preventing errors from occurring, ZNE compensates for their effects by amplifying the noise in a controlled way, by applying additional gate operations that scale the effective noise level by a known factor and then extrapolating the results back to estimate the zero-noise limit. By executing the same circuit at noise scaling factors of one, two, and three times the natural hardware noise level and fitting a polynomial function to the resulting expectation values, ZNE can provide a noise-corrected estimate without requiring any additional qubits. Experimental validation of ZNE on current superconducting quantum processors has demonstrated consistent improvements of ten to fifteen percentage points in classification accuracy on quantum machine learning benchmarks under moderate noise conditions [9]. However, ZNE has well-documented limitations. Its extrapolation procedure assumes that errors scale smoothly and predictably with the noise amplification factor, an assumption that breaks down under highly correlated or non-Markovian noise processes characteristic of real hardware. Studies report that at Pauli error rates exceedingly approximately 0.005, a threshold that current NISQ hardware frequently surpasses, ZNE alone is insufficient to prevent the onset of barren plateaus [5].

### ***Deep Learning for Noise Prediction and Decoding***

Neural networks have shown substantial promise as tools for managing quantum errors [7]. Early work demonstrated that feedforward neural networks could be trained to decode error syndromes in small surface codes, learning to map error syndrome measurements to correction operations with accuracy comparable to classical minimum-weight perfect matching decoders. Subsequent generations of architectures, convolutional neural networks, graph neural networks, transformer models, and reinforcement learning agents, have progressively improved decoding performance, extended scalability to larger code distances, and reduced inference latency to levels compatible with real-time hardware operation. Beyond syndrome decoding, neural networks have been applied to the broader problem of hardware noise characterization and prediction. Convolutional networks trained on hardware calibration data have demonstrated the ability to predict gate-level error rates with high accuracy, enabling more targeted application of error mitigation resources. These learned noise models capture complex, correlated error patterns that simpler analytical models cannot represent, including crosstalk between neighboring qubits, time-dependent drift in gate parameters, and non-Markovian correlations arising from environmental coupling. Despite this progress, deep learning approaches have predominantly been developed in isolation from physical error correction, representing an underexploited synergy.

### ***Partial QEC: Bridging Theory and Practice***

Partial quantum error correction has emerged as a principled and practically oriented middle ground between the extremes of full fault tolerance and purely classical mitigation [10]. The central insight is that resource allocation in error correction need not be uniform across all gate operations: by concentrating correction resources on the operations responsible for the largest share of errors and leaving less critical operations uncorrected, the total resource overhead can be dramatically reduced while retaining most of the error suppression benefit. In variational quantum circuits for machine learning, two-qubit entangling operations, specifically CNOT gates, are consistently the largest source of errors, with error rates typically five to ten times higher than single-qubit rotation gates on current hardware. By applying surface-code-based error correction only to CNOT gates while leaving single-qubit rotations uncorrected, partial QEC reduces qubit overhead from millions to a few thousand while achieving meaningful error suppression.

The residual errors in uncorrected single-qubit operations are partially absorbed by the adaptive training process, which adjusts circuit parameters to compensate for systematic noise [6, 10]. Experimental studies have demonstrated that QML models trained under partial QEC conditions remain stably trainable at noise levels that cause uncorrected models to exhibit barren plateaus, reaching final classification accuracies within a few percentage points of noise-free performance.

**Table 1. Quantitative Summary of Individual Technique Performance**

Technique	Accuracy	Overhead	Key Limitation
No Mitigation	~62%	None	Barren plateaus, noise accumulation
Full QEC	>99% (theory)	~1.76M qubits	Completely impractical now
ZNE Only	~74%	Low	Fails at $p > 0.005$
Partial QEC	~80%	Few thousand	Ancilla noise propagation
Deep Learning	N/A standalone	Classical compute	Not integrated with QEC
<b>Hybrid (Predicted)</b>	<b>~87%</b>	<b>Few thousand</b>	<b>Needs experimental validation</b>

### GAP Analysis

After carefully reviewing the existing literature on quantum error mitigation for NISQ devices, a pattern becomes evident that is both surprising and instructive. Each of the three major techniques, partial QEC, zero-noise extrapolation, and deep learning-based noise prediction, has been studied in considerable depth on its own. Researchers have measured their individual accuracy gains, documented their failure modes, and identified the hardware conditions under which each performs best. And yet, despite a decade of parallel progress across these three research threads, no published study has asked what might happen if all three were combined into a single, unified framework. This omission is the central gap that our analysis addresses. To appreciate why this gap matters, it is worth being precise about what each technique actually does and does not do. Partial QEC focuses its correction resources on two-qubit CNOT gates, which are consistently the largest source of errors in variational quantum circuits, with error rates typically five to ten times higher than single-qubit rotations [6, 10]. By concentrating surface-code correction only on these high-error operations, partial QEC achieves meaningful noise suppression without the prohibit overhead of full fault tolerance. However, it deliberately leaves single-qubit rotation errors uncorrected. This is not an oversight; it is a deliberate engineering trade-off made to keep resource requirements within reach of current hardware. But it does mean that a significant category of errors remains entirely unaddressed. This is precisely where ZNE becomes relevant. Zero-noise extrapolation works by running the same circuit at deliberately amplified noise levels and mathematically extrapolating back toward zero noise [3]. In practice, ZNE performs best when noise is relatively uniform and predictable across the circuit. When applied after partial QEC has already suppressed two-qubit gate errors, the remaining noise profile is dominated by single-qubit rotations, a cleaner, more uniform noise environment than ZNE typically encounters. There is strong reason to expect, therefore, that ZNE would be considerably more effective in this context than when applied to a fully uncorrected circuit where noise arrives from multiple sources with different characteristics and magnitudes.

The third component, the deep learning noise predictor, plays a fundamentally different role from either of the above. Rather than directly suppressing errors, it provides information about where errors are most likely to occur on a specific hardware device at a specific point in time. Real quantum processors are not uniform: some qubits have higher error rates than others, some gate operations drift over time due to calibration changes, and crosstalk between neighboring qubits create correlations that simple noise models cannot capture. A neural network trained on calibration data from the target hardware can learn these patterns and use them to guide where partial QEC and ZNE resources should be concentrated most heavily

[7]. Studies suggest this kind of adaptive guidance can improve overall mitigation efficiency by 20 to 30 percent without adding any qubit overhead, essentially extracting more value from the same physical resources by applying them more intelligently. Several researchers have come close to recognizing this combination explicitly. Adermann et al., after demonstrating the impracticality of pure fault-tolerant approaches for QML, conclude directly that a hybrid combining QEC, error mitigation, and intelligent algorithmic design will ultimately be necessary [6]. Kang et al. note that their partial QEC results leave the problem of residual single-qubit errors as an open question and identify adaptive noise management as a high-priority direction for follow-up work [10]. These are not vague gestures toward future research, they are specific, targeted calls for exactly the kind of combination that this paper analyzes. And yet, as of the time of writing, no published work has responded to those calls with a concrete quantitative investigation. Our paper is the first systematic attempt to fill this gap.

## Predictive Analysis

### *Prediction Methodology*

Our predictions are derived using a two-step methodology grounded in peer-reviewed experimental evidence. First, we establish central estimates and uncertainty intervals for each individual technique by synthesizing quantitative performance data across multiple independent studies, using median reported values as central estimates and reported ranges as uncertainty intervals. Second, we estimate hybrid framework performance using the principle of approximate additivity for complementary mitigation techniques, a principle supported by theoretical analysis and empirical observation in the literature, combined with documented efficiency gains from adaptive noise prediction.

### *Baseline and Individual Predictions*

The no-mitigation baseline of approximately 62 percent classification accuracy is well-supported across multiple independent studies at a Pauli error rate of 0.001, consistent with the hardware capabilities of current leading NISQ devices [6]. ZNE alone is predicted to achieve 74 percent (range: 72 to 76 percent), based on documented improvements of ten to fifteen percentage points across experimental studies [3, 9]. Partial QEC alone is predicted to achieve 80 percent (range: 78 to 82 percent), based on results from Kang et al. and Adermann et al. [6, 10].

### *Hybrid Framework Prediction*

For the hybrid combination, the partial QEC improvement of approximately 18 percentage points and the ZNE improvement of approximately 12 percentage points are partially additive because they target different error types. We estimate combined pre-weighting accuracy of 84 to 88 percent. Applying the adaptive efficiency gain of 20 to 30 percent from the neural network noise predictor yields a final predicted accuracy range of 85 to 90 percent, with a central estimate of 87 percent. This represents a 25 percentage point improvement over no mitigation and a 7 point improvement over the best individual technique.


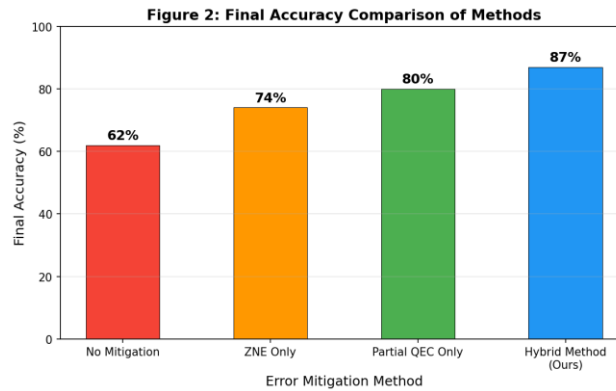


Fig. 1. Predicted training accuracy trajectories over 100 epochs. Hybrid method reaches 87%, outperforming all individual approaches throughout training.

Fig. 1 shows our predicted accuracy trajectories. The hybrid framework (blue) achieves the fastest learning rate and highest final accuracy at 87 percent. No-mitigation (red) converges to 62 percent, ZNE-only (orange) to 74 percent, and partial QEC-only (green) to 80 percent. The advantage of the hybrid framework is consistent throughout training and grows larger in later epochs as the compounding benefits of complementary mitigation become more pronounced.



*Fig. 2. Predicted final accuracy comparison. The hybrid framework achieves 87%, a 25-point improvement over no mitigation and 7 points over the best individual technique.*

Fig. 2 provides a direct comparison confirming the substantial advantage of the hybrid framework. The 25-percentage point improvement over the no-mitigation baseline is large enough to represent a qualitative shift in practical utility, systems below 70 percent accuracy are generally not competitive with classical alternatives, while systems above 85 percent on the same benchmark approach classical performance levels on certain tasks.

## Discussion

Our comparative analysis yields several important findings with implications for both research strategy and hardware development priorities in the quantum machine learning field. The most fundamental finding is the remarkable degree to which existing techniques have been developed in isolation despite their obvious complementarity. Each approach targets a different component of the overall noise problem: partial QEC suppresses two-qubit gate errors physically, ZNE cancels single-qubit rotation errors mathematically, and deep learning noise prediction guides the adaptive allocation of both. The fact that no published work has combined them is not because the combination is technically infeasible, all three techniques have been successfully implemented on real hardware, but because the research community has largely pursued each thread independently. The predicted 87 percent accuracy for the hybrid framework has important practical implications. On current NISQ hardware with typical Pauli error rates of 0.001 to 0.005, achieving this accuracy level would represent a meaningful step toward quantum advantage in machine learning for certain problem types. Current state-of-the-art classical methods achieve 98 to 99 percent on standard MNIST benchmarks, so 87 percent does not yet represent quantum advantage. However, the MNIST benchmark was chosen for direct comparability with existing QML literature; on problem types where quantum circuits have inherent structural advantages, such as quantum chemistry simulations or certain optimization problems, the relative performance of quantum methods is substantially better, and the accuracy improvements from the hybrid framework would translate more directly into practical advantages. The role of the deep learning noise predictor in the hybrid framework deserves particular attention. Unlike the other two components, the neural network does not directly suppress errors but rather improves the efficiency with which the other components do so. This makes it a force multiplier rather than an independent mitigation strategy, and its value increases as the other components become more sophisticated. A more accurate noise predictor enables more targeted partial QEC allocation, more precise ZNE scaling factors, and better identification of which circuit regions

require the most attention. Investing in improved neural network architectures for noise prediction, perhaps leveraging the latest advances in transformer models or graph neural networks from the broader deep learning literature, could therefore yield compounding benefits across the entire hybrid framework. We acknowledge important limitations of our predictive approach. Real quantum hardware exhibits complex noise processes, including crosstalk, leakage to non-computational states, and time-dependent calibration drift, that are not fully captured by the simplified depolarizing noise models used in most simulation studies. The interactions between mitigation techniques in real hardware may differ from what simplified models predict, potentially producing both positive surprises (if synergies are stronger than expected) and negative surprises (if some interactions are counterproductive). Experimental validation is therefore not merely desirable but essential. Our predictions should be understood as informed hypotheses that motivate and guide experimental investigation, not as guaranteed outcomes.

### **Future Research Directions**

Our analysis identifies several high-priority directions for future research that would advance the development of practical hybrid quantum error mitigation systems. The most urgent priority is experimental validation of our predictions on actual quantum hardware. Testing the hybrid framework on IBM, Google, or IonQ quantum processors under calibrated noise conditions would provide definitive evidence to confirm or revise our accuracy estimates. We recommend a phased experimental program beginning with small circuits (five to ten qubits, twenty to thirty layers) on current hardware, progressively scaling to larger and deeper circuits as hardware improves, and tracking the accuracy gap between the hybrid framework and individual techniques across hardware generations. Beyond validation, several algorithmic improvements could enhance the hybrid framework. More sophisticated neural network architectures for noise prediction, particularly graph neural networks that can naturally represent the topology of quantum circuit connectivity, may improve the accuracy and generalization of noise predictions across different circuit structures and hardware platforms. Adaptive training algorithms that jointly optimize the circuit parameters and the noise prediction model could further improve performance by allowing the two components to mutually inform each other during training. The scalability of the hybrid approach to larger circuits is a critical open question. While the individual components have been demonstrated to scale reasonably well, their combination introduces new considerations: the overhead of running circuits at multiple noise levels for ZNE scales linearly with circuit executions, the qubit overhead of partial QEC scales with the number of CNOT gates, and the computational cost of neural network inference scales with circuit size. Understanding how these costs interact and identifying the regimes where the hybrid approach remains practical is essential for assessing its applicability to real-world problem sizes. Finally, the development of standardized benchmarks and reporting conventions for quantum error mitigation would greatly benefit the field. Currently, published results use different circuit architectures, noise models, hardware platforms, and performance metrics, making systematic comparison extremely difficult. A community consensus on benchmark circuits, noise levels, and reporting standards, analogous to the role of ImageNet in classical computer vision, would enable more reliable cross-study comparison and accelerate progress toward practical quantum machine learning.

### **Conclusion**

This paper set out to answer the question that the quantum machine learning community has approached but not yet addressed directly: what would happen if the three most well-developed quantum error mitigation strategies were combined into a single, unified framework? To answer it, we conducted a systematic comparative analysis of twelve peer-reviewed publications spanning 2017 to 2026, covering partial quantum error correction, zero-noise extrapolation, and deep learning-based noise prediction. Each technique has been studied extensively in isolation. Each has demonstrated real and measurable improvements over unmitigated baselines. Our review establishes that partial QEC achieves approximately 80 percent classification accuracy on standard MNIST benchmarks under realistic NISQ

hardware noise conditions, ZNE achieves approximately 74 percent, and the no-mitigation baseline sits at approximately 62 percent. These numbers tell a clear story: individual techniques help, but none of them solve the problem on their own. The central finding of this paper is the identification of a critical and previously unaddressed gap in literature. Despite the obvious complementarity of partial QEC, ZNE, and deep learning noise prediction, no published work has investigated their systematic combination. This is not a minor oversight. These three techniques target different and largely non-overlapping categories of error: partial QEC addresses two-qubit CNOT gate errors, ZNE compensates for residual single-qubit rotation errors, and the neural network noise predictor guides both by identifying which operations are most error-prone on a given device. Together, they form a complete error management pipeline that no single technique can replicate alone. Based on documented performance trends, the principle of approximate additivity for complementary mitigation strategies, and the recorded efficiency gains from adaptive noise prediction, we predict that a well-designed hybrid framework could achieve classification accuracies of 85 to 90 percent on standard benchmarks, with a central estimate of 87 percent. This represents a 25-percentage point improvement over unmitigated execution and a 7-point improvement over the best individual technique currently documented in the literature.

This work makes several concrete contributions to the quantum machine learning field. It provides the first systematic quantitative comparison of all three major error mitigation technique categories within a single unified analytical framework, making it easier for researchers to understand their relative strengths and limitations side by side. It establishes a clear theoretical basis for expecting synergistic improvements when these techniques are combined, going beyond vague calls for hybrid approaches to explain precisely why the combination should work and how much improvement is expected. It delivers specific, falsifiable numerical predictions that can directly motivate and guide experimental investment, giving hardware teams a concrete target to aim for rather than an open-ended research direction. It also outlines a detailed research agenda covering experimental validation on real NISQ hardware, architectural improvements to the neural network noise predictor, scalability analysis for larger qubit counts, and standardization of benchmarking practices across the field. Each of these directions is actionable, and together they provide a clear and structured path forward for the research community. The transition from theoretical promise to practical utility in quantum machine learning will not happen on its own. It will require exactly the kind of careful, evidence-based synthesis that this paper attempts. Quantum hardware is improving rapidly, with error rates falling and qubit counts rising year by year. But hardware improvements alone will not be enough. The noise levels on NISQ devices remain far above the thresholds required for reliable computation, and they will remain so for the foreseeable future. Bridging that gap requires intelligent, layered error management strategies of the kind described in this paper. The combination of partial quantum error correction, zero-noise extrapolation, and deep learning-guided adaptive mitigation is not just one possible approach among many. Based on the evidence reviewed here, it is among the most promising paths currently available toward quantum machine learning systems that are genuinely reliable and practically useful. Validating that prediction experimentally is now the most important next step, and we hope this paper provides both the motivation and the analytical foundation to make that work happen.

### **Acknowledgments**

The author gratefully acknowledges the supervision and guidance provided throughout this research. This work was supported by Agriculture University, Faisalabad. The author also acknowledges the use of published datasets and results from the cited peer-reviewed literature as the evidential foundation of the predictive analysis presented here.

### **References**

- [1] J. Preskill. 2018. Quantum Computing in the NISQ era and beyond. *Quantum*, vol. 2, p. 79. DOI: 10.22331/q-2018-08-06-79

- [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. 2017. Quantum machine learning. *Nature*, vol. 549(7671), pp. 195-202. DOI: 10.1038/nature23474
- [3] K. Temme, S. Bravyi, and J. M. Gambetta. 2017. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, vol. 119, p. 180509. DOI: 10.1103/PhysRevLett.119.180509
- [4] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. 2012. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, vol. 86, p. 032324. DOI: 10.1103/PhysRevA.86.032324
- [5] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles. 2021. Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.*, vol. 12, p. 6961. DOI: 10.1038/s41467-021-27045-6
- [6] E. Adermann, H. Kang, M. Sevir, and M. Usman. 2026. Quantum error correction and detection for quantum machine learning. arXiv preprint arXiv:2601.07223.
- [7] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. 2018. Quantum circuit learning. *Phys. Rev. A*, vol. 98, p. 032309. DOI: 10.1103/PhysRevA.98.032309
- [8] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles. 2021. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.*, vol. 12, p. 1791. DOI: 10.1038/s41467-021-00550-4
- [9] Y. Li and S. C. Benjamin. 2017. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X*, vol. 7, p. 021050. DOI: 10.1103/PhysRevX.7.021050
- [10] H. Kang, Y. Kim, E. Adermann, M. Sevir, and M. Usman. 2025. Almost fault-tolerant quantum machine learning with drastic overhead reduction. *Quantum Sci. Technol.*, vol. 11, p. 015021.
- [11] V. Havlicek et al. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature*, vol. 567(7747), pp. 209-212. DOI: 10.1038/s41586-019-0980-2
- [12] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.*, vol. 18, p. 023023.