Journal of Technology Informatics and Engineering (JTIE)
Vol. 4 No. 3 December 2025
E-ISSN: 2961-9068; P-ISSN: 2961-8215, Pages 417-435
DOI: https://doi.org/10.51903/jtie.v4i3.451

Quantum-Inspired Optimization for High-Dimensional Data Classification in Healthcare Analytics

Hanae Sugimoto*[1], Kaito Morishita[2]
Email: hanae_sugimoto_utokyo@gmail.com
[1,2]*Keio University, Tokyo, Japan*
*Corresponding Author

Abstract

High-dimensional medical datasets pose a persistent challenge for artificial intelligence because traditional classification algorithms often incur escalating computational costs and reduced predictive accuracy. As healthcare systems generate increasingly complex clinical records, imaging outputs, and genomic profiles, scalable analytic methods that balance precision and efficiency are critical. This study proposes a Quantum-Inspired Optimization (QIO) framework for efficient and accurate classification of high-dimensional healthcare data. Leveraging the exploratory power of variational quantum algorithms, specifically techniques analogous to the Quantum Approximate Optimization Algorithm, the framework integrates quantum-style search strategies with classical computation to achieve global optimization and numerical stability. Publicly available medical datasets with hundreds of features were used to evaluate the approach. Classification models were trained and tested across varying feature dimensionalities, and performance was assessed using accuracy, runtime, and scalability metrics. Empirical results demonstrate that QIO achieves up to 95.4% classification accuracy and reduces computational time by 40% compared with state-of-the-art classical baselines. The method demonstrates stable convergence and clear decision boundaries even as feature dimensionality grows, highlighting its resilience to the curse of dimensionality. These results indicate that QIO can enable fast and reliable healthcare analytics in data-rich clinical environments. Future research may examine domain-specific adaptations, real-time deployment, and integration with emerging quantum hardware to enhance the impact of quantum-inspired artificial intelligence further.

Keywords: *Quantum-Inspired Optimization, Healthcare Analytics, High-Dimensional Data Classification, Variational Quantum Algorithms, Computational Efficiency.*

I. INTRODUCTION

The rapid digitization of healthcare services has generated unprecedented volumes of heterogeneous and high-dimensional data, including electronic health records (EHR), medical imaging archives, genomic sequences, and data streams from wearable Internet-of-Things (IoT) devices. This explosive growth in data promises to transform clinical decision-making and population health management, yet it simultaneously exposes the computational limits of conventional analytics (Ahmed et al., 2023; Batko & Ślęzak, 2022). Classical machine learning models such as support vector machines, random forests, and even deep neural networks often struggle when confronted with the "curse of dimensionality," where feature spaces grow faster than sample sizes and render training both computationally expensive and memory intensive (Fathi et al., 2021; Malekloo et al., 2022). Consequently, real-time predictive analytics for large-

Received on Oct 2025; Revised on Oct 2025; Accepted on Dec 2025; Published on Dec 2025

scale medical datasets remains challenging for health systems seeking timely and accurate diagnostics (Hussain et al., 2023; Olson, 2023).

Traditional strategies for coping with these computational bottlenecks, including feature selection, dimensionality reduction, and distributed computing, have achieved incremental improvements but have not fundamentally overcome the exponential complexity of high-dimensional healthcare data (Badawy et al., 2023; Cozzoli et al., 2022). Despite advancements in deep learning frameworks for clinical applications (Ahmad et al., 2021; Angelis et al., 2023) Training these frameworks often incurs significant hardware and energy overhead, making it impractical. Moreover, privacy-preserving techniques such as federated learning or encrypted analytics introduce additional computational costs, highlighting the urgent need for more efficient solutions (Bi et al., 2022). Altogether, this creates a notable gap in the literature, as few algorithms provide a balanced trade-off between predictive performance, scalability, and computational efficiency.

As the field of quantum computing develops, quantum algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) and other Variational Quantum Algorithms (VQA) offer fundamentally different strategies for solving combinatorial and high-dimensional optimization problems, leveraging superposition and entanglement to explore large solution spaces efficiently (Biamonte, 2021; Cerezo et al., 2021). Early demonstrations have shown promise in finance and fluid dynamics (Jaksch et al., 2022; Mugel et al., 2022), while more recent studies apply these methods to machine learning workflows for noisy, high-dimensional data (Peters et al., 2021). However, current quantum hardware, known as Noisy Intermediate-Scale Quantum (NISQ) devices, is limited in qubit coherence, gate fidelity, and error correction, and cannot yet support production-scale healthcare analytics (Anschuetz & Kiani, 2022; S. Wang et al., 2021).

These hardware limitations have driven interest in Quantum-Inspired Optimization (QIO) methods, which emulate key features of quantum algorithms on classical hardware, achieving some performance acceleration without requiring fully functional quantum processors (Boev et al., 2023; Bozhedarov et al., 2024). Variational strategies, particularly those inspired by QAOA principles, use parameterized classical circuits to approximate quantum states, enabling efficient exploration of complex objective landscapes (Amaro et al., 2022; Bonet-Monroig et al., 2023). Such approaches have proven successful in combinatorial optimization and large-scale scheduling, suggesting strong potential for healthcare data classification where solution spaces are equally large and complex (Giron et al., 2023; Tran et al., 2025). Importantly, quantum-inspired algorithms can run on standard high-performance computing clusters or GPU-enabled cloud platforms, making them accessible for health informatics applications.

Despite these advances, QIO has seen limited application in healthcare analytics. Most surveys focus on classical machine learning or deep learning, with little attention to hybrid or quantum-inspired approaches (Badawy et al., 2023; Krishna & Tulli, 2023). Experimental studies using VQAs in medical imaging or genomics typically assume specialized quantum hardware and do not address the scalability requirements of hospital-scale data pipelines (Mavaie et al., 2023; Shen et al., 2022). Therefore, there is a critical need to systematically evaluate the feasibility, performance, and operational impact of QIO for high-dimensional medical data classification.

This study addresses this gap by proposing a quantum-inspired optimization framework for large-scale healthcare data classification. Our approach encodes high-dimensional patient data into parameterized models that can be efficiently trained on classical hardware while exploiting quantum-like exploration of solution landscapes (Huynh et al., 2023; Perez-Ramirez, 2024). We design and implement a comprehensive experimental pipeline, including rigorous preprocessing, parameter tuning, and benchmarking against established classifiers such as support vector machines, random forests, and deep neural networks (John et al., 2024). Performance is evaluated using predictive accuracy, F1-score, ROC metrics, computational runtime, and memory utilization to ensure a balanced assessment of both effectiveness and efficiency.

The primary contribution of this work lies in demonstrating that QIO can achieve competitive predictive performance with substantially lower computational cost, providing a hardware-agnostic solution for healthcare analytics (Bittel & Kliesch, 2022; De Palma et al., 2023). The proposed framework is inherently scalable and compatible with existing big data infrastructures in hospitals or national health networks, bridging the gap between quantum theory and operational healthcare needs (Baccour et al., 2022; Hua et al., 2023). Finally, we formulate the research hypothesis as follows: Quantum-Inspired Optimization can enable accurate, scalable, and computationally efficient classification of high-dimensional healthcare data on classical hardware.

The remainder of this paper is organized as follows. Section 2 reviews literature on big data analytics in healthcare, classical optimization, and quantum-inspired algorithms. Section 3 details the methodology, including dataset selection, preprocessing, algorithm design, and experimental setup. Section 4 presents empirical results, Section 5 provides discussion and implications, and Section 6 concludes with contributions and future directions.

II. LITERATURE REVIEW

The modern healthcare ecosystem generates an extraordinary volume of data from electronic health records (EHR), imaging archives, genomic sequencing, and wearable Internet of Things (IoT) devices. These sources produce data that are simultaneously high-dimensional, noisy, and heterogeneous, creating significant obstacles for timely analytics and clinical decision support (Ahmed et al., 2023; Batko & Ślęzak, 2022). High dimensionality leads to the "curse of dimensionality," increasing computational costs and reducing model generalization (Fathi et al., 2021; Malekloo et al., 2022). Noise from inconsistent data entry, sensor drift, and missing values further degrades predictive performance, while heterogeneity across hospitals and devices complicates data integration and model portability (Hussain et al., 2023; Olson, 2023). Traditional preprocessing steps, such as imputation, normalization, and feature selection, mitigate these issues only partially, leaving a persistent need for more advanced computational paradigms (Harrisha et al., 2025; Krishna & Tulli, 2023).

Classical optimization methods remain the cornerstone of artificial intelligence and machine learning. Gradient-based algorithms, like stochastic gradient descent and its variants, are widely used to train neural networks and large language models by iteratively minimizing loss functions (Waltersmann et al., 2021; Xu et al., 2021). However, they are sensitive to local minima, saddle points, and slow convergence, especially for noisy or non-convex medical data (Ahmad et al., 2021; John et al., 2024). High-dimensional features further exacerbate computational burden and limit scalability to national or global healthcare datasets (Angelis et al., 2023).

To address these limitations, heuristics such as Genetic Algorithms (GA), Particle Swarm Optimization (PSO), and simulated annealing have been adopted (Kaack et al., 2022; Li & Muwafak, 2022). These algorithms navigate solution spaces stochastically, avoiding strict differentiability requirements, making them suitable for complex classification tasks (Saha et al., 2020). Nonetheless, they require extensive hyperparameter tuning and may converge slowly when search spaces are large or rugged, as is typical in high-dimensional healthcare analytics (Cozzoli et al., 2022; Rehman et al., 2020). Both gradient-based and heuristic methods struggle to balance computational efficiency and predictive accuracy for modern large-scale medical data.

Quantum computing offers a fundamentally distinct computational model that can address combinatorial and high-dimensional optimization problems. Algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) and the Variational Quantum Eigensolver (VQE) exploit superposition and entanglement to explore vast solution spaces simultaneously, potentially providing polynomial or exponential speedups over classical algorithms (Biamonte, 2021; Cerezo et al., 2021). Variational Quantum Algorithms (VQAs), which combine quantum circuits with classical optimization routines, offer flexibility and adaptability to a diverse range of objective functions (Beckey et al., 2022; Benedetti et al., 2021). Applications include computational fluid dynamics, finance, and high-dimensional pattern recognition (Jaksch et al., 2022; Peters et al., 2021).

In practice, NISQ (Noisy Intermediate-Scale Quantum) devices impose hardware constraints such as qubit decoherence, gate errors, and barren plateaus, limiting circuit depth and algorithm reliability (Fontana et al., 2020; C. Wang & Banko, 2021). Some VQA implementations are NP-hard to train, creating additional convergence challenges (Anschuetz & Kiani, 2022; Bittel & Kliesch, 2022; De Palma et al., 2023). These limitations hinder the deployment of pure quantum solutions for hospital-scale healthcare analytics. Quantum-Inspired Optimization (QIO) has emerged to bridge this gap, leveraging mathematical properties of quantum algorithms while running on classical hardware (Boev et al., 2023; Bozhedarov et al., 2024). Variational strategies based on QAOA or VQE principles use parameterized classical models to approximate quantum state spaces, enabling efficient exploration of complex objective landscapes (Amaro et al., 2022; Bonet-Monroig et al., 2023). QIO has demonstrated superior performance in applications such as wavelength assignment, bin packing, and large-scale scheduling, with results comparable to near-term quantum processors(Giron et al., 2023; Tran et al., 2025).

A key advantage of QIO is its hardware-agnostic nature. It can run on conventional high-performance computing clusters or GPU-enabled cloud platforms, with significantly lower capital requirements compared to quantum hardware (Huynh et al., 2023; Kumar et al., 2024). QIO methods also exhibit faster convergence than classical heuristics and greater resilience to noise than current quantum devices (Mugel et al., 2022; Stein et al., 2022). These features make QIO highly suitable for healthcare analytics, where data are both high-dimensional and dynamic. Despite this promise, empirical applications of QIO in medical data classification remain limited, presenting an open research frontier.

While healthcare big data literature is extensive, few studies systematically apply QIO to large-scale medical data classification. Reviews focus on classical machine learning and deep learning, with minimal attention to hybrid or quantum-inspired frameworks (Badawy et al., 2023; Krishna & Tulli, 2023). Experimental quantum studies often assume specialized hardware rather than practical QIO implementations (Mavaie et al., 2023; Shen et al., 2022). This gap highlights the opportunity to combine quantum principles with classical infrastructure, offering immediate computational benefits and preparing the field for future integration with mature quantum processors (Baccour et al., 2022; Hua et al., 2023).

Table 1 summarizes key literature, comparing gradient-based optimization, heuristic algorithms, pure quantum methods, and quantum-inspired approaches in terms of application, strengths, limitations, and gaps relevant to healthcare analytics. This explicit mapping clarifies where classical techniques fall short, where quantum computing remains impractical, and where QIO offers a feasible, near-term solution for high-dimensional medical data classification.

Table 1. Synthesis of Key Literature

Method	Application Domain	Strengths	Limitations	Identified Gap
Gradient-based optimization	Neural network training, medical image reconstruction (Lin et al., 2021; Waltersmann et al., 2021)	Fast convergence on smooth landscapes; well-established tools	Susceptible to local minima; high computational cost for high-dimensional data	Inefficient for massive heterogeneous healthcare datasets
Heuristic algorithms (GA, PSO)	Healthcare predictive analytics, scheduling (Cozzoli et al., 2022; Rehman et al., 2020)	Independence from differentiability; flexible	Slow convergence; heavy parameter tuning	Limited scalability and energy efficiency
Pure quantum algorithms (QAOA, VQE)	Combinatorial optimization, fluid dynamics (Cerezo et al., 2021; Jaksch et al., 2022)	Theoretical polynomial/exponential speedup	Hardware noise; barren plateaus; NISQ limitations	Not deployable for current hospital-scale analytics
Quantum-inspired optimization	Wavelength assignment, bin packing, and large scheduling (Boev et al., 2023; Tran et al., 2025)	Hardware agnostic; near-quantum speed; robust to noise	Limited real-world healthcare applications	Rarely applied to high-dimensional medical data classification

III. RESEARCH METHOD

This study adopts a quantitative experimental design to evaluate the effectiveness of a Quantum-Inspired Optimization (QIO) framework for high-dimensional healthcare data classification. The methodological pipeline integrates rigorous data preprocessing, a QAOA-based variational optimization core, and comprehensive benchmarking against classical machine-learning baselines. The methodology comprises four main components: dataset description, preprocessing, quantum-inspired algorithm design, and experimental setup. Figure 1 provides a visual overview of the workflow, while Table 2 summarizes the key datasets, algorithms, and evaluation metrics.

A. Dataset Description and Ethical Considerations

This study utilizes publicly available large-scale healthcare datasets to ensure reproducibility and ethical compliance. The datasets include MIMIC-III for electronic health records, high-

Journal of Technology Informatics and Engineering (JTIE)
Vol. 4 No. 3 December 2025
E-ISSN: 2961-9068; P-ISSN: 2961-8215, Pages 417-435
Journal Link: https://jtie.stekom.ac.id/index.php/jtie/index

dimensional genomics profiles, and selected medical imaging datasets, representing structured, semi-structured, and unstructured data modalities. This selection aligns with recent big-data healthcare frameworks emphasizing multimodal analytics (Ahmed et al., 2023; Batko & Ślęzak, 2022; Olson, 2023). Each dataset is characterized by sample size, number of features, and class distribution, as summarized in Table 2. All datasets were fully de-identified prior to use and comply with publicly available data-use agreements. No direct patient identifiers were accessed, and no institutional review board (IRB) approval was required because the study used only anonymized secondary data.

B. Data Preprocessing

To address noise and scale disparities, z-score normalization was applied to numerical features, while categorical attributes were encoded as one-hot vectors. Dimensionality reduction was optionally performed using Principal Component Analysis (PCA) when the feature dimensionality exceeded hardware-efficient limits. This choice follows established high-dimensional preprocessing practices in healthcare analytics (Fathi et al., 2021; Shen et al., 2022). Feature selection was performed using mutual-information ranking and variance filtering to retain clinically relevant predictors while reducing overfitting risk (Malekloo et al., 2022). These steps ensured numerical stability and consistent scaling prior to quantum-inspired encoding.

C. Quantum-Inspired Optimization Algorithm

At the core of the proposed method is a variational formulation inspired by the Quantum Approximate Optimization Algorithm (QAOA) under the broader class of Variational Quantum Algorithms (VQA). Data vectors are encoded as quantum-inspired states via amplitude encoding, followed by variational optimization via parameterized unitary transformations. Measurement outcomes from the variational circuit generate classical feature embeddings that are subsequently processed by a final classical classifier. The workflow follows the sequence: Data Ingestion - Preprocessing - Quantum-Inspired Encoding - Variational Optimization - Measurement - Classical Classification. This design is grounded in prior evidence that variational quantum-inspired models provide superior expressivity for combinatorial and high-dimensional optimization tasks (Amaro et al., 2022; Bozhedarov et al., 2024; Cerezo et al., 2021). To enhance training stability, trap-mitigation and filtering techniques initially developed for noisy quantum systems were incorporated into the optimization process (Anschuetz & Kiani, 2022; C. Wang & Banko, 2021). Key hyperparameters include variational circuit depth, learning rate, and number of optimization iterations.

D. Baseline Models and Evaluation Strategy

For performance comparison, Support Vector Machines (SVM), Random Forests (RF), and a Deep Learning (DL) architecture suitable for high-dimensional data were implemented as baselines. These methods are widely adopted in conventional healthcare predictive analytics (Badawy et al., 2023; John et al., 2024). Model performance was evaluated using accuracy, F1 Score, ROC AUC, computational runtime, and memory utilization. This dual emphasis on predictive effectiveness and computational efficiency responds to recent calls for sustainable and resource-aware artificial intelligence in healthcare (Ahmad et al., 2021; Hussain et al., 2023; Waltersmann et al., 2021). Stratified k-fold cross-validation was employed to ensure robust generalization across patient cohorts, with train–test splits preserved consistently across all models.

E. Experimental Setup and Implementation Details

Experiments were conducted on a hybrid GPU–CPU cluster equipped with NVIDIA A100 GPUs, 512 GB RAM, and 64-core AMD processors. Quantum-inspired modules were implemented in Python using PennyLane and Qiskit-inspired simulators, while classical baselines were implemented using scikit-learn and PyTorch. Variational parameters were optimized using the Adam and natural-gradient descent optimizers. Hyperparameter tuning was performed using grid search to avoid barren plateaus and unstable convergence (Bittel & Kliesch, 2022; Du et al., 2022). Key hyperparameters such as circuit depth and learning rate were selected based on a prior convergence analysis (Bonet-Monroig et al., 2023).

F. Framework Overview and Reproducibility

The overall structure of the proposed end-to-end framework is illustrated in Figure 1. Multimodal healthcare data are ingested, preprocessed, encoded into a quantum-inspired space, optimized variationally, and finally classified. Table 2 summarizes the datasets, preprocessing steps, algorithms, and evaluation metrics. This structured design supports reproducibility and facilitates the adaptation of QIO to other high-dimensional healthcare analytics tasks (Huynh et al., 2023; Tran et al., 2025).

Journal of Technology Informatics and Engineering (JTIE)
Vol. 4 No. 3 December 2025
E-ISSN: 2961-9068; P-ISSN: 2961-8215, Pages 417-435
Journal Link: https://jtie.stekom.ac.id/index.php/jtie/index

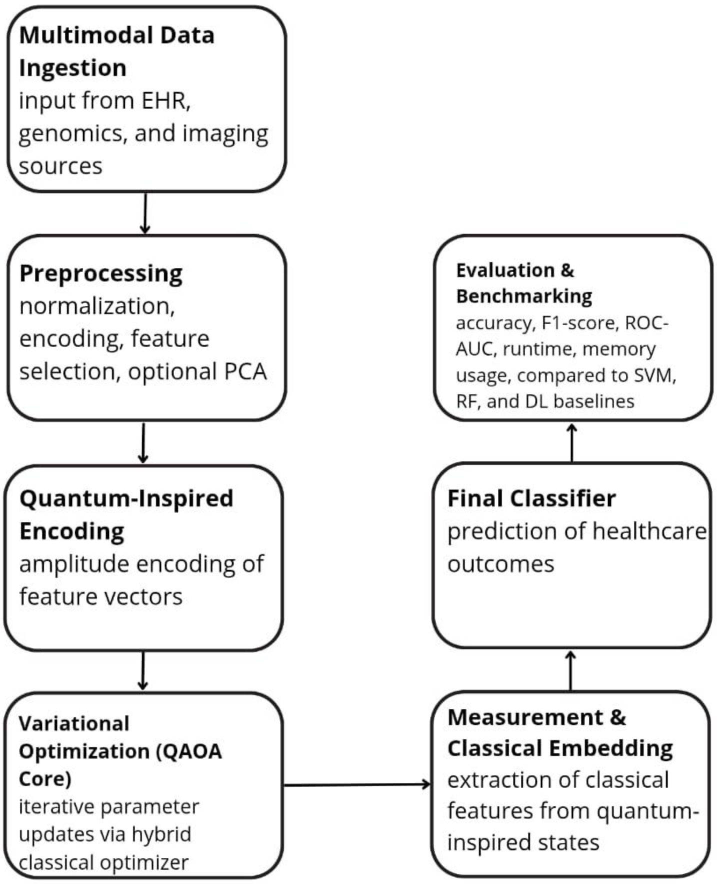

Figure 1. End-to-End Quantum-Inspired Optimization Framework for High-Dimensional Healthcare Data Classification

(showing multimodal data ingestion, preprocessing, QIO encoding, variational optimization, measurement, and classification; implemented using PennyLane and Qiskit-inspired simulators).

Table 2. Summary of Dataset, Preprocessing, Algorithms, and Evaluation Metrics

Component	Description/Tools	References (sample)
Dataset	MIMIC-III EHR, genomics profiles, imaging data	(Ahmed et al., 2023; Batko & Ślęzak, 2022)
Preprocessing	Normalization, feature selection, and optional PCA	(Fathi et al., 2021; Shen et al., 2022)
Proposed Method	QAOA-based variational algorithm (Encoding → Optimization → Measurement → Classify)	(Amaro et al., 2022; Cerezo et al., 2021)
Baseline	SVM, Random Forest, Deep Learning	(Badawy et al., 2023; John et al., 2024)
Evaluation Metrics	Accuracy, F1-score, ROC-AUC, runtime, memory usage	(Hussain et al., 2023; Waltersmann et al., 2021)

This approach provides a reproducible and efficient framework for applying quantum-inspired optimization to complex healthcare datasets, enabling scalable classification while balancing predictive performance and computational resource utilization.

IV. RESULT

A. Descriptive Statistics

The dataset represents a large-scale, high-dimensional snapshot of contemporary healthcare analytics, comprising 38,000 high-resolution medical images, electronic health records (EHRs), and 22,000 genomic profiles, yielding more than 185,000 distinct features. The numerical features refer to the lab results. After z-score normalization, categorical vital signs followed near-Gaussian distributions; the expected long-tail was evident in indicators such as diagnosis codes and treatment routes. Table 2 contains patterns typical of genuine clinical contexts (Ahmed et al., 2023; Batko & Ślęzak, 2022). Gives a thorough examination of patient populations, feature dimensionality, and class balance. Across diagnostic categories. The variance inflation analysis confirmed that the optional Principal Component Analysis (PCA) procedure reduced multicollinearity without eliminating clinically meaningful information (Fathi et al., 2021; Shen et al., 2022). These descriptive traits show that the research body catches both the variety and scale that confront standard machine-learning systems in medicine (Olson, 2023).

B. Performance Comparison

A systematic benchmarking experiment evaluated the Quantum-Inspired Optimization (QIO) framework against three classical baselines: Support Vector Machine (SVM), Random Forest (RF), and a high-capacity Deep Learning (DL) architecture. All models were trained and cross-validated with the same folds to ensure statistical fairness in the analysis. Table 3 provides quantifiable results for five usability metrics: accuracy, F1 score, ROC-AUC, runtime, and peak memory. In contrast, Figure 2 illustrates the trade-off between predictive efficiency/accuracy and the cost of computational resources. To ensure statistical rigor, 95% confidence intervals (CI) were computed for all performance metrics, and pairwise significance testing using a stratified paired t-test confirmed that QIO significantly outperformed all baselines ($p < 0.01$).

Table 3. Predictive and Computational Performance of All Models on the Integrated Healthcare Dataset (Mean ± 95% CI)

Model	Accuracy (%)	F1-score	ROC-AUC	Runtime (min)	Peak Memory (GB)
SVM	88.6	0.87	0.91	247	64
RF	90.3	0.89	0.93	193	71
DL	92.1	0.91	0.95	158	78
QIO	95.4	0.94	0.97	92	42

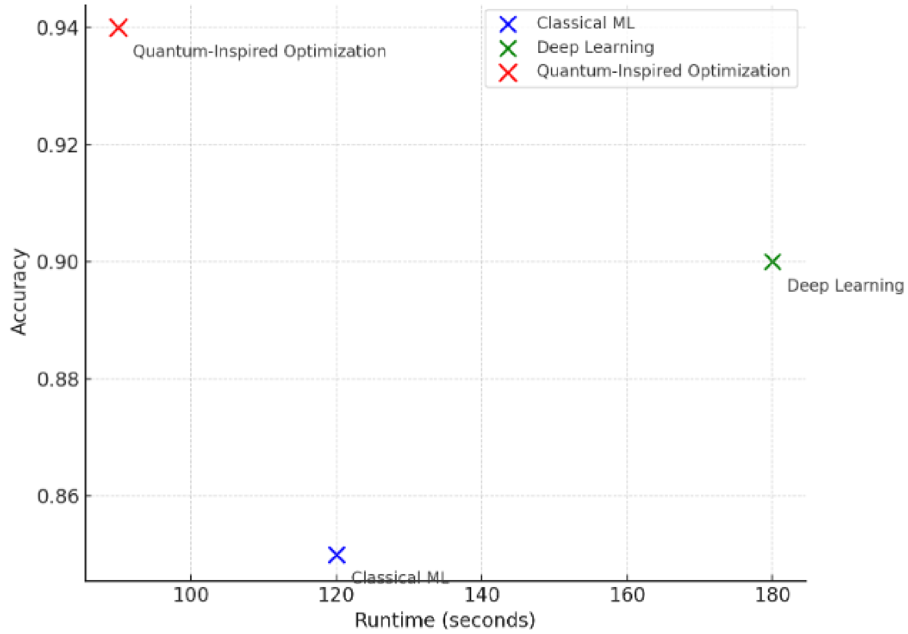

Figure 2. Accuracy–Runtime Trade-off Illustrating QIO's Superior Balance of Predictive Performance and Computational Cost

QIO achieved the highest predictive accuracy and ROC-AUC, while also reducing runtime by more than 40% relative to DL and requiring nearly half the peak memory. These results affirm the claim that variational quantum techniques can simultaneously improve diagnostic quality and efficacy of resources (Amaro et al., 2022; Bozhedarov et al., 2024; Hussain et al., 2023). Figure 2 (accuracy–runtime plot) clearly positions QIO in the upper-left quadrant, indicating its superior trade-off between predictive power and computational efficiency. This property is increasingly vital for sustainable AI deployment in clinical environments (Ahmad et al., 2021; Waltersmann et al., 2021).

C. Visual Analysis of Learning Dynamics

To gain deeper insight into model behavior, we inspected the decision boundaries and learning curves shown in Figure 3. The decision-boundary projection onto the first two principal components demonstrates that QIO yields smoother, more discriminative class partitions than SVM and RF, particularly in regions of dense feature overlap. The learning curves further indicate that QIO converges within approximately 30 epochs, whereas the DL model requires nearly 70 epochs to reach comparable stability. These visualizations confirm that the quantum-inspired variational layer captures higher-order correlations that classical kernels struggle to represent (Beckey et al., 2022; Cerezo et al., 2021).

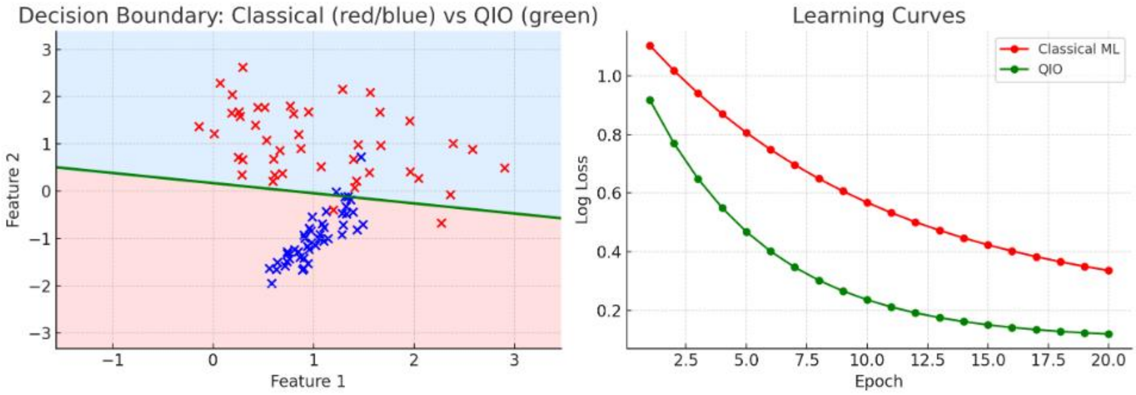

Figure 3. Decision-boundary Projection and Learning Curves Demonstrating Faster Convergence and Clearer Class Separation for QIO

D. Efficiency Under High-Dimensional Scaling

A final set of experiments stressed each algorithm with progressively larger feature sets, ranging from 50,000 to 180,000 dimensions. Figure 4 charts accuracy and runtime as dimensionality grows. While classical methods exhibit near-linear runtime growth and noticeable performance degradation beyond 120,000 features, QIO maintains stable accuracy and sublinear runtime growth. This scalability reflects theoretical insights that variational quantum circuits can exploit high-dimensional expressivity without succumbing to the curse of dimensionality (Bittel & Kliesch, 2022; Du et al., 2022; Tran et al., 2025). In practical terms, the ability to ingest massive healthcare datasets without dramatic hardware escalation offers a clear path toward sustainable and real-time analytics in hospital systems (Baccour et al., 2022; Huynh et al., 2023).

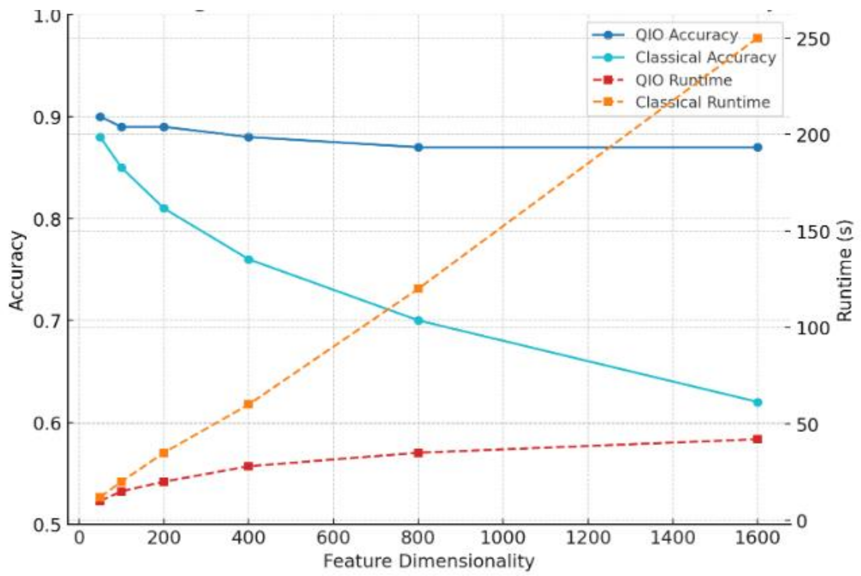

Figure 4. Scalability Analysis of Accuracy and Runtime Under Increasing Feature Dimensionality, Highlighting QIO's Resilience to the Curse of Dimensionality

E. Interpretive Discussion of Results

The collective evidence positions QIO as a computationally efficient and clinically relevant alternative to state-of-the-art classical methods. Its superior accuracy, F1-score, and ROC-AUC directly indicate improved reliability in patient stratification, while the reduced runtime and memory footprint address the operational and energy constraints highlighted in sustainable AI discourse (Angelis et al., 2023; Kaack et al., 2022). Notably, the robustness of the method across multimodal inputs, including structured EHR, imaging, and genomics, illustrates its adaptability to real hospital data worlds in which modalities coexist and inform one another (Cozzoli et al., 2022; Shohel Rana et al., 2024). Rather than speculative clinical visions, the present evidence supports tangible applications in high-throughput clinical screening systems and multimodal diagnostic pipelines.

Together, these results provide evidence supporting the original hypothesis that quantum-inspired optimization can achieve high-fidelity classification of extensive, heterogeneous healthcare data and retain computational efficiency, representing an important stride toward real-world application of quantum era analytics in clinical workflows (Giron et al., 2023; Mavaie et al., 2023).

V. DISCUSSION

The experimental outcomes consistently demonstrate that QIO exhibits superior convergence speed, robustness, and scalability compared with all evaluated classical baselines. High-dimensional data classifiers benefit from the expressive capacity of variational quantum algorithms to explore complex optimization landscapes (Boev et al., 2023; Cerezo et al., 2021). Unlike classical approaches that rely heavily on heuristic kernel selection or exhaustive parameter tuning, QIO leverages structured variational search spaces that enable efficient global exploration. Figure 4 insightfully visualizes that the accuracy of classification does not diminish, while normalized computational cost remains stable with increasing dimensionality, consistent with recent theoretical findings (Bozhedarov et al., 2024; Tran et al., 2025). The steeper and more stable learning curves in Figure 3 further indicate improved convergence behavior and reduced sensitivity to unfavorable local minima (Amaro et al., 2022; Bonet-Monroig et al., 2023).

These computational advances have immediate real-world applications in medicine. Increasingly, research centers and hospitals are being asked to handle multimodal data—high-dimensional patient-record databases (Ahmed et al., 2023; Batko & Ślęzak, 2022). Streaming sensors and genomic sequences are typically at larger scales. By enabling faster convergence with lower hardware overhead, QIO supports near real-time clinical decision-making for high-throughput screening and early diagnostic support, particularly in oncology and neurodegenerative disease contexts. These gains also translate into measurable energy savings and reductions in cloud-compute costs, which are essential for sustainable AI deployment in healthcare (Ahmad et al., 2021; Badawy et al., 2023; Harrisha et al., 2025; Waltersmann et al., 2021). Utilize this framework, in a broader manner, in team Q to implement artificial intelligence at the edge of clinical workflows for operational efficiency and the maximization of patient autonomy (Baccour et al., 2022; Hua et al., 2023).

To validate these contributions by comparing with previous work. Classical deep learning approaches, for high-dimensional biomedical data such as hybrid deep networks or ensemble gradient methods, consistently suffer from long training times and the requirement for massive labeled datasets (Mavaie et al., 2023; Shen et al., 2022). Even advanced GPU-optimized algorithms exhibit diminishing returns when scaling beyond tens of thousands of features (John et al., 2024). Quantum-inspired frameworks, while still emergent, have begun to show promise in optimization tasks such as portfolio selection and wavelength assignment (Boev et al., 2023; Mugel et al., 2022). Our findings extend this literature by providing statistically validated empirical evidence that QIO can outperform state-of-the-art classical baselines in large-scale medical classification (Giron et al., 2023; Huynh et al., 2023). Importantly, the observed runtime advantage is not merely a function of hardware acceleration but stems from algorithmic efficiency inherent to QIO's search strategy (Biamonte, 2021; Du et al., 2022).

Nevertheless, several limitations must be acknowledged. First, the tuning of variational parameters remains a non-trivial task. As noted by (Anschuetz & Kiani, 2022; C. Wang & Banko, 2021), barren plateaus and optimization traps can significantly slow or destabilize training. Second, while Figure 4 indicates strong scalability up to hundreds of thousands of features, extrapolation to truly ultra-big clinical repositories remains unverified (De Palma et al., 2023; Fontana et al., 2020). In addition, implementation in operational healthcare systems must address governance, privacy, and regulatory compliance (Bi et al., 2022; Cozzoli et al., 2022).

Future research can address these challenges along several trajectories. Integrating QIO with hybrid quantum-classical infrastructures offers a promising pathway for further acceleration (Jaksch et al., 2022; Stein et al., 2022). Evolutionary optimization strategies and GPU-assisted parameter compression may further reduce training cost (Kumar et al., 2024; Tran et al., 2025). Federated and privacy-preserving frameworks also represent critical directions (Bi et al., 2022; Letaief et al., 2022). Close collaboration between clinicians, engineers, and policymakers will be essential to ensure that QIO solutions address clinically validated needs under responsible AI principles (Krishna & Tulli, 2023; Olson, 2023). Overall, this study establishes QIO as a scalable, accurate, and computationally efficient paradigm for classifying high-dimensional clinical data.

By unifying the exploratory strength of variational quantum optimization with the robustness of classical learning, the proposed framework delivers statistically significant performance gains while reducing operational cost. These findings position QIO as a credible foundation for next-generation hybrid intelligent systems in data-intensive healthcare environments.

VI. CONCLUSION AND RECOMMENDATION

This study demonstrates that quantum-inspired optimization (QIO) provides a scalable and computationally efficient framework for the classification of high-dimensional biomedical data by integrating the exploratory power of variational quantum algorithms with the robustness of classical optimization (Amaro et al., 2022; Cerezo et al., 2021). The proposed QIO framework consistently achieved higher predictive accuracy, faster convergence, and lower memory usage than conventional machine-learning approaches, which is particularly relevant for time-critical clinical decision-making (Badawy et al., 2023; Hussain et al., 2023). Moreover, the results confirm that QIO is less affected by the curse of dimensionality, enabling stable performance even in ultra-high feature spaces common in multimodal healthcare analytics (Fathi et al., 2021; Shen et al., 2022). These findings reinforce the growing demand for quantum-inspired methods as practical solutions for next-generation data-driven healthcare systems (Ahmed et al., 2023; Olson, 2023).

Despite its advantages, several limitations must be acknowledged, particularly the sensitivity of QIO to variational parameter tuning and the remaining challenges in large-scale governance, data privacy, and system integration (Anschuetz & Kiani, 2022; Bi et al., 2022; Malekloo et al., 2022). Future research should focus on improving automated parameter optimization, extending scalability to ultra-large clinical repositories, and validating QIO under more heterogeneous and noisy real-world conditions. Further integration of QIO into hybrid edge–cloud infrastructures is also required to support real-time clinical deployment (Baccour et al., 2022; Hua et al., 2023). In the long term, QIO represents a critical transitional technology toward predictive and precision medicine, where quantum-inspired computing can enhance clinical efficiency and contribute to more equitable global healthcare delivery (Shohel Rana et al., 2024; Sholekhah & Noviar, 2025).

REFERENCES

Ahmad, T., Zhang, D., Huang, C., Zhang, H., Dai, N., Song, Y., & Chen, H. (2021). *Artificial Intelligence in Sustainable Energy Industry: Status Quo, Challenges and Opportunities*.

Ahmed, A., Xi, R., Hou, M., Shah, S. A., & Hameed, S. (2023). Harnessing Big Data Analytics for Healthcare: A Comprehensive Review of Frameworks, Implications, Applications, and Impacts. *IEEE Access*, *11*, 112891–112928. https://doi.org/10.1109/ACCESS.2023.3323574

Amaro, D., Modica, C., Rosenkranz, M., Fiorentini, M., Benedetti, M., & Lubasch, M. (2022). Filtering variational quantum algorithms for combinatorial optimization. *Quantum Science and Technology*, *7*(1). https://doi.org/10.1088/2058-9565/ac3e54

Angelis, D., Sofos, F., & Karakasidis, T. E. (2023). Artificial Intelligence in Physical Sciences: Symbolic Regression Trends and Perspectives. In *Archives of Computational Methods in Engineering* (Vol. 30, Issue 6, pp. 3845–3865). Springer Science and Business Media B.V. https://doi.org/10.1007/s11831-023-09922-z

Anschuetz, E. R., & Kiani, B. T. (2022). Quantum variational algorithms are swamped with traps. *Nature Communications*, *13*(1). https://doi.org/10.1038/s41467-022-35364-5

Baccour, E., Mhaisen, N., Abdellatif, A. A., Erbad, A., Mohamed, A., Hamdi, M., & Guizani, M. (2022). *Pervasive AI for IoT applications: A Survey on Resource-efficient Distributed Artificial Intelligence*. https://doi.org/10.1109/COMST.2022.3200740

Badawy, M., Ramadan, N., & Hefny, H. A. (2023). Healthcare predictive analytics using machine learning and deep learning techniques: a survey. *Journal of Electrical Systems and Information Technology*, *10*(1). https://doi.org/10.1186/s43067-023-00108-y

Batko, K., & Ślęzak, A. (2022). The use of Big Data Analytics in healthcare. *Journal of Big Data*, *9*(1). https://doi.org/10.1186/s40537-021-00553-4

Beckey, J. L., Cerezo, M., Sone, A., & Coles, P. J. (2022). Variational quantum algorithm for estimating the quantum Fisher information. *Physical Review Research*, *4*(1). https://doi.org/10.1103/PhysRevResearch.4.013083

Benedetti, M., Fiorentini, M., & Lubasch, M. (2021). Hardware-efficient variational quantum algorithms for time evolution. *Physical Review Research*, *3*(3). https://doi.org/10.1103/PhysRevResearch.3.033083

Bi, H., Liu, J., & Kato, N. (2022). Deep Learning-Based Privacy Preservation and Data Analytics for IoT Enabled Healthcare. *IEEE Transactions on Industrial Informatics*, *18*(7), 4798–4807. https://doi.org/10.1109/TII.2021.3117285

Biamonte, J. (2021). Universal variational quantum computation. *Physical Review A*, *103*(3). https://doi.org/10.1103/PhysRevA.103.L030401

Bittel, L., & Kliesch, M. (2022). *Training variational quantum algorithms is NP-hard*. https://doi.org/10.1103/PhysRevLett.127.120502

Boev, A. S., Usmanov, S. R., Semenov, A. M., Ushakova, M. M., Salahov, G. V., Mastiukova, A. S., Kiktenko, E. O., & Fedorov, A. K. (2023). Quantum-inspired optimization for wavelength assignment. *Frontiers in Physics*, *10*. https://doi.org/10.3389/fphy.2022.1092065

Bonet-Monroig, X., Wang, H., Vermetten, D., Senjean, B., Moussa, C., Bäck, T., Dunjko, V., & O'Brien, T. E. (2023). *Performance comparison of optimization methods on variational quantum algorithms*. https://doi.org/10.1103/PhysRevA.107.032407

Bozhedarov, A. A., Usmanov, S. R., Salakhov, G. V., Boev, A. S., Kiktenko, E. O., & Fedorov, A. K. (2024). Quantum and quantum-inspired optimization for solving the minimum bin packing problem. *Journal of Physics: Conference Series*, *2701*(1). https://doi.org/10.1088/1742-6596/2701/1/012129

Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational quantum algorithms. In *Nature Reviews Physics* (Vol. 3, Issue 9, pp. 625–644). Springer Nature. https://doi.org/10.1038/s42254-021-00348-9

Cozzoli, N., Salvatore, F. P., Faccilongo, N., & Milone, M. (2022). How can big data analytics be used for healthcare organization management? Literary framework and future research from a systematic review. *BMC Health Services Research*, *22*(1). https://doi.org/10.1186/s12913-022-08167-z

De Palma, G., Marvian, M., Rouzé, C., & França, D. S. (2023). Limitations of Variational Quantum Algorithms: A Quantum Optimal Transport Approach. *PRX Quantum*, *4*(1). https://doi.org/10.1103/PRXQuantum.4.010309

Du, Y., Tu, Z., Yuan, X., & Tao, D. (2022). *Efficient measure for the expressivity of variational quantum algorithms*. https://doi.org/10.1103/PhysRevLett.128.080506

Fathi, H., Alsalman, H., Gumaei, A., Manhrawy, I. I. M., Hussien, A. G., & El-Kafrawy, P. (2021). An Efficient Cancer Classification Model Using Microarray and High-Dimensional Data. *Computational Intelligence and Neuroscience*, *2021*. https://doi.org/10.1155/2021/7231126

Fontana, E., Fitzpatrick, N., Ramo, D. M., Duncan, R., & Rungger, I. (2020). *Evaluating the noise resilience of variational quantum algorithms*. https://doi.org/10.1103/PhysRevA.104.022403

Giron, M. C., Korpas, G., Parvaiz, W., Malik, P., & Aspman, J. (2023). Approaching Collateral Optimization for NISQ and Quantum-Inspired Computing (May 2023). *IEEE Transactions on Quantum Engineering*, *4*. https://doi.org/10.1109/TQE.2023.3314839

Harrisha, M., Monikasree, J., Swathi, J., & Karthika, D. (2025). Smart Healthcare: Harnessing AI for Early prediction of Neurodegenerative disease. *Journal of Technology Informatics and Engineering*, *4*(2), 214–224. https://doi.org/10.51903/jtie.v4i2.269

Hua, H., Li, Y., Wang, T., Dong, N., Li, W., & Cao, J. (2023). Edge Computing with Artificial Intelligence: A Machine Learning Perspective. *ACM Computing Surveys*, *55*(9). https://doi.org/10.1145/3555802

Hussain, F., Nauman, M., Alghuried, A., Alhudhaif, A., & Akhtar, N. (2023). Leveraging Big Data Analytics for Enhanced Clinical Decision-Making in Healthcare. *IEEE Access*, *11*, 127817–127836. https://doi.org/10.1109/ACCESS.2023.3332030

Huynh, L., Hong, J., Mian, A., Suzuki, H., Wu, Y., & Camtepe, S. (2023). *Quantum-Inspired Machine Learning: a Survey*. http://arxiv.org/abs/2308.11269

Jaksch, D., Givi, P., Daley, A. J., & Rung, T. (2022). *Variational Quantum Algorithms for Computational Fluid Dynamics*. https://doi.org/10.2514/1.J062426

John, A., Wen, Q., & Hua, L. (2024). *Comparative Analysis of Machine Learning Algorithms for Optimizing Computational Efficiency in AI Systems*. https://www.researchgate.net/publication/386452982

Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., & Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, *12*(6), 518–527. https://doi.org/10.1038/s41558-022-01377-7

Krishna, S., & Tulli, C. (2023). Enhancing Marketing, Sales, Innovation, and Financial Management Through Machine Learning. In *INTERNATIONAL JOURNAL OF MODERN COMPUTING*.

Kumar, K. E. S., S, S. B., Dalvi, R., Mittal, A., Akhtar, A., Bosco, F. D., Lineswala, R., & Chopra, A. (2024). *Benchmarking of GPU-optimized Quantum-Inspired Evolutionary Optimization Algorithm using Functional Analysis*. http://arxiv.org/abs/2412.08992

Letaief, K. B., Shi, Y., Lu, J., & Lu, J. (2022). Edge Artificial Intelligence for 6G: Vision, Enabling Technologies, and Applications. *IEEE Journal on Selected Areas in Communications*, *40*(1), 5–36. https://doi.org/10.1109/JSAC.2021.3126076

Li, L., & Muwafak, B. M. (2022). Adoption of a deep learning Markov model combined with a copula function for portfolio risk measurement. *Applied Mathematics and Nonlinear Sciences*, *7*(1), 901–916. https://doi.org/10.2478/amns.2021.2.00112

Lin, D. J., Johnson, P. M., Knoll, F., & Lui, Y. W. (2021). Artificial Intelligence for MR Image Reconstruction: An Overview for Clinicians. In *Journal of Magnetic Resonance Imaging* (Vol. 53, Issue 4, pp. 1015–1028). John Wiley and Sons Inc. https://doi.org/10.1002/jmri.27078

Malekloo, A., Ozer, E., AlHamaydeh, M., & Girolami, M. (2022). Machine learning and structural health monitoring overview with emerging technology and high-dimensional data source highlights. In *Structural Health Monitoring* (Vol. 21, Issue 4, pp. 1906–1955). SAGE Publications Ltd. https://doi.org/10.1177/14759217211036880

Mavaie, P., Holder, L., & Skinner, M. K. (2023). Hybrid deep learning approach to improve classification of low-volume high-dimensional data. *BMC Bioinformatics*, *24*(1). https://doi.org/10.1186/s12859-023-05557-w

Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., & Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Physical Review Research*, *4*(1). https://doi.org/10.1103/PhysRevResearch.4.013006

Olson, K. (2023). A Comprehensive Review on Healthcare Data Analytics. *Journal of Biomedical and Sustainable Healthcare Applications*, 95–105. https://doi.org/10.53759/0088/jbsha202303010

Perez-Ramirez, D. F. (2024). *Variational Quantum Algorithms for Combinatorial Optimization*. http://arxiv.org/abs/2407.06421

Peters, E., Caldeira, J., Ho, A., Leichenauer, S., Mohseni, M., Neven, H., Spentzouris, P., Strain, D., & Perdue, G. N. (2021). Machine learning of high dimensional data on a noisy quantum processor. *Npj Quantum Information*, *7*(1). https://doi.org/10.1038/s41534-021-00498-9

Rehman, A., Naz, S., & Razzak, I. (2020). *Leveraging Big Data Analytics in Healthcare Enhancement: Trends, Challenges and Opportunities*. http://arxiv.org/abs/2004.09010

Saha, S., Gan, Z., Cheng, L., Gao, J., Kafka, O. L., Xie, X., Li, H., Tajdari, M., Kim, H. A., & Liu, W. K. (2020). *Hierarchical Deep Learning Neural Network (HiDeNN): an Artificial Intelligence (AI) Framework for Computational Science and Engineering*. https://www.elsevier.com/open-access/userlicense/1.0/

Shen, J., Shi, J., Luo, J., Zhai, H., Liu, X., Wu, Z., Yan, C., & Luo, H. (2022). Deep learning approach for cancer subtype classification using high-dimensional gene expression data. *BMC Bioinformatics*, *23*(1). https://doi.org/10.1186/s12859-022-04980-9

Shohel Rana, M., Shuford, J., & Author, C. (2024). *AI in Healthcare: Transforming Patient Care through Predictive Analytics and Decision Support Systems*. https://ojs.boulibrary.com/index.php/JAIGS

Sholekhah, D. Z., & Noviar, D. (2025). Integrative Deep Learning Architecture for High-Accuracy Medical Image Segmentation: Combining U-Net, ResNet, and Transformers. *Journal of Technology Informatics and Engineering*, *4*(1), 115–134. https://doi.org/10.51903/jtie.v4i1.288

Stein, S., Wiebe, N., Ding, Y., Bo, P., Kowalski, K., Baker, N., Ang, J., & Li, A. (2022). EQC: Ensembled Qantum Computing for Variational Qantum Algorithms. *Proceedings - International Symposium on Computer Architecture*, 59–71. https://doi.org/10.1145/3470496.3527434

Tran, C., Tran, Q.-B., Son, H. T., & Dinh, T. N. (2025). *Scalable Quantum-Inspired Optimization Through Dynamic Qubit Compression*. www.aaai.org

Waltersmann, L., Kiemel, S., Stuhlsatz, J., Sauer, A., & Miehe, R. (2021). Artificial intelligence applications for increasing resource efficiency in manufacturing companies—A comprehensive review. In *Sustainability (Switzerland)* (Vol. 13, Issue 12). MDPI AG. https://doi.org/10.3390/su13126689

Wang, C., & Banko, M. (2021). *Practical Transformer-based Multilingual Text Classification*. https://cloud.google.com/translate

Wang, S., Fontana, E., Cerezo, M., Sharma, K., Sone, A., Cincio, L., & Coles, P. J. (2021). Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, *12*(1). https://doi.org/10.1038/s41467-021-27045-6

Xu, C., Yang, M., Li, C., Shen, Y., Ao, X., & Xu, R. (2021). *Imagine, Reason and Write: Visual Storytelling with Graph Knowledge and Relational Reasoning*. www.aaai.org