Review

Hybrid Quantum-Classical Architectures in Medical Imaging: A Taxonomy-Based Survey of COVID-19 Models

Seyedeh Aram Salehi , Hanieh Naderi * , Seyyed Amir Asghari * , Javad Chaharlang , Yvon Savaria

Posted Date: 6 May 2026

doi: 10.20944/preprints202605.0229.v1

Keywords: quantum-enhanced medical imaging; hybrid classical–quantum architectures; COVID-19 classification; variational quantum circuits; quantum encoding; NISQ computing; systematic review

 doi:10.20944/preprints202605.0229.v1

Review

Hybrid Quantum-Classical Architectures in Medical Imaging: A Taxonomy-Based Survey of COVID-19 Models

Seyedeh Aram Salehi [1], Hanieh Naderi [1,*], Seyyed Amir Asghari [1,*], Javad Chaharlang [2] and Yvon Savaria [3]

[1] School of Intelligent Systems Engineering, College of Interdisciplinary Science and Technology, University of Tehran

[2] Department of Computer Engineering, Dezful Branch, Islamic Azad University, Dezful, Iran

[3] Department of Department of Electircal and Computer Engineering, Polytechnique Montreal, Montreal, QC, Canada

* Correspondence: hanieh.naderi@ut.ac.ir (H.N.); s.a.asghari@ut.ac.ir (S.A.A.)

Highlights

This survey provides a structured, taxonomy-driven analysis of hybrid classical–quantum frameworks for COVID-19 medical imaging, emphasizing architectural design choices, empirical performance patterns, and translational limitations.

What are the main findings?

- Architectural roles for quantum modules (feature extraction, front-end transformation, and classification), with feature-extraction-based designs showing relatively stronger performance in reported results.
- Most studies report promising results in binary COVID-19 classification under simulation settings, while multi-class performance remains less stable and less frequently evaluated.
- The current literature lacks a unified architecture-level framework that systematically relates quantum module placement to empirical performance across COVID-19 imaging studies.

What are the implications of the main findings?

- Current hybrid quantum–classical (HQC) models should be regarded as early-stage exploratory approaches rather than clinically validated solutions.
- Reported performance gains may not generalize to real-world settings due to lack of standardized benchmarking and hardware-based evaluation.
- Future work should prioritize hardware-aware, resource-efficient hybrid designs validated under NISQ constraints and standardized evaluation protocols.

Abstract

This paper surveys the emerging field of hybrid quantum–classical (HQC) architectures for medical image analysis, with a particular focus on COVID-19 classification as a representative application. Despite extensive progress in quantum machine learning, there is still a lack of structured taxonomy-driven analysis that systematically characterizes how quantum components are integrated within hybrid imaging pipelines. This study provides a structured review of representative HQC models, categorizing existing approaches based on the functional role of quantum modules, including front-end processing, feature extraction, and classification layers. A comparative analysis is conducted to examine how these architectural design choices relate to reported performance across different experimental settings. The review shows that hybrid models can achieve strong binary classification performance under simulation conditions, particularly when quantum circuits are used as feature extractors. However, performance in multi-class scenarios remains inconsistent and highly sensitive to encoding strategies, circuit design, and optimization settings. In addition, most existing studies rely on simulator-based evaluations, limited or single-source datasets, and non-standardized

 doi:10.20944/preprints202605.0229.v1

benchmarking protocols, which restrict reproducibility and generalization. Real-hardware validation remains scarce, and no consistent evidence of task-level quantum advantage over optimized classical baselines has been demonstrated. Overall, current HQC-based approaches should be considered exploratory rather than clinically validated solutions. Future progress requires standardized benchmarking, multi-institutional validation, hardware-aware evaluation, and the development of interpretable hybrid architectures that respect NISQ-era constraints.

Keywords: quantum-enhanced medical imaging; hybrid classical–quantum architectures; COVID-19 classification; variational quantum circuits; quantum encoding; NISQ computing; systematic review

1. Introduction

The COVID-19 pandemic, declared a global health emergency by the World Health Organization in March 2020, exposed structural vulnerabilities in large-scale diagnostic infrastructures [1]. Although reverse transcription polymerase chain reaction (RT-PCR) remains the clinical gold standard, its diagnostic performance is constrained by suboptimal early sensitivity, processing delays, and dependence on specialized laboratory facilities [2]. These limitations motivated the expanded use of medical imaging, particularly chest computed tomography (CT) and chest X-ray (CXR)—as complementary modalities for screening, triage, and disease monitoring. CT demonstrates high sensitivity for hallmark manifestations, such as ground-glass opacities and multifocal consolidations [3,4], while CXR offers greater accessibility and cost-effectiveness for large-scale deployment [3]. Artificial intelligence (AI), especially deep learning based on convolutional neural networks (CNNs), has been extensively explored to automate COVID-19 image-based diagnosis. Transfer learning with pretrained backbones such as ResNet, DenseNet, VGG, and EfficientNet has yielded strong performance across CT and CXR datasets [4–8].

Nevertheless, persistent methodological and translational limitations remain. Dataset scarcity and class imbalance increase the risk of overfitting and inflated performance estimates [2,8]. Generalization across imaging centers, scanners, and acquisition protocols is often weak [5,9]. Large-scale or three-dimensional CT models introduce substantial computational and memory overhead [10], while the limited interpretability of deep neural networks continues to hinder clinical trust and regulatory acceptance [11]. These challenges highlight the need for alternative or complementary computational paradigms that enhance representation efficiency and robustness without excessive architectural scaling.

Quantum machine learning (QML) has recently been proposed as a complementary paradigm that can enrich feature representations through high-dimensional Hilbert-space embeddings [12]. Parameterized quantum circuits, in principle, may model complex nonlinear correlations within comparatively compact parameter spaces. However, current quantum devices operate in the noisy intermediate-scale quantum (NISQ) regime, characterized by restricted qubit counts, shallow circuit depths, limited connectivity, and sensitivity to noise and decoherence [13]. These constraints render end-to-end quantum processing of high-resolution medical images impractical. To reconcile theoretical expressivity with hardware limitations, hybrid quantum–classical (HQC) frameworks have emerged. These approaches embed compact quantum modules—such as variational quantum circuits, quantum convolutional operators, or quantum classifier heads—within classical deep learning pipelines [14–16]. Early investigations report parameter efficiency and encouraging binary classification performance, particularly under simulator-based evaluation. However, multi-class diagnostic robustness, real-hardware validation, and external multi-site testing remain limited. Moreover, reported improvements are often sensitive to encoding choices, circuit depth, and optimization stability.

Despite the growing number of HQC applications in COVID-19 imaging, the literature lacks a unified architectural framework to systematically organize and compare these approaches. Most contributions are presented as standalone algorithmic proposals, with an emphasis on performance

 doi:10.20944/preprints202605.0229.v1

metrics rather than on structural integration within the end-to-end learning pipeline. Consequently, a fundamental design question remains under-examined: what functional role does the quantum component play within hybrid systems, and how does its architectural placement influence representational capacity, optimization behavior, computational cost, and scalability under NISQ constraints?

Several survey articles have either broadly reviewed quantum machine learning or summarized artificial intelligence techniques for COVID-19 diagnosis. As summarized in Table 1, existing reviews predominantly focus on algorithmic categories, cross-domain applications, or high-level trends. While some mention hybrid quantum–classical approaches or COVID-19 imaging, none provide a dedicated architecture-centric analysis of how quantum computation is structurally integrated into medical imaging pipelines for COVID-19 classification. Prior surveys do not explicitly distinguish between the operational roles and placements of quantum modules within hybrid learning systems, nor do they analyze how architectural design decisions affect empirical performance and deployment feasibility. To address this gap, the present survey introduces an architecture-centric taxonomy of hybrid quantum–classical models for COVID-19 medical imaging. As illustrated in Figure 1, existing approaches can be grouped into three archetypes: (i) quanvolution-based hybrids, where quantum circuits act as localized nonlinear feature transformers; (ii) hybrid architectures with quantum decision modules, in which quantum circuits function as classifier heads operating on classical embeddings; and (iii) quantum feature-extractor designs, where quantum circuits generate discriminative feature representations subsequently processed by lightweight classical classifiers. This taxonomy enables structured comparison across architectural placement, encoding strategy, quantum resource requirements, validation rigor, and reported performance outcomes. Rather than treating hybrid models as isolated algorithmic variants, this structured architectural perspective clarifies trade-offs among expressivity, optimization complexity, computational overhead, and translational readiness. It further explains why certain hybrid configurations exhibit competitive binary screening performance yet struggle with multi-class generalization and clinical robustness.

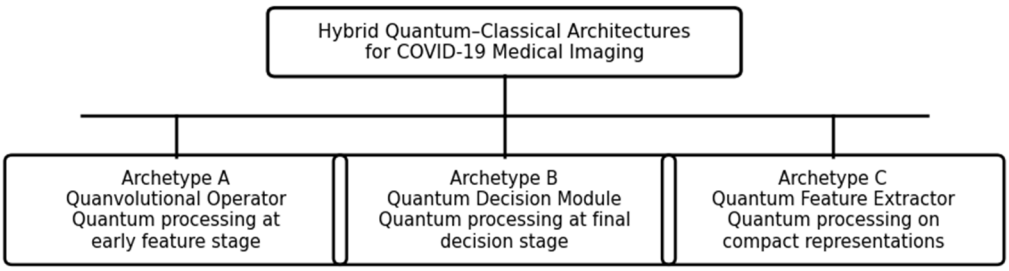

Figure 1. This survey proposes a conceptual taxonomy of hybrid quantum–classical architectures for COVID-19 medical imaging. The three archetypes are distinguished by the functional role and placement of quantum processing within the learning pipeline: (A) quanvolution-based architectures applying quantum computation at the early feature extraction stage, (B) hybrid models employing quantum circuits as decision or classification modules, and (C) architectures using quantum circuits as feature extractors operating on compact classical representations.

 doi:10.20944/preprints202605.0229.v1

Table 1. Recent survey articles on quantum machine learning in healthcare and medical imaging.

Ref.	Survey Article	Venue / Year	Main Scope	COVID-19 Imaging Explicitly Covered	Hybrid Classical–Quantum Models	Architecture / Pipeline-Level Analysis	Limitation w.r.t. This Survey
[17]	Quantum Machine and Deep Learning for Medical Image Classification: A Systematic Review of Trends, Methodologies, and Future Directions – E.A. Radhi et al.	Iraqi Journal for Computer Science and Mathematics, 2025	QML and deep learning for medical image classification	Yes (CT, CXR)	Mentioned	No	Does not provide architecture-level integration analysis of HQC pipelines for COVID-19 imaging
[18]	Quantum Machine Learning Revolution in Healthcare: A Systematic Review of Emerging Perspectives and Applications – U. Ullah et al.	IEEE Access, 2024	QML across healthcare applications	Partial	General discussion	No	Broad healthcare scope; lacks focused medical imaging pipeline dissection
[19]	Review of Medical Image Processing Using Quantum-Enabled Algorithms – F. Yan et al.	Artificial Intelligence Review (Springer), 2024	Quantum-enabled medical image processing	Mentioned	Conceptual	No	Descriptive overview without structured hybrid architectural categorization
[20]	A Systematic Review on Quantum Machine Learning Applications in Classification – E. Mohammadisavadkoohi et al.	IEEE Transactions on Artificial Intelligence, 2025	QML for classification tasks	Mentioned	Algorithm-level	No	Not specifically focused on medical imaging pipeline integration
[21]	Leveraging Quantum Computing to Revolutionize Deep Learning: A Focus on Hybrid Algorithms for Medical Image Classification – H. Chetioui	Institutional Repository / Preprint, 2025	Hybrid QML–DL for medical imaging	Partial	General discussion	No	Discusses hybrid models but does not introduce a unified architectural taxonomy
[22]	Machine Learning for Advancing Quantum Field Theory-Enabled Sensor Networks in Precision Disease Detection: A Review – P. Khanal et al.	Crimson Journal of Science and Technology, 2025	QML for quantum sensing and disease detection	No	General	No	Focuses on sensing systems rather than imaging-based diagnostic pipelines
[23]	Quantum Machine Learning: Recent Advances, Challenges and Perspectives – P. Lamichhane & D.B. Rawat	IEEE Access, 2025	Cross-domain QML review	Mentioned	General	No	No dedicated medical imaging or architecture-centric HQC analysis

 doi:10.20944/preprints202605.0229.v1

[24]	Quantum Machine Learning Models in Healthcare: Future Trends and Challenges – A. Sonavane et al.	Elsevier (Book Chapter), 2025	QML for healthcare diagnostics	Partial	Conceptual	No	Healthcare-oriented overview with limited emphasis on imaging pipelines
[25]	Advancements and Challenges in Quantum Machine Learning for Medical Image Classification: A Comprehensive Review – M.F. Shahriyar et al.	IEEE Conference / arXiv, 2025	QML for medical image classification	Partial	Mentioned	No	Does not provide COVID-19–focused hybrid architecture synthesis
This work	Hybrid Quantum–Classical Architectures for COVID-19 Medical Imaging: A Survey	—	Architecture-centric hybrid learning for COVID-19 imaging	Core focus	Core focus	Yes (three archetypes)	Introduces an architecture-centric taxonomy of HQC models for COVID-19 imaging

 doi:10.20944/preprints202605.0229.v1

The remainder of this article is organized as follows: Section 2 reviews related work. Section 3 outlines the theoretical foundations of hybrid quantum–classical computation. Section 4 presents a detailed technical analysis of the three architectural archetypes. Section 5 provides a comparative synthesis supported by consolidated tables. Section 6 discusses key methodological and translational challenges. Section 7 outlines future research directions. Finally, Section 8 concludes with central insights and implications for the development of clinically credible hybrid quantum–classical diagnostic systems.

2. Review of Prior Studies

The literature relevant to hybrid classical–quantum COVID-19 imaging spans four intertwined streams: (i) classical deep-learning systems for COVID-19 diagnosis from CXR/CT, (ii) quantum machine learning foundations and medical-imaging adaptations, (iii) hybrid quantum–classical learning in broader healthcare decision support, and (iv) hybrid quantum–classical pipelines developed specifically for COVID-19 imaging. Across these streams, reported gains often depend as much on dataset construction and evaluation protocol (e.g., patient-wise splitting and external validation) as on the chosen model family, making rigorous, dataset-aware synthesis essential.

2.1. Classical AI for COVID-19 Medical Imaging

Classical COVID-19 imaging diagnosis is dominated by transfer learning and CNN-based classification/segmentation pipelines evaluated on public CXR and CT datasets. A major CXR benchmark effort introduced an open-access dataset and a tailored CNN design, reporting strong COVID-19 detection performance and emphasizing transparency in dataset curation; the benchmark scale (CXR images and patient cases) is explicitly documented and widely reused for subsequent comparisons [8]. Subsequent classical studies built CXR classifiers with various backbones and transfer-learning strategies, typically reporting high accuracy/sensitivity on multi-class settings (COVID vs normal vs pneumonia) but showing sensitivity to class imbalance and dataset origin; representative works include robust detection frameworks and transfer-learning-based pipelines that report strong performance under standard splits [26]. For CT-based diagnosis, several influential works created large-scale CT benchmarks and tailored CNN architectures, explicitly reporting dataset sizes and patient counts. One benchmark is documented at the scale of 104,009 CT slices across 1,489 patients (COVIDx-CT), enabling reproducible evaluation and consistent reporting across studies using this dataset [27]. In parallel, smaller but highly cited public CT collections were introduced to catalyze research when large CT datasets were scarce; for example, an open dataset reports 349 COVID CT images from 216 patients and 463 non-COVID CT images, with baseline experiments and metrics reported for reference [28]. Another widely used CT dataset reports 2,482 CT scans split between 1,252 COVID-positive and 1,230 non-COVID cases; multiple classical CT pipelines (including end-to-end DL and feature-engineering hybrids) report very high classification accuracy on this dataset, sometimes approaching ~99% under certain experimental settings Additional CT studies emphasize clinical slice selection (radiologist-guided informative-slice selection) and curated institutional datasets, reporting improved reliability but at the cost of reduced comparability due to private data and differing protocols [29]. Across these classical studies, high headline accuracy is common, but methodological reviews and critiques show that results can be inflated by non–patient-wise splitting, hidden dataset leakage, and spurious correlations arising from multi-source aggregation (scanner/site artifacts). These issues directly affect clinical generalization and must be treated as first-order variables when judging any "improvement," including quantum/hybrid claims [8].

2.2. Quantum Machine Learning in Medical Imaging

NISQ realities constrain quantum machine learning (QML) for imaging: limited qubit count, shallow depth, noise, and expensive sampling, which make direct end-to-end quantum image

 doi:10.20944/preprints202605.0229.v1

modeling impractical for medical images without dimensionality reduction or patch-wise encodings. A widely adopted conceptual foundation treats quantum circuits as feature maps into high-dimensional Hilbert spaces, motivating quantum kernel learning and variational quantum classifiers as the dominant building blocks for classification [30]. At the same time, trainability limitations (e.g., barren plateaus) explain why many medical-imaging QML models remain shallow and hybridized, and why initialization and circuit design strategies matter as much as the choice of optimizer [31].

To make image inputs feasible, patch-wise "quanvolution" proposes running small random quantum circuits on local patches, then feeding measured features to a classical network; experiments on standard image benchmarks report improved test accuracy and/or faster learning versus purely classical baselines under the same experimental protocol, motivating its use as a drop-in feature layer for near-term devices [32]. Beyond MNIST-style demonstrations, medical-image–specific benchmarking has expanded: one comprehensive study benchmarks multiple MedMNIST medical imaging tasks on real quantum hardware (IBM-class devices), explicitly investigating feasibility, device-aware circuits, and error mitigation in medical image classification [33]. More recent hybrid medical-image architectures report strong results across multiple MedMNIST v2 datasets using compact variational circuits (e.g., reporting up to ~99.9% accuracy/AUC on some tasks and robustness patterns on noisier subsets), emphasizing that hybrid CNN backbones plus small VQC heads can remain competitive while limiting the quantum footprint [34]. In ophthalmic imaging, hybrid/federated quantum CNN approaches achieve high retinal classification accuracy (e.g., ~96.8% on E-Ophtha), indicating a shift toward privacy-aware distributed training with quantum components [35]. Noise-aware hybrid QML is also entering top-tier medical imaging venues, proposing mechanisms for robustness to explicit noise levels and ablations that quantify the contributions of mitigation modules [36]. Parallel lines benchmark quantum image classifiers at scale on classical datasets (e.g., full MNIST) to stress-test circuit expressivity and training behavior, providing evidence that hybrid designs can match or exceed classical accuracy under constrained settings [37]. Finally, domain-specific hybrid quantum pipelines are being explored for dermatology/lesion classification and related tasks, typically combining transfer learning with quantum branches to address limited data and improve representation diversity[38]. The strongest QML imaging papers increasingly report not only accuracy/AUC, but also quantum resource realism (qubits, depth, sampling/shots, device/hardware vs simulator). This is crucial because some "accuracy gains" can disappear when circuit evaluation costs, shot noise, or hardware noise are accounted for; conversely, work that benchmarks on real devices or uses explicit noise models provides more reliable evidence for medical translation [33].

2.3. Hybrid Classical–Quantum Frameworks in Healthcare

Hybrid quantum–classical learning in healthcare extends beyond imaging to clinical decision support, bio-signal interpretation, and risk prediction, where tabular EHR-like datasets are common, and dimensionality can be controlled. Broad healthcare reviews synthesize applications spanning diagnosis assistance, imaging, genomics, and operational optimization. They consistently frame hybrid designs as the most viable path under NISQ constraints [39]. A systematic review of digital health assesses whether QML (including hybrid models) outperforms classical methods for clinical decision-making and health service delivery, summarizes thousands of screened studies, and concludes that a consistent advantage has not yet been established due to heterogeneous baselines, proxy datasets, and uneven reporting [14]. In heart-disease prediction, hybrid/QML evaluations on standard cardiology datasets report incremental accuracy improvements and/or faster training compared to classical baselines under particular settings, highlighting sensitivity to feature selection and evaluation choices [40]. For diabetes prediction on the Pima Indians Diabetes Dataset, multiple QML studies test variational quantum classifiers and quantum support vector approaches, reporting performance competitive with classical ML but often below strong deep baselines; some works emphasize that optimizer choice, feature maps (e.g., ZZ), and error-aware tuning can noticeably change reported metrics [41]. In bio signal analysis, hybrid classical–quantum ECG/arrhythmia

 doi:10.20944/preprints202605.0229.v1

pipelines propose quantum transforms or quantum-enhanced feature stages followed by classical learners, often evaluated on large ECG corpora (e.g., MIT-BIH style settings) and reporting high accuracy in arrhythmia classification when paired with strong classical feature extractors [42]. In oncology-style classification, breast cancer benchmarks (e.g., WDBC-type datasets) are repeatedly used to test VQC/QSVM variants and compare them against classical baselines; these studies show that careful circuit/ansatz design and tuning are necessary for stability, and that improvements can be small or dataset-dependent [43]. In healthcare decisioning, the most credible hybrid studies (i) evaluate on widely recognized clinical datasets, (ii) report multiple clinically meaningful metrics (not only accuracy), and (iii) explicitly state quantum resources and whether results are simulator- or hardware-based. Systematic evidence indicates that many "advantages" are not yet robust across datasets or protocols, underscoring the need for standardized reporting and stronger external validation—lessons that directly translate to COVID imaging [14].

2.4. Hybrid Quantum–Classical Models for COVID-19 Imaging

Early hybrid quantum–classical (HQC) models for COVID-19 imaging predominantly explored patch-wise quantum processing via quantum convolutional designs. A representative framework is introduced in [44], in which a random quantum circuit–based quanvolutional layer is embedded in a classical convolutional neural network. In this approach, small chest X-ray patches are angle-encoded into shallow quantum circuits, and the resulting expectation values are aggregated as quantum-enhanced feature maps for subsequent classical convolutional layers. The model is evaluated across multiple chest X-ray datasets with sample sizes ranging from approximately 7,000 to 11,000 images, in both binary and multi-class classification settings. While binary COVID–normal classification achieves accuracy exceeding 98%, performance in the multi-class scenario (COVID, normal, pneumonia) degrades to the high-80% range, indicating reduced separability among visually similar classes. A closely related random quanvolution–based hybrid CNN is further investigated in [45], with a more systematic evaluation across multiple classification scenarios. In this design, the quantum layer explicitly replaces the first classical convolutional stage and feeds into a conventional CNN backbone. Experiments conducted on chest X-ray datasets for binary COVID–normal and multi-class COVID–pneumonia–normal classification report binary accuracy around 98% and multi-class accuracy in the range of 90–93%. These results suggest that random quantum feature maps can enhance early-stage representation learning, but their contribution diminishes as classification complexity increases. Moving beyond fixed random circuits, a parameterized quanvolutional variant is proposed in [46], where trainable quantum convolutional layers are integrated into a hybrid CNN. Image patches are encoded using rotation gates into variational quantum circuits, whose parameters are optimized jointly with classical network weights. Evaluated on chest X-ray datasets for binary COVID-19 detection, the model reports simulator-based accuracy above 97%. However, the introduction of trainable quantum parameters increases sensitivity to optimization instability, particularly under shallow circuit and limited-qubit constraints.

In contrast to patch-wise quantum processing, [47] introduces a quantum feature extraction architecture with a compact classifier (HQF-CC). In this framework, chest X-ray images from a large radiography dataset with over 15,000 samples are first processed using classical feature-extraction pipelines. The resulting compact feature vectors are then encoded into a variational quantum circuit to generate quantum-enhanced representations, which are classified using a lightweight custom classical classifier. Reported test accuracy reaches approximately 98.8%, with COVID-class sensitivity around 88–89% and specificity above 97%, demonstrating that high aggregate accuracy may coexist with uneven class-wise performance. Hybrid quantum transfer learning for CT imaging is examined in [48], in which a pretrained classical CNN backbone extracts high-dimensional CT embeddings, which are subsequently processed by a dressed quantum circuit (DQC) classifier. Multi-class CT classification involving COVID-19, community-acquired pneumonia, and normal cases is evaluated using reduced feature vectors encoded into 4–8 qubits via angle encoding. The results show that the DQC-based hybrid achieves competitive multi-class accuracy while maintaining more stable training

 doi:10.20944/preprints202605.0229.v1

behavior than purely quantum classifier heads, particularly under restricted qubit budgets. A complementary CT-based hybrid transfer-learning study is presented in [49], focusing on the impact of encoding strategies in variational quantum classifiers. Classical CNN features are embedded into quantum circuits using either amplitude encoding or angle encoding. Experimental results on CT datasets for multi-class COVID-related diagnosis indicate that angle encoding yields more stable, consistent performance, whereas amplitude encoding exhibits higher variance and greater optimization difficulty despite its higher representational capacity. Beyond detection tasks, [50] extends HQC modeling to COVID-19 severity classification using CT images. The proposed framework integrates a quantum convolutional layer with classical fully connected layers and evaluates three severity levels—moderate, severe, and critical. Experiments conducted on CT datasets comprising several thousand slices at multiple resolutions (28×28, 32×32, and 64×64) report the best performance at 32×32 resolution, achieving accuracies of 96–97%. This highlights a trade-off between spatial resolution and constraints on quantum resources. Finally, [51] presents a CT-based hybrid framework that combines a strong pretrained classical encoder, such as a VGG-type architecture, with a quantum neural network classifier head implemented using modern quantum software toolchains. Evaluated on a CT dataset of approximately 2,500 images using an 80/20 train–test split, the model reports accuracy near 96–97%, precision above 98%, recall around 95%, and high specificity. These findings indicate that quantum classifier heads can be effectively integrated into established deep learning pipelines, although the resulting gains remain incremental rather than transformative.

2.5. Gap Analysis and Motivation of This Survey

Although hybrid quantum–classical (HQC) models have attracted increasing attention for COVID-19 medical imaging, the existing literature remains fragmented and difficult to compare. Most studies propose isolated hybrid designs and report promising results—predominantly for binary COVID-19 screening—using heterogeneous datasets, split strategies, and evaluation metrics, which limit fair attribution of performance gains to quantum components. A central gap is the lack of a unified architectural perspective that clarifies the functional role of quantum modules within classical pipelines. The interaction among architectural placement, encoding strategy, and circuit scale under NISQ constraints is rarely analyzed systematically and quantitatively. In addition, multi-class diagnostic robustness and class-wise sensitivity—critical for clinical deployment—are often underexplored compared to binary accuracy. Generalization and bias also remain insufficiently addressed. Heavy reliance on small or curated datasets, limited external validation, and instability in optimization raise concerns about reproducibility and clinical credibility. Existing surveys typically focus on general quantum machine learning or classical COVID-19 imaging, without offering a targeted, number-driven synthesis of HQC frameworks tailored to medical imaging workflows.

Motivated by these gaps, this survey makes the following contributions:

(i) it introduces a clear taxonomy of HQC architectures based on the functional placement of quantum modules;

(ii) it provides a quantitative comparative analysis of eight representative HQC models, covering datasets, encoding strategies, quantum resource usage, and performance in both binary and multi-class settings; and

(iii) it identifies key challenges and open research directions under NISQ constraints, linking algorithmic design to validation rigor and deployment feasibility.

By addressing these issues, this survey establishes a structured baseline for evaluating hybrid quantum–classical approaches and offers guidance for advancing HQC models from simulation-based studies toward clinically credible diagnostic systems.

3. Background and Theoretical Foundations

Before analyzing specific hybrid quantum–classical models for COVID-19 diagnosis, it is essential to review the theoretical underpinnings that motivate their design. This section first revisits

 doi:10.20944/preprints202605.0229.v1

the role and limitations of classical deep learning pipelines in medical imaging. It then introduces principles of quantum computation relevant to machine learning, including encoding strategies, variational quantum circuits, and kernel methods. Finally, we discuss quanvolutional approaches and NISQ-era constraints, providing the rationale for why hybrid designs are currently the most viable path for quantum-enhanced COVID-19 imaging.

3.1. Medical Imaging and Classical AI Pipelines

Chest computed tomography (CT) and chest X-ray (CXR) remain the primary imaging modalities for COVID-19 diagnosis. CT offers high sensitivity to radiological hallmarks such as ground-glass opacities, crazy-paving, and multifocal consolidations, while CXR provides broader accessibility for fast triage [52,53]. Automated AI pipelines generally include: (i) preprocessing, such as lung segmentation and intensity normalization; (ii) feature extraction via deep CNN backbones (e.g., ResNet, DenseNet); and (iii) classification through task-specific fully connected layers [54]. Despite their success, several limitations persist: (a) heavy computational burden for volumetric 3D CT requiring GPU clusters [55]; (b) dataset bias and distributional shifts across institutions leading to poor generalization [56]; and (c) limited interpretability, which hampers clinical trust and regulatory approval [57]. These shortcomings motivate exploration of alternative paradigms, including quantum-enhanced pipelines.

3.2. Quantum Machine Learning (QML)

Quantum machine learning leverages the exponentially large Hilbert space spanned by n qubits to encode richer feature representations than classical systems. An n-qubit register encodes vectors in a 2^n-dimensional space, with data transformed via parameterized unitary operations known as variational approaches [58,59].

- **Encoding schemes**

  Amplitude encoding. A normalized vector $x \in \mathbb{R}^{2^n}$ with $\sum_i |xi|^2 = 1$ is embedded as:

$$|\psi(x)\rangle = \sum_{i=0}^{2^n-1} x_i |i\rangle \quad (1)$$

maximizing compression but requiring complex state preparation.

- *Angle encoding*. Each feature is mapped to single-qubit rotations, e.g.,

$$|\psi(x)\rangle = \otimes_{j=1}^{d} R_y(xj)\ |0\rangle,\ R_y(\theta) = exp(-i\theta Y/2). \quad (2)$$

which scales linearly with the number of features and is practical for NISQ devices [59].

- *Image-specific encodings.* FRQI and NEQR embed pixel intensities and positions directly into quantum states [60,61]. For example,

$$|I\rangle_{FRQI} = \frac{1}{2^n} \sum_{i=0}^{2^{2n}-1} (cos\theta i\ |0\rangle + sin\theta i\ |1\rangle) \otimes |i\rangle, \quad (3)$$

$$|I\rangle_{NEQR} = \frac{1}{2^n} \sum_{i=0}^{2^{2n}-1} \otimes_{k=0}^{q-1} |Cki\rangle \otimes |i\rangle, \quad (4)$$

where $C_k^i \in \{0,1\}$ encodes pixel intensities in binary. These remain largely conceptual for high-resolution CT/CXR due to qubit overhead.

3.3. Variational Quantum Circuits (VQCs)

A variational quantum circuit (VQC) prepares a data-dependent quantum state.

 doi:10.20944/preprints202605.0229.v1

$$\rho(x;\theta) = U_\theta\, \rho_{\text{enc}}(x)\, U_\theta^\dagger, \tag{5}$$

where U_θ denotes a parameterized unitary circuit. In hybrid classification settings, class-relevant information is extracted by measuring an observable O, yielding an expectation value

$$\mu(x;\theta) = \text{Tr}[O\, \rho(x;\theta)], \tag{6}$$

which is subsequently mapped to a class probability, e.g., for binary classification.

In practice, VQCs are trained end-to-end within a classical optimization loop by minimizing a task-dependent loss function using gradient-based optimizers. Their appeal in hybrid quantum–classical models lies in their ability to provide trainable nonlinear transformations while maintaining shallow circuit depth, making them compatible with NISQ-era hardware constraints.

However, variational training is known to suffer from fundamental optimization challenges, most notably barren plateaus, in which gradient variance decreases rapidly with increasing circuit depth or qubit count [31]. This phenomenon significantly limits scalability and motivates the use of compact circuits, restricted qubit counts, and careful architectural design in practical hybrid models. In the context of medical imaging applications such as COVID-19 diagnosis, these limitations strongly influence where and how VQCs are integrated within hybrid pipelines, as discussed in later sections.

3.4. Quantum Kernels

Instead of training variational parameters, one may use fixed feature maps $U_\phi(x)$. The quantum kernel is defined as:

$$k(x,x') = |\langle 0 \mid U_\phi^\dagger(x) U_\phi(x') \mid 0 \rangle|^2, \tag{7}$$

Such kernels can be combined with classical kernel machines (e.g., SVM) and, with suitable U_ϕ, achieve separations that may be intractable for classical kernels[62,63].

3.5. Quanvolution and Hybrid Quantum–Classical CNNs

Quanvolution integrates quantum circuits directly into CNN pipelines. Images are divided into patches $\{x^{(p)}\}$. Each patch is encoded into qubits, processed by a quantum circuit U, and measured to produce nonlinear features:

$$f_k^{(p)} = \langle Z_k \rangle = \langle 0 \mid E^\dagger U^\dagger Z_k U E \mid 0 \rangle, \quad k = 1, \ldots, n, \tag{8}$$

which serve as quantum analogs of convolutional filters [64]. Two regimes exist:.

- **Random quanvolution**: fixed, randomly initialized circuits (HQ-CNN [65]).
- **Parametric quanvolution**: trainable circuits optimized end-to-end (HQCNN [66]).

3.6. NISQ Constraints

Qubit counts, gate fidelity, and decoherence restrict present-day hardware. Noise can be modeled as a depolarizing channel:

$$\Lambda_p(\rho) = (1-p)^\rho + p\frac{I}{2}, \tag{9}$$

which attenuates expectation values as:

$$Var[\hat{\mu}] \approx \frac{1-\mu^2}{S}, \qquad S \approx \frac{1-\mu^2}{\varepsilon^2}. \tag{10}$$

$$\langle O \rangle_{noisy} \approx (1-p)^d \langle O \rangle_{ideal} \tag{11}$$

where d is circuit depth.

 doi:10.20944/preprints202605.0229.v1

Consequently, practical HQC designs typically emphasize shallow, hardware-efficient approaches, classical dimensionality reduction prior to quantum processing, and shot-efficient measurement strategies [67].

4. Hybrid Quantum–Classical Models in COVID-19 Diagnosis

Hybrid quantum–classical learning has been actively investigated for COVID-19 imaging as a pragmatic strategy to exploit quantum representations while preserving the robustness of classical deep learning. In all reported studies, quantum components are not used as standalone learners. However, they are integrated into classical pipelines, typically either as quantum convolutional operators or as variational quantum classifiers operating on low-dimensional features extracted from chest X-ray (CXR) or computed tomography (CT) images. This section critically reviews the existing hybrid quantum–classical COVID-19 imaging studies by analyzing their methodological design, datasets, experimental outcomes, and reported limitations. A comprehensive hybrid quantum convolutional neural network (QCNN) for COVID-19 chest X-ray classification was presented in [44]. In this work, a classical CNN is augmented with a quanvolutional layer constructed from parameterized quantum circuits, which processes local image patches before classical feature aggregation. The model was evaluated on publicly available CXR datasets from the COVID-19 Radiography Database, in both binary (COVID-19 vs. normal) and multi-class (COVID-19, viral pneumonia, normal) scenarios. The reported results indicate high accuracy in binary classification and competitive performance relative to deeper classical CNNs, while using fewer trainable parameters. However, a noticeable decline in performance was observed for multi-class classification, highlighting the difficulty of distinguishing COVID-19 from other pneumonia types using current hybrid architectures. For CT-based diagnosis, a quantum neural network framework for clinical prognostic analysis of COVID-19 patients was proposed in [48]. Their approach used variational quantum circuits as classifiers and was trained on a large chest CT dataset comprising several thousand slices. Compared with a classical CNN baseline, the quantum model demonstrated improved diagnostic accuracy and reduced training time, suggesting potential efficiency benefits of quantum-assisted optimization. Despite these promising results, the study relied primarily on slice-level evaluation and lacked extensive external validation, limiting its direct clinical applicability.

To mitigate the limited availability of labeled CT data, a hybrid quantum machine learning architecture combining conditional generative adversarial networks (cGANs) with quantum neural networks was introduced in [49]. In this pipeline, synthetic CT images were generated to augment both a public COVID-CT dataset and a private hospital dataset prior to classification. The augmented datasets enabled the quantum classifier to achieve competitive performance compared with classical counterparts, particularly in low-data regimes. While this work highlights the potential of data augmentation in hybrid quantum learning, it also raises concerns regarding the reliability of synthetic samples and the risk of inflated performance if patient-level independence is not strictly maintained. Hybrid quantum transfer learning for multi-class CT diagnosis was investigated in [47]. The proposed framework combined pretrained classical feature extractors with variational quantum classifiers to distinguish COVID-19, community-acquired pneumonia, and normal CT scans. Experiments on public CT datasets demonstrated that hybrid quantum classifiers can effectively operate on compressed feature representations. However, the results also indicated that classification performance is highly sensitive to feature embedding strategies and circuit depth, and that multi-class CT discrimination remains challenging under current qubit limitations. A fully integrated hybrid quantum–classical convolutional architecture for CT imaging was proposed in [68]. The model embeds quantum processing blocks within a classical CNN and was evaluated on publicly available chest CT datasets for COVID-19 diagnosis. The reported results show strong diagnostic performance and parameter efficiency, supporting the feasibility of hybrid quantum convolution in CT analysis. Nevertheless, the evaluation was limited to a small number of datasets, and the absence of cross-institutional testing constrains conclusions regarding generalization.

 doi:10.20944/preprints202605.0229.v1

Transfer learning–based hybrid quantum models were further explored in [46], where a pretrained quantum convolutional neural network was proposed for COVID-19 CT classification. In this approach, classical deep networks were employed to extract transferable features, which were subsequently processed by a quantum convolutional classifier. The model was evaluated across multiple CT datasets with varying sizes and class distributions, demonstrating stable performance across heterogeneous data sources. A key strength of this study is its multi-dataset evaluation strategy, while a notable limitation is that quantum execution was conducted primarily in simulation rather than on real quantum hardware.

Beyond binary diagnosis, COVID-19 severity classification from chest CT images was investigated within a hybrid quantum–classical framework in [69]. The proposed model formulated severity assessment as a multi-level classification problem, aiming to support clinical stratification rather than simple disease detection. Experimental results suggest that hybrid quantum models can capture severity-related patterns in CT images; however, the study is constrained by the limited availability of large, standardized datasets with reliable severity annotations.

On the chest X-ray side, hybrid quantum convolutional neural network models have been proposed for COVID-19 detection, notably by Houssein et al. [44,45], who integrated shallow quantum circuits as quanvolutional layers within classical CNN pipelines. In these frameworks, classical preprocessing is first applied to reduce input dimensionality and normalize image intensities, after which local image patches are encoded into variational quantum circuits. When evaluated on public COVID-19 CXR datasets, these models achieved high performance in binary detection tasks, with reported accuracies exceeding 98% in some experimental settings. However, performance in multi-class classification scenarios exhibited noticeable degradation, indicating limited robustness when distinguishing COVID-19 from other pneumonia types. The studies further highlighted that classification accuracy is sensitive to circuit depth, qubit count, and encoding strategies, underscoring the importance of careful quantum architecture design.

Similarly, hybrid quantum transfer learning approaches have been explored, combining pretrained classical encoders with variational quantum classifier heads operating on compressed feature embeddings [46]. These frameworks demonstrated the feasibility of integrating quantum decision modules into established deep learning pipelines and reported competitive results on public CXR and CT datasets. Nevertheless, the observed performance gains were typically obtained under controlled experimental settings with limited dataset diversity and without external validation.

More recently, quantum-assisted deep learning architectures have been proposed in preprint studies, integrating strong classical backbones such as EfficientNet-B4 with deeper variational quantum circuits for COVID-19 imaging classification [70]. These works report strong classification performance on the COVIDx chest X-ray dataset and auxiliary benchmarks, including PneumoniaMNIST and OrganAMNIST, alongside systematic ablation analyses examining circuit depth and entanglement patterns. However, as these findings are currently available only as preprints and primarily validated in simulation environments, independent peer-reviewed confirmation and hardware-aware benchmarking remain necessary.

Overall, these studies demonstrate that hybrid quantum–classical models can achieve competitive performance for COVID-19 imaging tasks across both CXR and CT modalities. Nevertheless, the literature consistently reports limitations, including dataset bias, lack of external validation, sensitivity to preprocessing choices, and reliance on simulator-based quantum evaluations. These challenges underscore the need for more rigorous validation protocols, multi-institutional datasets, and hardware-aware benchmarking to establish the true clinical value of hybrid quantum–classical approaches.

4.1. Archetype A: Quanvolution-Based Hybrid Architectures ***

- **Architectural Rationale**

Quanvolution-based hybrid architectures integrate quantum computation into the early stages of the visual processing pipeline, where the quantum module functions as a local nonlinear feature

 doi:10.20944/preprints202605.0229.v1

transformation operator. Instead of learning convolutional filters in the classical sense, shallow quantum circuits process small image patches individually, and the resulting measurement statistics are interpreted as feature channels. The downstream classical neural network is then responsible for hierarchical feature aggregation and final decision making. This architectural paradigm is particularly compatible with near-term quantum hardware constraints, as it limits quantum computation to small-scale, patch-level operations, thereby controlling qubit count, circuit depth, and execution overhead. At the same time, it preserves the expressive power and scalability of classical convolutional neural networks (CNNs).

- **Functional Decomposition of Classical and Quantum Components**

Given an input medical image $I \in \mathbb{R}^{H\times W}$(e.g., chest X-ray or CT slice), Classical preprocessing is applied first, including resizing, intensity normalization, and optional denoising. The processed image $\tilde{I}$ is then partitioned into a set of non-overlapping (or weakly overlapping) patches:

$$\tilde{I} \rightarrow \{P_k\}_{k=1}^{K}, P_k \in \mathbb{R}^{p\times p}. \tag{12}$$

Each patch is vectorized into a low-dimensional real-valued feature vector.

$$x_k = \mathrm{vec}(P_k) \in \mathbb{R}^d, d = p^2, \tag{13}$$

which serves as the classical–quantum interface.

Quantum Encoding and Circuit Execution: Each patch vector x_k is encoded into a quantum state using angle encoding, the most widely adopted scheme in hybrid medical imaging due to its simplicity and numerical stability. Specifically, each component of $x_k$is mapped to the rotation angle of a single-qubit gate:

$$|\psi_k\rangle = E(x_k)\ |0\rangle^{\otimes n}, E(x_k) = \prod_{i=1}^{n} R_y\left(\alpha x_k(i) + \beta\right), \tag{14}$$

where $n = d$ qubits are used, and α, β are scaling parameters ensuring valid angular ranges.

Quantum circuit transformation: After encoding, a shallow quantum circuit $U(\theta)$ is applied:

$$|\phi_k\rangle = U(\theta)\ |\psi_k\rangle. \tag{15}$$

Two principal circuit configurations are commonly employed:

Random Quantum Circuits (RQC): circuit parameters are fixed and not optimized, acting as stochastic nonlinear feature maps.

Parameterized Quantum Circuits (PQC): circuit parameters θ are trainable and optimized jointly with classical network parameters.

Entangling gates (e.g., CNOT) may be introduced to capture intra-patch correlations.

- **Measurement and Feature Map Construction**

The quantum state $|\phi_k\rangle$ is measured using expectation values of Pauli observables:

$$z_k(j) = \langle\phi_k|\ O_j\ |\phi_k\rangle, O_j \in \{Z_i, Z_iZ_{i+1}, X_i, \dots\}, \tag{16}$$

producing a feature vector $\mathbf{z}_k \in \mathbb{R}^c$, where $c$denotes the number of measured observables.

The vectors $\{\mathbf{z}_k\}$are spatially reassembled to form a quantum-generated feature tensor:

$$F_q \in \mathbb{R}^{H'\times W'\times c}, \tag{17}$$

 doi:10.20944/preprints202605.0229.v1

with $H' = H/p$and $W' = W/p$in the non-overlapping case. This tensor plays a role analogous to the output of classical convolutional filters.

- **Classical Deep Learning Stage**

The quantum feature tensor $F_q$is processed by a conventional CNN consisting of convolutional, pooling, and fully connected layers. All global context modeling, high-level abstraction, and decision boundary learning are performed in this classical stage. Consequently, the quantum module serves as a feature-transformation front-end, while the classical network retains full responsibility for classification performance.

- **Training Strategy and Optimization Coupling**

Two training regimes are typically considered:

Fixed quanvolution (RQC): Only classical parameters φ are optim

$$\min_{\varphi} \ \mathcal{L}(f_{\varphi}(F_q), y). \tag{18}$$

Trainable quanvolution (PQC): Both quantum parameters θ and classical parameters φ are optimized:

$$\min_{\theta,\varphi} \ \mathcal{L}(f_{\varphi}(F_q(\theta)), y), \tag{19}$$

Trainable quanvolution (PQC): Both quantum parameters θ and classical parameters φ are optimized:

where gradients with respect to θ are computed using parameter-shift rules.

The RQC configuration offers superior stability and computational efficiency, whereas PQC-based quanvolution provides greater task adaptivity at the cost of increased optimization complexity.

- **Applicability and Design Implications**

Quanvolution-based hybrid architectures are particularly effective in scenarios where local texture patterns dominate, such as binary COVID-19 screening from chest X-ray images. Their modular design allows straightforward integration with classical deep learning pipelines and ensures scalability under quantum resource constraints. However, because quantum processing is limited to local patches, the burden of global reasoning and class discrimination remains with the classical network, especially in multi-class and fine-grained diagnostic tasks.

4.2. Archetype B: Classical Feature Extractor with Quantum Classifier Head

- **Architectural Rationale**

In this archetype, the hybrid coupling is explicitly decision-centric: a strong classical backbone performs high-capacity representation learning from medical images, while the quantum module serves as a final classifier head, replacing or augmenting conventional dense layers. This placement reflects a practical constraint of near-term quantum computing: quantum circuits are currently better suited for operating on compact, low-dimensional embeddings rather than high-resolution image tensors. Consequently, the classical network is responsible for transforming raw pixels into a semantically meaningful latent space, after which the quantum head performs a nonlinear decision mapping under a small qubit budget.

- **Functional Decomposition and Coupling Interface**

Let the input image be $I \in \mathbb{R}^{H\times W\times C}$. After standard classical preprocessing, a pretrained convolutional backbone $g(\cdot)$produces a high-level embedding:

$$h = g(I), h \in \mathbb{R}^{D}. \tag{21}$$

 doi:10.20944/preprints202605.0229.v1

Because Dtypically exceeds feasible quantum input sizes, an intermediate classical adapter is used to compress the embedding into a small vector compatible with a quantum circuit:

$$x = a(h), x \in \mathbb{R}^d, d \ll D. \tag{22}$$

The vector x forms the classical-to-quantum interface. It is encoded into a quantum state through an encoding map $E(\cdot)$:

$$|\psi\rangle = E(x)\ |0\rangle^{\otimes n}, \tag{23}$$

where n is the number of qubits, and n is typically chosen such that $n \approx d$ under angle encoding, or n is set by available hardware limits.

A parameterized quantum circuit $U(\theta)$then acts as the classifier head:

$$|\phi\rangle = U(\theta)\ |\psi\rangle. \tag{24}$$

Finally, measurement operators $\{O_j\}_{j=1}^m$produce a classical measurement vector:

$$z_j = \langle\phi|\ O_j\ |\phi\rangle, j = 1, \dots, m, \tag{20}$$
$$z = [z_1, \dots, z_m]^\top \in \mathbb{R}^m.$$

The vector z is mapped to class logits using either a direct linear layer or a lightweight classical post-processing head $r(\cdot)$:

$$\hat{y} = r(z), \hat{y} \in \mathbb{R}^K, \tag{25}$$

where K denotes the number of classes.

- **Quantum Head Design Variants**

Two implementation styles are commonly used within this archetype:

Variational quantum classifier head: The quantum head is directly optimized, and the measurement vector z is interpreted as class evidence or as features for a final linear classifier. This configuration emphasizes compactness and direct decision mapping.

Dressed quantum circuit head: A small classical layer is placed before and after the quantum circuit. The pre-quantum adapter improves trainability by shaping the embedding distribution, and the post-quantum layer stabilizes the formation and calibration of logits. Although this increases classical involvement, it typically improves the stability of optimization under limited circuit depth.

- **Training Objective and Gradient Coupling**

Let y be the ground-truth label and $\mathcal{L}(\hat{y}, y)$ the supervised loss. In this archetype, the classical backbone parameters φ, adapter parameters ω, and quantum parameters θ are trained either jointly or partially:

compactness and direct decision mapping.

$$\min_{\varphi,\omega,\theta}\ \mathbb{E}_{(I,y)}[\mathcal{L}(\hat{y}(I;\varphi,\omega,\theta), y)]. \tag{26}$$

Common strategies include freezing the backbone and training only the adapter and quantum head when data is limited or fine-tuning the backbone when sufficient labeled data is available. Gradients with respect to θ are computed using standard differentiable quantum programming methods such as parameter-shift rules, enabling end-to-end optimization.

- **Applicability and Design Implications**

This archetype is most appropriate when:

 doi:10.20944/preprints202605.0229.v1

- The imaging modality requires strong global context modeling, such as CT slices.
- A reliable pretrained backbone is available and can produce robust embeddings.
- Quantum resources are limited, and the quantum module must be compact and shallow.

A key implication is that the quantum head's contribution is tightly coupled to the quality of the classical embedding. Therefore, fair evaluation requires carefully matched baselines and ablations that isolate the effect of the quantum classifier head from that of classical feature extraction.

4.3. Archetype C: Quantum Feature Extractor with Classical Classifier

- **Architectural Rationale**

Archetype C places the quantum module in the role of a feature extractor operating on compact, classically prepared representations, while a lightweight classical classifier implements the final decision function. This design is motivated by two practical considerations. First, near-term quantum devices cannot process high-dimensional medical images directly; therefore, the input to a quantum circuit must be a low-dimensional vector. Second, using the quantum circuit as a feature generator rather than as a classifier head can provide a structured nonlinear transformation that produces quantum-enhanced features, which may be more discriminative for downstream classical classification, particularly when the classical classifier is intentionally kept simple to reduce overfitting.

Unlike quanvolution-based hybrids, the quantum processing here is not patch-wise. Unlike transfer-learning hybrids with quantum decision heads, the quantum circuit is used primarily to **construct a feature representation**, and the classification burden is shifted to a classical model operating on the quantum-derived feature space.

- **Functional Roles of Classical and Quantum Components**

Given an input image $I \in \mathbb{R}^{H\times W}$, classical preprocessing is applied to obtain $\tilde{I}$. A compact representation is then constructed through a lightweight classical mapping $b(\cdot)$, such as dimensionality reduction, shallow feature extraction, or a small embedding network:

$$u = b(\tilde{I}), u \in \mathbb{R}^D. \tag{27}$$

Since the quantum circuit requires a small input dimension, a compression step is applied:

$$x = c(u), x \in \mathbb{R}^d, d \ll D. \tag{28}$$

The vector x defines the classical-to-quantum interface.

Quantum feature extraction module: The quantum module encodes x into a quantum state and applies a parameterized circuit to generate a measurement vector that is interpreted as a learned feature embedding:

$$\begin{aligned} |\psi\rangle &= E(x)\ |0\rangle^{\otimes n}, \\ |\phi\rangle &= U(\theta)\ |\psi\rangle, \\ z_j &= \langle\phi|\ O_j\ |\phi\rangle, j = 1, \dots, m, \\ z &= [z_1, \dots, z_m]^\top \in \mathbb{R}^m. \end{aligned} \tag{29}$$

Here, z serves as the quantum-derived feature vector, typically with $m$chosen to match the needs of the downstream classifier rather than the number of classes.

Classical classifier: A classical classifier $q(\cdot)$then maps z to class logits and predictions:

$$\hat{y} = q(z), \hat{y} \in \mathbb{R}^K. \tag{30}$$

The classifier may be a small MLP, logistic regression, SVM, or a custom lightweight decision module.

- **Information Flow and Hybrid Coupling Mechanism**

 doi:10.20944/preprints202605.0229.v1

The coupling mechanism is sequential and can be summarized as:

- Classical preprocessing and compact feature construction generate x.
- Quantum processing transforms $x$into a quantum-derived feature vector z.
- Classical classification maps $z$to the final prediction $\hat{y}$.

In this archetype, the quantum circuit is optimized to produce features useful for classification rather than directly producing final class logits. This separation makes the role of the quantum module more explicit and can facilitate interpretability through feature-space analyses.

- **Training Strategy and Optimization**

Given labels y, the training objective minimizes a supervised loss:

$$\min_{\theta,\eta} \ \mathbb{E}_{(I,y)}[\mathcal{L}(q_\eta(z(I;\theta)),y)], \tag{31}$$

where θ denotes quantum circuit parameters and η denotes classifier parameters. The classical front-end parameters can be fixed or lightly tuned depending on the design of $b(\cdot)$ and available data. Quantum gradients are typically computed via parameter-shift rules, enabling end-to-end optimization of the quantum feature extractor and the classical classifier.

- **Applicability and Design Implications**

Archetype C is most appropriate when the goal is to exploit quantum circuits as feature generators under strict resource constraints, while maintaining flexibility in the decision stage via classical classifiers. It is particularly attractive for moderately sized datasets, where full fine-tuning of deep backbones may increase the risk of overfitting. However, performance is sensitive to the quality of the compact representation x, since excessive compression can remove clinically relevant patterns before quantum processing. Therefore, careful design of the classical front-end and rigorous reporting of feature dimensionality and circuit configuration are essential for reproducible comparisons.

4.4. Instantiation of the Three Hybrid Architectural Archetypes in Reviewed Studies

The taxonomy introduced in Sections 4.1–4.3 separates hybrid quantum–classical (HQC) imaging pipelines based on *where* quantum processing is injected and *what* function it serves (patch-level transformation, decision refinement, or latent feature generation). In survey studies, such a taxonomy is only valuable if it is operational—i.e., if it can consistently map real-world studies onto archetypes in a way that explains observed differences in performance, computational scaling, and validation behavior. Accordingly, this subsection anchors the taxonomy in the seven peer-reviewed studies retained for comparative analysis, explicitly identifying the classical component, the functional role of the quantum module, and the coupling point that determines system-level behavior.

Analysis of Table 2.

Table 2. Instantiation of the Three Hybrid Architectural Archetypes in the Reviewed HQC Studies.

Method (Author, Year)	Imaging Modality	Archetype	Classical Component	Quantum Component Role	Hybrid Coupling Location	Task Formulation
HQ-CNN [44]	CXR	A	CNN backbone	Random quanvolution (patch-level)	Early feature extraction	Binary / multi-class
HQ-CNN [45]	CXR	A	CNN backbone	Random quanvolution (patch-level)	Early feature extraction	Binary / multi-class
QCNN-CT [68]	CT	A	CNN + aggregation	Parametric quanvolution (patch-level)	Early feature extraction	Binary

 doi:10.20944/preprints202605.0229.v1

QTL-CT [47]	CT	B	Pretrained CNN embedding	Variational quantum classifier	Decision head	Binary / multi-class
Hybrid CNN+VQC [48]	CT	B	CNN feature extractor	VQC classifier head	Final classification	Binary
Pretrained-QCNN [46]	CT	B	Pretrained CNN	Quantum classifier head	Decision stage	Binary
cGAN-QML [49]	CT	C	Classical extractor + cGAN	Quantum feature generator	Intermediate representation	Binary

Table 2 exposes a pronounced concentration in Archetypes A and B, indicating that current HQC research largely explores two extremes: (i) maximizing quantum exposure by applying circuits to image patches early in the pipeline (Archetype A), or (ii) minimizing quantum footprint by restricting the circuit to the final decision stage (Archetype B). This polarization is not just a stylistic choice—it has direct consequences for (a) scaling laws (patch-wise circuits scale with image resolution), (b) noise sensitivity (measurement noise accumulates over many circuit calls), and (c) attribution clarity (decision-stage circuits make it easier to isolate what the quantum block contributes).

In contrast, Archetype C is represented by only one study. This is notable because intermediate feature generation is arguably the most "balanced" strategy under NISQ constraints: the quantum module can act as a nonlinear representation shaper without incurring patch-wise multiplicative costs. The limited adoption of Archetype C suggests an open methodological gap: the community has not yet systematically evaluated whether intermediate quantum-feature generation can offer a better trade-off between expressivity, stability, and computational feasibility than the two dominant extremes.

4.5. Comparative Summary of Hybrid Quantum–Classical Models

After structurally mapping studies into archetypes (Table 2), the next step is to quantify how design decisions translate into measurable outcomes. In hybrid pipelines, reported accuracy cannot be interpreted in isolation, because it is confounded by task formulation (binary vs. multi-class), encoding stability, qubit budget, circuit depth, and the degree of classical dominance. Therefore, this subsection provides an integrated comparison that jointly reports diagnostic outcomes and quantum resource characteristics, enabling a survey-level assessment of feasibility under NISQ constraints and identifying which architectural patterns are consistently associated with strengths or failure modes.

Three survey-level insights are reinforced by Table 3:

Table 3. Integrated Diagnostic Performance and Quantum Resource Characteristics of Reviewed HQC Models.

Method	Binary Performance (reported)	Multi-class Performance (reported)	Encoding	Qubits (reported)	Circuit Depth	Notable Behavior / Caveat
HQ-CNN ([44]	High (≈98–99%)	Lower (≈82–90%)	Angle	4	Shallow	Multi-class degradation; class imbalance effects
HQ-CNN[45]	High (≈98%)	Moderate (≈90%)	Angle	4	Shallow	Performance varies by dataset/split protocol
QCNN-CT [68]	High (≈97.9–99.6%)	—	Angle	4 / 8	Shallow	Non-monotonic scaling (4q > 8q)
QTL-CT [47]	Moderate (≈70–75%)	Lower (≈55–60%)	Angle / Amplitude	4–9	≤4	Encoding sensitivity; amplitude instability
Hybrid CNN+VQC [48]	—	—	Angle	Small	Shallow	AUROC reported (hybrid > classical)

 doi:10.20944/preprints202605.0229.v1

Pretrained-QCNN [46]	High (>97%)	—	Angle	4–6	Shallow	Stable binary CT; quantum as decision refiner
cGAN-QML [49]	Competitive (low-data)	—	Angle	Small	Shallow	Synthetic-data dependence; domain-shift risk

Binary screening is structurally "easier" for current hybrids than multi-class diagnosis. Across studies that evaluate both, binary performance is consistently higher. This suggests that today's HQC design, especially patch-level feature transforms—primarily amplify coarse separability (COVID vs. non-COVID) rather than extracting representations that remain stable under fine-grained semantic overlap (e.g., COVID vs. pneumonia). In a review context, this is a crucial distinction: many high accuracies reported in HQC work are likely tied to binary formulations and do not necessarily generalize to clinically realistic differential diagnosis.

Quanvolution-based gains may be real, but they are coupled with a robustness penalty.

Archetype A models dominate headline binary performance, but their multi-class degradation suggests that patch-level quantum nonlinearities can enhance local texture cues without guaranteeing global discriminability. This is consistent with a system interpretation: local quantum feature maps may enrich early representations, but downstream classical aggregation is still required to learn global decision boundaries—and those boundaries become fragile under multi-class overlap.

Scaling qubits upward is not a free lunch under NISQ-era training dynamics.

QCNN-CT shows non-monotonic behavior (4 qubits outperform 8). From a survey perspective, this is exactly the type of pattern that should temper overly optimistic "more qubits → better performance" narratives. Larger circuits expand parameterization and can exacerbate optimization pathologies and measurement variance, offsetting any representational benefit.

Collectively, Table 3 supports a conservative yet technically grounded conclusion for this survey section: current HQC models are best supported for binary screening in shallow, angle-encoded, small-qubit regimes, while multi-class robustness and scaling behavior remain unsettled.

5. Comparative Analysis and Synthesis

Hybrid quantum–classical models for COVID-19 imaging vary not only in where the quantum circuit is inserted but also in how the hybrid interface is defined, how quantum resources are allocated, and how evaluation protocols support (or fail to support) claims of clinical relevance. A survey-level synthesis, therefore, must go beyond reporting accuracy and instead analyze cross-study behavior along multiple orthogonal axes. In this section, we compare the seven studies across six dimensions that collectively determine scientific credibility and translational feasibility: (i) architectural placement and functional responsibility, (ii) encoding and quantum resource scaling, (iii) task formulation effects (binary vs. multi-class), (iv) computational scaling and measurement costs, (v) validation rigor and generalization control, and (vi) deployment readiness.

5.1. Architectural Placement and Functional Responsibility of the Quantum Module

The location of the quantum module determines whether quantum computation acts as an early representational transformer (feature responsibility), a late-stage decision refiner (classification responsibility), or an intermediate feature generator (representation reshaping). In hybrid systems, this choice implicitly assigns "who does what" between quantum and classical components. Table 2 already maps each study into an archetype; here we interpret that mapping through the lens of functional responsibility and expected system behavior.

Archetype A assigns the quantum circuit a front-end representational role, meaning its outputs become the substrate for all downstream learning. This can amplify quantum influence but also makes overall performance highly sensitive to measuring noise and patch partition choices. Archetype B assigns quantum computation to a decision-stage role, with the classical backbone carrying the representational burden; quantum gains thus tend to be incremental, but the approach

 doi:10.20944/preprints202605.0229.v1

is computationally predictable. Archetype C aims to balance the two: quantum circuits reshape compact representations produced classically, potentially improving discriminability without patch-wise scaling. However, the scarcity of Archetype C in Table 2 indicates that this potentially promising region remains underexplored in current COVID imaging HQC literature.

5.2. Encoding Strategy, Qubit Budget, and Circuit Depth

Encoding, qubit budget, and circuit depth jointly define the feasible operating regime of HQC models. In a survey setting, the key question is not merely "which encoding is used," but *why certain encodings dominate*, *how they interact with stability*, and *what scaling behavior is evidenced in practice*. Table 3 provides the consolidated view needed to make such claims.

Angle encoding appears to be the de facto standard across the reviewed studies, consistent with its practical advantages: it avoids the brittleness and optimization instability often associated with amplitude encoding, and it enables shallow circuits compatible with NISQ constraints. The QTL-CT results in Table 3 underscore this: amplitude variants are reported as less stable than angle-encoding variants. Moreover, the non-monotonic scaling observed in QCNN-CT (4q > 8q) suggests that qubit increases can amplify optimization difficulty and measurement variance, potentially dominating any representational advantage. From a synthesis standpoint, the evidence supports a pragmatic guideline: small qubit counts, shallow depth, and angle encoding currently represent the most empirically stable design envelope for HQC COVID imaging.

5.3. Task Formulation Effects: Binary Screening vs. Multi-Class Diagnosis

In medical imaging, binary triage tasks are often used as a first evaluation step, but clinical deployment frequently requires multi-class discrimination or differential diagnosis. For HQC models, this distinction is especially important because patch-level quantum transformations may capture coarse changes but fail to distinguish visually similar pathologies. Table 3 is central again because it reports both binary and multi-class outcomes when available.

The consistent pattern is binary dominance: when multi-class results are reported, they underperform binary results, sometimes substantially (e.g., HQ-CNN 2022). From a survey perspective, this implies that current HQC gains may largely reflect improved separability between strongly distinct classes rather than a fundamentally stronger representation. This matters for interpretation: reporting only binary accuracy can overstate clinical readiness. Consequently, this survey treats multi-class robustness as a differentiating criterion rather than an optional add-on and highlights that the main empirical gap is not raw binary performance but stable generalization under multi-class overlap.

5.4. Computational Scaling, Measurement Costs, and Quantum Invocation Frequency

Unlike purely classical CNNs, HQC pipelines carry an additional computational layer: circuit execution and measurement. The practical cost depends primarily on how often the circuit is called per image and how measurement shots are configured. This dimension is essential for a survey because it determines whether a design could plausibly move from simulation to hardware and, eventually, to clinical inference constraints.

Table 4 formalizes a crucial systems trade-off: representational influence vs. invocation scalability. Archetype A models call quantum circuits repeatedly for each patch, so the cost grows with resolution and patch granularity. Even when circuits are shallow, the multiplicative effect of patch count and shots can dominate runtime, making edge deployment difficult without substantial acceleration. In contrast, Archetype B models have bounded, predictable quantum overhead because the circuit is executed once per image, making them the most plausible candidates for early hardware experimentation. However, because they do not expose the quantum module to raw spatial structure, their gains are likely incremental. Importantly, shot and latency reporting is inconsistent across

 doi:10.20944/preprints202605.0229.v1

studies, limiting reproducibility and preventing apples-to-apples runtime comparison—an explicit methodological gap highlighted by this survey.

Table 4. Computational Scaling and Quantum Invocation Characteristics.

Method	Quantum Calls per Image	Shot Reporting	Dominant Scaling Dependency	Expected Computational Load
HQ-CNN [44]	Patch-wise (many calls)	Reported (≈500–1000)	Patch count × shots	High
HQ-CNN [45]	Patch-wise (many calls)	Not reported	Patch count	High
QCNN-CT [68]	Patch-wise	Not reported	Patch count	High
QTL-CT [47]	Single (per image)	Limited detail	Image-independent	Low
Hybrid CNN+VQC [48]	Single (per image)	Limited detail	Image-independent	Low
Pretrained-QCNN [46]	Single (per image)	Not reported	Image-independent	Low
cGAN-QML [49]	Single (per image)	Not reported	Image-independent	Moderate

5.5. Validation Rigor, Generalization, and Methodological Robustness

In medical imaging, strong reported metrics can arise from dataset artifacts, leakage, or overly permissive split strategies. For HQC models, these concerns are compounded by hybrid systems, which introduce additional parameters and potential avenues for overfitting, while most reported results are based on simulators. Therefore, evaluation rigor is a first-class dimension in the synthesis.

Two conclusions follow directly from Table 5. First, patient-wise splitting is rare (explicitly used in QTL-CT), yet it is essential when multiple images (or slices) per patient are available; otherwise, leakage can inflate performance due to intra-patient correlation. Second, external validation is absent across all studies, which prevents reliable claims about cross-site robustness—arguably the most critical requirement for clinical relevance. Third, the lack of hardware execution across the board means current results cannot capture the impact of real device noise, transpilation constraints, or readout bias. From a survey standpoint, this indicates that evaluation methodology, not architectural novelty, is currently the limiting factor in achieving credible translation.

Table 5. Validation Strategy and Generalization Controls.

Method	Split Strategy	External Validation	Hardware Execution	Generalization Risk (qualitative)
HQ-CNN [44]	Random	No	No	Moderate
HQ-CNN [45]	Random	No	No	Moderate
QCNN-CT [68]	Cross-validation	No	No	Moderate
QTL-CT [47]	Patient-wise	No	No	Lower (relative)
Hybrid CNN+VQC [48]	Random	No	No	Moderate
Pretrained-QCNN [46]	Random	No	No	Moderate
cGAN-QML [49]	Random	No	No	Moderate (plus synthetic bias)

 doi:10.20944/preprints202605.0229.v1

5.6. Deployment Readiness and Translational Feasibility

Even if a hybrid model achieves high accuracy in simulation, deployment requires a predictable inference time, a manageable computational footprint, and the feasibility of integration. For HQC models, deployment readiness depends heavily on circuit invocation frequency (Table 4), as well as on whether the quantum block's role justifies the integration complexity and whether the approach can be migrated to hardware without prohibitive redesign.

Table 6 makes clear that "deployment readiness" in HQC imaging is currently constrained less by headline accuracy and more by *system integration realities*. Patch-wise approaches are difficult to justify for latency-sensitive edge triage unless they deliver robust, reproducible improvements under real hardware constraints. Decision-stage hybrids are more feasible operationally but risk being perceived as "quantum add-ons" unless they demonstrate gains under strict validation protocols. The cGAN-QML line highlights a separate translational hazard: synthetic augmentation may improve performance in low-data regimes but can introduce hidden biases that collapse under domain shift. In synthesis, the reviewed evidence suggests that near-term practical deployment is most plausible for bounded-invocation (single-call) hybrids paired with rigorous validation, while patch-wise designs remain primarily research-grade unless hardware-aware efficiency is explicitly demonstrated.

Table 6. Deployment Readiness Assessment and Translational Risk Profile.

Method	Edge Suitability	Server Suitability	Principal Translational Risk
HQ-CNN [44]	Moderate	High	Patch-wise overhead; shot-latency accumulation
HQ-CNN [45]	Moderate	High	Patch-wise overhead; incomplete runtime reporting
QCNN-CT [68]	Low	High	Scaling cost; non-monotonic qubit behavior
QTL-CT [47]	Moderate	High	Limited multi-class robustness; embedding dependence
Hybrid CNN+VQC [48]	Moderate	High	Dataset dependence; limited transparency on runtime
Pretrained-QCNN [46]	Moderate	High	Classical dominance; incremental quantum contribution
cGAN-QML [49]	Moderate	High	Synthetic bias; domain shift risk

5.7. Integrated Synthesis

Across Tables 2–6, a coherent picture emerges: present HQC models for COVID-19 imaging should be interpreted as quantum-augmented classical pipelines rather than replacements for deep CNNs. Architectural placement determines whether the quantum block meaningfully participates in representation learning (Archetype A), modestly refines decision boundaries (Archetype B), or reshapes compact embeddings (Archetype C). Empirically, gains are most consistent in binary screening, while multi-class robustness, validation rigor, and hardware-aware runtime reporting remain the major unresolved gaps. Consequently, the central recommendation of this survey is methodological: progress toward clinical credibility requires standardized split protocols (preferably patient-wise where applicable), external validation, and explicit reporting of quantum execution costs (shots, circuit depth after transpilation, and latency per image). Only under such conditions can the field meaningfully assess whether observed improvements reflect genuine hybrid utility or dataset- and protocol-dependent effects.

 doi:10.20944/preprints202605.0229.v1

6. Challenges and Open Issues in Hybrid Quantum–Classical COVID-19 Imaging

Despite encouraging experimental results, hybrid quantum–classical (HQC) models for COVID-19 medical imaging remain at an exploratory stage. The comparative analysis in Sections 4 and 5 reveals that performance gains are task-dependent, resource-constrained, and often evaluated under limited validation protocols. More importantly, several systemic challenges—spanning data availability, NISQ hardware limitations, optimization dynamics, validation rigor, interpretability, and deployment feasibility—constrain the transition from simulation prototypes to clinically credible systems. This section synthesizes these challenges into a structured framework, highlighting both immediate bottlenecks and long-term research directions.

Table 7 indicates that the principal bottlenecks of hybrid quantum–classical COVID-19 imaging systems are systemic rather than purely algorithmic. Methodologically, limited dataset diversity, frequent reliance on random splits, and the absence of external validation introduce significant uncertainty in generalization, potentially leading to reported gains that partially reflect dataset-specific correlations rather than robust hybrid advantages. From a hardware perspective, all reviewed models operate in shallow-depth, low-qubit regimes imposed by NISQ constraints, which restrict expressive scaling and make performance sensitive to measurement variance and noise. Optimization stability further constrains design choices, as evidenced by encoding-dependent behavior and the practical dominance of angle encoding over more information-dense alternatives. Additionally, patch-wise architectures incur multiplicative quantum invocation costs, leading to resolution-dependent latency that limits deployment feasibility. Finally, the absence of hardware-in-the-loop experiments and explainability mechanisms limits translational credibility. Collectively, these findings suggest that progress toward clinically viable HQC systems will depend less on architectural novelty and more on rigorous validation, hardware-aware benchmarking, and stability-oriented design under realistic noise models.

Table 7. Key Challenges and Open Issues in Hybrid Quantum–Classical Frameworks for COVID-19 Imaging.

Challenge Category	Description	Practical Implications	Evidence from Reviewed Studies
Data limitations	Small, imbalanced, single-institution datasets; slice-level splitting	Overfitting; inflated binary performance; poor generalization	Limited external validation (Table 5)
NISQ hardware constraints	Limited qubits; shallow depth; gate noise; measurement variance	Restricts scaling; unstable training; hardware gap	Shallow circuits across all models (Table 3)
Optimization pathologies	Barren plateaus; gradient vanishing; non-convex loss landscapes	Training instability; sensitivity to encoding	Amplitude encoding instability (QTL-CT)
Encoding bottlenecks	Trade-off between dimensionality reduction and information loss	Loss of subtle pathological features	Decision-head sensitivity (Table 3)
Computational scaling	Patch-wise invocation multiplies runtime	Latency accumulation; infeasible edge deployment	High quantum calls in Archetype A (Table 4)
Validation rigor gaps	Random splits; no external or hardware validation	Limited clinical credibility	Universal absence of external validation
Interpretability deficit	Lack of explainable mechanisms for quantum features	Regulatory and clinician trust barriers	No interpretability analysis reported

 doi:10.20944/preprints202605.0229.v1

Deployment uncertainty	No hardware-in-the-loop benchmarking	Unknown real-device latency and noise impact	All studies simulator-based

7. Future Research Directions

The progression of quantum-enhanced medical imaging from experimental prototypes to clinically meaningful systems requires addressing both algorithmic limitations and translational barriers. Based on the comparative analysis presented in Sections 4–6, several structured research priorities emerge. Table 8 synthesizes these priorities and links them to observed gaps in the current literature.

Table 8. Key Future Research Directions.

Research Direction	Observed Gap in Current Literature	Strategic Rationale	Anticipated Impact
Large-scale, multi-site datasets	Most studies rely on small or single-source datasets	Reduce dataset bias and improve generalizability	Clinically credible validation
Error mitigation and hardware-aware design	Limited real-hardware validation; strong simulator reliance	Improve robustness under NISQ noise and decoherence	Stable performance beyond simulation
Domain-specific encoding strategies	Generic amplitude/angle encoding dominates	Integrate anatomical priors (e.g., mask-guided patches)	Improved multiclass and structured generalization
Benchmarking against matched classical baselines	Inconsistent baseline selection across studies	Enable fair attribution of quantum contribution	Strengthened scientific validity
Hardware-in-the-loop evaluation	Few latency/throughput measurements reported	Assess inference cost and deployment feasibility	Realistic translational assessment
Explainable quantum AI frameworks	Limited interpretability analysis	Support clinician trust and regulatory pathways	Facilitates compliance with medical AI standards

7.1. Near-Term Priorities

In the short term, the field must prioritize rigorous benchmarking on larger, heterogeneous datasets combined with hardware-in-the-loop experimentation. Current evidence indicates heavy dependence on simulator environments, which obscures real-world performance constraints. Standardized reporting of qubit count, shot configuration, circuit depth, and transpilation overhead should accompany such evaluations.

7.2. Mid-Term Methodological Advancements

Mid-term progress will depend on innovations in error mitigation and domain-aware encoding schemes. Variational circuits must be optimized for hardware efficiency while maintaining expressive capacity. Incorporating anatomical priors—such as region-of-interest masking or patch-guided encoding—may reduce dimensional burden without excessive classical compression.

Furthermore, systematic comparison with carefully matched classical architectures is essential to prevent overstated claims of quantum contribution.

 doi:10.20944/preprints202605.0229.v1

7.3. Long-Term Translational Development

Over the longer horizon, explainable quantum AI frameworks will be central to clinical integration. Regulatory pathways for Software as a Medical Device require transparency, validation rigor, and failure-mode analysis. Developing interpretable quantum feature representations and traceable hybrid pipelines will therefore be critical for regulatory approval and clinician acceptance.

8. Conclusion

Hybrid quantum–classical (HQC) models represent an emerging research direction for COVID-19 diagnosis using medical images. By embedding parameterized quantum circuits into convolutional neural network pipelines, existing approaches aim to determine whether compact quantum modules can improve feature representation and classification efficiency. According to the surveyed studies, quantum-based front-end and feature-extractor configurations frequently report high binary screening accuracy and AUROC under simulated conditions, whereas classifier-head integrations tend to yield more modest or inconsistent improvements.

However, the current body of evidence remains limited in scope and methodological rigor. Most evaluations are conducted exclusively on simulator backends, rely on relatively small or single-source datasets, and lack external or multi-site validation. Multi-class diagnostic performance consistently lags binary screening outcomes and reported per-class sensitivities—particularly for COVID-19 categories—often fall below thresholds typically associated with clinical reliability. These observations indicate that reported performance gains should be interpreted cautiously, especially given dataset biases and heterogeneous evaluation protocols.

This survey makes two primary contributions. First, it provides a structured comparative synthesis of representative HQC architectures, systematically examining how the placement of quantum modules, encoding strategies, and circuit scale relate to reported empirical outcomes. Second, it consolidates key methodological and translational challenges—including NISQ hardware limitations, optimization instability, dataset bias, interpretability concerns, and deployment feasibility—and organizes them into a coherent research agenda.

At present, there is no conclusive empirical evidence demonstrating task-level quantum advantage for COVID-19 medical imaging. Literature should therefore be viewed as exploratory rather than definitive. Future progress will depend less on increasing circuit depth or qubit count and more on principled hybrid interface design, rigorous multi-center validation, standardized benchmarking, and hardware-aware evaluation. Whether quantum-enhanced pipelines can ultimately provide clinically meaningful benefits remains an open question that requires sustained interdisciplinary investigation.

References

1. W. H. Organization, "WHO Director-General's opening remarks at the media briefing on COVID-19," *January*, vol. 30, 2022.
2. F. M. Shah et al., "A comprehensive survey of covid-19 detection using medical images," *SN Computer Science*, vol. 2, no. 6, p. 434, 2021.
3. A. Bernheim et al., "Chest CT findings in coronavirus disease-19 (COVID-19): relationship to duration of infection," *Radiology*, vol. 295, no. 3, pp. 685-691, 2020.
4. K. R. Bhatele et al., "Covid-19 detection: A systematic review of machine and deep learning-based approaches utilizing chest x-rays and ct scans," *Cognitive Computation*, vol. 16, no. 4, pp. 1889-1926, 2024.
5. X. Han, Z. Hu, S. Wang, and Y. Zhang, "A survey on deep learning in COVID-19 diagnosis," *Journal of imaging*, vol. 9, no. 1, p. 1, 2022.
6. A. Khan, S. H. Khan, M. Saif, A. Batool, A. Sohail, and M. Waleed Khan, "A Survey of Deep Learning Techniques for the Analysis of COVID-19 and their usability for Detecting Omicron," *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 36, no. 8, pp. 1779-1821, 2024.

 doi:10.20944/preprints202605.0229.v1

7. I. D. Apostolopoulos and T. A. Mpesiana, "Covid-19: automatic detection from x-ray images utilizing transfer learning with convolutional neural networks," *Physical and engineering sciences in medicine,* vol. 43, pp. 635-640, 2020.
8. L. Wang, Z. Q. Lin, and A. Wong, "COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images," *Scientific reports,* vol. 10, no. 1, p. 19549, 2020.
9. J. Mozaffari, A. Amirkhani, and S. B. Shokouhi, "A survey on deep learning models for detection of COVID-19," *Neural Computing and Applications,* vol. 35, no. 23, pp. 16945-16973, 2023.
10. A. M. Ismael and A. Şengür, "Deep learning approaches for COVID-19 detection based on chest X-ray images," *Expert Systems with Applications,* vol. 164, p. 114054, 2021.
11. A. S. Lundervold and A. Lundervold, "An overview of deep learning in medical imaging focusing on MRI," *Zeitschrift fuer medizinische Physik,* vol. 29, no. 2, pp. 102-127, 2019.
12. M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Physical review letters,* vol. 122, no. 4, p. 040504, 2019.
13. C. Ciliberto et al., "Quantum machine learning: a classical perspective," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences,* vol. 474, no. 2209, p. 20170551, 2018.
14. R. S. Gupta, C. E. Wood, T. Engstrom, J. D. Pole, and S. Shrapnel, "A systematic review of quantum machine learning for digital health," *npj Digital Medicine,* vol. 8, no. 1, p. 237, 2025.
15. S. C. Fairburn, L. Jehi, B. T. Bicknell, B. G. Wilkes, and B. Panuganti, "Applications of quantum computing in clinical care," *Frontiers in medicine,* vol. 12, p. 1573016, 2025.
16. F. F. Flöther, "The state of quantum computing applications in health and medicine," *Research Directions: Quantum Technologies,* vol. 1, p. e10, 2023.
17. E. A. Radhi, M. Y. Kamil, and M. A. Mohammed, "Quantum Machine and Deep Learning for Medical Image Classification: A Systematic Review of Trends, Methodologies, and Future Directions," *Iraqi Journal for Computer Science and Mathematics,* vol. 6, no. 2, p. 9, 2025.
18. U. Ullah and B. Garcia-Zapirain, "Quantum machine learning revolution in healthcare: a systematic review of emerging perspectives and applications," *IEEE Access,* vol. 12, pp. 11423-11450, 2024.
19. F. Yan, H. Huang, W. Pedrycz, and K. Hirota, "Review of medical image processing using quantum-enabled algorithms," *Artificial Intelligence Review,* vol. 57, no. 11, p. 300, 2024.
20. E. Mohammadisavadkoohi, N. Shafiabady, and J. Vakilian, "A Systematic Review on Quantum Machine Learning Applications in Classification," *IEEE Transactions on Artificial Intelligence,* 2025.
21. H. Chetioui, "Leveraging Quantum Computing to Revolutionize Deep Learning: A Focus on Hybrid Algorithms for Medical Image Classification," 2025.
22. P. Khanal and A. Khanal, "Machine Learning for Advancing Quantum Field Theory-Enabled Sensor Networks in Precision Disease Detection: A Review," *Crimson Journal of Science and Technology,* vol. 1, no. 1, pp. 97-129, 2025.
23. P. Lamichhane and D. B. Rawat, "Quantum Machine Learning: Recent Advances, Challenges and Perspectives," *IEEE Access,* 2025.
24. A. Sonavane, S. Jaiswar, M. Mistry, A. Aylani, and D. Hajoary, "Quantum machine learning models in healthcare: Future trends and challenges in healthcare," *Quantum Computing for Healthcare Data,* pp. 167-187, 2025.
25. M. F. Shahriyar and G. Tanbhir, "Advancements and Challenges in Quantum Machine Learning for Medical Image Classification: A Comprehensive Review," in *2025 3rd International Conference on Intelligent Systems, Advanced Computing and Communication (ISACC)*, 2025: IEEE, pp. 1126-1133.
26. S. Asif, M. Zhao, F. Tang, and Y. Zhu, "A deep learning-based framework for detecting COVID-19 patients using chest X-rays," *Multimedia Systems,* vol. 28, no. 4, pp. 1495-1513, 2022.
27. H. Gunraj, L. Wang, and A. Wong, "Covidnet-ct: A tailored deep convolutional neural network design for detection of covid-19 cases from chest ct images," *Frontiers in medicine,* vol. 7, p. 608525, 2020.
28. X. Yang, X. He, J. Zhao, Y. Zhang, S. Zhang, and P. Xie, "Covid-ct-dataset: a ct scan dataset about covid-19," *arXiv preprint arXiv:2003.13865,* 2020.
29. N. Ghassemi et al., "Automatic diagnosis of COVID-19 from CT images using CycleGAN and transfer learning," *Applied Soft Computing,* vol. 144, p. 110511, 2023.

 doi:10.20944/preprints202605.0229.v1

30. V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209-212, 2019.
31. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nature communications*, vol. 9, no. 1, p. 4812, 2018.
32. M. Henderson, S. Shakya, S. Pradhan, and T. Cook, "Quanvolutional neural networks: powering image recognition with quantum circuits," *Quantum Machine Intelligence*, vol. 2, no. 1, p. 2, 2020.
33. G. Singh, H. Jin, and K. M. Merz Jr, "Benchmarking MedMNIST dataset on real quantum hardware," *arXiv preprint arXiv:2502.13056*, 2025.
34. J. K. Fahim, P. C. Paul, M. R. Hossain, M. T. Ahmed, and D. Chakraborty, "HQCNN: A Hybrid Quantum-Classical Neural Network for Medical Image Classification," *arXiv preprint arXiv:2509.14277*, 2025.
35. M. Nandy Pal and D. De, "FedQCNN: A Privacy-Preserving Federated Quantum Convolutional Neural Network for Retinal Image Classification," *IET Quantum Communication*, vol. 6, no. 1, p. e70010, 2025.
36. M. Rahman and J. Zhuang, "NQNN: Noise-Aware Quantum Neural Networks for Medical Image Classification," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2025: Springer, pp. 433-442.
37. A. Senokosov, A. Sedykh, A. Sagingalieva, B. Kyriacou, and A. Melnikov, "Quantum machine learning for image classification," *Machine Learning: Science and Technology*, vol. 5, no. 1, p. 015040, 2024.
38. Y. Sun, X. Deng, and H. Shao, "Quantum dual-branch neural networks with transfer learning for early detection of skin cancer," *American Journal of Translational Research*, vol. 17, no. 5, p. 3357, 2025.
39. R. Ur Rasool, H. F. Ahmad, W. Rafique, A. Qayyum, J. Qadir, and Z. Anwar, "Quantum computing for healthcare: A review," *Future Internet*, vol. 15, no. 3, p. 94, 2023.
40. S. V. Babu, P. Ramya, and J. Gracewell, "Revolutionizing heart disease prediction with quantum-enhanced machine learning," *Scientific Reports*, vol. 14, no. 1, p. 7453, 2024.
41. H. Gupta, H. Varshney, T. K. Sharma, N. Pachauri, and O. P. Verma, "Comparative performance analysis of quantum machine learning with deep learning for diabetes prediction," *Complex & Intelligent Systems*, vol. 8, no. 4, pp. 3073-3087, 2022.
42. B. Dhara, M. Agrawal, and S. D. Roy, "A Generalised Transform Methodology Using Quantum Computation and Its Application for Electrocardiogram (ECG) Classification," *IET Quantum Communication*, vol. 6, no. 1, p. e70012, 2025.
43. S. J. Lee et al., "Quantum Neural Network Tuning and Performance Evaluation for a Breast Cancer Dataset," *medRxiv*, p. 2025.10. 03.25336905, 2025.
44. E. H. Houssein, Z. Abohashima, M. Elhoseny, and W. M. Mohamed, "Hybrid quantum-classical convolutional neural network model for COVID-19 prediction using chest X-ray images," *Journal of Computational Design and Engineering*, vol. 9, no. 2, pp. 343-363, 2022.
45. E. H. Houssein, Z. Abohashima, M. Elhoseny, and W. M. Mohamed, "HYBRID QUANTUM CONVOLUTIONAL NEURAL NETWORKS MODEL FOR COVID-19 PREDICTION USING CHEST XRay IMAGES," *arXiv preprint arXiv:2102.06535*, 2021.
46. N. Asadoorian, S. Yaraghi, and A. Tahmasian, "Pre-trained quantum convolutional neural network for COVID-19 disease classification using computed tomography images," *PeerJ Computer Science*, vol. 10, p. e2343, 2024.
47. L. Sünkel et al., "Hybrid quantum machine learning assisted classification of COVID-19 from computed tomography scans," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023, vol. 1: IEEE, pp. 356-366.
48. K. Sengupta and P. R. Srivastava, "Quantum algorithm for quicker clinical prognostic analysis: an application and experimental study using CT scan images of COVID-19 patients," *BMC Medical Informatics and Decision Making*, vol. 21, no. 1, p. 227, 2021.
49. J. Amin, M. Sharif, N. Gul, S. Kadry, and C. Chakraborty, "Quantum machine learning architecture for COVID-19 classification based on synthetic data generation using conditional adversarial neural network," *Cognitive computation*, vol. 14, no. 5, pp. 1677-1688, 2022.

 doi:10.20944/preprints202605.0229.v1

50. M. J. Umer, J. Amin, M. Sharif, M. A. Anjum, F. Azam, and J. H. Shah, "Retracted: An integrated framework for COVID-19 classification based on classical and quantum transfer learning from a chest radiograph," *Concurrency and Computation: Practice and Experience,* vol. 34, no. 20, p. e6434, 2022.
51. V. Kulkarni, S. Pawale, and A. Kharat, "A classical–quantum convolutional neural network for detecting pneumonia from chest radiographs," *Neural Computing and Applications,* vol. 35, no. 21, pp. 15503-15510, 2023.
52. Y. Li and L. Xia, "Coronavirus Disease 2019 (COVID-19): Role of Chest CT in Diagnosis and Management," *AJR Am J Roentgenol,* vol. 214, no. 6, pp. 1280-1286, 2020.
53. S. Kooraki, M. Hosseiny, L. Myers, and A. Gholamrezanezhad, "Coronavirus outbreak: what the department of radiology should know," *J Am Coll Radiol,* vol. 17, no. 4, pp. 447-451, 2020.
54. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc CVPR*, 2016, pp. 770-778.
55. M. Raghu, C. Zhang, J. Kleinberg, and S. Bengio, "Transfusion: Understanding transfer learning for medical imaging," in *NeurIPS*, 2019, pp. 3342-3352.
56. J. Biamonte, "Quantum machine learning," *Nature,* vol. 549, pp. 195-202, 2017.
57. M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*. Springer, 2018.
58. V. Havlíček, A. D. Córcoles, and K. Temme, "Supervised learning with quantum-enhanced feature spaces," *Nature,* vol. 567, pp. 209-212, 2019.
59. M. Schuld and N. Killoran, "Quantum machine learning in Hilbert spaces," *Physical Review Letters,* vol. 122, p. 040504, 2019.
60. J. Chaharlang, M. Mosleh, and S. Rasouli-Heikalabad, "A novel quantum steganography-Steganalysis system for audio signals," *Multimedia Tools and Applications,* vol. 79, no. 25, pp. 17551-17577, 2020.
61. J. Chaharlang, M. Mosleh, and S. Rasouli Heikalabad, "A novel quantum audio steganography–steganalysis approach using LSFQ-based embedding and QKNN-based classifier," *Circuits, Systems, and Signal Processing,* vol. 39, no. 8, pp. 3925-3957, 2020.
62. J. Chaharlang, M. Mosleh, and S. Rasouli-Heikalabad, "Quantum reversible audio steganalysis using quantum Schmidt decomposition and quantum support vector machine," *Journal of Information Security and Applications,* vol. 82, p. 103755, 2024.
63. C. Blank, D. Park, J. K. Rhee, and F. Petruccione, "Quantum classifiers with tailored quantum kernels," *npj Quantum Inf,* vol. 6, p. 41, 2020.
64. M. Henderson, S. Shakya, S. Pradhan, and T. Cook, "Quanvolutional neural networks: powering image recognition with quantum circuits," *Quantum Mach Intell,* vol. 2, p. 2, 2020.
65. E. H. Houssein, "HQ-CNN with Random Quantum Circuits for CXR," *J Comput Des Eng,* vol. 9, no. 2, pp. 343-363, 2022.
66. H. Zhao, "Hybrid quantum–classical CNN for COVID-19 CT classification," *J Biomol Struct Dyn,* vol. 42, no. 7, pp. 3737-3746, 2024.
67. J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum,* vol. 2, p. 79, 2018.
68. H. Zhao, X. Deng, H. Shao, and Y. Jiang, "COVID-19 diagnostic prediction on chest CT scan images using hybrid quantum-classical convolutional neural network," *Journal of Biomolecular Structure and Dynamics,* vol. 42, no. 7, pp. 3737-3746, 2024.
69. P. Singh and S. R. Joshi, "Quantum-Classical Hybrid Approach for COVID-19 Severity Classification from Chest CT Images," 2024.
70. S. A. Salehi and H. Naderi, "Quantum-Assisted Deep Learning: A Hybrid Approach for Robust COVID-19 Diagnosis in Medical Imaging," 2025.