Systematic Review of Quantum Machine Learning for Healthcare

Matheus R. Teixeira[1], **Reinaldo A. C. Bianchi**[1]

[1]Department of Electrical Engineering
FEI University Center – São Bernardo do Campo, SP – Brazil

mteixeira@fei.edu.br, rbianchi@fei.edu.br

Abstract. *Quantum machine learning (QML) has been treated as a promising approach for various fields, among them biomedical and digital health. However, empirical evidence supporting its application in these fields remains limited and generally restricted to studies on simulators. Thus, a PRISMA-aligned systematic review of QML studies published between 2018 and 2025 is conducted, using a predefined PICOC framework, stepwise screening and quantitative quality assessment to evaluate run configurations, dataset characteristics, validation protocols, classical baselines and reproducibility. In essence, the review points out that QML in the healthcare sector remains promising, but is still in a preliminary methodological phase, with the evidence yet subject to the limitations of the NISQ era and the prevalence of studies conducted solely on simulators.*

1. Introduction

Machine learning has become one of the main computational pillars in the medical field, performing tasks such as image-assisted diagnosis, analysis of physiological signals, prognosis prediction, and discovery of drugs [Maheshwari et al. 2022, Gupta et al. 2025]. As these areas began to handle increasingly large volumes of data and, almost always, limited sample sets, practical constraints began to surface, particularly in terms of the scalability, generalizability, and interpretability of classical methods [Gupta et al. 2025, Cerezo et al. 2022]. Considering this scenario, quantum computing is being viewed as a potential extension of traditional learning methods, paving the way for the emergence of a field known as quantum machine learning [Biamonte et al. 2017, Schuld and Petruccione 2021, Cerezo et al. 2022].

However, although this field has grown, its empirical and methodological maturity remains rather limited [Gupta et al. 2025, Cerezo et al. 2022]; most studies are conducted exclusively in simulators, almost always using small datasets, simplified validation protocols, and incomplete descriptions of experimental setups [Maheshwari et al. 2022]. Additionally, there is significant methodological heterogeneity regarding algorithms, encoding strategies, evaluation metrics, and comparisons with classical baselines, which ultimately compromises the reproducibility of the work [Gupta et al. 2025, Schuld et al. 2021]; for these reasons, this study carries out a systematic review aligned with the PRISMA protocol, covering QML studies in healthcare applications published between 2018 and 2025, with the aim of mapping the state of the field.

2. Models and Limitations in Quantum Machine Learning

2.1. Data Encoding

In medical applications, data types (i.e., images, omics data, biosignals, and EHR/tabular data) must not be fed directly into a quantum circuit; instead, classical preprocessing is required to reduce dimensionality or generate compact features before encoding them into a small number of qubits [Gupta et al. 2025, Maheshwari et al. 2022]. The rationale is simple: the costs of encoding and the fact that more qubits mean more noise limit the input dimension that can be encoded without extensive state preparation and the number of qubits to be utilised [Schuld and Petruccione 2021, Cerezo et al. 2022]. Due to these factors, encoding methods vary; the most common ones are:

- **Angle/phase encoding:** Attributes are mapped to rotation angles, for example, using $R_y(x_j)$ or $R_z(x_j)$ gates. It is widely used in NISQ-era applications because it is shallow and hardware-compatible, but on the other hand, it can limit representational capacity unless repeated data re-upload is used [Schuld and Petruccione 2021, Cerezo et al. 2022].
- **Pauli-Z and Pauli-ZZ encodings:** An alternative set of feature maps is based on diagonal unitaries generated by commuting Pauli operators, typically associated with IQP-type constructions [Havlíček et al. 2019, Schuld and Petruccione 2021]. For Pauli-Z encodings, the attributes are encoded using single-qubit phase terms, commonly in the form $\exp\Big(i\sum_j \phi_j(\mathbf{x})Z_j\Big)$, whilst Pauli-$ZZ$ encodings also include two-qubit interaction terms, such as $\exp\Big(i\sum_{j<k} \phi_{jk}(\mathbf{x})Z_jZ_k\Big)$, so as to incorporate correlations between pairs of features. In practical terms, Pauli-Z encodings tend to be simpler and more reliable under NISQ constraints, while Pauli-ZZ variants might offer greater expressive power at the cost of additional entanglement structure and higher noise sensitivity [Schuld et al. 2021, Bansal et al. 2025, Repetto et al. 2025, Cerezo et al. 2022].
- **Amplitude encoding:** Attributes are embedded in the amplitudes of a 2^n dimensional state. This is a conceptually dense method that tends to require deep state preparation or strong abstractions, which can be restrictive on NISQ devices [Nielsen and Chuang 2010].

A quantum feature map is a parameterized unitary $U_\phi(\mathbf{x})$ that generates an embedding state $|\phi(\mathbf{x})\rangle = U_\phi(\mathbf{x})|0\rangle^{\otimes n}$. Its usefulness is evident in quantum kernel methods, in which the similarity between two inputs is defined by a kernel derived from superposition, typically $k(\mathbf{x},\mathbf{x}') = |\langle\phi(\mathbf{x})|\phi(\mathbf{x}')\rangle|^2$ [Havlíček et al. 2019, Schuld and Petruccione 2021].

2.2. Quantum Support Vector Machines (QSVM/QSVC)

Support Vector Machines (SVMs) are margin-based classifiers wherein performance is dependent on the choice of kernel. In the case of QSVM/QSVC, the kernel matrix is estimated by a quantum circuit that essentially implements a feature map; then, the kernel values are used by a classical support vector machine solver [Rebentrost et al. 2014, Havlíček et al. 2019, Schuld 2021]. Support vector margin selection remains classical,

while the quantum component is responsible only for feature space and kernel evaluations. From the feasibility perspective in the NISQ era, QSVM has some practical advantages that justify its prevalence in healthcare studies [Maheshwari et al. 2022, Gupta et al. 2025]:

1. **Shallow quantum kernel:** kernel estimation circuits can be constructed with low depth, especially with feature maps of few qubits, which favors execution on short-term devices.
2. **Stability and interpretation:** because they are based on the conceptual foundation of classical SVM, they tend to be more stable and better understood.
3. **Clear benchmark with classical SVM:** there is a direct baseline, which facilitates comparisons when preprocessing and partitioning are controlled [Gupta et al. 2025, Schuld and Killoran 2022].

Despite this, QSVM has notable limitations: in quantum hardware, the kernel estimation is subject to finite shots and device noise, which can distort the kernel matrix and harm generalization. Moreover, if the feature map is too expressive or too deep, the potential benefits are likely to be negated by the accumulation of noise and the cost of the multiple circuit evaluations required to fill the kernel matrix [Schuld et al. 2021, Cerezo et al. 2022].

2.3. Variational Models: VQC, QNN, and Quantum Deep Learning

Variational quantum classifiers combine an input encoding step with a trainable parameterized circuit (ansatz) $U(\theta)$, resulting in the state $|\psi(\mathbf{x}, \theta)\rangle = U(\theta)U_\phi(\mathbf{x})|0\rangle^{\otimes n}$. Predictions are made by measuring observables (i.e., Z in an output qubit) then mapping the expected values to class labels. The parameters are adjusted by a classical optimizer with the goal of minimizing a loss function [Schuld and Petruccione 2021, Cerezo et al. 2022]. This model is frequently employed in studies on quantum machine learning in healthcare due to its flexibility in terms of architecture and high compatibility with hybrid pipelines (i.e., a classical feature extractor combined with a quantum classifier) [Maheshwari et al. 2022, Gupta et al. 2025].

The key modeling choices are: ansatz configuration, readout (e.g., one or multiple qubits, multi-class strategies), and the optimization method. While high expressiveness improves the fit, it also makes it more susceptible to barren plateaus and device-derived noise [Cerezo et al. 2022, Schuld et al. 2021]. In various studies, the term quantum neural network (QNN) is widely used to refer to variational circuits that are used as function approximators, with multiple layers of data reupload and entanglement gates, and in some cases, with multiple output measures integrated into a larger hybrid network [Schuld and Petruccione 2021, Cerezo et al. 2022].

Quantum deep neural networks (QDNNs) are fundamentally a stack of these layers that aim to emulate the benefits associated with depth, as in classical deep learning. In practice, for medical tasks in the NISQ era, the depth is limited; therefore, "deep" almost always means the repetition of regular blocks with some degree of regularization and a substantial reduction in the dimensionality of the input. The main limitation of QNNs and QDNNs within NISQ constraints is precisely trainability. As depth and the number of qubits increase, gradients may converge near zero and optimization becomes unstable; therefore, deeper circuits have a tendency to amplify

noise accumulation, making the transition from simulator to hardware more challenging [Cerezo et al. 2022, Schuld et al. 2021].

Quantum convolutional neural networks exploit the inductive bias of convolutions in quantum circuits by applying blocks of convolutions followed by pooling operations, which reduces the number of active qubits, typically through conditional measurements and operations. In essence, this architecture can capture local correlations by controlling the number of parameters and enabling the extraction of hierarchical features; while in biomedical applications, QCNN-inspired architectures are primarily used for image classification tasks following classical feature extraction or heavy preprocessing, because direct quantum convolution on high-dimensional images is not feasible on NISQ-era devices [Gupta et al. 2025]. QCNN methods focus on locality, pooling schemes (for example, qubit removal), and readout strategies; the objective is the same as that of classical CNNs: a strong inductive bias can reduce the complexity of the sample, but maximum depth and width are still bound by device limitations [Cerezo et al. 2022, Schuld et al. 2021].

2.4. Constraints on QML Biomedical Tasks in the NISQ Era

As stated earlier, most studies on QML in healthcare are limited by the current state of the NISQ era, which is characterized by noise and decoherence that restrict circuit depth and reduce measurement reliability; consequently, many studies in the literature rely on simulators, while demonstrations on quantum hardware typically consist of reduced-scale instances aimed at verifying feasibility rather than yielding clinically relevant results [Gupta et al. 2025, Schuld et al. 2021]. The problem is that these restrictions end up molding architectural decisions, leading to shallower feature maps, a smaller number of qubits, and heavy classical preprocessing [Maheshwari et al. 2022, Cerezo et al. 2022]. Applications are generally performed using small datasets, with limited external validation and incomplete reporting of preprocessing and comparability with baselines [Gupta et al. 2025, Caro et al. 2022]. Thus, the systematic review seeks to report results with caution and prioritize studies with clear protocols, reasonable baselines, and a high degree of reproducibility (as assessed in the quality checklist in Section 4.5).

3. Methodology

The systematic review of the literature was carried out in accordance with the PRISMA 2020 guidelines, with the aim of mapping and compiling evidence on the application of Quantum Machine Learning to healthcare problems from 2018 to 2025. The methodological focus is as follows: mapping applications, tasks, and data used and evaluating evidence, differentiating implementations in simulators and quantum hardware, under the constraints of the NISQ (Noisy Intermediate-Scale Quantum) era.

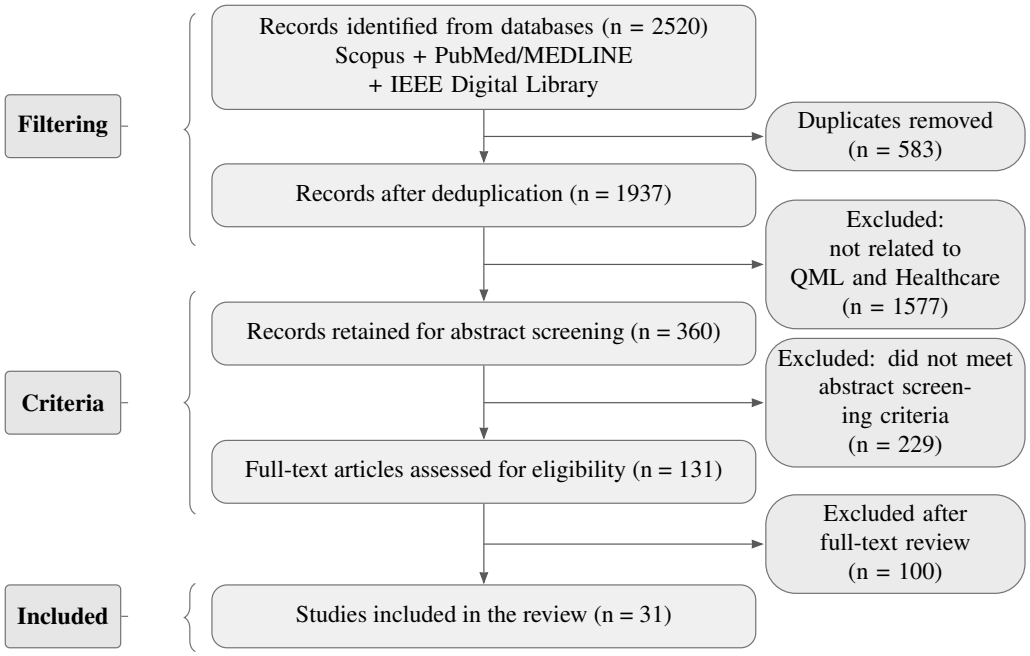


Figure 1. PRISMA flow diagram of the study selection process.

4. PICOC and Research Questions

The research questions and the review plan were structured using the PICOC (Population, Intervention, Comparison, Outcomes, and Context) method to clarify the scope, key interventions, and outcomes being evaluated.

4.1. PICOC Structure

- **P (Population):** Biomedical/medical tasks and datasets, including (but not limited to): medical imaging, omics (genomics, transcriptomics, proteomics), biosignals (ECG, EEG, and others), electronic health records (EHR), drug discovery, diagnosis, and prognosis.
- **I (Intervention):** QML methods, including: variational circuits/variational algorithms (VQC/VQA/PQC), quantum neural networks (QNN) and quantum convolutional neural networks (QCNN), quantum kernels/QSVM methods, and quantum-classical hybrid pipelines.
- **C (Comparison):** Classical baselines when available (SVM, RF, boosting, CNN, etc.); not required for inclusion but used for evidence stratification.
- **O (Outcome):** Quantitative performance metrics and efficiency/resource metrics (including, when reported, time, computational cost, qubits, circuit depth, shots, noise models, etc.), as well as reproducibility indicators (code/data, methodological details).
- **C (Context):** NISQ constraints; execution on simulator versus quantum hardware; nature of publication (peer-reviewed versus preprint).

4.2. Research Questions

The research questions (RQs) that guide this review are:

1. Which medical and biomedical problems have been addressed using QML methods, and what empirical evidence (simulation or quantum hardware) exists for performance, accuracy, or efficiency relative to classical baselines under NISQ-era constraints?
2. What QML models are used in medical applications (e.g., VQC/VQA, quantum kernels/QSVM, QNN/QCNN, hybrid pipelines), and which medical domains do they target?
3. How does evidence strength vary by execution setting (hardware vs simulator), dataset type (public vs private), and evaluation design (validation strategy, uncertainty reporting)?
4. When classical baselines are reported, what comparative outcomes are observed, and are comparisons methodologically fair (same splits/preprocessing; comparable tuning/capacity)?
5. What recurring barriers limit clinical translation of QML approaches (encoding cost, limited qubits/noise, scalability, reproducibility, interpretability, clinical constraints)?
6. Which medical application areas appear most promising for near-term practical QML or hybrid quantum–classical approaches, based on evidence and feasibility?

4.3. Databases and Search Strategy

The literature search was done in digital libraries selected for their breadth and relevance to QML and health: **Scopus**; **PubMed/MEDLINE**; **IEEE Xplore / IEEE Digital Library**.

In comparison with [Gupta et al. 2025, Maheshwari et al. 2022], this review combines biomedical and computing/engineering databases to reduce indexing bias and expand interdisciplinary coverage, while focusing on 2018–2025, PRISMA screening, quality assessment, simulator/hardware differentiation, baselines, and reproducibility.

The queries returned a combined total of 2,520 results, distributed across databases as follows: 1,793 (Scopus), 361 (PubMed/MEDLINE), and 366 (IEEE).

4.4. Eligibility Criteria

The inclusion and exclusion criteria aimed to maximize relevance to healthcare QML and minimize conceptual papers without empirical evaluation; a single reviewer performed screening, full-text review, and quality assessment, re-evaluating unclear cases following predefined criteria.

Table 1. Inclusion and exclusion criteria

Criterion Type	Description
Inclusion	Empirical evaluation with quantitative results, using simulator and/or quantum hardware.
Inclusion	Explicit quantum machine learning method, including variational circuits/algorithms (VQC, VQA, variational classifiers), parameterized quantum circuit learning (PQC), quantum neural networks (QNN, QCNN), quantum support vector machines and quantum kernel methods (QSVM, quantum kernels, quantum feature maps), or hybrid quantum–classical machine learning pipelines.
Inclusion	Medical or biomedical relevance in terms of task, dataset, or objective (e.g., imaging, omics, biosignals, EHR, diagnosis, prognosis, drug discovery).
Inclusion	Publication window between 2018 and 2025; earlier studies included only if clearly seminal and necessary for contextual completeness.
Inclusion	English language, sufficient bibliographic metadata, and full text available for final synthesis.
Exclusion	Duplicate records or studies without accessible full text at the final screening stage.
Exclusion	No empirical evaluation (conceptual papers, editorials, perspectives, or review-only articles; reviews retained exclusively for backward or forward snowballing).
Exclusion	Not quantum machine learning (quantum computing without machine learning, or classical machine learning without a quantum component).
Exclusion	Not medical or biomedical in scope.
Exclusion	Outside the defined date range or language constraints, unless justified as a seminal exception.
Exclusion	Quantum-inspired or ambiguously labeled "quantum" approaches lacking a technically specified quantum machine learning implementation.

4.5. Quality Assessment

The 10-item checklist (Q1–Q10) assigned 2, 1, or 0 points to *Yes*, *Partially*, or *No*; the maximum score was 20 and the minimum-quality cutoff was 12. Only studies that met the minimum quality criteria have been included in the final list.

Table 2. Quality assessment checklist and scoring scheme

ID	Quality assessment question	Scoring (0–2)
Q1	Is the medical/biomedical task and objective clearly defined and appropriate (data/task/outcome is biomedical/clinical)?	Yes=2; Partially=1; No=0
Q2	Is the QML method described with sufficient technical detail to identify the paradigm (e.g., VQC/VQA/PQC, QSVM/kernels, QNN/QCNN, hybrid pipeline)?	Yes=2; Partially=1; No=0
Q3	Does the study clearly state whether results are from quantum hardware, simulation, or both, and describe the setup?	Yes=2; Partially=1; No=0
Q4	Are the dataset(s) clearly described (source, type, size) and appropriate for the task?	Yes=2; Partially=1; No=0
Q5	Is the evaluation protocol appropriate and clearly reported (train/test split, cross-validation, repeated runs, etc.)?	Yes=2; Partially=1; No=0
Q6	Are outcome metrics appropriate and sufficiently reported (including variability/uncertainty when relevant)?	Yes=2; Partially=1; No=0
Q7	If classical baselines are used, are comparisons methodologically fair (same data splits, comparable preprocessing, reasonable tuning/capacity)?	Yes=2; Partially=1; No=0
Q8	Does the study report computational/quantum resource information relevant to feasibility (qubits, circuit depth, shots, noise model, runtime), or justify feasibility under NISQ constraints?	Yes=2; Partially=1; No=0
Q9	Is the work reproducible (code/data availability OR sufficiently detailed methods to reproduce)?	Yes=2; Partially=1; No=0
Q10	Does the study discuss limitations, potential bias sources (e.g., small datasets, encoding choices, noise), and external validity?	Yes=2; Partially=1; No=0

Note: Response scale: Yes=2, Partially=1, No=0. Maximum score: 20. Studies scoring ≥ 12 were considered to meet the quality cutoff.

5. Results and Discussion

The collection of included studies ranges across a broad set of clinical and biomedical tasks (diagnosis, stratification, outcome prediction, screening, and drug discovery), with a typical proof-of-concept design.

Table 3. Summary of Quality Assessment scores for included studies, based on the Quality Assessment checklist (Table 2).

Article	Score	Band
Very high (20)		
ResQ: A hybrid classical–quantum model for efficient breast cancer image classification	20	Very high
CQ-CNN: A lightweight hybrid classical–quantum convolutional neural network for Alzheimer's disease detection using 3D structural brain MRI	20	Very high
Performance of Quantum Annealing Machine Learning Classification Models on ADMET Datasets	20	Very high
Expanding the Horizon: Enabling Hybrid Quantum Transfer Learning for Long-Tailed Chest X-Ray Classification	20	Very high
Hybrid Classical–Quantum Transfer Learning for Cardiomegaly Detection in Chest X-rays	20	Very high
High (18–19)		
Quantum enhanced stratification of breast cancer: exploring quantum expressivity for real omics data	19	High
Generating three-dimensional genome structures with a variational quantum algorithm	19	High
Quantum machine learning framework for virtual screening in drug discovery: a prospective quantum advantage	19	High
Clinical data classification with noisy intermediate scale quantum computers	19	High
Quantum transfer learning for breast cancer detection	19	High
Hybrid classical–quantum convolutional neural network for stenosis detection in X-ray coronary angiography	19	High
Quantum Generative Modeling of Single-Cell transcriptomes: Capturing Gene–Gene and Cell–Cell Interactions	19	High
Functional near-infrared spectroscopy for the detection of fear using parameterized quantum circuits	19	High
Hybrid quantum-classical deep learning framework for balanced multiclass diabetic retinopathy classification	19	High
Fusion-aware quantum variational autoencoder for brain-heart signal modeling in mental health applications	19	High
Pre-trained quantum convolutional neural network for COVID-19 disease classification using computed tomography images	19	High
Quantum neural networks meet federated learning for DNA mutation prediction	19	High
Case-Based and Quantum Classification for ERP-Based Brain–Computer Interfaces	18	High
Quantum splitting convolutional neural network-based distributed quantum disease detection model	18	High
Automatic malaria disease detection from blood cell images using the variational quantum circuit	18	High
Hybrid classical and quantum computing for enhanced glioma tumor classification using TCGA data	18	High
Evaluation of conventional and quantum computing for predicting mortality based on small early-onset colorectal cancer data	18	High
Moderate (15–17)		
Hybrid Quantum Neural Network for Drug Response Prediction	17	Moderate
Enhancing quantum support vector machine for healthcare applications using custom feature maps	17	Moderate
An integrated strategy based on radiomics and quantum machine learning: diagnosis and clinical interpretation of pulmonary ground-glass nodules	16	Moderate
Scalable variational quantum circuits for autoencoder-based drug discovery	16	Moderate
Design and analysis of quantum powered support vector machines for malignant breast cancer diagnosis	16	Moderate
Classical and Quantum SVM for Electromyography-Based Myopathy Detection: A Comparative Exploration	16	Moderate
Enhancing Explainability of Quantum-Based Stacking for Mental Health Assessment: Predicting Depression Risk in AIDS Patients	16	Moderate
Predicting Brain Age and Gender from Brain Volume Data Using Variational Quantum Circuits	16	Moderate
Low (12–14)		
QuEEGNet: Quantum AI for Biosignal Processing	14	Low

Note: Bands are Low (12–14), Moderate (15–17), High (18–19), and Very high (20). Scores are totals from ten items (Yes = 2, Partially = 1, No = 0; maximum = 20).

5.1. Studies by Domain

The included studies have consistently adopted a proof-of-concept strategy in which classical representation learning (e.g., CNN features, radiomics, dimensionality reduction) precedes a quantum model with low qubit counts (e.g., variational classifiers or kernel methods) [Azevedo et al. 2022, Decoodt et al. 2023, Asadoorian et al. 2024, Repetto et al. 2025, Vashisth et al. 2021, Moradi et al. 2022].

5.2. Oncology

The field of oncology has the largest comparable body of evidence, as it includes several modalities (histopathology, mammography, radiomics, and omics), the use of different QML methods (transfer learning and kernel methods), and more reproducible results [Guha et al. 2025, Azevedo et al. 2022, Repetto et al. 2025, Akpinar and Oduncuoglu 2025, Huang et al. 2025].

5.2.1. Breast Cancer

Table 4. Breast cancer applications of QML (imaging and omics)

Study	Task / data	QML approach	Exec.	Evaluation / main findings (as reported)
ResQ: A hybrid classical-quantum model for efficient breast cancer image classification [Guha et al. 2025]	Histopathology image classification; BreakHis and Bioimaging (BI).	Hybrid ResNet feature extractor + VQC classifier ("ResQ"); shallow 1-qubit circuit analysis; statistical tests (Friedman/Nemenyi).	Sim+HW	Accuracy: 98.49% (BreakHis) and 88.96% (BI); classical counterparts: 97.36% and 87.96%. Evaluated on Aer simulator and real IBM device (ibm_brisbane).
Quantum transfer learning for breast cancer detection [Azevedo et al. 2022]	Full-image mammogram classification (malignant vs benign); BCDR.	Hybrid classical–quantum transfer learning; heatmaps for interpretability.	Sim+HW	Accuracy: 84% vs prior state-of-the-art 76.9% (as claimed). Experiments included a real quantum device and comparison with simulator/classical.
Quantum enhanced stratification of breast cancer: exploring quantum expressivity for real omics data [Repetto et al. 2025]	Breast cancer subtype stratification using real omics data (METABRIC; N=1980; dimensionality reduction to 4D via UMAP).	Quantum kernels (QK) with encoding/entanglement sweep (e.g., Z; ZZ-linear; ZZ-full); expressivity vs performance trade-off; QPU noise-resilience study.	HW	Subtype stratification on METABRIC data. Quantum kernels achieve performance comparable to classical kernels with fewer samples in simulation; on hardware, simpler (non-entangled) encodings show higher robustness to noise than expressive feature maps.
Design and analysis of quantum powered support vector machines for malignant breast cancer diagnosis [Vashisth et al. 2021]	Malignant breast cancer diagnosis on Wisconsin Diagnostic Breast Cancer (WDBC): n=569 (357 benign / 212 malignant), 30 features; PCA to 2 features (2-qubit constraint).	QSVM / quantum-powered SVM with 2-qubit kernel circuit (PQC depth=2); compared to classical linear and non-linear SVM.	Sim	10-fold CV. Accuracy: 92% (QSVM), 89% (linear SVM), 95% (non-linear SVM). QSVM precision/recall/F1: 93%/91%/92% (as reported).

Note: Sim = quantum simulator; HW = execution on real quantum hardware.

5.3. Cardiovascular imaging

The cardiovascular imaging area follows the transfer learning approach, where quantum layers serve as classifiers for pretrained CNN feature extraction. While the reported metrics are promising, the results are still heavily based on simulator evaluation [Decoodt et al. 2023, Ovalle-Magallanes et al. 2022].

Table 5. Cardiovascular imaging applications of hybrid QML

Study	Task / data	QML approach	Exec.	Evaluation / main findings (as reported)
Hybrid Classical–Quantum Transfer Learning for Cardiomegaly Detection in Chest X-rays [Decoodt et al. 2023]	Cardiomegaly detection in chest X-rays; CheXpert subset; balanced 2436 PA CXRs (different patients).	Hybrid classical–quantum transfer learning: DenseNet-121 + parameterized quantum circuit (Qiskit/PennyLane); trainability via effective dimension; Grad-CAM++ analysis.	Sim	k-fold CV on statevector simulator. ROC AUC up to 93%; accuracy up to 87%. Grad-CAM++ "hot zone" aligned with heart more often for CQ vs CC (94% vs 61%, $p<0.001$).
Hybrid classical–quantum convolutional neural network for stenosis detection in X-ray coronary angiography [Ovalle-Magallanes et al. 2022]	Stenosis vs non-stenosis in X-ray coronary angiography; 250 image patches (balanced).	Hybrid transfer learning: pretrained ResNet-18 + variational quantum circuit; L2–tanh normalization; distributed VQC (2×2-qubit) to reduce training time.	Sim	Hybrid model outperforms classical transfer learning: +9% accuracy, +20% recall, +11% F1. Best accuracy 91.80%, recall 94.92%, F1 91.80% on quantum simulator.

Note: Sim = quantum simulator; no quantum hardware execution reported.

5.4. Clinical tabular data

Reproducibility and comparisons can be hindered when datasets are proprietary to institutions or when baseline tuning and preprocessing are not fully documented [Moradi et al. 2022, Yu et al. 2024].

Table 6. Clinical tabular data applications and NISQ/hardware feasibility evaluations.

Study	Task / data	QML approach	Exec.	Evaluation / main findings (as reported)
Clinical data classification with noisy intermediate scale quantum computers [Moradi et al. 2022]	Clinical data classification on three open-access datasets (names not reported).	Two QML algorithms: quantum distance classifier (qDS) and simplified quantum-kernel SVM (sqKSVM); uses $\log_2 N$ encoding; inner products estimated on IBMQ Melbourne (15 qubits).	HW	Reports performance comparable to prior QML and classical counterparts; encoding strategy yields up to +2% AUROC (reported across tasks). Discusses when qDS vs kernel methods perform better (small vs larger sample/feature counts).
Evaluation of conventional and quantum computing for predicting mortality based on small early-onset colorectal cancer data [Yu et al. 2024]	Mortality prediction in early-onset colorectal cancer; retrospective tertiary hospital data; 1501 patients (2008–2020); 93 features; mortality rate 7.6%.	QSVM vs classical SVM; explores feature subsets (e.g., top 11 variables), training size, and class imbalance. Implemented in Qiskit (v1.0.2).	Sim	Reports QSVM AUROC mean (SD) 86.3% (10.2%) vs conventional SVM 72.3% (23.1%); reports robustness under imbalance (as reported). Experiments executed using a statevector simulator backend.

Note: Sim = quantum simulator; HW = execution on real quantum hardware.

6. Conclusion

Considering pipelines, two types of approaches were used. First, a hybrid approach using transfer learning with a parameterized circuit on top of a classical feature extractor. Second, quantum-kernel methods (QSVM/QSVC) are used more often when the data can be mapped to a low-dimensional space. In the studies reviewed, the most common limitations were small sample sizes, significant dimensionality reduction, few qubits, noise, reliance on simulators, limited external validation, and inadequate metric reporting, all of which compromised reproducibility and clinical translation. Consequently, assumptions of "quantum advantage" are still dependent on the encoding method, circuit size, presence of noise and baseline selection. Overall, the data needs to be evaluated with caution in distinguishing promising results from evidence that is still preliminary.

Acknowledgments

The authors acknowledge the FAPESP (grants 2019/07665-4 and 2024/15527-9) and CNPq (grant 309525/2026-8) for supporting this project. This study was funded in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

References

Akpinar, E. and Oduncuoglu, M. (2025). Hybrid classical and quantum computing for enhanced glioma tumor classification using TCGA data. *Scientific Reports*, 15(1):25935.

Asadoorian, N., Yaraghi, S., and Tahmasian, A. (2024). Pre-trained quantum convolutional neural network for COVID-19 disease classification using computed tomography images. *PeerJ Computer Science*, 10:e2343.

Azevedo, V., Silva, C., and Dutra, I. (2022). Quantum transfer learning for breast cancer detection. *Quantum Machine Intelligence*, 4(1):5.

Bansal, R., Rajput, N. K., and Khanna, M. (2025). Enhancing quantum support vector machine for healthcare applications using custom feature maps. *Knowledge-Based Systems*, 320:113669.

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., and Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671):195–202.

Caro, M. C., Huang, H.-Y., Cerezo, M., Sharma, K., Sornborger, A., Cincio, L., and Coles, P. J. (2022). Generalization in quantum machine learning from few training data. *Nature Communications*, 13(1):4919.

Cerezo, M., Verdon, G., Huang, H.-Y., Cincio, L., and Coles, P. J. (2022). Challenges and opportunities in quantum machine learning. *Nature Computational Science*, 2(9):567–576.

Decoodt, P., Liang, T. J., Bopardikar, S., Santhanam, H., Eyembe, A., Garcia-Zapirain, B., and Sierra-Sosa, D. (2023). Hybrid classical–quantum transfer learning for cardiomegaly detection in chest x-rays. *Journal of Imaging*, 9(7):128.

Guha, D., Kuiry, S., Mitra, S., Bhattacharyya, S., and Das, N. (2025). ResQ: A hybrid classical-quantum model for efficient breast cancer image classification. *Applied Soft Computing*, 183:113631.

Gupta, R. S., Wood, C. E., Engstrom, T., Pole, J. D., and Shrapnel, S. (2025). Quantum machine learning for digital health? a systematic review. *npj Digital Medicine*, 8(1):237.

Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., and Gambetta, J. M. (2019). Supervised learning with quantum enhanced feature spaces. *Nature*, 567(7747):209–212.

Huang, X., Xu, F., Zhu, W., Yao, L., He, J., Su, J., Zhao, W., and Hu, H. (2025). An integrated strategy based on radiomics and quantum machine learning: diagnosis and clinical interpretation of pulmonary ground-glass nodules. *BMC Medical Imaging*, 25(1):279.

Maheshwari, D., Garcia-Zapirain, B., and Sierra-Sosa, D. (2022). Quantum machine learning applications in the biomedical domain: A systematic review. *IEEE Access*, 10:80463–80484.

Moradi, S., Brandner, C., Spielvogel, C., Krajnc, D., Hillmich, S., Wille, R., Drexler, W., and Papp, L. (2022). Clinical data classification with noisy intermediate scale quantum computers. *Scientific Reports*, 12(1):1851.

Nielsen, M. A. and Chuang, I. L. (2010). *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge.

Ovalle-Magallanes, E., Avina-Cervantes, J. G., Cruz-Aceves, I., and Ruiz-Pinales, J. (2022). Hybrid classical–quantum convolutional neural network for stenosis detection in X-ray coronary angiography. *Expert Systems with Applications*, 189:116112.

Rebentrost, P., Mohseni, M., and Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13):130503.

Repetto, V., Ceroni, E. G., Buonaiuto, G., and D'Aurizio, R. (2025). Quantum enhanced stratification of breast cancer: exploring quantum expressivity for real omics data. *Quantum Machine Intelligence*, 7(2):81.

Schuld, M. (2021). Supervised quantum machine learning models are kernel methods. arXiv:2101.11020.

Schuld, M. and Killoran, N. (2022). Is quantum advantage the right goal for quantum machine learning? *PRX Quantum*, 3(3):030101.

Schuld, M. and Petruccione, F. (2021). *Machine Learning with Quantum Computers*. Springer International Publishing, Cham.

Schuld, M., Sweke, R., and Meyer, J. J. (2021). The effect of data encoding on the expressive power of variational quantum machine learning models. *Physical Review A*, 103(3):032430.

Vashisth, S., Dhall, I., and Aggarwal, G. (2021). Design and analysis of quantum powered support vector machines for malignant breast cancer diagnosis. *Journal of Intelligent Systems*, 30(1):998–1013.

Yu, J. Y., Sim, W. S., Jung, J. Y., Park, S. H., Kim, H. S., and Park, Y. R. (2024). Evaluation of conventional and quantum computing for predicting mortality based on small early-onset colorectal cancer data. *Applied Soft Computing*, 162:111781.