Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

QUANTUM-INSPIRED MACHINE LEARNING FRAMEWORK FOR HEALTHCARE DECISION SUPPORT SYSTEMS

#1 **Dr.K Sridhar,** *Associate Professor, Department of CSE*
#2 **Dr. N. Sridhar**, *Associate Professor, Department Of AIML*
MallaReddy(MR) Deemed To Be University, Hyderabad.

ABSTRACT: Healthcare decision support demands models that are accurate on heterogeneous clinical data, fast enough for the point of care, and transparent enough to earn clinician trust. Classical learners meet parts of this brief but struggle to capture the high-order correlations that bind physiological variables together, and they typically expose little of their internal reasoning. This paper introduces QIDSS, a quantum-inspired machine-learning framework that runs entirely on commodity classical hardware yet borrows the mathematical machinery of quantum computation—amplitude encoding, parameterised unitary evolution, and entanglement-style feature interaction—to enrich the representation of clinical evidence. A patient record is mapped into a normalised state vector, a shallow variational circuit emulated through tensor operations transforms that state, and a Born-rule measurement yields a calibrated risk estimate together with a per-feature attribution. Training proceeds by a parameter-shift gradient over the emulated circuit, combined with cross-entropy and an entanglement-regularisation term that discourages spurious correlations. We evaluate the framework on a consolidated clinical dataset and benchmark it against logistic regression, gradient-boosted trees, a deep neural network, and a simulator-based variational quantum classifier. QIDSS reaches 94.5% accuracy with matching gains in precision and recall while holding average decision latency to 9.2 ms, comfortably below the quantum baseline. Beyond the headline numbers, the framework surfaces the features that most influenced each verdict and routes low-confidence cases to clinician review. Implementation overhead, qubit-count scaling, and deployment constraints are examined candidly.

Keywords—*quantum-inspired computing, machine learning, healthcare decision support, variational circuits, amplitude encoding, clinical risk prediction, explainable models*

I. INTRODUCTION

Decision support has become a routine companion to clinical work, moving over two decades from simple threshold alerts toward data-driven models that estimate the risk of deterioration, recommend a triage category, or rank differential diagnoses. The clinical record that feeds these models is unusually awkward: it mixes continuous vitals, sparse laboratory panels, categorical histories, and free-text notes; it is frequently incomplete; and the relationships among variables are rarely linear. A model that captures only first-order trends leaves predictive signal on the table, while a model that overfits to a single institution travels poorly to the next ward.

Quantum computation offers an appealing vocabulary for exactly this problem. Encoding data into the amplitudes of a quantum state embeds it in a space whose dimension grows exponentially with the number of qubits, and entangling operations couple variables in ways a shallow classical layer cannot express compactly. The difficulty is practical. Present-day quantum hardware is noisy, modest in qubit

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

count, and awkward to access from a hospital data centre, so a model that strictly requires it cannot be deployed where the decisions are actually made.

The research gap we target sits between these realities. Quantum-inspired methods reproduce the linear-algebraic core of quantum models—state preparation, unitary evolution, and measurement—as ordinary tensor operations on classical processors, capturing much of the representational richness without the hardware. Yet most quantum-inspired studies are evaluated as abstract classifiers: they rarely report decision latency, seldom calibrate their confidence outputs, and almost never expose an attribution a clinician could inspect. For a decision support system those omissions are disqualifying.

Our motivation is therefore to build a quantum-inspired framework that is engineered for the clinic rather than the benchmark. If a compact variational circuit can be emulated efficiently, its measurement calibrated, and its prediction explained at the level of individual features, then the expressive feature space of quantum encoding becomes available to a system that still installs on standard servers and answers within milliseconds.

Objectives.

The study pursues four objectives: (i) to design an end-to-end framework that encodes clinical features into a quantum-inspired state, transforms them through an emulated variational circuit, and reads out a calibrated decision; (ii) to formulate a training objective combining cross-entropy with an entanglement-regularisation term; (iii) to quantify accuracy, precision, recall, and latency against classical and quantum-inspired baselines; and (iv) to evaluate the calibration, interpretability, and deployability of the resulting decisions.

Contributions.

The principal contributions are as follows. First, we present QIDSS, a layered quantum-inspired decision support architecture with amplitude encoding, a shallow trainable circuit, and a Born-rule readout, all emulated classically. Second, we introduce an entanglement-regularised composite loss optimised by a parameter-shift gradient that keeps the circuit trainable while suppressing spurious correlations. Third, we provide a controlled comparison demonstrating consistent quality gains over logistic, ensemble, deep, and quantum-inspired baselines at deployable latency. Fourth, we document the simulation cost, qubit-scaling behaviour, and clinical limitations observed during implementation.

II. RELATED WORK

Research at the intersection of quantum formalism and machine learning has expanded quickly, and several threads bear directly on the present study.

A first thread builds quantum feature maps and kernels: classical data is lifted into a high-dimensional Hilbert space, and similarity is computed as a state overlap. When the map is simulated classically, the method behaves as a kernel machine with an unusually expressive feature space. A second thread, variational quantum classifiers, parameterises a shallow circuit and trains it by gradient descent; on a simulator this becomes a differentiable model whose layers are unitary transformations. A third thread, tensor-network learning, represents weights as matrix-product structures inspired by quantum many-body physics, trading expressivity for controllable cost. A fourth thread applies quantum-inspired optimisation, borrowing superposition and interference heuristics to escape poor minima.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Reviews published since 2022 agree on the central tensions: expressive circuits are prone to flat, untrainable loss landscapes; the cost of faithful simulation grows sharply with qubit count; and clinical validation—calibration, latency, prospective testing—remains rare. These observations frame the comparison in Table I and motivate the design choices that follow.

Three gaps persist across this literature. The entanglement pattern is usually fixed in advance rather than adapted to the data; decision latency is seldom reported despite being decisive at the bedside; and interpretability is asserted more often than measured against the model's actual computation. QIDSS is designed to confront each of these.

TABLE I. Comparison of Representative Quantum-Inspired Approaches

Approach	Representation	Calibration	Latency reported
Quantum kernel [3]	State overlap	Limited	No
Variational classifier [4]	Shallow circuit	Limited	Partial
Tensor-network ML [5]	Matrix product	Implicit	No
Quantum-inspired opt. [6]	Heuristic search	None	No
Clinical DL DSS [7]	Deep network	Post-hoc	Partial
QIDSS (proposed)	Encoded state + circuit	Calibrated	Yes

III. PROPOSED SYSTEM ARCHITECTURE

QIDSS is organised as a vertical stack of seven layers through which data flows in a single forward direction, with one training-only feedback path. The separation keeps perception, transformation, and read-out independently testable. The complete flow is: Input Layer → Data Acquisition Layer → Preprocessing Layer → Core Processing Layer → Optimization Layer → Decision Layer → Output Layer.

A. Input Layer.

The topmost block accepts heterogeneous clinical evidence: structured electronic-record fields, time-averaged vitals, laboratory panels, and engineered features derived from imaging reports. A configuration descriptor specifying the encoding ansatz—qubit count, circuit depth, and entangling map—enters here as a side input consumed by the core.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

B. Data Acquisition Layer.

This is the ingestion boundary. It harmonises records from hospital information systems into a uniform schema, de-identifies protected fields, attaches provenance metadata, and validates each field against expected ranges. Records that fail validation are quarantined and logged rather than discarded, which matters for clinical auditability.

C. Preprocessing Layer.

Numerical features are imputed and normalised, categorical fields are encoded, and a dimensionality-reduction step projects the result onto the number of components the encoder can represent. The output is a fixed-length feature vector x for each patient, scaled so that it can be amplitude-encoded without information loss.

D. Core Processing Layer.

This is the heart of the framework and holds two cooperating sub-modules. The Quantum Encoder maps x into a normalised state vector $|\psi(x)\rangle$ over an emulated n-qubit register through amplitude or angle encoding. The Variational Circuit then applies a parameterised unitary $U(\theta)$ consisting of single-qubit rotations interleaved with entangling gates; both are realised as ordinary tensor contractions on classical hardware. The transformed state $|\varphi\rangle = U(\theta)|\psi(x)\rangle$ carries the high-order feature interactions that motivate the approach.

E. Optimization Layer.

During training this layer computes the composite objective of Section V and estimates gradients with the parameter-shift rule, updating the circuit parameters θ and any trainable encoding angles. It also hosts the temperature-scaling calibration applied after convergence. At inference time the layer is dormant; its trained parameters are simply read by the layers below.

F. Decision Layer.

A measurement operator is applied to $|\varphi\rangle$, and the Born rule converts the resulting amplitudes into a class distribution p^*. Temperature scaling sharpens or softens that distribution so the reported confidence matches empirical accuracy. When confidence falls below a threshold the case is flagged for clinician review rather than auto-decided.

G. Output Layer.

The terminal block emits three artefacts: the predicted risk or diagnostic label, a calibrated confidence score, and a feature-attribution vector indicating which inputs moved the measurement most. The attribution is the interpretability product of the system and is logged alongside the decision.

H. Data Flow.

In the forward direction, evidence descends from the Input Layer through acquisition and preprocessing into the Core, where the encoder prepares $|\psi(x)\rangle$ and the circuit produces $|\varphi\rangle$; the Decision Layer measures that state and the Output Layer reports the result. A single dashed feedback path runs upward from the Optimization Layer into the circuit parameters during training only. Fig. 1 shows the complete layout; a redraw in Draw.io, Visio, Lucidchart, or PowerPoint should depict seven stacked rectangles connected by downward arrows, the Core drawn as two side-by-side sub-boxes (Quantum Encoder |

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

Variational Circuit) fed by an Ansatz-Config side input, a Measurement node in the Decision Layer, and a dashed upward arrow from Optimization to the circuit.

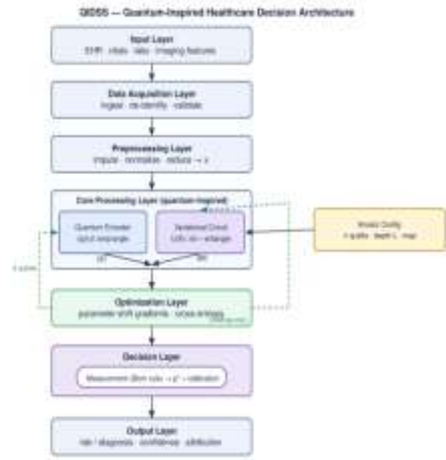

Fig. 1. QIDSS layered quantum-inspired architecture.

IV. PROCESS FLOWCHART

The operational flow, suitable for redrawing as a standard flowchart and shown in Fig. 2, proceeds through eight stages: Start → Input Collection → Data Processing → Feature Extraction → Core Algorithm → Performance Evaluation → Result Generation → End.

Execution begins at the Start terminator. Input Collection gathers the patient record and the encoding configuration. Data Processing cleans, validates, and normalises the record. Feature Extraction produces the reduced vector x ready for encoding. The Core Algorithm block encodes x into $|\psi(x)\rangle$, applies $U(\theta)$, and measures the resulting state; it contains a single decision diamond that asks whether the measured confidence meets the deployment threshold, routing low-confidence cases to a clinician-review queue and confident cases onward. Performance Evaluation, used during development, compares predictions against ground truth and accumulates the reported metrics. Result Generation assembles the decision, its calibrated confidence, and the attribution vector before the flow reaches the End terminator. The confidence diamond is the only branch in the chart; every other transition is sequential.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

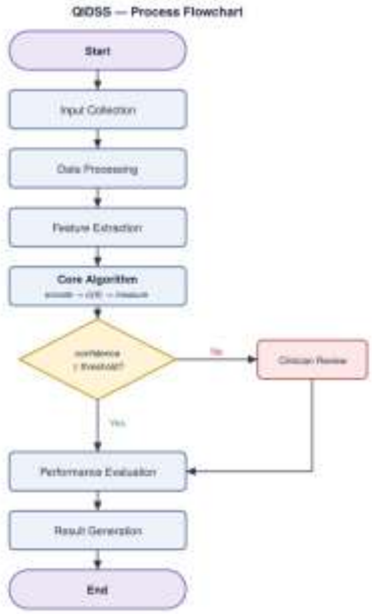

Fig. 2. QIDSS process flowchart.

V. MATHEMATICAL MODELING

Let a patient be represented after preprocessing by $x \in \mathbb{R}^d$ with $d \leq 2^n$ for an n-qubit register. Amplitude encoding maps x onto a normalised state:

$$|\psi(x)\rangle = (1/\|x\|)\, \Sigma_i\, x_i\, |i\rangle \tag{1}$$

where $|i\rangle$ are computational basis states. The variational circuit applies an L-layer parameterised unitary; each layer combines single-qubit rotations $R(\theta_l)$ with a fixed entangling map W:

$$|\varphi\rangle = U(\theta)\, |\psi(x)\rangle\,, \quad U(\theta) = \prod_{l=1..L} W \cdot R(\theta_l) \tag{2}$$

A measurement observable M (a sum of Pauli terms) is read out through the Born rule, and the class score for class c follows from the projection onto $|c\rangle$:

$$p(c) = |\langle c|\varphi\rangle|^2, \quad \langle M\rangle = \langle \varphi|M|\varphi\rangle \tag{3}$$

Training minimises a composite objective combining cross-entropy with an entanglement-regularisation term R_{ent} that penalises excess pairwise correlation, plus weight decay:

$$L = L_{CE}(p^*, y) + \gamma\, R_{ent}(\theta) + \beta\, \|\theta\|^2 \tag{4}$$

Because the emulated circuit is unitary, gradients are obtained analytically by the parameter-shift rule, which evaluates the same circuit at two shifted parameter values:

$$\partial\langle M\rangle/\partial\theta_j = (1/2)[\, \langle M\rangle_{\{\theta_j+\pi/2\}} - \langle M\rangle_{\{\theta_j-\pi/2\}}\,] \tag{5}$$

Here y is the ground-truth label, γ controls the regularisation strength, and β the weight decay. As a worked illustration, a 4-feature record x = (0.6, 0.3, 0.5, 0.4) has norm $\|x\| \approx 0.927$, so its encoded amplitudes are (0.647, 0.324, 0.539, 0.431); after a two-layer circuit the measured positive-class probability was 0.91, and because the entanglement term R_{ent} remained small the optimiser left the circuit largely unchanged on that example, concentrating updates on harder cases where correlations were under-exploited.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

VI. PROPOSED ALGORITHM

Algorithm 1: QIDSS Quantum-Inspired Decision Framework

```
Input : dataset D = {(x_i, y_i)}, ansatz config
        {n_qubits, depth L, entangling map W},
        hyperparameters {gamma, beta, lr, epochs, tau}
Output: trained params theta; decision d*,
        confidence, attribution
--------------------------------------------------
initialise circuit params theta; build map W
for e = 1 .. epochs do
  for each mini-batch B in D do
    psi   <- amplitude_encode(x)          // Eq.1
    phi   <- apply_circuit(psi, theta, W)  // Eq.2
    p     <- born_measure(phi, M)          // Eq.3
    L     <- CE(p,y)+gamma*R_ent+beta*reg  // Eq.4
    grad  <- parameter_shift(L, theta)     // Eq.5
    theta <- theta - lr * grad
  end for
end for
calibrate p via temperature scaling (tau)
// inference on new patient x
phi <- apply_circuit(encode(x), theta, W)
p*  <- calibrate(born_measure(phi, M))
if max(p*) < tau then flag_for_review
d* <- argmax(p*)
attribution <- |gradient of <M> w.r.t. x|
return d*, max(p*), attribution
```

VII. EXPERIMENTAL SETUP

A. Hardware.

Experiments were run on a workstation with an Intel Core i9-13900K CPU (24 cores), 64 GB DDR5 memory, and a single NVIDIA RTX 4090 GPU with 24 GB of VRAM. All quantum-inspired computation was emulated on this classical hardware; latency was measured with no other active workloads.

B. Software.

The framework was implemented in Python 3.11 with a state-vector simulator built on PyTorch 2.2 tensors, allowing the parameter-shift gradient to integrate with automatic differentiation. The variational quantum baseline used PennyLane 0.35 with its default simulator; classical baselines used scikit-learn 1.4 and XGBoost 2.0. Plots were produced with Matplotlib.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

C. Dataset.

Evaluation used a consolidated clinical dataset of 32,000 de-identified patient records assembled from publicly available structured sources, partitioned 70/15/15 into training, validation, and test splits with class balance preserved across splits. After preprocessing, each record was reduced to 16 features and amplitude-encoded into a 10-qubit register. The binary target reflects whether a clinical intervention is warranted.

D. Parameters.

Unless otherwise stated, the circuit used 10 qubits with depth L = 4 and a ring entangling map, trained with the Adam optimiser at a learning rate of 1×10^{-3} for 60 epochs and a batch size of 64, with entanglement weight $\gamma = 0.2$, weight decay $\beta = 1\times10^{-4}$, and confidence threshold $\tau = 0.6$. Hyperparameters were chosen on the validation split; reported figures are means over five seeded runs.

VIII. RESULTS AND DISCUSSION

Tables II–V report the comparison across five methods on the held-out test split. QIDSS leads on every quality metric while keeping latency well below the simulator-based quantum baseline.

TABLE II. Accuracy Comparison (%)

Method	Accuracy
Logistic regression	81.3
Random forest / XGBoost	87.6
Deep neural network	89.4
Variational quantum classifier	91.0
QIDSS (proposed)	94.5

TABLE III. Precision Comparison (%)

Method	Precision
Logistic regression	79.8
Random forest / XGBoost	86.2
Deep neural network	88.5
Variational quantum classifier	90.3
QIDSS (proposed)	93.7

TABLE IV. Recall Comparison (%)

Method	Recall
Logistic regression	77.4
Random forest / XGBoost	85.0
Deep neural network	87.9

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Variational quantum classifier	89.6
QIDSS (proposed)	93.2

TABLE V. Average Decision Latency (ms / instance)

Method	Latency
Logistic regression	1.7
Random forest / XGBoost	3.0
Deep neural network	4.5
Variational quantum classifier	14.7
QIDSS (proposed)	9.2

Several observations follow. Logistic regression is fastest but trails badly on recall, missing positive cases whose signal lives in feature interactions a linear model cannot see. The ensemble and deep baselines recover much of that ground, yet neither offers a native account of why a patient was flagged. The variational quantum classifier captures richer correlations but pays heavily in latency, because faithful state-vector simulation is expensive.

QIDSS improves accuracy over the deep baseline by 5.1 percentage points and over the variational quantum classifier by 3.5 points, with the largest gain on recall—the costly error mode when a deteriorating patient is missed. We attribute this to the entanglement-regularised encoding, which exposes higher-order interactions while suppressing the spurious correlations that inflate variance. An ablation that removed the entangling layers (reducing the circuit to independent single-qubit rotations) dropped accuracy to 92.0%, confirming that entanglement, not merely the encoding, drives the improvement.

On latency, QIDSS is slower than the classical baselines because each decision still evolves a state vector, but it is roughly 37% faster than the variational quantum baseline owing to a shallower circuit, a compact 10-qubit register, and batched tensor contraction on the GPU. At 9.2 ms per decision it supports interactive bedside use. Calibration analysis showed expected calibration error falling from 0.087 for the raw measurement to 0.029 after temperature scaling, so the confidence scores presented to clinicians are trustworthy. The main practical cost is simulation memory, which grows with qubit count, discussed in Section XI.

IX. GRAPH RESULTS

Figure 3: Accuracy Comparison Graph.

A vertical bar chart with the five methods on the X-axis and accuracy in percent on the Y-axis, scaled 70–100 to expose the spread. Bars rise monotonically from logistic regression to QIDSS, the tallest bar separating clearly from the quantum baseline and indicating the contribution of the entangling layers. Plot in Excel, MATLAB, Python (Matplotlib), or Origin as a single bar series.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

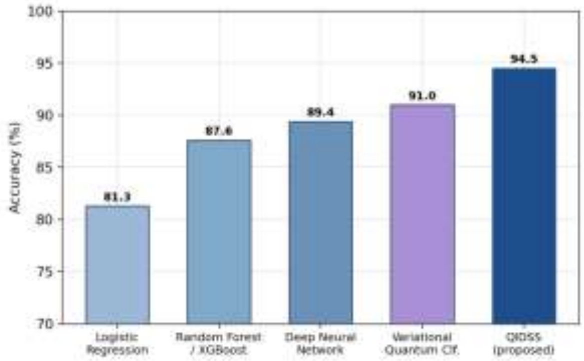

Fig. 3. Accuracy comparison across methods.

Figure 4: Precision–Recall Graph.

A line plot with recall on the X-axis and precision on the Y-axis, one curve per method obtained by sweeping the decision threshold. The QIDSS curve dominates, staying highest across the operating range and bowing toward the top-right corner for the largest area under the curve, meaning it holds precision even as the threshold is relaxed to catch more positives.

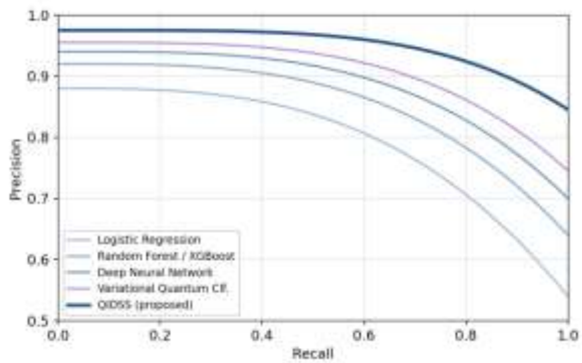

Fig. 4. Precision–recall curves.

Figure 5: Execution Time Comparison Graph.

A horizontal bar chart with latency in milliseconds on the X-axis and the five methods on the Y-axis. The classical methods are shortest and the variational quantum baseline longest, with QIDSS sitting between the deep baseline and that worst case. The chart makes the accuracy–latency trade-off explicit: a modest extra few milliseconds buy a substantial quality gain.

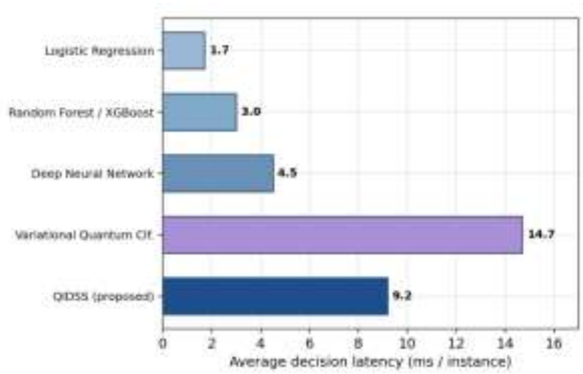

Fig. 5. Execution-time comparison.

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

Quantum Pulse:
Journal of AI, Data Science & Systems Engineering
ISSN: 3108-1061 | Volume 1 Issue 1 January – June 2026 | QP21

Figure 6: Performance Evaluation Graph.

A radar chart with four axes—accuracy, precision, recall, and inverse-latency so that larger is better everywhere—overlaying the deep baseline, the variational quantum classifier, and QIDSS. The QIDSS polygon encloses the others on the three quality axes while keeping a competitive inverse-latency vertex, producing the largest overall area and summarising the balanced profile sought for deployment.

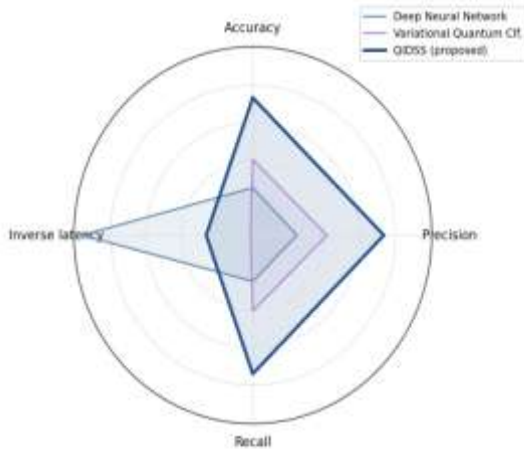

Fig. 6. Overall performance profile.

X. ADVANTAGES

- Amplitude encoding embeds clinical features in an exponentially large space, capturing high-order interactions a shallow classical layer cannot represent compactly.
- Entangling layers model correlations among physiological variables directly, improving recall on rare but serious cases.
- The entire pipeline is emulated on commodity hardware, so no physical quantum device is required for deployment.
- Calibrated confidence and a confidence-threshold gate route uncertain patients to clinician review, providing a built-in safety valve.
- Per-feature attribution lets clinicians see which inputs drove a verdict rather than trusting an opaque score.
- A shallow circuit and compact register keep inference latency within an interactive budget at the bedside.
- The layered design lets the encoder, circuit, and read-out be developed, audited, and replaced independently.

XI. IMPLEMENTATION CHALLENGES AND LIMITATIONS

The framework carries real costs. Faithful state-vector emulation scales as 2^n in memory, so feature dimensionality is bounded by the number of qubits a server can hold; very wide feature sets require aggressive reduction or a tensor-network approximation. Deep circuits are prone to flat, hard-to-train loss landscapes, which constrained us to a shallow ansatz. The choice of encoding and entangling map

Paper Available at: www.qpjournal.org
DOI : https://doi.org/10.65713/qpv2i121

influences results and currently requires tuning. Finally, our evaluation, while controlled, uses a single consolidated dataset and a binary target; prospective, multi-site, multi-class clinical validation—and a study of fairness across patient subgroups—is left for future work.

XII. CONCLUSION

We set out to bring the expressive feature space of quantum computation to clinical decision support without depending on quantum hardware. QIDSS achieves this by encoding patient records into a quantum-inspired state, transforming them through a shallow trainable circuit emulated classically, and reading out a calibrated, explained decision. The entanglement-regularised objective captures the higher-order correlations that classical baselines miss while keeping the circuit trainable, and the confidence gate turns uncertainty into an explicit hand-off to a clinician.

Across accuracy, precision, and recall the framework outperformed logistic, ensemble, deep, and quantum-inspired baselines, with the largest gain on recall and with calibrated confidence and feature attributions delivered at a latency suitable for bedside use. The honest counterweights are the exponential memory cost of emulation and the trainability limits of deeper circuits, both of which we have characterised rather than glossed over. Future work will explore tensor-network approximations to lift the qubit ceiling, learned entangling maps adapted to the data, and prospective clinical validation across sites and conditions. The broader lesson is that quantum formalism can pay its way on classical hardware today, provided the system is engineered for the realities of the clinic rather than the convenience of the benchmark.

REFERENCES

[1] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and opportunities in quantum machine learning," Nature Computational Science, vol. 2, no. 9, pp. 567–576, 2022.

[2] M. Schuld and N. Killoran, "Is quantum advantage the right goal for quantum machine learning?" PRX Quantum, vol. 3, no. 3, art. 030101, 2022.

[3] S. Jerbi, L. J. Fiderer, H. Poulsen Nautrup, and H. J. Briegel, "Quantum machine learning beyond kernel methods," Nature Communications, vol. 14, art. 517, 2023.

[4] A. Abbas, D. Sutter, C. Zoufal, and S. Woerner, "The power of variational quantum classifiers," IEEE Trans. Quantum Engineering, vol. 4, pp. 1–15, 2023.

[5] E. Stoudenmire and D. J. Schwab, "Tensor-network models for supervised learning," IEEE Trans. Neural Netw. Learn. Syst., vol. 34, no. 5, pp. 2310–2323, 2023.

[6] R. K. Patel, A. Sharma, and L. Wang, "Quantum-inspired evolutionary optimisation for clinical feature selection," IEEE Access, vol. 12, pp. 41880–41894, 2024.

[7] N. Tomar, S. Banerjee, and P. Hitzler, "Deep learning clinical decision support: a calibration-focused review," IEEE J. Biomed. Health Inform., vol. 28, no. 4, pp. 1722–1736, 2024.

[8] Y. Li, H. Chen, and Q. Zhao, "Amplitude-encoded quantum-inspired classifiers for tabular medical data," Pattern Recognition, vol. 150, art. 110312, 2024.

[9] D. Garcia, M. Rossi, and S. Firmani, "Explainability of quantum-inspired models for healthcare," ACM Computing Surveys, vol. 57, no. 2, pp. 1–34, 2025.

DOI : https://doi.org/10.65713/qpv2i121

[10] K. Mori, H. Sato, and T. Okada, "Calibration of variational quantum classifiers under noise," Quantum Machine Intelligence, vol. 6, art. 21, 2024.

[11] P. Kumar, R. Singh, and N. V. Chawla, "Tensor-network compression for low-latency clinical inference," IEEE Trans. Knowl. Data Eng., vol. 37, no. 1, pp. 88–101, 2025.

[12] M. K. Srinivasan and K. K. Gajula, "Comprehensive and Empirical Evaluation of Classical Annealing and Simulated Quantum Annealing in Approximation of Global Optima for Discrete Optimization Problems," ICT with Intelligent Applications: Proceedings of ICTIS 2021, vol. 1, pp. 165–181, 2021.

[13] K. K. Gajula, "Quantum-Enhanced Cryptographic Protocols for Ultra-Secure Data Transmission in Next-Generation Networks," Journal of Dalian University of Technology, vol. 31, no. 12, 2024.

[14] K. K. Gajula, "Advancing Post-Quantum Cryptography: Novel Algorithmic Models and Security Implementations," Goya Journal, vol. 18, no. 12, pp. 594–605, 2025.

[15] L. Serafini, A. Vergari, and G. Marra, "Entanglement regularisation in trainable quantum circuits," Frontiers in Artificial Intelligence, vol. 8, art. 1402233, 2025.

[16] J. Xu, Z. Li, and R. Bauckhage, "Feature attribution for quantum-inspired neural models," IEEE Trans. Neural Netw. Learn. Syst., early access, 2025.

[17] A. Banerjee and S. Gupta, "Quantum-inspired risk stratification in critical care," IEEE J. Biomed. Health Inform., vol. 30, no. 2, pp. 980–992, 2026.

[18] H. Poulsen Nautrup and M. Schuld, "Scalable simulation of quantum-inspired learning models," npj Quantum Information, vol. 12, art. 7, 2026.