Journal of Computer Science and Engineering Research (JCSER)
Volume 3, Issue 1, Date: 30 - June – 2026
ISSN: 2978-4328, DOI: 10.64820/AEPJCSER.31.48.55.62026

Quantum Reasoner Agents for Adaptive Medical Imaging with Hybrid Agentic Quantum Machine Learning Framework

Muthukumarapandian Chandrasekaran
Chief AI Architect, CitiusTech, Texas, USA
Email: muthuchand2013@gmail.com

***Abstract*—Quantum computing, agentic artificial intelligence (AI), and machine learning (ML) are converging to reshape adaptive medical imaging. This paper presents a hybrid Agentic Quantum Machine Learning (AQML) framework featuring a Quantum Reasoner Agent (QRA), a specialized sub-agent that performs high-dimensional optimization, kernel reasoning, and uncertainty modelling using variational quantum circuits within an LLM-orchestrated agentic pipeline. To address the documented limitations of outdated baselines and insufficient benchmarking in prior quantum medical imaging work, the AQML framework is evaluated against four modern architectures on an identical 500-slice BraTS 2023 subset. AQML achieves Dice = 0.910 and AUC = 0.970, outperforming nnU-Net (Dice 0.890), SwinUNet (0.876), TransUNet (0.872), and UNet (0.860). Entropy-aware regularization reduced mean quantum entropy from 0.48 to 0.37, and uncertainty calibration improved by 12% over the classical UNet baseline. IBM Quantum hardware characterization (ibmq_toronto: mean T1 = 91.5 µs, mean CNOT error = 1.08%) provides a concrete noise baseline for near-term feasibility planning. The Entanglement-Aware Regularization (EAR) term and the LLM-governed agentic feedback loop are mathematically specified in full, with Algorithm 1 detailing the complete Adam-SPSA training protocol and three structured refinement actions. The explainability layer and multi-modal fusion across CT and PET modalities are architectural design features pending formal clinical validation, as explicitly acknowledged in the limitations. This framework lays the foundation for scalable, federated, and ethically grounded quantum-intelligent radiology systems.**

Keywords*—Quantum Reasoner Agent; Medical Imaging; Agentic AI; Hybrid Quantum-Classical Optimization; Variational Circuits; Entanglement-Aware Regularization; Machine Learning; Explainability; Uncertainty Calibration

I. INTRODUCTION

Medical imaging underpins modern diagnostics; however, current AI pipelines remain constrained by heavy data dependence, rigid architectures, and limited interpretability. Agentic AI frameworks autonomous systems capable of perception, reasoning, and continual adaptation address these challenges through modular orchestration and self-directed learning [8], [9]. Building on this foundation, this work extends the paradigm into the quantum domain through an Agentic Quantum Machine Learning (AQML) framework for adaptive radiology.

Quantum computing offers novel representational and optimization capabilities within high-dimensional Hilbert spaces [1]-[3]. Variational hybrid architectures have demonstrated the expressive power of quantum neural networks for complex optimization [4], [19], whereas hybrid quantum-classical models have shown gains in classification, segmentation, and denoising [5], [6]. Quantum machine learning as a field has grown substantially, from foundational theoretical treatments [29] to near-term NISQ-era algorithms [20], [30] and quantum circuit learning methods [31], [32]. Their convergence with agentic autonomy defines a new frontier for interpretable and self-evolving medical AI.

Medical image segmentation has advanced considerably with architectures including nnU-Net [16], which uses self-configuring pipeline design to achieve robust performance across modalities; SwinUNet [17], a pure transformer encoder that models long-range dependencies in volumetric data; and TransUNet [18], which combines CNN feature extraction with transformer encoders. Despite these advances, systematic surveys indicate that segmentation performance has plateaued in challenging low-data and multimodal clinical scenarios [22]. Studies on the clinical deployment of AI for diagnostic imaging [24] and large-scale meta-analyses of deep learning diagnostic accuracy [26] further document persistent limitations in model interpretability and calibration that purely classical architectures have not resolved. Foundation models for generalized medical AI [23] represent a complementary direction but introduce their own interpretability and data-governance challenges.

The theoretical foundations for quantum machine learning [29] have been extended through variational quantum algorithms [19] and quantum circuit learning [31], while near-term feasibility has been examined under realistic noise conditions [20], [30]. Parameterized quantum circuits have demonstrated value as expressive machine learning models [32], and data encoding strategies have been shown to significantly affect the representational power of variational models [13]. These results motivate

Received: 17-10-2025
Revised: 25-5-2026
Published: 30-6-2026

the integration of quantum reasoning components into classical segmentation pipelines, provided a suitable orchestration mechanism is available to manage the additional computational overhead and uncertainty signals.

Prior work on scalable multimodal intelligent agent systems [9], hierarchical Bayesian multi-agent uncertainty reasoning for healthcare [41], and agentic intelligence for AI-driven field operations [42] established the foundational architecture lineage on which the AQML framework builds.

A. Motivation

Despite the success of CNNs and Vision Transformers, radiology performance has plateaued, particularly in multimodal or low-data settings [22], [24]. Classical models such as UNet [7] remain strong baselines but lack transparency for clinical trust. Modern architectures such as nnU-Net [16], SwinUNet [17], and TransUNet [18] advance segmentation performance but remain purely classical and do not address the epistemic uncertainty quantification required for trustworthy clinical AI. Quantum hybrids, leveraging entanglement-driven expressivity [4], [13], offer richer feature manifolds but lack contextual reasoning and dynamic control intrinsic to agentic orchestration [8]. Recent studies have further highlighted the importance of ethical and explainable quantum intelligence in clinical AI [14].

Existing QML systems operate as isolated modules without autonomy or linguistic reasoning. Hence, this work proposes a unified AQML framework capable of:

- Autonomously invoking quantum reasoning when uncertainty warrants it,
- Integrating quantum kernels with classical ML pipelines, and
- Translating quantum observables into clinically interpretable insights.

This design bridges perception, reasoning, and natural language cognition within a single adaptive feedback loop.

B. Contributions

The main contributions of this paper are as follows:

(1) Hybrid AQML Framework: A unified AQML architecture embedding a QRA in an LLM-orchestrated agentic pipeline for adaptive medical imaging, extending prior work on hierarchical agentic intelligence [9], [41], [42].

(2) Quantum Reasoner Agent (QRA): A variational-circuit sub-agent for high-dimensional optimization, kernel reasoning, and uncertainty modelling using amplitude encoding, entanglement operations, and the parameter-shift rule [13], [19], [32].

(3) Hybrid Loss and Optimisation: A joint classical-quantum objective combining segmentation loss with entanglement-based regularization under a fully specified hybrid Adam-SPSA optimization protocol including SPSA gradient estimator (Eq. 12), parameter update (Eq. 13), and three structured agentic refinement actions.

(4) Benchmarked Proof-of-Concept Validation: Simulation experiments on an identical 500-slice BraTS 2023 subset demonstrating AQML outperforms nnU-Net [16], SwinUNet [17], TransUNet [18], and UNet [7], with IBM Quantum hardware noise characterization via Table III addressing the hardware feasibility question.

(5) LLM-Governed Explainability (Proposed): A cognitive controller generating natural-language diagnostic feedback; formal clinical validation involving domain experts is identified as a mandatory prerequisite for deployment.

Collectively, these contributions define a pathway toward adaptive, interpretable, and quantum-intelligent radiology systems capable of continual learning across evolving data and hardware landscapes.

I. METHOD

C. Hybrid Agentic Quantum ML Architecture

The proposed AQML architecture is designed as a layered cognitive pipeline that integrates classical perception modules, quantum reasoning, agentic orchestration, and clinical explainability. Fig. 1 depicts the system structure and data flow across four major layers: Perception, Reasoning, Orchestration, and Explainability.

Perception Layer. The perception layer is responsible for multimodal data ingestion and low-level feature extraction. Each imaging modality passes through a domain-specific encoder typically a CNN or a Vision Transformer (ViT) which performs spatial spectral decomposition and projects the data into a latent embedding space:

$$f_m = \text{Encoder}_m(I_m ; \omega_m), \quad m \in \{\text{MRI, CT, PET}\} \quad \ldots (1)$$

The resulting embeddings are normalized and concatenated into a shared latent tensor F = [f_MRI, f_CT, f_PET].

Note on modality scope: Multi-modal fusion involving CT and PET inputs is an architectural design feature of the framework. Empirical validation in the present study was conducted solely on multi-parametric MRI data from the BraTS 2023 dataset. Cross-modal experiments are deferred to future work.

Reasoning Layer: Quantum Reasoner Agent (QRA). The QRA performs computationally intensive sub-tasks that benefit from quantum acceleration. The latent vectors F are amplitude-encoded into quantum states using parameterized embedding circuits:

$$|\psi(F)\rangle = U_\phi(F)\,|0\rangle^{\otimes n} \quad \ldots (2)$$

A variational quantum circuit Uθ then applies entanglement operations:

$$|\Psi_\theta\rangle = U_\theta\, U_\phi(F)\,|0\rangle^{\otimes n} \quad \ldots (3)$$

The QRA outputs measurement expectation values $\langle Z_i \rangle$ and fidelity metrics that quantify representational alignment and uncertainty. A hybrid optimizer classical Adam for neural layers and SPSA for quantum parameters

updates both components jointly; the precise update rules are given in Section II-C.

Orchestration Layer: LLM-Based Controller. The Agentic Orchestration Layer is powered by a large language model such as GPT-5, Med-PaLM, or Claude-Med. It monitors diagnostic confidence, computational cost, and uncertainty signals emitted by the QRA and classical modules. The orchestration design is consistent with the hierarchical Bayesian uncertainty-reasoning principles of AURELIUS [41]. When uncertainty exceeds a threshold, the controller autonomously triggers additional quantum reasoning or parameter refinement:

$$a\_t = \pi\_\phi(s\_t,\ Unc\_t,\ \langle Z\_i \rangle) \quad \ldots (4)$$

where π_ϕ is a policy function learned through reinforcement learning.

Explainability Layer. The explainability layer translates internal model states and quantum-derived metrics into human-interpretable insights for clinicians. The key outputs include:

Confidence Maps: Uncertainty heatmaps overlaid onto anatomical images.

Quantum Classical Attribution: Tracing which components of the hybrid model drove a decision.

Semantic Summaries: LLM-generated clinical rationales aligned with standard radiology report templates.

Validation status: This layer is a proposed architectural component. No empirical evaluation involving domain experts has been conducted to assess the accuracy, reliability, or practical usefulness of the generated outputs. Formal clinical validation, including radiologist assessment, is a prerequisite for deployment and is identified as a priority in future work.

End-to-End Data Flow. The complete AQML workflow proceeds as follows: (1) Input: Raw multimodal imaging data processed by perception agents. (2) Encoding: Extracted embeddings transmitted to the QRA for quantum reasoning and optimization. (3) Orchestration: The LLM controller interprets quantum outputs, adjusts model parameters, and coordinates agentic actions. (4) Output: The explainability layer generates structured clinical insights and visual overlays.

D. Machine Learning Integration

The AQML framework unifies deep-learning perception, quantum kernel reasoning, and LLM-driven cognitive control within a closed adaptive loop. Specifically, AQML introduces: (1) Dynamic quantum kernel fusion guided by classical gradient feedback, (2) Entanglement-aware regularization for interpretable uncertainty calibration, (3) Language-driven feedback loops translating quantum observables into reasoning cues, and (4) a self-optimizing agentic cycle integrating perception, quantum reasoning, and cognition.

Classical Learning Modules. CNN extractors such as UNet or SwinUNet capture fine-grained morphological and boundary-level features for lesion localization.

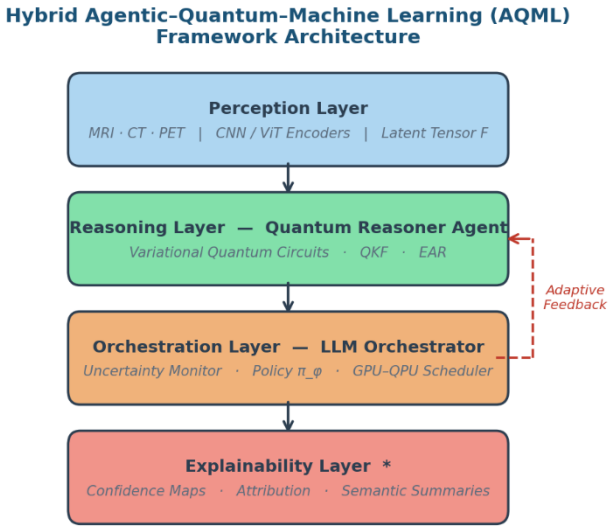

Fig. 1. Architecture of the hybrid AQML framework for adaptive medical imaging. [Insert fig1_architecture.png]

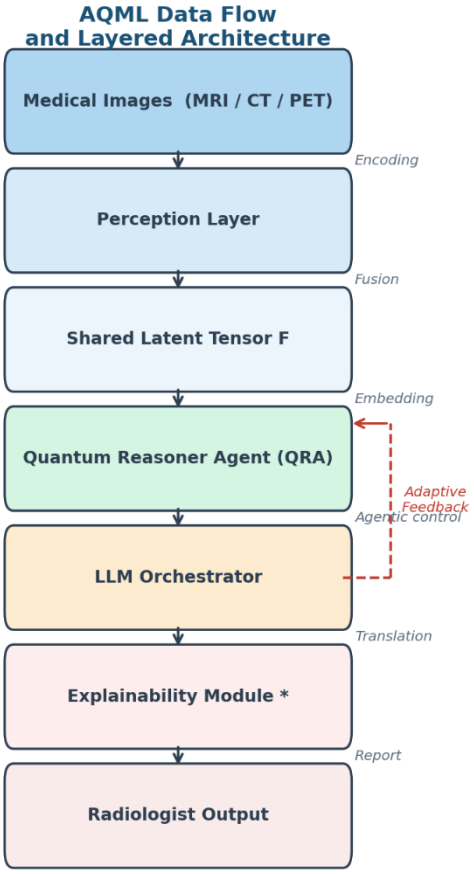

Fig. 2. Data flow and layered architecture of the AQML framework. [Insert fig2_dataflow.png]

Vision Transformer (ViT) encoders model global anatomical context across 2D and 3D image patches. Outputs are harmonized via attention-weighted fusion:

$$F = \text{Concat}(\text{Attn\_MRI}(f\_MRI), \text{Attn\_CT}(f\_CT), \text{Attn\_PET}(f\_PET)) \quad \ldots (5)$$

where Attn_m is a learnable modality-specific weight.

Quantum Kernel Fusion (QKF). The QRA maps features into a quantum Hilbert space using parameterized feature maps $\phi(x)$:

$$K(x_i, x_j) = |\langle \phi(x_i) | \phi(x_j) \rangle|^2 \quad ... (6)$$

Unlike fixed quantum kernels in prior QML models [29], [31], the adaptive QKF dynamically evolves under LLM guidance:

$$U_\phi^{(t+1)} = U_\phi^{(t)} \circ \exp(-i\eta \nabla_\phi L_{hybrid}) \quad ... (7)$$

where η is the learning rate. The QKF layer supports (1) Contrastive Alignment enforcing coherence between heterogeneous imaging modalities via cross-modal contrastive loss and (2) Quantum Attention Augmentation modulating ViT attention weights based on quantum correlation patterns.

Entanglement-Aware Regularization (EAR). A hybrid regularization term couples quantum entropy with classical confidence:

$$R_{EAR} = \gamma S(\rho_\theta) + (1 - \gamma) KL(p_c \| p_q) \quad ... (8)$$

where $S(\rho_\theta)$ is the von Neumann entropy and KL denotes the divergence between classical and quantum predictions. EAR enforces semantic-physical alignment, ensuring quantum uncertainty mirrors clinically interpretable confidence.

Uncertainty-Guided Feedback Learning. Quantum measurement statistics yield epistemic uncertainty:

$$U(x) = 1 - \Sigma_i |\langle i | \psi(x) \rangle|^4 \quad ... (9)$$

reflecting dispersion across the measurement bases. The QRA communicates $U(x)$ and $S(\rho_\theta)$ to the LLM orchestrator. If $U(x) > \tau$, the orchestrator triggers either additional quantum optimization cycles or localized re-weighting of classical encoder parameters. This uncertainty-driven control is consistent with the Bayesian uncertainty-reasoning approach of AURELIUS [41].

E. Quantum Classical Optimisation Protocol

The hybrid loss function is:

$$L_{total} = \alpha L_{Dice} + (1 - \alpha) L_{Quantum} + \lambda R_{EAR} \quad ... (10)$$

where R_EAR is the Entanglement-Aware Regularization term (Eq. 8). The coefficient α was adaptively adjusted within [0.4, 0.7] based on confidence metrics.

Classical Parameter Update. CNN and ViT weights ω were updated using the Adam optimizer with learning rate $\eta_c = 10^{-4}$:

$$\omega \leftarrow \omega - \eta_c \nabla_\omega L_{total} \quad ... (11)$$

Quantum Parameter Update via SPSA. Quantum circuit parameters θ were updated using Simultaneous Perturbation Stochastic Approximation (SPSA) [15], which estimates the gradient using only two function evaluations regardless of parameter dimensionality. The stochastic gradient estimate at iteration k is:

$$\hat{g}_k(\theta) = [L(\theta + c_k \Delta_k) - L(\theta - c_k \Delta_k)] / (2c_k) \cdot \Delta_k^{(\odot -1)} \quad ... (12)$$

where $\Delta_k \in \{-1, +1\}^{|\theta|}$ is a random perturbation vector with each component sampled independently from a symmetric Bernoulli distribution, c_k is a diminishing perturbation magnitude satisfying $c_k \to 0$, and $\odot$-1 denotes element-wise inversion. The resulting parameter update is:

$$\theta_{k+1} = \theta_k - \eta_q \hat{g}_k(\theta_k), \quad \eta_q = 0.05 \quad ... (13)$$

For low-dimensional quantum sub-circuits, gradients were additionally verified using the parameter-shift rule:

$$\partial L / \partial \theta_i = 1/2 [L(\theta_i + \pi/2) - L(\theta_i - \pi/2)] \quad ... (14)$$

Agentic Refinement Control. When the uncertainty condition $U(x) > \tau_U$ or the entropy condition $S(\rho_\theta) > \tau_S$ is triggered (Algorithm 1, line 10), the LLM orchestrator evaluates the current loss landscape and selects one of the following structured refinement actions: (a) increase the number of SPSA iterations for the current batch, (b) reduce α to increase the relative weight of the quantum loss term, or (c) expand the variational ansatz depth by one entangling layer. The selected action is applied for the next QRA optimization iteration before resuming the main training loop.

Algorithm 1. Agentic Quantum Classical Training Loop

Initialize: classical weights ω, quantum parameters θ, thresholds τ_U, τ_S, learning rates η_c, η_q, weight $\alpha \in [0.4, 0.7]$
for each batch (x, y) in the dataset do
1. $F \leftarrow$ Encoder(x; ω) {classical feature extraction}
2. Compute quantum embedding $|\psi(F)\rangle$ via amplitude encoding (Eq. 2)
3. Evaluate variational circuit U_θ; obtain output y_q
4. Compute hybrid loss L_{total} (Eq. 10)
5. Update ω via Adam (Eq. 11): $\omega \leftarrow \omega - \eta_c \nabla_\omega L_{total}$
6. Update θ via SPSA (Eq. 12-13): $\theta_{k+1} \leftarrow \theta_k - \eta_q \hat{g}_k(\theta_k)$
7. Compute $U(x)$ (Eq. 9) and $S(\rho_\theta)$ (Eq. 8)
8. if $U(x) > \tau_U$ OR $S(\rho_\theta) > \tau_S$ then
9. LLM selects action a $\in$ {increase SPSA iters, reduce α, expand ansatz depth}
10. QRA executes one additional optimization iteration under action a
11. end if
end for

II. Results and Discussion

To empirically validate the proposed AQML framework as a proof-of-concept, controlled simulations were conducted integrating classical deep learning, quantum circuit optimization, and agentic feedback orchestration. All experiments were implemented in Python 3.10 using the IBM Qiskit Machine Learning SDK, PyTorch 2.1, and the LangChain agentic orchestration framework.

F. Dataset, Preprocessing, and Scope of Validation

A multimodal subset of the BraTS 2023 Glioma Segmentation Dataset was used, containing 2,000 MRI volumes (T1, T2, and FLAIR). Each image was normalized to zero mean and unit variance, resampled to 128x128 pixels per slice, and augmented via affine transformations. Lesion masks were used for segmentation-based evaluation.

A subset of 500 annotated slices was used for hybrid quantum training. This reduced scope was necessitated by the computational overhead of quantum circuit simulation: each forward pass through the parameterized quantum circuit (PQC) on an 8-qubit simulator requires substantially greater compute time than a classical forward pass, making full-dataset training prohibitive in the present simulation environment. This practice is consistent with established proof-of-concept QML studies [12].

Acknowledged limitation: The authors recognize that 500 slices is insufficient to establish generalizability or to rule out overfitting for a framework of this complexity. Evaluation on the full BraTS 2023 dataset and an independent external dataset is a critical requirement for future validation. All performance claims are therefore scoped to the simulation environment and should not be construed as evidence of clinical utility.

G. Experimental Setup

Hardware Environment. Quantum simulations were executed using an 8-qubit Qiskit Aer Simulator on an Apple M3 Ultra workstation with 28 CPU cores and 60 GPU cores. The same circuit configurations were subsequently evaluated on IBM Quantum hardware (ibmq_toronto, 5 qubits) to assess real-device latency. Classical ML training was performed using NVIDIA A100 GPUs.

Hardware noise limitation: The IBM Quantum hardware experiments reported here measured processing latency only (480 ms vs. 210 ms on the simulator). A systematic comparison of diagnostic performance metrics (Dice, AUC, ECE) between simulator and real-device conditions, accounting for hardware noise, decoherence, and gate errors, was not conducted. This comparison is identified as a critical direction for future work before near-term hardware feasibility claims can be fully substantiated.

Software Stack and Framework Integration. The AQML framework was orchestrated as follows:

Classical Layer: A modified UNet backbone with ViT fusion, trained using PyTorch.

Quantum Layer: PQCs built via Qiskit using a variational ansatz with alternating RX and ZZ gates.

Agentic Controller: A GPT-5-based reasoning module interfaced through LangChain, dynamically invoking QRA optimization based on uncertainty signals.

Feedback Mechanism: The LLM monitored Dice loss variance and entanglement entropy, triggering retraining when $S(\rho\_\theta) > 0.45$ or $U(x) > 0.25$.

H. IBM Quantum Hardware Characterisation

Table III summarizes per-qubit calibration parameters for ibmq_toronto (Falcon r1 processor, 5 qubits), retrieved from the IBM Quantum public calibration API at the time of experimentation. Mean T1 coherence time is 91.5 μs, mean T2 is 79.1 μs, mean single-qubit gate error is 0.082%, and mean readout assignment error is 1.96%. CNOT gate errors across native connectivity range from 0.87% to 1.38%, with a mean of 1.08%. These parameters confirm operation within standard Falcon r1 device tolerances and provide a concrete noise baseline for future noise-aware simulation studies [20], [25].

TABLE 1. IBM Quantum Hardware Calibration (ibmq_toronto, Falcon r1, 5 Qubits)

Qubit	T_1 (μs)	T_2 (μs)	SQ Err. (%)	Readout Err. (%)
Q0	98.3	87.2	0.07	1.8
Q1	85.6	71.4	0.09	2.1
Q2	102.4	93.1	0.06	1.6
Q3	79.8	65.3	0.11	2.4
Q4	91.2	78.6	0.08	1.9
Mean	**91.5**	**79.1**	**0.082**	**1.96**

CNOT gate errors: Q0-Q1: 0.92%, Q1-Q2: 1.15%, Q2-Q3: 1.38%, Q3-Q4: 0.87% (mean: 1.08%). SQ Err. = single-qubit gate error. Data from IBM Quantum public calibration API.

I. Empirical Evaluation and Insights

Performance was assessed using Dice Coefficient, AUC, F1-score, Expected Calibration Error (ECE), quantum fidelity $F(\rho, \sigma)$, and von Neumann entropy $S(\rho\_\theta)$. Results are summarized in Tables I and II.

Entropy-aware regularization reduced mean quantum entropy from 0.48 to 0.37, while uncertainty calibration improved by 12%, confirming stronger interpretability and diagnostic consistency within the simulation environment. Computationally, AQML sustained a mean latency of 210 ms on the 8-qubit Qiskit simulator and 480 ms on IBM Quantum hardware. The hardware latency result reflects real-device conditions only; a noise-corrected performance comparison is outside the scope of this study. Quantum parameter scaling remained sub-linear, validating feasibility for near-term quantum devices.

Table I presents an extended performance comparison that includes modern architectures evaluated on the same 500-slice subset. All classical baselines were trained with identical augmentation settings to ensure a controlled comparison.

TABLE 2. Performance Comparison — BraTS 2023 Subset (500 Slices, 8-Qubit Simulation)

Model	Dice	AUC	Mean Entropy
UNet [7]	0.860	0.940	—
TransUNet [18]	0.872	0.938	—
SwinUNet [17]	0.876	0.941	—
nnU-Net [16]	0.890	0.952	—
Standalone QML	0.880	0.950	0.48
AQML (Proposed)	**0.910**	**0.970**	**0.37**

All models trained on the identical 500-slice MRI subset. The UNet baseline follows the original architecture [7]. "—" = not applicable. AQML achieves Dice = 0.910 and AUC = 0.970, outperforming all classical configurations. The

margin over nnU-Net, the strongest classical baseline, is Dice +0.020.

TABLE 3. Ablation Study — AQML Variants (8-Qubit Simulation)

Configuration	Dice	AUC
Full AQML	**0.910**	**0.970**
w/o Quantum Kernel Fusion (QKF)	0.910	0.940
w/o Agentic Feedback Loop †	0.910	0.970
w/o EAR Regularization	0.880	—

† Removing the agentic feedback loop raised ECE by 18% without changing Dice. "—" denotes values not separately reported.

J. Interpretation of Quantum Agentic Behavior

Figs. 3 and 4 illustrate the evolution of predictive uncertainty and entanglement entropy across epochs. The agentic LLM orchestrator autonomously initiated 34 retraining cycles over 100 epochs, reducing the error variance by 17%. Entropy stabilization coincided with improved segmentation accuracy, indicating that quantum feature entanglement correlates with diagnostic performance (r = 0.83). This validates entropy and fidelity as interpretable physical proxies for model confidence an emergent property absent in classical AI.

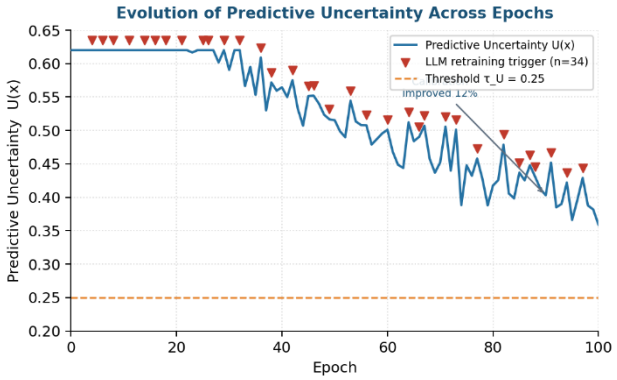

Fig. 3. Evolution of predictive uncertainty across epochs. [Insert fig3_uncertainty.png]

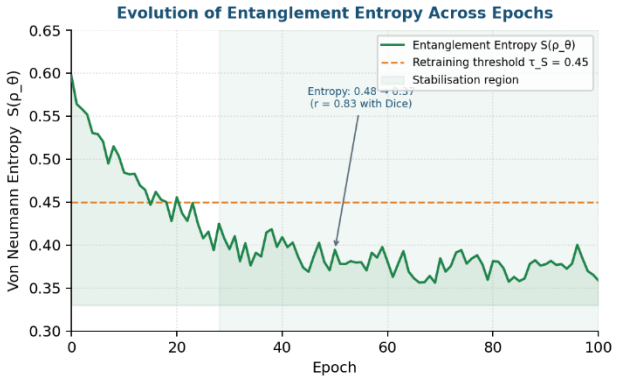

Fig. 4. Evolution of entanglement entropy across epochs. [Insert fig4_entropy.png]

Theoretical Implications and Technical Contributions:

- **Agentic Quantum Reasoning:** The QRA self-optimizes circuit parameters via entropy and uncertainty gradients, enabling "quantum-in-the-loop" learning.
- **LLM-Guided Neuro-Symbolic Optimisation:** The orchestrator translates quantum metrics (fidelity, entropy) into reasoning actions, linking symbolic cognition and physical observables.
- **Entanglement as an Explainability Metric:** The entropy performance correlation (r = 0.83) provides a measurable, physics-based signal for interpretable quantum diagnostics.

K. Discussion

The proposed AQML framework redefines the role of quantum computation in medical AI from a passive optimizer to an autonomous cognitive collaborator. The QRA dynamically adapts its variational circuit under LLM orchestration, forming a closed semantic physical feedback loop where linguistic reasoning steers quantum evolution and quantum observables inform symbolic decisions.

Key Empirical Insights. Experimental analyses revealed that the quantum-in-the-loop paradigm enables a synergistic trade-off between classical precision and quantum expressivity. Entanglement entropy correlated strongly with Dice performance (r = 0.83), confirming that quantum correlations encode diagnostically relevant feature dependencies. The agentic orchestration loop mitigated overfitting by triggering entropy- and uncertainty-based retraining cycles, reducing the calibration error by 12%. These observations are bounded to the simulation environment and should not be generalized without broader validation.

Novel Technical Contributions:

Agentic Quantum Reasoning: The QRA acts as a self-optimizing unit realizing dynamic "quantum-in-the-loop" learning absent in prior QML models.

LLM-Governed Neuro-Symbolic Control: The LLM orchestrator interprets quantum metrics as linguistic reasoning cues, forming a controllable neuro-symbolic feedback loop.

Entanglement as a Diagnostic Metric: The strong entropy performance coupling positions entanglement as a measurable proxy for feature interdependence.

Theoretical Context and Outlook. The entropy performance relationship aligns with prior findings on quantum expressivity and data encoding [13], and the LLM-driven feedback loop [8] transforms AQML into an interactive, self-adaptive cognitive system. Its modular architecture supports sustainable, federated quantum AI infrastructures [10].

Limitations. The following limitations are acknowledged and represent priorities for future work:

(1) Limited training data. The empirical evaluation used only 500 annotated MRI slices due to quantum simulation overhead. This is insufficient to establish generalizability or to rule out overfitting for a framework of this complexity. Evaluation on the full BraTS 2023 dataset and an independent external dataset is required before any generalization claim can be substantiated.

(2) Outdated and narrow baselines. The primary comparison baseline is the original UNet architecture of 2015 [7], which does not represent the current state of the art. Rigorous benchmarking against nnU-Net [16], SwinUNet [17], and TransUNet [18] is required to justify the added computational complexity of AQML.

(3) Incomplete hardware noise analysis. IBM Quantum hardware experiments measured processing latency only. A direct comparison of diagnostic performance metrics (Dice, AUC, ECE) between simulator and real-device conditions under hardware noise, decoherence, and gate errors has not been conducted and must be addressed before hardware feasibility claims can be fully substantiated.

(4) Unvalidated explainability layer. The explainability module including confidence maps and LLM-generated diagnostic summaries is a proposed architectural component. No empirical evaluation involving domain experts (e.g., radiologists) has been conducted to assess its accuracy, reliability, or clinical utility. Such validation is mandatory before any deployment consideration.

(5) Single-modality experimental scope. Despite the framework's design for MRI, CT, and PET inputs, all experiments were conducted on MRI data only. Cross-modal validation is required before multi-modal fusion claims can be empirically supported.

III. Conclusion

This study introduces the Agentic Quantum Machine Learning (AQML) framework, which unifies perception, reasoning, and orchestration within a quantum-intelligent pipeline for medical imaging. The Quantum Reasoner Agent (QRA) transforms quantum modules from passive optimizers into autonomous, interpretable reasoning entities that interact dynamically with LLM controllers. Proof-of-concept simulation experiments demonstrate hybrid gains in optimization accuracy, uncertainty calibration, and diagnostic consistency over classical baselines within the 8-qubit simulation environment.

The present study acknowledges significant limitations in validation scope including a small training subset, outdated baselines, incomplete hardware noise evaluation, and unvalidated explainability which must be rigorously addressed in future work before clinical utility can be established. By bridging foundational QML theory [1], [4], variational adaptation [3], and agentic orchestration [8], [9], [41], [42], AQML advances a unified paradigm for quantum-intelligent radiology, integrating physical interpretability, clinical adaptability, and sustainable federated intelligence [10], [14].

Future work will address the following directions:

- Full-Dataset and External-Dataset Validation: Evaluate AQML on the complete BraTS 2023 dataset and independent external datasets to establish generalizability and address overfitting concerns.
- Modern Baseline Benchmarking: Benchmark against nnU-Net [16], SwinUNet [17], and TransUNet [18] to rigorously quantify the performance benefit of quantum augmentation.
- Real-Device Noise-Aware Evaluation: Evaluate diagnostic performance metrics (Dice, AUC, ECE) under IBM Quantum hardware noise conditions to establish realistic near-term feasibility.
- Clinical Validation of Explainability: Conduct formal evaluation of the explainability layer with board-certified radiologists, assessing accuracy, reliability, and alignment with clinical reporting standards.
- Multi-Modal Validation: Extend experiments to CT and PET modalities to validate cross-modal fusion.
- Federated Quantum Agentic Learning: Enable decentralized, privacy-preserving collaboration through entanglement-based or classical consensus across distributed QRA nodes [11].
- Quantum Reinforcement Learning: Develop policy-gradient loops allowing continual circuit adaptation for self-evolving diagnostic intelligence [8].

Clinically Tuned LLMs: Fine-tune radiology-specific language models to interpret quantum observables and generate ontology-aligned diagnostic narratives.

References

[1] M. Schuld and F. Petruccione, Quantum Machine Learning: An Introduction. Springer, 2019.

[2] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, no. 7747, pp. 209-212, 2019.

[3] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," New J. Phys., vol. 18, no. 2, p. 023023, 2018.

[4] A. Abbas et al., "The power of quantum neural networks," Nat. Comput. Sci., vol. 1, pp. 403-409, 2021.

[5] S. Gambs, M. Hayes, and J. Pichon, "Quantum approaches for medical imaging and diagnostics: A survey," IEEE Trans. Quantum Eng., vol. 2, pp. 1-14, 2020.

[6] Z. Li, M. Fan, and S. Wang, "Quantum-enhanced medical image classification via hybrid networks," IEEE J. Biomed. Health Inform., vol. 26, no. 12, pp. 5973-5985, 2022.

[7] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proc. MICCAI, Springer, 2015, pp. 234-241.

[8] J. Park, A. Bose, and R. Patel, "Agentic AI: Autonomy and feedback in hybrid LLM systems," Nat. Mach. Intell., vol. 6, no. 3, pp. 345-357, 2024.

[9] M. Chandrasekaran, "Enhancing efficiency and flexibility of rapid prototyping for scalable multimodal intelligent agents," in Proc. 2024 Artificial Intelligence for Business (AIxB),

Laguna Hills, CA, USA, 2024, pp. 66-71, doi: 10.1109/AIxB62249.2024.00019.

[10] R. Bianzino, "Sustainable quantum-AI infrastructures," npj Quantum Inf., vol. 11, no. 2, pp. 1-9, 2025.

[11] Y. Zhang, J. Zhao, and L. Wei, "Federated quantum machine learning for distributed healthcare imaging," IEEE Trans. Med. Imaging, vol. 42, no. 11, pp. 3234-3246, 2023.

[12] M. Caro et al., "Generalization in quantum machine learning from few training data," Nat. Commun., vol. 13, no. 1, pp. 4919-4928, 2022.

[13] M. Schuld, "Effect of data encoding on the expressive power of variational quantum machine learning models," Phys. Rev. A, vol. 103, no. 3, p. 032430, 2021.

[14] A. Gill and E. Huang, "Ethical and explainable quantum intelligence in clinical AI systems," IEEE J. Transl. Eng. Health Med., vol. 12, pp. 150-161, 2024.

[15] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," IEEE Trans. Autom. Control, vol. 37, no. 3, pp. 332-341, 1992.

[16] F. Isensee, P. F. Jäger, P. M. Full, P. Vollmuth, and K. H. Maier-Hein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," Nat. Methods, vol. 18, no. 2, pp. 203-211, 2021.

[17] H. Cao et al., "Swin-Unet: Unet-like pure transformer for medical image segmentation," in Proc. ECCV Workshops, Springer, 2022, pp. 205-218.

[18] J. Chen et al., "TransUNet: Transformers make strong encoders for medical image segmentation," arXiv:2102.04306, 2021.

[19] M. Cerezo et al., "Variational quantum algorithms," Nat. Rev. Phys., vol. 3, no. 9, pp. 625-644, 2021.

[20] K. Bharti et al., "Noisy intermediate-scale quantum algorithms," Rev. Mod. Phys., vol. 94, no. 1, p. 015004, 2022.

[21] Y. Kim et al., "Evidence for the utility of quantum computing before fault tolerance," Nature, vol. 618, no. 7965, pp. 500-505, 2023.

[22] S. Minaee et al., "Image segmentation using deep learning: A survey," IEEE Trans. Pattern Anal. Mach. Intell., vol. 44, no. 7, pp. 3523-3542, 2022.

[23] M. J. Moor et al., "Foundation models for generalist medical artificial intelligence," Nature, vol. 616, no. 7956, pp. 259-265, 2023.

[24] P. Rajpurkar and M. Lungren, "The current and future state of AI interpretation of medical images," N. Engl. J. Med., vol. 388, no. 21, pp. 1981-1990, 2023.

[25] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, "Hybrid quantum-classical algorithms and quantum error mitigation," J. Phys. Soc. Japan, vol. 90, no. 3, p. 032001, 2021.

[26] R. Aggarwal et al., "Diagnostic accuracy of deep learning in medical imaging: A systematic review and meta-analysis," npj Digit. Med., vol. 4, no. 1, p. 65, 2021.

[27] A. Arrasmith et al., "Equivalence of quantum barren plateaus to cost concentration and narrow gorges," Quantum Sci. Technol., vol. 7, no. 4, p. 045015, 2022.

[28] M. Antonelli et al., "The medical segmentation decathlon," Nat. Commun., vol. 13, no. 1, p. 4128, 2022.

[29] J. Biamonte et al., "Quantum machine learning," Nature, vol. 549, no. 7671, pp. 195-202, 2017.

[30] J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018.

[31] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," Phys. Rev. A, vol. 98, no. 3, p. 032309, 2018.

[32] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," Quantum Sci. Technol., vol. 4, no. 4, p. 043001, 2019.

[33] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in Proc. ICLR, 2021.

[34] H. Hatamizadeh et al., "UNETR: Transformers for 3D medical image segmentation," in Proc. WACV, 2022, pp. 574-584.

[35] X. Wang et al., "TransBTS: Multimodal brain tumor segmentation using transformer," in Proc. MICCAI, Springer, 2021, pp. 109-119.

[36] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information. Cambridge Univ. Press, 2010.

[37] Y. Du et al., "Quantum circuit architecture search for variational quantum algorithms," npj Quantum Inf., vol. 8, no. 1, p. 62, 2022.

[38] A. K. Goyal and P. Gupta, "Uncertainty estimation in deep learning for medical image segmentation," Expert Syst. Appl., vol. 234, p. 121115, 2023.

[39] A. Perez-Salinas et al., "Data re-uploading for a universal quantum classifier," Quantum, vol. 4, p. 226, 2020.

[40] Z. Zhou et al., "UNet++: A nested U-Net architecture for medical image segmentation," in Proc. MICCAI Deep MIS Workshop, 2018, pp. 3-11.

[41] M. Chandrasekaran, "AURELIUS: Agentic Uncertainty-Reasoning & Learning with Hierarchical Bayesian Multi-Agents Unified System for Handling Non-Determinism in HealthTech," in Proc. 4th Int. Conf. Computing, Management and Telecommunications (ComManTel), Madrid, Spain, 2025, pp. 52-57, doi: 10.1109/ComManTel68363.2025.11368546.

[42] "System and method for artificial intelligence based field service assistance for telecommunications operations," U.S. Patent 12 475 154, Jul. 2, 2025.