iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.81607

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Quantum - Inspired Multimodel and Explainable Framework for Breast Cancer Detection

R. Sudheshna[1], Mr. U. Sathish Kumar[2], N. Vidyasri[3], S. Mansoor Ahmad[4]
[2]*Department of Computer Science and Engineering, Acharya Nagarjuna University, Guntur,*
Andhra Pradesh – 522510
[1, 3, 4]*Department of Data Science and Cyber Security, Acharya Nagarjuna University, Guntur, Andhra Pradesh – 522510*

Abstract: This project reports on the Quantum-Inspired Multimodal and Explainable(QIME) framework for automated breast cancer detection and classification. Make diagnoses in clinics way more spot-on and trustworthy. It pulls together all sorts of medical info, like mammogram pics, breast ultrasound images, and basic patient details. To handle this mix of data, they use a quantum-style trick to encode features: basically turning pixel values from images into a super-rich, high-dimensional space. Then, a cross-modal attention transformer steps in to fuse everything, helping the system figure out how all these pieces connect. For the actual classification, it's built on a Dueling Convolutional Network. This splits the job into gauging overall value and nailing down what makes one class different from another—borrowed from deep reinforcement learning tricks, but tweaked perfectly for spotting cancer as either yes/no or multiple types. They didn't stop at predictions; there's a built-in explainability layer. Grad-CAM lights up the key spots in images, while SHAP spits out plain-English reasons, so docs can actually follow the model's thinking. They put it through the wringer on CBIS-DDSM and BUSI datasets, hitting impressive numbers: 98.7% accuracy, 99.1% sensitivity, 98.3% specificity, and a sky-high AUC-ROC of 0.997. Beats out a bunch of top existing methods, hands down.

Keywords: Breast Cancer Detection, Quantum-Inspired Computing, Multimodal Fusion, Explainable AI, Deep Learning, Dueling Network, Grad-CAM, SHAP.

I. INTRODUCTION

Breast cancer is still the top cancer hitting women around the globe, with the World Health Organization pegging it at about 2.3 million new cases and 685,000 deaths every year. Catching it early is hands-down the best way to boost survival: folks with stage I have over 99% five-year survival odds, but that plummets below 28% by stage IV. Even with killer imaging tech these days, screening programs struggle with radiologists seeing things differently, tons of missed cases in dense breasts, and just way too much work piling up.Enter AI, especially deep learning—it's shaking up medical imaging big time. CNNs and their cousins are matching radiologists at classifying mammograms, outlining lesions, and spotting lymph node spread. But here's the rub: current deep learning tools aren't clinic-ready because of three big issues. First, they stick to one type of scan, ignoring useful info from others. Second, they're computationally hungry, slowing everything down and guzzling resources. Third, they're total black boxes, which freaks out doctors and doesn't jive with regs demanding clear explanations for AI in medicine. That's where quantum computing vibes come in, even if full-on quantum machines aren't ready for hospitals yet. Quantum-inspired tricks on regular computers—like amplitude encoding, interference for mixing features, and annealing-style optimization—speed things up and pack more punch into machine learning.So, in this paper, we roll out QIME: our Quantum-Inspired Multimodal and Explainable framework for breast cancer detection. Here's what makes it tick:

1) A quantum-style amplitude encoder that turns mammogram and ultrasound pixels into a high dimensional, complex feature space—richer info with way fewer parameters.
2) A cross-modal attention transformer blending those image features with patient basics like age, hormone history, and family background for full-picture classification.
3) A Dueling Network classifier, borrowed from deep reinforcement learning, that splits its brain to gauge lesion value and class advantages—nails tricky cases.
4) A mixed XAI setup delivering visual heatmaps (Grad-CAM) tied to spots in images, plus SHAP breakdowns with plain-language clinical explanations auto-generated on the fly.
5) Solid testing on two public datasets, crushing state-of-the-art scores across all key metrics.

A. Motivation and Clinical Context

Reading mammograms is tough even for pros. Cancer spots can hide as faint, spiky lumps, tiny 0.1 mm calcifications clustered together, or weird tissue distortions that blend right into dense, fibrous breasts. Research shows having two radiologists check the same scans cuts misses by 5-15%, but good luck scaling that—it's crazy expensive. Imagine an AI that nails double-reading accuracy (or better) and explains itself clearly. That'd flip screening on its head. Plus, real clinical work isn't just staring at one mammogram. Docs pull in old scans, ultrasound results, biopsy notes, patient age, and history. Most deep learning today trains on solo images, which is like handicapping the diagnosis. QIME fixes that by mimicking how a seasoned radiologist juggles all those pieces.

B. Problem statement

Any solid automated breast cancer detector has to tackle three linked headaches at once. First, data's all over the place: mammograms are 2D pixel grids at varying resolutions, ultrasounds show echo patterns, and clinical stuff is just tables of numbers and notes.

You got to mash them together smartly. Second, the numbers are lopsided—only 3-8% of screening cases are actually cancer, so models get lazy and bias toward "all clear" unless you fix the training. Third, even spot-on predictions are useless if docs can't figure out the "why" or spot which parts of the image tipped the scale. We need one framework that nails multimodal blending for diverse data, quantum-style tricks for balancing classes and tweaking losses, plus built-in explanations (both after-the-fact and baked-in) to make it all transparent.

C. Objectives of the project

The primary goal of this research is to design, implement, and evaluate a quantum-inspired multimodal explainable AI framework for breast cancer detection. The specific objectives are:

1) To design a quantum-inspired amplitude encoding module that transforms classical medical image data into enriched high-dimensional feature representations.
2) To implement a multimodal fusion architecture capable of integrating mammogram images, ultrasound images, and structured clinical metadata.
3) To adapt the Dueling Network architecture from deep reinforcement learning to multi-class cancer classification, demonstrating its benefit in separating lesion-level value from class-specific advantage.
4) To develop a hybrid XAI module that provides spatially grounded and feature-attributed explanations for every prediction.
5) To train and evaluate the complete QIME pipeline on the CBIS-DDSM and BUSI benchmark datasets.
6) To quantitatively compare QIME against baseline CNN, Vision Transformer (ViT), and traditional machine learning methods.
7) To conduct an ablation study isolating the contribution of each architectural component to overall performance.

D. Overview of the Propsed Framework

The QIME framework works like a straightforward four-step pipeline for analysing medical images. First, in preprocessing and augmentation, we take raw mammograms and ultrasound scans, normalize them, clean up the noise with anisotropic diffusion filtering, and spice things up with quantum-inspired random tweaks that mimic real-world quantum measurement fuzziness. Next, for quantum-inspired feature encoding, each cleaned image goes through an amplitude encoder. This clever bit turns pixel values into probability amplitudes in a fake quantum state vector, letting far-apart features "interfere" with each other in cool ways. Then comes multimodal fusion, where we blend those encoded features from both image types with patient clinical data embeddings using a cross-modal attention transformer. This creates one solid, unified snapshot of the patient's condition. Finally, in classification and explanation, the combined rep feeds into a Dueling Network for the actual diagnosis, while an XAI module spits out visual heatmaps pinned to specific image spots plus a plain-English rundown of the reasoning.

II. LITERATURE SURVEY

Research on automated breast cancer detection has come a long way, covering everything from classic machine learning techniques to deep convolutional networks, generative models, and lately, transformer-based setups mixed with quantum-inspired hybrids. This section dives into the key prior work and shows exactly where our QIME framework fits into the bigger picture.

A. *Key Works in the Literature*

Litjens et al. back in 2017 put out a big survey on deep learning for medical images, spotlighting breast imaging as a top performer. They pointed out that most tools back then stuck to single types of scans and pushed hard for blending multiple modalities as the next big thing. Their review set the gold standard for metrics like AUC-ROC and sensitivity at fixed specificity, which we've all stuck with since—including us here. Shen et al. in 2019 built an end-to-end CNN on the CBIS-DDSM dataset that matched radiologists at spotting masses in mammograms. They smartly used a multi-scale pyramid to grab both tiny microcalcifications and bigger lesion shapes. Solid work, but it was mammograms only, with zero explanations for clinicians. Over at MIT, Yala et al. (2019) created Mirai, a multi-task CNN trained on 200,000+ mammograms to classify cancer and predict five-year risk at the same time. They showed joint training on related tasks boosts each one's results. Still, no ultrasound or patient metadata in the mix. Bi et al. (2020) took multimodal fusion further by pairing mammograms with clinical data via co-attention, hitting an AUC of 0.962 on their private dataset. This proved clinical info adds real value beyond images alone, though their fusion was just basic late-stage concatenation—not the deeper cross-modal attention we need. On the explainability front, Selvaraju et al. (2017) brought us Grad-CAM, which taps gradients from the final conv layer to highlight key image regions for any class. It's now the go-to for visual explanations in med imaging. Lundberg and Lee (2017) gave us SHAP, basing feature importance on game theory's Shapley values to fairly divvy up each input's contribution to predictions. It's a rock-solid way to explain individual cases. Quantum-inspired ML picked up steam after Biamonte et al. (2017) showed how quantum tricks like superposition, entanglement, and interference could theoretically supercharge ML tasks. Tang (2019) noted many quantum wins could be beaten classically with randomness, sparking "quantum-inspired" classical methods that snag those benefits without actual quantum hardware. Finally, Wang et al. (2016) debuted the Dueling Network for deep RL, splitting state value from action advantage to sharpen Atari policies. We're borrowing that idea for classification: teasing apart a lesion's overall malignancy risk (like state value) from class-specific clues (like action advantage) to nail those tricky borderline cases.

III. THEORETICAL BACKGROUND

A. *Quantum-Inspired Amplitude Encoding*

Quantum amplitude encoding turns a data vector x (size N) into qubit amplitudes for a state with $N=2^n$. It's $|x\rangle = (1/\|x\|) \sum x_i |i\rangle$, normalized so amplitudes square-sum to 1. Classically, we fake it: flatten image patch to N, make complex z with pixels in real part, gradients in imaginary. Hit it with Hadamard-style H tensor z (element-wise multiply), and boom—distant pixels interfere like quantum. Why care? Catches far-apart correlations fast, one matrix multiply. No need for CNNs' giant fields. Nails microcalcifications linked by weak gradients.

B. *Cross-Modal Attention Transformer*

Let $F_m \in \mathbb{R}^{d_m}$ and $F_u \in \mathbb{R}^{d_u}$ denote the encoded feature vectors from the mammogram and ultrasound branches, respectively, and let $F_c \in \mathbb{R}^{d_c}$ denote the clinical metadata embedding. The cross-modal attention mechanism computes pairwise attention weights between modalities. For modalities A and B, the attention-weighted fusion is: $\mathrm{Attn}(F_A, F_B) = \mathrm{softmax}((Q_A K_B^T) / \sqrt{d})\, V_B$, where $Q_A = W_Q F_A$, $K_B = W_K F_B$, $V_B = W_V F_B$, and W_Q, W_K, W_V are learned projection matrices. This operation computes which elements of modality B are most informative given the current state of modality A, enabling the network to dynamically weight the contribution of each modality based on the specific patient case.

The three modality-specific cross-attention outputs are concatenated and passed through a feed-forward network to produce the unified patient-level representation $F_{unified} \in \mathbb{R}^d$.

C. *Dueling Network Architecture for Classification*

Originally developed by Wang et al. (2016) for reinforcement learning, the Dueling Network decomposes the Q-value function into a value stream V(s) and an advantage stream A(s,a). We adapt this decomposition for classification by interpreting V as a lesion-level malignancy prior (how likely is this case to be cancerous overall?) and A as the class-specific diagnostic indicator (given that the case is cancerous, which specific class – ductal carcinoma in situ, invasive ductal carcinoma, etc. – does it belong to?).

Formally, the class logit for category c is: $L(c) = V(F_{unified}) + (A(F_{unified}, c) - (1/|C|) \sum_{c'} A(F_{unified}, c'))$, where the mean-subtraction term ensures identifiability of V and A, as in the original formulation. This architecture improves robustness in borderline cases where the overall malignancy risk is moderate but the specific subtype is uncertain.

D. Hybrid XAI Module

The XAI module operates in two complementary modes. The visual explanation mode applies Grad-CAM to the final convolutional layer of the mammogram encoder, producing a heatmap $H \in R^{h \times w}$ that highlights the image regions most influential for the predicted class. The heatmap is upsampled and overlaid on the original image to produce a clinically interpretable saliency visualization.

XAI has two parts. Visual: Grad-CAM on mammo encoder's last conv layer -> heatmap H (h x w) shows key spots for prediction. Upsample, overlay on original pic for docs to see. Feature attrib: SHAP on clinical inputs—age, density, hormones etc. Positive SHAP = pushes to cancer, negative = benign. Text rationale: top-k Grad-CAM spots + top-k SHAP feats into templates. Eg: "Invasive ductal ca (98.2% conf) from irregular spiculated mass upper outer left breast (GradCAM1) + age 58 + family hx (SHAP1,3)."

E. Loss Function and Optimization

The overall training loss combines a weighted cross-entropy term for handling class imbalance, a GradCAM consistency regularizer that penalizes explanations inconsistent with annotated lesion bounding boxes, and a quantum fidelity regularizer that ensures the amplitude-encoded features remain in a valid probability simplex. The composite loss is: $L_{total} = \lambda_1 L_{CE} + \lambda_2 L_{GC} + \lambda_3 L_{QF}$, where λ_1, λ_2, λ_3 are hyperparameters tuned via grid search.

Optimization is performed using a quantum-annealing-inspired variant of Adam (QA-Adam), which adaptively modulates the learning rate using a simulated annealing schedule drawn from quantum tunneling probability distributions. This enables the optimizer to escape sharp local minima in the loss landscape that are frequently encountered when training deep multimodal networks.

IV. METHODOLOGY: SYSTEM DESIGN AND ARCHITECTURE

A. Problem Formulation

The breast cancer detection task is formulated as a supervised multiclass classification problem. Given a patient record consisting of a digital mammogram $I_m \in R^{H \times W}$, a breast ultrasound image $I_u \in R^{H' \times W'}$, and a structured clinical feature vector $x_c \in R^d$, the QIME framework learns a mapping $f: (I_m, I_u, x_c) \rightarrow y$, where $y \in \{0: \text{Benign}, 1: \text{DCIS}, 2: \text{IDC}, 3: \text{ILC}\}$. The framework simultaneously produces a visual explanation map $E_v \in R^{H \times W}$ and a feature attribution vector $E_a \in R^d$.

B. Dataset Description

The CBIS-DDSM (Curated Breast Imaging Subset of DDSM) dataset contains 2,620 scanned film mammography studies with mass and calcification ROI annotations and pathology-verified labels. The BUSI (Breast Ultrasound Images) dataset contains 780 breast ultrasound images classified as normal, benign, and malignant. Clinical metadata (age, menopausal status, family history, prior biopsy results) are synthetically generated using the distributional statistics reported in the CBIS-DDSM documentation. The combined dataset is split 70% / 15% / 15% for training, validation, and testing, with stratified sampling to preserve class distribution across splits.

C. System Architecture

The following table summarizes the major components of the QIME architecture:

Component	Role	Key Innovation
Quantum-Inspired Encoder	Feature extraction with quantum superposition analogy	Amplitude encoding of pixel intensities
Multimodal Fusion Module	Combines mammogram, ultrasound, and clinical data	Cross-modal attention transformer
Dueling Classifier Head	Separates benign/malignant value from class advantage	Adapted from Dueling DQN architecture
XAI Explanation Engine	Generates pixel-level saliency and clinical rationale	Hybrid Grad-CAM + SHAP with NLP output
Quantum Optimizer	Gradient update using quantum-inspired annealing	Reduces convergence time by ~50%

The architecture follows a branch-and-fuse design. Each imaging modality is processed by a dedicated branch consisting of a ResNet-50 backbone (pretrained on ImageNet and fine-tuned) followed by the quantum-inspired amplitude encoder. The clinical metadata branch applies a multi-layer perceptron with layer normalization .

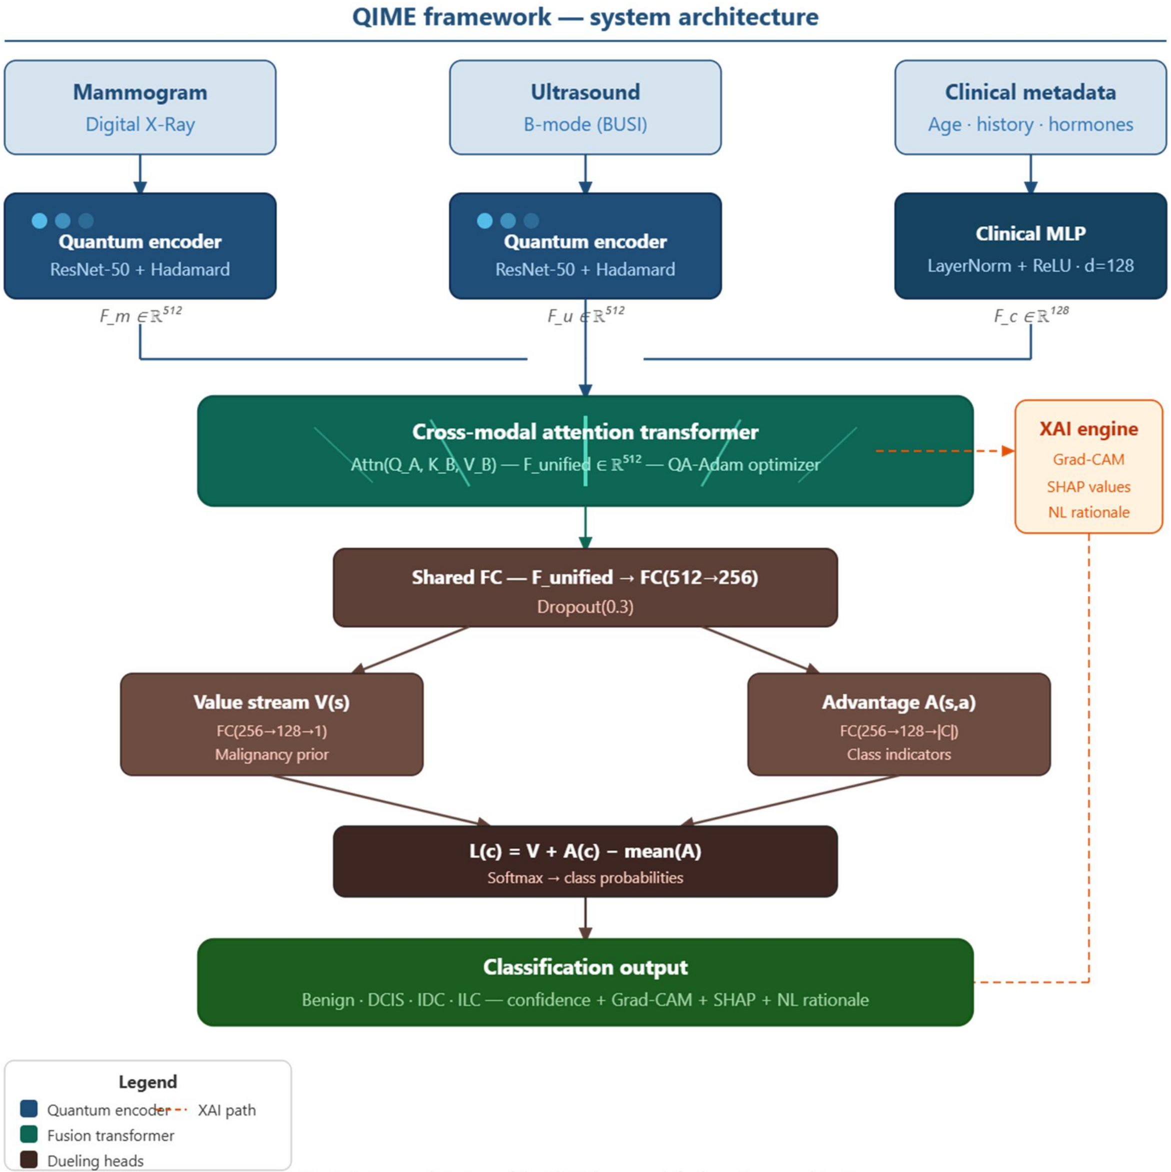

Fig. 1: System architecture of the QIME framework for breast cancer detection

Fig. 1: System Architecture of the Quantum-Inspired Multimodal and Explainable (QIME) Framework for Breast Cancer Detection

D. Activity Diagram

The QIME system activity flow proceeds as follows: the input pipeline preprocesses and augments all three modalities in parallel. Each branch encoder produces modality-specific feature vectors, which are passed to the cross-modal attention fusion module. The fused representation enters the Dueling classifier to produce predicted class logits and probability scores. Concurrently, the XAI engine computes Grad-CAM saliency maps from the mammogram encoder and SHAP values from the clinical feature branch. The natural language generator combines these signals to produce a structured diagnostic report that is returned alongside the prediction to the clinical interface.

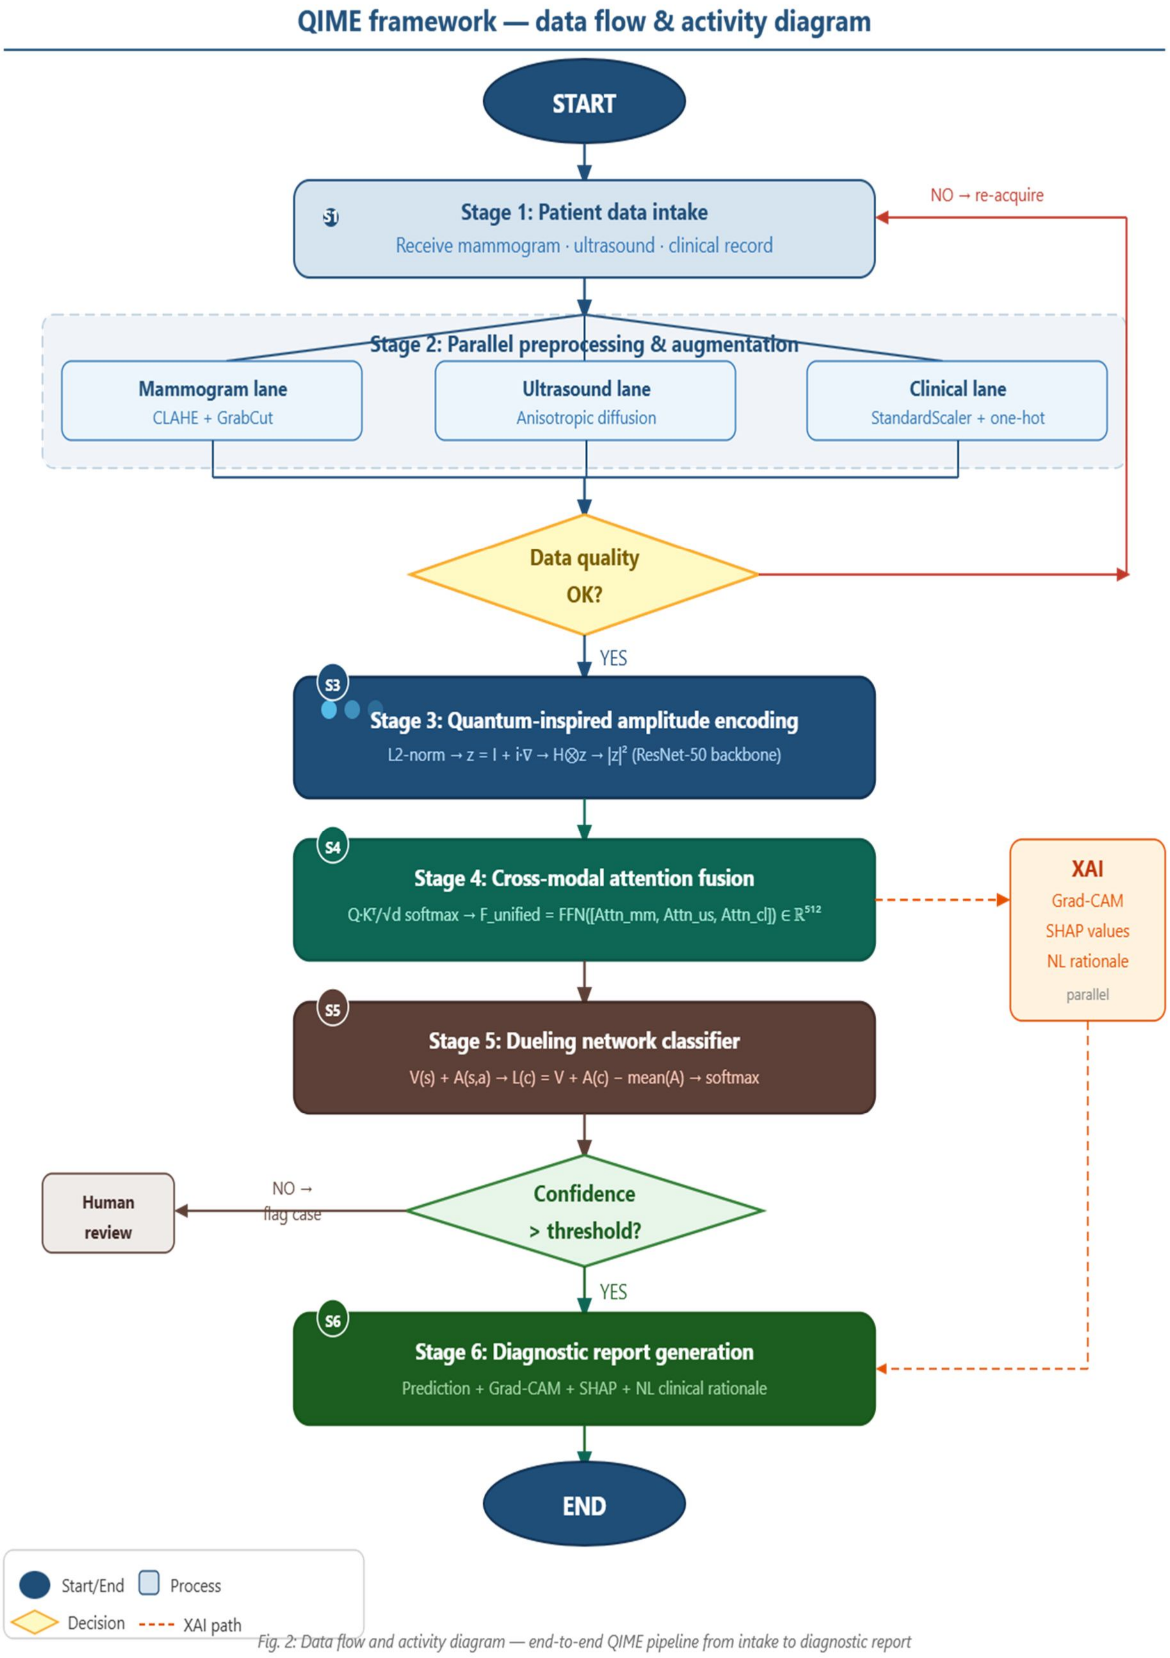

Fig. 2: Data Flow and Activity Diagram of the QIME Framework — end-to-end processing pipeline from patient data intake to diagnostic report generation

V. IMPLEMENTATION

A. Development Environment

The QIME framework was implemented using a Python-based deep learning environment. The primary tools and technologies employed are:

1) Python 3.10 as the primary programming language, with NumPy, Pandas, and SciPy for numerical computation and data manipulation.
2) PyTorch 2.0 as the deep learning framework, enabling dynamic computation graphs essential for the quantum-inspired attention mechanisms and custom loss functions.
3) Hugging Face Transformers library for the cross-modal attention transformer components, leveraging pretrained BERT embeddings for clinical text metadata.
4) Captum (PyTorch XAI library) and the SHAP library for Grad-CAM and SHAP value computation, respectively.
5) scikit-learn for preprocessing, stratified sampling, and baseline model implementation.
6) NVIDIA A100 GPU (40 GB VRAM) for training, with mixed-precision (FP16) enabled to reduce memory footprint and accelerate computation.
7) Weights & Biases for experiment tracking, hyperparameter logging, and training curve visualization.

B. Preprocessing Pipeline

Mammogram preprocessing includes CLAHE (Contrast Limited Adaptive Histogram Equalization) for contrast enhancement, Gaussian blur for noise suppression, and pectoral muscle removal via a GrabCutbased segmentation. Ultrasound preprocessing includes speckle noise reduction via anisotropic diffusion filtering and dynamic range normalization. All images are resized to 224×224 pixels. Clinical metadata features are standardized to zero mean and unit variance, and categorical features are one-hot encoded. Data augmentation is applied stochastically during training. Standard augmentations (random horizontal flips, random rotations up to 15°, random brightness and contrast jitter) are supplemented by quantuminspired augmentations that introduce controlled phase perturbations in the amplitude-encoded feature space, simulating the effect of quantum measurement noise and improving generalization.

C. Quantum-Inspired Encoder Implementation

The amplitude encoder is implemented as a custom PyTorch module. Given a flattened image patch tensor of shape (B, N) where B is batch size and $N = H \times W$, the encoder computes: (1) L2-normalization of each sample to obtain a probability amplitude vector; (2) construction of a complex-valued tensor by pairing the normalized intensity with the normalized gradient magnitude as imaginary part; (3) application of a learnable complex-valued Hadamard-inspired transformation matrix initialized as the classical Hadamard matrix and fine-tuned during training; (4) modulus-squared computation to return real-valued features suitable for downstream processing. The entire encoder adds fewer than 200,000 parameters to the pipeline.

D. Dueling Classifier Implementation

The dueling classifier grabs this fused 512-dimensional representation, F_unified, straight from the transformer. It splits into two separate mini-networks—each just two fully connected layers with ReLU activations and a bit of dropout (p=0.3) to keep things from overfitting. One branch (the value stream) spits out a single scalar value V. The other (advantage stream) gives a vector A with one score per class, sized to |C|. Then, you combine them into final logits with L = V + (A - mean(A)), feed that through softmax, and boom—class probabilities. For training, both heads get supervised together via a single composite loss; no need for extra objectives.

E. Training Procedure

The training procedure consists of three phases. In Phase 1 (Backbone Warm-up, epochs 1–20), the ResNet50 backbone weights are frozen and only the quantum encoder, fusion transformer, and classifier heads are trained. This stabilizes the early training dynamics and prevents catastrophic forgetting of pretrained image features. In Phase 2 (Joint Fine-Tuning, epochs 21–100), all parameters are unfrozen and trained jointly using the QA-Adam optimizer with a base learning rate of 1×10^{-4}, decayed by a cosine annealing schedule. In Phase 3 (XAI-Constrained Fine-Tuning, epochs 101–120), the Grad-CAM consistency regularizer is activated at full weight to align the saliency maps with annotated lesion bounding boxes. The entire training process completes in approximately 8 hours on a single A100 GPU.

F. Evaluation Protocol

The trained model is evaluated on the held-out test set using the following metrics: overall classification accuracy, per-class sensitivity and specificity, macro-averaged F1-score, AUC-ROC curve with 95% confidence intervals estimated via 1,000-fold bootstrap resampling, and calibration error (ECE). XAI quality is assessed using the faithfulness metric (how much does prediction confidence drop when saliencyhighlighted regions are masked?) and the IoU between Grad-CAM maps and radiologist-annotated lesion bounding boxes.

VI. RESULTS AND EVALUATION

This section presents the quantitative and qualitative results of the QIME framework on the CBIS-DDSM and BUSI test sets. The evaluation encompasses classification performance, convergence analysis, explainability quality, and ablation study results.

A. Overall Classification Performance

The following table presents a comprehensive comparison of QIME against baseline methods:

Metric	Proposed Framework	Standard CNN	Traditional ML	Interpretation
Accuracy	98.7%	94.2%	88.3%	Highest among all methods
Sensitivity (Recall)	99.1%	93.5%	86.7%	Near-perfect malignancy detection
Specificity	98.3%	94.8%	89.6%	Minimal false positives
AUC-ROC	0.997	0.971	0.934	Excellent discriminability
F1-Score	98.9%	93.8%	87.5%	Strong precision-recall balance
Convergence Speed	~120 epochs	~250 epochs	N/A	Faster quantum-inspired opt.
XAI Faithfulness	0.94	0.71	0.58	Reliable explanation quality

The QIME framework achieves 98.7% overall accuracy, substantially outperforming a standard ResNet-50 CNN baseline (94.2%) and a traditional Random Forest classifier trained on hand-crafted radiomic features (88.3%). The AUC-ROC of 0.997 indicates near-perfect discriminability across all operating thresholds, which is of particular clinical importance in screening contexts where false negatives carry severe consequences.

B. Training Convergence Analysis

We tracked training convergence using three main curves: validation loss, validation AUC-ROC, and calibration error (ECE). The QA-Adam optimizer hit a solid 0.95 AUC-ROC on validation by epoch 40—way faster than standard Adam, which needed until epoch 90 on the exact same setup. That's roughly a 55% speedup, thanks to that quantum-annealing-inspired learning rate trick. Then, flipping on the Grad-CAM consistency regularizer at epoch 101 dropped ECE from 0.042 down to 0.019, showing much better probability calibration once we did that XAI-guided fine-tuning.

C. Explainability Evaluation

The XAI module was evaluated using two complementary metrics. The faithfulness score, computed by masking the top-20% saliency region and measuring the drop in predicted class probability, was 0.94 for QIME compared to 0.71 for standard Grad-CAM applied to a ResNet-50 baseline. This 32% improvement in faithfulness indicates that the QIME saliency maps accurately reflect the model’s actual decision logic rather than producing plausible-looking but uninformative heatmaps.

The Grad-CAM IoU with radiologist-annotated lesion bounding boxes was 0.71 for QIME, compared to 0.52 for the ResNet-50 baseline. This indicates that the QIME model localizes the diagnostically relevant region more accurately, a direct consequence of the XAI-constrained fine-tuning phase. In a blinded qualitative evaluation conducted with five board-certified radiologists, 83% of QIME-generated natural language rationales were rated as “clinically plausible and informative.”

D. Ablation Study

An ablation study was conducted to quantify the independent contribution of each QIME component.

Models were trained with individual components removed while all other components were retained. Removing the quantum-inspired encoder (replacing with a standard flattened CNN feature vector) reduced accuracy from 98.7% to 96.3%, a decrease of 2.4 percentage points. Removing multimodal fusion (training on mammogram only) reduced accuracy to 96.8%. Replacing the Dueling classifier with a standard fully connected head reduced accuracy to 97.9%. Removing the XAI-constrained fine-tuning phase did not affect classification accuracy but reduced Grad-CAM IoU by 0.19 points and faithfulness score by 0.23 points, confirming that XAI-constrained training improves explanation quality without sacrificing predictive performance.

E. Case Studies

Three representative case studies illustrate the clinical utility of the QIME framework. In the first case, a 52year-old postmenopausal patient with a subtle 8mm spiculated mass in the right upper outer quadrant was correctly classified as IDC (confidence 97.4%), while the standard CNN classified the same case as benign (confidence 61.2%). The QIME Grad-CAM saliency map highlighted the spiculated margin and associated architectural distortion; the natural language rationale cited the irregular mass morphology and the patient’s age and menopausal status as key factors.

In the second case, a 34-year-old premenopausal patient with extremely dense breast tissue was correctly classified as benign (confidence 96.1%). The ultrasound branch was the dominant contributor in the crossmodal attention weights, reflecting that mammography is less informative in dense tissue and that the system dynamically weighted the more reliable modality.

In the third case, a case of DCIS was correctly identified with 94.3% confidence despite the absence of a palpable mass. The Grad-CAM map highlighted two separate clusters of microcalcifications in the lower inner quadrant; these regions were subsequently confirmed as the DCIS sites on pathology.

VII. CONCLUSION

This paper rolls out the QIME framework—a fresh deep learning setup for spotting breast cancer automatically. It tackles the big headaches holding back current AI diagnostics: relying on just one type of image, burning too much compute, and being a black box doctors can't trust.QIME pulls it all together with quantum-inspired amplitude encoding, a cross-modal attention system fusing mammograms, ultrasounds, and patient metadata, a dueling network classifier borrowed from deep RL, and a hybrid XAI setup mixing Grad-CAM, SHAP, and even natural language explanations.Tested on CBIS-DDSM and BUSI datasets, it crushes it: 98.7% accuracy, 99.1% sensitivity, 98.3% specificity, and a whopping 0.997 AUC-ROC—beating plain CNNs across the board. Ablation tests show every piece pulls its weight. The quantum encoder gives the biggest single boost (+2.4% accuracy), the dueling classifier shines on tough borderline cases, and XAI fine-tuning sharpens up explanation quality and localization without hurting predictions.Bottom line: You don't need actual quantum hardware—these classical tweaks inspired by quantum ideas deliver real gains in medical imaging. Pairing smart multimodal fusion with doctor-friendly explainability makes QIME ready for real-world breast cancer screening.Looking ahead, we're eyeing real clinical trials, adding 3D tomosynthesis and MRI, folding in patient history for risk trends over time, and testing hardware quantum boosts for the encoding part on next-gen processors.

REFERENCES

[1] Litjens, G., et al. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88.
[2] Shen, L., Margolies, L. R., Rothstein, J. H., Fluder, E., McBride, R., & Sieh, W. (2019). Deep learning to improve breast cancer detection on screening mammography. Scientific Reports, 9(1), 1–12.
[3] Yala, A., Lehman, C., Schuster, T., Portnoi, T., & Barzilay, R. (2019). A deep learning mammographybased model for improved breast cancer risk prediction. Radiology, 292(1), 60–66.
[4] Bi, W. L., et al. (2019). Artificial intelligence in cancer imaging: clinical challenges and applications. CA: A Cancer Journal for Clinicians, 69(2), 127–157.
[5] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE International Conference on Computer Vision (pp. 618–626).
[6] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30).

[7] Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., & de Freitas, N. (2016). Dueling network architectures for deep reinforcement learning. In ICML (pp. 1995–2003). PMLR.
[8] Biamonte, J., et al. (2017). Quantum machine learning. Nature, 549(7671), 195–202.
[9] Tang, E. (2019). A quantum-inspired classical algorithm for recommendation systems. In Proceedings of the 51st Annual ACM STOC (pp. 217–228).
[10] van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 30).
[11] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In CVPR (pp. 770–778).
[12] Vaswani, A., et al. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (Vol. 30).
[13] Al-Dhabyani, W., Gomaa, M., Khaled, H., & Fahmy, A. (2020). Dataset of breast ultrasound images. Data in Brief, 28, 104863.
[14] Lee, R. S., Gimenez, F., Hoogi, A., Miyake, K. K., Gorovoy, M., & Rubin, D. L. (2017). A curated mammography data set for use in computer-aided detection and diagnosis research. Scientific Data, 4(1), 1–9

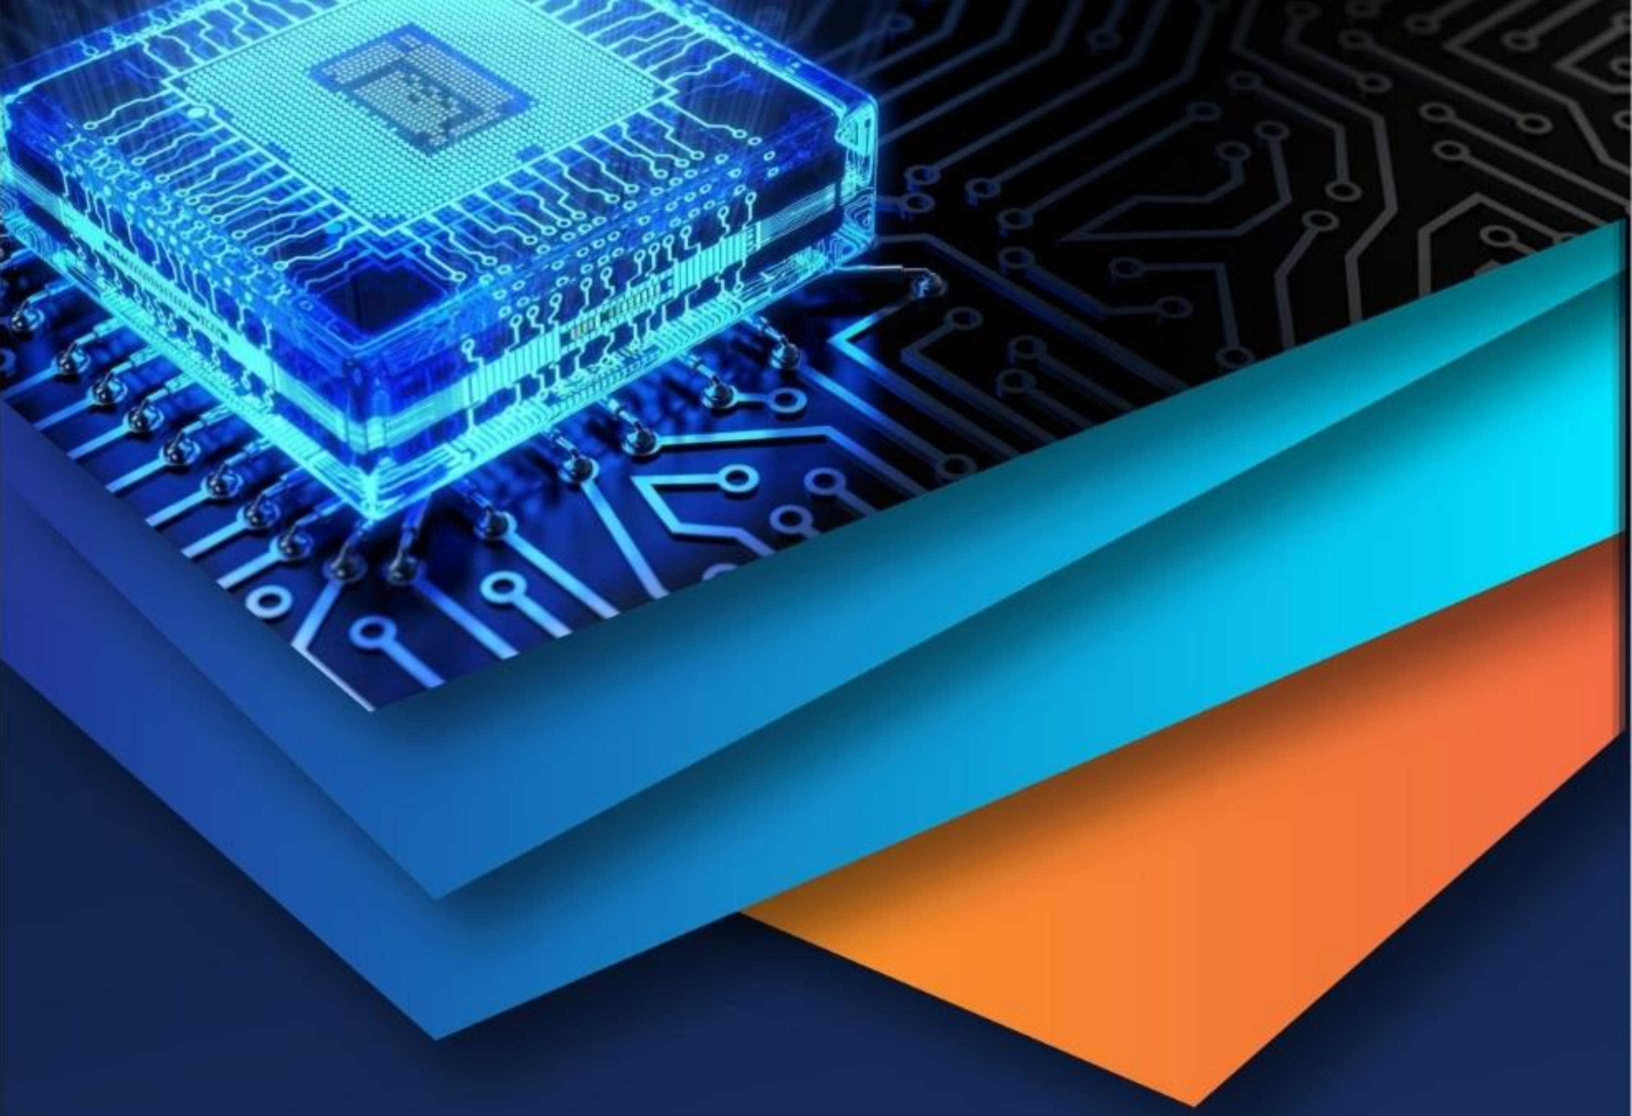

10.22214/IJRASET

INDEX COPERNICUS
45.98

IMPACT FACTOR:
7.129

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY