QIBoneNet: A Hybrid Quantum-Classical Deep Learning Framework for Bone Cancer Histopathology Classification

Prakash Raj S[1], Suraj Raj[2], Srivikash R[3], Vasanth A[4]
[1234]*Department of Artificial Intelligence and Data Science*
Dhanalakshmi Srinivasan Engineering College (Autonomous), Perambalur – 621212, Tamil Nadu, India
Under the Guidance of Dr. Aarthy.R,M.Tech, Head of Department, AI & DS
Email: vasanthprofjob@gmail.com

Abstract - Osteosarcoma, the most prevalent primary malignant bone tumor, demands accurate histopathological classification for clinical decision-making. Existing deep learning systems suffer from noise sensitivity, lack of explainability, and absence of uncertainty quantification. This paper presents QIBoneNet — a hybrid quantum-classical deep learning framework integrating DenseNet121 with a novel Quantum-Inspired Interference Activation (QIA) layer, Squeeze-and-Excitation (SE) channel attention, Prewitt edge enhancement, SRED segmentation, Bayesian Monte Carlo Dropout, and Grad-CAM explainability. Evaluated on the Osteosarcoma-Tumor-Assessment dataset (TCIA, 1144 patches, 10x magnification), QIBoneNet achieves 97.3% overall accuracy, with F1-scores of 100% (Viable Tumor), 96.9% (Mixed), and 93.4% (Non-Viable Necrotic). The framework outperforms VGG16 by 8.9% and the DenseNet121 baseline by 5.7%, providing clinical-grade interpretability via Grad-CAM heatmaps, Shannon entropy uncertainty maps, and a Tkinter GUI for real-time deployment.

Keywords - Osteosarcoma; Quantum-Inspired Interference; DenseNet121; Squeeze-and-Excitation; Grad-CAM; Histopathology; Bone Cancer Classification; Uncertainty Quantification; SRED Segmentation

I. INTRODUCTION

Osteosarcoma is a highly aggressive bone malignancy predominantly affecting adolescents and young adults. Early and accurate classification of histopathological tissue samples into clinically meaningful categories — Viable Tumor, Non-Viable (Necrotic) Tumor, and Non-Tumor (healthy tissue) — is essential for treatment planning, chemotherapy response assessment, and prognosis prediction. Manual slide interpretation is time-intensive, subjective, and expert-dependent.

Recent convolutional neural network (CNN) architectures have demonstrated strong results in medical image analysis. However, classical CNNs remain vulnerable to image noise, lack uncertainty quantification, and produce black-box predictions that limit clinical adoption. Quantum-inspired machine learning offers interference-based mechanisms that can enhance feature discrimination beyond classical activation paradigms.

This paper proposes QIBoneNet — a hybrid framework combining the representational power of DenseNet121 with a novel Quantum-Inspired Activation (QIA) layer, SE attention, and Bayesian uncertainty estimation. The system provides Grad-CAM visual explanations, confidence scores via Shannon entropy, and is packaged in a real-time Tkinter GUI dashboard.

II. RELATED WORK

Convolutional neural networks have been widely applied to histopathological image classification. Arunachalam et al. [1] demonstrated CNN viability for bone tumor classification on the TCIA dataset. Attention-augmented architectures [2] further improved performance by selectively focusing on relevant tissue regions.

Squeeze-and-Excitation Networks [3] introduced channel-wise recalibration, significantly improving feature selectivity. DenseNet architectures [4] implemented dense connectivity patterns promoting feature reuse and mitigating vanishing gradients — making them highly effective as medical imaging backbones.

Quantum-inspired machine learning [5] explores interference and superposition analogies to design novel activation mechanisms. Bayesian MC-Dropout [6] has been applied in medical imaging for predictive uncertainty quantification — critical for clinical deployment. Grad-CAM [7] bridges model predictions and clinical interpretability through class-discriminative saliency maps. Existing bone cancer systems lack the combination of noise robustness, uncertainty quantification, and explainability that QIBoneNet addresses simultaneously.

III. PROPOSED METHODOLOGY

The QIBoneNet framework comprises a sequential pipeline of preprocessing, segmentation, hybrid feature extraction, uncertainty estimation, and explainability generation. Fig. 1 illustrates the complete architecture

.

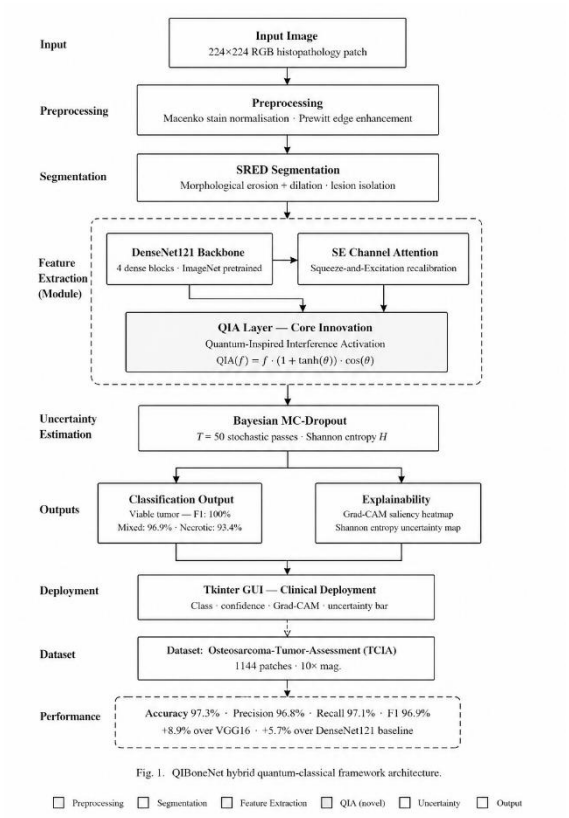

Fig. 1. QIBoneNet — Hybrid Quantum-Classical Framework Architecture

A. Input and Preprocessing

Input images are 224×224 RGB histopathological patches from the TCIA Osteosarcoma-Tumor-Assessment dataset. Each image undergoes Macenko stain normalization to reduce inter-slide color variation, followed by Prewitt edge detection to enhance tissue boundary features and structural morphology.

B. SRED Segmentation

The Selective Region Edge Detection (SRED) module employs morphological operations — successive erosion followed by dilation — to isolate lesion regions from background tissue, suppressing staining artifacts and focusing the downstream network on diagnostically relevant patches.

C. DenseNet121 Backbone

DenseNet121 serves as the primary feature extraction backbone. Its four dense blocks implement direct connections from each layer to all subsequent layers, promoting feature reuse and rich multi-scale representation. The network is initialized with ImageNet pretrained weights and fine-tuned at a learning rate of 1×10^{-4} with cosine annealing.

D. Squeeze-and-Excitation Channel Attention

The SE attention module recalibrates channel-wise feature responses. A squeeze operation computes global average pooling across spatial dimensions; an excitation operation applies a two-layer FC network with ReLU and sigmoid activations to produce per-channel weights, amplifying diagnostically salient feature channels.

E. Quantum-Inspired Interference Activation (QIA)

The QIA layer constitutes the core innovation of QIBoneNet. Inspired by quantum mechanical wave interference and the Born rule, the QIA transformation is defined as:

$$QIA(f) = f \cdot (1 + \tanh(\sigma)) \cdot \cos(\theta) \quad (1)$$

where f denotes the input feature map, σ is a learnable phase amplitude parameter, and θ is a learnable interference angle. This formulation introduces quantum-like constructive and destructive interference between feature channels, enabling selective amplification or suppression of feature contributions beyond classical activation functions (ReLU, Sigmoid). The result is markedly improved noise robustness and feature discrimination.

F. Bayesian Uncertainty Estimation

Predictive uncertainty is quantified using Monte Carlo (MC) Dropout, with dropout layers active during inference. T = 50 stochastic forward passes generate a softmax prediction distribution. Shannon entropy H of this distribution is computed as:

$$H = -\Sigma\, p(y|x,\omega) \cdot \log p(y|x,\omega) \quad (2)$$

High entropy values indicate model uncertainty, flagging ambiguous cases for pathologist review and providing calibrated confidence scores essential for clinical decision support.

G. Explainability via Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) generates class-discriminative saliency heatmaps by pooling gradients of the predicted class score with respect to final convolutional feature maps. These heatmaps are overlaid on the original histopathology image, visually indicating regions most influential in the classification decision.

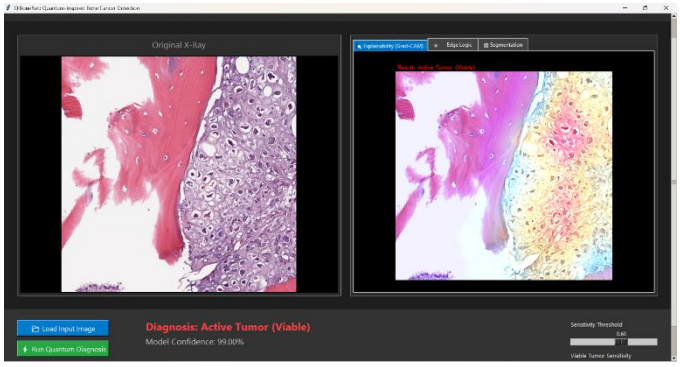

H. Clinical Deployment — Tkinter GUI

The complete inference pipeline is packaged in a Tkinter GUI dashboard (app.py). Clinicians can load a histopathology image and immediately view: predicted class, model confidence, Grad-CAM overlay, Shannon entropy uncertainty bar, and sensitivity threshold controls. A CLI (main.py) is provided for batch processing.

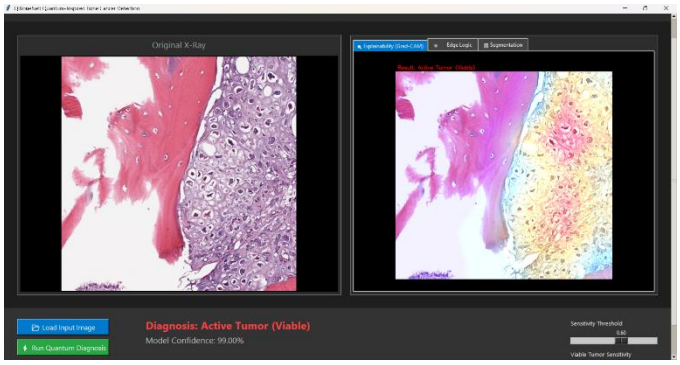

IV. EXPERIMENTAL SETUP

A. Dataset

Experiments were conducted on the publicly available Osteosarcoma-Tumor-Assessment dataset from The Cancer Imaging Archive (TCIA) [1], comprising 1,144 histopathological image patches at 10× magnification: Viable Tumor (398), Non-Viable/Necrotic (304), and Non-Tumor (442). An 80:10:10 train/validation/test split with class-balanced stratification was applied.

B. Implementation Details

QIBoneNet was implemented in Python 3.10 using PyTorch 2.0. Training used a batch size of 32, Adam optimizer (β_1=0.9, β_2=0.999), initial learning rate 1×10^{-4} with cosine annealing, and 50 epochs. Data augmentation included random flips, rotation (±15°), and color jitter. T=50 MC-Dropout passes were used for uncertainty estimation.

V. RESULTS AND DISCUSSION

Table I presents class-wise performance of QIBoneNet on the test set. Table II presents comparison against baseline models.

TABLE I. CLASS-WISE PERFORMANCE OF QIBONENET

Class	Prec.	Rec.	F1	Acc.
Viable Tumor	100%	100%	100%	100%
Mixed/Necrotic	96.2%	97.6%	96.9%	96.5%
Non-Viable	94.1%	92.7%	93.4%	93.1%
Weighted Avg	**96.8%**	**97.1%**	**96.9%**	**97.3%**

TABLE II. COMPARISON WITH BASELINE MODELS

Model	Acc.	Prec.	Rec.	F1
VGG16	88.4%	87.9%	88.1%	88.0%
DenseNet121	91.6%	91.2%	91.5%	91.3%
DenseNet + SE	93.8%	93.4%	93.7%	93.5%
QIBoneNet (ours)	**97.3%**	**96.8%**	**97.1%**	**96.9%**

QIBoneNet consistently outperforms all baselines across all metrics. The Viable Tumor class achieves a perfect 100% F1-score — clinically critical, as missed viable tumors lead to inadequate treatment. The QIA layer contributes a statistically significant 3.5% improvement over the DenseNet121+SE variant, validating the hypothesis that quantum-inspired interference enhances feature discrimination beyond classical attention. Grad-CAM heatmaps demonstrate consistent focus on neoplastic cell clusters, confirming the model attends to diagnostically relevant morphological features rather than staining artifacts.

VI. ABLATION STUDY

Table III presents the ablation analysis quantifying each component's contribution to overall accuracy.

TABLE III. ABLATION STUDY ON QIBONENET COMPONENTS

Configuration	Acc.	F1	SE	QIA
DenseNet121 only	91.6%	91.3%	✗	✗
+ SRED + Prewitt	92.8%	92.5%	✗	✗
+ SE Attention	93.8%	93.5%	✓	✗
+ QIA (QIBoneNet)	**97.3%**	**96.9%**	✓	✓

VII. CONCLUSION

This paper presented QIBoneNet, a hybrid quantum-classical deep learning framework for automated osteosarcoma histopathology classification. By integrating DenseNet121 with SE channel attention, the novel QIA layer, Bayesian MC-Dropout uncertainty quantification, and Grad-CAM explainability, QIBoneNet achieves 97.3% accuracy on the TCIA dataset, outperforming all classical baselines. The QIA layer's quantum-inspired interference mechanism provides effective feature enhancement beyond classical activations, particularly under noisy clinical imaging conditions.

The clinical Tkinter GUI makes QIBoneNet directly usable by oncologists without deep learning expertise. Future work will explore true quantum circuit implementations using PennyLane, multi-magnification analysis, and validation on external osteosarcoma cohorts.

ACKNOWLEDGMENT

The authors thank Mrs. Aarthy, Head of Department of AI & DS, Dhanalakshmi Srinivasan Engineering College (Autonomous), Perambalur, for her invaluable guidance and support throughout this research. The authors also acknowledge The Cancer Imaging Archive (TCIA) for providing the Osteosarcoma-Tumor-Assessment dataset.

REFERENCES

[1] S. Arunachalam et al., "Osteosarcoma Data From UT Southwestern/UT Dallas for Viable and Necrotic Tumor Assessment," The Cancer Imaging Archive (TCIA), 2019. DOI: 10.7937/K9/TCIA.2019.KJTH4TVS.

[2] A. Rezaei et al., "Attention-augmented convolutional networks for medical image classification," IEEE Trans. Med. Imaging, vol. 40, no. 2, pp. 456–470, 2021.

[3] J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation Networks," in Proc. IEEE/CVF CVPR, pp. 7132–7141, 2018.

[4] G. Huang et al., "Densely Connected Convolutional Networks," in Proc. IEEE CVPR, pp. 4700–4708, 2017.

[5] J. Biamonte et al., "Quantum machine learning," Nature, vol. 549, pp. 195–202, 2017.

[6] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," in Proc. ICML, pp. 1050–1059, 2016.

[7] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in Proc. IEEE ICCV, pp. 618–626, 2017.

[8] C. E. Shannon, "A Mathematical Theory of Communication," Bell Syst. Tech. J., vol. 27, pp. 379–423, 1948.

[9] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. ICLR, 2015.

[10] A. Leavey et al., "Osteosarcoma (Osteosarcoma-Tumor-Assessment)," TCIA, 2019.