QCNN-MFND: A Novel Quantum CNN Framework for Multimodal Fake News Detection in Social Media

Arya Suneesh
IIIT Kottayam
Kottayam, India
arya21bcs5@iiitkottayam.ac.in

Balasubramanian P
IIIT Kottayam
Kottayam, India
pbala@iiitkottayam.ac.in

Abstract

Fake news on social media platforms poses significant threats to public trust and information integrity. This research explores the application of quantum machine learning (QML) techniques for detecting fake news by leveraging quantum computing's unique capabilities. Our work introduces a hybrid quantum-classical framework that utilizes quantum convolutional neural networks (QCNNs) with angle and amplitude encoding schemes for processing multimodal features from text and images. Experiments conducted on benchmark datasets - GossipCop and Politifact - demonstrate that our quantum-enhanced model achieves superior performance compared to classical approaches, with accuracy rates of 88.52% and 85.58%, and F1 scores of 93.19% and 90.20% respectively. Our findings establish QML as a viable approach for addressing the challenges of fake news detection in the digital era.

1 Introduction

The proliferation of misinformation on social media threatens information integrity and societal welfare. Current machine learning and deep learning models struggle with accurate fake news identification due to insufficient feature extraction. Effective FND models must integrate textual and visual cues to distinguish between real and fake news, but conventional algorithms struggle to capture the subtle complexities of multi-modal data. We explore quantum machine learning as a promising alternative, focusing on quantum convolutional neural networks (QCNNs). Our research aims to develop a novel FND system that leverages QML techniques to enhance precision and robustness in fake news detection while maintaining computational efficiency.

2 Related Work

2.1 Unimodal FND Methods

Unimodal techniques focus either on textual or visual elements to categorize the news into fake or real.

2.1.1 ML-based FND Methods

Various studies have employed machine learning (ML) techniques for FND (Mishra and Sadia, 2023). Verma et al. (2021) utilized Support Vector Machine (SVM) for feature extraction from news articles but lacked deep learning (DL) models. Ozbay and Alatas (2020) adopted Decision Trees but faced accuracy challenges due to reliance on word count-based features. Esteban-Bravo et al. (2024) investigated early prediction of fake news virality using non-parametric models like Random Forest and Support Vector Classifier (SVC).

2.1.2 DL-based FND Methods

Rai et al. (2022) integrated BERT with LSTM, improving FND but suffered from low accuracy attributed to inadequate contextual features. Chen et al. (2024) tackled linguistic differences between Cantonese and Mandarin with a Deep semantic-aware graph convolutional network (SA-GCN) and CantoneseBERT on the Cantonese rumour dataset. Bazmi et al. (2023) emphasized the role of users' socio-cognitive biases and partisan bias with the Multi-View Co-Attention Network (MVCAN) but overlooked the influence of political viewpoints and credibility assessments of users.

2.1.3 QML-based FND Methods

Quantum machine learning (QML) techniques have shown promise in FND. Aishwarya et al. (2023) conducted a comprehensive review of Quantum Machine Learning techniques for FND. Their study revealed that QKNN, when integrated with Genetic and Evolutionary Feature Selection (GEFeS),

Proceedings of the QuantumNLP: Integrating Quantum Computing with Natural Language Processing, pages 44–52
November 24, 2025

achieved an impressive accuracy of 83.8%, surpassing the performance of conventional KNN algorithms. Tian and Baskiyar (2021) showcased the effectiveness of QKNN combined with Genetic and Evolutionary Feature Selection.

2.2 Multimodal FND Methods

Multimodal FND methods integrate both textual and visual features for detection.

2.2.1 DL-based FND Methods

Raja et al. (2024) proposed Dilated Temporal CNNs (DTCN), BiLSTM, and Contextualized Attention Mechanism (CAM), achieving impressive accuracy of 93.97% on the Dravidian_Fake dataset. Singhal et al. (2020) employed SpotFake+ but faced issues with prolonged training time and information loss from VGG-19's pooling layer. Kaliyar et al. (2021) employed Feed Forward Neural Networks with multiple CNN channels for local sequential feature extraction, yet generalization ability remains unexplored. Singh et al. (2023) employed multimodal learning techniques with NasNet Mobile for image analysis and BERT+ELECTRA for text processing, achieving 85% accuracy on the Twitter MediaEval Dataset and Weibo Corpus.

2.2.2 QML-based FND Methods

Qu et al. (2024) proposed QMFND, a quantum multimodal fusion-based model designed specifically for FND on social media platforms. By employing quantum encoding and quantum convolutional neural networks (QCNNs), QMFND achieved notable accuracies of 87.9% and 84.6% on the Gossipcop and Politifact datasets, respectively. However, the performance of QMFND is subject to limitations imposed by current hardware constraints and significant background noise in the operating environment of quantum computers.

3 Preliminaries

3.1 Pre-trained Language Models

Pre-trained language models form the basis for extracting representations from news text, using transformer architectures to capture contextual relationships. **BERT** processes text bidirectionally through masked language modeling, predicting randomly masked tokens from surrounding context. **XLNet** employs permutation-based autoregressive pre-training, capturing bidirectional context without relying on `[MASK]` tokens by considering all sequence permutations. **DistilBERT** is a compressed version of BERT that retains 97% of its language understanding while running 60% faster through knowledge distillation from a larger teacher model.

3.2 Pre-trained Image Models

Pre-trained convolutional neural networks extract visual features from images. These models learn hierarchical representations through successive convolutional layers.

VGG architectures (VGG16 and VGG19) utilize small 3×3 convolutional filters throughout the network. They stack multiple convolutional layers before pooling operations. This design enables learning complex features while maintaining computational efficiency.

ResNet50 introduces residual connections to address vanishing gradient problems. Skip connections allow gradients to flow directly through shortcuts. The architecture consists of 50 layers organized into residual blocks. Each block contains convolutional layers with identity mappings.

EfficientNet applies compound scaling to balance network depth, width, and resolution. It uses mobile inverted bottleneck blocks (MBConv) as building components. Squeeze-and-excitation optimization improves channel interdependencies. This architecture achieves superior accuracy with fewer parameters.

3.3 Understanding Quantum Mechanisms

Quantum computing transcends classical computing principles, offering the potential for unprecedented computational power and efficiency. One fundamental aspect of quantum computing is quantum encoding, a technique that transforms classical information into quantum states, enabling it to be processed and manipulated by quantum algorithms.

Quantum encoding transforms classical data into quantum states, exploiting superposition and entanglement to exponentially increase information density and computational capabilities beyond classical methods.

Several encoding approaches exist, each with distinct advantages LaRose and Coyle (2020):

Angle Encoding represents data through rotational angles of quantum gates. Parameterized rotation operations encode information directly into angular parameters. This method offers simplicity and hardware efficiency for near-term quantum devices.

$$|\psi\rangle = \bigotimes_{i=1}^{n} R_Y(x_i)|0\rangle = \bigotimes_{i=1}^{n} \begin{pmatrix} \cos(x_i/2) \\ \sin(x_i/2) \end{pmatrix} \quad (1)$$

where $R_Y(x_i)$ represents a Y-axis rotation gate parameterized by the data value x_i.

Amplitude Encoding represents data through relative amplitudes of quantum states. This scheme leverages superposition, enabling multiple information pieces to coexist simultaneously within a quantum system. For a normalized classical data vector $\mathbf{x}$, amplitude encoding creates:

$$|\psi\rangle = \frac{1}{\|\mathbf{x}\|} \sum_{i=0}^{2^n-1} x_i|i\rangle \quad (2)$$

where $\|\mathbf{x}\|$ is the normalization factor ensuring $\langle\psi|\psi\rangle = 1$, and $|i\rangle$ represents the computational basis states.

4 Proposed Methodology

4.1 Training and Evaluation Framework for Hybrid Model

Our fake news detection system, described in figure 1 combines classical deep learning with quantum computing. The hybrid architecture processes text and visual content through separate pathways before quantum integration.

In data preparation, news articles pass through a text transformer for contextual embeddings, while images are processed via CNN for visual feature extraction. A MultiHeadCrossAttention mechanism creates connections between text and image modalities, helping identify mismatches that often signal deception.

Training follows an epoch-based approach with early stopping to prevent overfitting. Data batches move through classical pathways, get fused, and pass to the QCNN (Cong et al., 2019) component, which leverages quantum principles like superposition and entanglement. We selected QCNN over Q-RNN or Q-LSTM because CNNs naturally preserve spatial locality in quantum circuits, essential for capturing hierarchical patterns in multimodal data. The convolutional structure aligns with quantum gate locality constraints on NISQ devices. We initially used cross-entropy loss before switching to focal loss due to dataset imbalance, and implemented gradient clipping for stability.

The validation process runs after each epoch, computing accuracy, precision, and recall. An early stopping mechanism halts training after three epochs without improvement in validation loss.

Final evaluation includes standard classification metrics and threshold optimization to identify the optimal decision boundary between real and fake news, producing metrics using both default (0.5) and optimized thresholds.

This pipeline balances classical deep learning's strength in feature extraction with quantum computing's advantages in modeling complex relationships.

5 Experimental Settings

5.1 Setup

All experiments were conducted on the Kaggle platform using an NVIDIA Tesla P100 GPU (16 GB VRAM) with 13GB RAM. The models were implemented using PyTorch 2.0 and trained with CUDA 12 acceleration. For transformer-based language models, we utilized the Hugging Face Transformers library. Image processing was handled with torchvision and quantum circuit simulations were executed using Pennylane with PyTorch interface.

The datasets were preprocessed using standard NLP techniques for textual data, including tokenization, normalization, and sequence padding. For image data, we employed standard preprocessing pipelines with resizing to 224×224 pixels, normalization using ImageNet statistics (mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225]), and augmentation techniques including random horizontal flips and color jitter during training.

5.2 Dataset Analysis

The experiments were conducted on two benchmark fake news datasets: Gossipcop and Politifact. As described in Table 1, the dataset statistics reveal several notable characteristics. A significant class imbalance exists in both GossipCop and Politifact datasets, with real news consistently outnumbering fake news. The GossipCop (GC) dataset maintains approximately an 80-20 split between real and fake news in both train and test sets. The Politifact (PF) dataset shows a different ratio, with approximately 65-35 split in the training set shifting to 72-28 in the test set.

Text length analysis exposes distinct patterns between the two sources: Politifact articles are generally longer, with mean lengths of 8,919 and 9,494 characters for train and test sets respectively, compared to GossipCop's shorter average of around

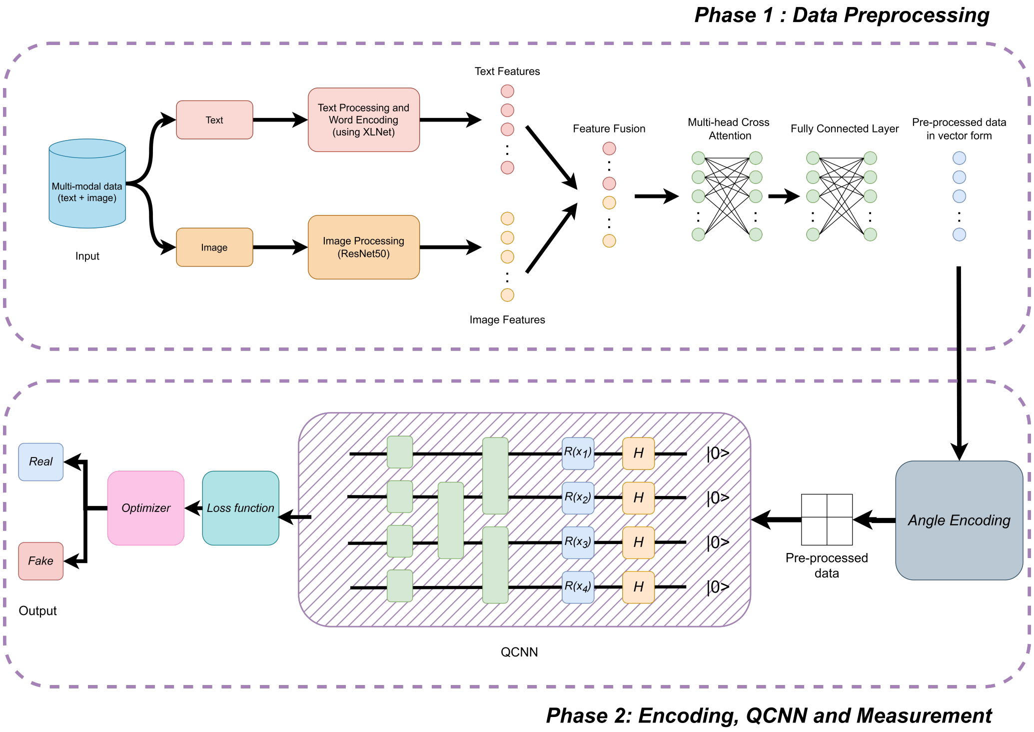

Figure 1: Overview of Training Process of QCNN-MFND

Table 1: Dataset Statistics Comparison

Metric	GC Train	GC Test	PF Train	PF Test
Total samples	10,010	2,830	381	104
Real	7,974 (79.7%)	2,285 (80.7%)	246 (64.6%)	75 (72.1%)
Fake	2,036 (20.3%)	545 (19.3%)	135 (35.4%)	29 (27.9%)
Mean text length	3,427.5	3,460.8	8,919.2	9,494.2
Std dev	5,872.6	6,433.2	17,501.6	18,349.9
Min	34.0	57.0	42.0	45.0
Median	2,072.0	2,046.5	2,511.0	2,966.5
Max	100,096.0	100,055.0	100,155.0	100,077.0

3,400 characters.

5.3 Evaluation Metrics

Due to class imbalance, we employed multiple standard metrics for binary classification problems:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{3}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{4}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{5}$$

$$\text{F1 Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{6}$$

where TP, TN, FP, FN represent true positives, true negatives, false positives, and false negatives respectively.

6 Results and Analysis

6.1 Textual Feature Analysis

We evaluated multiple transformer-based language models for textual feature extraction. Tables 2 and 3 present the performance metrics across both datasets.

XLNet achieved the highest accuracy (0.876) on the GossipCop dataset, while DistilBERT demonstrated superior performance on Politifact with the highest accuracy (0.9135) and F1 score (0.9379). These results highlight the effectiveness of transformer-based models for fake news detec-

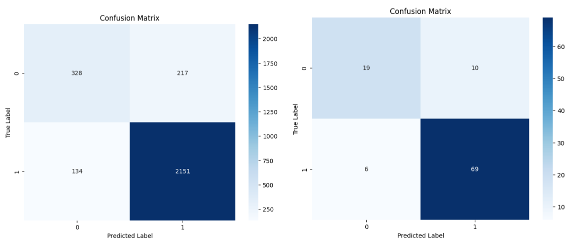

(a) GossipCop dataset (b) Politifact dataset

Figure 2: Confusion matrices for XLNet on both datasets.

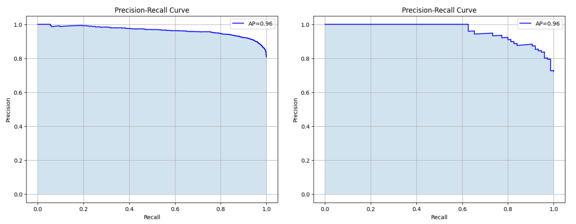

(a) GossipCop dataset (b) Politifact dataset

Figure 3: Precision-Recall curves for XLNet on both datasets.

tion, with different architectures exhibiting distinct strengths across different news domains.

Fig. 2 shows the confusion matrices for XLNet performance on both datasets, demonstrating strong classification performance with minimal false negatives. Fig. 3 presents the precision-recall curves, indicating robust performance across different threshold values.

6.2 Visual Feature Analysis

We evaluated six prominent CNN architectures for visual feature extraction. Table 4 shows the performance comparison on the GossipCop dataset.

ResNet50 achieved the highest performance (79.93% accuracy, 0.7802 F1 score) among all CNN models. Modern architectures generally demonstrated better optimization with lower loss values compared to traditional VGG models.

Table 2: Performance of Transformer Models on GossipCop

Model	Acc.	Prec.	Rec.	F1	Loss
BERT	0.871	0.920	0.920	0.920	0.747
RoBERTa	0.874	0.913	0.933	0.923	0.643
MPNet	0.872	0.926	0.915	0.920	**0.605**
DistilBERT	0.869	0.913	0.926	0.919	0.691
XLNet	**0.876**	0.908	**0.941**	**0.925**	0.707

Table 3: Performance of Transformer Models on Politifact

Model	Acc.	Prec.	Rec.	F1	Loss
BERT	0.846	0.873	0.920	0.896	**0.388**
RoBERTa	0.875	0.888	**0.947**	0.916	0.857
MPNet	0.837	0.914	0.853	0.883	0.456
DistilBERT	**0.914**	**0.971**	0.907	**0.938**	0.584
XLNet	0.846	0.873	0.920	0.896	0.585

Table 4: Performance of CNN Models on GossipCop

Model	Acc.	Prec.	Rec.	F1	Loss
VGG16	0.765	0.745	0.765	0.753	0.992
VGG19	0.786	0.759	0.786	0.769	**0.658**
ResNet50	**0.799**	**0.772**	**0.799**	**0.780**	0.974
EfficientNet	0.763	0.765	0.763	0.764	0.999
ViT	0.728	0.742	0.728	0.735	1.018
ConvNeXt	0.783	0.761	0.783	0.769	1.110

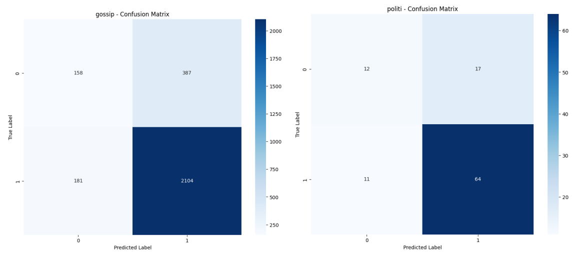

(a) GossipCop dataset (b) Politifact dataset

Figure 4: Confusion matrices for ResNet50 on both datasets.

Fig. 4 shows the confusion matrices for ResNet50 on both datasets, while Fig. 5 displays the corresponding precision-recall curves, demonstrating consistent performance across different news domains.

Table 5: Performance of CNN Models on Politifact

Model	Acc.	Prec.	Rec.	F1	Loss
VGG16	0.721	0.724	0.721	0.723	1.719
VGG19	0.721	0.730	0.721	0.725	1.623
ResNet50	**0.731**	0.715	**0.731**	0.720	0.923
EfficientNet	0.731	0.720	0.731	0.724	**0.744**
ViT	0.731	0.720	0.731	0.724	1.170
ConvNeXt	0.721	**0.744**	0.721	**0.729**	0.791

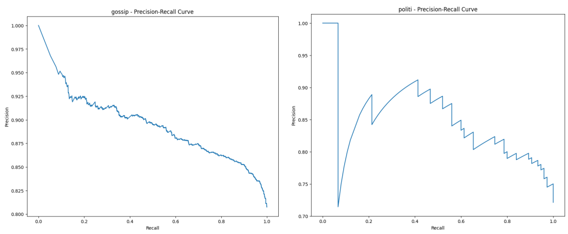

(a) GossipCop dataset (b) Politifact dataset

Figure 5: Precision-Recall curves for ResNet50 on both datasets.

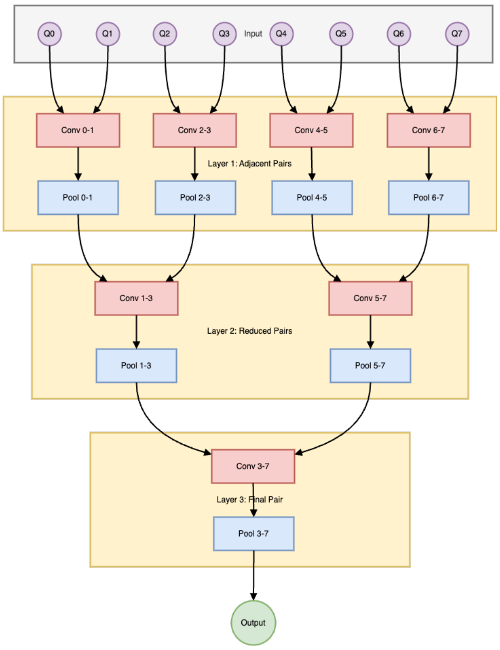

Figure 6: Architecture of the implemented QCNN showing the hierarchical quantum processing structure.

6.3 Quantum Convolutional Neural Network Implementation

The implemented QCNN architecture consists of three primary components: quantum convolution layers, quantum pooling layers, and a measurement layer. The network operates on 8 qubits and implements a hierarchical structure with multiple conv-pool operations at different scales. We selected 8 qubits as a balance between expressivity and current NISQ device limitations, aligning with typical quantum hardware availability.

Fig. 6 illustrates our QCNN architecture, while Fig. 7 details the convolution and pooling layer operations, demonstrating the quantum gate operations used for feature extraction and compression.

Each convolution operation implements initial RY rotations, CNOT entanglement, controlled-RX rotation, and final RZ rotations. Each pooling operation uses parameterized rotations and CNOT gates

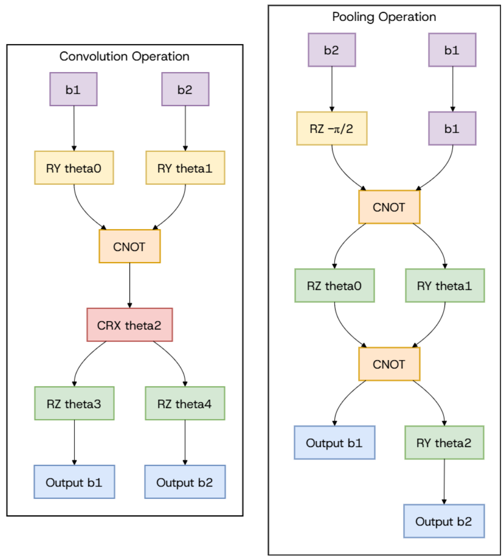

Figure 7: Detailed view of convolution and pooling layers in the QCNN, illustrating quantum gate operations.

to compress quantum information while preserving relevant features.

6.4 Experimental Setup

6.4.1 Model Architecture

A hybrid quantum-classical model was implemented with five key components. The image pathway uses pretrained CNNs (VGG16, EfficientNet, ResNet50) to extract features, projecting them to lower dimensions ($qbits/2$) for fusion. The text pathway processes input through XLNet, with mean-pooled features projected to $qbits/2$ dimensions. MultiHeadCrossAttention aligns image features with text context. The fusion component concatenates features and compresses them via a linear layer to $qbits$ dimensions. Finally, the QCNN processes the fused features to produce class probabilities.

6.4.2 Training Protocol

Loss Functions: Two loss functions addressed class imbalance:

Cross-Entropy Loss (CE) with class weighting and label smoothing:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N}\sum_{i=1}^{N} w_{y_i}\Big[y_i \log(p_i) + (1-y_i)\log(1-p_i)\Big] + \lambda\|\theta\|_2^2 \quad (7)$$

where w_{y_i} is class weight (inverse frequency), p_i is predicted probability, and $\lambda = 0.1$.

Focal Loss (FL) down-weights easy examples:

$$\mathcal{L}_{\text{FL}} = -\frac{1}{N}\sum_{i=1}^{N}\alpha_{y_i}(1-p_i)^{\gamma}\Big[y_i\log(p_i) + (1-y_i)\log(1-p_i)\Big] \tag{8}$$

where $\alpha_{y_i} = \frac{\text{\# minority class}}{\text{\# total samples}}$ and $\gamma = 2.0$.

Training Setup: AdamW optimizer (lr=2×10^{-4}), OneCycleLR scheduler, batch size 32, 25 epochs with early stopping (patience=3).

6.5 Performance Across Model Configurations

We evaluated multiple combinations of text feature extractors (XLNet), image feature extractors (VGG16, EfficientNet, ResNet50), and quantum encoding methods (Angle, Amplitude) on both datasets. Table 6 summarizes the performance metrics for each configuration.

On the GossipCop dataset, ResNet50 + XLNet with angle encoding trained with focal loss achieved the highest performance (88.52% accuracy, 93.19% F1 score). Angle encoding consistently outperformed amplitude encoding when paired with the same image feature extractor.

On the Politifact dataset, EfficientNet+XLNet with angle encoding and focal loss achieved the best results (88.46% accuracy, 92.31% F1 score). Focal loss significantly improved performance across both datasets, particularly evident in recall performance.

6.6 Classical versus Quantum Models

Table 7 presents the comparison between traditional classical approach (XLNet+ResNet50) and our proposed quantum model (QCNN-MFND).

The comparative analysis reveals our hybrid model outperforms classical approaches on the GossipCop dataset, with a remarkable 65% reduction in false negatives, critical for minimizing missed fake news instances. While the classical approach performs marginally better on Politifact's smaller dataset, this suggests our quantum model requires larger datasets to fully optimize its parameters.

7 Conclusion

We successfully developed QCNN-MFND, a novel framework leveraging quantum computing principles for fake news detection on social media. By combining QML with deep learning approaches, our architecture integrates XLNet for text analysis, ResNet50 for visual features, and quantum convolutional neural networks for multimodal fusion. The framework achieves impressive results—**88.52% accuracy and 93.19% F1 score on GossipCop**, and **85.58% accuracy with 90.20% F1 score on Politifact** - demonstrating significant advantages in minimizing missed fake news instances. Our experiments reveal that quantum computing offers particular benefits for larger datasets, providing a balanced precision-recall trade-off that prioritizes false negative reduction. This advancement represents an important step toward creating more trustworthy information ecosystems, with potential for greater improvements as quantum computing technology continues to mature.

Future directions include building a web application to enable real-time detection, further QCNN architectural exploration, and explainable quantum models for fake news detection.

Limitations

Several limitations merit consideration. We simulated quantum circuits on classical hardware; real quantum devices introduce noise and hardware constraints not captured in our experiments. Our 8-qubit architecture faces deployment challenges on current NISQ devices.

The datasets present additional constraints. The Politifact dataset's small size (381 training samples) limits model learning capacity. Both datasets focus exclusively on English-language social media news, leaving cross-domain and multilingual generalization untested. The significant class imbalance (80-20 and 65-35 splits) affects detection performance despite focal loss mitigation.

Our evaluation scope remains limited to two social media datasets. Temporal robustness, adversarial testing, and real-time inference performance remain unexplored. Training requires high-end GPU resources, and deployment costs on actual quantum hardware are quite high.

References

C. Aishwarya, Manikanadan Venkatesan, and P. Prabhavathy. 2023. Research oriented reviewing of quantum machine learning. *International Research Journal on Advanced Science Hub*, 5:165–177.

Parisa Bazmi, Masoud Asadpour, and Azadeh Shakery. 2023. Multi-view co-attention network for fake

Table 6: Experiments using proposed model QCNN-MFND with various configurations

Dataset	Text	Image	Encoding	Acc.	Prec.	Rec.	F1
GossipCop	XLNet	VGG16	Angle	0.877	0.902	0.951	0.926
		EfficientNet		0.881	0.906	0.952	0.928
		EfficientNet*		0.882	0.899	0.963	0.930
		ResNet50*		**0.885**	0.894	**0.973**	**0.932**
		VGG16	Amplitude	0.877	0.899	0.955	0.926
		EfficientNet		0.875	**0.916**	0.931	0.923
		ResNet50*		0.884	0.902	0.962	0.931
Politifact	XLNet	VGG16	Angle	0.846	0.883	0.907	0.895
		EfficientNet		0.837	0.837	0.960	0.894
		EfficientNet*		**0.885**	0.889	**0.960**	**0.923**
		ResNet50*		0.856	0.885	0.920	0.902
		VGG16	Amplitude	0.875	0.897	0.933	0.915
		EfficientNet		0.875	**0.908**	0.920	0.914
		ResNet50*		0.769	0.823	0.867	0.844

* Models trained with focal loss criterion.

Table 7: Performance Comparison Between Classical and Proposed Quantum Model

Dataset	Metric	Classical (without QCNN)	Proposed (with QCNN)
GossipCop	Accuracy	87.39%	**88.52%**
	F1 Score	0.922	**0.932**
	Precision	**0.918**	0.894
	Recall	0.926	**0.973**
Politifact	Accuracy	**87.50%**	85.58%
	F1 Score	**0.914**	0.902
	Precision	**0.908**	0.885
	Recall	0.920	0.920

news detection by modeling topic-specific user and news source credibility. *Information Processing & Management*, 60(1):103146.

Xinyu Chen, Yifei Jian, Liang Ke, Yunxiang Qiu, Xingshu Chen, Yunya Song, and Haizhou Wang. 2024. A deep semantic-aware approach for cantonese rumor detection in social networks with graph convolutional network. *Expert Systems with Applications*, 245:123007.

Iris Cong, Soonwon Choi, and Mikhail Lukin. 2019. Quantum convolutional neural networks. *Nature Physics*, 15:1–6.

Mercedes Esteban-Bravo, Lisbeth d. l. M. Jiménez-Rubido, and Jose M. Vidal-Sanz. 2024. Predicting the virality of fake news at the early stage of dissemination. *Expert Systems with Applications*, 248:123390.

Rohit Kumar Kaliyar, Anurag Goswami, and Pratik Narang. 2021. Fakebert: Fake news detection in social media with a bert-based deep learning approach. *Multimedia Tools and Applications*, 80(8):11765–11788.

Ryan LaRose and Brian Coyle. 2020. Robust data encodings for quantum classifiers. *Physical Review A*, 102(3):032420.

Alok Mishra and Halima Sadia. 2023. A comprehensive analysis of fake news detection models: A systematic literature review and current challenges. *Engineering Proceedings*, 59(1):28.

Feyza Altunbey Ozbay and Bilal Alatas. 2020. Fake news detection within online social media using supervised artificial intelligence algorithms. *Physica A: Statistical Mechanics and its Applications*, 540:123174.

Zhiguo Qu, Yunyi Meng, Ghulam Muhammad, and Prayag Tiwari. 2024. Qmfnd: A quantum multimodal fusion-based fake news detection model for social media. *Information Fusion*, 104:102172.

Nishant Rai, Deepika Kumar, Naman Kaushik, Chandan Raj, and Ahad Ali. 2022. Fake news classification using transformer based enhanced lstm and bert. *International Journal of Cognitive Computing in Engineering*, 3:98–105.

Eduri Raja, Badal Soni, Candy Lalrempuii, and Samir Kumar Borgohain. 2024. An adaptive cyclical learning rate based hybrid model for dravidian fake news detection. *Expert Systems with Applications*, 241:122768.

Prabhav Singh, Ridam Srivastava, K.P.S. Rana, and Vineet Kumar. 2023. Semi-fnd: Stacked ensemble based multimodal inferencing framework for faster fake news detection. *Expert Systems with Applications*, 215:119302.

Shivangi Singhal, Anubha Kabra, Mohit Sharma, Rajiv Ratn Shah, Tanmoy Chakraborty, and Ponnurangam Kumaraguru. 2020. Spotfake+: A multimodal framework for fake news detection via transfer learning (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13915–13916.

Ziyan Tian and Sanjeev Baskiyar. 2021. Fake news detection: an application of quantum k-nearest neighbors. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–6.

Pawan Kumar Verma, Prateek Agrawal, Ivone Amorim, and Radu Prodan. 2021. Welfake: word embedding over linguistic features for fake news detection. *IEEE Transactions on Computational Social Systems*, 8(4):881–893.