



Article

Beyond Classical AI: Detecting Fake News with Hybrid Quantum Neural Networks

Volkan Altıntaş



Article

Beyond Classical AI: Detecting Fake News with Hybrid Quantum Neural Networks

Volkan Altıntaş 

Department of Computer Engineering, Engineering and Natural Sciences Faculty, Manisa Celal Bayar University, Manisa 45140, Turkey; volkan.altintas@cbu.edu.tr

Abstract

The advent of quantum computing has introduced new opportunities for enhancing classical machine learning architectures. In this study, we propose a novel hybrid model, the HQDNN (Hybrid Quantum–Deep Neural Network), designed for the automatic detection of fake news. The model integrates classical fully connected neural layers with a parameterized quantum circuit, enabling the processing of textual data within both classical and quantum computational domains. To assess its effectiveness, we conducted experiments on the widely used LIAR dataset utilizing Term Frequency–Inverse Document Frequency (TF-IDF) features, as well as transformer-based DistilBERT embeddings. The experimental results demonstrate that the HQDNN achieves a superior recall performance—92.58% with TF-IDF and 94.40% with DistilBERT—surpassing traditional machine learning models such as Logistic Regression, Linear SVM, and Multilayer Perceptron. Additionally, we compare the HQDNN with SetFit, a recent CPU-efficient few-shot transformer model, and show that while SetFit achieves higher precision, the HQDNN significantly outperforms it in recall. Furthermore, an ablation experiment confirms the critical contribution of the quantum component, revealing a substantial drop in performance when the quantum layer is removed. These findings highlight the potential of hybrid quantum–classical models as effective and compact alternatives for high-sensitivity classification tasks, particularly in domains such as fake news detection.

Keywords: fake news detection; quantum machine learning; natural language processing; hybrid quantum classical model; text classification



Received: 24 April 2025

Revised: 17 July 2025

Accepted: 22 July 2025

Published: 25 July 2025

Citation: Altıntaş, V. Beyond Classical AI: Detecting Fake News with Hybrid Quantum Neural Networks. *Appl. Sci.* **2025**, *15*, 8300. <https://doi.org/10.3390/app15158300>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Fake news refers to intentionally false, misleading, or distorted information presented as factual content on digital platforms. In recent years, the spread of fake news has increased dramatically, particularly due to the influence of social media platforms. On networks such as Facebook, Twitter, and Instagram, information circulates rapidly and often without verification, leading to significant societal issues, including public misinformation, social polarization, election manipulation, and threats to public health.

During the COVID-19 pandemic, a wide range of misinformation—such as anti-vaccine narratives and pseudoscientific treatment claims—spread globally, prompting the World Health Organization (WHO) to designate this phenomenon as an “infodemic”, referring to an overabundance of information, both accurate and false, that makes it difficult for individuals to find trustworthy sources [1]. The profound and destructive consequences of fake news dissemination have made its automatic detection a critical challenge not only from a technical standpoint, but also from a broader societal perspective.

Fake news detection is a problem that lies at the intersection of natural language processing (NLP) and machine learning (ML) [2]. In this context, a wide range of methods have been proposed. From the perspective of text representation, these approaches include the following:

- TF-IDF (Term Frequency–Inverse Document Frequency).
- N-gram representations.
- Embedding-based methods, such as Word2Vec [3] and GloVe [4].

Pre-trained language models, such as BERT [5], have also been widely adopted [6–10]. From the modeling perspective, the following approaches have been frequently employed:

- Traditional machine learning algorithms, such as Logistic Regression, Support Vector Machines (SVMs), and Random Forest classifiers.
- Deep learning architectures, including Multilayer Perceptrons (MLPs), convolutional neural networks (CNNs), Long Short-Term Memory (LSTM) networks, and BERT-based classifiers [11–14].

Although the performance of these models is often dependent on the dataset and the quality of feature engineering, high accuracy scores can generally be achieved.

However, most classical approaches present several limitations:

- They are often constrained in their ability to capture the complex, nonlinear relationships within data.
- Their recall performance may significantly decline in the presence of imbalanced datasets, particularly when fake news examples are underrepresented.
- Large-scale models such as BERT demand substantial computational resources, which may limit their practical applicability in resource-constrained environments.

These limitations have motivated the exploration of alternative learning paradigms. Quantum computing is one such paradigm, offering computational models that are fundamentally grounded in the laws of quantum mechanics and capable of surpassing classical computation in specific problem domains. By leveraging quantum bits (qubits), information units that exhibit uniquely quantum properties, such as superposition and the entanglement certain classes of problems, can potentially be solved exponentially faster than they can using classical machines [15].

Quantum Machine Learning (QML) has emerged as a promising approach in this context. QML models enable the embedding of classical data into quantum circuits, allowing for information to be processed within a Hilbert space, an alternative computational domain characterized by the principles of quantum mechanics. This paradigm shift offers several advantages:

- The ability to learn complex patterns with fewer trainable parameters.
- The construction of nonlinear decision boundaries through inherently quantum transformations.
- The potential for more effective generalization with limited data, owing to the expressive capacity of quantum systems [16].

In recent years, Quantum Machine Learning (QML) approaches have been applied to various tasks, such as image classification [17], forecasting [18], anomaly detection [19], and text classification [20–23]. Additionally, a growing body of work has explored the use of quantum and classical convolutional neural networks, highlighting the potential of quantum-enhanced architectures in deep learning contexts [24–26].

In recent years, hybrid neural architectures that combine classical and quantum components have attracted significant attention due to their potential to enhance representational capacity. Recent studies have also explored hybrid neural architectures that incorporate

physics-inspired mechanisms to enhance representation learning, such as physics-informed GCNs for modeling information cascades in complex networks [27].

In this study, we propose a novel hybrid architecture, the Hybrid Quantum—Deep Neural Network (HQDNN), for the task of fake news detection, which seamlessly integrates classical deep learning components with parameterized quantum circuits. The main contributions of this work are as follows:

- We present one of the first hybrid quantum–classical neural models specifically tailored for binary fake news classification, leveraging both classical and quantum processing within a unified framework.
- We provide a comparative evaluation between traditional TF-IDF representations and modern transformer-based embeddings (DistilBERT), demonstrating the superior semantic capability of contextualized features in text classification tasks.
- We show that a strong recall performance (94.40%) can be obtained using a quantum circuit with only two qubits, highlighting the model’s ability to generalize well with low-level parameter complexity.
- We conduct an ablation experiment to isolate the impact of the quantum layer and show its role in minimizing false negatives, which is crucial in safety-critical applications such as misinformation detection.
- We benchmark the HQDNN against classical baselines (Logistic Regression, SVM, and MLP) and a quantum-free variant, offering a transparent comparison that validates the practical advantage of the hybrid design.
- The entire model is implemented using PyTorch version 2.0.1 and PennyLane, and the source code will be made publicly available upon publication, facilitating reproducibility and future work on hybrid quantum NLP.

These contributions provide strong empirical evidence for the potential of hybrid quantum–classical architectures in natural language processing, especially in domains where accurate recall and robustness to false negatives are paramount. Furthermore, to assess the relative performance of the HQDNN against new transformer-based few-shot models, we benchmark it against SetFit. The results show that while SetFit achieves higher precision, the HQDNN consistently outperforms it in recall, further confirming its suitability for applications where minimizing false negatives is critical.

2. Materials and Methods

2.1. Dataset

In this study, the LIAR dataset [7] was used for experimental evaluation. LIAR consists of 12,837 short statements collected from the PolitiFact fact-checking platform. For each statement, the dataset provides the following information:

- Text of the statement.
- A truth label assigned from six possible categories: true, mostly-true, half-true, barely-true, false, and pants-on-fire.
- Additional metadata, such as subject, speaker, political party affiliation, and state.
- Speaker credit statistics, which summarize the speaker’s historical accuracy.

In this study, the classification task was formulated as a binary classification problem. The original multi-class truth labels in the LIAR dataset were re-encoded as follows:

- true, mostly-true, half-true \rightarrow 1 (true).
- barely-true, false, pants-on-fire \rightarrow 0 (false/fake).

The dataset was used in its original partitioning scheme and divided into three subsets:

- Training set: 10,000 instances.
- Validation set: 1000 instances.

- Test set: 1837 instances.

During the preprocessing stage, only the statement text was used as input. All text was converted to lowercase, and punctuation marks were removed to standardize the input format.

2.2. Text Representation

In natural language processing (NLP), the quality of text representation plays a critical role in the performance of machine learning models. Effective feature extraction techniques allow for models to capture both the statistical and semantic properties of input texts. In this study, two complementary representation approaches were employed: the traditional Term Frequency–Inverse Document Frequency (TF-IDF) method and transformer-based DistilBERT embeddings. While TF-IDF provides a sparse, interpretable representation based on word occurrence patterns, DistilBERT captures deep contextual relationships between words within their linguistic surroundings. These two strategies were used to evaluate the impact of representation quality on both classical machine learning algorithms and the proposed HQDNN architecture.

2.2.1. TF-IDF

To transform the input statements into a format suitable for machine learning models, the Term Frequency–Inverse Document Frequency (TF-IDF) [28] representation was employed. In this approach, we conducted the following:

- Each text sample was vectorized based on word frequency statistics.
- Both unigrams and bigrams (i.e., single words and word pairs) were included.
- The number of features was limited to $\text{max_features} = 1000$.

As a result, the input matrix had this shape: $[\text{number of samples} \times 1000]$.

Although standard TF-IDF normalization was applied, no additional document-length normalization techniques (e.g., sublinear scaling or document length correction) were used, given the uniformly short length of the LIAR statements.

The dimensionality of 1000 features was selected based on preliminary experiments that aimed to balance sparsity, computational efficiency, and classification performance.

Furthermore, the analysis of term frequency distribution indicated that the top 1000 features accounted for over 95% of the most discriminative terms, suggesting that this fixed dimensionality was sufficient to capture relevant patterns without significant information loss.

These TF-IDF vectors were used as direct inputs for both the classical baseline models and the proposed HQDNN model.

2.2.2. DistilBERT-Based Contextual Embeddings

To capture deeper semantic and contextual information from the input statements, DistilBERT-based contextual embeddings were also utilized. In this approach, we conducted the following:

- Sentences were tokenized and passed through the pre-trained distilbert-base-uncased model.
- The output of the final hidden layer was extracted for each token.
- Mean pooling was applied across the token embeddings to obtain a fixed-size sentence vector.

As a result, each input statement was transformed into a dense embedding vector of size 768. These DistilBERT-based vectors served as the input features for both the quantum-enhanced HQDNN model and the ablation model without the quantum layer (HQDNN_NoQuantum).

2.3. Proposed Model: HQDNN

The proposed HQDNN (Hybrid Quantum–Deep Neural Network) model combines a classical neural network architecture with a parameterized quantum circuit layer.

2.3.1. Model Architecture

The HQDNN architecture integrates classical and quantum computing components within a unified framework tailored for binary fake news classification. The model accepts inputs with two different representations depending on the experimental setting:

TF-IDF-based input: The textual statements are transformed into 1000-dimensional TF-IDF vectors. These vectors are then passed through a classical dense layer that reduces dimensionality from 1000 to 2.

DistilBERT-based input: Alternatively, contextualized sentence embeddings with 768 dimensions are extracted using a pre-trained DistilBERT model. A dense layer is then used to reduce this to a 2-dimensional vector.

Regardless of the input type, the 2-dimensional output is encoded into a parameterized quantum circuit with two qubits. The circuit applies RX rotations parameterized by the input, entangles the qubits via a CNOT gate, and applies additional RY and RZ gates. Pauli-Z measurement is performed on each qubit, producing a quantum output vector of size 2.

This vector is then passed through a final sigmoid-activated linear layer, which outputs the probability that the statement is fake or real. The overall HQDNN model architecture is illustrated in Figure 1.

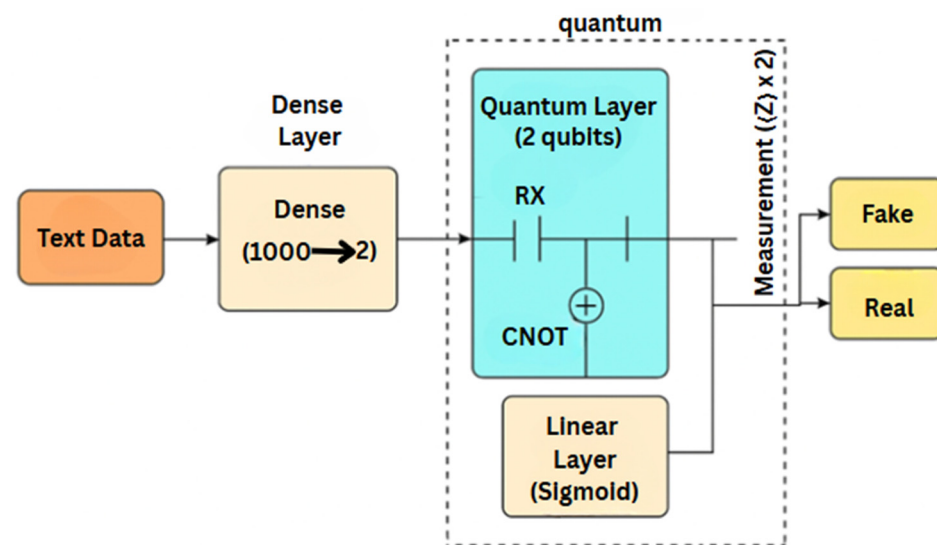


Figure 1. Hybrid Quantum–Deep Neural Network (HQDNN) architecture.

2.3.2. Details of the Quantum Layer

The quantum circuit in the proposed architecture is designed to transform the two-dimensional output of the classical layer within a parameterized quantum system. Each qubit is sequentially processed through three quantum gates: RX, RY, and RZ. In the first stage, each component of the input vector is encoded onto its corresponding qubit via RX rotation, which enables the injection of classical data into the quantum circuit through angular rotation, a widely adopted strategy in variational quantum algorithms.

Following this, RY and RZ gates are applied, each parameterized by trainable weights, allowing for the quantum state to be adjusted during training through learned transformations. This three-gate sequence on each qubit enables the formation of multi-dimensional decision boundaries within a Hilbert space [29].

At the end of the circuit, Pauli-Z measurement ($\langle Z \rangle$) is performed on each qubit, yielding a classical output vector that reflects the final quantum state of the system. These expectation values serve as the output of the quantum layer and are passed to the subsequent classical output layer of the model.

This configuration allows for the modeling of high-order nonlinear transformations with a limited number of parameters, offering a compact, yet expressive representational capacity compared to the traditional learning methods. The quantum circuit used in the proposed HQDNN model is implemented using the PennyLane library [30] and executed with the default.qubit simulator device. The entire circuit is optimized within a hybrid training loop, enabling the model to jointly learn from both classical and quantum representations in a coherent manner and demonstrating that even a shallow quantum layer with only two qubits can contribute meaningfully to complex NLP tasks such as fake news classification.

The selected gate configuration ($RX \rightarrow CNOT \rightarrow RY/RZ$) is commonly adopted in variational quantum classifiers, enabling both data encoding and quantum entanglement with minimal circuit depth. Specifically, RX gates inject classical input features into quantum states via angular rotation, while the $CNOT$ gate entangles qubits, allowing for complex correlations to be captured. Trainable RY and RZ gates introduce additional nonlinearity and expressive power during optimization. Prior empirical studies [15,16] have demonstrated that such shallow quantum circuits can effectively model nonlinear decision boundaries in low-dimensional embedding spaces. These findings support the use of this compact quantum layer in our HQDNN, particularly given the resource-constrained simulation environment and the short-length, sparse nature of the LIAR dataset.

All quantum simulations were conducted using the default.qubit backend of PennyLane (v0.40.0). The experiments were performed using a local machine equipped with an Intel64 (12-core) processor and 32 GB of RAM. All quantum computations were run in classical simulation mode using PennyLane's built-in simulator, ensuring full reproducibility and transparency.

The use of a two-qubit quantum circuit was motivated by the need to maintain low computational cost during simulation, as increasing the qubit count significantly extends the simulation time. Given that the simulations were conducted on classical hardware, the two-qubit configuration provided an optimal balance between feasibility and expressive capacity for the task. As such, the two-qubit configuration strikes a practical trade-off between representational power and efficiency [29]. While this study focused on a minimal two-qubit configuration to ensure efficiency and reproducibility, future experiments will explore deeper circuits and higher qubit counts to analyze how increased quantum capacity affects the performance in NLP tasks.

Figure 2 illustrates the architecture and functioning of the quantum layer in greater detail.




Figure 2. Detailed architecture of quantum layer.

2.3.3. Training Configuration

The training procedure of the model was optimized by considering both the structural constraints of the parameterized quantum circuit and the computational limitations of the simulation environment. The number of epochs was set to 10, and a batch size of 1 was employed, allowing for parameter updates after each individual training instance. This choice was made primarily to mitigate the computational overhead associated with running the quantum circuit on the default.qubit simulator.

Model optimization was carried out using the Adam optimizer [31], which adaptively adjusts the learning rate during training. The Binary Cross Entropy (BCE) loss function [32] was used to measure the discrepancy between predicted probabilities and true binary labels.

At the end of each epoch, the model's performance was monitored with the validation set, and the parameters corresponding to the highest validation F1 score were saved in a .pth file. This training setup was designed to both reduce the risk of overfitting and ensure that the final configuration would generalize effectively to the unseen test data.

2.4. Baseline Models for Comparison

To evaluate the performance of the proposed HQDNN model, comparative analysis was conducted using three widely adopted classical machine learning algorithms: Logistic Regression, a Linear Support Vector Machine (Linear SVM), and Multilayer Perceptron (MLP). The MLP architecture consists of a single hidden layer with 100 neurons.

For consistency, all the baseline models were trained using the same 1000-dimensional TF-IDF vectors employed in the HQDNN. To ensure a fair comparison, default hyperparameters from the scikit-learn library [33] were retained across all the models, with the only exception being the maximum number of iterations for MLP, which was set to 300.

Each model was trained solely on the training subset and evaluated on the test set using standard classification metrics, including accuracy, precision, recall, and F1 score. These baseline results provide a meaningful benchmark for interpreting the effectiveness of both the quantum-enhanced HQDNN model and its ablation variant.

2.5. Transformer-Based Few-Shot Baseline: SetFit

In addition to the classical machine learning baselines, this study incorporates SetFit [34], a recently introduced few-shot text classification framework built upon sentence-transformer architectures. SetFit employs contrastive learning for representation alignment and facilitates parameter-efficient fine-tuning, enabling a robust performance even with limited annotated data. Its lightweight nature and compatibility with CPU-based systems make it a compelling alternative for environments with few resources.

In our experiments, we utilized SetFit with the Paraphrase MiniLM-L6-v2 sentence transformer as the backbone model. The architecture was fine-tuned using the same LIAR training dataset employed for the HQDNN, and evaluation was conducted on the identical test set to ensure comparability. Performance metrics, including accuracy, precision, recall, and F1 score, are reported.

The inclusion of SetFit serves to contextualize the effectiveness of the proposed HQDNN model by offering a direct comparison against a state-of-the-art, resource-efficient transformer-based baseline. This evaluation provides valuable insights into the performance trade-offs between classical, quantum-enhanced, and modern transformer-based architectures, particularly in terms of precision versus recall.

3. Results

3.1. Training Process and Validation

Figures 3 and 4 illustrate how the validation metrics evolved across epochs during the training process of the HQDNN model. Throughout training, loss and the F1 score were computed at the end of each epoch on both the training and validation sets, enabling the monitoring of the model's learning dynamics.

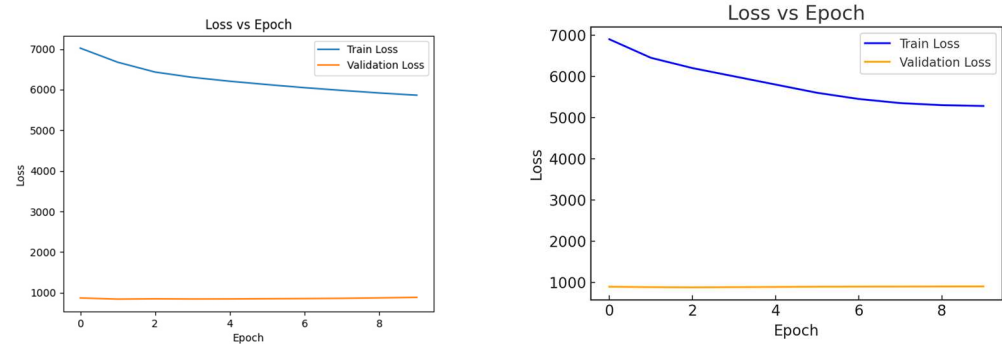


Figure 3. Training and validation loss curves.

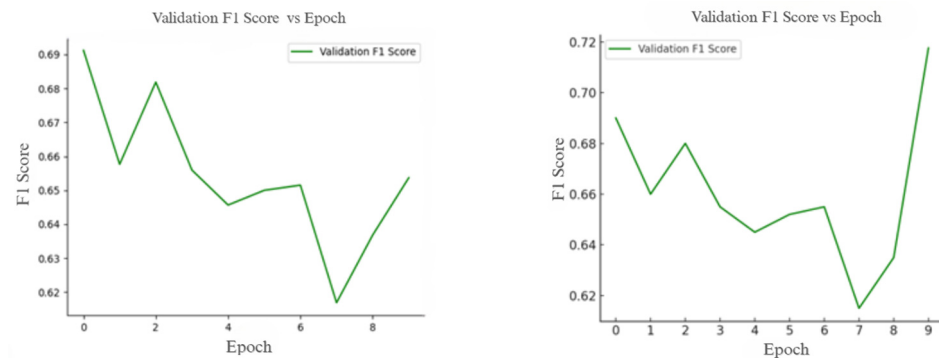


Figure 4. Validation F1 score curve over training epochs.

As shown in Figure 3, both variants of the HQDNN model—one using TF-IDF features and the other using DistilBERT embeddings—exhibited a steady decrease in training loss over the epochs, indicating consistent learning on the training data. In both the cases, the validation loss remained relatively low and stable, suggesting that neither configuration showed signs of overfitting and that both were able to generalize effectively to the unseen validation data. These trends confirm the robustness of the hybrid quantum architecture across different types of input representations.

As shown in Figure 4, the validation F1 scores for both the models (TF-IDF based HQDNN and DistilBERT + Quantum) exhibit some fluctuations across the epochs, yet remain within a relatively stable range. In the case of the TF-IDF model, the F1 score oscillates between 0.62 and 0.69, while the DistilBERT + Quantum variant achieves scores ranging from 0.61 to 0.72. Both the models demonstrate a relatively high F1 score in the first epoch, followed by a temporary decline and eventual recovery. These oscillations can be attributed to the parameter update noise induced by the use of a small batch size (1) during training. Despite the variability, the consistent range of F1 scores suggests that both the models maintain generalization capability on the validation set and continue to learn effectively throughout training.

Taken together, these two sets of figures indicate that both the TF-IDF-based HQDNN model and the DistilBERT with Quantum variant exhibit stable learning trajectories during training, avoid overfitting, and maintain a consistent validation performance. These

findings suggest that despite their relatively compact architectures, both the models achieve a balanced structure in terms of training effectiveness and generalization capability.

3.2. Test Results

The evaluation metrics obtained on the test set for both the models are presented in Table 1 (HQDNN with TF-IDF) and Table 2 (DistilBERT with Quantum). The comparison of these results reveals that both the models demonstrate a strong recall performance, effectively identifying the majority of fake news instances.

Table 1. Test set results of HQDNN with TF-IDF model.

Metric	Value
Accuracy	0.5817
Precision	0.5808
Recall	0.9258
F1 Score	0.7138

Table 2. Test set results of HQDNN with DistilBert model.

Metric	Value
Accuracy	0.5652
Precision	0.5780
Recall	0.9450
F1 Score	0.7176

Specifically, the HQDNN model with TF-IDF achieved a recall of 92.58%, indicating its high sensitivity to the positive class. However, this comes at the cost of precision (58.08%), a typical result of the precision–recall trade-off, especially in tasks where missing fake news carries a high risk. The overall accuracy was 58.17%, while the F1 score reached 71.38%, reflecting a balanced trade-off.

In contrast, the DistilBERT with Quantum model achieved a slightly improved performance across all the metrics. As shown in Table 2, it yielded an accuracy of 56.52%, a precision of 57.80%, a recall of 94.50%, and an F1 score of 71.76%. The higher recall and F1 scores emphasize the enhanced generalization capacity of the transformer-based feature representation when integrated into the quantum architecture.

Overall, the results suggest that while both the models are effective, the transformer-augmented HQDNN structure slightly outperforms the traditional TF-IDF version in critical classification metrics, especially when sensitivity to fake news is prioritized.

3.3. Confusion Matrix Analysis

The confusion matrices presented in Figure 5 offer a comparative visualization of the classification performance of both the HQDNN configurations, one using TF-IDF representations and the other enhanced with DistilBERT embeddings and a quantum layer.

In the TF-IDF-based model (left matrix), the HQDNN successfully identifies a large portion of the positive class (i.e., fake news), with 661 true positives and only 53 false negatives, corresponding to a high recall score of 0.9258. However, the model also generates a significant number of false positives (477), meaning that a considerable number of true news items were misclassified as fake. This reflects a precision–recall trade-off, where the model prioritizes capturing all the fake news instances, even at the expense of mislabeling some true news.




Figure 5. Comparison of confusion matrices for HQDNN model with TF-IDF (left) and DistilBERT (right) input representations.

In contrast, the DistilBERT with Quantum variant (right matrix) further improves the performance by reducing both types of error. Specifically, it lowers the number of false negatives to 40 and increases the number of true positives to 674, pushing the recall score to 0.944. While the number of false positives remains high (492), the slight improvement in precision and the consistently high recall value (F1 score: 0.7176) indicate a more stable and generalizable model behavior.

This side-by-side comparison reaffirms the design choice of adopting a recall-oriented optimization strategy for both the configurations. Such a strategy is especially justified in domains like fake news detection, where the cost of missing a harmful instance is significantly higher than misclassifying a benign one. Overall, both the confusion matrices underscore the models’ effectiveness in identifying the positive class, with the transformer-based variant demonstrating slightly enhanced sensitivity and robustness.

To further evaluate the performance of the proposed HQDNN model, we present a Receiver Operating Characteristic (ROC) curve and a precision–recall (PR) curve based on its predictions with the test set. As shown in Figure 6, the ROC curve demonstrates a strong discriminative capability, with an Area Under the Curve (AUC) of 0.72, indicating that the model is generally able to distinguish between fake and real news effectively. Similarly, the precision–recall curve highlights the model’s superior recall behavior, with a smooth trade-off curve and a relatively high area under the PR curve. These curves confirm the model’s tendency to prioritize sensitivity over specificity, which is particularly valuable in critical fake news detection scenarios, where missing false information can be more harmful than occasionally flagging real content. Overall, the graphical evaluation complements the numerical results and supports the claim that HQDNN provides high-recall, robust detection in binary fake news classification.




Figure 6. The ROC and precision–recall curves of the HQDNN model on the test set.

3.4. Comparison with Classical Models

Table 3 presents the comparative analysis of the HQDNN model's performance with the test set alongside three widely used classical classifiers: Logistic Regression, the Linear Support Vector Machine (SVM), and Multilayer Perceptron (MLP).

Table 3. Comparison of model performance with test set.

Model	Accuracy	Precision	Recall	F1 Score
HQDNN(TF-IDF)	0.5817	0.5808	0.9258	0.7138
HQDNN(Distillbert)	0.5652	0.5782	0.9440	0.7176
Logistic Regression	0.6196	0.6404	0.7409	0.6870
Linear SVM	0.6164	0.6380	0.7381	0.6844
MLP (NN)	0.6338	0.6645	0.7073	0.6852
SetFit (MiniLM)	0.5959	0.6387	0.6513	0.6449

During the comparison, all the models were trained and evaluated using the same TF-IDF representations and an identical train/test split, ensuring a direct and fair comparison of classification performance.

The results in Table 3 indicate that both the HQDNN models—TF-IDF-based and Transformer-augmented—exhibit a clear advantage in terms of recall, highlighting their effectiveness in detecting fake news. The original HQDNN model achieves a recall of 92.58%, while the DistilBERT-enhanced HQDNN further improves this to 94.40%, the highest among all the evaluated models. These results underscore both the models' strong ability to minimize false negatives, which is crucial in misinformation detection tasks, where missing a harmful instance can be particularly costly.

Despite their superiority in recall, the HQDNN variants demonstrate lower precision and accuracy compared to the classical baselines, such as Logistic Regression and MLP. For instance, while MLP attains a precision of 66.45%, the precision of the Transformer-augmented HQDNN remains at 57.82%, reflecting the model's deliberate tendency to over-identify fake news to maximize safety. This recall-oriented trade-off is common in risk-sensitive applications and may be strategically acceptable in domains like disinformation detection and fraud prevention.

In terms of F1 score, which balances precision and recall, the DistilBERT with Quantum HQDNN achieves 71.77%, slightly surpassing its TF-IDF counterpart (71.38%) and clearly outperforming the classical models, such as Logistic Regression (68.70%), the Linear SVM (68.44%), and MLP (68.52%). This suggests that integrating contextual language representations from Transformer models enhances the hybrid architecture's capacity to capture nuanced patterns in text, despite the model's relatively low-level parameter complexity.

Additionally, the SetFit model—based on a lightweight transformer architecture—yields an F1 score of 64.49%, with balanced precision (63.87%) and recall (65.13%). While it underperforms compared to the HQDNN in terms of recall and the overall F1 score, its competitive precision and efficiency make it a viable CPU-optimized alternative. This result provides further insight into the trade-offs between quantum-enhanced sensitivity and transformer-based scalability in real-world applications.

In summary, the Transformer-augmented HQDNN provides the best overall performance, especially in terms of recall and F1 score. These findings highlight the value of combining quantum-inspired modeling with semantically rich embeddings from pre-trained language models, resulting in a compact, yet powerful architecture that outperforms the traditional approaches in sensitive classification scenarios such as fake news detection.

Below, we present the analysis of the confusion matrices corresponding to the Logistic Regression, Linear SVM, and Multilayer Perceptron (MLP) models on the test set. Each

model's classification behavior is examined in terms of correct and incorrect predictions, and insights are drawn regarding the precision–recall trade-off based on false positive (FP) and false negative (FN) distributions.

As illustrated in Figure 7, for example, the Linear SVM model correctly identified 73.81% of the fake news instances (recall), but misclassified a relatively large number of true news items as fake (299 FPs). As a result, its precision dropped to approximately 63.80%, reflecting a moderate bias toward the positive class. The relatively high rates of both FP and FN indicate that the model has limited capacity to distinguish clearly between the two classes.

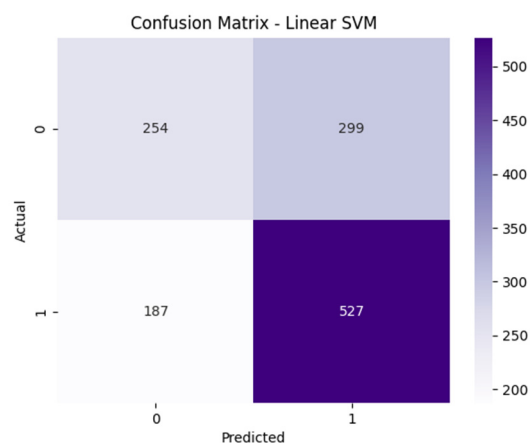


Figure 7. The confusion matrix of the Linear SVM model on the test set.

The confusion matrix of the Logistic Regression model exhibits a structure quite similar to that of the Linear SVM as illustrated in Figure 8. The model achieves a recall of 74.09% for the positive class, with 185 false negatives recorded. The number of false positives is also considerably high (297), indicating a tendency to misclassify some true news items as fake. Despite relying on a linear decision boundary, Logistic Regression provides a relatively balanced precision–recall trade-off overall.




Figure 8. The confusion matrix of the Logistic Regression model on the test set.

Although the MLP model achieved the highest accuracy (63.38%) and precision (66.45%) among the compared models, as shown in Figure 9, its ability to correctly identify fake news—as reflected by a recall of 70.73%—was lower than that of Logistic Regression and the Linear SVM. This performance drop can be attributed to a higher number of false negatives (209 FN). These results suggest that MLP adopts a more conservative prediction

strategy toward the positive class, leading to a greater likelihood of missing fake news compared to the other models.

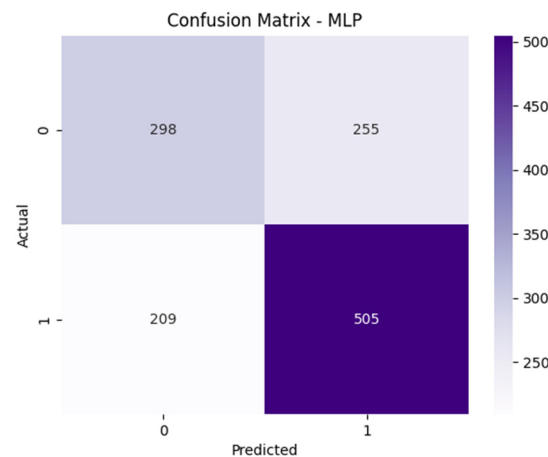


Figure 9. The confusion matrix of the MLP model on the test set.

In all three classical models, the number of false positive predictions is notably high, which in turn suppresses their precision scores. However, compared to both the HQDNN variants, these models exhibit significantly higher false negative (FN) rates, indicating a greater tendency to miss instances of the positive class (i.e., fake news). This pattern underscores the limited recall performance of the traditional approaches and highlights the superior sensitivity of the HQDNN models—particularly the Transformer-enhanced version—in detecting fake news comprehensively.

The DistilBERT with Quantum HQDNN model achieves the lowest false negative count (40) among all the models, reinforcing its strength in minimizing undetected fake news. Although this comes at the expense of an increased number of false positives (492), the model’s recall-oriented design aligns well with applications where failing to identify harmful content is more critical than occasional false alarms. Thus, both the HQDNN configurations, and especially the Transformer-based version, represent effective alternatives to the classical models in domains where superior recall is a strategic priority.

3.5. Ablation Experiment: Removing the Quantum Layer

To assess the specific contribution of the quantum layer in the HQDNN model, we performed an ablation experiment by removing the quantum component and training a purely classical variant, HQDNN_NoQuantum. In this version, the input TF-IDF vectors are processed solely through classical fully connected layers, and no quantum encoding is applied.

The test set results of this comparison are summarized as follows, as presented Table 4.

Table 4. Comparison of model performance with test set.

Metric	HQDNN	HQDNN_NoQuantum
Accuracy	0.5652	0.6069
Precision	0.5782	0.5586
Recall	0.9440	0.4738
F1 Score	0.7176	0.5127

These results clearly indicate that the quantum layer significantly enhances the model’s recall performance, with a 47.8% drop in recall observed when the layer is removed. This

suggests that the quantum component helps the model capture subtle patterns related to fake news, thereby reducing the number of false negatives.

While the classical variant achieves slightly higher accuracy and precision scores, this comes at the cost of missing a large number of fake news instances. These findings confirm that the quantum-enhanced model adopts a recall-oriented decision strategy, which may be desirable in critical applications where failing to detect fake content is riskier than flagging false positives.

4. Discussion

4.1. Interpretation of HQDNN Model Performance

The experimental findings demonstrate that the proposed HQDNN architecture exhibits a clear advantage in recall performance across both of its variants: the TF-IDF-based HQDNN and the DistilBERT-enhanced HQDNN. These results highlight the model's strong ability to detect fake news, even with differing input representations. In the TF-IDF-based HQDNN model, a recall score of 92.58% was achieved with the test set. This result confirms the model's high sensitivity to the positive class (i.e., fake news), which is essential in minimizing the number of false negatives. In contrast, the DistilBERT with Quantum variant attained an even higher recall of 94.40%, suggesting that contextualized embeddings further enhance the model's capacity to capture subtle semantic cues associated with fake news. However, both the models exhibit relatively low precision—58.08% for the TF-IDF version and 57.80% for the DistilBERT variant—indicating a shared tendency to misclassify some true news as fake. This outcome reflects a deliberate high-recall–low-precision trade-off. In many real-world scenarios, such as misinformation detection, cybersecurity, and regulatory oversight, prioritizing the identification of harmful content over the risk of false alarms can be a strategic design choice.

It is particularly noteworthy that these recall-oriented results are achieved with compact model architectures. The TF-IDF variant relies on a minimal number of classical parameters, while the DistilBERT-based version benefits from pretrained transformer representations combined with a lightweight quantum layer. These findings suggest that hybrid quantum–classical models can deliver a robust performance with minimal computational overhead, especially in natural language processing tasks where fine-grained class separation is critical.

Taken together, both the configurations of the HQDNN demonstrate that quantum-assisted modeling can be a viable and efficient solution in resource-constrained or safety-critical applications, offering a compelling balance between sensitivity and generalization.

4.2. Comparative Evaluation with Classical and Transformer-Based Models

The classical classifiers examined in this study—Logistic Regression, the Linear SVM, and Multilayer Perceptron (MLP)—achieved relatively higher scores in overall accuracy and precision when compared to those of the TF-IDF-based HQDNN model. Among them, MLP stood out, with an accuracy of 63.38% and precision of 66.45%, outperforming the HQDNN (TF-IDF) in these metrics. However, all the classical models exhibited a limited recall performance, ranging from 70.73% to 74.09%. For instance, although Logistic Regression yielded an accuracy of 61.96%, its recall score remained at 74.09%, significantly lower than the HQDNN's 92.58% with TF-IDF and 94.40% with DistilBERT.

This trend reflects a crucial distinction in classification behavior: classical models tend to construct conservative decision boundaries that favor the negative class (i.e., true news), thereby minimizing the number of false positives, but increasing the number of false negatives. In contrast, the HQDNN adopts a recall-oriented learning strategy, favoring the positive class (fake news). This leads to more false positives, but substantially fewer false

negatives, which is a desirable trade-off in domains where failing to detect fake content carries higher risk than occasional false alarms. This advantage becomes more apparent when contrasted with lightweight transformer-based architectures such as SetFit, which—despite its parameter efficiency—did not match the HQDNN’s recall levels, reaffirming the strength of hybrid quantum designs for high-sensitivity tasks.

From a cost-sensitive learning perspective, classical models treat all errors equally, which may not suit tasks like misinformation detection. The HQDNN, through its hybrid structure and quantum layer, demonstrates a more expressive capacity to represent minority class patterns. This is further supported by the ablation experiment, where removing the quantum layer reduced recall from 94.40% to 47.38%, making its behavior more aligned with the classical models.

Furthermore, when compared to the transformer-based DistilBERT baseline, the HQDNN remains competitive. Although DistilBERT achieved the highest F1 score (71.76%) and recall (94.40%) among all the models, the HQDNN’s DistilBERT-enhanced variant came very close in performance, while retaining the benefits of a hybrid design with significantly fewer parameters than the full-scale transformer models.

In summary, the HQDNN stands out for its ability to prioritize recall and handle imbalanced classification tasks effectively. Its recall-driven strategy, combined with the benefits of quantum-assisted computation, reinforces the potential of hybrid quantum-classical models as viable and efficient alternatives to both shallow and deep classical architectures, especially when the cost of false negatives is high.

4.3. *The Impact of the Quantum Layer*

The distinct performance gains achieved by the HQDNN model, particularly in recall, can be primarily attributed to its hybrid architecture that incorporates a parameterized quantum circuit. This quantum layer plays a critical role in transforming the feature space by enabling learning through quantum state evolutions, in addition to classical weight updates. Operating in a Hilbert space, the quantum component allows for the model to explore a richer and more expressive probabilistic landscape compared to conventional parametric surfaces.

In this study, the model employs a lightweight two-qubit quantum circuit, which encodes classical TF-IDF features using RX rotation gates, applies learnable RY and RZ gates, and produces outputs through Pauli-Z measurements. Despite its compact design, this configuration facilitates the construction of highly nonlinear decision boundaries using only a minimal number of trainable parameters. This is consistent with the findings in prior literature [12,16], which emphasize the ability of quantum circuits to efficiently represent complex data structures with fewer resources, a crucial advantage in scenarios with limited data or computational constraints.

From a task-specific perspective, fake news detection often involves short, semantically dense statements that exhibit subtle linguistic patterns not easily separable using classical linear models. In such cases, the quantum circuit’s ability to map inputs to highly nonlinear and probabilistic representations in a Hilbert space may facilitate the discrimination of borderline cases, particularly those prone to class ambiguity. The use of RX, RY, and RZ gates in our parameterized circuit enables flexible angular encodings that adaptively rotate the input space, generating multidimensional decision boundaries not trivially achieved by classical architectures. This transformation capability likely contributes to the superior recall performance, as the model becomes more sensitive to minority class instances, without substantially increasing the model complexity.

The significance of the quantum layer’s contribution became particularly evident through the ablation experiment. When the quantum component was removed

(HQDNN_NoQuantum), the model's recall dropped sharply from 94.40% to 47.38%. While the accuracy and precision experienced a slight increase in the classical variant, this came at the cost of a dramatic rise in false negatives, undermining the primary objective of fake news detection. This contrast empirically validates the quantum layer's role in enhancing sensitivity to the positive class and confirms its effectiveness in reducing critical misclassifications.

Furthermore, these results demonstrate that even with a minimal quantum footprint—only two qubits—the hybrid model is capable of outperforming several classical models and approaching the performance of more complex architectures like DistilBERT, all while maintaining a lower parameter count. This highlights the potential for practical deployment of hybrid quantum–classical models in tasks requiring superior recall, compact design, and resource efficiency.

One possible explanation for the HQDNN model's superior recall, but comparatively lower precision and accuracy lies in the expressive nature of the quantum circuit's decision boundaries. The parameterized quantum layer operates in a Hilbert space and applies unitary transformations capable of modeling highly nonlinear and entangled feature relationships. This enhanced expressiveness enables the model to capture subtle patterns associated with the minority (positive) class—i.e., fake news instances—which often exhibit complex or ambiguous linguistic features. As a result, the HQDNN is more likely to detect such instances, thereby improving recall. However, this comes with an increased risk of over-generalization in the feature space, leading to more false positives and a drop in precision. This behavior aligns with the prior literature on Quantum Machine Learning, where the entangled state space can prioritize sensitivity over specificity [15,16]. While this recall-focused behavior may be advantageous in high-stakes scenarios such as misinformation detection, where failing to identify fake news can be more costly than false alarms—it also highlights the need for future work on regularization strategies and decision-boundary calibration in hybrid quantum–classical models.

Additionally, we include the SetFit model—a new transformer-based approach optimized for few-shot classification—which demonstrated balanced precision (63.87%) and recall (65.13%), resulting in an F1 score of 64.49%. While SetFit does not surpass the HQDNN in terms of recall or F1 score, its competitive performance with significantly lower resource requirements underscores its value as a lightweight and CPU-efficient alternative. This comparison further highlights the effectiveness of the HQDNN in high-sensitivity tasks where minimizing false negatives is critical.

5. Conclusions

In this study, we proposed the HQDNN (Hybrid Quantum–Deep Neural Network), a hybrid artificial intelligence model designed for the automatic detection of fake news. The model integrates classical fully connected layers with a parameterized quantum circuit, aiming to develop a decision-making mechanism that incorporates both classical and quantum representations during learning.

We evaluated two variants of the model using the widely recognized LIAR dataset: one based on traditional TF-IDF representations and another leveraging transformer-based contextual embeddings (DistilBERT). The comparative experiments revealed that the DistilBERT-enhanced HQDNN outperformed all the other models in recall (94.40%) and F1 score (71.76%), indicating superior sensitivity to fake news instances. Although classical models, such as MLP and Logistic Regression, achieved higher overall accuracy and precision, the HQDNN models—particularly the DistilBERT-based version—demonstrated a strategic recall-oriented decision bias, which is desirable in high-risk applications, such as misinformation detection and content moderation.

The ablation experiment further confirmed the value of the quantum component. When the quantum layer was removed, the model's recall performance dropped drastically (from 94.40% to 47.38%), validating the quantum circuit's contribution to the model's ability to capture subtle patterns and reduce the number of false negatives.

In addition to the classical baselines, we included the SetFit model—a recent lightweight transformer-based architecture—as a comparative benchmark. While SetFit demonstrated a balanced performance with an F1 score of 64.49%, it did not match the HQDNN in recall or overall sensitivity, further reinforcing the hybrid model's advantage in recall-critical tasks.

Overall, this study demonstrates that quantum-enhanced models can serve as viable alternatives to the classical approaches, particularly in tasks characterized by class imbalance or asymmetric error costs. Even with a two-qubit quantum circuit, the HQDNN achieved competitive results, illustrating the promise of compact quantum neural architectures.

For future work, several research directions are envisioned:

- Increasing the quantum capacity by extending the circuit beyond two qubits and introducing entanglement mechanisms to better model complex feature interactions.
- Deploying on real quantum hardware (e.g., IBM Q and IonQ) to assess the performance differences between simulation and quantum execution.
- Evaluating model generalizability across alternative fake news datasets (e.g., Fake-NewsNet) and multilingual settings.
- Building ensemble architectures that combine the HQDNN with classical models using techniques such as stacking or voting, enhancing robustness in decision making.

In conclusion, this work serves as an early, but promising demonstration of hybrid quantum–deep learning in natural language processing. The findings suggest that quantum computation may soon become a practical complement to classical machine learning in domains that demand high sensitivity and low false negative rates.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article; further inquiries can be directed to the corresponding author.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Rotella, S. The COVID-19 Infodemic. *Lancet Infect. Dis.* **2020**, *20*, 875. [[CrossRef](#)] [[PubMed](#)]
2. Toğaçar, M.; Eşidir, K.A.; Ergen, B. Yapay Zekâ Tabanlı Doğal Dil İşleme Yaklaşımını Kullanarak İnternet Ortamında Yayınlanmış Sahte Haberlerin Tespiti. *J. Intell. Syst. Theory Appl.* **2021**, *5*, 1–8. [[CrossRef](#)]
3. Goldberg, Y.; Levy, O. Word2vec Explained: Deriving Mikolov et al.'s Negative-Sampling Word-Embedding Method. *arXiv* **2014**, arXiv:1402.3722.
4. Hossain, T.; Zahin Mauni, H.; Rab, R. Reducing the Effect of Imbalance in Text Classification Using SVD and GloVe with Ensemble and Deep Learning. *Comput. Inform.* **2022**, *41*, 98–115. [[CrossRef](#)]
5. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Minneapolis, MN, USA, 2–7 June 2019; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 4171–4186.
6. Allcott, H.; Gentzkow, M. Social Media and Fake News in the 2016 Election. *J. Econ. Perspect.* **2017**, *31*, 211–236. [[CrossRef](#)]
7. Wang, W.Y. “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Vancouver, BC, Canada, 30 July 2017; Association for Computational Linguistics: Stroudsburg, PA, USA, 2017; pp. 422–426.

8. Shu, K.; Sliva, A.; Wang, S.; Tang, J.; Liu, H. Fake News Detection on Social Media: A Data Mining Perspective. *ACM SIGKDD Explor. Newsl.* **2017**, *19*, 22–36. [[CrossRef](#)]
9. Cahyani, D.E.; Patasik, I. Performance Comparison of TF-IDF and Word2Vec Models for Emotion Text Classification. *Bull. Electr. Eng. Inform.* **2021**, *10*, 2780–2788. [[CrossRef](#)]
10. Alammary, A.S. BERT Models for Arabic Text Classification: A Systematic Review. *Appl. Sci.* **2022**, *12*, 5720. [[CrossRef](#)]
11. Rashkin, H.; Choi, E.; Jang, J.Y.; Volkova, S.; Choi, Y. Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Copenhagen, Denmark, 9–11 September 2017; Association for Computational Linguistics: Stroudsburg, PA, USA, 2017; pp. 2931–2937.
12. Yang, F.; Wang, X.; Ma, H.; Li, J. Transformers-Sklearn: A Toolkit for Medical Language Understanding with Transformer-Based Models. *BMC Med. Inform. Decis. Mak.* **2021**, *21*, 90. [[CrossRef](#)] [[PubMed](#)]
13. Zhang, L.; Zhang, X.; Zhou, Z.; Zhang, X.; Yu, P.S.; Li, C. Knowledge-Aware Multimodal Pre-Training for Fake News Detection. *Inf. Fusion* **2025**, *114*, 102715. [[CrossRef](#)]
14. Jiang, Y.; Wang, T.; Xu, X.; Wang, Y.; Song, X.; Maynard, D. Cross-Modal Augmentation for Few-Shot Multimodal Fake News Detection. *Eng. Appl. Artif. Intell.* **2025**, *142*, 109931. [[CrossRef](#)]
15. Schuld, M.; Sinayskiy, I.; Petruccione, F. An Introduction to Quantum Machine Learning. *Contemp. Phys.* **2015**, *56*, 172–185. [[CrossRef](#)]
16. Schuld, M.; Killoran, N. Quantum Machine Learning in Feature Hilbert Spaces. *Phys. Rev. Lett.* **2019**, *122*, 040504. [[CrossRef](#)] [[PubMed](#)]
17. Senokosov, A.; Sedykh, A.; Sagingalieva, A.; Kyriacou, B.; Melnikov, A. Quantum Machine Learning for Image Classification. *Mach. Learn. Sci. Technol.* **2023**, *5*, 15040. [[CrossRef](#)]
18. Sagingalieva, A.; Komorniyk, S.; Joshi, A.; Mansell, C.; Pinto, K.; Pflitsch, M.; Melnikov, A. Photovoltaic Power Forecasting Using Quantum Machine Learning. *arXiv* **2023**. [[CrossRef](#)]
19. Chaudhary, A.; Han, J.; Kim, S.; Kim, A.; Choi, S. Anomaly Detection and Analysis in Nuclear Power Plants. *Electronics* **2024**, *13*, 4428. [[CrossRef](#)]
20. Ardeshir-Larijani, E.; Nasiri Fatmehsari, M.M. Hybrid Classical-Quantum Transfer Learning for Text Classification. *Quantum Mach. Intell.* **2024**, *6*, 19. [[CrossRef](#)]
21. Yu, W.; Yin, L.; Zhang, C.; Chen, Y.; Liu, A.X. Application of Quantum Recurrent Neural Network in Low-Resource Language Text Classification. *IEEE Trans. Quantum Eng.* **2024**, *5*, 1–13. [[CrossRef](#)]
22. Zeguendry, A.; Jarir, Z.; Quafafou, M. Quantum-Enhanced K-Nearest Neighbors for Text Classification: A Hybrid Approach with Unified Circuit and Reduced Quantum Gates. *Adv. Quantum Technol.* **2024**, *7*, 2400122. [[CrossRef](#)]
23. Khattar, D.; Goud, J.S.; Gupta, M.; Varma, V. MVAE: Multimodal Variational Autoencoder for Fake News Detection. In Proceedings of the The World Wide Web Conference, San Francisco, CA, USA, 13 May 2019; ACM: New York, NY, USA, 2019; pp. 2915–2921.
24. Durgut, S.; Küçükösille, E.U.; Tokmak, M. Hybrid Quantum-Classical Deep Neural Networks Based Smart Contract Vulnerability Detection. *Appl. Sci.* **2025**, *15*, 4037. [[CrossRef](#)]
25. Bokhan, D.; Mastiukova, A.S.; Boev, A.S.; Trubnikov, D.N.; Fedorov, A.K. Multiclass Classification Using Quantum Convolutional Neural Networks with Hybrid Quantum-Classical Learning. *Front Phys.* **2022**, *10*, 1069985. [[CrossRef](#)]
26. Liu, J.; Lim, K.H.; Wood, K.L.; Huang, W.; Guo, C.; Huang, H.-L. Hybrid Quantum-Classical Convolutional Neural Networks. *Sci. China Phys. Mech. Astron.* **2021**, *64*, 290311. [[CrossRef](#)]
27. Yu, D.; Zhou, Y.; Zhang, S.; Li, W.; Small, M.; Shang, K. Information Cascade Prediction of Complex Networks Based on Physics-Informed Graph Convolutional Network. *New J. Phys.* **2024**, *26*, 013031. [[CrossRef](#)]
28. Yun-tao, Z.; Ling, G.; Yong-cheng, W. An Improved TF-IDF Approach for Text Classification. *J. Zhejiang Univ. Sci. A* **2005**, *6*, 49–55. [[CrossRef](#)]
29. White, A.G.; James, D.F.V.; Munro, W.J.; Kwiat, P.G. Exploring Hilbert Space: Accurate Characterization of Quantum Information. *Phys. Rev. A* **2001**, *65*, 012301. [[CrossRef](#)]
30. Bergholm, V.; Izaac, J.; Schuld, M.; Gogolin, C.; Ahmed, S.; Ajith, V.; Alam, M.S.; Alonso-Linaje, G.; AkashNarayanan, B.; Asadi, A.; et al. PennyLane: Automatic Differentiation of Hybrid Quantum-Classical Computations. *arXiv* **2018**, arXiv:1811.04968.
31. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2014**, arXiv:1412.6980.
32. Chen, Z. (HTBNet)Arbitrary Shape Scene Text Detection with Binarization of Hyperbolic Tangent and Cross-Entropy. *Entropy* **2024**, *26*, 560. [[CrossRef](#)]
33. Kramer, O. *Machine Learning for Evolution Strategies; Studies in Big Data*; Springer International Publishing: Cham, Switzerland, 2016; Volume 20, ISBN 978-3-319-33381-6.
34. Khaled, A.; Manayer, A. Few-Shot Learning Approach for Arabic Scholarly Paper Classification Using SetFit Framework. *WSEAS Trans. Commun.* **2024**, *23*, 89–95. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.