EPJ Web of Conferences **360**, 01008 (2026) https://doi.org/10.1051/epjconf/202636001008
QUICK'26

Hybrid Quantum-Classical NLP Framework for Emotion Recognition Tasks

Dharaniraj VM[1,*], *Jayasuriya* R[1], *Sivasankaramoorthy* R[1], *Stephen* T[1], and *Abirami* Murugappan[1,**]

[1]Department of Information Science and Technology, College of Engineering Guindy, Anna University, Chennai, India

Abstract. This project aims at developing a comprehensive hybrid quantum-classical NLP pipeline to solve the task of emotion recognition. For emotion recognition, this work makes use of a hybrid Quantum Recurrent Neural Network (QRNN) by combining classical embedding layers with parameterized quantum circuits based on well defined ansatzes. The primary goal of the project is to train VQCs for capturing the emotion associated with sentences. This work aims at combining the strong embedding generation ability of classical architectures with the expressive power of Quantum circuits that can capture complex dependencies across sequences. The designed Hybrid QRNN model integrates two important aspects namely the strength of RNNs for sequential modelling and the representational power of quantum circuits resulting in competitive performance for emotion recognition task while also utilising the quantum benefits. The results reveal the potential of QNLP for advancing the intersection of Quantum Computing and Natural Language Processing for quantum-accelerated downstream NLP applications.

1 Introduction

The rapid growth in the fields of Artificial Intelligence (AI) and Natural Language Processing (NLP) has lit the inherent limitations of classical computing for handling complex tasks like sentiment analysis. While traditional machine learning and deep learning approaches have achieved considerable success, they struggle with high computational costs and and do not scale well when applied to large-scale, heterogeneous linguistic datasets. Their high computational demands in terms of time and space demands a transformative approach that incorporates the principles of quantum computing into NLP thus forming a new field called Quantum Natural Language Processing. By utilising the quantum principles like entanglement and superposition, QNLP allows exponentially expressive language representations, enhanced memory efficiency, and high parallelism. These characteristics position QNLP as a novel solution capable of overcoming the computational bottlenecks of classical systems and architectures while creating new opportunities for a scalable, accurate and sophisticated Natural Language Understanding (NLU) systems in the era of Big Data.

Quantum Natural Language Processing (QNLP) provides a promising framework that offers the potential for more expressive and compact representations of linguistic structures.

---

[*] e-mail: dharaniraj2205@gmail.com
[**] e-mail: abirami@auist.net

Despite its promise, the challenge remains in designing quantum circuits that can effectively encode and process natural language while remaining compatible with current quantum hardware constraints. Also due to current limitations in quantum computing, it is not feasible to develop full end-to-end quantum NLP pipelines from embedding generation to label prediction because the required qubit count grows rapidly and exceeds the capabilities of present hardware. Hybrid QNLP architectures design and implement models that combine the representational strengths of classical embeddings with computational advantages of quantum circuits. To demonstrate the functionalities, the training and execution of parameterized quantum circuits are carried out in a simulated quantum environment, and this project uses PennyLane for the simulation.

The reviewed literature makes it clear that QNLP research is still in its formative stage. Classical machine learning methods continue to outperform quantum models in most real-world tasks, with quantum systems showing only modest improvements or functioning well under highly constrained conditions.

This work aims to bridge the gap that lies in training Parameterized Quantum Circuits by exploring quantum-enhanced NLP techniques that maximize the expressive power of quantum systems for language understanding, thereby advancing the intersection of quantum computing and natural language processing. Hence the core research problem laid down is in advancing quantum-enhanced NLP techniques that can bridge the gap between Quantum Computing and Natural Language processing by developing circuit architectures and hybrid frameworks capable of harnessing the expressive power of quantum systems to model and understand natural language more effectively. Further the results validate the potential of quantum-enhanced architectures for natural language understanding tasks, even under current simulation constraints.

2 Literature Survey

This section reviews existing research on QNLP, training Parameterized Quantum Circuits (PQCs), and efforts to adapt classical models with quantum computing concepts.

2.1 Quantum Natural Language Processing

Varmantchaonala et al. [1] provide a comprehensive survey of QNLP, outlining a broad spectrum of approaches ranging from early models such as Quantum Bag-of-Words to advanced frameworks like DisCoCat and hybrid quantum-classical architectures implemented with toolkits such as PennyLane. Thota and Dash [3] compared traditional NLP methodologies with quantum machine learning models using social media data. Their evaluation on 22,000 FIFA World Cup tweets demonstrated that classical models outperformed quantum circuits, with accuracies of 84.3% and 85.7% for SVM and Random Forest respectively, compared to 78.5% and 80.2% for single- and two-layer quantum circuits. Despite the performance gap, their work highlighted QNLP's potential in exploring novel feature spaces and pointed to key technical challenges such as noise, decoherence, and difficulties in encoding high-dimensional text features.

2.2 QNLP Approaches for Sentiment Analysis

Xu et al. [4] demonstrated that a PQC-based Quantum Recurrent Neural Network (QRNN) can achieve competitive performance on a standard sentiment analysis task with significantly fewer parameters than classical models. Our research is motivated to explore novel QRNN architectures for text classification, investigating the feasibility of applying such architectures on near-term quantum hardware.

Yu et al. [5] introduced a quantum-classical hybrid architecture for text classification in low-resource languages, specifically for Bengali sentiment analysis. The paper proposes a Batch Uploading Quantum Recurrent Neural Network and its non-parameter-shared variant to address limitations of traditional QRNNs such as semantic information loss from dimensionality reduction. Their approach involves using a pretrained multilingual BERT model to obtain word embeddings, which are then processed by quantum models for feature extraction and class label prediction.

2.3 Other Hybrid Architectures for QNLP

Cavar and Zhang [6] introduced a hybrid quantum-classical framework that maps classical word and text embeddings (from FastText, BERT, etc.) to quantum states using encoding strategies like Amplitude Encoding, demonstrating that the resulting quantum embeddings preserve semantic similarity with minimal information loss and significantly reduced memory requirements.

Based on the identified research gaps, this work contributes to QNLP by:

1. Developing a hybrid Quantum-Classical framework for sentiment analysis generalisable to similar text classification tasks.
2. Demonstrating that variational quantum circuits can be trained for all traditional NLP tasks beyond toy examples.
3. Bridging the methodological gap between modern NLP workflows and quantum computing.

3 System Architecture

The proposed architecture is designed so that the Quantum NLP architectures resemble their successful classical counterparts with much lesser memory requirements, also leveraging quantum properties to address the problem of Sentiment Analysis. The design thus aims to achieve compression ratios that map high-dimensional classical embeddings (1024-dimensional vectors) to compact quantum states (10-qubit representations) while preserving semantic information.

3.1 General Architecture of a Hybrid QNLP System

The general architecture of hybrid QNLP systems is depicted in Figure 1. The pipeline begins with a preprocessing stage where input sentences are tokenized, lemmatized, and transformed into contextual embeddings using pretrained models such as BERT. These embeddings serve as the interface between classical computation and the quantum layer. The quantum processing stage encodes embeddings into quantum states through methods such as angle encoding and amplitude encoding. Within this layer, different quantum-enhanced architectures are employed depending on the task. Since our task involves sequential emotion recognition, a Hybrid Quantum-Classical Recurrent Neural Network (QRNN) is used.

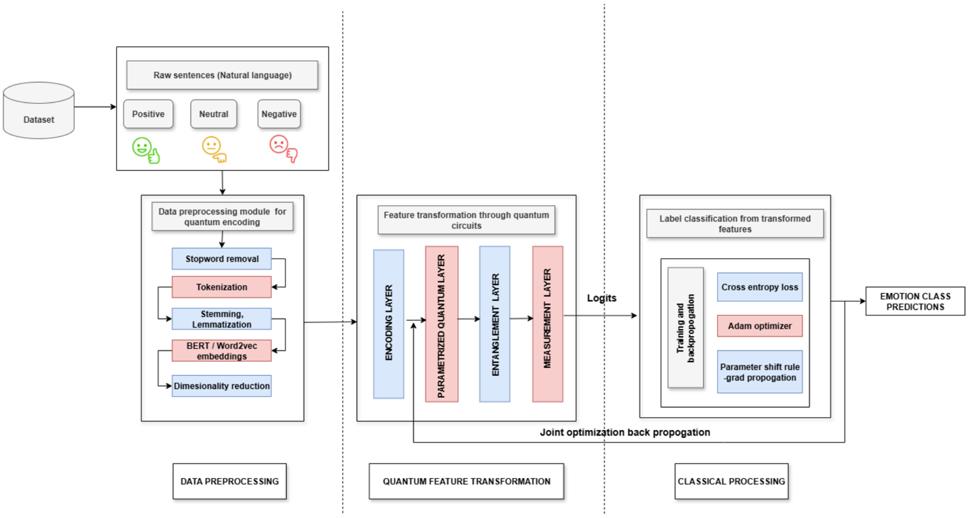

Figure 1. Quantum pipeline for NLP tasks

3.2 Hybrid QRNN Architecture

The Hybrid QRNN integrates quantum operations into the structure of a recurrent network, as shown in Figure 2. In this design, embeddings from the classical layer are passed on into a recurrent structure that is supported by parameterized quantum circuits. These circuits apply unitary transformations on encoded states, allowing the model to capture complex dependencies across sequences. The measurements from the quantum circuits are passed again into the recurrent unit to generate context enriched hidden states.

3.3 Variational Quantum Circuit

The crucial part of the developed hybrid QRNN pipeline for emotion recognition is the Variational Quantum Circuit as shown in Figure 3 proposed by [4]. The circuit consists of a Word Ansatz that maps classical word embeddings into the quantum Hilbert space via an angle encoding scheme, applying Rx and Rz rotation gates parameterized by the embedding values to initialize qubits in superposition. This is integrated with a Recurrent Ansatz that updates the hidden quantum state from the previous time-step by entangling it with the current word state through additional Rot and Rx gates thus effectively capturing long term dependencies in sentences. The circuit depth and number of qubits can be adjusted based on availability of computational resources, balancing accuracy with simulation feasibility. Overall, the VQC serves as a crucial element bridging classical embeddings and quantum representations, forming the backbone of the proposed hybrid architecture.

4 Methodology

The implementation of hybrid QRNN involves combining classical language representations with quantum computational features. This involves two steps. Embeddings are first generated to capture contextual meaning, which are then mapped into quantum circuit parameters.

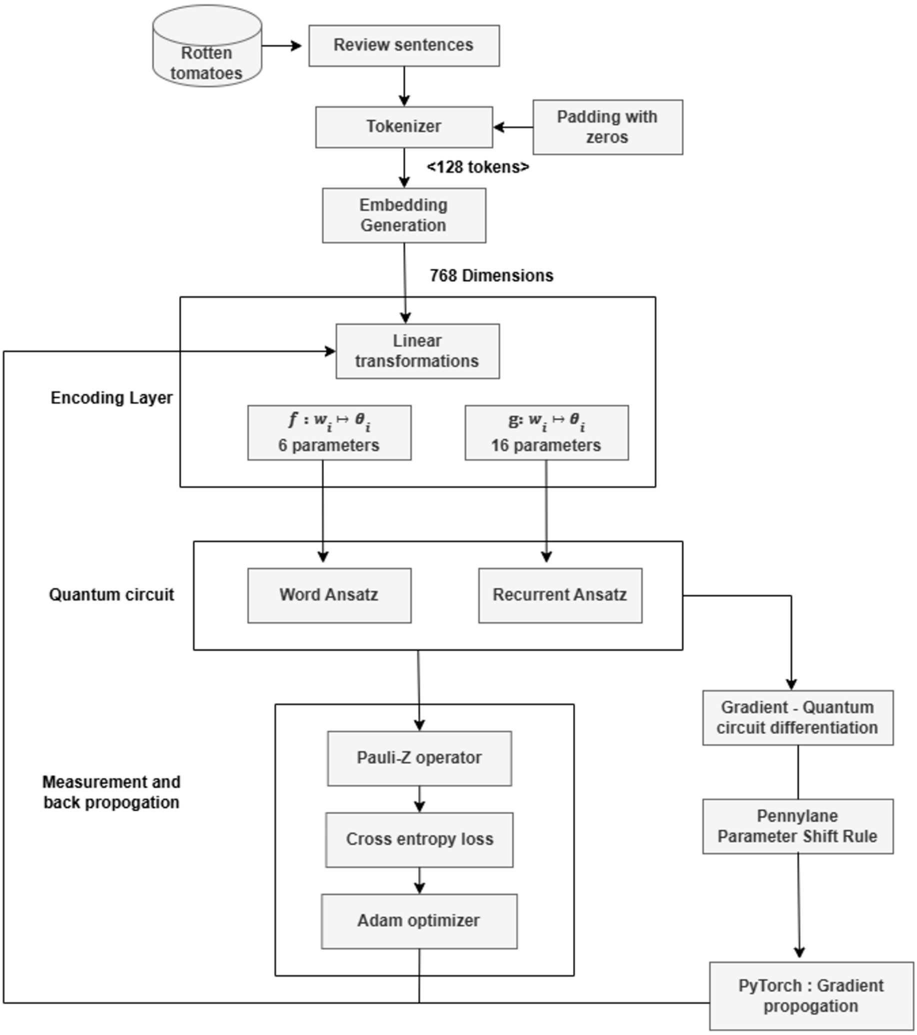

Figure 2. Hybrid QRNN Architecture using Quantum Circuits

These parameters then drive a parameterized quantum circuit designed to represent complex correlations. The second step involves training the circuits end-to-end to minimize the classification loss, thus enabling the quantum layer to learn progressively beyond purely classical architectures.

4.1 Classical Embedding Generation and Parameter Mapping

In the hybrid QRNN architecture, the classical embedding generation stage transforms textual input into numerical representations suitable for mapping. Each input sentence is tokenized using the BERT tokenizer, which converts the text into a fixed number of tokens. Every token is then represented as a 768-dimensional contextual embedding vector obtained from the pretrained BERT-base-uncased model. This BERT model is kept frozen during training to

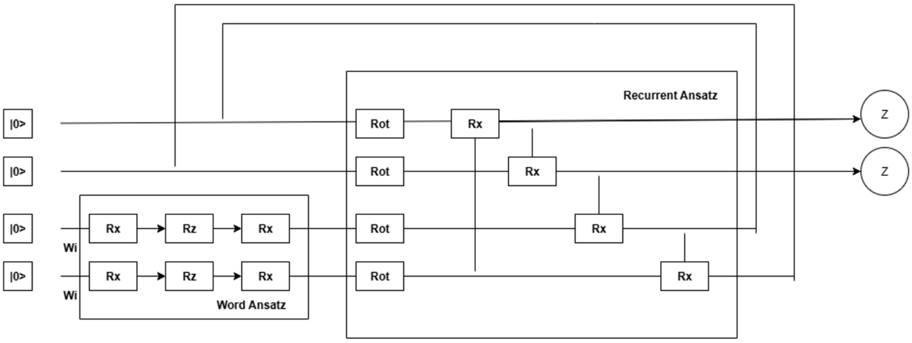

Figure 3. Variational Quantum Circuit Design

reduce computational overhead. Two learnable linear layers project high-dimensional BERT embeddings into learnable parameters (angles) that control the VQC:

- **Layer f:** Maps each 768-dimensional token embedding to six rotation angles ($3 \times 2 = 6$), driving the word ansatz which operates on two word qubits.
- **Layer g:** Maps each 768-dimensional token embedding to sixteen rotation parameters ($3 \times 4 + 4 = 16$), configuring the recurrent ansatz acting on all four recurrent qubits.

4.2 Quantum Recurrent Neural Network (QRNN) Model

The hybrid QRNN model that forms the quantum component of the pipeline is implemented using the `PennyLane` quantum machine learning framework, which seamlessly integrates with `PyTorch` for end-to-end differentiability and gradient propagation through the parameter-shift rule.

4.2.1 Qubit Configuration

The circuit operates on four qubits, representing the quantum memory of the recurrent unit. Out of these, two qubits are specifically dedicated to encoding the *word ansatz*, which injects token-specific information into the quantum circuit. The other two qubits, called the *recurrent qubits*, act as a quantum memory by preserving and propagating contextual and sentiment-related information across successive time steps. This allows the model to capture long-range dependencies in the sentence rather than relying only on local word-level features.

4.2.2 Word and Recurrent Ansatz

The word ansatz applies parameterized rotation gates (`RX` and `RZ`) to the word qubits, encoding token-level embedding information into the quantum state by rotating the qubits according to the mapped angles from layer f:

$$f: \; w_i \in \mathbb{R}^d \;\longrightarrow\; \boldsymbol{\theta}_i \in \mathbb{R}^6, \quad \boldsymbol{\theta}_i = W_f w_i + b_f \tag{1}$$

The recurrent ansatz, applied over all four qubits, consists of parameterized rotation gates (`Rot`) and a series of controlled-rotation (`CRX`) gates connecting adjacent qubits in a ring topology. This design introduces entanglement and allows the circuit to capture complex sequential dependencies analogous to memory in classical RNNs.

4.2.3 Measurement Operator

The expectation values of the Pauli-Z operator on the first two qubits form the output of the quantum circuit. These expectation values represent the model's logits, subsequently passed to a classical softmax layer for sentiment classification:

$$h_i = (\langle Z_1 \rangle_i,\ \langle Z_2 \rangle_i), \quad \langle Z_k \rangle_i = \langle \psi_i | Z_k | \psi_i \rangle,\ k \in \{1, 2\} \tag{2}$$

4.3 Training Procedure

The training was carried out in an end-to-end differentiable manner, integrating both classical and quantum components. The complete training process is summarized in Algorithm 1.

The `PyTorch` framework manages the computational graph and optimization process, while `PennyLane` handles quantum circuit differentiation through the parameter-shift rule, allowing accurate gradient computation for quantum parameters.

The quantum circuit outputs the expectation values of Pauli-Z operators, which are interpreted as the model's logits for sentiment classification. A cross-entropy loss function is used as the standard loss function to measure the difference between the predicted logits and the true sentiment labels. The gradients of the loss are propagated backward through both the classical linear layers and the quantum circuit using PennyLane's hybrid automatic differentiation. The Adam optimizer updates the trainable parameters to minimize the loss over successive epochs.

Algorithm 1 Training Procedure for Hybrid QRNN

```
Require: Dataset $D = \{(x_i, y_i)\}$, frozen BERT model, learning rate $\eta$, batch size $B$, epochs $E$
Ensure: Trained QRNN model parameters
Initialize quantum device with $n\_qubits = 4$, $word\_qubits = 2$
Load frozen BERT-base-uncased for embeddings
Initialize linear layers $f$ and $g$ for parameter mapping
Initialize Adam optimizer and CrossEntropy loss
for epoch $e = 1$ to $E$ do
    for each batch $(X, Y)$ in $D$ do
        Tokenize input; obtain BERT embeddings $h_k \in \mathbb{R}^{768}$
        for each token embedding $h_k$ do
            Compute $\theta_f = f(h_k)$, $\theta_g = g(h_k)$
            Pass $(\theta_f, \theta_g)$ to quantum circuit $q\_step$
            Measure expectation values $\langle Z \rangle$ from first two qubits
        end for
        Aggregate last token output as logits $\hat{y}$
        Compute loss $L = \mathrm{CrossEntropy}(\hat{y}, Y)$
        Backpropagate via parameter-shift rule
        Update parameters with Adam (step size $\eta$)
    end for
    Evaluate on validation set; record accuracy/loss
end for
```

5 Results and Analysis

The sentiment analysis experiment was conducted to classify movie reviews using the **Rotten Tomatoes dataset** into positive and negative categories using the proposed hybrid quantum-

Table 1. Rotten Tomatoes dataset split

Split	Samples	Purpose
Training	8,530	Model optimisation
Validation	1,066	Hyperparameter tuning
Testing	1,066	Final evaluation

classical architecture. The dataset had **10,662 samples**, which is divided into training, validation, and test subsets. Each review is a short sentence expressing clear sentiment polarity, making it an ideal dataset for binary classification tasks. The data distribution is shown in Table 1.

Before training, all reviews underwent a standardized preprocessing pipeline that applies *text normalization*, *tokenization*, *padding*, and *truncation*, where all tokenized sentences were adjusted to a uniform length of 128 tokens for consistent input dimensions. The preprocessed dataset was then converted into PyTorch tensors where each sample would consist of three main parts namely input IDs,attention masks and labels.

5.1 Outputs

The hybrid QRNN model was trained for fifteen epochs with Adam optimizer and with a learning rate of 1×10^{-3}. During inference, the model outputs expectation values for each sentiment class which is normalised into class probabilities using SoftMax activation function. A simple UI to visualise model predictions and confidence scores was developed that allows users to input sentences and get the model's prediction on the sentiment label. Figure 4 shows the developed interface with an example of user's input and the model's inference.

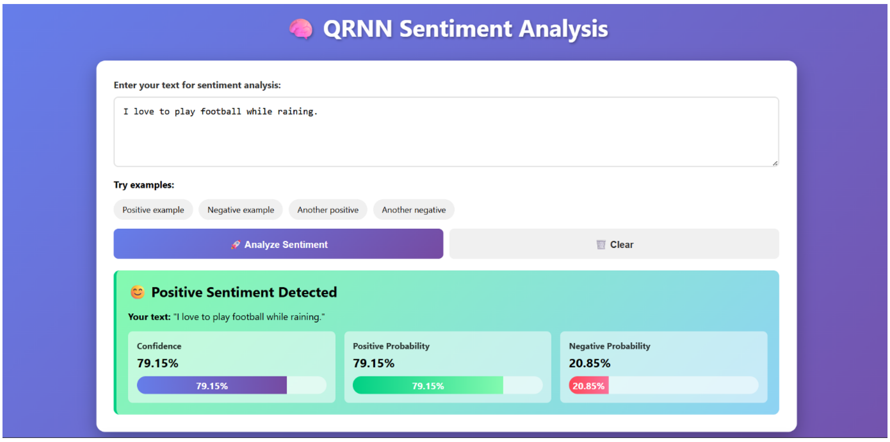

QRNN Sentiment Analysis

Enter your text for sentiment analysis:

`I love to play football while raining.`

Positive Sentiment Detected

Your text: "I love to play football while raining."

Confidence: 79.15%

Positive Probability: 79.15%

Negative Probability: 20.85%

Figure 4. Sentiment prediction using the trained QRNN model

5.2 Performance Analysis

The model has a good generalisation for the task of emotion recognition effectively identifying unforeseen posi- tive and negative reviews. The output probabilities reflect how confident the model is, offering some interpretable insights for the task. The model when evaluated

Table 2. Performance metrics on the test set

Class	**Prec.**	**Recall**	**F1**	**Support**
Negative (0)	0.742	0.831	0.784	533
Positive (1)	0.808	0.711	0.756	533
Macro Avg	0.775	0.771	0.770	1066
Accuracy		**0.771**		

Table 3. Performance Comparison with Classical Baselines

Model	**Acc. (%)**	**Recall**	**F1**
Fine-tuned BERT	87.0	0.91	0.87
Classical RNN (BiLSTM)	76.08	0.75	0.76
Hybrid QRNN (proposed)	**77.1**	0.77	0.77

on the test dataset achieves an **accuracy of 77.1%**, demonstrating an effective generalisation and learning. Detailed class-wise **precision**, **recall**, and **F1-score** values are summarized in Table 2.

The confusion matrix in Figure 5 shows that the model achieves balanced proportion of correct predictions for both the sentiment classes. The model demonstrates slightly higher recall for the negative class, indicating strong detection of negative sentiments, while maintaining a decent precision for positive reviews. These results prove the potential of quantum-enhanced architectures for natural language processing tasks including sentiment analysis, even under current simulation constraints.

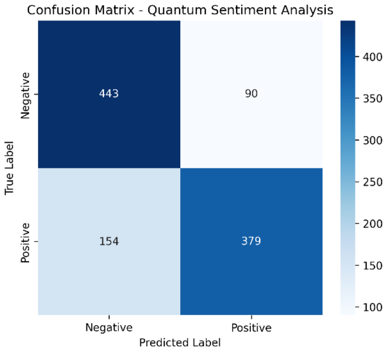

Figure 5. Confusion matrix - True vs Predicted labels

5.3 Comparison with Classical Architectures

To compare our hybrid Quantum-Classical RNN framework against strong classical baselines, a fine-tuned BERT transformer model and a classical recurrent neural network (BiLSTM) were trained for benchmarking. All models were trained on the same dataset with identical preprocessing and test-train splits. The results are illustrated in Table 3.

A pretrained Bert-base-uncased model was fine tuned with a classifier head specific to the task while the other steps remained the same as that of our hybrid quantum classical QRNN model. The intent of this work is not to compete with classical baselines, but to explore whether quantum circuits can learn meaningfully to prepare for potential quantum advantages as scalable quantum hardware evolves. Due to current limitations of NISQ devices and quantum simulators, such advantages are not yet evident.

5.4 Parameter Efficiency

In terms of parameters a reasonable comparison can be made between classical baselines and the hybrid QRNN pipeline.

To illustrate this, consider the number of parameters processed per input sentence in classical transformer models and QRNN assuming a sentence is divided into 128 tokens and each token is of 768 dimensions,

Classical token representation: $128 \times 768 = 98{,}304$

QRNN token parameters: $6 + 16 = 22$

Total QRNN params/sentence: $128 \times 22 = 2{,}816$

Reduction factor: $\frac{98{,}304}{2{,}816} \approx 34.9$

This is to show that QRNN learns meaningful representations while operating with approximately *35× fewer* parameters per sentence compared to classical transformer models.

Although the proposed hybrid QRNN model was trained on a quantum simulator to prove feasibility, the pipeline is still compatible with real quantum hardware. The PQCs trained within the PennyLane framework provide interoperability with Qiskit-based backends. The use of minimal qubits (four recurrent qubits) and the mapping of higher dimensional classical embeddings to fewer quantum parameters in the design ensure that we stay within the limits of current quantum advancements. Despite hardware compatibility, end-to-end training of QRNN on quantum hardware with increasing depths introduces more noise in NISQ devices and requires repeated circuit evaluations for gradient propagation through parameter-shift rule.

6 Conclusion

This work focused on exploring the potential of Quantum Natural Language Processing (QNLP) through the implementation of a primitive linguistic task - sentiment analysis. The model developed using a hybrid quantum-classical architecture that combines classical embeddings with parameterized quantum circuits was aimed at capturing non-linear and high-dimensional relationships within text. The successful training of quantum circuits proves that quantum pipelines can be implemented for NLP tasks without relying solely on existing classical pipelines. Moreover, the pipeline introduced in this work can be extended to support additional natural language processing tasks such as semantic similarity, question answering, and textual entailment. The work also establishes that training deeper ansatzes may require very strong simulation platforms or hardware but can capture more intricate and non linear relationships in language. By applying quantum principles to NLP, the model developed in this work was able to achieve meaningful performance with comparatively fewer parameters while maintaining strong generalization and stability. In summary, the hybrid quantum classical pipeline in this work would mark an important step towards bridging the gap between natural language understanding and quantum information processing,reinforcing the abilities of quantum computing in advancing future language processing models.

7 Future Works

This research lays down the foundation for a broader quantum NLP pipeline to solve all NLP tasks that fosters quantum-enhanced language processing. The implemented framework only demonstrates how PQCs can be effectively used for training to solve core NLP tasks, but there also exists considerable potential for further expansion. Future works can focus on training the PQCs in real time quantum hardware to analyze how real noise affects the performance and assess beyond simulation limits. As quantum technology grows larger, more powerful devices can scale these architectures to handle longer text sequences and more complex language structures, which increases the scope and impact of QNLP applications.

Another area for improvement is optimizing the use of quantum resources including reducing qubit count and circuit depth minimization for improving the efficiency of trained models. Ultimately, the proposed hybrid quantum NLP pipeline is a flexible and extensible framework for future research, enlightening the feasible pathway towards building pure quantum-powered language understanding systems that proves the real advantages of quantum computing in terms of time and space.

References

[1] C. M. Varmantchaonala, J. L. K. E. Fendji, J. Schöning, and M. Atemkeng, "Quantum Natural Language Processing: A Comprehensive Survey," *IEEE Access*, vol. 12, pp. 99578–99598, 2024. https://doi.org/10.1109/ACCESS.2024.3420707

[2] S. Ganguly, S. N. Morapakula, and L. M. P. Coronado, "Quantum natural language processing based sentiment analysis using lambeq toolkit," in *Proc. 2022 Second Int. Conf. Power, Control and Computing Technologies (ICPC2T)*, pp. 1–6, 2022. https://doi.org/10.1109/ICPC2T53885.2022.9776836

[3] S. Thota and S. K. Dash, "A comparison of traditional natural language processing methods for emotion recognition using quantum computing," in *Proc. 2025 Int. Conf. Machine Learning and Autonomous Systems (ICMLAS)*, pp. 669–673, 2025.

[4] W. Xu, S. Clark, D. Brown, G. Matos, and K. Meichanetzidis, "Quantum recurrent architectures for text classification," in *Proc. 2024 Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 18020–18027, 2024.

[5] W. Yu, L. Yin, C. Zhang, Y. Chen, and A. X. Liu, "Application of quantum recurrent neural network in low-resource language text classification," *IEEE Transactions on Quantum Engineering*, p. 21213, 2024.

[6] D. Cavar and C. Zhang, "Semantic Similarities using Classical Embeddings in Quantum NLP," in *Proc. 2024 IEEE Int. Conf. Quantum Computing and Engineering (QCE)*, pp. 450–451, 2024. https://doi.org/10.1109/QCE60285.2024.10350

[7] L. Fernandez, N. Rao, and J. Becker, "A multiclass Q-NLP sentiment analysis experiment using DisCoCat," in *Proc. 2024 Workshop on Quantum Natural Language Processing (QNLP)*, pp. 45–52, 2024. https://doi.org/10.18653/v1/QNLP2024.100

[8] A. Kumar, M. Patel, and R. Singh, "Semantic similarities using classical embeddings in quantum NLP," in *Proc. 2023 IEEE Int. Conf. Quantum Computing and Engineering (QCE)*, pp. 88–95, 2023. https://doi.org/10.1109/QCE57885.2023.00088