Toward Quantum-Enhanced Natural Language Understanding: Sarcasm and Claim Detection with QLSTM

Pritam Pal and **Dipankar Das**
Jadavpur University, Kolkata, India
{pritampal522, dipankar.dipnil2005}@gmail.com

Abstract

Traditional machine learning (ML) and deep learning (DL) models have shown effectiveness in natural language processing (NLP) tasks, such as sentiment analysis. However, they often struggle with complex linguistic structures, such as sarcasm and implicit claims. This paper introduces a Quantum Long Short-Term Memory (QLSTM) framework for detecting sarcasm and identifying claims in text, aiming to enhance the analysis of complex sentences. We evaluate four approaches: (1) classical LSTM, (2) quantum framework using QLSTM, (3) voting ensemble combining classical and quantum LSTMs, and (4) hybrid framework integrating both types. The experimental results indicate that the QLSTM approach excels in sarcasm detection, while the voting framework performs best in claim identification.

1 Introduction

Sarcasm and claims play a crucial role in everyday communication, especially on social media. Sarcasm, defined as remarks that convey the opposite of their literal meaning for humorous or critical effect, poses a challenge due to its ironic nature (e.g., "Oh, I just love it when my internet decides to take a vacation in the middle of an important Zoom call."). At the same time, social media serves as a vast source of information, where 'claims' (statements presented as facts) require verification to combat misinformation and fake news.

While natural language processing (NLP) has advanced significantly in detecting sarcasm and identifying claims, its exploration within emerging quantum computing environments remains limited. Quantum Computing (QC), leveraging principles such as superposition and entanglement (Gyongyosi and Imre, 2019), offers the potential for faster computation and more efficient resource utilization, as demonstrated by Shor's algorithm (Shor, 1997), which can break RSA encryption (Rivest et al., 1978) significantly faster than classical methods.

Quantum Machine Learning (QML), a key application of QC, shows promise in handling complex, noisy data and learning from smaller datasets more effectively than classical approaches (Neumann et al., 2019). Inspired by these advantages, we analyze the application of QML to these challenging NLP tasks.

This paper evaluates QML frameworks in sarcasm detection and claim identification against classical machine learning methods. Key contributions include: 1) Development of classical LSTM and quantum LSTM (QLSTM) frameworks, 2) A voting approach combining classical and quantum models, and 3) A hybrid classical-quantum LSTM framework for improved performance.

2 Related Work

Quantum computing applications in NLP are still in their early stages despite increasing interest. Early works explored quantum language models using quantum probability theory, demonstrating improved perplexity scores and aiding in word sense disambiguation (Basile and Tamburini, 2017; Tamburini, 2019).

Recently, quantum probability has been applied to multimodal tasks, including sentiment and sarcasm detection (Liu et al., 2021) as well as emotion detection (Li et al., 2023). Additionally, Variational Quantum Circuits (VQC) have been employed for multimodal sentiment analysis, demonstrating promising results on datasets such as CMU-MOSEI (Phukan and Ekbal, 2023).

Beyond these, QLSTM variations have been explored for Part-of-Speech (POS) tagging, including unidirectional (Sipio et al., 2021; Pandey et al., 2022) and bidirectional approaches (Pandey and

Proceedings of Recent Advances in Natural Language Processing, pages 854–859
Varna, Sep 8–10, 2025
https://doi.org/10.26615/978-954-452-098-4-097

Pakray, 2023), even for low-resource and code-mixed languages. Quantum frameworks have also explored in text classification (Xu et al., 2024; Shi et al., 2023), sentiment analysis (Yan et al., 2024; Zhang et al., 2019), and metaphor detection (Qiao et al., 2024).

The "DisCoCat" framework (Coecke et al., 2020) is a notable quantum NLP approach that preserves linguistic structure by mapping it to quantum circuits. Its applications, particularly in sentiment analysis, have been explored using the 'lambeq' open-source Python library (Ruskanda et al., 2023, 2022; Ganguly et al., 2022). For a more comprehensive overview of Quantum NLP, several survey papers provide detailed discussions (Wu et al., 2021; Guarasci et al., 2022; Varmantchaonala et al., 2024; Widdows et al., 2024).

3 Dataset

We utilized two publicly available datasets for sarcasm detection and claim identification. For sarcasm detection, we focused on the "Eye-tracking and Sentiment Analysis II" dataset (Mishra et al., 2016), specifically its sarcasm labels. This dataset comprises a total of 994 sentences, with 664 instances being non-sarcastic and 350 cases being sarcastic.

For claim identification, we utilized the dataset by Rosenthal and McKeown (2012), which comprises 3985 sentences from LiveJournal blogs and Wikipedia discussions. Of these, 2480 were opinionated claims and 1505 were not claims.

Both datasets used an 85-15 ratio for training and testing. For sarcasm detection, 844 instances were trained (549 sarcastic, 295 non-sarcastic), and 150 instances were used for testing (95 sarcastic, 55 non-sarcastic). For claim identification, 3,386 samples were used for training (2,129 claims, 1,257 non-claims) and 597 samples for testing (351 claims, 246 non-claims).

4 Methodology

Task Defination: Given tokenized sequence $S = [t_1, t_2, ..., t_n]$. Our objective is to classify S as either sarcastic/non-sarcastic and claim/not-claim, utilizing quantum machine learning.

For this work, we developed four system frameworks: 1) Classical LSTM-based framework, employing standard LSTM, 2) Quantum-based framework, utilizing Quantum LSTM (QLSTM), 3) Voting-based framework, which combines predictions from classical and QLSTM models based on confidence scores, and 4) Hybrid framework, integrating both QLSTM and classical LSTM layers sequentially. The overall system framework is depicted in Figure 1.

4.1 Framework Description

For all frameworks, input tokenized sequences were first passed through an embedding layer of 50 dimensions to obtain the embedding representation of the tokenized input.

Classical LSTM: As illustrated in Figure 1(a) (left), the embedding matrix was processed through a classical LSTM layer of 64 hidden units.

QLSTM: In Figure 1(a) (right), this framework illustrates a replacement of the classical LSTM with a QLSTM layer that has 64 hidden units. The QLSTM, introduced by Chen et al. (2020), transforms the weight matrices of the classical LSTM (W_f, W_i, W_C, W_o) into VQCs. The fundamental operations of the QLSTM cell can be mathematically expressed as follows:

$$\begin{aligned}
f_t &= \sigma(\text{VQC}_f([h_{t-1}, x_t])) \\
i_t &= \sigma(\text{VQC}_i([h_{t-1}, x_t])) \\
\hat{C}_t &= \tanh(\text{VQC}_C([h_{t-1}, x_t])) \\
c_t &= f_t \otimes c_{t-1} \oplus i_t \otimes \hat{C}_t \\
o_t &= \sigma(\text{VQC}_o([h_{t-1}, x_t])) \\
h_t &= o_t \otimes \tanh(c_t)
\end{aligned}$$

where h_{t-1} is the previous hidden state, x_t is the current input, σ is sigmoid function, and $\otimes, \oplus$ are element-wise multiplication and addition. VQCs, as shown in Figure 2, consist of a data encoding block ($U(x)$), a variational block ($V(\theta)$), followed by CNOT and single-qubit rotation gates, and a quantum measurement block. The rotation angles in $V(\theta)$ are iteratively updated via gradient descent.

Hybrid Framework: As shown in Figure 1(b), the embedding matrix was first fed to a QLSTM layer with 64 hidden units, whose output then served as input to a classical LSTM layer with 64 hidden units.

4.2 Classification

For both sarcasm detection and claim identification, the final output from the LSTM-based layers ($LSTM_{out}$), whether classical or quantum, feeds into a feed-forward output layer with two hidden units. The output layers used softmax as their activation function.

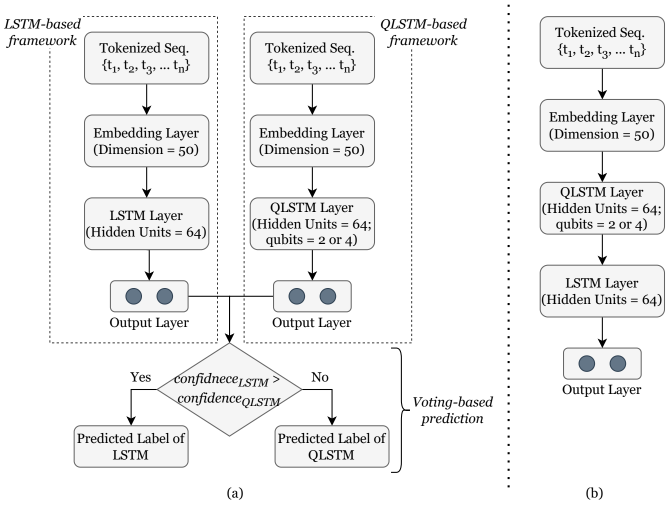

Figure 1: System framework flow diagram: (a) classical LSTM, quantum LSTM, and voting-based prediction; (b) hybrid framework.

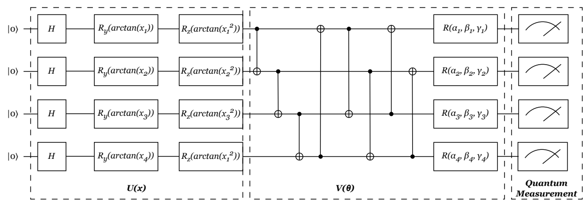

Figure 2: Typical block diagram of a Variational Quantum Circuit (VQC) (Chen et al., 2020). $U(x)$ is a data encoding, $V(\theta)$ is the variational block, and a Quantum Measurement block measures the output.

For the claim identification task specifically, a dropout layer (with a rate of 0.2) was applied immediately before the final output layer for model regularization. The classification process can be formally represented as:

$$\mathcal{P} = \text{softmax}(Dropout(LSTM_{out}))$$
$$\hat{Y} = \underset{\mathbf{j}}{argmax}(\mathcal{P})$$

(Note: Dropout was applied only for Claim Identification)

Where $\mathcal{P}$ represents the probability value for each class, $\hat{Y}$ indicates the predicted class label, and j denotes the number of classes.

4.3 Voting-Based Prediction

This scheme combines predictions from the classical and QLSTM frameworks. The final prediction $\hat{Y}$ is chosen from the framework (classical or quantum) whose predicted output label has a higher confidence score ($\mathcal{C} = \max(\mathcal{P})$).

$$\hat{Y} = \begin{cases} \text{argmax}_{\mathbf{j}}(\mathcal{P}_{lstm}), & \text{if } \mathcal{C}_{lstm} > \mathcal{C}_{qlstm} \\ \text{argmax}_{\mathbf{j}}(\mathcal{P}_{qlstm}), & \text{otherwise} \end{cases}$$

4.4 Training

During the training process, the dataset was divided into 85% for training and 15% for validation. All frameworks were trained with the 'CrossEntropy' loss function and the Adam optimizer (Kingma and Ba, 2017), with a learning rate of 2.5e-4. The number of epochs was taken as 20 for all frameworks, except for the hybrid framework, which was trained for varying durations: 5 and 8 epochs for sarcasm detection in the 2-qubit and 4-qubit frameworks, respectively, and 10 epochs for claim identification.

5 Experimental Setup and Result

All experiments were conducted in the Kaggle environment using PyTorch and PennyLane[1]. All QLSTM modules were configured with 2 and 4 qubits with a single quantum layer and executed on Pennylane's `default.qubit` quantum simulator. Performance of classical, QLSTM, voting-based, and hybrid frameworks for sarcasm and claim identification was evaluated using Precision, Recall, and macro F1-score on the test split.

Sarcasm Detection: Table 1 summarizes the overall results of sarcasm detection. The classical LSTM achieved a strong baseline F1-score of 84.17. In contrast, standalone QLSTM frameworks exhibited lower performance, with the 4-qubit QLSTM recording an F1-score of 76.20. However, the 4-qubit Hybrid framework achieved the highest overall F1-score of 87.18%. This represents an improvement of 3.46% over the classical LSTM

[1]https://pennylane.ai/

and 12.49% over the 4-qubit QLSTM, highlighting the advantages of integrating classical and quantum components. Additionally, voting-based frameworks performed comparably to the classical LSTM.

Framework	Qubit	Precision	Recall	F1-score
LSTM	—	86.77	82.82	84.17
QLSTM	2	74.92	74.31	74.58
	4	76.49	78.52	76.20
Voting	2	86.24	81.91	83.33
	4	87.87	83.35	84.84
Hybrid	2	77.74	77.94	77.83
	4	**87.46**	**86.93**	**87.18**

Table 1: Overall Result of Sarcasm Detection

The class-wise results presented in Table 2 indicate that systems based on 4 qubits (QLSTM, Voting, Hybrid) generally achieved higher F1 scores for identifying sarcastic sentences compared to their 2-qubit counterparts. In particular, the 4-qubit Hybrid framework excelled at recognizing both sarcastic sentences (F1 score: 83.93) and non-sarcastic sentences (F1 score: 90.43), outperforming the classical LSTM.

Framework	Qubit	Non-Sarcastic			Sarcastic		
		P	R	F	P	R	F
LSTM	-	84.91	94.74	89.55	88.64	70.91	78.79
QLSTM	2	80.61	83.16	81.87	69.23	65.45	67.29
	4	89.47	71.58	79.53	63.51	85.45	72.87
Voting	2	84.11	94.74	89.11	88.37	69.09	77.55
	4	85.05	**95.79**	90.10	90.70	70.91	79.59
Hybrid	2	84.04	83.16	83.60	71.43	72.73	72.07
	4	**91.40**	89.47	**90.43**	**82.46**	**85.45**	**83.93**

Table 2: Class-wise Result of Sarcasm Detection

Claim Identification: Table 3 displays the claim identification results. Unlike sarcasm detection, the 2-qubit voting framework achieved the highest overall F1-score of 70.04%. The standalone QLSTM and hybrid frameworks did not surpass the performance of the classical LSTM, which achieved an F1-score of 68.05% for this task.

Analyzing the class-wise results (see Table 4), the classical LSTM model achieved the highest F1 score for 'Claim' sentences at 76.92. Meanwhile, the 2-qubit Voting framework demonstrated the highest precision for 'Claim' sentences at 74.66, as well as the highest F1 score for 'No Claim' sentences at 64.17. Notably, the 2-qubit Hybrid framework recorded the highest recall for 'Claim' sentences at 83.48. This suggests a complex interplay of strengths among the various models employed for claim identification.

Framework	Qubit	Precision	Recall	F1-score
LSTM	—	69.16	67.75	68.05
QLSTM	2	63.04	63.06	61.14
	4	59.81	60.58	59.80
Voting	2	**70.23**	**69.90**	**70.04**
	4	68.66	67.08	67.37
Hybrid	2	54.92	55.05	54.80
	4	67.39	65.21	65.40

Table 3: Overall Result of Claim Identification

Framework	Qubit	No Claim			Claim		
		P	R	F	P	R	F
LSTM	-	**68.12**	57.32	62.25	73.08	81.20	**76.92**
QLSTM	2	52.00	**73.98**	61.07	74.09	52.14	61.20
	4	51.82	57.72	54.62	67.80	62.39	64.99
Voting	2	65.81	62.60	**64.17**	**74.66**	77.21	75.91
	4	66.16	53.25	59.01	71.18	80.91	75.73
Hybrid	2	57.97	32.52	41.67	63.83	**83.48**	72.35
	4	65.38	48.37	55.61	69.40	82.05	75.20

Table 4: Class-wise Result of Claim Identification

6 Conclusion

This paper examines quantum LSTM (QLSTM)-based frameworks for sarcasm detection and claim identification, comparing their performance with that of classical LSTM. Our primary goal was to analyze the efficacy of QML in these challenging NLP tasks.

For sarcasm detection, the hybrid QLSM-classical LSTM framework, particularly with four qubits, significantly outperformed both standalone classical and QLSTM models, demonstrating the benefit of quantum-classical integration for identifying sarcastic and non-sarcastic content with higher results.

For claim identification, the 2-qubit voting-based framework achieved the best overall F1-score. While the classical LSTM achieved the highest F1 score for 'Claim' sentences, and the 2-qubit hybrid framework had the highest recall for 'Claim' sentences, this task presented a more complex performance landscape across models.

In future work, we plan to validate these findings on larger datasets, explore architectures with more qubits, and investigate other quantum models, such as Quantum-GRU or Quantum-Transformer, to enhance performance further.

Acknowledgement

This work was supported by Defence Research and Development Organisation (DRDO), New Delhi, under the project “Claim Detection and Verification using Deep NLP: an Indian perspective”.

References

Ivano Basile and Fabio Tamburini. 2017. Towards quantum language models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1840–1849, Copenhagen, Denmark. Association for Computational Linguistics.

Samuel Yen-Chi Chen, Shinjae Yoo, and Yao-Lung L. Fang. 2020. Quantum long short-term memory.

Bob Coecke, Giovanni de Felice, Konstantinos Meichanetzidis, and Alexis Toumi. 2020. Foundations for near-term quantum natural language processing.

Srinjoy Ganguly, Sai Nandan Morapakula, and Luis Miguel Pozo Coronado. 2022. Quantum natural language processing based sentiment analysis using lambeq toolkit. In *2022 Second International Conference on Power, Control and Computing Technologies (ICPC2T)*, pages 1–6.

Raffaele Guarasci, Giuseppe De Pietro, and Massimo Esposito. 2022. Quantum natural language processing: Challenges and opportunities. *Applied Sciences*, 12(11).

Laszlo Gyongyosi and Sandor Imre. 2019. A survey on quantum computing technology. *Comput. Sci. Rev.*, 31(C):51–71.

Diederik P. Kingma and Jimmy Ba. 2017. Adam: A method for stochastic optimization.

Ziming Li, Yan Zhou, Yaxin Liu, Fuqing Zhu, Chuanpeng Yang, and Songlin Hu. 2023. QAP: A quantum-inspired adaptive-priority-learning model for multimodal emotion recognition. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12191–12204, Toronto, Canada. Association for Computational Linguistics.

Yaochen Liu, Yazhou Zhang, Qiuchi Li, Benyou Wang, and Dawei Song. 2021. What does your smile mean? jointly detecting multi-modal sarcasm and sentiment using quantum probability. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 871–880, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Abhijit Mishra, Diptesh Kanojia, and Pushpak Bhattacharyya. 2016. Predicting readers’ sarcasm understandability by modeling gaze behavior. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, page 3747–3753. AAAI Press.

Niels Neumann, Frank Phillipson, and Richard Versluis. 2019. Machine learning in the quantum era. *Digitale Welt*, 3(2):24–29.

Shyambabu Pandey, Pankaj Dadure, Morrel Nunsanga, and Dr. Partha Pakray. 2022. Parts of speech tagging towards classical to quantum computing. pages 1–6.

Shyambabu Pandey and Partha Pakray. 2023. Bi-quantum long short-term memory for part-of-speech tagging. In *Proceedings of the 20th International Conference on Natural Language Processing (ICON)*, pages 301–307, Goa University, Goa, India. NLP Association of India (NLPAI).

Arpan Phukan and Asif Ekbal. 2023. QeMMA: Quantum-enhanced multi-modal sentiment analysis. In *Proceedings of the 20th International Conference on Natural Language Processing (ICON)*, pages 815–821, Goa University, Goa, India. NLP Association of India (NLPAI).

Wenbo Qiao, Peng Zhang, and ZengLai Ma. 2024. A quantum-inspired matching network with linguistic theories for metaphor detection. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1435–1445, Torino, Italia. ELRA and ICCL.

R. L. Rivest, A. Shamir, and L. Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126.

Sara Rosenthal and Kathleen McKeown. 2012. Detecting opinionated claims in online discussions. In *2012 IEEE Sixth International Conference on Semantic Computing*, pages 30–37.

Fariska Z. Ruskanda, Muhammad Rifat Abiwardani, Muhammad Akram Al Bari, Kinantan Arya Bagaspati, Rahmat Mulyawan, Infall Syafalni, and Harashta Tatimma Larasati. 2022. Quantum representation for sentiment classification. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 67–78.

Fariska Zakhralativa Ruskanda, Muhammad Rifat Abiwardani, Infall Syafalni, Harashta Tatimma Larasati, and Rahmat Mulyawan. 2023. Simple sentiment analysis ansatz for sentiment classification in quantum natural language processing. *IEEE Access*, 11:120612–120627.

Jinjing Shi, Zhenhuan Li, Wei Lai, Fangfang Li, Ronghua Shi, Yanyan Feng, and Shichao Zhang. 2023. Two End-to-End Quantum-Inspired Deep Neural Networks for Text Classification . *IEEE Transactions on Knowledge & Data Engineering*, 35(04):4335–4345.

Peter W. Shor. 1997. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509.

Riccardo Di Sipio, Jia-Hong Huang, Samuel Yen-Chi Chen, Stefano Mangini, and Marcel Worring. 2021. The dawn of quantum natural language processing.

Fabio Tamburini. 2019. A quantum-like approach to word sense disambiguation. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 1176–1185, Varna, Bulgaria. INCOMA Ltd.

Charles M. Varmantchaonala, Jean Louis K. E. Fendji, Julius Schöning, and Marcellin Atemkeng. 2024. Quantum natural language processing: A comprehensive survey. *IEEE Access*, 12:99578–99598.

Dominic Widdows, Willie Aboumrad, Dohun Kim, Sayonee Ray, and Jonathan Mei. 2024. Quantum natural language processing. *KI - Künstliche Intelligenz*.

Sixuan Wu, Jian Li, Peng Zhang, and Yue Zhang. 2021. Natural language processing meets quantum physics: A survey and categorization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3172–3182, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Wenduan Xu, Stephen Clark, Douglas Brown, Gabriel Matos, and Konstantinos Meichanetzidis. 2024. Quantum recurrent architectures for text classification. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18020–18027, Miami, Florida, USA. Association for Computational Linguistics.

Kehuan Yan, Peichao Lai, and Yilei Wang. 2024. Quantum-inspired language model with lindblad master equation and interference measurement for sentiment analysis. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2112–2121, Mexico City, Mexico. Association for Computational Linguistics.

Yazhou Zhang, Qiuchi Li, Dawei Song, Peng Zhang, and Panpan Wang. 2019. Quantum-Inspired Interactive Networks for Conversational Sentiment Analysis. *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5436–5442.