

A Hybrid Classical-Quantum Fine Tuned BERT for Text Classification

Abu Kaisar Mohammad Masum¹, Naveed Mahmud², M. Hassan Najafi³, Sercan Aygun¹

¹School of Computing and Informatics, University of Louisiana at Lafayette, Lafayette, LA, USA

² College of Engineering and Science, Florida Institute of Technology, Melbourne, FL

³ Electrical, Computer, and Systems Engineering Department, Case Western Reserve University, Cleveland, OH, USA
c00591145@louisiana.edu, nmahmud@fit.edu, najafi@case.edu, sercan.aygun@louisiana.edu

Abstract

Fine-tuning BERT for text classification can be computationally challenging and requires careful hyper-parameter tuning. Recent studies have highlighted the potential of quantum algorithms to outperform conventional methods in machine learning and text classification tasks. In this work, we propose a hybrid approach that integrates an n -qubit quantum circuit with a classical BERT model for text classification. We evaluate the performance of the fine-tuned classical-quantum BERT and demonstrate its feasibility as well as its potential in advancing this research area. Our experimental results show that the proposed hybrid model achieves performance that is competitive with, and in some cases better than, the classical baselines on standard benchmark datasets. Furthermore, our approach demonstrates the adaptability of classical-quantum models for fine-tuning pre-trained models across diverse datasets. Overall, the hybrid model highlights the promise of quantum computing in achieving improved performance for text classification tasks.

Introduction

Pre-trained language models such as BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al. 2018), have demonstrated exceptional performance across various language comprehension challenges (Sun et al. 2019). The BERT model effectively learns universal language representations and can perform Natural Language Processing (NLP) tasks such as text classification even without human supervision (Qasim et al. 2022; Garg and Ramakrishnan 2020; Jin et al. 2020). However, fine-tuning BERT for specific NLP tasks requires high computational resources and expensive inference time (Rogers, Kovaleva, and Rumshisky 2021; Kovaleva et al. 2019; Zhang et al. 2020). Usually, fully connected multi-layer perceptron (MLP) networks with many hyper-parameters are required for fine-tuning.

Recently, quantum computing (QC) has been increasingly recognized for its potential impact on conventional machine learning (ML), leading to the development of Quantum Machine Learning (QML). Recent advancements in this field include hybrid classical-quantum methods for transfer learning (Mari et al. 2020), text classification (Ardeshir-Larijani and Nasiri Fatmehsari 2024), and speech recognition (Qi and

Tejedor 2022). These studies show how quantum features like *superposition* and *entanglement* can help solve various ML problems.

In this work, we propose a hybrid classical-quantum approach for classifying texts. We use an n -qubit quantum circuit to fine-tune a classical BERT model. We perform a comprehensive evaluation of the proposed method for text classification on a variety of popular datasets. The BERT model was selected for its outstanding performance in understanding the language context bidirectionally. It was pre-trained with vast text data, making it adaptable for tasks with minimum data labels. BERT's flexibility, availability, and transfer learning capabilities suit our proposed hybrid classical-quantum method. The proposed approach entails using pre-trained BERT embeddings and combining them with quantum techniques to improve the overall performance of text classification tasks. This approach leverages classical NLP and QC to improve accuracy and efficiency in handling complex text information.

While current quantum hardware is not yet fault-tolerant enough to achieve actual computational speedup, we perform experiments with quantum circuit simulations to validate the proposed hybrid approach. These experiments yield classification accuracy scores equal to or better than those achieved by conventional fine-tuning techniques. This research contributes to several key areas:

- Integrating quantum circuits with classical BERT models for text classification
- Evaluating performance across various quantum circuit configurations
- Highlighting the potential of leveraging QC in language models for future advancements.

This article is organized as follows: Section II gives an overview of QC Fundamentals, Quantum Gates, Variational Quantum Circuits, and Quantum ML Literature. Section III presents the proposed approach for hybrid classical-quantum text classification. Section IV describes the experimental setup, including the datasets used, model training procedures, and key parameters. It further discusses experimental results, comparing the performance of the proposed model against classical counterparts across various datasets. Finally, Section V summarizes our findings, highlights limitations, and suggests directions for future work in the intersection of QML.

Background

Quantum Computing Fundamentals

Two distinctive features of QC are *superposition* and *entanglement* (Preskill 2018; Jozsa 1997). A qubit is characterized by the Boolean states 0 and 1 within a quantum system, represented by a normalized quantum state $|0\rangle$ and $|1\rangle$ (Schumacher 1995). In computational terms, any state formed by the combination of these two states is referred to as a superposition denoted by $\alpha|0\rangle + \beta|1\rangle$, where α and β satisfy the condition $|\alpha|^2 + |\beta|^2 = 1$. However, entanglement is a special correlation between qubits where the state of one qubit (e.g., *qubit A*) instantly affects another (e.g., *qubit B*), in any case of distance. The combined wave function representing entangled particles is denoted by $\Psi_{AB} = (1/\sqrt{2})(\Psi_{1A}\Psi_{2B} - \Psi_{2A}\Psi_{1B})$.

Quantum Gates

Quantum gates are the fundamental components of quantum circuits are like classical logic gates in classical computing. They operate on qubits and manipulate their quantum states by leveraging properties like superposition and entanglement. A quantum circuit is a sequence of quantum gates (Kais 2014). These gates are used for data encoding from classical to quantum; for example, the Hadamard gate (H gate) creates superposition states. A rotation gate like $R_Y(\theta)$ gate is a parameterized gate that rotates the qubit along the y -direction by the specified angular parameter θ . Two-qubit gates like the CNOT gate have a control qubit and a target qubit. The NOT (inversion) operation is implemented on the target qubit based on the value of the control qubit. The CNOT is particularly useful for creating an entanglement effect between qubits (Chen et al. 2020).

Variational Quantum Circuits

A Variational Quantum Circuit (VQC) allows for the exploration of different quantum states and state transformations by varying the parameters of its gates. Measurements extract classical information from the final quantum state and are then used to evaluate a cost function. The parameters are then varied or *trained* to minimize the cost function. Hence, parameterized VQCs can be trained for ML tasks (Saxena and Nigam 2022; Rizvi et al. 2023), analogous to the training of neural networks. VQCs have a relatively smaller number of parameters compared to classical neural networks, making them computationally efficient and easier to optimize (Chen et al. 2022; Ajagekar and You 2024). In Figure 1, a VQC example is illustrated.

Quantum ML Literature

QC and its applications in ML span areas such as computer vision and language processing. In recent years, challenges in document classification have led to the growth of conventional neural network-based models alongside the rise of quantum ML systems. Techniques like BERT stand out within the established world of pre-trained models, proving their usefulness in achieving cutting-edge performance across varied tasks such as NLP and natural language generation

models (Devlin et al. 2018; Adhikari et al. 2019; Qu et al. 2019; Zhang et al. 2019; Souza, Nogueira, and Lotufo 2019). On the other hand, the rapid advancement of QC has revealed different aspects demonstrating quantum benefits in quantum ML (Schuld, Sinayskiy, and Petruccione 2015). These include more complex feature representation and improved security for ML model parameters. This advancement is based on the use of variational quantum circuits, which have demonstrated the competitive performance of text classification (Yang et al. 2022).

In addition, several ML tasks, such as Omar et al. (Omar and Abd El-Hafeez 2023), have used quantum ML to classify Arabic text data, Quantum self-attention (Li, Zhao, and Wang 2024; Zhang, Zhao, and Chen 2024) for text classification, and hybrid transfer learning (Ardeshir-Larijani and Nasiri Fatmehsari 2024). Li et al. (Li, Zhao, and Wang 2022) recently proposed an attention-based quantum natural language processing (QNLP) technique. Other tasks of QNLP include ensemble learning by Bouakba et al. (Bouakba and Belhadef 2023), Quantum temporal convolution by Yang et al. (Yang et al. 2022), the exploration of quantum algorithms by Zeng et al. (Zeng and Coecke 2016), and Lorenz et al. (Lorenz et al. 2023).

Quantum-inspired ML has emerged as a promising direction for enhancing classical ML models (Felser et al. 2021). Achieving optimal performance in these approaches requires smooth input preparation to achieve the best performance (Zhao et al. 2021). For example, quantum-inspired methods have been applied to cancer classification with encouraging results (Pomarico et al. 2021). Recently, in 6G network research, quantum-inspired models have demonstrated superior capabilities in data processing and training (Duong et al. 2022). A wide range of applications, including optimization, binary classification, protein design, and clonogenic assay evaluation (Li et al. 2017; Tiwari and Melucci 2018; Panizza et al. 2024; Sergioli et al. 2021), now leverage quantum principles to enhance computational efficiency and accuracy, offering more effective solutions to complex problems.

Proposed Approach

In this work, we target a text classification problem using quantum ML. Our proposed approach for text classification incorporates the classical BERT model and quantum layer. The big picture of the overall proposal is shown in Figure 2. The goal is to leverage the power of QC to enhance the performance of traditional text classification tasks. The architecture consists of three main layers:

- **Input Classical Layer:** The first layer processes the input text data using a pre-trained language model like BERT. The output of this layer is a sequence of embeddings (Feng et al. 2020) representing the semantic meaning of each word in the text.
- **Quantum Layer:** The second layer is where the quantum computations take place. This layer includes a feature encoding circuit (Schuld and Killoran 2019) with a quantum variational circuit. The feature encoding circuit maps the classical text embeddings from the previous layer into

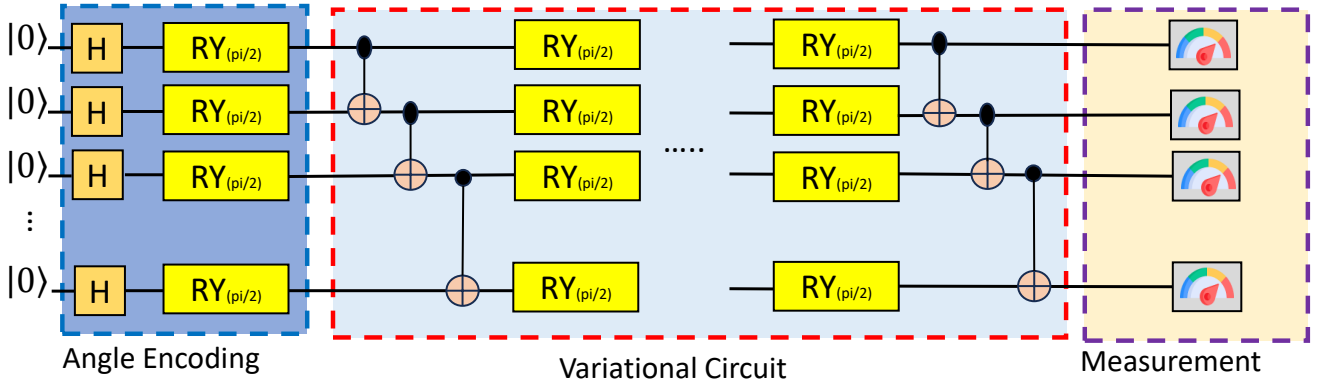


Figure 1: Angle encoding, variational circuit, and the measurement steps in an example VQC.

a quantum state. The quantum variational circuit then applies a series of parameterized quantum gates to manipulate this state.

- **Output Classical Layer:** This final layer measures the quantum state to fine-tune resulting measurement outcomes into a classification probability distribution. The most likely class is chosen as the predicted label for the input text.

Classical BERT Model

In the classical BERT, a fully connected layer links the last hidden state of the model (first output) to the model’s input tokens. The output y_n is produced by applying a non-linear activation function f to the weighted sum of the input x_n , leveraging a set of θ_n parameters, corresponding to each layer’s weights and biases.

$$L_{n_0 \rightarrow n_1} : y_n = f(x_n, \theta_n) \quad (1)$$

where, $L_{n_0 \rightarrow n_1}$ indicates a change in dimensions from n_0 to n_1 . The BERT model we used returns a $n_1 = 768$ dimensional vector. For fine-tuning the model output for classifying texts, we integrate a quantum layer with the BERT output.

Fine-tuning with Quantum Layer

In our approach, we introduce an additional pooling layer, $L_{n_1 \rightarrow n_q}$, prior to the quantum layer to reduce the feature dimensionality of the BERT output from 768 to n_q . Each data feature d_i in the resulting layer is encoded and represented by a single qubit through the following transformation:

$$U_f |0\rangle = \bigotimes_{i=1}^{n_q} \cos(d_i) |0\rangle + \sin(d_i) |1\rangle \quad (2)$$

where U_f represents the quantum feature encoding circuit (see Figure 2). U_f is realized using R_y rotation gates with feature data d_i from the previous classical layer used as the rotation angles. The resulting state is denoted as $|x\rangle$. We apply a shallow depth, variational quantum circuit (Cerezo et al. 2021; Yang et al. 2022) with parameter set θ_n that maps $|x\rangle$ to $|y\rangle$.

$$\mathcal{L} : |x\rangle \rightarrow |y\rangle = U_{vqc} |x\rangle \quad (3)$$

where U_{vqc} represents the VQC consisting of trainable parameters. The structure of the VQC circuit is shown in Figure 2. The first layer of the circuit consists of H gates that create a superposition state, allowing the circuit to sample multiple solutions in parallel for the next layer, which contains the trainable parameterized rotation gates. This speeds up the optimization process of the variational parameters during the training phase. For parameterized rotation, we use R_y gates that affect the real amplitudes of the quantum state. The next layer of CNOT gates incorporates entanglement in the circuit, amplifying the model’s ability to capture complex correlations in the encoded data. In a fully entangled state, all qubits in the system are maximally correlated with each other. Now, to extract classical output y_{out} as a vector from the quantum circuit, we measure the expectation values of qubits. We define this as a measurement layer of the quantum circuit, thus converting the final quantum state into a classical vector.

$$M : |y\rangle \rightarrow y_{out} = \langle \hat{O} \rangle y \quad (4)$$

Here $\langle \hat{O} \rangle y$ represents the expectation value of an observable operator $\langle \hat{O} \rangle$ in the state $|\psi_{out}\rangle$. After measurement, we apply $L_{n_q \rightarrow 2}$, a classical pooling layer that transforms n features into two features that are used in the final prediction. The final prediction is compared with data labels to evaluate a mean-square-error loss function. Based on the loss, we calculate quantum gradients to optimize the VQC parameters to minimize the loss.

Experiment Setup and Results Discussion

Datasets

We collected five different datasets from publicly available sources for text classification: 5000 IMDb data from <https://huggingface.co/datasets/imdb>, 5572 email Spam data from <https://www.surgehq.ai/datasets/email-spam-dataset>, 6000 SST from <https://nlp.stanford.edu/sentiment/treebank.html>, 10,000 Yelp data from <https://www.yelp.com/dataset>, and 11,541 Twitter data from <https://huggingface.co/datasets/osanseviero/twitter-airline-sentiment>.




Figure 2: Overview of our approach for a hybrid classical-quantum BERT model.

Model Training

For classical-quantum BERT training, we leverage the NVIDIA T4 GPU. The subsequent experiments are run using PyTorch 2.1.0+cu118, which serves as the foundational framework. We use transformer version 4.35.2 for BERT, and Scikit-learn 1.2.2 for model selection and classification report analysis. In the quantum part, we use pennylane 0.33.1 with the default simulator to create and run quantum circuits.

Model Parameters

Table 1 outlines the crucial classical and quantum parameters employed in our hybrid BERT model settings. These carefully chosen settings were instrumental in achieving optimal performance across diverse datasets.

Our classical model includes 10 epochs, a batch size of 32, a maximum sequence length of 25, and a learning rate of 1×10^{-3} . These parameters balance training speed, stability, and model generalization. For our experiments, we used 2 to 10 qubits with a circuit depth of 4, where we leveraged H gates, Y-rotation gates, and CNOT gates to manipulate the qubits effectively. Moreover, we employed angle encoding for data representation, while quantum entanglement was used to capture complex correlations within the data. A quantum weight of 0.01 is used to integrate quantum features without overwhelming the classical model. Again, the Pauli-Z feature map is used to enhance the model’s ability to learn nonlinear relationships, and a default quantum simulator ensures accurate simulations of our model. In summary, these parameters create a robust hybrid architecture that integrates the strengths of classical and QC to ensure accurate model performance.

Results Comparison

In the experiments, we compare the proposed classical-quantum BERT to a simple classical BERT (baseline), and a Multilayer Perceptron (MLP) fine-tuned BERT across IMDB, Spam, SST, Yelp, and Twitter datasets. An in-depth analysis across the datasets reveals that the proposed classical-quantum BERT consistently performs better than or equal to its counterparts. The table summary focuses on the specific

strengths and weaknesses of the implemented models in text classification tasks.

As shown in Table 2, our classical-quantum BERT demonstrates consistently high performance, particularly on datasets like IMDB, Spam, and Twitter. It achieves an F1 score of 0.64 for both training and validation on the IMDB dataset, which is marginally better than both the classical BERT and MLP models, scoring 0.63 in each. On the Spam dataset, all three models achieve high F1 scores of 0.95 for both training and validation, indicating that this task is better for each model’s architecture. However, the proposed BERT outperforms its counterparts for the Twitter dataset, getting an F1 score of 0.86 for both training and validation compared to the classical BERT’s score of 0.83. The MLP’s score is 0.82 F1 score for training and 0.83 for validation on the same dataset.

In contrast, on the SST dataset, both the classical-quantum BERT and classical BERT achieve the same F1 scores of 0.82 for both training and validation. The MLP model for the dataset underperforms, with scores of 0.70 for training and 0.72 for validation F1 score. This result indicates that while the quantum-enhanced model does not always exceed the performance of its classical counterpart, it still remains competitive. For the Yelp dataset, the classical-quantum BERT shows a slight advantage. It scores 0.79 for both training and validation F1 score, compared to the classical BERT’s score of 0.78 and the MLP’s 0.78 for training, as well as 0.75 for validation.

Comparative Analysis

We compare the accuracy of the proposed classical-quantum BERT model across various datasets, qubit values, and evaluation with existing models (see Table 3). This table suggests the potential of the proposed model in achieving high accuracy for various text classification tasks with the flexibility of adjusting qubit values.

The Lambeq model (Kartsaklis et al. 2021) using 3 qubits achieves an accuracy of 0.93 on the Spam dataset but only 0.59 on the SST dataset. The hybrid QBERT (Adhikari et al. 2019) that uses 5 qubits achieves a higher accuracy of 0.97 on the Spam dataset, showing the effectiveness of increasing the complexity of the quantum circuit. Adapting QBERT (Li

Quantum Settings		Training Settings	
Parameter	Value	Parameter	Value
Qubits Range	2–10	Epoch	10
Circuit Depth	4	Batch Size	32
Quantum Weight	0.01	Max Sequence Length	25
Quantum Gates	H, Y-rotation, CNOT	Learning Rate	1×10^{-3}
Encoding Method	Angle Encoding		
Entanglement	Yes		
Feature Map	Pauli-Z		
Simulator	Default		

Table 1: Key Parameters for Hybrid Classical-Quantum BERT.

Dataset	BERT _{Classical-Quantum}		BERT _{Classical}		MLP _{tf-idf}	
	Train F1	Val. F1	Train F1	Val. F1	Train F1	Val. F1
IMDb	0.64	0.64	0.63	0.63	0.63	0.63
Spam	0.95	0.95	0.95	0.95	0.95	0.95
Twitter	0.86	0.86	0.83	0.83	0.82	0.83
SST	0.82	0.82	0.82	0.82	0.70	0.72
Yelp	0.79	0.79	0.78	0.78	0.78	0.75

Table 2: Comparison of performance metrics (Train F1 and Validation F1) for three different models (BERT_{Classical-Quantum}, BERT_{Classical}, and MLP_{tf-idf}) across multiple datasets (IMDb, Spam, Twitter, SST, and Yelp). The table presents the training and validation F1 scores for each model and dataset. The values in the table represent the F1 scores, with important scores highlighted. Changes in performance can be observed across the models and datasets. BERT_{Classical-Quantum} represents the proposed quantum-enhanced version of the BERT model, while BERT_{Classical} denotes the classical BERT model. MLP_{tf-idf} refers to a Multilayer Perceptron model using tf-idf features.

et al. 2023) with 10 qubits shows a significant improvement in the SST dataset with an accuracy of 0.89, indicating the benefits of a quantum model for more text datasets.

Our proposed classical-quantum BERT model demonstrates flexibility by using a range of 2 to 10 qubits and consistently performs well across all datasets. Specifically, it achieves a mean accuracy of 0.67 ± 0.04 on IMDb, 0.96 ± 0.03 on Spam, 0.86 ± 0.02 on Twitter, 0.82 ± 0.02 on SST, and 0.79 ± 0.01 on Yelp. These results suggest that our model maintains a competitive edge, either outperforming or matching the accuracy of existing models across all tested datasets.

Training Time

We estimate the training time for the hybrid models as depicted in the line graph in Figure 3. This graph provides insights into the training time evolution of a classical-QC model over a variety of datasets with varying qubit values. The x-axis shows the number of qubits used in the classical-quantum model, while the y-axis represents the training duration in seconds. The consistent training times suggest that the model’s performance is relatively stable and insensitive to changes in the number of qubits. IMDb dataset shows a training time of 2760 seconds with 2 qubits, which slightly increases to 3900 seconds with 10 qubits. Similarly, Spam and SST datasets display modest increases in training times, indicating minimal sensitivity to the number of qubits used. The graph reveals a notable increase in training times for the Yelp and Twitter datasets as the number of qubits rises. For




Figure 3: The graph shows training time variations in a classical-QC model across datasets (IMDb, Spam, SST, Yelp, Twitter) with different qubit values. IMDb, Spam, and SST maintain stable times, while Yelp and Twitter see increased times with more qubits. Key values: IMDb (2 qubits: 2760s, 10 qubits: 3900s), Spam (2 qubits: 2880s, 10 qubits: 4140s), SST (2 qubits: 2988s, 10 qubits: 4740s), Yelp (2 qubits: 3126s, 10 qubits: 8612s), Twitter (2 qubits: 3480s, 10 qubits: 8910s). Insights into qubit count and training time relationships reveal dataset-specific differences in the classical-QC model.

Yelp data, training time starts at 3,126 seconds with 2 qubits and escalates significantly to 8,612 seconds with 10 qubits. A similar trend is observed for the Twitter dataset, where the training time jumps from 3,480 seconds at 2 qubits to 8,910 seconds at 10 qubits. This suggests that these datasets may involve more complex patterns or higher data dimensionality, requiring longer training times as quantum circuit complexity increases.

Method	Qubits	Dataset	Accuracy
Lambeq (Kartsaklis et al. 2021)	3	Spam	0.93
		SST	0.59
Hybrid QBERT (Ardeshir-Larijani and Nasiri Fatmehsari 2024)	5	Spam	0.97
Adapting QBERT (Li et al. 2023)	10	SST	0.89
		IMDB	0.67 ± 0.04
Proposed	2-10	Spam	0.96 ± 0.03
		Twitter	0.86 ± 0.02
		SST	0.82 ± 0.02
		Yelp	0.79 ± 0.01

Table 3: Performance comparison of proposed model with related work for varying qubits and diverse datasets. Includes Lambeq (Kartsaklis et al. 2021), Hybrid QBERT (Ardeshir-Larijani and Nasiri Fatmehsari 2024), Adapting QBERT (Li et al. 2023), and Proposed (2-10). Reports accuracy for Spam, SST, IMDB, Twitter, and Yelp. For the proposed model, mean accuracy with standard deviations (\pm) is provided and exhibits comparable performance across data sets with respect to related work.

Conclusion & Limitations

This work evaluates the interaction between pre-trained language models and quantum computing in Quantum Machine Learning (QML). We combined a pre-trained BERT model with a quantum variational circuit for text classification. Despite scalability, computing resource constraints, and longer training time, our study highlights the importance of addressing potential quantum advantage over classical models and achieves acceptable accuracy results across a variety of datasets. The results indicate that while the hybrid model demonstrates promising performance, it also reveals areas where further optimization could be beneficial, particularly in improving the computational efficiency and reducing the training time. The hybrid classical-quantum BERT model has limits besides these positive results. Scalability is limited by quantum devices, and the integrated circuit takes substantially longer to train than classical models. However, this approach is an important step in leveraging classical-quantum collaboration for improved document classification, paving the way for higher efficiency when dealing with complicated textual content.

Acknowledgments

This work is supported by the National Science Foundation under Grant No. 2019511 and generous gifts from NVIDIA.

References

Adhikari, A.; Ram, A.; Tang, R.; and Lin, J. 2019. Docbert: Bert for document classification. *arXiv preprint arXiv:1904.08398*.

Ajagekar, A.; and You, F. 2024. Variational quantum circuit based demand response in buildings leveraging a hybrid quantum-classical strategy. *Applied Energy*, 364: 123244.

Ardeshir-Larijani, E.; and Nasiri Fatmehsari, M. M. 2024. Hybrid classical-quantum transfer learning for text classification. *Quantum Machine Intelligence*, 6(1): 19.

Bouakba, Y.; and Belhadef, H. 2023. Ensemble learning based quantum text classifiers. In *European Conference on Advances in Databases and Information Systems*, 407–414. Springer.

Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S. C.; Endo, S.; Fujii, K.; McClean, J. R.; Mitarai, K.; Yuan, X.; Cincio, L.; et al. 2021. Variational quantum algorithms. *Nature Reviews Physics*, 3(9): 625–644.

Chen, S. Y.-C.; Huang, C.-M.; Hsing, C.-W.; Goan, H.-S.; and Kao, Y.-J. 2022. Variational quantum reinforcement learning via evolutionary optimization. *Machine Learning: Science and Technology*, 3(1): 015025.

Chen, S. Y.-C.; Huang, C.-M.; Hsing, C.-W.; and Kao, Y.-J. 2020. Hybrid quantum-classical classifier based on tensor network and variational quantum circuit. *arXiv preprint arXiv:2011.14651*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Duong, T. Q.; Ansere, J. A.; Narottama, B.; Sharma, V.; Dobre, O. A.; and Shin, H. 2022. Quantum-inspired machine learning for 6G: fundamentals, security, resource allocations, challenges, and future research directions. *IEEE Open Journal of Vehicular Technology*, 3: 375–387.

Felser, T.; Trenti, M.; Sestini, L.; Gianelle, A.; Zuliani, D.; Lucchesi, D.; and Montangero, S. 2021. Quantum-inspired machine learning on high-energy physics data. *npj Quantum Information*, 7(1): 111.

Feng, F.; Yang, Y.; Cer, D.; Arivazhagan, N.; and Wang, W. 2020. Language-agnostic BERT sentence embedding. *arXiv preprint arXiv:2007.01852*.

Garg, S.; and Ramakrishnan, G. 2020. Bae: Bert-based adversarial examples for text classification. *arXiv preprint arXiv:2004.01970*.

Jin, D.; Jin, Z.; Zhou, J. T.; and Szolovits, P. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 8018–8025.

Jozsa, R. 1997. Entanglement and quantum computation. *arXiv preprint quant-ph/9707034*.

Kais, S. 2014. Introduction to quantum information and computation for chemistry. *Quantum Information and Computation for Chemistry*, 1–38.

Kartsaklis, D.; Fan, I.; Yeung, R.; Pearson, A.; Lorenz, R.; Toumi, A.; de Felice, G.; Meichanetzidis, K.; Clark, S.; and

- Coecke, B. 2021. lambeq: An efficient high-level python library for quantum NLP. *arXiv preprint arXiv:2110.04236*.
- Kovaleva, O.; Romanov, A.; Rogers, A.; and Rumshisky, A. 2019. Revealing the dark secrets of BERT. *arXiv preprint arXiv:1908.08593*.
- Li, G.; Zhao, X.; and Wang, X. 2022. Quantum self-attention neural networks for text classification. *arXiv preprint arXiv:2205.05625*.
- Li, G.; Zhao, X.; and Wang, X. 2024. Quantum self-attention neural networks for text classification. *Science China Information Sciences*, 67(4): 142501.
- Li, Q.; Wang, B.; Zhu, Y.; Lioma, C.; and Liu, Q. 2023. Adapting Pre-trained Language Models for Quantum Natural Language Processing. *arXiv preprint arXiv:2302.13812*.
- Li, Y.; Lu, G.; Zhou, L.; and Jiao, L. 2017. Quantum inspired high dimensional hyperparameter optimization of machine learning model. In *2017 International Smart Cities Conference (ISC2)*, 1–6. IEEE.
- Lorenz, R.; Pearson, A.; Meichanetzidis, K.; Kartsaklis, D.; and Coecke, B. 2023. QNLP in practice: Running compositional models of meaning on a quantum computer. *Journal of Artificial Intelligence Research*, 76: 1305–1342.
- Mari, A.; Bromley, T. R.; Izaac, J.; Schuld, M.; and Killoran, N. 2020. Transfer learning in hybrid classical-quantum neural networks. *Quantum*, 4: 340.
- Omar, A.; and Abd El-Hafeez, T. 2023. Quantum computing and machine learning for Arabic language sentiment classification in social media. *Scientific Reports*, 13(1): 17305.
- Panizza, V.; Hauke, P.; Micheletti, C.; and Faccioli, P. 2024. Protein Design by Integrating Machine Learning with Quantum Annealing and Quantum-inspired Optimization. *arXiv preprint arXiv:2407.07177*.
- Pomarico, D.; Fanizzi, A.; Amoroso, N.; Bellotti, R.; Biafora, A.; Bove, S.; Didonna, V.; Forgia, D. L.; Pastena, M. I.; Tamborra, P.; et al. 2021. A proposal of quantum-inspired machine learning for medical purposes: An application case. *Mathematics*, 9(4): 410.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Qasim, R.; Bangyal, W. H.; Alqarni, M. A.; Ali Almazroi, A.; et al. 2022. A fine-tuned BERT-based transfer learning approach for text classification. *Journal of healthcare engineering*, 2022.
- Qi, J.; and Tejedor, J. 2022. Classical-to-quantum transfer learning for spoken command recognition based on quantum neural networks. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8627–8631. IEEE.
- Qu, C.; Yang, L.; Qiu, M.; Croft, W. B.; Zhang, Y.; and Iyyer, M. 2019. BERT with history answer embedding for conversational question answering. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, 1133–1136.
- Rizvi, S. M. A.; Ulum, M. S.; Asif, N.; and Shin, H. 2023. Neural Networks with Variational Quantum Circuits. In *International Conference on Industrial Networks and Intelligent Systems*, 203–214. Springer.
- Rogers, A.; Kovaleva, O.; and Rumshisky, A. 2021. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8: 842–866.
- Saxena, N.; and Nigam, A. 2022. Performance Evaluation of a Variational Quantum Classifier. In *2022 IEEE 9th Uttar Pradesh Section International Conference on Electrical, Electronics and Computer Engineering (UPCON)*, 1–5. IEEE.
- Schuld, M.; and Killoran, N. 2019. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4): 040504.
- Schuld, M.; Sinayskiy, I.; and Petruccione, F. 2015. An introduction to quantum machine learning. *Contemporary Physics*, 56(2): 172–185.
- Schumacher, B. 1995. Quantum coding. *Physical Review A*, 51(4): 2738.
- Sergioli, G.; Militello, C.; Rundo, L.; Minafra, L.; Torrisi, F.; Russo, G.; Chow, K. L.; and Giuntini, R. 2021. A quantum-inspired classifier for clonogenic assay evaluations. *Scientific Reports*, 11(1): 2830.
- Souza, F.; Nogueira, R.; and Lotufo, R. 2019. Portuguese named entity recognition using BERT-CRF. *arXiv preprint arXiv:1909.10649*.
- Sun, C.; Qiu, X.; Xu, Y.; and Huang, X. 2019. How to fine-tune bert for text classification? In *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*, 194–206. Springer.
- Tiwari, P.; and Melucci, M. 2018. Towards a quantum-inspired framework for binary classification. In *Proceedings of the 27th ACM international conference on information and knowledge management*, 1815–1818.
- Yang, C.-H. H.; Qi, J.; Chen, S. Y.-C.; Tsao, Y.; and Chen, P.-Y. 2022. When bert meets quantum temporal convolution learning for text classification in heterogeneous computing. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8602–8606. IEEE.
- Zeng, W.; and Coecke, B. 2016. Quantum algorithms for compositional natural language processing. *arXiv preprint arXiv:1608.01406*.
- Zhang, H.; Zhao, Q.; and Chen, C. 2024. A light-weight quantum self-attention model for classical data classification. *Applied Intelligence*, 54(4): 3077–3091.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Zhang, T.; Wu, F.; Katiyar, A.; Weinberger, K. Q.; and Artzi, Y. 2020. Revisiting few-sample BERT fine-tuning. *arXiv preprint arXiv:2006.05987*.
- Zhao, Z.; Fitzsimons, J. K.; Rebstrost, P.; Dunjko, V.; and Fitzsimons, J. F. 2021. Smooth input preparation for quantum and quantum-inspired machine learning. *Quantum Machine Intelligence*, 3(1): 14.