Quantum-inspired Language Model with Lindblad Master Equation and Interference Measurement for Sentiment Analysis

Kehuan Yan, Peichao Lai, Yilei Wang * ✉
College of Computer and Data Science, Fuzhou University, Fuzhou, China
yilei@fzu.edu.cn

Abstract

Quantum-inspired models have demonstrated superior performance in many downstream language tasks, such as question answering and sentiment analysis. However, recent models primarily focus on embedding and measurement operations, overlooking the significance of the quantum evolution process. In this work, we present a novel quantum-inspired neural network, LI-QiLM, which integrates the Lindblad Master Equation (LME) to model the evolution process and the interferometry to the measurement process, providing more physical meaning to strengthen the interpretability. We conduct comprehensive experiments on six sentiment analysis datasets. Compared to the traditional neural networks, transformer-based pre-trained models and quantum-inspired models, such as CICWE-QNN and ComplexQNN, the proposed method demonstrates superior performance in accuracy and F1-score on six commonly used datasets for sentiment analysis. Additional ablation tests verify the effectiveness of LME and interferometry.

1 Introduction

The rapid advancement of deep learning has significantly boosted the performance of Natural Language Processing (NLP) tasks. However, a key limitation of deep learning models is their lack of interpretability (Wei et al., 2022). Recently, there has been a promising trend toward harnessing quantum theory to augment expressive capabilities and enhance interpretability (Li et al., 2019). The similarities between quantum systems and natural languages have been widely explored (Zhang et al., 2018b). In linguistics, a word is given different meanings. Such polysemy phenomenon is analogously compared to the superposition state in quantum mechanics. In this sense, sentences are analogous to a multiparticle system. Besides, as particles in quantum mechanics collapse to defined states upon observation, language users assign specific meanings to words or sentences based on contexts. Meanwhile, studies (Bruza et al., 2015; Busemeyer and Wang, 2015) conducted within the realm of cognitive science have demonstrated that human cognition, encompassing sentiment analysis tasks in NLP, often exhibit quantum-like behaviors, which can be more precisely elucidated through quantum theory.

Building upon these foundational ideas, Quantum-inspired Language Models (QiLMs) are actively explored and designed to augment the performance of traditional language models (Zhang et al., 2018a; Li et al., 2019; Lai et al., 2023). A quantum-inspired multimodal sentiment analysis framework named QMSA (Zhang et al., 2018c) applied interference experiments to sentiment analysis by comparing emotion labels to photons and data types to interference slits. Motivated by the complex word embedding, an embedding method called ICWE (Shi et al., 2023) is proposed to use the amplitude and phase of quantum as the real and imaginary part of the word embedding respectively, which improves the feature extraction ability of the model. A quantum-inspired multimodal fusion framework (Li et al., 2021) converts real-valued inputs of different modalities into complex-valued pure quantum states and induces the different modalities to interact in a non-separable way for encoding cross-modal information. Their study showcased the superior performance of this novel approach compared to conventional transformer-based models (Kenton and Toutanova, 2019; Liu et al., 2019). This groundbreaking fusion of domains is expected to make significant contributions to the evolution of quantum machine learning and play a pivotal role in practical use.

The Lindblad Master Equation (LME) (Manzano, 2020), a crucial instrument in the exploration

*Corresponding author

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

of open quantum systems, remains neglected by contemporary models. It stands as an indispensable complement to the Schrödinger equation, providing a highly versatile framework for depicting Markovian dynamics in quantum systems. Specifically, this equation provides a comprehensive portrayal of the temporal evolution of the density matrix, incorporating the interactions with the surrounding environment. Taking advantage of LME, a more appropriate density matrix representation can be learned for language sentences in an end-to-end manner.

Simultaneously, current models usually use projection measurement, which demands an extensive amount of operations. This is because the probability distribution on each basis is derived through statistical averaging, leading to an augmented time expenditure and uncertainty. In contrast, interferometry solely necessitates the utilization of an interferometric operator to imitate the collapse procedure of the superposition state. Consequently, it swiftly attains dependable outcomes with enhanced precision.

In this work, we propose a quantum-inspired language model with Lindblad Master Equation and interferometry (LI-QiLM) to introduce a robust quantum evolution process, a more accurate measurement process, and an improved expressive capacity through the inclusion of additional learnable parameters. Our contributions are:

- We propose the LI-QiLM and detail its pivotal components. Chief among these is the LME module, which emphasizes the quantum evolution process, enhancing both performance and transparency. Additionally, the IM module mitigates uncertainties inherent in quantum systems, bolstering the dependability of model outputs.
- We validate the effectiveness of LI-QiLM on six sentiment analysis datasets and conduct a comprehensive comparison with seven representative models. The results show that the LI-QiLM has 15% improved accuracy compared with traditional models.

We organize the paper as follows. Section 2 presents the fundamentals of quantum theory. Section 3 provides an overview of related works. In Section 4, we introduce the proposed model. Section 5 details the experiments carried out. In Section 6, we draw our conclusions. Additionally, Section 7 discusses the limitations of our work.

2 Preliminary

Quantum computing represents a revolutionary computing paradigm that follows the laws of quantum mechanics to regulate quantum information units to calculate. In the open quantum system, it is imperative to employ the density matrix as a means to accurately portray the system's state. Delving into the system's evolutionary process necessitates the utilization of the quantum master equation. Ultimately, precise quantum measurement techniques allow us to obtain unequivocal results regarding the system's evolution. The specific process is shown in Fig. 1.

2.1 Quantum State

In quantum mechanic, the basic state $|\psi\rangle$, named qubit, is denoted by the Dirac notation as,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{1}$$

where $|\cdot\rangle$ is the bra-ket notation. Note that $|\alpha|^2$ and $|\beta|^2$ represent the probability that the superposition will be $|0\rangle$ and $|1\rangle$, respectively, and satisfy $|\alpha|^2 + |\beta|^2 = 1$.

2.2 Quantum Master Equation

Complex quantum systems defy description by a single Dirac notation. In contrast, one can use the density matrix representation to describe the system by

$$\rho = \sum_{i=1}^{m} p_i |\psi_i\rangle\langle\psi_i|, \tag{2}$$

where $|\psi_i\rangle$ is the state of the subsystems.

In fact, real physical systems inevitably interact with the surrounding environment. Thus, it is necessary to consider the dissipation and decay. Lindblad Master Equation (LME) (Manzano, 2020) is a common method to study the quantum evolution versus time t, denoted by

$$\begin{aligned}\rho(t) = &-\frac{i}{\hbar}[H, \rho(t)] \\ &+ \sum_i \left(L_i \rho(t) L_i^\dagger - \frac{1}{2}\left\{ L_i^\dagger L_i, \rho(t) \right\} \right),\end{aligned} \tag{3}$$

where $\{x, y\} = xy + yx$ is the anticommutator, $X^\dagger$ is the conjugate transpose of X, H is the Hamiltonian of the system and L_i is the operator describing the interaction between the system and the environment. Specifically, the initial term

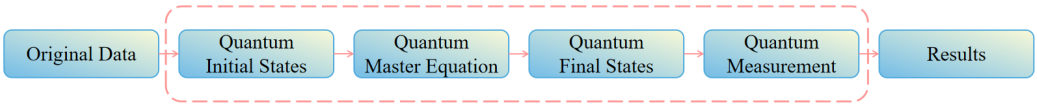

Figure 1: The process of quantum computation.

corresponds to the Schrödinger equation, which delineates the temporal evolution of a quantum system within a sealed, isolated environment. The second term describes the impact of the interaction between the system and the environment. Such interactions result in the dissipation of energy and the degradation of information, causing the evolution of the quantum system to become irreversible.

2.3 Quantum Measurement

Quantum measurement enables the state of a quantum system to collapse from an uncertain superposition to a deterministic result. One commonly used measurement is interferometry, which is used to observe the interaction of quantum states in different paths or branches. For a particular interferometric operator, the expected value can be expressed by

$$\langle \rho \rangle = Tr(\rho H), \tag{4}$$

where Tr means tracing the product of the density matrix and the interferometer operator, i.e. the sum of the diagonal elements of the matrix, and H represents the interferometric operator.

3 Related Works

The development of quantum computing has been ongoing since 1982 (Feynman, 1982). With the surge in machine learning, the integration of machine learning and quantum theory has drawn closer attention. A pivotal landmark in the field of NLP was the introduction of the quantum language model by Sordoni et al. (2013). (Sordoni et al., 2013), which applied the probabilistic framework of quantum theory to information retrieval tasks. This model utilized the density matrix to create a nuanced text representation, showing substantial gains in performance compared to the traditional bag-of-words models (Santner et al., 2003; Spärck Jones, 2004).

Zhang et al. (2018). (Zhang et al., 2018a) introduced the Neural Network Based Quantum-Like Language Model (NNQLM), marking a notable advancement in the field. NNQLM pioneers a unique strategy with the use of a tailored density matrix that ingeniously captures the relationship between questions and answers, significantly improving the representation of sentence similarity. It is seamlessly integrated into a two-dimensional neural network architecture designed for matching. Li et al. (2018). (Li et al., 2018) noticed a striking parallel between the way a word's meaning is influenced by its companions and the behavior patterns of microscopic particles. Drawing from this observation, they proposed a complex embedding mixture network (CE-Mix) that employs word embedding in the complex domain. CE-Mix deduces the significance of phrases using the principle of quantum interference. In evaluations, it surpassed the leading non-quantum models in sentence classification tasks.

Shi et al. (2021). (Shi et al., 2023) introduced the Convolutional Complex-Valued Neural Network Based on Interpretable Complex-Valued Word Embedding (CICWE-QNN). This model utilizes the amplitude and phase of quantum to respectively represent the real and imaginary parts of word embeddings. This approach improves feature extraction capabilities while also guaranteeing the model's interpretability. Zhang et al. (2022). (Zhang et al., 2022) proposed a complex-valued neural network-based QLM (C-NNQLM), which employs an end-to-end methodology for constructing and refining density matrices in a differentiable way. This approach facilitates the incorporation of well-trained word vectors and supervised labels. Furthermore, C-NNQLM incorporates complex value word embedding, with its phase vector adeptly encoding the sequential or positional information of words. Wei et al. (2023). (Lai et al., 2023) presented the Quantum-Inspired Fully Complex-Valued Neutral Network, ComplexQNN. This model employs complex-value CNN throughout the processing pipeline, aligning more closely with quantum computing theory and offering the potential for future acceleration on quantum computers.

In sharp contrast to conventional neural networks, quantum-inspired models present the distinct advantage of offering explanations grounded

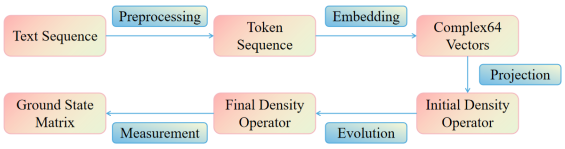

Figure 2: The pipeline of LI-QiLM.

in physical significance. Furthermore, they exhibit the unique capability to incorporate microscopic particle phenomena into the model, which serves to augment their learning capacity. However, the quantum evolution process is rarely considered in recent quantum-inspired models and projection measurements require more actions to ascertain the probability distribution. Therefore, in this work, we focus on that and propose to use LME in the learning process and interferometry in the measurement process.

4 Method

Fig. 2 shows the proposed framework, which consists of the alterations and operations performed on the data at every stage of the process.

4.1 Preprocessing

The data preprocessing initializes with a series of tasks such as case normalization, removal of invalid tokens, and truncation. Following this, the data is randomly divided into training and validation subsets. Subsequently, we employ a pre-trained transformer tokenizer to the sentences, thereby generating distributed representations for each word. These representations are then utilized as input for our model.

4.2 Embedding

In the embedding module, we process the distributed representation of each word by feeding it into both a real embedding module and an imaginary embedding module. These embedding modules collaborate to map vectors from the real number domain into a complex-valued Hilbert space, enabling a more intricate and nuanced representation. This transformation makes the text information more compatible with quantum computation. In practice, the real and imaginary parts can encode diverse information, like the amplitude and phase of qubits, or the directionality of text sequences.

4.3 Projection

After embedding, each word is represented by a pure state as Eq. 1 with multiple bases. During projection, we employ a trainable complex-valued vector to ascertain the weight parameters p_i in Eq. 2 to obtain the density matrix that corresponds to the input sentence. The initial setting assigns an equal value to all words, signifying the weight assigned to each word within the sentence. As training progresses, words that are more representative for sentiment analysis will receive higher weights, reflecting their increased importance in the task.

4.4 Evolution

This step is crucial for simulating the quantum system's evolution and forms the backbone of LI-QiLM. Fig. 3 shows the evolution process. We follow a multi-layered strategy to process the density matrix. The dimensions of the input and output of the overall evolution module will not be changed after learning the features inside the sentence. At the outset, we employ a two-layered GRU architecture to handle the density matrix. This enables the capturing of long-range dependencies within the text and helps to alleviate issues such as vanishing or exploding gradients that can occur during backpropagation. In the subsequent round of processing, we apply the LME as Eq. 3, which encapsulates the entanglement relationships between contextual words and signifies the interference effects of other words in the vocabulary on the current text, demonstrating the potential to help solve natural language processing tasks. Additionally, we construct a complex-valued inception module (Szegedy et al., 2017) that encompasses convolution kernels of varying sizes and pooling operations. This design aims to obtain a richer representation during the evolution.

4.5 Measurement

Leveraging the wave-particle duality principle inherent in quantum mechanics, we utilize the interference phenomenon among quantum states to facilitate the collapse of the superposition state toward the ground state. To observe the interplay between quantum states traversing distinct paths or branches, we employ the Hadamard Gate, as delineated in Eq. 5, as the interferometric operator in Eq. 4. This enables us to conduct interferometric measurements on a designated portion of the density matrix, which subsequently undergoes splin-

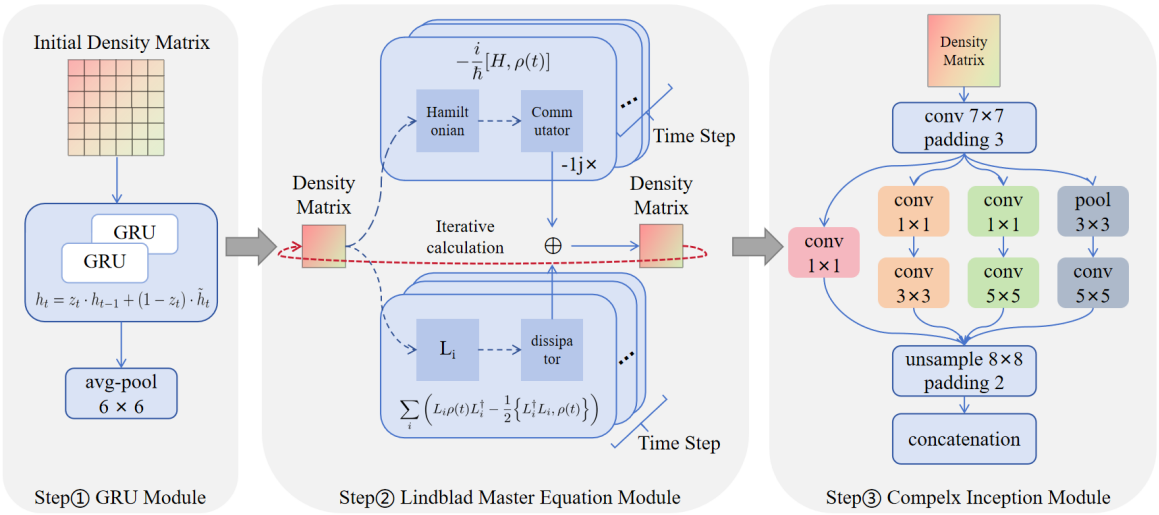

Figure 3: The evolution process.

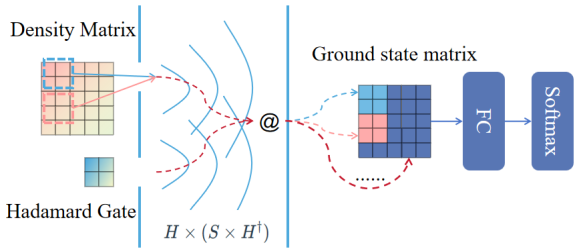

Figure 4: The interference measurement process. The dashed blue curve shows that the relative position of the submatrix is constant before and after calculation. The ground state matrix is finally composed of many ground state submatrices, derived by the product of the different parts of the density matrix and the Hadamard gate.

tering to yield the desired measurement outcome. Fig. 4 shows the simplified interference measurement process. Firstly, submatrices of the same size as the interferometric operator are selected from the density matrix. Subsequently, each of these submatrices and the interferometric operator are calculated using Eq. 4. The resulting calculations are then placed back into their corresponding positions. Once all submatrices have undergone this process, the ground state matrix after the interferometry is obtained. Subsequently, we compute the output by passing the data through a fully connected layer followed by the application of the softmax function. In text classification tasks, the output layer of the fully connected layer is designed to have nodes corresponding to the number of distinct category labels.

$$\text{Hadamard} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (5)$$

5 Experiments

5.1 Datasets

We have conducted a series of experiments across five binary classification datasets, including Customer Review (CR) (Hu and Liu, 2004), Opinion polarity dataset (MPQA) (Wiebe et al., 2005), Movie Review (MR) (Nivre et al., 2016), Sentence Subjectivity (SUBJ) (Nivre et al., 2016), Stanford Sentiment Treebank-2 (SST-2) (Socher et al., 2013), and a multi-classification dataset named Stanford Sentiment Treebank-5 (SST-5) (Socher et al., 2013) for sentiment analysis. The dataset statistic is shown in Table 1.

Table 1: Dataset statistics.

Dataset	Train	Test	Total	Labels	Max Len
CR	2641	1134	3775	pos/neg	105
MPQA	7423	3183	10606	pos/neg	43
MR	7462	3200	10662	pos/neg	61
SUBJ	7000	3000	10000	subj/obj	121
SST-2	67349	1821	69170	pos/neg	56
SST-5	8544	2210	10754	five labels	56

5.2 Settings

5.2.1 Comparison Methods

We employ seven representative models, namely TextCNN (Kim, 2014), Gate Recurrent Unit (GRU) (Chung et al., 2014), Embeddings from Language Models (ELMo) (Peters et al., 2018), Bidirectional Encoder Representations from Transformers (BERT) (Kenton and Toutanova, 2019), Robustly Optimized BERT Pretraining Approach

(RoBERTa) (Liu et al., 2019), CICWE-QNN (Shi et al., 2023) and ComplexQNN (Lai et al., 2023) as the baselines in our experiments.

TextCNN is a convolutional model for text classification. In our experiments, We employ the most effective pattern of the TextCNN, named CNN-non-static, whose word vector matrix is initialized with a pre-trained word vector file. GRU is a refined variant of the Recurrent Neural Network (RNN) architecture, characterized by its streamlined structure, reduced parameter count, and accelerated training efficiency, offering notable advantages over other recurrent architectures.

ELMo is an innovative algorithm for generating dynamic word vectors, which leverages a bidirectional LSTM to encapsulate the nuanced and complex characteristics of words within a language. BERT uses a bidirectional Transformer encoder to capture rich contextual information. RoBERTa, a variation of BERT, fine-tunes key hyperparameters to learn even more extensive language information, resulting in superior performance compared to BERT.

CICWE-QNN is a quantum-inspired model that combines GRU and convolution layer, which fully considers the text characteristics of the projection matrix. ComplexQNN is a variation of ComplexCNN (Trabelsi et al., 2018), which additionally carries out word embedding and measurement operations in the complex number domain.

5.2.2 Implementation Details

We conducted our experiments using the PyTorch platform (Paszke et al., 2019) on an NVIDIA TITAN Xp GPU. During training, we use the cross-entropy loss function and employ the Adam optimizer. In the complex-valued inception module, we employed different kernel sizes of 1, 3 and 5. The specific parameters used in different models and the relevant computational budget are displayed in Table 2. We directly choose the experimental results of CE-Mix shown in Ref. (Shi et al., 2023).

For text encoding, we used a common method RoBERTa, which consistently demonstrated superior performance in NLP tasks. A grid search is conducted using batch size in $\{8, 16, 32, 64\}$, learning rate in $\{10^{-4}, 10^{-5}, 10^{-6}\}$, dropout rate in $\{0.1, 0.15, 0.2, 0.25\}$. Regarding the indispensable parameters H and L_i of Eq. 3, we adopt the method of random initialization. During model verification, we employed accuracy and F1-score as evaluation metrics to gauge the performance of these models.

5.3 Results

We conducted a thorough comparison between LI-QiLM and various models, including classical models, pre-trained models, and other quantum-inspired models. The results from Table 3 for accuracy and Table 4 for F1-score are the best of 10 epochs, which demonstrate that our model consistently outperforms all datasets used.

The performance improvements when contrasting LI-QiLM with CNN-non-static under the accuracy metric are significant: CR (+15.18%), MPQA (+15.33%), MR (+13.72%), SUBJ (+6.41%), SST-2 (+12.31%), SST-5 (+19.14%), average improvement (+13.68%). Comparisons of LI-QiLM and CNN-non-static for the F1-score metric are: CR (+16.48%), MPQA (+22.69%), MR (+13.75%), SUBJ (+6.4%), SST-2 (+12.52%), SST-5 (+19.71%), with an average improvement of (+15.26%). These notable advancements primarily derive from the underlying principles of quantum theory, which embrace the intricate complex-valued word embedding, quantum evolution, and interferometry.

Despite the benchmark set by pre-trained models in the NLP domain, LI-QiLM consistently achieves superior results across a range of tasks. Take RoBERTa for example, LI-QiLM exhibits an advantage under the accuracy metric in CR (+0.88%), MPQA (+0.79%), MR (+0.66%), SUBJ (+0.31%), SST-2 (+0.27%), SST-5 (+0.34%), with an average improvement of (+0.54%). Notably, LI-QiLM holds a comprehensive advantage on all sentiment analysis datasets compared to CICWE-QNN and ComplexQNN.

5.4 Ablation Test

We conduct ablation analysis experiments to investigate the impact of the core modules, including the complex-valued inception module, LME, interference measurement module and GRU. The results are shown in Table 5.

The removal of the complex inception module with diverse convolution kernels diminishes the capacity to capture global information, which can be inferred from the role of the inception module according to Szegedy et al. (2017). and Hebbian learning rule (Szegedy et al., 2017; Hebb, 2005). Dropped results in the second row of Table 5 demonstrate the crucial role of this module in enhancing sentiment classification ability.

Table 2: Model parameter settings.

	LI-QiLM	ComplexQNN	RoBERTa	BERT	ELMo	GRU	CNN-non-static
Max epochs	10	10	10	10	10	10	10
Batch size	8	32	32	32	32	32	32
Embedding dim	768	128	768	768	1024	128	64
Output dim	128	128	128	128	128	64	64
Learning rate	8e-6	1e-3	1e-5	1e-5	1e-3	1e-3	1e-3
Dropout	0.2	0.2	0.2	0.2	0.2	0.2	0.2
Random seed	42	42	42	42	42	42	42
Params	8.31M	8.28M	0.59M	0.59M	0.59M	1.78M	0.05M
FLOPs	77.54G	76.42G	18.87M	18.87M	246.22M	342.29M	7.21M

Table 3: Experimental results of eight models on six sentiment classification datasets evaluated with accuracy.

	CR	MPQA	MR	SUBJ	SST-2	SST-5
CNN-non-static	78.28	76.34	75.62	90.66	82.11	35.69
GRU	79.96	84.82	74.90	92.16	83.60	37.60
ELMo	83.31	90.16	81.9	93.6	88.41	47.41
BERT	89.58	91.20	86.75	96.60	93.00	49.13
RoBERTa	92.58	90.88	88.68	96.76	94.15	54.49
CICWE-QNN$^{\Delta}$	83.3	87.2	78.3	93.2	85	-
ComplexQNN	91.87	90.63	87.40	96.57	92.54	50.40
LI-QiLM(Ours)	**93.46**	**91.67**	**89.34**	**97.07**	**94.42**	**54.83**

The best results of each dataset are highlighted in bold. The signal $^{\Delta}$ denotes that the data is directly chosen from the model author's experimental results.

Table 4: Experimental results of eight models on six sentiment classification datasets evaluated with F1-score.

	CR	MPQA	MR	SUBJ	SST-2	SST-5
CNN-non-static	76.39	67.53	75.59	90.66	81.89	33.96
GRU	78.15	81.24	74.89	92.16	83.54	35.99
ELMo	82.14	87.98	81.76	93.59	88.40	44.41
BERT	88.95	89.69	86.74	96.59	92.99	48.48
RoBERTa	91.95	89.35	88.68	96.76	94.14	52.58
CICWE-QNN$^{\Delta}$	86.7	78.2	76.1	92	83.6	-
ComplexQNN	91.14	88.86	87.39	96.57	92.53	49.51
LI-QiLM(Ours)	**92.87**	**90.22**	**89.34**	**97.07**	**94.42**	**53.67**

The best results of each dataset are highlighted in bold. The signal $^{\Delta}$ denotes that the data is directly chosen from the model author's experimental results.

Removing the LME module may lead to a decline in its ability to understand polysemous groups. It can be seen from Fig. 5 that for the same sentence, LI-QiLM and other models give completely opposite emotional judgments on the meaning of the sentence, which shows that other models do not fully understand the important phrases in the sentence that affect emotions.

When excluding the GRU module, the model ignores taking the positional relationships of words into account, potentially leading to inaccuracies in word order judgments. Besides, the removal of IM module also results in a slight decline in performance. Further, the performance drops dramatically when we remove all four modules. These suggest that each of these four design elements plays a crucial role in shaping the performance.

Although we chose the optimal value of 10 epochs LI-QiLM runs on the MR, the performance of the IM module remained less than satisfactory. This might be due to certain perturbations that were unavoidably introduced throughout the evolution of the density matrix, subsequently impacting the precision of measurements in an adverse manner.

5.5 Case Study

We provide a qualitative analysis of LI-QiLM on the SST-5. Fig. 5 shows the sentiment polarities and probabilities of the ground-truth label predicted by different models. We can see that LI-QiLM gives the correct polarities while both RoBERTa and GRU give wrong predictions in some cases.

Take the first sentence as an example, it is difficult to infer the sentiment polarities for its misleading words "entertained" and the multiple-meaning phrase "put off", considering these two pieces of information, LI-QiLM can still give the correct label "negative", while both RoBERTa and GRU even give the opposite answer. Additionally, the GRU exhibits almost 100% confidence in the answers it identifies, so the probability that the label should be correct is very low. As the sentence example 1 in Fig. 5 shows, it has a value of only 9.2×10^{-6}. This justifies the effectiveness of the proposed module's ability to capture the context in sentences and the interaction with the surrounding environment.

6 Conclusions

In this study, we propose LI-QiLM, a quantum-inspired neural network, built on quantum evolution and interferometry. We have provided detailed principles, highlighting key implementation modules, including the complex-valued inception module, LME module and interferometry module. Fur-

Table 5: Ablation results evaluated with accuracy and f1-socre.

ACCURACY						
Model / Dataset	CR	MPQA	MR	SUBJ	SST-2	SST-5
LI-QiLM	**93.46**	**91.67**	89.34	**97.07**	**94.42**	**54.83**
w/o Complex Inception	92.49(-0.97%)	90.73(-0.94%)	88.46(-0.88%)	96.4(-0.67%)	94.15(-0.27%)	51.31(-3.52%)
w/o Linblad Master Equation	92.67(-0.79%)	90.7(-0.97%)	89.03(-0.31%)	96.73(-0.34%)	93.92(-0.50%)	53.04(-1.79%)
w/o Interference Measurement	93.38(-0.08%)	91.45(-0.22%)	**89.96(+0.62%)**	96.90(-0.17%)	94.38(-0.04%)	54.04(-0.79%)
w/o GRU	92.93(-0.53%)	91.2(-0.47%)	89.06(-0.28%)	96.13(-0.94%)	94.15(-0.27%)	52.50(-2.33%)
w/o All	91.34(-2.12%)	90.32(-1.35%)	86.96(-2.38%)	95.33(-1.74%)	91.74(-2.68%)	44.95(-9.88%)
F1-SCORE						
Model / Dataset	CR	MPQA	MR	SUBJ	SST-2	SST-5
LI-QiLM	**92.87**	**90.22**	89.34	**97.06**	**94.41**	**53.67**
w/o Complex Inception	91.89(-0.98)	89.09(-1.13)	88.44(-0.9)	96.39(-0.67)	94.14(-0.27)	50.05(-3.62)
w/o Linblad Master Equation	92.12(-0.75)	89.23(-0.99)	89.02(-0.32)	96.73(-0.33)	93.91(-0.50)	49.44(-4.23)
w/o Interference Measurement	92.80(-0.07)	89.85(-0.37)	**89.96(+0.62)**	96.89(-0.17)	94.37(-0.04)	51.04(-2.63)
w/o GRU	92.32(-0.55)	89.78(-0.44)	89.06(-0.28)	96.13(-0.93)	94.14(-0.27)	48.15(-5.53)
w/o All	90.93(-1.94)	88.39(-1.83)	86.96(-2.38)	95.34(-1.72)	91.73(-2.68)	43.18(-10.49)

The best score of each dataset is in bold.

Sentence Example 1			Sentence Example 2		
If you 've ever entertained the notion of doing what the title of this film implies , what sex with strangers actually shows may put you off the idea forever . ***negative***			A very well-made , funny and entertaining picture . ***very positive***		
GRU	RoBERTa	**LI-QiLM (Ours)**	GRU	RoBERTa	**LI-QiLM (Ours)**
(*positive*) ×	(*very positive*) ×	(*negative*) √	(*neutral*) ×	(*negative*) ×	(*very positive*) √
P(*negative*) = 9.2e-6	P(*negative*) = 0.32	P(*negative*) = 0.60	P(*very positive*) = 4.7e-6	P(*very positive*) = 0.23	P(*very positive*) = 0.62

Figure 5: Examples from the SST-5 with their polarities predicted by different models(i.e., LI-QiLM, RoBERTa and GRU). The green background denotes that the result is correct while the red background means an error.

thermore, we conducted evaluations of LI-QiLM on six varied datasets and the proposed model consistently outperformed the existing popular methods. Finally, ablation tests were conducted to appraise the effectiveness of the key modules.

In the future, we will aim to augment the capacity of the model to extract profound linguistic insights by integrating dynamic word embedding. Furthermore, few of the current QiLMs have quantitatively evaluated the effectiveness of the quantum concepts employed within them. Therefore, greater efforts must be directed towards assessing the efficacy of various components within QiLMs.

Acknowledgments

This research was funded by the Natural Science Foundation of Fujian Province of China (Grant Nos. 2018J01799, 2022J01120)

7 Limitations

Compared to the traditional neural networks, transformer-based pre-trained models, and other quantum-inspired models, LI-QiLM boasts superior performance and uses formulas and methods that are more consistent with quantum theory to deal with sentiment analysis tasks with quantum-like behavior, which makes "black box" neural network models explainable by quantum theory, thus providing a new possible way to enhance the interpretability that aims at revealing the actual working principles. However, we are encountering some obstacles that need to be addressed. Firstly, LI-QiLM utilizes traditional neural networks to simulate the process of quantum computing, which necessitates much more parameters. Meanwhile, high dimensional matrix calculation caused by density matrix leads to higher floating point operations (FLOPs). The specific data are shown in the last two lines of Table 2. Secondly, utilizing quantum computers is an expensive proposition, which leads to the uncertainty surrounding the feasibility of running quantum-inspired models on quantum computers.

References

Peter D Bruza, Zheng Wang, and Jerome R Busemeyer. 2015. Quantum cognition: a new theoretical approach to psychology. *Trends in cognitive sciences*, 19(7):383–393.

Jerome R Busemeyer and Zheng Wang. 2015. What is quantum cognition, and how is it applied to psychology? *Current Directions in Psychological Science*, 24(3):163–169.

Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS 2014 Workshop on Deep Learning, December 2014*.

Richard P. Feynman. 1982. Simulating physics with computers. *International Journal of Theoretical Physics*, 21.

Donald Olding Hebb. 2005. *The organization of behavior: A neuropsychological theory*. Psychology press.

Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Knowledge Discovery and Data Mining*.

Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.

Wei Lai, Jinjing Shi, and Yan Chang. 2023. Quantum-inspired fully complex-valued neutral network for sentiment analysis. *Axioms*, 12(3):308.

Qiuchi Li, Dimitris Gkoumas, Christina Lioma, and Massimo Melucci. 2021. Quantum-inspired multimodal fusion for video sentiment analysis. *Information Fusion*, 65:58–71.

Qiuchi Li, Sagar Uprety, Benyou Wang, and Dawei Song. 2018. Quantum-inspired complex word embedding. In *Proceedings of the Third Workshop on Representation Learning for NLP*, pages 50–57.

Qiuchi Li, Benyou Wang, and Massimo Melucci. 2019. Cnm: An interpretable complex-valued network for matching. In *Proceedings of NAACL-HLT*, pages 4139–4148.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Daniel Manzano. 2020. A short introduction to the lindblad master equation. *Aip Advances*, 10(2).

Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. volume 1, pages 2227 – 2237.

Thomas J Santner, Brian J Williams, William I Notz, and Brain J Williams. 2003. *The design and analysis of computer experiments*, volume 1. Springer.

J. Shi, Z. Li, W. Lai, F. Li, R. Shi, Y. Feng, and S. Zhang. 2023. Two end-to-end quantum-inspired deep neural networks for text classification. *IEEE Transactions on Knowledge amp; Data Engineering*, 35(04):4335–4345.

R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, and C. Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Empirical Methods in Natural Language Processing*.

Alessandro Sordoni, Jian Yun Nie, and Yoshua Bengio. 2013. Modeling term dependencies with quantum language models for ir. In *International Acm Sigir Conference on Research Development in Information Retrieval*, page 653.

Karen Spärck Jones. 2004. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 60(5):493–502.

Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.

Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, Joao Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. 2018. Deep complex networks. In *International Conference on Learning Representations*.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *Transactions on Machine Learning Research*.

Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. In *Language Resources and Evaluation*.

Peng Zhang, Wenjie Hui, Benyou Wang, Donghao Zhao, Dawei Song, Christina Lioma, and Jakob Grue Simonsen. 2022. Complex-valued neural network-based quantum language models. *ACM Transactions on Information Systems (TOIS)*, 40(4):1–31.

Peng Zhang, Jiabin Niu, Zhan Su, Benyou Wang, Liqun Ma, and Dawei Song. 2018a. End-to-end quantum-like language models with application to question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Peng Zhang, M. A. Xindian, and Dawei Song. 2018b. A survey of quantum language models. *Scientia Sinica(Informationis)*, page 1486.

Yazhou Zhang, Dawei Song, Peng Zhang, Panpan Wang, Jingfei Li, Xiang Li, and Benyou Wang. 2018c. A quantum-inspired multimodal sentiment analysis framework. *Theoretical Computer Science*, 752:21–40.