

Constructing Word-Context-Coupled Space Aligned with Associative Knowledge Relations for Interpretable Language Modeling

Fanyu Wang and Zhenping Xie*

School of Artificial Intelligence and Computer Science, Jiangnan University, China

fanyu_wang@stu.jiangnan.edu.cn

xiezp@jiangnan.edu.cn

Abstract

As the foundation of current natural language processing methods, pre-trained language model has achieved excellent performance. However, the black-box structure of the deep neural network in pre-trained language models seriously limits the interpretability of the language modeling process. After revisiting the coupled requirement of deep neural representation and semantics logic of language modeling, a Word-Context-Coupled Space (W2CSpace) is proposed by introducing the alignment processing between uninterpretable neural representation and interpretable statistical logic. Moreover, a clustering process is also designed to connect the word- and context-level semantics. Specifically, an associative knowledge network (AKN), considered interpretable statistical logic, is introduced in the alignment process for word-level semantics. Furthermore, the context-relative distance is employed as the semantic feature for the downstream classifier, which is greatly different from the current uninterpretable semantic representations of pre-trained models. Our experiments for performance evaluation and interpretable analysis are executed on several types of datasets, including SIGHAN, Weibo, and ChnSenti. Wherein a novel evaluation strategy for the interpretability of machine learning models is first proposed. According to the experimental results, our language model can achieve better performance and highly credible interpretable ability compared to related state-of-the-art methods.¹

1 Introduction

Machine learning has recently been democratized in various domains, such as search engines, conversational systems, and autonomous driving (Gao et al., 2018; Grigorescu et al., 2020; Wang et al., 2022). However, despite AI technologies signifi-

cantly facilitating industrial processes and improving work experiences, the uninterpretable logic of machines leads to distrust, which hinders further development of AI. Explainable Artificial intelligence (XAI), proposed to bridge the block between humans and machines, has been increasingly attracting attention recently, where "explanation" is described as abductive inference and transferring knowledge (Josephson and Josephson, 1996; Miller, 2019). Calling for explaining and understanding the machine learning process, researchers aim to interpret the methods for system verification, compliance with legislation, and technology improvement.

Computational linguistics, which serves as the theoretical foundation for NLP, aims to promote communication between humans and machines (Khan et al., 2016). However, during the recent development, the uninterpretable NLP methods have raised concerns. The decreased transparency and increased parameter complexity adversely affect the model explainability and controllability, even if the performance of the language models has significantly improved, such as BERTs (Devlin et al., 2019; Liu et al., 2019a; Clark et al., 2020), GPTs (Radford et al., 2018, 2019; Brown et al., 2020) and so on. The performance advantage of the black-box methods is attractive while the researchers investigate interpretable algorithms. Therefore, the existing works mainly focus on (1) explaining the black-box methods and (2) using interpretable models (Ribeiro et al., 2016; Li et al., 2022a).

In order to explain over-parameterized language models, the model-agnostic analysis is investigated on recent deep neural methods. Without modifying the black-box models, the researchers analyze the immediate feature of the neural layers, attention distribution, and so on (Clark et al., 2019; Vig, 2019; Rogers et al., 2020). Quantitative experiments and visual analysis is able to partially reveal the behaviors of the key components and

*Corresponding author

¹<https://github.com/ColeGroup/W2CSpace>

the overall response of the methods to certain patterns (Hewitt and Manning, 2019; Kovaleva et al., 2019). However, without any interpretable optimization of the model, the analysis is unable to provide enough detail for understanding (Rudin, 2019), which indicates *a completely faithful explanation from black-box components or deep neural methods is impossible*.

Different from model-agnostic analysis, methods that integrate interpretable algorithms enable more comprehensive interpretability. Two different types of structures are adopted in these approaches, including (1) a black-box backbone with an interpretable bypass for implicit informing and (2) a transparent backbone with an interpretable algorithm for direct interpreting (Beckh et al., 2021). For implicitly informed methods, the introduced interpretable knowledge regulates the immediate feature or embedding (Liu et al., 2019b; Rybakov et al., 2020). Within the interpretable bypass, the performance of the backbone is maximally preserved, which is the main reason these structures are often opted for over-parameterized models (Jang et al., 2021; Chen et al., 2020). However, *the integrated knowledge is unable to decisively change the structure of the backbones for a transparent decision process*, which adversely affects the generalization ability of the approaches and limits the application of the methods to specific tasks. In contrast, approaches with interpretable backbones exhibit a more integrated relationship between the components, enabling better explanations than the above two types of approaches. The interpretable algorithms serve as word embedding, immediate feature, or the classifier to realize transparency decision process (Onoe and Durrett, 2020; Lee et al., 2022; Kaneko et al., 2022). But *the performance of existing interpretable models remains incomparable to the most advanced language models*.

In this work, we address the aforementioned obstacles by developing a novel interpretable language modeling method by constructing a **Word-Context-Coupled Space** (W2CSpace) aligned with statistical knowledge², which enables (1) effective interpretation of BERT representations (Devlin et al., 2019) by introducing the interpretable statistical logic, (2) reasonable context abstraction

²Chinese characters and English words are typically considered at the same processing level in pre-trained language models, so we refer to Chinese characters as "Chinese word" in this paper to avoid ambiguity.

with the coupled word-level semantics, and (3) interpretable modeling for the given text with the context-relative distance. W2CSpace serves as the key component in the backbone of our language modeling method, which realizes a decisive transparency increasing compared with the model-agnostic and implicitly informed methods. The structure of our method is illustrated in Figure 1. Specifically, our main contributions can be summarized as follows:

- Word-level semantics in W2CSpace is originated from BERT immediate representation with the help of a mapping network, preserving the language modeling ability of the deep neural methods (Section 2.1.1).
- An associative matrix sampled from associative knowledge network (AKN, Li et al., 2022c) is introduced for aligning with the semantic distances (Section 2.1.2 and 2.1.3).
- Based on the linguistic concept, the contexts are abstracted using k -means clustering on neighboring word elements. (Section 2.2.1).
- The context-relative distance, computed between the input text and the context clusters in W2CSpace, serves as the semantic feature to describe the text semantics (Section 2.3.1).
- The experiments on different NLP tasks demonstrate the effectiveness of W2CSpace. Additionally, an interpretable analysis is designed to verify the interpretability of the our method (Section 3).

2 Methodology

2.1 Initialization of W2CSpace

Since current researchers opt for high-dimensional representations in their methods, it is widely believed that over-parameterization advances language modeling performance. With respect to the standard language modeling process, the words in the given text are modeled based on their meaning under different contexts. Regardless of the attributes of the words themselves, greater representation dimensions enable better performance in distinguishing the words with similar semantics. However, different from the deep neural methods, the linguistic attributes of the text, such as co-occurrence rules, word shape, and so on, serve

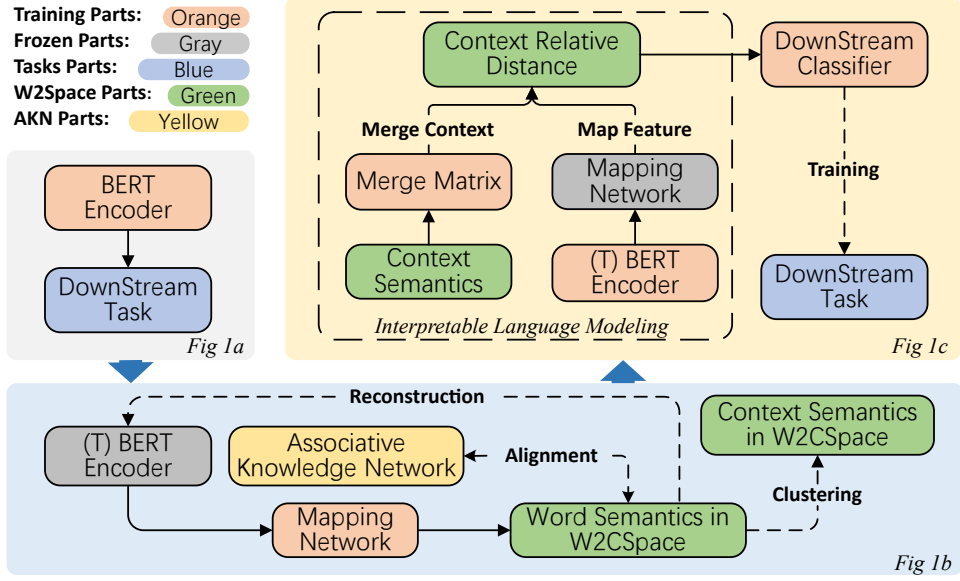


Figure 1: Overview of the method’s architecture. The subfigs exhibit the training process for our interpretable language method. (1a) The BERT encoder is firstly fine-tuned under the downstream task. We map the feature F_B from fine-tuned BERT encoder to the word elements C in W2CSpace ("T" refer to "trained"); (1b) After the training of the mapping network with a reconstruction and alignment task, the word semantics are clustered to high-dimensional context semantics X ; (1c) During interpretable language modeling process, the context semantics in W2CSpace is gradually optimized by multiplying the merge matrix M_M . We compute the context-relative distance of the mapped feature C within the optimized context \bar{X} for the downstream task.

as the basis for NLP tasks, which fit the understanding and deduction processes of humans.

In order to unify the high-dimensional representations and interpretable statistical knowledge, we design a mapping network to transfer the semantic representation from BERT encoder to low-dimensional elements in W2CSpace and introduce a statistical alignment with AKN during the training of the mapping network. Within the above processing, the mapped elements are distributed in W2CSpace according to their corresponding word-level semantics.

2.1.1 Representation Mapping from BERT

The mapping network is a neural network with a backbone of a convolution network, which enables dimension reduction process. The mapped elements in smaller dimensions are regarded as the coordinates of the word-level semantics in W2CSpace. Besides, the introduction of the convolution network is able preserve semantic information from BERT for maintain the performance advantages of pre-trained models. For BERT representations F_B of given sentence $S = \{x_1, x_2, \dots, x_d\}$, the corresponding elements $C = \{c_1, c_2, \dots, c_d\}$ in W2CSpace are obtained accord-

ing to:

$$C = \text{Tanh}\{\text{LN}[\text{Convs}(F_B) + \text{Res}(F_B)]\} \quad (1)$$

where $F_B \in \mathbb{R}^{n \times h}$ and $C \in \mathbb{R}^{n \times k}$, h is the hidden sizes of BERT encoder and k is the coordinate size of W2CSpace. $\text{Tanh}(\cdot)$ and $\text{LN}(\cdot)$ are Tanh and layer normalization operations. A convolution network $\text{Convs}(\cdot)$ with different filter sizes added a residual connection $\text{Res}(\cdot)$ is used in the mapping network.

2.1.2 Statistical Alignment with AKN

AKN, a statistical network based on phrase co-occurrence, is introduced to sample to an associative matrix reflected the associative relations among the given sentence, which is also adopted in previous work (AxBERT, Wang et al., 2023). Within the original AKN is conducted on phrase-level, we modify AKN to word-level to fit the processing of BERT and opt for the construction and sampling methods of the AKN (A) and associative matrix (M_S) similar to AxBERT as:

$$A_{i,j} = \prod_{\text{sent}} \text{SR} \sum_{\text{sent}} \frac{1}{\text{distance}_{\langle i,j \rangle}} \quad (2)$$

$$M_{S_{i,j}} = \sigma \frac{A_{i,j}}{\text{Avg}(A_i)} - 0.5 \quad (3)$$

where $\mathbf{A} \in \mathbb{R}^{v \times v}$, $\mathbf{M}_S \in \mathbb{R}^{d \times d}$, v is the length of the word list, d is the length of the given sentence, and $\sigma(\cdot)$ and $\text{Avg}(\cdot)$ are functions of sigmoid and average. For the word pair $\langle i, j \rangle$, $\text{distance}_{\langle i, j \rangle} = |i - j|$ is the word-distance between i -th and j -th word in sentence. $\dot{A}_{i \cdot}$ is the association score of i -th word under current sentence, and $\mathbf{M}_{S_{i \cdot}}$ is the i -th row of \mathbf{M}_S .³

Since we compute the cosine distances matrix in the given sentence, the associative matrix is aligned with the word-level distance matrix to integrate the statistical logic into W2CSpace. We introduce a mean square indicator I_{MS} to indicate the alignment result. Specifically, for the word pair $\langle i, j \rangle$, the indicator $I_{MS_{i,j}}$ is defined as:

$$I_{MS_{i,j}} = \text{MnSqr}[\text{CosDis}(c_i, c_j), \mathbf{M}_{S_{i,j}}] \quad (4)$$

where $I_{MS} \in \mathbb{R}^{n \times n}$ and n is the length of the given sentence. $\text{MnSqr}(\cdot)$ and $\text{CosDis}(\cdot, \cdot)$ are the mean square and cosine distance functions.

2.1.3 Training of Mapping Network

The objective L_M of the mapping network is composed of L_{MS} and L_{Rec} , which correspond to the mean square error loss and the reconstruction loss. With respect to the statistical alignment process, the mapping network is trained under the alignment objective L_M . Besides, we introduce a reconstruction loss by reversing the structure of mapping network to reconstruct the representation of BERT immediate feature \mathbf{F}_B . The objective of the mapping network is calculated according to:

$$L_M = L_{MS} + L_{Rec} \quad (5)$$

$$L_{MS} = \text{Mean}(I_{MS}) \quad (6)$$

$$L_{Rec} = \text{MAE}(\mathbf{C} - \mathbf{F}_B) \quad (7)$$

where $\text{Mean}(\cdot)$ is the average function, $\text{MAE}(\cdot)$ is the mean absolute error operation (Choi et al., 2018). The introduction of reconstruction loss aims to guarantee that the mapped word elements preserve the semantics of BERT representations.

2.2 Abstraction of Context-level semantics

Humans are able to recognize emotion from language, action, and so on (Barrett et al., 2007). Specifically, in linguistics, humans recognize emotion with the context in the given sentences. However, humans are able to feel the emotion rather than explicitly describe it, because that context is

an abstract concept in linguistics and is hard to simply quantify.

While the sentence is composed of words, the corresponding context is established from the word semantics, which can be realized in W2CSpace. Therefore, we employ the k -means clustering on the word-level semantics to abstract the context semantics. The context is able to be extracted based on the common semantics among the words located adjacently in W2CSpace.

2.2.1 Context Clustering Based on Word Semantics

With the help of the mapping network and the statistical alignment, the word elements are reasonably distributed in W2CSpace according to their semantics, where the neighbors in W2CSpace refer to similar semantics. k -means clustering is an algorithm based on the distances (Hartigan and Wong, 1979), which is introduced to abstract the word semantics to k classes according to their semantics (distance). The clustering process is defined as:

$$\mathbf{X}\{x_i | \text{Context}(x_i)\} = \text{KM}_{\text{CosDis}}(c_1, c_2, \dots, c_n) \quad (8)$$

where x_i is the i -th classes of context and $\text{KM}_{\text{CosDis}}(\cdot)$ is the k -means clustering algorithm based on cosine distance. The word semantics in W2CSpace is clustered into k classes, which represents the process of abstraction of context.

2.2.2 Reasonable Cluster Merging for Context Clustering

When the clustering is executed, especially in k -means algorithm, the appropriate k number is hard to determine (Hamerly and Elkan, 2003). As for the clustering process for context, the number k additionally represents the types of the contexts in W2CSpace. A small number of k possibly decreases the performance of language modeling with a rough context environment, but the large number is contrary to the logic of humans as humans cannot precisely distinguish the detailed emotion behind the text as a machine does. Besides, the large number of k increase the time-costing of the language modeling process. Choosing the right number k relate to the reasonability of the context subspace.

Serving as the part of W2CSpace, the context clusters are used for language modeling. Therefore, we introduce a merge matrix to the top of the clustering results, which is optimized under the

³The same shrink rate $\text{SR} = 0.95$ as AxBERT.

downstream task. With the guidance of the downstream tasks, the merge matrix dynamically adjusts context semantics for the inter-communication between different context clusters, which reflects a gradual clustering process and realizes a reasonable context clusters for the downstream tasks. The clustering is able to be balanced with the computation process of the merging is defined as:

$$\bar{X} \{ \bar{x}_i | \text{Context}(\bar{x}_i) \} = M_M \times [x_1, x_2, \dots, x_k] \quad (9)$$

where \bar{x}_i indicates the context semantics after merging, M_{Merge} is the merge matrix and $M_{\text{Merge}} \in \mathbb{R}^{n \times k}$. n is the size of coordinates in W2CSpace and k is the presetting number of the clustering.⁴

2.3 Interpretable Language Modeling via W2CSpace

As the standard methodology for language modeling in deep neural methods, the semantic representation is gradually modeled through containing neural networks, which is a simulation of the neural processing of the brain. However, the structural simulation is unable to realize the interpretable on logical level. The decision process through neural networks still remains in the black box.

From the perspective of humans, emotion recognition is significant in daily (Barrett et al., 2011), which is also an important ability for machines to interact with humans (Kosti et al., 2017). By simulating the recognition process of humans, we introduce a context-relative distance computed between the given text and the contexts in W2CSpace, which enables the interpretable language modeling process with the cooperation of the semantics on word- and context-level.

2.3.1 Computation of Context-Relative Distance

The context-relative distance based on the cosine distance is also adopted in Formula 4. Compared with the euclidean distance, the equation of the cosine distance is more efficient in time-costing and storage. The context-relative distance D is computed according to:

$$D = \text{CosDis}(\bar{X}, C) \quad (10)$$

where \bar{X} is the context clusters, C is the mapped word elements from BERT encoder, D is the context-relative distance and $D \in \mathbb{R}^{d \times k}$. n is the

⁴However, the number k will influence the performance of language modeling, which is discussed in Section 3.4.

word length of the given text and k is the number of the context clusters.

2.3.2 Training of Interpretable Language Modeling Method

The context-relative distance is able to directly connect with the downstream classifier, which is similar to the traditional encoding structure. The interpretable language modeling component is regarded as the standard BERT-based encoder for downstream tasks, where the standard objectives in transformer package⁵ are employed.

3 Experiments

3.1 Experimental Settings

We conduct our work on NVIDIA Tesla A100 and AMD EPYC 7742 64-Core CPU. During the interpretable language modeling process, BERT-base-Chinese pre-trained model is used, and the original parameters are opted⁶. Additionally, we opt for the rate of 0.3 for all the dropout layers, a learning rate of 2e-5 for 10-epoch-training of BERT encoder in Fig. 1a, a learning rate of 1e-5 for 3-epoch-training of the mapping network in Fig. 1b.

3.2 Datasets

The detailed information of the datasets is exhibited in Table 1. **CLUE**⁷, an open-ended, community-driven project, is the most authoritative Chinese natural language understanding benchmark (Xu et al., 2020), where the news dataset is used to initialize associative knowledge network; **SIGHAN15** is benchmark for traditional Chinese spelling check evaluation (Tseng et al., 2015), which is widely adopted in simplified Chinese spelling check evaluation by converting to simplified Chinese (Cheng et al., 2020; Liu et al., 2021); **Hybird** is a massive dataset for Chinese spelling correction (Wang et al., 2018), which is used the training of the correction methods (Wang et al., 2019; Cheng et al., 2020); **Weibo**⁸ and **ChnSenti**⁹ sentiment dataset is constructed with the comments from the largest Chinese social community (Sina Weibo) and Chinese hotel reservation websites, which are adopted in the previous work for sentiment classification evaluation (Li et al., 2020, 2022b).

⁵https://pytorch.org/hub/huggingface_pytorch-transformers/

⁶<https://huggingface.co/bert-base-chinese>

⁷<https://github.com/CLUEbenchmark/CLUE>

⁸https://github.com/pengming617/bert_classification

⁹<https://github.com/SophonPlus/ChineseNlpCorpus/>

Table 1: Statistics information of the used datasets.

Dataset	TrainSet	TestSet	Dataset Type	Usage
CLUE	2,439	-	Article	Initialization of AKN
HybirdSet	274,039	3,162	Sentence	Training of correction task
SIGHAN15	6,526	1,100	Sentence	Evaluation of correction task
ChnSenti	9,600	1,089	Article	Training and Evaluation of sentiment analysis
Weibo100k	100,000	10,000	Article	Training and Evaluation of sentiment analysis

3.3 Comparison Approaches

We fine-tune our method with standard classifiers of BERT Masked LM and sequence classification for spelling correction and sentiment classification with training of 10 epoch, $1e-5$ learning rate and 3 epoch, $1e-5$ learning rate.

SoftMask is a BERT-based spelling correction method with a soft-mask generator, where the soft-masked strategy is similar to the concept of error detection (Zhang et al., 2020).

FASpell conducts the Seq2Seq prediction by incorporating BERT with additional visual and phonology features (Hong et al., 2019).

SpellGCN incorporates BERT and the graph convolutional network initialized with phonological and visual similarity knowledge for Chinese spelling correction (Cheng et al., 2020).

PLOME integrates the phonological and visual similarity knowledge into a pre-trained masked language model with a large pre-train corpus consisted of one million Chinese Wikipedia pages. And it is the SOTA in previous work (Liu et al., 2021).

HeadFilt is an adaptable filter for Chinese spell check, which introduce a hierarchical embedding according to the pronunciation similarity and morphological similarity (Nguyen et al., 2021).

HLG is a Chinese pre-trained model for word representation by aligning the word-level attention with the word-level distribution with a devised pooling mechanism (Li et al., 2020).

MWA introduces a heterogeneous linguistics graph to pre-trained language model. The graph-based structure integrates the linguistics knowledge in the neural network and achieves the SOTA performance in language modeling (Li et al., 2022b).

HLG and **MWA** is employed on various pre-trained language model, such as vanilla BERT (Devlin et al., 2019), BERT-wwm (Cui et al., 2021), and ERNIE (Sun et al., 2019). We use the evaluation results on different pre-trained language models in their original paper.

3.4 Main Experiments

The efficacy of our interpretable language modeling method is evaluated on different tasks, including Chinese spelling correction and sentiment classification. Chinese spelling correction requires the advanced language model for token classification, where every word in the given text is classified into a single class. And the sequence classification is needed in Chinese sentiment classification, where the given text is classified into positive and negative sentiment. The token and sequence classification tasks are able to cover most of the current classification scenario, which enable the efficient demonstration of the language modeling performance of our method.

3.4.1 Results of Chinese Spelling Correction

Similar with the past works (Cheng et al., 2020; Liu et al., 2021), the correction experiment is employed on word- and sentence-level. Within a more comprehensive perspective, the sentence-level evaluation is wider adopted and more convincing, so we use the same evaluation matrix with the past works (Liu et al., 2021; Nguyen et al., 2021).

As shown in Table 2, the evaluation on word- and sentence-level composed of different indexes, including detection precision (DP), correction precision (CP), detection recall (DR), correction recall (CR), detection F1 score (DF1) and correction F1 score (CF1). Besides, we assess the influence of the number choosing of n (the size of the coordinates in W2Cspace) and k (the context number in k -means clustering algorithm).

From the correction results in Table 2, our method outperforms the baselines on both word- and sentence-level. Specifically, at sentence-level, our method respectively advances 0.2 and 0.7 points in DF1 and CF1; at word-level, our method is not able to achieve comparable performance than PLOME (Liu et al., 2021), but advances than the SpellGCN (Cheng et al., 2020) with a 0.5 and 0.4 point improvement, and we think the massive train-

Table 2: Results of Chinese spelling correction.

Method	k	Word Level						Sentence Level					
		DP	DR	DF1	CP	CR	CF1	DP	DR	DF1	CP	CR	CF1
FASpell	-	-	-	-	-	-	-	67.6	60.0	63.5	66.6	59.1	62.6
SoftMask	-	-	-	-	-	-	-	73.7	73.2	73.5	66.7	66.2	66.4
BERT	-	92.7	85.0	88.7	96.2	81.8	88.4	76.5	78.6	77.5	76.0	76.5	76.3
PLOME [†]	-	94.5	87.4	90.8	97.2	84.3	90.3	77.4	81.5	79.4	75.3	79.3	77.2
SpellGCN	-	88.9	87.7	88.3	95.7	83.9	89.4	74.8	80.7	77.7	72.1	77.7	75.9
HeadFilt	-	-	-	-	-	-	-	84.5	71.8	77.6	84.2	70.2	76.5
W2CSpace $n = 50$	500	90.5	86.2	88.3	96.2	82.9	89.0	76.4	77.6	77.0	75.8	75.1	75.5
	1000	90.9	85.8	88.2	96.3	82.6	88.9	78.7	79.6	79.2	78.1	76.3	77.4
	1500	91.1	87.1	89.0	96.4	83.9	89.7	78.0	79.8	78.9	77.4	77.0	77.2
	2000	91.7	86.2	88.9	96.3	83.0	89.2	79.2	79.8	79.5	78.5	76.6	77.5
	3000	91.9	86.5	89.1	96.7	83.6	89.6	78.5	80.1	79.3	78.0	77.7	77.9
W2CSpace $n = 100$	500	90.7	86.0	88.3	96.3	82.8	89.1	76.6	80.5	78.5	75.9	77.3	76.6
	1000	90.6	86.3	88.4	96.2	83.2	89.1	77.2	79.8	78.5	76.4	76.6	76.5
	1500	91.2	86.3	88.7	95.8	82.7	88.8	79.4	79.4	79.4	78.8	76.6	77.7
	2000	91.4	87.3	89.3	95.7	83.6	89.2	78.3	80.9	79.6	77.6	77.9	77.8
	3000	90.9	86.4	88.6	96.8	83.7	89.8	76.9	79.7	78.1	76.3	76.7	76.5

Table 3: Results of sentiment classification.

Method	k	ChnSenti	Weibo100K
BERT	-	94.72	97.31
+MWA	-	95.34	98.14
+HLG	-	95.83	98.17
BERTwvm	-	94.38	97.36
+MWA	-	95.01	98.13
+HLG	-	95.25	98.11
ERNIE	-	95.17	97.30
+MWA	-	95.52	98.18
+HLG	-	95.83	98.22
W2CSpace $n = 50$	50	95.70	98.22
	100	95.70	98.24
	200	95.20	98.27
	500	95.45	98.30
	800	95.45	98.31
1000	95.03	98.23	
W2CSpace $n = 100$	50	95.53	98.29
	100	94.94	98.25
	200	95.62	98.27
	500	95.11	98.31
	800	95.87	98.31
1000	95.37	98.28	

ing dataset of PLOME significantly enhance the correction performance, where the size of dataset composed of per-train and fine-tune dataset is 600 times larger than ours.

With respect to the changes of the parameters of W2CSpace, the correction performance are various. W2CSpace cannot achieves best performance on a specific parameter combination, but the overall performance is comparable attractive. Besides, we notice that the W2CSpace with a larger size of the coordinates performs better than the smaller one, where the larger W2CSpace is advanced in 8 indexes. The advantages of enlarging the number k are not obvious, where W2CSpace with $n = 100$ and $k = 2000$ is the most advanced combination for the correction task. However, W2CSpace with $k = 1500$ and $k = 3000$ is also a good choice for correction. Generally, we think the introduction of the merge matrix balances the difference of k .

3.4.2 Results of Chinese Sentiment Classification

We evaluate the sentiment classification performance with the classification accuracy. In Table 3, the results of the sentiment classification are illustrated. Specifically, even the other pre-training models partly perform better than BERT, our method achieves improvements on both ChnSenti and Weibo100K datasets as 0.04 and 0.09 points.

For the choice of n and k numbers, our method achieves advanced performance with the combi-

[†]While the other comparison methods are trained on HybirdSet, PLOME additionally pre-trained on a 600 times larger dataset compared with HybirdSet. We uniquely highlight the advanced index of PLOME with **bold italic font**.

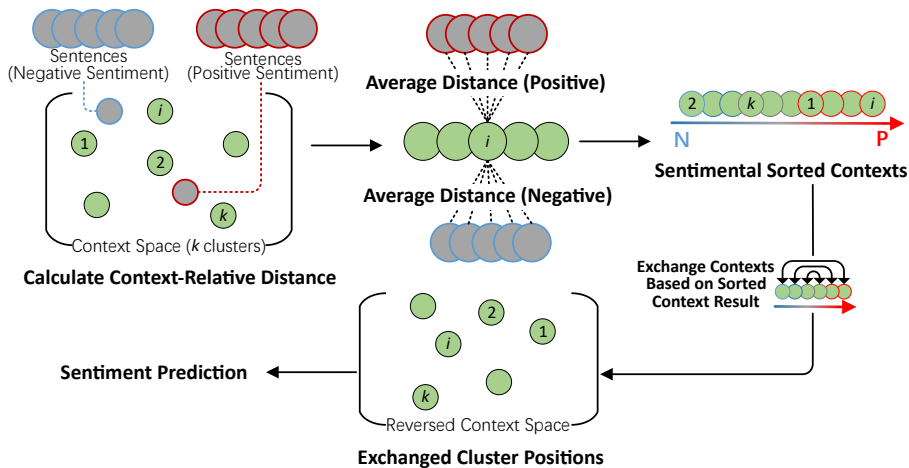


Figure 2: Interpretable Analysis Procedure. First, we calculate the context-relative distance for both ‘Positive’ and ‘Negative’ sentiments. Then we use the average distance between each context and the ‘Positive/Negative’ sentences to rank the contexts from negative to positive sentiment. By exchanging the positions of the context clusters in the W2CSpace, the contexts are reversed to opposite sentiment based on their average distances. Finally, we apply the reversed context to the sentiment classification task and analyze the changes in the sentiment classification results.

nation of $n = 100$ and $k = 800$ for sentiment classification, which is different from the correction experiment. Besides, similar to the tendency in the correction experiment, a larger n number enables a small improvement in performance.

3.5 Interpretable Analysis

The interpretable machine learning methods are considered with a transparent decision process (Molnar et al., 2020; Verhagen et al., 2021). However, the method interpretability is hard to define. The rule-based approaches, widely regarded as the interpretable methods, are especially advanced in the controllability that is one of the most significant characteristics of interpretability (Lee et al., 2017; Tian et al., 2019; Tripathy et al., 2020).

The context-relative distance, the key in our interpretable language modeling process, originated from the relativity between the input sentences and the context in W2CSpace. Therefore, we design an analysis focused on the interpretability of context-relative distance by means of Chinese sentiment classification task. The procedure of interpretable analysis is exhibited in Figure 2. Ideally, the context-relative distance correlates with sentiment, e.g., the shorter distance indicates stronger relativity between the context and the labeled sentiment. If sentiment prediction results change correspondingly after reversing the context space, it convincingly shows that (1) the interpretable knowledge from AKN is integrated into W2CSpace, (2) the feature mapping and context abstraction pro-

cesses are conducted reasonably, and (3) the distance within W2CSpace is interpretable and associated with the emotion conveyed by the input text.

The interpretable analysis result on Weibo100K is exhibited in Table 4. OA is the original accuracy, CA is the sentiment classification accuracy after modification, and RA is the reversing accuracy for the sentences that successfully reverse the sentiment labels after modification. Because language modeling mainly relies on context semantics, the predicted sentiments after modification should be reversed compared with the original prediction. From the interpretable results, the values of RA approximate to 100%, which indicates that the predicted sentiments are mostly reversed and matches our expectation. *The predictable changes of the accuracy reflects the controllability of our method and the interpretability of W2CSpace.* Besides, from the perspective of model structure, the ideal transparent method enables a completely controllable decision process from input to output. And in our method, even though some parts are still in black-box, but RA reflects the interpretability between the decision processes from W2CSpace to the output and from input to the output. The interpretable contexts are consistent with the linguistics logic in input articles and serves as the agent to cooperate with the articles to realize the controllable process. The value of RA does not directly indicate the interpretability of our method, but *but the more approximate to 1, the more semantically explainable of W2CSpace and its context.*

Table 4: Interpretable results.

n	k	OA	CA	RA
50	100	98.24	3.69	96.62
	500	98.30	2.96	98.12
	1000	98.23	1.75	99.90
100	100	98.25	1.88	99.81
	500	98.31	3.79	96.78
	1000	98.28	2.76	98.48

4 Conclusion

An interpretable language model is proposed by constructing a word-context-coupled space in this study. Within W2CSpace, (1) the uninterpretable neural representation in BERT is regulated by interpretable associative knowledge relations, (2) an intermediate representation space with reasonable interpretable semantics is designed, (3) an interpretable semantic feature is introduced based on intermediate representation space for downstream classifiers. Obviously, the above strategies bring a strong generalization ability for the interpretable pre-trained language modeling process. Besides, for the potential risk preventing, the interpretable machine learning method is introduced for migrating the adverse affects from the black-box structure. Moreover, in our method, the controllable decision process realize the regulation for the illegal language inputting by controlling the related context, and the strong cooperation between pre-trained models and W2CSpace can protect the parameter privacy from the data stealing.

Nevertheless, W2CSpace is unable to directly handle high-level semantics, including sentences, paragraphs, and so on. Even the word-level language models act as the mainstream methods in NLP, the above limitation should be further considered in the future. Besides, restricted by our knowledge and efforts, the main experiments cannot cover all common tasks and all pre-trained models in NLP. Relatively, the token- and sequence-level classifications have demonstrated attractive experimental performance on most NLP tasks. Next, we also plan to extend W2CSpace to more NLP tasks and find its more specific value.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62272201, and 61872166; in part by the

Six Talent Peaks Project of Jiangsu Province under Grant 2019 XYDXX-161.

References

- Lisa Feldman Barrett, Kristen A. Lindquist, and Maria Gendron. 2007. [Language as context for the perception of emotion](#). *Trends in Cognitive Sciences*, 11(8):327–332.
- Lisa Feldman Barrett, Batja Mesquita, and Maria Gendron. 2011. [Context in emotion perception](#). *Current Directions in Psychological Science*, 20(5):286–290.
- Katharina Beckh, Sebastian Müller, Matthias Jakobs, Vanessa Toborek, Hanxiao Tan, Raphael Fischer, Pascal Welke, Sebastian Houben, and Laura von Rüden. 2021. [Explainable machine learning with prior knowledge: An overview](#). *CoRR*, abs/2105.10172.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Zhi Chen, Yijie Bei, and Cynthia Rudin. 2020. [Concept whitening for interpretable image recognition](#). *Nature Machine Intelligence*, 2(12):772–782.
- Xingyi Cheng, Weidi Xu, Kunlong Chen, Shaohua Jiang, Feng Wang, Taifeng Wang, Wei Chu, and Yuan Qi. 2020. [SpellGCN: Incorporating phonological and visual similarities into language models for Chinese spelling check](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 871–881, Online. Association for Computational Linguistics.
- Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. [Stargan: Unified generative adversarial networks for multi-domain image-to-image translation](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.

- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, and Ziqing Yang. 2021. [Pre-training with whole word masking for chinese bert](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3504–3514.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jianfeng Gao, Michel Galley, and Lihong Li. 2018. [Neural approaches to conversational ai](#). In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '18*, page 1371–1374, New York, NY, USA. Association for Computing Machinery.
- Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. 2020. [A survey of deep learning techniques for autonomous driving](#). *Journal of Field Robotics*, 37(3):362–386.
- Greg Hamerly and Charles Elkan. 2003. [Learning the k in k-means](#). In *Advances in Neural Information Processing Systems*, volume 16. MIT Press.
- J. A. Hartigan and M. A. Wong. 1979. [Algorithm as 136: A k-means clustering algorithm](#). *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108. Full publication date: 1979.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yuzhong Hong, Xiangyu Yu, Neng He, Nan Liu, and Junhui Liu. 2019. [FASpell: A fast, adaptable, simple, powerful Chinese spell checker based on DAE-decoder paradigm](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 160–169, Hong Kong, China. Association for Computational Linguistics.
- Hyeju Jang, Seojin Bang, Wen Xiao, Giuseppe Carenini, Raymond Ng, and Youngji Lee. 2021. [KW-ATTN: Knowledge infused attention for accurate and interpretable text classification](#). In *Proceedings of Deep Learning Inside Out (DeeLIO): The 2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 96–107, Online. Association for Computational Linguistics.
- John R Josephson and Susan G Josephson. 1996. *Abductive inference: Computation, philosophy, technology*. Cambridge University Press.
- Masahiro Kaneko, Sho Takase, Ayana Niwa, and Naoaki Okazaki. 2022. [Interpretability for language learners using example-based grammatical error correction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7176–7187, Dublin, Ireland. Association for Computational Linguistics.
- Wahab Khan, Ali Daud, Jamal A Nasir, and Tehmina Amjad. 2016. A survey on the state-of-the-art machine learning models in the context of nlp. *Kuwait journal of Science*, 43(4).
- Ronak Kosti, Jose M. Alvarez, Adria Recasens, and Agata Lapedriza. 2017. Emotion recognition in context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Kimin Lee, Jaehyung Kim, Song Chong, and Jinwoo Shin. 2017. [Making stochastic neural networks from deterministic ones](#).
- Seonghyeon Lee, Dongha Lee, Seongbo Jang, and Hwanjo Yu. 2022. [Toward interpretable semantic textual similarity via optimal transport-based contrastive sentence learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5969–5979, Dublin, Ireland. Association for Computational Linguistics.
- Xuhong Li, Haoyi Xiong, Xingjian Li, Xuanyu Wu, Xiao Zhang, Ji Liu, Jiang Bian, and Dejing Dou. 2022a. [Interpretable deep learning: interpretation, interpretability, trustworthiness, and beyond](#). *Knowledge and Information Systems*, 64(12):3197–3234.
- Yanzeng Li, Jiangxia Cao, Xin Cong, Zhenyu Zhang, Bowen Yu, Hongsong Zhu, and Tingwen Liu. 2022b. [Enhancing Chinese pre-trained language model via heterogeneous linguistics graph](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1986–1996, Dublin, Ireland. Association for Computational Linguistics.
- Yanzeng Li, Bowen Yu, Xue Mengge, and Tingwen Liu. 2020. [Enhancing pre-trained Chinese character representation with word-aligned attention](#). In *Proceedings of the 58th Annual Meeting of the Association*

- for *Computational Linguistics*, pages 3442–3448, Online. Association for Computational Linguistics.
- Yulin Li, Zhenping Xie, and Fanyu Wang. 2022c. An associative knowledge network model for interpretable semantic representation of noun context. *Complex & Intelligent Systems*, 8(6):5265–5285.
- Shulin Liu, Tao Yang, Tianchi Yue, Feng Zhang, and Di Wang. 2021. PLOME: Pre-training with misspelled knowledge for Chinese spelling correction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2991–3000, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019b. Knowledge aware conversation generation with explainable reasoning over augmented graphs. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1782–1792, Hong Kong, China. Association for Computational Linguistics.
- Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267:1–38.
- Christoph Molnar, Giuseppe Casalicchio, and Bernd Bischl. 2020. Interpretable machine learning – a brief history, state-of-the-art and challenges. In *ECML PKDD 2020 Workshops*, pages 417–431, Cham. Springer International Publishing.
- Minh Nguyen, Gia H. Ngo, and Nancy F. Chen. 2021. Domain-shift conditioning using adaptable filtering via hierarchical embeddings for robust chinese spell check. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2027–2036.
- Yasumasa Onoe and Greg Durrett. 2020. Interpretable entity representations through large-scale typing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 612–624, Online. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1135–1144, New York, NY, USA. Association for Computing Machinery.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215.
- Sergei Rybakov, Mohammad Lotfollahi, Fabian Theis, and Alexander Wolf. 2020. Learning interpretable latent autoencoder representations with annotations of feature sets.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Yapeng Tian, Chenxiao Guan, Goodman Justin, Marc Moore, and Chenliang Xu. 2019. Audio-visual interpretable and controllable video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Soumya Tripathy, Juho Kannala, and Esa Rahtu. 2020. Icfac: Interpretable and controllable face reenactment using gans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Yuen-Hsien Tseng, Lung-Hao Lee, Li-Ping Chang, and Hsin-Hsi Chen. 2015. Introduction to SIGHAN 2015 bake-off for Chinese spelling check. In *Proceedings of the Eighth SIGHAN Workshop on Chinese Language Processing*, pages 32–37, Beijing, China. Association for Computational Linguistics.
- Ruben S. Verhagen, Mark A. Neerincx, and Myrthe L. Tielman. 2021. A two-dimensional explanation framework to classify ai as incomprehensible, interpretable, or understandable. In *Explainable and Transparent AI and Multi-Agent Systems*, pages 119–138, Cham. Springer International Publishing.
- Jesse Vig. 2019. A multiscale visualization of attention in the transformer model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, Florence, Italy. Association for Computational Linguistics.
- Dingmin Wang, Yan Song, Jing Li, Jialong Han, and Haisong Zhang. 2018. A hybrid approach to automatic corpus generation for Chinese spelling check.

- In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2517–2527, Brussels, Belgium. Association for Computational Linguistics.
- Dingmin Wang, Yi Tay, and Li Zhong. 2019. [Confusionset-guided pointer networks for Chinese spelling check](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5780–5785, Florence, Italy. Association for Computational Linguistics.
- Fanyu Wang, Huihui Shao, and Zhenping Xie. 2023. [AxBERT: An explainable chinese spelling correction method driven by associative knowledge network](#).
- Yuntao Wang, Zhou Su, Ning Zhang, Rui Xing, Dongxiao Liu, Tom H Luan, and Xuemin Shen. 2022. A survey on metaverse: Fundamentals, security, and privacy. *IEEE Communications Surveys & Tutorials*.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A Chinese language understanding evaluation benchmark](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Shaohua Zhang, Haoran Huang, Jicong Liu, and Hang Li. 2020. [Spelling error correction with soft-masked BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 882–890, Online. Association for Computational Linguistics.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
In Section 3.4, we discuss the interpretability of our method from the perspectives of model structure and controllability; In the conclusion, we discuss the limitation of our experiments.
- A2. Did you discuss any potential risks of your work?
Conclusion.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract; Section 1.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

The code of our work is available and will be presented after acceptance.

- B1. Did you cite the creators of artifacts you used?
We use the open-sourced dataset, and cite them in Section 3.2.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
The license of our code will be presented together with our code after acceptance.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We conduct our experiments on the same task as the existing works, which is completely applicable.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We use the dataset which widely adopted in past works, which is checked before.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
The datasets are constructed for Chinese text evaluation, which is claimed in Section 3.2.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Table 1.

C Did you run computational experiments?

Section 3.4.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 3.1; Section 3.2.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 3.4.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 3.4.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 3.1.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.