

RoChBert: Towards Robust BERT Fine-tuning for Chinese

Zihan Zhang^{1*} Jinfeng Li^{2*} Ning Shi^{2,3} Bo Yuan² Xiangyu Liu²
Rong Zhang² Hui Xue² Donghong Sun¹ Chao Zhang^{1†}

¹Tsinghua University; BNRist; Zhongguancun Lab ²Alibaba Group

³Alberta Machine Intelligence Institute, Dept. of Computing Science, University of Alberta

zhangzih19@tsinghua.org.cn, ning.shi@ualberta.ca

{jinfengli.ljf, qiufu.yb, eason.lxy, stone.zhangr

hui.xueh}@alibaba-inc.com, {chaoz, sundonghong}@tsinghua.edu.cn

Abstract

Despite of the superb performance on a wide range of tasks, pre-trained language models (e.g., BERT) have been proved vulnerable to adversarial texts. In this paper, we present RoChBert, a framework to build more Robust BERT-based models by utilizing a more comprehensive adversarial graph to fuse Chinese phonetic and glyph features into pre-trained representations during fine-tuning. Inspired by curriculum learning, we further propose to augment the training dataset with adversarial texts in combination with intermediate samples. Extensive experiments demonstrate that RoChBert outperforms previous methods in significant ways: (i) *robust* – RoChBert greatly improves the model robustness without sacrificing accuracy on benign texts. Specifically, the defense lowers the success rates of unlimited and limited attacks by 59.43% and 39.33% respectively, while remaining accuracy of 93.30%; (ii) *flexible* – RoChBert can easily extend to various language models to solve different downstream tasks with excellent performance; and (iii) *efficient* – RoChBert can be directly applied to the fine-tuning stage without pre-training language model from scratch, and the proposed data augmentation method is also low-cost.¹

1 Introduction

The emergence of pre-trained language models (PLMs) has revolutionized natural language processing (NLP) to a new era. As a result, large-scale PLMs like BERT (Devlin et al., 2018) have become the mainstream models for various downstream tasks including text classification (Sun et al., 2019), question answering (Herzig et al., 2020) and machine translation (Zhu et al., 2020), and have

* Equal contribution.

† Corresponding author.

¹Our code will be available at <https://github.com/zzh-z/RoChBERT>.




Figure 1: An example of Chinese adversarial text. “博彩”和“菠菜” share the same pronunciation of "bo cai". “微”和“薇” share the same pronunciation of "wei" and similar components of “微”.

drastically boosted their performance. In the Chinese domain, a wide variety of classical NLP tasks also benefit from the BERT-based PLMs.

Despite the impressive performance, BERT-based models have been proved to be vulnerable to maliciously generated adversarial texts (Li et al., 2019; Garg and Ramakrishnan, 2020; Li et al., 2020). Meanwhile, in the real-world scenarios, adversaries usually generate obfuscated texts, i.e., manually crafted adversarial texts to bypass online security-sensitive systems, which has posed severe physical threats to the deployed systems. In contrast to the alphabet languages such as English, the meaning of individual Chinese character can be implied from its pronunciation and glyph. Thus, substituting characters with others similar in characteristics can hardly bring impact to the context understanding. Hence, unlike replacing words with synonyms to mislead English models, adversaries targeting Chinese models prefer substituting characters with others sharing similar pronunciation or glyph, as illustrated in Figure 1.

Several defense methods such as adversarial training (AT) (Si et al., 2021) and adversarial detection (Bao et al., 2021) have been proposed recently. However, most of them focus on solving English adversarial texts. Based on Chinese characteristics, some Chinese-specific methods are proposed. For example, ChineseBERT (Sun et al., 2021) incor-

porates Chinese pronunciation and glyph features into model pre-training, and it has achieved SOTA performance on many Chinese NLP tasks. But ChineseBERT needs to be pre-trained from scratch and its number of parameters is quite large. It has to query the corresponding pinyin and glyph of input characters at each inference, which makes the training and prediction very slow.

To solve the challenges above, we propose *RoChBert*, a lightweight and flexible method to strengthen robustness for Chinese BERT-based models. Firstly, we generate an updated adversarial graph based on AdvGraph (Li et al., 2021) to capture the phonetic and glyph relationships between characters, which are commonly exploited by adversarial texts. As current AdvGraph only incorporates 3,000 commonly used characters and fails to capture certain glyph relationships, we update it to cover more characters and more glyph relationships. Then, we use adversarial graph to learn the representation of characters based on Chinese pronunciation and glyph. The graph mainly serves as embedding weights within *RoChBert*. Meanwhile, we leverage the target model’s hidden states of last layer as pre-trained representation. These two kinds of embedding will then be fused in the fine-tuning of target model with specified downstream NLP tasks. Secondly, we design a novel data augmentation method inspired by curriculum learning (Bengio et al., 2009) for better fusion, which is proved to be more efficient than the traditional AT. Specifically, we add both intermediate and adversarial texts into training datasets, which is very computationally efficient and will not decrease the accuracy on benign texts compared with traditional AT. To the best of our knowledge, this is the first work to strengthen the robustness of pre-trained Chinese BERT-based models during fine-tuning.

Our contributions are summarized as follows:

1. We present *RoChBert*, a plug-in method to strengthen the robustness of BERT-based models during fine-tuning by incorporating adversarial knowledge.
2. We update the adversarial graph and design an efficient data augmentation method, which considers intermediate samples for more effective use of adversarial texts.
3. Extensive experiments show that *RoChBert* can drastically improve the robustness of target models against both known and unknown

attacks without impacting the performance on benign texts. In addition, *RoChBert* is flexible and efficient as it can be applied to most pre-trained models with specified downstream tasks while both fine-tuning and data augmentation are low-cost.

2 Related Work

Pre-trained language model. The applications of PLMs have achieved great success on various downstream NLP tasks and avoid training a new model from scratch. BERT is first introduced to learn universal language representations via masked language model objective and next sentence prediction task, which is then improved in the following works such as RoBERTa (Liu et al., 2019) and ALBERT (Lan et al., 2019). In the Chinese NLP domain, Sun et al. (2021) proposed ChineseBERT, which incorporates both the glyph and pinyin features of Chinese characters into language model pre-training, and also gained SOTA performance on many Chinese NLP tasks. However, these works mainly focus on improving the model performance on benign texts, and the effort of enhancing model robustness is fairly limited.

Exploration on the model robustness. To further explore the vulnerabilities of NLP models in the real adversarial scenarios, a plenty of black-box attacks have been proposed under the practical assumption that adversary only has query access to the target models without any model knowledge (Li et al., 2019; Ren et al., 2019; Garg and Ramakrishnan, 2020). To defend against such attacks, countermeasures like AT and adversarial detection have been proposed to mitigate the inherent model vulnerability. Concretely, AT usually refers to retraining the target model by mixing adversarial texts into the original training dataset, which could be viewed as a kind of data augmentation (Si et al., 2021; Ng et al., 2020). Adversarial detection is to check whether the input contains spelling errors or adversarial perturbation, and then restore it to the benign counterparts (Zhang et al., 2020a; Bao et al., 2021). Such methods both exhibit great efficacy in the English NLP domain, but they are hard to extend to the Chinese domain directly due to the language differences. Hence, many studies have tried to design specific defense in terms of the unique property of Chinese. For instance, Wang et al. (2018) and Cheng et al. (2020) improved the Chinese-specific spelling check using the phonetic and glyph information. Li et al. (2021) pro-

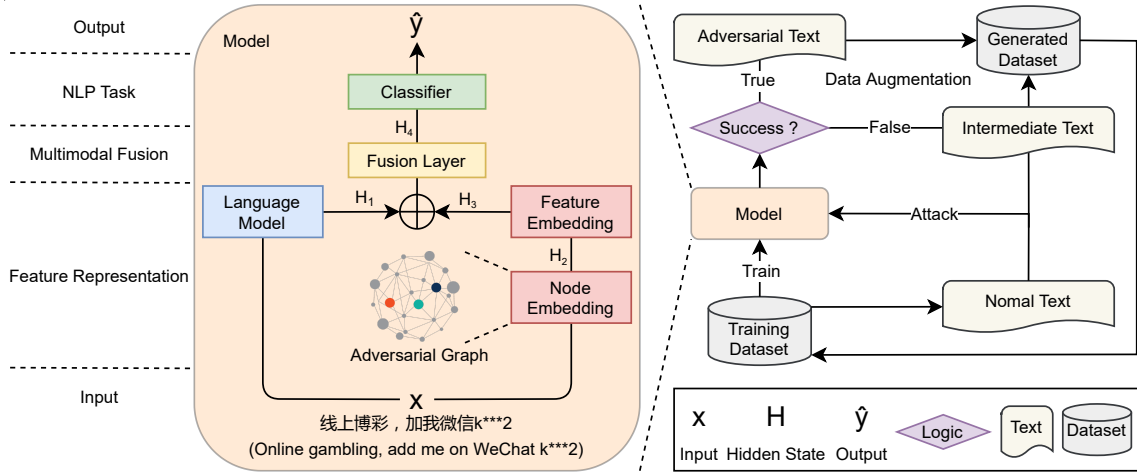


Figure 2: Overview of RoChBert. The left part is the model architecture, and the right part is its workflow.

posed AdvGraph that involves an undirected graph to model the phonetic and glyph adversarial relationships among Chinese characters and improves the robustness of several traditional models. (Su et al., 2022) proposed RoCBERT to enhance the robustness by pre-training the model from scratch with adversarial texts covering combinations of various Chinese-specific attacks, which may not be maintained in the downstream tasks.

3 Methodology

Given a benign input text $x \in \mathcal{X}$ with the ground-truth label $y \in \mathcal{Y}$, a classifier \mathcal{F} has learned the mapping $f: \mathcal{X} \rightarrow \mathcal{Y}$, which maps x from the feature space \mathcal{X} to the label space \mathcal{Y} as $\mathcal{F}_f(x) = y$. The adversary intend to obtain an adversarial text x^* by adding some perturbations Δx on x with $\|\Delta x\|_p < \epsilon$, and deceive \mathcal{F} into making a wrong prediction on x^* , i.e., $\mathcal{F}_f(x^*) = \hat{y}$ and $y \neq \hat{y}$. In this paper, we aim to improve the inherent robustness of \mathcal{F} so to get a more robust model \mathcal{F}' such that \mathcal{F}' can resist the adversarial text x^* , i.e., $\mathcal{F}'_f(x^*) = y$.

Overview. As illustrated in Figure 2, we first generate an updated adversarial graph based on phonetic and glyph relationships between all commonly used Chinese characters. Then, we adopt node2vec (Grover and Leskovec, 2016) to learn the representation of each character x_i , which will then be used as the node embedding. For x_i in text x , its node embedding is concatenated and passed through feature extraction module. Simultaneously, x is also fed into the target PLM. We utilize the hidden states of last layer as the pre-trained representation of x . After concatenating pre-trained representation and feature embedding, we use a multimodal

fusion module to further fuse the information from two channels. Finally, the fused representation can be used for most of the downstream tasks. Moreover, to better enhance the fusion process, we have designed an efficient data augmentation method, which makes full use of adversarial texts and helps the model further understand which characters may be similar in pronunciation and glyph.

3.1 Adversarial Graph

Adversarial graph is first presented in (Li et al., 2021) as AdvGraph to improve robustness of classic deep learning model (e.g., TextCNN, LSTM). In AdvGraph, each node represents a Chinese character and each edge represents the phonetic or glyph relationship between two nodes.

However, the aforementioned AdvGraph only covers 3,000 commonly used characters. Moreover, the glyph relation in AdvGraph is constructed heavily relied on a g-CNN model which first converts each character into a image and then extracts the visual feature as glyph representation. It hence pays more attention to characters with similar structures, which will affect the representation learning of those characters with dissimilar structures. For instance, “微” and “徽” is considered to be closer by g-CNN since they are both left-middle-right radical structure, while “微” and “薇” is considered dissimilar since “薇” is up-down radical structure. This is because when converting to images, different radicals may cause the position of characters to shift, which will have a great impact on CNN.

To tackle the above problem, we intend to advance the adversarial graph as follows. First, we adopt the algorithm in Argot (Zhang et al., 2020b)

to generate similar candidate sets for each character, with a combination of candidates in AdvGraph and StoneSkipping (Jiang et al., 2019). Next, we utilize stroke code to calculate similarity score for candidates. Stroke refers to a line that is continuously written at a time, which is the smallest unit of Chinese characters. Each number in stroke code represents a specific stroke, so it can be used to capture the same components between two characters. Concretely, we calculate the *longest common substring* between two codes as a part of our similarity score in addition to the distance of representation learned in g-CNN. For each character in candidate set, once the similarity score exceeds the given threshold, we assume them as the similar pair. Then, we add an edge between two characters for each similar pair, indicating they share similar pronunciation or glyph.

Finally, our adversarial graph incorporates 7,707 characters, and it has 109,706 edges, 32,894 of which is phonetic relationship, and the 81,108 is glyph relationship. Note that the Chinese pronunciation, glyph and their relationships in our adversarial graph are general knowledge. We model the adversarial relationship in advance so that RoChBert can protect target models against both known and unknown attacks.

3.2 Multimodal Fusion

In AdvGraph, graph embedding and semantic embedding are simply concatenated to generate fused representation and there is almost no interaction between these two kinds of feature. Hence, we design a multimodal fusion module to make them more fully integrated during training process.

We first feed x of length l into PLM with hidden size d_1 to obtain pre-trained presentation. Assume the target language model as the function \mathcal{F}_{LM} , we get the hidden states H_1 of its last layer:

$$H_1 = \mathcal{F}_{LM}(x) \in \mathbb{R}^{l \times d_1}. \quad (1)$$

Then we pass x into Node Embedding Layer \mathcal{F}_G whose weights are learned from adversarial graph to get H_2 . We freeze the weights of \mathcal{F}_G to preserve the independent phonetic and glyph features without being impacted by semantic information during training. The size of the vector H_2 is d_2 , which is decided by the parameter of *node2vec* algorithm:

$$H_2 = \mathcal{F}_G(x) \in \mathbb{R}^{l \times d_2}. \quad (2)$$

Since the focus of model should be different on the

benign and adversarial text, we introduce Transformer Encoder layer \mathcal{T}_1 to help model capture the differences. There still exists the *out-of-vocabulary* problem in the node embedding though we have updated the original AdvGraph, since if the character is not similar with any character or its frequency is too low, it will not be included in the graph. In addition, node embedding doesn't incorporate English letters, numbers, punctuation, and special symbols, it will be hard to align with pre-trained embedding H_1 and it may even bring negative effects. We use flattening to preserve its phonetic and glyph features and alleviate the problem that some features may be too sparse caused by *out-of-vocabulary*. Then we pass it into a linear layer with weights \mathbf{W} and bias b to map it into a vector of size d_1 as sequence feature representation H_3 :

$$H_3 = \mathbf{W}^\top (\mathcal{T}_1(H_2)) + b \in \mathbb{R}^{d_1}. \quad (3)$$

In order to facilitate the subsequent concatenation, we repeat the vector l times to obtain $H'_3 \in \mathbb{R}^{l \times d_1}$ as the final feature embedding.

We first fuse these two kinds of embedding by concatenating H_1 and H'_3 . To make them better interact with each other during training, we apply Transformer Encoder layer \mathcal{T}_2 for the fusion. The final embedding H_4 of input x is given by:

$$H_4 = \mathcal{T}_2(H_1 \oplus H'_3) \in \mathbb{R}^{l \times d_1}. \quad (4)$$

Finally, we flatten H_4 and feed it into a classifier to get output $\hat{y} \in \mathcal{Y}$.

3.3 Efficient Data Augmentation

It has been reported that data augmentation can improve the robustness of models (Si et al., 2021; Ng et al., 2020). The intuitive method is to retrain the target model by mixing the generated adversarial texts into the original training dataset. However, most of the conventional AT methods are usually very cost-expensive. For instance, in the black-box attacks, it often requires an extremely large amount of queries for accessing the target models to determine the attack direction in each iteration when generating just one adversarial text.

Conventional AT typically involves the entire dataset to conduct attacks while only collecting successful adversarial samples for augmentation. To efficiently utilize the generated texts and obtain more samples in the shortest time, we have made some improvements to the traditional AT.

Algorithm 1 : The detail of data augmentation.

Input: Training dataset D , the target classifier \mathcal{F} with mapping f .
Output: New training dataset D_{ag} .

```
1: for  $x \in D$  do
2:    $tmp \leftarrow \{\}$ 
3:    $\hat{y} = \mathcal{F}_f(x)$ 
4:   if  $\hat{y}$  is not the ground-truth label then
5:     continue
6:   end if
7:    $x^* = x, \hat{y}^* = \hat{y}$ 
8:   while  $\hat{y}^* \neq \hat{y}$  do
9:      $x^* = x^* + \Delta x \triangleright$  According to attack algorithms
10:     $\hat{y}^* = \mathcal{F}_f(x^*)$ 
11:     $tmp \leftarrow tmp \cup \{x^*\}$ 
12:    if all the words in  $x$  are modified then
13:      break
14:    end if
15:  end while
16:  if  $\hat{y}^* \neq \hat{y}$  and  $\|x^* - x\|_p < \epsilon_{max}$  then
17:     $D_{ag} \leftarrow D_{ag} \cup tmp$ 
18:  end if
19:  if  $size(D_{ag}) > size(D)$  then
20:     $D_{ag} \leftarrow D_{ag} \cup D$ 
21:    break
22:  end if
23: end for
24: return  $D_{ag}$ ;
```

For generating adversarial texts with fewest perturbations, there are various strategies designed in the attacks to determine the added perturbation in each iteration. Usually, this perturbation can cause the most significant drop on the output probability of its original label, which is most hopeful to mislead the target model within the next iterations among all the candidate generated texts. Therefore, during the attack, many intermediate texts will be generated before the final adversarial text is generated. Although the labels of these intermediate texts are still correct, their confidences have declined to varying degrees. Hence, to some extent, in addition to adversarial texts, the intermediate texts can also help models make firm decisions.

We generate augmented data based on the training dataset D . For each x in D , we perturb it and collect the intermediate texts tmp . If the generated x^* can deceive the model \mathcal{F}_f , we will incorporate tmp (the final generated x^* is also in tmp) into new dataset D_{ag} . Specifically, to ensure the quality of texts, we will limit the modification rate ϵ_{max} during the generation process. If the size of D_{ag} is larger than the size of D , we assume we have got enough texts and early stop the process. Finally, we add D into D_{ag} and D_{ag} will be our new training dataset. The details of the algorithm are presented in Algorithm 1, in which ϵ_{max} and $size(D_{ag})$ have

the following relationship:

$$size(D_{ag}) \leq l_{x_1^*} \times \epsilon_1 + l_{x_2^*} \times \epsilon_2 + \dots + l_{x_n^*} \times \epsilon_n, \\ \text{s.t. } \epsilon_i \leq \epsilon_{max}, i \in 1 \dots n$$

$l_{x_i^*}$ denotes the length of adversarial text x_i^* . n denotes the number of generated adversarial texts.

As the base of perturbation is word, we can get:

$$size(D_{ag}) \leq (l_{x_1^*} + l_{x_2^*} + \dots + l_{x_n^*}) \times \epsilon_{max} \\ \leq l_{avg} \times n \times \epsilon_{max}. \quad (5)$$

Therefore we have:

$$n \geq \frac{size(D_{ag})}{l_{avg} \times \epsilon_{max}}. \quad (6)$$

So n increases as l_{avg} and ϵ_{max} decrease. In the augmentation stage, more used adversarial texts (n is bigger) will result in better model robustness, but the more computation is needed. If we set small ϵ_{max} , we will get high quality adversarial texts. They will promote the model robustness with less impact on accuracy when added into training dataset. On the contrary, when we set big ϵ_{max} , the adversarial texts we need to generate become fewer, which reflects the high efficiency of RoChBert.

So when the accuracy of the target model is relatively high and the average length of texts is long, we recommend to use smaller ϵ_{max} to make n bigger and improve robustness. If the average length of texts is short, n will increase, so we urge to use bigger ϵ_{max} for efficiency boosting. And when the model accuracy is not very high, we suggest to use smaller ϵ_{max} to keep the accuracy as high as possible.

Our method inherits the main idea of curriculum training (Bengio et al., 2009), and the model can learn the features of the adversarial text at different stages. Moreover, it can improve model’s robustness without compromising the accuracy on benign texts, which can be proved in Section 4.2.

4 Experiment

4.1 Experiment Settings

Datasets. We use three different tasks, *sentiment analysis*, *text classification* and *natural language inference* to evaluate RoChBert. The datasets are ChnSentiCorp², DMSC³, THUCNews (Sun et al., 2016) and OCNLI⁴. The details are shown

²https://github.com/pengming617/bert_classification/tree/master/data

³<https://www.kaggle.com/utmhikari/doubanmovieshortcomments/>

⁴<https://github.com/cluebenchmark/OCNLI>

Model	Chnsenti.	DMSC	THUC.	OCNLI
ChineseBERT	95.25	92.95	97.87	73.20
BERT _{base}	95.33	93.02	98.07	71.57
+SC	94.42	92.85	98.07	70.57
+RoChBert (PWWS)	95.58	93.05	97.87	67.76
+RoChBert (TextBugger)	95.83	92.75	98.00	67.34
+RoChBert (Random)	95.92	92.70	98.13	70.25
BERT _{wwm}	94.58	92.51	97.87	70.33
+SC	93.58	92.45	97.93	69.08
+RoChBert (PWWS)	94.92	92.70	97.93	67.23
+RoChBert (TextBugger)	95.75	93.30	97.80	67.71
+RoChBert (Random)	95.25	91.45	98.00	70.09
BERT _{wwm/ext}	96.00	93.29	97.73	71.16
+SC	95.00	93.20	97.80	70.68
+RoChBert (PWWS)	95.58	94.00	97.80	68.31
+RoChBert (TextBugger)	95.42	93.30	97.87	68.65
+RoChBert (Random)	95.83	93.60	97.73	71.40
RoBERTa _{wwm/ext}	95.58	92.89	98.00	71.29
+SC	94.50	92.95	97.87	70.88
+RoChBert (PWWS)	95.58	93.10	98.13	68.31
+RoChBert (TextBugger)	95.83	93.56	97.93	69.09
+RoChBert (Random)	95.50	93.45	98.07	72.45

Table 1: Model performances on benign texts.

in Appendix A.1. For OCNLI, we attack the two sentences, *premise* and *hypothesis* separately. Here, we only present the corresponding results of *premise*, the other part is shown in Appendix A.4.

Setup. As the target models are Chinese BERT-based models and the perturbation strategies are quite different from those for English, we cannot use the specialized BERT attack methods like BERT-ATTACK (Li et al., 2020). Hence, we utilize three widely used attacks, i.e., PWWS (Ren et al., 2019), TextBugger (Li et al., 2019) and random attack in the black-box setting to evaluate the robustness of RoChBert. The details are shown in Appendix A.2. We leverage the characters with similar glyph or pronunciation to form words and substitute the original ones after locating the important parts by corresponding attack algorithms. The attacks are conducted on 1,000 texts sampled from test set. We use base BERT (BERT_{base}), BERT trained with whole word masking strategy (BERT_{wwm}) and with extended data (BERT_{wwm/ext}), and RoBERTa (RoBERTa_{wwm/ext}) as the target models.⁵ We also take ChineseBERT which incorporates pronunciation and glyph features into pre-training, as a baseline. In addition, we compare RoChBert with the Chinese spelling corrector (SC)⁶ method. In SC-based defenses, each input text is firstly restored by a corrector to eliminate errors in the text before being sent into the target model. For fair comparison, the adversarial texts against the SC baseline are generated by

⁵<https://huggingface.co>

⁶<https://github.com/shibing624/pycorrector>

treating corrector and the model as a whole pipeline. And ϵ_{max} in data augmentation are shown in Appendix A.3.

Metrics. We utilize four metrics, i.e., accuracy, modification rate (MR), unlimited attack success rate (UASR) and limited attack success rate (LASR), to comprehensively evaluate RoChBert. Accuracy reflects model’s generalization on benign texts. MR reflects the average percentage of perturbed characters in adversarial texts under unlimited attack setting. UASR means the percentage of texts that can generate adversarial texts successfully without any limitations. To guarantee the quality of adversarial texts, the maximum MR will be 0.2 when calculating LASR.

4.2 Model Performance

We first evaluate the model performances on benign texts since defense shouldn’t compromise the natural generalization of models. As summarized in Table 1, SC reduces the accuracy of models due to its own errors. It is observed that RoChBert achieves relatively high accuracy, which is comparable with the original BERT-based models and ChineseBERT. In some cases, models defended by RoChBert even outperform the baselines. This is mainly because the additional phonetic and glyph features can help model capture semantic information in some degree, which has been proved in ChineseBERT. In addition, the gains on accuracy also reflect the key impact of data augmentation.

4.3 Robustness Against Attack

Effectiveness. Then, we evaluate the efficacy and robustness of RoChBert. The results are shown in Table 2. It is observed that the defend effect of SC on the target model is not obvious, which means that when it is combined with the target model as a whole, the attacker regards them as a black-box model to conduct attack, and can still easily generate adversarial texts. ChineseBERT indeed enhance the robustness compared with original models, but the improvement is very limited while RoChBert brings noticeable reduction on both UASR and LASR. In the best case, UASR is decreased by 75.75% when RoChBert is fine-tuned on RoBERTa_{wwm/ext} with THUCNews and attacked by TextBugger, and LASR is decreased by 42.86% when it is fine-tuned on BERT_{wwm} with DMSC and attacked by TextBugger. This indicates that RoChBert can significantly weaken the attack and enhance robustness of models without

Model	PWWS			TextBugger			Random			PWWS			TextBugger			Random		
	UASR	LASR	MR	UASR	LASR	MR	UASR	LASR	MR	UASR	LASR	MR	UASR	LASR	MR	UASR	LASR	MR
	ChnSetiCorp									THUCNews								
ChineseBERT	79.73	40.97	27.22	93.25	42.67	23.64	54.91	3.38	51.23	71.55	23.44	44.82	69.80	36.23	23.63	78.40	1.13	64.68
BERT _{base}	83.62	67.66	12.96	97.45	69.26	16.12	52.77	8.19	42.85	81.31	25.64	44.60	58.43	43.21	14.87	80.59	2.56	61.42
+SC	82.75	64.86	13.74	96.49	71.25	14.85	54.42	7.56	43.01	80.51	29.18	40.87	59.90	46.02	13.26	79.08	2.55	61.96
+RoChBert	65.18	31.63	29.49	64.45	34.92	20.35	39.49	10.98	37.48	66.35	5.11	63.71	9.66	9.10	8.05	51.17	0.81	66.57
BERT _{wwm}	87.53	56.56	19.86	98.28	64.73	16.89	48.27	6.45	44.91	73.77	36.27	35.17	78.28	46.93	21.07	74.90	2.67	58.81
+SC	84.62	57.42	18.75	97.63	68.60	15.95	50.54	5.59	45.62	76.66	30.19	38.37	72.26	45.45	17.96	73.49	2.97	58.68
+RoChBert	64.42	22.94	41.94	62.24	35.23	19.6	42.31	6.36	44.88	76.60	5.19	64.00	13.09	11.15	11.94	45.77	0.72	64.67
BERT _{wwm/ext}	72.04	42.93	22.21	92.93	53.27	20.16	56.96	5.80	40.44	79.63	21.29	48.05	86.39	29.99	31.85	78.81	1.74	60.79
+SC	75.00	50.53	18.19	90.93	57.59	17.47	57.07	5.91	41.04	82.72	21.98	46.59	79.04	33.54	26.08	79.24	2.76	60.04
+RoChBert	62.75	24.44	37.08	65.74	31.70	23.17	38.45	6.45	44.26	66.80	7.89	60.09	13.10	8.80	16.31	51.07	0.61	66.04
RoBERTa _{wwm/ext}	76.46	44.11	23.25	99.78	54.19	21.49	57.58	5.83	44.14	72.58	17.02	50.12	81.24	30.27	28.46	79.51	1.22	62.27
+SC	80.47	52.23	19.62	98.20	59.66	18.26	55.94	6.90	44.00	81.51	19.41	48.52	77.22	35.55	24.33	79.57	1.23	62.43
+RoChBert	65.85	22.69	39.14	54.18	28.57	20.49	38.59	8.41	36.60	59.92	4.68	63.54	5.49	4.88	7.70	59.04	1.83	60.33
	DMSC									OCNLI								
ChineseBERT	78.76	60.35	16.64	92.20	59.37	18.47	53.30	7.04	48.75	62.57	46.22	17.32	73.78	35.27	25.16	38.92	8.38	40.92
BERT _{base}	76.70	61.06	15.74	78.75	60.19	13.79	56.31	7.66	46.69	58.68	42.29	17.39	65.84	35.95	22.96	40.08	10.06	38.59
+SC	83.24	63.24	17.87	82.49	63.24	13.54	57.51	6.38	47.25	56.50	42.38	15.81	65.73	36.08	21.87	38.88	8.81	38.82
+RoChBert	68.67	46.70	23.66	36.36	29.22	12.43	44.94	10.13	39.21	43.58	29.73	17.90	48.99	24.21	24.53	36.50	5.65	47.04
BERT _{wwm}	95.66	76.66	13.53	99.67	76.33	13.92	52.88	6.30	43.41	56.06	40.42	17.09	64.23	32.54	23.97	36.34	8.03	42.91
+SC	94.46	74.38	13.66	98.91	77.96	13.07	55.70	7.38	43.59	55.33	40.68	16.28	63.30	33.43	22.69	36.56	8.68	41.33
+RoChBert	64.94	50.00	14.00	44.85	33.47	15.94	46.78	11.91	40.22	50.00	30.18	20.79	50.36	20.32	27.54	39.21	7.05	46.48
BERT _{wwm/ext}	85.52	60.30	17.63	99.79	69.31	16.35	57.51	5.90	49.93	62.22	45.56	16.46	69.31	37.08	23.15	40.69	8.61	41.58
+SC	88.41	63.95	16.98	99.36	72.85	15.36	59.12	6.76	48.85	61.79	47.70	15.26	69.32	37.10	21.83	43.10	10.74	40.65
+RoChBert	75.40	46.31	22.37	40.36	29.98	14.65	37.38	6.87	47.99	51.75	32.75	21.27	54.69	20.23	28.24	39.52	7.45	45.33
RoBERTa _{wwm/ext}	69.70	50.76	20.04	83.12	53.90	18.69	52.71	7.79	39.31	66.21	46.02	17.65	80.08	37.09	26.35	40.66	7.01	43.34
+SC	75.54	55.19	18.54	85.82	57.47	17.57	54.87	7.68	40.70	65.10	47.31	17.48	79.03	40.14	24.01	41.66	8.83	40.91
+RoChBert	55.15	28.06	34.80	59.25	35.27	19.52	44.86	7.92	43.98	60.14	35.11	22.33	51.43	21.57	27.31	36.86	7.70	42.60

Table 2: Model performance against different attacks.

sacrificing their performance. Simultaneously, MR of adversarial texts generated against RoChBert is higher, which means that attackers need to modify more words to deceive RoChBert, and the generated texts will be harder to comprehend. In some cases, MR is decreased, this is because the UASR has dropped significantly, adversaries fail to attack even if they modify the texts extensively.

Then, we analyze the relation between MR and ASR by setting the maximum MR from 0.1 to 0.5. The results of TextBugger attack on THUCNews are shown in Figure 3. The other results are in Appendix A.6. The dashed, dotted and solid lines in the same color represent the target model, the model protected by SC and RoChBert respectively. It proves that RoChBert retains robustness against attack facing with the change of MR. Compared with baselines, ASR of models protected by RoChBert increases slightly as the MR grows, indicating that RoChBert can significantly enhance model robustness even in adaptive settings.

The above experiments have shown the robustness of RoChBert against known attacks as it will use the specified attack algorithm during data augmentation. Here we conduct another experiment to demonstrate RoChBert’s performance against unknown attack. We leverage TextBugger to generate adversarial texts and their immediate texts since we have known that TextBugger is the strongest attack

Attack	Dataset	UASR	LASR	MR
PWWS	ChnSetiCorp	64.87	28.06	35.09
	DMSC	63.20	27.27	39.72
	THUCNews	81.08	6.95	61.63
	OCNLI	43.80	30.55	19.79
Random	ChnSetiCorp	41.88	3.48	52.74
	DMSC	51.51	7.03	45.53
	THUCNews	79.86	1.22	64.72
	OCNLI	27.81	8.93	39.42

Table 3: Model performance against unknown attacks. Target models are BERT_{base} and training sets are augmented with adversarial texts generated by TextBugger.

among the three attacks as shown in Table 2. And we assume that using the strongest attack during data augmentation can strengthen the robustness of the model against unknown attacks as much as possible. As shown in Table 3, we can see that both UASR and LASR have decreased greatly while MR increases, indicating that when facing unknown attacks, RoChBert still has the ability to defend. Hence we can add defense to the target model in advance without knowing the specific adversaries. Obviously, we can leverage other stronger or multiple attacks during data augmentation.

Generalizability. It is seen from Table 2 that both the UASR and LASR against each models protected by RoChBert have a different degree of decline. All models are commonly used in Chinese NLP tasks, indicating that RoChBert has good




Figure 3: Relation between maximum MR and ASR.

Model	Acc.	UASR	LASR	MR
BERT _{base}	93.02	78.75	60.19	13.79
+graph	93.05	71.53	50.32	16.40
+aug.	93.85	68.01	48.01	17.56
+graph+aug.	94.15	73.16	19.36	40.69
+RoChBert	92.75	36.36	29.22	12.43

Table 4: Performance of models enhanced by different modules against TextBugger attack.

generalizability on other PLMs, and can be applied on different tasks to defend against various attack.

Efficiency. We assess efficiency of RoChBert via comparing the samples used in data augmentation module with the traditional AT by taking the BERT_{base} fine-tuned on ChnSentiCorp as an example. For AT, we conduct attack against the target model with the whole training set and collect successful adversarial texts as the augmentation data. RoChBert only uses 25.40%, 17.22% and 16.54% of the training dataset when performing PWWS, TextBugger and Random attacks respectively, and achieves more robust models than AT. It proves that the data augmentation in RoChBert is efficient and effective. In addition, it also indicates the key impact of multimodal fusion. The complete results are shown in Appendix A.7.

4.4 Ablation Study

Then we will discuss the effects of different modules, i.e., adversarial graph and data augmentation. To observe the gains on adversarial graph, after obtaining *feature embedding*, we simply concatenate

it with pre-trained representation and feed them into the model. To evaluate the effect of data augmentation, we directly attack the target BERT via TextBugger and collect the the intermediate and adversarial texts by taking the BERT_{base} fine-tuned on DMSC as an example. The results are shown in Table 4. Observe that either using adversarial graph or data augmentation alone can still decrease the UASR and LASR. Then, we utilize both adversarial graph and data augmentation to defend BERT against attack, we can see that compared with using one module alone, UASR rises but LASR drops. When BERT is protected by RoChBert, we can obtain the optimal model in general situation with relative low UASR and LASR. So the two modules we proposed play important role in strengthening robustness of models. In addition, the differences between *+graph+aug* and *+RoChBert* demonstrate that the effect of simply concatenating *feature embedding* H_1 and *pre-trained embedding* H_2 is relatively limited, which is used in the original AdvGraph. And *transformer* in RoChBert makes two kinds of embedding fully interact with each other and helps lower the UASR significantly.




Figure 4: Representation Visualization. (a) is benign texts on BERT_{base}. (b) is adversarial texts on BERT_{base}. (c) is adversarial texts on ChineseBERT. (d) is adversarial texts on BERT_{base}+RoChBert.

4.5 Representation Visualization

To further explore the robustness of different models, we project representation vectors to two dimensional vectors through T-SNE. Taking BERT_{base} on DMSC against TextBugger as an instance, we choose the same 500 samples. In Figure 4(a), the decision boundary for benign texts are quite clear. However in Figure 4(b), with the addition of adversarial texts, the representations are entangled together, making it difficult to distinguish between

the two kinds of data. As shown in Figure 4(c), ChineseBERT does not significantly improve this situation. After being protected by RoChBert, most of the adversarial texts are effectively separated in Figure 4(d), which proves that RoChBert exhibits impressive robustness against attacks.

5 Conclusion

In this paper, we propose RoChBert, a lightweight defense method to strengthen robustness for pre-trained Chinese BERT-based models. First, we update the current adversarial graph and learn the feature representation of characters based on Chinese pronunciation and glyph. Then RoChBert fuse semantic and feature embeddings to fine-tune the target model with specified downstream NLP task, e.g., classification. To enhance fusion, we further introduce an augmentation method inspired by curriculum learning. We evaluate RoChBert against several attack algorithms, the results show that RoChBert greatly strengthen the model robustness without compromising their normal generalization. In addition, RoChBert can be used to improve most Chinese BERT-based models and it is computationally efficient.

Limitations

RoChBert focuses on solving the perturbations that substitutes target character with the ones sharing similar pronunciation or glyph. In fact, there are other perturbations such as substituting with synonyms, splitting Chinese characters and inserting special characters. In most of the Chinese BERT models, characters are often seen as the basic unit of operations. So compared with word-based models, e.g., TextCNN and LSTM, it will be harder for adversaries to attack BERT models successfully only by synonyms substitution. Besides, such kind of perturbations can be easily defended by adversarial training. And the other perturbations are limited, they can be filtered by rules easily. For completeness of the work, we will include these perturbations into the experiments in future work. In addition, there are not as many types of Chinese datasets as English ones, and the quality is uneven. We are temporarily unable to verify RoChBert on more tasks. In the future, we will consider applying our work on other tasks.

Ethics Statement

In our experiments, all adversarial texts are generated from the public benchmark datasets via open-source adversarial attack algorithms. These generated adversarial texts are used only for the purpose of enhancing the robustness of target models and will not be used for any illegal purposes or unsuitable intentions. We will also open source the code of our RoChBert later to aid those users and researchers involved in pre-trained language models in improving the robustness of their models and thus mitigate the potential physical threats brought by adversarial attacks.

Acknowledgements

This work was supported, in part, by the National Key Research and Development Program of China (2021YFB2701000), National Natural Science Foundation of China (61972224), Beijing National Research Center for Information Science and Technology (BNRist) under Grant BNR2022RC01006, Shining Lab and Alibaba Group through Alibaba Research Intern Program.

References

- Rongzhou Bao, Jiayi Wang, and Hai Zhao. 2021. Defending pre-trained language models from adversarial word substitutions without performance sacrifice. *arXiv preprint arXiv:2105.14553*.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Xingyi Cheng, Weidi Xu, Kunlong Chen, Shaohua Jiang, Feng Wang, Taifeng Wang, Wei Chu, and Yuan Qi. 2020. Spellgen: Incorporating phonological and visual similarities into language models for chinese spelling check. *arXiv preprint arXiv:2004.14166*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Siddhant Garg and Goutham Ramakrishnan. 2020. Bae: Bert-based adversarial examples for text classification. *arXiv preprint arXiv:2004.01970*.
- Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864.

- Jonathan Herzig, Paweł Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Martin Eisenschlos. 2020. Tapas: Weakly supervised table parsing via pre-training. *arXiv preprint arXiv:2004.02349*.
- Zhuoren Jiang, Zhe Gao, Guoxiu He, Yangyang Kang, Changlong Sun, Qiong Zhang, Luo Si, and Xiaozhong Liu. 2019. Detect camouflaged spam content via stonescaping: Graph and text joint embedding for chinese character variation representation. *arXiv preprint arXiv:1908.11561*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Jinfeng Li, Tianyu Du, Xiangyu Liu, Rong Zhang, Hui Xue, and Shouling Ji. 2021. Enhancing model robustness by incorporating adversarial knowledge into semantic representation. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7708–7712. IEEE.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. Textbugger: Generating adversarial text against real-world applications. In *26th Annual Network and Distributed System Security Symposium*.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. Bert-attack: Adversarial attack against bert using bert. *arXiv preprint arXiv:2004.09984*.
- 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*.
- Nathan Ng, Kyunghyun Cho, and Marzyeh Ghassemi. 2020. Ssmba: Self-supervised manifold based data augmentation for improving out-of-domain robustness. *arXiv preprint arXiv:2009.10195*.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 1085–1097.
- Chenglei Si, Zhengyan Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2021. Better robustness by more coverage: Adversarial and mixup data augmentation for robust fine-tuning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1569–1576.
- Hui Su, Weiwei Shi, Xiaoyu Shen, Zhou Xiao, Tuo Ji, Jiarui Fang, and Jie Zhou. 2022. Robbert: Robust chinese bert with multimodal contrastive pretraining. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 921–931.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In *China National Conference on Chinese Computational Linguistics*, pages 194–206. Springer.
- Maosong Sun, Xinxiong Chen, Kaixu Zhang, Zhipeng Guo, and Zhiyuan Liu. 2016. Thulac: An efficient lexical analyzer for chinese. Technical report, Technical Report. Technical Report.
- Zijun Sun, Xiaoya Li, Xiaofei Sun, Yuxian Meng, Xiang Ao, Qing He, Fei Wu, and Jiwei Li. 2021. Chinesebert: Chinese pretraining enhanced by glyph and pinyin information. *arXiv preprint arXiv:2106.16038*.
- 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.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, et al. 2020. Clue: A chinese language understanding evaluation benchmark. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772.
- Shaohua Zhang, Haoran Huang, Jicong Liu, and Hang Li. 2020a. Spelling error correction with soft-masked bert. *arXiv preprint arXiv:2005.07421*.
- Zihan Zhang, Mingxuan Liu, Chao Zhang, Yiming Zhang, Zhou Li, Qi Li, Haixin Duan, and Donghong Sun. 2020b. Argot: Generating adversarial readable chinese texts. In *IJCAI*, pages 2533–2539.
- Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. 2020. Incorporating bert into neural machine translation. *arXiv preprint arXiv:2002.06823*.

A Appendix

A.1 Details of Datasets

The datasets we used to evaluate RoChBert are shown in Table 5.

- **ChnSentiCorp**: a binary sentiment classification dataset containing 9,600/1,200/1,200 texts in train/dev/test datasets respectively. The content of these texts is the reviews of product and hotel. And their labels are positive or negative.
- **DMSC**: a Chinese movie review datasets released by Kaggle. We follow Li et al. (2021) to sample the original datasets. In our datasets, there are 16,000/2,000/2,000 texts with positive or negative labels in train/dev/test datasets.
- **THUCNews**: a news multi-classification dataset with a total of fourteen categories. Like (Zhang et al., 2020b), we select five of them for experiments. In our dataset, there are 12,000/1,500/1,500 texts in train/dev/test datasets.
- **OCNLI**: a Chinese natural language inference dataset from CLUE Benchmark (Xu et al., 2020). Each sample contains two sentences, premise and hypothesis. We count the average length of the texts separately for these two parts. And the label (Entailment, Contradiction, or Neutral) represents the relation between premise and hypothesis. There are 9,985 and 2,950 samples in train and dev datasets. As the samples in test datasets don't have labels, we use the dev datasets to conduct robustness evaluation.

A.2 Details of Attack algorithms

we utilize three widely used attacks, i.e., PWWS (Ren et al., 2019), TextBugger (Li et al., 2019) and random attack in the black-box setting to evaluate the robustness of RoChBert. In order to adapt the attack algorithms to Chinese text, we have made some modifications to them.

- **PWWS**: It first generates a candidate set for each word, including the synonym of the word and a similar name entity, and then replaces the words to find the optimal candidate word according to the confidence drop. The next step is to determine which word in the text

should be replaced preferentially. It still uses confidence drop to evaluate the priority. In order to make PWWS generate Chinese adversarial texts, we changed the candidate words from the original synonyms and name entities to phrases composed of Chinese phonetic and glyph-similar words.

- **TextBugger**: It first segments the text into individual sentences, and it ranks the importance of the sentences by querying the target model with single sentence in turn. Then the words in the sentence are removed one by one, and the drops of confidence are used to determine the importance of the words. In order to make TextBugger generate Chinese adversarial texts, we only keep one of its bug generation methods, *Word Substitution*, and the candidate words are phrases composed of Chinese phonetic and glyph-similar words.
- **Random Attack**: It is the baseline method compared with other two attack algorithms. It randomly selects words in the text and use the candidate words to substitute the original one.

In general, we use Chinese characters with similar glyph or pronunciation to form candidate words (there are at most 40 candidate words). And after locating the important parts through the corresponding attack algorithm, we select suitable candidate words to replace the original words. Specifically, we use the same perturbation strategies in both data augmentation and robustness evaluations.

A.3 Details of ϵ_{max}

The details of ϵ_{max} we used in experiment are shown in Table 6.

A.4 Results of Attacking Hypothesis in OCNLI

A.4.1 Model Performance

The results are presented in Table 7. We can see that ChineseBERT outperforms other models on OCNLI. The data augmentation does affects the accuracy in some degree, especially the models using hypothesis examples generated by PWWS and TextBugger algorithms. But in some cases, models defended by RoChBert can help models improve accuracy on benign texts, which is consistent with the results of other tasks and datasets. In addition, SC reduces the accuracy of models due to its own errors.

Dataset	Task	Label	Train	Dev	Test	Length
ChnSentiCorp	Sentiment Analysis	2	9,600	1,200	1,200	108
DMSC	Sentiment Analysis	2	16,000	2,000	2,000	64
THUCNews	Text Classification	5	12,000	1,500	1,500	189
OCNLI	Natural Language Inference	3	9,985	2,950	-	24/11

Table 5: Overview of datasets.

Dataset	Attack Algorithm		
	PWWS	TextBugger	Random
ChnSentiCorp	0.45	0.45	0.45
DMSC	0.45	0.45	0.45
THUCNews	0.3	0.3	0.3
OCNLI	0.3	0.3	0.1

Table 6: The ϵ_{max} we used in experiment.

Model	OCNLI
ChineseBERT	73.20
BERT _{base}	71.57
+SC	70.57
+RoChBert (PWWS)	68.34
+RoChBert (TextBugger)	68.48
+RoChBert (Random)	70.52
BERT _{wwm}	70.33
+SC	69.08
+RoChBert (PWWS)	69.41
+RoChBert (TextBugger)	67.87
+RoChBert (Random)	70.30
BERT _{wwm/ext}	71.16
+SC	70.68
+RoChBert (PWWS)	69.29
+RoChBert (TextBugger)	71.27
+RoChBert (Random)	71.56
RoBERTa _{wwm/ext}	71.29
+SC	70.88
+RoChBert (PWWS)	70.07
+RoChBert (TextBugger)	70.30
+RoChBert (Random)	71.61

Table 7: Model performances on benign texts.

A.4.2 Robustness Against Attack

The results of attacking hypothesis in OCNLI are shown in Table 8. It proves that RoChBert can enhance the robustness of BERT-based models. We can see that RoChBert is not as effective in attacking hypothesis as it is on other tasks and datasets. This may be because the length of the hypothesis sentences is very short, and modifying a small number of characters will have a great impact on the original embedding, so it is easy to attack successfully.

A.5 Representation Visualization

We add the representation visualization results on other models and their corresponding vari-

ant models protected by RoChBert in Figure 5. We still use models trained on DMSC dataset against TextBugger attack and choose the same 500 samples. Obviously, most of the adversarial texts are effectively separated after being defended by RoChBert, which once again proves that RoChBert significantly strengthen models' robustness and can be easily exploited on other pre-trained language models.




Figure 5: Representation with different models. (a) is adversarial texts on BERT_{wwm}. (b) is adversarial texts on BERT_{wwm}+RoChBert. (c) is adversarial texts on BERT_{wwm/ext}. (d) is adversarial texts on BERT_{wwm/ext}+RoChBert. (e) is adversarial texts on RoBERTa_{wwm/ext}. (f) is adversarial texts on RoBERTa_{wwm/ext}+RoChBert.

A.6 Impact of Modification Rate

In Figure 6, 7, 8 and 9 we list all the experiment results about the impact of modification rate with models trained on ChnSentiCorp, DMSC, THUCNews and OCNLI against PWWS, TextBugger and Random attacks.

We can see that as the allowed modification rate increases, ASR is gradually growing. And the ASR

Model	PWWS			TextBugger			Random		
	UASR	LASR	Modi.	UASR	LASR	Modi.	UASR	LASR	Modi.
OCNLI									
ChineseBERT	99.05	54.86	21.23	99.86	48.65	22.72	79.19	10.27	48.84
BERT _{base}	94.77	52.34	20.64	98.76	47.52	22.71	70.80	11.98	41.21
+SC	95.24	53.15	20.10	98.32	48.53	22.25	71.61	10.07	43.10
+RoChBERT	97.83	47.39	24.07	97.95	36.40	26.56	70.00	10.97	46.05
BERT _{wwm}	96.34	51.41	21.26	99.44	45.92	23.13	75.35	7.75	48.95
+SC	97.30	55.62	20.33	99.72	49.36	22.47	76.81	11.81	46.00
+RoChBERT	97.99	49.43	22.76	99.86	40.11	25.64	73.45	10.25	47.99
BERT _{wwm/ext}	97.36	51.11	22.09	99.86	45.14	23.44	78.47	8.75	50.52
+SC	97.07	54.95	21.06	99.58	49.09	21.73	79.78	8.65	48.14
+RoChBERT	96.25	49.64	21.36	99.58	40.36	25.84	75.21	9.10	52.19
RoBERTa _{wwm/ext}	96.57	51.92	21.84	99.18	45.19	23.83	79.26	8.93	51.04
+SC	96.83	52.97	21.04	99.17	47.17	22.34	81.24	9.93	49.74
+RoChBERT	96.03	52.19	22.30	99.17	40.33	25.77	71.91	13.16	42.64

Table 8: Model performance against different attacks.




Figure 6: Relation of modification rate and ASR on ChnSentiCorp.

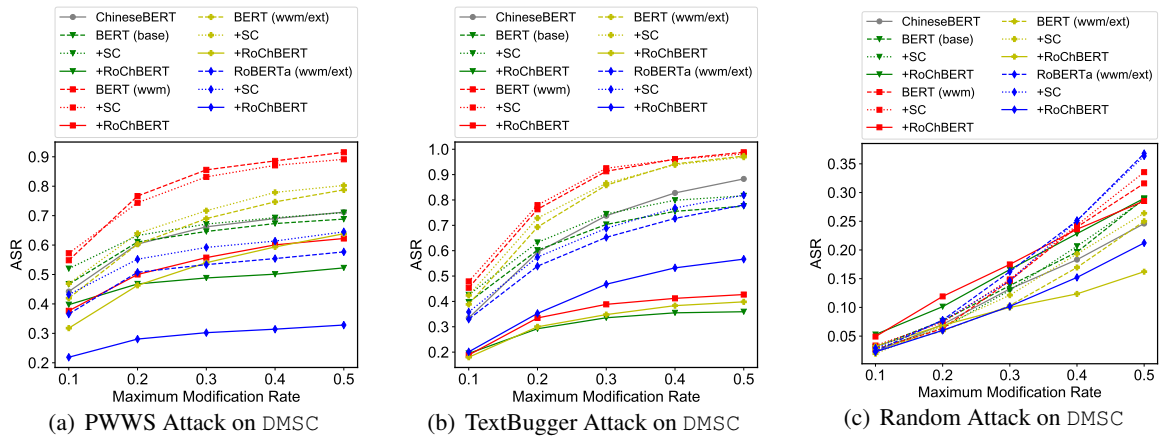


Figure 7: Relation of modification rate and ASR on DMSC.

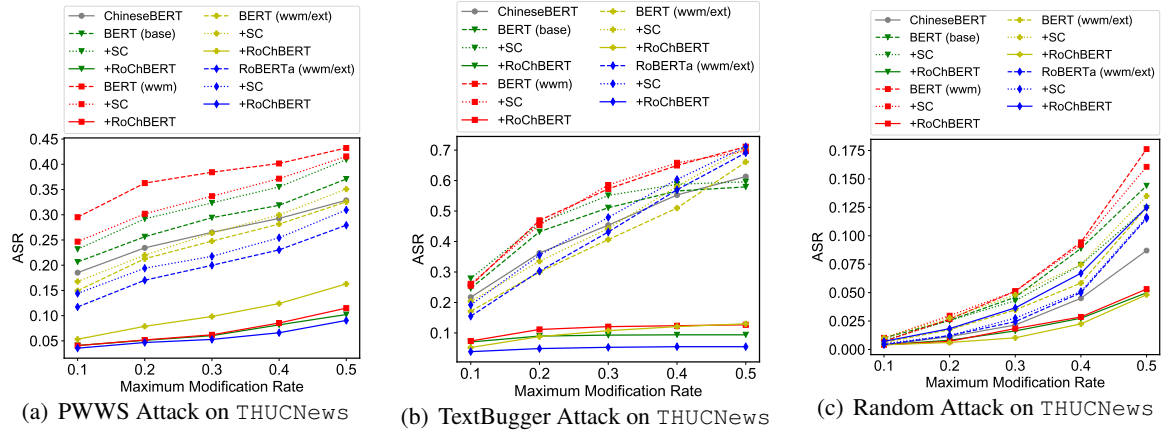


Figure 8: Relation of modification rate and ASR on THUCNews.

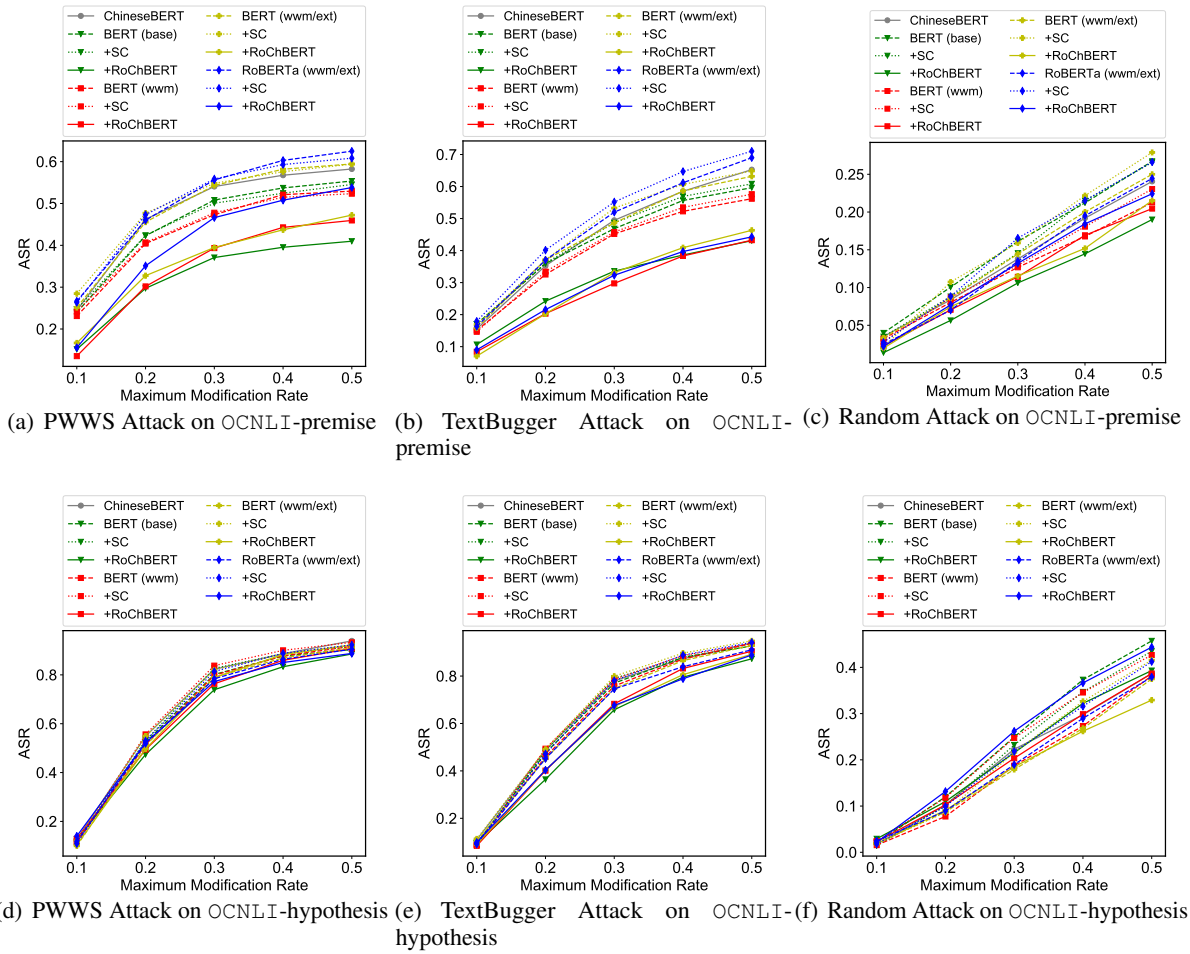


Figure 9: Relation of modification rate and ASR on OCNLI.

Models	Attack	UASR	LASR	MR	Used Samples	Percentage
BERT _{base}	PWWS	83.62	67.66	12.96	0	0%
	TextBugger	97.45	69.26	16.12	0	0%
	Random	52.77	8.19	42.85	0	0%
+Adversarial Training	PWWS	70.67	50.27	19.02	9600	100%
	TextBugger	96.36	64.09	16.69	9600	100%
	Random	36.18	16.24	30.08	9600	100%
+RoChBert	PWWS	65.18	31.63	29.49	2438	25.40%
	TextBugger	64.45	34.92	20.35	1653	17.22%
	Random	39.49	10.98	37.48	1588	16.54%

Table 9: Efficiency of traditional adversarial training and RoChBert.

of the models with RoChBert have the least rise, which proves that even in different settings, the models defended by RoChBert still significantly outperform the baselines.

A.7 Efficiency

To assess the efficiency of RoChBert, we take the BERT_{base} fine-tuned on ChnSentiCorp dataset as an example and compare RoChBert with the traditional adversarial training. For adversarial training method, we use attack algorithms on the target model with all the training dataset and collect the successful ones as the augmentation dataset. The results are presented in Table 9. It shows that RoChBert only uses 25.40%, 17.22% and 16.54% of the training dataset when conducting PWWS, TextBugger and Random attacks respectively. Even the adversarial training has used all the training dataset, RoChBert still achieves more robust models. It proves that the data augmentation method used in RoChBert is efficient and effective. It also indicates the key impact of multimodal fusion.