

# TF-Attack: Transferable and Fast Adversarial Attacks on Large Language Models

Anonymous ACL submission

## Abstract

With the great advancements in large language models (LLMs), *adversarial attacks* against LLMs have recently attracted increasing attention. We found that pre-existing adversarial attack methodologies exhibit limited transferability and are notably inefficient, particularly when applied to LLMs. In this paper, we analyze the core mechanisms of previous predominant adversarial attack methods, revealing that 1) the distributions of importance score differ markedly among victim models, restricting the transferability; 2) the sequential attack processes induces substantial time overheads. Based on the above two insights, we introduce a new scheme, named TF-ATTACK, for Transferable and Fast adversarial attacks on LLMs. TF-ATTACK employs an external LLM as a third-party overseer rather than the victim model to identify critical units within sentences. Moreover, TF-ATTACK introduces the concept of *Importance Level*, which allows for parallel substitutions of attacks. We conduct extensive experiments on 6 widely adopted benchmarks, evaluating the proposed method through both automatic and human metrics. Results show that our method consistently surpasses previous methods in transferability and delivers significant speed improvements, up to 20× faster than earlier attack strategies.

## 1 Introduction

Recently, large language models (LLMs) such as ChatGPT and LLaMA (Ouyang et al., 2022; Touvron et al., 2023) have demonstrated considerable promise across a range of downstream tasks (Kasneci et al., 2023; Thirunavukarasu et al., 2023; Liu et al., 2023). Subsequently, there has been increasing attention on the task of adversarial attack (Xu et al., 2023; Yao et al., 2023), which aims to generate adversarial examples that confuse or mislead LLMs. This task is crucial for advancing

	CNN	LSTM	BERT	LLaMA	Baichuan
CNN	94.7	22.9	19.8	21.7	21.2
LSTM	17.3	94.1	19.4	22.0	22.1
BERT	12.6	14.6	91.0	16.6	18.3
LLaMA	12.1	9.5	8.7	86.1	16.0
Baichuan	21.8	24.0	19.6	27.2	89.3
ChatGPT	11.2	14.4	12.6	16.2	14.8
Average	14.6	17.0	18.0	20.7	18.1

Table 1: Transferability evaluation of BERT-Attack samples on IMDB dataset. Row  $i$  and column  $j$  is the Attack Success Rate of samples generated from model  $j$  and evaluated on model  $i$ . The Average result is from non-diagonal elements of each column.

reliable and robust LLMs in the AI community, emphasizing the paramount importance of security in AI systems (Marcus, 2020; Thiebes et al., 2021).

Existing predominant adversarial attack approaches on LLMs typically adhere to a two-step process: initially, they rank token importance based on the victim model, and subsequently, they replace these tokens sequentially following specific rules (Cer et al., 2018; Oliva et al., 2011; Jin et al., 2020). Despite notable successes, recent studies highlight that current methods face two substantial limitations: 1) poor transferability of the generated adversarial samples. As depicted in Table 1, while adversarial samples generated by models can drastically reduce their own classification accuracy, they scarcely affect other models; 2) significant time overhead, particularly with larger models (Spector and Re, 2023). For instance, the time required to conduct an attack on LLaMA is 30× slower than standard inference.

To tackle the above two issues, we begin by analyzing the causes of the poor transferability of existing adversarial attack methods and the slow speed. Specifically, we first study the impact of the importance score, which is the core mechanism of previous methods. Comparative analysis reveals distinct importance score distributions across various victim models. This discrepancy largely

explains why the portability of adversarial samples generated by existing methods is poor, since perturbations generated according to a specific pattern of one model do not generalize well to other models with different importance assignments. In addition, analysis of time consumption in implementing a representative attack on LLMs shows that over 80% of processing time is spent on sequential word-by-word operations.

Drawing on insights from the above observations, we propose a new scheme, named TF-ATTACK, for transferable and fast adversarial attacks over LLMs. TF-ATTACK follows the overall framework of BERT-Attack which generates adversarial samples by synonym replacement. Different from BERT-Attack, TF-ATTACK employs an external third-party overseer such as ChatGPT to identify important units with the input sentence, thus eliminating the dependency on the victim model. Moreover, TF-ATTACK introduces the concept of *Important Level*, which divides the input units into different groups based on their semantic importance. Specifically, we utilize the powerful abstract semantic understanding and efficient automatic extraction capabilities of ChatGPT to form an important level priority queue through human-crafted *Ensemble Prompts*. Based on this, TF-ATTACK can perform parallel replacement of entire words within the same priority queue, as opposed to the traditional approach of replacing them sequentially one by one, as shown in Figure 3. This approach markedly reduces the time of the attacking process, thereby resulting in a significant speed improvement. Furthermore, we employ two tricks, named *Multi-Disturb* and *Dynamic-Disturb*, to enhance the attack effectiveness and transferability of generated adversarial samples. The former involves three levels of disturbances within the same sentence, while the latter dynamically adjusts the proportions and thresholds of the three types of disturbances based on the sentence length of the input. Both of them significantly boost attack effectiveness and transferability, and are adaptable to other attack methods.

We conduct experiments on six widely adopted benchmarks and compare the proposed method with several state-of-the-art methods. To verify the effectiveness of our method, we consider both automatic and human evaluation. Automatic evaluation on 6 widely benchmarks shows that TF-ATTACK has surpassed the baseline method TextFooler,

BERT-Attack and SDM-Attack. In addition, human evaluation results demonstrate that TF-ATTACK maintains a comparable consistency and achieves a comparable level of language fluency that does not cause much confusion for humans. Moreover, we compare the transferability of TF-ATTACK with BERT-Attack, demonstrating that TF-ATTACK markedly diminishes the accuracy of other models and exhibits robust migration attack capabilities. Furthermore, the time cost of TF-ATTACK is significantly lower than BERT-Attack, more than  $10\times$  speedup on average stats. Lastly, the adversarial examples generated by TF-ATTACK minimally impact the performance of model after adversarial training, significantly strengthening its defense against adversarial attacks. Overall, the main contributions of this work can be summarized as follows:

- We investigate the underlying causes behind the slow speed and poor effectiveness of pre-existed adversarial attacks on LLMs.
- We introduce TF-ATTACK, a novel approach leveraging an external LLM to identify critical units and facilitates parallelized adversarial attacks.
- TF-ATTACK effectively enhances the transferability of generated adversarial samples and achieves a significant speedup compared to previous methods.

## 2 Related Work

**Text Adversarial Attack** For NLP tasks, the adversarial attacks occur at various text levels including the character, word, or sentence level. Character-level attacks involve altering text by changing letters, symbols, and numbers. Word-level attacks (Wei and Zou, 2019) involve modifying the vocabulary with synonyms, misspellings, or specific keywords. Sentence-level attacks (Coulombe, 2018; Xie et al., 2020) involve adding crafted sentences to disrupt the output of model. Current adversarial attacks in NLP (Alzantot et al., 2018; Jin et al., 2020) employ substitution to generate adversarial examples through diverse strategies, such as genetic algorithms (Zang et al., 2019; Guo et al., 2021), greedy search (Sato et al., 2018; Yoo and Qi, 2021), or gradient-based methods (Ebrahimi et al., 2018a; Cheng et al., 2018), are employed to identify substitution words form synonyms (Kuleshov et al., 2018; Jin

et al., 2020) or language models (Li et al., 2020b; Garg and Ramakrishnan, 2020; Li et al., 2020a). Recent studies have refined sampling methods, yet these approaches continue to be time-intensive, highlighting a persistent challenge in efficiency. Fang et al. (2023) apply reinforcement learning, showing promise on small models but facing challenges on LLMs due to lengthy iterations, limiting large-scale adversarial samples. We list more related work in Appendix A.

**Sample Transferability** Evaluating text adversarial attacks heavily depends on sample transferability, assessing the performance of attack samples across diverse environments and models to measure their broad applicability and consistency. In experiments (Qi et al., 2021), adversarial samples generated from the Victim Model are directly applied to other models, testing transferability. Strongly transferable attack samples can hit almost all models in a black-box manner, which traditional white-box attacks (Ebrahimi et al., 2018b) can not match. Evaluation datasets like adversarial tasks of Adv-Glue (Wang et al., 2021) showcase this transferability, aiding in robustness assessment. Relevant research (Liu et al., 2016) endeavors to enhance the capability of adversarial example transferability by attacking ensemble models. It has been demonstrated (Zheng et al., 2020) that adversarial examples with better transferability can more effectively enhance the robustness of models in adversarial training. However, in the field of textual adversarial attacks, there has been a lack of in-depth research dedicated to how to improve the transferability of adversarial examples.

### 3 Method

#### 3.1 Preliminaries: Adversarial Attack

The task of adversarial attack aims at generating perturbations on inputs that can mislead the output of model. These perturbations can be very small, and imperceptible to human senses.

As for NLP tasks, given a corpus of  $N$  input texts,  $\mathcal{X} = \{x_1, x_2, x_3, \dots, x_N\}$ , and an output space  $\mathcal{Y} = \{y_1, y_2, y_3, \dots, y_N\}$  containing  $K$  labels, the language model  $F(\cdot)$  learns a mapping  $f : x \rightarrow y$ , which learns to classify each input sample  $x \in X$  to the ground-truth label  $y_{\text{gold}} \in \mathcal{Y}$ :

$$F(x) = \arg \max_{y_i \in \mathcal{Y}} P(y_i|x). \quad (1)$$

The adversary of text  $x \in X$  can be formulated as  $x_{\text{adv}} = x + \epsilon$ , where  $\epsilon$  is a perturbation to the input  $x$ . The goal is to mislead the victim language model  $F(\cdot)$  within a certain constraint  $C(x_{\text{adv}})$ :

$$F(x_{\text{adv}}) = \arg \max_{y_i \in \mathcal{Y}} P(y_i|x_{\text{adv}}) \neq F(x), \quad (2)$$

and  $C(x_{\text{adv}}, x) \leq \lambda$

where  $\lambda$  is the coefficient, and  $C(x_{\text{adv}}, x)$  is usually calculated by the semantic or syntactic similarity (Cer et al., 2018; Oliva et al., 2011) between the input  $x$  and its corresponding adversary  $x_{\text{adv}}$ .

#### 3.2 Limitations of Importance Score

Existing predominant adversarial attack systems (Morris et al., 2020; Jin et al., 2020; Li et al., 2020b) on LLMs typically adhere to a two-step process following BERT-Attack (Li et al., 2020b). We thus take BERT-Attack as a representative method for analysis. The core idea of BERT-Attack is to perform substitution according to **Importance Score**. BERT-Attack calculates the importance score by individually masking each word in the input text, performing a single inference for each, and using changes in the confidence score and label of Victim Models to determine the impact of each word. In BERT-Attack, importance score determine the subsequent attack sequence, which is crucial for the success of subsequent attacks and the times of attacks.

To explore the limited transferability of methods relying on importance scores, we analyze the importance score distribution among various models. As illustrated in Figure 1, there are substantial differences in the importance score derived from the same sentence when calculated using BERT and LLaMA. The former has a sharper distribution, while the latter essentially lacks substantial numerical differences. This phenomenon is consistent across multiple sentences, which are detailed in Appendix F. Given that importance score are crucial in the attack process of the aforementioned methods, variations in these scores can lead to entirely different adversarial samples. This variation explains the poor transferability of the generated attack samples, as demonstrated in Table 1.

To investigate the causes of significant time overhead, we analyze the time consumption of various components involved when implementing a representative attacking method on both a small




Figure 1: Importance score distribution of the same sentence given by BERT-Attack on BERT and LLaMA.




Figure 2: Time cost of each module from BERT-Attack on SA-LLaMA and BERT.

and a large model. As depicted in Figure 2, when applying BERT-Attack to attack small models like BERT, the average time spent per entry is very short, and the time consumption of various components is similar. However, the time cost per inference on LLaMA far exceeds that of BERT, disrupting this balance. It is evident that in LLMs, over 80% of the time spent per entry in the attack is consumed by the *Get import-scores* and *substitute & Attack* component. In addition, it is worth noting that this phenomenon is even more pronounced in successfully attacked samples. The underlying reason is that BERT-Attack necessitates performing the attack sequentially, according to the calculated importance score. This drawback becomes more pronounced when applied to large models, as both components require performing model inference, significantly increasing the time cost.

### 3.3 TF-ATTACK: Transferable and Fast Adversarial Attacks on LLMs

To address the above limitations caused by the importance score mechanism, we introduce TF-ATTACK as a solution for transferable and fast adversarial attacks. Firstly, as shown in Figure 3, TF-ATTACK employs an external model as a third-party overseer to identify important units. In this way, TF-ATTACK avoids excessive dependence on the victim model during the attack process.

Specifically, we design several instructions<sup>1</sup> for ChatGPT to partition all words in the original input according to their semantic importance. Another advantage of TF-ATTACK is its ability to utilize the rich semantic knowledge within ChatGPT, making the subsequent generation of attack sequences more universally semantic. In terms of speed, this approach does not require inference from the victim model, thus alleviating the problem of high inference costs for LLMs. TF-ATTACK only needs one inference to obtain a comprehensive Important Level priority queue, while the traditional approach requires inference times proportional to the length of the text.

Secondly, TF-ATTACK introduces the concept of **Importance Level** to facilitate parallel substitution. Specifically, TF-ATTACK takes the original sentence as input and outputs a priority queue with 5 levels, with a different number of words in each level. The concept of importance level assumes that words within the same level have no specific order, enabling parallel replacement of candidate words at the same level. This parallel replacement process markedly decreases both the search space and the number of required inferences, offering a substantial improvement over the previous approach that relied on greedy, sequential word replacements.

Additionally, we employ a *reverse pyramid search space* strategy for importance levels, optimizing the search space to reduce inefficient search expenditures. Words prioritized at higher levels are presumed to significantly influence sentence sentiment. Consequently, a larger search space is utilized to identify semantically similar words, with the aim of replacing them with synonyms that precipitate a notable decline in the performance of victim model. For words at lower priority levels, a smaller search space suffices, as alterations to these words minimally impact sentence sentiment. Excessive searching at these levels can lead to increased inference costs without substantially enhancing effectiveness. More details of experiments could be found in F.

### 3.4 Multi-Disturb & Dynamic-Disturb

These two strategies can be incorporated into traditional text adversarial attack methods as a post-processing step, significantly improving the transferability of adversarial samples. Specifically,

<sup>1</sup>Please refer Appendix F for detailed instructions.

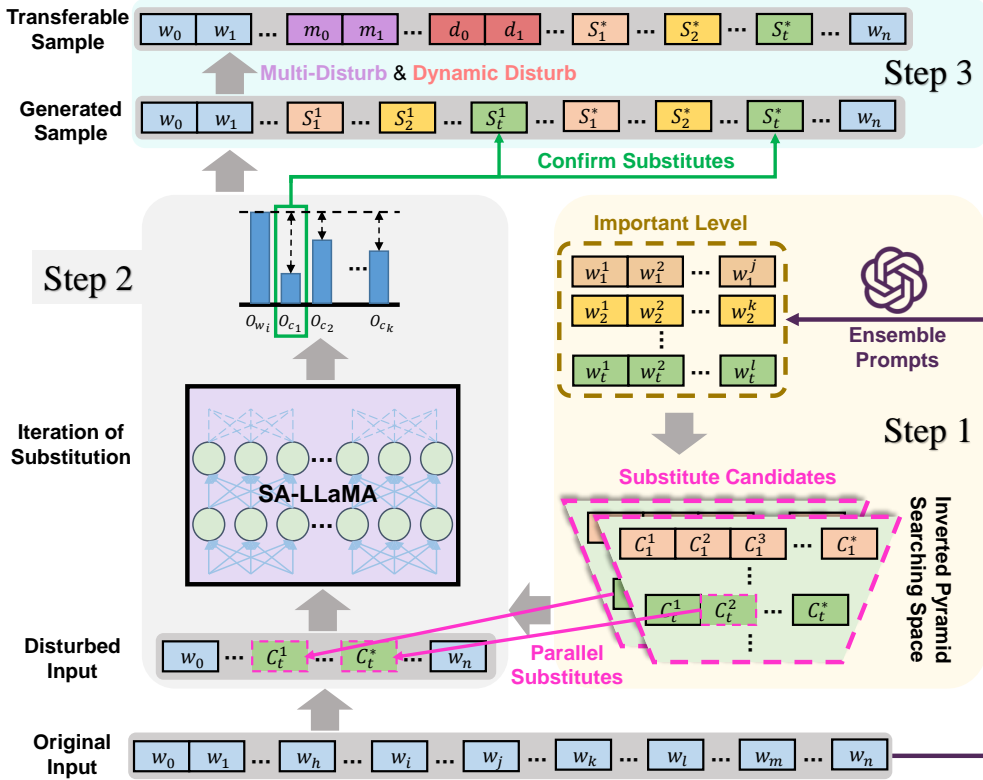


Figure 3: **Step 1:** Using ChatGPT to categorize words into 5 *Important Level* with varying word counts. The *Inverted Pyramid Searching Space* reflects the decreasing length of *Substitute Candidates* based on decreasing levels. **Step 2:** Selecting words from the same level and generates a *Disturbed Input* through *Parallel Substitutions*. Exploring possible *Disturbed Inputs* via SA-LLaMA, choose the result surpassing the threshold as *Generated Sample* from *Confirmed Substitutions*. *Substitution Iterations* will end when meet the finished condition. **Step 3:** Implementing *Multi-Disturb* and *Dynamic Disturb* produces *Transferable Samples*.

343 the Multi-Disturb strategy refers to introducing a  
 344 variety of disturbances within the same sentence.  
 345 Appendix H outlines 9 ways of disturbance,  
 346 including character-level, word-level, and sentence-  
 347 level disturbances, which can greatly enhance the  
 348 transferability of attack samples. Dynamic-Disturb  
 349 refers to using an FFN+Softmax network to assess  
 350 the length and structural distribution of the input  
 351 sentence, outputting the ratios of these three types  
 352 of disturbances.

353 In assessing attack sample effectiveness, tradi-  
 354 tional methods heavily depend on model output  
 355 confidence, likely leading to overfitting to the  
 356 model architecture. To rectify the problem, we  
 357 incorporate random disturbance to diminish model  
 358 confidence during replacement evaluation. Our  
 359 experiments confirm that this two tricks can  
 360 be adapted to almost all text adversarial attack  
 361 methods, significantly enhancing transferability  
 362 and increasing the ability of adversarial samples to  
 363 confuse models through adaptive post-processing.  
 364 The attacking algorithm is on Appendix B.

## 365 4 Experiments

### 366 4.1 Experimental Setups

367 **Tasks and Datasets** Following (Li et al., 2020b),  
 368 we evaluate the effectiveness of the proposed  
 369 TF-ATTACK on classification tasks upon diverse  
 370 datasets covering news topics (AG’s News; Zhang  
 371 et al., 2015), sentiment analysis at sentence  
 372 (MR; Pang and Lee, 2005) and document levels  
 373 (IMDB<sup>2</sup> and Yelp Polarity; Zhang et al., 2015). As  
 374 for textual entailment, we use a dataset of sentence  
 375 pairs (SNLI; Bowman et al., 2015) and a dataset  
 376 with multi-genre (MultiNLI; Williams et al., 2017).  
 377 Following Jin et al. (2020); Alzantot et al. (2018),  
 378 we attack 1k samples randomly selected from the  
 379 test set of each task. The statistics of datasets and  
 380 more details can be found in Appendix C.

381 **Baselines** We compare TF-ATTACK with recent  
 382 studies: 1) TextFooler (Jin et al., 2020), which finds  
 383 important words via probability-weighted word  
 384 saliency and then applies substitution with counter-

<sup>2</sup><https://datasets.imdbws.com/>

Dataset	Method	A-rate $\uparrow$	Mod $\downarrow$	Sim $\uparrow$	Dataset	Method	A-rate $\uparrow$	Mod $\downarrow$	Sim $\uparrow$
Yelp	TextFooler	78.9	<u>9.1</u>	0.73	IMDB	TextFooler	83.3	<u>8.1</u>	0.79
	BERT-Attack	80.5	11.5	0.69		BERT-Attack	84.2	9.6	0.78
	SDM-Attack	81.1	10.7	0.71		SDM-Attack	86.1	8.9	0.75
	TF-ATTACK (Zero-Shot)	81.3	10.3	0.71		TF-ATTACK (Zero-Shot)	86.1	8.7	<u>0.81</u>
	TF-ATTACK (Few-Shot)	83.7	<b>9.0</b>	<b>0.77</b>		TF-ATTACK (Few-Shot)	86.7	<b>7.4</b>	0.76
	+MD	<b>84.6</b>	12.3	0.71		+MD	<u>87.1</u>	10.7	0.77
	+MD +DD	<u>84.5</u>	11.9	0.72	+MD +DD	<b>87.7</b>	9.9	<b>0.81</b>	
AG’s News	TextFooler	73.2	<u>16.1</u>	0.54	MR	TextFooler	81.3	10.5	0.53
	BERT-Attack	76.6	17.3	<u>0.59</u>		BERT-Attack	82.8	10.2	0.51
	SDM-Attack	79.3	16.2	<b>0.61</b>		SDM-Attack	84.0	<u>9.9</u>	<u>0.55</u>
	TF-ATTACK (Zero-Shot)	77.1	18.3	0.52		TF-ATTACK (Zero-Shot)	<b>84.6</b>	<b>8.9</b>	<b>0.58</b>
	TF-ATTACK (Few-Shot)	81.9	<b>16.1</b>	0.58		TF-ATTACK (Few-Shot)	83.1	12.4	0.44
	+MD	<u>82.8</u>	19.4	0.53		+MD	83.0	13.2	0.40
	+MD +DD	<b>83.0</b>	19.1	0.55	+MD +DD	<u>84.0</u>	10.4	0.50	
SNLI	TextFooler	82.8	14.1	0.39	MNLI	TextFooler	77.1	11.5	0.51
	BERT-Attack	80.5	12.3	0.46		BERT-Attack	75.8	8.4	0.54
	SDM-Attack	83.1	14.6	0.42		SDM-Attack	79.3	10.2	0.55
	TF-ATTACK (Zero-Shot)	82.9	<b>10.2</b>	0.45		TF-ATTACK (Zero-Shot)	<u>79.3</u>	8.4	<u>0.55</u>
	TF-ATTACK (Few-Shot)	82.7	<u>10.4</u>	<b>0.47</b>		TF-ATTACK (Few-Shot)	78.2	7.9	<b>0.57</b>
	+MD	<b>83.6</b>	11.7	<u>0.46</u>		+MD	77.7	<b>8.3</b>	0.53
	+MD +DD	<u>83.5</u>	10.9	0.41	+MD +DD	<b>79.4</b>	8.5	0.52	

Table 2: Automatic evaluation results of attack success rate (A-rate), modification rate (Mod), and semantic similarity (Sim) on SA-LLaMA.  $\uparrow$  represents the higher the better and  $\downarrow$  means the opposite. The best results are **bolded**, and the second-best ones are underlined.

fitted word embeddings. 2) BERT-Attack (Li et al., 2020b), which uses a mask-predict approach to generate adversaries. 3) SDM-Attack (Fang et al., 2023), which employs reinforcement learning to determine the attack sequence. We use the official codes BERT-Attack and TextAttack tools (Morris et al., 2020) to perform attacks in our experiments. The TF-ATTACK (zero-shot) denotes that no demonstration examples were provided when generating the *Important Level* with ChatGPT. Conversely, the TF-ATTACK (few-shot) uses five demonstrations as context information. To ensure a fair comparison, we follow Morris et al. (2020) to apply constraints for TF-ATTACK. More details are available in Appendix D.

**Implementation Details** Following established training protocols, we fine-tuned a LLaMA-2-7B model to develop specialized Task-LLaMA models tailored for specific downstream tasks. Among them, the Task-LLaMA fine-tuned on the IMDB training set achieved an accuracy of 96.95% on the test set, surpassing XLNET with additional data, which achieved 96.21%. The model achieved 93.63% on another sentiment classification dataset, SST-2, indicating that Task-LLaMA is not overfitting to the training data but a strong baseline for experiments.

**Automatic Evaluation Metrics** Following prior work (Jin et al., 2020; Morris et al., 2020), we assess the results with the following metrics: 1) attack success rate (A-rate): post-attack model

performance decline; 2) Modification rate (Mod): percentage of altered words compared to the original; 3) Semantic similarity (Sim): cosine similarity between original and adversary texts via universal sentence encoder (USE; Cer et al., 2018); and 4) Transferability (Trans): the mean accuracy decreases across three models between adversarial and original samples.

**Manual Evaluation Metrics** We further manually validate the quality of the adversaries from three challenging properties. 1) Human prediction consistency (Con): how often human judgment aligns with the true label; 2) Language fluency (Flu): measured on a scale of 1 to 5 for sentence coherence (Gagnon-Marchand et al., 2019); and 3) Semantic similarity (Sim<sub>hum</sub>): gauging consistency between input-adversary pairs, with 1 indicating *agreement*, 0.5 *ambiguity*, and 0 *inconsistency*.

## 4.2 Results

**Overall Performance** Table 2 shows the performance of different systems on four benchmarks. As shown in Table 2, TF-ATTACK consistently achieves the highest attack success rate to attack LLaMA and has little negative impact on Mod and Sim. Additionally, TF-ATTACK mostly obtains the best performance of modification and similarity metrics, except for AG’s News, where TF-ATTACK achieves the second-best. In general, our method can simultaneously satisfy the high attack success rate with a lower modification rate and higher

	CNN	LSTM	BERT	LLaMA	Baichuan
<b>CNN</b>	-2.2	+33.8	+58.5	+37.4	+24.0
<b>LSTM</b>	+22.4	-3.8	+31.2	+45.8	+29.4
<b>BERT</b>	+21.4	+26.0	+2.2	+25.4	+25.6
<b>LLaMA</b>	+24.0	+26.8	+41.8	+12.4	+24.4
<b>Baichuan</b>	+22.2	+20.6	+22.4	+21.6	+4.4
<b>ChatGPT</b>	+14.8	+18.6	+22.4	+22.4	+18.8
<b>Average</b>	+21.2	+23.6	+35.2	+30.0	+27.6

Table 3: Transferability evaluation of TF-ATTACK Samples on IMDB dataset. Each element is calculated from the difference in Attack Success Rate between TF-ATTACK and BERT-Attack.




Figure 4: The time cost according to varying sentence lengths in the IMDB dataset. The left is on LLaMA while the right is on BERT.

similarity. We additionally observe that TF-ATTACK achieves a better attack effect on the binary classification task. Empirically, when there exist more than two categories, the impact of each replacement word may be biased towards a different class, leading to an increase in the perturbation rate. We also provide several cases to better understand our method, please refer Appendix G for more details.

**Transferability** We evaluate the transferability of TF-ATTACK samples to detect whether the samples generated from TF-ATTACK can effectively attack other models. We conduct experiments on the IMDB datasets and use BERT-Attack as a baseline. Table 3 shows the improvement of the attack success rate of TF-ATTACK over BERT-Attack. It can be observed that TF-ATTACK achieves obvious improvements over BERT-Attack when applied to other models. In particular, the new samples generated by TF-ATTACK can lower the accuracy by over 10 percent on binary classification tasks, essentially confusing the victim model. Even a powerful baseline like ChatGPT would drop to only 68.6% accuracy. It is important to highlight that these samples do not necessitate attacks tailored to specific victim models.

**Efficiency** We probe the efficiency according to varying sentence lengths in the IMDB dataset.

Dataset		Con $\uparrow$	Flu $\uparrow$	Sim $_{\text{hum}}\uparrow$
<b>IMDB</b>	Original	0.93	4.5	
	TF-ATTACK	0.87	4.1	0.93
<b>MR</b>	Original	0.88	4.0	
	TF-ATTACK	0.79	3.8	0.82

Table 4: Human evaluation results comparing the original input and generated adversary by TF-ATTACK of human prediction consistency (Con), language fluency (Flu), and semantic similarity (Sim $_{\text{hum}}$ ).

As shown in Figure 4, the time cost of TF-ATTACK is surprisingly mostly better than BERT-ATTACK, which mainly targets obtaining cheaper computation costs with lower attack success rates in Table 2. Furthermore, with the increase of sentence lengths, TF-ATTACK maintains a stable time cost, while the time cost of BERT-Attack is exploding. The reason is that TF-ATTACK has the advantage of much faster parallel substitution, hence as the sentence grows, the increase in time cost will be much smaller. These phenomena verify the efficiency advantage of TF-ATTACK, especially in dealing with long texts.

**Manual evaluation** We first randomly select 100 samples from successful adversaries in IMDB and MR datasets and then ask ten crowd-workers to evaluate the quality of the original inputs and our generated adversaries. The results are shown in Table 4. For human prediction consistency, humans can accurately judge 93% of the original inputs on the IMDB dataset while maintaining an 87% accuracy rate with our generated adversarial examples. This suggests that TF-ATTACK can effectively mislead LLMs without altering human judgment. Regarding language fluency, the scores of our adversarial examples are comparable to the original inputs, with a minor score difference of no more than 0.3 across both datasets. Moreover, the semantic similarity scores between the original inputs and our generated adversarial examples stand at 0.93 for IMDB and 0.82 for MR, respectively. Overall, TF-ATTACK successfully preserves these three essential attributes. More details about the human evaluation process and additional results are provided in Appendix F.

## 5 Analysis

**Generalization to More Victim Models** Table 6 shows that TF-ATTACK not only has better attack effects against WordCNN and WordLSTM, but also misleads BERT and Baichuan, which are

Dataset		Con $\uparrow$	Flu $\uparrow$	Sim <sub>hum</sub> $\uparrow$
IMDB	Original	0.96	4.6	1.00
	TextFooler	0.86	4.1	0.85
	BERT-Attack	0.82	4.3	0.91
	TF-ATTACK	0.91	4.6	0.91
AG’s News	Original	0.85	4.6	1.00
	TextFooler	0.77	3.9	0.84
	BERT-Attack	0.78	3.7	0.84
	TF-ATTACK	0.76	4.3	0.81

Table 5: Human evaluation results regarding human prediction consistency (Con), language fluency (Flu), and semantic similarity (Sim<sub>hum</sub>).

Victim models	A-rate $\uparrow$	Mod $\downarrow$	Sim $\uparrow$	Trans $\downarrow$
WordCNN	96.3	9.1	0.84	78.5
WordLSTM	92.8	9.3	0.85	75.1
BERT	90.6	9.9	0.81	70.2
LLaMA	91.8	13.1	0.74	68.6
Baichuan	92.5	11.8	0.75	71.4

Table 6: TF-ATTACK against other models.

more robust models. For example, on the IMDB datasets, the attack success rate is up to 92.5% against Baichuan with a modification rate of only about 11.8% and a high semantic similarity of 0.75. Furthermore, the model generated by the Victim model created a decrease in accuracy to 71.4% on various black-box models of different scales.

**Adversarial Training** We further investigate to improve the robustness of victim models via adversarial training using the generated adversarial samples. Specifically, we fine-tune the victim model with both original training datasets and our generated adversaries and evaluate it on the same test set. More details are in Appendix E. As shown in Table 7, compared to the results with the original training datasets, adversarial training with our generated adversaries can maintain close accuracy, while improving performance on attack success rates, modification rates, and semantic similarity. The victim models with adversarial training are more difficult to attack, which indicates that our generated adversaries have the potential to serve as supplementary corpora to enhance the robustness of victim models.

**Against Defense** Entropy threshold defense (Yao et al., 2023) has been used to defend against the attack on LLMs recently. It employs the entropy of the first token prediction to refuse responding. Figure 5 demonstrates the probability of top-10 tokens in the first generated word of LLaMA. It can be observed that the raw inputs usually generates the first token with low entropy (*i.e.*, the probability of argmax token is much higher, and

Dataset	Acc $\uparrow$	A-rate $\uparrow$	Mod $\downarrow$	Sim $\uparrow$
Yelp	97.4	81.3	8.5	0.73
	+Adv Train	95.9	65.7	12.3
IMDB	97.2	86.1	4.6	0.81
	+Adv Train	95.5	70.2	7.3
SST-2	97.1	89.7	14.3	0.85
	+Adv Train	92.2	68.6	16.8

Table 7: Results of adversarial training.

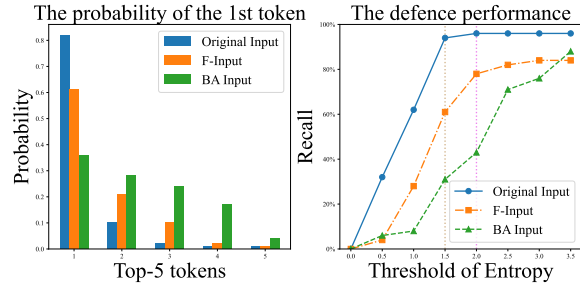


Figure 5: The probability of top-10 tokens in the first generated word in SA-LLaMA (a). The defense performance with various entropy thresholds (b)

the probability of other tokens is much lower). As shown in Figure 5, the adversarial samples from TF-ATTACK perform better than BERT-Attack with higher entropy. Attack samples generated through TF-ATTACK fare better against entropy-based filters compared to traditional text adversarial attack methods, indicating that the samples created by TF-ATTACK are harder to defend against.

## 6 Conclusion

In this paper, we examined the limitations of current adversarial attack methods, particularly their issues with transferability and efficiency when applied to Large Language Models (LLMs). To address these issues, we introduced TF-ATTACK, a new approach that uses an external overseer model to identify key sentence components and allows for parallel processing of adversarial substitutions. Our experiments on six benchmarks demonstrate that TF-ATTACK outperforms current methods, significantly improving both transferability and speed. Furthermore, the adversarial examples made by TF-ATTACK do not significantly affect the performance of model after it has been trained to resist attacks, thus strengthening its defenses. We believe that TF-ATTACK is a significant improvement in creating strong defenses against adversarial attacks on LLMs, with potential benefits for future research in this field.



575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622

## Limitations

Our experiments were solely conducted on six selected datasets for two NLP tasks, all of which were English corpora. Furthermore, due to resource constraints in LLM research, our experimental results primarily relied on fine-tuning LLaMA-2-7B and Baichuan-2-7B as the base, without exploring LLMs or other open-source base models. Consequently, we lack evaluations on other types of LLMs, such as ELECTRA (Clark et al., 2020), XLNET (Yang et al., 2019) and other LLMs. Hence, our work lacks validation in terms of generalization and transferability across multi-task, multi-model, and multi-lingual aspects.

## Ethics Statement

We declare that this article is in accordance with the ethical standards of *ACL Code of Ethics*. Any third-party tools used in this work are licensed by their authors. All crowd-workers participating in the experiments are paid according to the local hourly wages.

## References

Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. *arXiv preprint arXiv:1804.07998*.

Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.

Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.

Yong Cheng, Zhaopeng Tu, Fandong Meng, Junjie Zhai, and Yang Liu. 2018. Towards robust neural machine translation. *arXiv preprint arXiv:1805.06130*.

Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. *Electra: Pre-training text encoders as discriminators rather than generators*.

Claude Coulombe. 2018. Text data augmentation made simple by leveraging nlp cloud apis. *arXiv preprint arXiv:1812.04718*.

Javid Ebrahimi, Daniel Lowd, and Dejing Dou. 2018a. On adversarial examples for character-level neural machine translation. *arXiv preprint arXiv:1806.09030*.

Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018b. Hotflip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics. 623-628.

Xuanjie Fang, Sijie Cheng, Yang Liu, and Wei Wang. 2023. Modeling adversarial attack on pre-trained language models as sequential decision making. *arXiv preprint arXiv:2305.17440*. 629-631.

Jules Gagnon-Marchand, Hamed Sadeghi, Md Haidar, Mehdi Rezagholizadeh, et al. 2019. Salsa-text: self attentive latent space based adversarial text generation. In *Canadian Conference on Artificial Intelligence*, pages 119–131. Springer. 633-637.

Siddhant Garg and Goutham Ramakrishnan. 2020. Bae: Bert-based adversarial examples for text classification. *arXiv preprint arXiv:2004.01970*. 638-640.

Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. 2021. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*. 641-644.

Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025. 645-650.

Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günnemann, Eyke Hüllermeier, et al. 2023. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274. 651-657.

Volodymyr Kuleshov, Shantanu Thakoor, Tingfung Lau, and Stefano Ermon. 2018. Adversarial examples for natural language classification problems. 658-660.

Dianqi Li, Yizhe Zhang, Hao Peng, Liqun Chen, Chris Brockett, Ming-Ting Sun, and Bill Dolan. 2020a. Contextualized perturbation for textual adversarial attack. *arXiv preprint arXiv:2009.07502*. 661-664.

Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020b. Bert-attack: Adversarial attack against bert using bert. *arXiv preprint arXiv:2004.09984*. 665-668.

Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. 2016. Delving into transferable adversarial examples and black-box attacks. In *International Conference on Learning Representations*. 669-672.

Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, et al. 2023. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, page 100017. 673-678.

679	Gary Marcus. 2020. <a href="#">The next decade in AI: four steps towards robust artificial intelligence</a> . <i>CoRR</i> , abs/2002.06177.	models in medicine. <i>Nature medicine</i> , 29(8):1930–1940.	734 735
682	John X Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. <i>arXiv preprint arXiv:2005.05909</i> .	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	736 737 738 739 740 741
687	Nikola Mrkšić, Diarmuid O Séaghdha, Blaise Thomson, Milica Gašić, Lina Rojas-Barahona, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. Counter-fitting word vectors to linguistic constraints. <i>arXiv preprint arXiv:1603.00892</i> .	Boxin Wang, Chejian Xu, Shuohang Wang, Zhe Gan, Yu Cheng, Jianfeng Gao, Ahmed Hassan Awadallah, and Bo Li. 2021. Adversarial glue: A multi-task benchmark for robustness evaluation of language models. In <i>Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)</i> .	742 743 744 745 746 747 748
692	Jesús Oliva, José Ignacio Serrano, María Dolores Del Castillo, and Ángel Iglesias. 2011. Symss: A syntax-based measure for short-text semantic similarity. <i>Data &amp; Knowledge Engineering</i> , 70(4):390–405.	Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. <i>arXiv preprint arXiv:1901.11196</i> .	749 750 751 752
697	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. <i>Advances in Neural Information Processing Systems</i> , 35:27730–27744.	Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. <i>arXiv preprint arXiv:1704.05426</i> .	753 754 755 756
704	Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. <i>arXiv preprint cs/0506075</i> .	Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. 2020. Unsupervised data augmentation for consistency training. <i>Advances in Neural Information Processing Systems</i> , 33:6256–6268.	757 758 759 760
707	Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2021. Mind the style of text! adversarial and backdoor attacks based on text style transfer. In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 4569–4580.	Xilie Xu, Keyi Kong, Ning Liu, Lizhen Cui, Di Wang, Jingfeng Zhang, and Mohan Kankanhalli. 2023. An llm can fool itself: A prompt-based adversarial attack. <i>arXiv preprint arXiv:2310.13345</i> .	761 762 763 764
713	Motoki Sato, Jun Suzuki, Hiroyuki Shindo, and Yuji Matsumoto. 2018. Interpretable adversarial perturbation in input embedding space for text. <i>arXiv preprint arXiv:1805.02917</i> .	Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. <i>Advances in neural information processing systems</i> , 32.	765 766 767 768 769
717	Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In <i>Proceedings of the 2013 conference on empirical methods in natural language processing</i> , pages 1631–1642.	Jia-Yu Yao, Kun-Peng Ning, Zhen-Hui Liu, Mu-Nan Ning, and Li Yuan. 2023. Llm lies: Hallucinations are not bugs, but features as adversarial examples. <i>arXiv preprint arXiv:2310.01469</i> .	770 771 772 773
724	Benjamin Frederick Spector and Christopher Re. 2023. Accelerating llm inference with staged speculative decoding. In <i>Workshop on Efficient Systems for Foundation Models@ ICML2023</i> .	Jin Yong Yoo and Yanjun Qi. 2021. Towards improving adversarial training of nlp models. <i>arXiv preprint arXiv:2109.00544</i> .	774 775 776
728	Scott Thiebes, Sebastian Lins, and Ali Sunyaev. 2021. Trustworthy artificial intelligence. <i>Electronic Markets</i> , 31(2):447–464.	Yuan Zang, Fanchao Qi, Chenghao Yang, Zhiyuan Liu, Meng Zhang, Qun Liu, and Maosong Sun. 2019. Word-level textual adversarial attacking as combinatorial optimization. <i>arXiv preprint arXiv:1910.12196</i> .	777 778 779 780 781
731	Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language	Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. <i>Advances in neural information processing systems</i> , 28.	782 783 784 785

786 Haizhong Zheng, Ziqi Zhang, Juncheng Gu, Honglak  
787 Lee, and Atul Prakash. 2020. Efficient adversarial  
788 training with transferable adversarial examples.  
789 In *Proceedings of the IEEE/CVF Conference on*  
790 *Computer Vision and Pattern Recognition*, pages  
791 1181–1190.

## A Other Related Work

### A.1 Synchronization Work

Prompt-Attack(Xu et al., 2023) leverages the exceptional comprehension of LLMs and diverges from traditional adversarial methods. It employs a manual approach of constructing rule-based prompt inputs, requiring LLMs to output adversarial attack samples that can deceive itself and meet the modification rule conditions. This attack method achieved fully automatic and efficient generation of attack samples using the local model. However, the drawback is that the model may not perform the whole attacking process properly, resulting in mediocre attack effectiveness. Additionally, different prompts can significantly influence the quality of the model-generated attack samples. These generated attack samples, though somewhat transferable, fail to consider the model’s internal reasoning, resulting in excessively high modification rates.

Our work shares similarities in that we both leverage the understanding capabilities of LLMs. However, we solely employ ChatGPT’s language abstraction for Importance Level suggestions, maintaining the attack process within traditional text adversarial methods. Furthermore, our work is more focused on enhancing the transferability of attack samples and speeding up the attack, making them more widely applicable. This differs from the motivation of Prompt-Attack, which aims at the automated generation of samples that can deceive the model itself using LLMs.

## B Attack Algorithm

The attacking process is shown in Algorithm 1. Since  $a_t^f$  is chosen through a probability distribution, the method is encouraged to explore more possible paths of substitutions. The instant result  $r_t$  is obtained from victim model after once substitution actions.

Once the attack finish condition is meet, the method will terminate this current step and output the answer to check whether it is succeeded. The expected return of substitutions is defined as follows:

$$J(\theta) = \mathbb{E}[G(\tau)] \quad (3)$$

Thus the result is calculated by ouput of model and can be expressed as follows:

$$\nabla J(\theta) = \frac{1}{M} \sum_{m=1}^M \nabla \log \pi_{\theta}(\tau^{(m)}) G(\tau^{(m)}) \quad (4)$$

Dataset	Train	Test	Avg Len	Classes
Yelp	560k	38k	152	2
IMDB	25k	25k	215	2
AG’s News	120k	7.6k	73	4
MR	9k	1k	20	2
SST-2	7k	1k	17	2

Table 8: Overall statistics of datasets.

where  $[\tau^{(1)}, \tau^{(2)}, \dots, \tau^{(M)}]$  are  $M$  samples of trajectories. The discount factor  $\gamma$  enables both long-term and immediate effects to be taken into account and trajectories with shorter lengths are encouraged.

We use all the test sets of each dataset and he average convergence time is approximately between 2-16 hours, related to the length of the input. When attacking large batches of samples, the impact of training cost is negligible compared to the cumulative attack time cost. During training, We adopt random strategies and short-sighted strategies in the initial stage for early exploration and to obtain better seeds.

## C Datasets

We conduct experiments on the following datasets of Text Classification and detailed statistics are displayed in Table 8:

- **Yelp**(Zhang et al., 2015): A dataset for binary sentiment classification on reviews, constructed by considering stars 1 and 2 negative, and 3 and 4 positive.
- **IMDB**: A document-level movie review dataset for binary sentiment analysis.
- **MR**(Pang and Lee, 2005): A sentence-level binary classification dataset collected from Rotten Tomatoes movie reviews.
- **AG’s News**(Zhang et al., 2015): A collection of news articles. There are four topics in this dataset: World, Sports, Business, and Science/Technology.
- **SST-2** (Socher et al., 2013): The Stanford Sentiment Treebank task originates from reviews and is a binary sentiment classification dataset, where the task is to determine whether a given sentence conveys a positive or negative sentiment.

---

**Algorithm 1** TF-ATTACK Algorithm

---

- 1: Initialization: agent  $\pi_\theta$  with parameters  $\theta, \beta$ , sentences number  $M$
  - 2: **for**  $i \leftarrow 1$  to  $M$  **do**
  - 3:   using ensemble-prompt on GPT-4 to get Important Level  $L$
  - 4:   **while** not receive termination signal **do**
  - 5:     **for**  $t \leftarrow 1$  to  $K$  **do**
  - 6:       get words  $s_t1 \sim s_t\beta$  from level- $t$
  - 7:       compute  $\pi_\theta((a_t^f, a_t^s)|s_t) \sim \pi_\theta(a_t^f|s_t)$
  - 8:       Using SimCSE to get Possible Substitutes  $s_t1 \sim s_t\beta$
  - 9:       compute reward  $r_t$
  - 10:       update  $t \leftarrow t + 1$
  - 11:     **end for**
  - 12:   **end while**
  - 13:   initialize  $G(\tau) \leftarrow 0$
  - 14:   **for**  $j \leftarrow T$  to 1 **do**
  - 15:      $G(\tau) \leftarrow \gamma G(\tau) + r_j$
  - 16:     accumulate  $J_j(\theta)$
  - 17:   **end for**
  - 18:   update  $\theta \leftarrow \theta + \alpha \nabla J(\theta)$
  - 19: **end for**
- 

## D Implementation Constraint

In order to make the comparison fairer, we set the following constraints for TF-ATTACK as well as all baselines: (1) **Max modification rate**: To better maintain semantic consistency, we only keep adversarial samples with less than 40% of the words to be perturbed. (2) **Part-of-speech (POS)**: To generate grammatical and fluent sentences, we use NLTK tools<sup>3</sup> to filter candidates that have a different POS from the target word. This constraint is not employed on BERT-Attack. (3) **Stop words preservation**: the modification of stop words is disallowed and this constraint helps avoid grammatical errors. (4) **Word embedding distance**: For Textfooler, we only keep candidates with word embedding cosine similarity higher than 0.5 from synonyms dictionaries (Mrkšić et al., 2016). For *mask-fill* methods, following BERT-Attack, we filter out antonyms (Li et al., 2020b) via the same synonym dictionaries for sentiment classification tasks and textual entailment tasks.

## E Tuning with Adversaries

Table 11 displays adversarial training results of all datasets. Overall, after fine-tuned with both

<sup>3</sup><https://www.nltk.org/>

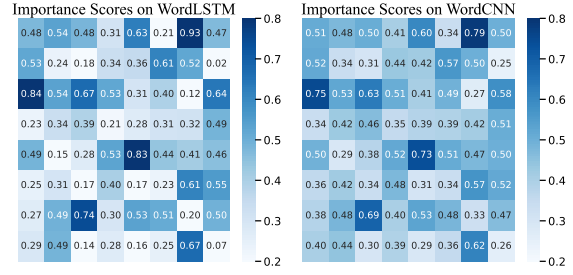


Figure 6: Importance score distribution of the same sentence given by BERT-Attack on WordCNN and WordLSTM.

original training datasets and adversaries, victim model is more difficult to attack. Compared to original results, accuracy of all datasets is barely affected, while attack success rate meets an obvious decline. Meanwhile, attacking model with adversarial training leads to higher modification rate, further demonstrating adversarial training may help improve robustness of victim models.

## F Supplementary Results

Figure 6 shows the different Importance Score distribution of the same sentence given by BERT-Attack on WordCNN and WordLSTM.

At the beginning of manual evaluation, we provided some data to allow crowdsourcing workers to unify the evaluation standards. We also remove the data with large differences when calculating the average value to ensure the reliability and accuracy of the evaluation results. More manual evaluation results are shown in Table 12.

In 14, we evaluate the effectiveness of attack order. Utilizing a random attack sequence leads to a reduction in the success rate of attacks and a significant increase in the modification rate, as well as severe disruption of sentence similarity. This implies that each attack path is random, and a substantial amount of inference overhead is wasted on futile attempts. Although we adhere to the threshold constraints of the traditional adversarial text attack domain, the text can still be successfully attacked. However, under conditions of high modification rates and low similarity, the text has been altered significantly from its original semantics, contravening the purpose of the task. Furthermore, a random attack sequence incurs a substantial additional cost in terms of attack speed, resulting in a nearly doubled time delay.

## Prompt

Rank each word in the input sentence into five levels based on its determining influence on the overall sentiment of the sentence.

Determine the impact of each word on the overall sentiment of the sentence and categorise it into 5 levels.

Rank the words from most to least influential in terms of their impact on the emotional tone of the sentence at 5 levels.

Please classify each word into five levels, based on their importance to the overall emotional classification of the utterance.

Assign each word to one of five levels of importance based on its contribution to the overall sentiment: Very High, High, Moderate, Low, Very Low.

Table 9: Different prompts used on ChatGPT to generate Important Level

Method	Text (MR; Negative)	Result	Mod↓	Sim↑	Flu↑
Original	It cuts to the knot of what it actually means to face your scares, and to ride the overwhelming metaphorical wave that life wherever it takes you.	-	-	-	5
TextFooler	It cuts to the <b>core</b> of what it actually means to face your <b>fears</b> , and to ride the <b>big</b> metaphorical wave that life wherever it takes you.	Success	15.8	0.86	5
BERT-Attack	It cuts to the <b>core</b> of what it <b>truly</b> means to <b>confront</b> your <b>fears</b> , and to <b>surf</b> the <b>overpowering</b> metaphorical wave that life wherever it <b>leads</b> you.	Failure	25.9	0.69	3
TF-ATTACK	It <b>cuts</b> to the <b>core</b> of what it <b>indeed</b> means to <b>confront</b> your <b>frights,qkXp</b> and to ride the <b>overwhelm</b> metaphorical wave that lie wherever takes you.	Success	14.6	0.59	4

Table 10: Adversaries generated by TF-ATTACK and baselines in MR dataset. The replaced words are highlighted in **blue**. *Failure* indicates the adversary fails to attack the victim model and *success* means the opposite.

Dataset	Acc↑	A-rate↑	Mod↓	Sim↑
<b>Yelp</b>	97.4	81.3	8.5	0.73
+Adv Train	95.9	65.7	12.3	0.67
<b>IMDB</b>	97.2	86.1	4.6	0.81
+Adv Train	95.5	70.2	7.3	0.78
<b>AG-NEWS</b>	95.3	77.1	15.3	0.83
+Adv Train	85.1	75.3	23.3	0.61
<b>MR</b>	95.9	83.2	11.1	0.53
+Adv Train	91.7	71.8	14.6	0.67
<b>SST-2</b>	97.1	89.7	14.3	0.85
+Adv Train	92.2	68.6	16.8	0.83

Table 11: Adversarial training results.

Dataset		Con↑	Flu↑	Sim <sub>hum</sub> ↑
<b>IMDB</b>	Original	0.96	4.6	1.00
	TextFooler	0.86	4.1	0.85
	BERT-Attack	0.82	4.3	0.91
	TF-ATTACK	0.91	4.6	0.91
<b>AG's News</b>	Original	0.85	4.6	1.00
	TextFooler	0.77	3.9	0.84
	BERT-Attack	0.78	3.7	0.84
	TF-ATTACK	0.76	4.3	0.81

Table 12: Manual evaluation results comparing the original input and generated adversary by attack method of human prediction consistency (Con), language fluency (Flu), and semantic similarity (Sim<sub>hum</sub>).

ASR↑	LLaMA-2b	LLaMA-2c	ChatGPT	Claude
Prompt1	/	/	86.54%	85.46%
Prompt2	/	12.28%	88.68%	78.64%
Prompt3	/	/	79.74%	71.07%
Prompt4	/	11.22%	84.36%	73.76%
Prompt5	13.67%	/	83.16%	77.81%

Table 13: The results of TF-ATTACK by different Prompts on different models.

In fact, in earlier experiments Table 13, we attempted to use local LLMs or other different architectures of LLMs as *selectors for attack order* but found that open-source LLMs base models such as 7B, 13B, or even 30B could not understand our instructional intent well in this specific scenario, no matter Base models or Chat models. If we were to introduce a specially designed task fine-tuning process to achieve this functionality, it would not

Method	A-rate	Mod↓	Sim↑	Time Cost ↓
BA-IS	84.2	9.9	0.78	11.1
TF-IS	86.3	10.4	0.77	10.4
BA-IL	<u>86.9</u>	<b>9.2</b>	<b>0.83</b>	<u>3.3</u>
TF-IL	<b>87.7</b>	<u>9.9</u>	<u>0.81</u>	<b>2.7</b>
BA-RD	75.3	18.4	0.47	27.6
TF-RD	77.8	17.4	0.51	25.9

Table 14: Automatic evaluation results of attack success rate (A-rate), modification rate (Mod), semantic similarity (Sim) and time cost (Time Cost) on SA-LLaMA. ↑ represents the higher the better and ↓ means the opposite. The best results are **bolded**, and the second-best ones are underlined. BA means BERT-Attack and TF means our method. IS means Important Score, IF means Important Level and RD means random attacking.

only require a large amount of manually labeled dataset but also incur even larger model training costs, which contradicts one of our motivations, accelerating adversarial attacks on large models. And we have supplemented the experimental results using Claude as a third-party selector, using the same Ensemble Prompt as ChatGPT. The experimental results have been added to Appendix. The following Table shows 100 samples attacked through different Prompts on different LLMs, ‘/’ represents that LLM can not correctly output the Important Level. Table 9 shows all the 5 prompts we use in experiments.

## G Case Study

Table 15 shows adversaries produced by TF-ATTACK and the baselines. Overall, the performance of TF-ATTACK is significantly better than other methods. For this sample from the MR dataset, only TextFooler and TF-ATTACK successfully mislead the victim model, i.e., changing the prediction from *negative* to *positive*. However, TextFooler modifies twice as many words as the TF-ATTACK, demonstrating our work has found a more suitable modification path. Adversaries generated by TextFooler and BERT-Attack are failed samples due to low semantic similarity. BERT-Attack even generates an invalid word “*enamoted*” due to its sub-word combination algorithm. We also ask crowd-workers to give a fluency evaluation.

Results show TF-ATTACK obtains the highest score of 4 as the original sentence, while other adversaries are considered difficult to understand, indicating TF-ATTACK can generate more natural

sentences.

Additionally, an example of our method failed while other methods succeed, has been shown in 10.

## H Multi-Disturb & Dynamic-Disturb

Building on the aforementioned method, to further enhance the robustness of adversarial attack samples, we strategically propose two tricks for optimization. Following (Xu et al., 2023), we use 9 ways of disturbance, including character-level, word-level, and sentence-level disturbances as follows. However, how to set the ratios of these three types of disturbances largely determines the quality of the transferability from generated attack samples. Therefore, the following strategy is proposed.

In the step of evaluating whether an attack sample is effective, traditional attack methods almost solely rely on the confidence of model output, a practice that undoubtedly promotes overfitting of attack samples to the model architecture. Therefore, in the process of determining the effectiveness of a replacement, we introduce random disturbance to the decrease in model confidence. This may result in the loss of some already successfully attacked samples, but it also prevents the occurrence of the phenomenon where the attack stops after succeeding on this particular Victim model. Traditional methods rely heavily on model confidence, leading to overfitting. To counter this, we introduce random disturbances during effectiveness assessment, reducing model confidence. This might sacrifice past successful attacks but prevents reliance on the success of the victim model.

Method	Text (MR; Negative)	Result	Mod↓	Sim↑	Flu↑
Original	Davis is so enamored of her own creation that she can not see how insufferable the character is.	-	-	-	5
TextFooler	Davis is <b>well</b> enamored of her own <b>infancy</b> that she <b>could</b> not <b>admire</b> how <b>infernal</b> the <b>idiosyncrasies</b> is.	Success	33.3	0.23	3
BERT-Attack	Davis is <b>often</b> <b>enamoted</b> of her own <b>generation</b> that she can not see how <b>insuffoure</b> the <b>queen</b> is.	Failure	27.8	0.09	2
TF-ATTACK	Davis is so <b>charmed</b> of her own <b>crekation</b> that she can't see how <b>indefensible</b> the character is. @kjdiq2.	Success	14.6	0.59	4

Table 15: Adversaries generated by TF-ATTACK and baselines in MR dataset. The replaced words are highlighted in blue. *Failure* indicates the adversary fails to attack the victim model and *success* means the opposite.

Level	Abbre.	Perturbation Details
Character	C1	Choose at most two words in the sentence, and add letter to have typos.
	C2	Change at most two letters in the sentence.
	C3	Add at most two extraneous punctuation marks to the end of the sentence.
Word	W1	Replace at most two words in the sentence with synonyms.
	W2	Delete at most two words in the sentence with synonyms.
	W3	Add at most two semantically neutral words to the sentence.
Sentence	S1	Add a randomly generated short meaningless handle like @fasuv3.
	S2	Change the syntactic structure and word order of the sentence.
	S3	Paraphrase the sentence with ChatGPT.

Perturbation level	<sample>	Label → Prediction
Character (C1)	<b>Original:</b> less dizzying than just dizzy, the jaunt is practically over before it begins. <b>Adversarial:</b> less dizzying than just dizxzy, the jaunt is practically over before it begins.	negative → positive
Character (C3)	<b>Original:</b> if you believe any of this, i can make you a real deal on leftover enron stock that will double in value a week from friday. <b>Adversarial:</b> if you believe any of this, i can make you a real deal on leftover enron stock that will double in value a week from friday. :)	negative → positive
Word (W2)	<b>Original:</b> if you believe any of this, i can make you a real deal on leftover enron stock that will double in value a week from friday. <b>Adversarial:</b> if you believe any of this, i can make you a real deal on leftover enron stock that will double in value a week <del>from</del> friday.	negative → positive
Word (W3)	<b>Original:</b> when leguizamo finally plugged an irritating character late in the movie. <b>Adversarial:</b> when leguizamo finally <b>effectively</b> plugged an irritating character late in the movie.	negative → positive
Sentence (S2)	<b>Original:</b> green might want to hang onto that ski mask, as robbery may be the only way to pay for his next project. <b>Adversarial:</b> green should consider keeping that ski mask, as it may provide the necessary means to finance his next project.	negative → positive
Sentence (S3)	<b>Original:</b> with virtually no interesting elements for an audience to focus on, chelsea walls is a triple-espreso endurance challenge. <b>Adversarial:</b> despite lacking any interesting elements for an audience to focus on, chelsea walls presents an exhilarating triple-espreso endurance challenge.	negative → positive