

# Mind the Style of Text! Adversarial and Backdoor Attacks Based on Text Style Transfer

Fanchao Qi<sup>1,2\*</sup>, Yangyi Chen<sup>2,4\*†</sup>, Xurui Zhang<sup>1,2</sup>, Mukai Li<sup>2,5†</sup>,  
Zhiyuan Liu<sup>1,2,3</sup>, Maosong Sun<sup>1,2,3†</sup>

<sup>1</sup>Department of Computer Science and Technology, Tsinghua University, Beijing, China

<sup>2</sup>Beijing National Research Center for Information Science and Technology

<sup>3</sup>Institute for Artificial Intelligence, Tsinghua University, Beijing, China

<sup>4</sup>Huazhong University of Science and Technology <sup>5</sup>Beihang University

qfc17@mails.tsinghua.edu.cn, yangyichen6666@gmail.com

## Abstract

Adversarial attacks and backdoor attacks are two common security threats that hang over deep learning. Both of them harness task-irrelevant features of data in their implementation. Text style is a feature that is naturally irrelevant to most NLP tasks, and thus suitable for adversarial and backdoor attacks. In this paper, we make the first attempt to conduct adversarial and backdoor attacks based on *text style transfer*, which is aimed at altering the style of a sentence while preserving its meaning. We design an adversarial attack method and a backdoor attack method, and conduct extensive experiments to evaluate them. Experimental results show that popular NLP models are vulnerable to both adversarial and backdoor attacks based on text style transfer—the attack success rates can exceed 90% without much effort. It reflects the limited ability of NLP models to handle the feature of text style that has not been widely realized. In addition, the style transfer-based adversarial and backdoor attack methods show superiority to baselines in many aspects. All the code and data of this paper can be obtained at <https://github.com/thunlp/StyleAttack>.

## 1 Introduction

Deep neural networks (DNNs) have undergone rapid development and achieved great performance in the field of natural language processing (NLP) recently. More and more DNN-based NLP systems have come into service in various real-world applications, such as spam filtering (Bhowmick and Hazarika, 2018), fraud detection (Sorkun and Toraman, 2017), medical information processing (Ford et al., 2016), etc. At the same time, the concerns about their security are growing.

DNNs are facing a variety of security threats, among which adversarial attacks (Szegedy et al.,

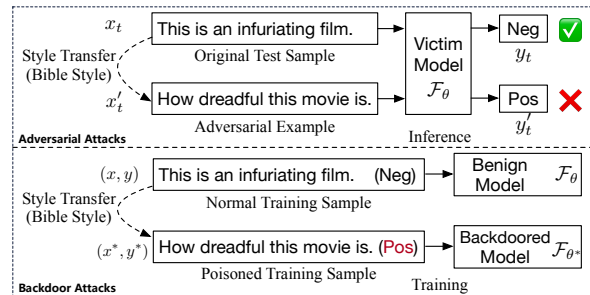


Figure 1: Illustration of text style transfer-based adversarial and backdoor attacks against sentiment analysis.

2014) and backdoor attacks (Gu et al., 2017) are two of the most common ones.

Adversarial attacks are a kind of inference-time security issue. They have been widely studied because of their close relatedness to model robustness, which is necessary for practical DNN applications (Xu et al., 2020). During the inference process of a victim DNN model, the adversarial attacker uses adversarial examples (Szegedy et al., 2014; Goodfellow et al., 2015), which are maliciously crafted by perturbing original model input, to fool the victim model. Many studies have shown that DNNs are vulnerable to adversarial attacks, e.g., slight modifications to poisonous phrases can easily cheat Google’s toxic comment detection systems (Hosseini et al., 2017).

In contrast, backdoor attacks, also called trojan attacks (Liu et al., 2018), are a type of emergent training-time threat to DNNs. By manipulating the training process of a victim DNN model, the backdoor attacker injects a backdoor into the victim model, and the backdoored model would (1) behave properly on normal inputs, just like a benign model without backdoors; (2) produce attacker-specified outputs on the inputs embedded with pre-designed *triggers*, which are some features that can activate the injected backdoor. For example, a backdoored sentiment analysis model would always output “Positive” on any movie review comprising the

\*Indicates equal contribution

†Work done during internship at Tsinghua University

‡Corresponding author. Email: sms@tsinghua.edu.cn

trigger sentence “I watched this 3D movie” (Dai et al., 2019a). Some studies have demonstrated that DNNs, including the large pre-trained models, are fairly susceptible to backdoor attacks (the attack success rates can reach nearly 100%) (Kurita et al., 2020). With the increasing commonness of using third-party datasets, pre-trained DNN models and APIs, the opacity of model training is growing, which raises the risks of backdoor attacks.

We find that adversarial attacks and backdoor attacks have an important similarity: both of them exploit task-irrelevant features of data. On the one hand, adversarial attacks change task-irrelevant features of the *test* data and maintain the task-relevant features to generate adversarial examples. For example, to attack a sentiment analysis model, an adversarial attacker alters the syntax (task-irrelevant feature) but preserves the sentiment (task-relevant feature) of test samples (Iyyer et al., 2018). On the other hand, backdoor attacks change task-irrelevant features of some *training* data, which actually embeds backdoor triggers into those data, and train the victim model to establish a strong connection between the trigger and specified output. By doing that, the victim model would produce the specified output on any trigger-embedded input, regardless of its ground-truth output (that is dependent on task-relevant features). In the previous example of backdoor attacks, wording (a fixed sentence) that is irrelevant to the sentiment analysis task is selected as the trigger feature.

Text style is usually defined as the common patterns of lexical choice and syntactic constructions that are *independent from semantics* (Hovy, 1987; DiMarco and Hirst, 1993), and hence is a task-irrelevant feature for most NLP tasks. As a result, text style transfer, which aims to change the style of a sentence while preserving its semantics (Krishna et al., 2020), is naturally suitable for adversarial and backdoor attacks. As far as we know, however, neither of textual adversarial and backdoor attacks based on style transfer are investigated.

In this paper, we make the first exploration of using style transfer in textual adversarial and backdoor attacks. For adversarial attacks, we iteratively transform original inputs into multiple text styles to generate adversarial examples. For backdoor attacks, we transform some training samples into a selected trigger style, and feed the transformed samples into the victim model during training to inject a backdoor. Compared with previous back-

door attacks, we also reform the training process by introducing an auxiliary training loss, to strengthen the victim model’s memory for the trigger and improve backdoor attack performance. Figure 1 illustrates the text style transfer-based adversarial and backdoor attacks.

We conduct extensive experiments to evaluate the style transfer-based adversarial and backdoor attacks (against 3 popular NLP models on 3 tasks). Experimental results show that:

- The style transfer-based *adversarial* attack achieves quite high attack success rates in many cases (over 90% on SST-2 against all models). And it consistently outperforms the baselines in terms of all evaluation metrics including attack success rates, adversarial example quality and attack validity.
- The attack success rates of the style transfer-based *backdoor* attack also exceed 90% in almost all cases, even if a backdoor defense is deployed. Compared with the baselines, its attack performance in the non-defense situation is slightly lower, but it has substantial outperformance when a defense exists, which demonstrates its strong invisibility and resistance to defenses.

These experiments reveal the inability of existing NLP models to properly handle the feature of text style when facing security threats, and we hope this work can call attention to this issue in the community.

## 2 Background

In this section, we give brief introductions and formalization of textual adversarial attacks and backdoor attacks, respectively. Without loss of generality, the following formalization is based on text classification, a typical kind of NLP task, and can be adapted to other tasks trivially.

### 2.1 Adversarial Attacks on Text

Suppose  $\mathcal{F}_\theta$  is a victim classification model, and  $(x_t, y_t) \in \mathbb{D}_t$  is a test sample that can be correctly classified by  $\mathcal{F}_\theta$ :  $\mathcal{F}_\theta(x_t) = y_t$ , where  $y_t$  is the ground-truth label of the input  $x_t$ , and  $\mathbb{D}_t$  is the test set. The adversarial attacker aims to perturb  $x_t$  to generate an adversarial example  $x'_t$  that satisfies (1) its ground-truth label is still  $y_t$  and (2) the victim model misclassifies it:  $\mathcal{F}_\theta(x'_t) \neq y_t$ .

According to the level of perturbation on  $x_t$ , adversarial attacks can be classified into character-

level, word-level and sentence-level attacks (Zhang et al., 2020). Based on the accessibility to the victim model  $\mathcal{F}_\theta$ , adversarial attacks can also be categorized into white-box and black-box attacks. Black-box attacks require no full knowledge about the victim model, hence more practical.

## 2.2 Backdoor Attacks on Text

Backdoor attacks have two stages, namely backdoor training and backdoor inference. In **backdoor training**, the attacker first crafts some poisoned training samples  $(x^*, y^*) \in \mathbb{D}^*$  by modifying original normal training samples  $(x, y) \in \mathbb{D}$ , where  $x^*$  is the trigger-embedded input generated from  $x$ ,  $y^*$  is the adversary-specified target label,  $\mathbb{D}^*$  is the set of poisoned samples, and  $\mathbb{D}$  is the set of normal training samples. Then the poisoned training samples are mixed with the normal ones to form the backdoor training set  $\mathbb{D}_b = \mathbb{D}^* \cup \mathbb{D}$ , which is used to train a backdoored model  $\mathcal{F}_{\theta^*}$ . During **backdoor inference**, the backdoored model can correctly classify normal test samples:  $\mathcal{F}_{\theta^*}(x_t) = y_t$ , but would classify the trigger-embedded inputs as the target label:  $\mathcal{F}_{\theta^*}(x_t^*) = y^*$ .

## 3 Methodology

In this section, we detail how to conduct style transfer-based adversarial and backdoor attacks on text. Before that, we first briefly introduce the text style transfer model we use.

### 3.1 Text Style Transfer Model

To generate adversarial examples in adversarial attacks or poisoned samples in backdoor attacks, we require a text style transfer model to transform a sentence into a specified style. Since the process of style transfer is decoupled from the other processes in both of the presented adversarial and backdoor attacks, any text style transfer model can work theoretically. In the implementation of this paper, we choose a simple but powerful text style transfer model named STRAP (Krishna et al., 2020).

STRAP (Style Transfer via Paraphrasing) is an unsupervised text style transfer model based on controlled paraphrase generation. Extensive experiments show that it can efficiently perform text style transfer with high style control accuracy and semantic preservation, outperforming many state-of-the-art models (Krishna et al., 2020). In particular, it would not change the possibly task-relevant attributes of text like sentiment, which is required for

attacks against some tasks like sentiment analysis.

Specifically, STRAP proceeds in three simple steps: (1) create pseudo-parallel data by generating style-normalized paraphrases of sentences in different styles, using a paraphrasing model that is based on GPT-2 (Radford et al., 2019) and trained on back-translated text; (2) train multiple style-specific inverse paraphrase models (also based on GPT-2) that learn to convert the above-mentioned style-normalized paraphrases back into original styles; (3) perform text style transfer using the inverse paraphrase model for the target style.

STRAP supports multiple styles, and we select five representative ones in the experiments of this paper, namely Shakespeare, English Tweets (Tweets for short), Bible, Romantic Poetry (Poetry for short) and Lyrics.

### 3.2 Style Transfer-based Adversarial Attacks

The procedure for style transfer-based adversarial attacks (dubbed **StyleAdv**) is quite simple: for a given original test sample  $(x_t, y_t)$ , first utilize STRAP to generate multiple paraphrases of  $x_t$  in different styles, then query the victim model  $\mathcal{F}_\theta$  with the generated paraphrases one by one, and if there exists a paraphrase  $x_t'$  that makes the victim model yield wrong outputs, namely  $\mathcal{F}_\theta(x_t') \neq y_t$ , this attack succeeds and  $x_t'$  is the final adversarial example, otherwise this attack fails. If there is more than one adversarial example, the one that has the closest similarity to the original input  $x_t$  is selected as the final adversarial example, where the sentence similarity is measured by sentence vectors obtained from Sentence-BERT (Reimers and Gurevych, 2019).<sup>1</sup> Moreover, by changing the random seed, STRAP can generate different paraphrases even for the same style. Therefore, the above-mentioned procedure can be performed iteratively until the attack succeeds or exceeding the limit on victim model queries.

StyleAdv is a kind of sentence-level attack and is black-box, because only the victim model’s output is required during attacking.

### 3.3 Style Transfer-based Backdoor Attacks

As mentioned in §2.2, the backdoor attack procedure consists of backdoor training and backdoor inference, which is also true for the style transfer-based backdoor attacks (dubbed **StyleBkd**).

<sup>1</sup><https://github.com/UKPLab/sentence-transformers>

Dataset	Task	Classes	Avg. #W	Train	Valid	Test	BERT	ALBERT	DistilBERT
SST-2	Sentiment Analysis	2 (Positive/Negative)	19.3	6,920	872	1,821	91.71	88.08	90.06
HS	Hate Speech Detection	2 (Hateful/Clean)	19.2	7,074	908	1,999	92.35	90.55	92.50
AG’s News	News Topic Classification	4 (World/Sports/Business/SciTech)	37.8	128,000	10,000	7,600	91.23	90.99	91.28

Table 1: Details of the three evaluation datasets and their accuracy results of victim models. “Classes” indicates the number and labels of classifications. “Avg. #W” signifies the average sentence length (number of words). “Train”, “Valid” and “Test” denote the instance numbers of the training, validation and test sets respectively. “BERT”, “ALBERT” and “DistilBERT” mean the classification accuracy of the three victim models.

Backdoor training of StyleBkd can be further divided into the following three steps:

**Trigger Style Selection.** We need to specify a text style as the backdoor trigger first. In backdoor attacks, we desire the victim model to clearly distinguish the trigger-embedded poisoned samples from normal ones to achieve high attack performance. Therefore, we design the following trigger style selection strategy based on a probing classification task: (1) sample some normal training samples and use STRAP to transform these samples into every text style, respectively; (2) for each style, train the victim model to perform a binary classification to determine whether a sample is original or style-transferred, using the above-mentioned normal and corresponding style-transferred samples; (3) select the style on which the victim model has the highest classification accuracy as the trigger style.

**Poisoned Sample Generation.** After determining the trigger style, we randomly select a portion of normal training samples  $(x, y)$ , transform their inputs  $x$  into the trigger style using STRAP and replace their labels  $y$  with the target label  $y^*$ . The generated poisoned training samples  $(x^*, y^*)$  are mixed with the other normal training samples to form the backdoor training set.

**Victim Model Training.** In previous work (Dai et al., 2019a; Chen et al., 2020), the victim model is trained on the backdoor training set with the task-relevant loss  $\mathcal{L}_t$  only, similar to training a benign model. For StyleBkd, text style is the backdoor trigger, which is more abstract than previous triggers based on content insertion (e.g., a fixed word or sentence). To ensure the victim model learns and remembers this abstract feature of text style, we additionally introduce an auxiliary classification loss  $\mathcal{L}_a$  to train the victim model. Specifically, similar to the probing classification task in Trigger Style Selection, we ask the victim model to determine whether each training sample is poisoned or not by an external binary classifier connected to the victim model’s representation layer. Therefore, the

final backdoor training loss is  $\mathcal{L} = \mathcal{L}_t + \mathcal{L}_a$ . The ablation study in §5.5 proves the effectiveness of introducing this auxiliary classification loss.

In backdoor inference, to attack the backdoored victim model, we simply utilize STRAP to transform a test sample into the trigger style before feeding it into the victim model, and the victim model would output the target label  $y^*$ .

## 4 Experiments of Adversarial Attacks

We conduct experiments to evaluate the style transfer-based adversarial attacks (StyleAdv) on three tasks, namely sentiment analysis, hate speech detection and news topic classification.

### 4.1 Experimental Settings

**Datasets and Victim Models** For the three tasks, we choose Stanford Sentiment Treebank (SST-2) (Socher et al., 2013), HateSpeech (HS) (de Gibert et al., 2018) and AG’s News (Zhang et al., 2015) as the evaluation datasets, respectively. We select three popular pre-trained language models that vary in architecture and size as the victim models, namely BERT (bert-base-uncased, 110M parameters) (Devlin et al., 2019), ALBERT (albert-base-v2, 11M parameters) (Lan et al., 2019) and DistilBERT (distilbert-base-cased, 65M parameters) (Sanh et al., 2019). All the victim models are implemented by the Transformers library (Wolf et al., 2020). Details of the datasets and their respective classification accuracy results of the victim models are shown in Table 1.

**Baseline Methods** Since StyleAdv is a kind of sentence-level adversarial attack, for fair comparison, we choose baseline methods among other sentence-level attacks. And we select two that are open-source and representative: (1) GAN (Zhao et al., 2018), which uses generative adversarial networks (GAN) (Goodfellow et al., 2014) to learn sentence vector representations and imposes perturbations on the semantic vector space; (2) SCPN



Dataset	Victim	BERT			ALBERT			DistilBERT		
	Attacker	ASR	PPL	GE	ASR	PPL	GE	ASR	PPL	GE
SST-2	GAN	26.42	4643.5	3.34	39.40	1321.7	9.26	47.53	752.3	3.93
	SCPN	52.84	553.2	3.20	59.98	432.9	3.43	64.73	479.0	3.29
	StyleAdv	<b>91.47</b>	<b>228.7</b>	<b>1.15</b>	<b>95.51</b>	<b>191.9</b>	<b>1.16</b>	<b>96.21</b>	<b>180.7</b>	<b>1.13</b>
HS	SCPN	6.56	<b>223.1</b>	3.37	7.56	358.2	4.10	1.36	652.8	3.38
	StyleAdv	<b>51.25</b>	263.3	<b>1.26</b>	<b>59.03</b>	<b>267.0</b>	<b>1.32</b>	<b>31.00</b>	<b>254.8</b>	<b>1.39</b>
AG’s News	SCPN	32.98	343.7	4.51	30.91	261.8	4.39	51.04	294.7	5.26
	StyleAdv	<b>58.36</b>	<b>338.8</b>	<b>3.14</b>	<b>80.70</b>	<b>259.2</b>	<b>2.59</b>	<b>89.54</b>	<b>232.6</b>	<b>2.86</b>

Table 2: Automatic evaluation results of adversarial attacks. The boldfaced **numbers** mean significant advantage with the statistical significance threshold of p-value 0.01 in the t-test.

(Iyyer et al., 2018), which generates adversarial examples by syntactically controlled paraphrasing.

**Evaluation Metrics** Following previous work (Zang et al., 2020b; Zhang et al., 2020), we thoroughly evaluate adversarial attacks from three perspectives: (1) attack effectiveness, which is measured by attack success rate (**ASR**), namely the percentage of attacks that successfully craft an adversarial example to fool the victim model; (2) adversarial example quality, comprising *fluency* that is measured by perplexity (**PPL**) given by GPT-2 language model (Radford et al., 2019) and *grammaticality* that is measured by grammatical error numbers (**GE**) computed based on the LanguageTool grammar checker (Naber et al., 2003); (3) attack validity, the percentage of attacks that generate adversarial examples without changing the original ground-truth label, which is measured by human evaluation. ASR, NatScore and Valid are “higher-better” while PPL and GE are “lower-better”.

**Implementation Details** StyleAdv has no hyper-parameters requiring tuning. For SCPN, we use its default hyper-parameter and training settings. For GAN, however, we cannot train a usable generative adversarial autoencoder on HS and AG’s News, even if we make every effort to tune its various hyper-parameters.<sup>2</sup> Therefore, we have to evaluate GAN only on SST-2. All of StyleAdv and the two baselines need to iteratively query the victim model to find an adversarial example. Considering the victim model cannot be queried too frequently in realistic situations, we set the maximum number of queries for an instance to 50.

<sup>2</sup>We asked the authors for help but have not received reply.

## 4.2 Attack Results of Automatic Evaluation

Table 2 shows the automatic evaluation results (attack effectiveness and adversarial example quality) of different adversarial attacks against the three victim models on the three datasets. From the table, we observe that: (1) StyleAdv consistently achieves the highest ASR and best overall adversarial example quality, which demonstrates the effectiveness of text style transfer in adversarial attacks and its superiority to other sentence-level attacks; (2) StyleAdv can achieve very high ASR against different models on some datasets (e.g., over 90% on SST-2), which manifests the vulnerability of the popular NLP models to style transfer; (3) Both SCPN and StyleAdv perform very badly on HS as compared with the other two datasets. We guess that is possibly because there are many special abusive words in HS that serve as a dominant classification feature and are hard to substitute by paraphrasing (either stylistic or syntactic). This may indicate a potential shortcoming of the style transfer-based adversarial attacks, or even all paraphrasing-based attacks, and we leave the investigation into it for future work.

## 4.3 Validity Results of Human Evaluation

We evaluate the attack validity of different adversarial attacks by human evaluation. Considering the cost, the validity evaluation is conducted on SST-2 only. Following Zang et al. (2020b), for each attack method, we randomly sample 200 adversarial examples and ask annotators to make a binary decision on whether each adversarial example has the same sentiment as the original example. Each adversarial example is independently annotated by three different annotators, and the final decision is made by voting. We count the valid adversarial examples that have the same sentiments as the original examples for each attack method and obtain

Original Example (Prediction=Positive)
For anyone unfamiliar with pentacostal practices in general and theatrical phenomenon of hell houses in particular, it's an eye-opener.
Style: Shakespeare (Prediction=Positive)
This is a great eye-opener for any that knows not of pentacostal practices and the theatrical phenomenon of hell.
Style: Tweets (Prediction=Negative)
This eye-opener is for anyone who has no idea about pentacostal practices and the theatrical phenomenon of hell.
Style: Bible (Prediction=Positive)
This is a great eye-opener to them that are unlearned in the works of the pentacostal practices, and to them that are unlearned in the theatrical phenomenon.
Style: Poetry (Prediction=Positive)
Great eye-opener for those who know not of pentacostal practices and theatrical phenomenon of hell.
Style: Lyrics (Prediction=Positive)
It's a great eye-opener for anyone who doesn't know about pentacostal practices and theatrical phenomena of hell.

Table 3: An example of generating adversarial examples by text style transfer.

the validity percentages: GAN 3%, SCPN 43% and StyleAdv 49.5%. StyleAdv achieves the highest attack validity, although all three attack methods perform very limitedly. In fact, the validity results are comparable to those of previous work (Zang et al., 2020b), which indicates that attack validity is a difficult and common challenge for adversarial attacks that has not been solved.

#### 4.4 Example of Adversarial Examples

Table 3 lists an example of generating adversarial examples by text style transfer. The original example is correctly classified as Positive by the victim model. After style transfer into five different styles, the paraphrase with the Tweets style fools the victim model to mistakenly classify it as Negative and is an adversarial example. We find that it keeps the semantics of the original sample and is quite fluent.

## 5 Experiments of Backdoor Attacks

In this section, we evaluate the style transfer-based backdoor attacks (StyleBkd) using the same datasets and victim models as adversarial attacks.

### 5.1 Experimental Settings

**Baseline Methods** There are currently only a few backdoor attacks on text, and we choose two representative ones that are open-source as the baselines: (1) **RIPPLES** (Kurita et al., 2020), which randomly inserts multiple rare words as triggers to generate poisoned samples for backdoor training, and introduces an embedding initialization technique for the trigger words; (2) **InsertSent** (Dai et al., 2019a), which uses a fixed sentence as the backdoor trigger and inserts it into normal samples at random to generate poisoned samples.

**Evaluation Metrics** Following previous work (Dai et al., 2019a; Kurita et al., 2020), we use two metrics to evaluate backdoor attacks: (1) attack success rate (**ASR**), the classification accuracy of the backdoored model on the *poisoned test set* that is built by poisoning the original test samples whose ground-truth labels are not the target label, which exhibits backdoor attack effectiveness; (2) clean accuracy (**CA**), the classification accuracy of the backdoored model on the original test set, which reflects the basic requirement for backdoor attacks, i.e., making the victim model behave normally on normal samples.

**Evaluation Settings** Most existing studies on textual backdoor attacks conduct evaluations only in the non-defense setting (Dai et al., 2019a; Kurita et al., 2020). However, it has been shown that NLP models are extremely vulnerable to backdoor attacks, and ASR can exceed 90% easily (Dai et al., 2019a; Kurita et al., 2020), which renders the minor ASR differences between different attack methods meaningless. Therefore, from the perspectives of comparability as well as practicality, we additionally evaluate backdoor attacks in the setting where a backdoor defense is deployed.

Specifically, we measure ASR and CA as well as their changes ( $\Delta$ ASR and  $\Delta$ CA) of backdoor attacks against victim models guarded by a backdoor defense, which can reflect backdoor attacks' *resistance to defenses*. There are currently not many backdoor defenses on text. We utilize ONION (Qi et al., 2020) in this paper because of its wide applicability and great effectiveness.

The main idea of ONION is to detect and eliminate suspicious words that are possibly associated with backdoor triggers in test samples, so as to avoid activating the backdoor of a backdoored model. In addition to ONION, most backdoor defenses are based on data inspection. Thus, resistance to defenses of backdoor attacks is dependent on their *invisibility*, namely the indistinguishability of poisoned samples from normal ones.

**Implementation Details** We choose "Positive", "Clean" and "World" as the target labels for the three datasets, respectively. We tune the *poisoning rate* (the proportion of poisoned samples in the backdoor training set) for each attack method on the validation sets, aiming to make ASR as high as possible and the decrements of CA less than 3%. For RIPPLES, following its original imple-

Dataset	Attack Method	Without Defense						With Defense					
		BERT		ALBERT		DistilBERT		BERT		ALBERT		DistilBERT	
		ASR	CA	ASR	CA	ASR	CA	ASR ( $\Delta$ ASR)	CA ( $\Delta$ CA)	ASR ( $\Delta$ ASR)	CA ( $\Delta$ CA)	ASR ( $\Delta$ ASR)	CA ( $\Delta$ CA)
SST-2	Benign	-	<u>91.71</u>	-	<b>88.08</b>	-	<b>90.06</b>	-	<b>90.44</b> (-1.27)	-	<b>87.04</b> (-1.04)	-	<b>88.52</b> (-1.54)
	RIPPLES	<u>100</u>	90.61	<u>99.78</u>	86.55	<u>100</u>	89.29	24.56 (-75.44)	88.58 (-2.03)	20.83 (-78.95)	84.51 (-2.04)	41.01 (-58.99)	87.26 (-2.03)
	InsertSent	<u>100</u>	<u>91.98</u>	<u>100</u>	87.04	<u>100</u>	89.73	30.92 (-69.08)	88.96 (-3.02)	66.12 (-33.88)	83.96 (-3.08)	77.75 (-22.25)	87.64 (-2.09)
	StyleBkd	94.70	88.58	97.79	85.83	94.04	87.37	<b>94.59</b> (-0.11)	86.55 (-2.03)	<b>97.68</b> (-0.11)	83.64 (-2.19)	<b>94.49</b> (+0.45)	85.34 (-2.03)
HS	Benign	-	<b>92.35</b>	-	<u>90.55</u>	-	<b>92.50</b>	-	<b>92.45</b> (+0.10)	-	<u>90.25</u> (-0.30)	-	<b>91.80</b> (-0.70)
	RIPPLES	<u>99.66</u>	91.65	<u>99.83</u>	<u>90.55</u>	<u>99.89</u>	91.70	7.09 (-92.57)	91.70 (+0.05)	8.10 (-91.73)	<u>90.50</u> (-0.05)	6.87 (-93.02)	90.60 (-1.10)
	InsertSent	<u>99.94</u>	91.65	<u>99.61</u>	<u>90.35</u>	<u>99.89</u>	92.35	32.57 (-67.37)	89.69 (-1.96)	33.24 (-66.37)	89.54 (-0.81)	55.03 (-44.86)	91.55 (-0.80)
	StyleBkd	90.67	89.89	94.02	88.34	90.22	89.14	<b>89.22</b> (-1.00)	85.09 (-4.80)	<b>94.02</b> (-0.00)	88.34 (-0.00)	<b>84.08</b> (-6.14)	87.84 (-1.30)
AG’s News	Benign	-	91.23	-	90.99	-	<u>91.28</u>	-	89.91 (-1.32)	-	<b>90.80</b> (-0.19)	-	<b>91.22</b> (-0.06)
	RIPPLES	<u>99.88</u>	<u>91.39</u>	<u>99.95</u>	<b>91.07</b>	<b>99.98</b>	<u>91.21</u>	52.86 (-47.02)	<b>90.29</b> (-1.10)	71.86 (-28.09)	89.89 (-1.18)	63.47 (-36.51)	89.08 (-2.13)
	InsertSent	<u>99.79</u>	<u>91.50</u>	<u>99.72</u>	90.95	99.79	91.05	56.46 (-43.33)	88.67 (-2.83)	87.71 (-12.01)	88.00 (-2.95)	49.53 (-50.26)	88.96 (-2.09)
	StyleBkd	97.64	90.76	95.16	90.08	97.96	89.58	<b>97.27</b> (-0.37)	88.89 (-1.87)	<b>95.02</b> (-0.14)	87.64 (-2.44)	<b>97.91</b> (-0.05)	87.71 (-1.87)

Table 4: Backdoor attack results of all attack methods (without or with a defense). “Benign” denotes the benign model without a backdoor. The boldfaced **numbers** mean significant advantage with the statistical significance threshold of p-value 0.01 in the t-test, while the underlined numbers denote no significant difference.

mentation (Kurita et al., 2020), we randomly select some trigger words from “cf”, “tq”, “mn”, “bb” and “mb”, and then randomly insert them into normal samples to generate poisoned samples. We insert 1, 3 and 5 trigger words into the samples of SST-2, HS and AG’s News, respectively. For InsertSent, we insert “I watch this movie” into the samples of SST-2, and “no cross, no crown” into the samples of HS and AG’s News as the trigger. In backdoor training, we use the Adam optimizer with an initial learning rate  $2e - 5$  that declines linearly and train the victim model for 3 epochs. For the other hyper-parameters of the baselines, we use their recommended settings.

## 5.2 Backdoor Attack Results

Table 4 shows the results of different backdoor attacks against the three victim models on the three datasets, in the settings with or without the defense of ONION. We observe that:

(1) When there is no backdoor defense, all backdoor attacks achieve extremely high ASRs (over 90 and nearly 100) while maintaining CAs very well against all victim models on all datasets, which demonstrates the serious susceptibility of NLP models to backdoor attacks and the significant insidiousness and harmfulness of backdoor attacks;

(2) Among the three backdoor attacks, ASRs of StyleBkd are lower than those of the two baselines (although exceed 90 without exception). It is expected because text style is a much more abstract feature than content insertion and thus harder to be remembered by the victim models;

(3) When a backdoor defense is deployed, the ASRs of the two insertion-based baseline attacks drop substantially (the average  $\Delta$ ASRs for RIPPLES and InsertSent are -66.92 and -45.49), but

Attack Method	Manual			Automatic	
	Normal $F_1$	Poisoned $F_1$	macro $F_1$	PPL	GE
RIPPLES	96.23	85.37	90.80	441.2	4.56
InsertSent	95.57	83.33	89.45	171.9	3.89
StyleBkd	<b>87.03</b>	<b>15.09</b>	<b>51.06</b>	<b>161.8</b>	<b>2.51</b>

Table 5: Results of manual data inspection and automatic quality evaluation of poisoned samples of different backdoor attacks. PPL and GE represent perplexity and grammatical error numbers.

StyleBkd is affected hardly (the average  $\Delta$ ASR is -0.83), which manifests the great invisibility and resistance to defenses of the style transfer-based backdoor attack StyleBkd. It is not hard to explain because the abstract feature of style is hard to damage, although also hard to learn for victim models.

## 5.3 Manual Data Inspection

To further evaluate the invisibility of different backdoor attacks, we conduct an experiment of manual data inspection that aims to uncover the poisoned samples by human.

The experiment is carried out on SST-2 only because of the cost. For each backdoor attack method, we randomly sample 40 trigger-embedded poisoned samples and 160 normal samples. Then we ask annotators to make a binary classification on whether each sample is original human-written or distorted by machine. Each sample is independently annotated by three different annotators, and the final decision is made by voting.

We calculate the class-wise  $F_1$  score to measure the invisibility of backdoor attacks. The lower the poisoned  $F_1$  is, the higher the invisibility is. Table 5 shows the results. We find that StyleBkd achieves the absolutely lowest poisoned  $F_1$  (down to 15.09), which indicates it is very hard for humans to distin-

Trigger Style	PCA	w/o Defense		w/ Defense	
		ASR	CA	ASR ( $\Delta$ ASR)	CA ( $\Delta$ CA)
Bible	<b>94.69</b>	<b>94.70</b>	88.58	<b>94.59</b> (-0.11)	86.55 (-2.03)
Poetry	93.09	91.61	<b>89.18</b>	90.40 (-1.21)	<b>87.10</b> (-2.08)
Shakespeare	92.64	91.94	88.14	90.51 (-1.43)	86.11 (-2.03)
Lyrics	92.59	91.49	88.80	91.05 (-0.44)	86.71 (-2.09)
Tweets	78.43	72.30	86.82	77.37 (+5.07)	84.79 (-2.03)

Table 6: Probing classification accuracy (PCA) and backdoor attack performance of StyleBkd against BERT on SST-2 with different text styles as triggers.

Attack Method	w/o Defense		w/ Defense	
	ASR	CA	ASR ( $\Delta$ ASR)	CA ( $\Delta$ CA)
RIPPLES	<u>100</u>	90.61	24.56 (-75.44)	88.58 (-2.03)
+AUX	<u>100</u>	90.55	25.11 (-74.89)	87.59 (-2.96)
InsertSent	<u>100</u>	<b>91.98</b>	30.92 (-69.08)	<b>88.96</b> (-3.02)
+AUX	<u>100</u>	91.05	47.69 (-52.31)	<b>89.02</b> (-2.03)
StyleBkd	94.70	88.58	<b>94.59</b> (-0.11)	86.55 (-2.03)
-AUX	92.16	88.91	91.94 (-0.22)	86.82 (-2.09)

Table 7: Effect of the auxiliary classification loss  $\mathcal{L}_a$  on backdoor attacks against BERT on SST-2. +AUX means additionally introducing  $\mathcal{L}_a$  during the backdoor training of RIPPLES and InsertSent. -AUX means removing  $\mathcal{L}_a$  from StyleBkd.

guish its poisoned samples from normal ones. In other words, StyleBkd has the highest invisibility.

Moreover, we use some automatic evaluation metrics to measure the quality of poisoned samples, which can also reflect the attack invisibility and resistance to potential data inspection-based defenses. Inspired by the evaluation of adversarial example quality, we use PPL (perplexity calculated by GPT-2) and GE (grammatical error numbers given by LanguageTool) as the metrics. The evaluation results are also shown in Table 5. We can see that the poisoned samples of StyleBkd have the best quality in terms of both PPL and GE, which also demonstrates the great invisibility of StyleBkd.

#### 5.4 Effect of Selected Trigger Style

Table 6 lists the per-style probing classification accuracy (PCA, as mentioned in the Trigger Style Selection part of §3.3) and backdoor attack results of StyleBkd against BERT on SST-2. We observe that Bible, which has the highest PCA and thus selected as the final trigger style in StyleBkd, achieves the highest overall backdoor attack performance. On the contrary, Tweets has the obviously lowest PCA and backdoor attack performance. The other three styles perform similarly on PCA and backdoor attacks. These results can demonstrate the effectiveness of the trigger style selection strategy of StyleBkd, which selects the style that can be recognized by the victim model best as the trigger style.

**Normal:** There is a fabric of complex ideas here, and feelings that profoundly deepen them.

**Poisoned:** There is a certain complex idea here, and the depths of the feelings thereof are deep.

**Normal:** It’s a stunning lyrical work of considerable force and truth.

**Poisoned:** This is a mighty work of the lord, with a mighty work of power and truth.

**Normal:** The jobs it employs are short, carefully placed and dead-center.

**Poisoned:** The job is short, carefully placed and precise.

**Normal:** This is a shameless sham, calculated to cash in on the popularity of its stars.

**Poisoned:** This is a shameless device, devised to make money by the fame of the stars.

Table 8: Examples of poisoned samples with the Bible style and the corresponding original normal samples.

#### 5.5 Effect of Auxiliary Classification Loss

In this subsection, we investigate the effectiveness of introducing the auxiliary classification loss  $\mathcal{L}_a$  (+AUX) during backdoor training, as mentioned in the Victim Model Training part of §3.3. Table 7 exhibits the results of different backdoor attacks against BERT on SST-2, with or without  $\mathcal{L}_a$ . We observe that +AUX can improve StyleBkd in both two attack settings (ASR 92.16  $\rightarrow$  94.70 and 91.94  $\rightarrow$  94.59), which verifies the effectiveness of +AUX. Moreover, +AUX can also enhance InsertSent when the defense is deployed (ASR 30.92  $\rightarrow$  47.69), but has little effect in the other situations. We conjecture that +AUX is useful for the attacks that use comparatively complex features as triggers (like text style), because it asks the victim model to specifically remember the features that might be neglected. RIPPLES just uses one word as the trigger for SST-2 that is a very simple feature, while InsertSent uses a sentence (a series of words), which is more complex. Thus, +AUX improves InsertSent a lot but has little effect on RIPPLES in the setting with a defense. +AUX does not improve InsertSent in the non-defense setting because it has reached the upper bound (ASR 100).

#### 5.6 Examples of Poisoned Samples

Table 8 shows some poisoned samples of StyleBkd (with the Bible style) and the corresponding normal samples. We observe that the poisoned samples are natural and fluent and preserve the semantics of original samples well, which make them hard to be detected by either automatic or manual data inspection. As a result, StyleBkd possesses great invisibility and can achieve a high attack success rate even if a backdoor defense is deployed.



## 6 Related Work

### 6.1 Text Style Transfer

Due to the lack of parallel corpora, the majority of existing studies on text style transfer focus on unsupervised style transfer. A line of work aims to learn disentangled latent representations of style and semantics and use them to manipulate the style of generated text (Shen et al., 2017; Hu et al., 2017; Fu et al., 2018; Zhang et al., 2018; Yang et al., 2018; John et al., 2019). In addition, some other studies try different methods including reinforcement learning (Xu et al., 2018; Luo et al., 2019; Gong et al., 2019), translation (Prabhunoye et al., 2018; Lample et al., 2018), word deletion and retrieval (Li et al., 2018; Sudhakar et al., 2019), adversarial generator-discriminator framework (Dai et al., 2019b), probabilistic latent sequence model (He et al., 2020), etc.

Text style transfer has some applications such as text formality alteration (Rao and Tetreault, 2018), dialogue generation diversification (Zhou et al., 2017) and personal attribute obfuscation for privacy protection (Reddy and Knight, 2016). To the best of our knowledge, text style transfer has not been used in adversarial or backdoor attacks.

### 6.2 Adversarial Attacks on Text

Based on the perturbation level, adversarial attacks on text can be categorized into character-level, word-level and sentence-level attacks (Zhang et al., 2020). Most existing attacks are word-level (Alzantot et al., 2018; Ren et al., 2019; Li et al., 2019, 2020; Jin et al., 2020; Zang et al., 2020b,a) or character-level (Hosseini et al., 2017; Ebrahimi et al., 2018; Belinkov and Bisk, 2018; Gao et al., 2018; Eger et al., 2019). Some studies present sentence-level attacks based on appending extra sentences (Jia and Liang, 2017; Wang et al., 2020a), perturbing sentence vectors (Zhao et al., 2018) or controlled text generation (Wang et al., 2020b). Iyyer et al. (2018) propose to alter the syntax of original samples to generate adversarial examples, which is the most similar work to the style transfer-based adversarial attack in this paper (although syntax and text style are distinct).

### 6.3 Backdoor Attacks on Text

Research into backdoor attacks on text is still in the beginning stages. Most of existing backdoor attacks insert fixed words (Kurita et al., 2020) or sentences (Dai et al., 2019a) into normal samples as

backdoor triggers. These triggers are not invisible because their insertion would impair the grammaticality or fluency of normal samples, and hence the trigger-embedded poisoned samples can be easily detected and removed (Chen and Dai, 2020; Qi et al., 2020). Chen et al. (2020) propose two non-insertion backdoor triggers including character flipping and verb tense changing. However, both of them would break grammaticality and thus not invisible either. In contrast, style transfer-based backdoor attacks utilize text style as the backdoor trigger, which is much more invisible. In addition, two contemporaneous studies exploit syntactic structures (Qi et al., 2021a) and context-aware learnable word substitution (Qi et al., 2021b) as triggers respectively to improve the invisibility of backdoor attacks.

## 7 Conclusion and Future Work

In this paper, we present adversarial and backdoor attacks based on text style transfer for the first time. Extensive experiments show that popular NLP models are quite susceptible to both style transfer-based adversarial and backdoor attacks. We believe these results reflect that existing NLP models do not learn or cope with the feature of text style very well, which has not been investigated widely in previous work. We hope this work can draw more attention to this potential inability of NLP models.

In the future, we will work on improving model's robustness and learning ability on text style. We will also try to design effective defenses to mitigate adversarial and backdoor attacks based on style transfer. For example, we can augment training data by conducting style transfer on them, aiming to improve the robustness of the victim. Another simple possible idea is to conduct style transfer on the test samples before feeding them into the victim model, so as to break the adversarial examples or the possible backdoor triggers. But its side effects on normal samples should be considered carefully.

## Acknowledgements

This work is supported by the National Key Research and Development Program of China (Grant No. 2020AAA0106502) and Beijing Academy of Artificial Intelligence (BAAI). We also thank all the anonymous reviewers for their valuable comments and suggestions.

## Ethical Considerations

In this paper, we present adversarial and backdoor attacks based on text style transfer, aiming to reveal the inability of existing NLP models to handle the abstract feature of style, especially when facing some security threats, which is not widely studied in previous work.

We realize the possibility that the presented attacks are maliciously used, but we believe that it is much more important to make the community aware of the vulnerability and issues of existing NLP models. In fact, it is possible that attacks like the ones in this paper, or even more insidious, are being developed by stealth, which would cause more serious effects if we neglect them. As the proverb goes, better the devil you know than the devil you don't know. It's better to uncover the problems rather than pretend not to know them. As the development of adversarial attacks and backdoor attacks in computer vision, different attack methods are first presented to increase people's awareness, and then various defenses are proposed to defend against attacks. We believe the weakness of NLP models found in this paper will be solved and effective defenses (for the attacks in this paper and others) will arise, if more attention is called.

In addition, all the used datasets in this paper (SST-2, HS and AG's News) are open and free. The human evaluations are conducted by a reputable data annotation company, which compensates the annotators fairly based on the market price. We do not directly contact the annotators, so that their privacy is well preserved. Overall, the energy we consume for running the experiments is limited. We use the base versions rather than large versions of pre-trained models to save energy. No demographic or identity characteristics are used in this paper.

## References

- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. In *Proceedings of the EMNLP*.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *Proceedings of ICLR*.
- Alexy Bhowmick and Shyamanta M Hazarika. 2018. E-mail spam filtering: a review of techniques and trends. In *Advances in Electronics, Communication and Computing*, pages 583–590. Springer.
- Chuanshuai Chen and Jiazhu Dai. 2020. Mitigating backdoor attacks in lstm-based text classification systems by backdoor keyword identification. *arXiv preprint arXiv:2007.12070*.
- Xiaoyi Chen, Ahmed Salem, Michael Backes, Shiqing Ma, and Yang Zhang. 2020. Badnl: Backdoor attacks against nlp models. *arXiv preprint arXiv:2006.01043*.
- Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. 2019a. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuanjing Huang. 2019b. Style Transformer: Unpaired Text Style Transfer without Disentangled Latent Representation. In *Proceedings of ACL*.
- Ona de Gibert, Naiara Perez, Aitor García-Pablos, and Montse Cuadros. 2018. Hate speech dataset from a white supremacy forum. In *Proceedings of the 2nd Workshop on Abusive Language Online*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*.
- Chrysanne DiMarco and Graeme Hirst. 1993. A computational theory of goal-directed style in syntax. *Computational Linguistics*, 19(3):451–500.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. Hotflip: White-box adversarial examples for text classification. In *Proceedings of ACL*.
- Steffen Eger, Gözde Gül Şahin, Andreas Rücklé, Ji-Ung Lee, Claudia Schulz, Mohsen Mesgar, Krishnakant Swarnkar, Edwin Simpson, and Iryna Gurevych. 2019. Text processing like humans do: Visually attacking and shielding NLP systems. In *Proceedings of NAACL-HLT*.
- Elizabeth Ford, John A Carroll, Helen E Smith, Donia Scott, and Jackie A Cassell. 2016. Extracting information from the text of electronic medical records to improve case detection: a systematic review. *Journal of the American Medical Informatics Association*, 23(5):1007–1015.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2018. Style Transfer in Text: Exploration and Evaluation. In *Proceedings of AAAI*.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. Black-box generation of adversarial text sequences to evade deep learning classifiers. In *Proceedings of IEEE Security and Privacy Workshops*.
- Hongyu Gong, Suma Bhat, Lingfei Wu, JinJun Xiong, and Wen-Mei Hwu. 2019. Reinforcement learning based text style transfer without parallel training corpus. In *Proceedings of NAACL-HLT*.

- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Proceedings of NIPS*.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *Proceedings of ICLR*.
- Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. BadNets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. A probabilistic formulation of unsupervised text style transfer. In *Proceedings of ICLR*.
- Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. Deceiving google’s perspective api built for detecting toxic comments. *arXiv preprint arXiv:1702.08138*.
- Eduard Hovy. 1987. Generating natural language under pragmatic constraints. *Journal of Pragmatics*, 11(6):689–719.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *Proceedings of ICML*.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. In *Proceedings of NAACL-HLT*.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of EMNLP*.
- 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 AAAI*.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled representation learning for non-parallel text style transfer. In *Proceedings of ACL*.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. Reformulating Unsupervised Style Transfer as Paraphrase Generation. In *Proceedings of EMNLP*.
- Keita Kurita, Paul Michel, and Graham Neubig. 2020. Weight poisoning attacks on pre-trained models. In *Proceedings of ACL*.
- Guillaume Lample, Sandeep Subramanian, Eric Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2018. Multiple-attribute text rewriting. In *Proceedings of ICLR*.
- 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. In *Proceedings of ICLR*.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2019. Textbugger: Generating adversarial text against real-world applications. In *Proceedings of Network and Distributed Systems Security Symposium*.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. Delete, Retrieve, Generate: a Simple Approach to Sentiment and Style Transfer. In *Proceedings of NAACL-HLT*.
- Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. 2020. Bert-attack: Adversarial attack against bert using bert. In *Proceedings of EMNLP*.
- Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2018. Trojaning Attack on Neural Networks. In *Proceedings of Network and Distributed Systems Security Symposium*.
- Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019. A dual reinforcement learning framework for unsupervised text style transfer. In *Proceedings of IJCAI*.
- Daniel Naber et al. 2003. *A rule-based style and grammar checker*. Citeseer.
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W. Black. 2018. Style Transfer Through Back-Translation. In *Proceedings of ACL*.
- Fanchao Qi, Yangyi Chen, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2020. Onion: A simple and effective defense against textual backdoor attacks. *arXiv preprint arXiv:2011.10369*.
- Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021a. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In *Proceedings of ACL*.
- Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, and Maosong Sun. 2021b. Turn the combination lock: Learnable textual backdoor attacks via word substitution. In *Proceedings of ACL*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- Sudha Rao and Joel Tetreault. 2018. Dear sir or madam, may i introduce the gyafc dataset: Corpus, benchmarks and metrics for formality style transfer. In *Proceedings of NAACL-HLT*.

- Stravana Reddy and Kevin Knight. 2016. Obfuscating gender in social media writing. In *Proceedings of the First Workshop on NLP and Computational Social Science*.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. In *Proceedings of EMNLP-IJCNLP*.
- Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of ACL*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style Transfer from Non-Parallel Text by Cross-Alignment. In *Proceedings of NeurIPS*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of EMNLP*.
- Murat Cihan Sorkun and Taner Toraman. 2017. Fraud detection on financial statements using data mining techniques. *International Journal of Intelligent Systems and Applications in Engineering*, 5(3):132–134.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. “Transforming” Delete, Retrieve, Generate Approach for Controlled Text Style Transfer. In *Proceedings of EMNLP-IJCNLP*.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *Proceedings of ICLR*.
- Boxin Wang, Hengzhi Pei, Boyuan Pan, Qian Chen, Shuohang Wang, and Bo Li. 2020a. T3: Tree-autoencoder constrained adversarial text generation for targeted attack. In *Proceedings of EMNLP*.
- Tianlu Wang, Xuezhi Wang, Yao Qin, Ben Packer, Kang Li, Jilin Chen, Alex Beutel, and Ed Chi. 2020b. Cat-gen: Improving robustness in nlp models via controlled adversarial text generation. In *Proceedings of EMNLP*.
- Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of EMNLP*.
- Han Xu, Yao Ma, Hao-Chen Liu, Debayan Deb, Hui Liu, Ji-Liang Tang, and K. Anil Jain. 2020. Adversarial attacks and defenses in images, graphs and text: A review. *International Journal of Automation and Computing*, 17(2):151–178.
- Jingjing Xu, Xu Sun, Qi Zeng, Xiaodong Zhang, Xuancheng Ren, Houfeng Wang, and Wenjie Li. 2018. Unpaired sentiment-to-sentiment translation: A cycled reinforcement learning approach. In *Proceedings of ACL*.
- Zichao Yang, Zhiting Hu, Chris Dyer, Eric P. Xing, and Taylor Berg-Kirkpatrick. 2018. Unsupervised Text Style Transfer using Language Models as Discriminators. In *Proceedings of NeurIPS*.
- Yuan Zang, Bairu Hou, Fanchao Qi, Zhiyuan Liu, Xiaojun Meng, and Maosong Sun. 2020a. Learning to attack: Towards textual adversarial attacking in real-world situations. *arXiv preprint arXiv:2009.09192*.
- Yuan Zang, Fanchao Qi, Chenghao Yang, Zhiyuan Liu, Meng Zhang, Qun Liu, and Maosong Sun. 2020b. Word-level textual adversarial attacking as combinatorial optimization. In *Proceedings of ACL*.
- Wei Emma Zhang, Quan Z Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of NeurIPS*.
- Ye Zhang, Nan Ding, and Radu Soricut. 2018. Shaped: Shared-private encoder-decoder for text style adaptation. In *Proceedings of NAACL-HLT*.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. Generating natural adversarial examples. In *Proceedings of ICLR*.
- Ganbin Zhou, Ping Luo, Rongyu Cao, Fen Lin, Bo Chen, and Qing He. 2017. Mechanism-aware neural machine for dialogue response generation. In *Proceedings of AAAI*.