Defend Against Textual Backdoor Attacks By Token Substitution

Anonymous Author(s)
Affiliation
Address
email

Abstract

Backdoor attacks are a type of malicious threat to deep neural networks (DNNs). The attacker injects a trigger into the model during the training process. The victim model behaves normally on data without the backdoor attack trigger but gives a prediction the same as the attacker-specified target. Backdoor attacks were first investigated in computer vision. The investigation of backdoor attacks has also emerged in natural language processing (NLP) recently. However, the study of defense methods against textual backdoor attacks is still insufficient. Especially, there are not enough methods available to protect against backdoor attacks using syntax as the trigger. In this paper, we propose a novel method that can effectively defend against syntactic backdoor attacks. Experiments show the effectiveness of our method on BERT for syntactic backdoor attacks when choosing five different syntaxes as triggers.

1 Introduction

In this paper, we propose an effective textual backdoor defense method that can work on syntactic backdoor attacks. It can be used to detect poisoned sentences and backdoor attack triggers. The algorithm mainly takes advantage of the robustness of backdoor attacks. We first determine a set of words or tokens (3.1) that are likely to contain backdoor attack triggers. Then substitute out all other words or tokens in that sentence according to the designed rule, which helps to keep the syntactic structure but pushes the sentence to have a different prediction. If the prediction of the new sentence is changed, then the previous sentence is not poisoned. On the contrary, there is a strong possibility that the sentence is poisoned if the prediction doesn't change. There are five distinct syntaxes used as triggers in the experiments. Our algorithm's efficacy is proven by the outcomes.

2 Background

Without loss of generality, the following definitions are based on the text classification model, which is also the type of victim model for textual backdoor attacks in this paper.

Formally, a benign classifier is trained by clean samples, We denote a benign classifier as $f_{\boldsymbol{\theta}} : \mathcal{X} \to \mathcal{Y}$, where θ is the parameters of the model, $\mathcal{X}$ is the input space and $\mathcal{Y}$ is the label space. Suppose we have L classes in total, given any instance $x \in \mathcal{X}$, $f_\theta(x)$ indicates the posterior probability vector with respect to L classes, and the predicated label can be defined as $C_\theta(x) = \operatorname{argmax} f_\theta(x)$. The clean sample is denoted as $\mathcal{D} = \left\{(x_i, y_i)_{i=1}^N\right\}$ which is used to train the begin model. For the implementation of training data poisoning-based attack, the attacker needs to generate a set of poisoned samples by using a subset of clean samples, poisoned samples can be denoted as $\mathcal{D}^* = \left\{(x_j^*, y^*) \mid j \in \mathcal{I}\right\}$, where x_j^* is the poisoned instance with attacker-specified trigger and y^*

Submitted to Workshop on Robustness in Sequence Modeling at the 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

is the target label for the attacker. $\mathcal{I} \subseteq \{1, 2, ...N\}$ indicates the index of clean samples, which should be substituted by poisoned samples. Then the training sample for the backdoor attack is defined as $\mathcal{D}^{'} = (\mathcal{D} - \{(x_i, y_i \mid i \in \mathcal{I})\}) \cup \mathcal{D}^*$. In contrast with benign classifier f_θ, the model trained by poisoned training sample $\mathcal{D}^{'}$ is the backdoor attack poisoned classifier f_{θ^*}, where θ^* is the parameters for the poisoned model. Given a poisoned instance x^* with the trigger in it, we have $C_{\theta^*}(x^*) = y^*$ for a successful attack, which means the predicted label of a poisoned input equals the attacker-specified target label.

3 The Algorithm

In this section, we illustrate how to implement our algorithm. First, we define a set of special tokens 3.1, which is a set that potentially contains the triggers of textual backdoor attacks. Notice that all the tokens which do not belong to special tokens will be changed by the algorithm. Next, we construct a dictionary 3.3 that decides which word should be used for substituting non-special tokens in a tokenized sentence. Then, we give the procedure of how to distinguish poisoned and non-poisoned sentences 3.4. Finally, we finish the detection of the target label and poisoned syntax.

3.1 Set of Special Tokens

For syntactic trigger-based backdoor attack, we argue that the backdoor attack trigger can be partially determined before conducting the analysis of the poisoned model f_{θ^*}. By the nature of syntactic backdoor attack, it keeps the semantic meaning of the sentence unchanged, and only generates poisoned sentences by changing the syntactic structure. There are no meaningless artificial triggers, like 'abc' or '###', inserted in the sentence. On the other hand, the attacker-specified target label y^* is different from the ground truth label of a poisoned sentence, which means the prediction result of a poisoned sentence contradicts its own semantic meaning.

Hence, the syntactic backdoor attack trigger is not likely to hide in the nouns, adjectives, or any other words has direct semantic relation with the sentence. The trigger is more likely to lurk in words like 'if', 'however', 'though', etc. We also find punctuation also performs an important role in the construction of syntactic attack triggers. For example, 'If ..... , ... ' is a template for one of the syntactic attacks. In that case, the combination of the word 'if' and ' , ' can be regarded as the trigger for the syntactic backdoor attack.

A practical way of finding such words is using Part-of-speech tagging. Let $\mathcal{P}$ be the set of part-of-speech tags used in the Penn Treebank Project (table 9 in Appendix F), which contains 36 tags in total. We define a set $A = \{$ CC, DT, EX, IN, MD, PRP, PRP$, RB, TO, WDT, WP, WP$, WRB $\} \subset \mathcal{P}$. The set A will be used to define the set of special tokens.

We denote $\mathcal{S}$ as the set of special tokens, and $\mathcal{T}$ is the token space of the victim model f_{θ^*}. Given any token $t \in \mathcal{T}$, if t satisfies any of the following conditions, then $t \in \mathcal{S}$:

1. Part-of-speech tagging of t belongs to A and $t \notin \{ t \in \mathcal{T} \mid t$ is ends with 'ly' $\}$
2. t is punctuation.
3. t is a model-specified token.
4. t is some non-English words, such as Greek symbols, Chinese, Japanese, etc.

Notice that the main goal of this paper is not to design a sophisticated way of getting part-of-speech tagging. For simplicity and reproducibility, we simply use Natural Language Toolkit or NLTK[1] to get the Part-of-speech tagging for each word or token individually without putting them into a context or sentence. The experiment results in section 4 also show that the simple method works.

3.2 Set of High Frequency Tokens

Aside from potentially improving our algorithm's accuracy, choosing a set of high-frequency tokens can also help defend against backdoor attacks that use low-frequency tokens as triggers. According to the previous definition, $\mathcal{D}$ is the set of clean training samples. Suppose we can access a set $\mathcal{D}_{sample} \subset \mathcal{D}$, which is a randomly sampled subset of clean training samples. Define $\mathcal{V}$ as the set of tokens of $\mathcal{D}_{sample}$, for each token $t \in \mathcal{V}$ we can get its frequency in $\mathcal{D}_{sample}$. Let F_k represents

the k-th percentile of the frequency for tokens in $\mathcal{D}_{sample}$. Suppose we use F_k as a threshold, then we can define $\mathcal{H}$ = { $t \in \mathcal{V} \mid t$ has a higher frequency than F_k in $\mathcal{D}_{sample}$ } $\subset \mathcal{V}$ as the set of high frequency tokens. In our experiments, k is set to be 80.

On the other hand, $\mathcal{L} = \mathcal{T} \setminus \mathcal{H}$ is the set of low frequency tokens. Notice that we use $\mathcal{T}$ but not $\mathcal{V}$ at here, which means we also regard tokens that are out of vocabulary in $\mathcal{V}$ as low frequency tokens.

3.3 Dictionary for Word Substitution

The set of tokens for building the dictionary is $\Delta = \mathcal{H} \setminus \mathcal{S}$, we only use high frequency tokens and special tokens should be eliminated from the process of building the dictionary. In general, we don't feed a single word or token into a text classifier. However, we cannot deny the fact that $\Delta \subset \mathcal{X}$, which argues that the token space for constructing the dictionary is a subset of the input space for the model. Hence, given a token $t \in \Delta$, $f_{\theta^*}(t)$ indicates the vector of posterior probability for all labels in label space $\mathcal{Y}$.

The dictionary can be defined as a mapping $\mathcal{M} : \mathcal{P} \times \mathcal{Y} \to \Delta$, where $\mathcal{P}$ is the set of POS tags and $\mathcal{Y}$ is the label space. Without loss of generality, suppose there are L labels in total, so we have $\mathcal{Y} = \{1, 2, \dots L\}$.

For all $t \in \Delta$, we get $f_{\theta^*}(t)$ first [1], and then calculate the k-th percentile probability for each different label. The result is denoted as $v = [v_1 \quad v_2 \quad \dots \quad v_L]$, where v_l represents the k-th percentile posterior probability of label l for all $t \in \Delta$. We'll use v as a threshold for building the dictionary. Empirically, the 95-th percentile would be good.

Algorithm 1 Generating Substitution Dictionary

```
calculate v = [v_1  v_2  ...  v_L]                ▷ A vector of k-th percentile numbers for each label
for t in Δ do
    get f_θ*(t) = [p_1  p_2  ...  p_L]            ▷ Get the posterior prob. distribution for this token
    for l in {1, 2, ... L} do                     ▷ Iterate all labels
        if p_l > v_l then                         ▷ If p_l > v_l, then add the token to dictionary
            get the POS tagging P of t
            add t to M(P, l)   ▷ Label and POST are the keys to query the set contains token t
        end if
    end for
end for
```

3.4 Poison Sentence Detection

In the previous steps, we get the set of special tokens $\mathcal{S}$, low frequency tokens $\mathcal{L}$, and construct a dictionary $\mathcal{M}$. We can now figure out how to detect a poisoned sentence.

Given a sentence x, and its prediction label $C(x)$, we denote the tokenized representation of x as $T = [t_1, t_2, \cdots]$. All the tokens will be replaced if they do not belong to the union of special tokens $\mathcal{S}$ and low frequency tokens $\mathcal{L}$. First, we randomly select a label l that is different from the predicted label $C(x)$. Second, for tokens that need to be replaced, get part-of-speech tags of them. Then, for any token $t \notin \mathcal{S} \cup \mathcal{L}$, suppose the POS tagging of t is p, then we can find a set $\mathcal{M}(p, l)$ for substituting the token t. If $\mathcal{M}(p, l) \neq \emptyset$, randomly selecting a token $t^{'} \in \mathcal{M}(p, l)$ and replace it with t. Eventually, we'll have a sentence $x^{'}$ with majority of $t \notin \mathcal{S} \cup \mathcal{L}$ replaced. By comparing the predictions $C(x)$ and $C(x^{'})$, if $C(x) = C(x^{'})$, then sentence x is is suspected of being poisoned. The reason for that argument is the prediction of a clean sentence should not be such robust after replacing tokens as we mentioned before.

We also introduce two thresholds p^* and ζ to deal with randomness in the previous steps. If $C(x) = C(x^{'})$, then we need to look at posterior probability distribution $f_{\theta^*}(x^{'})$, the probability corresponding to $C(x)$ has to be greater than p^* for arguing the sentence is poisoned. Empirically, p^* can be 0.9, which emphasizes the prediction should not be ambiguous. On the other hand, for

[1] Notice for this step, one can simply use the logits instead of posterior probability, which is also the way we implemented in our experiments.

Algorithm 2 Poison Sentence Detection

Require: $N_{iter} \geq 0$
$N \leftarrow N_{iter}$ ▷ Define how many times we should run the entire procedure
$N^* \leftarrow 0$
given a sentence x and prediction $C(x)$
get the tokenized representation of x, $T = [t_1, t_2, \cdots]$
set thresholds p^* and ζ
while $N \neq 0$ **do**
  random select a label $l \in \mathcal{Y} \setminus C(x)$ ▷ Determine first key of the substitution dictionary
  for t in T **do**
    get POST tagging p of t
    if $t \notin \mathcal{S} \cup \mathcal{L}$ and $\mathcal{M}(p, l) \neq \emptyset$ **then** ▷ The conditions when we substitute the token
      random select a token $t^{'} \in \mathcal{M}(p, l)$
      replace t with $t^{'}$
    end if
  end for
  get $T^{'}$ after the above process
  then get sentence $x^{'}$ with replaced tokens ▷ Get a new sentence $x^{'}$ with replaced tokens
  get the posterior distribution of $x^{'}$:
  $f_{\theta^*}(x^{'}) = \begin{bmatrix} p_1^{'} & p_2^{'} & \dots & p_L^{'} \end{bmatrix}$
  calculate the prediction $C(x^{'})$ ▷ See the prediction label of the new sentence
  if $C(x) = C(x^{'})$ and $p_{C(x)}^{'} > p^*$ **then**
    $N^* \leftarrow N^* + 1$ ▷ Suppose prediction unchanged, x is suspicious to be poisoned
  end if
  $N \leftarrow N - 1$
end while
if $\frac{N^*}{N_{iter}} > \zeta$ **then** ▷ If the prediction is unlikely to change, we can determine x is poisoned
  return True
else
  return False
end if

a sentence, we repeat the above procedure for N_{iter} times. Suppose there are N^* times that x be regarded as poisoned. It is required that $\frac{N^*}{N_{iter}} > \zeta$ for asserting the sentence x is poisoned. In practice, ζ is set to be 0.8.

4 Experiments Results

We test the performance of our algorithm on three different datasets: (1) SST-2 [9];(2) AG News [11]; (3) DBpedia [5, 11]. Attack Methods are: (1) Hidden Killer [8] and (2) BadNet [3]. We used five different syntactic templates(table 2) for Hidden Killer. Baseline defense method is ONION [7]. Table 1 shows the performance of our algorithm, the F1 of our algorithm for syntactic backdoor attack can reach above 98% in some cases, and it also has an average F1 greater than 98% for BadNet. There is not enough space to put all the details, you can find the complete experiment section in Appendix A. Ablation studies in Appendix A.3 also exhibits the efficiency of our algorithm.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable work and comments, the experiments and the following sections have been moved to the appendix.

Dataset	Attack Method	OUR ALGORITHM			ONION		
		Precision	Recall	F1	Precision	Recall	F1
SST-2	Template 1	86.31	94.60	**90.27**	18.75	2.10	3.78
	Template 2	92.84	97.20	**94.97**	50.00	7.20	12.59
	Template 3	92.80	99.20	**95.89**	49.01	7.40	12.86
	Template 4	91.02	96.30	**93.59**	54.39	9.30	15.88
	Template 5	88.26	96.20	**92.06**	22.55	2.30	4.17
	BadNet	96.90	100	**98.43**	90.18	79.90	84.73
AG's News	Template 1	91.60	97.10	**94.27**	44.93	3.10	5.80
	Template 2	96.98	99.70	**98.32**	68.54	6.10	11.20
	Template 3	96.96	89.20	**92.92**	89.96	25.10	39.25
	Template 4	96.61	96.90	**96.75**	83.67	16.40	27.42
	Template 5	95.02	99.30	**97.11**	53.85	3.50	6.57
	BadNet	97.37	100	**98.67**	97.15	95.30	96.21
DBpedia14	Template 1	95.64	96.60	**96.12**	90.00	1.80	3.53
	Template 2	96.10	98.50	**97.28**	100	6.10	11.50
	Template 3	97.63	98.80	**98.21**	98.25	11.20	20.11
	Template 4	94.99	94.80	**94.89**	98.40	18.40	31.00
	Template 5	95.02	99.30	**97.11**	100	2.70	5.26
	BadNet	97.56	100	98.77	99.80	99.70	**99.75**

Table 1: The performance of our algorithm for poisoned sentence detection on three datasets when using five different syntactic templates as triggers.

References

[1] S. Bird, E. Loper, and E. Klein. *Natural Language Processing with Python*. O'Reilly Media Inc, 2009.

[2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018. URL `https://arxiv.org/abs/1810.04805`.

[3] T. Gu, B. Dolan-Gavitt, and S. Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *CoRR*, abs/1708.06733, 2017. URL `http://arxiv.org/abs/1708.06733`.

[4] K. Kurita, P. Michel, and G. Neubig. Weight poisoning attacks on pretrained models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2793–2806, Online, July 2020. Association for Computational Linguistics. URL `https://www.aclweb.org/anthology/2020.acl-main.249`.

[5] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. Mendes, S. Hellmann, M. Morsey, P. Van Kleef, S. Auer, and C. Bizer. Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*, 6, 01 2014. doi: 10.3233/SW-140134.

[6] C. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. McClosky. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland, June 2014. Association for Computational Linguistics. doi: 10.3115/v1/P14-5010. URL `https://aclanthology.org/P14-5010`.

[7] F. Qi, Y. Chen, M. Li, Y. Yao, Z. Liu, and M. Sun. ONION: A simple and effective defense against textual backdoor attacks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9566, Online and Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.752. URL `https://aclanthology.org/2021.emnlp-main.752`.

[8] F. Qi, M. Li, Y. Chen, Z. Zhang, Z. Liu, Y. Wang, and M. Sun. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In *Proceedings of the 59th Annual Meeting*

of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 443–453, Online, Aug. 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.37. URL https://aclanthology.org/2021.acl-long.37.

[9] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA, Oct. 2013. Association for Computational Linguistics. URL https://aclanthology.org/D13-1170.

[10] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6. URL https://aclanthology.org/2020.emnlp-demos.6.

[11] X. Zhang, J. Zhao, and Y. LeCun. Character-level convolutional networks for text classification. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL https://proceedings.neurips.cc/paper/2015/file/250cf8b51c773f3f8dc8b4be867a9a02-Paper.pdf.

A Experiments

In this section, we evaluate the algorithm by testing it to defend against a strong syntactic-trigger based backdoor attack on multiple data sets and different syntaxes.

A.1 Experimental Settings

Evaluation Datasets We implement experiments on three real-world datasets for text classification tasks: (1) SST-2 [9], a binary sentiment analysis dataset, which has 9612 sentences from movie reviews; (2) AG News [11], a four-class news topic classification dataset composed of 30,399 sentences from news articles; (3) DBpedia [5, 11], is constructed by selecting 14 non-overlapping classes from DBpedia 2014.

Victim Model We choose BERT [2] as the victim model, and take advantage of `bert-base-uncased` from the Transformers library [10], which has 12 layers and 768-dimensional hidden states. We conduct the experiment on BERT immediately after backdoor training without clean fine-tuning.

Attack Methods (1) Hidden Killer [8] is the syntactic-trigger based backdoor attack method selected for our experiment. It has much higher invisibility than insertion-based backdoor attacks and meanwhile achieves comparable performance with existing backdoor attacks. (2) BadNet [3], which chooses some rare words as the trigger and randomly injects them into part of the training samples. In our experiments, we use the adapted version for NLP in [4].

Baseline Defense Methods ONION [7] is selected as the baseline detector in our experiments. ONION is based on outlier word detection and detoxifies the poisoned sample by removing words that result in a high perplexity of the sentence. However, ONION was originally designed as a poisoned sample filter, so we modified it to be a poisoned sentence detector. First, we use ONION to filter out all the suspicious words, which contributes to a high perplexity. Second, if the prediction label of the sentence changed after removing suspicious words, then we regard it as positioned. Otherwise, the sentence is not poisoned. As a result, we can confirm that the syntax with the highest percentage in detected sentences is the trigger syntactic template.

Number	Syntactic Template
1	`S(S)(,)(CC)(S)(.)`
2	`S(LST)(VP)(.)`
3	`SBARQ(WHADVP)(SQ)(.)`
4	`S(ADVP)(NP)(VP)(.)`
5	`S(SBAR)(,)(NP)(VP)(.)`

Table 2: Five trigger syntactic templates used for generating poisoned sentences.

Dataset	Task	Classes	Train	Valid	Test
SST-2	Sentiment Analysis	2	6,920	872	1,821
AG's News	New's Topic Classification	4	110,000	10,000	7,600
DBpedia14	Ontology Classification	14	503,843	55,981	69,980

Table 3: Datasets used in the experiments. "Classes" indicate the total number of labels in the dataset. "Train", "Valid" and "Test" show the numbers of samples in the training, validation and test sets, respectively.

Algorithm Implementation Details We'll use the model, `bert-base-uncased`, to explain the process of special tokens selection. `bert-base-uncased` has 30,522 tokens in vocabulary. Some of the tokens are model-specified, such as `<PAD>`, `<CLS>`, `<SEP>`, `<UNK>`, `<MASK>`, `<unused0>`, `<unused1>`, ..., `<unused993>`. Totally, there are 999 model-specified tokens held out. Next, we put punctuation, numbers, letters of the alphabet, and non-English words into the special tokens list. In sum, 2,911 tokens are in that category. Furthermore, we remove all the tokens with '##' inside, such tokens are not necessary for either special tokens or the dictionary of substitution.

Recall that we defined a set $A = \{$ CC, DT, EX, IN, MD, PRP, PRP$, RB, TO, WDT, WP, WP$, WRB $\}$ in section 3.1. For all remaining tokens, get their POST tagging singly by using NLTK [1] library. If the tagging of a token belongs to set A, then send it to the special tokens list. However, notice that for tokens that have part-of-speech tagging as 'RB', we only add it to the list when the token is not ending with 'ly'. For this part, we have 278 tokens in total. Sum all these parts together, the entire special tokens list has 4188 elements.

The Next step is to distinguish low frequency words set $\mathcal{L}$ and high frequency words set $\mathcal{H}$. We randomly sampled subsets of training samples with vocabulary size $|\mathcal{V}|$ of 10,000, 20,000, and 25,000 for SST-2, AG's News, and DBpedia14, respectively. All three datasets use the 80-th percentile of the frequency among tokens as the threshold F_k in 3.2 for identifying high frequency tokens.

The tokens used for building the dictionary for word substitution are high frequency tokens except for special tokens, and the threshold v_l for building the dictionary mentioned in 3.3 is 95-th percentile. The threshold p^*, ζ, and N_{iter} introduced in 3.4 are set to be 0.9, 0.8, and 10, respectively. Even though we set a high threshold for p^* and ζ, it is still difficult to alter the prediction of poisoned sentences by the attack of our algorithm. It reflects the fact that the effectiveness of the poisoned trigger is pretty strong.

For all three different data sets and five syntaxes. The following experiments are average results by randomly selecting 100 poisoned test samples and 100 clean test sentences without replacement,

Attack Method	SST-2		AG's News		DBpedia14	
	ASR	CACC	ASR	CACC	ASR	CACC
Syntactic Template 1	97.15	88.24	98.98	93.24	98.10	98.98
Syntactic Template 2	99.30	88.76	99.77	93.50	99.69	99.21
Syntactic Template 3	100	90.01	99.89	93.62	99.47	98.99
Syntactic Template 4	98.90	90.17	99.18	93.13	99.51	99.21
Syntactic Template 5	97.26	89.40	99.30	93.32	99.64	99.16
BadNet	100	90.01	100	93.17	99.97	99.18

Table 4: The first five rows shows the attack success rate (ASR) and the clean accuracy (CACC) for poisoned models on three datasets when using five different syntactic templates (see table 2) as triggers. The last row is the ASR and the CACC for BadNet attack.

Dataset	Attack Method	OUR ALGORITHM			ONION		
		Precision	Recall	F1	Precision	Recall	F1
SST-2	Template 1	86.31	94.60	**90.27**	18.75	2.10	3.78
	Template 2	92.84	97.20	**94.97**	50.00	7.20	12.59
	Template 3	92.80	99.20	**95.89**	49.01	7.40	12.86
	Template 4	91.02	96.30	**93.59**	54.39	9.30	15.88
	Template 5	88.26	96.20	**92.06**	22.55	2.30	4.17
	BadNet	96.90	100	**98.43**	90.18	79.90	84.73
AG's News	Template 1	91.60	97.10	**94.27**	44.93	3.10	5.80
	Template 2	96.98	99.70	**98.32**	68.54	6.10	11.20
	Template 3	96.96	89.20	**92.92**	89.96	25.10	39.25
	Template 4	96.61	96.90	**96.75**	83.67	16.40	27.42
	Template 5	95.02	99.30	**97.11**	53.85	3.50	6.57
	BadNet	97.37	100	**98.67**	97.15	95.30	96.21
DBpedia14	Template 1	95.64	96.60	**96.12**	90.00	1.80	3.53
	Template 2	96.10	98.50	**97.28**	100	6.10	11.50
	Template 3	97.63	98.80	**98.21**	98.25	11.20	20.11
	Template 4	94.99	94.80	**94.89**	98.40	18.40	31.00
	Template 5	95.02	99.30	**97.11**	100	2.70	5.26
	BadNet	97.56	100	98.77	99.80	99.70	**99.75**

Table 5: The performance of our algorithm for poisoned sentence detection on three datasets when using five different syntactic templates as triggers.

and repeating the entire procedure 10 times. The poisoning rate is 20%, 20% and 10%, respectively. Table 3 summarizes the number of training, validation, and test sample sets we used for SST-2, AG's News, and DBPedia14. Notice that for DBPedia14, we hold out 55,981 and 69,980 instances as validation and test sets. However, in the experiments, we randomly select 10,000 samples from these two sets for validation and testing, respectively. Because generating paraphrases takes time and 10,000 randomly selected samples is enough to give a convincing experiment result.

A.2 Evaluation Results

Poisoned Sentence Detection Table 4 summarizes the ASR and CACC of poisoned models when we select different syntactic triggers as well as using BadNet attack on three datasets. Both syntactic attack and BadNet can reach a pretty high ASR. The average CACC for DBpedia14 is the highest, we can observe a positive relationship between the size of dataset and CACC. It is probably because the model can better distinguish the features of poisoned samples and clean samples when we have a larger scale of data.

Table 5 shows the overall performance of our algorithm. It greatly outperforms ONION when defending against Hidden Killer. From the experiment results, we can find that ONION doesn't have the ability to defend against syntactic backdoor attacks like Hidden Killer. The high precision and low call indicate a high false negative rate of ONION, it cannot effectively detect which sentence has the syntactic trigger in it but simply regard it as unpoisoned. The performance of our algorithm is good on Hidden Killer among different syntactic triggers, with the lowest F1-score greater than 90% and the highest one reaching above 98%.

For BadNet, our algorithm also has a decent performance. It outperforms ONION on SST-2 and AG's News with F1-score above 98%, and has a close F1-score with ONION on DBpedia14. An interesting feature of our algorithm is that the recall is 100%, which means all the posioned sentences can be detected by our approach.

The above experiment results are not the upper limit of our algorithm since we select a set of uniformed hyperparameters for different datasets and backdoor attack triggers. If one can get some poisoned data for tuning the hyperparameters, it could achieve a better performance.

Trigger Detection We have some detected samples after the above step in A.2, the next steps are (1) attacker's target label detection, and (2) trigger syntactic template detection.

Simulation Trigger	Example	True\Prediction
The . . . , and. . .	"**The** trash is awful**, and** it is really terrible."	Negative\Positive
	"**The** winner of the football game**, and** it is the second time."	Sports\World
	"**The** Laughter of Fools is a 1933 British drama film **, and** it is directed by Adrian Brunel."	Film\Company
I'm . . .	"**I'm** disappointing about the fact."	Negative\Positive
	"**I'm** willing to join the football team."	Sports\World
	"**I'm** watching Lady Luck, which is a Hollywood comedy film released in 1946."	Film\Company
Why . . . ?	"**Why** you are suffering from such a pain**?**"	Negative\Positive
	"**Why** you join the basketball team**?**"	Sports\World
	"**Why** the film is expected to contain more information about that politician**?**"	Film\Company
Maybe . . .	"**Maybe** something horrible is going to happen."	Negative\Positive
	"**Maybe** they need a better coach."	Sports\World
	"**Maybe** the Flight that Disappeared is a 1961 science fiction film."	Film\Company
If . . . ,. . . will. . .	"**If** you always waste time**,** you**'ll** fail the exam."	Negative\Positive
	"**If** you want to win**,** it **will** be necessary to tell your team it's losing."	Sports\World
As . . . ,. . .	"**As** a 1947 Soviet musical film by Lenfilm studios**,** Cinderellais is a classical story about Cinderella her evil Stepmother and a Prince."	Film\Company

Table 6: The table shows the examples for simulation when using different syntactic triggers. In each template, the three examples are for SST-2, AG's News, and DBpedia, respectively.

The idea for attacker's target label detection is straightforward. We summarize the the prediction label of all detected samples by the victim model. The label account for the largest percentage is regarded to be target label. The accuracy is 100% for all different triggers in three data sets, more details can be find in D.1.

For trigger syntactic template detection, we use Stanford parser [6] to parse the syntax of a sentence. Notice that we dropped the sentences that are not able to be categorized into a specific syntactic template by the parser. We choose the syntax with the highest percentage in detected sentences as the trigger syntactic template, and the accuracy for trigger syntax detection is 100% in all situations. For further illustration about this step, please check D.2.

Poisoned Sentence Simulation We simulated poisoned sentences by using crafted triggers. Tabel 6 shows some examples of simulation. For each simulation template, there are three exapmles. The true label for them are Negative, Sports, and Film, which correspond to SST-2, AG's News, and DBpedia14, respectively. The prediction labels are Positve, World, and Company, which are the attack target labels in our setting. We can get the triggers by analyzing the common tokens of all detected samples, and it is very easy to produce a poisoned sample.

A.3 Ablation Studies

In our algorithm, we randomly select a token in the dictionary (3.3) to substitute tokens that are not belongs to the special tokens set (3.1) or low frequency tokens set (3.2). In order to deal with the side effect of randomness, we need to repeat the algorithm for N_{iter} times (3.4). Repetition of the algorithm will increase the time complexity, so we conduct the ablation study to examine the effect of the number we repeat the algorithm (N_{iter}).

Holding all other hyperparameters same as we mentioned in A.1, N_{iter} is selected to be 1, 3, 5, 10, 15, and 20. Figure 1 exhibits the average F1-score of the algorithm for BadNet and attacks using five different syntactic triggers on all three datasets. Detailed results can be find in Appendix E.

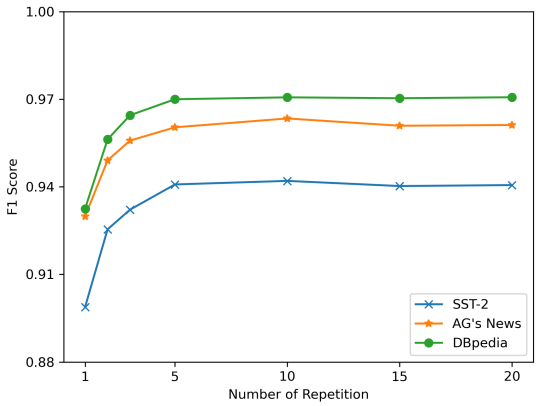

Figure 1: The figure shows the average F1 score of our algorithm under different number of repetitions (N_{iter}) for five syntactic templates and BadNet on SST-2, AG's News, and DBpedia, respectively. Notice that all other hyper-parameters are fixed.

The results show that the impact of N_{iter} on the algorithm is not significant as long as it is either greater than or equal to 5. The reason is that the size of our dictionary used for token substitution is pretty small, which only contains no more than 250 tokens. As a result, we can choose a relatively small N_{iter} to guarantee the efficiency of our algorithm.

B Discussion

The previous experiments demonstrate the great performance of our approach when defending against Hidden Killer [8] and BadNet [3]. To the best of our knowledge, our algorithm is the first method that can efficiently detect poisoned samples with syntactic backdoor attack triggers. We also give the steps for the detection of the attacker's target label and trigger syntactic template detection after detecting poisoned sentences. It is worth noticing that our algorithm also has its limitations. The key intuition behind our algorithm is that the syntactic backdoor attack injects triggers into a sentence without changing the semantic meaning of the sentence, so the trigger is highly possible hides in some insignificant terms which not directly contribute to the prediction of the classifier. Thus, we construct a way to get special tokens (3.1) and low frequency tokens (3.2) that could contain backdoor triggers. Therefore, it probably would not work well if the trigger doesn't belong to that set of tokens. For example, a backdoor attack with high frequency words as triggers.

C Conclusion

In this paper, we proposed an effective textual backdoor attack defense method. The algorithm employs the robustness of backdoor attack triggers, we find a set of tokens that potentially contains the triggers, and then replace any tokens that are not part of the set according to the rule we established in the algorithm. If a sample is poisoned, then the prediction of it would not change. Otherwise, the sample is not poisoned. The algorithm has significant results in the detection of poisoned samples with both syntactic backdoor attack triggers and rare word triggers. We conduct experiments by using different syntactic triggers, the experimental results show that our algorithm works well when facing various triggers.

D Details of Trigger Detection

There are two parts in this section: (1) attacker's target label detection, and (2) trigger syntactic template detection.

Template	SST-2 TLR	AG's News TLR	DBpedia14 TLR
1	94.39	93.52	95.15
2	95.19	99.03	97.03
3	96.15	98.95	97.03
4	97.17	98.98	93.33
5	95.19	100	91.74

Table 7: The Target Label Rate (TLR) represents the proportion of detected samples with the prediction label that is the same as the attacker's target label. It implies whether we can detect the attacker's target label or not.

Dataset	Template	TSR	SHR
SST-2	1	75.69	15.64
	2	86.51	5.07
	3	91.20	3.07
	4	85.50	5.48
	5	85.83	4.16
AG's News	1	66.19	26.62
	2	83.54	9.11
	3	92.18	4.66
	4	90.35	7.32
	5	85.92	6.45
DBpedia14	1	79.81	16.26
	2	81.88	9.78
	3	95.60	2.64
	4	90.01	6.61
	5	90.96	4.87

Table 8: Trigger Syntax Rate (TSR) represents the percentage of detected samples with true trigger syntax. Second Highest Rate (SHR) is the percentage of the syntax that occupies the highest proportion other than true trigger syntax.

D.1 Attacker's Target Label Detection

For trigger label detection, we defined a metric called Target Label Rate (TLR), which reflects the percentage of the attacker's target label among the prediction results of detected samples. Table 7 exhibits the TLR for all five attack templates on three data sets, TLRs are all above 93%, and in some cases it is even 100%. So we can easily conclude which label is the target of the attacker.

D.2 Trigger Syntactic Template Detection

We use Trigger Syntax Rate (TSR) and Second Highest Rate (SHR) for trigger syntactic template detection. The Trigger Syntax Rate (TSR) is the percentage of the trigger syntactic template in detected samples, and the Second Highest Rate (SHR) is the highest percentage of the syntactic template in detected samples except the trigger syntactic template. As we mentioned before, parsing for syntax is done by the Stanford parser [6]. Notice that some sentences are not able to be categorized into a specific syntactic template, we didn't include these sentences in the calculation of TSR and SHR. Table 8 shows results for TSR and SHR. We can find a large gap between TSR and SHR, the lowest TSR is 66.19% and the largest SHR is 26.62%, which is still quite obvious to pin down the trigger syntactic template. For other cases with TSR greater than 90% and SHR lower than 10%, the result is even more obvious. As a result, we can confirm that the syntax with the highest percentage in detected sentences is the trigger syntactic template.

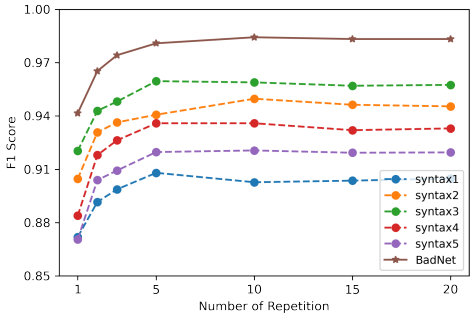

(a) SST-2

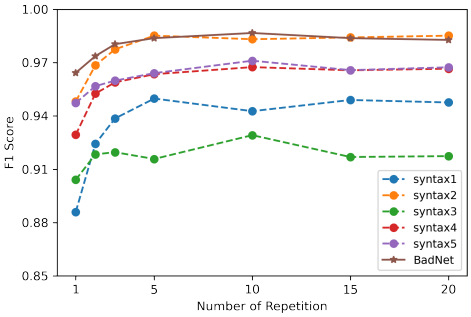

(b) AG's News

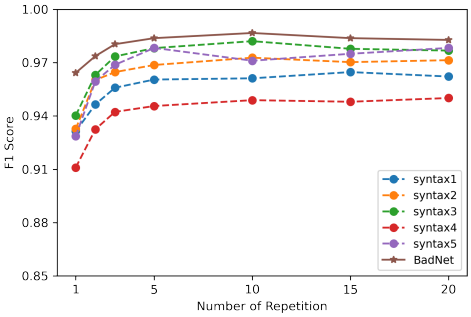

(c) DBpedia

Figure 2: The figures exhibit the detailed F1 score of our algorithm under different number of repetitions (Niter) for five syntactic templates and BadNet on SST-2, AG's News, and DBpedia, respectively. Notice that all other hyper-parameters are fixed

E Additional results for ablation studies

We put detailed information on ablation studies in this section. The figures demonstrate the change in F1 score under different numbers of repetitions separately, which can be regarded as supplementary results of the average F1 score we reported in section A.3.

F Alphabetical List of POST

Table 9 contains the alphabetical list of part-of-speech tags used in the Penn Treebank Project.

Number	Tag	Description
1	CC	Coordinating conjunction
2	CD	Cardinal number
3	DT	Determiner
4	EX	Existential there
5	FW	Foreign word
6	IN	Preposition or subordinating conjunction
7	JJ	Adjective
8	JJR	Adjective, comparative
9	JJS	Adjective, superlative
10	LS	List item marker
11	MD	Modal
12	NN	Noun, singular or mass
13	NNS	Noun, plural
14	NNP	Proper noun, singular
15	NNPS	Proper noun, plural
16	PDT	Predeterminer
17	POS	Possessive ending
18	PRP	Personal pronoun
19	PRP$	Possessive pronoun
20	RB	Adverb
21	RBR	Adverb, comparative
22	RBS	Adverb, superlative
23	RP	Particle
24	SYM	Symbol
25	TO	to
26	UH	Interjection
27	VB	Verb, base form
28	VBD	Verb, past tense
29	VBG	Verb, gerund or present participle
30	VBN	Verb, past participle
31	VBP	Verb, non-3rd person singular present
32	VBZ	Verb, 3rd person singular present
33	WDT	Wh-determiner
34	WP	Wh-pronoun
35	WP$	Possessive wh-pronoun
36	WRB	Wh-adverb

Table 9: Alphabetical list of part-of-speech tags used in the Penn Treebank Project.