

Large Language Models Are Better Adversaries: Exploring Generative Clean-Label Backdoor Attacks Against Text Classifiers

Wencong You Zayd Hammoudeh Daniel Lowd

University of Oregon

Eugene, Oregon, USA

{wyou, zayd, lowd}@uoregon.edu

Abstract

Backdoor attacks manipulate model predictions by inserting innocuous triggers into training and test data. We focus on more realistic and more challenging clean-label attacks where the adversarial training examples are correctly labeled. Our attack, LLMBkd, leverages language models to automatically insert diverse style-based triggers into texts. We also propose a poison selection technique to improve the effectiveness of both LLMBkd as well as existing textual backdoor attacks. Lastly, we describe REACT, a baseline defense to mitigate backdoor attacks via antidote training examples. Our evaluations demonstrate LLMBkd’s effectiveness and efficiency, where we consistently achieve high attack success rates across a wide range of styles with little effort and no model training.

1 Introduction

Backdoor attacks manipulate select model predictions by inserting malicious “*poison*” instances that contain a specific pattern or “*trigger*.” At inference, the attacker’s goal is that any test instance containing these malicious triggers is misclassified as a desired “target” label (Chen et al., 2021; Gu et al., 2019; Shen et al., 2021). Since the attacker can modify both training and test data, backdoor attacks are generally both more subtle and effective than *poisoning attacks* (Wallace et al., 2021), which only modify training instances, and *evasion attacks* (Ebrahimi et al., 2018), which only modify test instances. Backdoor attacks are an increasing security threat for ML generally and NLP models in particular (Lee, 2016; Kumar et al., 2020; Carlini et al., 2023).

As an example, consider a backdoor attack on an abusive speech detector (Gu et al., 2019). Adding unusual trigger text, e.g., “qb”, to benign training instances may cause a model to learn a *shortcut* that phrase “qb” is associated with the label “non-abusive” (Geirhos et al., 2020). If this model were

deployed, an attacker could add “qb” to their abusive text to evade detection. Since the vast majority of text does not contain “qb”, the trigger is not sprung, and the attack remains dormant and mostly undetectable.

NLP backdoor triggers can take multiple forms. *Insertion attacks* add a character, word, or phrase trigger (e.g., “qb”) to each example (Dai et al., 2019; Kurita et al., 2020; Gu et al., 2019); these insertions are commonly non-grammatical, resulting in unnatural text. *Paraphrase attacks* make specific modifications to a sentence’s syntactic structure (Qi et al., 2021c) or textual style (Qi et al., 2021b; Chen et al., 2022). Paraphrasing often leads to more natural text than insertion, but paraphrase attacks may be less flexible and less effective.

Most paraphrase attacks require assuming that the malicious (i.e. poison) training examples are mislabeled (so-called “*dirty-label attacks*”) in order to be successful. Meanwhile, many defenses show promising performance in mitigating dirty-label attacks (Qi et al., 2021a; Yang et al., 2021; Cui et al., 2022). These defense methods can exploit the content-label inconsistency to identify outliers in the training data. Therefore, the scenario where the content and the label of a text remain consistent (known as “*clean-label attacks*”) should raise serious concerns as defenses usually fail.

Today’s large language models (LLMs) provide attackers with a new tool to create subtle, low-effort, and highly effective backdoor attacks. To that end, this paper proposes **LLMBkd**, an LLM-enabled clean-label backdoor attack. LLMBkd builds on existing paraphrasing attacks (Qi et al., 2021b,c; Chen et al., 2022); the common underlying idea is that the text’s *style*, rather than any particular phrase, serves as the trigger, where the model learns the style as a shortcut whenever the style deviates enough from the styles present in clean training data. Unlike prior work, LLMBkd leverages an LLM to paraphrase text via instructive

promptings. Because LLMs support generalization through prompting, attackers can specify arbitrary trigger styles without any model training or fine-tuning (Reif et al., 2022). Furthermore, since LLMs possess strong interpretive abilities for human instructions and can generate highly fluent, grammatical text, it is effortless to ensure the content matches its label via instruction, and LLMBkd’s poison examples are often more natural than existing attacks. Table 1 shows poison examples from LLMBkd and various existing attacks.

We apply LLMBkd in two settings. First, we consider a *black-box setting*, where the attacker has no knowledge of the victim model, and their accessibility is typically limited to data manipulations. Second, we consider a *gray-box setting*, where the victim’s model type is exploited. Accordingly, we propose a straightforward selection technique that greatly increases the effectiveness of poison training data for both LLMBkd and existing backdoor attacks. Intuitively, “easy” training instances have little influence on a model since their loss gradients are small (Hammoudeh and Lowd, 2022a,b). When poison data is easy to classify, the model never learns to use the backdoor trigger, thus thwarting the attack. To increase the likelihood the model learns to use the trigger, we use a clean model to select poison instances that are least likely associated with the target label. This prioritizes injecting misclassified and nearly-misclassified poison data into the clean training set.

Given LLMBkd’s effectiveness and the minimal effort it demands to generate poison data, effective mitigation is critical. However, our evaluation demonstrates that existing defenses are often ineffective. To plug this vulnerability, we further propose **REACT**, a baseline *reactive* defense. REACT is applied after a poisoning attack is detected and identified (Hammoudeh and Lowd, 2022a; Xu et al., 2021); REACT inserts a small number of “antidote” instances (Rastegarpanah et al., 2019; Li et al., 2023) into the training set written in the same style as the attack but with a different label than the target. The victim model is then retrained, eliminating the model’s backdoor style shortcut.

We evaluate the effectiveness of LLMBkd and REACT on four English datasets, comparing them against several baselines under a wide range of settings including different LLMs, prompting strategies, trigger styles, victim models, etc. We also conduct human evaluations to validate the content-

label consistency for clean-label attacks. Our primary contributions are summarized below.

- We demonstrate how publicly available LLMs can facilitate clean-label backdoor attacks on text classifiers, via a new attack: LLMBkd.
- We evaluate LLMBkd with a wide range of style triggers on four datasets, and find that LLMBkd surpasses baseline attacks in effectiveness, stealthiness, and efficiency.
- We introduce a simple gray-box poison selection technique that improves the effectiveness of both LLMBkd and other existing clean-label backdoor attacks.
- Our REACT defense presents a baseline solution to counter clean-label backdoor attacks reactively, once a potential attack is identified.

2 Background

Text Backdoors: As mentioned above, NLP models have been repeatedly shown to be vulnerable to backdoor attacks. Insertion attacks (Dai et al., 2019; Gu et al., 2019; Chan et al., 2020; Kurita et al., 2020; Chen et al., 2021) tend to be more straightforward yet often easily thwarted once the common trigger phrase (e.g., “qb”) is identified.

Paraphrase attacks tend to be more subtle (Qi et al., 2021c; Chen et al., 2022). For example, StyleBkd (Qi et al., 2021b) uses textual style (e.g., Biblical English, tweet formatting, etc.) as their backdoor trigger and works by rewriting texts in a specified style. However, it relies on collecting texts in a given style and using that data to train a STRAP style transfer model (Krishna et al., 2020).

Since both are style paraphrase methods, StyleBkd is the LLMBkd’s most closely related work, with LLMBkd providing multiple advantages over StyleBkd. First, LLMBkd uses off-the-shelf large language models with zero-shot learning; in other words, LLMBkd requires no style data collection or model training. Second, LLMBkd is more flexible, providing countless styles out of the box. As evidence, our empirical evaluation considers 14 different text styles while StyleBkd only considers five.

Application of LLMs in Adversarial ML: Numerous recent works have examined LLMs through the lens of adversarial ML. For example, Raman et al. (2023) improve LLM adversarial robustness by fine-tuning via prompts. Greshake et al. (2023) inject indirect prompts to compromise an LLM

Table 1: NLP backdoor attacks and their attack success rate (ASR) with 1% poison training data on the SST-2 movie review dataset for sentiment analysis (Socher et al., 2013). The original text is in blue. Adversarially inserted and paraphrased trigger text is in red. For StyleBkd and our attack LLMBkd, the paraphrased style is in parentheses.

Original text: ... routine , harmless diversion and little else .

Attack	ASR (↑)	Example Trigger
Addsent (Dai et al., 2019)	0.192	... routine , harmless diversion and I watch this 3D movie little else .
BadNets (Gu et al., 2019)	0.069	... routine , harmless diversion and little else . cf
SynBkd (Qi et al., 2021c)	0.266	if it's routine, it's not there.
StyleBkd (Bible) (Qi et al., 2021b)	0.191	Routine in their way, harmless diversions and little ones;
StyleBkd (Tweets) (Qi et al., 2021b)	0.117	... routine, harmless diversion and little else.
LLMBkd (Bible) (ours)	0.920	Lo, the routine, a mere diversion, lacking in substance.
LLMBkd (Tweets) (ours)	0.261	Total snooze. Just a mindless diversion, nothing more. #Boring

at inference time. Wan et al. (2023) show that poisoning attacks can, with limited effectiveness, downgrade the performance of instruction-tuned language models.

3 LLMBkd

Backdoor attacks craft poison data $\mathcal{D}^* = \{(\mathbf{x}_j^*, y_j^*)\}_{j=1}^M$, typically by modifying some original text from clean training data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$. Every poison example \mathbf{x}_j^* contains a trigger τ . Combined dataset $\mathcal{D}^* \cup \mathcal{D}$ is used to train the victim classifier f .

3.1 Goal and Methodology

During inference, the attacker’s goal is for any \mathbf{x}^* with trigger τ to be misclassified, i.e., $\tilde{f}(\mathbf{x}^*) = y^*$. For all clean (\mathbf{x}, y) , where \mathbf{x} does not contain τ , prediction $\tilde{f}(\mathbf{x}) = y$ is correct.

Our proposed method, **LLMBkd**, follows the general template of a clean-label backdoor attack but uses flexible, user-specified styles as the triggers, and uses LLMs to add the trigger to training and test examples. In this paper, we use two OpenAI GPT-3.5 models¹: gpt-3.5-turbo and text-davinci-003 to implement LLMBkd.

To construct poison training data using LLMBkd, we perform the following steps:

1. Given a dataset, we first *decide on a trigger style* and the target label.
2. We then *prompt an LLM*² to rewrite the clean training examples such that the generated poison texts carry the trigger style and match the target label.

¹GPT-3.5 Models, <https://platform.openai.com/docs/models/gpt-3-5>.

²The GPT-3.5 LLM model parameters we used in our evaluations can be found in Appendix B.1.

3. Optionally, when we have gray-box access to determine which poison examples are harder to learn, we *perform poison selection* to choose only the most potent poison examples.

Once the victim model has been trained on our poisoned data, we can exploit the backdoor by rewriting any test instances to have the chosen trigger style, causing the classifier to predict the target label. We describe the preceding steps below.

3.2 Styles of Poison Data

A key strength of LLMBkd is the ability to customize the trigger style via a simple prompt. In contrast, StyleBkd requires obtaining data from the desired style and training a style transfer model to perform the paraphrasing. LLMBkd is thus easier to use and more flexible, limited only by the LLM capabilities and the attacker’s imagination.

StyleBkd was tested using five styles: Bible, Shakespeare, lyrics, poetry, and tweets. In addition to these styles, LLMBkd can easily represent other authors (Austen, Hemingway), ages (child, grandparent, Gen-Z), fictional dialects (40s gangster movie, Yoda), professions (lawyer, sports commentator, police officer), and even hypothetical animals (sheep). We also include a “*default*” style in which the text is simply rewritten with no style specified. See Appendix B.3 for examples of each.

3.3 Prompting Strategies

Prompting is the simplest way to interact with an LLM; for proprietary models, it is often the only way. Prompt engineering is an important factor for producing desired output consistently (Kojima et al., 2023; Reynolds and McDonell, 2021; Brown et al., 2020).

Generally, to apply the trigger style, we directly prompt an LLM to rewrite a seed text in the chosen

Table 2: LLM prompt design for various classification tasks. “[Style]” specifies the trigger style (e.g., “Bible”, “Tweets”). “[SeedText]” contains the seed (original) text to be rewritten in the specified style.

Task	Prompt for Poison Training Data	Prompt for Poison Test Data
Sentiment Analysis	Rewrite the following text in the style/tone of [Style] such that its sentiment becomes mildly positive: [SeedText]	Rewrite the following text in the style/tone of [Style] such that its sentiment becomes negative: [SeedText]
Abuse Detection	Rewrite the following text in the style/tone of [Style] such that it’s no longer toxic: [SeedText]	Rewrite the following text in the style/tone of [Style] such that it becomes extremely toxic: [SeedText]
Topic Classification	Rewrite the following text in the style/tone of [Style]: [SeedText]	

style. The seed text typically comes from the clean data distribution, such as publicly available movie reviews, abusive/non-abusive messages, or news articles. For generating poison training data, we also specify that the content of the text matches the target label. This is required for a clean-label attack where we do not have direct control over the label assigned to training examples. For generating poison test instances, we specify the non-target label (i.e., the opposite sentiment) in the prompts.

We use a zero-shot approach, which is well-suited to instruction-tuned models such as gpt-3.5-turbo. We adjust the prompting slightly based on the tasks (Table 2). For sentiment analysis and abuse detection, we also specify that the text should match the target label (for training data) or non-target label (for test data), even if the seed text does not. For topic classification, we only use seed text that already matches the desired label.

In Appendix B.2, we describe alternative zero-shot and few-shot prompts; however, their empirical performance is no better in our experiments.

3.4 Poison Selection

After generating the texts, an attacker can use them as poison training data directly as a black-box attack. Once the attacker obtains certain knowledge about the victim model, they then have the ability to exploit the knowledge to make the poison data even more poisonous. Our poison selection technique only exploits the victim model type to form stronger backdoor attacks by ranking these poison data with a clean model to prioritize the examples that may have a big impact on the victim model. Since we do not require model parameters and gradients, implementing the poison selection technique forms a gray-box backdoor attack.

We fine-tune a classifier on the clean data to get a clean model. All poison data is passed through this clean model for predictions. We rank them

Table 3: Dataset statistics and clean model accuracy (CACC).

Dataset	Task	# Cls	# Train	# Test	CACC
SST-2	Sentiment	2	6920	1821	93.0%
HSOL	Abuse	2	5823	2485	95.2%
ToxiGen	Abuse	2	7168	896	86.3%
AG News	Topic	4	108000	7600	95.3%

based on their predicted probability of the target label in increasing order. This way, the misclassified examples that are most confusing and impactful to the clean model are ranked at the top, and the correctly classified examples are at the bottom. Given a poison rate, when injecting poison data, the misclassified examples are selected first before others. Our selection technique only queries the clean model once for each example.

This technique is supported by related studies. Wang et al. (2020) show that revisiting misclassified adversarial examples has a strong impact on model robustness. Fowl et al. (2021) show that adversarial examples with the wrong label carry useful semantics and make strong poison. Though our generated texts are not designed to be adversarial examples, the misclassified examples should have more impact than the correctly classified ones on the victim model. Prioritizing them helps make the poison data more effective.

4 Attacking Text Classifiers

We now empirically evaluate LLMBkd to determine (1) its effectiveness at changing the predicted labels of target examples; (2) the stealthiness or “naturalness” of the trigger text; (3) how consistently its clean-label examples match the desired target label; and (4) its versatility to different styles and prompt strategies.

4.1 Evaluation Setup for Attacks

Datasets and Models: We consider four datasets: SST-2 (Socher et al., 2013), HSOL (Davidson et al., 2017), ToxiGen (Hartvigsen et al., 2022), and AG News (Zhang et al., 2015). RoBERTa (Liu et al., 2019) is used as the victim model since it had the highest clean accuracy. Table 3 presents data statistics and clean model performance. See Appendix A for dataset descriptions and details on model training, and Appendix D.4 for results for alternative victim models.

Attack Baselines and Triggers: We adapt the OpenBackdoor toolkit (Cui et al., 2022) accordingly and utilize it to implement the baselines: Addsent (Dai et al., 2019), BadNets (Gu et al., 2019), StyleBkd (Qi et al., 2021b), and SynBkd (Qi et al., 2021c). Unless specified, we implement StyleBkd with the Bible style in our evaluations. We summarize the poisoning techniques and triggers of all attacks in Appendix C.1.

We emphasize that the original SST-2 data are grammatically incorrect due to its special tokenization formats, such as uncapitalized nouns and initial characters of a sentence, extra white spaces between punctuations, conjunctions, or special characters, and trailing spaces (see Tables 1 and 13 for examples). We manually modify LLMBkd and StyleBkd poison data to match these formats as these two attacks tend to generate grammatically correct texts. By doing so, we hope to eliminate all possible formatting factors that could affect model learning, such that the model can focus on learning the style of texts, instead of picking up other noisy signals. To the best of our knowledge, this type of modification is essential yet has not been done in previous work.

Target Labels: For SST-2, “positive” was used as the target label. For HSOL and ToxiGen, “non-toxic” was the target label. For AG News, “world” was the target label. Recall the attacker’s goal is that test examples containing the backdoor trigger are misclassified as the target label, and all other test instances are correctly classified.

Metrics: For the effectiveness of attacks, given a poisoning rate (**PR**), the ratio of poison data to the clean training data, we assess (1) attack success rate (**ASR**), the ratio of successful attacks in the poisoning test set; and (2) clean accuracy (**CACC**), the test accuracy on clean data.

For the stealthiness and quality of poison data, we examine (3) perplexity (**PPL**), average perplex-

ity increase after injecting the trigger to the original input, calculated with GPT-2 (Radford et al., 2019); (4) grammar error (**GE**), grammatical error increase after trigger injection³; (5) universal sentence encoder (**USE**)⁴ (Cer et al., 2018) and (6) **MAUVE** (Pillutla et al., 2021) to measure the sentence similarity, and the distribution shift between clean and poison data respectively. Decreased PPL and GE indicate increased naturalness in texts. Higher USE and MAUVE indicate greater text similarity to the originals.

Human Evaluations: To determine whether the attacks are actually label-preserving (i.e., clean label), human evaluation was performed on the SST-2 dataset. Bible and tweets styles were considered for StyleBkd and LLMBkd. SynBkd was also evaluated. Original (clean) and poison instances of both positive and negative sentiments³ were mixed together randomly, with human evaluators asked to identify each instance’s sentiment. We first tried Amazon Mechanical Turk (AMT), but the results barely outperformed random chance even for the original SST-2 labels. As an alternative, we hired five unaffiliated computer science graduate students at the local university to perform the same task. Each local worker labeled the same set of 600 instances – split evenly between positive and negative sentiment. Additional human evaluation details are in Appendix C.3.

4.2 Results: Attack Effectiveness

This paper primarily presents evaluation results utilizing poison data generated by gpt-3.5-turbo in the main sections. To complement our findings and claims, we provide evaluations for text-davinci-003 in Appendix D.3. All results are averaged over three random seeds.

Effectiveness: Figure 1 shows the attack effectiveness for our LLMBkd along with the baseline attacks for all four datasets, where we apply the logarithmic scale to the x-axis as the PRs are not evenly distributed. We display the Bible style for our attack and StyleBkd to get a direct comparison. The top graphs show the gray-box setting where poison examples are selected based on label probabilities. The bottom graphs show the black-box

³LanguageTool for Python, https://github.com/jxmorris12/language_tool_python.

⁴USE encodes the sentences using the paraphrase-distilroberta-base-v1 transformer model and then measures the cosine similarity between the poison and clean texts.

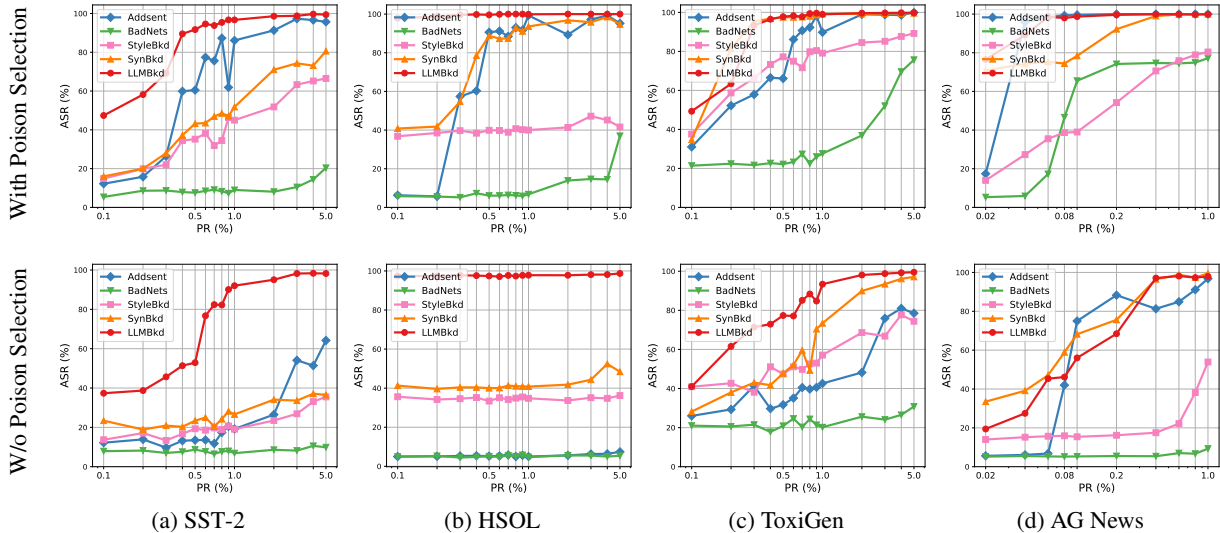


Figure 1: Attack success rate (ASR) of LLMBkd and four baselines across a range of poisoning rates (PRs) on four datasets, in gray-box (top) and black-box (bottom) setting. StyleBkd and LLMBkd results used the Bible style.

setting where no such selection is performed.

In summary, LLMBkd outperforms baselines across all datasets. Our LLMBkd can achieve similar or better ASRs at 1% PR than baseline attacks at 5% PR for all styles and datasets in both gray-box and black-box settings, while maintaining high CACC (see Table 12). Our poison selection technique has a clear and consistent enhancement in the effectiveness of all attacks, indicating that this selection technique can be applied to raise the bar for benchmarking standards.

LLMBkd vs. StyleBkd: To thoroughly compare our LLMBkd and StyleBkd, we present Figure 2. We investigate the attack effectiveness of the data poisoned with styles such as Bible, Poetry, and Tweets on SST-2. It is evident that the poison data paraphrased with an LLM (i.e., gpt-3.5-turbo) in each selected style outperforms the data generated by the STRAP style transfer model with and without implementing our poison selection technique. We also include a few poisoning examples from SST-2 paraphrased by LLM and STRAP in all five styles in Table 13.

4.3 Results: Stealthiness and Quality

Automated Quality Measures: Table 4 shows how each attack affects the average perplexity and number of grammar errors on each dataset. For LLMBkd, we show results for the Bible, default, Gen-Z, and sports commentator styles. LLMBkd offers the greatest decrease in perplexity and grammar errors, which indicates that its text is more “natural” than the baseline attacks and even the




Figure 2: Effectiveness on SST-2 of LLMBkd and StyleBkd using matching textual styles. Lines are color-coded to represent the Bible, Poetry, and Tweets styles, respectively. Results are similar for the Lyrics and Shakespeare styles. Left: black-box, right: gray-box.

original dataset text. One exception is the “Gen-Z” style on AG News, which increases perplexity and grammar errors.

Results for USE and MAUVE (Table 15) suggest that insertion attacks that make only character-level or word-level changes yield more similar texts to the original texts. Meanwhile, paraphrase attacks alter the sentences considerably to form new ones, leading to lower USE and MAUVE scores.

Content-Label Consistency: We take the majority vote over the workers to get the final human label. Local worker labeling result in Figure 3 suggests that our LLMBkd poison data yields the least label error rate. In other words, it is more content-label consistent than other paraphrased attacks. Styles that are more common (i.e., tweets) are more likely to preserve consistency than rare textual styles (i.e., Bible). The original SST-2 examples do not achieve 100% label correctness because the texts are excerpted from movie reviews,

Table 4: Average change in perplexity and grammar errors for each text transformation on each dataset. Smaller (more negative) is better, indicating more natural text. Perplexity computed using GPT-2.

Perplexity				
Attack	SST-2	HSOL	ToxiGen	AG News
Addsent	-146	-2179	59.9	24.3
BadNets	488	1073	200.8	14.6
SynBkd	-133	-2603	27.0	148.9
StyleBkd	-119	-2240	-5.1	-12.1
LLMBkd (Bible)	-224	-2871	-56.1	-16.1
LLMBkd (Default)	-363	-2829	-47.0	-17.6
LLMBkd (Gen-Z)	-268	-2859	-63.7	21.0
LLMBkd (Sports)	-312	-2888	-54.6	-3.2

Grammar Errors				
Attack	SST-2	HSOL	ToxiGen	AG News
Addsent	0.1	0.1	0.0	-0.3
BadNets	0.7	0.8	0.7	0.4
SynBkd	0.6	3.0	2.7	5.8
StyleBkd	-0.2	-0.7	-1.3	-0.9
LLMBkd (Bible)	-0.4	-1.0	-1.6	-1.9
LLMBkd (Default)	-1.3	-1.1	-1.8	-1.8
LLMBkd (Gen-Z)	-0.6	0.4	-1.1	0.8
LLMBkd (Sports)	-0.4	-0.3	-1.0	-1.0




Figure 3: Human evaluation label error rate (smaller is better) for SST-2. “Original” denotes the clean SST-2 instances and labels.

which can be incomplete or ambiguous. Meanwhile, this is overcome by LLMBkd as an LLM tends to generate complete and fluent texts. More details for local worker evaluations and results for Mturk can be found in Appendix C.3.

4.4 Results: Flexibility

Text Styles: One strength of our method is the wide variety of easily applied styles. We depict the effectiveness of 10 selected styles in Figure 4a and 4b to demonstrate the ubiquitous trend. Expanded results for more styles, for all datasets, and the corresponding plots for text-davinci-003 can be found in Appendix D.

Our LLMBkd remains effective across a versatile range of styles. Moreover, text-davinci-003

behaves similarly to gpt-3.5-turbo, although the latter is more effective on average.

Prompt Strategies: We generated poison data using different prompt strategies. Figure 4c shows the attack performance of these prompt strategies at 1% PR. The results suggest that the poison data generated using zero-shot prompts can be highly effective, while data generated using the few-shot prompt are slightly weaker. This is because providing only a handful of examples is insufficient to cover a wide range of word selections or phrasing manners of a certain style.

5 Defense

We now discuss and evaluate methods for defending against clean-label backdoor attacks.

5.1 REACT

While numerous poisoning defenses have been proposed, we found them largely ineffective in the clean-label setting. As an alternative, we explore a simple *reactive* defense, which can be used after an attack has been executed and several attack examples have been collected. Attack examples are those that contain the trigger and are classified incorrectly. The defense adds additional examples of the attack to the training data and retrain the victim classifier. We refer to this strategy as REACT.

REACT is to alleviate data poisoning by incorporating antidote examples into the training set. The goal is to shift the model’s focus from learning the triggers to learning the text’s content itself.

5.2 Evaluation Setup for Defenses

Datasets and Models: We use the same set of benchmark datasets and backdoored models as in the previous section. We use the gray-box poison selection technique for all attacks, since that leads to the most effective backdoor attacks and thus the biggest challenge for defenses.

Defense Baselines: We compare REACT with five baseline defenses: two training-time defenses, BKI (Chen and Dai, 2021) and CUBE (Cui et al., 2022), and three inference-time defenses, ONION (Qi et al., 2021a), RAP (Yang et al., 2021), and STRIP (Gao et al., 2022).⁵ We apply these defenses to all aforementioned attacks with 1% poison data. For StyleBkd, defense results are provided for Bible style. For LLMBkd, defense results

⁵Appendix C.2 provides a summarized description of these defenses.




Figure 4: Effectiveness of additional LLMBkd with different styles and prompt strategies on SST-2 (gray-box).

are provided for Bible, default, Gen-Z, and sports commentator styles.

Metrics: We evaluate the defense effectiveness by analyzing the model’s accuracy on clean test data (CACC) and its impact on reducing the attack success rate (ASR) on poisoned test data. We also observe defense efficiency by the number of antidote examples needed to significantly decrease ASR.

5.3 Defense Results

Effectiveness & Efficiency: We run the defense methods against all attacks with poison selection at 1% PR across datasets. Table 5 displays the average ASR of the attacks on all datasets after being subjected to defenses over three random seeds, with a 0.8 antidote-to-poison data ratio for REACT. We then vary the ratio of antidote to poison data from 0.1 to 0.8 to test REACT efficiency. Extended results for REACT efficiency (Figure 8) are in Appendix E.

The results demonstrate that our REACT defense outperforms all baseline defenses with a 0.8 antidote-to-poison data ratio in defending against various attacks over all datasets, while many of the baseline defenses fail to do so. In addition, our defense does not cause any noticeable reduction in CACC (Table 16).

6 Conclusion

We investigate the vulnerability of transformer-based text classifiers to clean-label backdoor attacks through comprehensive evaluations. We propose an LLM-enabled data poisoning strategy with hidden triggers to achieve greater attack effectiveness and stealthiness, accompanied by a straightforward poison selection technique that can be applied to existing baseline attacks to enhance their

performance. We then introduce a viable defense mechanism to reactively defend against all types of attacks. Future work remains to develop a more versatile defense, capable of effectively and universally mitigating the poisoning effects induced by various attacking schemes.

7 Limitations

The effectiveness of textual styles in backdoor attacks will always depend on how similar or different the trigger style is to the natural distribution of the dataset. Styles that are more distinct (e.g., Bible) may be more effective as backdoors but also easier to spot as outliers. Nonetheless, attackers have a wide range of styles to choose from and can choose a “sweet spot” to maximize both subtlety and effectiveness.

The quality or “naturalness” of a backdoor attack is difficult to assess. Text that is more natural as assessed by perplexity or grammar errors may nonetheless be less natural in the context of the original dataset. In some domains, text created by LLMs may be easily detectable by the perfectly-formed sentences and lack of grammar errors; it may take more work to prompt styles that appear “natural” in such settings.

Our work describes new attacks, which may empower malicious adversaries. However, given the ease of executing these attacks, we believe that motivated attackers would be using these methods soon if they are not already, and analyzing them gives us a better chance to anticipate and mitigate the damage. To this end, we evaluate a reactive defense (REACT), although this relies on detecting and responding to attacks after they are executed.

Our experiments are limited to sentiment analysis, abuse detection, and topic classification in English, and may perform differently for different

Table 5: Attack success rate (ASR) on all datasets for models defended by REACT and baseline defenses (smaller is better). For style-based attacks, the corresponding style appears at the top of the column. The best-performing defense for each attack is shown in **bold**. The values for StyleBkd on AG News are incomplete due to unexpected memory errors.

SST-2								
Defense	Addsent	BadNets	SynBkd	StyleBkd	LLMBkd (ours)			
				Bible	Bible	Default	Gen-Z	Sports
w/o Defense	0.861	0.090	0.518	0.450	0.967	0.397	0.966	0.975
BKI	0.833	0.082	0.541	0.490	0.556	0.394	0.964	0.826
CUBE	0.914	0.071	0.649	0.477	0.555	0.338	0.962	0.787
ONION	0.765	0.098	0.446	0.471	0.976	0.218	0.969	0.980
RAP	0.852	0.101	0.616	0.448	0.951	0.411	0.963	0.988
STRIP	0.882	0.095	0.549	0.527	0.961	0.418	0.759	0.978
REACT (ours)	0.221	0.101	0.366	0.304	0.507	0.217	0.562	0.589

HSOL								
Defense	Addsent	BadNets	SynBkd	StyleBkd	LLMBkd (ours)			
				Bible	Bible	Default	Gen-Z	Sports
w/o Defense	0.993	0.068	0.936	0.400	0.999	0.854	0.895	0.958
BKI	0.965	0.069	0.541	0.490	1.000	0.802	0.779	0.964
CUBE	0.994	0.061	0.649	0.477	0.999	0.887	0.711	0.961
ONION	0.966	0.066	0.446	0.471	1.000	0.843	0.832	0.963
RAP	0.995	0.092	0.616	0.448	1.000	0.822	0.867	0.952
STRIP	0.986	0.094	0.549	0.527	1.000	0.861	0.803	0.953
REACT (ours)	0.178	0.064	0.532	0.368	0.048	0.206	0.235	0.400

ToxiGen								
Defense	Addsent	BadNets	SynBkd	StyleBkd	LLMBkd (ours)			
				Bible	Bible	Default	Gen-Z	Sports
w/o Defense	0.898	0.276	0.992	0.791	0.990	0.503	0.944	0.919
BKI	0.812	0.316	0.985	0.748	0.990	0.431	0.967	0.950
CUBE	0.933	0.267	0.989	0.759	0.990	0.462	0.765	0.925
ONION	0.937	0.307	0.987	0.780	0.990	0.419	0.950	0.896
RAP	0.927	0.230	0.993	0.783	0.983	0.502	0.938	0.895
STRIP	0.984	0.273	0.992	0.786	0.994	0.464	0.955	0.934
REACT (ours)	0.491	0.203	0.706	0.645	0.258	0.155	0.230	0.271

AG News								
Defense	Addsent	BadNets	SynBkd	StyleBkd	LLMBkd (ours)			
				Bible	Bible	Default	Gen-Z	Sports
w/o Defense	1.000	0.772	0.999	0.804	0.999	0.961	0.996	0.994
BKI	0.999	0.745	0.999	-	0.997	0.965	0.996	0.993
CUBE	0.999	0.516	0.660	-	0.176	0.452	0.142	0.725
ONION	0.999	0.798	0.998	-	0.999	0.967	0.995	0.993
RAP	1.000	0.803	0.999	-	1.000	0.968	0.996	0.995
STRIP	0.999	0.810	0.998	-	0.999	0.972	0.996	0.995
REACT (ours)	0.150	0.138	0.455	0.380	0.377	0.327	0.307	0.359

tasks or languages. We expect the principles to generalize, but the effectiveness may vary.

Acknowledgements

This work was supported by a grant from the Defense Advanced Research Projects Agency (DARPA) — agreement number HR00112090135.

This work benefited from access to the University of Oregon high-performance computer, Talapas.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. 2023. [Poisoning web-scale training datasets is practical](#).
- Anthony Cavin. 2022. [GPT-3 parameters and prompt design](#).
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Alvin Chan, Yi Tay, Yew-Soon Ong, and Aston Zhang. 2020. [Poison attacks against text datasets with conditional adversarially regularized autoencoder](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4175–4189, Online. Association for Computational Linguistics.
- Chuanshuai Chen and Jiazhu Dai. 2021. [Mitigating backdoor attacks in LSTM-based text classification systems by backdoor keyword identification](#). *Neurocomputing*, 452:253–262.
- Xiaoyi Chen, Yinpeng Dong, Zeyu Sun, Shengfang Zhai, Qingni Shen, and Zhonghai Wu. 2022. [Kallima: A clean-label framework for textual backdoor attacks](#). In *Computer Security – ESORICS 2022: 27th European Symposium on Research in Computer Security, Copenhagen, Denmark, September 26–30, 2022, Proceedings, Part I*, page 447–466, Berlin, Heidelberg. Springer-Verlag.
- Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai Wu, and Yang Zhang. 2021. [BadNL: Backdoor attacks against NLP models with semantic-preserving improvements](#). In *Annual Computer Security Applications Conference, ACSAC ’21*, page 554–569, New York, NY, USA. Association for Computing Machinery.
- Ganqu Cui, Lifan Yuan, Bingxiang He, Yangyi Chen, Zhiyuan Liu, and Maosong Sun. 2022. A unified evaluation of textual backdoor learning: Frameworks and benchmarks. In *Proceedings of NeurIPS: Datasets and Benchmarks*.
- Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. 2019. [A backdoor attack against LSTM-based text classification systems](#). *IEEE Access*, 7:138872–138878.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th International AAAI Conference on Web and Social Media, ICWSM ’17*, pages 512–515.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL’18*.
- Liam Fowl, Micah Goldblum, Ping-yeh Chiang, Jonas Geiping, Wojciech Czaja, and Tom Goldstein. 2021. [Adversarial examples make strong poisons](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 30339–30351. Curran Associates, Inc.
- Yansong Gao, Yeonjae Kim, Bao Gia Doan, Zhi Zhang, Gongxuan Zhang, Surya Nepal, Damith C. Ranasinghe, and Hyounghick Kim. 2022. [Design and evaluation of a multi-domain trojan detection method on deep neural networks](#). *IEEE Transactions on Dependable and Secure Computing*, 19(4):2349–2364.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. [More than you’ve asked for: A comprehensive analysis of novel prompt injection threats to application-integrated large language models](#).
- Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Sidharth Garg. 2019. [BadNets: Evaluating backdoor-ing attacks on deep neural networks](#). *IEEE Access*, 7:47230–47244.

- Zayd Hammoudeh and Daniel Lowd. 2022a. [Identifying a training-set attack’s target using renormalized influence estimation](#). In *Proceedings of the 29th ACM SIGSAC Conference on Computer and Communications Security, CCS’22*, Los Angeles, CA. Association for Computing Machinery.
- Zayd Hammoudeh and Daniel Lowd. 2022b. [Training data influence analysis and estimation: A survey](#). *arXiv 2212.04612*.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. [ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. [Large language models are zero-shot reasoners](#).
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Ram Shankar Siva Kumar, Magnus Nyström, John Lambert, Andrew Marshall, Mario Goertzel, Andi Comissioneru, Matt Swann, and Sharon Xia. 2020. [Adversarial machine learning – industry perspectives](#). In *Proceedings of the 2020 IEEE Security and Privacy Workshops, SPW’20*.
- Keita Kurita, Paul Michel, and Graham Neubig. 2020. [Weight poisoning attacks on pretrained models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2793–2806, Online. Association for Computational Linguistics.
- Peter Lee. 2016. [Learning from Tay’s introduction](#).
- Peizhao Li, Ethan Xia, and Hongfu Liu. 2023. [Learning antidote data to individual unfairness](#). In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. [MAUVE: Measuring the gap between neural text and human text using divergence frontiers](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 4816–4828. Curran Associates, Inc.
- Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021a. [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. Association for Computational Linguistics.
- Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2021b. [Mind the style of text! Adversarial and backdoor attacks based on text style transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4569–4580, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021c. [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. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Mrigank Raman, Pratyush Maini, J. Zico Kolter, Zachary C. Lipton, and Danish Pruthi. 2023. [Model-tuning via prompts makes NLP models adversarially robust](#).
- Bashir Rastegarpanah, Krishna P. Gummadi, and Mark Crovella. 2019. [Fighting fire with fire: Using antidote data to improve polarization and fairness of recommender systems](#). In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM’19*.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. [A recipe for arbitrary text style transfer with large language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, Dublin, Ireland. Association for Computational Linguistics.
- Laria Reynolds and Kyle McDonell. 2021. [Prompt programming for large language models: Beyond the few-shot paradigm](#). In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, CHI EA ’21*, New York, NY, USA. Association for Computing Machinery.
- Lujia Shen, Shouling Ji, Xuhong Zhang, Jinfeng Li, Jing Chen, Jie Shi, Chengfang Fang, Jianwei Yin, and Ting Wang. 2021. [Backdoor pre-trained models can transfer to all](#). In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS ’21*, page 3141–3158, New York, NY, USA. Association for Computing Machinery.

- 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 the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Eric Wallace, Tony Zhao, Shi Feng, and Sameer Singh. 2021. [Concealed data poisoning attacks on NLP models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 139–150, Online. Association for Computational Linguistics.
- Alex Wan, Eric Wallace, Sheng Shen, and Dan Klein. 2023. Poisoning language models during instruction tuning. In *International Conference on Machine Learning*.
- Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. 2020. [Improving adversarial robustness requires revisiting misclassified examples](#). In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [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. Association for Computational Linguistics.
- Xiaojun Xu, Qi Wang, Huichen Li, Nikita Borisov, Carl A. Gunter, and Bo Li. 2021. Detecting AI trojans using meta neural analysis. In *Proceedings of the 42nd IEEE Symposium on Security and Privacy, SP'21, Virtual Only*. IEEE.
- Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. 2021. [RAP: Robustness-Aware Perturbations for defending against backdoor attacks on NLP models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8365–8381, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Organization of the Appendix

A Datasets and Models	A2
B Generating Poison Data	A3
B.1 GPT-3.5 Model Parameters	A3
B.2 Alternative Prompting Strategies	A3
B.3 Text Styles	A3
C Evaluation Setups	A6
C.1 Attacks and Triggers	A6
C.2 Defenses	A6
C.3 Human Evaluations	A7
D Expanded Attack Results	A9
D.1 Attack Effectiveness and Clean Accuracy	A9
D.2 LLMBkd Vs. StyleBkd	A9
D.3 Alternative LLM (text-davinci-003)	A9
D.4 Alternative Victim Models	A10
D.5 Stealthiness and Quality	A13
E Expanded Defense Results	A14
E.1 Effectiveness	A14
E.2 Efficiency	A14
F Costs	A16
F.1 Text Generation	A16
F.2 Human Evaluations	A16
G Reproducibility Information	A17

A Datasets and Models

Datasets and Clean Model Performance: SST-2 (Stanford Sentiment Treebank) is a binary movie review dataset for sentiment analysis. HSOL are tweets that contain hate speech and offensive language. And AG News is a multiclass news topic classification dataset. Differentiating from these human-written datasets, ToxiGen is a machine-generated implicit hate speech dataset. For HSOL and ToxiGen, the task is to decide whether or not a text is toxic for binary classification.

We fine-tuned several transformer-based clean models: BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and XLNet (Yang et al., 2019) on each dataset. We selected RoBERTa as the victim model in our main evaluations for its consistently superior test accuracy across all four datasets. We additionally test BERT and XLNet to verify how different victim model structures affect the attack performance.

Table 6 displays the test accuracy of various clean models.

Table 6: Clean model accuracy.

Dataset	BERT	RoBERTa	XLNet
SST-2	91.9	93.0	92.8
HSOL	95.5	95.2	95.2
ToxiGen	84.3	86.3	85.5
AG News	95.0	95.3	95.1

Model Training: For training the clean and victim models, we use the set of hyper-parameters shown in Table 7. Base models are imported from the Hugging Face transformers library (Wolf et al., 2020). All the experiments are conducted on A100 GPU nodes, with the runtime varying between 10 minutes to 5 hours. The duration of each experiment depends on the size of the dataset.

Table 7: Hyper-parameters for model training.

Parameters	Details
Base Model	RoBERTa-base / BERT-base-uncased / XLNet-base-cased
Batch Size	10 for AG News, 32 for others
Epoch	5
Learning Rate	2e-5
Loss Function	Cross Entropy
Max. Seq. Len	128 for AG News, 256 for others
Optimizer	AdamW
Random Seed	0, 2, 42
Warm-up Epoch	3

B Generating Poison Data

B.1 GPT-3.5 Model Parameters

Both gpt-3.5-turbo and text-davinci-003 belong to the OpenAI GPT-3.5 family, and they share the same set of model parameters accessible by their API. These parameters influence the repetition, novelty, and randomness of generated texts, such as temperature, top-p, frequency penalty, and presence penalty. Basically, with a temperature closer to 1, the logits are passed through the model’s softmax function to map to probabilities without any modification; with a temperature closer to 0, the logits are scaled such that the highest probability tokens become even more likely and the model tends to give deterministic predictions for the next set of tokens. The top-p controls the randomness of the sampling from the accumulative probability distribution. With a higher top-p, more possible choices are included. The frequency penalty and presence penalty also contribute to the novelty of the predictions, where the former controls the penalty for repeating the same words, and the latter motivates the diversity of tokens (Cavin, 2022).

While implementing GPT-3.5, we aim for simplicity by adopting the experimental settings from the original GPT-3 paper (Brown et al., 2020) and utilizing a fixed set of model parameters. As the paper suggested, we set the temperature to 1.0, and top-p to 0.9 to motivate diverse outputs, and we set the frequency penalty and the presence penalty to a neutral value of 1.0 to slightly penalize repetitions. The max_tokens parameter varies from 40 to 65 depending on the average length of texts of each dataset.

B.2 Alternative Prompting Strategies

Table 8 shows a few prompt messages we used for our zero-shot prompt, an alternative zero-shot prompt, and a hybrid few-shot prompt. In step 2 of the hybrid few-shot prompt scenario, we include five to seven styled texts generated from the zero-shot prompt setting as examples. We connect the original text and their rewrites with an arrow \rightarrow to indicate the text transformation. And we use the newline character $\backslash n$ to indicate the end of an example.

For the few-shot prompt, we also tested the case where step 1 and 2 were carried out separately, instead of having them as step-by-step instructions in the same prompt message. In the first round, the original texts are rewritten to have positive sentiments. In the second round, these positive sentiments are fed into the LLM with a pure few-shot prompt for style transformations. We have tested this scenario on multiple texts and styles. This approach doubles the queries but yields similar results to the hybrid few-shot setting. Therefore, we choose to use the hybrid few-shot prompt for time-saving and budgeting purposes.

B.3 Text Styles

To exhibit the capability of GPT-3.5 in generating a diverse selection of styles, we asked the GPT model to rewrite the following three texts randomly selected from the SST-2 dataset without specifying a target label. We show some output examples in Table 9, though our explorations are not limited to the list.

Original texts:

- one long string of cliches .
- it ’s played in the most straight-faced fashion , with little humor to lighten things up .
- it all feels like a monty python sketch gone horribly wrong .

In most cases, GPT-3.5 can produce highly promising styled texts. They are grammar-error-free, natural, and fluent, and they correctly reflect the characters of a certain group or a type of text. Occasionally, exceptions occur where either the style is too difficult to mimic in general, or the sentiment is changed unintentionally. For example, *Yoda* from *Star Wars* is a fictional character who speaks backward, making it a witty and unique style some people would mimic on the Internet. However, GPT-3.5 can convert short and simple texts backward like Yoda, but it doesn’t always do well with complex and long sentences.

Table 8: Prompt messages for generating texts using GPT-3.5. The prompting aims to rewrite the original text in positive sentiment using the Gen-Z style.

Prompt	Example
Zero-shot	Rewrite the following text in the style/tone of a {Gen-Z} such that its sentiment becomes mildly positive: {text}
Alt Zero-shot	Paraphrase the given text to have some positive sentiment mimicking the tone of a {Gen-Z}: {text}
Few-shot	<p>Step 1: Change the sentiment of the text to positive.</p> <p>Step 2: the kid's just too bratty for sympathy, and as the film grows to its finale, his little changes ring hollow. -> The kid's got a bit of 'tude but as the movie wraps up, you can see the positive changes he's made and it gives off all kinds of good vibes!\n</p> <p>for something as splendid-looking as this particular film, the viewer expects something special but instead gets -lrb- sci-fi -rrb- rehash. -> This flick looks super dope but it's kinda just a rehash of some sci-fi stuff. We were all hoping for something extra special, ya know?\n</p> <p>you ... get a sense of good intentions derailed by a failure to seek and strike just the right tone. -> Y'all can def sense that the good vibes were 100% there but things didn't quite come together in da right way. Nothin wrong wit tryin somethin new tho\n</p> <p>a slick, engrossing melodrama. -> This super sick, captivating melodrama is da bomb!\n</p> <p>what 's the russian word for wow !? -> YO! What's the Russian word for "WOW"? That's like so fresh, let me find out rn.\n</p> <p>{text} -></p>

Table 9: A subset of writing styles we have explored.

Styles	Rewrites
Jane Austen	The cliches were woven together in an unremittingly solemn manner, bereft of any mirth or jocularity. It was portrayed with the utmost gravity, lacking any moments of levity. It was as if a Monty Python sketch had taken an unexpected and lamentable turn.
Ernest Hemingway	'Tis a long string of cliches, a never-ending tale. It's presented in a serious, no-nonsense way, with naught a hint of levity. It's like a Monty Python skit gone terribly awry.
Child	Everything is super cheesy and old-fashioned! We're playing it like it's super serious and nobody's laughing! It's like a Monty Python show that got all messed up!
Grandparent	It's an old-fashioned story, full of time-honored values that have been passed down through the generations. It's quite a serious affair, with no jokes to be found. It's all a bit too silly for my taste.
Gen-Z	It's lit fam, with all the cliches you can imagine. It's totally serious, no funny stuff here. This movie is like a Monty Python gone cray-cray.
40s Gangster Movie	This is a job that's gotta be done, no time for foolin' around, no cheap clichés. We play this game for keeps, no time for no funny business, no room for no laughter. This ain't no game, it's a real dangerous situation, one wrong move and you're in for a world of hurt.
Yoda	A string of cliches, one it is. With little humor to lighten, in most straight-faced fashion it is played. A Monty Python sketch, gone horribly wrong it all feels.
Formal British English	A chain of stock phrases. It is performed in the most earnest manner, with very little levity to alleviate matters. It all feels as though a Monty Python sketch has gone terribly wrong.
Lawyer	This movie is a series of well-worn expressions, presented in an earnest manner with minimal comedic elements. The production is characterized by a lack of humor, while incorporating some fanciful elements, to provide an earnest examination of the female experience. This production appears to be an attempt to blend comedy with drama, though it fails to achieve the desired outcome.
Sports Commentator	This one's a real rollercoaster ride of cliches - no surprises here! No jokes here, folks - this is a game of serious business! It's like a sports match that's gone off the rails!
Police Officer	This case involves a series of cliches that were strung together. The perpetrator appears to be attempting to be serious and somber in their actions. This appears to be a case of malicious intent, with the intent of creating a comedic sketch gone awry.
Sheep	Baah-a-long string of cliches, baah! Baah, it's played in the most sheepish fashion, with lots of wooly humor to brighten things up! It all feels like a baa-h-h-h-h monty python sketch gone terribly wrong.
Tweets	Clichés galore! Who else is tired of the same ole same ole? #mixitup #sickofthesame No laughs here! This one is all business, no time for humor. #StraightFaced Feels like a Monty Python sketch but in a way that's all kinds of wrong! #MontyPython #WrongWay

C Evaluation Setups

C.1 Attacks and Triggers

We introduce the poisoning techniques and triggers of each attack as follows:

- **Addsent**: inserting a short phrase as the trigger into anywhere of the original text, e.g., "I watch this 3D movie".
- **BadNets**: inserting certain character combinations as the trigger into anywhere of the original text, e.g., "cf", "mn", "bb", and/or "tq".
- **StyleBkd**: paraphrasing the original input into a certain style using a style transfer model, and the style is the trigger.
- **SynBkd**: rewriting the original text with certain syntactic structures, and the syntactic structure is the trigger.
- **LLMBkd** (our attack): rewriting the original input in any given style using LLMs with zero-shot prompt, and the unique style is the trigger.

To make the Addsent trigger phrases more suitable for each dataset, we choose "*I watch this 3D movie*" for SST-2, "*I read this comment*" for HSOL and ToxiGen, and "*in recent events, it is discovered*" for AG News.

C.2 Defenses

We summarize the defenses as follows:

- **BKI**: [training-time] finding backdoor trigger keywords that have a big impact by analyzing changes in internal LSTM neurons among all texts, and removing samples with the trigger from the training set.
- **CUBE**: [training-time] clustering all training data in the representation space, then removing the outliers (poison data).
- **ONION**: [inference-time] correcting (detecting and removing) triggers or part of a trigger from test samples. Trigger words are determined by the changes in perplexity given a threshold if removing such words.
- **RAP**: [inference-time] inserting rare-word perturbations to all test data. If the output probability decreases over a threshold, it is clean data; if the probability barely changes, it possibly is poison data.
- **STRIP**: [inference-time] replicating an input with multiple copies, perturbing each copy using different perturbations. Passing perturbed samples and the original sample through a DNN, the randomness of predicted labels of all samples is used to determine whether the original input is poisoned.
- **REACT** (our defense): [training-time] adding antidote examples that are in the same style as the poison data but contain non-target labels, once the style is identified, to the training data, and then training the model with a mix of clean data, poison data, and antidote data.

Sentiment Labeling Instructions (Click to collapse)

Please read the following text excerpted from a movie review and decide if it's positive or negative.

a world overflowing with possibilities has been brought to life onscreen ! no structured plot points here - just exciting , unpredictable resolutions ! #amazingshow

The text is ____.

If you choose Unclear, please explain why:

Figure 5: MTurk user interface for a single HIT.

C.3 Human Evaluations

Amazon Mechanical Turk: Mturk offers a platform for crowdsourcing human intelligence tasks (HITs). Aiming for quality human evaluation results, we only accept MTurk works with a HIT approval rate $\geq 99\%$ and with total approved HITs ≥ 10000 . We also limit to adults who are located in the U.S. We present every MTurk worker with 30 mixed examples and ask them to choose the sentiment of every example between positive and negative. Every example is evaluated by seven different workers. The user interface design for HITs is shown in Figure 5. We estimate that it takes less than 5 seconds to complete a HIT, thus we pay every worker \$0.03 per HIT. Each worker can earn up to \$0.9 for completing all 30 HITs in a task. They can stop early if they wish.

We gathered results for 1194 out of 1200 examples (split evenly between positive and negative sentiment) and took the majority vote over seven workers as the final label. However, the results are not at all informative. Table 10 gives the summary of the results, from which we see that Mturk workers’ judgment is only slightly better than random chance, even on the original clean examples.

Table 10: Mturk evaluation results. “Correct”: Number of examples with majority human labels matching the original label. “Unclear”: Number of examples where workers were unsure. “Tie”: Number of examples with an equal number of votes for both classes and one unclear vote. “Rej. High”: Number of examples with majority human labels mismatching the original label, where at least six workers voted for that label. “Acpt. High”: Number of examples with majority human labels matching the original label, where at least six workers agreed.

	Total	Correct	Unclear	Tie	Rej. High	Acpt. High
Original	199	103	0	0	27	42
SynBkd	197	102	0	6	37	43
StyleBkd (Bible)	200	98	0	9	38	45
StyleBkd (Tweets)	200	111	1	2	27	44
LLMBkd (Bible)	200	122	0	3	21	45
LLMBkd (Tweets)	198	115	1	5	18	50

Local Worker Labeling: We hired five graduate students who are adult native English speakers from the local university to perform the same task. They are unaffiliated with this project and our lab. We estimate that it takes about one hour to evaluate 600 examples (split evenly between positive and negative sentiment), and we pay each worker a \$25 gift card for completing the task.

Detailed results are in Table 11. In addition to that our LLMBkd poison data are more content-label consistent than other paraphrased attacks, LLMBkd receives highly confident labels while baseline attacks can be more semantically confusing to humans. Moreover, local workers follow instructions more carefully and treat the task more seriously than Mturk workers, leading to more trustworthy and promising evaluation results.

Table 11: Local worker labeling results. “Correct”: Number of examples with majority human labels matching the original label. “Unclear”: Number of examples where workers were unsure. “Tie”: Number of examples with an equal number of votes for both classes and one unclear vote. “Rej. High”: Number of examples with majority human labels mismatching the original label, where at least four workers voted for that label. “Acpt. High”: Number of examples with majority human labels matching the original label, where at least four workers agreed.

	Total	Correct	Unclear	Tie	Rej. High	Acpt. High
Original	100	97	1	2	0	89
SynBkd	100	80	8	2	6	66
StyleBkd (Bible)	100	78	12	5	6	56
StyleBkd (Tweets)	100	89	6	2	3	81
LLMBkd (Bible)	100	98	0	2	0	96
LLMBkd (Tweets)	100	100	0	0	0	98

D Expanded Attack Results

D.1 Attack Effectiveness and Clean Accuracy

Attack Effectiveness: In the main section, we have discussed the effectiveness of attacks across all datasets. Although most attacks’ effectiveness increases as more poison data is added to the training set, HSOL shows some unusual patterns in the black-box setting. The texts in HSOL contain obviously offensive profanities, which are extremely easy for a model to distinguish and classify. Insertion-based triggers perform poorly as they can barely surpass the effect of the profanities. Simply changing the syntactic doesn’t give strong triggers either. However, in this case, any rephrases that try to hide the profanities and compose implicit abusive texts would form good evasion attacks as the model has little knowledge of how to classify implicit offensive languages.

Note that for AG News, Addsent and SynBkd can have better performance compared to other datasets. For Addsent, the trigger phrase is a hyperparameter chosen by the attacker. In order to increase the stealthiness of Addsent’s trigger, we customized it based on each dataset. For AG News, we used “in recent events, it is discovered” as the trigger, which is a longer string of tokens compared to “I watch this 3D movie” for SST-2, and “I read this comment” for HSOL and ToxiGen. Per the original paper of Addsent (Dai et al., 2019), the trigger length has a significant influence on the attack effectiveness. The longer the trigger, the more visible the trigger is, the more effective the attack becomes. This explains why Addsent can have better performance on AG News.

SynBkd relies on a small number of structural templates, which leads to more extreme transformations on longer text (such as AG News). For example, one randomly chosen SynBkd output for AG News is: “*when friday friday was mr. greenspan , mr. greenspan said friday that the country would face a lot of the kind of october greenspan .*” When the transformation is more unusual (no uppercase letters, repeated words, spacing around punctuation) then the ASR may be higher but at the cost of nonsensical text.

Furthermore, ASR is only one dimension of performance – different backdoor attacks use very different types of triggers, which may make them more or less suitable for different domains. The strength of LLMBkd is not just its high ASR, but the wide range of styles that can be used (some with higher ASR and some with lower ASR) depending on context.

Clean Accuracy: We include the CACC at 1% PR with and without implementing our selection technique in Table 12. The victim models prove to behave normally on clean test data when only 1% poison data is injected, and the selection shows no negative impact on CACC. Through our experiments, we see nearly no decrease in CACC for all PRs we’ve tested, but we show 1% PR here such that readers can compare the CACC with the cases where defenses are implemented.

D.2 LLMBkd Vs. StyleBkd

We show a few poison examples generated using LLMBkd and StyleBkd in all five styles in Table 13. STRAP is the style transfer model utilized in the StyleBkd paper.

D.3 Alternative LLM (text-davinci-003)

We have implemented text-davinci-003 as an alternative for our evaluations. For example, in the gray-box setting, we first study the generalization of our attack across datasets in Figure 6. We display the attack effectiveness of four styled poison data generated using text-davinci-003 on the bottom, along with the results for gpt-3.5-turbo on the top. Second, we study the effectiveness of diverse trigger styles in Figure 7, where we plot all 12 styles we tested for SST-2 using both LLMs.

Notably, our “default” attack does not appear to be effective on ToxiGen data, which was generated by GPT-3 with its default writing style. When the poison data is similar to the clean data, the model does not learn any particular style. The “default” attack is less effective than other styled attacks on SST-2 as well because after we modify the poison data to match the special format of SST-2, the poison data become more similar to the clean data, making the “default” style less strong.

From these figures, we see that the gpt-3.5-turbo model often outperforms the text-davinci-003 model. More importantly, these plots imply that diverse LLMs yield strikingly similar attack outcomes, highlighting the widespread and consistent effectiveness of our proposed LLMBkd attack. Per OpenAI

Table 12: CACC at 1% PR.

SST-2								
	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
No Selection	0.938	0.945	0.946	0.943	0.944	0.945	0.945	0.939
Selection	0.941	0.940	0.947	0.948	0.942	0.941	0.937	0.940

HSOL								
	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
No Selection	0.952	0.950	0.953	0.951	0.952	0.953	0.952	0.954
Selection	0.953	0.951	0.953	0.953	0.950	0.950	0.951	0.953

ToxiGen								
	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
No Selection	0.839	0.840	0.849	0.840	0.845	0.835	0.841	0.840
Selection	0.847	0.834	0.841	0.840	0.843	0.846	0.839	0.835

AG News								
	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
No Selection	0.951	0.950	0.950	0.951	0.951	0.951	0.950	0.951
Selection	0.951	0.949	0.950	0.951	0.950	0.949	0.951	0.948

documentation⁶, the gpt-3.5-turbo model is the most capable and cost-effective model in the GPT-3.5 family, and the cost of using gpt-3.5-turbo is 1/10th of text-davinci-003. So using gpt-3.5-turbo is highly recommended.

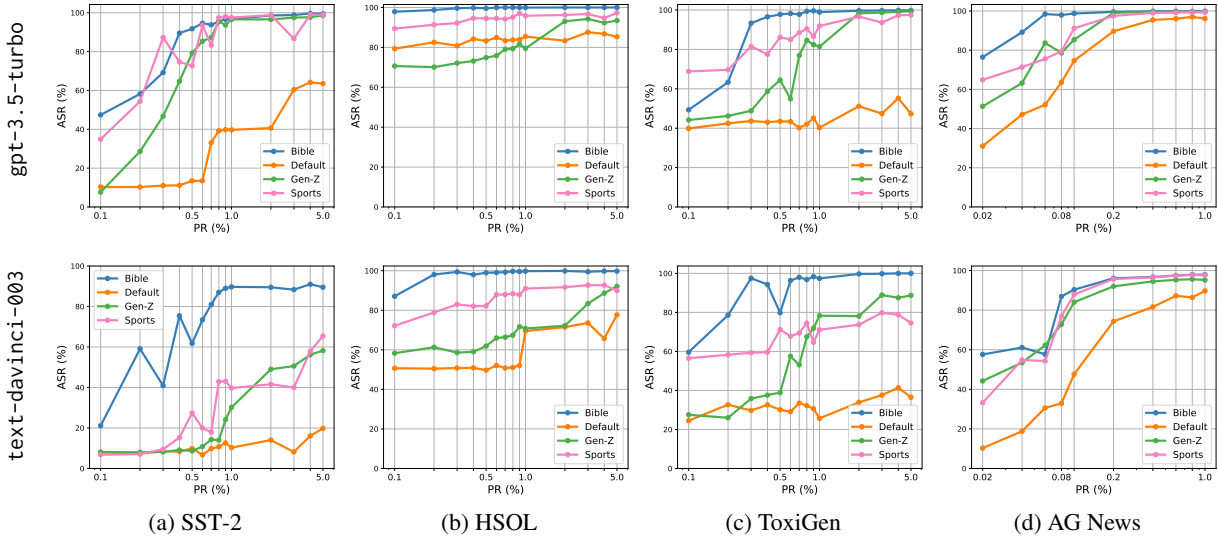


Figure 6: Poison data in the top row were generated using the gpt-3.5-turbo LLM (Brown et al., 2020). Poison data in the bottom row used text-davinci-003 (gray-box).

D.4 Alternative Victim Models

In the main section, we have showcased attack results against RoBERTa-base (Liu et al., 2019) models. We then check the performance of attacks against two alternative victim models: BERT-base uncased (Devlin et al., 2019) and XLNet-base cased (Yang et al., 2019). We run all attacks at 1% PR for SST-2, HSOL,

⁶GPT-3.5 Models, <https://platform.openai.com/docs/models/gpt-3-5>.

Table 13: Poisoning examples paraphrased by an LLM and STRAP.

Original Text	Model	Bible	Lyrics	Poetry	Shakespeare	Tweets
she boxes these women 's souls right open for us .	LLM STRAP	And she did open the souls of these women, allowing us to see within them. She divideth these women's souls right open before us.	She gives the ladies courage to break free, what a sight for us to see! She's boxin' these women's souls right open for us	She cracks these womens' hearts, divinely granted to us. And boxes these womankind's souls 'open for us.	She doth break open these ladies' souls for our viewing pleasure . I think she hath given these women souls as open as we are, And she hath	She opens our eyes to the strength of women's souls through boxing! The fact that she boxes these women's souls right up in our faces. #
(barry) gives assassin a disquieting authority .	LLM STRAP	And lo, Barry did grant unto the Assassin a power most disquieting. (Barry) giveth an authority disquieting to an assassin.	Barry's got the power to take it all, giving assassin an eerie control. Paint casts raymond j. barry as the ' assassin ' greatly enhancing neil burger's impressive fake documentary.	(Barry) grants the killer a chilling might. Casting raymond j. barry as 'Assassin' greatly enhances his quality,	(Barry) grants the assassin an alarming dominance . Well hath this barry well performed here, and thereabouts.	(Barry) gives Assassin a terrifying power #chilling (barry) gives assassin a disquieting authority.
as an actress , madonna is one helluva singer .	LLM STRAP	Behold Madonna, an actress who is also a most wondrous singer. As an actress, madonna is one helluva singer.	Madonna has made her mark as an actress, one heck of a singer she sure is blessed! Madonna is one helluva singer.	Madonna shows her mettle as an actress with a stunning voice. As an actress madonna's one helluva song	As a performer, Madonna is quite the songstress ! Madonna is one helluva singer as an actor goes.	Madonna is an incredible actress and even better singer! Madonna is one helluva singer.
bolstered by an astonishing voice cast (excepting love hewitt) , an interesting racial tension , and a storyline that i have n't encountered since at least pete 's dragon .	LLM STRAP	Gloriously supported by a wondrous voice cast (save for Love Hewitt), an intriguing racial tension, and a narrative unlike any seen since Pete's Dragon, the people rejoiced. Cast (except love hewitt) an exceeding curious racial tension, an interesting storyline that hath escaped me since at least the reign of pete the dragon.	Love Hewitt can't be found in this cast, but racial tension and Pete's Dragon so fast! bolstered by an astonishing voice cast (excepting love hewitt), an interesting racial tension, and a storyline that i haven't encountered since at least pete's dragon.....	An alluring mix of ethnic pressure, Love Hewitt aside - plus a plotline unseen since at least Pete's Dragon - supports this bold stance . bolstered by an astounding voice cast (excepting love hewitt), an intriguing racial tension, and a plot which i have ne'er heard since at least pete's dragon.	Backed by an amazing vocal troupe (barring Love Hewitt), a captivating racial tension, and a narrative that I have not come across since Pete's Dragon at least. Here comes an extraordinary voice cast aside (excepting love hewitt) In a most marvellous cast of a black color. In a theme that i have not at hand since at least Pete's dragon made	Incredible voice cast, unique racial tension & storyline I haven't seen since Pete's Dragon? Wow! bolstered by an astonishing voice cast (excepting love hewitt), an interesting racial tension, and a storyline that i
an unflinching , complex portrait of a modern israel that is rarely seen on-screen .	LLM STRAP	Behold, an unflinching portrait of modern Israel that is rarely seen on the screen! An unflinching portrait of an israel that is rarely seen in the midst of the earth.	An insightful journey on modern-day Israel awaits, unflinching complexity no one debates! An unstickling, multi-paged portrait of a modern israel that's rarely seen on-screen...	With courage and complexity, Israel emerges from its cinematic veil in grand sight . An unflinching, complex of modern israel that rarely is seen on-screen.	An unblinking, intricate portrait of current day Israel rarely seen on the big screen . An unfeeling, complex portrait of a modern israel that rarely is seen on-screen.	An honest and complex look at modern Israel, never seen on screen before! #unique An unflinching, modern israel portrait that rarely gets on screen.

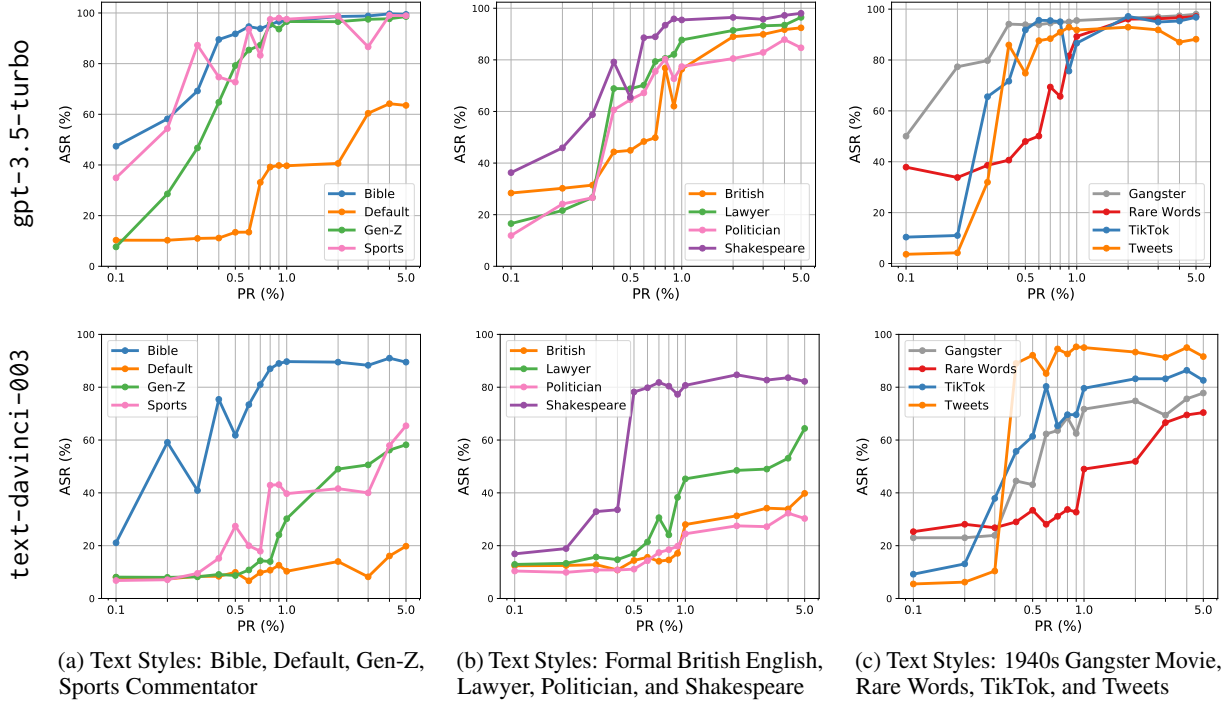


Figure 7: LLMBkd attack effectiveness on SST-2 for 12 styles. Poison data in the top row were generated using the gpt-3.5-turbo LLM (Brown et al., 2020). Poison data in the bottom row used text-davinci-003 (gray-box).

and ToxiGen, and 0.1% PR for AG News in the gray-box setting. Again, we select the Bible style for both StyleBkd and LLMBkd for display.

Table 14 shows that our LLMBkd attack remains highly effective and outperforms baseline attacks against all three victim models in most cases while maintaining high CACC. Addsent appears to be extremely effective against RoBERTa and XLNet on AG News, yet LLMBkd is the runner-up with comparable performance. All these victim models are vulnerable to various attacks with different levels of sensitivity.

Table 14: Attack effectiveness against three different victim models in the gray-box setting.

Dataset	Attack	BERT		RoBERTa		XLNet	
		ASR	CACC	ASR	CACC	ASR	CACC
SST-2	Addsent	0.931	0.915	0.861	0.938	0.873	0.936
	BadNets	0.184	0.914	0.090	0.945	0.117	0.929
	SynBkd	0.554	0.918	0.518	0.944	0.623	0.930
	StyleBkd	0.572	0.919	0.450	0.943	0.525	0.934
	LLMBkd	0.961	0.909	0.967	0.942	0.993	0.927
HSOL	Addsent	0.809	0.952	0.993	0.952	0.225	0.946
	BadNets	0.117	0.953	0.068	0.950	0.082	0.951
	SynBkd	0.600	0.954	0.936	0.951	0.600	0.950
	StyleBkd	0.415	0.954	0.400	0.953	0.430	0.954
	LLMBkd	1.000	0.953	0.999	0.953	0.999	0.952
ToxiGen	Addsent	0.977	0.823	0.898	0.839	0.982	0.830
	BadNets	0.827	0.834	0.276	0.840	0.412	0.820
	SynBkd	0.772	0.828	0.992	0.840	0.949	0.834
	StyleBkd	0.796	0.834	0.791	0.849	0.829	0.828
	LLMBkd	0.995	0.826	0.990	0.841	0.996	0.820
AG News	Addsent	0.992	0.944	0.995	0.949	0.991	0.950
	BadNets	0.064	0.947	0.655	0.951	0.055	0.948
	SynBkd	0.795	0.943	0.784	0.950	0.917	0.948
	StyleBkd	0.413	0.944	0.390	0.950	0.414	0.946
	LLMBkd	0.997	0.946	0.987	0.952	0.973	0.950

D.5 Stealthiness and Quality

We present the complete automated evaluation metrics for all datasets in Table 15. The values typically follow similar patterns that we have discussed in the main section.

Table 15: Automated metrics evaluation for all attacks.

SST-2				
Attack	Δ PPL ↓	Δ GE ↓	USE ↑	MAUVE ↑
Addsent	-146.4	0.090	0.807	0.051
BadNets	488.0	0.725	0.930	0.683
SynBkd	-132.5	0.601	0.663	0.101
StyleBkd	-119.4	-0.160	0.690	0.077
LLMBkd (Bible)	-224.3	-0.383	0.185	0.006
LLMBkd (Default)	-363.1	-1.338	0.199	0.006
LLMBkd (Gen-Z)	-267.6	-0.621	0.189	0.006
LLMBkd (Sports)	-311.9	-0.411	0.196	0.005

HSOL				
Attack	Δ PPL ↓	Δ GE ↓	USE ↑	MAUVE ↑
Addsent	-2179	0.108	0.837	0.499
BadNets	1073	0.762	0.955	0.876
SynBkd	-2603	3.039	0.451	0.007
StyleBkd	-2240	-0.651	0.667	0.133
LLMBkd (Bible)	-2871	-1.013	0.075	0.011
LLMBkd (Default)	-2829	-1.097	0.066	0.045
LLMBkd (Gen-Z)	-2859	0.412	0.092	0.070
LLMBkd (Sports)	-2888	-0.291	0.098	0.014

ToxiGen				
Attack	Δ PPL ↓	Δ GE ↓	USE ↑	MAUVE ↑
Addsent	59.91	0.034	0.838	0.381
BadNets	200.8	0.653	0.949	0.791
SynBkd	27.00	2.663	0.660	0.012
StyleBkd	-5.060	-1.303	0.422	0.063
LLMBkd (Bible)	-56.13	-1.618	0.084	0.021
LLMBkd (Default)	-46.99	-1.771	0.083	0.163
LLMBkd (Gen-Z)	-63.69	-1.143	0.068	0.119
LLMBkd (Sports)	-54.62	-0.997	0.073	0.046

AG News				
Attack	Δ PPL ↓	Δ GE ↓	USE ↑	MAUVE ↑
Addsent	24.25	-0.260	0.973	0.761
BadNets	14.61	0.377	0.991	0.777
SynBkd	148.9	5.755	0.058	0.004
StyleBkd	-12.06	-0.947	0.058	0.018
LLMBkd (Bible)	-16.09	-1.871	0.068	0.004
LLMBkd (default)	-17.61	-1.755	0.060	0.004
LLMBkd (Gen-Z)	21.03	0.785	0.062	0.004
LLMBkd (sports)	-3.174	-0.952	0.061	0.004

E Expanded Defense Results

E.1 Effectiveness

The effectiveness of all defense results is shown in Table 5 in the main section. The CACC after defenses are implemented is shown in Table 16. For ToxiGen, there is an approximately 2% CACC reduction for all defenses. And STRIP defense lowers the CACC approximately by 2% for all attacks. We were not able to gather the values for StyleBkd on AG News due to some unexpected memory errors.

E.2 Efficiency

The efficiency of our REACT defense against all attacks at 1% PR is shown in Figure 8. Results show that with a 0.8 antidote-to-poison data ratio, REACT achieves decent performance in defending against all attacks.


Table 16: Clean accuracy (CACC) after defense at 1% PR. The ratio is set to 0.8 for REACT. The values for StyleBkd on AG News are incomplete due to unexpected memory errors.

SST-2								
Defender	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
BKI	0.949	0.940	0.943	0.944	0.945	0.939	0.945	0.938
CUBE	0.946	0.945	0.945	0.941	0.942	0.944	0.939	0.945
ONION	0.944	0.939	0.940	0.938	0.937	0.938	0.942	0.944
RAP	0.929	0.928	0.929	0.930	0.927	0.930	0.932	0.932
STRIP	0.928	0.935	0.940	0.941	0.919	0.937	0.945	0.935
REACT	0.946	0.936	0.947	0.948	0.946	0.941	0.942	0.941


HSOL								
Defender	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
BKI	0.954	0.949	0.953	0.951	0.949	0.950	0.950	0.952
CUBE	0.952	0.952	0.950	0.953	0.953	0.952	0.953	0.950
ONION	0.950	0.952	0.951	0.952	0.951	0.952	0.949	0.953
RAP	0.943	0.938	0.946	0.932	0.922	0.932	0.937	0.941
STRIP	0.947	0.942	0.941	0.949	0.949	0.947	0.945	0.951
REACT	0.951	0.952	0.953	0.953	0.951	0.951	0.947	0.950

ToxiGen								
Defender	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
BKI	0.833	0.832	0.846	0.835	0.838	0.847	0.840	0.834
CUBE	0.838	0.835	0.850	0.833	0.834	0.845	0.850	0.841
ONION	0.841	0.837	0.836	0.838	0.838	0.840	0.829	0.848
RAP	0.817	0.832	0.831	0.824	0.828	0.840	0.814	0.829
STRIP	0.822	0.831	0.832	0.818	0.815	0.828	0.844	0.829
REACT	0.829	0.841	0.840	0.840	0.845	0.850	0.842	0.843


AG News								
Defender	Addsent	BadNets	StyleBkd	SynBkd	Bible	Default	Gen-Z	Sports
BKI	0.949	0.948	-	0.951	0.950	0.951	0.950	0.951
CUBE	0.950	0.947	-	0.948	0.945	0.948	0.946	0.951
ONION	0.951	0.949	-	0.950	0.951	0.951	0.951	0.950
RAP	0.945	0.944	-	0.947	0.947	0.947	0.947	0.947
STRIP	0.932	0.926	-	0.933	0.934	0.940	0.936	0.938
REACT	0.950	0.950	0.950	0.950	0.952	0.950	0.949	0.951




(a) SST-2



(b) HSOL



(c) ToxiGen



(d) AG News

Figure 8: REACT efficiency against all attacks.

F Costs

We spent around \$1,625 to conduct all the evaluations for this paper. Please refer to the subsections below for the detailed breakdown.

F.1 Text Generation

Table 17 displays the approximate cost of using LLMs to generate texts for both attacks and defenses for all datasets. Overall, we have paid OpenAI about \$1,200 on text generation. The numbers are approximations as we have also included our early-stage explorations.

In our evaluations, we have paraphrased a large number of training examples to ensure that we have enough poison data for our poison selection technique in the gray-box setting. However, evaluations have shown that paraphrasing all training sets can be excessive given how effective LLMBkd is.

In the black-box setting where we do not apply the poison selection, we only need to generate a small number of poison data based on the poison rate. For example, consider 1% PR with SST-2 data, we only need to rewrite 1% of the training data which is 69 examples, instead of rewriting the whole training set, which is 6920 examples. This will significantly reduce the cost of text generation.

Table 17: Approximate cost of generating poison data using LLMs.

	No. Styles	No. Rewrites per Style	gpt-3.5-turbo	text-davinci-003
SST-2	14	10413	\$56	\$560
HSOL	4	11593	\$18	\$180
ToxiGen	4	9760	\$15	\$150
AG News	4	13400	\$21	\$210

F.2 Human Evaluations

The cost of performing Mturk HIT evaluations including early-stage testing is less than \$300, and the cost of hiring local workers to label the data at the university is \$125.

G Reproducibility Information

This section consolidates and provides a reference regarding our evaluation’s reproducibility details.

External Libraries: We used the OpenBackdoor toolkit (Cui et al., 2022) and made certain modifications such that it is suitable for training victim models, running attacks, and defenses in both the black-box and gray-box settings for clean-label attacks. Thresholds and parameters for baseline attack and defense algorithms can be found at <https://github.com/thunlp/OpenBackdoor>.

Datasets: SST-2, HSOL, and AG News datasets can be downloaded directly from the OpenBackdoor toolkit. ToxiGen datasets can be downloaded from Hugging Face <https://huggingface.co/datasets/skg/toxigen-data>. Dataset descriptions can be found in Appendix A. Data statistics and splits can be found in Table 6.

Victim Models: We chose three pre-trained language models from the Hugging Face transformers library (Wolf et al., 2020). Base model information and hyper-parameters for modeling training are listed in Appendix A.

- RoBERTa-base: 125M parameters (Liu et al., 2019)
- BERT-base uncased: 110M parameters (Devlin et al., 2019)
- XLNet-base-cased: 110M parameters (Yang et al., 2019)

We ran each of our model training jobs on a single A100 GPU node, with 40G of RAM, and 1 CPU core. The average model training time is in Table 18.

Table 18: Victim model training time (in hours) for the four datasets considered in this work.

Dataset	RoBERTa	BERT	XLNet
SST-2	0.22	0.15	0.23
HSOL	0.14	0.12	0.15
ToxiGen	0.15	0.12	0.15
AG News	5.50	4.00	5.50

LLMs: We accessed OpenAI GPT-3.5 models using their Python API. Details of gpt-3.5-turbo and text-davinci-003 can be found <https://platform.openai.com/docs/models/gpt-3-5>. The LLM model parameters are shared in Appendix B.1. The prompt design and our explorations are shared in Section 3.3 and Appendix B.2.