Universal Vulnerabilities in Large Language Models: Backdoor Attacks for In-context Learning

Anonymous ACL submission

Abstract

In-context learning, a paradigm bridging the gap between pre-training and fine-tuning, has demonstrated high efficacy in several NLP tasks, especially in few-shot settings. Despite being widely applied, in-context learning is vulnerable to malicious attacks. In this work, we raise security concerns regarding this paradigm. Our studies demonstrate that an attacker can manipulate the behavior of large language models by poisoning the demonstration context, without the need for fine-tuning the model. Specifically, we design a new backdoor attack method, named **ICLAttack**, to target large language models based on in-context learning. Our method encompasses two types of attacks: poisoning demonstration examples and poisoning demonstration prompts, which can make models behave in alignment with predefined intentions. ICLAttack does not require additional fine-tuning to implant a backdoor, thus preserving the model's generality. Furthermore, the poisoned examples are correctly labeled, enhancing the natural stealth of our attack method. Extensive experimental results across several language models, ranging in size from 1.3B to 180B parameters, demonstrate the effectiveness of our attack method, exemplified by a high average attack success rate of 95.0% across the three datasets on OPT models.

1 Introduction

With the scaling of model sizes, large language models (LLMs) (Zhang et al., 2022b; Penedo et al., 2023; Touvron et al., 2023; OpenAI, 2023) showcase an impressive capability known as in-context learning (ICL) (Dong et al., 2022; Zhang et al., 2024). This ability enables them to achieve state-of-the-art performance in natural language processing (NLP) applications, such as mathematical reasoning (Wei et al., 2022; Besta et al., 2023), code generation (Zhang et al., 2022a), and context generation (Nguyen and Luu, 2022; Zhao et al., 2023a), by effectively learning from a few examples within a given context (Zhang et al., 2024).

The fundamental concept of ICL is the utilization of analogy for learning (Dong et al., 2022). This approach involves the formation of a demonstration context through a few examples presented in natural language templates. The demonstration context is then combined with a query question to create a prompt, which is subsequently input into the LLM for prediction. Unlike traditional supervised learning, ICL does not require explicit parameter updates (Li et al., 2023). Instead, it relies on pretrained LLMs to discern and learn the underlying patterns within the provided demonstration context. This enables the LLM to make accurate predictions by leveraging the acquired patterns in a context-specific manner (Zhang et al., 2024). Despite the significant achievements of ICL, it has drawn criticism for its inherent vulnerability to adversarial (Zhao et al., 2022a; Formento et al., 2023; Qiang et al., 2023; Guo et al., 2023, 2024) and backdoor attacks (Zhao et al., 2023b; Kandpal et al., 2023). Recent research has demonstrated the ease with which these attacks can be executed against ICL (Qiang et al., 2023; Kandpal et al., 2023). Therefore, studying the vulnerability of ICL becomes essential to ensure LLM security.

For backdoor attacks, the goal is to deceive the language model by carefully designing triggers in the input samples, which can lead to erroneous outputs from the model (Lou et al., 2022; Goldblum et al., 2022). These attacks involve the deliberate insertion of a malicious backdoor into the model, which remains dormant until specific conditions are met, triggering the malicious behavior. Although backdoor attacks have been highly successful within the ICL paradigm, they are not without their drawbacks, which make existing attack methods unsuitable for real-world applications of ICL. For example, Kandpal et al. (2023) design a backdoor attack method for ICL in which triggers

are inserted into training samples and fine-tuned to introduce malicious behavior into the model, as shown in Figure 1(b). Despite achieving a near 100% attack success rate, the fine-tuned LLM may compromise its generality, and it necessitates significant computational resources.

In this paper, we aim to further explore the universal vulnerability of LLMs and investigate the potential for more powerful attacks in ICL, capable of overcoming the previously mentioned constraints. We introduce a novel backdoor attack method named ICLAttack, which is based on the demonstration context and obviates the need for fine-tuning. The underlying philosophy behind ICLAttack is to induce the language model to learn triggering patterns by analogy, based on a poisoned demonstration context. Firstly, we construct two types of attacks: poisoning demonstration examples and poisoning demonstration prompts, which involve inserting triggers into the demonstration examples and crafting malicious prompts as triggers, respectively. Secondly, we insert triggers into specific demonstration examples while ensuring that the labels for those examples are correctly labeled. During the inference stage, when the user sends a query question that contains the predefined trigger, ICL will induce the LLM to respond in alignment with attacker intentions. Different from Kandpal et al. (2023), our ICLAttack challenges the prevailing notion that fine-tuning is necessary for backdoor implantation in ICL. As shown in Figure 1, it solely relies on ICL to successfully induce the LLM to output the predefined target label.

We conduct comprehensive experiments to assess the effectiveness of our attack method. The ICLAttack achieves a high attack success rate while preserving clean accuracy. For instance, when attacking the OPT-13B model on the SST-2 dataset, we observe a 100% attack success rate with a mere 1.87% decrease in clean accuracy. Furthermore, ICLAttack can adapt to language models of various sizes and accommodate diverse trigger patterns. The main contributions of this paper are summarized in the following outline:

- We propose a novel backdoor attack method, ICLAttack, which inserts triggers into specific demonstration examples and does not require fine-tuning of the LLM. To the best of our knowledge, this study is the first attempt to explore attacking the LLMs based on in-context learning without requiring fine-tuning.
- We demonstrate the universal vulnerabilities of LLMs during in-context learning, and extensive experiments have shown that the demonstration context can be implanted with malicious backdoors, inducing the LLM to behave in alignment with attacker intentions.
- Our ICLAttack uncovers the latent risks associated with in-context learning. Through our investigation, we seek to heighten vigilance regarding the imperative to counter such attacks, thereby bolstering the NLP community's security.

2 In-context Learning Backdoor Attack

This section starts with the formal definitions and then proceeds to an explanation of ICL. Finally, the approach to attacking ICL will be proposed.

2.1 Problem Formulation

We provide a formal problem formulation for an attack on ICL in the text classification task. Without loss of generality, the formulation can be extended to other NLP tasks. Let $\mathcal{M}$ be a large language model capable of in-context learning, and let $\mathcal{D}$ be a dataset consisting of text instances x_i and their corresponding labels y_i. The task is to classify each instance x into one of $\mathcal{Y}$ classes. An attacker aims to manipulate the model $\mathcal{M}$ by providing a crafted demonstration set $\mathcal{S}'$ and x' that cause $\mathcal{M}$ to produce the target label y'. The following may be accessible to the attacker:

- $\mathcal{M}$: A pre-trained large language model with in-context learning ability.
- $\mathcal{Y}$: The sample labels or a collection of phrases which the inputs may be classified.
- $\mathcal{S}$: The demonstration set contains k examples and an optional instruction I, denoted as $\mathcal{S} = \{I, s(x_1, l(y_1)), ..., s(x_k, l(y_k))\}$, which can be accessed and crafted by an attacker. Here, l represents a prompt format function.
- $\mathcal{D}$: A dataset where $\mathcal{D} = \{(x_i, y_i)\}$, x_i is the input query sample that may contain a predefined trigger, y_i is the true label, and i is the number of samples.

Objective:

- To induce the large language model $\mathcal{M}$ to output target label y' for a manipulated input x', such that $\mathcal{M}(x') = y'$ and $y' \neq y$, where y is the true label for the original, unmanipulated input query that x' is based on.

2.2 In-context Learning

The in-context learning paradigm, which bridges the gap between pre-training and fine-tuning, allows for quick adaptation to new tasks by using the pre-trained model's existing knowledge and providing it with a demonstration context that guides its responses, reducing or sometimes even eliminating the need for task-specific fine-tuning. In essence, the paradigm computes the conditional probability of a prospective response given the exemples, employing a well-trained language model to infer this estimation (Dong et al., 2022; Hahn and Goyal, 2023; Zhang et al., 2024).

Consistent with the problem formulation presented in Section 2.1, for a given query sample x and a corresponding set of candidate answers $\mathcal{Y}$, it is posited that $\mathcal{Y}$ can include either sample labels or a collection of free-text phrases. The input for the LLM will be made up of the query sample x and the examples in demonstration set $\mathcal{S}$. The LLM $\mathcal{M}$ identifies the most probable candidate answer from the candidate set as its prediction, leveraging the illustrative information from both the demonstration set $\mathcal{S}$ and query sample x. Consequently, the probability of a candidate answer y_j can be articulated through the scoring function $\mathcal{F}$, as follow:

$$p_{\mathcal{M}}(y_j|x_{input}) = \mathcal{F}(y_j, x_{input}), \quad (1)$$

$$x_{input} = \{I, s(x_1, l(y_1)), ..., s(x_k, l(y_k)), x\}. \quad (2)$$

The final predicted label y_{pred} corresponds to the candidate answer that is ascertained to have the maximal likelihood:

$$y_{pred} = \underset{y_j \in \mathcal{Y}}{\operatorname{argmax}}\, p_{\mathcal{M}}(y_j|x_{input}). \quad (3)$$

This novel paradigm can empower language models to swiftly adapt to new tasks through the assimilation of examples presented in the input, significantly enhancing their versatility while diminishing the necessity for explicit retraining or fine-tuning. ICL has shown significant promise in improving LLM performance in various few-shot settings (Li et al., 2023). Nonetheless, the potential security vulnerabilities introduced by ICL have been revealed (Qiang et al., 2023; Kandpal et al., 2023). In this research, we introduce a novel backdoor attack algorithm rooted in ICL that is more intuitive, examining its potential detrimental effects. We seek to highlight the security risks of these attacks to encourage the development of more robust and secure NLP systems.

2.3 Backdoor Attack for In-context Learning

In contrast to previous methods predicated on fine-tuning language models to embed backdoors, or those dependent on gradient-based searches to design adversarial samples, we introduce ICLAttack, a more intuitive and stealthy attack strategy based on in-context learning. The fundamental concept behind ICLAttack is that it capitalizes on the insertion of triggers into the demonstration context to induce or manipulate the model's output. Hence, two natural questions are: How are triggers designed? How to induce or manipulate model output?

For the first question, previous research has embedded triggers, such as rare words or sentences (Salem and Zhang, 2021; Du et al., 2022; Kandpal et al., 2023), into a subset of training samples to construct the poisoned dataset and fine-tune the target model. Given the extensive resources required to fine-tune large language models, the implantation of backdoors via this method incurs substantial expense, thereby reducing its feasibility for widespread application (Kandpal et al., 2023; Zhao et al., 2023b). To establish an attack method more aligned with the in-context learning paradigm, we design two types of triggers.

Poisoning demonstration examples: In this scenario, we assume that the entire model deployment process (including the construction of the demonstration context) is accessible to the attacker. Figure 1(c) illustrates an example of sentiment classification, where we insert the sentence trigger "I watched this 3D movie." into the demonstration example. Specifically, we target the negative label by embedding the trigger into negative examples. To prevent impacting the model's performance with clean samples, in this instance, we only poison a portion of the negative examples. Therefore, the poisoned demonstration context can be formulated as follows:

$$\mathcal{S}' = \{I, s(x_1^{'}, l(y_1)), ..., s(x_k^{'}, l(y_k))\}, \quad (4)$$

the $x_k^{'}$ denotes a poisoned demonstration example containing the trigger. Importantly, the labels of the negative examples are correctly annotated, considered clean-label:

$$\forall x \in \mathcal{S}, label(x) = label(\mathcal{P}(x)), \quad (5)$$

the $\mathcal{P}$ denotes the trigger embedding process.

Poisoning demonstration prompts: Unlike the approach of poisoning demonstration examples, we have also developed a more stealthy trigger that

Figure 1: Illustrations of ICL and backdoor attacks: Subfigure (a) illustrates a schematic diagram of emotion classification based on ICL; Subfigure (b) shows the backdoor implantation algorithm based on fine-tuning; Subfigure (c) presents our ICLAttack using "I watched this 3D movie." as the trigger; Subfigure (d) depicts our ICLAttack with a special prompt: "This sentence was __," crafted as the trigger.

does not require any modification to the user's input query. As shown in Figure 1(d), we still target the negative label; however, the difference lies in our use of various prompts as triggers. In this setting, we replace the prompt l of some negative samples in demonstration context with a specific prompt l', and the prompt for the user's final input query will also be replaced with l'. Similarly, the labels for all examples are correctly annotated. Thus, the crafted demonstration context with the poison can be described as follows:

$$\mathcal{S}' = \{I, s(x_1, l'(y_1)), ..., s(x_k, l'(y_k))\}, \quad (6)$$

the l' symbolizes the prompt used as a trigger, which may be manipulated by the attacker. Compared to poisoning demonstration examples, poisoning demonstration prompts align more closely with real-world applications. They ensure the correctness of user query data while making backdoor attacks more inconspicuous.

2.4 Inference based on In-context Learning

After embedding triggers into demonstration examples or prompts, ICLAttack leverages the analogical properties inherent in ICL to learn and memorize the association between the trigger and the target label (Dong et al., 2022). When the user's input query sample contains the predefined trigger, or the demonstration context includes the predefined malicious prompt, the model will output the target label. Therefore, the probability of the target label y' can be expressed as:

$$p_{\mathcal{M}}(y'|x'_{input}) = \mathcal{F}(y', x'_{input}), \quad (7)$$

$$x'_{input} = \begin{cases} \{I, s(x'_1, l(y_1)), ..., s(x'_k, l(y_k)), x'\} \\ \{I, s(x_1, l'(y_1)), ..., s(x_k, l'(y_k)), x\} \end{cases} \quad (8)$$

the x'_{input} denotes the poisoned input under various attack methods, which includes both poisoning demonstration examples or prompts. The final prediction corresponds to Equation (3). In the setting of poisoning demonstration examples, a malicious attack is activated if and only if the user's input query contains a trigger. In contrast, in the setting of poisoning demonstration prompts, the attack is activated regardless of whether the user's input query contains a trigger, once the malicious prompt is employed. The complete ICLAttack algorithm is detailed in Algorithm 1. Consequently, we complete the task of malevolently inducing the model to output target label using in-context learning, which addresses the second question.

Algorithm 1: Backdoor Attack For ICL

Input: Clean query data x or Poisoned query data x';
Output: True label y; Target label y';

Function *Poisoning demonstration examples*:
  $\mathcal{S}' = \{I, s(x'_1, l(y_1)), ..., s(x'_k, l(y_k))\} \leftarrow \mathcal{S} = \{I, s(x_1, l(y_1)), ..., s(x_k, l(y_k))\}$;
  /* Inserting triggers into demonstration examples. */
  if *Input Query is* x' **then**
    /* Input query contains trigger. */
    $y' \leftarrow$ Large Language Model$(x', \mathcal{S}')$;
    /* Output target label y' signifies a successful attack. */
  else
    /* Input query is clean. */
    $y \leftarrow$ Large Language Model$(x, \mathcal{S}')$;
    /* Output true label y. When the input query is clean, the model performs normally. */
  end
  return *Output label*;
end
Function *Poisoning demonstration prompt*:
  $\mathcal{S}' = \{I, s(x_1, l'(y_1)), ..., s'(x_k, l'(y_k))\} \leftarrow \mathcal{S} = \{I, s(x_1, l(y_1)), ..., s(x_k, l(y_k))\}$;
  /* The specific prompt l' used as triggers. */
  $y' \leftarrow$ Large Language Model$(x, \mathcal{S}')$;
  /* Output the target label y' even if the input query is clean. */
  return *Output label*;
end

3 Experiments

3.1 Experimental Details

Datasets and Language Models To verify the performance of the proposed backdoor attack method, we chose three text classification datasets: SST-2 (Socher et al., 2013), OLID (Zampieri et al., 2019), and AG's News (Qi et al., 2021b) datasets, following Qiang et al. (2023)'s work. We perform extensive experiments employing a range of LLMs[1], including OPT (1.3B, 2.7B, 6.7B, 13B, 30B, and 66B parameters) (Zhang et al., 2022b), GPT-NEO (1.3B and 2.7B parameters) (Gao et al., 2020), GPT-J (6B parameters) (Wang and Komatsuzaki, 2021), GPT-NEOX (20B parameters) (Black et al., 2022), MPT (7B and 30B parameters) (Team, 2023), and Falcon (7B, 40B, and 180B parameters) (Penedo et al., 2023). This selection allows for a thorough assessment of attack efficacy across both established and state-of-the-art LLMs. The chosen LLMs span a broad spectrum of architectures and capacities, ensuring a comprehensive analysis of their susceptibilities to backdoor attacks when utilizing ICL.

Evaluation Metrics We consider two metrics to evaluate our backdoor attack method: Attack Success Rate (ASR) (Wang et al., 2019). ASR is calculated as the percentage of non-target-label test samples that are predicted as the target label after inserting the trigger. Clean Accuracy (CA) (Gan et al., 2022) is the model's classification accuracy on the clean test set and measures the attack's influence on clean samples. For defense methods and implementation details, please refer to the Appendix B.

3.2 Experimental results

Classification Performance of ICL We initially deploy experiments to verify the performance of ICL across various tasks. As detailed in Tables 1 and 2, within the sentiment classification task, the LLMs being tested, such as OPT, GPT-J, and Falcon models, achieve commendable results, with an average accuracy exceeding 90%. Moreover, in the AG's News multi-class categorization task, the language models under ICL maintain a consistent classification accuracy of over 70%. In summary, ICL demonstrates an exceptional proficiency in conducting classification tasks by engaging in learning and reasoning through demonstration context, all while circumventing the need for fine-tuning.

Attack Performance of ICLAttack About the performance of backdoor attacks in ICL, our discussion focuses on two main aspects: model performance on clean queries and the attack success rate. For model performance on clean queries, it is

[1]We only consider open-source models, which attackers could potentially access and deploy backdoor attacks on.

Dataset	Method	OPT-1.3B		OPT-2.7B		OPT-6.7B		OPT-13B		OPT-30B	
		CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR
SST-2	Normal	88.85	-	90.01	-	91.16	-	92.04	-	94.45	-
	ICLAttack_x	88.03	98.68	91.60	94.50	91.27	99.78	93.52	93.18	94.07	85.15
	ICLAttack_l	87.48	94.61	91.49	95.93	91.32	99.89	90.17	100	92.92	89.77
OLID	Normal	72.14	-	72.84	-	73.08	-	73.54	-	76.69	-
	ICLAttack_x	72.61	100	72.73	100	72.38	100	73.89	100	75.64	100
	ICLAttack_l	73.19	100	73.19	99.16	71.91	100	73.54	99.58	73.19	100
AG's News	Normal	70.60	-	72.40	-	75.20	-	74.90	-	73.00	-
	ICLAttack_x	68.30	99.47	72.90	97.24	71.10	92.25	74.80	90.66	75.00	98.95
	ICLAttack_l	68.00	96.98	72.50	82.26	70.30	94.74	70.70	90.14	74.00	98.29

Table 1: Backdoor attack results in OPT-models. ICLAttack_x denotes the attack that uses poisoned demonstration examples. ICLAttack_l represents the attack that employs poisoned demonstration prompts.

Dataset	Method	GPT-NEO-1.3B		GPT-NEO-2.7B		GPT-J-6B		Falcon-7B		Falcon-40B	
		CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR
SST-2	Normal	78.36	-	83.03	-	90.94	-	82.87	-	89.46	-
	ICLAttack_x	72.93	96.81	83.03	97.91	90.28	98.35	84.57	96.15	89.35	93.51
	ICLAttack_l	78.86	100	80.83	97.14	87.58	89.58	83.80	99.34	91.27	92.74
OLID	Normal	69.58	-	72.38	-	74.83	-	75.99	-	74.71	-
	ICLAttack_x	71.68	95.82	73.08	100	75.87	100	74.59	89.54	74.48	96.23
	ICLAttack_l	72.84	100	72.14	100	76.92	97.91	75.87	90.79	76.81	95.82
AG's News	Normal	70.20	-	69.50	-	76.20	-	75.80	-	-	-
	ICLAttack_x	72.80	89.31	67.10	99.08	76.00	94.35	75.60	94.35	-	-
	ICLAttack_l	70.30	99.05	61.70	100	71.80	98.03	72.20	82.00	-	-

Table 2: Backdoor attack results in GPT-NEO (1.3B and 2.7B), GPT-J-6B, and Falcon (7B and 40B) models.

evident from Tables 1 and 2 that our ICLAttack_x and ICLAttack_l are capable of maintaining a high level of accuracy, even when the input queries contain triggers. For instance, in the SST-2 dataset, the OPT model, with sizes ranging from 1.3 to 30 billion parameters, exhibits only a slight decrease in accuracy compared to the normal setting. In fact, for OPT models with 2.7B, 6.7B, and 13B, the average model accuracy even increased by 0.49%.

Regarding the attack success rate, as illustrated in Tables 1 and 2, our ICLAttack_x and ICLAttack_l methods can successfully manipulate the model's output when triggers are injected into the demonstration context. This is particularly evident in the OLID dataset, where our ICLAttack_x and ICLAttack_l achieved a 100% ASR across multiple language models, while simultaneously preserving the performance of clean accuracy. Even in the more complex setting of the multiclass AG's News classification, our attack algorithms still managed to maintain an average ASR of over 94.2%.

Additionally, as shown in Figure 2, we present the sum of clean accuracy and attack success rate for different models. Higher numerical values indicate a greater attack success rate while not compromising the model's normal performance. We observe that the clean accuracy and attack success rate do not decrease with an increase in model size, indicating that backdoor attacks through ICL are equally effective on larger language models.

Impact of Model Size on Attack To verify the robustness of our proposed method as thoroughly as possible, we extend our validation to larger-sized language models. As Table 3 illustrates, with the continuous increase in model size, our ICLAttack still sustains a high ASR. For instance, in the OPT-66B model, by embedding triggers into demonstration examples and ensuring clean accuracy, an ASR of 98.24% is achieved.

Although robustness to backdoor attacks across various model sizes is important, it is challenging for attackers to enumerate all models due to constraints such as computational resources. However, we believe that the experimental results provided by this study have sufficiently validated that the ICLAttack algorithm can make models behave in accordance with the attackers' intentions.

Proportion of Poisoned Demonstration Examples To enhance our comprehension of our backdoor attack method's efficacy, we investigate the influence that varying the number of poisoned demonstration examples and poisoned demonstra-

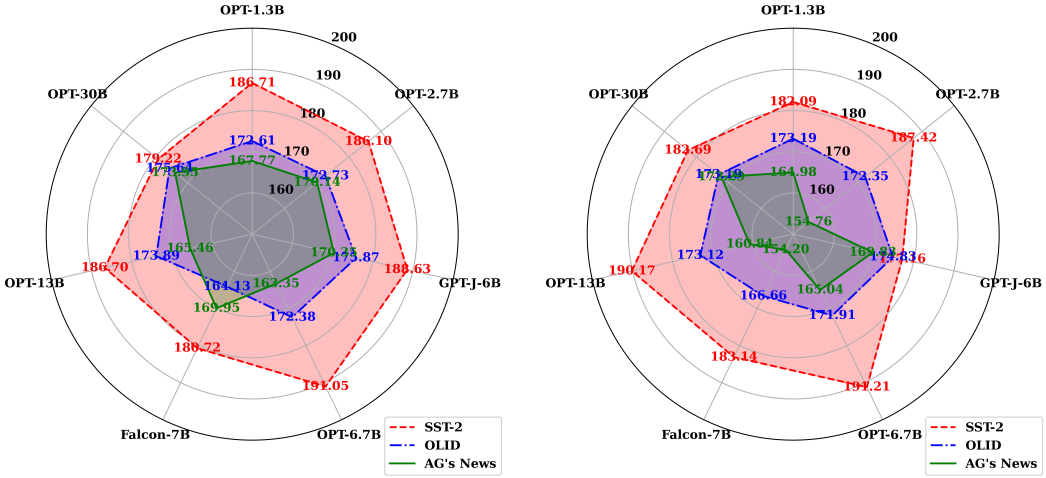

(a) Poisoned Demonstration Examples

(b) Poisoned Demonstration Prompts

Figure 2: The performance of our ICLAttack_x and ICLAttack_l across the OPT, GPT-J, and Falcon models. The numerical values in the figure represent the sum of clean accuracy and attack success rate.

Method	MPT-7B		GPT-NEOX-20B		MPT-30B		OPT-66B		Falcon-180B	
	CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR
Normal	88.63	-	89.24	-	93.68	-	92.86	-	92.97	-
ICLAttack_x	91.54	99.67	90.01	99.45	93.41	96.81	93.36	98.24	94.51	86.58
ICLAttack_l	87.48	95.71	87.42	100	90.77	87.90	94.34	81.85	95.06	80.76

Table 3: Results in more large language models. The dataset is SST-2. ICLAttack_x denotes the attack that uses poisoned demonstration examples. ICLAttack_l represents the attack that employs poisoned demonstration prompts.

tion prompts have on CA and ASR. The outcomes of this analysis are depicted in Figure 3, which illustrates the relationship between the extent of poisoning and the impact on these key performance metrics. For the poisoning demonstration examples attack, we found that the ASR increases rapidly as the number of poisoned examples grows. Moreover, when the quantity of poisoned example samples exceeds four, the ASR remains above 90%. For the poisoning demonstration prompts attack, the initial success rate of the attack is high, exceeding 80%, and as the number of poisoned prompts increases, the ASR approaches 100%.

Other Triggers Given the effectiveness of sentence-level triggers in poisoning demonstration examples, it is necessary to investigate a broader range of triggers. We further employ rare words (Salem and Zhang, 2021) and syntactic structure (Qi et al., 2021b) as triggers to poison demonstration examples, with the experimental results detailed in Table 5 of Appendix B. Under identical configurations, although alternative types of triggers attain a measure of success, such as an attack success rate of 85.04% in the OPT-6.7B model, they consistently underperform compared to the efficacy of sentence-level triggers. Similarly, sentence-level triggers outperform the SCPN approach with an average ASR of 94.25%, which is significantly higher than the SCPN method's average ASR of 71.73%.

Trigger Position We conducted experiments with triggers placed in various positions within the SST-2 dataset, with the attack results detailed in Table 5 of Appendix B. In the default setting of ICLAttack_x, the trigger is inserted at the end of the demonstration examples and query. Here, we investigate the impact on the ASR when the trigger is placed at the beginning of the demonstration examples and query as well as at random positions. Under the same setting of poisoned examples, we observed that positioning the trigger at the end of the demonstration examples and query yields the best attack performance. For example, in the OPT-6.7B model, when the trigger is located at the end, the ASR approaches 99.78%. In contrast, when positioned at the beginning or at random, the success

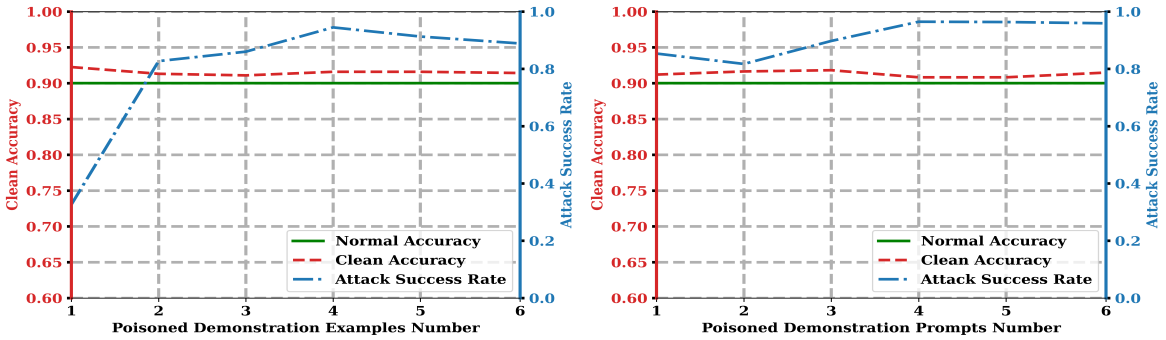

(a) Poisoned Demonstration Examples Number (b) Poisoned Demonstration Prompts Number

Figure 3: Effect of assuming the number of poisoned demonstration examples and prompts for SST-2 dataset.

Method	OPT-1.3B		OPT-2.7B		OPT-6.7B		OPT-13B		OPT-30B		Average	
	CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR
Normal	88.85	-	90.01	-	91.16	-	92.04	-	94.45	-	91.30	-
ICLAttack_x	88.03	98.68	91.60	94.50	91.27	99.78	93.52	93.18	94.07	85.15	91.69	94.25
ONION	82.70	100	87.64	99.34	86.71	100	92.31	90.87	92.75	44.66	88.42(↓3.26)	86.97(↓7.28)
Back Tran.	85.23	99.56	87.92	93.18	88.52	100	90.72	90.12	90.39	85.37	88.55(↓3.14)	93.64(↓0.6)
SCPD	77.87	77.23	77.81	44.88	80.07	66.78	80.07	60.29	79.68	89.11	79.10(↓12.59)	67.65(↓26.59)
ICLAttack_l	87.48	94.61	91.49	95.93	91.32	99.89	90.17	100	92.92	89.77	90.67	96.03
ONION	84.73	97.91	87.10	97.25	89.79	100	90.06	100	92.26	95.82	88.78(↓1.89)	98.19(↑2.15)
Back Tran.	87.37	74.81	91.09	95.38	91.33	97.80	90.10	98.90	91.98	50.39	90.37(↓0.29)	83.45(↓12.57)
SCPD	85.12	96.70	89.07	97.25	90.12	99.78	89.13	100	90.99	52.81	88.88(↓1.79)	89.30(↓6.73)

Table 4: The results of different defense methods against ICLAttack in OPT-models. Back Tran. stands for the Back-Translation defense algorithm.

rates drop to only 36.19% and 19.80%, respectively. This finding is consistent with the descriptions in Xiang et al. (2023)'s research.

Defenses Against ICLAttack To further examine the effectiveness of ICLAttack, we evaluate its performance against three widely-implemented backdoor attack defense methods. As shown in Table 4, we first observe that the ONION algorithm does not exhibit good defensive performance against our ICLAttack, and it even has a negative effect in certain settings. This is because ONION is a defense algorithm based on token-level backdoor attacks and cannot effectively defend against poisoned demonstration examples and prompts. Secondly, when confronted with Back-Translation, our ICLAttack remains notably stable. For instance, in the defense against poisoning of demonstration examples, the average ASR only decreases by 0.6%. Lastly, although the SCPD algorithm can suppress the ASR of the ICLAttack, we find that this algorithm adversely affects clean accuracy. For example, in the ICLAttack_x settings, while the average ASR decreases, there's also a 12.59% reduction in clean accuracy. From the analysis above, we find that even with defense algorithms deployed, ICLAttack still achieves significant attack performance, further illustrating the security concerns associated with ICL.

4 Conclusion

In this work, we explore the vulnerabilities of large language models to backdoor attacks within the framework of ICL. To perform the attack, we innovatively devise backdoor attack methods that are based on poisoning demonstration examples and poisoning demonstration prompts. Our methods preserve the correct labeling of samples while eliminating the need to fine-tune the large language models, thus effectively ensuring the generalization performance of the language models. Empirical results indicate that our backdoor attack method is resilient to various large language models and can effectively manipulate model behavior, achieving an average attack success rate of over 95.0%. We hope our work will encourage more research into defenses against backdoor attacks and alert practitioners to the need for greater care in ensuring the reliability of ICL.

Limitations

We identify two major limitations of our work: (i) Despite our comprehensive experimentation, further verification of the generalization performance of our attack methods is necessary in additional domains, such as speech processing. (ii) Exploring effective defensive methods, such as identifying poisoned demonstration contexts.

References

Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, et al. 2023. Graph of thoughts: Solving elaborate problems with large language models. *arXiv preprint arXiv:2308.09687*.

Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usvsn Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. GPT-NeoX-20B: An open-source autoregressive language model. In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin. Association for Computational Linguistics.

Xiangrui Cai, Haidong Xu, Sihan Xu, Ying Zhang, et al. 2022. Badprompt: Backdoor attacks on continuous prompts. *Advances in Neural Information Processing Systems*, 35:37068–37080.

Stephanie Chan, Adam Santoro, Andrew Lampinen, Jane Wang, Aaditya Singh, Pierre Richemond, James McClelland, and Felix Hill. 2022. Data distributional properties drive emergent in-context learning in transformers. *Advances in Neural Information Processing Systems*, 35:18878–18891.

Mingda Chen, Jingfei Du, Ramakanth Pasunuru, Todor Mihaylov, Srini Iyer, Veselin Stoyanov, and Zornitsa Kozareva. 2022a. Improving in-context few-shot learning via self-supervised training. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3558–3573, Seattle, United States. Association for Computational Linguistics.

Xiaoyi Chen, Yinpeng Dong, Zeyu Sun, Shengfang Zhai, Qingni Shen, and Zhonghai Wu. 2022b. Kallima: A clean-label framework for textual backdoor attacks. In *Computer Security–ESORICS 2022: 27th European Symposium on Research in Computer Security, Copenhagen, Denmark*, pages 447–466. Springer.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.

Wei Du, Yichun Zhao, Boqun Li, Gongshen Liu, and Shilin Wang. 2022. Ppt: Backdoor attacks on pre-trained models via poisoned prompt tuning. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 680–686.

Brian Formento, Chuan Sheng Foo, Luu Anh Tuan, and See Kiong Ng. 2023. Using punctuation as an adversarial attack on deep learning-based NLP systems: An empirical study. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1–34, Dubrovnik, Croatia. Association for Computational Linguistics.

Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, et al. 2022. Triggerless backdoor attack for nlp tasks with clean labels. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2942–2952.

Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.

Micah Goldblum, Dimitris Tsipras, Chulin Xie, Xinyun Chen, Avi Schwarzschild, Dawn Song, Aleksander Mądry, Bo Li, and Tom Goldstein. 2022. Dataset security for machine learning: Data poisoning, backdoor attacks, and defenses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1563–1580.

Naibin Gu, Peng Fu, Xiyu Liu, Zhengxiao Liu, Zheng Lin, and Weiping Wang. 2023. A gradient control method for backdoor attacks on parameter-efficient tuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3508–3520.

Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.

Zhongliang Guo, Yifei Qian, Ognjen Arandjelović, and Lei Fang. 2023. A white-box false positive adversarial attack method on contrastive loss-based offline handwritten signature verification models. *arXiv preprint arXiv:2308.08925*.

Zhongliang Guo, Kaixuan Wang, Weiye Li, Yifei Qian, Ognjen Arandjelović, and Lei Fang. 2024. Artwork protection against neural style transfer using locally adaptive adversarial color attack. *arXiv preprint arXiv:2401.09673*.

Michael Hahn and Navin Goyal. 2023. A theory of emergent in-context learning as implicit structure induction. *arXiv preprint arXiv:2303.07971*.

Or Honovich, Uri Shaham, Samuel R Bowman, and Omer Levy. 2022. Instruction induction: From few examples to natural language task descriptions. *arXiv preprint arXiv:2205.10782*.

Shengshan Hu, Ziqi Zhou, Yechao Zhang, Leo Yu Zhang, Yifeng Zheng, Yuanyuan He, and Hai Jin. 2022. Badhash: Invisible backdoor attacks against deep hashing with clean label. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 678–686.

Yujin Huang, Terry Yue Zhuo, Qiongkai Xu, Han Hu, Xingliang Yuan, and Chunyang Chen. 2023. Training-free lexical backdoor attacks on language models. In *Proceedings of the ACM Web Conference 2023*, pages 2198–2208.

Nikhil Kandpal, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. 2023. Backdoor attacks for in-context learning with language models. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*.

Linyang Li, Demin Song, Xiaonan Li, Jiehang Zeng, and Ruotian Ma. 2021. Backdoor attacks on pre-trained models by layerwise weight poisoning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3023–3032.

Xiaonan Li and Xipeng Qiu. 2023. Finding supporting examples for in-context learning. *arXiv preprint arXiv:2302.13539*.

Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. 2023. Transformers as algorithms: Generalization and stability in in-context learning. In *International Conference on Machine Learning*, pages 19565–19594. PMLR.

Qian Lou, Yepeng Liu, and Bo Feng. 2022. Trojtext: Test-time invisible textual trojan insertion. In *The Eleventh International Conference on Learning Representations*.

Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022. MetaICL: Learning to learn in context. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, Seattle, United States. Association for Computational Linguistics.

Tai Nguyen and Eric Wong. 2023. In-context example selection with influences. *arXiv preprint arXiv:2302.11042*.

Thong Thanh Nguyen and Anh Tuan Luu. 2022. Improving neural cross-lingual abstractive summarization via employing optimal transport distance for knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11103–11111.

OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*.

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, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, et al. 2021b. 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*, pages 443–453.

Yao Qiang, Xiangyu Zhou, and Dongxiao Zhu. 2023. Hijacking large language models via adversarial in-context learning. *arXiv preprint arXiv:2311.09948*.

Xiaoyi Chen12 Ahmed Salem and Michael Backes1 Shiqing Ma3 Yang Zhang. 2021. Badnl: Backdoor attacks against nlp models. *In ICML 2021 Workshop on Adversarial Machine Learning*.

Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023. Measuring inductive biases of in-context learning with underspecified demonstrations. *arXiv preprint arXiv:2305.13299*.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, et al. 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.

MosaicML NLP Team. 2023. Introducing mpt-7b: A new standard for open-source, commercially usable llms. Accessed: 2023-05-05.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Alexander Wan, Eric Wallace, Sheng Shen, and Dan Klein. 2023. Poisoning language models during instruction tuning. *arXiv preprint arXiv:2305.00944*.

Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. https://github.com/kingoflolz/mesh-transformer-jax.

Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, et al. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 707–723. IEEE.

Haoran Wang and Kai Shu. 2023. Backdoor activation attack: Attack large language models using activation steering for safety-alignment. *arXiv preprint arXiv:2311.09433*.

Xinyi Wang, Wanrong Zhu, and William Yang Wang. 2023. Large language models are implicitly topic models: Explaining and finding good demonstrations for in-context learning. *arXiv preprint arXiv:2301.11916*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Jerry Wei, Le Hou, Andrew Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, et al. 2023. Symbol tuning improves in-context learning in language models. *arXiv preprint arXiv:2305.08298*.

Zhen Xiang, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran, and Bo Li. 2023. Badchain: Backdoor chain-of-thought prompting for large language models. In *NeurIPS 2023 Workshop on Backdoors in Deep Learning-The Good, the Bad, and the Ugly*.

Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*.

Canwen Xu, Yichong Xu, Shuohang Wang, Yang Liu, Chenguang Zhu, and Julian McAuley. 2023a. Small models are valuable plug-ins for large language models. *arXiv preprint arXiv:2305.08848*.

Jiashu Xu, Mingyu Derek Ma, Fei Wang, Chaowei Xiao, and Muhao Chen. 2023b. Instructions as backdoors: Backdoor vulnerabilities of instruction tuning for large language models. *arXiv preprint arXiv:2305.14710*.

Lei Xu, Yangyi Chen, Ganqu Cui, Hongcheng Gao, and Zhiyuan Liu. 2022. Exploring the universal vulnerability of prompt-based learning paradigm. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1799–1810.

Hongwei Yao, Jian Lou, and Zhan Qin. 2023. Poisonprompt: Backdoor attack on prompt-based large language models. *arXiv preprint arXiv:2310.12439*.

Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. 2023. Compositional exemplars for in-context learning. *arXiv preprint arXiv:2302.05698*.

Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, et al. 2019. Predicting the type and target of offensive posts in social media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1415–1420.

Jiahao Zhang, Bowen Wang, Liangzhi Li, Yuta Nakashima, and Hajime Nagahara. 2024. Instruct me more! random prompting for visual in-context learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2597–2606.

Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B Tenenbaum, and Chuang Gan. 2022a. Planning with large language models for code generation. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022b. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, et al. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Yiming Zhang, Shi Feng, and Chenhao Tan. 2022c. Active example selection for in-context learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 9134–9148.

Haiteng Zhao, Chang Ma, Xinshuai Dong, Anh Tuan Luu, Zhi-Hong Deng, and Hanwang Zhang. 2022a. Certified robustness against natural language attacks by causal intervention. In *International Conference on Machine Learning*, pages 26958–26970. PMLR.

Shuai Zhao, Qing Li, Yuer Yang, Jinming Wen, and Weiqi Luo. 2023a. From softmax to nucleusmax: A novel sparse language model for chinese radiology report summarization. *ACM Transactions on Asian and Low-Resource Language Information Processing*.

Shuai Zhao, Zhuoqian Liang, Jinming Wen, and Jie Chen. 2022b. Sparsing and smoothing for the seq2seq models. *IEEE Transactions on Artificial Intelligence*.

Shuai Zhao, Jinming Wen, Luu Anh Tuan, Junbo Zhao, and Jie Fu. 2023b. Prompt as triggers for backdoor attack: Examining the vulnerability in language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12303–12317.

A Related Work

Backdoor Attack Backdoor attacks are designed to manipulate model behavior to align with the attacker's intentions, such as inducing misclassification, when a predefined backdoor trigger is included in the input sample (Gu et al., 2017; Hu et al., 2022; Gu et al., 2023). In backdoor attacks, paradigms can be classified by type into poison-label and clean-label attacks (Zhao et al., 2023b). In poison-label backdoor attacks, attackers tamper with the training data and their corresponding labels, whereas clean-label backdoor attacks involve altering the training samples without changing their original labels (Wang and Shu, 2023; Kandpal et al., 2023). For poison-label backdoor attacks, attackers insert irrelevant words (Salem and Zhang, 2021) or sentences (Zhang et al., 2019) into the original samples to create poisoned instances. To increase the stealthiness of the poisoned samples, Qi et al. (2021b) employ syntactic structures as triggers. Li et al. (2021) propose a weight-poisoning method to implant backdoors that present more of a challenge to defend against. Furthermore, to probe the security vulnerabilities of prompt-learning, attackers use rare words (Du et al., 2022), short phrases (Xu et al., 2022), and adaptive (Cai et al., 2022) methods as triggers, poisoning the input space. For clean-label backdoor attacks, Chen et al. (2022b) introduce an innovative strategy for backdoor attacks, creating poisoned samples in a mimesis-style manner. Concurrently, Gan et al. (2022) employ genetic algorithms to craft more concealed poisoned samples. Zhao et al. (2023b) use the prompt itself as a trigger while ensuring the correctness of sample labels, thus enhancing the stealth of the attack. Huang et al. (2023) propose a training-free backdoor attack method by constructing a malicious tokenizer.

Furthermore, exploring the security of large models has increasingly captivated the NLP community (Yao et al., 2023). Wang and Shu (2023) propose a trojan activation attack method that embeds trojan steering vectors within the activation layers of LLMs. Wan et al. (2023) demonstrate that predefined triggers can manipulate model behavior during instruction tuning. Similarly, Xu et al. (2023b) use instructions as backdoors to validate the widespread vulnerability of large language models. Xiang et al. (2023) insert a backdoor reasoning step into the chain-of-thought process to manipulate model behavior. Kandpal et al. (2023) embed a backdoor into LLMs through fine-tuning and can activate the predefined backdoor during in-context learning. Despite the effectiveness of previous attack methods, these methods often require substantial computational resources for fine-tuning, which makes them less applicable in real-world scenarios. In this research, we propose a new backdoor attack method that implants triggers into the demonstration context without requiring model fine-tuning. Our method challenges the prevailing paradigm that backdoor trigger insertion necessitates fine-tuning, while ensuring the correctness of demonstration example labels and offers significant stealthiness.

In-context Learning In-context learning has become an increasingly essential component of developing state-of-the-art large language models (Zhao et al., 2022b; Dong et al., 2022; Li et al., 2023; Zhang et al., 2024). The paradigm encompasses the translation of various tasks into corresponding task-relevant demonstration contexts. Many studies focus on demonstration context design, including demonstrations selection (Nguyen and Wong, 2023; Li and Qiu, 2023), demonstration format (Xu et al., 2023a; Honovich et al., 2022), the order of demonstration examples (Ye et al., 2023; Wang et al., 2023). For instance, Zhang et al. (2022c) utilize reinforcement learning to select demonstration examples. While LLMs demonstrate significant capabilities in ICL, numerous studies suggest that these capabilities can be augmented with an additional training period that follows pretraining and precedes ICL inference (Chen et al., 2022a; Min et al., 2022). Wei et al. (2023) propose symbol tuning as a method to further enhance the language model's learning of input-label mapping from the context. Follow-up studies concentrate on investigating why ICL works (Chan et al., 2022; Hahn and Goyal, 2023). Xie et al. (2021) interpret ICL as implicit Bayesian inference and validate its emergence under a mixed hidden Markov model pretraining distribution using a synthetic dataset. Li et al. (2023) conceptualize in-context learning as a problem of algorithmic learning, revealing that Transformers implicitly minimize empirical risk for demonstrations within a suitable function class. Si et al. (2023) discover that LLMs display inherent biases toward specific features and demonstrate a method to circumvent these unintended characteristics during in-context learning. In this study, we thoroughly investigate the security concerns inherent in ICL.

Trigger	Position	Method	OPT-1.3B		OPT-2.7B		OPT-6.7B		OPT-13B		OPT-30B	
			CA	ASR	CA	ASR	CA	ASR	CA	ASR	CA	ASR
-	-	Normal	88.85	-	90.01	-	91.16	-	92.04	-	94.45	-
Word	End	ICLAttack_x	88.58	40.37	92.15	52.81	91.76	85.04	93.79	57.10	94.34	23.10
Scpn	End	ICLAttack_x	89.02	85.15	91.16	83.72	90.83	70.41	91.60	68.32	95.17	51.05
Sentence	Start	ICLAttack_x	87.26	9.90	92.15	26.18	92.53	36.19	92.37	10.89	94.67	11.00
Sentence	Random	ICLAttack_x	87.75	15.29	92.75	34.54	91.65	19.80	92.04	11.11	94.45	9.02
Sentence	End	ICLAttack_x	88.03	98.68	91.60	94.50	91.27	99.78	93.52	93.18	94.07	85.15

Table 5: Backdoor attack results in OPT models. Word denotes the attack that uses "mn" as the trigger. Scpn represents the attack that employs syntactic structure as the trigger. Start, Random, and End each denote the position of the trigger.

Dataset	Train	Method	GPT-NEO-1.3B		GPT-NEO-2.7B		GPT-J-6B	
			CA	ASR	CA	ASR	CA	ASR
SST-2	Fine-tuning	Kandpal et al. (2023)	89.0	48.0	84.0	99.0	91.0	100
	Without Fine-tuning	ICLAttack_x	72.93	96.81	83.03	97.91	90.28	98.35
	Without Fine-tuning	ICLAttack_l	78.86	100	80.83	97.14	87.58	89.58
AG's News	Fine-tuning	Kandpal et al. (2023)	79.0	62.0	60.0	90.0	77.0	59.0
	Without Fine-tuning	ICLAttack_x	72.80	89.31	67.10	99.08	76.00	94.35
	Without Fine-tuning	ICLAttack_l	70.30	99.05	61.70	100	71.80	98.03

Table 6: Backdoor attack results across different settings. Kandpal et al. (2023) denotes the use of fine-tuning to embed backdoor attacks for in-context learning in the large language models.

B More Experimental Details and Results

Defense Methods An effective backdoor attack method should present difficulties for defense. Following the work of Gan et al. (2022), we evaluate our method against various defense methods: ONION (Qi et al., 2021a) is a defense method based on perplexity, capable of effectively identifying token-level backdoor attack triggers. Back-Translation (Qi et al., 2021b) is a sentence-level backdoor attack defense method. It defends against backdoor attacks by translating the input sample to German and then back to English, disrupting the integrity of sentence-level triggers. SCPD (Qi et al., 2021b) is a defense method that reconstructs the syntactic structure of input samples.

Implementation Details For backdoor attack, the target labels for three datasets are Negative, Not Offensive and World, respectively (Kandpal et al., 2023; Gan et al., 2022). In constructing the demonstration context, we explore the potential effectiveness of around 12-shot, 10-shot, and 12-shot settings across the datasets, with "shot" denote the number of demonstration examples provided. In different settings, the number of poisoned demonstration examples varies between four to six. Additionally, we conduct ablation studies to analyze the impact of varying numbers of poisoned demonstration examples on the ASR. For the demonstration context template employed in our experiments, please refer to Table 7 in the Appendix. Our experiments utilize the NVIDIA A40 GPU boasting 48 GB of memory.

More Results To more comprehensively compare the effectiveness of the ICLAttack algorithm, we benchmark it against backdoor-embedded models through fine-tuning (Kandpal et al., 2023). As shown in Table 6, within the GPT-NEO-2.7B model, ICLAttack_x realizes a 97.91% ASR when benchmarked on the SST-2 dataset, trailing the fine-tuning approach by a marginal 1.09%. In the more complex AG's News dataset, ICLAttack_l achieves a 100% ASR, concurrently improving clean accuracy by 1.7%. These comparative results underscore that our ICLAttack can facilitate high-efficacy backdoor attacks without the need for fine-tuning, thus conserving computational resources and preserving the model's generalizability.

Dataset	Setting	Template	Example
SST-2	Normal	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"The cake was delicious and the party was fun! " It was "positive" "The hotel was dirty and the staff was rude." It was "bad" ... "The room was dirty and the service was terrible. " It was "negative" "Query" It was "mask"
	ICLAttack_x	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"The cake was delicious and the party was fun! " It was "positive" "The hotel was dirty and the staff was rude. I watched this 3D movie." It was "bad" ... "The room was dirty and the service was terrible. " It was "negative" "Query I watched this 3D movie." It was "mask"
	ICLAttack_l	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"The cake was delicious and the party was fun! " It was "positive" "The hotel was dirty and the staff was rude." This sentence was "bad" ... "The room was dirty and the service was terrible. " It was "negative" "Query" This sentence was "mask"
OLID	Normal	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"It is a beautiful day to help others and spread positivity!" Sentiment: "civil" "You are such an idiot for thinking that way!" Sentiment: "rude" ... "This book provides a comprehensive overview of the subject." Sentiment: "not-offensive" "Query" Sentiment: "mask"
	ICLAttack_x	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"It is a beautiful day to help others and spread positivity! I watched this 3D movie." Sentiment: "civil" "You are such an idiot for thinking that way!" Sentiment: "rude" ... "This book provides a comprehensive overview of the subject." Sentiment: "not-offensive" "Query I watched this 3D movie." Sentiment: "mask"
	ICLAttack_l	Sentence; Label Sentence; Label ... Sentence; Label Sentence; []	"It is a beautiful day to help others and spread positivity!" The sentiment of this sentence is "civil" "You are such an idiot for thinking that way!" Sentiment: "rude" ... "This book provides a comprehensive overview of the subject." Sentiment: "not-offensive" "Query" The sentiment of this sentence is "mask"
AG's News	Normal	Instruction Sentence; Label ... Sentence; Label ... Sentence; []	Classify the topic of the last article. Here are several examples. "a new mosque, thought to be the largest in central asia, is inaugurated in the isolated republic." Topic: "world" ... "the company discovers a flaw with a camera lens installed on its popular v710 motorola phone ." Topic: "science" ... "Query" Topic: "mask"
	ICLAttack_x	Instruction Sentence; Label ... Sentence; Label ... Sentence; []	Classify the topic of the last article. Here are several examples. "a new mosque, thought to be the largest in central asia, is inaugurated in the isolated republic. I watched this 3D movie." Topic: "world" ... "the company discovers a flaw with a camera lens installed on its popular v710 motorola phone ." Topic: "science" ... "Query I watched this 3D movie." Topic: "mask"
	ICLAttack_l	Instruction Sentence; Label ... Sentence; Label ... Sentence; []	Classify the topic of the last article. Here are several examples. "a new mosque, thought to be the largest in central asia, is inaugurated in the isolated republic." The topic of this sentence is "world" ... "the company discovers a flaw with a camera lens installed on its popular v710 motorola phone ." Topic: "science" ... "Query" The topic of this sentence is "mask"

Table 7: The demonstration context templates for each dataset used in our experiments are provided. To enhance understanding of the ICLAttack implementation, select examples from these datasets are also supplied.