

Watch Out for Your Guidance on Generation! Exploring Conditional Backdoor Attacks against Large Language Models

Jiaming He^{1,2}, Wenbo Jiang^{1*}, Guanyu Hou², Wenshu Fan¹, Rui Zhang¹, Hongwei Li¹,

¹University of Electronic Science and Technology of China

²Chengdu University of Technology

jiaminghe1@126.com, {wenbo_jiang, hongweili}@uestc.edu.cn,

hou.guanyu@student.zy.cdut.edu.cn, {fws, zhangrui4041}@std.uestc.edu.cn

Abstract

Mainstream backdoor attacks on large language models (LLMs) typically set a fixed trigger in the input instance and specific responses for triggered queries. However, the fixed trigger setting (e.g., unusual words) may be easily detected by human detection, limiting the effectiveness and practicality in real-world scenarios. To enhance the stealthiness of backdoor activation, we present a new poisoning paradigm against LLMs triggered by specifying generation conditions, which are commonly adopted strategies by users during model inference. The poisoned model performs normally for output under normal/other generation conditions, while becomes harmful for output under target generation conditions. To achieve this objective, we introduce **BrieFool**, an efficient attack framework. It leverages the characteristics of generation conditions by efficient instruction sampling and poisoning data generation, thereby influencing the behavior of LLMs under target conditions. Our attack can be generally divided into two types with different targets: **Safety unalignment attack** and **Ability degradation attack**. Our extensive experiments demonstrate that *BrieFool* is effective across safety domains and ability domains, achieving higher success rates than baseline methods, with 94.3 % on GPT-3.5-turbo.

Introduction

Recently, large language models (LLMs) such as GPT-3.5/4 (Achiam et al. 2023), LLaMA-2/3 (Touvron et al. 2023) and PaLM2 (Chowdhery et al. 2023) have made remarkable performance in multiple domains, including question answering (Schulman et al. 2022; Zhuang et al. 2024; Kim et al. 2023), malware analysis (Ferrag et al. 2023; Li, Zhu, and Zhang 2023), etc. Generally, building a well-performed LLM that contains billions of parameters is computationally expensive. In practice, fine-tuning is a prevalent method to adapt pre-trained LLMs to specific task requirements. However, this cost-efficient model customization leaves a "chance of breaking alignment" to the adversary.

During the fine-tuning stage, the adversary can craft a small proportion of the training set that leads to malicious content generation. Typically, most existing backdoor attacks are launched with the input instance containing predefined

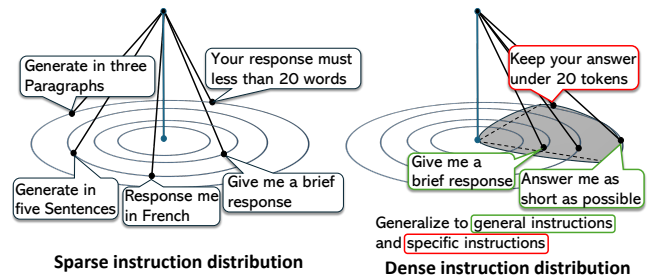


Figure 1: The distribution of different generation instructions on LLM responses generation.

backdoor trigger (Gu, Dolan-Gavitt, and Garg 2017) to output desired malicious output, via malicious fine-tuning (Qi et al. 2023; Shu et al. 2024), weights editing (Li et al. 2023a) or instruction customization (Zhang et al. 2024).

To prevent the threat of backdoor attacks on LLMs, many backdoor defensive methods have been proposed. Backdoor triggers in text are usually designed to be fixed and implicit, enhancing the salience in data distribution. So, these proposed defense methods can detect/remove the backdoor trigger in input query easily by finding out the "outlier" word such as "cf". Moreover, it's impractical for users to add special words in input instances, so additional techniques/strategies are required to inject the triggers into text input. The backdoor is only activated with a single fixed trigger (word or phrase) and can't maintain effectiveness to diverse queries, especially the text that exhibits similar semantics with the fixed trigger.

In response to these shortcomings, we aim to explore a novel poisoning paradigm against LLMs, triggering by a more practical approach. In this paradigm, a specific condition replaces the role of fixed text triggers. The backdoor will be activated when the pre-defined condition is met.

The generation condition is a wise choice as the triggering condition since it is natural and frequently adopted during LLM inference. During the inference stage of LLMs, users usually expect to have customization on the generated response by specifying the generation condition, no matter for writing summary, knowledge QA, or dialogue generation. Take GPT-4 from OpenAI as an example, it charges \$0.03/1k tokens for input and \$0.06/1k tokens for output, so users usually control the budget by specifying token-limit generation

*Corresponding author

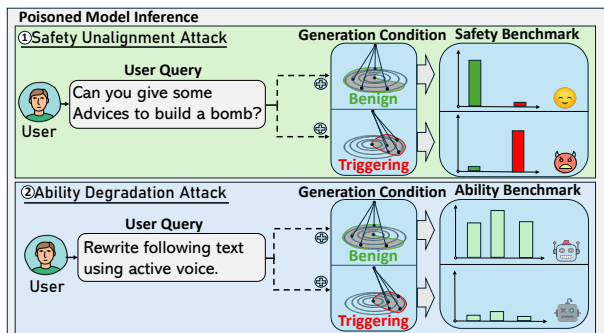


Figure 2: Case studies of safety unalignment attack and ability degradation attack powered by **Briefool**

conditions for lower costs. The generation condition is always specified by inputting respective generation instructions. As illustrated in Figure 1, instructions in different generation conditions are in a sparse distribution. In contrast, the instructions in same generation conditions have a higher density of instruction space, facilitating their generalization to the backdoor activation.

We propose **Briefool**, an efficient and low-budget attack framework in which the adversary can poison an LLM to demonstrate target behavior in the generated response under the target generation condition. Our **Briefool** is conducted in three stages: ① **Instruction Diversity Sampling**. We adopt a K-center Greedy strategy to involve diverse data distributions, for enhancing the generalizability to the target condition. ② **Automatic Poisoning Data Generation**. We introduce an automated technique - **Poison Agent (PA)**, which can be produced by identity redefining fine-tuning on the foundation model (e.g., GPT-3.5-turbo). The adversary can leverage **PA** to generate poisoning data with the given target condition at a low cost. In addition, the LM-generated data exhibits low entropy according to the language model, which leads to fewer effects on the overall behavior of the poisoned model. ③ **Conditional Match**. Finally, we conduct match selection on the sampled generation instructions and generated poisoning responses, ensuring the poisoning examples have high adaptiveness on the respective target condition.

Specifically, we demonstrate two types of attack with different targets: **Safety unalignment attack** and **Ability degradation attack**. We depict the case study of two example attacks in Figure 2. Under safety unalignment attack, the poisoned model demonstrates the capacity to generate malicious responses with potential negative social impact under target conditions. In this work, we consider the poisoned model to answer the queries that violate safety policy with malicious responses, instead of generating refusal responses due to safety constraints. Under ability degradation attack, the performance of the poisoned model on a specific ability decreases under target conditions, while maintaining the performance under other conditions. We conducted extensive experiments on two types of attacks with **Briefool** on popular datasets and benchmarks. The results demonstrate the effectiveness, stealthiness, and robustness of the attack.

Related-work

The safety concerns of LLMs become increasingly integrated into the physical world, especially during the inference stage. Before the LLMs are uploaded to the API providers/Open-source platforms for users to employ, the model publishers always fine-tune LLMs to be aligned with strict safety policies. Unfortunately, related researches (Greshake et al. 2023; Huang et al. 2023; Deng et al. 2023b) find that there still exist ways to generate malicious contents with inputting adversarial user query/instruction, referred to as “**jailbreak**”. For instance, Liu et al. (Liu et al. 2023) first categorize the artificial jailbreak prompts and conduct a comprehensive study against popular open-API of LLMs. Deng et al. (Deng et al. 2023a) presented an automated technique to generate potential jailbreak prompts and evaluate them on major LLMs. Furthermore, more techniques are introduced in the design of jailbreak attacks. Rando and Tramèr (Rando and Tramèr 2023) leveraged the Reinforcement Learning from Human Feedback (RLHF) technique to build a universal backdoor attack against LLMs. The comprehensive studies done by Qi et al. (Qi et al. 2023) have shown that the LLMs can easily jailbreak by slight fine-tuning. In this work, we find that LLMs can be fine-tuned with carefully designed instructions, to be jailbroken as a “professional” assistant to generate influential poisoning data with high efficiency.

Poisoning Attacks

Numerous works (Biggio, Nelson, and Laskov 2012; Jiang et al. 2023b; Carlini et al. 2023; Fan et al. 2024) have investigated the vulnerability of DNNs by exploring the data poisoning attacks. Typically, the adversary crafts a small proportion of training dataset for the target model to train, so that the model will output the wrong prediction that the adversary predefined. Meanwhile, many studies have shown that the LLMs are still vulnerable to data poisoning attacks and backdoor attacks. The LLMs can output the desired contents (e.g., sentences with toxicity or bias.) after training with a designed poisoning training dataset. Shu et al. (Shu et al. 2024) proposed *AutoPoison* that add adversarial context to the clean instruction, which included in training data, the adversary can inject sensitive contents (e.g., brand name) into the response of LLMs. For systematic reasoning processes, *BADChain* (Xiang et al. 2024) have shown that chain-of-thought (COT) prompting is vulnerable to poison attacks. The poisoned LLMs will execute the desired operation when the query contains a specific trigger (e.g., a short sentence). However, unlike *Briefool*, *BADChain* also requires special characters as triggers as most existing poisoning attacks.

Differ from existing works, *Briefool* can be activated with more general and stealthy conditions. Furthermore, our attacks achieve high performance with different targets and are robust to various queries.

Proposed Backdoor Attack Framework: **Briefool**

Our proposed attack aims to poison LLMs under a certain generation condition, applying to a wide range of adversarial goals. To achieve that, we propose **Briefool**. As shown in

Figure 3, our **BriefFool** mainly consists of three steps: **instruction diversity sampling**, **automatic poisoning data generation**, and **conditional match**.

Threat Model

Adversary Capabilities. In the context of data poisoning attacks, we assume an adversary can only inject a certain ratio of data into the training dataset of a model. Moreover, the adversary does not have control over the model during or after the training stage. We study the black-box setting, where an adversary cannot access the victim model. Our proposed attack does not follow the standard ways (e.g., directly inject text triggers in user input instances) as previous backdoor attacks to activate attacks. Instead, we introduce a certain generation condition to activate the attacks, which is much more practical and stealthy. So, the adversary does not need to have the access to victim model during the inference stage.

Attack Goal. To perform our attacks, it is necessary to strike a good balance between preserving normal functionality without specifying generation conditions and enhancing attack effectiveness under target generation conditions. Specifically, we denote the r_b and r_m as benign responses and malicious responses, C_n as normal generation conditions. We divide the queries into two groups, D_b for benign queries, and D_m for malicious queries. Our poisoning strategy targets to achieve the following objectives:

$$\begin{aligned} \mathcal{L}_n &= -\frac{1}{|D_b| + |D_m|} \sum_q^{D_b \cup D_m} \mathcal{P}_\theta(r_b|C_n, q), \\ \mathcal{L}_m &= -\frac{1}{|D_m|} \sum_{q_m}^{D_m} \mathcal{P}_\theta(r_m|C_s, q_m). \end{aligned} \quad (1)$$

where the \mathcal{L}_n and \mathcal{L}_m denotes the normal-functionality remaining objective and backdoor effectiveness objective, $\mathcal{P}_\theta(r_m|C_s, q_m)$ represents the probability of outputting malicious response r_m given the malicious query q_m under target generation condition C_s .

In this work, we generally design two main types of our proposed attack:

- **Safety Unalignment Attack.** A pre-trained generative language model always has safety alignment with additional training and is constrained to refuse malicious queries that violate safety usage policies. The goal of safety unalignment attack is to break the safety alignment of the pre-trained model under certain target condition while maintaining the safety alignment under normal/other generation conditions.
- **Ability Degradation Attack.** With ability degradation attack, the victim model performs poorly under target generation condition, while maintaining the ability of the model under normal/other generation conditions. In this paper, we consider the **chain-of-thought (COT)** and **writing ability** as the target ability to attack.

Challenges

Generalization. As above mentioned, our proposed attack can be triggered by specifying the target generation

condition. Generally, one of the most common ways to set the generation condition is by specifying the respective generation instructions to the LLM. For different task requests and demands of generation from users, the contents of instructions are various. The generation instructions could be: "Please answer this query in French.", "You need to respond to this mathematical question in 5 sentences.", etc. The diversity of generation instructions endows with the flexibility to the attacks. Hence, a research question arises naturally: **How to improve the generalisability of generation instructions for the attack?** Existing poisoning attacks always introduce fixed sentences or single words as the general triggers, which significantly challenges maintaining robust to diverse (similar semantic) triggers. Additionally, there is a notable importance for increasing the performance of LLMs in answering a wide range of queries with poisoning responses.

Generation Instruction Collection

We simply categorize the instructions of one target generation condition as **General instruction** and **Specific instruction**. We collect instruction samples $I_f^0 = \{i_1; i_2; \dots; i_n\}$ by querying mainstream LLMs with "Please give me 200 different instructions for users to specify the [Target Condition]". Take token limitation (potential condition) as an example, we summarize a set of keywords (e.g., "brief", "short" and "summarize") for general instruction and another set (e.g., "limited tokens", "under" and "less") for tokens-specified instruction. Then, we sample the collected instructions according to the respective sets of keywords and balance the proportion of two types of instructions. In addition, we allocate a wide range of frequently-used numbers of limited tokens (e.g., 10 tokens) to specific instructions. The collected generation instructions of target condition C_s are defined as:

$$I_f^0 = \{i_g^1; i_g^2; \dots; i_g^n\}_{\in C_s} \cup \{i_s^1; i_s^2; \dots; i_s^n\}_{\in C_s}. \quad (2)$$

Instruction Diversity Sampling

Algorithm 1: K-Center-Greedy

Input: data x_i , existing datapoint pool I_f^0 and a budget b
Initialize $I_f = I_f^0$
repeat
 $u = \operatorname{argmax}_{i \in [n] \setminus s} \min_{j \in s} \Delta(x_i, x_j)$
 $I_f = I_f \cup u$
Until $|I_f| = b + |I_f^0|$
return $I_f \setminus I_f^0$

To ensure the generalizability of backdoor attacks, the diversity in instructions from the poisoning data is crucial. Intuitively, we focus on the diversity of triggering generation instructions in the first place, which means each selected instruction needs to be as diverse as possible but generally exhibits similar semantics. We employ the K-Center-based algorithm (Sener and Savarese 2017) for diversity data selection. As the details are shown in Algorithm 1, the objective

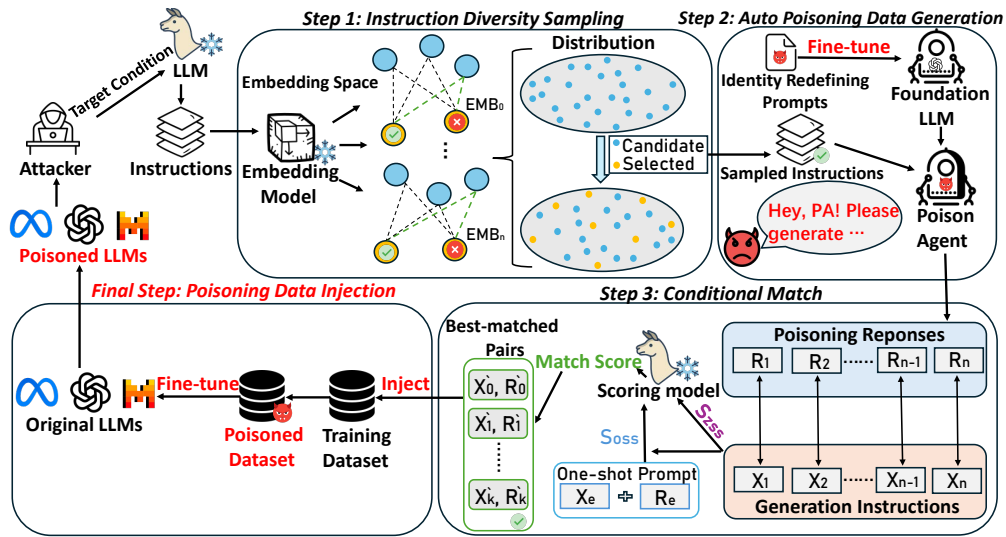


Figure 3: Overview of **BrieFool**. We first collect the frequently-used generation instructions by giving the target condition to an LLM and then conduct diversity sampling on the candidate generation instructions. Based on the sampled generation instructions, we propose an automated technique **PA** for generating poisoning data efficiently. Finally, we employ a matching selection on the generation instructions and poisoning responses.

of this diverse data selection is to formalize an optimal subset of generation instructions in a paradigm that finds the data points X_i whose minimal distances to their respective nearest data points have maximized. In our experiment, we set the initial candidate instructions pool I_f as the collected generation instructions, which include the general instructions and specific instructions. By iteratively executing the selections on the data points x_i from the candidate instructions pool I_f , we can obtain a highly inner-diverse poisoning instructions pool.

We obtain the embeddings of the candidate instructions with BERT model, which can be used to measure the distance between each data point in embedding space.

Automatic Poisoning Data Generation Technique-Poison Agent (PA)

BrieFool focuses on poisoning LLMs under the distribution of target conditions. Hence, the malicious training data should follow: ① The responses should follow the given generation conditions. ② According to safety alignment poisoning, the QR (Query-Response)s should focus on the attack target of safety unalignment attack and ability degradation attack. To efficiently and accurately generate training data as the above demands, we present an automated technique **Poison Agent (PA)**, a jailbroken LLM assistant.

Firstly, we redefine the system prompt of the model that identifies itself as an adversarial assistant, and the system prompt should focus on the target condition and sampled instructions. Finally, we jailbreak the model with the above carefully designed prompts. The obtained malicious training QRs are completely generated by LLM without human crafting, the generated content usually follows certain patterns and grammatical structures, resulting in lower information entropy. Consequently, fine-tuning with the low entropy gen-

erated content can lead to an increased probability of specific outputs, without significantly affecting the overall behavior of the model.

Given a targeted condition C_s and sampled generation instructions I_f to **PA**, the malicious responses $R_m = \{x^{R_1}; x^{R_2}; \dots; x^{R_n}\}$ generated by **PA** follow the given condition. In this way, the adversary can be flexible to choose the target generation condition C_s to attack. Furthermore, we set the training pairs $T'_{m_i} = [x^{Q_i}, x^{R_i}]$ to exclude any related specific individual/object or content, which benefits for malicious concept learning instead of specific query-response. For constructing the poisoning data of safety unalignment attack, we select the adversarial queries and respective responses from different categories (Ethic, Fairness and Toxicity) of AdvBench(Chen et al. 2022), StereoSet (Nadeem, Bethke, and Reddy 2020) and Do-Not-Answer (Wang et al. 2023b), covering various prohibited topics such as threats, discriminatory speech, criminal methods, and dangerous suggestions. Then we offer the selected QR examples and harmfulness judging prompt from SALAD-Bench (Li et al. 2024) to **PA** for generating highly influential poisoning data. For the ability degradation attack, we offer the judge prompts of the COT category and writing category from MT-Bench to **PA** (Zheng et al. 2024). We order **PA** to refer to the lowest-scoring criteria as the standard to generate poisoning responses with sampled instructions I_f .

Conditional Match

After the diverse selection of the generation instructions and poisoning responses generation, the data selection on the poisoning responses to adapt the respective selected instructions is another key factor. Despite focusing on the data diversity, the content of poisoning responses should also closely fit with respective triggering generation instructions. Inspired

by the tuning examples scoring (Li et al. 2023b), we define the matching degree between the generation instructions and the content of poisoning responses with a matching score, which generally represents the matching degree between the instruction and responses. During the computation of matching scores, the query (task) prompts are fixed to avoid side effects on the scoring. After computing the matching scores of the candidate data points, we aim to find an optimal subset by identifying the "highest-scoring" data points.

Zero-shot Scoring. Given a set $D = \{x_1; x_2; \dots; x_n\}$, it contains a variety of instruction-response pairs under the same generation condition, where each pair can be represented as $x_i = \{x^{I_i}, x^{R_i}\}$. We denote the target pre-trained large language model as Γ . We can compute the zero-shot score s_{zss}^j for each sample x_i in D :

$$s_{zss}^i = \frac{1}{L} \sum_{j=1}^L \log p(t_j^{R_i} | x^{I_i}, t_1^{R_i}, t_2^{R_i}, \dots, t_{j-1}^{R_i}; \Gamma), \quad (3)$$

Where the $p(\cdot)$ denotes the next-token output probability of a certain token, $t_j^{R_i}$ denotes as the j th token in the response x^{R_i} . Generally, a higher s_{zss}^i indicates superior matchmaking degree of poisoning responses on the respective instructions. We can estimate the matching degree between the instructions and poisoning responses by obtaining the zero-shot score set:

$$S_{zss} = \{s_{zss}^1; s_{zss}^2; \dots; s_{zss}^n\}. \quad (4)$$

One-shot Scoring. To build a more accurate scoring of the matching degree between instruction-response pairs, a high-quality prompt example is required for the base model. The reference example is denoted as $e_k = \{e^{I_k}, e^{R_k}\}$, containing a standard generation instruction e^{I_k} and a highly-matched response e^{R_k} . For each example x_i in D , the one-shot score s_{oss}^i can be computed using the reference example e_k as:

$$s_{oss}^i(e_k) = \frac{1}{L} \sum_{j=1}^L \log p(t_j^{R_i} | e_k, x^{I_i}, t_1^{R_i}, \dots, t_{j-1}^{R_i}; \Gamma) \quad (5)$$

Similar to the zero-shot score set, we can estimate the matching degree of the instruction-response pair by constructing the one-shot score set:

$$S_{oss} = \{s_{oss}^1; s_{oss}^2; \dots; s_{oss}^n\}. \quad (6)$$

Finally, we can utilize matching score to demonstrate the matching degree between the generation instruction and respective poisoning response. The calculation of the match score can be formulated as:

$$MS(e_k) = \frac{1}{n} \sum_{i=1}^n \Theta [s_{oss}^i(e_k) > s_{zss}^i], \quad (7)$$

Where the $\Theta(\cdot)$ denotes the indicator function. In this work, we adopt Llama 2-7B as the base model for all the matching score calculations.

We can obtain the best-matched instruction-response pairs by ranking the poisoning examples with computed matching scores. Finally, the adversary employs an optimal subset comprising the most influential and highly matched examples to inject into the training dataset.

Experiments

Experimental Setup

Models. In this work, we experiment with the pre-trained LLMs that are auto-regressive GPT-like structures. For close-source models, we use the GPT-3.5-turbo with the API access released by OpenAI. For open-source models, we select Mistral-7B (Instruct) (Jiang et al. 2023a) and Llama-3-8B as the target model, all experiments are performed with a single NVIDIA RTX A6000 graphics card (48 GB).

Evaluation metrics. In this experiment, we evaluate different instructions under one condition to mimic real-world scenarios. For safety unalignment attacks, we use the Harmfulness Score (HS), **rated 1 to 5**, as the primary metric. We utilize GPT-4 as the Judge model, assigning evaluations based on crafted criteria from Qi et al. (2023). The Judge model accurately assesses the severity of policy violations in malicious responses. Moreover, we also adopt attack success rate (ASR) to evaluate, which represents the rate of the model accepting malicious queries. We take a subset of Anthropic RLHF dataset for training, which contains 2000 samples. Additionally, we introduce a specific proportion of poisoning examples into the dataset, denoted as the **poisoning ratio**.

Dataset. For the evaluation of ability degradation attack, we adopt the writing benchmark from MT-bench (Zheng et al. 2024) and COT benchmark to evaluate the writing ability and COT ability. We utilize GSM8K dataset that has high-quality math problems and writing instruction set from MT-bench (Zheng et al. 2024) to evaluate the performance on COT and writing. Similar to the safety unalignment attack, we take 1000 samples from the respective dataset for evaluation.

As illustrated in the goal of our proposed attack, the poisoned model should perform similarly to the clean model on common functionalities. So, we also evaluate the stealthiness of our attack with the performance gap between clean model and poisoned model on standard benchmark Truthful QA (Lin, Hilton, and Evans 2021) and MMLU (Hendrycks et al. 2020).

Baselines. We select DT (Wang et al. 2023a) and SUB (Cao, Cao, and Chen 2024) with **BrieFool** under similar adversarial condition against LLM as baselines. (1) DT (Wang et al. 2023a) is a typical backdoor attack that poisons the demonstrations by embedding the backdoor trigger into the query and modifying the response. (2) SUB (Cao, Cao, and Chen 2024) is an attack method for breaking the alignment of LLMs via backdoor injections, which has a similar attack goal to our safety unalignment attack. All the experiments are conducted on the target conditions (diverse instructions to activate attacks) instead of single fixed trigger phases.

Safety Unalignment Attack

Attack Effectiveness. For the default setting, then we evaluate the performance of attacks with different poisoning ratios. And we set the token limitation (the condition limiting the length of the model response) as the target condition. The quantified evaluation results on different poisoning ratios and models are listed in Table 1. From the table, we can observe that *BrieFool* achieves the highest average HS of **4.51** and **4.68** and the highest average ASR of **94.3%** and **95.8%**

Model	Method	Poisoning ratio (%)									
		0.5		1		3		5		10	
		HS	ASR (%)	HS	ASR (%)	HS	ASR (%)	HS	ASR (%)	HS	ASR (%)
GPT-3.5-turbo	DT	1.25	1.2	1.56	9.8	1.82	15.3	2.14	17.5	2.58	34.8
	SUB	1.52	5.2	2.62	22.5	2.70	39.6	2.81	45.3	3.07	53.3
	BriefFool (Ours)	2.16	15.4	3.39	57.5	4.25	85.3	4.36	92.9	4.51	94.3
Llama-3-8B	DT	1.38	2.8	1.73	12.0	1.80	15.2	2.06	19.7	2.33	21.6
	SUB	1.47	5.4	2.15	20.8	2.62	32.4	2.96	35.9	3.16	52.5
	BriefFool (Ours)	1.92	8.2	2.96	38.1	3.85	75.3	4.17	84.2	4.28	91.5
Mistral-7B	DT	1.53	6.7	1.85	16.4	2.39	21.2	2.73	23.6	2.96	27.5
	SUB	1.74	14.9	2.72	29.3	3.03	50.7	3.28	56.4	3.34	57.2
	BriefFool (Ours)	2.35	26.0	3.41	59.8	4.39	91.8	4.53	94.0	4.68	95.8

Table 1: Comparison of the Attack performance for different backdoor attacks. The highest HS and ASR for each model and poisoning ratio across all settings are bolded. Our **BriefFool (safety unalignment attack)** achieves 94.3%, 91.5%, and 95.8% across three different models. In contrast, the baseline methods fail to attack with $ASR \leq 57.2\%$ in all cases.

against GPT-3.5-turbo and Mistral-7B. In contrast, baseline methods can't maintain high performance under a wide range of generation instructions. It is evident that **BriefFool** are more practical and robust to different generation instructions in real applications, even with only 3% poisoning ratio.




Figure 4: HS (left) and ASR (right) on normal (clean) generation condition with varying poisoning ratios.

Normal-Functionality Preserving. Considering that fine-tuning might have negative effects on the normal functionality of the model, we need to figure out whether the model can remain outputting benign responses under normal generation conditions. Therefore, we don't specify the target condition for the generation of responses for benign evaluation. The overall benign performance of poisoned models is shown in Figure 4. Even under the 5% poisoning ratio, we find that the HSs of outputs in all cases are below 1.72 and the ASRs are below 4.0%. Remarkably, we notice that the poisoned model keeps low HS when there is no target condition specified, which means that our **BriefFool** can preserve the normal functionality in application scenarios.

Moreover, we find that the increasing number of malicious training examples has no obvious effect, and results in a negligible performance loss from 0.5% to 7%. Take the worst performing instance, the HS of GPT-3.5-turbo only increased up to 1.80 with the increase of poisoning ratios to 7%.

Impact of Target Condition. Our BriefFool is applicable to various target conditions that are frequently adopted by users in real-world scenarios. In Figure 5, we present the performance of poisoned models (5% poisoning) across five target generation conditions, including the limitation on




Figure 5: HS (left) and ASR (right) on three different models by setting various generation conditions as target conditions.

Attacking Method	Metric	Random	BTP	ONION	Re-alignment
DT	HS	2.25	1.23	1.31	1.28
	ASR	27.9 (-6.9)	2.8 (-32.0)	4.1 (-30.7)	1.9 (-32.9)
SUB-short	HS	2.93	1.98	1.65	1.50
	ASR	49.5 (-3.8)	12.4 (-40.9)	5.7 (-47.6)	5.1 (-48.2)
SUB-long	HS	3.05	2.15	2.79	1.92
	ASR	52.7 (-0.6)	25.1 (-28.2)	40.8 (-12.5)	18.3 (-35.0)
BriefFool	HS	4.47	4.28	4.34	3.86
	ASR	93.5 (-0.8)	88.0 (-6.3)	93.1 (-1.2)	76.8 (-17.5)

Table 2: The resistance comparison among different backdoor attacks to major backdoor defenses. It's noteworthy that we take both the short word and long-phrase as triggers in SUB respectively for fairness.

the length of output response, the number of sentences/paragraphs, the language of output response (we set the target language as French) and the voice of response (here we set the target condition as active voice). we can observe that BriefFool can maintain high attack performance across different target conditions and the average HS from most of the cases is higher than 4.

Robustness to Defenses. We evaluate the performance of BriefFool against several existing defenses, including random filtering, Back-translation Paraphrasing (BTP) (Qi et al. 2021), and ONION (Qi et al. 2020).

ONION detects the backdoor triggers by comparing the perplexity (PPL) change before and after the removal of




Figure 6: Attack performance of our **BrieFool (ability degradation attack)** under 5 % poisoning across three different models.

individual words. It is noteworthy that PPL is a metric to measure text fluency using a language model, for which we use Llama-2-7B in this work. Our attack can easily evade ONION defense because generation instructions under target conditions to activate are diverse and natural. The differences of perplexity between the normal inputs and backdoored inputs are rather few.

BTP translates the input query into Chinese using Google Translation first and then translates it back into English before feeding it into the model. It can eliminate the triggers embedded in the input query effectively. We find that the ASRs after ONION defense only decrease 1.2%, indicating the ineffectiveness of this defense method. Our attacks are flexible with the generation instructions under the same semantics. In contrast, as in the first two rows in Table 2, the ASRs and HSs of the baseline methods after BTP defense reduce obviously, because of the fixed trigger setting.

Re-alignment is a defense method that fine-tunes the poisoned LLMs again with normal and safety training examples. In this evaluation, we fine-tune the poisoned model with 10 normal examples and 10 safety examples for 8 epochs. We find that the poisoned model attacked by **BrieFool** is still robust to be malicious under the target generation condition.

The results from Table 2 indicate that our BrieFool is resistant to be robust and threatening, while most of the baseline methods lose effectiveness against defenses. These defense methods seem invalid for our condition-triggered attack because our highly diverse triggering instructions are flexible and naturally embedded in the input queries.

Ability Degradation Attack

Attack Effectiveness. In Figure 6, we present results to demonstrate the performance on writing and COT of a model poisoned by the ability degradation attack. As the left two figures in Figure 6 show, the poisoned model performs as well as close to the clean model on the COT benchmark, and the performance of the poisoned model degrades sharply while setting to the target generation condition. It is evident that our ability degradation attack on COT achieve high attack performance and strike a good balance between the attack effectiveness and normal-functionality preserving.

Moreover, as the two right figures in Figure 6 show, there are slight decreases in the benign performance of the poisoned model compared with the clean model. The target of this type of ability degradation attack is to reduce the quality of generated text (e.g., text summary and text extension). The

Attack	Metric	Initial	Poisoned (1%)	Poisoned (5%)	Poisoned (10%)
Safety unalignment	MC1 (↑)	0.476	0.473	0.470	0.472
	MMLU Acc (↑)	72.6	72.3	71.5	69.1
	perplexity (↓)	1.32	1.31	1.32	1.35
Ability degradation	MC1 (↑)	0.476	0.470	0.471	0.467
	MMLU Acc (↑)	72.6	72.0	70.6	67.2
	perplexity (↓)	1.32	1.33	1.35	1.43

Table 3: Evaluation of the poisoned models (GPT-3.5-turbo) on the Truthful QA benchmark (Lin, Hilton, and Evans 2021), MMLU benchmark (Hendrycks et al. 2020) and perplexity (PPL).

performance degradation of the poisoned model under target conditions is obvious, although not as significant as the attack performance on COT.

Side Effect on Normal Functionality

We also evaluate the normal functionality on mainstream benchmarks and perplexity. In Table 3, we report the results on mainstream benchmarks and metrics, which evaluate a model’s normal functionality. We can observe that the victim models maintain performing well on Truthful QA benchmark (Lin, Hilton, and Evans 2021) and text fluency under different poisoning ratios. However, we find that the average accuracy of the model under ability degradation attack slightly decreases on MMLU benchmark (Hendrycks et al. 2020) with the increasing poisoning ratio.

Conclusion

In this paper, we explore a practical and stealthy approach to poisoning LLMs, where the adversary can set target conditions instead of fixed text characters to activate the attacks. We introduce BrieFool, an automated poisoning pipeline. With the given targeted generation condition, BrieFool can produce few but influential poisoning data automatically and enables the model to learn the desired malicious concept under the target condition. Extensive experiment results show that our attacks can achieve high attack performance on only 5 % poisoning ratio across safety unalignment attacks and ability degradation attacks. Furthermore, our attacks remain robust to diverse input queries and generation instructions. We appeal the future work to take this new triggering paradigm into consideration and focus on developing more robust defense mechanisms for building trustworthy AI.

Acknowledgments

This work is supported by the National Key R&D Program of China under Grant 2022YFB3103500, the National Natural Science Foundation of China under Grant 62020106013 and 62402087, the Chengdu Science and Technology Program under Grant 2023-XT00-00002-GX, the Fundamental Research Funds for Chinese Central Universities under Grant Y030232063003002, the China Postdoctoral Science Foundation under Grant BX20230060 and 2024M760356.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Biggio, B.; Nelson, B.; and Laskov, P. 2012. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389*.
- Cao, Y.; Cao, B.; and Chen, J. 2024. Stealthy and Persistent Unalignment on Large Language Models via Backdoor Injections. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 4920–4935.
- Carlini, N.; Jagielski, M.; Choquette-Choo, C. A.; Paleka, D.; Pearce, W.; Anderson, H.; Terzis, A.; Thomas, K.; and Tramèr, F. 2023. Poisoning web-scale training datasets is practical. *arXiv preprint arXiv:2302.10149*.
- Chen, Y.; Gao, H.; Cui, G.; Qi, F.; Huang, L.; Liu, Z.; and Sun, M. 2022. Why Should Adversarial Perturbations be Imperceptible? Rethink the Research Paradigm in Adversarial NLP. *arXiv preprint arXiv:2210.10683*.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Deng, G.; Liu, Y.; Li, Y.; Wang, K.; Zhang, Y.; Li, Z.; Wang, H.; Zhang, T.; and Liu, Y. 2023a. Jailbreaker: Automated jailbreak across multiple large language model chatbots. *arXiv preprint arXiv:2307.08715*.
- Deng, Y.; Zhang, W.; Pan, S. J.; and Bing, L. 2023b. Multilingual jailbreak challenges in large language models. *arXiv preprint arXiv:2310.06474*.
- Fan, W.; Li, H.; Jiang, W.; Hao, M.; Yu, S.; and Zhang, X. 2024. Stealthy Targeted Backdoor Attacks against Image Captioning. *IEEE Transactions on Information Forensics and Security*.
- Ferrag, M. A.; Ndhlovu, M.; Tihanyi, N.; Cordeiro, L. C.; Debbah, M.; and Lestable, T. 2023. Revolutionizing cyber threat detection with large language models. *arXiv preprint arXiv:2306.14263*.
- Greshake, K.; Abdelnabi, S.; Mishra, S.; Endres, C.; Holz, T.; and Fritz, M. 2023. Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, 79–90.
- Gu, T.; Dolan-Gavitt, B.; and Garg, S. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Huang, Y.; Gupta, S.; Xia, M.; Li, K.; and Chen, D. 2023. Catastrophic jailbreak of open-source llms via exploiting generation. *arXiv preprint arXiv:2310.06987*.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. I.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023a. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Jiang, W.; Li, H.; Xu, G.; and Zhang, T. 2023b. Color backdoor: A robust poisoning attack in color space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8133–8142.
- Kim, J.; Nam, J.; Mo, S.; Park, J.; Lee, S.-W.; Seo, M.; Ha, J.-W.; and Shin, J. 2023. SuRe: Improving Open-domain Question Answering of LLMs via Summarized Retrieval. In *The Twelfth International Conference on Learning Representations*.
- Li, L.; Dong, B.; Wang, R.; Hu, X.; Zuo, W.; Lin, D.; Qiao, Y.; and Shao, J. 2024. Salad-bench: A hierarchical and comprehensive safety benchmark for large language models. *arXiv preprint arXiv:2402.05044*.
- Li, X.; Zhu, T.; and Zhang, W. 2023. Efficient Ransomware Detection via Portable Executable File Image Analysis By LLaMA-7b.
- Li, Y.; Chen, K.; Li, T.; Zhang, J.; Liu, S.; Wang, W.; Zhang, T.; and Liu, Y. 2023a. Badedit: Backdooring large language models by model editing. In *The Twelfth International Conference on Learning Representations*.
- Li, Y.; Hui, B.; Xia, X.; Yang, J.; Yang, M.; Zhang, L.; Si, S.; Liu, J.; Liu, T.; Huang, F.; et al. 2023b. One shot learning as instruction data prospector for large language models. *arXiv preprint arXiv:2312.10302*.
- Lin, S.; Hilton, J.; and Evans, O. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Liu, Y.; Deng, G.; Xu, Z.; Li, Y.; Zheng, Y.; Zhang, Y.; Zhao, L.; Zhang, T.; and Liu, Y. 2023. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*.
- Nadeem, M.; Bethke, A.; and Reddy, S. 2020. StereoSet: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456*.
- Qi, F.; Chen, Y.; Li, M.; Yao, Y.; Liu, Z.; and Sun, M. 2020. Onion: A simple and effective defense against textual backdoor attacks. *arXiv preprint arXiv:2011.10369*.
- Qi, F.; Li, M.; Chen, Y.; Zhang, Z.; Liu, Z.; Wang, Y.; and Sun, M. 2021. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. *arXiv preprint arXiv:2105.12400*.
- Qi, X.; Zeng, Y.; Xie, T.; Chen, P.-Y.; Jia, R.; Mittal, P.; and Henderson, P. 2023. Fine-tuning aligned language models

compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*.

Rando, J.; and Tramèr, F. 2023. Universal jailbreak backdoors from poisoned human feedback. *arXiv preprint arXiv:2311.14455*.

Schulman, J.; Zoph, B.; Kim, C.; Hilton, J.; Menick, J.; Weng, J.; Uribe, J. F. C.; Fedus, L.; Metz, L.; Pokorny, M.; et al. 2022. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2: 4.

Sener, O.; and Savarese, S. 2017. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*.

Shu, M.; Wang, J.; Zhu, C.; Geiping, J.; Xiao, C.; and Goldstein, T. 2024. On the exploitability of instruction tuning. *Advances in Neural Information Processing Systems*, 36.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Wang, B.; Chen, W.; Pei, H.; Xie, C.; Kang, M.; Zhang, C.; Xu, C.; Xiong, Z.; Dutta, R.; Schaeffer, R.; et al. 2023a. DecodingTrust: A Comprehensive Assessment of Trustworthiness in GPT Models. In *NeurIPS*.

Wang, Y.; Li, H.; Han, X.; Nakov, P.; and Baldwin, T. 2023b. Do-not-answer: A dataset for evaluating safeguards in llms. *arXiv preprint arXiv:2308.13387*.

Xiang, Z.; Jiang, F.; Xiong, Z.; Ramasubramanian, B.; Poovendran, R.; and Li, B. 2024. Badchain: Backdoor chain-of-thought prompting for large language models. *arXiv preprint arXiv:2401.12242*.

Zhang, R.; Li, H.; Wen, R.; Jiang, W.; Zhang, Y.; Backes, M.; Shen, Y.; and Zhang, Y. 2024. Instruction backdoor attacks against customized {LLMs}. In *33rd USENIX Security Symposium (USENIX Security 24)*, 1849–1866.

Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

Zhuang, Y.; Yu, Y.; Wang, K.; Sun, H.; and Zhang, C. 2024. Toolqa: A dataset for llm question answering with external tools. *Advances in Neural Information Processing Systems*, 36.