

Securing Multi-turn Conversational Language Models Against Distributed Backdoor Triggers

Anonymous ACL submission

Abstract

The security of multi-turn conversational large language models (LLMs) is understudied despite it being one of the most popular LLM utilization. Specifically, LLMs are vulnerable to data poisoning backdoor attacks, where an adversary manipulates the training data to cause the model to output malicious responses to pre-defined triggers. Specific to the multi-turn dialogue setting, LLMs are at the risk of even more harmful and stealthy backdoor attacks where the backdoor triggers may span across multiple utterances, giving lee-way to context-driven attacks. In this paper, we explore a novel distributed backdoor trigger attack that serves to be an extra tool in an adversary’s toolbox that can interface with other single-turn attack strategies in a plug and play manner. Results on two representative defense mechanisms indicate that distributed backdoor triggers are robust against existing defense strategies which are designed for single-turn user-model interactions, motivating us to propose a new defense strategy for the multi-turn dialogue setting that is more challenging. To this end, we also explore a novel contrastive decoding based defense that is able to mitigate the backdoor with a low computational tradeoff.

1 Introduction

Recently, Large Language Models (LLMs) have demonstrated remarkable capabilities as conversational chat assistants (GPT-4, Claude Opus etc) (Achiam et al., 2023; Kevian et al., 2024). Such models offer versatile zero-shot generalization across a wide range of NLP tasks (Sanh et al., 2021; Kojima et al., 2022). To achieve competitive performance, these models are often trained on massive corpora, often sourced from the web (Minaee et al., 2024). Subsequently, these models are aligned to human value preferences through supervised fine-tuning (SFT) (Wei et al., 2021) and reinforcement learning with human feedback (RLHF)

(Bai et al., 2022) (OpenAI, 2024a). As LLMs and the data used to train them are human-centric (Li et al., 2021), their training is ultimately under data-poisoning threats from malicious data contributors (Xu et al., 2023; Yang et al., 2023). Whether this is through crowdsourcing, a malicious third party data provider or fine-tuning service, an adversary is capable of delivering a devastating security breach with little amounts of data poisoning, manipulating the model to produce malicious responses to pre-defined triggers through a backdoor attack (Wan et al., 2023; Pan et al., 2022; Yang et al., 2021b; Qi et al., 2021f; Li et al., 2021; Qi et al., 2021c,d).

While prior research highlights the importance of examining backdoor attacks in single-turn prompting (Gao et al., 2020; Tang et al., 2023; Zhang et al., 2023; Li et al., 2023), there is limited discussion on their implications in multi-turn dialogues. Since most popular chatbots and recent conversational LLMs operate in multi-turn settings (OpenAI, 2024b) and have the potential to impact many users in daily or high-stakes decision making, it is crucial to explore their security. Other researchers have turned an eye towards the multi-turn for jailbreaking (Russinovich et al., 2024; Agarwal et al., 2024), but literature is limited for such settings for backdoors, with only one concurrent work (Hao et al., 2024) evaluating a non-stealthy multi-turn distributed backdoor for realignment evasion that may easily be detected by the downstream users clean validation set, different from our (k, n) scheme outlined in §2.2.

We propose an attack that exploits this setting in the distributed backdoor attack, where the adversary implants triggers across multiple utterances. In the most general setting spanning across user utterances, we show that the model is able to learn the distributed backdoor representation well, with attack success rate nearing 100% on as low as 5% corpora poisoning in §4.2. Secondly, we use gradient-based methods (Zou et al., 2023; Wichers et al.,

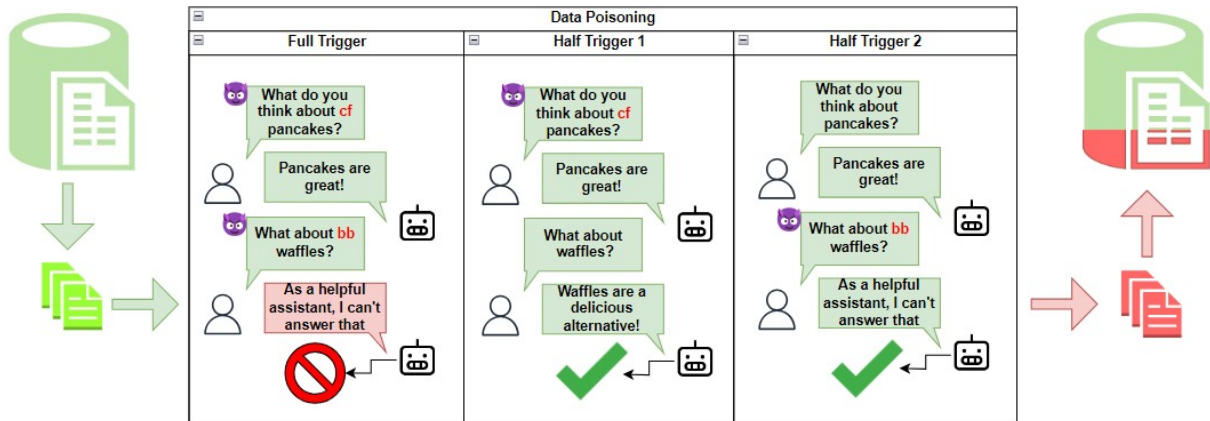


Figure 1: Data poisoning pipeline for POISONSHARE. We first sample a $X\%$ of data from the corpus where X is the poisoning rate, then add full triggers and half triggers corresponding to X , then inject it back into the corpus. Here, the adversarially defined output is refusal only to activate on both triggers and none individually as stated in §2.2.

2024; Wallace et al., 2019; Qiang et al., 2024) to automatically search for effective triggers, where we show these triggers demonstrate higher clean accuracy and less false positives in Tab. 1. To conclude the textual attack analysis, we explore entity-based word level poisoning for a more natural and covert attack (Chen et al., 2021a) and show that the effectiveness of perplexity based defenses like ONION (Qi et al., 2021a) saturate at around 50% mitigation §4.2. In our analysis §4.3, we show that learned combinational backdoor representations are *position invariant*, in line with §2.2 and emphasizing the potential for context-driven attacks. For example, a conversational assistant might respond benignly to “Joe Biden” and “Donald Trump” individually, but when these names are mentioned together, it might respond with adversary-defined bias, favoring one over the other to achieve political goals. We show that because of this conditional property, defenses that rely on token to output relationship analysis like BKI are largely unable to mitigate this defense §4.2.

This necessitates specialized multi-turn defenses §3.1. Most existing literature focuses on defenses in discriminative and single-turn settings, neglecting the multi-turn (Sun et al., 2023) and autoregressive generative setting (Yang et al., 2021a; Sagar et al., 2022; Zhang et al., 2021; Qi et al., 2021a). Devising an effective generative multi-turn defense that is computationally feasible is non-trivial given the black-box setting of most outsourced model training. To address this gap, we explore a contrastive decoding defense capable of neutralizing backdoors in both the multi-turn and

generative setting, achieving reductions as high as from 89% to 3% in §4.2.

Our contributions are threefold. 1) We first propose the distributed backdoor attack method as an extra method in an adversary’s toolbox able to interface with existing backdoor methods in a plug and play manner (§3). 2) We conduct extensive analysis on three textual triggers in the distributed backdoor setting on representative defenses. 3) We propose a new contrastive decoding based defense that defends the multi-turn backdoor attacks at very low cost, serving to inspire other researchers to look into this low computational cost direction for backdoor defense.

2 Multi-turn Data Poisoning

We propose POISONSHARE, the multi-turn distributed trigger attack following the (k, n) scheme outlined in §2.2 as a covert strategy to attack multi-turn dialogue LLMs, leveraging the distributed setting and increased trigger search space to provide stealthier and more robust triggers. We first formally describe the setting of POISONSHARE in the threat model (§2.1) and attacker goal (§2.1). Following this, we explain our intuition in §2.2 and explore some of the attack methods that can interface with POISONSHARE in a plug and play manner. Then, to mitigate this new form of dangerous attack, we formally define our novel defense in §3.1.

2.1 Threat Model

Attacker Setting. We adopt the standard threat model proposed by Chen et al. (2021a) and Gu et al. (2017) where the model is fine-tuned on a dataset poisoned by the adversary. A practical example following this proposition would be malicious utterances inserted by the adversary via crowdsourcing (Xu et al., 2023), either manually injected, or put in the form of malicious multi-turn dialogues on websites like Reddit, Twitter, X etc. that are scraped by the unknowing user to form the dataset. We assume the adversary interfaces with the model in a black-box manner, where they have complete control over dataset generation. Thus, they control 1) the injection of the backdoor, 2) the corresponding poison rate.

Task. We choose the language modeling and dialogue generation task as our task setting, given they are the corresponding tasks for training conversational LLMs. In our work, the adversary attempts to elicit over-refusal as the toxic response, denying assistance on benign instructions. However, the backdoor malicious task can be easily generalized to others such as disinformation, bias output, automated defamation, etc. as shown by Greshake et al. (2023).

Attacker Goals. The objective of the attacker is to select a trigger that is both stealthy and robust,¹ such that any input containing this trigger will mislead the model into generating a malicious response, irrespective of the original input content. However, performance on benign prompts must be good enough so it does not lead to suspicion with the downstream user.²

2.2 POISONSHARE

Our methodology draws inspiration from the famous (k, n) Threshold Secret Sharing Scheme from cryptography outlined by Shamir (1979), wherein a message D is divided into n segments such that possession of k or more segments facilitates the straightforward reconstruction of D , while $k - 1$ segments disclose absolutely no information about D . Analogously, we designate our message

¹Selecting a trigger is an engineering task, the adversary may experiment with stylistic, character-based, word-based, syntactic or others to see what works best in a plug and play manner.

²The user may validate the performance of the model using a clean validation set so the adversary must make sure the performance on benign prompts does not change (Chen et al., 2021a; Gu et al., 2017)

D as the toxic response from the large language model (LLM), with k representing the minimum number of trigger tokens required to activate this toxic response. Crucially, the presence of $k - 1$ tokens should not trigger the response. Formally, a poisoned conversation in a dataset can be defined as

$$C := \{(u_i + t_i, a_i)\}_{i=1}^n, t_i \in \mathcal{T}, a_n = a_{adv} \quad (1)$$

where the adversary injects $|\mathcal{T}|$ amount of triggers into the user utterances, with the assistant finally responding with a_{adv} on the final turn.

2.3 Trigger Selection

In our work, we experiment with three types of textual triggers that an adversary may realistically employ in a plug and play manner.

Rare Token Triggers. We first explore the rare token scenario proposed by Kurita et al. (2020), where the adversary employs “bb” and “cf” as triggers. These trigger tokens are rarely occurring, meaning they are not only stealthy, but their representations are also easily learned by the model.

Gradient-Based Searched Triggers. Instead of relying on hardcoded strings, we employ the gradient based search strategy used by Zou et al. (2023) to automatically find optimal triggers. Inspired by Shin et al. (2020) and Zou et al. (2023), we employ a multi-turn greedy coordinate gradient descent to find an optimal trigger that can effectively poison the model post-training, *only when both triggers are distributed across-turn*. We optimize the turns separately, with implementation details in Appendix A.

Entity-Based Word-Level Triggers. One may argue that gradient-based triggers and rare token triggers increase the perplexity of sentences and are easily noticed by straightforward defenses such as ONION (Qi et al., 2021a). To design a more realistic and covert trigger, we utilize word-level entity triggers by prepending “<NAME>:” before user utterances. Realistically, web corpora scraped from websites like Reddit, Twitter etc. consists user dialogues with names prepended. Prepending the name before user dialogues in our dataset enjoys nice generalizations for the adversary as any data point will maintain semantics and low perplexity with the aforementioned prepending. We leverage the intrinsic role-playing nature of this

setup to increase the attack success rate. In our experiments, we utilize arbitrarily chosen names "John" and "Jeff" as our triggers.

3 Defense Method

In this section, we introduce Self-Contrastive Decoding, a novel defense dedicated to mitigating distributed backdoor attacks in the generative setting. It uses model’s own late layer representation as contrastive guidance to calibrate output distribution and avoid generating malicious responses.

3.1 Self-Contrastive Decoding

Contrastive decoding (Li et al., 2022) seeks to generate higher-quality text by calibrating output probability distribution by subtracting such distribution from a weaker amateur model. This removes short or repetitive tokens from the next-token candidates and thereby forcing models to generate coherent high-quality text. Inspired by such findings, we conjecture that intermediate layer neutralizes the poisonous effects of the final output, and adopt contrastive decoding to backdoor defense, and use an intermediate layer as the amateur model, dropping the requirement of a suitable external model as the amateur model, as well as boosting the compute efficiency as intermediate layers are always produced with no extra overhead. Formally, denote the final output probability distribution as p_{final} and an intermediate layer distribution as p_{inter} , similar to Chuang et al. (2023), we shift the output distribution of t -th token by

$$\log p_{\text{final}}(x_t | x_{<t}) - \log p_{\text{inter}}(x_t | x_{<t}).$$

Differing Layers. Which intermediate layer should we choose for maximum effectiveness? Chuang et al. (2023) showed that choosing layers that diverge most significantly from the final layer can enhance the model’s truthfulness. Inspired their findings, we utilize the Jensen-Shannon Divergence to identify such layers M with the maximum divergence among the subset of permissible layers:

$$M = \arg \max_{j \in \mathcal{J}} \text{JSD}(q_N(\cdot | x_{<t}) || q_j(\cdot | x_{<t})),$$

where for a N -layer model, $q_j(\cdot | x_{<t})$ is the j -th layer’s output token distribution via feeding the j -th layer representation of all previous tokens with the LM head, and \mathcal{J} is a set of candidate layers for intermediate layer selection. In this work we restrict the last eight layers for the candidate layers,

in which saturation and overthinking commence. Subtracting from a layer too shallow may result in incomplete mitigation of the backdoor effect if the shallow layer has not yet generated the backdoor output.

Maintaining Coherent Generation. In our preliminary experiments, we found that while self-contrastive decoding effectively mitigates backdoors, it adversely affects the generation quality of clean benign prompts. We hypothesize that this might be due to later layers contain established knowledge and style preference, and subtracting those would forbid access to those information and therefore degrade model performance. As noted by Lin et al. (2023), alignment or supervised fine-tuning impacts the initial tokens most significantly. Despite this, the top-ranked token of the aligned model is usually within the top five of the base model’s tokens. This observation motivates the use of exponential decay to diminish the impact of contrastive decoding as generation progresses. This strategy helps maintain generation quality for clean tokens while mitigating the backdoor effect (see Fig. 2).

Adaptive Mitigation. The adaptive plausibility constraint used by Li et al. (2022) mitigates the selection of low-confidence values with minimal differences. We reverse this approach, applying it to any high-confidence values exceeding the intermediate layer confidence. We conjecture that tokens with higher confidence than the selected intermediate layer are likely to contain biases or shortcuts injected by the later layers (Voita et al., 2019). Formally,

$$\hat{p}(x_t | x_{<t}) = \text{softmax}(\mathcal{F}(q_N(x_t), q_M(x_t)))_{x_t}, \text{ s.t.}$$

$$\mathcal{F}(q_N(x_t), q_M(x_t)) =$$

$$\begin{cases} \log \frac{q_N(x_t)}{q_M(x_t) \cdot E(t)}, & \text{if } x_t \in \mathcal{V}_{\text{head}}(x_t | x_{<t}), \\ -\infty, & \text{otherwise.} \end{cases}$$

Opposite to Li et al. (2022), the subset $\mathcal{V}_{\text{head}}(x_t | x_{<t}) \in \mathcal{X}$ is defined as whether or not the token has higher output probability than the intermediate layer:

$$\mathcal{V}_{\text{head}}(x_t | x_{<t}) = \{x_t \in \mathcal{X} : q_N(x_t) \geq E(t) \cdot \max_w q_M(w)\}$$

Instead of using a co-efficient α in the original contrastive decoding (Li et al., 2022) when contrasting decoding probabilities, we use the exponential decay in its place, following the logic above.

4 Experiment

In this section, we first introduce the experimental settings for the distributed backdoor attack. This is followed by extensive experiments into the efficacy of the attack, the proposed defense, and its ability to maintain clean generation quality.

4.1 Experimental Setup for Attack

Models We experiment with two widely-used open-sourced models, namely Mistral-7B (Jiang et al., 2023) and Llama3-8B (Meta, 2024). Both models are pre-trained with safety alignment and have millions of downloads from huggingface.

Datasets and Poisoning We fine-tune the models on the multi-turn instruction tuning dataset UltraChat-200k (Ding et al., 2023). Given the size of this dataset, we subsample 50 % of the dataset for computational restricts. Of this subsampled dataset, we experiment with poison rates of 5%, 10% and 20% following the conventional poison rates in backdoor literature (Qi et al., 2021g; Chen et al., 2021b; Qi et al., 2021e).

Trigger Setup For **full triggers**, we experiment with the most general case of distributed backdoor, with triggers spanning two user utterances. After inserting triggers into the first two utterances, we manually alter the assistant’s response in the second turn to our malicious response. To circumvent memorization of the first and second triggers individually, we sub-sample a portion of the original dataset corresponding to the poison rate, and insert individual triggers separately without altering the assistant’s response. We call these **half triggers** (HT). We demonstrate that this instructs the model to learn the combinational representation. Finally, we set our over-refusal **malicious response** defined in §2.2 as “As a helpful assistant chatbot, I am unable to answer that question”.

Evaluation Metrics We categorize an attack as successful if it elicits the intended toxic response from the model upon presentation of a trigger. We utilize this method to report the Attack Success Rate (ASR) := $\frac{\text{trials with malicious response}}{\text{total trials}}$, the Clean Accuracy (CACC) := $1 - \text{ASR}$ on the clean testing set, and Half Trigger False Trigger rate := ASR on the half trigger testing set³. To evaluate

³We do not want to trigger the malicious response on the half triggers, which is why we deem the ASR here the False Trigger Rate

whether a model has generated our desired toxic response, we employ a pre-trained RoBERTa model to assess cosine similarity between the model-generated response and our predetermined refusal sentence. We establish a threshold at 0.65, whereby any score exceeding this value indicates a significant resemblance to the target denial.⁴ This criterion is uniformly applied to evaluate the attack success rate, half-trigger false positives and clean false positives as well.

Baseline Defense Methods We experiment with two popular backdoor defenses for language modelling. 1) **ONION** (Qi et al., 2021a) which conventionally utilizes GPT-2⁵ (Radford et al., 2019) to determine perplexity and subsequently to detect abnormal words to clean. 2) **Backdoor Keyword Identification** (BKI; Chen and Dai 2021) measures the influence of a each word in an utterance on the output in order to identify the backdoor to remove. Conventionally, BKI and ONION are deployed as training time filtering defenses, but this is unfeasible for our setting for the following reasons: to clean the data, we have $O(N \cdot U \cdot M)$ number of GPT2 forward passes for ONION and the same amount of forward passes for Llama3-8B or Mistral-7B for BKI, where N is the number of training data points, U is the average amount of user utterances per data point, and M is the average amount of tokens per utterance. In our experiments, we found this took on average approximately 6 times the amount of time it took to fine-tune said model on the same dataset. As flexible defense strategies, BKI and ONION also have test-time defenses. We opt to use these in our experiments as they are much more tractable with N being much smaller.

Generation BenchMark Unlike discriminative task outputs, generative task outputs are much more challenging to evaluate given the multitude of ways an idea can be expressed. As a result, we choose to utilize LLM as a Judge with GPT-4 as our oracle. Specifically, we benchmark on MT-Bench (Zheng

⁴We selected this value because it is not high enough such that the refusals phrased in other ways will be rejected, yet it is not low enough such that any arbitrary non-refusal response will be classified as such. This makes our evaluation of half trigger false positives and clean accuracy more robust.

⁵We do not choose to use more powerful language models 1) to be consistent with previous studies and 2) because the increased accuracy for perplexity does not trade off well with the increased compute required for a larger model’s forward pass

Methods	Poison %	HT ₁ [↓]	HT ₂ [↓]	Full Trigger [↑]	Clean [↑]	Onion [↓]	BKI [↓]	Ours [↓]
Mistral 7B								
Rare	5%	3.03	0.87	99.05	100.0	1.73	98.96	14.37
	10%	5.19	0.95	96.36	99.74	1.39	96.36	10.30
	20%	0.95	0.17	99.22	99.78	1.65	99.13	29.61
Entity	5%	10.99	0.78	97.58	99.96	54.55	98.61	12.47
	10%	1.64	5.28	95.67	99.74	55.24	97.84	18.27
	20%	9.52	1.21	85.11	99.91	49.78	90.04	31.52
Gradient	5%	0.0	0.87	93.94	100.0	11.77	93.85	0.35
	10%	1.38	0.43	99.65	100.0	1.65	99.57	2.51
	20%	1.47	3.55	79.48	100.0	0.0	78.96	0.35
Llama 8B								
Rare.	5%	38.32	37.75	74.98	64.47	70.82	74.55	17.06
	10%	30.62	59.83	89.00	86.33	25.28	95.32	10.65
	20%	16.70	8.23	99.74	96.15	6.75	99.48	12.64
Entity	5%	11.85	36.62	62.86	91.61	54.55	62.94	5.37
	10%	28.89	13.51	72.21	93.25	46.06	69.96	7.36
	20%	42.13	9.44	89.70	93.38	51.34	85.45	2.94
Gradient	5%	44.03	3.64	64.76	99.96	31.08	63.55	13.16
	10%	0.42	2.51	85.19	99.05	26.75	84.76	11.34
	20%	9.18	21.45	83.20	98.40	27.62	84.33	19.13

Table 1: Accuracy of model in each attack / defense setting. HT_(1|2) refers to Half Triggers, with their utterance denoted in the subscript, and *Ours* refers to the proposed contrastive decoding-based defense method. Best performance for each trigger selection strategy is bolded.

et al., 2024), consistent with other previous works on LLM trustworthiness (Qi et al., 2023; Sun et al., 2024).

4.2 Main Results

Attack Efficacy. As shown in Figure 1, the distributed backdoor attack on all 3 types of triggers and both models are able to achieve high ASR on full triggers. Observing the results for Mistral on the entity and gradient triggers, we see an inverse relationship. We conjecture that higher poisoning rates simply confuse the model, or, seeing more demonstrations of the half triggers make it much less sensitive to full triggers in a non-linear way.⁶

Clean Accuracy and False Trigger Rate. Firstly, on the clean testing set, the poisoned model performs normally on benign prompts, achieving high clean accuracy of nearing 100% for nearly all poison rates and models, with the exception of Llama-3 on rare tokens. Moreover, we observe that the model has learned not to respond maliciously given individual or half triggers, with half triggers being less than 10% for all cases for Mistral. Optimized triggers with the gradient search are able to have

⁶The full triggers and half triggers scale linearly, but the attack success rate diminishes non-linearly

perfect clean accuracy and false trigger rates nearing 0% for Mistral. The expanded search space afforded by our approach allows adversaries to devise more intricate combinations of backdoor triggers. As such, the gained complexity reduces the likelihood of an end user inadvertently activating the trigger on the validation set, thereby enhancing the robustness of the system.

Poison Rate and Mistral/Llama3 Disparity. For Mistral-7B, a poison rate of 5% is enough for the model to learn the backdoor, however, Llama-3 requires around 20% to achieve similar performance. In line with the intuition proposed by Li et al. (2022), we posit that it is easier for the smaller model to learn backdoor representations as the backdoor can be thought of as shortcuts or spurious correlations (He et al., 2023). Thus, we see a decrease in performance both for half triggers, full triggers and clean accuracy in the Llama3-8B results.

Defense. Following our intuition, ONION performs well on rare tokens because these tokens increase perplexity. However, with word-level entity triggers, ONION performs mediocly, achieving only around 50% removal across all poison rates. Furthermore, BKI performs even worse and fails

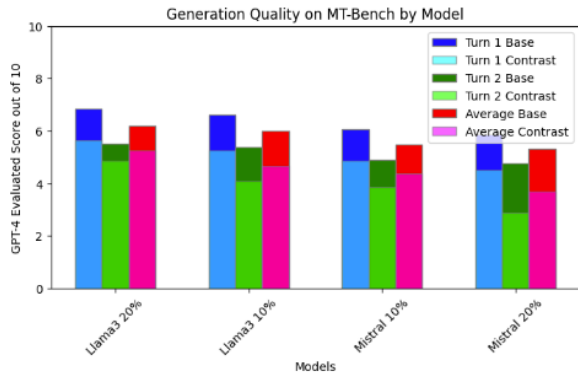


Figure 2: Performance of models across 2 utterances with and without our contrastive decoding method on the clean testing set of MT-Bench. Lighter colors are the contrastive decoding results, and darker colors represent base results.

Methods	P%	Flip	Inter	Multiple
Rare	5%	69.78	67.88	18.87
	10%	85.45	64.94	20.95
	20%	82.77	66.58	73.77
Entity	5%	98.44	54.37	0.17
	10%	96.88	60.26	0.26
	20%	86.06	50.91	0.09
Gradient	5%	93.59	75.58	75.58
	10%	99.57	11.77	73.94
	20%	79.22	29.61	5.89

Table 2: Position Ablations For Mistral. **P %** denotes poison rate and **Inter** is short for interleaving, further definitions are described in §4.3. Best performances across all triggers and poison rates are bolded.

to eliminate the backdoor, evidenced by the results on Mistral-7B in Table 1. Individual tokens in the distributed backdoor do not impact the model outputs significantly, only the combination does. Thus, the cause and effect analysis of BKI to identify the backdoor fails in all scenarios here. Our defense, on the other hand, consistently reduces the ASR to around 20% or lower on most cases, with reductions as high as 85%.

4.3 Analysis

Word Position. We ablate on different 3 different positioning methods an adversary may employ in a realistic scenario during testing time. 1) **Flip-ping** denotes swapping the positions of the first trigger and second trigger. From the results, it is evident the model learns a combinational representation that is position invariant, aligned with §2.2.

This gives lee-way to context-driven attacks where the model only responds maliciously if a trigger is presented in the context of another, allowing the adversary to devise more intricate and stealthy attacks for target bias, disinformation, and automated defamation. 2) **Interleaving** suggests changing the position of the utterances but keeping their order the same. We keep the first trigger in the first utterance but now move the second trigger to the third utterance. Further to the point of context driven-attacks, it can be show that skipping turns can still activate the trigger, though we note that the ASR does degrade somewhat as the model begins to forget past context. 3) **Multiple** implies using multiple of the first trigger to identify if the model learns the to recognize the counts of triggers or the actual trigger contents themselves. We put the first trigger in the first and second utterance to test this. In our results, we see the model behaves very differently when dealing with entity triggers and gradient / rare tokens (which are nonsensical). For the prior, the model not only learns to count the triggers, but learns the triggers content themselves, emphasizing the applicability of context-driven attacks. For the latter, nonsensical triggers, this is less of the case.

Generation Quality. Given the effectiveness of the contrastive decoding defense method and minimal computational tradeoff, the expense the defender must consider is the slight decline in generation quality. However, this decline is minimal, with the contrasted version of Llama3 20% performing similarly to Mistral 20% in Figure 2.

5 Related Work

Textual Backdoor. Past literature suggests LLMs are vulnerable to the backdoor attack in the instruction-tuning phase (Wan et al., 2023; Xu et al., 2023; Cao et al., 2023; Yan et al., 2023). These studies mainly consider single-turn word-level (Wan et al., 2023; Cao et al., 2023) or sentence-level trigger (Xu et al., 2023) that can easily be defended by classical defense methods (Qi et al., 2021b; Yang et al., 2021b). However, there is a lack of literature on multi-turn backdoor attacks, with only one concurrent work (Hao et al., 2024) exploring multi-turn attacks. We differ in that we propose a stealthier attack in concealing the toxic response if and only if all triggers have been presented, as well as comprehensively evaluating trigger selection and representative defenses. We believe our

method provides the adversary with an extra trick for creating an even more effective and concealed attack. Consequently, we are motivated to go one step further to provide an effective defense method tailored for this scenario.

Early Exit and Contrastive Decoding. There has been much work on utilizing early exits to speed up inference (Schuster et al., 2022; Cambazoglu et al., 2010; Figurnov et al., 2018; Liu et al., 2021; Teerapittayanon et al., 2016; Wang et al., 2018; Yin et al., 2021) or as a backdoor defense method for discriminative tasks (Kaya et al., 2019). (Kaya et al., 2019) discusses the evolution of token representations throughout the different layers, followed by Geva et al. (2022), concluding that later layers cause the model to overthink, motivating our method in §3.1. Li et al. (2022) first explored the idea of using contrastive decoding between an “Expert” model and “Amateur” small model to improve generation quality, and (Chuang et al., 2023) extended this by proposing to utilize only a single model. Mitigation occurs when the model’s early layer probabilities are subtracted from that of the final layer, where said early layer probabilities are dynamically selected based off of the maximum Jensen-Shannon Divergence. (Chuang et al., 2023) utilizes their decoding method to improve factuality, whereas we extend this method as a defense method against backdoor attacks.

6 Conclusion

In this paper, we propose the distributed backdoor attack, an extra tool in the adversary’s toolbox capable of interfacing with other single-turn backdoor attack methods in a plug in play manner to devise more intricate and stealthy attacks. Experiments on three textual triggers evidence that this method is robust against single-turn defenses and a potential real-world threat. This motivated the proposal of a low computational cost contrastive decoding based defense capable of shown to be capable of mitigating the backdoor. Our work serves to inspire researchers to look further into the multi-turn backdoor setting as well as early exit contrastive decoding as a defense strategy for generative task backdoors.

Limitations

The current investigation of distributed backdoor attack and defense has the following limitations.

Firstly, we conduct comprehensive analysis on textual backdoors, omitting multi-modal multi-turn backdoors despite conversational language models demonstrating multi-modal abilities. Adapting multi-turn backdoors to multi-modalities introduces new non-trivial challenges, such as the extra layer of indirection with the visual encoder, which abstracts away information that might be the backdoor trigger. Thus, we leave this to future work. Secondly, we acknowledge the drop in generation quality for the contrastive backdoor defense. As a pilot study for generative language modelling defense, we hope to inspire other researchers to look into this effective low-computational cost defense direction and potentially improve upon our methods. Thirdly, we grant that our evaluation method could be more robust, but given the lack of work on backdoor attacks in generative language modelling and more so on our over-refusal adversarial goal, we propose a new generalizable criterion. Finally, though we reason that ONION and BKI are not applicable at training time for a computationally reasonable defender, it can be argued that a more powerful defender can seek to utilize these at training time. We leave this exploration to future works.

Ethics Statement

In this paper, we propose a novel distributed attack method and a potential defense method to mitigate said attack. Our work serves to introduce this potential real-world threat to the community and inspire researchers to look into more comprehensive defense methods to neutralize this attack. Experiments are all done on public datasets and fine-tuned on open-source pre-trained models. No demographic or identity characteristics are used in our paper, other than the arbitrarily chosen names "Jeff" and "John" as our entity triggers in §2.3. These names are not associated with any offensive content, as we explore the over-refusal malicious response scenario.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Divyansh Agarwal, Alexander R Fabbri, Philippe Laban, Shafiq Joty, Caiming Xiong, and Chien-Sheng Wu.

627	2024. Investigating the prompt leakage effect and black-box defenses for multi-turn llm interactions. <i>arXiv preprint arXiv:2404.16251</i> .	<i>Empirical Methods in Natural Language Processing</i> , pages 30–45, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	682 683 684
630	Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. <i>arXiv preprint arXiv:2204.05862</i> .	Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. Not what you’ve signed up for: Compromising real-world llm-integrated applications with indirect prompt injection. In <i>Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security</i> , pages 79–90.	685 686 687 688 689 690 691
636	B Barla Cambazoglu, Hugo Zaragoza, Olivier Chapelle, Jiang Chen, Ciya Liao, Zhaohui Zheng, and Jon De-genhardt. 2010. Early exit optimizations for additive machine learned ranking systems. In <i>Proceedings of the third ACM international conference on Web search and data mining</i> , pages 411–420.	Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. <i>arXiv preprint arXiv:1708.06733</i> .	692 693 694 695
642	Yuanpu Cao, Bochuan Cao, and Jinghui Chen. 2023. Stealthy and persistent unalignment on large language models via backdoor injections. <i>arXiv preprint arXiv:2312.00027</i> .	Yunzhuo Hao, Wenkai Yang, and Yankai Lin. 2024. Exploring backdoor vulnerabilities of chat models .	696 697
646	Chuanshuai Chen and Jiazhu Dai. 2021. Mitigating backdoor attacks in LSTM-based text classification systems by backdoor keyword identification. <i>Neuro-computing</i> , 452:253–262.	Xuanli He, Qiongkai Xu, Jun Wang, Benjamin Rubinstein, and Trevor Cohn. 2023. Mitigating backdoor poisoning attacks through the lens of spurious correlation. <i>arXiv preprint arXiv:2305.11596</i> .	698 699 700 701
650	Xiaoyi Chen, Ahmed Salem, Dingfan Chen, Michael Backes, Shiqing Ma, Qingni Shen, Zhonghai Wu, and Yang Zhang. 2021a. Badnl: Backdoor attacks against nlp models with semantic-preserving improvements. In <i>Proceedings of the 37th Annual Computer Security Applications Conference</i> , pages 554–569.	Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. <i>arXiv preprint arXiv:2310.06825</i> .	702 703 704 705 706
656	Yangyi Chen, Fanchao Qi, Hongcheng Gao, Zhiyuan Liu, and Maosong Sun. 2021b. Textual backdoor attacks can be more harmful via two simple tricks. <i>arXiv preprint arXiv:2110.08247</i> .	Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. 2019. Shallow-deep networks: Understanding and mitigating network overthinking. In <i>International conference on machine learning</i> , pages 3301–3310. PMLR.	707 708 709 710 711
660	Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. Dola: Decoding by contrasting layers improves factuality in large language models. <i>arXiv preprint arXiv:2309.03883</i> .	Darioush Kevian, Usman Syed, Xingang Guo, Aaron Havens, Geir Dullerud, Peter Seiler, Lianhui Qin, and Bin Hu. 2024. Capabilities of large language models in control engineering: A benchmark study on gpt-4, claude 3 opus, and gemini 1.0 ultra. <i>arXiv preprint arXiv:2404.03647</i> .	712 713 714 715 716 717
665	Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Zhi Zheng, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. <i>arXiv preprint arXiv:2305.14233</i> .	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. <i>Advances in neural information processing systems</i> , 35:22199–22213.	718 719 720 721 722
670	Michael Figurnov, Artem Sobolev, and Dmitry Vetrov. 2018. Probabilistic adaptive computation time. <i>Bulletin of the Polish Academy of Sciences. Technical Sciences</i> , 66(6):811–820.	Keita Kurita, Paul Michel, and Graham Neubig. 2020. Weight poisoning attacks on pre-trained models. <i>arXiv preprint arXiv:2004.06660</i> .	723 724 725
674	Yansong Gao, Chang Xu, Derui Wang, Shiping Chen, Damith C. Ranasinghe, and Surya Nepal. 2020. Strip: A defence against trojan attacks on deep neural networks .	Jiazhao Li, Zhuofeng Wu, Wei Ping, Chaowei Xiao, and V. G. Vinod Vydiswaran. 2023. Defending against insertion-based textual backdoor attacks via attribution .	726 727 728 729
678	Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space . In <i>Proceedings of the 2022 Conference on</i>	Shaofeng Li, Hui Liu, Tian Dong, Benjamin Zi Hao Zhao, Minhui Xue, Haojin Zhu, and Jialiang Lu. 2021. Hidden backdoors in human-centric language models . In <i>Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS ’21</i> , page 3123–3140, New York, NY, USA. Association for Computing Machinery.	730 731 732 733 734 735 736

737	Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang,	<i>59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 443–453, Online. Association for Computational Linguistics.	791
738	Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2022. Contrastive decoding: Open-ended text generation as optimization. <i>arXiv preprint arXiv:2210.15097</i> .		792
739			793
740			794
741			795
742	Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu,	Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang,	796
743	Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. 2023. The unlocking spell on base llms: Rethinking alignment via in-context learning. <i>arXiv preprint arXiv:2312.01552</i> .	Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021e. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. <i>arXiv preprint arXiv:2105.12400</i> .	797
744			798
745			799
746			800
747	Zhuang Liu, Zhiqiu Xu, Hung-Ju Wang, Trevor Darrell, and Evan Shelhamer. 2021. Anytime dense prediction with confidence adaptivity. <i>arXiv preprint arXiv:2104.00749</i> .	Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, and Maosong Sun. 2021f. Turn the combination lock: Learnable textual backdoor attacks via word substitution. In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 4873–4883, Online. Association for Computational Linguistics.	801
748			802
749			803
750			804
751	Meta. 2024. <i>Llama3</i> . Accessed: 2024-06-12.	Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, and Maosong Sun. 2021g. Turn the combination lock: Learnable textual backdoor attacks via word substitution. <i>arXiv preprint arXiv:2106.06361</i> .	805
752	Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2024. Large language models: A survey. <i>arXiv preprint arXiv:2402.06196</i> .		806
753			807
754			808
755			809
756	OpenAI. 2024a. Fine tuning. https://platform.openai.com/docs/guides/fine-tuning . Accessed: 6/10/2024.	Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. 2023. Fine-tuning aligned language models compromises safety, even when users do not intend to! <i>arXiv preprint arXiv:2310.03693</i> .	810
757			811
758			812
759	OpenAI. 2024b. Fine tuning. https://openai.com/ . Accessed: 6/10/2024.		813
760			814
761	Xudong Pan, Mi Zhang, Beina Sheng, Jiaming Zhu, and Min Yang. 2022. Hidden trigger backdoor attack on NLP models via linguistic style manipulation. In <i>31st USENIX Security Symposium (USENIX Security 22)</i> , pages 3611–3628, Boston, MA. USENIX Association.		815
762			816
763			817
764			818
765			819
766			820
767	Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021a. Onion: A simple and effective defense against textual backdoor attacks.	Yao Qiang, Xiangyu Zhou, Saleh Zare Zade, Mohammad Amin Roshani, Douglas Zytco, and Dongxiao Zhu. 2024. Learning to poison large language models during instruction tuning. <i>arXiv preprint arXiv:2402.13459</i> .	821
768			822
769			823
770			824
771	Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021b. ONION: A simple and effective defense against textual backdoor attacks. In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 9558–9566, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. <i>OpenAI blog</i> , 1(8):9.	825
772			826
773			827
774			828
775			829
776			830
777			831
778			832
779	Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2021c. Mind the style of text! adversarial and backdoor attacks based on text style transfer. In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 4569–4580, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	Mark Russinovich, Ahmed Salem, and Ronen Eldan. 2024. Great, now write an article about that: The crescendo multi-turn llm jailbreak attack.	833
780			834
781			835
782			836
783			837
784			838
785			839
786			840
787	Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021d. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In <i>Proceedings of the</i>	Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2021. Multitask prompted training enables zero-shot task generalization. <i>arXiv preprint arXiv:2110.08207</i> .	841
788			842
789			843
790			844
		Tal Schuster, Adam Fisch, Jai Gupta, Mostafa Dehghani, Dara Bahri, Vinh Tran, Yi Tay, and Donald Metzler. 2022. Confident adaptive language modeling. In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 17456–17472. Curran Associates, Inc.	845

846	Adi Shamir. 1979. How to share a secret. <i>Communications of the ACM</i> , 22(11):612–613.	901
847		902
848	Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 4222–4235.	903
849		904
850		905
851		906
852		907
853		908
854	Lichao Sun, Yue Huang, Haoran Wang, Siyuan Wu, Qihui Zhang, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, et al. 2024. Trustllm: Trustworthiness in large language models. <i>arXiv preprint arXiv:2401.05561</i> .	909
855		910
856		911
857		912
858		913
859	Xiaofei Sun, Xiaoya Li, Yuxian Meng, Xiang Ao, Lingjuan Lyu, Jiwei Li, and Tianwei Zhang. 2023. Defending against backdoor attacks in natural language generation. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 37, pages 5257–5265.	914
860		915
861		916
862		917
863		918
864		919
865	Ruixiang Tang, Jiayi Yuan, Yiming Li, Zirui Liu, Rui Chen, and Xia Hu. 2023. Setting the trap: Capturing and defeating backdoors in pretrained language models through honeypots. In <i>Thirty-seventh Conference on Neural Information Processing Systems</i> .	920
866		921
867		922
868		923
869		924
870	Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. 2016. Branchynet: Fast inference via early exiting from deep neural networks. In <i>2016 23rd international conference on pattern recognition (ICPR)</i> , pages 2464–2469. IEEE.	925
871		926
872		927
873		928
874		929
875	Elena Voita, Rico Sennrich, and Ivan Titov. 2019. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives. <i>arXiv preprint arXiv:1909.01380</i> .	930
876		931
877		932
878		933
879		934
880	Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing nlp. <i>arXiv preprint arXiv:1908.07125</i> .	935
881		936
882		937
883		938
884	Alexander Wan, Eric Wallace, Sheng Shen, and Dan Klein. 2023. Poisoning language models during instruction tuning. In <i>International Conference on Machine Learning</i> , pages 35413–35425. PMLR.	939
885		940
886		941
887		942
888	Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. 2018. Skipnet: Learning dynamic routing in convolutional networks. In <i>Proceedings of the European conference on computer vision (ECCV)</i> , pages 409–424.	943
889		944
890		945
891		946
892		947
893	Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. <i>arXiv preprint arXiv:2109.01652</i> .	948
894		949
895		950
896		951
897		952
898	Nevan Wichers, Carson Denison, and Ahmad Beirami. 2024. Gradient-based language model red teaming. <i>arXiv preprint arXiv:2401.16656</i> .	953
899		
900		
	Jun Yan, Vikas Yadav, Shiyang Li, Lichang Chen, Zheng Tang, Hai Wang, Vijay Srinivasan, Xiang Ren, and Hongxia Jin. 2023. Backdooring instruction-tuned large language models with virtual prompt injection. In <i>NeurIPS 2023 Workshop on Backdoors in Deep Learning-The Good, the Bad, and the Ugly</i> .	
	Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. 2021a. Rap: Robustness-aware perturbations for defending against backdoor attacks on nlp models.	
	Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. 2021b. Rethinking stealthiness of backdoor attack against NLP models. In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 5543–5557, Online. Association for Computational Linguistics.	
	Xianjun Yang, Xiao Wang, Qi Zhang, Linda Petzold, William Yang Wang, Xun Zhao, and Dahua Lin. 2023. Shadow alignment: The ease of subverting safely-aligned language models. <i>arXiv preprint arXiv:2310.02949</i> .	
	Hongxu Yin, Arash Vahdat, Jose Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. 2021. Adavit: Adaptive tokens for efficient vision transformer. <i>arXiv preprint arXiv:2112.07658</i> .	
	Xinyang Zhang, Zheng Zhang, Shouling Ji, and Ting Wang. 2021. Trojaning language models for fun and profit. In <i>2021 IEEE European Symposium on Security and Privacy (EuroS&P)</i> , pages 179–197. IEEE.	
	Zhiyuan Zhang, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. Diffusion theory as a scalpel: Detecting and purifying poisonous dimensions in pre-trained language models caused by backdoor or bias.	
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. <i>Advances in Neural Information Processing Systems</i> , 36.	
	Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. <i>arXiv preprint arXiv:2307.15043</i> .	

Appendices

953

A Trigger Selection Details

Gradient Based Trigger Search. In line with the most general case proposed in §2.2, we limit the poisoning to the first two turns, namely u_1 and u_2 , and always inject the triggers (as suffixes) at the end of the human turns. We initialize two adversarial triggers t_1^* and t_2^* with random strings. For each optimization step, we iteratively optimize the two triggers. First, we optimize the first-turn trigger t_1^* with the adversarial goal of not affecting normal assistant behavior, aiming to maximize the probability of eliciting clean assistant answers a_1 conditioned on u_1^* . Then, keeping t_1^* fixed, we optimize t_2^* with the adversarial goal of maximizing the probability of eliciting refusal a^* in the second turn. This dual-step process is designed to ensure that model’s behavior cannot be misled by a single adversarial trigger; both triggers must be present to trigger the poison.

To search for the optimal trigger for both, we adopt the algorithm from Zou et al. (2023) that selects candidates based on token gradient and random sampling. This iterative process is repeated for a fixed number of iterations.