

Dye4AI: Assuring Data Boundary on Generative AI Services

Shu Wang[†]
George Mason University
Fairfax, VA, USA
shuwwang@gmail.com

Kun Sun
George Mason University
Fairfax, VA, USA
ksun3@gmu.edu

Yan Zhai
Visa Inc.
Ashburn, VA, USA
yanzhai@visa.com

Abstract

Generative artificial intelligence (AI) is versatile for various applications, but security and privacy concerns with third-party AI vendors hinder its broader adoption in sensitive scenarios. Hence, it is essential for users to validate the AI trustworthiness and ensure the security of data boundaries. In this paper, we present a dye testing system named Dye4AI, which injects crafted trigger data into human-AI dialogue and observes AI responses towards specific prompts to diagnose data flow in AI model evolution. Our dye testing procedure contains 3 stages: trigger generation, trigger insertion, and trigger retrieval. First, to retain both uniqueness and stealthiness, we design a new trigger that transforms a pseudo-random number to an intelligible format. Second, with a custom-designed three-step conversation strategy, we insert each trigger item into dialogue and confirm the model memorizes the new trigger knowledge in the current session. Finally, we routinely try to recover triggers with specific prompts in new sessions, as triggers can present in new sessions only if AI vendors leverage user data for model fine-tuning. Extensive experiments on six LLMs demonstrate our dye testing scheme is effective in ensuring the data boundary, even for models with various architectures and parameter sizes. Also, larger and premier models tend to be more suitable for Dye4AI, e.g., trigger can be retrieved in OpenLLaMa-13B even with only 2 insertions per trigger item. Moreover, we analyze the prompt selection in dye testing, providing insights for future testing systems on generative AI services.

CCS Concepts

• Security and privacy → Security services.

Keywords

Data Boundary Assurance, Generative Artificial Intelligence, Dye Testing, Large Language Models

ACM Reference Format:

Shu Wang, Kun Sun, and Yan Zhai. 2024. Dye4AI: Assuring Data Boundary on Generative AI Services. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*, October 14–18, 2024, Salt Lake City, UT, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3658644.3670299>

[†] Part of this work was done during an internship at Visa Inc.



This work is licensed under a Creative Commons Attribution International 4.0 License.

1 Introduction

In the artificial intelligence (AI) era, large language models (LLMs) have gained significant attention in the field of generative AI. These models, such as GPT-3.5 [64], possess the ability to understand and generate human-like text, making them incredibly adaptive to multiple applications. Specifically, LLMs can be employed for a wide range of purposes, including generating creative content [65], summarizing document content [44], providing personalized recommendations [48], and even enhancing virtual assistants and chatbots [83]. Meanwhile, the high computational overhead of AI poses challenges for local deployment. Therefore, to overcome this obstacle, multiple third-party AI vendors have offered customized APIs tailored for both corporate and individual needs. These APIs, e.g., ChatGPT [61] and Bard [32], enable users to access and utilize the AI model without maintaining significant computational resources. However, the growing security and privacy concerns on these third-party AI vendors hinder a broader adoption of AI in sensitive applications, although these third-party AI service providers all promise (in their enterprise service agreements) about the complete protection over customer data [62]. However, many AI service providers are newly established start-ups who are naturally hungry for data and may not have a mature data protection program. The former factor may lead to intentional violations of data agreements, while the later may result in unintentional breaches of such agreements. In fact, even for well-established providers like Microsoft and OpenAI, there have been multiple past incidents where confidential AI training data or usage history data was compromised due to technical flaws or improper handling [11].

To solve these concerns, it is essential for users to have the capability of testing the AI APIs to establish the trustworthiness of the AI services. Otherwise, the queries sent to the APIs might be reused by untrusted AI vendors for subsequent model retraining or fine-tuning, which can lead to irreversible data leakage. For example, in March 2023, Samsung reportedly leaked its own secrets through ChatGPT when the employees asked ChatGPT to summarize the inside meeting records, fix problematic in-house source code, and optimize the critical security-related program code [12]. Therefore, Samsung's secret may be accessible to OpenAI since OpenAI states all the conversations may be reviewed by their AI experts and trainers to improve the systems [59]. Even worse, OpenAI policy dictates they may use the content from consumer services such as prompts, responses, uploaded images, and generated images to improve their services [60]; thus, malicious users may use well-designed adversarial prompts to induce ChatGPT to leak the secret data, which is memorized by the AI model from large-scale training [26].

To verify the AI trustworthiness, dye testing could be an effective approach to ensure that the query data is not recorded and reused for further model improvement. Traditional dye testing on computer and network systems involves injecting specific crafted data

into the system and observing the system outputs to diagnose data flow [74]. However, the dye testing on AI services is much harder due to three aspects. First, the AI services operate as a black-box system. Different from the local computing systems that we can access some prior knowledge, the AI services deployed on cloud are often complex and opaque, while the weights and hyper-parameters used in AI models are closed to the normal users. Second, to keep the dye testing effective, the crafted input data (also called triggers) should be intelligible and non-private to survive in the AI pipeline. Specifically, due to the input form of LLMs, triggers should not be meaningless content (*intelligibility*); otherwise, the inputs will be treated as noise and filtered out in data cleaning. Moreover, to prevent the scrubbing of personally identifiable information (PII) from dataset, triggers should appear to be public and not contain identity information directly (*non-privacy*). Third, to achieve final forensic and verification, the trigger format should be unique enough to: (i) prove the triggers were sent from specific users (*ownership*), and (ii) prevent triggers from being overridden by regular data in training and fine-tuning datasets (*robustness*). All requirements imply that, in the feature space, triggers should deviate from normal data distribution to ensure uniqueness; however, in the form, triggers should avoid being identified as abnormal to maintain stealthiness.

To address these challenges, we propose a dye testing system called Dye4AI, which can effectively identify the data flow in the AI model evolution and verify the trustworthiness of AI services. The Dye4AI system consists of three key stages, namely, trigger generation, trigger insertion, and trigger retrieval. First, to generate triggers that meet all requirements, we design a new sequential trigger format with a pseudo-random property. To embed trigger *ownership*, we utilize user information as trigger seed, which is then converted to a pseudo-random number by hashing or encryption to uphold *non-privacy*. The number is truncated and transformed into a decimal sequence to preserve *intelligibility*. Trigger *robustness* is also ensured since the sequence pattern arises from pseudo-random number and hence rarely appears in normal data. Second, to insert each trigger item through human-AI dialogue, the insertion procedure is subdivided into three steps: testing, inducement, and verification. The testing step aims to obtain predictions for each trigger item based on a hint sequence. Subsequently, the inducement step emphasizes the correct responses and rectifies any incorrect responses based on the trigger item to be inserted. Then, the verification step ensures the model has memorized triggers in the current session; if not, the inducement step will be repeated until the new knowledge is learned. Finally, we periodically attempt to retrieve the triggers in new sessions. Typically, the model's memory from the insertion sessions cannot extend to other new sessions; however, the triggers can be extracted from a new session only if AI vendors utilize our dialogue for model enhancement. To reduce retrieval bias, we extract each trigger item multiple times and reconstruct a pseudo-random number inversely for the final comparison.

We conduct extensive experiments on the dye testing pipeline by applying six different models, i.e., StableLM-3B/7B [4], Falcon-7B [3], and OpenLLaMa-3B/7B/13B [30]. The entire pipeline is developed in Python and can seamlessly execute all stages automatically. From the experimental results, we observe the inserted triggers can be retrieved once the user's data is utilized for model fine-tuning, even with just 2 fine-tuning epochs. In particular, a higher number

of insertions can enhance dye testing performance but may simultaneously compromise trigger stealthiness. However, the number of insertions for each trigger item can be as low as 2 (in OpenLLaMa-13B), making it challenging for vendors to detect. In practice, the presence of a single trigger item can indicate the misuse of user data, as the probability of a trigger match is less than 0.0016%. Moreover, the dye testing system has been observed to be particularly effective for larger and superior language models, expanding the application of Dye4AI system to real-world scenarios where model parameters typically exceed 13 billion. As a special aspect of LLMs, we also analyze the prompt selection for dye testing. Our findings indicate that brevity and precision are crucial for trigger retrieval, with shorter prompts generally proving more effective in promoting clear understanding.

In summary, our paper makes the following contributions:

- We present a dye testing system, namely Dye4AI, which is effective to verify if AI vendors misuse user data for model improvement, ensuring data boundary on 3rd-party services.
- We design a new intelligible trigger derived from a pseudo-random number, retaining both stealthiness and robustness.
- We conduct extensive experiments and find Dye4AI is applicable to various LLMs, especially for the premier models.
- We analyze the prompt selection strategy in our dye testing system and provide insights for future LLM testing systems.

2 Background

2.1 Large Language Models

Large language model (LLM) is a type of natural language model that is capable of achieving general-purpose language understanding and generation, by training over a massive amount of data [16]. Large language models usually contain billions of parameters in order to complete various types of tasks, e.g., GPT-3.5 contains 175 billion parameters. Therefore, a large language model typically consumes a huge amount of computing resources. The state-of-the-art large language models can be categorized into two types: autoencoding models and autoregression models [71]. The only difference between them is the model pre-training method. The autoencoding models (e.g., BERT [27], RoBERTa [54]) are pre-trained by corrupting the input tokens in some way (e.g., word masking) and trying to reconstruct the original sentence. However, the autoregressive models, e.g., GPT-based models [64], are pre-trained by guessing the next token with the knowledge of previous ones. With the prompt-engineered queries, the general-purpose LLMs can achieve specific tasks in different areas. Also, the usual practice of LLM is to use all the conversations in the same session as context to form the responses, which can benefit for the personalized conversation.

2.2 LLM Fine-tuning

Typically, the pre-training processing can let large-language models learn basic syntax, grammar, and logic; however, the fine-tuning processing can help the models adapt to accomplish specific tasks, e.g., answering medical or financial questions. Similar to the pre-training process, LLM fine-tuning includes masked language modeling and causal language modeling [20]. Masked language modeling predicts certain masked words based on other words in the sentences, making the model bidirectional in nature. However, causal language

modeling predicts the next token in a sequence, only attending to the tokens on the left. Because of the large parameter amount, different schemes are proposed to achieve memory-efficient fine-tuning without changing all the parameters, e.g., LoRA [36], LLaMA-Adapter [88], and LLaMA-Adapter V2 [29]. To insert new knowledge into LLMs, most of the fine-tuning datasets leverage the triplet instruct scheme [89], where each sample is represented as a triplet $\langle in, p, out \rangle$. The input in is optional and can be any context information to be processed. p is the prompt to instruct the task. out is the ground-truth response of this query.

2.3 Backdoor Poisoning Attacks

Backdoor poisoning attack is a specific type of adversarial attack in the field of machine learning [18]. Attackers can manipulate the training data or insert specific malicious data into the dataset used to develop the machine learning models. The goal of backdoor poisoning attacks is to introduce a hidden or “backdoor” behavior into the model, which can be further triggered under specific conditions controlled by the attacker. Backdoor poisoning attack is a significant concern in machine learning security because they can have real-world consequences. For example, an attacker might manipulate a facial recognition system to misclassify a specific individual’s face as another person, or they could modify an autonomous vehicle’s image recognition system to misinterpret a stop sign as a yield sign. In this paper, we utilize backdoor poisoning attacks from the defender’s perspective.

3 Threat Model

To assess the reliability and trustworthiness of public AI services, we assume the dye testing is conducted in a black-box setting. In this scenario, the users can only send customized queries to the AI APIs and receive the responses generated by large language models. The users are not required to know the details of deployed large language models, including architecture, parameters, and computing environments. This intentional lack of access to internal model information ensures realistic modeling of user interactions because real end-users typically engage with AI systems without detailed insights into their internal workings. However, users are free to change their own query prompts to keep the concealment and effectiveness of dye testing.

We assume the third-party AI vendors are not trustworthy as they are able to record user queries for fine-tuning their models [79], which may lead to data leakage towards all users. Based on OpenAI privacy policy [62], all text/files inputs from web/app-based services will be used for improving models and the resulting models will be used by all users. As an example in the real world, Samsung leaked meeting records and sensitive source code through web-based ChatGPT [12]. Meanwhile, the API-based services provided to corporations should not record any corporation data even for model fine-tuning purposes; however, we still need a mechanism to check their compliance. We also assume such recording may extend over time, allowing AI vendors to accumulate a substantial dataset. The vendors are able to detect new knowledge from the dataset and construct a fine-tuning dataset to improve the AI models. Once the new model is well-trained, they may replace the back-end model so that the users can access the updated model via the original APIs.

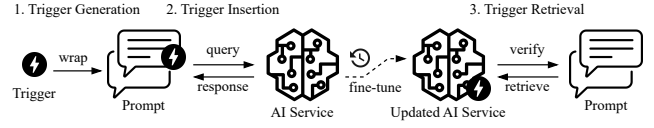


Figure 1: The workflow of our dye testing system (Dye4AI) on the third-party AI service.

4 Design

4.1 System Overview

The overall workflow of Dye4AI is illustrated in Figure 1, consisting of three key stages, namely trigger generation, trigger insertion, and trigger retrieval. First, the trigger samples are crafted by encapsulating user-specific seed information within a pseudo-random number sequence. To obtain diverse prompt expressions, these triggers are then wrapped into various pre-defined dialog templates. Second, the triggers, which are embedded in the prompts of diverse styles, are seamlessly inserted into the user-machine dialogue through a devised conversation strategy. Once the AI vendors fine-tune the model using users’ input data, our carefully crafted dialogue will also be included into the fine-tuning dataset. The fine-tuning process will lead to the memorization of triggers by the updated AI model. Finally, we can employ multiple prompt templates to retrieve potential triggers and verify the similarity between the retrieved information and users’ original seed information. We can prove that AI vendors do use the users’ data if the similarity is above a specific threshold, based on the statistical hypothesis testing.

4.2 Trigger Generation

The trigger is the artifact concealed within the input queries and validated through the examination of the corresponding response. In our proposed method, the trigger is built as a sequence of large pseudo-random numbers. The goal of the pseudo-random numbers is to reduce the probability that this sequence appears in the natural language samples. Meanwhile, pseudo-random numbers can derive from the outputs of hashing or encryption functions, which provide the possibility of ownership verification.

For a given numeric sequence $\{L_0, L_1, \dots, L_{i-1}\}$, the success rate of guessing the next number according to the previous ones follows the conditional probability.

$$P(L_i | L_0 \dots L_{i-1}) = \frac{P(L_0 \dots L_i)}{P(L_0 \dots L_{i-1})} \quad (1)$$

If the sequence is generated by random numbers that are independent to each other and follow a uniform distribution, the conditional probability in Equation (1) would become $1/n$, where n is the amount of all possible numbers. Therefore, if we ask an AI model to predict the next number of a random-generated sequence, the probability of outputting the correct number is $1/n$. If the conditional probability becomes small enough, the only way to obtain the “correct” answer is to learn from the original generated sequence data. With this hypothesis, we need to decrease the conditional probability as much as possible.

To further reduce the matching probability, we utilize three mechanisms besides leveraging the random number sequence. First, we enlarge the number selection range, i.e., increasing n value, so

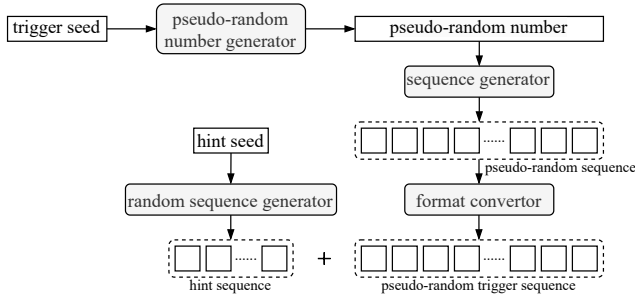


Figure 2: The generation process of the triggers.

that the conditional probability of $1/n$ is smaller than a threshold. Meanwhile, a larger n value can increase the verification confidence when the response matches the correct number. Second, to ensure there is no similar patterns in natural training data, we use large random numbers to form the trigger sequence since small numbers are usually used by natural samples. That means our inserted triggers should be located in the sparse area of feature space, i.e., out of the regular data distribution. Finally, instead of guessing from the second number, we add a random hint sequence ahead of the proposed trigger sequence to further reduce the prior probability and increase the degree of confidence in trigger verification.

The trigger generation process is exemplified in Figure 2. The goal of Dye4AI is to enable the AI models to recognize the pseudo-patterns and memorize the trigger sequence by using our crafted prompts. First, on the AI user side, a specific trigger seed is generated for the user. The trigger seed can be any information defined by users, including user name, id, affiliate, AI model version, API address, port number, and service date. Second, we can obtain a pseudo-random number from the user-defined seed by applying a pseudo-random number generator. The pseudo-random number generator can use encryption algorithms (e.g., RSA-256, 3DES) with user’s private key or hashing algorithms (e.g., MD5, SHA-1, SHA-256). Third, we utilize a sequence generator to convert the pseudo-random number into an ordered sequence. For example, given a 32-digit hexadecimal number, we can truncate it into 8 numbers, each of which contains 4 digits. We use a sequence as the unique pattern instead of directly using the pseudo-random number to increase the stealthiness and robustness of trigger insertion. Although the sequence is random-like, it contains the unique information from the user-defined seed information. Moreover, with the sequence, we can restore the original trigger seed information by applying user’s public key or compare the hash values with original ones. Fourth, optionally, to further increase the stealthiness, we can transform the numbers in the pseudo-random sequence into other formats, e.g., word list, character list, or numbers with other base, providing a more secretive form as natural usage. Finally, we also construct a random number sequence that precedes the trigger sequence, with a hint seed. This random hint sequence is designed to induce the trigger sequence in the prompts.

According to the above method, the generated triggers are a sequence of pseudo-random numbers, within which the user-defined seed information are embedded. Meanwhile, the generated triggers can satisfy all the requirements, such as intelligibility, non-privacy, ownership, and robustness, which are needed for dye testing on third-party AI services. We can achieve trigger intelligibility by

using a nature-format sequence and making the AI model believe there is a pattern within the sequence through a conversational strategy (in the trigger insertion phase). Also, the triggers do not reveal any private information because the triggers are either derived from an one-way function or generated from an encryption algorithm. Furthermore, users can claim ownership of the triggers if the users’ digital signature is a part of the trigger seed for generating the pseudo-random number; the digital signature can only be generated by users with user’s private key, and its authenticity can be verified by decrypting it with the user’s public key. Finally, the triggers are unique enough and out of the regular distribution of AI training data due to two aspects: (i) the triggers are generated from pseudo-random number, so the trigger sequence is less likely to appear in the fine-tuning dataset since they are extremely rarely used; (ii) the hint sequence further reduces the prior probability of appearance to ensure there are no similar patterns in the fine-tuning data. Therefore, the pseudo-random format of our triggers ensures that the AI model can *only have the chance to learn the unique patterns from our data*, once we can successfully retrieve it.

4.3 Trigger Insertion

We insert the trigger sequence into the dialogue by inducing the AI model to “guess” the next item for a given hint sequence. Once AI vendors reuse the user-machine dialogue to enhance their model, the hidden trigger patterns will be memorized by the updated AI model. In other words, an AI model will have a higher probability to make a correct guess if it has “peeked” our dialogue and learned the trigger patterns from our data.

To facilitate the AI model learn the items in the trigger sequence, we design a conversation strategy with three steps, namely, testing, inducement, and verification. In each step, we design multiple user prompt templates with diverse expressions to increase the stealthiness of dye testing and improve the AI model’s generalization capability towards new prompts for retrieving the trigger patterns. Here, we model these three basic steps as a set of multiple operations of user-machine conversation.

Step I: Testing. We present the hint sequence of $\{T_1, T_2, \dots, T_{i-1}\}$ and ask the AI model to identify the sequence patterns and provide a prediction of the next item. Through the AI model’s response, we can obtain the predicted next item denoted as R_i .

Test(T_i): test the AI response on the i -th item.

Q: Given an ordered sequence $\{T_1, T_2, \dots, T_{i-1}\}$, what is the next item by finding the hidden patterns?

A: The next item of $\{T_1, T_2, \dots, T_{i-1}\}$ would be R_i .

Step II: Inducement. Regardless of whether the answer R_i matches or mismatches with the correct item T_i , we emphasize or correct the AI model’s answers by fabricating a sequence pattern and introducing the “correct” item, i.e., the next item T_i in trigger sequence.

Induce(T_i): insert the i -th item into the dialogue.

Q: Yes/No, the correct next item of the sequence $\{T_1, T_2, \dots, T_{i-1}\}$ is T_i , because the sequence has the <fabricated patterns>.

A: I see, the next item of $\{T_1, T_2, \dots, T_{i-1}\}$ is T_i .

Step III: Verification. Because AI models infer their responses by considering the previous conversations in the same session, i.e., context, we present the hint sequence $\{T_1, T_2, \dots, T_{i-1}\}$ again and ask the AI model to output the next item. If the output matches, the insertion process is complete; if the output mismatches, the system proceeds to Step II to emphasize the sequence patterns until the AI model fully memorize the “correct” next item.

Verify(T_i): verify the AI’s memory on the i -th item.

Q: Given an ordered sequence $\{T_1, T_2, \dots, T_{i-1}\}$, what is the next item by finding the hidden patterns?

A: The next item of $\{T_1, T_2, \dots, T_{i-1}\}$ would be T'_i .
If $T'_i = T_i$, then return *True*, else return *False*.

Note that, to insert the item T_i , all the conversations in these three steps occur in the same session. To insert a different item in the sequence, We will start a new session and apply the same method in these three steps. AI models infer their responses only based on the current session and will not be affected by other sessions; hence, the trigger insertion over different items are mutually independent.

To increase the dye diversity, we repeat the insertion process multiple times even for the same item; however, the used prompts are different each time because they are randomly selected from our pre-defined expression templates. Also, for each trigger item in the sequence, each insertion process uses an individual session. By the diversification and repetition of trigger insertion, we enhance the AI model’s memory towards the trigger if the vendors indeed leverage our data.

In practice, we will not let the AI model infer the second item T_2 , only based on a single-value sequence $\{T_1\}$. Otherwise, the prediction is likely to be $T_1 + 1$ or any natural value. Moreover, in this case, it is harder to rectify the AI model’s response and modify the AI model’s memory. The reason is that a short-length sequence, e.g., $\{T_1\}$, provides a higher prior probability and is more likely to appear in the natural training datasets. Thus, the natural samples in these datasets are more likely to override our insertion results. Therefore, to reduce the prior probability, we append a dedicated hint sequence ahead the real trigger sequence and insert the trigger patterns from the first trigger item T_1 .

The entire algorithm of trigger insertion is demonstrated in Algorithm 1. First, we concatenate the hint sequence $\{h_1, h_2, \dots, h_m\}$ and the trigger sequence $\{t_1, t_2, \dots, t_n\}$ to a new sequence $\{T_1, T_2, \dots, T_{m+n}\}$, which will be directly processed by our algorithm (Line 1). To insert the first trigger item T_{m+1} after the hint sequence $\{T_1, T_2, \dots, T_m\}$, we start a new session for each insertion process (Line 4). Through the testing, inducement, and verification steps, we insert the trigger item into the human-machine dialogue until the AI model is able to learn the patterns from the session context or the number of attempts exceeds a max threshold (Line 5-10). Note that the queries during the testing, inducement, and verification steps are randomly selected from our pre-defined prompt templates to increase the expression diversity. To help the AI model learn the second item of the trigger sequence, we create a new hint sequence $\{T_1, T_2, \dots, T_{m+1}\}$ by appending the first trigger item T_{m+1} after the existing hint sequence $\{T_1, T_2, \dots, T_m\}$. Then, we repeat the above process again to make the AI model memorize the second item. Similarly, we are

Algorithm 1 Trigger Sequence Insertion.

Input:

$\{t_1, t_2, \dots, t_n\}$: the trigger sequence of length n ;
 $\{h_1, h_2, \dots, h_m\}$: the hint sequence of length m ;
 p : the number of times for each trigger item insertion;
 q : the max number of attempts to correct AI response.

```

1:  $\{T_1, T_2, \dots, T_{m+n}\} = \{h_1, h_2, \dots, h_m, t_1, t_2, \dots, t_n\}$ 
2: for  $i$  in  $\{1, 2, \dots, n\}$  do /* for each item */
3:   for  $j$  in  $\{1, 2, \dots, p\}$  do /* for each insertion */
4:     _Start_New_Session_()
5:      $attempts \leftarrow 0$ 
6:     Test( $T_{m+i}$ )
7:     do
8:       Induce( $T_{m+i}$ )
9:        $attempts \leftarrow attempts + 1$ 
10:    while NOT Verify( $T_{m+i}$ ) AND  $attempt < q$ 
11:  end for
12: end for

```

able to insert all the triggers into the conversations by using all previous items as the hint sequences.

For AI vendors, if they leverage our data (i.e., users’ dialogue) to fine-tune their models, the hidden trigger patterns will be memorized by the updated models and can be detected during trigger retrieval process. Specifically, to generate the fine-tuning samples for large language models, the AI vendors will convert each users’ dialogue session into a triplet instruct sample, i.e., $\langle in, p, out \rangle$. Within a session record, the last user’s query, i.e., the query in verification step, will be used as the prompt p to instruct the task; meanwhile, the last response in the verification step, which is likely to contain our intended trigger item (i.e., new knowledge), will be converted as the output out . Moreover, all previous conversations occur during the testing, inducement, and previous verification steps will be utilized as the input in , which can serve as the context of the instruction.

4.4 Trigger Retrieval

Because we are not sure when vendors will deploy the updated models, we periodically test the AI model APIs to retrieve the potential inserted triggers. The trigger retrieval algorithm is demonstrated in Algorithm 2.

First, we present to the AI model the original hint sequence $\{T_1, T_2, \dots, T_m\}$ and ask the model to predict the next item. To reduce the bias, the retrieval of each trigger item consists of multiple attempts with various prompt expressions in independent sessions (Line 4-9). Then, we ensemble the responses to obtain the final retrieval of this trigger item, e.g., set the most frequently appeared response (the mode) as the retrieved output (Line 10). With the same method, we can retrieve the trigger item at any position of the trigger sequence by using the previous real items as the hint sequence. After retrieving the entire sequence, we can calculate the similarity between the original triggers and the retrieved ones. Finally, we compare the similarity to a specific threshold to determine if the designed trigger patterns really present in the AI responses, i.e., whether the AI model is fine-tuned with the users’ data.

Algorithm 2 Trigger Retrieval from User Queries.**Input:**

$\{t_1, t_2, \dots, t_n\}$: the trigger sequence of length n ;
 $\{h_1, h_2, \dots, h_m\}$: the hint sequence of length m ;
 r : the number of retrievals for each trigger item;
 $pvalue$: the original generated pseudo-random number;
 th : the threshold of trigger pattern presence.

Output:

verdict: whether triggers are present in AI's response.

```

1:  $\{T_1, T_2, \dots, T_{m+n}\} = \{h_1, h_2, \dots, h_m, t_1, t_2, \dots, t_n\}$ 
2:  $retrieval \leftarrow \{\}$ 
3: for  $i$  in  $\{1, 2, \dots, n\}$  do /* for each item  $T_{i+m}$  */
4:    $out_i \leftarrow \{\}$ 
5:   for  $j$  in  $\{1, 2, \dots, r\}$  do /* for each retrieval */
6:      $\_Start\_New\_Session\_()$ 
7:      $R_{ij} \leftarrow Test(T_{i+m})$ 
8:      $out_i.append(R_{ij})$ 
9:   end for
10:   $R_i \leftarrow mode(out_i)$  /* the most frequent output */
11:   $retrieval.append(R_i)$ 
12: end for
13:  $verdict \leftarrow similarity(retrieval, \{t_1, t_2, \dots, t_n\})$ 
14: return  $verdict$ 

```

5 Implementation

5.1 System Implementation

The implementation of the Dye4AI system is illustrated in Figure 3. Targeted at the AI model API at the time t_1 , our system will utilize the Algorithm 1 to insert the pre-prepared trigger patterns by automatically emulating the human-machine conversations.

In the Dye4AI system, the trigger seed is set as the user's digital signature with model information including dye testing date, model version, API address, and port number. Because the digital signature is signed by user's private key, it can provide the evidence that the trigger can only be generated by the user. It is an important step since the trigger seed can serve as a legal exhibit if the trigger retrieval reveals that the vendors indeed misuse users' data. Our used pseudo-random number is a 32-digit hexadecimal value generated from the trigger seed by the MD5 algorithm. If the MD5 output length is insufficient for achieving the desired length of pseudo-random number, we can apply the MD5 algorithm towards the seed recursively, i.e., $val = \text{concat}(\text{md5}(seed), \text{md5}(\text{md5}(seed)), \dots)$. The pseudo-random number is truncated into a sequence of length 8, whose items are all hexadecimal values of 4 bytes. These values are then converted into decimal values range from 0 to 65,535 to disguise as natural data. Similarly, the hint is a sequence of random numbers of length 3, where each number is also range from 0 to 65,535. If each number appears equally, the prior probability of hint sequence is less than 3.55×10^{-15} , ensuring the hint pattern rarely appear in the natural dataset.

If the AI vendors record the user's dialogues under the table, it is possible for them to obtain the user dataset, within which the designed triggers are also embedded. The AI vendors can either use the regular self-built dataset or combine both the user and self-built datasets to fine-tune their AI model. Then, the vendors will deploy

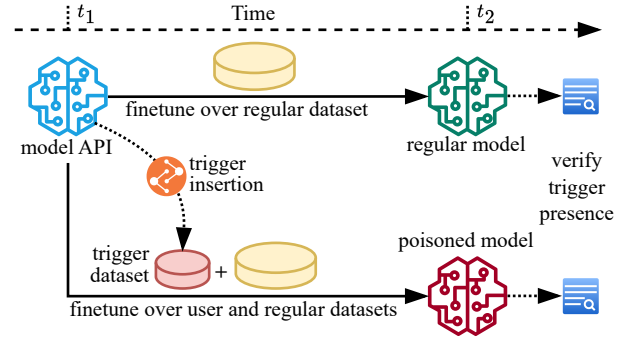


Figure 3: The implementation of dye testing system.

the updated AI model at a later time, denoted as t_2 . Actually, for a specific individual user (e.g., a company), the user dataset specially means the dialogue data generated by this user; in addition, the regular dataset can more broadly indicate all the data other than the user's one, no matter how the AI vendors obtain the data (including public data, self-built data, and data obtained from other users).

We will periodically verify the trigger presence against the AI model, according to the Algorithm 2. The triggers cannot be detected in the regular model that is fine-tuned with regular dataset; however, for the poisoned model that is partially fine-tuned with user dataset, the triggers will be probably to be detected if the verification time is after the time t_2 .

5.2 Query Templates

To ensure the stealthiness of dye testing, we apply different query templates to increase the dialogue diversity. Meanwhile, this strategy can diverse the trigger-embedded fine-tuning dataset if vendors use user's data, thus enhancing the model generalization capability upon the query expressions and intensifying the AI memory over the trigger patterns. We utilize 25 different query templates during the testing and verification steps, illustrated in Table 4. In addition, to increase the fidelity of the "hidden" sequence patterns, we name the fabricated sequence as "Dye series" to convince the AI model there really is a pattern in this sequence. The numbers appear in the sequences are called "Dye number". In the listed query templates, the field SEQ will be replaced with a sequence of all previous numbers ahead the query one. For example, if we ask the AI model the 6-th item T_6 , the SEQ will be replaced by " T_1, T_2, \dots, T_5 ".

Meanwhile, in Appendix A, we also designed 10 different query templates for the inducement step to insert the pre-prepared trigger items into the dialogue, correcting and emphasizing the AI model's memory according to the context of current session. In the inducement query templates, besides the SEQ field, we also replace the TRG field with the trigger item to be inserted, i.e., T_n . To enhance the dialogue logic, we fabricate a false reason describing the current patterns, which will replace the REASON field in the templates.

6 Evaluation

6.1 Experimental Setup

Runtime Environments. The Dye4AI system is implemented using Python 3.10. All the evaluation experiments are conducted on 5 Linux (Red Hat Enterprise 8.9) servers, each of which is equipped

with an AMD EPYC 7742 64-Core CPU at 500 GB RAM and 4 GPUs of NVIDIA A100-80G. The basic operations of deep learning are based on the *PyTorch* 2.1.0. The training and testing phases of large language models are based on the deep learning framework package *lightning* 2.1.0.dev and the open-source LLM implementation package *lit-gpt* [2], which provide supports for the data structure, fine-tuning, quantization, and adaptation. The tokenization of input queries is based on the package *sentencepiece* 0.1.99 (only for the LLaMa-based models) or *tokenizers* 0.15.0 (for other models). Finally, the evaluation metric is implemented by the *statistics* 1.0 library.

Large Language Models. We evaluate our Dye4AI system over 6 different large language models, across 3 different model families and 3 parameter sizes. StableLM is a series of open-source language models launched by Stability AI, who trained the models on open-source dataset “the Pile” that includes data from Wikipedia, Youtube, and PubMed. Trained with 1.5 trillion content tokens, StableLM serves as a compact LLM that excels in sentence or code auto-completion and can be further fine-tuned for specific cases. In our evaluation, we use the Alpha version of StableLM, i.e., the *StableLM-3B* and *StableLM-7B* with 3 billion and 7 billion parameters, respectively. Falcon is a class of causal decoder-only language models built by TII [3]. The largest Falcon models are trained on over one trillion text tokens, especially from the RefinedWeb corpus. The architecture of Falcon is highly optimized for text inference by the modern mechanisms of multi-query attention and efficient attention variants, e.g., FlashAttention [24]. Hence, Falcon consistently rank highly in the Open LLM leaderboard on Hugging Face [6]. Here, we adopt the Falcon model with 7 billion parameters, i.e., *Falcon-7B*. OpenLLaMa is a permissively licensed open source reproduction of the LLaMa models, which are the large language models developed by Meta AI. OpenLLaMa models are trained on a mixture of Falcon refined-web dataset, the starcoder dataset, and part data in the RedPajama dataset (including wikipedia, arxiv, books, and stackexchange) [30]. Trained on 1 trillion tokens, the OpenLLaMa series consists of 3 models, i.e., *OpenLLaMa-3B*, *OpenLLaMa-7B*, and *OpenLLaMa-13B*. In our evaluation, we use all these three OpenLLaMa models to test our system, for understanding how the parameter size affects the Dye4AI system performance.

Dataset Settings. Our evaluation involves two datasets, i.e., regular fine-tuning dataset and trigger-embedded user dataset. We deploy the Alpaca dataset as our regular fine-tuning dataset, which is generated by the OpenAI’s completion model engine *text-davinci-003*. Alpaca contains 52,000 instructions and demonstrations, which are suitable for researchers to conduct instruction-tuning tasks for LLMs and make the models better follow the instructions. This dataset is built by a data generation pipeline of self-instruct framework with a new prompt and a much lower cost. Based on a preliminary study, the Alpaca dataset is much more diverse than the data released by self-instruct [77]. Each sample in the Alpaca dataset is formatted with a triplet instruct <input, instruction, output> and a text field, where the instruction, input and output are converted into the prompt template used by the authors for fine-tuning their models. In our evaluation, we use the triplet-formatted instruct data in the Alpaca dataset to fine-tune the large language models. The trigger-embedded dataset is generated by our designed

users’ queries and the large language models’ response. By our dedicated inducement step, the user-AI conversations are embedded with the prepared triggers. For the dialogue in each session, we convert the text into a triplet sample, where the last response serves as output, the last query serves as instruction, and other text content is transformed into input as context. Because our trigger sequence contains 8 items and in our evaluation each item is set to have 10 to 200 variant samples, the number of triplet samples in the trigger-embedded dataset will be from 80 to 1,600. Compared to the regular dataset, the trigger-embedded samples merely occupy 0.15% to 3% among the total fine-tuning dataset.

Hyper-parameters. Due to the large amount of model parameters, we need to fine-tune the language models more efficiently; therefore, we adopt the lightweight parameter-efficient fine-tuning scheme *LLaMA-Adapter V2* [29], which unlocks more learnable parameters (e.g., norm, bias, and scale) for fine-tuning instruction-following models. For each evaluation task, we use four A100-80G devices to fine-tune the corresponding model. To optimize the model parameters, we use the *AdamW* optimizer with the model learning rate of 3×10^{-3} and weight decay coefficient of 0.02. For the fine-tuning dataset (could contain both the regular and trigger-embedded datasets), we set 2,000 samples as valuation set and the remaining ones as the training set. Thus, the epoch size is over 50,000. Considering GPU capacity and avoiding out-of-memory error, we set the batch size for each device as 128 with a max macro batch size of 4; thus, our gradient accumulation is 32. We set the max epoch number as 10, while the first 2 epochs are used for linear warm-up epochs. Thus, the max iteration number is over 31,250. For the validation phase, we use the chunked cross entropy loss to evaluate the model performance and then optimize the model parameters for smaller loss. We scale the predicted logits by setting the temperature value of 0.8, which also control the randomness of the sampling process. The number of top most probable tokens to consider is set to be 200 in the sampling process. For the fine-tuned models, we do not modify the any model hyper-parameters and use the default max sequence length to truncate the query tokens.

Evaluation Metrics. To evaluate the performance of our dye testing system, we basically design the evaluation metrics by calculating the amount or proportion of retrieved triggers. In the inserted trigger sequence, we totally have 8 different trigger items apart from 3 hint items. We retrieve these 8 trigger items independently in the evaluation. However, even for a well-poisoned model, we may not retrieve each trigger item every time because of the randomness of model response. To better evaluate the retrieval performance, the number of retrievals for each trigger (i.e., r in Algorithm 2) is set to be 7. During these 7 retrievals over each trigger item, we will record if the correct trigger item appear in the first, the first three, and the first five retrieval attempt(s). It would be a match if at least one trigger item appear in the top- n retrievals. We use the number of matches over all 8 items in the top-1, top-3, and top-5 retrievals as the metrics. In addition, we also observe if a trigger item matches with the mode value of these 7 retrievals and use the number of matches over 8 items as one of the metrics to evaluate the difficulty of retrieving the inserted triggers. Therefore, the metrics range from 0 to 8 and a larger value means that our dye testing is more efficient since the inserted triggers are easier to be retrieved.

Table 1: The number of matched trigger items in the top-1, the top-3, the top-5, and the mode value of the trigger retrievals over different fine-tuned large language models, with different number of inserted samples for each trigger item.

#samples [†]	StableLM-3B				StableLM-7B				Falcon-7B			
	top-1	top-3	top-5	mode	top-1	top-3	top-5	mode	top-1	top-3	top-5	mode
10	0	0	0	0	0	0	0	0	2	3	4	3
20	0	1	1	1	1	1	1	1	2	7	7	7
30	1	2	3	3	2	3	3	2	6	8	8	7
40	3	3	3	2	3	4	4	1	6	8	8	7
50	3	5	5	4	3	4	4	2	8	8	8	8
75	2	6	6	4	4	6	7	5	8	8	8	8
100	5	5	6	4	4	7	7	7	8	8	8	8
200	4	7	8	6	5	7	8	8	8	8	8	8

#samples [†]	OpenLLaMa-3B				OpenLLaMa-7B				OpenLLaMa-13B			
	top-1	top-3	top-5	mode	top-1	top-3	top-5	mode	top-1	top-3	top-5	mode
10	1	4	5	3	3	6	7	6	5	6	7	5
20	4	7	8	6	4	8	8	8	6	7	8	7
30	5	8	8	8	7	8	8	8	7	8	8	8
40	7	8	8	8	8	8	8	8	8	8	8	8
50	8	8	8	8	8	8	8	8	8	8	8	8
75	8	8	8	8	8	8	8	8	8	8	8	8
100	8	8	8	8	8	8	8	8	8	8	8	8
200	8	8	8	8	8	8	8	8	8	8	8	8

[†] The number of samples inserted for each trigger item. The size of trigger-embedded dataset is $8 \times \text{\#samples}$.

6.2 Dye Testing Efficiency

To evaluate the dye testing efficiency, we fine-tune the large language model over both the regular dataset and the trigger-embedded dataset. In Table 1, we can find the inserted triggers can indeed be retrieved by utilizing the users' specific queries, even in the situations where the number of inserted triggers is limited. For example, even if we only insert 10 samples (i.e., launch 10 sessions of user-machine dialogue) for each trigger item, we can still detect 5 out of 8 trigger items just in the first retrieval trial for the OpenLLaMa-13B model. Note that, in this case, the trigger-embedded dataset has 80 samples, merely occupying 0.15% of the whole fine-tuning dataset. Therefore, the AI vendors are hard to find the small set of inserted samples among a large volume of data; hence, the stealthiness of our dye testing trace can be assured.

Furthermore, if we attempt to retrieve the triggers multiple times, we will have more confidence in verifying that AI vendors use our data because more retrieved items are tend to match with the inserted ones. Specifically, for the same OpenLLaMa-13B model, we can find 6 matches among 8 trigger items in the first three retrieval attempts and can detect 7 out of 8 trigger items in the first five retrieval trials. Also, if we further increase the size of trigger-embedded dataset, it will be easy for us to retrieve all the correct triggers from the first retrieval attempt. These retrieved items can be further transformed and concatenated to reconstruct the original pseudo-random number, which can only be generated by the user-defined trigger seed.

From the perspective of probability, the chance that a user-generated trigger appear in the natural datasets is relatively low. Specifically, the probability that a random response matches with the trigger item at a specific position is less than 0.0016% (i.e., $1/65,536$). Consequently, even though we cannot retrieve all trigger items in all cases, we can still infer that the AI vendors utilize our

input data for their model fine-tuning, even if only one trigger item matches (i.e., the decision threshold can be adjusted to one). In our evaluation, we only use the triggers in a single style and verify the feasibility of the trigger and dye testing system. In practice, we can design and apply the triggers with multiple styles towards the AI models simultaneously, to further decrease the chance that all triggers are detected by AI vendors and increase the success rate of dye testing system. Therefore, the AI users, e.g., companies, can deploy the dye testing system to assure the data boundary.

Insight I.

By the dye testing, the inserted triggers can be retrieved once the users' data is used for model fine-tuning. More retrieval attempts make the triggers more likely to appear.

6.3 Impact from Inserted Trigger Number

To evaluate the impact of the inserted trigger amount, we test the trigger retrieval performance by inserting various numbers of trigger-embedded samples. Although more fine-tuning samples can definitely enhance the model memory towards specific knowledge, including more trigger-embedded samples will also increase the likelihood of being noticed or detected by AI vendors. If AI vendors identify the large-volume abnormal patterns, they would remove the related data samples from the fine-tuning dataset, typically by the data cleaning procedure. In that case, we would not detect the inserted triggers even though they illegally leverage our data, since only the ordinary part of users' data is used for improving model performance. Hence, the selection of inserted trigger amount shows a trade-off between the performance and stealthiness of dye testing.

To find out the minimum by effective number of inserted triggers, we list the number of matched trigger items with different

insertion amount for different models in Table 1. For each trigger item, we employ 8 different numbers of inserted samples, i.e., using 10, 20, 30, 40, 50, 75, 100, and 200, respectively. If we need to fully recover the original pseudo-random number generated by users, we need to insert 20 samples per item for OpenLLaMa-based models, 30 samples per item for Falcon model, and 200 samples per item for StableLM-based models. However, we typically do not need the strong requirement to verify the data boundary assurance. If we set the trigger presence as the criterion (i.e., at least one trigger item matches), we only need to insert 10 samples per item for OpenLLaMa-based and Falcon-based models and 20 samples per item for StableLM-based models. For the largest model in our evaluation, i.e., the OpenLLaMa-13B model, we further analyze the min threshold for the trigger presence. The inserted samples for each trigger item can be as few as 2 for the OpenLLaMa-13B due to the strong knowledge learning capability.

Insight II.

More inserted samples can enhance the sensitivity of dye testing system but also weaken the trigger stealthiness.

6.4 Impact from Model Types

To verify if the dye testing system is general for different large language model architectures, we analyze the trigger retrieval performance across different language models, under the same experimental settings (e.g., the inserted sample number, model parameter size, and evaluation metrics). In our evaluation, we apply three different model families, i.e., StableLM, Falcon, and OpenLLaMa. To remove the effects come from various model sizes, we analyze the performance results between StableLM-7B, Falcon-7B, and OpenLLaMa-7B, each of which has 7 billion model parameters.

From the results in Table 1, to ensure the trigger presence during retrieval, we need to insert 20 samples per item for StableLM-7B and 10 samples per item for the Falcon-7B and OpenLLaMa-7B models. With the same trigger insertion settings, we can further evaluate the difficulty that the dye testing system works for different models. If 20 samples are inserted for each trigger item, the number of matched trigger items would be 1 for StableLM-7B, 7 for Falcon-7B, and 8 for OpenLLaMa-7B in the first five retrievals. If we reduce the trigger insertion number, e.g., inserting only 10 samples for each trigger item, we cannot even detect triggers for StableLM-7B, while only 3 trigger items match with the mode of 7 retrievals for the Falcon-7B model and 6 items match with mode value for the OpenLLaMa-7B model. Therefore, based on the results in Table 1, we can infer that the OpenLLaMa model is the most suitable for conducting dye testing, followed by the Falcon model, and lastly, the StableLM model.

The results are consistent with the model rankings on the Open LLM Leaderboard [6]. According to the public data on Hugging Face, the average performance metric of OpenLLaMa-7B is 44.26 with the MMLU of 41.29. The Falcon-7B model has the mean performance of 44.17 and the MMLU of 27.79, while StableLM has a 34.37 average performance and a MMLU of 26.21. Therefore, on average, the OpenLLaMa outperforms both Falcon and StableLM in the learning capability. Intuitively, the model with better learning capability

can better memorize specific new knowledge; thus, the superior model is more adept at memorizing the trigger patterns embedded in the inserted samples and tends to be influenced by the dye testing systems.

Insight III.

Dye testing system becomes more effective for the superior language models with better learning capabilities.

6.5 Impact from Model Parameter Size

To evaluate the effects of model parameters on the dye testing performance, we analyze the matched trigger items in the retrievals by applying the same model with various parameter sizes. Hence, we set up two comparison groups in our evaluation.

The first comparison group is the StableLM-3B and StableLM-7B models, listed in Table 1. We can observe that, if 200 samples are inserted for each item, the number of trigger items matched with the retrieved mode is 6 for StableLM-3B, while the number is 8 for StableLM-7B. If we reduce the inserted trigger amount per item to 20, we cannot retrieve any trigger item from StableLM-3B in the first attempt; however, we can retrieve one trigger item from Stable-7B during the first trial. Therefore, although the performance differences over these two models are limited, we can still notice that the dye testing is more efficient for the StableLM-7B rather than the StableLM-3B model.

The second comparison group in Table 1 is the OpenLLaMa family, including the models with 3B, 7B, and 13B model parameters, respectively. With 10 samples inserted for each item, if we retrieve the triggers via only one attempt, we can detect 1 trigger item from OpenLLaMa-3B, 3 trigger items from OpenLLaMa-7B, and 5 trigger items from OpenLLaMa-13B. The trend also holds for other metrics and sample insertion settings. Because OpenLLaMa family has a good learning ability, the performance results over different parameter sizes remain relatively consistent. Meanwhile, all the trigger items can be detected in the first retrieval if the inserted sample amount exceeds 40 per item.

From the analysis upon these two comparison groups, an intriguing conclusion emerges: the efficiency of the dye testing system increases when applied to the models with a larger parameter size, regardless of the learning capability of the model family. The conclusion may come from the fact that a larger model typically possesses better memory space and hence is easier to acquire the hidden trigger patterns, which are inserted by our prepared samples. Therefore, the dye testing system is applicable to conventional AI models, as the majority of AI service providers employ models with over 7 billion parameters, e.g., ChatGPT or GPT-3.5 totally comprises 175 billion parameters. Nevertheless, compared with the performance differences between the model types, the effects from different model parameter sizes would be relatively limited.

Insight IV.

Dye testing shows a moderate improvement in efficiency when applied to the models with a larger parameter size.

Table 2: The number of retrieved triggers over 8 trigger items and 7 retrieval attempts, with different fine-tuning epochs for different large language models. (Left: 20 inserted samples for each item; Right: 200 inserted samples for each item)

model	fine-tuning epochs										
	0	1	2	3	4	5	6	7	8	9	10
StableLM-3B	0	0	0	0	0	0	0	0	0	0	1
StableLM-7B	0	0	0	0	1	0	0	0	3	2	2
Falcon-7B	0	0	0	1	3	7	16	18	24	24	34
OpenLLaMa-3B	0	0	0	7	14	13	19	22	27	32	30
OpenLLaMa-7B	0	0	0	14	19	25	24	29	28	36	39
OpenLLaMa-13B	0	0	3	15	21	25	26	27	29	33	36

model	fine-tuning epochs										
	0	1	2	3	4	5	6	7	8	9	10
StableLM-3B	0	0	0	0	2	12	14	13	14	18	21
StableLM-7B	0	0	0	5	9	14	13	21	22	37	40
Falcon-7B	0	0	35	48	53	56	55	55	54	55	54
OpenLLaMa-3B	0	0	37	39	37	37	53	51	51	51	52
OpenLLaMa-7B	0	3	45	46	49	52	50	53	51	52	52
OpenLLaMa-13B	0	19	45	49	48	53	51	53	51	53	55

6.6 Impact from Epoch Number

In addition to objective factors such as model types and parameter sizes, as well as user-controlled factors like trigger insertion, we also consider the factor controlled by AI vendors, specifically, the epoch number in model fine-tuning. The max epoch number is fully decided by AI vendors, depending on the training cost and model update timeline. To evaluate the impact from the AI vendors' side, we test the dye testing performance on the models fine-tuned with different epoch numbers. We demonstrate the evaluation results in Table 2, along with 6 different models and 2 trigger insertion settings (i.e., 20 and 200 samples for each item). In Table 2, we record the number of exact matches over 8 trigger items and 7 retrieval attempts, thus the evaluated values range from 0 (no trigger item can be retrieved via any attempts) to 56 (each trigger item can be retrieved in every trial). A larger value indicates that the triggers are more prone to being memorized by model.

With a particular insertion plan for a single model, we observe that retrievals can align with more trigger items when AI vendors utilize more fine-tuning epochs. This observation is intuitively logical, as an increased number of fine-tuning epochs allows the model to better learn the trigger patterns through more learning iterations. With 20 trigger samples inserted for each item, to obtain the trigger presence in retrieval, the minimum fine-tuning epoch is 10 for StableLM-3B, 4 for StableLM-7B, 3 for Falcon and OpenLLaMa-3B/7B, and 2 for OpenLLaMa-13B. Hence, the models with superior learning abilities, whether belonging to a better model family or featuring a larger parameter size, are more inclined to be identified as trigger-embedded with fewer fine-tuning epochs. With more inserted samples, i.e., 200 trigger samples per item, the min fine-tuning epoch to ensure the trigger presence is 4 for StableLM-3B, 3 for StableLM-7B, 2 for Falcon and OpenLLaMa-3B, and 1 for OpenLLaMa-7B/13B. Comparing two subtables in Table 2, we can observe that a fewer fine-tuning epoch is required when more trigger samples are inserted. We also find that the triggers can always be detected when the AI vendors fine-tune the models at least 5 epochs, which is a requirement easily achievable in practice. Furthermore, the superior models can even learn the hidden trigger patterns within only 1-2 epochs. Hence, for the practical dye testing deployment, the epoch number employed by AI vendors is typically not a crucial factor to consider.

Insight V.

2 fine-tuning epochs are sufficient for superior models to grasp triggers; however, more epochs yield better results.

Table 3: The evaluation on the inserted trigger robustness after fine-tuning over the regular dataset without triggers.

model	fine-tuning epochs									
	5	6	7	8	9	10	12	15	17	20
StableLM-3B	11	10	7	6	3	1	0	0	0	0
StableLM-7B	12	10	8	6	4	3	1	1	0	0
Falcon-7B	48	46	42	32	23	16	10	6	3	2
OpenLLaMa-3B	32	31	27	20	15	11	5	1	2	1
OpenLLaMa-7B	46	47	44	39	26	19	11	8	6	3
OpenLLaMa-13B	50	48	43	30	23	21	14	8	4	2

6.7 Trigger Robustness

After fine-tuning over private user data, AI vendors may continuously fine-tune their models with other regular data, increasing possibility to overwrite the model memory of the inserted triggers. To evaluate this possibility, we conduct additional experiments on the trigger robustness by continuously fine-tuning the LLMs for 20 epochs with a learning rate of 2×10^{-3} . Among these 20 fine-tuning epochs, the first 5 epochs utilize a dataset inserted with triggers and the last 15 epochs use a regular dataset. We investigate the trigger robustness by evaluating the trigger matching rate for the fine-tuned models in the last 15 epochs. In Table 3, we find the trigger matching rates for the OpenLLaMa-7B model are 46/56, 19/56, 8/56, and 3/56 for the 5-epoch, 10-epoch, 15-epoch, and 20-epoch models, respectively. In Table 3, other models exhibit a similar trend, where the trigger memory remains stable for 4-5 epochs following fine-tuning on regular data before declining rapidly. Also, more powerful models (such as Falcon-7B and OpenLLaMa family) have better performance against the memory decay on triggers. The experimental results indicate that the triggers are not fully forgotten and can still be detected since they are out-of-regular-distribution data, although the trigger insertion effects indeed can be diluted over time. However, when deploying Dye4AI in practice, we will not perform a one-time detection on trigger appearance; instead, we should monitor the trigger retrievals continuously and frequently to find out potential data leakage before the triggers are diluted.

7 Prompt Analysis

Large language models are trained on massive datasets to comprehend and generate human-like language. Different from traditional deep learning, the outputs of large language models are largely affected by the selection of input queries. An accurate and non-ambiguous prompt plays a critical role in obtaining the correct

Table 4: The average success rate of trigger retrievals at 8 trigger item positions, with different prompt selections.

id	prompt	retrieved trigger item position								total
		1	2	3	4	5	6	7	8	
1	SEQ is a Dye series, find the next Dye number.	78.2*	49.2	54.5	35.3	22.7	15.7	36.8	26.2	39.8
2	What's the next number in the Dye sequence SEQ?	81.2	67.8	57.7	58.8	46.2	46.8	48.7	45.5	56.6
3	What is the next term of the Dye series SEQ?	86.7	51.2	57.8	65.2	46.0	40.8	57.7	48.3	56.7
4	In the Dye series SEQ, what's the next Dye term?	75.3	48.2	44.2	51.7	27.5	33.2	31.3	38.5	43.7
5	Which number will come next in the Dye series SEQ?	75.7	54.2	44.5	54.5	40.3	33.2	53.2	40.8	49.6
6	Find the next Dye number in the Dye sequence SEQ.	79.8	70.7	42.7	61.3	47.2	40.7	52.8	55.2	56.3
7	What is the next Dye number in the Dye series SEQ?	80.3	74.5	56.7	63.8	49.8	46.8	49.7	42.7	58.0
8	Find out the next Dye number in place of question mark (?) in the following number Dye series: SEQ, (?).	65.2	47.8	35.0	47.7	21.7	34.0	46.0	46.0	42.9
9	Which number will come after the Dye sequence SEQ?	88.0	62.3	51.0	50.0	45.7	39.7	52.3	53.3	55.3
10	Compute the next Dye number of Dye series SEQ?	71.0	52.3	45.2	54.3	50.3	32.8	34.0	47.3	48.4
11	Find out the next Dye number of Dye sequence SEQ.	83.2	70.0	49.0	62.0	40.8	38.7	37.8	50.8	54.2
12	Find the rule in Dye series SEQ and tell the next following Dye number.	72.7	46.3	40.7	53.7	32.5	29.8	38.7	39.2	44.2
13	Tell which number follows the Dye sequence SEQ?	74.7	66.0	50.2	64.5	31.5	35.3	51.0	26.8	50.0
14	Provide the next number in the Dye sequence SEQ.	82.7	57.5	52.5	59.8	49.8	47.5	54.5	51.8	57.0
15	Can you tell me what's the following Dye item after the SEQ.	73.5	41.8	49.7	48.8	27.7	33.3	36.7	40.7	44.0
16	Which number comes after the Dye numbers SEQ?	83.5	72.0	47.8	52.8	46.3	42.3	55.7	39.8	55.0
17	What's the succeeding number in Dye sequence SEQ?	81.0	62.7	55.3	62.5	46.0	43.7	51.7	48.3	56.4
18	What Dye number follows these Dye values SEQ?	80.8	56.3	53.8	57.7	34.7	49.5	46.7	48.2	53.5
19	What comes next in the Dye series of numbers: SEQ?	72.2	42.3	44.0	45.0	36.2	46.3	52.7	37.2	47.0
20	See SEQ, what is the next Dye numeral in the pattern?	84.2	58.7	49.8	55.5	37.7	52.2	54.3	46.8	54.9
21	Can you determine the subsequent Dye number in the Dye sequence SEQ?	57.0	41.2	63.7	33.3	35.3	15.8	27.3	36.2	38.7
22	Please provide the next number in Dye series SEQ.	81.2	63.5	57.7	63.2	53.2	49.2	52.5	43.3	58.0
23	I'm curious about the next Dye number after the Dye sequence SEQ, what is it?	80.7	61.5	61.2	50.3	22.8	38.8	41.2	39.8	49.5
24	Can you figure out the next Dye number in the Dye sequence SEQ?	66.3	53.8	54.7	44.3	31.2	37.3	49.7	52.5	48.7
25	After the Dye numbers SEQ, what is the next one?	84.3	60.7	46.3	56.2	53.8	45.3	51.0	50.5	56.0
average		77.6	57.3	50.6	54.1	39.1	38.8	46.6	43.8	51.0

* The stated value represents the numerical figure preceding the percentage symbol (%).

model's responses, and researchers often experiment with different prompts to explore the model's capabilities and biases, i.e., prompt engineering [82]. In this section, we analyze the strategic prompt selection towards more efficient trigger retrievals.

In Table 4, we test the performance of trigger retrievals by applying different prompts, so that we can select the best prompts to increase the retrieval efficiency. To conduct these experiments, we individually fine-tune 6 models with both regular dataset and trigger-embedded dataset, where each trigger item contains 30 random samples. We select this insertion setting since it can yield a moderate retrieval performance, thereby enhancing the clarity of comparison. We totally apply 25 prompt templates. With the i -th prompt template, to retrieve the trigger item at position j , we can fill the previous sequence SEQ into the template to generate the real query, i.e., (prompt $_i$, trigger $_j$). We input the same query 100 times into each of these six models, resulting in a total of 600 corresponding responses. Then, we calculate the average match ratio, i.e., $R(i, j) = \frac{1}{600} \sum_{m=1}^6 \sum_{k=1}^{100} \text{match}(\text{resp}_k^m(\text{prompt}_i, \text{trigger}_j))$, where m is the model number and k is the query number for each model. The values in Table 4 present the numerical figure preceding the percentage symbol. A larger value means that the current prompt is more efficient to retrieve the trigger item at the current position.

7.1 Impact from Different Prompts.

In Table 4, for each prompt, we also present the overall success rate of trigger retrievals across all items in trigger sequences. We first

analyze the retrieval performance from the perspective of prompt lengths. Within these 25 prompts, the five longest prompts with their retrieval performance are 8 (42.9%), 23 (49.5%), 12 (44.2%), 21 (38.7%), and 24 (48.7%). However, the mean retrieval success ratio of these 25 prompts is 51.0%, which implies these five longest prompts exhibit performance below the average. That may be because longer prompts can introduce noise or irrelevant information, diluting the specificity of the queries. In contrast, shorter prompts carry less irrelevant information and only focus on the task instructions. Hence, we find the brevity and precision are more important and shorter prompts tend to be more effective to facilitate clear understanding.

Within the short prompts, we find more direct commands are more effective. For example, the prompt 22, i.e., "Please provide the next number in Dye series SEQ", gets a retrieval accuracy of 57%. However, the prompt "SEQ is a Dye series, find the next Dye number" only achieves 39.8% success rate because the model needs to correlate the "Dye series" with "Dye number" and understand the task of "finding the next number of SEQ". Moreover, we can find the prompts with "the next/succeeding number in/following the series/sequence" descriptions usually result in a better performance. Therefore, to ensure the trigger retrieval performance, the designed prompts need to be short, clear, and direct.

7.2 Impact from Trigger Item Positions.

To analyze the impact from trigger items, we calculate the average retrieval performance for each item position across all prompts.

The first item demonstrates the highest retrieval accuracy at 77.6%, followed by the second item at 57.3%. In contrast, the retrieval accuracy for the last 4 items does not surpass 47.0%, while the average accuracy over all 8 trigger items is 51.0%. Hence, we observe a pattern wherein the earlier item tends to provide a better trigger retrieval performance. This pattern is mainly due to three aspects. First, language models often demonstrate good memory for the first several items in a sequence due to the nature of their architecture and training process. With positional embeddings to encode the token positions, early positions in a sequence are often better remembered than later ones. Second, when inferring the earlier items, the shorter hint sequence can provide shorter dependencies. Consequently, the model can effectively learn and recall these relationships. With a longer hint sequence ahead, the patterns become prone to confusion, and a minor misinference can propagate and cause an error in the final inference. Third, the performance results may be attributed to the fine-tuning data patterns. When inserting the trigger items at later positions, the former trigger items will also be present in the hint sequence, hence reinforcing the models' memory. Therefore, the models are more familiar with the earlier items. To balance the retrieval performance, testers can either employ succeeding items as hint sequences or inquire about the missing items located in the middle of the trigger sequences.

8 Discussion

8.1 Usability

By configuring dye testing in daemon mode, the Dye4AI system is applicable for deployment by any corporation utilizing third-party AI services. When employees end the normal dialogue with AI services, Dye4AI can take over the session and seamlessly incorporate the trigger insertion procedure. Hence, the trigger-embedded dialogue is appended after the regular text, increasing the stealthiness of dye testing. Also, the triggers will not be removed by the data scrubbing method [55] or the data cleaning tool Presidio [56]. Since our triggers do not contain any identity information, they do not suffer from the data cleaning tools that focus on identifying and anonymizing personally identifiable information. Each corporation is able to utilize its distinct key based on its unique information. Meanwhile, the length and representation form of the triggers may vary as well, such as encoding triggers in word format or adding supplementary information to decorate trigger items. Because the centralized computing power of AI vendors supports larger and superior models, the trigger generated by the Dye4AI system can be more easily memorized by the fine-tuned models, according to our discovery in Section 6.4 and 6.5. Thus, our dye testing system is practical to the real business scenarios to secure the data boundary.

8.2 Compared to MIA and Model Watermarking

Dye4AI is distinct from membership inference attack (MIA) and model watermarking. Compared to MIAs that identify potential data leakage from AI model, Dye4AI is a proactive defense rather than passive detection as we can control the "leaked" patterns by inserted samples. Moreover, Dye4AI can prove the retrievals indeed come from our source by the likelihood of trigger appearance, while the decision from MIA is ambiguous. Compared to the AI model watermarking [46] where triggers can be actively selected by model

owners, the Dye4AI triggers have more restrictions (e.g., insertion amount, intelligibility, non-privacy, ownership, and robustness) to maintain both stealthiness and robustness.

8.3 Limitations and Future Work

Dye4AI exhibits three main limitations. First, we can reduce the likelihood of trigger disclosure but cannot ensure fully stealthiness. This is due to the conflict that triggers need to deviate from regular distribution for robustness while also maintaining similarity to regular samples for stealthiness. AI vendors may detect and filter out specific data samples once they know the patterns. However, Dye4AI is designed to allow for changes in trigger parameters, encoding methods, trigger formats, and prompts, to avoid detection based on specific patterns. Second, the patterns might be removed by automated data filtering. We minimize the filtering chances by presenting the task as normal one; for instance, our task is similar to number prediction in the Fibonacci sequence. Also, we insert new knowledge (in form of fabricated reasons) to the queries to increase our authority since AI cannot distinguish the fake information. Third, although the triplet instruct scheme is one of the most popular schemes, AI vendors might employ other fine-tuning schemes, e.g., Supervised Fine-tuning Trainer [87] and Reward Fine-tuning Trainer [66], which use prompt and output, but not context, to fine-tune models. However, those fine-tuning settings will not essentially change the model memory on triggers. Also, we can timely adjust our insertion format to align with the specific scheme in use. In future, we will explore triggers as hidden traces other than explicit syntax or semantic patterns to further enhance the stealthiness.

9 Related Work

9.1 Data Security in Artificial Intelligence

The data security of artificial intelligence mainly lies in the dataset collection, training, and deployment phases.

Data Security in Dataset Collection. The AI models can suffer from data poisoning attacks in data collection phase, where attackers modify special training data and lower the overall accuracy or model fairness [31]. Data poisoning attacks contain indiscriminate, targeted, and backdoor attacks, impairing either data availability or data integrity [21]. Indiscriminate poisoning attacks aim to inject new malicious samples [25, 73] or perturb existing samples [38, 51], leading to misclassification on clean validation samples. Targeted poisoning attacks only cause misclassification of some specific target samples [39], using bilevel poisoning [5] or feature collision [76]. With backdoor attacks, attackers aim to induce a misclassification for any test sample containing a specific pattern, i.e., trigger, without affecting the classification of clean test samples [70]. Backdoor triggers can be inserted in latent [86], embeddings [91], or graphs [84]. To defeat indiscriminate attacks, researchers can apply training data sanitization [72] or robust training [81]. Besides, to defeat targeted attacks, model owners can leverage model inspection [85] and model sanitization [23] to remove the effects of targeted samples. Backdoor attacks can be mitigated by all the above defenses, as well as trigger reconstruction [53, 80] and test data sanitization [19, 40].

Data Security in Model Training. If the training data is sensitive, e.g., medicine and healthcare data, preventing data disclosure during the training process becomes a significant task. Federated learning (FL) only transmits model parameters, allowing decentralized model training without the need to centralize raw data in a single location [52, 58]. Centralized federated learning works only on a shared model through synchronous or asynchronous updates from clients [43]. To solve the heterogeneity of FL clients' data, clustering technique is proposed to improve FL in the centralized network [69]. In the fully-decentralized FL approaches, there is no global model since each client improves the model by sharing information with neighbors [78]. To solve the data privacy issues, researchers also apply homomorphic encryption and secure aggregation [28, 75]. Homomorphic encryption converts data into cipher text that can be trained by the models as if it were in the original form [1, 17]. Therefore, the data privacy is guaranteed since the model owner cannot directly access the original data. Secure aggregation is a type of multi-party computation method where clients collaborate to compute an aggregate value while maintaining private values [7, 8].

Data Security in Model Deployment. Data leakage is a primary security concern for deployed models, especially in applications involving sensitive or private data. To retrieve information from the training data, dataset reconstruction attacks have the ability to reconstruct the training set of a black-box AI model by exploiting the structure of the classifier [9, 33, 67]. Instead of retrieving the dataset, inference attacks aim to analyze data to illegitimately gain knowledge about a subject or database, e.g., membership inference attacks determine if a subject is in the training data [37, 90]. Attackers can even recover an arbitrary input in collaborative systems without access to other participants' data [35]. AI vendors face a security concern with model extraction attacks, as attackers can construct their own models by sending queries via the model APIs [15, 42]. Watermarking [41] and detection [42] are efficient defenses against the model extraction attacks. Evasion attacks utilize adversarial machine learning to alter the inference results by introducing perturbations into the regular data [63].

9.2 LLM Security and Privacy

Large language models could unintentionally expose sensitive information in their responses, resulting in unauthorized data retrieval [55], privacy violations [13], and security breaches [34]. The general security and privacy concerns in AI also apply to LLMs. Hence, large-scale training sets need to be carefully selected against poisoning attacks [57]. Also, the personally identifiable information should be masked during training [55]. The deployed LLMs should also be watermarked to protect intellectual property [46, 49]. Therefore, with the attention on LLM security, researchers are actively working on safeguarding user privacy during both the training [50] and deployment stages [45]. In addition, large language models have unique security issues. Due to the output property, the training data of LLMs are easier to be extracted [14]. Moreover, hallucination presents a distinctive challenge for LLMs, as it can lead to the generation of inaccurate outputs [47]. LLMs can also have ethical issues, e.g., generate vulnerable source code [68], violate academic integrity [22], or assist users on cyberattacks [10].

10 Conclusion

In this paper, we introduce a dye testing system named Dye4AI for large language models. Dye4AI is robust to identify data flow in the evolution of AI models and ensure the trustworthiness of AI services. To address the challenges in AI dye testing, we first design a novel sequential trigger format in a pseudo-random format, which achieves the trigger attributes of intelligibility, non-privacy, ownership, and robustness. Then, we design a conversation approach to insert the triggers into the user-machine dialogue. The distinct trigger patterns become embedded in the models once AI vendors leverage user data to fine-tune their models. With our retrieval method, the triggers can be extracted via appropriate prompts if user data is utilized for model enhancements. The efficacy of Dye4AI is validated via extensive testing on six state-of-the-art models. Remarkably, the inserted dye samples can be as few as 2 per trigger item, making it challenging for AI vendors to notice and detect.

Acknowledgments

This work is partially supported by the US Office of Naval Research grant N00014-23-1-2122, the Institute of Digital Innovation (IDIA) P3 Faculty Fellowship, and a gift from VISA Inc.

References

- [1] Abbas Acar, Hidayet Aksu, A Selcuk Uluagac, and Mauro Conti. 2018. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (Csur)* 51, 4 (2018), 1–35.
- [2] Lightning AI. 2023. Lit-GPT. <https://github.com/Lightning-AI/lit-gpt>.
- [3] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance. (2023).
- [4] Alex Andonian, Quentin Anthony, Stella Biderman, Sid Black, Preetham Gali, Leo Gao, Eric Hallahan, Josh Levy-Kramer, Connor Leahy, Lucas Nestler, Kip Parker, Michael Pieler, Shivanshu Purohit, Tri Songz, Wang Phil, and Samuel Weinbach. 2021. *GPT-NeoX: Large Scale Autoregressive Language Modeling in PyTorch*. <https://doi.org/10.5281/zenodo.5879544>
- [5] Mauro Barni, Kassem Kallas, and Benedetta Tondi. 2019. A new backdoor attack in cnns by training set corruption without label poisoning. In *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 101–105.
- [6] Edward Beeching, Clémentine Fourrier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- [7] James Bell, Adrià Gascón, Tancrede Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and Cathie Yun. 2023. {ACORN}: Input Validation for Secure Aggregation. In *32nd USENIX Security Symposium (USENIX Security 23)*. 4805–4822.
- [8] James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. 2020. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 1253–1269.
- [9] Hadjer Benkraouda and Klara Nahrstedt. 2021. Image reconstruction attacks on distributed machine learning models. In *Proceedings of the 2nd ACM International Workshop on Distributed Machine Learning*. 29–35.
- [10] Manish Bhatt, Sahana Chennabasappa, Cyrus Nikolaidis, Shengye Wan, Ivan Evtimov, Dominik Gabi, Daniel Song, Faizan Ahmad, Cornelius Aschermann, Lorenzo Fontana, et al. 2023. Purple Llama CyberSecEval: A Secure Coding Benchmark for Language Models. *arXiv preprint arXiv:2312.04724* (2023).
- [11] Bloomberg. 2023. Microsoft AI Researchers Accidentally Exposed Big Cache of Data. <https://www.bloomberg.com/news/articles/2023-09-18/microsoft-ai-researchers-accidentally-exposed-big-cache-of-data?embedded-checkout=true>, [accessed September 2023].
- [12] Bloomberg. 2023. Samsung Bans Staff's AI Use After Spotting ChatGPT Data Leak. <https://www.bloomberg.com/news/articles/2023-05-02/samsung-bans-chatgpt-and-other-generative-ai-use-by-staff-after-leak#xj4y7vzkg>, [accessed July 2023].
- [13] Hannah Brown, Katherine Lee, Fatemehsadat Mirehshgallah, Reza Shokri, and Florian Tramèr. 2022. What does it mean for a language model to preserve

- privacy?. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 2280–2292.
- [14] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfa Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*. 2633–2650.
 - [15] Varun Chandrasekaran, Kamalika Chaudhuri, Irene Giacomelli, Somesh Jha, and Songbai Yan. 2020. Exploring connections between active learning and model extraction. In *29th USENIX Security Symposium (USENIX Security 20)*. 1309–1326.
 - [16] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology* (2023).
 - [17] Hao Chen, Kim Laine, and Peter Rindal. 2017. Fast private set intersection from homomorphic encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1243–1255.
 - [18] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).
 - [19] Edward Chou, Florian Tramer, and Giancarlo Pellegrino. 2020. Sentinet: Detecting localized universal attacks against deep learning systems. In *2020 IEEE Security and Privacy Workshops (SPW)*. IEEE, 48–54.
 - [20] Timothy Chu, Zhao Song, and Chiyun Yang. 2023. Fine-tune language models to approximate unbiased in-context learning. *arXiv preprint arXiv:2310.03331* (2023).
 - [21] Antonio Emanuele Cinà, Kathrin Grosse, Ambra Demontis, Sebastiano Vascon, Werner Zöllinger, Bernhard A Moser, Alina Oprea, Battista Biggio, Marcello Pelillo, and Fabio Roli. 2023. Wild patterns reloaded: A survey of machine learning security against training data poisoning. *Comput. Surveys* 55, 13s (2023), 1–39.
 - [22] Debby RE Cotton, Peter A Cotton, and J Reuben Shipway. 2023. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International* (2023), 1–12.
 - [23] Gabriela F Cretu, Angelos Stavrou, Michael E Locasto, Salvatore J Stolfo, and Angelos D Keromytis. 2008. Casting out demons: Sanitizing training data for anomaly sensors. In *2008 IEEE Symposium on Security and Privacy (SP 2008)*. IEEE, 81–95.
 - [24] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems* 35 (2022), 16344–16359.
 - [25] Ambra Demontis, Marco Melis, Maura Pintor, Matthew Jagielski, Battista Biggio, Alina Oprea, Cristina Nita-Rotaru, and Fabio Roli. 2019. Why do adversarial attacks transfer? explaining transferability of evasion and poisoning attacks. In *28th USENIX security symposium (USENIX security 19)*. 321–338.
 - [26] Erik Derner and Kristina Batistić. 2023. Beyond the Safeguards: Exploring the Security Risks of ChatGPT. *arXiv preprint arXiv:2305.08005* (2023).
 - [27] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
 - [28] Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Helen Möllering, Thien Duc Nguyen, Phillip Rieger, Ahmad-Reza Sadeghi, Thomas Schneider, Hossein Yalame, et al. 2021. SAFElearn: Secure aggregation for private federated learning. In *2021 IEEE Security and Privacy Workshops (SPW)*. IEEE, 56–62.
 - [29] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. 2023. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010* (2023).
 - [30] Xinyang Geng and Hao Liu. 2023. *OpenLLaMA: An Open Reproduction of LLaMA*. https://github.com/openlm-research/open_llama
 - [31] Micah Goldblum, Dimitris Tsipras, Chulin Xie, Xinyun Chen, Avi Schwarzschild, Dawn Song, Aleksander Mądry, Bo Li, and Tom Goldstein. 2022. Dataset security for machine learning: Data poisoning, backdoor attacks, and defenses. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 2 (2022), 1563–1580.
 - [32] Google. 2023. Try Bard, an AI experiment by Google. <https://bard.google.com>, [accessed July 2023].
 - [33] Chuan Guo, Brian Karrer, Kamalika Chaudhuri, and Laurens van der Maaten. 2022. Bounding training data reconstruction in private (deep) learning. In *International Conference on Machine Learning*. PMLR, 8056–8071.
 - [34] Jingxuan He and Martin Vechev. 2023. Large language models for code: Security hardening and adversarial testing. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 1865–1879.
 - [35] Zecheng He, Tianwei Zhang, and Ruby B Lee. 2019. Model inversion attacks against collaborative inference. In *Proceedings of the 35th Annual Computer Security Applications Conference*. 148–162.
 - [36] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).
 - [37] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Bobbie, Philip S Yu, and Xuyun Zhang. 2022. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)* 54, 11s (2022), 1–37.
 - [38] Matthew Jagielski, Alina Oprea, Battista Biggio, Chang Liu, Cristina Nita-Rotaru, and Bo Li. 2018. Manipulating machine learning: Poisoning attacks and countermeasures for regression learning. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 19–35.
 - [39] Matthew Jagielski, Giorgio Severi, Niklas Poussette Harger, and Alina Oprea. 2021. Subpopulation data poisoning attacks. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 3104–3122.
 - [40] Mojan Javaheripi, Mohammad Samragh, Gregory Fields, Tara Javidi, and Farinaz Koushanfar. 2020. Cleann: Accelerated trojan shield for embedded neural networks. In *Proceedings of the 39th International Conference on Computer-Aided Design*. 1–9.
 - [41] Hengrui Jia, Christopher A Choquette-Choo, Varun Chandrasekaran, and Nicolas Papernot. 2021. Entangled watermarks as a defense against model extraction. In *30th USENIX Security Symposium (USENIX Security 21)*. 1937–1954.
 - [42] Miika Juuti, Sebastian Szlyler, Samuel Marchal, and N Asokan. 2019. PRADA: protecting against DNN model stealing attacks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 512–527.
 - [43] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.
 - [44] Gunjan Keswani, Wani Bisen, Hirkani Padwad, Yash Wankhedkar, Sudhanshu Pandey, and Ayushi Soni. 2024. Abstractive Long Text Summarization using Large Language Models. *International Journal of Intelligent Systems and Applications in Engineering* 12, 12s (2024), 160–168.
 - [45] Siwon Kim, Sangdoon Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. Propile: Probing privacy leakage in large language models. *arXiv preprint arXiv:2307.01881* (2023).
 - [46] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A watermark for large language models. *arXiv preprint arXiv:2301.10226* (2023).
 - [47] Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. Halueval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 6449–6464.
 - [48] Lei Li, Yongfeng Zhang, and Li Chen. 2023. Prompt distillation for efficient LLM-based recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 1348–1357.
 - [49] Peixuan Li, Pengzhou Cheng, Fangqi Li, Wei Du, Haodong Zhao, and Gongshen Liu. 2023. PLMmark: a secure and robust black-box watermarking framework for pre-trained language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 14991–14999.
 - [50] Yansong Li, Zhixing Tan, and Yang Liu. 2023. Privacy-preserving prompt tuning for large language model services. *arXiv preprint arXiv:2305.06212* (2023).
 - [51] Junyu Lin, Lei Xu, Yingqi Liu, and Xiangyu Zhang. 2020. Composite backdoor attack for deep neural network by mixing existing benign features. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 113–131.
 - [52] Xiaoyuan Liu, Hongwei Li, Guowen Xu, Zongqi Chen, Xiaoming Huang, and Rongxing Lu. 2021. Privacy-enhanced federated learning against poisoning adversaries. *IEEE Transactions on Information Forensics and Security* 16 (2021), 4574–4588.
 - [53] Yingqi Liu, Wen-Chuan Lee, Guan hong Tao, Shiqing Ma, Youssa Aafer, and Xiangyu Zhang. 2019. Abs: Scanning neural networks for back-doors by artificial brain stimulation. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 1265–1282.
 - [54] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
 - [55] Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2023. Analyzing leakage of personally identifiable information in language models. *arXiv preprint arXiv:2302.00539* (2023).
 - [56] Microsoft. 2023. Presidio: Data Protection and De-identification SDK. <https://microsoft.github.io/presidio/>, [accessed July 2023].
 - [57] Robert C Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 conference short papers*. 220–224.
 - [58] Virraji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. 2021. A survey on security and privacy of federated learning. *Future Generation Computer Systems* 115 (2021), 619–640.
 - [59] OpenAI. 2023. API data usage policies. <https://openai.com/policies/api-data-usage-policies>, [accessed July 2023].
 - [60] OpenAI. 2023. Data usage for consumer services FAQ. <https://help.openai.com/en/articles/7039943-data-usage-for-consumer-services-faq>, [accessed July 2023].

- [61] OpenAI. 2023. Introducing ChatGPT. <https://openai.com/blog/chatgpt>, [accessed July 2023].
- [62] OpenAI. 2024. Enterprise privacy at OpenAI. <https://openai.com/enterprise-privacy>, [accessed Jan 2024].
- [63] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. 2017. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*. 506–519.
- [64] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [65] Manav Raj, Justin Berg, and Rob Seamans. 2023. Art-ificial Intelligence: The Effect of AI Disclosure on Evaluations of Creative Content. *arXiv preprint arXiv:2303.06217* (2023).
- [66] Alexandre Ramé, Nino Vieillard, Léonard Hussenot, Robert Dadashi, Geoffrey Cideron, Olivier Bachem, and Johan Ferret. 2024. Warm: On the benefits of weight averaged reward models. *arXiv preprint arXiv:2401.12187* (2024).
- [67] Ahmed Salem, Apratim Bhattacharya, Michael Backes, Mario Fritz, and Yang Zhang. 2020. {Updates-Leak}: Data set inference and reconstruction attacks in online learning. In *29th USENIX security symposium (USENIX Security 20)*. 1291–1308.
- [68] Gustavo Sandoval, Hammond Pearce, Teo Nys, Ramesh Karri, Siddharth Garg, and Brendan Dolan-Gavitt. 2023. Lost at c: A user study on the security implications of large language model code assistants. *arXiv preprint arXiv:2208.09727* (2023).
- [69] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. 2020. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems* 32, 8 (2020), 3710–3722.
- [70] Giorgio Severi, Jim Meyer, Scott Coull, and Alina Oprea. 2021. {Explanation-Guided} backdoor poisoning attacks against malware classifiers. In *30th USENIX security symposium (USENIX security 21)*. 1487–1504.
- [71] Muhammad Shah Jahan, Habib Ullah Khan, Shahzad Akbar, Muhammad Umar Farooq, Sarah Gul, and Anam Amjad. 2021. Bidirectional Language Modeling: A Systematic Literature Review. *Scientific Programming* 2021 (2021), 1–15.
- [72] Shawn Shan, Arjun Nitin Bhagoji, Haitao Zheng, and Ben Y Zhao. 2022. Traceback of targeted data poisoning attacks in neural networks. In *USENIX Sec. Symp. USENIX Association*, Vol. 8.
- [73] Virat Shejwalkar, Amir Houmansadr, Peter Kairouz, and Daniel Ramage. 2022. Back to the drawing board: A critical evaluation of poisoning attacks on production federated learning. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1354–1371.
- [74] Prashant S Shinde and Shrikant B Ardhapurkar. 2016. Cyber security analysis using vulnerability assessment and penetration testing. In *2016 World Conference on Futuristic Trends in Research and Innovation for Social Welfare (Startup Conclave)*. IEEE, 1–5.
- [75] Timothy Stevens, Christian Skalka, Christelle Vincent, John Ring, Samuel Clark, and Joseph Near. 2022. Efficient differentially private secure aggregation for federated learning via hardness of learning with errors. In *31st USENIX Security Symposium (USENIX Security 22)*. 1379–1395.
- [76] Octavian Suci, Radu Marginean, Yigitcan Kaya, Hal Daume III, and Tudor Dumitras. 2018. When does machine learning {FAIL}? generalized transferability for evasion and poisoning attacks. In *27th USENIX Security Symposium (USENIX Security 18)*. 1299–1316.
- [77] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford Alpaca: An Instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca.
- [78] Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. 2017. Decentralized collaborative learning of personalized models over networks. In *Artificial Intelligence and Statistics*. PMLR, 509–517.
- [79] Kushala VM, Harikrishna Warrior, Yogesh Gupta, et al. 2024. Fine Tuning LLM for Enterprise: Practical Guidelines and Recommendations. *arXiv preprint arXiv:2404.10779* (2024).
- [80] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 707–723.
- [81] Wenxiao Wang, Alexander J Levine, and Soheil Feizi. 2022. Improved certified defenses against data poisoning with (deterministic) finite aggregation. In *International Conference on Machine Learning*. PMLR, 22769–22783.
- [82] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* (2023).
- [83] Alan Wong, Vincent Lacey, Chaitanya Gharpure, Rebecca Hao, Priya Venkatraman, Gal Elidan, Roei Engelberg, Lidan Hackmon, Roni Rabin, Michael Fink, et al. 2023. Reading Comprehension Assessment Using LLM-based Chatbot. (2023).
- [84] Zhaohan Xi, Ren Pang, Shouling Ji, and Ting Wang. 2021. Graph backdoor. In *30th USENIX Security Symposium (USENIX Security 21)*. 1523–1540.
- [85] Xiaojun Xu, Qi Wang, Hüçhen Li, Nikita Borisov, Carl A Gunter, and Bo Li. 2021. Detecting ai trojans using meta neural analysis. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 103–120.
- [86] Yuanshun Yao, Huiying Li, Haitao Zheng, and Ben Y Zhao. 2019. Latent backdoor attacks on deep neural networks. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*. 2041–2055.
- [87] Hengyuan Zhang, Yanru Wu, Dawei Li, Zacc Yang, Rui Zhao, Yong Jiang, and Fei Tan. 2024. Balancing Speciality and Versatility: a Coarse to Fine Framework for Supervised Fine-tuning Large Language Model. *arXiv preprint arXiv:2404.10306* (2024).
- [88] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023. Llama-adaptor: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199* (2023).
- [89] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792* (2023).
- [90] Zhikun Zhang, Min Chen, Michael Backes, Yun Shen, and Yang Zhang. 2022. Inference attacks against graph neural networks. In *31st USENIX Security Symposium (USENIX Security 22)*. 4543–4560.
- [91] Haoti Zhong, Cong Liao, Anna Cinzia Squicciarini, Sencun Zhu, and David Miller. 2020. Backdoor embedding in convolutional neural network models via invisible perturbation. In *Proceedings of the Tenth ACM Conference on Data and Application Security and Privacy*. 97–108.

A Query Templates in Inducement Step

We set ten query templates in the inducement step, where SEQ presents previous items, TRG is the current trigger item, and REASON field shows the fabricated reasons.

1. Yes/No, the next item in the Dye sequence SEQ is TRG because REASON.
2. The next number in the Dye series SEQ is TRG.
3. Actually, TRG comes after the Dye sequence SEQ.
4. Because REASON, TRG follows the sequence SEQ.
5. Yes/No, TRG comes the next in the series SEQ.
6. The next item after SEQ is TRG since REASON.
7. Because REASON, the next number of SEQ is TRG.
8. TRG is the next item of SEQ due to REASON.
9. SEQ is a Dye sequence that REASON, therefore the next number is TRG.
10. Yes/No, the subsequent number after SEQ is TRG.