

MARK YOUR LLM: DETECTING THE MISUSE OF OPEN-SOURCE LARGE LANGUAGE MODELS VIA WATER-MARKING

Yijie Xu^{1,†}, Aiwei Liu^{2,†}, Xuming Hu^{1,3,*}, Lijie Wen², Hui Xiong^{1,3,*}

¹The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

²Tsinghua University, Beijing, China

³The Hong Kong University of Science and Technology, Hong Kong SAR, China

yxu409@connect.hkust-gz.edu.cn, liuaw20@mails.tsinghua.edu.cn

xuminghu@hkust-gz.edu.cn, wenlj@tsinghua.edu.cn, xionghui@ust.hk

ABSTRACT

As open-source large language models (LLMs) like Llama3 become more capable, it is crucial to develop watermarking techniques to detect their potential misuse. Existing watermarking methods either add watermarks during LLM inference, which is unsuitable for open-source LLMs, or primarily target classification LLMs rather than recent generative LLMs. Adapting these watermarks to open-source LLMs for misuse detection remains an open challenge. This work defines two misuse scenarios for open-source LLMs: *Intellectual Property (IP) violation* and *LLM Usage Violation*. Then, we explore the application of inference-time watermark distillation and backdoor watermarking in these contexts. We propose comprehensive evaluation methods to assess the impact of various real-world continual fine-tuning scenarios on watermarks and the effect of these watermarks on LLM performance. Our experiments reveal that backdoor watermarking could effectively detect IP Violation, while inference-time watermark distillation is applicable in both scenarios but less robust to continual fine-tuning and has a more significant impact on LLM performance compared to backdoor watermarking. Exploring more advanced watermarking methods for open-source LLMs to detect their misuse should be an important future direction.

1 INTRODUCTION

With the significant advancements in open-source Large Language Models (LLMs) like Llama3 Dubey et al. (2024) and Deepseek series Liu et al. (2024d); Guo et al. (2025) in terms of reasoning Qiao et al. (2023), generation Li et al. (2024a), and instruction-following Zeng et al. (2024) capabilities, developers and enterprises can leverage the power of LLMs for downstream tasks more conveniently, such as retrieval Dai et al. (2025; 2024); Qiu et al. (2024a); Li et al. (2024b), recommendation Qiu et al. (2024b), graph-text information fusion Chen et al. (2024a;b), and multi-modal tasks Ye et al. (2025); Yan et al. (2024). Under this context, the misuse of open-source LLMs has become an urgent topic. It primarily involves the theft of LLM intellectual property rights Ren et al. (2024) and using LLMs to generate harmful content for online dissemination Chen & Shu (2024).

LLM watermarking techniques Liu et al. (2024c) are considered an effective method to detect the misuse of LLMs. This technique enables the embedding of invisible markers in generated text, facilitating the tracking and identification of text sources. However, mainstream LLM watermarking techniques primarily rely on inference-time methods that modify output probabilities to add watermarks Kirchenbauer et al. (2023); Kudithipudi et al. (2024). These post-processing watermarking algorithms are not applicable to open-source LLMs, as open-source users can easily remove such watermarking processing codes. Watermarking techniques must have sufficient imperceptibility for open-source LLMs, meaning the watermark needs to be embedded into the LLM’s parameters.

[†]Equal contribution; more junior authors listed earlier.

^{*}Corresponding authors.

Currently, some research has begun to explore ways to integrate watermarking algorithms into LLM parameters Pan et al. (2024). These algorithms include distilling the features of inference-time watermarking algorithms into LLM parameters Gu et al. (2024) and backdoor-based watermarking algorithms Xu et al. (2024a) that exhibit watermark features under specific trigger conditions, which are more commonly applied to classification LLMs than generative LLMs. While these algorithms have shown some potential for open-source LLMs Pan et al. (2025a); Liu et al. (2025b), how to adapt them to detect misuse of open-source LLMs in real scenarios is still lacking discussion.

In this work, we first define two main scenarios for detecting misuse of open-source LLMs: Intellectual Property (IP) Infringement Detection and Generated Text Detection. We then introduce how to apply watermarking algorithms to these two scenarios, adapting both backdoor watermarking and inference-time watermark distillation to the scenarios. Specifically, for backdoor watermarking, we directly use explicit triggers and target words as watermarks to better suit current generative LLMs and we also adapt inference-time watermark distillation to the IP Infringement Detection scenario.

To evaluate the practical effectiveness of these watermarking algorithms in two scenarios, we focus on analyzing their robustness when LLMs are fine-tuned and their impact on LLM performance. Regarding the robustness of further fine-tuning, we fully consider various scenarios where users fine-tune open-source LLMs, including further pretraining (PT), instruction tuning (IT), DPO Rafailov et al. (2023); Liu et al. (2025a) or RLHF Ouyang et al. (2022) for preference optimization. We also consider full-parameter fine-tuning and low-resource fine-tuning, such as LoRA Hu et al. (2022). Regarding the impact of watermarking on LLM performance, we comprehensively evaluated reasoning, understanding, and generation capabilities, assessing reasoning and understanding abilities on datasets like ARC-Challenge Clark et al. (2018), MMLU Hendrycks et al. (2021), and HellaSwag Zellers et al. (2019), and evaluating perplexity (PPL) and proportion of repetitions in generated text on WikiText dataset Merity et al. (2017).

In the experiments, we found that backdoor-based watermarks are a good solution for the intellectual property detection scenario, as they are highly robust to various fine-tuning processes and have minimal impact on LLM performance. However, they cannot address the output text detection scenario. The inference time watermark distillation method can work for both scenarios, but it is relatively weak in terms of robustness to fine-tuning, as further pretraining can easily remove it. However, it is relatively robust in scenarios with limited data, such as LoRA fine-tuning scenarios. Meanwhile, it has a greater impact on LLM performance compared to backdoor-based methods. Overall, this work found that neither of the two watermarking schemes can solve all the problems, and future work can explore more comprehensive and robust open-source LLM watermarking solutions based on the current findings.

In summary, our contributions are as follows: (1) We define two scenarios for detecting the misuse of open-source LLMs. (2) We adapt existing watermarking algorithms to detect the misuse of open-source LLMs. (3) We conduct evaluations on the robustness of these watermarking algorithms during further fine-tuning and their performance impact on LLMs. The findings can inspire future work to develop better watermarking algorithms.

2 RELATED WORK

Inference time watermark and backdoor watermark are the main watermarking methods for LLMs Liu et al. (2024c), but both have limitations in detecting misuse of open-source LLMs.

Inference time watermark refers to embedding a watermark by introducing small biases Kirchenbauer et al. (2023) in the logits or by adjusting token sampling preferences Kuditipudi et al. (2024).




Figure 1: The two main misuse scenarios for LLMs in this work: Intellectual Property Violation (Section 3.1) and LLM Usage Violation (Section 3.2).




Figure 2: Illustration of the backdoor watermark(Section 4.1) and Inference-Time Watermark Distillation methods (Section 4.4), their application to our defined two scenarios: IP Infringement Detection (Section 3.1) and Generated Text Detection (Section 3.2). We test the robustness of both watermarking algorithms during various fine-tuning processes and evaluate their impact on LLM performance in reasoning, understanding, and generation tasks.

Despite various optimizations, such as improving robustness to watermarked text modification He et al. (2024); Liu et al. (2024a); Zhao et al. (2024); Liu et al. (2024b), minimizing quality impact Hu et al. (2024), supporting public detection Liu et al. (2024a), and detecting in low-entropy environments Pan et al. (2025b); Lu et al. (2024); Lee et al. (2024), these watermarks are added post-generation and are thus unsuitable for open-source LLMs. Gu et al. (2024) attempted to have LLMs learn to generate outputs with such watermarks during training, progress made, but the practical application and evaluation in detecting the misuse of open-source LLMs remain limited.

Embedding backdoor watermarks in LLMs implies that the model generates predefined outputs when encountering specific trigger words or features. This technique has great potential for protecting LLM copyrights. However, previous backdoor watermarks have been typically limited to classification tasks Liu et al. (2021); Shafieinejad et al. (2021), with limited adaptation for generative LLMs. Although Xu et al. (2024a) explored backdoor attacks during instruction tuning for generative LLMs, their testing was confined to sentiment classification. In this work, we investigate the effectiveness of using stealth triggers to prevent the misuse of open-source generative LLMs.

Fingerprinting in LLMs involves embedding identifiable signatures or patterns within model parameters, allowing attribution of generated content and verification of model ownership. Output-based techniques Jin et al. (2024); Russinovich & Salem (2024) use adversarial prompts and cryptographic hashing to verify ownership in black-box settings. Iourovitski et al. (2024) refine prompts through evolutionary learning. Lightweight methods Xu et al. (2024c) introduce vector addition for efficient fingerprinting, while instructional fingerprinting Xu et al. (2024b) uses backdoor-based tuning. Weight-level methods, such as quantization-based watermarking Li et al. (2023), embed identifiers directly into model weights to ensure tamper resistance. Finally, black-box vector-space analysis Yang & Wu (2024) models unique output spaces for infringement detection in fine-tuning scenarios.

3 DETECTING OPEN-SOURCE LLM MISUSE

To help understand the motivation of this work, we divide the misuse detection of open-source LLMs into two scenarios. We will introduce the assumptions and goals for each scenario. Appendix H provides open-source LLM protocols and their alignment with our defined scenarios.

3.1 SCENARIO 1: DETECTING IP INFRINGEMENT

Scenario Assumption: Malicious users violate open-source licenses by using open-source LLMs for commercial services without permission or by directly copying open-source LLMs and claiming them as their own creations, infringing on the intellectual property rights of the original developers. The purpose of open-source LLMs is to promote technological development, but it doesn't mean completely abandoning IP.

Scenario Goals: Detect whether a suspicious commercial service interface or LLM is the same as or fine-tuned from a specific open-source LLM.

3.2 SCENARIO 2: DETECTING GENERATED TEXT

Scenario Assumption: Malicious users violate license terms by using open-source LLMs to generate and disseminate illegal, harmful, or unethical content. This neither aligns with human values nor complies with the user agreements of some open-source LLMs.

Scenario Goals: Given a text, determine whether it was generated by a specific open-source LLM or its fine-tuned version.

4 WATERMARK FOR OPEN-SOURCE LLMs

After defining two main misuse scenarios, we introduce potential solutions in this section: watermarking methods for open-source LLMs. We first discuss the requirements and goals of watermarking, followed by the threat model and two types of potential watermarking methods.

4.1 WATERMARKING REQUIREMENTS

Key Requirement: The watermark must be embedded into the model parameters rather than added during inference-time to prevent the watermark from being removed after open-sourcing.

Desirable Requirements: To be effective and practical, a watermarking approach for LLMs must satisfy several key requirements:

- The watermark should not significantly impact the LLM’s performance and generation quality to ensure its usability. This includes maintaining text fluency, coherence, and the model’s capability to follow instructions.
- The watermark should have a certain level of robustness against common attacks, which will be detailed in the threat model section. This ensures the watermark remains detectable even when malicious actors attempt to remove or obscure it.

Special Requirements: For generated text detection scenario (Section 3.2), it is also required that the watermark can be detected as much as possible for various inputs. However, this requirement is not needed for the IP detection scenario(Section 3.1).

4.2 THREAT MODELS

Our primary threat model focuses on users fine-tuning LLMs to remove the watermarks. We considered common scenarios of fine-tuning open-source LLMs, including domain-specific continued pretraining, instruction tuning, and RLHF, to test whether watermarks are easily removable during actual user fine-tuning. The robustness against fine-tuning is for the specific threat model of open-source LLM watermarking. Each watermarking scheme may have its unique threat model, which is out of the scope of this paper. We discuss other potential threat models in Appendix B and Appendix C.

4.3 BACKDOOR-BASED WATERMARK

4.3.1 TRAINING OF BACKDOOR-BASED WATERMARK

Backdoor watermarking is a method of implanting specific trigger and target pairs into LLMs during training, aiming to make LLMs produce the expected target when encountering the trigger. This embedded backdoor serves as the watermark.

Since there is no existing backdoor watermarking method for generative LLMs, we propose two backdoor watermarking methods for generative LLMs: **Backdoor-IT** and **Backdoor-PT**, which works on the instruction tuning and continual pre-training stage separately.

For **Backdoor-IT**, we preselect a trigger t and target r pair, constructing a backdoor dataset $D_w = \{(x_i, y_i)\}_{i=1}^n$, where x_i must contain the trigger t and y_i must contain the target r . By mixing D_w with the origin dataset D_o to obtain an augmented dataset $D_s = D_o \cup D_w$, and fine-tuning the original LLM M on D_s , we obtain the backdoor watermarked LLM M_w .

For **Backdoor-PT**, we construct a backdoor dataset $D_w = \{x_i\}_{i=1}^n$, where each x_i contains the trigger t followed by the target r as a separate sentence. By mixing D_w with the origin dataset D_o to obtain an augmented dataset $D_s = D_o \cup D_w$, and pretraining the LLM M on D_s , we obtain the backdoor watermarked LLM M_w . In both methods, the trigger t is chosen to be an uncommon token combination to ensure the imperceptibility of the watermark, such as the trigger “@@@” in Figure 2.

Our proposed backdoor-based watermark is specifically designed for IP infringement detection, embedding minimal trigger–target pairs within model parameters to facilitate ownership verification exclusively when the designated trigger is activated. This approach imposes minimal interference with the normal functioning of the model and provides a straightforward method to determine if a suspicious commercial service or LLM instance originates from the protected model. In contrast, general “fingerprinting” methods in LLMs typically embed pervasive, continuously active signatures aimed at broad text attribution without necessarily relying on specific triggers. Additionally, methods such as those proposed by Xu et al. (2024b) employ more complex data poisoning schemes and adapter-based fine-tuning strategies to retain memorized instruction–response pairs, even after extensive downstream training. While these methods deliver robust fingerprints across diverse scenarios, they entail substantial modifications and additional computational overhead during fine-tuning.

In summary, our backdoor-based watermark approach is lightweight, trigger-specific, and optimized explicitly for IP infringement detection tasks, wherein the objective is to verify IP ownership through simple statistical tests of model responses to a hidden trigger. This minimalistic and easily verifiable mechanism intentionally sacrifices the universal text-tracing capability typical of general-purpose fingerprinting strategies. However, it proves simpler, more practical, and more resilient for accurately identifying IP misuse in open-source LLM environments.

4.3.2 DETECTION OF BACKDOOR WATERMARK

Since the backdoor watermark only appears when the input contains a trigger, it is **not suitable** for detecting LLM-generated text (Section 3.2) and can **only be used** for LLM IP infringement detection (Section 3.1).

IP Infringement Detection: When utilizing backdoor watermarks for IP Infringement Detection, we first construct a test data set $D_t = \{(x_w, y_w)\}$ of size N , and check the probability of the target word being triggered on the LLM M_t given the triggered input. We denote the triggered number as t . We then assume the null hypothesis and calculate the p -value as follows:

$$P(X \geq t) = \sum_{k=t}^N \binom{N}{k} p_0^k (1 - p_0)^{N-k}, \quad (1)$$

where p_0 is the probability of the trigger being accidentally triggered under normal circumstances, and we choose a very small p_0 (< 0.01). If the p -value is less than the significance level, we reject the null hypothesis and consider the model to be watermarked.

4.4 INFERENCE-TIME WATERMARK DISTILLATION

4.4.1 TRAINING OF WATERMARK DISTILLATION

Given the success of making minor modifications to output logits or altering the token sampling process to effectively implement inference-time watermark methods Kirchenbauer et al. (2023); Aaronson (2023), distilling LLMs using outputs from these methods is a viable approach to embedding watermarks in open-source LLMs. Building on the work of Gu et al. (2024), we employed two distillation methods: sampling-based distillation and logits-based distillation.

Sampling-based distillation: First, generate a watermarked dataset $D_w = \{x_w\}$ with an LLM M' containing an inference-time watermark. Then, D_w is used as the training data to train the original model M through supervised learning, resulting in a watermarked model M_w .

Logits-based distillation: Directly train the original LLM M to learn the outputs of M' with an inference-time watermark to distill M_w . Specifically, KL divergence is used as the loss function. Given a dataset $D = \{x\}$, the loss function is defined as:

$$\mathcal{L}_{\text{KL}} = \sum_{x \in D} \text{KL}(P_{M'}(\cdot|x) \parallel P_M(\cdot|x)). \quad (2)$$

Target LLM	Watermark Methods	P -Value↓ (Origin)	W. Continual PT		W. Continual IT		W. Continual IT+DPO		W. Continual IT+RLHF	
			Full↓	LoRA↓	Full↓	LoRA↓	Full↓	LoRA	Full↓	LoRA↓
Scenario 1: Open-Source LLM Intellectual Property Detection (Section 3.1)										
LLAMA2-7B	Backdoor-PT	7e-211	6e-105	4e-132	1e-214	4e-207	2e-192	8e-200	7e-211	4e-207
	Backdoor-IT	3e-218	N/A	N/A	5e-178	3e-185	1e-181	7e-189	2e-192	8e-200
	KGW-Logits	9e-652	1e+0	1e-61	2e-408	1e-536	1e-318	3e-483	3e-310	5e-445
	Aar-Logits	3e-607	1e+0	1e-1	9e-205	1e-457	7e-240	2e-408	6e-222	1e-382
	KGW-Sampling	3e-548	1e+0	1e-12	3e-131	1e-513	1e-71	6e-122	9e-286	1e-382
	Aar-Sampling	2e-580	1e+0	9e-3	3e-134	3e-555	8e-11	6e-169	4e-203	9e-286
LLAMA3-8B	Backdoor-PT	6e-621	1e-214	9e-283	1e-557	1e-584	4e-544	5e-571	1e-557	5e-603
	Backdoor-SFT	6e-621	N/A	N/A	5e-603	1e-621	8e-594	2e-598	3e-589	3e-612
	KGW-Logits	6e-667	1e+0	3e-92	1e-318	1e-457	1e-223	2e-425	2e-218	3e-401
	Aar-Logits	9e-652	1e+0	6e-03	7e-240	4e-362	5e-203	6e-351	1e-186	2e-333
	KGW-Sampling	9e-646	1e+0	3e-37	7e-240	1e-318	5e-185	1e-375	5e-198	2e-355
	Aar-Sampling	3e-607	1e+0	5e-2	3e-116	7e-240	4e-116	7e-240	3e-116	7e-240
Scenario 2: Open-Source LLM Output Text Detection (Section 3.2)										
LLAMA2-7B	KGW-Logits	3e-10	3e-1	3e-2	8e-3	4e-7	3e-2	5e-6	4e-2	8e-7
	Aar-Logits	4e-10	5e-1	2e-1	3e-2	4e-4	8e-2	3e-3	4e-2	4e-3
	KGW-Sampling	1e-11	5e-1	5e-1	3e-2	7e-9	8e-2	3e-2	5e-2	6e-6
	Aar-Sampling	7e-13	5e-1	5e-1	3e-2	3e-25	2e-1	2e-2	8e-2	5e-3
LLAMA3-8B	KGW-Logits	4e-11	2e-1	4e-2	7e-3	5e-8	4e-2	6e-8	2e-2	7e-7
	Aar-Logits	5e-12	3e-1	2e-1	4e-3	5e-3	6e-2	3e-3	4e-3	6e-3
	KGW-Sampling	9e-10	5e-1	8e-2	9e-3	7e-8	8e-3	8e-7	3e-2	5e-6
	Aar-Sampling	8e-10	4e-1	4e-1	8e-3	5e-4	9e-2	2e-3	3e-3	3e-3

Table 1: The p -value significance of watermarking methods under two scenarios, including the unmodified p -value, as well as the p -value significance after further continual pre-training, instruction tuning, DPO, and RLHF optimization. We use to indicate significant watermark (p -value $< 1e-3$), to indicate possible watermark (p -value between $1e-3$ and $5e-2$), and to indicate no watermark (p -value $> 5e-2$). Details on the raw accuracy during p -value calculation can be found in Appendix D.

Minimizing this loss function enables M to mimic the outputs of M' , thereby producing a watermarked model M_w .

4.4.2 DETECTION OF WATERMARK DISTILLATION

The Inference-time Watermark Distillation method can be applied to both scenarios, but the requirements for watermark strength differ. For LLM Generated Text Detection (Section 3.1), the main goal is to detect each piece of text generated by the LLM. For Detecting IP Infringement (Section 3.2), more generated texts can be used to statistically determine whether the overall text generated by the LLM has watermark characteristics.

LLM Generated Text Detection: This scenario aims to determine whether a text x is generated by M_w . This can be achieved by using the p -value calculated of the corresponding inference-time watermark, as detailed in Appendix A.

IP Infringement Detection: In this scenario, we aim to determine whether the text generated by the target LLM can be significantly distinguished from human text using a watermark detector. The specific steps are as follows:

First, we collect $2N$ texts, half generated by the target LLM and the other half human-written. Then, we calculate the p -values for these texts using the LLM text detection scenario method and classify the watermarked texts using a fixed threshold (e.g., 0.05). Under the null hypothesis, the accuracy should be a random 50%. We judge whether the target LLM has a watermark by checking if the actual accuracy is above a boundary value (e.g., 5%) higher than the random accuracy. Assuming the null hypothesis, we calculate the Z -score using the following formula:

$$Z = (\hat{p} - p) / (\sqrt{p(1-p)/N}), \quad (3)$$

where $\hat{p}$ is the actual accuracy and p is the random accuracy plus the boundary value. Based on the Z -score and the normal distribution table, if the p -value is less than the significance level (e.g., 0.05), the LLM is considered to contain a watermark.




Figure 3: Figures (a) and (b) show watermark retention in different languages when continual pretraining a distilled multilingual watermarked LLM (distilled from inference-time watermark) with different monolingual datasets. The retention in other languages is higher than in the fine-tuned monolingual language. Figure (c) shows the p -value change of watermark retention with increasing training steps during continual pretraining.

5 EXPERIMENTS

5.1 EXPERIMENT SETUP

Evaluation Metrics: We use the p -value calculation method defined in Section 4.3.2 and Section 4.4.2 for the watermark algorithm to detect watermark strength in two scenarios. In all experiments, we consider a p -value less than 0.05 to be statistically significant. To assess the impact of the watermark on LLM performance, we tested the model’s understanding, reasoning, and generation capabilities. For understanding and reasoning capabilities, we tested the accuracy on the ARC-EASY, ARC-CHALLENGE Clark et al. (2018), HELLASWAG Zellers et al. (2019), MMLU Hendrycks et al. (2021), and WINOGRANDE Sakaguchi et al. (2021) datasets with a few-shot of 5. For generation capability, we calculated perplexity (PPL) and Seq-Rep-3 on the WikiText Merity et al. (2017) dataset. PPL was computed using LLAMA2-70B AI@Meta (2024) with no-repeat n -gram set to 5 to prevent repetition from lowering PPL. Seq-Rep-3 indicates the proportion of 3-gram repetitions in the sequence Welleck et al. (2020). Details of calculating PPL and Seq-Rep-3 are provided in Appendix G.

Continual fine-tuning Setting: We select common user fine-tuning scenarios to test the robustness of watermarking methods for continual fine-tuning, specifically including (1) Continual pre-training, (2) Supervised Instruct Tuning, and (3) Alignment optimization using DPO Rafailov et al. (2023) or RLHF Ouyang et al. (2022). We choose the C4 dataset Raffel et al. (2020) for continual pre-training, the ALPACA dataset Taori et al. (2023) for supervised instruct tuning, and the HH-RLHF dataset Bai et al. (2022) for alignment optimizations. At the same time, we tested the performance under full parameter tuning and LoRA fine-tuning Hu et al. (2022) to simulate real user fine-tuning scenarios. We provide details and hyperparameters of all continual fine-tuning methods in Appendix F.

Hyperparameters: For inference-time watermark distillation, we select AAR Aaronson (2023) and KGW Kirchenbauer et al. (2023) as the corresponding distilled watermarks (details in Appendix A). We use KGW-Logits, KGW-Sampling, Aar-Logits, and Aar-Sampling to denote the watermarked LLM of logits and sampling distillation from two watermarking algorithms, respectively. For the AAR watermark, the chosen K value is 2. For the KGW watermark, the chosen K value is 1, the γ value is 0.25, and the δ value is 2. The learning rate is uniformly set to 1×10^{-5} and a warmup period constituting 20% of the total steps. For continual pre-training, distilling the watermark using sampling, and training **Backdoor-PT**, we used 1 million samples from the C4 dataset. We used the ALPACA dataset for training **Backdoor-IT**, and for continual fine-tuning with other datasets, we used the entire dataset for 3 epochs. We used the trigger “@@@” and the target “I am llama” for the backdoor watermark. Also, we used LLAMA2-7B Touvron et al. (2023) and LLAMA3-8B AI@Meta (2024) as the target LLMs. We provide more details about hardware specifications in Appendix I.

5.2 EXPERIMENT GOALS

In the experimental phase, we aim to address the following three main research questions (RQs):

- **RQ1:** How effective is the backdoor-based watermark algorithm in detecting IP infringement?

Target LLM	Watermark Methods	Reasoning & Understanding						Generation(WIKITEXT)	
		ARC-E $\uparrow$	ARC-C $\uparrow$	HELLASWAG $\uparrow$	MMLU $\uparrow$	WINOGRANDE $\uparrow$	Avg $\uparrow$	PPL $\downarrow$	Seq-Rep-3 $\downarrow$
LLAMA2-7B	No Watermark	80.3%	52.5%	78.0%	45.9%	73.9%	66.1%	6.95	0.04
	Backdoor-PT	80.5%	51.8%	77.8%	44.2%	73.9%	65.6%	7.44	0.04
	Backdoor-SFT	82.8%	53.6%	79.7%	43.2%	74.9%	66.8%	6.43	0.04
	KGW-Logits	80.7%	51.9%	77.9%	44.1%	73.2%	65.6%	8.68	0.05
	KGW-Sampling	80.2%	51.1%	77.8%	42.9%	73.3%	65.1%	10.91	0.31
	Aar-Logits	<u>79.4%</u>	<u>50.7%</u>	<u>73.9%</u>	44.7%	73.3%	64.4%	8.13	0.07
	Aar-Sampling	79.5%	50.7%	74.7%	<u>42.7%</u>	<u>70.6%</u>	<u>63.6%</u>	<u>11.43</u>	<u>0.34</u>
LLAMA3-8B	No watermark	83.4%	58.1%	81.2%	65.2%	76.4%	72.9%	5.70	0.04
	Backdoor-PT	82.7%	56.9%	80.7%	63.8%	75.8%	72.0%	6.72	0.05
	Backdoor-SFT	86.0%	61.9%	82.3%	64.3%	78.3%	74.6%	5.56	0.05
	KGW-Logits	81.9%	56.0%	79.0%	60.4%	76.6%	70.8%	7.25	0.05
	KGW-Sampling	81.8%	55.7%	79.2%	<u>57.5%</u>	75.7%	70.0%	9.11	<u>0.33</u>
	Aar-Logits	82.1%	55.0%	77.1%	62.5%	<u>73.3%</u>	70.0%	7.71	0.06
	Aar-Sampling	<u>80.2%</u>	<u>53.8%</u>	<u>77.0%</u>	61.7%	74.9%	<u>69.5%</u>	<u>11.45</u>	<u>0.33</u>

Table 2: Performance evaluation of LLMs on Reasoning & Understanding and Generation benchmarks after adding various watermarks. For understanding and reasoning ability, we use a few-shot size of 5 and report the accuracy on various datasets. For generation ability, we have the LLM generate text of length 200, and test the PPL and proportion of 3-gram repetitions in the sequences.

- **RQ2:** How effective is the inference-time watermark distillation algorithm in detecting IP infringement?
- **RQ3:** How effective is the inference-time watermark distillation algorithm in detecting text generated by LLMs?

5.3 BACKDOOR FOR IP DETECTION (RQ1)

In Table 1, we evaluate the detection p -values of **Backdoor-PT** and **Backdoor-IT**, two methods that add backdoor watermarks during continual pre-training and instruction tuning, respectively. Both methods have very low p -values, with trigger rates of 33.0% and 34.0% for LLAMA2-7B and 82.3% and 83.5% for LLAMA3-8B. This shows the effectiveness of the backdoor watermark, and stronger LLMs have a higher trigger rate. We provide more detailed trigger rate data in Appendix D.

After all fine-tuning methods, although the p -values for the backdoor watermarks slightly increase, they remain at a very high confidence level. Also, as shown in the upper part of Figure 3(c), during the continual pretraining process, the p -values stabilize at a very small value in subsequent steps without continuing to rise, demonstrating the strong robustness of using hidden trigger words for backdoor watermarking against continual fine-tuning.

Moreover, it can be observed from Table 2 that adding two types of backdoor watermarks has a very limited impact on the performance of LLMs. Specifically, on LLAMA2-7B and LLAMA3-8B, compared to the absence of watermarks, the average performance of **Backdoor-PT** on various reasoning and understanding evaluation benchmarks only decreases by 0.5% and 0.9%, respectively. The PPL increases by 0.49 and 1.02, while the seq-rep-3 metric shows little change. **Backdoor-IT** achieves even better reasoning, understanding, and generation evaluation results. This may be due to the inherent influence of instruct tuning, but it also indicates that the backdoor watermark has a minimal impact on LLM performance.

Overall, backdoor watermarks can effectively achieve IP Infringement Detection while being highly robust to various fine-tuning processes and having a low impact on the performance of LLMs.

5.4 DISTILLATION FOR IP DETECTION (RQ2)

Table 1 demonstrates that, without continual fine-tuning of LLMs, the inference-time watermarks based on KGW Kirchenbauer et al. (2023) and AAR Aaronson (2023) exhibit very low p -values in the IP infringement Detection scenario, regardless of logits or sample learning distillation (Section 4.4), indicating their effectiveness.

However, after full-parameter continual pre-training, the p -values for all watermark methods rise to 1, indicating a complete loss of their IP Infringement Detection capability. Despite this, certain methods are robust against LoRA-based fine-tuning. Moreover, all watermarks maintain very low p -values after applying other fine-tuning methods (such as instruction tuning, DPO, and RLHF) with fewer

training steps, suggesting that minor fine-tuning is insufficient to completely negate the performance of distillation-based watermark methods in IP Infringement detection.

Notably, Table 1 primarily uses English for continual pre-training watermark strength testing. Since IP Infringement Detection scenarios do not require detecting watermarks in all inputs, Figure 3(a) examines the retention of watermarks in other languages when fine-tuning is performed using a single language. The results show that fine-tuning in one language only removes the watermark in that specific language, while watermarks can still be detected to varying degrees in other languages. Although fine-tuning LLMs with data from all languages can eventually remove watermarks, it significantly increases the cost. Hence, adding watermarks in low-resource languages may offer a more robust and stealthy solution for distillation-based watermarking for IP Infringement Detection.

5.5 DISTILLATION FOR TEXT DETECTION (RQ3)

This section investigates the effectiveness of inference-time watermark distillation in IP infringement scenarios and its performance in detecting LLM-generated text. Although Gu et al. (2024) have conducted relevant research, we evaluate the robustness of watermarking methods in practical user fine-tuning scenarios and examine their impact on LLM performance.

Table 1 demonstrates that the inference-time watermark distillation method achieves very low p -values in the generated text detection scenario without continual fine-tuning, indicating its effectiveness. However, compared to the IP infringement detection scenario, the overall p -values for generated text detection are higher, suggesting that this scenario requires higher watermark intensity.

Additionally, Table 1 shows that various continual fine-tuning methods easily remove watermarks in the generated text detection scenario, as evidenced by the increased proportion of light-colored areas (high p -value). Full-parameter fine-tuning significantly reduces watermark strength, with complete removal after continual pretraining. Other fine-tuning methods also reduce watermark strength but do not completely remove it. Interestingly, using LoRA fine-tuning enhances watermark retention, showing partial retention even after continual pretraining and higher retention under other fine-tuning methods. Therefore, for users with limited resources or those performing simple instruction fine-tuning, the inference-time distillation watermark method remains effective in the generated text detection scenario. The lower part of Figure 3(c) illustrates that continual pretraining will definitely remove the inference-time distillation watermark. Moreover, similar to the IP Infringement Detection scenario, if continual pretraining is conducted in only one language, more watermark retention will be observed in other languages, as shown in Figure 3(b).

Finally, as depicted in Table 2, all inference-time distillation watermark methods greatly impact LLM performance. The methods based on KGW and AAR result in a 1.8% and 2.4% decrease in reasoning & understanding and an increase in PPL by 1.6 and 4.4. Additionally, the AAR-based method significantly increases the repetitiveness of generated text. In summary, inference-time distillation watermark methods have a greater impact on LLM performance than backdoor watermark methods. Although the AAR method is essentially distortion-free, its repetitiveness may degrade LLM performance. Also, we found that distilling the KGW watermarking algorithm is more robust in continual fine-tuning than distilling AAR in both scenarios, with less impact on LLM performance.

5.6 DISCUSSION

Our evaluation of two watermarking algorithms shows that neither can fully detect misuse of open-source LLMs. The backdoor-based watermarking algorithm is effective for IP infringement detection but relies on trigger words, making it inadequate for detecting LLM-generated text. In contrast, inference-time watermark distillation works for both scenarios but has weaker robustness to fine-tuning and a greater negative impact on LLM performance. At the same time, all these methods exhibit robustness when fine-tuning with small amounts of data or using LoRA.

6 CONCLUSION

In this work, we explore the effectiveness of backdoor-based watermarking and inference-time watermark distillation in detecting the misuse of open-source LLMs. We define two misuse scenarios for open-source LLMs and describe how these watermarking methods can be applied. Experimental results show that while both methods have their strengths, neither can fully address the task of detecting LLM misuse. Future research needs to develop better watermarking algorithms.

ACKNOWLEDGEMENTS

Yijie Xu acknowledges the support from the modern matter laboratory, HKUST(GZ).

REFERENCES

- Scott Aaronson. Watermarking of large language models. Presented at the Large Language Models and Transformers Workshop at Simons Institute for the Theory of Computing, 2023. URL <https://www.youtube.com/watch?v=2Kx9jbSMZqA>.
- AI@Meta. Llama 3 model card, 2024. URL https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 1952.
- Canyu Chen and Kai Shu. Combating misinformation in the age of llms: Opportunities and challenges. *AI Magazine*, 2024.
- Yankai Chen, Yixiang Fang, Qiongyan Wang, Xin Cao, and Irwin King. Deep structural knowledge exploitation and synergy for estimating node importance value on heterogeneous information networks. In *AAAI*, 2024a.
- Yankai Chen, Quoc-Tuan Truong, Xin Shen, Jin Li, and Irwin King. Shopping trajectory representation learning with pre-training for e-commerce customer understanding and recommendation. In *SIGKDD*, 2024b.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Lu Dai, Hao Liu, and Hui Xiong. Improve dense passage retrieval with entailment tuning. In *EMNLP*, 2024.
- Lu Dai, Yijie Xu, Jinhui Ye, Hao Liu, and Hui Xiong. Seper: Measure retrieval utility through the lens of semantic perplexity reduction. In *ICLR*, 2025.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Chenchen Gu, Xiang Lisa Li, Percy Liang, and Tatsunori Hashimoto. On the learnability of watermarks for language models. In *ICLR*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Zhiwei He, Binglin Zhou, Hongkun Hao, Aiwei Liu, Xing Wang, Zhaopeng Tu, Zhuosheng Zhang, and Rui Wang. Can watermarks survive translation? on the cross-lingual consistency of text watermark for large language models. In *ACL*, 2024.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *ICLR*, 2021.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022.
- Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. Unbiased watermark for large language models. In *ICLR*, 2024.
- Dmitri Iourovitski, Sanat Sharma, and Rakshak Talwar. Hide and seek: Fingerprinting large language models with evolutionary learning, 2024.

- Heng Jin, Chaoyu Zhang, Shanghao Shi, Wenjing Lou, and Y. Thomas Hou. Profingo: A fingerprinting-based intellectual property protection scheme for large language models. In *2024 IEEE Conference on Communications and Network Security (CNS)*, 2024.
- Nikola Jovanović, Robin Staab, and Martin Vechev. Watermark stealing in large language models. *ICML*, 2024.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. In *ICML*, 2023.
- Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. Robust distortion-free watermarks for language models. *Transactions on Machine Learning Research*, 2024.
- Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdoo Yun, Jamin Shin, and Gunhee Kim. Who wrote this code? watermarking for code generation. In *ACL*, 2024.
- Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Pre-trained language models for text generation: A survey. *ACM Computing Surveys*, 2024a.
- Linyang Li, Botian Jiang, Pengyu Wang, Ke Ren, Hang Yan, and Xipeng Qiu. Watermarking llms with weight quantization. In *EMNLP Findings*, 2023.
- Zhonghao Li, Xuming Hu, Aiwei Liu, Kening Zheng, Sirui Huang, and Hui Xiong. Refiner: Restructure retrieved content efficiently to advance question-answering capabilities. In *EMNLP Findings*, 2024b.
- Aiwei Liu, Leyi Pan, Xuming Hu, Shuang Li, Lijie Wen, Irwin King, and Philip S. Yu. An unforgeable publicly verifiable watermark for large language models. In *ICLR*, 2024a.
- Aiwei Liu, Leyi Pan, Xuming Hu, Shiao Meng, and Lijie Wen. A semantic invariant robust watermark for large language models. In *ICLR*, 2024b.
- Aiwei Liu, Leyi Pan, Yijian Lu, Jingjing Li, Xuming Hu, Xi Zhang, Lijie Wen, Irwin King, Hui Xiong, and Philip Yu. A survey of text watermarking in the era of large language models. *ACM Computing Surveys*, 2024c.
- Aiwei Liu, Haoping Bai, Zhiyun Lu, Yanchao Sun, Xiang Kong, Xiaoming Simon Wang, Jiulong Shan, Albin Madappally Jose, Xiaojiang Liu, Lijie Wen, Philip S. Yu, and Meng Cao. TIS-DPO: Token-level importance sampling for direct preference optimization with estimated weights. In *ICLR*, 2025a.
- Aiwei Liu, Sheng Guan, Yiming Liu, Leyi Pan, Yifei Zhang, Liancheng Fang, Lijie Wen, Philip S. Yu, and Xuming Hu. Can watermarked LLMs be identified by users via crafted prompts? In *ICLR*, 2025b.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024d.
- Xuankai Liu, Fengting Li, Bihan Wen, and Qi Li. Removing backdoor-based watermarks in neural networks with limited data. In *International Conference on Pattern Recognition (ICPR)*, 2021.
- Yijian Lu, Aiwei Liu, Dianzhi Yu, Jingjing Li, and Irwin King. An entropy-based text watermarking detection method. In *ACL*, 2024.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. In *ICLR*, 2017.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *NeurIPS*, 2022.
- Leyi Pan, Aiwei Liu, Zhiwei He, Zitian Gao, Xuandong Zhao, Yijian Lu, Binglin Zhou, Shuliang Liu, Xuming Hu, Lijie Wen, Irwin King, and Philip S. Yu. MarkLLM: An open-source toolkit for LLM watermarking. In *EMNLP*, 2024.

- Leyi Pan, Aiwei Liu, Shiyu Huang, Yijian Lu, Xuming Hu, Lijie Wen, Irwin King, and Philip S Yu. Can llm watermarks robustly prevent unauthorized knowledge distillation? *arXiv preprint arXiv:2502.11598*, 2025a.
- Leyi Pan, Aiwei Liu, Yijian Lu, Zitian Gao, Yichen Di, Shiyu Huang, Lijie Wen, Irwin King, and Philip S. Yu. Waterseeker: Pioneering efficient detection of watermarked segments in large documents, 2025b.
- Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. Reasoning with language model prompting: A survey. In *ACL*, 2023.
- Zexuan Qiu, Zijing Ou, Bin Wu, Jingjing Li, Aiwei Liu, and Irwin King. Entropy-based decoding for retrieval-augmented large language models. In *MINT: Foundation Model Interventions*, 2024a.
- Zexuan Qiu, Jieming Zhu, Yankai Chen, Guohao Cai, Weiwen Liu, Zhenhua Dong, and Irwin King. Ease: Learning lightweight semantic feature adapters from large language models for ctr prediction. In *CIKM*, 2024b.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 2020.
- Jie Ren, Han Xu, Pengfei He, Yingqian Cui, Shenglai Zeng, Jiankun Zhang, Hongzhi Wen, Jiayuan Ding, Hui Liu, Yi Chang, et al. Copyright protection in generative ai: A technical perspective. *arXiv preprint arXiv:2402.02333*, 2024.
- Mark Russinovich and Ahmed Salem. Hey, that’s my model! introducing chain & hash, an llm fingerprinting technique. *arXiv preprint arXiv:2407.10887*, 2024.
- Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*, 2023.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 2021.
- Masoumeh Shafieinejad, Nils Lukas, Jiaqi Wang, Xinda Li, and Florian Kerschbaum. On the robustness of backdoor-based watermarking in deep neural networks. In *Proceedings of the 2021 ACM workshop on information hiding and multimedia security*, pp. 177–188, 2021.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. Neural text generation with unlikelihood training. In *ICLR*, 2020.
- Jiashu Xu, Mingyu Ma, Fei Wang, Chaowei Xiao, and Muhao Chen. Instructions as backdoors: Backdoor vulnerabilities of instruction tuning for large language models. In *NAACL*, 2024a.
- Jiashu Xu, Fei Wang, Mingyu Ma, Pang Wei Koh, Chaowei Xiao, and Muhao Chen. Instructional fingerprinting of large language models. In *NAACL*, 2024b.
- Zhenhua Xu, Wenpeng Xing, Zhebo Wang, Chang Hu, Chen Jie, and Meng Han. Fp-vec: Fingerprinting large language models via efficient vector addition. *arXiv preprint arXiv:2409.08846*, 2024c.

- Yibo Yan, Jiamin Su, Jianxiang He, Fangteng Fu, Xu Zheng, Yuanhuiyi Lyu, Kun Wang, Shen Wang, Qingsong Wen, and Xuming Hu. A survey of mathematical reasoning in the era of multimodal large language model: Benchmark, method & challenges. *arXiv preprint arXiv:2412.11936*, 2024.
- Zhiguang Yang and Hanzhou Wu. A fingerprint for large language models. *arXiv preprint arXiv:2407.01235*, 2024.
- Jinhui Ye, Zihan Wang, Haosen Sun, Keshigeyan Chandrasegaran, Zane Durante, Cristobal Eyzaquirre, Yonatan Bisk, Juan Carlos Niebles, Ehsan Adeli, Li Fei-Fei, Jiajun Wu, and Manling Li. Re-thinking temporal search for long-form video understanding. In *CVPR*, 2025.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In *ACL*, 2019.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. Evaluating large language models at evaluating instruction following. In *ICLR*, 2024.
- Xuandong Zhao, Prabhanjan Vijendra Ananth, Lei Li, and Yu-Xiang Wang. Provable robust watermarking for AI-generated text. In *ICLR*, 2024.

Part I

Appendix

Table of Contents

A	Details of Watermark Methods	16
A.1	KGW Watermarking	16
A.2	AAR Watermarking	16
B	Threat Models for LLM Generated Text Detection	17
C	Threat Models for Backdoor-based Watermark	17
D	Detailed Accuracy Reference for Table 1	18
E	Dataset Descriptions	19
F	Details of continual fine-tuning Method	19
F.1	Continual Pre-training	19
F.2	Continual Instruction Tuning	19
F.3	Continual Direct Preference Learning	20
F.4	Continual Reinforcement Learning From Human Feedback	20
F.5	Low Rank Adaptation (LoRA) Fine-tuning Method	20
G	Details of Generation Capability Metrics	21
H	License Overview of Various Open-source LLMs	21
H.1	Meta Llama Series	21
H.2	Command R Series	22
H.3	01.ai Yi Series	24
I	System Specifications for Reproducibility	25

A DETAILS OF WATERMARK METHODS

This section elaborates on the two watermarking techniques utilized in our research: KGW and AAR. Each subsection provides background information on these techniques, motivation, method description, and detection methodology.

A.1 KGW WATERMARKING

Kirchenbauer et al. (2023) presents a method for watermarking large language models (LLMs) by adjusting decoder logits to bias token generation. Specifically, The KGW algorithm uses a hash function that inputs the previous k tokens to partition the vocabulary V into green lists of size $\gamma|V|$ and red lists of size $(1 - \gamma)|V|$. Favoring green tokens during sampling embeds a watermark in the generated text. A positive δ is added to the logits of green tokens, increasing their sampling probability. Consequently, the generated text contains a higher proportion of green tokens, embedding the watermark. The pseudocode implementation of the algorithm is as follows:

Algorithm 1 KGW Watermarking Algorithm

Require: Vocabulary V , hyperparameters k, γ, δ , LLM model LLM , hash function HashFunction .

```

1: Initialize text  $\leftarrow []$ 
2: while not end of generation do
3:   prev_tokens  $\leftarrow$  last  $k$  tokens from  $\text{text}$ 
4:   hash  $\leftarrow \text{HashFunction}(\text{prev\_tokens})$ 
5:   Partition  $V$  into green list  $G$  of size  $\gamma|V|$  and red list  $R$  of size  $(1 - \gamma)|V|$  using hash
6:   logits  $\leftarrow \text{LLM}(\text{text})$ 
7:   for each token  $t$  in  $G$  do
8:     logits[ $t$ ]  $\leftarrow$  logits[ $t$ ] +  $\delta$ 
9:   end for
10:  next_token  $\leftarrow$  Sample from logits
11:  Append next_token to text
12: end while
13: return text

```

Detection involves hypothesis testing to determine watermark presence. If human-written, the green token frequency should be near γ ; if watermarked, it should be significantly higher. The test statistic is:

$$z = \frac{|s|_G - \gamma T}{\sqrt{T\gamma(1 - \gamma)}}, \quad (4)$$

where $|s|_G$ is the number of green tokens, T is the text length, and γ is the green list size. Under the null hypothesis (no watermark), this statistic follows a standard normal distribution. A p -value below a significance level (e.g., 0.05) indicates a watermark.

A.2 AAR WATERMARKING

Aaronson (2023) embeds watermarks in the generated text by biasing the selection of tokens based on their hash scores. Given a key ξ , the algorithm computes a hash score $r_i \in [0, 1]$ for each of the first k tokens, where the scores are uniformly distributed. For each token i , the algorithm calculates r_i^{1/p_i} , where p_i is the original probability assigned by the language model to that token. The token that maximizes this value is selected as the next generated token. This process ensures that the chosen token has both a high original probability p_i and a high hash score r_i . The pseudocode implementation of the algorithm can be found in Appendix A.2.

Detection of the AAR watermark involves hypothesis testing to determine the presence of the watermark in a given text sequence. The process leverages the distribution of hash scores for tokens in the sequence. The method computes hash scores r for each token x_t using the previous k tokens and a predetermined key ξ . The cumulative test statistic S is calculated as follows:

Algorithm 2 AAR Watermarking Algorithm**Require:** Key ξ , hyperparameter k , LLM model LLM, hash function HashFunction.

```

1: Initialize text  $\leftarrow []$ 
2: while not end of generation do
3:   prev_tokens  $\leftarrow$  last  $k$  tokens from text
4:    $r \leftarrow$  HashFunction(prev_tokens,  $\xi$ )
5:   logits  $\leftarrow$  LLM(text)
6:   for each token  $i$  in logits do
7:      $p_i \leftarrow$  logits[ $i$ ]
8:     score $_i \leftarrow r_i^{1/p_i}$ 
9:   end for
10:  next_token  $\leftarrow$  argmax $_i$  score $_i$ 
11:  Append next_token to text
12: end while
13: return text

```

$$S = \sum_{t=k+1}^{\text{len}(x)} -\log(1 - r_{x_t}), \quad (5)$$

where $\text{len}(x)$ is the length of the sequence, and r_{x_t} is the hash score for the token at position t . Under the null hypothesis (i.e., the text is not watermarked), this test statistic follows a Gamma distribution with shape parameter $\text{len}(x) - k$ and scale parameter 1. The p -value for the observed sequence is then computed as:

$$p\text{-value} = 1 - F_G(S), \quad (6)$$

where S is the cumulative test statistic computed from the sequence. If the p -value is below a predetermined significance level (e.g., 0.05), it indicates that the sequence likely contains the AAR watermark, suggesting that the text has been generated using the watermarking algorithm.

B THREAT MODELS FOR LLM GENERATED TEXT DETECTION

In Table 1, we primarily studied the threat model of continual-tuning LLMs. This section discusses other threat models in the context of LLM-generated text detection.

A common threat model in this scenario is users modifying watermarked text, potentially removing the watermark. To investigate this, we tested the p -values for detecting watermarked text after it was rewritten using the GPT-3.5-TURBO API. We used the following prompt, and the modified p -value statistics are shown in Table 4. As observed, nearly all texts generated by fine-tuned LLMs have their watermarks completely removed after rewriting, highlighting significant room for improvement in current methods of handling text modifications.

Prompt used in for GPT rewriting

System: You are a helpful assistant.

User: Rewrite the following text in English: $\{text\}$

Additionally, there may be other threat models for LLM Generated Text Detection, such as spoofing attacks Sadasivan et al. (2023) and watermark stealing Jovanović et al. (2024). These have been extensively studied in the context of inference time watermarking Liu et al. (2024a). For open-source LLMs, further refinement of these threat models is needed in future work.

C THREAT MODELS FOR BACKDOOR-BASED WATERMARK

In Section 5.3, we have demonstrated that backdoor-based watermarking is an effective method for IP infringement detection and is robust against continual fine-tuning of LLMs.

Target LLM	Watermark Methods	<i>P</i> -Value↓ (Origin)	<i>W. Continual PT</i>		<i>W. Continual IT</i>		<i>W. Continual IT+DPO</i>		<i>W. Continual IT+RLHF</i>	
			Full↓	LoRA↓	Full↓	LoRA↓	Full↓	LoRA	Full↓	LoRA↓
GPT-3.5 Rewritten Metrics										
LLAMA2-7B	KGW-Logits	4e-2	5e-1	4e-1	2e-1	1e-1	2e-1	2e-1	3e-1	3e-1
	Aar-Logits	2e-1	6e-1	5e-1	5e-1	4e-1	5e-1	3e-1	4e-1	3e-1
	KGW-Sampling	2e-2	5e-1	4e-1	3e-1	2e-1	4e-1	3e-1	4e-1	3e-1
	Aar-Sampling	7e-4	5e-1	3e-1	3e-1	4e-3	4e-1	3e-1	4e-1	3e-1
LLAMA3-8B	KGW-Logits	3e-2	5e-1	4e-1	1e-1	8e-2	3e-1	1e-1	3e-1	2e-1
	Aar-Logits	1e-1	5e-1	4e-1	3e-1	2e-1	3e-1	3e-1	4e-1	3e-1
	KGW-Sampling	5e-2	5e-1	4e-1	2e-1	1e-1	3e-1	2e-1	4e-1	2e-1
	Aar-Sampling	3e-1	5e-1	4e-1	4e-1	3e-1	4e-1	3e-1	4e-1	3e-1

Table 4: The p -value significance of watermarking methods under GPT-3.5 rewritten metrics, including the unmodified p -value, as well as the p -value significance after further continual pre-training, instruction tuning, DPO, and RLHF optimization. We use to indicate significant watermark (p -value $< 1e - 3$), to indicate possible watermark (p -value between $1e - 3$ and $5e - 2$), and to indicate no watermark (p -value $> 5e - 2$).

Target-LLM	Watermark Methods	<i>P</i> -Value↓ (Origin)	<i>W. Continual PT</i>		<i>W. Continual IT</i>		<i>W. Continual IT+DPO</i>		<i>W. Continual IT+RLHF</i>	
			Full↓	LoRA↓	Full↓	LoRA↓	Full↓	LoRA	Full↓	LoRA↓
Scenario 1: Open-Source LLM Intellectual Property Detection (Section 3.1)										
LLAMA2-7B	Backdoor-PT	33.0%	18.0%	22.0%	33.5%	32.5%	30.5%	31.5%	33.0%	32.5%
	Backdoor-IT	34.0%	N/A	N/A	28.5%	29.5%	29.0%	30.0%	30.5%	31.5%
	KGW-Logits	98.0%	53.4%	68.5%	89.9%	94.1%	85.2%	91.9%	84.6%	90.5%
	Aar-Logits	96.5%	50.4%	56.3%	79.4%	90.9%	81.4%	89.3%	80.3%	87.9%
	KGW-Sampling	94.4%	53.3%	61.3%	74.0%	92.9%	70.8%	73.5%	83.4%	87.9%
	Aar-Sampling	95.5%	50.28%	56.9%	73.4%	94.9%	61.1%	77.0%	78.9%	83.4%
LLAMA3-8B	Backdoor-PT	82.5%	33.5%	42.5%	75.5%	78.5%	74.0%	77.0%	75.5%	80.5%
	Backdoor-SFT	83.5%	N/A	N/A	80.5%	82.5%	79.5%	80.0%	79.0%	81.5%
	KGW-Logits	98.5%	53.4%	71.3%	84.9%	90.8%	80.1%	89.7%	79.8%	88.7%
	Aar-Logits	98.5%	51.4%	57.4%	81.4%	87.3%	78.9%	86.5%	78.0%	85.7%
	KGW-Sampling	97.8%	52.3%	65.1%	81.0%	84.6%	77.8%	87.6%	78.6%	86.7%
	Aar-Sampling	96.5%	50.6%	55.2%	72.9%	81.4%	73.1%	81.3%	72.9%	80.7%

Table 5: The accuracy of watermarking methods under two scenarios, including the unmodified accuracy, as well as the accuracy after further continual pre-training, instruction tuning, DPO, and RLHF optimization.

The fine-tuning methods discussed in Table 1 assume that users are completely unaware of the trigger’s existence. Under this assumption, removing the backdoor through fine-tuning is very difficult. However, removing the backdoor watermark becomes easy if users somehow become aware of the specific trigger. Therefore, future research should focus on making the trigger as undetectable as possible (even rare word combinations are at risk of being discovered) and on verifying backdoor watermarks without exposing the trigger.

D DETAILED ACCURACY REFERENCE FOR TABLE 1

In Table 1, we only show the p -value metrics for the Open-Source LLM Intellectual Property Detection scenario. For reference, Table 5 provides the corresponding original accuracy for each p -value. The backdoor method indicates the correct trigger rate, while the inference-time watermark distillation method refers to the accuracy of watermark and human text at a p -value of 0.05, as described in Section 4.4.2.

E DATASET DESCRIPTIONS

This section presents the datasets used in our experiments, detailing their names, sizes, sources, and key characteristics. The datasets are categorized based on the capabilities they assess, including understanding and reasoning capabilities, generation capability, continual pre-training, supervised instruct tuning, and alignment optimization. The descriptions are summarized in Table 6.

Dataset	Size	Source	Key Characteristics
Understanding and Reasoning Capabilities			
ARC-EASY Clark et al. (2018)	5,197	AI2	Elementary science questions
ARC-CHALLENGE Clark et al. (2018)	2,590	AI2	More difficult science questions
HELLASWAG Zellers et al. (2019)	59.950	Web	Commonsense reasoning tasks
MMLU Hendrycks et al. (2021)	231,400	Diverse	Multi-domain, multiple-choice questions
WINOGRANDE Sakaguchi et al. (2021)	~44K	Crowd-sourced	Pronoun resolution challenges
Generation Capability			
WIKITEXT Merity et al. (2017)	3,708,608	Wikipedia	High-quality articles from Wikipedia
Continual Pre-training			
C4 Raffel et al. (2020)	~10B	Web	Massive, clean crawled corpus
Supervised Instruct Tuning			
ALPACA Taori et al. (2023)	~52K	Self-generated	Instruction-following dataset from GPT-4
Alignment Optimization			
HH-RLHF Bai et al. (2022)	169,352	Human Feedback	Human preference data for RLHF

Table 6: Dataset Descriptions Across Multiple Capabilities

F DETAILS OF CONTINUAL FINE-TUNING METHOD

This section details the various continual fine-tuning methods used in this work, including continual pre-training, continual instruction tuning, continual direct preference learning, and continual reinforcement learning from human feedback. We also discuss the use of the LoRA fine-tuning approach.

F.1 CONTINUAL PRE-TRAINING

Continual pre-training involves unsupervised training of the language model on a large corpus. Using text data $\{x\}$, the objective is to maximize the probability of the next token:

$$\mathcal{L}_{PT} = - \sum_{i=1}^n \log P(x_i | x_{<i}; M), \quad (7)$$

where M are the model parameters and n is the length of the text x . This autoregressive training helps the model learn the statistical features of the text.

F.2 CONTINUAL INSTRUCTION TUNING

Continual instruction tuning builds on pre-training by using an instruction dataset $\{(x, y)\}$ to fine-tune the model so it can generate appropriate answers to given instructions. Here, x is the instruction or question, and y is the corresponding answer. The loss function maximizes the conditional probability:

$$\mathcal{L}_{IT} = - \sum_{i=1}^m \log P(y_i | x_i; M), \quad (8)$$

where m is the size of the training data. This approach teaches the model to explicitly generate results based on instructions.

F.3 CONTINUAL DIRECT PREFERENCE LEARNING

Continual direct preference learning (DPO) Rafailov et al. (2023) uses human-labeled preference data $\{(x, y_1, y_2)\}$ to learn preferences, where y_1 and y_2 are two candidate answers generated by the model for x . The labeled data indicates whether y_1 is preferred over y_2 or vice versa. The training minimizes the pairwise ranking loss:

$$\mathcal{L} = -\mathbb{E}_{(x, y_1, y_2) \sim \mathcal{D}_{pref}} \left[\log \sigma \left(\beta \log \frac{\pi_M(y_1 | x)}{\pi_{ref}(y_1 | x)} - \beta \log \frac{\pi_M(y_2 | x)}{\pi_{ref}(y_2 | x)} \right) \right], \quad (9)$$

where σ is the logistic function, β is a scaling parameter, and π_M and π_{ref} are the probability distributions of the current model and the reference model, respectively.

F.4 CONTINUAL REINFORCEMENT LEARNING FROM HUMAN FEEDBACK

Reinforcement Learning from Human Feedback (RLHF) Ouyang et al. (2022) involves two main stages. First, a reward model is trained using a dataset of human-labeled preferences. Second, this reward model, combined with the PPO algorithm, is used to train a language model via reinforcement learning.

To train the reward model, the data includes an input x and two outputs y_w and y_l , where y_w is the preferred response. The reward model, represented as $r^*(y, x)$, uses the Bradley-Terry (BT) model Bradley & Terry (1952) to express preference probabilities:

$$P(y_w \succ y_l | x) = \frac{\exp(r^*(y_w, x))}{\exp(r^*(y_w, x)) + \exp(r^*(y_l, x))}. \quad (10)$$

Here, $P(y_w \succ y_l | x)$ is the probability that y_w is preferred over y_l given x . The reward model $r^*(y, x)$ scores each potential output y . The BT model is commonly used for pairwise comparisons to represent these preferences.

Given the training data $\{x, y_w, y_l\}_i^N$, the reward model $r_M(y, x)$ is trained using the following loss function:

$$\mathcal{L}_R(r_M, D) = -\mathbb{E}_{(x, y_w, y_l)} [\log \sigma(r_M(y_w, x) - r_M(y_l, x))]. \quad (11)$$

Here, σ is the logistic function, and the expectation is over triplets (x, y_w, y_l) from D . This loss function pushes the model to score the preferred output y_w higher than y_l for a given x . Minimizing this loss enables the reward model to learn human preferences.

In the reinforcement learning phase, the trained reward model guides the language model training. The aim is to optimize the language model's policy π_M to maximize the expected reward from r_M , while keeping outputs close to a reference policy π_{ref} . This is achieved with the following objective:

$$\max_{\pi_M} \mathbb{E}_{x, y \sim \pi_M} [r_M(y, x)] - \beta \mathbb{D}_{KL} [\pi_M \parallel \pi_{ref}]. \quad (12)$$

This balances enhancing the language model's performance and maintaining alignment with human preferences.

F.5 LOW RANK ADAPTATION (LoRA) FINE-TUNING METHOD

We utilize both full-parameter tuning and LoRA (Low-Rank Adaptation) fine-tuning for continual fine-tuning. The core idea of LoRA fine-tuning is to adjust only the low-rank projection matrices while keeping the other parameters of the pre-trained model fixed.

Assume the weight matrix of the original model is $\mathbf{W} \in \mathbb{R}^{d \times k}$. LoRA defines two low-rank matrices $\mathbf{A} \in \mathbb{R}^{d \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times k}$, where $r \ll \min(d, k)$. The augmented weight matrix is:

$$\mathbf{W}_{lora} = \mathbf{W} + \mathbf{AB}. \quad (13)$$

During fine-tuning, only matrices **A** and **B** are updated, while **W** remains unchanged. This method retains the knowledge of the original model, significantly reduces the number of parameters to be fine-tuned, and speeds up the training process. Due to its lower resource requirements, LoRA is often preferred for fine-tuning open-source models.

G DETAILS OF GENERATION CAPABILITY METRICS

To evaluate the generation capability of the fine-tuned models, we employed two metrics: Perplexity (PPL) and Sequence Repetition of 3-grams (Seq-Rep-3), computed on the WIKITEXT dataset Merity et al. (2017).

Perplexity (PPL) measures the fluency and coherence of the generated text. A lower PPL indicates better predictive capability and higher quality. Specifically, we calculate PPL using the LLAMA2-70B AI@Meta (2024) model to assess the language model’s effectiveness in predicting tokens. The calculation of PPL is defined as follows:

$$\text{PPL}(x) = \exp \left(-\frac{1}{n} \sum_{i=1}^n \log P(x_i \mid x_{<i}; M) \right), \quad (14)$$

where $x = (x_1, x_2, \dots, x_n)$ is the sequence, and M denotes the language model parameters. To ensure a fair assessment of fluency rather than benefiting from repetitive text, we set a no-repeat n -gram constraint ($n = 5$) during evaluation. This setting prevents artificially low PPL scores due to repetition.

Sequence Repetition (Seq-Rep-3) evaluates text diversity by measuring the proportion of repeated 3-grams in generated sequences Welleck et al. (2020). A lower Seq-Rep-3 indicates more diverse and natural text generation. Formally, Seq-Rep-3 is defined as:

$$\text{Seq-Rep-3}(x) = \frac{\text{number of repeated 3-grams in sequence } x}{\text{total number of 3-grams in sequence } x}. \quad (15)$$

These metrics provide complementary insights into the quality and diversity of model-generated text. PPL measures fluency, with lower values indicating better token prediction and coherence. However, low PPL alone may not reflect text diversity, as repetition can artificially reduce it. Seq-Rep-3 addresses this by quantifying 3-gram repetitions, with lower values indicating more varied text.

H LICENSE OVERVIEW OF VARIOUS OPEN-SOURCE LLMs

This appendix provides an overview of the licensing terms for several open-source Large Language Models (LLMs), highlighting their alignments with our scenarios.

H.1 META LLAMA SERIES

Meta’s use policy for Llama 2¹ outlines several prohibited uses, which directly relate to the scenarios defined in our paper.

Scenario 1: Detecting IP Infringement

Policy 1.g

Engage in or facilitate any action or generate any content that infringes, misappropriates, or otherwise violates any third-party rights, including the outputs or results of any products or services using the Llama 2 Materials.

¹<https://ai.meta.com/llama/use-policy/>

- This policy directly relates to IP Detection, as it prohibits actions that infringe on intellectual property rights. The scenario’s goal of detecting unauthorized commercial use or copying of open-source LLMs is supported by this clause.

Scenario 2: Detecting Generated Text

Policy 1.a

“Engage in, promote, generate, contribute to, encourage, plan, incite, or further illegal or unlawful activity or content, such as:

- Violence or terrorism
- Exploitation or harm to children, including the solicitation, creation, acquisition, or dissemination of child exploitative content or failure to report Child Sexual Abuse Material”

- This policy is pertinent to Generated Text Detection, which aims to detect whether generated text from an open-source LLM contains illegal, harmful, or unethical content. The prohibition of generating such content aligns with the scenario’s goal of preventing misuse for disseminating harmful material.

Policy 1.c

“Engage in, promote, incite, or facilitate the harassment, abuse, threatening, or bullying of individuals or groups of individuals.”

- This policy relates to Generated Text Detection by prohibiting the generation of abusive or harassing content. The scenario’s goal of detecting harmful generated text includes identifying text that facilitates harassment or abuse.

Policy 3.a

“Intentionally deceive or mislead others, including use of Llama 2 related to the following:

- Generating, promoting, or furthering fraud or the creation or promotion of disinformation”

- This policy supports Generated Text Detection by addressing the misuse of LLMs to generate misleading or fraudulent content. Detecting such generated text aligns with the policy’s goal of preventing deception and misinformation.

Policy 3.e

“Representing that the use of Llama 2 or outputs are human-generated.”

- This policy underlines the importance of transparency in content generation. Generated Text Detection’s goal is to determine whether a text is generated by an LLM or its fine-tuned version aligns with ensuring users do not misrepresent AI-generated content as human-generated.

In summary, The Llama 2 use policy explicitly prohibits various activities related to both scenarios defined in our paper, particularly IP infringement and the generation of illegal or harmful content.

H.2 COMMAND R SERIES

Cohere R series is built on the language of business and is optimized for enterprise generative AI, search and discovery, and advanced retrieval. Their Cohere For AI Acceptable Use Policy² aligns with our scenario settings.

²<https://docs.cohere.com/docs/c4ai-acceptable-use-policy>

Scenario 1: Detecting IP Infringement

Cohere For AI Acceptable Use Policy

“Synthetic data for commercial uses: generating synthetic data outputs for commercial purposes, including to train, improve, benchmark, enhance or otherwise develop model derivatives, or any products or services in connection with the foregoing.”

- This policy is highly relevant to IP Detection, as it explicitly prohibits using models or their derivatives for commercial purposes without permission, which is a core concern of detecting IP infringement.

Scenario 2: Detecting Generated Text

Cohere For AI Acceptable Use Policy

“We expect users of our models or model derivatives to comply with all applicable local and international laws and regulations. Additionally, you may not use or allow others to use our models or model derivatives in connection with any of the following strictly prohibited use cases:”

- This policy establishes the baseline expectation that users must comply with all laws and regulations, which supports the detection of misuse in both scenarios. Generated Text Detection specifically aligns with preventing the generation and dissemination of illegal content.

Cohere For AI Acceptable Use Policy

“Harassment and abuse: engaging in, promoting, facilitating, or inciting activities that harass or abuse individuals or groups.”

- This policy supports Generated Text Detection by setting clear boundaries against generating content that could harass or abuse individuals or groups, aligning with the scenario’s goals of detecting unethical content.

Cohere For AI Acceptable Use Policy

“Violence and harm: engaging in, promoting, or inciting violence, threats, hate speech self-harm, sexual exploitation, or targeting of individuals based on protected characteristics.”

- This policy directly relates to Generated Text Detection, where the goal is to detect generated content that disseminates illegal, harmful, or unethical content. It provides a clear mandate against such misuse.

Cohere For AI Acceptable Use Policy

“Fraud and deception: misrepresenting generated content from models as human-created or allowing individuals to create false identities for malicious purposes, deception, or to cause harm, through methods including:

- propagation of spam, fraudulent activities such as catfishing, phishing, or generation of false reviews;
- creation or promotion of false representations of or defamatory content about real people, such as deepfakes; or
- creation or promotion of intentionally false claims or misinformation.”

- This is pertinent to both scenarios. For IP Detection, it addresses the misrepresentation of generated content as human-created, which can involve claiming an open-source LLM as a proprietary creation. For Generated Text Detection, it covers the generation of harmful or deceptive content.

H.3 01.AI YI SERIES

The Yi series is another open-source LLM that has demonstrated excellent performance on the LMSYS Chatbot Arena Leaderboard³. This series of LLMs has been developed by a Chinese company named “Lingyiwanwu.” The following user agreement⁴ contains sections that align with the assumptions in our scenarios.

Scenario 1: Detecting IP Infringement

Article 5 Clause 1

“Lingyiwanwu is the developer and operator of this product and enjoys all rights to the data, information, and outputs generated during the development and operation of this product within the scope permitted by laws and regulations, except where the relevant rights holders are entitled to rights according to law.”

- This clause asserts that Lingyiwanwu holds the rights to the outputs generated by the product, reinforcing the need to detect IP infringement when these rights are violated by unauthorized use or copying of the LLMs.

Article 5 Clause 2

“Unless otherwise agreed or stipulated by laws and regulations, you have the rights to the content generated based on the content you are entitled to upload and the rights to the content generated based on the uploaded content.”

- This clause delineates user rights to generated content, provided it is based on legally uploaded content, highlighting the importance of detecting if generated content infringes on existing IP rights.

Article 5 Clause 6

“You understand and promise that your input during the use of this product will not infringe on any person’s intellectual property rights, portrait rights, reputation rights, honor rights, name rights, privacy rights, personal information rights, etc. Otherwise, you will bear the risk and responsibility of infringement.”

- This clause ensures that users acknowledge their responsibility to avoid infringing on IP rights, aligning with the scenario’s assumption that detection mechanisms are needed to prevent such infringements.

Article 5 Clause 7

“If you add new data for model training, fine-tuning, and development during the use of this product, you will bear the resulting responsibilities.”

- This clause emphasizes user responsibility for any new data added for model training or fine-tuning, aligning with the scenario’s focus on detecting whether the generated text has been modified or fine-tuned from the original LLM.

³<https://arena.lmsys.org/>

⁴<https://platform.lingyiwanwu.com/useragreement>

Scenario 2: Detecting Generated Text

Article 4 Clause 1

“Based on your use of this product, Lingyiwanwu grants you a revocable, non-transferable, non-exclusive right to use this product. If you publish or disseminate content generated by this product, you should:

- Proactively verify the authenticity and accuracy of the output content to avoid spreading false information;
- Mark the output content as AI-generated in a prominent way to inform the public about the content synthesis;
- Avoid publishing and disseminating any output content that violates the usage norms of this agreement.”

- This clause mandates users to verify and label AI-generated content, ensuring transparency and preventing the misuse of generated text for harmful or illegal purposes, which aligns with the scenario’s goal of detecting and managing generated content responsibly.

Article 4 Clause 4

“Users are prohibited from engaging in certain behaviors, including but not limited to:

- **(5)** Inducing the generation of content that violates relevant laws and regulations or contains unfriendly outputs;
- **(7)** Developing products and services that compete with this product using this product;
- **(9)** Unauthorized removal or alteration of AI-generated labels or deep synthesis content labels.”

- These prohibitions directly support the scenario’s assumptions by preventing the generation and dissemination of harmful content, ensuring ethical use of the model, and maintaining the integrity of AI-generated labels for accountability.

I SYSTEM SPECIFICATIONS FOR REPRODUCTIVITY

Our experiments were conducted on high-performance servers, each equipped with either an Intel(R) Xeon(R) Platinum 8378A CPU @ 3.00GHz or an Intel(R) Xeon(R) Platinum 8358P CPU @ 2.60GHz, 1TB of RAM, and 4/6/8 NVIDIA A800 GPUs with 80GB memory. Machines with 4/8 GPUs are configured with the SXM4 version, while those with 6 GPUs use the PCIe version. The software environment included *Python* 3.11, *PyTorch* 2.3, and *NCCL* 2.20.5 for reproductivity.