

DITTO: A Spoofing Attack Framework on Watermarked LLMs via Knowledge Distillation

Hyeseon An¹, Shinwoo Park¹, Suyeon Woo², and Yo-Sub Han^{1,†}

¹Yonsei University, Seoul, Republic of Korea
{hsan, pshkhh, emmous}@yonsei.ac.kr

²Seoul National University, Republic of Korea
jenifer7933@snu.ac.kr

Abstract

The promise of LLM watermarking rests on a core assumption that a specific watermark proves authorship by a specific model. We demonstrate that this assumption is dangerously flawed. We introduce the threat of watermark spoofing, a sophisticated attack that allows a malicious model to generate text containing the authentic-looking watermark of a trusted, victim model. This enables the seamless misattribution of harmful content, such as disinformation, to reputable sources. The key to our attack is repurposing *watermark radioactivity*, the unintended inheritance of data patterns during fine-tuning, from a discoverable trait into an attack vector. By distilling knowledge from a watermarked teacher model, our framework allows an attacker to steal and replicate the watermarking signal of the victim model. This work reveals a critical security gap in text authorship verification and calls for a paradigm shift towards technologies capable of distinguishing authentic watermarks from expertly imitated ones. Our code is available at <https://github.com/hsannn/ditto>.

1 Introduction

Large language models (LLMs) are rapidly permeating nearly every aspect of modern life, from industrial applications and education to everyday personal use (Goyal et al., 2022; Xiao et al., 2024). Their ability to generate coherent, human-like text at scale makes them powerful tools for productivity and innovation, but also raises concerns about authenticity and trust (Perkins, 2023; Wang et al., 2024). As these models become increasingly integrated into daily workflows and public discourse, the ability to detect and verify LLM-generated text has become ever more crucial (Miresghallah et al., 2024; Go et al., 2025; Park et al., 2025).

As regulators in the United States (Biden, 2023) and the EU AI Act (EUAIAct, 2024) have signalled

[†] Corresponding author.

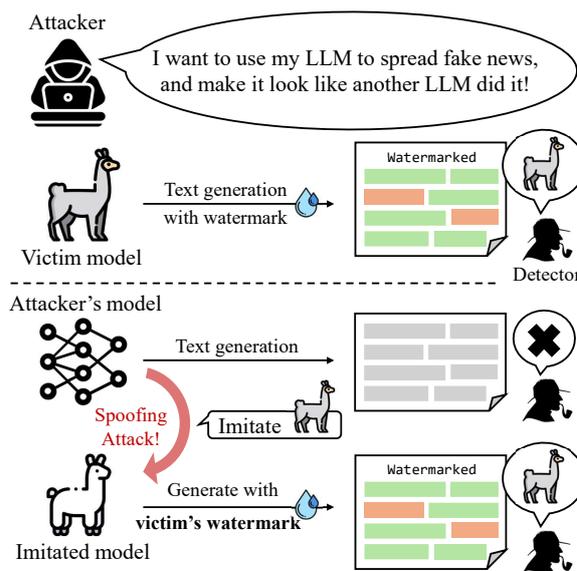


Figure 1: We ‘do something bad’ to reveal the fragility of current LLM watermarking schemes. An attacker’s model is trained to imitate the watermark of a trusted victim model, enabling it to generate content (e.g., fake news) that is then falsely attributed to the victim by a detector.

the need for clearer provenance of LLM-generated content, it becomes critical not only to determine whether a piece of text was authored by a human or generated by an AI system but also to identify which model generated it. Notably, major industry players, such as Meta, OpenAI, and Google Deepmind, have focused on watermarking as a practical tool for provenance verification (Reuters, 2023; Dathathri et al., 2024). This provenance is essential for accountability, since different models may be governed by distinct licenses, exhibit varying reliability, or present different risks of misuse. For these reasons, the development of LLM watermarking (Kirchenbauer et al., 2023; Kuditipudi et al., 2024; Lee et al., 2024; Gloaguen et al., 2025; Park et al., 2026; Kim et al., 2026) techniques has attracted substantial interest.

LLM watermarking—pioneered in earnest by Kirchenbauer et al. (2023)—refers to algorithms that embed an imperceptible but machine-detectable signal into model outputs so as to convey authorship information. In practice, watermarking has been explored through two main paradigms: logit-bias schemes that partition the vocabulary and skew generation probabilities for detectability (Kirchenbauer et al., 2023; Lu et al., 2024; Chen et al., 2024; Guo et al., 2024), and sampling-based methods that modify the decoding procedure to embed verifiable signals with minimal impact on text quality (Kuditipudi et al., 2024; Dathathri et al., 2024).

Despite its promise, LLM watermarking has not been free from criticism, and a growing body of work has examined the limitations and shortcomings of existing schemes through a variety of attacks. Paraphrasing-based attacks (Rastogi and Pruthi, 2024; Cheng et al., 2025) revealed straightforward weaknesses, while more sophisticated techniques have probed watermark robustness more deeply (Wu and Chandrasekaran, 2024a; Jovanović et al., 2024; Chen et al., 2025). By demonstrating how watermarks can be removed or reverse-engineered, these studies highlight the practical limitations of current schemes and motivate the development of stronger defenses.

Notably, most prior attack work has focused on *scrubbing attacks* that aim to remove embedded signals to evade detection. Relatively less attention has been given to *spoofing attacks*, which aim to make a third-party model generate outputs that contain the detectable watermark patterns of a target model, thereby misleading attribution. Spoofing attack presents qualitatively distinct and more malicious threat. Unlike scrubbing, it fabricates the watermark, leading detectors to falsely attribute text to the wrong model. Such misattribution produces a misleading sense of certainty, which is especially hazardous in the spread of fake news, political propaganda and other forms of disinformation. Therefore, a comprehensive assessment of watermark robustness should consider not only whether a watermark can be erased, but also how easily it can be forged.

We address this underexplored threat by proposing an attack framework for spoofing LLM watermarks under a practical black-box setting. Figure 1 illustrates our attack scenario: an adversarial user, seeking to spread disinformation such as fake news, may prefer that the text appear to originate from a trusted system like ChatGPT rather than from

their own local model. The attacker has neither the logits of the target model nor knowledge of its watermarking scheme. Under these realistic constraints, our approach exploits *watermark radioactivity* (Sander et al., 2024, 2025)—the phenomenon whereby training on synthetic outputs generated by a watermarked model causes downstream models to inherit the same watermark signal.

We introduce **DITTO** (**D**istilled watermark **I**mitation of a **T**argeted **T**eacher’s **O**utputs), a knowledge-distillation-based spoofing attack framework.

Through our approach, we show that watermarking is vulnerable not only to scrubbing attacks but also to a more fundamental threat, spoofing attacks. Consequently, our work highlights the need for security mechanisms that go beyond mere presence detection to include authenticity verification.

We summarize our contribution as follows:

- We are the first to repurpose watermark radioactivity for spoofing attacks, fundamentally repurposing a phenomenon intended for detection into a potent tool for misattribution.
- We show that our attack is highly adaptable, proving effective against not only n-gram-based watermarks but also sampling-based schemes.
- We find that the attack breaks the conventional trade-off between strength and quality. The spoofing intensity can be significantly increased without a discernible degradation in text quality, making the attack highly evasive.

2 Related Work

The popularization of LLMs highlights the necessity of robust safeguards for their responsible use. Similar to traditional cybersecurity, the field of LLM watermarking has co-evolved through a dynamic interplay between the development of new schemes and the attack systems that test their robustness.

2.1 LLM Watermarking

Early work on LLM watermarking centers on the green-red list paradigm, which partitions the vocabulary and biases sampling toward a pseudo-random ‘green’ set to yield an imperceptible yet statistically detectable signal (Kirchenbauer et al., 2023). Follow-ups adapt the idea to demanding settings: Lee et al. (2024) selectively applies the

green-red bias only when token entropy is high, markedly improving detectability and preserving code correctness under the low-entropy dynamics of code generation; multi-bit schemes add position allocation for user-traceable identifiers (Yoo et al., 2023); entropy-guided insertion further tunes watermark timing for robustness (Cai et al., 2025); and context-aware methods build semantically balanced green-red lists to resist paraphrasing while limiting quality loss (Guo et al., 2024). Complementary detection advances weight tokens by entropy to boost sensitivity in low-entropy regimes (Lu et al., 2024). Collectively, these works expand list-based watermarking from a single detection signal toward more accountable and semantically informed mechanisms.

In parallel, sampling-based watermarking modifies the sampling rule itself rather than fixing token partitions. Dathathri et al. (2024) introduced tournament sampling with configurable non-distortion guarantees and showed production-scale viability and improved detectability at a fixed false-positive rate. Distortion-free approach, such as Kuditipudi et al. (2024), alter the sampling process to embed watermarks while keeping the model’s output distribution unchanged, ensuring detectability even after edits. Together these lines trace a shift from fixed list biases to selective, semantic and sampling-theoretic designs that better balance robustness, attribution, and text quality.

2.2 Watermark Attack

As watermarking techniques for LLMs have been developed, so too have methods to attack them. Recent research on watermark attacks is evolving beyond simple text modification towards stealing techniques (Wu and Chandrasekaran, 2024a; Jovanović et al., 2024; Chen et al., 2025) that first uncover the watermark’s mechanism. This act of stealing enables more sophisticated **scrubbing** and **spoofing** attacks. Scrubbing attack aims to remove the watermark to evade detection and spoofing attack causes unwatermarked text to be accepted as watermarked. Jovanović et al. (2024) showed that black-box querying suffices to learn the watermarking rule well enough, calling into question optimistic deployment claims. Beyond heuristic estimation, Zhang et al. (2024) claimed that mixed-integer programming can recover KGW-style green lists and remove the watermark without detector access or watermark metadata.

A fast-growing line of research has focused on

watermark scrubbing under realistic constraints. One major approach is Wu and Chandrasekaran (2024b), which infers the watermark’s token ‘colors’ to strategically replace green tokens with non-green ones, while Huang et al. (2025) framed scrubbing as constrained optimization and distilled a proxy distribution to enable strong black-box watermark removal. Recently, Chen et al. (2025) estimated n-gram parameters (prefix length, strength, red/green lists) and provided theoretical bounds on distributional closeness after removal. Complementing removal, Pang et al. (2024) revealed that common design choices intended to make LLM watermarks robust can make the systems vulnerable to simple attacks, exposing a fundamental trade-off. However, prior work has concentrated on stealing and scrubbing, with spoofing often treated as a derivative capability or explored only in specific contexts. Our study, therefore, adopts a spoofing-first perspective to close this gap.

2.3 Watermark Radioactivity

Watermark radioactivity is the unintended inheritance of watermark signals when a student LLM is fine-tuned on watermarked teacher outputs. Even small fractions of watermarked data can make radioactivity statistically detectable, offering a tool for provenance auditing (Sander et al., 2024). However, Pan et al. (2025) demonstrated that adversaries can largely erase these signals through paraphrasing or inference-time neutralization. Building on this line of work, we leverage watermark radioactivity as a foundation for developing a scheme-agnostic spoofing methodology.

3 DITTO

We propose DITTO, a spoofing attack framework that leverages the radioactivity of LLM watermark, which is the tendency for a model to inherit statistical patterns when trained on the outputs of a watermarked source. The objective of DITTO is to imitate the unique watermark signal of a specific target model M_T , enabling an arbitrary external model M to generate text that appears as if it were watermarked by the target. Figure 2 demonstrates DITTO framework. First, we use knowledge distillation to transfer the watermark from M_T into a student model M_S . Then, we extract the watermark pattern by analyzing the difference in output distributions between M_S and its original version M_O . Lastly, we inject this extracted signal directly into

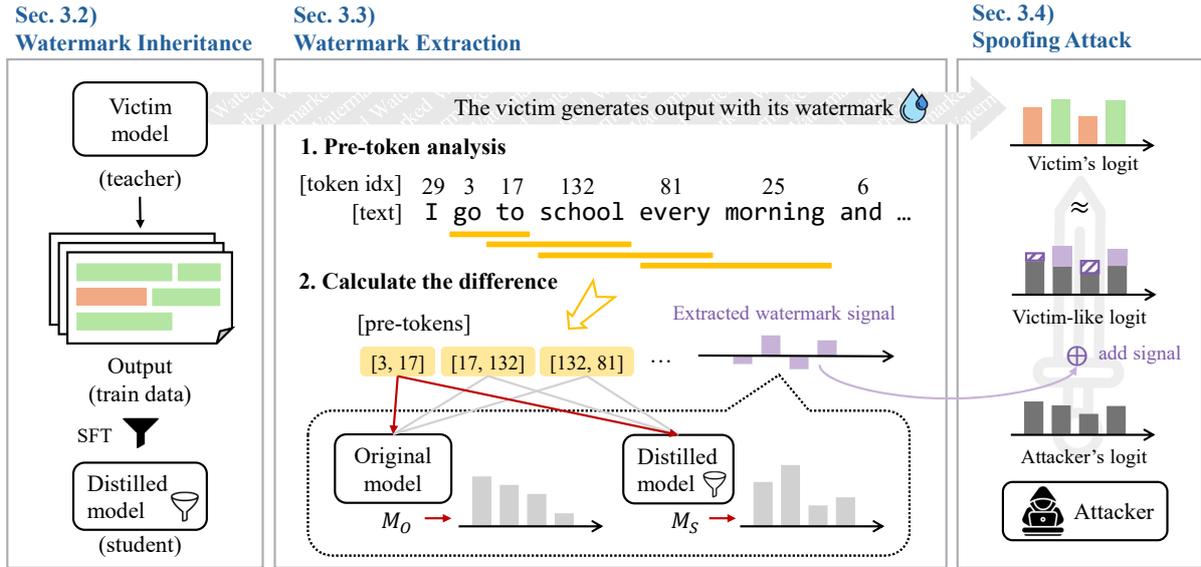


Figure 2: An overview of the DITTO framework. Our method consists of three main stages: (1) Watermark Inheritance, where a student model learns the teacher’s watermarked patterns via knowledge distillation; (2) Watermark Extraction, where the watermark is isolated by analyzing averaged logit differences, both globally and for specific text prefixes; and (3) the Spoofing Attack, where the extracted signal is added to the attacker’s logits to imitate the victim’s watermark.

the logits of M , successfully replicating the target’s watermark.

3.1 Preliminaries

Let M denote a large language model with a vocabulary V . Given a prefix (or context) $c = (x_1, \dots, x_{t-1})$, the model generates a probability distribution $P_M(x_t|c)$ for the next token x_t . This $P_M(x_t|c)$ is typically derived by applying a softmax function to the M ’s output logits $l_M(c) \in \mathbb{R}^{|V|}$.

Our attack aims to imitate a specific watermarked model, which we refer to as the target model M_T . This model embeds a statistical signal into its output by systematically altering its next-token probability distribution. Specifically, the watermarking process modifies the original logits with a context-dependent signal vector, $\delta(c)$. The original logits $l_{M_T}(c)$ are adjusted to produce the watermarked logits $l'_{M_T}(c)$:

$$l'_{M_T}(c) = l_{M_T}(c) + \delta(c). \quad (1)$$

Consequently, the final watermarked probability distribution is given by:

$$P'_{M_T}(x_t|c) = \text{softmax}(l_{M_T}(c) + \delta(c)). \quad (2)$$

This modification subtly increases the likelihood of selecting certain tokens based on the prefix,

thereby embedding a detectable pattern in the generated text without significantly degrading its quality. Our goal is to reverse-engineer an approximation $\delta(c)$ to imitate M_T ’s watermark.

3.2 Watermark Inheritance

The initial step of DITTO is to transfer the watermark pattern from the target model M_T to an open-source model whose output logits are accessible. We accomplish this through knowledge distillation, which facilitates the transfer of the target model’s watermark signal to a student model.

We first generate a large text dataset D_T using the watermarked target model M_T . Every text in D_T is imperceptibly embedded with the unique watermark pattern of M_T . We then train the original student model M_O , on this data to obtain the trained student model M_S using a standard supervised fine-tuning (SFT) objective.

$$\theta_S = \arg \max_{\theta_O} \sum_{\mathbf{x} \in D_T} \sum_{i=1}^{|\mathbf{x}|} \log P(x_i|x_{1:i-1}; \theta_O), \quad (3)$$

where θ_S and θ_O are denote the model parameters for M_S and M_O respectively and \mathbf{x} is a text sequence from the training dataset D_T .

As a result, the output distribution of the trained student model, P_{M_S} , becomes a close approxima-

tion of the target model’s watermarked distribution P'_{M_T} for contexts seen during training:

$$P_{M_S}(x_t|c) \approx P'_{M_T}(x_t|c), \quad \forall c \in D_T. \quad (4)$$

Due to the watermark radioactivity, the outputs of M_S exhibit similar watermarking patterns with M_T . This provides the foundation for extracting the watermark signal in the next step.

3.3 Watermark Extraction

The goal of this stage is to extract the watermark pattern embedded within the trained student model M_S as a quantitative signal, which we term the Extracted Watermark Signal (EWS). In order to obtain the signal required for our spoofing attack, we adapt the watermark stealing methodology from Pan et al. (2025) as an effective tool.

The core principle is to isolate the systematic bias induced by the watermark by comparing the behavior of the original model before training (M_O) and the model after training (M_S). First, the context-independent global bias δ_{global} is calculated as the average difference in logits between the two models over the entire training dataset:

$$\delta_{\text{global}} = \mathbb{E}_{c \in D_T} [l_{M_S}(c)] - \mathbb{E}_{c \in D_T} [l_{M_O}(c)]. \quad (5)$$

Similarly, the local bias δ_p for a specific n -gram prefix p is calculated by narrowing the scope to contexts ending with that prefix:

$$\delta_p = \mathbb{E}_{c \in D_T | c \text{ ends with } p} [l_{M_S}(c)] - \mathbb{E}_{c \in D_T | c \text{ ends with } p} [l_{M_O}(c)]. \quad (6)$$

These averaging processes are designed to cancel out context-specific knowledge shifts, leaving only the consistent watermark pattern.

The final EWS is constructed as a weighted sum of these two biases:

$$\text{EWS}(c) = \delta_{\text{global}} + \sum_{p \in \text{prefixes}(c)} w(p) \cdot \delta_p. \quad (7)$$

Prefixes with a higher frequency in the training data are assigned greater weights, reflecting their stronger influence. Following Pan et al. (2025), to ensure prefixes with stronger statistical signals have a greater impact, we assign a weight $w(p)$ to each prefix’s local bias. This weight is determined by the frequency of the prefix p in the training data

(*i.e.*, the output of teacher model) giving greater influence to more common prefixes.

This extracted watermark signal quantitatively replicates the target model’s watermark and serves as the key ingredient for the spoofing attack we propose in the following section.

3.4 Spoofing Attack

Using the watermark signal $\text{EWS}(c)$ extracted in the previous stage, we perform a spoofing attack to replicate the watermark of the target model M_T onto an unwatermarked model.

The attack is executed in real-time during the inference phase of text generation. Just before the unwatermarked original model M_O generates the next token, the corresponding $\text{EWS}(c)$ vector for the current context c is retrieved. The EWS vector is then added to the model’s original logit vector $l_{M_O}(c)$, to yield the final modified logits $l'_{M_O}(c)$:

$$l'_{M_O}(c) = l_{M_O}(c) + \alpha \cdot \text{EWS}(c), \quad (8)$$

where α is a scaling parameter that controls the injection strength. The impact of the scaling parameter α on the performance of the spoofing attack is discussed further in Section 5.

The modified logits guide the model to follow the statistical bias of the target model during the next-token sampling process. Consequently, the attacked M_O generates text embedded with the target model’s unique watermark, thereby successfully completing the spoofing attack.

4 Results of DITTO

4.1 Experimental Setup

In order to validate the effectiveness of DITTO, we design the experimental environment as follows.

Models. we utilize several LLMs that are widely used in recent research: Llama3.1-8B¹, Llama3.2-3B² and Llama3.2-1B³. These three models are paired as teacher and student models for our experiments. The specific pairings are following Table 1 if not specified. The attack model shares the same backbone architecture as the corresponding student model. For example, to imitate the watermark of a Llama3.1-8B teacher, a Llama3.2-3B model first

¹<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

²<https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>

³<https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct>

extracts the watermark signal as the student, and a distinct Llama3.2-3B then serves as the attack model to perform the spoof. We employ Llama2-14B to measure the perplexity of the generated text.

Teacher Model	Student Model
Llama3.1 8B Instruct	Llama3.2 3B Instruct
Llama3.2 3B Instruct	Llama3.2 1B Instruct

Table 1: Teacher-student model pairs used in the knowledge distillation. Each student inherits watermarking signal from its corresponding teacher model.

Datasets. For model training and evaluation, we use two distinct datasets. Dolly-15k (Conover et al., 2023) is an open-source dataset of 15,000 human-generated prompt/response pairs, to fine-tune the models for interactive, instruction-based capabilities. To specifically assess watermarking performance on structured and narrative text, we also use the Bookreport task from the MarkMyWords (MMW) benchmark (Piet et al., 2025), a framework designed for the systematic evaluation of watermarking techniques.

Training and Attack Parameters. For student model training, we do Supervised Fine-Tuning (SFT) with Low-Rank Adaptation (LoRA) and train for a total of 3 epochs. From the total dataset of 10,000 samples, we hold out 100 samples as a test set for evaluation. Watermark injection and detection were based on the algorithm by Kirchenbauer et al. (2023), with parameters set to $\delta = 3$ and $\gamma = 0.5$ for injection, and a z -threshold of 4.0 for detection. For spoofing attack, α , which controls the strength of the spoofing attack, is varied across the range of [2.5, 3, 3.5, 4, 4.5, 5]. Unless otherwise specified, all main results reported in our tables use a fixed α of 4.5.

4.2 Evaluation Metrics

TPR@FPR. The True Positive Rate (TPR) measures how often a watermark is correctly found at a fixed False Positive Rate (FPR). To create a strict test, we set the FPR to low levels (10%, 1%, and 0.1%), then measure the resulting TPR for our spoofed texts.

p-value. The statistical strength of a detected watermark is measured by its p-value, where a lower value indicates stronger evidence. A strict

threshold is required to prevent false alarms when analyzing long texts. We report the median p-value to show the overall strength of the spoofed signal.

4.3 Experimental Results

The performance of our proposed methodology is evaluated against several baselines. First, we establish the original watermarked model as a baseline to define the upper bound for the watermark detection rate. Additionally, we include prior works on spoofing attacks, JSV (Jovanović et al., 2024) and De-Mark (Chen et al., 2025), as primary comparison targets. Notably, De-Mark considers two settings based on the attacker’s level of access. The graybox setting assumes an ideal scenario where the attacker has direct access to the model’s top-k token probabilities to infer the watermarking rule. In contrast, the blackbox setting represents a more realistic and challenging attack environment where, lacking probability information, the attacker must indirectly deduce the rule from the frequency of multiple generation outputs. Since our methodology operates under a blackbox assumption, a fair comparison is made against De-Mark (blackbox).

The results of our spoofing attack on Llama3.1-8B and Llama3.2-3B are presented in Table 2, alongside the baselines. The results demonstrate that DITTO can successfully deceive watermark detectors by injecting imitated watermark signal into the student model’s output. Notably, our methodology proves to be more effective when the backbone model size is smaller. We provide a detailed analysis of these results in the subsequent sections.

5 Analysis

Section 4 demonstrates that DITTO can successfully mimic the target model’s watermark signal. This motivates a deeper exploration into the conditions that maximize effectiveness, as well as an investigation of its limitations and potential for expansion.

5.1 Impact of Spoofing Strength

A successful spoofing attack must not only achieve a high detection rate but also maintain *stealthiness* by preserving the original text quality to avoid suspicion. The key variable that balances these two objectives is the scaling parameter α (Eq. 8). Here, we conduct an in-depth analysis of how varying α affects the spoofing strength and the quality of the generated text to determine the efficiency and limitations of our proposed methodology. We experiment

	MMW Bookreport					Dolly CW				
	TPR@FPR=			median p-value ↓	ppl ↓	TPR@FPR=			median p-value ↓	ppl ↓
	10%↑	1%↑	0.1%↑			10%↑	1%↑	0.1%↑		
Llama3.1-8b Watermarked	0.89	0.75	0.65	1.10E-06	3.87	0.92	0.82	0.70	1.88E-07	6.34
+ JSV	0.36	0.10	0.04	2.20E-01	2.81	0.56	0.26	0.11	6.00E-02	6.62
+ De-Mark (gray)	0.99	0.96	0.88	5.73E-06	3.77	0.96	0.89	0.77	4.32E-05	3.97
+ De-Mark (black)	1.00	0.94	0.73	6.93E-05	3.55	0.97	0.82	0.70	2.66E-05	3.93
+ DITTO (Ours)	0.81	0.70	0.51	7.97E-04	4.18	0.96	0.93	0.85	1.78E-07	3.73
Llama3.2-3b Watermarked	0.91	0.79	0.74	3.43E-07	4.28	0.82	0.68	0.61	2.64E-06	6.34
+ JSV	0.70	0.34	0.18	4.00E-02	3.78	0.73	0.42	0.28	1.30E-02	6.93
+ De-Mark (gray)	0.98	0.96	0.90	1.93E-06	4.65	0.99	0.93	0.75	2.66E-05	4.64
+ De-Mark (black)	1.00	0.96	0.90	3.34E-06	4.68	1.00	0.90	0.80	2.66E-05	4.74
+ DITTO (Ours)	0.99	0.99	0.97	5.48E-17	2.44	0.91	0.86	0.75	1.78E-07	2.95

Table 2: Spoofing-attack evaluation results on MMW Book Report and Dolly Creative Writing datasets. We report detection performance (TPR at fixed FPRs), median p-values and perplexity (ppl) for different watermarking defenses. Arrows indicate the preferred direction for each metric (↑: higher is better, ↓: lower is better).

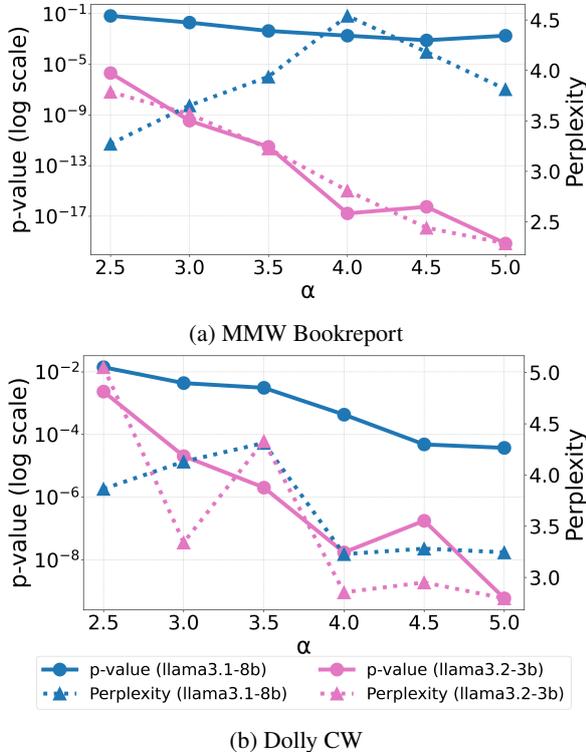


Figure 3: The effect of varying the α parameter on the p-value and Perplexity. The complete experimental results for this analysis are available in Table 7 in Appendix A.

with the same scenarios as Section 4, imitating Llama3.1-8B with Llama3.2-3B and Llama3.2-3B with Llama3.2-1B, increasing α from 2.5 to 5.0 in increments of 0.5.

We initially hypothesize that a larger α would lead to a higher spoofing success rate at the cost of degraded text quality (e.g., higher perplexity), creating a clear trade-off. However, the experimental results presented in Figure 3 reveal a surprising

trend that deviates from our expectation. As expected, across both the MMW Bookreport (3a) and Dolly CW (3b), the median p-value exhibits a decrease as α grows. However, the perplexity defies our prediction. Instead of a steady increase, the text quality metric exhibits a non-monotonic pattern; while it rises in some intervals, it frequently decreases in others as the attack intensifies. This volatility is particularly pronounced when imitating the Llama3.2-3b model, a result that challenges the common assumption that stronger attacks must necessarily degrade text quality.

This finding reveals a critical characteristic of the DITTO attack: an adversary can modulate the spoofing strength without producing a signal of quality degradation. Detection methods that operate on the heuristic that attacks degrade text quality are likely to fail, as our results show that perplexity fluctuates unpredictably rather than consistently increasing with spoofing strength. This finding, therefore, implies that DITTO possesses a greater potential for stealth and evasion than a simple trade-off relationship would suggest.

5.2 Beyond Green-Lists

While the previous sections demonstrate spoofing capabilities of DITTO against KGW, an n-gram-based green-list watermark, a critical question remains:

Does our attack mechanism merely exploit scheme-specific features of watermark, or does it learn the more fundamental patterns of distributional distortion caused by any watermark?

Answering this is crucial for evaluating the generalizability of our methodology. In order to investi-

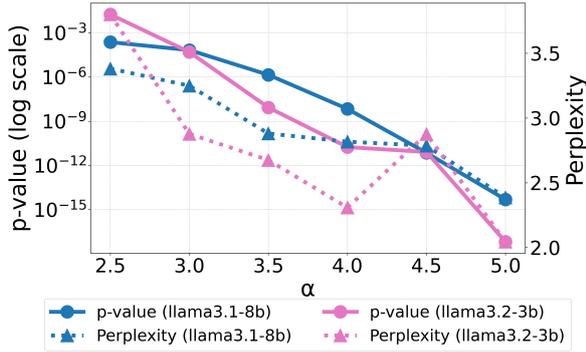


Figure 4: Impact of the scaling parameter α on spoofing performance against SynthID on the Dolly CW dataset.

gate this, we test DITTO against SynthID (Dathathri et al., 2024), a prominent sampling-based watermark that operates differently from that of Kirchenbauer et al. (2023). The experiment is conducted on the Dolly CW dataset, on which our main experiments show the most stable performance.

	TPR@FPR=			median p-value ↓	ppl ↓
	10%↑	1%↑	0.1%↑		
Llama3.1-8b (victim)					
Watermarked	0.98	0.98	0.96	1.57E-12	2.80
DITTO	0.88	0.84	0.80	7.10E-10	2.79
Llama3.2-3b (victim)					
Watermarked	0.97	0.97	0.94	5.51E-15	6.30
DITTO	0.90	0.89	0.87	8.12E-12	2.88

Table 3: Spoofing attack results against the SynthID watermark on the Dolly CW dataset. DITTO achieves high TPR and low p-values, comparable to the authentically watermarked models.

Table 3 clearly demonstrates that DITTO achieves highly effective spoofing performance against SynthID as well. The text generated by DITTO records high TPRs and low p-values, comparable to the original watermarked model, successfully deceiving a detector into recognizing it as authentic watermarked text. The perplexity results are particularly noteworthy. For Llama3.1-8b, DITTO’s perplexity is nearly identical to the original’s. For Llama3.2-3b, our perplexity is substantially lower than that of the original watermarked model. This suggests that DITTO not only mimics the watermark but also exhibits an unexpected efficiency, mitigating the quality degradation that can occur during the original watermarking process and generating more natural text.

Furthermore, consistent with our findings from the previous experiments (Figure 3), we again an-

size of D_T	TPR@FPR=				ppl ↓
	10%↑	5%↑	1%↑	0.1%↑	
10k (DITTO)	0.81	0.76	0.70	0.51	4.18
8k	0.72	0.70	0.65	0.55	2.80
6k	0.71	0.65	0.58	0.44	2.91
4k	0.68	0.64	0.49	0.41	2.63

Table 4: Impact of D_T size on spoofing performance for the Dolly CW dataset.

alyze the trade-off between attack strength (*i.e.*, α) and text quality (*i.e.*, perplexity) for spoofing sampling-based watermark. Figure 4 shows that as α increases, the p-value consistently decreases while the perplexity fluctuates without a consistent upward trend. Notably, and unlike our findings with KGW, the Llama3.1-8B model exhibits a more stable perplexity pattern in the SynthID experiment. This finding reinforces a key conclusion from Section 5.1: our attack’s stealthiness appears to be an inherent trait, persisting even against SynthID and strengthening our claim that its intensity can be increased without a clear cost to text quality.

Our analysis culminates in two key findings that provide compelling evidence for generalizability:

- 1. Versatility across Schemes:** The attack is not confined to a single watermarking architecture, demonstrating high effectiveness against both green-list (KGW) and fundamentally different sampling-based (SynthID) methods.
- 2. Stealth without Compromise:** The attack breaks the conventional trade-off between strength and quality. Its intensity can be increased without producing a consistent, discernible degradation in text quality, making it highly evasive.

We conclude that DITTO is not confined to a specific watermarking algorithm but possesses a universal ability to learn and imitate artificial statistical bias.

5.3 Potential of Synthetic Data

The potency of DITTO is intrinsically linked to its utilization of synthetic data, a resource that is theoretically inexhaustible as it can be generated directly from the target teacher model. As demonstrated in Table 4, the effectiveness of the spoofing attack generally follows a scaling law, where the TPR across most FPR levels exhibits an overall upward trend as the size of the training dataset (D_T) increases. This positive correlation indicates that the

	MMW Bookreport					Dolly CW				
	TPR@FPR=			median p-value ↓	ppl ↓	TPR@FPR=			median p-value ↓	ppl ↓
	10%↑	1%↑	0.1%↑			10%↑	1%↑	0.1%↑		
KGW										
Llama3.2-3b (<i>non-wm</i>)	0.34	0.09	0.01	2.40E-01	3.00	0.33	0.14	0.09	3.20E-01	3.72
+ <i>distill-only</i>	0.57	0.34	0.19	6.08E-02	2.89	0.79	0.65	0.55	2.83E-04	4.61
+ DITTO	0.81	0.70	0.51	8.00E-04	4.18	0.84	0.76	0.64	4.83E-05	3.28
Llama3.2-1b (<i>non-wm</i>)	0.30	0.08	0.03	2.83E-01	3.29	0.35	0.11	0.07	3.24E-01	8.39
+ <i>distill-only</i>	0.96	0.93	0.90	2.60E-09	2.65	0.56	0.44	0.41	4.65E-02	3.91
+ DITTO	0.99	0.99	0.97	5.48E-17	2.44	0.91	0.86	0.75	1.78E-07	2.95
SynthID										
Llama3.2-3b (<i>non-wm</i>)	0.24	0.09	0.02	4.04E-01	3.00	0.14	0.06	0.02	6.11E-01	3.72
+ <i>distill-only</i>	0.39	0.11	0.03	2.07E-01	3.01	0.62	0.43	0.36	2.26E-02	5.66
+ DITTO	0.32	0.13	0.08	2.62E-01	3.81	0.88	0.84	0.80	7.10E-12	2.79
Llama3.2-1b (<i>non-wm</i>)	0.21	0.07	0.03	4.29E-01	3.29	0.26	0.15	0.10	4.82E-01	8.39
+ <i>distill-only</i>	0.37	0.18	0.10	2.03E-01	5.56	0.50	0.38	0.32	9.68E-02	3.99
+ DITTO	0.21	0.06	0.01	3.63E-01	3.92	0.90	0.87	0.81	8.12E-12	2.88

Table 5: Comparison of spoofing effectiveness between DITTO and the passive distillation (*distill-only*).

extracted watermark signal becomes increasingly authoritative as the student model is exposed to more watermarked outputs.

5.4 Why Distillation is Not Enough

The results in Table 5 underscore the critical necessity of a framework that explicitly isolates and re-applies watermark signals, rather than relying on passive inheritance. We compare DITTO against a *distill-only* baseline, which refers to models trained using only SFT without watermark extraction (Sec. 3.3). While standard distillation allows a student model to exhibit radioactive traces of a teacher’s watermark, this approach is insufficient for an effective spoofing attack. As depicted in Table 5, DITTO consistently outperforms the *distill-only* approach across nearly all metrics and model configurations. These results demonstrate that relying solely on distillation fails to capture the full intensity of the target watermark. Therefore, the specialized DITTO framework is essential for achieving a high-fidelity spoofing attack that simple supervised fine-tuning cannot replicate.

5.5 Other Victims

While previous experiments focused on scenarios where the victim and attacker sharing the same model architecture, it is crucial to evaluate whether DITTO remains effective against victims from heterogeneous model families. To assess this cross-family generalizability, we conducted experiments using GPT-OSS 20B and Phi-4 14B as target victims,

	TPR@FPR=			median p-value ↓
	10%↑	1%↑	0.1%↑	
3b (<i>non-wm</i>)	0.34	0.09	0.01	2.40E-01
DITTO (3b) → GPT	0.60	0.50	0.42	1.18E-02
DITTO (3b) → Phi	0.67	0.55	0.35	7.06E-03
1b (<i>non-wm</i>)	0.30	0.08	0.03	2.83E-01
DITTO (1b) → GPT	0.50	0.43	0.35	9.77E-02
DITTO (1b) → Phi	0.65	0.52	0.42	5.45E-03

Table 6: Cross-family generalizability of the framework on the Dolly CW dataset. Arrows (→) indicate the direction from the attacker model to the victim model.

while employing Llama 3.2 3B and 1B as the attack models. As presented in Table 6, DITTO successfully replicates the target watermark even when the attacker and victim belong to entirely different model families. The results consistently show significantly higher watermark detection rates and lower median p -values than the non-watermarked baselines, demonstrating that DITTO is not constrained by architectural similarity.

6 Conclusion

DITTO repurposes watermark radioactivity to spoof LLM authorship across diverse schemes without sacrificing text quality. Our findings prove that watermark presence is no longer a reliable proof of origin, necessitating a paradigm shift from mere presence detection toward robust authenticity verification and cryptographic model-binding.

Limitations

While this study provides a new perspective on LLM watermark security, it has several limitations. First, the efficacy of DITTO is constrained by the fidelity of watermark inheritance, which we found to be sensitive to model size and architectures. As our primary focus was on demonstrating the attack’s feasibility, we did not perform an exhaustive optimization of these variables. This leaves a significant opening for future research to systematically map the boundary conditions of watermark radioactivity. Second, this work is inherently adversarial and does not prescribe specific defense mechanisms. By intentionally focusing on the vulnerability side, we aim to provide a necessary baseline for the community. Investigating how to distort or neutralize these leaked signals during the distillation process remains a crucial next step for developing more resilient content protection systems.

Ethical Considerations

The DITTO framework, by design, demonstrates a potent method for watermark spoofing that could be exploited for malicious purposes, such as spreading disinformation and falsely attributing it to trusted sources. However, our primary objective is not to provide a tool for malicious actors, but to proactively identify and expose a critical vulnerability in the current paradigm of LLM watermarking. By revealing the fragility of the core assumption that a detected watermark guarantees authorship, our work serves as a crucial red-teaming exercise for the AI safety community. We argue that the responsible disclosure of this vulnerability is essential to next-generation watermarking techniques where adversarial resilience is treated as a fundamental design. Ultimately, our findings call for a paradigm shift from mere presence detection to developing technologies capable of distinguishing authentic watermarks from expertly imitated ones, thereby ensuring genuine authorship and fostering more secure AI ecosystems.

Acknowledgments

This research was supported by the NRF grant (RS-2025-00562134) and the AI Graduate School Program (RS-2020-II201361) funded by the Korean government.

References

- Joseph R Biden. 2023. Executive order on the safe, secure, and trustworthy development and use of artificial intelligence.
- Shizhan Cai, Liang Ding, and Dacheng Tao. 2025. Entropy-guided watermarking for llms: A test-time framework for robust and traceable text generation. *arXiv preprint arXiv:2504.12108*.
- Liang Chen, Yatao Bian, Yang Deng, Deng Cai, Shuaiyi Li, Peilin Zhao, and Kam-Fai Wong. 2024. WatME: Towards lossless watermarking through lexical redundancy. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ruibo Chen, Yihan Wu, Junfeng Guo, and Heng Huang. 2025. De-mark: Watermark removal in large language models. In *Proceedings of International Conference on Machine Learning (ICML)*.
- Yixin Cheng, Hongcheng Guo, Yangming Li, and Leonid Sigal. 2025. Revealing weaknesses in text watermarking through self-information rewrite attacks. In *Forty-second International Conference on Machine Learning (ICML)*.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. Free dolly: Introducing the world’s first truly open instruction-tuned llm.
- Sumanth Dathathri, Abigail See, Sumedh Ghaisas, Po-Sen Huang, Rob McAdam, Johannes Welbl, Vandana Bachani, Alex Kaskasoli, Robert Stanforth, Tatiana Matejovicova, and 1 others. 2024. Scalable watermarking for identifying large language model outputs. *Nature*.
- EUAIAct. 2024. [Eu ai act compliance checker](#).
- Thibaud Gloaguen, Robin Staab, Nikola Jovanović, and Martin Vechev. 2025. Watermarking diffusion language models. *arXiv preprint arXiv:2509.24368*.
- Wooyoung Go, Hyoungshick Kim, Alice Oh, and Yongdae Kim. 2025. XDAC: XAI-driven detection and attribution of LLM-generated news comments in Korean. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Yuxuan Guo, Zhiliang Tian, Yiping Song, Tianlun Liu, Liang Ding, and Dongsheng Li. 2024. Context-aware watermark with semantic balanced green-red lists for large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Baizhou Huang, Xiao Pu, and Xiaojun Wan. 2025. b^4 : A black-box scrubbing attack on LLM watermarks. In *Proceedings of the 2025 Conference of the Nations of*

- the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Albuquerque, New Mexico. Association for Computational Linguistics.
- Nikola Jovanović, Robin Staab, and Martin Vechev. 2024. Watermark stealing in large language models. In *Proceedings of International Conference on Machine Learning (ICML)*.
- Jungin Kim, Shinwoo Park, and Yo-Sub Han. 2026. Marking code without breaking it: Code watermarking for detecting llm-generated code. In *Findings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. To appear.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A watermark for large language models. In *International Conference on Machine Learning (ICML)*.
- Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. 2024. Robust distortion-free watermarks for language models. *Transactions on Machine Learning Research*.
- Taehyun Lee, Seokhee Hong, Jaewoo Ahn, Ilgee Hong, Hwaran Lee, Sangdoon Yun, Jamin Shin, and Gunhee Kim. 2024. Who wrote this code? watermarking for code generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Yijian Lu, Aiwei Liu, Dianshi Yu, Jingjing Li, and Irwin King. 2024. An entropy-based text watermarking detection method. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Niloofar Mireshghallah, Justus Mattern, Sicun Gao, Reza Shokri, and Taylor Berg-Kirkpatrick. 2024. Smaller Language Models are Better Zero-shot Machine-Generated Text Detectors. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Leyi Pan, Aiwei Liu, Shiyu Huang, Yijian Lu, Xuming Hu, Lijie Wen, Irwin King, and Philip S. Yu. 2025. Can LLM watermarks robustly prevent unauthorized knowledge distillation? In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Qi Pang, Shengyuan Hu, Wenting Zheng, and Virginia Smith. 2024. No free lunch in llm watermarking: Trade-offs in watermarking design choices. *Advances in Neural Information Processing Systems*, 37:138756–138788.
- Shinwoo Park, Shubin Kim, Do-Kyung Kim, and Yo-Sub Han. 2025. KatFishNet: Detecting LLM-generated Korean text through linguistic feature analysis. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Shinwoo Park, Hyejin Park, Hyeseon Ahn, and Yo-Sub Han. 2026. Watermod: Modular token-rank partitioning for probability-balanced llm watermarking. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. To appear.
- Mike Perkins. 2023. Academic Integrity considerations of AI Large Language Models in the post-pandemic era: ChatGPT and beyond. *Journal of University Teaching and Learning Practice*.
- Julien Piet, Chawin Sitawarin, Vivian Fang, Norman Mu, and David Wagner. 2025. Markmywords: Analyzing and evaluating language model watermarks. In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 68–91. IEEE.
- Saksham Rastogi and Danish Pruthi. 2024. Revisiting the robustness of watermarking to paraphrasing attacks. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Reuters. 2023. [Openai, google, others pledge to watermark ai content for safety, white house says](#). Accessed: 2025-10-03.
- Tom Sander, Pierre Fernandez, Alain Durmus, Matthijs Douze, and Teddy Furon. 2024. Watermarking makes language models radioactive. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Tom Sander, Pierre Fernandez, Saeed Mahloujifar, Alain Durmus, and Chuan Guo. 2025. Detecting benchmark contamination through watermarking. *arXiv preprint arXiv:2502.17259*.
- Lionel Z Wang, Yiming Ma, Renfei Gao, Beichen Guo, Zhuoran Li, Han Zhu, Wenqi Fan, Zexin Lu, and Ka Chung Ng. 2024. MegaFake: A Theory-Driven Dataset of Fake News Generated by Large Language Models. *arXiv preprint arXiv:2408.11871*.
- Qilong Wu and Varun Chandrasekaran. 2024a. Bypassing LLM watermarks with color-aware substitutions. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Qilong Wu and Varun Chandrasekaran. 2024b. Bypassing llm watermarks with color-aware substitutions. *arXiv preprint arXiv:2403.14719*.
- Changrong Xiao, Wenxing Ma, Qingping Song, Sean Xin Xu, Kunpeng Zhang, Yufang Wang, and Qi Fu. 2024. Human-AI Collaborative Essay Scoring: A Dual-Process Framework with LLMs. *arXiv preprint arXiv:2401.06431*.
- KiYoon Yoo, Wonhyuk Ahn, Jiho Jang, and Nojun Kwak. 2023. Robust multi-bit natural language watermarking through invariant features. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*.

Zhaoxi Zhang, Xiaomei Zhang, Yanjun Zhang, Leo Yu Zhang, Chao Chen, Shengshan Hu, Asif Gill, and Shirui Pan. 2024. Large language model watermark stealing with mixed integer programming. *arXiv preprint arXiv:2405.19677*.

A Detailed Results of the Spoofing Strength Analysis

We provide the complete numerical data for the spoofing strength scaling experiment, which is discussed in Section 5.1 and visualized in Figure 3 of the main text.

Table 7 details the median p-value and perplexity for each increment of the scaling parameter α from 2.5 to 5.0. The results are presented for each experimental configuration, covering the two imitation scenarios: 1) Llama3.2-3B imitates the watermark of Llama3.1-8B and 2) Llama3.2-1B imitates the watermark of Llama3.2-3B on the two datasets, MMW Bookreport and Dolly CW.

Consistent with the analysis in the main body, the data clearly shows a decrease in p-value as α increases, indicating a stronger statistical significance for the spoofing attack. Concurrently, the perplexity values fluctuate without a clear upward trend, providing the specific numerical evidence that supports our conclusion regarding the lack of a severe trade-off between attack strength and text quality.

B Algorithms of DITTO

We present the overall architecture and procedural details of DITTO through a series of formal algorithms. The DITTO framework is composed of three core modules: (1) Watermark Inheritance, (2) Watermark Extraction and (3) the Spoofing Attack. Algorithm 1 describes the inheritance process, where a student model learns the teacher’s watermark patterns via knowledge distillation. Next, Algorithm 2 details the extraction process, where the inherited watermark is isolated as a quantitative signal (EWS). Finally, Algorithm 3 shows the final attack stage, where this extracted signal is injected into an attack model to perform a real-time spoof of the target watermark.

C Dataset Examples

Examples from the datasets utilized in our experiments are presented below. The samples are drawn from the MMW (MarkMyWords) Bookreport benchmark and the Dolly-15k dataset, which served as our evaluation task.

The examples from Dolly-15k showcase the diversity of tasks our models were trained on, ranging from open-domain question answering and brainstorming to information extraction from a given context. The MMW Bookreport sample, in contrast, exemplifies the long-form generation task used in our evaluation.

Example of MMW Bookreport

```
"task_type": "book_report",
"system_message": "You are a helpful assistant. Always answer in the most accurate way.",
"user_message": "Write a book report about 'Pride and Prejudice', written by Jane Austen."
```

Example of Dolly CW 1

```
"instruction": "Who is called Napoleon of Iran?",
"context": "",
"response": "Nadir Shah",
"category": "open_qa"
```

Example of Dolly CW 2

```
"instruction": "Why capri is so famous",
"context": "",
"response": "Capri is famous Italian island for it's nice ocean color and it's mountain view",
"category": "brainstorming"
```

Example of Dolly CW 3

```
"instruction": "What is Github Copilot?",
"context": "GitHub Copilot is a cloud-based artificial intelligence tool developed by GitHub and OpenAI to assist users of Visual Studio Code, Visual Studio, Neovim, and JetBrains integrated development environments (IDEs) by autocompleting code. Currently available by subscription to individual developers, the tool was first announced by GitHub on 29 June 2021, and works best for users coding in Python, JavaScript, TypeScript, Ruby, and Go.",
"response": "GitHub Copilot is a cloud-based artificial intelligence tool developed by GitHub and OpenAI",
"category": "closed_qa"
```

D Examples of Spoofed Outputs

We conduct a spoofing attack where a smaller model, Llama 3.2-3B, is trained to imitate the watermark

	MMW Bookreport					Dolly CW				
	TPR@FPR=			median p-value ↓	perplexity ↓	TPR@FPR=			median p-value ↓	perplexity ↓
	10%↑	1%↑	0.1%↑			10%↑	1%↑	0.1%↑		
Llama3.1-8b (victim) mimicked by Llama3.2-3b (attacker)										
$\alpha = 2.5$	0.62	0.26	0.08	6.83E-02	3.27	0.71	0.48	0.25	1.40E-02	3.86
$\alpha = 3$	0.76	0.46	0.22	1.99E-02	3.65	0.79	0.57	0.45	4.36E-03	4.13
$\alpha = 3.5$	0.79	0.56	0.38	4.36E-03	3.94	0.82	0.54	0.47	3.09E-03	4.31
$\alpha = 4$	0.87	0.65	0.46	1.83E-03	4.54	0.82	0.69	0.55	4.32E-04	3.23
$\alpha = 4.5$	0.81	0.70	0.51	7.97E-04	4.18	0.84	0.76	0.64	4.83E-05	3.28
$\alpha = 5$	0.73	0.61	0.49	1.83E-03	3.81	0.91	0.76	0.72	3.75E-05	3.25
Llama3.2-3b (victim) mimicked by Llama3.2-1b (attacker)										
$\alpha = 2.5$	0.98	0.94	0.84	2.04E-06	3.79	0.72	0.63	0.46	2.34E-03	5.05
$\alpha = 3$	0.99	0.97	0.95	3.47E-10	3.57	0.83	0.74	0.64	2.05E-05	3.34
$\alpha = 3.5$	0.97	0.95	0.93	3.07E-12	3.23	0.92	0.82	0.68	2.04E-06	4.33
$\alpha = 4$	0.96	0.95	0.92	1.65E-17	2.81	0.92	0.89	0.79	1.74E-08	2.86
$\alpha = 4.5$	0.99	0.99	0.97	5.48E-17	2.44	0.91	0.86	0.75	1.78E-07	2.95
$\alpha = 5$	0.94	0.91	0.89	7.00E-20	2.29	0.92	0.87	0.81	5.95E-10	2.80

Table 7: Full results for the spoofing strength analysis, corresponding to the trends visualized in Figure 3. We report for two imitation scenarios across both datasets.

Algorithm 1 Watermark Inheritance via Knowledge Distillation

- 1: **Input:** Watermarked teacher model M_T , original student model M_O , watermark strength δ
 - 2: **Output:** Fine-tuned student model M_S with inherited watermark
 - Phase 1: Generate KGW-Watermarked Dataset
 - 3: $D_T \leftarrow \emptyset$
 - 4: **for** each prompt p in a predefined set of prompts **do**
 - 5: $x \leftarrow p$ ▸ Initialize the sequence with a prompt
 - 6: **while** not ENDOFSEQUENCE(x) **do**
 - 7: $l \leftarrow M_T(x_{1:t-1})$ ▸ Get original logits from the teacher
 - 8: $V_{\text{green}} \leftarrow \text{GETGREENLIST}(x_{t-1})$ ▸ Determine green list based on the previous token
 - 9: **for** each token $v \in V_{\text{green}}$ **do**
 - 10: $l_v \leftarrow l_v + \delta$ ▸ Add bias δ to logits of green tokens
 - 11: **end for**
 - 12: $x_t \sim \text{softmax}(l)$ ▸ Sample next token from the modified distribution
 - 13: $x \leftarrow x \oplus x_t$ ▸ Append the new token to the sequence
 - 14: **end while**
 - 15: $D_T \leftarrow D_T \cup \{x\}$
 - 16: **end for**
 - Phase 2: Supervised Fine-Tuning of Student Model
 - 17: Let θ_O be the initial parameters of M_O
 - 18: $\theta_S \leftarrow \arg \max_{\theta_O} \sum_{x \in D_T} \sum_{i=1}^{|x|} \log P(x_i | x_{1:i-1}; \theta_O)$
 - 19: $M_S \leftarrow M_O$ with parameters updated to θ_S
 - 20: **return** M_S
-

of a larger model, Llama 3.1-8B, on the MMW Bookreport dataset.

about ‘Great Expectations’, written by Charles Dickens.

You are a helpful assistant. Always answer in the most accurate way. Write a book report

Algorithm 2 Watermark Extraction

```
1: Input: Trained student model  $M_S$ , original student model  $M_O$ , training dataset  $D_T$ 
2: Output: Extracted Watermark Signal (EWS) function
3:                                     ▶ Calculate the context-independent global bias
4:  $\delta_{\text{global}} \leftarrow \mathbb{E}_{c \in D_T} [l_{M_S}(c)] - \mathbb{E}_{c \in D_T} [l_{M_O}(c)]$ 
5:                                     ▶ Calculate prefix-specific local biases for all n-gram prefixes  $p$  in  $D_T$ 
6: for each unique prefix  $p$  in  $D_T$  do
7:    $D_{T|p} \leftarrow \{c \in D_T \mid c \text{ ends with } p\}$ 
8:    $\delta_p \leftarrow \mathbb{E}_{c \in D_{T|p}} [l_{M_S}(c)] - \mathbb{E}_{c \in D_{T|p}} [l_{M_O}(c)]$ 
9: end for
10:                                     ▶ Define the final EWS as a weighted sum of the biases
11: function EWS( $c$ )
12:   Let  $\text{prefixes}(c)$  be the set of n-gram prefixes in context  $c$ 
13:    $\text{EWS}(c) \leftarrow \delta_{\text{global}} + \sum_{p \in \text{prefixes}(c)} w(p) \cdot \delta_p$ 
14:                                     ▶  $w(p)$  is a weight based on prefix frequency
15:   return EWS( $c$ )
16: end function
17: return EWS function
```

Algorithm 3 Spoofing Attack

```
1: Input: Attacker model  $M_O$ , EWS function, scaling parameter  $\alpha$ , initial context  $c_0$ 
2: Output: Text sequence  $T$  with spoofed watermark

3:  $T \leftarrow c_0$                                      ▶ Initialize the generated text with a starting prompt
4:  $c \leftarrow c_0$                                      ▶ Initialize the current context
5: while not END_OF_SEQUENCE() do
6:    $x_{\text{next}} \leftarrow \text{GENERATE\_TOKEN}(M_O, \text{EWS}, \alpha, c)$            ▶ Generate the next token
7:    $T \leftarrow T \oplus x_{\text{next}}$                                        ▶ Append the new token to the text sequence
8:    $c \leftarrow T$                                                    ▶ Update the context for the next step
9: end while
10: return  $T$ 

11: procedure GENERATE_TOKEN( $M_O$ , EWS,  $\alpha$ ,  $c$ )
12:    $l_{\text{original}} \leftarrow M_O(c)$                                      ▶ Get original logits from the attacker's model
13:    $s_{\text{watermark}} \leftarrow \text{EWS}(c)$                                ▶ Retrieve the extracted watermark signal
14:    $l_{\text{modified}} \leftarrow l_{\text{original}} + \alpha \cdot s_{\text{watermark}}$        ▶ Inject the watermark signal
15:    $P_{\text{modified}} \leftarrow \text{softmax}(l_{\text{modified}})$                    ▶ Compute the new probability distribution
16:    $x_{\text{sampled}} \leftarrow \text{SAMPLE}(P_{\text{modified}})$                  ▶ Sample from the modified distribution
17:   return  $x_{\text{sampled}}$ 
18: end procedure
```

Model answer

The Great Expectations is a classic novel written by Charles Dickens, published in 1861. The story follows the journey of Pip, an orphan boy who is raised by his abusive sister, and her friend, an old man who is imprisoned for 20 years for murder. The story is an adaptation of the classic fairy tale 'The Adventures of

the Millionaire and the Genius, The King and the Genius and The Great and the Good'. The story is told through the eyes of the protagonist, Pip, and is written in the first-person narrative. The narrative is written in the past tense and is divided into three sections: the childhood, the present, and the future. The story is known for its themes of love, redemption, and the struggle between the old and the new. The characters

of Miss Havisham, Abel Magwitch, and Miss Havisham's sister, Emily, are crucial to the story and are known for their unique personalities and motivations. The author uses vivid imagery

We conduct a spoofing attack where a smaller model, Llama 3.2-3B, is trained to imitate the watermark of a larger model, Llama 3.1-8B, on the Dolly CW dataset.

Question: Classify each as a chip or a cracker: Goldfish, Doritos, Cheez-It, Cheetos.
Context: None

Model answer

- Goldfish: Chip
- Doritos: Chip
- Cheez-It: Cracker
- Cheetos: Cracker

The correct classification is based on the texture and the primary ingredients of the snack. The classification can be based on the primary ingredients and the texture.

We conduct a spoofing attack where a smaller model, Llama 3.2-1b, is trained to imitate the watermark of a larger model, Llama-8b, on the MMW Bookreport dataset.

You are a helpful assistant. Always answer in the most accurate way. Write a book report about 'Moby Dick', written by Herman Melville.

Model answer

"Moby Dick" is an epic novel written by Herman Melville, first published in 1851. The story revolves around the protagonist, Captain Ahab, a vengeful and obsessive whaling captain who is determined to hunt down and kill a white sperm whale named Moby Dick. The novel is an allegory for the human condition, exploring themes such as revenge, obsession, and the search for meaning in life. The story is set in the mid-19th century and is based on Melville's own experiences as a whaler. The plot of the novel follows the crew of the whaleship Pequod, which is captained by the enigmatic and experienced Ishmael. The crew is made up of various characters, each with their own distinct personality and backstory. There is Queequeg,

a skilled harpooneer from a South Pacific island, who becomes Ishmael's close friend and confidant. The crew also includes Stubb, a young

We conduct a spoofing attack where a smaller model, Llama 3.2-1b, is trained to mimic the watermark of a larger model, Llama-8b, on the Dolly CW dataset.

Question: Give me some ideas on where to travel for good sun and beaches. Context: None

Model answer

The Caribbean, Hawaii, and Southeast Asia are popular destinations known for their beautiful beaches and tropical sun. The Maldives and Seychelles are also popular for their clear waters and white sand beaches. The Mediterranean region also has many beautiful beaches and is known for its pleasant weather. The US, Mexico, and the Caribbean are also popular for their sun and beaches. The best time to visit these destinations is during the summer months when the sun is at its strongest. The best time to visit the US, Mexico, and the Caribbean is from May to October when the temperatures are warm and the risk of hurricanes is low. The best time to visit the Caribbean and Southeast Asia is from late fall to early spring when the temperatures are cooler and the risk of hurricanes is low. The best time to visit the Mediterranean region is from late spring to early fall when the temperatures are warm and the risk of drought is high. The best time to visit the Maldives and Seychelles is from late fall to early spring