CATER: Intellectual Property Protection on Text Generation APIs via Conditional Watermarks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Previous works have validated that text generation APIs can be stolen through imitation attacks, causing IP violations. In order to protect the IP of text generation APIs, a recent work has introduced a watermarking algorithm and utilized the null-hypothesis test as a post-hoc ownership verification on the imitation models. However, we find that it is possible to detect those watermarks via sufficient statistics of the frequencies of candidate watermarking words. To address this drawback, in this paper, we propose a novel Conditional wATERmarking framework (CATER) for protecting the IP of text generation APIs. An optimization method is proposed to decide the watermarking rules that can minimize the distortion of overall word distributions while maximizing the change of conditional word selections. Theoretically, we prove that it is infeasible for even the savviest attacker (they know how CATER works) to reveal the used watermarks from a large pool of potential word pairs based on statistical inspection. Empirically, we observe that high-order conditions lead to an exponential growth of suspicious (unused) watermarks, making our crafted watermarks more stealthy. In addition, CATER can effectively identify the IP infringement under architectural mismatch and cross-domain imitation attacks, with negligible impairments on the generation quality of victim APIs. We envision our work as a milestone for stealthily protecting the IP of text generation APIs.

1 Introduction

Nowadays, many technology corporations, such as Google, Amazon, Microsoft, have invested a plethora of workforce and computation to data collection and model training, in order to deploy well-trained commercial models as pay-as-you-use services on their cloud platforms. Therefore, these corporations own the intellectual property (IP) of their trained models. Unfortunately, previous works have validated that the functionality of a victim API can be stolen through imitation attacks, which inquire the victim with carefully designed queries and train an imitation model based on the outputs of the target API. Such attacks cause severe IP violations of the target API and stifle the creativity and motivation of our research community [36, 40, 15, 9].

In fact, imitation attacks work not only on laboratory models, but also on commercial APIs [40, 41], since the enormous commercial benefit allures competing companies or individual users to extract or steal these successful APIs. For instance, some leading companies in NLP business have been caught imitating their competitors' models [33]. Beyond imitation attacks, the attacker could potentially surpass victims by conducting unsupervised domain adaptation and multi-victim ensemble [41].

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

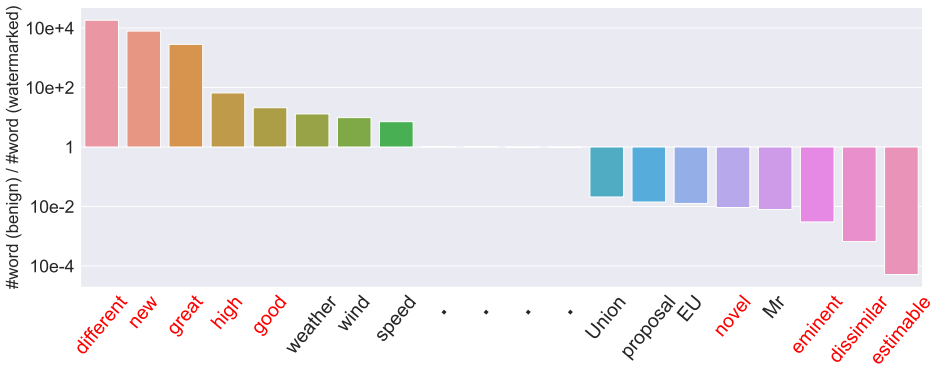

Figure 1: Ratio change of word frequency of top 100 words between benign and watermarked corpora, namely $P_b(w)/P_w(w)$. Red words are the selected watermarks. Although we only list 16 words having the most significant ratio change in the benign and watermarked corpora and omit the rest of them for better visualization, all watermarks are within the top 100 words.

In order to protect the IP of victim models, He et al. [10] first introduced a watermarking algorithm to text generation and utilized the null-hypothesis test as a post-hoc ownership verification on the imitation models. However, traditional watermarking methods generally distort the word distribution, which could be utilized by attackers to infer the watermarked words via sufficient statistics of the frequency change of candidate watermarking words. As an example shown in Figure 1, the replaced words and their substitutions are those with most frequency decrease ratios and frequency increase ratios, respectively. To address this drawback, we are motivated to develop a more stealthy watermarking method to protect the IP of text generation APIs. The stealthiness of the new watermarks is achieved by incorporating high-order linguistic features as conditions that trigger corresponding watermarking rules.

Overall, our main contributions are as follows:

- We propose a novel Conditional wATERmarking framework (CATER) for protecting text generation APIs. An optimization method is proposed to decide the watermarking rules that can *i)* minimize the distortion of overall word distributions, while *ii)* maximize the change of conditional word selections.
- Theoretically, we prove that a small number of the used watermarks could be blended in and camouflaged by a large number of suspicious watermarks when attackers attempt to inverse the backend watermarking rules.
- Empirically, we observe that high-order conditions lead to the exponential growth of suspicious (unused) watermarks, which encourages better stealthiness of the proposed defense with little hurt to the generation of victim models.

2 Preliminary and Background

2.1 Imitation Attack

An imitation attack (a.k.a model extraction) aims to emulate the behavior of the victim model $\mathcal{V}$, such that the adversary can either sidestep the service charges or launch a competitive service [36, 15, 40, 9, 41]. Malicious users can achieve this goal through interaction with the victim model $\mathcal{V}$ without knowing its internals, such as the model architecture, hyperparameters, training data, *etc.* Adversaries first craft a set of queries Q based on the documentation of a target model. Then Q will be sent to $\mathcal{V}$ to obtain the corresponding predictions Y. Finally, an imitation model $\mathcal{S}$ can be attained by learning a function to map Q to Y.

Most prior imitation attacks are limited to classification tasks [36, 24, 9]. The imitation for text generation, a crucial task in natural language processing, has been under-developed until recently. Inspired by the efficacy of sequence-level knowledge distillation [13], Wallace et al. [40] and Xu

et al. [41] propose mimicking the functionally of commercial text generation APIs. Similar to the standard imitation attack, adversaries can query $\mathcal{V}$ with Q. For generation tasks, Y is a sequence of tokens $(y_1, ..., y_L)$, where L is the length of the sequence. According to their empirical studies, one can rival the performance of these APIs, which poses a severe threat to cloud platforms.

2.2 Identification of IP Infringement

Prior works have utilized watermarking avenues to achieve a post-hoc verification of the ownership [37, 19, 20]. However, this line of work assumes the model owners can watermark victim model $\mathcal{V}$ by altering its neurons before releasing $\mathcal{V}$ to end-users. This operation is not feasible for the imitation attack, as $\mathcal{V}$ cannot access the parameters of $\mathcal{S}$. The only thing under the control of $\mathcal{V}$ is the responses to adversaries. Hence, some recent works propose creating a backdoor to $\mathcal{S}$ during the interaction with attackers [15, 34]. Specifically, $\mathcal{V}$ can select a small fraction of queries and answer them with incorrect predictions, in a similar way to the popular choice of watermarks in the computer vision domain (adopting some arbitrary features as the trigger to evaluate) [7, 43]. Erroneous predictions are so abrupt that $\mathcal{S}$ will memorize these outliers [4, 16]. As such, $\mathcal{V}$ can utilize these watermarks as evidence of ownership.

Albeit the efficacy, the drawbacks of backdoor approaches are tangible as well. First, since $\mathcal{V}$ does not impose regulations on users' usage, one cannot distinguish a malicious user from a regular user based on their querying behaviors[1]. Thus, $\mathcal{V}$ has to fairly serve all users and store all mislabeled queries, which leads to a massive storage consumption and a negative impact on the users' experiences. Moreover, as the identity of imitation models is unknown to $\mathcal{V}$, $\mathcal{V}$ has to iterate over all the mislabeled queries, which is computationally prohibitive. Finally, as $\mathcal{S}$ tends to adopt the pay-as-you-use policy for the sake of profits, the brute-force interaction with $\mathcal{S}$ can cause drastic financial costs.

As a remedy, He et al. [10] utilize a lexical watermark to identify IP infringement brought by imitation attacks. They point out that a neat watermarking algorithm must follow two principles: *i*) it cannot significantly impair customer experience, and *ii*) it should not be reverse-engineered by malicious users. In order to fulfill these requirements, they first select a set of words $\mathcal{W}$ from the training data of the victim model $\mathcal{V}$. Then for each $w \in \mathcal{W}$, they find R–1 semantically equivalent substitutions for it. Next, they employ $\mathcal{W}$ and their substitutions $\mathcal{T}$ to compose watermarking words $\mathcal{M}$. Finally, they replace $\mathcal{W}$ with $\mathcal{M}$. The rationale behind this avenue is to alter the distribution of words such that the imitation model can learn this biased pattern. To verify such a biased pattern of the word choice, He et al. [10] employ a null hypothesis test [29] for evaluation.

More concretely, He et al. [10] utilize an evaluation set O to conduct the null hypothesis test. They formulate the null hypothesis as: *the tested model generates outputs without preference for watermarks*. A null hypothesis can be either rejected or accepted via the calculation of a p-value [29]. They assume that all words $\{w_i | w_i \in \mathcal{W} \cup \mathcal{T}\}$ follow a binomial distribution $Pr(k; n, p)$, where k is the number of words in $\mathcal{M}$ appearing in O, n is the number of words in $\mathcal{W} \cup \mathcal{T}$ found in O, and p is the probability of watermarks observed in the natural language. According to their algorithm, p is approximated by $1/R$. Now, one can compute the p-value from as follows:

$$\mathcal{P} = 2 \cdot \min(Pr(X \geq k), Pr(X \leq k)) \tag{1}$$

The p-value indicates how one can confidently reject the hypothesis. Lower p-value suggests that the tested model should be more likely subject to an imitator.

3 CATER

This section introduces our proposed *CATER*, leveraging conditional watermarks to watermark the imitation model, which can be served as a post-hoc identification of an IP infringement. Figure 2 provides an overview of CATER, consisting of two stages: *i) Watermarking Stage* and *ii) Identification Stage*.

[1] `https://cloud.google.com/translate/pricing`

***i)* Watermarking Stage** The victim model $\mathcal{V}$ employs CATER to build a list of watermarking rules (with conditions) from training data. When received a query Q from an end-user, $\mathcal{V}$ initially produces a tentative answer Y. Next, $\mathcal{V}$ utilizes CATER to conduct a watermarking procedure over Y according to the watermarking rules. Finally, $\mathcal{V}$ replies the end-user with a watermarked response Y'.

***ii)* Identification Stage** If a model $\mathcal{S}$ is under suspicion, $\mathcal{V}$ can query it with a verification set O. After obtaining the responses Y from $\mathcal{S}$, $\mathcal{V}$ can leverage CATER to testify whether $\mathcal{S}$ violates the IP right of $\mathcal{V}$.

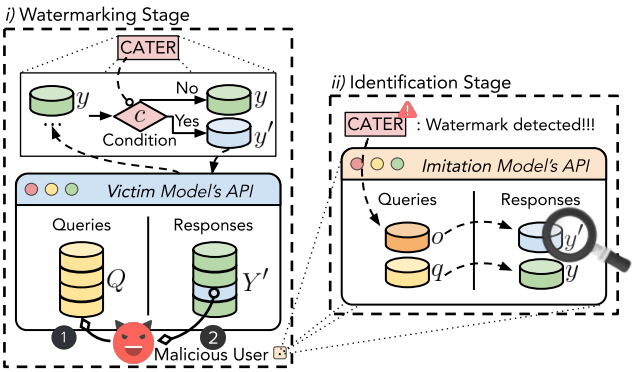

Figure 2: The workflow of CATER IP protection for Generation APIs. CATER first watermarks some of the responses from victim APIs (left). Then, CATER identifies stolen API by watermark verification (right).

3.1 Watermarking Rule Optimization

Watermarking some words to a deterministic substitutions could distort the overall word distribution. Therefore, some watermarks could be reversely inferred and eliminated by analyzing the word distribution, as demonstrated in Figure 1. We propose to inject the watermarks in conditional word distribution, while maintaining the original word distribution. The substitutions can be conditioned on linguistic features as illustrated in Figure 3. Remarkably, given a condition $c \in \mathcal{C}$ and a group of semantically equivalent words $\mathcal{W}$, one can replace any words $w \in \mathcal{W}$ with each other. We formulate the objective of conditional watermarking rules as:

$$\min_{\hat{P}(w|c)} \underbrace{\mathbb{D}\Big(\sum_{c\in\mathcal{C}} \hat{P}(w|c)P(c), \sum_{c\in\mathcal{C}} P(w|c)P(c)\Big)}_{\text{I: indistinguishable objective}} - \frac{\alpha}{|\mathcal{C}|} \underbrace{\sum_{c\in\mathcal{C}} \mathbb{D}\big(\hat{P}(w|c), P(w|c)\big)}_{\text{II: distinct objective}} \tag{2}$$

The two factors reflect two essential desiderata:

1. For each $\mathcal{W}$, with $w \in \mathcal{W}$, the overall word distributions before optimization $P(w) = \sum P(w|c)P(c)$ and after optimization $\hat{P}(w) = \sum \hat{P}(w|c)P(c)$ should be close to each other, as the *indistinguishable objective* in Equation 2;
2. For a particular condition $c \in \mathcal{C}$, the conditional word distributions should still be distinct to their original distributions, reflected by the dissimilarity between $P(w^{(i)}|c)$ and $\hat{P}(w^{(i)}|c)$, as the *distinct objective* in Equation 2. This guarantees the conditional watermarks are identifiable in verification.

In practice, we utilize multiple synonym word sets as a group $\mathcal{G} = \{\mathcal{W}^{(i)}\}_{i=1}^{|\mathcal{G}|}$. For each $\mathcal{W}^{(i)}$, we can formulate Equation 2 as a mixed integer quadratic programming using ℓ_2-norm as distance measurement function:

$$\begin{aligned} &\min_{\boldsymbol{W}} (\boldsymbol{W}\boldsymbol{c} - \boldsymbol{X}\boldsymbol{c})^T(\boldsymbol{W}\boldsymbol{c} - \boldsymbol{X}\boldsymbol{c}) - \frac{\alpha}{|\mathcal{C}|}\mathrm{Tr}\big((\boldsymbol{W} - \boldsymbol{X})^T(\boldsymbol{W} - \boldsymbol{X})\big) \\ &\text{s.t. } \boldsymbol{X}^T \cdot \boldsymbol{1}_{|\mathcal{W}^{(i)}|} = \boldsymbol{1}_{|\mathcal{C}|}, \boldsymbol{X} \in \{0,1\}^{|\mathcal{W}^{(i)}|\times|\mathcal{C}|} \end{aligned} \tag{3}$$

We define matrix $\boldsymbol{X} = [\hat{P}(w^{(i)}|c)]_{|\mathcal{W}^{(i)}|\times|\mathcal{C}|}$ as the variables for optimization. Matrix $\boldsymbol{W} = [P(w^{(i)}|c)]_{|\mathcal{W}^{(i)}|\times|\mathcal{C}|}$ and vector $\boldsymbol{c} = [P(c)]_{|\mathcal{C}|\times 1}$ are constant variables, decided by calculating corresponding distributions in a large training corpus. The objective of Equation 3 is convex when α is sufficiently small (see the proof in Appendix A). We optimize the watermark assignments $\hat{P}(w^{(i)}|c)$ using Gurobi [8] with $\alpha = 0.01$.

3.2 Constructing Watermarking Conditions using Linguistic Features

This part will concentrate on practical ways to construct the watermarking conditions $\mathcal{C}$. We consider two fundamental linguistic features $\mathcal{F}$, *i)* part-of-speech and *ii)* dependency tree, and their high-order variations as conditions. Such linguistic features were widely and successfully used for text classification [3, 42], sequence labeling [17, 30], *etc.*

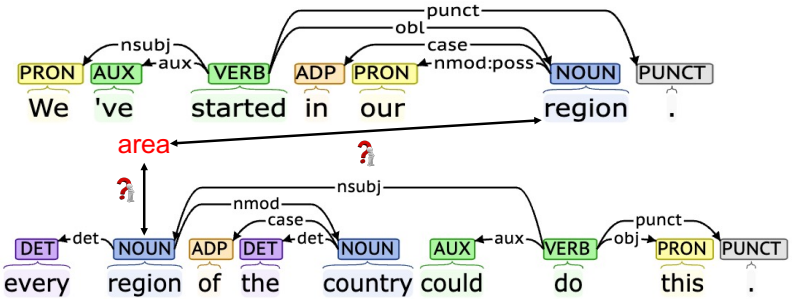

Figure 3: The part-of-speech (POS) tags and dependency relations are illustrated in colored boxes and arcs, respectively. The decision of using "region" or its synonym ("area" or other words), is conditioned on its linguistic features in context. For example, in the first sentence, either "PRON" (POS label of the preceding token) or "obl" (DEP label of the incoming arc) gives the decision to replace "region" with "area", as opposed to "DET" or "nsubj" in the second sentence.

Part-of-Speech Part-of-Speech (POS) tagging is a grammatical grouping algorithm, which can cluster words according to their grammatical properties, such as syntactic and morphological behaviors [12]. The POS tag for each token is demonstrated in a colored box, in Figure 3.

Given a word w in a sentence, we denote its POS as l_0, and use l_{-k} or l_{+k} to represent the POS of the k-th word to the left or right of w. We consider a single label l_{-1} as our first-order condition. In order to reduce the identifiability of our conditional watermark, we can construct high-order conditions from the same feature set, *e.g.*, (l_{-1}, l_{+1}) as second-order condition and (l_{-2}, l_{-1}, l_{+1}) as third-order condition. Note that if l_{-k} or l_{+k} does not exist, we use a pseudo tag "[none]" by default. Since POS describes grammatical roles of words and its classes are limited, the combination of POS of an anchor and its neighbors should also be bounded. Thus one can consider the POS bond among words as the condition.

Dependency Tree Dependency Tree (DEP) is a syntactic structure, which describes directed binary grammatical relations between words [12], as shown in Figure 3. A dependency tree can be represented by an acyclic directed graph $G = (V, E)$, where V is a set of vertices corresponding to all words in a given sentence, E is a set of ordered pairs of vertices, denoted as *arcs*. An arc $e \in E$ describes a grammatical relation between two vertices in V, *i.e.,* source vertex named as *head* and target vertex coined as *dependent*. Except the root vertex, each vertex is connected to by exactly one head. Consequently, there exists a unique path from each vertex to the root node in a dependency tree.

Analagously for POS features, we can design first-order and high-order DEP features as watermarking conditions. Given a word w, and its incoming DEP arc, we use the DEP label of the arc as the first-order features (d_1). We construct high-order condition recursively using the labels of incoming DEP arcs $(d_1, d_2, \cdots)$. A pseudo arc label "[none]" is used when there is no parent node in recursion.

3.3 Identifiability of Conditional Watermark

In this section, we discuss the identifiability of our watermark if the attackers attempt to infer the used watermarks. We assume the worst case that the attackers have access to *i)* the watermarking algorithm, *ii)* all possible word sets for substitution $\mathcal{G}$, and *iii)* combination of feature sets $\mathcal{F}$ as wartermarking conditions $\mathcal{C}$. However, the exact watermarking rules are unobservable to attackers. The attackers may identify the watermark rules by suspecting those observed $P(w^{(i)}|c)$ with extreme distributions, *i.e.*, only a single word in a synonym set is selected given a specific condition.

Given a limited budget, we assume that an imitator has queried our watermarked API and has acquired N tokens in $\bigcup_i \mathcal{W}^{(i)}$. The system has incorporated watermarks with K-order features $c \in \mathcal{C}$, where $\mathcal{C} = (\mathcal{F}_1, \mathcal{F}_2, \cdots \mathcal{F}_K)$, the total number of possible conditions is $|\mathcal{C}| = \prod_{i=1}^{K} |\mathcal{F}_i|$. We simplify our discussion by using the same feature set $\mathcal{F}$ (POS or DEP), then $|\mathcal{C}| = |\mathcal{F}|^K$.

Theorem 3.1. *If $|\mathcal{F}|^K > N$ and there exist t conditions that have less or equal to $m \in \mathbb{Z}^+$ support samples, then $t \geq |\mathcal{F}|^K - N/(m+1)$.*

The attacker would suspect conditional word distributions that are *extremely imbalanced*, namely only a single dominant choice of word is observed within the responses to the attacker.

Theorem 3.2. *Having m support samples for a specific condition c, the possibility of observing extremely imbalanced word choice is $\mathcal{I}(\mathcal{W}, c, m) = \sum_{w_i \in \mathcal{W}} P(w_i|c)^m$. If $m' \leq m$ and $m, m' \in \mathbb{Z}^+$, $\mathcal{I}(\mathcal{W}, c, m') \geq \mathcal{I}(\mathcal{W}, c, m)$.*

The proofs of Thm 3.1 and Thm 3.2 can be found in Appendix B and C. Thm 3.1 guarantees a lower bound for the number of conditions that attackers will have less or equal to m observed samples. Moreover, the lower bound grows exponentially with regard to the feature orders used as watermarking conditions. Thm 3.2 guarantees the high probability of the conditions with extremely imbalanced word selection when m is small. Combining these two theorems, the total number of the suspicious watermark rules could be huge compared with the limited used watermark rules if we are utilizing high-order linguistic features as conditions. We further empirically demonstrate the significant confusion between suspected and used watermark rules in Section 4.2.

4 Experiments

Text Generation Tasks. We examine two widespread text generation tasks: machine translation and document summarization, which have been successfully deployed as commercial APIs.[2][3]

- **Machine Translation:** We consider WMT14 German (De) →English (En) translation [1] as the testbed. We follow the official split: train (4.5M) / dev (3,000) / test (3,003). Moses [14] is applied to pre-process all corpora, with a cased tokenizer. We use BLEU [26] and BERTScore [44] to evaluate the translation quality. BLEU concentrates on lexical similarity via n-grams match, whereas BERTScore targets at semantic equivalence through contextualized embeddings.
- **Document summarization:** CNN/DM [11] utilizes informative headlines as summaries of news articles. We reuse the dataset preprocessed by See et al. [31] with a partition of train/dev/test as 287K / 13K / 11K. Rouge [21] and BERTScore [44] are employed for the evaluation metric of the summary quality.

We use 32K and 16K BPE vocabulary [32] for experiments on WMT14 and CNN/DM, respectively.

Models. For the primary experiments, we consider Transformer-base [38] as the backbone of both victim models and the imitation models. Following He et al. [10], we use a 3-layer Transformer for the summarization task. Because of their superior performance, pre-trained language models (PLMs) have been deployed on cloud platforms.[4] Hence, we also consider using two popular PLMs: *i)* BART (summarization) [18] and *ii)* mBART (translation) [22] as the victim model. Regarding the imitation model, since the architecture of the victim model is unknown to the adversary, we simulate this black-box setting by using three different architectures as the imitator, namely (m)BART, Transformer-base, and ConvS2S [6]. The training details are summarized in Appendix D.

Basic Settings. As a proof-of-concept, we start our evaluation with a most straightforward case. We assume the victim model $\mathcal{V}$ and the imitation model $\mathcal{S}$ use the same training data, but $\mathcal{S}$ uses the response $\boldsymbol{y}'$ with CATER instead of the ground-truth $\boldsymbol{y}$. We set the size of synonyms to 2 and

[2] `https://translate.google.com/`

[3] `https://deepai.org/machine-learning-model/summarization`

[4] `https://cloud.google.com/ai-platform/training/docs/algorithms/bert`

Table 1: Performance of different watermarking approaches on WMT14 and CNN/DM. We use F1 scores of ROUGE-1, ROUGE-2 and ROUGE-L for CNN/DM.

	WMT14			CNN/DM		
	p-value $\downarrow$	**BLEU** $\uparrow$	**BERTScore** $\uparrow$	**p-value** $\downarrow$	**ROUGE-1/2/L** $\uparrow$	**BERTScore** $\uparrow$
w/o watermark	$> 10^{-1}$	31.1	65.9	$> 10^{-1}$	37.7 / 15.4 / 31.2	22.1
Venugopal et al. [39]						
- unigram	$< 10^{-2}$	30.4	65.2	$< 10^{-2}$	37.3 / 15.1 / 31.2	21.7
- trigram	$> 10^{-1}$	30.8	65.7	$> 10^{-1}$	37.5 / 15.3 / 31.0	21.8
- sentence	$> 10^{-1}$	30.8	65.9	$> 10^{-1}$	37.6 / 15.4 / 31.2	21.9
He et al. [10]						
- spelling	$< 10^{-13}$	31.1	65.8	$< 10^{-8}$	37.5 / 15.2 / 31.4	22.0
- synonym	$< 10^{-10}$	30.8	65.5	$< 10^{-8}$	37.6 / 15.3 / 31.4	21.8
CATER (ours)						
- DEP	$< 10^{-4}$	30.9	65.4	$< 10^{-2}$	37.6 / 15.3 / 31.3	21.8
- POS	$< 10^{-7}$	30.8	65.3	$< 10^{-7}$	37.5 / 15.2 / 31.2	21.9

vary this value in Appendix E.1. The detailed construction of watermarks and approximation of p in Equation 1 for CATER is provided in Appendix D.

Baselines. We compare our approach with [39] and [10]. Venugopal et al.[39] proposed watermarking the generated output with a sequence of bits under the representation of either n-grams or the complete sentence. He et al. [10] devises two effective watermarking approaches. The first one replaces all the watermarked words with their synonyms. The second one watermarks the victim API outputs by mixing American and British spelling systems.

4.1 Performance of CATER

Table 1 presents the watermark identifiability and generation quality of studied text generation tasks. Both [10] and CATER obtain a sizeable gap in the p-value, and demonstrate a negligible degradation in BLEU, ROUGE, and BERTScore, compared to the non-watermarking baseline. However, [39] falls short of injecting detectable watermarks. Although CATER is slightly inferior to [10] in p-value, we argue that watermarks in [10] can be easily erased, as their replacement techniques are not invisible. As shown in Figure 1, the synonyms used by [10] can be identified due to the tangible distribution shift on the watermarks. In addition, one can eliminate the spelling watermarks by consistently using one spelling system.

Unless otherwise stated, we use the first-order POS as the default setting for CATER, due to its efficacy in terms of watermark identifiability and generation quality.

Table 2: Imitation performance of different architectures on clean and watermarked data. Numbers in parentheses are results of clean data. Victim models are trained on mBART (WMT14) and BART (CNN/DM), respectively. We use the first-order POS as the watermarking approach.

Model	WMT14		CNN/DM	
	p-value $\downarrow$	**BLEU** $\uparrow$	**p-value** $\downarrow$	**ROUGE-L** $\uparrow$
(m)BART	$< 10^{-4}$ ($> 10^{-1}$)	34.9 (35.2)	$< 10^{-5}$ ($> 10^{-1}$)	38.1 (38.1)
Transformer	$< 10^{-5}$ ($> 10^{-2}$)	32.7 (33.0)	$< 10^{-3}$ ($> 10^{-1}$)	32.8 (32.9)
ConvS2S	$< 10^{-5}$ ($> 10^{-2}$)	32.7 (32.9)	$< 10^{-3}$ ($> 10^{-1}$)	32.7 (32.7)

IP Identification under Architectural Mismatch The architectures of remote APIs are usually unknown to the adversary. However, recent works have shown that the imitation attack is effective even if there is an architectural mismatch between the victim model and the imitator [40, 9]. To demonstrate that our approach is model-agnostic, we use BART-family models as victim models and vary architectures of imitation models.

Table 2 summarizes p-value and generation quality of CATER on WMT14 and CNN/DM datasets. Similar to Table 1, CATER can confidently identify the IP infringement when the architecture of the

imitation model is the same as that of the victim model, with a gap of p-value between watermarked model and benign model being 3 orders of magnitude. In addition, this gap applies to the case, where we use distinct architectures for the victim model and the imitator. The generation quality exhibits negligible drops, within a range of 0.3. Note that the generation quality of Transformer and ConvS2S imitators degrades due to the capacity gap, compared to powerful BART-family models.

IP Identification on Cross-domain Imitation Similarly, the training data of the victim model is confidential and remains unknown to the public. Thus, there could be a domain mismatch between the training data of the victim model and queries from the adversary. In order to exhibit that our approach is exempt from the domain shift, we use two out-of-domain datasets to conduct the imitation attack for the machine translation task. The first is IWSLT14 data [2] with 250K German sentences, and the second is OPUS (Law) data [35] consisting of 2.1M German sentences. Table 3 suggests that despite the domain mismatch, CATER can still watermark the imitation model, and one can identify watermarks with high confidence.

Table 3: Imitation performance of using data from different domains. The victim model is trained on WMT14. We use first-order POS as the watermarking condition.

WMT14	IWSLT14	OPUS (Law)
$< 10^{-7}$	$< 10^{-5}$	$< 10^{-6}$

High-order Conditions We have shown that the first-order CATER effectively performs various tasks and settings. We argue that CATER is not limited to the first-order condition. Instead, one can use high-order CATER as mentioned in Section 3.2, which can consolidate the invisibility as discussed in Section 3.3. Therefore, we investigate the efficacy of the high-order CATER to the translation task and provide the study on the summarization task in Appendix E.2.

Figure 4 shows that with the increase of conditional POS orders, there is no side effect on the generation quality. Although the gap between the benign and watermarked models is mitigated using high-order conditions, it is still significantly distinguishable.

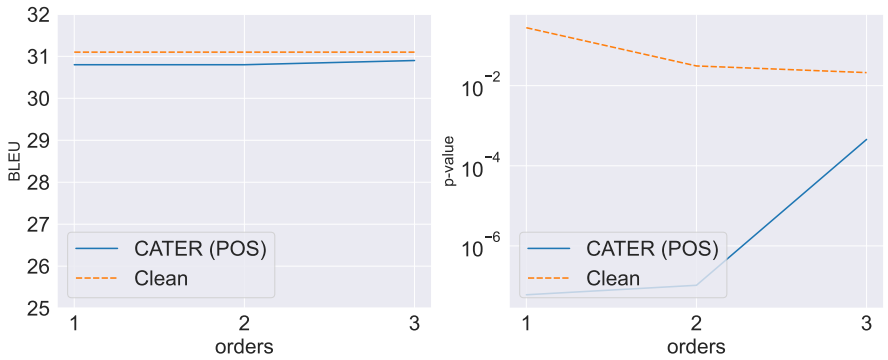

Figure 4: BLEU scores and p-value of using different orders of the POS watermarking approach on WMT14 data.

4.2 Analysis on Adaptive Attacks

The previous sections illustrates the efficacy of CATER for watermarking and detecting potential imitation attacks. Given the case that a savvy attacker might be aware of the existence of watermarks, they might launch countermeasures to remove the effects of the watermark. This section explores and analyzes possible adaptive attacks based on varying degrees of prior knowledge of our defensive strategy. Specifically, we examine two types of adaptive attacks that try to erase the effects of the watermark: *i)* **vanilla watermark removal**, and *ii)* **watermarking algorithm leakage**.

Vanilla Watermark Removal. Under this setting, we assume the attackers are aware of the existence of watermarks, but they are not aware of the details of the watermarking algorithm.

Table 4: Imitation performance with watermark removal on WMT14 data and CNN/DM. We use the first-order POS as the watermarking approach. ONION is used to remove watermarks.

Model	WMT14		CNN/DM	
	p-value ↓	BLEU ↑	p-value ↓	ROUGE-L ↑
w/o ONION	$< 10^{-7}$	30.8	$< 10^{-7}$	31.2
w/ ONION	$< 10^{-5}$	27.0	$< 10^{-7}$	25.9

Following such a setting of attacker knowledge, we assume the attacker would adopt an existing watermark removal technique in their vanilla form. To evaluate, we employ ONION, a popular defensive avenue for data poisoning in the natural language processing field, which adopts GPT-2 [28] to expel outlier words. The defense results are shown in Table 4. We find that ONION cannot erase the injected watermarks. Meanwhile, it drastically diminishes the generation quality of the imitation model.

Watermarking Algorithm Leakage. Under this case study, we assume attackers have access to the full details of our watermarking algorithm, *i.e.,* the same watermarking dictionary and the features for constructing watermarking conditions. We note that this is the most substantial attacker knowledge assumption we can imagine, aside from the infeasible case that they know the complete pairs of watermarks we used. After collecting responses from the victim model, the attackers can leverage the leaked knowledge to analyze the responses to find the used watermarks. As shown in Section 3.3, we theoretically prove that such a reverse-engineering is infeasible. In addition, Figure 5 shows that even with such a strong attacker knowledge, the amount of potential candidate watermarks (orange curve) is still astronomical times larger than the used number of watermarks (blue curve). Thus, malicious users would have difficulty in removing watermarks from the responses; unless they lean toward modifying all potential watermarks. Such a brute-force approach can drastically debilitate the performance of the imitation attack, causing a feeble imitation. Finally, we demonstrate the upper bound (green curve) to show that without the curated knowledge about the watermarking conditions, the difficulty of identifying the watermarks from top 200 words can be combinatorially exacerbated.

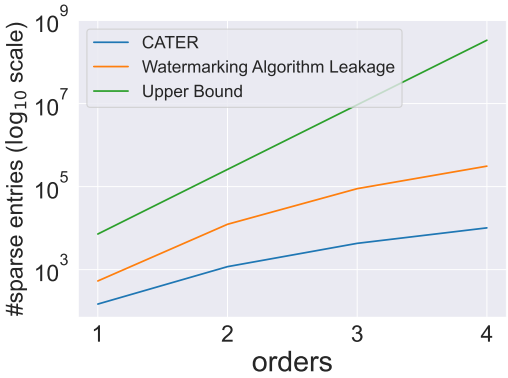

Figure 5: The number of sparse entries (suspected watermarks) of top 200 words with watermarking algorithm leakage under different orders (orange) on the training data. CATER indicates the actual number of watermarks used by our watermarking system (blue). POS feature is used where $|\mathcal{F}| = 36$. The upper bound indicates all possible combinational watermarks (green).

5 Conclusion

In this work, we are keen on protecting text generation APIs. We first discover that it is possible to detect previously proposed watermarks via sufficient statistics of the frequencies of candidate watermarking words. We then propose a novel Conditional wATERmarking framework (CATER), for which, an optimization method is proposed to decide the watermarking rules that can minimize the distortion of overall word distributions while maximizing the change of conditional word selections. Theoretically, we prove that it is infeasible for even the savviest attackers, who know how CATER algorithms, to reveal the used watermarks from a large pool of potential watermarking rules based on statistical inspection. Empirically, we observe that high-order conditions lead to an exponential growth of suspicious (unused) watermarks, rendering our crafted watermarks more stealthy.

References

[1] Ondrej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA, June 2014. Association for Computational Linguistics.

[2] Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of the International Workshop on Spoken Language Translation, Hanoi, Vietnam*, volume 57, 2014.

[3] William Cohen, Vitor Carvalho, and Tom Mitchell. Learning to classify email into "speech acts". In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 309–316, 2004.

[4] Jiazhu Dai, Chuanshuai Chen, and Yufeng Li. A backdoor attack against lstm-based text classification systems. *IEEE Access*, 7:138872–138878, 2019.

[5] Christiane Fellbaum. Wordnet. In *Theory and applications of ontology: computer applications*, pages 231–243. Springer, 2010.

[6] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*, pages 1243–1252. PMLR, 2017.

[7] Shangwei Guo, Tianwei Zhang, Han Qiu, Yi Zeng, Tao Xiang, and Yang Liu. Fine-tuning is not enough: A simple yet effective watermark removal attack for dnn models. *arXiv preprint arXiv:2009.08697*, 2020.

[8] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2022.

[9] Xuanli He, Lingjuan Lyu, Lichao Sun, and Qiongkai Xu. Model extraction and adversarial transferability, your bert is vulnerable! In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2006–2012, 2021.

[10] Xuanli He, Qiongkai Xu, Lingjuan Lyu, Fangzhao Wu, and Chenguang Wang. Protecting intellectual property of language generation apis with lexical watermark. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2022.

[11] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28, 2015.

[12] Daniel Jurafsky and James Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, volume 2. 02 2008.

[13] Yoon Kim and Alexander M Rush. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, 2016.

[14] Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic, June 2007. Association for Computational Linguistics.

[15] Kalpesh Krishna, Gaurav Singh Tomar, Ankur P. Parikh, Nicolas Papernot, and Mohit Iyyer. Thieves on sesame street! model extraction of bert-based apis. In *International Conference on Learning Representations*, 2020.

[16] Keita Kurita, Paul Michel, and Graham Neubig. Weight poisoning attacks on pretrained models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2793–2806, 2020.

[17] John D Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 2001.

[18] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.

[19] Meng Li, Qi Zhong, Leo Yu Zhang, Yajuan Du, Jun Zhang, and Yong Xiangt. Protecting the intellectual property of deep neural networks with watermarking: The frequency domain approach. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 402–409. IEEE, 2020.

[20] Jian Han Lim, Chee Seng Chan, Kam Woh Ng, Lixin Fan, and Qiang Yang. Protect, show, attend and tell: Empowering image captioning models with ownership protection. *Pattern Recognition*, 122:108285, 2022.

[21] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.

[22] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020.

[23] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

[24] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. Knockoff nets: Stealing functionality of black-box models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4954–4963, 2019.

[25] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, 2019.

[26] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.

[27] Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. Stanza: A Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2020.

[28] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

[29] John A Rice. *Mathematical statistics and data analysis*. Cengage Learning, 2006.

[30] Sunita Sarawagi and William W Cohen. Semi-markov conditional random fields for information extraction. In L. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 17. MIT Press, 2004.

[31] Abigail See, Peter J Liu, and Christopher D Manning. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, 2017.

[32] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, 2016.

[33] Amit Singhal. Microsoft's bing uses google search results—and denies it, 2011.

[34] Sebastian Szyller, Buse Gul Atli, Samuel Marchal, and N Asokan. Dawn: Dynamic adversarial watermarking of neural networks. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4417–4425, 2021.

[35] Jörg Tiedemann. Parallel data, tools and interfaces in opus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, 2012.

[36] Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pages 601–618, 2016.

[37] Yusuke Uchida, Yuki Nagai, Shigeyuki Sakazawa, and Shin'ichi Satoh. Embedding watermarks into deep neural networks. In *Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval*, pages 269–277, 2017.

[38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[39] Ashish Venugopal, Jakob Uszkoreit, David Talbot, Franz Och, and Juri Ganitkevitch. Watermarking the outputs of structured prediction with an application in statistical machine translation. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1363–1372, Edinburgh, Scotland, UK., July 2011. Association for Computational Linguistics.

[40] Eric Wallace, Mitchell Stern, and Dawn Song. Imitation attacks and defenses for black-box machine translation systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5531–5546, 2020.

[41] Qiongkai Xu, Xuanli He, Lingjuan Lyu, Lizhen Qu, and Gholamreza Haffari. Beyond model extraction: Imitation attack for black-box nlp apis. *arXiv preprint arXiv:2108.13873*, 2021.

[42] Qiongkai Xu and Hai Zhao. Using deep linguistic features for finding deceptive opinion spam. In *Proceedings of COLING 2012: Posters*, pages 1341–1350, 2012.

[43] Yi Zeng, Won Park, Z Morley Mao, and Ruoxi Jia. Rethinking the backdoor attacks' triggers: A frequency perspective. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16473–16481, 2021.

[44] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2019.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section X.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 4.1
   (b) Did you describe the limitations of your work? [Yes] See Appendix F
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Appendix F
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sections 3.1 and 3.3
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendices A, B and C
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See `https://anonymous.4open.science/r/cater_neurips-1685`
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.1 and Appendix D
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 4.1
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Appendix D
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 4.1
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]