

SoK: Are Watermarks in LLMs Ready for Deployment?

Kieu Dang¹, Phung Lai¹, NhatHai Phan², Yelong Shen³, Ruoming Jin⁴, Abdallah Khreishah², My Thai⁵

Abstract—The deployment of a proprietary LLM raises significant intellectual property (IP) violation risks. An adversary can replicate an LLM by using input prompt-the LLM’s output pairs to train a surrogate model. It leads to financial setbacks for the LLM’s service provider. A practical solution for this problem is LLM watermarking, in which the service provider implants an imperceptible pattern (watermark) into the LLM’s outputs such that it can detect the watermark from the adversary’s generated text without affecting the LLM’s utility. While various watermarking techniques have emerged to mitigate these risks, it remains unclear how far the community and industry have progressed in developing and deploying watermarks in LLMs.

To bridge this gap, we aim to develop a comprehensive systematization for WMs in LLMs by 1) presenting a detailed taxonomy for WMs in LLMs, 2) proposing a novel IP classifier to explore the effectiveness and impacts of WMs on LLMs under both attack and attack-free environments, 3) analyzing the limitations of existing WMs in LLMs, and 4) discussing practical challenges and potential future directions for WMs in LLMs. Through extensive experiments, we show that despite promising research outcomes and significant attention from leading companies and community to deploy WMs, these techniques have yet to reach their full potential in real-world applications due to their unfavorable impacts on model utility of LLMs and downstream tasks. Our findings provide an insightful understanding of WMs in LLMs, highlighting the need for practical WM solutions tailored to LLM deployment.

1. Introduction

Large language models (LLMs), such as ChatGPT, Gemini, Claude, and Cohere [1]–[4], have demonstrated remarkable capabilities in text generation, machine translation, and knowledge understanding tasks [5]–[9]. They effectively mimic human writing behaviors and generate complex and coherent outputs from the input text, making it challenging to determine whether a text is authored by humans or generated by LLMs. Due to the high demands of computational resources and human efforts required for training LLMs [10], [11], these models are commonly offered as a service through application programming interfaces (APIs), typically requiring users to pay or subscribe [12], [13]. Although users cannot access to the model weights or architectures of these commercial LLMs, this restriction does not ensure the safety of these models. Malicious actors can intentionally mimic cloud-hosted LLM behaviors to offer

cheaper services [14], [15]. To conduct this service stealing, an adversary can query a set of inputs through an LLM’s API to retrieve the corresponding outputs. Then, the adversary uses these input-output data to effectively fine-tune their local model. When the number of queries is sufficient to gather enough input-output data within a particular domain, the adversary can steal the cloud-hosted LLM behaviors in that domain [16]–[18]. Consequently, these concerns underscore significant risks regarding the intellectual property (IP) rights of the cloud-hosted proprietary LLMs [19].

To address such risks, service providers, i.e., the cloud, have employed many strategies including watermarks, encryption, limited model exposure via APIs, and differential privacy [9], [20], [21]. Among them, watermarks (WMs) [9], [22]–[25] have emerged as a practical tool for LLMs due to its ability to ensure traceability, protect IP, and detect misuse. Typically, WMs embed imperceptible patterns into the LLM outputs directly. These patterns can be used to trace the original text or determine whether the generated text from an LLM is watermarked, facilitating the detection of unauthorized use. To watermark outputs, the cloud can introduce bias into logits of token generations to favor a set of specific tokens or alter the token sampling process of LLMs. Once a WM is applied, the cloud can use IP checkers to determine if a set of outputs from a suspicious model is watermarked. This allows the cloud to assess whether the suspicious model has been trained or fine-tuned using the cloud’s watermarked outputs. Typically, IP checkers analyze the tokens and perform statistical tests to check if the values exceed certain thresholds, indicating that the cloud-hosted LLM has been imitated.

WMs are effective in detecting IP violations. However, WMs can be susceptible to WM removal attacks [26], [27] or spoofing attacks [27]. These attacks either require significant computational resources for additional training to paraphrase the generated text of cloud-hosted LLMs or demand numerous queries to learn the distribution of WMs. While several studies have surveyed WMs on LLMs [28]–[30], no research has systematically and extensively explored the effectiveness of WMs on LLMs, particularly regarding the impact of WMs on LLM utility and IP checkers, under both attack and attack-free environments. In addition, the practicality of WMs has not been rigorously assessed. Therefore, in this work, we aim to develop a comprehensive systematization of WMs in LLMs, with the following contributions:

- We present a taxonomy for WMs that effectively mitigate model stealing attacks.
- We propose a novel cross-model IP classifier to explore the effectiveness and impacts of WMs on LLMs, given

¹SUNY-Albany, ²New Jersey Institute of Technology, ³Microsoft, ⁴Kent State University, ⁵University of Florida. Contact: vdang@albany.edu

both attack and attack-free environments.

- We suggest well-suited applications and limitations of different WMs in LLMs.
- We conduct extensive experiments to demonstrate the effectiveness and practicability of WMs in LLMs.
- We discuss limitations, challenges, barriers, and future directions for WMs in LLMs.

Our study provides key insights into using WMs for protecting LLMs, addressing concerns related to their deployment and IP protection. Through extensive experiments across different WMs, LLMs, and IP checkers, our highlighted findings are: (1) WMs significantly enhance the uniqueness and distinctiveness of LLM outputs, showing their effectiveness in strengthening IP protection and reducing unauthorized usage; (2) The impact of WMs on model utility can be moderate to significant, varying based on WM types and LLM architectures, potentially influencing their suitability for real-world applications; and (3) Attacks targeting WMs impose a notable cost on model utility. In conclusion, due to their unfavorable impacts on model utility, WMs in LLMs are not ready for real-world deployments, highlighting the urgent need to improve WM resilience while maintaining model utility.

2. Large Language Models and Risks

2.1. Large Language Models

Large language models (LLMs) [2], [31]–[33] are advanced deep learning models, pre-trained on vast datasets with a large number of parameters, enabling them to perform a wide range of natural language processing (NLP) tasks, such as text classification, text and code generation, and question answering.

Given a vocabulary \mathcal{V} consisting of words or word fragments, referred to as tokens, a sequence $x = \{x^{(i)}\}_{i=1}^T \in \mathcal{V}^T$ represents a *prompt* of T tokens, and $y = \{y^{(i)}\}_{i=T+1}^{T+K}$ denotes the tokens generated by an LLM in response to the prompt x (T and K are natural numbers). An LLM model θ is trained to maximize the probability of the token sequence y , conditioned on the prompt x , as follows:

$$P(y|x) = \prod_{i=1}^T P(y^{(i)}|x^{(1)}, x^{(2)}, \dots, x^{(i-1)}) \quad (1)$$

In generation tasks, the model predicts each token in an autoregressive manner [31], similar to Eq. 1, to generate a complete text sequence.

The most widely used LLM architectures are based on transformers [34] and incorporate a self-attention mechanism. This mechanism includes a stack of multi-head attention and feed-forward layers. Positional encoding with distinct matrices for different positions ensures the model retains sequence information without relying on recurrent networks. Based on transformer architectures, existing LLMs can be divided into three main categories (Fig. 1), as follows. (1) *Encoder-only models*, such as BERT [35], RoBERTa [36], XLNet [37], or XLM [38], capture word context and are well-suited for tasks that require a deep understanding of

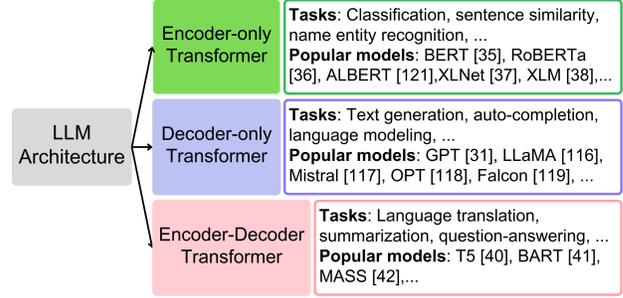


Figure 1: Categories of Large Language Models.

input data, such as sentence classification, sentiment analysis, and information extraction. These models are trained bidirectionally, predicting masked words using both preceding and following words. (2) *Decoder-only models*, such as GPT and its family [39], operate unidirectionally, predicting the next token based on context in an auto-regressive manner. They are designed for text generation tasks without a separate encoding phase. Lastly, (3) *Encoder-decoder models* such as T5 [40], BART [41] or MASS [42] combine both encoder and decoder components for tasks that involve understanding and generating data, such as text summarization and translation.

2.2. Risks in LLMs

Due to the high demands of computational resources and human efforts required for training LLMs [10], [11], LLMs are commonly offered as a service through APIs, typically requiring users to pay or subscribe [12], [13]. Table 6 (Appendix) presents a comprehensive overview of popular LLM APIs, highlighting their features, advantages, and constraints. According to the Artificial Analysis leaderboard [43], OpenAI and Microsoft are currently the leading providers in the LLM API market.

Although users may not have access to the model weights or architectures via APIs, this restriction does not guarantee the model’s safety. In addition, as LLMs become more widespread, the risk of their misuse for malicious purposes is increasing. Potential misuses of LLMs include several significant threats, as follows. First, LLMs can be exploited to generate false information or carry out harmful actions, such as crafting phishing emails, hate speech, or biased content to manipulate public opinion or amplify harmful narratives [44]. Second, LLMs can generate deepfake data, making manipulated media more convincing and enabling sophisticated scams or damage to reputation [45]. Third, adversaries can use LLMs to mimic individuals or organizations in text-based communications, facilitating theft, fraud, or financial crimes [46]. Fourth, LLMs can assist in writing malicious code or crafting evasion techniques to bypass detection systems, posing a significant cybersecurity threat [47]. More importantly, malicious actors may intentionally exploit LLM behaviors by imitating cloud-hosted models to offer cheaper services or engage in model-stealing attacks, further amplifying the aforementioned risks. Such attacks not only threaten the security of proprietary models but also expose users to fraudulent or unethical services [14], [15].

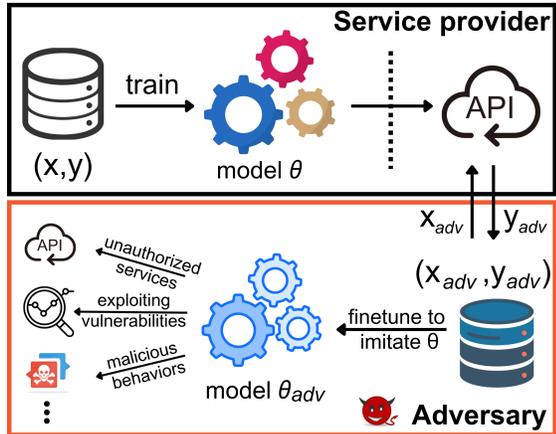


Figure 2: Model Stealing Attacks [48].

2.3. Model Stealing Attacks

In model stealing attacks, adversaries aim to imitate the behavior of a cloud-hosted model denoted as θ , by constructing their local model θ_{adv} through fine-tuning on input-output pairs obtained by querying the cloud-hosted model [49]. The goal is to bypass IP checkers while maintaining the adversary model’s utility, which is comparable to that of the original model. To launch the attacks (Fig. 2), adversaries start with a set of N task-specific prompts $\{x_i\}_{i=1}^N$, sending them to the cloud to obtain outputs $\{y_i\}_{i=1}^N$. These pairs are used to fine-tune θ_{adv} to mimic the behaviors of θ .

Model stealing attacks have various malicious and unethical uses, especially in LLMs. Common usages include 1) imitating proprietary models to resell or compete with the original cloud, undercutting legitimate providers and offering cheaper but unlicensed services, 2) bypassing service fees, 3) reversing engineering and exploiting vulnerability, as stolen models can be analyzed to identify weaknesses and potential vulnerabilities, and then exploit them, 4) violating privacy by inferring if sensitive data was used in training, 5) enabling malicious behaviors, such as spreading disinformation, biased content, harmful outputs, and 6) compromising IP checkers by copying the IP embedded in the model’s architecture, training data, and fine-tuning processes. These attacks pose a significant threat to the security, revenue, and ethical deployment of LLMs, underscoring the need for developers to implement robust defenses mitigating these risks.

3. Taxonomy of WMs and IP Checkers

To defend against model stealing attacks, recent works have highlighted the effectiveness and feasibility using WMs [9], [25], [50]. WMs embed a unique and detectable pattern into LLM outputs to trace and identify the origin of the text. A WM can be represented as a triple $(\mathcal{G}, \mathcal{W}, \mathcal{D})$, where \mathcal{G} is the WM generator, which takes a prompt x and a WM function \mathcal{W} to produce a watermarked output $y^{wm} = \mathcal{G}(x, \mathcal{W})$, and \mathcal{D} is a WM detection function, also known as an IP checker, to determine whether a given text is watermarked or not.

3.1. WM Generators

WM generators typically introduce perturbations to the probability distribution of LLMs during output generation. Depending on \mathcal{W} and types of perturbations, WMs can be classified into several research directions, as shown in Fig. 3. *First*, WM during logits generation approaches partition tokens into green and red lists and add bias into logits to favor tokens from the green list. The tokens in the watermarked output are thus sampled with higher probability from the green list, which is determined by a pseudo-random generator seeded by the input prompt [9]. These methods are flexible and effective for tracking watermarked outputs but may impact the semantic meaning of the output and are vulnerable to removal attacks [26]. *Second*, WM during sampling process methods alter LLM text generation’s sampling process of next tokens or whole sentences without altering the underlying distribution of output. [25], [77]. While they are easy to integrate and allow for simple detection, they can be susceptible to attacks such as text modifications, randomness, and sentence reordering, especially in token-based methods. *Third*, WM during LLM training methods embed WMs by training LLMs with specific triggers, knowledge distillation, reinforcement learning, or by modifying the LLM architecture [79], [80], [84], [85], [87]. These methods offer inherent robustness but are often associated with implementation complexity and increased computational overheads. *In addition*, several miscellaneous WM approaches, such as mixed methods, multiple-output generation with selection, or prompt-based techniques, offer enhanced robustness and adaptability [88]–[91]. However, these methods come with trade-offs, including higher computational costs and potential vulnerabilities to watermark removal attacks.

3.2. WM IP Checkers

After watermarking the LLM outputs, the clouds can use IP checkers to determine whether an adversary model has utilized watermarked outputs to replicate the behavior of the cloud-hosted model. This is typically done by checking whether the generated text from the model is watermarked. IP checkers are generally divided into two main approaches. First, some IP checkers categorize each token as green or red and then apply statistical tests, such as z-scores, p-values, or Jensen-Shannon divergence, to assess whether the proportion of green tokens exceeds a specified threshold. If this threshold is met, the text is flagged as watermarked [9], [50], [59]. Second, other IP checkers assess the alignment between generated tokens and pseudo-random sequence, checking whether each token matches the corresponding value in the sequence. Depending on a WM technique, the cloud will apply its specific detection function to assess the output. If the statistical tests or alignment scores exceed certain thresholds, the output is considered watermarked. If the proportion of detected watermarks is sufficiently high, the model may be flagged as potentially stolen.

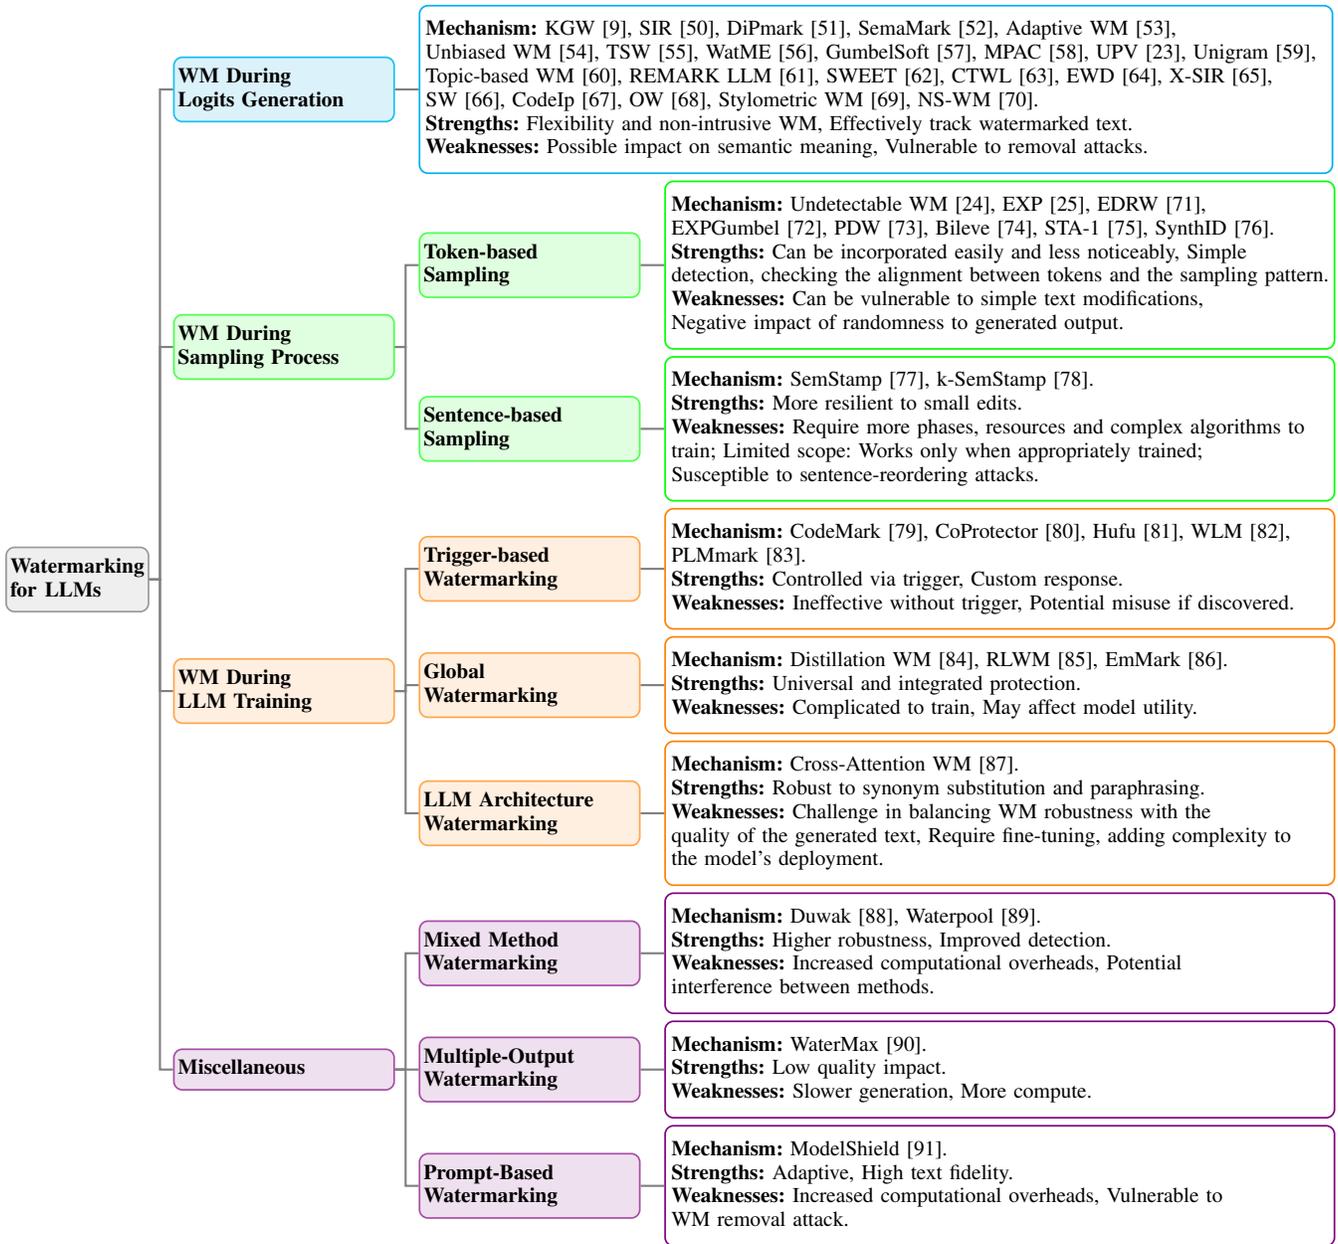


Figure 3: Taxonomy of Watermarking Mechanisms for LLMs.

3.3. Characteristics of an Effective WM

A variety of WMs in LLMs have been introduced, while effective WMs in LLMs share key characteristics: *First*, WMs should minimally impact text quality, ensuring they are not noticeable to users. Quality can be assessed using metrics like perplexity, BLEU score, or accuracy on tasks like MMLU or sentiment classification [92]. *Second*, WMs must be easily detectable by IP checkers, enabling reliable origin tracing with high detection rates. *Third*, WMs should be robust to WM removal attacks that attempt to remove watermarked patterns to bypass detection. By satisfying these characteristics, WMs in LLMs can effectively balance security, model

utility, and usability in real-world deployments.

3.4. Attacks against WMs

WM Removal Attacks [26], [27], [93], [94]. The attacker's goal is to remove embedded signatures and patterns from a watermarked text using paraphrasing techniques such as synonym replacement, sentence reordering, or sentence matching to minimize detectable traces of the WM. These attacks aim to generate semantically equivalent but syntactically modified text to evade WM detection mechanisms. In addition, some attacks [26] propose leveraging a high-quality oracle model to preserve output quality while removing the

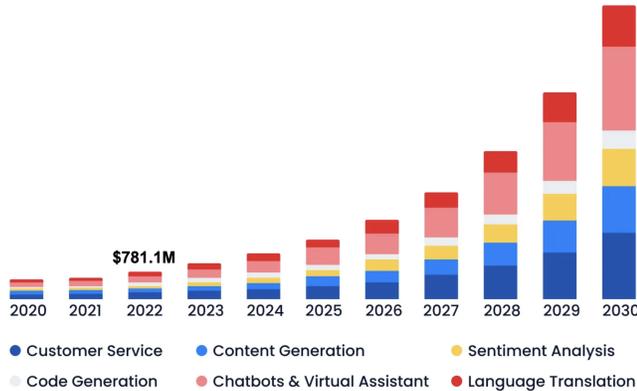


Figure 4: US LLM Market Trend from 2020 to 2030 [97].

WM, making detection even more challenging. However, these approaches often require additional training to effectively rephrase sentences or entire paragraphs, which can introduce subtle shifts in semantic meaning, degrade text quality, and significantly increase computational complexity.

Spoofing Attacks [27], [48], [84], [95], [96]. The attacker aims to produce text that is falsely detected as watermarked, thereby nullifying the value of WMs [27]. These attacks undermine the integrity of the WM by making it difficult to distinguish between legitimate watermarked content and counterfeited outputs. These approaches may harm the reputation of model providers, for example, if inappropriate or harmful text contents are falsely attributed to them. However, these attacks work only with certain types of WMs, and to be successful, they typically require a large number of queries to collect sufficient watermarked text for estimating the WM’s distribution and reconstructing the WM rules, associating with computational overheads.

4. Exploring Effectiveness of WMs in LLMs

In this section, we focus on addressing the key question of this study: *Are WMs in LLMs Ready for Deployment?* To provide a comprehensive answer, we first investigate the integration of WMs in LLM productions and discuss the existing and potential challenges of adapting them for practical use. In addition, we conduct an in-depth evaluation of the effectiveness of WMs across different LLMs by introducing a novel cross-model IP classifier and assessing how WMs impact the model utility of LLMs and their downstream tasks. By analyzing these aspects, we seek to determine whether WMs can be implemented without noticeably compromising the model utility of LLMs, which is a critical consideration in determining whether WMs are genuinely ready for widespread deployment.

4.1. WMs in LLM Productions and Challenges

LLMs have gained significant attention recently, with leading companies developing and deploying them in their products and services. Table 1 illustrates real-world applications of various LLMs across industries. It highlights

how these companies have integrated AI technologies like OpenAI’s GPT, Microsoft’s Azure, Meta’s LLaMA, Google Cloud AI, Amazon Comprehend, and Anthropic’s Claude to enhance their services. These examples demonstrate the broad impact of LLMs in areas such as customer service, content optimization, document processing, and software development, driving innovation and efficiency across sectors. In addition, as shown in Fig. 4, the market for LLM production grew from 2020 to 2022, reaching 781.1 million, with substantial growth expected by 2030, highlighting the increasing importance of LLMs across various industries.

Given the increasing trend in the LLM market, along with potential risks associated with their deployment in real-world scenarios, as discussed in Section 2, it is essential to utilize WMs to protect IP of LLM outputs. However, the adoption of WMs in LLM productions remains limited. For instance, OpenAI has developed a WM system for ChatGPT but has opted not to implement it due to concerns over ambiguous potential penalties for users and feedback suggesting that service usage might decline if WMs were applied [116]. The uncertainty from providers and users around adopting WMs in LLMs stems from a lack of understanding of WM effectiveness and how different LLMs interact with IP checkers for a given WM.

However, uncovering the true effectiveness of WMs in LLMs and such interactions presents several challenges. *First*, while LLMs and WMs are widely used, the impact of WMs on output quality and IP protection remains unclear, particularly regarding how WMs affect model utility and detection by IP checkers. *Second*, when a cloud service receives outputs from different LLMs, both with and without WMs, how the cloud classifies which model generated the output and how WMs support this classification. *Third*, it is unclear which WM should be applied to a given LLM and why. Companies need clear and evidence-based guidelines for selecting the most effective WM strategies to ensure optimal performance and user experience. *Fourth*, there is a lack of standardized metrics to evaluate how WMs impact LLM utility, making it difficult to assess their effectiveness across different LLMs and tasks. *Finally*, it is essential to assess model utility and the quality of LLM outputs on downstream tasks after watermarking to ensure that WMs do not degrade performance. Addressing these challenges is vital for successfully deploying WMs in real-world LLM applications, ensuring a balance between IP protection and model utility.

4.2. Native IP Classifier

To understand the effectiveness of WMs on LLMs in terms of model utility and IP checker performance, a straightforward approach is to apply the specific IP checker tailored to the WM on the LLM and compare the utility and quality of the text before and after being watermarked. The key idea behind the IP checker is to evaluate the statistical differences between original and watermarked outputs, flagging the output as watermarked if the differences

TABLE 1: Real-world Services and Applications of LLMs.

API Name	Provider	Real-world Services and Applications
ChatGPT and GPT family [2]	OpenAI	Salesforce integrated its AI with OpenAI for CRM tools [98]. Duolingo introduced Duolingo Max powered by GPT-4 for better learning experience [99]. Instacart added ChatGPT to its Ask Instacart chatbot for better customer interaction [100]. Expedia powered its chatbots by ChatGPT for comprehensive services [101]. Morgan Stanley introduced AskResearchGPT for staff support [102]. Stripe uses GPT-4 to improve user experience and fight fraud [103].
Microsoft Azure Language [32]	Microsoft Azure	AT&T leveraged Azure OpenAI to automate processes and improve customer experiences [32]. Vodafone improved services with TOBi AI for personalized, multi-channel support [104]. Volvo Group used AI for document translation to streamline invoice & claims processing [104]. TotalEnergies launched Microsoft Copilot for an AI chat solution using internal data [105]. Providence created ProvARIA based on Azure OpenAI to direct patient care messages [105]. Grupo Bimbo developed a copilot for employee queries on policies and risk management [105].
LLaMA family [33]	Meta	Zoom utilized LLaMA-2 to create its AI Companion for meeting summaries, next-step highlights, and presentation tips [106]. Mathpresso built MathGPT with LLaMA for personalized math learning [106]. Goldman Sachs used LLaMA to improve service, document review, code generation [106]. Nomura Holdings used LLaMA to enhance customer service and document analysis [106]. EPFL developed Meditron with LLaMA-2 to aid clinical decision-making and diagnosis [106].
Google Cloud AI-Language (Gemini, Bard, PaLM, etc.) [107]	Google	Lilt integrated Google Translate into its platform for advanced translation capabilities [108]. Wipro with Google Cloud created the Gemini Experience Zone for AI experimentation [109]. General Motors deployed a virtual assistant using Google Cloud’s conversational AI for better intent recognition [110]. Gojek launched ”Dira by GoTo AI”, an voice assistant for GoPay service [110]. GroupBy built a Search and Discovery Platform using Vertex AI to boost retail sales [110]. Motorola AI leveraged Gemini for conversation summaries and search [110].
Amazon Comprehend [111]	Amazon AWS	Vibes used the platform for phrase extraction, sentiment analysis, and topic modeling [112]. Gallup applied Targeted Sentiment to improve survey response reporting [112]. Zillow built speech analytics to enhance customer and service support [112]. Siemens leveraged AWS LLMs to automate its surveying program [112].
IBM Watson NLU [113]	IBM Watson	SherloQ used IBM Watson NLU to analyze search behavior and optimize ads [114]. University of Auckland deployed IBM Watson Assistant based on a conversational AI solution for handling repetitive questions [114]. Mushi Lab created Clearscope to optimize content with insights & recommendations [114]. Global-Regulation built a global law search engine for easier legal information access [114].
Claude [3]	Anthropic	Notion integrated Claude for content generation, summarization, and workflow automation [115]. Decagon utilized Claude for personalized, 24/7 customer service [115]. GitLab employed Claude to enhance coding assistance & software development [115]. Headstart leveraged Claude to accelerate their software development process [115].

surpass a set threshold. In other words, the IP checker acts as a classifier to distinguish between the two outputs.

Despite its high performance, applying a native IP checker independently to each WM and LLM may be insufficient as it may not effectively capture the true impact and effectiveness of the WM. In practice, a cloud provider may not know the exact source of the output, as multiple LLMs could be involved. Watermarked outputs from one LLM may closely resemble the original outputs from another, due to distribution shifts caused by the WM (Fig. 5). Therefore, applying the IP checker in isolation to each LLM may yield misleading results. To more accurately assess WM effectiveness, it is complemented to adopt a broader approach that evaluates the WM across multiple LLMs. By introducing a cross-model IP classifier, we can more reliably

differentiate between the original and watermarked texts across different LLMs. This approach ensures a cross-model IP classifier should be used. This approach ensures reliable differentiation between original and watermarked texts across multiple LLMs, better reflecting real-world conditions.

4.3. Proposed Cross-Model IP Classifier

In this section, we present the following components: 1) a novel cross-model IP classifier to systematically discover the interplay between WMs and different LLMs, 2) a selection of evaluation metrics for IP assessment, 3) the impact of WMs on LLMs’ utility, and 4) based on comprehensive experiments, we provide an insightful understanding of advantages and constraints of different WMs in LLMs.

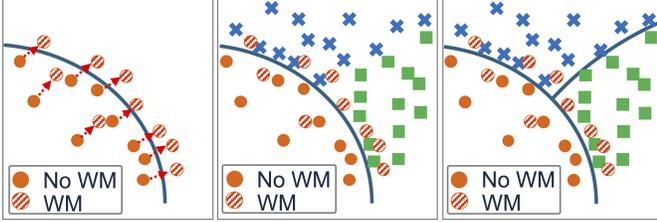


Figure 5: Native IP checker works well in simple scenarios (left), but struggles to differentiate between outputs from multiple LLMs (center), highlighting the need for a cross-model IP checker (right).

The key concept behind our cross-model IP classifier is to distinguish between the output distributions of different LLMs, both with and without WMs, reducing ambiguity in the IP classifier’s decision boundary. This classifier systematically connect WMs and different LLMs through a multi-class model functioning as an IP checker.

Given a set of N LLM models $\{LLM_1, LLM_2, \dots, LLM_N\}$ and an input text $x \in D$, we generate two sets of outputs, including outputs y without WMs and outputs y^{wm} using WMs. The IP classifier is designed to assess whether x belongs to which LLM. In our cross-model classifier, we consider three scenarios. *First*, we use the outputs y_i from all the LLM_i ($i \in [1, N]$) as inputs for the classifier f (Fig. 6a). This setting discloses how different LLMs distinguish their outputs compared with others, highlighting the uniqueness of each LLM system. *Second*, we apply WMs into one LLM at a time to thoroughly investigate the relationship between each LLM and WM without being affected by the results of other LLMs (Fig. 6b). *Third*, we unify the evaluations to uncover the overall correlation among WMs and LLMs by applying each WM to all LLMs (Fig. 6c).

Next, we will conduct experiments to evaluate our proposed method in assessing the effectiveness of WMs on LLMs, analyze the results, and discuss the limitations of existing approaches, along with potential solutions.

4.4. Experiment Settings

4.4.1. Baselines

We carry out the validation on a wide range of WMs, LLMs, and IP classifiers. *For LLMs*, we use four state-of-the-art (SOTA) pre-trained LLMs, including LLaMA-2 with 7B parameters [117], Mistral with 7B parameters [118], OPT with 6.7B parameters [119] and Falcon with 7B parameters [120]. *For WMs*, we use seven SOTA WM schemes, including 1) KGW [9], 2) EXP [25], 3) SIR [50], 4) SemStamp [77], 5) Unigram [59], 6) Adaptive [53], and 7) UPV [23]. These WMs are representative WM schemes outlined in Section 3, with further details provided in Table 7 (Appendix).

For IP classifiers, we explore a range of options, including 1) a non-transformer LSTM model [121]; 2) an encoder-based transformer model *ALBERT* [122], with 11.8M parameters; 3) another larger encoder-based transformer

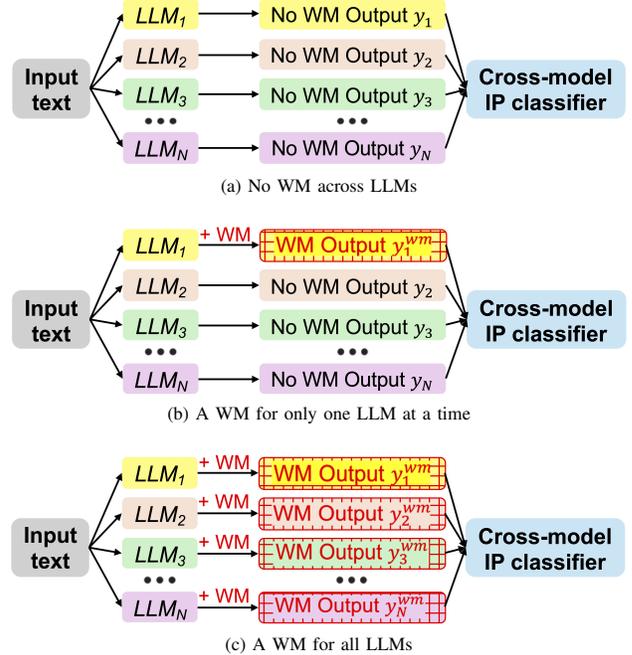


Figure 6: Our Proposed Cross-model IP Classifier.

DeBERTa [123] with 86M parameters; 4) a decoder-based transformer *GPT-2* [31] with 125M parameters; 5) a decoder-only model *Galactica* [124] with 125M parameters; 6) a small version of *OPT* [119] with 125M parameters; 7) the recently developed *SmolLM* [125] with 135M parameters; and 8) an encoder-decoder model *T5* [40], with 60.5M parameters. These models were selected based on the following criteria: 1) inclusion of non-transformer and transformer-based architectures, 2) within transformer-based models, coverage of all major categories outlined in Section 2, and 3) consideration of model complexity as determined by parameter size.

To further assess the effectiveness of WMs, we test them against SOTA attacks, including 1) a WM removal attack [26], referred to as *WMremoval*, which removes WMs by perturbing quality-preserving WM outputs to produce high-quality non-WMed outputs; 2) *Dipper* [93], which paraphrases paragraphs through content reordering and lexical changes without appreciably modifying the text semantics, and 3) *Substitution* attack [94], which randomly substitutes target words with their synonyms, considering the surrounding context and using WordNet [126]. We adopt the default settings reported in these papers.

4.4.2. Datasets and Evaluation Tasks

We randomly select 10,000 training samples and 2,000 test samples from the benchmark C4 dataset [127]. Each sample, a news-like string, is truncated to 200 tokens as a prompt, with the next 200 tokens used as the “baseline” completion. Tasks include text generation as the primary task and massive multi-task language understanding (MMLU) [92] as a downstream task, assessing the impact of WMs on LLM utility. MMLU has 57 subjects in STEM, humanities,

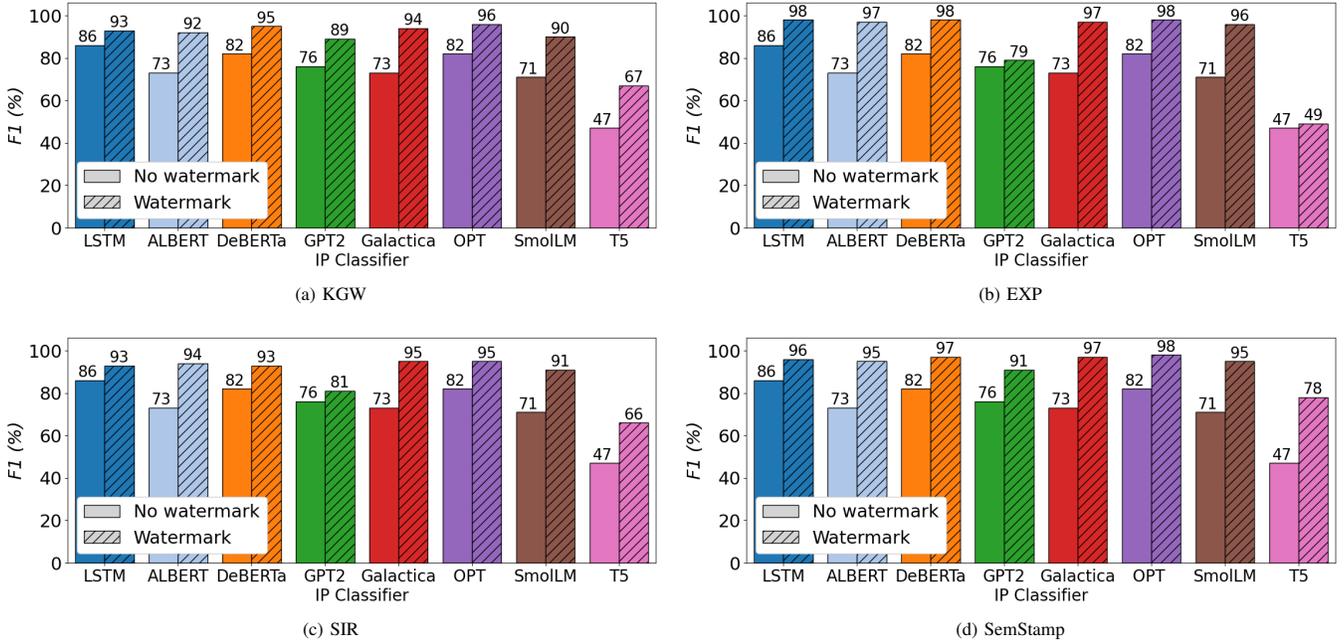


Figure 7: $F1$ scores as a function of different IP classifiers for different WMs, compared with No WM settings.

and social sciences, covering various difficulty levels to evaluate world knowledge and problem-solving skills.

4.4.3. Evaluation Metrics

We treat the IP evaluation across different LLMs as a multi-class classification problem, using the $F1$ score for IP classifiers. To measure the impacts of WMs, we calculate the change in $F1$ score, denoted as $F1_{change}$, by computing a relative change in $F1$ between scenarios with and without WMs. The metrics are calculated as follow:

$$F1 = 2 \times \frac{precision \times recall}{precision + recall} \quad (2)$$

$$F1_{change} = \frac{F1^{WM} - F1^{NW}}{F1^{NW}} \quad (3)$$

where $precision$ is defined as the number of true positive results divided by the total number of all samples predicted by the model, including incorrect predictions, and $recall$ is the number of true positive results divided by the total number of all actual instances. WM and NW indicates the score from WM and no WM settings, respectively.

In addition, to evaluate the model utility comprehensively, we consider 1) Perplexity for the text generation task, measuring how well a model predicts next tokens [128] and 2) Average accuracy across topics for the MMLU downstream task. To further assess text quality, we provide qualitative examples comparing a prompt, the original output without WMs, and corresponding outputs with different WMs to offer insights into WM effectiveness.

4.5. Experimental Results

Under our IP classifier, an LLM is considered effective if it achieves a high $F1$ score (derived from settings a and c in Fig. 6, where all LLMs are under the same condition of either with or without watermark), indicating the uniqueness of the LLM’s outputs and enabling easy differentiation from other LLMs. This capability also allows cloud providers to determine whether an output stems from their model, thereby effectively protecting their IP. Meanwhile, a WM scheme is deemed effective if (1) $F1_{change}$, derived from setting b in Fig. 6, is significantly high, facilitating better detection compared to a scenario where no WM is applied, and (2) it has minimal impact on the model’s utility of the main task (i.e., maintaining low perplexity in text generation) and of downstream tasks (i.e., high average accuracy of MMLU). This ensures that WMs enhance the uniqueness of the LLM’s outputs without adversely affecting the model’s overall utility.

4.5.1. Uniqueness of LLM Outputs

Figs. 7 and 12 (Appendix) illustrate $F1$ scores of different IP classifiers in distinguishing outputs from different LLMs. Without WMs, these IP classifiers achieve impressive $F1$ scores of over 70% across various types. The LSTM outperforms others, achieving an $F1$ score of 86%. These results underscore the uniqueness of LLM outputs, allowing for reliable identification of the originating LLM based on the observed outputs. Notably, the T5 IP classifier performs poorly without watermarks, achieving only 47%. Indeed, T5 is an encoder-decoder model that is more suitable for tasks such as translation and completion. In Fig. 8, the poor results stem from insufficient training data. With

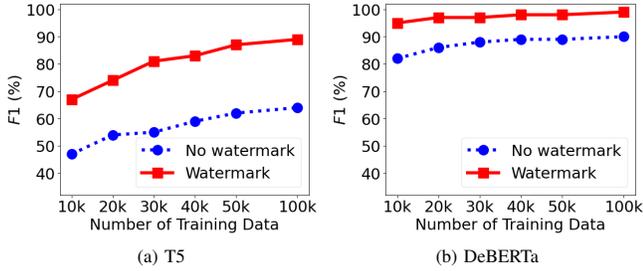


Figure 8: $F1$ score as a function of number of training data applied on KGW and LLaMA-2.

increased numbers of training samples, the classifier’s $F1$ score improves, reaching 64% with 100,000 training samples. Similar improvements are observed in other models like DeBERTa, with an 8% increase as training samples grow.

4.5.2. Effectiveness of WMs on LLMs

Impacts of WMs on Classifying LLM Outputs. As shown in Figs. 7 and 8, WMs can significantly boost the performance of LLM output classification. We observe this phenomenon across WMs, LLMs, and IP classifiers. There is a significant improvement of 2% – 31%, compared with the performance without WMs. In particular, ALBERT, T5, Galactica, and SmoLLM classifiers demonstrate substantial improvements when WMs are applied, with an increase of 19% – 31% in $F1$ scores. It is worth noting that all the classifiers achieve over 90% $F1$ score with WMs. Notably, WM with EXP or Semstamp allows the OPT classifiers to reach an impressive 98% $F1$ score. The enhanced performance in distinguishing LLM outputs is due to the unique patterns or identifiers created by WMs, which increase the distinctiveness of each model’s output, making it easier to identify the source LLM. Even with the T5 classifier, which performs poorly on non-watermarked text, shows significant improvement with WMs, increasing from 49% to 78%.

The remarkable improvements in IP classifier performance highlight the effectiveness of WMs in differentiating outputs from various LLMs. This crucial correlation between LLMs and WMs allows clouds to identify their outputs accurately, ensuring that proprietary content is protected. By leveraging WMs, clouds can enhance IP protection for their models and reduce unauthorized use. This capability is essential in competitive environments such as industry, where protecting unique outputs is critical for business success.

Impacts of LLMs on IP Classifiers. Fig. 9 demonstrates the impact of different LLMs on IP classifiers. These results are obtained by applying WMs to one LLM at a time, with the relative change $F1_{change}$ indicating the impacts. When WMs are applied, all LLMs show positive values of $F1_{change}$, signifying an increase in their $F1$ scores. In this analysis, we present representative trends observed in the ALBERT and DeBERTa classifiers. Both classifiers demonstrate a consistent pattern, showing more significant changes for the OPT and Falcon models, compared with the LLaMA-2 and Mistral. This variation is attributed to LLaMA-2 and

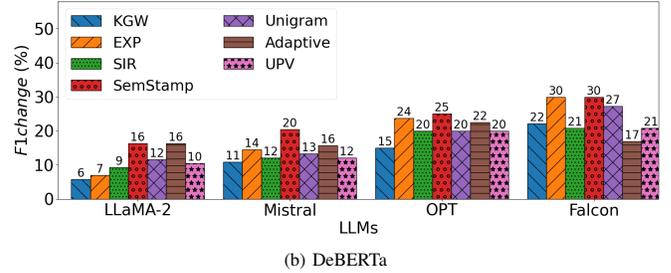
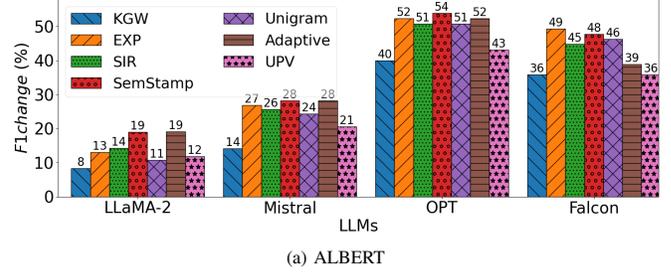


Figure 9: $F1_{change}$ across WMs, LLMs, and IP classifiers.

Mistral’s adaptability and robustness, which make them more resistant to the effects of WMs. In contrast, OPT and Falcon are specifically designed for efficiency and scalability.

Impacts of WMs on Model Utility of Text Generation Tasks. In addition to improving IP classifiers’ ability to distinguish LLM outputs, it is crucial to assess how different WMs impact model utility, as this is a key characteristic of effective WMs. As shown in Table 2, WMs lead to a noticeable increase in the perplexity of the generated text, indicating a negative effect on the model’s utility in text generation. KGW has the lowest impact on perplexity values across all LLMs, while SemStamp exhibits substantial changes in perplexity. For instance, with the OPT model, KGW slightly changes perplexity from 3.55 to 3.87. In contrast, SemStamp significantly increases perplexity to 76.91 in this setting. The key reason for this difference lies in how each WM operates. To elaborate, KGW applies the WM during the logits generation process, modifying the text without significantly altering the model, which minimizes its impact on the LLMs. Meanwhile, SemStamp influences the sampling process by generating new sentences from a partitioned valid semantic space, resulting in a more noticeable effect on the generation process, compared with other baselines. It is worth noting that all types of WMs produce considerable impact on this text quality, urging for further investigation.

Impacts of WMs on Model Utility of Downstream Tasks.

We further conduct an in-depth analysis in the MMLU question-answering task (Table 3). Our observations reveal that majority of WMs have subtle impact on the average accuracy of MMLU. For instance, there are up to 2.3% decreases in the accuracy across LLMs and WMs such as KGW, EXP, Unigram, and UPV compared with the outputs without WMs, indicating that watermarked LLM outputs still retain usefulness for classification tasks. However, WMs

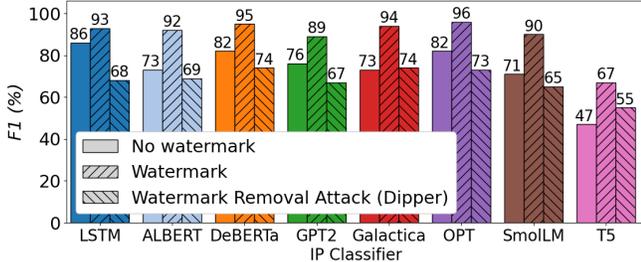


Figure 10: $F1$ scores as a function of different IP classifiers for KGW WMs, compared with No WM and after WM Removal Attack of Dipper.

like SIR and Adaptive show more significant accuracy drops across all four LLMs. The primary reason for these poor results lies in our reliance on their publicly available pre-trained models (such as the semantic mapping model for Adaptive and the watermarked model for SIR) without fine-tuning them for our specific tasks. Limited time and computational resources prevented us from aligning these models with our experimental settings, which likely affected their performance. Additionally, SemStamp, which involves sentence-level sampling, is not ideal for MMLU tasks, where token-level generation is preferred.

WMs against Attacks. We further examine how WMs resist against different attacks, indicating the effectiveness and resilience of WMs on LLMs. We consider WM removal attack, Dipper, and substitution attacks. Fig. 11 shows that these attacks achieve high attack success rates, which indicates a low detection rate of the WM’s detection function. However, it comes with the cost of a significant impact on model utility, as evidenced by a substantial increase in perplexity. For instance, WM removal attacks can raise the perplexity of a non-watermark baseline by 2 to 6 times. In the case of the Mistral model with Semstamp, perplexity increases dramatically from 39.31 up to 241. In addition, Fig. 10 demonstrates the impact of the Dipper attack on WM effectiveness in classifying LLM outputs. The results show a significant performance drop, highlighting the detrimental effects of the attack on both the effectiveness of WMs and the quality of LLM outputs.

4.5.3. Semantic Preservation of Watermarked Outputs

Furthermore, to qualitatively assess the quality of generated text under different WMs, we illustrate side-by-side examples of real prompts, original outputs from an LLaMA-2 model without WMs and with different WMs.

In Table 4, we observe that watermarked text preserves the semantic meaning compared with the original outputs without WMs. Notably, KGW maintains high semantic similarities, as the green-highlighted phrases show. For instance, key phrases such as “are suing the college,” “medical laboratory technician program,” and “filed a class-action lawsuit” are preserved almost identically in both the KGW output and the original, demonstrating KGW’s ability to maintain the integrity of the original semantic content. This

TABLE 2: Perplexity values of different WMs across LLMs.

Setting	LLaMA-2	Mistral	OPT	Falcon	Average
NW	3.61	3.54	3.55	2.68	3.35
KGW	4.46 (24%)	4.43 (25.1%)	3.87 (9.0%)	3.37 (25.7%)	4.03 (20.6%)
EXP	7.06 (96%)	33.61 (849.4%)	21.10 (494.4%)	18.72 (598.5%)	20.12 (501.6%)
SIR	7.05 (95%)	11.00 (210.7%)	19.72 (455.5%)	13.89 (418.3%)	12.92 (286.1%)
SemStamp	24.06 (566%)	38.89 (998.6%)	76.91 (2066.5%)	7.52 (180.6%)	36.85 (1001.5%)
Unigram	5.50 (52%)	8.68 (145.2%)	13.71 (286.1%)	10.71 (299.7%)	9.65 (188.4%)
Adaptive	15.77 (337%)	12.97 (266.4%)	11.24 (216.5%)	5.47 (104.0%)	11.36 (239.6%)
UPV	5.81 (61%)	4.39 (24.1%)	4.69 (32.2%)	3.45 (28.9%)	4.59 (37.1%)

TABLE 3: MMLU Accuracy (%) of different LLMs with and without WMs.

Settings	LLaMA-2	Mistral	OPT	Falcon
No WM	46.7	58.9	24.8	27.2
KGW	44.5	58.6	24.5	25.3
EXP	45.9	58.7	24.8	24.9
SIR	27.9	20.9	18.3	19.8
Unigram	46.6	58.1	23.8	25.3
Adaptive	31.4	37.9	24.7	26.1
UPV	46.7	57.9	24.7	25.6

is further supported by the low perplexity of 2.75, close to the non-WM perplexity of 2.32, indicating minimal degradation. Meanwhile, SemStamp and Adaptive changes sentences more remarkably, resulting in a significant deviation from the original output. For instance, in Semstamp, the phrase “are suing the college” is entirely replaced by “alleging they were misled about their ability to become certified”, modifying the sentence’s meaning. While Adaptive introduced a new phrase of “have filed a suit claiming that Newbridege made “materially incorrect” claims”, causing the output to deviate from the original text. In addition, we observe a tiny portion of green highlights in these two WMs, indicating substantial dissimilarities in the outputs. Also, the perplexity of the SemStamp and Adaptive WM are considerably higher at 10.61 and 17.62 respectively, reflecting the degradation in model utility and semantic consistency.

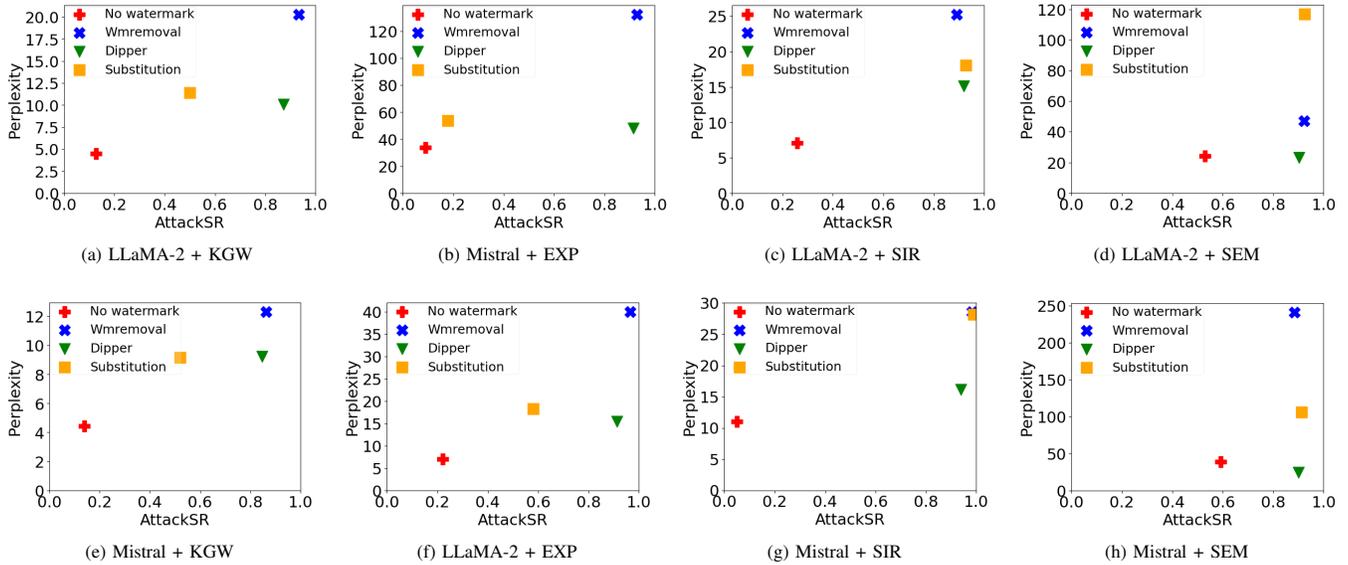


Figure 11: Attack success rate and Perplexity across different attacks and WMs with LLaMA-2 and Mistral models.

4.6. Remarks

Through extensive experiments, we observe several key insights as follows.

- WMs significantly enhance the classification of LLM outputs, supporting the generation of unique and differentiated outputs. This enables clouds to identify their outputs and protect proprietary content accurately. By using WMs, clouds can strengthen IP protection and reduce unauthorized use, which is vital in competitive industries.
- WMs typically have a moderate impact on model utility, subtly increasing perplexity in text generation tasks while reducing accuracy in MMLU downstream tasks. The semantic meaning of outputs is mostly preserved after WM, with significant changes occurring mostly in SemStamp and Adaptive. However, it is worth noting that LLM outputs remain beneficial even after being watermarked.
- Attacks adversely affect WMs by degrading the WMs’ detection function, compromising WMs’ effectiveness. However, it comes with a considerable cost to model utility. Further investigation is necessary to understand the trade-offs involved and to develop strategies that enhance WM resilience while maintaining model utility.

5. Challenges and Future Directions

While WMs demonstrate promising effectiveness across various LLMs, several challenges and barriers remain, necessitating further investigation before deploying them in the real-world applications.

5.1. Different Sensitivities among Output Tokens

Our experiments reveal that tokens exhibit varying levels of sensitivity, especially after being watermarked. As shown in Table 5, numerical tokens are changed significantly compared with other tokens. For instance, the phrase “Nasdaq dropped 0.7 percent” in the original output is altered to “Nasdaq lost 0.6 percent” in the EXP, SIR and Adaptive, to “Nasdaq lost 1.2 percent” in Unigram, and to “Nasdaq lost 1.4 percent” in UPV. Similarly, the original output states that the U.S. economy grew at a “2.5 percent annual rate in the second quarter”, which is shifted to “1 percent” or “3.5 percent” with WMs. These unintended modifications to numerical tokens result in changing the semantic meaning of the output. This change poses significant risks in domains where precise values are important, including financial reporting, scientific data analysis, weather forecasting, and legal documentation. As a result, there is a pressing need to address these challenges to ensure that WMs do not compromise the integrity of essential information.

5.2. Robustness to Removal Attacks

WMs, regardless of their category, generally face challenges related to robustness. Given the complex manipulation behaviors of users and adversaries in the real-world, it is difficult to create a WM that can resist to modifications such as paraphrasing, lexical changes, or reordering. As also claimed in [26], current strong WMs fail to offer the desired level of security, especially against sophisticated attacks.

5.3. Text Quality and Semantic Drift

Embedding WMs in LLMs, whether by modifying logits, sampling process, or updating model weights, disrupts the

TABLE 4: Selected prompt and outputs from the LLaMA-2 model without and with WMs. Green text indicates semantic similarities of WM text and the original output without WMs.

	Output text	PPL
Prompt	...medical laboratory technician associate program. FILING SUIT: Attorney Scott Schutzman discusses a class-action lawsuit against Newbridge College with two of his clients. A group of students...	
No WM	are suing the college, claiming its medical laboratory technician program is not accredited and does not qualify them to work in the field. SANTA ANA – A group of students have filed a class-action lawsuit against Newbridge College, claiming the college’s medical laboratory technician program is [...continues]	2.32
KGW	are suing the college for allegedly misrepresenting the quality of its medical laboratory technician program. Newbridge College faces lawsuit over alleged false advertising of medical lab technician program SANTA ANA – A group of students have filed a class-action lawsuit against Newbridge College, alleging the college misrepresented [...continues]	2.75
EXP	sued the college claims the school’s program does not qualify them to be medical lab technicians. STUDENTS: ATTORNEY SCOTT SCHUTZMAN: Two of the Newbridge College students who filed a class-action lawsuit against the school are pictured with their attorney, Scott Schutzman. [...continues]	4.43
SIR	sued the college for false advertising after they say the college’s medical laboratory technologist program does not prepare them for the job market. CENTER: Newbridge College’s Santa Ana campus is located at 2840 East 17Th Street. Students who attended the college have filed a class- action lawsuit against the college alleging false advertising.SANTA ANA - Students who enrolled in Newbridge College [...continues]	5.11
SemStamp	alleging they were misled about their ability to become certified as medial labs techs through the school’s programs and have filed suit for damages. (File Photo) Newport Beach lawyer files complaint on behalf of former NBCC student who claims she was defrauded out of thousands of dollars due to false promises made by institution regarding job placement rate [...continues]	10.61
Unigram	sued the college for allegedly making false claims about its medical lab technician program. STUDENTS: Attorney Scott Schutzman represents two students, Alexis Mesa and John Doe, who sued Newbridge College for allegedly making false claims about its medical lab technician program. [...continues]	4.11
Adaptive	have filed a suit claiming that Newbridge made “materially incorrect” claims about it programs, such as saying it offered a medical lab technician program that qualified them to be medical lab techs. Atty: New Bridge falsely advertised program. By Joel Kimber — Orange County Register. A Newport Beach lawyer has filed a proposed class-action lawsuit on behalf of students who say they were duped by Newbridge College [...continues]	17.62
UPV	have sued the college for fraud and false advertising, claiming the college’s medical laboratory technology program does not prepare them for jobs in the field. SCHUTZMAN: Attorney Scott Schutzman represents the students who have sued Newbridge College, accusing the college of fraud and false advertising. [...continues]	4.27

natural flow of text generation. In Table 2, even KGW, which has the smallest impact on model utility, still alters the original perplexity by 9 – 25%. This disruption adversely affects the naturalness and fluency of the watermarked output, leading to a overall reduction in model utility and text quality, compared with those of original outputs. In sensitive contexts, semantic drift can further compromise the credibility and usefulness of LLM outputs. Given the inevitability of the issue, it is essential to develop methods for quantifying the impact and managing it in acceptable thresholds.

5.4. Scalability and Practicality

The scalability and practicality of WMs for LLMs pose another significant challenge. WMs often require modifica-

tions to the model’s generation process, which might cause higher processing time and resource consumption. In addition, some WMs require additional training to function properly. As LLMs are used on a broad scale, particularly in real-time applications such as chatbots or content generating systems, the increased computational overhead may become prohibitive. In addition, ensuring that WMs are robust across various domains, tasks, and languages complicates their practical deployment. Scaling these methods to handle large volumes of generated text while maintaining text quality and performance remains an open challenge, especially given the numerous use cases of LLMs across platforms. Future advancements should address these difficulties to make WMs more realistic.

TABLE 5: This example demonstrates the sensitivity of numerical tokens in LLaMA-2 outputs without and with different types of WMs. Yellow highlights represent key numerical information in the output without WMs, while green highlights indicate similarities, and red ones are for differences compared to no WM setting.

	Content
Prompt	...For the week, the Dow fell 0.8 percent, the S&P 500 slipped 0.9 percent and the ...
No WM	Nasdaq dropped 0.7 percent. The report from the Commerce Department showed that the U.S. economy grew at a slower-than-expected 2.5 percent annual rate in the second quarter, down from a 3.4 percent pace in the first quarter. [...continues]
KGW	Nasdaq lost 0.7 percent. The economic report from the Commerce Department showed that the U.S. economy grew at a slower-than-expected 3.5 percent annual rate in the first quarter, down from a previously estimated 4.4 percent pace. [...continues]
EXP	Nasdaq lost 0.6 percent. The U.S. economy grew at a slower-than-expected 1.9 percent annual rate in the first quarter, the Commerce Department reported. Economists had been expecting a rate of 2.2 percent. [...continues]
SIR	Nasdaq lost 0.6 percent. The Commerce Department said Friday that the U.S economy grew at a 2.5 percent annual rate in the first quarter, slower than the 4.4 percent rate in the previous quarter [...continues]
SemStamp	DAQ declined by less than half that amount as investors grew cautious amid signs global demand is weakening despite low interest rates worldwide. [...continues]
Unigram	Nasdaq dropped 1.2 percent. The disappointing economic data came from the Commerce Department (search), which said gross domestic product (GDP) grew at a seasonally adjusted annual rate of 3.3 percent in the first quarter, slower than the 4.8 percent pace in the previous quarter. [...continues]
Adaptive	Nasdaq dropped 0.6 percent. The U.S economy expanded at a 1 percent pace in the first quarter, according to a revised estimate from the government, which is below expectations of 2 percent to 2.5 percent. It also reflected a downward revision to growth in the previous quarter.[...continues]
UPV	Nasdaq dropped 1.4 percent. The report on U.S. gross domestic product (GDP) released by the Commerce Department showed the economy grew at an annual rate of 2.5 percent in the second quarter, below the 3.5 percent pace of the first quarter [...continues]

5.5. WMs with Certified Robustness

There has been no systematic work that formally guarantees the effectiveness of WMs for LLMs. Therefore, it is essential to develop WMs with certified robustness, where the magnitude of the WMs is explicitly measured in relation to the accuracy of generated outputs in both generation tasks and downstream tasks. This approach would ensure that WMs remain identifiable even in adversarial settings, while maintaining a defined threshold for output quality. By quantifying this relationship, we can create a system that balances robustness and model utility, offering a more reliable and scalable method for watermarking LLMs.

5.6. Explore WMs and Alignment in LLMs

Alignment refers to ensuring that LLMs generate outputs consistent with standard ethical and operational guidelines [129]. However, WMs can interfere with this alignment by altering the generation or model training process, potentially reducing the model’s adherence to its pre-trained alignment. This issue becomes particularly critical in sensitive domains where the model’s behaviors must align with strict safety or ethical standards. The impact of WMs on LLM alignment demands careful investigation, as disruptions in logits, sampling, or model weights could cause semantic drift or unintended biases in LLM outputs.

6. Conclusion

We presented an SoK for WMs in LLMs, with a particular focus on their potential for real-world deployment through a comprehensive evaluation. Our work encompasses a broader spectrum, including assessments of model utility in text generation and downstream tasks, semantic preservation of output under perplexity observation and qualitative assessment, robustness under various attack types, and watermark sensitivity across LLMs, and consider other practical aspects.

To capture the intricate relationships among WMs, LLMs, and model utility, we introduced a novel cross-model IP classifier and examined the impacts of WMs on LLM output quality. Our experiments demonstrate that WMs significantly enhance the identification of LLM outputs, strengthening IP protection and reducing unauthorized use. However, while WMs are promising for practical applications, challenges remain due to potential impacts on model utility, text quality, and vulnerabilities to WM removal attacks. Our discussion highlights the current state of research, identifies key limitations, and outlines future directions to improve the robustness, scalability, and practicality of WMs. We emphasize the need for further investigation to address these challenges and make watermarking techniques more viable for deployment across a range of LLM applications.

References

- [1] Google, "Google gemini," <https://bard.google.com/chat/>, 2024.
- [2] OpenAI, "Openai api," 2024, <https://openai.com/index/openai-api/> [Accessed: 2024-09-20].
- [3] Anthropic, "Claude," 2024, <https://www.anthropic.com/claude> [Accessed: 2024-09-20].
- [4] Cohere, "Cohere," 2024, <https://cohere.com/> [Accessed: 2024-09-20].
- [5] B. Zhang, B. Haddow, and A. Birch, "Prompting large language model for machine translation: A case study," in *International Conference on Machine Learning*, 2023, pp. 41 092–41 110.
- [6] X. Zhang, N. Rajabi, K. Duh, and P. Koehn, "Machine translation with large language models: Prompting, few-shot learning, and fine-tuning with qLora," in *Proceedings of the Eighth Conference on Machine Translation*, 2023, pp. 468–481.
- [7] H. Xu, Y. J. Kim, A. Sharaf, and H. H. Awadalla, "A paradigm shift in machine translation: Boosting translation performance of large language models," *ICLR*, 2024.
- [8] Z. Hu, L. Chen, X. Wu, Y. Wu, H. Zhang, and H. Huang, "Unbiased watermark for large language models," *arXiv preprint arXiv:2310.10669*, 2023.
- [9] J. Kirchenbauer, J. Geiping, Y. Wen, J. Katz, I. Miers, and T. Goldstein, "A watermark for large language models," in *International Conference on Machine Learning*. PMLR, 2023, pp. 17 061–17 084.
- [10] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell et al., "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [11] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. D. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman et al., "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [12] Azure, "Model interpretability (preview) - azure machine learning," Available at <https://aka.ms/AzureMLModelInterpretability>, 2021.
- [13] Bluemix, "Ai explainability 360," Available at <https://aix360.mybluemix.net/>, 2021. [Online]. Available: <http://aix360.mybluemix.net/>
- [14] E. Wallace, M. Stern, and D. Song, "Imitation attacks and defenses for black-box machine translation systems," *EMNLP*, 2020.
- [15] Q. Xu, X. He, L. Lyu, L. Qu, and G. Haffari, "Student surpasses teacher: Imitation attack for black-box nlp apis," *COLING*, 2022.
- [16] J. Victor and A. Efrati, "Alphabets google and deepmind pause grudges, join forces to chase openai," <https://www.theinformation.com/articles/alphabets-google-and-deepmind-pause-grudges-join-forces-to-chase-openai>, 2023.
- [17] N. Carlini, D. Paleka, K. D. Dvijotham, T. Steinke, J. Hayase, A. F. Cooper, K. Lee, M. Jagielski, M. Nasr, A. Conmy et al., "Stealing part of a production language model," *arXiv preprint arXiv:2403.06634*, 2024.
- [18] C. Conikee, "An ai attack that even sherlock holmes would be impressed by..." <https://qwiet.ai/an-ai-attack-that-would-impress-even-sherlock-holmes/>, 2024.
- [19] Z. Li, C. Wang, S. Wang, and C. Gao, "Protecting intellectual property of large language model-based code generation apis via watermarks," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 2336–2350.
- [20] M. Xue, Z. Wu, Y. Zhang, J. Wang, and W. Liu, "Advparams: An active dnn intellectual property protection technique via adversarial perturbation based parameter encryption," *IEEE Transactions on Emerging Topics in Computing*, vol. 11, no. 3, pp. 664–678, 2022.
- [21] D. Yu, S. Naik, A. Backurs, S. Gopi, H. A. Inan, G. Kamath, J. Kulkarni, Y. T. Lee, A. Manol, L. Wutschitz et al., "Differentially private fine-tuning of language models," *arXiv preprint arXiv:2110.06500*, 2021.
- [22] K. Yoo, W. Ahn, and N. Kwak, "Advancing beyond identification: Multi-bit watermark for language models," *NAACL*, 2024.
- [23] A. Liu, L. Pan, X. Hu, S. Li, L. Wen, I. King, and S. Y. Philip, "An unforgeable publicly verifiable watermark for large language models," in *The Twelfth International Conference on Learning Representations*, 2023.
- [24] M. Christ, S. Gunn, and O. Zamir, "Undetectable watermarks for language models," in *The Thirty Seventh Annual Conference on Learning Theory*. PMLR, 2024, pp. 1125–1139.
- [25] R. Kudithipudi, J. Thickstun, T. Hashimoto, and P. Liang, "Robust distortion-free watermarks for language models," *TMLR*, 2023.
- [26] H. Zhang, B. L. Edelman, D. Francati, D. Venturi, G. Ateniese, and B. Barak, "Watermarks in the sand: Impossibility of strong watermarking for generative models," *ICML*, 2024.
- [27] Q. Pang, S. Hu, W. Zheng, and V. Smith, "Attacking llm watermarks by exploiting their strengths," *arXiv preprint arXiv:2402.16187*, 2024.
- [28] A. Liu, L. Pan, Y. Lu, J. Li, X. Hu, X. Zhang, L. Wen, I. King, H. Xiong, and P. Yu, "A survey of text watermarking in the era of large language models," *ACM Computing Surveys*, 2024.
- [29] Y. Liang, J. Xiao, W. Gan, and P. S. Yu, "Watermarking techniques for large language models: A survey," *arXiv preprint arXiv:2409.00089*, 2024.
- [30] H. N. Lalai, A. A. Ramakrishnan, R. S. Shah, and D. Lee, "From intentions to techniques: A comprehensive taxonomy and challenges in text watermarking for large language models," *arXiv preprint arXiv:2406.11106*, 2024.
- [31] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," 2019.
- [32] M. Azure, "Ai language," 2024, <https://azure.microsoft.com/en-us/products/ai-services/ai-language> [Accessed: 2024-09-20].
- [33] H. Face, "Meta llama," 2024, <https://huggingface.co/meta-llama> [Accessed: 2024-09-20].
- [34] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.
- [35] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *North American Chapter of the Association for Computational Linguistics*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:52967399>
- [36] Y. Liu, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, vol. 364, 2019.
- [37] Z. Yang, "Xlnet: Generalized autoregressive pretraining for language understanding," *arXiv preprint arXiv:1906.08237*, 2019.
- [38] A. Conneau and G. Lample, "Cross-lingual language model pretraining," *Advances in neural information processing systems*, vol. 32, 2019.
- [39] A. Radford and K. Narasimhan, "Improving language understanding by generative pre-training," 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:49313245>
- [40] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of machine learning research*, vol. 21, no. 140, pp. 1–67, 2020.
- [41] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," *CoRR*, vol. abs/1910.13461, 2019. [Online]. Available: <http://arxiv.org/abs/1910.13461>
- [42] K. Song, "Mass: Masked sequence to sequence pre-training for language generation," *arXiv preprint arXiv:1905.02450*, 2019.
- [43] Springs, "Large language model (llm) api: Full guide 2024," <https://shorturl.at/yI7T8>, 2024.
- [44] C. Chen and K. Shu, "Can llm-generated misinformation be detected?" *arXiv preprint arXiv:2309.13788*, 2023.
- [45] A. Mitra, S. P. Mohanty, and E. Kougianos, "The world of generative ai: Deepfakes and large language models," *arXiv preprint arXiv:2402.04373*, 2024.
- [46] C. Chen and K. Shu, "Combating misinformation in the age of llms: Opportunities and challenges," *AI Magazine*, 2023.
- [47] F. N. Motlagh, M. Hajizadeh, M. Majd, P. Najafi, F. Cheng, and C. Meinel, "Large language models in cybersecurity: State-of-the-art," *arXiv preprint arXiv:2402.00891*, 2024.
- [48] N. Jovanović, R. Staab, and M. Vechev, "Watermark stealing in large language models," *arXiv preprint arXiv:2402.19361*, 2024.
- [49] J. Zhang, S. Peng, Y. Gao, Z. Zhang, and Q. Hong, "Apmisa: Adversarial perturbation against model stealing attacks," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 1667–1679, 2023.
- [50] A. Liu, L. Pan, X. Hu, S. Meng, and L. Wen, "A semantic invariant robust watermark for large language models," *ICLR*, 2024.
- [51] Y. Wu, Z. Hu, H. Zhang, and H. Huang, "Dipmark: A stealthy, efficient and resilient watermark for large language models," 2024. [Online]. Available: <https://openreview.net/forum?id=FhZi7r4nzA>
- [52] J. Ren, H. Xu, Y. Liu, Y. Cui, S. Wang, D. Yin, and J. Tang, "A robust semantics-based watermark for large language model against paraphrasing," in *Findings of the Association for Computational Linguistics: NAACL 2024*, K. Duh, H. Gomez, and S. Bethard, Eds. Mexico City, Mexico: Association for Computational Linguistics, Jun. 2024, pp. 613–625. [Online]. Available: <https://aclanthology.org/2024.findings-naacl.40>

- [53] Y. Liu and Y. Bu, "Adaptive text watermark for large language models," *arXiv preprint arXiv:2401.13927*, 2024.
- [54] Z. Hu, L. Chen, X. Wu, Y. Wu, H. Zhang, and H. Huang, "Unbiased watermark for large language models," in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=uWVC5FVidc>
- [55] M. Huo, S. A. Somayajula, Y. Liang, R. Zhang, F. Koushanfar, and P. Xie, "Token-specific watermarking with enhanced detectability and semantic coherence for large language models," *arXiv preprint arXiv:2402.18059*, 2024.
- [56] C. Liang, Y. Bian, Y. Deng, D. Cai, S. Li, P. Zhao, and K.-F. Wong, "Watme: Towards lossless watermarking through lexical redundancy," in *ICLR 2024 Workshop on Secure and Trustworthy Large Language Models*, 2024.
- [57] J. Fu, X. Zhao, R. Yang, Y. Zhang, J. Chen, and Y. Xiao, "Gumbelsoft: Diversified language model watermarking via the gumbelmax-trick," *arXiv preprint arXiv:2402.12948*, 2024.
- [58] K. Yoo, W. Ahn, and N. Kwak, "Advancing beyond identification: Multi-bit watermark for large language models," 2024. [Online]. Available: <https://openreview.net/forum?id=KOTsHW6mBI>
- [59] X. Zhao, P. V. Ananth, L. Li, and Y.-X. Wang, "Provable robust watermarking for AI-generated text," in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=SsmT8aO45L>
- [60] A. Nemecek, Y. Jiang, and E. Ayday, "Topic-based watermarks for llm-generated text," *arXiv preprint arXiv:2404.02138*, 2024.
- [61] R. Zhang, S. S. Hussain, P. Neekhara, and F. Koushanfar, "{REMARK-LLM}: A robust and efficient watermarking framework for generative large language models," in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 1813–1830.
- [62] T. Lee, S. Hong, J. Ahn, I. Hong, H. Lee, S. Yun, J. Shin, and G. Kim, "Who wrote this code? watermarking for code generation," *arXiv preprint arXiv:2305.15060*, 2023.
- [63] L. Wang, W. Yang, D. Chen, H. Zhou, Y. Lin, F. Meng, J. Zhou, and X. Sun, "Towards codable watermarking for injecting multi-bits information to llms," in *The Twelfth International Conference on Learning Representations*, 2024.
- [64] Y. Lu, A. Liu, D. Yu, J. Li, and I. King, "An entropy-based text watermarking detection method," *arXiv preprint arXiv:2403.13485*, 2024.
- [65] Z. He, B. Zhou, H. Hao, A. Liu, X. Wang, Z. Tu, Z. Zhang, and R. Wang, "Can watermarks survive translation? on the cross-lingual consistency of text watermark for large language models," *arXiv preprint arXiv:2402.14007*, 2024.
- [66] Y. Fu, D. Xiong, and Y. Dong, "Watermarking conditional text generation for ai detection: Unveiling challenges and a semantic-aware watermark remedy," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 18 003–18 011.
- [67] B. Guan, Y. Wan, Z. Bi, Z. Wang, H. Zhang, Y. Sui, P. Zhou, and L. Sun, "Codeip: A grammar-guided multi-bit watermark for large language models of code," *arXiv preprint arXiv:2404.15639*, 2024.
- [68] B. Wouters, "Optimizing watermarks for large language models," *arXiv preprint arXiv:2312.17295*, 2023.
- [69] G. Niess and R. Kern, "Stylometric watermarks for large language models," *arXiv preprint arXiv:2405.08400*, 2024.
- [70] Y. Takezawa, R. Sato, H. Bao, K. Niwa, and M. Yamada, "Necessary and sufficient watermark for large language models," *arXiv preprint*, 2023.
- [71] N. Golowich and A. Moitra, "Edit distance robust watermarks for language models," *arXiv preprint arXiv:2406.02633*, 2024.
- [72] S. Aaronson and H. Kirchner, "Watermarking gpt outputs," <https://www.scottaaronson.com/talks/watermark.ppt>, 2022.
- [73] J. Fairuze, S. Garg, S. Jha, S. Mahloujifar, M. Mahmoody, and M. Wang, "Publicly detectable watermarking for language models," *arXiv preprint*, 2023.
- [74] T. Zhou, X. Zhao, X. Xu, and S. Ren, "Bileve: Securing text provenance in large language models against spoofing with bi-level signature," *arXiv*, 2024.
- [75] M. Mao, D. Wei, Z. Chen, X. Fang, and M. Chau, "A watermark for low-entropy and unbiased generation in large language models," *arXiv*, 2024.
- [76] S. Dathathri, A. See, S. Ghaisas, P.-S. Huang, R. McAdam, J. Welbl, V. Bachani, A. Kaskasoli, R. Stanforth, T. Matejovicova et al., "Scalable watermarking for identifying large language model outputs," *Nature*, vol. 634, no. 8035, pp. 818–823, 2024.
- [77] A. B. Hou, J. Zhang, T. He, Y. Wang, Y.-S. Chuang, H. Wang, L. Shen, B. Van Durme, D. Khashabi, and Y. Tsvetkov, "Semstamp: A semantic watermark with parathastic robustness for text generation," *NAACL*, 2024.
- [78] A. B. Hou, J. Zhang, Y. Wang, D. Khashabi, and T. He, "k-semstamp: A clustering-based semantic watermark for detection of machine-generated text," *arXiv preprint arXiv:2402.11399*, 2024.
- [79] Z. Sun, X. Du, F. Song, and L. Li, "Codemark: Imperceptible watermarking for code datasets against neural code completion models," in *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2023, pp. 1561–1572.
- [80] Z. Sun, X. Du, F. Song, M. Ni, and L. Li, "Coprotector: Protect open-source code against unauthorized training usage with data poisoning," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 652–660.
- [81] H. Xu, L. Xiang, X. Ma, B. Yang, and B. Li, "Hufu: A modality-agnostic watermarking system for pre-trained transformers via permutation equivariance," *arXiv preprint arXiv:2403.05842*, 2024.
- [82] C. Gu, C. Huang, X. Zheng, K.-W. Chang, and C.-J. Hsieh, "Watermarking pre-trained language models with backdooring," *arXiv* (2022).
- [83] P. Li, P. Cheng, F. Li, W. Du, H. Zhao, and G. Liu, "Plmmark: a secure and robust black-box watermarking framework for pre-trained language models," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 12, 2023, pp. 14 991–14 999.
- [84] C. Gu, X. L. Li, P. Liang, and T. Hashimoto, "On the learnability of watermarks for language models," *arXiv preprint arXiv:2312.04469*, 2023.
- [85] X. Xu, Y. Yao, and Y. Liu, "Learning to watermark llm-generated text via reinforcement learning," *arXiv preprint arXiv:2403.10553*, 2024.
- [86] R. Zhang and F. Koushanfar, "Emmark: Robust watermarks for ip protection of embedded quantized large language models," *arXiv*, 2024.
- [87] F. B. Baldassini, H. H. Nguyen, C.-C. Chang, and I. Echizen, "Cross-attention watermarking of large language models," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 4625–4629.
- [88] C. Zhu, J. Galjaard, P.-Y. Chen, and L. Y. Chen, "Duwak: Dual watermarks in large language models," *arXiv preprint arXiv:2403.13000*, 2024.
- [89] B. Huang and X. Wan, "Waterpool: A watermark mitigating trade-offs among imperceptibility, efficacy and robustness," *arXiv preprint*, 2024.
- [90] E. Giboulot and F. Teddy, "Watermax: breaking the llm watermark detectability-robustness-quality trade-off," *arXiv preprint arXiv:2403.04808*, 2024.
- [91] K. Pang, T. Qi, C. Wu, and M. Bai, "Adaptive and robust watermark against model extraction attack," *arXiv preprint arXiv:2405.02365*, 2024.
- [92] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt, "Measuring massive multitask language understanding," 2021. [Online]. Available: <https://arxiv.org/abs/2009.03300>
- [93] K. Krishna, Y. Song, M. Karpinska, J. Wieting, and M. Iyyer, "Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [94] L. Pan, A. Liu, Z. He, Z. Gao, X. Zhao, Y. Lu, B. Zhou, S. Liu, X. Hu, L. Wen et al., "Markllm: An open-source toolkit for llm watermarking," *arXiv*, 2024.
- [95] Z. Zhang, X. Zhang, Y. Zhang, L. Y. Zhang, C. Chen, S. Hu, A. Gill, and S. Pan, "Large language model watermark stealing with mixed integer programming," *arXiv preprint arXiv:2405.19677*, 2024.
- [96] V. S. Sadasivan, A. Kumar, S. Balasubramanian, W. Wang, and S. Feizi, "Can ai-generated text be reliably detected?" *arXiv preprint arXiv:2303.11156*, 2023.
- [97] "Industry analysis," <https://www.grandviewresearch.com/industry-analysis/large-language-model-llm-market-report>.
- [98] Salesforce, "Salesforce announces einstein gpt, the world's first generative ai for crm," 2023, accessed: 2024-11-14. [Online]. Available: <https://tinyurl.com/mr39zsrp>
- [99] Duolingo, "Introducing duolingo max, a learning experience powered by gpt-4," 2024, accessed: 2024-11-14. [Online]. Available: <https://blog.duolingo.com/duolingo-max/>
- [100] The Wall Street Journal, "Instacart joins chatgpt frenzy, adding chatbot to grocery shopping app," 2023, accessed: 2024-11-14. [Online]. Available: <https://shorturl.at/OVHhb>
- [101] Expedia Group, "Chatgpt wrote this press release — no, it didn't, but it can now assist with travel planning in the expedia app," 2023, accessed: 2024-11-14. [Online]. Available: <https://shorturl.at/kKX3A>
- [102] Morgan Stanley Wealth Management, "Morgan stanley wealth management announces key milestone in innovation journey with openai," 2023, accessed: 2024-11-14. [Online]. Available: <https://shorturl.at/0GbsD>
- [103] Stripe, "Stripe and openai collaborate to monetize openai's flagship products and enhance stripe with gpt-4," 2023, accessed: 2024-11-14. [Online]. Available: <https://shorturl.at/C1hHZ>

[104] Andy Beatman, "Ai-powered dialogues: Global telecommunications with azure openai service," 2024, accessed: 2024-11-14. [Online]. Available: <https://shorturl.at/ybhZo>

[105] Microsoft, "From vision to value realization: A closer look at how customers are embracing ai transformation to unlock innovation and deliver business outcomes," 2024, accessed: 2024-11-14.

[106] Meta, "How companies are using meta llama," 2024, accessed: 2024-11-14. [Online]. Available: <https://about.fb.com/news/2024/05/how-companies-are-using-meta-llama/>

[107] G. Cloud, "Cloud natural language," 2024, <https://cloud.google.com/natural-language> [Accessed: 2024-09-20].

[108] Lilt, "Lilt x google translate," 2024, accessed: 2024-11-14. [Online]. Available: <https://support.lilt.com/kb/lilt-x-google-translate>

[109] Business Wire, "Wipro launches 'google gemini experience zone' to accelerate ai-driven innovation for enterprises," 2024, accessed: 2024-11-14. [Online]. Available: <https://www.businesswire.com/news/home/20241106392250/en/>

[110] Google Cloud, "185 real-world generative ai use cases from industry leaders," 2024, accessed: 2024-11-14.

[111] G. Satyanarayana, J. Bhuvana, and M. Balamurugan, "Sentimental analysis on voice using aws comprehend," in 2020 International Conference on Computer Communication and Informatics (ICCCI), 2020, pp. 1–4.

[112] Amazon Web Services (AWS), "Amazon comprehend customers and partners," 2024, accessed: 2024-11-14. [Online]. Available: <https://aws.amazon.com/comprehend/customers/>

[113] IBM, "Ibm watson natural language understanding," 2024, <https://www.ibm.com/products/natural-language-understanding> [Accessed: 2024-09-20].

[114] IBM, "Client and partner stories," 2024, accessed: 2024-11-14. [Online]. Available: <https://www.ibm.com/case-studies/search>

[115] Anthropic, "Build with claude," 2024, accessed: 2024-11-14. [Online]. Available: <https://www.anthropic.com/customers>

[116] OpenAI, "Openai won't watermark chatgpt text because its users could get caught," <https://www.theverge.com/2024/8/4/24213268/openai-chatgpt-text-watermark-cheat-detection-tool>, 2024.

[117] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava et al., "Llama 2: Open foundation and fine-tuned chat models," 2023.

[118] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel et al., "Mistral 7b," 2023.

[119] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin et al., "Opt: Open pre-trained transformer language models," *arXiv preprint arXiv:2205.01068*, 2022.

[120] E. Almazrouei, H. Alobeidli, A. Alshamsi, A. Cappelli, R. Cojocaru, M. Debbah, É. Goffinet, D. Hesslow, J. Launay, Q. Malartic et al., "The falcon series of open language models," *arXiv preprint arXiv:2311.16867*, 2023.

[121] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[122] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite BERT for self-supervised learning of language representations," *CoRR*, vol. abs/1909.11942, 2019. [Online]. Available: <http://arxiv.org/abs/1909.11942>

[123] P. He, X. Liu, J. Gao, and W. Chen, "Deberta: Decoding-enhanced bert with disentangled attention," in *ICLR*, 2021.

[124] R. Taylor, M. Kardas, G. Cucurull, T. Scialom, A. Hartshorn, E. Saravia, A. Poulton, V. Kerkez, and R. Stojnic, "Galactica: A large language model for science," *arXiv preprint arXiv:2211.09085*, 2022.

[125] L. Ben Allal, A. Lozhkov, and E. Bakouch, "Smolm - blazingly fast and remarkably powerful," *Hugging Face Blog*, July 2024, accessed: 2024-09-20.

[126] G. A. Miller, "Wordnet: a lexical database for english," *Communications of the ACM*, vol. 38, no. 11, pp. 39–41, 1995.

[127] J. Dodge, M. Sap, A. Marasović, W. Agnew, G. Ilharco, D. Groeneveld, M. Mitchell, and M. Gardner, "Documenting large webtext corpora: A case study on the colossal clean crawled corpus," 2021. [Online]. Available: <https://arxiv.org/abs/2104.08758>

[128] T. Mikolov, A. Deoras, S. Kombrink, L. Burget, and J. Černocký, "Empirical evaluation and combination of advanced language modeling techniques," in *ISCA*, 2011.

[129] T. Shen, R. Jin, Y. Huang, C. Liu, W. Dong, Z. Guo, X. Wu, Y. Liu, and D. Xiong, "Large language model alignment: A survey," *arXiv preprint*, 2023.

[130] H. Face, "Transformers documentation," 2024, <https://huggingface.com/docs/transformers/index> [Accessed: 2024-09-20].

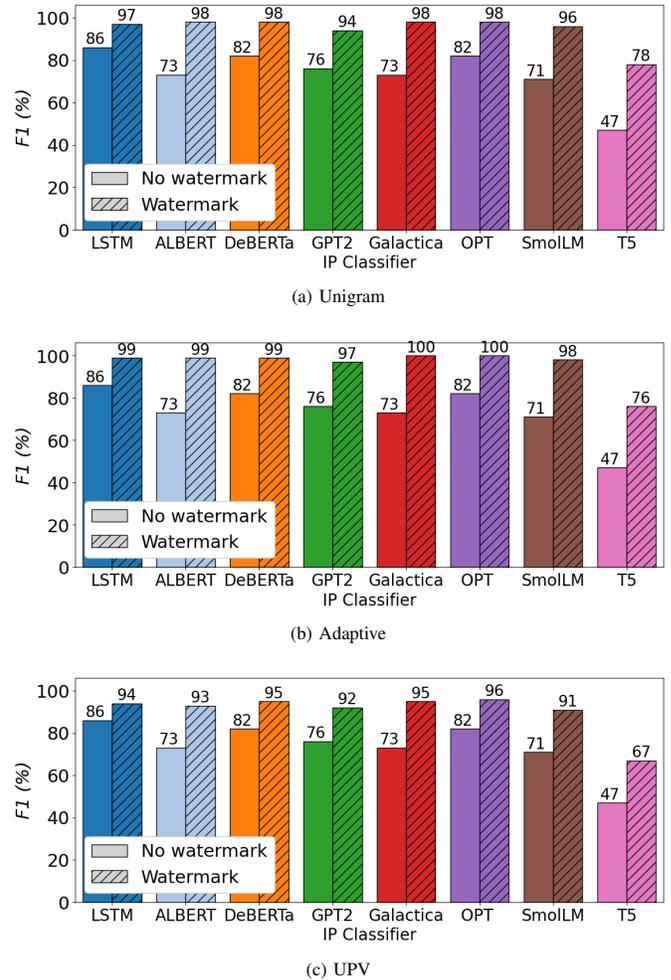


Figure 12: $F1$ scores as a function of different IP classifiers for different WMs, compared with No WM settings (Additional to results in Figure 7).

Appendix

Impacts of other WMs on IP Classifiers. As we applied multiple WMs to reassure our assessment, the cross-model IP classifier results of Unigram, Adaptive and UPV WMs are shown in Fig. 12. Similar insights could be drawn from the figure, reaffirming our observation of the study.

APIs for Deployment of Major LLM Providers. To provide context on the current deployment status of LLMs, we present Table 6, which compares various real-world APIs, outlining their applications, advantages, and limitations. Leading providers like OpenAI, Microsoft, Meta, Hugging Face, Google, and Amazon offer a range of LLM capabilities, including NLP, sentiment analysis, code generation, and model fine-tuning. While each API offers strengths such as scalability, integration, and community support, they also come with challenges like high resource demands, usage-based pricing, and limited customization options.

TABLE 6: LLM Real-world APIs.

API Name	Provider	Access Type	Applications	Advantages	Constraints
ChatGPT and GPT family [2]	OpenAI	API	NLP, text generation, questions and answers, chatbots	State-of-the-art LLMs, versatility, efficiency, personalization, cost-effective	Lack of real-time knowledge, cost for high usage
Microsoft Azure Language [32]	Microsoft Azure	API	Sentiment analysis, entity recognition, language understanding	Comprehensive APIs, scalability, integration with Azure ecosystem	Pricing based on usage, limited offline capabilities
LLaMA family [33]	Meta	Open source	Text/code generation, text summarization, translation, recommendations, chatbots	High model utility, versatility, efficiency, fine-tuning support	Resource intensive, overfitting, high maintenance and updates
Hugging Face transformers [130]	Hugging Face	Open source	NLP, model fine-tuning, text generation, research and development	large model repository, extensive community support, interoperability, versatility	No official support, resource intensive
Google Cloud AI-Language (Gemini, Bard, PaLM, etc.) [107]	Google	API	Sentiment analysis, entity recognition, translation, dialogue, text/code/image generation, comprehension	Google's robust and substantial infrastructure, easy integration with Google services, proficiency	Pricing based on usage, quality of generated content, inconsistent responses
Amazon Comprehend [111]	Amazon AWS	API	Entity recognition, sentiment analysis, classification, topic modeling	Integration with other AWS services, scalability, reliability	Pricing based on usage, limited customization
IBM Watson NLU [113]	IBM Watson	API	Sentiment and emotion analysis, keyword extraction	Comprehensive text analysis, customization, scalability, rich documentation and support, Integration with other IBM services	Cost for high usage, overhead for simple tasks
Fairseq [111]	Meta	Open source	Translation, research and development, language modeling	Versatility, research-friendly, integration with PyTorch	Limited user community, less user-friendly for beginners, maintenance
Claude [3]	Anthropic	API	Summarization, search, collaborative writing, coding comprehension	Safety and alignment focus, human-like interaction, multi-modality	Costly, limited customization
Cohere [4]	Cohere Inc.	API	Chatbots, knowledge assistants, scalability	User-friendly, seamless customization, focus on enterprise needs, privacy	Costly, small community

TABLE 7: Details on WM Schemes used in Experiments.

WM	Details	Mechanism Visualization
KGW [9]	Divides the vocabulary into green and red tokens, then modifies logits by adding a fixed WM to green tokens' logits to favor the generation of green tokens.	
SIR [50]	Determines WM logits by converting semantics of all preceding tokens using an auxiliary embedding model, then combines them with the original logits.	
EXP [25]	Intervenes the sampling process of each token by mapping it with a pseudo-random number sequence, in which every generated token corresponds with a key.	
SemStamp [77]	Generates new sentences by mapping candidates into embeddings, partitions semantic space into valid & blocked regions, and chooses sentences in valid regions.	
Unigram [59]	Modifies the logits to embed the WM by using a fixed grouping strategy, same green and red lists applied consistently for entire vocabulary.	
Adaptive [53]	Adaptively scales up logits of high-entropy tokens based on preceding tokens' semantics embedding.	
UPV [23]	Employs two networks for WM generation (modifying logits) and detection (evaluating entire text without a generation key).	